

# Moab Workload Manager

## Administrator Guide 10.0.1

September 2023



## Legal Notices

© 2018, 2023 Adaptive Computing Enterprises, Inc. All rights reserved.

This documentation and related software are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

This documentation and related software may provide access to or information about content, products, and services from third-parties. Adaptive Computing is not responsible for and expressly disclaims all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Adaptive Computing. Adaptive Computing will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Adaptive Computing.

Distribution of this document for commercial purposes in either hard or soft copy form is strictly prohibited without prior written consent from Adaptive Computing Enterprises, Inc..

Adaptive Computing, Cluster Resources, Moab, Moab Workload Manager, Moab Viewpoint, Moab Cluster Manager, Moab Cluster Suite, Moab Grid Scheduler, Moab Grid Suite, Moab Access Portal, and other Adaptive Computing products are either registered trademarks or trademarks of Adaptive Computing Enterprises, Inc. The Adaptive Computing logo and the Cluster Resources logo are trademarks of Adaptive Computing Enterprises, Inc. All other company and product names may be trademarks of their respective companies.

The information contained herein is subject to change without notice and is not warranted to be error free. If you find any errors, please report them to us in writing.

Adaptive Computing Enterprises, Inc.  
1100 5th Ave S, Suite #201  
Naples, FL 34102  
+1 (239) 330-6093  
[www.adaptivecomputing.com](http://www.adaptivecomputing.com)

# Contents

<b>Moab Workload Manager Overview</b> .....	<b>23</b>
<b>Chapter 1: Philosophy and Goals of Moab Workload Manager</b> .....	<b>24</b>
1.1 Value of a Batch System .....	25
1.1.1 Traffic Control .....	25
1.1.2 Mission Policies .....	25
1.1.3 Optimizations .....	25
1.2 Philosophy and Goals .....	26
1.2.1 Management Goals .....	26
1.2.2 Administration Goals .....	26
1.2.3 End User Goals .....	27
1.3 Workload .....	27
1.3.1 Batch Workload .....	27
1.3.2 Interactive Workload .....	28
1.3.3 Calendar Workload .....	28
1.3.4 Service Workload .....	28
<b>Chapter 2: Scheduler Basics</b> .....	<b>30</b>
2.1 Initial Moab Configuration .....	31
2.2 Layout of Scheduler Components .....	32
2.2.1 Layout of Scheduler Components .....	32
2.2.2 Layout of Scheduler Components with Integrated Database .....	34
2.3 Scheduling Environment .....	35
2.3.1 Jobs .....	35
2.3.2 Nodes .....	37
2.3.3 Advance Reservations .....	38
2.3.4 Policies .....	38
2.3.5 Resources .....	38
2.3.6 Class (or Queue) .....	40
2.3.7 Resource Manager (RM) .....	41
2.4 Scheduling Dictionary .....	42
2.5 Scheduling Iterations and Job Flow .....	52
2.5.1 Scheduling Iterations .....	52
2.5.2 Detailed Job Flow .....	53
2.6 Configuring the Scheduler .....	54
2.6.1 Adjusting Server Behavior .....	55
2.7 Credential Overview .....	58

2.7.1 General Credential Attributes .....	58
2.7.2 User Credential .....	64
2.7.3 Group Credential .....	65
2.7.4 Account (or Project) Credential .....	66
2.7.5 Class (or Queue) Credential .....	66
2.7.6 QoS Credential .....	88
2.8 Job Flags .....	89
<b>Chapter 3: Scheduler Commands .....</b>	<b>97</b>
3.1 Moab Command Overview .....	98
3.1.1 Moab Commands .....	98
3.1.2 Moab Command Options .....	100
3.1.3 Commands Providing Maui Compatibility .....	101
3.2 Status Commands .....	101
3.3 Job Management Commands .....	102
3.4 Reservation Management Commands .....	103
3.5 Policy/Configuration Management Commands .....	104
3.6 End-User Commands .....	104
3.7 Moab Commands .....	105
3.7.1 checkjob .....	107
3.7.2 checknode .....	119
3.7.3 mcredctl .....	123
3.7.4 mdiag .....	127
3.7.5 mdiag -a .....	132
3.7.6 mdiag -b .....	133
3.7.7 mdiag -c .....	133
3.7.8 mdiag -f .....	137
3.7.9 mdiag -j .....	139
3.7.10 mdiag -n .....	141
3.7.11 mdiag -p .....	147
3.7.12 mdiag -q .....	150
3.7.13 mdiag -r .....	151
3.7.14 mdiag -R .....	155
3.7.15 mdiag -s .....	156
3.7.16 mdiag -S .....	157
3.7.17 mdiag -t .....	158
3.7.18 mdiag -T .....	159
3.7.19 mdiag -u .....	161
3.7.20 mjobctl .....	162
3.7.21 mnodectl .....	179



3.7.22 moab .....	185
3.7.23 mrmctl .....	186
3.7.24 mrsvctl .....	190
3.7.26 mschedctl .....	216
3.7.27 mshow .....	225
3.7.28 mshow -a .....	227
3.7.29 mshow -a (mshow in a Hosting Environment) .....	237
3.7.30 msub .....	239
3.7.31 mvcctl (Moab Virtual Container Control) .....	268
3.7.32 showbf .....	273
3.7.33 showq .....	277
3.7.34 showhist.moab.pl .....	287
3.7.35 showres .....	291
3.7.36 showstart .....	297
3.7.37 showstate .....	300
3.7.38 showstats .....	302
3.7.39 showstats -f .....	315
3.7.40 Deprecated Commands .....	318

<b>Chapter 4: Prioritizing Jobs and Allocating Resources .....</b>	<b>334</b>
4.1 Job Prioritization .....	335
4.1.1 Priority Overview .....	335
4.1.2 Job Priority Factors .....	336
4.1.3 Fairshare Job Priority Example .....	348
4.1.4 Common Priority Usage .....	349
4.1.5 Prioritization Strategies .....	352
4.1.6 Manual Job Priority Adjustment .....	353
4.2 Node Allocation Policies .....	354
4.2.1 Node Allocation Overview .....	354
4.2.2 Node Selection Factors .....	358
4.2.3 Resource-Based Algorithms .....	358
4.2.4 User-Defined Algorithms .....	363
4.2.5 Specifying Per Job Resource Preferences .....	364
4.3 Node Access Policies .....	365
4.3.1 Node Access Policy Descriptions .....	365
4.3.2 Configuring Node Access Policies .....	366
4.4 Node Availability Policies .....	367
4.4.1 Node Resource Availability Policies .....	368
4.4.2 Node Categorization .....	369
4.4.3 Node Failure/Performance Based Notification .....	371

4.4.4 Node Failure/Performance Based Triggers .....	371
4.4.5 Handling Transient Node Failures .....	371
4.4.6 Allocated Resource Failure Policy for Jobs .....	373
<b>Chapter 5: Managing Fairness - Throttling Policies, Fairshare, Allocation Management .....</b>	<b>376</b>
5.1 Fairness Overview .....	377
5.1.1 Fairness Facilities .....	377
5.1.2 Selecting the Correct Policy Approach .....	380
5.2 Usage Limits/Throttling Policies .....	381
5.2.1 Fairness via Throttling Policies .....	381
5.2.2 Override Limits .....	393
5.2.3 Idle Job Limits .....	394
5.2.4 Hard and Soft Limits .....	396
5.2.5 Per-partition Limits .....	397
5.2.6 Usage-based limits .....	398
5.3 Fairshare .....	401
5.3.1 Fairshare Parameters .....	402
5.3.2 Using Fairshare Information .....	406
5.3.3 Hierarchical Fairshare/Share Trees .....	412
5.4 Sample FairShare Data File .....	417
5.5 Accounting, Charging, and Allocation Management .....	418
5.5.1 Accounting Manager Overview .....	418
5.5.2 Accounting Mode .....	419
5.5.3 Accounting Manager Interface Types .....	420
5.5.4 Charging for Jobs .....	423
5.5.5 Charging for Reservations .....	424
5.5.6 Accounting Properties Reported to the Accounting Manager .....	425
5.5.7 Accounting Stages .....	430
5.5.8 Accounting Events .....	433
5.5.9 Blocking Versus Non-Blocking Accounting Actions .....	433
5.5.10 Retrying Failed Charges .....	434
5.6 AMCFG Parameters and Flags .....	437
5.6.1 AMCFG Parameters .....	437
5.6.2 AMCFG Flags .....	460
<b>Chapter 6: Controlling Resource Access - Reservations, Partitions, and QoS Facilities .....</b>	<b>461</b>
6.1 Advance Reservations .....	462
6.1.1 Reservation Overview .....	462
6.1.2 Administrative Reservations .....	468
6.1.3 Standing Reservations .....	469
6.1.4 Reservation Policies .....	470

6.1.5 Configuring and Managing Reservations .....	474
6.1.6 Personal/User Reservations - Enabling Reservations for End Users .....	515
6.2 Partitions .....	518
6.2.1 Partition Overview .....	519
6.2.2 Defining Partitions .....	520
6.2.3 Managing Partition Access .....	520
6.2.4 Requesting Partitions .....	521
6.2.5 Per-Partition Settings .....	522
6.2.6 Miscellaneous Partition Issues .....	523
6.3 Quality of Service (QoS) Facilities .....	523
6.3.1 QoS Overview .....	524
6.3.2 QoS Enabled Privileges .....	525
6.3.3 Managing QoS Access .....	533
6.3.4 Requesting QoS Services at Job Submission .....	534
6.3.5 Restricting Access to Special Attributes .....	534
<b>Chapter 7: Optimizing Scheduling Behavior – Backfill and Node Sets .....</b>	<b>535</b>
7.1 Optimization Overview .....	535
7.2 Backfill .....	536
7.2.1 Backfill Overview .....	536
7.2.2 Backfill Algorithms .....	538
7.2.3 Configuring Backfill .....	540
7.3 Node Set Overview .....	543
7.3.1 Node Set Usage Overview .....	543
7.3.2 Node Set Configuration Examples .....	544
7.3.3 Requesting Node Sets for Job Submission .....	550
7.3.4 Configuring Node Sets for Classes .....	550
<b>Chapter 8: Evaluating System Performance - Statistics, Profiling and Testing .....</b>	<b>552</b>
8.1 Moab Performance Evaluation Overview .....	552
8.2 Accounting: Job and System Statistics .....	552
8.2.1 Accounting Overview .....	553
8.2.2 Real-Time Statistics .....	554
8.3 Testing New Versions and Configurations .....	555
8.3.1 MONITOR Mode .....	555
8.3.2 INTERACTIVE Mode .....	556
<b>Chapter 9: General Job Administration .....</b>	<b>557</b>
9.1 Job Holds .....	558
9.1.1 Holds and Deferred Jobs .....	558
9.1.2 User Holds .....	558

9.1.3 System Holds .....	558
9.1.4 Batch Holds .....	559
9.1.5 Job Defer .....	559
9.2 Job Priority Management .....	560
9.3 Suspend/Resume Handling .....	560
9.4 Checkpoint/Restart Facilities .....	561
9.5 Job Dependencies .....	562
9.5.1 Basic Job Dependency Support .....	562
9.5.2 Job Dependency Syntax .....	563
9.6 Job Defaults and Per Job Limits .....	564
9.6.1 Job Defaults .....	564
9.6.2 Per Job Maximum Limits .....	565
9.6.3 Per Job Minimum Limits .....	565
9.7 General Job Policies .....	566
9.7.1 Multi-Node Support .....	566
9.7.2 Multi-Req Support .....	566
9.7.3 Malleable Job Support .....	567
9.7.4 Enabling Job User Proxy .....	567
9.8 Using a Local Queue .....	568
9.9 Job Deadlines .....	571
9.9.1 Deadline Overview .....	571
9.9.2 Setting Job Deadlines via QoS .....	571
9.9.3 Job Termination Date .....	573
9.9.4 Conflict Policies .....	573
9.10 Job Arrays .....	574
9.10.1 Job Array Overview .....	574
9.10.2 Enabling Job Arrays .....	575
9.10.3 Subjob Definitions .....	575
9.10.4 Using Environment Variables to Specify Array Index Values .....	575
9.10.5 Job Array Cancellation Policies .....	578
9.10.6 Minimizing the Impact of Very Large Job Arrays .....	579
9.10.7 Examples .....	581
<b>Chapter 10: General Node Administration .....</b>	<b>584</b>
10.1 Node Attribute Types .....	585
10.1.1 Resource Manager Specified 'Opaque' Attributes .....	585
10.1.2 Scheduler Specified Default Node Attributes .....	585
10.1.3 Scheduler Specified Node Attributes .....	586
10.2 Node Location .....	586
10.2.1 Partitions .....	586

10.2.2 Racks .....	586
10.2.3 Queues .....	587
10.2.4 Node Selection .....	588
10.3 Node Attributes .....	590
10.3.1 Configurable Node Attributes .....	590
10.3.2 Node Features/Node Properties .....	599
10.4 Node Specific Policies .....	600
10.4.1 Node Usage/Throttling Policies .....	600
10.4.2 Node Access Policies .....	602
10.5 Managing Shared Cluster Resources (Floating Resources) .....	602
10.5.1 Shared Cluster Resource Overview .....	602
10.5.2 Configuring Generic Consumable Floating Resources .....	603
10.5.3 Configuring Cluster File Systems .....	604
10.5.4 Configuring Cluster Licenses .....	604
10.5.5 Configuring Generic Resources as Features .....	604
10.5.6 Configuring Generic Resources as Licenses .....	606
10.6 Managing Node State .....	607
10.6.1 Node State Definitions .....	607
10.6.2 Specifying Node States within Native Resource Managers .....	608
10.6.3 Moab Based Node State Adjustment .....	608
10.6.4 Adjusting Scheduling Behavior Based on Reported Node State .....	608
10.6.5 Adding or Removing Nodes .....	609
10.7 Managing Consumable Generic Resources .....	609
10.7.1 Differences Between Node Features and Consumable Resources .....	610
10.7.2 Configuring Node-locked Consumable Generic Resources .....	610
10.7.3 Managing Generic Resource Race Conditions .....	612
10.8 Enabling Generic Metrics .....	613
10.8.1 Configuring Generic Metrics .....	613
10.8.2 Example Generic Metric Usage .....	615
10.9 Enabling Generic Events .....	616
10.9.1 Configuring Generic Events .....	617
10.9.2 Reporting Generic Events .....	620
10.9.3 Generic Events Attributes .....	621
10.9.4 Manually Creating Generic Events .....	621
<b>Chapter 11: Resource Managers and Interfaces .....</b>	<b>622</b>
11.1 Resource Manager Overview .....	624
11.1.1 Scheduler/Resource Manager Interactions .....	625
11.1.2 Resource Manager Specific Details (Limitations/Special Features) .....	626
11.1.3 Synchronizing Conflicting Information .....	626

11.1.4 Evaluating Resource Manager Availability and Performance .....	627
11.2 Resource Manager Configuration .....	627
11.2.1 Defining and Configuring Resource Manager Interfaces .....	627
11.2.2 Resource Manager Configuration Details .....	653
11.2.3 Scheduler/Resource Manager Interactions .....	658
11.3 Resource Manager Extensions .....	659
11.3.1 Resource Manager Extension Specification .....	660
11.3.2 Resource Manager Extension Values .....	660
11.3.3 Resource Manager Extension Examples .....	687
11.3.4 Configuring dynamic features in Torque and Moab .....	688
11.4 Adding New Resource Manager Interfaces .....	689
11.4.1 Resource Manager Specific Interfaces .....	689
11.4.2 Wiki Interface .....	689
11.4.3 SSS Interface .....	689
11.5 Managing Resources Directly with the Native Interface .....	690
11.5.1 Native Interface Overview .....	690
11.5.2 Configuring the Native Interface .....	691
11.5.3 Generating Cluster Query Data .....	692
11.5.4 Interfacing with FlexNet (Formerly FLEXlm) .....	692
11.5.5 Interfacing to Nagios .....	694
11.5.6 Configuring Resource Types .....	695
11.5.7 Creating New Tools to Manage the Cluster .....	696
11.6 Utilizing Multiple Resource Managers .....	701
11.6.1 Multi-RM Overview .....	701
11.6.2 Configuring Multiple Independent Resource Manager Partitions .....	701
11.6.3 Migrating Jobs between Resource Managers .....	701
11.6.4 Aggregating Information into a Cohesive Node View .....	702
11.7 License Management .....	703
11.7.1 License Management Overview .....	703
11.7.2 Controlling and Monitoring License Availability .....	704
11.7.3 Requesting Licenses within Jobs .....	705
11.8 Resource Provisioning .....	706
11.8.1 Resource Provisioning Overview .....	706
11.8.2 Configuring Provisioning .....	706
11.9 Managing Networks .....	707
11.9.1 Network Management Overview .....	707
11.9.2 Dynamic VLAN Creation .....	707
11.9.3 Network Load and Health Monitoring .....	708
11.9.4 Creating a Resource Management Interface for a New Network .....	709
11.9.5 Per-Job Network Monitoring .....	710

11.10 Intelligent Platform Management Interface .....	710
11.10.1 IPMI Overview .....	710
11.10.2 Node IPMI Configuration .....	711
11.10.3 Installing IPMItool .....	711
11.10.4 [Optional] Creating the IPMI BMC-Node Map File .....	711
11.10.5 Configuring the Moab IPMI Tools .....	712
11.10.6 Configuring Moab .....	712
11.10.7 Ensuring Proper Setup .....	713
11.11 Resource Manager Translation .....	714
11.11.1 Translation Overview .....	714
11.11.2 Translation Enablement Steps .....	714
<b>Chapter 12: Troubleshooting and System Maintenance .....</b>	<b>715</b>
12.1 Internal Diagnostics/Diagnosing System Behavior and Problems .....	716
12.1.1 The mdiag Command .....	716
12.1.2 Other Diagnostic Commands .....	718
12.1.3 Using Moab Logs for Troubleshooting .....	718
12.1.4 Automating Recovery Actions after a Failure .....	718
12.2 Logging Overview .....	719
12.2.1 Log Facility Configuration .....	720
12.2.2 Standard Log Format .....	721
12.2.3 Searching Moab Logs .....	722
12.2.4 Event Logs .....	722
12.2.5 Enabling Syslog .....	725
12.2.6 Managing Verbosity .....	725
12.3 Object Messages .....	726
12.3.1 Object Message Overview .....	726
12.3.2 Viewing Messages .....	726
12.3.3 Creating Messages .....	726
12.4 Notifying Administrators of Failures .....	727
12.4.1 Enabling Administrator Email .....	727
12.4.2 Handling Events with the Notification Routine .....	727
12.5 Issues with Client Commands .....	729
12.5.1 Client Overview .....	729
12.5.2 Diagnosing Client Problems .....	729
12.6 Tracking System Failures .....	730
12.6.1 System Failures .....	730
12.6.2 Internal Errors .....	732
12.6.3 Reporting Failures .....	732
12.7 Problems with Individual Jobs .....	733

12.8 Diagnostic Scripts .....	733
12.8.1 support-diag.py .....	734
12.8.2 support.diag.pl .....	735
<b>Chapter 13: Improving User Effectiveness .....</b>	<b>737</b>
13.1 User Feedback Loops .....	737
13.1.1 Improving Job Size/Duration Requests .....	737
13.1.2 Improving Resource Requirement Specification .....	738
13.2 User Level Statistics .....	738
13.3 Enhancing Wallclock Limit Estimates .....	739
13.4 Job Start Time Estimates .....	739
13.4.1 Example .....	739
13.4.2 Estimation Types .....	739
13.5 Providing Resource Availability Information .....	740
13.6 Collecting Performance Information on Individual Jobs .....	741
<b>Chapter 14: Cluster Analysis and Testing .....</b>	<b>742</b>
14.1 Testing New Releases and Policies .....	742
14.1.1 Moab Evaluation Modes .....	743
14.1.2 Testing New Releases .....	744
14.1.3 Testing New Policies .....	745
14.1.4 Moab Side-by-Side .....	746
14.2 Testing New Middleware .....	746
14.2.1 Analysis Aspects .....	747
14.2.2 General Analysis .....	748
14.2.3 Native Mode Analysis .....	748
14.3 Workload Event Format .....	749
14.3.1 Workload Event Record Format .....	750
14.3.2 Reservation Event Records .....	759
14.3.3 Recording Job Events .....	760
<b>Chapter 15: Green Computing .....</b>	<b>762</b>
15.1 Green Computing Methods .....	763
15.1.1 Moab Edition Green Features .....	763
15.1.2 Moab Power Management Methods .....	764
15.1.3 Theory of Operation .....	766
15.1.4 Active Node Power Management .....	772
15.1.5 Idle Node Power Management .....	775
15.1.6 Green Policy Configuration .....	776
15.2 Deploying Adaptive Computing IPMI Scripts .....	776
15.2.1 Prerequisites .....	776



15.2.2 To Deploy the Adaptive Computing IPMI Scripts .....	777
15.3 Choosing which Nodes Moab Powers On or Off .....	778
15.4 Enabling Green Computing .....	779
15.5 Adjusting Green Pool Size .....	782
15.6 Handling Power-Related Events .....	783
15.7 Maximizing Scheduling Efficiency .....	783
15.8 Putting Idle Nodes in Power-Saving States .....	784
15.9 Troubleshooting Green Computing .....	785
<b>Chapter 16: Elastic Computing Overview .....</b>	<b>788</b>
16.1 About Elastic Computing .....	790
16.2 Configuring Elastic Computing .....	791
16.2.1 To Configure Elastic Computing .....	791
16.2.2 Sample moab.cfg File Excerpt .....	794
16.3 Elastic Trigger .....	794
16.4 Integration with a Private OpenStack Cloud .....	795
16.4.1 Configuring Moab to Talk to OpenStack Integration Scripts .....	795
16.4.2 Verification .....	796
16.4.3 Troubleshooting .....	797
16.5 Dynamic Nodes .....	797
16.5.1 Dynamic Node Parameters .....	797
16.5.2 Dynamic Node Events .....	798
16.5.3 Configuring Dynamic Nodes .....	799
16.6 Viewing Node and Trigger Information .....	801
16.6.1 mdiag -n -v --xml .....	801
16.6.2 mdiag -T .....	802
16.6.3 checknode -v <node name> .....	802
16.7 Usage Policies .....	803
16.7.1 Available Policies .....	803
16.7.2 Policy Levels .....	804
<b>Chapter 17: Object Triggers .....</b>	<b>805</b>
17.1 About Object Triggers .....	805
17.2 Object Trigger Tasks .....	806
17.2.1 Creating a Trigger .....	806
17.2.2 Using a Trigger to Send Email .....	809
17.2.3 Using a Trigger to Execute a Script .....	810
17.2.4 Using a Trigger to Perform Internal Moab Actions .....	811
17.2.5 Requiring an Object Threshold for Trigger Execution .....	812
17.2.6 Enabling Job Triggers .....	812
17.2.7 Modifying a Trigger .....	813

17.2.8 Viewing a Trigger .....	814
17.2.9 Checkpointing a Trigger .....	815
17.3 Object Trigger Reference .....	815
17.3.1 Job Triggers .....	816
17.3.2 Node Triggers .....	818
17.3.3 Reservation Triggers .....	819
17.3.4 Resource Manager Triggers .....	821
17.3.5 Scheduler Triggers .....	822
17.3.6 Threshold Triggers .....	823
17.3.7 Trigger Components .....	824
17.3.8 Trigger Exit Codes .....	833
17.3.9 Node Maintenance Example .....	833
17.3.10 Environment Creation Example .....	835
17.4 About Trigger Variables .....	836
17.4.1 Trigger Variable Tasks .....	837
17.4.2 Trigger Variable Reference .....	841
<b>Chapter 18: Miscellaneous .....</b>	<b>845</b>
18.1 User Feedback Overview .....	845
18.2 Enabling High Availability Features .....	847
18.2.1 High Availability Overview .....	847
18.2.2 Configuring High Availability on a Networked File System .....	848
18.2.3 Confirming High Availability on a Networked File System .....	849
18.2.4 Other High Availability Configuration .....	849
18.3 Malleable Jobs .....	850
18.4 Identity Managers .....	850
18.4.1 Identity Manager Overview .....	851
18.4.2 Basic Configuration .....	851
18.4.3 Importing Credential Fairness Policies .....	854
18.4.4 Identity Manager Data Format .....	854
18.4.5 Identity Manager Conflicts .....	855
18.4.6 Refreshing Identity Manager Data .....	855
18.5 Generic System Jobs .....	856
18.5.1 Creating a Generic System Job .....	856
18.5.2 Workflows Using Job Template Dependencies .....	857
<b>Chapter 19: Database Configuration .....</b>	<b>859</b>
19.1 SQLite3 .....	859
19.2 Connecting to a MySQL Database with an ODBC Driver .....	860
19.3 Connecting to a PostgreSQL Database with an ODBC Driver .....	863
19.4 Connecting to an Oracle Database with an ODBC Driver .....	865

19.4.1 To Connect to an Oracle Database with an ODBC Driver .....	866
19.4.2 Installing the Oracle Instant Client .....	872
19.5 Migrating Your Database to Newer Versions of Moab .....	875
19.5.1 Migrate from Moab 9.1 to Moab 10.0 .....	876
19.5.2 Migrate from Moab 9.0 to Moab 9.1 .....	876
19.5.3 Migrate from Moab 8.1 to Moab 9.0 .....	877
19.5.4 Migrate from Moab 8.0 to Moab 8.1 .....	877
19.5.5 Migrate from Moab 7.5 to Moab 8.0 .....	877
19.5.6 Migrate from Moab 7.2.6-7.2.10 to Moab 7.5 .....	878
19.5.7 Migrate from Moab 7.2.0-7.2.5 to Moab 7.2.6 .....	878
19.6 Importing Statistics from stats/DAY.* to the Moab Database .....	878
<b>Chapter 20: Accelerators .....</b>	<b>880</b>
20.1 Scheduling GPUs .....	880
20.1.1 Deploying and Configuring GPUs .....	881
20.1.2 Using GPUs with Minimal Configuration .....	882
20.2 Using GPUs with NUMA .....	882
20.3 NVIDIA GPUs .....	884
20.3.1 Using NVIDIA GPUs .....	884
20.3.2 Package Installation/Upgrade .....	885
20.3.3 Torque Configuration .....	886
20.3.4 GPU Modes for NVIDIA 260.x Driver .....	887
20.3.5 GPU Modes for NVIDIA 270.x Driver .....	887
20.3.6 gpu_status .....	887
20.3.7 Enabling Persistence Mode .....	888
20.3.8 Requesting GPUs and Setting GPU Mode .....	888
20.4 GPU Metrics .....	889
20.5 Intel® Xeon Phi™ Coprocessor Configuration .....	891
20.5.1 Intel Many-Integrated Cores (MIC) Architecture Configuration .....	892
20.5.2 Validating the Configuration .....	893
20.5.3 Job Submission .....	894
20.6 Intel® Xeon Phi™ Co-processor Metrics .....	896
<b>Chapter 21: Preemption .....</b>	<b>898</b>
21.1 Preemption Tasks .....	899
21.1.1 Canceling Jobs with Preemption .....	899
21.1.2 Checkpointing Jobs with Preemption .....	902
21.1.3 Requeuing Jobs with Preemption .....	904
21.1.4 Suspending Jobs with Preemption .....	906
21.1.5 Using Owner Preemption .....	910
21.1.6 Using QoS Preemption .....	914

21.2 Preemption Reference .....	915
21.2.1 Manual Preemption Commands .....	916
21.2.2 Preemption Flags .....	917
21.2.3 PREEMTPOLICY Types .....	918
21.2.4 Simple Example of Preemption .....	919
21.2.5 Testing and Troubleshooting Preemption .....	923
<b>Chapter 22: About Job Templates .....</b>	<b>926</b>
22.1 Job Template Tasks .....	926
22.1.1 Creating Job Templates .....	927
22.1.2 Viewing Job Templates .....	928
22.1.3 Applying Templates Based on Job Attributes .....	928
22.1.4 Requesting Job Templates Directly .....	929
22.1.5 Creating Workflows with Job Templates .....	930
22.2 Job Template Reference .....	931
22.2.1 Job Template Extension Attributes .....	931
22.2.2 Job Template Matching Attributes .....	944
22.2.3 Job Template Examples .....	945
22.2.4 Job Template Workflow Examples .....	946
<b>Chapter 23: Moab Workload Manager for Grids .....</b>	<b>948</b>
23.1 Grid Basics .....	950
23.1.1 Grid Overview .....	950
23.1.2 Grid Benefits .....	950
23.1.3 Management-Scalability .....	951
23.1.4 Resource Access .....	951
23.1.5 Load-Balancing .....	952
23.1.6 Single System Image (SSI) .....	952
23.1.7 High Availability .....	953
23.1.8 Grid Relationships .....	953
23.1.9 Submitting Jobs to the Grid .....	958
23.1.10 Viewing Jobs and Resources .....	959
23.2 Grid Configuration Basics .....	960
23.2.1 Peer Configuration Overview .....	960
23.2.2 Initial Configuration .....	960
23.3 Centralized Grid Management (Moab Grid Control / Moab Grid Member) .....	961
23.3.1 Moab Grid Control Configuration .....	961
23.3.2 Moab Grid Member Configuration .....	962
23.4 Hierarchical Grid Management .....	962
23.4.1 Configuring a Peer Server (Source) .....	962
23.4.2 Simple Hierarchical Grid .....	963

23.5 Localized Grid Management .....	964
23.5.1 Enabling Bi-Directional Job Flow .....	964
23.5.2 True Peer-to-Peer Grid .....	965
23.6 Resource Control and Access .....	965
23.6.1 Controlling Resource Information .....	966
23.6.2 Managing Resources with Grid Sandboxes .....	968
23.7 Workload Submission and Control .....	970
23.8 Reservations in the Grid .....	970
23.9 Grid Usage Policies .....	971
23.9.1 Grid Usage Policy Overview .....	971
23.9.2 Peer Job Resource Limits .....	972
23.9.3 Usage Limits via Peer Credentials .....	972
23.9.4 Using General Policies in a Grid Environment .....	973
23.10 Grid Scheduling Policies .....	974
23.10.1 Peer-to-Peer Resource Affinity Overview .....	974
23.10.2 Peer Allocation Policies .....	975
23.10.3 Per-partition Scheduling .....	975
23.11 Grid Credential Management .....	976
23.11.1 Peer Credential Management Overview .....	976
23.11.2 Peer Credential Mapping .....	976
23.11.3 Source and Destination Side Credential Mapping .....	978
23.11.4 Preventing User Space Collisions .....	978
23.12 Grid Data Management .....	978
23.12.1 Grid Data Management Overview .....	979
23.12.2 Configuring Peer Data Staging .....	979
23.12.3 Peer-to-Peer SCP Key Authentication .....	981
23.12.4 Diagnostics .....	982
23.13 Accounting and Allocation Management .....	984
23.13.1 Peer-to-Peer Accounting Overview .....	984
23.13.2 Peer-to-Peer Allocation Management .....	985
23.14 Grid Security .....	986
23.15 Grid Diagnostics and Validation .....	986
23.15.1 Peer Management Overview .....	986
23.15.2 Peer Diagnostic Overview .....	986
<b>Chapter 24: Data Staging .....</b>	<b>987</b>
24.1 Data Staging Example .....	987
24.2 Data Staging Tasks .....	988
24.2.1 Configuring the SSH Keys for the Data Staging Transfer Script .....	989
24.2.2 Configuring Data Staging .....	992

24.2.3 Staging Data to or from a Shared File System .....	994
24.2.4 Staging Data to or from a Shared File System in a Grid .....	999
24.2.5 Staging Data To or From a Compute Node .....	1005
24.2.6 Configuring Data Staging with Advanced Options .....	1009
24.3 Data Staging References .....	1012
24.3.1 Sample User Job Script .....	1012
<b>Chapter 25: Using NUMA with Moab .....</b>	<b>1014</b>
25.1 Using NUMA-Aware with Moab .....	1014
25.1.1 NUMA Process .....	1015
25.1.2 Installation and Configuration .....	1016
25.1.3 Moab and NUMA Resources .....	1016
25.1.4 Track Dedicated NUMA Resources .....	1017
25.2 Using NUMA-Support with Moab .....	1018
<b>Appendix A: Moab Parameters .....</b>	<b>1019</b>
<b>Appendix B: Multi-OS Provisioning .....</b>	<b>1184</b>
B.1 xCAT Plug-in Configuration Parameters .....	1185
B.2 Configuration Validation .....	1192
B.3 Deploying Images with Torque .....	1193
B.4 Installing Moab on the Management Node .....	1193
B.5 Integrating MSM and xCAT .....	1194
B.6 Moab Configuration File Example .....	1195
B.7 MSM Configuration .....	1196
B.8 MSM Installation .....	1196
B.9 Troubleshooting .....	1197
B.10 Verifying the Installation .....	1197
B.11 xCAT Configuration Requirements .....	1200
<b>Appendix C: Moab Event Dictionary .....</b>	<b>1201</b>
<b>Appendix D: Adjusting Default Limits .....</b>	<b>1353</b>
<b>Appendix E: Security .....</b>	<b>1359</b>
E.1 Authentication (Interface Security) .....	1359
E.1.1 Mauth Authentication .....	1359
E.1.2 Munge Authentication .....	1362
E.1.3 Server Response Control .....	1363
E.1.4 Checksum Algorithm for Client Authentication .....	1363
E.1.5 Interface Development Notes .....	1363
E.2 Authorization .....	1364

E.2.1 Role Based Authorization Security Configuration .....	1364
E.3 Host Security for Compute Resources .....	1367
E.3.1 Minimal Host Security Enforcement .....	1367
E.3.2 Medium Host Security Enforcement .....	1367
E.3.3 Strict Host Security Enforcement .....	1368
E.4 Securing Sensitive Configuration Information .....	1368
<b>Appendix F: Initial Moab Testing .....</b>	<b>1369</b>
F.0.1 Scheduler Modes .....	1369
F.0.2 Normal Mode .....	1369
F.0.3 Monitor Mode (or Test Mode) .....	1369
F.0.4 Interactive Mode .....	1370
<b>Appendix G: Integrating Other Resources with Moab .....</b>	<b>1372</b>
G.1 Compute Resource Managers .....	1372
G.1.1 Moab-Torque Integration Guide .....	1372
G.2 Hardware Integration .....	1376
G.2.1 Moab-NUMA-Support Integration Guide .....	1376
G.3 Torque/PBS Integration Guide - RM Access Control .....	1380
G.3.1 Server Configuration .....	1380
G.3.2 (Optional) MOM Configuration .....	1381
G.4 Torque/PBS Config - Default Queue Settings .....	1381
G.4.1 Default Queue .....	1381
G.4.2 Queue Default Node and Walltime Attributes .....	1382
G.4.3 System-Wide Default Node and Walltime Attributes .....	1382
G.5 Provisioning Resource Managers .....	1382
G.5.1 Validating an xCAT Installation for Use with Moab .....	1382
<b>Appendix H: Interfacing with Moab (APIs) .....</b>	<b>1385</b>
H.1 Accounting Interfaces .....	1385
H.2 Grid Interfaces .....	1385
H.2.1 Services Utilized .....	1385
H.2.2 Services Provided .....	1386
H.3 Identity and Credential Management Interfaces .....	1386
H.4 Job Submission and Management Interface .....	1386
H.5 Query and Control APIs .....	1387
H.5.1 CLI (Command Line Interface) XML API .....	1388
H.6 Resource Management Interfaces .....	1389
<b>Appendix I: Considerations for Large Clusters .....</b>	<b>1390</b>
I.1 Handling Large Jobs .....	1390

I.2 Handling Large Numbers of Jobs .....	1390
I.2.1 Set a Minimum RMPOLLINTERVAL .....	1391
I.2.2 Reduce Command Processing Time .....	1391
I.2.3 Minimize Job Processing Time .....	1392
I.2.4 Load All Non-Completed Jobs at Startup .....	1393
I.2.5 Reducing Job Start Time .....	1393
I.2.6 Reducing Job Reservation Creation Time .....	1393
I.2.7 Optimizing Backfill Time .....	1394
I.2.8 Constraining Moab Logging - LOGLEVEL .....	1394
I.2.9 Preemption .....	1394
I.2.10 Handling Transient Resource Manager Failures .....	1394
I.2.11 Constrain the Number of Jobs Preempted Per Iteration .....	1394
I.2.12 Scheduler Settings .....	1395
I.2.13 Configure Torque for Large Job Numbers .....	1395
I.3 Handling Large Numbers of Nodes .....	1395
I.4 Handling Large SMP Systems .....	1396
I.5 Resource Manager Scaling .....	1396
I.6 Server Sizing .....	1397
<b>Appendix J: Configuring Moab as a Service .....</b>	<b>1398</b>
J.1 Moab Grid Scheduler Service Script .....	1398
J.2 Moab Workload Manager Service Scripts .....	1399
<b>Appendix K: Migrating from Maui 3.2 .....</b>	<b>1402</b>
K.1 Migrating from Maui to Moab .....	1402
K.2 Other Notes .....	1403
K.2.1 File Naming .....	1403
K.2.2 Statistics and Checkpointing .....	1403
K.2.3 Verify Configuration File Compatibility .....	1404
K.2.4 Environment Variables .....	1404
K.3 Running Maui and Moab Side-By-Side .....	1404
<b>Appendix L: Node Allocation Plug-in Developer Kit .....</b>	<b>1406</b>
L.1 Moab Configuration .....	1406
L.1.1 Moab.cfg .....	1406
L.1.2 Syntax Rules .....	1407
L.1.3 Troubleshooting .....	1408
L.2 Writing the Plug-In .....	1408
L.2.1 Node Allocation Plug-in .....	1408
L.2.2 API and Data Structures .....	1409



<b>Appendix M: Scalable Systems Software Specification</b> .....	<b>1414</b>
M.1 Scalable Systems Software Job Object Specification .....	1414
Status of this Memo .....	1414
Abstract .....	1414
Table of Contents .....	1415
1.0 Introduction .....	1417
2.0 Conventions Used in this Document .....	1420
3.0 The Job Model .....	1422
4.0 JobGroup Element .....	1423
5.0 Job and JobDefaults Element .....	1425
6.0 TaskGroup and TaskGroupDefaults Element .....	1442
7.0 Task and TaskDefaults Element .....	1444
8.0 Property Categories .....	1445
9.0 AwarenessPolicy Attribute .....	1448
10.0 References .....	1449
11.0 Units of Measure Abbreviations .....	1449
M.2 Scalable Systems Software Resource Management and Accounting Protocol (SSSRMAP) Mes- sage Format .....	1450
Status of this Memo .....	1450
Abstract .....	1450
Table of Contents .....	1451
1.0 Introduction .....	1452
2.0 Conventions Used in this Document .....	1452
3.0 Encoding .....	1453
4.0 Error Reporting .....	1466
5.0 References .....	1475
M.3 Scalable Systems Software Node Object Specification .....	1476
Status of this Memo .....	1476
Abstract .....	1476
Table of Contents .....	1476
1.0 Introduction .....	1477
2.0 Conventions Used in this Document .....	1478
3.0 The Node Model .....	1480
4.0 Node Element .....	1480
5.0 Units of Measure Abbreviations .....	1485
M.4 Scalable Systems Software Resource Management and Accounting Protocol (SSSRMAP) Wire Protocol .....	1486
Status of this Memo .....	1486
Abstract .....	1487
Table of Contents .....	1487

1.0 Introduction .....	1488
2.0 Conventions Used in this Document .....	1489
3.0 Encoding .....	1490
4.0 Transport Layer .....	1491
5.0 Framing .....	1491
6.0 Asynchrony .....	1493
7.0 Security .....	1494
8.0 References .....	1503
<b>Appendix N: Moab Resource Manager Language Interface Overview .....</b>	<b>1505</b>
N.1 Moab RM Language Socket Protocol Description .....	1505
N.1.1 RM Language Overview .....	1505
N.1.2 Checksum Algorithm ('C' Version) .....	1506
N.1.3 Header Creation (PERL Code) .....	1508
N.1.4 Header Processing (PERL Code) .....	1508
N.2 Moab Resource Manager Language Data Format .....	1509
N.2.1 Query Resources Data Format .....	1509
N.2.2 Query Workload Data Format .....	1514
<b>Appendix O: SCHEDCFG Flags .....</b>	<b>1520</b>

## Moab Workload Manager Overview

### Welcome to the *Moab Workload Manager 10.0.1 Administrator Guide*

This guide is intended for Moab Workload Manager system administrators.

Moab Workload Manager is a scheduling and management system designed for clusters, grids, and on-demand/utility computing systems. Moab:

- Applies site policies and extensive optimizations to orchestrate jobs, services, and other workload across the ideal combination of network, compute, and storage resources.
- Enables Adaptive Computing; allowing compute resources to be customized to changing needs and failed systems to be automatically fixed or replaced.
- Increases system resource availability, offers extensive cluster diagnostics, delivers powerful quality of service (QoS) and service level agreement (SLA) features, and it provides rich visualization of cluster performance through advanced statistics, reports, and charts. In addition, the Elastic Computing feature allows Moab to temporarily utilize systems that can provide additional resources to take care of increased workload demand (caused by high job backlog) in a more timely manner.

Moab also works with major resource management and resource monitoring tools. From hardware monitoring systems such as IPMI to provisioning systems and storage managers, Moab takes advantage of domain expertise to allow these systems to do what they do best, importing their state information and providing them with the information necessary to do their job better. Moab uses its global information to coordinate the activities of both resources and services, which optimizes overall performance in-line with high-level mission objectives.

# Chapter 1: Philosophy and Goals of Moab Workload Manager

The scheduler's purpose is to optimally use resources in a convenient and manageable way. System users want to specify resources, obtain quick turnaround on their jobs, and have reliable resource allocation. On the other hand, administrators want to understand both the workload and the resources available. This includes current state, problems, and statistics—information about what is happening that is transparent to the end user. Administrators need an extensive set of options to enable management enforced policies and tune the system to obtain desired statistics.

There are other systems that provide batch management; however, Moab is unique in many respects. Moab matches jobs to nodes, dynamically reprovisions nodes to satisfy workload, and dynamically modifies workload to better take advantage of available nodes. Moab allows sites to fully visualize cluster and user behavior. It can integrate and orchestrate resource monitors, databases, identity managers, license managers, networks, and storage systems, therefore providing a cohesive view of the cluster—a cluster that fully acts and responds according to site mission objectives.

Moab can dynamically adjust security to meet specific job needs. Moab can create real and virtual clusters on demand and from scratch that are custom-tailored to a specific request. Moab can integrate visualization services, web farms, and application servers; it can also create powerful grids of disparate clusters. Moab maintains complete accounting and auditing records, exporting this data to information services on command, even providing professional billing statements to cover all used resources and services.

Moab provides user- and application-centric web portals and powerful graphical tools for monitoring and controlling every conceivable aspect of a cluster's objectives, performance, workload, and usage. Moab is unique in its ability to deliver a powerful user-centric cluster with little effort. Its design is focused on ROI, better use of resources, increased user effectiveness, and reduced staffing requirements.

In this chapter:

1.1 Value of a Batch System .....	25
1.1.1 Traffic Control .....	25
1.1.2 Mission Policies .....	25
1.1.3 Optimizations .....	25
1.2 Philosophy and Goals .....	26
1.2.1 Management Goals .....	26
1.2.2 Administration Goals .....	26

1.2.3 End User Goals .....	27
1.3 Workload .....	27
1.3.1 Batch Workload .....	27
1.3.2 Interactive Workload .....	28
1.3.3 Calendar Workload .....	28
1.3.4 Service Workload .....	28

## 1.1 Value of a Batch System

Batch systems provide centralized access to distributed resources through mechanisms for submitting, launching, and tracking jobs on a shared resource. This greatly simplifies use of the cluster's distributed resources, allowing users a *single system image* in terms of managing jobs and aggregate compute resources available. Batch systems should do much more than just provide a global view of the cluster, though. Using compute resources in a fair and effective manner is complex, so a scheduler is necessary to determine when, where, and how to run jobs to optimize the cluster. This topic describes the categories of scheduling decisions.

### 1.1.1 Traffic Control

A scheduler must prevent jobs from interfering. If jobs contend for resources, cluster performance decreases, job execution is delayed, and jobs may fail. Therefore, the scheduler tracks resources and dedicates requested resources to a particular job, which prevents use of such resources by other jobs.

### 1.1.2 Mission Policies

Clusters and other HPC platforms typically have specific purposes; to fulfill these purposes, or mission goals, there are usually rules about system use pertaining to who or what is allowed to use the system. To be effective, a scheduler must provide a suite of policies allowing a site to *map* site mission policies into scheduling behavior.

### 1.1.3 Optimizations

The compute power of a cluster is a limited resource; over time, demand inevitably exceeds supply. Intelligent scheduling decisions facilitate higher job volume and faster job

completion. Though subject to the constraints of the traffic control and mission policies, the scheduler must use whatever freedom is available to maximize cluster performance.

## 1.2 Philosophy and Goals

Managers want high system utilization and the ability to deliver various qualities of service to various users and groups. They need to understand how available resources are delivered to users over time. They also need administrators to tune *cycle delivery* to satisfy the current site mission objectives.

Determining a scheduler's success is contingent upon establishing metrics and a means to measure them. The value of statistics is best understood if optimal statistical values are known for a given environment, including workload, resources, and policies. That is, if an administrator could determine that a site's typical workload obtained an average queue time of 3.0 hours on a particular system, that would be a useful *statistic*; however, if an administrator knew that through proper tuning the system could deliver an average queue time of 1.2 hours with minimal negative side effects, that would be valuable *knowledge*.

Moab development relies on extensive feedback from users, administrators, and managers. At its core, it is a tool designed to *manage* resources and provide meaningful information about what is actually happening on the system.

### 1.2.1 Management Goals

A manager must ensure that a cluster fulfills the purpose for which it was purchased, so a manager must deliver cycles to those projects that are most critical to the success of the funding organizations. Management tasks to fulfill this role may include the following:

- Define cluster mission objectives and performance criteria
- Evaluate current and historical cluster performance
- Instantly graph delivered service

### 1.2.2 Administration Goals

An administrator must ensure that a cluster is effectively functioning within the bounds of the established mission goals. Administrators translate goals into cluster policies, identify and correct cluster failures, and train users in best practices. Given these objectives, an administrator may be tasked with each of the following:

- Maximize utilization and cluster responsiveness
- Tune fairness policies and workload distribution

- Automate time-consuming tasks
- Troubleshoot job and resource failures
- Instruct users of available policies and in their use regarding the cluster
- Integrate new hardware and cluster services into the batch system

### 1.2.3 End User Goals

End users are responsible for learning about the resources available, the requirements of their workload, and the policies to which they are subject. Using this understanding and the available tools, they find ways to obtain the best possible responsiveness for their own jobs. A typical end user may have the following tasks:

- Manage current workload
- Identify available resources
- Minimize workload response time
- Track historical usage
- Identify effectiveness of prior submissions

## 1.3 Workload

Moab can manage a broad spectrum of compute workload types, and it can optimize all workload types within the same cluster simultaneously, delivering on the objectives most important to each workload type, as described in this topic.

In this topic:

[1.3.1 Batch Workload - page 27](#)

[1.3.2 Interactive Workload - page 28](#)

[1.3.3 Calendar Workload - page 28](#)

[1.3.4 Service Workload - page 28](#)

### 1.3.1 Batch Workload

Batch workload is characterized by a *job* command file that typically describes all critical aspects of the needed compute resources and execution environment. With a batch job, the job is submitted to a job queue and runs somewhere on the cluster as resources become

available. In most cases, the submitter submits multiple batch jobs with no execution time constraints and processes job results as they become available.

Moab can enforce rich policies defining how, when, and where batch jobs run to deliver compute resources to the most important workload and provide general SLA guarantees while maximizing system utilization and minimizing average response time.

### 1.3.2 Interactive Workload

Interactive workload differs from batch in that requestors are interested in immediate response and are generally waiting for the interactive request to be executed before going on to other activities. In many cases, interactive submitters will continue to be *attached* to the interactive job, routing keystrokes and other input into the job and seeing both output and error information in real-time. While interactive workload can be submitted within a job file, commonly, it is routed into the cluster via a web or other graphical terminal and the end user may never even be aware of the underlying use of the batch system.

For managing interactive jobs, the focus is usually on setting aside resources to guarantee immediate execution or at least a minimal wait time for interactive jobs. Targeted service levels require management when mixing batch and interactive jobs. Interactive and other job types can be dynamically steered in terms of what they are executing, as well as in terms of the quantity of resources required by the application.

### 1.3.3 Calendar Workload

Calendar workload must be executed at a particular time and possibly in a regular periodic manner. For such jobs, time constraints range from flexible to rigid. For example, some calendar jobs may need to complete by a certain time, while others must run exactly at a given time each day or each week.

Moab can schedule the future and can therefore guarantee resource availability at needed times to allow calendar jobs to run as required. Furthermore, Moab provisioning features can locate or temporarily create the needed compute environment to properly execute the target applications.

### 1.3.4 Service Workload

Moab can schedule and manage both individual applications and long-running or persistent services. Service workload processes externally-generated transaction requests while Moab provides the distributed service with needed resources to meet target backlog or response targets to the service. Examples of service workload include parallel databases, web farms, and visualization services. Moab can apply cluster, [grid](#), or dynamically-generated on-demand resources to the service.



When handling service workload, Moab observes the application in a highly abstract manner. Using the **JOBCFG** parameter, aspects of the service jobs can be discovered or configured with attributes describing them as resource consumers possessing response time, backlog, state metrics, and associated QoS targets. In addition, each application can specify the type of compute resource required (OS, arch, memory, disk, network adapter, data store, and so forth), as well as the support environment (network, storage, external services, and so forth).

If the QoS response time/backlog targets of the application are not being satisfied by the current resource allocation, Moab evaluates the needs of this application against all other site mission objectives and workload needs and determines what it must do to locate or create (that is, provision, customize, secure) the needed resources. With the application resource requirement specification, a site can also indicate proximity/locality constraints, partition policies, ramp-up/ramp-down rules, and so forth.

Once Moab identifies and creates appropriate resources, it hands these resources to the application via a site customized URL. This URL can be responsible for whatever application-specific handshaking must be done to launch and initialize the needed components of the distributed application upon the new resources. Moab engages in the hand-off by providing needed context and resource information and by launching the URL at the appropriate time.

---

### Related Topics

- [Malleable Jobs](#)
- [QOS/SLA Enforcement](#)

## Chapter 2: Scheduler Basics

In this chapter:


2.1	Initial Moab Configuration .....	31
2.2	Layout of Scheduler Components .....	32
2.2.1	Layout of Scheduler Components .....	32
2.2.2	Layout of Scheduler Components with Integrated Database .....	34
2.3	Scheduling Environment .....	35
2.3.1	Jobs .....	35
2.3.2	Nodes .....	37
2.3.3	Advance Reservations .....	38
2.3.4	Policies .....	38
2.3.5	Resources .....	38
2.3.6	Class (or Queue) .....	40
2.3.7	Resource Manager (RM) .....	41
2.4	Scheduling Dictionary .....	42
2.5	Scheduling Iterations and Job Flow .....	52
2.5.1	Scheduling Iterations .....	52
2.5.2	Detailed Job Flow .....	53
2.6	Configuring the Scheduler .....	54
2.6.1	Adjusting Server Behavior .....	55
2.7	Credential Overview .....	58
2.7.1	General Credential Attributes .....	58
2.7.2	User Credential .....	64
2.7.3	Group Credential .....	65
2.7.4	Account (or Project) Credential .....	66
2.7.5	Class (or Queue) Credential .....	66
2.7.6	QoS Credential .....	88
2.8	Job Flags .....	89

## 2.1 Initial Moab Configuration

After Moab is installed, there may be minor configuration remaining within the primary configuration file, `moab.cfg`. While the `configure` script automatically sets these parameters, sites can choose to specify additional parameters. If the values selected in `configure` are satisfactory, then this section can be safely ignored.

The parameters needed for proper initial startup include the following:

Parameter	Instructions
<b>SCHEDCFG</b>	<p>The <code>SCHEDCFG</code> parameter specifies how the Moab server will execute and communicate with client requests. The <code>SERVER</code> attribute allows Moab client commands to locate the Moab server and is specified as a URL or in <code>&lt;HOST&gt;[:&lt;PORT&gt;]</code> format. For example:</p> <pre>SCHEDCFG[orion] SERVER=cw.psu.edu</pre> <p>Specifying the server in the Moab configuration file is optional. If nothing is specified, <code>gethostname()</code> is called. You can restart Moab and run <code>mdiag -S</code> to confirm that the correct host name is specified.</p> <p><b>i</b> The <code>SERVER</code> attribute can also be set using the environment variable <code>\$MOABSERVER</code>. Using this variable allows you to quickly change to the Moab server that client commands will connect to.</p> <pre>&gt; export MOABSERVER=cluster2:12221</pre>
<b>ADMINCFG</b>	<p>Moab provides role-based security enabled via multiple levels of admin access. Users who are to be granted full control of all Moab functions should be indicated by setting the <code>ADMINCFG[1]</code> parameter. The first user in this <code>USERS</code> attribute list is considered the <i>primary</i> administrator. It is the ID under which Moab will execute. For example, the following can be used to enable users <code>greg</code> and <code>thomas</code> as level 1 admins:</p> <pre>ADMINCFG[1] USERS=greg,thomas</pre> <p><b>i</b> Moab can only be launched by the primary administrator user ID.</p> <p><b>i</b> The primary administrator should be configured as a manager/operator/administrator in every resource manager with which Moab will interface.</p> <p><b>i</b> If the <code>msub</code> command will be used, then 'root' <i>must</i> be the primary administrator.</p>

Parameter	Instructions
	<div style="border: 1px solid #004a87; border-radius: 10px; padding: 10px; background-color: #e6f2ff;">  Moab's home directory and contents should be owned by the primary administrator. </div>
<b>RMCFG</b>	<p>For Moab to properly interact with a resource manager, the interface to this resource manager must be defined as described in the <a href="#">Resource Manager Configuration Overview</a>. Further, it is important that the primary Moab administrator also be a resource manager administrator within each of those systems. For example, to interface to a Torque resource manager, the following can be used:</p> <div style="border: 1px dashed #ccc; border-radius: 10px; padding: 5px; margin-top: 10px;"> <pre>RMCFG[torque1] TYPE=pbs</pre> </div>

---

### Related Topics

- [Parameter Overview](#)
- `mdiag -C` command (for diagnosing current Moab configuration)

## 2.2 Layout of Scheduler Components

In this topic:

- [2.2.1 Layout of Scheduler Components - page 32](#)
- [2.2.2 Layout of Scheduler Components with Integrated Database - page 34](#)

### 2.2.1 Layout of Scheduler Components

Moab is initially unpacked into a simple one-deep directory structure. What follows demonstrates the default layout of scheduler components; some of the files (such as log and statistics files) are created while Moab runs.

- `$(MOABHOMEDIR)` Default is `/opt/moab`, which can be modified via the `--with-homedir` parameter during `./configure`. `$(MOABHOMEDIR)` contains the files shown in the table below.

Filename	Description
<b>contrib/</b>	Directory containing contributed code and plug-ins
<b>.counters</b>	File containing last 3 counters for InsightIDs, jobs, and reservations respectively. Created during installation and required for Moab operation.
<b>docs/</b>	Directory for documentation
<b>etc/</b>	Directory for configuration files
<b>lib/</b>	Directory for library files (primarily for <code>tools/</code> )
<b>log/</b>	Directory for log files
<b>[etc/]moab.cfg</b>	General configuration file (can be located in <code>\$(MOABHOMEDIR)</code> or <code>\$(MOABHOMEDIR)/etc</code> ).
<b>.moab.ck</b>	Checkpoint file
<b>[etc/].moab.key</b>	Secret key used in authentication (can be located in <code>\$(MOABHOMEDIR)</code> or <code>\$(MOABHOMEDIR)/etc</code> ).
<b>moab.dat</b>	Configuration file generated by Moab Cluster Manager
<b>[etc/]moab-client.cfg</b>	Client configuration file (can be located in <code>\$(MOABHOMEDIR)</code> or <code>\$(MOABHOMEDIR)/etc</code> ).
<b>moab.lic</b>	License file
<b>moab.log</b>	Log file
<b>moab.log.1</b>	Previous log file
<b>.moab.pid</b>	Lock file
<b>[etc/]moab-private.cfg</b>	Secure configuration file containing private information (can be located in <code>\$(MOABHOMEDIR)</code> or <code>\$(MOABHOMEDIR)/etc</code> ).

Filename	Description
<b>stats/</b>	Directory for statistics files: <ul style="list-style-type: none"> <li>◦ <code>events.&lt;date&gt;</code> – event files</li> <li>◦ <code>{DAY WEEK MONTH YEAR}.&lt;date&gt;</code> – usage profiling data</li> <li>◦ <code>FS.&lt;PARTITION&gt;.&lt;epochtime&gt;</code> – fairshare usage data</li> </ul>

- `$(MOABINSTDIR)` Default is `/opt/moab`, which can be modified via the `--prefix` parameter during `./configure`. `$(MOABINSTDIR)` contains the files shown in the table below.

Filename	Description
<b>bin/</b>	Directory for client commands (for example, <code>showq</code> , <code>setres</code> , etc.)
<b>moab</b>	Moab binary
<b>sbin/</b>	Directory for server daemons
<b>tools/</b>	Directory for resource manager interfaces and local scripts

- `/etc/moab.cfg` If the Moab home directory cannot be found at startup, this file is checked to see if it declares the Moab home directory. If a declaration exists, the system checks the declared directory to find Moab. The syntax is:  
`MOABHOMEDIR=<DIRECTORY>`.

If you want to run Moab from a different directory other than `/opt/moab` but did not use the `--with-homedir` parameter during `./configure`, you can set the `$MOABHOMEDIR` environment variable, declare the home directory in the `/etc/moab.cfg` file, or use the `-C` command line option when using the Moab server or client commands to specify the configuration file location.

When Moab runs, it creates a log file, `moab.log`, in the `log/` directory and creates a statistics file in the `stats/` directory with the naming convention `events.WWW_MMM_DD_YYYY` (for example, `events.Sat_Oct_8_2022`). Additionally, a checkpoint file, `.moab.ck`, and lock file, `.moab.pid`, are maintained in the Moab home directory.

## 2.2.2 Layout of Scheduler Components with Integrated Database

### Layout of Scheduler Components with Integrated Database Enabled

If `USEDATABASE INTERNAL` is configured, the layout of scheduler components varies slightly. The `.moab.ck` file and usage profiling data (`stat/`

{ DAY | WEEK | MONTH | YEAR } .<date>) are stored in the `moab.db` database. In addition, the event information is stored in both event files: (`stat/events.<date>`) and `moab.db`.

---

## Related Topics

- [Commands Overview](#)

## 2.3 Scheduling Environment

Moab functions by manipulating a number of elementary objects, including jobs, nodes, reservations, QoS structures, resource managers, and policies. Multiple minor elementary objects and composite objects are also used; these objects are defined in [2.4 Scheduling Dictionary - page 42](#).

In this topic:

- [2.3.1 Jobs - page 35](#)
- [2.3.2 Nodes - page 37](#)
- [2.3.3 Advance Reservations - page 38](#)
- [2.3.4 Policies - page 38](#)
- [2.3.5 Resources - page 38](#)
- [2.3.6 Class \(or Queue\) - page 40](#)
- [2.3.7 Resource Manager \(RM\) - page 41](#)

Moab functions by manipulating a number of elementary objects, including jobs, nodes, reservations, QoS structures, resource managers, and policies. Multiple minor elementary objects and composite objects are also used; these objects are defined in the [scheduling dictionary](#).

### 2.3.1 Jobs

Job information is provided to the Moab scheduler from a resource manager such as PBS or Wiki. Job attributes include ownership of the job, [job state](#), amount and type of resources required by the job, and a wallclock limit indicating how long the resources are required. A job consists of one or more [task groups](#), each of which requests a number of resources of a given type; for example, a job may consist of two task groups, the first asking for a single master task consisting of *1 IBM SP node with at least 512 MB of RAM* and the second asking for a set of slave tasks such as *24 IBM SP nodes with at least 128 MB of RAM*. Each task

group consists of one or more [tasks](#) where a task is defined as the minimal independent unit of resources. By default, each task is equivalent to one processor. In SMP environments, however, users might want to tie one or more processors together with a certain amount of memory and other resources.

### 2.3.1.A Job States

The job's *state* indicates its current status and eligibility for execution and can be any of the values listed in the following tables:

**Table 2-1: Pre-execution states**

State	Definition
<b>Deferred</b>	Job that has been held by Moab due to an inability to schedule the job under current conditions. Deferred jobs are held for <a href="#">DEFERTIME</a> before being placed in the idle queue. This process is repeated <a href="#">DEFERCOUNT</a> times before the job is placed in batch hold.
<b>Hold</b>	Job is idle and is not eligible to run due to a user, (system) administrator, or batch system <i>hold</i> (also, <a href="#">batchhold</a> , <a href="#">systemhold</a> , <a href="#">userhold</a> ).
<b>Idle</b>	Job is currently queued and eligible to run but is not executing (also, <b>notqueued</b> ).
<b>NotQueued</b>	The job has not been queued.
<b>Unknown</b>	Moab cannot determine the state of the job.

**Table 2-2: Execution states**

State	Definition
<b>Starting</b>	Batch system has attempted to start the job and the job is currently performing <i>prestart</i> tasks that may include provisioning resources, staging data, or executing system pre-launch scripts.
<b>Running</b>	Job is currently executing the user application.
<b>Suspended</b>	Job was running but has been suspended by the scheduler or an administrator; user application is still in place on the allocated compute resources, but it is not executing.



**Table 2-3: Post-execution states**

State	Definition
<b>Completed</b>	Job has completed running without failure.
<b>Removed</b>	Job has run to its requested walltime successfully but has been canceled by the scheduler or resource manager due to exceeding its walltime or violating another policy; includes jobs canceled by users or administrators either before or after a job has started.
<b>Vacated</b>	Job canceled after partial execution due to a system failure.

### 2.3.1.B Task Group (or Req)

A *job task group* (or req) consists of a request for a single type of resources. Each task group consists of the following components:

Component	Description
<b>Task Definition</b>	A specification of the elementary resources that compose an individual task.
<b>Resource Constraints</b>	A specification of conditions that must be met for resource matching to occur. Only resources from nodes that meet <i>all</i> resource constraints can be allocated to the job task group.
<b>Task Count</b>	The number of task instances required by the task group.
<b>Task List</b>	The list of nodes on which the task instances are located.
<b>Task Group Statistics</b>	Statistics tracking resource utilization.

### 2.3.2 Nodes

Moab recognizes a node as a collection of resources with a particular set of associated attributes. This definition is similar to the traditional notion of a node found in a Linux cluster or supercomputer wherein a node is defined as one or more CPUs, associated memory, and possibly other compute resources such as local disk, swap, network adapters, and software licenses. Additionally, this node is described by various attributes such as an

architecture type or operating system. Nodes range in size from small uniprocessor PCs to large symmetric multiprocessing (SMP) systems where a single node may consist of hundreds of CPUs and massive amounts of memory.

In many cluster environments, the primary source of information about the configuration and status of a compute node is the [resource manager](#). This information can be augmented by additional information sources including node monitors and information services. Further, extensive node policy and node configuration information can be specified within Moab via the graphical tools or the configuration file. Moab aggregates this information and presents a comprehensive view of the node configuration, usages, and state.

While a node in Moab in most cases represents a standard compute host, nodes can also be used to represent more generalized resources. The GLOBAL node possesses floating resources that are available cluster wide, and created virtual nodes (such as network, software, and data nodes) track and allocate resource usage for other resource types.

For additional node information, see [General Node Administration](#).

### 2.3.3 Advance Reservations

An advance reservation dedicates a block of specific resources for a particular use. Each reservation consists of a list of resources, an access control list, and a time range for enforcing the access control list. The reservation ensures the matching nodes are used according to the access controls and policy constraints within the time frame specified. For example, a reservation could reserve 20 processors and 10 GB of memory for users Bob and John from Friday 6:00 a.m. to Saturday 10:00 p.m. Moab uses advance reservations extensively to manage backfill, guarantee resource availability for active jobs, allow service guarantees, support deadlines, and enable metascheduling. Moab also supports both regularly recurring reservations and the creation of dynamic one-time reservations for special needs. Advance reservations are described in detail in the [Advance Reservations](#) overview.

### 2.3.4 Policies

A configuration file specifies policies and controls how and when jobs start. Policies include job prioritization, fairness policies, fairshare configuration policies, and scheduling policies.

### 2.3.5 Resources

Jobs, nodes, and reservations all deal with the abstract concept of a resource. A resource in the Moab world is one of the following:

Resource	Description
<b>processors</b>	Specify with a simple count value
<b>memory</b>	Specify real memory or RAM in megabytes (MB)
<b>swap</b>	Specify virtual memory or <i>swap</i> in megabytes (MB)
<b>disk</b>	Specify local disk in megabytes (MB)

In addition to these elementary resource types, there are two higher level resource concepts used within Moab: [Task](#) and the processor equivalent (or [PE \(Processor Equivalent\) Calculation](#)), as explained below.

### 2.3.5.A Task

A task is a collection of elementary resources that must be allocated together within a single [node](#). For example, a task may consist of one processor, 512 MB of RAM, and 2 GB of local disk. A key aspect of a task is that the resources associated with the task must be allocated as an atomic unit, without spanning node boundaries. A task requesting 2 processors cannot be satisfied by allocating 2 uniprocessor nodes, nor can a task requesting 1 processor and 1 GB of memory be satisfied by allocating 1 processor on 1 node and memory on another.

In Moab, when jobs or reservations request resources, they do so in terms of tasks typically using a task count and a task definition. By default, a task maps directly to a single processor within a job and maps to a full node within reservations. In all cases, this default definition can be overridden by specifying a new task definition.

Within both jobs and reservations, depending on task definition, it is possible to have multiple tasks from the same job mapped to the same node. For example, a job requesting 4 tasks using the default task definition of 1 processor per task, can be satisfied by 2 dual processor nodes.

### 2.3.5.B PE (Processor Equivalent) Calculation

The concept of the processor equivalent, or PE, arose out of the need to translate multi-resource consumption requests into a scalar value. It is not an elementary resource but rather a derived resource metric. It is a measure of the actual impact of a set of requested resources by a job on the total resources available system wide. It is calculated as follows:

$$PE = \text{MAX}(\text{ProcsRequestedByJob} / \text{TotalOnlineProcs}, \text{MemoryRequestedByJob} / \text{TotalOnlineMemory},$$

```
DiskRequestedByJob / TotalOnlineDisk,
SwapRequestedByJob / TotalOnlineSwap) * TotalOnlineProcs
```

For example, if a job requested 20% of the total processors and 50% of the total memory of a 128-processor MPP system, only two such jobs could be supported by this system. The job is essentially using 50% of all available resources since the system can only be scheduled to its most constrained resource - memory in this case. The processor equivalents for this job should be 50% of the processors, or PE = 64.

Another example: Assume a homogeneous 100-node system with 4 processors and 1 GB of memory per node. A job is submitted requesting 2 processors and 768 MB of memory. The PE for this job would be calculated as follows:

$$PE = \text{MAX}(2 / (100 * 4), 768 / (100 * 1024)) * (100 * 4) = 3.$$

This result makes sense since the job would be consuming 3/4 of the memory on a 4-processor node.

The calculation works equally well on homogeneous or heterogeneous systems, uniprocessor or large SMP systems.

### 2.3.6 Class (or Queue)

A class (or queue) is a logical container object that implicitly or explicitly applies policies to jobs. In most cases, a class is defined and configured within the resource manager and associated with one or more of the following attributes or constraints:

Attribute	Description
<b>Default Job Attributes</b>	A queue can be associated with a default job duration, default size, or default resource requirements.
<b>Host Constraints</b>	A queue can constrain job execution to a particular set of hosts.
<b>Job Constraints</b>	A queue can constrain the attributes of jobs that can be submitted, including setting limits such as max wallclock time and minimum number of processors.
<b>Access List</b>	A queue can constrain who can submit jobs into it based on such things as user lists and group lists.
<b>Special Access</b>	A queue can associate special privileges with jobs including adjusted job priority.

As stated previously, most resource managers allow full class configuration within the resource manager. Where additional class configuration is required, the `CLASSCFG` parameter can be used.

Moab tracks class usage as a consumable resource allowing sites to limit the number of jobs using a particular class. This is done by monitoring class initiators that may be considered to be a ticket to run in a particular class. Any compute node can simultaneously support several types of classes and any number of initiators of each type. By default, nodes will have a one-to-one mapping between class initiators and configured processors. For every job task run on the node, one class initiator of the appropriate type is consumed. For example, a 3-processor job submitted to the class 'batch' consumes three batch class initiators on the nodes where it runs.

Using queues as consumable resources allows sites to specify various policies by adjusting the class initiator to node mapping. For example, a site running serial jobs might want to allow a particular 8-processor node to run any combination of batch and special jobs subject to the following constraints:

- Only 8 jobs of any type allowed simultaneously.
- No more than 4 special jobs allowed simultaneously.

To enable this policy, the site can set the node's `MAXJOB` policy to 8 and configure the node with 4 special class initiators and 8 batch class initiators.

In virtually all cases, jobs have a one-to-one correspondence between processors requested and class initiators required. However, this is not a requirement, and with special configuration, sites can choose to associate job tasks with arbitrary combinations of class initiator requirements.

In displaying class initiator status, Moab signifies the type and number of class initiators available using the format [`<CLASSNAME>:<CLASSCOUNT>`]. This is most commonly seen in the output of node status commands indicating the number of configured and available class initiators, or in job status commands when displaying class initiator requirements.

### 2.3.7 Resource Manager (RM)

While other systems may have more strict interpretations of a resource manager and its responsibilities, Moab's multi-resource manager support allows a much more liberal interpretation. In essence, any object that can provide environmental information and environmental control can be used as a resource manager, including sources of resource, workload, credential, or policy information such as scripts, peer services, databases, web services, hardware monitors, or even flat files. Likewise, Moab considers to be a resource manager any tool that provides control over the cluster environment whether that be a license manager, queue manager, checkpoint facility, provisioning manager, network manager, or storage manager.

Moab aggregates information from multiple unrelated sources into a larger more complete world view of the cluster that includes all the information and control found within a standard resource manager such as Torque, including node, job, and queue management services. For more information, see the [Resource Managers and Interfaces](#) overview.

### 2.3.7.A Arbitrary Resource

Nodes can also be configured to support various arbitrary resources. Use the `NODECFG` parameter to specify information about such resources. For example, you could configure a node to have *256 MB RAM, 4 processors, 1 GB Swap, and 2 tape drives*.

## 2.4 Scheduling Dictionary

Index: [A](#) [B](#) [C](#) [D](#) [E](#) [F](#) [G](#) [H](#) [I](#) [J](#) [K](#) [L](#) [M](#) [N](#) [O](#) [P](#) [Q](#) [R](#) [S](#) [T](#) [U](#) [V](#) [W](#) [X](#) [Y](#) [Z](#)

### A

Account	
<b>Definition</b>	A credential also known as 'project ID.' Multiple users can be associated a single account ID and each user can have access to multiple accounts. See <a href="#">credential</a> definition and <code>ACCOUNTCFG</code> parameter.
<b>Example</b>	<code>ACCOUNT=hgc13</code>

ACL (Access Control List)	
<b>Definition</b>	In the context of scheduling, an access control list is used and applied much as it is elsewhere. An ACL defines what credentials are required to access or use particular objects. The principal objects to which ACLs are applied are <a href="#">reservations</a> and <a href="#">QoS</a> s. ACLs can contain both allow and deny statements, include wildcards, and contain rules based on multiple object types.
<b>Example</b>	Reservation META1 contains 4 access statements: <ul style="list-style-type: none"> <li>• Allow jobs owned by user "john" or "bob"</li> <li>• Allow jobs with QoS "premium"</li> <li>• Deny jobs in class "debug"</li> <li>• Allow jobs with a duration of less than 1 hour</li> </ul>

Allocation	
<b>Definition</b>	A logical, scalar unit assigned to users on a credential basis, providing access to a particular quantity of compute resources. Allocations are consumed by jobs associated with those credentials.
<b>Example</b>	<code>ALLOCATION=30000</code>

B

C

Class	
<b>Definition</b>	(See <a href="#">Queue</a> ) A class is a logical container object that holds jobs allowing a site to associate various constraints and defaults to these jobs. Class access can also be tied to individual nodes defining whether a particular node will accept a job associated with a given class. Class based access to a node is denied unless explicitly allowed via resource manager configuration. Within Moab, classes are tied to jobs as a <a href="#">credential</a> .
<b>Example</b>	job "cw.073" is submitted to class batch node "cl02" accepts jobs in class batch reservation weekend allows access to jobs in class batch

CPU	
<b>Definition</b>	A single processing unit. A CPU is a consumable resource. Nodes typically consist of one or more CPUs. (same as <a href="#">processor</a> )

Credential	
<b>Definition</b>	An attribute associated with <a href="#">jobs</a> and other objects that determines object identity. In the case of schedulers and resource managers, credential based policies and limits are often established. At submit time, jobs are associated with a number of credentials such as user, <a href="#">group</a> , <a href="#">account</a> , <a href="#">QoS</a> , and <a href="#">class</a> . These job credentials subject the job to various polices and grant it various types of access. In most cases, credentials set both the privileges of the job and the ID of the actual job <a href="#">executable</a> .
<b>Example</b>	Job "cw.24001" possesses the following credentials:

## Credential

```
USER=john;GROUP=staff;ACCOUNT=[NONE];
QOS=[DEFAULT];CLASS=batch
```

## D

### Disk

**Definition** A quantity of local disk available for use by batch jobs. Disk is a [consumable resource](#).

## E

### Execution Environment

**Definition** A description of the environment where the executable is launched. This environment can include attributes such as the following:

- an executable
- command line arguments
- input file
- output file
- local user ID
- local group ID
- process resource limits

**Example** Job "cw.24001" possesses the following execution environment:

```
EXEC=/bin/sleep;ARGS="60";
INPUT=[NONE];OUTPUT=[NONE];
USER=loadl;GROUP=staff;
```

## F

### Fairshare

**Definition** A mechanism that allows historical resource utilization information to be incorporated into job priority decisions.



## Fairness

**Definition** The access to shared compute resources that each user is granted. Access can be equal or based on factors such as historical resource usage, political issues, and job value.

## G

### Group

**Definition** A [credential](#) typically directly mapping to a user's UNIX group ID.

## H

## I

## J

### Job

**Definition** The fundamental object of resource consumption. A job contains the following components:

- A list of required [consumable resources](#)
- A list of [resource constraints](#) controlling which resources can be allocated to the job
- A list of [job constraints](#) controlling where, when, and how the job should run
- A list of [credentials](#)
- An [execution environment](#)

### Job Constraints

**Definition** A set of conditions that must be fulfilled for the job to start. These conditions are far reaching and can include one or more of the following:

- When the job can run. After time X, within Y minutes.
- Which resources can be allocated. For example, node must possess at least 512 MB of RAM, run only in partition or Partition C, or run on HostA and HostB.
- Starting job relative to a particular event. Start after job X successfully completes.

## Job Constraints

### Example

```
RELEASETIME>='Tue Feb 12, 11:00AM'
DEPEND=AFTERANY:cw.2004
NODEMEMORY==256MB
```

## K

## L

## M

## Memory

### Definition

A quantity of physical memory (RAM). Memory is provided by compute nodes. It is required as a constraint or consumed as a consumable resource by jobs. Within Moab, memory is tracked and reported in megabytes (MB).

### Example

Node "node001" provides the following resources:

```
PROCS=1, MEMORY=512, SWAP=1024
```

"Job cw.24004" consumes the following resources per task:

```
PROCS=1, MEMORY=256
```

## N

## Node

### Definition

A node is the fundamental object associated with compute resources. Each node contains the following components:

- A list of [consumable resources](#)
- A list of [node attributes](#)

## Node Attribute

### Definition

A node attribute is a non-quantitative aspect of a node. Attributes typically describe the node itself or possibly aspects of various node resources such as processors or memory. While it is probably not optimal to aggregate node and resource attributes together in this manner, it is common practice. Common node attributes include processor architecture, operating system, and processor speed. Jobs often specify that resources be allocated from nodes possessing certain node attributes.

Node Attribute	
<b>Example</b>	<code>ARCH=AMD, OS=LINUX24, PROCSPEED=950</code>

Node Feature	
<b>Definition</b>	A node feature is a <a href="#">node attribute</a> that is typically specified locally via a configuration file. Node features are opaque strings associated with the node by the resource manager that generally only have meaning to the end-user, or possibly to the scheduler. A node feature is commonly associated with a subset of nodes allowing end-users to request use of this subset by requiring that resources be allocated from nodes with this feature present. In many cases, node features are used to extend the information provided by the resource manager.
<b>Example</b>	<div style="border: 1px dashed gray; padding: 5px; margin-bottom: 5px;"><code>FEATURE=s950,pIII,geology</code></div> <div style="border: 1px dashed gray; padding: 5px;"><i>This can be used to indicate that the node possesses a 950 MHz Pentium III processor and that the node is owned by the Geology department.</i></div>

## O

## P

Processor	
<b>Definition</b>	A processing unit. A processor is a consumable resource. Nodes typically consist of one or more processors. (same as CPU)

## Q

Quality of Service (QoS)	
<b>Definition</b>	An object that provides special services, resources, and so forth.

Queue	
<b>Definition</b>	(see <a href="#">Class</a> )

## R

Reservation	
<b>Definition</b>	An object that reserves a specific collection of resources for a specific timeframe for use by jobs that meet specific conditions. Each reservation consists of three major components: (1) a set of resources, (2) a time frame, and (3) an access control list. It is a scheduler role to ensure that the access control list is not violated during the reservation's lifetime (that is, its time frame) on the resources listed. For example, a reservation may specify that node002 is reserved for user Tom on Friday. The scheduler is therefore constrained to make certain that only Tom's jobs can use node002 at any time on Friday.
<b>Example</b>	Reserve 24 processors and 8 GB of memory from time T1 to time T2 for use by user X or jobs in the class batch.
Resource	
<b>Definition</b>	Hardware, generic resources such as software, and features available on a node, including memory, disk, swap, and processors.
Resource, Available	
<b>Definition</b>	A compute node's <a href="#">configured</a> resources minus the <i>maximum</i> of the sum of the resources <a href="#">utilized</a> by all job tasks running on the node and the resources <a href="#">dedicated</a> ; that is, $R.Available = R.Configure - MAX(R.Dedicated, R.Utilized)$ . In most cases, resources utilized will be associated with compute jobs that the batch system has started on the compute nodes, although resource consumption may also come from the operating system or <i>rogue</i> processes outside of the batch system's knowledge or control. Further, in a well-managed system, utilized resources are less than or equal to dedicated resources and when exceptions are detected, one or more <a href="#">usage-based limits</a> are activated to <a href="#">preempt</a> the jobs violating their requested resource usage.
<b>Example</b>	Node "cl003" has 4 processors and 512 MB of memory. It is executing 2 tasks of job "clserver.0041" that are using 1 processor and 60 MB of memory each. One processor and 250 MB of memory are reserved for user "jsmith" but are not currently in use. Resources available to user jsmith on node "cl003": <ul style="list-style-type: none"> <li>• 2 processors</li> <li>• 392 MB memory</li> </ul> Resources available to a user other than jsmith on node "cl003":

Resource, Available	
	<ul style="list-style-type: none"> <li>• 1 processor</li> <li>• 142 MB memory</li> </ul>

Resource, Configured	
<b>Definition</b>	The total amount of <b>consumable resources</b> that are available on a compute node for use by job tasks.
<b>Example</b>	<p>Node "cl003" has 4 processors and 512 MB of memory. It is executing 2 tasks of job "clserver.0041" that are using 1 processor and 60 MB of memory each. One processor and 250 MB of memory are reserved for user "jsmith" but are not currently in use.</p> <p>Configured resources for node "cl003":</p> <ul style="list-style-type: none"> <li>• 4 processors</li> <li>• 512 MB memory</li> </ul>

Resource, Consumable	
<b>Definition</b>	<p>Any object that can be used (that is, consumed and therefore made unavailable to another job) by, or dedicated to a job is considered to be a resource. Common examples of resources are a node's physical memory or local disk. As these resources can be given to one job and therefore become unavailable to another, they are considered to be consumable. Other aspects of a node, such as its operating system, are not considered to be consumable since its use by one job does not preclude its use by another. Note that some node objects, such as a network adapter, can be dedicated under some operating systems and resource managers and not under others. On systems where the network adapter cannot be dedicated and the network usage per job cannot be specified or tracked, network adapters are not considered to be resources, but rather attributes.</p> <p>Nodes possess a specific quantity of consumable resources such as real memory, local disk, or processors. In a resource management system, the node manager can choose to report only those configured resources available to batch jobs. For example, a node may possess an 80 GB hard drive but may have only 20 GB dedicated to batch jobs. Consequently, the resource manager may report that the node has 20 GB of local disk available when idle. Jobs can explicitly request a certain quantity of consumable resources.</p>

Resource, Constraint	
<b>Definition</b>	A resource constraint imposes a rule on which resources can be used to match a resource request. Resource constraints either specify a required quantity and

### Resource, Constraint

type of resource or a required node attribute. All resource constraints must be met by any given node to establish a match.

### Resource, Dedicated

**Definition** A job may request that a block of resources be dedicated while the job is executing. At other times, a certain number of resources can be reserved for use by a particular user or group. In these cases, the scheduler is responsible for guaranteeing that these resources, [utilized](#) or not, are set aside and made unavailable to other jobs.

**Example** Node "cl003" has 4 processors and 512 MB of memory. It is executing 2 tasks of job "clserver.0041" that are using 1 processor and 60 MB of memory each. One processor and 250 MB of memory are reserved for user "jsmith" but are not currently in use.

Dedicated resources for node "cl003":

- 1 processor
- 250 MB memory

### Resource, Utilized

**Definition** All [consumable](#) resources actually used by all job tasks running on the compute node.

**Example** Node "cl003" has 4 processors and 512 MB of memory. It is executing 2 tasks of job "clserver.0041" that are using 1 processor and 60 MB of memory each. One processor and 250 MB of memory are reserved for user "jsmith" but are not currently in use.

Utilized resources for node "cl003":

- 2 processors
- 120 MB memory

## S

### Swap

**Definition** A quantity of virtual memory available for use by batch jobs. Swap is a consumable resource provided by nodes and consumed by jobs.

## T

### Task

**Definition** An atomic collection of consumable resources.

### Time to Live (TTL)

**Definition** Specifies the time that a node is supposed to be retired by Moab. Moab will not schedule any jobs on a node after its time to live has passed.

## U

### User, Global

**Definition** The user credential used to provide access to functions and resources. In local scheduling, global user IDs map directly to local user IDs.

### User, Local

**Definition** The user credential under which the job executable will be launched.

## V

## W

### Workload

**Definition** A set of tasks to be performed or services to be provided by a system that comprises a set of resources.

## X

## Y

## Z

## 2.5 Scheduling Iterations and Job Flow

In this topic:

[2.5.1 Scheduling Iterations - page 52](#)

[2.5.2 Detailed Job Flow - page 53](#)

### 2.5.1 Scheduling Iterations

In any given scheduling iteration, many activities take place, examples of which are listed below:

- [Update State Information](#)
- [Handle User Requests](#)
- [Perform Next Scheduling Cycle](#)
- [Refresh reservations](#)
- [Schedule reserved jobs](#)
- [Schedule priority jobs](#)
- [Backfill jobs](#)
- [Update statistics](#)

#### 2.5.1.A Update State Information

Each iteration, the scheduler contacts the resource manager(s) and requests up-to-date information on compute resources, workload, and policy configuration. On most systems, these calls are to a centralized resource manager daemon that possesses all information. Jobs can be reported as being in any of the following states listed in the [job state](#) table.

#### 2.5.1.B Handle User Requests

User requests include any call requesting state information, configuration changes, or job or resource manipulation commands. These requests can come in the form of user client calls, peer daemon calls, or process signals.



### 2.5.1.C Perform Next Scheduling Cycle

Moab operates on a polling/event driven basis. When all scheduling activities complete, Moab processes user requests until a new resource manager event is received or an internal event is generated. Resource manager events include activities such as a new job submission or completion of an active job, addition of new node resources, or changes in resource manager policies. Internal events include administrator [schedule](#) requests, reservation activation/deactivation, or the expiration of the [RMPOLLINTERVAL](#) timer.

## 2.5.2 Detailed Job Flow

### 2.5.2.A Determine Basic Job Feasibility

The first step in scheduling is determining which jobs are feasible. This step eliminates jobs that have job holds in place, invalid job states (such as Completed, Not Queued, Deferred), or unsatisfied preconditions. Preconditions can include stage-in files or completion of preliminary job steps.

### 2.5.2.B Prioritize Jobs

With a list of feasible jobs created, the next step involves [determining the relative priority](#) of all jobs within that list. A priority for each job is calculated based on job attributes such as job owner, job size, and length of time the job has been queued.

### 2.5.2.C Enforce Configured Throttling Policies

Any configured [throttling policies](#) are then applied constraining how many jobs, nodes, processors, and so forth are allowed on a per credential basis. Jobs that violate these policies are not considered for scheduling.

### 2.5.2.D Determine Resource Availability

For each job, Moab attempts to locate the required compute resources needed by the job. For a match to be made, the node must possess all node attributes specified by the job and possess adequate available resources to meet the 'TasksPerNode' job constraint. (Default 'TasksPerNode' is 1.) Normally, Moab determines that a node has adequate resources if the resources are *neither utilized by nor dedicated to* another job using the calculation.

$$R.Available = R.Configured - \text{MAX}(R.Dedicated, R.Utilized).$$

The [NODEAVAILABILITYPOLICY](#) parameter can be modified to adjust this behavior.

### 2.5.2.E Allocate Resources to Job

If adequate resources can be found for a job, the [node allocation policy](#) is then applied to select the best set of resources. These allocation policies allow selection criteria such as speed of node, type of reservations, or excess node resources to be figured into the allocation decision to improve the performance of the job and maximize the freedom of the scheduler in making future scheduling decisions.

### 2.5.2.F Launch Job

With the resources selected and task distribution mapped, the scheduler then contacts the resource manager and informs it where and how to launch the job. The resource manager then initiates the actual job executable.

## 2.6 Configuring the Scheduler

Scheduler configuration is maintained using the flat text configuration file `moab.cfg`. All configuration file entries consist of simple `<PARAMETER> <VALUE>` pairs that are whitespace delimited. Parameter names are not case sensitive but `<VALUE>` settings are. Some parameters are array values and should be specified as `<PARAMETER> [<INDEX>]` (Example: `QOSCFG[hiprio] PRIORITY=1000`); the `<VALUE>` settings can be integers, floats, strings, or arrays of these. Some parameters can be specified as arrays wherein index values can be numeric or alphanumeric strings. If no array index is specified for an array parameter, an index of zero (0) is assumed. The example below includes both array based and non-array based parameters:

```
SCHEDCFG[cluster2] SERVER=head.c2.org MODE=NORMAL
LOGLEVEL 6
LOGDIR /var/tmp/moablog
```

See the [parameters](#) documentation for information on specific parameters.

The `moab.cfg` file is read when Moab is started up or recycled. Also, the `mschedctl -m` command can be used to reconfigure the scheduler at any time, updating some or all of the configurable parameters dynamically. This command can be used to modify parameters either permanently or temporarily. For example, the command `mschedctl -m LOGLEVEL 3` will temporarily adjust the scheduler log level. When the scheduler restarts, the log level restores to the value stored in the Moab configuration files. To adjust a parameter permanently, the option `--flags=persistent` should be set.

At any time, the current server parameter settings can be viewed using the `mschedctl -l` command.

## 2.6.1 Adjusting Server Behavior

Most aspects of Moab behavior are configurable. This includes both scheduling policy behavior and daemon behavior. In terms of configuring server behavior, the following realms are most commonly modified.

### 2.6.1.A Logging

Moab provides extensive and highly configurable logging facilities controlled by parameters.

Parameter	Description
<b>LOGDIR</b>	Indicates directory for log files.
<b>LOGFACILITY</b>	Indicates scheduling facilities to track.
<b>LOGFILE</b>	Indicates path name of log file.
<b>LOGFILEMAXSIZE</b>	Indicates maximum size of log file before rolling.
<b>LOGFILEROLLDEPTH</b>	Indicates maximum number of log files to maintain.
<b>LOGLEVEL</b>	Indicates verbosity of logging.

### 2.6.1.B Checkpointing

Moab checkpoints its internal state. The checkpoint file records statistics and attributes for jobs, nodes, reservations, users, groups, classes, and almost every other scheduling object.

Parameter	Description
<b>CHECKPOINTEXPIRATIONTIME</b>	Indicates how long unmodified data should be kept after the associated object has disappeared; that is, job priority for a job no longer detected.
<b>CHECKPOINTFILE</b>	Indicates path name of checkpoint file.
<b>CHECKPOINTINTERVAL</b>	Indicates interval between subsequent checkpoints.

### 2.6.1.C Client Interface

Clients will read from the client configuration file (`moab-client.cfg`), if present, and then from the server configuration file (`moab.cfg`), if present. First, clients will search for the presence of a `moab-client.cfg` file, loading client parameters from the first file detected in `$MOABHOMEDIR` or `$MOABHOMEDIR/etc`. Next, clients will search for the presence of a `moab.cfg` file, loading client parameters from the first file detected in `$MOABHOMEDIR` or `$MOABHOMEDIR/etc`, overriding any parameter values read from the client configuration file. If both files are present on a host, it is safe to remove the `moab-client.cfg` file after merging the client parameters into the `moab.cfg` file.

The Client interface is configured using the **SCHEDCFG** parameter. Most commonly, the attributes `SERVER` and `PORT` must be set to point client commands to the appropriate Moab server. Other parameters such as **CLIENTTIMEOUT** can also be set.

By default, Moab listens on all the interfaces of the machine on which it is installed. To bind Moab to a specific address use `'SCHEDCFG[] BINDADDRESS=<IPV4>'` and specify the specific IPv4 address of the interface on which Moab should listen. By default, Moab also verifies that the `SERVER` parameter matches the output of the `'gethostbyname'` system call. To configure Moab to use a different alias (on multi-homed hosts for example) you can specify the valid server aliases using `'SCHEDCFG[] SERVERALIAS=<alias1>,<alias2>...'`

### 2.6.1.D Scheduler Mode

The scheduler mode of operation is controlled by setting the `MODE` attribute of the **SCHEDCFG** parameter. The following modes are allowed:

Mode	Description
<b>INTERACTIVE</b>	Moab interactively confirms each scheduling action before taking any steps. See <a href="#">interactive mode overview</a> for more information.
<b>MONITOR</b>	Moab observes cluster and workload performance, collects statistics, interacts with allocation management services, and evaluates failures, but it does not actively alter the cluster, including workload scheduling, and resource provisioning. See <a href="#">monitor mode overview</a> for more information.
<b>NORMAL</b>	Moab actively schedules workload according to mission objectives and policies; it creates reservations; starts, cancels, preempts, and modifies jobs; and takes other scheduling actions.
<b>SINGLESTEP</b>	Moab behaves as in <b>NORMAL</b> mode but will only schedule a single iteration and then exit.

Mode	Description
<b>SLAVE</b>	Moab behaves as in <a href="#">NORMAL</a> mode but will only start a job when explicitly requested by a trusted <a href="#">grid peer</a> service or <a href="#">administrator</a> .
<b>TEST</b>	Moab behaves as in <a href="#">NORMAL</a> mode, will make reservations, and scheduling decisions, but will then only log scheduling actions it would have taken if running in <a href="#">NORMAL</a> mode. In most cases, "TEST" mode is identical to <a href="#">MONITOR</a> mode. See <a href="#">test mode overview</a> for more information.

### 2.6.1.E Configuring a Job ID Offset

Moab assigns job IDs as integers in numeric order as jobs are submitted, starting with 1. In some situations, you might want to offset the integer at which Moab starts to assign job IDs in your system.

This example describes how you would offset the job IDs in a compound system consisting of Site A, Site B, and Site C, each of which runs its own instance of Moab. Users belonging to any of the sites can submit jobs to their own site and to the other two. To simplify aggregation of usage records from the three sites, offset the job IDs for Site B to a starting value higher than the expected total lifetime value for the system; in this example, to 20000000. Likewise, set Site C to 20,000,000 more, or 40000000. To do so, set the `MINJOBID` attribute of `SCHEDCFG` in each system's `moab.cfg` to the offset value. To ensure that Moab will never use the same job ID for two different sites, also set `MAXJOBID`. If the Moab job naming process ever reaches the `MAXJOBID`, it will start over again with the `MINJOBID`.

```
SCHEDCFG[moab] SERVER=moab_siteA:4244 MAXJOBID=19999999
```

```
SCHEDCFG[moab] SERVER=moab_siteB:4344 MINJOBID=20000000 MAXJOBID=39999999
```

```
SCHEDCFG[moab] SERVER=moab_siteC:4444 MINJOBID=40000000 MAXJOBID=59999999
```

When users submit jobs to Moab using `msub`, Moab selects the job ID in numeric order, starting with 1 in Site A, 20000000 in Site B, and 40000000 in Site C.

If the compound system in this example uses Torque as its resource manager and users submit jobs directly to Torque using `qsub`, Torque assigns the job ID instead of Moab. In this case, you should also offset the Torque job IDs by setting the `next_job_number` server parameter of Site B and Site C to 20000000 and 40000000, respectively.

```
$user qmgr "set server next_job_number=20000000"
```

```
$user qmgr "set server next_job_number=40000000"
```

**i** Torque job ID limits will allow you to use the 20,000,000 offset scheme for up to 4 sites.

---

### Related Topics

- [Initial Configuration](#)
- Adding `#INCLUDE` files to `moab.cfg`

## 2.7 Credential Overview

Moab supports the concept of credentials, which provide a means of attributing policy and resource access to entities such as users and groups. These credentials allow specification of job ownership, tracking of resource usage, enforcement of policies, and many other features. There are five types of credentials: [user](#), [group](#), [account](#), [class](#), and [QoS](#). While the credentials have many similarities, each plays a slightly different role.

### In this topic:

- [2.7.1 General Credential Attributes - page 58](#)
- [2.7.2 User Credential - page 64](#)
- [2.7.3 Group Credential - page 65](#)
- [2.7.4 Account \(or Project\) Credential - page 66](#)
- [2.7.5 Class \(or Queue\) Credential - page 66](#)
- [2.7.6 QoS Credential - page 88](#)

### 2.7.1 General Credential Attributes

Internally, credentials are maintained as objects. Credentials can be created, destroyed, queried, and modified. They are associated with jobs and requests providing access and privileges. Each credential type has the following attributes:

- [Priority Settings](#)
- [Usage Limits](#)
- [Service Targets](#)
- [Credential and Partition Access](#)

- [Statistics](#)
- [Credential Defaults, State and Configuration Information](#)

All credentials represent a form of identity, and when applied to a job, express ownership. Consequently, jobs are subject to policies and limits associated with their owners.

### 2.7.1.A Credential Priority Settings

Each credential can be assigned a priority using the `PRIORITY` attribute. This priority affects a job's total credential priority factor as described in the [Priority Factors](#) section. In addition, each credential can also specify priority weight offsets, which adjust priority weights that apply to associated jobs. These priority weight offsets include [FSWEIGHT](#) (see [Priority-Based Fairshare](#) for more information), [QTWEIGHT](#), and [XFWWEIGHT](#).

For example:

```
# set priority weights
CREDWEIGHT      1
USERWEIGHT      1
CLASSWEIGHT     1
SERVICEWEIGHT  1
XFACTORWEIGHT  10
QUEUETIMEWEIGHT 1000
# set credential priorities
USERCFG[john]   PRIORITY=200
CLASSCFG[batch] PRIORITY=15
CLASSCFG[debug] PRIORITY=100
QOSCFG[bottomfeeder] QTWEIGHT=-50 XFWWEIGHT=100
ACCOUNTCFG[topfeeder] PRIORITY=100
```

### 2.7.1.B Credential Usage Limits

Usage limits constrain which jobs can run, which jobs can be considered for scheduling, and what quantity of resources each individual job can consume. With usage limits, policies such as [MAXJOB](#), [MAXNODE](#), and [MAXMEM](#) can be enforced against both idle and active jobs. Limits can be applied in any combination as shown in the example below where usage limits include 32 active processors per group and 12 active jobs for user `john`. For a job to run, it must satisfy the most limiting policies of all associated credentials. The [Throttling Policy](#) section documents credential usage limits in detail.

```
GROUPCFG[DEFAULT] MAXPROC=32 MAXNODE=100
GROUPCFG[staff]   MAXNODE=200
USERCFG[john]     MAXJOB=12
```

### 2.7.1.C Service Targets

Credential service targets allow jobs to obtain special treatment to meet usage or response time based metrics. Additional information about service targets can be found in the

[Fairshare](#) section.

### 2.7.1.D Credential and Partition Access

Access to partitions and to other credentials can be specified on a per credential basis with credential [access lists](#), [default credentials](#), and credential [membership lists](#).

#### Credential Access Lists

You can use the `ALIST`, `PLIST`, and `QLIST` attributes (shown in the following table) to specify the list of credentials or partitions that a given credential can access.

Credential	Attribute
<b>Account</b>	<code>ALIST</code> (allows credential to access specified list of accounts)
<b>Partition</b>	<code>PLIST</code> (allows credential to access specified list of partitions)
<b>QoS</b>	<code>QLIST</code> (allows credential to access specified list of <a href="#">QoSes</a> )

*Example 2-1:*

```
USERCFG[bob]    ALIST=jupiter,quantum
USERCFG[steve] ALIST=quantum
```

**i** Account-based access lists are only enforced if using an [accounting manager](#) or if the `ENFORCEACCOUNTACCESS` parameter is set to 'TRUE.'

#### Assigning Default Credentials

Use the `*DEF` attribute (shown in the following table) to specify the default credential or partition for a particular credential.

Credential	Attribute
<b>Account</b>	<code>ADEF</code> (specifies default account)
<b>Class</b>	<code>CDEF</code> (specifies default class)
<b>QoS</b>	<code>QDEF</code> (specifies default QoS)



*Example 2-2:*

```
# user bob can access accounts a2, a3, and a6. If no account is explicitly requested,
# his job will be assigned to account a3
USERCFG[bob] ALIST=a2,a3,a6 ADEF=a3
# user steve can access accounts a14, a7, a2, a6, and a1. If no account is explicitly
# requested, his job will be assigned to account a2
USERCFG[steve] ALIST=a14,a7,a2,a6,a1 ADEF=a2
```

## Specifying Credential Membership Lists

As an alternate to specifying access lists, administrators can also specify membership lists. This allows a credential to specify who can access it rather than allowing each credential to specify which credentials it can access. Membership lists are controlled using the `MEMBERULIST`, `EXCLUDEUSERLIST` and `REQUIREDUSERLIST` attributes, shown in the following table:

Credential	Attribute
User	---
Account, Group, QoS	<code>MEMBERULIST</code>
Class	<code>EXCLUDEUSERLIST</code> and <code>REQUIREDUSERLIST</code>

*Example 2-3:*

```
# account omega3 can only be accessed by users johnh, stevek, jenp
ACCOUNTCFG[omega3] MEMBERULIST=johnh,stevek,jenp
```

*Example 2-4: Controlling Partition Access on a Per User Basis*

A site may specify the user `john` can access partitions `atlas`, `pluto`, and `zeus` and will default to partition `pluto`. To do this, include the following line in the configuration file:

```
USERCFG[john] PLIST=atlas,pluto,zeus
```

*Example 2-5: Controlling QoS Access on a Per Group Basis*

A site may also choose to allow everyone in the group `staff` to access QoS `standard` and `special` with a default QoS of `standard`. To do this, include the following line in the configuration file:

```
GROUPCFG[staff] QLIST=standard,special QDEF=standard
```

*Example 2-6: Controlling Resource Access on a Per Account Basis*

An organization wants to allow everyone in the account `omega3` to access nodes 20 through 24. To do this, include the following in the configuration file:

```
ACCOUNTCFG [omega3] MEMBERULIST=johnh,stevek,jenp
SRCFG [omega3] HOSTLIST=r:20-24 ACCOUNTLIST=omega3
```

### 2.7.1.E Credential Statistics

Full statistics are maintained for each credential instance. These statistics record current and historical resource usage, level of service delivered, accuracy of requests, and many other aspects of workload. Note, though, that you must explicitly enable credential statistics as they are not tracked by default. You can enable credential statistics by including the following in the configuration file:

```
USERCFG [DEFAULT]          ENABLEPROFILING=TRUE
GROUPCFG [DEFAULT]        ENABLEPROFILING=TRUE
ACCOUNTCFG [DEFAULT]      ENABLEPROFILING=TRUE
CLASSCFG [DEFAULT]        ENABLEPROFILING=TRUE
QOSCFG [DEFAULT]          ENABLEPROFILING=TRUE
```

### 2.7.1.F Job Defaults, Credential State, and General Configuration

Credentials can apply defaults and force job configuration settings via the following parameters:

COMMENT	
<b>Description</b>	Associates a comment string with the target credential.
<b>Example</b>	<pre>USERCFG[steve] COMMENT='works for boss, provides good service' CLASSCFG[i3] COMMENT='queue for I/O intensive workload'</pre>

HOLD	
<b>Description</b>	<p>Specifies a hold should be placed on all jobs associated with the target credential. Any job associated with the target credential will remain in the hold state (i.e., the only way to remove the hold is to remove or disassociate the target credential from the job).</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> The order in which this HOLD attribute is evaluated depends on the following credential precedence: USERCFG, GROUPCFG, ACCOUNTCFG, CLASSCFG, QOSCFG, USERCFG [DEFAULT], GROUPCFG [DEFAULT], ACCOUNTCFG [DEFAULT], CLASSCFG [DEFAULT], QOSCFG [DEFAULT].</p> </div>
<b>Example</b>	<pre>USERCFG[user1] HOLD=false</pre>

HOLD	
	<pre>GROUPCFG[user1] HOLD=true</pre> <p><i>Moab evaluates the user hold first, sees that it should not put a hold on the job, and moves on with scheduling.</i></p> <pre>GROUPCFG[user1] HOLD=true CLASSCFG[user1] HOLD=false</pre> <p><i>Moab evaluates the group first, puts a hold on the job, and moves on.</i></p>

JOBFLAGS	
<b>Description</b>	Assigns the specified <a href="#">job flag</a> to all jobs with the associated credential.
<b>Example</b>	<pre>CLASSCFG[batch] JOBFLAGS=suspendable QOSCFG[special] JOBFLAGS=restartable</pre>

NOSUBMIT	
<b>Description</b>	Specifies whether jobs belonging to this credential can submit jobs using <code>msub</code> .
<b>Example</b>	<pre>ACCOUNTCFG[general] NOSUBMIT=TRUE CLASSCFG[special] NOSUBMIT=TRUE</pre>

OVERRUN	
<b>Description</b>	The amount of time a job can exceed its wallclock limit before being terminated. (Only applies to user and class credentials.)
<b>Example</b>	<pre>CLASSCFG[bigmem] OVERRUN=00:15:00</pre>

VARIABLE	
<b>Description</b>	Specifies attribute-value pairs associated with the specified credential. These variables can be used in triggers and other interfaces to modify system behavior.

## VARIABLE

### Example

```
GROUPCFG[staff] VARIABLE='nocharge=true'
```

Credentials can carry additional configuration information. They can specify that detailed statistical profiling should occur, that submitted jobs should be held, or that corresponding jobs should be marked as preemptible.

## 2.7.2 User Credential

The user credential is the fundamental credential within a workload manager; each job requires an association with exactly one user. In fact, the user credential is the only required credential in Moab; all others are optional. In most cases, the job's user credential is configured within or managed by the operating system itself, although Moab can be configured to obtain this information from an independent security and identity management service.

As the fundamental credential, the user credential has a number of unique attributes:

- [Role](#)
- [Privileges](#)
- [Email Address](#)
- [Disable Moab User Email](#)
- [Disable Memory Enforcement in RESOURCELIMITPOLICY](#)

### 2.7.2.A Role

Moab supports role-based authorization, mapping particular roles to collections of specific users. See the [Security](#) section for more information.

### 2.7.2.B Privileges

Moab supports the ability to configure which 'mdiag' commands a user can run.

Give all users as default:

```
USERCFG[DEFAULT] PRIVILEGES=RM:diagnose;NODE:diagnose
```

Users without any specific PRIVILEGES can run 'mdiag -R' and 'mdiag -n'.

Give specific PRIVILEGES:

```
USERCFG[carol] PRIVILEGES=SCHED:diagnose;NODE:diagnose
```

User 'carol' can run 'mdiag -S' and 'mdiag -n' but NOT 'mdiag -R'.

### 2.7.2.C Email Address

Facilities exist to allow user notification in the event of job or system failures or under other general conditions. This attribute allows these notifications to be mailed directly to the target user.

```
USERCFG[sally] EMAILADDRESS=sally@acme.com
```

### 2.7.2.D Disable Moab User Email

You can disable Moab email notifications for a specific user.

```
USERCFG[john] NOEMAIL=TRUE
```

### 2.7.2.E Disable Memory Enforcement in RESOURCELIMITPOLICY

You can disable memory enforcement for a specific user.

```
USERCFG[doug] FLAGS=DisableMemEnforcement
```

## 2.7.3 Group Credential

The group credential represents an aggregation of users. User-to-group mappings are often specified by the operating system or resource manager and typically map to a user's UNIX group ID. However, user-to-group mappings can also be provided by a security and identity management service, or you can specify such directly within Moab.

With many resource managers such as Torque and PBSPro, the group associated with a job is either the user's active primary group as specified within the operating system or a group that is explicitly requested at job submission time. When a secondary group is requested, the user's default group and associated policies are not taken into account. Also note that a job can only run under one group. If more constraining policies are required for these systems, an alternate aggregation scheme such as the use of [Account](#) or [QOS](#) credentials is recommended.

To enable support for secondary groups, add a [SCHEDCFG](#) line to moab.cfg with [FLAGS=EXTENDEDGROUPSUPPORT](#).

To submit a job as a secondary group, refer to your local resource manager's job submission options. For Torque users, see the `group_list=g_list` option of the `qsub -W` command.

## 2.7.4 Account (or Project) Credential

The account credential is also referred to as the project. This credential is generally associated with a group of users along the lines of a particular project for accounting and billing purposes. User-to-accounting mapping can be obtained from a resource manager or [accounting manager](#), or you can configure it directly within Moab. Access to an account can be controlled via the `ALIST` and `ADEF` credential attributes specified via the [Identity Manager](#) or the `moab.cfg` file.

The `MANAGERS` attribute (applicable only to the account and [class](#) credentials) allows an administrator to assign a user the ability to manage jobs inside the credential, as if the user is the job owner.

### Example 2-7: MANAGERS Attribute

```
ACCOUNTCFG[general]  MANAGERS=ops
ACCOUNTCFG[special]  MANAGERS=stevep
```

If a user is able to access more than one account, the desired account can be specified at job submission time using the resource-manager specific attribute. For example, with Torque this is accomplished using the `-A` argument to the `qsub` command.

### Example 2-8: Enforcing Account Usage

Job-to-account mapping can be enforced using the `ALIST` attribute and the [ENFORCEACCOUNTACCESS](#) parameter.

```
USERCFG[john]      ALIST=proj1,proj3
USERCFG[steve]     ALIST=proj2,proj3,proj4
USERCFG[brad]      ALIST=proj1
USERCFG[DEFAULT]  ALIST=proj2
ENFORCEACCOUNTACCESS TRUE
...
```

## 2.7.5 Class (or Queue) Credential

- [Class Job Defaults](#)
- [Per Job Min/Max Limits](#)
- [Resource Access](#)
- [Class Membership Constraints](#)
- [Attributes Enabling Class Access to Other Credentials](#)

- [Special Class Attributes \(such as Managers and Job Prologs\)](#)
- [Setting Default Classes](#)
- [Creating a Remap Class](#)
- [Class Attribute Overview](#)
- [Enabling Queue Complex Functionality](#)

The concept of the class credential is derived from the resource manager class or queue object. Classes differ from other credentials in that they more directly impact job attributes. In standard HPC usage, a user submits a job to a class and this class imposes a number of factors on the job. The attributes of a class can be specified within the resource manager or directly within Moab. Class attributes include the following:

- [Job Defaults](#)
- [Per Job Min/Max Limits](#)
- [Resource Access Constraints](#)
- [Class Membership Constraints](#)
- [Attributes Enabling Class Access to Other Credentials](#)
- [Special Class Attributes](#)

**i** For all classes configured in Moab, a resource manager queue with the same name should be created.

**i** When Torque reports a new queue to Moab, a class of the same name is automatically applied to all nodes (the same goes for existing queues when adding nodes). To associate nodes to only specific classes, add `CLASSCFG` entries for every Torque queue, and define the nodes linked to each queue/class via `HOSTLIST` expressions and/or `REMAPCLASS`. (This augments the optional `resources_default.neednodes` queue setting in `qmgr`.)

### 2.7.5.A Class Job Defaults

Classes can be assigned to a default [job template](#) that can apply values to job attributes not explicitly specified by the submitter. Additionally, you can specify shortcut attributes from the table that follows:

Attribute	Description
<b>DEFAULT.ATTR</b>	Job Attribute
<b>DEFAULT.DISK</b>	Required Disk (in MB)
<b>DEFAULT.EXT</b>	Job RM Extension
<b>DEFAULT.FEATURES</b>	Required Node Features/Properties
<b>DEFAULT.GRES</b>	Required Consumable Generic Resources
<b>DEFAULT.MEM</b>	Required Memory/RAM (in MB)
<b>DEFAULT.NODESET</b>	Node Set Specification
<b>DEFAULT.PROC</b>	Required Processor Count
<b>DEFAULT.TPN</b>	Tasks Per Node
<b>DEFAULT.WCLIMIT</b>	Wallclock Limit

**i** Defaults set in a class/queue of the resource manager will override the default values of the corresponding class/queue specified in Moab.

**i** **RESOURCELIMITPOLICY** must be configured in order for the **CLASSCFG** limits to take effect.

*Example 2-9:*

```
CLASSCFG[batch] DEFAULT.DISK=200MB DEFAULT.FEATURES=prod DEFAULT.WCLIMIT=1:00:00
CLASSCFG[debug] DEFAULT.FEATURES=debug DEFAULT.WCLIMIT=00:05:00
```

### 2.7.5.B Per Job Min/Max Limits

Classes can be assigned a minimum and a maximum **job template** that constrains resource requests. Jobs submitted to a particular queue must meet the resource request constraints of these templates. If a job submission exceeds these limits, the entire job submission fails.



Limit	Description
<b>MAX.ARRAYSUBJOBS</b>	Max Allowed Jobs in an Array
<b>MAX.CPUTIME</b>	Max Allowed Utilized CPU Time
<b>MAX.NODE</b>	Max Allowed Node Count
<b>MAX.PROC</b>	Max Allowed Processor Count
<b>MAX.PS</b>	Max Requested Processor-Seconds
<b>MIN.NODE</b>	Min Allowed Node Count
<b>MIN.PROC</b>	Min Allowed Processor Count
<b>MIN.PS</b>	Min Requested Processor-Seconds
<b>MIN.TPN</b>	Min Tasks Per Node
<b>MIN.WCLIMIT</b>	Min Requested Wallclock Limit
<b>MAX.WCLIMIT</b>	Max Requested Wallclock Limit

**i** The parameters listed in the preceding table are for classes and PARCFG only, not users, accounts, groups or QoSes, and they function on a per-job basis. The `MAX.*` and `MIN.*` parameters are different from the `MAXJOB`, `MAXNODE`, and `MAXMEM` parameters described earlier in [Credential Usage Limits](#).

### 2.7.5.C Resource Access

Classes can be associated with a particular set of compute resources. Consequently, jobs submitted to a given class can only use listed resources. This can be handled at the [resource manager](#) level or via the `CLASSCFG HOSTLIST` attribute.

### 2.7.5.D Class Membership Constraints

Classes can be configured at either the resource manager or scheduler level to only allow select users and groups to access them. Jobs that do not meet these criteria are rejected. If specifying class membership/access at the resource manager level, see the respective

resource manager documentation. Moab automatically detects and enforces these constraints. If specifying class membership/access at the scheduler level, use the `REQUIREDUSERLIST` or `EXCLUDEUSERLIST` attributes of the `CLASSCFG` parameter.

**i** Under most resource managers, jobs must always be a member of one and only one class.

### 2.7.5.E Attributes Enabling Class Access to Other Credentials

Classes can be configured to allow jobs to access other credentials such as QoSs and Accounts. This is accomplished using the `QDEF`, `QLIST`, `ADEF`, and `ALIST` attributes.

### 2.7.5.F Special Class Attributes

The class object also possesses a few unique attributes including `JOBPROLOG`, `JOBPILOG`, `RESFAILPOLICY`, and `DISABLEAM` attributes described in what follows.

## MANAGERS

Users listed via the `MANAGERS` parameter are granted full control over all jobs submitted to or running within the specified class.

```
# allow john and steve to cancel and modify all jobs submitted to the class/queue
special
CLASSCFG[special] MANAGERS=john,steve
```

In particular, a class manager can perform the following actions on jobs within a class/queue:

- view/diagnose job (`checkjob`)
- cancel, requeue, suspend, resume, and checkpoint job (`mjobctl`)
- modify job (`mjobctl`)

## JOBPROLOG

The `JOBPROLOG` class performs a function similar to the resource manager level job prolog feature; however, there are some key differences:

- Moab prologs execute on the head node; resource manager prologs execute on the nodes allocated to the job.
- Moab prologs execute as the primary Moab administrator, resource manager prologs execute as root.

- Moab prologs can incorporate cluster environment information into their decisions and actions. See [Valid Variables](#).
- Unique Moab prologs can be specified on a per class basis.
- Job start requests are not sent to the resource manager until the Moab job prolog is successfully completed.
- Error messages generated by a Moab prolog are attached to jobs and associated objects; stderr from prolog script is attached to job.
- Moab prologs have access to Moab internal and peer services.

Valid epilog and prolog variables are:

Variable	Description
<b>\$TIME</b>	Time that the trigger launches
<b>\$HOME</b>	Moab home directory
<b>\$USER</b>	User name the job is running under
<b>\$JOBID</b>	Unique job identifier
<b>\$HOSTLIST</b>	Entire host list for job
<b>\$MASTERHOST</b>	Master host for job

The `JOBPROLOG` class attribute allows a site to specify a unique per-class action to take before a job is allowed to start. This can be used for environmental provisioning, pre-execution resource checking, security management, and other functions. Sample uses may include enabling a VLAN, mounting a global file system, installing a new application or virtual node image, creating dynamic storage partitions, or activating job specific software services.

**i** A prolog is considered to have failed if it returns a negative number. If a prolog fails, the associated job will not start.

**i** If a prolog executes successfully, the associated epilog is guaranteed to start, even if the job fails for any reason. This allows the epilog to undo any changes made to the system by the prolog.

## Job Prolog Examples

```
# explicitly specify prolog arguments for special epilog
CLASSCFG[special] JOBPROLOG='$TOOLS DIR/specialprolog.pl $JOBID $HOSTLIST'
# use default prolog arguments for batch prolog
CLASSCFG[batch]   JOBPROLOG=$TOOLS DIR/batchprolog.pl
```

## JOBEPILOG

The Moab epilog is nearly identical to the prolog in functionality except that it runs after the job completes within the resource manager but before the scheduler releases the allocated resources for use by subsequent jobs. It is commonly used for job clean-up, file transfers, signaling peer services, and undoing other forms of resource customization.

**i** An epilog is considered to have failed if it returns a negative number. If an epilog fails, the associated job will be annotated and a message will be sent to administrators.

## RESFAILPOLICY

This policy allows specification of the action to take on a per-class basis when a failure occurs on a node allocated to an actively running job. See the [Node Availability Overview](#) for more information.

## DISABLEAM

You can disable [allocation management](#) for jobs in specific classes by setting the DISABLEAM class attribute to TRUE. For all jobs outside of the specified classes, allocation enforcement will continue to be enforced.

```
# do not enforce allocations on low priority and debug jobs
CLASSCFG[lowprio]  DISABLEAM=TRUE
CLASSCFG[debug]   DISABLEAM=TRUE
```

### 2.7.5.G Setting Default Classes

In many cases, end-users do not want to be concerned with specifying a job class/queue. This is often handled by defining a default class. Whenever a user does not explicitly submit a job to a particular class, a default class, if specified, is used. In resource managers such as Torque, this can be done at the resource manager level and its impact is transparent to the scheduler. The default class can also be enabled within the scheduler on a per resource manager or per user basis. To set a resource manager default class within Moab, use the DEFAULTCLASS attribute of the [RMCFG](#) parameter. For per user defaults, use the CDEF attribute of the [USERCFG](#) parameter.

## 2.7.5.H Creating a Remap Class

If a single default class is not adequate, Moab provides more flexible options with the `REMAPCLASS` parameter. If this parameter is set and a job is submitted to the remap class, Moab attempts to determine the final class to which a job belongs based on the resources requested. If a remap class is specified, Moab compares the job's requested nodes, processors, memory, and node features with the class's corresponding minimum and maximum resource limits. Classes are searched in the order in which they are defined; when the first match is found, Moab assigns the job to that class.



You should not use remap classes to route jobs to queues/nodes in conjunction with a Torque routing queue. You should select only one of the two methods.

Because Moab remaps at job submission, updates you make to job requirements after submission will not cause any class changes. Moab does not restart the process.



In order to use `REMAPCLASS`, you must specify a `DEFAULTCLASS`. For example:

```
RMCFG[internal] DEFAULTCLASS=batch
```

In the example that follows, a job requesting 4 processors and the node feature `fast` are assigned to the class `quick`:

```
# You must specify a default class in order to use remap classes
RMCFG[internal] DEFAULTCLASS=batch

# Jobs submitted to "batch" should be remapped
REMAPCLASS batch

# Stevens only queue
CLASSCFG[stevens] REQ.FEATURES=stevens REQUIREDUSERLIST=stevens,stevens2

# Special queue for I/O nodes
CLASSCFG[io] MAX.PROC=8 REQ.FEATURES=io

# General access queues
CLASSCFG[quick] MIN.PROC=2 MAX.PROC=8 REQ.FEATURES=fast|short
CLASSCFG[medium] MIN.PROC=2 MAX.PROC=8
CLASSCFG[DEFAULT] MAX.PROC=64
...
```

The following parameters can be used to remap jobs to different classes:

- `MIN.PROC`
- `MAX.PROC`
- `MIN.TPN`
- `MAX.TPN`
- `MIN.WCLIMIT`

- `MAX.WCLIMIT`
- `REQ.FEATURES`
- `REQ.FLAGS=INTERACTIVE`
- `REQUIREDUSERLIST`

If the parameter `REMAPCLASSLIST` is set, then only the listed classes are searched and they are searched in the order specified by this parameter. If none of the listed classes are valid for a particular job, that job retains its original class.

**i** The remap class only works with resource managers that allow dynamic modification of a job's assigned class/queue.

**i** If default credentials are specified on a remap class, a job submitted to that class will inherit those credentials. If the destination class has different default credentials, the new defaults override the original settings. If the destination class does not have default credentials, the job maintains the defaults inherited from the remap class.

### 2.7.5.1 Class Attribute Overview

The following table enumerates the different attributes for `CLASSCFG`.

**i** Setting `DEFAULT.*` on a class does not assign resources or features to that class. Rather, it specifies resources that jobs will inherit when they are submitted to the class without their own resource requests. To configure features, use `NODECFG`.

*CLASSCFG Parameters*

DEFAULT.ATTR	
<b>Format</b>	<ATTRIBUTE>[,<ATTRIBUTE>]...
<b>Description</b>	One or more comma-delimited generic job attributes.
<b>Example</b>	---

DEFAULT.DISK	
<b>Format</b>	<INTEGER>
<b>Description</b>	Default amount of requested disk space.
<b>Example</b>	---

DEFAULT.EXT	
<b>Format</b>	<STRING>
<b>Description</b>	Default job RM extension.
<b>Example</b>	---

DEFAULT.FEATURESDEFAULT.EXT	
<b>Format</b>	Comma-delimited list of features.
<b>Description</b>	Default list of requested <a href="#">node features</a> (a.k.a. node properties). This only applies to compute resource reqs.
<b>Example</b>	---

DEFAULT.GRES	
<b>Format</b>	<STRING>[<COUNT>][,<STRING>[<COUNT>]]...

DEFAULT.GRES	
<b>Description</b>	Default list of per task required <a href="#">consumable generic resources</a> .
<b>Example</b>	<pre>CLASSCFG[viz] DEFAULT.GRES=viz:2</pre>

DEFAULT.MEM	
<b>Format</b>	<INTEGER> (in MB)
<b>Description</b>	Default amount of requested memory.
<b>Example</b>	---

DEFAULT.NODE	
<b>Format</b>	<INTEGER>
<b>Description</b>	Default required node count.
<b>Example</b>	<pre>CLASSCFG[viz] DEFAULT.NODE=5</pre> <p>When a user submits a job to the <code>viz</code> class without a specified node count, the job is assigned 5 nodes.</p>

DEFAULT.NODESET	
<b>Format</b>	<SETTYPE>:<SETATTR>[:<SETLIST>[,<SETLIST>]...]
<b>Description</b>	Default <a href="#">node set</a> .
<b>Example</b>	<pre>CLASSCFG[amd] DEFAULT.NODESET=ONEOF:FEATURE:ATHLON,OPTERON</pre>

DEFAULT.PROC	
<b>Format</b>	<INTEGER>



DEFAULT.PROC	
<b>Description</b>	Default number of requested processors.
<b>Example</b>	---

DEFAULT.TPN	
<b>Format</b>	<INTEGER>
<b>Description</b>	Default number of tasks per node.
<b>Example</b>	---

DEFAULT.WCLIMIT	
<b>Format</b>	<INTEGER>
<b>Description</b>	Default wallclock limit.
<b>Example</b>	---

EXCL.FEATURES	
<b>Format</b>	Comma- or pipe-delimited list of node features.
<b>Description</b>	Set of excluded (disallowed) features. If delimited by commas, reject job if all features are requested; if delimited by the pipe symbol ( ), reject job if at least one feature is requested.
<b>Example</b>	<code>CLASSCFG[intel] EXCL.FEATURES=ATHLON,AMD</code>

EXCL.FLAGS	
<b>Format</b>	Comma-delimited list of <a href="#">job flags</a> .
<b>Description</b>	Set of excluded (disallowed) job flags. Reject job if any listed flags are set.

## EXCL.FLAGS

### Example

```
CLASSCFG [batch] EXCL.FLAGS=INTERACTIVE
```

## EXCLUDEUSERLIST

### Format

Comma-delimited list of users.

### Description

List of users not permitted access to class.

### Example

---

## FLAGS

### Format

NoBackfill

### Description

Disable jobs from this class from backfilling.

### Example

```
CLASSCFG [batch] FLAGS=NoBackfill
```

## FORCENODEACCESSPOLICY

### Format

One of SINGLETASK, SINGLEJOB, SINGLEUSER, or SHARED

### Description

Node access policy associated with queue. If set, this value overrides any per job settings specified by the user at the job level. See [Node Access Policy](#) overview for more information.

### Example

```
CLASSCFG [batch] FORCENODEACCESSPOLICY=SINGLEJOB
```

## FSCAP

### Format

<DOUBLE>[%]

### Description

See [fairshare policies](#) specification.

### Example

---

FSTARGET	
<b>Format</b>	<DOUBLE>[%]
<b>Description</b>	See <a href="#">fairshare policies</a> specification.
<b>Example</b>	---

HOSTLIST	
<b>Format</b>	<a href="#">Host expression</a> , or comma-delimited list of hosts or host ranges.
<b>Description</b>	List of hosts associated with a class. If specified, Moab constrains the availability of a class to only nodes listed in the class host list.
<b>Example</b>	<pre>CLASSCFG[batch] HOSTLIST=r:abs[45-113]</pre>

IGNHOSTLIST	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	If set to TRUE, any job submitted to the class will have its requested hostlist ignored by the scheduler.
<b>Example</b>	<pre>CLASSCFG[batch] IGNHOSTLIST=TRUE</pre>

JOBPILOG	
<b>Format</b>	<STRING>
<b>Description</b>	Scheduler level job epilog to be run after job is completed by resource manager. See <a href="#">special class attributes</a> .
<b>Example</b>	---

JOBFLAGS	
<b>Format</b>	Comma-delimited list of job flags.
<b>Description</b>	See the <a href="#">flag overview</a> for a description of legal flag values.
<b>Example</b>	<code>CLASSCFG[batch] JOBFLAGS=restartable</code>

JOBPROLOG	
<b>Format</b>	<STRING>
<b>Description</b>	Scheduler level job prolog to be run before job is started by resource manager. See <a href="#">special class attributes</a> .
<b>Example</b>	---

JOBTRIGGER	
<b>Format</b>	<STRING>
<b>Description</b>	Job trigger associated with the class. See <a href="#">Job Triggers</a> .
<b>Example</b>	<pre>CLASSCFG[batch] JOBTRIGGER=atype=exec,etype=create,action="/opt/moab/tools/job_trigger.pl"</pre> <p><i>Execute /opt/moab/tools/job_trigger.pl when a job of class batch is created.</i></p>

MANAGERS	
<b>Format</b>	<USER>[,<USER>]...
<b>Description</b>	Users allowed to control, cancel, preempt, and modify jobs within class/queue. See <a href="#">special class attributes</a> .
<b>Example</b>	<code>CLASSCFG[fast] MANAGERS=root,kerry,e43</code>

MAXJOB	
<b>Format</b>	<INTEGER>
<b>Description</b>	Maximum number of active (starting or running) jobs allowed in the class.
<b>Example</b>	---

MAXPROCPERNODE	
<b>Format</b>	<INTEGER>
<b>Description</b>	Maximum number of processors requested per node. May optionally include node names to articulate which nodes have a specific limit.
<b>Example</b>	<pre>CLASSCFG[cpu] MAXPROCPERNODE=20 # When using this class, limit 20 for all nodes</pre> <pre>CLASSCFG[cpu] MAXPROCPERNODE[n1,n2]=20 MAXPROCPERNODE[n3]=10 # When using this class, limit 20 for n1 &amp; n2 and limit 10 for n3</pre> <pre>CLASSCFG[cpu] MAXPROCPERNODE[n1,n2]=20 MAXPROCPERNODE=10 # When using this class, limit 20 for n1 &amp; n2 and limit 10 for all other nodes</pre>

MAX.CPUTIME	
<b>Format</b>	<INTEGER>
<b>Description</b>	Maximum allowed utilized CPU time.
<b>Example</b>	---

MAX.NODE	
<b>Format</b>	<INTEGER>
<b>Description</b>	Maximum number of requested nodes per job. Also used when <a href="#">REMAPCLASS</a> is set to correctly route the job.

MAX.NODE	
<b>Example</b>	<pre>CLASSCFG[batch] MAX.NODE=64</pre> <p>Deny jobs requesting over 64 nodes access to the class <code>batch</code>.</p>

MAX.PROC	
<b>Format</b>	<INTEGER>
<b>Description</b>	<p>Maximum number of requested processors per job. Also used when <a href="#">REMAPCLASS</a> is set to correctly route the job.</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> This enforces the requested processors, not the actual processors dedicated to a job. When enforcing limits for <code>NODEACCESSPOLICY SINGLEJOB</code>, use <code>MAX.NODE</code> instead.</p> </div>
<b>Example</b>	<pre>CLASSCFG[small] MAX.PROC[USER]=3,6</pre>

MAX.PS	
<b>Format</b>	<INTEGER>
<b>Description</b>	Maximum requested processor-seconds.
<b>Example</b>	---

MAX.TPN	
<b>Format</b>	<INTEGER>
<b>Description</b>	Maximum required tasks per node per job. Also used when <a href="#">REMAPCLASS</a> is set to correctly route the job.
<b>Example</b>	---

MAX.WCLIMIT	
<b>Format</b>	[[[DD:]HH:]MM:]SS
<b>Description</b>	Maximum allowed wallclock limit per job. Also used when <a href="#">REMAPCLASS</a> is set to correctly route the job.
<b>Example</b>	<pre>CLASSCFG[long] MAX.WCLIMIT=96:00:00</pre>

MIN.NODE	
<b>Format</b>	<INTEGER>
<b>Description</b>	Minimum number of requested nodes per job. Also used when <a href="#">REMAPCLASS</a> is set to correctly route the job.
<b>Example</b>	<pre>CLASSCFG[dev] MIN.NODE=16</pre> Jobs must request at least 16 nodes to be allowed to access the class.

MIN.PROC	
<b>Format</b>	<INTEGER>
<b>Description</b>	Minimum number of requested processors per job. Also used when <a href="#">REMAPCLASS</a> is set to correctly route the job.
<b>Example</b>	<pre>CLASSCFG[dev] MIN.PROC=32</pre> Jobs must request at least 32 processors to be allowed to access the class.

MIN.PS	
<b>Format</b>	<INTEGER>
<b>Description</b>	Minimum requested processor-seconds.
<b>Example</b>	---

MIN.TPN	
<b>Format</b>	<INTEGER>
<b>Description</b>	Minimum required tasks per node per job. Also used when <a href="#">REMAPCLASS</a> is set to correctly route the job.
<b>Example</b>	---

MIN.WCLIMIT	
<b>Format</b>	[[[DD:]HH:]MM:]SS
<b>Description</b>	Minimum required wallclock limit per job. Also used when <a href="#">REMAPCLASS</a> is set to correctly route the job.
<b>Example</b>	---

NODEACCESSPOLICY	
<b>Format</b>	One of SINGLETASK, SINGLEJOB, SINGLEUSER, or SHARED
<b>Description</b>	Default node access policy associated with queue. This value will be overridden by any per job settings specified by the user at the job level. See <a href="#">Node Access Policy</a> overview.
<b>Example</b>	<pre>CLASSCFG[batch] NODEACCESSPOLICY=SINGLEJOB</pre>

PARTITION	
<b>Format</b>	<STRING>
<b>Description</b>	Partition name where jobs associated with this class must run.
<b>Example</b>	<pre>CLASSCFG[batch] PARTITION=p12</pre>



PRIORITY	
<b>Format</b>	<INTEGER>
<b>Description</b>	Priority associated with the class. See <a href="#">Priority</a> overview.
<b>Example</b>	<code>CLASSCFG[batch] PRIORITY=1000</code>

QDEF	
<b>Format</b>	<QOSID>
<b>Description</b>	<p>Default QoS for jobs submitted to this class. You can specify a maximum of four QDEF entries per credential. Any QoSes specified after the fourth will not be accepted.</p> <div style="border: 1px solid #005596; padding: 5px; background-color: #e6f2ff;"> <p><b>i</b> In addition to classes, you can also specify QDEF for accounts, groups, and users.</p> </div>
<b>Example</b>	<p><code>CLASSCFG[batch] QDEF=base</code></p> <p>Jobs submitted to class <code>batch</code> that do not explicitly request a QoS will have the QoS <code>base</code> assigned.</p>

QLIST	
<b>Format</b>	<QOSID>[,<QOSID>]...
<b>Description</b>	List of accessible QoSs for jobs submitted to this class.
<b>Example</b>	<code>CLASSCFG[batch] QDEF=base QLIST=base,fast,special,bigio</code>

REQ.FEATURES	
<b>Format</b>	Comma- or pipe-delimited list of node features.
<b>Description</b>	Set of required features. If delimited by commas, all features are required; if delimited by the pipe symbol ( ), at least one feature is required.

REQ.FEATURES	
<b>Example</b>	<code>CLASSCFG[amd] REQ.FEATURES=ATHLON,AMD</code>

REQ.FLAGS	
<b>Format</b>	<code>REQ.FLAGS</code> can be used with only the <code>INTERACTIVE</code> flag.
<b>Description</b>	Sets the <code>INTERACTIVE</code> flag on jobs in this class.
<b>Example</b>	<code>CLASSCFG[orion] REQ.FLAGS=INTERACTIVE</code>

REQUIREDACCOUNTLIST	
<b>Format</b>	Comma-delimited list of accounts.
<b>Description</b>	List of accounts allowed to access and use a class (analogous to <code>*LIST</code> for other credentials).
<b>Example</b>	<code>CLASSCFG[jasper] REQUIREDACCOUNTLIST=testers,development</code>

REQUIREDUSERLIST	
<b>Format</b>	Comma-delimited list of users.
<b>Description</b>	List of users allowed to access and use a class (analogous to <code>*LIST</code> for other credentials).
<b>Example</b>	<code>CLASSCFG[jasper] REQUIREDUSERLIST=john,u13,steve,guest</code>

REQUIREDQOSLIST	
<b>Format</b>	Comma-delimited list of QoSs
<b>Description</b>	List of QoSs allowed to access and use a class (analogous to <code>*LIST</code> for other credentials).

REQUIREDQOSLIST	
	<div style="border: 1px solid #004a7c; padding: 5px; background-color: #e6f2ff;"> <p><b>i</b> The number of unique QoSs is limited by the Moab Maximum ACL limit, which defaults to 32.</p> </div>
<b>Example</b>	<pre>CLASSCFG[jasper] REQUIREDQOSLIST=hi,lo</pre>

SYSPRIO	
<b>Format</b>	<INTEGER>
<b>Description</b>	<p>Value of <a href="#">system priority</a> applied to every job submitted to this class.</p> <div style="border: 1px solid #004a7c; padding: 5px; background-color: #e6f2ff;"> <p><b>i</b> Once a system priority has been added to a job, either manually or through configuration, it can only be removed manually.</p> </div>
<b>Example</b>	<pre>CLASSCFG[special] SYSPRIO=100</pre>

WCOVERRUN	
<b>Format</b>	[[[DD:]HH:]MM:]SS
<b>Description</b>	Tolerated amount of time beyond the specified wallclock limit.
<b>Example</b>	---

### 2.7.5.J Enabling Queue Complex Functionality

Queue complexes allow an organization to build a hierarchy of queues and apply certain limits and rules to collections of these queues. Moab supports this functionality in two ways. The first way, queue mapping, is very simple but limited in functionality. The second method provides very rich functionality but requires more extensive configuration using the Moab hierarchical fairshare facility.

#### Queue Mapping

Queue mapping allows collections of queues to be mapped to a parent credential object against which various limits and policies can be applied, as in the following example:

```

QOSCFG[general]   MAXIJOB[USER]=14  PRIORITY=20
QOSCFG[prio]     MAXIJOB[USER]=8   PRIORITY=2000
# group short, med, and long jobs into 'general' QOS
CLASSCFG[short]  QDEF=general FSTARGET=30
CLASSCFG[med]    QDEF=general FSTARGET=40
CLASSCFG[long]   QDEF=general FSTARGET=30 MAXPROC=200
# group interactive and debug jobs into 'prio' QOS
CLASSCFG[inter]  QDEF=prio
CLASSCFG[debug]  QDEF=prio
CLASSCFG[premier] PRIORITY=10000

```

## 2.7.6 QoS Credential

The concept of a quality of service (QoS) credential is unique to Moab and is not derived from any underlying concept or peer service. In most cases, the QoS credential is used to allow a site to set up a selection of service levels for end-users to choose from on a long-term or job-by-job basis. QoSs differ from other credentials in that they are centered around special access where this access may allow use of additional services, additional resources, or improved responsiveness. Unique to this credential, organizations can also choose to apply different charge rates to the varying levels of service available within each QoS. As QoS is an internal credential, all QoS configuration occurs within Moab.

QoS access and QoS defaults can be mapped to users, groups, accounts, and classes, allowing limited service offering for key users. As mentioned, these services focus around increasing access to special scheduling capabilities & additional resources and improving job responsiveness. At a high level, unique QoS attributes can be broken down into the following:

- [Usage Limit Overrides](#)
- [Service Targets](#)
- [Privilege Flags](#)
- [Charge Rate](#)
- [Access Controls](#)

### 2.7.6.A QoS Usage Limit Overrides

All credentials allow specification of job limits. In such cases, jobs are constrained by the most limiting of all applicable policies. With QoS override limits, however, jobs are limited by the override, regardless of other limits specified.

### 2.7.6.B QoS Service Targets

Service targets cause the scheduler to take certain job-related actions as various responsiveness targets are met. Targets can be set for either job queue time or job

expansion factor and cause priority adjustments, reservation enforcement, or preemption activation. In strict service centric organizations, Moab can be configured to trigger various events and notifications in the case of failure by the cluster to meet responsiveness targets.

### 2.7.6.C QoS Privilege Flags

QoSs can provide access to special capabilities. These capabilities include preemption, job deadline support, backfill, next to run priority, guaranteed resource reservation, resource provisioning, dedicated resource access, and many others. See the complete list in the [QoS Facility Overview](#) section.

### 2.7.6.D QoS Charge Rate

Associated with the QoSs many privileges is the ability to assign end-users costs for the use of these services. This charging can be done on a per-QoS basis and can be specified for both dedicated and use-based resource consumption. The [Quality of Service \(QoS\) Facilities](#) section covers more details on QoS level costing configuration while the [Charging and Allocation Management](#) section provides more details regarding general single cluster and multi-cluster charging capabilities.

### 2.7.6.E QoS Access Controls

QoS access control can be enabled on a per QoS basis using the [MEMBERULIST](#) attribute or specified on a *per-requestor* basis using the [QDEF](#) and [QLIST](#) attributes of the [USERCFG](#), [GROUPCFG](#), [ACCOUNTCFG](#), and [CLASSCFG](#) parameters. See [Managing QoS Access](#) for more detail.

---

#### Related Topics

- [Identity Manager Interface](#)
- [Usage Limits](#)

## 2.8 Job Flags

ADVRES

Format

ADVRES[:<RESID>]

ADVRES	
<b>Default</b>	Use available resources where ever found, whether inside a reservation or not.
<b>Description</b>	Specifies the job can only utilize accessible, reserved resources. If <RESID> is specified, only resources in the specified reservation can be utilized.
<b>Example</b>	<pre>FLAGS=ADVRES:META.1</pre> <p><i>The job can only utilize resources located in the META.1 reservation.</i></p>

ALLPROCS	
<b>Format</b>	---
<b>Default</b>	---
<b>Description</b>	<p>Each task should occupy all the processors on the node.</p> <ul style="list-style-type: none"> <li><b>i</b> Incompatible with ppn and non-Torque systems.</li> <li><b>i</b> ALLPROCS is scheduled to be deprecated in a future Moab version where it will be replaced with the new NUMA job submission syntax (place=node in this particular case).</li> </ul>
<b>Example</b>	<pre>msub -l nodes=6 -l flags=allprocs</pre> <p><i>Each of the 6 tasks will occupy all the processors on the node and the job will launch enough processes to occupy each of those processors.</i></p>

ARRAYJOBPARLOCK	
<b>Format</b>	---
<b>Default</b>	---
<b>Description</b>	Specifies that the job array being submitted should not span across multiple partitions. This locks all subjobs of the array to a single partition. If you want to lock all job arrays to a single partition, specify the <a href="#">ARRAYJOBPARLOCK</a> parameter in <code>moab.cfg</code> to force this behavior on a global scale.

## ARRAYJOBPARLOCK

### Example

```
> msub -t moab.[1-5]%3 -l walltime=30,flags=arrayjobparlock
```

## ARRAYJOBPARSPAN

### Format

---

### Default

---

### Description

Specifies that the job array being submitted should span across multiple partitions. This is the default behavior in Moab, unless the [ARRAYJOBPARLOCK](#) parameter is specified in `moab.cfg`. This job flag overrides the `ARRAYJOBPARLOCK` parameter so that job arrays can be allowed to span multiple partitions at submit time.

### Example

```
> msub -t moab.[1-5]%3 -l walltime=30,flags=arrayjobparspan
```

## FORCEPROVISION

### Format

FORCEPROVISION

### Default

---

### Description

A job will provision nodes whether or not they already have the requested OS. When provisioning is enabled (on KNL systems, for example) and this flag is present, the default provisioning behavior (where Moab does not provision a node if the current OS already matches the one being requested) is overridden.

### Example

```
msub -l os=RHEL,flags=forceprovision
```

## GRESONLY

### Format

GRESONLY

### Default

False

### Description

Uses no compute resources such as processors, memory, and so forth; uses

GRESONLY	
	only generic resources.
<b>Example</b>	<pre>&gt; msub -l gres=matlab,walltime=300</pre>

IGNIDLEJOBRSV	
<b>Format</b>	IGNIDLEJOBRSV
<b>Default</b>	N/A
<b>Description</b>	Only applies to QOS. IGNIDLEJOBRSV allows jobs to start without a guaranteed walltime. Instead, it overlaps the idle reservations of real jobs and is preempted 2 minutes before the real job starts.
<b>Example</b>	<pre>QOSCFG[standby] JOBFLAGS=IGNIDLEJOBRSV</pre>

NOQUEUE	
<b>Format</b>	NOQUEUE
<b>Default</b>	Jobs remain queued until they are able to run.
<b>Description</b>	Specifies that the job should be removed if it is unable to allocate resources and start execution immediately.
<b>Example</b>	<pre>FLAGS=NOQUEUE</pre> <p><i>The job should be removed unless it can start running at submit time.</i></p> <p>This functionality is identical to the resource manager extension <a href="#">QUEUEJOB:FALSE</a>.</p>

NORMSTART	
<b>Format:</b>	NORMSTART



NORMSTART	
<b>Default:</b>	Moab passes jobs to a resource manager to schedule.
<b>Description:</b>	Specifies that the job is an internal system job and will not be started via an RM.
<b>Example:</b>	<pre>FLAGS=NORMSTART</pre> <p><i>The job begins running in Moab without a corresponding RM job.</i></p>

PREEMPTEE	
<b>Format:</b>	PREEMPTEE
<b>Default:</b>	Jobs cannot be preempted by other jobs
<b>Description:</b>	Specifies that the job can be <a href="#">preempted</a> by other jobs that have the <code>PREEMPTOR</code> flag set.
<b>Example:</b>	<pre>FLAGS=PREEMPTEE</pre> <p><i>The job can be preempted by other jobs that have the <code>PREEMPTOR</code> flag set.</i></p>

PREEMPTOR	
<b>Format</b>	PREEMPTOR
<b>Default</b>	Jobs cannot preempt other jobs.
<b>Description</b>	Specifies that the job can <a href="#">preempt</a> other jobs that have the <code>PREEMPTEE</code> flag set.
<b>Example</b>	<pre>FLAGS=PREEMPTOR</pre> <p><i>The job may preempt other jobs that have the <code>PREEMPTEE</code> flag set.</i></p>

PURGEONSUCCESSONLY	
<b>Format</b>	PURGEONSUCCESSONLY
<b>Default</b>	Completed jobs are sent to a queue for a short period of time before Moab purges them from the system.
<b>Description</b>	Specifies that Moab should only purge the job from the completed queue if it completed successfully. If the job failed, Moab will keep it in the queue indefinitely to allow you to restart it at any time. This flag is particularly useful for setup and take down jobs in job workflows. See <a href="#">Creating Workflows with Job Templates</a> for more information.
<b>Example</b>	<pre>FLAGS=PURGEONSUCCESSONLY</pre> <p><i>If the job fails, Moab will not purge it from the completed job queue.</i></p>

RESTARTABLE	
<b>Format</b>	RESTARTABLE
<b>Default</b>	Jobs cannot be restarted if preempted.
<b>Description</b>	Specifies jobs can be <i>requeued</i> and later restarted if <a href="#">preempted</a> .
<b>Example</b>	<pre>FLAGS=RESTARTABLE</pre> <p><i>The associated job can be preempted and restarted at a later date.</i></p>

SUSPENDABLE	
<b>Format</b>	SUSPENDABLE
<b>Default</b>	Jobs cannot be suspended if preempted.
<b>Description</b>	Specifies jobs can be <i>suspended</i> and later resumed if <a href="#">preempted</a> .
<b>Example</b>	<pre>FLAGS=SUSPENDABLE</pre>

**SUSPENDABLE**

*The associated job can be suspended and resumed at a later date.*

**SYSTEMJOB**

<b>Format</b>	SYSTEMJOB
<b>Default</b>	N/A
<b>Description</b>	Creates an internal system job that does not require resources.
<b>Example</b>	FLAGS=SYSTEMJOB

**USEMOABJOBID**

<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	<p>Specifies whether to return the Moab job ID when running 'msub', or the resource manager's job ID if it is available.</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> Setting USEMOABJOBID here overrides the global setting for USEMOABJOBID in moabcfg. See <a href="#">USEMOABJOBID</a> for more information.</p> </div>
<b>Example</b>	FLAGS=USEMOABJOBID SELECT=TRUE

**WIDERSVSEARCHALGO**

<b>Format</b>	<BOOLEAN>
<b>Default</b>	---
<b>Description</b>	<p>When Moab is determining when and where a job can run, it either searches for the most resources or the longest range of resources. In almost all cases searching for the longest range is ideal and returns the soonest starttime. In some rare cases, however, a particular job may need to search for the most</p>

**WIDERSVSEARCHALGO**

resources. In those cases this flag can be used to have the job find the soonest starttime. The flag can be specified at submit time, or you can use [mjobctl -m](#) to modify the job after it has been submitted. See the [RSVSEARCHALGO](#) parameter.

**Example**

```
> msub -l flags=widersvsearchalgo
> mjobctl -m flags+=widersvsearchalgo job.1
```

---

**Related Topics**

- [Setting Per-Credential Job Flags](#)

## Chapter 3: Scheduler Commands

In this chapter:

3.1 Moab Command Overview .....	98
3.1.1 Moab Commands .....	98
3.1.2 Moab Command Options .....	100
3.1.3 Commands Providing Maui Compatibility .....	101
3.2 Status Commands .....	101
3.3 Job Management Commands .....	102
3.4 Reservation Management Commands .....	103
3.5 Policy/Configuration Management Commands .....	104
3.6 End-User Commands .....	104
3.7 Moab Commands .....	105
3.7.1 checkjob .....	107
3.7.2 checknode .....	119
3.7.3 mcredctl .....	123
3.7.4 mdiag .....	127
3.7.5 mdiag -a .....	132
3.7.6 mdiag -b .....	133
3.7.7 mdiag -c .....	133
3.7.8 mdiag -f .....	137
3.7.9 mdiag -j .....	139
3.7.10 mdiag -n .....	141
3.7.11 mdiag -p .....	147
3.7.12 mdiag -q .....	150
3.7.13 mdiag -r .....	151
3.7.14 mdiag -R .....	155
3.7.15 mdiag -s .....	156
3.7.16 mdiag -S .....	157
3.7.17 mdiag -t .....	158
3.7.18 mdiag -T .....	159
3.7.19 mdiag -u .....	161
3.7.20 mjobctl .....	162
3.7.21 mnodectl .....	179

3.7.22 moab .....	185
3.7.23 mrmctl .....	186
3.7.24 mrvctl .....	190
3.7.26 mschedctl .....	216
3.7.27 mshow .....	225
3.7.28 mshow -a .....	227
3.7.29 mshow -a (mshow in a Hosting Environment) .....	237
3.7.30 msub .....	239
3.7.31 mvcctl (Moab Virtual Container Control) .....	268
3.7.32 showbf .....	273
3.7.33 showq .....	277
3.7.34 showhist.moab.pl .....	287
3.7.35 showres .....	291
3.7.36 showstart .....	297
3.7.37 showstate .....	300
3.7.38 showstats .....	302
3.7.39 showstats -f .....	315
3.7.40 Deprecated Commands .....	318

## 3.1 Moab Command Overview

In this topic:

- [3.1.1 Moab Commands - page 98](#)
- [3.1.2 Moab Command Options - page 100](#)
- [3.1.3 Commands Providing Maui Compatibility - page 101](#)

### 3.1.1 Moab Commands

Command	Description
<b>checkjob</b>	Provide detailed status report for specified job

Command	Description
<b>checknode</b>	Provide detailed status report for specified node
<b>mcredctl</b>	Controls various aspects about the credential objects within Moab
<b>mdiag</b>	Provide diagnostic reports for resources, workload, and scheduling
<b>mjobctl</b>	Control and modify job
<b>mnodectl</b>	Control and modify nodes
<b>moab</b>	Control the Moab daemon
<b>mrmctl</b>	Query and control resource managers
<b>mrsvctl</b>	Create, control and modify reservations
<b>mschedctl</b>	Modify scheduler state and behavior
<b>mshow</b>	Displays various diagnostic messages about the system and job queues
<b>mshow -a</b>	Query and show available system resources
<b>msub</b>	Scheduler job submission
<b>mvctl</b>	Create, modify, and delete VCs
<b>showbf</b>	Show current resource availability
<b>showhist.moab.pl</b>	Show past job information
<b>showq</b>	Show queued jobs
<b>showres</b>	Show existing reservations
<b>showstart</b>	Show estimates of when job can/will start
<b>showstate</b>	Show current state of resources

Command	Description
<b>showstats</b>	Show usage statistics
<b>showstats -f</b>	Show various tables of scheduling/system performance

### 3.1.2 Moab Command Options

For many Moab commands, you can use the following options to specify that Moab will run the command in a different way or different location from the configured default. These options do not change your settings in the configuration file; they override the settings for this single instance of the command.

Option	Description
<b>--about</b>	Displays build and version information and the status of your Moab license.
<b>--help</b>	Displays usage information about the command.
<b>--host= &lt;serverHostName&gt;</b>	Causes Moab to run the client command on the specified host.
<b>-- loglevel= &lt;logLevel&gt;</b>	Causes Moab to write log information to STDERR as the client command is running. For more information, see <a href="#">Logging Overview</a> .
<b>--msg=&lt;message&gt;</b>	Causes Moab to annotate the action in the <a href="#">event log</a> .
<b>-- port= &lt;serverPort&gt;</b>	Causes Moab to run the command using the port specified.
<b>-- timeout=&lt;seconds&gt;</b>	Sets the maximum time that the client command will wait for a response from the Moab server.
<b>--version</b>	Displays version information.
<b>--xml</b>	Causes Moab to return the command output in XML format.



### 3.1.3 Commands Providing Maui Compatibility



The following commands are deprecated. Click the link for respective deprecated commands to see the updated replacement command for each.

Command	Description
<a href="#">canceljob</a>	Cancel job
<a href="#">changeparam</a>	Change in memory parameter settings
<a href="#">diagnose</a>	Provide diagnostic report for various aspects of resources, workload, and scheduling
<a href="#">releasehold</a>	Release job defers and holds
<a href="#">releaseres</a>	Release reservations
<a href="#">runjob</a>	Force a job to run immediately
<a href="#">sethold</a>	Set job holds
<a href="#">setqos</a>	Modify job QOS settings
<a href="#">setres</a>	Set an admin/user reservation
<a href="#">setspri</a>	Adjust job/system priority of job
<a href="#">showconfig</a>	Show current scheduler configuration

## 3.2 Status Commands

The status commands organize and present information about the current state and historical statistics of the scheduler, jobs, resources, users, and accounts. The following table presents the primary status commands and flags:

Command	Description
<b>checkjob</b>	Displays detailed job information such as job state, resource requirements, environment, constraints, credentials, history, allocated resources, and resource utilization.
<b>checknode</b>	Displays detailed node information such as node state, resources, attributes, reservations, history, and statistics.
<b>mdiag -f</b>	Displays summarized fairshare information and any unexpected fairshare configuration.
<b>mdiag -j</b>	Displays summarized job information and any unexpected job state.
<b>mdiag -n</b>	Displays summarized node information and any unexpected node state.
<b>mdiag -p</b>	Displays summarized job priority information.
<b>mschedctl -f</b>	Resets internal statistics.
<b>showstats -f</b>	Displays various aspects of scheduling performance across a job duration/job size matrix.
<b>showq [-r -i]</b>	Displays various views of currently queued active, idle, and non-eligible jobs.
<b>showstats -g</b>	Displays current and historical usage on a per group basis.
<b>showstats -u</b>	Displays current and historical usage on a per user basis.
<b>showstats -v</b>	Displays high level current and historical scheduling statistics.

## 3.3 Job Management Commands

Moab shares job management tasks with the resource manager. Typically, the scheduler only modifies scheduling relevant aspects of the job such as partition access, job priority, charge account, and hold state. The following table covers the available job management commands. The [Commands Overview](#) lists all available commands.

Command	Description
<b>canceljob</b>	Cancels existing job.
<b>checkjob</b>	Displays job state, resource requirements, environment, constraints, credentials, history, allocated resources, and resource utilization.
<b>mdiag -j</b>	Displays summarized job information and any unexpected job state.
<b>releasehold -a</b>	Removes job holds or deferrals.
<b>runjob</b>	Starts job immediately, if possible.
<b>sethold</b>	Sets hold on job.
<b>setqos</b>	Sets/modifies QoS of existing job.
<b>setspri</b>	Adjusts job/system priority of job.

---

### Related Topics

- [Job State Definitions](#)

## 3.4 Reservation Management Commands

Moab exclusively controls and manages all advance reservation features including both standing and administrative reservations. The following table covers the available reservation management commands:

Command	Description
<b>mdiag -r</b>	Displays summarized reservation information and any unexpected state.
<b>mrsvctl</b>	Reservation control.
<b>mrsvctl -r</b>	Removes reservations.
<b>mrsvctl -c</b>	Creates an administrative reservation.

Command	Description
<b>showres</b>	Displays information regarding location and state of reservations.

## 3.5 Policy/Configuration Management Commands

Moab allows dynamic modification of most scheduling parameters allowing new scheduling policies, algorithms, constraints, and permissions to be set at any time. Changes made via Moab client commands are temporary and are overridden by values specified in Moab configuration files the next time Moab is shut down and restarted. The following table covers the available configuration management commands:

Command	Description
<b>mschedctl -l</b>	Displays triggers, messages, and settings of all configuration parameters.
<b>mschedctl</b>	Controls the scheduler (behavior, parameters, triggers, messages).
<b>mschedctl -m</b>	Modifies system values.

## 3.6 End-User Commands

While the majority of Moab commands are tailored for use by system administrators, a number of commands are designed to extend the knowledge and capabilities of end-users. The following table covers the commands available to end-users.

**i** When using Active Directory as a central authentication mechanism, all nodes must be reported with a different name when booted in both Linux and Windows (for instance, `node01-1` for Linux and `node01` for Windows). If a machine account with the same name is created for each OS, the most recent OS will remove the previously-joined machine account. The nodes must report to Moab with the same hostname. This can be done by using aliases (adding all node names to the `/etc/hosts` file on the system where Moab is running) and ensuring that the Linux resource manager reports the node with its global name rather than the Linux-specific one (`node01` rather than `node01-1`).

Command	Description
<b>canceljob</b>	Cancels existing job.
<b>checkjob</b>	Displays job state, resource requirements, environment, constraints, credentials, history, allocated resources, and resource utilization.
<b>msub</b>	Submit a new job.
<b>releaseres</b>	Releases a <a href="#">user reservation</a> .
<b>setres</b>	Create a <a href="#">user reservation</a> .
<b>showbf</b>	Shows resource availability for jobs with specific resource requirements.
<b>showq</b>	Displays detailed prioritized list of active and idle jobs.
<b>showstart</b>	Shows estimated start time of idle jobs.
<b>showstats</b>	Shows detailed usage statistics for users, groups, and accounts, to which the end-user has access.

---

### Related Topics

- [Commands Overview](#)

## 3.7 Moab Commands

See the Moab commands below.

**checkjob**

**checknode**

**mcredctl**

**mdiag**

<b>mjobctl</b>
<b>mnodectl</b>
<b>moab</b>
<b>mrnctl</b>
<b>mrsvctl</b>
<b>mschedctl</b>
<b>mshow</b>
<b>mshow -a</b>
<b>msub</b>
<b>mvcctl</b>
<b>showbf</b>
<b>showhist.moab.pl</b>
<b>showq</b>
<b>showres</b>
<b>showstart</b>
<b>showstate</b>
<b>showstats</b>
<b>showstats -f</b>

## 3.7.1 checkjob

### 3.7.1.A Synopsis

```
checkjob [exact:jobid] [jobname:jobname] [-l policylevel] [-n nodeid] [-q qosid] [-r
reservationid] [-v] [--flags=future | complete] [--blocking] jobid [--about] [--help]
[--host=<serverHostName>] [--loglevel=<logLevel>] [--
msg=<message>] [--port=<serverPort>] [--timeout=<seconds>] [--
version] [--xml]
```

### 3.7.1.B Overview

`checkjob` displays detailed job [state](#) information and diagnostic output for a specified job. Detailed information is available for queued, blocked, active, and recently completed jobs. The `checkjob` command shows the master job of an array, as well as a summary of array subjobs but does not display all subjobs. Use `checkjob -v` to display all job-array subjobs.

### 3.7.1.C Access

This command can be run by level 1-3 Moab administrators for any job. Also, end users can use `checkjob` to view the status of their own jobs.

### 3.7.1.D Options

#### --blocking

**Format**      `--blocking`

**Description**      Do not use cache information in the output. The `--blocking` flag retrieves results exclusively from the scheduler.

**Example**

```
> checkjob -v --blocking 1234
```

```
Display real time data about job 1234.
```

#### --flags

**Format**      `--flags=future | complete`

--flags	
<b>Description</b>	<ul style="list-style-type: none"> <li>• <code>future</code> - Evaluates future eligibility of job (ignore current resource state and usage limitations).</li> <li>• <code>complete</code> - Queries details for jobs that have already terminated.</li> </ul>
<b>Example</b>	<pre>&gt; checkjob -v --flags=future 6235</pre> <p><i>Display reasons why idle job is blocked ignoring node state and current node utilization constraints.</i></p>

exact	
<b>Format</b>	<code>exact:&lt;JOBID&gt;</code>
<b>Description</b>	Searches for and returns the exact job ID.
<b>Example</b>	<pre>&gt; checkjob exact:1.job_dependency1</pre>

jobname	
<b>Format</b>	<code>jobname:&lt;JOBNAME&gt;</code>
<b>Description</b>	Searches for and returns the first job with the matching <code>&lt;JOBNAME&gt;</code> .
<b>Example</b>	<pre>&gt; checkjob jobname:STEP4</pre>

-l (Policy level)	
<b>Format</b>	<code>&lt;POLICYLEVEL&gt;</code> HARD, SOFT, or OFF
<b>Description</b>	Reports job start eligibility subject to specified throttling policy level.
<b>Example</b>	<pre>&gt; checkjob -l SOFT 6235 &gt; checkjob -l HARD 6235</pre>



<b>-n (NodeID)</b>	
<b>Format</b>	<NODEID>
<b>Description</b>	Checks job access to specified node and <a href="#">preemption</a> status with regards to jobs located on that node.
<b>Example</b>	<pre>&gt; checkjob -n node113 6235</pre>

<b>-q (QoS)</b>	
<b>Format</b>	<QOSID>
<b>Description</b>	Checks job access to specified QoS <QOSID>.
<b>Example</b>	<pre>&gt; checkjob -q special 6235</pre>

<b>-r (Reservation)</b>	
<b>Format</b>	<RSVID>
<b>Description</b>	Checks job access to specified reservation <RSVID>.
<b>Example</b>	<pre>&gt; checkjob -r orion.1 6235</pre>

<b>-v (Verbose)</b>	
<b>Description</b>	<p>Sets verbose mode. If the job is part of an array, the <code>-v</code> option shows pertinent array information before the job-specific information (see <a href="#">Example 2</a> and <a href="#">Example 3</a> for differences between standard output and <code>-v</code> output).</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> Specifying the double verbose (<code>-v -v</code>) displays additional information about the job. See <a href="#">the Output table</a> for details.</p> </div>
<b>Example</b>	<pre>&gt; checkjob -v 6235</pre>

### 3.7.1.E Details

This command allows any Moab administrator to check the detailed status and resource requirements of an active, queued, or recently **completed** job. Additionally, this command performs numerous diagnostic checks and determines if and where the job could potentially run. Diagnostic checks include **policy** violations, reservation constraints, preemption status, and job to resource mapping. If a job cannot run, a text reason is provided along with a summary of how many nodes are and are not available. If the `-v` flag is specified, a node by node summary of resource availability will be displayed for idle jobs.

#### Job Eligibility

If a job cannot run, a text reason is provided along with a summary of how many nodes are and are not available. If the `-v` flag is specified, a node by node summary of resource availability will be displayed for idle jobs. For job level eligibility issues, one of the following reasons will be given:

Reason	Description
<b>Job has hold in place</b>	One or more job holds are currently in place
<b>Insufficient idle procs</b>	There are currently not adequate processor resources available to start the job
<b>Idle procs do not meet requirements</b>	adequate idle processors are available but these do not meet job requirements
<b>Start date not reached</b>	Job has specified a minimum <i>start date</i> that is still in the future
<b>Expected state is not idle</b>	Job is in an unexpected state
<b>State is not idle</b>	Job is not in the idle state
<b>Dependency is not met</b>	Job depends on another job reaching a certain state
<b>Rejected by policy</b>	Job start is prevented by a throttling policy

If a job cannot run on a particular node, one of the following 'per node' reasons will be given:

Reason	Description
<b>Class</b>	Node does not allow required job class/queue
<b>CPU</b>	Node does not possess required processors
<b>Disk</b>	Node does not possess required local disk
<b>Features</b>	Node does not possess required node features
<b>Memory</b>	Node does not possess required real memory
<b>Network</b>	Node does not possess required network interface
<b>State</b>	Node is not Idle or Running

## Reservation Access

The `-r` flag can be used to provide detailed information about job access to a specific reservation

## Preemption Status

If a job is marked as a `preemptor` and the `-v` and `-n` flags are specified, `checkjob` will perform a job by job analysis for all jobs on the specified node to determine if they can be preempted.

### 3.7.1.F Output

The `checkjob` command displays the following job attributes:

Attribute	Value	Description
<b>Account</b>	<STRING>	Name of account associated with job.
<b>Allocated Nodes</b>	Square bracket delimited list of node and processor IDs	List of nodes and processors allocated to job.
<b>Applied Nodeset**</b>	<STRING>	Nodeset used for job's node allocation.

Attribute	Value	Description
<b>Arch</b>	<STRING>	Node architecture required by job.
<b>Attr</b>	square bracket delimited list of job attributes	Job Attributes (i.e., [BACKFILL] [PREEMPTEE]).
<b>Available Memory**</b>	<INTEGER>	The available memory requested by job. Moab displays the relative or exact value by returning a comparison symbol (>, <, >=, <=, or ==) with the value (i.e., Available Memory <= 2048).
<b>Available Swap**</b>	<INTEGER>	The available swap requested by job. Moab displays the relative or exact value by returning a comparison symbol (>, <, >=, <=, or ==) with the value (i.e., Available Swap >= 1024).
<b>Average Utilized Procs*</b>	<FLOAT>	Average load balance for a job.
<b>Avg Util Resources Per Task*</b>	<FLOAT>	
<b>BecameEligible</b>	<TIMESTAMP>	The date and time when the job moved from Blocked to Eligible.
<b>Bypass</b>	<INTEGER>	Number of times a lower priority job with a later submit time ran before the job.
<b>CheckpointStartTime**</b>	[ [ [DD:] HH:] MM:] SS	The time the job was first checkpointed.
<b>Class</b>	[<CLASS NAME> <CLASS COUNT>]	Name of class/queue required by job and number of class initiators required per task.
<b>Dedicated Resources Per Task*</b>	Space-delimited list of <STRING> ;<INTEGER>	Resources dedicated to a job on a per-task basis.

Attribute	Value	Description
<b>Disk</b>	<INTEGER>	Amount of local disk required by job (in MB).
<b>Estimated Walltime</b>	[[[DD:]HH:]MM:]SS	The scheduler's estimated walltime. <div style="border: 1px solid black; padding: 5px; margin-top: 10px;">  In simulation mode, it is the actual walltime. </div>
<b>EnvVariables**</b>	Comma-delimited list of <STRING>	List of environment variables assigned to job.
<b>Exec Size*</b>	<INTEGER>	Size of job executable (in MB).
<b>Executable</b>	<STRING>	Name of command to run.
<b>Features</b>	Square bracket delimited list of <STRING> <sub>s</sub>	Node features required by job.
<b>Flags</b>		
<b>Group</b>	<STRING>	Name of UNIX group associated with job.
<b>Holds</b>	Zero or more of User, System, and Batch	Types of job holds currently applied to job.
<b>Image Size</b>	<INTEGER>	Size of job data (in MB).
<b>IWD (Initial Working Directory)</b>	<DIR>	Directory to run the executable in.
<b>Job Messages**</b>	<STRING>	Messages attached to a job.
<b>Job Submission**</b>	<STRING>	Job script submitted to RM.
<b>Memory</b>	<INTEGER>	Amount of real memory required per node (in MB).
<b>Max Util Resources Per Task*</b>	<FLOAT>	

Attribute	Value	Description
<b>NodeAccess*</b>		
<b>Nodecount</b>	<INTEGER>	Number of nodes required by job.
<b>Opsys</b>	<STRING>	Node operating system required by job.
<b>Partition Mask</b>	ALL or colon delimited list of partitions	List of <a href="#">partitions</a> the job has access to.
<b>PE</b>	<FLOAT>	Number of processor-equivalents requested by job.
<b>Per Partition Priority**</b>	Tabular	Table showing job template priority for each partition.
<b>Priority Analysis**</b>	Tabular	Table showing how job's priority was calculated: Job PRIORITY* Cred (User:Group:Class) Serv (QTime)
<b>QOS</b>	<STRING>	Quality of Service associated with job.
<b>Reservation</b>	<RSVID> (<TIME1> - <TIME2> Duration: <TIME3>)	RESID specifies the reservation ID, TIME1 is the relative start time, TIME2 the relative end time, TIME3 the duration of the reservation.
<b>Req</b>	[<INTEGER>] TaskCount: <INTEGER> Partition: <partition>	A <a href="#">job requirement</a> for a single type of resource followed by the number of tasks instances required and the appropriate <a href="#">partition</a> .
<b>StageIn</b>	<SOURCE> %<DESTINATION>	The <SOURCE> is the username, hostname, directory and file name of origin for the file(s) that Moab will stage in for this job. The <DESTINATION> is the username, hostname, directory and file name where Moab will place the file during this job. See <a href="#">Data Staging Example</a> for more information.

Attribute	Value	Description
<b>StageInSize</b>	<INTEGER><UNIT>	The size of the file Moab will stage in for this job. <UNIT> can be KB, MB, GB, or TB. See <a href="#">Data Staging Example</a> for more information.
<b>StageOut</b>	<SOURCE> %<DESTINATION>	The <SOURCE> is the username, hostname, directory and file name of origin for the file(s) that Moab will stage out for this job. The <DESTINATION> is the username, hostname, directory and file name where Moab will place the file during this job. See <a href="#">Data Staging Example</a> for more information.
<b>StageOutSize</b>	<INTEGER><UNIT>	The size of the file Moab will stage out for this job. <UNIT> can be KB, MB, GB, or TB. See <a href="#">Data Staging Example</a> for more information.
<b>StartCount</b>	<INTEGER>	Number of times job has been started by Moab.
<b>StartPriority</b>	<INTEGER>	Start priority of job.
<b>StartTime</b>	<TIME>	Time job was started by the resource management system.
<b>State</b>	One of Idle, Starting, Running, etc. See <a href="#">Job States</a> for all possible values.	Current Job State.
<b>SubmitTime</b>	<TIME>	Time job was submitted to resource management system.
<b>Swap</b>	<INTEGER>	Amount of swap disk required by job (in MB).
<b>Task Distribution*</b>	Square bracket delimited list of nodes	

Attribute	Value	Description
<b>Time Queued</b>		
<b>Total Requested Nodes**</b>	<INTEGER>	Number of nodes the job requested.
<b>Total Requested Tasks</b>	<INTEGER>	Number of tasks requested by job.
<b>User</b>	<STRING>	Name of user submitting job.
<b>Utilized Resources Per Task*</b>	<FLOAT>	
<b>WallTime</b>	[[[DD:]HH:]MM:]SS of [[[DD:]HH:]MM:]SS	Length of time job has been running out of the specified limit.

In the above table, fields marked with an asterisk (\*) are only displayed when set or when the `-v` flag is specified. Fields marked with two asterisks (\*\*) are only displayed when set or when the `-v -v` flag is specified.

#### Example 3-1: checkjob 717

```
> checkjob 717
job 717
State: Idle
Creds: user:jacksond group:jacksond class:batch
WallTime: 00:00:00 of 00:01:40
SubmitTime: Mon Aug 15 20:49:41
  (Time Queued Total: 3:12:23:13 Eligible: 3:12:23:11)
TerminationDate: INFINITY Sat Oct 24 06:26:40
Total Tasks: 1
Req[0] TaskCount: 1 Partition: ALL
Network: --- Memory >= 0 Disk >= 0 Swap >= 0
Opsys: --- Arch: --- Features: ---

IWD:          /home/jacksond/moab/moab-4.2.3
Executable:   STDIN
Flags:        RESTARTABLE,NORMSTART
StartPriority: 5063
Reservation '717' ( INFINITY -> INFINITY Duration: 00:01:40)
Note: job cannot run in partition base (idle procs do not meet requirements : 0 of 1
procs found)
idle procs:   4 feasible procs:   0
Rejection Reasons: [State      : 3][ReserveTime : 1]
cannot select job 717 for partition GM (partition GM does not support requested class
batch)
```

*The example job cannot be started for two different reasons.*

- *It is temporarily blocked from partition base because of node state and node reservation*



*conflicts.*

- *It is permanently blocked from partition GM because the requested class batch is not supported in that partition.*

**Example 3-2: Using `checkjob` (no `-v`) on a job array master job:**

```
checkjob array.1
job array.1

AName: array
Job Array Info:
  Name: array.1

Sub-jobs:      10
Active:        6 ( 60.0%)
Eligible:      2 ( 20.0%)
Blocked:       2 ( 20.0%)
Complete:      0 (  0.0%)
```

**Example 3-3: Using `checkjob -v` on a job array master job:**

```
$ checkjob -v array.1
job array.1

AName: array
Job Array Info:
  Name: array.1
  1 : array.1.1 : Running
  2 : array.1.2 : Running
  3 : array.1.3 : Running
  4 : array.1.4 : Running
  5 : array.1.5 : Running
  6 : array.1.6 : Running
  7 : array.1.7 : Idle
  8 : array.1.8 : Idle
  9 : array.1.9 : Blocked
 10 : array.1.10 : Blocked

Sub-jobs:      10
Active:        6 ( 60.0%)
Eligible:      2 ( 20.0%)
Blocked:       2 ( 20.0%)
Complete:      0 (  0.0%)
```

**Example 3-4: Using `checkjob -v` on a data staging job**

```
$ checkjob -v moab.14.dsin
job moab.14.dsin

AName: moab.14.dsin
State: Running
Creds: user:fred group:company
WallTime: 00:00:00 of 00:01:01
SubmitTime: Wed Apr 16 10:07:19
  (Time Queued Total: 00:00:00 Eligible: 00:00:00)

StartTime: Wed Apr 16 10:07:19
```

```

TemplateSets: dsin
Triggers: 78$start+0@0.000000:exec@/opt/moab/tools/datastaging/ds_move_rsync --
stagein:FALSE
Total Requested Tasks: 1

Req[0] TaskCount: 1 Partition: SHARED
Dedicated Resources Per Task: bandwidth: 1
NodeAccess: SHARED

Allocated Nodes:
[GLOBAL:1]

Job Group: moab.14
SystemID: moab
SystemJID: moab.14.dsin
Task Distribution: GLOBAL
IWD: $HOME/test/datastaging
SubmitDir: $HOME/test/datastaging
StartCount: 1
Parent VCs: vc11
User Specified Partition List: local
Partition List: local
SrcRM: internal
Flags: NORMSTART,GRESONLY,TEMPLATESAPPLIED
Attr: dsin
StageInSize: 386MB
StageOutSize: 100MB
StageIn: fred@remotelab:/home/fred/input1/%fred@scratch:/home/fred/input1/
StageIn: fred@remotelab:/home/fred/input2/%fred@scratch:/home/fred/input2/
StageIn: fred@remotelab:/home/fred/input3/%fred@scratch:/home/fred/input3/
StageOut: fred@scratch:/home/fred/output/%fred@remotelab:/home/fred/output/
StartPriority: 1
  SJob Type: datastaging
  Completion Policy: datastaging
PE: 0.00
Reservation 'moab.14.dsin' (-00:00:06 -> 00:00:55 Duration: 00:01:01)

```

---

## Related Topics

- [showhist.moab.pl](#) - explains how to query for past job information
- (Optional) Install Moab Client - in the *Moab HPC Suite Installation and Configuration Guide* explains how to distribute this command to client nodes
- [mdiag -j](#) command - display additional detailed information regarding jobs
- [showq](#) command - showq high-level job summaries
- [JOBPURGETIME](#) parameter - specify how long information regarding completed jobs is maintained
- diagnosing job [preemption](#)

## 3.7.2 checknode

### 3.7.2.A Synopsis

```
checknode flags [nodeID | ALL] [--about] [--help] [--
host=<serverHostName>] [--loglevel=<logLevel>] [--
msg=<message>] [--port=<serverPort>] [--timeout=<seconds>] [--
version] [--xml]
```

### 3.7.2.B Overview

This command shows detailed state information and statistics for nodes that run jobs.

The following information is returned by this command:

Name	Description
<b>ACL</b>	Node Access Control List (displayed only if set)
<b>ActiveTime</b>	Total time node has been busy (allocated to active jobs) since statistics initialization expressed in HH:MM:SS notation (percent of time busy: BusyTime/TotalTime)
<b>Adapters</b>	Network adapters available
<b>Arch</b>	Architecture
<b>Classes</b>	Classes available
<b>Disk</b>	Disk space available
<b>Downtime</b>	Displayed only if downtime is scheduled
<b>EffNodeAccessPolicy</b>	Configured effective node access policy
<b>Features</b>	Features available
<b>Load</b>	CPU Load (Berkley one-minute load average)
<b>Memory</b>	Memory available

Name	Description
<b>Opsys</b>	Operating system
<b>RequestID</b>	Dynamic Node RequestID set by the RM (displayed only if set)
<b>State</b>	Node state
<b>StateTime</b>	Time node has been in current state in HH:MM:SS notation
<b>Swap</b>	Swap space available
<b>TotalTime</b>	Total time node has been detected since statistics initialization expressed in HH:MM:SS notation
<b>TTL</b>	Dynamic Node Time To Live set by the RM (expiration date, displayed only if set)
<b>UpTime</b>	Total time node has been in an available (Non-Down) state since statistics initialization expressed in HH:MM:SS notation (percent of time up: UpTime/TotalTime)

After displaying this information, some analysis is performed and any unusual conditions are reported.

### 3.7.2.C Access

By default, this command can be run by any Moab Administrator (see [ADMINCFG](#)).

### 3.7.2.D Parameters

Name	Description
<b>NODE</b>	Node name you want to check. Moab uses regular expressions to return any node that contains the provided argument. For example, if you ran <code>checknode node1</code> , Moab would return information about <code>node1</code> , <code>node10</code> , <code>node100</code> , etc. If you want to limit the results to <code>node1</code> only, you would run <code>checknode "^node1\$"</code> .

### 3.7.2.E Flags

Name	Description
<b>ALL</b>	Returns checknode output on all nodes in the cluster.
<b>-h</b>	Help for this command.
<b>-v</b>	Returns verbose output.
<b>--xml</b>	Output in XML format. Same as <a href="#">mdiag -n --xml</a> .

#### Example 3-5: checknode

```
> checknode P690-032
node P690-032

State:      Busy (in current state for 11:31:10)
Configured Resources: PROCS: 1 MEM: 16G SWAP: 2000M DISK: 500G
Utilized Resources: PROCS: 1
Dedicated Resources: PROCS: 1
Opsys:      AIX      Arch:      P690
Speed:      1.00     CPULoad:  1.000
Network:    InfiniBand, Myrinet
Features:   Myrinet
Attributes: [Batch]
Classes:    [batch]

Total Time: 5:23:28:36 Up: 5:23:28:36 (100.00%) Active: 5:19:44:22 (97.40%)

Reservations:
  Job '13678'(x1) 10:16:12:22 -> 12:16:12:22 (2:00:00:00)
  Job '13186'(x1) -11:31:10 -> 1:12:28:50 (2:00:00:00)
Jobs: 13186
```

#### Example 3-6: checknode ALL

```
> checknode ALL
node ahe

State:      Idle (in current state for 00:00:30)
Configured Resources: PROCS: 12 MEM: 8004M SWAP: 26G DISK: 1M
Utilized Resources: PROCS: 1 SWAP: 4106M
Dedicated Resources: ---
  MTBF(longterm): INFINITY MTBF(24h): INFINITY
Opsys:      linux     Arch:      ---
Speed:      1.00     CPULoad:  1.400
Flags:      rmdetected
Classes:    [batch]
RM[ahe]*:   TYPE=PBS
EffNodeAccessPolicy: SHARED

Total Time: 00:01:44 Up: 00:01:44 (100.00%) Active: 00:00:00 (0.00%)
```

```

Reservations: ---
node ahe-ubuntu32

State:   Running (in current state for 00:00:05)
Configured Resources: PROCS: 12 MEM: 2013M SWAP: 3405M DISK: 1M
Utilized Resources: PROCS: 6 SWAP: 55M
Dedicated Resources: PROCS: 6
  MTBF(longterm): INFINITY MTBF(24h): INFINITY
Opsys:   linux Arch: ---
Speed:   1.00 CPUload: 2.000
Flags:   rmdetected
Classes: [batch]
RM[ahe]*: TYPE=PBS
EffNodeAccessPolicy: SHARED

Total Time: 00:01:44 Up: 00:01:44 (100.00%) Active: 00:00:02 (1.92%)

Reservations:
  6x2 Job:Running -00:00:07 -> 00:01:53 (00:02:00)
  7x2 Job:Running -00:00:06 -> 00:01:54 (00:02:00)
  8x2 Job:Running -00:00:05 -> 00:01:55 (00:02:00)
Jobs:    6,7,8
node ahe-ubuntu64

State:   Busy (in current state for 00:00:06)
Configured Resources: PROCS: 12 MEM: 2008M SWAP: 3317M DISK: 1M
Utilized Resources: PROCS: 12 SWAP: 359M
Dedicated Resources: PROCS: 12
  MTBF(longterm): INFINITY MTBF(24h): INFINITY
Opsys:   linux Arch: ---
Speed:   1.00 CPUload: 0.000
Flags:   rmdetected
Classes: [batch]
RM[ahe]*: TYPE=PBS
EffNodeAccessPolicy: SHARED

Total Time: 00:01:44 Up: 00:01:44 (100.00%) Active: 00:00:55 (52.88%)

Reservations:
  0x2 Job:Running -00:01:10 -> 00:00:50 (00:02:00)
  1x2 Job:Running -00:00:20 -> 00:01:40 (00:02:00)
  2x2 Job:Running -00:00:20 -> 00:01:40 (00:02:00)
  3x2 Job:Running -00:00:17 -> 00:01:43 (00:02:00)
  4x2 Job:Running -00:00:13 -> 00:01:47 (00:02:00)
  5x2 Job:Running -00:00:07 -> 00:01:53 (00:02:00)
Jobs:    0,1,2,3,4,5
ALERT: node is in state Busy but load is low (0.000)

```

**Example 3-7: checknode n001 (Dynamic Node)**

```

> checknode node001
node node001

State:   Idle (in current state for 00:13:50)
Configured Resources: PROCS: 2 MEM: 4096M
Utilized Resources: PROCS: 2
Dedicated Resources: ---
ACL:     USER==FRED+==BOB+ GROUP==DEV+
  MTBF(longterm): INFINITY MTBF(24h): INFINITY
Opsys:   --- Arch: ---
Speed:   1.00 CPUload: 2.000

```

```

Partition: local Rack/Slot: --- NodeIndex: 1
RM[local]*: TYPE=NATIVE:AGFULL
EffNodeAccessPolicy: SHARED
RequestID: 1234
TTL: Tue Nov 10 00:00:00 2022
Total Time: 2:21:19:05 Up: 2:21:19:05 (100.00%) Active: 00:00:00 (0.00%)

Reservations:
  node001-TTL-1234x1 User      441days -> INFINITY ( INFINITY)
    Blocked Resources@ 441days Procs: 2/2 (100.00%) Mem: 4096/4096 (100.00%)
  Swap: 1/1 (100.00%) Disk: 1/1 (100.00%)
  ALERT: node is in state Idle but load is high (2.000)

```

## Related Topics

- (Optional) Install Moab Client - in the *Moab HPC Suite Installation and Configuration Guide* explains how to distribute this command to client nodes
- [mdiag -n](#)
- [showstate](#)

## 3.7.3 mcredctl

### 3.7.3.A Synopsis

```

mcredctl [-d credtype[:credid]] [-h credtype:credid] [-l credtype] [-q
{role|limit|profile|accessfrom|accessto|policies} credtype[:credid]] [--format=xml] [-r
{stats|fairshare|uid} <type>[:<ID>] [-t <STARTTIME>[,<ENDTIME>] [--about] [--help]
[--host=<serverHostName>] [--loglevel=<logLevel>] [--msg=<message>]
[--port=<serverPort>] [--timeout=<seconds>] [--version] [--xml]

```

### 3.7.3.B Overview

The `mcredctl` command controls various aspects about the credential objects within Moab. It can be used to display configuration, limits, roles, and relationships for various Moab credential objects.

**i** If using Insight, you must restart Moab to view credential modifications.

### 3.7.3.C Options

**i** In all cases `<CREDTYPE>` is one of `acct`, `group`, `user`, `class`, or `qos`.

**i** In most cases it is necessary to use the `--format=xml` flag in order to print the output (see examples below for specific syntax requirements).

#### -d - DESTROY

<b>Format</b>	<code>&lt;TYPE&gt;:&lt;VAL&gt;</code>
<b>Description</b>	Purge a credential from <code>moab.cfg</code> (does not delete credential from memory).
<b>Example</b>	<pre>&gt; mcredctl -d user:john</pre> <p><i>All references to <code>USERCFG[john]</code> will be commented out of <code>moab.cfg</code></i></p>

#### -h - HOLD


<b>Format</b>	<code>&lt;TYPE&gt;:&lt;VAL&gt;</code>
<b>Description</b>	Toggles whether a given credential's jobs should be placed on hold or not.
<b>Example</b>	<pre>&gt; mcredctl -h user:john</pre> <p><i>User <code>[john]</code> will be put on hold.</i></p>

#### -l - LIST

<b>Format</b>	<code>&lt;TYPE&gt;</code>
<b>Description</b>	List the various sub-objects of the specified credential.
<b>Example</b>	<pre>&gt; mcredctl -l user --format=xml</pre> <p><i>List all users within Moab in XML.</i></p> <pre>&gt; mcredctl -l group --format=xml</pre> <p><i>List all groups within Moab in XML.</i></p>



<b>-q - QUERY</b>	
<b>Format</b>	<pre>{role   accessfrom   accessto   limit  profile   policies} limit &lt;TYPE&gt; policies &lt;TYPE&gt; role &lt;USER&gt;:&lt;USERID&gt; profile &lt;TYPE&gt;[:&lt;VAL&gt;] accessfrom &lt;TYPE&gt;[:&lt;VAL&gt;] accessto &lt;TYPE&gt;[:&lt;VAL&gt;]</pre>
<b>Description</b>	Display various aspects of a credential (formatted in XML)
<b>Example</b>	<pre>&gt; mcredctl -q role user:bob --format=xml</pre> <p style="border: 1px dashed gray; border-radius: 10px; padding: 5px; margin: 5px 0;"><i>View user bob's administrative role within Moab in XML</i></p> <pre>&gt; mcredctl -q limit acct --format=xml</pre> <p style="border: 1px dashed gray; border-radius: 10px; padding: 5px; margin: 5px 0;"><i>Display limits for all accounts in XML</i></p> <pre>&gt; mcredctl -q policies user:bob</pre> <p style="border: 1px dashed gray; border-radius: 10px; padding: 5px; margin: 5px 0;"><i>View limits organized by credential for user bob on each partition and resource manager</i></p> <pre>&gt; mcredctl -q profile group --format=xml --timeout=00:10:00 -o time:1388590200,1431529200,types:TPSD</pre> <p style="border: 1px dashed gray; border-radius: 10px; padding: 5px; margin: 5px 0;"><i>Generates a report of processor hours used by groups per month. TPSD represents total proc-seconds dedicated by this credential in the profiling interval.</i></p>

<b>-r - RESET</b>	
<b>Format</b>	<pre>{stats fairshare uid} &lt;TYPE&gt; [:&lt;ID&gt;]</pre>
<b>Description</b>	Reset the stats, fairshare, or UID/GID of a given credential. <div style="border: 1px solid #4a6984; border-radius: 10px; padding: 5px; margin-top: 10px; background-color: #e6f2ff;"> <p> When resetting UID, only a type of user is supported.</p> </div>
<b>Example</b>	<pre>&gt; mcredctl -r uid user:john</pre> <p style="border: 1px dashed gray; border-radius: 10px; padding: 5px; margin-top: 10px;"><i>Resets the UID/GID for the user named john.</i></p>

<b>-t - TIMEFRAME</b>	
<b>Format</b>	<STARTTIME>[,<ENDTIME>]
<b>Description</b>	Can be used in conjunction with the <code>-q</code> profile option to display profiling information for the specified timeframe.
<b>Example</b>	<pre>&gt; mcredctl -q profile user -t 14:30_06/20</pre>

### 3.7.3.D Credential Statistics XML Output

Credential statistics can be requested as XML (via the `--format=xml` argument) and will be written to `STDOUT` in the following format:

```
> mcredctl -q profile user --format=xml -o time:1182927600,1183013999
<Data>
  <user ...>
    <Profile ...>
    </Profile>
  </user>
</Data>
```

#### Example 3-8: Deleting a group

```
> mcredctl -d group:john
GROUPCFG[john] Successfully purged from config files
```

#### Example 3-9: List users in XML format

```
> mcredctl -l user --format=xml
<Data><user ID="john"></user><user ID="john"></user><user ID="root"></user><user
ID="dev"></user></Data>
```

#### Example 3-10: Display information about a user

```
> mcredctl -q role user:john --format=xml
<Data><user ID="test" role="admin5"></user></Data>
```

## Related Topics

- (Optional) Install Moab Client - in the *Moab HPC Suite Installation and Configuration Guide* explains how to distribute this command to client nodes

## 3.7.4 mdiag

### 3.7.4.A Synopsis

```

mdiag -a [accountid]
mdiag -b [-l policylevel] [-t partition] [-v]
mdiag -c [classid] [-v]
mdiag -C [configfile]
mdiag -e [-w
<starttime>|<endtime>|<eventtypes>|<oidlist>|<eidlist>|<object
list>] --xml
mdiag -f [-o user|group|acct|qos|class] [-v] [--
flags=relative]
mdiag -g [groupid]
mdiag -G [Green]
mdiag -j [jobid] [-t <partition>] [-v][-w
state|user|account|class|group|qos=VALUE] [--flags=policy] [--
blocking]
mdiag -l
mdiag -L [-v]
mdiag -n [-A <creds>] [-t partition] [nodeid] [-v]
mdiag -p [-t partition] [-v] [-v]
mdiag -P [-v] [-v]
mdiag -q [qosid] [-v]
mdiag -r [reservationid] [-v] [--blocking]
mdiag -R [resourcemanagename] [-v] [-v]
mdiag -s [standingreservationid] [--blocking]
mdiag -S [-v] [-v]
mdiag -t [-v] [-v] [partitionid]
mdiag -T [triggerid] [-v] [--blocking]
mdiag -u [userid] [-v]
mdiag [--format=xml]

```

```
[--about] [--help] [--host=<serverHostName>] [--loglevel=<logLevel>] [--msg=<message>] [--port=<serverPort>]
[--timeout=<seconds>] [--version] [--xml]
```

### 3.7.4.B Overview

The *mdiag* command is used to display information about various aspects of the cluster and the results of internal diagnostic tests. In summary, it provides the following:

- Current object health and state information
- Current object configuration (resources, policies, attributes, etc.)
- Current and historical performance/utilization information
- Reports on recent failure
- Object messages

Some *mdiag* options gather information from the Moab cache, which prevents them from interrupting the scheduler, but the `--blocking` option can be used to bypass the cache and interrupt the scheduler.

### 3.7.4.C Arguments

Option	Description
<b>-a [accountid]</b>	Display <a href="#">account</a> information.
<b>-b [-l policylevel] [-t partition] [-v]</b>	Display information on jobs blocked by policies, holds, or other factors. <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> If blocked job diagnostics are specified, the <code>-t</code> option is also available to constrain the report to analysis of particular partition. Also, with blocked job diagnosis, the <code>-l</code> option can be used to specify the analysis policy level.</p> </div>
<b>-c [classid] [-v]</b>	Display <a href="#">class</a> information.

Option	Description
<b>-C [file]</b>	<p>With the vast array of options in the configuration file, the <code>-C</code> option does not validate function, but it does analyze the configuration file for syntax errors including use of invalid parameters, deprecated parameters, and some illegal values. If you start Moab with the <code>-e</code> flag, Moab evaluates the configuration file at startup and quits if an error exists.</p> <div data-bbox="1019 709 1414 961" style="border: 1px solid black; padding: 5px;"> <p><b>i</b> <code>mdiag -C</code> does not print out any <code>#INCLUDE</code> lines listed in <code>moab.cfg</code> (and <code>moab.dat</code>), but it does evaluate and print out the lines found in those included files.</p> </div>
<b>-e [-w starttime endtime eventtypes oidlist eidlist objectlist] --xml</b>	<p>Moab will do a query for all events whose <code>eventtime</code> starts at <code>&lt;starttime&gt;</code> and matches the search criteria. This works only when Moab is configured with ODBC MySQL. The syntax is:</p> <pre>mdiag -e[-w &lt;starttime&gt;  &lt;eventtypes&gt;  &lt;oidlist&gt;   &lt;eidlist&gt; &lt;objectlist&gt;] --xml</pre> <ul style="list-style-type: none"> <li>• <code>starttime</code> default is -</li> <li>• <code>eventtypes</code> default is command delimited, the default is all event types (possible values can be found in the <code>EventType</code> table in the Moab database)</li> <li>• <code>oidlist</code> is a comma-delimited list of object IDs,</li> </ul>

Option	Description
	<p>the default is all objects IDs</p> <ul style="list-style-type: none"> <li>• <code>eidlist</code> is a comma-delimited list of specific event IDs, the default is all event IDs</li> <li>• <code>objectlist</code> is a comma-delimited list of object types, the default is all object types (possible values can be found in the ObjectType table in the Moab database)</li> </ul>
<b>-f [-o user group acct qos class] [-v] [--flags=relative]</b>	Display <a href="#">fairshare</a> information.
<b>-g [groupid]</b>	Display <a href="#">group</a> information.
<b>-G [Green]</b>	Display <a href="#">green computing</a> information.
<b>-j [jobid] [-t partition] [-v] [-w state user account class group qos=VALUE] [--flags=policy] [--blocking]</b>	Display job information.
<b>-l</b>	Diagnose license information contained in the moab.lic file.
<b>-L [-v]</b>	Display limits.
<b>-n [-A creds] [-t partition] [nodeid] [-v]</b>	Display nodes. <div data-bbox="1019 1415 1414 1604" style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> If node diagnostics are specified, the <code>-t</code> option is also available to constrain the report to a particular partition.</p> </div>
<b>-p [-t partition] [-v] [-v]</b>	Display <a href="#">job priority</a> .

Option	Description
	<div style="border: 1px solid black; padding: 5px;"> <p><b>i</b> If priority diagnostics are specified, the <code>-t</code> option is also available to constrain the report to a particular partition.</p> </div>
<code>-P [-v] [-v]</code>	Display partition information.
<code>-q [qosid] [-v]</code>	Display <code>qos</code> information.
<code>-r [reservationid] [-v] [--blocking]</code>	Display reservation information.
<code>-R [rmid] [-v] [-v]</code>	Display resource manager information.
<code>-s [srsv] [--blocking]</code>	Display <code>standing reservation</code> information.
<code>-S [-v] [-v]</code>	Display general scheduler information.
<code>-t [-v] [-v] [partitionid]</code>	Display configuration, usage, health, and diagnostic information about partitions maintained by Moab.
<code>-T [triggerid] [-v] [--blocking]</code>	Display trigger information.
<code>-u [userid] [-v]</code>	Display <code>user</code> information.
<code>--format=xml</code>	Display output in XML format.

## XML Output

Information for most of the options can be reported as XML as well. This is done with the command `mdiag -<option> <CLASS_ID> --format=xml`. For example, XML-based class information will be written to `STDOUT` in the following format:

```

<Data>
  <class <ATTR>="<VAL>" ... >

```

```

<stats <ATTR>="<VAL>" ... >
  <Profile <ATTR>="<VAL>" ... >
  </Profile>
</stats>
</class>
<Data>
...
</Data>

```

Of the *mdiag* options, only `-G` and `-L` cannot be reported as XML.

## Related Topics

- (Optional) Install Moab Client - in the *Moab HPC Suite Installation and Configuration Guide* explains how to distribute this command to client nodes
- [checkjob](#)
- [checknode](#)

## 3.7.5 mdiag -a

### 3.7.5.A Synopsis

```
mdiag -a [accountid]
```

### 3.7.5.B Overview

The *mdiag -a* command provides detailed information about the [accounts](#) (a.k.a. projects) Moab is currently tracking. This command also allows an administrator to verify correct throttling policies and access provided to and from other credentials.

*Example 3-11: Generating information about accounts*

```

> mdiag -a
evaluating acct information
Name      Priority  Flags      QDef      QOSList*
PartitionList Target  Limits
engineering 100      -          high     high,urgent,low [A]
[B]        30.00   MAXJOB=50,75  MAXPROC=400,500
marketing  1        -          low      low [A]
          5.00   MAXJOB=100,110  MAXPS=54000,54500
it        10      -          DEFAULT  DEFAULT,high,urgent,low [A]
          100.00  MAXPROC=100,1250  MAXPS=12000,12500
          FSWEIGHT=1000
development 100      -          high     high,urgent,low [A]
[B]        30.00   MAXJOB=50,75  MAXNODE=100,120

```



```

research      100      -      high      DEFAULT,high,low      [A]
[B]    30.00  MAXNODE=400,500  MAXPS=900000,1000000
DEFAULT      0      -      -      -      -
0.00  -

```

## Related Topics

- [Account](#) credential

## 3.7.6 mdiag -b

### 3.7.6.A Synopsis

```
mdiag -b [-l policylevel] [-t partition] [-v]
```

### 3.7.6.B Overview

The *mdiag -b* command returns information about blocked jobs.

## 3.7.7 mdiag -c

### 3.7.7.A Synopsis

```
mdiag -c [classid] [-v]
```

### 3.7.7.B Overview

The *mdiag -c* command provides detailed information about the classes Moab is currently tracking. This command also allows an administrator to verify correct throttling policies and access provided to and from other credentials.

**i** The term [class](#) is used interchangeably with the term queue and generally refers to resource manager queue.

## 3.7.7.C XML Attributes

Name	Description
<b>ADEF</b>	Accounts a class has access to.
<b>CAPACITY</b>	Number of procs available to the class.
<b>DEFAULT.ATTR</b>	Default attributes attached to a job.
<b>DEFAULT.DISK</b>	Default required disk attached to a job.
<b>DEFAULT.FEATURES</b>	Default required node features attached to a job.
<b>DEFAULT.GRES</b>	Default generic resources attached to a job.
<b>DEFAULT.MEM</b>	Default required memory attached to a job.
<b>DEFAULT.NODESET</b>	Default specified nodeset attached to a job.
<b>DEFAULT.WCLIMIT</b>	Default wallclock limit attached to a job.
<b>EXCL.FEATURES</b>	List of excluded (disallowed) node features.
<b>EXCL.FLAGS</b>	List of excluded (disallowed) job flags.
<b>FSTARGET</b>	The class' fairshare target.
<b>HOLD</b>	If TRUE this credential has a hold on it, FALSE otherwise.
<b>HOSTLIST</b>	The list of hosts in this class.
<b>JOBEPILOG</b>	Scheduler level job epilog to be run after job is completed by resource manager (script path).
<b>JOBFLAGS</b>	Default flags attached to jobs in the class.
<b>JOBPROLOG</b>	Scheduler level job prolog to be run before job is started by resource manager (script path).
<b>ID</b>	The unique ID of this class.

Name	Description
<b>LOGLEVEL</b>	The log level attached to jobs in the class.
<b>MAX.PROC</b>	The max processors per job in the class.
<b>MAX.PS</b>	The max processor-seconds per job in the class.
<b>MAX.WCLIMIT</b>	The max wallclock limit per job in the class.
<b>MAXIJOB</b>	The max idle jobs in the class.
<b>MAXIPROC</b>	The max idle processors in the class.
<b>MAXJOBPERUSER</b>	The max jobs per user.
<b>MAXNODEPERJOB</b>	The max nodes per job.
<b>MAXNODEPERUSER</b>	The max nodes per user.
<b>MAXPROCERJOB</b>	The max processors per job.
<b>MAXPROCERNODE</b>	The max processors per node.
<b>MAXPROCERUSER</b>	The max processors per user.
<b>MIN.NODE</b>	The minimum nodes per job in the class.
<b>MIN.PROC</b>	The minimum processors per job in the class.
<b>MIN.WCLIMIT</b>	The minimum wallclock limit per job in the class.
<b>NODEACCESSPOLICY</b>	The node access policy associated with jobs in the class.
<b>OCDPROCFACOR</b>	Dedicated processor factor.
<b>OCNODE</b>	Overcommit node.
<b>PRIORITY</b>	The class' associated priority.

Name	Description
<b>PRIORITYF</b>	Priority calculation function.
<b>REQ.FEATURES</b>	Required features for a job to be considered in the class.
<b>REQ.FLAGS</b>	Required flags for a job to be considered in the class.
<b>REQ.IMAGE</b>	Required image for a job to be considered in the class.
<b>REQUIREDUSERLIST</b>	The list of users who have access to the class.
<b>RM</b>	The resource manager reporting the class.
<b>STATE</b>	The class' state.
<b>WCOVERRUN</b>	Tolerated amount of time beyond the specified wallclock limit.

Example 3-12: Generating information about classes

```

> mdiag -c
Class/Queue Status
ClassID      Priority Flags      QDef          QOSList* PartitionList
Target Limits
DEFAULT      0 ---      ---          --- ---
0.00 ---
batch        1 ---      ---          --- [A] [B]
70.00 MAXJOB=33:200,250
      MAX.WCLIMIT=10:00:00 MAXPROCPERJOB=128
long         1 ---      low          low [A]
10.00 MAXJOB=3:100,200
      MAX.WCLIMIT=1:00:00:00 MAXPROCPERJOB=128
fast         100 ---      high         high [B]
10.00 MAXJOB=8:100,150
      MAX.WCLIMIT=00:30:00 MAXPROCPERJOB=128
bigmem       1 ---      low,high     low ---
10.00 MAXJOB=1:100,200
      MAXPROCPERJOB=128
    
```

*In the example above, class fast has MAXJOB soft and hard limits of 100 and 150 respectively and is currently running 8 jobs.*

 The Limits column will display limits in the following format:  
 <USAGE>:<HARDLIMIT>[,<SOFTLIMIT>]

**Related Topics**

- [showstats](#) command - display general statistics

## 3.7.8 mdiag -f

### 3.7.8.A Synopsis

```
mdiag -f [-o user|group|acct|qos|class] [--flags=relative]
```

### 3.7.8.B Overview

The `mdiag -f` command is used to display at a glance information about the fairshare configuration and historic resource utilization. The fairshare usage might impact job prioritization, job eligibility, or both based on the credential `FSTARGET` and `FSCAP` attributes and by the fairshare priority weights as described in the [Job Prioritization Overview](#). The information presented by this command includes fairshare configuration and credential fairshare usage over time.

The command hides information about credentials that have no fairshare target and no fairshare cap.

If an object type (`<OTYPE>`) is specified, then only information for that credential type (user, group, acct, class, or qos) will be displayed. If the `relative` flag is set, then per user fairshare usage will be displayed relative to each non-user credential (see the second example below).

**i** Relative output is only displayed for credentials that have user mappings. For example, if there is no association between classes and users, no relative per user fairshare usage class breakdown will be provided.

#### Example 3-13: Standard Fairshare Output

```
> mdiag -f
FairShare Information
Depth: 6 intervals   Interval Length: 00:20:00   Decay Rate: 0.50
FS Policy: DEDICATEDPES
System FS Settings:  Target Usage: 0.00
FSInterval          %      Target      0      1      2      3      4      5
FSWeight             -----
TotalUsage          100.00 -----
USER
-----
mattp                2.51 -----      2.20      2.69      2.21      2.65      2.65      3.01
jsmith              12.82 -----      12.66     15.36     10.96     8.74     8.15     13.85
kyliem              3.44 -----      3.93      2.78      4.36      3.11     3.94     4.25
tgh                 4.94 -----      4.44      5.12      5.52      3.95     4.66     4.76
walex               1.51 -----      3.14      1.15      1.05      1.61     1.22     1.60
jimf                4.73 -----      4.67      4.31      5.67      4.49     4.93     4.92
```

poy	4.64	-----	4.43	4.61	4.58	4.76	5.36	4.90
mjackson	0.66	-----	0.35	0.78	0.67	0.77	0.55	0.43
tfw	17.44	-----	16.45	15.59	19.93	19.72	21.38	15.68
gjohn	2.81	-----	1.66	3.00	3.16	3.06	2.41	3.33
ljill	10.85	-----	18.09	7.23	13.28	9.24	14.76	6.67
kbill	11.10	-----	7.31	14.94	4.70	15.49	5.42	16.61
stevei	1.58	-----	1.41	1.34	2.09	0.75	3.30	2.15
gms	1.54	-----	1.15	1.74	1.63	1.40	1.38	0.90
patw	5.11	-----	5.22	5.11	4.85	5.20	5.28	5.78
wer	6.65	-----	5.04	7.03	7.52	6.80	6.43	2.83
anna	1.97	-----	2.29	1.68	2.27	1.80	2.37	2.17
susieb	5.69	-----	5.58	5.55	5.57	6.48	5.83	6.16
GROUP								
-----								
dallas	13.25	15.00	14.61	12.41	13.19	13.29	15.37	15.09
sanjose*	8.86	15.00	6.54	9.55	9.81	8.97	8.35	4.16
seattle	10.05	15.00	9.66	10.23	10.37	9.15	9.94	10.54
austin*	30.26	15.00	29.10	30.95	30.89	28.45	29.53	29.54
boston*	3.44	15.00	3.93	2.78	4.36	3.11	3.94	4.25
orlando*	26.59	15.00	29.83	26.77	22.56	29.49	25.53	28.18
newyork*	7.54	15.00	6.33	7.31	8.83	7.54	7.34	8.24
ACCT								
-----								
engineering	31.76	30.00	32.25	32.10	31.94	30.07	30.74	31.14
marketing	8.86	5.00	6.54	9.55	9.81	8.97	8.35	4.16
it	9.12	5.00	7.74	8.65	10.92	8.29	10.64	10.40
development*	24.86	30.00	24.15	24.76	25.00	24.84	26.15	26.78
research	25.40	30.00	29.32	24.94	22.33	27.84	24.11	27.53
QOS								
-----								
DEFAULT*	0.00	50.00	-----	-----	-----	-----	-----	-----
high*	83.69	90.00	86.76	83.20	81.71	84.35	83.19	88.02
urgent	0.00	5.00	-----	-----	-----	-----	-----	-----
low*	12.00	5.00	7.34	12.70	14.02	12.51	12.86	7.48
CLASS								
-----								
batch*	51.69	70.00	53.87	52.01	50.80	50.38	48.67	52.65
long*	18.75	10.00	16.54	18.36	20.89	18.36	21.53	16.28
fast*	15.29	10.00	18.41	14.98	12.58	16.80	15.15	18.21
bigmem	14.27	10.00	11.17	14.65	15.73	14.46	14.65	12.87

**i** An asterisk (\*) next to a credential name indicates that that credential has exceeded its fairshare target.

**Example 3-14: Grouping User Output by Account**

```
> mdiag -f -o acct --flags=relative
FairShare Information
Depth: 6 intervals   Interval Length: 00:20:00   Decay Rate: 0.50
FS Policy: DEDICATEDPES
System FS Settings: Target Usage: 0.00
FSInterval      %      Target      0      1      2      3      4      5
FSWeight        -----
TotalUsage      100.00 -----
ACCOUNT
-----
dallas          13.12 15.00    15.42 12.41 13.19 13.29 15.37 15.09
  mattp         19.47 -----    15.00 21.66 16.75 19.93 17.26 19.95
  walex          9.93 -----    20.91 9.28 7.97 12.14 7.91 10.59
```

stevei	12.19	-----	9.09	10.78	15.85	5.64	21.46	14.28
anna	14.77	-----	16.36	13.54	17.18	13.55	15.44	14.37
susieb	43.64	-----	38.64	44.74	42.25	48.74	37.92	40.81
sanjose*	9.26	15.00	8.69	9.55	9.81	8.97	8.35	4.16
mjackson	7.71	-----	6.45	8.14	6.81	8.62	6.54	10.29
gms	17.61	-----	21.77	18.25	16.57	15.58	16.51	21.74
wer	74.68	-----	71.77	73.61	76.62	75.80	76.95	67.97
seattle	10.12	15.00	10.16	10.23	10.37	9.15	9.94	10.54
tgh	49.56	-----	46.21	50.05	53.26	43.14	46.91	45.13
patw	50.44	-----	53.79	49.95	46.74	56.86	53.09	54.87
austin*	30.23	15.00	25.58	30.95	30.89	28.45	29.53	29.54
jsmith	42.44	-----	48.77	49.62	35.47	30.70	27.59	46.90
tfw	57.56	-----	51.23	50.38	64.53	69.30	72.41	53.10
boston*	3.38	15.00	3.78	2.78	4.36	3.11	3.94	4.25
kyliem	100.00	-----	100.00	100.00	100.00	100.00	100.00	100.00
orlando*	26.20	15.00	30.13	26.77	22.56	29.49	25.53	28.18
poy	17.90	-----	16.28	17.22	20.30	16.15	20.98	17.39
ljill	37.85	-----	58.60	26.99	58.87	31.33	57.79	23.67
kbill	44.25	-----	25.12	55.79	20.83	52.52	21.23	58.94
newyork*	7.69	15.00	6.24	7.31	8.83	7.54	7.34	8.24
jimf	61.42	-----	69.66	58.94	64.20	59.46	67.21	59.64
gjohn	38.58	-----	30.34	41.06	35.80	40.54	32.79	40.36

---

## Related Topics

- [Fairshare Overview](#)

## 3.7.9 mdiag -j

### 3.7.9.A Synopsis

```
mdiag -j [jobid] [-t <partition>] [-v] [-w] [--flags=policy] [-  
-xml] [--blocking]
```

### 3.7.9.B Overview

The *mdiag -j* command provides detailed information about the state of jobs Moab is currently tracking. This command also performs a large number of sanity and state checks. The job configuration and status information, as well as the results of the various checks, are presented by this command. The command gathers information from the Moab cache that prevents it from interrupting the scheduler, but the `--blocking` option can be used to bypass the cache and interrupt the scheduler. If the `-v` (verbose) flag is specified, additional information about less common job attributes is displayed. If `--flags=policy` is specified, information about job templates is displayed.

If used with the `-t <partition>` option on a running job, the only thing `mdiag -j` shows is if the job is running on the specified partition. If used on job that is not running, it shows if the job is able to run on the specified partition.

The `-w` flag enables you to specify specific job states, such as Running, Completed, Idle, or ALL (see [Job States](#) for all valid options), or jobs associated with a given credential (user, acct, class, group, jobgroup, qos). For example:

```
mdiag -j -w user=david           # Displays only David's jobs
mdiag -j -w state=Idle,Running  # Displays only idle or running jobs
mdiag -j -w jobgroup=workflow1  # displays jobs in jobgroup workflow1
```

**i** The `mdiag -j` command does not show all subjobs of an array unless you use `mdiag -j --xml`. In the XML, the master job element contains a child element called `ArraySubJobs` that contains the subjobs in the array. Using `mdiag -j -v --xml` shows the completed subjobs as well.

### 3.7.9.C XML Output

If XML output is requested (via the `--format=xml` argument), XML based node information will be written to STDOUT in the following format:

```
<Data>
  <job ATTR="VALUE" ... > </job>
  ...
</Data>
```

For information about legal attributes, refer to the [XML Attributes](#) table.

**i** To show jobs in XML, use `mdiag -j --xml -w [completed=true|system=true|ALL=true]` to limit or filter jobs. This is for XML use only.

---

#### Related Topics

- [checkjob](#)
- [mdiag](#)



## 3.7.10 `mdiag -n`

### 3.7.10.A Synopsis

```
mdiag -n [-t partitionid] [-A creds] [-v] [nodeid]
```

### 3.7.10.B Overview

The `mdiag -n` command provides detailed information about the state of nodes Moab is currently tracking. This command also performs a large number of sanity and state checks. The node configuration and status information and the results of the various checks are presented by this command.

### 3.7.10.C Arguments

Flag	Argument	Description
<b>[-A]</b>	{user group account qos class job}: <OBJECTID>	Report if each node is accessible by requested job or credential.
	[nodeid]	Report on the specified node (default is all nodes).
<b>[-t]</b>	<partitionid>	Report only nodes from specified partition.
<b>[-v]</b>	---	Show verbose output (do not truncate columns and add columns for additional node attributes).

### 3.7.10.D Output

This command presents detailed node information in whitespace-delineated fields.

The output of this command can be extensive and the values for a number of fields can be truncated. If truncated, the `-v` flag can be used to display full field content.

Column	Format
<b>Name</b>	<NODE NAME>

Column	Format
<b>State</b>	<NODE STATE>
<b>Procs</b>	<AVAILABLE PROCS>;<CONFIGURED PROCS>
<b>Memory</b>	<AVAILABLE MEMORY>;<CONFIGURED MEMORY>
<b>Disk</b>	<AVAILABLE DISK>;<CONFIGURED DISK>
<b>Swap</b>	<AVAILABLE SWAP>;<CONFIGURED SWAP>
<b>Speed</b>	<RELATIVE MACHINE SPEED>
<b>Opsys</b>	<NODE OPERATING SYSTEM>
<b>Arch</b>	<NODE HARDWARE ARCHITECTURE>
<b>Par</b>	<PARTITION NODE IS ASSIGNED TO>
<b>Load</b>	<CURRENT 1 MINUTE BSD LOAD>
<b>Rsv</b>	<NUMBER OF RESERVATIONS ON NODE>
<b>Classes</b>	<CLASS NAME>
<b>Network</b>	<NETWORK NAME>...
<b>Features</b>	<NODE FEATURE>...

### 3.7.10.E Examples

```

> mdiag -n

compute node summary
Name                State   Procs   Memory   Opsys
opt-001             Busy   0:2    2048:2048 SUSE
opt-002             Busy   0:2    2048:2048 SUSE
opt-003             Busy   0:2    2048:2048 SUSE
opt-004             Busy   0:2    2048:2048 SUSE
opt-005             Busy   0:2    2048:2048 SUSE
opt-006             Busy   0:2    2048:2048 SUSE
WARNING: swap is low on node opt-006
    
```

```

opt-007          Busy  0:2    2048:2048      SUSE
opt-008          Busy  0:2    2048:2048      SUSE
opt-009          Busy  0:2    2048:2048      SUSE
opt-010          Busy  0:2    2048:2048      SUSE
opt-011          Busy  0:2    2048:2048      SUSE
opt-012          Busy  0:2    2048:2048      SUSE
opt-013          Busy  0:2    2048:2048      SUSE
opt-014          Busy  0:2    2048:2048      SUSE
opt-015          Busy  0:2    2048:2048      SUSE
opt-016          Busy  0:2    2048:2048      SUSE
x86-001          Busy  0:1    512:512        Redhat
x86-002          Busy  0:1    512:512        Redhat
x86-003          Busy  0:1    512:512        Redhat
x86-004          Busy  0:1    512:512        Redhat
x86-005          Idle  1:1    512:512        Redhat
x86-006          Idle  1:1    512:512        Redhat
x86-007          Idle  1:1    512:512        Redhat
x86-008          Busy  0:1    512:512        Redhat
x86-009          Down  1:1    512:512        Redhat
x86-010          Busy  0:1    512:512        Redhat
x86-011          Busy  0:1    512:512        Redhat
x86-012          Busy  0:1    512:512        Redhat
x86-013          Busy  0:1    512:512        Redhat
x86-014          Busy  0:1    512:512        Redhat
x86-015          Busy  0:1    512:512        Redhat
x86-016          Busy  0:1    512:512        Redhat
P690-001         Busy  0:1    16384:16384    AIX
P690-002         Busy  0:1    16384:16384    AIX
P690-003         Busy  0:1    16384:16384    AIX
P690-004         Busy  0:1    16384:16384    AIX
P690-005         Busy  0:1    16384:16384    AIX
P690-006         Busy  0:1    16384:16384    AIX
P690-007         Idle  1:1    16384:16384    AIX
P690-008         Idle  1:1    16384:16384    AIX
WARNING: node P690-008 is missing ethernet adapter
P690-009         Busy  0:1    16384:16384    AIX
P690-010         Busy  0:1    16384:16384    AIX
P690-011         Busy  0:1    16384:16384    AIX
P690-012         Busy  0:1    16384:16384    AIX
P690-013         Busy  0:1    16384:16384    AIX
P690-014         Busy  0:1    16384:16384    AIX
P690-015         Busy  0:1    16384:16384    AIX
P690-016         Busy  0:1    16384:16384    AIX
-----          ---   6:64   745472:745472  -----

Total Nodes: 36 (Active: 30 Idle: 5 Down: 1)

```

**i** Warning messages are interspersed with the node configuration information with all warnings preceded by the keyword **WARNING**.

### 3.7.10.F XML Output

If XML output is requested (via the `--format=xml` argument), XML based node information will be written to **STDOUT** in the following format:

```
mdiag -n --format=xml
```

```

<Data>
  <node <ATTR>="<VAL>" ... </node>
  ...
</Data>

```

### 3.7.10.G XML Attributes

Name	Description
<b>ACL</b>	Node Access Control List.
<b>AGRES</b>	Available generic resources.
<b>ALLOCRES</b>	Special allocated resources (like VLANs).
<b>ARCH</b>	The node's processor architecture.
<b>AVLCLASS</b>	Classes available on the node.
<b>AVLETIME</b>	Time when the node will no longer be available (used in Utility centers).
<b>AVLSTIME</b>	Time when the node will be available (used in Utility centers).
<b>CFGCLASS</b>	Classes configured on the node.
<b>ENABLEPROFILING</b>	If true, a node's state and usage is tracked over time.
<b>FEATURES</b>	A list of comma-separated custom features describing a node.
<b>GEVENT</b>	A user-defined event that allows Moab to perform some action.
<b>GMETRIC</b>	A list of comma-separated consumable resources associated with a node.
<b>GRES</b>	generic resources on the node.
<b>HOPCOUNT</b>	How many hops the node took to reach this Moab (used in hierarchical grids).
<b>ISDELETED</b>	Node has been deleted.

Name	Description
<b>ISDYNAMIC</b>	Node is dynamic (used in Utility centers).
<b>JOBLIST</b>	The list of jobs currently running on a node.
<b>LOAD</b>	Current load as reported by the resource manager.
<b>LOADWEIGHT</b>	Load weight used when calculating node priority.
<b>MAXJOB</b>	See <a href="#">Node Policies</a> for details.
<b>MAXJOBPERUSER</b>	See <a href="#">Node Policies</a> for details.
<b>MAXLOAD</b>	See <a href="#">Node Policies</a> for details.
<b>MAXPROC</b>	See <a href="#">Node Policies</a> for details.
<b>MAXPROCERUSER</b>	See <a href="#">Node Policies</a> for details.
<b>NETWORK</b>	The ability to specify which networks are available to a given node is limited to only a few resource managers. Using the <code>NETWORK</code> attribute, administrators can establish this node to network connection directly through the scheduler. The <code>NODECFG</code> parameter allows this list to be specified in a comma-delimited list.
<b>NODEID</b>	The unique identifier for a node.
<b>NODESTATE</b>	The state of a node.
<b>OS</b>	A node's operating system.
<b>OSLIST</b>	Operating systems the node can run.
<b>OSMODACTION</b>	URL for changing the operating system.
<b>OWNER</b>	Credential type and name of owner.
<b>PARTITION</b>	The partition a node belongs to. See <a href="#">Node Location</a> for details.

Name	Description
<b>POWER</b>	The state of the node's power. Either ON or OFF.
<b>PRIORITY</b>	The fixed node priority relative to other nodes.
<b>PROCSPEED</b>	A node's processor speed information specified in MHz.
<b>RACK</b>	The rack associated with a node's physical location.
<b>RADISK</b>	The total available disk on a node.
<b>RAMEM</b>	The total available memory on a node.
<b>RAPROC</b>	The total number of processors available on a node.
<b>RASWAP</b>	The total available swap on a node.
<b>RCMEM</b>	The total configured memory on a node.
<b>RCPROC</b>	The total configured processors on a node.
<b>RCSWAP</b>	The total configured swap on a node.
<b>RequestID</b>	Dynamic Node RequestID set by the RM.
<b>RESCOUNT</b>	Number of reservations on the node.
<b>RESOURCES</b>	Deprecated (use GRES).
<b>RSVLIST</b>	List of reservations on the node.
<b>RMACCESSLIST</b>	A comma-separated list of resource managers who have access to a node.
<b>SIZE</b>	The number of slots or size units consumed by the node.
<b>SLOT</b>	The first slot in the rack associated with the node's physical location.
<b>SPEED</b>	A node's relative speed.

Name	Description
<b>SPEEDWEIGHT</b>	Speed weight used to calculate node's priority.
<b>STATACTIVETIME</b>	Time node was active.
<b>STATMODIFYTIME</b>	Time node's state was modified.
<b>STATTOTALTIME</b>	Time node has been monitored.
<b>STATUPTIME</b>	Time node has been up.
<b>TASKCOUNT</b>	The number of tasks on a node.
<b>TTL</b>	Dynamic Node Time To Live set by the RM (expiration date in epoch format).

---

## Related Topics

- [checknode](#)

### 3.7.11 mdiag -p

#### 3.7.11.A Synopsis

```
mdiag -p [-t partition] [-v] [-v]
```

#### 3.7.11.B Overview

The *mdiag -p* command is used to display at a glance information about the job priority configuration and its effects on the current eligible jobs. The information presented by this command includes priority weights, priority components, and the percentage contribution of each component to the total job priority.

The command hides information about priority components that have been deactivated (i.e., by setting the corresponding component priority weight to 0). For each displayed priority component, this command gives a small amount of context sensitive information. The following table documents this information. In all cases, the output is of the form

<PERCENT> (<CONTEXT INFO>) where <PERCENT> is the percentage contribution of the associated priority component to the job's total priority.

**i** By default, this command only shows information for jobs that are eligible for immediate execution. Jobs that violate soft or hard policies, or have holds, job dependencies, or other job constraints in place will not be displayed. If priority information is needed for any of these jobs, use the **-v** flag or the [checkjob](#) command.

### 3.7.11.C Format

Flag	Name	Format	Default	Description	Example
<b>-t</b>		<b>-t</b> partition	all partitions	Constrain the report to a particular partition.	<pre>&gt; mdiag -p -t partition1</pre> <p><i>Display priority summary information for jobs in partition1.</i></p>
<b>-v</b>	VERBOSE	---	---	Display verbose priority information. If specified, display priority breakdown information for blocked, eligible, and active jobs. <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> By default, only information for eligible jobs is displayed. To view blocked jobs in addition to eligible, run <code>mdiag -p -v -v</code>.</p> </div>	<pre>&gt; mdiag -p -v</pre> <p><i>Display priority summary information for eligible and active jobs.</i></p>



## 3.7.11.D Output

Priority Component	Format	Description
<b>Target</b>	<PERCENT>()	
<b>QOS</b>	<PERCENT> (<QOS>:<QOSPRI>)	QOS — QOS associated with job QOSPRI — Priority assigned to the QOS
<b>FairShare</b>	<PERCENT> ( <USR> : <GRP> :<ACC>:<QOS>:<CLS>)	USR — user fs usage - user fs target GRP — group fs usage - group fs target ACC — account fs usage - account fs target QOS — QOS fs usage - QOS fs target CLS — class fs usage - class fs target
<b>Service</b>	<PERCENT> (<QT>:<XF>:<Byp>)	QTime — job queue time that is applicable towards priority (in minutes) XF — current theoretical minimum XFactor is job were to start immediately Byp — number of times job was bypassed by lower priority jobs via backfill
<b>Resource</b>	<PERCENT> ( <NDE>:<PE>:<PRC>:<MEM>)	NDE — nodes requested by job PE — Processor Equivalents as calculated by all resources requested by job PRC — processors requested by job MEM — real memory requested by job

## 3.7.11.E Examples

*Example 3-15: mdiag -p*

```
diagnosing job priority information (partition: ALL)
Job
Weights  -----  PRIORITY*  Cred(  QOS)  FS(Accnt)  Serv(QTime)
          1(    1)    1(    1)    1(    1)
13678          1321*    7.6(100.0)  0.2(  2.7)  92.2(1218.)
13698          235*    42.6(100.0) 1.1(  2.7)  56.3(132.3)
13019          8699    0.6( 50.0)  0.3( 25.4)  99.1(8674.)
13030          8699    0.6( 50.0)  0.3( 25.4)  99.1(8674.)
13099          8537    0.6( 50.0)  0.3( 25.4)  99.1(8512.)
13141          8438    0.6( 50.0)  0.2( 17.6)  99.2(8370.)
13146          8428    0.6( 50.0)  0.2( 17.6)  99.2(8360.)
13153          8360    0.0(  1.0)  0.1( 11.6)  99.8(8347.)
```

```

13177          8216    0.0 ( 1.0)    0.1 ( 11.6)   99.8 (8203.)
13203          8127    0.6 ( 50.0)    0.3 ( 25.4)   99.1 (8102.)
13211          8098    0.0 ( 1.0)    0.1 ( 11.6)   99.8 (8085.)
...
13703          137    36.6 ( 50.0)   12.8 ( 17.6)   50.6 ( 69.2)
13702           79    1.3 ( 1.0)     5.7 ( 4.5)    93.0 ( 73.4)

Percent Contribution  -----    0.9 ( 0.9)    0.4 ( 0.4)   98.7 ( 98.7)

```

\* indicates system prio set on job

The `mdiag -p` command only displays information for priority components actually utilized. In the above example, QoS, Account Fairshare, and QueueTime components are utilized in determining a job's priority. Other components, such as Service Targets, and Bypass are not used and therefore are not displayed. See the [Priority Overview](#) for more information. The output consists of a header, a job by job analysis of jobs, and a summary section.

The header provides column labeling and provides configured priority component and subcomponent weights. In the above example, `QOSWEIGHT` is set to 1000 and `FSWEIGHT` is set to 100. When configuring fairshare, a site also has the option of weighting the individual components of a job's overall fairshare, including its user, group, and account fairshare components. In this output, the QoS and account fairshare weights are set to 1.

The job by job analysis displays a job's total priority and the percentage contribution to that priority of each of the priority components. In this example, job 13019 has a total priority of 8699. Both QoS and Fairshare contribute to the job's total priority although these factors are quite small, contributing 0.6% and 0.3% respectively with the fairshare factor being contributed by an account fairshare target. For this job, the dominant factor is the service subcomponent `qtime`, which is contributing 99.1% of the total priority since the job has been in the queue for approximately 8600 minutes.

At the end of the job by job description, a Totals line is displayed, which documents the average percentage contributions of each priority component to the current idle jobs. In this example, the QoS, Fairshare, and Service components contributed an average of 0.9%, 0.4%, and 98.7% to the jobs' total priorities.

## Related Topics

- [Job Priority Overview](#)

## 3.7.12 mdiag -q

### 3.7.12.A Synopsis

```
mdiag -q [qosid]
```

### 3.7.12.B Overview

The `mdiag -q` command is used to present information about each QoS maintained by Moab. The information presented includes QoS name, membership, scheduling priority, weights and flags.

### 3.7.12.C Examples

#### Example 3-16: Standard QOS Diagnostics

```

> mdiag -q
QOS Status
System QOS Settings:  QList: DEFAULT (Def: DEFAULT)  Flags: 0
Name                  * Priority QTWeight QTTarget XFWeight XFTarget   QFlags
JobFlags Limits
DEFAULT              1         1         3         1         5.00  PREEMPTEE
[NONE] [NONE]
  Accounts:  it research
  Classes:  batch
[ALL]                0         0         0         0         0.00  [NONE]
[NONE] [NONE]
high                 1000      1         2         1         10.00  PREEMPTOR
[NONE] [NONE]
  Accounts:  engineering it development research
  Classes:  fast
urgent              10000     1         1         1         7.00  PREEMPTOR
[NONE] [NONE]
  Accounts:  engineering it development
low                 100       1         5         1         1.00  PREEMPTEE
[NONE] [NONE]
  Accounts:  engineering marketing it development research
  Classes:  long bigmem

```

### 3.7.13 mdiag -r

#### 3.7.13.A Synopsis

```
mdiag -r [reservationid] [-v] [--blocking]
```

#### 3.7.13.B Overview

The *mdiag -r* command allows administrators to look at detailed reservation information. It provides the name, type, partition, starttime and endtime, proc and node counts, as well as actual utilization figures. It also provides detailed information about which resources are being used, how many nodes, how much memory, swap, and processors are being associated with each task. Administrators can also view the Access Control Lists for each reservation, as well as any flags that may be active in the reservation. The command gathers information from the Moab cache, which prevents it from waiting for the scheduler, but the *--blocking* option can be used to bypass the cache and allow waiting for the scheduler.

## 3.7.13.C Examples

```

> mdiag -r
Diagnosing Reservations
RsvID          Type Par   StartTime      EndTime        Duration Node Task
Proc
-----
-
engineer.0.1    User  A     -6:29:00      INFINITY      INFINITY      0    0
7
  Flags: STANDINGRSV IGNSTATE OWNERPREEMPT
  ACL:   CLASS==batch+==long+==fast+==bigmem+ QOS==low-==high+ JATTR==PREEMPTEE+
  CL:   RSV==engineer.0.1
  Task Resources: PROCS: [ALL]
  Attributes (HostExp='fr10n01 fr10n03 fr10n05 fr10n07 fr10n09 fr10n11 fr10n13
fr10n15')
  PH Allocated to Jobs: 43.77/45.44 (96.31%)
  SRAttributes (TaskCount: 0 StartTime: 00:00:00 EndTime: 1:00:00:00 Days: ALL)
research.0.2    User  A     -6:29:00      INFINITY      INFINITY      0    0
8
  Flags: STANDINGRSV IGNSTATE OWNERPREEMPT
  ACL:   CLASS==batch+==long+==fast+==bigmem+ QOS==high+==low- JATTR==PREEMPTEE+
  CL:   RSV==research.0.2
  Task Resources: PROCS: [ALL]
  Attributes (HostExp='fr3n01 fr3n03 fr3n05 fr3n07 fr3n07 fr3n09 fr3n11 fr3n13
fr3n15')
  PH Allocated to Jobs: 51.60/51.93 (99.36%)
  SRAttributes (TaskCount: 0 StartTime: 00:00:00 EndTime: 1:00:00:00 Days: ALL)
fast.0.3        User  A      00:14:05      5:14:05       5:00:00       0    0
16
  Flags: STANDINGRSV IGNSTATE OWNERPREEMPT
  ACL:   CLASS==fast+ QOS==high+==low+==urgent+==DEFAULT+ JATTR==PREEMPTEE+
  CL:   RSV==fast.0.3
  Task Resources: PROCS: [ALL]
  Attributes (HostExp='fr12n01 fr12n02 fr12n03 fr12n04 fr12n05 fr12n06 fr12n07
fr12n08 fr12n09 fr12n10 fr12n11 fr12n12 fr12n13 fr12n14 fr12n15 fr12n16')
  SRAttributes (TaskCount: 0 StartTime: 00:00:00 EndTime: 5:00:00 Days:
Mon,Tue,Wed,Thu,Fri)
fast.1.4        User  A      1:00:14:05    1:05:14:05    5:00:00       0    0
16
  Flags: STANDINGRSV IGNSTATE OWNERPREEMPT
  ACL:   CLASS==fast+ QOS==high+==low+==urgent+==DEFAULT+ JATTR==PREEMPTEE+
  CL:   RSV==fast.1.4
  Task Resources: PROCS: [ALL]
  Attributes (HostExp='fr12n01 fr12n02 fr12n03 fr12n04 fr12n05 fr12n06 fr12n07
fr12n08 fr12n09 fr12n10 fr12n11 fr12n12 fr12n13 fr12n14 fr12n15 fr12n16')
  SRAttributes (TaskCount: 0 StartTime: 00:00:00 EndTime: 5:00:00 Days:
Mon,Tue,Wed,Thu,Fri)
job2411         Job   A      -00:01:00     00:06:30      Each tile contains a
summary information about the service it represents, including the following:
  ACL:   JOB==job2411=
  CL:   JOB==job2411 USER==jimf GROUP==newyork ACCT==it CLASS==bigmem QOS==low
JATTR==PREEMPTEE DURATION==00:07:30 PROC==6 PS==2700
job1292         Job   A       00:00:00     00:07:30      00:07:30       0    0
4
  ACL:   JOB==job1292=
  CL:   JOB==job1292 USER==jimf GROUP==newyork ACCT==it CLASS==batch QOS==DEFAULT
JATTR==PREEMPTEE DURATION==00:07:30 PROC==4 PS==1800

```

**Example 3-17:**

With the `-v` option, a nodes line is included for each reservation and shows how many nodes are in the reservation, as well as how many tasks are on each node.

```

> mdiag -r -v
Diagnosing Reservations
RsvID          Type Par   StartTime   EndTime     Duration Node Task
Proc          -----
-----
-
Moab.6          Job  B    -00:01:05   00:00:35   00:01:40   1    1
1
  Flags: ISACTIVE
  ACL:   JOB==Moab.6=
  CL:   JOB==Moab.6 USER==tuser1 GROUP==tgroup1 CLASS==fast QOS==starter
  JPRIORITY<=0 DURATION==00:01:40 PROC==1 PS==100
  SubType: JobReservation
  Nodes='node002:1'
  Rsv-Group: Moab.6

Moab.4          Job  B    -00:01:05   00:00:35   00:01:40   1    1
1
  Flags: ISACTIVE
  ACL:   JOB==Moab.4=
  CL:   JOB==Moab.4 USER==tuser1 GROUP==tgroup1 CLASS==batch QOS==starter
  JPRIORITY<=0 DURATION==00:01:40 PROC==1 PS==100
  SubType: JobReservation
  Nodes='node002:1'
  Rsv-Group: Moab.4

Moab.5          Job  A    -00:01:05   00:00:35   00:01:40   3    3
6
  Flags: ISACTIVE
  ACL:   JOB==Moab.5=
  CL:   JOB==Moab.5 USER==tuser1 GROUP==tgroup1 ACCT==marketing CLASS==long
  QOS==low JPRIORITY<=0 DURATION==00:01:40 PROC==6 PS==600
  Task Resources: PROCS: [ALL]
  SubType: JobReservation
  Nodes='node008:1,node007:1,node006:1'
  Rsv-Group: Moab.5

Moab.7          Job  A    -00:01:04   00:00:36   00:01:40   1    1
1
  Flags: ISACTIVE
  ACL:   JOB==Moab.7=
  CL:   JOB==Moab.7 USER==tuser1 GROUP==tgroup1 CLASS==bigmen QOS==starter
  JPRIORITY<=0 DURATION==00:01:40 PROC==1 PS==100
  SubType: JobReservation
  Nodes='node005:1'
  Rsv-Group: Moab.7

Moab.2          Job  A    -00:01:07   3:58:53    4:00:00    1    2
2
  Flags: ISACTIVE
  ACL:   JOB==Moab.2=
  CL:   JOB==Moab.2 USER==tuser1 GROUP==tgroup1 QOS==starter JPRIORITY<=0
  DURATION==4:00:00 PROC==2 PS==28800
  SubType: JobReservation
  Nodes='node009:1'
  Rsv-Group: Moab.2

```

```

Moab.8                Job  A      3:58:53      7:58:53      4:00:00      8      16
16
  Flags: PREEMPTTEE
  ACL:   JOB==Moab.8=
  CL:    JOB==Moab.8 USER==tuser1 GROUP==tgroup1 ACCT==development CLASS==bigmen
QOS==starter JPRIORITY<=0 DURATION==4:00:00 PROC==16 PS==230400
  SubType: JobReservation

Nodes='node009:1,node008:1,node007:1,node006:1,node005:1,node004:1,node003:1,node001:
1'
  Attributes (Priority=148)
  Rsv-Group: idle

system.3              User bas  -00:01:08      INFINITY      INFINITY      1      1
2
  Flags: ISCLOSED,ISACTIVE
  ACL:   RSV==system.3=
  CL:    RSV==system.3
  Accounting Creds: User:root
  Task Resources: PROCS: [ALL]
  SubType: Other
  Nodes='node254:1'
  Attributes (HostExp='node254')
  PH Allocated to Jobs: 0.00/0.01 (0.00%)
  History: 1322773208:PROCS=2

system.2              User bas  -00:01:08      INFINITY      INFINITY      1      1
2
  Flags: ISCLOSED,ISACTIVE
  ACL:   RSV==system.2=
  CL:    RSV==system.2
  Accounting Creds: User:root
  Task Resources: PROCS: [ALL]
  SubType: Other
  Nodes='node255:1'
  Attributes (HostExp='node255')
  PH Allocated to Jobs: 0.00/0.01 (0.00%)
  History: 1322773208:PROCS=2

system.1              User bas  -00:01:08      INFINITY      INFINITY      1      1
2
  Flags: ISCLOSED,ISACTIVE
  ACL:   RSV==system.1=
  CL:    RSV==system.1
  Accounting Creds: User:root
  Task Resources: PROCS: [ALL]
  SubType: Other
  Nodes='node256:1'
  Attributes (HostExp='node256')
  PH Allocated to Jobs: 0.00/0.01 (0.00%)
  History: 1322773208:PROCS=2

```

## 3.7.14 mdiag -R

### 3.7.14.A Synopsis

```
mdiag -R [-v] [-v] [resourcemanagerid]
```

### 3.7.14.B Overview

The *mdiag -R* command is used to present information about configured resource managers. The information presented includes name, host, port, state, type, performance statistics and failure notifications.

### 3.7.14.C Examples

```
> $ mdiag -R -v
diagnosing resource managers

RM[internal] State: --- Type: SSS ResourceType: COMPUTE
Max Fail/Iteration: 0
JobCounter: 6
Partition: SHARED
RM Performance: AvgTime=0.00s MaxTime=0.00s (55353 samples)
RM Languages: -
RM Sub-Languages: -

RM[torque] State: Active Type: PBS ResourceType: COMPUTE
Timeout: 30000.00 ms
Version: '4.2.4'
Job Submit URL: exec:///opt/torque-4.2/bin/qsub
Objects Reported: Nodes=1 (12 procs) Jobs=1
Nodes Reported: 1 (N/A)
Flags: executionServer
Partition: torque
Event Management: EPORT=15004 (last event: 00:03:07)
NOTE: SSS protocol enabled
Submit Command: /opt/torque-4.2/bin/qsub
DefaultClass: batch
Total Jobs Started: 1
RM Performance: AvgTime=0.00s MaxTime=35.00s (220097 samples)
RM Languages: PBS
RM Sub-Languages: PBS

RM[torque] Failures:
clusterquery (683 of 55349 failed)
-12days 'cannot connect to PBS server '' (pbs_errno=15033, 'Batch protocol
error')'

NOTE: use 'mrmctl -f messages <RMID>' to clear stats/failures

RM[FLEXlm] State: Active Type: NATIVE ResourceType: LICENSE
Timeout: 30000.00 ms
Cluster Query URL: exec://$TOOLSDIR/flexlm/license.mon.flexLM.pl
Licenses Reported: 6 types (250 of 282 available)
Partition: SHARED
License Stats: Avg License Avail: 239.01 (978 iterations)
Iteration Summary: Idle: 396.42 Active: 150.92 Busy: -447.34
License biocol 50 of 50 available (Idle: 100.00% Active: 0.00%)
```

```

License cloudform      100 of 100 available (Idle: 100.00% Active: 0.00%)
License mathworks     8 of 25 available (Idle: 52.00% Active: 48.00%)
License verity        25 of 25 available (Idle: 100.00% Active: 0.00%)
Event Management:    (event interface disabled)
RM Performance:      AvgTime=0.00s MaxTime=0.61s (1307618 samples)
  clusterquery:     AvgTime=0.02s MaxTime=0.61s (9465 samples)
  queuequery:       AvgTime=0.00s MaxTime=0.00s (1 samples)
  rinitialize:      AvgTime=0.00s MaxTime=0.00s (1 samples)
  getdata:          AvgTime=0.17s MaxTime=0.60s (978 samples)
RM Languages:        NATIVE
RM Sub-Languages:    NATIVE

AM[mam] Type: MAM State: 'Active'
Host:                localhost
Port:                7112
Timeout:             15
Accounting Mode:     strict-allocation
Job Charge Policy:   All
Reservation Charge Policy: Select
Retry Failed Charges: TRUE

AM[mam] Failures:
Thu Jun 21 14:32:45 Create          'Failure registering job Create (1) with
accounting manager -- server rejected request with status code 740 - Insufficient
funds: There are no valid allocations to satisfy the quote'
```

## 3.7.15 `mdiag -s`

### 3.7.15.A Synopsis

```
mdiag -s [reservationid] [-v]>
```

### 3.7.15.B Overview

The `mdiag -s` command allows administrators to look at detailed standing reservation information. It provides the name, type, partition, starttime and endtime, period, task count, host list, and a list of child instances.

### 3.7.15.C Examples

```
> mdiag -s
```



```

standing reservation overview
RsvID          Type    Par    StartTime    EndTime    Duration    Period
-----
TestSR          User    ---    00:00:00     ---       00:00:00     DAY
  Days:         ALL
  Depth:        2
  RsvList:      testSR.1,testSR.2,testSR.3
  HostExp:      'node1,node2,node4,node8'

test2           User    ---    00:00:00     ---       00:00:00     DAY
  Days:         ALL
  TaskCount:    4
  Depth:        1
  RsvList:      test2.4,test2.5

```

## 3.7.16 mdiag -S

### 3.7.16.A Synopsis

```
mdiag -S [-v] [-v]
```

### 3.7.16.B Overview

The *mdiag -S* command is used to present information about the status of the scheduler and grid interface.

This command will report on the following aspects of scheduling:

- General Scheduler Configuration
  - Reports short and long term scheduler load
  - Reports detected overflows of node, job, reservation, partition, and other scheduler object tables
- High Availability
  - Configuration
  - Reports health of HA primary
  - Reports health of HA backup
- Scheduling Status
  - Reports if scheduling is paused
  - Reports if scheduling is stopped

- System Reservation Status
  - Reports if global system reservation is active
- Message Profiling/Statistics Status
- Moab scheduling activities (only with `mdiag -S -v -v`)
  - Activity[JobStart]: Time Moab spends telling the RM to start a job and waiting for a response.
  - Activity[RMResourceLoad]: Time Moab spends querying license managers and nodes.
  - Activity[RMWorkloadLoad]: Time Moab spends querying resource managers about jobs (as opposed to nodes)
  - Activity[Schedule]: Time Moab spends prioritizing jobs and scheduling them onto nodes.
  - Activity[UIProcess]: Time Moab spends handling client commands.

### 3.7.16.C Examples

*Example 3-18:*

```
> mdiag -S
Moab Server running on orion-1:43225 (Mode: NORMAL)
Load(5m) Sched: 12.27% RMAction: 1.16% RMQuery: 75.30% User: 0.29% Idle: 10.98%
Load(24h) Sched: 10.14% RMAction: 0.93% RMQuery: 74.02% User: 0.11% Idle: 13.80%
HA Fallback Server: orion-2:43225 (Fallback is Ready)
Note: system reservation blocking all nodes
Message: profiling enabled (531 of 600 samples/5:00 interval)
```

## 3.7.17 mdiag -t

### 3.7.17.A Synopsis

```
mdiag -t [-v] [-v] [partitionid]
```

### 3.7.17.B Overview

The `mdiag -t` command is used to present configuration, usage, health, and diagnostic information about partitions maintained by Moab. The information presented includes partition name, limits, configured and available resources, allocation weights and policies.

### 3.7.17.C Examples

*Example 3-19: Standard partition diagnostics*

```
> mdiag -t
Partition Status
...
```

## 3.7.18 mdiag -T

### 3.7.18.A Synopsis

```
mdiag -T [triggerid] [-v] [--blocking]
```

### 3.7.18.B Overview

The *mdiag -T* command is used to present information about each Trigger. The information presented includes TrigID, Object ID, Event (Etype) TType, AType, ActionDate, State. The command gathers information from the Moab cache, which prevents it from waiting for the scheduler, but the *--blocking* option can be used to bypass the cache and allow waiting for the scheduler.

### 3.7.18.C Examples

```
> mdiag -T
TrigID          Object ID          Event TType      AType ActionDate
State
-----
sched_trig.0    sched:Moab         end generic    exec      -
Blocked
3              node:node010      threshol generic    exec      -
Blocked
5              job:Moab.7        preempt generic    exec      -
Blocked
6              job:Moab.8        preempt generic    exec      -
Blocked
7              qos:HIGH          threshol elastic    exec      -
Blocked
4*             job:Moab.5        start generic    exec      0:00:36
Failure
* indicates trigger has completed
```

*Example 3-20:*

```
> mdiag -T -v
```

```

TrigID          Object ID          Event TType      AType
ActionDate      State
-----
sched_trig.0    sched:Moab          end generic      exec
- Blocked
  Name:          sched_trig
  Flags:         globaltrig
  BlockUntil:    INFINITY ActiveTime: ---
  Action Data:   date
  NOTE: trigger can launch

3              node:node010        threshol generic  exec
- Blocked
  Flags:         globaltrig
  BlockUntil:    INFINITY ActiveTime: ---
  Threshold:     CPULoad > 3.00 (current value: 0.00)
  Action Data:   date
  NOTE: trigger cannot launch - threshold not satisfied - threshold type not
supported

5              job:Moab.7          preempt generic   exec
- Blocked
  Flags:         user,globaltrig
  BlockUntil:    INFINITY ActiveTime: ---
  Action Data:   $HOME/tools/preemptnotify.pl $OID $OWNER $HOSTNAME

6              job:Moab.8          preempt generic   exec
- Blocked
  Flags:         user,globaltrig
  BlockUntil:    INFINITY ActiveTime: ---
  Action Data:   $HOME/tools/preemptnotify.pl $OID $OWNER $HOSTNAME
  NOTE: trigger cannot launch - parent job Moab.8 is in state Idle

7              qos:HIGH          threshol elastic  exec
- Blocked
  Flags:         multifire,globaltrig
  BlockUntil:    INFINITY ActiveTime: ---
  Timeout:       00:05:00
  Threshold:     BacklogCompletionTime > 500.00 (current value: 0.00)
  Trigger Type:  elastic
  RearmTime:     00:00:10
  Action Data:   $HOME/geometry.pl $REQUESTGEOMETRY
  NOTE: trigger cannot launch - threshold not satisfied -
requires usage 0.000000 > 500.000000

4*              job:Moab.5          start generic     exec   Mon Jan 16
12:33:00      Failure
  Launch Time:   -00:02:17
  Flags:         globaltrig
  Last Execution State: Failure (ExitCode: 0)
  BlockUntil:    00:00:00 ActiveTime: 00:00:00
  Action Data:   $HOME/tools/preemptnotify.pl $OID $OWNER $HOSTNAME
  ALERT: trigger failure detected
  Message:       'exec '/usr/test/moab/tools/preemptnotify.pl' cannot be located or is
not executable'

* indicates trigger has completed

```

## 3.7.19 mdiag -u

### 3.7.19.A Synopsis

```
mdiag -u [userid]
```

### 3.7.19.B Overview

The *mdiag -u* command is used to present information about user records maintained by Moab. The information presented includes user name, UID, scheduling priority, default job flags, default QOS level, List of accessible QOS levels, and list of accessible partitions.

### 3.7.19.C Examples

```
> mdiag -u
evaluating user information
Name          Priority    Flags          QDef          QOSList*      PartitionList
Target Limits
jvella        0          [NONE]         [NONE]        [NONE]         [NONE]
0.00 [NONE]
  ALIST=Engineering
  Message: profiling enabled (597 of 3000 samples/00:15:00 interval)
[NONE]        0          [NONE]         [NONE]        [NONE]         [NONE]
0.00 [NONE]
reynolds      0          [NONE]         [NONE]        [NONE]         [NONE]
0.00 [NONE]
  ALIST=Administration
  Message: profiling enabled (597 of 3000 samples/00:15:00 interval)
mshaw        0          [NONE]         [NONE]        [NONE]         [NONE]
0.00 [NONE]
  ALIST=Test
  Message: profiling enabled (584 of 3000 samples/00:15:00 interval)
kforbes      0          [NONE]         [NONE]        [NONE]         [NONE]
0.00 [NONE]
  ALIST=Shared
  Message: profiling enabled (597 of 3000 samples/00:15:00 interval)
gastor       0          [NONE]         [NONE]        [NONE]         [NONE]
0.00 [NONE]
  ALIST=Engineering
  Message: profiling enabled (597 of 3000 samples/00:15:00 interval)
```

Note that only users who have jobs that are currently queued or have been queued since Moab was most recently started are listed.

#### Related Topics

- [showstats](#) command (display user statistics)

## 3.7.20 mjobctl

### 3.7.20.A Synopsis

```

mjobctl -c jobexp
mjobctl -c -w [jobexp] attr=val
mjobctl -C jobexp
mjobctl -e jobid
mjobctl -F jobexp
mjobctl -h [User|System|Batch|Defer|All] jobexp
mjobctl -m attr{+=|-=|=}val jobexp [--flags=force]
mjobctl -N [<SIGNO>] jobexp
mjobctl -p <PRIORITY> jobexp
mjobctl -q {diag|starttime|hostlist} jobexp
mjobctl -r jobexp
mjobctl -R jobexp [--flags=force | unmigrate
mjobctl -s jobexp
mjobctl -u jobexp
mjobctl -w attr{+=|-=|=}val jobexp
mjobctl -x [-w flags=val jobexp

[--about] [--help] [--host=<serverHostName>] [--
loglevel=<logLevel>] [--msg=<message>] [--port=<serverPort>]
[--timeout=<seconds>] [--version] [--xml]

```

### 3.7.20.B Overview

The *mjobctl* command controls various aspects of jobs. It is used to submit, cancel, execute, and checkpoint jobs. It can also display diagnostic information about each job. The *mjobctl* command enables the Moab administrator to control almost all aspects of job behavior. See [General Job Administration](#) for more details on jobs and their attributes.

## 3.7.20.C Options

<b>-c - Cancel</b>	
<b>Format</b>	<b>JOBEXP</b>
<b>Description</b>	<p>Cancel a job.</p> <div style="border: 1px solid black; padding: 5px; margin: 5px 0;"> <p><b>i</b> Use <b>-w</b> (following a <b>-c</b> flag) to specify job cancellation according to given credentials or job attributes. See <b>-c -w</b> for more information.</p> </div> <p>You can use <code>mjobctl -c flags=follow-dependency &lt;job_id&gt;</code> to cancel all jobs that the <code>&lt;job_id&gt;</code> depends on.</p> <div style="border: 1px solid black; padding: 5px; margin: 5px 0;"> <p><b>i</b> If you want to cancel all jobs that depend on this <code>&lt;job_id&gt;</code>, add <code>FLAGS=CANCELFAILEDDEPENDENCYJOBS</code> to your <code>SCHEDCFG</code> entry in <code>moab.cfg</code> file. See <a href="#">CANCELFAILEDDEPENDENCYJOBS</a> for more information.</p> </div>
<b>Example</b>	<pre>&gt; mjobctl -c job1045</pre> <p><i>Cancel job job1045.</i></p>

<b>-c -w - Cancel Where</b>	
<b>Format</b>	<p><b>[JOBEXP] &lt;ATTR&gt;=&lt;VALUE&gt;</b></p> <p>where <code>&lt;ATTR&gt;=[ user   account   qos   class   reqreservation (RsvName)   state (JobState)   jobname (JobName, not job ID)]   partition</code></p>
<b>Description</b>	<p>Cancel a job based on a given credential or job attribute. Use <b>-w</b> following a <b>-c</b> flag to specify job cancellation according to credentials or job attributes (see examples) .</p> <p>See <a href="#">Job States</a> for a list of all valid job states.</p> <p>Also, you can cancel jobs from given partitions using <code>-w partition=&lt;PAR1&gt;[&lt;PAR2&gt;...]</code>; however, you must also either use another <code>-w</code> flag to specify a job or use the standard job expression.</p>
<b>Example</b>	<pre>&gt; mjobctl -c -w state=USERHOLD</pre> <p><i>Cancels all jobs that currently have a USERHOLD on them.</i></p> <pre>&gt; mjobctl -c -w user=user1 -w acct=acct1</pre>

**-c -w - Cancel Where**

*Cancels all jobs assigned to user1 or acct1.*

```
> mjobctl -c moab.48655 -w state=IDLE
```

*Cancels job moab.48655, if it is idle.*

**-C - Checkpoint****Format**

JOBEXP

**Description**

Checkpoint a job. See [Checkpoint/Restart Facilities](#) for more information.

**Example**

```
> mjobctl -C job1045
```

*Checkpoint job job1045.*

**-e - Rerun****Format**

JOBID

**Description**

Rerun the completed Torque job. This works only for jobs that are completed and show up in Torque as completed. This flag does not work with other resource managers.

**Example**

```
> mjobctl -e job1045
```

*Rerun job job1045.*

**-F - Force Cancel****Format**


JOBEXP

**Description**

Forces a job to cancel and ignores previous cancellation attempts.



**-F - Force Cancel**

 Specifying this option tells Moab to purge a job from Torque (equivalent to `qdel -p`). This only tells `pbs_server` to remove any knowledge of the job from its internal memory. If the job is actually running, this will not cause `pbs_server` to tell the nodes with the job to cancel it. Therefore, users and administrators should only use this form of `mjobctl` when they've confirmed that the job no longer exists on any compute nodes, and want to force Torque to stop tracking the job.

**Example**

```
> mjobctl -F job1045
```

*Force cancel job job1045.*

**-h - Hold****Format**

<HOLDTYPE><JOBEXP>

<HOLDTYPE> = { user | batch | system | defer | ALL }

**Default**

user

**Description**

Set or release a job hold.

See [Job Holds](#) for more information.

**Example**

```
> mjobctl -h user job1045
```


*Set a user hold on job job1045.*

```
> mjobctl -u all job1045
```

*Unset all holds on job job1045.*

**-m - Modify****Format**

<ATTR>{ += | = | -- } <VAL>

 When using `mjobctl -m` with the `hostlist` attribute, only "=" is supported.

**-m - Modify**

**i** If using Torque and `mjobctl -m` with the partition attribute, only "=" is supported. "+=", "-=", and "=" are supported with other resource managers (Native).

<ATTR>={ account | advres | arraylimit | awduration | class | cpuclock | deadline | depend | eeduration | env | features | feature | flags | gres | group | hold | hostlist | jobdisk | jobmem | jobname | jobswap | loglevel | maxmem | messages | minstarttime | nodeaccess | nodecount | notificationaddress | partition | priority | queue | qos | reqreservation | rmxstring | reqattr | reqawduration | sysprio | tpn | trig | trigvar | user | userprio | var | wclimit}

**Description**

Modify a specific job attribute.

**i** If an `mjobctl -m` attribute can affect how a job starts, then it generally cannot affect a job that is already running. For example, it is not feasible to change the `hostlist` of a job that is already running.

The `userprio` attribute allows you to specify user priority. For job priority, use the `'-p'` flag.

Modification of the job dependency is also communicated to the resource manager in the case of PBS/Torque.

Adding `--flags=warnifcompleted` causes a warning message to print when a job completes.

To define values for `awduration`, `eeduration`, `minstarttime` (Note that the `minstarttime` attribute performs the same function as `msub -a.`), `reqawduration`, and `wclimit`, use the [time spec](#) format.

A *non-active* job's partition list can be modified. If using Torque, only "=" (set) is supported. If using a Native resource manager you can add or subtract partitions, even multiple partitions. When adding or subtracting *multiple* partitions, each partition must have its own `-m partition{+= | = | -=}name` on the command line. An example for adding multiple partitions is provided in the list of examples.

To modify a job's generic resources, use the following format: `gres{ += | = | -= } <gresName>[:<count>].<gresName>` is a single resource, not a list. `<count>` is an integer that, if not specified, is assumed to be 1. Modifying a job's generic resources causes Moab to append the new gres (`+=`), subtract the specified gres (`-=`), or clear out all existing generic resources attached to the job and override them with the newly-specified one (`=`). If `<gresName>` is an empty string, all generic resources will be removed from the job.

To modify the node access policy for a *queued* job, use `nodeaccess=[<policy>]`. See [4.3 Node Access Policies - page 365](#) for a list of supported node access policies.

**-m - Modify****Example**

```
> mjobctl -m messages+="Adding a message" --flags=completed 1664
```

*Set the message on the job, even if the job is completed.*

```
> mjobctl -m reqawduration+=600 1664
```

*Add 10 minutes to the job walltime.*

```
> mjobctl -m eeduration=-1 1664
```

*Reset job's effective queue time, to when the job was submitted.*

```
> mjobctl -m var=Flag1=TRUE 1664
```

*Set the job variable `Flag1` to `TRUE`.*

```
> mjobctl -m notificationaddress="name@server.com"
```

*Sets the notification email address associated with a job to `name@server.com`.*

```
> mjobctl -m partition+=p3 -m partition+=p4 Moab.5
```

*Adds multiple partitions (`p3` and `p4`) to job `Moab.5`.  
Torque only supports `"="`. `"+="`, `"-="`, and `"="` are supported with other resource managers (Native).*

```
> mjobctl -m arraylimit=10 sim.25
```

*Changes the concurrently running subjob limit to 10 for array `sim.25`.*

```
> mjobctl -m gres=matlab:1 job0201
```

*Overrides all generic resources applied to job `job0201` and replaces them with 1 `matlab`.*

```
> mjobctl -m user=user.job
```

*Modifies the user of a job that was submitted directly to moab (`msub`) and has not yet been migrated.*

```
> mjobctl -m userprio=-100 Moab.4
```

*Reduces the user priority of `Moab.4` by 100.*

```
> mjobctl -m tpn=2 Moab.128
```

*Changes the requested "tasks per node" for job `Moab.128` to 2.*

**-m - Modify**

```
> mjobctl -m maxmem=80mb 157
```

*Modifies the total job memory of job 157. See [MAXMEM - page 671](#) for more information.*

**-N - Notify**

**Format** [signal=]<SIGID>[JOBEXP](#)

**Description** Send a signal to all jobs matching the job expression.

**Example**

```
> mjobctl -N INT 1664
```

*Send an interrupt signal to job 1664.*

```
> mjobctl -N 47 1664
```

*Send signal 47 to job 1664.*

**-p - Priority**

**Format** [+|=|-]=<VAL><JOBID> [--flags=relative]

**Description** Modify a job's system priority.

**Example**

Priority is the job priority plus the system priority. Each format affects the job and system priorities differently. Using the format <VAL><JOBID> or +<VAL><JOBID> will set the system priority to the maximum system priority plus the specified value. Using +=<VAL><JOBID> or <VAL><JOBID> --flags=relative will relatively increase the job's priority and set the system priority. Using the format -=<VAL> <JOBID> sets the system priority to 0, and does not change priority based on <VAL> (it will not decrease priority by that number).

For the following example, job1045 has a priority of 10, which is composed of a job priority of 10 and a system priority of 0.

```
> mjobctl -p +1000 job1045
```

*The system priority changes to the max system priority plus 1000 points, ensuring that this job will be higher priority than all normal jobs. In this case, the job priority of 10 is not added, so the priority of job1045 is now 1000001000.*

**-p - Priority**

```
> mjobctl -p --=1 job1045
```

*The system priority of job1045 resets to 0. The job priority is still 10, so the overall priority becomes 10.*

```
> mjobctl -p 3 job1045 --flags=relative
```

*Adds 3 points to the relative system priority. The priority for job1045 changes from 10 to 13.*

**-q - Query**

**Format** [ diag( ALL)| hostlist | starttime| template] <JOBEXP>

**Description** Query a job.

**Example**

```
> mjobctl -q diag job1045
```

*Query job job1045.*

```
> mjobctl -q diag ALL --format=xml
```

*Query all jobs and return the output in machine-readable XML.*

```
> mjobctl -q starttime job1045
```

*Query the estimated starttime of job job1045. The method used to estimate the start time can be specified by adding the --flags option with a value of prio, rsv, hist or all. If the --flags option is not specified, the value of the DEFAULTSTARTTIMEQUERY parameter will determine the default estimation method to be used -- which defaults to PRIORITY.*

```
> mjobctl -q template <job>
```

*Query job templates. If the <job> is set to ALL or empty, it will return information for all job templates.*

```
> mjobctl -q wiki <jobName>
```

*Query a job with the output displayed in a WIKI string. The job's name can be replaced with ALL.*



**--flags=completed will only work with the diag option.**

<b>-r - Resume</b>	
<b>Format</b>	<b>JOBEXP</b>
<b>Description</b>	Resume a job.
<b>Example</b>	<pre>&gt; mjobctl -r job1045</pre> <pre>Resume job job1045.</pre>

<b>-R - Requeue</b>	
<b>Format</b>	<b>JOBEXP</b> [--flags=force unmigrate]
<b>Description</b>	<p>Requeue a job.</p> <p>Adding <code>--flags=force</code> forces an asynchronous requeue on Torque systems.</p> <p>Adding <code>--flags=unmigrate</code> causes Moab to pull a grid job back to the central scheduler for further evaluation on all valid partitions.</p>
<b>Example</b>	<pre>&gt; mjobctl -R job1045</pre> <pre>Requeue job job1045.</pre>

<b>-s - Suspend</b>	
<b>Format</b>	<b>JOBEXP</b>
<b>Description</b>	Suspend a job. For more information, see <a href="#">Suspend/Resume Handling</a> .
<b>Example</b>	<pre>&gt; mjobctl -s job1045</pre> <pre>Suspend job job1045.</pre>

<b>-u - Unhold</b>	
<b>Format</b>	<p>[&lt;TYPE&gt;[,&lt;TYPE&gt;]]<b>JOBEXP</b></p> <p>&lt;TYPE&gt; = [ user   system   batch   defer   ALL ]</p>

<b>-u - Unhold</b>	
<b>Default</b>	ALL
<b>Description</b>	Release a hold on a job. See <a href="#">Job Holds</a> for more information.
<b>Example</b>	<pre>&gt; mjobctl -u user,system scrib.1045</pre> <p><i>Release user and system holds on job scrib.1045.</i></p>

<b>-w - Where</b>	
<b>Format</b>	[CompletionTime   StartTime][<=   =   >=]<EPOCH_TIME>
<b>Description</b>	Add a where constraint clause to the current command. As it pertains to CompletionTime   StartTime, the where constraint only works for completed jobs. CompletionTime filters according to the completed jobs' completion times; StartTime filters according to the completed jobs' start times.
<b>Example</b>	<pre>&gt; mjobctl -q diag ALL --flags=COMPLETED --format=xml -w CompletionTime&gt;=1246428000 -w CompletionTime&lt;=1254376800</pre> <p><i>Prints all completed jobs still in memory that completed between July 1, 2009 and October 1, 2009.</i></p>

<b>-x - Execute</b>	
<b>Format</b>	<b>JOBEXP</b>
<b>Description</b>	Execute a job. The <code>-w</code> option allows flags to be set for the job. Allowable flags are, ignorepolicies, ignorenodestate, and ignorersv.
<b>Example</b>	<pre>&gt; mjobctl -x job1045</pre> <p><i>Execute job job1045.</i></p> <pre>&gt; mjobctl -x -w flags=ignorepolicies job1046</pre> <p><i>Execute job job1046 and ignore policies, such as MaxJobPerUser.</i></p>

## 3.7.20.D Parameters

JOB EXPRESSION	
<b>Format</b>	<STRING>
<b>Description</b>	<p>The name of a job or a regular expression for several jobs. The flags that support job expressions can use node expression syntax as described in <a href="#">Node Selection</a>. Using <code>x:</code> indicates the following string is to be interpreted as a regular expression, and using <code>r:</code> indicates the following string is to be interpreted as a range. Job expressions do not work for array subjobs.</p> <div style="border: 1px solid black; padding: 5px; margin-bottom: 5px;"> <p><b>i</b> Moab uses regular expressions conforming to the POSIX 1003.2 standard. This standard is somewhat different than the regular expressions commonly used for filename matching in UNIX environments (see <code>man 7 regex</code>). To interpret a job expression as a regular expression, use <code>x:.</code></p> </div> <div style="border: 1px solid black; padding: 5px; margin-bottom: 5px;"> <p><b>i</b> In most cases, it is necessary to quote the job expression (for example, <code>job13[5-9]</code>) to prevent the shell from intercepting and interpreting the special characters.</p> </div> <div style="border: 1px solid black; padding: 5px;"> <p><b>i</b> The <code>mjobctl</code> command accepts a comma-delimited list of job expressions. Example usage might be <code>mjobctl -r job[1-2],job4</code> or <code>mjobctl -c job1,job2,job4</code>.</p> </div>
<b>Example</b>	<pre>&gt; mjobctl -c "x:80.*" job '802' cancelled job '803' cancelled job '804' cancelled job '805' cancelled job '806' cancelled job '807' cancelled job '808' cancelled job '809' cancelled</pre> <p style="text-align: center;"><i>Cancel all jobs starting with 80.</i></p> <pre>&gt; mjobctl -m priority+=200 "x:74[3-5]" job '743' system priority modified job '744' system priority modified job '745' system priority modified</pre> <pre>&gt; mjobctl -h x:17.* # This puts a hold on any job that has a 17 that is followed by an unlimited # amount of any # character and includes jobs 1701, 17mjk10, and 17DjN_JW-07  &gt; mjobctl -h r:1-17 # This puts a hold on jobs 1 through 17.</pre>



### 3.7.20.E XML Output

*mjobctl* information can be reported as XML as well. This is done with the command `mjobctl -q diag <JOB_ID>`.

#### XML Attributes

Name	Description
<b>Account</b>	The account assigned to the job
<b>AllocNodeList</b>	The nodes allocated to the job
<b>Args</b>	The job's executable arguments
<b>AWDuration</b>	The active wall time consumed
<b>BlockReason</b>	The block message index for the reason the job is not eligible
<b>Bypass</b>	Number of times the job has been bypassed by other jobs
<b>Calendar</b>	The job's timeframe constraint calendar
<b>Class</b>	The class assigned to the job
<b>CmdFile</b>	The command file path
<b>CompletionCode</b>	The return code of the job as extracted from the RM
<b>CompletionTime</b>	The time of the job's completion
<b>Cost</b>	The cost of executing the job relative to an accounting manager
<b>CPULimit</b>	The CPU limit for the job
<b>Depend</b>	Any dependencies on the status of other jobs
<b>DRM</b>	The master destination RM
<b>DRMJID</b>	The master destination RM job ID

Name	Description
<b>EEDuration</b>	The duration of time the job has been eligible for scheduling
<b>EFile</b>	The stderr file
<b>Env</b>	The job's environment variables set for execution
<b>EnvOverride</b>	The job's overriding environment variables set for execution
<b>EState</b>	The expected state of the job
<b>EstHistStartTime</b>	The estimated historical start time
<b>EstPrioStartTime</b>	The estimated priority start time
<b>EstRsvStartTime</b>	The estimated reservation start time
<b>ExcHList</b>	The excluded host list
<b>Flags</b>	Command delimited list of Moab flags on the job
<b>GAttr</b>	The requested generic attributes
<b>GJID</b>	The global job ID
<b>Group</b>	The group assigned to the job
<b>Hold</b>	The hold list
<b>Holdtime</b>	The time the job was put on hold
<b>HopCount</b>	The hop count between the job's peers
<b>HostList</b>	The requested host list
<b>IFlags</b>	The internal flags for the job
<b>IsInteractive</b>	If set, the job is interactive

Name	Description
<b>IsRestartable</b>	If set, the job is restartable
<b>IsSuspendable</b>	If set, the job is suspendable
<b>IWD</b>	The directory where the job is executed
<b>JobID</b>	The job's batch ID.
<b>JobName</b>	The user-specified name for the job
<b>JobGroup</b>	The job ID relative to its group
<b>LogLevel</b>	The individual log level for the job
<b>MasterHost</b>	The specified host to run primary tasks on
<b>Messages</b>	Any messages reported by Moab regarding the job
<b>MinPreemptTime</b>	The minimum amount of time the job must run before being eligible for preemption
<b>Notification</b>	Any events generated to notify the job's user
<b>OFile</b>	The stdout file
<b>OldMessages</b>	Any messages reported by Moab in the old message style regarding the job
<b>OWCLimit</b>	The original wallclock limit
<b>PAL</b>	The partition access list relative to the job
<b>QueueStatus</b>	The job's queue status as generated this iteration
<b>QOS</b>	The QoS assigned to the job
<b>QOSReq</b>	The requested QoS for the job
<b>ReqAWDuration</b>	The requested active walltime duration

Name	Description
<b>ReqCMaxTime</b>	The requested latest allowed completion time
<b>ReqMem</b>	The total memory requested/dedicated to the job
<b>ReqNodes</b>	The number of requested nodes for the job
<b>ReqProcs</b>	The number of requested procs for the job
<b>ReqReservation</b>	The required reservation for the job
<b>ReqRMType</b>	The required RM type
<b>ReqSMinTime</b>	The requested earliest start time
<b>RM</b>	The master source resource manager
<b>RMXString</b>	The resource manager extension string
<b>RsvAccess</b>	The list of reservations accessible by the job
<b>RsvStartTime</b>	The reservation start time
<b>RunPriority</b>	The effective job priority
<b>Shell</b>	The execution shell's output
<b>SID</b>	The job's system ID (parent cluster)
<b>Size</b>	The job's computational size
<b>STotCPU</b>	The average CPU load tracked across all nodes
<b>SMaxCPU</b>	The max CPU load tracked across all nodes
<b>STotMem</b>	The average memory usage tracked across all nodes
<b>SMaxMem</b>	The max memory usage tracked across all nodes

Name	Description
<b>SRMJID</b>	The source RM's ID for the job
<b>StartCount</b>	The number of the times the job has tried to start
<b>StartPriority</b>	The effective job priority
<b>StartTime</b>	The most recent time the job started executing
<b>State</b>	The state of the job as reported by Moab
<b>StatMSUtl</b>	The total number of memory seconds utilized
<b>StatPSDed</b>	The total number of processor seconds dedicated to the job
<b>StatPSUtl</b>	The total number of processor seconds utilized by the job
<b>StdErr</b>	The path to the stderr file
<b>StdIn</b>	The path to the stdin file
<b>StdOut</b>	The path to the stdout file
<b>StepID</b>	StepID of the job (used with LoadLeveler systems)
<b>SubmitHost</b>	The host where the job was submitted
<b>SubmitLanguage</b>	The RM language that the submission request was performed
<b>SubmitString</b>	The string containing the entire submission request
<b>SubmissionTime</b>	The time the job was submitted
<b>SuspendDuration</b>	The amount of time the job has been suspended
<b>SysPrio</b>	The admin specified job priority
<b>SysSMinTime</b>	The system specified min. start time

Name	Description
<b>TaskMap</b>	The allocation taskmap for the job
<b>TermTime</b>	The time the job was terminated
<b>User</b>	The user assigned to the job
<b>UserPrio</b>	The user specified job priority
<b>UtlMem</b>	The utilized memory of the job
<b>UtlProcs</b>	The number of utilized processors by the job
<b>Variable</b>	
<b>VWCTime</b>	The virtual wallclock limit

### 3.7.20.F Examples

#### Example 3-21:

```

> mjobctl -q diag ALL --format=xml
<Data><job AWDuration="346" Class="batch" CmdFile="jobsleep.sh" EEDuration="0"
EState="Running" Flags="RESTARTABLE" Group="test" IWD="/home/test" JobID="11578"
QOS="high"
RMJID="11578.lolo.icluster.org" ReqAWDuration="00:10:00" ReqNodes="1" ReqProcs="1"
StartCount="1"
StartPriority="1" StartTime="1083861225" StatMSUtl="903.570" StatPSDed="364.610"
StatPSUtl="364.610"
State="Running" SubmissionTime="1083861225" SuspendDuration="0" SysPrio="0"
SysSMinTime="00:00:00"
User="test"><req AllocNodeList="hana" AllocPartition="access" ReqNodeFeature="[NONE]"
ReqPartition="access"></req></job><job AWDuration="346" Class="batch"
CmdFile="jobsleep.sh"
EEDuration="0" EState="Running" Flags="RESTARTABLE" Group="test" IWD="/home/test"
JobID="11579"
QOS="high" RMJID="11579.lolo.icluster.org" ReqAWDuration="00:10:00" ReqNodes="1"
ReqProcs="1"
StartCount="1" StartPriority="1" StartTime="1083861225" StatMSUtl="602.380"
StatPSDed="364.610"
StatPSUtl="364.610" State="Running" SubmissionTime="1083861225" SuspendDuration="0"
SysPrio="0"
SysSMinTime="00:00:00" User="test"><req AllocNodeList="lolo" AllocPartition="access"
ReqNodeFeature="[NONE]" ReqPartition="access"></req></job></Data>

```

## Related Topics

- (Optional) Install Moab Client - in the *Moab HPC Suite Installation and Configuration Guide* explains how to distribute this command to client nodes
- [setspri](#)
- [canceljob](#)
- [runjob](#)

## 3.7.21 mnodectl

### 3.7.21.A Synopsis

```
mnodectl -m attr{=|-|=+}val nodeexp
mnodectl -q [cat|diag|profile|wiki] nodeexp

[--about] [--help] [--host=<serverHostName>] [--
loglevel=<logLevel>] [--msg=<message>] [--port=<serverPort>]
[--timeout=<seconds>] [--version] [--xml]
```

### 3.7.21.B Overview

Change specified attributes for a given [node expression](#).

### 3.7.21.C Access

By default, this command can be run by any Moab Administrator.

### 3.7.21.D Options


#### -m - Modify


##### Format

```
<ATTR>{=|-|=+}<VAL>
```

Where <ATTR> is one of the following:

```
CFGCLASS,
FEATURES,
GEVENT,
```

-m - Modify	
	<p><b>GMETRIC,</b> <b>MESSAGE,</b> <b>OS,</b> <b>POWER,</b> <b>STATE,</b> <b>VARIABLE</b></p> <p>and -=, except when used for features, clears the attribute instead of decrementing the attribute's value and = indicates that you are specifying a new value to replace the old one(s), if any.</p> <p>When the -= option is used to modify features, it removes the specified features from the node. The += option, which is only available for features, allows you to append additional features to the current list rather than replacing the current list entirely.</p> <div style="border: 1px solid #4a6984; border-radius: 10px; padding: 5px; margin-top: 10px;"> <p> Changing OS and POWER require a Moab Adaptive Computing Suite license and a provisioning resource manager.</p> </div>
<b>Description</b>	Modify the state or attribute of <b>specified</b> node(s).
<b>Example</b>	<pre style="border: 1px dashed #ccc; padding: 10px;"> &gt; mnodectl -m cfgclass==debug node1 &gt; mnodectl -m features+=fastio,highmem node1 &gt; mnodectl -m gevent=cpufail:'cpu02 has failed w/ec:0317' node1 &gt; mnodectl -m gmetric=temp:131.2 node1 &gt; mnodectl -m message='cpufailure:cpu02 has failed w/ec:0317' node1 &gt; mnodectl -m OS=RHAS30 node1 &gt; mnodectl -m power=off node1 &gt; mnodectl -m state=idle node1 &gt; mnodectl -m variable=IP=10.10.10.100,Location=R1S2 node1 </pre>

-q - Query	
<b>Format</b>	{cat   diag   profile   wiki}
<b>Description</b>	<p>Query node categories or node profile information (see <a href="#">ENABLEPROFILING</a> for nodes).</p> <div style="border: 1px solid #4a6984; border-radius: 10px; padding: 5px; margin-top: 10px;"> <p> The diag and profile options must use --xml.</p> </div>
<b>Example</b>	



**-q - Query**

```
> mnodectl -q cat ALL
node categorization stats from Mon Jul 10 00:00:00 to Mon Jul 10 15:30:00
Node: moab
Categories:
                busy: 96.88%
                idle: 3.12%
Node: maka
Categories:
                busy: 96.88%
                idle: 3.12%
Node: pau
Categories:
                busy: 96.88%
                idle: 3.12%
Node: maowu
Categories:
                busy: 96.88%
                down-hw: 3.12%
Cluster Summary:
                busy: 96.88%
                down-hw: 0.78%
                idle: 2.34%
```

```
> mnodectl -v -q profile
...
```

```
> mnodectl -q wiki <ALL>
GLOBAL STATE=Idle PARTITION=SHARED
n0 STATE=Idle PARTITION=base APROC=4 CPROC=4 RM=base NODEACCESSPOLICY=SHARED
n1 STATE=Idle PARTITION=base APROC=4 CPROC=4 RM=base NODEACCESSPOLICY=SHARED
n2 STATE=Idle PARTITION=base APROC=4 CPROC=4 RM=base NODEACCESSPOLICY=SHARED
n3 STATE=Idle PARTITION=base APROC=4 CPROC=4 RM=base NODEACCESSPOLICY=SHARED
n4 STATE=Idle PARTITION=base APROC=4 CPROC=4 RM=base NODEACCESSPOLICY=SHARED
n5 STATE=Idle PARTITION=base APROC=4 CPROC=4 RM=base NODEACCESSPOLICY=SHARED
n6 STATE=Idle PARTITION=base APROC=4 CPROC=4 RM=base NODEACCESSPOLICY=SHARED
n7 STATE=Idle PARTITION=base APROC=4 CPROC=4 RM=base NODEACCESSPOLICY=SHARED
n8 STATE=Idle PARTITION=base APROC=4 CPROC=4 RM=base NODEACCESSPOLICY=SHARED
n9 STATE=Idle PARTITION=base APROC=4 CPROC=4 RM=base NODEACCESSPOLICY=SHARED
```

*Query a node with the output displayed in a WIKI string.*

**Parameters****CFGCLASS****Format**

&lt;STRING&gt;

**Description**

Class name.

**i** Only "-" is supported when modifying cfgclass on a node. To add or set classes on a node, see [HOSTLIST](#)

CFGCLASS	
<b>Example</b>	<pre>&gt; mnodectl -m cfgclass==debug node1</pre>
FEATURES	
<b>Format</b>	<p>&lt;STRING&gt; One of the following:</p> <ul style="list-style-type: none"> <li>• a comma-delimited list of features</li> <li>• [NONE] (to clear features on the node)</li> </ul>
<b>Description</b>	<p>Sets the features on a node.</p> <div style="border: 1px solid #005596; border-radius: 10px; padding: 5px; background-color: #e6f2ff;"> <p><b>i</b> These node features will be overwritten when an RM reports features.</p> </div>
<b>Example</b>	<pre>mnodectl -m features=fastio,highmem node1 mnodectl -m features=[NONE] node1</pre>
GEVENT	
<b>Format</b>	<EVENT>:<MESSAGE>
<b>Description</b>	Creates a generic event on the node to which Moab can respond (see <a href="#">Enabling Generic Events</a> ).
<b>Example</b>	<pre>mnodectl -m gevent=powerfail:'power has failed' node1</pre>
GMETRIC	
<b>Format</b>	<ATTR>:<VALUE>
<b>Description</b>	<p>Sets the value for a generic metric on the node (see <a href="#">Enabling Generic Metrics</a>).</p> <div style="border: 1px solid #005596; border-radius: 10px; padding: 5px; background-color: #e6f2ff;"> <p><b>i</b> When a gmetric set in Moab conflicts with what the resource manager reports, Moab uses the set gmetric until the next time the resource manager reports a different number.</p> </div>

## GMETRIC

### Example

```
mnodectl -m gmetric=temp:120 node1
```

## MESSAGE

### Format

'<MESSAGE>'

### Description

Sets a message to be displayed on the node.

### Example

```
mnodectl -m message='powerfailure: power has failed'
node1
```

## NODEEXP

### Format

<STRING>

Where <NODEEXP> is a node name, regex or ALL

**i** Node regex has the potential to unintentionally match many nodes (for example, specifying n1 will match n10, n11, n12, n100, etc). To ensure correct matching, explicitly use the "x:<node\_regex>" when modifying multiple nodes in one command. Currently, this is supported for features.

### Description

Identifies one or more nodes.

### Example

node1 - applies only to node1  
 fr10n\* - all nodes starting with fr10n  
 ALL - all known nodes

## OS

### Format

<STRING>

### Description

Operating System (see [Resource Provisioning](#)).

### Example

```
mnodectl node1 -m OS=RHELAS30
```

POWER	
<b>Format</b>	{off on}
<b>Description</b>	<p>Set the power state of a node. Action will NOT be taken if the node is already in the specified state.</p> <div style="border: 1px solid #004a7c; padding: 5px; margin-bottom: 10px;"> <p><b>i</b> If you power off a node, a green policy will try to turn it back on. If you want the node to remain powered off, you must associate a reservation with it.</p> </div> <div style="border: 1px solid #004a7c; padding: 5px;"> <p><b>i</b> If you request to power off a node that has active work on it, Moab returns a status indicating that the node is busy with a job and will not be powered off. You will see one of these messages:</p> <ul style="list-style-type: none"> <li>• Ignoring node &lt;name&gt;: power ON in process (indicates node is currently powering on)</li> <li>• Ignoring node &lt;name&gt;: power OFF in process (indicates node is currently powering off)</li> <li>• Ignoring node &lt;name&gt;: has active jobs running (indicates the node is currently running active jobs)</li> </ul> <p>Once you resolve the activity on the node (by preempting or migrating the jobs, for example), you can attempt to power the node off again. You can use the <code>--flags=force</code> option to cause a force override. However, doing this will power off the node regardless of whether or not its jobs get migrated or preempted (i.e., you run the risk of losing the jobs entirely). For example:</p> <div style="border: 1px dashed #ccc; padding: 5px; margin-top: 10px;"> <pre>&gt; mnodectl node1 -m power=off --flags=force</pre> </div> </div>
<b>Example</b>	<div style="border: 1px dashed #ccc; padding: 5px;"> <pre>&gt; mnodectl node1 -m power=off</pre> </div>

STATE	
<b>Format</b>	{drained idle}
<b>Description</b>	Remove (drained) or add (idle) a node from scheduling.
<b>Example</b>	<div style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 10px;"> <pre>mnodectl node1 -m state=drained</pre> </div> <div style="border: 1px dashed #ccc; padding: 5px;"> <p><i>Moab ignores node1 when scheduling.</i></p> </div>

VARIABLE	
<b>Format</b>	<name>[=<value>],<name>[=<value>]...
<b>Description</b>	Set a list of variables for a node.
<b>Example</b>	<pre>&gt; mnodectl node1 -m variable=IP=10.10.10.100,Location=R1S2</pre>

### Related Topics

- (Optional) Install Moab Client - in the *Moab HPC Suite Installation and Configuration Guide* explains how to distribute this command to client nodes
- [mdiag -n](#)
- [showres -n](#)
- [checknode](#)
- [showstats -n](#) — report current and historical node statistics

## 3.7.22 moab

### 3.7.22.A Synopsis

```
moab --about --help --loglevel=<LOGLEVEL> --version [-c <CONFIG_FILE>] [-C] [-d] [-e] [-h] [-P [<PAUSEDURATION>]] [-R <RECYCLEDURATION>] [-s] [-S [<STOPITERATION>]] [-v]
```

### 3.7.22.B Options

Option	Description
<b>--about</b>	Displays build environment and version information.
<b>--loglevel</b>	Sets the server loglevel to the specified value.

Option	Description
<b>--version</b>	Displays version information.
<b>-c</b>	Configuration file the server should use.
<b>-C</b>	Clears checkpoint files ( <code>.moab.ck</code> , <code>.moab.ck.1</code> ).
<b>-d</b>	Debug mode (does not background itself).
<b>-e</b>	Forces Moab to exit if there are any errors in the configuration file, if it can't connect to the configured database, or if it can't find these directories: <ul style="list-style-type: none"> <li>• <code>statdir</code></li> <li>• <code>logdir</code></li> <li>• <code>spooldir</code></li> <li>• <code>toolsdir</code></li> </ul>
<b>-P</b>	Starts Moab in a paused state for the duration specified (default: pause indefinitely; resume with <code>mschedctl -r</code> (or <code>-R</code> ), or a service restart).
<b>-R</b>	Causes Moab to automatically <b>recycle</b> every time the specified duration transpires.
<b>-s</b>	Starts Moab in the state that was most recently checkpointed.
<b>-S</b>	Suspends/stops scheduling at specified iteration (or at startup if no iteration is specified).
<b>-v</b>	Same as <code>--version</code> .

## 3.7.23 mrmctl

### 3.7.23.A Synopsis

```
mrmctl -f [<fobject>] {<rmid> | AM[:<amid>] | ID[:<imid>]}
```

```
mrmctl -l [<rmid> | AM[:<amid>]]
```

```

mrmctl -m <attr>=<value> [<rmid>]
mrmctl -p {<rmid> | AM[:<amid>]}
mrmctl -q AccountBalanceCache AM[:<amid>]
mrmctl -R {AM[:<amid>] | ID[:<imid>]}
[--about] [--help] [--host=<serverHostName>] [--loglevel=<logLevel>] [--msg=<message>] [--port=<serverPort>] [--timeout=<seconds>] [--version] [--xml]

```

### 3.7.23.B Overview

`mrmctl` allows an admin to query, list, modify, and ping the [resource managers](#) and [accounting managers](#) in Moab. `mrmctl` also allows for a queue (often referred to as a class) to be created for a resource manager.

### 3.7.23.C Access

By default, this command can be run by level 1 and level 2 Moab administrators (see [ADMINCFG](#)).

### 3.7.23.D Options

-f	
<b>Format</b>	-f [<fobject>] where fobject is optional and one of messages or stats.
<b>Default</b>	If no <code>fobject</code> is specified, then reported failures and performance data will be flushed. If no resource manager ID is specified, the first resource manager will be flushed.
<b>Description</b>	Clears resource manager statistics. If <code>messages</code> is specified, then reported failures, performance data, and messages will be flushed.
<b>Example</b>	<pre>&gt; mrmctl -f base</pre> <p><i>Moab will clear the statistics for RM base.</i></p>

-l	
<b>Format</b>	-l
<b>Default</b>	All RMs and AMs (when no RM/AM is specified).
<b>Description</b>	List Resource and Accounting Manager(s).
<b>Example</b>	<pre>&gt; mrmctl -l</pre> <p><i>Moab will list all resource and accounting managers.</i></p>

-m	
<b>Format</b>	-m <attr>=<val>
<b>Default</b>	All RMs and AMs (when no RM/AM is specified).
<b>Description</b>	Modify Resource and Accounting Manager(s).
<b>Example</b>	<pre>&gt; mrmctl -m state=disabled peer13</pre>

-p	
<b>Format</b>	-p
<b>Default</b>	First RM configured.
<b>Description</b>	Ping Resource Manager.
<b>Example</b>	<pre>&gt; mrmctl -p base</pre> <p><i>Moab will ping RM base.</i></p>

-q	
<b>Format</b>	-q AccountBalanceCache



-q	
<b>Default</b>	---
<b>Description</b>	When an accounting manager is being used and the fast-allocation accounting mode is configured, this option queries Moab's internal cache of account balances. See <a href="#">5.5.2 Accounting Mode - page 419</a> . Also see 'Select an Appropriate Accounting Mode' in the <i>Moab Accounting Manager Administrator Guide</i> for more information.
<b>Example</b>	<pre>&gt; mrmctl -q AccountBalanceCache AM</pre>

-R	
<b>Format</b>	-R
<b>Description</b>	Dynamically reloads <a href="#">server</a> information for the <a href="#">identity manager</a> service if ID is specified; if AM is specified, reloads the accounting manager service.
<b>Example</b>	<pre>&gt; mrmctl -R ID</pre> Reloads the identity manager on demand.

**i** Resource manager interfaces can be enabled/disabled using the modify operation to change the resource manager state as in the following example:

```
# disable active resource manager interface
> mrmctl -m state=disabled torque
# restore disabled resource manager interface
> mrmctl -m state=enabled torque
```

## Related Topics

- (Optional) Install Moab Client - in the *Moab HPC Suite Installation and Configuration Guide* explains how to distribute this command to client nodes
- [mdiag -R](#)
- [mdiag -c](#)

## 3.7.24 mrsvctl

### 3.7.24.A Synopsis

`mrsvctl -BSRVSID`

`mrsvctl -c [-a ACL] [-b SUBTYPE] [-d DURATION] [-D DESCRIPTION] [-e ENDTIME] [-E EXCLUSIVE] [-f FEATURES] [-F FLAGS] [-g RSVGROU] [-h HOSTLIST] [-n NAME] [-o OWNER] [-p PARTITION] [-P PROFILE] [-R RESOURCES] [-s STARTTIME] [-S SET ATTRIBUTE] [-t TASKS] [-T TRIGGER] [-V VARIABLE] [-x JOBLIST]`

`mrsvctl -C [-g SRSVID] {RESERVATION PATTERN}`

`mrsvctl -l [{RESERVATION PATTERN | -i INDEX}]`

`mrsvctl -m`

`<duration|endtime|hostexp|loglevel|reqtaskcount|rsvaccesslist| rsvgroup|starttime|variable>{=|+=|-=}<VAL> <hostexp>{+=|-=}<VAL> <variable>{+=KEY=VAL|-KEY_TO_REMOVE} {RESERVATION PATTERN | -i INDEX}`

`mrsvctl -q {RESERVATION PATTERN | -i INDEX} [--blocking]`

`mrsvctl -r {RESERVATION PATTERN | -i INDEX}`

`[--about] [--help] [--host=<serverHostName>] [-- loglevel=<logLevel>] [--msg=<message>] [--port=<serverPort>] [--timeout=<seconds>] [--version] [--xml]`

### 3.7.24.B Overview

`mrsvctl` controls the creation, modification, querying, and releasing of reservations.

The timeframe covered by the reservation can be specified on either an absolute or relative basis. Only jobs with credentials listed in the reservation's access control list can utilize the reserved resources. However, these jobs still have the freedom to utilize resources outside of the reservation. The reservation will be assigned a name derived from the ACL specified. If no reservation ACL is specified, the reservation is created as a system reservation and no jobs will be allowed access to the resources during the specified timeframe (valuable for system maintenance, etc.). See the [Reservation Overview](#) for more information.

Reservations can be viewed using the `-q` flag and can be released using the `-r` flag.

**i** By default, reservations are not exclusive and can overlap with other reservations and jobs. Use the `'-E'` flag to adjust this behavior.

### 3.7.24.C Access

By default, this command can be run by level 1 and level 2 Moab administrators (see [ADMINCFG](#)).

### 3.7.24.D Options

-a	
<b>Name</b>	ACL
<b>Format</b>	<p>&lt;TYPE&gt;==&lt;VAL&gt; [ , &lt;TYPE&gt;==&lt;VAL&gt; ] . . .</p> <p>Where &lt;TYPE&gt; is one of the following:  ACCT,  CLASS,  DURATION,  GROUP,  JATTR,  PROC,  QOS,  USER</p>
<b>Description</b>	List of limitations for access to the reserved resources (see also <a href="#">ACL Modifiers</a> ).
<b>Example</b>	<pre>&gt; mrsvctl -c -h node01 -a USER==john+,CLASS==batch-</pre> <p><i>Moab will make a reservation on node01 allowing access to user john and restricting access from class batch when other resources are available to class batch</i></p> <pre>&gt; mrsvctl -m -a USER==john system.1</pre> <p><i>Moab will remove user john from the system.1 reservation</i></p>
<b>Notes</b>	<ul style="list-style-type: none"> <li>When you specify multiple credentials, a user must only match one of them in order to access the reservation. To require one or more of the listed limitations for reservation access, each required specification must end with an asterisk (*). If a user meets the required limitation(s), that user has access to the reservation (without meeting any that are not marked required).</li> <li>There are three different assignment operators that can be used for modifying most credentials in the ACL. The operator == will reassess the</li> </ul>

-a

list for that particular credential type. The += operator will append to the list for that credential type, and -= will remove from the list. Two other operators are used to specify DURATION and PROC: >= (greater than) and <= (less than).

- To add multiple credentials of the same type with one command, use a colon to separate them. To separate lists of different credential types, use commas. For example, to reassign the user list to consist of users Joe and Bob, and to append the group MyGroup to the groups list on the `system.1` reservation, you could use the command `mrsvctl -m -a USER==Joe:Bob,GROUP+=MyGroup system.1`.
- Any of the ACL modifiers can be used. When using them, it is often useful to put single quotes on either side of the assignment command. For example, `mrsvctl -m -a 'USER==&Joe' system.1`.
- Some flags are mutually exclusive. For example, the ! modifier means that the credential is blocked from the reservation and the & modifier means that the credential must run on that reservation. Moab will take the most recently parsed modifier. Modifiers can be placed on either the left or the right of the argument, so `USER==&JOE` and `USER==JOE&` are equivalent. Moab parses each argument starting from right to left on the right side of the argument, then from left to right on the left side. So, if the command was `USER==!Joe&`, Moab would keep the equivalent of `USER==!Joe` because the ! would be the last one parsed.
- You can set a reservation to have a time limit for submitted jobs using DURATION and the \* modifier. For example, `mrsvctl -m -a 'DURATION<=*1:00:00' system.1` would cause the `system.1` reservation to not accept any jobs with a walltime greater than one hour. Similarly, you can set a reservation to have a processor limit using PROC and the \* modifier. `mrsvctl -a 'PROC>=2*' system.2` would cause the `system.2` reservation to only allow jobs requesting more than 2 procs to run on it.
- You can verify the ACL of a reservation using the `mdiag -r` command.

```
mrsvctl -m -a 'USER==Joe:Bob,GROUP-=BadGroup,ACCT+=GoodAccount,DURATION<=*1:00:00' system.1
```

*Moab will reassign the USER list to be Joe and Bob, will remove BadGroup from the GROUP list, append GoodAccount to the ACCT list, and only allow jobs that have a submitted walltime of an hour or less on the system.1 reservation.*

```
mrsvctl -m -a 'USER==Joe,USER==Bob' system.1
```

*Moab will assign the USER list to Joe, and then reassign it again to Bob. The final result will be that the USER list will just be Bob. To add Joe and Bob, use `mrsvctl -m -a USER==Joe:Bob system.1` or `mrsvctl -m -a USER==Joe,USER+=Bob`*

-a	
	<i>system.1.</i>

-b	
<b>Name</b>	SUBTYPE
<b>Format</b>	One of the <a href="#">node category</a> values or node category shortcuts.
<b>Description</b>	Add subtype to reservation.
<b>Example</b>	<div style="border: 1px dashed gray; padding: 5px; margin-bottom: 5px;"> <pre>&gt; mrsvctl -c -b SoftwareMaintenance -t ALL</pre> </div> <div style="border: 1px dashed gray; padding: 5px;"> <p><i>Moab will associate the reserved nodes with the <a href="#">node category</a> SoftwareMaintenance.</i></p> </div>

-B	
<b>Name</b>	REBUILD
<b>Format</b>	<SRSVID>
<b>Description</b>	Rebuilds standing reservations while Moab is running.
<b>Example</b>	<div style="border: 1px dashed gray; padding: 5px;"> <pre>&gt; mrsvctl -B &lt;SRSVID&gt;</pre> </div>

-c	
<b>Name</b>	CREATE
<b>Format</b>	<ARGUMENTS>
<b>Description</b>	Creates a reservation.

-C	
<b>Examples</b>	<div style="border: 1px solid #ccc; padding: 10px; margin-bottom: 10px;"> <p><b>i</b> If a created reservation has a given duration but the start time is in the past, one of the following actions occur depending on whether the present time falls within the reservation's given duration:</p> <ul style="list-style-type: none"> <li>• If the present time is still within the reservation's duration time frame, the start time does not change and the reservation shows however long is left in the reservation (present time minus the duration time).</li> <li>• If present time is outside of the reservation's duration time frame, the reservation start time automatically sets to the present time and the reservation continues for its full given duration.</li> </ul> </div> <div style="border: 1px solid #ccc; padding: 10px; margin-bottom: 10px;"> <p><b>i</b> The <code>-x</code> flag, when used with <code>-F ignjobrsv</code>, lets users create reservations but exclude certain nodes from being part of the reservation because they are running specific jobs. The <code>-F</code> flag instructs <code>mrsvctl</code> to still consider nodes with current running jobs.</p> </div> <div style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 5px;"> <pre>&gt; mrsvctl -c -t ALL</pre> </div> <div style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 5px;"> <p><i>Moab will create a reservation across all system resources.</i></p> </div> <div style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 5px;"> <pre>&gt; mrsvctl -c -t 5 -F ignjobrsv -x moab.5,moab.6</pre> </div> <div style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 5px;"> <p><i>Moab will create the reservation while assigning the nodes. Nodes running jobs moab5 and moab6 will not be assigned to the reservation.</i></p> </div> <div style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 5px;"> <pre>&gt; mrsvctl -c -t 1 -d INFINITY</pre> </div> <div style="border: 1px dashed #ccc; padding: 5px;"> <p><i>Moab will create an infinite reservation.</i></p> </div>

-C	
<b>Name</b>	CLEAR
<b>Format</b>	<RSVID>   -g <SRSVID>
<b>Description</b>	Clears any disabled time slots from standing reservations and allows the recreation of disabled reservations.
<b>Example</b>	<div style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 5px;"> <pre>&gt; mrsvctl -C -g testing</pre> </div> <div style="border: 1px dashed #ccc; padding: 5px;"> <p><i>Moab will clear any disabled timeslots from the standing</i></p> </div>

-C	
	<i>reservation testing.</i>

-d	
<b>Name</b>	DURATION
<b>Format</b>	[ [ [DD:]HH:]MM:]SS
<b>Default</b>	INFINITY
<b>Description</b>	Duration of the reservation (not needed if ENDTIME is specified).
<b>Example</b>	<pre style="border: 1px dashed black; padding: 5px;">&gt; mrsvctl -c -h node01 -d 5:00:00</pre> <p style="border: 1px dashed black; padding: 5px; margin-left: 20px;"><i>Moab will create a reservation on node01 lasting 5 hours.</i></p> <pre style="border: 1px dashed black; padding: 5px; margin-top: 10px;">&gt; mrsvctl -c -h node01 -d INFINITY</pre> <p style="border: 1px dashed black; padding: 5px; margin-left: 20px;"><i>Moab will create a reservation with a duration of INFINITY (no endtime).</i></p>

-D	
<b>Name</b>	DESCRIPTION
<b>Format</b>	<STRING>
<b>Description</b>	Human-readable description of reservation or purpose.
<b>Example</b>	<pre style="border: 1px dashed black; padding: 5px;">&gt; mrsvctl -c -h node01 -d 5:00:00 -D 'system maintenance to test network'</pre> <p style="border: 1px dashed black; padding: 5px; margin-left: 20px;"><i>Moab will create a reservation on node01 lasting 5 hours.</i></p>

-e	
<b>Name</b>	ENDTIME
<b>Format</b>	[HH[:MM[:SS]]][_MO[/DD[/YY]]] or +[[DD:]HH:]MM:]SS
<b>Default</b>	INFINITY
<b>Description</b>	Absolute or relative time reservation will end (not required if Duration specified). ENDTIME also supports an epoch timestamp.
<b>Example</b>	<pre>&gt; mrvsctl -c -h node01 -e +3:00:00</pre> <p><i>Moab will create a reservation on node01 ending in 3 hours.</i></p>

-E	
<b>Name</b>	EXCLUSIVE
<b>Description</b>	When specified, Moab will only create a reservation if there are no other reservations (exclusive or otherwise) that would conflict with the time and space constraints of this reservation. If exceptions are desired, the <a href="#">rsvaccesslist</a> attribute can be set or the <a href="#">ignrsv</a> flag can be used.
<b>Example</b>	<pre>&gt; mrvsctl -c -h node01 -E</pre> <p><i>Moab will only create a reservation on node01 if no conflicting reservations are found.</i></p> <div style="border: 1px solid #4a6984; padding: 5px; margin-top: 10px;"> <p><b>i</b> This flag is only used at the time of reservation creation. Once the reservation is created, Moab allows jobs into the reservation based on the ACL. Also, once the exclusive reservation is created, it is possible that Moab will overlap it with jobs that match the ACL.</p> </div>

-f	
<b>Name</b>	FEATURES



-f	
<b>Format</b>	<STRING> [:<STRING>] . . .
<b>Description</b>	List of node features that must be possessed by the reserved resources. You can use a backslash and pipe to delimit features to indicate that Moab can use one or the other.
<b>Example</b>	<pre>&gt; mrsvctl -c -h node[0-9] -f fast\ slow</pre> <p><i>Moab will create a reservation on nodes matching the expression and which also have either the feature <code>fast</code> or the feature <code>slow</code>.</i></p>


-F	
<b>Name</b>	FLAGS
<b>Format</b>	<flag> [[, <flag>] . . .]
<b>Description</b>	Comma-delimited list of flags to set for the reservation (see <a href="#">Managing Reservations</a> for flags).
<b>Example</b>	<pre>&gt; mrsvctl -c -h node01 -F ignstate</pre> <p><i>Moab will create a reservation on <code>node01</code> ignoring any conflicting node states.</i></p>

-g	
<b>Name</b>	RSVGROUP
<b>Format</b>	<STRING>
<b>Description</b>	For a create operation, create a reservation in this reservation group. For list and modify operations, take actions on all reservations in the specified reservation group. The <code>-g</code> option can also be used in conjunction with the <code>-r</code> option to release a reservation associated with a specified group. See <a href="#">Reservation Group</a> for more information.
<b>Example</b>	<pre>&gt; mrsvctl -c -g staff -h 'node0[1-9]'</pre>

**-g**

*Moab will create a reservation on nodes matching the **node expression** given and assign it to the reservation group *staff*.*

**-h**

<b>Name</b>	HOSTLIST
<b>Format</b>	class:<classname>[,<classname>]... or <STRING> or 'r:<nodeNameStart>[<beginRange>-<endRange>]' or ALL
<b>Description</b>	<p><b>Host expression</b> or a class mapping indicating the nodes that the reservation will allocate.</p> <div style="border: 1px solid red; padding: 5px; background-color: #ffe6e6;">  When you specify a &lt;STRING&gt;, the HOSTLIST attribute is always treated as a regular expression. <code>foo10</code> will map to <code>foo10</code>, <code>foo101</code>, <code>foo1006</code>, etc. To request an exact host match, the expression can be bounded by the carat and dollar op expression markers as in <code>^foo10\$</code>. </div>
<b>Example</b>	<pre>&gt; mrvctl -c -h 'r:node0[1-9]'</pre> <p><i>Moab will create a reservation on nodes <code>node01</code>, <code>node02</code>, <code>node03</code>, <code>node04</code>, <code>node05</code>, <code>node06</code>, <code>node07</code>, <code>node08</code>, and <code>node09</code>.</i></p> <pre>&gt; mrvctl -c -h class:batch</pre> <p><i>Moab will create a reservation on all nodes that support class/queue <code>batch</code>.</i></p>

**-i**

<b>Name</b>	INDEX
-------------	-------

-i	
<b>Format</b>	<STRING>
<b>Description</b>	Use the reservation index instead of full reservation ID.
<b>Example</b>	<pre>&gt; mrsvctl -m -i 1 starttime=+5:00</pre> <p><i>Moab will create a reservation on nodes matching the expression given.</i></p>

-l	
<b>Name</b>	LIST
<b>Format</b>	<RSV_ID> or ALL RSV_ID can be the name of a reservation or a regular expression.
<b>Default</b>	ALL
<b>Description</b>	List reservation(s).
<b>Example</b>	<pre>&gt; mrsvctl -l system*</pre> <p><i>Moab will list all of the reservations whose names start with system.</i></p>

-m	
<b>Name</b>	MODIFY

-m																					
<b>Format</b>	<p>&lt;ATTR&gt;=&lt;VAL&gt;[-m &lt;ATTR2&gt;=&lt;VAL2&gt;]...</p> <p>Where &lt;ATTR&gt; is one of the following:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th colspan="2" style="background-color: #4a6984; color: white; text-align: left; padding: 5px;">flags</th> </tr> </thead> <tbody> <tr> <td style="background-color: #d3d3d3; padding: 5px;"><b>duration</b></td> <td style="padding: 5px;">duration{+ = =}&lt;RELTIME&gt;</td> </tr> <tr> <td style="background-color: #d3d3d3; padding: 5px;"><b>endtime</b></td> <td style="padding: 5px;">endtime{+ = =}&lt;RELTIME&gt; or endtime=&lt;ABSTIME&gt;</td> </tr> <tr> <td style="background-color: #d3d3d3; padding: 5px;"><b>hostexp</b></td> <td style="padding: 5px;">hostexp{+ = =}&lt;node&gt;[,&lt;node&gt;]</td> </tr> <tr> <td style="background-color: #d3d3d3; padding: 5px;"><b>loglevel</b></td> <td style="padding: 5px;">loglevel[=]&lt;loglevel&gt;</td> </tr> <tr> <td style="background-color: #d3d3d3; padding: 5px;"><b>reqtaskcount</b></td> <td style="padding: 5px;">reqtaskcount{+ = =}&lt;TASKCOUNT&gt;</td> </tr> <tr> <td style="background-color: #d3d3d3; padding: 5px;"><b>rsvaccesslist</b></td> <td style="padding: 5px;">rsvaccesslist[=]&lt;reservation&gt;</td> </tr> <tr> <td style="background-color: #d3d3d3; padding: 5px;"><b>rsvgroup</b></td> <td style="padding: 5px;">rsvgroup[=]&lt;rsvgroup&gt;</td> </tr> <tr> <td style="background-color: #d3d3d3; padding: 5px;"><b>starttime</b></td> <td style="padding: 5px;">starttime{+ = =}&lt;RELTIME&gt; or starttime=&lt;ABSTIME&gt;</td> </tr> <tr> <td style="background-color: #d3d3d3; padding: 5px;"><b>variable</b></td> <td style="padding: 5px;">variable[+=key1=val1 -key_to_remove]</td> </tr> </tbody> </table>	flags		<b>duration</b>	duration{+ = =}<RELTIME>	<b>endtime</b>	endtime{+ = =}<RELTIME> or endtime=<ABSTIME>	<b>hostexp</b>	hostexp{+ = =}<node>[,<node>]	<b>loglevel</b>	loglevel[=]<loglevel>	<b>reqtaskcount</b>	reqtaskcount{+ = =}<TASKCOUNT>	<b>rsvaccesslist</b>	rsvaccesslist[=]<reservation>	<b>rsvgroup</b>	rsvgroup[=]<rsvgroup>	<b>starttime</b>	starttime{+ = =}<RELTIME> or starttime=<ABSTIME>	<b>variable</b>	variable[+=key1=val1 -key_to_remove]
flags																					
<b>duration</b>	duration{+ = =}<RELTIME>																				
<b>endtime</b>	endtime{+ = =}<RELTIME> or endtime=<ABSTIME>																				
<b>hostexp</b>	hostexp{+ = =}<node>[,<node>]																				
<b>loglevel</b>	loglevel[=]<loglevel>																				
<b>reqtaskcount</b>	reqtaskcount{+ = =}<TASKCOUNT>																				
<b>rsvaccesslist</b>	rsvaccesslist[=]<reservation>																				
<b>rsvgroup</b>	rsvgroup[=]<rsvgroup>																				
<b>starttime</b>	starttime{+ = =}<RELTIME> or starttime=<ABSTIME>																				
<b>variable</b>	variable[+=key1=val1 -key_to_remove]																				
<b>Description</b>	<p>Modify aspects of a reservation.</p> <div style="border: 1px solid #4a6984; padding: 10px; margin-top: 10px;"> <p><b>i</b> Moab is constantly scheduling and updating reservations. Before modifying a reservation it is recommended that you first stop the scheduler (<code>mschedctl -s</code>) so that the scheduler and reservation are in a stable and steady state. Once the reservation has been modified, resume the scheduler with <code>mschedctl -r</code>.</p> </div>																				
<b>Example</b>	<pre style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 10px;">&gt; mrsvctl -m duration=2:00:00 system.1</pre> <div style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 10px; text-align: center;"> <p><i>Moab sets the duration of reservation <code>system.1</code> to be exactly two hours, therefore modifying the endtime of the reservation.</i></p> </div> <pre style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 10px;">&gt; mrsvctl -m duration+=5:00:00 system.1</pre> <div style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 10px; text-align: center;"> <p><i>Moab extends the duration of <code>system.1</code> by five hours.</i></p> </div>																				

**-m**

```
> mrsvctl -m endtime+=5:00:00 system.1
```

*Moab moves the endtime of reservation `system.1` ahead by five hours.*

```
> mrsvctl -m endtime-=5:00:00 system.1
```

*Moab moves the endtime of reservation `system.1` five hours from its current endtime (without modifying the starttime; therefore, this action is equivalent to modifying the duration of the reservation).*

```
> mrsvctl -m hostexp+=node02 system.1
```

*Moab adds `node02` to the hostlist for reservation `system.1`.*

```
> mrsvctl -m loglevel=2 system.1
```

*Overrides the global `LOGLEVEL` parameter when dealing with events related to the reservation. `LOGLEVEL` values are 0-9, where 9 is most verbose.*

```
> mrsvctl -m reqtaskcount+=5 system.1
```

*Increases the `TASKCOUNT` for the `system.1` reservation by 5.*

```
> mrsvctl -m rsvaccesslist=network system.1
```

*Gives the `system.1` reservation access to the network reservation.*

```
> mrsvctl -m rsvgroup=network system.1
```

*Changes reservation `system.1` to the network `RSVGROUP`.*

```
> mrsvctl -m starttime+=5:00:00 system.1
```

*Moab advances the starttime of `system.1` five hours from its current starttime (without modifying the duration of the reservation).*

```
> mrsvctl -m starttime=15:00:00_7/6/22 system.1
```

*Moab sets the starttime of reservation `system.1` to 3:00 p.m. on July 6, 2022.*

```
> mrsvctl -m starttime-=5:00:00 system.1
```

*Moab moves the starttime of reservation `system.1` ahead five hours.*

**-m**

```
> mrsvctl -m variable+=key1=val1 system.1
```

*Moab adds the variable `key1` with the value `key2` to `system.1`.*

```
> mrsvctl -m variable+=key1=val1 variable+=key2=val2 system.1
```

*Moab adds the variable `key1` with the value `val1`, and variable `key2` with `val2` to `system.1`. (Note that each variable flag requires a distinct `-m` entry.)*

```
> mrsvctl -m variable-=key1 system.1
```

*Moab deletes the variable `key1` from `system.1`.*

```
> mrsvctl -m variable-=key1 -m variable-=key2 system.1
```

*Moab deletes the variables `key1` and `key2` from `system.1`.*

**Notes:**

- Modifying the starttime does not change the duration of the reservation, so the endtime changes as well. The starttime can be changed to be before the current time, but if the change causes the endtime to be before the current time, the change is not allowed.
- Modifying the endtime changes the duration of the reservation as well (and vice versa). An endtime *cannot* be placed before the starttime or before the current time.
- Duration cannot be negative.
- The += and -= operators operate on the time of the reservation (`starttime+=5` adds five seconds to the current reservation starttime), while + and - operate on the current time (`starttime+5` sets the starttime to five seconds from now).
- If the starttime or endtime specified is before the current time without a date specified, it is set to the next time that fits the command. To force the date, add the date as well. For the following examples, assume that the current time is 9:00 a.m. on March 1, 2022.

```
> mrsvctl -m starttime=8:00:00_3/1/22 system.1
```

*Moab moves `system.1`'s starttime to 8:00 a.m., March 1.*

```
> mrsvctl -m starttime=8:00:00 system.1
```

*Moab moves `system.1`'s starttime to 8:00 a.m., March 2.*

```
> mrsvctl -m endtime=7:00:00 system.1
```

*Moab moves `system.1`'s endtime to 7:00 a.m., March 3. This happens because the endtime must also be after the starttime, so*

**-m**

*Moab continues searching until it has found a valid time that is in the future and after the starttime.*

```
> mrsvctl -m endtime=7:00:00_3/2/22 system.1
```

*Moab will return an error because the endtime cannot be before the starttime.*

**-n**

<b>Name</b>	NAME
<b>Format</b>	<STRING>
<b>Description</b>	<p>Name for new reservation.</p> <ul style="list-style-type: none"> <li><b>i</b> If no name is specified, the reservation name is set to first name listed in ACL or SYSTEM if no ACL is specified.</li> <li><b>i</b> Reservation names cannot contain whitespace.</li> </ul>
<b>Example</b>	<pre>mrsvctl -c -h node01 -n John</pre> <p><i>Moab will create a reservation on node01 with the name John.</i></p>

**-o**

<b>Name</b>	OWNER
<b>Format</b>	<CREDTYPE>:<CREDID>
<b>Description</b>	The owner of a reservation. See <a href="#">Reservation Ownership</a> for more information.
<b>Example</b>	<pre>mrsvctl -c -h node01 -o USER:user1</pre> <p><i>Moab creates a reservation on node01 owned by user1.</i></p>

<b>-p</b>	
<b>Name</b>	PARTITION
<b>Format</b>	<STRING>
<b>Description</b>	Only allocate resources from the specified partition.
<b>Example</b>	<pre>mrsvctl -c -p switchB -t 14</pre> <p><i>Moab will allocate 14 tasks from the switchB partition.</i></p>

<b>-P</b>	
<b>Name</b>	PROFILE
<b>Format</b>	<STRING>
<b>Description</b>	Indicates the <a href="#">reservation profile</a> to load when creating this reservation.
<b>Example</b>	<pre>mrsvctl -c -P testing2 -t 14</pre> <p><i>Moab will allocate 14 tasks to a reservation defined by the testing2 reservation profile.</i></p>

<b>-q</b>	
<b>Name</b>	QUERY
<b>Format</b>	<RSV_ID> — The <code>-r</code> option accepts x: node regular expressions and r: node range expressions (asterisks (*) are supported wildcards as well).
<b>Description</b>	Get diagnostic information or list all completed reservations. The command gathers information from the Moab cache, which prevents it from interrupting the scheduler, but the <code>--blocking</code> option can be used to bypass the cache and interrupt the scheduler.
<b>Example</b>	<pre>mrsvctl -q ALL</pre>



-q	
	<i>Moab will query reservations.</i>
	<code>mrsvctl -q system.1</code>
	<i>Moab will query the reservation system.1.</i>

-r	
<b>Name</b>	RELEASE
<b>Format</b>	<RSV_ID> — The <code>-r</code> option accepts x: node regular expressions and r: node range expressions (asterisks (*) are supported wildcards as well).
<b>Description</b>	Releases the specified reservation. <div style="border: 1px solid #4a6984; padding: 10px; margin-top: 10px;"> <p><b>i</b> When you release an instance of a standing reservation, Moab will remember that and prevent a reservation from being created for that same period (even after a restart of Moab). When Moab reaches the end of the period, it will still create new reservations in the future to meet the reservation depth requirement.</p> </div>
<b>Example</b>	<div style="border: 1px dashed gray; padding: 5px; margin-bottom: 10px;"><code>&gt; mrsvctl -r system.1</code></div> <div style="border: 1px dashed gray; padding: 5px; margin-bottom: 10px;"><i>Moab will release reservation system.1.</i></div> <div style="border: 1px dashed gray; padding: 5px; margin-bottom: 10px;"><code>&gt; mrsvctl -r -g idle</code></div> <div style="border: 1px dashed gray; padding: 5px;"><i>Moab will release all idle job reservations.</i></div>

-R	
<b>Name</b>	RESOURCES
<b>Format</b>	<tid> or <RES>=<VAL>[{ + ;}<RES>=<VAL>]...

<b>-R</b>	
	Where <RES> is one of the following: PROCS, MEM, DISK, SWAP, GRES
<b>Default</b>	PROCS=-1
<b>Description</b>	<p>The resources to be reserved per task (-1 indicates all resources on node).</p> <div style="border: 1px solid #4a6984; padding: 5px; margin: 5px 0;"> <p><b>i</b> When specifying multiple resources, enclose the resource list in single quotes and separate the resource identifiers with semicolons (example: 'MEM=100;PROCS=1'). Alternatively, you can omit the single quotes and separate the resource identifiers with escaped semicolons (example: MEM=100\;PROCS=1).</p> </div> <div style="border: 1px solid #4a6984; padding: 5px; margin: 5px 0;"> <p><b>i</b> For GRES resources, &lt;VAL&gt; is specified in the format &lt;GRESNAME&gt; [ :&lt;COUNT&gt; ]</p> </div>
<b>Example</b>	<pre style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 10px;">&gt; mrsvctl -c -R 'MEM=100;PROCS=2' -t 2</pre> <p style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 10px;"><i>Moab will create a reservation for two tasks with the specified resources.</i></p> <pre style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 10px;">&gt; mrsvctl -c -R GRES:licenseA:100 -t 1</pre> <p style="border: 1px dashed #ccc; padding: 5px;"><i>Moab will reserve 100 instances of "licenseA".</i></p>

<b>-S</b>	
<b>Name</b>	STARTTIME
<b>Format</b>	[HH[:MM[:SS]]][_MO[/DD[/YY]]] or + [[DD:]HH:]MM:]SS
<b>Default</b>	[NOW]
<b>Description</b>	Absolute or relative time reservation will start. STARTTIME also supports an epoch timestamp.

-S	
<b>Example</b>	<pre style="border: 1px dashed gray; padding: 5px; margin-bottom: 10px;">&gt; mrsvctl -c -t ALL -s 3:00:00_4/4/22</pre> <p style="border: 1px dashed gray; padding: 5px; margin-bottom: 10px;"><i>Moab will create a reservation on all system resources at 3:00 am on April 4, 2022.</i></p> <pre style="border: 1px dashed gray; padding: 5px; margin-bottom: 10px;">&gt; mrsvctl -c -h node01 -s +5:00</pre> <p style="border: 1px dashed gray; padding: 5px; margin-bottom: 10px;"><i>Moab will create a reservation in 5 minutes on node01.</i></p> <pre style="border: 1px dashed gray; padding: 5px; margin-bottom: 10px;">&gt; mrsvctl -m -s -=5:00 system.1</pre> <p style="border: 1px dashed gray; padding: 5px;"><i>This will decrement the start time by 5 minutes.</i></p>

-S	
<b>Name</b>	SET ATTRIBUTE
<b>Format</b>	<p>&lt;ATTR&gt;=&lt;VALUE&gt; where &lt;ATTR&gt; is one of:</p> <ul style="list-style-type: none"> <li>• aaccount - accountable account</li> <li>• agroup - accountable group</li> <li>• aqos - accountable QoS</li> <li>• auser - accountable user</li> <li>• reqarch - required architecture</li> <li>• reqmemory - required node memory (in MB)</li> <li>• reqos - required operating system</li> <li>• rsvaccesslist - comma-delimited list of reservations or reservation groups that can be accessed by this reservation request. Because each reservation can access all other reservations by default, you should make any reservation with a specified rsvaccesslist exclusive by setting the <b>-E</b> flag. This setting gives the otherwise exclusive reservation access to reservations specified in the list.</li> </ul>
<b>Description</b>	A reservation attribute will be used to create this reservation.
<b>Example</b>	<pre style="border: 1px dashed gray; padding: 5px; margin-bottom: 10px;">&gt; mrsvctl -c -h node01 -S aqos=high</pre> <p style="border: 1px dashed gray; padding: 5px;"><i>Moab will create a reservation on node01 and will use the QOS high as the accountable credential.</i></p>

-t	
<b>Name</b>	TASKS
<b>Format</b>	<INTEGER>[-<INTEGER>]
<b>Description</b>	<p>The number of tasks to reserve. ALL indicates all resources available should be reserved.</p> <div style="border: 1px solid #4a6984; padding: 5px; margin-top: 10px;"> <p><b>i</b> If the task value is set to ALL, Moab applies the reservation regardless of existing reservations and exclusive issues. If an integer is used, Moab only allocates accessible resources. If a range is specified Moab attempts to reserve the maximum number of tasks, or at least the minimum.</p> </div>
<b>Example</b>	<pre style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 10px;">&gt; mrsvctl -c -t ALL</pre> <p style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 10px;"><i>Moab will create a reservation on all resources.</i></p> <pre style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 10px;">&gt; mrsvctl -c -t 3</pre> <p style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 10px;"><i>Moab will create a reservation for three tasks.</i></p> <pre style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 10px;">&gt; mrsvctl -c -t 3-10 -E</pre> <p style="border: 1px dashed #ccc; padding: 5px;"><i>Moab will attempt to reserve 10 tasks but will fail if it cannot get at least three.</i></p>

-T	
<b>Name</b>	TRIGGER
<b>Format</b>	<STRING>
<b>Description</b>	<p>Comma-delimited reservation trigger list following format described in the trigger format section of the reservation configuration overview. See <a href="#">Creating a Trigger</a> for more information.</p> <div style="border: 1px solid #4a6984; padding: 5px; margin-top: 10px;"> <p><b>i</b> To cancel a standing reservation with a trigger, the SRCFG parameter's attribute <a href="#">DEPTH</a> must be set to 0.</p> </div>
<b>Example</b>	<pre style="border: 1px dashed #ccc; padding: 5px;">&gt; mrsvctl -c -h node01 -T offset=200,etype=start,atype=exec,action=/tmp/email.sh</pre>

-T	
	<p><i>Moab will create a reservation on <code>node01</code> and fire the script <code>/tmp/email.sh</code> 200 seconds after it starts</i></p>
-V	
<b>Name</b>	VARIABLE
<b>Format</b>	<name>[=<value>][[:<name>[=<value>]]...]
<b>Description</b>	Semicolon-delimited list of variables that will be set when the reservation is created (see <a href="#">About Trigger Variables</a> for more information). Names with no values will simply be set to <code>TRUE</code> .
<b>Example</b>	<pre>&gt; mrsvctl -c -h node01 -V \$T1=mac;var2=18.19</pre> <div style="border: 1px dashed gray; padding: 10px; margin-top: 10px;"> <p><i>Moab will create a reservation on <code>node01</code> and set <code>\$T1</code> to <code>mac</code> and <code>var2</code> to <code>18.19</code>.</i></p> </div> <div style="border: 1px solid gray; padding: 5px; margin-top: 10px; display: flex; align-items: center;"> <span style="font-size: 1.2em; margin-right: 5px;">i</span> <span>For information on modifying a variable on a reservation, see <a href="#">MODIFY</a>.</span> </div>
-x	
<b>Name</b>	JOBLIST
<b>Format</b>	-x<jobs to be excluded>
<b>Description</b>	The <code>-x</code> flag, when used with <code>-F ignjobrsv</code> , lets users create reservations but exclude certain nodes that are running the listed jobs. The <code>-F</code> flag instructs <code>mrsvctl</code> to still consider nodes with current running jobs. The nodes are not listed directly.
<b>Example</b>	<pre>&gt; mrsvctl -c -t 5 -F ignjobrsv -x moab.5,moab.6</pre> <div style="border: 1px dashed gray; padding: 10px; margin-top: 10px;"> <p><i>Moab will create the reservation while assigning the nodes. Nodes running jobs <code>moab5</code> and <code>moab6</code> will not be assigned to the reservation.</i></p> </div>

### 3.7.24.E Parameters

RESERVATION PATTERN	
<b>Format</b>	<STRING>
<b>Description</b>	<p>A pattern specifying the reservation(s) to be affected by this action consisting of a space-delimited list of one or more of the following reservation expressions:</p> <ul style="list-style-type: none"> <li>• The name of a reservation.</li> <li>• The string "ALL", which matches all reservations.</li> <li>• A regular expression matching zero or more reservations. A reservation expression is treated as a regular expression if it has a prefix of "x:" or if it contains one of the characters in "[ ] * ^ \$ ,". Moab does a case-insensitive match using POSIX extended regular expressions and will match any part of the reservation name unless anchored with "^" or "\$".</li> </ul>
<b>Example</b>	<pre>'^system'</pre> <p><i>Specifies all reservations starting with the word "system".</i></p>

### 3.7.24.F Resource Allocation Details

When allocating resources, the following rules apply:

- When specifying tasks, each task defaults to one full compute node unless otherwise specified using the **-R** specification.
- When specifying tasks, the reservation will not be created unless all requested resources can be allocated (this behavior can be changed by specifying **-F** besteffort).
- When specifying tasks or hosts, only nodes in an idle or running state will be considered (this behavior can be changed by specifying **-F** ignstate).

### 3.7.24.G Reservation Timeframe Modification

Moab supports dynamically modifying the timeframe of existing reservations. This can be accomplished using the `mrsvctl -m` flag. By default, Moab will perform advanced boundary and resource access to verify that the modification does not result in an invalid scheduler state. However, in certain circumstances administrators may want to **FORCE** the modification in spite of any access violations. This can be done using the switch `mrsvctl -m --flags=force`, which forces Moab to bypass any access verification and force the change through.

### 3.7.24.H Extending a Reservation by Modifying the Endtime

The following increases the endtime of a reservation using the += tag:

```
$> showres
ReservationID      Type S      Start      End      Duration  N/P      StartTime
system.1          User -     11:35:57   1:11:35:57  1:00:00:00  1/2      Sat Nov 18
00:00:00
1 reservation located
$> mrsvctl -m endtime+=24:00:00 system.1
endtime for rsv 'system.1' changed
$> showres
ReservationID      Type S      Start      End      Duration  N/P      StartTime
system.1          User -     11:35:22   2:11:35:22  2:00:00:00  1/2      Sat Nov 18
00:00:00
1 reservation located
```

The following increases the endtime of a reservation by setting the endtime to an absolute time:

```
$> showres
ReservationID      Type S      Start      End      Duration  N/P      StartTime
system.1          User -     11:33:18   1:11:33:18  1:00:00:00  1/2      Sat Nov 18
00:00:00
1 reservation located
$> mrsvctl -m endtime=0_11/20 system.1
endtime for rsv 'system.1' changed
$> showres
ReservationID      Type S      Start      End      Duration  N/P      StartTime
system.1          User -     11:33:05   2:11:33:05  2:00:00:00  1/2      Sat Nov 18
00:00:00
1 reservation located
```

### 3.7.24.I Extending a Reservation by Modifying the Duration

The following increases the duration of a reservation using the += tag:

```
$> showres
ReservationID      Type S      Start      End      Duration  N/P      StartTime
system.1          User -     11:28:46   1:11:28:46  1:00:00:00  1/2      Sat Nov 18
00:00:00
1 reservation located
$> mrsvctl -m duration+=24:00:00 system.1
duration for rsv 'system.1' changed
>$ showres
ReservationID      Type S      Start      End      Duration  N/P      StartTime
system.1          User -     11:28:42   2:11:28:42  2:00:00:00  1/2      Sat Nov 18
00:00:00
1 reservation located
```

The following increases the duration of a reservation by setting the duration to an absolute time:

```
$> showres
ReservationID      Type S      Start      End      Duration  N/P      StartTime
system.1          User -     11:26:41   1:11:26:41  1:00:00:00  1/2      Sat Nov 18
```

```
00:00:00
1 reservation located
$> mrsvctl -m duration=48:00:00 system.1
duration for rsv 'system.1' changed
$> showres
ReservationID      Type S      Start      End      Duration  N/P      StartTime
system.1           User -     11:26:33  2:11:26:33  2:00:00:00  1/2      Sat Nov 18
00:00:00
1 reservation located
```

### 3.7.24.J Shortening a Reservation by Modifying the Endtime

The following modifies the endtime of a reservation using the `--` tag:

```
$> showres
ReservationID      Type S      Start      End      Duration  N/P      StartTime
system.1           User -     11:15:51  2:11:15:51  2:00:00:00  1/2      Sat Nov 18
00:00:00
1 reservation located
$> mrsvctl -m endtime--24:00:00 system.1
endtime for rsv 'system.1' changed
$> showres
ReservationID      Type S      Start      End      Duration  N/P      StartTime
system.1           User -     11:15:48  1:11:15:48  1:00:00:00  1/2      Sat Nov 18
00:00:00
1 reservation located
```

The following modifies the endtime of a reservation by setting the endtime to an absolute time:

```
$ showres
ReservationID      Type S      Start      End      Duration  N/P      StartTime
system.1           User -     11:14:00  2:11:14:00  2:00:00:00  1/2      Sat Nov 18
00:00:00
1 reservation located
$> mrsvctl -m endtime=0_11/19 system.1
endtime for rsv 'system.1' changed
$> showres
ReservationID      Type S      Start      End      Duration  N/P      StartTime
system.1           User -     11:13:48  1:11:13:48  1:00:00:00  1/2      Sat Nov 18
00:00:00
1 reservation located
```

### 3.7.24.K Shortening a Reservation by Modifying the Duration

The following modifies the duration of a reservation using the `--` tag:

```
$> showres
ReservationID      Type S      Start      End      Duration  N/P      StartTime
system.1           User -     11:12:20  2:11:12:20  2:00:00:00  1/2      Sat Nov 18
00:00:00
1 reservation located
$> mrsvctl -m duration--24:00:00 system.1
duration for rsv 'system.1' changed
$> showres
```



```

ReservationID      Type S      Start      End      Duration  N/P      StartTime
system.1          User -     11:12:07   1:11:12:07 1:00:00:00 1/2      Sat Nov 18
00:00:00
1 reservation located

```

The following modifies the duration of a reservation by setting the duration to an absolute time:

```

$> showres
ReservationID      Type S      Start      End      Duration  N/P      StartTime
system.1          User -     11:10:57   2:11:10:57 2:00:00:00 1/2      Sat Nov 18
00:00:00
1 reservation located
$> mrsvctl -m duration=24:00:00 system.1
duration for rsv 'system.1' changed
$> showres
ReservationID      Type S      Start      End      Duration  N/P      StartTime
system.1          User -     11:10:50   1:11:10:50 1:00:00:00 1/2      Sat Nov 18
00:00:00
1 reservation located

```

### 3.7.24.L Modifying the Starttime of a Reservation

The following increases the starttime of a reservation using the += tag:

```

$> showres
ReservationID      Type S      Start      End      Duration  N/P      StartTime
system.1          User -     11:08:30   2:11:08:30 2:00:00:00 1/2      Sat Nov 18
00:00:00
1 reservation located
$> mrsvctl -m starttime+=24:00:00 system.1
starttime for rsv 'system.1' changed
$> showres
ReservationID      Type S      Start      End      Duration  N/P      StartTime
system.1          User -     1:11:08:22 3:11:08:22 2:00:00:00 1/2      Sun Nov 19
00:00:00
1 reservation located

```

The following decreases the starttime of a reservation using the -= tag:

```

$> showres
ReservationID      Type S      Start      End      Duration  N/P      StartTime
system.1          User -     11:07:04   2:11:07:04 2:00:00:00 1/2      Sat Nov 18
00:00:00
1 reservation located
$> mrsvctl -m starttime-=24:00:00 system.1
starttime for rsv 'system.1' changed
$> showres
ReservationID      Type S      Start      End      Duration  N/P      StartTime
system.1          User -     -12:53:04  1:11:06:56 2:00:00:00 1/2      Fri Nov 17
00:00:00
1 reservation located

```

The following modifies the starttime of a reservation using an absolute time:

```

$> showres
ReservationID      Type S      Start      End      Duration  N/P      StartTime

```

```

system.1          User -    11:05:31  2:11:05:31  2:00:00:00    1/2    Sat Nov 18
00:00:00
1 reservation located
$> mrsvctl -m starttime=0_11/19 system.1
starttime for rsv 'system.1' changed
$> showres
ReservationID      Type S      Start      End      Duration      N/P      StartTime
system.1          User -    1:11:05:18  3:11:05:18  2:00:00:00    1/2      Sun Nov 19
00:00:00
1 reservation located

```

The following modifies the starttime of a reservation using an absolute time:

```

$> showres
ReservationID      Type S      Start      End      Duration      N/P      StartTime
system.1          User -    11:04:04  2:11:04:04  2:00:00:00    1/2      Sat Nov 18
00:00:00
1 reservation located
$> mrsvctl -m starttime=0_11/17 system.1
starttime for rsv 'system.1' changed
$> showres
ReservationID      Type S      Start      End      Duration      N/P      StartTime
system.1          User -    -12:56:02  1:11:03:58  2:00:00:00    1/2      Fri Nov 17
00:00:00
1 reservation located

```

## 3.7.25 Examples

- [Basic Reservation](#)
- [System Maintenance Reservation](#)
- [Explicit Task Description](#)
- [Dynamic Reservation Modification](#)
- [Reservation Modification](#)
- [Allocating Reserved Resources](#)
- [Modifying an Existing Reservation](#)

### **Example 3-22: Basic Reservation**

Reserve two nodes for use by users `john` and `mary` for a period of 8 hours starting in 24 hours.

```

> mrsvctl -c -a USER=john,USER=mary -starttime +24:00:00 -duration 8:00:00 -t 2
reservation 'system.1' created

```

### **Example 3-23: System Maintenance Reservation**

Schedule a system wide reservation to allow a system maintenance on Jun 20, 8:00 AM until Jun 22, 5:00 PM.

```
% mrsvctl -c -s 8:00:00_06/20 -e 17:00:00_06/22 -h ALL
reservation 'system.1' created
```

**Example 3-24: Explicit Task Description**

Reserve one processor and 512 MB of memory on nodes `node003` through `node006` for members of the group `staff` and jobs in the interactive class.

```
> mrsvctl -c -R PROCS=1,MEM=512 -a GROUP=staff,CLASS=interactive -h 'node00[3-6]'
reservation 'system.1' created
```

**Example 3-25: Dynamic Reservation Modification**

Modify reservation `john.1` to start in 2 hours, run for 2 hours, and include `node02` in the hostlist.

```
> mrsvctl -m starttime=+2:00:00,duration=2:00:00,HostExp+=node02
Note: hosts added to rsv system.3
```

**Example 3-26: Reservation Modification**

Remove user `John`'s access to reservation `system.1`

```
> mrsvctl -m -a USER=John system.1 --flags=unset
successfully changed ACL for rsv system.1
```

**Example 3-27: Allocating Reserved Resources**

Allocate resources for group `dev` that are **exclusive** except for resources found within reservations `myrinet.3` or `john.6`

```
> mrsvctl -c -E -a group=dev,rsv=myrinet.3,rsv=john.6 -h 'node00[3-6]'
reservation 'dev.14' created
```

Create exclusive network reservation on racks 3 and 4.

```
> mrsvctl -c -E -a group=ops -g network -f rack3 -h ALL
reservation 'ops.1' created
> mrsvctl -c -E -a group=ops -g network -f rack4 -h ALL
reservation 'ops.2' created
```

Allocate 64 nodes for 2 hours to new reservation and grant access to reservation `system.3` and all reservations in the reservation group `network`.

```
> mrsvctl -c -E -d 2:00:00 -a group=dev -t 64 -S rsvaccesslist=system.3,network
reservation 'system.23' created
```

Allocate 4 nodes for 1 hour to new reservation and grant access to idle job reservations.

```
> mrsvctl -c -E -d 1:00:00 -t 4 -S rsvaccesslist=idle
reservation 'system.24' created
```

**Example 3-28: Modifying an Existing Reservation**

Remove user `john` from reservation ACL.

```
> mrsvctl -m -a USER=john system.1 --flags=unset
successfully changed ACL for rsv system.1
```

### Change reservation group.

```
> mrsvctl -m RSVGROU=network ops.4
successfully changed RSVGROU for rsv ops.4
```

---

## Related Topics

- (Optional) Install Moab Client - in the *Moab HPC Suite Installation and Configuration Guide* explains how to distribute this command to client nodes
- [Admin Reservation Overview](#)
- [showres](#)
- [mdiag -r](#)
- [mshow -a](#) command to identify available resources
- [job to rsv binding](#)

## 3.7.26 mschedctl

### 3.7.26.A Synopsis

```
mschedctl -A '<MESSAGE>'
```

```
mschedctl -c message messagestring [-o type:val]
```

```
mschedctl -c trigger triggerid -o type:val
```

```
mschedctl -d trigger:triggerid
```

```
mschedctl -d message:index:wq
```

```
mschedctl -f {all|fairshare|usage}
```

```
mschedctl -k
```

```
mschedctl -l {config|feature|gmetric|gres|message|opsys|trigger|trans} [-v] [--xml]
```

```
mschedctl -L [<LOGLEVEL>[:<LOG_FILE>]]
```

```
mschedctl -m config string [-e]
```

```
mschedctl -m trigger triggerid attr=val[,attr=val...]
```

```
mschedctl -q mschedctl -q pactions --xml
```

`mschedctl -p``mschedctl -r [resumetime]``mschedctl -R``mschedctl -s [STOPIITERATION]``mschedctl -S [STEPITERATION]``mschedctl -W`

```
[--about] [--help] [--host=<serverHostName>] [--
loglevel=<logLevel>] [--msg=<message>] [--port=<serverPort>]
[--timeout=<seconds>] [--version] [--xml]
```

### 3.7.26.B Overview

The `mschedctl` command controls various aspects of scheduling behavior. It is used to manage scheduling activity, shutdown the scheduler, and create resource trace files. It can also evaluate, modify, and create parameters, triggers, and messages.

**i** With many flags, the `--msg=<MSG>` option can be specified to annotate the action in the [event log](#).

### 3.7.26.C Options

#### -A - ANNOTATE

<b>Format</b>	<STRING>
<b>Description</b>	Report the specified parameter modification to the event log and annotate it with the specified message. The <code>RECORDEVENTLIST</code> parameter must be set in order for this to work.
<b>Example</b>	<pre>mschedctl -A 'increase logging' -m 'LOGLEVEL 6'</pre> <p><i>Adjust the LOGLEVEL parameter and record an associated message.</i></p>

#### -c - CREATE

<b>Format</b>	One of:
---------------	---------

-c - CREATE	
	<ul style="list-style-type: none"> <li>• <b>message</b> &lt;STRING&gt; [-o &lt;TYPE&gt;:&lt;VAL&gt;]</li> <li>• <b>trigger</b>&lt;TRIGSPEC&gt; -o &lt;OBJECTTYPE&gt;:&lt;OBJECTID&gt;</li> <li>• <b>gevent</b> -n &lt;NAME&gt; [-m &lt;message&gt;]</li> </ul> <p>Where &lt;ATTR&gt; is one of account, duration, ID, messages, profile, reqresources, resources, rsvprofile, starttime, user, or variables.</p>
<b>Description</b>	Create a message, trigger, or gevent and attach it to the specified object. To create a trigger on a default object, use the Moab configuration file ( <code>moab.cfg</code> ) rather than the <code>mschedctl</code> command.
<b>Example</b>	<pre>mschedctl -c message tell the admin to be nice</pre> <p><i>Create a message on the system table.</i></p> <pre>mschedctl -c trigger EType=start,AType=exec,Action="/tmp/email \$OWNER \$TIME" -o rsv:system.1</pre> <p><i>Create a trigger linked to system.1.</i></p> <div style="border: 1px solid #004a7c; padding: 5px; margin: 10px 0;"> <p><b>i</b> Creating triggers on default objects via <code>mschedctl -c trigger</code> does not propagate the triggers to individual objects. To propagate triggers to all objects, the triggers must be created within the <code>moab.cfg</code> file; for example: <code>NODECFG [DEFAULT] TRIGGER.</code></p> </div> <pre>mschedctl -c gevent -n diskfailure -m "node=n4"</pre> <p><i>Create a gevent indicating a disk failure on the node labeled n4.</i></p>

-d - DESTROY	
<b>Format</b>	One of: <ul style="list-style-type: none"> <li>• <b>trigger</b>:&lt;TRIGID&gt;</li> <li>• <b>message</b>:&lt;INDEX&gt;</li> </ul>
<b>Description</b>	Delete a trigger or message.
<b>Example</b>	<pre>mschedctl -d trigger:3</pre>

**-d - DESTROY***Delete trigger 3.*`mschedctl -d message:5`*Delete message  
with index 5.***-f - FLUSH****Format** {all|fairshare|usage}**Description** Reset all internally-stored Moab Scheduler statistics to the initial start-up state as of the time the command was executed.

**i** Flushing should only be used if you experience corrupt statistics. The best practice is to pause the Moab scheduler with `mschedctl -p` before running the flush command. After running the flush command, unpause the Moab scheduler with `mschedctl -r` and the jobs will start flowing again. For all external observers this will be a transparent flush unless they are watching the stats.

**Example**`mschedctl -f usage`*Flush usage statistics.***-k - KILL****Description** Stop scheduling and exit the scheduler.**Example**`mschedctl -k`*Kill the scheduler.***-l - LIST****Format** {config|feature|gmetric|gres|message|opsys|trans|trigger} [-v] [--xml]

-l - LIST	
	<p><b>i</b> Using the <code>--xml</code> argument with the <code>trans</code> option returns XML that states if the queried TID is valid or not.</p>
<b>Default</b>	config
<b>Description</b>	<p>List the generic metrics, generic resources, scheduler configuration, system messages, operating systems, triggers, transactions, or node features recognized by Moab.</p> <p><b>i</b> This command does not show credential parameters (such as user, group class, QoS, account).</p>
<b>Example</b>	<pre>mschedctl -l config</pre> <p><i>List system parameters.</i></p> <p><b>i</b> The <code>config</code> command without the <code>-v</code> flag does not show the settings of all scheduling parameters. To show the settings of all scheduling parameters, use the <code>-v</code> flag. This will provide an extended output. This output is often best used in conjunction with the <code>grep</code> command as the output can be voluminous.</p> <pre>mschedctl -l feature</pre> <p><i>List all node features recognized by Moab.</i></p> <pre>mschedctl -l gmetric</pre> <p><i>List all configured generic metrics.</i></p> <pre>mschedctl -l gres</pre> <p><i>List all configured generic resources. Use the -v flag to display generic resource traits (such as license or numa).</i></p> <pre>mschedctl -l message</pre> <p><i>List all system messages.</i></p> <pre>mschedctl -l opsys</pre> <p><i>List all recognized operating systems.</i></p> <pre>mschedctl -l trans 1</pre> <p><i>List transaction id 1.</i></p>



**-l - LIST**

```
mschedctl -l trigger
```

*List triggers.*

**-L - LOG**

**Format** [`<LOGLEVEL>[: <LOG_FILE>]`]

**Default** 7 \$MOABHOMEDIR/log/moab.log

**Description** Create a temporary log file with the specified loglevel. If no log file is given, Moab creates a log file in the log directory whose filename extension is the timestamp of when the command was run (for example, "/opt/moab/log/moab.log.20220405081227").

**Example**

```
mschedctl -L7:/tmp/moab.log
```

**-m - MODIFY****Format**

One of:

- **config** [`<STRING>`]  
`[-e]`  
`<STRING>` is any string that would be acceptable in `moab.cfg`
  - If no string is specified, `<STRING>` is read from STDIN.
  - If `-e` is specified, the configuration string will be evaluated for correctness but no configuration changes will take place. Any issues with the provided string will be reported to STDERR.

**i** Use of `mschedctl --flags=persistent -m <config>` has been deprecated; use the following method instead:

1. Run `mschedctl -m <config>` to put the change into effect dynamically.
2. Manually add the settings to the `moab.cfg` file, so that it always goes into effect after any future Moab restarts/recycles.

<b>-m - MODIFY</b>	
	<div style="border: 1px solid #000; padding: 10px; margin-bottom: 10px;"> <p><b>i</b> Dynamically modifying classes is not recommended. Moab should be restarted whenever classes are modified. This is especially true given the fact that sometimes the classes/queues/partitions are under control of a resource manager. For dynamic operations, use node sets/features or reservations.</p> </div> <ul style="list-style-type: none"> <li>• <b>trigger:</b>&lt;TRIGID&gt; &lt;ATTR&gt;=&lt;VAL&gt; Where &lt;ATTR&gt; is one of action, atype, etype, iscomplete, oid, otype, offset, or threshold.</li> </ul>
<b>Description</b>	<p>Modify a system parameter or trigger.</p> <div style="border: 1px solid #000; padding: 10px; margin-top: 10px;"> <p><b>i</b> Moab only loads the following list of parameters when first starting up. Therefore, to change any of these, you must edit the setting in moab.cfg and then restart/recycle with mschedctl -R.</p> <ul style="list-style-type: none"> <li>• JOBMAXNODECOUNT</li> <li>• MAXGMETRIC</li> <li>• MAXGRES</li> <li>• MAXJOB</li> <li>• MAXNODE</li> <li>• MAXRSVPERNODE</li> <li>• STATPROC*</li> <li>• STATTIME*</li> </ul> </div>
<b>Example</b>	<pre style="border: 1px dashed #000; padding: 5px; margin-bottom: 10px;">mschedctl -m config LOGLEVEL 9</pre> <p style="border: 1px dashed #000; padding: 5px; margin-bottom: 10px;"><i>Change the system loglevel to 9.</i></p> <pre style="border: 1px dashed #000; padding: 5px; margin-bottom: 10px;">mschedctl -m trigger:2 AType=exec,Offset=200,OID=system.1</pre> <p style="border: 1px dashed #000; padding: 5px;"><i>Change aspects of trigger 2.</i></p>
<b>-p - PAUSE</b>	
<b>Description</b>	Disable scheduling but allow the scheduler to update its cluster and workload state information.
<b>Example</b>	<pre style="border: 1px dashed #000; padding: 5px;">mschedctl -p</pre>

-q QUERY PENDING ACTIONS	
<b>Default</b>	<code>mschedctl -q pactions --xml</code>
<b>Description</b>	A way to view pending actions. Only an XML request is valid. Pending actions can be system jobs.
<b>Example</b>	<pre>mschedctl -q pactions --xml</pre>
-r - RESUME	
<b>Format</b>	<code>mschedctl -r [[HH:[MM:]]SS]</code>
<b>Default</b>	0
<b>Description</b>	Resume scheduling in the specified amount of time (or immediately if none is specified).
<b>Example</b>	<pre>mschedctl -r</pre> <i>Resume scheduling immediately.</i>
-R - RECYCLE	
<b>Description</b>	Recycle scheduler immediately (shut it down and restart it using the original execution environment and command line arguments). <div style="border: 1px solid #4a7c9c; padding: 5px; margin-top: 10px;"><b>i</b> If Moab has been started under systemd, use <code>systemctl restart moab.service</code> instead of using this option.</div>
<b>Example</b>	<pre>mschedctl -R</pre> <i>Recycle scheduler immediately.</i> <div style="border: 1px solid #4a7c9c; padding: 5px; margin-top: 10px;"><b>i</b> To restart Moab with its last known scheduler state, use: <code>mschedctl -R savestate</code></div>

<b>-s - STOP</b>	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	Suspend/stop scheduling at specified iteration (or at the end of the current iteration if none is specified). If the letter I follows <ITERATION>, Moab will not process client requests until this iteration is reached.
<b>Example</b>	<pre>mschedctl -s 100I</pre> <p><i>Stop scheduling at iteration 100 and ignore all client requests until then.</i></p>

<b>-S - STEP</b>	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	Step the specified number of iterations (or to the next iteration if none is specified) and suspend scheduling. If the letter I follows <ITERATION>, Moab will not process client requests until this iteration is reached.
<b>Example</b>	<pre>mschedctl -S</pre> <p><i>Step to the next iteration and stop scheduling.</i></p>

<b>-W</b>	
<b>Description</b>	Preform a manual checkpoint file write.
<b>Example</b>	<pre>mschedctl -W</pre>

### 3.7.26.D Examples

*Example 3-29: Shutting down the Scheduler*

```
mschedctl -k
scheduler will be shutdown immediately
```

#### Related Topics

- (Optional) Install Moab Client - in the *Moab HPC Suite Installation and Configuration Guide* explains how to distribute this command to client nodes

## 3.7.27 mshow

### 3.7.27.A Synopsis

```
mshow [-a] [-q jobqueue=active] [--about] [--help] [--
host=<serverHostName>] [--loglevel=<logLevel>] [--
msg=<message>] [--port=<serverPort>] [--timeout=<seconds>] [--
version] [--xml]
```

### 3.7.27.B Overview

The *mshow* command displays various diagnostic messages about the system and job queues.

### 3.7.27.C Arguments

Flag	Description
<b>-a</b>	<b>AVAILABLE RESOURCES</b>
<b>-q [&lt;QUEUENAME&gt;]</b>	Displays the job queues.

## 3.7.27.D Format

## AVAILABLE RESOURCES

<b>Format</b>	Can be combined with <code>--flags=[tid verbose future]</code> <code>--format=xml</code> and/or <code>-w</code>
<b>Description</b>	Display available resources.
<b>Example</b>	<pre>&gt; mshow -a -w user=john --flags=tid --format=xml</pre> <p><i>Show resources available to john in XML format with a transaction id. See <a href="#">mshow -a</a> for details.</i></p>

## JOB QUEUE

<b>Format</b>	<QUEUENAME>, where the queue name is one of: active, eligible, or blocked. Job queue names can be delimited by a comma to display multiple queues. If no job queue name is specified, <code>mshow</code> displays all job queues.
<b>Description</b>	Displays the job queues. If a job queue name is specified, <code>mshow</code> shows only that job queue.
<b>Example</b>	<pre>&gt; mshow -q active,blocked [Displays all jobs in the active and blocked queues] ...</pre>

## Related Topics

- (Optional) Install Moab Client - in the *Moab HPC Suite Installation and Configuration Guide* explains how to distribute this command to client nodes
- [mshow -a](#) command to show available resources

## 3.7.28 mshow -a

### 3.7.28.A Synopsis

```
mshow -a [-i] [-o] [-T] [-w where] [-x] [--about] [--help] [--
host=<serverHostName>] [--loglevel=<logLevel>] [--
msg=<message>] [--port=<serverPort>] [--timeout=<seconds>] [--
version] [--xml]
```

### 3.7.28.B Overview

The `mshow -a` command allows for querying of available system resources.

### 3.7.28.C Arguments

<b>[-i]</b>	INTERSECTION
<b>[-o]</b>	NO AGGREGATE
<b>[-T]</b>	TIMELOCK
<b>[-w]</b>	WHERE
<b>[-x]</b>	EXCLUSIVE

*Table 3-1: Argument Format*

--flags	
<b>Name</b>	Flags
<b>Format</b>	--flags=[ future   policy   tid   summary   verbose ]
<b>Description</b>	<p><b>future</b> will return resources available immediately and available in the future.</p> <p><b>policy (Deprecated. May be removed in a future release.)</b> will apply charging policies to determine the total cost of each reported solution (only enabled for XML responses).</p> <p><b>summary</b> will assign all jointly allocated transactions as dependencies of the</p>

--flags	
	<p>first transaction reported.</p> <p><b>tid</b> will associate a transaction ID with the reported results.</p> <p><b>verbose</b> will return diagnostic information.</p>
<b>Example</b>	<pre>&gt; mshow -a -w user=john --flags=tid --xml</pre> <p><i>Show resources available to john in XML format with a transaction ID.</i></p>

--xml	
<b>Name</b>	XML
<b>Format</b>	--xml
<b>Description</b>	Report results in XML format.
<b>Example</b>	<pre>&gt; mshow -a -w user=john --flags=tid --xml</pre> <p><i>Show resources available to john in XML format with a transaction ID.</i></p>

-i	
<b>Name</b>	INTERSECTION
<b>Description</b>	Specifies that an intersection should be performed during an <code>mshow -a</code> command with multiple requirements.

-o	
<b>Name</b>	NO AGGREGATE
<b>Description</b>	Specifies that the results of the command <code>mshow -a</code> with multiple requirements should not be aggregated together.



-T	
<b>Name</b>	TIMELOCK
<b>Description</b>	Specifies that the multiple requirements of an <code>mshow -a</code> command should be timelocked.
<b>Example</b>	<pre>&gt; mshow -a -w minprocs=1,os=linux,duration=1:00:00 \ -w minprocs=1,os=aix,duration=10:00 \ --flags=tid,future -x -T</pre>

-w	
<b>Name</b>	WHERE
<b>Format</b>	<p>Comma-delimited list of <code>&lt;ATTR&gt;=&lt;VAL&gt;</code> pairs:  <code>&lt;ATTR&gt;=&lt;VAL&gt; [,&lt;ATTR&gt;=&lt;VAL&gt;]...</code></p> <div style="border: 1px solid #4a6984; padding: 5px; margin: 5px 0;"> <p><b>i</b> If any of the <code>&lt;ATTR&gt;=&lt;VAL&gt;</code> pairs contains a sub-list that is also comma-delimited, the entire <code>-w</code> string must be wrapped in single quotations with the sub-list expression wrapped in double quotations. See the example below.</p> </div> <p>Attributes are listed below in <a href="#">table 2</a>.</p>
<b>Description</b>	Add a Where clause to the current command (currently supports up to six co-allocation clauses).
<b>Example</b>	<pre>&gt; mshow -a -w minprocs=2,duration=1:00:00 -w nodemem=512,duration=1:00:00</pre> <p><i>Moab returns a list of all nodes with at least 2 processors and one hour duration or with a memory of 512 and a duration of one hour.</i></p> <pre>&gt; mshow -a -w nodefeature=!vmware:gpfs --flags=future</pre> <p><i>Moab returns a list of all nodes that do not contain the <code>vmware</code> feature but that do contain the <code>gpfs</code> feature.</i></p> <pre>&gt; mshow -a -w 'duration=INFINITY,"excludehostlist=n01,n12,n23"'</pre> <p><i>Moab returns a list of all nodes with a duration of <code>INFINITY</code>, except for nodes named <code>n01</code>, <code>n12</code>, and <code>n23</code>. Note the use of single quotations containing the entire <code>-w</code> string and the use of double quotations containing the <code>excludehostlist</code> attribute.</i></p>

-x	
<b>Name</b>	EXCLUSIVE
<b>Description</b>	Specifies that the multiple requirements of an <code>mshow -a</code> command should be exclusive (i.e., each node can only be allocated to a single requirement).
<b>Example</b>	<pre>&gt; mshow -a -w minprocs=1,os=linux -w minprocs=1,os=aix --flags=tid -x</pre>

Table 3-2: Request Attributes

Name	Description
<b>account</b>	The account credential of the requestor.
<b>acl</b>	ACL to attach to the reservation.  This ACL must be enclosed in quotation marks. For example: <code>\$ mshow -a ... -w acl=\"user=john\" ...</code>
<b>arch</b>	Select only nodes with the specified architecture.
<b>class</b>	The class credential of the requestor.
<b>coalloc</b>	The co-allocation group of the specific Where request (can be any string but must match co-allocation group of at least one other Where request).  <div style="border: 1px solid #4a6984; padding: 5px; background-color: #e6f2ff;"><b>i</b> The number of tasks requested in each Where request must be equal whether this taskcount is specified via <code>minprocs</code>, <code>mintasks</code>, or <code>gres</code>.</div>
<b>count</b>	The number of profiles to apply to the resource request.
<b>displaymode</b>	Possible value is <code>future</code> . (Example: <code>displaymode=future</code> ). Constrains how results are presented; setting <code>future</code> evaluates which resources are available now and which resources will be available in the future that match the requested attributes.
<b>duration</b>	The duration for which the resources will be required in format <code>[ [DD:]HH:]MM:]SS</code>
<b>excludehostlist</b>	Do not select any nodes from the given list. The list must be comma-delimited.

Name	Description
	<pre>&gt; mshow -a -w 'duration=INFINITY,"excludehostlist=n01,n12,n23"'</pre> <p><i>Moab returns a list of all nodes with a duration of INFINITY, except for nodes named n01, n12, and n23. Note the use of single quotations to contain the entire -w string, and the use of double quotations containing the excludehostlist attribute.</i></p>
<b>gres</b>	Select only nodes that possess the specified generic resource.
<b>group</b>	The group credential of the requestor.
<b>hostlist</b>	Select only the specified resources. The list must be comma-delimited. <pre>&gt; mshow -a -w 'duration=INFINITY,"hostlist=n01,n12,n23"'</pre> <p><i>Moab returns a list of nodes from the selected hostlist that have a duration of INFINITY. Note the use of single quotations to contain the entire -w string, and the use of double quotations containing the hostlist attribute.</i></p>
<b>job</b>	Use the resource, duration, and credential information for the job specified as a resource request template.
<b>jobfeature</b>	Select only resources that would allow access to jobs with the specified job features.
<b>jobflags</b>	Select only resources that would allow access to jobs with the specified job flags. The <code>jobflags</code> attribute accepts a colon delimited list of multiple flags.
<b>minnodes</b>	Return only results with at least the number of nodes specified. If used with TIDs, return only solutions with exactly minnodes nodes available.
<b>minprocs</b>	Return only results with at least the number of processors specified. If used with TIDs, return only solutions with exactly minprocs processors available.
<b>mintasks</b>	<b>FORMAT:</b> <TASKCOUNT>[@<RESTYPE>:<COUNT> [+<RESTYPE>:<COUNT>]...] where <RESTYPE> is one of <code>procs</code> , <code>mem</code> , <code>disk</code> , or <code>swap</code> . Return only results with at least the number of tasks specified. If used with TIDs, return only solutions with exactly mintasks

Name	Description
	available.
<b>nodedisk</b>	Select only nodes with at least <code>nodedisk</code> MB of local disk configured.
<b>nodefeature</b>	Select only nodes with all specified features present and nodes without all \! specified features using format <code>[\!]&lt;feature&gt;[: [\!]&lt;feature&gt;] . . .</code> You must set the <b>future</b> flag when specifying node features.
<b>nodemem</b>	Select only nodes with at least <code>nodemem</code> MB of memory configured.
<b>offset</b>	Select only resources that can be co-allocated with the specified time offset where offset is specified in the format <code>[ [ DD: ] HH: ] MM: ] SS</code>
<b>os</b>	Select only nodes with have, or can be provisioned to have, the specified operating system.
<b>partition</b>	The partition where the resources must be located.
<b>policylevel</b>	Enable policy enforcement at the specified policy constraint level.
<b>qos</b>	The qos credential of the requestor.
<b>rsvprofile</b>	Use the specified profile if committing a resulting transaction ID directly to a reservation.
<b>starttime</b>	<p>Constrain the timeframe for the returned results by specifying one or more ranges using the format <code>&lt;STIME&gt;[-&lt;ENDTIME&gt;];&lt;STIME&gt;[-&lt;ENDTIME&gt;]]</code> where each time is specified in the format in absolute, relative, or epoch time format <code>( [ HH [ : MM [ : SS ] ] ] [ _MO [ / DD [ / YY ] ] ] ] or + [ [ DD: ] HH: ] MM: ] SS or &lt;EPOCHTIME&gt;)</code>.</p> <div style="border: 1px solid black; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p><b>i</b> The <code>starttime</code> specified is not the exact time at which the returned range must start, but is rather the earliest possible time the range can start.</p> </div>
<b>taskmem</b>	Require <code>taskmem</code> MB of memory per task located.
<b>tpn</b>	Require exactly <code>tpn</code> tasks per node on all discovered resources.

Name	Description
<b>user</b>	The user credential of the requestor.
<b>var</b>	Use associated variables in generating per transaction charging quotes.
<b>variables</b>	Takes a string of the format <code>variables='var [=attr] ' [; 'var [=attr] '</code> and passes the variables onto the reservation when used in conjunction with <code>--flags=tid</code> and <code>mrsvctl -c -R &lt;tid&gt;</code> .

### 3.7.28.D Usage Notes

The `mshow -a` command allows for querying of available system resources. When combined with the `--flags=tid` option these available resources can then be placed into a packaged reservation (using `mrsvctl -c -R`). This allows system administrators to grab and reserve available resources for whatever reason, without conflicting with jobs or reservations that may be holding certain resources.

There are a few restrictions on which `<ATTR>` from the `-w` command can be placed in the same `req:minprocs,minnodes`, and `gres` are all mutually exclusive, only one can be used per `-w` request.

The allocation of available nodes will follow the global [NODEALLOCATIONPOLICY](#).

When the `'-o'` flag is not used, multi-request results will be aggregated. This aggregation will negate the use of offsets and request-specific starttimes.

The config parameter [RESOURCEQUERYDEPTH](#) controls the maximum number of options that will be returned in response to a resource query.

### 3.7.28.E Examples

#### Example 3-30: Basic Compute Node Query and Reservation

```
> mshow -a -w duration=10:00:00,minprocs=1,os=AIX53,jobfeature=shared --
flags=tid,future
```

Partition	Tasks	Nodes	Duration	StartOffset	StartDate	TID	ReqID
ALL	1	1	10:00:00	00:00:00	13:28:09_04/27	TID=4	ReqID=0
ALL	1	1	10:00:00	10:00:00	17:14:48_04/28	TID=5	ReqID=0
ALL	1	1	10:00:00	20:00:00	21:01:27_04/29	TID=6	ReqID=0

```
> mrsvctl -c -R 4
Note: reservation system.2 created
```

**Example 3-31: Mixed Processor and License Query**

Select one node with 4 processors and 1 matlab license where the matlab license is only available for the last hour of the reservation. Also, select 16 additional processors that are available during the same timeframe but which can be located anywhere in the cluster. Group the resulting transactions together using transaction dependencies so only the first transaction needs to be committed to reserve all associated resources.

```
> mshow -a -i -o -x -w mintasks=1@PROCS:4,duration=10:00:00,coalloc=a \
-w gres=matlab,offset=9:00:00,duration=1:00:00,coalloc=a \
-w minprocs=16,duration=10:00:00 --flags=tid,future,summary
Partition      Tasks  Nodes  Duration  StartOffset  StartDate
-----
ALL            1      1    10:00:00    00:00:00    13:28:09_04/27  TID=4  ReqID=0
ALL            1      1    10:00:00    10:00:00    17:14:48_04/28  TID=5  ReqID=0
ALL            1      1    10:00:00    20:00:00    21:01:27_04/29  TID=6  ReqID=0
> mrsvctl -c -R 4

Note: reservation system.2 created
Note: reservation system.3 created
Note: reservation system.4 created
```

**Example 3-32: Request for Generic Resources**

Query for a generic resource on a specific host (no processors, only a generic resource).

```
> mshow -a -i -x -o -w gres=dvd,duration=10:00,hostlist=node03 --flags=tid,future
Partition      Tasks  Nodes  StartOffset  Duration  StartDate
-----
ALL            1      1    00:00:00    00:10:00    11:33:25_07/27  TID=16
ReqID=0
ALL            1      1    00:10:00    00:10:00    11:43:25_07/27  TID=17
ReqID=0
ALL            1      1    00:20:00    00:10:00    11:53:25_07/27  TID=18
ReqID=0
> mrsvctl -c -R 16
Note: reservation system.6 created
> mdiag -r system.6
Diagnosing Reservations
RsvID          Type  Par  StartTime  EndTime  Duration  Node  Task
-----
-
system.6      User  loc  -00:01:02  00:08:35  00:09:37  1    1
0
Flags: ISCLOSED
ACL:   RSV==system.6=
CL:   RSV==system.6
Accounting Creds: User:test
Task Resources: dvd: 1
Attributes (HostExp='^node03$')
Rsv-Group: system.6
```

**Example 3-33: Allocation of Shared Resources**

This example walks through a relatively complicated example where a set of resources can be reserved to be allocated for shared requests. In the example below, the first *mshow* query looks for resources within an existing shared reservation. In the example, this first

query fails because there is now existing reservation. The second query looks for resources within an existing shared reservation. In the example, this first query fails because there is now existing reservation. The second *mshow* request asks for resources outside of a shared reservation and finds the desired resources. These resources are then reserved as a shared pool. The third *mshow* request again asks for resources inside of a shared reservation and this time finds the desired resources.

```
> mshow -a -w duration=10:00:00,minprocs=1,os=AIX53,jobflags=ADVRES,jobfeature=shared
--flags=tid

Partition      Tasks  Nodes      Duration      StartOffset      StartDate
-----
> mshow -a -w duration=100:00:00,minprocs=1,os=AIX53,jobfeature=shared --flags=tid
Partition      Tasks  Nodes      Duration      StartOffset      StartDate
-----
ALL              1      1      100:00:00      00:00:00      13:20:23_04/27  TID=1  ReqID=0
> mrsvctl -c -R 1
Note: reservation system.1 created
> mshow -a -w duration=10:00:00,minprocs=1,os=AIX53,jobflags=ADVRES,jobfeature=shared
--flags=tid
Partition      Tasks  Nodes      Duration      StartOffset      StartDate
-----
ALL              1      1      10:00:00      00:00:00      13:20:36_04/27  TID=2  ReqID=0
> mrsvctl -c -R 2
Note: reservation system.2 created
```

### Example 3-34: Full Resource Query in XML Format

The following command will report information on all available resources that meet at least the minimum specified processor and walltime constraints and which are available to the specified user. The results will be reported in XML to allow for easy system processing.

```
> mshow -a -w class=grid,minprocs=8,duration=20:00 --format=xml --flags=future,verbose

<Data>
  <Object>cluster</Object>
  <job User="john" time="1162407604"></job>
  <par Name="template">
    <range duration="Duration" nodecount="Nodes" proccount="Procs"
starttime="StartTime"></range>
  </par>
  <par Name="ALL" feasibleNodeCount="131" feasibleTaskCount="163">
    <range duration="1200" hostlist="opt-001:1,opt-024:1,opt-025:1,opt-027:2,opt-
041:1,opt-042:1,x86-001:1,P690-001:1,P690-021:1,P690-022:1"
      index="0" nodecount="10" proccount="8" reqid="0"
starttime="1162407604"></range>
    <range duration="1200" hostlist="opt-001:1,opt-024:1,opt-025:1,opt-027:2,opt-
039:1,opt-041:1,opt-042:1,x86-001:1,P690-001:1,P690-021:1,P690-022:1"
      index="0" nodecount="11" proccount="8" reqid="0"
starttime="1162411204"></range>
    <range duration="1200" hostlist="opt-001:1,opt-024:1,opt-025:1,opt-027:2,opt-
039:1,opt-041:1,opt-042:1,x86-001:1,x86-002:1,x86-004:1,
x86-006:1,x86-013:1,x86-014:1,x86-015:1,x86-016:1,x86-037:1,P690-001:1,P690-
021:1,P690-022:1"
      index="0" nodecount="19" proccount="8" reqid="0"
starttime="1162425519"></range>
  </par>
  <par Name="SharedMem">
```

```

    <range duration="1200" hostlist="P690-001:1,P690-002:1,P690-003:1,P690-004:1,P690-
005:1,P690-006:1,P690-007:1,P690-008:1,P690-009:1,
    P690-010:1,P690-011:1,P690-012:1,P690-013:1,P690-014:1,P690-015:1,P690-
016:1,P690-017:1,P690-018:1,P690-019:1,P690-020:1,P690-021:1,
    P690-022:1,P690-023:1,P690-024:1,P690-025:1,P690-026:1,P690-027:1,P690-
028:1,P690-029:1,P690-030:1,P690-031:1,P690-032:1"
    index="0" nodecount="32" proccount="8" reqid="0"
starttime="1163122507"></range>
</par>
<par Name="64Bit">
    <range duration="1200" hostlist="opt-001:1,opt-024:1,opt-025:1,opt-027:2,opt-
039:1,opt-041:1,opt-042:1"
    index="0" nodecount="7" proccount="8" reqid="0"
starttime="1162411204"></range>
    <range duration="1200" hostlist="opt-001:1,opt-024:1,opt-025:1,opt-027:2,opt-
039:1,opt-041:1,opt-042:1,opt-043:1,opt-044:1,opt-045:1,
    opt-046:1,opt-047:1,opt-048:1,opt-049:1,opt-050:1"
    index="0" nodecount="15" proccount="8" reqid="0"
starttime="1162428996"></range>
    <range duration="1200" hostlist="opt-001:1,opt-006:1,opt-007:2,opt-008:2,opt-
009:2,opt-010:2,opt-011:2,opt-012:2,opt-013:2,opt-014:2,
    opt-015:2,opt-016:2,opt-017:2,opt-018:2,opt-019:2,opt-020:2,opt-021:2,opt-
022:2,opt-023:2,opt-024:2,opt-025:1,opt-027:2,opt-039:1,
    opt-041:1,opt-042:1,opt-043:1,opt-044:1,opt-045:1,opt-046:1,opt-047:1,opt-
048:1,opt-049:1,opt-050:1"
    index="0" nodecount="33" proccount="8" reqid="0"
starttime="1162876617"></range>
</par>
<par Name="32Bit">
    <range duration="1200" hostlist="x86-001:1,x86-002:1,x86-004:1,x86-006:1,x86-
013:1,x86-014:1,x86-015:1,x86-016:1,x86-037:1"
    index="0" nodecount="9" proccount="8" reqid="0"
starttime="1162425519"></range>
    <range duration="1200" hostlist="x86-001:1,x86-002:1,x86-004:1,x86-006:1,x86-
013:1,x86-014:1,x86-015:1,x86-016:1,x86-037:1,x86-042:1,x86-043:1"
    index="0" nodecount="11" proccount="8" reqid="0"
starttime="1162956803"></range>
    <range duration="1200" hostlist="x86-001:1,x86-002:1,x86-004:1,x86-006:1,x86-
013:1,x86-014:1,x86-015:1,x86-016:1,x86-027:1,x86-028:1,
    x86-029:1,x86-030:1,x86-037:1,x86-041:1,x86-042:1,x86-043:1,x86-046:1,x86-
047:1,x86-048:1,x86-049:1"
    index="0" nodecount="20" proccount="8" reqid="0"
starttime="1163053393"></range>
</par>
</Data>

```

**i** This command reports the original query, and the timeframe, resource size, and hostlist associated with each possible time slot.

## Related Topics

- (Optional) Install Moab Client - in the *Moab HPC Suite Installation and Configuration Guide* explains how to distribute this command to client nodes
- [mshow in a hosting environment](#)



## 3.7.29 mshow -a (mshow in a Hosting Environment)

### 3.7.29.A Basic Current and Future Requests

The *mshow* command can report information on many aspects of the scheduling environment. To request information on available resources, the `-a` flag should be used. By default, the *mshow* command resource availability query only reports resources that are immediately available. To request information on specific resources, the type of resources required can be specified using the `-w` flag as in the following example:

```
> mshow -a -w taskmem=1500,duration=600
...
```

To view current and future resource availability, the `future` flag should be set as in the following example:

```
> mshow -a -w taskmem=1500,duration=600 --flags=future
...
```

### 3.7.29.B Co-Allocation Resources Queries

In many cases, a particular request will need simultaneous access to resources of different types. The *mshow* command supports a co-allocation request specified by using multiple `-w` arguments. For example, to request 16 nodes with feature `fastcpu` and 2 nodes with feature `fastio`, the following request might be used:

```
> mshow -a -w minprocs=16,duration=1:00:00,nodefeature=fastcpu -w
minprocs=2,nodefeature=fastio,duration=1:00:00 --flags=future
Partition      Procs  Nodes  StartOffset  Duration  StartDate  ReqID=
-----
ALL            16     8      00:00:00     1:00:00  13:00:18_08/25  ReqID=0
ALL             2     1      00:00:00     1:00:00  13:00:18_08/25  ReqID=1
```

The `mshow -a` documentation contains a list of the different resources that can be queried, as well as examples on using *mshow*.

### 3.7.29.C Using Transaction IDs

By default, the *mshow* command reports simply when and where the requested resources are available. However, when the `tid` flag is specified, the *mshow* command returns both resource availability information and a handle to these resources called a Transaction ID as in the following example:

```
> mshow -a -w minprocs=16,nodefeature=fastcpu,duration=2:00:00 --flags=future,tid
Partition      Procs  Nodes  StartOffset  Duration  StartDate  TID=26 ReqID=0
-----
ALL            16    16      00:00:00     2:00:00  13:00:18_08/25  TID=26 ReqID=0
```

In the preceding example, the returned transaction ID (TID) can then be used to reserve the available resources using the `mrsvctl -c -R` command:

```
> mrsvctl -c -R 26
reservation system.1 successfully created
```

Any TID can be printed out using the `mschedctl -l trans` command:

```
Code example (replace with your own content)
```

```
> mschedctl -l trans 26 TID[26] A1='node01' A2='600' A3='1093465728' A4='ADVRES' A5='fastio'
```

Where A1 is the hostlist, A2 is the duration, A3 is the starttime, A4 are any flags, and A5 are any features.

### 3.7.29.D Using Reservation Profiles

Reservation profiles (`RSVPROFILE`) stand as templates against which reservations can be created. They can contain a hostlist, starttime, endtime, duration, access-control list, flags, triggers, variables, and most other attributes of an Administrative Reservation. The following example illustrates how to create a reservation with the exact same trigger-set:

```
-----
# moab.cfg
-----
RSVPROFILE[test1] TRIGGER=Sets=$Var1.$Var2.$Var3.!Net,EType=start,AType=exec,
    Action=/tmp/host/triggers/Net.sh,
    Timeout=1:00:00
RSVPROFILE[test1]                                TRIGGER=Requires=$Var1.$Var2.$Var3,
    Sets=$Var4.$Var5,EType=start,
    AType=exec,Action=/tmp/host/triggers/
    FS.sh+$Var1:$Var2:$Var3,Timeout=20:00
RSVPROFILE[test1]
TRIGGER=Requires=$Var1.$Var2.$Var3.$Var4.$Var5,
    Sets=!NOOSinit.OSinit,Etype=start,
    AType=exec,
    Action=/tmp/host/triggers/
    OS.sh+$Var1:$Var2:$Var3:$Var4:$Var5
RSVPROFILE[test1]
TRIGGER=Requires=NOOSini,AType=cancel,EType=start
RSVPROFILE[test1]
TRIGGER=EType=start,Requires=OSinit,AType=exec,
    Action=/tmp/host/triggers/success.sh
...
-----
```

To create a reservation with this profile the `mrsvctl -c -P` command is used:

```
> mrsvctl -c -P test1
reservation system.1 successfully created
```

### 3.7.29.E Using Reservation Groups

Reservation groups are a way for Moab to tie reservations together. When a reservation is created using multiple Transaction IDs, these transactions and their resulting reservations are tied together into one group.

```
> mrsvctl -c -R 34,35,36
reservation system.99 successfully created
reservation system.100 successfully created
reservation system.101 successfully created
```

In the preceding example, these three reservations would be tied together into a single group. The `mdiag -r` command can be used to see which group a reservation belongs to. The `mrsvctl -q diag -g` command can also be used to print out a specific group of reservations. The `mrsvctl -c -g` command can also be used to release a group of reservations.

---

#### Related Topics

- [mshow](#)

## 3.7.30 msub

### 3.7.30.A Synopsis

```
msub [-a datetime] [-A account] [-b retry_count] [-c interval] [-C directive_prefix] [-d path]
[-e path] [-E] [-F] [-h] [-I] [-j join] [-k keep] [-K] [-l resourcelist] [-L NUMA_resourcelist] [-m
mailoptions] [-M user_list] [-n node_exclusive] [-N name] [-o path] [-p priority] [-P <user>]
[-q destination] [-r yn] [-S pathlist] [-t jobarrays] [-u userlist] [-v variablelist] [-V] [-w
<path>] [-W additionalattributes] [-x] [-z] [--stagein] [--stageout] [--stageinfile] [--
stageoutfile] [--stageinsize] [--stageoutsize] [--workflowjobids] [script] [--about] [--
help] [--host=<serverHostName>] [--loglevel=<logLevel>] [--
msg=<message>] [--port=<serverPort>] [--timeout=<seconds>] [--
version] [--xml]
```

### 3.7.30.B Overview

`msub` allows users to submit jobs directly to Moab. When a job is submitted directly to a resource manager (such as Torque), it is constrained to run on only those nodes that the resource manager is directly monitoring. In many instances, a site may be controlling multiple resource managers. When a job is submitted to Moab rather than to a specific

resource manager, it is not constrained as to what nodes it is executed on. *msub* can accept command line arguments (with the same syntax as *qsub*), job scripts (in either PBS or LoadLeveler syntax), or the SSS Job XML specification.

**i** Moab must run as a root user in order for *msub* submissions to work. Workload submitted via *msub* when Moab is running as a non-root user fail immediately.

Submitted jobs can then be viewed and controlled via the [mjobctl](#) command.

**i** Flags specified in the following table are not necessarily supported by all resource managers.

### 3.7.30.C Access

When Moab is configured to run as root, any user can submit jobs via *msub*.

### 3.7.30.D Options

-a	
<b>Name</b>	Eligible Date
<b>Format</b>	[ [ [ [CC] YY] MM] DD] hhmm [ . SS ]
<b>Description</b>	Declares the time after which the job is eligible for execution.
<b>Example</b>	<pre>&gt; msub -a 12041300 cmd.pbs</pre> <p><i>Moab will not schedule the job until 1:00 pm on December 4, of the current year.</i></p>

-A	
<b>Name</b>	Account
<b>Format</b>	<ACCOUNT NAME>
<b>Description</b>	Defines the account associated with the job.

**-A****Example**

```
> msub -A research cmd.pbs
```

*Moab will associate this job with account research.*

**-b****Name**

Retry count

**Format**

&lt;retry\_count&gt;

**Description**

Defines the number of times msub should retry connecting to the server.

**Example**

```
> msub -b 5 cmd.pbs
```

*Moab will attempt to retry connecting to the server 5 times.*

**-c****Name**

Checkpoint Interval

**Format**

[n|s|c|c=&lt;minutes&gt;]

**Description**

Checkpoint of the job will occur at the specified interval.

**n** — No Checkpoint is to be performed.

**s** — Checkpointing is to be performed only when the server executing the job is shut down.

**c** — Checkpoint is to be performed at the default minimum time for the server executing the job.

**c=<minutes>** — Checkpoint is to be performed at an interval of minutes.

**Example**

```
> msub -c c=12 cmd.pbs
```

*The job will be checkpointed every 12 minutes.*


-C	
<b>Name</b>	Directive Prefix
<b>Format</b>	'<PREFIX NAME>'
<b>Default</b>	First known prefix (#PBS, #@, #BSUB, #!, #MOAB, #MSUB)
<b>Description</b>	<p>Specifies which directive prefix should be used from a job script:</p> <ul style="list-style-type: none"> <li>• It is best to submit with single quotes. '#PBS'</li> <li>• An empty prefix will cause Moab to not search for any prefix. -C "</li> <li>• Command line arguments have precedence over script arguments.</li> <li>• Custom prefixes can be used with the -C flag. -C '#MYPREFIX'</li> <li>• Custom directive prefixes must use PBS syntax.</li> <li>• If the -C flag is not given, Moab will take the first default prefix found. Once a directive is found, others are ignored.</li> </ul>
<b>Example</b>	<pre style="border: 1px dashed gray; padding: 5px;">&gt; msub -C '#MYPREFIX' cmd.pbs #MYPREFIX -l walltime=5:00:00 (in cmd.pbs)</pre> <p style="border: 1px dashed gray; padding: 5px; margin-top: 5px;"><i>Moab will use the #MYPREFIX directive specified in cmd.pbs, setting the wallclock limit to five hours.</i></p>

-d	
<b>Name</b>	Initial Working Directory
<b>Format</b>	<path>
<b>Default</b>	Depends on the RM being used. If using Torque, the default is \$HOME.
<b>Description</b>	Specifies which directory the job should execute in.
<b>Example</b>	<pre style="border: 1px dashed gray; padding: 5px;">&gt; msub -d /home/test/job12 cmd.pbs</pre> <p style="border: 1px dashed gray; padding: 5px; margin-top: 5px;"><i>The job will begin execution in the /home/test/job12 directory.</i></p>

-e	
<b>Name</b>	Error Path
<b>Format</b>	[<hostname>:]<path>
<b>Default</b>	\$SUBMISSIONDIR/\$JOBNAME.e\$JOBID
<b>Description</b>	Defines the path to be used for the standard error stream of the batch job.
<b>Example</b>	<pre>&gt; msub -e test12/stderr.txt</pre> <p><i>The STDERR stream of the job will be placed in the relative (to execution) directory specified.</i></p>

-E	
<b>Name</b>	Environment Variables
<b>Description</b>	<p>Moab adds the following variables, if populated, to the job's environment:</p> <ul style="list-style-type: none"> <li>• MOAB_ACCOUNT — Account name.</li> <li>• MOAB_BATCH — Set if a batch job (non-interactive).</li> <li>• MOAB_CLASS — Class name.</li> <li>• MOAB_DEPEND — Job dependency string.</li> <li>• MOAB_GROUP — Group name.</li> <li>• MOAB_JOBARRAYINDEX — For a job in an array, the index of the job.</li> <li>• MOAB_JOBARRAYRANGE — For a system with job arrays, the range of all job arrays.</li> <li>• MOAB_JOBID — Job ID. If submitted from the grid, grid jobid.</li> <li>• MOAB_JOBNAME — Job name.</li> <li>• MOAB_MACHINE — Name of the machine (i.e., Destination RM) that the job is running on.</li> <li>• MOAB_NODECOUNT — Number of nodes allocated to job.</li> <li>• MOAB_NODELIST — Comma-separated list of nodes (listed singly with no ppn info).</li> <li>• MOAB_PARTITION — Partition name the job is running in. If grid job, cluster scheduler's name.</li> <li>• MOAB_PROCCOUNT — Number of processors allocated to job.</li> <li>• MOAB_QOS — QOS name.</li> <li>• MOAB_SUBMITDIR — Directory from which the job was submitted.</li> </ul>

-E	
	<ul style="list-style-type: none"> <li>MOAB_TASKMAP — Node list with procs per node listed. &lt;nodename&gt;.&lt;procs&gt;</li> <li>MOAB_USER — User name.</li> </ul> <p>This feature only works with Torque/PBS.</p>
<b>Example</b>	<pre>&gt; msub -E mySim.cmd</pre> <p><i>The job mySim will be submitted with extra environment variables.</i></p>

-F	
<b>Name</b>	Script Flags
<b>Format</b>	"<STRING>"
<b>Description</b>	<p>Specifies the flags Torque will pass to the job script at execution time.</p> <div style="border: 1px solid #4a6982; border-radius: 5px; padding: 5px; background-color: #d9e1f2;"> <p> The <code>-F</code> flag is only compatible with Torque resource managers.</p> </div>
<b>Example</b>	<pre>&gt; msub -F "arg1 arg2" -1 nodes=1,walltime=60 files/job.sh</pre> <p><i>Torque will pass parameters arg1 and arg2 to the job.sh script when the job executes.</i></p>

-h	
<b>Name</b>	Hold
<b>Description</b>	Specifies that a user hold be applied to the job at submission time.
<b>Example</b>	<pre>&gt; msub -h cmd.ll</pre> <p><i>The job will be submitted with a user hold on it.</i></p>

-i	
<b>Name</b>	Interactive



-I	
<b>Description</b>	<p>Declares the job is to be run interactively.</p> <div style="border: 1px solid #4a6982; padding: 5px; margin-top: 10px;"> <p><b>i</b> <i>qsub</i> must exist on the same host as <i>msub</i> if the interactive job is destined for a Torque cluster, because the interactive <i>msub</i> request will be converted to a <i>qsub -I</i> request.</p> </div>
<b>Example</b>	<pre style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 5px;">&gt; msub -I job117.sh</pre> <p style="border: 1px dashed #ccc; padding: 5px; margin-top: 5px;"><i>The job will be submitted in interactive mode.</i></p>

-j	
<b>Name</b>	Join
<b>Format</b>	[eo oe n]
<b>Default</b>	n (not merged)
<b>Description</b>	<p>If <i>eo</i> is specified, the error and output streams are merged into the <i>error</i> stream. If <i>oe</i> is specified, the error and output streams will be merged into the <i>output</i> stream.</p> <div style="border: 1px solid #4a6982; padding: 5px; margin-top: 10px;"> <p><b>i</b> If using either the <i>-e</i> or the <i>-o</i> option and the <i>-j eo oe</i> option, the <i>-j</i> option takes precedence and all standard error and output messages go to the chosen output file.</p> </div>
<b>Example</b>	<pre style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 5px;">&gt; msub -j oe cmd.sh</pre> <p style="border: 1px dashed #ccc; padding: 5px; margin-top: 5px;"><i>STDOUT and STDERR will be merged into one file.</i></p>

-k	
<b>Name</b>	Keep
<b>Format</b>	[e o eo oe n]
<b>Default</b>	n (not retained)
<b>Description</b>	Defines which (if either) of output and error streams will be retained on the

-k	
	execution host (overrides path for stream).
<b>Example</b>	<pre>&gt; msub -k oe myjob.sh</pre> <p><i>STDOUT and STDERR for the job will be retained on the execution host.</i></p>

-K	
<b>Name</b>	Continue Running
<b>Format</b>	N/A
<b>Description</b>	<p>Tells the client to continue running until the submitted job is completed. The client will query the status of the job every 5 seconds. The time interval between queries can be specified or disabled via <a href="#">MSUBQUERYINTERVAL</a>.</p> <div style="border: 1px solid #004a7c; padding: 5px; background-color: #e6f2ff;"> <p><b>i</b> Use the <code>-K</code> option sparingly (if at all) as it slows down the Moab scheduler with frequent queries. Running ten jobs with the <code>-K</code> option creates an additional fifty queries per minute for the scheduler.</p> </div>
<b>Example</b>	<pre>&gt; msub -K newjob.sh 3 Job 3 completed*</pre> <p><i>*Only shows up after job completion.</i></p>

-l	
<b>Name</b>	Resource List
<b>Format</b>	<p>&lt;STRING&gt;  <code>-l [BANDWIDTH   DDISK   DEADLINE   DEPEND   DMEM   EXCLUDENODES   FEATURE...   ]</code>            Additional options can be referenced on the <a href="#">resource manager extensions</a> page.</p>
<b>Description</b>	<p>Defines the resources that are required by the job and establishes a limit to the amount of resource that can be consumed. Resources native to the resource manager, scheduler <a href="#">resource manager extensions</a>, or <a href="#">job flags</a> can be specified.</p>

-l	
	<p>Note that resource lists are dependent on the resource manager in use. For information on specifying multiple types of resources for allocation, see <a href="#">Multi-Req Support</a>.</p> <div style="border: 1px solid #4a6984; padding: 5px; margin-top: 10px;"> <p><b>i</b> Moab does not support the combination of <code>msub -l excludenodes</code> and <code>ENABLEHIGHTHROUGHPUT TRUE</code>.</p> </div>
<b>Example</b>	<div style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 10px;"> <pre>&gt; msub -l nodes=32:ppn=2,pmem=1800mb,walltime=3600,VAR=testvar:myvalue cmd.sh</pre> </div> <div style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 10px;"> <p><i>The job requires 32 nodes with 2 processors each, 1800 MB per task, a walltime of 3600 seconds, and a variable named testvar with a value of myvalue.</i></p> </div> <div style="border: 1px solid #4a6984; padding: 5px; margin-bottom: 10px;"> <p><b>i</b> If <code>JOBNODEMATCHPOLICY</code> is not set, Moab does not reserve the requested number of processors on the requested number of nodes. It reserves the total number of requested processors (nodes x ppn) on any number of nodes. Rather than setting <code>nodes=&lt;value&gt;:ppn=&lt;value&gt;</code>, set <code>procs=&lt;value&gt;</code>, replacing <code>&lt;value&gt;</code> with the total number of processors the job requires. Note that <code>JOBNODEMATCHPOLICY</code> is not set by default.</p> </div> <div style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 10px;"> <pre>&gt; msub -l nodes=32:ppn=2 -l advres=!&lt;resvid&gt;</pre> </div> <div style="border: 1px dashed #ccc; padding: 5px;"> <p><i>This entry would tell Moab to only consider resources other than the specified &lt;reservation id&gt;.</i></p> </div>
-L	
<b>Name</b>	NUMA req_information
<b>Description</b>	<div style="border: 1px solid #4a6984; padding: 5px; margin-bottom: 10px;"> <p><b>i</b> Available with Moab 9.0 or later with Torque 6.0 or later. This uses a different syntax than the <code>-l</code> resource_list option.</p> </div> <p>Defines the NUMA-aware resource requests for NUMA hardware. This option will work with non-NUMA hardware.</p> <p>See '-L NUMA Resource Request' in the <i>Torque Resource Manager Administrator Guide</i> for the syntax and valid values.</p>
-m	
<b>Name</b>	Mail Options

-m	
<b>Format</b>	<STRING> (either n or one or more of the characters a, b, and e)
<b>Description</b>	Defines the set of conditions (abort,begin,end) when the server will send a mail message about the job to the user.
<b>Example</b>	<pre>&gt; msub -m be cmd.sh</pre> <p><i>Mail notifications will be sent when the job begins and ends.</i></p>

-M	
<b>Name</b>	Mail List
<b>Format</b>	<user>[@<host>] [,<user>[@<host>], ...]
<b>Default</b>	\$JOBOWNER
<b>Description</b>	Specifies the list of users to whom mail is sent by the execution server. Overrides the EMAILADDRESS specified on the <a href="#">USERCFG</a> credential.
<b>Example</b>	<pre>&gt; msub -M jon@node01,bill@node01,jill@node02 cmd.sh</pre> <p><i>Mail will be sent to the specified users if the job is aborted.</i></p>


-n	
<b>Name</b>	Node Exclusive
<b>Description</b>	Allows a user to specify an exclusive-node access/allocation request for the job. See <a href="#">SINGLEJOB - page 365</a> for more information.
<b>Example</b>	<pre>&gt; msub -n job1187.sh</pre> <p><i>Job will have exclusive access to each node on which it runs.</i></p>

-N	
<b>Name</b>	Name
<b>Format</b>	<STRING>
<b>Default</b>	STDIN or name of job script
<b>Description</b>	Specifies the user-specified job name attribute.
<b>Example</b>	<pre>&gt; msub -N chemjob3 cmd.sh</pre> <p><i>Job will be associated with the name chemjob3.</i></p>

-o	
<b>Name</b>	Output Path
<b>Format</b>	[<hostname>:]<path> - %J and %I are acceptable variables. %J is the master array name and %I is the array member index in the array.
<b>Default</b>	\$(SUBMISSIONDIR)/\$(JOBNAME).o\$(JOBID)
<b>Description</b>	<p>Defines the path to be used for the standard output stream of the batch job. More variables are allowed when they are used in the job script instead of <code>msub -o</code>. In the job script, specify a <code>#PBS -o</code> line and input your desired variables. The allowable variables are:</p> <ul style="list-style-type: none"> <li>• OID</li> <li>• OTYPE</li> <li>• USER</li> <li>• OWNER</li> <li>• JOBID</li> <li>• JOBNAME</li> </ul> <p>Submitting a job script that has the line <code>#PBS -o \$(USER)_\$(JOBID)_\$(JOBNAME).txt</code> results in a file called <code>&lt;username&gt;_&lt;jobID&gt;_&lt;jobName&gt;.txt</code>.</p> <p>Do not use <code>msub -o</code> when submitting a job script that has a <code>#PBS -o</code> line defined.</p>

-o	
<b>Example</b>	<pre>&gt; msub -o test12/stdout.txt</pre> <p><i>The STDOUT stream of the job will be placed in the relative (to execution) directory specified.</i></p> <pre>&gt; msub -t 1-2 -o /home/jsmith/simulations/%J-%I.out ~/sim5.sh</pre> <p><i>A job array is submitted and the name of the output files includes the master array index and the array member index.</i></p>

-p	
<b>Name</b>	Priority
<b>Format</b>	<INTEGER> (between -1024 and 0)
<b>Default</b>	0
<b>Description</b>	Defines the priority of the job. To enable priority range from -1024 to +1023, see <a href="#">ENABLEPOSUSERPRIORITY</a> .
<b>Example</b>	<pre>&gt; msub -p 25 cmd.sh</pre> <p><i>The job will have a user priority of 25.</i></p>

-P	
<b>Name</b>	Proxy User
<b>Format</b>	<user>[:<group>]
<b>Description</b>	<p>Allows a root user or manager to submit a job as another user. Moab treats proxy jobs as though the jobs were submitted by the supplied username.</p> <div style="border: 1px solid #000; padding: 5px; background-color: #e6f2ff;"> <p> This option can only be used by users in the ADMINCFG[1] security level.</p> </div>
<b>Example</b>	<pre>msub -P user1 cmd.pbs</pre>

-q	
<b>Name</b>	Destination Queue (Class)
<b>Format</b>	[<queue>] [@<server>]
<b>Default</b>	[<DEFAULT>]
<b>Description</b>	<p>Defines the destination of the job.</p> <div style="border: 1px solid #003366; padding: 5px; margin-top: 10px;"> <p><b>i</b> If no destination queue is specified and the environment variable MOAB_DEFAULTQUEUE is present, msub will use the environment variable when submitting the job.</p> </div>
<b>Example</b>	<pre style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 5px;">&gt; msub -q priority cmd.sh</pre> <p style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 5px;"><i>The job will be submitted to the priority queue.</i></p>

-r	
<b>Name</b>	Rerunable
<b>Format</b>	[y n]
<b>Default</b>	n
<b>Description:</b>	Declares whether the job is rerunable.
<b>Example</b>	<pre style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 5px;">&gt; msub -r n cmd.sh</pre> <p style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 5px;"><i>The job cannot be rerun.</i></p> <div style="border: 1px solid #003366; padding: 5px; margin-top: 10px;"> <p><b>i</b> The default for qsub -r is 'y' (yes), which is opposite from msub -r. For better clarity, use the following instead.</p> <pre style="border: 1px dashed #ccc; padding: 5px; margin-top: 5px;">msub -l [flags jobflags]=restartable]</pre> </div>

-S	
<b>Name</b>	Shell

-S	
<b>Format</b>	<path>[@<host>] [,<path>[@<host>], ...]
<b>Default</b>	\$SHELL
<b>Description</b>	Declares the shell that interprets the job script.
<b>Example</b>	<pre>&gt; msub -S /bin/bash</pre> <p><i>The job script will be interpreted by the /bin/bash shell.</i></p>

-t	
<b>Name</b>	Job Arrays
<b>Format</b>	<name>[<indexlist>]%<limit>
<b>Description</b>	Starts a job array with the jobs in the index list. The limit variable specifies how many jobs can run at a time. For more information, see <a href="#">Submitting Job Arrays</a> .
	<div style="border: 1px solid #4a6982; padding: 5px; background-color: #e6f2ff;"> <p><b>i</b> Moab enforces an internal limit of 100,000 subjobs that a single array job submission can specify.</p> </div>
<b>Example</b>	<pre>&gt; msub -t myarray[1-1000]%4</pre>

-u	
<b>Name</b>	User List
<b>Format</b>	<user>[@<host>[,<user>[@<host>], ...]
<b>Default</b>	UID of <i>msub</i> command.
<b>Description</b>	Defines the user name under which the job is to run on the execution system.
<b>Example</b>	<pre>&gt; msub -u bill@node01 cmd.sh</pre> <p><i>On node01 the job will run under Bill's UID, if permitted.</i></p>



-v	
<b>Name</b>	Variable List
<b>Format</b>	<string>[,<string>,...]
<b>Description</b>	Retrieves the values of the included environment variables on the job submission node (if no value is provided) or defines a name and value and exports these variables to the job's compute node(s).
<b>Example</b>	<pre>&gt; msub -v DEBUG cmd.sh</pre> <p><i>The DEBUG environment variable on the job submission node will be defined for the job.</i></p> <pre>&gt; msub -v VAR1=xxx cmd.sh</pre> <p><i>The VAR1 environment variable will be defined for the job, with a value of xxx.</i></p>

-V	
<b>Name</b>	All Variables
<b>Description</b>	Declares that all environment variables in the msub environment are exported to the batch job.
<b>Example</b>	<pre>&gt; msub -V cmd.sh</pre> <p><i>All environment variables will be exported to the job.</i></p>

-w	
<b>Name</b>	Working Directory
<b>Format</b>	<path>
<b>Description</b>	Defines the working directory path to be used for the job. If the -w option is not specified, the default working directory is the current directory. This option sets the environment variable PBS_O_WORKDIR.
<b>Example</b>	<pre>&gt; msub -l -w /tmp</pre>

-W	
<b>Name</b>	Additional Attributes
<b>Format</b>	<string>
<b>Description</b>	Allows for specification of additional job attributes (see <a href="#">Resource Manager Extension</a> ).
<b>Example</b>	<pre style="border: 1px dashed gray; padding: 5px;">&gt; msub -W x=GRES:matlab:1 cmd.sh</pre> <p style="border: 1px dashed gray; padding: 5px; margin: 5px 0;"><i>The job requires one resource of matlab.</i></p> <p>This flag can be used to set a filter for what namespaces will be passed from a job to a trigger using a comma-delimited list. This limits the trigger's action to objects contained in certain workflows. For more information, see <a href="#">Requesting Name Space Variables</a>.</p> <pre style="border: 1px dashed gray; padding: 5px;">&gt; msub -W x="trigns=vc1,vc2"</pre> <p style="border: 1px dashed gray; padding: 5px; margin: 5px 0;"><i>The job passes namespaces vc1 and vc2 to triggers.</i></p>

-x	
<b>Format</b>	<script> or <command>
<b>Description</b>	<p>When running an interactive job, the <code>-x</code> flag makes it so that the corresponding script won't be parsed for PBS directives, but is instead a command that is launched once the interactive job has started. The job terminates at the completion of this command. This option works only when using Torque.</p> <div style="border: 1px solid #4a6984; padding: 5px; margin: 5px 0;"> <p><b>i</b> The <code>-x</code> option for <code>msub</code> differs from <code>qsub</code> in that <code>qsub</code> does not require the script name to come directly after the flag. The <code>msub</code> command requires a script or command immediately after the <code>-x</code> declaration.</p> </div>
<b>Example</b>	<pre style="border: 1px dashed gray; padding: 5px;">&gt; msub -I -x ./script.pl &gt; msub -I -x /tmp/command</pre>

-z	
<b>Name</b>	Silent Mode

-z	
<b>Description</b>	The job's identifier will not be printed to stdout upon submission.
<b>Example</b>	<pre>&gt; msub -z cmd.sh</pre> <p><i>No job identifier will be printout the stdout upon successful submission.</i></p>

## Staging Data

Data staging, or the ability to copy data required for a job from one location to another or to copy resulting data to a new location (see [Data Staging Example](#) for more information), must be specified at job submission. To stage data in, you would use the `msub --stagein` and/or `--stageinfile` option, optionally with `--stageinsize`. You would use similar options the same way for staging out: `--stageout`, `--stageoutfile`, and `--stageoutsize`. `--stagein` and `--stageout`, which you can use multiple times in the same `msub` command, allow you to specify a single file or directory to stage in or out. `--stageinfile` and `--stageoutfile` allow you to specify a text file that lists the files to stage in or out. The `--stageinsize` and `[--stageoutsize]` options allow you to estimate the total size of the files and directories that you want to stage in or out, which can help Moab make an intelligent guess about how long it will take to stage the data in or out, therefore ensuring that the job can start as soon as possible after the staging has occurred.

## Staging a File or Directory

The `--stagein` and `--stageout` options use the same format.

```
--<stagein|stageout><=<| ><source>%<destination>
```

Where `<source>` and `<destination>` take on the following format:

```
[<user>@]<host>: /<path> [ /<fileName> ]
```

Specifying a user and file name are optional. If you do not specify a file name, Moab will assume a directory.

```
> msub ... --stagein=student@biology:/stats/file001%admin@moab:/tmp/staging
<jobScript>
```

*This msub commands tells Moab that the job requires file001 from student's stats directory on the biology server to be staged to admin's staging directory on the moab server prior to the job's starting.*

You can specify the option multiple times for the same `msub` command; however, staging large number of files is easier with `--stageinfile` or `--stageoutfile`.

You can also use `#MSUB` or `#PBS` within a job script to specify data staging options. For example:

```
#MSUB --stageinsize=1gb
#MSUB --stagein=...
```

See [Sample User Job Script](#) for more information. Note that the data staging options are not compatible with `qsub`.

## Staging Multiple Files or Directories

The `--stageinfile` and `--stageoutfile` options use the same format. You must include the path to a text file that lists each file to stage in or out on its own line. Each file specification follows the same format as a `--stagein` or `--stageout` specification as described above. The format of the command options looks like this:

```
--<stageinfile|stageoutfile><=<| ><path>/<fileName>
```

The file contains multiple lines with the following format:

```
[<user>@]<host>:/<path>[/<fileName>]%[<user>@]<host>:/<path>
[/<fileName>]
```

...

Moab ignores blank lines in the file. You can comment out lines by preceding them with a pound sign (`#`). The following examples demonstrate what the `--stageinfile` option looks like on the command line and what the file it specifies might look like.

```
> msub ... --stageinfile=/tmp/myStagingFile <jobScript>
```

```
/tmp/myStagingFile:
```

```
student@biology:/stats/file001%moab:/tmp/staging
student@biology:/stats/file002%moab:/tmp/staging
student@biology:/stats/file003%moab:/tmp/staging
#student@biology:/stats/file004%moab:/tmp/staging
student@biology:/stats/file005%moab:/tmp/staging

student@biology:/stats/file006%moab:/tmp/staging
student@biology:/stats/file007%moab:/tmp/staging
student@biology:/stats/file008%moab:/tmp/staging
student@biology:/stats/file009%moab:/tmp/staging
student@biology:/stats/file010%moab:/tmp/staging
```

*Moab stages in each file listed in myStagingFile to the /tmp/staging directory. Each file resides on the biology host as the student user. Moab ignores the blank line and the line specifying file004.*

## Stage in or out File Size

The optional `--stageinsize` and `--stageoutsize` options give you the opportunity to estimate the size of the file(s) or directory(-ies) being staged to aid Moab in choosing an appropriate start time. Both options use the same format:

```
--<stageinsize|stageoutsize>=<integer>[unit]
```

The integer indicates the size of the file(s) and directory(-ies) in megabytes unless you specify a different unit. Moab accepts the follow case-insensitive suffixes: KB, MB, GB, or TB.

```
> msub --stageinfile=/stats/file003 --stageinsize=100 <jobScript>
```

*Moab copies the /davidharris/research/recordlist file, which is approximately 100 megabytes, from the biology node to the host where the job will run prior to job start.*

```
> msub --stageinfile=/stats/file002 --stageinsize=1gb <jobScript>
```

*Moab copies all files specified in the /davidharris/research/recordlist file, which add up to approximately 1 gigabyte, to the host where the job will run prior to job start.*

## Return all the Job IDs in the Workflow at Submission Time

By default, `msub` will print the job ID to stdout at the time of submission. If you want `msub` to print all of the jobs that are created as part of the workflow template, you can use the `msub --workflowjobids` option to show all the job IDs at submission time:

```
$ echo sleep 60 | msub -l walltime=15 --workflowjobids
```

```
MoabA.3.dsin MoabA.3 MoabA.3.dsout
```

## Job Script

The `msub` command supports job scripts written in any one of the following languages:

Language	Notes
<b>PBS/Torque Job Submission Language</b>	---
<b>SSS XML Job Object Specification</b>	---

## Low Latency

The `msub` can be configured to return a job ID very quickly by eliminating the processing of some job attributes, filters, remap classes, job arrays, templates, workflows, limits and other information when a job is submitted. This can be done globally by configuring `DISPLAYFLAGS USENOBLOCKMSUB` or on the individual job submission by appending "--noblock" to the command line.

**i** It is recommended that when using a non-blocking `msub` that `JOBIDFORMAT` be configured (and `PROXYJOB SUBMISSION` if desired).

### 3.7.30.E /etc/msubrc

Sites that want to automatically add parameters to every job submission can populate the file `/etc/msubrc` with global parameters that every job submission will inherit.

For example, if a site wanted every job to request a particular generic resource the following `/etc/msubrc` could be used:

```
-W x=GRES:matlab:2
```

### 3.7.30.F Usage Notes

`msub` is designed to be as flexible as possible, allowing users accustomed to PBS or LoadLeveler syntax, to continue submitting jobs as they normally would. It is not recommended that different styles be mixed together in the same `msub` command.

When only one resource manager is configured inside of Moab, all jobs are immediately staged to the only resource manager available. However, when multiple resource managers are configured Moab will determine which resource manager can run the job soonest. Once this has been determined, Moab will stage the job to the resource manager.

It is possible to have Moab take a best effort approach at submission time using the `forward` flag. When this flag is specified, Moab will do a quick check and make an intelligent guess as to which resource manager can run the job soonest and then immediately stage the job.

Moab can be configured to instantly stage a job to the underlying resource manager (like Torque/LOADLEVELER) through the parameter `INSTANTSTAGE`. When set inside `moab.cfg`, Moab will migrate the job instantly to an appropriate resource manager. Once migrated, Moab will destroy all knowledge of the job and refresh itself based on the information given to it from the underlying resource manager.

In most instances Moab can determine what syntax style the job belongs to (PBS or LoadLeveler); if Moab is unable to make a guess, it will default the style to whatever resource manager was configured at compile time. If LoadLeveler and PBS were both compiled then LoadLeveler takes precedence.

Moab can translate a subset of job attributes from one syntax to another. It is therefore possible to submit a PBS style job to a LoadLeveler resource manager, and vice versa, though not all job attributes will be translated.

### 3.7.30.G Examples

```
> msub -l nodes=3:ppn=2,walltime=1:00:00,pmem=100kb script2.pbs.cmd
4364.orion
```

This example is the XML-formatted version of the above example. See [Submitting Jobs via msub in XML](#) for more information.

```
<job>
  <InitialWorkingDirectory>/home/user/test/perlAPI
</InitialWorkingDirectory>
  <Executable>/home/user/test/perlAPI/script2.pbs.cmd
</Executable>
  <SubmitLanguage>PBS</SubmitLanguage>
  <Requested>
    <Feature>ppn2</Feature>
    <Processors>3</Processors>
    <WallclockDuration>3600</WallclockDuration>
  </Requested>
</job>
```


---

## Related Topics

- (Optional) Install Moab Client - in the *Moab HPC Suite Installation and Configuration Guide* explains how to distribute this command to client nodes
- [mjobctl](#) command to view, modify, and cancel jobs
- [checkjob](#) command to view detailed information about the job
- [mshow](#) command to view all jobs in the queue
- [MSUBQUERYINTERVAL](#) parameter
- [SUBMITFILTER](#) parameter
- [Applying the msub Submit Filter](#) for job script sample

### 3.7.30.H Applying the msub Submit Filter

When you use [msub](#) to submit a job, *msub* processes the input, converts it to XML, and sends the job specification XML to the Moab scheduler. You can create a submission filter to modify the job XML based on the criteria you set before Moab receives and processes it.

*Image 3-1: Job submission process*

The filter gives you the ability to customize the submission process, which is helpful if jobs should have certain defaults assigned to them, if you want to keep detailed submission statistics, or if you want to change job requests based on custom needs.

The submit filter, is a simple executable or script that receives XML via its standard input and returns the modified XML in its standard output. It modifies the attributes of the job specification XML based on policies you specify. It can perform various other actions at your request, too; for instance, logging. Once the submit filter has modified the job XML based on your criteria, it writes the XML representing the actual job submission to stdout. The new XML could potentially match the original XML, depending on whether the job met the criteria for modification set in the job submit filter script. Job submissions you want to proceed will leave the filter with an exit code of 0 and continue to Moab for scheduling. If the job meets the filter's specified criteria for rejection, it exits with a non-zero value, aborting the job submission process. You can configure the filter script to write a descriptive rejection message to stderr.

Job submit filters follow these rejection rules: 1) *msub* will reject job XML with an exit code of anything other than zero, 2) the *msub* command displays filter's error output on the command line, 3) *msub* will reject the job if the filter outputs invalid job XML, and 4) *msub* will reject the job if it violates any policies in your general Moab configuration; you cannot use a submit filter to bypass other policies.

To see the schema for job submission XML, refer to [Submitting Jobs via msub in XML](#).

## Submit Filter Types

You can implement submit filters on either the client or server side of a job submission. The primary differences between the two submit filter types are the location from which the filter runs, the powers and privileges of the user running the filter, and whether a user can bypass the filter. Client-based submit filters run from the *msub* client as the user who submits the job and can be bypassed, and server-based submit filters run from the Moab server as the user as which the server is running and cannot be bypassed.



### Client-Based Submit Filter

Client-based filters run from the `msub` client as the user who is submitting the job. Because they do not have elevated privileges, the risk of client-based submit filters' being abused is low; however, it is possible for the client to specify its own configuration file and bypass the filter or substitute its own filter. Job submissions do not even reach the server if a client-based submit filter rejects it.

To configure `msub` to use the submit filter, give each submission host access to the submit filter script and add a `SUBMITFILTER` parameter to the Moab configuration file (`moab.cfg`) on each submission host. The following example demonstrates how you might modify the `moab.cfg` file:

```
SUBMITFILTER /home/submitfilter/filter.pl
```

If you experience problems with your submit filter and want to debug its interaction with `msub`, enter `msub --loglevel=9`. This will cause `msub` to print verbose log messages to the terminal.

### Server-Based Submit Filter

Server-based submit filters run from the Moab server as the user as which the server is running. Because it runs as a privileged user, you must evaluate the script closely for security implications. A client configuration cannot bypass the filter.

To configure Moab to automatically apply a filter to all job submissions, use the `SERVERSUBMITFILTER` parameter. `SERVERSUBMITFILTER` specifies the path to a global job submit filter script, which Moab will run on the head node and apply to every job submitted.

```
SERVERSUBMITFILTER /opt/moab/scripts/jobFilter.pl
```

*Moab runs `jobFilter.pl`, located in the `/opt/moab/scripts` directory, on the head node, applying the filter to all jobs submitted.*

### OutputFormat XML Tag

The 'OutputFormat' element is used by a job submit filter to alter the output of the `msub` command when it reports the submitted job's job ID. For example, if a job submit filter performs a complex procedure on behalf of the user, such as submitting system jobs for a predefined workflow to accomplish some function, the filter can set this element to a value that permits it to return the job IDs of the system jobs it submitted in addition to the user's job ID the `msub` command returns.

To illustrate this element's functionality using the Moab/DataWarp integration example, a DataWarp job submit filter submits a 'DataWarp instance creation/input data staging' script as a system job and a corresponding 'output data staging/DataWarp instance destruction' script as another system job, and then ties them together with job dependencies in a 'DataWarp job workflow' that causes the user job's execution to depend on the successful completion of the DataWarp creation/input staging job and the DataWarp output staging/DataWarp Destruction system job to depend on the user job, regardless whether it completes successfully or not, or is cancelled. This DataWarp 3-job workflow guarantees the proper creation and destruction of job-based DataWarp storage; all set up and accomplished by a job submit filter.

However, users often create job workflows that have dependencies between their own jobs and may require the job IDs of all jobs to be made available in order to build a desired job workflow (i.e., 'jobB' may require 'jobA' to complete before 'jobB' is able to run). For example, if jobA was a DataWarp job and jobB should not run unless JobA successfully completes, but not until JobA's output data files are successfully staged, jobB must depend on jobA's job ID, as well as jobA's 'output data staging/DataWarp instance destruction' system job's job ID. The user can indicate jobB's job dependencies when jobA is a DataWarp job using the job submission option:

```
-l depend=afterok:<jobAid>:<jobAoutputSystemJobId>.
```

The OutputFormat XML tag provides a way for a job submit filter to pass the job IDs of additional jobs it submitted to perform a service on behalf of the user's job.

**i** The <OutputFormat> tag must be added to the job tag. If it is outside, it is treated as an invalid XML.

For example, you might submit a job and a job submit filter submits two additional jobs to assist it; the first additional job, 'job11', will run before your job, and the second additional job, 'job12', needs to run after your job finishes. If the job submit filter requires them to output in the order of 'pre', 'user', and 'post' job IDs (which is the same order Moab outputs job IDs for user jobs with input and output data-staging options), it would return the following OutputFormat element as the user's job ID string:

```
<OutputFormat>moab.11 %s moab.12</OutputFormat>
```

*msub displays the user ID string as "Moab.11 Moab.13 Moab.12"*

This means that you can have all three job IDs delivered to the end user, or a job workflow generation script in an easy to read format.

## Sample Submit Filter Script

The following example is a trivial implementation that will not affect whether a job is submitted. Use it as reference to verify that you are writing your filter properly.

```
#!/usr/bin/perl
use strict;

## Simple filter example that re-directs the output to a file.

my $file = "xmllog.out";

open FILE, ">>$file" or die "Couldn't open $file: $!";
while (<>)
{
print FILE;
print;
}
close FILE;
```

### 3.7.30.I Submitting Jobs Via *msub* in XML

The following describes the XML format used with the *msub* command to submit a job to a Moab server. This information can be used to implement a filter and modify the XML normally generated by the *msub* command. The XML format described in what follows is based on a variant of the [Scalable Systems Software Job Object Specification](#).

## Overall XML Format

The overall format of an XML request to submit a job can be shown through the following example:

```
<job>
**job attribute children**
</job>
```

An example of a simple job element with all the required children for a job submission is as follows:

```
<job>
  <Owner>user</Owner>
  <UserId>user</UserId>
  <GroupId>group</GroupId>
  <InitialWorkingDirectory>/home/user/directory</InitialWorkingDirectory>
  <UMask>18</UMask>
  <Executable>/full/path/to/script/or/first/line/of/stdin</Executable>
  <SubmitLanguage>Resource Manager Type</SubmitLanguage>
  <SubmitString>\START\23!/usr/bin/ruby\0contents\20of\20script</SubmitString>
</job>
```

The section that follows entitled Job Element Format describes the possible attributes and their meanings in detail. In actuality, all that is needed to run a job in Moab is something similar to the following:

```
<job>
  <SubmitString>\START\23!/bin/sh\0asleep\201000</SubmitString>
</job>
```

*This piece of XML requests Moab to submit a job using the contents of the SubmitString tag as a script, which is in this case a simple sh script to sleep for 1000 seconds. The msub command will create default values for all other needed attributes.*

## Job Element Format

The job element of the submission request contains a list of children and string values inside the children that represent the attribute/value pairs for the job. The earlier section, Overall XML Format, gives an example of this format. This section explains these attributes in detail.

**Arguments** — The arguments to be passed to the program are normally specified as arguments after the first argument specifying the script to be executed.

**EligibleTime** — The minimum time after which the job is eligible. This is the equivalent of the `-a` option in `msub`. Format: [ [ [ [CC] YY] MM] DD] hhmm [ . SS ]

**Environment** — The semi-colon list of environment variables that are exported to the job (taken from the `msub` command environment). The `-V msub` flag, for example, adds all the environment variables present at the time `msub` is invoked. Environment variables are delimited by the `~rs;` characters. Following is an example of the results of the `msub -v arg1=1, arg2=2` command:

```
<Environment>arg1=1~rs;arg2=2~rs;</Environment>
```

**ErrorFile** — Defines the path to be used for the standard error stream of the batch job. This is equivalent to the `-e` flag in `msub`.

**Executable** — This is normally either the name of the script to be executed, or the first line of the script if it is passed to `msub` through standard input.

**Extension** — The resource manager extension string. This can be specified via the command line in a number of ways, including the `-W x=` directive. Some other requests, such as some extensions used in the `-l` flag, are also converted to an extension string. The element has the following format:

```
<Extension>x=extension</Extension>
```

See [Using the Extension Element to Submit Triggers](#) for additional information on the extension element.

**GroupId** — The string name of the group of the user submitting the job. This will correspond to the user's primary group on the operating system.

**Hold** — Specifies that a user hold be applied to the job at submission time. This is the equivalent to the *msub* flag `-h`. It will have the form:

```
<Hold>User</Hold>
```

**InitialWorkingDirectory** — Specifies in which directory the job should begin executing. This is equivalent to the `-d` flag in the *msub* command.

```
<InitialWorkingDirectory>/home/user/directory</InitialWorkingDirectory>
```

**Interactive** — Specifies that the job is to be interactive. This is the equivalent of the `-I` flag in *msub*.

```
<Interactive>TRUE</Interactive>
```

**JobName** — The user-specified job name attribute. This is equivalent to the `-N` flag in *msub*.

**NotificationList** — The job states after which an email should be sent and also specifies the users to be emailed. This is the equivalent of the `-m` and `-M` options in *msub*.

```
<NotificationList URI=user1:user2>JobFail,JobStart,JobEnd</NotificationList>
```

*In this example, the command `msub -m abc -M user1:user2` ran indicating that emails should be sent when a job fails, starts, or ends, and that they should be sent to `user1` and `user2`.*

**OutputFile** — Defines the path to be used for the standard output stream of the batch job. This is the equivalent of the `-o` flag in *msub*.

**Priority** — A user-requested priority value. This is the equivalent to the *msub* `-p` flag.

**ProjectId** — Defines the account associated with the job. This is equivalent to the `-A` *msub* flag.

**QueueName** — The requested class of the job. This is the equivalent of the *msub* `-q` flag.

**Requested** — Specifies resources and attributes the job specifically requests and has the following form:

```
<Requested>
  <... requested attributes>
</Requested>
```

See the section dedicated to requestable attributes in this element.

**RMFlags** — Flags that will get passed directly to the resource manager on job submission. This is equivalent to any arguments listed after the `-l` *msub* flag.

```
<RMFlags>arg1 arg2 arg3</RMFlags>
```

**ShellName** — Declares the shell that interprets the job script. This is equivalent to the *msub* flag *-S*.

**SubmitLanguage** — Resource manager whose language the job is using. Use Torque to specify a Torque resource manager.

**SubmitString** — Contains the contents of the script to be run, retrieved either from an actual script or from standard input. This also includes all resource manager specific directives that may have been in the script already or added as a result of other command line arguments.

**TaskGroup** — Groups a set of requested resources together. It does so by encapsulating a Requested element. For example, the command `msub -l nodes=2+nodes=3:ppn=2` generates the following XML:

```
<TaskGroup>
  <Requested>
    <Processors>2</Processors>
    <TPN>2</TPN>
  </Requested>
</TaskGroup>
<TaskGroup>
  <Requested>
    <Processors>2</Processors>
  </Requested>
</TaskGroup>
```

**UserId** — The string value of the user ID of the job owner. This will correspond to the user's name on the operating system.

## Using the Extension Element to Submit Triggers

Use the Extension element to submit triggers. With the exception of certain characters, the syntax for [trigger creation](#) is the same for non-XML trigger submission. See [About Object Triggers](#) for detailed information on triggers. The ampersand (&) and less than sign (<) characters must be replaced for the XML to be valid. The following example shows how the Extension element is used to submit multiple triggers (separated by a semi-colon). Note that ampersand characters are replaced with `&amp;` in the example:

```
<Job>
  <UserId>user1</UserId>
  <GroupId>user1</GroupId>
  <Arguments>60</Arguments>
  <Executable>/bin/sleep</Executable>

  <Extension>x=trig:AType=exec&amp;Action="env"&amp;EType=start;trig:AType=exec&amp;Action="trig2.sh"&amp;EType=end</Extension>
  <Processors>3</Processors>
  <Disk>500</Disk>
  <Memory>1024</Memory>
  <Swap>600</Swap>
  <WallclockDuration>300</WallclockDuration>
  <Environment>PERL5LIB=/perl5:</Environment>
```

```
</Job>
```

## Elements Found in Requested Element

The following describes the tags that can be found in the Requested sub-element of the job element in a job submission request.

**Nodes** — A list of nodes that the job requests to be run on. This is the equivalent of the `-l hosts=<host-list> msub` directive.

```
<Requested>
  <Nodes>
    <Node>n1:n2</Node>
  </Nodes>
</Requested>
```

*In this example, the users requested the hosts n1 and n2 with the command `msub -l host=n1:n2`.*

**Processors** — The number of processors requested by the job. The following example was generated with the command `msub -l nodes=5`:

```
<Requested>
  <Processors>5</Processors>
</Requested>
```

**TPN** — Tasks per node. This is generated using the ppn resource manager extensions. For example, from `msub -l nodes=3:ppn=2`, the following results:

```
<Requested>
  <Processors>6</Processors>
  <TPN>2</TPN>
</Requested>
```

**WallclockDuration** — The requested wallclock duration of the job. This attribute is specified in the Requested element.

```
<Requested>
  <WallclockDuration>3600</WallclockDuration>
</Requested>
```

---

## Related Topics

- [Applying the msub Submit Filter](#)
- `SUBMITFILTER` parameter

## 3.7.31 mvctl (Moab Virtual Container Control)

### 3.7.31.A Synopsis

```

mvctl -a <OType>:<OName>[,<OType>:<OName>] <name>
mvctl -c [<description>]
mvctl -d <name>
mvctl -m <ATTR>=VAL[,<ATTR>=VAL] <name>
mvctl -q [<name>|ALL] [--xml][--blocking][--flags=fullxml]
mvctl -r <OType>:<OName>[,<OType>:<OName>] <name>
mvctl -x <action><name>

[--about] [--help] [--host=<serverHostName>] [--
loglevel=<logLevel>] [--msg=<message>] [--port=<serverPort>]
[--timeout=<seconds>] [--version] [--xml]

```

### 3.7.31.B Overview

A virtual container (VC) is a logical grouping of objects with a shared variable space and applied policies. Containers can hold virtual machines, jobs, reservations, and nodes. Containers can also be nested inside other containers.

A VC can be owned by a user, group, or account. Users can only view VCs to which they have access. Level 1 administrators (Admin1) can view and modify all VCs. The owner can also be changed. When modifying the owner, you must also specify the owner type:

```
mvctl -m OWNER=acct:bob myvc
```

Adding objects to VCs at submission: You associate jobs and reservations with a specified VC upon submission. For example:

- `mrsvctl -c ... -H <VC>`
- `msub ... -W x="vc=<VC>"`

**i** The user who submits objects must have access to the VC or the command is rejected.

### 3.7.31.C FullXML flag

The FullXML flag will cause the `mvctl -q` command to show VCs in a hierarchical manner. If doing a non-XML (plaintext) query, sub-VCs will be listed inside their parent



VCS. Each VC will be indented more than its parent.

```

VC[vc1] (vc1)
  Owner: user:jason
  VCs:
    VC[vc2] (vc2)
      Owner: user:jason
      Jobs: Moab.1
      Rsvs: system.1
      VCs:
        VC[vc3] (vc3)
          Owner: user:jason
    VC[vc4] (vc4)
      Owner: user:jason

```

If doing an XML query, the XML for all sub-objects (VCS, but also reservations, jobs, etc.) will also be included in the VC.

```

<Data>
  <vcs>
    <vc CREATETIME="1460666817" CREATOR="tshaw" DESCRIPTION="vc1"
      NAME="vc1" OWNER="user:tshaw" VCS="vc2,vc4">
      <ACL aff="positive" cmp="%" name="tshaw" type="USER" />
    </vc>
    <vc CREATETIME="1460666818" CREATOR="tshaw" DESCRIPTION="vc2"
      JOBS="moab.1" NAME="vc2" OWNER="user:tshaw" RSVS="system.2"
      VCS="vc3">
      <ACL aff="positive" cmp="%" name="tshaw" type="USER" />
    </vc>
    <vc CREATETIME="1460666818" CREATOR="tshaw" DESCRIPTION="vc3"
      NAME="vc3" OWNER="user:tshaw">
      <ACL aff="positive" cmp="%" name="tshaw" type="USER" />
    </vc>
    <vc CREATETIME="1460666818" CREATOR="tshaw" DESCRIPTION="vc4"
      NAME="vc4" OWNER="user:tshaw">
      <ACL aff="positive" cmp="%" name="tshaw" type="USER" />
    </vc>
  </vcs>
</Data>

```

*Note that the XML from the blocking and non-blocking commands may differ.*

### 3.7.31.D Virtual Container Flags

The following table indicates available virtual container (VC) flags and associated descriptions. Note that the `Deleting`, `HasStarted`, and `Workflow` flags cannot be set by a user but are helpful indicators of status.

VC Flags	
<b>DestroyObjects</b>	When the VC is destroyed, any reservations and jobs in the VC are also destroyed. This is recursive, so any objects in sub-VCS are also destroyed. Nodes are not removed.

VC Flags	
<b>DestroyWhenEmpty</b>	When the VC is empty, it is destroyed.
<b>Deleting</b>	Set by the scheduler when the VC has been instructed to be removed. <div style="border: 1px solid #0056b3; border-radius: 5px; padding: 5px; margin-top: 10px;"> <span style="font-size: 1.2em; font-weight: bold; color: #0056b3;">i</span> Internal flag. Administrators cannot set or clear this flag.                 </div>
<b>HasStarted</b>	This flag is set on a VC workflow where at least one job has started. <div style="border: 1px solid #0056b3; border-radius: 5px; padding: 5px; margin-top: 10px;"> <span style="font-size: 1.2em; font-weight: bold; color: #0056b3;">i</span> Internal flag. Administrators cannot set or clear this flag.                 </div>
<b>HoldJobs</b>	This flag will place a hold on any job that is submitted to the VC while this flag is set. It is not applied for already existing jobs that are added into the VC. If a job with a workflow is submitted to the VC, all jobs within the workflow are placed on hold.
<b>NoReleaseWhenScheduled</b>	Prevents Moab from lifting the UserHold on the workflow when it is scheduled. This enables an approval method where an administrator must release the hold manually before the service is allowed to start as scheduled.
<b>Workflow</b>	Designates this VC as a VC that is for workflows. This flag is set when generated by a job template workflow. Workflow jobs can only be attached to one workflow VC. <div style="border: 1px solid #0056b3; border-radius: 5px; padding: 5px; margin-top: 10px;"> <span style="font-size: 1.2em; font-weight: bold; color: #0056b3;">i</span> Internal flag. Administrators cannot set or clear this flag.                 </div>

### 3.7.31.E Options

<b>-a</b>	
<b>Format</b>	<code>mvctl -a&lt;OType&gt;:&lt;OName&gt;[,&lt;OType&gt;:&lt;OName&gt;] &lt;name&gt;</code> Where <OType> is one of JOB, RSV, NODE, or VC.

<b>-a</b>	
<b>Description</b>	Add the given object(s).
<b>Example</b>	<pre>mvcctl -a JOB:Moab.45 vc13 &gt;&gt;job 'Moab.45' added to VC 'vc13'</pre>

<b>-c</b>	
<b>Format</b>	<code><i>mvcctl -c</i> [&lt;description&gt;]</code>
<b>Description</b>	Create a virtual container (VC). The VC name is auto-generated. It is recommended that you supply a description; otherwise the description is the same as the auto-generated name.
<b>Example</b>	<pre>mvcctl -c "Linux testing machine" &gt;&gt;VC 'vc13' created</pre>

<b>-d</b>	
<b>Format</b>	<code><i>mvcctl -d</i>&lt;lab01&gt;</code>
<b>Description</b>	Destroy the VC.
<b>Example</b>	<pre>mvcctl -d vc13 &gt;&gt;VC 'vc13' destroyed</pre>

<b>-m</b>	
<b>Format</b>	<code><i>mvcctl -m</i>&lt;ATTR&gt;=VAL[,&lt;ATTR&gt;=&lt;VAL&gt;] &lt;name&gt;</code>
<b>Description</b>	Modify the VC. Attributes are flags, owner, reqstarttime, reqnodeset, variables, and owner; note that only the owner can modify owner. Use reqstarttime when implementing guaranteed start time to specify when jobs should start. The reqnodeset attribute indicates the node set that jobs should run in that are submitted to a virtual container.

**-m****Example**

```

mvcctl -m variables+=HV=node8 vc13
>>VC 'vc13' successfully modified

mvcctl -m flags+=DESTROYWHENEMPTY vc1
>>VC 'vc1' successfully modified

mvcctl -m messages="\"This VC is for internal use, etc.\"" vc5
>>VC 'vc5' successfully modified

```

**-q****Format**

```
mvcctl -q [<name>|ALL] [--xml][--blocking][--flags=fullxml]
```

**Description**

Query VCs.

**Example**

```

mvcctl -q ALL
VC[vc13] (Linux testing machine)
Create Time: 1311027343    Creator: jdoe
Owner: user:jdoe
ACL:  USER=%=jdoe+;
Jobs: Moab.45
Vars: HV=node88
Flags: DESTROYWHENEMPTY

```

**-r****Format**

```
mvcctl -r<OType>:<OName>[,<OType>:<OName>] <name>
```

Where <OType> is one of JOB, RSV, NODE, or VC.

**Description**

Remove the given object(s) from the VC.

**Example**

```

mvcctl -r JOB:Moab.45 vc13
>>job 'Moab.45' removed from VC 'vc13'

```

**-x****Format**

```
mvcctl -x<action><name>
```

**Description**

Executes the given action on the virtual container (VC).

**Example**

```
mvcctl -x schedulevc vc1
```

## 3.7.32 showbf

### 3.7.32.A Synopsis

```
showbf [-a account] [-A] [-c class] [-d duration] [-D] [-f features] [-g group] [-h] [-L] [-m
[=>|>|=<|<=] memory] [-n nodecount] [-p partition] [-q qos] [-r processorcount] [-u user]
[-v] [--blocking] [--about] [--help] [--host=<serverHostName>] [-
-loglevel=<logLevel>] [--msg=<message>] [--port=<serverPort>]
[--timeout=<seconds>] [--version] [--xml]
```

### 3.7.32.B Overview

Shows what resources are available for immediate use.

**i** The results Moab returns do not include resources that may be freed due to preemption.

This command can be used by any user to find out how many processors are available for immediate use on the system. It is anticipated that users will use this information to submit jobs that meet these criteria and therefore obtain quick job turnaround times. This command incorporates down time, reservations, and node state information in determining the available backfill window.

**i** If specific information is not specified, *showbf* will return information for the user and group running but with global access for other credentials. For example, if *-q qos* is not specified, Moab will return resource availability information for a job as if it were entitled to access all QOS based resources (i.e., resources covered by reservations with a QOS based ACL), if *-c class* is not specified, the command will return information for resources accessible by any class.

**i** The *showbf* command incorporates node configuration, node utilization, node state, and node reservation information into the results it reports. This command does not incorporate constraints imposed by credential based fairness policies on the results it reports.

### 3.7.32.C Access

By default, this command can be used by any user or administrator.

### 3.7.32.D Parameters

Parameter	Description
<b>ACCOUNT</b>	Account name.
<b>CLASS</b>	Class/queue required.
<b>DURATION</b>	Time duration specified as the number of seconds or in [DD:]HH:MM:SS notation.
<b>FEATURELIST</b>	Colon separated list of node features required.
<b>GROUP</b>	Specify particular group.
<b>MEMCMP</b>	Memory comparison used with the <code>-m</code> flag. Valid signs are <code>&gt;</code> , <code>&gt;=</code> , <code>==</code> , <code>&lt;=</code> , and <code>&lt;</code> .
<b>MEMORY</b>	Specifies the amount of required real memory configured on the node, (in MB), used with the <code>-m</code> flag.
<b>NODECOUNT</b>	Specify number of nodes for inquiry with <code>-n</code> flag.
<b>PARTITION</b>	Specify partition to check with <code>-p</code> flag.
<b>PROCESSORCOUNT</b>	Specify number of processors required.
<b>QOS</b>	Specify QOS to check with <code>-q</code> flag.
<b>USER</b>	Specify particular user to check with <code>-u</code> flag.

## 3.7.32.E Options

Option	Description
<b>-a</b>	Show resource availability information only for the specified account.
<b>-A</b>	Show resource availability information for all users, groups, and accounts. By default, <i>showbf</i> uses the default user, group, and account ID of the user issuing the command.
<b>--blocking</b>	Do not use cache information in the output. The <code>--blocking</code> flag retrieves results exclusively from the scheduler.
<b>-c</b>	Show resource availability only for the specified class.
<b>-d</b>	Show resource availability information for specified duration.
<b>-D</b>	Display current and future resource availability notation.
<b>-f</b>	Display availability for the specified colon-separated list of node features.
<b>-g</b>	Show resource availability information only for specified group.
<b>-h</b>	Help for this command.
<b>-L</b>	Enforce Hard limits when showing available resources.
<b>-m</b>	Allows user to specify the memory requirements for the backfill nodes of interest. It is important to note that if the optional MEMCMP and MEMORY parameters are used, they must be enclosed in single ticks (') to avoid interpretation by the shell. For example, enter <code>showbf -m '==256'</code> to request nodes with 256 MB memory.
<b>-n</b>	Show resource availability information for a specified number of nodes. That is, this flag can be used to force <i>showbf</i> to display only blocks of resources with at least this many nodes available.
<b>-p</b>	Show resource availability information for the specified partition.
<b>-q</b>	Show information for the specified QOS.

Option	Description
<b>-r</b>	Show resource availability for the specified processor count.
<b>-u</b>	Show resource availability information only for specified user.
<b>-v</b>	Displays verbose information.

### 3.7.32.F Examples

In this example, a job requiring up to 2 processors could be submitted for immediate execution in partition `ClusterA` for any duration. Additionally, a job requiring 1 processor could be submitted for immediate execution in partition `ClusterB`. Note that by default, each task is tracked and reported as a request for a single processor.

```
> showbf
Partition      Tasks  Nodes  StartOffset  Duration  StartDate
-----
ALL            3      3      00:00:00     INFINITY  11:32:38_08/19
ReqID=0
ClusterA       1      1      00:00:00     INFINITY  11:32:38_08/19
ReqID=0
ClusterB       2      2      00:00:00     INFINITY  11:32:38_08/19
ReqID=0
```

**i** StartOffset is the amount of time remaining before resources will be available.

#### Example 3-35:

In this example, the output verifies that a backfill window exists for jobs requiring a 3 hour runtime and at least 16 processors. Specifying job duration is of value when time based access is assigned to reservations (i.e., using the `SRCFG TIMELIMIT ACL`).

```
> showbf -r 16 -d 3:00:00
Partition      Tasks  Nodes  Duration  StartOffset  StartDate
-----
ALL            20     20     INFINITY  00:00:00     09:22:25_07/19
```

#### Example 3-36:

In this example, a resource availability window is requested for processors located only on nodes with at least 512 MB of memory.

```
> showbf -m '=512'
Partition      Tasks  Nodes  Duration  StartOffset  StartDate
-----
ALL            20     20     INFINITY  00:00:00     09:23:23_07/19
ClusterA       10     10     INFINITY  00:00:00     09:23:23_07/19
```



```
ClusterB      10      10      INFINITY      00:00:00      09:23:23_07/19
```

## Related Topics

- (Optional) Install Moab Client - in the *Moab HPC Suite Installation and Configuration Guide* explains how to distribute this command to client nodes
- [showq](#)
- [mdiag -t](#)

## 3.7.33 showq

### 3.7.33.A Synopsis

```
showq [-b] [-g] [-l] [-c|-i|-r] [-n] [-N] [-o] [-p partition] [-R rsvid] [-s] [-S] [-u] [-v] [-w
<CONSTRAINT>] [--blocking] [--noblock] [--about] [--help] [--
host=<serverHostName>] [--loglevel=<logLevel>] [--
msg=<message>] [--port=<serverPort>] [--timeout=<seconds>] [--
version] [--xml]
```

### 3.7.33.B Overview

Displays information about active, eligible, blocked, and/or recently completed jobs. Since the resource manager is not actually scheduling jobs, the job ordering it displays is not valid. The `showq` command displays the actual job ordering under the Moab Workload Manager. When used without flags, this command displays all jobs in active, idle, and non-queued states.

### 3.7.33.C Access

By default, this command can be run by any user. However, the `-c`, `-i`, and `-r` flags can only be used by level 1, 2, or 3 Moab administrators.

## 3.7.33.D Options

Option	Argument	Description
<b>-b</b>	---	Display blocked jobs only.
<b>-c</b>	---	Display details about recently <a href="#">completed</a> jobs (see <a href="#">example</a> , <a href="#">JOB</a> <a href="#">PURGETIME</a> ).
<b>-g</b>	---	Display grid job and system IDs for all jobs.
<b>-i</b>	---	Display extended details about idle jobs.
<b>-l</b>	---	Display local/remote view. For use in a <a href="#">Grid</a> environment, displays job usage of both local and remote compute resources.
<b>-n</b>	---	Displays normal showq output, but lists job names under JOBID.
<b>-N</b>	---	Show the node/task allocation of the job.
<b>-o</b>	---	Displays jobs in the active queue in the order specified (uses format <code>showq -o &lt;specifiedOrder&gt;</code> ). Valid options include REMAINING, REVERSEREMAINING, JOB, USER, STATE, and STARTTIME. The default is REMAINING.
<b>-p</b>	partition	Display only jobs assigned to the specified partition.
<b>-r</b>	---	Display extended details about active (running) jobs (see <a href="#">example</a> ).
<b>-R</b>	rsvid	Display only jobs that overlap the specified reservation.
<b>-s</b>	---	Display workload summary.
<b>-S</b>	---	Display system jobs.
<b>-u</b>	---	Display all running jobs for a particular user.

Option	Argument	Description
<b>-v</b>	---	Display local and full resource manager job IDs, as well as partitions. If specified with the <i>-i</i> option, will display job reservation time. To see a summary of array subjobs, run <i>checkjob -v &lt;jobID&gt;</i> . To see array subjobs in <i>showq</i> , include the <i>--blocking</i> option.
<b>-w</b>	constraint	Display only jobs associated with the specified constraint. Valid constraints include user, group, jobgroup, acct, nodefeature, class, and qos (see <a href="#">showq -w</a> examples).
<b>--blocking</b>	---	Do not use cache information in the output. The <i>--blocking</i> flag retrieves results exclusively from the scheduler. This option also causes <i>showq</i> to display an individual line for each array subjob.
<b>--noblock</b>	---	Use cache information for a faster response.

### 3.7.33.E Details

Beyond job information, the *showq* command will also report if the scheduler is stopped or paused or if a system reservation is in place. Further, the *showq* command will also report public system messages.

### 3.7.33.F Examples

- [Default Report](#)
  - [Detailed Active/Running Job Report](#)
  - [Eligible Jobs](#)
  - [Detailed Completed Job Report](#)
- [Filtered Job Report](#)

#### Default Report

The output of this command is divided into three parts, [Active Jobs](#), [Eligible Jobs](#), and [Blocked Jobs](#).

```
> showq
```

```

active jobs-----
JOBID USERNAME STATE PROCS REMAINING          STARTTIME
12941          sartois   Running   25    2:44:11 Thu Sep  1 15:02:50
12954          t gates   Running    4    2:57:33 Thu Sep  1 15:02:52
12944          evall    Running   16    6:37:31 Thu Sep  1 15:02:50
12946          t gates   Running    2    1:05:57:31 Thu Sep  1 15:02:50

4 active jobs          47 of 48 processors active (97.92%)
                      32 of 32 nodes active      (100.00%)

eligible jobs-----
JOBID          USERNAME          STATE  PROCS    WCLIMIT          QUEUETIME
12956          cfosdyke         Idle   32      6:40:00 Thu Sep  1 15:02:50
12969          cfosdyke         Idle    4      6:40:00 Thu Sep  1 15:03:23
12939          evall            Idle   16      3:00:00 Thu Sep  1 15:02:50
12940          mwillis          Idle    2      3:00:00 Thu Sep  1 15:02:50
12947          mwillis          Idle    2      3:00:00 Thu Sep  1 15:02:50
12949          evall            Idle    2      3:00:00 Thu Sep  1 15:02:50
12953          t gates          Idle   10      4:26:40 Thu Sep  1 15:02:50
12955          evall            Idle    2      4:26:40 Thu Sep  1 15:02:50
12957          t gates          Idle   16      3:00:00 Thu Sep  1 15:02:50
12963          evall            Idle   16    1:06:00:00 Thu Sep  1 15:02:52
12964          t gates          Idle   16    1:00:00:00 Thu Sep  1 15:02:52
12937          allendr          Idle    9    1:00:00:00 Thu Sep  1 15:02:50
12962          aacker           Idle    6      00:26:40 Thu Sep  1 15:02:50
12968          tamaker          Idle    1      4:26:40 Thu Sep  1 15:02:52

14 eligible jobs

blocked jobs-----
JOBID          USERNAME          STATE  PROCS    WCLIMIT          QUEUETIME

0 blocked jobs

Total jobs: 18

```

The fields are as follows:

Column	Description
<b>JOBID</b>	Job identifier.
<b>USERNAME</b>	User owning job.
<b>STATE</b>	Job State. Current batch state of the job.
<b>PROCS</b>	Number of processors being used by the job.
<b>REMAINING/WCLIMIT</b>	For active jobs, the time the job has until it has reached its wallclock limit or for idle/blocked jobs, the amount of time requested by the job. Time specified in [DD:]HH:MM:SS notation.

Column	Description
<b>STARTTIME</b>	Time job started running.

### Active Jobs

Active jobs are those that are **Running** or **Starting** and consuming resources. Displayed are the job ID\*, the job's owner, and the job state. Also displayed are the number of processors allocated to the job, the amount of time remaining until the job completes (given in HH:MM:SS notation), and the time the job started. All active jobs are sorted in 'Earliest Completion Time First' order.

**i** \*Job IDs can be marked with a single character to specify the following conditions:

Character	Description
<b>_ (underbar)</b>	job violates usage limit
<b>* (asterisk)</b>	job is backfilled AND is preemptible
<b>+ (plus)</b>	job is backfilled AND is NOT preemptible
<b>- (hyphen)</b>	job is NOT backfilled AND is preemptible

**i** Detailed active job information can be obtained using the `-r` flag.

### Eligible Jobs

Eligible Jobs are those that are queued and eligible to be scheduled. They are all in the Idle job state and do not violate any fairness policies or have any job holds in place. The jobs in the Idle section display the same information as the Active Jobs section except that the wallclock CPULIMIT is specified rather than job time REMAINING, and job QUEUE TIME is displayed rather than job STARTTIME. The jobs in this section are ordered by job priority. Jobs in this queue are considered eligible for both scheduling and backfilling.

**i** Detailed eligible job information can be obtained using the `-i` flag.

## Blocked Jobs

Blocked jobs are those that are ineligible to be run or queued. Jobs listed here could be in a number of states for the following reasons:

State	Description
<b>Idle</b>	Job violates a fairness policy. Use <i>diagnose -q</i> for more information.
<b>UserHold</b>	A user hold is in place.
<b>SystemHold</b>	An administrative or system hold is in place.
<b>BatchHold</b>	A scheduler batch hold is in place (used when the job cannot be run because the requested resources are not available in the system or because the resource manager has repeatedly failed in attempts to start the job).
<b>Deferred</b>	A scheduler defer hold is in place (a temporary hold used when a job has been unable to start after a specified number of attempts. This hold is automatically removed after a short period of time).
<b>NotQueued</b>	Job is in the resource manager state NQ (indicating the job's controlling scheduling daemon is unavailable).

A summary of the job queue's status is provided at the end of the output.

### Example 3-37: Detailed Active/Running Job Report

```

> showq -r

active jobs-----
JOBID          S  PAR  EFFIC  XFACTOR  Q      USER   GROUP      MHOST  PROCS
REMAINING      STARTTIME
12941          R   3  100.00    1.0  -   sartois  Arches      G5-014  25
2:43:31 Thu Sep 1 15:02:50
12954          R   3  100.00    1.0  Hi   tgates   Arches      G5-016   4
2:56:54 Thu Sep 1 15:02:52
12944          R   2  100.00    1.0  De   eval1   RedRock    P690-016  16
6:36:51 Thu Sep 1 15:02:50
12946          R   3  100.00    1.0  -   tgates   Arches      G5-001   2
1:05:56:51 Thu Sep 1 15:02:50

4 active jobs          47 of 48 processors active (97.92%)
                       32 of 32 nodes active      (100.00%)

Total jobs:  4

```

After displaying the running jobs, a summary is provided indicating the number of jobs, the number of allocated processors, and the system utilization.

Column	Description
<b>JOBID</b>	Name of active job.
<b>S</b>	<a href="#">Job State</a> . Either R for Running or S for Starting.
<b>PAR</b>	Partition in which job is running.
<b>EFFIC</b>	CPU efficiency of job.
<b>XFACTOR</b>	See <a href="#">Expansion Factor (XFACTOR) Subcomponent</a> for a detailed description.
<b>Q</b>	Quality Of Service specified for job.
<b>USER</b>	User owning job.
<b>GROUP</b>	Primary group of job owner.
<b>MHOST</b>	Master Host running primary task of job.
<b>PROCS</b>	Number of processors being used by the job.
<b>REMAINING</b>	Time the job has until it has reached its wallclock limit. Time specified in HH:MM:SS notation.
<b>STARTTIME</b>	Time job started running.

```

> showq -i
eligible jobs-----
JOBID          PRIORITY  XFACTOR  Q      USER   GROUP  PROCS  WCLIMIT
CLASS         SYSTEMQUEUE TIME
12956*
batch Thu Sep 1 15:02:50 20      1.0 -   cfosdyke RedRock 32      6:40:00
12969*
batch Thu Sep 1 15:03:23 19      1.0 -   cfosdyke RedRock 4       6:40:00
12939
batch Thu Sep 1 15:02:50 16      1.0 -   eval1   RedRock 16      3:00:00
12940
batch Thu Sep 1 15:02:50 16      1.0 -   mwillis Arches  2       3:00:00
12947
batch Thu Sep 1 15:02:50 16      1.0 -   mwillis Arches  2       3:00:00
12949
batch Thu Sep 1 15:02:50 16      1.0 -   eval1   RedRock 2       3:00:00
12953
batch Thu Sep 1 15:02:50 16      1.0 -   tgates  Arches 10      4:26:40

```

```

batch Thu Sep 1 15:02:50
12955 16 1.0 - evall RedRock 2 4:26:40
batch Thu Sep 1 15:02:50
12957 16 1.0 - tgates Arches 16 3:00:00
batch Thu Sep 1 15:02:50
12963 16 1.0 - evall RedRock 16 1:06:00:00
batch Thu Sep 1 15:02:52
12964 16 1.0 - tgates Arches 16 1:00:00:00
batch Thu Sep 1 15:02:52
12937 1 1.0 - allendr RedRock 9 1:00:00:00
batch Thu Sep 1 15:02:50
12962 1 1.2 - aacker RedRock 6 00:26:40
batch Thu Sep 1 15:02:50
12968 1 1.0 - tamaker RedRock 1 4:26:40
batch Thu Sep 1 15:02:52

14 eligible jobs

Total jobs: 14

```

The fields are as follows:

Column	Description
<b>JOBID</b>	Name of job.
<b>PRIORITY</b>	Calculated job priority.
<b>XFACTOR</b>	See <a href="#">Expansion Factor (XFACTOR) Subcomponent</a> for a detailed description.
<b>Q</b>	Quality Of Service specified for job.
<b>USER</b>	<a href="#">User</a> owning job.
<b>GROUP</b>	Primary <a href="#">group</a> of job owner.
<b>PROCS</b>	Minimum number of processors required to run job.
<b>WCLIMIT</b>	Wallclock limit specified for job. Time specified in HH:MM:SS notation.
<b>CLASS</b>	<a href="#">Class</a> requested by job.
<b>SYSTEMQUEUEETIME</b>	Time job was admitted into the system queue.



**i** An asterisk at the end of a job (job 12956\* in this example) indicates that the job has a job [reservation](#) created for it. The details of this reservation can be displayed using the [checkjob](#) command.

### Example 3-38: Detailed Completed Job Report

```
> showq -c
completed jobs-----
JOBID          SCCODE  PAR  EFFIC  XFACTOR  Q  USERNAME  GROUP  MHOST
PROC  WALLTIME  STARTTIME
13098          C      0  bas  93.17    1.0 -  sartois  Arches  G5-014
25  2:43:31  Thu Sep 1 15:02:50
13102          C      0  bas  99.55    2.2 Hi   t gates  Arches  G5-016
4  2:56:54  Thu Sep 1 15:02:52
13103          C      2  tes  99.30    2.9 De   e val1  RedRock P690-016
16  6:36:51  Thu Sep 1 15:02:50
13115          C      0  tes  97.04    1.0 -   t gates  Arches  G5-001
2  1:05:56:51 Thu Sep 1 15:02:50
3 completed jobs
```

The fields are as follows:

Column	Description
<b>JOBID</b>	job ID for completed job.
<b>S</b>	<b>Job State</b> . Either C for <a href="#">Completed</a> or V for <a href="#">Vacated</a> .
<b>CCODE</b>	Completion code reported by the job.
<b>PAR</b>	Partition in which job ran.
<b>EFFIC</b>	CPU efficiency of job.
<b>XFACTOR</b>	See <a href="#">Expansion Factor (XFACTOR) Subcomponent</a> for a detailed description.
<b>Q</b>	<a href="#">Quality of Service</a> specified for job.
<b>USERNAME</b>	<a href="#">User</a> owning job.
<b>GROUP</b>	Primary <a href="#">group</a> of job owner.
<b>MHOST</b>	Master Host that ran the primary task of job.

Column	Description
<b>PROCS</b>	Number of processors being used by the job.
<b>WALLTIME</b>	Wallclock time used by the job. Time specified in [DD:]HH:MM:SS notation.
<b>STARTTIME</b>	Time job started running.

After displaying the active jobs, a summary is provided indicating the number of jobs, the number of allocated processors, and the system utilization.

**i** If the **DISPLAYFLAGS** parameter is set to **ACCOUNTCENTRIC**, job group information will be replaced with job account information.

#### Example 3-39: Filtered Job Report

Show only jobs associated with user `john`, class `benchmark`, and nodefeature `bigmem`.

```
> showq -w class=benchmark -w user=john -w nodefeature=bigmem
...
```

#### Example 3-40: Filtered Job Report

Show only jobs associated with jobgroup `workflow1`.

```
> showq -w jobgroup=workflow1
...
```

### 3.7.33.G Job Array

Job arrays show the name of the job array and then in parenthesis, the number of subjobs in the job array that are in the specified state.

```
> showq
active jobs-----
JOBID            USERNAME        STATE  PROCS   REMAINING        STARTTIME
Moab.1(14)      aesplin        Running  14     00:59:41  Fri May 27 14:58:57
14 active jobs          14 of 14 processors in use by local jobs (100.00%)
2 of 2 nodes active      (100.00%)

eligible jobs-----
JOBID            USERNAME        STATE  PROCS   WCLIMIT        QUEUETIME
```

```

Moab.1 (4)          aesplin      Idle      4      1:00:00  Fri May 27 14:58:52
4 eligible jobs

blocked jobs-----
JOBID              USERNAME      STATE  PROCS      WCLIMIT      QUEUE TIME
Moab.1 (2)          aesplin      Blocked  2      1:00:00  Fri May 27 14:58:52
2 blocked jobs

Total jobs: 20

```

### Related Topics

- (Optional) Install Moab Client - in the *Moab HPC Suite Installation and Configuration Guide* explains how to distribute this command to client nodes
- [showbf](#) - command to display resource availability.
- [mdiag -j](#) - command to display detailed job diagnostics.
- [checkjob](#) - command to check the status of a particular job.
- [JOBPCPURGETIME](#) - parameter to adjust the duration of time Moab preserves information about completed jobs
- [DISPLAYFLAGS](#) - parameter to control what job information is displayed

## 3.7.34 showhist.moab.pl

### 3.7.34.A Synopsis

```

showhist.moab.pl [-a accountname]
                 [-c classname] [-e enddate]
                 [-g groupname] [-j jobid] [-n days]
                 [-q qosname] [-s startdate]
                 [-u username]

```

### 3.7.34.B Overview

The `showhist.moab.pl` script displays historical job information. Its purpose is similar to the [checkjob](#) command's, but `showhist.moab.pl` displays information about jobs that have already completed.

### 3.7.34.C Access

By default, this script's use is limited to administrators on the head node; however, end users can also be given power to run the script. To grant access to the script to end users, move `showhist.moab.pl` from the `tools` directory to the `bin` directory.

### 3.7.34.D Options

#### -a (Account)

**Format** <ACCOUNTNAME>

**Description** Displays job records matching the specified account.

**Example**

```
> showhist.moab.pl -a myAccount
```

*Information about jobs related to the account `myAccount` is displayed.*

#### -c (Class)

**Format** <CLASSNAME>

**Description** Displays job records matching the specified class (queue).

**Example**

```
> showhist.moab.pl -c newClass
```

*Information about jobs related to the class `newClass` is displayed.*

#### -e (End Date)

**Format** YYYY-MM-DD

**Description** Displays the records of jobs recorded before or on the specified date.

**Example**

```
> showhist.moab.pl -e 2022-01-03
```

*Information about all jobs recorded on or before January 3, 2022 is displayed.*

**-e (End Date)**

```
> showhist.moab.pl -s 2022-01-01 -e 2022-01-31
```

*Information is displayed about all jobs recorded in January 2022.*

**-g (Group)**

**Format** <GROUPNAME>

**Description** Displays job records matching the specified group.

**Example**

```
> showhist.moab.pl -g admins
```

*Information about jobs related to the group `admins` is displayed.*

**-j (Job ID)**

**Format** <JOBID>

**Description** Displays job records matching the specified job ID.

**Example**

```
> showhist.moab.pl -j moab01
```

*Information about job `moab01` is displayed.*

**-n (Number of Days)**

**Format** <INTEGER>

**Description** Restricts the number of past jobs to search by a specified number of days relative to today.

**Example**

```
> showhist.moab.pl -n 90 -j moab924
```

*Displays job information for job `moab924`. The search is restricted to the last 90 days.*

<b>-q (QoS)</b>	
<b>Format</b>	<QOSNAME>
<b>Description</b>	Displays job records matching the specified quality of service.
<b>Example</b>	<pre>&gt; showhist.moab.pl -q myQos</pre> <p><i>Information about jobs related to the QoS myQos is displayed.</i></p>

<b>-s (Start Date)</b>	
<b>Format</b>	YYYY-MM-DD
<b>Description</b>	Displays the records of jobs that recorded on the specified date and later.
<b>Example</b>	<pre>&gt; showhist.moab.pl -s 1776-07-04</pre> <p><i>Information about all jobs recorded on July 4, 1776 and later is displayed.</i></p> <pre>&gt; showhist.moab.pl -s 2020-07-05 -e 2022-07-05</pre> <p><i>Information is displayed about all jobs recorded between July 5, 2020 and July 5, 2022.</i></p>

<b>-u (User)</b>	
<b>Format</b>	<USERNAME>
<b>Description</b>	Displays job records matching the specified user.
<b>Example</b>	<pre>&gt; showhist.moab.pl -u bob</pre> <p><i>Information about user bob's jobs is displayed.</i></p>

## Sample Output

```
> showhist.moab.pl
```

```

Job Id           : Moab.4
User Name       : user1
Group Name      : company
Queue Name      : NONE
Processor Count : 4
Wallclock Duration: 00:00:00
Submit Time     : Mon Nov 21 10:48:32 2022
Start Time      : Mon Nov 21 10:49:37 2022
End Time        : Mon Nov 21 10:49:37 2022
Exit Code       : 0
Allocated Nodelist: 10.10.10.3

Job Id           : Moab.1
Executable      : 4
User Name       : user1
Group Name      : company
Account Name    : 1321897709
Queue Name      : NONE
Quality Of Service: 0M
Processor Count : -0
Wallclock Duration: 00:01:05
Submit Time     : Mon Nov 21 10:48:29 2022
Start Time      : Mon Nov 21 10:48:32 2022
End Time        : Mon Nov 21 10:49:37 2022
Exit Code       : 0
Allocated Nodelist: 512M

```

Information is displayed for all completed jobs.

**i** When a job's Start Time and End Time are the same, the job is infinite and still running.

## Related Topics

- [checkjob](#) - explains how to query for a status report for a specified job
- [mdiag -j](#) command - display additional detailed information regarding jobs
- [showq](#) command - showq high-level job summaries

## 3.7.35 showres

### 3.7.35.A Synopsis

```

showres [-f] [-n [-g]] [-o] [-r] [-v] [reservationid] [--blocking] [--about] [--help]
[--host=<serverHostName>] [--loglevel=<logLevel>] [--msg=<message>] [--port=<serverPort>]
[--timeout=<seconds>] [--version] [--xml]

```


### 3.7.35.B Overview

This command displays all reservations currently in place within Moab. The default behavior is to display reservations on a reservation-by-reservation basis.

### 3.7.35.C Access

By default, this command can be run by any Moab administrator.

### 3.7.35.D Options

Option	Description
<b>-f</b>	Show free (unreserved) resources rather than reserved resources. The <code>-f</code> flag cannot be used in conjunction with the any other flag.
<b>-g</b>	When used with the <code>-n</code> flag, shows <i>grep</i> -able output with nodename on every line.
<b>-n</b>	Display information regarding all nodes reserved by <code>&lt;RSVID&gt;</code>
<b>-o</b>	Display all reservations that overlap <code>&lt;RSVID&gt;</code> (in time and space). <div style="border: 1px solid black; border-radius: 5px; padding: 5px; margin-top: 5px;">  Not supported with <code>-n</code> flag         </div>
<b>-r</b>	Display reservation timeframes in relative time mode.
<b>-v</b>	Show verbose output. If used with the <code>-n</code> flag, the command will display all reservations found on nodes contained in <code>&lt;RSVID&gt;</code> . Otherwise, it will show long reservation start dates including the reservation year.

Parameter	Description
<b>RSVID</b>	ID of reservation of interest — optional

### 3.7.35.E Examples

```
> showres
ReservationID      Type S      Start      End      Duration  N/P      StartTime
12941              Job R      -00:05:01  2:41:39  2:46:40   13/25   Thu Sep 1
```



```

15:02:50
12944          Job R   -00:05:01      6:34:59      6:40:00      16/16      Thu Sep  1
15:02:50
12946          Job R   -00:05:01      1:05:54:59   1:06:00:00    1/2        Thu Sep  1
15:02:50
12954          Job R   -00:04:59      2:55:01      3:00:00      2/4        Thu Sep  1
15:02:52
12956          Job I   1:05:54:59    1:12:34:59   6:40:00      16/32      Fri Sep  2
21:02:50
12969          Job I       6:34:59      13:14:59     6:40:00      4/4        Thu Sep  1
21:42:50

6 reservations located

```

*The above example shows all reservations on the system.*

The fields are as follows:

Column	Description
<b>Type</b>	Reservation Type. This will be one of the following: Job or User.
<b>ReservationID</b>	This is the name of the reservation. Job reservation names are identical to the job name. User, Group, or Account reservations are the user, group, or account name followed by a number. System reservations are given the name SYSTEM followed by a number.
<b>S</b>	State. This field is valid only for job reservations. It indicates whether the job is (S)tarting, (R)unning, or (I)dle.
<b>Start</b>	Relative start time of the reservation. Time is displayed in HH:MM:SS notation and is relative to the present time.
<b>End</b>	Relative end time of the reservation. Time is displayed in HH:MM:SS notation and is relative to the present time. Reservations that will not complete in 1,000 hours are marked with the keyword INFINITY.
<b>Duration</b>	Duration of the reservation in HH:MM:SS notation. Reservations lasting more than 1,000 hours are marked with the keyword INFINITY.
<b>Nodes</b>	Number of nodes involved in reservation.
<b>StartTime</b>	Time Reservation became active.

*Example 3-41:*

```

> showres -n
reservations on Thu Sep  1 16:49:59

```

NodeName StartTime	Type	ReservationID	JobState	Task	Start	Duration	
G5-001	Job	12946	Running	2	-1:47:09	1:06:00:00	Thu
Sep 1 15:02:50							
G5-001	Job	12956	Idle	2	1:04:12:51	6:40:00	Fri
Sep 2 21:02:50							
G5-002	Job	12956	Idle	2	1:04:12:51	6:40:00	Fri
Sep 2 21:02:50							
G5-002	Job	12953	Running	2	-00:29:37	4:26:40	Thu
Sep 1 16:20:22							
G5-003	Job	12956	Idle	2	1:04:12:51	6:40:00	Fri
Sep 2 21:02:50							
G5-003	Job	12953	Running	2	-00:29:37	4:26:40	Thu
Sep 1 16:20:22							
G5-004	Job	12956	Idle	2	1:04:12:51	6:40:00	Fri
Sep 2 21:02:50							
G5-004	Job	12953	Running	2	-00:29:37	4:26:40	Thu
Sep 1 16:20:22							
G5-005	Job	12956	Idle	2	1:04:12:51	6:40:00	Fri
Sep 2 21:02:50							
G5-005	Job	12953	Running	2	-00:29:37	4:26:40	Thu
Sep 1 16:20:22							
G5-006	Job	12956	Idle	2	1:04:12:51	6:40:00	Fri
Sep 2 21:02:50							
G5-006	Job	12953	Running	2	-00:29:37	4:26:40	Thu
Sep 1 16:20:22							
G5-007	Job	12956	Idle	2	1:04:12:51	6:40:00	Fri
Sep 2 21:02:50							
G5-007	Job	12939	Running	2	-00:29:37	3:00:00	Thu
Sep 1 16:20:22							
G5-008	Job	12956	Idle	2	1:04:12:51	6:40:00	Fri
Sep 2 21:02:50							
G5-008	Job	12939	Running	2	-00:29:37	3:00:00	Thu
Sep 1 16:20:22							
G5-009	Job	12956	Idle	2	1:04:12:51	6:40:00	Fri
Sep 2 21:02:50							
G5-009	Job	12939	Running	2	-00:29:37	3:00:00	Thu
Sep 1 16:20:22							
G5-010	Job	12956	Idle	2	1:04:12:51	6:40:00	Fri
Sep 2 21:02:50							
G5-010	Job	12939	Running	2	-00:29:37	3:00:00	Thu
Sep 1 16:20:22							
G5-011	Job	12956	Idle	2	1:04:12:51	6:40:00	Fri
Sep 2 21:02:50							
G5-011	Job	12939	Running	2	-00:29:37	3:00:00	Thu
Sep 1 16:20:22							
G5-012	Job	12956	Idle	2	1:04:12:51	6:40:00	Fri
Sep 2 21:02:50							
G5-012	Job	12939	Running	2	-00:29:37	3:00:00	Thu
Sep 1 16:20:22							
G5-013	Job	12956	Idle	2	1:04:12:51	6:40:00	Fri
Sep 2 21:02:50							
G5-013	Job	12939	Running	2	-00:29:37	3:00:00	Thu
Sep 1 16:20:22							
G5-014	Job	12956	Idle	2	1:04:12:51	6:40:00	Fri
Sep 2 21:02:50							
G5-014	Job	12939	Running	2	-00:29:37	3:00:00	Thu
Sep 1 16:20:22							
G5-015	Job	12956	Idle	2	1:04:12:51	6:40:00	Fri
Sep 2 21:02:50							
G5-015	Job	12949	Running	2	-00:08:57	3:00:00	Thu

```

Sep  1 16:41:02
G5-016      Job          12956      Idle      2  1:04:12:51  6:40:00  Fri
Sep  2 21:02:50
G5-016      Job          12947      Running   2  -00:08:57   3:00:00  Thu
Sep  1 16:41:02
P690-001    Job          12944      Running   1  -1:47:09    6:40:00  Thu
Sep  1 15:02:50
P690-002    Job          12944      Running   1  -1:47:09    6:40:00  Thu
Sep  1 15:02:50
P690-003    Job          12944      Running   1  -1:47:09    6:40:00  Thu
Sep  1 15:02:50
P690-004    Job          12944      Running   1  -1:47:09    6:40:00  Thu
Sep  1 15:02:50
P690-005    Job          12944      Running   1  -1:47:09    6:40:00  Thu
Sep  1 15:02:50
P690-006    Job          12944      Running   1  -1:47:09    6:40:00  Thu
Sep  1 15:02:50
P690-007    Job          12944      Running   1  -1:47:09    6:40:00  Thu
Sep  1 15:02:50
P690-008    Job          12944      Running   1  -1:47:09    6:40:00  Thu
Sep  1 15:02:50
P690-009    Job          12944      Running   1  -1:47:09    6:40:00  Thu
Sep  1 15:02:50
P690-010    Job          12944      Running   1  -1:47:09    6:40:00  Thu
Sep  1 15:02:50
P690-011    Job          12944      Running   1  -1:47:09    6:40:00  Thu
Sep  1 15:02:50
P690-012    Job          12944      Running   1  -1:47:09    6:40:00  Thu
Sep  1 15:02:50
P690-013    Job          12944      Running   1  -1:47:09    6:40:00  Thu
Sep  1 15:02:50
P690-013    Job          12969      Idle      1  4:52:51     6:40:00  Thu
Sep  1 21:42:50
P690-014    Job          12944      Running   1  -1:47:09    6:40:00  Thu
Sep  1 15:02:50
P690-014    Job          12969      Idle      1  4:52:51     6:40:00  Thu
Sep  1 21:42:50
P690-015    Job          12944      Running   1  -1:47:09    6:40:00  Thu
Sep  1 15:02:50
P690-015    Job          12969      Idle      1  4:52:51     6:40:00  Thu
Sep  1 21:42:50
P690-016    Job          12944      Running   1  -1:47:09    6:40:00  Thu
Sep  1 15:02:50
P690-016    Job          12969      Idle      1  4:52:51     6:40:00  Thu
Sep  1 21:42:50

```

52 nodes reserved

*This example shows reservations for nodes.*

The fields are as follows:

Column	Description
<b>NodeName</b>	Node on which reservation is placed.
<b>Type</b>	Reservation Type. This will be one of the following: Job or User.

Column	Description
<b>ReservationID</b>	This is the name of the reservation. Job reservation names are identical to the job name. User, Group, or Account reservations are the user, group, or account name followed by a number. System reservations are given the name SYSTEM followed by a number.
<b>JobState</b>	This field is valid only for job reservations. It indicates the state of the job associated with the reservation.
<b>Start</b>	Relative start time of the reservation. Time is displayed in HH:MM:SS notation and is relative to the present time.
<b>Duration</b>	Duration of the reservation in HH:MM:SS notation. Reservations lasting more than 1000 hours are marked with the keyword INFINITY.
<b>StartTime</b>	Time Reservation became active.

*Example 3-42:*

```

> showres 12956
ReservationID      Type S      Start      End      Duration  N/P      StartTime
12956             Job I      1:04:09:32 1:10:49:32 6:40:00   16/32    Fri Sep 2
21:02:50
1 reservation located

```

*In this example, information for a specific reservation (job) is displayed.*

**Related Topics**

- (Optional) Install Moab Client - in the *Moab HPC Suite Installation and Configuration Guide* explains how to distribute this command to client nodes
- [mrsvctl -c](#) - create new reservations
- [mrsvctl -r](#) - release existing reservations
- [mdiag -r](#) - diagnose/view the state of existing reservations
- [Reservation Overview](#) - description of reservations and their use


## 3.7.36 showstart


### 3.7.36.A Synopsis


```
showstart {jobid|proccount[@duration]|s3jobspec} [-e {all|hist|prio|rsv}] [-f] [-g [peer]]
[-l qos=<QOS>] [--blocking] [--format=xml] [-v] [--about] [--help] [--
host=<serverHostName>] [--loglevel=<logLevel>] [--
msg=<message>] [--port=<serverPort>] [--timeout=<seconds>] [--
version] [--xml]
```

### 3.7.36.B Overview

This command displays the estimated start time of a job based a number of analysis types. This analysis may include information based on historical usage, earliest available reservable resources, and priority based backlog analysis. Each type of analysis will provide somewhat different estimates based on current cluster environmental conditions. The default estimation method used is determined by the value of the `DEFAULTSTARTTIMEQUERY` parameter, which defaults to `PRIORITY`.

 **showstart** is a processor-intensive command. Multiple submissions per iteration may slow Moab's scheduling, especially on larger/busier systems.

 The start time estimate Moab returns does not account for resources that will become available due to preemption.

 **showstart** *only* determines where a job would run if it were to run next, taking into account all currently running jobs, queued idle jobs with advance reservations, and all current standing and administrative reservations in the system.

For example, assume `RESERVATIONDEPTH` is set to 1 (the default value), job 12300 is at the top of the idle queue and has an advance reservation to run next, and job 12312 is in position 12 in the idle queue. If the owner of job 12312 runs `showstart 12312`, in calculating where the job will run, Moab does not consider jobs 12301-12311. It only estimates where and when job 12312 would be scheduled to run after job 12300 starts.

**Historical** analysis utilizes historical queue times for jobs that match a similar processor count and job duration profile. This information is updated on a sliding window, which is configurable within `moab.cfg`

**Reservation** based start time estimation incorporates information regarding current administrative, user, and job reservations to determine the earliest time the specified job could allocate the needed resources and start running. In essence, this estimate will indicate the earliest time the job would start assuming this job was the highest priority job in the queue.

**Priority** based job start analysis determines when the queried job would fit in the queue and determines the estimated amount of time required to complete the jobs that are currently running or scheduled to run before this job can start.

In all cases, if the job is running, this command will return the time the job started. If the job already has a reservation, this command will return the start time of the reservation.

### 3.7.36.C Access

By default, this command can be run by any user.

### 3.7.36.D Parameters

Parameters	Description
<b>DURATION</b>	Duration of pseudo-job to be checked in format [[[DD:]HH:]MM:]SS (default duration is 1 second).
<b>QOS</b>	Specifies what QOS the job must start under, using the same syntax as the msub command. Currently, no other resource manager extensions are supported. This flag only applies to hypothetical jobs by using the proccount [@duration] syntax.
<b>JOBID</b>	Job to be checked.
<b>PROCCOUNT</b>	Number of processors in pseudo-job to be checked.
<b>S3JOBSPEC</b>	XML describing the job according to the Dept. of Energy <a href="#">Scalable Systems Software/S3</a> job specification.

## 3.7.36.E Options

Parameters	Description
<b>--blocking</b>	Do not use cache information in the output. The <code>--blocking</code> flag retrieves results exclusively from the scheduler.
<b>-e</b>	Estimate method. If not specified, Moab will use the value of the <code>DEFAULTSTARTTIMEQUERY</code> parameter, which defaults to <code>PRIORITY</code> .
<b>-f</b>	Use feedback. If specified, Moab will apply historical accuracy information to improve the quality of the estimate.
<b>-g</b>	Grid mode. Obtain showstart information from remote resource managers. If <code>-g</code> is not used and Moab determines that job is already migrated, Moab obtains showstart information from the remote Moab where the job was migrated to. All resource managers can be queried by using the keyword 'all', which returns all information in a table.  <pre style="border: 1px dashed black; padding: 5px;">\$ showstart -g all head.1 Estimated Start Times [ Remote RM ] [ Reservation ] [ Priority ] [ Historical ] [ c1 ] [ 00:15:35 ] [ ] [ ] [ c2 ] [ 3:15:38 ] [ ] [ ]</pre>
<b>-l qos=&lt;QOS&gt;</b>	Specifies what QOS the job must start under, using the same syntax as the <code>msub</code> command. Currently, no other resource manager extensions are supported. This flag only applies to hypothetical jobs by using the <code>proccount[@duration]</code> syntax.
<b>-v</b>	Displays verbose information.

## 3.7.36.F Examples

```
> showstart orion.13762
job orion.13762 requires 2 procs for 0:33:20
Estimated Rsv based start in          1:04:55 on Fri Jul 15 12:53:40
Estimated Rsv based completion in     2:44:55 on Fri Jul 15 14:33:40
Estimated Priority based start in      5:14:55 on Fri Jul 15 17:03:40
Estimated Priority based completion in  6:54:55 on Fri Jul 15 18:43:40
Estimated Historical based start in    00:00:00 on Fri Jul 15 11:48:45
Estimated Historical based completion in 1:40:00 on Fri Jul 15 13:28:45
Best Partition: fast
```

**Example 3-43:**

```
> showstart 12@3600
job 12@3600 requires 12 procs for 1:00:00
```

```

Earliest start in      00:01:39 on Wed Aug 31 16:30:45
Earliest completion in 1:01:39 on Wed Aug 31 17:30:45
Best Partition: 32Bit

```

**i** You cannot specify job flags when running `showstart`, and since a job by default can only run on one partition, `showstart` fails when querying for a job requiring more nodes than the largest partition available.

### 3.7.36.G Additional Information

For reservation based estimates, the information provided by this command is more highly accurate if the job is highest priority, if the job has a reservation, or if the majority of the jobs that are of higher priority have reservations. Consequently, sites wanting to make decisions based on this information might want to consider using the `RESERVATIONDEPTH` parameter to increase the number of priority based reservations. This can be set so that most or even all idle jobs receive priority reservations and make the results of this command generally useful. The only caution of this approach is that increasing the `RESERVATIONDEPTH` parameter more tightly constrains the decisions of the scheduler and might result in slightly lower system utilization (typically less than 8% reduction).

#### Related Topics

- (Optional) Install Moab Client - in the *Moab HPC Suite Installation and Configuration Guide* explains how to distribute this command to client nodes
- [checkjob](#)
- [showres](#)
- [showstats -f eststarttime](#)
- [showstats -f avgqtime](#)
- [Job Start Estimates](#)

## 3.7.37 showstate

### 3.7.37.A Synopsis

```

showstate [--about] [--help] [--host=<serverHostName>] [--
loglevel=<logLevel>] [--msg=<message>] [--port=<serverPort>]
[--timeout=<seconds>] [--version] [--xml]

```



### 3.7.37.B Overview

This command provides a summary of the state of the system. It displays a list of all active jobs and a text-based map of the status of all nodes and the jobs they are servicing. Basic diagnostic tests are also performed and any problems found are reported.

### 3.7.37.C Access

By default, this command can be run by any Moab Administrator.

### 3.7.37.D Examples

```
> showstate
cluster state summary for Wed Nov 23 12:00:21
```

	JobID	S	User	Group	Procs	Remaining	StartTime
(A)	fr17n11.942.0	R	johns	staff	16	13:21:15	Nov 22 12:00:21
(B)	fr17n12.942.0	S	johns	staff	32	13:07:11	Nov 22 12:00:21
(C)	fr17n13.942.0	R	johns	staff	8	11:22:25	Nov 22 12:00:21
(D)	fr17n14.942.0	S	johns	staff	8	10:43:43	Nov 22 12:01:21
(E)	fr17n15.942.0	S	johns	staff	8	9:19:25	Nov 22 12:01:21
(F)	fr17n16.942.0	R	johns	staff	8	9:01:16	Nov 22 12:01:21
(G)	fr17n17.942.0	R	johns	staff	1	7:28:25	Nov 22 12:03:22
(H)	fr17n18.942.0	R	johns	staff	1	3:05:17	Nov 22 12:04:22
(I)	fr17n19.942.0	S	johns	staff	24	0:54:38	Nov 22 12:00:22

```
Usage Summary: 9 Active Jobs 106 Active Nodes
```

```

[0] [0] [0] [0] [0] [0] [0] [0] [0] [0] [1] [1] [1] [1] [1] [1] [1] [1]
[1] [2] [3] [4] [5] [6] [7] [8] [9] [0] [1] [2] [3] [4] [5] [6]
Frame 2: XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX[ ] [A] [C] [ ] [A] [C] [C] [A]
Frame 3: [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [A] [ ] [I] [ ] [ ] [I] [ ] [ ] [ ] [ ] [ ]
Frame 4: [ ] [ ] [I] [ ] [ ] [ ] [ ] [ ] [A] [ ] [ ] [I] [ ] [ ] [ ] [ ] [E] [ ] [I] [ ] [E]
Frame 5: [F] [ ] [E] [ ] [ ] [ ] [ ] [ ] [F] [F] [F] [I] [ ] [ ] [ ] [E] [ ] [E] [E]
Frame 6: [ ] [ ] [I] [I] [E] [I] [ ] [ ] [I] [I] [ ] [I] [F] [I] [I] [I] [I] [F]
Frame 7: [ ] XXXX[ ] XXXX[ ] XXXX[ ] XXXX[b] XXXX[ ] XXXX[ ] XXXX[#] XXXX
Frame 9: [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [E] [ ]
Frame 11: [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [I] [F] [E] [ ] [ ] [A] [I] [ ] [ ] [F] [ ] [A]
Frame 12: [A] [ ] [ ] [ ] [A] [ ] [ ] [ ] [C] [A] [ ] [ ] [C] [A] [A] [ ] [ ] [ ] [ ] [ ]
Frame 13: [D] XXXX[I] XXXX[ ] XXXX[ ] XXXX[ ] XXXX[ ] XXXX[I] XXXX[I] XXXX
Frame 14: [D] XXXX[I] XXXX[I] XXXX[D] XXXX[ ] XXXX[H] XXXX[I] XXXX[ ] XXXX
Frame 15: [b] XXXX[b] XXXX[b] XXXX[b] XXXX[D] XXXX[b] XXXX[b] XXXX[b] XXXX
Frame 16: [b] XXXX[ ] XXXX[b] XXXX[ ] XXXX[b] XXXX[b] XXXX[ ] XXXX[b] XXXX
Frame 17: [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ] [ ]
Frame 21: [ ] XXXX[b] XXXX[b] XXXX[ ] XXXX[b] XXXX[b] XXXX[b] XXXX[b] XXXX
Frame 22: [b] XXXX[b] XXXX[b] XXXX[ ] XXXX[b] XXXX[ ] XXXX[b] XXXX[b] XXXX
Frame 27: [b] XXXX[b] XXXX[ ] XXXX[b] XXXX[b] XXXX[b] XXXX[b] XXXX[b] XXXX
Frame 28: [G] XXXX[ ] XXXX[D] XXXX[ ] XXXX[D] XXXX[D] XXXX[D] XXXX[ ] XXXX
Frame 29: [A] [C] [A] [A] [C] [ ] [A] [C] XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
Key: XXX:Unknown [*]:Down w/Job [#]:Down [']:Idle w/Job [ ]:Idle [@]:Busy w/No Job
[!]:Drained
Key: [a]:(Any lower case letter indicates an idle node that is assigned to a job)

Check Memory on Node fr3n07
Check Memory on Node fr4n06
Check Memory on Node fr4n09
```

*In this example, nine active jobs are running on the system. Each job listed in the top of the output is associated with a letter. For example, job fr17n11.942.0 is associated with the letter A. This letter can now be used to determine*

where the job is currently running. By looking at the system map, it can be found that job `fr17n11.942.0` (job A) is running on nodes `fr2n10`, `fr2n13`, `fr2n16`, `fr3n07` ...

The key at the bottom of the system map can be used to determine unusual node states. For example, `fr7n15` is currently in the state down.

After the key, a series of warning messages may be displayed indicating possible system problems. In this case, warning messages indicate that there are memory problems on three nodes, `fr3n07`, `fr4n06`, and `fr4n09`. Also, warning messages indicate that job `fr15n09.1097.0` is having difficulty starting. Node `fr11n08` is in state `BUSY` but has no job assigned to it (it possibly has a runaway job running on it).

## Related Topics

- (Optional) Install Moab Client - in the *Moab HPC Suite Installation and Configuration Guide* explains how to distribute this command to client nodes
- [Specifying Node Rack/Slot Location](#)

## 3.7.38 showstats

### 3.7.38.A Synopsis

`showstats`

`showstats -a` [*accountid*] [-v] [-t <TIMESPEC>]

`showstats -c` [*classid*] [-v] [-t <TIMESPEC>]

`showstats -f` <statistictype>

`showstats -g` [*groupid*] [-v] [-t <TIMESPEC>]

`showstats -j` [*jobtemplate*] [-t <TIMESPEC>]

`showstats -n` [*nodeid*] [-t <TIMESPEC>]

`showstats -q` [*qosid*] [-v] [-t <TIMESPEC>]

`showstats -s`

`showstats -T` [*leafid* | *tree-level*]

`showstats -u` [*userid*] [-v] [-t <TIMESPEC>]

[--about] [--help] [--host=<serverHostName>] [--loglevel=<logLevel>] [--msg=<message>] [--port=<serverPort>] [--timeout=<seconds>] [--version] [--xml]

### 3.7.38.B Overview

This command shows various accounting and resource usage statistics for the system. Historical statistics cover the timeframe from the most recent execution of the `mschedctl -f` command.

### 3.7.38.C Access

By default, this command can be run by any Moab level 1, 2, or 3 Administrator.

### 3.7.38.D Options

Option	Description
<code>-a[&lt;ACCOUNTID&gt;]</code>	Display account statistics. See <a href="#">Account Statistics</a> for an example.
<code>-c[&lt;CLASSID&gt;]</code>	Display class statistics.
<code>-f &lt;statistictype&gt;</code>	Display full matrix statistics (see <a href="#">showstats -f</a> for full details).
<code>-g[&lt;GROUPID&gt;]</code>	Display group statistics. See <a href="#">Group Statistics</a> for an example.
<code>-j[&lt;JOBTEMPLATE&gt;]</code>	Display template statistics.
<code>-n[&lt;NODEID&gt;]</code>	Display node statistics ( <code>ENABLEPROFILING</code> must be set). See <a href="#">Node Statistics</a> for an example.
<code>-q [ &lt;QOSID&gt;]</code>	Display QoS statistics.
<code>-s</code>	Display general scheduler statistics.
<code>-t</code>	Display statistical information from the specified timeframe: <div style="border: 1px dashed black; padding: 5px; margin: 5px 0;"> <pre>&lt;START_TIME&gt;[, &lt;END_TIME&gt;]   (ABSTIME: [HH[:MM[:SS]]][_MO[/DD[/YY]]] ie 14:30_06/20)   (RELTIME: -[[[DD:]HH:]MM:]SS)</pre> </div> See <a href="#">Statistics from an Absolute Time Frame</a> and <a href="#">Statistics from a Relative Time Frame</a> for examples.

Option	Description
	<div style="border: 1px solid black; padding: 5px;"> <p><b>i</b> Profiling must be enabled for the credential type you want statistics for. See <a href="#">Credential Statistics</a> for information on how to enable profiling. Also, <b>-t</b> is not a stand-alone option. It must be used in conjunction with the <b>-a</b>, <b>-c</b>, <b>-g</b>, <b>-n</b>, <b>-q</b>, or <b>-u</b> flag.</p> </div>
<b>-T</b>	Display fairshare tree statistics. See <a href="#">Fairshare Tree Statistics</a> for an example.
<b>-u[&lt;USERID&gt;]</b>	Display user statistics. See <a href="#">User Statistics</a> for an example.
<b>-v</b>	Display verbose information. See <a href="#">Verbose Statistics</a> for an example.

### 3.7.38.E Examples

#### Example 3-44: Account Statistics


```

> showstats -a
Account Statistics Initialized Tue Aug 26 14:32:39
      |----- Running -----|----- Completed -----
      |-----|
Account  Jobs Procs ProcHours  Jobs  %  PHReq  %  PHDed  %  FSTgt  AvgXF
MaxXF  AvgQH  Effic  WCAcc
137651  16  92  1394.52  229  39.15  18486  45.26  7003.5  41.54  40.00  0.77
8.15  5.21  90.70  34.69
462212  11  63  855.27  43  7.35  6028  14.76  3448.4  20.45  6.25  0.71
5.40  3.14  98.64  40.83
462213  6  72  728.12  90  15.38  5974  14.63  3170.7  18.81  6.25  0.37
4.88  0.52  82.01  24.14
005810  3  24  220.72  77  13.16  2537  6.21  1526.6  9.06  -----  1.53
14.81  0.42  98.73  28.40
175436  0  0  0.00  12  2.05  6013  14.72  958.6  5.69  2.50  1.78
8.61  5.60  83.64  17.04
000102  0  0  0.00  1  0.17  64  0.16  5.1  0.03  -----  10.85
10.85  10.77  27.90  7.40
000023  0  0  0.00  1  0.17  12  0.03  0.2  0.00  -----  0.04
0.04  0.19  21.21  1.20
    
```

*This example shows a statistical listing of all active accounts. The top line (Account Statistics Initialized...) of the output indicates the beginning of the timeframe covered by the displayed statistics. The statistical output is divided into two categories, Running and Completed. Running statistics include information about jobs that are currently running. Completed statistics are compiled using historical information from both running and completed jobs.*

The fields are as follows:

Column	Description
<b>Account</b>	Account Number.
<b>Jobs</b>	Number of running jobs.
<b>Procs</b>	Number of processors allocated to running jobs.
<b>ProcHours</b>	Number of proc-hours required to complete running jobs.
<b>Jobs*</b>	Number of jobs completed.
<b>%</b>	Percentage of total jobs that were completed by account.
<b>PHReq*</b>	Total proc-hours requested by completed jobs.
<b>%</b>	Percentage of total proc-hours requested by completed jobs that were requested by account.
<b>PHDed</b>	Total proc-hours dedicated to active and completed jobs. The proc-hours dedicated to a job are calculated by multiplying the number of allocated procs by the length of time the procs were allocated, regardless of the job's CPU usage.
<b>%</b>	Percentage of total proc-hours dedicated that were dedicated by account.
<b>FSTgt</b>	Fairshare target. An account's fairshare target is specified in the <code>fs.cfg</code> file. This value should be compared to the account's node-hour dedicated percentage to determine if the target is being met.
<b>AvgXF*</b>	Average expansion factor for jobs completed. A job's XFactor (expansion factor) is calculated by the following formula: $(\text{QueuedTime} + \text{RunTime}) / \text{WallClockLimit}$ .
<b>MaxXF*</b>	Highest expansion factor received by jobs completed.
<b>AvgQH*</b>	Average queue time (in hours) of jobs.
<b>Effic</b>	Average job efficiency. Job efficiency is calculated by dividing the actual node-hours of CPU time used by the job by the node-hours allocated to the job.
<b>WCAcc*</b>	Average wallclock accuracy for jobs completed. Wallclock accuracy is calculated by dividing a job's actual run time by its specified wallclock limit.

Column	Description
	 A job's wallclock accuracy is capped at 100% so even if a job exceeds its requested walltime it will report an accuracy of 100%.

\* These fields are empty until an account has completed at least one job.

**Example 3-45: Group Statistics**

```


> showstats -g
Group Statistics Initialized Tue Aug 26 14:32:39
      |----- Running -----|----- Completed -----
      |-----|
GroupName  GID  Jobs  Procs  ProcHours  Jobs  %  PHReq  %  PHDed  %  FSTgt
AvgXF  MaxXF  AvgQH  Effic  WCacc
univ  214  16  92  1394.52  229  39.15  18486  45.26  7003.5  41.54  40.00
0.77  8.15  5.21  90.70  34.69
daf  204  11  63  855.27  43  7.35  6028  14.76  3448.4  20.45  6.25
0.71  5.40  3.14  98.64  40.83
dnavy 207  6  72  728.12  90  15.38  5974  14.63  3170.7  18.81  6.25
0.37  4.88  0.52  82.01  24.14
govt  232  3  24  220.72  77  13.16  2537  6.21  1526.6  9.06  -----
1.53  14.81  0.42  98.73  28.40
asp  227  0  0  0.00  12  2.05  6013  14.72  958.6  5.69  2.50
1.78  8.61  5.60  83.64  17.04
derim 229  0  0  0.00  74  12.65  669  1.64  352.5  2.09  -----
0.50  1.93  0.51  96.03  32.60
dchall 274  0  0  0.00  3  0.51  447  1.10  169.2  1.00  25.00
0.52  0.88  2.49  95.82  33.67
nih  239  0  0  0.00  17  2.91  170  0.42  148.1  0.88  -----
0.95  1.83  0.14  97.59  84.31
darmy 205  0  0  0.00  31  5.30  366  0.90  53.9  0.32  6.25
0.14  0.59  0.07  81.33  12.73
systems 80  0  0  0.00  6  1.03  67  0.16  22.4  0.13  -----
4.07  8.49  1.23  28.68  37.34
pdc  252  0  0  0.00  1  0.17  64  0.16  5.1  0.03  -----
10.85 10.85 10.77 27.90 7.40
staff  1  0  0  0.00  1  0.17  12  0.03  0.2  0.00  -----
0.04  0.04  0.19 21.21 1.20
    
```

*This example shows a statistical listing of all active groups. The top line (Group Statistics Initialized...) of the output indicates the beginning of the timeframe covered by the displayed statistics. The statistical output is divided into two categories, Running and Completed. Running statistics include information about jobs that are currently running. Completed statistics are compiled using historical information from both running and completed jobs.*

The fields are as follows:

Column	Description
<b>GroupName</b>	Name of group.

Column	Description
<b>GID</b>	Group ID of group.
<b>Jobs</b>	Number of running jobs.
<b>Procs</b>	Number of procs allocated to running jobs.
<b>ProcHours</b>	Number of proc hours required to complete running jobs.
<b>Jobs*</b>	Number of jobs completed.
<b>%</b>	Percentage of total jobs that were completed by group.
<b>PHReq*</b>	Total proc-hours requested by completed jobs.
<b>%</b>	Percentage of total proc-hours requested by completed jobs that were requested by group.
<b>PHDed</b>	Total proc-hours dedicated to active and completed jobs. The proc-hours dedicated to a job are calculated by multiplying the number of allocated procs by the length of time the procs were allocated, regardless of the job's CPU usage.
<b>%</b>	Percentage of total proc-hours dedicated that were dedicated by group.
<b>FSTgt</b>	Fairshare target. A group's fairshare target is specified in the <code>fs.cfg</code> file. This value should be compared to the group's node-hour dedicated percentage to determine if the target is being met.
<b>AvgXF*</b>	Average expansion factor for jobs completed. A job's XFactor (expansion factor) is calculated by the following formula: $(\text{QueuedTime} + \text{RunTime}) / \text{WallClockLimit}$ .
<b>MaxXF*</b>	Highest expansion factor received by jobs completed.
<b>AvgQH*</b>	Average queue time (in hours) of jobs.
<b>Effic</b>	Average job efficiency. Job efficiency is calculated by dividing the actual node-hours of CPU time used by the job by the node-hours allocated to the job.
<b>WCAcc*</b>	Average wallclock accuracy for jobs completed. Wallclock accuracy is calculated by dividing a job's actual run time by its specified wallclock limit.

Column	Description
	 A job's wallclock accuracy is capped at 100% so even if a job exceeds its requested walltime it will report an accuracy of 100%.

\* These fields are empty until a group has completed at least one job.

**Example 3-46: Node Statistics**

```

> showstats -n
node stats from Mon Jul 10 00:00:00 to Mon Jul 10 16:30:00
node      CfgMem MinMem MaxMem AvgMem | CfgProcs MinLoad MaxLoad AvgLoad
node01    58368    0 21122  5841    32   0.00  32.76  27.62
node02   122880    0 19466   220    30   0.00  33.98  29.54
node03   18432    0  9533  2135    24   0.00  25.10  18.64
node04   60440    0 17531  4468    32   0.00  30.55  24.61
node05   13312    0  2597  1189     8   0.00   9.85   8.45
node06   13312    0  3800  1112     8   0.00   8.66   5.27
node07   13312    0  2179  1210     8   0.00   9.62   8.27
node08   13312    0  3243  1995     8   0.00  11.71   8.02
node09   13312    0  2287  1943     8   0.00  10.26   7.58
node10   13312    0  2183  1505     8   0.00  13.12   9.28
node11   13312    0  3269  2448     8   0.00   8.93   6.71
node12   13312    0 10114  6900     8   0.00  13.13   8.44
node13   13312    0  2616  2501     8   0.00   9.24   8.21
node14   13312    0  3888   869     8   0.00   8.10   3.85
node15   13312    0  3788   308     8   0.00   8.40   4.67
node16   13312    0  4386  2191     7   0.00  18.37   8.36
node17   13312    0  3158  1870     8   0.00   8.95   5.91
node18   13312    0  5022  2397     8   0.00  19.25   8.19
node19   13312    0  2437  1371     8   0.00   8.98   7.09
node20   13312    0  4474  2486     8   0.00   8.51   7.11
node21   13312    0  4111  2056     8   0.00   8.93   6.68
node22   13312    0  5136  2313     8   0.00   8.61   5.75
node23   13312    0  1850  1752     8   0.00   8.39   5.71
node24   13312    0  3850  2539     8   0.00   8.94   7.80
node25   13312    0  3789  3702     8   0.00  21.22  12.83
node26   13312    0  3809  1653     8   0.00   9.34   4.91
node27   13312    0  5637    70     4   0.00  17.97   2.46
node28   13312    0  3076  2864     8   0.00  22.91  10.33
    
```

**Example 3-47: Verbose Statistics**

```

> showstats -v
current scheduler time: Sat Aug 18 18:23:02 2007
moab active for      00:00:01  started on Wed Dec 31 17:00:00
statistics for iteration 0  initialized on Sat Aug 11 23:55:25
Eligible/Idle Jobs:      6/8      (75.000%)
Active Jobs:              13
Successful/Completed Jobs: 167/167  (100.000%)
Preempt Jobs:             0
Avg/Max QTime (Hours):   0.34/2.07
    
```



```

Avg/Max XFactor:          1.165/3.26
Avg/Max Bypass:          0.40/8.00
Dedicated/Total ProcHours: 4.46K/4.47K (99.789%)
Preempt/Dedicated ProcHours: 0.00/4.46K (0.000%)
Current Active/Total Procs: 32/32 (100.0%)
Current Active/Total Nodes: 16/16 (100.0%)
Avg WallClock Accuracy:  64.919%
Avg Job Proc Efficiency:  99.683%
Min System Utilization:  87.323% (on iteration 46)
Est/Avg Backlog:         02:14:06/03:02:567

```

*This example shows a concise summary of the system scheduling state. Note that `showstats` and `showstats -s` are equivalent.*

*The first line of output indicates the number of scheduling iterations performed by the current scheduling process, followed by the time the scheduler started. The second line indicates the amount of time the Moab Scheduler has been scheduling in HH:MM:SS notation followed by the statistics initialization time.*

The fields are as follows:

Column	Description
<b>Active Jobs</b>	Number of jobs currently active (Running or Starting).
<b>Eligible Jobs</b>	Number of jobs in the system queue (jobs that are considered when scheduling).
<b>Idle Jobs</b>	Number of jobs both in and out of the system queue that are in the LoadLeveler Idle state.
<b>Completed Jobs</b>	Number of jobs completed since statistics were initialized.
<b>Successful Jobs</b>	Jobs that completed successfully without abnormal termination.
<b>XFactor</b>	Average expansion factor of all completed jobs.
<b>Max XFactor</b>	Maximum expansion factor of completed jobs.
<b>Max Bypass</b>	Maximum bypass of completed jobs.
<b>Available ProcHours</b>	Total proc-hours available to the scheduler.
<b>Dedicated ProcHours</b>	Total proc-hours made available to jobs.
<b>Effic</b>	Scheduling efficiency (DedicatedProcHours / Available ProcHours).

Column	Description
<b>Min Efficiency</b>	Minimum scheduling efficiency obtained since scheduler was started.
<b>Iteration</b>	Iteration on which the minimum scheduling efficiency occurred.
<b>Available Procs</b>	Number of procs currently available.
<b>Busy Procs</b>	Number of procs currently busy.
<b>Effic</b>	Current system efficiency (BusyProcs/AvailableProcs).
<b>WallClock Accuracy</b>	Average wallclock accuracy of completed jobs (job-weighted average).
<b>Job Efficiency</b>	Average job efficiency (UtilizedTime / DedicatedTime).
<b>Est Backlog</b>	Estimated backlog of queued work in hours.
<b>Avg Backlog</b>	Average backlog of queued work in hours.

**Example 3-48: User Statistics**

```

> showstats -u
User Statistics Initialized Tue Aug 26 14:32:39
----- Running -----|----- Completed -----
-----|-----
UserName  UID Jobs Procs ProcHours Jobs  %  PHReq  %  PHDed  %  FSTgt
AvgXF  MaxXF  AvgQH  Effic  WCAcc
moorej  2617   1   16   58.80   2  0.34  221  0.54 1896.6 11.25 -----
1.02   1.04  0.14  99.52 100.00
zhong   1767   3   24   220.72  20  3.42 2306  5.65 1511.3  8.96 -----
0.71   0.96  0.49  99.37  67.48
lui     2467   0    0    0.00   16  2.74 1970  4.82 1505.1  8.93 -----
1.02   6.33  0.25  98.96  57.72
evans   3092   0    0    0.00   62 10.60 4960 12.14 1464.3  8.69   5.0
0.62   1.64  5.04  87.64  30.62
wengel  2430   2   64   824.90   1  0.17  767  1.88  630.3  3.74 -----
0.18   0.18  4.26  99.63  0.40
mukho   2961   2   16   71.06   6  1.03  776  1.90  563.5  3.34 -----
0.31   0.82  0.20  93.15  30.28
jimenez 1449   1   16   302.29   2  0.34  768  1.88  458.3  2.72 -----
0.80   0.98  2.31  97.99  70.30
neff    3194   0    0    0.00   74 12.65  669  1.64  352.5  2.09 10.0
0.50   1.93  0.51  96.03  32.60
cholik  1303   0    0    0.00   2  0.34  552  1.35  281.9  1.67 -----
1.72   3.07 25.35  99.69  66.70
jshoema 2508   1   24   572.22   1  0.17  576  1.41  229.1  1.36 -----

```

```

0.55  0.55  3.74  99.20  39.20
      kudo 2324    1    8    163.35    6    1.03  1152    2.82  211.1    1.25  -----
0.12  0.34  1.54  96.77  5.67
      xztang 1835    1    8    18.99  -----  -----  -----  176.3    1.05  10.0  -----
-----  99.62  -----
      feller 1880    0    0    0.00   17    2.91   170    0.42  148.1    0.88  -----
0.95  1.83  0.14  97.59  84.31
      maxia 2936    0    0    0.00    1    0.17   191    0.47  129.1    0.77   7.5
0.88  0.88  4.49  99.84  69.10
      ktgnov71 2838    0    0    0.00    1    0.17   192    0.47   95.5    0.57  -----
0.53  0.53  0.34  90.07  51.20
    
```

*This example shows a statistical listing of all active users. The top line (User Statistics Initialized...) of the output indicates the timeframe covered by the displayed statistics.*

*The statistical output is divided into two statistics categories, Running and Completed. Running statistics include information about jobs that are currently running. Completed statistics are compiled using historical information from both running and completed jobs.*

The fields are as follows:

Column	Description
<b>UserName</b>	Name of user.
<b>UID</b>	User ID of user.
<b>Jobs</b>	Number of running jobs.
<b>Procs</b>	Number of procs allocated to running jobs.
<b>ProcHours</b>	Number of proc-hours required to complete running jobs.
<b>Jobs*</b>	Number of jobs completed.
<b>%</b>	Percentage of total jobs that were completed by user.
<b>PHReq*</b>	Total proc-hours requested by completed jobs.
<b>%</b>	Percentage of total proc-hours requested by completed jobs that were requested by user.
<b>PHDed</b>	Total proc-hours dedicated to active and completed jobs. The proc-hours dedicated to a job are calculated by multiplying the number of allocated procs by the length of time the procs were allocated, regardless of the job's CPU usage.
<b>%</b>	Percentage of total proc-hours dedicated that were dedicated by user.
<b>FSTgt</b>	Fairshare target. A user's fairshare target is specified in the <code>fs.cfg</code> file. This

Column	Description
	value should be compared to the user's node-hour dedicated percentage to determine if the target is being met.
<b>AvgXF*</b>	Average expansion factor for jobs completed. A job's XFactor (expansion factor) is calculated by the following formula: $(\text{QueuedTime} + \text{RunTime}) / \text{WallClockLimit}$ .
<b>MaxXF*</b>	Highest expansion factor received by jobs completed.
<b>AvgQH*</b>	Average queue time (in hours) of jobs.
<b>Effic</b>	Average job efficiency. Job efficiency is calculated by dividing the actual node-hours of CPU time used by the job by the node-hours allocated to the job.
<b>WCAcc*</b>	Average wallclock accuracy for jobs completed. Wallclock accuracy is calculated by dividing a job's actual run time by its specified wallclock limit. <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> A job's wallclock accuracy is capped at 100% so even if a job exceeds its requested walltime it will report an accuracy of 100%.</p> </div>

\* These fields are empty until a user has completed at least one job.

#### Example 3-49: Fairshare Tree Statistics

```

> showstats -T
statistics initialized Mon Jul 10 15:29:41
----- Active ----- Completed
-----|-----
user      Jobs Procs ProcHours  Mem Jobs  %   PHReq  %   PHDed  %   FSTgt
  AvgXF  MaxXF  AvgQH  Effic  WCAcc
root      0      0      0.00   0      56 100.00 2.47K 100.00 1.58K 48.87 -----
  1.22   0.00   0.24 100.00 58.84
11.1      0      0      0.00   0      25 44.64 845.77 34.31 730.25 22.54 -----
  1.97   0.00   0.20 100.00 65.50
Administrati 0      0      0.00   0      10 17.86 433.57 17.59 197.17 6.09 -----
  3.67   0.00   0.25 100.00 62.74
Engineering 0      0      0.00   0      15 26.79 412.20 16.72 533.08 16.45 -----
  0.83   0.00   0.17 100.00 67.35
11.2      0      0      0.00   0      31 55.36 1.62K 65.69 853.00 26.33 -----
  0.62   0.00   0.27 100.00 53.46
Shared    0      0      0.00   0      3  5.36 97.17  3.94 44.92  1.39 -----
  0.58   0.00   0.56 100.00 31.73
Test      0      0      0.00   0      3  5.36 14.44  0.59 14.58  0.45 -----
  0.43   0.00   0.17 100.00 30.57
Research  0      0      0.00   0      25 44.64 1.51K 61.16 793.50 24.49 -----
  0.65   0.00   0.24 100.00 58.82

```

```

> showstats -T 2
statistics initialized Mon Jul 10 15:29:41
----- Active -----|----- Completed -----
user                Jobs Procs ProcHours Mem Jobs    %    PHReq    %    PHDed    %    FSTgt
AvgXF  MaxXF  AvgQH  Effic  WCAcc
Test                0    0    0.00    0    22    4.99 271.27    0.55 167.42    0.19 -----
3.86   0.00   2.89 100.00    60.76
Shared              0    0    0.00    0    59   13.38 12.30K   24.75  4.46K    5.16 -----
6.24   0.00  10.73 100.00    49.87
Research            0    0    0.00    0   140   31.75  9.54K   19.19  5.40K    6.25 -----
2.84   0.00   5.52 100.00    57.86
Administrati        0    0    0.00    0    84   19.05  7.94K   15.96  4.24K    4.91 -----
4.77   0.00   0.34 100.00    62.31
Engineering          0    0    0.00    0   136   30.84 19.67K   39.56 28.77K   33.27 -----
3.01   0.00   3.66 100.00    63.70

> showstats -T 11.1
statistics initialized Mon Jul 10 15:29:41
----- Active -----|----- Completed -----
user                Jobs Procs ProcHours Mem Jobs    %    PHReq    %    PHDed    %    FSTgt
AvgXF  MaxXF  AvgQH  Effic  WCAcc
11.1                0    0    0.00    0   220   49.89 27.60K   55.52 33.01K   38.17 -----
3.68   0.00   2.39 100.00    63.17
Administrati        0    0    0.00    0    84   19.05  7.94K   15.96  4.24K    4.91 -----
4.77   0.00   0.34 100.00    62.31
Engineering          0    0    0.00    0   136   30.84 19.67K   39.56 28.77K   33.27 -----
3.01   0.00   3.66 100.00    63.70

```

**Example 3-50: Statistics from an Absolute Time Frame**

```

> showstats -c batch -v -t 00:00:01_01/01/22,23:59:59_12/31/22
statistics initialized Wed Jan 1 00:00:00
----- Active -----|----- Completed -----
class Jobs Procs ProcHours Mem Jobs    %    PHReq    %    PHDed    %    FSTgt AvgXF
MaxXF  AvgQH  Effic  WCAcc
batch    0    0    0.00    0    23   100.00    15 100.00    1 100.00 ----- 0.40
5.01   0.00  88.94  39.87

```

*Moab returns information about the class `batch` from January 1, 2022 to December 31, 2022. For more information about specifying absolute dates, see "Absolute Time Format" in [TIMESPEC](#).*

**Example 3-51: Statistics from a Relative Time Frame**

```

> showstats -u bob -v -t -30:00:00:00
statistics initialized Mon Nov 11 15:30:00
----- Active -----|----- Completed -----
user Jobs Procs ProcHours Mem Jobs    %    PHReq    %    PHDed    %    FSTgt AvgXF
MaxXF  AvgQH  Effic  WCAcc
bob     0    0    0.00    0    23   100.00    15 100.00    1 100.00 ----- 0.40
5.01   0.00  88.94  39.87

```

*Moab returns information about user `bob` from the past 30 days. For more information about specifying relative*

*dates, see "Relative Time Format" in [TIMESPEC](#).*

## Related Topics

- (Optional) Install Moab Client - in the *Moab HPC Suite Installation and Configuration Guide* explains how to distribute this command to client nodes
- [mschedctl -f](#) command - re-initialize statistics
- [showstats -f](#) command - display full matrix statistics

## 3.7.38.F TIMESPEC

### Relative Time Format

The relative time format specifies a time by using the current time as a reference and specifying a time offset.

#### Format

+[[[DD:]HH:]MM:]SS

#### Examples

2 days, 3 hours and 57 seconds in the future:

+02:03:0:57

21 days (3 weeks) in the future:

+21:0:0:0

30 seconds in the future:

+30

### Absolute Time Format

The absolute time format specifies a specific time in the future.

#### Format

[HH[:MM[:SS]]][\_MO[/DD[/YY]]] (i.e., 14:30\_06/20)

#### Examples

1 PM, March 1 (this year)

13:00\_03/01

## 3.7.39 showstats -f

### 3.7.39.A Synopsis

```
showstats -f<statistictype> [--about] [--help] [--
host=<serverHostName>] [--loglevel=<logLevel>] [--
msg=<message>] [--port=<serverPort>] [--timeout=<seconds>] [--
version] [--xml]
```

### 3.7.39.B Overview

Shows table of various scheduler statistics.

This command displays a table of the selected Moab Scheduler statistics, such as expansion factor, bypass count, jobs, proc-hours, wallclock accuracy, and backfill information.

**i** Statistics are aggregated over time. This means statistical information is not available for time frames and the `-t` option is not supported with `showstats -f`.

### 3.7.39.C Access

This command can be run by any Moab Scheduler Administrator.

### 3.7.39.D Options

Options	Description
<b>AVGBYPASS</b>	The number of times a priority job has been 'bypassed' by backfill, allowing a lower priority job to run ahead of it. See <a href="#">Example 3-52</a> for more information.
<b>AVGQTIME</b>	Average queue time. Includes summary of job-weighted queue time and total samples.
<b>AVGXFACTOR</b>	Average expansion factor. Includes summary of job-weighted expansion factor, processor-weighted expansion factor, processor-hour-weighted

Options	Description
	expansion factor, and total number of samples.
<b>BFCOUNT</b>	Number of jobs backfilled. Includes summary of job-weighted backfill job percent and total samples.
<b>BFPHRUN</b>	Number of proc-hours backfilled. Includes summary of job-weighted backfill proc-hour percentage and total samples.
<b>ESTSTARTTIME</b>	Job start time estimate for jobs meeting specified processor/duration criteria. This estimate is based on the <a href="#">reservation start time analysis</a> algorithm.
<b>JOBCOUNT</b>	Number of jobs. Includes summary of total jobs and total samples.
<b>MAXBYPASS</b>	Maximum bypass count. Includes summary of overall maximum bypass and total samples.
<b>MAXXFACTOR</b>	Maximum expansion factor. Includes summary of overall maximum expansion factor and total samples.
<b>PHREQUEST</b>	proc-hours requested. Includes summary of total proc-hours requested and total samples.
<b>PHRUN</b>	proc-hours run. Includes summary of total proc-hours run and total samples.
<b>QOSDELIVERED</b>	Quality of service delivered. Includes summary of job-weighted quality of service success rate and total samples.
<b>WCACCURACY</b>	Wallclock accuracy. Includes summary of overall wall clock accuracy and total samples.

### 3.7.39.E Examples

```

> showstats -f AVGXFACTOR
Average XFactor Grid
[ NODES ][ 00:02:00 ][ 00:04:00 ][ 00:08:00 ][ 00:16:00 ][ 00:32:00 ][ 01:04:00 ][
02:08:00 ][ 04:16:00 ][ 08:32:00 ][ 17:04:00 ][ 34:08:00 ][ TOTAL ][
[ 1 ][ ----- ][ ----- ][ ----- ][ ----- ][ ----- ][ ----- ][ ---
----- ][ ----- ][ ----- ][ ----- ][ ----- ][ ----- ][
[ 2 ][ ----- ][ ----- ][ ----- ][ ----- ][ ----- ][ ----- ][ ---
----- ][ ----- ][ ----- ][ ----- ][ ----- ][ ----- ][
[ 4 ][ ----- ][ ----- ][ ----- ][ ----- ][ ----- ][ ----- ][


```



```

1.00  1][ ----- ][ 1.12  2][ ----- ][ ----- ][ 1.10  3]
[   8 ][ ----- ][ ----- ][ ----- ][ ----- ][ ----- ][
1.00  2][ 1.24  2][ ----- ][ ----- ][ ----- ][ 1.15  4]
[  16 ][ ----- ][ ----- ][ ----- ][ ----- ][ ----- ][ 1.01  2][ ---
----- ][ ----- ][ ----- ][ ----- ][ ----- ][ 1.01  2]
[  32 ][ ----- ][ ----- ][ ----- ][ ----- ][ ----- ][ ----- ][ ---
----- ][ ----- ][ ----- ][ ----- ][ ----- ][ ----- ][
[  64 ][ ----- ][ ----- ][ ----- ][ ----- ][ ----- ][ ----- ][ ---
----- ][ ----- ][ ----- ][ ----- ][ ----- ][ ----- ][
[ 128 ][ ----- ][ ----- ][ ----- ][ ----- ][ ----- ][ ----- ][ ---
----- ][ ----- ][ ----- ][ ----- ][ ----- ][ ----- ][
[ 256 ][ ----- ][ ----- ][ ----- ][ ----- ][ ----- ][ ----- ][ ---
----- ][ ----- ][ ----- ][ ----- ][ ----- ][ ----- ][
[ T TOT ][ ----- ][ ----- ][ ----- ][ ----- ][ ----- ][ ----- ][ 1.01  2][
1.00  3][ 1.24  2][ 1.12  2][ ----- ][ ----- ][
Job Weighted X Factor:    1.0888
Node Weighted X Factor:  1.1147
NS Weighted X Factor:    1.1900
Total Samples:           9
    
```

The `showstats -f` command returns a table with data for the specified `STATISTICSCTYPE` parameter. The left-most column shows the maximum number of processors required by the jobs shown in the other columns. The column headers indicate the maximum wallclock time (in HH:MM:SS notation) requested by the jobs shown in the columns. The data returned in the table varies by the `STATISTICSCTYPE` requested. For table entries with one number, it is of the data requested. For table entries with two numbers, the left number is the data requested and the right number is the number of jobs used to calculate the average. Table entries that contain only dashes (-----) indicate no job has completed that matches the profile associated for this inquiry. The bottom row shows the totals for each column. Following each table is a summary, which varies by the `STATISTICSCTYPE` requested.

 The column and row break down can be adjusted using the `STATPROC*` and `STATTIME*` parameters respectively.

This particular example shows the average expansion factor grid. Each table entry indicates two pieces of information — the average expansion factor for all jobs that meet this slot's profile and the number of jobs that were used to calculate this average. For example, the XFactors of two jobs were averaged to obtain an average XFactor of 1.24 for jobs requiring over 2 hours 8 minutes, but not more than 4 hours 16 minutes and between 5 and 8 processors. Totals along the bottom provide overall XFactor averages weighted by job, processors, and processor-hours.

**Example 3-52:**

```

> showstats -f AVGBYPASS
Average Bypass (bypass count)
[ PROCS ][ 0:15:00 ][ 1:00:00 ][ 4:00:00 ][ 16:00:00 ][ 64:00:00 ][ 256:00:00 ][ TOTAL
]
[ 1 ][ 0.00 10][ 0.00 70][ 0.00 31][ 0.00 34][ 0.00 6][-----][ 0.00 150]
[ 4 ][-----][-----][-----][-----][-----][-----][-----]
--][-----]
[ 16 ][ 0.08 37687][ 0.08 164307][ 0.32 117767][ 0.10 34073][ 0.58 1282][-----]
[ 0.16 355116]
[ 64 ][ 0.18 769][ 0.13 1839][ 0.18 8084][ 0.82 2812][ 0.00 34][-----][ 0.31
13538]
[ 256 ][ 0.39 316][ 1.40 778][ 4.40 494][ 1.77 28917][ 0.33 6][-----][ 1.79
30511]
[ TOTAL ][ 0.08 38782][ 0.09 166994][ 0.33 126376][ 0.86 65835][ 0.57 1328][-----]
--]
Job Weighted X Bypass: 0.2932
Total Samples: 399315
    
```

The `showstats -f` command returns a table with data for the specified `STATISTICSCTYPE` parameter, in this case for `AVGBYPASS`. In this particular example, the upper left cell indicates that 10 jobs were run by Moab, which had 0-15 minutes of requested walltime and 0-1 procs allocated. The 0.00 indicates that of the 10 jobs, the average number

*of times the jobs were bypassed was 0, meaning it did not occur. Further, looking at row 256 and column 4:00, we see that 494 jobs have been run by Moab that meet this criteria. On average, these jobs were each bypassed 4.40 times.*

---

### Related Topics


- (Optional) Install Moab Client - in the *Moab HPC Suite Installation and Configuration Guide* explains how to distribute this command to client nodes
- [mschedctl -f](#) command
- [showstats](#) command
- [STATPROCMIN](#) parameter
- [STATPROCSTEPCOUNT](#) parameter
- [STATPROCSTEPSIZE](#) parameter
- [STATTIMEMIN](#) parameter
- [STATTIMESTEPCOUNT](#) parameter
- [STATTIMESTEPSIZE](#) parameter

## 3.7.40 Deprecated Commands

In this topic:

- [canceljob](#)
- [changeparam](#)
- [diagnose](#)
- [releasehold](#)
- [releaseres](#)
- [resetstats](#)
- [runjob](#)
- [sethold](#)
- [setqos](#)
- [setres](#)
- [setspri](#)
- [showconfig](#)

### 3.7.40.A canceljob

 This command is deprecated. Use [mjobctl -c](#) instead.

#### Synopsis

`canceljob jobid [jobid]...`

#### Overview

The `canceljob` command is used to selectively cancel the specified job(s) (active, idle, or non-queued) from the queue.

#### Access

This command can be run by any Moab Administrator and by the owner of the job (see [ADMINCFG](#)).


Flag	Name	Format	Default	Description	Example
<b>-h</b>	HELP		N/A	Display usage information.	<pre>&gt; canceljob -h</pre>
	JOB ID	<STRING>	---	A jobid, a job expression, or the keyword ALL.	<pre>&gt; canceljob 13001 13003</pre>

#### Examples

*Example 3-53: Cancel job 6397*

```
> canceljob 6397
```

### 3.7.40.B changeparam

 This command is deprecated. Use [mschedctl -m](#) instead.

#### Synopsis

`changeparam parameter value`

## Overview

The `changeparam` command is used to dynamically change the value of any parameter that can be specified in the `moab.cfg` file. The changes take effect at the beginning of the next scheduling iteration. They are not persistent, only lasting until Moab is shut down.

`changeparam` is a compact command of `mschedctl -m`.

## Access

This command can be run by a level 1 Moab administrator.

### 3.7.40.C diagnose



This command is deprecated. Use `mdiag` instead.

## Synopsis

```
diagnose -a [accountid]
diagnose -b [-l policylevel] [-t partition]
diagnose -c [classid]
diagnose -C [configfile]
diagnose -f [-o user|group|account|qos|class]
diagnose -g [groupid]
diagnose -j [jobid]
diagnose -L
diagnose -m [rackid]
diagnose -n [-t partition] [nodeid]
diagnose -p [-t partition]
diagnose -q [qosid]
diagnose -r [reservationid]
diagnose -R [resourcemanagename]
diagnose -s [standingreservationid]
diagnose -S diagnose -u [userid]
diagnose -v
```

diagnose -x

## Overview

The *diagnose* command is used to display information about various aspects of scheduling and the results of internal diagnostic tests.

### 3.7.40.D releasehold



This command is deprecated. Use [mjobctl -u](#) instead.

## Synopsis

```
releasehold [-a|-b] jobexp
```

## Overview

Release hold on specified job(s).

This command allows you to release batch holds or all holds (system, user, and batch) on specified jobs. Any number of jobs can be released with this command.

## Access

By default, this command can be run by any Moab Scheduler Administrator.

## Parameters

JOBEXP	Job expression of job(s) to release.
--------	--------------------------------------

## Flags

-a	Release all types of holds (user, system, batch) for specified job(s).
-b	Release batch hold from specified job(s).
-h	Help for this command.

## Examples

### Example 3-54: `releasehold -b`

```
> releasehold -b 6443
batch hold released for job 6443
```

*In this example, a batch hold was released from this one job.*

### Example 3-55: `releasehold -a`

```
> releasehold -a "81[1-6]"
holds modified for job 811
holds modified for job 812
holds modified for job 813
holds modified for job 814
holds modified for job 815
holds modified for job 816
```

*In this example, all holds were released from the specified jobs.*

---

## Related Topics

- [sethold](#)
- [mjobctl](#)

### 3.7.40.E `releaseres`



This command is deprecated. Use [mrsvctl -r](#) instead.

## Synopsis

```
releaseres [arguments] reservationid [reservationid...]
```

## Overview

Release existing reservation.

This command allows Moab Scheduler Administrators to release any user, group, account, job, or system reservation. Users are allowed to release reservations on jobs they own.

Note that releasing a reservation on an active job has no effect since the reservation will be automatically recreated.

## Access

Users can use this command to release any reservation they own. Level 1 and level 2 Moab administrators can use this command to release any reservation.

## Parameters


RESERVATION ID	Name of reservation to release.
----------------	---------------------------------

## Examples

*Example 3-56: Release two existing reservations*

```
> releaseres system.1 bob.2
released User reservation 'system.1'
released User reservation 'bob.2'
```

### 3.7.40.F resetstats

 This command is deprecated. Use [mschedctl -f](#) instead.

## Synopsis

*resetstats*

## Overview

This command resets all internally-stored Moab Scheduler statistics to the initial start-up state as of the time the command was executed.


## Access

By default, this command can be run by level 1 scheduler administrators.

## Examples

```
> resetstats Statistics Reset at time Wed Feb 25 23:24:55 2022
```

### 3.7.40.G runjob

 This command is deprecated. Use [mjobctl -x](#) instead.

## Synopsis

*runjob* [-c|-f|-n *nodelist*|-p *partition*|-s|-x] *jobid*

## Overview

This command will attempt to immediately start the specified job.

`runjob` is a deprecated command, replaced by [mjobctl](#).

## Access

By default, this command can be run by any Moab administrator.

## Parameters

JOBID	Name of the job to run.
-------	-------------------------

Args	Description
<b>-c</b>	<i>Clear</i> job parameters from previous runs (used to clear PBS neednodes attribute after PBS job launch failure)
<b>-f</b>	Attempt to <i>force</i> the job to run, ignoring throttling policies
<b>-n &lt;NODELIST&gt;</b>	Attempt to start the job using the specified <i>nodelist</i> where nodenames are comma or colon delimited
<b>-p &lt;PARTITION&gt;</b>	Attempt to start the job in the specified <i>partition</i>
<b>-s</b>	Attempt to <i>suspend</i> the job
<b>-x</b>	Attempt to force the job to run, ignoring throttling policies, QoS constraints, and reservations

## Examples

*Example 3-57: Run job cluster.231*

```
> runjob cluster.231
job cluster.231 successfully started
```


## See Also

- [mjobctl](#)
- [canceljob](#) - cancel a job



- `checkjob` - show detailed status of a job
- `showq` - list queued jobs

### 3.7.40.H sethold

 This command is deprecated. Use `mjobctl -h` instead.

#### Synopsis

```
sethold [-b] jobid [jobid...]
```

#### Overview

Set hold on specified job(s).

#### Permissions

This command can be run by any Moab Scheduler Administrator.

#### Parameters

JOB	Job number of job to hold.
-----	----------------------------

#### Flags


<b>-b</b>	Set a batch hold. Typically, only the scheduler places batch holds. This flag allows an administrator to manually set a batch hold.
<b>-h</b>	Help for this command.

#### Examples

```
> sethold -b fr17n02.1072.0 fr15n03.1017.0
Batch Hold Placed on All Specified Jobs
```

*In this example, a batch hold is placed on job `fr17n02.1072.0` and job `fr15n03.1017.0`.*

### 3.7.40.I setqos

 This command is deprecated. Use [mjobctl -m](#) instead.

#### Synopsis

```
setqos qosid jobid
```

#### Overview

Set Quality Of Service for a specified job.

This command allows users to change the QOS of their own jobs.

#### Access

This command can be run by any user.

#### Parameters

JOBID	Job name.
QOSID	QOS name.


#### Examples

```
> setqos high_priority moab.3
```

```
Job QOS Adjusted
```

*This example sets the Quality Of Service to a value of high\_priority for job moab.3.*

### 3.7.40.J setres

 This command is deprecated. Use [mrsvctl -c](#) instead.

#### Synopsis

```
setres [arguments] resourceexpression
[ -a <ACCOUNT_LIST> ]
[ -b <SUBTYPE> ]
[ -c <CHARGE_SPEC> ]
[ -d <DURATION> ]
[ -e <ENDTIME> ]
```

```

[ -E ] // EXCLUSIVE
[ -f <FEATURE_LIST> ]
[ -g <GROUP_LIST> ]
[ -n <NAME> ]
[ -o <OWNER> ]
[ -p <PARTITION> ]
[ -q <QUEUE_LIST> ] // (i.e., CLASS_LIST)
[ -Q <QOSLIST> ]
[ -r <RESOURCE_DESCRIPTION> ]
[ -R <RESERVATION_PROFILE> ]
[ -s <STARTTIME> ]
[ -T <TRIGGER> ]
[ -u <USER_LIST> ]
[ -x <FLAGS> ]

```

## Overview

Reserve resources for use by jobs with particular credentials or attributes.

## Access

This command can be run by level 1 and level 2 Moab administrators.

## Parameters

Name	Format	Default	Description
<b>ACCOUNT_LIST</b>	<STRING> [:<STRING>]...	---	List of accounts that will be allowed access to the reserved resources
<b>SUBTYPE</b>	<STRING>	---	Specify the subtype for a reservation
<b>CHARGE_SPEC</b>	<ACCOUNT> [,<GROUP> [,<USER>]]	---	Specifies which credentials will be accountable for unused resources dedicated to the reservation
<b>CLASS_LIST</b>	<STRING> [:<STRING>]...	---	List of classes that will be allowed access to the reserved resource
<b>DURATION</b>	[[[DD:]HH:]MM:]SS	INFINITY	Duration of the reservation (not needed if ENDTIME is specified)
<b>ENDTIME</b>	[HH[:MM[:SS]]][_MO[/DD[/YY]]]	INFINITY	Absolute or relative time reservation will end (not required if Duration specified)

Name	Format	Default	Description
	or + [[[DD:]HH:]MM:]SS		
<b>EXCLUSIVE</b>	N/A	N/A	Requests exclusive access to resources
<b>FEATURE_LIST</b>	<STRING> [:<STRING>]...	---	List of node features that must be possessed by the reserved resources
<b>FLAGS</b>	<STRING> [:<STRING>]...	---	List of reservation flags (see <a href="#">Managing Reservations</a> for details)
<b>GROUP_LIST</b>	<STRING> [:<STRING>]...	---	List of groups that will be allowed access to the reserved resources
<b>NAME</b>	<STRING>	Name set to first name listed in ACL or SYSTEM if no ACL specified	Name for new reservation
<b>OWNER</b>	<CREDTYPE> :<CREDID> where CREDTYPE is one of user, group, acct, class, or qos	N/A	Specifies which credential is granted reservation ownership privileges
<b>PARTITION</b>	<STRING>	[ANY]	Partition where resources must be located
<b>QOS_LIST</b>	<STRING> [:<STRING>]...	---	List of QOSs that will be allowed access to the reserved resource
<b>RESERVATION_PROFILE</b>	Existing reservation profile ID	N/A	Requests that default reservation attributes be loaded from the specified reservation profile (see <a href="#">RSVPROFILE</a> )

Name	Format	Default	Description
<b>RESOURCE_DESCRIPTION</b>	Colon delimited list of zero or more of the following <ATTR> =<VALUE> pairs PROCS= <INTEGER> MEM= <INTEGER> DISK=<INTEGER> SWAP= <INTEGER> GRES=<STRING>	PROCS=-1	The resources to be reserved per task. (-1 indicates all resources on node)
<b>RESOURCE_EXPRESSION</b>	ALL or TASKS { ==   >= }<TASKCOUNT> or <HOST_REGEX>	Required Field. No Default	The tasks to reserve. ALL indicates all resources available should be reserved. <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> If ALL or a host expression is specified, Moab will apply the reservation regardless of existing reservations and exclusive issues. If TASKS is used, Moab will only allocate accessible resources.</p> </div>
<b>STARTTIME</b>	[HH[:MM[:SS]]][_MO[/DD[/YY]]] or + [[[DD:]HH:]MM:]SS	NOW	Absolute or relative time reservation will start
<b>TRIGGER</b>	<STRING>	N/A	Comma-delimited reservation trigger list following format described in the <a href="#">trigger format</a> section of the reservation configuration overview
<b>USER_LIST</b>	<STRING> [:<STRING>]...	---	List of users that will be allowed access to the reserved resources

## Description

The *setres* command allows an arbitrary block of resources to be reserved for use by jobs that meet the specified access constraints. The timeframe covered by the reservation

can be specified on either an absolute or relative basis. Only jobs with credentials listed in the reservation ACL (i.e., USERLIST, GROUPLIST,...) can utilize the reserved resources. However, these jobs still have the freedom to utilize resources outside of the reservation. The reservation will be assigned a name derived from the ACL specified. If no reservation ACL is specified, the reservation is created as a system reservation and no jobs will be allowed access to the resources during the specified timeframe (valuable for system maintenance, etc.). See the [Reservation Overview](#) for more information.

Reservations can be viewed using the [showres](#) command and can be released using the [releaseres](#) command.

## Examples

```
> setres -u john:mary -s +24:00:00 -d 8:00:00 TASKS==2
reservation 'john.1' created on 2 nodes (2 tasks)
node001:1
node005:1
```

*Reserve two nodes for use by users john and mary for a period of 8 hours starting in 24 hours.*

### Example 3-58:

```
> setres -s 8:00:00_06/20 -e 17:00:00_06/22 ALL
reservation 'system.1' created on 8 nodes (8 tasks)
node001:1
node002:1
node003:1
node004:1
node005:1
node006:1
node007:1
node008:1
```

*Schedule a system wide reservation to allow system maintenance on Jun 20, 8:00 AM until Jun 22, 5:00 PM.*

### Example 3-59:

```
> setres -r PROCS=1:MEM=512 -g staff -l interactive 'node00[3-6]'
reservation 'staff.1' created on 4 nodes (4 tasks)
node003:1
node004:1
node005:1
node006:1
```

*Reserve one processor and 512 MB of memory on nodes node003 through node node006 for members of the group staff and jobs in the interactive class.*

## 3.7.40.K setspri



This command is deprecated. Use [mjobctl -p](#) instead.

## Synopsis

```
setspri [-r] priorityjobid
```

## Overview

(This command is deprecated by the [mjobctl command](#))

Set or remove absolute or relative system priorities for a specified job.

This command allows you to set or remove a system priority level for a specified job. Any job with a system priority level set is guaranteed a higher priority than jobs without a system priority. Jobs with higher system priority settings have priority over jobs with lower system priority settings.

## Access

This command can be run by any Moab Scheduler Administrator.

## Parameters

JOB	Name of job.
PRIORITY	System priority level. By default, this priority is an absolute priority overriding the policy generated priority value. Range is 0 to clear, 1 for lowest, 1000 for highest. The given value is added onto the system priority (see 32-bit and 64-bit values below), except for a given value of zero. If the '-r' flag is specified, the system priority is relative, adding or subtracting the specified value from the policy generated priority.  If a relative priority is specified, any value in the range +/- 1,000,000,000 is acceptable.

## Flags

**-r** Set relative system priority on job.

## Examples

```
> setspri 10 orion.4752
job system priority adjusted
```

*In this example, a system priority of 10 is set for job orion.4752.*

**Example 3-60:**

```
> setspri 0 clusterB.1102
job system priority adjusted
```

*In this example, system priority is cleared for job clusterB.1102.*

**Example 3-61:**

```
> setspri -r 100000 job.00001
job system priority adjusted
```

*In this example, the job's priority will be increased by 100000 over the value determine by configured priority policy.*

**i** This command is deprecated. Use [mjobctl](#) instead.

### 3.7.40.L showconfig

#### Synopsis

```
showconfig [-v] [--about] [--help] [--host=<serverHostName>] [--loglevel=<logLevel>] [--msg=<message>] [--port=<serverPort>] [--timeout=<seconds>] [--version] [--xml]
```

#### Overview

View the current configurable parameters of the Moab Scheduler.

The `showconfig` command shows the current scheduler version and all scheduler parameters. These parameters are set via internal defaults, command line arguments, environment variable settings, parameters in the `moab.cfg` file, and via the [mschedctl -m](#) command. Because of the many sources of configuration settings, the output may differ from the contents of the `moab.cfg` file. The output is such that it can be saved and used as the contents of the `moab.cfg` file if desired.

**i** The `showconfig` command does not show credential parameters (such as user, group class, QoS, account).

#### Access

This command can be run by a level 1, 2, or 3 Moab administrator.

#### Flags

<b>-h</b>	Help for this command.
<b>-v</b>	This optional flag turns on verbose mode, which shows all possible Moab Scheduler parameters and their current settings. If this flag is not used, this command operates in context-sensitive terse mode, which shows only certain parameter settings.



## Examples

### Example 3-62: *showconfig*

```
> showconfig
# moab scheduler version 4.2.4 (PID: 11080)
BACKFILLPOLICY          FIRSTFIT
BACKFILLMETRIC          NODES
ALLOCATIONPOLICY         MINRESOURCE
RESERVATIONPOLICY      CURRENTHIGHEST
...
```

**i** The *showconfig* command without the `-v` flag does not show the settings of all scheduling parameters. To show the settings of all scheduling parameters, use the `-v` (verbose) flag. This will provide an extended output. This output is often best used in conjunction with the *grep* command as the output can be voluminous.

---

### Related Topics

- Use the `mschedctl -m` command to change the various Moab Scheduler parameters
- See the [Moab Parameters](#) appendix for details about configurable parameters

## Chapter 4: Prioritizing Jobs and Allocating Resources

In this section:

- [4.1 Job Prioritization - page 335](#)
- [4.2 Node Allocation Policies - page 354](#)
- [4.3 Node Access Policies - page 365](#)
- [4.4 Node Availability Policies - page 367](#)

## 4.1 Job Prioritization

In general, prioritization is the process of determining which of many options best fulfills overall goals. In the case of scheduling, a site will often have multiple, independent goals that may include maximizing system utilization, giving preference to users in specific projects, or making certain that no job sits in the queue for more than a given period of time. The approach used by Moab in representing a multi-faceted set of site goals is to assign weights to the various objectives so an overall value or priority can be associated with each potential scheduling decision. With the jobs prioritized, the scheduler can roughly fulfill site objectives by starting the jobs in priority order.

In this topic:

- [4.1.1 Priority Overview - page 335](#)
- [4.1.2 Job Priority Factors - page 336](#)
- [4.1.3 Fairshare Job Priority Example - page 348](#)
- [4.1.4 Common Priority Usage - page 349](#)
- [4.1.5 Prioritization Strategies - page 352](#)
- [4.1.6 Manual Job Priority Adjustment - page 353](#)

---

### Related Topics

- [mdiag -p](#) (Priority Diagnostics)

### 4.1.1 Priority Overview

Moab's prioritization mechanism allows component and subcomponent weights to be associated with many aspects of a job to enable fine-grained control over this aspect of scheduling. To allow this level of control, Moab uses a simple priority-weighting hierarchy where the contribution of each priority subcomponent is calculated as follows:

$$\langle \text{COMPONENT WEIGHT} \rangle * \langle \text{SUBCOMPONENT WEIGHT} \rangle * \langle \text{PRIORITY SUBCOMPONENT VALUE} \rangle$$

Each priority component contains one or more subcomponents as described in the section titled [Job Priority Factors](#). For example, the Resource component consists of Node, Processor, Memory, Swap, Disk, Walltime, and PE subcomponents. While there are numerous priority components and many more subcomponents, a site need only focus on and configure the subset of components related to their particular priority needs. In actual

usage, few sites use more than a small fraction (usually 5 or fewer) of the available priority subcomponents. This results in fairly straightforward priority configurations and tuning. By mixing and matching priority weights, sites can generally obtain the desired job-start behavior. At any time, you can issue the [mdiag -p](#) command to determine the impact of the current priority-weight settings on idle jobs. Likewise, the command [showstats -f](#) can assist the administrator in evaluating priority effectiveness on historical system usage metrics such as queue time or expansion factor.

As mentioned above, a job's priority is the weighted sum of its activated subcomponents. By default, the value of all component and subcomponent weights is set to 1 and 0 respectively. The one exception is the 'QUEUE TIME' subcomponent weight that is set to 1. This results in a total job priority equal to the period of time the job has been queued, causing Moab to act as a simple FIFO. Once the summed component weight is determined, this value is then bounded resulting in a priority ranging between 0 and MAX\_PRIO\_VAL, which is currently defined as 1000000000 (one billion). In no case will a job obtain a priority in excess of MAX\_PRIO\_VAL through its priority subcomponent values.

**i** Negative priority jobs can be allowed if desired; see [ENABLENEGJOBPRIORITY](#) and [REJECTNEGPRIOJOBS](#) for more information.

Using the [mjobctl -p](#) command, site administrators can adjust the base calculated job priority by either assigning a relative priority adjustment or an absolute system priority. A relative priority adjustment causes the base priority to be increased or decreased by a specified value. Setting an absolute system priority, SPRIO, causes the job to receive a priority equal to MAX\_PRIO\_VAL + SPRIO, and therefore guaranteed to be of higher value than any naturally occurring job priority.

---

## Related Topics

- [REJECTNEGPRIOJOBS](#) parameter

## 4.1.2 Job Priority Factors

In this topic:

- [4.1.2.A Job Priority Factors and Subfactors - page 337](#)
- [4.1.2.B Credential \(CRED\) Component - page 340](#)
- [4.1.2.C Fairshare \(FS\) Component - page 341](#)

- [4.1.2.D Resource \(RES\) Component - page 342](#)
- [4.1.2.E Service \(SERVICE\) Component - page 343](#)
- [4.1.2.F Target Service \(TARG\) Component - page 346](#)
- [4.1.2.G Usage \(USAGE\) Component - page 346](#)
- [4.1.2.H Job Attribute \(ATTR\) Component - page 347](#)

### 4.1.2.A Job Priority Factors and Subfactors

Moab allows jobs to be prioritized based on a range of job related factors. These factors are broken down into a two-tier hierarchy of priority factors and subfactors, each of which can be independently assigned a weight. This approach provides the administrator with detailed yet straightforward control of the job selection process.

Each factor and subfactor can be configured with independent priority weight and priority cap values (described later). In addition, per credential and per QoS priority weight adjustments can be specified for a subset of the priority factors. For example, QoS credentials can adjust the queue time subfactor weight and group credentials can adjust fairshare subfactor weight.

The following table highlights the factors and subfactors that make up a job's total priority:

Factor	SubFactor	Metric
<b>CRED</b> (job credentials)	USER	user-specific priority (see <a href="#">USERCFG</a> )
	GROUP	group-specific priority (see <a href="#">GROUPCFG</a> )
	ACCOUNT	account-specific priority (see <a href="#">ACCOUNTCFG</a> )
	QOS	QoS-specific priority (see <a href="#">QOSCFG</a> )
	CLASS	class/queue-specific priority (see <a href="#">CLASSCFG</a> )
<b>FS</b> (fairshare usage)	FSUSER	user-based historical usage (see <a href="#">Fairshare Overview</a> )
	FSGROUP	group-based historical usage (see <a href="#">Fairshare Overview</a> )
	FSACCOUNT	account-based historical usage (see <a href="#">Fairshare</a> )

Factor	SubFactor	Metric
		<a href="#">Overview</a> )
	FSQOS	QoS-based historical usage (see <a href="#">Fairshare Overview</a> )
	FSCLASS	class/queue-based historical usage (see <a href="#">Fairshare Overview</a> )
	FSGUSER	imported global user-based historical usage (see <a href="#">ID Manager</a> and <a href="#">Fairshare Overview</a> )
	FSGGROUP	imported global group-based historical usage (see <a href="#">ID Manager</a> and <a href="#">Fairshare Overview</a> )
	FSGACCOUNT	imported global account-based historical usage (see <a href="#">ID Manager</a> and <a href="#">Fairshare Overview</a> )
	FSJPU	current active jobs associated with job user
	FSPPU	current number of processors allocated to active jobs associated with job user
	FSPSPU	current number of processor-seconds allocated to active jobs associated with job user
	WCACCURACY	user's current historical job wallclock accuracy calculated as total processor-seconds dedicated / total processor-seconds requested
		 Factor values are in the range of 0.0 to 1.0.

Factor	SubFactor	Metric
<b>RES</b> (requested job resources)	NODE	number of nodes requested
	PROC	number of processors requested
	MEM	total real memory requested (in MB)
	SWAP	total virtual memory requested (in MB)
	DISK	total local disk requested (in MB)
	PS	total processor-seconds requested
	PE	total processor-equivalent requested
	WALLTIME	total walltime requested (in seconds)
<b>SERV</b> (current service levels)	QUEUE TIME	time job has been queued (in minutes)
	XFACTOR	minimum job expansion factor
	BYPASS	number of times job has been bypassed by backfill
	STARTCOUNT	number of times job has been restarted
	DEADLINE	proximity to job deadline
	SPVIOLATION	Boolean indicating whether the active job violates a <b>soft</b> usage limit
	USERPRIO	user-specified job priority
<b>TARGET</b> (target service levels)	TARGETQUEUE TIME	time until queue time target is reached (exponential)
	TARGETXFACTOR	distance to target expansion factor (exponential)

Factor	SubFactor	Metric
<b>USAGE</b> (consumed resources -- active jobs only)	CONSUMED	processor-seconds dedicated to date
	REMAINING	processor-seconds outstanding
	PERCENT	percent of required walltime consumed
	EXECUTIONTIME	seconds since job started
<b>ATTR</b> (job attribute-based prioritization)	ATTRATTR	<a href="#">Attribute priority</a> if specified job attribute is set (attributes can be user-defined or one of <b>preemptor</b> , or <b>preemptee</b> ). Default is 0.
	ATTRSTATE	<a href="#">Attribute priority</a> if job is in specified state (see <a href="#">Job States</a> ). Default is 0.
	ATTRGRES	<a href="#">Attribute priority</a> if a <a href="#">generic resource</a> is requested. Default is 0.

**i** \*CAP parameters (FSCAP, for example) are available to limit the maximum absolute value of each priority component and subcomponent. If set to a positive value, a priority cap will bound priority component values in both the positive and negative directions.

**i** All \*CAP and \*WEIGHT parameters are specified as positive or negative integers. Non-integer values are not supported.

#### 4.1.2.B Credential (CRED) Component

The credential component allows a site to prioritize jobs based on political issues such as the relative importance of certain groups or accounts. This allows direct political priorities to be applied to jobs.

The priority calculation for the credential component is as follows:

$$\text{Priority} += \text{CREDWEIGHT} * ( \\ \text{USERWEIGHT} * \text{Job.User.Priority} + \\ \text{GROUPWEIGHT} * \text{Job.Group.Priority} + \\ \text{ACCOUNTWEIGHT} * \text{Job.Account.Priority} + \\ \text{QOSWEIGHT} * \text{Job.Qos.Priority} + \\ \text{CLASSWEIGHT} * \text{Job.Class.Priority})$$



All user, group, account, QoS, and class weights are specified by setting the `PRIORITY` attribute of using the respective `*CFG` parameter (namely, `USERCFG`, `GROUPCFG`, `ACCOUNTCFG`, `QOSCFG`, and `CLASSCFG`).

For example, to set user and group priorities, you could use the following:

```
CREDWEIGHT      1
USERWEIGHT      1
GROUPWEIGHT     1
USERCFG[john]   PRIORITY=2000
USERCFG[paul]   PRIORITY=-1000
GROUPCFG[staff] PRIORITY=10000
```

**i** Class (or queue) priority can also be specified via the resource manager where supported (as in PBS queue priorities). However, if Moab class priority values are also specified, the resource manager priority values will be overwritten.

All priorities can be positive or negative.

### 4.1.2.C Fairshare (FS) Component

Fairshare components allow a site to favor jobs based on short-term historical usage. The [Fairshare Overview](#) describes the configuration and use of fairshare in detail.

The fairshare factor is used to adjust a job's priority based on current and historical percentage system utilization of the job's user, group, account, class, or QoS. This allows sites to steer workload toward a particular usage mix across user, group, account, class, and QoS dimensions.

The fairshare priority factor calculation is as follows:

$$\text{Priority} += \text{FSWEIGHT} * \text{MIN}(\text{FSCAP}, ($$

$$\text{FSUSERWEIGHT} * \text{DeltaUserFSUsage} +$$

$$\text{FSGROUPWEIGHT} * \text{DeltaGroupFSUsage} +$$

$$\text{FSACCOUNTWEIGHT} * \text{DeltaAccountFSUsage} +$$

$$\text{FSQOSWEIGHT} * \text{DeltaQOSFSUsage} +$$

$$\text{FSCLASSWEIGHT} * \text{DeltaClassFSUsage} +$$

$$\text{FSJPUWEIGHT} * \text{ActiveUserJobs} +$$

$$\text{FSPPUWEIGHT} * \text{ActiveUserProcs} +$$

$$\text{FSPSPUWEIGHT} * \text{ActiveUserPS} +$$

$$\text{WCACCURACYWEIGHT} * \text{UserWCAccuracy} ))$$

All `*WEIGHT` parameters just listed are specified on a per partition basis in the `moab.cfg` file. The `Delta*Usage` components represent the difference in actual fairshare usage from the corresponding fairshare usage target. Actual fairshare usage is determined based on historical usage over the time frame specified in the fairshare configuration. The target usage can be a target, floor, or ceiling value as specified in the

fairshare configuration file. See the [Fairshare Overview](#) for further information on configuring and tuning fairshare. Additional insight may be available in the [fairshare usage example](#). The `ActiveUser*` components represent current usage by the job's user credential.

### How violated ceilings and floors affect fairshare-based priority

Moab determines `FSUsageWeight` in the previous section. In order to account for violated ceilings and floors, Moab multiplies that number by the `FSUsagePriority` as demonstrated in the following formula:

$$\text{FSPriority} = \text{FSUsagePriority} * \text{FSUsageWeight}$$

When a ceiling or floor is violated, `FSUsagePriority = 0`, so `FSPriority = 0`. This means the job will gain no priority because of fairshare. If fairshare is the only component of priority, then violation takes the priority to 0. For more information, see [Priority-Based Fairshare](#) and [Fairshare Targets](#).

#### 4.1.2.D Resource (RES) Component

Weighting jobs by the amount of resources requested allows a site to favor particular types of jobs. Such prioritization may allow a site to better meet site mission objectives, improve fairness, or even improve overall system utilization.

Resource based prioritization is valuable when you want to favor jobs based on the resources requested. This is good in three main scenarios: (1) when you need to favor large resource jobs because it's part of your site's mission statement, (2) when you want to level the response time distribution across large and small jobs (small jobs are more easily backfilled and therefore generally have better turnaround time), and (3) when you want to improve system utilization. While this might be surprising, system utilization actually increases as large resource jobs are pushed to the front of the queue. This keeps the smaller jobs in the back where they can be selected for backfill and therefore increase overall system utilization. The situation is like the story about filling a cup with golf balls and sand. If you put the sand in first, it gets in the way and you are unable to put in as many golf balls. However, if you put in the golf balls first, the sand can easily be poured in around them completely filling the cup.

The calculation for determining the total resource priority factor is as follows:

```
Priority += RESWEIGHT* MIN(RESCAP, (
  NODEWEIGHT * TotalNodesRequested +
  PROCWEIGHT * TotalProcessorsRequested +
  MEMWEIGHT * TotalMemoryRequested +
  SWAPWEIGHT * TotalSwapRequested +
  DISKWEIGHT * TotalDiskRequested +
  WALLTIMEWEIGHT* TotalWalltimeRequested +
  PEWEIGHT * TotalPERequested))
```

The sum of all weighted resources components is then multiplied by the `RESWEIGHT` parameter and capped by the `RESCAP` parameter. Memory, Swap, and Disk are all measured in megabytes (MB). The final resource component, PE, represents **Processor Equivalents**. This component can be viewed as a processor-weighted maximum *percentage of total resources* factor.

For example, if a job requested 25% of the processors and 50% of the total memory on a 128-processor system, it would have a PE value of  $\text{MAX}(25,50) * 128$ , or 64. The concept of PEs is a highly effective metric in shared resource systems.

**i** Ideal values for requested job processor count and walltime can be specified using `PRIORITYTARGETPROCCOUNT` and `PRIORITYTARGETDURATION`.

#### 4.1.2.E Service (SERVICE) Component

The Service component specifies which service metrics are of greatest value to the site. Favoring one service subcomponent over another generally improves that service metric.

The priority calculation for the service priority factor is as follows:

```
Priority += SERVICEWEIGHT * (
  QUEUETIMEWEIGHT * <QUEUETIME> +
  XFACTORWEIGHT * <XFACTOR> +
  BYPASSWEIGHT * <BYPASSCOUNT> +
  STARTCOUNTWEIGHT * <STARTCOUNT> +
  DEADLINEWEIGHT * <DEADLINE> +
  SPVIOLATIONWEIGHT * <SPBOOLEAN> +
  USERPRIOWEIGHT * <USERPRIO> )
```

#### QueueTime (QUEUETIME) Subcomponent

In the priority calculation, a job's queue time is a duration measured in minutes. Using this subcomponent tends to prioritize jobs in a FIFO order. Favoring queue time improves queue time based fairness metrics and is probably the most widely used single job priority metric. In fact, under the initial default configuration, this is the only priority subcomponent enabled within Moab. It is important to note that within Moab, a job's queue time is not necessarily the amount of time since the job was submitted. The parameter `JOBPRIOACCRUALPOLICY` allows a site to select how a job will accrue queue time based on meeting various **throttling policies**. Regardless of the policy used to determine a job's queue time, this effective queue time is used in the calculation of the `QUEUETIME`, `XFACTOR`, `TARGETQUEUETIME`, and `TARGETXFACTOR` priority subcomponent values.

The need for a distinct effective queue time is necessitated by the fact that many sites have users who like to work the system, whatever system it happens to be. A common practice at some long existent sites is for some users to submit a large number of jobs and then place

them on hold. These jobs remain with a hold in place for an extended period of time and when the user is ready to run a job, the needed executable and data files are linked into place and the hold released on one of these presubmitted jobs. The extended hold time guarantees that this job is now the highest priority job and will be the next to run. The use of the `JOBPRIOACCRUALPOLICY` parameter can prevent this practice and prevent 'queue stuffers' from doing similar things on a shorter time scale. These 'queue stuffer' users submit hundreds of jobs at once to swamp the machine and consume use of the available compute resources. This parameter prevents the user from gaining any advantage from stuffing the queue by not allowing these jobs to accumulate any queue time based priority until they meet certain idle and active Moab fairness policies (such as max job per user and max idle job per user).

As a final note, you can adjust the `QUEUETIMEWEIGHT` parameter on a per QoS basis using the `QOSCFG` parameter and the `QTWEIGHT` attribute. For example, the line `QOSCFG [special] QTWEIGHT=5000` causes jobs using the QoS `special` to have their queue time subcomponent weight increased by 5000.

## Expansion Factor (XFACTOR) Subcomponent

The expansion factor subcomponent has an effect similar to the queue time factor but favors shorter jobs based on their requested wallclock run time. In its traditional form, the expansion factor (XFactor) metric is calculated as follows:

$$\text{XFACTOR} = 1 + \langle \text{QUEUETIME} \rangle / \langle \text{EXECUTIONTIME} \rangle$$

However, a couple of aspects of this calculation make its use more difficult. First, the length of time the job will actually run—`<EXECUTIONTIME>`—is not actually known until the job completes. All that is known is how much time the job requests. Secondly, as described in the [Queue Time Subcomponent](#) section, Moab does not necessarily use the raw time since job submission to determine `<QUEUETIME>` to prevent various scheduler abuses. Consequently, Moab uses the following modified equation:

$$\text{XFACTOR} = 1 + \langle \text{EFFQUEUETIME} \rangle / \langle \text{WALLCLOCKLIMIT} \rangle$$

In the equation Moab uses, `<EFFQUEUETIME>` is the effective queue time subject to the `JOBPRIOACCRUALPOLICY` parameter and `<WALLCLOCKLIMIT>` is the user—or system—specified job wallclock limit.

Using this equation, it can be seen that short running jobs will have an XFactor that will grow much faster over time than the xfactor associated with long running jobs. The following table demonstrates this favoring of short running jobs:

Job Queue Time	1 hour	2 hours	4 hours	8 hours	16 hours
<b>XFactor for 1 hour job</b>	$1 + (1 / 1)$ = 2.00	$1 + (2 / 1)$ = 3.00	$1 + (4 / 1)$ = 5.00	$1 + (8 / 1)$ = 9.00	$1 + (16 / 1)$ = 17.0

Job Queue Time	1 hour	2 hours	4 hours	8 hours	16 hours
<b>XFactor for 4 hour job</b>	$1 + (1 / 4)$ = 1.25	$1 + (2 / 4)$ = 1.50	$1 + (4 / 4)$ = 2.00	$1 + (8 / 4)$ = 3.00	$1 + (16 / 4)$ = 5.0

Since XFactor is calculated as a ratio of two values, it is possible for this subcomponent to be almost arbitrarily large, potentially swamping the value of other priority subcomponents. This can be addressed either by using the subcomponent cap [XFACTORCAP](#), or by using the [XFMINWCLIMIT](#) parameter. If the latter is used, the calculation for the XFactor subcomponent value becomes:

$$\text{XFACTOR} = 1 + \langle \text{EFFQUEUE TIME} \rangle / \text{MAX}(\langle \text{XFMINWCLIMIT} \rangle, \langle \text{WALLCLOCKLIMIT} \rangle)$$

Using the [XFMINWCLIMIT](#) parameter allows a site to prevent very short jobs from causing the XFactor subcomponent to grow inordinately.

Some sites consider XFactor to be a more fair scheduling performance metric than queue time. At these sites, job XFactor is given far more weight than job queue time when calculating job priority and job XFactor distribution consequently tends to be fairly level across a wide range of job durations. (That is, a flat XFactor distribution of 1.0 would result in a one-minute job being queued on average one minute, while a 24-hour job would be queued an average of 24 hours.)

Like queue time, the effective XFactor subcomponent weight is the sum of two weights, the [XFACTORWEIGHT](#) parameter and the QoS-specific [XFWEIGHT](#) setting. For example, the line `QOSCFG[special] XFWEIGHT=5000` causes jobs using the QoS *special* to increase their expansion factor subcomponent weight by 5000.

## Bypass (BYPASS) Subcomponent

The bypass factor is based on the bypass count of a job where the bypass count is increased by one every time the job is bypassed by a lower priority job via backfill. Backfill starvation has never been reported, but if encountered, use the [BYPASS](#) subcomponent.

## StartCount (STARTCOUNT) Subcomponent

Apply the startcount factor to sites with trouble starting or completing due to policies or failures. The primary causes of an idle job having a startcount greater than zero are resource manager level job start failure, administrator based requeue, or requeue based preemption.

## Deadline (DEADLINE) Subcomponent

The deadline factor allows sites to take into consideration the proximity of a job to its [DEADLINE](#). As a jobs moves closer to its deadline its priority increases linearly. This is an alternative to the strict deadline discussed in [QOS SERVICE](#).

## Soft Policy Violation (SPVIOLATION) Subcomponent

The soft policy violation factor allows sites to favor jobs that do not violate their associated [soft resource limit policies](#).

## User Priority (USERPRIO) Subcomponent

The user priority subcomponent allows sites to consider end-user specified job priority in making the overall job priority calculation. Under Moab, end-user specified priorities can only be negative and are bounded in the range 0 to -1024. See [Manual Priority Usage](#) and [Enabling End-user Priorities](#) for more information.

**i** User priorities can be positive, ranging from -1024 to 1023, if `ENABLEPOSUSERPRIORITY TRUE` is specified in `moab.cfg`.

### 4.1.2.F Target Service (TARG) Component

The target factor component of priority takes into account job scheduling performance targets. Currently, this is limited to target expansion factor and target queue time. Unlike the expansion factor and queue time factors described earlier that increase gradually over time, the target factor component is designed to grow exponentially as the target metric is approached. This behavior causes the scheduler to do essentially all in its power to make certain the scheduling targets are met.

The priority calculation for the target factor is as follows:

$$\text{Priority} += \text{TARGETWEIGHT} * ( \text{TARGETQUEUE TIMEWEIGHT} * \text{QueueTimeComponent} + \text{TARGETXFACTORWEIGHT} * \text{XFactorComponent} )$$

The queue time and expansion factor target are specified on a per QoS basis using the `XFTARGET` and `QTTARGET` attributes with the `QOSCFG` parameter. The QueueTime and XFactor component calculations are designed to produce small values until the target value begins to approach, at which point these components grow very rapidly. If the target is missed, this component remains high and continues to grow, but it does not grow exponentially.

### 4.1.2.G Usage (USAGE) Component

The Usage component applies to active jobs only. The priority calculation for the usage priority factor is as follows:

$$\text{Priority} += \text{USAGEWEIGHT} * ( \text{USAGECONSUMEDWEIGHT} * \text{ProcSecondsConsumed} + \text{USAGEHUNGERWEIGHT} * \text{ProcNeededToBalanceDynamicJob} + \text{USAGEREMAININGWEIGHT} * \text{ProcSecRemaining} + )$$

`USAGEEXECUTIONTIMEWEIGHT * SecondsSinceStart +  
USAGEPERCENTWEIGHT * WalltimePercent )`

#### 4.1.2.H Job Attribute (ATTR) Component

The Attribute component allows the incorporation of job attributes into a job's priority. The most common usage for this capability is to do one of the following:

- adjust priority based on a job's state (favor suspended jobs)
- adjust priority based on a job's requested node features (favor jobs that request attribute `pvfs`)
- adjust priority based on internal job attributes (disfavor `backfill` or `preemptee` jobs)
- adjust priority based on a job's requested licenses, network consumption, or generic resource requirements

To use job attribute based prioritization, the `JOBPRIOF` parameter must be specified to set corresponding attribute priorities. To favor jobs based on node feature requirements, the parameter `NODETOJOBATTRMAP` must be set to map node feature requests to job attributes.

The priority calculation for the attribute priority factor is as follows:

```
Priority += ATTRWEIGHT * (
  ATTRATTRWEIGHT * <ATTRPRIORITY> +
  ATTRSTATEWEIGHT * <STATEPRIORITY> +
  ATTRGRESWEIGHT * <GRESPRIORITY>
  JOBIDWEIGHT * <JOBID> +
  JOBNAMWEIGHT * <JOBNAME_INTEGER> )
```

*Example 4-1:*

```
ATTRWEIGHT      100
ATTRATTRWEIGHT  1
ATTRSTATEWEIGHT 1
ATTRGRESWEIGHT  5
# favor suspended jobs
# disfavor preemptible jobs
# favor jobs requesting 'matlab'

JOBPRIOF STATE[Running]=100 STATE[Suspended]=1000 ATTR[PREEMPTEE]=-200 ATTR
[gpfs]=30 GRES[matlab]=400
# map node features to job features

NODETOJOBATTRMAP gpfs,pvfs
...
```

---

## Related Topics

- [Node Allocation Priority](#)
- [Per Credential Priority Weight Offsets](#)
- [Managing Consumable Generic Resources](#)

### 4.1.3 Fairshare Job Priority Example

Consider the following information associated with calculating the fairshare factor for job X.

#### Job X

User A  
Group B  
Account C  
QOS D  
Class E

#### User A

Fairshare Target: 50.0  
Current Fairshare Usage: 45.0

#### Group B

Fairshare Target: [NONE]  
Current Fairshare Usage: 65.0

#### Account C

Fairshare Target: 25.0  
Current Fairshare Usage: 35.0

#### QOS D

Fairshare Target: 10.0+  
Current Fairshare Usage: 25.0

#### Class E

Fairshare Target: [NONE]  
Current Fairshare Usage: 20.0

#### Priority Weights:

FSWEIGHT	100
FSUSERWEIGHT	10
FSGROUPWEIGHT	20
FSACCOUNTWEIGHT	30
FSQOSWEIGHT	40
FSCCLASSWEIGHT	0



In this example, the Fairshare component calculation would be as follows:

$$\begin{aligned} \text{Priority} &+= 100 * ( \\ &10 * 5 + \\ &20 * 0 + \\ &30 * (-10) + \\ &40 * 0 + \\ &0 * 0) \end{aligned}$$

User A is 5% below his target so fairshare increases the total fairshare factor accordingly. Group B has no target so group fairshare usage is ignored. Account C is above its 10% above its fairshare usage target so this component decreases the job's total fairshare factor. QOS D is 15% over its target but the '+' in the target specification indicates that this is a 'floor' target, only influencing priority when fairshare usage drops below the target value. Therefore, the QOS D fairshare usage delta does not influence the fairshare factor.

Fairshare is a great mechanism for influencing job turnaround time via priority to favor a particular distribution of jobs. However, it is important to realize that fairshare can only favor a particular distribution of jobs, it cannot force it. If user X has a fairshare target of 50% of the machine but does not submit enough jobs, no amount of priority favoring will get user X's usage up to 50%.

---

### Related Topics

- [5.3 Fairshare - page 401](#)

## 4.1.4 Common Priority Usage

In this topic:

- [4.1.4.A Credential Priority Factors - page 350](#)
- [4.1.4.B Service Level Priority Factors - page 350](#)
- [4.1.4.C Priority Factor Caps - page 351](#)
- [4.1.4.D User Selectable Prioritization - page 352](#)

Site administrators vary widely in their preferred manner of prioritizing jobs. Moab's scheduling hierarchy allows sites to meet job control needs without requiring adjustments to dozens of parameters. Some choose to use numerous subcomponents, others a few, and still others are content with the default FIFO behavior. Any subcomponent that is not of interest can be safely ignored.

#### 4.1.4.A Credential Priority Factors

To help clarify the use of priority weights, a brief example may help. Suppose a site wanted to maintain the FIFO behavior but also incorporate some credential based prioritization to favor a special user. Particularly, the site would like the user `john` to receive a higher initial priority than all other users. Configuring this behavior requires two steps. First, the user credential subcomponent must be enabled and second, `john` must have his relative priority specified. Take a look at the sample `moab.cfg` file:

```
USERWEIGHT      1
USERCFG[john]   PRIORITY=300
```

**i** The 'USER' priority subcomponent was enabled by setting the `USERWEIGHT` parameter. In fact, the parameters used to specify the weights of all components and subcomponents follow this same '\*WEIGHT' naming convention (as in `RESWEIGHT` and `TARGETQUEUEWEIGHT`).

The second part of the example involves specifying the actual user priority for the user `john`. This is accomplished using the `USERCFG` parameter. Why was the priority 300 selected and not some other value? Is this value arbitrary? As in any priority system, actual priority values are meaningless, only relative values are important. In this case, we are required to balance user priorities with the default queue time based priorities. Since queue time priority is measured in minutes queued, the user priority of 300 places a job by user `john` on par with a job submitted 5 minutes earlier by another user.

Is this what the site wants? Maybe, maybe not. At the onset, most sites are uncertain what they want in prioritization. Often, an estimate initiates prioritization and adjustments occur over time. Cluster resources evolve, the workload evolves, and even site policies evolve, resulting in changing priority needs over time. Anecdotal evidence indicates that most sites establish a relatively stable priority policy within a few iterations and make only occasional adjustments to priority weights from that point.

#### 4.1.4.B Service Level Priority Factors

In another example, suppose a site administrator wants to do the following:

- Favor jobs in the low, medium, and high QoSs so they will run in QoS order
- Balance job expansion factor
- Use job queue time to prevent jobs from starving

Under such conditions, the sample `moab.cfg` file might appear as follows:

```
QOSWEIGHT      1
XFACTORWEIGHT  1
```

```

QUEUETIMEWEIGHT      10
TARGETQUEUETIMEWEIGHT  1
QOSCFG[low]          PRIORITY=1000
QOSCFG[medium]       PRIORITY=10000
QOSCFG[high]         PRIORITY=100000
QOSCFG[DEFAULT]     QTTARGET=4:00:00

```

This example is a bit more complicated but is more typical of the needs of many sites. The desired QoS weightings are established by enabling the QoS subfactor using the `QOSWEIGHT` parameter while the various QoS priorities are specified using `QOSCFG`. `XFACTORWEIGHT` is then set as this subcomponent tends to establish a balanced distribution of expansion factors across all jobs. Next, the `queuetime` component is used to gradually raise the priority of all jobs based on the length of time they have been queued. Note that in this case, `QUEUETIMEWEIGHT` was explicitly set to 10, overriding its default value of 1. Finally, the `TARGETQUEUETIMEWEIGHT` parameter is used in conjunction with the `USERCFG` line to specify a queue time target of 4 hours.

#### 4.1.4.C Priority Factor Caps

Assume now that the site administrator is content with this priority mix but has a problem with users submitting large numbers of very short jobs. Very short jobs would tend to have rapidly growing XFactor values and would consequently quickly jump to the head of the queue. In this case, a factor cap would be appropriate. Such caps allow a site to limit the contribution of a job's priority factor to be within a defined range. This prevents certain priority factors from swamping others. Caps can be applied to either priority components or subcomponents and are specified using the `<COMPONENTNAME>CAP` parameter (such as `QUEUETIMECAP`, `RESCAP`, and `SERVCAP`). Note that both component and subcomponent caps apply to the preweighted value, as in the following equation:

```

Priority =
  C1WEIGHT * MIN(C1CAP, SUM(
    S11WEIGHT * MIN(S11CAP, S11S) +
    S12WEIGHT * MIN(S12CAP, S12S) +
    ...)) +
  C2WEIGHT * MIN(C2CAP, SUM(
    S21WEIGHT * MIN(S21CAP, S21S) +
    S22WEIGHT * MIN(S22CAP, S22S) +
    ...)) +
  ...

```

#### Example 4-2: Priority cap

```

QOSWEIGHT      1
QOSCAP         10000
XFACTORWEIGHT  1
XFACTORCAP     1000
QUEUETIMEWEIGHT 10
QUEUETIMECAP   1000

```

#### 4.1.4.D User Selectable Prioritization

Moab allows users to specify a job priority to jobs they own or manage. This priority can be set at job submission time or it can be dynamically modified (using [setspri](#) or [mjobctl](#)) after submitting the job. For fairness reasons, users can only apply a negative priority to their job and therefore slide it further back in the queue. This enables users to allow their more important jobs to run before their less important ones without gaining unfair advantage over other users.

**i** User priorities can be positive if `ENABLEPOSUSERPRIORITY TRUE` is specified in `moab.cfg`.

In order to set `ENABLEPOSUSERPRIORITY`, you must change the `USERPRIOWEIGHT` from its default value of 0. For example:

```
USERPRIOWEIGHT 100
```

```
> setspri -r 100 332411
successfully modified job priority
```

**i** Specifying a user priority at job submission time is resource manager specific. See the associated resource manager documentation for more information.

### User Selectable Priority w/QoS

Using the [QoS](#) facility, organizations can set up an environment where users can more freely select the desired priority of a given job. Organizations can enable access to a number of QoSs each with its own charging rate, priority, and target service levels. Users can then assign job importance by selecting the appropriate QoS. If desired, this can allow a user to jump ahead of other users in the queue if they are willing to pay the associated costs.

---

#### Related Topics

- [User Selectable Priority](#)

### 4.1.5 Prioritization Strategies

Each component or subcomponent can be used to accomplish different objectives. `WALLTIME` can be used to favor (or disfavor) jobs based on their duration. Likewise,

ACCOUNT can be used to favor jobs associated with a particular project while QUEUETIME can be used to favor those jobs waiting the longest.

- Queue Time
- Expansion Factor
- Resource
- Fairshare
- Credential
- Target Metrics

Each priority factor group can contain one or more subfactors. For example, the Resource factor consists of Node, Processor, Memory, Swap, Disk, and PE components. From the table in [Job Priority Factors](#) section, it is apparent that the prioritization problem is fairly complex since every site needs to prioritize a bit differently. When calculating a priority, the various priority factors are summed and then bounded between 0 and MAX\_PRIO\_VAL, which is currently defined as 100000000 (one billion).

The `mdiag -p` command assists with visualizing the priority distribution resulting from the current job priority configuration. Also, the `showstats -f` command helps indicate the impact of the current priority settings on scheduler service distributions.

## 4.1.6 Manual Job Priority Adjustment

Batch administrators regularly find a need to adjust the calculated priority of a job to meet current needs. Current needs often are broken into two categories:

1. The need to run an administrator test job as soon as possible.
2. The need to pacify a disserved user.

You can use the `setspri` command to handle these issues in one of two ways; this command allows the specification of either a relative priority adjustment or the specification of an absolute priority. Using absolute priority specification, administrators can set a job priority guaranteed to be higher than any calculated value. Where Moab-calculated job priorities are in the range of 0 to 1 billion, system administrator assigned absolute priorities start at 1 billion and go up. Issuing the `setspri <PRIO> <JOBID>` command, for example, assigns a priority of 1 billion + <PRIO> to the job. Therefore, `setspri 5 job.1294` sets the priority of 'job.1294' to 1000000005.

For more information, see [Common Priority Usage - End-user Adjustment](#).

## 4.2 Node Allocation Policies

While job prioritization allows a site to determine which job to run, node allocation policies allow a site to specify how available resources should be allocated to each job. The algorithm used is specified by the parameter `NODEALLOCATIONPOLICY`. There are multiple node allocation policies to choose from allowing selection based on reservation constraints, node configuration, resource usage, preferences, and other factors. You can specify these policies with a system-wide default value, on a per-partition basis, or on a per-job basis. Note that `LASTAVAILABLE` is the default policy.

Available algorithms are described in detail in the following sections and include `CONTIGUOUS`, `CPULOAD`, `FIRSTAVAILABLE`, `LASTAVAILABLE`, `MINRESOURCE`, `MAXBALANCE`, `PLUGIN`, `PRIORITY`.

In this topic:

- [4.2.1 Node Allocation Overview - page 354](#)
- [4.2.2 Node Selection Factors - page 358](#)
- [4.2.3 Resource-Based Algorithms - page 358](#)
- [4.2.4 User-Defined Algorithms - page 363](#)
- [4.2.5 Specifying Per Job Resource Preferences - page 364](#)

### 4.2.1 Node Allocation Overview

Node allocation is the process of selecting the best resources to allocate to a job from a list of available resources. Making this decision intelligently is important in an environment that possesses one or more of the following attributes:

- Heterogeneous resources (resources which vary from node to node in terms of quantity or quality)
- Shared nodes (nodes can be utilized by more than one job)
- Reservations or service guarantees
- Non-flat network (a network where a perceptible performance degradation may potentially exist depending on workload placement)

#### 4.2.1.A Heterogeneous Resources

Moab analyzes job processing requirements and assigns resources to maximize hardware utility.

For example, suppose two nodes are available in a system, A and B. Node A has 768 MB of RAM and node B has 512 MB. The next two jobs in the queue are X and Y. Job X requests 256 MB and job Y requests 640 MB. Job X is next in the queue and can fit on either node, but Moab recognizes that job Y (640 MB) can only fit on node A (768 MB). Instead of putting job X on node A and blocking job Y, Moab can put job X on node B and job Y on node A.

### 4.2.1.B Shared Nodes

#### Symmetric Multiprocessing (SMP)

When sharing SMP-based compute resources amongst tasks from more than one job, resource contention and fragmentation issues arise. In SMP environments, the general goal is to deliver maximum system utilization for a combination of compute-intensive and memory-intensive jobs while preventing overcommitment of resources.

By default, most current systems do not do a good job of logically partitioning the resources (such as CPU, memory, and network bandwidth) available on a given node. Consequently contention often arises between tasks of independent jobs on the node. This can result in a slowdown for all jobs involved, which can have significant ramifications if large-way parallel jobs are involved. Virtualization, CPU sets, and other techniques are maturing quickly as methods to provide logical partitioning within shared resources.

On large-way SMP systems (> 32 processors/node), job packing can result in intra-node fragmentation. For example, take two nodes, A and B, each with 64 processors. Assume they are currently loaded with various jobs and A has 24 and B has 12 processors free. Two jobs are submitted; job X requests 10 processors and job Y requests 20 processors. Job X can start on either node but starting it on node A prevents job Y from running. An algorithm to handle intra-node fragmentation is straightforward for a single resource case, but the algorithm becomes more involved when jobs request a combination of processors, memory, and local disk. These workload factors should be considered when selecting a site's node allocation policy, as well as identifying appropriate policies for handling resource utilization limit violations.

#### Interactive Nodes

In many cases, sites are interested in allowing multiple users to simultaneously use one or more nodes for interactive purposes. Workload is commonly not compute intensive consisting of intermittent tasks including coding, compiling, and testing. Because these jobs are highly variant in terms of resource usage over time, sites are able to pack a larger number of these jobs onto the same node. Consequently, a common practice is to restrict job scheduling based on utilized, rather than dedicated resources.

## Interactive Node Example

The example configuration files that follow show one method by which node sharing can be accomplished within a Torque + Moab environment. This example is based on a hypothetical cluster composed of 4 nodes each with 4 cores. For the compute nodes, job tasks are limited to actual cores preventing overcommitment of resources. For the interactive nodes, up to 32 job tasks are allowed, but the node also stops allowing additional tasks if either memory is fully utilized or if the CPU load exceeds 4.0. Therefore, Moab continues packing the interactive nodes with jobs until carrying capacity is reached.

*Example 4-3: /opt/moab/etc/moab.cfg*

```
# constrain interactive jobs to interactive nodes
# constrain interactive jobs to 900 proc-seconds
CLASSCFG[interactive]  HOSTLIST=interactive01,interactive02
CLASSCFG[interactive]  MAX.CPUTIME=900
RESOURCELIMITPOLICY   CPUTIME:ALWAYS:CANCEL
# base interactive node allocation on load and jobs
NODEALLOCATIONPOLICY   PRIORITY
NODECFG[interactive01] PRIORITYF='-20*LOAD - JOBCOUNT'
NODECFG[interactive02] PRIORITYF='-20*LOAD - JOBCOUNT'
```

*Example 4-4: /var/spool/torque/server\_priv/nodes*

```
interactive01 np=32
interactive02 np=32
compute01    np=4
compute02    np=4
```

*Example 4-5: /var/spool/torque/mom\_priv/config on "interactive01"*

```
# interactive01
$max_load 4.0
```

*Example 4-6: /var/spool/torque/mom\_priv/config on "interactive02"*

```
# interactive02
$max_load 4.0
```


### 4.2.1.C Reservations or Service Guarantees

A reservation-based system adds the time dimension into the node allocation decision. With reservations, node resources must be viewed in a type of two dimension node-time space. Allocating nodes to jobs fragments this node-time space and makes it more difficult to schedule jobs in the remaining, more constrained node-time slots. Allocation decisions should be made in such a way as to minimize this fragmentation and maximize the scheduler's ability to continue to start jobs in existing slots. The following figure shows that job A and job B are running. A reservation, X, is created some time in the future. Assume that job A is 2 hours long and job B is 3 hours long. Again, two new single-processor jobs are submitted, C and D; job C requires 3 hours of compute time while job D requires 5



hours. Either job will just fit in the free space located above job A or in the free space located below job B. If job C is placed above job A, job D, requiring 5 hours of time will be prevented from running by the presence of reservation X. However, if job C is placed below job B, job D can still start immediately above job A.

*Image 4-1: Job A, Job B, and Reservation X scheduled on nodes*



The preceding example demonstrates the importance of time based reservation information in making node allocation decisions, both at the time of starting jobs and at the time of creating reservations. The impact of time based issues grows significantly with the number of reservations in place on a given system. The `LASTAVAILABLE` algorithm works on this premise, locating resources that have the smallest space between the end of a job under consideration and the start of a future reservation.

#### 4.2.1.D Non-Flat Network

On systems where network connections do not resemble a flat all-to-all topology, task placement may impact performance of communication intensive parallel jobs. If latencies and network bandwidth between any two nodes vary significantly, the node allocation algorithm should attempt to pack tasks of a given job as close to each other as possible to minimize impact of bandwidth and latency differences.

## 4.2.2 Node Selection Factors

While the node allocation policy determines which nodes a job will use, other factors narrow the options before the policy makes the final decision. The following process demonstrates how Moab executes its node allocation process and how other policies affect the decision:

1. Moab eliminates nodes that do not meet the hard resource requirements set by the job.
2. Moab gathers affinity information, first from workload proximity rules and then from reservation affinity rules (see [Affinity](#) for more information). Reservation affinity rules trump workload proximity rules.
3. Moab allocates nodes using the allocation policy:
  - If more than enough nodes with Required affinity exist, only they are passed down for the final sort by the node allocation policy.
  - If the number of nodes with Required affinity matches the number of nodes requested exactly, then the node allocation policy is skipped entirely and all of those nodes are assigned to the job.
  - If too few nodes have Required affinity, all of them are assigned to the job, then the node allocation policy is applied to the remaining eligible nodes (after Required, Moab will use Positive, then Neutral, then Negative).

## 4.2.3 Resource-Based Algorithms

Moab contains a number of allocation algorithms that address some of the needs described earlier. You can also create allocation algorithms and interface them with the Moab scheduling system. Each of these policies has a name and descriptive alias. They can be configured using either one, but Moab will only report their names.

**i** If `ENABLEHIGHTHROUGHPUT` is `TRUE`, you must set `NODEALLOCATIONPOLICY` to `FIRSTAVAILABLE`.

The current suite of algorithms is described in the table below:

*Allocation Algorithms*

Allocation Algorithm Name	Alias	Description
<b>CONTIGUOUS</b>	Contiguous	Allocates nodes in contiguous (linear) blocks as required by the Compaq RMS system.
<b>CPULOAD</b>	ProcessorLoad	Nodes are selected that have the maximum amount of available, unused CPU power (<#of CPUs> - <CPU load>). CPULOAD is a good algorithm for timesharing node systems and applies to jobs starting immediately. For the purpose of future reservations, the MINRESOURCE algorithm is used.
<b>FIRSTAVAILABLE</b>	InReportedOrder	Simple first come, first served algorithm where nodes are allocated in the order they are presented by the resource manager. This is a very simple, and very fast algorithm.
<b>LASTAVAILABLE</b>	InReserveReportedOrder	Nodes are allocated in descending order that they are presented by the resource manager, or the reverse of FIRSTAVAILABLE.
<b>MAXBALANCE</b>	ProcessorSpeedBalance	Attempts to allocate the most balanced set of nodes possible to a job. In most cases, but not all, the metric for balance of the nodes is node procspeed. Therefore, if possible, nodes with identical procspeeds are allocated to the job. If identical procspeed nodes cannot be found, the algorithm allocates the set of nodes with the minimum node procspeed span or range.
<b>MINRESOURCE</b>	MinimumConfiguredResources	Prioritizes nodes according to the configured memory resources on each node. Those nodes with the fewest configured memory resources, that still meet the job's resource constraints, are selected.
<b>PRIORITY</b>	CustomPriority	Allows a site to specify the priority of various static and dynamic aspects of compute nodes and allocate them with preference for higher priority nodes. It is highly flexible allowing node attribute and usage information to be combined

Allocation Algorithm Name	Alias	Description
		<p>with reservation affinity. Using node allocation priority, you can specify the following priority components:</p> <ul style="list-style-type: none"> <li>• ADISK - Local disk currently available to batch jobs in MB.</li> <li>• AMEM - Real memory currently available to batch jobs in MB.</li> <li>• APROCS - Processors currently available to batch jobs on node (configured procs - dedicated procs).</li> <li>• ARCH [&lt;ARCH&gt;] - Processor architecture.</li> </ul>
		<ul style="list-style-type: none"> <li>• ASWAP - Virtual memory currently available to batch jobs in MB.</li> <li>• CDISK - Total local disk allocated for use by batch jobs in MB.</li> <li>• CMEM - Total real memory on node in MB.</li> <li>• COST - Based on node CHARGERATE.</li> <li>• CPROCS - Total processors on node.</li> </ul>
		<ul style="list-style-type: none"> <li>• CSWAP - Total virtual memory configured on node in MB.</li> <li>• FEATURE [&lt;FNAME&gt;] - Boolean; specified feature is present on node.</li> <li>• FREETIME - FREETIME is calculated as the time during which there is no reservation on the machine. It uses either the job wallclock limit (if there is a job), or 2 months. The more free time a node has within either the job wallclock limit or 2 months, the higher this value will be.</li> <li>• GMETRIC [&lt;GMNAME&gt;] - Current value of specified generic metric on node.</li> <li>• JOBCOUNT - Number of jobs currently running on node.</li> </ul>
		<ul style="list-style-type: none"> <li>• JOBFREETIME - The number of seconds that the node is idle between now and when the job is scheduled to start.</li> <li>• LOAD - Current 1 minute load average.</li> <li>• MTBF - Mean time between failures (in</li> </ul>

Allocation Algorithm Name	Alias	Description
		<p>seconds).</p> <ul style="list-style-type: none"> <li>• NODEINDEX - Node's nodeindex as specified by the resource manager.</li> <li>• OS - True if job compute requirements match node operating system.</li> </ul>
		<ul style="list-style-type: none"> <li>• PARAPROCS - Processors currently available to batch jobs within partition (configured procs - dedicated procs).</li> <li>• POWER - TRUE if node is ON.</li> <li>• PREF - Boolean; node meets job specific resource preferences.</li> <li>• PRIORITY - Administrator specified node priority.</li> <li>• RANDOM - Per iteration random value between 0 and 1. (Allows introduction of random allocation factor.)</li> </ul> <div data-bbox="868 961 1414 1287" style="border: 1px solid #0056b3; border-radius: 10px; padding: 10px; background-color: #e6f2ff;"> <p><b>i</b> Regardless of coefficient, the contribution of this weighted factor cannot exceed 32768.</p> <p>The coefficient, if any, of the RANDOM component must precede, not follow, the component in order to work correctly. For example:</p> <div style="border: 1px dashed #ccc; border-radius: 10px; padding: 5px; width: fit-content; margin: 5px auto;"> <math>100 * \text{RANDOM}</math> </div> </div>
		<ul style="list-style-type: none"> <li>• SPEED - If set, node processor speed (procspeed); otherwise, relative node speed.</li> <li>• SUSPENDEDJCOUNT - Number of suspended jobs currently on the node.</li> <li>• USAGE - Percentage of time node has been running batch jobs since the last statistics initialization.</li> <li>• WINDOWTIME - The window of time between the end of one reservation and the beginning of another. This algorithm, given a negative value, can be used to pack reservations as close together on a node as possible.</li> </ul>

Allocation Algorithm Name	Alias	Description
		<p>The node allocation priority function can be specified on a node by node or cluster wide basis. In both cases, the recommended approach is to specify the <b>PRIORITYF</b> attribute with the <b>NODECFG</b> parameter. Some examples follow.</p>
		<p><i>Example 1:</i> Favor the fastest nodes with the most available memory that are running the fewest jobs.</p> <pre data-bbox="812 703 1412 819"> NODEALLOCATIONPOLICY PRIORITY NODECFG[DEFAULT] PRIORITYF='SPEED + .01 * AMEM - 10 * JOBCOUNT' ... </pre> <div data-bbox="812 840 1412 997" style="border: 1px solid black; padding: 5px;"> <p><b>i</b> If spaces are placed within the priority function for readability, the priority function value must be quoted to allow proper parsing.</p> </div>
		<p><i>Example 2:</i> Favor the nodes with the least amount of idle time between now and the job's scheduled start time.</p> <pre data-bbox="812 1165 1412 1228"> NODEALLOCATIONPOLICY PRIORITY NODECFG[DEFAULT] PRIORITYF=-JOBFREETIME </pre> <div data-bbox="860 1239 1364 1354" style="border: 1px dashed black; padding: 5px;"> <p><i>Moab stacks jobs on the nodes that are busiest between now and the job's scheduled start time.</i></p> </div>
		<p><i>Example 3:</i> A site has a batch system consisting of two dedicated 'batchX' nodes, as well as numerous desktop systems. The allocation function should favor batch nodes first, followed by desktop systems that are the least loaded and have received the least historical usage.</p> <pre data-bbox="812 1606 1412 1795"> NODEALLOCATIONPOLICY PRIORITY NODECFG[DEFAULT] PRIORITYF='-LOAD - 5*USAGE' NODECFG[batch1] PRIORITY=1000 PRIORITYF='PRIORITY + APROCS' NODECFG[batch2] PRIORITY=1000 PRIORITYF='PRIORITY + APROCS' ... </pre>

Allocation Algorithm Name	Alias	Description
		<p><i>Example 4: Pack tasks onto loaded nodes first.</i></p> <pre>NODEALLOCATIONPOLICY PRIORITY NODECFG[DEFAULT] PRIORITYF=JOBCOUNT ...</pre>
		<p><i>Example 5: Pack tasks onto nodes with the most processors available and the lowest CPU temperature.</i></p> <pre>RMCFG[torque] TYPE=pbs RMCFG[temp] TYPE=NATIVE CLUSTERQUERYURL=exec://\$TOOLS_DIR/hwmon.pl NODEALLOCATIONPOLICY PRIORITY NODECFG[DEFAULT] PRIORITYF='100*APROCS - GMETRIC [temp] ' ...</pre>

## 4.2.4 User-Defined Algorithms

User-defined algorithms allow administrators to define their own algorithms based on factors such as their system's network topology. When node allocation is based on topology, jobs finish faster, administrators see better cluster productivity and users pay less for resources.

### 4.2.4.A PLUGIN

This algorithm allows administrators to define their own node allocation policy and create a plug-in that allocates nodes based on factors such as a cluster's network topology. This has the following advantages:

- plug-ins keep the source code of the cluster's interconnect network for node allocation separate from Moab's source code (customers can implement plug-ins independent of Moab's release schedule)
- plug-ins can be independently created and tailored to specific hardware and network topology
- plug-ins can be modified without assistance from Adaptive Computing

## 4.2.5 Specifying *Per Job* Resource Preferences

While the resource based node allocation algorithms can make a good guess at what compute resources would best satisfy a job, sites often possess a subset of jobs that benefit from more explicit resource allocation specification. For example, one job might perform best on a particular subset of nodes due to direct access to a tape drive, another might be very memory intensive. Resource preferences are distinct from node requirements. While the former describes what a job needs to run at all, the latter describes what the job needs to run well. In general, a scheduler must satisfy a job's node requirement specification and then satisfy the job's resource preferences as well as possible.

### 4.2.5.A Specifying Resource Preferences

A number of resource managers natively support the concept of resource preferences (such as Loadleveler). When using these systems, the language specific preferences keywords can be used. For systems that do not support resource preferences natively, Moab provides a [resource manager extension](#) keyword 'PREF', which you can use to specify desired resources. This extension allows specification of node features, memory, swap, and disk space conditions that define whether the node is considered preferred.

**i** Moab 5.2 (and earlier) only supports feature-based preferences.

### 4.2.5.B Selecting Preferred Resources

Enforcing resource preferences is not completely straightforward. A site might have a number of potentially conflicting requirements that the scheduler is asked to simultaneously satisfy. For example, a scheduler may be asked to maximize the proximity of the allocated nodes at the same time it is supposed to satisfy resource preferences and minimize node overcommitment. To allow site specific weighting of these varying requirements, Moab allows resource preferences to be enabled through the [PRIORITY](#) node allocation algorithm. For example, to use resource preferences together with node load, the following configuration might be used:

```
NODEALLOCATIONPOLICY PRIORITY
NODECFG[DEFAULT]    PRIORITYF='5 * PREF - LOAD'
...
```

To request specific resource preferences, a user could then submit a job indicating those preferences. In the case of a PBS job, the following can be used:

```
> qsub -l nodes=4,walltime=1:00:00,pref=feature:fast
```



**Related Topics**

- [Generic Metrics](#)
- Per Job Node Allocation Policy Specification via [Resource Manager Extensions](#)

## 4.3 Node Access Policies

Moab allocates resources to jobs on the basis of a job task—an atomic collection of resources that must be co-located on a single compute node. A given job may request 20 tasks where each task is defined as one processor and 128 MB of RAM. Compute nodes with multiple processors often possess enough resources to support more than one task simultaneously. When it is possible for more than one task to run on a node, node access policies determine which tasks can share the compute node's resources.

In this topic:

[4.3.1 Node Access Policy Descriptions - page 365](#)

[4.3.2 Configuring Node Access Policies - page 366](#)

### 4.3.1 Node Access Policy Descriptions

Moab supports a distinct number of node access policies that are listed in the following table:

Policy	Description
<b>SHARED</b>	Tasks from any combination of jobs can use available resources.
<b>SHAREDONLY</b>	Only jobs requesting shared node access can use available resources.
<b>SINGLEACCOUNT</b>	Tasks from any jobs owned by the same account can use available resources.
<b>SINGLECLASS</b>	Tasks from any jobs owned by the same class can use available resources.
<b>SINGLEGROUP</b>	Tasks from any jobs owned by the same group can use available resources.
<b>SINGLEJOB</b>	Only tasks from a single job can use the node's resources.

Policy	Description
	<p><b>i</b> When enforcing limits using <code>CLASSCFG</code> attributes, use <code>MAX.NODE</code> instead of <code>MAX.PROC</code>. <code>MAX.PROC</code> enforces the requested processors, not the actual processors dedicated to the job.</p>
<b>SINGLETASK</b>	Only a single task from a single job can run on the node.
<b>SINGLEUSER</b>	Tasks from any jobs owned by the same user can use available resources.
<b>UNIQUEUSER</b>	<p>Any number of tasks from a single job can allocate resources from a node but only if the user has no other jobs running on that node. <code>UNIQUEUSER</code> limits the number of jobs a single user can run on a node, allowing other users to run jobs with the remaining resources.</p> <p><b>i</b> This policy is useful in environments where job epilog/prologs scripts are used to clean up processes based on <code>userid</code>.</p>

### 4.3.2 Configuring Node Access Policies

The global node access policies can be specified via the parameter `NODEACCESSPOLICY`. This global default can be overridden on a per node basis with the `ACCESS` attribute of the `NODECFG` parameter or on a per job basis using the resource manager extension `NACCESSPOLICY`. Finally, a per queue node access policy can also be specified by setting either the `NODEACCESSPOLICY` or `FORCENODEACCESSPOLICY` attributes of the `CLASSCFG` parameter. `FORCENODEACCESSPOLICY` overrides any per job specification in all cases, whereas `NODEACCESSPOLICY` is overridden by per job specification.

**i** When multiple node access policies apply to a given job or node (for example `SINGLEJOB` is configured globally but the class is configured as `SHARED`) then the more restrictive policy applies. The most restrictive policy is `SINGLETASK`, followed by `SINGLEJOB`, the single credentials, and `SHARED` being the least restrictive.

By default, nodes are accessible using the setting of the system wide `NODEACCESSPOLICY` parameter unless a specific `ACCESS` policy is specified on a per node basis using the `NODECFG` parameter. Jobs can override this policy and subsequent jobs are bound to conform to the access policies of all jobs currently running on a given node. For example, if the `NODEACCESSPOLICY` parameter is set to `SHARED`, a new job can be launched on an idle node with a job specific access policy of `SINGLEUSER`. While this job runs, the effective node access policy changes to `SINGLEUSER` and subsequent job tasks can only be launched on this node provided they are submitted by the same user.

When all single user jobs have completed on that node, the effective node access policy reverts back to `SHARED` and the node can again be used in `SHARED` mode.

For example, to set a global policy of `SINGLETASK` on all nodes except nodes 13 and 14, use the following:

```
# by default, enforce dedicated node access on all nodes
NODEACCESSPOLICY SINGLETASK
# allow nodes 13 and 14 to be shared
NODECFG[node13] ACCESS=SHARED
NODECFG[node14] ACCESS=SHARED
```

You can also set `SINGLEJOB` using the `qsub` `node-exclusive` option (`-n`). For example:

```
qsub -n jobscript.sh
```

*This will set `node_exclusive = True` in the output for `qstat -f <job id>`.*

Alternatively, you could also use either of the following:

```
qsub -l naccesspolicy=singlejob jobscript.sh
qsub -W x=naccesspolicy:singlejob jobscript.sh
```

---

## Related Topics

- Per job [naccesspolicy](#) specification via [Resource Manager Extensions](#)
- [JOBNODEMATCHPOLICY](#) parameter
- [NODEAVAILABILITYPOLICY](#) parameter

## 4.4 Node Availability Policies

In this topic:

- [4.4.1 Node Resource Availability Policies - page 368](#)
- [4.4.2 Node Categorization - page 369](#)
- [4.4.3 Node Failure/Performance Based Notification - page 371](#)
- [4.4.4 Node Failure/Performance Based Triggers - page 371](#)
- [4.4.5 Handling Transient Node Failures - page 371](#)
- [4.4.6 Allocated Resource Failure Policy for Jobs - page 373](#)

Moab enables several features relating to node availability. These include policies that determine how per node resource availability should be reported, how node failures are detected, and what should be done in the event of a node failure.

## 4.4.1 Node Resource Availability Policies

Moab allows a job to be launched on a given compute node as long as the node is not full or busy. The `NODEAVAILABILITYPOLICY` parameter allows a site to determine what criteria constitute a node being busy. The legal settings are listed in the following table:

Availability Policy	Description
<b>DEDICATED</b>	The node is considered busy if dedicated resources equal or exceed configured resources.
<b>UTILIZED</b>	The node is considered busy if utilized resources equal or exceed configured resources.
<b>COMBINED</b>	The node is considered busy if either dedicated or utilized resources equal or exceed configured resources.

The default setting for all nodes is `COMBINED`, indicating that a node can accept workload so long as the jobs that the node was allocated to do not request or use more resources than the node has available. In a load balancing environment, this might not be the desired behavior. Setting the `NODEAVAILABILITYPOLICY` parameter to `UTILIZED` allows jobs to be packed onto a node even if the aggregate resources requested exceed the resources configured. For example, assume a scenario with a 4-processor compute node and 8 jobs requesting 1 processor each. If the resource availability policy was set to `COMBINED`, this node would only allow 4 jobs to start on this node even if the jobs induced a load of less than 1.0 each. With the resource availability policy set to `UTILIZED`, the scheduler continues allowing jobs to start on the node until the node's load average exceeds a per processor load value of 1.0 (in this case, a total load of 4.0). To prevent a node from being over populated within a single scheduling iteration, Moab artificially raises the node's load for one scheduling iteration when starting a new job. On subsequent iterations, the actual measured node load information is used.

### 4.4.1.A Per Resource Availability Policies

By default, the `NODEAVAILABILITYPOLICY` sets a global per node resource availability policy. This policy applies to all resource types on each node such as processors, memory, swap, and local disk. However, the syntax of this parameter is as follows:

```
<POLICY> [ : <RESOURCETYPE> ] . . .
```

This syntax allows per resource availability specification. For example, consider the following:

```
NODEAVAILABILITYPOLICY DEDICATED:PROC COMBINED:MEM COMBINED:DISK
...
```

This configuration causes Moab to only consider the quantity of processing resources actually dedicated to active jobs running on each node and ignore utilized processor information (such as CPU load). For memory and disk, both utilized resource information and dedicated resource information should be combined to determine what resources are actually available for new jobs.

## 4.4.2 Node Categorization

Moab allows organizations to detect and use far richer information regarding node status than the standard batch 'idle,' 'busy,' 'down states' commonly found. Using node categorization, organizations can record, track, and report on per node and cluster level status including the following categories:

Category	Description
<b>Active</b>	Node is healthy and currently executing batch workload.
<b>BatchFailure</b>	Node is unavailable due to a failure in the underlying batch system (such as a resource manager server or resource manager node daemon).
<b>Benchmark</b>	Node is reserved for benchmarking.
<b>EmergencyMaintenance</b>	Node is reserved for unscheduled system maintenance.
<b>GridReservation</b>	Node is reserved for grid use.
<b>HardwareFailure</b>	Node is unavailable due to a failure in one or more aspects of its hardware configuration (such as a power failure, excessive temperature, memory, processor, or swap failure).
<b>HardwareMaintenance</b>	Node is reserved for scheduled system maintenance.
<b>Idle</b>	Node is healthy and is currently not executing batch workload.
<b>JobReservation</b>	Node is reserved for job use.
<b>NetworkFailure</b>	Node is unavailable due to a failure in its network adapter or in the switch.

Category	Description
<b>Other</b>	Node is in an uncategorized state.
<b>OtherFailure</b>	Node is unavailable due to a general failure.
<b>PersonalReservation</b>	Node is reserved for dedicated use by a personal reservation.
<b>Site[1-8]</b>	Site specified usage categorization.
<b>SoftwareFailure</b>	Node is unavailable due to a failure in a local software service (such as automounter, security or information service such as NIS, local databases, or other required software services).
<b>SoftwareMaintenance</b>	Node is reserved for software maintenance.
<b>StandingReservation</b>	Node is reserved by a standing reservation.
<b>StorageFailure</b>	Node is unavailable due to a failure in the cluster storage system or local storage infrastructure (such as failures in Lustre, GPFS, PVFS, or SAN).
<b>UserReservation</b>	Node is reserved for dedicated use by a particular user or group and may or may not be actively executing jobs.

Node categories can be explicitly assigned by cluster administrators using the `mrsvctl -c` command to create a reservation and associate a category with that node for a specified timeframe. Further, outside of this explicit specification, Moab automatically mines all configured interfaces to learn about its environment and the health of the resources it is managing. Consequently, Moab can identify many hardware failures, software failures, and batch failures without any additional configuration. However, it is often desirable to make additional information available to Moab to allow it to integrate this information into reports; automatically notify managers, users, and administrators; adjust internal policies to steer workload around failures; and launch various custom [triggers](#) to rectify or mitigate the problem.

**i** You can specify the `FORCERSVSTYPE` parameter to require all administrative reservations be associated with a node category at reservation creation time. For example:

```
NODECFG [DEFAULT] ENABLEPROFILING=TRUE
FORCERSVSTYPE TRUE
```

Node health and performance information from external systems can be imported into Moab using the [native resource manager interface](#). This is commonly done using [generic metrics](#) or [consumable generic resources](#) for performance and node categories or node variables for status information. Combined with arbitrary node messaging information, Moab can combine detailed information from remote services and report this to other external services.

**i** Use the [NODECATCREDLIST](#) parameter to generate extended node category based statistics.

### 4.4.3 Node Failure/Performance Based Notification

Moab can be configured to cause node failures and node performance levels that cross specified thresholds to trigger notification events. This is accomplished using the [GEVENTCFG](#) parameter as described in the [Generic Event Overview](#) section. For example, the following configuration can be used to trigger an email to administrators each time a node is marked down:

```
GEVENTCFG[nodedown] ACTION=notify REARM=00:20:00
...
```

### 4.4.4 Node Failure/Performance Based Triggers

Moab supports per node triggers that can be configured to fire when specific events are fired or specific thresholds are met. These triggers can be used to modify internal policies or take external actions. A few examples follow:

- decrease node allocation priority if node throughput drops below threshold X
- launch local diagnostic/recovery script if parallel file system mounts become stale
- reset high performance network adapters if high speed network connectivity fails
- create general system reservation on node if processor or memory failure occurs

As mentioned, Moab triggers can be used to initiate almost any action, from sending mail to updating a database, to publishing data for an SNMP trap, to driving a web service.

### 4.4.5 Handling Transient Node Failures

Since Moab actively schedules both current and future actions of the cluster, it is often important for it to have a reasonable estimate of when failed nodes will be again available for use. This knowledge is particularly useful for proper scheduling of new jobs and management of resources in regard to [backfill](#). With backfill, Moab determines which

resources are available for priority jobs and when the highest priority idle jobs can run. If a node experiences a failure, Moab should have a concept of when this node will be restored.

When Moab analyzes [down](#) nodes for allocation, one of two issues may occur with the highest priority jobs. If Moab believes that down nodes will not be recovered for an extended period of time, a transient node failure within a reservation for a priority job might cause the reservation to slide far into the future allowing other lower priority jobs to allocate and launch on nodes previously reserved for it. Moments later, when the transient node failures are resolved, Moab might be unable to restore the early reservation start time as other jobs may already have been launched on previously available nodes.

In the reverse scenario, if Moab recognizes a likelihood that down nodes will be restored too quickly, it might make reservations for top priority jobs that allocate those nodes. Over time, Moab slides those reservations further into the future as it determines that the reserved nodes are not being recovered. While this does not delay the start of the top priority jobs, these unfulfilled reservations can end up blocking other jobs that should have properly been backfilled and executed.

#### 4.4.5.A Creating Automatic Reservations

If a node experiences occasional transient failures (often not associated with a node state of down), Moab can automatically create a temporary reservation over the node to allow the transient failure time to clear and prevent Moab from attempting to re-use the node while the failure is active. This reservation behavior is controlled using the [NODEFAILURERESERVETIME](#) parameter as in the following example:

```
# reserve nodes for 1 minute if transient failures are detected
NODEFAILURERESERVETIME 00:01:00
```

#### 4.4.5.B Blocking Out Down Nodes

If one or more resource managers identify failures and mark nodes as down, Moab can be configured to associate a default *unavailability* time with this failure and the node state *down*. This is accomplished using the [NODEDOWNSTATEDELAYTIME](#) parameter. This delay time floats and is measured as a fixed time into the future from the time 'NOW'; it is not associated with the time the node was originally marked down. For example, if the delay time was set to 10 minutes, and a node was marked down 20 minutes ago, Moab would still consider the node unavailable until 10 minutes into the future.

While it is difficult to select a good default value that works for all clusters, the following is a general rule of thumb:



- Increase `NODEDOWNSTATEDELAYTIME` if jobs are getting blocked due to priority reservations sliding as down nodes are not recovered.
- Decrease `NODEDOWNSTATEDELAYTIME` if high priority job reservations are getting regularly delayed due to transient node failures.

```
# assume down nodes will not be recovered for one hour
NODEDOWNSTATEDELAYTIME 01:00:00
```

## 4.4.6 Allocated Resource Failure Policy for Jobs

If a failure occurs within a collection of nodes allocated to a job, Moab can automatically reallocate replacement resources. This can be configured with `JOBACTIONONNODEFAILURE`.

How an active job behaves when one or more of its allocated resources fail depends on the allocated resource failure policy. Depending on the type of job, type of resources, and type of middleware infrastructure, a site may choose to have different responses based on the job, the resource, and the type of failure.

### 4.4.6.A Failure Responses

By default, Moab cancels a job when an allocated resource failure is detected. However, you can specify the following actions:

Option	Policy Action
<b>CANCEL</b>	Cancels the job
<b>FAIL</b>	Terminates the job as a failed job
<b>HOLD</b>	Places a hold on the job. This option is only applicable if you are using checkpointing
<b>IGNORE</b>	Ignores the failed node, allowing the job to proceed
<b>NOTIFY</b>	Notifies the administrator and user of failure but takes no further action
<b>REQUEUE</b>	Requeues job and allows it to run when alternate resources become available

### 4.4.6.B Policy Precedence

For a given job, the applied policy can be set at various levels with policy precedence applied in the job, class/queue, partition, and then system level. The following table

indicates the available methods for setting this policy:

Object	Parameter	Example
<b>Job</b>	<b>RESFAILPOLICY</b> resource manager extension	<pre>&gt; qsub -l resfailpolicy=requeue</pre>
<b>Class/Queue</b>	<b>RESFAILPOLICY</b> attribute of <b>CLASSCFG</b> parameter	<pre>CLASSCFG[batch] RESFAILPOLICY=CANCEL</pre>
<b>Partition</b>	<b>JOBACTIIONONNODE FAILURE</b> attribute of <b>PARCFG</b> parameter	<pre>PARCFG[web3] JOBACTIIONONNODEFAILURE=NOTIFY</pre>
<b>System</b>	<b>NODEALLOCRESFAILURE POLICY</b> parameter	<pre>NODEALLOCRESFAILUREPOLICY=MIGRATE</pre>

#### 4.4.6.C Failure Definition

Any allocated node going down constitutes a failure. However, for certain types of workload, responses to failures may be different depending on whether it is the master task (task 0) or a slave task that fails. To indicate that the associated policy should only take effect if the master task fails, the allocated resource failure policy should be specified with a trailing asterisk (\*), as in the following example:

```
CLASSCFG[virtual_services] RESFAILPOLICY=requeue*
```

#### 4.4.6.D Torque Failure Details

When a node fails to send a status update within a configurable time frame (default 600 seconds, see `node_check_rate` in the *Torque Resource Manager Administrator Guide*), `pbs_server` determines that the node is down. Depending on the `JOBACTIIONONNODEFAILURE` parameter setting, Moab may then notify administrators, hold the job, requeue the job, allocate replacement resources to the job, or cancel the job. If Moab requests that Torque cancel or requeue the job, Torque immediately frees all non-failed resources, making them available for use by other jobs. `pbs_mom` also cleans up parallel jobs after a configurable time frame (default 600 seconds, see `$job_exit_wait_time` in the *Torque Resource Manager Administrator Guide*). Once the failed node is recovered, it contacts the resource manager, determines that the associated job has been canceled/requeued, cleans up, and makes itself available for new workload.

### Related Topics

- [Node State Overview](#)
- [JOBACTIONONNODEFAILURE](#) parameter
- [NODEFAILURERESERVETIME](#) parameter
- [NODEDOWNSTATEDELAYTIME](#) parameter (down nodes will be marked unavailable for the specified duration)
- [NODEDRAINSTATEDELAYTIME](#) parameter (offline nodes will be marked unavailable for the specified duration)
- [NODEBUSYSTATEDELAYTIME](#) parameter (nodes with unexpected background load will be marked unavailable for the specified duration)
- [NODEALLOCRESFILUREPOLICY](#) parameter (action to take if executing jobs have one or more allocated nodes fail)

## Chapter 5: Managing Fairness - Throttling Policies, Fairshare, Allocation Management

In this chapter:

5.1	Fairness Overview .....	377
5.1.1	Fairness Facilities .....	377
5.1.2	Selecting the Correct Policy Approach .....	380
5.2	Usage Limits/Throttling Policies .....	381
5.2.1	Fairness via Throttling Policies .....	381
5.2.2	Override Limits .....	393
5.2.3	Idle Job Limits .....	394
5.2.4	Hard and Soft Limits .....	396
5.2.5	Per-partition Limits .....	397
5.2.6	Usage-based limits .....	398
5.3	Fairshare .....	401
5.3.1	Fairshare Parameters .....	402
5.3.2	Using Fairshare Information .....	406
5.3.3	Hierarchical Fairshare/Share Trees .....	412
5.4	Sample FairShare Data File .....	417
5.5	Accounting, Charging, and Allocation Management .....	418
5.5.1	Accounting Manager Overview .....	418
5.5.2	Accounting Mode .....	419
5.5.3	Accounting Manager Interface Types .....	420
5.5.4	Charging for Jobs .....	423
5.5.5	Charging for Reservations .....	424
5.5.6	Accounting Properties Reported to the Accounting Manager .....	425
5.5.7	Accounting Stages .....	430
5.5.8	Accounting Events .....	433
5.5.9	Blocking Versus Non-Blocking Accounting Actions .....	433
5.5.10	Retrying Failed Charges .....	434
5.6	AMCFG Parameters and Flags .....	437
5.6.1	AMCFG Parameters .....	437
5.6.2	AMCFG Flags .....	460

## 5.1 Fairness Overview

The concept of cluster fairness varies widely from person to person and site to site. While some interpret it as giving all users equal access to compute resources, more complicated concepts incorporating historical resource usage, political issues, and job value are equally valid. While no scheduler can address all possible definitions of fair, Moab provides one of the industry's most comprehensive and flexible set of tools allowing most sites the ability to address their many and varied fairness management needs.

In this topic:

[5.1.1 Fairness Facilities - page 377](#)

[5.1.2 Selecting the Correct Policy Approach - page 380](#)

### 5.1.1 Fairness Facilities

Under Moab, most fairness policies are addressed by a combination of the facilities described in the following table:

## Fairness Facilities

Job Prioritization	
Description	Specifies what is most important to the scheduler. Using service based priority factors allows a site to balance job turnaround time, expansion factor, or other scheduling performance metrics.
Example	<pre>SERVICEWEIGHT 1 QUEUETIMEWEIGHT 10</pre> <p><i>Causes jobs to increase in priority by 10 points for every minute they remain in the queue.</i></p>
Usage Limits (Throttling Policies)	
Description	Specifies limits on exactly what resources can be used at any given instant.
Example	<pre>USERCFG[john] MAXJOB=3 GROUPCFG[DEFAULT] MAXPROC=64 GROUPCFG[staff] MAXPROC=128</pre> <p><i>Allows john to only run 3 jobs at a time. Allows the group staff to use up to 128 total processors and all other groups to use up to 64 processors.</i></p>
Fairshare	
Description	Specifies usage targets to limit resource access or adjust priority based on historical cluster and grid level resource usage.
Example	<pre>USERCFG[steve] FSTARGET=25.0+ FSWEIGHT 1 FSUSERWEIGHT 10</pre> <p><i>Enables priority based fairshare and specifies a fairshare target for user steve such that his jobs are favored in an attempt to keep his jobs using at least 25.0% of delivered compute cycles.</i></p>
Allocation Management	
Description	Specifies long term, credential-based resource usage limits.

### Allocation Management

<p>Example</p>	<pre>AMCFG[mam] TYPE=MAM HOST=server.sys.net</pre> <p><i>Enables the Moab Accounting Manager allocation management interface. Within the accounting manager, project or account based allocations can be configured. These allocations can, for example, do such things as allow project X to use up to 100,000 processor-hours per quarter, provide various QoS sensitive charge rates, and share allocation access.</i></p>
----------------	---

### Quality of Service

<p>Description</p>	<p>Specifies additional resource and service access for particular users, groups, and accounts. QoS facilities can provide special priorities, policy exemptions, reservation access, and other benefits (as well as special charge rates).</p>
<p>Example</p>	<pre>QOSCFG[orion] PRIORITY=1000 XFTARGET=1.2 QOSCFG[orion] QFLAGS=PREEMPTOR, IGNSYSTEM, RESERVEALWAYS</pre> <p><i>Enables jobs requesting the orion QoS a priority increase, an expansion factor target to improve response time, the ability to preempt other jobs, an exemption from system level job size policies, and the ability to always reserve needed resources if it cannot start immediately.</i></p>

### Standing Reservations

<p>Description</p>	<p>Reserves blocks of resources within the cluster for specific, periodic time frames under the constraints of a flexible access control list.</p>
<p>Example</p>	<pre>SRCFG[jupiter] HOSTLIST=node01[1-4] SRCFG[jupiter] STARTTIME=9:00:00 ENDTIME=17:00:00 SRCFG[jupiter] USERLIST=john,steve ACCOUNTLIST=jupiter</pre> <p><i>Reserve nodes node011 through node014 from 9:00 AM until 5:00 PM for use by jobs from user john or steve or from the project jupiter.</i></p>

### Class/Queue Constraints

<p>Description</p>	<p>Associates users, resources, priorities, and limits with cluster classes or</p>
--------------------	--

Class/Queue Constraints	
	cluster queues that can be assigned to or selected by end-users.
Example	<pre>CLASSCFG[long] HOSTLIST=acn[1-4][0-9] CLASSCFG[long] MIN.WCLIMIT=24:00:00 SRCFG[jupiter] PRIORITY=10000 SRCFG[jupiter] CLASSLIST=long&amp;</pre> <p><i>Assigns long jobs a high priority but only allows them to run on certain nodes.</i></p>

## 5.1.2 Selecting the Correct Policy Approach

Moab supports a rich set of policy controls in some cases allowing a particular policy to be enforced in more than one way. For example, cycle distribution can be controlled using usage limits, fairshare, or even queue definitions. Selecting the most correct policy depends on site objectives and needs; consider the following when making such a decision:

- Minimal end-user training
  - Does the solution use an approach familiar to or easily learned by existing users?
- End-user transparency
  - Can the configuration be enabled or disabled without impacting user behavior or job submission?
- Impact on system utilization and system responsiveness
- Solution complexity
  - Is the impact of the configuration readily intuitive, and is it easy to identify possible side effects?
- Solution extensibility and flexibility
  - Will the proposed approach allow the solution to be easily tuned and extended as cluster needs evolve?

---

### Related Topics

- [Job Prioritization](#)
- [Usage Limits \(Throttling Policies\)](#)
- [Fairshare](#)



- [Allocation Management](#)
- [Quality of Service](#)
- [Standing Reservations](#)
- [Class/Queue Constraints](#)

## 5.2 Usage Limits/Throttling Policies

A number of Moab policies allow an administrator to control job flow through the system. These throttling policies work as filters allowing or disallowing a job to be considered for scheduling by specifying limits regarding system usage for any given moment. These policies can be specified as global or specific constraints specified on a per user, group, account, QoS, or class basis.

In this topic:

- [5.2.1 Fairness via Throttling Policies - page 381](#)
- [5.2.2 Override Limits - page 393](#)
- [5.2.3 Idle Job Limits - page 394](#)
- [5.2.4 Hard and Soft Limits - page 396](#)
- [5.2.5 Per-partition Limits - page 397](#)
- [5.2.6 Usage-based limits - page 398](#)

### 5.2.1 Fairness via Throttling Policies

Moab allows significant flexibility with usage limits, or throttling policies. At a high level, Moab allows resource usage limits to be specified in three primary workload categories: (1) active, (2) idle, and (3) system job limits.

#### 5.2.1.A Basic Fairness Policies

Workload category	Description
<b>Active job limits</b>	Constrain the total cumulative resources available to active jobs at a given time.
<b>Idle job limits</b>	Constrain the total cumulative resources available to idle jobs at a given

Workload category	Description
	time.
<b>System job limits</b>	Constrain the maximum resource requirements of any single job.

These limits can be applied to any job credential (user, group, account, QoS, and class), or on a system-wide basis. Using the keyword `DEFAULT`, a site can also specify the default setting for the desired user, group, account, QoS, and class. Additionally, you can configure QoS to allow limit overrides to any particular policy.

To run, a job must meet all policy limits. Limits are applied using the `*CFG` set of parameters, particularly `USERCFG`, `GROUPCFG`, `ACCOUNTCFG`, `QOSCFG`, `CLASSCFG`, and `SYSCFG`. Limits are specified by associating the desired limit to the individual or default object. The usage limits currently supported are listed in the following table.

### Usage Limits

MAXARRAYJOB	
<b>Units</b>	Number of simultaneous active array job subjobs.
<b>Description</b>	Limits the number of simultaneously active (starting or running) array subjobs a credential can have.
<b>Example</b>	<pre>USERCFG[gertrude] MAXARRAYJOB=10</pre> <p><i>Gertrude can have a maximum of 10 active job array subjobs.</i></p>

MAXGRES	
<b>Units</b>	# of concurrent uses of a generic resource
<b>Description</b>	Limits the concurrent usage of a generic resource to a specific quantity or quantity range.
<b>Example</b>	<pre>USERCFG[joe] MAXGRES[matlab]=2 USERCFG[jim] MAXGRES[matlab]=2,4</pre>

MAXJOB	
<b>Units</b>	# of jobs
<b>Description</b>	<p>Limits the number of jobs a credential can have active (starting or running) at any given time. Moab places a hold on all new jobs submitted by that credential once it has reached its maximum number of allowable jobs.</p> <div style="border: 1px solid #0070C0; padding: 5px; margin: 5px 0;"> <p><b>i</b> MAXJOB=0 is not supported. You can, however, achieve similar results by using the <b>HOLD</b> attribute of the <b>USERCFG</b> parameter:</p> <pre>USERCFG[john] HOLD=yes</pre> </div>
<b>Example</b>	<pre>USERCFG[DEFAULT] MAXJOB=8 GROUPCFG[staff] MAXJOB=2,4</pre>

MAXMEM	
<b>Units</b>	total memory in MB
<b>Description</b>	Limits the total amount of dedicated memory (in MB) that can be allocated by a credential's active jobs at any given time.
<b>Example</b>	<code>ACCOUNTCFG[jasper] MAXMEM=2048</code>

MAXNODE	
<b>Units</b>	# of nodes
<b>Description</b>	<p>Limits the total number of compute nodes that can be in use by active jobs at any given time.</p> <div style="border: 1px solid #005596; padding: 5px; margin: 5px 0;"> <p><b>i</b> Adaptive Computing recommends that you set <code>JOBNODEMATCHPOLICY EXACTNODE</code> when using <b>MAXNODE</b>. This ensures jobs submitted using the <code>msub/qsub "-l nodes=#"</code> syntax will have a node count associated with the request.</p> </div> <div style="border: 1px solid #005596; padding: 5px; margin: 5px 0;"> <p><b>i</b> On some systems (including Torque/PBS), nodes have been softly defined rather than strictly defined; that is, a job may request 2 nodes but Torque will translate this request into 1 node with 2 processors. This can prevent Moab from enforcing a <b>MAXNODE</b> policy correctly for a single job. Correct behavior can be achieved using <b>MAXPROC</b>.</p> </div>
<b>Example</b>	<code>CLASSCFG[batch] MAXNODE=64</code>

MAXPE	
<b>Units</b>	# of <a href="#">processor equivalents</a>
<b>Description</b>	Limits the total number of dedicated processor-equivalents that can be allocated by active jobs at any given time.
<b>Example</b>	<code>QOSCFG[base] MAXPE=128</code>

MAXPROC	
<b>Units</b>	# of processors
<b>Description</b>	Limits the total number of dedicated processors that can be allocated by active jobs at any given time per credential. To set <code>MAXPROC</code> per job, use <code>msub -W</code> .
<b>Example</b>	<pre>CLASSCFG[debug] MAXPROC=32</pre>

MAXPS	
<b>Units</b>	<# of processors> * <walltime>
<b>Description</b>	Limits the number of outstanding processor-seconds a credential can have allocated at any given time. For example, if a user has a 4-processor job that will complete in 1 hour and a 2-processor job that will complete in 6 hours, they have $4 * 1 * 3600 + 2 * 6 * 3600 = 16 * 3600$ outstanding processor-seconds. The outstanding processor-second usage of each credential is updated each scheduling iteration, decreasing as jobs approach their completion time.
<b>Example</b>	<pre>USERCFG[DEFAULT] MAXPS=720000</pre>

MAXSUBMITJOBS	
<b>Units</b>	# of jobs
<b>Description</b>	Limits the number of jobs a credential can submit and have in the system at once. Moab will reject any job submitted beyond this limit. If you use a Torque resource manager, you should also set <code>max_user_queueable</code> in case the user submits jobs via <code>qsub</code> instead of <code>msub</code> . See 'Queue Attributes' in the <i>Torque Administrator Guide</i> for more information.
<b>Example</b>	<pre>USERCFG[DEFAULT] MAXSUBMITJOBS=5</pre>

MAXWC	
<b>Units</b>	job duration [[[DD:]HH:]MM:]SS
<b>Description</b>	Limits the cumulative remaining walltime a credential can have associated with active jobs. It behaves identically to the <a href="#">MAXPS</a> limit (listed earlier) only lacking the processor weighting. Like MAXPS, the cumulative remaining walltime of each credential is also updated each scheduling iteration. <div style="border: 1px solid #004a87; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p><b>i</b> MAXWC does not limit the maximum wallclock limit per job. For this capability, use <a href="#">MAX.WCLIMIT</a>.</p> </div>
<b>Example</b>	<pre>USERCFG[ops] MAXWC=72:00:00</pre>

The following example demonstrates a simple limit specification:

```
USERCFG[DEFAULT] MAXJOB=4
USERCFG[john] MAXJOB=8
```

*This example allows user john to run up to 8 jobs while all other users can only run up to 4.*

Simultaneous limits of different types can be applied per credential and multiple types of credentials can have limits specified. The next example demonstrates this mixing of limits and is a bit more complicated.

```
USERCFG[steve] MAXJOB=2 MAXNODE=30
GROUPCFG[staff] MAXJOB=5
CLASSCFG[DEFAULT] MAXNODE=16
CLASSCFG[batch] MAXNODE=32
```

This configuration may potentially apply multiple limits to a single job. As discussed previously, a job can only run if it satisfies all applicable limits. Thus, in this example, the scheduler will be constrained to allow at most 2 simultaneous user `steve` jobs with an aggregate node consumption of no more than 30 nodes. However, if the job is submitted to a class other than `batch`, it may be limited further. Here, only 16 total nodes can be used simultaneously by jobs running in any given class with the exception of the class `batch`. If `steve` submitted a job to run in the class `interactive`, for example, and there were jobs already running in this class using a total of 14 nodes, his job would be blocked unless it requested 2 or fewer nodes by the default limit of 16 nodes per class.

### 5.2.1.B Multi-Dimension Fairness Policies and Per Credential Overrides

Multi-dimensional fairness policies allow a site to specify policies based on combinations of job credentials. A common example might be setting a maximum number of jobs allowed per queue per user or a total number of processors per group per QoS. As with basic

fairness policies, multi-dimension policies are specified using the `*CFG` parameters or through the [identity manager interface](#). Moab supports the most commonly used multi-dimensional fairness policies (listed in the table below) using the following format:

```
*CFG [X] <LIMITTYPE> [<CRED>] =<LIMITVALUE>
```

`*CFG` is one of `USERCFG`, `GROUPCFG`, `ACCOUNTCFG`, `QOSCFG`, or `CLASSCFG`, the `<LIMITTYPE>` policy is one of the policies listed in the table in section [6.2.1.1](#), and `<CRED>` is of the format `<CREDTYPE>[:<VALUE>]` with `CREDTYPE` being one of `USER`, `GROUP`, `ACCT`, `QoS`, or `CLASS`. The optional `<VALUE>` setting can be used to specify that the policy only applies to a specific credential value. For example, the following configuration sets limits on the class `fast`, controlling the maximum number of jobs any group can have active at any given time and the number of processors in use at any given time for user `steve`.

```
CLASSCFG[fast] MAXJOB[GROUP]=12
CLASSCFG[fast] MAXPROC[USER:steve]=50
CLASSCFG[fast] MAXIJOB[USER]=10
```

The following example configuration may clarify further:

```
# allow class batch to run up the 3 simultaneous jobs
# allow any user to use up to 8 total nodes within class
CLASSCFG[batch] MAXJOB=3 MAXNODE[USER]=8
# allow users steve and bob to use up to 3 and 4 total processors respectively within
class
CLASSCFG[fast] MAXPROC[USER:steve]=3 MAXPROC[USER:bob]=4
```

**i** Multi-dimensional policies cannot be applied on `DEFAULT` credentials.

The table below lists the currently implemented, multi-dimensional usage limit permutations. The "slmt" stands for "Soft Limit" and "hlmt" stands for "Hard Limit."

*Multi-dimensional Usage Limit Permutations*

Multi-dimension usage limit permutations	
ACCOUNTCFG[name]	MAXIJOB[QOS]=hlmt MAXIJOB[QOS:qosname]=hlmt
	MAXIPROC[QOS]=hlmt MAXIPROC[QOS:qosname]=hlmt
	MAXJOB[QOS]=slmt,hlmt MAXJOB[QOS:qosname]=slmt,hlmt
	MAXJOB[USER]=slmt,hlmt MAXJOB[USER:username]=slmt,hlmt
	MAXMEM[USER]=slmt,hlmt MAXMEM[USER:username]=slmt,hlmt
	MAXNODE[USER]=slmt,hlmt MAXNODE[USER:username]=slmt,hlmt
	MAXPE[QOS]=slmt,hlmt MAXPE[QOS:qosname]=slmt,hlmt
	MAXPROC[USER]=slmt,hlmt MAXPROC[USER:username]=slmt,hlmt
	MAXPROC[QOS]=slmt,hlmt MAXPROC[QOS:qosname]=slmt,hlmt
	MAXPROC[USER]=slmt,hlmt MAXPROC[USER:username]=slmt,hlmt
	MAXPS[QOS]=slmt,hlmt MAXPS[QOS:qosname]=slmt,hlmt
	MAXPS[USER]=slmt,hlmt MAXPS[USER:username]=slmt,hlmt



Multi-dimension usage limit permutations	
	MAXWC[USER]=slmt,hlmt MAXWC[USER:username]=slmt,hlmt
CLASSCFG[name]	MAXJOB[USER]=hlmt  MAXJOB[GROUP]=slmt,hlmt MAXJOB[GROUP:groupname]=slmt,hlmt
	MAXJOB[QOS:qosname]=hlmt
	MAXJOB[USER]=slmt,hlmt MAXJOB[USER:username]=slmt,hlmt
	MAXMEM[GROUP]=slmt,hlmt MAXMEM[GROUP]=slmt,hlmt
	MAXMEM[GROUP]=slmt,hlmt MAXMEM[GROUP:groupname]=slmt,hlmt
	MAXMEM[QOS:qosname]=hlmt
	MAXMEM[USER]=slmt,hlmt MAXMEM[USER:username]=slmt,hlmt
	MAXNODE[GROUP]=slmt,hlmt MAXNODE[GROUP:groupname]=slmt,hlmt
	MAXNODE[QOS:qosname]=hlmt
	MAXNODE[USER]=slmt,hlmt MAXNODE[USER:username]=slmt,hlmt
	MAXPE[GROUP]=slmt,hlmt MAXPE[GROUP:groupname]=slmt,hlmt
	MAXPE[QOS:qosname]=hlmt

Multi-dimension usage limit permutations	
	MAXPE[USER]=slmt,hlmt MAXPE[USER:username]=slmt,hlmt
	MAXPROC[GROUP]=slmt,hlmt MAXPROC[GROUP:groupname]=slmt,hlmt
	MAXPROC[QOS:qosname]=hlmt
	MAXPROC[USER]=slmt,hlmt MAXPROC[USER:username]=slmt,hlmt
	MAXPS[GROUP]=slmt,hlmt MAXPS[GROUP:groupname]=slmt,hlmt
	MAXPS[QOS:qosname]=hlmt
	MAXPS[USER]=slmt,hlmt MAXPS[USER:username]=slmt,hlmt
	MAXWC[GROUP]=slmt,hlmt MAXWC[GROUP:groupname]=slmt,hlmt
	MAXWC[QOS:qosname]=hlmt
	MAXWC[USER]=slmt,hlmt MAXWC[USER:username]=slmt,hlmt

Multi-dimension usage limit permutations	
GROUPCFG[name]	MAXJOB[CLASS:classname]=slmt,hlmt
	MAXJOB[USER]=slmt,hlmt MAXJOB[USER:username]=slmt,hlmt
	MAXMEM[CLASS:classname]=slmt,hlmt
	MAXMEM[USER]=slmt,hlmt MAXMEM[USER:username]=slmt,hlmt
	MAXNODE[CLASS:classname]=slmt,hlmt
	MAXNODE[USER]=slmt,hlmt MAXNODE[USER:username]=slmt,hlmt
	MAXPE[CLASS:classname]=slmt,hlmt
	MAXPE[USER]=slmt,hlmt MAXPE[USER:username]=slmt,hlmt
	MAXPROC[CLASS:classname]=slmt,hlmt
	MAXPROC[USER]=slmt,hlmt MAXPROC[USER:username]=slmt,hlmt
	MAXPS[CLASS:classname]=slmt,hlmt
	MAXPS[USER]=slmt,hlmt MAXPS[USER:username]=slmt,hlmt
	MAXWC[CLASS:classname]=slmt,hlmt
	MAXWC[USER]=slmt,hlmt MAXWC[USER:username]=slmt,hlmt
QOSCFG[name]	MAXIJOB[ACCT]=hlmt

Multi-dimension usage limit permutations	
	MAXIJOB[ACCT:accountname]=hlmt
	MAXIJOB[USER]=hlmt MAXIJOB[USER:class+classname]=hlmt
	MAXINODE[ACCT]=slmt,hlmt MAXINODE[ACCT:accountname]=slmt,hlmt
	MAXINODE[USER]=hlmt MAXINODE[USER:username]=slmt,hlmt
	MAXIPROC[ACCT]=hlmt MAXIPROC[ACCT:accountname]=hlmt
	MAXJOB[ACCT]=slmt,hlmt MAXJOB[ACCT:accountname]=slmt,hlmt
	MAXJOB[USER]=slmt,hlmt MAXJOB[USER:username]=slmt,hlmt
	MAXMEM[USER]=slmt,hlmt MAXMEM[USER:username]=slmt,hlmt
	MAXNODE[USER]=slmt,hlmt MAXNODE[USER:username]=slmt,hlmt
	MAXPE[ACCT]=slmt,hlmt MAXPE[ACCT:accountname]=slmt,hlmt
	MAXPE[USER]=slmt,hlmt MAXPE[USER:username]=slmt,hlmt
	MAXPROC[ACCT]=slmt,hlmt MAXPROC[ACCT:accountname]=slmt,hlmt
	MAXPROC[USER]=slmt,hlmt

Multi-dimension usage limit permutations	
	MAXPROC[USER:username]=slmt,hlmt
	MAXPS[ACCT]=slmt,hlmt MAXPS[ACCT:accountname]=slmt,hlmt
	MAXPS[USER]=slmt,hlmt MAXPS[USER:username]=slmt,hlmt
	MAXWC[USER]=slmt,hlmt MAXWC[USER:username]=slmt,hlmt
USERCFG[name]	MAXJOB[GROUP]=slmt,hlmt MAXJOB[GROUP:groupname]=slmt,hlmt
	MAXMEM[GROUP]=slmt,hlmt MAXMEM[GROUP:groupname]=slmt,hlmt
	MAXNODE[GROUP]=slmt,hlmt MAXNODE[GROUP:groupname]=slmt,hlmt
	MAXPE[GROUP]=slmt,hlmt MAXPE[GROUP:groupname]=slmt,hlmt
	MAXPROC[GROUP]=slmt,hlmt MAXPROC[GROUP:groupname]=slmt,hlmt
	MAXPS[GROUP]=slmt,hlmt MAXPS[GROUP:groupname]=slmt,hlmt
	MAXWC[GROUP]=slmt,hlmt MAXWC[GROUP:groupname]=slmt,hlmt

## 5.2.2 Override Limits

Like all job credentials, the QoS object can be associated with resource usage limits. However, this credential can also be given special override limits that supersede the limits of other credentials, effectively causing all other limits of the same type to be ignored. See

[QoS Usage Limits and Overrides](#) for a complete list of policies that can be overridden. The following configuration provides an example of this in the last line:

```
USERCFG[steve]      MAXJOB=2    MAXNODE=30
GROUPCFG[staff]    MAXJOB=5
CLASSCFG[DEFAULT] MAXNODE=16
CLASSCFG[batch]    MAXNODE=32
QOSCFG[hiprio]     OMAXJOB=3   OMAXNODE=64
```

*Only 3 hiprio QoS jobs can run simultaneously and hiprio QoS jobs can run with up to 64 nodes per credential ignoring other credential MAXNODE limits.*

Given the preceding configuration, assume a job is submitted with the credentials, user `steve`, group `staff`, class `batch`, and QoS `hiprio`.

Such a job will start so long as running it does not lead to any of the following conditions:

- Total nodes used by user `steve` does not exceed 64.
- Total active jobs associated with user `steve` does not exceed 2.
- Total active jobs associated with group `staff` does not exceed 5.
- Total nodes dedicated to class `batch` does not exceed 64.
- Total active jobs associated with QoS `hiprio` does not exceed 3.

While the preceding example is a bit complicated for most sites, similar combinations may be required to enforce policies found on many systems.

### 5.2.3 Idle Job Limits

Idle (or queued) job limits control which jobs are eligible for scheduling. To be eligible for scheduling, a job must meet the following conditions:

- Be idle as far as the resource manager is concerned (no holds).
- Have all job prerequisites satisfied (no outstanding job or data dependencies).
- Meet all idle job throttling policies.

If a job fails to meet any of these conditions, it will not be considered for scheduling and will not accrue service based job prioritization (see [Service \(SERVICE\) Component](#) and [JOBPRIOACCRUALPOLICY](#).) The primary purpose of idle job limits is to ensure fairness among competing users by preventing queue stuffing and other similar abuses. Queue stuffing occurs when a single entity submits large numbers of jobs, perhaps thousands, all at once so they begin accruing queue time based priority and remain first to run despite subsequent submissions by other users.

Idle limits are specified in a manner almost identical to active job limits with the insertion of the capital letter *I* into the middle of the limit name. The following tables describe the

MAXIARRAYJOB, MAXIJOB, and MAXINODE limits, which are idle limit equivalents to MAXARRAYJOB, MAXJOB, and MAXNODE limits, respectively.

MAXIARRAYJOB	
<b>Units</b>	Number of simultaneous idle array job subjobs.
<b>Description</b>	Limits the number of simultaneously idle (eligible) job array subjobs across <i>all</i> job arrays submitted by a credential.
<b>Example</b>	<pre>USERCFG[gertrude] MAXARRAYJOB=10 MAXIARRAYJOB=5</pre> <p><i>Gertrude can have a maximum of 10 active job array subjobs and 5 eligible job array subjobs.</i></p>

MAXIJOB	
<b>Units</b>	# of jobs
<b>Description</b>	Limits the number of idle (eligible) jobs a credential can have at any given time.
<b>Example</b>	<pre>USERCFG[DEFAULT] MAXIJOB=8 GROUPCFG[staff] MAXIJOB=4</pre>

MAXINODE	
<b>Units</b>	# of nodes
<b>Description</b>	Limits the total number of compute nodes that can be requested by jobs in the eligible/idle queue at any time. Once the limit is exceeded, the remaining jobs will be placed in the blocked queue. The number of nodes is determined by <code>&lt;tasks&gt; / &lt;maximumProcsOnOneNode&gt;</code> or, if using <a href="#">JOBNODEMATCHPOLICY EXACTNODE</a> , by the number of nodes requested.
<b>Example</b>	<pre>USERCFG[DEFAULT] MAXINODE=2</pre>

Idle limits can constrain the total number of jobs considered to be eligible on a per credential basis. Further, like active job limits, idle job limits can also constrain eligible jobs based on aggregate requested resources. This could, for example, allow a site to indicate that for a given user, only jobs requesting up to a total of 64 processors, or 3200 processor-seconds would be considered at any given time. Which jobs to select is

accomplished by prioritizing all idle jobs and then adding jobs to the eligible list one at a time in priority order until jobs can no longer be added. This eligible job selection is done only once per scheduling iteration, so, consequently, idle job limits only support a single hard limit specification. Any specified soft limit is ignored.

All single dimensional job limit types supported as active job limits are also supported as idle job limits. In addition, Moab also supports `MAXIJOB[USER]` and `MAXIPROC[USER]` policies on a per class basis (see [Basic Fairness Policies](#).)

### Example:

```
USERCFG[steve]      MAXIJOB=2
GROUPCFG[staff]    MAXIJOB=5
CLASSCFG[batch]    MAXIJOB[USER]=2 MAXIJOB[USER:john]=6
QOSCFG[hiprio]     MAXIJOB=3
```

## 5.2.4 Hard and Soft Limits

Hard and soft limit specification allows a site to balance both fairness and utilization on a given system. Typically, throttling limits are used to constrain the quantity of resources a given credential (such as user or group) is allowed to consume. These limits can be very effective in enforcing fair usage among a group of users. However, in a lightly loaded system, or one in which there are significant swings in usage from project to project, these limits can reduce system utilization by blocking jobs even when no competing jobs are queued.

Soft limits help address this problem by providing additional scheduling flexibility. They allow sites to specify two tiers of limits; the more constraining limits soft limits are in effect in heavily loaded situations and reflect tight fairness constraints. The more flexible hard limits specify how flexible the scheduler can be in selecting jobs when there are idle resources available after all jobs meeting the tighter soft limits have started. Soft and hard limits are specified in the format `[<SOFTLIMIT>, ]<HARDLIMIT>`. For example, a given site may want to use the following configuration:

```
USERCFG[DEFAULT]  MAXJOB=2, 8
```

*With this configuration, the scheduler would select all jobs that meet the per user MAXJOB limit of 2. It would then attempt to start and reserve resources for all of these selected jobs. If after doing so there still remain available resources, the scheduler would then select all jobs that meet the less constraining hard per user MAXJOB limit of 8 jobs. These jobs would then be scheduled and reserved as available resources allow.*

*If no soft limit is specified or the soft limit is less constraining than the hard limit, the soft limit is set equal to the hard limit.*

### Example:

```
USERCFG[steve]    MAXJOB=2, 4 MAXNODE=15, 30
GROUPCFG[staff]  MAXJOB=2, 5
CLASSCFG[DEFAULT] MAXNODE=16, 32
CLASSCFG[batch]  MAXNODE=12, 32
QOSCFG[hiprio]   MAXJOB=3, 5 MAXNODE=32, 64
```



**i** Job **preemption** status can be adjusted based on whether the job violates a soft policy using the **ENABLESPVIOLATIONPREEMPTION** parameter.

## 5.2.5 Per-partition Limits

Per-partition scheduling can set limits and enforce credentials and policies on a per-partition basis. Configuration for per-partition scheduling is done on the Moab Grid Control. In a grid, each Moab cluster is considered a partition. Per-partition scheduling is typically used in a Moab Grid Control / Moab Grid Member grid.

To enable per-partition scheduling, add the following to `moab.cfg`:

```
PERPARTITIONSCHEDULING TRUE
JOBMIGRATEPOLICY JUSTINTIME
```

**i** With per-partition scheduling, it is recommended that limits go on the specific partitions and not on the global level. If limits are specified on both levels, Moab will take the more constricting of the limits. Also, note that a **DEFAULT** policy on the global partition is not overridden by any policy on a specific partition.

### 5.2.5.A Configuring Per-partition Limits

You can configure per-job limits and credential usage limits on a per-partition basis in the `moab.cfg` file. Here is a sample configuration for partitions `g02` and `g03` in `moab.cfg`.

```
PARCFG[g02]    CONFIGFILE=/opt/moab/parg02.cfg
PARCFG[g03]    CONFIGFILE=/opt/moab/parg03.cfg
```

You can then add per-partition limits in each partition configuration file:

```
# /opt/moab/parg02.cfg
CLASSCFG[pbatch]    MAXJOB=5
```

```
# /opt/moab/parg03.cfg
CLASSCFG[pbatch]    MAXJOB=10
```

You can configure Moab so that jobs submitted to any partition besides `g02` and `g03` get the default limits in `moab.cfg`:

```
stl
CLASSCFG[pbatch]    MAXJOB=2
```

### 5.2.5.B Supported Credentials and Limits

The user, group, account, QoS, and class credentials are supported in per-partition scheduling.

The following per-job limits are supported:

- `MAX.NODE`
- `MAX.WCLIMIT`
- `MAX.PROC`

The following credential usage limits are supported:

- `MAXJOB`
- `MAXNODE`
- `MAXPROC`
- `MAXWC`
- `MAXSUBMITJOBS`

**Multi-dimensional limits** are supported for the listed credentials and per-job limits. For example:

```
CLASSCFG[pbatch]    MAXJOB[user:frank]=10
```

## 5.2.6 Usage-based limits

Resource usage limits constrain the amount of resources a given job can consume. These limits are generally proportional to the resources requested and can include walltime, any standard resource, or any specified generic resource. The parameter `RESOURCELIMITPOLICY` controls which resources are limited, what limit policy is enforced per resource, and what actions the scheduler should take in the event of a policy violation.


### 5.2.6.A Configuring Actions

The `RESOURCELIMITPOLICY` parameter accepts a number of policies, resources, and actions using the format and values defined below.

**i** If walltime is the resource to be limited, be sure that the resource manager is configured to not interfere if a job surpasses its given walltime. For Torque, this is done by using `$ignwalltime` in the configuration on each MOM node.

### 5.2.6.B Format

```
RESOURCELIMITPOLICY<RESOURCE>: [<SPOLICY>, ]<HPOLICY>:
[<SACTION>, ]<HACTION> [ : [<SVIOLATIONTIME>, ]<HVIOLATIONTIME>] . . .
```

Resource	Description
<b>CPUTIME</b>	Maximum total job proc-seconds used by any single job (allows scheduler enforcement of cpulimit).
<b>DISK</b>	Local disk space (in MB) used by any single job task.
<b>JOBMEM</b>	Maximum real memory/RAM (in MB) used by any single job. <div style="border: 1px solid black; padding: 5px; margin-top: 5px;">  JOBMEM will only work with the <a href="#">MAXMEM</a> flag. </div>
<b>JOBPROC</b>	Maximum processor load associated with any single job. You must set <a href="#">MAXPROC</a> to use <code>JOBPROC</code> .
<b>MEM</b>	Maximum real memory/RAM (in MB) used by any single job task.
<b>MINJOBPROC</b>	Minimum processor load associated with any single job (action taken if job is using 5% or less of potential CPU usage).
<b>NETWORK</b>	Maximum network load associated with any single job task.
<b>PROC</b>	Maximum processor load associated with any single job task.
<b>SWAP</b>	Maximum virtual memory/SWAP (in MB) used by any single job task.
<b>WALLTIME</b>	Requested job walltime.

Policy	Description
<b>ALWAYS</b>	take action whenever a violation is detected
<b>EXTENDEDVIOLATION</b>	take action only if a violation is detected and persists for greater than the specified time limit
<b>BLOCKEDWORKLOADONLY</b>	take action only if a violation is detected and the constrained resource is required by another job

Action	Description
<b>CANCEL</b>	terminate the job
<b>CHECKPOINT</b>	<a href="#">checkpoint</a> and terminate job
<b>MIGRATE</b>	requeue the job and require a different set of hosts for execution
<b>NOTIFY</b>	notify admins and job owner regarding violation
<b>REQUEUE</b>	terminate and requeue the job
<b>SUSPEND</b>	suspend the job and leave it suspended for an amount of time defined by the <a href="#">MINADMINSTIME</a> parameter

*Example 5-1: Notify and then cancel job if requested memory is exceeded*

```
# if job exceeds memory usage, immediately notify owner
# if job exceeds memory usage for more than 5 minutes, cancel the job
RESOURCELIMITPOLICY MEM:ALWAYS,EXTENDEDVIOLATION:NOTIFY,CANCEL:00:05:00
```

*Example 5-2: Checkpoint job on walltime violations*

```
# if job exceeds requested walltime, checkpoint job
RESOURCELIMITPOLICY WALLTIME:ALWAYS:CHECKPOINT
# when checkpointing, send term signal, followed by kill 1 minute later
RMCFG[base] TYPE=PBS CHECKPOINTTIMEOUT=00:01:00 CHECKPOINTSIG=SIGTERM
```

*Example 5-3: Cancel jobs that use 5% or less of potential CPU usage for more than 5 minutes*

```
RESOURCELIMITPOLICY MINJOBPROC:EXTENDEDVIOLATION:CANCEL:5:00
```

*Example 5-4: Migrating a job when it blocks other workload*

```
RESOURCELIMITPOLICY JOBPROC:BLOCKEDWORKLOADONLY:MIGRATE
```

### 5.2.6.C Specifying Hard and Soft Policy Violations

Moab is able to perform different actions for both hard and soft policy violations. In most resource management systems, a mechanism does not exist to allow the user to specify both hard and soft limits. To address this, Moab provides the [RESOURCELIMITMULTIPLIER](#) parameter that allows per partition and per resource multiplier factors to be specified to generate the actual hard and soft limits to be used. If the factor is less than one, the soft limit will be lower than the specified value and a Moab action will be taken before the specified limit is reached. If the factor is greater than one, the hard limit will be set higher than the specified limit allowing a buffer space before the hard limit action is taken.

In the following example, job owners will be notified by email when their memory reaches 100% of the target, and the job will be canceled if it reaches 125% of the target. For wallclock usage, the job will be requeued when it reaches 90% of the specified limit if another job is waiting for its resources, and it will be checkpointed when it reaches the full limit.

```
RESOURCELIMITPOLICY      MEM:ALWAYS,ALWAYS:NOTIFY,CANCEL
RESOURCELIMITPOLICY      WALLTIME:BLOCKEDWORKLOADONLY,ALWAYS:REQUEUE,CHECKPOINT
RESOURCELIMITMULTIPLIER  MEM:1.25,WALLTIME:0.9
```

### 5.2.6.D Constraining Walltime Usage

While Moab constrains walltime using the parameter `RESOURCELIMITPOLICY` like other resources, it also allows walltime exception policies that are not available with other resources. In particular, Moab allows jobs to exceed the requested wallclock limit by an amount specified on a global basis using the `JOBMAXOVERRUN` parameter or on a per credential basis using the `WCOVERRUN` attribute of the `CLASSCFG` parameter.

```
JOBMAXOVERRUN      00:10:00
CLASSCFG[debug]    wcoverrun=00:00:30
```

#### Related Topics

- [RESOURCELIMITPOLICY](#) parameter
- [FSTREE](#) parameter (set usage limits within share tree hierarchy)
- [Credential Overview](#)
- [JOBMAXOVERRUN](#) parameter
- [WCVIOLATIONACTION](#) parameter
- [RESOURCELIMITMULTIPLIER](#) parameter

## 5.3 Fairshare

Fairshare allows historical resource utilization information to be incorporated into job feasibility and priority decisions. This feature allows site administrators to set system utilization targets for users, groups, accounts, classes, and QoS levels. Administrators can also specify the time frame over which resource utilization is evaluated in determining whether the goal is being reached. Parameters allow sites to specify the utilization metric, how historical information is aggregated, and the effect of fairshare state on scheduling behavior. You can specify fairshare targets for any credentials (such as user, group, and class) that administrators want such information to affect.

In this topic:

[5.3.1 Fairshare Parameters - page 402](#)

[5.3.2 Using Fairshare Information - page 406](#)

[5.3.3 Hierarchical Fairshare/Share Trees - page 412](#)

## 5.3.1 Fairshare Parameters

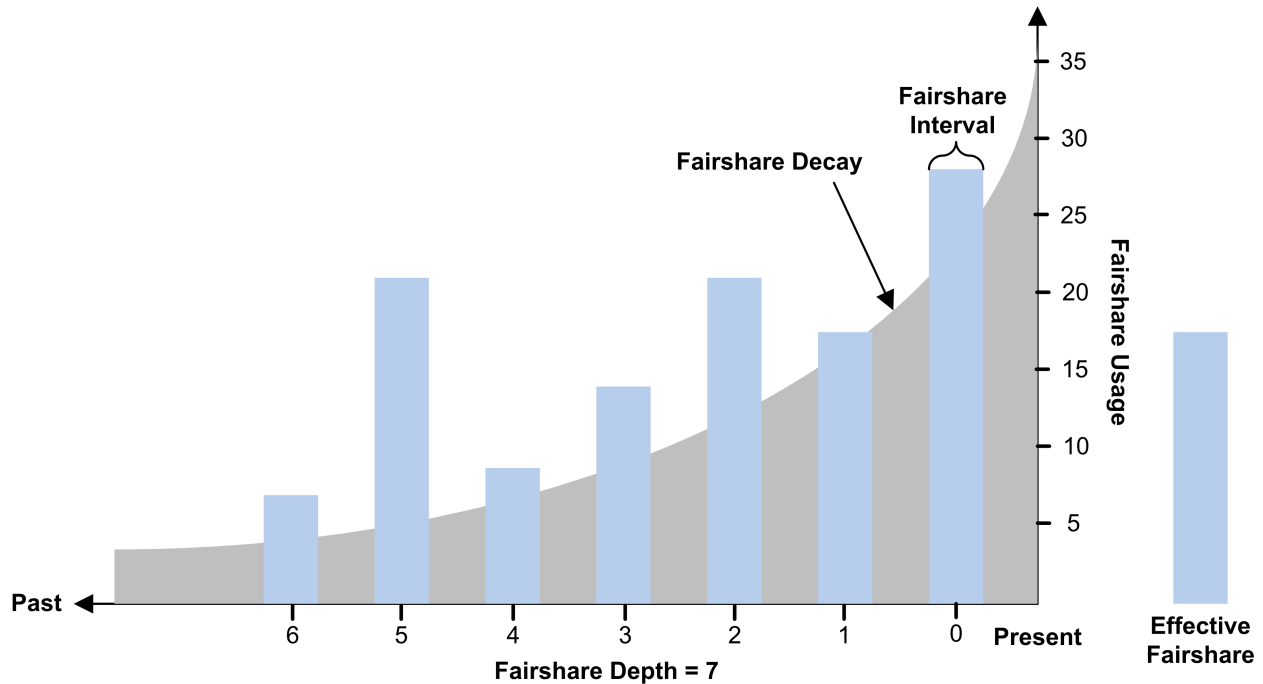
Fairshare is configured at two levels. First, at a system level, configuration is required to determine how fairshare usage information is to be collected and processed. Second, some configuration is required at the credential level to determine how this fairshare information affects particular jobs. The following are system level parameters:

Parameter	Description
<b>FSINTERVAL</b>	Duration of each fairshare window.
<b>FSDEPTH</b>	Number of fairshare windows factored into current fairshare utilization.
<b>FSDECAY</b>	Decay factor applied to weighting the contribution of each fairshare window.
<b>FSPOLICY</b>	Metric to use when tracking fairshare usage.

Credential level configuration consists of specifying fairshare utilization targets using the \*CFG suite of parameters, including [ACCOUNTCFG](#), [CLASSCFG](#), [GROUPECFG](#), [QOSCFG](#), and [USERCFG](#).

If global (multi-cluster) fairshare is used, Moab must be configured to synchronize this information with an [identity manager](#).

Image 5-1: Effective fairshare over 7 days



### 5.3.1.A FSPOLICY - Specifying the Metric of Consumption

As Moab runs, it records how available resources are used. Each iteration (RMPOLLINTERVAL seconds) it updates fairshare resource utilization statistics. Resource utilization is tracked in accordance with the **FSPOLICY** parameter allowing various aspects of resource consumption information to be measured. This parameter allows selection of both the types of resources to be tracked and the method of tracking. It provides the option of tracking usage by dedicated or consumed resources, where dedicated usage tracks what the scheduler assigns to the job and consumed usage tracks what the job actually uses.

Metric	Description
<b>DEDICATEDPES</b>	Usage tracked by <b>processor-equivalent</b> seconds dedicated to each job. This is based on the total number of dedicated processor-equivalent seconds delivered in the system. Useful in dedicated and shared nodes environments.
<b>DEDICATEDPS</b>	Usage tracked by processor seconds dedicated to each job. This is based on the total number of dedicated processor seconds delivered in the system. Useful in dedicated node environments.
<b>DEDICATEDPS%</b>	Usage tracked by processor seconds dedicated to each job. This is based on the total number of dedicated processor seconds <i>available</i> in the

Metric	Description
	system.
[NONE]	Disables fairshare.
UTILIZEDDPS	Usage tracked by processor seconds used by each job. This is based on the total number of utilized processor seconds delivered in the system. Useful in shared node/SMP environments.

*Example 5-5:*

An example may clarify the use of the `FSPOLICY` parameter. Assume a 4-processor job is running a parallel `/bin/sleep` for 15 minutes. It will have a dedicated fairshare usage of 1 processor-hour but a consumed fairshare usage of essentially nothing since it did not consume anything. Most often, dedicated fairshare usage is used on dedicated resource platforms while consumed tracking is used in shared SMP environments.

```

FSPOLICY      DEDICATEDDPS%
FSINTERVAL    24:00:00
FSDEPTH       28
FSDECAY       0.75

```

### 5.3.1.B Specifying Fairshare Timeframe

When configuring fairshare, it is important to determine the proper timeframe that should be considered. Many sites choose to incorporate historical usage information from the last one to two weeks while others are only concerned about the events of the last few hours. The correct setting is very site dependent and usually incorporates both average job turnaround time and site mission policies.

With Moab's fairshare system, time is broken into a number of distinct fairshare windows. Sites configure the amount of time they want to consider by specifying two parameters, `FSINTERVAL` and `FSDEPTH`. The `FSINTERVAL` parameter specifies the duration of each window while the `FSDEPTH` parameter indicates the number of windows to consider. Therefore, the total time evaluated by fairshare is simply `FSINTERVAL * FSDEPTH`.

Many sites want to limit the impact of fairshare data according to its age. The `FSDECAY` parameter allows this, causing the most recent fairshare data to contribute more to a credential's total fairshare usage than older data. This parameter is specified as a standard decay factor, which is applied to the fairshare data. Generally, decay factors are specified as a value between 1 and 0 where a value of 1 (the default) indicates no decay should be specified. The smaller the number, the more rapid the decay using the calculation  $\text{WeightedValue} = \text{Value} * \langle \text{DECAY} \rangle ^ \langle \text{N} \rangle$  where  $\langle \text{N} \rangle$  is the window number.



The following table shows the impact of a number of commonly used decay factors on the percentage contribution of each fairshare window:

Decay Factor	Win0	Win1	Win2	Win3	Win4	Win5	Win6	Win7
<b>1.00</b>	100%	100%	100%	100%	100%	100%	100%	100%
<b>0.80</b>	100%	80%	64%	51%	41%	33%	26%	21%
<b>0.75</b>	100%	75%	56%	42%	31%	23%	17%	12%
<b>0.50</b>	100%	50%	25%	13%	6%	3%	2%	1%

While selecting how the total fairshare time frame is broken up between the number and length of windows is a matter of preference, it is important to note that more windows will cause the decay factor to degrade the contribution of aged data more quickly.

### 5.3.1.C Managing Fairshare Data

Using the selected fairshare usage metric, Moab continues to update the current fairshare window until it reaches a fairshare window boundary, at which point it rolls the fairshare window and begins updating the new window. The information for each window is stored in its own file located in the Moab statistics directory. Each file is named `FS.<EPOCHTIME>[.<PNAME>]` where `<EPOCHTIME>` is the time the new fairshare window became active (see [sample data file](#)) and `<PNAME>` is only used if per-partition [share trees](#) are configured. Each window contains utilization information for each entity, as well as for total usage.

**i** Historical fairshare data is recorded in the fairshare file using the metric specified by the `FSPOLICY` parameter. By default, this metric is processor-seconds.

**i** Historical fairshare data can be directly analyzed and reported using the `mdiag -f -v` command.

When Moab needs to determine current fairshare usage for a particular credential, it calculates a decay-weighted average of the usage information for that credential using the most recent fairshare intervals where the number of windows evaluated is controlled by the `FSDEPTH` parameter. For example, assume the credential of interest is user `john` and the following parameters are set:

```

FSINTERVAL 12:00:00
FSDEPTH    4
FSDECAY    0.5

```

Further assume that the fairshare usage intervals have the following usage amounts:

Fairshare interval	Total user <code>john</code> usage	Total cluster usage
0	60	110
1	0	125
2	10	100
3	50	150

Based on this information, the current fairshare usage for user `john` would be calculated as follows:

$$Usage = (60 * 1 + .5^1 * 0 + .5^2 * 10 + .5^3 * 50) / (110 + .5^1 * 125 + .5^2 * 100 + .5^3 * 150)$$

**i** The current fairshare usage is relative to the actual resources delivered by the system over the timeframe evaluated, not the resources available or configured during that time.

**i** Historical fairshare data is organized into a number of data files, each file containing the information for a length of time as specified by the `FSINTERVAL` parameter. Although `FSDEPTH`, `FSINTERVAL`, and `FSDECAY` can be freely and dynamically modified, such changes might result in unexpected fairshare status for a period of time as the fairshare data files with the old `FSINTERVAL` setting are rolled out.

## 5.3.2 Using Fairshare Information

### 5.3.2.A Fairshare Targets

Once the global fairshare policies have been configured, the next step involves applying resulting fairshare usage information to affect scheduling behavior. As mentioned in the Fairshare Overview, by specifying fairshare targets, site administrators can configure how fairshare information impacts scheduling behavior. The targets can be applied to user, group, account, QoS, or class credentials using the `FSTARGET` attribute of `*CFG` credential

parameters. These targets allow fairshare information to affect job priority and each target can be independently selected to be one of the types documented in the following table:

Target type - Ceiling	
<b>Target modifier</b>	-
<b>Job impact</b>	Priority
<b>Format</b>	Percentage Usage
<b>Description</b>	Adjusts job priority down when usage exceeds target. See <a href="#">How violated ceilings and floors affect fairshare-based priority</a> for more information on how ceilings affect job priority.

Target type - Floor	
<b>Target modifier</b>	+
<b>Job impact</b>	Priority
<b>Format</b>	Percentage Usage
<b>Description</b>	Adjusts job priority up when usage falls below target. See <a href="#">How violated ceilings and floors affect fairshare-based priority</a> for more information on how floors affect job priority.

Target type - Target	
<b>Target modifier</b>	N/A
<b>Job impact</b>	Priority
<b>Format</b>	Percentage Usage
<b>Description</b>	Adjusts job priority when usage does not meet target.

**i** Setting a fairshare target value of 0 indicates that there is no target and that the priority of jobs associated with that credential should not be affected by the credential's previous fairshare target. If you want a credential's cluster usage near 0%, set the target to a very small value, such as 0.001.

**Example**

The following example increases the priority of jobs belonging to user `john` until he reaches 16.5% of total cluster usage. All other users have priority adjusted both up and down to bring them to their target usage of 10%.

```
FSPOLICY          DEDICATEDPS
FSWEIGHT          1
FSUSERWEIGHT      100
USERCFG[john]     FSTARGET=16.5+
USERCFG[DEFAULT] FSTARGET=10
...
```

**5.3.2.B Fairshare Caps**

Where fairshare targets affect a job's priority and position in the eligible queue, fairshare caps affect a job's eligibility. Caps can be applied to users, accounts, groups, classes, and QoSs using the `FSCAP` attribute of `*CFG` credential parameters and can be configured to modify scheduling behavior. Unlike fairshare targets, if a credential reaches its fairshare cap, its jobs can no longer run and are therefore removed from the eligible queue and placed in the blocked queue. In this respect, fairshare targets behave like soft limits and fairshare caps behave like hard limits. Fairshare caps can be absolute or relative as described in the following table. If no modifier is specified, the cap is interpreted as relative.

Absolute Cap	
<b>Cap Modifier</b>	^
<b>Job Impact</b>	Feasibility
<b>Format</b>	Absolute Usage
<b>Description</b>	Constrains job eligibility as an absolute quantity measured according to the scheduler charge metric as defined by the <a href="#">FSPOLICY</a> parameter.

Relative Cap	
<b>Cap Modifier</b>	%
<b>Job Impact</b>	Feasibility
<b>Format</b>	Percentage Usage
<b>Description</b>	Constrains job eligibility as a percentage of total delivered cycles measured according to the scheduler charge metric as defined by the <a href="#">FSPOLICY</a> parameter.

### Example

The following example constrains the `marketing` account to use no more than 16,500 processor seconds during any given floating one week window. At the same time, all other accounts are constrained to use no more than 10% of the total delivered processor seconds during any given one week window.

```
FSPOLICY          DEDICATEDDPS
FSINTERVAL        12:00:00
FSDEPTH           14
ACCOUNTCFG[marketing] FSCAP=16500^
ACCOUNTCFG[DEFAULT]  FSCAP=10
...
```

### 5.3.2.C Priority-Based Fairshare

The most commonly used type of fairshare is priority based fairshare. In this mode, fairshare information does not affect whether a job can run, but rather only the job's priority relative to other jobs. In most cases, this is the desired behavior. Using the standard fairshare target, the priority of jobs of a particular user who has used too many resources over the specified fairshare window is lowered. Also, the standard fairshare target increases the priority of jobs that have not received enough resources.

While the standard fairshare target is the most commonly used, Moab can also specify fairshare ceilings and floors. These targets are like the default target; however, ceilings only adjust priority down when usage is too high and floors only adjust priority up when usage is too low.

Since fairshare usage information must be integrated with Moab's overall priority mechanism, it is critical that the corresponding fairshare priority weights be set. Specifically, the [FSWEIGHT](#) component weight parameter and the target type subcomponent weight (such as [FSACCOUNTWEIGHT](#), [FSCCLASSWEIGHT](#), [FSGROUPWEIGHT](#), [FSQOSWEIGHT](#), and [FSUSERWEIGHT](#)) be specified.

**i** If these weights are not set, the fairshare mechanism will be enabled but have no effect on scheduling behavior. See the [Job Priority Factor Overview](#) for more information on setting priority weights.

### Example

```
# set relative component weighting
FSWEIGHT      1
FSUSERWEIGHT  10
FSGROUPWEIGHT 50

FSINTERVAL    12:00:00
FSDEPTH       4
FSDECAY       0.5
FSPOLICY      DEDICATEDPS
# all users should have a FS target of 10%
USERCFG[DEFAULT] FSTARGET=10.0
# user john gets extra cycles
USERCFG[john]    FSTARGET=20.0
# reduce staff priority if group usage exceed 15%
GROUPCFG[staff] FSTARGET=15.0-
# give group orion additional priority if usage drops below 25.7%
GROUPCFG[orion] FSTARGET=25.7+
```

**i** Job preemption status can be adjusted based on whether the job violates a fairshare target using the [ENABLEFSVIOLATIONPREEMPTION](#) parameter.

### 5.3.2.D Credential-Specific Fairshare Weights

Credential-specific fairshare weights can be set using the **FSWEIGHT** attribute of the **ACCOUNT**, **GROUP**, and **QOS** credentials as in the following example:

```
FSWEIGHT 1000
ACCOUNTCFG[orion1] FSWEIGHT=100
ACCOUNTCFG[orion2] FSWEIGHT=200
ACCOUNTCFG[orion3] FSWEIGHT=-100
GROUPCFG[staff] FSWEIGHT=10
```

If specified, a per-credential fairshare weight is added to the global component fairshare weight.

**i** The **FSWEIGHT** attribute is only enabled for **ACCOUNT**, **GROUP**, and **QOS** credentials.

### 5.3.2.E Fairshare Usage Scaling

Moab uses the **FSSCALINGFACTOR** attribute for **QOS** credentials to get the calculated fairshare usage of a job.

```
QOSCFG[qos1] FSSCALINGFACTOR=<double>
```

Moab will multiply the actual fairshare usage by this value to get the calculated fairshare usage of a job. The actual fairshare usage is calculated based on the [FSPOLICY](#) parameter.

For an example, if FSPOLICY is set to DEDICATEDPS and a job runs on two processors for 100 seconds then the actual fairshare usage would be 200. If the job ran on a qos with FSSCALINGFACTOR=.5 then Moab would multiply  $200 * .5 = 100$ . If the job ran on a partition with FSSCALINGFACTOR=2 then Moab would multiply  $200 * 2 = 400$ .

**i** PARCFG also lets you specify the FSSCALINGFACTOR for partitions. See [6.2.5 Per-Partition Settings - page 522](#).

### 5.3.2.F Extended Fairshare Examples

#### Example 5-6: Multi-Cred Cycle Distribution

Example 1 represents a university setting where different schools have access to a cluster. The Engineering department has put the most money into the cluster and therefore has greater access to the cluster. The Math, Computer Science, and Physics departments have also pooled their money into the cluster and have reduced relative access. A support group also has access to the cluster, but since they only require minimal compute time and shouldn't block the higher-paying departments, they are constrained to five percent of the cluster. At this time, users Tom and John have specific high-priority projects that need increased cycles.

```
#global general usage limits - negative priority jobs are considered in scheduling
ENABLENEGJOBPRIORITY TRUE
# site policy - no job can last longer than 8 hours
USERCFG[DEFAULT] MAX.WCLIMIT=8:00:00
# Note: default user FS target only specified to apply default user-to-user balance
USERCFG[DEFAULT] FSTARGET=1
# high-level fairshare config
FSPOLICY DEDICATEDPS
FSINTERVAL 12:00:00
FSDEPTH 32 #recycle FS every 16 days
FSDECAY 0.8 #favor more recent usage info
# qos config
QOSCFG[inst] FSTARGET=25
QOSCFG[supp] FSTARGET=5
QOSCFG[premium] FSTARGET=70
# account config (QoS access and fstargets)
# Note: user-to-account mapping handled via accounting manager
# Note: FS targets are percentage of total cluster, not percentage of QOS
ACCOUNTCFG[cs] QLIST=inst FSTARGET=10
ACCOUNTCFG[math] QLIST=inst FSTARGET=15

ACCOUNTCFG[phys] QLIST=supp FSTARGET=5
ACCOUNTCFG[eng] QLIST=premium FSTARGET=70
# handle per-user priority exceptions
USERCFG[tom] PRIORITY=100
USERCFG[john] PRIORITY=35
```

```
# define overall job priority
USERWEIGHT      10 # user exceptions
# relative FS weights (Note: QOS overrides ACCOUNT which overrides USER)
FSUSERWEIGHT    1
FSACCOUNTWEIGHT 10
FSQOSWEIGHT     100
# apply XFactor to balance cycle delivery by job size fairly
# Note: queuetime factor also on by default (use QUEUEWEIGHT to adjust)
XFACTORWEIGHT   100
# enable preemption
PREEMPTPOLICY   REQUEUE
# temporarily allow phys to preempt math
ACCOUNTCFG[phys] JOBFLAGS=PREEMPTOR PRIORITY=1000
ACCOUNTCFG[math] JOBFLAGS=PREEMPTTEE
```

### 5.3.3 Hierarchical Fairshare/Share Trees

Moab supports arbitrary depth hierarchical fairshare based on a share tree. In this model, users, groups, classes, and accounts can be arbitrarily organized and their usage tracked and limited. Moab extends common share tree concepts to allow mixing of credential types, enforcement of ceiling and floor style usage targets, and mixing of hierarchical fairshare state with other priority components.

**i** You can terminate your tnode with '</tnode>' or '<tnode />'.

#### 5.3.3.A Defining the Tree

The **FSTREE** parameter can be used to define and configure the share tree used in fairshare configuration. This parameter supports the following attributes:

SHARES	
<b>Format</b>	<COUNT>[@<PARTITION>][,<COUNT>[@<PARTITION>]]... where <COUNT> is a double and <PARTITION> is a specified partition name.
<b>Description</b>	The node target usage or share.
<b>Example</b>	<pre>FSTREE[Eng]    SHARES=1500.5 FSTREE[Sales] SHARES=2800</pre>
MEMBERLIST	
<b>Format</b>	Comma-delimited list of child nodes of the format [ <b>&lt;OBJECT_TYPE&gt;</b> ] : <b>&lt;OBJECT_ID&gt;</b> where object types are only specified for <i>leaf nodes</i>



MEMBERLIST	
	associated with <b>user</b> , <b>group</b> , <b>class</b> , <b>qos</b> , or <b>acct</b> credentials.
<b>Description</b>	The tree objects associated with this node.
<b>Example</b>	<pre> FSTREE[root]    SHARES=100    MEMBERLIST=Eng, Sales FSTREE[Eng]    SHARES=1500.5  MEMBERLIST=user:john,user:steve,user:bob FSTREE[Sales]  SHARES=2800    MEMBERLIST=Sales1, Sales2, Sales3 FSTREE[Sales1] SHARES=30     MEMBERLIST=user:kellyp,user:sam FSTREE[Sales2] SHARES=10     MEMBERLIST=user:ux43,user:ux44,user:ux45 FSTREE[Sales3] SHARES=60     MEMBERLIST=user:robert,user:tjackson </pre>

Current tree configuration and monitored usage distribution is available using the `mdiag -f -v` commands.

### 5.3.3.B Controlling Tree Evaluation

Moab provides multiple policies to customize how the share tree is evaluated.

Policy	Description
<b>FSTREETIERMULTIPLIER</b>	Decreases the value of sub-level usage discrepancies. It can be a positive or negative value. When positive, the parent's usage in the tree takes precedence; when negative, the child's usage takes precedence. The usage amount is not changed, only the coefficient used when calculating the value of fstree usage in priority. When using this parameter, it is recommended that you research how it changes the values in <code>mdiag -p</code> to determine the appropriate use.
<b>FSTREECAP</b>	Caps lower level usage factors to prevent them from exceeding upper tier discrepancies.

## Using FS Floors and Ceilings with Hierarchical Fairshare

All standard fairshare facilities including target floors, target ceilings, and target caps are supported when using hierarchical fairshare.

### Multi-Partition Fairshare

Moab supports independent, per-partition hierarchical fairshare targets allowing each partition to possess independent prioritization and usage constraint settings. This is accomplished by setting the `PERPARTITIONSCHEDULING` attribute of the `FSTREE`

parameter to `TRUE` in `moab.cfg` and setting `partition="name"` in your `<fstree>` leaf.

```
FSTREE[tree]
<fstree>
  <tnode partition="slave1" name="root" type="acct" share="100" limits="MAXJOB=6">
    <tnode name="accta" type="acct" share="50" limits="MAXSUBMITJOBS=2 MAXJOB=1">
      <tnode name="fred" type="user" share="1" limits="MAXWC=1:00:00">
      </tnode>
    </tnode>
    <tnode name="acctb" type="acct" share="50" limits="MAXSUBMITJOBS=4 MAXJOB=3">
      <tnode name="george" type="user" share="1" >
      </tnode>
    </tnode>
  </tnode>
  <tnode partition="slave2" name="root" type="acct" share="100"
limits="MAXSUBMITJOBS=6 MAXJOB=5">
    <tnode name="accta" type="acct" share="50">
      <tnode name="paul" type="user" share="1">
      </tnode>
    </tnode>
    <tnode name="acctb" type="acct" share="50">
      <tnode name="ringo" type="user" share="1">
      </tnode>
    </tnode>
  </tnode>
</fstree>
```

**i** If no partition is specified for a given share value, then this value is assigned to the global partition. If a partition exists for which there are no explicitly specified shares for any node, this partition will use the share distribution assigned to the global partition.

## Dynamically Importing Share Tree Data

Share trees can be centrally defined within a database, flat file, information service, or other system and this information can be dynamically imported and used within Moab by setting the `FSTREE` parameter within the [Identity Managers](#). This interface can be used to load current information at startup and periodically synchronize this information with the master source.

## To Create a Fairshare Tree in a Separate XML File and Import it into Moab

1. Create a file to store your fair share tree specification. Give it a descriptive name and store it in your Moab home directory (`$MOABHOMEDIR` or `$MOABHOMEDIR/etc`). In this example, the file is called `fstree.dat`.
2. In the first line of `fstree.dat`, set `FSTREE[myTree]` to indicate that this is a fairshare file.
3. Build a tree in XML to match your needs. For example:

```
FSTREE [myTree]
<fstree>
<tnode name="root" share="100">
<tnode name="john" type="user" share="50" limits="MAXJOB=8 MAXPROC=24
MAXWC=01:00:00"></tnode>
<tnode name="jane" type="user" share="50" limits="MAXJOB=5"></tnode>
</tnode>
</fstree>
```

*This configuration creates a fairshare tree where users share a value of 100. Users john and jane share the value equally, because each has been given 50.*

Because 100 is an arbitrary number, users john and jane could be assigned 10000 and 10000 respectively and still have a 50% share under the parent leaf. To keep the example simple, however, it is recommended that you use 100 as your arbitrary share value and distribute the share as percentages. In this case, john and jane each have 50%.

If the users' numbers do not add up to at least the fairshare value of 100, the remaining value is shared among all users under the tree. For instance, if the tree had a value of 100, user john had a value of 50, and user jane had a value of 25, then 25% of the fairshare tree value would belong to all other users associated with the tree. By default, tree leaves do not limit who can run under them.

**i** Each value specified in the `tnode` elements must be contained in quotation marks.

- Optional: Share trees defined within a flat file can be cumbersome; consider running `tidy` for `xml` to improve readability. Sample usage:

```
> tidy -i -xml mam-tiy.cfg <filename> <output file>

# Sample output

FSTREE [myTree]
<fstree>
  <tnode name="root" share="100">
    <tnode name="john" type="user" share="50" limits="MAXJOB=8
      MAXPROC=24 MAXWC=01:00:00">
    </tnode>
    <tnode name="jane" type="user" share="50" limits="MAXJOB=5">
    </tnode>
  </tnode>
</fstree>
```

- Link the new file to Moab using the `IDCFG` parameter in your Moab configuration file:

```
IDCFG [myTree] server="FILE:/// $MOABHOMEDIR/etc/fstree.dat" REFRESHPERIOD=INFINITY
```

*Moab imports the myTree fairshare tree from the `fstree.dat` file. Setting `REFRESHPERIOD` to `INFINITY` causes Moab to read the file each time it starts or restarts, but setting a positive interval (e.g., `4:00:00`) cause Moab to read the file more often. See [Refreshing Identity Manager Data](#) for more information.*

6. To view your fairshare tree configuration, run `mdiag -f`. If it is configured correctly, the tree information will appear beneath all the information about your fairshare settings configured in `moab.cfg`.

```
> mdiag -f
Share Tree Overview for partition 'ALL'
Name          Usage      Target      (FSFACTOR)
-----
root          100.00    100.00 of 100.00 (node: 1171.81) (0.00)
- john        16.44     50.00 of 100.00 (user: 192.65) (302.04) MAXJOB=8
MAXPROC=24 MAXWC=3600
- jane        83.56     50.00 of 100.00 (user: 979.16) (-302.04) MAXJOB=5
```

*The settings you configured in `fstree.dat` appear in the output. The tree of 100 is shared equally between users `john` and `jane`.*

## Specifying Share Tree Based Limits

Limits can be specified on internal nodes of the share tree using standard [credential limit semantics](#). The following credential usage limits are valid:

- `MAXIJOB` (Maximum number of idle jobs allowed for the credential)
- `MAXJOB`
- `MAXMEM`
- `MAXNODE`
- `MAXPROC`
- `MAXSUBMITJOBS`
- `MAXWC`

*Example 5-7: FSTREE limits example*

```
FSTREE[myTree]
<fstree>
  <tnode name="root" share="100">
    <tnode name="john" type="user" share="50" limits="MAXJOB=8
      MAXPROC=24 MAXWC=01:00:00">
    </tnode>
    <tnode name="jane" type="user" share="50" limits="MAXJOB=5">
    </tnode>
  </tnode>
</fstree>
```

## Specifying a Default Account in a Fair Share Tree

The `adef` attribute in a fair share tree can be used to specify a default account for a credential and its children. This is useful for sites with many users who need access to an account and who use an identity manager to import credentials.

The rules are as follows:

- When you define an adef attribute on a credential such as a user or qos, then the child under the tnode inherits the credential.
- When a child has an adef, then that adef overrides the parent.

To define an adef, add a qdef attribute to the tnode for the user:

```
<tnode name="jane" type="user" adef="acct2" share="1"
qlist="batch, special, standby, test, exempt, expedite, super"> </tnode>
```

## Other Uses of Share Trees

If a share tree is defined, it can be used for purposes beyond fairshare, including organizing general usage and performance statistics for reporting purposes (see [showstats -T](#)), enforcement of tree node based usage limits, and specification of resource access policies.

---

### Related Topics

- [mdiag -f](#) command (provides diagnosis and monitoring of the fairshare facility)
- [FSENABLECAPRIORITY](#) parameter
- [ENABLEFSPREEMPTION](#) parameter
- [FSTARGETISABSOLUTE](#) parameter

## 5.4 Sample FairShare Data File

FS.<EPOCHTIME>

```
# FS Data File (Duration: 43200 seconds) Starting: Sat Jul 8 06:00:20
user          jvella      134087.910
user          reynolds    98283.840
user          gastor      18751.770
user          uannan     145551.260
user          mwillis    149279.140
...
group         DEFAULT    411628.980
group         RedRock    3121560.280
group         Summit     500327.640
group         Arches     3047918.940
acct          Administration 653559.290
acct          Engineering 4746858.620
acct          Shared     75033.020
acct          Research   1605984.910
qos           Deadline   2727971.100
qos           HighPriority 4278431.720
qos           STANDARD   75033.020
class         batch      7081435.840
sched         iCluster   7081435.840
```

*The total usage consumed in this time interval is 7081435.840 processor-seconds. Since every job in this example*

scenario had a user, group, account, and QOS assigned to it, the sum of the usage of all members of each category should equal the total usage value:  $USERA + USERB + USERC + USERD = GROUPA + GROUPB = ACCTA + ACCTB + ACCTC = QOS0 + QOS1 + QOS2 = SCHED$ .

## 5.5 Accounting, Charging, and Allocation Management

In this topic:

- [5.5.1 Accounting Manager Overview - page 418](#)
- [5.5.2 Accounting Mode - page 419](#)
- [5.5.3 Accounting Manager Interface Types - page 420](#)
- [5.5.4 Charging for Jobs - page 423](#)
- [5.5.5 Charging for Reservations - page 424](#)
- [5.5.6 Accounting Properties Reported to the Accounting Manager - page 425](#)
- [5.5.7 Accounting Stages - page 430](#)
- [5.5.8 Accounting Events - page 433](#)
- [5.5.9 Blocking Versus Non-Blocking Accounting Actions - page 433](#)
- [5.5.10 Retrying Failed Charges - page 434](#)

For a complete list of and additional information on the AMCFG parameters and flags, see [5.6 AMCFG Parameters and Flags - page 437](#).

### 5.5.1 Accounting Manager Overview

An accounting manager is a software system that enables tracking and charging for job resource usage. Moab Accounting Manager is a commercial charge-back accounting system that has built-in integration with Moab Workload Manager. Moab Accounting Manager can be used in a variety of accounting modes such as for usage tracking, notional charging or allocation enforcement.

When used for usage tracking only, the accounting manager simply records workload usage details. When configured additionally to perform charging, resource charge rates are used to impute a charge for each job. When configured to enforce resource allocation limits, jobs are charged against allocations and new jobs may be blocked from running if their account runs out of funds. See [Accounting Mode](#) and see 'Select an Appropriate Accounting Mode' in the *Moab Accounting Manager Administrator Guide* for more details on supported accounting modes.

In a typical allocation enforcement use case, credits are allocated to accounts for designated time periods; establishing limits on the use of compute resources. The base


currency credits can be defined in terms of system resource units (e.g., Processor-Seconds) or a real currency (e.g., U.S. dollars). Charge rates are established for the use of resources. Accounts are created and users are given access to the appropriate accounts. Deposits are made into funds associated with the account's creating allocations. An allocation cycle can be established whereby funds are reset on a regular periodic basis (such as yearly, quarterly, or monthly) and where allocations are renewed for accepted accounts. Before a job is started, Moab Workload Manager will verify that the user has sufficient credits to run the job by attempting to place a hold against their funds (referred to as a lien). When a job completes, the user's funds will be debited via a charge, usage information will be recorded for the job, and the lien will be removed.

## 5.5.2 Accounting Mode

The accounting mode (specified via the `AMCFG[] MODE` parameter) modifies the way in which accounting-relevant job and reservation stages (e.g., create, start, end, etc.) are processed. See [5.5.7 Accounting Stages - page 430](#) for more information on the behaviors of the different values of the accounting mode.

The following table describes the valid values for the accounting mode:

Value	Description
<b>strict-allocation</b>	Use this mode if you want to strictly enforce allocation limits. Under this mode, holds (called liens) will be placed against allocations in order to prevent multiple jobs from starting up on the same funds. Jobs and reservations can be prevented from running if the end-users do not have sufficient funds. This is the default.
<b>fast-allocation</b>	Use this mode if you want to debit allocations, but need higher throughput by eliminating the lien and quote operations of strict-allocation mode. Under this mode, jobs and reservations check a cached account balance, and can be prevented from running after the balance has become zero or negative.

 The cached account balances in Moab can be viewed by running `'mrmctl -q AccountBalanceCache AM'`.

Value	Description
	<p><b>i</b> If you are using fast-allocation, funds are assumed to have account-based constraints only. Moab will reject funds having no constraints or having non-account constraints. It is highly recommended that you enable <code>ENFORCEACCOUNTACCESS</code> to <code>TRUE</code> and <code>AMCFG[] CREATECRED=TRUE</code> with an appropriate refresh period (via <code>AMCFG[] REFRESHPERIOD</code>) so that Moab can prevent jobs from running under accounts that the user does not belong to (this is enforced via liens in the strict-allocation accounting mode). Also, the configured refresh period will apply to both credential updates and account balance updates. See <a href="#">Moab Parameters</a> for more information on the <code>ENFORCEACCOUNTACCESS</code> parameter.</p>
<b>notional-charging</b>	Use this mode if you want to calculate and record charges for workload usage, but not keep track of fund balances or allocation limits.
<b>usage-tracking</b>	Use this mode if you want to record workload usage details, but not to calculate a charge nor keep track of fund balances or allocation limits.

### 5.5.3 Accounting Manager Interface Types

Moab Workload Manager supports two accounting manager interface types:

- **MAM** - When using the MAM interface type, Moab communicates directly over the network with Moab Accounting Manager using the SSS wire protocol.
- **Native** - When using the Native accounting manager interface type, Moab invokes scripts that can be customized to interact with Moab Accounting Manager or other third party accounting systems.

#### 5.5.3.A MAM

The MAM accounting manager interface type enables direct communication between Moab Workload Manager and Moab Accounting Manager. This often results in the fastest accounting performance. Use this interface type if you do not need to customize the interaction with the accounting manager.

To configure Moab to use the MAM accounting manager interface, run `configure` using the `--with-am` option.



**Example 5-8:**

```
./configure --with-am=mam ...
```

Consequently, `make install` will add the essential configuration and connection entries into the `moab.cfg` and `moab-private.cfg` files.

The following are typical entries in the Moab configuration files for using the MAM interface:

- `moab.cfg`

```
AMCFG[mam] TYPE=MAM HOST=localhost
```

- `moab-private.cfg`

```
CLIENTCFG[AM:mam] KEY=UiW7EihzKyUyVQg6dKirDhV3
```

Synchronize the secret key with Moab Accounting Manager by copying the value of the `token.value` parameter from the `MAM_PREFIX/etc/mam-site.conf` file, which is randomly generated during the Moab Accounting Manager install process.

When using the MAM accounting manager interface, by default Moab will communicate directly with Moab Accounting Manager via the SSS wire protocol. However, it is possible to enable a hybrid model and override individual accounting actions by specifying the `exec` protocol and the path of a custom script to the appropriate `AMCFG[] *URL` parameters.

Moab Accounting Manager should be installed, started, and initialized. See 'Initial Setup' in the *Moab Accounting Manager Administrator Guide* for examples of how to initialize MAM for your initial mode of operation.

### 5.5.3.B Native

The Native accounting manager interface type provides a customization layer between Moab Workload Manager and Moab Accounting Manager. This interface can be used where greater accounting customization is required. The native interface can also be customized to interact with third-party accounting manager systems. Moab passes job accounting details to scripts that handle the interaction with the external system.

To configure Moab to use the MAM accounting manager interface, run `configure` using the `--with-am=native` option.

Additionally, you might need to use the `--with-am-dir` `configure` option to specify the prefix directory for Moab Accounting Manager if MAM has been installed in a non-default location.

**Example 5-9:**

```
./configure --with-am=native ...
```

Consequently, make `install` will add the essential accounting manager entries into `moab.cfg` and install the accounting-related scripts (`$(PREFIX)/tools/mam/usage.*.mam.pl`) in the correct locations..

Moab will default to using a set of stock scripts for the accounting stages. To view the scripts that are currently in use, run `mdiag -R -v am:` (even more information may be available in `mdiag -R -v --xml`). The following shows sample output from running the `mdiag -R -v am:` command:

```
AM[mam] Type: Native State: 'Active'
Timeout: 15
Accounting Mode: strict-allocation
Job Charge Policy: All
Reservation Charge Policy: Select
Create URL: exec:///opt/moab/tools/mam/usage.quote.mam.pl
Start URL: exec:///opt/moab/tools/mam/usage.reserve.mam.pl
Pause URL: exec:///opt/moab/tools/mam/usage.charge.mam.pl
Resume URL: exec:///opt/moab/tools/mam/usage.reserve.mam.pl
Update URL: exec:///opt/moab/tools/mam/usage.charge.mam.pl
Continue URL: exec:///opt/moab/tools/mam/usage.reserve.mam.pl
End URL: exec:///opt/moab/tools/mam/usage.charge.mam.pl
Delete URL: exec:///opt/moab/tools/mam/lien.delete.mam.pl
Query URL: exec:///opt/moab/tools/mam/account.query.mam.pl
Retry Failed Charges: TRUE
```

Moab will invoke the native accounting manager scripts by passing the job or reservation information via XML to the standard input of the script. You can override any of the default scripts with a custom script by specifying the appropriate AMCFG URL parameter in the moab server configuration file. See [5.6 AMCFG Parameters and Flags - page 437](#) for `CREATEURL`, `STARTURL`, `PAUSEURL`, `RESUMEURL`, `UPDATEURL`, `CONTINUEURL`, `ENDURL`, `DELETEURL`, and `QUERYURL` values for more information.

The XML sent to the scripts is in the form of an SSS Request that is identical to the Request sent to MAM when you use the MAM Accounting Manager Interface type. For example, the XML sent to the `usage.charge.mam.pl` script in a final charge consists of an encapsulating Request element with an action attribute that has a value of 'Charge'; an object element with a value of 'UsageRecord'; one or more optional Option elements; and a Data element. The Data element has a single UsageRecord element with property elements describing the job or reservation properties. For example:

```
<Request action="Charge"><Object>UsageRecord</Object><Option
name="Duration">1234</Option><Data><UsageRecord><Type>Job</Type><Instance>Moab.165</In
stance><User>amy</User><Group>staff</Group><Account>chemistry</Account><Class>batch</C
lass><QualityOfService>high</QualityOfService><Machine>colony</Machine><Nodes>1</Nodes
><NodeType>Fast</NodeType><NodeCharge>2.000000</NodeCharge><Partition>Torque</Partitio
n><Processors
consumptionRate="0.50">2</Processors><Memory>2048</Memory><Matlab>2</Matlab><StartTime
>1398805354</StartTime><EndTime>1398805357</EndTime><CompletionCode>0</CompletionCode>
<OpSys>CentOS</OpSys><Temp>87.00</Temp></UsageRecord></Data></Request> *
```

In the sample XML above, Matlab is an example of a generic resource, OpSys is an example of a job variable, and Temp is an example of a generic metric.

A reservation charge, or quote or lien, is very similar. For example:

```
<Request action="Charge"><Object>UsageRecord</Object><Option
name="Duration">7200</Option><Data><UsageRecord><Type>Reservation</Type><Instance>rese
rvation.7</Instance><User>amy</User><Machine>colony</Machine><Nodes>1</Nodes><Processo
rsconsumptionRate="0.76">12</Processors><Duration>7200</Duration><StartTime>1398797430
</StartTime><EndTime>1398804630</EndTime></UsageRecord></Data></Request>
```

The majority of the scripts use this same basic XML format; for instance, `usage.quote.mam.pl`, `usage.reserve.mam.pl`, and `usage.charge.mam.pl`.

The XML sent to the `lien.delete.mam.pl` script to clean up after a failure consists of an encapsulating Request element with an action attribute that has a value of 'Delete'; an object element with the value of 'Lien'; and a condition (Where) element indicating the lien instance to delete. For example:

```
<Request action="Delete"><Object>Lien</Object><Where
name="Instance">Moab.127</Where></Request>
```

The script should return a return code (zero for success), data on standard out and messages on standard error. A failure in `CREATEURL`, `STARTURL`, `RESUMEURL`, or `CONTINUEURL` should result in the application of the `CREATEFAILUREACTION`, `STARTFAILUREACTION`, `RESUMEFAILUREACTION`, or `CONTINUEFAILUREACTION` respectively.

Moab Accounting Manager should be installed, started, and initialized. The simplest procedure is to install it on the same server as Moab Workload Manager so that the Moab Accounting Manager can share libraries and configuration files with the Moab Workload Manager and Moab Accounting Manager scripts. See 'Initial Setup' in the *Moab Accounting Manager Administrator Guide* for examples of how to initialize MAM for your initial mode of operation.

## 5.5.4 Charging for Jobs

Moab can be configured to charge for the resources used in jobs.

Job tracking and charging via an accounting manager is enabled or disabled by the `AMCFG [ ] JOBCHARGEPOLICY` parameter. By default, if an accounting manager is defined in Moab, Moab will charge for all jobs, independent of the job's exit status (derived from an internal default of `AMCFG [ ] JOBCHARGEPOLICY=All`). This is the policy that is most frequently used because of the unfortunate fact that a clever user can fake a job failure to avoid being charged for their job. Jobs that are discovered to have failed due to system issues can be proactively refunded by the system admin, or the user may be required to apply for a refund. Charges can be restricted to jobs that complete successfully by setting `AMCFG [ ] JOBCHARGEPOLICY=Successful`. To disable charging for jobs entirely, set `AMCFG [ ] JOBCHARGEPOLICY=None`.

When using the default accounting mode of strict-allocation, before Moab starts a job, it contacts the accounting manager and requests an allocation reservation (or lien) be placed on the associated account. The lien amount is equivalent to the total cost of resources that

could be consumed by the job (based on the job's wallclock limit) and is used to prevent the possibility of allocation oversubscription. Moab then starts the job. When the job completes, Moab debits the allocation by the amount actually consumed by the job and then releases the lien.

These steps should be transparent to users. Only when an account has insufficient allocations to run a requested job will the presence of the accounting manager be noticed. The policies guiding what action should be taken if a user is out of funds is dictated by the `AMCFG[] STARTFAILUREACTION` (and related) parameters. If desired, a fallback account can be specified for use when a job's primary account is out of allocations. This account, specified using the `AMCFG` parameter's `FALLBACKACCOUNT` attribute, is often associated with a low QoS privilege, priority, and cost and is often configured to run only when no other jobs are present.

The actual policies that determine what resources are charged for and in what amounts are specified by the charge rates defined in the accounting manager. Moab will pass the job properties (shown in [5.5.6 Accounting Properties Reported to the Accounting Manager - page 425](#)) to the accounting manager. It is the task of the accounting manager to record and charge for the job according to site objectives.

## 5.5.5 Charging for Reservations

Moab can be configured to charge for the unused cycles in reservations. One of the hesitations with dedicating resources to a particular group is that if the resources are not used by that group, they go idle and are wasted. By configuring a reservation to be chargeable, sites can charge every idle cycle of the reservation to a particular account. When the reservation is in use, the consumed resources will be charged to the jobs using the resources. When the resources are idle, the resources will be charged to the reservation's charge account. See `AMCFG[] RESERVATIONCHARGEPOLICY`.

If `RESERVATIONCHARGEPOLICY` is set to `Select` (the default), charging can be enabled for select reservations by specifying the `CHARGEACCOUNT` and `CHARGEUSER` attributes for the reservation. For standing reservations, these are set via the `SRCFG[X] CHARGEACCOUNT` and `CHARGEUSER` parameters. For administrative reservations, these are set via the `-S` `aaccount` and `auser` options.

*Example 5-10: Enabling charging in a standing reservation*

```
SRCFG[foo] PERIOD=DAY DAYS=Mon,Tue,Wed,Thu,Fri DEPTH=1 USERLIST=amy
CHARGEACCOUNT=chemistry CHARGEUSER=amy RESOURCES=PROCS:1 TASKCOUNT=2
```

*Example 5-11: Enabling charging in an administrative reservation*

```
mrsvctl -c -a USER=amy -S aaccount=chemistry -S auser=amy -R procs=1 -t 1 -d 7200
```

If `RESERVATIONCHARGEPOLICY` is set to `All`, idle cycles will be charged for all reservations unless disabled for individual reservations by specifying the reservation `Charge` attribute with a value of `False`. For standing reservations, this are set via the `SRCFG CHARGE` parameter. For administrative reservations, this is set via the `-S charge` options.

*Example 5-12: Disabling charging in a standing reservation*

```
SRCFG[foo] PERIOD=DAY DAYS=Mon,Tue,Wed,Thu,Fri DEPTH=1 USERLIST=amy CHARGE=False
RESOURCES=PROCS:1 TASKCOUNT=2
```

*Example 5-13: Disabling charging in an administrative reservation*

```
mrsvctl -c -a USER=amy -S charge=False -R procs=1 -t 1 -d 7200
```

If `RESERVATIONCHARGEPOLICY` is set to `None`, idle cycles will not be charged for any reservations, regardless of what attributes are specified with the reservation.

When utilizing the accounting manager to track or charge for idle cycles in any reservations, the accounting manager must be configured to track and charge for reservation-relevant properties. See the examples for 'Enabling Reservation Statistics' and 'Charging for the unused cycles in reservations' in the *Moab Accounting Manager Administrator Guide* for steps on how to do this for the Moab Accounting Manager.

## 5.5.6 Accounting Properties Reported to the Accounting Manager

Moab can send the following information to the accounting manager via charging actions.

### 5.5.6.A For Jobs

Property name in the Accounting Manager Usage Record	Description of property value recorded in the Accounting Manager Usage Record
<b>Account</b>	Account name.
<b>BlockedProcessors*</b>	Number of processors blocked from use to other jobs. This might be more than the allocated or requested processors if entire nodes are given exclusive access to the job (e.g., from a node access policy or node exclusivity flag).

Property name in the Accounting Manager Usage Record	Description of property value recorded in the Accounting Manager Usage Record
	<div style="border: 1px solid black; padding: 10px;"> <p><b>i</b> Tracking or charging for BlockedProcessors should only be used with policies that allow only a single job to dedicate a node such as with a Node Access Policy of SINGLEJOB or SINGLETASK, or using a QOS with the DEDICATED flag. Using BlockedProcessors with any policy allowing more than one job to dedicate a node (such as a Node Access Policy of SINGLEUSER, SINGLECLASS, SINGLEACCOUNT or UNIQUEUSER) will yield inconsistent results and is not recommended or supported.</p> </div>
<b>Charge</b>	If the AMCFG <b>LOCALCOST</b> flag is set, Moab will calculate and pass the Charge amount to MAM. If it is not, MAM will calculate the charge based on the transmitted job properties.
<b>Class</b>	Class/queue name.
<b>Cores*</b>	<p>NUMA cores allocated to the job.</p> <p>Cores will be reported to the accounting manager when using the -L NUMA Resource Request syntax and specifying place=socket, place=numanode or place=core.</p>
<b>CPUTime</b>	<p>CPU time. The value sent by Moab to the accounting manager is the cumulative CPU time for the job.</p> <div style="border: 1px solid black; padding: 10px;"> <p><b>i</b> Using this value as the basis of charging is not compatible with periodic charging.</p> </div>
<b>Duration</b>	Moab sends the wallclock time for the job charge(s) in seconds. This is aggregated in MAM as Duration.
<b>EndTime</b>	Job end time.
<b>Exit Code</b> (formerly CompletionCode*)	Exit code.
<b>Features*</b>	Allocated node features.

Property name in the Accounting Manager Usage Record	Description of property value recorded in the Accounting Manager Usage Record
	<p>Node features are sent in the JSON object format { "&lt;node_feature_or_combo_name&gt;":&lt;node_feature_or_combo_count&gt;, ... }.</p> <p>Node features are not passed to the accounting manager unless configured to do so via the AMCFG[] IncludeFeatures parameter.</p>
<b>GPUs*</b>	Number of GPUs allocated to the job.
<b>Group</b>	Group name.
<b>Instance</b>	Job ID.
<b>Licenses</b>	<p>License generic resources allocated to the job. License generic resources are sent in the JSON object format { "&lt;license_generic_resource_name&gt;":&lt;license_generic_resource_value&gt;, ... }. The license generic resource value is the number of this license generic resource consumed by the job.</p> <p>License generic resources are distinguished by those generic resources having the license trait. Generic resources can be marked as license generic resources in two primary ways:</p> <ul style="list-style-type: none"> <li>• Any generic resources seen on nodes reported via a native resource manager having the server configuration parameter RMCFG[] RESOURCETYPE=License will be marked as a license generic resource.</li> <li>• Generic resources defined in Moab can be identified as license types by using the Moab configuration parameter GRESCFG[] LICENSE=TRUE.</li> </ul> <p>You can distinguish the generic resources having the license trait via the command `mschedctl -l gres -v`.</p>
<b>Machine</b>	Cluster (RM) name.
<b>Memory</b>	Dedicated or utilized memory in megabytes.
<b>Metrics</b>	<p>Generic metrics are sent in the JSON object format { "&lt;generic_metric_name&gt;":&lt;generic_metric_value&gt;, ... } where the generic metric value is the average value of the generic metric across the nodes of the job and across time.</p>

Property name in the Accounting Manager Usage Record	Description of property value recorded in the Accounting Manager Usage Record
<b>MICs*</b>	Number of MICs allocated to the job.
<b>NodeCharge*</b>	Aggregate node charge rate. See <a href="#">NODECHARGEPOLICY</a> and <a href="#">CHARGERATE</a> for more information.
<b>Nodes</b>	Node count.
<b>NodeType*</b>	Node type. See <a href="#">NODETYPE</a> for more information.
<b>NumaNodes*</b>	NUMA nodes allocated to the job. NumaNodes will be reported to the accounting manager when using the -L NUMA Resource Request syntax and specifying place=socket or place=numanode.
<b>Partition*</b>	Partition name.
<b>ProcessorEquivalents*</b>	Processor Equivalents.
<b>Processors</b>	Number of processors allocated to the job. This is normally equivalent to the requested processors.
<b>QualityOfService</b>	QoS name.
<b>QueueDuration*</b>	Effective duration the job was in the idle state.
<b>RequestedDuration</b>	Requested wallclock limit.
<b>Resources</b>	Generic resources are sent in the JSON object format <code>{"&lt;generic_resource_name&gt;":&lt;generic_resource_value&gt;, ...}</code> . The generic resource value is the number of this generic resource consumed by the job.
<b>Sockets*</b>	NUMA sockets allocated to the job. Sockets will be reported to the accounting manager when using the -L NUMA Resource Request syntax and specifying place=socket.
<b>Stage</b>	Accounting stage.



Property name in the Accounting Manager Usage Record	Description of property value recorded in the Accounting Manager Usage Record
<b>StartTime</b>	Job start time.
<b>SubmitTime</b>	Job submission time.
<b>Threads*</b>	NUMA threads allocated to the job. Threads will be reported to the accounting manager when using the -L NUMA Resource Request syntax and specifying place=socket, place=numanode, place=core or place=thread.
<b>Type</b>	Set to 'Job'.
<b>User</b>	User name.
<b>Variables</b>	Job variables are sent in the JSON object format {"<job_variable_name>": "<job_variable_value>", ...}.

\* For this property to be recorded in the MAM Usage Record, you must define a custom usage record attribute in MAM for it. See 'Customizing the Usage Record Object' in the *Moab Accounting Manager Administrator Guide* for more information.

### 5.5.6.B For Reservations

Property name in MAM Usage Record	Description of property value recorded in MAM Usage Record
<b>Account</b>	Charge account.
<b>Duration</b>	Moab sends the wallclock time for the reservation in seconds. This is aggregated in MAM as Duration.
<b>EndTime</b>	Reservation end time.
<b>IdleProcessorSeconds</b>	Processor seconds not blocked by jobs within the reservation. This is the metric typically used to charge for spare cycles not attributed to jobs within the reservation.
<b>Instance</b>	Reservation ID.

Property name in MAM Usage Record	Description of property value recorded in MAM Usage Record
<b>Machine</b>	Cluster (RM) name.
<b>Nodes</b>	Number of node allocated to the reservation.
<b>Partition</b>	Partition name.
<b>Processors</b>	Number of processors allocated to the reservation.
<b>ReservedProcessorSeconds</b>	Total processor seconds included within the reservation (to-date)
<b>Stage</b>	Accounting stage.
<b>StartTime</b>	Reservation start time.
<b>Type</b>	Set to 'Reservation'.
<b>User</b>	Charge user or reservation owner.

## 5.5.7 Accounting Stages

The accounting manager performs various actions throughout different stages of a job or reservation lifetime. For a stock configuration (meaning you have not overridden the accounting actions with custom scripts), the following describes the stages and the respective actions that occur at these stages depending on the accounting mode:

- **Create stage** – When a job is submitted or a chargeable reservation is created and either `AMCFG[] VALIDATEJOB SUBMISSION` is `TRUE` or an `AMCFG[] FALLBACKACCOUNT` or `FALLBACKQOS` are specified:
  - If the accounting mode is `strict-allocation`, Moab will check with the accounting manager to verify that sufficient funds exist for the job or reservation to run.
  - If the accounting mode is `fast-allocation`, Moab will check its cached balance for the job's or reservation's account, to verify that sufficient funds exist for the fund or reservation to run.
  - Otherwise, it does nothing.

- **Start stage** – When a job or a chargeable reservation is about to start:
  - If the accounting mode is `strict-allocation`, Moab will attempt to place a hold against the allocation in the accounting manager in order to prevent multiple jobs or reservations from starting on the same funds.
  - If the accounting mode is `fast-allocation`, Moab will check its cached balance for the job's or reservation's account, to verify that sufficient funds exist for the job or reservation to run.\*
  - Otherwise, it does nothing.
- **Delete stage** – If a job or chargeable reservation fails to start:
  - If the accounting mode is `strict-allocation` and Moab has already placed a hold on an allocation for the job or reservation, Moab will contact the accounting manager to remove the lien.
  - Otherwise, it does nothing.
- **Pause stage** – If a job becomes suspended, Moab will make a charge for the resources used for the time the job has run this far:
  - If the accounting mode is `strict-allocation`, the usage record will be updated with resource usage and charge amounts, the allocation will be debited, and the lien will be reduced.
  - If the accounting mode is `fast-allocation`, the usage record will be updated with resource usage and charge amounts, and the allocation will be debited.
  - If the accounting mode is `notional-charging`, the usage record will be updated with resource usage and charge amounts.
  - If the accounting mode is `usage-tracking`, the usage record will be updated with resource usage.
- **Resume stage** – If a suspended job is resumed:
  - If the accounting mode is `strict-allocation`, Moab will attempt to place a hold against the funds in the accounting manager for the smaller of (the duration of the next charge period, or the remaining duration of the job or reservation).
  - If the accounting mode is `fast-allocation`, Moab will check its cached balance for the job's or reservation's account, to verify that sufficient funds exist for the job or reservation to run for the smaller of (the duration of the next charge period, or the remaining duration of the job or reservation).\*
  - Otherwise, it does nothing.
- **Update stage** – If `AMCFG[] FLUSHINTERVAL` is set and Moab has reached the end of a charge period, Moab will make an incremental charge for all running jobs and

active chargeable reservations for the resources used during the last charge period:

- If the accounting mode is `strict-allocation`, the usage record will be updated with resource usage and charge amounts, the allocation will be debited, and the lien will be reduced.
  - If the accounting mode is `fast-allocation`, the usage record will be updated with resource usage and charge amounts, and the allocation will be debited.
  - If the accounting mode is `notional-charging`, the usage record will be updated with resource usage and charge amounts.
  - If the accounting mode is `usage-tracking`, the usage record will be updated with resource usage.
- **Continue stage** – If `AMCFG[] FLUSHINTERVAL` is set and Moab is beginning a new charge period for a job or reservation:
    - If the accounting mode is `strict-allocation`, Moab will attempt to place a hold against the funds in the accounting manager for the smaller of (the duration of the next charge period, or the remaining duration of the job or reservation).
    - If the accounting mode is `fast-allocation`, Moab will check its cached balance for the job's or reservation's account, to verify that sufficient funds exist for the job or reservation to run for the smaller of (the duration of the next charge period, or the remaining duration of the job or reservation)\*
    - Otherwise, it does nothing.
  - **End stage** – If a job or chargeable reservation ends, Moab will make a final charge for the remainder of the resources used by the job or reservation:
    - If the accounting mode is `strict-allocation`, the usage record will be updated with resource usage and charge amounts, the allocation will be debited, and the lien will be removed.
    - If the accounting mode is `fast-allocation`, the usage record will be updated with resource usage and charge amounts, and the allocation will be debited.
    - If the accounting mode is `notional-charging`, the usage record will be updated with resource usage and charge amounts.
    - If the accounting mode is `usage-tracking`, the usage record will be updated with resource usage.

\* The cached account balances in Moab can be viewed by running `'mrmctl -q AccountBalanceCache AM'`.

## 5.5.8 Accounting Events

You can add accounting events to the event log by specifying one or more of the following with `RECORDEVENTLIST`:

Event	Description
<b>AMCREATE</b>	Record accounting events trigger when an object is created; for example, when a balance check occurs at job submission.
<b>AMDELETE</b>	Record accounting events triggered when an object's normal accounting lifecycle is interrupted; for example, when the lifecycle is interrupted to clean up reservations for a failed job start.
<b>AMEND</b>	Record accounting events triggered when an object ends; for example, when a charge occurs at the end of a job.
<b>AMPAUSE</b>	Record accounting events triggered when an object is paused; for example, when a partial charge occurs when a job is paused.
<b>AMQUOTE</b>	Record accounting events triggered when an object requires a quote amount.
<b>AMRESUME</b>	Record accounting events triggered when an object is resumed; for example, when a lien is made when a job is resumed.
<b>AMSTART</b>	Record accounting events triggered when an object is started; for example, when a lien is made when a job starts.
<b>AMUPDATE</b>	Record accounting events triggered when an object continues past a flush interval; for example, when a partial charge occurs and new lien is made for a job.

## 5.5.9 Blocking Versus Non-Blocking Accounting Actions

Moab uses a thread pool to perform non-blocking actions. Instead of blocking the scheduling thread, the request is added to a queue that is serviced by the accounting thread pool. Using the thread pool to perform non-blocking accounting actions can result in faster aggregate scheduling and better client response times, though individual actions can, in some cases, be shortly delayed. By default, Moab uses non-blocking calls for the final charge *only*. The default behavior for individual accounting actions (such as Create, Start, End) can be overridden via the associated parameter (`CONTINUEISBLOCKING`, `CREATEISBLOCKING`, `DELETEISBLOCKING`, `ENDISBLOCKING`, `PAUSEISBLOCKING`, `RESUMEISBLOCKING`, `STARTISBLOCKING`).

**i** For best performance when using non-blocking accounting actions, it is recommended to specify an RM poll interval with a minimum poll time of zero (such as `RM_POLLINTERVAL=0, 30`). Setting a non-zero minimum poll time can prevent Moab from responding quickly to accounting actions and can result in increased latency in job scheduling.

**i** When using the `fast-allocation` accounting mode, if the charge action is set to be non-blocking (which is the default), Moab's account balance cache is not updated with the effects of the charge until the iteration after the charge is issued.

### 5.5.10 Retrying Failed Charges

If the `AMCFG[] RETRYFAILEDCHARGES` parameter is set to true (this is the default), job charges will be retried if they have failed due to a connection failure. When a job charge or usage record update (such as might occur when a job is suspended, at the periodic charge interval, or when a job completes) results in a connection failure between Moab and the accounting manager, then the charge request will be saved to a file in `SPOOLDIR/am/retrying/`. Once Moab detects that the connection with the accounting manager has been restored, the charge will be retried up to `CHARGERETRYCOUNT` times.

Charges that fail due to reasons other than a connection failure, or connection failures that surpass the `CHARGERETRYCOUNT`, will be saved to files in `SPOOLDIR/am/failed/`. Although these failures generally represent permanent failures, in some cases it might be possible to reissue some of these charges with a slight modification. For example, a user might have been moved from one account to another after the job started causing the final charge to fail. For such circumstances, a script has been provided (`TOOLS_DIR/mam/mam-charge-retry.pl`) to facilitate the re-issuance of a failed usage charge from a failed charge retry file.

```
[root]# /software/moab-accounting/tools/mam/mam-charge-retry.pl --help
```

The `mam-charge-retry.pl` script mimics the `mam-charge` command in making a charge to MAM. The specified command-line options will override the original values contained in the failed charge file. The `--dry-run` option can be used to issue the retry as a quote rather than a charge in order to see if the charge would be successful. The `--delete-on-success` option can be used to delete the retry file after a successful charge. This script cannot be used to rerun a command when the accounting action uses a native script. In such cases, the modified request XML from the charge retry file can be passed as the standard input to the native script to reissue a charge.

### 5.5.10.A Using the Script

This section provides synopsis information and an example on using the `TOOLSDIR/mam/mam-charge-retry.pl` script

#### Synopsis

```
mam-charge-retry {[--filename] <retry_filename>} [-J <instance_name>] [-j
<usage_record_id>] [-q <quote_id>] [-l <lien_id>] [-T <usage_record_
type>] [-u <user_name>] [-g <group_name>] [-a <account_name>] [-o
<organization_name>] [-c <class_name>] [-Q <quality_of_service>] [-m
<machine_name>] [-N <nodes>] [-P <processors>] [-C <cpu_time>] [-M
<memory>] [-D <disk>] [--stage <lifecycle_stage>] [-X, --extension
<property>=<value>]... [-t <charge_duration>] [-s <charge_start_time>]
[-e <charge_end_time>] [-d <charge_description>] [-z <charge_amount>]]
[-f <fund_id>]] [--incremental] [-R <charge_rate_name>][{<charge_rate_
value>}] = <charge_rate_amount>]... [--hours] [--itemize] [--delete-on-success] [-
-dry-run] [--debug] [--site <site_name>]] [--help] [--man] [--quiet] [--verbose] [--
version]
```

\* For this property to be recorded in the MAM Usage Record, you must define a custom usage record attribute in MAM for it. See 'Customizing the Usage Record Object' in the *Moab Accounting Manager Administrator Guide* for more information.

#### Reissuing a Charge that has Failed Example

First we will list the files in the `SPOOLDIR/am/failed` directory to see if there are any 'permanently' failed charges that we might want to reissue.

```
[root]# cd /opt/moab/spool/am/failed
[root]# ls
job.250
```

We see there is a failed charge for job 250. It might be useful to check the charge file and examine the message to see what went wrong.

```
[root]# cat job.250
{"action":"End","message":"Failure registering job End (250) with accounting manager -
- Unable to invoke AM request - server rejected request with status code 740 - Failed
charging 1.00 credits for instance 250 and created usage record
25\nUser amy is not a valid member of Account biology","request":{"<Request
action=\<Charge\><Object>UsageRecord</Object><Option name=\<AccountingMode\>strict-
allocation</Option><Option name=\<StartTime\>1432070300</Option><Option
name=\<Duration\>300</Option><Data><UsageRecord><Stage>End</Stage><Type>Job</Type><In
stance>
250</Instance><User>amy</User><Group>staff</Group><Account>biology</Account><Class>bat
ch</Class><QualityOfService>premium</QualityOfService>
```

```
<Machine>colony</Machine><Nodes>1</Nodes><Partition>colony</Partition><Processors
consumptionRate=\"1.00\">12</Processors><StartTime>1432070300</StartTime><SubmitTime>1
432070300</SubmitTime><EndTime>1432070600</EndTime>
<CompletionCode>0</CompletionCode></UsageRecord></Data></Request>"}

```

We can see that this charge failed because the user (amy) was not a member of the specified account (biology). In this case, the user was a member of the biology account when the job started, but had been moved to the account chemistry by the time the job ended, resulting in a charge failure.

If we were to reissue the charge without modification, it would fail again, as we can see by using the script with the `--dry-run` option.

```
[root]# /opt/moab/tools/mam/mam-charge-retry.pl job.250 --dry-run
User amy is not a valid member of Account biology

```

We can reissue the charge after changing the request to use her new chemistry account.

```
[root]# /opt/moab/tools/mam/mam-charge-retry.pl job.250 -a chemistry --dry-run
Successfully quoted 1.00 credits for instance 250

```

Since that looks like it will work correctly, we'll issue the corrected charge request and delete the charge file.

```
[root]# /opt/moab/tools/mam/mam-charge-retry.pl job.250 -a chemistry --delete-on-
success
Successfully charged 1.00 credits for instance 250 and created usage record 35

```

---

## Related Topics

- [5.6 AMCFG Parameters and Flags - page 437](#)
- Per Class `DISABLEAM` - page 72 attribute
- `ENFORCEACCOUNTACCESS` parameter



## 5.6 AMCFG Parameters and Flags

In this topic:

- [5.6.1 AMCFG Parameters - page 437](#)
- [5.6.2 AMCFG Flags - page 460](#)


### 5.6.1 AMCFG Parameters

Moab's accounting manager policies are defined using the `AMCFG[]` parameter. All AMCFG parameters must use the same accounting manager name between the square brackets (e.g., `AMCFG[mam]`).

The following AMCFG parameter values are supported:


<a href="#">ALWAYSCHARGERESERVATIONS</a>	<a href="#">ENDISBLOCKING</a>	<a href="#">QUERYURL</a>
<a href="#">BACKUPHOST</a>	<a href="#">ENDURL</a>	<a href="#">REFRESHPERIOD</a>
<a href="#">BLOCKINGACTIONS</a>	<a href="#">FALLBACKACCOUNT</a>	<a href="#">RESERVATIONCHARGEPOLICY</a>
<a href="#">CHARGEPOLICY</a>	<a href="#">FALLBACKQOS</a>	<a href="#">RESUMEFailureReaction</a>
<a href="#">CHARGERetryCount</a>	<a href="#">FLAGS</a>	<a href="#">RESUMEISBLOCKING</a>
<a href="#">CONTINUEFailureReaction</a>	<a href="#">FLUSHINTERVAL</a>	<a href="#">RESUMEURL</a>
<a href="#">CONTINUEISBLOCKING</a>	<a href="#">HOST</a>	<a href="#">RETRYFailedCharges</a>
<a href="#">CONTINUEURL</a>	<a href="#">INCLUDEFEATURES</a>	<a href="#">SERVER</a>
<a href="#">CREATECRED</a>	<a href="#">JOBCHARGEPOLICY</a>	<a href="#">STARTFailureReaction</a>
<a href="#">CREATEFailureReaction</a>	<a href="#">LIENGRANULARITY</a>	<a href="#">STARTISBLOCKING</a>
<a href="#">CREATEISBLOCKING</a>	<a href="#">MODE</a>	<a href="#">STARTURL</a>
<a href="#">CREATEURL</a>	<a href="#">NODECHARGEPOLICY</a>	<a href="#">TIMEOUT</a>
<a href="#">DELETEISBLOCKING</a>	<a href="#">PAUSEISBLOCKING</a>	<a href="#">TYPE</a>
<a href="#">DELETEURL</a>	<a href="#">PAUSEURL</a>	<a href="#">UPDATEURL</a>
<a href="#">DISABLEDACTIONS</a>	<a href="#">PORT</a>	<a href="#">VALIDATEJOBSubmission</a>


#### ALWAYSCHARGERESERVATIONS

<b>Description</b>	<div style="border: 1px solid red; padding: 5px; margin-bottom: 5px;">  This parameter is deprecated beginning with the Moab 9.1.0 release and may be removed in future releases.                 </div> <p>Use <a href="#">RESERVATIONCHARGEPOLICY</a> instead.</p> <ul style="list-style-type: none"> <li>• If you were using <code>AMCFG[] ALWAYSCHARGERESERVATIONS=True</code>,</li> </ul>
--------------------	---


ALWAYSCHARGERESERVATIONS	
	<p>comparable functionality can be obtained by using the new AMCFG [] RESERVATIONCHARGEPOLICY=All.</p> <ul style="list-style-type: none"> <li>If you were using AMCFG[] ALWAYSCHARGERESERVATIONS=False, comparable functionality can be obtained by using the new AMCFG [] RESERVATIONCHARGEPOLICY=Select.</li> </ul>


BACKUPHOST	
<b>Format</b>	STRING
<b>Default</b>	---
<b>Description</b>	The backup host name for the accounting manager server daemon.
<b>Example</b>	<pre>AMCFG [mam] BACKUPHOST=headnode2</pre> <p><i>Use the backup accounting manager server on headnode2 if the connection fails to the primary accounting manager server.</i></p>

BLOCKINGACTIONS	
<b>Description</b>	<div style="border: 1px solid red; padding: 5px; display: inline-block;">  This parameter is deprecated and may be removed in a future release.         </div> <p>Instead, specify the corresponding AMCFG[] <a href="#">CREATEISBLOCKING</a>, <a href="#">DELETEISBLOCKING</a>, <a href="#">ENDISBLOCKING</a>, <a href="#">PAUSEISBLOCKING</a>, <a href="#">RESUMEISBLOCKING</a>, and <a href="#">STARTISBLOCKING</a> parameters.</p>

CHARGERETRYCOUNT	
<b>Format</b>	<INTEGER> (non-negative)
<b>Default</b>	24
<b>Description</b>	<div style="border: 1px solid blue; padding: 5px; display: inline-block;">  Only applicable if <a href="#">RETRYFAILEDCHARGES</a> is enabled.         </div> <p>The maximum number of times that Moab will retry failed connection charges. Moab will continue to retry until the charge succeeds, the charge fails due to a</p>

CHARGERETRYCOUNT	
	non-connection failure, or until the <code>CHARGERETRYCOUNT</code> limit is reached. If set to zero, no retries will be performed, and all charge failures will be written to files in the <code>SPOOL/am/failed/</code> directory.
<b>Example</b>	<pre>AMCFG[mam] RETRYFAILEDCHARGES=TRUE CHARGERETRYCOUNT=12</pre> <p><i>Moab will retry connection-oriented charge failures up to 12 times.</i></p>

CHARGEPOLICY	
<b>Description</b>	<div style="border: 1px solid red; padding: 5px; margin-bottom: 10px;">  This parameter is deprecated beginning with the Moab 9.1.0 release and may be removed in a future release. </div> <p>The policy guiding whether to charge for all jobs or just successful jobs has been subsumed into the new <code>AMCFG[] JOBCHARGEPOLICY</code> parameter. Rather than specifying in Moab whether you want to charge for allocated processors, blocked processors, processor equivalents, or CPU time, Moab <i>now</i> sends all of these resource usage properties to the accounting manager so that they can be tracked and charged independently.</p>

CONTINUEFAILUREACTION	
<b>Format</b>	<GeneralFailureAction>[,<FundsFailureAction> [,<ConnectionFailureAction>]] where the action is one of <code>CANCEL</code> or <code>IGNORE</code>
<b>Default</b>	<code>IGNORE, IGNORE, IGNORE</code>
<b>Description</b>	<div style="border: 1px solid red; padding: 5px; margin-bottom: 10px;">  Setting this parameter to <code>IGNORE</code> can result in under-charging, and possibly over-charging of jobs. We recommend using <code>DEFER, CANCEL, DEFER</code> in most cases, unless it is acceptable to have occasional errors in the job charges. </div> <p>If periodic charging is enabled (via the <code>AMCFG[] FLUSHINTERVAL</code> parameter), this parameter specifies the action to be taken if a failure is detected when Moab performs its periodic accounting update (e.g., to determine whether the job should be continued):</p> <ul style="list-style-type: none"> <li>• Moab applies <code>&lt;ConnectionFailureAction&gt;</code> to a job if it is rejected due to a connection failure to MAM.</li> <li>• Moab applies <code>&lt;FundsFailureAction&gt;</code> to a job if it is rejected due to insufficient funds.</li> </ul>

CONTINUEFAILUREACTION	
	<ul style="list-style-type: none"> <li>Moab applies &lt;GeneralFailureAction&gt; to a job if the accounting manager rejects it for any other reason.</li> <li>If you do not specify a &lt;ConnectionFailureAction&gt;, or if you do not specify a &lt;FundsFailureAction&gt;, then Moab will apply the &lt;GeneralFailureAction&gt; for the unspecified case.</li> </ul> <p>If the action is set to CANCEL, Moab cancels the job; for IGNORE, Moab ignores the failure and continues running the job.</p>
<b>Example</b>	<pre>AMCFG [mam] CONTINUEFAILUREACTION=IGNORE, CANCEL, IGNORE</pre> <p><i>A job will be canceled if there are insufficient funds when Moab performs its periodic accounting update; but will be allowed to continue running if MAM is down, or for any other reason.</i></p>

CONTINUEISBLOCKING	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	TRUE
<b>Description</b>	If set to TRUE, the scheduler will block while authorizing the continuation of a job with the accounting manager. If set to FALSE, the accounting operation will be queued to the accounting thread pool and scheduling will continue; but application of the failure action will be delayed until a response is received.
<b>Example</b>	<pre>AMCFG [mam] CONTINUEISBLOCKING=FALSE</pre> <p><i>Specifies that Moab should use non-blocking calls with the accounting manager when checking to see if a job should be continued after a periodic accounting update.</i></p>

CONTINUEURL	
<b>Format</b>	exec://<fullPathToContinueScript> or null:
<b>Default</b>	exec://\$TOOLSDIR/mam/usage.reserve.mam.pl if TYPE=Native; otherwise it will make a direct call to MAM (mam:)

CONTINUEURL	
<b>Description</b>	<p>If periodic charging is enabled (via <code>AMCFG[] FLUSHINTERVAL</code>), when Moab performs a periodic accounting update for a job, this script is invoked to determine whether there are sufficient allocations for it to continue running for another period.</p> <p>For jobs, the <code>CONTINUEFAILUREACTION</code> attribute specifies the action that Moab should take if the authorization fails (such as for insufficient funds). If you use a job charge policy of Successful, Moab will not call the script because it does not yet know the completion status of the job.</p> <p>To disable a script from being run at this stage, use 'null:' as the parameter value.</p>
<b>Example</b>	<pre>AMCFG[mam] CONTINUEURL=exec://\$TOOLSDIR/mam/usage.continue.custom.pl</pre> <p><i>Moab calls the usage.continue.custom.pl script for authorization when checking to see if a job should be continued after a periodic accounting update.</i></p>

CREATECRED	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	<p>If set to <code>TRUE</code>, Moab will be enabled to query accounts, users, user membership in accounts, and users' default accounts from Moab Accounting Manager and define them in Moab. These credentials can be manually updated by running <code>mrmctl -R AM</code> or automatically updated by setting the <code>AMCFG[] REFRESHPERIOD</code> parameter.</p> <p>If you want Moab to enforce the imported account-user memberships, you will need to set the <code>ENFORCEACCOUNTACCESS</code> parameter to <code>TRUE</code>. See <a href="#">Moab Parameters</a> for more information on the <code>ENFORCEACCOUNTACCESS</code> parameter.</p>
<b>Example</b>	<pre>AMCFG[mam] CREATECRED=TRUE REFRESHPERIOD=30:00</pre> <p><i>Moab will automatically update account credential information from MAM every half hour.</i></p>

CREATEFAILUREACTION	
<b>Format</b>	<GeneralFailureAction>[,<FundsFailureAction>[,<ConnectionFailureAction>]]

CREATEFAILUREACTION	
	where the action is one of CANCEL, DEFER, HOLD, or IGNORE
<b>Default</b>	IGNORE, IGNORE, IGNORE
<b>Description</b>	<p>Before creating a job that should be tracked or charged within the accounting manager, Moab contacts the accounting manager for authorization:</p> <ul style="list-style-type: none"> <li>• If the job creation is rejected due to a connection failure with MAM, Moab applies the &lt;ConnectionFailureAction&gt; to the job.</li> <li>• If the job creation is rejected due to lack of funds, Moab applies the &lt;FundsFailureAction&gt; to the job.</li> <li>• For any other rejection reason, Moab applies the &lt;GeneralFailureAction&gt; to the job.</li> <li>• If you do not specify a &lt;ConnectionFailureAction&gt;, or if you do not specify a &lt;FundsFailureAction&gt;, then Moab will apply the &lt;GeneralFailureAction&gt; for the unspecified case.</li> </ul> <p>If the action is set to CANCEL, Moab cancels the job; DEFER, Moab defers the job; HOLD, Moab puts the job on hold; IGNORE, Moab ignores the failure and continues to start the job.</p> <div style="border: 1px solid #005596; border-radius: 10px; padding: 5px; margin-bottom: 5px;"> <p><b>i</b> In order for the CREATEFAILUREACTION policy to be applied, the AMCFG[] VALIDATEJOB SUBMISSION parameter must be set to true.</p> </div> <div style="border: 1px solid #005596; border-radius: 10px; padding: 5px;"> <p><b>i</b> If you have FALLBACKQOS or FALLBACKACCOUNT defined and a user requests an account with insufficient funds, the job will still be moved to the fallback credential, regardless of the action defined for CREATEFAILUREACTION.</p> </div>
<b>Example</b>	<pre>AMCFG [mam] CREATEFAILUREACTION=CANCEL, HOLD, IGNORE</pre> <div style="border: 1px dashed #ccc; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p><i>A job will be placed on hold when submitted if there are insufficient funds for it to start. However, it will be allowed to be submitted if there is a connection problem with MAM. The job will be canceled if there is any other failure (e.g., the user does not belong to the specified account).</i></p> </div>

CREATEISBLOCKING	
<b>Format</b>	<BOOLEAN>


CREATEISBLOCKING	
<b>Default</b>	TRUE
<b>Description</b>	If set to <code>TRUE</code> , the scheduler will block while authorizing the creation of a job with the accounting manager. If set to <code>FALSE</code> , the accounting operation will be queued to the accounting thread pool and scheduling will continue, but further consideration for the job will be delayed until a response is received.
<b>Example</b>	<pre>AMCFG[mam] CREATEISBLOCKING=FALSE</pre> <p><i>Specifies that Moab should use non-blocking calls with the accounting manager when creating jobs.</i></p>

CREATEURL	
<b>Format</b>	exec://<fullPathToCreateScript> or null:
<b>Default</b>	exec://\$TOOLSDIR/mam/usage.quote.mam.pl if <code>TYPE=Native</code> ; otherwise it will make a direct call to MAM (mam:)
<b>Description</b>	Moab runs this script at the time a job or reservation is being created. For jobs, the <code>CREATEFAILUREACTION</code> attribute specifies the action that should be taken if the authorization fails (such as for insufficient funds). To disable a script from being run at this stage, use 'null:' as the parameter value.
<b>Example</b>	<pre>AMCFG[mam] CREATEURL=exec://\$TOOLSDIR/mam/usage.create.custom.pl</pre> <p><i>Moab calls the <code>usage.create.custom.pl</code> script for authorization before starting a job or reservation.</i></p>

DELETEISBLOCKING	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	TRUE
<b>Description</b>	If set to <code>TRUE</code> , the scheduler will block while contacting the accounting manager to clean up after a failed job start. If set to <code>FALSE</code> , the accounting operation will be queued to the accounting thread pool and scheduling will

DELETEISBLOCKING	
	continue.
<b>Example</b>	<pre>AMCFG[mam] DELETEISBLOCKING=FALSE</pre> <p><i>Specifies that Moab should use non-blocking calls with the accounting manager when cleaning up after failed job starts.</i></p>

DELETEURL	
<b>Format</b>	exec://<fullPathToDeleteScript> or null:
<b>Default</b>	exec://\$TOOLSDIR/mam/lien.delete.mam.pl if TYPE=Native; otherwise it will make a direct call to MAM (mam:)
<b>Description</b>	<p>Moab runs this script to clean up after an interrupted job or reservation life-cycle. The default behavior is to remove outstanding liens.</p> <p>To disable a script from being run at this stage, use 'null:' as the parameter value.</p>
<b>Example</b>	<pre>AMCFG[mam] DELETEURL=exec://\$TOOLSDIR/mam/usage.delete.custom.pl</pre> <p><i>Moab calls the usage.delete.custom.pl script to clean up after an interrupted job or reservation.</i></p>

DISABLEDACTIONS	
<b>Description</b>	<div style="border: 1px solid red; padding: 5px; display: inline-block;">  This parameter is deprecated and may be removed in a future release.         </div> <p>Instead, specify an empty value or a protocol of 'null:' for the corresponding AMCFG[] <a href="#">CREATEURL</a>, <a href="#">DELETEURL</a>, <a href="#">ENDURL</a>, <a href="#">PAUSEURL</a>, <a href="#">RESUMEURL</a>, <a href="#">STARTURL</a>, and <a href="#">UPDATEURL</a> parameters.</p>

ENDISBLOCKING	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE



ENDISBLOCKING	
<b>Description</b>	If set to <code>TRUE</code> , the scheduler will block while registering the end of a job with the accounting manager. If set to <code>FALSE</code> , the accounting operation will be queued to the accounting thread pool and scheduling will continue.
<b>Example</b>	<pre>AMCFG [mam] ENDISBLOCKING=FALSE</pre> <p><i>Specifies that Moab should use non-blocking calls with the accounting manager when a job ends.</i></p>
ENDURL	
<b>Format</b>	exec://<fullPathToEndScript> or null:
<b>Default</b>	exec://\$TOOLSDIR/mam/usage.charge.mam.pl if <code>TYPE=Native</code> ; otherwise it will make a direct call to MAM (mam:)
<b>Description</b>	Moab runs this script after the end of a chargeable job or reservation in order to make a final charge or update the accounting record. The default behavior is to make a prorated charge for the job or reservation. To disable a script from being run at this stage, use 'null:' as the parameter value.
<b>Example</b>	<pre>AMCFG [mam] ENDURL=exec://\$TOOLSDIR/mam/usage.end.custom.pl</pre> <p><i>Moab calls the <code>usage.end.custom.pl</code> script to make the final charge for a job or reservation.</i></p>
FALLBACKACCOUNT	
<b>Format</b>	<STRING>
<b>Default</b>	---
<b>Description</b>	If specified, Moab verifies adequate allocations for all new jobs. If adequate allocations are not available in the job's primary account, Moab changes the job's credentials to use the fallback account. If not specified, Moab places a hold on jobs that do not have adequate allocations in their primary account.
<b>Example</b>	<pre>AMCFG [mam] FALLBACKACCOUNT=freecycle</pre>

**FALLBACKACCOUNT**

*Moab assigns the account `freecycle` to jobs that do not have adequate allocations in their primary account.*

**FALLBACKQOS**

**Format** <STRING>

**Default** ---

**Description** If specified, Moab verifies adequate allocations for all new jobs. If adequate allocations are not available in the job's primary QoS, Moab changes the job's credentials to use the fallback QoS. If not specified, Moab places a hold on jobs that do not have adequate allocations in their primary QoS.

**Example** `AMCFG[mam] FALLBACKQOS=freecycle`

*Moab assigns the QoS `freecycle` to jobs that do not have adequate allocations in their primary QoS.*

**FLAGS**


**Format** <STRING>

**Default** ---

**Description** AMCFG [flags](#) are used to enable special services.

**Example** `AMCFG[mam] FLAGS=LOCALCOST`

*Moab calculates the charge for the job locally and sends that as a charge to the accounting manager, which then charges that amount for the job.*

 The LOCALCOST flag is deprecated beginning in Moab 9.1.0 and may be removed in a future release.

FLUSHINTERVAL	
<b>Format</b>	[[[DD:]HH:]MM:]SS or INFINITY <div style="border: 1px solid black; padding: 5px; margin-top: 5px;"> <p><b>i</b> The former values of HOUR, DAY, WEEK, MONTH, or NONE are deprecated and may be removed in a future release.</p> </div>
<b>Default</b>	INFINITY
<b>Description</b>	Indicates the amount of time between accounting manager updates for long running reservations and jobs. If FLUSHINTERVAL is set to a positive time period, Moab will update the accounting manager (e.g., make an incremental charge) on the specified period relative to the start of the job or reservation. If FLUSHINTERVAL is set to INFINITY, the update will only occur at the end of the job or reservation.
<b>Example</b>	<pre>AMCFG[mam] FLUSHINTERVAL=1:00:00:00</pre> <div style="border: 1px dashed black; padding: 5px; margin-top: 5px;"> <p><i>Moab will make periodic accounting updates every 24 hours for long running jobs and reservations.</i></p> </div>

HOST	
<b>Format</b>	<STRING>
<b>Default</b>	localhost
<b>Description</b>	The host name for the accounting manager server daemon.
<b>Example</b>	<pre>AMCFG[mam] HOST=my-mam-server</pre> <div style="border: 1px dashed black; padding: 5px; margin-top: 5px;"> <p><i>Moab will communicate with the MAM server running on my-mam-server.</i></p> </div>

INCLUDEFEATURES	
<b>Format</b>	One of NodeCombination, NodeCount, TaskCombination, TaskCount, or NONE
<b>Default</b>	NONE


INCLUDEFEATURES	
<b>Description</b>	<p>Moab will pass allocated node features to the accounting manager under the Features property according to the specified policy:</p> <ul style="list-style-type: none"> <li>• If set to <code>NodeCombination</code>, Moab will pass a list of aggregate combinations of features for each node in the job.</li> <li>• If set to <code>NodeCount</code>, Moab will pass a list of features counted over each node in the job.</li> <li>• If set to <code>TaskCombination</code>, Moab will pass a list of aggregate combinations of features for each node in each task in the job.</li> <li>• If set to <code>TaskCount</code>, Moab will pass a list of features counted over each node in each task in the job.</li> </ul>
<b>Example</b>	<p>For example, assuming a job that runs with two tasks on each of eight nodes (e.g., <code>-l nodes=8:ppn=2</code>); four nodes; two having a single <code>bigmem</code> feature, two having a single <code>fastcpu</code> feature, two having both the <code>bigmem</code> and the <code>fastcpu</code> feature, and two having no features. The following parameter values would pass the accompanying Features value to the accounting manager.</p> <pre> AMCFG[mam] IncludeFeatures=NodeCombination            Features=bigmem:2,bigmem+fastcpu:2,fastcpu:2,None:2 AMCFG[mam] IncludeFeatures=NodeCount            Features=bigmem:4,fastcpu:4 AMCFG[mam] IncludeFeatures=TaskCombination            Features=bigmem:4,bigmem+fastcpu:4,fastcpu:4,None:4 AMCFG[mam] IncludeFeatures=TaskCount            Features=bigmem:8,fastcpu:8 </pre>

JOBCHARGEPOLICY	
<b>Format</b>	One of <code>All</code> , <code>None</code> , or <code>Successful</code>
<b>Default</b>	<code>All</code>
<b>Description</b>	<p>Specifies whether all, successful only, or no jobs should be charged for their resource usage:</p> <ul style="list-style-type: none"> <li>• If set to <code>All</code> (the default), all jobs are charged for their resource usage, independent of their completion status.</li> <li>• If set to <code>None</code>, no jobs will be tracked or charged for their resource usage.</li> <li>• If set to <code>Successful</code>, only jobs having a successful exit status (0) are charged for their resource usage.</li> </ul>

JOBCHARGEPOLICY	
	<p><b>i</b> If the LOCALCOST flag (AMCFG[] FLAGS=LOCALCOST) is set, Moab overrides the charge sent to the accounting manager and instead charges a value of the processors times duration (additionally factoring in NODECHARGE is defined). If LOCALCOST is <i>not</i> set, Moab allows the accounting manager to calculate the job charges.</p> <p><b>i</b> The LOCALCOST flag is deprecated beginning with the Moab 9.1.0 release and may be removed in a future release.</p>
<b>Example</b>	<pre>AMCFG[mam] JOBCHARGEPOLICY=Successful</pre> <p><i>Charge only for jobs having a successful exit status.</i></p>

LIENGRANULARITY	
<b>Format</b>	One of Partial or Combined
<b>Default</b>	Partial
<b>Description</b>	<p>When periodic charging is enabled via AMCFG[] FLUSHINTERVAL, lien granularity controls whether a combined lien is sought for the duration of the entire job (Combined) or whether partial liens are sought for the duration of each periodic charge interval (Partial):</p> <ul style="list-style-type: none"> <li>• When using a lien granularity of Partial, a job or reservation may get started if it has enough funds to run for the FLUSHINTERVAL, but it might trigger a CONTINUEFAILUREACTION if it runs out of funds before completion.</li> <li>• When using a lien granularity of Combined, the funds for the entire job or reservation must be available before it starts, but the funds will be protected by the lien and consumed on a periodic interval.</li> </ul>
<b>Example</b>	<pre>AMCFG[mam] LIENGRANULARITY=Combined</pre> <p><i>When using periodic charging, Moab will seek to obtain a lien for the entire duration of the job or reservation before starting it.</i></p>

MODE	
<b>Format</b>	One of <code>strict-allocation</code> , <code>fast-allocation</code> , <code>notional-charging</code> , or <code>usage-tracking</code>
<b>Default</b>	<code>strict-allocation</code>
<b>Description</b>	The accounting mode. The accounting mode modifies the way in which accounting-relevant job stages (e.g., create, start, end, etc.) are processed. See <a href="#">5.5.2 Accounting Mode - page 419</a> for details on the behavior of the accounting modes.
<b>Example</b>	<pre>AMCFG[mam] MODE=notional-charging</pre> <p><i>Configures Moab to use the notional-charging accounting mode when interacting with the accounting manager.</i></p>

NODECHARGEPOLICY	
<b>Format</b>	One of <code>AVG</code> , <code>MAX</code> , or <code>MIN</code>
<b>Default</b>	<code>MIN</code>
<b>Description</b>	<p>When charging for resource usage, the accounting manager will charge by node allocation according to the specified policy. For <code>AVG</code>, <code>MAX</code>, and <code>MIN</code>, the accounting manager will charge by the average, maximum, and minimum node charge rate of all allocated nodes.</p> <p>If you use this feature in conjunction with the <code>AMCFG[] LOCALCOST</code> flag, Moab will include the calculation of the node charge value sent to MAM. See <a href="#">LOCALCOST</a>.</p> <div style="border: 1px solid #0070C0; border-radius: 5px; padding: 5px; margin: 5px 0;"> <p> The <code>LOCALCOST</code> flag is deprecated beginning with the Moab 9.1.0 release and may be removed in a future release.</p> </div> <p>If you do not use this feature in conjunction with the <code>AMCFG[] LOCALCOST</code> flag, you must perform the following MAM commands to include node charges in charge calculations:</p> <ol style="list-style-type: none"> <li>Add <code>NodeCharge</code> as a usage record property. <pre>mam-shell Attribute Create Object=UsageRecord Name=NodeCharge DataType=Float Description="\Node Charge\""</pre> </li> <li>Add <code>NodeCharge</code> as a multiplier charge rate. <pre>mam-create-chargerate -n NodeCharge -z "*1" -d "Node Charge Multiplier"</pre> </li> </ol>

## NODECHARGEPOLICY

### Example

```
NODECFG[node01] CHARGERATE=1.5
NODECFG[node02] CHARGERATE=1.75
AMCFG[mam] NODECHARGEPOLICY=MAX
```

*Charge jobs by the maximum allocated node's charge rate.*

## PAUSEISBLOCKING

### Format

<BOOLEAN>

### Default

TRUE

### Description

If set to `TRUE`, the scheduler will block while registering the suspension of a job with the accounting manager. If set to `FALSE`, the accounting operation will be queued to the accounting thread pool and scheduling will continue.

### Example

```
AMCFG[mam] PAUSEISBLOCKING=FALSE
```

*Specifies that Moab should use non-blocking calls with the accounting manager when suspending jobs.*

## PAUSEURL

### Format

exec://<fullPathToPauseScript> or null:

### Default

exec://\$TOOLSDIR/mam/usage.charge.mam.pl if `TYPE=Native`; otherwise it will make a direct call to MAM (mam:)

### Description

Moab runs this script after preempting a job that might be resumed later. The default behavior is to make an incremental charge but not create a fresh lien. If you use a job charge policy of `Successful`, Moab will not call the script because it does not yet know the completion status of the job.

To disable a script from being run at this stage, use 'null:' as the parameter value.

### Example

```
AMCFG[mam] PAUSEURL=exec://$TOOLSDIR/mam/usage.pause.custom.pl
```

*Moab calls the `usage.pause.custom.pl` script after pausing a job.*

PORT	
<b>Format</b>	<INTEGER>
<b>Default</b>	7112
<b>Description</b>	The listening port for the accounting manager server daemon.
<b>Example</b>	<pre>AMCFG[mam] PORT=7731</pre> <p><i>Moab will communicate with the MAM server listening on port 7731.</i></p>


QUERYURL	
<b>Format</b>	exec://<fullPathToQueryScript> or null:
<b>Default</b>	exec://\$TOOLSDIR/mam/account.query.mam.pl if TYPE=Native; otherwise it will make a direct call to MAM (mam:)
<b>Description</b>	<p>Moab runs this script to customize and forward the Moab query to the accounting manager. The standard input to the script will be an XML Request in SSS format and is used directly between Moab and Moab Accounting Manager. Its primary purpose is to synchronize accounts and user information with the accounting manager if the <a href="#">CREATECRED</a> parameter is specified.</p> <p>To disable a script from being run at this stage, use 'null:' as the parameter value.</p>
<b>Example</b>	<pre>AMCFG[mam] QUERYURL=exec://\$TOOLSDIR/mam/cred.query.custom.pl</pre> <p><i>Moab calls the cred.query.custom.pl script in order to obtain account and user information from the accounting manager.</i></p>

REFRESHPERIOD	
<b>Format</b>	<pre>[[[DD:]HH:]MM:]SS or INFINITY</pre> <div style="border: 1px solid black; padding: 5px; margin-top: 5px;"> <p><b>i</b> The former values of MINUTE, HOUR, DAY or NONE are deprecated and may be removed in a future release.</p> </div>
<b>Default</b>	INFINITY



REFRESHPERIOD	
<b>Description</b>	<p>Indicates the period at which Moab will poll for updated information from Moab Accounting Manager (MAM):</p> <ul style="list-style-type: none"> <li>• If <code>AMCFG[] CREATECRED</code> is set to <code>TRUE</code>, Moab will update the accounting credentials from MAM on the specified period.</li> <li>• If <code>AMCFG[] MODE</code> is set to <code>fast-allocation</code>, Moab will update the account balance cache from MAM on the specified period.</li> </ul> <p>Moab will poll MAM for updated information when it first starts up unless <code>REFRESHPERIOD</code> is set to 0. If <code>REFRESHPERIOD</code> is set to a positive time period, Moab will refresh the accounting credentials on the specified period relative to the scheduler start time. If <code>REFRESHPERIOD</code> is set to <code>INFINITY</code>, Moab will only request updated information from MAM when first started. Use <code>mrmctl -R am</code> to force an immediate refresh.</p>
<b>Example</b>	<pre>AMCFG[mam] REFRESHPERIOD=2:00:00</pre> <p><i>Moab will request an update from MAM every two hours.</i></p>

RESERVATIONCHARGEPOLICY	
<b>Format</b>	One of All, None, or Select
<b>Default</b>	Select
<b>Description</b>	<p>If set to <code>All</code>, idle cycles in reservations will be charged to the accounting manager by default, even if the <code>ChargeAccount</code> and <code>ChargeUser</code> are not specified for the reservation. For reservations that you do not want to be charged with the accounting manager, specify the reservation <code>Charge</code> attribute with a value of <code>False</code>.</p> <p>If set to <code>None</code>, idle cycles in reservations will never be charged to the accounting manager, even if you have specified the <code>ChargeAccount</code>, <code>ChargeUser</code> or the <code>Charge</code> attribute with a value of <code>True</code>.</p> <p>If set to <code>Select</code> (the default), idle cycles in reservations will not be charged to the accounting manager unless you specify the reservation <code>ChargeAccount</code> or <code>ChargeUser</code> attributes or set the reservation <code>Charge</code> attribute with a value of <code>True</code>.</p>
<b>Example</b>	<pre>AMCFG[mam] RESERVATIONCHARGEPOLICY=All</pre> <p><i>By default, Moab will charge for idle cycles in reservations unless overridden with <code>Charge=False</code>.</i></p>

RESUMEFailureACTION	
<b>Format</b>	<GeneralFailureAction>[,<FundsFailureAction>[,<ConnectionFailureAction>]] where the action is one of CANCEL, DEFER, HOLD, IGNORE, or RETRY
<b>Default</b>	IGNORE, IGNORE, IGNORE
<b>Description</b>	<div style="border: 1px solid #c00000; padding: 5px; margin-bottom: 10px;">  Setting this parameter to IGNORE can result in under-charging, and possibly over-charging of jobs. We recommend using DEFER, CANCEL, DEFER in most cases, unless it is acceptable to have occasional errors in the job charges.                 </div> <p>This action is applied after a failure with the accounting manager when a job is being resumed (e.g., after being suspended):</p> <ul style="list-style-type: none"> <li>• Moab will apply &lt;ConnectionFailureAction&gt; to a job if there is a connection failure between Moab and the accounting manager.</li> <li>• Moab applies &lt;FundsFailureAction&gt; to the job if it is rejected due to insufficient funds.</li> <li>• Moab applies &lt;GeneralFailureAction&gt; to a job if the accounting manager rejects it for any other reason.</li> <li>• If you do not specify a &lt;ConnectionFailureAction&gt;, or if you do not specify a &lt;FundsFailureAction&gt;, then Moab will apply the &lt;GeneralFailureAction&gt; for the unspecified case.</li> </ul> <p>If the action is set to CANCEL, Moab cancels the job; DEFER, Moab defers the job; HOLD, Moab puts the job on hold; IGNORE, Moab ignores the failure and continues to resume the job; RETRY, Moab does not resume the job on this attempt but will continue to try to resume the job at the next opportunity.</p>
<b>Example</b>	<pre>AMCFG [mam] RESUMEFailureACTION=HOLD, HOLD, IGNORE</pre> <div style="border: 1px dashed #ccc; padding: 10px; margin-top: 10px;"> <p><i>A job will be resumed if Moab is unable to contact the accounting manager. Otherwise, the job will be placed on hold if there is any other failure with the accounting manager when Moab tries to resume it.</i></p> </div>

RESUMEISBLOCKING	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	TRUE


RESUMEISBLOCKING	
<b>Description</b>	If set to <code>TRUE</code> , the scheduler will block while authorizing the resumption of a job with the accounting manager. If set to <code>FALSE</code> , the accounting operation will be queued to the accounting thread pool and scheduling will continue, but resumption of the job will be delayed until a response is received.
<b>Example</b>	<pre>AMCFG[mam] RESUMEISBLOCKING=FALSE</pre> <p><i>Specifies that Moab should use non-blocking calls with the accounting manager when resuming jobs.</i></p>

RESUMEURL	
<b>Format</b>	exec://<fullPathToResumeScript> or null:
<b>Default</b>	exec://\$TOOLSDIR/mam/usage.reserve.mam.pl if <code>TYPE=Native</code> ; otherwise it will make a direct call to MAM (mam:)
<b>Description</b>	<p>Moab runs this script before resuming a suspended job to determine whether there it has authorization to resume (e.g., has sufficient funds).</p> <p>For jobs, the <code>RESUMEFAILUREACTION</code> attribute specifies the action that Moab should take if the authorization fails (such as for insufficient funds). If you use a job charge policy of <code>Successful</code>, Moab will not call the script because it does not yet know the completion status of the job.</p> <p>To disable a script from being run at this stage, use 'null:' as the parameter value.</p>
<b>Example</b>	<pre>AMCFG[mam] RESUMEURL=exec://\$TOOLSDIR/mam/usage.resume.custom.pl</pre> <p><i>Moab calls the <code>usage.resume.custom.pl</code> script for authorization before resuming a suspended job.</i></p>

RETRYFAILEDCHARGES	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	TRUE
<b>Description</b>	If set to <code>TRUE</code> , job charges will be retried if they have failed due to a connection failure. When a job charge or usage record update (such as might occur when a job is suspended, at the periodic charge interval, or when a job

RETRYFAILEDCHARGES	
	<p>completes) results in a connection failure between Moab and the accounting manager, then the charge request will be saved to a file in <code>SPOOLDIR/am/retrying/</code>. Once Moab detects that the connection with the accounting manager has been restored, the charge will be retried up to <code>CHARGERETRYCOUNT</code> times. Charges that fail due to reasons other than a connection failure, or connection failures that surpass the <code>CHARGERETRYCOUNT</code>, will be saved to files in <code>SPOOLDIR/am/failed/</code>.</p>
<b>Example</b>	<pre>AMCFG[mam] RETRYFAILEDCHARGES=TRUE</pre> <p><i>Moab will retry connection-oriented charge failures.</i></p>


SERVER	
<b>Format</b>	<URL>
<b>Default</b>	N/A
<b>Description</b>	The type and location of the accounting manager service.
<b>Example</b>	<pre>AMCFG[mam] SERVER=mam://tiny.supercluster.org:4368</pre>

STARTFAILUREACTION	
<b>Format</b>	<p>&lt;GeneralFailureAction&gt;[,&lt;FundsFailureAction&gt;[,&lt;ConnectionFailureAction&gt;]]            where the action is one of CANCEL, DEFER, HOLD, IGNORE, or RETRY</p>
<b>Default</b>	IGNORE, IGNORE, IGNORE
<b>Description</b>	<div style="border: 1px solid red; padding: 5px; margin-bottom: 10px;"> <p> Setting this parameter to <code>IGNORE</code> can result in under-charging, and possibly over-charging of jobs. We recommend using <code>DEFER</code>, <code>CANCEL</code>, <code>DEFER</code> in most cases, unless it is acceptable to have occasional errors in the job charges.</p> </div> <p>Moab applies the appropriate failure action if there is a failure when registering the job start with the accounting manager:</p> <ul style="list-style-type: none"> <li>• Moab applies &lt;ConnectionFailureAction&gt; to the job if there is a communication problem with the accounting manager.</li> <li>• Moab applies &lt;FundsFailureAction&gt; to the job if it is rejected due to</li> </ul>

STARTFAILUREACTION	
	<p>insufficient funds.</p> <ul style="list-style-type: none"> <li>• Moab applies &lt;GeneralFailureAction&gt; to a job if the accounting manager rejects it for any other reason.</li> <li>• If you do not specify a &lt;ConnectionFailureAction&gt;, or if you do not specify a &lt;FundsFailureAction&gt;, then Moab will apply the &lt;GeneralFailureAction&gt; for the unspecified case.</li> </ul> <p>If the action is set to CANCEL, Moab cancels the job; DEFER, Moab defers the job; HOLD, Moab puts the job on hold; IGNORE, Moab ignores the failure and continues to start the job; and RETRY, Moab does not start the job on this attempt but attempts to start the job at the next opportunity.</p>
<b>Example</b>	<pre>AMCFG[mam] STARTFAILUREACTION=CANCEL, HOLD, IGNORE</pre> <p><i>A job will be placed on hold if there are insufficient funds when it is time for it to start. It will be allowed to start if Moab is unable to reach the accounting manager. For all other failures with the accounting manager, the job will be canceled.</i></p>

STARTISBLOCKING	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	TRUE
<b>Description</b>	<p>If set to TRUE, the scheduler will block while authorizing the starting of a job with the accounting manager. If set to FALSE, the accounting operation will be queued to the accounting thread pool and scheduling will continue, but the start of the job will be delayed until a response is received.</p> <p>If using Moab in a Peer-to-Peer grid, do not set this parameter to FALSE. The Start action is not supported as a non-blocking action in Peer-to-Peer grids.</p>
<b>Example</b>	<pre>AMCFG[mam] STARTISBLOCKING=FALSE</pre> <p><i>Specifies that Moab should use non-blocking calls with the accounting manager when starting jobs.</i></p>

STARTURL	
<b>Format</b>	exec://<fullPathToStartScript> or null:
<b>Default</b>	exec://\$TOOLS DIR/mam/usage.reserve.mam.pl if <code>TYPE=Native</code> ; otherwise it will make a direct call to MAM (mam:)
<b>Description</b>	<p>Moab runs this script on a chargeable job or reservation to determine whether it should start.</p> <p>For jobs, the <code>STARTFAILUREACTION</code> attribute specifies the action that Moab should take if the authorization fails (such as for insufficient funds).</p> <p>To disable a script from being run at this stage, use 'null:' as the parameter value.</p>
<b>Example</b>	<pre>AMCFG[mam] STARTURL=exec://\$TOOLS DIR/mam/usage.start.custom.pl</pre> <p><i>Moab calls the <code>usage.start.custom.pl</code> script for authorization before starting a job or reservation.</i></p>

THREADPOOLSIZE	
<b>Description</b>	 This parameter is undocumented in 9.0.

TIMEOUT	
<b>Format</b>	[[[DD:]HH:]MM:]SS
<b>Default</b>	15
<b>Description</b>	The maximum delay allowed for communications with the accounting manager.
<b>Example</b>	<pre>AMCFG[mam] TIMEOUT=30</pre>

TYPE	
<b>Format</b>	One of MAM or Native
<b>Default</b>	MAM

TYPE	
<b>Description</b>	The accounting manager interface type.
<b>Example</b>	<pre>AMCFG [mam] TYPE=MAM</pre> <p><i>Configures Moab to interact with MAM using the direct SSS wire protocol.</i></p>


UPDATEURL	
<b>Format</b>	exec://<fullPathToUpdateScript> or null:
<b>Default</b>	exec://\$TOOLSDIR/mam/usage.charge.mam.pl if TYPE=Native; otherwise it will make a direct call to MAM (mam:)
<b>Description</b>	<p>If you have FLUSHINTERVAL set, Moab runs this script every flush interval for each chargeable job or reservation to charge for the previous interval. This call is usually followed by a call to the CONTINUEURL script, if defined, to check whether there are sufficient funds to run for the next interval. If you use a job charge policy of Successful, Moab will not call the script because it does not yet know the completion status of the job.</p> <p>To disable a script from being run at this stage, use 'null:' as the parameter value.</p>
<b>Example</b>	<pre>AMCFG [mam] UPDATEURL=exec://\$TOOLSDIR/mam/usage.update.custom.pl</pre> <p><i>Moab calls the usage.update.custom.pl script for authorization to continue a job or reservation.</i></p>

VALIDATEJOB SUBMISSION	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	<p>If set to TRUE, when a new job is submitted, Moab will execute the CREATEURL script (for TYPE=Native) or seek a job quote from Moab Accounting Manager (TYPE=MAM) before allowing the job to be submitted. Otherwise, the fund validation step is just utilized by reservations and fallback account checks. If the call fails (for example, if the user's account does not have sufficient funds or specifies an invalid account), Moab applies</p>

VALIDATEJOBSUBMISSION	
	the CREATEFAILUREACTION.
<b>Example</b>	<pre>AMCFG [mam] VALIDATEJOBSUBMISSION=True CREATEFAILUREACTION=Hold</pre> <p><i>Verify jobs have sufficient funds to run at the time they are submitted.</i></p>

## 5.6.2 AMCFG Flags

AMCFG flags can be used to enable special services and to disable default services. These services are enabled/disabled by setting the AMCFG FLAGS attribute (see [FLAGS](#)).

Flag Name	Description
<b>ACCOUNTFAILASFUNDS</b>	When this flag is set, logic failures within the accounting manager are treated as fund failures and are canceled. When ACCOUNTFAILASFUNDS is not set, accounting manager failures are treated as a server failure and the result is a job that requests an account to which the user does not have access.
<b>LOCALCOST</b>	<div style="border: 1px solid red; padding: 5px; background-color: #ffe6e6;">  This flag is deprecated beginning with the Moab 9.1.0 release and may be removed in a future release.                 </div>
<b>STRICTQUOTE</b>	Sends an estimated process count to the accounting manager when an initial quote is requested for a newly-submitted job.

### Related Topics

- [5.5 Accounting, Charging, and Allocation Management - page 418](#)



## Chapter 6: Controlling Resource Access - Reservations, Partitions, and QoS Facilities

In this section:

[6.1 Advance Reservations - page 462](#)

[6.2 Partitions - page 518](#)

[6.3 Quality of Service \(QoS\) Facilities - page 523](#)

## 6.1 Advance Reservations

An advance reservation is the mechanism by which Moab guarantees the availability of a set of resources at a particular time. Each reservation consists of three major components: (1) a set of resources, (2) a time frame, and (3) an access control list. It is a scheduler role to ensure that the access control list is not violated during the reservation's lifetime (that is, its time frame) on the resources listed. For example, a reservation may specify that node002 is reserved for user Tom on Friday. The scheduler is therefore constrained to make certain that only Tom's jobs can use node002 at any time on Friday. Advance reservation technology enables many features including [backfill](#), [deadline](#) based scheduling, [grid scheduling](#), and [QOS](#) support.

The [mrsvctl](#) command is used to [create](#), [modify](#), [query](#), and [release](#) reservations.

In this topic:

- [6.1.1 Reservation Overview - page 462](#)
- [6.1.2 Administrative Reservations - page 468](#)
- [6.1.3 Standing Reservations - page 469](#)
- [6.1.4 Reservation Policies - page 470](#)
- [6.1.5 Configuring and Managing Reservations - page 474](#)
- [6.1.6 Personal/User Reservations - Enabling Reservations for End Users - page 515](#)

### 6.1.1 Reservation Overview

Every reservation consists of three major components: (1) a set of resources, (2) a time frame, and (3) an access control list. Additionally, a reservation can also have a number of optional attributes controlling its behavior and interaction with other aspects of scheduling. Reservation attribute descriptions follow.

In this topic:

- [6.1.1.A Resources - page 463](#)
- [6.1.1.B Time Frame - page 464](#)
- [6.1.1.C Access Control List - page 464](#)

[6.1.1.D Job to Reservation Binding - page 464](#)

[6.1.1.E Reservation Specification - page 466](#)

[6.1.1.F Reservation Behavior - page 466](#)

[6.1.1.G Reservation Group - page 467](#)

[6.1.1.H Infinite Jobs and Reservations - page 467](#)

## 6.1.1.A Resources

Under Moab, the resources specified for a reservation are specified by way of a [task](#) description. Conceptually, a task can be thought of as an atomic, or indivisible, collection of resources. If reservation resources are unspecified, a task is a node by default. To define a task, specify resources. The resources can include processors, memory, swap, local disk, and so forth. For example, a single task may consist of one processor, 2 GB of memory, and 10 GB of local disk.

A reservation consists of one or more tasks. In attempting to locate the resources required for a particular reservation, Moab examines all feasible resources and locates the needed resources in groups specified by the task description. An example may help clarify this concept:

Reservation A requires four tasks. Each task is defined as 1 processor and 1 GB of memory.

Node X has 2 processors and 3 GB of memory available

Node Y has 2 processors and 1 GB of memory available

Node Z has 2 processors and 2 GB of memory available

When collecting the resources needed for the reservation, Moab examines each node in turn. Moab finds that Node X can support 2 of the 4 tasks needed by reserving 2 processors and 2 GB of memory, leaving 1 GB of memory unreserved. Analysis of Node Y shows that it can only support 1 task reserving 1 processor and 1 GB of memory, leaving 1 processor unreserved. Note that the unreserved memory on Node X cannot be combined with the unreserved processor on Node Y to satisfy the needs of another task because a task requires all resources to be located on the same node. Finally, analysis finds that node Z can support 2 tasks, fully reserving all of its resources.

Both reservations and jobs use the concept of a task description in specifying how resources should be allocated. It is important to note that although a task description is used to allocate resources to a reservation, this description does not in any way constrain the use of those resources by a job. In the above example, a job requesting resources simply sees 4 processors and 4 GB of memory available in reservation A. If the job has access to the reserved resources and the resources meet the other requirements of the job, the job could use these resources according to its own task description and needs.

Currently, the resources that can be associated with reservations include processors, memory, swap, local disk, initiator classes, and any number of arbitrary resources. Arbitrary resources can include peripherals such as tape drives, software licenses, or any other site specific resource.

### 6.1.1.B Time Frame

Associated with each reservation is a time frame. This specifies when the resources will be reserved or dedicated to jobs that meet the reservation's access control list (ACL). The time frame simply consists of a start time and an end time. When configuring a reservation, this information can be specified as a start time together with either an end time or a duration.


### 6.1.1.C Access Control List

A reservation's access control list specifies which jobs can use a reservation. Only jobs that meet one or more of a reservation's access criteria are allowed to use the reserved resources during the reservation time frame. Currently, the reservation access criteria include the following: users, groups, accounts, classes, QOS, job attributes, job duration, and job templates.

### 6.1.1.D Job to Reservation Binding

While a reservation's ACL will allow particular jobs to use reserved resources, it does not force any job to use these resources. With each job, Moab attempts to locate the best possible combination of available resources whether these are reserved or unreserved. For example, in the following figure, note that job X, which meets access criteria for both reservation A and B, allocates a portion of its resources from each reservation and the remainder from resources outside of both reservations.

Image 6-1: Job X uses resources from reservations A and B



Although by default, reservations make resources available to jobs that meet particular criteria, Moab can be configured to constrain jobs to only run within accessible reservations. This can be requested by the user on a job by job basis using a resource manager extension flag, or it can be enabled administratively via a QoS flag. For example, assume two reservations were created as follows:

```
> mrsvctl -c -a GROUP==staff -d 8:00:00 -h 'node[1-4]'
reservation staff.1 created
```

```
> mrsvctl -c -a USER==john -t 2
reservation john.2 created
```

If the user 'john,' who happened to also be a member of the group 'staff,' wanted to force a job to run within a particular reservation, 'john' could do so using the `FLAGS resource manager extension`. Specifically, in the case of a PBS job, the following submission would force the job to run within the 'staff.1' reservation:

```
> msub -l nodes=1,walltime=1:00:00,flags=ADVRES:staff.1 testjob.cmd
```

Note that for this to work, PBS needs to have resource manager extensions enabled as described in the PBS Resource Manager Extension Overview. (Torque has resource manager extensions enabled by default.) If the user wants the job to run on reserved resources but does not care which, the user could submit the job with the following:

```
> msub -l nodes=1,walltime=1:00:00,flags=ADVRES testjob.cmd
```

To enable job to reservation mapping via QoS, the QoS flag [USERRESERVED](#) should be set in a similar manner.

**i** Use the reservation [BYNAME](#) flag to require explicit binding for reservation access.

To lock jobs linked to a particular QoS into a reservation or reservation group, use the [REQRID](#) attribute.

### 6.1.1.E Reservation Specification

There are two main types of reservations that sites typically deal with. The first, administrative reservations, are typically one-time reservations created for special purposes and projects. These reservations are created using the [mrsvctl](#) or [setres](#) commands. These reservations provide an integrated mechanism to allow graceful management of unexpected system maintenance, temporary projects, and time critical demonstrations. This command allows an administrator to select a particular set of resources or just specify the quantity of resources needed. For example an administrator could use a regular expression to request a reservation be created on the nodes 'blue0[1-9]' or could simply request that the reservation locate the needed resources by specifying a quantity based request such as 'TASKS==20.'

The second type of reservation is called a [standing reservation](#). It is specified using the [SRCFG](#) parameter and is of use when there is a recurring need for a particular type of resource distribution. Standing reservations are a powerful, flexible, and efficient means for enabling persistent or periodic policies such as those often enabled using [classes](#) or [queues](#). For example, a site could use a standing reservation to reserve a subset of its compute resources for quick turnaround jobs during business hours on Monday thru Friday. The [Standing Reservation Overview](#) provides more information about configuring and using these reservations.

### 6.1.1.F Reservation Behavior

As previously mentioned, a given reservation can have one or more access criteria. A job can use the reserved resources if it meets at least one of these access criteria. It is possible to stack multiple reservations on the same node. In such a situation, a job can only use the given node if it has access to each active reservation on the node.

### 6.1.1.G Reservation Group

Reservations groups are ways of associating multiple reservations. This association is useful for variable namespace and [reservation requests](#). The reservations in a group inherit the variables from the reservation group head, but if the same variable is set locally on a reservation in the group, the local variable overrides the inherited variable. Variable inheritance is useful for [triggers](#) as it provides greater flexibility with automating certain tasks and system behaviors.

Jobs can be bound to a reservation group (instead of a single reservation) by using the resource manager extension [ADVRES](#).

### 6.1.1.H Infinite Jobs and Reservations

To allow infinite walltime jobs, you must have the following scheduler flag set:

```
SCHEDCFG[Moab] FLAGS=allowinfinitejobs
```

You can submit an infinite job by completing:

```
msub -l walltime=INFINITY
```

Or an infinite reservation by completing:

```
mrsvctl -c -d INFINITY
```

Infinite jobs can run in infinite reservations. Infinite walltime also works with job templates and advres.

Output XML for infinite jobs will print 'INFINITY' in the ReqAWDDuration, and XML for infinite rsvs will print 'INFINITY' in duration and endtime:

```
<Data>
  <rsv AUser="jgardner" AllocNodeCount="1" AllocNodeList="n5"
    AllocProcCount="4" AllocTaskCount="1" HostExp="n5"
    LastChargeTime="0" Name="jgardner.1" Partition="base"
    ReqNodeList="n5:1" Resources="PROCS=[ALL]" StatCBPS="0"
    StatCRPS="800" StatTBPS="0" StatTRPS="0" SubType="Other"
    Type="User" cost="0.000000" ctime="1302127058"
    duration="INFINITY" endtime="INFINITY" starttime="1302127058">
    <ACL aff="neutral" cmp="%" name="jgardner.1" type="RSV"></ACL>
    <ACL cmp="%" name="jgardner" type="USER"></ACL>
    <ACL cmp="%" name="company" type="GROUP"></ACL>
    <ACL aff="neutral" cmp="%" name="jgardner.1" type="RSV"></ACL>
    <History>
      <event state="PROCS=4" time="1302127058"></event>
    </History>
  </rsv>
</Data>
```

## Related Topics

- [Reservation Allocation Policies](#)

## 6.1.2 Administrative Reservations

Administrative reservations behave much like standing reservations but are generally created to address non-periodic, one-time issues. All administrative reservations are created using the `mrsvctl -c` (or `setres`) command and are persistent until they expire or are removed using the `mrsvctl -r` (or `releaseres`) command.

### 6.1.2.A Annotating Administrative Reservations

Reservations can be labeled and annotated using comments allowing other administrators, local users, portals and other services to obtain more detailed information regarding the reservations. Naming and annotations are configured using the `-n` and `-D` options of the `mrsvctl` command respectively, as in the following example:

```
> mrsvctl -c -D 'testing infiniband performance' -n nettest -h 'r:agt[15-245]'
```

### 6.1.2.B Using Reservation Profiles

You can set up reservation profiles to avoid manually and repetitively inputting standard reservation attributes. Profiles can specify reservation names, descriptions, ACLs, durations, hostlists, triggers, flags, and other aspects that are commonly used. With a reservation profile defined, a new administrative reservation can be created that uses this profile by specifying the `-P` flag as in the following example:

```
RSVPROFILE [mtn1] TRIGGER=Atype=exec,Action="/tmp/trigger1.sh",EType=start
RSVPROFILE [mtn1] USERLIST=steve,marym
RSVPROFILE [mtn1] HOSTEXP="r:50-250"
```

```
> mrsvctl -c -P mtn1 -s 12:00:00_10/03 -d 2:00:00
```

*Example 6-1: Non-Blocking System Reservations with Scheduler Pause:*

```
RSVPROFILE [pause] TRIGGER=atype=exec,etype=start,action="/opt/moab/bin/mschedctl -p"
RSVPROFILE [pause] TRIGGER=atype=exec,etype=cancel,action="/opt/moab/bin/mschedctl -r"
RSVPROFILE [pause] TRIGGER=atype=exec,etype=end,action="/opt/moab/bin/mschedctl -r"
```

```
> mrsvctl -c -P pause -s 12:00:00_10/03 -d 2:00:00
```



### 6.1.2.C Optimizing Maintenance Reservations

Any reservation causes some negative impact on cluster performance as it further limits the scheduler's ability to optimize scheduling decisions. You can mitigate this impact by using flexible ACLs and triggers.

In particular, a maintenance reservation can be configured to reduce its effective reservation shadow by allowing overlap with checkpointable/preemptible jobs until the time the reservation becomes active. This can be done using a series of triggers that perform the following actions:

- Modify the reservation to disable preemption access.
- Preempt jobs that may overlap the reservation.
- Cancel any jobs that failed to properly checkpoint and exit.

The following example highlights one possible configuration:

```
RSVPROFILE [adm1] JOBATTRLIST=PREEMPTEE
RSVPROFILE [adm1] DESCRIPTION="regular system maintenance"
RSVPROFILE [adm1] TRIGGER=EType=start,Offset=-300,AType=internal,Action="rsv:-:modify:acl:jattr==PREEMPTEE"
RSVPROFILE [adm1] TRIGGER=EType=start,Offset=-240,AType=jobpreempt,Action="checkpoint"
RSVPROFILE [adm1] TRIGGER=EType=start,Offset=-60,AType=jobpreempt,Action="cancel"
```

```
> mrsvctl -c -P adm1 -s 12:00:00_10/03 -d 8:00:00 -h ALL
```

This reservation reserves all nodes in the cluster for a period of eight hours. Five minutes before the reservation starts, the reservation is modified to remove access to new preemptible jobs. Four minutes before the reservation starts, preemptible jobs that overlap the reservation are checkpointed. One minute before the reservation, all remaining jobs that overlap the reservation are canceled.

---

#### Related Topics

- [Backfill](#)
- [Preemption](#)
- [mrsvctl](#) command

### 6.1.3 Standing Reservations

Standing reservations build upon the capabilities of advance reservations to enable a site to enforce advanced usage policies in an efficient manner. Standing reservations provide a

superset of the capabilities typically found in a batch queuing system's class or queue architecture. For example, queues can be used to allow only particular types of jobs access to certain compute resources. Also, some batch systems allow these queues to be configured so that they only allow this access during certain times of the day or week. Standing reservations allow these same capabilities but with greater flexibility and efficiency than is typically found in a normal queue management system.

Standing reservations provide a mechanism by which a site can dedicate a particular block of resources for a special use on a regular daily or weekly basis. For example, node X could be dedicated to running jobs only from users in the accounting group every Friday from 4 to 10 p.m. See the [Reservation Overview](#) for more information about the use of reservations. The [Managing Reservations](#) section provides a detailed explanation of the concepts and steps involved in the creation and configuration of standing reservations.

A standing reservation is a powerful means of doing the following:

- Controlling local credential based access to resources.
- Controlling external peer and grid based access to resources.
- Controlling job responsiveness and turnaround.

---

#### Related Topics

- [SRCFG](#)
- [Moab Workload Manager for Grids](#)
- [mdiag -s](#) (diagnose standing reservations)

### 6.1.4 Reservation Policies

In this topic:

- [6.1.4.A Controlling Priority Reservation Creation - page 471](#)
- [6.1.4.B Priority Reservation Creation Policy - page 471](#)
- [6.1.4.C Managing Resource Failures - page 473](#)
- [6.1.4.D Resource Allocation Policy - page 474](#)
- [Accounting for Reserved Resources](#)

### 6.1.4.A Controlling Priority Reservation Creation

In addition to standing and administrative reservations, Moab can also create priority reservations. These reservations are used to allow the benefits of out-of-order execution (such as is available with [backfill](#)) without the side effect of job starvation. Starvation can occur in any system where the potential exists for a job to be overlooked by the scheduler for an indefinite period. In the case of backfill, small jobs may continue to run on available resources as they become available while a large job sits in the queue, never able to find enough nodes available simultaneously on which to run.

To avoid such situations, priority reservations are created for high priority jobs that cannot run immediately. When making these reservations, the scheduler determines the earliest time the job could start and then reserves these resources for use by this job at that future time.

### 6.1.4.B Priority Reservation Creation Policy

Organizations have the ability to control how priority reservations are created and maintained. It is possible that one job can be at the top of the priority queue for a time and then get bypassed by another job submitted later. The parameter [RESERVATIONPOLICY](#) allows a site to determine how existing reservations should be handled when new reservations are made.

Value	Description
<b>HIGHEST</b>	<p>All jobs that have ever received a priority reservation up to the <code>RESERVATIONDEPTH</code> number will maintain that reservation until they run, even if other jobs later bypass them in priority value. For example, if there are four jobs with priorities of 8, 10,12, and 20.</p> <div style="border: 1px dashed black; padding: 5px; margin: 10px 0;"> <pre>RESERVATIONPOLICY HIGHEST RESERVATIONDEPTH 3</pre> </div> <p>Only jobs 20, 12, and 10 get priority reservations. Later, if a job with priority higher than 20 is submitted into the queue, it will also get a priority reservation along with the jobs listed previously. If four jobs higher than 20 were to be submitted into the queue, only three would get priority reservations, in accordance with the condition set in the <code>RESERVATIONDEPTH</code> policy.</p> <p>With <code>HIGHEST</code>, Moab may appear to exceed the <code>RESERVATIONDEPTH</code> if it has already scheduled the maximum number of priority reservations and then users submit jobs with higher priority than those already given a priority reservation. Moab keeps all of the previously-created priority reservations and creates new ones for jobs with higher priority (again up to the quantity specified with <code>RESERVATIONDEPTH</code>). This means that, if your</p>

Value	Description
	RESERVATIONDEPTH is set to 3, Moab can potentially schedule up to 3 new priority reservations each scheduling iteration, as long as new higher-priority jobs are continually submitted. This behavior ensures that the highest-priority jobs receive attention while the former highest-priority jobs do not lose their priority reservation.
<b>CURRENTHIGHEST</b>	Only the current top <RESERVATIONDEPTH> priority jobs receive reservations. Under this policy, all job reservations are destroyed each iteration when the queue is re-prioritized. The top jobs in the queue are then given new reservations.
<b>NEVER</b>	No priority reservations are made.

## Priority Reservation Depth

By default, only the highest priority job receives a priority reservation. However, this behavior is configurable via the [RESERVATIONDEPTH](#) policy. Moab's default behavior of only reserving the highest priority job allows backfill to be used in a form known as liberal backfill. Liberal backfill tends to maximize system utilization and minimize overall average job turnaround time. However, it does lead to the potential of some lower priority jobs being indirectly delayed and might lead to greater variance in job turnaround time. The [RESERVATIONDEPTH](#) parameter can be set to a very large value, essentially enabling what is called conservative backfill where every job that cannot run is given a reservation. Most sites prefer the liberal backfill approach associated with the default [RESERVATIONDEPTH](#) of 1 or else select a slightly higher value. It is important to note that to prevent starvation in conjunction with reservations, monotonically increasing priority factors such as queue time or job XFactor should be enabled. See [4.1.1 Priority Overview - page 335](#) for more information on priority factors.

Another important consequence of backfill and reservation depth is how they affect job priority. In Moab, all jobs are prioritized. Backfill allows jobs to be run out of order and therefore, to some extent, job priority to be ignored. This effect, known as priority dilution, can cause many site policies implemented via Moab prioritization policies to be ineffective. Setting the [RESERVATIONDEPTH](#) parameter to a higher value gives job priority more teeth at the cost of slightly lower system utilization. This lower utilization results from the constraints of these additional reservations, decreasing the scheduler's freedom and its ability to find additional optimizing schedules. Anecdotal evidence indicates that these utilization losses are fairly minor, rarely exceeding 8%.

It is difficult to know the right setting for the [RESERVATIONDEPTH](#) parameter. Surveys indicate that the vast majority of sites use the default value of 1. Sites that do modify this value typically set it somewhere in the range of 2 to 10. The following guidelines may be useful in determining if and how to adjust this parameter:

## Reasons to Increase RESERVATIONDEPTH

- The estimated job start time information provided by the `showstart` command is heavily used and the accuracy needs to be increased.
- Priority dilution prevents certain key mission objectives from being fulfilled.
- Users are more interested in knowing when their job will run than in having it run sooner.

## Reasons to Decrease RESERVATIONDEPTH

- Scheduling efficiency and job throughput need to be increased.

## Assigning Per-QoS Reservation Creation Rules

QoS based reservation depths can be enabled via the `RESERVATIONQOSLIST` parameter. This parameter allows varying reservation depths to be associated with different sets of job QoSs. For example, the following configuration creates two reservation depth groupings:

```
RESERVATIONDEPTH[0]      8
RESERVATIONQOSLIST[0]   highprio, interactive, debug
RESERVATIONDEPTH[1]      2
RESERVATIONQOSLIST[1]   batch
```

*This example causes that the top 8 jobs belonging to the aggregate group of highprio, interactive, and debug QoS jobs will receive priority reservations. Additionally, the top two batch QoS jobs will also receive priority reservations. Use of this feature allows sites to maintain high throughput for important jobs by guaranteeing that a significant proportion of these jobs progress toward starting through use of the priority reservation.*

By default, the following parameters are set inside Moab:

```
RESERVATIONDEPTH[DEFAULT] 1
RESERVATIONQOSLIST[DEFAULT] ALL
```

*This allows one job with the highest priority to get a reservation. These values can be overwritten by modifying the DEFAULT policy.*

### 6.1.4.C Managing Resource Failures

Moab allows organizations to control how to best respond to a number of real-world issues. Occasionally when a reservation becomes active and a job attempts to start, various resource manager race conditions or corrupt state situations will prevent the job from starting. By default, Moab assumes the resource manager is corrupt, releases the reservation, and attempts to re-create the reservation after a short timeout. However, in the interval between the reservation release and the re-creation timeout, other priority reservations may allocate the newly available resources, reserving them before the original reservation gets an opportunity to reallocate them. Therefore, when the original job reservation is re-established, its original resource may be unavailable and the resulting new reservation might be delayed several hours from the earlier start time. The parameter

[RESERVATIONRETRYTIME](#) allows a site that is experiencing frequent resource manager race conditions and/or corruption situations to tell Moab to hold on to the reserved resource for a period of time in an attempt to allow the resource manager to correct its state.

#### 6.1.4.D Resource Allocation Policy

By default, when a standing or administrative reservation is created, Moab allocates nodes in accordance with the specified taskcount, node expression, node constraints, and the [MINRESOURCE](#) node allocation policy.

#### 6.1.4.E Accounting for Reserved Resources

If an accounting manager is configured within Moab, resources consumed by jobs are tracked and charged by default. However, resources dedicated to a reservation are not charged by default although they are recorded within the reservation event record. In particular, total processor seconds reserved by the reservation and total processor seconds blocked by jobs are among the statistics recorded. While some of this information is available in real-time using the [mdiag -r](#) command (see the 'PH Allocated to Jobs' field), it is not written to the event log until reservation completion.

It is possible to track or charge for the total and unused cycles in a reservation with the accounting manager. See [5.5.5 Charging for Reservations - page 424](#) for details on configuring Moab to use the accounting manager to track or charge for unused processors in a reservation.

---

#### Related Topics

- [Reservation Overview](#)
- [Backfill](#)

### 6.1.5 Configuring and Managing Reservations

In this topic:

[6.1.5.A Reservation Attributes - page 475](#)

[6.1.5.B Configuring Standing Reservations - page 483](#)

## 6.1.5.C Managing Administrative Reservations - page 514

## 6.1.5.A Reservation Attributes

All reservations possess a time frame of activity, an access control list (ACL), and a list of resources to be reserved. Additionally, reservations can also possess a number of extension attributes including epilog/prolog specification, reservation ownership and accountability attributes, and special flags that modify the reservation's behavior.

**Start/End Time**

All reservations possess a start and an end time that define the reservation's active time. During this active time, the resources within the reservation can only be used as specified by the reservation access control list (ACL). This active time can be specified as either a start/end pair or a start/duration pair. Reservations exist and are visible from the time they are created until the active time ends at which point they are automatically removed.

**Access Control List (ACL)**

For a reservation to be useful, it must be able to limit who or what can access the resources it has reserved.

**i** By default, a reservation can allocate resources that possess credentials that meet the submitter's ACL. In other words, a user's reservation won't necessarily allocate only free and idle nodes. If a reservation exists that coincides with the submitter's ACL, the nodes under that reservation are also considered for allocation. This is referred to as ACL overlap. To make new reservations allocate *only* free and idle nodes, you must use the `NOACLOVERLAP` flag.

This is handled by way of an ACL. With reservations, ACLs can be based on credentials, resources requested, or performance metrics. In particular, with a standing reservation, the attributes `USERLIST`, `GROUPLIST`, `ACCOUNTLIST`, `CLASSLIST`, `QOSLIST`, `JOBATTRLIST`, `PROCLIMIT`, `MAXTIME`, or `TIMELIMIT` can be specified. See [Affinity and Modifiers](#).

**i** Reservation access can be adjusted based on a job's requested node features by mapping node feature requests to job attributes as in the following example:

```
NODECFG[DEFAULT]   FEATURES=ia64
NODETOJOBATTRMAP  ia64,ia32
SRCFG[pgs]        JOBATTRLIST=ia32
```

```
> mrsvctl -c -a jattr=gpbs\! -h "r:13-500"
```

## Selecting Resources

When specifying which resources to reserve, the administrator has a number of options. These options allow control over how many resources are reserved and where they are reserved. The following reservation attributes allow the administrator to define resources.

### Task Description

Moab uses the task concept extensively for its job and reservation management. A task is simply an atomic collection of resources, such as processors, memory, or local disk, which must be found on the same node. For example, if a task requires 4 processors and 2 GB of memory, the scheduler must find all processors AND memory on the same node; it cannot allocate 3 processors and 1 GB on one node and 1 processor and 1 GB of memory on another node to satisfy this task. Tasks constrain how the scheduler must collect resources for use in a standing reservation; however, they do not constrain the way in which the scheduler makes these cumulative resources available to jobs. A job can use the resources covered by an accessible reservation in whatever way it needs. If reservation X allocates 6 tasks with 2 processors and 512 MB of memory each, it could support job Y that requires 10 tasks of 1 processor and 128 MB of memory or job Z that requires 2 tasks of 4 processors and 1 GB of memory each. The task constraints used to acquire a reservation's resources are transparent to a job requesting use of these resources.

```
SRCFG[test] RESOURCES=PROCS:2,MEM:1024
```

### Taskcount

Using the task description, the taskcount attribute defines how many tasks must be allocated to satisfy the reservation request. To create a reservation, a taskcount and/or a hostlist must be specified.

```
SRCFG[test] TASKCOUNT=256
```

### Hostlist



A hostlist constrains the set of resources available to a reservation. If no taskcount is specified, the reservation attempts to reserve one task on each of the listed resources. If a taskcount is specified that requests fewer resources than listed in the hostlist, the scheduler reserves only the number of tasks from the hostlist specified by the taskcount attribute. If a taskcount is specified that requests more resources than listed in the hostlist, the scheduler reserves the hostlist nodes first and then seeks additional resources outside of this list.

When specifying resources for a hostlist, you can specify *exact set*, *superset*, or *subset* of nodes on which the job must run. Use the caret (^) or asterisk (\*) characters to specify a hostlist as *superset* or *subset* respectively.

- An exact set is defined without a caret or asterisk. An exact set means *all* the hosts in the specified hostlist must be selected for the job.
- A subset means the specified hostlist is used first to select hosts for the job. If the job requires more hosts than are in the subset hostlist, they will be obtained from elsewhere if possible. If the job does not require all of the nodes in the subset hostlist, it will use only the ones it needs.
- A superset means the hostlist is the *only* source of hosts that should be considered for running the job. If the job can't find the necessary resources in the superset hostlist it should *not* run. No other hosts should be considered in allocating the job.

```
SRCFG[test] HOSTLIST=node01,node1[3-5]
```

#### Example 6-2: Subset

```
SRCFG[one] HOSTLIST=node1,node5* TASKCOUNT=5 PERIOD=DAY USERLIST=user1
```

#### Example 6-3: Superset

```
SRCFG[two] HOSTLIST=node1,node2,node3,node4,node5^ TASKCOUNT=3 PERIOD=DAY  
USERLIST=user1
```

## Node Features

Node features can be specified to constrain which resources are considered.

#### Example 6-4:

```
SRCFG[test] NODEFEATURES=fastos
```

## Partition

A partition can be specified to constrain which resources are considered.

```
SRCFG[test] PARTITION=core3
```

## Flags

Reservation flags allow specification of special reservation attributes or behaviors. Supported flags are listed in the following table:

## Flags

Flag Name	Description
<b>ACLOVERLAP</b>	<b>Deprecated (this is now a default flag).</b> In addition to free or idle nodes, a reservation can also reserve resources with job reservations possessing credentials that meet the reservation's ACL. To change this behavior, set the <a href="#">NOACLOVERLAP</a> flag.
<b>ADVRESJOBDESTROY</b>	All jobs that have an <a href="#">ADVRES</a> matching this reservation are canceled when the reservation is destroyed.
<b>ALLOWGRID</b>	By default, jobs migrated from one Moab to another Moab in a grid are not allowed within local reservations. This flag allows migrated jobs to access local reservations when they match the ACL.
<b>ALLOWJOBOverlap</b>	A job is allowed to start in a reservation that may end before the job completes. When the reservation ends before the job completes, the job will not be canceled but will continue to run.
<b>BESTEFFORT</b>	Reservation is placed, even if only some of the specified resources are available.
<b>BYNAME</b>	Reservation only allows access to jobs that meet reservation ACLs and explicitly request the resources of this reservation using the job <a href="#">ADVRES</a> flag. See <a href="#">Job to Reservation Binding</a> .
<b>DEDICATEDRESOURCE (a.k.a. EXCLUSIVE)</b>	Reservation placed only on resources that are not reserved by any other reservation including job, system, and user reservation. There are two exceptions to this: <ol style="list-style-type: none"> <li>1. Reserved resources could be allocated when DEDICATEDRESOURCE is combined with <a href="#">IGNJOBRSV*</a></li> <li>2. Reserved resources could be allocated when a reservation matches the submitter's ACL. In this case, to make DEDICATEDRESOURCE <i>truly</i> exclusive, use the <a href="#">NOACLOVERLAP</a> flag.</li> </ol>

Flag Name	Description
	<p><b>i</b> The order that SRCFG reservations are listed in the configuration is important when using <code>DEDICATEDRESOURCE</code>, because reservations made afterwards can steal resources later. During configuration, list <code>DEDICATEDRESOURCE</code> reservations last to guarantee exclusiveness.</p>
<b>ENFORCENODESET</b>	<p>Moab will ensure global NODESET rules are followed when initially placing the reservation.</p>
<b>EVACVMS</b>	<p>Reservation will automatically evacuate virtual machines from the reservation nodelist.</p> <p><b>i</b> The same action can be accomplished by using reservation profiles. For more information, see <a href="#">Optimizing Maintenance Reservations</a>.</p>
<b>IGNIDLEJOBS*</b>	<p>Reservation can be placed on top of idle job reservations.</p> <p><b>i</b> This flag is meant to be used in conjunction with <code>DEDICATEDRESOURCE</code>.</p>
<b>IGNJOBRSV*</b>	<p>Ignores existing job reservations, allowing the reservation to be forced onto available resources even if it conflicts with existing job reservations. User and system reservation conflicts are still valid. It functions the same as <a href="#">IGNIDLEJOBS</a> plus allows a reservation to be placed on top of an existing running job's reservation.</p> <p><b>i</b> This flag is meant to be used in conjunction with <code>DEDICATEDRESOURCE</code>.</p>
<b>IGNRSV*</b>	<p>Request ignores existing resource reservations allowing the reservation to be forced onto available resources even if this conflicts with other reservations. It functions the same as <a href="#">IGNJOBRSV</a> plus allows the reservation to be placed on top of</p>

Flag Name	Description
	<p>the system reservations.</p> <p><b>i</b> This flag is meant to be used in conjunction with <code>DEDICATEDRESOURCE</code>.</p>
<b>IGNSTATE*</b>	<p>Reservation ignores node state when assigning nodes. It functions the same as <code>IGNRSV</code> plus allows the reservation to be placed on nodes that are not currently available. Also ignores resource availability on nodes.</p> <p><b>i</b> <code>IGNSTATE</code> is specified by default when using a <code>HOSTLIST</code> to define nodes. However, if using a <code>HOSTLIST</code> and a <code>TASKCOUNT</code>, you need to specify <code>IGNSTATE</code> if you want Moab to ignore the node state when assigning nodes to the reservation.</p>
<b>NOACLOVERLAP</b>	<p>All resources must be free or idle, with no existing reservations. Moab will not allocate in-use resources even if they match the reservation's ACL.</p> <pre data-bbox="776 1062 1414 1108">mrsvctl -c -t 12 -E -F noacloverlap -a user==john</pre> <p><i>Moab looks for resources that are exclusive (free). Without the flag, Moab would look for resources that are exclusive or that are already running john's jobs.</i></p> <p><b>i</b> This flag is meant to be used in conjunction with <code>DEDICATEDRESOURCE</code>.</p>
<b>OWNEREXCLUSIVEBF</b>	<p>When the owner of the reservation has an idle job in the queue only owner jobs will be allowed to backfill into the reservation. This blocks non-owner jobs from backfilling into the reservation.</p> <p><b>i</b> <code>ENABLEPROFILING</code> must be set for the owner credential.</p>
<b>OWNERPREEMPT</b>	<p>Jobs by the reservation <code>owner</code> are allowed to</p>

Flag Name	Description
	preempt non-owner jobs using reservation resources.
<b>OWNERPREEMPTIGNOREMINTIME</b>	<p>Allows the OWNERPREEMPT flag to "trump" the PREEMPTMINTIME setting for jobs already running on a reservation when the owner of the reservation submits a job. For example: without the OWNERPREEMPTIGNOREMINTIME flag set, a job submitted by the owner of a reservation will not preempt non-owner jobs already running on the reservation until the PREEMPTMINTIME setting (if set) for those jobs is passed.</p> <p>With the OWNERPREEMPTIGNOREMINTIME flag set, a job submitted by the owner of a reservation immediately preempts non-owner jobs already running on the reservation, regardless of whether PREEMPTMINTIME is set for the non-owner jobs.</p>
<b>OWNERPREEMPTQT</b>	<p>Specifies how much time a job from OWNER must wait in the queue before preempting jobs within the standing reservation.</p> <div style="border: 1px dashed black; padding: 5px; margin: 10px 0;"> <p>SRCFG[test] OWNERPREEMPTQT=2:00:00</p> </div> <div style="border: 1px dashed black; padding: 5px; margin: 10px 0;"> <p><i>OWNER jobs must wait 2 hours in the queue before preempting.</i></p> </div>
<b>REQFULL</b>	Reservation is only created when all resources can be allocated.
<b>SINGLEUSE</b>	Reservation is automatically removed after completion of the first job to use the reserved resources.
<b>SPACEFLEX</b>	<p><b>Deprecated (this is now a default flag).</b> Reservation is allowed to adjust resources allocated over time in an attempt to optimize resource utilization.</p>

**i** \* `IGNIDLEJOBS`, `IGNJOBRSV`, `IGNRSV`, and `IGNSTATE` flags are built on one another and form a hierarchy. `IGNJOBRSV` performs the function of `IGNIDLEJOBS` plus its own functions. `IGNRSV` performs the function of `IGNJOBRSV` and `IGNIDLEJOBS` plus its own functions. `IGNSTATE` performs the function of `IGNRSV`, `IGNJOBRSV`, and `IGNIDLEJOBS` plus its own functions. While you can use combinations of these flags, it is not necessary. If you set one flag, you do not need to set other flags that fall beneath it in the hierarchy.

Most flags can be associated with a reservation via the `mrsvctl -c -F` command or the `SRCFG` parameter.

### 6.1.5.B Configuring Standing Reservations

Standing reservations allow resources to be dedicated for particular uses. This dedication can be configured to be permanent or periodic, recurring at a regular time of day and/or time of week. There is extensive applicability of standing reservations for everything from daily dedicated job runs to improved use of resources on weekends. By default, standing reservations can overlap other reservations. Unless you set an ignore-type flag (`ACLOVERLAP`, `DEDICATEDRESOURCE`, `IGNIDLEJOBS`, or `IGNJOBRSV`), they are automatically given the `IGNRSV` flag. All standing reservation attributes are specified via the `SRCFG` parameter using the attributes listed in the table below.

#### Standing Reservation Attributes

### Standing Reservation Attributes

ACCESS	
<b>Format</b>	DEDICATED or SHARED
<b>Default</b>	---
<b>Description</b>	If set to SHARED, allows a standing reservation to use resources already allocated to other non-job reservations. Otherwise, these other reservations block resource access.
<b>Example</b>	<pre>SRCFG[test] ACCESS=SHARED</pre> <p><i>Standing reservation test can access resources allocated to existing standing and administrative reservations.</i></p> <p><b>i</b> The order that SRCFG reservations are listed in the configuration are important when using DEDICATED, because reservations made afterwards can steal resources later. During configuration, list DEDICATED reservations last to guarantee exclusiveness.</p>


ACCOUNTLIST	
<b>Format</b>	List of valid, comma-delimited account names (see <a href="#">ACL Modifiers</a> ).
<b>Default</b>	---
<b>Description</b>	Specifies that jobs with the associated accounts can use the resources contained within this reservation.
<b>Example</b>	<pre>SRCFG[test] ACCOUNTLIST=ops,staff</pre> <p><i>Jobs using the account ops or staff are granted access to the resources in standing reservation test.</i></p>

CHARGE	
<b>Format</b>	<BOOLEAN>



CHARGE	
<b>Default</b>	---
<b>Description</b>	<p>Overrides the default charging behavior:</p> <ul style="list-style-type: none"> <li>• If set to <code>True</code>, overrides <code>AMCFG[] RESERVATIONCHARGEPOLICY=Select</code> to indicate that this reservation should be charged, even if no <code>ChargeAccount</code> or <code>ChargeUser</code> are specified.</li> <li>• If set to <code>False</code>, overrides <code>AMCFG[] RESERVATIONCHARGEPOLICY=All</code> to indicate that this reservation should not be charged.</li> </ul> <p>If <code>AMCFG[] RESERVATIONCHARGEPOLICY=None</code>, this parameter has no effect.</p>
<b>Example</b>	<pre>SRCFG[sr_mam1] CHARGE=False</pre> <p><i>Prevents charges to this reservation (e.g., when <code>AMCFG[] RESERVATIONCHARGEPOLICY=ALL</code>).</i></p>

CHARGEACCOUNT	
<b>Format</b>	Any valid account name.
<b>Default</b>	---
<b>Description</b>	<p>Specifies that idle cycles for this reservation should be charged against the specified account (via the Accounting Manager).</p> <p><b>i</b> <code>CHARGEACCOUNT</code> must be used in conjunction with <code>CHARGEUSER</code>.</p>
<b>Example</b>	<pre>SRCFG[sr_mam1] CHARGEACCOUNT=math SRCFG[sr_mam1] CHARGEUSER=john</pre> <p><i>Moab charges all idle cycles within reservations supporting standing reservation <code>sr_mam1</code> to account <code>math</code>.</i></p>

CHARGEUSER	
<b>Format</b>	Any valid username.
<b>Default</b>	---
<b>Description</b>	<p>Specifies that idle cycles for this reservation should be charged against the specified user (via the Accounting Manager).</p> <div style="border: 1px solid #0056b3; border-radius: 10px; padding: 5px; display: inline-block;">  CHARGEUSER must be used in conjunction with <a href="#">CHARGEACCOUNT</a>.         </div>
<b>Example</b>	<pre>SRCFG[sr_mam1] CHARGEACCOUNT=math SRCFG[sr_mam1] CHARGEUSER=john</pre> <div style="border: 1px dashed #ccc; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p><i>Moab charges all idle cycles within reservations supporting standing reservation <code>sr_mam1</code> to user <code>john</code>.</i></p> </div>

CLASSLIST	
<b>Format</b>	List of valid, comma-delimited classes/queues (see <a href="#">ACL Modifiers</a> ).
<b>Default</b>	---
<b>Description</b>	Specifies that jobs with the associated classes/queues can use the resources contained within this reservation.
<b>Example</b>	<pre>SRCFG[test] CLASSLIST=!interactive</pre> <div style="border: 1px dashed #ccc; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p><i>Jobs not using the class <code>interactive</code> are granted access to the resources in standing reservation <code>test</code>.</i></p> </div>

CLUSTERLIST	
<b>Format</b>	List of valid, comma-delimited peer clusters (see <a href="#">Moab Workload Manager for Grids</a> ).
<b>Default</b>	---
<b>Description</b>	Specifies that jobs originating within the listed clusters can use the resources contained within this reservation.

CLUSTERLIST	
<b>Example</b>	<pre>SRCFG[test] CLUSTERLIST=orion2,orion7</pre> <p><i>Moab grants jobs from the listed peer clusters access to the reserved resources.</i></p>
COMMENT	
<b>Format</b>	<p>&lt;STRING&gt;</p> <p><b>i</b> If the string contains whitespace, it should be enclosed in single (') or double quotes (").</p>
<b>Default</b>	---
<b>Description</b>	Specifies a descriptive message associated with the standing reservation and all child reservations.
<b>Example</b>	<pre>SRCFG[test] COMMENT='rsv for network testing'</pre> <p><i>Moab annotates the standing reservation <code>test</code> and all child reservations with the specified message. These messages show up within Moab client commands, Moab web tools, and graphical administrator tools.</i></p>
DAYS	
<b>Format</b>	<p>One or more of the following (comma-delimited):</p> <ul style="list-style-type: none"> <li>• Mon</li> <li>• Tue</li> <li>• Wed</li> <li>• Thu</li> <li>• Fri</li> <li>• Sat</li> <li>• Sun</li> <li>• [ALL]</li> </ul>
<b>Default</b>	[ALL]

DAYS	
<b>Description</b>	Specifies which days of the week the standing reservation is active.
<b>Example</b>	<pre>SRCFG[test] DAYS=Mon,Tue,Wed,Thu,Fri</pre> <p><i>Standing reservation test is active Monday through Friday.</i></p>

DEPTH	
<b>Format</b>	<INTEGER>
<b>Default</b>	2
<b>Description</b>	<p>Specifies the depth of standing reservations to be created (one per <a href="#">period</a>).</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> To satisfy the <code>DEPTH</code>, Moab creates new reservations at the beginning of the specified <code>PERIOD</code>. If your reservation ends at the same time that a new <code>PERIOD</code> begins, the number of reservations might not match the requested <code>DEPTH</code>. To prevent or resolve this issue, set the <code>ENDTIME</code> a couple minutes before the beginning of the next <code>PERIOD</code>. For example, set the <code>ENDTIME</code> to <code>23:58</code> instead of <code>00:00</code>.</p> </div>
<b>Example</b>	<pre>SRCFG[test] PERIOD=DAY DEPTH=6</pre> <p><i>Specifies that six reservations will be created for standing reservation test.</i></p>

DISABLE	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	Specifies that the standing reservation should no longer spawn child reservations.
<b>Example</b>	<pre>SRCFG[test] PERIOD=DAY DEPTH=7 DISABLE=TRUE</pre>

DISABLE	
	<i>Specifies that reservations are created for standing reservation <code>test</code> for today and the next six days.</i>

ENDTIME	
<b>Format</b>	[[[DD:]HH:]MM:]SS
<b>Default</b>	24:00:00
<b>Description</b>	Specifies the time of day the standing reservation period ends (end of day or end of week depending on <a href="#">PERIOD</a> ).
<b>Example</b>	<pre>SRCFG[test] STARTTIME=8:00:00 SRCFG[test] ENDTIME=17:00:00 SRCFG[test] PERIOD=DAY</pre> <p><i>Standing reservation <code>test</code> is active from 8:00 AM until 5:00 PM.</i></p>

FLAGS	
<b>Format</b>	Comma-delimited list of zero or more flags listed in the <a href="#">reservation flags overview</a> .
<b>Default</b>	---
<b>Description</b>	Specifies special reservation attributes.
<b>Example</b>	<pre>SRCFG[test] FLAGS=BYNAME, DEDICATEDRESOURCE</pre> <p><i>Jobs can only access the resources within this reservation if they explicitly request the reservation by name. Further, the reservation is created to not overlap with other reservations.</i></p>

GROUPLIST	
<b>Format</b>	One or more comma-delimited group names.

GROUPLIST	
<b>Default</b>	[ALL]
<b>Description</b>	Specifies the groups allowed access to this standing reservation (see <a href="#">ACL Modifiers</a> ).
<b>Example</b>	<pre>SRCFG[test] GROUPLIST=staff,ops,special SRCFG[test] CLASSLIST=interactive</pre> <p><i>Moab allows jobs with the listed group IDs or which request the job class <code>interactive</code> to use the resources covered by the standing reservation.</i></p>

HOSTLIST	
<b>Format</b>	One or more comma-delimited host names or <i>host expressions</i> or the string "class:<classname>".
<b>Default</b>	---
<b>Description</b>	<p>Specifies the set of hosts that the scheduler can search for resources to satisfy the reservation. If specified using the "class:X" format, Moab only selects hosts that support the specified class. If TASKCOUNT is also specified, only TASKCOUNT tasks are reserved. Otherwise, all matching hosts are reserved.</p> <p><b>i</b> The HOSTLIST attribute is treated as host regular expression so <code>foo10</code> will map to <code>foo10</code>, <code>foo101</code>, <code>foo1006</code>, and so forth. To request an exact host match, the expression can be bounded by the caret and dollar symbol expression markers as in <code>^foo10\$</code>.</p> <p><b>i</b> When specifying resources for a hostlist, you can specify exact set, superset, or subset of nodes on which the job must run. Use the caret (^) or asterisk (*) characters to specify a hostlist as superset or subset respectively. See hostlist in <a href="#">Selecting Resources</a> for more information.</p> <p><b>i</b> When using <code>r:</code> ensure your node indexes are correct by customizing the NODEIDFORMAT parameter. See <a href="#">NODEIDFORMAT</a> for more information.</p>

HOSTLIST	
<b>Example</b>	<pre>SRCFG[test] HOSTLIST=node001,node002,node003 SRCFG[test] RESOURCES=PROCS:2;MEM:512 SRCFG[test] TASKCOUNT=2</pre> <p><i>Moab reserves a total of two tasks with 2 processors and 512 MB each, using resources located on node001, node002, and/or node003.</i></p>
	<pre>SRCFG[test] HOSTLIST=node01,node1[3-5]</pre> <p><i>The reservation will consume all nodes that have "node01" somewhere in their names and all nodes that have both "node1" and either a "3," "4," or "5" in their names.</i></p>
	<pre>SRCFG[test] HOSTLIST=r:node[1-6]</pre> <p><i>The reservation will consume all nodes with names that begin with "node" and end with any number 1 through 6. In other words, it will reserve node1, node2, node3, node4, node5, and node6.</i></p>

JOBATTRLIST	
<b>Format</b>	<p>Comma-delimited list of one or more of the following job attributes:</p> <ul style="list-style-type: none"> <li>• PREEMPTEE</li> <li>• INTERACTIVE</li> <li>• Any generic attribute configured through NODECFG</li> </ul>
<b>Default</b>	---
<b>Description</b>	<p>Specifies job attributes that grant a job access to the reservation.</p> <div style="border: 1px solid black; padding: 5px; margin-bottom: 5px;"> <p><b>i</b> Values can be specified with a "!=" assignment to only allow jobs NOT requesting a certain feature inside the reservation.</p> </div> <div style="border: 1px solid black; padding: 5px; margin-bottom: 5px;"> <p><b>i</b> To enable/disable reservation access based on requested node features, use the parameter NODETOJOBATTRMAP. This applies to any attribute, including PREEMPTEE and INTERACTIVE (which require uppercase).</p> </div> <div style="border: 1px solid black; padding: 5px;"> <p><b>i</b> JOBATTRLIST can serve as an ACL on its own. See <a href="#">Access Control List (ACL)</a> for more information.</p> </div>

JOBATTRLIST	
<b>Example</b>	<pre>SRCFG[test] JOBATTRLIST=PREEMPTEE NODETOJOBATTRMAP PREEMPTEE</pre> <p><i>Preemptible jobs can access the resources reserved within this reservation.</i></p>

MAXJOB	
<b>Format</b>	<INTEGER>
<b>Default</b>	---
<b>Description</b>	Specifies the maximum number of jobs that can run in the reservation.
<b>Example</b>	<pre>SRCFG[test] MAXJOB=1</pre> <p><i>Only one job will be allowed to run in this reservation.</i></p>

MAXTIME	
<b>Format</b>	[[[DD:]HH:]MM:]SS[+]
<b>Default</b>	---
<b>Description</b>	<p>Specifies the maximum time for jobs allowable. Can be used with Affinity to attract jobs with same MAXTIME.</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> The MAXTIME and TIMELIMIT reservation attributes perform the same function, and are therefore interchangeable. Also, either will serve as an ACL, as described in <a href="#">Access Control List (ACL)</a>, above.</p> </div>
<b>Example</b>	<pre>SRCFG[test] MAXTIME=1:00:00+</pre> <p><i>Jobs with a time of 1:00:00 are attracted to this reservation.</i></p>



NODEFEATURES	
<b>Format</b>	Comma-delimited list of node features.
<b>Default</b>	---
<b>Description</b>	Specifies the required node features for nodes that are part of the standing reservation.
<b>Example</b>	<pre>SRCFG[test] NODEFEATURES=wide, fddi</pre> <p><i>All nodes allocated to the standing reservation must have both the wide and fddi node attributes.</i></p>

OWNER	
<b>Format</b>	<p>&lt;CREDTYPE&gt;:&lt;CREDID&gt;</p> <p>Where &lt;CREDTYPE&gt; is one of USER, GROUP, ACCT, QoS, CLASS or CLUSTER and &lt;CREDTYPE&gt; is a valid credential ID of that type.</p>
<b>Default</b>	---
<b>Description</b>	<p>Specifies the owner of the reservation. Setting ownership for a reservation grants the user management privileges, including the power to release it.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-bottom: 5px;"> <p><b>i</b> Setting a USER as the OWNER of a reservation gives that user privileges to query and release the reservation.</p> </div> <div style="border: 1px solid #ccc; padding: 5px;"> <p><b>i</b> For sandbox reservations, sandboxes are applied to a specific peer only if OWNER is set to CLUSTER:&lt;PEERNAME&gt;.</p> </div>
<b>Example</b>	<pre>SRCFG[test] OWNER=ACCT:jupiter</pre> <p><i>User jupiter owns the reservation and can be granted special privileges associated with that ownership.</i></p>

PARTITION	
<b>Format</b>	Valid partition name.

PARTITION	
<b>Default</b>	[ALL]
<b>Description</b>	Specifies the partition in which to create the standing reservation.
<b>Example</b>	<pre>SRCFG[test] PARTITION=OLD</pre> <p><i>The standing reservation will only select resources from partition OLD.</i></p>

PERIOD	
<b>Format</b>	One of DAY, WEEK, or INFINITY
<b>Default</b>	DAY
<b>Description</b>	Specifies the period of the standing reservation.
<b>Example</b>	<pre>SRCFG[test] PERIOD=WEEK</pre> <p>Each standing reservation covers a one week period.</p>

PROCLIMIT	
<b>Format</b>	<p>&lt;QUALIFIER&gt;&lt;INTEGER&gt;</p> <p>&lt;QUALIFIER&gt; can be one of the following &lt;, &lt;=, ==, &gt;=, &gt;</p>
<b>Default</b>	---
<b>Description</b>	<p>Specifies the processor limit for jobs requesting access to this standing reservation.</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> PROCLIMIT can serve as an ACL on its own. See <a href="#">Access Control List (ACL)</a> for more information.</p> </div>
<b>Example</b>	<pre>SRCFG[test] PROCLIMIT&lt;=4</pre> <p><i>Jobs requesting 4 or fewer processors are allowed to run.</i></p>

PSLIMIT	
<b>Format</b>	<QUALIFIER><INTEGER> <QUALIFIER> can be one of the following <, <=, ==, >=, >
<b>Default</b>	---
<b>Description</b>	Specifies the processor-second limit for jobs requesting access to this standing reservation.
<b>Example</b>	<pre>SRCFG[test] PSLIMIT&lt;=40000</pre> <p><i>Jobs requesting 40000 or fewer processor-seconds are allowed to run.</i></p>

QOSLIST	
<b>Format</b>	Zero or more valid, comma-delimited QoS names.
<b>Default</b>	---
<b>Description</b>	Specifies that jobs with the listed QoS names can access the reserved resources.
<b>Example</b>	<pre>SRCFG[test] QOSLIST=hi,low,special</pre> <p><i>Moab allows jobs using the listed QoS access to the reserved resources.</i></p>

REQUIREDACCTLIST	
<b>Format</b>	One or more comma-delimited accounts.
<b>Default</b>	---
<b>Description</b>	<p>When present, any jobs in the reservation must match one of the listed accounts.</p> <p>This attribute can also be used in conjunction with <a href="#">REQUIREDUSERLIST</a>. If both <a href="#">REQUIREDACCTLIST</a> and <a href="#">REQUIREDUSERLIST</a> are specified, all jobs in the reservation must match both.</p>

REQUIREDACCTLIST	
	<p><b>i</b> It is recommended that any entries in the <code>REQUIREDACCTLIST</code> be present in the <code>ACCOUNTLIST</code> attribute to handle reservation affinities.</p>
<b>Example</b>	<pre>SRCFG[test] REQUIREDUSERLIST=john,bob USERLIST=john,bob SRCFG[test] REQUIREDACCTLIST=eng,chem ACCOUNTLIST=eng,chem</pre> <p><i>A job must belong to either user "john" or "bob" AND either account "eng" or "chem".</i></p>

REQUIREDTPN	
<b>Format</b>	<p>&lt;QUALIFIER&gt;&lt;INTEGER&gt;            &lt;QUALIFIER&gt; can be one of the following &lt;, &lt;=, ==, &gt;=, &gt;</p>
<b>Default</b>	---
<b>Description</b>	Restricts access to reservations based on the job's TPN (tasks per node).
<b>Example</b>	<pre>SRCFG[test] REQUIREDTPN==4</pre> <p><i>Jobs with tpn=4 or ppn=4 would be allowed within the reservation, but any other TPN value would not. (For more information, see <a href="#">TPN (Exact Tasks Per Node)</a>.)</i></p>

REQUIREDUSERLIST	
<b>Format</b>	One or more comma-delimited accounts.
<b>Default</b>	---
<b>Description</b>	<p>When present, any jobs in the reservation must match one of the listed users.</p> <p>This attribute also be used in conjunction with <code>REQUIREDACCTLIST</code>. If both <code>REQUIREDACCTLIST</code> and <code>REQUIREDUSERLIST</code> are specified, all jobs in the reservation must match both.</p>

REQUIREDUSERLIST	
	<p><b>i</b> It is recommended that any entries in the <code>REQUIREDUSERLIST</code> be present in the <code>USERLIST</code> attribute to handle reservation affinities.</p>
<b>Example</b>	<pre>SRCFG[test] REQUIREDUSERLIST=john,bob USERLIST=john,bob SRCFG[test] REQUIREDACCTLIST=eng,chem ACCOUNTLIST=eng,chem</pre> <p><i>A job must belong to either user "john" or "bob" AND either account "eng" or "chem".</i></p>

RESOURCES	
<b>Format</b>	Semicolon delimited <code>&lt;ATTR&gt;:&lt;VALUE&gt;</code> pairs, where <code>&lt;ATTR&gt;</code> can be one of <code>PROCS</code> , <code>MEM</code> , <code>SWAP</code> , <code>DISK</code> , or <code>GRES</code> .
<b>Default</b>	<code>PROCS : -1</code> (All processors available on node)
<b>Description</b>	<p>Specifies what resources constitute a single standing reservation task. (Each task must be able to obtain all of its resources as an atomic unit on a single node.) Supported resources currently include the following:</p> <ul style="list-style-type: none"> <li>• <code>PROCS</code> (number of processors)</li> <li>• <code>MEM</code> (real memory in MB)</li> <li>• <code>SWAP</code> (virtual memory in MB)</li> <li>• <code>DISK</code> (local disk in MB)</li> <li>• <code>GRES</code> (generic resource specified in the format <code>GRES:&lt;GRESNAME&gt;[:&lt;COUNT&gt;]</code>)</li> </ul>
<b>Example</b>	<pre>SRCFG[test] RESOURCES=PROCS:1;MEM:512;GRES=matlab:3;GRES=fluent:12</pre> <p><i>Each standing reservation task reserves one processor, 512 MB of real memory, 3 matlab generic resources and 12 fluent generic resources.</i></p>

ROLLBACKOFFSET	
<b>Format</b>	<code>[[[DD:]HH:]MM:]SS</code>
<b>Default</b>	---

ROLLBACKOFFSET	
<b>Description</b>	<p>Specifies the minimum time in the future at which the reservation can start. This offset is rollback meaning the start time of the reservation will continuously roll back into the future to maintain this offset. Rollback offsets are a good way of providing guaranteed resource access to users under the conditions that they must commit their resources in the future or lose dedicated access. See <a href="#">QoS</a> for more info about quality of service and service level agreements; also see <a href="#">Rollback Reservation Overview</a>.</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> Neither credlock nor advres is compatible on the jobs submitted for this reservation.</p> </div>
<b>Example</b>	<pre style="border: 1px dashed gray; padding: 5px;">SRCFG[ajax] ROLLBACKOFFSET=24:00:00 TASKCOUNT=32 SRCFG[ajax] PERIOD=INFINITY ACCOUNTLIST=ajax</pre> <p style="border: 1px dashed gray; padding: 5px; margin-top: 10px;"><i>The standing reservation guarantees access to up to 32 processors within 24 hours to jobs from the ajax account.</i></p> <p>Adding an asterisk to the ROLLBACKOFFSET value pins rollback reservation start times when an idle reservation is created in the rollback reservation. For example:</p> <pre style="border: 1px dashed gray; padding: 5px; margin-top: 10px;">SRCFG[staff] ROLLBACKOFFSET=18:00:00* PERIOD=INFINITY</pre>

RSVACCESSLIST	
<b>Format</b>	<RESERVATION>[,...]
<b>Default</b>	---
<b>Description</b>	A list of reservations to which the specified reservation has access.
<b>Example</b>	<pre style="border: 1px dashed gray; padding: 5px;">SRCFG[test] RSVACCESSLIST=rsv1,rsv2,rsv3</pre>

RSVGROUP	
<b>Format</b>	<STRING>
<b>Default</b>	---

RSVGROUP	
<b>Description</b>	See section <a href="#">Reservation Group</a> for a detailed description.
<b>Example</b>	<pre> SRCFG[test] RSVGROUP=rsvgrp1 SRCFG[ajax] RSVGROUP=rsvgrp1 </pre>

STARTTIME	
<b>Format</b>	[[[DD:]HH:]MM:]SS
<b>Default</b>	00:00:00:00 (midnight)
<b>Description</b>	<p>Specifies the time of day/week the standing reservation becomes active. Whether this indicates a time of day or time of week depends on the setting of the <a href="#">PERIOD</a> attribute.</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> If specified within a <a href="#">reservation profile</a>, a value of 0 indicates the reservation should start at the earliest opportunity.</p> </div>
<b>Example</b>	<pre> SRCFG[test] STARTTIME=08:00:00 SRCFG[test] ENDTIME=17:00:00 SRCFG[test] PERIOD=DAY </pre> <p><i>The standing reservation will be active from 8:00 a.m. until 5:00 p.m. each day.</i></p>

TASKCOUNT	
<b>Format</b>	<INTEGER>
<b>Default</b>	0 (unlimited tasks)
<b>Description</b>	Specifies how many tasks should be reserved for the reservation.
<b>Example</b>	<pre> SRCFG[test] RESOURCES=PROCS:1;MEM:256 SRCFG[test] TASKCOUNT=16 </pre> <p><i>Standing reservation test reserves 16 tasks worth of resources; in this case, 16 processors and 4 GB of real memory.</i></p>

TIMELIMIT	
<b>Format</b>	[[[DD:]HH:]MM:]SS
<b>Default</b>	-1 (no time based access)
<b>Description</b>	<p>Specifies the maximum allowed overlap between the standing reservation and a job requesting resource access.</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> The MAXTIME and TIMELIMIT reservation attributes perform the same function, and are therefore interchangeable. Also, either will serve as an ACL, as described in <a href="#">Access Control List (ACL)</a>, above.</p> </div>
<b>Example</b>	<pre>SRCFG[test] TIMELIMIT=1:00:00</pre> <p><i>Moab allows jobs to access up to one hour of resources in the standing reservation.</i></p>

TPN (Exact Tasks Per Node)	
<b>Format</b>	<INTEGER>
<b>Default</b>	0 (no TPN constraint)
<b>Description</b>	Specifies the exact number of tasks per node that must be available on eligible nodes.
<b>Example</b>	<pre>SRCFG[2] TPN=4 SRCFG[2] RESOURCES=PROCS:2;MEM:256</pre> <p><i>Moab must locate four tasks on each node that is to be part of the reservation. That is, each node included in standing reservation 2 must have 8 processors and 1 GB of memory available.</i></p>

TRIGGER	
<b>Format</b>	See <a href="#">Creating a Trigger</a> for syntax.
<b>Default</b>	N/A



TRIGGER	
<b>Description</b>	Specifies event triggers to be launched by the scheduler under the scheduler's ID. These triggers can be used to conditionally cancel reservations, or launch various actions at specified event offsets. See <a href="#">About Object Triggers</a> for more detail.
<b>Example</b>	<pre>SRCFG[fast] TRIGGER=EType=start,Offset=5:00:00,AType=exec,Action="/usr/local/ domail.pl"</pre> <p><i>Moab launches the domail.pl script 5 hours after any fast reservation starts.</i></p>

USERLIST	
<b>Format</b>	Comma-delimited list of users.
<b>Default</b>	---
<b>Description</b>	Specifies which users have access to the resources reserved by this reservation (see <a href="#">ACL Modifiers</a> ).
<b>Example</b>	<pre>SRCFG[test] USERLIST=bob,joe,mary</pre> <p><i>Users bob, joe and mary can all access the resources reserved within this reservation.</i></p>

## Standing Reservation Overview

A standing reservation is similar to a normal administrative reservation in that it also places an access control list on a specified set of resources. Resources are specified on a per-task basis and currently include processors, local disk, real memory, and swap. The access control list supported for standing reservations includes users, groups, accounts, job classes, and QoS levels. Standing reservations can be configured to be permanent or periodic on a daily or weekly basis and can accept a daily or weekly start and end time. Regardless of whether permanent or recurring on a daily or weekly basis, standing reservations are enforced using a series of reservations, extending a number of periods into the future as controlled by the `DEPTH` attribute of the `SRCFG` parameter.

The following examples demonstrate possible configurations specified with the `SRCFG` parameter.

### Basic Business Hour Standing Reservation

```

SRCFG[interactive] TASKCOUNT=6 RESOURCES=PROCS:1, MEM:512
SRCFG[interactive] PERIOD=DAY DAYS=MON, TUE, WED, THU, FRI
SRCFG[interactive] STARTTIME=9:00:00 ENDTIME=17:00:00
SRCFG[interactive] CLASSLIST=interactive

```

**i** When using the SRCFG parameter, attribute lists must be delimited using the comma (,), pipe (|), or colon (:) characters; they cannot be space delimited. For example, to specify a multi-class ACL, specify:

```
SRCFG[test] CLASSLIST=classA,classB
```

**i** Only one STARTTIME and one ENDTIME value can be specified per reservation. If varied start and end times are desired throughout the week, complementary standing reservations should be created. For example, to establish a reservation from 8:00 p.m. until 6:00 a.m. the next day during business days, two reservations should be created—one from 8:00 p.m. until midnight, and the other from midnight until 6:00 a.m. Jobs can run across reservation boundaries allowing these two reservations to function as a single reservation that spans the night. The following example demonstrates how to span a reservation across 2 days on the same nodes:

```

SRCFG[Sun] PERIOD=WEEK
SRCFG[Sun] STARTTIME=00:20:00:00 ENDTIME=01:00:00:00
SRCFG[Sun] HOSTLIST=node01,node02,node03

SRCFG[Mon] PERIOD=WEEK
SRCFG[Mon] STARTTIME=01:00:00:00 ENDTIME=01:06:00:00
SRCFG[Sun] HOSTLIST=node01,node02,node03

```

The preceding example fully specifies a reservation including the quantity of resources requested using the TASKCOUNT and RESOURCES attributes. In all cases, resources are allocated to a reservation in units called tasks where a task is a collection of resources that must be allocated together on a single node. The TASKCOUNT attribute specifies the number of these tasks that should be reserved by the reservation. In conjunction with this attribute, the RESOURCES attribute defines the reservation task by indicating what resources must be included in each task. In this case, the scheduler must locate and reserve 1 processor and 512 MB of memory together on the same node for each task requested.

As mentioned previously, a standing reservation reserves resources over a given time frame. The PERIOD attribute can be set to a value of DAY, WEEK, or INFINITY to indicate the period over which this reservation should recur. If not specified, a standing reservation recurs on a daily basis. If a standing reservation is configured to recur daily, the attribute DAYS can be specified to indicate which days of the week the reservation should exist. This attribute takes a comma-delimited list of days where each day is specified as the first three letters of the day in all capital letters: MON or FRI. The preceding

example specifies that this reservation is periodic on a daily basis and should only exist on business days.

The time of day during which the requested tasks are to be reserved is specified using the `STARTTIME` and `ENDTIME` attributes. These attributes are specified in standard military time `HH:MM:SS` format and both `STARTTIME` and `ENDTIME` specification is optional defaulting to midnight at the beginning and end of the day respectively. In the preceding example, resources are reserved from 9:00 a.m. until 5:00 p.m. on business days.

The final aspect of any reservation is the access control list indicating who or what can use the reserved resources. In the preceding example, the `CLASSLIST` attribute is used to indicate that jobs requesting the class "interactive" should be allowed to use this reservation.

## Specifying Reservation Resources

In most cases, only a small subset of standing reservation attributes must be specified in any given case. For example, by default, `RESOURCES` is set to `PROCS=-1`, which indicates that each task should reserve all of the processors on the node on which it is located. This, in essence, creates a one task equals one node mapping. In many cases, particularly on uniprocessor systems, this default behavior might be easiest to work with. However, in SMP environments, the `RESOURCES` attribute provides a powerful means of specifying an exact, multi-dimensional resource set.

**i** An examination of the parameters documentation shows that the default value of `PERIOD` is `DAYS`. Therefore, specifying this parameter in the preceding above was unnecessary. It was used only to introduce this parameter and indicate that other options exist beyond daily standing reservations.

## Host Constrained Standing Reservation

Although the first example did specify a quantity of resources to reserve, it did not specify where the needed tasks were to be located. If this information is not specified, Moab attempts to locate the needed resources anywhere it can find them. The Example 1 reservation essentially discovers hosts where the needed resources can be found. If the `SPACEFLEX` reservation flag is set, then the reservation continues to float to the best hosts over the life of the reservation. Otherwise, it will be locked to the initial set of allocated hosts.

If a site wanted to constrain a reservation to a subset of available resources, this could be accomplished using the `HOSTLIST` attribute. The `HOSTLIST` attribute is specified as a comma-separated list of hostnames and constrains the scheduler to only select tasks from the specified list. This attribute can exactly specify hosts or specify them using host regular expressions. The following example demonstrates a possible use of the `HOSTLIST` attribute:

```

SRCFG[interactive] DAYS=MON,TUE,WED,THU,FRI
SRCFG[interactive] PERIOD=DAY
SRCFG[interactive] STARTTIME=10:00:00 ENDTIME=15:00:00
SRCFG[interactive] RESOURCES=PROCS:2, MEM:256
SRCFG[interactive] HOSTLIST=node001,node002,node005,node020
SRCFG[interactive] TASKCOUNT=6
SRCFG[interactive] CLASSLIST=interactive

```

*Note that the `HOSTLIST` attribute specifies a non-contiguous list of hosts. Any combination of hosts can be specified and hosts can be specified in any order. In this example, the `TASKCOUNT` attribute is also specified. These two attributes both apply constraints on the scheduler with `HOSTLIST` specifying where the tasks can be located and `TASKCOUNT` indicating how many total tasks can be allocated. In this example, six tasks are requested but only four hosts are specified. To handle this, if adequate resources are available, the scheduler may attempt to allocate more than one task per host. For example, assume that each host is a quad-processor system with 1 GB of memory. In such a case, the scheduler could allocate up to two tasks per host and even satisfy the `TASKCOUNT` constraint without using all of the hosts in the hostlist.*

**i** It is important to note that even if there is a one to one mapping between the value of `TASKCOUNT` and the number of hosts in `HOSTLIST`, the scheduler will not necessarily place one task on each host. If, for example, `node001` and `node002` were 8 processor SMP hosts with 1 GB of memory, the scheduler could locate up to four tasks on each of these hosts fully satisfying the reservation taskcount without even partially using the remaining hosts. (Moab will place tasks on hosts according to the policy specified with the `NODEALLOCATIONPOLICY` parameter.) If the hostlist provides more resources than what is required by the reservation as specified via `TASKCOUNT`, the scheduler will simply select the needed resources within the set of hosts listed.

## Enforcing Policies Via Multiple Reservations

Single reservations enable multiple capabilities. Combinations of reservations can further extend a site's capabilities to impose specific policies.

### Reservation Stacking

If `HOSTLIST` is specified but `TASKCOUNT` is not, the scheduler will pack as many tasks as possible onto all of the listed hosts. For example, assume the site added a second standing reservation named *debug* to its configuration that reserved resources for use by certain members of its staff using the following configuration:

```

SRCFG[interactive] DAYS=MON,TUE,WED,THU,FRI
SRCFG[interactive] PERIOD=DAY
SRCFG[interactive] STARTTIME=10:00:00 ENDTIME=15:00:00
SRCFG[interactive] RESOURCES=PROCS:2, MEM:256
SRCFG[interactive] HOSTLIST=node001,node002,node005,node020
SRCFG[interactive] TASKCOUNT=6
SRCFG[interactive] CLASSLIST=interactive
SRCFG[debug] HOSTLIST=node001,node002,node003,node004
SRCFG[debug] USERLIST=helpdesk
SRCFG[debug] GROUPLIST=operations,sysadmin

```

```
SRCFG[debug]          PERIOD=INFINITY
```

The new standing reservation is quite simple. Since `RESOURCES` is not specified, it will allocate all processors on each host that is allocated. Since `TASKCOUNT` is not specified, it will allocate every host listed in `HOSTLIST`. Since `PERIOD` is set to `INFINITY`, the reservation is always in force and there is no need to specify `STARTTIME`, `ENDTIME`, or `DAYS`.

The standing reservation has two access parameters set using the attributes `USERLIST` and `GROUPLIST`. This configuration indicates that the reservation can be accessed if any one of the access lists specified is satisfied by the job. In essence, reservation access is logically `ORed` allowing access if the requester meets any of the access constraints specified. In this example, jobs submitted by either user `helpdesk` or any member of the groups `operations` or `sysadmin` can use the reserved resources (see [ACL Modifiers](#)).

Unless [ACL Modifiers](#) are specified, access is granted to the logical `OR` of access lists specified within a standing reservation and granted to the logical `AND` of access lists across different standing reservations. A comparison of the standing reservations `interactive` and `debug` in the preceding example indicates that they both can allocate hosts `node001` and `node002`. If `node001` had both of these reservations in place simultaneously and a job attempted to access this host during business hours when standing reservation `interactive` was active. The job could only use the *doubly* reserved resources if it requests the run class `interactive` and it meets the constraints of reservation `debug`—that is, that it is submitted by user `helpdesk` or by a member of the group `operations` or `sysadmin`.

As a rule, the scheduler does not stack reservations unless it must. If adequate resources exist, it can allocate reserved resources side by side in a single SMP host rather than on top of each other. In the case of a 16 processor SMP host with two 8 processor standing reservations, 8 of the processors on this host will be allocated to the first reservation, and 8 to the next. Any configuration is possible. The 16 processor hosts can also have 4 processors reserved for user "John," 10 processors reserved for group "Staff," with the remaining 2 processors available for use by any job.

Stacking reservations is not usually required but some site administrators choose to do it to enforce elaborate policies. There is no problem with doing so as long as you can keep things straight. It really is not too difficult a concept; it just takes a little getting used to. See the [Reservation Overview](#) section for a more detailed description of reservation use and constraints.

As mentioned earlier, by default the scheduler enforces standing reservations by creating a number of reservations where the number created is controlled by the `DEPTH` attribute. Each night at midnight, the scheduler updates its periodic non-floating standing reservations. By default, `DEPTH` is set to 2, meaning when the scheduler starts up, it will create two 24-hour reservations covering a total of two days' worth of time—a reservation for today and one for tomorrow. For daily reservations, at midnight, the reservations roll,

meaning today's reservation expires and is removed, tomorrow's reservation becomes today's, and the scheduler creates a new reservation for the next day.

With this model, the scheduler continues creating new reservations in the future as time moves forward. Each day, the needed resources are always reserved. At first, all appears automatic but the standing reservation `DEPTH` attribute is in fact an important aspect of reservation rollback, which helps address certain site specific environmental factors. This attribute remedies a situation that might occur when a job is submitted and cannot run immediately because the system is backlogged with jobs. In such a case, available resources might not exist for several days out and the scheduler must reserve these future resources for this job. With the default `DEPTH` setting of two, when midnight arrives, the scheduler attempts to roll its standing reservations but a problem arises in that the job has now allocated the resources needed for the standing reservation two days out. Moab cannot reserve the resources for the standing reservation because they are already claimed by the job. The standing reservation reserves what it can but because all needed resources are not available, the resulting reservation is now smaller than it should be, or is possibly even empty.

If a standing reservation is smaller than it should be, the scheduler will attempt to add resources each iteration until it is fully populated. However, in the case of this job, the job is not going to release its reserved resources until it completes and the standing reservation cannot claim them until this time. The `DEPTH` attribute allows a site to specify how deep into the future a standing reservation should reserve its resources allowing it to claim the resources first and prevent this problem. If a partial standing reservation is detected on a system, it might be an indication that the reservation's `DEPTH` attribute should be increased.

In Example 3, the `PERIOD` attribute is set to `INFINITY`. With this setting, a single, permanent standing reservation is created and the issues of resource contention do not exist. While this eliminates the contention issue, infinite length standing reservations cannot be made periodic.

## Multiple ACL Types

In most cases, access lists within a reservation are logically ORed together to determine reservation access. However, exceptions to this rule can be specified by using the required ACL marker-the asterisk (\*). Any ACL marked with this symbol is required and a job is only allowed to use a reservation if it meets all required ACLs and at least one non-required ACL (if specified). A common use for this facility is in conjunction with the `TIMELIMIT` attribute. This attribute controls the length of time a job can use the resources within a standing reservation. This access mechanism can be ANDed or ORed to the cumulative set of all other access lists as specified by the required ACL marker. Consider the following example configuration:

```

SRCFG[special] TASKCOUNT=32
SRCFG[special] PERIOD=WEEK
SRCFG[special] STARTTIME=1:08:00:00
SRCFG[special] ENDTIME=5:17:00:00
SRCFG[special] NODEFEATURES=largememory
SRCFG[special] TIMELIMIT=1:00:00*
SRCFG[special] QOSLIST=high,low,special-
SRCFG[special] ACCOUNTLIST=!projectX,!projectY

```

The above configuration requests 32 tasks, which translate to 32 nodes. The `PERIOD` attribute makes this reservation periodic on a weekly basis while the attributes `STARTTIME` and `ENDTIME` specify the week offsets when this reservation is to start and end (note that the specification format has changed to DD:HH:MM:SS). In this case, the reservation starts on Monday at 8:00 a.m. and runs until Friday at 5:00 p.m. The reservation is enforced as a series of weekly reservations that only cover the specified time frame. The `NODEFEATURES` attribute indicates that each of the reserved nodes must have the node feature "largememory" configured.

As described earlier, `TIMELIMIT` indicates that jobs using this reservation can only use it for one hour. This means the job and the reservation can only overlap for one hour. Clearly jobs requiring an hour or less of wallclock time meet this constraint. However, a four-hour job that starts on Monday at 5:00 a.m. or a 12-hour job that starts on Friday at 4:00 p.m. also satisfies this constraint. Also, note the `TIMELIMIT` required ACL marker, `*`; it is set indicating that jobs must not only meet the `TIMELIMIT` access constraint but must also meet one or more of the other access constraints. In this example, the job can use this reservation if it can use the access specified via `QOSLIST` or `ACCOUNTLIST`; that is, it is assigned a QoS of `high`, `low`, or `special`, or the submitter of the job has an account that satisfies the `!projectX` and `!projectY` criteria. See the [QoS Overview](#) for more info about QoS configuration and usage.

## Affinity

Reservation ACLs allow or deny access to reserved resources but they can be configured to also impact a job's affinity for a particular reservation. By default, jobs gravitate toward reservations through a mechanism known as positive affinity. This mechanism allows jobs to run on the most constrained resources leaving other, unreserved resources free for use by other jobs that might not be able to access the reserved resources. Normally this is a desired behavior. However, sometimes, it is desirable to reserve resources for use only as a last resort—using the reserved resources only when there are no other resources available. This last resort behavior is known as negative affinity. Note the `'-` (hyphen or negative sign) following the `special` in the `QOSLIST` values. This special mark indicates that QoS `special` should be granted access to this reservation but should be assigned negative affinity. Therefore, the `QOSLIST` attribute specifies that QoS `high` and `low` should be granted access with positive affinity (use the reservation first where possible) and QoS `special` granted access with negative affinity (use the reservation only when no other resources are available).

Affinity status is granted on a per access object basis rather than a per access list basis and always defaults to positive affinity. In addition to negative affinity, neutral affinity can also be specified using the equal sign (=) as in `QOSLIST[0] normal= high debug= low-`.

When a job matches multiple ACLs for a reservation, the final node affinity for the node, job, and reservation combination is based on the last matching ACL entry found in the configuration file.

For example, given the following reservation ACLs, a job matching both will receive a negative affinity:

```
SRCFG[res1] USERLIST=joe+ MAXTIME<=4:00:00-
```

With the following reservation ACLs, a job matching both will receive a positive affinity:

```
SRCFG[res1] MAXTIME<=4:00:00- USERLIST=joe+
```

**i** To configure the behavior when multiple reservations with varying affinities are on the same node, see [NODEAFFINITYPOLICY](#).

## ACL Modifiers

ACL modifiers allow a site to change the default behavior of ACL processing. By default, a reservation can be accessed if one or more of its ACLs can be met by the requestor. This behavior can be changed using the following modifiers.

Not	
<b>Symbol:</b>	! (exclamation point)
<b>Description</b>	If attribute is met, the requestor is denied access regardless of any other satisfied ACLs.
<b>Example</b>	<pre>SRCFG[test] GROUPLIST=staff USERLIST=!steve</pre> <p><i>Allow access to all staff members other than steve.</i></p>
Required	
<b>Symbol:</b>	* (asterisk)
<b>Description</b>	All required ACLs must be satisfied for requestor access to be granted.



Required	
<b>Example</b>	<pre>SRCFG[test] QOSLIST=*high MAXTIME=*2:00:00</pre> <p><i>Only jobs in QoS high that request less than 2 hours of walltime are granted access.</i></p>

XOR	
<b>Symbol:</b>	^ (carat)
<b>Description</b>	All attributes of the type specified other than the ones listed in the ACL satisfy the ACL.
<b>Example</b>	<pre>SRCFG[test] QOSLIST ^=high</pre> <p><i>All jobs other than those requesting QoS high are granted access.</i></p>

CredLock	
<b>Symbol:</b>	& (ampersand)
<b>Description</b>	Matching jobs will be required to run on the resources reserved by this reservation. You can use this modifier on accounts, classes, groups, qualities of service, and users.
<b>Example</b>	<pre>SRCFG[test] USERLIST=&amp;john</pre> <p><i>All of user john's jobs must run in this reservation.</i></p>

HPEnable (hard policy enable)	
<b>Symbol:</b>	~ (tilde)
<b>Description</b>	ACLs marked with this modifier are ignored during soft policy scheduling and are only considered for hard policy scheduling once all eligible soft policy jobs start.
<b>Example</b>	<pre>SRCFG[johnspace] USERLIST=john CLASSLIST=~debug</pre>

**HPEnable (hard policy enable)**

*All of user john's jobs are allowed to run in the reservation at any time. Debug jobs are also allowed to run in this reservation but are only considered after all of John's jobs are given an opportunity to start. User john's jobs are considered before debug jobs regardless of job priority.*

**i** If HPEnable and **Not** markers are used in conjunction, then specified credentials are *blocked-out* of the reservation during soft-policy scheduling.

Note the ACCOUNTLIST values in are preceded with an exclamation point, or NOT symbol. This indicates that all jobs with accounts other than `projectX` and `projectY` meet the account ACL. Note that if a `!<X>` value (`!projectX`) appears in an ACL line, that ACL is satisfied by any object not explicitly listed by a NOT entry. Also, if an object matches a NOT entry, the associated job is excluded from the reservation even if it meets other ACL requirements. For example, a QoS 3 job requesting account `projectX` is denied access to the reservation even though the job QoS matches the QoS ACL.

**Example 6-5: Binding Users to Reservations at Reservation Creation**

```
# create a 4 node reservation for john and bind all of john's jobs to that reservation
> mrsvctl -c -a user=&john -t 4
```

**Reservation Ownership**

Reservation ownership allows a site to control who owns the reserved resources during the reservation time frame. Depending on needs, this ownership can be identical to, a subset of, or completely distinct from the reservation ACL. By default, reservation ownership implies resource accountability and resources not consumed by jobs are accounted against the reservation owner. In addition, ownership can also be associated with special privileges within the reservation.

Ownership is specified using the OWNER attribute in the format `<CREDTYPE>:<CREDID>`, as in `OWNER=USER:john`. To enable john's jobs to preempt other jobs using resources within the reservation, the SRCFG attribute FLAG should be set to **OWNERPREEMPT**. In the example below, the `jupiter` project chooses to share resources with the `saturn` project but only when it does not currently need them.

**Limited Shared Access**

```
ACCOUNTCFG[jupiter] PRIORITY=10000
SRCFG[jupiter] HOSTLIST=node0[1-9]
SRCFG[jupiter] PERIOD=INFINITY
SRCFG[jupiter] ACCOUNTLIST=jupiter,saturn-
SRCFG[jupiter] OWNER=ACCT:jupiter
```

```
SRCFG[jupiter] FLAGS=OWNERPREEMPT
```

## Partitions

A reservation can be used in conjunction with a partition. Configuring a standing reservation on a partition allows constraints to be (indirectly) applied to a partition.

### Example 6-6: Time Constraints by Partition

The following example places a 3-day wall-clock limit on two partitions and a 64 processor-hour limit on jobs running on partition `small`:

```
SRCFG[smallrsv] PARTITION=small MAXTIME=3:00:00:00 PSLIMIT<=230400 HOSTLIST=ALL
SRCFG[bigrsv] PARTITION=big MAXTIME=3:00:00:00 HOSTLIST=ALL
```

## Resource Allocation Behavior

As mentioned, standing reservations can operate in one of two modes, floating, or non-floating (essentially node-locked). A floating reservation is created when the flag `SPACEFLEX` is specified. If a reservation is non-floating, the scheduler allocates all resources specified by the `HOSTLIST` parameter regardless of node state, job load, or even the presence of other standing reservations. Moab interprets the request for a non-floating reservation as, "I want a reservation on these exact nodes, no matter what!"

If a reservation is configured to be floating, the scheduler takes a more relaxed stand, searching through all possible nodes to find resources meeting standing reservation constraints. Only **Idle**, **Running**, or **Busy** nodes are considered and further, only considered if no reservation conflict is detected. The reservation attribute `ACCESS` modifies this behavior slightly and allows the reservation to allocate resources even if reservation conflicts exist.

**i** If a `TASKCOUNT` is specified with or without a `HOSTEXPRESSION`, Moab will, by default, only consider "up" nodes for allocation. To change this behavior, the reservation flag `IGNSTATE` can be specified as in the following example:

```
SRCFG[nettest] GROUPLIST=sysadm
SRCFG[nettest] FLAGS=IGNSTATE
SRCFG[nettest] HOSTLIST=node1[3-8]
SRCFG[nettest] STARTTIME=9:00:00
SRCFG[nettest] ENDTIME=17:00:00
```

**i** Access to existing reservations can be controlled using the reservation flag `IGNRSV`.

Other standing reservation attributes not covered here include `PARTITION` and `CHARGEACCOUNT`. These parameters are described in some detail in the [parameters](#) documentation.

## Using Reservations to Guarantee Turnover

In some cases, it is desirable to make certain a portion of a cluster's resources are available within a specific time frame. The following example creates a floating reservation belonging to the `jupiter` account that guarantees 16 tasks for use by jobs requesting up to one hour:

```

SRCFG[shortpool] OWNER=ACCT:jupiter
SRCFG[shortpool] FLAGS=SPACEFLEX
SRCFG[shortpool] MAXTIME=1:00:00
SRCFG[shortpool] TASKCOUNT=16
SRCFG[shortpool] STARTTIME=9:00:00
SRCFG[shortpool] ENDTIME=17:00:00
SRCFG[shortpool] DAYS=Mon,Tue,Wed,Thu,Fri

```

This reservation enables a capability similar to what was known in early Maui releases as 'shortpool'. The reservation covers every weekday from 9:00 a.m. to 5:00 p.m., reserving 16 tasks and allowing jobs to overlap the reservation for up to one hour. The `SPACEFLEX` flag indicates that the reservation can be dynamically modified--over time to re-locate to more optimal resources. In the case of a reservation with the `MAXTIME` ACL, this would include migrating to resources that are in use but that free up within the `MAXTIME` time frame. Additionally, because the `MAXTIME` ACL defaults to positive [affinity](#), any jobs that fit the ACL attempt to use available reserved resources first before looking elsewhere.

## Rollback Reservations

Rollback reservations are enabled using the [ROLLBACKOFFSET](#) attribute and can be used to allow users guaranteed access to resources, but the guaranteed access is limited to a time-window in the future. This functionality forces users to commit their resources in the future or lose access. In Iteration 1 of the diagram below, a rollback reservation is in place for nodes 2 and 3 for a 6.5 hour block, 2 hours in the future. If the user for whom the reservation is made does not make use of the reserved nodes, other jobs are scheduled on the nodes. In Iteration 2, the rollback reservation remains in place for the same nodes, 2 hours in the future.

Image 6-2: Rollback reservation Iteration 1


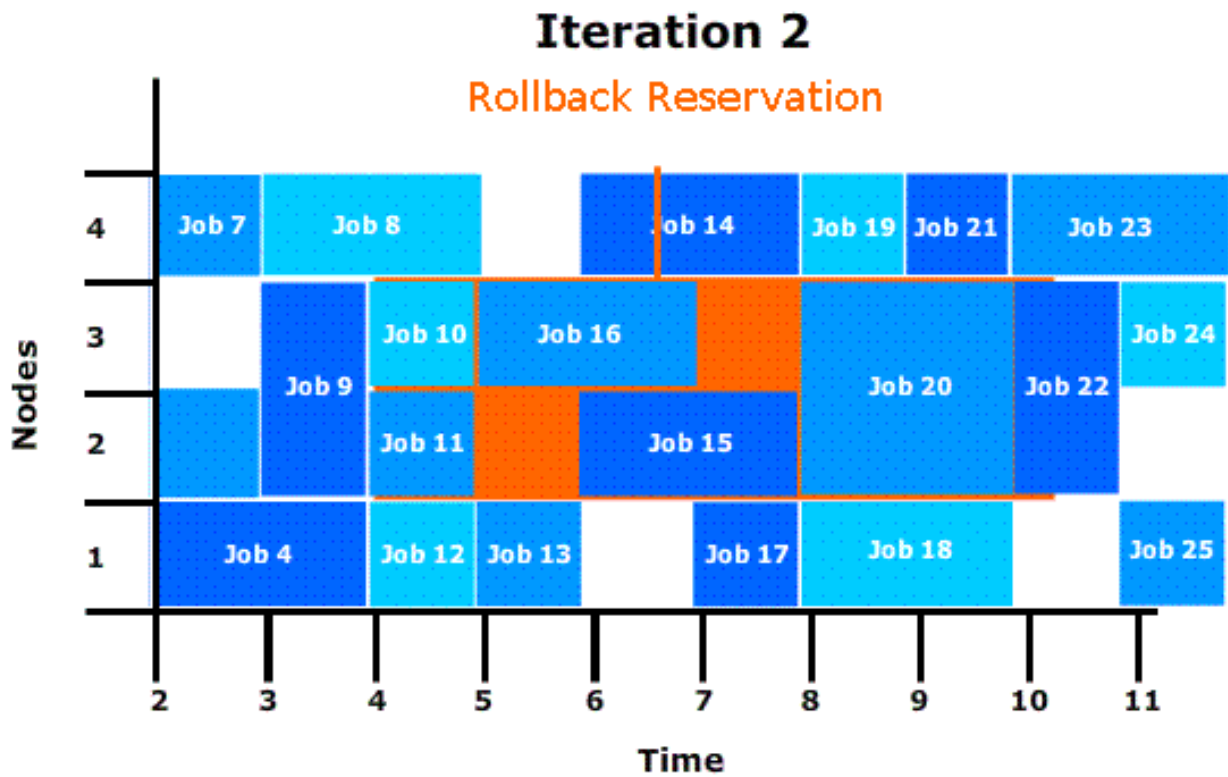


Image 6-3: Rollback reservation Iteration 2



### Rollback Reservations

```

SRCFG[ajax] ROLLBACKOFFSET=24:00:00 TASKCOUNT=32
SRCFG[ajax] PERIOD=INFINITY ACCOUNTLIST=ajax
    
```

Adding an asterisk to the ROLLBACKOFFSET value pins rollback reservation start times when an idle reservation is created in the rollback reservation. For example: SRCFG [staff] ROLLBACKOFFSET=18:00:00\* PERIOD=INFINITY.

#### 6.1.5.C Managing Administrative Reservations

A default reservation with no ACL is termed an *administrative* reservation, but is occasionally referred to as a *system* reservation. It blocks access to all jobs because it possesses an empty access control list. It is often useful when performing administrative tasks but cannot be used for enforcing resource usage policies.

Administrative reservations are created and managed using the `mrsvctl` command. With this command, all aspects of reservation time frame, resource selection, and access control can be dynamically modified. The `mdiag -r` command can be used to view configuration,

state, allocated resource information, as well as identify any potential problems with the reservation. The following table briefly summarizes commands used for common actions. More detailed information is available in the command summaries.

Action	Command
<b>create reservation</b>	<code>mrsvctl -c &lt;RSV_DESCRIPTION&gt;</code>
<b>list reservations</b>	<code>mrsvctl -l</code>
<b>release reservation</b>	<code>mrsvctl -r &lt;RSVID&gt;</code>
<b>modify reservation</b>	<code>mrsvctl -m &lt;ATTR&gt;=&lt;VAL&gt; &lt;RSVID&gt;</code>
<b>query reservation configuration</b>	<code>mdiag -r &lt;RSVID&gt;</code>
<b>display reservation hostlist</b>	<code>mrsvctl -q resources &lt;RSVID&gt;</code>

---

### Related Topics

- [SRCFG](#) (configure standing reservations)
- [RSVPROFILE](#) (create reservation profiles)

## 6.1.6 Personal/User Reservations - Enabling Reservations for End Users

In this topic:

- [6.1.6.A Enabling Personal Reservation Management - page 516](#)
- [6.1.6.B Reservation Accountability - page 516](#)
- [6.1.6.C Reservation Limits - page 517](#)
- [6.1.6.D Reservation and Job Binding - page 517](#)

By default, advance reservations are only available to scheduler administrators. While administrators can create and manage reservations to provide resource access to end-users, end-users cannot create, modify, or destroy these reservations. Moab extends the ability to manage reservations to end-users and provides control facilities to keep these

features manageable. Reservations created by end-users are called personal reservations or user reservations.

### 6.1.6.A Enabling Personal Reservation Management

User, or personal, reservations can be enabled on a per QoS basis by setting the `ENABLEUSERRSV` flag as in the following example:

```
QOSCFG[titan]    QFLAGS=ENABLEUSERRSV # allow 'titan' QOS jobs to create user
reservations
USERCFG[DEFAULT] QDEF=titan          # allow all users to access 'titan' QOS
...
```

If set, end-users are allowed to create, modify, cancel, and query reservations they own. As with jobs, users can associate a personal reservation with any QoS or account to which they have access. This is accomplished by specifying per reservation accountable credentials as in the following example:

```
> mrsvctl -c -S AQOS=titan -h node01 -d 1:00:00 -s 1:30:00
Note: reservation test.126 created
```

As in the preceding example, a non-administrator user who wants to create a reservation must *ALWAYS* specify an accountable QoS with the `mrsvctl -S` flag. This specified QoS must have the `ENABLEUSERRSV` flag. By default, a personal reservation is created with an ACL of only the user who created it.

*Example 6-7: Allow All Users in Engineering Group to Create Personal Reservations*

```
QOSCFG[rsv]      QFLAGS=ENABLEUSERRSV # allow 'rsv' QOS jobs to create user
reservations
GROUPCFG[sales] QDEF=rsv             # allow all users in group sales to access 'rsv'
QOS
...
```

*Example 6-8: Allow Specific Users to Create Personal Reservations*

```
# special qos has higher job priority and ability to create user reservations
QOSCFG[special] QFLAGS=ENABLEUSERRSV
QOSCFG[special] PRIORITY=1000
# allow betty and steve to use the special qos
USERCFG[betty]  QDEF=special
USERCFG[steve]  QLIST=fast,special,basic QDEF=rsv
...
```

### 6.1.6.B Reservation Accountability

Personal reservations must be configured with a set of accountable credentials. These credentials (user, group, account, and so forth) indicate who is responsible for the resources dedicated by the reservation. If resources are dedicated by a reservation but not consumed by a job, these resources can be charged against the specified accountable



credentials. Administrators are allowed to create reservations and specify any accountable credentials for that reservation. While end-users can also be allowed to create and otherwise modify personal reservations, they are only allowed to create reservations with accountable credentials to which they have access. Further, while administrators can manage any reservation, end-users can only control reservations they own.

Like jobs, reservation accountable credentials specify which credentials are charged for reservation usage and what policies are enforced as far as usage limits and allocation management is concerned. See the [mrvsctl](#) command documentation for more information on setting personal reservation credentials. While similar to jobs, personal reservations do have a separate set of usage limits and different allocation charging policies.

## Setting Reservation Default Attributes

Organizations can use [reservation profiles](#) to set default attributes for personal reservations. These attributes can include reservation aspects such as management policies, charging credentials, ACLs, host constraints, and time frame settings.

### 6.1.6.C Reservation Limits

Allowing end-users the ability to create advance reservations can lead to potentially unfair and unproductive resource usage. This results from the fact that by default, there is nothing to prevent a user from reserving all resources in a given system or reserving resources during time slots that would greatly impede the scheduler's ability to schedule jobs efficiently. Because of this, it is highly advised that sites initially place either usage or allocation based constraints on the use of personal reservations. This can be achieved using Moab Accounting Manager (see the *Moab Accounting Manager Administrator Guide*).

### 6.1.6.D Reservation and Job Binding

Moab allows job-to-reservation binding to be configured at an administrator or end-user level. This binding constrains how job to reservation mapping is allowed.

## Constraining a Job to Only Run in a Particular Reservation

Jobs can be bound to a particular reservation at submit time (using the RM extension [ADVRES](#)) or dynamically using the [mjobctl](#) command (see [Job to Reservation Mapping](#)). In either case, once bound to a reservation, a job can only run in that reservation even if other resources can be found outside of that reservation. The *mjobctl* command can also be used to dynamically release a job from reservation binding.

*Example 6-9: Bind job to reservation*

```
> mjobctl -m flags+=advres:grid.3 job1352
```

*Example 6-10: Release job from reservation binding*

```
> mjobctl -m flags==advres job1352
```

## Constraining a Reservation to Only Accept Certain Jobs

Binding a job to a reservation is independent of binding a reservation to a job. For example, a reservation may be created for user 'steve.' User 'steve' may then submit a number of jobs including one that is bound to that reservation using the `ADVRES` attribute. However, this binding simply forces that one job to use the reservation, it does not prevent the reservation from accepting other jobs submitted by user 'steve.' To prevent these other jobs from using the reserved resources, reservation to job binding must occur. This binding is accomplished by specifying either general job binding or specific job binding.

General job binding is the most flexible form of binding. Using the `BYNAME` attribute, a reservation can be created that only accepts jobs specifically bound to it.

Specific job binding is more constraining. This form of binding causes the reservation to only accept specific jobs, regardless of other job attributes and is set using the `JOB` reservation ACL.

*Example 6-11: Configure a reservation to accept only jobs that are bound to it*

```
> mrsvctl -m flags+=byname grid.3
```

*Example 6-12: Remove general reservation to job binding*

```
> mrsvctl -m flags-=byname grid.3
```

*Example 6-13: Configure a reservation to accept a specific job*

```
> mrsvctl -m -a JOB=3456 grid.3
```

*Example 6-14: Remove a specific reservation to job binding*

```
> mrsvctl -m -a JOB=3456 grid.3 --flags=unset
```

## 6.2 Partitions

In this topic:

[6.2.1 Partition Overview - page 519](#)

[6.2.2 Defining Partitions - page 520](#)

[6.2.3 Managing Partition Access - page 520](#)

[6.2.4 Requesting Partitions - page 521](#)

[6.2.5 Per-Partition Settings - page 522](#)

[6.2.6 Miscellaneous Partition Issues - page 523](#)

## 6.2.1 Partition Overview

Partitions are a logical construct that divide available resources. Any single resource (compute node) can only belong to a single partition. Often, natural hardware or resource manager bounds delimit partitions such as in the case of disjoint networks and diverse processor configurations within a cluster. For example, a cluster might consist of 256 nodes containing four 64 port switches. This cluster might receive excellent interprocess communication speeds for parallel job tasks located within the same switch but sub-stellar performance for tasks that span switches. To handle this, the site may choose to create four partitions, allowing jobs to run within any of the four partitions but not span them.

While partitions do have value, it is important to note that within Moab, the [standing reservation](#) facility provides significantly improved flexibility and should be used in the vast majority of politically motivated cases where partitions may be required under other resource management systems. Standing reservations provide time flexibility, improved access control features, and more extended resource specification options. Also, another Moab facility called [Node Sets](#) allows intelligent aggregation of resources to improve per job node allocation decisions. In cases where system partitioning is considered for such reasons, node sets may be able to provide a better solution.

Still, one key advantage of partitions over standing reservations and node sets is the ability to specify partition specific policies, limits, priorities, and scheduling algorithms although this feature is rarely required. An example of this need may be a cluster consisting of 48 nodes owned by the Astronomy Department and 16 nodes owned by the Mathematics Department. Each department may be willing to allow sharing of resources but wants to specify how their partition will be used. As mentioned, many of Moab's scheduling policies can be specified on a per partition basis allowing each department to control the scheduling goals within their partition.

The partition associated with each node should be specified as indicated in the [Node Location](#) section. With this done, partition access lists can be specified on a per job or per QoS basis to constrain which resources a job can have access to. See the [QoS Overview](#) for more information. By default, QoSs and jobs allow global partition access. Note that by default, a job can only use resources within a single partition.

If no partition is specified, Moab creates one partition per resource manager into which all resources corresponding to that resource manager are placed. This partition is given the same name as the resource manager.

**i** A partition can not span multiple resource managers. In addition to these resource manager partitions, a pseudo-partition named '[ALL]' is created that contains the aggregate resources of all partitions.

**i** While the resource manager partitions are real partitions containing resources not explicitly assigned to other partitions, the '[ALL]' partition is only a convenience object and is not a real partition; therefore it cannot be requested by jobs or included in configuration ACLs.

## 6.2.2 Defining Partitions

Node to partition mappings can be established directly using the `NODECFG` parameter or indirectly using the `FEATUREPARTITIONHEADER` parameter. If using direct mapping, this is accomplished as shown in the example that follows.

```
NODECFG [node001]    PARTITION=astronomy
NODECFG [node002]    PARTITION=astronomy
...
NODECFG [node049]    PARTITION=math
...
```

**i** By default, Moab creates two partitions, 'DEFAULT' and '[ALL]'. These are used internally, and consume spots in the 31-partition maximum defined in the `M_MAX_PAR` parameter. If more partitions are needed, you can adjust the maximum partition count. See [Adjusting Default Limits](#) for information on increasing the maximum number of partitions.

## 6.2.3 Managing Partition Access

Partition access can be constrained by credential ACLs and by limits based on job resource requirements.

### 6.2.3.A Credential Based Access

Determining who can use which partition is specified using the `*CFG` parameters (`USERCFG`, `GROUPCFG`, `ACCOUNTCFG`, `QOSCFG`, `CLASSCFG`, and `SYSCFG`). These parameters allow you to select a partition access list on a credential or system wide basis using the `PLIST` attribute. By default, the access associated with any given job is the logical OR of all partition access lists assigned to the job's credentials.

For example, assume a site with two partitions, `general`, and `test`. The site management would like everybody to use the `general` partition by default. However, one user, Steve, needs to perform the majority of his work on the `test` partition. Two special groups, `staff` and `management` will also need access to use the `test` partition from time to time but will perform most of their work in the `general` partition. The following example configuration enables the needed user and group access and defaults for this site:

```
SYSCFG[base]      PLIST=general:test
USERCFG[DEFAULT] PLIST=general
USERCFG[steve]   PLIST=general:test
GROUPCFG[staff]  PLIST=general:test
GROUPCFG[mgmt]   PLIST=general:test
```

While using a logical OR approach allows sites to add access to certain jobs, some sites prefer to work the other way around. In these cases, access is granted by default and certain credentials are then restricted from accessing various partitions. To use this model, a system partition list must be specified as in the following example:

```
SYSCFG[base]      PLIST=general,test&
USERCFG[demo]     PLIST=test&
GROUPCFG[staff]   PLIST=general&
```

In the preceding example, note the ampersand (&). This character, which can be located anywhere in the `PLIST` line, indicates that the specified partition list should be logically ANDed with other partition access lists. In this case, the configuration limits jobs from user `demo` to running in partition `test` and jobs from group `staff` to running in partition `general`. All other jobs are allowed to run in either partition.

**i** When using AND-based partition access lists, the base system access list must be specified with `SYSCFG`.

### 6.2.3.B Per Job Resource Limits

Access to partitions can be constrained based on the resources requested on a per job basis with limits on both minimum and maximum resources requested. All limits are specified using `PARCFG`. See [Usage Limits](#) for more information on the available limits.

```
PARCFG[pIII]      PARCFG[amd]      MAX.PROC=16
PARCFG[aix]      MAX.WCLIMIT=12:00:00 MIN.PROC=4
                  MIN.NODE=12
```

## 6.2.4 Requesting Partitions



Users can request to use any partition they have access to on a per job basis. This is accomplished using the resource manager extensions since most native batch systems do not support the partition concept. For example, on a Torque system, a job submitted by a

member of the group `staff` could request that the job run in the `test` partition by adding the line `-l partition=test` to the `qsub` command line. See the [resource manager extension overview](#) for more information on configuring and using resource manager extensions.

## 6.2.5 Per-Partition Settings

The following settings can be specified on a per-partition basis using the `PARCFG` parameter:

Setting	Description
<b>FSSCALINGFACTOR</b>	<p>Moab will multiple the actual fairshare usage by this value to get the calculated fairshare usage of a job. The actual fairshare usage is calculated based on the <a href="#">FSPOLICY</a> parameter.</p> <p>For an example, if <code>FSPOLICY</code> is set to <code>DEDICATEDPS</code> and a job runs on two processors for 100 seconds then the actual fairshare usage would be 200. If the job ran on a partition with <code>FSSCALINGFACTOR=.5</code> then Moab would multiply <math>200 * .5 = 100</math>. If the job ran on a partition with <code>FSSCALINGFACTOR=2</code> then Moab would multiply <math>200 * 2 = 400</math>.</p> <pre>PARCFG[par1] FSSCALINGFACTOR=&lt;double&gt;</pre>
<b>FSSECONDARYGROUPS</b>	Map unix groups to fairshare groups.
<b>GMETRIC</b>	<p>Specifies a <a href="#">generic metric</a> to apply to the partition. It is configured like a Moab parameter, with the gmetric name inside square brackets. Specify multiple gmetrics by separating each configuration with a space. For example:</p> <pre>PARCFG[par1] GMETRIC[GM1]=20 GMETRIC[GM2]=10</pre> <p><i>Partition <code>par1</code> has a <code>GM1</code> metric of 20 and a <code>GM2</code> metric of 10.</i></p>
<b>JOBNODEMATCHPOLICY</b>	Specifies the <a href="#">JOBNODEMATCHPOLICY</a> to be applied to jobs that run in the specified partition.
<b>NODEACCESSPOLICY</b>	Specifies the <a href="#">NODEACCESSPOLICY</a> to be applied to jobs that run in the specified partition.
<b>NODEALLOCATIONPOLICY</b>	Specifies the <a href="#">NODEALLOCATIONPOLICY</a> to be applied to jobs that run in the specified partition.

Setting	Description
<b>RESOURCELIMITMULTIPLIER</b>	<p>Specifies the <a href="#">RESOURCELIMITMULTIPLIER[&lt;PARID&gt;]</a> to be applied to jobs that run in the specified partition.</p> <p> This can only be viewed with 'showconfig -v'</p> <pre>PARCFG [A] RESOURCELIMITMULTIPLIER=PROC:1.1 RESOURCELIMITMULTIPLIER=MEM:2.0</pre>
<b>RESOURCELIMITPOLICY</b>	<p>Specifies the <a href="#">RESOURCELIMITPOLICY</a> to be applied to jobs that run in the specified partition.</p> <p> This can only be viewed with 'showconfig -v'</p> <pre>PARCFG [A] RESOURCELIMITPOLICY=WALLTIME:ALWAYS:CANCEL PARCFG [B] RESOURCELIMITPOLICY=WALLTIME:ALWAYS:REQUEUE</pre>

## 6.2.6 Miscellaneous Partition Issues

A brief caution: Use of partitions has been quite limited in recent years as other, more effective approaches are selected for site scheduling policies. Consequently, some aspects of partitions have received only minor testing. Still, note that partitions are fully supported and any problem found will be rectified.

### Related Topics

- [Standing Reservations](#)
- [Node Sets](#)
- [FEATUREPARTITIONHEADER](#) parameter
- [PARCFG](#) parameter

## 6.3 Quality of Service (QoS) Facilities

This section describes how to do the following:

- Allow key projects access to special services (such as preemption, resource dedication, and advance reservations).
- Provide access to special resources by requested QoS.

- Enable special treatment within priority and fairshare facilities by requested QoS.
- Provide exemptions to usage limits and other policies by requested QoS.
- Specify delivered service and response time targets.
- Enable job deadline guarantees.
- Control the list of QoSs available to each user and job.
- Enable special charging rates based on requested or delivered QoS levels.
- Enable limits on the extent of use for each defined QoS.
- Monitor current and historical usage for each defined QoS.

In this topic:

[6.3.1 QoS Overview - page 524](#)

[6.3.2 QoS Enabled Privileges - page 525](#)

[6.3.3 Managing QoS Access - page 533](#)

[6.3.4 Requesting QoS Services at Job Submission - page 534](#)

[6.3.5 Restricting Access to Special Attributes - page 534](#)

## 6.3.1 QoS Overview

Moab's QoS facility allows a site to give special treatment to various classes of jobs, users, groups, and so forth. Each QoS object can be thought of as a container of special privileges ranging from fairness policy exemptions, to special job prioritization, to special functionality access. Each QoS object also has an extensive access list of users, groups, and accounts that can access these privileges.

Sites can configure various QoSs each with its own set of priorities, policy exemptions, and special resource access settings. They can then configure user, group, account, and class access to these QoSs. A given job will have a default QoS and may have access to several additional QoSs. When the job is submitted, the submitter can request a specific QoS or just allow the default QoS to be used. Once a job is submitted, a user can adjust the QoS of the job at any time using the `setqos` command. The `setqos` command will only allow the user to modify the QoS of that user's jobs and only change the QoS to a QoS that this user has access to. Moab administrators can change the QoS of any job to any value.

Jobs can be granted access to QoS privileges if the QoS is listed in the system default configuration `QDEF` (QoS default) or `QLIST` (QoS access list), or if the QoS is specified in the



**QDEF** or **QLIST** of a **user**, **group**, **account**, or **class** associated with that job. Alternatively, a user can access QoS privileges if that user is listed in the QoSs **MEMBERULIST** attribute.

The **mdiag -q** command can be used to obtain information about the current QoS configuration including specified credential access.

## 6.3.2 QoS Enabled Privileges

The privileges enabled via QoS settings can be broken into the following categories:

- [Special Prioritization](#)
- [Service Access and Constraints](#)
- [Usage Limits and Overrides](#)
- [Service Access Thresholds](#)
- [Preemption Management](#)

All privileges are managed via the **QOSCFG** parameter.

### 6.3.2.A Special Prioritization

Attribute Name	Description
<b>FSTARGET</b>	Specifies QoS <b>fairshare</b> target.
<b>FSWEIGHT</b>	Sets QoS fairshare weight offset affecting a job's <b>fairshare</b> priority component.
<b>PRIORITY</b>	Assigns priority to all jobs requesting particular QoS.
<b>PROCWEIGHT</b>	Sets QoS PROCWEIGHT weight offset affecting a job's <b>resource</b> priority component.
<b>QTTARGET</b>	Sets QoS queuetime target affecting a job's <b>target</b> priority component and QoS delivered.
<b>QTWEIGHT</b>	Sets QoS queuetime weight offset affecting a job's <b>service</b> priority component.
<b>XFTARGET</b>	Sets QoS XFactor target affecting a job's <b>target</b> priority component and QoS delivered.

Attribute Name	Description
<b>XFWEIGHT</b>	Sets QoS XFactor weight offset affecting a job's <a href="#">service</a> priority component.

*Example 6-15:*

```
# assign priority for all qos geo jobs
QOSCFG[geo] PRIORITY=10000
```

### 6.3.2.B Service Access and Constraints

The QoS facility can be used to enable special services and to disable default services. These services are enabled/disabled by setting the QoS `QFLAGS` attribute.

Flag Name	Description
<b>DEADLINE</b>	Job can request an absolute or relative completion <a href="#">deadline</a> and Moab will reserve resources to meet that deadline. (An alternative priority based deadline behavior is discussed in the <a href="#">PRIORITY FACTORS</a> section.)
<b>DEDICATED</b>	Moab dedicates all resources of an allocated node to the job meaning that the job will not share a node's compute resources with any other job.
<b>ENABLEUSERRSV</b>	Allow user or personal <a href="#">reservations</a> to be created and managed.
<b>IGNALL</b>	Scheduler ignores all resource usage policies for jobs associated with this QoS.
<b>JOBPRIOACCRUALPOLICY</b>	Specifies how Moab should track the dynamic aspects of a job's priority. The two valid values are <code>ACCRUE</code> and <code>RESET</code> : <ul style="list-style-type: none"> <li><code>ACCRUE</code> indicates that the job will accrue queue time based priority from the time it is submitted unless it violates any of the policies not specified in <a href="#">JOBPRIOEXCEPTIONS</a>.</li> <li><code>RESET</code> indicates that it will accrue priority from the time it is submitted unless it violates any of the <a href="#">JOBPRIOEXCEPTIONS</a>. However, with <code>RESET</code>, if the job does violate <a href="#">JOBPRIOEXCEPTIONS</a> then its queue time based priority will be reset to 0.</li> </ul>

Flag Name	Description
	<div data-bbox="630 323 1409 474" style="border: 1px solid #ccc; padding: 5px; margin-bottom: 10px;"> <p><b>i</b> <b>JOBPRIOACCRUALPOLICY</b> is a global parameter, but can be configured to work only in QOSCFG:</p> <pre style="border: 1px dashed #ccc; padding: 2px; display: inline-block;">QOSCFG[arrays] JOBPRIOACCRUALPOLICY=ACCRUE</pre> </div> <p>The following old <b>JOBPRIOACCRUALPOLICY</b> values have been deprecated and should be adjusted to the following values:</p> <ul style="list-style-type: none"> <li>• <b>QUEUEPOLICY = ACCRUE</b> and <b>JOBPRIOEXCEPTIONS SOFTPOLICY,HARDPOLICY</b></li> <li>• <b>QUEUEPOLICYRESET = RESET</b> and <b>JOBPRIOEXCEPTIONS SOFTPOLICY,HARDPOLICY</b></li> <li>• <b>ALWAYS = ACCRUE</b> and <b>JOBPRIOEXCEPTIONS ALL</b></li> <li>• <b>FULLPOLICY = ACCRUE</b> and <b>JOBPRIOEXCEPTIONS NONE</b></li> <li>• <b>FULLPOLICYRESET = RESET</b> and <b>JOBPRIOEXCEPTIONS NONE</b></li> </ul>
<b>JOBPRIOEXCEPTIONS</b>	<p>Specifies exceptions for calculating a job's dynamic priority (<b>QUEUETIME</b>, <b>XFACTOR</b>, <b>TARGETQUEUETIME</b>). Valid values are a comma-delimited list of any of the following: <b>DEFER</b>, <b>DEPENDS</b>, <b>SOFTPOLICY</b>, <b>HARDPOLICY</b>, <b>IDLEPOLICY</b>, <b>USERHOLD</b>, <b>BATCHHOLD</b>, and <b>SYSTEMHOLD</b> (<b>ALL</b> or <b>NONE</b> can also be specified on their own).</p> <p>Normally, when a job violates a policy, is placed on hold, or has an unsatisfied dependency, it will not accrue priority. Exceptions can be configured to allow a job to accrue priority in spite of any of these violations. With <b>DEPENDS</b> a job will increase in priority even if there exists an unsatisfied dependency. With <b>SOFTPOLICY</b>, <b>HARDPOLICY</b>, or <b>IDLEPOLICY</b> a job can accrue priority despite violating a specific limit. With <b>DEFER</b>, <b>USERHOLD</b>, <b>BATCHHOLD</b>, or <b>SYSTEMHOLD</b> a job can accrue priority despite being on hold.</p> <div data-bbox="630 1411 1409 1562" style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p><b>i</b> <b>JOBPRIOEXCEPTIONS</b> is a global parameter, but can be configured to work only in QOSCFG:</p> <pre style="border: 1px dashed #ccc; padding: 2px; display: inline-block;">QOSCFG[arrays] JOBPRIOEXCEPTIONS=IDLEPOLICY</pre> </div>
<b>NOBF</b>	Job is not considered for backfill.
<b>NORESERVATION</b>	Job should never reserve resources regardless of priority.
<b>NTR</b>	Job is prioritized as next to run (NTR) and backfill is disabled to

Flag Name	Description
	<p>prevent other jobs from jumping in front of ones with the NTR flag.</p> <div style="border: 1px solid black; padding: 5px;"> <p><b>i</b> It is important to note that jobs marked with this flag should not be blocked. If they are, Moab will stop scheduling because if a job is marked with this flag, no other jobs will be run until the flagged NTR (Next to Run) job starts. Consider using the <code>PRIORITY</code> attribute of the <code>QOSCFG[&lt;QOSID&gt;]</code> parameter instead, when possible. Or, as you may encounter a scheduling delay for NTR-flagged jobs to start, consider using the <code>RESERVATIONDEPTH</code> and <code>RESERVATIONQOSLIST</code> parameters to provide better scheduling flow. See <a href="#">Reservation Policies</a> (especially the section on Assigning Per-QoS Reservation Creation Rules) for more information.</p> </div>
<b>PREEMPTCONFIG</b>	User jobs can specify options to alter how preemption impacts the job such as <a href="#">minpreempttime</a> .
<b>PREEMPTEE</b>	Job can be preempted by higher priority <code>PREEMPTOR</code> jobs.
<b>PREEMPTFSV</b>	Job can be preempted by higher priority <code>PREEMPTOR</code> jobs if it exceeds its <a href="#">fairshare</a> target when started.
<b>PREEMPTOR</b>	Job can preempt lower priority <code>PREEMPTEE</code> jobs.
<b>PREEMPTSPV</b>	Job can be preempted by higher priority <code>PREEMPTOR</code> jobs if it currently violates a <a href="#">soft usage policy</a> limit.
<b>PROVISION</b>	If the job cannot locate available resources with the needed OS or software, the scheduler can provision a number of nodes to meet the needed OS or software requirements.
<b>RESERVEALWAYS</b>	Job should create resource reservation regardless of job priority.
<b>RUNNOW</b>	<p>Boosts a job's system priority and makes the job a preemptor.</p> <div style="border: 1px solid black; padding: 5px;"> <p><b>i</b> <code>RUNNOW</code> overrides resource restrictions such as <code>MAXJOB</code> or <code>MAXPROC</code>.</p> </div>
<b>TRIGGER</b>	The job is able to directly specify triggers.

Flag Name	Description
<b>USERRESERVED[:&lt;RSVID&gt;]</b>	Job can only use resources within accessible reservations. If <RSVID> is specified, job can only use resources within the specified reservation.

*Example 6-16: For lowprio QoS job, disable backfill and make job preemptible*

```
QOSCFG[lowprio] QFLAGS=NOBF,PREEMPTEE
```

*Example 6-17: Bind all jobs to chemistry reservation*

```
QOSCFG[chem-b] QFLAGS=USERRESERVED:chemistry
```

## Other QoS Attributes

In addition to the flags, there are attributes that alter service access.

Attribute name	Description
<b>SYSPRIO</b>	<p>Sets the system priority on jobs associated with this QoS.</p> <p>Example: All jobs submitted under a QoS sample receive a system priority of 1.</p> <pre>QOSCFG[sample] SYSPRIO=1</pre> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> Once a system priority has been added to a job, either manually or through configuration, it can only be removed manually.</p> </div>
<b>REQUESTGEOMETRY</b>	<p>Defines the size that is requested when Elastic Computing occurs. Potential values are 'PRIORITYJOBSIZE' or '&lt;NODECOUNT&gt;@&lt;DURATION&gt;'. If PRIORITYJOBSIZE is set, then the nodecount and duration for Elastic Computing is set in realtime to whatever is the size of the highest priority idle job.</p> <p>Example:</p> <pre>QOSCFG[sample] REQUESTGEOMETRY=12@4:00:00:00</pre>

## Per QoS Required Reservations

If desired, jobs associated with a particular QoS can be locked into a reservation or reservation group using the `REQRID` attribute. For example, to force jobs using QoS `jasper` to only use the resources within the `failsafe` standing reservation, use the following:

```
QOSCFG[jasper] REQRID=failsafe
...
```

### 6.3.2.C Usage Limits and Overrides

All credentials, including QoS, allow specification of job usage limits as described in the [5.2.1.A Basic Fairness Policies - page 381](#) overview. In such cases, jobs are constrained by the most limiting of all applicable policies. With QoSs, an override limit can also be specified and with this limit, jobs are constrained by the override, regardless of other limits specified. Using override limits, you can create custom QoSs that allow for more jobs or more processors.

*Example 6-18:*

```
# staff QoS should have a limit of 48 jobs, ignoring the user limit
USERCFG[DEFAULT]    MAXJOB=10
QOSCFG[staff]       OMAXJOB=48
```

(See [5.2.2 Override Limits - page 393](#).)

The following parameters can override the throttling policies from other credentials:

Parameter Name	Description
OMAXJOB	Overrides a credential's limit on the number of jobs the credential can have active (starting or running) at any given time. Moab places a hold on all new jobs submitted by that credential once it has reached its maximum number of allowable jobs. Overrides a limit set using the <a href="#">MAXJOB</a> parameter).
OMAXNODE	Overrides a credential's limit on the total number of compute nodes that can be in use by active jobs at any given time. Overrides a limit set using the <a href="#">MAXNODE</a> parameter.
OMAXPE	Overrides a credential's limit on the total number of dedicated processor-equivalents the credential can have allocated by active jobs at any given time. Overrides a limit set using the <a href="#">MAXPE</a> parameter.
OMAXPROC	Overrides a credential's limit on the total number of dedicated processors the credential can have allocated by active jobs at any given time. Overrides a limit set using the <a href="#">MAXPROC</a> parameter.
OMAXPS	Overrides a credential's limit on the number of outstanding processor-seconds the credential can have allocated at any given time. Overrides a limit set using the <a href="#">MAXPS</a> parameter.

Parameter Name	Description
OMAXJPROC	Overrides a limit on the total number of dedicated processors that can be allocated to an active job at a given time. Overrides limits set using <code>msub -W x=MAXPROC</code> or the <code>CLASSCFG MAX.PROC</code> attribute.
OMAXJPS	Limits the number of outstanding processor-seconds allocated to a job at any given time. Overrides a limit set using the <code>CLASSCFG MAX.PS</code> attribute.
OMAXJWC	Overrides a maximum wallclock limit per job. Overrides a limit set using the <code>CLASSCFG MAX.WCLIMIT</code> attribute.
OMAXJNODE	Overrides a limit on the total number of compute nodes that can be in use by a job at any given time. Overrides a limit set using the <code>CLASSCFG MAX.NODE</code> attribute.

### 6.3.2.D Service Access Thresholds

Jobs can be granted access to services such as [preemption](#) and [reservation creation](#), and they can be granted access to resource reservations. However, with QoS thresholds, this access can be made conditional on the current queuetime and XFactor metrics of an idle job. The following table lists the available QoS service thresholds:

Threshold attribute	Description
<b>PREEMPTQTHRESHOLD</b>	A job with this QoS becomes a preemptor if the specified queuetime threshold is reached.
<b>PREEMPTXFTHRESHOLD</b>	A job with this QoS becomes a preemptor if the specified XFactor threshold is reached.
<b>RSVQTHRESHOLD</b>	A job with this QoS can create a <a href="#">job reservation</a> to guarantee resource access if the specified queuetime threshold is reached.
<b>RSVXFTHRESHOLD</b>	A job with this QoS can create a <a href="#">job reservation</a> to guarantee resource access if the specified XFactor threshold is reached.
<b>ACLQTHRESHOLD</b>	A job with this QoS can access reservations with a corresponding QoS ACL only if the specified queuetime threshold is reached.
<b>ACLXFTHRESHOLD</b>	A job with this QoS can access reservations with a corresponding

Threshold attribute	Description
	QoS ACL only if the specified XFactor threshold is reached.
<b>TRIGGERQTTHRESHOLD</b>	If a job with this QoS fails to run before this threshold is reached, any failure triggers associated with this QoS will fire.

### 6.3.2.E QoS Metrics

Metric Name	Description
<b>BACKLOGCOMPLETIONTIME</b>	<p>The estimated run-time to all idle jobs for a certain QoS. More specifically, it is the processor second count of all the idle jobs in the QoS, divided by the total processors on the system.</p> <pre> QOOSCFG[HIGH TRIGGER=EType=threshold,AType=exec,TType=elastic,threshold=BACKLOGCOMPLETIONTIME&gt;1,Action="\$HOME/geometry.pl blah",timeout=5:00 </pre> <p><b>i</b> In order to calculate the BacklogCompletionTime, the QoS must have <code>ENABLEPROFILING=TRUE</code>, either on the QoS itself or on the <code>DEFAULT</code> QoS.</p>

### 6.3.2.F Preemption Management

Job [preemption](#) facilities can be controlled on a per-QoS basis using the [PREEMPTTEE](#) and [PREEMPTOR](#) flags. Jobs that are preemptible can optionally be constrained to only be preempted in a particular manner by specifying the QoS `PREEMPTPOLICY` attribute as in the following example:

```
QOOSCFG[special] QFLAGS=PREEMPTTEE PREEMPTPOLICY=CHECKPOINT
```

For preemption to be effective, a job must be marked as a preemptee and must be enabled for the requested preemption type. For example, if the `PREEMPTPOLICY` is set to `suspend`, a potential target job must be both a preemptee and marked with the job flag `SUSPENDABLE`. See [suspension](#) for more information. If the target job is not suspendable, it will be either requeued or canceled. Likewise, if the `PREEMPTPOLICY` is set to `requeue`, the job will be requeued if it is marked `restartable`. Otherwise, it will be canceled.

The minimum time a job must run before being considered eligible for preemption can also be configured on a per-QoS basis using the `PREEMPTMINTIME` parameter, which is analogous to the [JOBPREEMPTMINACTIVETIME](#). Conversely, `PREEMPTMAXTIME` sets a



threshold for which a job is no longer eligible for preemption; see [JOBPREEMPTMAXACTIVETIME](#) for analogous details.

The `PREEMPTTEES` attribute allows you to specify which QoSs that a job in a specific QoS is allowed to preempt. The `PREEMPTTEES` list is a comma-delimited list of QoS IDs. When a `PREEMPTTEES` attribute is specified, a job using that QoS can only preempt jobs using QoSs listed in the `PREEMPTTEES` list. In turn, those QoSs must be flagged as `PREEMPTTEE` as in the following example:

```
QOSCFG[a] QFLAGS=PREEMPTOR PREEMPTTEES=b,c
QOSCFG[b] QFLAGS=PREEMPTTEE
QOSCFG[c] QFLAGS=PREEMPTTEE
```

*In the example, jobs in the 'a' QoS can only preempt jobs in the b and c QoSs.*

## 6.3.3 Managing QoS Access

### 6.3.3.A Specifying Credential Based QoS Access

You can define the privileges allowed within a QoS by using the `QOSCFG` parameter; however, in most cases access to the QoS is enabled via credential specific `*CFG` parameters, specifically the [USERCFG](#), [GROUPCFG](#), [ACCOUNTCFG](#), and [CLASSCFG](#) parameters, which allow defining QoS access lists and QoS defaults. Specify credential specific QoS access by using the `QLIST` and/or `QDEF` attributes of the associated credential parameter.

### 6.3.3.B QOS Access via Logical OR

To enable QoS access, the `QLIST` and/or `QDEF` attributes of the appropriate user, group, account, or class/queue should be specified as in the following example:

```
# user john's jobs can access QOS geo, chem, or staff with geo as default
USERCFG[john] QDEF=geo QLIST=geo,chem,staff
# group system jobs can access the development qos
GROUPCFG[systems] QDEF=development
# class batch jobs can access the normal qos
CLASSCFG[batch] QDEF=normal
```

By default, jobs can request a QoS if access to that QoS is allowed by any of the job's credentials. (In the previous example, a job from user `john` submitted to the class `batch` could request QoSs `geo`, `chem`, `staff`, or `normal`).

### 6.3.3.C QOS Access via Logical AND

If desired, QoS access can be masked or logically ANDed if the QoS access list is specified with a terminating ampersand (&) as in the following example:

```
# user john's jobs can access QoS geo, chem, or staff with geo as default
USERCFG[john]      QDEF=geo  QLIST=geo,chem,staff
# group system jobs can access the development qos
GROUPCFG[systems] QDEF=development
# class batch jobs can access the normal qos
CLASSCFG[batch]   QDEF=normal
# class debug jobs can only access the development or lowpri QoSs regardless of other
credentials
CLASSCFG[debug]   QLIST=development,lowpri&
```

## Specifying QoS Based Access

QoS access can also be specified from within the QoS object using the QoS `MEMBERLIST` attribute as in the following example:

```
# define qos premiere and grant access to users steve and john
QOSCFG[premiere]  PRIORITY=1000  QFLAGS=PREEMPTOR  MEMBERLIST=steve,john
```

**i** By default, if a job requests a QoS that it cannot access, Moab places a hold on that job. The `QOSREJECTPOLICY` can be used to modify this behavior.

### 6.3.4 Requesting QoS Services at Job Submission

By default, jobs inherit a default QoS based on the user, group, class, and account associated with the job. If a job has access to multiple QoS levels, the submitter can explicitly request a particular QoS using the QoS resource manager `extension` as in the following example:

```
> msub -l nodes=1,walltime=100,qos=special3 job.cmd
```

### 6.3.5 Restricting Access to Special Attributes

This feature is removed for Moab 9.0 and later. You can achieve the same results using job templates.

---

#### Related Topics

- [Credential Overview](#)
- [Allocation Management Overview](#)
- [Rollback Reservations](#)
- [Job Deadlines](#)
- [Using QoS preemption](#)

## Chapter 7: Optimizing Scheduling Behavior – Backfill and Node Sets

In this chapter:

7.1 Optimization Overview .....	535
7.2 Backfill .....	536
7.2.1 Backfill Overview .....	536
7.2.2 Backfill Algorithms .....	538
7.2.3 Configuring Backfill .....	540
7.3 Node Set Overview .....	543
7.3.1 Node Set Usage Overview .....	543
7.3.2 Node Set Configuration Examples .....	544
7.3.3 Requesting Node Sets for Job Submission .....	550
7.3.4 Configuring Node Sets for Classes .....	550

### 7.1 Optimization Overview

Moab optimizes cluster performance. Every policy, limit, and feature is designed to allow maximum scheduling flexibility while enforcing the required constraints. A driving responsibility of the scheduler is to do all in its power to maximize system use and to minimize job response time while honoring the policies that make up the site's mission goals.

However, as all jobs are not created equal, optimization must be abstracted slightly further to incorporate this fact. Cluster optimization must also focus on targeted cycle delivery. In the scientific HPC community, the true goal of a cluster is to maximize delivered research. For businesses and other organizations, the purposes may be slightly different, but all organizations agree on the simple tenet that the cluster should optimize the site's mission goals.

To obtain this goal, the scheduler has several levels of optimization it performs:

Level	Description
<b>Workload Ordering</b>	Prioritizing workload and utilizing backfill

Level	Description
<b>Intelligent Resource Allocation</b>	Selecting those resources that best meet the job's needs or best enable future jobs to run (see <a href="#">node allocation</a> )
<b>Maximizing Intra-Job Efficiency</b>	Selecting the type of nodes, collection of nodes, and proximity of nodes required to maximize job performance by minimizing both job compute and inter-process communication time (see <a href="#">node sets</a> and <a href="#">node allocation</a> )
<b>Job Preemption</b>	Preempting jobs to allow the most important jobs to receive the best response time (see <a href="#">preemption</a> )
<b>Utilizing Flexible Policies</b>	Using policies that minimize blocking and resource fragmentation while enforcing needed constraints (see <a href="#">soft throttling policies</a> and <a href="#">reservations</a> )

## 7.2 Backfill

In this topic:

- [7.2.1 Backfill Overview - page 536](#)
- [7.2.2 Backfill Algorithms - page 538](#)
- [7.2.3 Configuring Backfill - page 540](#)

### 7.2.1 Backfill Overview

Backfill is a scheduling optimization that allows a scheduler to make better use of available resources by running jobs out of order. When Moab schedules, it prioritizes the jobs in the queue according to a number of factors and then orders the jobs into a highest priority first (or priority FIFO) sorted list. It starts the jobs one by one stepping through the priority list until it reaches a job it cannot start. Because all jobs and reservations possess a start time and a wallclock limit, Moab can determine the completion time of all jobs in the queue. Consequently, Moab can also determine the earliest the needed resources will become available for the highest priority job to start.

Backfill operates based on this earliest job start information. Because Moab knows the earliest the highest priority job can start, and which resources it will need at that time, it can also determine which jobs can be started without delaying this job. Enabling backfill allows the scheduler to start other, lower-priority jobs so long as they do not delay the highest priority job. If backfill is enabled, Moab protects the highest priority job's start time by creating a job reservation to reserve the needed resources at the appropriate time. Moab then can start any job that will not interfere with this reservation.

Backfill offers significant scheduler performance improvement. In a typical large system, enabling backfill increases system utilization by about 20% and improves turnaround time by an even greater amount. Because of the way it works, essentially filling in holes in node space, backfill tends to favor smaller and shorter running jobs more than larger and longer running ones. It is common to see over 90% of these small and short jobs backfilled. Consequently, sites will see marked improvement in the level of service delivered to the small, short jobs and moderate to little improvement for the larger, long ones.

With most algorithms and policies, there is a trade-off. Backfill is not an exception but the negative effects are minor. Because backfill locates jobs to run from throughout the idle job queue, it tends to diminish the influence of the job prioritization a site has chosen and therefore may negate some desired workload steering attempts through this prioritization. Although by default the start time of the highest priority job is protected by a reservation, there is nothing to prevent the third priority job from starting early and possibly delaying the start of the second priority job. This issue is addressed along with its trade-offs [later](#) in this section.

Another problem is a little more subtle. Consider the following scenario involving a two-processor cluster. Job A has a four-hour wallclock limit and requires one processor. It started one hour ago (time zero) and will reach its wallclock limit in three more hours. Job B is the highest priority idle job and requires two processors for one hour. Job C is the next highest priority job and requires one processor for two hours. Moab examines the jobs and correctly determines that job A must finish in three hours and therefore, the earliest job B can start is in three hours. Moab also determines that job C can start and finish in less than this amount of time. Consequently, Moab starts job C on the idle processor at time one. One hour later (time two), job A completes early. Apparently, the user overestimated the amount of time job A would need by a few hours. Since job B is now the highest priority job, it should be able to run. However, job C, a lower priority job was started an hour ago and the resources needed for job B are not available. Moab re-evaluates job B's reservation and determines that it can slide forward an hour. At time three, job B starts.

In review, backfill provided positive benefits. Job A successfully ran to completion. Job C was started immediately. Job B was able to start one hour sooner than its original target time, although, had backfill not been enabled, job B would have been able to run two hours earlier.

The scenario just described occurs quite frequently because user estimates for job duration are generally inaccurate. Job wallclock estimate accuracy, or wallclock accuracy, is

defined as the ratio of wall time required to actually run the job divided by the wall time requested for the job. Wallclock accuracy varies from site to site but the site average is rarely better than 50%. Because the quality of the walltime estimate provided by the user is so low, job reservations for high priority jobs are often later than they need to be.

Although there do exist some minor drawbacks with backfill, its net performance impact on a site's workload is very positive. While a few of the highest priority jobs may get temporarily delayed, their position as highest priority was most likely accelerated by the fact that jobs in front of them were able to start earlier due to backfill. Studies have shown that only a very small number of jobs are truly delayed and when they are, it is only by a fraction of their total queue time. At the same time, many jobs are started significantly earlier than would have occurred without backfill.


## 7.2.2 Backfill Algorithms

**i** `BACKFILLPOLICY` controls which job gets selected first to be backfilled. Backfill jobs are still placed on nodes according to the `NODEALLOCATIONPOLICY`.

The algorithm behind Moab backfill scheduling is straightforward, although there are a number of issues and parameters that should be highlighted. First of all, Moab makes two backfill scheduling passes. For each pass, Moab selects a list of jobs that are eligible for backfill. On the first pass, only those jobs that meet the constraints of the soft fairness throttling policies are considered and scheduled. The second pass expands this list of jobs to include those that meet the hard (less constrained) fairness throttling policies.

The second important concept regarding Moab backfill is the concept of backfill windows. The figure below shows a simple batch environment containing two running jobs and a reservation for a third job. The present time is represented by the leftmost end of the box with the future moving to the right. The light gray boxes represent currently idle nodes that are eligible for backfill. For this example, let's assume that the space represented covers 8 nodes and a 3 hour time frame. To determine backfill windows, Moab analyzes the idle nodes essentially looking for largest node-time rectangles. It determines that there are two backfill windows. The first window, Window 1, consists of 4 nodes that are available for only one hour (because some of the nodes are blocked by the reservation for Job 3). The second window contains only one node but has no time limit because this node is not blocked by the reservation for Job 3. It is important to note that these backfill windows overlap.

Image 7-1: Backfillable nodes create backfill windows 1 and 2



Once the backfill windows have been determined, Moab begins to traverse them. The current behavior is to traverse these windows widest window first (most nodes to fewest

nodes). As each backfill window is evaluated, Moab applies the backfill algorithm specified by the `BACKFILLPOLICY` parameter.

If the `FIRSTFIT` algorithm is applied, the following steps are taken:

1. The list of feasible backfill jobs is filtered, selecting only those that will actually fit in the current backfill window.
2. The first job is started.
3. While backfill jobs and idle resources remain, repeat step 1.

If the `BESTFIT` algorithm is applied, the following steps are taken:

1. The list of feasible backfill jobs is filtered, selecting only those that actually fit in the current backfill window.
2. The degree of fit of each job is determined based on the `BACKFILLMETRIC` parameter (processors, seconds, processor-seconds).
3. The job with the best fit starts.
4. While backfill jobs and idle resources remain, repeat step 1.

If `NONE` is set, the backfill policy is disabled.

Other backfill policies behave in a generally similar manner. The [parameters](#) documentation provides further details.

### 7.2.2.A Liberal Versus Conservative Backfill

By default, Moab reserves only the highest priority job resulting in a liberal and aggressive backfill. This reservation guarantees that backfilled jobs will not delay the highest priority job, although they might delay other jobs. The parameter `RESERVATIONDEPTH` controls how conservative or liberal the backfill policy is. This parameter controls how deep down the queue priority reservations will be made. While increasing this parameter improves guarantees that priority jobs will not be bypassed, it reduces the freedom of the scheduler to backfill resulting in somewhat lower system utilization. The significance of the trade-offs should be evaluated on a site by site basis.

## 7.2.3 Configuring Backfill

### 7.2.3.A Backfill Policies

Backfill is enabled in Moab by specifying the `BACKFILLPOLICY` parameter. The `BACKFILLPOLICY` parameter is used to control which job gets selected first to be backfilled. Once the job has been selected, it is still placed on nodes according to the



`NODEALLOCATIONPOLICY` you have defined. By default, backfill is enabled in Moab using the `FIRSTFIT` algorithm. However, this parameter can also be set to `NONE` (disabled).

The number of reservations that protect the resources required by priority jobs can be controlled using `RESERVATIONDEPTH`. This depth can be distributed across job QoS levels using `RESERVATIONQOSLIST`.

### 7.2.3.B Backfill Chunking

In a batch environment saturated with serial jobs, serial jobs will, over time, dominate the resources available for backfill at the expense of other jobs. This is due to the time-dimension fragmentation associated with running serial jobs. For example, given an environment with an abundance of serial jobs, if a multi-processor job completes freeing processors, one of three things will happen:

1. The freed resources are allocated to another job requiring the same number of processors.
2. Additional jobs may complete at the same time allowing a larger job to allocate the aggregate resources.
3. The freed resources are allocated to one or more smaller jobs.

In environments where the scheduling iteration is much higher than the average time between completing jobs, case 3 occurs far more often than case 2, leading to smaller and smaller jobs populating the system over time.

To address this issue, the scheduler incorporates the concept of chunking. Chunking allows the scheduler to favor case 2 maintaining a more controlled balance between large and small jobs. The idea of chunking involves establishing a time-based threshold during which resources available for backfill are aggregated. This threshold is set using the parameter `BFCHUNKDURATION`. When resources are freed, they are made available only to jobs of a certain size (set using the parameter `BFCHUNKSIZE`) or larger. These resources remain protected from smaller jobs until either additional resources are freed up and a larger job can use the aggregate resources, or until the `BFCHUNKDURATION` threshold time expires.

**i** Backfill chunking is only activated when a job of size `BFCHUNKSIZE` or larger is blocked in backfill due to lack of resources.

It is important to note that the optimal settings for these parameters is very site-specific and will depend on the workload (including the average job turnaround time, job size, and mix of large to small jobs), cluster resources, and other scheduling environmental factors. Setting too restrictive values needlessly reduces utilization while settings that are too relaxed do not allowed the desired aggregation to occur.

**i** Backfill chunking is only enabled in conjunction with the `FIRSTFIT` backfill policy.

The current implementation of backfill chunking in Moab behaves as follows:

- When Moab starts, if backfill chunking is enabled, it starts immediately in a chunking window (with duration dictated by `BFChunkDuration`).
- Each time a job completes, the chunking window is re-enabled and extended to the chunking duration (current time + `BFChunkDuration`).
- While within this potentially continuously extending chunking window, the first backfill job, whether large or small, will always be evaluated without hindrance. Jobs will continue to be evaluated with the behavior that the first large job (based on `BFChunkSize`) to be encountered (including the first) will prevent subsequent small jobs from being backfilled.
- If a timeframe occurs where a job has not completed within `BFChunkDuration` of the previous job completion, then all job sizes can be freely backfilled.
- Backfill chunking does not kick in again until `BFChunkDuration` past the next job completion.

### 7.2.3.C Virtual Wallclock Time Scaling

In most environments, users submit jobs with rough estimations of the wallclock times. Within the HPC industry, a job typically runs for 40% of its specified wallclock time. Virtual Wallclock Time Scaling takes advantage of this fact to implement a form of optimistic backfilling. Jobs that are eligible for backfilling and not restricted by other policies are virtually scaled by the `BFVIRTUALWALLTIMESCALINGFACTOR` (assuming that the jobs finish before this new virtual wallclock limit). The scaled jobs are then compared to backfill windows to see if there is space and time for them to be scheduled. The scaled jobs are only scheduled if there is no possibility that it will conflict with a standing or administrator reservation. Conflicts with such reservations occur if the virtual wallclock time overlaps a reservation, or if the original non-virtual wallclock time overlaps a standing or administrator reservation. Jobs that can fit into an available backfill window without having their walltime scaled are backfilled 'as-is' (meaning, without virtually scaling the original walltime).

**i** Virtual Wallclock Time Scaling is only enabled when the `BFVIRTUALWALLTIMESCALINGFACTOR` parameter is defined.

If a virtually-scaled job fits into a window, and is backfilled, it will run until completion or until it comes within one scheduling iteration (`RMPOLLINTERVAL` defines the exact time of an iteration) of the virtual wallclock time expiration. In the latter case the job's wallclock time is restored to its original time and Moab checks and resolves conflicts caused by this

'expansion.' Conflicts may occur when the backfilled job is restored to its full duration resulting in reservation overlap. The [BFVIRTUALWALLTIMECONFLICTPOLICY](#) parameter controls how Moab handles these conflicts.

If the [BFVIRTUALWALLTIMECONFLICTPOLICY](#) parameter is set to NONE or is not specified, the overlapped job reservations are rescheduled.

---

### Related Topics

- [BACKFILLDEPTH](#) Parameter
- [BACKFILLPOLICY](#) Parameter
- [BFMINVIRTUALWALLTIME](#)
- [Reservation Policy Overview](#)

## 7.3 Node Set Overview

In this topic:

[7.3.1 Node Set Usage Overview - page 543](#)

[7.3.2 Node Set Configuration Examples - page 544](#)

[7.3.3 Requesting Node Sets for Job Submission - page 550](#)

[7.3.4 Configuring Node Sets for Classes - page 550](#)

### 7.3.1 Node Set Usage Overview

While backfill improves the scheduler's performance, this is only half the battle. The efficiency of a cluster, in terms of actual work accomplished, is a function of both scheduling performance and individual job efficiency. In many clusters, job efficiency can vary from node to node, as well as with the node mix allocated. Most parallel jobs written in popular languages such as MPI or PVM do not internally load balance their workload and therefore run only as fast as the slowest node allocated. Consequently, these jobs run most effectively on homogeneous sets of nodes. However, while many clusters start out as homogeneous, they quickly evolve as new generations of compute nodes are integrated into the system. Research has shown that this integration, while improving scheduling performance due to increased scheduler selection, can actually decrease average job efficiency.

A feature called node sets allows jobs to request sets of common resources without specifying exactly what resources are required. Node set policy can be specified globally or

on a per-job basis. In addition to their use in forcing jobs onto homogeneous nodes, these policies can also be used to guide jobs to one or more types of nodes on which a particular job performs best, similar to job preferences available in other systems. For example, an I/O intensive job might run best on a certain range of processor speeds, running slower on slower nodes, while wasting cycles on faster nodes. A job can specify `ANYOF:FEATURE:bigmem,fastos` to request nodes with the `bigmem` or `fastos` feature. Alternatively, if a simple feature-homogeneous node set is desired, `ONEOF:FEATURE` can be specified. On the other hand, a job may request a feature based node set with the configuration `ONEOF:FEATURE:bigmem,fastos`, in which case Moab will first attempt to locate adequate nodes where all nodes contain the `bigmem` feature. If such a set cannot be found, Moab will look for sets of nodes containing the other specified features. In highly heterogeneous clusters, the use of node sets improves job throughput by 10 to 15%.

Node sets can be requested on a system wide or per job basis. System wide configuration is accomplished via the `NODESET*` parameters while per job specification occurs via the [resource manager extensions](#).

**i** The GLOBAL node is included in all feature node sets.

When creating node sets, you have the option of using a fixed configuration or of creating node sets dynamically (by using the `msub` command). This topic explains how to set up both node set use cases.

## 7.3.2 Node Set Configuration Examples

Global node sets are defined using the [NODESETPOLICY](#), [NODESETATTRIBUTE](#), [NODESETLIST](#), and [NODESETISOPTIONAL](#) parameters. As stated before, you can create node sets dynamically (see [Dynamic Example](#)) or with a fixed configuration (see [Fixed Configuration Example](#)). The use of these parameters can be best highlighted with two examples.

### 7.3.2.A Fixed Configuration Example

In this example, a large site possesses a Myrinet based interconnect and wants to, whenever possible, allocate nodes within Myrinet switch boundaries. To accomplish this, they could assign node attributes to each node indicating which switch it was associated with (`switchA`, `switchB`, and so forth) and then use the following system wide node set configuration:

```
NODESETPOLICY      ONEOF
NODESETATTRIBUTE   FEATURE
NODESETISOPTIONAL TRUE
```

```
NODESETLIST      switchA,switchB,switchC,switchD
...
```

## Node Set Policy

In the preceding example, the `NODESETPOLICY` parameter is set to the policy `ONEOF` and tells Moab to allocate nodes within a single attribute set. Other node set policies are listed in the following table:

Policy	Description
<b>ANYOF</b>	Select resources from all sets contained in node set list. The job could span multiple node sets.
<b>FIRSTOF</b>	Select resources from first set to match specified constraints.
<b>ONEOF</b>	Select a single set that contains adequate resources to support job.

## Node Set Attribute

The example's `NODESETATTRIBUTE` parameter is set to `FEATURE`, specifying that the node sets are to be constructed along node feature boundaries.

You could also set the `NODESETATTRIBUTE` to `VARATTR`, specifying that node sets are to be constructed according to `VARATTR` values on the job.

You could also set the `NODESETATTRIBUTE` to `RESERVATION`, specifying that node sets are to be constructed according to the reservations (or reservation groups) specified in the `NODESETLIST` parameter.

## Node Set Constraint Handling

The next parameter, `NODESETISOPTIONAL`, indicates that Moab should not delay the start time of a job if the desired node set is not available but adequate idle resources exist outside of the set. Setting this parameter to `TRUE` basically tells Moab to attempt to use a node set if it is available, but if not, run the job as soon as possible anyway.

**i** Setting `NODESETISOPTIONAL` to `FALSE` will force the job to always run in a complete nodeset regardless of any start delay this imposes.

## Node Set List

Finally, the `NODESETLIST` value of `switchA switchB...` tells Moab to only use node sets based on the listed feature values. This is necessary since sites will often use node features for many purposes and the resulting node sets would be of little use for switch

proximity if they were generated based on irrelevant node features indicating things such as processor speed or node architecture.

To add nodes to the `NODESETLIST`, you must configure features on your nodes using the `NODECFG FEATURES` attribute.

```
NODECFG[node01] FEATURES=switchA
NODECFG[node02] FEATURES=switchA
NODECFG[node03] FEATURES=switchB
```

*Nodes node01 and node02 contain the switchA feature, and node node03 contains the switchB feature.*

## Node Set Priority

When resources are available in more than one resource set, the `NODESETPRIORITYTYPE` parameter allows control over how the best resource set is selected. Legal values for this parameter are described in the following table:

Priority Type	Description	Details
<b>AFFINITY</b>	Avoid a resource set with negative <a href="#">affinity</a> .	Choosing this type causes Moab to select a node set with no negative affinity nodes (nodes that have a reservation that with negative affinity). If all node sets have negative affinity, then Moab will select the first matching node set.
<b>BESTFIT</b>	Select the smallest resource set possible.	Choosing this type causes Moab, when selecting a node set, to eliminate sets that do not have all the required resources. From the remaining sets, Moab chooses the set with the least amount of resources. This priority type most closely matches the job requirements in order to waste the least amount of resources.  This type minimizes fragmentation of larger resource sets.
<b>FIRSTFIT</b>	Select the first set with enough resources.	Moab will select the first nodeset with enough resources to satisfy the job. This is the fastest of the priority types.
<b>MINLOSS</b>	Select the resource set that results in the minimal wasted resources assuming no internal job load balancing is available. (Assumes parallel jobs only run as fast as the slowest	Choosing this type works only when using the following configuration: <code>NODESETATTRIBUTE FEATURE</code> In a <code>SHAREDMEM</code> environment (see <a href="#">Moab-NUMA-Support Integration Guide</a> for more information), Moab will select the node set based on NUMA properties (the smallest

Priority Type	Description	Details
	allocated node.)	feasible node set).
<b>WORSTFIT</b>	Select the largest resource set possible.	<p>This type causes Moab, when choosing a node set, to eliminate sets that do not have all the required resources. From the remaining sets, Moab chooses the set with the greatest amount of resources.</p> <p>This type minimizes fragmentation of smaller resource sets, but increases fragmentation of larger resource sets.</p>

### 7.3.2.B Dynamic Example

In this example, a site wants to be able to dynamically specify which `VARATTR` values the node set will be based on. To accomplish this, the following configuration in the `moab.cfg` file could be used:

```

NODESETISOPTIONAL FALSE
NODESETPOLICY      FIRSTOF
NODESETATTRIBUTE   VARATTR

```

#### Node Set Attribute

The example's `NODESETATTRIBUTE` parameter is set to `FEATURE`, specifying that the node sets are to be constructed along node feature boundaries.

You could also set the `NODESETATTRIBUTE` to `VARATTR`, specifying that node sets are to be constructed according to `VARATTR` values on the job.

You could also set the `NODESETATTRIBUTE` to `RESERVATION`, specifying that node sets are to be constructed according to the reservations (or reservation groups) specified in the `NODESETLIST` parameter.

#### Node Set Policy

In the preceding example, the `NODESETPOLICY` parameter is set to the policy `FIRSTOF` and tells Moab to allocate nodes from the first set that matches specified constraints.

#### Node Set Constraint Handling

The parameter, `NODESETISOPTIONAL`, indicates that Moab should not delay the start time of a job if the desired node set is not available but adequate idle resources exist outside of

the set. Setting this parameter to `FALSE` will force the job to always run in a complete node set regardless of any start delay this imposes.

## msub Example

With the configuration (above) set in the `moab.cfg`, Moab is configured for dynamic node sets. You can create node sets dynamically by using the `msub -l` command. For more information, see [Resource Manager Extensions](#). Use the following format:

```
msub -l nodeset=FIRSTOF:VARATTR:<var> [=<value>], ...
```

For example, if you wanted to create a dynamic node set for the Provo datacenter:

```
msub -l nodeset=FIRSTOF:VARATTR: datacenter=Provo
```

*This command causes Moab to set `datacenter=Provo` as the node set.*

**i** You can specify more than one VARATTR in the command. For example, if you want to create a dynamic node set for the Provo datacenter and the SaltLake datacenter:

```
msub -l nodeset=FIRSTOF:VARATTR: datacenter=Provo: datacenter=SaltLake
```

If you specify only `datacenter` (without specifying a value, such as `=Provo`), Moab will look up all possible values (values reported on the node for that VARATTR), and then choose one. So if, for example, you have nodes that have VARATTRs `datacenter=Provo`, `datacenter=SaltLake`, and `datacenter=StGeorge`, then specifying `msub -l nodeset=FIRSTOF:VARATTR:datacenter` will cause the job to run in Provo *or* SaltLake *or* StGeorge.

You should also note that Moab also adds the VARATTR (whether you specify it or if Moab chooses it) to the required attribute (REQATTR) of the job. For example, if you specify `datacenter=Provo` as the VARATTR, `datacenter=Provo` will also be added to the job REQATTR. Likewise, if you specify only `datacenter`, and Moab chooses `datacenter=SaltLake`, then `datacenter=SaltLake` will be added to the job REQATTR.

If you do not request a VARATTR in the nodeset of the `msub -l` command, the job will run as if it did not use node sets at all, and nothing will be added to its REQATTR.

**i** If you manually specify a different REQATTR on a job (for example, `datacenter=SaltLake`) from the node set VARATTR (for example, `datacenter=Provo`), the job will never run.



### 7.3.2.C NODESETPLUS

Moab supports additional NodeSet behavior by specifying the **NODESETPLUS** parameter. Possible values when specifying this parameter are `SPANEVENLY` and `DELAY`.

**i** Neither `SPANEVENLY` nor `DELAY` will work with multi-req jobs or preemption.


Value	Description
<b>SPANEVENLY</b>	Moab attempts to fit all jobs within one node set, or it spans any number of node sets evenly. When a job specifies a <b>NODESETDELAY</b> , Moab attempts to contain the job within a single node set; if unable to do so, it spans node sets evenly, unless doing so would delay the job beyond the requested <code>NODESETDELAY</code> .
<b>DELAY</b>	Moab attempts to schedule the job within a nodeset for the configured <code>NODESETDELAY</code> . If Moab cannot find space for the job to start within <code>NODESETDELAY</code> (Moab considers future workload to determine if space will open up in time and might create a future reservation), then Moab schedules the job and ignores the nodeset requirement.

### 7.3.2.D Nested Node Sets

Moab attempts to fit jobs on node sets in the order they are specified in the `NODESETLIST`. You can create nested node sets by listing your node sets in a specific order. Here is an example of a 'smallest to largest' nested node set:

```
NODESETPOLICY ONEOF
NODESETATTRIBUTE FEATURE
NODESETISOPTIONAL FALSE
NODESETLIST
blade1a,blade1b,blade2a,blade2b,blade3a,blade3b,blade4a,blade4b,quad1a,quad1b,quad2a,q
uad2b,octet1,octet2,sixteen
```

The accompanying cluster would look like this:

*Image 7-2: Octet, quad, and blade node sets on a cluster*

In this example, Moab tries to fit the job on the nodes in the blade sets first. If that doesn't work, it moves up to the nodes in the quad sets (a set of four blade sets). If the quads are insufficient, it tries the nodes in the octet sets (a set of four quad node sets).

### 7.3.3 Requesting Node Sets for Job Submission

On a per job basis, each user can specify the equivalent of all parameters except `NODESETDELAY`. As mentioned previously, this is accomplished using the [resource manager extensions](#).

### 7.3.4 Configuring Node Sets for Classes

Classes can be configured with a default node set. In the configuration file, specify `DEFAULT.NODESET` with the following syntax:

```
DEFAULT.NODESET=<SETTYPE>:<SETATTR>[:<SETLIST>[,<SETLIST>]...].
```

For example, in a heterogeneous cluster with two different types of processors, the following configuration confines jobs assigned to the `amd` class to run on either `ATHLON` or `OPTERON` processors:

```
CLASSCFG[amd] DEFAULT.NODESET=ONEOF:FEATURE:ATHLON,OPTERON
...
```

---

### Related Topics

- [Resource Manager Extensions](#)
- [CLASSCFG](#)
- [Partition Overview](#)

## Chapter 8: Evaluating System Performance - Statistics, Profiling and Testing

In this chapter:

8.1 Moab Performance Evaluation Overview .....	552
8.2 Accounting: Job and System Statistics .....	552
8.2.1 Accounting Overview .....	553
8.2.2 Real-Time Statistics .....	554
8.3 Testing New Versions and Configurations .....	555
8.3.1 MONITOR Mode .....	555
8.3.2 INTERACTIVE Mode .....	556

### 8.1 Moab Performance Evaluation Overview

Moab Workload Manager tracks numerous performance statistics for jobs, accounting, users, groups, accounts, classes, QoS, the system, and so forth. These statistics can be accessed through various commands or Moab Cluster Manager/Monitor.

### 8.2 Accounting: Job and System Statistics

Moab provides extensive accounting facilities that allow resource usage to be tracked by resources (compute nodes), jobs, users, and other objects. The accounting facilities can be used in conjunction with, and correlated with, the accounting records provided by the resource and accounting manager.

Moab maintains both raw persistent data and a large number of processed in memory statistics allowing instant summaries of cycle delivery and system utilization. With this information, Moab can assist in accomplishing any of the following tasks:

- Determining cumulative cluster performance over a fixed time frame.
- Graphing changes in cluster utilization and responsiveness over time.
- Identifying which compute resources are most heavily used.
- Charting resource usage distribution among users, groups, projects, and classes.

- Determining allocated resources, responsiveness, and failure conditions for jobs completed in the past.
- Providing real-time statistics updates to external accounting systems.

This topic describes how to accomplish each of these tasks using Moab tools and accounting information.

In this topic:

[8.2.1 Accounting Overview - page 553](#)

[8.2.2 Real-Time Statistics - page 554](#)

## 8.2.1 Accounting Overview

Moab provides accounting data correlated to most major objects used within the cluster scheduling environment. These records provide job and reservation accounting, resource accounting, and credential-based accounting.

### 8.2.1.A Job and Reservation Accounting

As each job or reservation completes, Moab creates a complete persistent trace record containing information about who ran, the time frame of all significant events, and what resources were allocated. In addition, actual execution environment, failure reports, requested service levels, and other pieces of key information are also recorded. A complete description of each accounting data field can be found within section [Workload Traces](#).

### 8.2.1.B Resource Accounting

The load on any given node is available historically allowing identification of not only its usage at any point in time, but the actual jobs that were running on it. Moab Cluster Manager can show load information (assuming load is configured as a generic metric), but not the individual jobs that were running on a node at some point in the past. For aggregated, historical statistics covering node usage and availability, the `showstats` command can be run with the `-n` flag.

### 8.2.1.C Credential Accounting

Current and historical usage for users, groups, account, QoSs, and classes are determined in a manner similar to that available for evaluating nodes. For aggregated, historical statistics covering credential usage and availability, the `showstats` command can be run with the corresponding credential flag.

If needed, detailed credential accounting can also be enabled globally or on a credential by credential basis. With detailed credential accounting enabled, real-time information regarding per-credential usage over time can be displayed. To enable detailed per credential accounting, the `ENABLEPROFILING` attribute must be specified for credentials that are to be monitored. For example, to track detailed credentials, the following should be used:

```
USERCFG [DEFAULT]      ENABLEPROFILING=TRUE
QOSCFG [DEFAULT]      ENABLEPROFILING=TRUE
CLASSCFG [DEFAULT]    ENABLEPROFILING=TRUE
GROUPCFG [DEFAULT]    ENABLEPROFILING=TRUE
ACCOUNTCFG [DEFAULT]  ENABLEPROFILING=TRUE
```

Credential level profiling operates by maintaining a number of time-based statistical records for each credential. The parameters `PROFILECOUNT` and `PROFILEDURATION` control the number and duration of the statistical records.

## 8.2.2 Real-Time Statistics

Moab provides real-time statistical information about how the machine is running from a scheduling point of view. The `showstats` command is actually a suite of commands providing detailed information on an overall scheduling basis, as well as a per user, group, account and node basis. This command gets its information from in memory statistics that are loaded at scheduler start time from the scheduler checkpoint file. See [Checkpoint/Restart](#) for more information. This checkpoint file is updated periodically and when the scheduler is shut down allowing statistics to be collected over an extended time frame. At any time, real-time statistics can be reset using the `mschedctl -f` command.

In addition to the `showstats` command, the `showstats -f` command also obtains its information from the in memory statistics and checkpoint file. This command displays a processor-time based matrix of scheduling performance for a wide variety of metrics. Information such as backfill effectiveness or average job queue time can be determined on a job size/duration basis.

### 8.2.2.A FairShare Usage Statistics

Regardless of whether fairshare is enabled, detailed credential based fairshare statistics are maintained. Like job traces, these statistics are stored in the directory pointed to by the `STATDIR` parameter. Fairshare stats are maintained in a separate statistics file using the format `FS.<EPOCHTIME>` (`FS.982713600`, for example) with one file created per fairshare window. See the [Fairshare Overview](#) for more information. These files are also flat text and record credential based usage statistics. Information from these files can be seen via the `mdiag -f` command.

---

## Related Topics

- [Generic Consumable Resources](#)
- [Object Variables](#)

## 8.3 Testing New Versions and Configurations

### 8.3.1 MONITOR Mode

Moab supports a scheduling mode called `MONITOR`. In this mode, the scheduler initializes, contacts the resource manager and other peer services, and conducts scheduling cycles exactly as it would if running in `NORMAL` or production mode. Jobs are prioritized, reservations created, policies and limits enforced, and administrator and end-user commands enabled. The key difference is that although live resource management information is loaded, `MONITOR` mode disables Moab's ability to start, preempt, cancel, or otherwise modify jobs or resources. Moab continues to attempt to schedule exactly as it would in `NORMAL` mode but its ability to actually impact the system is disabled. Using this mode, a site can quickly verify correct resource manager configuration and scheduler operation. This mode can also be used to validate new policies and constraints. In fact, Moab can be run in `MONITOR` mode on a production system while another scheduler or even another version of Moab is running on the same system. This unique ability can allow new versions and configurations to be fully tested without any exposure to potential failures and with no cluster downtime.

To run Moab in `MONITOR` mode, simply set the `MODE` attribute of the `SCHEDCFG` parameter to `MONITOR` and start Moab. Normal scheduler commands can be used to evaluate configuration and performance. [Diagnostic commands](#) can be used to look for any potential issues. Further, the Moab log file can be used to determine which jobs Moab attempted to start, and which resources Moab attempted to allocate.

If another instance of Moab is running in production and a site administrator wants to evaluate an alternate configuration or new version, this is easily done but care should be taken to avoid conflicts with the primary scheduler. Potential conflicts include statistics files, logs, checkpoint files, and user interface ports. One of the easiest ways to avoid these conflicts is to create a new test directory with its own log and stats subdirectories. The new `moab.cfg` file can be created from scratch or based on the existing `moab.cfg` file already in use. In either case, make certain that the `PORT` attribute of the `SCHEDCFG` parameter differs from that used by the production scheduler by at least two ports. If testing with the production binary executable, the `MOABHOMEDIR` environment variable should be set to point to the new test directory to prevent Moab from loading the production `moab.cfg` file.

## 8.3.2 INTERACTIVE Mode

INTERACTIVE mode allows for evaluation of new versions and configurations in a manner different from MONITOR mode. Instead of disabling all resource and job control functions, Moab sends the desired change request to the screen and asks for permission to complete it. For example, before starting a job, Moab might post something like the following to the screen:

```
Command:  start job 1139.ncsa.edu on node list test013,test017,test018,test021
Accept:   (y/n) [default: n]?
```

The administrator must specifically accept each command request after verifying it correctly meets desired site policies. Moab then executes the specified command. This mode is highly useful in validating scheduler behavior and can be used until configuration is appropriately tuned and all parties are comfortable with the scheduler's performance. In most cases, sites will want to set the scheduling mode to NORMAL after verifying correct behavior.

---

### Related Topics

- [Testing New Releases and Policies](#)
- [Side-by-Side Mode](#)



## Chapter 9: General Job Administration

In this chapter:

9.1	Job Holds .....	558
9.1.1	Holds and Deferred Jobs .....	558
9.1.2	User Holds .....	558
9.1.3	System Holds .....	558
9.1.4	Batch Holds .....	559
9.1.5	Job Defer .....	559
9.2	Job Priority Management .....	560
9.3	Suspend/Resume Handling .....	560
9.4	Checkpoint/Restart Facilities .....	561
9.5	Job Dependencies .....	562
9.5.1	Basic Job Dependency Support .....	562
9.5.2	Job Dependency Syntax .....	563
9.6	Job Defaults and Per Job Limits .....	564
9.6.1	Job Defaults .....	564
9.6.2	Per Job Maximum Limits .....	565
9.6.3	Per Job Minimum Limits .....	565
9.7	General Job Policies .....	566
9.7.1	Multi-Node Support .....	566
9.7.2	Multi-Req Support .....	566
9.7.3	Malleable Job Support .....	567
9.7.4	Enabling Job User Proxy .....	567
9.8	Using a Local Queue .....	568
9.9	Job Deadlines .....	571
9.9.1	Deadline Overview .....	571
9.9.2	Setting Job Deadlines via QoS .....	571
9.9.3	Job Termination Date .....	573
9.9.4	Conflict Policies .....	573
9.10	Job Arrays .....	574
9.10.1	Job Array Overview .....	574
9.10.2	Enabling Job Arrays .....	575
9.10.3	Subjob Definitions .....	575

9.10.4 Using Environment Variables to Specify Array Index Values .....	575
9.10.5 Job Array Cancellation Policies .....	578
9.10.6 Minimizing the Impact of Very Large Job Arrays .....	579
9.10.7 Examples .....	581

## 9.1 Job Holds

### 9.1.1 Holds and Deferred Jobs

Moab supports job holds applied by users ([user holds](#)), administrators ([system holds](#)), and resource managers ([batch holds](#)). There is also a temporary hold known as a [job defer](#).

### 9.1.2 User Holds

User holds are very straightforward. Many, if not most, resource managers provide interfaces by which users can place a hold on their own job that tells the scheduler not to run the job while the hold is in place. Users can use this capability because the job's data is not yet ready, or they want to be present when the job runs to monitor results. Such user holds are created by, and under the control of a non-privileged user and can be removed at any time by that user. As would be expected, users can only place holds on their jobs. Jobs with a user hold in place will have a Moab state of `Hold` or `UserHold` depending on the resource manager being used.

### 9.1.3 System Holds

The system hold is put in place by a system administrator either manually or by way of an automated tool. As with all holds, the job is not allowed to run so long as this hold is in place. A batch administrator can place and release system holds on any job regardless of job ownership. However, unlike a user hold, normal users cannot release a system hold even on their own jobs. System holds are often used during system maintenance and to prevent particular jobs from running in accordance with current system needs. Jobs with a system hold in place will have a Moab state of `Hold` or `SystemHold` depending on the resource manager being used.

## 9.1.4 Batch Holds

Batch holds are placed on a job by the scheduler itself when it determines that a job cannot run. The reasons for this vary but can be displayed by issuing the `checkjob<JOBID>` command. Possible reasons are included in the following list:

- No Resources — The job requests resources of a type or amount that do not exist on the system.
- System Limits — The job is larger or longer than what is allowed by the specified system policies.
- Bank Failure — The allocations bank is experiencing failures.
- No Allocations — The job requests use of an account that is out of allocations and no fallback account has been specified.
- RM Reject — The resource manager refuses to start the job.
- RM Failure — The resource manager is experiencing failures.
- Policy Violation — The job violates certain throttling policies preventing it from running now and in the future.
- No QoS Access — The job does not have access to the QoS level it requests.

Jobs that are placed in a batch hold will show up within Moab in the state `BatchHold`.

## 9.1.5 Job Defer

In most cases, a job violating these policies is not placed into a batch hold immediately; rather, it is deferred. The parameter `DEFERTIME` indicates how long it is deferred. At this time, it is allowed back into the idle queue and again considered for scheduling. If it again is unable to run at that time or at any time in the future, it is again deferred for the timeframe specified by `DEFERTIME`. A job is released and deferred up to `DEFERCOUNT` times at which point the scheduler places a batch hold on the job and waits for a system administrator to determine the correct course of action. Deferred jobs have a Moab state of `Deferred`. As with jobs in the `BatchHold` state, the reason the job was deferred can be determined by use of the `checkjob` command.

At any time, a job can be released from any hold or deferred state using the `releasehold` command. The Moab logs should provide detailed information about the cause of any batch hold or job deferral.

**i** Under Moab, the reason a job is deferred or placed in a batch hold is stored in memory but is not checkpointed. Therefore this information is available only until Moab is recycled at which point the checkjob command no longer displays this reason information.

### Related Topics

- [DEFERSTARTCOUNT](#) - number of job start failures allowed before job is deferred

## 9.2 Job Priority Management

Job priority management is controlled via both configured and manual intervention mechanisms:

- Priority Configuration - see [Job Prioritization](#)
- Manual Intervention with [setspri](#)

## 9.3 Suspend/Resume Handling

When supported by the resource manager, Moab can suspend and resume jobs. By default, a job is suspended for one minute before it can be resumed by Moab. You can modify this default time using the [MINADMINSTIME](#) parameter.

**i** Moab schedules suspended jobs each iteration to see if they can be resumed. If the node the jobs are running on is free, then Moab automatically resumes the job.

Alternatively, a user can suspend his/her own jobs, but only an administrator can resume them. The administrator can resume jobs before the time set for Moab to resume.

A job must be marked as `suspendable` for Moab to suspend and resume it. To do so, either submit the job with the `suspendable` flag attached to it or configure a credential to pass the flag to its associated jobs. These methods are demonstrated in the examples below:

```
msub -l flags=suspendable
```

```
GROUPCFG[default] JOBFLAGS=SUSPENDABLE
```

Once the job is suspendable, Moab allows you to suspend jobs using the two following methods: (1) manually on the command line and (2) automatically in the `moab.cfg` file.

To manually suspend jobs, use the `mjobctl` command as demonstrated in the following example:

```
> mjobctl -s job05
```

*Moab suspends job05, preventing it from running immediately in the job queue.*

If you are an administrator and want to resume a job, use the `mjobctl` command as demonstrated in the following example:

```
> mjobctl -r job05
```

*Moab removes job05 from a suspended state and allows it to run.*

You can also configure the Moab preemption policy to suspend and resume jobs automatically by setting the `PREEMTPOLICY` parameter to `SUSPEND`. A sample Moab configuration looks like this:

```
PREEMTPOLICY SUSPEND
...
USERCFG[tom] JOBFLAGS=SUSPENDABLE
```

*Moab suspends jobs submitted by user tom if necessary to make resources available for jobs with higher priority.*

**i** If your resource manager has a native interface, you must configure `JOBSUSPENDURL` to suspend and resume jobs.

For more information about suspending and resuming jobs in Moab, see the following sections:

- manual preemption with the `mjobctl` command
- [Job preemption](#)

## 9.4 Checkpoint/Restart Facilities

Checkpointing records the state of a job, allowing for it to restart later without interruption to the job's execution. Checkpointing can be performed manually, as the result of [triggers](#) or [events](#), or in conjunction with various [QoS](#) policies.

Moab's ability to checkpoint is dependent upon both the cluster's [resource manager](#) and operating system. In most cases, two types of checkpoint are enabled, including (1) checkpoint and continue and (2) checkpoint and terminate. While either checkpointing method can be activated using the `mjobctl` command, only the checkpoint and terminate type is used by internal scheduling and event managements facilities.

Checkpointing behavior can be configured on a per-resource manager basis using various attributes of the [RMCFG](#) parameter.

---

### Related Topics

- [Job Preemption Overview](#)
- [PREEMTPOLICY](#) Parameter
- Resource Manager [CHECKPOINTSIG](#) Attribute
- Resource Manager [CHECKPOINTTIMEOUT](#) Attribute

## 9.5 Job Dependencies

### 9.5.1 Basic Job Dependency Support

By default, basic single step job dependencies are supported through completed/failed step evaluation. Basic dependency support does not require special configuration and is activated by default. Dependent jobs are only supported through a resource manager and therefore submission methods depend upon the specific resource manager being used.

Use the `-l depend=<STRING>` flag for the Torque `qsub` command and the Moab `msub` command. See documentation for the [DEPEND](#) resource manager extension value for more information.

**i** Torque `qsub` also supports the `-W x=depend=<STRING>` or `-W depend=<STRING>` flag. Moab `msub` command also supports the `-W x=depend=<STRING>` flag.

**i** Dependencies submitted with `-w` are handled by Torque, while dependencies submitted with `-l` are handled by Moab. For array dependencies, if any subjob fails then the status of the entire job array will be marked with a failure and any `afterok` dependencies will not be satisfied. Moab does not use the Torque `afteranyarray/afterokarray` syntax but instead uses the `after/afterok` syntax for both normal jobs and job arrays.

For other resource managers, consult the resource manager specific documentation.

## 9.5.2 Job Dependency Syntax

Dependency	Format	Description
<b>after</b>	after:<job> [:<job>]...	Job can start at any time after specified jobs have started execution.
<b>afterany</b>	afterany:<job> [:<job>]...	Job can start at any time after all specified jobs have completed regardless of completion status.
<b>afterok</b>	afterok:<job> [:<job>]...	Job can start at any time after all specified jobs have successfully completed.
<b>afternotok</b>	afternotok:<job> [:<job>]...	Job can start at any time after all specified jobs have completed unsuccessfully.
<b>before</b>	before:<job> [:<job>]...	Job can start at any time before specified jobs have started execution.
<b>beforeany</b>	beforeany:<job> [:<job>]...	Job can start at any time before all specified jobs have completed regardless of completion status.
<b>beforeok</b>	beforeok:<job> [:<job>]...	Job can start at any time before all specified jobs have successfully completed.
<b>beforenotok</b>	beforenotok:<job> [:<job>]...	Job can start at any time before any specified jobs have completed unsuccessfully.
<b>on</b>	on:<count>	Job can start after <count> dependencies on other jobs have been satisfied.
<b>synccount</b>	synccount:<count>	Job is the first in a set of jobs to be executed at the same time. <count> is the number of additional jobs in the set, which can be up to 5. synccount is valid for single-request jobs with Torque as the resource manager.
<b>syncwith</b>	syncwith:<job>	Job is an additional member of a set of jobs to be executed at the same time. Moab supports up to 5 jobs. syncwith is valid for single-request jobs with Torque as the resource manager.

**i** `<job>={jobname.jobname|jobid}`

When using JobName dependencies, prepend 'jobname.' to avoid ambiguity.

**i** The `before*`, `synccount`, and `syncwith` dependencies do not work with jobs submitted with `msub`; they work only with `qsub`.

**i** `before*` - we do not recommend using these dependencies because the job numbers of follow-up jobs would not be known yet.

Any of the dependencies containing `before` must be used in conjunction with the `on` dependency. So, if job A must run before job B, job B must be submitted with `depend=on:1`, as well as job A having `depend=before:A`. This means job B cannot run until one dependency of another job on job B has been fulfilled. This prevents job B from running until job A can be successfully submitted.

When you submit a dependency job and the dependency is not met, the job will remain idle in the queue indefinitely. To configure Moab to automatically cancel these failed dependency jobs, set the `CANCELFAILEDDEPENDENCYJOBS` scheduler flag. Moab also lets you cancel all jobs that a specified `<job_id>` depends on using `mjobctl -c flags=follow-dependency <job_id>`.

---

## Related Topics

- [Job Deadlines](#)

## 9.6 Job Defaults and Per Job Limits

In this topic:

- [9.6.1 Job Defaults - page 564](#)
- [9.6.2 Per Job Maximum Limits - page 565](#)
- [9.6.3 Per Job Minimum Limits - page 565](#)

### 9.6.1 Job Defaults

Job defaults can be specified on a per queue basis. These defaults are specified using the `CLASSCFG` parameter. The following table shows the applicable attributes:



Attribute	Format	Example
<b>DEFAULT.FEATURES</b>	comma-delimited list of node features	<pre>CLASSCFG[batch] DEFAULT.FEATURES=fast,io</pre> <p><i>Jobs submitted to class <code>batch</code> will request nodes features <code>fast</code> and <code>io</code>.</i></p>
<b>DEFAULT.WCLIMIT</b>	[ [ [DD:] HH:] MM:] SS	<pre>CLASSCFG[batch] DEFAULT.WCLIMIT=1:00:00</pre> <p><i>Jobs submitted to class <code>batch</code> will request one hour of walltime by default.</i></p>

## 9.6.2 Per Job Maximum Limits

Job maximum limits can be specified on a per queue basis. These defaults are specified using the **CLASSCFG** parameter. The following table shows the applicable attributes:

Attribute	Format	Example
<b>MAX.WCLIMIT</b>	[ [ [DD:] HH:] MM:] SS	<pre>CLASSCFG[batch] MAX.WCLIMIT=1:00:00</pre> <p><i>Jobs submitted to class <code>batch</code> can request no more than one hour of walltime.</i></p>

## 9.6.3 Per Job Minimum Limits

Furthermore, minimum job defaults can be specified with the **CLASSCFG** parameter. The following table shows the applicable attributes:

Attribute	Format	Example
<b>MIN.PROC</b>	<integer>	<pre>CLASSCFG[batch] MIN.PROC=10</pre> <p><i>Jobs submitted to class <code>batch</code> can request no less than ten processors.</i></p>

## Related Topics

- [Usage-based Limits](#)

## 9.7 General Job Policies

In this topic:

- [9.7.1 Multi-Node Support - page 566](#)
- [9.7.2 Multi-Req Support - page 566](#)
- [9.7.3 Malleable Job Support - page 567](#)
- [9.7.4 Enabling Job User Proxy - page 567](#)

There are a number of configurable policies that help control advanced job functions. These policies help determine allowable job sizes and structures.

### 9.7.1 Multi-Node Support

You can configure the ability to allocate resources from multiple nodes to a job with the `MAX.NODE` limit.

### 9.7.2 Multi-Req Support

Jobs can specify multiple types of resources for allocation. For example, a job could request 4 nodes with 256 MB of memory and 8 nodes with feature `fast` present.

Resources specified in a multi-req job are delimited with a plus sign (+).

**i** Neither `SPANEVENLY` nor `DELAY` values of the `NODESETPLUS` parameter will work with multi-req jobs or preemption.

*Example 9-1:*

```
-l nodes=4:ppn=1+10:ppn=5+2:ppn=2
```

*This example requests 4 nodes with 1 proc each, 10 nodes with 5 procs each, and 2 nodes with 2 procs each. The total number of processors requested is  $(4*1) + (10*5) + (2*2)$ , or 58 processors.*

**Example 9-2:**

```
-l nodes=15+1:ppn=4
```

*The job submitted in this example requests a total of 16 nodes. 15 of these nodes have no specific requirements, but the remaining node must have 4 processors.*

**Example 9-3:**

```
-l nodes=3:fast+1:io
```

*The job requests a total of 4 nodes: 3 nodes with the `fast` feature and 1 node with the `io` feature.*

## 9.7.3 Malleable Job Support

A job can specify whether it is able to use more processors or less processors and what effect, if any, that has on its wallclock time. For example, a job may run for 10 minutes on 1 processor, 5 minutes on 2 processors and 3 minutes on 3 processors. When a job is submitted with a task request list attached, Moab determines which task request fits best and molds the job based on its specifications. To submit a job with a task request list and allow Moab to mold it based on the current scheduler environment, use the [TRL \(Format 1\)](#) or the [TRL \(Format 2\)](#) flag in the Resource Manager Extension.

## 9.7.4 Enabling Job User Proxy

By default, user proxying is disabled. To be enabled, it must be authorized using the `PROXYLIST` attribute of the [USERCFG](#) parameter. This parameter can be specified either as a comma-delimited list of users or as the keyword `validate`. If the keyword `validate` is specified, the [RMCFG](#) attribute `JOBVALIDATEURL` should be set and used to confirm that the job's owner can proxy to the job's execution user. An example script performing this check for ssh-based systems is provided in the `tools` directory (see [Job Validate Tool Overview](#)).

For some resource managers (RM), proxying must also be enabled at the RM level. The following example shows how ssh-based proxying can be accomplished in a Moab+Torque with SSH environment.

**i** To validate proxy users, Moab must be running as root.

**Example 9-4: SSH Proxy Settings**

```
USERCFG[DEFAULT] PROXYLIST=validate
RMCFG[base] TYPE=<resource manager>
JOBVALIDATEURL=exec://$HOME/tools/job.validate.sshproxy.pl
```

```

> qmgr -c 's s allow_proxy_user=true'
> su - testuser
> qsub -I -u testuser2
qsub: waiting for job 533.igt.org to start
qsub: job 533.igt.org ready
testuser2@igt:~$

```

*In this example, the validate tool, 'job.validate.sshproxy.pl', can verify proxying is allowed by becoming the submit user and determining if the submit user can achieve passwordless access to the specified execution user. However, site-specific tools can use any method to determine proxy access including a flat file look-up, database lookup, querying of an information service such as NIS or LDAP, or other local or remote tests. For example, if proxy validation is required but end-user accounts are not available on the management node running Moab, the job validate service could perform the validation test on a representative remote host such as a login host.*

 This feature supports *qsub* only.

The job validate tool is highly flexible allowing any combination of job attributes to be evaluated and tested using either local or remote validation tests. The validate tool allows not only pass/fail responses but also allows the job to be modified, or rejected in a custom manner depending on the site or the nature of the failure.

### Related Topics

- [Usage Limits](#)
- [qmgr -c](#) in the Torque Administrator Guide
- [qsub -I](#) in the Torque Administrator Guide

## 9.8 Using a Local Queue

Moab allows jobs to be submitted directly to the scheduler. With a local queue, Moab is able to directly manage the job or translate it for resubmission to a standard resource manager queue. There are multiple advantages to using a local queue:

- Jobs can be translated from one resource manager job submission language to another (such as submitting a PBS job and running it on a Loadleveler cluster).
- Jobs can be migrated from one local resource manager to another.
- Jobs can be migrated to remote systems using Moab peer-to-peer functionality.
- Jobs can be dynamically modified and optimized by Moab to improve response time and system utilization.
- Jobs can be dynamically modified to account for system hardware failures or other issues.

- Jobs can be dynamically modified to conform to site policies and constraints.
- Grid jobs are supported.

## Local Queue Configuration

A local queue is configured just like a standard resource manager queue. It can have defaults, limits, resource mapping, and credential access constraints. The following table describes the most common settings:

Default queue	
<b>Format</b>	<code>RMCFG[internal] DEFAULTCLASS=&lt;CLASSID&gt;</code>
<b>Description</b>	<p>The job class/queue assigned to the job if one is not explicitly requested by the submitter.</p> <div style="border: 1px solid #0070C0; padding: 5px; margin-top: 10px;"> <p><b>i</b> All jobs submitted directly to Moab are initially received by the pseudo-resource manager <code>internal</code>. Therefore, default queue configuration can only be applied to it.</p> </div>
<b>Example</b>	<code>RMCFG[internal] DEFAULTCLASS=batch</code>

Class default resource requirements	
<b>Format</b>	<pre>CLASSCFG[&lt;CLASSID&gt;] DEFAULT.FEATURES=&lt;X&gt; CLASSCFG [&lt;CLASSID&gt;] DEFAULT.MEM=&lt;X&gt; CLASSCFG[&lt;CLASSID&gt;] DEFAULT.NODE=&lt;X&gt; CLASSCFG[&lt;CLASSID&gt;] DEFAULT.NODESET=&lt;X&gt; CLASSCFG[&lt;CLASSID&gt;] DEFAULT.PROC=&lt;X&gt; CLASSCFG[&lt;CLASSID&gt;] DEFAULT.WCLIMIT=&lt;X&gt;</pre>
<b>Description</b>	The settings assigned to the job if not explicitly set by the submitter. Default values are available for node features, per task memory, node count, nodeset configuration, processor count, and wallclock limit.
<b>Example</b>	<pre>CLASSCFG[batch] DEFAULT.WCLIMIT=4 DEFAULT.FEATURES=matlab</pre> <p>or</p> <pre>CLASSCFG[batch] DEFAULT.WCLIMIT=4 CLASSCFG[batch] DEFAULT.FEATURES=matlab</pre>

Class maximum resource limits	
<b>Format</b>	CLASSCFG [<CLASSID>] MAX.FEATURES=<X> CLASSCFG [<CLASSID>] MAX.NODE=<X> CLASSCFG [<CLASSID>] MAX.PROC=<X> CLASSCFG [<CLASSID>] MAX.WCLIMIT=<X>
<b>Description</b>	The maximum node features, node count, processor count, and wallclock limit allowed for a job submitted to the class/queue. If these limits are not satisfied, the job is not accepted and the submit request fails. MAX.FEATURES indicates that only the listed features can be requested by a job.
<b>Example</b>	<pre>CLASSCFG[smalljob] MAX.PROC=4 MAX.FEATURES=slow,matlab</pre> <p>or</p> <pre>CLASSCFG[smalljob] MAX.PROC=4 CLASSCFG[smalljob] MAX.FEATURES=slow,matlab</pre>
Class minimum resource limits	
<b>Format</b>	CLASSCFG [<CLASSID>] MIN.FEATURES=<X> CLASSCFG [<CLASSID>] MIN.NODE=<X> CLASSCFG [<CLASSID>] MIN.PROC=<X> CLASSCFG [<CLASSID>] MIN.WCLIMIT=<X>
<b>Description</b>	The minimum node features, node count, processor count, and wallclock limit allowed for a job submitted to the class/queue. If these limits are not satisfied, the job is not accepted and the submit request fails. MIN.FEATURES indicates that only the listed features can be requested by a job.
<b>Example</b>	<pre>CLASSCFG[bigjob] MIN.PROC=4 MIN.WCLIMIT=1:00:00</pre> <p>or</p> <pre>CLASSCFG[bigjob] MIN.PROC=4 CLASSCFG[bigjob] MIN.WCLIMIT=1:00:00</pre>
Class access	
<b>Format</b>	CLASSCFG [<CLASSID>] REQUIREDUSERLIST=<USERID> [,<USERID>] . . .
<b>Description</b>	The list of users who can submit jobs to the queue.
<b>Example</b>	<pre>CLASSCFG[math] REQUIREDUSERLIST=john,steve</pre>

Available resources	
<b>Format</b>	CLASSCFG [<CLASSID>] HOSTLIST=<HOSTID> [, <HOSTID>] . . .
<b>Description</b>	The list of nodes that jobs in the queue can use.
<b>Example</b>	<code>CLASSCFG[<i>special</i>] HOSTLIST=node001,node003,node13</code>

Class mapping between multiple sites is described in the section on Moab grid facilities.

If a job is submitted directly to the resource manager used by the local queue, the class default resource requirements are not applied. Also, if the job violates a local queue limitation, the job is accepted by the resource manager, but placed in the Blocked state.

## 9.9 Job Deadlines

In this topic:

- [9.9.1 Deadline Overview - page 571](#)
- [9.9.2 Setting Job Deadlines via QoS - page 571](#)
- [9.9.3 Job Termination Date - page 573](#)
- [9.9.4 Conflict Policies - page 573](#)

### 9.9.1 Deadline Overview

Job deadlines can be specified on a per job and per credential basis and are also supported using both absolute and QoS based specifications. A job requesting a deadline is first evaluated to determine if the deadline is acceptable. If so, Moab adds it to the list of deadline jobs and allocates resources to guarantee that all accepted deadline jobs are able to complete on or before their requested deadline. Once the scheduler confirms that all deadlines can be satisfied, it then optimizes resource allocation (in priority order) attempting to execute all jobs at the earliest possible time.

### 9.9.2 Setting Job Deadlines via QoS

Two types of job deadlines exist in Moab. The priority-based deadline linearly increases a job's priority as its deadline approaches (see [Deadline \(DEADLINE\) Subcomponent](#) for more information). The QoS method allows you to set a job completion time on job

submission if, and only if, it requests and is allowed to access a QoS with the `DEADLINE QFLAG` set. This method is more powerful than the priority method, because Moab will attempt to make a reservation for the job as soon as the job enters the queue in order to meet the deadline, essentially bumping it to the front of the queue.

When a job is submitted to a QoS with the `DEADLINE` flag set, the job's `-l deadline` attribute is honored. If such QoS access is not available, or if resources do not exist at job submission time to allow the deadline to be satisfied, the job's deadline request is ignored.

Two methods exist for setting deadlines with a QoS:

- Submitting a job to a deadline-enabled QoS and specifying a deadline using `msub -l`.
- Submitting a job to a deadline-enabled QoS with a `QTTARGET` specified.

### 9.9.2.A Setting Job Deadlines at Job Submission

This method of setting a job deadline allows you to specify a job deadline as you submit the job. You can set the deadline as either an exact date and time or as an amount of time after job submission (i.e., three hours after submission).

#### To specify a deadline on job submission

1. In `moab.cfg`, create a QoS with the `DEADLINE` flag enabled:

```
...
QOSCFG[special] QFLAGS=DEADLINE
```

*Jobs requesting the QoS special may submit jobs with a deadline that Moab will honor.*

2. Submit a job to the QoS and set a deadline. This can be either absolute or relative.
  - a. For an absolute deadline, use the format `hh:mm:ss_mm/dd/yy`. The following configuration sets a deadline for a job to finish by 8 a.m. on March 15th, 2022:

```
msub -l qos=special deadline=08:00:00_03/15/22 job.sh
```

*The job must finish running by 8 A.M. on March 15, 2022.*

- b. For a relative deadline, or the completion deadline of the job relative to its submission time, use the time format `[ [ DD : ] HH : ] MM : ] SS :`

```
msub -l qos=special deadline=5:00:00 job.sh
```

*The job's deadline is 5 hours after its submission.*



### 9.9.2.B Submitting a Job to a QoS with a Preconfigured Deadline

You can also set a relative job deadline by limiting the job's queue time. This method allows you to pre-configure the deadline rather than giving the power to specify a deadline to the user submitting the job. For jobs requesting these QoSes, Moab identifies and sets job deadlines to satisfy the corresponding response time targets.

#### To submit a job to a QoS with a preconfigured deadline

1. In `moab.cfg`, create a QoS with both the `DEADLINE_QFLAG` and a response time target (`QTTARGET`). The `QTTARGET` is the maximum amount of time that Moab should allow the job to be idle in the queue.

```
...
QOSCFG[special2] QFLAGS=DEADLINE QTTARGET=1:00:00
```

*Given this configuration, a job requesting QoS `special2` must spend a maximum of one hour in the queue.*

2. Submit a job requesting the `special2` quality of service:

```
msub -l qos=special2 walltime=2:00:00 job.sh
```

*This two-hour job has a completion time deadline set to three hours after its submission (one hour of target queue time and two hours of run time).*

### 9.9.3 Job Termination Date

In addition to job completion targets, jobs can also be submitted with a `TERMTIME` attribute. The scheduler attempts to complete the job prior to the termination date, but if it is unsuccessful, it will terminate (cancel) the job once the termination date is reached.

### 9.9.4 Conflict Policies

The specific policy can be configured using the `DEADLINEPOLICY` parameter. Moab does not have a default policy for this parameter.

Policy	Description
<b>CANCEL</b>	The job is canceled and the user is notified that the deadline could not be satisfied.
<b>HOLD</b>	The job has a batch hold placed on it indefinitely. The administrator can then decide what action to take.
<b>RETRY</b>	The job continually retries each iteration to meet its deadline; note that when used with <code>QTTARGET</code> the job's deadline continues to slide with relative time.

Policy	Description
<b>IGNORE</b>	The job has its request ignored and is scheduled as normal.

**i** Deadline scheduling might not function properly with per partition scheduling enabled. Check that `PARALLOCATIONPOLICY` is disabled to ensure `DEADLINEPOLICY` will work correctly.

### Related Topics

- [QoS Facilities](#)
- Job Submission [Eligible Start Time](#) constraints

## 9.10 Job Arrays

In this topic:

- [9.10.1 Job Array Overview - page 574](#)
- [9.10.2 Enabling Job Arrays - page 575](#)
- [9.10.3 Subjob Definitions - page 575](#)
- [9.10.4 Using Environment Variables to Specify Array Index Values - page 575](#)
- [9.10.5 Job Array Cancellation Policies - page 578](#)
- [9.10.6 Minimizing the Impact of Very Large Job Arrays - page 579](#)
- [9.10.7 Examples - page 581](#)

### 9.10.1 Job Array Overview

You can submit an array of jobs to Moab via the `msub` command. Array jobs are an easy way to submit many subjobs that perform the same work using the same script, but operate on different sets of data. Subjobs are the jobs created by a job array and are identified by the job array ID and an index; for example, if `235 [1]` is an identifier, the number `235` is a job array ID, and `1` is the subjob.

Subjobs of an array are executed in subjob index order.

**i** Moab job arrays are different from Torque job arrays.

## 9.10.2 Enabling Job Arrays

To enable job arrays, include the `ENABLEJOBARRAYS` parameter in the Moab configuration file (`moab.cfg`).

## 9.10.3 Subjob Definitions

Like a normal job, an array job submits a job script, but it additionally has a start index (`sidx`) and an end index (`eidx`); array jobs also have increment (`incr`) values, which Moab uses to create subjobs, all executing the same script. The model for subjob creation follows the formula of end index minus start index plus increment divided by the increment value:  $(eidx - sidx + incr) / incr$ .

To illustrate, suppose an array job has a start index of 1, an end index of 100, and an increment of 1. This is an array job that creates  $(100 - 1 + 1) / 1 = 100$  subjobs with indexes of 1, 2, 3, ..., 100. An increment of 2 produces  $(100 - 1 + 2) / 2 = 50$  subjobs with indexes of 1, 3, 5, ..., 99. An increment of 2 with a start index of 2 produces  $(100 - 2 + 2) / 2 = 50$  subjobs with indexes of 2, 4, 6, ..., 100. Again, subjobs are jobs in their own right that have a slightly different job naming convention `jobID[subJobIndex]` (e.g., `mycluster.45[37]` or `45[37]`).

## 9.10.4 Using Environment Variables to Specify Array Index Values

The script can use an environment variable to obtain the array index value to form data file and/or directory names unique to a job array's particular subjob. The following two environment variables are supplied so job scripts can recognize what index in the array they are in; use the `msub` command with the `-V` option to pass the environment parameters to the resource manager, or include the parameters in a job script; for example: `#PBS -V MOAB_JOBARRAYRANGE`.

Environment Parameter	Description
<b>MOAB_JOBARRAYINDEX</b>	Used to create dataset file names, directory names, and so forth, when splitting up a single problem into multiple jobs. For example, a user may split up a problem into 20 separate jobs, each with its own input and output data files whose names contain the numbers 1-20. To illustrate, assume a user submits the 20 subjobs using two <code>msub</code>

Environment Parameter	Description
	<p>commands; one to submit the ten even-numbered jobs and one to submit the ten odd-numbered jobs.</p> <pre>msub -t job1.[1-20:2] msub -t job2.[2-20:2]</pre> <p>The MOAB_JOBARRAYINDEX environment variable value would populate each of the two job arrays' ten subjobs as 1, 3, 5, 7, 9, 11, 13, 15, 17 and 19 for the first job array's ten subjobs, and 2, 4, 6, 8, 10, 12, 14, 16, 18, and 20 for the second job array's ten subjobs.</p>
<b>MOAB_JOBARRAYRANGE</b>	The count of jobs in the array.

### 9.10.4.A Control

Users can control individual subjobs in the same manner as normal jobs. In addition, a job array represents its group of subjobs and any user or administrator commands performed on a job array apply to its subjobs; for example, the command `canceljob <arrayJobId>` cancels all subjobs that belong to the job array. For more information about job control, see the documentation for the `mjobctl` command.

### 9.10.4.B Reporting

In the first example below, the parts unique to array subjobs are in bold:

```
$ checkjob -v Moab.1[1]
job Moab.1[1]

AName: Moab
State: Running
Creds: user:user1 group:usergroup1
WallTime: 00:00:17 of 8:20:00
SubmitTime: Thu Nov 4 11:50:03
(Time Queued Total: 00:00:00 Eligible: INFINITY)
StartTime: Thu Nov 4 11:50:03
Total Requested Tasks: 1
Req[0] TaskCount: 1 Partition: base
Average Utilized Procs: 0.96
NodeCount: 1
Allocated Nodes:
[node010:1]

Job Group: Moab.1
```

```

Parent Array ID: Moab.1
Array Index: 1
Array Range: 10
SystemID: Moab
SystemJID: Moab.1[1]
Task Distribution: node010
IWD: /home/user1
UMask: 0000
Executable: /opt/moab/spool/moab.job.3CvNj1
StartCount: 1
Partition List: base
SrcRM: internal DstRM: base DstRMJID: Moab.1[1]
Flags: ARRAYJOB,GLOBALQUEUE
StartPriority: 1
PE: 1.00
Reservation 'Moab.1[1]' (-00:00:19 -> 8:19:41 Duration: 8:20:00)

```

If the array range is not provided, the output displays all the jobs in the array:

```

$ checkjob -v Moab.1
job Moab.1

AName: Moab
Job Array Info:
  Name: Moab.1
  1 : Moab.1[1] : Running
  2 : Moab.1[2] : Running
  3 : Moab.1[3] : Running
  4 : Moab.1[4] : Running
  5 : Moab.1[5] : Running
  6 : Moab.1[6] : Running
  7 : Moab.1[7] : Running
  8 : Moab.1[8] : Running
  9 : Moab.1[9] : Running
 10 : Moab.1[10] : Running
 11 : Moab.1[11] : Running
 12 : Moab.1[12] : Running
 13 : Moab.1[13] : Running
 14 : Moab.1[14] : Running
 15 : Moab.1[15] : Running
 16 : Moab.1[16] : Running
 17 : Moab.1[17] : Running
 18 : Moab.1[18] : Running
 19 : Moab.1[19] : Running
 20 : Moab.1[20] : Running
Totals:
  Active: 20
  Idle: 0
  Complete: 0

```

You can also use `showq`. This displays the array master job with a count of how many subjobs are in each queue:

```

$ showq

active jobs-----
JOBID                USERNAME          STATE  PROCS   REMAINING          STARTTIME
Moab.1(5)            aesplin          Running  5      00:52:41  Thu Jun 23 17:05:56
Moab.2(1)            aesplin          Running  1      00:53:41  Thu Jun 23 17:06:56

```

```

6 active jobs                6 of 6 processors in use by local jobs (100.00%)
1 of 1 nodes active          (100.00%)

eligible jobs-----
JOBID            USERNAME      STATE  PROCS    WCLIMIT      QUEUETIME
Moab.2(4)        aesplin        Idle   4         1:00:00      Thu Jun 23 17:06:56

4 eligible jobs

blocked jobs-----
JOBID            USERNAME      STATE  PROCS    WCLIMIT      QUEUETIME
Moab.2(1)        aesplin        Blocked 1         1:00:00      Thu Jun 23 17:06:56

1 blocked job

Total jobs: 11
    
```

*Moab.1 has five subjobs running. Moab.2 has one subjob running, four waiting to run, and one that is currently blocked.*

### 9.10.5 Job Array Cancellation Policies

Job arrays can be canceled based on the success or failure of the first subjob, the first success or failure of any subjob, or if any subjob exits with a specified exit code. The job array cancellation policies are:

Cancel Policy	Description	Exclusivity
<b>CancelOnFirstFailure</b>	Cancels the job array if the first subjob (JOBARRAYINDEX = 1) fails. <pre>&gt; msub -t myarray[1-1000]%50 -1 ... ,flags=CancelOnFirstFailure</pre>	Mutually exclusive
<b>CancelOnFirstSuccess</b>	Cancels the job array if the first subjob (JOBARRAYINDEX = 1) succeeds. <pre>&gt; msub -t myarray[1-1000]%50 -1 ... ,flags=CancelOnFirstSuccess</pre>	
<b>CancelOnAnyFailure</b>	Cancels the job array if any subjob fails. <pre>&gt; msub -t myarray[1-1000]%50 -1 ... ,flags=CancelOnAnyFailure</pre>	
<b>CancelOnAnySuccess</b>	Cancels the job array if any subjob succeeds. <pre>&gt; msub -t myarray[1-1000]%50 -1 ... ,flags=CancelOnAnySuccess</pre>	

Cancel Policy	Description	Exclusivity
<b>CancelOnExitCode</b>	<p>Cancels the job array if any subjob returns the specified exit code.</p> <pre>&gt; msub -t myarray[1-1000%50] -1 ...,flags=CancelOnExitCode:&lt;error code list&gt;</pre> <p>The syntax for the error code list are ranges specified with a dash and individual codes delimited by a plus (+) sign, such as: 1-4+9+15</p> <p>Exit codes 1-387 are accepted.</p>	

Up to two cancellation policies can be specified for an array and the two policies must be delimited by a colon (:). The two 'first subjob' policies are mutually exclusive, as are the three 'any subjob' policies. You can use either 'first subjob' policy with one of the 'any subjob' policies, as shown in this example:

```
> msub -t myarray[1-1000]%50 -1 ...,flags=CancelOnFirstFailure:CancelOnExitCode:3-7+11
```

## 9.10.6 Minimizing the Impact of Very Large Job Arrays

### 9.10.6.A Potential Problems with Very Large Job Arrays

Since Moab uses multiple files to track jobs, including each subjob of a job array, a very large job array can add thousands of files that Moab has to create and manage. On each iteration, all of those subjobs are evaluated, and those eligible to run are considered for scheduling.

The jobs also need to be migrated to the workload manager (for this discussion, we'll just use Torque). That also adds additional work to both Moab and Torque. With default settings, it's possible that very large job arrays with a large number of job files, can cause problems in both Moab and Torque.

### 9.10.6.B Mitigating the Impact of Job Arrays

Fortunately, there are some settings that can be made to make Moab's work of handling a large job array much easier. This section will help you choose the best configuration options for your site.

For the settings in this section to be most effective, users must submit large job arrays using `msub` instead of `qsub`. The reason for this is when a job array is submitted using `qsub`, Torque creates all of the subjobs in its own queues, then passes that information

onto Moab during the cluster query process. This prevents Moab from managing the creation of subjobs in both Moab and Torque.

## **JOBMIGRATEPOLICY**

This option tells Moab how to handle the migration of jobs to the resource manager, typically Torque. This option applies to both normal jobs, as well as array subjobs. The default for this setting, if not specified, is `AUTO`, which allows Moab to decide how to migrate each job based on its current situation. If set to `IMMEDIATE`, that tells Moab to migrate all jobs to the RM as soon as they are created. Both of these will result in Torque having the files for all array subjobs soon after the job has been submitted, even if it does not run for days.

The best option is to configure Moab with `JOBMIGRATEPOLICY JUSTINTIME`. This option tells Moab to only migrate jobs once they have been scheduled to run, which minimizes the job count in the resource manager, as well as simplifying the management of jobs. It also simplifies cancellation of array jobs, as Torque only needs to cancel the ones it knows about, and Moab does not need to tell the RM about jobs it has not yet migrated.

## **MAXARRAYJOB**

This option sets an overall limit for the total number of running, plus idle jobs that can simultaneously exist. Jobs that exceed that limit remain in the blocked state until other jobs free up space. If the available resources allow more than this number of jobs to be running, then Moab will only run that number of jobs from the array and fill in the remaining resources with other jobs. In this situation, the Idle (or 'eligible') queue will be empty.

This limit can be applied to multiple credentials, as described in 'Credentials for `MAXARRAYJOB` and `MAXIARRAYJOB`' below.

## **MAXIARRAYJOB**

The default behavior for Moab is to migrate all of the subjobs to the Idle (or 'eligible') queue as soon as possible. Since with large job arrays it's highly unlikely that all of the subjobs can be run at once, this results in a lot of extra overhead during Moab's scheduling. Moab is continually checking and sometimes updating the job status, and updating the job files.

By setting the option `MAXIARRAYJOB`, the number of subjobs being considered at once for scheduling can be controlled. The option specifies how many jobs will be simultaneously allowed in the Idle (or 'eligible') queue. The array subjobs that don't fit remain in the blocked state. As jobs migrate to the running state, Moab will migrate additional subjobs to Idle state until the limit is again reached. For example, if the setting was 200, a job array with 10,000 subjobs would keep 200 subjobs in the idle queue, run as many as could be run at once, and the remaining subjobs remain blocked until there's room for them. By



selecting a value slightly greater than the maximum number of jobs that could be expected to be started at once, the huge job array should complete in the same time.

If the MAXARRAYJOB option is also configured, and Moab is able to run that limit of jobs at once, then the Idle queue will be empty no matter what this setting is.

This limit can be applied to multiple credentials, as described in 'Credentials for MAXARRAYJOB and MAXIARRAYJOB' below.

## Credentials for MAXARRAYJOB and MAXIARRAYJOB

Both of these limits can be applied to USERCFG, QOSCFG, CLASSCFG, and even JOBCFG. For QOSCFG and CLASSCFG, the limit will apply cumulatively to all jobs running under that credential, which can result in some users being blocked. In the case of USERCFG, the limit is applied to the specified user, or if DEFAULT is used, it will apply the limit to each user that does not have the limit defined in their own USERCFG.

When using JOBCFG, you are defining a job template. This makes understanding the results a lot more complicated, and for this reason we recommend avoiding this unless there is a specific reason. The problem is all jobs that match a job template are aggregated under this single limit, so a higher priority large array could potentially block all other arrays from running.

## The Nitro Product

Finally, if job arrays often contain serial workload (no MPI/parallel jobs) that can all run on a single node, the best solution may be the Nitro product. Instead of having a Moab job for each of these serial tasks, this product can run any number of tasks within a single Moab job. A task file is used to specify the command lines that need to be executed, and they are run with no Moab scheduling overhead. Nitro tracks the progress, and can even resume an interrupted run. A single Nitro job running trivial workload (hostname, for example) can easily run well over a thousand tasks per second.

### 9.10.7 Examples

Operations can be performed on individual jobs, a selection of jobs in a job array, or on the entire array.

#### 9.10.7.A Submitting Job Arrays

The syntax for submitting job arrays is: `msub -t [<jobname>]<indexlist> [%<limit>] arrayscript.sh`

The `<jobname>` and `<limit>` are optional. The jobname does not override the `jobID` Moab assigns to the array. When submitting an array with a jobname, Moab returns the `jobID`, which is the scheduler name followed by a unique ID.

For example, if the scheduler name in `moab.cfg` is `Moab (SCHEDCFG[Moab])`, submitting an array with a jobname responds like this:

```
> msub -t myarray[1-10] job.sh
Moab.6
```

To specify that only a certain number of subjobs in the array can run at a time, use the percent sign (%) delimiter. In this example, only five subjobs in the array can run at a time:

```
> msub -t myarray[1-1000]%5
```

To submit a specific set of array subjobs, use the comma delimiter in the array index list:

```
> msub -t myarray[1,2,3,4]
> msub -t myarray[1-5,7,10]
```

You can use the `checkjob` command on either the `jobID` or the `jobname` you specified:

```
> msub -t myarray[1-2] job.sh
Moab.10
$ checkjob -v myarray
job Moab.10
AName: myarray
Job Array Info:
  Name: Moab.10
  1 : Moab.10[1] : Running
  2 : Moab.10[2] : Running
Sub-jobs:          2
Active:            2 ( 100.0% )
Eligible:          0 ( 0.0% )
Blocked:           0 ( 0.0% )
Completed:         0 ( 0.0% )
State: Idle
Creds: user:tuser1 group:tgroup1
WallTime: 00:00:00 of 99:23:59:59
SubmitTime: Thu Jun  2 16:37:17
          (Time Queued Total: 00:00:33 Eligible: 00:00:00)
Total Requested Tasks: 1
Req[0] TaskCount: 1 Partition: ALL
```

To submit a job with a step size, use a colon in the array range and specify how many jobs to step. In the example below, a step size of 2 is requested. The subjobs will be numbered according to the step size inside the index limit. The array master job name will be the same as explained above.

```
$ msub -t myarray[2-10:2] job.sh

job Moab.15

$ checkjob -v myarray #or you could use 'checkjob -v Moab.15'
job Moab.15

AName: myarray
Job Array Info:
  Name: Moab.15
  2 : Moab.15[2] : Running
  4 : Moab.15[4] : Running
  6 : Moab.15[6] : Running
  8 : Moab.15[8] : Running
 10 : Moab.15[10] : Running

Sub-jobs:          5
  Active:          5 ( 100.0% )
  Eligible:        0 ( 0.0% )
  Blocked:         0 ( 0.0% )
  Completed:       0 ( 0.0% )

State: Idle
Creds: user:tuser1 group:tgroup1
WallTime: 00:00:00 of 99:23:59:59
SubmitTime: Thu Jun  2 16:37:17
  (Time Queued Total: 00:00:33 Eligible: 00:00:00)

Total Requested Tasks: 1

Req[0] TaskCount: 1 Partition: ALL
```

---

## Related Topics

- [Moab Workload Manager for Grids](#)
- [Job Dependencies](#)

## Chapter 10: General Node Administration

Moab has a very flexible and generalized definition of a [node](#). This flexible definition, together with the fact that Moab must inter-operate with many resource managers of varying capacities, requires that Moab must possess a complete set of mechanisms for managing nodes that in some cases might be redundant with resource manager facilities.

In this chapter:

10.1 Node Attribute Types .....	585
10.1.1 Resource Manager Specified 'Opaque' Attributes .....	585
10.1.2 Scheduler Specified Default Node Attributes .....	585
10.1.3 Scheduler Specified Node Attributes .....	586
10.2 Node Location .....	586
10.2.1 Partitions .....	586
10.2.2 Racks .....	586
10.2.3 Queues .....	587
10.2.4 Node Selection .....	588
10.3 Node Attributes .....	590
10.3.1 Configurable Node Attributes .....	590
10.3.2 Node Features/Node Properties .....	599
10.4 Node Specific Policies .....	600
10.4.1 Node Usage/Throttling Policies .....	600
10.4.2 Node Access Policies .....	602
10.5 Managing Shared Cluster Resources (Floating Resources) .....	602
10.5.1 Shared Cluster Resource Overview .....	602
10.5.2 Configuring Generic Consumable Floating Resources .....	603
10.5.3 Configuring Cluster File Systems .....	604
10.5.4 Configuring Cluster Licenses .....	604
10.5.5 Configuring Generic Resources as Features .....	604
10.5.6 Configuring Generic Resources as Licenses .....	606
10.6 Managing Node State .....	607
10.6.1 Node State Definitions .....	607
10.6.2 Specifying Node States within Native Resource Managers .....	608
10.6.3 Moab Based Node State Adjustment .....	608
10.6.4 Adjusting Scheduling Behavior Based on Reported Node State .....	608

10.6.5 Adding or Removing Nodes .....	609
10.7 Managing Consumable Generic Resources .....	609
10.7.1 Differences Between Node Features and Consumable Resources .....	610
10.7.2 Configuring Node-locked Consumable Generic Resources .....	610
10.7.3 Managing Generic Resource Race Conditions .....	612
10.8 Enabling Generic Metrics .....	613
10.8.1 Configuring Generic Metrics .....	613
10.8.2 Example Generic Metric Usage .....	615
10.9 Enabling Generic Events .....	616
10.9.1 Configuring Generic Events .....	617
10.9.2 Reporting Generic Events .....	620
10.9.3 Generic Events Attributes .....	621
10.9.4 Manually Creating Generic Events .....	621

## 10.1 Node Attribute Types

### 10.1.1 Resource Manager Specified 'Opaque' Attributes

Many resource managers support the concept of opaque node attributes, allowing a site to assign arbitrary strings to a node. These strings are opaque in the sense that the resource manager passes them along to the scheduler without assigning any meaning to them. Nodes possessing these opaque attributes can then be requested by various jobs. Using certain Moab parameters, sites can assign a meaning within Moab to these opaque node attributes and extract specific node information. For example, setting the parameter `FEATUREPROCSPEEDHEADER xps` causes a node with the opaque string `xps950` to be assigned a processor speed of 950 MHz within Moab.

### 10.1.2 Scheduler Specified Default Node Attributes

Some default node attributes can be assigned on a rack or partition basis. In addition, many node attributes can be specified globally by configuring the `DEFAULT` node template using the `NODECFG` parameter (i.e., `NODECFG[DEFAULT] PROCSPEED=3200`). Unless explicitly specified otherwise, nodes inherit node attributes from the associated rack or partition or from the default node template.

### 10.1.3 Scheduler Specified Node Attributes

The `NODECFG` parameter also allows direct per-node specification of virtually all node attributes supported via other mechanisms and also provides a number of additional attributes not found elsewhere. For example, a site administrator might want to specify something such as the following:

```
NODECFG[node031] MAXJOB=2 PROCSPEED=600 PARTITION=small
```

**i** These approaches can be mixed and matched according to the site's local needs. Precedence for the approaches generally follows the order listed earlier in cases where conflicting node configuration information is specified through one or more mechanisms.

## 10.2 Node Location

Nodes can be assigned three types of location information based on partitions, racks, and queues.

In this topic:

[10.2.1 Partitions - page 586](#)

[10.2.2 Racks - page 586](#)

[10.2.3 Queues - page 587](#)

[10.2.4 Node Selection - page 588](#)

### 10.2.1 Partitions

The first form of location assignment, the partition, allows nodes to be grouped according to physical resource constraints or policy needs. By default, jobs are not allowed to span more than one partition so partition boundaries are often valuable if an underlying network topology make certain resource allocations undesirable. Additionally, per-partition policies can be specified to grant control over how scheduling is handled on a partition by partition basis. See the [Partition Overview](#) for more information.

### 10.2.2 Racks

Rack-based location information is orthogonal to the partition based configuration and is mainly an organizational construct. In general rack based location usage, a node is assigned

both a rack and a slot number. This approach has descended from the IBM SP2 organizational approach where a rack can contain any number of slots, but typically contains between 1 and 99. Using the rack and slot number combo, individual compute nodes can be grouped and displayed in a more ordered manner in certain Moab commands (i.e., [showstate](#)). Currently, rack information can only be specified directly by the system via the SDR interface on SP2/Loadleveler systems. In all other systems, this information must be specified using an information service or specified manually using the [RACK](#), [SLOT](#), and [SIZE](#) attributes of the [NODECFG](#) parameter.

**i** Sites can arbitrarily assign nodes to racks and rack slots without impacting scheduling behavior. Neither rack numbers nor rack slot numbers need to be contiguous; their use is simply for convenience purposes in displaying and analyzing compute resources.

#### Example 10-1:

```
NODECFG[node024] RACK=1 SLOT=1
NODECFG[node025] RACK=1 SLOT=2
NODECFG[node026] RACK=2 SLOT=1 PARTITION=special
...
```

When specifying node and rack information, slot values must be in the range of 1 to 99, and racks must be in the range of 1 to 399.

## 10.2.3 Queues

Some resource managers allow queues (or classes) to be defined and then associated with a subset of available compute resources. With systems such as Loadleveler or PBSPro these queue to node mappings are automatically detected. On resource managers that do not provide this service, Moab provides alternative mechanisms for enabling this feature.

### 10.2.3.A Torque Queue to Node Mapping

Under Torque, queue to node mapping can be accomplished by using the `qmgr` command to set the `acl_hosts` parameter to the mapping hostlist desired. Further, the `acl_host_enable` parameter should be set to `False`.

**i** Setting `acl_hosts` and then setting `acl_host_enable` to `True` constrains the list of hosts from which jobs can be submitted to the queue.

The following example highlights this process and maps the queue `debug` to the nodes `host14` through `host17`:

```

> qmgr
Max open servers: 4
Qmgr: set queue debug acl_hosts = "host14,host15,host16,host17"
Qmgr: set queue debug acl_host_enable = false
Qmgr: quit

```

**i** All queues that do not have `acl_hosts` specified are global; that is, they show up on every node. To constrain these queues to a subset of nodes, each queue requires its own `acl_hosts` parameter setting.

## 10.2.4 Node Selection

When selecting or specifying nodes either via command line tools or via configuration file based lists, Moab offers three types of node expressions that can be based on node lists, exact lists, node ranges, or regular expressions.

### 10.2.4.A Node Lists

Node lists can be specified as one or more comma or whitespace delimited node IDs. Specified node IDs can be based on either short or fully qualified hostnames. Each element will be interpreted as a regular expression.

```

SRCFG[basic]  HOSTLIST=cl37.icluster,ax45,ax46
...

```

### 10.2.4.B Exact Lists

When Moab receives a list of nodes it will, by default, interpret each element as a regular expression. To disable this and have each element interpreted as a string node name, the `l:` can be used as in the following example:

```

> setres l:n00,n01,n02

```

### 10.2.4.C Node Range

Node lists can be specified as one or more comma or whitespace delimited node ranges. Each node range can be based using either `<STARTINDEX>-<ENDINDEX>` or `<HEADER> [<STARTINDEX>-<ENDINDEX>]` format. To explicitly request a range, the node expression must be preceded with the string `r:` as in the following example:

```

> setres r:37-472,513,516-855

```

When you specify a `<HEADER>` for the range, note that it must only contain alphabetical characters. As always, the range must be numeric.



```
CLASSCFG[long] HOSTLIST=r:anc-b[37-472]
```

**i** Only one expression is allowed with node ranges.

**i** By default, Moab attempts to extract a node's node index assuming this information is built into the node's naming convention. If needed, this information can be explicitly specified in the Moab configuration file using `NODECFG`'s `NODEINDEX` attribute, or it can be extracted from alternatively formatted node IDs by specifying the `NODEIDFORMAT` parameter.

### 10.2.4.D Node Regular Expression

Node lists can also be specified as one or more comma or whitespace delimited regular expressions. Each node regular expression must be specified in a format acceptable by the standard C regular expression libraries that allow support for wildcard and other special characters such as the following:

- \* (asterisk)
- . (period)
- [ ] (left and right bracket)
- ^ (caret)
- \$ (dollar)

Node lists are by default interpreted as a regular expression but can also be explicitly requested with the string `x:` as in the following examples:

```
# select nodes c130 thru c155
SRCFG[basic] HOSTLIST=x:c1[34],c15[0-5]
...
```

```
# select nodes c130 thru c155
SRCFG[basic] HOSTLIST=c1[34],c15[0-5]
...
```

**i** To control node selection search ordering, set the `OBJECTELIST` parameter to one of the following options: `exact`, `range`, `regex`, `rangere`, or `rerange`.

## 10.3 Node Attributes

In this topic:

[10.3.1 Configurable Node Attributes - page 590](#)

[10.3.2 Node Features/Node Properties - page 599](#)

### 10.3.1 Configurable Node Attributes


Nodes can possess a large number of attributes describing their configuration, which are specified using the `NODECFG` parameter. The majority of these attributes such as operating system or configured network interfaces can only be specified by the direct resource manager interface. However, the number and detail of node attributes varies widely from resource manager to resource manager. Sites often have interest in making scheduling decisions based on scheduling attributes not directly supplied by the resource manager. Configurable node attributes are listed in the following table; click an attribute for more detailed information:


*Node Attributes*

ACCESS	NODEAVAILABILITYPOLICY	PROVRM
ARCH	NODEINDEX	RACK
CHARGERATE	NODETYPE	RADISK
COMMENT	OS	RAMEM
ENABLEPROFILING	OSLIST	RCDISK
FEATURES	PARTITION	RCMEM
GRES	POWERPOLICY	RCPROC
MAXJOB	PREEMPTMAXCPULOAD	RCSWAP
MAXJOBPERUSER	PREEMPTMINMEMAVAIL	SIZE
MAXLOAD	PREEMTPOLICY	SLOT
MAXPE	PRIORITY	SPEED
MAXPEPERJOB	PRIORITYF	TRIGGER
MAXPROC	PROCSPEED	VARIABLE

Attribute	Description
<b>ACCESS</b>	<p>The node access policy that can be one of <code>SHARED</code>, <code>SHAREDONLY</code>, <code>SINGLEJOB</code>, <code>SINGLETASK</code>, or <code>SINGLEUSER</code>. See <a href="#">Node Access Policies</a> for more details.</p> <pre>NODECFG[node013] ACCESS=singlejob</pre>
<b>ARCH</b>	<p>The node's processor architecture.</p> <pre>NODECFG[node013] ARCH=opteron</pre>
<b>CHARGERATE</b>	<p>Allows a site to assign specific charging rates to the usage of particular resources. The <code>CHARGERATE</code> value can be specified as a floating point value and is integrated into a job's total charge (as documented in the <a href="#">Charging and Allocation Management</a> section).</p> <pre>NODECFG[DEFAULT] CHARGERATE=1.0 NODECFG[node003] CHARGERATE=1.5 NODECFG[node022] CHARGERATE=2.5</pre>
<b>COMMENT</b>	<p>Allows an organization to annotate a node via the configuration file to indicate special information regarding this node to both users and administrators. The <code>COMMENT</code> value can be specified as a quote delimited string as shown in the example that follows. Comment information is visible using <a href="#">checknode</a>, <a href="#">mdiag</a>, and Moab Cluster Manager.</p> <pre>NODECFG[node013] COMMENT="Login Node"</pre>

Attribute	Description
<b>ENABLEPROFILING</b>	<p>Allows an organization to track node state over time. This information is available using <a href="#">showstats -n</a>.</p> <pre data-bbox="618 407 1414 447">NODECFG[DEFAULT] ENABLEPROFILING=TRUE</pre>
<b>FEATURES</b>	<p>Not all resource managers allow specification of opaque <a href="#">node features</a> (also known as node properties). For these systems, the <a href="#">NODECFG</a> parameter can be used to directly assign a list of node features to individual nodes. To append node features, use <code>FEATURES=&lt;X&gt;</code>; to overwrite or remove a node's features, you must update them in your Moab configuration file or resource manager.</p> <pre data-bbox="618 737 1414 777">NODECFG[node013] FEATURES=gpfs,fastio</pre> <p data-bbox="667 789 1365 898"><i>Node node013 now has features gpfs and fastio in addition to any other features configured in this file or the resource manager.</i></p> <p data-bbox="618 919 1414 1010"><b>i</b> The total number of supported node features is limited as described in the <a href="#">Adjusting Default Limits</a> section.</p> <p data-bbox="618 1024 1414 1310"><b>i</b> If supported by the resource manager, the resource manager specific manner of requesting node features/properties within a job can be used. (Within Torque, use <code>qsub -l nodes=&lt;NODECOUNT&gt; :&lt;NODEFEATURE&gt;</code>.) However, if either not supported within the resource manager or if support is limited, the Moab <a href="#">feature</a> resource manager extension can be used.</p>
<b>GRES</b>	<p>Many resource managers do not allow specification of consumable generic node resources. For these systems, the <a href="#">NODECFG</a> parameter can be used to directly assign a list of consumable generic attributes to individual nodes or to the special pseudo-node <a href="#">global</a>, which provides shared cluster (floating) consumable resources. To set/overwrite a node's generic resources, use <code>GRES=&lt;NAME&gt; [ :&lt;COUNT&gt; ]</code>. See <a href="#">Managing Consumable Generic Resources</a>.</p> <pre data-bbox="618 1633 1414 1673">NODECFG[node013] GRES=quickcalc:20</pre>
<b>MAXJOB</b>	See <a href="#">Node Policies</a> for details.

Attribute	Description
<b>MAXJOBPERUSER</b>	See <a href="#">Node Policies</a> for details.
<b>MAXLOAD</b>	See <a href="#">Node Policies</a> for details.
<b>MAXPE</b>	See <a href="#">Node Policies</a> for details.
<b>MAXPEPERJOB</b>	<p>Maximum allowed Processor Equivalent per job on this node. A job will not be allowed to run on this node if its PE exceeds this number.</p> <pre>NODECFG[node024] MAXPEPERJOB=10000 ...</pre>
<b>MAXPROC</b>	<p>Maximum dedicated processors allowed on this node. No jobs are scheduled on this node when this number is reached. See <a href="#">Node Policies</a> for more information.</p> <pre>NODECFG[node024] MAXPROC=8 ...</pre>
<b>NODEAVAILABILITYPOLICY</b>	<p>Specifies how available node resources are reported.</p> <div style="border: 1px solid #0070C0; padding: 5px; margin: 5px 0;"> <p> This sets the NODEAVAILABILITYPOLICY at the local level and uses a different format from the NODEAVAILABILITYPOLICY server parameter. See <a href="#">NODEAVAILABILITYPOLICY</a>.</p> </div> <pre>NODECFG[node00] NODEAVAILABILITYPOLICY=DEDICATED:PROC,UTILIZED:MEM,COMBINED:DISK</pre>
<b>NODEINDEX</b>	The node's index. See <a href="#">Node Location</a> for details.
<b>NODETYPE</b>	<p>The NODETYPE attribute is most commonly used in conjunction with an accounting manager such as Moab Accounting Manager. In these cases, each node is assigned a node type and within the accounting manager, each node type is assigned a charge rate. For example, a site administrator might want to charge users more for using large memory nodes and might assign a node type of BIGMEM to these nodes. The accounting manager would then charge a premium rate for jobs using BIGMEM nodes. See the <a href="#">Accounting, Charging, and Allocation Management</a> for more information.</p>

Attribute	Description
	<p>Node types are specified as simple strings. If no node type is explicitly set, the node will possess the default node type of <code>DEFAULT</code>. Node type information can be specified directly using <code>NODECFG</code> or through use of the <a href="#">FEATURENODETYPEHEADER</a> parameter.</p> <pre data-bbox="618 499 1414 541">NODECFG[node024] NODETYPE=BIGMEM</pre>
<b>OS</b>	<p>This attribute specifies the node's operating system.</p> <pre data-bbox="618 642 1414 684">NODECFG[node013] OS=suse15</pre> <div data-bbox="618 701 1414 894" style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p> Because the Torque operating system overwrites the Moab operating system, change the operating system with <code>opsys</code> instead of <b>OS</b> if you are using Torque, and/or you can use <code>RMCFG[torque] FLAGS=IgnOS</code> to override that.</p> </div>
<b>OSLIST</b>	<p>This attribute specifies the list of operating systems the node can run.</p> <pre data-bbox="618 1024 1414 1066">NODECFG[compute002] OSLIST=linux,windows</pre>
<b>PARTITION</b>	See <a href="#">Node Location</a> for details.
<b>POWERPOLICY</b>	<p>The <code>POWERPOLICY</code> can be set to <code>OnDemand</code> or <code>STATIC</code>. It defaults to <code>STATIC</code> if not set. If set to <code>STATIC</code>, Moab will never automatically change the power status of a node. If set to <code>OnDemand</code>, Moab will turn the machine off and on based on workload and global settings. See <a href="#">Chapter 15: Green Computing - page 762</a> for further details.</p>
<b>PREEMPTMAXCPULOAD</b>	<p>If the node CPU load exceeds the specified value, any batch jobs running on the node are preempted using the preemption policy specified with the node's <a href="#">PREEMTPOLICY</a> attribute. If this attribute is not specified, the global default policy specified with <a href="#">PREEMTPOLICY</a> parameter is used. See <a href="#">Sharing Server Resources</a> for further details.</p> <pre data-bbox="618 1650 1414 1734">NODECFG[node024] PRIORITY=-150 COMMENT="NFS Server Node" NODECFG[node024] PREEMTPOLICY=CANCEL PREEMPTMAXCPULOAD=1.2 ...</pre>

Attribute	Description
<b>PREEMPTMINMEMAVAIL</b>	<p>If the available node memory drops below the specified value, any batch jobs running on the node are preempted using the preemption policy specified with the node's <a href="#">PREEMTPOLICY</a> attribute. If this attribute is not specified, the global default policy specified with <a href="#">PREEMTPOLICY</a> parameter is used. See <a href="#">Sharing Server Resources</a> for further details.</p> <pre data-bbox="618 527 1414 621"> NODECFG[node024] PRIORITY=-150 COMMENT="NFS Server Node" NODECFG[node024] PREEMTPOLICY=CANCEL PREEMPTMINMEMAVAIL=256 ... </pre>
<b>PREEMTPOLICY</b>	<p>If any node preemption policies are triggered (such as <a href="#">PREEMTMAXCPULOAD</a> or <a href="#">PREEMPTMINMEMAVAIL</a>) any batch jobs running on the node are preempted using this preemption policy if specified. If not specified, the global default preemption policy specified with <a href="#">PREEMTPOLICY</a> parameter is used. See <a href="#">Sharing Server Resources</a> for further details.</p> <pre data-bbox="618 877 1414 972"> NODECFG[node024] PRIORITY=-150 COMMENT="NFS Server Node" NODECFG[node024] PREEMTPOLICY=CANCEL PREEMTMAXCPULOAD=1.2 ... </pre>
<b>PRIORITY</b>	<p>The <code>PRIORITY</code> attribute specifies the fixed node priority relative to other nodes. It is only used if <a href="#">NODEALLOCATIONPOLICY</a> is set to <code>PRIORITY</code>. The default node priority is 0. A default cluster-wide node priority can be set by configuring the <code>PRIORITY</code> attribute of the <code>DEFAULT</code> node. See <a href="#">Priority Node Allocation</a> for more details.</p> <pre data-bbox="618 1226 1414 1320"> NODEALLOCATIONPOLICY PRIORITY NODECFG[node024] PRIORITY=120 ... </pre>
<b>PRIORITYF</b>	<p>The <code>PRIORITYF</code> attribute specifies the function to use when calculating a node's allocation priority specific to a particular job. It is only used if <a href="#">NODEALLOCATIONPOLICY</a> is set to <code>PRIORITY</code>. The default node priority function sets a node's priority exactly equal to the configured node priority. The priority function allows a site to indicate that various environmental considerations such as node load, reservation affinity, and ownership be taken into account as well using the following format:</p> <pre data-bbox="618 1654 1187 1717"> &lt;COEFFICIENT&gt; * &lt;ATTRIBUTE&gt; [ + &lt;COEFFICIENT&gt; * &lt;ATTRIBUTE&gt; ] ... </pre> <p><code>&lt;ATTRIBUTE&gt;</code> is an attribute from the table found in the <a href="#">Priority Node Allocation</a> section.</p>

Attribute	Description
	<p>A default cluster-wide node priority function can be set by configuring the <code>PRIORITYF</code> attribute of the <code>DEFAULT</code> node. See <a href="#">Priority Node Allocation</a> for more details.</p> <pre data-bbox="618 436 1414 527" style="border: 1px dashed black; padding: 5px;"> NODEALLOCATIONPOLICY PRIORITY NODECFG[node024] PRIORITYF='APROC + .01 * AMEM - 10 * JOBCOUNT' ... </pre>
<b>PROCSPEED</b>	<p>Knowing a node's processor speed can help the scheduler improve intra-job efficiencies by allocating nodes of similar speeds together. This helps reduce losses due to poor internal job load balancing. Moab's node set scheduling policies allow a site to control processor speed based allocation behavior.</p> <p>Processor speed information is specified in MHz and can be indicated directly using <code>NODECFG</code> or through use of the <code>FEATUREPROCSPEEDHEADER</code> parameter.</p>
<b>PROVRM</b>	<p>Provisioning resource managers can be specified on a per node basis. This allows flexibility in mixed environments. If the node does not have a provisioning resource manager, the default provisioning resource manager will be used. The default is always the first one listed in <code>moab.cfg</code>.</p> <pre data-bbox="618 1066 1414 1205" style="border: 1px dashed black; padding: 5px;"> RMCFG[prov] TYPE=NATIVE RESOURCETYPE=PROV RMCFG[prov] PROVDURATION=10:00 RMCFG[prov] NODEMODIFYURL=exec://\$HOME/tools/os.switch.pl ... NODECFG[node024] PROVRM=prov </pre>
<b>RACK</b>	<p>The rack associated with the node's physical location. Valid values range from 1 to 400. See <a href="#">Node Location</a> for details.</p>
<b>RADISK</b>	<p>Jobs can request a certain amount of disk space through the <a href="#">RM Extension String's</a> <code>DDISK</code> parameter. When done this way, Moab can track the amount of disk space available for other jobs. <code>RADISK</code> is used as an initial value and is subsequently determined by <code>RCDISK - &lt;JOB USAGE&gt;</code>.</p>
<b>RAMEM</b>	<p>The total available memory on a node. Jobs can request a certain amount of real memory (RAM) in MB through the <a href="#">RM Extension String's</a> <code>DMEM</code> parameter. When done this way, Moab can track the amount of memory available for other jobs. <code>RAMEM</code> is used as an initial value and is subsequently determined by <code>RCMEM - &lt;JOB USAGE&gt;</code>.</p>



Attribute	Description
<b>RCDISK</b>	<p>Jobs can request a certain amount of disk space (in MB) through the <a href="#">RM Extension String's</a> <code>DDISK</code> parameter. When done this way, Moab can track the amount of disk space available for other jobs. The <code>RCDISK</code> attribute constrains the amount of disk reported by a resource manager while the <code>RADISK</code> attribute specifies the amount of disk available to jobs. If the resource manager does not report available disk, the <code>RADISK</code> attribute should be used.</p>
<b>RCMEM</b>	<p>Jobs can request a certain amount of real memory (RAM) in MB through the <a href="#">RM Extension String's</a> <code>DMEM</code> parameter. When done this way, Moab can track the amount of memory available for other jobs. The <code>RCMEM</code> attribute constrains the amount of RAM reported by a resource manager while the <code>RAMEM</code> attribute specifies the amount of RAM available to jobs. If the resource manager does not report available memory, the <code>RAMEM</code> attribute should be used.</p> <p>Note that memory reported by the resource manager will override the configured value unless a trailing caret (^) is used.</p> <pre data-bbox="617 966 1412 1039">NODECFG[node024] RCMEM=2048 ...</pre> <p data-bbox="665 1039 1364 1155"><i>If the resource manager does not report any memory, then Moab will assign node024 2048 MB of memory.</i></p> <pre data-bbox="617 1165 1412 1239">NODECFG[node024] RCMEM=2048^ ...</pre> <p data-bbox="665 1239 1364 1333"><i>Moab will assign 2048 MB of memory to node024 regardless of what the resource manager reports.</i></p>
<b>RCPROC</b>	<p>The <code>RCPROC</code> specifies the number of processors available on a compute node.</p> <pre data-bbox="617 1459 1412 1533">NODECFG[node024] RCPROC=8 ...</pre>
<b>RCSWAP</b>	<p>Jobs can request a certain amount of swap space in MB.</p>

Attribute	Description
	<p><b>i</b> <b>RCSWAP</b> works similarly to <b>RCMEM</b>. Setting <b>RCSWAP</b> on a node will set the swap but can be overridden by swap reported by the resource manager. If the trailing caret (^) is used, Moab will ignore the swap reported by the resource manager and use the configured amount.</p> <pre>NODECFG[node024] RCSWAP=2048 ...</pre> <p><i>If the resource manager does not report any memory, Moab will assign node024 2048 MB of swap.</i></p> <pre>NODECFG[node024] RCSWAP=2048^ ...</pre> <p><i>Moab will assign 2048 MB of swap to node024 regardless of what the resource manager reports.</i></p>
<b>SIZE</b>	<p>The number of slots or size units consumed by the node. This value is used in graphically representing the cluster using <a href="#">showstate</a> or Moab Cluster Manager. See <a href="#">Node Location</a> for details. For display purposes, legal size values include 1, 2, 3, 4, 6, 8, 12, and 16.</p> <pre>NODECFG[node024] SIZE=2 ...</pre>
<b>SLOT</b>	<p>The first slot in the rack associated with the node's physical location. Valid values range from 1 to <code>M_MAX_RACKSIZE</code> (default=64). See <a href="#">Node Location</a> for details.</p>
<b>SPEED</b>	<p>Because today's processors have multiple cores and adjustable clock frequency, this feature has no meaning and will be deprecated.</p> <p><b>i</b> The <b>SPEED</b> specification must be in the range of 0.01 to 100.0.</p>
<b>TRIGGER</b>	<p>See <a href="#">About Object Triggers</a> for information.</p>
<b>VARIABLE</b>	<p>Variables associated with the given node, which can be used in job scheduling. See <code>-l PREF</code>.</p>

Attribute	Description
	<pre>NODECFG[node024] VARIABLE=var1 ...</pre>

### 10.3.2 Node Features/Node Properties

A node feature (or node property) is an opaque string label that is associated with a compute node. Each compute node can have any number of node features assigned to it, and jobs can request allocation of nodes that have specific features assigned. Node features are labels and their association with a compute node is not conditional, meaning they cannot be consumed or exhausted.

Node features can be assigned by the resource manager, and this information can be imported by Moab or node features can be specified within Moab directly. Moab supports hyphens and underscores in node feature names.

As a convenience feature, certain node attributes can be specified via node features using the parameters listed in the following table:

Parameter	Description
<b>FEATURENODETYPEHEADER</b>	Set <a href="#">Node Type</a>
<b>FEATUREPARTITIONHEADER</b>	Set <a href="#">Partition</a>
<b>FEATUREPROCSPEEDHEADER</b>	Set <a href="#">Processor Speed</a>
<b>FEATURERACKHEADER</b>	Set <a href="#">Rack</a>
<b>FEATURESLOTHEADER</b>	Set <a href="#">Slot</a>

*Example 10-2:*

```
FEATUREPARTITIONHEADER par
FEATUREPROCSPEEDHEADER cpu
```

#### Related Topics

- [Job Preferences](#)
- [Specifying Node Features \(Node Properties\) in the Torque Administrator Guide](#)

- [Configuring Node Features in Moab with NODECFG](#)
- [Viewing Feature Availability Breakdown with mdiag -t](#)
- [Differences between Node Features and Managing Consumable Generic Resources](#)

## 10.4 Node Specific Policies

Node policies within Moab allow specification of not only how the node's load should be managed, but who can use the node, and how the node and jobs should respond to various events. These policies allow a site administrator to specify on a node by node basis what the node will and will not support. Node policies can be applied to specific nodes or applied system-wide using the specification `NODECFG [DEFAULT] . . . .`

In this topic:

[10.4.1 Node Usage/Throttling Policies - page 600](#)

[10.4.2 Node Access Policies - page 602](#)

### 10.4.1 Node Usage/Throttling Policies

#### 10.4.1.A MAXJOB

This policy constrains the number of total independent jobs a given node can run simultaneously. It can only be specified via the `NODECFG` parameter.

#### 10.4.1.B MAXJOBPERUSER

Constrains the number of total independent jobs a given node can run simultaneously associated with any single user. It can only be specified via the `NODECFG` parameter.

#### 10.4.1.C MAXJOBPERGROUP

Constrains the number of total independent jobs a given node can run simultaneously associated with any single group. It can only be specified via the `NODECFG` parameter. Setting `MAXLOAD` here to -1 unsets the `NODEMAXLOAD` parameter setting.

#### 10.4.1.D MAXLOAD

`MAXLOAD` constrains the CPU load the node will support as opposed to the number of jobs. This maximum load policy can also be applied system wide using the parameter

## NODEMAXLOAD.

### 10.4.1.E MAXPE

This policy constrains the number of total dedicated processor-equivalents a given node can support simultaneously. It can only be specified via the `NODECFG` parameter. For more information, see the subsection 'PE (Processor Equivalent) Calculation' under [Scheduling Environment](#).

### 10.4.1.F MAXPROC

This policy constrains the number of total dedicated processors a given node can support simultaneously. It can only be specified via the `NODECFG` parameter.

### 10.4.1.G MAXPROCUSER

This policy constrains the number of total processors a given node can have dedicated to any single user. It can only be specified via the `NODECFG` parameter.

### 10.4.1.H MAXPROCGRP

This policy constrains the number of total processors a given node can have dedicated to any single group. It can only be specified via the `NODECFG` parameter.

**i** Node throttling policies are used strictly as constraints. If a node is defined as having a single processor or the `NODEACCESSPOLICY` is set to `SINGLETASK`, and a `MAXPROC` policy of 4 is specified, Moab will not run more than one task per node. A node's configured processors must be specified so that multiple jobs can run and then the `MAXJOB` policy will be effective. The number of configured processors per node is specified on a resource manager specific basis. PBS, for example, allows this to be adjusted by setting the number of virtual processors with the `np` parameter for each node in the PBS `nodes` file.

#### Example 10-3:

```
NODECFG [node024] MAXJOB=4 MAXJOBPERUSER=2
NODECFG [node025] MAXJOB=2
NODECFG [node026] MAXJOBPERUSER=1
NODECFG [DEFAULT] MAXLOAD=2.5
...
```

## 10.4.2 Node Access Policies

While most sites require only a single cluster wide node access policy (commonly set using [NODEACCESSPOLICY](#)), it is possible to specify this policy on a node by node basis using the [ACCESS](#) attributes of the [NODECFG](#) parameter. This attribute can be set to any of the valid node access policy values listed in the [Node Access Policies](#) section.

*Example 10-4:*

To set a global policy of `SINGLETASK` on all nodes except nodes 13 and 14, use the following:

```
# by default, enforce dedicated node access on all nodes
NODEACCESSPOLICY SINGLETASK
# allow nodes 13 and 14 to be shared
NODECFG[node13] ACCESS=SHARED
NODECFG[node14] ACCESS=SHARED
```

### Related Topics

- [mnodectl](#)

## 10.5 Managing Shared Cluster Resources (Floating Resources)

This section describes how to configure, request, and reserve cluster file system space and bandwidth, [software licenses](#), and generic cluster resources.

In this topic:

- [10.5.1 Shared Cluster Resource Overview - page 602](#)
- [10.5.2 Configuring Generic Consumable Floating Resources - page 603](#)
- [10.5.3 Configuring Cluster File Systems - page 604](#)
- [10.5.4 Configuring Cluster Licenses - page 604](#)
- [10.5.5 Configuring Generic Resources as Features - page 604](#)
- [10.5.6 Configuring Generic Resources as Licenses - page 606](#)

### 10.5.1 Shared Cluster Resource Overview

Shared cluster resources such as file systems, networks, and licenses can be managed through creating a pseudo-node. You can configure a pseudo-node via the [NODECFG](#)

parameter much as a normal node would be but additional information is required to allow the scheduler to contact and synchronize state with the resource.

In the following example, a license manager is added as a cluster resource by defining the GLOBAL pseudo-node and specifying how the scheduler should query and modify its state.

```
NODECFG[GLOBAL] RMLIST=NATIVE
NODECFG[GLOBAL] QUERYCMD=/usr/local/bin/flquery.sh
NODECFG[GLOBAL] MODIFYCMD=/usr/local/bin/flmodify.sh
```

In some cases, pseudo-node resources might be very comparable to node-locked [generic resources](#) however there are a few fundamental differences that determine when one method of describing resources should be used over the other. The following table contrasts the two resource types:

Attribute	Pseudo-Node	Generic Resource
<b>Node-Locked</b>	No - Resources can be encapsulated as an independent node.	Yes - Must be associated with an existing compute node.
<b>Requires exclusive batch system control over resource</b>	No - Resources (such as file systems and licenses) can be consumed both inside and outside of batch system workload.	Yes - Resources must only be consumed by batch workload. Use outside of batch control results in loss of resource synchronization.
<b>Allows scheduler level allocation of resources</b>	Yes - If required, the scheduler can take external administrative action to allocate the resource to the job.	No - The scheduler can only maintain logical allocation information and cannot take any external action to allocate resources to the job.

## 10.5.2 Configuring Generic Consumable Floating Resources

Consumable floating resources are configured in the same way as node-locked [generic resources](#) with the exception of using the GLOBAL node instead of a particular node.

```
NODECFG[GLOBAL] GRES=tape:4,matlab:2
...
```

*In this setup, four resources of type `tape` and 2 of type `matlab` are floating and available across all nodes.*

### 10.5.2.A Requesting Consumable Floating Resources

Floating resources are requested on a per task basis using native resource manager job submission methods or using the [GRES](#) resource manager extensions.

### 10.5.3 Configuring Cluster File Systems

Moab allows both the file space and bandwidth attributes of a cluster file system to be tracked, reserved, and scheduled. With this capability, a job or reservation can request a particular quantity of file space and a required amount of I/O bandwidth to this file system. While file system resources are managed as a cluster generic resource, they are specified using the `FS` attribute of the `NODECFG` parameter as in the following example:

```
NODECFG[GLOBAL] FS=PV1:10000@100,PV2:5000@100
...
```

*In this example, PV1 defines a 10 GB file system with a maximum throughput of 100 MB/s while PV2 defines a 5 GB file system also possessing a maximum throughput of 100 MB/s.*

A job can request cluster file system resources using the `fs` resource manager extension. For a Torque based system, the following could be used:

```
>qsub -l nodes=1,walltime=1:00:00 -W x=fs:10@50
```

### 10.5.4 Configuring Cluster Licenses

Jobs can request and reserve software licenses using native methods or using the `GRES` resource manager extension. If the cluster license manager does not support a query interface, license availability can be specified within Moab using the `GRES` attribute of the `NODECFG` parameter.

*Example 10-5: Configure Moab to support four floating `quickcalc` and two floating `matlab` licenses.*

```
NODECFG[GLOBAL] GRES=quickcalc:4,matlab:2
...
```

*Example 10-6: Submit a Torque job requesting a node-locked or floating `quickcalc` license.*

```
> qsub -l nodes=1,software=quickcalc,walltime=72000 testjob.cmd
```

### 10.5.5 Configuring Generic Resources as Features

Moab can be configured to treat generic resources as features in order to provide more control over server access. For instance, if a node is configured with a certain `GRES` and that `GRES` is turned off, jobs requesting the node will not run. To turn a `GRES` into a feature, set the `FEATUREGRES` attribute of `GRES_CFG` to `TRUE` in the `moab.cfg` file.

```
GRES_CFG[gres1] FEATUREGRES=TRUE
```

*Moab now treats `gres1` as a scheduler-wide feature rather than a normal generic resource.*

Note that jobs are submitted normally using the same `GRES` syntax.



**i** If you are running a grid, verify that `FEATUREGRES=TRUE` is set on all members of the grid.

**i** You can safely upgrade an existing cluster to use the feature while jobs are running. If you are in a grid, upgrade all clusters at the same time.

Two methods exist for managing GRES features: via Moab commands and via the resource manager. Using Moab commands means that feature changes are not checkpointed; they do not remain in place when Moab restarts. Using the resource manager causes changes to be reported by the RM, so any changes made before a Moab restart are still present after it.

These methods are mutually exclusive. Use one or the other, but do not mix methods.

### 10.5.5.A Managing Feature GRES via Moab Commands

In the following example, `gres1` and `gres2` are configured in the `moab.cfg` file. `gres1` is not currently functioning correctly, so it is set to 0, turning the feature off. Values above 0 and non-specified values turn the feature on.

```
NODECFG[GLOBAL] GRES=gres1:0
NODECFG[GLOBAL] GRES=gres2:10000
GRESCFG[gres1] FEATUREGRES=TRUE
GRESCFG[gres2] FEATUREGRES=TRUE
```

*Moab now treats gres1 and gres2 as features.*

To verify that this is set up correctly, run `mdiag -S -v`. It returns the following:

```
> mdiag -S -v
...
Scheduler FeatureGres: gres1:off,gres2:on
```

Once Moab has started, use `mschedctl -m` to modify whether the feature is turned on or off.

```
mschedctl -m sched featuregres:gres1=on
INFO: FeatureGres 'gres1' turned on
```

You can verify that the feature turned on or off by once again running `mdiag -S -v`.

**i** If Moab restarts, it will not checkpoint the state of these changed feature generic resources. Instead, it will read the `moab.cfg` file to determine whether the feature GRES is on or off.

With feature GRES configured, jobs are submitted normally, requesting GRES type `gres1` and `gres2`. Moab ignores GRES counts and reads the feature simply as on or off.

```

> msub -l nodes=1,walltime=600,gres=gres1

1012
> checkjob 1012
job 1012

AName: STDIN
State: Running
.....
StartTime: Tue Jul 3 15:33:28
Feature GRES: gres1
Total Requested Tasks: 1

```

If you request a feature that is currently turned off, the state is not reported as `Running`, but as `Idle`. A message such as the following returns:

```
BLOCK MSG: requested feature gres 'gres2' is off
```

### 10.5.5.B Managing Feature GRES via the Resource Manager

You can automate the process of having a feature GRES turn on and off by setting up an external tool and configuring Moab to query the tool the same way that Moab queries a license manager. For example:

```

RMCFG[myRM] CLUSTERQUERYURL=file:/// $HOME/tools/myRM.dat TYPE=NATIVE
RESOURCETYPE=LICENSE

GRESCFG[gres1] FEATUREGRES=TRUE
GRESCFG[gres2] FEATUREGRES=TRUE

```

*LICENSE means that the RM does not contain any compute resources and that Moab should not attempt to use it to manage any jobs (start, cancel, submit, etc.).*

The `myRM.dat` file should contain something such as the following:

```
GLOBAL state=Idle cres=gres1:0,gres2:10
```

External tools can easily update the file based on filesystem availability. Switching any of the feature GRES to 0 turns it off and switching it to a positive value turns it on. If you use this external mechanism, you do not need to use `mschedctl -m` to turn a feature GRES on or off. You also do not need to worry about whether Moab has checkpointed the information or not, since the information is provided by the RM and not by any external commands.

### 10.5.6 Configuring Generic Resources as Licenses

Moab can be configured to treat generic resources as licenses in order to distinguish them as licenses in terms of tracking and charging with the accounting manager. To turn a GRES into a license, set the `LICENSE` attribute of `GRESCFG` to `TRUE` in the `moab.cfg` file.

For example:

```
GRESCFG[matlab] LICENSE=TRUE
```

*Moab will pass the matlab generic resource to the accounting manager in the Licenses property rather than the Resources property.*

## Related Topics

- [Managing Resources Directly with the Native Interface](#)

## 10.6 Managing Node State

There are multiple models where Moab can operate allowing it to either honor the node state set by an external service or locally determine and set the node state. This section covers the following:

- Identifying meanings of particular node states
- Specifying node states within locally developed services and resource managers
- Adjusting node state within Moab based on load, policies, and events

In this topic:

[10.6.1 Node State Definitions - page 607](#)

[10.6.2 Specifying Node States within Native Resource Managers - page 608](#)

[10.6.3 Moab Based Node State Adjustment - page 608](#)

[10.6.4 Adjusting Scheduling Behavior Based on Reported Node State - page 608](#)

[10.6.5 Adding or Removing Nodes - page 609](#)

### 10.6.1 Node State Definitions

State	Definition
<b>Down</b>	Node is either not reporting status, is reporting status but failures are detected, or is reporting status but has been marked down by an administrator.
<b>Idle</b>	Node is reporting status, currently is not executing any workload, and is ready to accept additional workload.
<b>Busy</b>	Node is reporting status, currently is executing workload, and cannot accept

State	Definition
	additional workload due to load.
<b>Running</b>	Node is reporting status, currently is executing workload, and can accept additional workload.
<b>Drained</b>	Node is reporting status, currently is not executing workload, and cannot accept additional workload due to administrative action.
<b>Draining</b>	Node is reporting status, currently is executing workload, and cannot accept additional workload due to administrative action.

## 10.6.2 Specifying Node States within Native Resource Managers

Native resource managers can report node state implicitly and explicitly, using `NODESTATE`, `LOAD`, and other attributes. See [Managing Resources Directly with the Native Interface](#) for more information.

## 10.6.3 Moab Based Node State Adjustment

Node state can be adjusted based on reported processor, memory, or other load factors. It can also be adjusted based on reports of one or more resource managers in a multi-resource manager configuration. Also, both generic events and generic metrics can be used to adjust node state.

- Torque health scripts (allow compute nodes to detect and report site specific failures).

## 10.6.4 Adjusting Scheduling Behavior Based on Reported Node State

Based on reported node state, Moab can support various policies to make better use of available resources. For more information, see the [Green Computing](#).

### 10.6.4.A Down State

- `JOBACTIONONNODEFAILURE` parameter (cancel/requeue jobs if allocated nodes fail).
- `Triggers` (take specified action if failure is detected).

## 10.6.5 Adding or Removing Nodes

When a node has been deleted by a resource manager and the resource manager no longer reports data for the node, the node continues to exist in Moab until the next restart.

As a best practice, Adaptive Computing recommends adding or removing nodes only during cluster maintenance, rather than during periods of production activity. A restart of Moab must follow the addition and/or removal of nodes. This guarantees that Moab will handle nodes in a reliable, predictable way. If you want to remove nodes from service, but cannot immediately restart Moab after doing so, we recommend marking the nodes offline (for example, with `pbsnodes -o <nodeID>` or `mnodectl -m state=down <nodeID>`) and/or placing an administrative reservation over the nodes, until such time as you can follow the recommended removal procedure during a planned maintenance window.

---

### Related Topics

- [Managing Resources Directly with the Native Interface](#)
- [License Management](#)
- [Adjusting Node Availability](#)
- `NODEMAXLOAD` parameter
- [Green computing overview](#)

## 10.7 Managing Consumable Generic Resources

Each time a job is allocated to a compute node, it consumes one or more types of resources. Standard resources such as CPU, memory, disk, network adapter bandwidth, and swap are automatically tracked and consumed by Moab. However, in many cases, additional resources may be provided by nodes and consumed by jobs that must be tracked. The purpose of this tracking may include accounting, billing, or the prevention of resource over-subscription. Generic consumable resources can be used to manage software licenses, I/O usage, bandwidth, application connections, or any other aspect of the larger compute environment; they can be associated with compute nodes, networks, storage systems, or other real or virtual resources.

These additional resources can be managed within Moab by defining one or more generic resources. The first step in defining a generic resource involves naming the resource. Generic resource availability can then be associated with various compute nodes and generic resource usage requirements can be associated with jobs.

In this topic:

[10.7.1 Differences Between Node Features and Consumable Resources - page 610](#)

[10.7.2 Configuring Node-locked Consumable Generic Resources - page 610](#)

[10.7.3 Managing Generic Resource Race Conditions - page 612](#)

## 10.7.1 Differences Between Node Features and Consumable Resources

A [node feature](#) (or node property) is an opaque string label that is associated with a compute node. Each compute node can have any number of node features assigned to it and jobs can request allocation of nodes that have specific features assigned. Node features are labels and their association with a compute node is not conditional, meaning they cannot be consumed or exhausted.

## 10.7.2 Configuring Node-locked Consumable Generic Resources

Consumable generic resources are supported within Moab using either direct configuration or resource manager auto-detect (as when using Torque and [accelerator hardware](#)). For direct configuration, node-locked consumable generic resources (or generic resources) are specified using the `NODECFG` parameter's `GRES` attribute. This attribute is specified using the format `<ATTR> : <COUNT>` as in the following example:

```
NODECFG[titan001] GRES=tape:4
NODECFG[login32] GRES=matlab:2,prime:4
NODECFG[login33] GRES=matlab:2
...
```

**i** By default, Moab supports up to 128 independent generic resource types.

### 10.7.2.A Requesting Consumable Generic Resources

Generic resources can be requested on a per task or per job basis using the [GRES resource manager extension](#). If the generic resource is located on a compute node, requests are by default interpreted as a per task request. If the generic resource is located on a shared, cluster-level resource (such as a network or storage system), then the request defaults to a per job interpretation.

**i** Generic resources are specified per task, not per node. When you submit a job, each processor becomes a task. For example, a job asking for `nodes=3:ppn=4,gres=test:5` asks for 60 gres of type test  $((3*4 \text{ processors}) * 5)$ .

If using Torque, the **GRES** or **software** resource can be requested as in the following examples:

#### Example 10-7: Per Task Requests

```
NODECFG[compute001] GRES=dvd:2 SPEED=2200
NODECFG[compute002] GRES=dvd:2 SPEED=2200
NODECFG[compute003] GRES=dvd:2 SPEED=2200
NODECFG[compute004] GRES=dvd:2 SPEED=2200
NODECFG[compute005] SPEED=2200
NODECFG[compute006] SPEED=2200
NODECFG[compute007] SPEED=2200
NODECFG[compute008] SPEED=2200
```

```
# submit job which will allocate only from nodes 1 through 4 requesting one dvd per
task
> qsub -l nodes=2,walltime=100,gres=dvd job.cmd
```

*In this example, Moab determines that compute nodes exist that possess the requested generic resource. A compute node is a node object that possesses processors on which compute jobs actually execute. License server, network, and storage resources are typically represented by non-compute nodes. Because compute nodes exist with the requested generic resource, Moab interprets this job as requesting two compute nodes each of which must also possess a DVD generic resource.*

#### Example 10-8: Per Job Requests

```
NODECFG[network] PARTITION=shared GRES=bandwidth:2000000
```

```
# submit job which will allocate 2 nodes and 10000 units of network bandwidth
> qsub -l nodes=2,walltime=100,gres=bandwidth:10000 job.cmd
```

*In this example, Moab determines that there exist no compute nodes that also possess the generic resource bandwidth so this job is translated into a multiple-requirement—multi-req—job. Moab creates a job that has a requirement for two compute nodes and a second requirement for 10000 bandwidth generic resources. Because this is a multi-req job, Moab knows that it can locate these needed resources separately.*

### 10.7.2.B Using Generic Resource Requests in Conjunction with other Constraints

Jobs can explicitly specify generic resource constraints. However, if a job also specifies a hostlist, the hostlist constraint overrides the generic resource constraint if the request is for per task allocation. In the Per Task Requests example, if the job also specified a hostlist, the DVD request is ignored.

### 10.7.2.C Requesting Resources with No Generic Resources

In some cases, it is valuable to allocate nodes that currently have no generic resources available. This can be done using the special value `none` as in the following example:

```
> qsub -l nodes=2,walltime=100,gres=none job.cmd
```

*In this case, the job only allocates compute nodes that have no generic resources associated with them.*

### 10.7.2.D Requesting Generic Resources Automatically within a Queue/Class

Generic resource constraints can be assigned to a queue or class and inherited by any jobs that do not have a `gres` request. This allows targeting of specific resources, automation of co-allocation requests, and other uses. To enable this, use the `DEFAULT.GRES` attribute of the `CLASSCFG` parameter as in the following example:

```
CLASSCFG[viz] DEFAULT.GRES=graphics:2
```

For each node requested by a `viz` job, also request two graphics cards.

## 10.7.3 Managing Generic Resource Race Conditions

A software license race condition 'window of opportunity' opens when Moab checks a license server for sufficient available licenses and closes when the user's software actually checks out the software licenses. The time between these two events can be seconds to many minutes depending on overhead factors such as node OS provisioning, job startup, licensed software startup, and so forth.

During this window, another Moab-scheduled job or a user or job external to the cluster or cloud can obtain enough software licenses that by the time the job attempts to obtain its software licenses, there are an insufficient quantity of available licenses. In such cases a job will sit and wait for the license, and while it waits it occupies but does not use resources that another job could have used. Use the `STARTDELAY` parameter to prevent such a situation.

```
GRESCFG[<license>] STARTDELAY=<window_of_opportunity>
```

With the `STARTDELAY` parameter enabled (on a per generic resource basis) Moab blocks any idle jobs requesting the same generic resource from starting until the `<window_of_opportunity>` passes. The window is defined by the customer on a per generic resource basis.



---

## Related Topics

- [GRES CFG parameter](#)
- [Generic Events](#)
- [General Node Attributes](#)
- [Floating Generic Resources](#)
- [Per Class Assignment of Generic Resource Consumption](#)
- [mnodectl -m command to dynamically modify node resources](#)
- [Favoring Jobs Based On Generic Resource Requirements](#)

## 10.8 Enabling Generic Metrics

Moab allows organizations to enable generic performance metrics. These metrics allow decisions to be made and reports to be generated based on site specific environmental factors. This increases Moab's awareness of what is occurring within a given cluster environment, and allows arbitrary information to be associated with resources and the workload within the cluster. Uses of these metrics are widespread and can cover anything from tracking node temperature, to memory faults, to application effectiveness.

- Execute triggers when specified thresholds are reached
- Modify node allocation affinity for specific jobs
- Initiate automated notifications when thresholds are reached
- Display current, average, maximum, and minimum metrics values in reports and charts within Moab Cluster Manager

In this topic:

[10.8.1 Configuring Generic Metrics - page 613](#)

[10.8.2 Example Generic Metric Usage - page 615](#)

### 10.8.1 Configuring Generic Metrics

A new generic metric is automatically created and tracked at the server level if it is reported by either a node or a job.

To associate a generic metric with a job or node, a [native resource manager](#) must be set up and the GMETRIC attribute must be specified. For example, to associate a generic metric of

`temp` with each node in a Torque cluster, the following could be reported by a native resource manager:

```
# temperature output
node001 GMETRIC[temp]=113
node002 GMETRIC[temp]=107
node003 GMETRIC[temp]=83
node004 GMETRIC[temp]=85
...
```

**i** Generic metrics are tracked as floating point values allowing virtually any number to be reported.

In the preceding example, the new metric, `temp`, can now be used to monitor system usage and performance or to allow the scheduler to take action should certain thresholds be reached. Some uses include the following:

- Executing [triggers](#) based on generic metric thresholds
- Adjust a node's [availability](#) for accepting additional workload
- Adjust a node's [allocation priority](#)
- Initiate administrator [notification](#) of current, minimum, maximum, or average generic metric values
- Use metrics to report resource and job performance
- Use metrics to report resource and job failures
- Using job profiles to allow Moab to learn which resources best run which applications
- Tracking effective application efficiency to identify resource brown outseven when no node failure is obvious
- Viewing current and [historical](#) cluster-wide generic metric values to identify failure, performance, and usage
- Enable charging policies based on consumption of generic metrics patterns
- View changes in generic metrics on nodes, jobs, and cluster wide over time
- Submit jobs with generic metric based [node-allocation requirements](#)

Generic metric values can be viewed using [checkjob](#), [checknode](#), [mdiag -n](#), [mdiag -j](#), or Moab Cluster Manager Charting and Reporting Features.

**i** Historical job and node generic metric statistics can be cleared using the [mjobctl](#) and [mnodectl](#) commands.

## 10.8.2 Example Generic Metric Usage

As an example, consider a cluster with two primary purposes for generic metrics. The first purpose is to track and adjust scheduling behavior based on node temperature to mitigate overheating nodes. The second purpose is to track and charge for utilization of a locally developed data staging service.

The first step in enabling a generic metric is to create probes to monitor and report this information. Depending on the environment, this information may be distributed or centralized. In the case of temperature monitoring, this information is often centralized by a hardware monitoring service and available via command line or an API. If monitoring a locally developed data staging service, this information might need to be collected from multiple remote nodes and aggregated to a central location. The following are popular freely available monitoring tools:

Tool	Link
<b>Ganglia</b>	<a href="https://ganglia.sourceforge.net">https://ganglia.sourceforge.net</a>
<b>Monit</b>	<a href="https://www.tildeslash.com/monit">https://www.tildeslash.com/monit</a>
<b>Nagios</b>	<a href="https://www.nagios.org">https://www.nagios.org</a>

Once the needed probes are in place, a [native resource manager](#) interface must be created to report this information to Moab. Creating a native resource manager interface should be very simple, and in most cases a script similar to those found in the `$TOOLS_DIR` (`$PREFIX/tools`) directory can be used as a template. For this example, we will assume centralized information and will use the RM script that follows:

```
#!/usr/bin/perl
# 'hwctl outputs information in format '<NODEID> <TEMP>'
open(TQUERY,"/usr/sbin/hwctl -q temp |");
while (<TQUERY>)
{
    my $nodeid,$temp = split /\w+;/;
    $dstage=GetDSUsage($nodeid);
    print "$nodeid GMETRIC[temp]=$temp GMETRIC[dstage]=$dstage
";
}
```

With the script complete, the next step is to integrate this information into Moab. This is accomplished with the following configuration line:

```
RMCFG[local] TYPE=NATIVE CLUSTERQUERYURL=file://$TOOLS_DIR/node.query.local.pl
...
```

*Moab can now be recycled and temperature and data staging usage information will be integrated into Moab compute node reports.*

If the `checknode` command is run, output similar to the following is reported:

```
> checknode cluster013
...
Generic Metrics:  temp=113.2,dstage=23748
...
```

*Moab Cluster Manager reports full current and historical generic metric information in its visual cluster overview screen.*

The next step in configuring Moab is to inform Moab to take certain actions based on the new information it is tracking. For this example, there are two purposes. The first purpose is to get jobs to avoid hot nodes when possible. This is accomplished using the `GMETRIC` attribute of the [Node Allocation Priority](#) function as in the following example:

```
NODEALLOCATIONPOLICY PRIORITY
NODECFG[DEFAULT] PRIORITYF=PRIORITY-10*GMETRIC[temp]
...
```

This simple priority function reduces the priority of the hottest nodes making such less likely to be allocated.

The example cluster is also interested in notifying administrators if the temperature of a given node ever exceeds a critical threshold. This is accomplished using a [trigger](#). The following line will send email to administrators any time the temperature of a node exceeds 120 degrees:

```
NODECFG[DEFAULT] TRIGGER=atype=mail,etype=threshold,threshold=gmetric
[temp]>120,action='warning: node $OID temp high'
...
```

---

## Related Topics

- [Generic Consumable Resources](#)
- [Object Variables](#)

## 10.9 Enabling Generic Events

Generic events are used to identify failures and other occurrences that Moab or other systems must be made aware. This information may result in automated resource recovery, notifications, adjustments to statistics, or changes in policy. Generic events also have the ability to carry an arbitrary human readable message that may be attached to associated objects or passed to administrators or external systems. Generic events typically signify the occurrence of a specific event as opposed to [generic metrics](#), which indicate a change in a measured value.

Using generic events, Moab can be configured to automatically address many failures and environmental changes improving the overall performance. Some sample events that sites might be interested in monitoring, recording, and taking action on include:

- Machine Room Status
  - Excessive Room Temperature
  - Power Failure or Power Fluctuation
  - Chiller Health
- Network File Server Status
  - Failed Network Connectivity
  - Server Hardware Failure
  - Full Network File System
- Compute Node Status
  - Machine Check Event (MCE)
  - Network Card (NIC) Failure
  - Excessive Motherboard/CPU Temperature
  - Hard Drive Failures

In this topic:

[10.9.1 Configuring Generic Events - page 617](#)

[10.9.2 Reporting Generic Events - page 620](#)

[10.9.3 Generic Events Attributes - page 621](#)



[10.9.4 Manually Creating Generic Events - page 621](#)

## 10.9.1 Configuring Generic Events

Generic events are defined in the `moab.cfg` file and have several different configuration options. The only required option is `action`.

The full list of configurable options for generic events is contained in the following table:

Attribute	Description
<b>ACTION</b>	Comma-delimited list of actions to be processed when a new event is received.

Attribute	Description
<b>ECOUNT</b>	Number of events that must occur before launching action.  <div style="border: 1px solid #005596; border-radius: 10px; padding: 5px; background-color: #e6f2ff;">  Action will be launched each &lt;ECOUNT&gt; event if rearm is set. </div>
<b>REARM</b>	Minimum time between events specified in [ [ [DD:]HH:]MM:]SS format.
<b>SEVERITY</b>	An arbitrary severity level from 1 through 4, inclusive. SEVERITY appears in the output of <code>mdiag -n -v -v --xml</code> .  <div style="border: 1px solid #005596; border-radius: 10px; padding: 5px; background-color: #e6f2ff;">  The severity level will not be used for any other purpose. </div>

### 10.9.1.A Action Types

The impact of the event is controlled using the `ACTION` attribute of the `GEVENTCFG` parameter. The `ACTION` attribute is comma-delimited and can include any combination of the actions in the following table:

Value	Description
<b>DISABLE</b> [:<OTYPE>:<OID>]	Marks event object (or specified object) down until event report is cleared.
<b>EXECUTE</b>	Executes a script at the provided path. The value of <code>EXECUTE</code> is not contained in quotation marks. Arguments are allowed at the end of the path and are separated by question marks (?). Trigger variables (such as <code>\$OID</code> ) are allowed.
<b>NOTIFY</b>	Notifies administrators of the event occurrence.
<b>OBJECTXMLSTDIN</b>	If the <code>EXECUTE</code> action type is also specified, this flag passes an XML description of the firing event to the script.
<b>OFF</b>	Powers off node or resource.
<b>ON</b>	Powers on node or resource.
<b>PREEMPT</b> [:<POLICY>]	Preempts workload associated with object (valid for node, job, reservation, partition, resource manager, user, group, account, class, QoS, and cluster objects).

Value	Description
<b>RECORD</b>	Records events to the event log. The record action causes a line to be added to the event log regardless of whether or not <code>RECORDEVENTLIST</code> includes <code>GEVENT</code> .
<b>RESERVE</b> <b>[:&lt;DURATION&gt;]</b>	Reserves node for specified duration (default: 24 hours).
<b>RESET</b>	Resets object (valid for nodes - causes reboot).
<b>SIGNAL</b> [:<SIGNO>]	Sends signal to associated jobs or services (valid for node, job, reservation, partition, resource manager, user, group, account, class, QoS, and cluster objects).

This is an example of using `objectxmlstdin` with a gevent:

```
<gevent name="bob" statuscode="0" time="1320334763">Testing</gevent>
```

### 10.9.1.B Named Events

In general, generic events are named, with the exception of those based on [generic metrics](#). Names are used primarily to differentiate between different events and do not have any intrinsic meaning to Moab. It is suggested that the administrator choose names that denote specific meanings within the organization.

*Example 10-9:*

```
# Note: cpu failures require admin attention, create maintenance reservation
GEVENTCFG[cpufail] action=notify,record,disable,reserve rearm=01:00:00# Note: power
failures are transient, minimize future use
GEVENTCFG[powerfail] action=notify,record, rearm=00:05:00
# Note: fs full can be automatically fixed
GEVENTCFG[fsfull] action=notify,execute:/home/jason/MyPython/cleartmp.py?${OID}?nodefix
# Note: memory errors can cause invalid job results, clear node immediately
GEVENTCFG[badmem] action=notify,record,preempt,disable,reserve
```

### 10.9.1.C Generic Metric (GMetric) Events

GMetric events are generic events based on [generic metrics](#). They are used for executing an action when a generic metric passes a defined threshold. Unlike named events, GMetric events are not named and use the following format:

```
GEVENTCFG [GMETRIC<COMPARISON>VALUE] ACTION=...
```

*Example 10-10:*

```
GEVENTCFG[cputemp>150] action=off
```

This form of generic events uses the GMetric name, as returned by a GMETRIC attribute in a [native Resource Manager](#) interface.

**i** Only one generic event can be specified for any given generic metric.

Valid comparative operators are shown in the following table:

Type	Comparison	Notes
>	greater than	Numeric values only
> =	greater than or equal to	Numeric values only
= =	equal to	Numeric values only
<	less than	Numeric values only
< =	less than or equal to	Numeric values only
< >	not equal	Numeric values only

## 10.9.2 Reporting Generic Events

Unlike [generic metrics](#), generic events can be optionally configured at the global level to adjust rearm policies, and other behaviors. In all cases, this is accomplished using the [GEVENTCFG](#) parameter.

To report an event associated with a job or node, use the [native Resource Manager](#) interface or the [mjobctl](#) or [mnodectl](#) commands. You can report generic events on the scheduler with the [mschedctl](#) command.

If using the native Resource Manager interface, use the GEVENT attribute as in the following example:

```
node001 GEVENT[hitemp]='temperature exceeds 150 degrees'
node017 GEVENT[fullfs]='/var/tmp is full'
```



**i** The time at which the event occurred can be passed to Moab to prevent multiple processing of the same event. This is accomplished by specifying the event type in the format `<GEVENTID> [:<EVENTTIME>]` as in what follows:

```
node001 GEVENT[hitemp:1130325993]='temperature exceeds 150 degrees'
node017 GEVENT[fullfs:1130325142]='/var/tmp is full'
```

### 10.9.3 Generic Events Attributes

Each node will record the following about reported generic events:

- status - is event active
- message - human readable message associated with event
- count - number of event incidences reported since statistics were cleared
- time - time of most recent event

Each event can be individually cleared, annotated, or deleted by cluster administrators using a `mnodectl` command.

**i** Generic events are only available in Moab 4.5.0 and later.

### 10.9.4 Manually Creating Generic Events

Generic events can be manually created on a physical node.

To add GEVENT event with message "hello" to node02, do the following:

```
> mnodectl -m gevent=event:"hello" node02
```

#### Related Topics

- [Generic Consumable Resources](#)
- [Object Variables](#)

## Chapter 11: Resource Managers and Interfaces

In this chapter:

11.1 Resource Manager Overview .....	624
11.1.1 Scheduler/Resource Manager Interactions .....	625
11.1.2 Resource Manager Specific Details (Limitations/Special Features) .....	626
11.1.3 Synchronizing Conflicting Information .....	626
11.1.4 Evaluating Resource Manager Availability and Performance .....	627
11.2 Resource Manager Configuration .....	627
11.2.1 Defining and Configuring Resource Manager Interfaces .....	627
11.2.2 Resource Manager Configuration Details .....	653
11.2.3 Scheduler/Resource Manager Interactions .....	658
11.3 Resource Manager Extensions .....	659
11.3.1 Resource Manager Extension Specification .....	660
11.3.2 Resource Manager Extension Values .....	660
11.3.3 Resource Manager Extension Examples .....	687
11.3.4 Configuring dynamic features in Torque and Moab .....	688
11.4 Adding New Resource Manager Interfaces .....	689
11.4.1 Resource Manager Specific Interfaces .....	689
11.4.2 Wiki Interface .....	689
11.4.3 SSS Interface .....	689
11.5 Managing Resources Directly with the Native Interface .....	690
11.5.1 Native Interface Overview .....	690
11.5.2 Configuring the Native Interface .....	691
11.5.3 Generating Cluster Query Data .....	692
11.5.4 Interfacing with FlexNet (Formerly FLEXlm) .....	692
11.5.5 Interfacing to Nagios .....	694
11.5.6 Configuring Resource Types .....	695
11.5.7 Creating New Tools to Manage the Cluster .....	696
11.6 Utilizing Multiple Resource Managers .....	701
11.6.1 Multi-RM Overview .....	701
11.6.2 Configuring Multiple Independent Resource Manager Partitions .....	701
11.6.3 Migrating Jobs between Resource Managers .....	701

11.6.4 Aggregating Information into a Cohesive Node View .....	702
11.7 License Management .....	703
11.7.1 License Management Overview .....	703
11.7.2 Controlling and Monitoring License Availability .....	704
11.7.3 Requesting Licenses within Jobs .....	705
11.8 Resource Provisioning .....	706
11.8.1 Resource Provisioning Overview .....	706
11.8.2 Configuring Provisioning .....	706
11.9 Managing Networks .....	707
11.9.1 Network Management Overview .....	707
11.9.2 Dynamic VLAN Creation .....	707
11.9.3 Network Load and Health Monitoring .....	708
11.9.4 Creating a Resource Management Interface for a New Network ..	709
11.9.5 Per-Job Network Monitoring .....	710
11.10 Intelligent Platform Management Interface .....	710
11.10.1 IPMI Overview .....	710
11.10.2 Node IPMI Configuration .....	711
11.10.3 Installing IPMItool .....	711
11.10.4 [Optional] Creating the IPMI BMC-Node Map File .....	711
11.10.5 Configuring the Moab IPMI Tools .....	712
11.10.6 Configuring Moab .....	712
11.10.7 Ensuring Proper Setup .....	713
11.11 Resource Manager Translation .....	714
11.11.1 Translation Overview .....	714
11.11.2 Translation Enablement Steps .....	714

Moab provides a powerful resource management interface that enables significant flexibility in how resources and workloads are managed. Highlights of this interface are listed in what follows:

Highlight	Description
<b>Support for Generic Resource Manager Interfaces</b>	Manage cluster resources securely via locally developed or open source projects using simple flat text interfaces or XML over HTTP.
<b>Support for Multiple</b>	Integrate resource and workload streams from multiple

Highlight	Description
<b>Simultaneous Resource Managers</b>	independent sources reporting disjoint sets of resources.
<b>Independent Workload and Resource Management</b>	Allow one system to manage your workload (queue manager) and another to manage your resources.
<b>Support for Rapid Development Interfaces</b>	Load resource and workload information directly from a file, a URL, or from the output of a configurable script or other executable.
<b>Resource Extension Information</b>	Integrate information from multiple sources to obtain a cohesive view of a compute resource. (That is, mix information from a resource manager and a cluster performance monitor to obtain a single node image with a coordinated state and a more extensive list of node configuration and utilization attributes.)

## 11.1 Resource Manager Overview

For most installations, the Moab Workload Manager uses the services of a resource manager to obtain information about the state of compute resources (nodes) and workload (jobs). Moab also uses the resource manager to manage jobs, passing instructions regarding when, where, and how to start or otherwise manipulate jobs.

Moab can be configured to manage more than one resource manager simultaneously, even resource managers of different types. Using a local queue, jobs can even be migrated from one resource manager to another. However, there are currently limitations regarding jobs submitted directly to a resource manager (not to the local queue.) In such cases, the job is constrained to only run within the bound of the resource manager to which it was submitted.

In this topic:

[11.1.1 Scheduler/Resource Manager Interactions - page 625](#)

[11.1.2 Resource Manager Specific Details \(Limitations/Special Features\) - page 626](#)

[11.1.3 Synchronizing Conflicting Information - page 626](#)

[11.1.4 Evaluating Resource Manager Availability and Performance - page 627](#)

## 11.1.1 Scheduler/Resource Manager Interactions

Moab interacts with all resource managers using a common set of commands and objects. Each resource manager interfaces, obtains, and translates Moab concepts regarding workload and resources into native resource manager objects, attributes, and commands.

Information on creating a new scheduler resource manager interface can be found in the [Adding New Resource Manager Interfaces](#) section.

### 11.1.1.A Resource Manager Commands

For many environments, Moab interaction with the resource manager is limited to the following objects and functions:

Object	Function	Details
Job	Query	Collect detailed state, requirement, and utilization information about jobs
	Modify	Change job state and/or attributes
	Start	Execute a job on a specified set of resources
	Cancel	Cancel an existing job
	Preempt/Resume	Suspend, resume, checkpoint, restart, or requeue a job
Node	Query	Collect detailed state, configuration, and utilization information about compute resources
	Modify	Change node state and/or attributes
Queue	Query	Collect detailed policy and configuration information from the resource manager

Using these functions, Moab is able to fully manage workload, resources, and cluster policies. More detailed information about resource manager specific capabilities and limitations for each of these functions can be found in the individual resource manager overviews (PBS or [WIKI](#)).

Beyond these base functions, other commands exist to support advanced features such as provisioning and cluster level resource management.

### 11.1.1.B Resource Manager Flow

In general, Moab interacts with resource managers in a sequence of steps each scheduling iteration. These steps are outlined in what follows:

1. load global resource information
2. load node specific information (optional)
3. load job information
4. load queue/policy information (optional)
5. cancel/preempt/modify jobs according to cluster policies
6. start jobs in accordance with available resources and policy constraints
7. handle user commands

Typically, each step completes before the next step is started. However, with current systems, size and complexity mandate a more advanced parallel approach providing benefits in the areas of reliability, concurrency, and responsiveness.

### 11.1.2 Resource Manager Specific Details (Limitations/Special Features)

- Torque
  - Torque Homepage
- Wiki
  - [Wiki Overview](#)

### 11.1.3 Synchronizing Conflicting Information

Moab does not trust resource manager information. Node, job, and policy information is reloaded on each iteration and discrepancies are detected. Synchronization issues and allocation conflicts are logged and handled where possible. To assist sites in minimizing stale information and conflicts, a number of policies and parameters are available.

- Node State Synchronization Policies
- Stale Data Purging (see [JOBPURGETIME](#))
- Thread Management (preventing resource manager failures from affecting scheduler operation)
- Resource Manager Poll Interval (see [RMPOLLINTERVAL](#))

## 11.1.4 Evaluating Resource Manager Availability and Performance

Each resource manager is individually tracked and evaluated by Moab. Using the `mdiag -R` command, a site can determine how a resource manager is configured, how heavily it is loaded, what failures, if any, have occurred in the recent past, and how responsive it is to requests.

### Related Topics

- [Resource Manager Configuration](#)
- [Resource Manager Extensions](#)

## 11.2 Resource Manager Configuration

In this topic:

[11.2.1 Defining and Configuring Resource Manager Interfaces - page 627](#)

[11.2.2 Resource Manager Configuration Details - page 653](#)

[11.2.3 Scheduler/Resource Manager Interactions - page 658](#)

### 11.2.1 Defining and Configuring Resource Manager Interfaces

Moab resource manager interfaces are defined using the `RMCFG` parameter. This parameter allows specification of key aspects of the interface. In most cases, only the `TYPE` attribute needs to be specified and Moab determines the needed defaults required to activate and use the selected interface. In the following example, an interface to a Loadleveler resource manager is defined:

```
RMCFG[orion] TYPE=LL...
```

Note that the resource manager is given a label of `orion`. This label can be any arbitrary site-selected string and is for local usage only. For sites with multiple active resource managers, the labels can be used to distinguish between them for resource manager specific queries and commands.

#### 11.2.1.A Resource Manager Attributes


The following table lists the possible resource manager attributes that can be configured:

*Resource Manager Attributes*

ADMINEXEC	JOBMODIFYURL	SBINDIR
AUTHTYPE	JOBRSVRECREATE	SERVER
BANDWIDTH	JOBSTARTURL	STAGETHRESHOLD
CHECKPOINTSIG	JOBSUBMITURL	STARTCMD
CHECKPOINTTIMEOUT	JOBSUSPENDURL	SUBMITCMD
CLIENT	JOBVALIDATEURL	SUBMITPOLICY
CLUSTERQUERYURL	MAXDSOP	SUSPENDSIG
CONFIGFILE	MAXITERATIONFAILURECOU	SYNCJOBID
DEFAULTCLASS	NT	SYSTEMMODIFYURL
DEFAULTHIGHSPEEDADAPT	MAXJOBPERMINUTE	SYSTEMQUERYURL
ER	MAXJOBS	TARGETUSAGE
DESCRIPTION	MINETIME	TIMEOUT
ENV	NMPORT	TRIGGER
EPORT	NODEFAILURERSVPROFILE	TYPE
FAILTIME	NODESTATEPOLICY	VARIABLES
FBSERVER	OMAP	VERSION
FLAGS	PORT	WORKLOADQUERYU
FNLIST	PROVDURATION	RL
HOST	RESOURCECREATEURL	
JOBCANCELURL	RESOURCECTYPE	
JOBEXTENDDURATION	RMSTARTURL	
	RMSTOPURL	

ADMINEXEC	
<b>Format</b>	"jobsubmit"
<b>Default</b>	NONE
<b>Description</b>	Normally, when the JOBSUBMITURL is executed, Moab will drop to the UID and GID of the user submitting the job. Specifying an ADMINEXEC of <code>jobsubmit</code> causes Moab to use its own UID and GID instead (usually root). This is useful for some native resource managers where the JOBSUBMITURL is not a user command (such as <code>qsub</code> ) but a script that interfaces directly with the resource manager.
<b>Example</b>	<pre>RMCFG[base] ADMINEXEC=jobsubmit</pre> <p><i>Moab will not use the user's UID and GID for executing the JOBSUBMITURL.</i></p>



AUTHTYPE	
<b>Format</b>	One of CHECKSUM, OTHER, PKI, SECUREPORT, or NONE
<b>Default</b>	CHECKSUM
<b>Description</b>	<p>The security protocol to be used in scheduler-resource manager communication.</p> <div style="border: 1px solid #0070C0; border-radius: 5px; padding: 5px; margin-top: 10px;">  Only valid with WIKI based interfaces.         </div>
<b>Example</b>	<div style="border: 1px dashed #000; border-radius: 10px; padding: 10px; margin-bottom: 5px;"> <pre>RMCFG [base] AUTHTYPE=CHECKSUM</pre> </div> <div style="border: 1px dashed #000; border-radius: 10px; padding: 10px;"> <p><i>Moab requires a secret key-based checksum associated with each resource manager message.</i></p> </div>

BANDWIDTH	
<b>Format</b>	<FLOAT> [ {M   G   T} ]
<b>Default</b>	-1 (unlimited)
<b>Description</b>	<p>The maximum deliverable bandwidth between the Moab server and the resource manager for staging jobs and data. Bandwidth is specified in units per second and defaults to a unit of MB/s. If a unit modifier is specified, the value is interpreted accordingly (M - megabytes/sec, G - gigabytes/sec, T - terabytes/sec).</p>
<b>Example</b>	<div style="border: 1px dashed #000; border-radius: 10px; padding: 10px; margin-bottom: 5px;"> <pre>RMCFG [base] BANDWIDTH=340G</pre> </div> <div style="border: 1px dashed #000; border-radius: 10px; padding: 10px;"> <p><i>Moab will reserve up to 340 GB of network bandwidth when scheduling job and data staging operations to and from this resource manager.</i></p> </div>

CHECKPOINTSIG	
<b>Format</b>	One of suspend, <INTEGER> or SIG<X>

CHECKPOINTSIG	
<b>Description</b>	Specifies what signal to send the resource manager when a job is checkpointed. See <a href="#">Checkpoint Overview</a> .
<b>Example</b>	<pre>RMCFG[base] CHECKPOINTSIG=SIGKILL</pre> <p><i>Moab routes the signal SIGKILL through the resource manager to the job when a job is checkpointed.</i></p>

CHECKPOINTTIMEOUT	
<b>Format</b>	[ [ [DD:]HH:]MM:]SS
<b>Default</b>	0 (no timeout)
<b>Description</b>	Specifies how long Moab waits for a job to checkpoint before canceling it. If set to 0, Moab does not cancel the job if it fails to checkpoint. See <a href="#">Checkpoint Overview</a> .
<b>Example</b>	<pre>RMCFG[base] CHECKPOINTTIMEOUT=5:00</pre> <p><i>Moab cancels any job that has not exited 5 minutes after receiving a checkpoint request.</i></p>

CLIENT	
<b>Format</b>	<PEER>
<b>Default</b>	Use name of resource manager for peer client lookup
<b>Description</b>	If specified, the resource manager will use the peer value to authenticate remote connections. See configuring peers. If not specified, the resource manager will search for a <a href="#">CLIENTCFG[&lt;X&gt;]</a> entry of RM:<RMNAME> in the moab-private.cfg file.
<b>Example</b>	<pre>RMCFG[clusterBI] CLIENT=clusterB</pre>

CLIENT	
	Moab will look up and use information for peer <code>clusterB</code> when authenticating the <code>clusterBI</code> resource manager.

CLUSTERQUERYURL	
<b>Format</b>	[<protocol>://][<path>]
<b>Description</b>	Specifies how Moab queries the resource manager. See <a href="#">Native RM</a> , <a href="#">URL Notes</a> , and <a href="#">interface details</a> .
<b>Example</b>	<pre>RMCFG[base] CLUSTERQUERYURL=file:///tmp/cluster.config</pre> <p><i>Moab reads /tmp/cluster.config when it queries base resource manager.</i></p>

CONFIGFILE	
<b>Format</b>	<STRING>
<b>Description</b>	<p>The resource manager specific configuration file that must be used to enable correct API communication.</p> <div style="border: 1px solid black; padding: 5px; background-color: #e6f2ff;"> <p><b>i</b> Only valid with LL-based interfaces. This parameter is deprecated and may be removed in a future release.</p> </div>
<b>Example</b>	<pre>RMCFG[base] TYPE=LL CONFIGFILE=/home/loadl/loadl_config</pre> <p><i>The scheduler uses the specified file when establishing the resource manager/scheduler interface connection.</i></p>

DEFAULTCLASS	
<b>Format</b>	<STRING>
<b>Description</b>	The class to use if jobs submitted via this resource manager interface do not have an associated class.
<b>Example</b>	<pre>RMCFG[internal] DEFAULTCLASS=batch</pre> <p><i>Moab assigns the class <code>batch</code> to all jobs from the resource manager <code>internal</code> that do not have a class assigned.</i></p> <p><b>i</b> If you are using PBS as the resource manager, a job will never come from PBS without a class, and the default will never apply.</p>

DEFAULTHIGHSPEEDADAPTER	
<b>Format</b>	<STRING>
<b>Default</b>	sn0
<b>Description</b>	The default high speed switch adapter to use when starting LoadLeveler jobs (supported in version 4.2.2 and higher of Moab and 3.2 of LoadLeveler).
<b>Example</b>	<pre>RMCFG[base] DEFAULTHIGHSPEEDADAPTER=sn1</pre> <p><i>The scheduler will start jobs requesting a high speed adapter on <code>sn1</code>.</i></p>

DESCRIPTION	
<b>Format</b>	<STRING>
<b>Description</b>	The human-readable description for the resource manager interface. If white space is used, the description should be quoted.

DESCRIPTION	
<b>Example</b>	<pre>RMCFG[torque] DESCRIPTION='Torque RM for launching jobs'</pre> <p><i>Moab annotates the Torque resource manager accordingly.</i></p>

ENV	
<b>Format</b>	Semi-colon-delimited (;) list of <KEY>=<VALUE> pairs
<b>Default</b>	MOABHOMEDIR=<MOABHOMEDIR>
<b>Description</b>	Specifies a list of environment variables that will be passed to URLs of type <code>exec://</code> for that resource manager.
<b>Example</b>	<pre>RMCFG[base] ENV=HOST=node001;RETRYTIME=50 RMCFG[base] CLUSTERQUERYURL=exec:///opt/moab/tools/cluster.query.pl RMCFG[base] WORKLOADQUERYURL=exec:///opt/moab/tools/workload.query.pl</pre> <p><i>The environment variables HOST and RETRYTIME (with values node001 and 50 respectively) are passed to the /opt/moab/tools/cluster.query.pl and /opt/moab/tools/workload.query.pl when they are executed.</i></p>

EPORT	
<b>Format</b>	<INTEGER>
<b>Description</b>	The event port to use to receive resource manager based scheduling events.
<b>Example</b>	<pre>RMCFG[base] EPORT=15017</pre> <p><i>The scheduler will look for scheduling events from the resource manager host at port 15017.</i></p>

FAILTIME	
<b>Format</b>	[ [ [DD:] HH:] MM:] SS
<b>Description</b>	Specifies how long a resource manager must be down before any failure triggers associated with the resource manager fire.
<b>Example</b>	<pre>RMCFG[base] FAILTIME=3:00</pre> <p><i>If the base resource manager is down for three minutes, any resource manager failure triggers fire.</i></p>

FBSEVER	
<b>Format</b>	<RMNAME>
<b>Description</b>	The fallback server to use when talking to Moab in an HA configuration.
<b>Example</b>	<pre>RMCFG[base] TYPE=MOAB SERVER=server1 FBSEVER=server1-ha</pre>

FLAGS	
<b>Format</b>	Comma-delimited list of zero or selected resource manger flags. See <a href="#">11.2.2.D Resource Manager Flags - page 655</a> for valid values.
<b>Description</b>	Specifies various attributes of the resource manager.
<b>Example</b>	<pre>RMCFG[base] FLAGS=asyncstart</pre> <p><i>Moab directs the resource manager to start the job asynchronously.</i></p>

FNLIST	
<b>Format</b>	Comma-delimited list of zero or more of the following: <code>clusterquery</code> , <code>jobcancel</code> , <code>jobrequeue</code> , <code>jobresume</code> , <code>jobstart</code> , <code>jobsuspend</code> , <code>queuequery</code> , <code>resourcequery</code> or <code>workloadquery</code>
<b>Description</b>	By default, a resource manager utilizes all functions supported to query and control batch objects. If this parameter is specified, only the listed functions are used.
<b>Example</b>	<pre>RMCFG[base] FNLIST=queuequery</pre> <p><i>Moab only uses this resource manager interface to load queue configuration information.</i></p>

HOST	
<b>Format</b>	<STRING>
<b>Default</b>	localhost
<b>Description</b>	The host name of the machine on which the resource manager server is running.
<b>Example</b>	<pre>RMCFG[base] host=server1</pre>

JOBCANCELURL	
<b>Format</b>	[<protocol>://][<path>]
<b>Default</b>	---
<b>Description</b>	Specifies how Moab cancels jobs via the resource manager. See <a href="#">URL Notes</a> .
<b>Example</b>	<pre>RMCFG[base] JOBCANCELURL=exec:///opt/moab/job.cancel.lsf.pl</pre> <p><i>Moab executes /opt/moab/job.cancel.lsf.pl to cancel</i></p>

JOBCANCELURL	
	<i>specific jobs.</i>

JOBEXTENDDURATION	
<b>Format</b>	[[ [DD:]HH:]MM:]SS [, [[ [DD:]HH:]MM:]SS ] [!] [<] (or <MIN TIME> [, <MAX TIME>] [!])
<b>Default</b>	---
<b>Description</b>	<p>The minimum and maximum amount of time that can be added to a job's walltime if it is possible for the job to be extended. See <a href="#">MINWCLIMIT</a>. As the job runs longer than its current specified minimum wallclock limit (<code>-l minwclimit</code>, for example), Moab attempts to extend the job's limit by the minimum <code>JOBEXTENDDURATION</code>. This continues until either the extension can no longer occur (it is blocked by a reservation or job), the maximum <code>JOBEXTENDDURATION</code> is reached, or the user's specified wallclock limit (<code>-l walltime</code>) is reached. When a job is extended, it is marked as <code>PREEMPTIBLE</code>, unless the <code>!</code> is appended to the end of the configuration string. If the less than sign (<code>&lt;</code>) is appended to the end of the expression, the values are taken to be a per-iteration minimum and maximum extension with the overall maximum extension being the job's walltime limit.</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> <code>JOBEXTENDDURATION</code> and <code>JOBEXTENDSTARTWALLTIME TRUE</code> cannot be configured together. If they are in the same <code>moab.cfg</code> or are both active, then the <code>JOBEXTENDDURATION</code> will not be honored.</p> <p>For example, comment out the <code>JOBEXTENDSTARTWALLTIME</code>.</p> <pre style="border: 1px dashed black; padding: 5px; margin-top: 5px;"> RMCFG [base] JOBEXTENDDURATION=30,1:00:00 #JOBEXTENDSTARTWALLTIME TRUE                     </pre> </div>
<b>Example</b>	<pre style="border: 1px dashed black; padding: 5px; margin-bottom: 5px;"> RMCFG [base] JOBEXTENDDURATION=30,1:00:00                     </pre> <p><i>Moab extends a job's walltime by 30 seconds each time the job is about to run out of walltime until it is bound by one hour, a reservation/job, or the job's original 'maximum' wallclock limit.</i></p>



JOBIDFORMAT	
<b>Format</b>	INTEGER
<b>Default</b>	---
<b>Description</b>	Specifies that Moab should use numbers to create job IDs. This eliminates multiple job IDs associated with a single job.
<b>Example</b>	<pre>RMCFG[base] JOBIDFORMAT=INTEGER</pre> <p><i>Job IDs are generated as numbers.</i></p>

JOBMODIFYURL	
<b>Format</b>	[<protocol>://][<path>]
<b>Default</b>	---
<b>Description</b>	Specifies how Moab modifies jobs via the resource manager. See <a href="#">URL Notes</a> , and <a href="#">interface details</a> .
<b>Example</b>	<pre>RMCFG[base] JOBMODIFYURL=exec://\$TOOLSDIR/job.modify.dyn.pl</pre> <p><i>Moab executes /opt/moab/job.modify.dyn.pl to modify specific jobs.</i></p>

JOBRSVRECREATE	
<b>Format</b>	Boolean
<b>Default</b>	TRUE
<b>Description</b>	Specifies whether Moab will re-create a job reservation each time job information is updated by a resource manager. See <a href="#">Considerations for Large Clusters</a> for more information.

## JOBRSVRECREATE

### Example

```
RMCFG[base] JOBRSVRECREATE=FALSE
```

*Moab only creates a job reservation once when the job first starts.*

## JOBSTARTURL

### Format

```
[<protocol>://][<path>]
```

### Default

```
TRUE
```

### Description

Specifies how Moab starts jobs via the resource manager. See [URL Notes](#).

### Example

```
RMCFG[base] JOBSTARTURL=exec://$TOOLSDIR/job.submit.slurm.py
```

*Moab will execute the `job.submit.slurm.py` script when it wants to start a job.*

## JOBSUBMITURL

### Format

```
[<protocol>://][<path>]
```

### Description

Moab submits jobs to the script `job.submit.dyn.pl` located in the `TOOLSDIR`.

### Example

```
RMCFG[base] JOBSUBMITURL=exec://$TOOLSDIR/job.submit.dyn.pl
```

*Moab submits jobs directly to the database located on host `dbserver.flc.com`.*

## JOBSPENDURL

### Format

```
[<protocol>://][<path>]
```

JOBSUSPENDURL	
<b>Description</b>	Specifies how Moab suspends jobs via the resource manager. See <a href="#">URL Notes</a> .
<b>Example</b>	<pre>RMCFG[base] JOBSUSPENDURL=EXEC://\$HOME/scripts/job.suspend</pre> <p><i>Moab executes the <code>job.suspend</code> script when jobs are suspended.</i></p>


JOBVALIDATEURL	
<b>Format</b>	[<protocol>://][<path>]
<b>Description</b>	Specifies how Moab validates newly submitted jobs. See <a href="#">URL Notes</a> . If the script returns with a non-zero exit code, the job is rejected. See <a href="#">User Proxying/Alternate Credentials</a> .
<b>Example</b>	<pre>RMCFG[base] JOBVALIDATEURL=exec://\$TOOLS/job.validate.pl</pre> <p><i>Moab executes the <code>'job.validate.pl'</code> script when jobs are submitted to verify they are acceptable.</i></p>

MAXDSOP	
<b>Format</b>	<INTEGER>
<b>Default</b>	-1 (unlimited)
<b>Description</b>	The maximum number of data staging operations that can be simultaneously active.
<b>Example</b>	<pre>RMCFG[ds] MAXDSOP=16</pre>

MAXITERATIONFAILURECOUNT	
<b>Format</b>	<INTEGER>
<b>Default</b>	80
<b>Description</b>	The number of times the RM must fail within a certain iteration before Moab considers it down or corrupt. When an RM is down or corrupt, Moab will not attempt to interact with it.
<b>Example</b>	<pre>RMCFG[base] MAXITERATIONFAILURECOUNT=25</pre> <p><i>The RM base must fail 25 times in a single iteration for Moab to consider it down and cease interacting with it.</i></p>

MAXJOBPERMINUTE	
<b>Format</b>	<INTEGER>
<b>Default</b>	-1 (unlimited)
<b>Description</b>	The maximum number of jobs allowed to start per minute via the resource manager.
<b>Example</b>	<pre>RMCFG[base] MAXJOBPERMINUTE=5</pre> <p><i>The scheduler only allows five jobs per minute to launch via the resource manager base.</i></p>

MAXJOBS	
<b>Format</b>	<INTEGER>
<b>Default</b>	0 (limited only by the Moab MAXJOB setting)
<b>Description</b>	The maximum number of active jobs that this interface is allowed to load from the resource manager.

MAXJOBS	
	 Only works with Moab peer resource managers at this time.
<b>Example</b>	<pre>RMCFG[cluster1] SERVER=moab://cluster1 MAXJOBS=200</pre> <p><i>The scheduler loads up to 200 active jobs from the remote Moab peer cluster1.</i></p>

MINETIME	
<b>Format</b>	<INTEGER>
<b>Default</b>	1
<b>Description</b>	The minimum time in seconds between processing subsequent scheduling events.
<b>Example</b>	<pre>RMCFG[base] MINETIME=5</pre> <p><i>The scheduler batch-processes scheduling events that occur less than five seconds apart.</i></p>

NMPORT	
<b>Format</b>	<INTEGER>
<b>Default</b>	(any valid port number)
<b>Description</b>	Allows specification of the resource manager's node manager port and is only required when this port has been set to a non-default value.
<b>Example</b>	<pre>RMCFG[base] NMPORT=13001</pre> <p><i>The scheduler contacts the node manager located on each compute node at port 13001.</i></p>

NODEFAILURERSVPROFILE	
<b>Format</b>	<STRING>
<b>Description</b>	The rsv template to use when placing a reservation onto failed nodes. See also <a href="#">NODEFAILURERESERVETIME</a> .
<b>Example</b>	<pre># moab.cfg RMCFG[base] NODEFAILURERSVPROFILE=long RSVPROFILE[long] DURATION=25:00RSVPROFILE[long] USERLIST=john</pre> <p><i>The scheduler will use the long rsv profile when creating reservations over failed nodes belonging to base.</i></p>

NODESTATEPOLICY	
<b>Format</b>	One of OPTIMISTIC or PESSIMISTIC
<b>Default</b>	PESSIMISTIC
<b>Description</b>	Specifies how Moab should determine the state of a node when multiple resource managers are reporting state: OPTIMISTIC specifies that if any resource manager reports a state of up, that state will be used. PESSIMISTIC specifies that if any resource manager reports a state of down, that state will be used.
<b>Example</b>	<pre># moab.cfg RMCFG[native] TYPE=NATIVE NODESTATEPOLICY=OPTIMISTIC</pre>

OMAP	
<b>Format</b>	[<protocol>://][<path>]
<b>Description</b>	Specifies an object map file that is used to map credentials and other objects when using this resource manager peer. See <a href="#">Grid Credential</a>

OMAP	
	<a href="#">Management</a> for full details.
<b>Example</b>	<pre>moab.cfg RMCFG[peer1] OMAP=file:///opt/moab/omap.dat</pre> <p><i>When communicating with the resource manager <code>peer1</code>, objects are mapped according to the rules defined in the <code>/opt/moab/omap.dat</code> file.</i></p>

PORT	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	The port on which the scheduler should contact the associated resource manager. The value 0 specifies that the resource manager default port should be used.
<b>Example</b>	<pre>RMCFG[base] TYPE=PBS HOST=cws PORT=20001</pre> <p><i>Moab attempts to contact the PBS server daemon on host <code>cws</code>, port <code>20001</code>.</i></p>

PROVDURATION	
<b>Format</b>	[ [ [DD:] HH:] MM:] SS
<b>Default</b>	2:30
<b>Description</b>	The upper bound (walltime) of a provisioning request. After this duration, Moab will consider the provisioning attempt failed.
<b>Example</b>	<pre>RMCFG[base] PROVDURATION=5:00</pre> <p><i>When RM <code>base</code> provisions a node for more than 5 minutes,</i></p>

## PROVDURATION

*Moab considers the provisioning as having failed.*

## RESOURCECREATEURL

## Format

[<protocol>://][<path>]

## Description

Specifies a script or method that can be used by Moab to create resources dynamically, such as creating a virtual machine on a hypervisor.

## Example

```
RMCFG[base] RESOURCECREATEURL=exec:///opt/script/vm.provision.py
```

*Moab invokes the `vm.provision.py` script, passing in data as command line arguments, to request a creation of new resources.*

## RESOURCETYPE

## Format

{COMPUTE|FS|LICENSE|NETWORK|PROV}

## Description

Specifies which type of resource this resource manager is configured to control. See [Native Resource Managers](#) for more information.



If LICENSE is specified, *all* generic resources reported by the resource manager will be marked as a license; causing them to be tracked by the accounting manager under the Licenses property. See [Licenses - page 427](#) for more information on the Licenses property.

## Example

```
RMCFG[base] TYPE=NATIVE RESOURCETYPE=FS
```

*Resource manager `base` will function as a NATIVE resource manager and control file systems.*




RMSTARTURL	
<b>Format</b>	[<protocol>://][<path>]
<b>Description</b>	Specifies how Moab starts the resource manager.
<b>Example</b>	<pre>RMCFG[base] RMSTARTURL=exec:///tmp/nat.start.pl</pre> <p><i>Moab executes /tmp/nat.start.pl to start the resource manager base.</i></p>

RMSTOPURL	
<b>Format</b>	[<protocol>://][<path>]
<b>Description</b>	Specifies how Moab stops the resource manager.
<b>Example</b>	<pre>RMCFG[base] RMSTOPURL=exec:///tmp/nat.stop.pl</pre> <p><i>Moab executes /tmp/nat.stop.pl to stop the resource manager base.</i></p>

SBINDIR	
<b>Format</b>	<PATH>
<b>Description</b>	For use with Torque; specifies the location of the Torque system binaries (supported in Torque 1.2.0p4 and higher).
<b>Example</b>	<pre>RMCFG[base] TYPE=pbs SBINDIR=/usr/local/torque/sbin</pre> <p><i>Moab tells Torque that its system binaries are located in /usr/local/torque/sbin.</i></p>

SERVER	
<b>Format</b>	<URL>
<b>Description</b>	The resource management service to use. If not specified, the scheduler locates the resource manager via built-in defaults or, if available, with an information service.
<b>Example</b>	<pre>RMCFG[base] server=ll://supercluster.org:9705</pre> <p><i>Moab attempts to use the Loadleveler scheduling API at the specified location.</i></p>

STAGETHRESHOLD	
<b>Format</b>	[ [ [DD:]HH:]MM:]SS
<b>Description</b>	<p>The maximum time a job waits to start locally before considering being migrated to a remote peer. In other words, if a job's start time on a remote cluster is less than the start time on the local cluster, but the difference between the two is less than STAGETHRESHOLD, then the job is scheduled locally. The aim is to avoid job/data staging overhead if the difference in start times is minimal.</p> <p> If this attribute is used, backfill is disabled for the associated resource manager.</p>
<b>Example</b>	<pre>RMCFG[remote_cluster] STAGETHRESHOLD=00:05:00</pre> <p><i>Moab only migrates jobs to remote_cluster if the jobs can start five minutes sooner on the remote cluster than they could on the local cluster.</i></p>

STARTCMD	
<b>Format</b>	<STRING>
<b>Description</b>	The full path to the resource manager job start client. If the resource

STARTCMD	
	<p>manager API fails, Moab executes the specified start command in a second attempt to start the job.</p> <div style="border: 1px solid black; padding: 5px;"> <p><b>i</b> Moab calls the start command with the format <code>&lt;CMD&gt;&lt;JOBID&gt; -H &lt;HOSTLIST&gt;</code> unless the environment variable <code>MOABNOHOSTLIST</code> is set, in which case Moab will only pass the job ID.</p> </div>
<b>Example</b>	<pre>RMCFG[base] STARTCMD=/usr/local/bin/qrun</pre> <div style="border: 1px dashed black; padding: 5px; margin-top: 10px;"> <p><i>Moab uses the specified start command if API failures occur when launching jobs.</i></p> </div>

SUBMITCMD	
<b>Format</b>	<STRING>
<b>Description</b>	The full path to the resource manager job submission client.
<b>Example</b>	<pre>RMCFG[base] SUBMITCMD=/usr/local/bin/qsub</pre> <div style="border: 1px dashed black; padding: 5px; margin-top: 10px;"> <p><i>Moab uses the specified submit command when migrating jobs.</i></p> </div>

SUBMITPOLICY	
<b>Format</b>	One of <code>NODECENTRIC</code> or <code>PROCCENTRIC</code>
<b>Default</b>	<code>PROCCENTRIC</code>
<b>Description</b>	If set to <code>NODECENTRIC</code> , each specified node requested by the job is interpreted as a true compute host, not as a task or processor.
<b>Example</b>	<pre>RMCFG[base] SUBMITPOLICY=NODECENTRIC</pre> <div style="border: 1px dashed black; padding: 5px; margin-top: 10px;"> <p><i>Moab uses the specified submit policy when migrating jobs.</i></p> </div>

SUSPENDSIG	
<b>Format</b>	<INTEGER> (valid UNIX signal between 1 and 64)
<b>Default</b>	RM-specific default
<b>Description</b>	If set, Moab sends the specified signal to a job when a job suspend request is issued.
<b>Example</b>	<pre>RMCFG[base] SUSPENDSIG=19</pre> <p><i>Moab uses the specified suspend signal when suspending jobs within the base resource manager.</i></p> <p><b>i</b> SUSPENDSIG should not be used with Torque or other PBS-based resource managers.</p>

SYNCJOBID	
<b>Format</b>	<BOOLEAN>
<b>Description</b>	Specifies that Moab should migrate jobs to the local resource manager with the job's Moab-assigned job ID. In a grid, the grid-head will only pass dependencies to the underlying Moab if SYNCJOBID is set. This attribute can be used with the <a href="#">JOBIDFORMAT</a> attribute and <a href="#">PROXYJOB SUBMISSION</a> flag in order to assign job IDs from Moab to the resource manager. For more information about all the steps necessary for using one job ID between Moab and Torque, see <a href="#">Synchronizing Job IDs in Torque and Moab</a> .
<b>Example</b>	<pre>RMCFG[torque] TYPE=PBS SYNCJOBID=TRUE</pre>

SYSTEMMODIFYURL	
<b>Format</b>	[<protocol>://][<path>]

SYSTEMMODIFYURL	
<b>Description</b>	Specifies how Moab modifies attributes of the system. This interface is used in <a href="#">data staging</a> .
<b>Example</b>	<pre>RMCFG[base] SYSTEMMODIFYURL=exec:///tmp/system.modify.pl</pre> <p><i>Moab executes /tmp/system.modify.pl when it modifies system attributes in conjunction with the resource manager base.</i></p>

SYSTEMQUERYURL	
<b>Format</b>	[<protocol>://][<path>]
<b>Description</b>	Specifies how Moab queries attributes of the system. This interface is used in <a href="#">data staging</a> .
<b>Example</b>	<pre>RMCFG[base] SYSTEMQUERYURL=file:///tmp/system.query</pre> <p><i>Moab reads /tmp/system.query when it queries the system in conjunction with base resource manager.</i></p>

TARGETUSAGE	
<b>Format</b>	<INTEGER> [%]
<b>Default</b>	90%
<b>Description</b>	Amount of resource manager resources to explicitly use. In the case of a storage resource manager, indicates the target usage of data storage resources to dedicate to active data migration requests. If the specified value contains a percent sign (%), the target value is a percent of the configured value. Otherwise, the target value is considered to be an absolute value measured in megabytes (MB).
<b>Example</b>	<pre>RMCFG[storage] TYPE=NATIVE RESOURCETYPE=storage RMCFG[storage] TARGETUSAGE=80%</pre>

## TARGETUSAGE

*Moab schedules data migration requests to never exceed 80% usage of the storage resource manager's disk cache and network resources.*

## TIMEOUT

**Format** <INTEGER>

**Default** 30

**Description** Time (in seconds) the scheduler waits for a response from the resource manager.

**Example**

```
RMCFG[base] TIMEOUT=40
```

*Moab waits 40 seconds to receive a response from the resource manager before timing out and giving up. Moab tries again on the next iteration.*

## TRIGGER

**Format** <TRIG\_SPEC>

**Description** A [trigger](#) specification indicating behaviors to enforce in the event of certain events associated with the resource manager, including resource manager start, stop, and failure.

**Example**

```
RMCFG[base] TRIGGER=<X>
```

## TYPE

**Format** <RMTYPE> [ :<RMSUBTYPE> ] where <RMTYPE> is one of the following:

TYPE	
	Torque, <a href="#">NATIVE</a> , PBS, RMS, SSS, or <a href="#">WIKI</a> and the optional <RMSUBTYPE> value is one of RMS.
<b>Default</b>	PBS
<b>Description</b>	<p>Specifies type of resource manager to be contacted by the scheduler.</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> For TYPE WIKI, AUTHTYPE must be set to CHECKSUM. The &lt;RMSUBTYPE&gt; option is currently only used to support Compaq's RMS resource manager in conjunction with PBS. In this case, the value PBS :RMS should be specified.</p> </div>
<b>Example</b>	<pre>RMCFG[clusterA] TYPE=PBS HOST=clusterA PORT=15003 RMCFG[clusterB] TYPE=PBS HOST=clusterB PORT=15005</pre> <p><i>Moab interfaces to two different PBS resource managers, one located on server clusterA at port 15003 and one located on server clusterB at port 15005.</i></p>

VARIABLES	
<b>Format</b>	<VAR>=<VAL> [, <VAR>=<VAL> ]
<b>Description</b>	Opaque resource manager variables.
<b>Example</b>	<pre>RMCFG[base] VARIABLES=SCHEDDHOST=head1</pre> <p><i>Moab associates the variable SCHEDDHOST with the value head1 on resource manager base.</i></p>

VERSION	
<b>Format</b>	<STRING>
<b>Description</b>	Resource manager-specific version string.

VERSION	
<b>Example</b>	<pre>RMCFG[base] VERSION=10124</pre> <p><i>Moab assumes that resource manager base has a version number of 1.1.24.</i></p>

WORKLOADQUERYURL	
<b>Format</b>	[<protocol>://][<path>]
<b>Description</b>	Specifies how Moab queries the resource manager for workload information. See <a href="#">Native RM</a> , <a href="#">URL Notes</a> , and <a href="#">interface details</a> .
<b>Example</b>	<pre>RMCFG[Torque] WORKLOADQUERYURL=exec://\$TOOLS_DIR/job.query.dyn.pl</pre> <p><i>Moab executes /opt/moab/tools/job.query.dyn.pl to obtain updated workload information from resource manager Torque.</i></p>

## URL Notes

URL parameters can load data by using the `exec` or `file` protocols.

For the protocol `file`, Moab loads the data directly from the text file pointed to by path.

```
RMCFG[base] SYSTEMQUERYURL=file:///tmp/system.query
```

For the protocol `exec`, Moab executes the file pointed to by path and loads the output written to `STDOUT`. If the script requires arguments, you can use a question mark (?) between the script name and the arguments, and an ampersand (&) for each space.

```
RMCFG[base] JOBVALIDATEURL=exec://$TOOLS/job.validate.pl
RMCFG[native] CLUSTERQUERYURL=exec://opt/moab/tools/cluster.query.pl?-group=group1&-arch=x86
```

## Synchronizing Job IDs in Torque and Moab

**i** Unless you use an `msub` submit filter or you're in a grid, it is recommended that you use your RM-specific job submission command (for instance, `qsub`).

In order to synchronize your job IDs between Torque and Moab you must perform the following steps:



1. Verify that you are using Torque version 2.5.6 or later.
2. Set `SYNCJOBID` to `TRUE` in all resource managers.

```
RMCFG[torque] TYPE=PBS SYNCJOBID=TRUE
```

3. Set the `PROXYJOBSUBMISSION` flag. With `PROXYJOBSUBMISSION` enabled, you must run Moab as a Torque manager or operator. Verify that other users can submit jobs using `msub`. Moab, as a non-root user, should still be able to submit jobs to Torque and synchronize job IDs.

```
RMCFG[torque] TYPE=PBS SYNCJOBID=TRUE
RMCFG[torque] FLAGS=PROXYJOBSUBMISSION
```

4. Add `JOBIDFORMAT=INTEGER` to the internal RM. Adding this parameter forces Moab to only use numbers as job IDs and those numbers to use across Moab, Torque, and the entire grid. This enhances the end-user experience as it eliminates multiple job IDs associated with a single job.

```
RMCFG[torque] TYPE=PBS SYNCJOBID=TRUE
RMCFG[torque] FLAGS=PROXYJOBSUBMISSION

RMCFG[internal] JOBIDFORMAT=INTEGER
```

## 11.2.2 Resource Manager Configuration Details

As with all scheduler parameters, the `RMCFG` parameter follows the syntax described in [2.6 Configuring the Scheduler - page 54](#).

### 11.2.2.A Resource Manager Types

The `RMCFG` parameter allows the scheduler to interface to multiple types of resource managers using the `TYPE` or `SERVER` attributes. Specifying these attributes, any of the following listed resource managers can be supported.

Type	Resource Managers	Details
<b>Moab</b>	Moab Workload Manager	Use the Moab peer-to-peer (grid) capabilities to enable grids and other configurations. See <a href="#">Grid Configuration</a> .
<b>MWS</b>	Moab Web Services	The MWS resource manager type is a native integration between Moab and MWS. Resource manager data is passed directly between Moab and MWS using JSON (rather than Moab's native WIKI syntax). This simplifies RM configuration for systems where one or more MWS plugins are acting as resource managers. See the 'Moab

Type	Resource Managers	Details
		Workload Manager Resource Manager Integration' section in the <i>Moab Web Services Reference Guide</i> for more information.
<b>Native</b>	Moab <i>Native</i> Interface	Used for connecting directly to scripts, files, and databases. See <a href="#">Managing Resources Directly with the Native Interface</a> .
<b>PBS</b>	Torque (all versions)	N/A
<b>SSS</b>	Scalable Systems Software Project version 2.0 and higher	N/A
<b>WIKI</b>	<a href="#">Wiki</a> interface specification version 1.0 and higher	Used for most non-PBS resource managers.

### 11.2.2.B Resource Manager Name

Moab can support more than one resource manager simultaneously. Consequently, the `RMCFG` parameter takes an index value such as `RMCFG[clusterA]`. This index value essentially names the resource manager (as done by the deprecated parameter `RMNAME`). The resource manager name is used by the scheduler in diagnostic displays, logging, and in reporting resource consumption to the accounting manager. For most environments, the selection of the resource manager name can be arbitrary.

### 11.2.2.C Resource Manager Location

The `HOST`, `PORT`, and `SERVER` attributes can be used to specify how the resource manager should be contacted. For many resource managers the interface correctly establishes contact using default values. These parameters need only to be specified for resource managers such as the `WIKI` interface (that do not include defaults) or with resources managers that can be configured to run at non-standard locations (such as `PBS`). In all other cases, the resource manager is automatically located.

### 11.2.2.D Resource Manager Flags

The `FLAGS` attribute can be used to modify many aspects of a resources manager's behavior.

**i** `AUTOSYNC`, `COLLAPSEVIEW`, `HOSTINGCENTER`, `PRIVATE`, `REPORT`, `SHARED`, and `STATIC` are deprecated.

Flag	Description
<b>ASYNCDELETE</b>	<p>Moab directs the resource manager to not wait for confirmation that the job correctly cancels before the API call returns. See <a href="#">Large Cluster Tuning</a> for more information.</p> <p><b>i</b> This flag is only applicable for Torque or Moab Native resource managers.</p>
<b>ASYNCSTART</b>	<p>Jobs started on this resource manager start asynchronously. In this case, the scheduler does not wait for confirmation that the job correctly starts before proceeding. See <a href="#">Large Cluster Tuning</a> for more information.</p> <p><b>i</b> This flag is only applicable for Torque or Moab Native resource managers.</p>
<b>AUTOSTART</b>	<p>Jobs staged to this resource manager do not need to be explicitly started by the scheduler. The resource manager itself handles job launch.</p>
<b>BECOMEMASTER</b>	<p>Nodes reported by this resource manager will transfer ownership to this resource manager if they are currently owned by another resource manager that does not have this flag set.</p>
<b>CLIENT</b>	<p>A client resource manager object is created for diagnostic/statistical purposes or to configure Moab's interaction with this resource manager. It represents an external entity that consumes server resources or services, allows a local administrator to track this usage, and configures specific policies related to that resource manager. A client resource manager object loads no data and provides no services.</p>
<b>CLOCKSKEWCHECKING</b>	<p>Setting <code>CLOCKSKEWCHECKING</code> allows you to configure</p>

Flag	Description
	clock skew adjustments. Most of the time it is sufficient to use an NTP server to keep the clocks in your system synchronized.
<b>DYNAMICCRED</b>	The resource manager creates credentials within the cluster as needed to support workload.
<b>EnableCondensedQuery</b>	Enables the condensed workload query. <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> Only applies if the Torque parameter <code>job_full_report_time</code> is used (Torque Resource Manager version 5.1.x or later). See Server Parameters in the <i>Torque Resource Manager Administrator Guide</i>.</p> </div>
<b>EXECUTIONSERVER</b>	The resource manager is capable of launching and executing batch workload.
<b>FSISREMOTE</b>	Add this flag if the working file system doesn't exist on the server to prevent Moab from validating files and directories at migration.
<b>FULLCP</b>	Always checkpoint full job information (useful with <a href="#">Native</a> resource managers).
<b>IgnOS</b>	Ignore the operating system reported by the resource manager on each node and use the OS in Moab's configuration files. See <a href="#">OS</a> for more information.
<b>IGNQUEUESTATE</b>	The queue state reported by the resource manager should be ignored. Can be used if queues must be disabled inside of a particular resource manager to allow an external scheduler to properly operate.
<b>IGNWORKLOADSTATE</b>	When this flag is applied to a native resource manager, any jobs that are reported via that resource manager's 'workload query URL' have their reported state ignored. For example, if an RM has the <code>IgnWorkloadState</code> flag and it reports that a set of jobs have a state of 'Running,' this state is ignored and the jobs will either have a default state set or will inherit the state from another RM reporting on that same set of jobs. This flag only changes the behavior of RMs of type

Flag	Description
	NATIVE.
<b>LOCALWORKLOADEXPORT</b>	When set, destination peers share information about local and remote jobs, allowing job management of different clusters at a single peer. For more information, see <a href="#">Workload Submission and Control</a> .
<b>MIGRATEALLJOBATTRIBUTES</b>	When set, this flag causes additional job information to be migrated to the resource manager; additional job information includes things such as node features applied via <code>CLASSCFG[name] DEFAULT.FEATURES</code> , the account to which the job was submitted, job walltime limit, and node exclusivity.
<b>NOAUTORES</b>	If the resource manager does not report CPU usage to Moab because CPU usage is at 0%, Moab assumes full CPU usage. When set, Moab recognizes the resource manager report as 0% usage. This is only valid for PBS.
<b>NoCondensedQuery</b>	Disables the condensed workload query. This is the default for Moab 9.0 and later. <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> Only applies if the Torque parameter <code>job_full_report_time</code> is used (Torque Resource Manager version 5.1.x or later). See Server Parameters in the <i>Torque Resource Manager Administrator Guide</i>.</p> </div>
<b>NOCREATERESOURCE</b>	To use resources discovered from this resource manager, they must be created by another resource manager first. For example, if you set <code>NOCREATERESOURCE</code> on RM A, which reports nodes 1 and 2, and RM B only reports node 1, then node 2 will not be created because RM B did not report it.
<b>PROXYJOBSUBMISSION</b>	Enables Admin proxy job submission, which means administrators can submit jobs in behalf of other users.
<b>PUSHSLAVEJOBUPDATES</b>	Enables job changes made on a Moab Grid Member to be pushed to the Moab Grid Control. Without this flag, jobs being reported to the Moab Grid Control do not show any changes made on the remote Moab server (via <code>mjobctl</code> and so forth).

Flag	Description
<b>RECORDGPUMETRICS</b>	Enables the recording of GPU metrics for nodes.
<b>RECORDMICMETRICS</b>	Enables the recording of MIC metrics for nodes.
<b>SLAVEPEER</b>	Information from this resource manager cannot be used to identify new jobs or nodes. Instead, this information can only be used to update jobs and nodes discovered and loaded from other non-slave resource managers.
<b>THREADEDQUERIES</b>	When this flag is set for an individual RM, the queries that Moab performs to get information from the RM is done in a separate thread from the main Moab process. This allows Moab to remain responsive during the query and ultimately reduces the time spent in a scheduling cycle. If multiple RMs are being used the effect can be more significant because all RMs will be queried in parallel.
<b>USEPHYSICALMEMORY</b>	Tells Moab to use a node's physical memory instead of the swap space.  For example, if a node has 12 GB of RAM and an additional 12 GB of swap space, it has 24 GB of virtual memory. If a 4 GB job is assigned to that node, the reported available memory shows 12 GB because the job is using the swap space not the physical memory. The reported available memory doesn't decrease until the swap space is used up. When this flag is set, the 4 GB job immediately reduces the available memory to 8 GB (physical memory - used memory).
<b>USERSPACEISSEPARATE</b>	Tells Moab to ignore validating the user's UID and GID in the case that information doesn't exist on the Moab server.

### Example

```
# resource manager 'torque' should use asynchronous job start
RMCFG[torque] FLAGS=asyncstart
```

## 11.2.3 Scheduler/Resource Manager Interactions

In the simplest configuration, Moab interacts with the resource manager using the following four primary functions:

Function	Description
<b>GETJOBINFO</b>	Collect detailed state and requirement information about idle, running, and recently completed jobs.
<b>GETNODEINFO</b>	Collect detailed state information about idle, busy, and defined nodes.
<b>STARTJOB</b>	Immediately start a specific job on a particular set of nodes.
<b>CANCELJOB</b>	Immediately cancel a specific job regardless of job state.

Using these four simple commands, Moab enables nearly its entire suite of scheduling functions. More detailed information about resource manager specific requirements and semantics for each of these commands can be found in the specific resource manager (such as [WIKI](#)) overviews.

In addition to these base commands, other commands are required to support advanced features such as suspend/resume, gang scheduling, and scheduler initiated checkpoint restart.

Information on creating a new scheduler resource manager interface can be found in the [Adding New Resource Manager Interfaces](#) section.

## 11.3 Resource Manager Extensions

In this topic:

[11.3.1 Resource Manager Extension Specification - page 660](#)

[11.3.2 Resource Manager Extension Values - page 660](#)

[11.3.3 Resource Manager Extension Examples - page 687](#)

[11.3.4 Configuring dynamic features in Torque and Moab - page 688](#)

All resource managers are not created equal. There is a wide range in what capabilities are available from system to system. Additionally, there is a large body of functionality that many, if not all, resource managers have no concept of. A good example of this is job QoS. Since most resource managers do not have a concept of quality of service, they do not provide a mechanism for users to specify this information. In many cases, Moab is able to add capabilities at a global level. However, a number of features require a *per job* specification. Resource manager extensions allow this information to be associated with the job.

### 11.3.1 Resource Manager Extension Specification

Specifying resource manager extensions varies by resource manager. Torque and Wiki each allow the specification of an *extension* field as described in the following table:

Resource Manager	Specification Method
<b>Torque 2.0+</b>	-l <pre>&gt; qsub -l nodes=3,qos=high sleepy.cmd</pre>
<b>Torque 1.x</b>	-W x= <pre>&gt; qsub -l nodes=3 -W x=qos:high sleepy.cmd</pre>
<b>Wiki</b>	comment <pre>comment=qos:high</pre>

### 11.3.2 Resource Manager Extension Values

All of the following job extensions will work with `msub -l` (or `msub -W x=. . .`). However, `qsub -l` only provides legacy support for a subset of these extensions; see 'Requesting Resources' in the *Torque Resource Manager Administrator Guide* for the list.

If your configuration primarily uses `qsub` to submit jobs, Adaptive Computing recommends you use the `qsub -W x=` syntax for all submissions with Moab job extensions to avoid `qsub` rejection for any unsupported (non-legacy) extensions.

The following job extensions are supported when using the resource manager-specific method:



*Resource Manager Extension Values*

ADVRES	MEM	PROLOGUE
CPUCLOCK	MICs	PVMEM
DDISK	MINPREEMPTTIME	QoS
DEADLINE	MINPROCSPEED	QUEUEJOB
DEPEND	MINWCLIMIT	REQATTR
DMEM	MSTAGEIN	RESFAILPOLICY
EPILOGUE	MSTAGEOUT	RMTYPE
EXCLUDENODES	NACCESSPOLICY	SIGNAL
FEATURE	NALLOCPOLICY	GRES and SOFTWARE
GATTR	NCPUS	SPRIORITY
GMETRIC	NMATCHPOLICY	TEMPLATE
GPU <sub>s</sub>	NODESET	TERMTIME
GRES and SOFTWARE	NODESETCOUNT	TPN
HOSTLIST	NODESETDELAY	TRIG
JGROUP	NODESETISOPTIONAL	TRL (Format 1)
JOBFLAGS (a.k.a. FLAGS)	OPSYS	TRL (Format 2)
JOBREJECTPOLICY	PARTITION	VAR
MAXMEM	PMEM	VC
MAXPROC	PREF	VMEM
	PROCS	

ADVRES	
<b>Format</b>	[!]<RSVID>
<b>Description</b>	<p>Specifies that reserved resources are required to run the job. If &lt;RSVID&gt; is specified, then only resources within the specified reservation can be allocated (see <a href="#">Job to Reservation Binding</a>).</p> <p>You can request to not use a specific reservation by using <code>advres=!&lt;reservationname&gt;</code>.</p>
<b>Example</b>	<pre>&gt; qsub -l advres=grid.3</pre> <p><i>Resources for the job must come from grid.3.</i></p> <pre>&gt; qsub -l advres=!grid.5</pre> <p><i>Resources for the job must not come from grid.5</i></p>

CPUCLOCK	
<b>Format</b>	<STRING>

## CPUCLOCK

**Description**

Specify the CPU clock frequency for each node requested for this job. A `cpuclock` request applies to every processor on every node in the request. Specifying varying CPU frequencies for different nodes or different processors on nodes in a single job request is not supported.

Not all CPUs support all possible frequencies or ACPI states. If the requested frequency is not supported by the CPU, the nearest frequency is used.

 If a job does not place any load on the node then some OSs will drop the frequency below the requested frequency.

Using `cpuclock` sets `NODEACCESSPOLICY` to `SINGLEJOB`.

The clock frequency can be specified via:

- a number that indicates the clock frequency (with or without the SI unit suffix).
- a Linux power governor policy name. The governor names are:
  - `performance`: This governor instructs Linux to operate each logical processor at its maximum clock frequency. This setting consumes the most power and workload executes at the fastest possible speed.
  - `powersave`: This governor instructs Linux to operate each logical processor at its minimum clock frequency. This setting executes workload at the slowest possible speed. This setting does not necessarily consume the least amount of power since applications execute slower, and may actually consume more energy because of the additional time needed to complete the workload's execution.
  - `ondemand`: This governor dynamically switches the logical processor's clock frequency to the maximum value when system load is high and to the minimum value when the system load is low. This setting causes workload to execute at the fastest possible speed or the slowest possible speed, depending on OS load. The system switches between consuming the most power and the least power.

## CPUCLOCK

**i** The power saving benefits of `ondemand` might be non-existent due to frequency switching latency if the system load causes clock frequency changes too often.

This has been true for older processors since changing the clock frequency required putting the processor into the C3 'sleep' state, changing its clock frequency, and then waking it up, all of which required a significant amount of time.

Newer processors, such as the Intel Xeon E5-2600 Sandy Bridge processors, can change clock frequency dynamically and much faster.

- `conservative`: This governor operates similar to the `ondemand` governor but is more conservative in switching between frequencies. It switches more gradually and uses all possible clock frequencies.
- This governor can switch to an intermediate clock frequency if it seems appropriate to the system load and usage, which the `ondemand` governor does not do.
- an ACPI performance state (or P-state) with or without the P prefix. P-states are a special range of values (0-15) that map to specific frequencies. Not all processors support all 16 states, however, they all start at P0. P0 sets the CPU clock frequency to the highest performance state, which runs at the maximum frequency. P15 sets the CPU clock frequency to the lowest performance state, which runs at the lowest frequency.

When reviewing job or node properties when `cpuclock` was used, be mindful of unit conversion. The OS reports frequency in Hz, not MHz or GHz.

**i** If a job does not place any load on the node then some OSs will drop the frequency below the requested frequency.

## Example

```
msub -l cpuclock=1800,nodes=2 script.sh
msub -l cpuclock=1800mhz,nodes=2 script.sh
```

*This job requests 2 nodes and specifies their CPU frequencies should be set to 1800 MHz.*

```
msub -l cpuclock=performance,nodes=2 script.sh
```

*This job requests 2 nodes and specifies their CPU frequencies should be set to the performance power governor policy.*

```
msub -l cpuclock=3,nodes=2 script.sh
```

CPUCLOCK	
	<pre>mbsub -l cpuclock=<b>p3</b>,nodes=2 script.sh</pre> <p><i>This job requests 2 nodes and specifies their CPU frequencies should be set to a performance state of 3.</i></p>

DDISK	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	Dedicated disk per task in MB.
<b>Example</b>	<pre>&gt; qsub -l ddisk=2000</pre>

DEADLINE	
<b>Format</b>	Relative time: [ [ [DD:]HH:]MM:]SS Absolute time: hh:mm:ss_mm/dd/yy
<b>Description</b>	Either the relative completion deadline of job (from job submission time) or an absolute deadline where you specify the date and time the job will finish.
<b>Example</b>	<pre>&gt; qsub -l deadline=2:00:00,nodes=4 /tmp/bio3.cmd</pre> <p><i>The job's deadline is 2 hours after its submission.</i></p>

DEPEND	
<b>Format</b>	[<DEPENDTYPE>:] [{jobname jobid}.]<ID>[: [{jobname jobid}.]<ID>]...
<b>Description</b>	Allows specification of job dependencies for compute or system jobs. If no ID prefix (jobname or jobid) is specified, the ID value is interpreted as a job ID. See <a href="#">9.5 Job Dependencies - page 562</a> for more information.

DEPEND	
<b>Example</b>	<pre># submit job which will run after job 1301 and 1304 complete &gt; msub -l depend=orion.1301:orion.1304 test.cmd orion.1322 # submit jobname-based dependency job &gt; msub -l depend=jobname.data1005 dataetl.cmd orion.1428</pre>
DMEM	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	Dedicated memory per task in bytes.
<b>Example</b>	<pre>&gt; msub -l dmem=20480</pre> <p><i>Moab will dedicate 20 MB of memory to the task.</i></p>
EPILOGUE	
<b>Format</b>	<STRING>
<b>Description</b>	<p>Specifies a user owned epilogue script, which is run before the system epilogue and <code>epilogue.user</code> scripts at the completion of a job. The syntax is <code>epilogue=&lt;file&gt;</code>. The file can be designated with an absolute or relative path.</p> <p><b>i</b> This parameter works only with Torque.</p>
<b>Example</b>	<pre>&gt; msub -l epilogue=epilogue_script.sh job.sh</pre>
EXCLUDENODES	
<b>Format</b>	{<nodeid> <node_range>}[:...]
<b>Description</b>	Specifies nodes that should not be considered for the given job.

EXCLUDENODES	
	<p><b>i</b> Moab does not support the combination of <code>msub -l excludenodes</code> and <code>ENABLEHIGHTHROUGHPUT TRUE</code>.</p>
<b>Example</b>	<pre>&gt; msub -l excludenodes=k1:k2:k[5-8]</pre>
FEATURE	
<b>Format</b>	<FEATURE> [ { :   } <FEATURE> ] . . .
<b>Description</b>	<p>Required list of node attribute/node <a href="#">features</a>.</p> <p><b>i</b> If the <i>pipe</i> ( ) character is used as a delimiter, the features are logically ORed together and the associated job can use resources that match any of the specified features.</p> <p><b>i</b> Requesting node names as features will result in the job being blocked from running.</p>
<b>Example</b>	<pre>&gt; qsub -l feature='fastos:bigio' testjob.cmd</pre> <p><i>Submits testjob.cmd with fastos:bigio as a required feature.</i></p> <pre>&gt; qsub -l feature=!bigmem testjob.cmd &gt; qsub -l feature='!bigmem' testjob.cmd</pre> <p><i>Submits testjob.cmd with a requirement that bigmem is <b>not</b> a node feature. (The exclamation point must either be escaped (!bigmem) or quoted ('!bigmem')).</i></p>
GATTR	
<b>Format</b>	<STRING>
<b>Description</b>	Generic job attribute associated with job. The maximum size for an attribute is 63 bytes (the core Moab size limit of 64, including a null byte).
<b>Example</b>	<pre>&gt; qsub -l gattr=bigjob</pre>

GMETRIC	
<b>Format</b>	Generic metric requirement for allocated nodes where the requirement is specified using the format <code>&lt;GMNAME&gt; [: {lt:, le:, eq:, ge:, gt:, ne: } &lt;VALUE&gt;]</code>
<b>Description</b>	Indicates generic constraints that must be found on all allocated nodes. If a <code>&lt;VALUE&gt;</code> is not specified, the node must simply possess the generic metric (See <a href="#">Generic Metrics</a> for more information).
<b>Example</b>	<pre>&gt; qsub -l gmetric=bioversion:ge:133244 testj.txt</pre>

GPUs	
<b>Format</b>	<pre>msub -l nodes=&lt;VALUE&gt;:ppn=&lt;VALUE&gt;:gpus=&lt;VALUE&gt; [:mode] [:reseterrr]</pre> <p>Where <code>mode</code> is one of:</p> <ul style="list-style-type: none"> <li><i>exclusive</i> - The default setting. The GPU is used exclusively by one process thread.</li> <li><i>exclusive_thread</i> - The GPU is used exclusively by one process thread.</li> <li><i>exclusive_process</i> - The GPU is used exclusively by one process regardless of process thread.</li> </ul> <p>If present, <code>reseterrr</code> resets the ECC memory bit error counters. This only resets the volatile error counts, or errors since the last reboot. The permanent error counts are not affected.</p> <p>Moab passes the <code>mode</code> and <code>reseterrr</code> portion of the request to Torque for processing.</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> Moab does not support requesting GPUs as a GRES. Submitting <code>msub -l gres=gpus:x</code> does not work.</p> </div>
<b>Description</b>	Moab schedules GPUs as a special type of <a href="#">node-locked generic resources</a> . When <a href="#">Torque reports GPUs</a> to Moab, Moab can schedule jobs and correctly assign GPUs to ensure that jobs are scheduled efficiently. To have Moab schedule GPUs, configure them in Torque then submit jobs using the 'GPU' attribute. Moab automatically parses the 'GPU' attribute and assigns them in the correct manner. For information about GPU metrics, see <a href="#">GPGPUMetrics</a> .
<b>Examples</b>	<pre>&gt; msub -l nodes=2:ppn=2:gpus=1:exclusive_process:reseterrr</pre>

GPUs	
	<p><i>Submits a job that requests 2 tasks, 2 processors and 1 GPU per task (2 GPUs total). Each GPU runs only threads related to the task and resets the volatile ECC memory bit error counts at job start time.</i></p> <pre>&gt; msub -l nodes=4:gpus=1,tpn=2</pre>
	<p><i>Submits a job that requests 4 tasks, 1 GPU per node (4 GPUs total), and 2 tasks per node. Each GPU is dedicated exclusively to one task process and the ECC memory bit error counters are not reset.</i></p> <pre>&gt; msub -l nodes=4:gpus=1:reseterr</pre>
	<p><i>Submits a job that requests 4 tasks, 1 processor and 1 GPU per task (4 GPUs total). Each GPU is dedicated exclusively to one task process and resets the volatile ECC memory bit error counts at job start time.</i></p> <pre>&gt; msub -l nodes=4:gpus=2+1:ppn=2,walltime=600</pre>
	<p><i>Submits a job that requests two different types of tasks, the first is 4 tasks, each with 1 processor and 2 gpus, and the second is 1 task with 2 processors. Each GPU is dedicated exclusively to one task process and the ECC memory bit error counters are not reset.</i></p>

GRES and SOFTWARE	
<b>Format</b>	Percent sign (%) delimited list of generic resources where each resource is specified using the format <RESTYPE> [ {+   : } <COUNT> ]
<b>Description</b>	Indicates generic resources required by the job. If the generic resource is node-locked, it is a per-task count. If a <COUNT> is not specified, the resource count defaults to 1.
<b>Example</b>	<pre>&gt; qsub -W x=GRES:tape+2%matlab+3 testj.txt</pre> <p><b>i</b> When specifying more than one generic resource with -l, use the percent (%) character to delimit them.</p> <pre>&gt; qsub -l gres=tape+2%matlab+3 testj.txt &gt; qsub -l software=matlab:2 testj.txt</pre>



HOSTLIST	
<b>Format</b>	Comma (,) or plus (+) delimited list of hostnames. Ranges and regular expressions are supported in <i>msub</i> only.
<b>Description</b>	<p>Indicates an <i>exact set</i>, <i>superset</i>, or <i>subset</i> of nodes on which the job must run. Use the caret (^) or asterisk (*) characters to specify a host list as <i>superset</i> or <i>subset</i> respectively.</p> <p>An exact set is defined without a caret or asterisk. An exact set means <i>all</i> the hosts in the specified hostlist must be selected for the job.</p> <p>A subset means the specified hostlist is used first to select hosts for the job. If the job requires more hosts than are in the subset hostlist, they will be obtained from elsewhere if possible. If the job does not require all of the nodes in the subset hostlist, it will use only the ones it needs.</p> <p>A superset means the hostlist is the <i>only</i> source of hosts that should be considered for running the job. If the job can't find the necessary resources in the superset hostlist it should <i>not</i> run. No other hosts should be considered in allocating the job.</p>
<b>Examples</b>	<pre>&gt; msub -l hostlist=nodeA+nodeB+nodeE</pre> <pre>hostlist=foo[1-5]</pre> <p><i>This is an exact set of (foo1,foo2,...,foo5). The job must run on all these nodes.</i></p> <pre>hostlist=foo1+foo[3-9]</pre> <p><i>This is an exact set of (foo1,foo3,foo4,...,foo9). The job must run on all these nodes.</i></p> <pre>hostlist=foo[1,3-9]</pre> <p><i>This is an exact set of the same nodes as the previous example.</i></p> <pre>hostlist=foo[1-3]+bar[72-79]</pre> <p><i>This is an exact set of (foo1,foo2,foo3,bar72,bar73,...,bar79). The job must run on all these nodes.</i></p> <pre>hostlist=^node[1-50]</pre> <p><i>This is a superset of (node1,node2,...,node50). These are the only</i></p>

HOSTLIST	
	<p><i>nodes that can be considered for the job. If the necessary resources for the job are not in this hostlist, the job is not run. If the job does not require all the nodes in this hostlist, it will use only the ones that it needs.</i></p>
	<pre>hostlist=*node[15-25]</pre>
	<p><i>This is a subset of (node15,node16,...,node25). The nodes in this hostlist are considered first for the job. If the necessary resources for the job are not in this hostlist, Moab tries to obtain the necessary resources from elsewhere. If the job does not require all the nodes in this hostlist, it will use only the ones that it needs.</i></p>
JGROUP	
<b>Format</b>	<JOBGROUPID>
<b>Description</b>	ID of job group to which this job belongs (different from the GID of the user running the job).
<b>Example</b>	<pre>&gt; msub -l JGROUP=bluegroup</pre>
JOBFLAGS (a.k.a. FLAGS)	
<b>Format</b>	One or more of the following colon delimited job flags including ADVRES [:RSVID], NOQUEUE, NORMSTART, PREEMPTEE, PREEMPTOR, RESTARTABLE, or SUSPENDABLE (see <a href="#">job flag overview</a> for a complete listing).
<b>Description</b>	Associates various flags with the job.
<b>Example</b>	<pre>&gt; qsub -l nodes=1,walltime=3600,jobflags=advres myjob.py</pre>
JOBREJECTPOLICY	
<b>Format</b>	One or more of CANCEL, HOLD, IGNORE, MAIL, or RETRY

JOBREJECTPOLICY	
<b>Default</b>	HOLD
<b>Details</b>	<p>The action to take when the scheduler determines that a job can never run. CANCEL issues a call to the resource manager to cancel the job. HOLD places a batch hold on the job preventing the job from being further evaluated until released by an administrator.</p> <div style="border: 1px solid #004a87; border-radius: 10px; padding: 10px; margin: 10px 0;"> <p><b>i</b> Administrators can dynamically alter job attributes and possibly <i>fix</i> the job with <code>mjobctl -m</code>.</p> </div> <p>With IGNORE, the scheduler will allow the job to exist within the resource manager queue but will neither process it nor report it. MAIL will send email to both the admin and the user when rejected jobs are detected. If RETRY is set, then Moab will allow the job to remain idle and will only attempt to start the job when the policy violation is resolved. Any combination of attributes can be specified.</p> <p>This is a per-job policy specified with <code>msub -l</code>. JOBREJECTPOLICY also exists as a global parameter.</p> <p>Also see <a href="#">QOSREJECTPOLICY</a>.</p>
<b>Example</b>	<pre>&gt; msub -l jobrejectpolicy=cancel:mail</pre>

MAXMEM	
<b>Format</b>	<INTEGER> (in megabytes)
<b>Description</b>	Maximum amount of memory the job can consume across all tasks before the <a href="#">JOBMEM</a> action is taken.
<b>Example</b>	<pre>&gt; qsub -l x=MAXMEM:1000mb bw.cmd</pre> <div style="border: 1px dashed #ccc; border-radius: 10px; padding: 10px; margin: 10px 0;"> <p><i>If a <a href="#">RESOURCELIMITPOLICY</a> is set for per-job memory utilization, its action will be taken when this value is reached.</i></p> </div>

MAXPROC	
<b>Format</b>	<INTEGER>

MAXPROC	
<b>Description</b>	Maximum CPU load the job can consume across all tasks before the <b>JOBPROC</b> action is taken.
<b>Example</b>	<pre>&gt; qsub -W x=MAXPROC:4 bw.cmd</pre> <p><i>If a <b>RESOURCELIMITPOLICY</b> is set for per-job processor utilization, its action will be taken when this value is reached.</i></p>

MEM	
<b>Format</b>	<INTEGER>
<b>Description</b>	Specify the maximum amount of physical memory used by the job. If you do not specify MB or GB, Moab uses bytes if your resource manager is Torque, and MB if your resource manager is Native.
<b>Example</b>	<pre>&gt; msub -l nodes=4:ppn=2,mem=1024mb</pre> <p><i>The job must have 4 compute nodes with 2 processors per node. The job is limited to 1024 MB of memory.</i></p>

MICs	
<b>Format</b>	<pre>msub -l nodes=&lt;VALUE&gt;:ppn=&lt;VALUE&gt;:mics=&lt;VALUE&gt; [:mode]</pre> <p>Where <code>mode</code> is one of:</p> <ul style="list-style-type: none"> <li><i>exclusive</i> - The default setting. The MIC is used exclusively by one process thread.</li> <li><i>exclusive_thread</i> - The MIC is used exclusively by one process thread.</li> <li><i>exclusive_process</i> - The MIC is used exclusively by one process regardless of process thread.</li> </ul> <p>Moab passes the <code>mode</code> portion of the request to Torque for processing.</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> Moab does not support requesting MICs as a GRES. Submitting <code>msub -l gres=mics:x</code> does not work.</p> </div>
<b>Description</b>	Moab schedules MICs as a special type of <a href="#">node-locked generic resources</a> . When <a href="#">Torque reports MICs</a> to Moab, Moab can schedule jobs and correctly

MICs	
	assign MICs to ensure that jobs are scheduled efficiently. To have Moab schedule MICs, configure them in Torque then submit jobs using the 'MIC' attribute. Moab automatically parses the 'MIC' attribute and assigns them in the correct manner.
<b>Examples</b>	<pre>&gt; msub -l nodes=2:ppn=2:mics=1:exclusive_process</pre> <p><i>Submits a job that requests 2 tasks, 2 processors and 1 MIC per task (2 MICs total). Each MIC runs only threads related to the task.</i></p> <pre>&gt; msub -l nodes=4:mics=1,tpn=2</pre> <p><i>Submits a job that requests 4 tasks, 1 MIC per node (4 MICs total), and 2 tasks per node. Each MIC is dedicated exclusively to one task process.</i></p> <pre>&gt; msub -l nodes=4:mics=1</pre> <p><i>Submits a job that requests 4 tasks, 1 processor and 1 MIC per task (4 MICs total). Each MIC is dedicated exclusively to one task process.</i></p> <pre>&gt; msub -l nodes=4:mics=2+1:ppn=2,walltime=600</pre> <p><i>Submits a job that requests two different types of tasks, the first is 4 tasks, each with 1 processor and 2 MICs, and the second is 1 task with 2 processors. Each MIC is dedicated exclusively to one task process.</i></p>
MINPREEMPTTIME	
<b>Format</b>	[ [DD:]HH:]MM:]SS
<b>Description</b>	<p>Minimum time job must run before being eligible for preemption.</p> <p> Can only be specified if associated QoS allows per-job preemption configuration by setting the <a href="#">preemptconfig</a> flag.</p>
<b>Example</b>	<pre>&gt; qsub -l minpreempttime=900 bw.cmd</pre> <p><i>Job cannot be preempted until it has run for 15 minutes.</i></p>

MINPROCSPEED	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	Minimum <a href="#">processor speed</a> (in MHz) for every node that this job will run on.
<b>Example</b>	<pre>&gt; qsub -W x=MINPROCSPEED:2000 bw.cmd</pre> <p><i>Every node that runs this job must have a processor speed of at least 2000 MHz.</i></p>

MINWCLIMIT	
<b>Format</b>	[ [DD:]HH:]MM:]SS
<b>Default</b>	---
<b>Description</b>	Minimum wallclock limit job must run before being eligible for extension (see <a href="#">JOBEXTENDDURATION</a> or <a href="#">JOBEXTENDSTARTWALLTIME</a> ).
<b>Example</b>	<pre>&gt; qsub -l minwclimit=300,walltime=16000 bw.cmd</pre> <p><i>Job will run for at least 300 seconds but up to 16,000 seconds if possible (without interfering with other jobs).</i></p>

MSTAGEIN	
<b>Format</b>	[<SRCURL> [   <SRCRUL> . . . ] %] <DSTURL>
<b>Description</b>	<p>Indicates a job has <a href="#">data staging</a> requirements. The source URL(s) listed will be transferred to the execution system for use by the job. If more than one source URL is specified, the destination URL must be a directory.</p> <p>The format of &lt;SRCURL&gt; is: [PROTO://] [HOST] [:PORT] [/PATH] where the path is local.</p> <p>The format of &lt;DSTURL&gt; is: [PROTO://] [HOST] [:PORT] [/PATH] where the path is remote.</p>

## MSTAGEIN

### Example

```
> msub -Wx='mstagein=file://$HOME/helperscript.sh|file:///home/dev/
datafile.txt%ssh://host/home/dev/' script.sh
```

Copy helperscript.sh and datafile.txt from the local machine to /home/dev/ on host for use in execution of script.sh. \$HOME is a path containing a preceding / (i.e., /home/adaptive)

## MSTAGEOUT

### Format

```
[<SRCURL> [ | <SRCRUL> . . . ] % ] <DSTURL>
```

### Description

Indicates whether a job has [data staging](#) requirements. The source URLs listed will be transferred from the execution system after the completion of the job. If more than one source URL is specified, the destination URL must be a directory.

The format of <SRCURL> is: [PROTO://] [HOST] [:PORT] [/PATH] where the path is remote.

The format of <DSTURL> is: [PROTO://] [HOST] [:PORT] [/PATH] where the path is local.

PROTO can be any of the following protocols: ssh, file, or gsiftp.

HOST is the name of the host where the file resides.

PATH is the path of the source or destination file. The destination path can be a directory when sending a single file and must be a directory when sending multiple files. If a directory is specified, it must end with a forward slash (/).

Valid variables include:

\$JOBID

\$HOME - Path the script was run from

\$RHOME - Home dir of the user on the remote system

\$SUBMITHOST

\$DEST - This is the Moab where the job will run

\$LOCALDATASTAGEHEAD



If no destination is given, the protocol and file name will be set to the same as the source.



The \$RHOME (remote home directory) variable is for when a user's home directory on the compute node is different than on the submission host.

## MSTAGEOUT

### Example

```
> msub -W x='mstageout=ssh://$DEST/$HOME/resultfile1.txt|
ssh://host/home/dev/resultscript.sh%file:///home/dev/' script.sh
```

*Copy resultfile1.txt and resultscript.sh from the execution system to /home/dev/ after the execution of script.sh is complete. \$HOME is a path containing a preceding / (i.e., /home/adaptive).*

## NACCESSPOLICY

### Format

One of **SHARED**, **SINGLEJOB**, **SINGLETASK**, **SINGLEUSER**, or **UNIQUEUSER**

### Description

Specifies how node resources should be accessed. See [Node Access Policies](#) for more information.

**i** The `naccesspolicy` option can only be used to make node access more constraining than is specified by the system, partition, or node policies. If the effective node access policy is `shared`, `naccesspolicy` can be set to `singleuser`, if the effective node access policy is `singlejob`, `naccesspolicy` can be set to `singletask`.

### Example

```
> qsub -l naccesspolicy=singleuser bw.cmd
```

```
> bsub -ext naccesspolicy=singleuser lancer.cmd
```

*Job can only allocate free nodes or nodes running jobs by same user.*

```
> qsub -l naccesspolicy=singlejob jobscript.sh
# OR
> qsub -W x=naccesspolicy:singlejob jobscript.sh
```

*Jobs can only run on specific nodes; regardless if the machine has free cores.*

## NALLOCPOLICY

### Format


One of the valid settings for the parameter **NODEALLOCATIONPOLICY**

### Description

Specifies how node resources should be selected and allocated to the job.



NALLOCPOLICY	
	See <a href="#">Node Allocation Policies</a> for more information.
<b>Example</b>	<pre>&gt; qsub -l nallocpolicy=minresource bw.cmd</pre> <p><i>Job should use the minresource node allocation policy.</i></p>

NCPUS	
<b>Format</b>	<INTEGER>
<b>Description</b>	<p>The number of processors in one task where a task cannot span nodes. If NCPUS is used, then the resource manager's <a href="#">SUBMITPOLICY</a> should be set to NODECENTRIC to get correct behavior. <code>-l ncpus=&lt;#&gt;</code> is equivalent to <code>-l nodes=1:ppn=&lt;#&gt;</code> when <a href="#">JOBNODEMATCHPOLICY</a> is set to EXACTNODE. NCPUS is used when submitting jobs to an SMP. When using <a href="#">GPUs</a> to submit to an SMP, use <code>-l ncpus=&lt;#&gt;:GPUs=&lt;#&gt;</code>.</p> <p> You cannot request both ncpus and nodes in the same job.</p>

NMATCHPOLICY	
<b>Format</b>	One of the valid settings for the parameter <a href="#">JOBNODEMATCHPOLICY</a>
<b>Description</b>	Specifies how node resources should be selected and allocated to the job.
<b>Example</b>	<pre>&gt; qsub -l nodes=2 -W x=nmatchpolicy:exactnode bw.cmd</pre> <p><i>Job should use the EXACTNODEJOBNODEMATCHPOLICY.</i></p>

NODESET	
<b>Format</b>	<SETTYPE>:<SETATTR>[:<SETLIST>]
<b>Description</b>	Specifies nodeset constraints for job resource allocation See the <a href="#">NodeSet Overview</a> for more information.
<b>Example</b>	<pre>&gt; qsub -l nodeset=ONEOF:FEATURE:fastos:hiprio:bigmem bw.cmd</pre>

NODESETCOUNT	
<b>Format</b>	<INTEGER>
<b>Description</b>	Specifies how many node sets a job uses.
<b>Example</b>	<pre>&gt; msub -l nodesetcount=2</pre>

NODESETDELAY	
<b>Format</b>	[ [ [DD:] HH:] MM:] SS
<b>Description</b>	Causes Moab to attempt to span a job evenly across nodesets unless doing so delays the job beyond the requested NODESETDELAY.
<b>Example</b>	<pre>&gt; qsub -l nodesetdelay=300,walltime=16000 bw.cmd</pre>

NODESETOPTIONAL	
<b>Format</b>	<BOOLEAN>
<b>Description</b>	Specifies whether the nodeset constraint is optional. See the <a href="#">NodeSet Overview</a> for more information. <div style="border: 1px solid black; border-radius: 5px; padding: 5px; margin-top: 5px;">  Requires SCHEDCFG[] <b>FLAGS=allowperjobnodesetisoptional</b>.         </div>
<b>Example</b>	<pre>&gt; msub -l nodesetisoptional=true bw.cmd</pre>

OPSYS	
<b>Format</b>	<OperatingSystem>
<b>Description</b>	The job's required operating system.
<b>Example</b>	<pre>&gt; qsub -l nodes=1,opsys=rh73 chem92.cmd</pre>

PARTITION	
<b>Format</b>	<STRING> [:<STRING>] . . .
<b>Description</b>	<p>The <a href="#">partition</a> (or partitions) where the job must run.</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> The job must have access to this partition based on system wide or credential based partition access lists.</p> </div>
<b>Example</b>	<pre>&gt; qsub -l nodes=1,partition=math:geology</pre> <div style="border: 1px dashed black; padding: 5px; margin-top: 10px;"> <p><i>The job must only run in the math partition or the geology partition.</i></p> </div>

PMEM	
<b>Format</b>	<INTEGER>
<b>Description</b>	The maximum amount of physical memory used by any single process of the job.
<b>Example</b>	<pre>&gt; msub -l nodes=4:ppn=2,pmem=1024mb</pre> <div style="border: 1px dashed black; padding: 5px; margin-top: 10px;"> <p><i>The job must have 4 compute nodes with 2 processors per node, and each process of the job is limited to 1024 MB of physical memory.</i></p> </div>

PREF	
<b>Format</b>	[{feature variable}:]<STRING>[:<STRING>] . . .
	<div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> If feature or variable are not specified, then feature is assumed.</p> </div>
<b>Description</b>	<p>Specifies which <a href="#">node features</a> are preferred by the job and should be allocated if available. If preferred node criteria are specified, Moab favors the allocation of matching resources but is not bound to only consider these resources.</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> Preferences are not honored unless the <a href="#">node allocation</a> policy is set to <b>PRIORITY</b> and the <b>PREF</b> priority component is set within the node's <b>PRIORITYF</b> attribute.</p> </div>

PREF	
<b>Example</b>	<pre>&gt; qsub -l nodes=1,pref=bigmem</pre> <p><i>The job can run on any nodes but prefers to allocate nodes with the bigmem feature.</i></p>
PROCS	
<b>Format</b>	<INTEGER>
<b>Description</b>	<p>Requests a specific amount of processors for the job. Instead of users trying to determine the amount of nodes they need, they can instead decide how many processors they need and Moab will automatically request the appropriate amount of nodes from the RM. This also works with feature requests, such as <code>procs=12[:feature1[:feature2[-]]]</code>.</p> <p><b>i</b> Using this resource request overrides any other processor or node related request, such as <code>nodes=4</code>.</p>
<b>Example</b>	<pre>&gt; msub -l procs=32 myjob.pl</pre> <p><i>Moab will request as many nodes as is necessary to meet the 32-processor requirement for the job.</i></p>
PROLOGUE	
<b>Format</b>	<STRING>
<b>Description</b>	<p>Specifies a user owned prologue script, which will be run after the system prologue and <code>prologue.user</code> scripts at the beginning of a job. The syntax is <code>prologue=&lt;file&gt;</code>. The file can be designated with an absolute or relative path.</p> <p><b>i</b> This parameter works only with Torque.</p>
<b>Example</b>	<pre>&gt; msub -l prologue=prologue_script.sh job.s</pre>

PVMEM	
<b>Format</b>	<INTEGER>
<b>Description</b>	Specify the maximum amount of virtual memory used by any single process in the job.
<b>Example</b>	<pre>&gt; msub -l nodes=4:ppn=2,pvmem=1024mb</pre> <p><i>The job must have 4 compute nodes with 2 processors per node, and each process of the job is limited to 1024 MB of virtual memory.</i></p>

QoS	
<b>Format</b>	<STRING>
<b>Description</b>	Requests the specified QoS for the job.
<b>Example</b>	<pre>&gt; qsub -l walltime=1000,qos=highprio biojob.cmd</pre>


QUEUEJOB	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	TRUE
<b>Description</b>	Indicates whether or not the scheduler should queue the job if resources are not available to run the job immediately.
<b>Example</b>	<pre>&gt; msub -l nodes=1,queuejob=false test.cmd</pre>

REQATTR	
<b>Format</b>	Required node attributes with version number support: reqattr=[<must must not should should not>]:<ATTRIBUTE> [{>= > <= < =}<VERSION>]

REQATTR	
<b>Description</b>	<p>Indicates required node attributes. Values can include letters, numbers, dashes, underscores, periods, and spaces.</p> <p>You can choose one of four requirement types for each node attribute you request:</p> <ul style="list-style-type: none"> <li>• <code>must</code> – The node on which this job runs must include the attribute at the value specified. If no node matches this requirement, Moab will not schedule the job.</li> <li>• <code>must not</code> – The node on which this job runs must not include the attribute at the value specified. If no node matches this requirement, Moab will not schedule the job.</li> <li>• <code>should</code> – If possible, the node on which this job runs should include the attribute at the value specified. If no node matches this requirement, Moab selects a node without it.</li> <li>• <code>should not</code> – If possible, the node on which this job runs should not include the attribute at the value specified. If no node matches this requirement, Moab selects a node without it.</li> </ul> <p>If you do not specify a requirement type, Moab assumes 'must.'</p> <p>For information about using <code>reqattr</code> to request dynamic features, see <a href="#">Configuring dynamic features in Torque and Moab</a>.</p>
<b>Example</b>	<pre>&gt; qsub -l reqattr=matlab=7.1 testj.txt</pre>

RESFAILPOLICY	
<b>Format</b>	One of CANCEL, HOLD, IGNORE, NOTIFY, or REQUEUE
<b>Description</b>	<p>The action to take on an executing job if one or more allocated nodes fail. This setting overrides the global value specified with the <a href="#">NODEALLOCRESFAILUREPOLICY</a> parameter.</p>
<b>Example</b>	<pre>&gt; msub -l resfailpolicy=ignore</pre> <p><i>For this particular job, ignore node failures.</i></p>

RMTYPE	
<b>Format</b>	<STRING>
<b>Description</b>	One of the resource manager types currently available within the cluster or grid. Typically, this is PBS.
<b>.Example</b>	<pre>&gt; msub -l rmttype=pbs</pre> <p>Only run job on a PBS destination resource manager.</p>

SIGNAL	
<b>Format</b>	<INTEGER> [ @<OFFSET> ]
<b>Description</b>	<p>The pretermination signal to be sent to a job prior to it reaching its walltime limit or being terminated by Moab. The optional offset value specifies how long before job termination the signal should be sent. By default, the pretermination signal is sent one minute before a job is terminated.</p> <div style="border: 1px solid #004a7c; border-radius: 5px; padding: 5px; background-color: #e6f2ff;"> <p> Setting an offset value on a job-end <i>event</i> is not supported.</p> </div>
<b>Example</b>	<pre>&gt; msub -l signal=32@120 bio45.cmd</pre>

SPRIORITY	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	Allows Moab administrators to set a system priority on a job (similar to <a href="#">setspri</a> ). This only works if the job submitter is an administrator.
<b>Example</b>	<pre>&gt; qsub -l nodes=16,spriority=100 job.cmd</pre>

TEMPLATE	
<b>Format</b>	<STRING>
<b>Description</b>	Specifies a job template to be used as a set template. The requested template must have <code>SELECT=TRUE</code> (see <a href="#">Job Templates</a> ).
<b>Example</b>	<pre>&gt; msub -l walltime=1000,nodes=16,template=biojob job.cmd</pre>

TERMTIME	
<b>Format</b>	<TIMESPEC>
<b>Default</b>	0
<b>Description</b>	The time at which Moab should cancel a queued or active job (see <a href="#">Job Deadline Support</a> ).
<b>Example</b>	<pre>&gt; msub -l nodes=10,walltime=600,termtime=12:00_Jun/14 job.cmd</pre>

TPN	
<b>Format</b>	<INTEGER> [+]
<b>Default</b>	0
<b>Description</b>	<p>Tasks per node allowed on allocated hosts. If the plus (+) character is specified, the tasks per node value is interpreted as a minimum tasks per node constraint; otherwise it is interpreted as an exact tasks per node constraint.</p> <p><b>Differences between TPN and PPN:</b></p> <p>There are two key differences between the following: (A) <code>qsub -l nodes=12:ppn=3</code> and (B) <code>qsub -l nodes=12,tpn=3</code>.</p> <p>The first difference is that <code>ppn</code> is interpreted as the <i>minimum</i> required tasks per node while <code>tpn</code> defaults to exact tasks per node; case (B) executes the job with exactly 3 tasks on each allocated node while case (A) executes the job with at least 3 tasks on each allocated node-  <code>nodeA: 4, nodeB: 3, nodeC: 5</code></p> <p>The second major difference is that the line, <code>nodes=X:ppn=Y</code> actually</p>



TPN	
	<p>requests <math>X*Y</math> tasks, whereas <code>nodes=X, tpn=Y</code> requests only <math>X</math> tasks.</p> <p><b>TPN with Torque as an RM:</b></p> <p>Moab interprets <code>nodes</code> loosely as <code>procs</code>. Torque interprets <code>nodes</code> as the number of nodes from the actual number of nodes that you have in your <code>nodes</code> file, not your total number of <code>procs</code>. This means that if Torque is your resource manager and you specify <code>msub -l nodes=16:tpn=8</code> but do not have 16 nodes, Torque will not run the job. Instead, you should specify <code>msub -l procs=16:tpn=8</code>.</p> <p>To resolve the problem long term, you can also set <code>server resources_available.nodect</code> to the total number of <code>procs</code> in your system and use <code>msub -l nodes=16:tpn=8</code> as you would in a non-Torque Moab environment. See 'resources_available' in the <i>Torque Administrator Guide</i> for more information.</p>
<b>Example</b>	<pre>&gt; msub -l nodes=10,walltime=600,tpn=4 job.cmd</pre>

TRIG	
<b>Format</b>	<TRIGSPEC>
<b>Description</b>	<p>Adds triggers to the job (see <a href="#">Creating a Trigger</a> for specific syntax).</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> Job triggers can only be specified if allowed by the QoS flag <a href="#">trigger</a>. See <a href="#">Enabling Job Triggers</a> for more information.</p> </div>
<b>Example</b>	<pre>&gt; qsub -l trig=etype=start\&amp;atype=exec\&amp;action="/tmp/email.sh job.cmd"</pre>

TRL (Format 1)	
<b>Format</b>	<INTEGER> [ @<INTEGER> ] [ :<INTEGER> [ @<INTEGER> ] ] . . .
<b>Default</b>	0
<b>Description</b>	<p>Specifies alternate task requests with their optional walltimes. See <a href="#">Malleable Jobs</a>.</p>
<b>Example</b>	<pre>&gt; msub -l trl=2@500:4@250:8@125:16@62 job.cmd</pre>

TRL (Format 1)	
	<p>or</p> <pre>&gt; qsub -l trl=2:3:4</pre>

TRL (Format 2)	
<b>Format</b>	<INTEGER>-<INTEGER>
<b>Default</b>	0
<b>Description</b>	Specifies a range of task requests that require the same walltime. See <a href="#">Malleable Jobs</a> .
<b>Example</b>	<pre>&gt; msub -l trl=32-64 job.cmd</pre> <div style="border: 1px solid #005596; padding: 5px; margin-top: 10px;"> <p><b>i</b> For optimization purposes Moab does not perform an exhaustive search of all possible values but will at least do the beginning, the end, and 4 equally distributed choices in between.</p> </div>

VAR	
<b>Format</b>	<ATTR> [ :<VALUE> ]
<b>Description</b>	Adds a generic variable or variables to the job.
<b>Example</b>	<pre>&gt; msub -l VAR=testvar1:testvalue1</pre> <p style="text-align: center;"><i>Single variable</i></p> <pre>&gt; msub -l VAR=testvar1:testvalue1+testvar2:testvalue2+testvar3:testvalue3</pre> <p style="text-align: center;"><i>Multiple variables</i></p>

VC	
<b>Format</b>	VC=<NAME>

VC	
<b>Description</b>	Submits the job or workflow to a virtual container (VC).
<b>Example</b>	<code>vc=vc13</code>

VMEM	
<b>Format</b>	<INTEGER>
<b>Description</b>	Specify the maximum amount of virtual memory used by all concurrent processes in the job.
<b>Example</b>	<pre>&gt; msub -l nodes=4:ppn=2,vmem=1024mb</pre> <p><i>The job must have 4 compute nodes with 2 processors per node, and the job is limited to 1024 MB of virtual memory.</i></p>

### 11.3.3 Resource Manager Extension Examples

If more than one extension is required in a given job, extensions can be concatenated with a semicolon separator using the format <ATTR>:<VALUE> [ ;<ATTR>:<VALUE> ] . . .

Example 11-1:

```
#@comment="HOSTLIST:node1,node2;QOS:special;SID:silverA"
```

*Job must run on nodes node1 and node2 using the QoS special. The job is also associated with the system ID silverA allowing the silver daemon to monitor and control the job.*

Example 11-2:

```
# PBS -W x="\NODESET:ONEOF:NETWORK;DMEM:64\""
```

*Job will have resources allocated subject to network based nodeset constraints. Further, each task will dedicate 64 MB of memory.*

Example 11-3:

```
> qsub -l nodes=4,walltime=1:00:00 -W x="FLAGS:ADVRES:john.1"
```

*Job will be forced to run within the john.1 reservation.*

### 11.3.4 Configuring dynamic features in Torque and Moab

Used together, the `reqattr` RM extension and Torque `$varattr` parameter allow you to create jobs that request resources that may change or disappear. For example, if you wanted a job to request a certain version of Octave but different versions are configured on each node and updated at any time, you can create a script that searches for the feature and version on the nodes at a specified interval. Your Moab job can then retrieve the dynamic node attributes from the latest poll and use them for scheduling.

This functionality is available when you use the Torque `$varattr` parameter to configure a script that regularly retrieves updates on the nodes' feature(s) and the `reqattr` RM extension to require a feature with a certain value.

#### To set up a dynamic feature in Torque and Moab

1. Create a script that pulls the information you need. For instance, the following script pulls the version of Octave on each node and prints it:

```
#!/bin/bash
# pull the version string for octave and print it for $varattr
version_str='octave -v | grep version'
[[ $version_str =~ ([[[:digit:]]][[:digit:]]][[:digit:]] ) ]]
echo "octave=${BASH_REMATCH[1]}"
```

2. Use the Torque `$varattr` parameter to configure the script. Specify both the number of seconds between each time Torque runs the script and the path to the script. If you set the seconds to `-1`, the script will run just once. You can include arguments if desired. In the following example, the `varattr` parameter specifies that Torque calls the Octave script every 30 seconds:

```
$varattr 30 /usr/local/scripts/octave.sh
```

3. Submit your job in Moab, specifying `reqattr` as a resource. In this example, the job requests a node where the octave feature has a value of 3.2.4 (that the node has Octave version 3.2.4 installed):

```
> msub -l reqattr=octave=3.2.4 myJob.sh
```

*Your job requests a node with Octave version 3.2.4. Torque passes the most recent (pulled within the last 30 seconds) version of Octave on each node. Moab then schedules the job on a node that currently has Octave 3.2.4.*

---

#### Related Topics

- [Resource Manager Overview](#)
- `$varattr` in the Torque Administrator Guide

## 11.4 Adding New Resource Manager Interfaces

Moab is designed to interface with multiple resource management systems. Some of these interact through a resource manager specific interface (Torque), while others interact through generalized interfaces such as SSS or Wiki (see the [Wiki Overview](#)). For most resource managers, either route is possible depending on where it is easiest to focus development effort. Use of Wiki generally requires modifications on the resource manager side while creation of a new resource manager specific Moab interface would require more changes to Moab modules.

Regardless of the interface approach selected, adding support for a new resource manager is typically a straightforward process for about 95% of all supported features. The final 5% of features usually requires a bit more effort as each resource manager has a number of distinct concepts that must be addressed.

### 11.4.1 Resource Manager Specific Interfaces

If you require tighter integration and need additional instruction, see [Managing Resources Directly with the Native Interface](#). If you would like consultation on support for a new resource manager type, please [contact](#) Adaptive Computing.

### 11.4.2 Wiki Interface

The Wiki interface is already defined as a resource manager type, so no modifications are required within Moab. Additionally, no resource manager specific library or header file is required. However, within the resource manager, internal job and node objects and attributes must be manipulated and placed within Wiki based interface concepts as defined in the [Wiki Overview](#). Additionally, resource manager parameters must be created to allow a site to configure this interface appropriately.

### 11.4.3 SSS Interface

The SSS interface is an XML based generalized resource manager interface. It provides an extensible, scalable, and secure method of querying and modifying general workload and resource information.

---

#### Related Topics

- [Creating New Tools within the Native Resource Manager Interface](#)

## 11.5 Managing Resources Directly with the Native Interface

In this topic:

- [11.5.1 Native Interface Overview - page 690](#)
- [11.5.2 Configuring the Native Interface - page 691](#)
- [11.5.3 Generating Cluster Query Data - page 692](#)
- [11.5.4 Interfacing with FlexNet \(Formerly FLEXlm\) - page 692](#)
- [11.5.5 Interfacing to Nagios - page 694](#)
- [11.5.6 Configuring Resource Types - page 695](#)
- [11.5.7 Creating New Tools to Manage the Cluster - page 696](#)

### 11.5.1 Native Interface Overview

The Native interface allows a site to augment or even fully replace a resource manager for managing resources. In some situations, the full capabilities of the resource manager are not needed and a lower cost or lower overhead alternative is preferred. In other cases, the nature of the environment might make use of a resource manager impossible due to lack of support. Still, in other situations it is desirable to provide information about additional resource attributes, constraints, or state from alternate sources.

In any case, Moab provides the ability to directly query and manage resources along side of or without the use of a resource manager. This interface, called the NATIVE interface can also be used to launch, cancel, and otherwise manage jobs. This NATIVE interface offers several advantages including the following:

- No cost associated with purchasing a resource manager
- No effort required to install or configure the resource manager
- Ability to support abstract resources
- Ability to support abstract jobs
- Ability to integrate node availability information from multiple sources
- Ability to augment node configuration and utilization information provided by a resource manager

However, the NATIVE interface might also have some drawbacks:

- No support for standard job submission languages
- Limited default configured and utilized resource tracking (additional resource tracking available with additional effort)

At a high level, the native interface works by launching threaded calls to perform standard resource manager activities such as managing resources and jobs. The desired calls are configured within Moab and used whenever an action or updated information is required.

## 11.5.2 Configuring the Native Interface

Using the native interface consists of defining the interface type and location. As mentioned earlier, a single object can be fully defined by multiple interfaces simultaneously with each interface updating a particular aspect of the object.

### 11.5.2.A Configuring the Resource Manager

The Native resource manager must be configured using the **RMCFG** parameter. To specify the native interface, the **TYPE** attribute must be set to **NATIVE**.

```
RMCFG[local] TYPE=NATIVE
RMCFG[local] CLUSTERQUERYURL=exec:///tmp/query.sh
```

### 11.5.2.B Reporting Resources

To indicate the source of the resource information, the **CLUSTERQUERYURL** attribute of the **RMCFG** parameter should be specified. This attribute is specified as a URL where the protocols **EXEC** and **FILE** are allowed. If a protocol is not specified, the protocol **EXEC** is assumed.

Format	Description
<b>EXEC</b>	Execute the script specified by the URL path. Use the script stdout as data.
<b>FILE</b>	Load the file specified by the URL path. Use the file contents as data.

Moab considers a NativeRM script to have failed if it returns with a non-zero exit code or if the **CHILDSTDERRCHECK** parameter is set and its appropriate conditions are met. In addition, the NativeRM script associated with a job submit URL will be considered as having failed if its standard output stream contains the text **ERROR**.

This simple example queries a file on the server for information about every node in the cluster. This differs from Moab remotely querying the status of each node individually.

```

RMCFG[local]    TYPE=NATIVE
RMCFG[local]    CLUSTERQUERYURL=file:///tmp/query.txt

```

## 11.5.3 Generating Cluster Query Data

### 11.5.3.A Flat Cluster Query Data

If the EXEC or FILE protocol is specified in the CLUSTERQUERYURL attribute, the data should provide flat text strings indicating the state and attributes of the node. The format follows the [Moab Resource Manager Language Data Format](#) where attributes are delimited by white space rather than ';' (see [Query Resources Data Format](#)):

Describes any set of node attributes with format: <NAME><ATTR>=<VAL>  
[<ATTR>=<VAL>] . . .

<NAME>	Name of node
<ATTR>	Node attribute
<VAL>	Value of node attribute

```
n17 CPROC=4 AMEMORY=100980 STATE=idle
```

## 11.5.4 Interfacing with FlexNet (Formerly FLEXlm)

Moab can interface with FlexNet to provide scheduling based on [License Management](#) availability. Informing Moab of license dependencies can reduce the number of costly licenses required by your cluster by allowing Moab to intelligently schedule around license limitations.

Provided with Moab in the tools directory is a Perl script, `license.mon.flexLM.pl`. This script queries a FlexNet license server and gathers data about available licenses. This script then formats this data for Moab to read through a native interface. This script can easily be used by any site to help facilitate FlexNet integration—the only modification necessary to the script is setting the `@FLEXlmCmd` to specify the local command to query FlexNet. To make this change, edit `license.mon.flexLM.pl` and, near the top of the file, look for the line:

```
my @FLEXlmCmd = ("SETME");
```

Set the `@FLEXlmCmd` to the appropriate value for your system to query a license server and license file (if applicable). If `lmutil` is not in the `PATH` variable, specify its full path.



Using the `lmutil -a` argument will cause it to report all licenses. The `-c` option is used to specify a license file on a remote server.

```
<path_to_lmstat>/lmstat -c port@host -a
my @FLEXlmCmd = ("<path_to_lmstat>/lmstat -c port@host -a");
```

**i** The `@` specifying the port `@servername` must be escaped.

To test this script, run it manually. If working correctly, it will produce output similar to the following:

```
> ./license.mon.flexLM.pl
GLOBAL UPDATETIME=1104688300 STATE=idle ARES=autoCAD:130,idl_mpeg:160
CRES=autoCAD:200,idl_mpeg:330
```

If the output looks incorrect, set the `$LOGLEVEL` variable inside of `license.mon.flexLM.pl`, run it again, and address the reported failure.

Once the license interface script is properly configured, the next step is to add a *license* native resource manager to Moab via the `moab.cfg` file:

```
RMCFG[FLEXlm] TYPE=NATIVE RESOURCETYPE=LICENSE
RMCFG[FLEXlm] CLUSTERQUERYURL=exec://$TOOLSDIR/flexlm/license.mon.flexLM.pl
...
```

Once this change is made, restart Moab. The command `mdiag -R` can be used to verify that the resource manager is properly configured and is in the state `Active`. Detailed information regarding configured and utilized licenses can be viewed by issuing the `mdiag -n`. Floating licenses (non-node-locked) will be reported as belonging to the `GLOBAL` node.

**i** Due to the inherent conflict with the plus sign (+), the provided license manager script replaces occurrences of the plus sign in license names with the underscore symbol (`_`). This replacement requires that licenses with a plus sign in their names be requested with an underscore in place of any plus signs.

#### 11.5.4.A Interfacing to Multiple License Managers Simultaneously

If multiple license managers are used within a cluster, Moab can interface to each of them to obtain the needed license information. In the case of FlexNet, this can be done by making one copy of the `license.mon.flexLM.pl` script for each license manager and configuring each copy to point to a different license manager. Then, within Moab, create one native resource manager interface for each license manager and point it to the corresponding script as in the following example:

```
RMCFG[FLEXlm1] TYPE=NATIVE RESOURCETYPE=LICENSE
RMCFG[FLEXlm1] CLUSTERQUERYURL=exec://$TOOLSDIR/flexlm/license.mon.flexLM1.pl
```

```

RMCFG [FLEXlm2]    TYPE=NATIVE RESOURCETYPE=LICENSE
RMCFG [FLEXlm2]    CLUSTERQUERYURL=exec://$TOOLS DIR/flexlm/license.mon.flexLM2.pl
RMCFG [FLEXlm3]    TYPE=NATIVE RESOURCETYPE=LICENSE
RMCFG [FLEXlm3]    CLUSTERQUERYURL=exec://$TOOLS DIR/flexlm/license.mon.flexLM3.pl
...

```

**i** For an overview of license management, including job submission syntax, see [License Management](#).

**i** It may be necessary to increase the default limit, `MMAX_GRES`. See [Adjusting Default Limits](#) for more implementation details.

## 11.5.5 Interfacing to Nagios

Moab can interface with Nagios to provide scheduling based on network hosts and services availability.

Nagios installation and configuration documentation can be found at [Nagios.org](#).

Provided with Moab in the `tools` directory is a Perl script, `node.query.nagios.pl`. This script reads the Nagios `status.dat` file and gathers data about network hosts and services. This script then formats data for Moab to read through a native interface. This script can be used by any site to help facilitate Nagios integration. To customize the data that will be formatted for Moab, make the changes in this script.

You may need to customize the associated configuration file in the `etc` directory, `config.nagios.pl`. The `statusFile` line in this script tells Moab where the Nagios `status.dat` file is located. Make sure that the path name specified is correct for your site. Note that the interval that Nagios updates the Nagios `status.dat` file is specified in the Nagios `nagios.cfg` file. Refer to Nagios documentation for further details.

To make these changes, familiarize yourself with the format of the Nagios `status.dat` file and make the appropriate additions to the script to include the desired Moab RM language (formerly WIKI) Interface attributes in the Moab output.

To test this script, run it manually. If working correctly, it will produce output similar to the following:

```

> ./node.query.nagios.pl
gateway STATE=Running
localhost STATE=Running CPULOAD=1.22 ADISK=75332

```

Once the Nagios interface script is properly configured, the next step is to add a Nagios native resource manager to Moab via the `moab.cfg` file:

```

RMCFG [nagios]    TYPE=NATIVE
RMCFG [nagios]    CLUSTERQUERYURL=exec://$TOOLS DIR/node.query.nagios.pl

```


...

Once this change is made, restart Moab. The command `mdiag -R` can be used to verify that the resource manager is properly configured and is in the state `Active`. Detailed information regarding configured Nagios node information can be viewed by issuing the `mdiag -n`.

```
> mdiag -n -v
compute node summary
Name           State  Procs      Memory      Disk      Swap
Speed  Opsys  Arch Par   Load Rsv Classes
Features
gateway        Running  0:0        0:0        0:0        0:0
1.00           -      - dav  0.00  0 -
-
WARNING: node 'gateway' is busy/running but not assigned to an active job
WARNING: node 'gateway' has no configured processors
localhost      Running  0:0        0:0      75343:75347  0:0
1.00           -      - dav  0.48  0 -
-
WARNING: node 'localhost' is busy/running but not assigned to an active job
WARNING: node 'localhost' has no configured processors
-----
Total Nodes: 2 (Active: 2 Idle: 0 Down: 0)
```

## 11.5.6 Configuring Resource Types

Native Resource managers can also perform special tasks when they are given a specific resource type. These types are specified using the `RESOURCETYPE` attribute of the `RMCFG` parameter.

Type	Description
<b>COMPUTE</b>	Normal compute resources (no special handling)
<b>FS</b>	File system resource manager (see <a href="#">Utilizing Multiple Resource Managers</a> for an example)
<b>LICENSE</b>	Software license manager (see <a href="#">Interfacing with FlexNet (Formerly FLEXlm) and License Management</a> ) <div style="border: 1px solid red; padding: 5px; margin-top: 10px;">  If <code>LICENSE</code> is specified, <i>all</i> generic resources reported by the resource manager will be marked as a license; causing them to be tracked by the accounting manager under the Licenses property. See <a href="#">Licenses - page 427</a> for more information on the Licenses property. </div>
<b>NETWORK</b>	Network resource manager

Type	Description
<b>PROV</b>	Provisioning resource manager. This is the RM Moab uses to modify the OS of a node and to power a node on or off.

## 11.5.7 Creating New Tools to Manage the Cluster

Using the scripts found in the `$TOOLS_DIR` (`$INSTDIR/tools`) directory as a template, new tools can be quickly created to monitor or manage most any resource. Each tool should be associated with a particular resource manager service and specified using one of the following resource manager URL attributes.

CLUSTERQUERYURL	
<b>Description</b>	Queries resource state, configuration, and utilization information for compute nodes, networks, storage systems, software licenses, and other resources. For more details, see <a href="#">CLUSTERQUERYURL</a> .
<b>Output</b>	Node status and configuration for one or more nodes. See <a href="#">Query Resources Data Format</a> .
<b>Example</b>	<pre>RMCFG[v-stor] CLUSTERQUERYURL=exec://\$HOME/storquery.pl</pre> <p><i>Moab will execute the <code>storquery.pl</code> script to obtain information about 'v-stor' resources.</i></p>

JOBCANCELURL	
<b>Description</b>	Cancels a job.
<b>Input</b>	<JOBID>
<b>Example</b>	<pre>RMCFG[v-stor] JOBCANCELURL=exec://\$HOME/cancel.pl</pre> <p><i>Moab will execute the <code>cancel.pl</code> script to cancel jobs.</i></p>

JOBMODIFYURL	
<b>Description</b>	Modifies a job or application. For more details, see <a href="#">JOBMODIFYURL</a> .

JOBMODIFYURL	
<b>Input</b>	<code>[-j &lt;JOBEXPR&gt;] [--s[et] --u[nset] --c[lear] --i[ncrement] --d[ecrement]] &lt;ATTR&gt;[=&lt;VALUE&gt;] [&lt;ATTR&gt;[=&lt;VALUE&gt;]]...</code>
<b>Example</b>	<pre>RMCFG[v-stor] JOBMODIFYURL=exec://\$HOME/jobmodify.pl</pre> <p><i>Moab will execute the jobmodify.pl script to modify the specified job.</i></p>

JOBQUEUEURL	
<b>Description</b>	Requeues a job.
<b>Input</b>	<JOBID>
<b>Example</b>	<pre>RMCFG[v-stor] JOBQUEUEURL=exec://\$HOME/requeue.pl</pre> <p><i>Moab will execute the requeue.pl script to requeue jobs.</i></p>

JOBRESUMEURL	
<b>Description</b>	Resumes a suspended job or application.
<b>Input</b>	<JOBID>
<b>Example</b>	<pre>RMCFG[v-stor] JOBRESUMEURL=exec://\$HOME/jobresume.pl</pre> <p><i>Moab will execute the jobresume.pl script to resume suspended jobs.</i></p>

JOBSTARTURL	
<b>Description</b>	Launches a job or application on a specified set of resources.
<b>Input</b>	<JOBID><TASKLIST><USERNAME> [ARCH=<ARCH>] [OS=<OPSYS>] [IDATA=<STAGEINFILEPATH>[, <STAGEINFILEPATH>]...]

JOBSTARTURL	
	[EXEC=<EXECUTABLEPATH>]
<b>Example</b>	<pre>RMCFG[v-stor] JOBSTARTURL=exec://\$HOME/jobstart.pl</pre> <p><i>Moab will execute the jobstart.pl script to execute jobs.</i></p>

JOBSUBMITURL	
<b>Description</b>	Submits a job to the resource manager, but it does not execute the job. The job executes when the JOBSTARTURL is called.
<b>Input</b>	<pre>[ACCOUNT=&lt;ACCOUNT&gt;] [ERROR=&lt;ERROR&gt;] [GATTR=&lt;GATTR&gt;] [GNAM=&lt;GNAM&gt;] [GRES=&lt;GRES&gt;:&lt;Value&gt;[, &lt;GRES&gt;:&lt;Value&gt;]*] [HOSTLIST=&lt;HOSTLIST&gt;] [INPUT=&lt;INPUT&gt;] [IWD=&lt;IWD&gt;] [NAME=&lt;NAME&gt;] [OUTPUT=&lt;OUTPUT&gt;] [RCLASS=&lt;RCLASS&gt;] [REQUEST=&lt;REQUEST&gt;] [RFEATURES=&lt;RFEATURES&gt;] [RMFLAGS=&lt;RMFLAGS&gt;] [SHELL=&lt;SHELL&gt;] [TASKLIST=&lt;TASKLIST&gt;] [TASKS=&lt;TASKS&gt;] [TEMPLATE=&lt;TEMPLATE&gt;] [UNAME=&lt;UNAME&gt;] [VARIABLE=&lt;VARIABLE&gt;] [WCLIMIT=&lt;WCLIMIT&gt;] [ARGS=&lt;Value&gt; [ &lt;Value&gt;]*]</pre> <p><b>i</b> ARGS must be the last submitted attribute because there can be multiple space-separated values for ARGS.</p>
<b>Example</b>	<pre>RMCFG[v-stor] JOBSUBMITURL=exec://\$HOME/jobsubmit.pl</pre> <p><i>Moab submits the job to the jobsubmit.pl script for future job execution.</i></p>

JOBSUSPENDURL	
<b>Description</b>	Suspends in memory an active job or application.
<b>Input</b>	<JOBID>
<b>Example</b>	<pre>RMCFG[v-stor] JOBSUSPENDURL=exec://\$HOME/jobsuspend.pl</pre> <p><i>Moab will execute the jobsuspend.pl script to suspend active jobs.</i></p>

## JOBSUSPENDURL

## NODEMODIFYURL

### Description

Provide method to dynamically modify/provision compute resources including operating system, applications, queues, node features, power states, etc.

**i** Moab blocks scheduling when invoking the script defined by NODEMODIFYURL, If your NODEMODIFYURL script is a long running process (i.e., > 30 seconds or, if defined, longer than the RMCFG TIMEOUT setting), then you must background the process and return quickly.

### Input

<NODEID> [, <NODEID>] [--force] [--set <ATTR>=<VAL> | --clear <ATTR>}

ATTR is one of the node attributes listed in [Query Resources Data Format](#).

### Example

```
RMCFG[warewulf] NODEMODIFYURL=exec://$HOME/provision.pl
```

*Moab will reprovision compute nodes using the provision.pl script.*

## NODEPOWERURL

### Description

Allows Moab to issue IPMI power commands.

### Input

<NODEID> [, <NODEID>] ON | OFF

### Example

```
RMCFG[node17rm] NODEPOWERURL=exec://$TOOLSDIR/ipmi.power.pl
```

*Moab will issue a power command contained in the ipmi.power.pl script.*

## SYSTEMMODIFYURL

### Description

Provide method to dynamically modify aspects of the compute environment that are directly associated with cluster resources. For more details, see [SYSTEMMODIFYURL](#).

SYSTEMQUERYURL	
<b>Description</b>	Provide method to dynamically query aspects of the compute environment that are directly associated with cluster resources. For more details, see <a href="#">SYSTEMQUERYURL</a> .
<b>Input</b>	Default <ATTR> ATTR is one of <code>images</code>
<b>Output</b>	<STRING>
<b>Example</b>	<pre>RMCFG[warewulf] SYSTEMQUERYURL=exec://\$HOME/checkimage.pl</pre> <p><i>Moab will load the list of images available from warewulf using the <code>checkimage.pl</code> script.</i></p>

WORKLOADQUERYURL	
<b>Description:</b>	<p>Provide method to dynamically query the system workload (jobs, services, etc.) of the compute environment that are associated with managed resources.</p> <div style="border: 1px solid black; padding: 5px; margin: 5px 0;"> <p><b>i</b> Job/workload information should be reported back from the URL (script, file, web service, etc.) using the <a href="#">Moab Resource Manager Language Data Format</a> (formerly WIKI).</p> </div> <p>For more details, see <a href="#">WORKLOADQUERYURL</a>.</p>
<b>Output:</b>	<STRING>
<b>Example</b>	<pre>RMCFG[xt] WORKLOADQUERYURL=exec://\$HOME/job.query.xt3.pl</pre> <p><i>Moab will load job/workload information by executing the <code>job.query.xt3.pl</code> script.</i></p>

---

### Related Topics

- [3.7.14 mdiag -R - page 155](#) command (evaluate resource managers)
- [11.7 License Management - page 703](#)
- [Moab Resource Manager Language Data Format](#)



## 11.6 Utilizing Multiple Resource Managers

In this topic:

[11.6.1 Multi-RM Overview - page 701](#)

[11.6.2 Configuring Multiple Independent Resource Manager Partitions - page 701](#)

[11.6.3 Migrating Jobs between Resource Managers - page 701](#)

[11.6.4 Aggregating Information into a Cohesive Node View - page 702](#)

### 11.6.1 Multi-RM Overview

In many instances a site may have certain resources controlled by different resource managers. For example, a site may use a particular resource manager for licensing software for jobs, another resource manager for managing file systems, another resource manager for job control, and another for node monitoring. Moab can be configured to communicate with each of these resource managers, gathering all their data and incorporating such into scheduling decisions. With a more distributed approach to resource handling, failures are more contained and scheduling decisions can be more intelligent.

### 11.6.2 Configuring Multiple Independent Resource Manager Partitions

Moab must know how to communicate with each resource manager. In most instances, this is simply done by configuring a [query command](#).

### 11.6.3 Migrating Jobs between Resource Managers

With multi-resource manager support, a job can be submitted either to a local resource manager queue or to the Moab global queue. In most cases, submitting a job to a resource manager queue constrains the job to only run within the resources controlled by that resource manager. However, if the job is submitted to the Moab global queue, it can use resources of any active resource manager. This is accomplished through job translation and staging.

When Moab evaluates resource availability, it determines the cost in terms of both data and job staging. If staging a job's executable or input data requires a significant amount of time, Moab integrates data and compute resource availability to determine a job's earliest

potential start time on a per resource manager basis and makes an optimal scheduling decision accordingly. If the optimal decision requires a data stage operation, Moab reserves the required compute resources, stages the data, and then starts the job when the required data and compute resources are available.

## 11.6.4 Aggregating Information into a Cohesive Node View

Using the native interface, Moab can actually perform most of these functions without the need for an external resource manager. First, configure the native resource managers:

```
RMCFG[base]      TYPE=PBS
RMCFG[network]  TYPE=NATIVE
RMCFG[network]  CLUSTERQUERYURL=/tmp/network.sh
RMCFG[fs]       TYPE=NATIVE
RMCFG[fs]       CLUSTERQUERYURL=/tmp/fs.sh
```

The network script can be as simple as the following:

```
> _RX=`/sbin/ifconfig eth0 | grep "RX by" | cut -d: -f2 | cut -d' ' -f1`; \
> _TX=`/sbin/ifconfig eth0 | grep "TX by" | cut -d: -f3 | cut -d' ' -f1`; \
> echo `hostname` GMETRIC[netusage]=`echo "$_RX + $_TX" | bc`;
```

The preceding script would output something such as the following:

```
node01 GMETRIC[netusage]=10928374
```

Moab grabs information from each resource manager and includes its data in the final view of the node.

```
> checknode node01
node node01
State:   Running (in current state for 00:00:20)
Configured Resources: PROCS: 2 MEM: 949M SWAP: 2000M disk: 1000000
Utilized Resources: SWAP: 9M
Dedicated Resources: PROCS: 1 disk: 1000
Opsys:   Linux-2.6.5-1.358 Arch:   linux
Speed:   1.00 CPUload:   0.320
Location: Partition: DEFAULT Rack/Slot: NA
Network Load: 464.11 b/s
Network:  DEFAULT
Features: fast
Classes:  [batch 1:2][serial 2:2]
Total Time: 00:30:39 Up: 00:30:39 (100.00%) Active: 00:09:57 (32.46%)
Reservations:
  Job '5452'(x1) -00:00:20 -> 00:09:40 (00:10:00)
JobList: 5452
```

Notice that the Network Load is now being reported along with disk usage.

### 11.6.4.A Example File System Utilization Tracker (per user)

The following configuration can be used to track file system usage on a per user basis:

```

.....
RMCFG[file]      TYPE=NATIVE
RMCFG[file]      RESOURCETYPE=FS
RMCFG[file]      CLUSTERQUERYURL=/tmp/fs.pl
.....

```

Assuming that `/tmp/fs.pl` outputs something of the following **format**:

```

DEFAULT STATE=idle AFS=<fs id="user1" size="789456"></fs><fs
id="user2" size="123456"></fs>

```

This will track disk usage for users *user1* and *user2* every 24 hours.

## 11.7 License Management

In this topic:

[11.7.1 License Management Overview - page 703](#)


[11.7.2 Controlling and Monitoring License Availability - page 704](#)

[11.7.3 Requesting Licenses within Jobs - page 705](#)

### 11.7.1 License Management Overview

Software license management is typically enabled in one of two models: node-locked and floating. Under a node-locked license, use of a given application is constrained to certain hosts. For example, `node013` can support up to two simultaneous jobs accessing application `matlab`. In a floating license model, a limited number of software licenses are made available cluster wide, and these licenses can be used on any combination of compute hosts. In each case, these licenses are consumable and application access is denied once they are gone.

Moab supports both node-locked and floating license models and even allows mixing the two models simultaneously. Moab monitors license usage and only launches an application when required software license availability is guaranteed. In addition, Moab also reserves licenses in conjunction with future jobs to ensure these jobs can run at the appropriate time.

 By default, Moab supports up to 128 independent license types.

**i** Moab license recognition is case insensitive. This means that two licenses with the same spelling and different capitalization are still recognized as the same license. When this occurs, Moab considers the license invalid.

## 11.7.2 Controlling and Monitoring License Availability

Moab can use one of three methods to determine license availability. These methods include locally specifying [consumable generic resources](#), obtaining consumable generic resource information from the [resource manager](#), and interfacing directly with a [license manager](#).

### 11.7.2.A Local Consumable Resources

Both node-locked and floating licenses can be locally specified within Moab using the [NODECFG](#) parameter. In all cases, this is accomplished by associating the license with a node using the [GRES](#) (or generic resource) attribute. If floating, the total cluster-wide license count should be associated with the GLOBAL node. If node-locked, the per node license count should be associated with each compute host (or globally using the [DEFAULT](#) node). For example, if a site has two node-locked licenses for application `EvalA` and six floating licenses for application `EvalB`, the following configuration could be used:

```
NODECFG[node001]  GRES=EvalA:2
NODECFG[node002]  GRES=EvalA:2
NODECFG[GLOBAL]  GRES=EvalB:6
...
```

If you are using an accounting manager and want to distinguish certain generic resources as 'Licenses', you can use the [GRES CFG\[\] LICENSE=TRUE](#) parameter in the Moab server configuration. See [GRES CFG\[<GRES>\]](#) for more information.

### 11.7.2.B Resource Manager Based Consumable Resources

Some resource managers support the ability to define and track generic resource usage at a per node level. In such cases, support for node-locked licenses can be enabled by specifying this information within the resource manager. Moab automatically detects and schedules these resources. For example, in the case of Torque, this can be accomplished by adding generic resource specification lines to the MOM configuration file.

### 11.7.2.C Interfacing to an External License Manager

Moab can also obtain live software license information from a running license manager. Direct interfaces to supported license managers such as FlexNet (formerly FLEXlm) can be

created using the [Native Resource Manager](#) feature. A complete example on interfacing to an external license manager is provided in the [FlexNet](#) section of the native resource manager overview.

## Interfacing to Multiple License Managers

Moab can interface to multiple external license managers simultaneously simply by defining additional native resource manager interfaces. See the [Native Resource Manager Overview](#) for more information.

### 11.7.3 Requesting Licenses within Jobs

Requesting use of software licenses within jobs is typically done in one of two ways. In most cases, the native resource manager job submission language provides a direct method of license specification; for example, in the case of Torque, the software argument could be specified using the format `<SOFTWARE_NAME> [+<LICENSE_COUNT>]` as in the following example:

```
> qsub -l nodes=2,software=blast cmdscript.txt
```

**i** Known issues have been reported using 'software'. The '-l software' syntax is scheduled to be deprecated. Adaptive Computing recommends using 'gres' instead. For example:

```
> qsub -l nodes=2,gres=blast cmdscript.txt
```

**i** The license count is a job total, not a per task total, and the license count value defaults to 1.

An alternative to direct specification is the use of the Moab [resource manager extensions](#). With these extensions, licenses can be requested as generic resources, using the [GRES](#) attribute. The job in the preceding example could also be requested using the following syntax:

```
> qsub -l nodes=2 -W x=GRES:blast cmdscript.txt
```

In each case, Moab automatically determines if the software licenses are node-locked or floating and applies resource requirements accordingly.

If a job requires multiple software licenses, whether of the same or different types, a user would use the following syntax:

```
> qsub -l nodes=2 -W x=GRES:blast+2 cmdscript.txt # two 'blast' licenses required
> qsub -l nodes=2 -W x=GRES:blast+2%bkeep+3 cmdscript.txt # two 'blast' and three
```

```
'bkeep' licenses are required
```

## Related Topics

- [Native Resource Manager License Configuration](#)
- License Ownership with [Advance Reservations](#)
- Multi-Cluster License Sharing with [Moab Workload Manager for Grids](#) Interfaces


# 11.8 Resource Provisioning

## 11.8.1 Resource Provisioning Overview

When processing a resource request, Moab attempts to match the request to an existing available resource. However, if the scheduler determines that the resource is not available or will not be available due to load or policy for an appreciable amount of time, it can select a resource to modify to meet the needs of the current requests. This process of modifying resources to meet existing needs is called provisioning.

Moab evaluates the costs of making the provisioning change in terms of time and other resources consumed before making the decision. Only if the benefits outweigh the costs will the scheduler initiate the change required to support the current workload.

When Moab provisions an environment, it provisions (on a per node basis) the OS and its associated libraries, applications, etc. as a single environment.

 Preemption (requeuing) does not work with dynamic provisioning.

## 11.8.2 Configuring Provisioning

Enabling provisioning consists of configuring an interface to a provisioning manager, specifying which nodes can take advantage of this service, and what the estimated cost and duration of each change will be. This interface can be used to contact provisioning software such as [xCat](#) or HP's Server Automation tool. Additionally, locally developed systems can be interfaced via a script or web service.

---

## Related Topics

- [Native Resource Manager Overview](#)
- [Appendix O: Resource Manager Integration](#)

# 11.9 Managing Networks

## In this topic:

- [11.9.1 Network Management Overview - page 707](#)
- [11.9.2 Dynamic VLAN Creation - page 707](#)
- [11.9.3 Network Load and Health Monitoring - page 708](#)
- [11.9.4 Creating a Resource Management Interface for a New Network - page 709](#)
- [11.9.5 Per-Job Network Monitoring - page 710](#)

## 11.9.1 Network Management Overview

Network resources can be tightly integrated with the rest of a compute cluster using the Moab multi-resource manager management interface. This interface has the following capabilities:

- Dynamic per job and per partition **VLAN** creation and management
- Monitoring and reporting of network health and failure events
- Monitoring and reporting of network load
- Creation of subnets with guaranteed performance criteria
- Automated workload-aware configuration and router maintenance
- Intelligent network-aware scheduling algorithms

## 11.9.2 Dynamic VLAN Creation

Most sites using dynamic VLANs operate under the following assumptions:

- Each compute node has access to two or more networks, one of which is the compute network, and another which is the administrator network.
- Each compute node can only access other compute nodes via the compute network.

- Each compute node can only communicate with the head node via the administrator network.
- Logins on the head node cannot be requested from a compute node.

In this environment, organizations may choose to have VLANs automatically configured that encapsulate individual jobs. These VLANs essentially disconnect the job from either incoming or outgoing communication with other compute nodes.

### 11.9.2.A Configuring VLANs

Automated VLAN management can be enabled by setting up a network resource manager that supports dynamic VLAN configuration and a QoS to request this feature. The example configuration highlights this setup:

```
...
RMCFG[cisco] TYPE=NATIVE RESOURCETYPE=NETWORK FLAGS=VLAN
RMCFG[cisco] CLUSTERQUERYURL=exec://$TOOLS DIR/node.query.cisco.pl
RMCFG[cisco] SYSTEMMODIFYURL=exec://$TOOLS DIR/system.modify.cisco.pl
QOSCFG[netsecure] SECURITY=VLAN
```

### 11.9.2.B Requesting a VLAN

VLANs can be requested on a per job basis directly using the associated resource manager extension or indirectly by requesting a QoS with a VLAN security requirement.

```
> qsub -l nodes=256,walltime=24:00:00,qos=netsecure biojob.cmd
143325.umc.com submitted
```

## 11.9.3 Network Load and Health Monitoring

Network-level load and health monitoring is enabled by supporting the cluster query action in the network resource manager and specifying the appropriate `CLUSTERQUERYURL` attribute in the associated resource manager interface. Node (virtual node) query commands ([mnodectl](#), [checknode](#)) can be used to view this load and health information that will also be correlated with associated workload and written to persistent accounting records. Network load and health based event information can also be fed into [generic events](#) and used to drive appropriate event based [triggers](#).

At present, load and health attributes such as fan speed, temperature, port failures, and various core switch failures can be monitored and reported. Additional failure events are monitored and reported as support is added within the network management system.



## 11.9.4 Creating a Resource Management Interface for a New Network

Many popular networks are supported using interfaces provided in the Moab `tools` directory. If a required network interface is not available, a new one can be created using the following guidelines.

### 11.9.4.A General Requirements

In all cases, a network resource manager should respond to a cluster query request by reporting a single node with a node name that will not conflict with any existing compute nodes. This node should report as a minimum the `state` attribute.

### 11.9.4.B Monitoring Load

Network load is reported to Moab using the generic resource bandwidth. For greatest value, both configured and used bandwidth (in megabytes per second) should be reported as in the following example:

```
force10 state=idle ares=bandwidth:5466 cres=bandwidth:10000
```

### 11.9.4.C Monitoring Failures

Network warning and failure events can be reported to Moab using the `gevent` metric. If automated responses are enabled, embedded epochtime information should be included.

```
force10 state=idle gevent[checksum]='ECC failure detected on port 13'
```

### 11.9.4.D Controlling Router State

Router power state can be controlled as a system modify interface is created that supports the commands `on`, `off`, and `reset`.

### 11.9.4.E Creating VLANs

VLAN creation, management, and reporting is more advanced requiring persistent VLAN ID tracking, global pool creation, and other features. Use of existing routing interface tools as templates is highly advised. VLAN management requires use of both the cluster query interface and the system modify interface.

## 11.9.5 Per-Job Network Monitoring

It is possible to gather network usage on a per job basis using the [Native](#) Interface. When the native interface has been configured to report [netin](#) and [netout](#) Moab automatically gathers this data through the life of a job and reports total usage statistics upon job completion.

```
...
node99 netin=78658 netout=1256
...
```

This information is visible to users and administrators via command-line utilities, the web portal, and the desktop graphical interfaces.

### Related Topics

- [Native Resource Manager Overview](#)
- [Network Utilization Statistics](#)

## 11.10 Intelligent Platform Management Interface

### In this topic:

- [11.10.1 IPMI Overview - page 710](#)
- [11.10.2 Node IPMI Configuration - page 711](#)
- [11.10.3 Installing IPMItool - page 711](#)
- [11.10.4 \[Optional\] Creating the IPMI BMC-Node Map File - page 711](#)
- [11.10.5 Configuring the Moab IPMI Tools - page 712](#)
- [11.10.6 Configuring Moab - page 712](#)
- [11.10.7 Ensuring Proper Setup - page 713](#)

### 11.10.1 IPMI Overview

The Intelligent Platform Management Interface (IPMI) specification defines a set of common interfaces system administrators can use to monitor system health and manage the system. The IPMI interface can monitor temperature and other sensor information, query platform status and power-on/power-off compute nodes. As IPMI operates independently of the node's OS interaction with the node can happen even when powered

down. Moab can use IPMI to monitor temperature information, check power status, power-up, power-down, and reboot compute nodes.

## 11.10.2 Node IPMI Configuration

IPMI must be enabled on each node in the compute cluster. This is usually done either through the node's BIOS or by using a boot CD containing IPMI utilities provided by the manufacturer. With regard to configuring IPMI on the nodes, be sure to enable IPMI-over-LAN and set a common login and password on all the nodes. Additionally, you must set a unique IP address for each node's BMC. Take note of these addresses as you will need them when reviewing the [Creating the IPMI BMC-Node Map File](#) section.

## 11.10.3 Installing IPMItool

[IPMItool](#) is an open-source tool used to retrieve sensor information from the IPMI Baseboard Management Controller (BMC) or to send remote chassis power control commands. The IPMItool developer provides Fedora Core binary packages, as well as a source tarball on the [IPMItool download page](#). Download and install IPMItool on the Moab head node and make sure the `ipmitool` binary is in the current shell PATH.

Proper IPMI setup and IPMItool configuration can be confirmed by issuing the following command on the Moab head node:

```
> ipmitool -I lan -U username -P password -H BMC IP chassis status
```

The output of this command should be similar to the following:

```
System Power          : off
Power Overload        : false
Power Interlock       : inactive
Main Power Fault      : false
Power Control Fault   : false
Power Restore Policy  : previous
Last Power Event      :
Chassis Intrusion     : inactive
Front-Panel Lockout   : inactive
Drive Fault           : false
Cooling/Fan Fault     : false
```

## 11.10.4 [Optional] Creating the IPMI BMC-Node Map File

Since the BMC can be controlled via LAN, it is possible for the BMC to have its own unique IP address. Since this IP address is separate from the IP address of the node, a simple mapping file is required for Moab to know each node's BMC address. The file is a flat text file and should be stored in the Moab home directory. If a mapping file is needed, specify

the name in the `config.ipmi.pl` configuration file in the `etc/` directory. The following is an example of the mapping file:

```
#<NodeID> <BMC IP>
node01 10.10.10.101
node02 10.10.10.102
node03 10.10.10.103
node04 10.10.10.104
node05 10.10.10.105
# NodeID = the name of the nodes returned with "mdiag -n"
# BMC IP = the IP address of the IPMI BMC network interface
```

Note that only the nodes specified in this file are queried for IPMI information. Also note that the mapping file is disabled by default and the nodes that are returned from Moab with `mdiag -n` are the ones that are queried for IPMI sensor data.

### 11.10.5 Configuring the Moab IPMI Tools

The `tools/` subdirectory in the install directory already contains the Perl scripts needed to interface with IPMI. The following is a list of the Perl scripts that should be in the `tools/` directory; confirm these are present and executable:

```
ipmi.mon.pl      # The daemon front-end called by Moab
ipmi.power.pl   # The power control script called by Moab
__mon.ipmi.pl   # The IPMI monitor daemon that updates and caches IPMI data from nodes
```

Next, a few configuration settings need to be adjusted in the `config.ipmi.pl` file found in the `etc` subdirectory. The IPMI-over-LAN username and password need to be set to the values that were set in the [Node IPMI Configuration](#) section. Also, the IPMI query daemon's polling interval can be modified by adjusting `$pollInterval`. This specifies how often the IPMI-enabled nodes are queried to retrieve sensor data.

### 11.10.6 Configuring Moab

To allow Moab to use the IPMI tools, a native resource manager is configured. To do this, the following lines must be added to `moab.cfg`:

```
...
# IPMI - Node monitor script
RMCFG[ipminative] TYPE=NATIVE CLUSTERQUERYURL=exec://$TOOLSDIR/ipmi.mon.pl
...
```

Next, the following lines can be added to allow Moab to issue IPMI power commands:

```
...
# IPMI - Power on/off/reboot script
RMCFG[ipminative] NODEPOWERURL=exec://$TOOLSDIR/ipmi.power.pl
...
```

Moab can be configured to perform actions based on sensor data. For example, Moab can shut down a compute node if its CPU temperature exceeds 100 degrees Celsius, or it can power down idle compute nodes if workload is low. Generic event thresholds are used to tell Moab to perform certain duties given certain conditions. The following example is of a way for Moab to recognize it should power off a compute node if its CPU0 temperature exceeds 100 degrees Celsius:

```
...
# IPMI - Power off compute node if its CPU0 temperature exceeds 100 degrees Celsius.
GEVENTCFG[CPU0_TEMP>100] action=off
...
```

### 11.10.7 Ensuring Proper Setup

Once the preceding steps have been taken, Moab should be started as normal. The IPMI monitoring daemon should start automatically, which can be confirmed with the following:

```
moab@headnode:~/ $ ps aux | grep _mon
moab 11444 0.0 0.3 6204 3172 pts/3 S 10:54 0:00 /usr/bin/perl -w
/opt/moab/tools/_mon.ipmi.pl --start
```

After a few minutes, IPMI data should be retrieved and cached. This can be confirmed with the following command:

```
moab@headnode:~/ $ cat spool/ipmicache.gm
node01 GMETRIC[CPU0_TEMP]=49
node01 GMETRIC[CPU1_TEMP]=32
node01 GMETRIC[SYS_TEMP]=31
node01 POWER=ON
```

Finally, issue the following to ensure Moab is grabbing the IPMI data. Temperature data should be present in the `Generic Metrics` row.

```
moab@headnode:~/ $ checknode node01
node node01
State: Idle (in current state for 00:03:12)
Configured Resources: PROCS: 1 MEM: 2000M SWAP: 3952M DISK: 1M
Utilized Resources: ---
Dedicated Resources: ---
Generic Metrics: CPU0_TEMP=42.00,CPU1_TEMP=30.00,SYS_TEMP=29.00
...
```

## 11.11 Resource Manager Translation

### 11.11.1 Translation Overview

Resource manager translation allows end-users to continue to use existing job command scripts and familiar job management and resource query commands. This is accomplished by emulating external commands, routing the underlying queries to Moab, and then formatting the responses in a familiar manner. Using translation, job submission clients, job query clients, job control clients, and resource query clients can be emulated making switching from one resource manager to another transparent and preserving investment in legacy scripts, tools, and experience.

### 11.11.2 Translation Enablement Steps

To enable translation, you must:

- Edit the Moab tools configuration file.
- Copy, rename, and link the emulation scripts to a shorter, easier-to-use name.

#### 11.11.2.A Configure Translation Tools

Located in the `$MOABHOMEDIR/etc` directory are tools-specific configuration files. For each resource manager that has installed translation tools, edit the Moab tools configuration file in the `etc` directory. For example, do the following:

```
> vi $MOABHOMEDIR/etc/config.moab.pl
# Set the PATH to include directories for moab client commands - mjobctl, etc.
$ENV{PATH} = "/opt/moab/bin:$ENV{PATH}";
```

#### 11.11.2.B Add Translation Tools

In a directory accessible to users, create links to (or copy) the emulation scripts you want your users to use. For example, the emulation script `tools/bjobs.lsf.pl` could be copied to `bin/bjobs`, or, a symbolic link could be created in `bin/bjobs` that points to `tools/bjobs.lsf.pl`.

```
> ln -s tools/bjobs.lsf.pl bin/bjobs
> ln -s tools/bhosts.lsf.pl bin/bhosts
```

## Chapter 12: Troubleshooting and System Maintenance

In this chapter:

12.1 Internal Diagnostics/Diagnosing System Behavior and Problems .....	716
12.1.1 The mdiag Command .....	716
12.1.2 Other Diagnostic Commands .....	718
12.1.3 Using Moab Logs for Troubleshooting .....	718
12.1.4 Automating Recovery Actions after a Failure .....	718
12.2 Logging Overview .....	719
12.2.1 Log Facility Configuration .....	720
12.2.2 Standard Log Format .....	721
12.2.3 Searching Moab Logs .....	722
12.2.4 Event Logs .....	722
12.2.5 Enabling Syslog .....	725
12.2.6 Managing Verbosity .....	725
12.3 Object Messages .....	726
12.3.1 Object Message Overview .....	726
12.3.2 Viewing Messages .....	726
12.3.3 Creating Messages .....	726
12.4 Notifying Administrators of Failures .....	727
12.4.1 Enabling Administrator Email .....	727
12.4.2 Handling Events with the Notification Routine .....	727
12.5 Issues with Client Commands .....	729
12.5.1 Client Overview .....	729
12.5.2 Diagnosing Client Problems .....	729
12.6 Tracking System Failures .....	730
12.6.1 System Failures .....	730
12.6.2 Internal Errors .....	732
12.6.3 Reporting Failures .....	732
12.7 Problems with Individual Jobs .....	733
12.8 Diagnostic Scripts .....	733
12.8.1 support-diag.py .....	734
12.8.2 support.diag.pl .....	735

## 12.1 Internal Diagnostics/Diagnosing System Behavior and Problems

Moab provides a number of commands for diagnosing system behavior. These diagnostic commands present detailed state information about various aspects of the scheduling problem, summarize performance, and evaluate current operation reporting on any unexpected or potentially erroneous conditions found. Where possible, Moab's diagnostic commands even correct detected problems if desired.

At a high level, the diagnostic commands are organized along functionality and object based delineations. Diagnostic commands exist to help prioritize workload, evaluate fairness, and determine effectiveness of scheduling optimizations. Commands are also available to evaluate reservations reporting state information, potential reservation conflicts, and possible corruption issues. Scheduling is a complicated task. Failures and unexpected conditions can occur as a result of resource failures, job failures, or conflicting policies.

Moab's diagnostics can intelligently organize information to help isolate these failures and allow them to be resolved quickly. Another powerful use of the diagnostic commands is to address the situation where there are no hard failures. In these cases, the jobs, compute nodes, and scheduler are all functioning properly, but the cluster is not behaving exactly as desired. Moab diagnostics can help a site determine how the current configuration is performing and how it can be changed to obtain the desired behavior.

In this topic:

[12.1.1 The `mdiag` Command - page 716](#)

[12.1.2 Other Diagnostic Commands - page 718](#)

[12.1.3 Using Moab Logs for Troubleshooting - page 718](#)

[12.1.4 Automating Recovery Actions after a Failure - page 718](#)

### 12.1.1 The `mdiag` Command

The cornerstone of Moab's diagnostics is the `mdiag` command. This command provides detailed information about scheduler state and also performs a large number of internal sanity checks presenting problems it finds as warning messages.

Currently, the `mdiag` command provides in-depth analysis of the following objects and subsystems:



Object/Subsystem	mdiag Flag	Use
<b>Account</b>	-a	Shows detailed account configuration information.
<b>Blocked</b>	-b	Indicates why blocked (ineligible) jobs are not allowed to run.
<b>Class</b>	-c	Shows detailed class configuration information.
<b>Config</b>	-C	Shows configuration lines from <code>moab.cfg</code> and whether or not they are valid.
<b>FairShare</b>	-f	Shows detailed fairshare configuration information, as well as current fairshare usage.
<b>Group</b>	-g	Shows detailed group information.
<b>Job</b>	-j	Shows detailed job information. Reports corrupt job attributes, unexpected states, and excessive job failures.
<b>Frame/Rack</b>	-m	Shows detailed frame/rack information.
<b>Node</b>	-n	Shows detailed node information. Reports unexpected node states and resource allocation conditions.
<b>Priority</b>	-p	Shows detailed job priority information including priority factor contributions to all idle jobs.
<b>QoS</b>	-q	Shows detailed QoS information.
<b>Reservation</b>	-r	Shows detailed reservation information. Reports reservation corruption and unexpected reservation conditions.
<b>Resource Manager</b>	-R	Shows detailed resource manager information. Reports configured and detected state, configuration, performance, and failures of all configured resource manager interfaces.
<b>Standing Reservations</b>	-s	Shows detailed standing reservation information. Reports reservation corruption and unexpected reservation conditions.
<b>Scheduler</b>	-S	Shows detailed scheduler state information. Indicates if

Object/Subsystem	mdiag Flag	Use
		scheduler is stopped, reports status of grid interface, and identifies and reports high-level scheduler failures.
<b>Partition</b>	-t	Shows detailed partition information.
<b>User</b>	-u	Shows detailed user information.

### 12.1.2 Other Diagnostic Commands

Beyond *mdiag*, the [checkjob](#) and [checknode](#) commands also provide detailed information and sanity checking on individual jobs and nodes respectively. These commands can indicate why a job cannot start, which nodes can be available, and information regarding the recent events impacting current job or nodes state.

### 12.1.3 Using Moab Logs for Troubleshooting

Moab logging is extremely useful in determining the cause of a problem. Where other systems may be criticized for not providing adequate logging to diagnose a problem, Moab may be criticized for the opposite reason. If the logging level is configured too high, huge volumes of log output might be recorded, potentially obscuring the problems in a flood of data. Intelligent searching combined with the use of the [LOGLEVEL](#) and [LOGFACILITY](#) parameters can mine out the needed information. Key information associated with various problems is generally marked with the keywords WARNING, ALERT, or ERROR. See the [Logging Overview](#) for further information.

### 12.1.4 Automating Recovery Actions after a Failure

The [RECOVERYACTION](#) parameter of [SCHEDCFG](#) can be used to control scheduler action in the case of a catastrophic internal failure. Valid actions include die, ignore, restart, and trap.

Recovery Mode	Description
<b>die</b>	Moab will exit and, if core files are externally enabled, create a core file for analysis (this is the default behavior).

Recovery Mode	Description
<b>ignore</b>	Moab will ignore the signal and continue processing. This might cause Moab to continue running with corrupt data, which might be dangerous. Use this setting with caution.
<b>restart</b>	When a SIGSEGV is received, Moab will relaunch using the current checkpoint file, the original launch environment, and the original command line flags. The receipt of the signal will be logged but Moab will continue scheduling. Because the scheduler is restarted with a new memory image, no corrupt scheduler data should exist. One caution with this mode is that it might mask underlying system failures by allowing Moab to overcome them. If used, the event log should be checked occasionally to determine if failures are being detected.
<b>trap</b>	When a SIGSEGV is received, Moab stays alive but enters diagnostic mode. In this mode, Moab stops scheduling but responds to client requests allowing analysis of the failure to occur using internal diagnostics available via the <a href="#">mddiag</a> command.

---

### Related Topics

- [Troubleshooting Individual Jobs](#)

## 12.2 Logging Overview

The Moab Workload Manager provides the ability to produce detailed logging of all of its activities. This is accomplished using verbose server logging, event logging, and system logging facilities.

### In this topic:

- [12.2.1 Log Facility Configuration - page 720](#)
- [12.2.2 Standard Log Format - page 721](#)
- [12.2.3 Searching Moab Logs - page 722](#)
- [12.2.4 Event Logs - page 722](#)
- [12.2.5 Enabling Syslog - page 725](#)
- [12.2.6 Managing Verbosity - page 725](#)

## 12.2.1 Log Facility Configuration

The **LOGFILE** and/or **LOGDIR** parameters within the `moab.cfg` file specify the destination of this logging information. Logging information will be written in the file `<MOABHOMEDIR>/<LOGDIR><LOGFILE>` unless `<LOGDIR>` or `<LOGFILE>` is specified using an absolute path. If the log file is not specified or points to an invalid file, all logging information is directed to `STDERR`. However, because of the sheer volume of information that can be logged, it is not recommended that this be done while in production. By default, `LOGDIR` and `LOGFILE` are set to `log` and `moab.log` respectively, resulting in scheduler logs being written to `<MOABHOMEDIR>/log/moab.log`.

The parameter **LOGFILEMAXSIZE** determines how large the log file is allowed to become before it is rolled and is set to 10 MB by default. When the log file reaches this specified size, the log file is rolled. The parameter **LOGFILEROLLDEPTH** controls the number of old logs maintained and defaults to 3. Rolled log files have a numeric suffix appended indicating their order.

The **LOGLEVEL** parameter controls the verbosity of the messages recorded in logs. `LOGLEVEL` can be set to a value between 0 and 9, with 0 being the least verbose and 9 being the most verbose.

If a problem is detected, you might want to increase the `LOGLEVEL` value to get more details. However, doing so will cause the logs to roll faster and will also cause a lot of possibly unrelated information to clutter up the logs. Also be aware of the fact that high `LOGLEVEL` values results in large volumes of possibly unnecessary file I/O to occur on the scheduling machine. Consequently, it is not recommended that high `LOGLEVEL` values be used unless tracking a problem or similar circumstances warrant the I/O cost.

**i** If high log levels are desired for an extended period of time and your Moab home directory is located on a network file system, performance may be improved by moving your log directory to a local file system using the `LOGDIR` parameter.

A final log related parameter is **LOGFACILITY**. This parameter can be used to focus logging on a subset of scheduler activities. This parameter is specified as a list of one or more scheduling facilities as listed in the parameters documentation.

### Example 12-1:

```
# moab.cfg
# allow up to 30 100MB logfiles
LOGLEVEL          3
LOGDIR            /var/tmp/moab
LOGFILEMAXSIZE   100000000
LOGFILEROLLDEPTH 30
```

## 12.2.2 Standard Log Format

Each log event line follows a standard, tab-delimited log format:

```
timestamp <tab> thread ID <tab> visibility <tab> origin <tab> event
code <tab> scope IDs <tab> message
```

Field	Description
<b>timestamp</b>	Timestamps are given in local time, in <a href="#">ISO 8601 format</a> , with a 4-digit timezone offset suffix. For example, 2022-01-27T15:18:30.000-0700.
<b>thread ID</b>	The ID of the thread that is producing the log output.
<b>visibility</b>	Visibility is either a severity (FATAL, ERROR, WARNING, INFO) or a trace level (TRACE1, TRACE2, TRACE3).
<b>origin</b>	Origin is where the log event came from.
<b>event code</b>	The event code provides a way to determine what kind of event happened. For a full list of event codes, see <a href="#">Moab Event Dictionary</a> . When there is no matching event, this field will instead show the hex value of the log level required to display that log statement.
<b>scope IDs</b>	The scope ID associates the event with a specific job or service.
<b>message</b>	Messages can give details about the event and possibly some action information to resolve issues.

*Example 12-2:*

```
2022-08-15T05:26:18.108-0600 846 TRACE1 MQueue.c:MQueueCheckStatus:3081 0
MQueueCheckStatus()
2022-08-15T05:26:18.108-0600 846 TRACE1 MNode.c:MNodeCheckStatus:949 0
MNodeCheckStatus()
2022-08-15T05:26:18.108-0600 846 TRACE1 MVC.c:MVCHarvestVCs:2911 0
Checking for VCs to harvest
2022-08-15T05:26:18.108-0600 846 TRACE1 MSU.c:MUClearChild:5301 0
MUClearChild(PID)
2022-08-15T05:26:18.108-0600 846 INFO MSysMainLoop.c:MSysMainLoop:1059
0x1002a14 Scheduling complete. Sleeping for 60 seconds.
2022-08-15T05:26:18.108-0600 846 TRACE1 MSchedStats.c:MSchedUpdateStats:36 0
MSchedUpdateStats()
2022-08-15T05:26:18.108-0600 846 INFO MSchedStats.c:MSchedUpdateStats:45
0x100a9da Iteration: 23; scheduling time: 0.00 seconds.
2022-08-15T05:26:18.108-0600 846 TRACE1 MRsv.c:MRsvUpdateStats:605 0
MRsvUpdateStats()
2022-08-15T05:26:18.108-0600 846 TRACE1 MSchedStats.c:MSchedUpdateStats:164 0
current util[23]: 0/1d (0.002f%) PH: 0.072f% active jobs: 0 of 0 (completed: 6217)
2022-08-15T05:26:18.109-0600 846 INFO MSysMainLoop.c:MSysMainLoop:1138
```

```
0x1000193 scheduler:Moab A scheduler iteration is ending.
```

## 12.2.3 Searching Moab Logs

While major failures are reported via the `mdiag -S` command, these failures can also be uncovered by searching the logs using the `grep` command as in the following:

```
> grep -E "WARNING|ALERT|ERROR" moab.log
```

On a production system working normally, this list usually includes some ALERT and WARNING messages. The messages are usually self-explanatory, but if not, viewing the log can give context to the message.

If a problem is occurring early when starting the Moab scheduler (before the configuration file is read) Moab can be started up using the `-L <LOGLEVEL>` flag. If this is the first flag on the command line, then the LOGLEVEL is set to the specified level immediately before any setup processing is done and additional logging is recorded.

If problems are detected in the use of one of the client commands, the client command can be re-issued with the `--loglevel=<LOGLEVEL>` command line argument specified. This argument causes log information to be written to STDERR as the client command is running. As with the server, <LOGLEVEL> values from 0 to 9 are supported.

The LOGLEVEL can be changed dynamically by use of the `mschedctl -m` command, or by modifying the `moab.cfg` file and restarting the scheduler. Also, if the scheduler appears to be hung or is not properly responding, the log level can be incremented by one by sending a SIGUSR1 signal to the scheduler process. Repeated SIGUSR1 signals continue to increase the log level. The SIGUSR2 signal can be used to decrease the log level by one.

If an unexpected problem does occur, save the log file as it is often very helpful in isolating and correcting the problem.

## 12.2.4 Event Logs

Major events are reported to both the Moab log file and the Moab event log. By default, the event log is maintained in the statistics directory and rolls on a daily basis, using the naming convention `events.WWW_MMM_DD_YYYY` as in `events.Tue_Oct_18_2022`.

### 12.2.4.A Event Log Format

The event log contains information about major job, reservation, node, and scheduler events and failures and reports this information in the following format:

```
<EVENTTIME> <EPOCHTIME>:<EID> <OBJECT> <OBJECTID> <EVENT> <DETAILS>
```

**Example 12-3:**

```


VERSION 500
07:03:21 110244322:0 sched clusterA start
07:03:26 110244327:1 rsv system.1 start 1124142432 1324142432 2 2 0.0 2342155.3
node1|node2 NA RSV=%=system.1=
07:03:54 110244355:2 job 1413 end 8 16 llw mcc 432000 Completed [batch:1]
11 08708752 1108703981 ...
07:04:59 110244410:3 rm base failure cannot connect to RM
07:05:20 110244431:4 sched clusterA stop admin
...

```

The parameter [RECORDEVENTLIST](#) can be used to control which events are reported to the event log. See the sections on [job](#) and [reservation](#) trace format for more information regarding the values reported in the details section for those records.

## Record Type Specific Details Format

The format for each record type is unique and is described in the following table:

Record Type	Event Types	Description
<b>gevent</b>	See <a href="#">Enabling Generic Events</a> for gevent information.	 Generic events are included within node records. See node detail format that follows.
<b>job</b>	JOBCANCEL, JOBCHECKPOINT, JOBEND, JOBHOLD, JOBMIGRATE, JOBMODIFY, JOBPREEMPT, JOBREJECT, JOBRESUME, JOBSTART, JOBSUBMIT AM*	See <a href="#">Workload Accounting Records</a> . See <a href="#">Accounting Events</a> for information about the AM* event types.
<b>node</b>	NODEDOWN, NODEFAILURE, NODEUP	The following fields are displayed in the event file in a space-delimited line as long as Moab has information pertaining to it: state, partition, disk, memory, maxprocs, swap, os, rm, nodeaccesspolicy, class, and message, where state is the node's current state and message is a human readable message indicating reason for node state change.
<b>rm</b>	RMDOWN, RMPOLLEND, RMPOLLSTART, RMUP	Human readable message indicating reason for resource manager state change.

Record Type	Event Types	Description
		<p><b>i</b> For <code>SCHEDCOMMAND</code>, only create/modify commands are recorded. No record is created for general list/query commands. <code>ALLSCHEDCOMMAND</code> does the same thing as <code>SCHEDCOMMAND</code>, but it also logs info query commands.</p>
<b>trigger</b>	TRIGEND, TRIGFAILURE, TRIGSTART	<p><code>&lt;ATTR&gt;="&lt;VALUE&gt;" [ &lt;ATTR&gt;="&lt;VALUE&gt;" ] . . .</code>            where <code>&lt;ATTR&gt;</code> is one of the following: <code>actiondata</code>, <code>actiontype</code>, <code>description</code>, <code>ebuf</code>, <code>eventtime</code>, <code>eventtype</code>, <code>flags</code>, <code>name</code>, <code>objectid</code>, <code>objecttype</code>, <code>obuf</code>, <code>offset</code>, <code>period</code>, <code>requires</code>, <code>sets</code>, <code>threshold</code>, <code>timeout</code>, and so forth.            See <a href="#">About Object Triggers</a> for more information.</p>

### 12.2.4.B Exporting Events in Real-Time

Moab event information can be exported to external systems in real-time using the [ACCOUNTINGINTERFACEURL](#) parameter. When set, Moab activates this URL each time one of the default events or one of the events specified by the [RECORDEVENTLIST](#) occurs.

While various protocols can be used, the most common protocol is `exec`, which indicates that Moab should launch the specified tool or script and pass in event information as command line arguments. This tool can then select those events and fields of interest and re-direct them as appropriate providing significant flexibility and control to the organization.

#### Exec Protocol Format

When a URL with an `exec` protocol is specified, the target is launched with the event fields passed in as STDIN. These fields appear exactly as they do in the [event logs](#) with the same values and order.

**i** The `tools/sql` directory included with the Moab distribution contains `event.create.sql.pl`, a sample accounting interface processing script that can be used as a template.



## 12.2.5 Enabling Syslog

In addition to the log file, the Moab scheduler can report events it determines to be critical to the UNIX syslog facility via the daemon facility using priorities ranging from `INFO` to `ERROR`. See `USESYSLOG`. The verbosity of this logging is not affected by the `LOGLEVEL` parameter. In addition to errors and critical events, user commands that affect the state of the jobs, nodes, or the scheduler may also be logged to syslog. Moab syslog messages are reported using the `INFO`, `NOTICE`, and `ERR` syslog priorities.

By default, messages are logged to syslog's user facility. However, using the `USESYSLOG` parameter, Moab can be configured to use any of the following:

- `user`
- `daemon`
- `local0`
- `local1`
- `local2`
- `local3`
- `local4`
- `local5`
- `local6`
- `local7`

## 12.2.6 Managing Verbosity

In very large systems, a highly verbose log may roll too quickly to be of use in tracking specific targeted behaviors. In these cases, one or more of the following approaches may be of use:

- Use the `LOGFACILITY` parameter to log only functions and services of interest.
- Use `syslog` to maintain a permanent record of critical events and failures.
- Specify higher object loglevels on jobs, nodes, and reservations of interest (such as `NODECFG[orion13] LOGLEVEL=6`).
- Increase the range of events reported to the event log using the `RECORDEVENTLIST` parameter.
- Review object messages for required details.
- Run Moab in `monitor` mode using `IGNOREUSERS`, `IGNOREJOBS`, `IGNORECLASSES`, or `IGNORENODES`.

## Related Topics

- [RECORDEVENTLIST](#) parameter
- [USESYSLOG](#) parameter
- [Notifying Administrators of Failures](#)
- [Workload Event Format](#)
- [mschedctl -L](#) command

## 12.3 Object Messages

### 12.3.1 Object Message Overview

Messages can be associated with the scheduler, jobs, and nodes. Their primary use is a line of communication between resource managers, the scheduler, and end-users. When a node goes offline, or when a job fails to run, both the resource manager and the scheduler will post messages to the object's message buffer, giving the administrators and end-users a reason for the failure. They can also be used as a way for different administrators and users to send messages associated with the various objects. For example, an administrator can set the message `Node going down for maintenance Apr/6/08 12pm, " on node node01`, which would then be visible to other administrators.

### 12.3.2 Viewing Messages

To view messages associated with a job (either from users, the resource manager, or Moab), run the [checkjob](#) command.

To view messages associated with a node (either from users, the resource manager, or Moab), run the [checknode](#) command.

To view system messages, use the [mschedctl -l](#) message command.

To view the messages associated with a credential, run the [mcredctl -c](#) command.

### 12.3.3 Creating Messages

To create a message use the [mschedctl -c](#) message `<STRING> [-o <OBJECTTYPE>:<OBJECTID>] [-w <ATTRIBUTE>=<VALUE> [-w ...]]` command.

The `<OBJECTTYPE>` can be one of the following:

- node
- job
- rsv
- user
- acct
- qos
- class
- group

The <ATTRIBUTE> can be one of the following:

- owner
- priority
- expirertime
- type

Valid types include:

- annotation
- other
- hold
- pendactionerror

## 12.4 Notifying Administrators of Failures

### 12.4.1 Enabling Administrator Email

In the case of certain events, Moab can automatically send email to administrators. To enable mail notification, the [MAILPROGRAM](#) parameter must be set to `DEFAULT` or point to the locally available mail client. With this set, policies such as [JOBREJECTPOLICY](#) will send email to administrators if set to a value of `MAIL`.

### 12.4.2 Handling Events with the Notification Routine

Moab possesses a primitive event management system through the use of the `notify` program. The program is called each time an event of interest occurs. Currently, most

events are associated with failures of some sort but use of this facility need not be limited in this way. The **NOTIFICATIONPROGRAM** parameter allows a site to specify the name of the program to run. This program is most often locally developed and designed to take action based on the event that has occurred. The location of the notification program can be specified as a relative or absolute path. If a relative path is specified, Moab looks for the notification relative to the `$(INSTDIR)/tools` directory. In all cases, Moab verifies the existence of the notification program at start up and disables it if it cannot be found or is not executable.

The notification program's action can include steps such as reporting the event via email, adjusting scheduling parameters, rebooting a node, or even recycling the scheduler.

For most events, the notification program is called with command line arguments in a simple `<EVENTTYPE>: <MESSAGE>` format. The following event types are currently enabled:

Event Type	Format	Description
<b>JOBCORRUPTION</b>	<MESSAGE>	An active job is in an unexpected state or has one or more allocated nodes that are in unexpected states.
<b>JOBHOLD</b>	<MESSAGE>	A job hold has been placed on a job.
<b>JOBWCVIOLATION</b>	<MESSAGE>	A job has exceeded its wallclock limit.
<b>RESERVATIONCORRUPTION</b>	<MESSAGE>	Reservation corruption has been detected.
<b>RESERVATIONCREATED</b>	<RSVNAME> <RSVTYPE> <NAME> <PRESENTTIME> <STARTTIME> <ENDTIME> <NODECOUNT>	A new reservation has been created.
<b>RESERVATIONDESTROYED</b>	<RSVNAME> <RSVTYPE> <PRESENTTIME> <STARTTIME> <ENDTIME> <NODECOUNT>	A reservation has been destroyed.
<b>RMFAILURE</b>	<MESSAGE>	The interface to the resource manager has failed.

Perhaps the most valuable use of the notify program stems from the fact that additional notifications can be easily inserted into Moab to handle site specific issues. To do this, locate the proper block routine, specify the correct conditional statement, and add a call to the routine `notify(<MESSAGE>);`.

---

### Related Topics

- [JOBREJECTPOLICY](#) parameter
- [MAILPROGRAM](#) parameter
- [Event Log Overview](#)

## 12.5 Issues with Client Commands

### 12.5.1 Client Overview

Moab client commands are implemented as links to the executable `mclient`. When a Moab client command runs, the client executable determines the name under which it runs and behaves accordingly. At the time Moab was configured, a home directory was specified. The Moab client attempts to open the configuration file, `moab.cfg`, in the `etc/` folder of this home directory on the node where the client command executes. This means that the home directory specified at configure time must be available on all hosts where the Moab client commands are executed. This also means that a `moab.cfg` file must be available in the `etc/` folder of this home directory. When the clients open this file, they will try to load the `SCHEDCFG` parameter to determine how to contact the Moab server.

**i** The home directory value specified at configure time can be overridden by creating an `/etc/moab.cfg` file or by setting the `MOABHOMEDIR` environment variable.

Once the client has determined where the Moab server is located, it creates a message, adds an encrypted checksum, and sends the message to the server. The Moab client and Moab server must use a shared secret key for this to work. When the Moab server receives the client request and verifies the message, it processes the command and returns a reply.

### 12.5.2 Diagnosing Client Problems

The easiest way to determine where client failures are occurring is to use built-in Moab logging. On the client side, use the `--loglevel` flag. For example:

```
> showq --loglevel=9
```

*This will display verbose logging information regarding the loading of the configuration file, connecting to the Moab server, sending the request, and receiving a response.*

This information almost always reveals the source of the problem. If it does not, the next step is to look at the Moab server side logs; this is done using the following steps:

- Stop Moab scheduling so that the only activity is handling Moab client requests.

```
> mschedctl -s
```

- Set the logging level to *very verbose*.

```
> mschedctl -m loglevel 7
```

- Watch Moab activity.

```
> tail -f log/moab.log | more
```

Now, in a second window, issue any failing client command, such as [showq](#).

The `moab.log` file will record the client request and any reasons it was rejected.

## 12.6 Tracking System Failures

In this topic:

[12.6.1 System Failures - page 730](#)

[12.6.2 Internal Errors - page 732](#)

[12.6.3 Reporting Failures - page 732](#)

### 12.6.1 System Failures

The scheduler has a number of dependencies that may cause failures if not satisfied. These dependencies are in the areas of disk space, network access, memory, and processor utilization.

#### 12.6.1.A Disk Space

The scheduler uses a number of files. If the file system is full or otherwise inaccessible, the following behaviors might be noted:

Unavailable File	Behavior
<b>moab.pid</b>	Scheduler cannot perform single instance check.
<b>moab.ck*</b>	Scheduler cannot store persistent record of reservations, jobs, policies, summary statistics, and so forth.
<b>moab.cfg</b> <b>/moab.dat</b>	Scheduler cannot load local configuration.
<b>log/*</b>	Scheduler cannot log activities.
<b>stats/*</b>	Scheduler cannot write job records.

**i** When possible, configure Moab to use local disk space for configuration files, statistics files, and logs files. If any of these files are located in a networked file system (such as NFS, DFS, or AFS) and the network or file server experience heavy loads or failures, Moab server may appear sluggish or unresponsive and client command might fail. Use of local disk space eliminates susceptibility to this potential issue.

### 12.6.1.B Network

The scheduler uses a number of socket connections to perform basic functions. Network failures may affect the following facilities:

Network Connection	Behavior
<b>scheduler client</b>	Scheduler client commands fail.
<b>resource manager</b>	Scheduler is unable to load/update information regarding nodes and jobs.
<b>accounting manager</b>	Scheduler is unable to validate account access or reserve/debit account balances.

### 12.6.1.C Memory

Depending on cluster size and configuration, the scheduler may require up to 120 MB of memory on the server host. If inadequate memory is available, multiple aspects of scheduling might be negatively affected. The scheduler log files should indicate if memory failures are detected and mark any such messages with the ERROR or ALERT keywords.

### 12.6.1.D Processor Utilization

On a heavily loaded system, the scheduler may appear sluggish and unresponsive. However, no direct failures should result from this slowdown. Indirect failures may include timeouts of peer services (such as the resource manager or accounting manager) or timeouts of client commands. All timeouts should be recorded in the scheduler log files.

## 12.6.2 Internal Errors

The Moab scheduling system contains features to assist in diagnosing internal failures. If the scheduler exits unexpectedly, the scheduler logs may provide information regarding the cause. If no reason can be determined, use of a debugger might be required.

### 12.6.2.A Logs

The first step in any exit failure is to check the last few lines of the scheduler log. In many cases, the scheduler may have exited due to misconfiguration or detected system failures. The last few lines of the log should indicate why the scheduler exited and what changes would be required to correct the situation. If the scheduler did not intentionally exit, increasing the `LOGLEVEL` parameter to 7, or higher, may help isolate the problem.

## 12.6.3 Reporting Failures

If an internal failure is detected on your system, the information of greatest value to developers in isolating the problem will be the output of the `gdb` where subcommand and a printout of all variables associated with the failure. In addition, a level 7 log covering the failure can also help in determining the environment that caused the failure. If you encounter such and require assistance, please submit a ticket at the following address:

<https://support.adaptivecomputing.com/hpc-cloud-support-portal-2/>

**i** If you do not already have a support username and password, create a free account to request a support ticket



## 12.7 Problems with Individual Jobs

To determine why a particular job will not start, there are several helpful commands:

### checkjob -v

*checkjob* evaluates the ability of a job to start immediately. Tests include resource access, node state, job constraints (such as startdate, taskspnode, and QoS). Additionally, command line flags can be specified to provide further information.

Flag	Description
<b>-l</b> <POLICYLEVEL>	Evaluates impact of throttling policies on job feasibility.
<b>-n</b> <NODENAME>	Evaluates resource access on specific node.
<b>-r</b> <RESERVATION_LIST>	Evaluates access to specified reservations.

### checknode

Displays detailed status of node.

### mdiag -b

Displays various reasons job is considered blocked or non-queued.

### mdiag -j

Displays high level summary of job attributes and performs sanity check on job attributes/state.

### showbf -v

Determines general resource availability subject to specified constraints.

## 12.8 Diagnostic Scripts

Moab Workload Manager provides diagnostic scripts that can help aid in monitoring the state of the scheduler, resource managers, and other important components of the cluster software stack. These scripts can also be used to help diagnose issues that might need to be resolved with the help of Adaptive Computing support staff. This section introduces available diagnostic scripts.

In this topic:

[12.8.1 support-diag.py - page 734](#)

[12.8.2 support.diag.pl - page 735](#)

## 12.8.1 support-diag.py

The `support-diag.py` script has a two-fold purpose. First, it can be used by a Moab trigger or cron job to create a regular snapshot of the state of Moab. The script captures the output of several Moab diagnostic commands (such as `showq`, `mdiag -n`, and `mdiag -S`), gathers configuration/log files, and records pertinent operating system information. This data is then compressed in a time-stamped tarball for easy long-term storage.

Second, the script provides Adaptive Computing support personnel with a complete package of information that can be used to help diagnose configuration issues or system bugs. After capturing the state of Moab, the resulting tarball can be sent to your Adaptive Computing support contact for further diagnosis.

The script asks you for the trouble ticket number, `-t <TICKET#>`, or `-n`. If you chose to enter `-t <TICKET#>` the script uploads your support diagnostic output to Adaptive Computing Customer Support. The upload and ticket number request can be prevented using the `-n` option.

### 12.8.1.A Synopsis

```
support-diag.py [<options>]
```

### 12.8.1.B Arguments

Argument	Description
<code>-h, --help</code>	Show this help message and exit.
<code>-q, --diag-torque-off, --without-torque</code>	Disable Torque diagnostics.
<code>-p TMPDIR, --tmp-dir=TMPDIR</code>	Use a different tmp directory to store output.
<code>-n, --no-upload</code>	Do not upload to Adaptive Computing.
<code>-t TICKET#</code>	Support ticket number.

Argument	Description
<code>-f, --full-mode</code>	Gather additional logs, stats and, <code>moab.db</code> files.
<code>-u TIMEOUT, --moab-timeout=TIMEOUT</code>	Define Moab command timeout (default: 300 seconds).
<code>-d, --debug-mode</code>	<code>support-diag</code> print debug variables.
<code>-o, --offline-mode</code>	Gather offline logging only.
<code>-r, --ftp</code>	Use ftp instead of scp.
<code>-V, --version</code>	Print version information.

## 12.8.2 support.diag.pl

**i** This script is deprecated with the 8.0 release. Use the `support-diag.py` script instead.

The `support.diag.pl` script has a two-fold purpose. First, it can be used by a Moab trigger or cron job to create a regular snapshot of the state of Moab. The script captures the output of several Moab diagnostic commands (such as `showq`, `mdiag -n`, and `mdiag -S`), gathers configuration/log files, and records pertinent operating system information. This data is then compressed in a time-stamped tarball for easy long-term storage.

The second purpose of the `support.diag.pl` script is to provide Adaptive Computing support personnel with a complete package of information that can be used to help diagnose configuration issues or system bugs. After capturing the state of Moab, the resulting tarball could be sent to your Adaptive Computing support contact for further diagnosis.

The `support.diag.pl` will ask you for the trouble ticket number then guide you through the process of uploading the data to Adaptive Computing Customer Support. The uploading and ticket number request can be prevented using the `--no-upload` and `--support-ticket=<SUPPORT_TICKET_ID>` flags detailed in the Arguments table that follows.

### 12.8.2.A Synopsis

```
support.diag.pl [--include-log-lines=<NUM>] [--diag-torque]
```

### 12.8.2.B Arguments

Argument	Description
<code>--include-log-lines=&lt;NUM&gt;</code>	Instead of including the entire <code>moab.log</code> file, only the last <code>&lt;NUM&gt;</code> lines are captured in the diagnostics.
<code>--diag-torque</code>	Diagnostic commands pertinent to the Torque resource manager are included.
<code>--no-upload</code>	Prevents the system from asking the user if they want to upload the tarball to Adaptive Computing Customer Support.
<code>--support-ticket=&lt;SUPPORT_TICKET_ID&gt;</code>	Prevents the system from asking the user for a support ticket number.

## Chapter 13: Improving User Effectiveness

In this chapter:

13.1 User Feedback Loops .....	737
13.1.1 Improving Job Size/Duration Requests .....	737
13.1.2 Improving Resource Requirement Specification .....	738
13.2 User Level Statistics .....	738
13.3 Enhancing Wallclock Limit Estimates .....	739
13.4 Job Start Time Estimates .....	739
13.4.1 Example .....	739
13.4.2 Estimation Types .....	739
13.5 Providing Resource Availability Information .....	740
13.6 Collecting Performance Information on Individual Jobs .....	741

### 13.1 User Feedback Loops

Almost invariably, real world systems outperform simulated systems, even when all policies, reservations, workload, and resource distributions are fully captured and emulated. What is it about real world usage that is not emulated via a simulation? The answer is the user feedback loop, the impact of users making decisions to optimize their level of service based on real time information.

A user feedback loop is created any time information is provided to a user that modifies job submission or job management behavior. As in a market economy, the cumulative effect of many users taking steps to improve their individual scheduling performance results in better job packing, lower queue time, and better overall system utilization. Because this behavior is beneficial to the system at large, system administrators and management should encourage this behavior and provide the best possible information to them.

There are two primary types of information that help users make improved decisions: cluster wide resource availability information and per-job resource utilization information.

#### 13.1.1 Improving Job Size/Duration Requests

Moab provides a number of informational commands that help users make improved job management decisions based on real-time cluster wide resource availability information.

These commands include [showbf](#), [showstats -f](#), and [showq](#). Using these commands, a user can determine what resources are available and what job configurations statistically receive the best scheduling performance.

## 13.1.2 Improving Resource Requirement Specification

A job's resource requirement specification tells the scheduler what type of compute nodes are required to run the job. These requirements may state that a certain amount of memory is required per node or that a node has a minimum processor speed. At many sites, users will determine the resource requirements needed to run an initial job. Then, for the next several years, they will use the same basic batch command file to run all of their remaining jobs even though the resource requirements of their subsequent jobs may be very different from their initial run. Users often do not update their batch command files even though these constraints may be unnecessarily limiting the resources available to their jobs for two reasons: (1) users do not know how much their performance will improve if better information were provided and (2) users do not know exactly what resources their jobs are using and are afraid to lower their job's resource requirements since doing so might cause their job to fail.

To help with determining accurate per job resource utilization information, Moab provides the [FEEDBACKPROGRAM](#) facility. This tool allows sites to send detailed resource utilization information back to users via email, to store it in a centralized database for report preparation, or use it in other ways to help users refine their batch jobs.

## 13.2 User Level Statistics

Besides displaying job queues, end-users can display a number of their own statistics. The [showstats -u <USER\\_ID>](#) command displays current and historical statistics for a user as seen in what follows:

```
$ showstats -u john
statistics initialized Wed Dec 31 17:00:00

----- Active -----|----- Completed -----
-----|
user      Jobs Procs ProcHours Jobs    %    PHReq    %    PHDed    %    FSTgt  AvgXF
MaxXF    AvgQH  Effic  WCAcc
john     1      1      30.96   9    0.00  300.0   0.00  148.9   0.00  -----  0.62
0.00    4.33  100.00  48.87
```

Users can query available system resources with the [showbf](#) command. This can aid users in requesting node configurations that are idle. Also, users can use the [checkjob](#) command to determine what parameter(s) are restricting their job from running. Moab performs better with more accurate wallclock estimates.

**i** Moab must use an ODBC-compliant database to report statistics with Viewpoint reports.

## 13.3 Enhancing Wallclock Limit Estimates

As explained in the previous section, `showstats -u <USER_ID>` reports statistics for a given user. The `showstats -u` command can be accessed by all users. They can use fields such as `PHReq`, `PHDed`, or `WCacc` to gauge wallclock estimates. Accurate wallclock estimates allow a job to be scheduled as soon as possible in a slot that it will fit in. Low or high estimates can cause a job to be scheduled in a less favorable position.

## 13.4 Job Start Time Estimates

### 13.4.1 Example

Each user can use the `showstart` command to display estimated start and completion times. The following example illustrates a typical response from issuing this command:

```
> showstart orion.13762
job orion.13762 requires 2 procs for 0:33:20
Estimated Rsv based start in      1:04:55 on Fri Jul 15 12:53:40
Estimated Rsv based completion in  2:44:55 on Fri Jul 15 14:33:40
Estimated Priority based start in   5:14:55 on Fri Jul 15 17:03:40
Estimated Priority based completion in 6:54:55 on Fri Jul 15 18:43:40
Estimated Historical based start in  00:00:00 on Fri Jul 15 11:48:45
Estimated Historical based completion in 1:40:00 on Fri Jul 15 13:28:45
Best Partition: fast
```

### 13.4.2 Estimation Types

#### 13.4.2.A Reservation-Based Estimates

Reservation-based start time estimation incorporates information regarding current administrative, user, and job reservations to determine the earliest time the specified job can allocate the needed resources and start running. In essence, this estimate indicates the earliest time the job will start, assuming this job is the highest priority job in the queue.

**i** For reservation-based estimates, the information provided by this command is more highly accurate if the job is highest priority, if the job has a reservation, or if the majority of the jobs that are of higher priority have reservations. Consequently, site administrators wanting to make decisions based on this information might want to consider using the `RESERVATIONDEPTH` parameter to increase the number of priority-based reservations. This can be set so that most, or even all, idle jobs receive priority reservations and make the results of this command generally useful. The only caution of this approach is that increasing the `RESERVATIONDEPTH` parameter more tightly constrains the decisions of the scheduler and may result in slightly lower system utilization (typically less than 8% reduction).

### 13.4.2.B Backlog/Priority Estimates

Priority-based job start analysis determines when the queried job will fit in the queue and determines the estimated amount of time required to complete the jobs currently running or scheduled to run before this job can start.

In all cases, if the job is running, this command returns the time the job starts. If the job already has a reservation, this command returns the start time of the reservation.

### 13.4.2.C Historical Estimates

Historical analysis uses historical queue times for jobs that match a similar processor count and job duration profile. This information is updated on a sliding window that is configurable within `moab.cfg`.

---

#### Related Topics

- [showstart](#) command

## 13.5 Providing Resource Availability Information

Moab provides commands to allow the user to query available resources. The [showbf command](#) displays what resources are available for immediate use. Using different command line parameters, such as `-m`, `-n`, and `-q` allows the user to query resources based on memory, nodecount, or QoS respectively.



## 13.6 Collecting Performance Information on Individual Jobs

Individual job information can be collected from the statistics file in [STATDIR](#), which contains start time, end time, end state, QoS requested, QoS delivered, and so forth for different jobs. Also, Moab optionally provides similar information to a site's feedback program. See the section [User Feedback Overview](#) for more information about the feedback program.

## Chapter 14: Cluster Analysis and Testing

In this chapter:

14.1 Testing New Releases and Policies .....	742
14.1.1 Moab Evaluation Modes .....	743
14.1.2 Testing New Releases .....	744
14.1.3 Testing New Policies .....	745
14.1.4 Moab Side-by-Side .....	746
14.2 Testing New Middleware .....	746
14.2.1 Analysis Aspects .....	747
14.2.2 General Analysis .....	748
14.2.3 Native Mode Analysis .....	748
14.3 Workload Event Format .....	749
14.3.1 Workload Event Record Format .....	750
14.3.2 Reservation Event Records .....	759
14.3.3 Recording Job Events .....	760

Moab has a number of unique features that allow site administrators to visualize current cluster behavior and performance, safely evaluate changes on production systems, and analyze probable future behaviors within a variety of environments.

These capabilities are enabled through a number of Moab facilities that might not appear to be closely related at first. However, taken together, these facilities allow organizations the ability to analyze their cluster without the losses associated with policy conflicts, unnecessary downtime, and faulty systems middleware.

### 14.1 Testing New Releases and Policies

In this topic:

- 14.1.1 Moab Evaluation Modes - page 743
- 14.1.2 Testing New Releases - page 744
- 14.1.3 Testing New Policies - page 745
- 14.1.4 Moab Side-by-Side - page 746

## 14.1.1 Moab Evaluation Modes

### 14.1.1.A MONITOR Mode

Moab supports a scheduling mode called `MONITOR`. In this mode, the scheduler initializes, contacts the resource manager and other peer services, and conducts scheduling cycles exactly as it would if running in `NORMAL` or production mode. Jobs are prioritized, reservations created, policies and limits enforced, and administrator and end-user commands enabled. The key difference is that although live resource management information is loaded, `MONITOR` mode disables Moab's ability to start, preempt, cancel, or otherwise modify jobs or resources. Moab continues to attempt to schedule exactly as it would in `NORMAL` mode, but its ability to actually impact the system is disabled. Using this mode, a site can quickly verify correct resource manager configuration and scheduler operation. This mode can also be used to validate new policies and constraints. In fact, Moab can be run in `MONITOR` mode on a production system while another scheduler or even another version of Moab is running on the same system. This unique ability can allow new versions and configurations to be fully tested without any exposure to potential failures and with no cluster downtime.

To run Moab in `MONITOR` mode, simply set the `MODE` attribute of the `SCHEDCFG` parameter to `MONITOR` and start Moab. Normal scheduler commands can be used to evaluate configuration and performance. [Diagnostic commands](#) can be used to look for any potential issues. Further, the Moab log file can be used to determine which jobs Moab attempted to start, and which resources Moab attempted to allocate.

If another instance of Moab is running in production and a site administrator wants to evaluate an alternate configuration or new version, this is easily done but care should be taken to avoid conflicts with the primary scheduler. Potential conflicts include statistics files, logs, checkpoint files, and user interface ports. One of the easiest ways to avoid these conflicts is to create a new test directory with its own log and statistics subdirectories. The new `moab.cfg` file can be created from scratch or based on the existing `moab.cfg` file already in use. In either case, make certain that the `PORT` attribute of the `SCHEDCFG` parameter differs from that used by the production scheduler by at least two ports. If testing with the production binary executable, the `MOABHOMEDIR` environment variable should be set to point to the new test directory to prevent Moab from loading the production `moab.cfg` file.

### 14.1.1.B TEST Mode

`TEST` mode behaves much like `MONITOR` mode with the exception that Moab will log the scheduling actions it would have taken to the `stats/<DAY>.events` file. Using this file, sites can determine the actions Moab would have taken if running in `NORMAL` mode and verify all actions are in agreement with expected behavior.

### 14.1.1.C INTERACTIVE Mode

INTERACTIVE mode allows for evaluation of new versions and configurations in a manner different from MONITOR mode. Instead of disabling all resource and job control functions, Moab sends the desired change request to the screen and requests permission to complete it. For example, before starting a job, Moab might post something like the following to the screen:

```
Command:  start job 1139.ncsa.edu on node list test013,test017,test018,test021
Accept:   (y/n) [default: n]?
```

The administrator must specifically accept each command request after verifying it correctly meets desired site policies. Moab will then execute the specified command. This mode is highly useful in validating scheduler behavior and can be used until configuration is appropriately tuned and all parties are comfortable with the scheduler's performance. In most cases, sites will want to set the scheduling mode to NORMAL after verifying correct behavior.

## 14.1.2 Testing New Releases

By default, Moab runs in a **mode** called NORMAL, which indicates that it is responsible for the cluster. It loads workload and resource information, and is responsible for managing that workload according to mission objectives and policies. It starts, cancels, preempts, and modifies jobs according to these policies.

If Moab is configured to use a mode called TEST, it loads all information, performs all analysis, but, instead of actually starting or modifying a job, it merely logs the fact that it would have done so. A test instance of Moab can run at the same time as a production instance of Moab. A test instance of Moab can also run while a production scheduler of another type (such as PBS) is simultaneously running. This multi-scheduler ability allows stability and performance tests to be conducted that can help answer the following questions:

- What impact do Moab services have on network, processor, and memory load?
- What impact do Moab services have on the underlying resource manager?
- Is Moab able to correctly import resource, workload, policy, and credential information from the underlying resource manager?
- Are Moab's logged scheduling decisions in line with mission objectives?

In test mode, all of Moab's commands and services operate normally allowing the use of client commands to perform analysis. In most cases, the **mdiag** command is of greatest value, displaying loaded values, as well as reporting detected failures, inconsistencies, and object corruption. The following table highlights the most common diagnostics performed:

Command	Object
<code>mdiag -n</code>	Compute nodes, storage systems, network systems, and generic resources
<code>mdiag -j</code>	Applications and static jobs
<code>mdiag -u</code> <code>mdiag -g</code> <code>mdiag -a</code>	User, group, and account credentials
<code>mdiag -c</code>	Queues and policies
<code>mdiag -R</code>	Resource manager interface and performance
<code>mdiag -S</code>	Scheduler/system level failures introduced by corrupt information

These commands will not only verify proper scheduling objects but will also analyze the behavior of each resource manager, recording failures, and delivered performance. If any misconfiguration, corruption, interface failure, or internal failure is detected, it can be addressed in the test mode instance of Moab with no urgency or risk to production cluster activities.

## 14.1.3 Testing New Policies

### 14.1.3.A Verifying Correct Specification of New Policies

The first aspect of verifying a new policy is verifying correct syntax and semantics. If manually editing the `moab.cfg` file, the following command can be used for validation:

```
> mdiag -C
```

*This command will validate the configuration file and report any misconfiguration.*

### 14.1.3.B Verifying Correct Behavior of New Policies

If concern exists over the impact of a new policy, an administrator can babysit Moab by putting it into **INTERACTIVE** mode. In this mode, Moab will schedule according to all mission objectives and policies, but before taking any action, it will request that the administrator confirm the action. See the [interactive mode overview](#) for more information.

In this mode, only actions approved by the administrator will be carried out. Once proper behavior is verified, the Moab mode can be set to `NORMAL`.

## 14.1.4 Moab Side-by-Side

Moab provides an additional evaluation method that allows a production cluster or other resource to be logically partitioned along resource and workload boundaries and allows different instances of Moab to schedule different partitions. The parameters [IGNORENODES](#), [IGNORECLASSES](#), [IGNOREJOBS](#), and [IGNOREUSERS](#) are used to specify how the system is to be partitioned. In the following example, a small portion of an existing cluster is partitioned for temporary grid testing so that there is no impact on the production workload:

```
SCHEDCFG[prod] MODE=NORMAL SERVER=orion.cxz.com:42020
RMCFG[Torque]  TYPE=PBS
IGNORENODES    node61,node62,node63,node64
IGNOREUSERS    gridtest1,gridtest2
...
SCHEDCFG[prod] MODE=NORMAL SERVER=orion.cxz.com:42030
RMCFG[Torque]  TYPE=PBS
IGNORENODES    !node61,node62,node63,node64
IGNOREUSERS    !gridtest1,gridtest2
...
```

*Two completely independent Moab servers schedule the cluster. The first server handles all jobs and nodes except for the ones involved in the test. The second server handles only test nodes and test jobs. While both servers actively talk and interact with a single Torque resource manager, the `IGNORE*` parameters cause them to not schedule, nor even see the other partition and its associated workload.*

**i** When enabling Moab side-by-side, each Moab server should have an independent home directory to prevent logging and statistics conflicts. Also, in this environment, each Moab server should communicate with its client commands using a different port as shown in the previous example.

**i** When specifying the `IGNORENODES` parameter, the exact node names, as returned by the resource manager, should be specified.

### Related Topics

- [Testing New Versions and Configurations](#)

## 14.2 Testing New Middleware

Moab can be used to drive new middleware stress testing resource management systems, information services, allocation services, security services, data staging services, and other aspects. Moab is unique when compared to other stress testing tools as it can perform the

tests in response to actual or recorded workload traces, performing a playback of events and driving the underlying system as if it were part of the production environment.

This feature can be used to identify scalability issues, pathological use cases, and accounting irregularities in anything from LDAP, to NIS, and NFS.

Using Moab's time management facilities, Moab can drive the underlying systems in accordance with the real recorded distribution of time, at a multiplier of real time, or as fast as possible.

In this topic:

[14.2.1 Analysis Aspects - page 747](#)

[14.2.2 General Analysis - page 748](#)

[14.2.3 Native Mode Analysis - page 748](#)

## 14.2.1 Analysis Aspects

The following table describes some aspects of cluster analysis that can be driven by Moab:

System	Details
<b>Accounting Manager</b>	Use <code>test</code> mode to drive scheduling queries, allocation debits, and reservations to accounting packages. Verify synchronization of cluster statistics and stress test interfaces and underlying databases.
<b>On-Demand/Provisioning Services</b>	Use native resource manager mode to drive triggers and resource management interfaces to enable dynamic provisioning of hardware, operating systems, application software, and services. Test reliability and scalability of data servers, networks, and provisioning software, as well as the interfaces and business logic coordinating these changes.
<b>Resource Monitoring</b>	Use <code>test</code> or native resource manager mode to actively load information from compute, network, storage, and software license managers confirming validity of data, availability during failures, and scalability.

With each evaluation, the following tests can be enabled:

- functionality
- reliability
  - hard failure
    - hardware failure - compute, network, and data failures
    - software failure - loss of software services (NIS, LDAP, NFS, database)
    - soft failure
    - network delays, full file system, dropped network packets
  - corrupt data
- performance
- determine peak responsiveness in seconds/request
- determine peak throughput in requests/second
- determine responsiveness under heavy load conditions
- determine throughput under external load conditions
  - large user base (many users, groups, accounts)
  - large workload (many jobs)
  - large cluster (many nodes)
- manageability
  - full accounting for all actions/events
  - actions/failures can be easily and fully diagnosed

**i** If using a native resource manager and you do not want to actually submit real workload, you can set the environment variable `MFORCESUBMIT` to allow virtual workload to be managed without ever launching a real process.

## 14.2.2 General Analysis

For all middleware interfaces, Moab provides built-in performance analysis and failure reporting. Diagnostics for these interfaces are available via the [mddiag](#) command.

## 14.2.3 Native Mode Analysis

Using [native mode](#) analysis, organizations can run Moab in `normal` mode with all facilities fully enabled, but with the resource manager fully emulated. With a native resource



manager interface, any arbitrary cluster can be emulated with a simple script or flat text file. Artificial failures can be introduced, jobs can be virtually running, and artificial performance information generated and reported.

In the simplest case, emulation can be accomplished using the following configuration:

```
SCHEDCFG[natcluster] MODE=NORMAL SERVER=test1.bbli.com
ADMINCFG[1] USERS=dev
RMCFG[natcluster] TYPE=NATIVE CLUSTERQUERYURL=file://$HOME/cluster.dat
```

The preceding configuration will load cluster resource information from the file `cluster.dat`. An example resource information file follows:

```
node01 state=idle cproc=2
node02 state=idle cproc=2
node03 state=idle cproc=2
node04 state=idle cproc=2
node05 state=idle cproc=2
node06 state=idle cproc=2
node07 state=idle cproc=2
node08 state=idle cproc=2
```

In actual usage, any number of node attributes can be specified to customize these nodes, but in this example, only the node state and node configured processors attributes are specified.

The `RMCFG` flag `NORMSTART` indicates that Moab should not actually issue a job start command to an external entity to start the job, but rather start the job logically internally only.

If it is desirable to take an arbitrary action at the start of a job, end of a job, or anywhere in between, the `JOBCFG` parameter can be used to create one or more arbitrary [triggers](#) to initiate internal or external events. The triggers can do anything from executing a script, to updating a database, to using a Web service.

Using native resource manager mode, jobs can be introduced using the [msub](#) command according to any arbitrary schedule. Moab will load them, schedule them, and start them according to all site mission objectives and policies and drive all interfaced services as if running in a full production environment.

## 14.3 Workload Event Format

Moab workload accounting records fully describe all scheduling relevant aspects of batch jobs including resources requested and used, time of all major scheduling events (such as submission time and start time), the job credentials used, and the job execution environment. Each job trace is composed of a single line consisting of either attribute=value pairs or whitespace-delimited fields as shown in the following table. The attribute=value pairs format can be found at [Query Workload Data Format](#).

**i** Moab can be configured to provide this information in flat text tabular form or in XML format conforming to the SSS 1.0 job description specification.

In this topic:

[14.3.1 Workload Event Record Format - page 750](#)

[14.3.2 Reservation Event Records - page 759](#)

[14.3.3 Recording Job Events - page 760](#)


## 14.3.1 Workload Event Record Format

All job events (`JOBSUBMIT`, `JOBSTART`, `JOBEND`, and so forth) provide job data in a standard format as described in the following table:

*Workload Event Record Format*

Field Name	Field Index	Data Format	Default Value	Details
<b>Event Time (Human Readable)</b>	1	HH:MM:SS	-	Specifies time event occurred.
<b>Event Time (Epoch)</b>	2	<epochtime>:<eventID>	-	Specifies time event occurred and the unique event ID.
<b>Object Type</b>	3	job	-	Specifies record object type.
<b>Object ID</b>	4	<STRING>	-	Unique object identifier.
<b>Object Event</b>	5	One of jobcancel, jobcheckpoint, jobend, jobfailure, jobhold, jobmigrate, jobpreempt, jobreject, jobresume, jobstart or jobsubmit	-	Specifies record <a href="#">event type</a> .
<b>Nodes Requested</b>	6	<INTEGER>	0	Number of nodes requested (0 = no node request count specified).
<b>Tasks Requested</b>	7	<INTEGER>	1	Number of tasks requested.
<b>User Name</b>	8	<STRING>	-	Name of user submitting job.
<b>Group Name</b>	9	<STRING>	-	Primary group of user submitting job.
<b>Wallclock</b>	10	<INTEGER>	1	Maximum allowed job

Field Name	Field Index	Data Format	Default Value	Details
<b>Limit</b>				duration (in seconds).
<b>Job Event State</b>	11	<STRING>	-	Job <a href="#">state</a> at time of event.
<b>Required Class</b>	12	<STRING>	[DEFAULT: 1]	<a href="#">Class/queue</a> required by job specified as square bracket list of <QUEUE>[:<QUEUEINSTANCE>] requirements. (For example: [batch:1]).
<b>Submission Time</b>	13	<INTEGER>	0	Epoch time when job was submitted.
<b>Dispatch Time</b>	14	<INTEGER>	0	Epoch time when scheduler requested job begin executing.
<b>Start Time</b>	15	<INTEGER>	0	Epoch time when job began executing. This is usually identical to Dispatch Time.
<b>Completion Time</b>	16	<INTEGER>	0	Epoch time when job completed execution.
<b>Required Node Architecture</b>	17	<STRING>	-	Required node architecture if specified.
<b>Required Node Operating System</b>	18	<STRING>	-	Required node operating system if specified.
<b>Required Node Memory</b>	19	One of >, >=, =, <=, <	>=	Comparison for determining compliance with required node memory.


Field Name	Field Index	Data Format	Default Value	Details
<b>Comparison</b>				
<b>Required Node Memory</b>	20	<INTEGER>	0	Amount of required configured RAM (in MB) on each node.
<b>Required Node Disk Comparison</b>	21	One of >, >=, =, <=, <	>=	Comparison for determining compliance with required node disk.
<b>Required Node Disk</b>	22	<INTEGER>	0	Amount of required configured local disk (in MB) on each node.
<b>Required Node Attributes/Features</b>	23	<STRING>	-	Square bracket enclosed list of node features required by job if specified. (For example: [fast][ethernet])
<b>System Queue Time</b>	24	<INTEGER>	0	Epoch time when job met all fairness policies.
<b>Tasks Allocated</b>	25	<INTEGER>	<TASKS REQUESTED>	Number of tasks actually allocated to job. <div style="border: 1px solid black; padding: 5px; margin-top: 10px;">  In most cases, this field is identical to field #7, Tasks Requested. </div>
<b>Required Tasks Per Node</b>	26	<INTEGER>	-1	Number of Tasks Per Node required by job or '-1' if no requirement specified.
<b>QoS</b>	27	<STRING> [ :<STRING> ]	-	QoS requested/assigned using the format <QoS_

Field Name	Field Index	Data Format	Default Value	Details
				REQUESTED> [:<QOS_DELIVERED>]. (For example: hipriority:bottomfeeder)
<b>JobFlags</b>	28	<STRING> [:<STRING>] . . .	-	Square bracket delimited list of job attributes. (For example: [BACKFILL] [PREEMPTEE]).
<b>Account Name</b>	29	<STRING>	-	Name of account associated with job if specified.
<b>Executable</b>	30	<STRING>	-	Name of job executable if specified.
<b>Resource Manager Extension String</b>	31	<STRING>	-	Resource manager specific list of job attributes if specified. See the <a href="#">Resource Manager Extension Overview</a> for more information.
<b>Bypass Count</b>	32	<INTEGER>	-1	Number of times job was bypassed by lower priority jobs via backfill or '-1' if not specified.
<b>ProcSeconds Utilized</b>	33	<DOUBLE>	0	Number of processor seconds actually used by job.
<b>Partition Name</b>	34	<STRING>	[DEFAULT]	Name of partition where job ran.
<b>Dedicated Processors per Task</b>	35	<INTEGER>	1	Number of processors required per task.

Field Name	Field Index	Data Format	Default Value	Details
<b>Dedicated Memory per Task</b>	36	<INTEGER>	0	Amount of RAM (in MB) required per task.
<b>Dedicated Disk per Task</b>	37	<INTEGER>	0	Amount of local disk (in MB) required per task.
<b>Dedicated Swap per Task</b>	38	<INTEGER>	0	Amount of virtual memory (in MB) required per task.
<b>Start Date</b>	39	<INTEGER>	0	Epoch time indicating earliest time job can start.
<b>End Date</b>	40	<INTEGER>	0	Epoch time indicating latest time by which job must complete.
<b>Allocated Host List</b>	41	<hostname> [, <hostname>]...	-	Comma-delimited list of hosts allocated to job. (For example: node001,node004).
<b>Resource Manager Name</b>	42	<STRING>	-	Name of resource manager if specified.
<b>Required Host List</b>	43	<hostname> [, <hostname>]...	-	List of hosts required by job. (If the job's taskcount is greater than the specified number of hosts, the scheduler must use these nodes in addition to others; if the job's taskcount is less than the specified number of hosts, the scheduler must select needed hosts from this list.)

Field Name	Field Index	Data Format	Default Value	Details
<b>Reservation</b>	44	<STRING>	-	Name of reservation required by job if specified.
<b>Application Simulator Data</b>	45	<STRING> [:<STRING>]	-	Name of application simulator module and associated configuration data. (For example:  HSM:IN=infile.txt:140000;OUT=outfile.txt:500000).
<b>Set Description</b>	46	<STRING>: <STRING> [:<STRING>]	-	Set constraints required by node in the form <SetConstraint> :<SetType> [:<SetList>] where SetConstraint is one of ONEOF, FIRSTOF, or ANYOF, SetType is one of PROCSPEED, FEATURE, or NETWORK, and SetList is an optional colon delimited list of allowed set attributes. (For example: ONEOF:PROCSPEED:350:450:500)
<b>Job Message</b>	47	<STRING>	-	Job messages including resource manager, scheduler, and administrator messages if specified.
<b>Job Cost</b>	48	<DOUBLE>	0.0	Cost of executing job incorporating resource consumption metric, resource quantity consumed, and credential, allocated resource, and delivered QoS charge rates.



Field Name	Field Index	Data Format	Default Value	Details
<b>History</b>	49	<STRING>	-	List of job events impacting resource allocation (XML). <div style="border: 1px solid black; padding: 5px; margin-top: 10px;">  History information is only reported in Moab 5.1.0 and higher. </div>
<b>Utilization</b>	50	Comma-delimited list of one or more of the following: <ATTR>= <VALUE> pairs, where <VALUE> is a double and <ATTR> is one of the following: network (in MB transferred), license (in license-seconds), storage (in MB-seconds stored), or gmetric:<TYPE>.	-	Cumulative resources used over life of job.
<b>Estimate Data</b>	51	<STRING>	-	List of job estimate usage.
<b>Completion Code</b>	52	<INTEGER>	-	Job exit status/completion code.
<b>Extended Memory Load Information</b>	53	<STRING>	-	<b>Deprecated.</b> Extended memory usage statistics (max, mem, avg, and so forth).
<b>Extended CPU Load Information</b>	54	<STRING>	-	Extended CPU usage statistics (max, mem, avg, and so forth).

Field Name	Field Index	Data Format	Default Value	Details
<b>Generic Metric Averages</b>	55	<STRING>	-1	Generic metric averages.
<b>Effective Queue Duration</b>	56	<INTEGER>	-1	The amount of time, in seconds, that the job was eligible for scheduling.
<b>Job Submission Arguments</b>	57	<STRING>	-	The job's submit arguments and script. This field is enabled by setting <a href="#">STOREJOB SUBMISSION</a> to TRUE.

**i** If a field has an empty value, Moab will use a single dash ( - ) as a placeholder in the event record.

**i** Fields that contain a description string such as Job Message use a packed string format. The packed string format replaces white space characters such as spaces and carriage returns with a hex character representation. For example a blank space is represented as \20. Since fields in the event record are space delimited, this preserves the correct order and spacing of fields in the record.

### 14.3.1.A Sample Workload Event

```
13:21:05 110244355 job 1413 JOBEND 20 20 josh staff 86400 Removed [batch:1] 887343658
889585185 \
889585185 889585411 ethernet R6000 AIX53 >= 256 >= 0 - 889584538 20 0 0 2 0 test.cmd \
1001 6 678.08 0 1 0 0 0 0 0 - 0 - - - - - 0.0 - - - 0 - -
```

JOBSUBMISSIONPOLICY Value	Critical Time Based Fields
<b>NORMAL</b>	WallClock Limit Submission Time StartTime Completion Time

JOB SUBMISSION POLICY Value	Critical Time Based Fields
CONSTANTJOBDEPTH CONSTANTPSDEPTH	WallClock Limit StartTime Completion Time

### 14.3.2 Reservation Event Records

All reservation events provide reservation data in a standard format as described in the following table:

Field Name	Field Index	Data Format	Default Value	Details
<b>Event Time (Human)</b>	0	[HH:MM:SS]	-	Specifies time event occurred.
<b>Event Time (Epoch)</b>	1	<epochtime>	-	Specifies time event occurred.
<b>Object Type</b>	2	rsv	-	Specifies record object type.
<b>Object ID</b>	3	<STRING>	-	Unique object identifier.
<b>Object Event</b>	4	One of rsvcreate, rsvstart, rsvmodify, rsvfail or rsvend	-	Specifies record event type.
<b>Creation Time</b>	5	<EPOCHTIME>	-	Specifies epoch time of reservation start date.
<b>Start Time</b>	6	<EPOCHTIME>	-	Specifies epoch time of reservation start date.
<b>End Time</b>	7	<EPOCHTIME>	-	Specifies epoch time of reservation end date.
<b>Tasks Allocated</b>	8	<INTEGER>	-	Specifies number of tasks allocated to reservation at event

Field Name	Field Index	Data Format	Default Value	Details
				time.
<b>Nodes Allocated</b>	9	<INTEGER>	-	Specifies number of nodes allocated to reservation at event time.
<b>Total Active Proc-Seconds</b>	10	<INTEGER>	-	Specifies proc-seconds reserved resources were dedicated to one or more job at event time.
<b>Total Proc-Seconds</b>	11	<INTEGER>	-	Specifies proc-seconds resources were reserved at event time.
<b>Hostlist</b>	12	<comma-delimited list of hostnames>	-	Specifies list of hosts reserved at event time.
<b>Owner</b>	13	<STRING>	-	Specifies reservation <a href="#">ownership</a> credentials.
<b>ACL</b>	14	<STRING>	-	Specifies reservation <a href="#">access control list</a> .
<b>Comment</b>	15	<STRING>	-	Specifies general human readable event message.
<b>Command Line</b>	16	<STRING>	-	Displays the command line arguments used to create the reservation (only shows on the rsvcreate event).

### 14.3.3 Recording Job Events

Job events occur when a job undergoes a definitive change in state. Job events include submission, starting, cancellation, migration, and completion. Some site administrators do not want to use an external accounting system and use these logged events to determine their clusters' accounting statistics. Moab can be configured to record these events in the appropriate event file found in the Moab `stats/` directory. To enable job event recording

for both local and remotely staged jobs, use the [RECORDEVENTLIST](#) parameter. For example:

```
RECORDEVENTLIST JOBCANCEL, JOBCOMPLETE, JOBSTART, JOBSUBMIT  
...
```

*This configuration records an event each time both remote and/or local jobs are canceled, run to completion, started, or submitted. The [Event Logs](#) section details the format of these records.*

---

### Related Topics

- [Event Logging Overview](#)

## Chapter 15: Green Computing

SearchDataCenter.com defines green computing as the environmentally responsible use of computers and related resources. Such practices include the implementation of energy-efficient central processing units (CPUs), servers, and peripherals, as well as reduced resource consumption and proper disposal of electronic waste (e-waste).

The Moab HPC Suite contains power management features that give a Moab administrator the ability to implement policies that can conserve energy and save on operational costs, often without affecting an HPC system's performance with regard to job execution times.

Effective power management means managing power or energy consumption while a compute node is actively running jobs, and when a compute node is idle. Both scenarios require different tools and policies:

- Active compute node power management is mainly performed through control of the clock frequency of the processor(s) on a compute node while a job is executing. Decreasing the clock frequency can reduce energy usage.
- Idle compute node power management is mainly performed by placing a compute node into different low-power states, such as standby and suspend, or no-power states, such as hibernate and shutdown.

### In this chapter:

15.1 Green Computing Methods .....	763
15.1.1 Moab Edition Green Features .....	763
15.1.2 Moab Power Management Methods .....	764
15.1.3 Theory of Operation .....	766
15.1.4 Active Node Power Management .....	772
15.1.5 Idle Node Power Management .....	775
15.1.6 Green Policy Configuration .....	776
15.2 Deploying Adaptive Computing IPMI Scripts .....	776
15.2.1 Prerequisites .....	776
15.2.2 To Deploy the Adaptive Computing IPMI Scripts .....	777
15.3 Choosing which Nodes Moab Powers On or Off .....	778
15.4 Enabling Green Computing .....	779
15.5 Adjusting Green Pool Size .....	782
15.6 Handling Power-Related Events .....	783
15.7 Maximizing Scheduling Efficiency .....	783

<a href="#">15.8 Putting Idle Nodes in Power-Saving States</a> .....	784
<a href="#">15.9 Troubleshooting Green Computing</a> .....	785

## 15.1 Green Computing Methods

In this topic:

- [15.1.1 Moab Edition Green Features](#) - page 763
- [15.1.2 Moab Power Management Methods](#) - page 764
- [15.1.3 Theory of Operation](#) - page 766
- [15.1.4 Active Node Power Management](#) - page 772
- [15.1.5 Idle Node Power Management](#) - page 775
- [15.1.6 Green Policy Configuration](#) - page 776

### 15.1.1 Moab Edition Green Features

The table below identifies the Moab power management features and/or methods available in the Moab HPC Suite.

Feature or Method
<p>CPU Clock Frequency Control:</p> <ul style="list-style-type: none"> <li>• Moab Job Submission Option</li> <li>• Torque Job Submission Option</li> <li>• Moab Job Template Option</li> </ul>
<p>Manual Power Management:</p> <ul style="list-style-type: none"> <li>• Moab-based <code>on</code> and <code>off</code> states</li> <li>• Torque-based low-power and no-power states</li> </ul>
<p>Automated Power Management and Green Policies:</p> <ul style="list-style-type: none"> <li>• Moab-only global-level policies and power management for <code>on</code> and <code>off</code> states</li> <li>• Moab/Moab Web Services-based global, partition, and node-level policies and power management for low-power and no-power states</li> <li>• Green Idle Node Pool Management Policies</li> </ul>

## 15.1.2 Moab Power Management Methods

Moab supports two separate and mutually-exclusive methods for managing the power state of compute nodes, which affects energy consumption. The first method, introduced in Moab 7.2, allows an administrator to manually power on and power off compute nodes and to create a global set of green policies that automatically perform these two functions based on specific conditions involving idle compute nodes. The second method, introduced in Moab 8.0 and Torque 5.0, give an administrator additional power states besides `on` and `off` and offer finer control of green policies at the global, partition, and node levels. Before delving into the theory of operation of these two separate methods, an administrator must understand how Moab views power management regardless of which method is used.

### 15.1.2.A Moab View of Power Management

Moab is not aware of the actual power state of nodes. From Moab's perspective, nodes are only on or off. If Moab needs a node that is off, it issues a power-on job prior to scheduling the incoming job.

In addition, in order to schedule a job to a compute node, Moab requires the compute node's workload resource manager, which in our example is Torque, to report the compute node's state is `idle`. When the compute node's binary power state indicates `on` and the RM indicates the compute node's state is `idle`, Moab will schedule jobs to the compute node. Any value other than `idle` for the node's state and Moab will not schedule a job to the node. If the power state is `off`, Moab issues a power-on job as a dependency to the regular job.

Moab performs compute node power management entirely through power management resource managers, or Power RMs. Each of the two power management methods mentioned above has its own Power RM implementation. The older Moab-only method uses Python-based scripts to implement a power RM while the newer Moab+Moab Web Services (MWS)-based method uses a Java-based MWS RM power management plug-in that runs much simpler Python-based scripts.

These Power RMs perform all power-related management and monitoring, meaning power state control and power state query, respectively, and only report back to Moab whether a compute node is in a state where it can run jobs (`on`) or not (`off`). All actual power state-aware control and management is performed by the power RMs.

### 15.1.2.B Moab Power RMs

Adaptive Computing provides two power management methods to handle different site scenarios; mainly for site-specific security policies. The older method handles sites with a security policy that does not permit web service-based services, which can be an attack vector, or sites that do not want to run an MWS service.



The newer method uses the MWS RM plug-in feature, which allows an administrator to instantiate a separate RM power management plug-in instance for different partitions, or different compute nodes for situations where different compute node hardware requires the use of different power management commands run from Python scripts.

### 15.1.2.C Power Management Scripts

Each power management method, old or new, employs at some point a script that allows the administrator to customize power management for a site, which may be required because the working reference scripts provided by Adaptive Computing (based on OpenIPMI tools) do not use the power management commands specific to the site's vendor-provided hardware.

### 15.1.2.D Moab System Jobs

Moab performs power management functions through a mechanism known as system jobs. A Moab system job is a special, separately scheduled job that performs some Moab system function (e.g., power management, data-staging) that Moab executes on the Moab head node and not on a compute node. This allows Moab to apply policies such as a job wallclock estimate, etc. to system-related functions, which can aid error recovery procedures, etc.

System jobs perform internal Moab-related functions on Moab's behalf, are nearly always script-based, and usually require some customization by the Moab administrator in order to perform the needed function for the HPC system site. For example, the administrator might have to modify power management scripts so they use a site's hardware vendor-specific power management commands to effect power state changes in compute nodes.

To create a system job, Moab internally submits an administrator-defined script, with a path typically specified as a Moab `*URL` parameter, to itself, which it flags as a system job. Moab schedules the job and because it is flagged as a system job, executes the script on the head node. Moab submits a system job whenever it needs to send a power on or off command to a Power RM. Administrators can easily recognize queued and running power management system jobs in the `showq` command output as their job ID has the format `id.poweron` and `id.poweroff`, where `id` is the internally generated Moab job ID number and `.poweron` and `.poweroff` are suffixes appended to the job ID number that represent Moab's on and off commands sent to Power RMs.

### 15.1.2.E Green Policies

Moab provides green policies that automate power management for idle compute nodes, which an administrator can modify and/or configure to control the power state of compute nodes not always in use. These policies allow Moab to dynamically control the power state of compute nodes between the active running state or power-on nodes that may be

needed. It also allows Moab to power-off nodes that are idle and wasting energy. Which power state such compute nodes enter depends entirely on the commands the administrator configures and/or modifies in a power RM's scripts and, for the newer Moab+MWS method, on configuration information specified for each MWS RM power management plug-in instance.

The green policies maintain a green idle node pool, the size of which the administrator configures. As jobs start and use idle nodes from the pool, Moab replenishes the pool by performing an `on` command on those compute nodes on which it previously had performed an `off` command, therefore bringing them into the idle node pool as they enter into an active running state. When jobs finish and the pool has excess idle nodes, Moab performs an `off` command on the excess nodes, which removes them from the idle pool. Therefore, Moab maintains a pool of available idle nodes for immediate use by submitted jobs and reduces energy consumption by powering off any idle nodes in excess of the pool size.

### 15.1.3 Theory of Operation

Moab itself operates the same regardless of the method of power management, Moab-only or Moab+MWS, chosen. This is especially true for the green policies as Moab simply uses the configured power management method to carry out the policies. In order to know how to configure the different parts and components of each power management method so they work well together, it is necessary for a site administrator to understand how the power management methods work; that is, how the components work together to implement a power management method.


#### 15.1.3.A Moab-Only Method

The Moab-only method has a Power RM composed entirely of Python-based scripts. The script must maintain a Power Query daemon that queries the power state of all compute nodes and saves their state for Moab to query, the actual power state query Moab runs to find out the current power state of all compute nodes, and a power state control that places compute nodes into the state of `on` so Moab can schedule jobs to them or into the state of `off` so energy consumption is minimized and operational costs reduced. The administrator determines what the actual power state Moab's `off` represents by configuring the `off` command in the power management control script with the actual hardware vendor-supplied command that effects the desired power state (remember, Moab is not aware of actual power states).

The list below enumerates the advantages and disadvantages of the Moab-only method:

- Advantages
  - Do not have to run the MWS service and its MongoDB database.
  - Power management command scripts execute as Moab system jobs.
  - Ability to customize the node power and cluster query power management scripts
    - For more information on how to specify the node power control script, see the [NODEPOWERURL](#) parameter.
    - For more information on how to specify the power cluster query script, see the [CLUSTERQUERYURL](#) parameter.
  - Moab power control using `mnodectl -mpower=[on|off] <nodelist>`.
    - For more information on how to diagnose power states, see [mdiag -n](#).
- Disadvantages
  - More complex scripts to customize.
  - Only global power management control (no partition-based or node-based).
  - Heterogeneous compute node hardware from different vendors requires more modification of the control and query scripts.
  - Reference scripts not scalable (did not take advantage of Python multi-threading).
  - Administrator must maintain complex scripts that must maintain the entire cluster query information.

The following diagram shows the Moab-only Method Architecture and what occurs between its components:



The Python-based IPMI Monitor daemon script running in the background periodically polls the power state of all compute nodes through IPMI using the command customized by the administrator. As it gathers power state information, it saves the information in a text file in a specific format understood by Moab (binary power state). In order to prevent race conditions, it actually writes to a temporary file and then moves the temporary file on top of the permanent file (not shown).

When Moab starts a scheduling cycle/iteration, it directly executes the power RM's Python-based Cluster Query script that reads the permanent text file and delivers the compute node power states to Moab. Moab then performs the scheduling cycle and based on green policies and the state of the HPC cluster will run the IPMI Node Power script as a Moab system job to perform an `on` or `off` (which may be something different than a `power off`) command using the actual commands customized by the administrator in the script.

### 15.1.3.B Moab+MWS Method

The Moab+MWS method has a Power RM composed of a MWS RM plug-in that encapsulates all power management logic, which itself uses the Torque `pbsnodes` command to effect compute node power state changes into low-power and no-power states of standby and suspend, and hibernate and shutdown, respectively, as well as the IPMI Node Power script to effect compute node power on, power off (pull the plug) and awaken (resume active running state from low-power state). The Power RM Power Management plug-in also performs the power query daemon function identified in the Moab-only

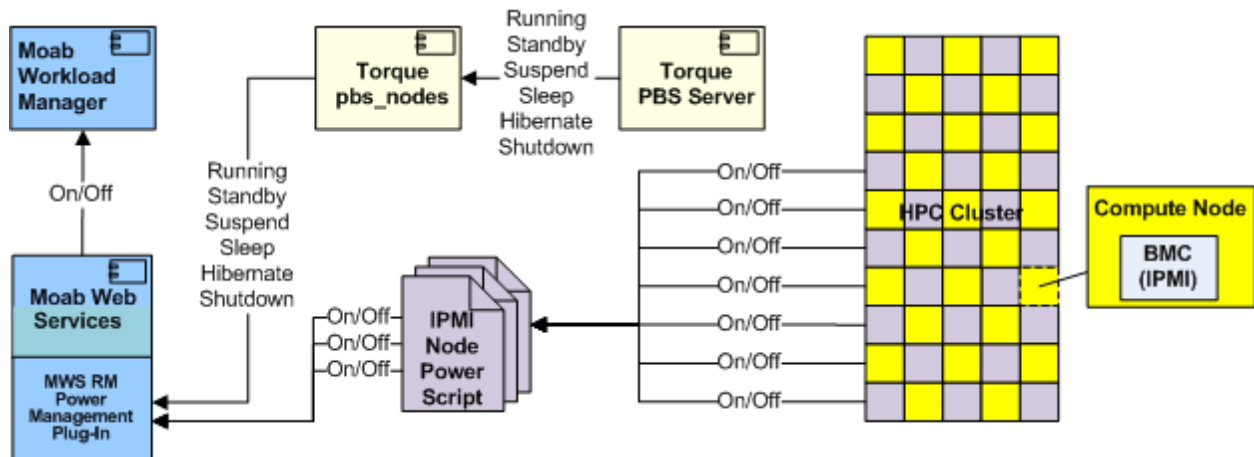
method using its built-in power management logic, therefore handling more actual power states and allowing much better power control than the Moab-only method offers.

The advantages and disadvantages of the Moab+MWS-based method are enumerated below:

- Advantages
  - More power states to choose from.
    - Low-power states of standby and suspend.
    - No-power states of hibernate and shutdown.
    - On and Off (pull the plug) power states still available.
  - Torque power control of low-power and no-power states using `pbsnodes -m <state> <nodelist>`.
    - You can view node power states with the Torque `pbsnodes` command.
  - Power management command scripts execute as Moab system jobs.
  - Much simpler `moab.cfg` [customization and maintenance](#).
  - Global, partition-based, and node-based granularity for power management control.
  - Heterogeneous compute node hardware from different vendors handled by creating multiple instances of MWS RM power management plug-in with different configurations.
  - Reference scripts are scalable (use Python multi-threading).
  - The MWS RM architecture is easier to support DRAC, ILO, and other protocols.
- Disadvantages
  - Must run the MWS service and its MongoDB database.
  - Configuration of the MWS RM Power Management plug-in and possible multiple instances.

The following architecture diagrams show the Moab+MWS-based method architecture and what occurs between its components.

The diagram below illustrates power state query:

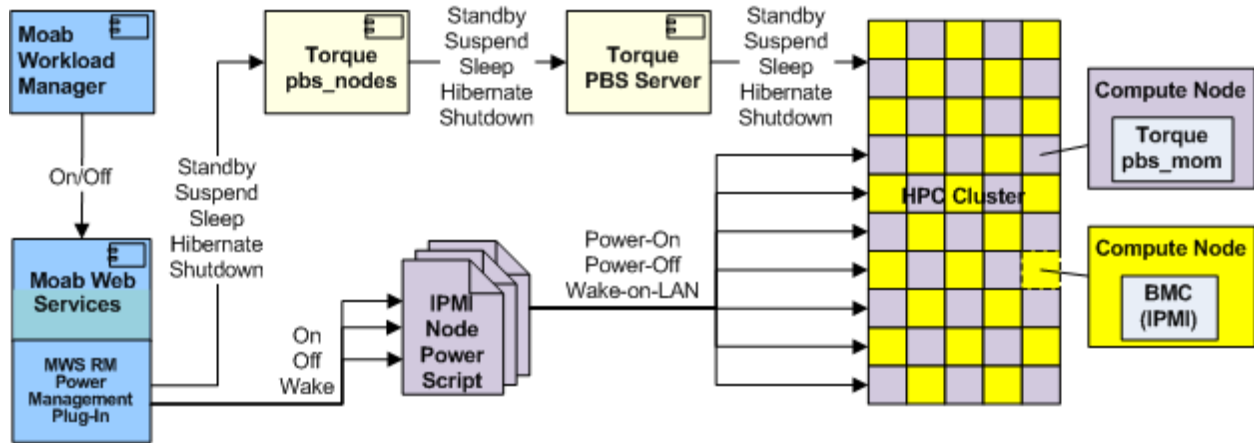


The MWS RM power management plug-in runs the multi-threaded Power Query script for sets of compute nodes that obtain their actual power state through IPMI, or more specifically, a hardware vendor's IPMI implementation (e.g., Dell DRAC, HP iLO, etc), which the RM plug-in saves. It also runs the Torque `pbsnodes` command to obtain the low-power or no-power states that may have been set via Torque earlier (`pbs_server` retains knowledge of any previous command to set a node's power state to one of the low-power or no-power states).

Note it is quite possible for IPMI to report `off` and Torque to report hibernate or shutdown, both of which indicate a compute node has no power, and for IPMI to report `on` and Torque to report standby or suspend, both of which indicate a compute node is in a low-power state from which it can be quickly awakened. It is also possible for IPMI to report `on` and Torque to report hibernate or shutdown, which can indicate a booting node that has not yet started the Torque `pbs_mom` daemon or a node hibernating or shutting down that has not yet powered off. The MWS plug-in's power management logic reconciles the IPMI and Torque reports to produce a single `on` or `off` understood by Moab, which it passes to MWS.

When Moab queries MWS for the current state information of compute nodes at the start of a scheduling cycle/iteration, MWS passes all node information including the binary power on/off Moab understands and the Torque node state, at which point Moab has the information it needs to perform green policy-based automated power management.

The diagram below illustrates Moab+MWS power state control interactions.



When Moab detects a condition that requires changing the power state of a compute node, usually as a result of green policies, it performs the appropriate `on` or `off` command as a system job that sends the command to MWS with a list of the host names of compute nodes that should enter an appropriate power state.

MWS interacts with the appropriate MWS RM power management plug-in for each compute node and passes it the `on` or `off` command. For the `off` command, the plug-in examines its configuration of what `off` means and passes the configured standby, suspend, hibernate, or shutdown command to the Torque `pbsnodes` command, or passes the configured `off` command to the Node Power script.

If the RM plug-in executes the Torque `pbsnodes` command for the configured power state and requested list of compute node host names, it sends the command to the `pbs_server`, which passes the command to each compute node's `pbs_mom` daemon. The `pbs_mom` executes software to place the node into the requested state. The `pbs_server` daemon keeps the requested state in a file for each compute node, which it passes on to the MWS RM power management plug-in as part of a node update report.

**i** In clusters where there is a Torque *pbs\_server* and *pbs\_mom* on the same machine, the administrator should set the `POWERPOLICY` to `STATIC` on this node, because the *pbs\_server* should not be powered down. If the *pbs\_server* is powered down, Moab will be unable to get cluster query updates from all *pbs\_moms* managed by that *pbs\_server*.

On all Torque nodes where *pbs\_moms* are running, the *pbs\_mom* must be configured to auto-start after being rebooted. If the *pbs\_mom* isn't auto-started, the *pbs\_server* will not be able to determine when it has been powered up and entered an idle state, and therefore won't have the ability to inform Moab on a cluster query the node is idle. Refer to Startup/Shutdown Service Script for /Moab (OPTIONAL) in the Torque Administrator Guide for details on how to have the *pbs\_mom* auto-start on boot.

When the RM plug-in executes the Node Power script for the configured `off` power state and requested list of compute node host names, the script executes its `IPMI on` command (whatever the administrator configured in the script) that tells the node's baseboard management controller (BMC) to power off the node.

When the RM plug-in receives the `on` command from Moab via MWS, it checks the internal power state of each compute node in the requested list of compute node host names. If the internal power state is `standby` or `suspend`, the script executes its `IPMI wake` command (whatever the administrator configured in the script) that tells the node's BMC to bump the node into the active running state; otherwise, the script executes its `IPMI off` command (whatever the administrator configured in the script) that tells the node's BMC to power on the node.

**i** Some operating systems require the Wake-on-LAN bit to be enabled using a tool like *ethtool*. Also, Wake-on-LAN packets might be blocked by the router, but not always.

In this manner, the MWS RM power management plug-in queries the actual power state of individual compute nodes and returns to Moab the simple binary on/off state it understands for scheduling jobs to compute nodes. Likewise, Moab controls the actual power state of individual compute nodes using only its simple binary on/off command. This method of simple command and simple job-scheduling-ability state enables Moab to remain scalable and responsive for automatic power management control using green policies.

### 15.1.4 Active Node Power Management

Moab 8.0 and Torque 5.0 introduced support for active node power management; that is, the management of energy consumption while a compute node is running a job, which the new CPU Clock Frequency Control feature provides.

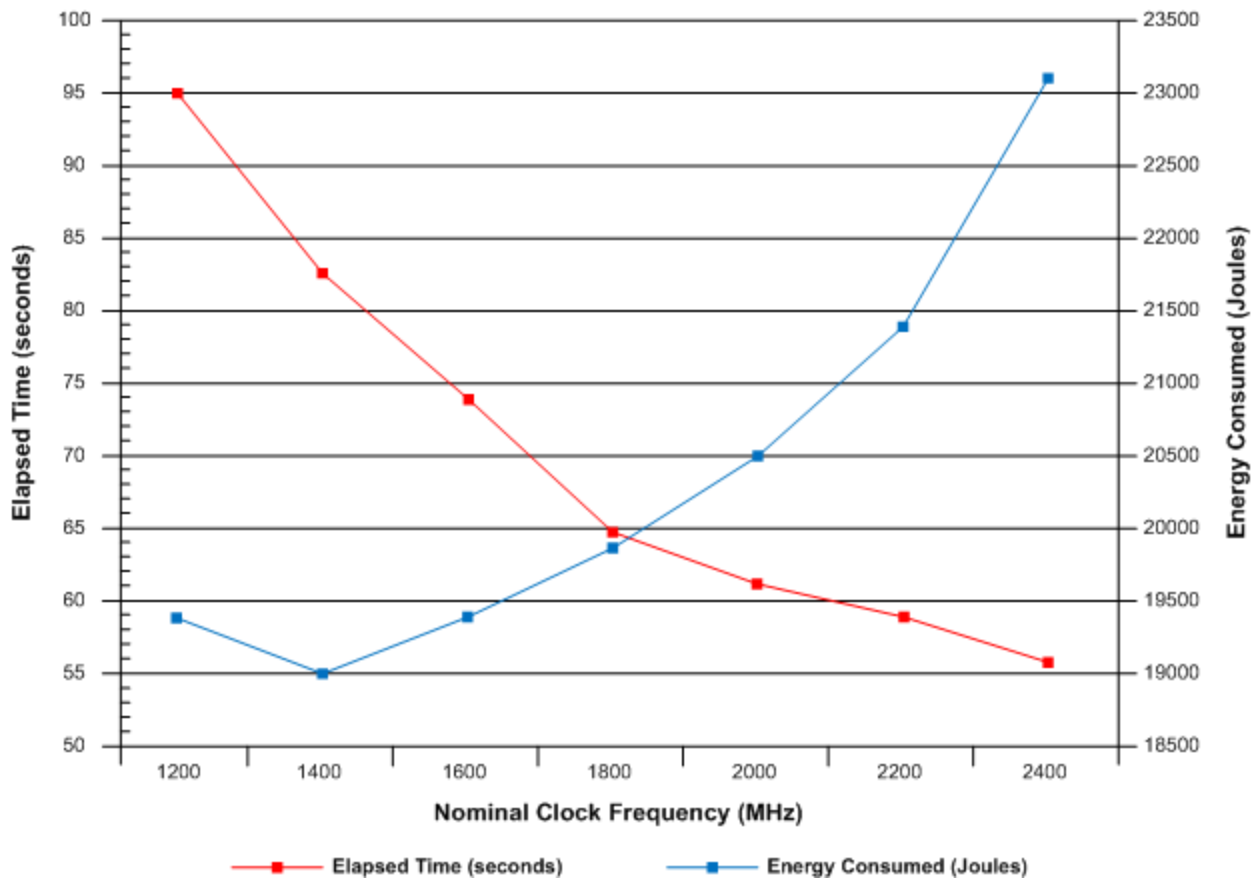


The amount of energy consumption savings achievable through the CPU Clock Frequency Control feature is application-dependent. For example, memory, I/O, and/or network-bound applications, especially memory-bound applications, can often drop the clock frequency of their compute nodes' processors and still have the same execution time even though the compute nodes consume less power. Several studies have shown common power savings of 18-20% and one study showed one application saving 30% on power consumption, all of which translate directly into operational cost savings.


### 15.1.4.A Power/Performance Profiling

To determine whether a lower clock frequency will produce energy consumption savings, applications must be profiled; that is, a job running a particular application with the same or equivalent data must be run at different clock frequencies while measuring the energy consumption of the job's compute node. Each pair of frequency/energy consumption data points are plotted in a chart to show the application's power performance profile. The charts below are an example of two such profiles for two NAS benchmark HPC applications.

**NAS sp.C.64 Power/Performance Profile**



NAS bt.C.64 Power/Performance Profile



**i** The intersection of the two lines has no particular meaning, as each line has its own vertical scale, either on the left or the right as noted.

Note both applications do not consume the least energy (vertical dashed green line) when running at the lowest clock frequency, which demonstrates the importance of profiling applications to determine the nominal clock frequency at which energy consumption is the lowest. The charts amply illustrate why a simplistic policy of using the lowest clock frequency is not the best policy when a site's objective is the least energy consumption possible.

If the least energy consumption is not a site's primary objective, but running jobs in a manner that balances energy consumption and job execution time, a power/performance profile chart is very useful to determine the clock frequency that meets a balanced objective. For example, the vertical dashed purple line in the second chart shows that running the bt.C.64 application at 1800 MHz has an increase in energy consumption of ~1% over the minimal energy consumption possible (vertical dashed green line) but results in a ~10% drop in execution time; a possibly very good trade-off!

Obviously, if a site's primary objective is to complete a job as fast as possible but do so saving energy where possible, profiling memory-bound and other bound applications can clearly show the lowest clock frequency at which the application takes longer to execute. The site would then institute a policy that the application should run at the next highest frequency to fulfill the twin objectives of job performance and energy consumption minimization.

For more information about the CPU clock frequency job submission option, see [CPUCLOCK resource manager extension](#) of `msub -l`.

## Job Templates

Most users will not care or want to know about clock frequency control, so administrators can use a job template to specify the CPU clock frequency at which a particular recurring job should execute. A clock frequency specified on a job template overrides a clock frequency given on the job submission command line or inside a job script file with Torque PBS commands. This order of precedence allows an administrator to control clock frequency for commonly used applications and jobs based on site policies and objectives.

For more information about using a CPU clock frequency job submission option in job templates, see the [CPUCLOCK job template extension attribute](#).

### 15.1.5 Idle Node Power Management

Moab has so-called green policies that together configure Moab to manage and maintain a pool of idle nodes in an active running state so it can immediately schedule jobs to them. When Moab does so and diminishes the pool's idle compute node quantity, it powers on compute nodes by performing an `on` command for nodes in a powered-down state (actually, in a low-power or no-power state) to bring them on-line in order to replenish the pool of idle nodes up to its configured size. When jobs end and the idle node exceed the configured idle node pool size and there are no jobs to run on the now-idle nodes, Moab will power off excess idle nodes by performing an `off` command. In this manner, Moab achieves a site's power management and energy consumption objectives through the configured green policies.

See the [The following diagram shows the Moab-only Method Architecture and what occurs between its components](#): diagram above to see the color-coded compute nodes in the diagram's cluster illustrating Moab's green idle node pool management. The green nodes represent nodes running jobs, the yellow nodes are idle nodes in a green pool of size 12, and the gray nodes represent `off` nodes. Note Moab does not know what actual power state `off` means; what it means will be based on command customization inside Moab-only method scripts or Moab+MWS plug-in configuration information.

In order to perform green policy management of an idle node pool, Moab must first be configured to use either the Moab-only or the Moab+MWS method of power management.

It is *best practice* to configure power management first and test its configuration before configuring green policies. Therefore, if power management is misconfigured, an administrator will know it is the power management configuration and/or scripts and not the green computing policies that are incorrect. If the manual power management commands for the configured power management method work, green computing will work using the configured power management method. For information on how to configure each power management method in Moab, see [Enabling Green Computing](#).

### 15.1.6 Green Policy Configuration

There are several green policies that affect how Moab performs green idle node pool management using automated power management operations. The policies are configured in the same manner regardless of the power management method used, whether Moab-only or Moab+MWS. The other sections of this chapter describe how to configure green policies that manage the idle node pool for site energy management objectives.

---

#### Related Topics

- [15.4 Enabling Green Computing - page 779](#)
- [15.2 Deploying Adaptive Computing IPMI Scripts - page 776](#)
- pbsnodes in the *Torque Administrator Guide*

## 15.2 Deploying Adaptive Computing IPMI Scripts

If you want to enable green computing on your system using the Adaptive Computing supplied IPMI reference scripts, follow the steps here. The IPMI scripts provided are meant as a reference for you to configure the solution to your environment, but can also be used as-is.

### 15.2.1 Prerequisites

- OpenIPMI and ipmitool must be installed and working.
- All nodes must have the same IPMI username and password.
- You must know the IPMI host names and/or IPMI IP addresses of your nodes.
- Python must be installed. The provided IPMI scripts were developed using Python 2.6.5.

- You must identify your Moab home directory. These instructions assume the default Moab home directory of `/opt/moab`.
- You must identify your Moab tools directory. These instructions assume the default Moab tools directory of `/opt/moab/tools`.

## 15.2.2 To Deploy the Adaptive Computing IPMI Scripts

1. Edit the `/opt/moab/tools/ipmi/config.py` script:
  - a. Set `self.ipmiuser` to the IPMI username for your nodes.
  - b. Set `self.ipmipass` to the location of the IPMI password file (`/opt/moab/passfile.txt` by default).

**i** The permissions for the directory and the password file itself should be set so that they can be read only by root or the Moab user running the script.

  - c. Set `self.homeDir` to your Moab home directory.
  - d. If desired, change the `self.pollInterval` value. This is the interval, in seconds, between polls from the IPMI monitoring script.
  - e. The `self.ipmifile` value is the name of a temporary file where the cluster query information is stored. You can change this or leave it alone.
  - f. The `self.bmcaddrmap` value is the filename for the Moab node name/IPMI mapping. The file must exist in the Moab home directory and will be created in the next step.
2. Create a `node-bmc.txt` file in the Moab home directory. The file must contain a space-delimited list of Moab node names that map to the IPMI host names or IP address. For Example:
 

```
node01 node01_ipmi # For all three of these entries, the first value is the
node02 node02_ipmi # node name as Moab knows it. The second value is either
node03 10.1.1.1    # the node IPMI name or IPMI IP address.
```
3. Configure the `moab.cfg` file for green computing as described in [Enabling green computing](#). Use the `ipmi.mon.py` script for the `CLUSTERQUERYURL` and the `ipmi.power.py` script for the `NODEPOWERURL`.
4. Restart Moab and verify green computing is working correctly. If you encounter trouble, see the [Troubleshooting green computing](#) topic for help.

---

## Related Topics

- [15.4 Enabling Green Computing - page 779](#)
- [15.9 Troubleshooting Green Computing - page 785](#)
- [15.5 Adjusting Green Pool Size - page 782](#)
- [15.6 Handling Power-Related Events - page 783](#)
- [15.7 Maximizing Scheduling Efficiency - page 783](#)

## 15.3 Choosing which Nodes Moab Powers On or Off

Moab can use the GREENPOOLPRIORITYF function to determine which nodes to power on or off. The [PRIORITY node allocation policy](#) is used to determine which nodes to allocate workload to. When Moab can no longer allocate workload to available nodes, it begins to power nodes on in the order specified by the GREENPOOLPRIORITYF function.

### To Choose which Nodes Moab Powers On or Off

Set a GREENPOOLPRIORITYF function to describe which order nodes should be selected for power on/off actions. GREENPOOLPRIORITYF uses the [PRIORITY node allocation policy](#) options and syntax.

```
GREENPOOLPRIORITYF '10*RANDOM'
```

*This tells Moab to randomly choose a node to power on to meet workload demands, and to randomly choose an idle node to power off to meet the [MAXGREENSTANDBYPOOLSIZ](#) goal.*

### To Choose which Nodes Moab Allocates Jobs to

Set a PRIORITY node allocation policy that uses power as the major factor. This causes Moab to allocate jobs to nodes that are already powered on. When no nodes are available to meet this policy, Moab uses the GREENPOOLPRIORITYF function to turn on nodes that are powered off.

```
NODEALLOCATIONPOLICY PRIORITY
NODECFG[DEFAULT] PRIORITYF='10000*POWER + 10*RANDOM'
```

*The nodes with the highest priority for workload are the nodes that are powered on. After that, Moab randomly allocates workload.*

---

## Related Topics

- [15.5 Adjusting Green Pool Size - page 782](#)
- [15.7 Maximizing Scheduling Efficiency - page 783](#)

## 15.4 Enabling Green Computing

There are two ways to do green computing in Moab. With just Moab, nodes can be turned on or off. With MWS, however, you can put nodes into several low-power states. The MWS solution is also more scalable. The supported low-power states are:

- Running
- Standby
- Suspend
- Hibernate
- Shutdown

Nodes cannot be moved from one low-power state to another. The node must go from low-power to running, and then to the new low-power state.

### To Enable Green computing with Moab and MWS

1. Edit `moab.cfg` to use MWS for green computing:
  - a. Configure the **POWERPOLICY** attribute of the **NODECFG** parameter. The default value is `STATIC`. Set it to `OnDemand`.
  - b. Set the resource manager type as `MWS`
  - c. Set `FLAGS=UserSpaceIsSeparate` for the MWS resource manager.
  - d. Point `BASEURL` to your MWS server.

```

NODECFG [DEFAULT]      POWERPOLICY=OnDemand
RMCFG [mws]           TYPE=MWS
RMCFG [mws]           FLAGS=UserSpaceIsSeparate
RMCFG [mws]           BASEURL=https://localhost:8080/mws

```

2. Configure the MWS Power Management Plugin.

### To Enable Green Computing with just Moab

1. Edit `moab.cfg` to enable green computing. There are four things you must configure for basic functionality of green computing:
  - a. Configure the **POWERPOLICY** attribute of the **NODECFG** parameter. The default value is `STATIC`. Set it to `OnDemand`.

- b. Configure a power provisioning resource manager to be `TYPE=NATIVE` and `RESOURCETYPE=PROV`. The resource type of `PROV` means the RM works only with node hardware and not workloads.
- c. Configure a `CLUSTERQUERYURL` attribute of the power provisioning RM to point to the power query script you'd like to use. Moab uses this script to query the current power state of the nodes. `CLUSTERQUERYURL` is traditionally used as a workload query but is also used by green computing for the node power state query. Adaptive Computing provides a reference [IPMI script](#) you can use.
- d. Configure a `NODEPOWERURL` attribute of the power provisioning RM to point to the power action script you'd like to use. Moab uses this script to turn nodes on or off. Adaptive Computing provides a reference [IPMI script](#) you can use.

```
NODECFG[DEFAULT] POWERPOLICY=OnDemand
RMCFG[ipmi] TYPE=NATIVE RESOURCETYPE=PROV
RMCFG[ipmi] CLUSTERQUERYURL=exec://$TOOLS DIR/ipmi/ipmi.mon.py
RMCFG[ipmi] NODEPOWERURL=exec://$TOOLS DIR/ipmi/ipmi.power.py
```

## Sample moab.cfg for Green Computing

Below is a sample `moab.cfg` configuration file of a green computing setup using the Adaptive Computing IPMI scripts:

```
#####
#
# Use 'mdiag -C' to validate config file parameters
#
#####

SCHEDCFG[Moab]          SERVER=myhostname:5150
ADMINCFG[1]             USERS=myusername,root
TOOLS DIR              /$HOME/tools
LOGLEVEL                1

#####
#
# Basic Resource Manager configuration
#
# For more information on configuring a Resource Manager, see:
# docs.adaptivecomputing.com
#
#####

RMCFG[local]           TYPE=NATIVE
RMCFG[local]           CLUSTERQUERYURL=exec://$HOME/scripts/query.resource
RMCFG[local]           WORKLOADQUERYURL=exec://$HOME/scripts/query.workload

RMCFG[local]           JOBSUBMITURL=exec://$HOME/scripts/submit.pl
RMCFG[local]           JOBSTARTURL=exec://$HOME/scripts/job.start
RMCFG[local]           JOBCANCELURL=exec://$HOME/scripts/job.cancel
RMCFG[local]           JOBMODIFYURL=exec://$HOME/scripts/job.modify
RMCFG[local]           JOBREQUEUEURL=exec://$HOME/scripts/job.requeue
RMCFG[local]           JOBSUSPENDURL=exec://$HOME/scripts/job.suspend
RMCFG[local]           JOBRRESUMEURL=exec://$HOME/scripts/job.resume
```



```
#####
# GREEN configuration:
#####
# Turn on "green" policy. (This is the policy that enables green computing).
# Here we are doing it for all nodes, but it can be controlled on a node-by-node basis
# Default is STATIC, which means green computing is disabled.
#NODECFG[DEFAULT] POWERPOLICY=STATIC
NODECFG[DEFAULT] POWERPOLICY=OnDemand

# Configure the power provisioning and power state query scripts for the power
# management system.
# Note that this is an entirely different RM (with a name of power in this case
# and a type of 'PROV').
# The PROV type RM is the only one that uses a NODEPOWERURL. Additionally, the
# output of the CLUSTERQUERYURL for this type of RM is different. (See docs)
RMCFG[mws] TYPE=MWS
RMCFG[mws] FLAGS=UserSpaceIsSeparate
RMCFG[mws] BASEURL=https://localhost:8080/mws

# We want green policy to work so it allocates jobs to compute nodes already
# powered on and will power on powered-off compute nodes only when there are
# no powered-on compute nodes available. This requires using the PRIORITY
# node allocation policy with a PRIORITYF function that has the POWER variable
# as the greatest contributing factor to the function (1 = powered-on,
# 0 = powered-off).
# If we want all compute nodes to operate under green policy, we can assign
# the PRIORITYF function to the default node configuration, which is easier
# than assigning it to individual compute nodes. If only some compute nodes
# should operate under green policy, then the PRIORITYF function must be
# configured for the individual nodes. Note the POWER variable must be the
# largest factor in the function below; it is assigned the largest multiplier,
# which should be greater than the sum of all other factors! Doing so forces
# Moab to use all eligible powered-on nodes for workload placement before
# powering on any eligible powered-off nodes.

# Enable PRIORITYF functionality
NODEALLOCATIONPOLICY PRIORITY

# Use a priority function that uses power as the major factor (plus some other
# imaginary factors)
#NODECFG[DEFAULT] PRIORITYF='1000000*POWER + 1000*factor2 + 100*factor3...'
# Use a priority function where power is the only factor.
#NODECFG[DEFAULT] PRIORITYF='10000*POWER'
# Use a priority function that adds some randomness but uses power as the major
# factor.
NODECFG[DEFAULT] PRIORITYF='10000*POWER + 10*RANDOM'

# Set a priority function that specifies the order nodes should be chosen to power
# up/down. By default, Moab will start at the top of the node list and go down. Some
# installations want to rotate power cycles among nodes in a different order.
# The configuration below forces Moab to power on/off random nodes, which
# eventually guarantees all nodes occasionally go through a power cycle.
#GREENPOOLPRIORITYF '10*RANDOM'

# Ensure we are recording power management events
# (powering on and off nodes are recorded as "node modification" events).
#RECORDEVENTLIST +NODEMODIFY

# Set the size of the standby pool. This is the number of idle nodes that will
# be powered on and idle. As the workload changes, Moab turns nodes on
# or off to try to meet this goal.
# Default value is -1 (all idle nodes are left powered on by default)
```

**MAXGREENSTANDBYPOOLSIZE 5**

```
# Set the length of time that it takes to power a node on/off. This will be the
# walltime of the system job that performs the power operation and should be the
# maximum expected time. If Moab detects (via the power RM) that the power
# operations have all completed, the system job will finish early.
# Default value is 10 minutes (600)
PARCFG[ALL]          NODEPOWEROFFDURATION=600
PARCFG[ALL]          NODEPOWERONDURATION=600
# Set the length of time a node should remain idle before it is powered off.
# This prevents Moab from immediately powering off nodes that have just finished
# a job. Increasing this number should decrease power on/off thrashing
# This should be set higher than NODEPOWEROFFDURATION and/or NODEPOWERONDURATION
NODEIDLEPOWERTHRESHOLD 660

# If a node fails to power on, we need to remove it from the available nodes so
# Moab won't keep [re-]trying to power it on. Do this by setting a reservation
# on the failed node to give time for manual investigation.
#RMCFG[torque] NODEFAILURERSVPROFILE=failure
#RSVPROFILE[failure] DURATION=3600
```

**Related Topics**

- [15.2 Deploying Adaptive Computing IPMI Scripts - page 776](#)
- [15.3 Choosing which Nodes Moab Powers On or Off - page 778](#)
- [15.5 Adjusting Green Pool Size - page 782](#)
- [15.6 Handling Power-Related Events - page 783](#)
- [15.7 Maximizing Scheduling Efficiency - page 783](#)
- [15.9 Troubleshooting Green Computing - page 785](#)
- Power Management Plugin in the *Moab Web Services Reference Guide*

## 15.5 Adjusting Green Pool Size

The `MAXGREENSTANDBYPOOLSIZE` parameter allows you to allocate the number of nodes to keep powered on in the standby pool. This is the number of idle nodes that are allowed to be powered on and idle. As the workload changes, Moab turns nodes on or off to try to meet this goal. The default value is -1, which disables the standby pool.

### To Adjust the Green Pool Size

Modify the `MAXGREENSTANDBYPOOLSIZE` parameter with the number of nodes you want Moab to keep powered on for the standby pool:

```
MAXGREENSTANDBYPOOLSIZE 10
```

*Moab keeps up to 10 idle nodes powered on to be kept on standby.*

**Related Topics**

- [15.7 Maximizing Scheduling Efficiency - page 783](#)
- [15.3 Choosing which Nodes Moab Powers On or Off - page 778](#)

## 15.6 Handling Power-Related Events

Power actions are considered `NODEMODIFYURL` events and are not recorded by default, but you can configure Moab to include power-related events in the logs. Also, if a node fails to turn on (or off), it's best to associate a reservation on the failed node so that Moab won't keep trying to perform the power action over and over.

### To Configure Moab to Record Power-Related Events

Modify the `RECORDEVENTLIST` parameter:

```
RECORDEVENTLIST +NODEMODIFY
```

*Power-related events are logged to the Moab log file.*

### To Put a Reservation on a Node that Fails to Perform a Power Action

Configure the `NODEFAILURERSVPROFILE` attribute of `RMCFG` and create an `RSVPROFILE` with a high duration:

```
RMCFG[torque] NODEFAILURERSVPROFILE=failure
RSVPROFILE[failure] DURATION=3600
```

*Nodes that fail to power on or off have a 1-hour reservation placed on them.*

**Related Topics**

- [RECORDEVENTLIST](#)
- [12.2.4 Event Logs - page 722](#)

## 15.7 Maximizing Scheduling Efficiency

When considering whether to power a node on or off, Moab can take into account the amount of time that it takes to power on or power off the node. With this information, Moab can keep an idle node powered on if it knows that workload in the queue will be ready for the node in less time that it takes to power off/power on the node.

Moab can also wait to shut down nodes after they've been idle for a specific amount of time.

## To Specify Node Power On/Power Off Duration

Modify the `NODEPOWERONDURATION` and `NODEPOWEROFFDURATION` attributes of `PARCFG` with the maximum amount of time it takes for your nodes to power on/power off. Make sure to use the keyword `ALL` for the resource manager name to avoid cases where Moab won't consider the power on/off duration for a node before making a power action decision.

```
PARCFG [ALL] NODEPOWERONDURATION=2:00
PARCFG [ALL] NODEPOWEROFFDURATION=2:00
```

*If a node goes idle and has to wait for workload, Moab will not power off the node if the workload will be available within 4 minutes or less.*

## To Shut Down on Nodes after they've been Idle for a Specified Time

Modify the `NODEIDLEPOWERTHRESHOLD` parameter with the duration (in seconds) you want Moab to wait before shutting down an idle node. The default value is 60 seconds. Increasing the number should decrease power on/off thrashing. This should be set higher than `NODEPOWERONDURATION` and/or `NODEPOWEROFFDURATION`.

```
NODEIDLEPOWERTHRESHOLD 300
```

*Moab will wait 5 minutes before shutting down a node that has become idle.*

---

### Related Topics

- [15.5 Adjusting Green Pool Size - page 782](#)
- [15.3 Choosing which Nodes Moab Powers On or Off - page 778](#)

## 15.8 Putting Idle Nodes in Power-Saving States

When nodes exceed their idle threshold limits, the default behavior is to turn the nodes off. With the `NODEIDLEPOWERACTION` parameter, you can choose which power-saving state to put idle nodes into. This parameter is configured at the partition level. Configuring it for the `ALL` partition effectively makes it a global parameter.

### To specify what to do with idle nodes

Modify the `NODEIDLEPOWERACTION` parameter.

```
NODEIDLEPOWERTHRESHOLD 300
PARCFG[ALL] NODEIDLEPOWERACTION SLEEP
```

*All nodes that are idle for more than 5 minutes are put into a sleep state.*

## Related Topics

- [15.3 Choosing which Nodes Moab Powers On or Off - page 778](#)

## 15.9 Troubleshooting Green Computing

If you've enabled green computing and are having trouble, here are some tips that can help you determine the cause of the issues you encounter. These tips are specifically for the [Adaptive Computing supplied IPMI scripts](#), but can be generalized for whatever power management solution you use. Simply substitute your power management system, power query script (as specified by `CLUSTERQUERYURL`), and power action script (as specified by `NODEPOWERURL`) where appropriate.

### Verify your IPMI Access

Use the `ipmitool` command to verify you have access to the IPMI interface of your nodes. Try getting the current power state of a node. The syntax is `ipmitool -I lan -H <host> -U <IPMI username> -P <IPMI password> chassis power status`.

```
$ ipmitool -I lan -H qt06 -U ADMIN -P ADMIN chassis power status
Chassis Power is off
```

### Verify the Power Query (`CLUSTERQUERYURL`) Script is Working

1. Execute the `impi.mon.py` script (should be found in `/<MOABHOMEDIR>/tools/ipmi`) to start the monitor:

```
$ cd /opt/moab/tools/ipmi
$ ./impi.mon.py
```

2. Execute the script again. The following is an example of the expected output:

```
$ ./impi.mon.py

qt09  GMETRIC[System_Temp]=27 GMETRIC[CPU_Temp]=25 POWER=on State=Unknown
qt08  GMETRIC[System_Temp]=31 GMETRIC[CPU_Temp]=25 POWER=on State=Unknown
qt07  GMETRIC[System_Temp]=30 GMETRIC[CPU_Temp]=29 POWER=on State=Unknown
qt06  GMETRIC[System_Temp]=Disabled GMETRIC[CPU_Temp]=Disabled POWER=off
State=Unknown
```

*If the **POWER** attribute is not present the script is not working correctly.*

## Verify the Power Action (NODEPOWERURL) Script is Working

1. Execute the `ipmi.power.py` script (should be found in `/<MOABHOMEDIR>/tools/ipmi`) to see if you can force a node to power on or off. The syntax is `ipmi.power.py <node>, <node>, <node>... [off|on]`

```
$ /opt/moab/tools/ipmi/ipmi.power.py qt06 off
```

*This example is trying to power off a node named qt06.*

2. Verify the machine's power state was changed to what you attempted in the previous step. You can do this remotely via two methods:
  - a. If the [cluster query script](#) is working, you can use that to verify the current power state of the node.
  - b. If you have [IPMI access](#), you can use the `ipmitool` command to verify the current power state of the node.

## Verify the Scripts are Configured Correctly

1. Run the `mdiag -R` command to verify your IPMI resource manager configuration:

```
$ mdiag -R -v
RM[ipmi]      State: Active  Type: NATIVE  ResourceType: PROV
Timeout:      30000.00 ms
Cluster Query URL:  exec://$TOOLS DIR/ipmi/ipmi.mon.py
Node Power URL:   exec://$TOOLS DIR/ipmi/ipmi.power.py
Objects Reported:  Nodes=3 (0 procs)  Jobs=0
Nodes Reported:  3 (N/A)
Partition:      SHARED
Event Management: (event interface disabled)
RM Performance:  AvgTime=0.05s  MaxTime=0.06s  (176 samples)
RM Languages:    NATIVE
RM Sub-Languages: NATIVE
```

2. Run the `mdiag -G` command to verify that power information is being reported correctly:

```
$ mdiag -G

NodeID      State      Power  Watts  PWatts
qt09        Idle      On      0.00   0.00
qt08        Idle      On      0.00   0.00
qt07        Idle      Off     0.00   0.00
```

## Verify the Scripts are Running

Once green is configured and Moab is running, Moab should start the power query script automatically. Use the `ps` command to verify the script is running:

```
$ ps -ef | grep <CLUSTERQUERYURL script name>
```

*If this command does not show the power query script running then your settings in `moab.cfg` aren't working.*

## Verify Moab can Power Nodes On or Off

Use the `mnodectl` command to turn a node on or off. The syntax is `mnodectl -m power=[off|on] <node>`.

```
mnodectl -m power=off qt06
```

*Moab should turn off the node named qt06.*

1. Moab generates a system job called `poweron-<num>` or `poweroff-<num>` job as shown in [showq](#). The system job calls the `ipmi.power.py` (NODEPOWERURL) script to execute the command.
2. Moab waits until the cluster query reports the correct data. In this case, the `ipmi.power.py` script reports that the power attribute has changed.
3. Moab does not change the power status based on the power script return code. Rather, Moab completes the system power job when it detects the power attribute has changed as indicated by the cluster query script.

---

### Related Topics

- [15.4 Enabling Green Computing - page 779](#)
- [15.2 Deploying Adaptive Computing IPMI Scripts - page 776](#)

## Chapter 16: Elastic Computing Overview

**i** Elastic Computing is an add-on package for Moab Workload Manager. Contact your Adaptive Computing account manager for more information.

**i** Elastic Computing is only applicable for Torque Resource Manager and Native RMs with QoS triggers.

**i** Elastic Computing is not supported on Ubuntu.

During the course of operation, the number of job requests will go up and down. Under some circumstances the job backlog might increase to the point where additional resource are required to complete the job backlog in a reasonable time frame. In this scenario, the job will be held until resources become available. The Elastic Computing feature in Moab allows the Moab scheduler to take advantage of systems that can temporarily provide additional nodes (for example, to create new virtual machines or borrow physical nodes from another system) to fulfill the workload demand in a more timely manner.

Moab's Elastic Computing framework serves as a basis for Moab/ODDC Cloud Bursting, which can be configured to access multiple cloud providers either on-demand or based on a job backlog. For more information about Moab/ODDC Cloud Bursting, see the *ODDC Connect User Guide*.

**⚠** This chapter provides examples of the Elastic Computing and node end scripts. Your scripts will vary based on your system configuration. Contact your Adaptive Computing account manager for suggestions and options to configure Elastic Computing.

In this chapter:

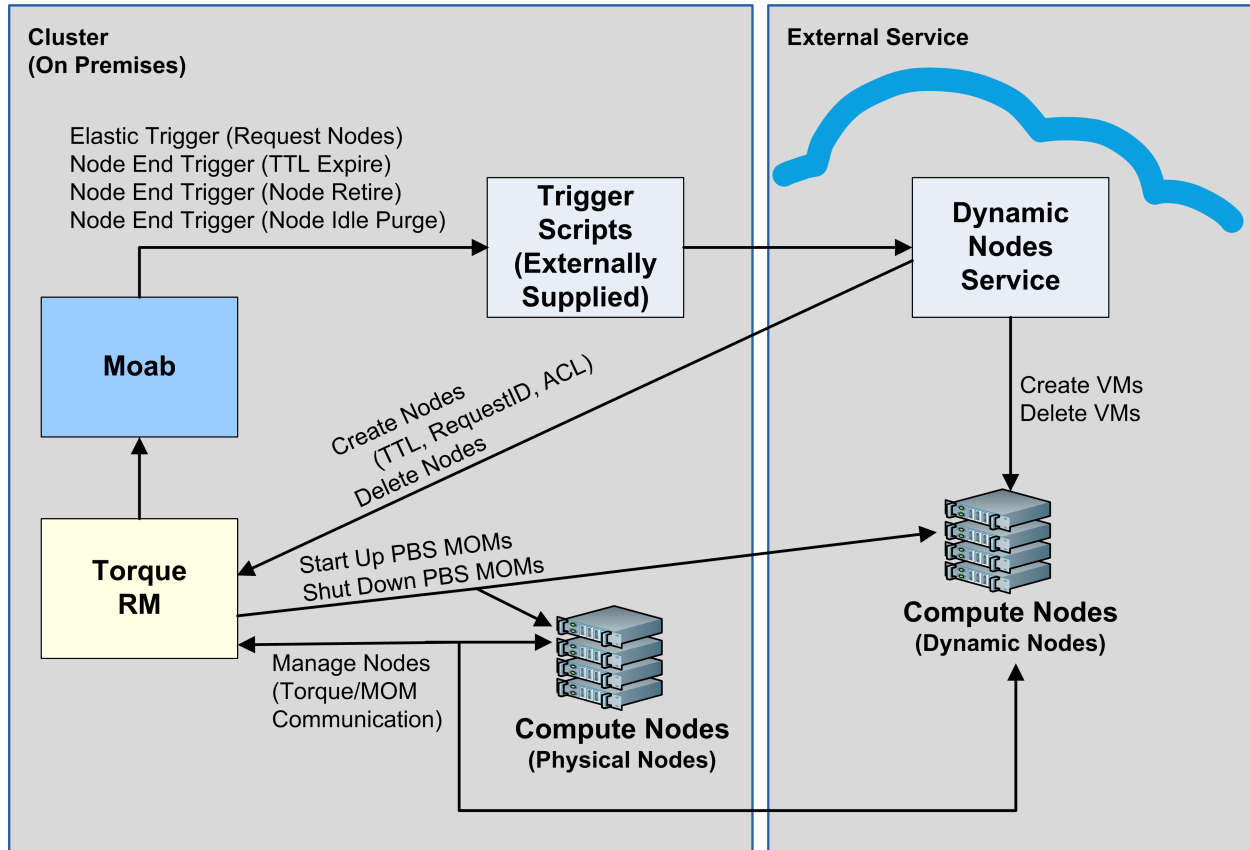
16.1 About Elastic Computing .....	790
16.2 Configuring Elastic Computing .....	791
16.2.1 To Configure Elastic Computing .....	791
16.2.2 Sample moab.cfg File Excerpt .....	794
16.3 Elastic Trigger .....	794
16.4 Integration with a Private OpenStack Cloud .....	795
16.4.1 Configuring Moab to Talk to OpenStack Integration Scripts .....	795



16.4.2 Verification .....	796
16.4.3 Troubleshooting .....	797
16.5 Dynamic Nodes .....	797
16.5.1 Dynamic Node Parameters .....	797
16.5.2 Dynamic Node Events .....	798
16.5.3 Configuring Dynamic Nodes .....	799
16.6 Viewing Node and Trigger Information .....	801
16.6.1 mdiag -n -v --xml .....	801
16.6.2 mdiag -T .....	802
16.6.3 checknode -v <node name> .....	802
16.7 Usage Policies .....	803
16.7.1 Available Policies .....	803
16.7.2 Policy Levels .....	804

## 16.1 About Elastic Computing

The diagram below depicts Moab's Elastic Computing feature:





With the Elastic Computing feature enabled and configured:

1. Moab monitors the job backlog and, when a predefined threshold is reached, fires the elastic trigger. The elastic trigger calls a script to request additional nodes (dynamic nodes) from an external service.
2. The procured dynamic nodes are then added to the resource manager (RM). For example, via `qmgr` in Torque.
3. Moab then begins scheduling jobs for the allocated dynamic nodes.
4. When the job backlog is cleared and the dynamic nodes become idle for a specified amount of time (for example, defined by the `NODEIDLEPURGETIME` parameter), Moab fires an end node trigger to remove the nodes from the RM and deprovision the virtual machine or physical nodes.

Alternatively, you can utilize the Elastic Computing feature for interaction with a private OpenStack cloud. See [16.4 Integration with a Private OpenStack Cloud - page 795](#).

## 16.2 Configuring Elastic Computing

 This topic provides examples of the Elastic Computing and node end scripts. Your scripts will vary based on your system configuration. Please contact your Adaptive Computing account manager for suggestions and options to configure Elastic Computing.

 If you are using Elastic Computing with Torque, you cannot have a `mom_hierarchy` file in the `$PBS_HOME/server_priv` directory.

In this topic:

- [16.2.1 To Configure Elastic Computing - page 791](#)
- [16.2.2 Sample moab.cfg File Excerpt - page 794](#)

### 16.2.1 To Configure Elastic Computing

1. If you installed Moab Workload Manager from the tarball (Manual Installation), ensure you installed `acpython-base` RPM on the Moab Head Node:

```
[root]# rpm -qa | grep acpython-base
```

If it is not, follow the instructions in Install Moab Server in the *Moab 10.0.1 Installation and Configuration Guide*.

2. Enable dynamic nodes in the `moab.cfg` file:

```
SCHEDCFG[] FLAGS=enabledynamicnodes
```

A sample excerpt from a `moab.cfg` file is shown below.

3. If you want to be able to view node and trigger information, use one of these Moab tools:

- `mdiag -n -v -xml`
- `mdiag -T`
- `checknode -v <node name>`

See [16.6 Viewing Node and Trigger Information - page 801](#) for more information.

4. If you want to record dynamic node activity, enable `NODEADD` and/or `NODEREMOVE` for `RECORDEVENTLIST` in the `moag.cfg` file. See [16.5 Dynamic Nodes - page 797](#) for more information.
5. In the `moab.cfg` file, make these changes for QoS triggers:
  - a. Add the elastic trigger: `TType=elastic`. See [16.3 Elastic Trigger - page 794](#) for more information.
  - b. Specify how nodes are requested when the trigger fires, using one of these two options:

- `QOSCFG[xyz] REQUESTGEOMETRY=12@4:00:00:00`

When the elastic trigger fires, request 12 additional nodes for 4 days, 0 hours, 0 minutes, and 0 seconds.

 The `REQUESTGEOMETRY` values shown are just an example.

- `QOSCFG[xyz] REQUESTGEOMETRY=PRIORITYJOBSize`

When the elastic trigger fires, request enough nodes to run the highest priority job in the backlog for the amount of walltime specified by the highest priority job.

6. Use `BACKLOGCOMPLETIONTIME` to specify when the elastic trigger fires (adding nodes).

 The trigger that contains the `BACKLOGCOMPLETIONTIME` threshold can only be used when profiling is enabled.

The `BACKLOGCOMPLETIONTIME` is calculated by Moab as follows: (The maximum number of processor seconds in the QoS) divided by (The total number of processors in the system). See [BACKLOGCOMPLETIONTIME](#) for more information.

```
NODECFG[DEFAULT] ENABLEPROFILING=TRUE
QOSCFG[xyz] TRIGGER=EType=threshold,AType=exec,Action="$HOME/tools/elastic.py
$REQUESTGEOMETRY",Threshold=BACKLOGCOMPLETIONTIME>1800,RearmTime=05:00
```

In the above example, when the `BACKLOGCOMPLETIONTIME` is more than 1800 seconds, the `QOSCFG` threshold trigger will fire. When the `QOSCFG` trigger is fired, the `$HOME/tools/elastic.py` script is executed. This is a user-supplied script that needs to create virtual machines or provision physical hardware and add these dynamic nodes to the RM.

The following examples show commands that the script will run in order to create a node on Torque:

```
qmgr -c "create node node01 np=4,TTL=2022-09-26T12:00:00Z,acl='user==user1',requestid=1234"
```

**i** Once the `BACKLOGCOMPLETIONTIME` threshold is reached, the trigger will begin firing. The administrator can configure the trigger to fire once only or periodically until the node is deleted from Torque by the external service.

7. Determine how the dynamic nodes will be removed. See [16.5 Dynamic Nodes - page 797](#) for more information.

Use one or both of these methods:

- Set the `TTL` when creating the node via the RM. This parameter tells Moab to remove the node when the `TTL` has passed.
- Add the `NODEIDLEPURGETIME` parameter to `moab.cfg`. To turn off the purging of *individual* dynamic nodes output, specify "noidlepurge" in the `varattr` output of the node using Torque. See '\$varattr' in the *Torque Resource Manager Administrator Guide*. Alternatively, you can use the `varattr` output from the wiki interface. See [VARATTR](#).

You can optionally report a `requestid` on each node in the same group.

**i** Nodes without a `requestid` that hit the configured idle purge time are immediately purged. Whereas, nodes with a `requestid` that hit the configured idle purge time are only purged when all the nodes that have the same `requestid` hit the configured idle purge time.

Configure the node end trigger in `moab.cfg`.

```
NODECFG[DEFAULT]
TRIGGER=EType=end,TType=elastic,AType=exec,Action="/$HOME/tools/nodeend.sh $OID"
```

In this example, the `nodeend.sh` trigger will be called with the name of each node in the `requestid` group.

The node end trigger notifies the external service that this node (along with all the other nodes with the same `requestid`) has met the node idle purge time. The external service may then choose to remove the node from Torque (which in turn removes it from Moab).

The following is an example of the command that a service would run to remove a node from Torque.

```
qmgr -c 'delete node node01'
```

8. If you want to set limits on whether bursting is available, specify the limits using the usage policies. You can set these limits at the global partition or QoS level. See [16.7 Usage Policies - page 803](#).

## 16.2.2 Sample moab.cfg File Excerpt

```

NODECFG[DEFAULT] ENABLEPROFILING=TRUE
SCHEDCFG[moab] FLAGS=enabledynamicnodes
QOSCFG[xyz] REQUESTGEOMETRY=12@4:00:00:00
QOSCFG[xyz] TRIGGER=EType=threshold,AType=exec,Action="$HOME/tools/elastic.py
$REQUESTGEOMETRY",Threshold=BACKLOGCOMPLETIONTIME>1800,RearmTime=05:00
NODEIDLEPURGETIME 3600
NODECFG[DEFAULT]
TRIGGER=EType=end,TType=elastic,AType=exec,Action="/$HOME/tools/nodeend.sh $OID"

```

## 16.3 Elastic Trigger

When enabled, the elastic trigger allows the Moab scheduler to take advantage of systems that can temporarily provide additional nodes to fulfill the backlog in a reasonable time frame.

The elastic trigger is added to `moab.cfg` when the `TType` trigger component is set to 'elastic'. See [17.3.7 Trigger Components - page 824](#) for more information.

When configured and enabled, this trigger:

- Takes the `REQUESTGEOMETRY` parameter and creates nodes in provider.
- Adds nodes to Torque using the create node `qmgr` command (or optionally add it to their RM's cluster query).
- Makes sure that `TTL` is set correctly on the new nodes.
- Optionally adds a request ID (generated by the script) and/or ACL to the nodes.

**i** Elastic Computing scripts should only service one request at a time, it is recommend to not return until the node is up and free in the RM.

### Example:

```

NODECFG[DEFAULT] ENABLEPROFILING=TRUE
QOSCFG[xyz]
TRIGGER=EType=threshold,AType=exec,TType=elastic,Action="$HOME/tools/elastic.py
$REQUESTGEOMETRY",Threshold=BACKLOGCOMPLETIONTIME>1800,RearmTime=05:00

```

**i** The BACKLOGCOMPLETION time trigger threshold can only be used when profiling is enabled.

## 16.4 Integration with a Private OpenStack Cloud

Adaptive Computing has provided a services-enabled integration with a private OpenStack cloud. This consists of the MWS OpenStack plugin, integration scripts, and Moab configuration.

This topic provides instructions to configure your system to send elastic compute requests to OpenStack.

In this topic:

[16.4.1 Configuring Moab to Talk to OpenStack Integration Scripts - page 795](#)

[16.4.2 Verification - page 796](#)

[16.4.3 Troubleshooting - page 797](#)

Parameter	Description
OpenStack Customization Script (cloudinit)	This can be used to perform post-creation actions on the provisioned machines.
OpenStack VLAN Name	Use this if the correct IP address is not returned from the OpenStack plugin after provisioning a machine.
OpenStack Keypair Name	Use this parameter to gain access to the machine after provisioning.

### 16.4.1 Configuring Moab to Talk to OpenStack Integration Scripts

The following configuration snippet shows how to configure Moab to use the OpenStack plugin:

```
NODECFG[DEFAULT] ENABLEPROFILING=TRUE
SCHEDCFG[moab] FLAGS=enabledynamicnodes
QOSCFG[xyz] REQUESTGEOMETRY=12@4:00:00:00

QOSCFG[xyz]
```

```

TRIGGER=EType=threshold,AType=exec,TType=elastic,Action="$HOME/tools/openstack/openstack_elastic.py $REQUESTGEOMETRY",Threshold=BACKLOGCOMPLETIONTIME>1800,RearmTime=05:00
NODEIDLEPURGETIME 3600
NODECFG[DEFAULT]
TRIGGER=EType=end,TType=elastic,AType=exec,Action="/$HOME/tools/openstack/openstack_delete.py $OID"

```

**i** See [16.2 Configuring Elastic Computing - page 791](#) for more information on how to configure elastic triggers and thresholds.

These options are also available on the OpenStack elastic script to fine-tune the configuration:

Option	Sample Value	Description
--acl	user=bob	Set ACLs on the provisioned machines to a specific value.
--ttl-pad	360	By default, the requested TTL is 'padded' by 3 minutes to allow for provisioning time. This parameter can be used to increase or decrease the padding time.

For example, to set an ACL and increase the TTL padding to 4 minutes, use the following elastic trigger definition:

```

QOSCFG[xyz]
TRIGGER=EType=threshold,AType=exec,TType=elastic,Action="$HOME/tools/openstack/openstack_elastic.py --acl=user=alice --ttl-pad=240 $REQUESTGEOMETRY",Threshold=BACKLOGCOMPLETIONTIME>1800,RearmTime=05:00

```

## 16.4.2 Verification

The following methods can be used to verify that the configuration is correct:

- The triggerElastic and triggerNodeEnd web services can be called directly through a browser or command line utility to ensure that the plugin can interact with OpenStack correctly.
- The openstack\_elastic.py and openstack\_delete.py scripts can be called directly through the command line. Make sure to match the parameters used in your Moab configuration file.
- Submit enough workload to cause Moab to fire the elastic trigger and observe that the process works end-to-end.



### 16.4.3 Troubleshooting

The following methods can be used to troubleshoot the OpenStack integration:

- Check the output of `mdiag -T` for information on the configured triggers and to see whether they are firing and results of their execution.
- Check the trigger script log files located in the Moab log directory. By default, these are located at `/opt/moab/log/openstack_elastic.log` and `opt/moab/log/openstack_delete.log`.
- Check the MWS log for information on provisioning and deleting OpenStack machines.
- If the triggers are not firing, check the Moab Workload Manager log files for additional information.

## 16.5 Dynamic Nodes

Dynamic nodes are nodes that can be added and removed from Torque at any time. Specifically, any node that has a TTL (time to live) is considered a dynamic node. The following section explains how to add and delete nodes via `qmgr`.

**i** As of Moab version 9.1.2, dynamic node procs are no longer counted against the total procs listed in the Moab license. This allows you to do as many bursts as you desire without exceeding the total procs used for on-premises nodes. If your version of Moab is before 9.1.2, please contact your Adaptive Computing sales representative.

In this topic:

[16.5.1 Dynamic Node Parameters - page 797](#)

[16.5.2 Dynamic Node Events - page 798](#)

[16.5.3 Configuring Dynamic Nodes - page 799](#)

### 16.5.1 Dynamic Node Parameters

The table below describes the parameters that are used while adding and removing dynamic nodes:

Parameter Name	Required/Optional	Data Format	Description
TTL	Optional	yyyy-mm-ddThh:mm:ss±hh OR yyyy-mm-ddThh:mm:ss±hhmm OR yyyy-mm-ddThh:mm:ssZ	Time, given as a UTC time, for the node to be removed. The time is Greenwich Mean Time with either an offset or a Z to indicate zero offset.
requestid	Optional	Any sequence of non-white-space characters	Identifier used by Moab to identify a group of nodes. See <a href="#">requestid Parameter (Adding or Removing Nodes)</a> for more information.
acl	optional	user==user1:user2,host==host1	List of credentials that can run jobs on this dynamic node.

## 16.5.2 Dynamic Node Events

You can record dynamic node activity using `RECORDEVENTLIST` in the `moab.cfg` using one or both of these events:

- [NODEADD](#)
- [NODEREMOVE](#)

### 16.5.2.A NODEADD

The `NODEADD` event is generated when the RM first reports a new node to Moab.

The following is an example from the `event_XXX` file in the `$MOAB_HOME/stats` directory:

```
16:22:32 1412202152:359437 node nuc2 NODEADD nuc2 STATE=Idle
PARTITION=bdaw ADISK=1 AMEMORY=15193 APROC=4 ASWAP=16717 CDISK=1 CMEMORY=15918 CPROC=4
CSWAP=17442 OS=linux RM=bdaw NODEACCESSPOLICY=SHARED CCLASS=[DevQ] [batch] MSG='Node
'nuc2' was newly reported in the last cluster query. RequestID = 1234, TTL =
1420070400'
```

### 16.5.2.B NODEREMOVE

The `NODEREMOVE` event is generated when Moab removes a dynamic node after `TTL` has expired, or if the node is no longer reported to Moab by the RM.

The following is an example from the `event_XXX` file in the `$MOAB_HOME/stats` directory:

```
16:21:44 1412202104:359401 node      nuc2          NODEREMOVE   nuc2 STATE=Idle
PARTITION=bdaw ADISK=1 AMEMORY=15192 APROC=4 ASWAP=16716 CDISK=1 CMEMORY=15918 CPROC=4
CSWAP=17442 OS=linux RM=bdaw NODEACCESSPOLICY=SHARED FEATURE=[DEV] CCLASS=[DevQ]
[batch] MSG='Dynamic node 'nuc2' is being removed. RequestID = 1234, TTL =
1420070400, Reason = node removed because the RM did not report it in the cluster
query'
```

### 16.5.3 Configuring Dynamic Nodes

This section contains information on configuration options when adding or removing nodes:

- [TTL Parameter \(Creating Nodes\)](#)
- [requestid Parameter \(Adding or Removing Nodes\)](#)
- [NODEIDLEPURGETIME Parameter \(Removing Nodes\)](#)

**i** During the creation of a dynamic node, the `pbs_server` will attempt to resolve the node name to an IP address. If `pbs_server` is unable to resolve the name, it will not create the node; nor will it retry the creation later.

**i** Immediately after a dynamic node is created, it is assigned a state of `'down|MOM-list-not-sent'`. Once the new node has received the list of all moms, it will be assigned a state of `'free'` and be available for job scheduling.

#### 16.5.3.A TTL Parameter (Creating Nodes)

The dynamic nodes are added to the RM with a `TTL` parameter. The `TTL` parameter is passed to Moab by the RM. Moab does not schedule workload for a node beyond the `TTL` assigned to it. Moab removes a dynamic node when it reaches its expiration date as set by `TTL`. A node end trigger will then fire to notify the service that the dynamic node has been removed in Moab and the service may destroy the virtual machine or deprovision the physical nodes at its convenience.

The following is an example of a node being created with a `TTL` parameter:

```
qmgr -c 'create node node003[,node004,node005...] [np=n,][TTL=2022-05-16T05:26:30Z,]
[acl="user==user1:user2:user3",][requestid=n]'
```

In the above example, `node003` is created with `TTL=2022-05-16T05:26:30Z` as the `TTL` parameter. The dynamic node will be removed when the `TTL` is expired.

### 16.5.3.B requestid Parameter (Adding or Removing Nodes)

The dynamic nodes are added to the RM with a `requestid` parameter that is passed to Moab by the RM. Moab reports the `requestid` parameter along with the node ID in Moab logs, events, and node end triggers. This allows the external service to tag the nodes allocated together in a block. The tagged nodes are then associated as events, and are reported on a node-by-node basis by Moab.

The `requestid` can also be used by the external service to de-allocate nodes together in the same block as they were created by the service. For example, a group of nodes has their node end trigger fired due to node idle purge time or `TTL` expiration.

The `requestid` is useful if nodes are dynamically added, removed, and then re-added at some later time with the same node ID. Using a `requestid` when a node is re-added, will help identify each unique instance of a dynamic node's lifetime in logs, events, etc.

Moab also uses the `requestid` with the `NODEIDLEPURGETIME` parameter. The `requestid` parameter groups the nodes and then references the `NODEIDLEPURGETIME` information, if specified, to determine when to remove the group of nodes. When all the nodes associated with the `requestid` have reached the idle purge time threshold defined by the `NODEIDLEPURGETIME` parameter, Moab fires the node end trigger for all the nodes with the same `requestid`.

**i** When `requestid` is configured with `NODEIDLEPURGETIME`, *all* of the nodes must be idle.

### 16.5.3.C NODEIDLEPURGETIME Parameter (Removing Nodes)

The `NODEIDLEPURGETIME` parameter instructs Moab to fire a node end trigger when all the nodes in the `requestid` group have been idle for the time period specified by `NODEIDLEPURGETIME`.

Setting the `NODEIDLEPURGETIME` to 0 effectively disables the `NODEIDLEPURGETIME`. The default value is 0 if `NODEIDLEPURGETIME` is not configured in the `moab.cfg` file. See [NODEIDLEPURGETIME](#) for more information.

The following is an example of configuring the node end trigger in `moab.cfg`:

```
NODECFG[DEFAULT]
TRIGGER=EType=end,TType=elastic,AType=exec,Action="/$HOME/tools/nodeend.sh $OID"
```

In this example, the `nodeend.sh` trigger will be called with the name of each node in the `requestid` group.

The node end trigger notifies the external service that the node (along with all the other nodes with the same `requestid`) has met the node idle purge time set by the `NODEIDLEPURGETIME` parameter. The external service may then choose to remove the node from Torque (which in turn removes it from Moab).

The following is an example of the command that a service will run to remove a node from Torque:

```
qmgr -c 'delete node node003'
```

**i** If a job is running on a node when it is deleted, the job will be requeued if the job is requeueable or deleted if it is not. If the node has already been shut down, any jobs running on the node will be immediately purged.

## 16.6 Viewing Node and Trigger Information

You can optionally configure Elastic Computing to allow you to view the node and trigger information using the Moab commands described in this topic.

### 16.6.1 `mdiag -n -v --xml`

The `mdiag -n -v --xml` command provides detailed information about the state of nodes that Moab is currently tracking. See `mdiag -n` for more information.

In the following example, the `mdiag -n -v --xml` command shows the current list of nodes including dynamic node parameters `TTL` and `REQUESTID` in the XML format:

```
$ mdiag -n -v --xml | xmllint --format -
<?xml version="1.0"?>
<Data>
  <node ACL="USER=%=bdaw+:%=adaptive+;" AVLCLASS="[DevQ] [batch]" CFGCLASS="[DevQ]
[batch]" FEATURES="DEV" LASTUPDATETIME="1412200545" LOAD="0.330000" MAXJOB="0"
MAXJOBPERUSER="0" MAXLOAD="0.000000" NODEID="bdaw" NODEINDEX="0" NODESTATE="Idle"
OS="linux" OSLIST="linux" PARTITION="bdaw" PRIORITY="0" PROCSPEED="0" RADISK="1"
RAMEM="9746" RAPROC="1" RASWAP="26128" RCDISK="1" RCMEM="16050" RCPROC="1"
RCSWAP="32432" REQUESTID="1234" RESCOUNT="1" RMACCESSLIST="bdaw" RSVLIST="bdaw-TTL-
1234" SPEED="1.000000" STATACTIVETIME="2109" STATMODIFYTIME="1412181806"
STATTOTALTIME="2164684" STATUPTIME="2164668" TTL="1441778400" VARATTR="DEV"/>
  <node AVLCLASS="[DevQ] [batch]" CFGCLASS="[DevQ] [batch]" CPUCLOCK="OnDemand:800mhz"
FEATURES="DEV" LASTUPDATETIME="1412200545" MAXJOB="0" MAXJOBPERUSER="0"
MAXLOAD="0.000000" NODEID="nuc2" NODEINDEX="2" NODESTATE="Idle" OS="linux"
OSLIST="linux" PARTITION="bdaw" PRIORITY="0" PROCSPEED="0" RADISK="1" RAMEM="15193"
RAPROC="4" RASWAP="16717" RCDISK="1" RCMEM="15918" RCPROC="4" RCSWAP="17442"
RMACCESSLIST="bdaw" SPEED="1.000000" STATACTIVETIME="34" STATMODIFYTIME="1412114379"
```

```
STATTOTALTIME="86507" STATUPTIME="86475" VARATTR="DEV"/>
</Data>
```

## 16.6.2 mdiag -T

The `mdiag -T` command is used to display information about each trigger. See [mdiag -T](#) for more information.

In the following example, the current list of triggers is displayed using the `mdiag -T` command. Notice the node end triggers associated with nodes.

```
$ mdiag -T
TrigID          Object ID          Event  AType          ActionDate
State
-----
83              node:DEFAULT      end    exec           -
Blocked
85*            node:nuc2         end    exec           -1:00:01:10
Successful
84*            node:bdaw        end    exec           -5:17:23
Active
* indicates trigger has completed
```

## 16.6.3 checknode -v <node name>

The `checknode -v <node name>` command shows detailed state information and statistics including the TTL, the access control list (ACL) and the requestid for nodes that run jobs. See [checknode](#) for more information.

In the following example, a reservation is created on the node at the TTL so that the jobs are not scheduled on the node beyond the TTL. Also, a node end trigger is configured on this node that will fire when the node is removed.

```
$ checknode -v bdaw
node bdaw

State:      Idle (in current state for 5:18:38)
Configured Resources: PROCS: 1 MEM: 15G SWAP: 31G DISK: 1M
Utilized Resources: MEM: 6230M SWAP: 6230M
Dedicated Resources: ---
Attributes: DEV
ACL:        USER==bdaw+::==adaptive+
           MTBF(longterm): 1:00:31:02 MTBF(24h): INFINITY
Opsys:      linux Arch: ---
Speed:      1.00 CPULoad: 0.340
Partition:  bdaw Rack/Slot: ---
Features:   DEV
IdleTime:   23:38:11
Classes:    [DevQ][batch]
RM[bdaw]*:  TYPE=PBS
EffNodeAccessPolicy: SHARED
RequestID:  1234
```

```

TTL: Wed Sep 9 00:00:00 2022

Total Time: 25:01:24:09 Up: 25:01:23:53 (100.00%) Active: 00:35:09 (0.10%)

Reservations:
  bdaw-TTL-1234x1 User      342days -> INFINITY ( INFINITY)
    Blocked Resources@ 342days Procs: 1/1 (100.00%) Mem: 16050/16050 (100.00%)
Swap: 32432/32432 (100.00%) Disk: 1/1 (100.00%)
TrigID          Object ID          Event  AType          ActionDate
State
-----
84*              node:bdaw                end    exec    Wed Oct 1 10:43:26
Active
Launch Time:    -00:00:14
Flags:          globaltrig
Last Execution State: Active (ExitCode: 0)
BlockUntil:    5:18:24 ActiveTime: -1:00:29:57
PID:           7088
Action Data:   /home/bdaw/nodeend.sh $OID
StdOut:       /opt/moab/spool/nodeend.sh.oMnNwKU
StdErr:       /opt/moab/spool/nodeend.sh.enNubAp

* indicates trigger has completed

```

## 16.7 Usage Policies

As part of your Elastic Computing solution, you can keep track of processor seconds on all dynamic nodes to limit over-bursting. For example, if your configuration allows 1000 processor seconds of use every day, then if a job needs to burst (and the used processor seconds reaches 1000 before the job can burst), the trigger to burst the job will not fire, and an error message is generated. You can view the error message using 'mdiag -T -v'.

### 16.7.1 Available Policies

There are four different values you can set: day, month, quarter, or year. The second count resets at the beginning of each period. For ease of use, you can choose to set the limits based on processor hours, and the system will automatically convert the hours to seconds.

These are the available policies you can set in the moab.cfg file to limit over-bursting:

- To specify by processor seconds, use:
  - MAXDAILYELASTICPROCSECONDS
  - MAXMONTHLYELASTICPROCSECONDS

- MAXQUARTERLYELASTICPROCSECONDS
- MAXYEARLYELASTICPROCSECONDS
- To specify by processor hours, use:
  - MAXDAILYELASTICPROCHOURS
  - MAXMONTHLYELASTICPROCHOURS
  - MAXQUARTERLYELASTICPROCHOURS
  - MAXYEARLYELASTICPROCHOURS

## 16.7.2 Policy Levels

You can set the usage policies at the global partition or QoS level:

- Global Partition – Once the elastic node first appears, Moab will begin keeping track of its processor seconds or hours. If the processor seconds reaches the limit, it will not fire off the elastic trigger so no new nodes will come in. For example:

```
PARCFG [ALL]                MAXDAILYELASTICPROCSECONDS=1000
```

You can view the used and remaining limits using 'showstats -v'.

- QoS – Processor seconds or hours start being counted once a job is submitting using that particular QoS, not from when the node first appears. For example:

```
QOSCFG [HIGH]                MAXDAILYELASTICPROCSECONDS=500
```

*A job is submitted requesting the "HIGH" QOS; the processor seconds begin ticking up for that QOS.*

You can view the used and remaining limits using 'mdiag -q -v'.



## Chapter 17: Object Triggers

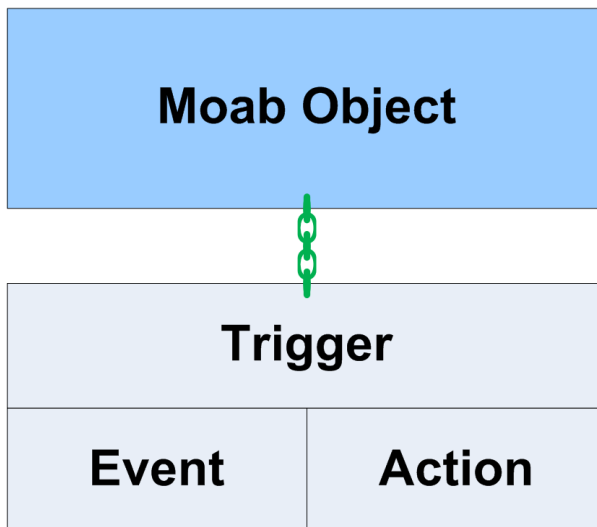
In this chapter:

- [17.1 About Object Triggers - page 805](#)
- [17.2 Object Trigger Tasks - page 806](#)
- [17.3 Object Trigger Reference - page 815](#)
- [17.4 About Trigger Variables - page 836](#)

### 17.1 About Object Triggers

Moab triggers are configurable actions that respond to an event occurring on a Moab object. A trigger is attached to an object and consists of both an event that can take place on the object and the action that the trigger will take.

*Image 17-1: Trigger attachment*



**i** Triggers are a powerful tool. Extreme caution should be taken when using them. They are useful in creating automatic responses to well-understood Moab events; however, by default triggers run as root and do exactly as they are told, meaning they require great thought and consideration to ensure that they act appropriately in response to the event.

**Use Case**

An administrator wants to create the following setup in Moab:

When a node's temperature exceeds 34°C, Moab reserves it. If the temperature increases to more than 40°C, Moab requeues all jobs on the node. If the node's temperature exceeds 50°C, Moab shuts it down. Moab removes the node's reservation and unsets the variables when the node cools to less than 25°C.

The administrator wants to receive an email whenever any of these events occur. All of this can be configured in Moab using triggers. To see a full example for this use case, see [Node Maintenance Example](#).

## 17.2 Object Trigger Tasks

In this section:

- [17.2.1 Creating a Trigger - page 806](#)
- [17.2.2 Using a Trigger to Send Email - page 809](#)
- [17.2.3 Using a Trigger to Execute a Script - page 810](#)
- [17.2.4 Using a Trigger to Perform Internal Moab Actions - page 811](#)
- [17.2.5 Requiring an Object Threshold for Trigger Execution - page 812](#)
- [17.2.6 Enabling Job Triggers - page 812](#)
- [17.2.7 Modifying a Trigger - page 813](#)
- [17.2.8 Viewing a Trigger - page 814](#)
- [17.2.9 Checkpointing a Trigger - page 815](#)

### 17.2.1 Creating a Trigger

Three methods exist for attaching a trigger to an object:

- Directly to the object via the command line
- Directly to the object via the configuration file
- As part of a template via the configuration file

`<attr>=<val>` pair delimiters, quotation marks, and other elements of the syntax might differ slightly from one method/object combination to another, but creating any trigger follows the same basic format:

```
<attr>=<val>[[{&, }<attr>=<val>]...]
```

The beginning of the trigger is set off by the keyword *trigger*. It is followed by a delimited list (typically by commas) of `<attr>=<val>` pairs.

Each method of trigger creation can only be used for certain Moab objects. The following table displays which objects can receive triggers via each method. The links contain examples.

Method	Objects
<b>Command line</b>	job, reservation; a trigger can be attached to any existing object using <a href="#">mschedctl -c</a>
<b>Configuration file</b>	node, reservation, RM, scheduler
<b>Template</b>	job, reservation

Triggers are composed of attributes. Only three are required for each trigger: an [EType](#) (event type), an [AType](#) (action type), and an [Action](#).

*Image 17-2: Required trigger attributes*

## Required Trigger Attributes

<b>EType=start</b>	<b>AType=exec</b>	<b>Action="/tmp/report.pl"</b>
<b>EType</b>	<b>Event Type - The event that fires the trigger</b>	
<b>AType</b>	<b>Action Type – The type of action the trigger will perform</b>	
<b>Action</b>	<b>Action – The actual action the trigger will perform</b>	

Other attributes exist to further customize triggers. See [Trigger Components](#) for more information.

## To Create a Moab Trigger

1. Choose an object to which, and a method by which, you will attach the trigger. Use the format and examples described in its corresponding documentation:

- [Job Triggers](#)
- [Node Triggers](#)
- [Reservation Triggers](#)
- [Resource Manager Triggers](#)
- [Scheduler Triggers](#)

**i** If the trigger is to be attached to a job, you must first enable job triggers (see [Enabling Job Triggers](#) for more information). Carefully review the warning before doing so.

2. Decide whether to attach the trigger via the command line or configuration file. Verify the correct syntax.
3. Set the `EType` equal to whichever event will launch the trigger if and when it occurs on the object.

Each object has a different lifecycle, so not every event type will occur on every object. For a list of valid `ETypes` for your selected object, see the corresponding object reference page linked in step 1.

- a. To modify the timing of the trigger in any of the following ways, see [Event-Modifying Trigger Components](#):
  - To set the trigger as rearmable and specify the amount of time the trigger must wait before firing again.
  - To set an amount of time before or after the event that the trigger will fire (see [Offset](#) for restrictions).
  - To set a specific threshold and the amount of time that the object must meet that threshold before the trigger will fire.
4. Configure the action that the trigger will take when the event happens. To do so, you must set the `AType` to a valid value for your object and specify the action. For instance, to execute a script, set the `AType` to `exec` and the `Action` to the location of the script in quotation marks. Include the name of the object on which the script will run.

```
NODECFG[node01] TRIGGER=EType=fail,AType=exec,Action="node.fail.sh node01"
```

- a. To modify the action in any of the following ways, see [Action-Modifying Trigger Components](#):

- To specify environment variables available to the trigger
  - To set a flag on the trigger:
    - To attach any stderr output generated by the trigger to the parent object
    - To destroy the trigger if its object ends or cancels
    - To tell Moab to checkpoint the trigger
    - To set the trigger as periodic
    - To pass the object's XML information to the trigger's stdin
    - To set the trigger to reset if its object is modified
    - To set the trigger to fire under the user ID of the object's owner
  - To specify an amount of time that Moab will suspend normal operation to wait for the trigger to execute
  - To allot an amount of time that the trigger will attempt to run before it is marked as unsuccessful and the process, if any exists, is killed
  - Set a maximum number of times that a trigger will attempt to fire before it fails
- b. To give the trigger a name or description, see [Organizational Trigger Components](#).
- c. To configure the trigger to set or unset a variable when it fires or to require a variable to fire, see [Setting and Receiving Trigger Variables](#).

## 17.2.2 Using a Trigger to Send Email

Mail triggers can be attached to nodes, jobs, reservations, and the scheduler. The recipient of the email depends on the object to which the trigger is attached. To select different recipient(s) and add flexibility to formatting, send email via a script [using an exec trigger](#).

### To Use a Trigger to Send Email

1. For objects that send mail to the primary user, you must configure the Moab administrator email using the `MAILPROGRAM` parameter.
2. Create a trigger on one of the four valid objects listed below, setting the `AType` to `mail` and the `Action` to the body of the message inside of quotation marks.

Object	Recipient
<b>Node</b>	The primary user (the first user listed in <a href="#">ADMINCFG[1]</a> , typically root)
<b>Job</b>	The job's owner
<b>Reservation</b>	The primary user
<b>Scheduler</b>	The primary user

3. When attaching a mail trigger to all objects of a certain type, use internal variables in the `Action` to add information that is specific to an object, such as the ID, owner, time the event occurred, etc. A variable must be preceded by a dollar sign (\$).

Variable	Description
<b>\$OID</b>	Name of the object to which the trigger is attached
<b>\$OTYPE</b>	The type of object to which the trigger is attached
<b>\$TIME</b>	Time the trigger launched
<b>\$HOSTLIST</b>	Hostlist of the trigger's object (jobs and reservations)
<b>\$OWNER</b>	Owner of the trigger's object (jobs and reservations)
<b>\$USER</b>	User (jobs and reservations)

The variable is replaced with the information described above. For example, the following trigger is configured on all nodes:

```
NODECFG[DEFAULT] EType=fail,AType=mail,Action="node $OID failed at $TIME"
```

*When, for example, node node03 fails, an email is sent to the primary user with a message with the subject line "node node03 started on Sat Aug 18 11:42:00".*

### 17.2.3 Using a Trigger to Execute a Script

Exec triggers launch a program or script when the event occurs. A few examples of what a script might do in response to an event include:

- Execute an external program
- Send a complex email to any desired recipient(s)
- Collect diagnostics

**i** It is important to note that when a script runs via a trigger, Moab forks and performs a direct OS exec, meaning there will be no pre-processing of the command by the shell. In addition, the script runs in a new, reduced environment without the same settings and variables as the environment from which it stemmed. The script must be able to run in the reduced environment.

## To Use a Trigger to Execute a Script

1. Create or locate the script and note its location.
2. Create a trigger on the desired object, setting the `AType` to `exec` and the `Action` to location of the script or program.

```
JOBCFG[temp1] TRIGGER=EType=start,AType=exec,Offset=03:00,Action="/tmp/monitor.pl"
```

*Jobs with the temp1 template receive a trigger that executes monitor.pl three minutes after the job starts.*

## 17.2.4 Using a Trigger to Perform Internal Moab Actions

### To Perform Internal Actions in Moab with a Trigger

Create a trigger on a job, node, or reservation, setting the `AType` to `internal` and the `Action` to one of the following:

- `node:-:reserve` - reserves the node to which the trigger is attached
- `job:-:cancel` - cancels the job to which the trigger is attached
- `reservation:-:cancel` - cancels the reservation to which the trigger is attached

The specified object reserves or cancels itself once the event occurs. See [Internal Action](#) for examples.

## 17.2.5 Requiring an Object Threshold for Trigger Execution

Threshold triggers allow sites to configure triggers to launch based on internal scheduler statistics, such as generic metrics. For example, you might configure a trigger to warn the administrator when the percentage of nodes available is less than 25.

### To Configure a Threshold Trigger

1. Create a trigger. Set its `EType` to `threshold`. Configure the `AType`, `Action`, and `Threshold` attributes' values based on the valid thresholds per object listed in the table found in [Threshold Triggers](#).

```
NODECFG[node04] TRIGGER=EType=threshold,AType=exec,Action="$HOME/hightemp.py
$OID",Threshold=gmetric
```

2. Insert the `gmetric` name between brackets (such as `gmetric[temp]`). Provide a comparison operator. For valid options, see the [comparison operators table](#).
3. Provide a number or string to match against the threshold.

```
NODECFG[node04] TRIGGER=EType=threshold,AType=exec,Action="$HOME/hightemp.py
$OID",Threshold=gmetric[TEMP]>70,RearmTime=5:00
```

*Moab launches a script that warns the administrator when node04's gmetric temp exceeds 70. Moab rearms the trigger five minutes after it fires.*

## 17.2.6 Enabling Job Triggers

By default, common users cannot create most objects, and as a result, common users also cannot create triggers. The exception, however, is jobs. Because common users can create jobs and triggers generally run as root, additional security is necessary to ensure that not all users can create triggers. For this reason, job triggers are disabled by default.



Because triggers generally run as root, any user given the power to attach triggers has the power to run scripts and commands as root. It is recommended that you only enable job triggers on closed systems where human users do not have access to directly submit jobs.

To give specific users permission to create job triggers, you must create a QoS, set the `trigger` flag, and add users to it.



## To Enable Job Triggers

1. In the `moab.cfg` file, create a QoS and set the `trigger` flag.

```
QOSCFG[triggerok] QFLAGS=trigger
```

2. Add users to the QoS who should be allowed to add triggers to jobs.

```
USERCFG[joe] QDEF=triggerok
```

User `joe` is added to the `triggerok` QoS, giving him both the power to create job triggers and root access to the machine.

## 17.2.7 Modifying a Trigger

You can modify a trigger at any time by updating its settings in the Moab configuration file (`moab.cfg`). This will update most triggers at the beginning of the next Moab iteration; however, modifying template triggers (configured using [RSVPROFILE](#) or [JOB\\_CFG](#)) will not update the instances of the trigger that were attached to individual reservations or jobs on creation. The modification will only affect the triggers that the template attaches to future objects.

Any trigger with a specified name can be modified using the `mschedctl -m` command in the following format:

```
mschedctl -m trigger: <triggerID><attr1>=<val1><attr2>=<val2>
```

**i** Modifying triggers on the command line does not change their configuration in `moab.cfg`. Except for reservations that are checkpointed, changes made dynamically are lost when Moab restarts.

For example, the procedure below demonstrates how to modify the following trigger so that the offset is 10 minutes instead of 5 and so that Moab will attempt to fire the trigger up to 10 times if it fails. Assume your trigger currently looks like this:

```
NODECFG[DEFAULT] EType=fail,AType=exec,Action="/scripts/node_fail.pl",Name=nodeFailTrig,Offset=00:05:00,MultiFire=TRUE,RearmTime=01:00:00
```

## To Modify a Trigger

1. Type `mschedctl -m` into the command line and set off the trigger modification with `trigger:<id>`. Use the trigger's assigned ID or specified name to state which trigger will receive the modification.

```
> mschedctl -m trigger:nodeFailTrig
```

2. Type any changing attributes equal to the new value. Separate multiple modifications with a space between each `<attr>=<val>` pair. In this case, set the `Offset` and `MaxRetry` attributes the following way:

```
> mschedctl -m trigger:nodeFailTrig Offset=00:10:00 MaxRetry=10
```

The newly-specified attributes replace the original ones. Trigger `nodeFailTrig` now has an offset of 10 minutes and will try to fire a maximum of 10 times if it fails. The new trigger has the following attributes:

```
EType=fail,AType=exec,Action="/scripts/start_rsv.pl",Name=nodeFailTrig,Offset=00:10:00,MultiFire=TRUE,RearmTime=01:00:00,MaxRetry=10
```

## 17.2.8 Viewing a Trigger

Moab provides a list of triggers when you run the `mdiag -T` command. You can view a specific trigger by running `mdiag -T` in the following format:

```
mdiag -T [<triggerID>|<objectID>|<triggerName>|<objectType>]
```

### To View a Trigger

1. Type `mdiag -T` in the command line.
2. Specify either the trigger ID, the trigger name, the name of the object to which the trigger is attached, or the type of object to which the trigger is attached. For example, if you wanted to view information about a trigger with ID `trigger.34` and name `jobFailTrigger`, which is attached to job `job.493`, you could run any of the following commands:

```
> mdiag -T trigger.34
> mdiag -T job.493
> mdiag -T jobFailTrigger
> mdiag -T job
```

*The output of the first command would provide basic information about `trigger.34`; the second command, information about all triggers attached to `job.493` that the user can access; the third command, basic information about `jobFailTrigger`; and the fourth command, basic information about all triggers attached to jobs that the user can access.*

3. Optional: to view additional information about the trigger, run the same command with the `-v` flag specified after `-T`.

```
> mdiag -T -v job.493
```

This mode outputs information in multiple lines.

4. Optional: to view detailed information about all triggers available to you, use the `mdiag -T -v` command. This outputs all triggers available to the user in a single line for each trigger. It provides additional state information about triggers, including reasons triggers are currently blocked.

```
> mdiag -T -v
```

## 17.2.9 Checkpointing a Trigger

Checkpointing is the process of saving state information when Moab is shut down. In general, triggers defined in the `moab.cfg` file are not checkpointed but are recreated when Moab starts. The exception is the `JOBCFG` parameter, which attaches triggers to jobs as they are created. There are two cases where you might want to tell Moab to checkpoint a trigger:

- If a trigger is defined in the `moab.cfg` file but was created at the command line
- When creating a trigger using the `mschedctl` command

### To Checkpoint a Trigger

1. Locate the trigger to be checkpointed in the `moab.cfg` file, create one on the command line, or modify a trigger dynamically (see [Modifying a Trigger](#) for more information). Attach the `checkpoint` flag using the `FLAGS` attribute. For more information about flags, see [Flags](#).

```
FLAGS=checkpoint
```

2. If you are working in the configuration file, save the changes. Moab will now checkpoint your trigger.

## 17.3 Object Trigger Reference

In this section:

- [17.3.1 Job Triggers - page 816](#)
- [17.3.2 Node Triggers - page 818](#)
- [17.3.3 Reservation Triggers - page 819](#)
- [17.3.4 Resource Manager Triggers - page 821](#)
- [17.3.5 Scheduler Triggers - page 822](#)
- [17.3.6 Threshold Triggers - page 823](#)
- [17.3.7 Trigger Components - page 824](#)
- [17.3.8 Trigger Exit Codes - page 833](#)
- [17.3.9 Node Maintenance Example - page 833](#)
- [17.3.10 Environment Creation Example - page 835](#)

## 17.3.1 Job Triggers

For security reasons, job triggers are disabled by default. They must be enabled in order to successfully attach triggers to jobs (see [Enabling Job Triggers](#) for more information).

Triggers attached to jobs follow the same basic rules and formats as attaching them to other objects; however, not all attribute options are valid for each object. Jobs, like other objects, have a unique set of trigger rules. The table below details the methods, options, and other notable details associated with attaching triggers to jobs.

### 17.3.1.A Creation Methods

Method	Format	Example
<b>Command line on job creation:</b> <b>msub -l</b>	<code>msub &lt;jobName&gt; -l 'trig=&lt;trigSpec&gt;'</code> Attributes are delimited by backslash ampersand (&).	<pre>&gt; msub my.job -l 'trig=EType=create&amp;AType=exec&amp;Action= "/jobs/my_trigger.pl"\&amp;Offset=10:00'</pre>
<b>Command line on existing job:</b> <b>mschedctl -c</b>	<code>mschedctl -c trigger &lt;trigSpec&gt; -o job:&lt;jobID&gt;</code>	<pre>&gt; mschedctl -c trigger EType=end,AType=mail,Action= "Job moab.54 has ended" -o job:moab.54</pre>
<b>Job template in moab.cfg</b>	<code>JOBCFG[&lt;templateName&gt;] TRIGGER=&lt;trigSpec&gt;</code>	<pre>JOBCFG[vmcreate] TRIGGER=,EType=end,AType=exec,Action= "/tmp/jobEnd.sh"</pre>

Method	Format	Example
: <b>JOBCFG</b>		
<b>Class event:</b> <b>CLASSCFG</b>	CLASSCFG [<classID>] JOBTRIGGER=<trigSpec>	<pre>CLASSCFG[batch] JOBTRIGGER=atype=exec,etype=create,action= "/opt/moab/tools/job_trigger.pl"</pre>

### 17.3.1.B Valid Event Types

- [cancel](#)
- [checkpoint](#)
- [create](#)
- [end](#)
- [fail](#)
- [hold](#)
- [modify](#)
- [preempt](#)
- [start](#)

### 17.3.1.C Valid Action Types

- [changeparam](#)
- [exec](#)
- [internal](#)
- [mail](#)

### 17.3.1.D Mail Recipient

The job's owner,

See [Using a Trigger to Send Email](#) for more information.

## 17.3.2 Node Triggers

Triggers attached to nodes follow the same basic rules and formats as attaching them to other objects; however, not all attribute options are valid for each object. Nodes, like the other objects, have a unique set of trigger rules. The table below details the methods, options, and other notable details that come with attaching triggers to nodes.

### 17.3.2.A Creation Methods

Method	Format	Example
<b>Command line on existing node:</b> <b>mschedctl -c</b>	mschedctl -c trigger <trigSpec> -o node: <nodeID>	<pre>&gt; mschedctl -c trigger EType=fail,AType=exec,Action="/tmp/nodeFailure.sh" -o node:node01</pre>
<b>Node configuration in moab.cfg:</b> <b>NODECFG</b>	NODECFG [<name>] TRIGGER= <trigSpec>	<pre>NODECFG[node04] TRIGGER=EType=threshold,AType=exec,Action="\$HOME/hightemp .py \$OID",Threshold=gmetric[TEMP]&gt;70</pre>

### 17.3.2.B Valid Event Types

- create
- discover
- end
- fail
- standing
- threshold

### 17.3.2.C Valid Action Types

- changeparam
- exec

- [internal](#)
- [mail](#)

### 17.3.2.D Thresholds

Node Threshold Settings	
Valid ETypes	threshold
Valid Threshold Types	gmetric

### 17.3.2.E Mail Recipient

The user listed first in [ADMINCFG\[1\]](#) (usually `root`).

See [Using a Trigger to Send Email](#) for more information.

## 17.3.3 Reservation Triggers

Triggers attached to reservations follow the same basic rules and formats as attaching them to other objects; however, not all attribute options are valid for each object. Reservations, like the other objects, have a unique set of trigger rules. The table below details the methods, options, and other notable details that come with attaching triggers to reservations.

### 17.3.3.A Creation Methods

Method	Format	Example
<b>Command line on reservation creation:</b> <a href="#">mrsvctl -T</a>	<pre>mrsvctl - c -h &lt;hostlis t&gt; -T &lt;trigSpe c&gt;</pre>	<pre>&gt; mrsvctl -c -h node01 -T EType=start,AType=exec, Action="/scripts/node_start.pl"</pre>
<b>Command line on</b>	<pre>mschedctl -c</pre>	<pre>&gt; mschedctl -c trigger EType=modify,AType=mail,Action="Reservation system.4 has</pre>

Method	Format	Example
<b>existing reservation:</b> <b>mschedctl -c</b>	trigger <trigSpec> -o rsv: <rsvID>	<pre>been modified" -o rsv:system.4</pre>
<b>Standing reservation configuration in moab.cfg:</b> <b>SRCFG</b>	SRCFG [<name>] TRIGGER= <trigSpec>	<pre>SRCFG[Mail2] TRIGGER=EType=start,Offset=200,AType=exec,Action="/tmp/email.sh"</pre>
<b>Reservation template in moab.cfg:</b> <b>RSVPROFILE</b>	RSVPROFILE E [<name>] TRIGGER= <trigSpec>	<pre>RSVPROFILE[rsvtest] TRIGGER=EType=cancel,AType=exec,Action="\$HOME/logdate.pl TEST CANCEL \$VPCHOSTLIST \$OID \$HOSTLIST \$ACTIVE"</pre>

### 17.3.3.B Valid Event Types

- create
- end
- modify
- standing
- start
- threshold

### 17.3.3.C Valid Action Types

- cancel
- changeparam
- exec
- internal
- jobpreempt
- mail



### 17.3.3.D Thresholds

#### Node Threshold Settings

<b>Valid ETypes</b>	threshold
<b>Valid Threshold types</b>	usage

### 17.3.3.E Mail Recipient

The owner of the reservation. If the owner is unknown or not a user, the first user listed first in [ADMINCFG](#) (usually `root`).

See [Using a Trigger to Send Email](#) for more information.

## 17.3.4 Resource Manager Triggers

Triggers attached to the resource manager follow the same basic rules and formats as attaching them to other objects; however, not all attribute options are valid for each object. The resource manager, like other objects, has a unique set of trigger rules. The table below details the methods, options, and other notable details that come with attaching triggers to RMs.

### 17.3.4.A Creation Methods

Method	Format	Example
<b>Command line on existing RM:</b> <a href="#">mschedctl -c</a>	mschedctl -c trigger <trigSpec> c> -o rm:<rmID>	<pre>&gt; mschedctl -c trigger EType=start,AType=exec,Action="/tmp/rmStart.sh" -o rm:torque</pre>
<b>RM configuration in moab.cfg:</b> <a href="#">RMCFG</a>	RMCFG [<name>] TRIGGER= <trigSpec> c>	<pre>RMCFG[base] TRIGGER=EType=fail,AType=exec,Action="/opt/moab/tools/diagnose_rm.pl \$OID"</pre>

### 17.3.4.B Valid Event Types

- [fail](#)
- [threshold](#)

### 17.3.4.C Valid Action Types

- [changeparam](#)
- [exec](#)
- [internal](#)

## 17.3.5 Scheduler Triggers

Triggers attached to the scheduler follow the same basic rules and formats as attaching them to other objects; however, not all attribute options are valid for each object. The scheduler, like the other objects, has a unique set of trigger rules. The table below details the methods, options, and other notable details associated with attaching triggers to the scheduler.

### 17.3.5.A Creation Methods

Method	Format	Example
<b>Command line on existing scheduler:</b> <a href="#">mschedctl - c</a>	<code>mschedctl -c trigger &lt;trigSpec&gt; -o sched:&lt;schedID&gt;</code>	<pre>&gt; mschedctl -c trigger EType=end,AType=exec,Action="/tmp/startRsvs. sh" -o sched:moab</pre>
<b>Scheduler configuration in moab.cfg</b> : <a href="#">SCHEDCFG</a>	<code>SCHEDCFG [&lt;name&gt;] TRIGGER= &lt;trigSpec&gt;</code>	<pre>SCHEDCFG [MyCluster] TRIGGER=EType=fail,AType=mail,Action="schedu ler failure detected on \$TIME",RearmTime=15:00</pre>

### 17.3.5.B Valid Event Types

- [create](#)
- [end](#)

- [fail](#)
- [modify](#)
- [standing](#)
- [start](#)

### 17.3.5.C Valid Action Types

- [changeparam](#)
- [exec](#)
- [internal](#)
- [mail](#)

### 17.3.5.D Mail Recipient

The user listed first in [ADMINCFG](#) (usually `root`).

See [Using a Trigger to Send Email](#) for more information.

## 17.3.6 Threshold Triggers

The following table identifies the object event, and usage types with which the threshold event/action type feature works:

Object Type	Event Type	Usage Types
<b>Node</b>	Threshold	<code>gmetric</code>
<b>Reservation</b>	Threshold	<code>usage</code>

The following table defines each of the usage types:

Usage Type	Description
<b><code>gmetric</code></b>	Generic performance metrics configured in Moab (see <a href="#">Enabling Generic Metrics</a> for more information).

Usage Type	Description
<b>usage</b>	The percentage of the resource being used (not idle).

The following table defines each of the threshold trigger comparison operators:

Comparison Operator	Value
>	Greater than
>=	Greater than or equal to
<	Less than
<=	Less than or equal to
==	Equal to

### 17.3.6.A Examples

*Example 17-1: Reservation usage threshold*

```
SRCFG[res1] TRIGGER=EType=threshold,AType=mail,Action="More than 75% of reservation
res1 is being used",Threshold=usage>75,FailOffset=1:00
```

*When more than 75% of the reservation has been in use for at least a minute, Moab fires a trigger to notify the primary user.*

## 17.3.7 Trigger Components

In this topic:

[17.3.7.A Required Trigger Components - page 825](#)

[17.3.7.B Event-Modifying Trigger Components - page 829](#)

[Action-Modifying Trigger Components](#)

[17.3.7.D Organizational Trigger Components - page 832](#)

### 17.3.7.A Required Trigger Components

#### AType

Action Type	Description
<b>cancel</b>	Cancels the object.
<b>changeparam</b>	Causes Moab to give a parameter to a new value.
<b>exec</b>	Launches an external program or script on the command line when the dependencies are fulfilled. See <a href="#">Using a Trigger to Execute a Script</a> for more information.
<b>internal</b>	Modifies Moab without using the command line. See <a href="#">Using a Trigger to Perform Internal Moab Actions</a> for more information.
<b>jobpreempt</b>	Indicates the preempt policy to apply to all jobs currently allocated resources assigned to the trigger's parent reservation.
<b>mail</b>	Causes Moab to send mail. See <a href="#">Using a Trigger to Send Email</a> for more information.

#### Action

Cancel Action	
<b>Format</b>	NONE
<b>Description</b>	Indicates that Moab should cancel the reservation when the event occurs. No action should be specified.
<b>Example</b>	<pre>Etype=threshold,Threshold=usage&lt;10,FailOffset=1:00,AType=cancel</pre> <p><i>When less than 10% of the reservation has been in use for a minute, Moab cancels it.</i></p>
Changeparam Action	
<b>Format</b>	Action="<STRING>"

Changeparam Action	
<b>Description</b>	The parameter to change and its new value (using the same syntax and behavior as the <a href="#">changeparam</a> command).
<b>Example</b>	<pre>Atype=changeparam,Action="JOBPURGETIME 02:00:00"</pre> <p><i>Moab maintains detailed job information for two hours after a job has completed.</i></p>

Jobpreempt Action	
<b>Format</b>	Action="cancel checkpoint requeue suspend"
<b>Description</b>	Signifies <a href="#">PREEMTPOLICY</a> to apply to jobs that are running on allocated resources.
<b>Example</b>	<pre>RSVPROFILE[adm1] TRIGGER=EType=start,Offset=-240,Atype=jobpreempt,Action="cancel"</pre> <p><i>40 minutes after the reservation adm1 starts, all jobs using the reservation's resources adopt a <a href="#">PREEMTPOLICY</a> of cancel.</i></p>

Mail Action	
<b>Format</b>	Action="<MESSAGE>"
<b>Description</b>	<p>When <code>Atype=mail</code>, the <code>Action</code> parameter contains the message body of the email. This can be configured to include certain variables. See <a href="#">Using a Trigger to Send Email</a> for details.</p> <p>Mail triggers can be configured to launch for node failures, reservation creation or release, scheduler failures, and even job events. In this way, site administrators can keep track of scheduler events through email.</p> <p>The email comes from <i>moabadmin</i>, has a subject of <i>moab update</i>, and has a body of whatever you specified in the <code>Action</code> attribute. The recipient list depends on the type of object the trigger is attached to.</p> <ul style="list-style-type: none"> <li>• <b>Node</b> - The primary user (first listed in <a href="#">ADMINCFG[1]</a>), typically <code>root</code></li> <li>• <b>Scheduler</b> - The primary user</li> <li>• <b>Job</b> - The user who owns the job</li> <li>• <b>Reservation</b> - The primary user</li> </ul>

Mail Action	
<b>Example</b>	<pre>NODECFG[DEFAULT] TRIGGER=EType=fail,AType=mail,Action="node \$OID will failed .",Offset=05:00:00</pre> <p><i>This example sends an email to the primary administrator informing him/her that the node (including the node ID) has failed.</i></p>

Exec Action	
<b>Format</b>	Action="<script>"
<b>Description</b>	Exec triggers will launch an external program or script when their dependencies are fulfilled. The following example will submit <code>job.cmd</code> and then execute <code>monitor.pl</code> three minutes after the job is started. See <a href="#">Using a Trigger to Execute a Script</a> for more information.
<b>Example</b>	<pre>&gt; msub -l trig=EType=start\&amp;AType=exec\&amp;Action="/tmp/monitor.pl" job.cmd\&amp;Offset=03:00</pre>


Internal Action	
<b>Format</b>	Action="<objectType>:-:<cancel reserve>"
<b>Description</b>	<p>A couple different actions are valid depending on what type of object the internal trigger is acting upon. The following list shows the available actions:</p> <ul style="list-style-type: none"> <li>• Reserve a node</li> <li>• Cancel a job</li> <li>• Cancel a reservation</li> </ul> <p>See <a href="#">Using a Trigger to Perform Internal Moab Actions</a> for more information.</p>
<b>Example</b>	<pre>NODECFG[node01] TRIGGER=EType=start,AType=internal,Action="node:-:reserve"</pre> <p><i>When node01 starts, it becomes a reservation.</i></p> <pre>&gt; msub moab.3 -l 'trig=EType=fail\&amp;AType=internal\&amp;Action="job:-:cancel"</pre> <p><i>If moab.3 fails, Moab cancels it.</i></p> <pre>&gt; mrvctl -c -a user==joe -h node50 -T EType=start,AType=internal,Action=" reservation:-:cancel",Offset=10:00</pre>

Internal Action	
	<i>User joe's jobs are given a ten-minute window to start, then the reservation cancels.</i>

## EType

Event Type	Description
<b>cancel</b>	The event is triggered when the parent object is either canceled or deleted.
<b>checkpoint</b>	Triggers fire when the job is checkpointed. <code>checkpoint</code> triggers can only be attached to jobs.
<b>create</b>	Triggers fire when the parent object is created. <code>create</code> triggers can be attached to nodes, jobs, reservations, classes, and the scheduler (when attached to the scheduler, triggers fire when Moab starts).
<b>discover</b>	Triggers fire when the node is loaded from a resource manager and Moab cannot recognize it nor find it in the checkpoint file.
<b>end</b>	Triggers fire when the parent object ends. <code>end</code> triggers can be attached to nodes, jobs, reservations, and the scheduler. When attached to the scheduler, triggers fire when Moab shuts down. When attached to jobs, triggers fire when jobs return successfully (completion code of 0).
<b>fail</b>	<code>fail</code> triggers can be attached to jobs, nodes, resource managers, and the scheduler. Triggers fire when the resource manager is in a corrupt or down state for longer than the configured fail time, when Moab detects a corruption in a node's reservation table, or when the job returns an unsuccessful completion code.
<b>hold</b>	Triggers fire when the job is put on hold. <code>hold</code> triggers can only be attached to jobs.
<b>modify</b>	Triggers fire when the parent object is modified. <code>modify</code> triggers can be attached to jobs and reservations.
<b>preempt</b>	Triggers fire when the job is preempted. <code>preempt</code> triggers can only be attached to jobs.



Event Type	Description
<b>standing</b>	Triggers fire multiple times based on a certain period. They can be used with <a href="#">Period</a> and <a href="#">Offset</a> attributes. <code>standing</code> triggers can be attached to nodes and the scheduler.
<b>start</b>	Triggers fire when the parent object or Moab starts. <code>start</code> triggers can be attached to jobs, reservations, resource managers, and the scheduler (when Moab starts and at the beginning of Moab's first iteration).
<b>threshold</b>	Triggers fire when a threshold, such as usage or a gmetric comparison, is true. <code>threshold</code> triggers can be attached to nodes and reservations. <div style="border: 1px solid black; padding: 5px; margin-top: 10px;">  Triggers with <code>ETypes</code> set to <code>threshold</code> must include the <a href="#">Threshold</a> attribute. </div>

## TType

Trigger Type	Description
<b>elastic</b>	Specifies trigger for Elastic Computing. Only applicable for QoS triggers.
<b>generic</b>	Default trigger type. Not applicable for Elastic Computing.

### 17.3.7.B Event-Modifying Trigger Components

The following trigger attributes modify the event that causes the trigger to fire.

RearmTime	
<b>Possible Values</b>	[ [HH:]MM:]SS
<b>Description</b>	The amount of time that must pass before a trigger can fire again. <code>RearmTime</code> is enforced from the trigger event time.
<b>Usage Notes</b>	---

Offset	
<b>Possible Values</b>	<code>[-] [[HH:]MM:]SS</code>
<b>Description</b>	The relative time offset from event when trigger can fire.
<b>Usage Notes</b>	<ul style="list-style-type: none"> <li>• Only end triggers can have a negative value for Offset.</li> <li>• Offset cannot be used with cancel.</li> </ul>

Period	
<b>Possible Values</b>	Minute, Hour, Day, Week, Month, Infinity
<b>Description</b>	The period at which the trigger will regularly fire.
<b>Usage Notes</b>	---

Threshold	
<b>Possible Values</b>	Threshold={<metric>[<metricName>]}{> >= < <= ==}<FLOAT> Where <metric> is one of: <ul style="list-style-type: none"> <li>• <b>gmetric</b></li> <li>• <b>usage</b></li> </ul>
<b>Description</b>	When the object meets, drops below, or increases past the configured Threshold, the trigger will fire.
<b>Usage Notes</b>	Threshold triggers allow sites to configure triggers to launch based on internal scheduler statistics, such as the usage of a reservation.

FailOffset	
<b>Possible Values</b>	<code>[[HH:]MM:]SS</code>
<b>Description</b>	The time that the threshold condition must exist before the trigger fires.
<b>Usage Notes</b>	Use with fail triggers to avoid transient triggers.

### 17.3.7.C Action-Modifying Trigger Components

Flags	
<b>Possible Values</b>	<p>Flags=&lt;flag&gt;[:&lt;flag&gt;] or Flags=[&lt;flag&gt;][ [&lt;flag&gt;]]</p> <p>attacherror - If the trigger outputs anything to stderr, Moab attaches it as a message to the trigger object.</p> <p>cleanup - If the trigger is still running when the parent object completes or is canceled, Moab kills the trigger.</p> <p>checkpoint - Moab always checkpoints this trigger. For more information, see <a href="#">Checkpointing a Trigger</a>.</p> <p>objectxmlstdin - Trigger passes its parent's object XML information into the trigger's stdin. This only works for exec triggers with reservation type parents.</p> <p>removestdfiles - When the trigger is deleted, Moab will remove the stdin, stdout, and stderr files used by the trigger.</p> <p>resetonmodify - The trigger resets if its object is modified, even if RearmTime is not set.</p> <p>user - The trigger executes under the user ID of the object's owner. If the parent object is the scheduler, you can explicitly specify the user using the format user+&lt;username&gt;. For example: Flags=user+john.</p>
<b>Description</b>	Specifies various trigger behaviors and actions.
<b>Usage Notes</b>	<p>When specifying multiple flags, each flag can be delimited by colons (:) or with square brackets; for example:</p> <p>Flags=[user][cleanup] or Flags=user:cleanup</p>

BlockTime	
<b>Possible Values</b>	[ [HH:]MM:]SS
<b>Description</b>	The amount of time Moab will suspend normal operation to wait for trigger execution to finish.
<b>Usage Notes</b>	Use caution; Moab will completely stop normal operation until BlockTime expires.

ExpireTime	
<b>Possible</b>	<INTEGER>

ExpireTime	
<b>Values</b>	
<b>Description</b>	The time at which trigger should be terminated if it has not already been activated.
<b>Usage Notes</b>	---

Timeout	
<b>Possible Values</b>	[+ -] [[HH:]MM:]SS
<b>Description</b>	The time allotted to this trigger before it is marked as unsuccessful and its process (if any) killed.
<b>Usage Notes</b>	---

MaxRetry	
<b>Possible Values</b>	MaxRetry=<INTEGER>
<b>Description</b>	The number of times Action will be attempted before the trigger is designated a failure.
<b>Usage Notes</b>	If Action fails, the trigger will restart immediately (up to MaxRetry times). If it fails more than MaxRetry times, the trigger has failed. This restart ignores FailOffset and RearmTime.

### 17.3.7.D Organizational Trigger Components

Name	
<b>Possible Values</b>	Name=<STRING>
<b>Description</b>	Name of the trigger.

Name	
<b>Usage Notes</b>	Because Moab uses its own internal ID to distinguish triggers, the <code>Name</code> need not be unique. Only the first 16 characters of <code>Name</code> are stored by Moab.
Description	
<b>Possible Values</b>	<code>Description=&lt;STRING&gt;</code>
<b>Description</b>	Description of the trigger.
<b>Usage Notes</b>	---

### 17.3.8 Trigger Exit Codes

By default, Moab considers any non-zero exit code as a failure and marks the trigger as having failed. If a trigger is killed by a signal outside of Moab, Moab treats the signal as the exit code and (in almost all cases) marks the trigger as having failed. Only exec triggers that exit with an exit code of 0 are marked as successful.

### 17.3.9 Node Maintenance Example

#### Example Scenario

An administrator wants to create the following setup in Moab:

When a node's temperature exceeds 34°C, Moab reserves it. If the temperature increases to more than 40°C, Moab requeues all jobs on the node. If the node's temperature exceeds 50°C, Moab shuts it down. Moab removes the node's reservation and unsets the variables when the node cools to less than 25°C. The administrator wants to receive an email whenever any of these events occur.

The first trigger reserves the node when its reported temperature exceeds 34°C. Note that the gmetric name in the trigger must match the name of the configured gmetric exactly, including its case (see [Enabling Generic Metrics](#) for more information).

```
NODECFG[DEFAULT] TRIGGER=Description="ThresholdA",EType=threshold,Threshold=gmetric
```

```
[temp]>34,AType=internal,Action="node--:reserve",RearmTime=30,Offset=2:00,Sets=temp_rsv
```

The administrator wants the trigger to fire any time a node overheats, so it must be rearmable. It also needs to specify that the node must be over 34°C for at least two minutes for Moab to reserve it. If the trigger succeeds, it will set a variable to be received by the next trigger in order to make them sequential.

The administrator wants to know when this trigger has fired, so another trigger will send an email once the first trigger has fired and the `temp_rsv` variable is set. This one does so via a script:

```
NODECFG[DEFAULT] Trigger=Description="Email on
Reservation",EType=start,AType=exec,Action="$TOOLS_DIR/node_temp_emailReserve.pl
$OID",RearmTime=3:00,Requires=temp_rsv
```

The second threshold trigger requeues the node's jobs if the node exceeds 40°C and the `temp_rsv` variable is set. It uses a script to do so. It sets `node_evac` variable when it fires, regardless of whether it succeeds or fails.

```
NODECFG[DEFAULT] Trigger=Description="Threshold B",EType=threshold,Threshold=gmetric
[temp]>40,AType=exec,Action="$TOOLS_DIR/node_evacuate.pl
$OID",RearmTime=3:00,requires=temp_rsv,Sets=node_evac,!node_evac
```

The administrator wants another email to inform him that the node is still overheating and has been evacuated. Another email trigger fires once it receives the `node_evac` variable.

```
NODECFG[DEFAULT] Trigger=Description="Email on
Evacuation",EType=start,AType=exec,Action="$TOOLS_DIR/node_temp_emailEvac.pl
$OID",RearmTime=3:00,Requires=node_evac
```

The third threshold trigger uses a script to shut down the node if the temp gmetric exceeds 50 and the `node_evac` variable is set. It sets a `node_shutdown` variable to be received by the notification email.

```
NODECFG[DEFAULT TRIGGER=Description="Threshold C",EType=threshold,Threshold=gmetric
[temp]>50,AType=exec,Action="$TOOLS_DIR/node_shutdown.pl
$OID",RearmTime=3:00,Requires=node_evac,Sets=node_shutdown

NODECFG[DEFAULT] Trigger=Description="Email on
Shutdown",EType=start,AType=exec,Action="$TOOLS_DIR/node_temp_emailShutdown.pl
$OID",RearmTime=3:00,Requires=node_shutdown
```

The final trigger removes the reservation and unsets the variables once the node's temp gmetric is less than 25.

```
NODECFG[DEFAULT] Trigger=Description="Remove
Reservation",EType=threshold,Threshold=gmetric
[temp]<25,AType=exec,Action="opt/moab/bin/mrsvctl -r
r:$OID",RearmTime=3:00,Requires=temp_rsv,unsets=temp_rsv.node_evac.node_shutdown
```

## 17.3.10 Environment Creation Example

### Example Scenario

An administrator wants to create the following setup in Moab:

If a user requests an environment, she must have the permission of her two managers and the administrator. If all three approve, then the environment builds. The user is sent email messages informing her of the environment's end date in case she would like an extension. These are sent 7, 3, and 1 days prior to the environment's ending.

The administrator wants to require his and the managers' approval of any modifications the user makes to her environment so that it cannot be extended without consent.

The first trigger requests manager and administrator approval in response to the user's environment request. So in the event of a reservation's creation, a script is used to send messages to the administrator and manager. The internal variable `OWNER` is used to indicate to the recipients (via the script) which user is requesting the environment.

```
RSVPROFILE [envSetup] TRIGGER=EType=create,AType=exec,Action="envRequest.sh $OWNER"
```

The managers and administrator use an external program to approve or reject the request. On approval, a variable is sent back to Moab (to the reservation specifically). Once all three variables are set, the environment can start. In this example, the variables are called `approval1`, `approval2`, and `approval3`.

```
RSVPROFILE [envSetup]
TRIGGER=EType=start,AType=exec,Action="buildScript",Requires=approval1.approval2.approval3
```

As it is configured now, the reservation will continue to reserve the requested resources regardless of whether all three approvals are given. So, in case approval is not given, the next trigger cancels the reservation 7 days after its creation if the three variables are not set.

```
RSVPROFILE [envSetup]
TRIGGER=EType=create,Offset=7:00:00,AType=internal,Action="rsv:-:cancel",Requires=!approval1.!approval2.!approval3
```

Every remaining trigger in this series is meant to fire for an approved environment and must require the approval variables. Otherwise these notifications would be sent to users who do not have the environment they requested. The next triggers must be rearmable so that it can fire again if necessary; however, they should be set to just over the amount of time left on the reservation so that it doesn't fire again for the same environment. The notification triggers use the `Offset` attribute to fire at the administrator's requested times (7, 3, and 1 day(s) prior to the environment's end).

```
RSVPROFILE [envSetup] TRIGGER=EType=end,Offset=-
7:00:00,AType=exec,Action="weekNotification.sh",RearmTime=7:00:00:02,Requires=approval
1.approval2.approval3

RSVPROFILE [envSetup] TRIGGER=EType=end,Offset=-
3:00:00,AType=exec,Action="3dayNotification.sh",RearmTime=3:00:00:02,Requires=approval
1.approval2.approval3

RSVPROFILE [envSetup] TRIGGER=EType=end,Offset=-
1:00:00,AType=exec,Action="dayNotification.sh",RearmTime=1:00:00:02,Requires=approval1
.approval2.approval3
```

The next trigger requests administrator and manager approval when the environment is modified. The problem is that the trigger must be rearmable in case of multiple modifications and each time the **RearmTime** is reached, Moab will fire the trigger based on the *first* instance of modification. To resolve this issue, this modification trigger requires a `modify` variable. When the reservation is modified, the `modify` variable is set.

```
RSVPROFILE [envSetup]
TRIGGER=EType=modify,AType=exec,Action="modify.sh",RearmTime=1:00:00,Requires=approval
1.approval2.approval3.!modify,Sets=modify
RSVPROFILE [envSetup]
TRIGGER=EType=modify,AType=exec,Action="modificationRequest.sh",RearmTime=5:00,Require
s=approval1.approval2.approval3.modify,Unsets=modify
```

The final triggers notify the user of the end of the environment.

```
RSVPROFILE [envSetup]
TRIGGER=EType=end,AType=exec,Action="end.sh",Requires=approval1.approval2.approval3
```

The same trigger is repeated for the `cancelEType` in case the environment ends unexpectedly.

```
RSVPROFILE [envSetup]
TRIGGER=EType=cancel,AType=exec,Action="end.sh",Requires=approval1.approval2.approval3
```

## 17.4 About Trigger Variables

Trigger variables are pieces of information that pass from trigger to trigger. They allow triggers to fire based on another trigger's behavior, state, and/or output. A variable can be a required condition for a trigger to fire; for instance, a trigger might be set to launch when a reservation starts, but only if it has received a variable from another trigger indicating that a specific node has started first. Variables give greater flexibility and power to a site administrator who wants to automate certain tasks and system behaviors.

Variables can be used to define under what circumstances the trigger will fire. Many Moab objects have their own variables and each object's variable name space is unique. Triggers can use their own variables or the variables attached to their parent objects. A trigger's



variable name space is limited to itself and its parent object. Variables do not have to be unique across all objects.

In this section:

[17.4.1 Trigger Variable Tasks - page 837](#)

[17.4.2 Trigger Variable Reference - page 841](#)

## 17.4.1 Trigger Variable Tasks

### 17.4.1.A Setting and Receiving Trigger Variables

Following is an example of how comparative dependencies can be expressed when creating a trigger.

#### To Set and Require Variables

1. Create a trigger.

```
EType=start, AType=exec, Action="/tmp/trigger1.sh"
```

2. Use the `Sets` attribute to set a variable if the trigger succeeds. You can precede the variable with `!"` to indicate that the variable should be set if the trigger fails. You can specify more than one variable by separating them with a period.

```
AType=exec, Action="/tmp/trigger1.sh", EType=start, Sets=!Var1.Var2
```

*The trigger sets variable `Var2` when it succeeds and variable `Var1` when it fails.*

3. Set up the recipient trigger(s). Use the `Requires` attribute to receive the variable(s). Note that preceding the variable with `!"` means that the variable must not be set in order for the trigger to fire.

```
AType=exec, Action="/tmp/trigger1.sh", EType=start, Sets=!Var1.Var2
AType=exec, Action="/tmp/trigger2.sh", EType=start, Requires=Var1
AType=exec, Action="/tmp/trigger3.sh", EType=start, Requires=Var2
```

*The second trigger will launch if `Var1` has been set (the first trigger failed), and the third trigger will launch if `Var2` is set (the first trigger succeeded).*

4. Refine the requirement with comparisons.

- a. Use the following format:

```
<varID>[:<type>[:<varVal>]]
```

- b. Change `<varID>` to the variable name.
- c. Use any of the comparisons found on the page in place of `<type>`:
- d. Set the value that the variable will be compared against.

```
AType=exec,Action="/tmp/trigger2.sh",EType=start,Requires=Var1:eq:45
AType=exec,Action="/tmp/trigger3.sh",EType=start,Requires=Var2:ne:failure1
```

*The first trigger fires if Var1 exists and has a value of 45. The second trigger fires if Var2 does not have a string value of failure1.*

### 17.4.1.B Externally Injecting Variables Into Job Triggers

Job triggers are able to see the variables in the job object to which it is attached. This means that, for triggers that are attached to job objects, another method for supplying variables exists. Updating the job object's variables effectively updates the variable for the trigger.

#### To Externally Inject Variables into Job Triggers

Use the `mjobctl -m` command to set a variable to attach to a job.

```
> mjobctl -m var=Flag1=TRUE 1664
```

*The variable Flag1 is set. This will be available to any trigger attached to job 1664.*

### 17.4.1.C Exporting Variables to Parent Objects

#### To Export Variables to Parent Objects

1. When setting a variable, indicate that the variable is to be exported to the parent object by using a caret (^).

```
AType=exec,Action="/tmp/trigger1.sh",EType=start,Sets=Var1.^Var2
Atype=exec,Action="/tmp/trigger2.sh",EType=start,Requires=Var1
AType=exec,Action="/tmp/trigger3.sh",EType=start,Requires=Var2
```

`Var2` is exported to the parent object if the trigger fails. It can be used by job and reservation triggers at the same level or by parent objects.

2. Optional: if running a script, you can set a variable as a string to pass up to the parent object.
  - a. Set the variable to pass up to the parent object with the caret (^). Use the `execAtype` to run a script.

```
AType=exec,Action="/tmp/trigger.sh",EType=start,Sets=^Var1
```

The trigger sets `Var1` when it completes successfully. Because the trigger launches a script, a string value can be set for `Var1`.

- b. Declare the variable's string value on its own line in the trigger stdout.

```
EXITCODE=15
Var1=linux
```

*Var1 has the value of `linux` and is passed up to the parent object. This is useful in workflows where a trigger might depend on the value given by a previous trigger.*

**i** To return multiple variables, simply print out one per line.

#### 17.4.1.D Requiring Variables from Generations of Parent Objects

By default, triggers look for variables to fulfill dependencies in the object to which they are directly attached. If they are attached to a job object, they will also look in the job group, if defined. However, it is not uncommon for objects to have multiple generations of parent objects. If the desired behavior is to search through all parent objects, do the following task.

#### To Require Variables from Generations of Parent Objects

Set the `Requires` attribute in the trigger to the required variable, preceded by a caret (^).

```
EType=start,AType=exec,Action="/tmp/trigger2.sh",Requires=^Var1
```

*The trigger searches through the parent objects in which it resides for the variable `Var1`.*

#### 17.4.1.E Requesting Name Space Variables

#### To Request a Name Space Variable in a Trigger

1. **Configure the trigger.** If it is attached to a generic system job, verify that it meets all .
2. Create an argument list in the `Action` attribute (after the script path and before the closing quotes) and request the desired variable with an asterisk (\*) in place of the name space.

```
...Action="$HOME/myTrig.py $*.IPAddr"...
```

*Each applicable name space variable is added to the argument list in the format `<varName>=<val>`.*

For instance, the example above would cause the script to run the following way:

```
> myTrig.py vc1.IPAddr=/tmp/dir1 vc2.IPAddr=/tmp/dir2 vc4.IPAddr=/tmp/dir3
```

*Any other arguments provided here without name spaces will not change.*

- Filter which name spaces are passed down to a job trigger by setting `trigns` when you submit the job. Its value is a comma-delimited list of the desired name spaces.

```
mbsub -l ... -W x="trigns=vc2,vc4"
```

If the new job is applied to the example in step 2, the script's arguments include `vc2.IPAddr` and `vc4.Addr` and exclude `vc1.IPAddr`. The script runs as follows:

```
> myTrig.py vc2.IPAddr=/tmp/dir1 vc4.IPAddr=/tmp/dir2
```

### 17.4.1.F Generic System Job Trigger Requirements

A generic system job specifies one trigger that must meet all of the following criteria:

- The `EType` is `start`.
- The `AType` is `exec`.
- The `Timeout` attribute is the desired walltime of the job. Moab ignores walltime requests when you submit a generic system job, using the trigger `Timeout` instead.

The trigger fires when the system job begins, and, because the trigger's `Timeout` doubles as the job's walltime, both complete at the same time. The job and trigger have the same completion code.

```
JOBCFG[gen] GENERICSYSJOB=EType=start,AType=exec,Action="$HOME/installVM.py $HOSTLIST",Timeout=1:00:00,Flags=objectxmlstdin
```

*The job template `gen` creates a job with a walltime of 1 hour.*

Sometimes the trigger will set a variable on completion or require a variable to run at all. For information about setting variables, passing them up to parent objects, and requiring variables on parent objects, see [About Trigger Variables](#).

You can attach additional triggers using the `TRIGGER` attribute and delimit them with semicolons.

```
JOBCFG[gen] GENERICSYSJOB=<genericSystemJobTriggerSpecs>
JOBCFG[gen] TRIGGER=<triggerSpecs>;TRIGGER=<triggerSpecs>;TRIGGER=<triggerSpecs>
```

*The job template `gen` creates a job with a walltime of 1 hour.*

## 17.4.2 Trigger Variable Reference

### 17.4.2.A Dependency Trigger Components

Sets	
<b>Possible values</b>	' ' delimited string
<b>Description</b>	Variable values this trigger sets upon success or failure.
<b>Usage notes</b>	Preceding the string with an exclamation mark (!) indicates this variable is set upon trigger failure. Preceding the string with a caret (^) indicates this variable is to be exported to the parent object when the trigger completes and satisfies all its set conditions. Used in conjunction with <a href="#">Requires</a> to create trigger dependencies.

Unsets	
<b>Possible values</b>	' ' delimited string
<b>Description</b>	Variable this trigger destroys upon success or failure.
<b>Usage notes</b>	Preceding the string with an exclamation mark (!) indicates this variable is unset upon trigger failure. Used in conjunction with <a href="#">Requires</a> to create trigger dependencies.

Requires	
<b>Possible values</b>	' ' delimited string
<b>Description</b>	Variables this trigger requires to be set or not set before it will fire.
<b>Usage notes</b>	Preceding the string with an exclamation mark (!) indicates this variable must not be set. Preceding the string with a caret (^) indicates that the variable can come from a parent object (see <a href="#">Requiring Variables from Generations of Parent Objects</a> for more information). Used in conjunction with <a href="#">Sets</a> to create trigger dependencies.

### 17.4.2.B Trigger Variable Comparison Types

The following table describes the valid types of comparisons you can use to express the relationship of a trigger variable to its value:

Type	Comparison	Notes
<b>set</b>	is set (exists)	Default
<b>notset</b>	not set (does not exist)	Same as specifying '!' before a variable
<b>eq</b>	equals	
<b>ne</b>	not equal	
<b>gt</b>	greater than	Integer values only
<b>lt</b>	less than	Integer values only
<b>ge</b>	greater than or equal to	Integer values only
<b>le</b>	less than or equal to	Integer values only

### 17.4.2.C Internal Variables


#### Internal Trigger Variables

Several internal variables are available for use in trigger scripts. These can be accessed using `$(varName)`.

Internal Variables	
<b>COMPLETION CODE</b>	When used in conjunction with the config parameter 'JOBCFG' and a trigger event type of 'end', it provides the trigger with the return code of the job.
<b>ETYPE</b>	The type of event that signals that the trigger can fire. ETYPE values include cancel, checkpoint, create, end, fail, hold, migrate, preempt, standing, start, and threshold.
<b>OID</b>	The name of the object to which the trigger was attached.
<b>OTYPE</b>	The type of object to which the trigger is attached; can be rsv, job, node, or

Internal Variables	
	sched.
<b>OWNERMAIL</b>	A variable that is populated only if the trigger's parent object has a user associated with it and that user has an email address associated with it.
<b>REQUESTGEOMETRY</b>	<p>Passed to the Elastic Computing trigger script as a variable. Action="\$HOME/geometry.pl \$REQUESTGEOMETRY</p> <p>Example:</p> <pre>QOSCFG [sample] TRIGGER=EType=start, AType=exec, TType=elastic, Action="\$HOME/geometry.pl \$REQUESTGEOMETRY", timeout=5:00</pre>
<b>TIME</b>	<p>The time of the trigger launch in the following format: Wed Mar 10 12:35:12 2022</p>
<b>USER</b>	The user (when applicable).

## Object-Specific Internal Variables

Job Variables	
<b>HOSTLIST</b>	The entire hostlist of the job
<b>JOBID</b>	<p>String that identifies the job</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p> You must also have REQUESTGEOMETRY=PRIORITYJOBSize to get the JOBID.</p> </div>
<b>MASTERHOST</b>	The primary node for the job

Reservation Variables	
<b>HOSTLIST</b>	The entire hostlist for the reservation
<b>OBJECTXML</b>	The XML representation of an object output is the same that is generated by <code>mdiag -r --xml</code>
<b>OS</b>	The operating system on the first node of the reservation

### Reservation Variables

<b>OWNER</b>	The owner of the reservation
--------------	------------------------------

*Example 17-2: Internal variable example*

```
AType=exec,Action="/tmp/trigger.sh $OID $HOSTLIST",EType=start
```

*The object ID (\$OID) and hostlist (\$HOSTLIST) will be passed to /tmp/trigger.sh as command line arguments when the trigger executes the script. The script can then process this information as needed.*



## Chapter 18: Miscellaneous

In this chapter:

18.1 User Feedback Overview .....	845
18.2 Enabling High Availability Features .....	847
18.2.1 High Availability Overview .....	847
18.2.2 Configuring High Availability on a Networked File System .....	848
18.2.3 Confirming High Availability on a Networked File System .....	849
18.2.4 Other High Availability Configuration .....	849
18.3 Malleable Jobs .....	850
18.4 Identity Managers .....	850
18.4.1 Identity Manager Overview .....	851
18.4.2 Basic Configuration .....	851
18.4.3 Importing Credential Fairness Policies .....	854
18.4.4 Identity Manager Data Format .....	854
18.4.5 Identity Manager Conflicts .....	855
18.4.6 Refreshing Identity Manager Data .....	855
18.5 Generic System Jobs .....	856
18.5.1 Creating a Generic System Job .....	856
18.5.2 Workflows Using Job Template Dependencies .....	857

### 18.1 User Feedback Overview

The Feedback facility allows a site administrator to provide job performance information to users at job completion time. When a job completes, the program pointed to by the `FEEDBACKPROGRAM` parameter is called with a number of command line arguments. The site administrator is responsible for creating a program capable of processing and acting upon the contents of the command line. The command line arguments passed are as follows:

1. job id
2. user name
3. user email
4. final job state

5. QoS requested
6. epoch time job was submitted
7. epoch time job started
8. epoch time job completed
9. job XFactor
10. job wallclock limit
11. processors requested
12. memory requested
13. average per task cpu load
14. maximum per task cpu load
15. average per task memory usage
16. maximum per task memory usage
17. messages associated with the job (if none, [NONE])
18. hostlist (comma-delimited)
19. gres requests (<GRES>:<COUNT> [ ,<GRES>:<COUNT> . . . ])

For many sites, the feedback script is useful as a means of letting users know the accuracy of their wallclock limit estimate, as well as the CPU efficiency, and memory usage pattern of their job. The feedback script can be used as a mechanism to do any of the following:

- email users regarding statistics of all completed jobs
- email users only when certain criteria are met (such as "Job 14991 has just completed, which requested 128 MB of memory per task. During execution, it used 253 MB of memory per task potentially conflicting with other jobs. Please improve your resource usage estimates in future jobs.")
- update system databases
- take system actions based on job completion statistics

**i** Some of these fields may be set to zero if the underlying OS/resource manager does not support the necessary data collection.

*Example 18-1:*

```
FEEDBACKPROGRAM /opt/moab/tools/fb.pl
```

## 18.2 Enabling High Availability Features

This topic provides information and instructions for Adaptive Computing's HA solution with failover. If you want to use Linux HA or any other software stack, please contact your account manager.

In this topic:

[18.2.1 High Availability Overview - page 847](#)

[18.2.2 Configuring High Availability on a Networked File System - page 848](#)

[18.2.3 Confirming High Availability on a Networked File System - page 849](#)

[18.2.4 Other High Availability Configuration - page 849](#)

### 18.2.1 High Availability Overview

High availability allows Moab to run on two different machines: a primary and secondary server. The configuration method to achieve this behavior takes advantage of a networked file system to configure two Moab servers with only one operating at a time.

**i** If you use a shared file system for high availability and Moab is configured to use a database, Moab must be an ODBC build, not SQLite.

**i** It is recommended that you define LOGDIR to be a directory that exists on each server, but isn't a part of the NFS share.

When configured to run on a networked file system — any networked file system that supports file locking is supported — the first Moab server that starts locks a particular file. The second Moab server waits on that lock and only begins scheduling when it gains control of the lock on the file. This method achieves near instantaneous turnover between failures and eliminates the need for two Moab servers to synchronize information periodically as the two Moab servers access the same database/checkpoint file.

**i** As Moab uses timestamping in the lock file to implement high availability, the clocks on both servers require synchronization; all machines in a cluster must be synchronized to the same time server.

Moab high availability and Torque high availability operate independently of each other. If a job is submitted with *msub* and the primary Moab server is down, *msub* tries to connect

to the fallback Moab server. Once the job is given to Torque, if Torque can't connect to the primary `pbs_server`, it tries to connect to the fallback `pbs_server`. For example:

A job is submitted with `msub`, but Moab is down on `server01`, so `msub` contacts Moab running on `server02`.

A job is submitted with `msub` and Moab hands it off to Torque, but `pbs_server` is down on `server01`, so `qsub` contacts `pbs_server` running on `server02`.

When you shut down or restart Moab on both servers, you must run the command twice. A single shutdown (`mschedctl -k`) or restart (`mschedctl -R`) command will go to the primary server and kill it, causing the secondary server to fall back and start operating. To kill the secondary server, resubmit the command.

**i** Do not use anything but a plain simple NFS fileshare that is not used by anybody or anything else (i.e., only Moab can use the fileshare).

**i** Do not use any general-purpose NAS, do not use any parallel file system, and do not use company-wide shared infrastructure to set up Moab high availability using 'native' high availability.

**i** When mounting a remote machine to the NFS server using the `/etc/fstab` file, do not use the option `'noatime'`. This option disables access time updates on the file system that Moab relies on to prevent scheduling conflicts between multiple Moab servers.

## 18.2.2 Configuring High Availability on a Networked File System

Because the two Moab servers access the same files, configuration is only required in the `moab.cfg` file. The two hosts that run Moab must be configured with the `SERVER` and `FBSERVER` parameters. File lock is turned on using the `FLAGS=filelockha` flag. Specify the lock file with the `HALOCKFILE` parameter. The following example illustrates a possible configuration:

```
SCHEDCFG [Moab]  SERVER=host1:42559
SCHEDCFG [Moab]  FBSERVER=host2
SCHEDCFG [Moab]  FLAGS=filelockha
SCHEDCFG [Moab]  HALOCKFILE=/opt/moab/.moab_lock
```

Use the `HALOCKUPDATETIME` parameter to specify how frequently the primary server updates the timestamp on the lock file. Use the `HALOCKCHECKTIME` parameter to specify how frequently the secondary server checks the timestamp on the lock file.

```
HALOCKCHECKTIME 9
```

```
HALOCKUPDATETIME 3
```

*In the preceding example, the secondary server checks the lock file for updates every 9 seconds. The `HALOCKUPDATETIME` parameter is set to 3 seconds, permitting the primary server three opportunities to update the timestamp for each time the secondary server checks the timestamp on the lock file.*

**i** `FBSERVER` does not take a port number. The primary server's port is used for both the primary server and the fallback server.

### 18.2.3 Confirming High Availability on a Networked File System

Administrators can run the `mdiag -S -v` command to view which Moab server is currently scheduling and responding to client requests.

### 18.2.4 Other High Availability Configuration

Moab has many features to improve the availability of a cluster beyond the ability to automatically relocate to another execution server. The following table describes some of these features:

Feature	Description
<b>AMCFG[] BACKUPHOST</b>	<p>If using the Moab Accounting Manager, you can enable high availability with the accounting manager by specifying a backup server as in the following example:</p> <pre>AMCFG[mam] BACKUPHOST=headnode2</pre>
<b>FBSERVER</b>	<p>If you are communicating to an HA configuration from the Moab Grid Control, you can use this parameter to point Moab to the fallback server.</p>
<b>JOBACTIONONNODEFAILURE</b>	<p>If a node allocated to an active job fails, it is possible for the job to continue running indefinitely even though the output it produces is of no value. Setting this parameter allows the scheduler to automatically preempt these jobs when a node failure is detected, possibly allowing the job to run elsewhere and also allowing other allocated nodes to be used by other jobs.</p>
<b>SCHEDCFG[] RECOVERYACTION</b>	<p>If a catastrophic failure event occurs (SIGSEGV or SIGKILL signal is triggered), Moab can be configured to automatically restart, trap the failure, ignore the failure, or behave in the default manner for the specified signal. These actions are specified using the values <code>RESTART</code>, <code>TRAP</code>, <code>IGNORE</code>, or <code>DIE</code>, as in the following</p>

Feature	Description
	example: <pre>SCHEDCFG [bas] MODE=NORMAL RECOVERYACTION=RESTART</pre>
<b>SCHEDCFG[] &lt;Failover Trigger Definition&gt;</b>	HA failover trigger. <pre>SCHEDCFG [Moab] TRIGGER=atype=exec,etype=failure,action="/opt/moab/hafailover.sh"</pre>

## 18.3 Malleable Jobs

Malleable jobs are jobs that can be adjusted in terms of resources and duration required, and which allow the scheduler to maximize job responsiveness by selecting a job's resource shape or footprint prior to job execution. Once a job has started, however, its resource footprint is fixed until job completion.

To enable malleable jobs, the underlying resource manager must support dynamic modification of resource requirements prior to execution (i.e., Torque) and the jobs must be submitted using the TRL (task request list) resource manager extension string. With the TRL attribute specified, Moab will attempt to select a start time and resource footprint to minimize job completion time and maximize overall effective system utilization (i.e., `<AverageJobEfficiency> * <AverageSystemUtilization>`).

*Example 18-2:*

With the following job submission, Moab will execute the job in one of the following configurations: 1 node for 1 hour, 2 nodes for 30 minutes, or 4 nodes for 15 minutes.

```
> qsub -l nodes=1,trl=1@3600:2@1800:4@900 testjob.cmd
job 72436.orion submitted
```

## 18.4 Identity Managers

The Moab identity manager interface can be used to coordinate global and local information regarding users, groups, accounts, and classes associated with compute resources. The identity manager interface can also be used to allow Moab to automatically and dynamically create and modify user accounts and credential attributes according to current workload needs.

**i** Only one identity manager can be configured at a time.

In this topic:

- [18.4.1 Identity Manager Overview - page 851](#)
- [18.4.2 Basic Configuration - page 851](#)
- [18.4.3 Importing Credential Fairness Policies - page 854](#)
- [18.4.4 Identity Manager Data Format - page 854](#)
- [18.4.5 Identity Manager Conflicts - page 855](#)
- [18.4.6 Refreshing Identity Manager Data - page 855](#)

## 18.4.1 Identity Manager Overview

Moab allows sites extensive flexibility when it comes to defining credential access, attributes, and relationships. In most cases, use of the [USERCFG](#), [GROUPECFG](#), [ACCOUNTCFG](#), [CLASSCFG](#), and [QOSCFG](#) parameters is adequate to specify the needed configuration. However, in certain cases such as the following, this approach might not be ideal or even adequate:

- Environments with very large user sets
- Environments with very dynamic credential configurations in terms of fairshare targets, priorities, service access constraints, and credential relationships
- Grid environments with external credential mapping information services
- Enterprise environments with fairness policies based on multi-cluster usage

Moab addresses these and similar issues through the use of an identity manager. An identity manager is configured with the [IDCFG](#) parameter and allows Moab to exchange information with an external identity management service. As with Moab resource manager interfaces, this service can be a full commercial package designed for this purpose, or something far simpler such as a web service, text file, or database.

## 18.4.2 Basic Configuration

Configuring an identity manager in basic read-only mode can be accomplished by simply setting the `SERVER` attribute. If Moab is to interact with the identity manager in read/write mode, some additional configuration might be required.

BLOCKCREDLIST	
<b>Format</b>	One or more comma-delimited object types from the following list: <code>acct</code> , <code>group</code> , or <code>user</code>
<b>Details</b>	<p>If specified, Moab will block all jobs associated with credentials not explicitly reported in the most recent identity manager update. If the credential appears on subsequent updates, resource access will be immediately restored.</p> <div style="border: 1px solid #0070C0; padding: 5px;"> <p><b>i</b> Jobs will only be blocked if <a href="#">fairshare</a> is enabled. This can be accomplished by setting the <code>FSPOLICY</code> parameter to any value such as in the following example:</p> <pre style="border: 1px dashed #0070C0; padding: 5px; width: fit-content; margin: 5px auto;">FSPOLICY DEDICATEDPS</pre> </div>
<b>Example</b>	<pre>IDCFG[test01] BLOCKCREDLIST=acct,user,groups</pre> <div style="border: 1px dashed #0070C0; padding: 5px; margin-top: 5px;"> <p><i>Moab will block any jobs associated with accounts, users, or groups not in the most recent identity manager update.</i></p> </div>

CREATECREC	
<b>Format</b>	<BOOLEAN> (default is FALSE)
<b>Details</b>	Specifies whether Moab should create credentials reported by the identity manager that have not yet been locally discovered or loaded via the resource manager. By default, Moab will only load information for credentials that have been discovered outside of the identity manager.
<b>Example</b>	<pre>IDCFG[test01] CREATECREC=TRUE</pre> <div style="border: 1px dashed #0070C0; padding: 5px; margin-top: 5px;"> <p><i>Moab will create credentials from test01 that have not been previously loaded.</i></p> </div>

REFRESHPERIOD	
<b>Format</b>	[[[DD:]HH:]MM:]SS or INFINITY (default is INFINITY)
	<div style="border: 1px solid #0070C0; padding: 5px;"> <p><b>i Note:</b> The former values of MINUTE, HOUR, DAY or NONE are deprecated and may be removed in a future release.</p> </div>
<b>Details</b>	Moab will refresh the identity manager information on the specified period



REFRESHPERIOD	
	relative to the scheduler start time. If INFINITY is specified, the information is updated only at Moab start up.
<b>Example</b>	<pre>IDCFG[test01] REFRESHPERIOD=4:00:00</pre> <p><i>Moab queries the identity manager every four hours.</i></p>

REQUIREDUSERLIST	
<b>Format</b>	One or more comma-delimited object types from the <code>user</code> list.
<b>Details</b>	<p>Lets you dynamically set the user list with a class for jobs.</p> <div style="border: 1px solid #005596; padding: 5px;"> <p><b>i</b> Removing a user from the REQUIREDUSERLIST will not affect the user's running jobs. However, the user's idle jobs will become blocked because the user no longer has access to the class requested.</p> </div>
<b>Example</b>	<pre>IDCFG[test01] CLASS:&lt;name&gt; REQUIREDUSERLIST=&lt;user&gt;</pre>

RESETCREDLIST	
<b>Format</b>	One or more comma-delimited object types from the following list: <code>acct</code> , <code>group</code> , or <code>user</code> .
<b>Details</b>	If specified, Moab will reset the account access list and fairshare cap and target for all credentials of the specified type(s) regardless of whether they are included in the current info manager report. Moab will then load information for the specified credentials.
<b>Example</b>	<pre>IDCFG[test01] RESETCREDLIST=group</pre> <p><i>Moab will reset the account access list and fairshare target for all groups.</i></p>

SERVER	
<b>Format</b>	<URL>

SERVER	
<b>Details</b>	Specifies the protocol/interface to use to contact the identity manager.
<b>Example</b>	<pre>IDCFG[test01] SERVER=exec://\$HOME/example.pl</pre> <p><i>Moab will use <code>example.pl</code> to communicate with the identity manager.</i></p>

UPDATEREFRESHONFAILURE	
<b>Format</b>	<BOOLEAN> (default is FALSE)
<b>Details</b>	When an IDCFG script fails, it retries almost immediately and continuously until it succeeds. When UPDATEREFRESHONFAILURE is set to TRUE, a failed script does not attempt to rerun immediately, but instead follows the specified REFRESHPERIOD schedule. When set to TRUE, UPDATEREFRESHONFAILURE updates the script execution timestamp, even if the script does not end successfully.
<b>Example</b>	<pre>IDCFG[info] SERVER=exec:///home/tshaw/test/1447/bad_script.pl REFRESHPERIOD=hour UPDATEREFRESHONFAILURE=TRUE</pre>

### 18.4.3 Importing Credential Fairness Policies

One common use for an identity manager is to import fairness data from a global external information service. As an example, assume a site needed to coordinate Moab group level fairshare targets with an allocation database that constrains total allocations available to any given group. To enable this, a configuration like the following might be used:

<pre>IDCFG[alloc] SERVER=exec://\$TOOLS DIR/idquery.pl ...</pre>	<p><i>The <code>tools/idquery.pl</code> script could be set up to query a local database and report its results to Moab. Each iteration, Moab will then import this information, adjust its internal configuration, and immediately respect the new fairness policies.</i></p>
--	--

### 18.4.4 Identity Manager Data Format

When an identity manager outputs credential information either through an `exec` or `file` based interface, the data should be organized in the following format:

```
<CREDTYPE>:<CREDID> <ATTR>=<VALUE>
```

Where:

- `<CREDTYPE>` is one of `user`, `group`, `account`, `class`, or `qos`.
- `<CREDID>` is the name of the credential.
- `<ATTR>` is one of `adminlevel`, `alist`, `chargerate`, `comment`, `emailaddress`, `fstarget`, `globalfstarget`, `globalfsusage`, `maxjob`, `maxmem`, `maxnode`, `maxpe`, `maxproc`, `maxps`, `maxwc`, `MAX.WCLIMIT`, `plist`, `priority`, `qlist`, or `role`. **Multi-dimensional policies** work here as well.
- `<VALUE>` is the value for the specified attribute.

**i** To clear a comment, set its value to `""`; for example: `comment=""`.

*Example 18-3:*

The following output may be generated by an `exec` based identity manager:

```
group:financial fstarget=16.3 alist=project2
group:marketing fstarget=2.5
group:engineering fstarget=36.7
group:dm fstarget=42.5
user:jason adminlevel=3
account:sales maxnode=128 maxjob=8,16
```

The following example limits user `bob` to 8 `matlab` generic resources:

```
user:bob MAXGRES[matlab]=8
```

**i** To specify unlimited use of generic resources, set the value to `-1`.

## 18.4.5 Identity Manager Conflicts

When local credential configuration (as specified via `moab.cfg`) conflicts with identity manager configuration, the identity manager value takes precedence and the local values are overwritten.

## 18.4.6 Refreshing Identity Manager Data

By default, Moab only loads identity manager information once when it is first started up. If the identity manager data is dynamic, then you might want Moab to periodically update its information. To do this, set the `REFRESHPERIOD` attribute of the `IDCFG` parameter. Legal values are documented in the following table:

Value	Description
<b>minute</b>	Update identity information once per minute
<b>hour</b>	Update identity information once per hour
<b>day</b>	Update identity information once per day
<b>infinity</b>	Update identity information only at start-up (default)

*Example 18-4:*

```
IDCFG[hq] SERVER=exec://$TOOLSDIR/updatepolicy.sh REFRESHPERIOD=hour
```

 Job credential feasibility is evaluated at job submission and start time.

## Related Topics

- [Credential Overview](#)
- [Usage Limits/Throttling Policies](#)

## 18.5 Generic System Jobs

Generic system jobs are system jobs with a trigger. They are useful for specifying steps in a workflow.

In this topic:

- [18.5.1 Creating a Generic System Job - page 856](#)
  - [18.5.1.A The Trigger - page 857](#)
- [18.5.2 Workflows Using Job Template Dependencies - page 857](#)
  - [18.5.2.A Inheriting Resources in Workflows - page 858](#)

### 18.5.1 Creating a Generic System Job

Generic system jobs are specified via a job template. The template can be selectable and you must use the `GENERICSYSJOB` attribute to let Moab know that this job template describes a generic system job and to specify a trigger, as shown in the following example:

```

JOBCFG[gen]
GENERICSYSJOB=EType=start,AType=exec,Action="$HOME/genericTrig.py",Timeout=5:00

```

### 18.5.1.A The Trigger

The generic system job's trigger that meets certain criteria. This trigger must have a timeout, an `AType=Exec`, and the `EType` must equal 'start'. The timeout of the trigger will be used as the walltime for the job. The trigger will begin when the system job begins and the job will be considered completed when the trigger completes. The job will have the same completion code as the trigger. The walltime on the job template is not applicable in this case since the timeout of the trigger will be the walltime.

If the trigger fails, an error message will be attached to all of the job's parent VCs. You can view this in the `--xml` output of the VC query. The message includes the location of `STDIN`, `STDOUT`, and `STDERR` files. For example:

```

mvcctl -q ALL --xml

<Data>
<vc CREATETIME="1320184350" DESCRIPTION="Moab.1"
  FLAGS="DESTROYOBJECTS,DESTROYWHENEMPTY,HASSTARTED,WORKFLOW"
  JOBS="Moab.1" NAME="vc1" OWNER="user:frank">
<ACL aff="positive" cmp="%" name="frank" type="USER"></ACL>
<MESSAGES>
<message COUNT="1" CTIME="1320184362"
  DATA="Trigger 10 failed on job Moab.1.setup- STDIN:
/tmp/ByLLl2wv/spool/vm.py.ieWPPS5 STDOUT:
/tmp/ByLLl2wv/spool/vm.py.ODMIXAW STDERR /tmp/ByLLl2wv/spool/vm.py.e2jD5iN"
  EXPIRETIME="1322776362" OWNER="frank" PRIORITY="0"
  TYPE="other" index="0"></message>
</MESSAGES>
<Variables>
<Variable name="VMID">vm1</Variable>
<Variable name="HV">TRUE</Variable>
</Variables>
</vc>
</Data>

```

You can specify other triggers on a generic system job using the `TRIGGER` attribute and delimiting them with semicolons. For example:

```

JOBCFG[gen]  GENERICSYSJOB=<genericSystemJobTriggerSpecs>
JOBCFG[gen]  TRIGGER=<triggerSpecs>;TRIGGER=<triggerSpecs>

```

### 18.5.2 Workflows Using Job Template Dependencies

To create workflows, use the following format:

```

JOBCFG[gen]  TEMPLATEDEPEND=AFTERANY:otherTemplate

```

This will create a job based on the template `otherTemplate`. The generic job will run after the `otherTemplate` job has finished. **Afterany** in the example means after all other jobs have completed, regardless of success.

### 18.5.2.A Inheriting Resources in Workflows

The `INHERITRES` flag can be used to cause the same resources in one step of a workflow to be passed to the next step:

```
JOBCFG[gen]  TEMPLATEDEPEND=AFTERANY:otherTemplate
JOBCFG[otherTemplate]  INHERITRES=TRUE
```

*This example forces the job based on `otherTemplate` to have the same resource requirements as its parent. When the `otherTemplate` job is finished, the `INHERITRES` flag will cause the parent to run on the same resources as the child.*

The job that finishes first will pass its allocation up.

Any variables on the original job will be passed to the other jobs in the workflow. Variables can be added by other jobs in the workflow via the `sets` attribute in the generic system job's trigger. Other triggers must then request that variable name in the command line options.

**i** You will need to set the carat (^) in order for the variable to be sent up to the job group.

If you set the variable, you need to set it in the `STDOUT` of the trigger script. See the example below:

```
JOBCFG[W1]  GENERICSYSJOB=...,action='$HOME/W1.py $ipaddress'  TEMPLATEDEPEND=AFTER:W2
JOBCFG[W2]  TRIGGER=...,action='$HOME/W2.py',sets=^ipaddress
```

*If a variable value is not set in `STDOUT`, it will be set to `TRUE`.*

To set the variable to a specific value, the `W2 .py` script must set the value in its `STDOUT`:

```
print "ipaddress=10.10.10.1" #This will be parsed by Moab and set as the value of the
"ipaddress" variable
```

*Example 18-5:*

## Chapter 19: Database Configuration

Moab supports connecting to a database via native SQLite3, and it can also connect to other databases using the ODBC driver. These optional external databases store some additional information that the MongoDB database does not and allow you to query them directly using SQL. These databases are slower, however, and only SQLite3, which does not allow external queries, is supported.

The SQLite3 connection is for storing statistics. Consider reviewing the SQLite web page [Appropriate Uses for SQLite](#) for information regarding the suitability of using SQLite3 on your system.

Connecting to an external database makes Moab more searchable, allowing you to run queries for statistics and events rather than using regular expressions to draw the information from the Moab flat files.

**i** Moab must use an ODBC-compliant database to report statistics with Viewpoint reports.

In this section:

- [19.1 SQLite3 - page 859](#)
- [19.2 Connecting to a MySQL Database with an ODBC Driver - page 860](#)
- [19.3 Connecting to a PostgreSQL Database with an ODBC Driver - page 863](#)
- [19.4 Connecting to an Oracle Database with an ODBC Driver - page 865](#)
- [19.5 Migrating Your Database to Newer Versions of Moab - page 875](#)
- [19.6 Importing Statistics from stats/DAY.\\* to the Moab Database - page 878](#)

### 19.1 SQLite3

Moab supports connecting to a database via native SQLite3. Database installation and configuration occurs automatically during normal Moab installation (configure, make install). If you did not follow the normal process to install Moab and need to install the database, do the following to manually install and configure Moab database support:

1. Create the database file `moab.db` in your moab home directory by running the following command from the root of your unzipped Moab build directory:
 

```
perl buildutils/install.sqlite3.pl <moab-home-directory>
```

- Verify that the command worked by running `lib/sqlite3 <moab-home-directory>/moab.db`; at the resulting prompt, type `.tables` and press ENTER. You should see several tables such as `mcheckpoint` listed. Exit from this program with the `.quit` command.
- The `perl buildutils/install.sqlite3.pl <moab-home-directory>` command could fail if your operating system cannot find the SQLite3 libraries. Also, Moab fails if unable to identify the libraries. To temporarily force the libraries to be found, run the following command:  

```
export LD_LIBRARY_PATH=<location where libraries were copied>
```

2. In the `moab.cfg` file in the `etc/` folder of the home directory, add the following line:

```
USEDATABASE INTERNAL
```

To verify that Moab is running with SQLite3 support, start Moab and run the `mdiag -S -v` command. If there are no database-related error messages displayed, then Moab should be successfully connected to a database.

**i** > `moabd` is a safe and recommended method of starting Moab if things are not installed in their default locations.

## 19.2 Connecting to a MySQL Database with an ODBC Driver

This documentation shows how to set up and configure Moab to connect to a MySQL database using the MySQL ODBC driver. This document assumes the necessary MySQL and ODBC drivers have already been installed and configured.

To set up and configure Moab to connect to a MySQL database using the MySQL ODBC driver, do the following.

**i** This solution has been tested and works with these versions:

- `libmyodbc` - 5.1.5
- MySQL 5.1

1. Download and install Moab. Install and configure Moab as normal but add the following in the Moab configuration file (`moab.cfg`):



```

USEDATABASE          ODBC
# Turn on stat profiling
USERCFG[DEFAULT]     ENABLEPROFILING=TRUE
GROUPCFG[DEFAULT]    ENABLEPROFILING=TRUE
QOSCFG[DEFAULT]      ENABLEPROFILING=TRUE
CLASSCFG[DEFAULT]    ENABLEPROFILING=TRUE
ACCOUNTCFG[DEFAULT]  ENABLEPROFILING=TRUE
NODECFG[DEFAULT]     ENABLEPROFILING=TRUE

```

2. Create the database in MySQL using the MySQL database dump contained in the `moab-db-mysql-create.sql` file. This file is located in the `contrib/sql` directory.

**i** This `contrib/sql` directory is in the expanded tarball directory.

- Run the following command:

```
mysql -u root -p < moab-db-mysql-create.sql
```

3. Configure the MySQL and ODBC driver. The `odbcinst.ini` file must be contained in `/etc`.

**i** Run the following command to find the MySQL ODBC client driver. You could also query the `libmyodbc` package that was installed.

```
[root]# updatedb
[root]# locate libmyodbc
```

```
[MySQL]
Description = ODBC for MySQL
Driver = /usr/lib/odbc/libmyodbc.so
```

4. Configure Moab to use the MySQL ODBC driver. Moab uses an ODBC datastore file to connect to MySQL using ODBC. This file must be located in the Moab home directory (`/opt/moab` by default) and be named `dsninfo.dsn`, which is used by Moab. You need to have the following data in both `/etc/odbc.ini` and `$MOABHOMEDIR/dsninfo.dsn`:

```
[ODBC]
Driver = MySQL
USER = <username>
PASSWORD = <password>
Server = localhost
Database = Moab
Port = 3306
```

**i** The user should have read/write privileges on the Moab database.

The preceding example file tells ODBC to use the MySQL driver, username `<username>`, password `<password>`, and to connect to MySQL running on the

localhost on port 3306. ODBC uses this information to connect to the database called Moab.

5. Test the ODBC to MySQL connection by running the `isql` command, which reads the `/etc/odbc.ini` file:

```

$ isql -v ODBC
+-----+
| Connected! |
| |
| sql-statement |
| help [tablename] |
| quit |
| |
+-----+
SQL> show tables;
+-----+
| Tables_in_Moab |
+-----+
| EventType |
| Events |
| GeneralStats |
| GenericMetrics |
| Moab |
| NodeStats |
| NodeStatsGenericResources |
| ObjectType |
| mcheckpoint |
+-----+
SQLRowCount returns 10
10 rows fetched
SQL>

```

If you encounter any errors using the `isql` command, there was a problem setting up the ODBC to MySQL connection. Try the following debugging steps to resolve the issue:

- a. The `odbcinst.ini` and `odbc.ini` files are usually assumed to be located in `/etc`, but that is not always true. Use the `odbcinst -j` command to determine the assumed location of the files in your configuration.

```

[root#] odbcinst -j
unixODBC 2.2.12
DRIVERS.....: /etc/unixODBC/odbcinst.ini
SYSTEM DATA SOURCES: /etc/unixODBC/odbc.ini
USER DATA SOURCES...: /home/adaptive/.odbc.ini

```

- b. Because `odbcinst.ini` and `odbc.ini` are expected in `/etc/unixODBC`, not `/etc`, move them from `/etc` to `/etc/unixODBC`.
- c. Use the `strace` command to determine where `isql` expects the `odbc.ini` and `odbcinst.ini` files. Note the location where `isql` expects these files.

```

$ strace isql -v ODBC

```

6. With the ODBC driver configured, the database created, and Moab configured to use the database, start Moab for it to begin storing information in the created database.

**i** > `moabd` is a safe and recommended method of starting Moab if things are not installed in their default locations.

## Related Topics

- [Importing Statistics to the Moab Database](#)

## 19.3 Connecting to a PostgreSQL Database with an ODBC Driver

This documentation shows how to set up and configure Moab to connect to a PostgreSQL database using the ODBC driver. This document assumes the necessary ODBC drivers have already been installed and configured.

**i** Occasionally vacuuming your PostgreSQL database could improve Moab performance. See the PostgreSQL documentation for information on how to vacuum your database.

To set up and configure Moab to connect to a PostgreSQL database using the ODBC driver, do the following:

**i** This solution has been tested and works with the following file version:

- `odbc-postgresql - 1:08.03.0200-1.2`

1. Configure the PostgreSQL and ODBC driver. `odbcinst.ini` file must be contained in `/etc`.

**i** Run the following commands to find the PostgreSQL ODBC client driver and setup file. You could also query the `libodbcpsql` package that was installed.

```
[root]# updatedb
[root]# locate psqldbca
[root]# locate libodbcpsql
```

```
[PostgreSQL]
Description = PostgreSQL ODBC driver
Driver = /usr/lib/odbc/psqldbca.so
Setup = /usr/lib/odbc/libodbcpsqlS.so
Debug = 0
CommLog = 1
```

```
UsageCount = 2
```

2. Configure Moab to use the PostgreSQL ODBC driver. Moab uses an ODBC datastore file to connect to PostgreSQL using ODBC. This file must be located in the Moab home directory (`/opt/moab` by default) and be named `dsninfo.dsn`, which is used by Moab. If the following content, which follows the standard ODBC driver file syntax, is not already included in the `/etc/odbc.ini` file, make sure that you include it. Also, include the same content in the `dsninfo.dsn` file.

```
[ODBC]
Driver = PostgreSQL
Description = PostgreSQL Data Source
Servername = localhost
Port = 5432
Protocol = 8.4
UserName = postgres
Password = moab
Database = Moab
```

**i** The user should have read/write privileges on the Moab database.

The preceding example file tells ODBC to use the PostgreSQL driver, `postgres` user, `moab` password, and to connect to PostgreSQL running on the localhost on port 5432. ODBC uses this information and connects to the database called `Moab`.

3. Test the ODBC to PostgreSQL connection by running the `isql` command, which reads the `/etc/odbc.ini` file. If connected, you should be able to run the `help` command.

If you encounter any errors using the `isql` command, there was a problem setting up the ODBC to MySQL connection. Try the following debugging steps to resolve the issue:

- a. The `odbcinst.ini` and `odbc.ini` files are usually assumed to be located in `/etc`, but that is not always true. Use the `odbcinst -j` command to determine the assumed location of the files in your configuration.

```
[root#] odbcinst -j
unixODBC 2.2.12
DRIVERS.....: /etc/unixODBC/odbcinst.ini
SYSTEM DATA SOURCES: /etc/unixODBC/odbc.ini
USER DATA SOURCES...: /home/adaptive/.odbc.ini
```

- b. Because `odbcinst.ini` and `odbc.ini` are expected in `/etc/unixODBC`, not `/etc`, move them from `/etc` to `/etc/unixODBC`.
- c. Use the `strace` command to determine where `isql` expects the `odbc.ini` and `odbcinst.ini` files. Note the location where `isql` expects these files.

```
$ strace isql -v ODBC
```

4. Create the database in PostgreSQL using the `moab-db-postgresql.sh` setup script contained in the `contrib/sql` directory.

**i** This contrib/sql directory is in the expanded tarball directory.

- Run the script and provide the DB username that will attach to the Moab database (you must supply a DB username or the script will exit). The default admin user is `postgres`, but you can make a new user at this time:

```
> ./moab-db-postgresql.sh postgres
Create db user "postgres" in postgresQL? (y/n)>
```

- The script asks if you want to create the DB user you specified in postgresQL. If the DB user already exists, answer 'n'. Otherwise, the DB user is created and it asks for the new user's password.
  - The script then creates the database "Moab".
  - Finally, as the DB user you provided, the script imports the DB schema from `moab-db-postgresql-create.sql` into the Moab database.
5. Download and install Moab. Install and configure Moab as normal but add the following in the Moab configuration file (`moab.cfg`):

```
USEDATABASE          ODBC
# Turn on stat profiling
USERCFG [DEFAULT]    ENABLEPROFILING=TRUE
GROUPCFG [DEFAULT]   ENABLEPROFILING=TRUE
QOSCFG [DEFAULT]     ENABLEPROFILING=TRUE
CLASSCFG [DEFAULT]   ENABLEPROFILING=TRUE
ACCOUNTCFG [DEFAULT] ENABLEPROFILING=TRUE
NODECFG [DEFAULT]    ENABLEPROFILING=TRUE
```

6. With the ODBC driver configured, the database created, and Moab configured to use the database, start Moab for it to begin storing information in the created database.

**i** `> moabd` is a safe and recommended method of starting Moab if things are not installed in their default locations.

## Related Topics

- [Importing Statistics to the Moab Database](#)

## 19.4 Connecting to an Oracle Database with an ODBC Driver

This documentation shows how to set up and configure Moab to connect to an Oracle database using the ODBC driver.

## 19.4.1 To Connect to an Oracle Database with an ODBC Driver

1. Install and configure the Oracle Instant Client with ODBC supporting libraries. For instructions, see [Installing the Oracle Instant Client](#).
2. Open your Moab configuration file (`$MOABHOMEDIR/moab.cfg`) and add the following lines to the end of the file:

```
USEDATABASE ODBC

# Turn on stat profiling
USERCFG[DEFAULT]      ENABLEPROFILING=TRUE
GROUPCFG[DEFAULT]    ENABLEPROFILING=TRUE
QOSCFG[DEFAULT]      ENABLEPROFILING=TRUE
CLASSCFG[DEFAULT]    ENABLEPROFILING=TRUE
ACCOUNTCFG[DEFAULT]  ENABLEPROFILING=TRUE
NODECFG[DEFAULT]     ENABLEPROFILING=TRUE
```

3. Configure the Oracle ODBC Driver. The `odbcinst.ini` file must be contained in `/etc`.

```
[root]# vim /etc/odbcinst.ini
```

**i** Run the following command to find the Oracle Instant Client driver. You could also query the Oracle Instant Client package that was installed.

```
[root]# updatedb && locate libsqora
```

Add the following text to the file:

```
[Oracle 11g ODBC driver]
Description      = Oracle ODBC driver for Oracle 11g
Driver           = /usr/lib/oracle/11.2/client64/lib/libsqora.so.11.1
Setup           =
FileUsage       =
CPTimeout       =
CPReuse         =
Driver Logging   = 7

[ODBC]
Trace = Yes
TraceFile = /tmp/odbc.log
ForceTrace = Yes
Pooling = No
DEBUG = 1
```

**i** Driver Logging is set high (level 7) so that you can debug during the installation and configuration process if necessary. You can decrease the setting or remove the directive once you finish the process.

**i** To configure the location of the ODBC log (`/tmp/odbc.log`), set the `TraceFile` attribute shown in the example above. See [unixODBC without the GUI](#) on the [unixODBC website](#) for more information.

4. Because the driver installed in step 1 is a shared library, run `ldd` to verify that it and all of its dependencies are installed and working:

```
[root]# ldd /usr/lib/oracle/11.2/client64/lib/libsqora.so.11.1
linux-vdso.so.1 => (0x00007fff631ff000)
libdl.so.2 => /lib64/libdl.so.2 (0x00007f8afbe83000)
libm.so.6 => /lib64/libm.so.6 (0x00007f8afbbff000)
libpthread.so.0 => /lib64/libpthread.so.0 (0x00007f8afb9e1000)
libnsl.so.1 => /lib64/libnsl.so.1 (0x00007f8afb7c8000)
libclntsh.so.11.1 =>
/usr/lib/oracle/11.2/client64/lib/libclntsh.so.11.1 (0x00007f8af8e59000)
libodbcinst.so.1 => not found
libc.so.6 => /lib64/libc.so.6 (0x00007f8af8ac5000)
/lib64/ld-linux-x86-64.so.2 (0x0000003bdb000000)
libnnz11.so => /usr/lib/oracle/11.2/client64/lib/libnnz11.so
(0x00007f8af86f8000)
libaio.so.1 => /lib64/libaio.so.1 (0x00007f8af84f6000)
```

**i** If the command returns `libodbcinst.so.1 => not found`, create a symbolic link from `/usr/lib64/libodbcinst.so.1` to `/usr/lib64/libodbcinst.so.2`. This is a known Red Hat issue. See [Red Hat Bugzilla](#) for more information.

```
[root]# locate libodbcinst
/usr/local/lib/libodbcinst.so.2

[root]# cd /usr/lib64
[root]# ln -s libodbcinst.so.2 libodbcinst.so.1
```

Rerun `ldd`. It should load `libsqora.so.11.1` without error, as shown in the `ldd` example above.

**i** If the `ldd` command returns a warning like this: "ldd: warning: you do not have execution permission for `/usr/lib/oracle/11.2/client64/lib/libsqora.so.11.1`", run the following command:

```
[root]# chmod 755 /usr/lib/oracle/11.2/client64/lib/lib*
```

Rerun `ldd`. It should load `libsqora.so.11.1` without error, as shown in the `ldd` example above.

5. Configure Moab to use the Oracle ODBC driver. This example assumes that a Moab user exists and has been granted read and write privileges to the MOAB database instance referred to on the page.

```
[root]# vim $MOABHOMEDIR/dsninfo.dsn
```

Add the following lines the file, but change `ServerName`, `UserName`, and `Password` to suit your own system. `ServerName` is the name of the Oracle database instance. `Username` and `Password` are the credentials used to connect to that instance.

```
[ODBC]
Application Attributes = T
Attributes = W
BatchAutocommitMode = IfAllSuccessful
BindAsFLOAT = F
CloseCursor = F
DisableDPM = F
DisableMTS = T
Driver = Oracle 11g ODBC driver
DSN = ODBC
EXECSchemaOpt =
EXECSyntax = T
Failover = T
FailoverDelay = 10
FailoverRetryCount = 10
FetchBufferSize = 64000
ForceWCHAR = F
Lobs = T
Longs = T
MaxLargeData = 0
MetadataIdDefault = F
QueryTimeout = T
ResultSets = T
ServerName = MOAB
SQLGetData extensions = F
Translation DLL =
Translation Option = 0
DisableRULEHint = T
UserID = moab
Password = moab
StatementCache=F
CacheBufferSize=20
UseOCIDescribeAny=F
MaxTokenSize=8192
```

6. Add the contents of the `dsninfo.dsn` file to `/etc/odbc.ini`. Because the contents of `dsninfo.dsn` are required in both files, use the following command to concatenate the contents of `dsninfo.dsn` to `/etc/odbc.ini`. If the `odbc.ini` file already has content, verify that there are no conflicts.

```
[root]# cat $MOABHOMDIR/dsninfo.dsn >> /etc/odbc.ini
```

7. Create a directory to store the `tnsnames.ora` file you will create in the next step:

```
[root]# mkdir /etc/oracle
```

8. Create the `tnsnames.ora` file. The `ServerName` in `$MOABHOMEDIR/dsninfo.dsn` tells the Oracle ODBC driver what `tnsnames.ora` entry to use (MOAB). The MOAB `tnsnames` entry tells the Oracle ODBC driver to connect to server `adaptive-oracle` on the local domain (`ac`) on port 1561 using TCP and to connect to the Oracle instance named MOAB (the SID is the unique name of



the instance).

```
[root]# cat >/etc/oracle/tnsnames.ora <<EOL
MOAB =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP) (HOST = adaptive-oracle) (PORT = 1561))
    )
    (CONNECT_DATA =
      (SID = MOAB)
    )
  )
EOL
```

9. Create a profile script (`oracle-instant-client.sh`) to be invoked by the operating system at startup. This script will set the `ORACLE_HOME`, `TWO_TASK`, and `TNS_ADMIN` environment variables required by Oracle and will amend the `LD_LIBRARY_PATH` to include required Oracle client libraries in the library search path.

```
[root]# cat >/etc/profile.d/oracle-instant-client.sh <<EOL
# Set ORACLE_HOME to the directory where the bin and lib directories are located
for the oracle client
export ORACLE_HOME=/usr/lib/oracle/11.2/client64

# No need to add ORACLE_HOME to the linker search path. oracle-instant-client.conf
in
# /etc/ld.so.conf.d should already contain /usr/lib/oracle/11.2/client64.
# Alternatively, you can set it here by uncommenting the following line:
# export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$ORACLE_HOME/lib

# Define the default location where Oracle should look for the server
export TWO_TASK=//adaptive-oracle:1561/listener

# Define where to find the tnsnames.ora file
export TNS_ADMIN=/etc/oracle
EOL
```

10. Source the `oracle-instant-client.sh` script and verify that each environment variable is set correctly:

```
[root]# source /etc/profile.d/oracle-instant-client.sh
[root]# echo $ORACLE_HOME
[root]# echo $LD_LIBRARY_PATH
[root]# echo $TWO_TASK
[root]# echo $TNS_ADMIN
```

11. Modify either the Moab startup script (`/etc/init.d/moab`) – recommended – or the `moabd` script (`/opt/moab/sbin/moabd`) to source `oracle-instant-client.sh`.

- Moab startup script (recommended): the following example suggests a location to source the `oracle-instant-client.sh` script within the Moab startup script.

```
...
# Export all environment variables required by the Oracle Instant Client
. /etc/profile.d/oracle-instant-client.sh
```

```
export MOABHOMEDIR=/opt/moab
...
```

- **moabd shell script:** the following example will resemble the `moabd` script in `/opt/moab/sbin`. Note that the `moabd` script is not invoked by the Moab startup script; The Moab startup script invokes the Moab binary (`/opt/moab/sbin/moab`) by default.

```
#!/bin/sh
#
# Copyright (C) 2022 by Adaptive Computing Enterprises, Inc. All Rights
Reserved.
#
# Export all environment variables required by the Oracle Instant Client
. /etc/profile.d/oracle-instant-client.sh
MOABHOMEDIR="/opt/moab" LD_LIBRARY_PATH="/opt/moab/lib:$LD_LIBRARY_PATH" moab
"$@"
```

## 12. Verify the Oracle ODBC driver is working:

```
isql -v ODBC
+-----+
| Connected! |
+-----+
| sql-statement |
| help [tablename] |
| quit |
+-----+
```

If you encounter any errors using the `isql` command, there was a problem setting up the ODBC to Oracle connection. Try the following debugging steps to resolve the issue:

- The `odbcinst.ini` and `odbc.ini` files are usually assumed to be located in `/etc`, but that is not always true. Use the `odbcinst -j` command to determine the assumed location of the files in your configuration.

```
[root#] odbcinst -j
unixODBC 2.2.12
DRIVERS.....: /etc/unixODBC/odbcinst.ini
SYSTEM DATA SOURCES: /etc/unixODBC/odbc.ini
USER DATA SOURCES...: /home/adaptive/.odbc.ini
```

- Because `odbcinst.ini` and `odbc.ini` are expected in `/etc/unixODBC`, not `/etc`, move them from `/etc` to `/etc/unixODBC`.
- Use the `strace` command to determine where `isql` expects the `odbc.ini` and `odbcinst.ini` files. Note the location where `isql` expects these files.

```
$ strace isql -v ODBC
```

13. If you have not already done so, create the database tables in Oracle using the `moab-db-oracle-create.sql` script located in the `contrib/sql` directory.

**i** This `contrib/sql` directory is in the expanded tarball directory.

This example assumes that you are logged into the MOAB database instance (referred to on the page) as Moab user with read and write privileges.

```
SQL> @./contrib/sql/moab-db-oracle-create.sql
```

14. Verify that the database schema installed correctly by listing the tables. Your results should look like this:

```
SQL> select table_name from all_tables where owner = 'MOAB';
+-----+
| TABLE_NAME |
+-----+
| TRIGGERS    |
| MOAB       |
| OBJECTTYPE  |
| VCS        |
| EVENTTYPE   |
| JOBHISTORY  |
| MCHECKPOINT |
| NODES      |
| EVENTS     |
| NODESTATSGENERICRESOURCES |
| JOBS       |
| RESERVATIONS |
| GENERICMETRICS |
| REQUESTS   |
| GENERALSTATS |
| NODESTATS  |
+-----+
SQLRowCount returns -1
16 rows fetched
```

15. Restart Moab:

```
[root]# mschedctl -R
```

16. Verify Moab is correctly configured to write to the Oracle database by doing each of the following steps:

- a. Tail the `moab.log` file for ODBC errors.

```
# Check the $MOABHOMEDIR/log/moab.log file for ODBC errors. You should see a few hits even if there are no errors.
[root]# tail -f $MOABHOMEDIR/log/moab.log | grep -i odbc
```

- b. Log in to the Moab Oracle database.

In the first example below, `isql` will search `/etc/odbc.ini` for "[ODBC]". `unixODBC` will then use the Oracle 11g ODBC driver defined in `/etc/odbcinst.ini` to establish a connection. The `ServerName` in

`/etc/odbc.ini` tells the Oracle driver to reference the MOAB `tnsnames` entry in `/etc/oracle/tnsnames.ora` for connection parameters.

The second example uses `sqlplus` and a connect string to connect.

Try both connection methods.

```
# Log in to Oracle. Try both isql and sqlplus64 clients.
[root]# isql -v ODBC
```

```
[root]# sqlplus64 moab/moab@adaptive-oracle:1561/MOAB
```

- c. Select some data from one or more of the tables (Nodes, Events, and the like) to verify that data is being stored in the Moab Oracle instance:

```
# sqlplus64 moab/moab@adaptive-oracle:1561/MOAB
SQL*Plus: Release 11.2.0.4.0 Production on Fri Oct 4 14:59:02 2013
Copyright (c) 1982, 2013, Oracle. All rights reserved.

Connected to:
Oracle Database 11g Release 11.2.0.1.0 - 64bit Production

SQL> select table_name from user_tables;

TABLE_NAME
-----
JOBS
REQUESTS
RESERVATIONS
VCS
EVENTTYPE
GENERALSTATS
GENERICMETRICS
NODESTATS
NODESTATSGENERICRESOURCES
EVENTS
JOBHISTORY
MCHECKPOINT
NODES
TRIGGERS
MOAB
OBJECTTYPE

16 rows selected.
```

## 19.4.2 Installing the Oracle Instant Client

The following procedure demonstrates how to install the correct ODBC drivers for your Oracle database. This guide is a prerequisite for the [Connecting to an Oracle Database with an ODBC Driver](#) task. Each step must be performed as root.

## To Install the Oracle Instant Client

1. Go to the [Install Client Downloads](#) page on the Oracle website. Choose the link that matches your system type (for instance, Instant Client for Linux x86-64). Choose Accept License Agreement at the top of the page and download the following RPM or zip files for your target version (such as 11.2):

**i** The process of connecting Oracle to Moab Workload Manager has been tested on Oracle Instant Client version 11.2. The process might work with other versions, but they are not supported.

- Basic (oracle-instantclient11.2-basic-11.2.0.4.0-1.x86\_64.rpm)
- SQL Plus (oracle-instantclient11.2-sqlplus-11.2.0.4.0-1.x86\_64.rpm)
- ODBC (oracle-instantclient11.2-odbc-11.2.0.4.0-1.x86\_64.rpm)

2. Install the packages. This example installs the RPMs:

```
[root]# rpm -i ./oracle-instantclient11.2-basic-11.2.0.4.0-1.x86_64.rpm
[root]# rpm -i ./oracle-instantclient11.2-sqlplus-11.2.0.4.0-1.x86_64.rpm
[root]# rpm -i ./oracle-instantclient11.2-odbc-11.2.0.4.0-1.x86_64.rpm
```

3. Create a configuration file in `/etc/ld.so.conf.d` to add the Oracle client libraries to the `LD_LIBRARY_PATH`.

To confirm where the RPMs installed the libraries, run `rpm -qlp <rpmFileName>`.

```
[root]# cat >/etc/ld.so.conf.d/oracle-instant-client.conf <<EOL
/usr/lib/oracle/11.2/client64/lib
EOL
```

**i** If you installed Oracle Instant Client from a repository, run `repoquery -ql <rpmName>` instead.

Rebuild the `LD_LIBRARY_PATH`:

```
[root]# ldconfig
```

4. Connect to the database using `sqlplus`. If you used RPMs to install the client, the 32-bit and 64-bit clients are already in your `PATH`.

```
[root]# sqlplus64 moab/moab@adaptive-oracle:1561/MOAB
SQL*Plus: Release 11.2.0.4.0 Production on Mon Sep 30 14:35:10 2013
Copyright (c) 1982, 2013, Oracle. All rights reserved.
```

```
Connected to:
Oracle Database 11g Release 11.2.0.1.0 - 64bit Production
```

*The 64-bit sqlplus client was used to connect to a 64-bit 11g instance called MOAB, which is hosted on adaptive-oracle.ac.*

5. Verify that you are logged in to the correct database:

```
SQL> select name from v$database
2 ;
NAME
-----
MOAB
```

6. Create the database in Oracle using the `moab-db-oracle-create.sh` script located in the `contrib/sql` directory.

**i** This `contrib/sql` directory is in the expanded tarball directory.

**i** Useful comments are at the top of the script. Read the comments before running the script.

```
[root]# ./moab-db-oracle-create.sh
```

7. Display all of user's MOAB tables:

```
SQL> select table_name from all_tables where owner = 'MOAB';
TABLE_NAME
-----
TRIGGERS
MOAB
OBJECTTYPE
VCS
EVENTTYPE
JOBHISTORY
MCHECKPOINT
NODES
EVENTS
NODESTATSGENERICRESOURCES
JOBS
RESERVATIONS
GENERICMETRICS
REQUESTS
GENERALSTATS
NODESTATS

16 rows selected.
SQL>
```

8. Generate a script to describe all of user's MOAB tables. Cut and paste the following into a terminal that is *not* logged in to SQLPlus:

```
[root]# cat > /tmp/generateDescribe.sql <<EOL
```

```

SET HEADING OFF
SET FEEDBACK OFF
SET ECHO OFF
SET PAGESIZE 0
SPOOL /tmp/describeAllUserTables.sql
select 'desc '||owner||'.'||table_name||';' from all_tables where owner = 'MOAB';
SPOOL OFF
EOL

```

### 9. Run describeAllUserTables.sql:

```
[root]# SQL> start /tmp/describeAllUserTables.sql
```

### Related Topics

- 
- [19.2 Connecting to a MySQL Database with an ODBC Driver - page 860](#)
- [19.3 Connecting to a PostgreSQL Database with an ODBC Driver - page 863](#)
- [Chapter 19: Database Configuration - page 859](#)

## 19.5 Migrating Your Database to Newer Versions of Moab

Sometimes when upgrading from an older version of Moab to a newer version, you must update your database schema. If the schema Moab expects to operate against is different from the actual schema of the database Moab is connected to, Moab might not be able to use the database properly and data might be lost.

When upgrading the Moab database schema from an old version, you must perform each version upgrade in order. **You cannot skip versions.**

**i** If upgrading a Moab version prior to 7.2.0, contact Adaptive Computing.

#### In this section:

<a href="#">19.5.1 Migrate from Moab 9.1 to Moab 10.0</a>	876
<a href="#">19.5.2 Migrate from Moab 9.0 to Moab 9.1</a>	876
<a href="#">19.5.3 Migrate from Moab 8.1 to Moab 9.0</a>	877
<a href="#">19.5.4 Migrate from Moab 8.0 to Moab 8.1</a>	877

19.5.5 Migrate from Moab 7.5 to Moab 8.0 .....	877
19.5.6 Migrate from Moab 7.2.6-7.2.10 to Moab 7.5 .....	878
19.5.7 Migrate from Moab 7.2.0-7.2.5 to Moab 7.2.6 .....	878

## 19.5.1 Migrate from Moab 9.1 to Moab 10.0

In Moab Workload Manager 10.0, some obsolete node attributes were removed (e.g., `ResOvercommitFactor`) from the ODBC database schema. To upgrade your database with these changes, use the `moab-db-<database>-upgrade10_0.sql` file located in the `contrib/sql` directory.

**i** This `contrib/sql` directory is in the expanded tarball directory.

For example, to migrate your PostgreSQL database from the 9.1 schema to the 10.0 schema, run the following:

```
[postgres]$ psql Moab < moab-db-postgresql-upgrade10_0.sql
```

The database name is usually `Moab`.

Similar migration scripts exist for Oracle and MySQL.

## 19.5.2 Migrate from Moab 9.0 to Moab 9.1

In Moab Workload Manager 9.1, the former reservation statistics (`STATCAPS`, `STATCIPS`, `STATTAPS` and `STATTIPS`) were replaced with (`STATCBPS`, `STATCRPS`, `STATTBPS` and `STATTRPS`) in the ODBC database schema. To upgrade your database with these changes, use the `moab-db-<database>-upgrade9_1.sql` file located in the `contrib/sql` directory.

**i** This `contrib/sql` directory is in the expanded tarball directory.

For example, to migrate your PostgreSQL database from the 9.0 schema, run the following:

```
[postgres]$ psql Moab < moab-db-postgresql-upgrade9_1.sql
```

*The database name is usually "Moab".*

Similar migration scripts exist for Oracle and MySQL.



### 19.5.3 Migrate from Moab 8.1 to Moab 9.0

There were no schema changes between Moab 8.1 and Moab 9.0, therefore there is no migration script that needs to be run to adapt your database from Moab 8.1 to Moab 9.0.

### 19.5.4 Migrate from Moab 8.0 to Moab 8.1

In Moab Workload Manager 8.1, a new accounting event 'AMCONTINUE' was added and the datatypes of some reservation statistics were changed. To upgrade your database with these changes, use the `moab-db-<database>-upgrade8_1.sql` file located in the `contrib/sql` directory.

**i** This `contrib/sql` directory is in the expanded tarball directory.

For example, to migrate your MySQL database from the 8.0 (or later) schema, run the following:

```
[root@]# mysql -u root -D <database name> -p < moab-db-mysql-upgrade8_1.sql
```

The database name is usually 'Moab'.

Similar migration scripts exist for Oracle and PostgreSQL.

### 19.5.5 Migrate from Moab 7.5 to Moab 8.0

In Moab Workload Manager 8.0, column names that have become reserved words in newer versions of MySQL, PostgreSQL, and Oracle were renamed to eliminate the need to quote column names in SQL statements. Also, a few additional columns were added to existing tables to support Moab's Green feature. To upgrade your database with these changes, use the `moab-db-<database>-upgrade8_0.sql` file located in the `contrib/sql` directory.

**i** This `contrib/sql` directory is in the expanded tarball directory.

For example, to migrate your MySQL database from the 7.5 (or later) schema, run the following:

```
[root@]# mysql -u root -D <database name> -p < moab-db-mysql-upgrade8_0.sql
```

*The database name is usually "Moab".*

Similar migration scripts exist for Oracle and PostgreSQL.

## 19.5.6 Migrate from Moab 7.2.6-7.2.10 to Moab 7.5

In Moab Workload Manager 7.5, column names that are reserved words in databases supported by Adaptive Computing were renamed to eliminate the need to quote column names in SQL statements. To upgrade your database with these changes, use the `moab-db-<database>-upgrade7_5.sql` file located in the `contrib/sql` directory.

**i** This `contrib/sql` directory is in the expanded tarball directory.

For example, to migrate your MySQL database from the 7.2.6 (or later) schema, run the following:

```
[root@]# mysql -u root -D <database name> -p < moab-db-mysql-upgrade7_5.sql
```

*The database name is usually "Moab".*

Similar migration scripts exist for Oracle and PostgreSQL.

## 19.5.7 Migrate from Moab 7.2.0-7.2.5 to Moab 7.2.6

In Moab Workload Manager 7.2.6, several columns were extended and the primary key on the Triggers table changed. To upgrade your database with these changes, use the `moab-db-<database>-upgrade7_2_6.sql` file located in the `contrib/sql` directory.

**i** This `contrib/sql` directory is in the expanded tarball directory.

For example, to migrate your MySQL database from the 7.2.x (pre-7.2.6) schema to the 7.2.6 schema, run the following:

```
[root@]# mysql -u root -D <database name> -p < moab-db-mysql-upgrade7_2_6.sql
```

*The database name is usually "Moab".*

Similar migration scripts exist for Oracle and PostgreSQL.

**i** The 7.2.6 database upgrade is compatible with all earlier versions of 7.2.

## 19.6 Importing Statistics from stats/DAY.\* to the Moab Database

The `contrib/stat_converter` folder contains the files to build `mstat_converter`, an executable that reads file-based statistics in a Moab `stats` directory and

dumps them into a database. It also reads the Moab checkpoint file (`.moab.ck`) and dumps that to the database as well. It uses the `$MOABHOMEDIR/moab.cfg` file to connect to the appropriate database and reads the statistics files from `$MOABHOMEDIR/stats`.

To run, execute the program `mstat_converter` with no arguments.

The statistics converter program does not clear the database before converting. However, if there are statistics in the database and the statistics files from the same period, the converter reports duplicate data errors.

## Chapter 20: Accelerators

Moab can integrate with the Torque resource manager to discover, report, schedule, and submit workload to various accelerator architectures (such as NVIDIA GPUs or Intel® Xeon Phi™ co-processor architecture) for parallel processing. See the topics below for specific information.

In this chapter:

20.1 Scheduling GPUs .....	880
20.1.1 Deploying and Configuring GPUs .....	881
20.1.2 Using GPUs with Minimal Configuration .....	882
20.2 Using GPUs with NUMA .....	882
20.3 NVIDIA GPUs .....	884
20.3.1 Using NVIDIA GPUs .....	884
20.3.2 Package Installation/Upgrade .....	885
20.3.3 Torque Configuration .....	886
20.3.4 GPU Modes for NVIDIA 260.x Driver .....	887
20.3.5 GPU Modes for NVIDIA 270.x Driver .....	887
20.3.6 gpu_status .....	887
20.3.7 Enabling Persistence Mode .....	888
20.3.8 Requesting GPUs and Setting GPU Mode .....	888
20.4 GPU Metrics .....	889
20.5 Intel® Xeon Phi™ Coprocessor Configuration .....	891
20.5.1 Intel Many-Integrated Cores (MIC) Architecture Configuration .....	892
20.5.2 Validating the Configuration .....	893
20.5.3 Job Submission .....	894
20.6 Intel® Xeon Phi™ Co-processor Metrics .....	896

### 20.1 Scheduling GPUs

In this topic:

[20.1.1 Deploying and Configuring GPUs - page 881](#)  
[20.1.2 Using GPUs with Minimal Configuration - page 882](#)

## 20.1.1 Deploying and Configuring GPUs

There are several ways a site can deploy and configure Moab/Torque systems with GPUs.

### Method 1 (recommended): GPUs with NVIDIA/NVML and cgroups

- The only method that guarantees that users do not use more GPUs than they request.
- Ensures that multiple users will not attempt to use the same GPU.
- Torque will attempt to use CPUs and memory that are close to the GPUs to allow the jobs to execute faster and more consistently.
- Provides all of the information on the GPUs that configuring with NVML does (method 2).
- In addition to configuring NVIDIA GPU support (see [20.3 NVIDIA GPUs - page 884](#)), enabling cgroups requires installing hwloc libraries. See 'Torque NUMA-Aware Configuration' in the *Torque Administrator Guide* for instructions.

### Method 2: GPUs with NVIDIA/NVML, but without cgroups (see [20.3 NVIDIA GPUs - page 884](#))

- Guarantees that you detect an accurate number of GPUs without manual configuration.
- Provides reporting information on the state of each GPU.
- Allows users to set a mode for each GPU in use by their job, as needed.

**i** When using this method, `pbs_server` automatically appends `'gpus=<count>'` to the end of the line in `TORQUE_HOME/server_priv/nodes` for any node with a GPU, overriding any such manual configuration.

### Method 3: Configuring GPUs in the nodes file (see [20.1.2 Using GPUs with Minimal Configuration - page 882](#))

- Allows jobs to request GPUs.
- Requires manual configuration.
- Does not guarantee accuracy.

## 20.1.2 Using GPUs with Minimal Configuration

In Torque 2.5.4 and later, users can request GPUs on a node at job submission by specifying a nodes resource request, using the `qsub -l` option. The number of GPUs a node has must be specified in the nodes file. The GPU is then reported in the output of `pbsnodes`:

```
napali
state = free
np = 2
ntype = cluster
status = rectime=1288888871,varattr=,jobs=,state=free,netload=1606207294,gres=tom:!/home/dbeer/dev/scripts/dynamic_resc.sh,loadave=0.10,ncpus=2,physmem=3091140kb,availmem=32788032348kb,totmem=34653576492kb,idletime=4983,nusers=3,nsessions=14,sessions=3136 1805 2380 2428 1161 3174 3184
3191 3209 3228 3272 3333 20560 32371,uname=Linux napali 2.6.32-25-generic #45-Ubuntu SMP Sat Oct 16 19:52:42 UTC 2021 x86_64,opsys=linux
mom_service_port = 15002
mom_manager_port = 15003
gpus = 1
```

The `$PBS_GPUFILE` has been created to include GPU awareness. The GPU appears as a separate line in `$PBS_GPUFILE` and follows this syntax:

```
<hostname>-gpu<index>
```

If a job were submitted to run on a server called 'napali' (the submit command would look something like: `qsub test.sh -l nodes=1:ppn=2:gpus=1`), the `$PBS_GPUFILE` would contain:

```
napali-gpu0
```

It is left up to the job's owner to make sure that the job executes properly on the GPU. By default, Torque treats GPUs exactly the same as ppn (which corresponds to CPUs).

### Related Topics

- [Using GPUs with NUMA](#)
- [NVIDIA GPUs](#)

## 20.2 Using GPUs with NUMA

The `pbs_server` requires awareness of how the MOM is reporting nodes since there is only one MOM daemon and multiple MOM nodes. Configure the `server_priv/nodes` file

with the `num_node_boards` and `numa_gpu_node_str` attributes. The attribute `num_node_boards` tells `pbs_server` how many NUMA nodes are reported by the MOM. If each NUMA node has the same number of GPUs, add the total number of GPUs to the nodes file. Following is an example of how to configure the nodes file with `num_node_boards`:

```
numahost gpus=12 num_node_boards=6
```

This line in the nodes file tells `pbs_server` there is a host named `numahost` and that it has 12 GPUs and 6 nodes. The `pbs_server` divides the value of GPUs (12) by the value for `num_node_boards` (6) and determines there are 2 GPUs per NUMA node.

In this example, the NUMA system is uniform in its configuration of GPUs per node board, but a system does not have to be configured with the same number of GPUs per node board. For systems with non-uniform GPU distributions, use the attribute `numa_gpu_node_str` to let `pbs_server` know where GPUs are located in the cluster.

If there are equal numbers of GPUs on each NUMA node, you can specify them with a string. For example, if there are 3 NUMA nodes and the first has 0 GPUs, the second has 3, and the third has 5, you would add this to the nodes file entry:

```
numa_gpu_node_str=0,3,5
```

In this configuration, `pbs_server` knows it has three MOM nodes and the nodes have 0, 3s, and 5 GPUs respectively. Note that the attribute `gpus` is not used. The `gpus` attribute is ignored because the number of GPUs per node is specifically given.

In Torque 3.0.2 or later, `qsub` supports the mapping of `-l gpus=X` to `-l gres=gpus:X`. This allows users who are using NUMA systems to make requests such as `-l ncpus=20, gpus=5` (or `-l ncpus=20: gpus=5`) indicating they are not concerned with the GPUs in relation to the NUMA nodes they request; they only want a total of 20 cores and 5 GPUs.

**i** The `qsub -l gpus=X` option is deprecated. It is recommended that you request GPUs using the resource request 2.0 syntax (`-L`, see '`-L NUMA Resource Request`' in the *Torque Resource Manager Administrator Guide*).

---

## Related Topics


- [Scheduling GPUs](#)
- [NVIDIA GPUs](#)


## 20.3 NVIDIA GPUs


In this topic:

- [20.3.1 Using NVIDIA GPUs - page 884](#)
- [20.3.2 Package Installation/Upgrade - page 885](#)
- [20.3.3 Torque Configuration - page 886](#)
- [20.3.4 GPU Modes for NVIDIA 260.x Driver - page 887](#)
- [20.3.5 GPU Modes for NVIDIA 270.x Driver - page 887](#)
- [20.3.6 gpu\\_status - page 887](#)
- [20.3.7 Enabling Persistence Mode - page 888](#)
- [20.3.8 Requesting GPUs and Setting GPU Mode - page 888](#)

### 20.3.1 Using NVIDIA GPUs

 This document assumes that you have installed the NVIDIA CUDA ToolKit and the NVIDIA development drivers on a compute node with an NVIDIA GPU. (Both can be downloaded from <https://developer.nvidia.com/cuda-downloads>).

 Severe scheduling performance problems have been observed in systems with GPUs that do not have persistence mode enabled. We strongly recommend doing this on all GPU nodes. See [20.3.7 Enabling Persistence Mode - page 888](#) for more information.

 CUDA version 6.0 or later is recommended for Torque 6.0 or later. CUDA version 4.1 is the minimum required.

 The recommended method for deploying Moab/Torque systems with GPUs is to include NVIDIA/NVML options and cgroups. See [20.1 Scheduling GPUs - page 880](#).

As of version 2.5.6, `pbs_mom` can query for GPU hardware information and report that status to `pbs_server`, adding a `gpustatus` line in the output for `pbsnodes`.

`qsub` includes options for setting the GPU mode and for resetting GPU ECC error counts.

To generate MOM binaries with GPU support, you must build on a system that has the CUDA libraries for that specific GPU hardware and operating system. Because the CUDA toolkit installer refuses to run on a system without a GPU card (and the server typically lacks that hardware), the usual method of building for GPUs involves putting the source on



a node with a GPU and compiling there with NVIDIA/NVML options. This requires `gcc`, `libtool`, and other build utilities, so first you must follow the 'Install Packages' instructions for the Torque server in the Moab Installation and Configuration Guide for your version and OS on that host. Once you've configured and built, you can generate the MOM installer by running `make packages`, as described below. `pbs_server` can communicate with `pbs_mom` binaries configured for GPU support regardless the server's build options.

To configure for NVIDIA GPU support, include these options:

- `--enable-nvidia-gpus`
- `--with-nvml-lib=DIR` (library path for `libnvidia-ml.so`)
- `--with-nvml-include=DIR` (include path for `nvml.h`)

**i** `nvml.h` is only found in the NVIDIA CUDA ToolKit.

Example:

```
./configure --with-debug --enable-nvidia-gpus --with-nvml-lib=/usr/lib64 --with-nvml-include=/cuda/NVML --with-hwloc-path=/usr/lib64/
```

## 20.3.2 Package Installation/Upgrade

The package files are self-extracting packages that can be copied and executed on your production nodes to do a new installation, or overlay and upgrade existing installations in the same locations. Example:

```
> make packages
Building ./torque-package-clients-linux-x86_64.sh ...
Building ./torque-package-mom-linux-x86_64.sh ...
Building ./torque-package-server-linux-x86_64.sh ...
Building ./torque-package-gui-linux-x86_64.sh ...
Building ./torque-package-devel-linux-x86_64.sh ...
Done.
$
$ ls -l torque-package-*
-rwxr-xr-x 1 root root 2180510 May 19 15:05 torque-package-clients-linux-x86_64.sh
-rwxr-xr-x 1 root root 4066774 May 19 15:05 torque-package-devel-linux-x86_64.sh
-rwxr-xr-x 1 root root 163505 May 19 15:05 torque-package-doc-linux-x86_64.sh
-rwxr-xr-x 1 root root 4813027 May 19 15:05 torque-package-mom-linux-x86_64.sh
-rwxr-xr-x 1 root root 8168502 May 19 15:05 torque-package-server-linux-x86_64.sh
$
$ ./torque-package-clients-linux-x86_64.sh --install

Installing TORQUE archive...

Done.
$
```

**i** When updating, it is good practice to stop the `pbs_server` and make a backup of the Torque home directory. You will also want to backup the output of `qmgr -c "print server"`. The update will only overwrite the binaries. To do a 'rolling upgrade' and have the MOMs automatically start at a safe point (when `pbs_mom` is between jobs), install the new version and then toggle `enablemomrestart` for the existing MOM processes. Example:

```
# pdsh -w node[01-99] /usr/local/sbin/momctl -q enablemomrestart=1
```

or

```
momctl -q enablemomrestart=1 -h :ALL
```

Refer to the Rolling Upgrade subsection of Appendix E: Considerations Before Upgrading in the *Torque Administrator Guide* for more details and suggestions.

**i** If you move GPU cards to different slots, you must restart `pbs_server` in order for Torque to recognize the drivers as the same ones in different locations rather than two new, additional drivers.

For further details, see these topics:

- [Torque Configuration](#)
- [GPU Modes for NVIDIA 260.x Driver](#)
- [GPU Modes for NVIDIA 270.x Driver](#)
- [gpu\\_status](#)
- [Enabling Persistence Mode](#)
- [Requesting GPUs and Setting GPU Mode](#)

### 20.3.3 Torque Configuration

**i** CUDA 6.0 or later is recommended for Torque 6.0 or later. CUDA version 4.1 is the minimum required. Using `nvidia-smi` to configure Torque for NVIDIA GPU support is deprecated.

To use the NVML (NVIDIA Management Library) API instead of `nvidia-smi`, configure Torque using `--with-nvml-lib=DIR` and `--with-nvml-include=DIR`. These commands specify the location of the `libnvidia-ml` library and the location of the `nvml.h` include file.

```
./configure --with-nvml-lib=/usr/lib --with-nvml-include=/usr/local/cuda/Tools/NVML
server_priv/nodes:
node001  gpus=1
node002  gpus=4
...
pbsnodes -a
node001
...
      gpus = 1
...
```

By default, when Torque is configured with `--enable-nvidia-gpus` the `TORQUE_HOME/nodes` file is automatically updated with the correct GPU count for each MOM node.

### 20.3.4 GPU Modes for NVIDIA 260.x Driver

- 0 – Default - Shared mode available for multiple processes
- 1 – Exclusive - Only one COMPUTE thread is allowed to run on the GPU
- 2 – Prohibited - No COMPUTE contexts are allowed to run on the GPU



Effective with 5.1.3, 6.0.2, and later, prohibited mode is not allowed for user jobs.

### 20.3.5 GPU Modes for NVIDIA 270.x Driver

- 0 – Default - Shared mode available for multiple processes
- 1 – Exclusive Thread - Only one COMPUTE thread is allowed to run on the GPU (v260 exclusive)
- 2 – Prohibited - No COMPUTE contexts are allowed to run on the GPU



Effective with 5.1.3, 6.0.2, and later, prohibited mode is not allowed for user jobs.

- 3 – Exclusive Process - Only one COMPUTE process is allowed to run on the GPU

### 20.3.6 gpu\_status

```
root@gpu:~# pbsnodes gpu
gpu
...
gpus = 2
gpu_status = gpu[1]=gpu_id=0:6:0;gpu_product_name=Tesla
C2050;gpu_display=Disabled;gpu_pci_device_id=6D110DE;gpu_pci_location_id=0:6:0;
gpu_fan_speed=54 %;gpu_memory_total=2687 Mb;gpu_memory_used=74
Mb;gpu_mode=Default;gpu_state=Unallocated;gpu_utilization=96
%;gpu_memory_utilization=10
```

```

%;gpu_ecc_mode=Enabled;gpu_single_bit_ecc_errors=0;gpu_double_bit_ecc_errors=
0;gpu_temperature=88 C,gpu[0]=gpu_id=0:5:0;gpu_product_name=Tesla
C2050;gpu_display=Enabled;gpu_pci_device_id=6D110DE;gpu_pci_location_id=0:5:0;
gpu_fan_speed=66 %;gpu_memory_total=2687 Mb;gpu_memory_used=136
Mb;gpu_mode=Default;gpu_state=Unallocated;gpu_utilization=96
%;gpu_memory_utilization=10
%;gpu_ecc_mode=Enabled;gpu_single_bit_ecc_errors=0;
gpu_double_bit_ecc_errors=0;gpu_temperature=86 C,driver_ver=270.41.06,timestamp=Wed
May 4 13:00:35
2022

```

### 20.3.7 Enabling Persistence Mode

When `pbs_mom` has been built with NVIDIA support and linked to NVML libraries, `pbs_mom` regularly calls `nvmlInit()`. `nvmlInit` can take a long time to complete when persistence mode has not been enabled for NVIDIA devices. It is recommended that you enable persistence mode for NVIDIA GPUs. You can enable persistence mode using `nvidia-smi` (as root):

```
nvidia-smi -i <target gpu> -pm ENABLED
```

or

```
nvidia-smi -i <target gpu> --persistence-mode ENABLED
```

When the `-i` option is not specified, `nvidia-smi` enables persistence mode for all NVIDIA GPUs. The effect of this operation is immediate, however, it does not persist across reboots. After each reboot, persistence mode will default to `DISABLED`. To have persistence mode persist across reboots, run the NVIDIA Persistence Daemon. See [NVIDIA GPU Management and Deployment](https://docs.nvidia.com/deploy/driver-persistence/index.html) at <https://docs.nvidia.com/deploy/driver-persistence/index.html> for more information on persistence mode and [Persistence Daemon](https://docs.nvidia.com/deploy/driver-persistence/index.html#security) at <https://docs.nvidia.com/deploy/driver-persistence/index.html#security> for more information on the Persistence Daemon.

### 20.3.8 Requesting GPUs and Setting GPU Mode

**i** If you are using CUDA 8 or newer, the default of `exclusive_thread` is no longer supported. If the server specifies an `exclusive_thread` setting, the MOM will substitute an `exclusive_process` mode setting. We recommend that you set the default to `exclusive_process`.

`qsub` allows specifying required compute mode when requesting GPUs. If no GPU mode is requested, it will default to 'exclusive' for NVIDIA driver version 260 or 'exclusive\_thread' for NVIDIA driver version 270 and above.

- `qsub -l nodes=1:ppn=1:gpus=1`
- `qsub -l nodes=1:gpus=1`
- `qsub -l nodes=1:gpus=1:default`
- `qsub -l nodes=1:gpus=1:shared` ('shared' and 'default' are both recognized and are equivalent.)
- `qsub -l nodes=1:gpus=1:exclusive_thread`
- `qsub -l nodes=1:gpus=1:exclusive_process`
- `qsub -l nodes=1:gpus=1:reseterr`
- `qsub -l nodes=1:gpus=1:reseterr:exclusive_thread`  
(`exclusive_thread:reseterr`)
- `qsub -l nodes=1:gpus=1:reseterr:exclusive_process`

---

### Related Topics

- [20.1 Scheduling GPUs - page 880](#)
- [20.2 Using GPUs with NUMA - page 882](#)

## 20.4 GPU Metrics

GPU metrics can be collected for nodes that:

- Have one or more GPUs.
- Run Torque 2.5.x or later.
- Use NVIDIA drivers v260.x or v270.x.


GPU metric tracking must be enabled in `moab.cfg`:

```
RMCFG[torque] flags=RECORDGPUMETRICS
```

**i** There is one GPU metric for all GPU devices within a node (`gpu_timestamp`) and nine GPU metrics for each GPU device within a node. If the maximum GPU devices within a node is 4, you must increase the `MAXGMETRIC` value in `moab.cfg` by  $(\text{maxgpudevices} \times \text{gpumetrics}) + 1$ . In this case, the formula is  $(4 \times 9) + 1 = 37$ , so whatever the `MAXGMETRIC` value is, it must be increased by 37. This way, when enabling GPU metrics recording, Moab has enough `GMETRIC` types to accommodate the GPU metrics.

## GPU Metrics Map

The GPU metric names map is as follows (where X is the GPU number):

Metric name as returned by pbsnodes	GMETRIC name as stored in Moab	Metric output
<b>timestamp</b>	gpu_timestamp <div style="border: 1px solid black; padding: 5px; margin-top: 10px;">  The gpu_timestamp metric is global to all GPUs on the node and indicates the last time the driver collected information on the GPUs. </div>	The time data was collected in epoch time
<b>gpu_fan_speed</b>	gpuX_fan	The current fan speed as a percentage
<b>gpu_memory_total</b>	gpuX_mem	The total GPU memory in megabytes
<b>gpu_memory_used</b>	gpuX_usedmem	The total used GPU memory in megabytes
<b>gpu_utilization</b>	gpuX_util	The GPU capability currently in use as a percentage
<b>gpu_memory_utilization</b>	gpuX_memutil	The GPU memory currently in use as a percentage
<b>gpu_ecc_mode</b>	gpuX_ecc	Whether ECC is enabled or disabled
<b>gpu_single_bit_ecc_errors</b>	gpuX_ecc1err	The total number of EEC single-bit errors since the last counter reset
<b>gpu_double_bit_ecc_errors</b>	gpuX_ecc2err	The total number of EEC double-bit errors since the last counter reset

Metric name as returned by pbsnodes	GMETRIC name as stored in Moab	Metric output
<b>gpu_temperature</b>	gpuX_temp	The GPU current temperature in Celsius

*Example 20-1: GPU example*

```

$ mdiag -n -v --xml

<Data>
<node AGRES="GPUS=2;"
AVLCLASS="[test 8][batch 8]"
CFGCLASS="[test 8][batch 8]"
GMETRIC="gpu1_fan:59.00,gpu1_mem:2687.00,gpu1_usedmem:74.00,gpu1_util:94.00,gpu1_
memutil:9.00,gpu1_ecc:0.00,gpu1_ecc1err:0.00,gpu1_ecc2err:0.00,gpu1_temp:89.00,gpu0_
fan:70.00,gpu0_mem:2687.00,gpu0_usedmem:136.00,gpu0_util:94.00,gpu0_memutil:9.00,gpu0_
ecc:0.00,gpu0_ecc1err:0.00,gpu0_ecc2err:0.00,gpu0_temp:89.00,gpu_
timestamp:1304526680.00"
GRES="GPUS=2;"
LASTUPDATETIME="1304526518" LOAD="1.050000"
MAXJOB="0" MAXJOBPERUSER="0" MAXLOAD="0.000000" NODEID="gpu"
NODEINDEX="0" NODESTATE="Idle" OS="linux" OSLIST="linux"
PARTITION="makai" PRIORITY="0" PROCSPEED="0" RADISK="1"
RAMEM="5978" RAPROC="7" RASWAP="22722" RCDISK="1" RCMEM="5978"
RCPROC="8" RCSWAP="23493" RMACCESSLIST="makai" SPEED="1.000000"
STATMODIFYTIME="1304525679" STATTOTALTIME="315649"
STATUPTIME="315649"></node>
</Data>

```

## 20.5 Intel® Xeon Phi™ Coprocessor Configuration

### In this topic:

- [20.5.1 Intel Many-Integrated Cores \(MIC\) Architecture Configuration - page 892](#)
  - [20.5.1.A Prerequisites - page 892](#)
  - [20.5.1.B Setup Options - page 892](#)
- [20.5.2 Validating the Configuration - page 893](#)
  - [20.5.2.A Torque - page 893](#)
  - [20.5.2.B Moab - page 893](#)
- [20.5.3 Job Submission - page 894](#)
  - [20.5.3.A Syntax - page 894](#)
  - [20.5.3.B qstat -f - page 894](#)

### 20.5.3.C checkjob -v - page 895

## 20.5.1 Intel Many-Integrated Cores (MIC) Architecture Configuration

If you use an Intel Many-Integrated Cores (MIC) architecture-based product (e.g., Intel Xeon Phi™) in your cluster for parallel processing, you must configure Torque to detect them.

### 20.5.1.A Prerequisites

- Torque 4.2 or later
- If you set up Torque using auto-detection and intend to get the MIC-based device status report, you must build `pbs_mom` on a system that has the lower-level API libraries for the MIC-based device(s) installed. Additionally, every MOM built with `--enable-mics` and running on a compute node must already have the lower-level API libraries installed on the node. Note that the library is called `coi_host`. You must obtain the API libraries from [Intel](#).

### 20.5.1.B Setup Options

There are two ways to configure MIC-based devices with Torque: (1) manually and (2) by auto-detection.

#### Manual Configuration

- Add `mics=X` to the `nodes` file for the appropriate nodes. See 'Specifying Compute Nodes' in the *Torque Resource Manager Administrator Guide* for more information.

```
napali np=12 mics=2
```

#### Auto-Detect

When you use auto-detection, `pbs_mom` discovers the MIC-based devices and reports them to `pbs_server`.

- At build time, add `--enable-mics` to the configure line:

```
./configure --enable-mics <other configure options>
```



## 20.5.2 Validating the Configuration

### 20.5.2.A Torque

#### pbsnodes

*Example 20-2: pbsnodes output*

```
slesmic
  state = free
  np = 100
  ntype = cluster
  status =
rectime=1347634381,varattr=,jobs=,state=free,netload=7442004852,gres=,loadave=0.00,ncpus=32,physmem=65925692kb,availmem=66531344kb,totmem=68028984kb,idletime=59059,nusers=2,nsessions=8,sessions=4387 4391 4392 4436 4439 4443 4459 100395,uname=Linux slesmic 3.0.13-0.27-default #1 SMP Wed June 15 13:33:49 UTC 2022 (d73692b) x86_64,opsys=linux
  mom_service_port = 15002
  mom_manager_port = 15003
  mics = 2
  mic_status = mic[1]=mic_id=8796;num_cores=61;num_threads=244;physmem=8065748992;free_physmem=7854972928;swap=0;free_swap=0;max_frequency=1090;isa=COI_ISA_KNC;load=0.000000;normalized_load=0.000000;,mic[0]=mic_id=8796;num_cores=61;num_threads=244;physmem=8065748992;free_physmem=7872712704;swap=0;free_swap=0;max_frequency=1090;isa=COI_ISA_KNC;load=0.540000;normalized_load=0.008852;

rhmic.ac
  state = free
  np = 100
  ntype = cluster
  status =
rectime=1347634381,varattr=,jobs=,state=free,netload=3006171583,gres=,loadave=0.00,ncpus=32,physmem=65918268kb,availmem=66901588kb,totmem=67982644kb,idletime=59477,nusers=2,nsessions=2,sessions=3401 29320,uname=Linux rhmic.ac 2.6.32-220.el6.x86_64 #1 SMP Tue Dec 6 19:48:22 GMT 2022 x86_64,opsys=linux
  mom_service_port = 15002
  mom_manager_port = 15003
  mics = 1
  mic_status = mic[0]=mic_id=8796;num_cores=61;num_threads=244;physmem=8065748992;free_physmem=7872032768;swap=0;free_swap=0;max_frequency=1090;isa=COI_ISA_KNC;load=0.540000;normalized_load=0.008852;<mic_status>;
```

### 20.5.2.B Moab

#### mdiag -n -v

*Example 20-3: mdiag -n -v output*

```
$ mdiag -n -v
compute node summary
Name          State  Procs  Memory  Disk  Swap
Speed  Opsys  Arch Par  Load Classes  Features
```

```

hola      Idle    4:4    8002:8002    1:1    10236:13723
1.00  linux    - hol    0.24 [batch]
GRES=MICS:2,
-----
          ---    4:4    8002:8002    1:1    10236:13723

Total Nodes: 1  (Active: 0  Idle: 1  Down: 0)

```

## checknode -v

### Example 20-4: checknode output

```

$ checknode slesmic
node slesmic

State:      Idle (in current state for 00:00:16)
Configured Resources: PROCS: 100 MEM: 62G SWAP: 64G DISK: 1M MICS: 2
Utilized Resources: SWAP: 1581M
Dedicated Resources: ---
Generic Metrics:  mic1_mic_id=8796.00,mic1_num_cores=61.00,mic1_num_
threads=244.00,mic1_physmem=8065748992.00,mic1_free_physmem=7854972928.00,mic1_
swap=0.00,mic1_free_swap=0.00,mic1_max_frequency=1090.00,mic1_load=0.12,mic1_
normalized_load=0.00,mic0_mic_id=8796.00,mic0_num_cores=61.00,mic0_num_
threads=244.00,mic0_physmem=8065748992.00,mic0_free_physmem=7872679936.00,mic0_
swap=0.00,mic0_free_swap=0.00,mic0_max_frequency=1090.00
  MTBF(longterm):  INFINITY MTBF(24h):  INFINITY
Opsys:      linux    Arch:      ---
Speed:      1.00    CPULoad:   0.000
Classes:    [batch]
RM[napali]* TYPE=PBS
EffNodeAccessPolicy: SHARED

Total Time: 3:45:43 Up: 3:45:43 (100.00%) Active: 00:00:00 (0.00%)

Reservations:
---
```

## 20.5.3 Job Submission

### 20.5.3.A Syntax

#### Example 20-5: Request MIC-based device(s) in qsub

```
qsub .... -l nodes=X:mics=Y
```

### 20.5.3.B qstat -f

#### Example 20-6: qstat -f output

```

Job Id: 5271.napali
Job_Name = STDIN
Job_Owner = dbeer@napali

```

```

job_state = Q
queue = batch
server = napali
Checkpoint = u
ctime = Fri Oct 14 08:56:33 2022
Error_Path = napali:/home/dbeer/dev/private-torque/trunk/STDIN.e5271
Hold_Types = n
Join_Path = oe
Keep_Files = n
Mail_Points = a
mtime = Fri Oct 14 08:56:33 2022
Output_Path = napali:/home/dbeer/dev/private-torque/trunk/STDIN.o5271
Priority = 0
qtime = Fri Oct 14 08:56:33 2022
Rerunable = True
Resource_List.nodect = 1
Resource_List.nodes = 1:mics=1
substate = 10
Variable_List = PBS_O_QUEUE=batch,PBS_O_HOME=/home/dbeer,
                PBS_O_LOGNAME=dbeer,
                PBS_O_PATH=/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/b
                in:/usr/games,PBS_O_MAIL=/var/mail/dbeer,PBS_O_SHELL=/bin/bash,
                PBS_O_LANG=en_US.UTF-8,
                PBS_O_SUBMIT_FILTER=/usr/local/sbin/torque_submitfilter,
                PBS_O_WORKDIR=/home/dbeer/dev/private-torque/trunk,PBS_O_HOST=napali,
                PBS_O_SERVER=napali
euser = dbeer
egroup = company
queue_rank = 3
queue_type = E
etime = Fri Oct 14 08:56:33 2022
submit_args = -l nodes=1:mics=1
fault_tolerant = False
job_radix = 0
submit_host = napali

```

### 20.5.3.C checkjob -v

#### Example 20-7: checkjob -v output

```

dthompson@mahalo:~/dev/moab-test/trunk$ checkjob -v 2
job 2 (RM job '2.mahalo')

AName: STDIN
State: Idle
Creds: user:dthompson group:dthompson class:batch
WallTime: 00:00:00 of 1:00:00
SubmitTime: Thu Sep 13 17:06:06
(Time Queued Total: 00:00:24 Eligible: 00:00:02)

TemplateSets: DEFAULT
Total Requested Tasks: 1

Req[0] TaskCount: 1 Partition: ALL
Dedicated Resources Per Task: PROCS: 1 MICS: 1
...

```

## 20.6 Intel® Xeon Phi™ Co-processor Metrics

Intel Many-Integrated Cores (MIC) architecture-based device (e.g., Intel Xeon Phi™) metrics can be collected for nodes that:

- Have one or more MIC-based devices.
- Run Torque 4.2.x or later.
- Run Moab 7.2 or later.

MIC-based device metric tracking must be enabled in `moab.cfg`:

```
RMCFG[torque] flags=RECORDMICMETRICS
```

**i** There are 11 metrics for each MIC-based device within a node. If the maximum MIC-based devices within a node is 4, you must increase the `MAXGMETRIC` value in `moab.cfg` by `(maxmicdevices x micmetrics)`. In this case, the formula is  $(4 \times 11) = 44$ , so whatever the `MAXGMETRIC` value is, it must be increased by 44. This way, when enabling MIC-based device metrics recording, Moab has enough `GMETRIC` types to accommodate the additional metrics.

### MIC-Based Metrics Map

The MIC-based metric names map is as follows (where `X` is the MIC-based device number):

Metric name as returned by pbsnodes	GMETRIC name as stored in Moab	Metric output
<code>mic_id</code>	<code>micX_mic_id</code>	The ID of the MIC-based device
<code>num_cores</code>	<code>micX_num_cores</code>	The number of cores in the MIC-based device
<code>num_threads</code>	<code>micX_num_threads</code>	The number of hardware threads on the MIC-based device
<code>physmem</code>	<code>micX_physmem</code>	The total physical memory in the MIC-based device
<code>free_physmem</code>	<code>micX_free_physmem</code>	The available physical memory in the MIC-based device
<code>swap</code>	<code>micX_swap</code>	The total swap space on the MIC-based device

Metric name as returned by pbsnodes	GMETRIC name as stored in Moab	Metric output
free_swap	micX_free_swap	The unused swap space on the MIC-based device
max_frequency	micX_max_frequency	The maximum frequency speed of any core in the MIC-based device
isa	micX_isa	The hardware interface type of the MIC-based device
load	micX_load	The total current load of the MIC-based device
normalized_load	micX_normalized_load	The normalized load of the MIC-based device (total load divided by number of cores in the MIC-based device)

## Chapter 21: Preemption


Sites possess workloads of varying importance, and users might want to run jobs with higher priorities before jobs with lower priorities. This can be done by using preemption. Preemption is simply the process by which a higher-priority job can take the place of a lower-priority job. You can also use preemption for optimistic scheduling and development job support.

This section explains how to configure and use preemption. [Simple Example of Preemption](#) offers a basic introduction and contains examples to help you get started using preemption. The other sections offer more explanation and information about what you can do with preemption and contain some best practices that will help you avoid the need for troubleshooting in the future.

While this section does not explain every possible preemption configuration, it does prescribe the best practices for setting up and using preemption with your system. It is recommended that you follow the established instructions contained in this section.

 Preemption does not work with dynamic provisioning.

 Neither `SPANEVENLY` nor `DELAY` values of the `NODESETPLUS` parameter will work with multi-req jobs or preemption.

 Do not allow preemption with interactive jobs unless `PREEMPTPOLICY` is set to `CANCEL`. (For more information, see [Canceling Jobs with Preemption](#).)

In this section:

- [21.1 Preemption Tasks - page 899](#)
- [21.2 Preemption Reference - page 915](#)

---

### Related Topics

- [Chapter 7: Optimizing Scheduling Behavior – Backfill and Node Sets - page 535](#)

## 21.1 Preemption Tasks

In this section:

- [21.1.1 Canceling Jobs with Preemption - page 899](#)
- [21.1.2 Checkpointing Jobs with Preemption - page 902](#)
- [21.1.3 Requeuing Jobs with Preemption - page 904](#)
- [21.1.4 Suspending Jobs with Preemption - page 906](#)
- [21.1.5 Using Owner Preemption - page 910](#)
- [21.1.6 Using QoS Preemption - page 914](#)

### 21.1.1 Canceling Jobs with Preemption

CANCEL is one of the [PREEMTPOLICY](#) types (for more information, see [PREEMTPOLICY Types](#)). The CANCEL attribute cancels active jobs, regardless of any [JOBFLAGS](#) (such as [REQUEUEABLE](#) or [SUSPENDABLE](#)).

For information about [PREEMPTTEE](#) and [PREEMPTOR](#) flags, see [Preemption Flags](#)

**i** You should not allow preemption with interactive jobs unless [PREEMTPOLICY](#) is set to CANCEL.

The following outlines some benefits of using CANCEL and also lists some things you should be aware of if you choose to use it:

- Advantages - This attribute is the easiest to configure and use.
- Cautions - Canceled jobs are not automatically restarted or requeued. Users must resubmit canceled jobs.

#### To Preempt Jobs Using CANCEL

1. Make the following configurations to the `moab.cfg` file:
  - a. Set [GUARANTEEDPREEMPTION](#) to TRUE. (This causes Moab to lock [PREEMPTOR](#) jobs until [JOBRETRYTIME](#) expires.)
  - b. Set [PREEMPTIONALGORITHM](#) to specify how Moab handles preemption scheduling policies.

**i** If you use `JOBNODEMATCHPOLICY EXACTNODE`, you must also add `PREEMPTIONALGORITHM PREEMPTORCENTRIC` in order for preemption to function reliably.

- c. Set `PREEMPTPOLICY` to `CANCEL` (for more information, see [PREEMPTPOLICY Types](#)).
- d. Make sure that the `PREEMPTTEE` job has a lower priority than the `PREEMPTOR` job (for more information, see [Preemption Flags](#)).

For example:

```
GUARANTEEDPREEMPTION TRUE
PREEMPTPOLICY CANCEL

QOSCFG[test1] QFLAGS=PREEMPTTEE MEMBERULIST=john PRIORITY=100
QOSCFG[test2] QFLAGS=PREEMPTOR MEMBERULIST=john PRIORITY=10000
```

2. Submit a job to the preemptee QoS (`test1`). For example:

```
[john@g06]# echo sleep 600 | msub -l walltime=600 -l qos=test1 -l procs=128
```

(Optional) Examine the following output for `showq`:

```
Moab.7
[john@g06]# showq

active jobs-----
JOBID      USERNAME    STATE      PROCS      REMAINING   STARTTIME
Moab.7     john        Running    128        00:01:59    Thu Nov 10 12:28:44

1 active job      128 of 128 processors in use by local jobs (100.00%)
                   2 of 2 nodes active (100.00%)

eligible jobs-----
JOBID      USERNAME    STATE      PROCS      WCLIMIT     QUEUE TIME
0 eligible jobs

blocked jobs-----
JOBID      USERNAME    STATE      PROCS      WCLIMIT     QUEUE TIME
0 blocked jobs

Total job: 1
```

3. Now submit a job to the preemptor QoS (`test2`). For example:

```
[john@g06]$ echo sleep 120 | msub -l procs=128,walltime=120 -l qos=test2
```

(Optional) Examine the following output for `showq`:

```
Moab.8
[john@g06]# showq
```



```

active jobs-----
JOBID  USERNAME  STATE      PROCS    REMAINING  STARTTIME
Moab.7  john      Canceling 128      00:01:56  Thu Nov 10 12:28:44
Moab.8  john      Running    128      00:02:00  Thu Nov 10 12:28:48

2 active jobs 128 of 128 processors in use by local jobs (100.00%)
      2 of 2 nodes active (100.00%)

eligible jobs-----
JOBID  USERNAME  STATE      PROCS    WCLIMIT    QUEUE TIME

0 eligible jobs

blocked jobs-----
JOBID  USERNAME  STATE      PROCS    WCLIMIT    QUEUE TIME

0 blocked jobs

Total jobs: 2

```

*Note that test1 is canceled when test2 is submitted.*

(Optional) Examine the `checkjob` outputs for these two jobs:

```

[john@g06]$ checkjob Moab.9
job Moab.9

State: Removed
Completion Code: -1 Time: Thu Nov 10 12:28:48
Creds: user:john group:john qos:test1
WallTime: 00:00:02 of 00:02:00
SubmitTime: Thu Nov 10 12:28:44
(Time Queued Total: 00:00:07 Eligible: 00:00:00)

Total Requested Tasks: 128

Req[0] TaskCount: 128 Partition: licenses
NodeCount: 2

Allocated Nodes:
node[01-02]*64

IWD: /opt/native
SubmitDir: /opt/native
Executable: /opt/native/spool/moab.job.UFe8sQ

StartCount: 1
Flags: GLOBALQUEUE, PROCSPECIFIED
Attr: PREEMPTEE
StartPriority: 100

```

*Note that the preempted job has been removed.*

```

[john@g06]$ checkjob Moab.10
job Moab.10

State: Running
Creds: user:john group:john qos:test2
WallTime: 00:00:00 of 00:02:00

```

```

SubmitTime: Thu Nov 10 12:36:31
(Time Queued Total: 00:00:00 Eligible: 00:00:00)

StartTime: Thu Nov 10 12:28:48
Total Requested Tasks: 128

Req[0] TaskCount: 128 Partition: licenses

Allocated Nodes:
node[01-02]*64

IWD: /opt/native
SubmitDir: /opt/native
Executable: /opt/native/spool/moab.job.CZavjU

StartCount: 1
Flags: HASPREEMPTED, PREEMPTOR, GLOBALQUEUE, PROCSPECIFIED
StartPriority: 10000
Reservation 'Moab.10' (-00:00:07 -> 00:01:53 Duration: 00:02:00)

```

---

## Related Topics

- [21.1.4 Suspending Jobs with Preemption - page 906](#)
- [21.1.2 Checkpointing Jobs with Preemption - page 902](#)
- [21.1.3 Requeuing Jobs with Preemption - page 904](#)
- [21.2.2 Preemption Flags - page 917](#)
- [Chapter 21: Preemption - page 898](#)
- [21.2.3 PREEMTPOLICY Types - page 918](#)
- [21.2.5 Testing and Troubleshooting Preemption - page 923](#)

## 21.1.2 Checkpointing Jobs with Preemption

CHECKPOINT is one of the [PREEMTPOLICY](#) types (for more information, see [PREEMTPOLICY Types](#)). For systems that allow checkpointing, the CHECKPOINT attribute allows a job to save its current state and either terminate or continue running. A checkpointed job can restart at any time and resume execution from its most recent checkpoint.

You can tune checkpointing behavior on a per-resource manager-basis by setting the CHECKPOINTSIG and CHECKPOINTTIMEOUT attributes of the [RMCFG](#) parameter.

For information about PREEMPTTEE and PREEMPTOR flags, see [Preemption Flags](#)

The following outlines some benefits of using `CHECKPOINT` and also lists some things you should be aware of if you choose to use it:

- Advantages - This attribute allows you to restart a job from its last checkpoint.
- Cautions - Jobs tend to take longer to complete when you use `CHECKPOINT`.

## To Preempt Jobs Using `CHECKPOINT`

Make the following configurations to the `moab.cfg` file:

1. Set `GUARANTEEDPREEMPTION` to `TRUE`. (This causes Moab to lock `PREEMPTOR` jobs until `JOBRETRYTIME` expires. This locks the job on a node and keeps trying to preempt.)
2. Set `PREEMPTIONALGORITHM` to specify how Moab handles preemption scheduling policies.

**i** If you use `JOBNODEMATCHPOLICY EXACTNODE`, you must also add `PREEMPTIONALGORITHM PREEMPTORCENTRIC` in order for preemption to function reliably.

3. Set `PREEMPTPOLICY` to `CHECKPOINT` (for more information, see [PREEMPTPOLICY Types](#)).
4. Make sure that the `PREEMPTEE` job has a lower priority than the `PREEMPTOR` job (for more information, see [Preemption Flags](#)).

For example:

```
GUARANTEEDPREEMPTION TRUE
PREEMPTPOLICY CHECKPOINT

QOSCFG[test1] QFLAGS=PREEMPTEE MEMBERULIST=john PRIORITY=100
QOSCFG[test2] QFLAGS=PREEMPTOR MEMBERULIST=john PRIORITY=10000
```

---

### Related Topics

- [21.1.4 Suspending Jobs with Preemption - page 906](#)
- [21.1.3 Requeuing Jobs with Preemption - page 904](#)
- [Canceling Jobs with Preemption](#)
- [21.2.2 Preemption Flags - page 917](#)
- [Chapter 21: Preemption - page 898](#)
- [21.2.3 PREEMPTPOLICY Types - page 918](#)
- [21.2.5 Testing and Troubleshooting Preemption - page 923](#)

## 21.1.3 Requeuing Jobs with Preemption

REQUEUE is one of the [PREEMTPOLICY](#) types (for more information, see [PREEMTPOLICY Types](#)). The REQUEUE value terminates active jobs and returns them to the job queue in an idle state.

For information about [PREEMPTTEE](#) and [PREEMPTOR](#) flags, see [Preemption Flags](#)

The following outlines some benefits of using REQUEUE and also lists some things you should be aware of if you choose to use it:

- Advantages - Jobs are automatically resubmitted into the job queue.
- Cautions - A job gets resubmitted in the job queue at the same priority it had when Moab originally started it (i.e., the job does not jump ahead in the queue). Jobs start over from the beginning.

**i** You must mark a job as `RESTARTABLE` if you want it to requeue. If you do not, the job will be canceled when it is preempted.

If supported by the resource manager, you can set the `RESTARTABLE` job flag when submitting the job by using the `msub -r` option. Otherwise, use the `JOBFLAGS` attribute of the associated class or QoS credential, as in this example:

```
CLASSCFG[low] JOBFLAGS=RESTARTABLE
```

### To Preempt Jobs Using REQUEUE

1. Make the following configurations to the `moab.cfg` file:
  - a. Set [GUARANTEEDPREEMPTION](#) to `TRUE`. (This causes Moab to lock [PREEMPTOR](#) jobs until [JOBRETRYTIME](#) expires.)
  - b. Set [PREEMPTIONALGORITHM](#) to specify how Moab handles preemption scheduling policies.

**i** If you use [JOBNODEMATCHPOLICY EXACTNODE](#), you must also add [PREEMPTIONALGORITHM PREEMPTORCENTRIC](#) in order for preemption to function reliably.

- c. Set [PREEMTPOLICY](#) to `REQUEUE` (for more information, see [PREEMTPOLICY Types](#)).
- d. Make sure that the [PREEMPTTEE](#) job has a lower priority than the [PREEMPTOR](#) job (for more information, see [Preemption Flags](#)).

For example:

```

GUARANTEEDPREEMPTION TRUE
PREEMPTPOLICY REQUEUE

QOSCFG[test1] QFLAGS=PREEMPTTEE JOBFLAGS=RESTARTABLE MEMBERULIST=john PRIORITY=100
QOSCFG[test2] QFLAGS=PREEMPTOR MEMBERULIST=john PRIORITY=10000

```

## 2. Submit a job to the preemptee QoS (`test1`). For example:

```
[john@g06]# echo sleep 600 | msub -l walltime=600 -l qos=test1 -l procs=128
```

(Optional) Examine the following output for `showq`:

```

Moab.1
[john@g06]# showq

active jobs-----
JOBID      USERNAME    STATE      PROCS      REMAINING    STARTTIME
Moab.1     john        Running    128        00:09:59     Wed Nov 9 15:56:33

1 active job      128 of 128 processors in use by local jobs (100.00%)
                    2 of 2 nodes active (100.00%)

eligible jobs-----
JOBID      USERNAME    STATE      PROCS      WCLIMIT      QUEUE TIME
0 eligible jobs

blocked jobs-----
JOBID      USERNAME    STATE      PROCS      WCLIMIT      QUEUE TIME
0 blocked jobs

Total job: 1

```

## 3. Now submit a job to the preemptor QoS (`test2`). For example:

```
[john@g06]# echo sleep 600 | msub -l walltime=600 -l qos=test2 -l procs=128
```

(Optional) Examine the following output for `showq` and `checkjob`:

```

Moab.2
[john@g06]# showq

active jobs-----
JOBID      USERNAME    STATE      PROCS      REMAINING    STARTTIME
Moab.2     john        Running    128        00:09:59     Wed Nov 9 15:56:47

1 active job 128 of 128 processors in use by local jobs (100.00%)
                    2 of 2 nodes active (100.00%)

eligible jobs-----
JOBID      USERNAME    STATE      PROCS      WCLIMIT      QUEUE TIME
Moab.1     john        Idle       128        00:10:00     Wed Nov 9 15:56:33

1 eligible job

blocked jobs-----
JOBID      USERNAME    STATE      PROCS      WCLIMIT      QUEUE TIME

```

```
0 blocked jobs
```

```
Total jobs: 2
```

```
[john@g06]# checkjob Moab.2
job Moab.2

State: Running
Creds: user:john group:john qos:test2
WallTime: 00:02:04 of 00:10:00
SubmitTime: Wed Nov 9 15:56:46
(Time Queued Total: 00:00:01 Eligible: 00:00:00)

StartTime: Wed Nov 9 15:56:47
Total Requested Tasks: 128

Req[0] TaskCount: 128 Partition: licenses
NodeCount: 2

Allocated Nodes:
node[01-02]*64

IWD: /opt/native
SubmitDir: /opt/native
Executable: /opt/native/spool/moab.job.ELoX5Q

StartCount: 1
Flags: HASPREEMPTED, PREEMPTOR, GLOBALQUEUE, PROCSPECIFIED
StartPriority: 10000
Reservation 'Moab.2' (-00:02:21 -> 00:07:39 Duration: 00:10:00)
```

---

## Related Topics

- [21.1.4 Suspending Jobs with Preemption - page 906](#)
- [21.1.2 Checkpointing Jobs with Preemption - page 902](#)
- [Canceling Jobs with Preemption](#)
- [21.2.2 Preemption Flags - page 917](#)
- [Chapter 21: Preemption - page 898](#)
- [21.2.3 PREEMTPOLICY Types - page 918](#)
- [21.2.5 Testing and Troubleshooting Preemption - page 923](#)

## 21.1.4 Suspending Jobs with Preemption

`SUSPEND` is one of the `PREEMTPOLICY` types (for more information, see [PREEMTPOLICY Types](#)). The `SUSPEND` attribute causes active jobs to stop executing, but

to remain in memory on the allocated compute nodes.

For information about `PREEMPTEE` and `PREEMPTOR` flags, see [Preemption Flags](#)

The following outlines some benefits of using `SUSPEND` and also lists some things you should be aware of if you choose to use it:

- **Advantages** - The job remains in memory on the allocated compute nodes. Using `SUSPEND` frees up processor resources. The job can restart where it left off before it was suspended.
- **Cautions** - There is a possibility that having multiple suspended jobs on a compute node will crash the swap. Moab tracks only *requested* memory of active jobs (not *used* memory). The swap can crash if the job uses a lot of memory and Moab starts other jobs. Suspended jobs do not relinquish their licenses.

**i** You must mark a job as `SUSPENDABLE` if you want it to suspend. If you do not, the job will be requeued or canceled when it is preempted.

If supported by the resource manager, you can set the job `SUSPENDABLE` flag when submitting the job by using the `msub -r` option. Otherwise, use the `JOBFLAGS` attribute of the associated class or QoS credential, as in this example:

```
CLASSCFG[low] JOBFLAGS=SUSPENDABLE
```

## To Preempt Jobs Using `SUSPEND`

When you use `SUSPEND`, you must increase your `JOBRETRYTIME`. By default, `JOBRETRYTIME` is set to 60 seconds, but when you use `SUSPEND`, it is recommended that you increase the time to 300 seconds (5 minutes).

1. Make the following configurations to the `moab.cfg` file:
  - a. Set `GUARANTEEDPREEMPTION` to `TRUE`. (This causes Moab to lock `PREEMPTOR` jobs until `JOBRETRYTIME` expires.)
  - b. Set `PREEMPTIONALGORITHM` to specify how Moab handles preemption scheduling policies.

**i** If you use `JOBNODEMATCHPOLICY EXACTNODE`, you must also add `PREEMPTIONALGORITHM PREEMPTORCENTRIC` in order for preemption to function reliably.

- c. Set `PREEMTPOLICY` to `SUSPEND` (for more information, see [PREEMTPOLICY Types](#)).

- d. For the PREEMPTEE job, set `JOBFLAGS=RESTARTABLE, SUSPENDABLE`.
- e. Make sure that the PREEMPTEE job has a lower priority than the PREEMPTOR job (for more information, see [Preemption Flags](#)).

For example:

```

GUARANTEEDPREEMPTION TRUE
PREEMPTPOLICY SUSPEND

QOSCFG[test1] QFLAGS=PREEMPTEE JOBFLAGS=RESTARTABLE, SUSPENDABLE MEMBERULIST=john
PRIORITY=100
QOSCFG[test2] QFLAGS=PREEMPTOR MEMBERULIST=john PRIORITY=10000

```

2. Submit a job to the preemptee QoS (`test1`). For example:

```
[john@g06]$ echo sleep 120 | msub -l procs=128,walltime=120 -l qos=test1
```

(Optional) Examine the output for `showq`:

```

Moab.7
[john@g06]# showq

active jobs-----
JOBID   USERNAME   STATE      PROCS      REMAINING   STARTTIME
Moab.7   john       Running    128        00:01:59    Thu Nov 10 12:28:44

1 active job      128 of 128 processors in use by local jobs (100.00%)
                    2 of 2 nodes active (100.00%)

eligible jobs-----
JOBID   USERNAME   STATE      PROCS      WCLIMIT     QUEUETIME

0 eligible jobs

blocked jobs-----
JOBID   USERNAME   STATE      PROCS      WCLIMIT     QUEUETIME

0 blocked jobs

Total job: 1

```

3. Now submit a job to the preemptor QoS (`test2`). For example:

```
[john@g06]$ echo sleep 120 | msub -l procs=128,walltime=120 -l qos=test2
```

(Optional) Examine the output for `showq`:

```

Moab.8
[john@g06]# showq

active jobs-----
JOBID   USERNAME   STATE      PROCS      REMAINING   STARTTIME
Moab.7   john       Suspended  128        00:01:56    Thu Nov 10 12:28:44
Moab.8   john       Running    128        00:02:00    Thu Nov 10 12:28:48

2 active jobs 128 of 128 processors in use by local jobs (100.00%)

```



```

                2 of 2 nodes active (100.00%)

eligible jobs-----
JOBID      USERNAME      STATE      PROCS      WCLIMIT      QUEUE TIME
0 eligible jobs

blocked jobs-----
JOBID      USERNAME      STATE      PROCS      WCLIMIT      QUEUE TIME
0 blocked jobs

Total jobs: 2

```

*Note that when a job is suspended, it stays in the output of showq. This is normal behavior for a suspended job. Moab should only suspend a job once.*

#### 4. (Optional) Examine the `checkjob` outputs for these two jobs:

```

[john@g06]$ checkjob Moab.9
job Moab.9

State: Suspended
Creds: user:john group:john qos:test1
WallTime: 00:00:02 of 00:02:00
SubmitTime: Thu Nov 10 12:36:29
(Time Queued Total: 00:00:07 Eligible: 00:00:00)

Total Requested Tasks: 128

Req[0] TaskCount: 128 Partition: licenses
NodeCount: 2

Allocated Nodes:
node[01-02]*64

IWD: /opt/native
SubmitDir: /opt/native
Executable: /opt/native/spool/moab.job.UFe8sQ

StartCount: 1
Flags: RESTARTABLE,SUSPENDABLE,PREEMPTEE,GLOBALQUEUE,PROCSPECIFIED
Attr: PREEMPTEE
StartPriority: 100
job cannot be resumed: preemption required but job is conditional preemptor with no
targets
BLOCK MSG: non-idle state 'Running' (recorded at last scheduling iteration)

```

```

[john@g06]$ checkjob Moab.10
job Moab.10

State: Running
Creds: user:john group:john qos:test2
WallTime: 00:00:00 of 00:02:00
SubmitTime: Thu Nov 10 12:36:31
(Time Queued Total: 00:00:00 Eligible: 00:00:00)

StartTime: Thu Nov 10 12:36:31
Total Requested Tasks: 128

```

```

Req[0] TaskCount: 128 Partition: licenses

Allocated Nodes:
node[01-02]*64

IWD: /opt/native
SubmitDir: /opt/native
Executable: /opt/native/spool/moab.job.CZavjU

StartCount: 1
Flags: HASPREEMPTED, PREEMPTOR, GLOBALQUEUE, PROCSPECIFIED
StartPriority: 10000
Reservation 'Moab.10' (-00:00:07 -> 00:01:53 Duration: 00:02:00)

```

**i** Occasionally, Moab will keep a job from restarting, holding it in a suspended state for a long period of time, if it thinks the job cannot restart. For example, if a job could write to I/O before it was suspended, and now it cannot, Moab would realize the job is unable to start and would leave it in a suspended state.

---

## Related Topics

- [21.1.2 Checkpointing Jobs with Preemption - page 902](#)
- [21.1.3 Requeuing Jobs with Preemption - page 904](#)
- [Canceling Jobs with Preemption](#)
- [21.2.2 Preemption Flags - page 917](#)
- [Chapter 21: Preemption - page 898](#)
- [21.2.3 PREEMTPOLICY Types - page 918](#)
- [21.2.5 Testing and Troubleshooting Preemption - page 923](#)

## 21.1.5 Using Owner Preemption

Owner preemption allows jobs submitted by a reservation owner to preempt jobs submitted by other users (for more information, see [Configuring and Managing Reservations](#)).

Owner preemption is enabled with the [OWNERPREEMPT](#) reservation flag.

For information about [PREEMPTED](#) and [PREEMPTOR](#) flags, see [Preemption Flags](#)

## To Enable Owner Preemption

1. Make the following configurations to the `moab.cfg` file:
  - a. Set `GUARANTEEDPREEMPTION` to `TRUE`. (This causes Moab to lock `PREEMPTOR` jobs until `JOBRETRYTIME` expires.)
  - b. Make sure that `JOBNODEMATCHPOLICY` is *not* set to `EXACTNODE`, which is not currently supported for preemption (for more information, see [Testing and Troubleshooting Preemption](#)).
  - c. Set the `PREEMPTPOLICY` type (for more information, see [PREEMPTPOLICY Types](#)).
  - d. Set the `OWNERPREEMPT` flag.

**i** Optionally, if you want the owner preemption to override any `PREEMPTMINTIME` settings for `PREEMPTEE` jobs, you can set the `OWNERPREEMPTIGNOREMINTIME` flag as well.

- e. Specify an owner.

**i** If the non-owner job does not have a `RESTARTABLE` or `REQUEUEABLE` flag set, the job will cancel.

For example:

```

GUARANTEEDPREEMPTION TRUE
PREEMPTPOLICY <policy>

SRCFG[myrez]  FLAGS=OWNERPREEMPT HOSTLIST=node01
SRCFG[myrez]  OWNER=USER:john
SRCFG[myrez]  USERLIST=jane,john PERIOD=INFINITY

QOSCFG[test1] QFLAGS=PREEMPTEE JOBFLAGS=restartable MEMBERULIST=john PRIORITY=100
QOSCFG[test2] QFLAGS=PREEMPTOR MEMBERULIST=john PRIORITY=10000

```

2. Submit a job to a user who is not the owner (in this example, `jane`):

```
[jane@g06]$ echo sleep 600 | msub -l walltime=600 -l procs=64
```

(Optional) Examine the following output for `showq` and `checkjob` for `jane`'s job:

```

Moab.1
[jane@g06]$ showq

active jobs-----
JOBID      USERNAME    STATE      PROCS      REMAINING    STARTTIME
Moab.1     jane        Running    64          00:09:57     Mon Nov 14 12:07:52

1 active job      64 of 64 processors in use by local jobs (100.00%)
                  1 of 1 nodes active (100.00%)

```

```

eligible jobs-----
JOBID      USERNAME    STATE      PROCS      WCLIMIT     QUEUETIME

0 eligible jobs

blocked jobs-----
JOBID      USERNAME    STATE      PROCS      WCLIMIT     QUEUETIME

0 blocked jobs

Total job: 1

```

```

root@g06]# checkjob Moab.1
job Moab.1

State: Running
Creds: user:jane group:jane
WallTime: 00:01:02 of 00:10:00
SubmitTime: Mon Nov 14 12:07:52
(Time Queued Total: 00:00:00 Eligible: 00:00:00)

StartTime: Mon Nov 14 12:07:52
Total Requested Tasks: 64

Req[0] TaskCount: 64 Partition: FLEXlm
NodeCount: 1

Allocated Nodes:
[node01:64]

IWD: /opt/native
SubmitDir: /opt/native
Executable: /opt/native/spool/moab.job.FoZfIU

StartCount: 1
Flags: GLOBALQUEUE,PROCSPECIFIED
StartPriority: 1
Reservation 'Moab.1' (-00:01:24 -> 00:08:36 Duration: 00:10:00)

```

### 3. Now submit a job for the owner (in this example, john):

```
[john@g06]$ echo sleep 600 | msub -l walltime=600 -l procs=50
```

```
[john@g06]$ echo sleep 600 | msub -l walltime=600 -l procs=50
```

(Optional) Examine the following output for *showq* and *checkjob* for john's job:

```

Moab.2
[john@g06]$ showq

active jobs-----
JOBID      USERNAME    STATE      PROCS      REMAINING    STARTTIME
Moab.1     jane        Canceling  64         00:07:43     Mon Nov 14 12:07:52
Moab.2     john        Running    50         00:09:59     Mon Nov 14 12:10:08

2 active jobs      64 of 64 processors in use by local jobs (100.00%)
                   1 of 1 nodes active (100.00%)

```

```

eligible jobs-----
JOBID      USERNAME    STATE      PROCS      WCLIMIT    QUEUE TIME
0 eligible jobs

blocked jobs-----
JOBID      USERNAME    STATE      PROCS      WCLIMIT    QUEUE TIME
0 blocked jobs

Total jobs: 2

```

*Note that jane's job is canceled once john's job is submitted.*

```

[john@g06]$ checkjob Moab.2
job Moab.2

State: Running
Creds: user:john group:john
WallTime: 00:00:31 of 00:10:00
SubmitTime: Mon Nov 14 12:10:08
(Time Queued Total: 00:00:00 Eligible: 00:00:00)

StartTime: Mon Nov 14 12:10:08
Total Requested Tasks: 50

Req[0] TaskCount: 50 Partition: FLEXlm
NodeCount: 1

Allocated Nodes:
[node01:50]

IWD: /opt/native
SubmitDir: /opt/native
Executable: /opt/native/spool/moab.job.jf1N4a

StartCount: 1
Flags: HASPREEMPTED, GLOBALQUEUE, PROCSPECIFIED
StartPriority: 1
Reservation 'Moab.2' (-00:00:48 -> 00:09:12 Duration: 00:10:00)

```

*Note the new **HASPREEMPTED** flag.*

(Optional) Now look at the `showq` for jane's job (after):

```

[root@g06]# checkjob Moab.1
job Moab.1

State: Removed
Completion Code: -1 Time: Mon Nov 14 12:10:08
Creds: user:jane group:jane
WallTime: 00:02:47 of 00:10:00
SubmitTime: Mon Nov 14 12:07:52
(Time Queued Total: 00:00:00 Eligible: 00:00:00)

Total Requested Tasks: 64

Req[0] TaskCount: 64 Partition: FLEXlm

```

```

NodeCount: 1

Allocated Nodes:
[node01:64]

IWD: /opt/native
Executable: /opt/native/spool/moab.job.FoZfIU

Execution Partition: FLEXlm
Flags: GLOBALQUEUE, PROCSPECIFIED
StartPriority: 0

```

*Note that the state is now **Removed**.*

## Related Topics

- [21.2.2 Preemption Flags - page 917](#)
- [Chapter 21: Preemption - page 898](#)
- [21.2.3 PREEMTPOLICY Types - page 918](#)
- [21.2.5 Testing and Troubleshooting Preemption - page 923](#)

## 21.1.6 Using QoS Preemption

This section breaks down how to configure the `moab.cfg` file to set up preemption with QoS. Using QoS, you can specify preemption rules and control access to preemption privileges by using the `QFLAGS` `PREEMPTTEE` and `PREEMPTOR` credentials. For information about the `PREEMPTTEE` and `PREEMPTOR` flags, see [Preemption Flags](#).

QoS-based preemption only occurs when the following three conditions are satisfied:

- The preemptor job has the `PREEMPTOR` attribute set.
- The preemptee job has the `PREEMPTTEE` attribute set.
- The preemptor job has a higher priority than the preemptee job.

### To Configure `moab.cfg` for QoS Preemption

1. Set `GUARANTEEDPREEMPTION` to `TRUE`. (This causes Moab to lock `PREEMPTOR` jobs until `JOBRETRYTIME` expires.)
2. Make sure that `JOBNODEMATCHPOLICY` is *not* set to `EXACTNODE`, which is not currently supported for preemption (for more information, see [Testing and Troubleshooting Preemption](#)).

3. If it is not already, set `NODEACCESSPOLICY` to `SHARED`.
4. Set the `PREEMTPOLICY` policy type (for more information, see [PREEMTPOLICY Types](#)).
5. Set up `QFLAGS` to mark jobs as `PREEMPTTEE` (a lower-priority job that can be preempted by a higher-priority job), or as `PREEMPTOR` (a higher-priority job that can preempt a lower-priority job). As in the example:

```
QOSCFG[test1] QFLAGS=PREEMPTTEE MEMBERULIST=<user> PRIORITY=100
QOSCFG[test2] QFLAGS=PREEMPTOR MEMBERULIST=<user> PRIORITY=10000
```

For more information, see [Preemption Flags](#).

6. Make sure that the `PREEMPTTEE` job has a lower priority than the `PREEMPTOR` job. As in the example:

```
QOSCFG[test1] QFLAGS=PREEMPTTEE MEMBERULIST=<user> PRIORITY=100
QOSCFG[test2] QFLAGS=PREEMPTOR MEMBERULIST=<user> PRIORITY=10000
```

For example:

```
GUARANTEEDPREEMPTION TRUE
PREEMTPOLICY <policy>

QOSCFG[test1] QFLAGS=PREEMPTTEE MEMBERULIST=<user> PRIORITY=100
QOSCFG[test2] QFLAGS=PREEMPTOR MEMBERULIST=<user> PRIORITY=10000
```

---

## Related Topics

- [Chapter 21: Preemption - page 898](#)
- [Preemption Management](#)
- [21.2.2 Preemption Flags - page 917](#)
- [21.2.3 PREEMTPOLICY Types - page 918](#)
- [21.2.4 Simple Example of Preemption - page 919](#)
- [21.2.5 Testing and Troubleshooting Preemption - page 923](#)

## 21.2 Preemption Reference

In this section:

- [21.2.1 Manual Preemption Commands - page 916](#)
- [21.2.2 Preemption Flags - page 917](#)
- [21.2.3 PREEMTPOLICY Types - page 918](#)
- [21.2.4 Simple Example of Preemption - page 919](#)
- [21.2.5 Testing and Troubleshooting Preemption - page 923](#)

## 21.2.1 Manual Preemption Commands

You can use the [mjobctl](#) command to manually preempt jobs. The command can modify a job's execution state in the following ways:

Action	Flag	Details
<b>Cancel</b>	-c	Terminate job; remove from queue
<b>Checkpoint</b>	-C	Terminate and checkpoint job leaving job in queue
<b>Requeue</b>	-R	Terminate job; leave in queue
<b>Resume</b>	-r	Resume suspended job
<b>Start (execute)</b>	-x	Start idle job
<b>Suspend</b>	-s	Suspend active job

In general, users are allowed to suspend or terminate jobs they own. Administrators are allowed to suspend, terminate, resume, and execute any queued jobs.

---

### Related Topics

- [Chapter 21: Preemption - page 898](#)
- [21.2.5 Testing and Troubleshooting Preemption - page 923](#)



## 21.2.2 Preemption Flags

Using QoS, you can specify preemption rules and control access to preemption privileges. This allows you to increase system throughput, improve job response time for specific classes of jobs, or enable various political policies. You enable all policies by specifying some QoS credentials with the `QFLAGS` `PREEMPTTEE`, and others with `PREEMPTOR`.

PREEMPTTEE	
<b>Description</b>	Indicates that the job can be preempted by a higher-priority job.
<b>Use</b>	Use for lower-priority jobs that can be preempted.
<b>Notes</b>	<div style="border: 1px solid black; padding: 5px;"> <p><b>i</b> This might delay some node actions. When reprovisioning, the system job might expire before the provision action occurs; while the action will still occur, the job will not show it.</p> </div>
<b>Example</b>	<pre>QOSCFG[test1] QFLAGS=PREEMPTTEE MEMBERULIST=&lt;user&gt; PRIORITY=100</pre>

PREEMPTOR	
<b>Description</b>	Indicates that the job should take priority and preempt any <code>PREEMPTTEE</code> jobs.
<b>Use</b>	Use for jobs that need to take precedence over lower-priority jobs.
<b>Notes</b>	<div style="border: 1px solid black; padding: 5px;"> <p><b>i</b> <code>PREEMPTOR</code> jobs, either queued or running, must have a higher priority than <code>PREEMPTTEE</code> jobs.</p> <p>When you configure job as a <code>PREEMPTOR</code>, you should also increase its priority (for details, see <a href="#">PREEMPTPRIOJOBSELECTWEIGHT</a> and <a href="#">PREEMPTRTIMEWEIGHT</a>).</p> </div>
<b>Example</b>	<pre>QOSCFG[test2] QFLAGS=PREEMPTOR MEMBERULIST=&lt;user&gt; PRIORITY=10000</pre>

### Additional Preemptor and Preemptee Information

**i** Preemptor priority plays a big role in preemption. Generally, you should assign the preemptor job a higher priority value than any other queued jobs so that it will move to (or near to) the top of the eligible queue.

You can set the [RESERVATIONPOLICY](#) parameter to `NEVER`. With this configuration, preemptee jobs can start whenever idle resources become available. These jobs will be allowed to run until a preemptor job arrives, at which point the preemptee jobs are preempted, freeing the resource. This configuration allows near immediate resource access for the preemptor jobs. Using this approach, a cluster can maintain near 100% system utilization while still delivering excellent turnaround time to the most important jobs.

In environments where job checkpointing or job suspension incur significant overhead, you might want to constrain the rate at which job preemption is allowed. You can use the [JOBPREEMPTMINACTIVETIME](#) parameter to throttle job preemption. In essence, this parameter prevents a newly started or newly resumed job from being eligible for preemption until it has executed for a specified amount of time. Conversely, you can exclude jobs from preemption after they have run for a certain amount of time by using the [JOBPREEMPTMAXACTIVETIME](#) parameter.

---

### Related Topics

- [Chapter 21: Preemption - page 898](#)
- [21.1.6 Using QoS Preemption - page 914](#)
- [21.2.1 Manual Preemption Commands - page 916](#)
- [21.2.3 PREEMTPOLICY Types - page 918](#)
- [21.2.5 Testing and Troubleshooting Preemption - page 923](#)

## 21.2.3 PREEMTPOLICY Types

You can use the [PREEMTPOLICY](#) parameter to control how the scheduler preempts a job. This parameter enforces preemption using one of the following methods:

PREEMTPOLICY type	Description
<b>SUSPEND</b>	Causes active jobs to stop executing, but to remain in memory on the allocated compute nodes.
<b>CHECKPOINT</b>	Saves the current job state and either terminates or continues running the job. A checkpointed job can restart at any time and resume execution from its most recent checkpoint.

PREEMTPOLICY type	Description
<b>REQUEUE</b>	Terminates active jobs and returns them to the job queue in an idle state.
<b>CANCEL</b>	Cancels active jobs.

Each of these methods varies in the level of disruption to the job, `SUSPEND` being the least disruptive and `CANCEL` being the most disruptive.

Moab uses preemption escalation to free up resources. So for example, if the `PREEMTPOLICY` is set to `SUSPEND`, Moab uses this method if it is available; however, Moab will escalate it to something potentially more disruptive if necessary to preempt and free up resources.

---

### Related Topics

- [21.1.4 Suspending Jobs with Preemption - page 906](#)
- [21.1.2 Checkpointing Jobs with Preemption - page 902](#)
- [21.1.3 Requeuing Jobs with Preemption - page 904](#)
- [Canceling Jobs with Preemption](#)
- [Chapter 21: Preemption - page 898](#)
- [21.2.2 Preemption Flags - page 917](#)

## 21.2.4 Simple Example of Preemption

This section illustrates the process of setting up preemption on your system from beginning to end and contains examples of what actions to take and what you should see as you go.

In this topic:

- [21.2.4.A Scenario - page 920](#)
- [21.2.4.B Configuring moab.cfg - page 920](#)

[21.2.4.C Submitting a Job to the Preemptee QoS - page 921](#)

[21.2.4.D Submitting a Job to the Preemptor QoS - page 922](#)

## 21.2.4.A Scenario

### Example Scenario

For this basic setup example, we will have a user who can submit to either a 'test1' or 'test2' QoS. This example will use a `REQUEUE` preemption type.

We will go through three parts to set up this preemption:

- Configuring the `moab.cfg` file
- Submitting a job to the `PREEMPTEE` QoS
- Submitting a job to the `PREEMPTOR` QoS

Okay, let's get started!

## 21.2.4.B Configuring `moab.cfg`

First, you will need to make some configurations to the `moab.cfg` file.

1. Set `GUARANTEEDPREEMPTION` to `TRUE` (this causes Moab to lock `PREEMPTOR` jobs until `JOBRETRYTIME` expires).
2. Make sure that `JOBNODEMATCHPOLICY` is *not* set to `EXACTNODE`, which is not currently supported for preemption (for more information, see [Testing and Troubleshooting Preemption](#)).
3. Set the `PREEMPTPOLICY` type. In this example, `PREEMPTPOLICY` is set to `REQUEUE`. For more information, see [PREEMPTPOLICY Types](#).
4. Set up `QFLAGS` to mark jobs as `PREEMPTEE` (a lower-priority job that can be preempted by a higher-priority job), or as `PREEMPTOR` (a higher-priority job that can preempt a lower-priority job). For more information, see [Preemption Flags](#).



For this example, we also set `JOBFLAGS=RESTARTABLE` (because this example uses `REQUEUE`). For more information, see [Requeuing Jobs with Preemption](#).

5. Make sure that the `PREEMPTEE` job has a lower priority than the `PREEMPTOR` job.

Here is an example of how that would all look in a `moab.cfg` file (text marked **bold** for emphasis):

```

GUARANTEEDPREEMPTION TRUE
#should not be JOBNODEMATCHPOLICY EXACTNODE as it causes problems when starting jobs

PREEMTPOLICY REQUEUE

QOSCFG[test1] QFLAGS=PREEMPTEE JOBFLAGS=RESTARTABLE MEMBERULIST=john PRIORITY=100
QOSCFG[test2] QFLAGS=PREEMPTOR MEMBERULIST=john PRIORITY=1000

```

Now you can submit a job to the preemptee QoS (`test1`).

### 21.2.4.C Submitting a Job to the Preemptee QoS

Let's submit a job to the preemptee QoS (`test1`), requesting all processor cores in the cluster:

```
[john@g06]# echo sleep 600 | msub -l walltime=600 -l qos=test1 -l procs=128
```

Take a look at the `showq` and `checkjob` output:

```

Moab.1
[john@g06]# showq

active jobs-----
JOBID      USERNAME    STATE      PROCS      REMAINING    STARTTIME
Moab.1     john        Running    128        00:09:59     Wed Nov 9 15:56:33

1 active job      128 of 128 processors in use by local jobs (100.00%)
                    2 of 2 nodes active (100.00%)

eligible jobs-----
JOBID      USERNAME    STATE      PROCS      WCLIMIT      QUEUEETIME

0 eligible jobs

blocked jobs-----
JOBID      USERNAME    STATE      PROCS      WCLIMIT      QUEUEETIME

0 blocked jobs

Total job: 1

```

```

[john@g06]# checkjob Moab.1
job Moab.1

State: Running
Creds: user:john group:john qos:test1
WallTime: 00:00:00 of 00:10:00
SubmitTime: Wed Nov 9 15:56:33
(Time Queued Total: 00:00:00 Eligible: 00:00:00)

StartTime: Wed Nov 9 15:56:33
Total Requested Tasks: 128

Req[0] TaskCount: 128 Partition: licenses

```

```

Allocated Nodes:
node[01-02]*64

IWD: /opt/native/
SubmitDir: /opt/native/
Executable: /opt/native/spool/moab.job.zOyf1N

StartCount: 1
Flags: RESTARTABLE,PREEMPTEE,GLOBALQUEUE,PROCSPECIFIED
Attr: PREEMPTEE
StartPriority: 100
Reservation 'Moab.1' (-00:00:03 -> 00:09:57 Duration: 00:10:00

```

### 21.2.4.D Submitting a Job to the Preemptor QoS

Now we will submit a preemptor QoS job (`test2`) to preempt the first job (`test1`):

```
[john@g06]# echo sleep 600 | msub -l walltime=600 -l qos=test2 -l procs=128
```

Examine the following output for `showq` and `checkjob`:

```

Moab.2
[john@g06]# showq

active jobs-----
JOBID      USERNAME    STATE      PROCS      REMAINING   STARTTIME
Moab.2    john        Running   128        00:09:59    Wed Nov 9 15:56:47

1 active job 128 of 128 processors in use by local jobs (100.00%)
  2 of 2 nodes active (100.00%)

eligible jobs-----
JOBID      USERNAME    STATE      PROCS      WCLIMIT     QUEUETIME
Moab.1    john        Idle       128        00:10:00    Wed Nov 9 15:56:33

1 eligible job

blocked jobs-----
JOBID      USERNAME    STATE      PROCS      WCLIMIT     QUEUETIME

0 blocked jobs

Total jobs: 2

```

*Note that the preemptor job (Moab.2) moved to **Running**, while the preemptee job (Moab.1) was queued.*

```

[john@g06]# checkjob Moab.2
job Moab.2

State: Running
Creds: user:john group:john qos:test2
WallTime: 00:02:04 of 00:10:00
SubmitTime: Wed Nov 9 15:56:46
(Time Queued Total: 00:00:01 Eligible: 00:00:00)

StartTime: Wed Nov 9 15:56:47
Total Requested Tasks: 128

```

```

Req[0] TaskCount: 128 Partition: licenses
NodeCount: 2

Allocated Nodes:
node[01-02]*64

IWD: /opt/native/
SubmitDir: /opt/native/
Executable: /opt/native/spool/moab.job.ELoX5Q

StartCount: 1
Flags: HASPREEMPTED, PREEMPTOR, GLOBALQUEUE, PROCSPECIFIED
StartPriority: 10000
Reservation 'Moab.2' (-00:02:21 -> 00:07:39 Duration: 00:10:00)

```

*Note the flag, **HASPREEMPTED**. **HASPREEMPTED** is set when the **PREEMPTOR** job has preempted the **PREEMPTEE** job. Also note that the preemptor job priority plays a very big role in preemption. Generally, you should assign the preemptor a higher priority than any other queued jobs so that it will move to (or near to) the top of the eligible queue.*

---

## Related Topics

- [Chapter 21: Preemption - page 898](#)
- [21.2.2 Preemption Flags - page 917](#)
- [21.2.3 PREEMTPOLICY Types - page 918](#)
- [21.2.1 Manual Preemption Commands - page 916](#)
- [21.2.5 Testing and Troubleshooting Preemption - page 923](#)

## 21.2.5 Testing and Troubleshooting Preemption

There are multiple steps associated with setting up a working preemption policy. With preemption, issues arise because it appears that Moab is not allowing preemptor jobs to preempt preemptee jobs in the right way. To diagnose this, use the following checklist:

Verify that preemptor jobs are marked with the `PREEMPTOR` flag. Verify with `checkjob <JOBID> | grep Flags`.

Verify that preemptee jobs are marked with the `PREEMPTEE` flag. Verify with `checkjob <JOBID> | grep Flags`.

Verify that the start priority of the preemptor job is higher than the priority of the preemptee job. Verify with `checkjob <JOBID> | grep Priority`.

Verify that the resources allocated to the preemptee job match those requested by the preemptor job.
Verify that there are no policies preventing preemption from occurring. Verify with <code>checkjob -v -n &lt;NODEID&gt; &lt;JOBID&gt;</code> .
Verify that the <code>PREEMTPOLICY</code> parameter is properly set. See <a href="#">PREEMTPOLICY Types</a> .
Verify that the preemptee job is properly marked as restartable, suspendable, or checkpointable. Verify with <code>checkjob &lt;JOBID&gt;   grep Flags</code> .
Verify that <code>GUARANTEEDPREEMPTION</code> is set to <code>TRUE</code> .
If <code>JOBNODEMATCHPOLICY</code> is set to <code>EXACTNODE</code> , verify <code>PREEMPTIONALGORITHM</code> is set to <code>PREEMPTORCENTRIC</code> .
Verify that <code>NODEACCESSPOLICY</code> is <i>not</i> set to <code>SINGLEUSER</code> . <code>SHARED</code> is recommended.
Verify that <code>BACKFILLPOLICY</code> is set to <code>FIRSTFIT</code> .
Verify that the resource manager is properly responding to preemption requests. (Use <code>mdiag -R</code> .)
If there is a resource manager level race condition, verify that Moab is properly holding target resources. Verify with <code>mdiag -S</code> and set <code>RESERVATIONRETRYTIME</code> if needed.

---

## Related Topics

- [Chapter 21: Preemption - page 898](#)
- [Quality of Service \(QoS\) Facilities](#)
- [Managing QoS Access](#)
- [JOBMAXPREEMTPERITERATION](#)
- [17.3.7 Trigger Components - page 824](#)
- [9.4 Checkpoint/Restart Facilities - page 561](#)
- [ENABLEFSVIOLATIONPREEMPTION](#)
- [PREEMTPRIOJOBSELECTWEIGHT](#)
- [PREEMPTSEARCHDEPTH](#)



- **USAGEEXECUTIONTIMEWEIGHT** (control priority of suspended jobs)
- **IGNOREPREEMPTTEEPRIORITY** (relative job priority is ignored in preemption decisions)
- **DISABLESAMECREDPREEMPTION** (jobs cannot preempt other jobs with the same credential)
- **PREEMPTRTIMEWEIGHT** (add remaining time of jobs to preemption calculation)

## Chapter 22: About Job Templates

A Moab job template is a set of preconfigured settings, attributes, and resources that Moab applies to jobs that match certain criteria or to which you manually apply it. They perform three primary functions:

1. They generically match and categorize jobs.
2. They set arbitrary default or forced attributes for certain jobs.
3. They generate workflows that create and maintain user-requested services in a cloud environment.

You can use job templates in many aspects of scheduling, including [Peer-Based Grid](#) usage policies. Job templates are defined using the [JOB\\_CFG](#) configuration parameter.

Two methods exist for applying job templates to jobs. You can use the [JOBMATCHCFG](#) parameter to mark a template that contains the criteria a job must meet for eligibility and another template as the one to be applied to the job if it is eligible. This allows you to automate the use of templates. For example, to force all interactive jobs to run on a certain set of nodes, you can set one template (the criteria template) to have the `interactive` flag, then give the other template the desired host list. You can also apply a template directly to a job at submission if that ability is enabled for that template.

In this section:

- [22.1 Job Template Tasks - page 926](#)
- [22.2 Job Template Reference - page 931](#)

### 22.1 Job Template Tasks

In this section:

- [22.1.1 Creating Job Templates - page 927](#)
- [22.1.2 Viewing Job Templates - page 928](#)
- [22.1.3 Applying Templates Based on Job Attributes - page 928](#)
- [22.1.4 Requesting Job Templates Directly - page 929](#)
- [22.1.5 Creating Workflows with Job Templates - page 930](#)

## 22.1.1 Creating Job Templates

Job templates are created in the Moab configure file using the `JOBCFG` parameter.

### To Create a Job Template

1. Open `moab.cfg`. Add the `JOBCFG` parameter and give the new job template a unique name:

```
JOBCFG[newtemplate]
```

2. Configure any desired attributes (see [Job Template Extension Attributes](#)). Some of the important attributes include:

- **FLAGS** - Lets you specify any job flags that should be applied.

```
JOBCFG[newtemplate] FLAGS=SUSPENDABLE
```

*When Moab applies `newtemplate` to a job, the job is marked as suspendable.*

- **SELECT** - Lets you apply the template directly at job submission.

```
JOBCFG[newtemplate] FLAGS=SUSPENDABLE SELECT=TRUE
```

*When you submit a job via `msub`, you can specify that your job has `newtemplate` applied to it. When Moab applies the template to a job, that job is marked as suspendable.*

- **TEMPLATEDPEND** - Lets you create dependencies when you create a job template workflow (see [Creating Workflows with Job Templates](#)).

```
JOBCFG[newtemplate] FLAGS=SUSPENDABLE SELECT=TRUE TEMPLATEDPEND=AFTER:job1.pre
```

*When Moab applies `newtemplate` to a job, the job cannot run until job `job1.pre` has finished running; the job is also marked as suspendable. You can specify that Moab apply this template to a job as you submit it.*

3. If you want to automate job template application, see [Applying Templates Based on Job Attributes](#) for instructions. If you want to apply the template manually on job submission, see [Requesting Job Templates Directly](#) for instructions.

---

### Related Topics

- [22.2.1 Job Template Extension Attributes](#) - page 931
- [22.2.3 Job Template Examples](#) - page 945

## 22.1.2 Viewing Job Templates

### To View a Job Template

Run the `mdiag -j` command with the `policy` flag:

```
> mdiag -j --flags=policy --blocking
```

Moab returns a list of job templates configured in `moab.cfg`.

## 22.1.3 Applying Templates Based on Job Attributes

The `JOBMATCHCFG` parameter allows you to establish relationships between a number of job templates. `JMAX` and `JMIN` function as filters to determine whether a job is eligible for a subsequent template to be applied to the job. If a job is eligible, `JDEF` and `JSET` templates apply attributes to the job. See [Job Template Extension Attributes](#) for more information about the `JOBMATCHCFG` attributes. The table on that page indicates which job template types are compatible with which job template extension attributes.

**i** `JSETS` and `JDEFS` have only been tested using `msub` as the job submission command.

### To Apply a Job Template Based on Job Attributes

1. In the Moab configuration file, create a job template with a set of criteria that a job must meet in order for Moab to apply the template. In the following example, Moab will apply a template to all interactive jobs, so the first template sets the `interactive` flag.

```
JOBCFG[inter.min] FLAGS=interactive
```

2. Create the job template that Moab should apply to the job if it meets the requirements set in the first template. In this example, Moab ignores all configured policies, so the second template sets the `ignpolicies` flag.

```
JOBCFG[inter.set] FLAGS=ignpolicies
```

3. Use the `JOBMATCHCFG` parameter and its `JMAX` or `JMIN` (specify the template specifying maximum or minimum requirements) and `JDEF` or `JSET` (specify the template to be applied) attributes to demonstrate the relationship between the two templates (see [Job Template Matching Attributes](#) for more information). In this case, all interactive jobs ignore policies; in other words, if a submitted job has at least the

`inter.min` template settings, Moab applies the `inter.set` template settings to the job.

```
JOBMATCHCFG[interactive] JMIN=inter.min JSET=inter.set
```

*Moab applies the `inter.set` template to all jobs with the `interactive` flag set, causing them to ignore Moab's configured policies.*

- To control which job template is applied to a job that matches multiple templates, use `FLAGS=BREAK`. Job templates are processed in the order they are listed in the configuration file and using the `BREAK` flag causes Moab to stop evaluating `JOBMATCHCFG` entries that occur after the current match.

```
JOBMATCHCFG[small] JMIN=small.min JMAX=small.max JSET.set=small.set FLAGS=BREAK
JOBMATCHCFG[large] JMIN=large.min JMAX=large.max JSET=large.set
```

*In this case, the large template would not be applied when a job matches both the small and large templates. The small template matches first, and because of `FLAGS=BREAK`, Moab stops evaluating further `JOBMATCHCFG` entries for the job.*

---

## Related Topics

- [22.1.4 Requesting Job Templates Directly - page 929](#)
- [22.2.3 Job Template Examples - page 945](#)

## 22.1.4 Requesting Job Templates Directly

When a job template has its `SELECT` attribute set to `TRUE`, you can request that template directly on job submission.

### To Directly Request Job Templates

- Set the `SELECT` attribute on the template in `moab.cfg`:

```
JOBCFG[medium.set] NODESET=ONEOF:FEATURE:fast,slow SELECT=true
```

- Submit a job with a resource list (`msub -l`), requesting the template using the format `template=<templateName>`:

```
> msub -l template=medium.set
```

Moab creates a job with the `medium.set` job template created in step 1.

**i** Attributes set in the template are evaluated as if they were part of the job submission. They are still subject to all of the same ACLs and policies.

## Related Topics

- [22.1.3 Applying Templates Based on Job Attributes - page 928](#)

## 22.1.5 Creating Workflows with Job Templates

Moab can create workflows from individual jobs using job templates.

### To Build a Workflow with Job Templates

1. Create the jobs in the workflow using the `JOBCFG` parameter (see [Creating Job Templates](#) for more information). It might be useful to add the `PURGEONSUCCESSONLY` flag to your setup or destroy jobs; it will allow you to restart the jobs easily if they fail. Specify the order in which they should run with the `TEMPLATEDPEND` attribute. See the [Job dependency syntax table](#) for a list of valid dependency options.

```
JOBCFG[setup.pre]    TASKS=2 WCLIMIT=00:01:00 SELECT=TRUE
EXEC=/usr/tools/setup.pre.sh
JOBCFG[setup.pre2]  TEMPLATEDPEND=AFTER:setup.pre SELECT=TRUE
EXEC=/nfs/tools/setup.pre2.sh
JOBCFG[engineering] TEMPLATEDPEND=AFTER:setup.pre2
```

*When Moab applies the engineering template to a qualifying job, the job will not run until template job setup.pre and then setup.pre2 are created from the specified EXEC strings and finish running.*

**i** The Moab naming convention for jobs created with job templates is `<moabId>.<templateName>`. By default, when Moab submits jobs to only one resource manager, the job IDs are synchronized with the resource manager's job IDs. You can use the parameter `USEMOABJOBID` so that a template-created job is easily associated with its parent job (such as `moab.1, moab.1.setup.pre`).

2. Create the job template that will act as the criteria a job must meet for Moab to apply the engineering template. In this situation, the job must be submitted with the account name engineering:

```
JOBCFG[engineering.match] ACCOUNT=engineering
```

3. Create the `JOBMATCHCFG` configuration to tell Moab that when a job matches the `engineering.match` template, it should apply the engineering template:

```
JOBMATCHCFG[engineering.job] JMIN=engineering.match JSET=engineering
```

**Related Topics**

- [22.2.1 Job Template Extension Attributes - page 931](#)
- [22.2.4 Job Template Workflow Examples - page 946](#)
- [22.1.1 Creating Job Templates - page 927](#)

## 22.2 Job Template Reference

In this section:

- [22.2.1 Job Template Extension Attributes - page 931](#)
- [22.2.2 Job Template Matching Attributes - page 944](#)
- [22.2.3 Job Template Examples - page 945](#)
- [22.2.4 Job Template Workflow Examples - page 946](#)

### 22.2.1 Job Template Extension Attributes

When creating a job template, you can use any attribute acceptable within the [WIKI](#) workload query data format. In addition, job templates can use any of the extension attributes in the following table. Note that the Template type (JMIN, JMAX, JDEF, JSET) row indicates compatibility with the associated attribute (see [Applying Templates Based on Job Attributes](#) for more information).

**i** Attributes set in a template are evaluated as if they were part of the original job submission. Their jobs are still subject to all the same ACLs and policies.

**Attributes**

ACCOUNT	GNAME	PRIORITY	TASKS
CLASS	GRES	PRIORITYF	TASKPERNODE
CPUCLOCK	GROUP	QOS	TEMPLATEDPEND
CPULIMIT	MEM	RARCH	UNAME
DESCRIPTION	NODEACCESSPOLICY	RFEATURES	USER

DPROCS ENV EXEC FLAGS	NODES NODESET PARTITION PREF	RM ROPSYS SELECT SYSTEMJOBTYPE	VARIABLE WCLIMIT
--------------------------------	---------------------------------------	---	---------------------

## ACCOUNT

<b>Format</b>	<ACCOUNT>[,<ACCOUNT>]...
<b>Template type</b>	JMIN JDEF JSET
<b>Description</b>	Account credentials associated with job. This is used for job template matching.
<b>Example</b>	<pre> JOBcfg[public] FLAGS=preemptee JOBcfg[public.min] ACCOUNT=public_acct JOBMATCHCFG[public] JMIN=public.min JSET=public </pre>

## CLASS

<b>Format</b>	<CLASS>[,<CLASS>]...
<b>Template type</b>	JMIN JDEF JSET
<b>Description</b>	Class credentials associated with job. This is used for job template matching.
<b>Example</b>	<pre> JOBcfg[night] FLAGS=preemptor JOBcfg[night.min] CLASS=night_class JOBMATCHCFG[night] JMIN=night.min JSET=night </pre>

## CPUCLOCK

<b>Format</b>	<STRING>
<b>Template type</b>	JMIN JMAX



CPUCLOCK	
	JSET
<b>Description</b>	CPU clock frequency for all CPUs of a job. For more information, see <a href="#">CPUCLOCK</a> . The job template extension overrides the job script.
<b>Example</b>	<pre> JOBCFG[slow] SELECT=TRUE cpuclock=1400 JOBCFG[fast] SELECT=TRUE cpuclock=3200  JOBCFG[cpu.min] CPUCLOCK=1000 JOBCFG[cpu.max] CPUCLOCK=2000 JOBCFG[cpu.set] CPUCLOCK=1500  JOBMATCHCFG[cpu] JMIN=cpu.min JMAX=cpu.max JSET=cpu.set </pre>

CPULIMIT	
<b>Format</b>	[[[DD:]HH:]MM:]SS
<b>Template type</b>	JMIN JMAX JDEF JSET
<b>Description</b>	Maximum amount of CPU time used by all processes in the job.
<b>Example</b>	<pre> JOBCFG[job.min] CPULIMIT=1:00:00:00 JOBCFG[job.max] CPULIMIT=2:00:00:00 </pre>

DESCRIPTION	
<b>Format</b>	<STRING>
<b>Template type</b>	JMAX JDEF
<b>Description</b>	Description of the job. When you run the <a href="#">checkjob</a> command, the description appears as Description.
<b>Example</b>	<pre> JOBCFG[webdb] DESCRIPTION="Template job" </pre>


DPROCS	
<b>Format</b>	<INTEGER>
<b>Template type</b>	JMIN JMAX JSET
<b>Description</b>	Number of processors dedicated per task. The default is 1.
<b>Example</b>	<pre>JOBCFG[job.min] DPROCS=2 JOBCFG[job.max] DPROCS=4</pre>

ENV	
<b>Format</b>	<STRING>
<b>Template type</b>	JSET
<b>Description</b>	Adds the specified job environment variables to the job.
<b>Example</b>	<pre>JOBCFG[container] ENV=PBS_CONTAINERINFO=centos</pre>


EXEC	
<b>Format</b>	<STRING>
<b>Template type</b>	JSET
<b>Description</b>	Specifies what the job runs, regardless of what the user set.
<b>Example</b>	<pre>JOBCFG[setup.pre] EXEC=nfs/tools/setup.pre.sh</pre>

FLAGS	
<b>Format</b>	<JOBFLAG>[,<JOBFLAG>]...
<b>Template type</b>	JADMIN

FLAGS	
	JDEF JSET
<b>Description</b>	One or more legal job flag values.
<b>Example</b>	<code>JOBCFG[webdb] FLAGS=NORMSTART</code>

GNAME	
<b>Format</b>	<STRING>
<b>Template type</b>	JDEF JSET
<b>Description</b>	Group credential associated with job.
<b>Example</b>	<code>JOBCFG[webserv] GNAME=service</code>
	 For matching the group, see the <a href="#">GROUP</a> attribute.

GRES	
<b>Format</b>	<genericResource>[:<COUNT>][,<genericResource>[:<COUNT>]]...
<b>Template type</b>	JMIN JMAX JDEF
<b>Description</b>	Consumable generic attributes associated with individual nodes or the special pseudo-node <a href="#">global</a> , which provides shared cluster (floating) consumable resources. Use the <a href="#">NODECFG</a> parameter to configure such resources.
<b>Example</b>	<code>JOBCFG[gres.set] GRES=abaqus:2</code>
	<i>In this example, the <code>gres.set</code> template applies two Abaqus licenses per task to a matched job.</i>

GROUP	
<b>Format</b>	<GROUP>[,<GROUP>]...
<b>Template type</b>	JMIN
<b>Description</b>	Group credentials associated with job. This is used for job template matching.
<b>Example</b>	<pre>JOBCFG[webserv] GROUP=service</pre> <p> For information about setting the group, see the <a href="#">GNAME</a> attribute.</p>

MEM	
<b>Format</b>	<INTEGER>
<b>Template type</b>	JMIN JMAX JDEF JSET
<b>Description</b>	Maximum amount of physical memory per task used by the job in megabytes. You can optionally specify other units with your integer (300kb or 2gb, for example). See 'Requesting Resources' in the <i>Torque Administrator Guide</i> for more information.
<b>Example</b>	<pre>JOBCFG[smalljobs] MEM=25</pre>

NODEACCESSPOLICY	
<b>Format</b>	One of the following: <a href="#">SHARED</a> , <a href="#">SHAREDONLY</a> , <a href="#">SINGLEJOB</a> , <a href="#">SINGLETASK</a> , <a href="#">SINGLEUSER</a> , or <a href="#">UNIQUEUSER</a>
<b>Template type</b>	JDEF JSET
<b>Description</b>	Specifies how node resources will be shared by a job. See the <a href="#">Node Access Policies</a> for more information.

## NODEACCESSPOLICY

<b>Example</b>	<code>JOB_CFG[serverapp] NODEACCESSPOLICY=SINGLEJOB</code>
----------------	--

## NODES

<b>Format</b>	<INTEGER>
<b>Template type</b>	JMIN JMAX JSET
<b>Description</b>	<p>Number of nodes required by the job. The default is 1. See <a href="#">Node Definition</a> for more information.</p> <p>When using JSET:</p> <ul style="list-style-type: none"> <li>• If the taskcount of the job is <i>less</i> than the NODES value, Moab will modify the taskcount to match the NODES value.</li> <li>• If the taskcount of the job is <i>greater</i> than the NODES value, Moab will attempt to evenly divide the tasks. If the taskcount is not evenly divisible by the NODES value, the job is rejected.</li> </ul>
<b>Example</b>	<code>JOB_CFG[job.min] NODES=2</code> <code>JOB_CFG[job.max] NODES=4</code>

## NODESET


<b>Format</b>	<STRING>
<b>Template type</b>	JSET
<b>Description</b>	See <a href="#">Node Set Overview</a> for more information.
<b>Example</b>	<code>JOB_CFG[medium.set]</code> <code>NODESET=ONEOF:FEATURE:fast,slow</code>

## PARTITION

<b>Format</b>	<PARTITION>[:<PARTITION>]...
---------------	------------------------------

PARTITION	
<b>Template type</b>	JMIN JDEF JSET
<b>Description</b>	The partition (or partitions) where a job must run.
<b>Example</b>	<code>JOBCFG[meis] PARTITION=math:geology</code>

PREF	
<b>Format</b>	<FEATURE>[,<FEATURE>]...
<b>Template type</b>	JDEF JSET
<b>Description</b>	Specifies which node features are preferred by the job and should be allocated if available. See <a href="#">PREF</a> for more information.
<b>Example</b>	<code>JOBCFG[meis] PREF=bigmem</code>

PRIORITY	
<b>Format</b>	<INTEGER>
<b>Template type</b>	JMAX JDEF
<b>Description</b>	System job priority. <div style="border: 1px solid #0070C0; border-radius: 10px; padding: 5px; margin-top: 10px;"> <b>PRIORITY</b> works only as a default setting and not as an override (JSET) setting.</div>
<b>Example</b>	<code>JOBCFG[meis] PRIORITY=25000</code>

PRIORITYF	
<b>Format</b>	PRIORITYF='<VALUE>'
<b>Template type</b>	JSET
<b>Description</b>	Applicable only when using NODEALLOCATIONPOLICY with the PRIORITY format. Lets you change the priority function used to allocate nodes for the job. See <a href="#">Node Allocation Policies</a> for available PRIORITY values.
<b>Example</b>	<pre>JOBCFG[limit.set] PRIORITYF='NODEINDEX'</pre>

QOS	
<b>Format</b>	<QOS>[,<QOS>]...
<b>Template type</b>	JMIN JDEF JSET
<b>Description</b>	QoS credentials associated with job. This is used for job template matching.
<b>Example</b>	<pre>JOBCFG[admin] RFEATURES=bigmem JOBCFG[admin.min] QOS=admin_qos JOBMATCHCFG[admin] JMIN=admin.min JSET=admin</pre>

RARCH	
<b>Format</b>	<STRING>
<b>Template type</b>	JSET
<b>Description</b>	Architecture required by job.
<b>Example</b>	<pre>JOBCFG[servapp] RARCH=i386</pre>

RFEATURES	
<b>Format</b>	<FEATURE>[,<FEATURE>]...
<b>Template type</b>	JMIN JDEF JSET
<b>Description</b>	List of features required by job.
<b>Example</b>	<pre>JOBCFG[servapp] RFEATURES=fast,bigmem</pre>

RM	
<b>Format</b>	<STRING>
<b>Template type</b>	JDEF JSET
<b>Description</b>	Destination resource manager to be associated with job.
<b>Example</b>	<pre>JOBCFG[webdb] RM=torque</pre>

ROPSYS	
<b>Format</b>	<STRING>
<b>Template type</b>	JDEF JSET
<b>Description</b>	Operating system required by job.
<b>Example</b>	<pre>JOBCFG[test.set] ROPSYS=windows</pre>

SELECT	
<b>Format</b>	<BOOLEAN> : TRUE   FALSE



SELECT	
<b>Description</b>	Job template can be directly requested by job at submission.
<b>Example</b>	<pre>JOBCFG[servapp] SELECT=TRUE</pre>

SYSTEMJOBTYPE	
<b>Template type</b>	JMIN
<b>Description</b>	System job type (example: vmcreate).
<b>Example</b>	<pre>JOBCFG[vmcreate.min] SYSTEMJOBTYPE=vmcreate JOBCFG[vmcreate.set] TRIGGER=atype=reserve,action="00:05:00",etype=end JOBMATCHCFG[vmcreate] JMIN=vmcreate.min JSET=vmcreate.set</pre>

TASKS	
<b>Format</b>	<INTEGER>
<b>Template type</b>	JMIN JMAX JSET
<b>Description</b>	Number of tasks required by job. The default is 1. See <a href="#">Task Definition</a> for more information.
<b>Example</b>	<pre>JOBCFG[job.min] TASKS=4 JOBCFG[job.max] TASKS=8</pre>


TASKPERNODE	
<b>Format</b>	<INTEGER>
<b>Template type</b>	JMIN JMAX JDEF
<b>Description</b>	Exact number of tasks required per node. The default is 0.


TASKPERNODE	
	<p><b>i</b> TASKPERNODE works only as a default setting and not as an override (JSET) setting.</p>
<b>Example</b>	<pre>JOBCFG[job.min] TASKPERNODE=2 JOBCFG[job.max] TASKPERNODE=4</pre>

TEMPLATEDEPEND	
<b>Format</b>	<TYPE>:<TEMPLATE_NAME>
<b>Description</b>	<p>Create another job from the &lt;TEMPLATE_NAME&gt; job template, on which any jobs using this template will depend. This is used for dynamically creating workflows. See <a href="#">Job Dependencies</a> for more information.</p> <p><b>i</b> SYNCWITH only supports one dependency with Torque as the resource manager.</p>
<b>Example</b>	<pre>JOBCFG[engineering] TEMPLATEDEPEND=AFTER:setup.pre JOBCFG[setup.pre] SELECT=TRUE EXEC=/tools/setup.pre.sh</pre>

UNAME	
<b>Format</b>	<STRING>
<b>Default</b>	JDEF JSET
<b>Description</b>	User credential associated with job.
<b>Example</b>	<pre>JOBCFG[webserv] UNAME=service</pre> <p><b>i</b> For matching the user, see the <a href="#">USER</a> attribute.</p>

USER	
<b>Format</b>	<USER>[,<USER>]...

USER	
<b>Template type</b>	JMIN JMAX
<b>Description</b>	User credentials associated with job.
<b>Example</b>	<pre>JOBCFG[webserv] USER=service</pre> <p> For setting the user, see the <a href="#">UNAME</a> attribute.</p>

VARIABLE	
<b>Format</b>	<NAME>[:<VAL>]
<b>Template type</b>	JMIN JSET
<b>Description</b>	Variables attached to the job template.
<b>Example</b>	<pre>JOBCFG[this] VARIABLE=var1:1 VARIABLE=var2:1</pre> <p> Variables are set upon successful completion of the job.</p>

WCLIMIT	
<b>Format</b>	[[HH:]MM:]SS
<b>Template type</b>	JMIN JMAX JDEF JSET
<b>Description</b>	Walltime required by job. The default is 8640000 (100 days).
<b>Example</b>	<pre>JOBCFG[job.min] WCLIMIT=2:00:00 JOBCFG[job.max] WCLIMIT=12:00:00</pre>

---

**Related Topics**

- [22.2.3 Job Template Examples - page 945](#)
- [22.1.1 Creating Job Templates - page 927](#)

## 22.2.2 Job Template Matching Attributes

The `JOBMATCHCFG` parameter allows you to establish relationships between a number of job templates. The table in [Job Template Extension Attributes](#) indicates which job template types are compatible with which job template extension attributes. The following types of templates can be specified with the `JOBMATCHCFG` parameter:

Attribute	Description
<b>JMAX</b>	<p>A potential job is rejected if it has matching attributes set or has resource requests that exceed those specified in this template.</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> For <code>JMAX</code>, a job template can specify only positive non-zero numbers as maximum limits for generic resources. If a job requests a generic resource that is not limited by the template, then the template can still be used.</p> </div>
<b>JMIN</b>	A potential job is rejected if it does not have matching attributes set or has resource requests that do not meet or exceed those specified in this template.
<b>JDEF</b>	A matching job has the specified attributes set as defaults but all values can be overridden by the user if the matching attribute is explicitly set at job submission time.
<b>JSET</b>	A matching job has the specified attributes forced to these values and these values override any values specified by the submitter at job submission time.
<b>JSTAT</b>	A matching job has its usage statistics reported into this template.

---

**Related Topics**

- [22.2.1 Job Template Extension Attributes - page 931](#)
- [22.2.3 Job Template Examples - page 945](#)
- [22.1.3 Applying Templates Based on Job Attributes - page 928](#)

## 22.2.3 Job Template Examples

Job templates can be used for a wide range of purposes including enabling automated learning, setting up custom application environments, imposing special account constraints, and applying group default settings. The following examples highlight some of these uses:

### Example 22-1: Setting up Application-Specific Environments

```
JOBCFG[xxx] EXEC=*app* JOBPROLOG=/usr/local/appprolog.x
```

### Example 22-2: Applying job preferences and defaults

```
JOBCFG[xxx] CLASS=appq EXEC=*app* PREF=clearspeed
NODEALLOCATIONPOLICY PRIORITY
NODECFG[DEFAULT] PRIORITYF=5.0*PREF
```

### Example 22-3: Applying Resource Constraints to Fuzzy Collections

In the following example, a job template match is set up. Using the [JOBMATCHCFG](#) parameter, Moab is configured to apply all attributes of the `inter.set` job template to all jobs that match the constraints of the `inter.min` job template. In this example, all interactive jobs are assigned the `ignpolicies` flag that allows them to ignore active, idle, system, and partition level policies. Interactive jobs are also locked into the test standing reservation and therefore only allowed to run on the associated nodes.

```
# limit all users to a total of two non-interactive jobs
USERCFG[DEFAULT] MAXJOB=2
SRCFG[test] DESCRIPTION="compute pool for interactive and short duration jobs"
SRCFG[test] JOBATTRLIST=INTERACTIVE
SRCFG[test] MAXTIME=1:00:00
SRCFG[test] HOSTLIST=R:atl[16-63]
JOBCFG[inter.min] FLAGS=interactive
JOBCFG[inter.set] FLAGS=ignpolicies
JOBMATCHCFG[interactive] JMIN=inter.min JSET=inter.set
```

### Example 22-4: Resource Manager Templates

In the following example, interactive jobs are not allowed to enter through this resource manager and any job that does route in from this resource manager interface has the `preemptee` flag set.

```
JOBCFG[no_inter] FLAGS=interactive
JOBCFG[preempt_job] FLAGS=preemptee
RMCFG[gridA.in] MAX.JOB=no_inter SET.JOB=preempt_job
```

## Related Topics

- [22.2.1 Job Template Extension Attributes - page 931](#)
- [22.2.4 Job Template Workflow Examples - page 946](#)
- [22.1.1 Creating Job Templates - page 927](#)

## 22.2.4 Job Template Workflow Examples

### Example 22-5: A Workflow with Multiple Dependencies

In this example the job will depend on the completion of two other jobs Moab creates. Both jobs execute at the same time.

```
# Engineering2
JOBCFG[engineering2] TEMPLATEDEPEND=AFTER:engineering2.pre2
TEMPLATEDPEND=AFTER:engineering2.pre
JOBCFG[engineering2.pre2] TASKS=2 WCLIMIT=00:01:00 SELECT=TRUE
EXEC=/usr/tools/engineering2.pre2.sh
JOBCFG[engineering2.pre] TASKS=2 WCLIMIT=00:01:00 SELECT=TRUE
EXEC=/usr/tools/engineering2.pre.sh
JOBCFG[engineering2.match] ACCOUNT=engineering2
JOBMATCHCFG[engineering2.job] JMIN=engineering2.match JSET=engineering2
```

### Example 22-6: Jobs that Run After the Submission Job

Three additional jobs are created that depend on the submitted job:

```
# Workflow 2
JOBCFG[workflow2] TEMPLATEDEPEND=BEFORE:workflow2.post1
TEMPLATEDPEND=BEFORE:workflow2.post2 TEMPLATEDPEND=BEFORE:workflow2.post3
JOBCFG[workflow2.post1] TASKS=2 WCLIMIT=00:01:00 SELECT=TRUE
EXEC=/usr/tools/workflow2.post1.sh
JOBCFG[workflow2.post2] TASKS=2 WCLIMIT=00:01:00 SELECT=TRUE
EXEC=/usr/tools/workflow2.post2.sh
JOBCFG[workflow2.post3] TASKS=2 WCLIMIT=00:01:00 SELECT=TRUE
EXEC=/usr/tools/workflow2.post3.sh
JOBCFG[workflow2.match] ACCOUNT=workflow2
JOBMATCHCFG[workflow2.job] JMIN=workflow2.match JSET=workflow2
```

### Example 22-7: A Complex Workflow

A complex workflow that handles failures:

```
# Workflow 4
JOBCFG[workflow4.step1] TASKS=1 WCLIMIT=00:01:00 SELECT=TRUE
EXEC=/usr/tools/workflow.step1.sh TEMPLATEDPEND=BEFOREFAIL:workflow4.fail1
JOBCFG[workflow4.fail1] TASKS=1 WCLIMIT=00:00:30 SELECT=TRUE
EXEC=/usr/tools/workflow.fail.1.sh TEMPLATEDPEND=BEFOREANY:workflow4.fail2
JOBCFG[workflow4.fail2] TASKS=1 WCLIMIT=00:00:30 SELECT=TRUE
EXEC=/usr/tools/workflow.fail.2.sh
```

```
# Submission job
JOBCFG[workflow4.step2] TEMPLATEDEPEND=AFTEROK:workflow4.step1
TEMPLATEDEPEND=BEFOREOK:workflow4.step3.1 TEMPLATEDEPEND=BEFOREOK:workflow4.step3.2
JOBCFG[workflow4.step3.1] TASKS=1 WCLIMIT=00:01:00 SELECT=TRUE
EXEC=/usr/tools/workflow.step3.1.sh
JOBCFG[workflow4.step3.2] TASKS=1 WCLIMIT=00:01:00 SELECT=TRUE
EXEC=/usr/tools/workflow.step3.2.sh TEMPLATEDEPEND=BEFOREOK:workflow4.step4
JOBCFG[workflow4.step4] TASKS=1 WCLIMIT=00:01:00 SELECT=TRUE
EXEC=/usr/tools/workflow.step4.sh
JOBCFG[workflow4.step4] TEMPLATEDEPEND=BEFOREOK:workflow4.step5.1
TEMPLATEDEPEND=BEFOREOK:workflow4.step5.2 TEMPLATEDEPEND=BEFORENOTOK:workflow4.step5.3
JOBCFG[workflow4.step5.1] TASKS=1 WCLIMIT=00:01:00 SELECT=TRUE
EXEC=/usr/tools/workflow.step5.1.sh
JOBCFG[workflow4.step5.2] TASKS=1 WCLIMIT=00:01:00 SELECT=TRUE
EXEC=/usr/tools/workflow.step5.2.sh
JOBCFG[workflow4.step5.3] TASKS=1 WCLIMIT=00:00:30 SELECT=TRUE
EXEC=/usr/tools/workflow.step5.3.sh
JOBCFG[workflow4.match] ACCOUNT=workflow4
```

---

## Related Topics

- [22.1.5 Creating Workflows with Job Templates - page 930](#)
- [22.1.3 Applying Templates Based on Job Attributes - page 928](#)
- [22.2.3 Job Template Examples - page 945](#)
- [22.2.1 Job Template Extension Attributes - page 931](#)

## Chapter 23: Moab Workload Manager for Grids

Moab Grid Scheduler allows sites to establish relationships among multiple clusters. There are three types of relationships you can implement within the grid: (1) centralized management, (2) hierarchical management, and (3) localized management. These relationships provide access to additional resources, improve load-balancing, provide single system images, and offer other benefits. The grid interface is flexible allowing sites to establish the needed relationship.

### In this chapter:

23.1	Grid Basics .....	950
23.1.1	Grid Overview .....	950
23.1.2	Grid Benefits .....	950
23.1.3	Management-Scalability .....	951
23.1.4	Resource Access .....	951
23.1.5	Load-Balancing .....	952
23.1.6	Single System Image (SSI) .....	952
23.1.7	High Availability .....	953
23.1.8	Grid Relationships .....	953
23.1.9	Submitting Jobs to the Grid .....	958
23.1.10	Viewing Jobs and Resources .....	959
23.2	Grid Configuration Basics .....	960
23.2.1	Peer Configuration Overview .....	960
23.2.2	Initial Configuration .....	960
23.3	Centralized Grid Management (Moab Grid Control / Moab Grid Member) .....	961
23.3.1	Moab Grid Control Configuration .....	961
23.3.2	Moab Grid Member Configuration .....	962
23.4	Hierarchical Grid Management .....	962
23.4.1	Configuring a Peer Server (Source) .....	962
23.4.2	Simple Hierarchical Grid .....	963
23.5	Localized Grid Management .....	964
23.5.1	Enabling Bi-Directional Job Flow .....	964
23.5.2	True Peer-to-Peer Grid .....	965
23.6	Resource Control and Access .....	965
23.6.1	Controlling Resource Information .....	966



23.6.2 Managing Resources with Grid Sandboxes .....	968
23.7 Workload Submission and Control .....	970
23.8 Reservations in the Grid .....	970
23.9 Grid Usage Policies .....	971
23.9.1 Grid Usage Policy Overview .....	971
23.9.2 Peer Job Resource Limits .....	972
23.9.3 Usage Limits via Peer Credentials .....	972
23.9.4 Using General Policies in a Grid Environment .....	973
23.10 Grid Scheduling Policies .....	974
23.10.1 Peer-to-Peer Resource Affinity Overview .....	974
23.10.2 Peer Allocation Policies .....	975
23.10.3 Per-partition Scheduling .....	975
23.11 Grid Credential Management .....	976
23.11.1 Peer Credential Management Overview .....	976
23.11.2 Peer Credential Mapping .....	976
23.11.3 Source and Destination Side Credential Mapping .....	978
23.11.4 Preventing User Space Collisions .....	978
23.12 Grid Data Management .....	978
23.12.1 Grid Data Management Overview .....	979
23.12.2 Configuring Peer Data Staging .....	979
23.12.3 Peer-to-Peer SCP Key Authentication .....	981
23.12.4 Diagnostics .....	982
23.13 Accounting and Allocation Management .....	984
23.13.1 Peer-to-Peer Accounting Overview .....	984
23.13.2 Peer-to-Peer Allocation Management .....	985
23.14 Grid Security .....	986
23.15 Grid Diagnostics and Validation .....	986
23.15.1 Peer Management Overview .....	986
23.15.2 Peer Diagnostic Overview .....	986

## 23.1 Grid Basics

In this topic:

- [23.1.1 Grid Overview - page 950](#)
- [23.1.2 Grid Benefits - page 950](#)
- [23.1.3 Management-Scalability - page 951](#)
- [23.1.4 Resource Access - page 951](#)
- [23.1.5 Load-Balancing - page 952](#)
- [23.1.6 Single System Image \(SSI\) - page 952](#)
- [23.1.7 High Availability - page 953](#)
- [23.1.8 Grid Relationships - page 953](#)
  - [23.1.8.A Centralized Management \(Moab Grid Control / Moab Grid Member\) - page 953](#)
  - [23.1.8.B Centralized/Localized Management \(Hierarchical\) - page 955](#)
  - [23.1.8.C Localized Management \(Peer-to-Peer\) - page 957](#)
- [23.1.9 Submitting Jobs to the Grid - page 958](#)
- [23.1.10 Viewing Jobs and Resources - page 959](#)

### 23.1.1 Grid Overview

A grid enables you to exchange workload and resource status information and to distribute jobs and data among clusters in an established relationship. In addition, you can use resource reservations to mask reported resources, coordinate requests for consumable resources, and quality of service guarantees.

In a grid, some servers running Moab are a source for jobs (that is, where users, portals, and other systems submit jobs), while other servers running Moab are a destination for these jobs (that is, where the jobs execute). Thus, jobs originate from a source server and move to a destination server. For a source server to make an intelligent decision, though, resource availability information must flow from a destination server to that source server.

Because you can manage workload on both the source and destination side of a grid relationship, you have a high degree of control over exactly when, how, and where to execute workload.

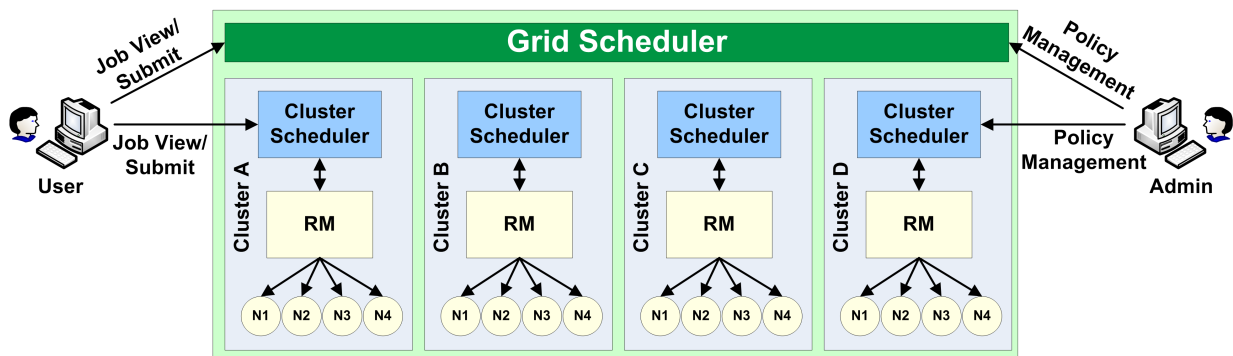
### 23.1.2 Grid Benefits

Moab's peer-to-peer capabilities can be used for multiple purposes, including any of the following:

- manage access to external shared resources
- enable cluster monitoring information services
- enable massive-scalability clusters
- enable distributed grid computing

Of these, the most common use is the creation of grids to join multiple centrally managed, partially autonomous, or fully autonomous clusters. The purpose of this section is to highlight the most common uses of grid technology and provide references to sections that further detail their configuration and management. Other sections cover the standard aspects of grid creation including configuring [peer relationships](#), enabling data staging, [credential management](#), [usage policies](#), and other factors.

*Image 23-1: Jobs Submitted to Grid Scheduler, Then Cluster Schedulers*



### 23.1.3 Management-Scalability

Much like a massive-scalability cluster, a massive-scalability grid allows organizations to overcome scalability limitations in resource managers, networks, message passing libraries, security middleware, file systems, and other forms of software and hardware infrastructure. Moab does this by allowing a single large set of resources to be broken into multiple smaller, more manageable clusters, and then virtually re-assembling them using Moab. Moab becomes responsible for integrating the seams between the clusters and presenting a single-system image back to the end-users, administrators, and managers.

**i** Jobs cannot span clusters.

### 23.1.4 Resource Access

In some cases, the primary motivation for creating a grid is to aggregate resources of different types into a single system. This aggregation allows for multi-step jobs to run a

portion of the job on one architecture, and a portion on another.

A common example of a multi-architecture parameter-sweep job would be a batch regression test suite that requires a portion of the tests running on Redhat 7.2, a portion on SUSE 9.1, a portion on Myrinet nodes, and a portion on Infiniband nodes. While it would be very difficult to create and manage a single cluster that simultaneously provided all of these configurations, Moab can be used to create and manage a single grid that spans multiple clusters as needed.

### 23.1.5 Load-Balancing

While grids often have additional motivations, it is rare to have a grid created where increased total system utilization is not an objective. By aggregating the total pool of jobs requesting resources and increasing the pool of resources available to each job, Moab is able to improve overall system utilization, sometimes significantly. The biggest difficulty in managing multiple clusters is preventing inter-cluster policies and the cost of migration from overwhelming the benefits of decreased fragmentation losses. Even though remote resources may be available for immediate usage, migration costs can occur in the form of credential, job, or data staging and impose a noticeable loss in responsiveness on grid workload.

Moab provides tools to allow these costs to be monitored and managed and both cluster and grid level performance to be reported.

### 23.1.6 Single System Image (SSI)

Another common benefit of grids is the simplicity associated with a single system image-based resource pool. This simplicity generally increases productivity for end-users, administrators, and managers.

An SSI environment tends to increase the efficiency of end-users by minimizing human errors associated with porting a request from a known system to a less known system. Additionally, the single point of access grid reduces human overhead associated with monitoring and managing workload within multiple independent systems.

For system administrators, a single system image can reduce overhead, training time, and diagnostic time associated with managing a cluster. Furthermore, with Moab's peer-to-peer technology, no additional software layer is required to enable the grid and no new tools must be learned. No additional layers means no additional failure points, and that is good for everyone involved.

Managers benefit from SSI by being able to pursue organization mission objectives globally in a more coordinated and unified manner. They are also able to monitor progress toward those objectives and effectiveness of resources in general.

## 23.1.7 High Availability

A final benefit of grids is their ability to decrease the impact of failures. Grids add another layer of high availability to the cluster-level high availability. For some organizations, this benefit is a primary motivation, pulling together additional resources to allow workload to continue to be processed even in the event that some nodes, or even an entire cluster, become unavailable. Whether the resource unavailability is based on node failures, network failures, systems middleware, systems maintenance, or other factors, a properly configured grid can reroute priority workload throughout the grid to execute on other compatible resources.

With grids, there are a number of important factors in high availability that should be considered:

- enabling highly available job submission/job management interfaces
- avoiding network failures with redundant routes to compute resources
- handling partial failures
- dynamically restarting failed jobs

## 23.1.8 Grid Relationships

There are three types of relationships you can implement within the grid:

- [Centralized Management \(Moab Grid Control / Moab Grid Member\)](#)
- [Centralized/Localized Management \(Hierarchical\)](#)
- [Localized Management \(Peer-to-Peer\)](#)

### 23.1.8.A Centralized Management (Moab Grid Control / Moab Grid Member)

**Note:** Moab Grid Control / Moab Grid Member are also known as MGC/MGM.

The centralized management model (MGC/MGM) allows users to submit jobs to a centralized source server running Moab. The source Moab server obtains full resource information from all clusters and makes intelligent scheduling decisions across all clusters. Jobs (and data when configured to do so) are distributed to the remote clusters as needed. The centralized management model is recommended for intra-organization grid environments when cluster autonomy is not as necessary.


In the centralized management (Moab Grid Control / Moab Grid Member) configuration, roles are clear. In other configurations, individual Moab servers may simultaneously act as sources to some clusters and destinations to others or as both a source and a destination to another cluster.

**Example of the Centralized Management (Moab Grid Control / Moab Grid Member) Model**

XYZ Research has three clusters - MOAB1, MOAB2, and MOAB3--running Moab and the Torque resource manager. They would like to submit jobs at a single location (cluster MOAB1) and have the jobs run on whichever cluster can provide the best responsiveness.

The desired behavior is essentially a *master-slave* relationship. MOAB1 is the central, or MGC cluster. On MOAB1, resource managers point to the local Torque resource manager and to the Moab servers on cluster MOAB2 and cluster MOAB3. The Moab servers on MOAB2 and MOAB3 are configured to trust cluster MOAB1 and to execute in SLAVE mode.

*Image 23-2: Centralized Management*



With this configuration, XYZ Research can submit jobs to the MGC Moab server running on cluster MOAB1 and can, as stated earlier, submit jobs from the MGM nodes as well. However, only the MGC Moab server can schedule jobs. For example, cluster MOAB2 and

cluster MOAB3 cannot schedule a job, but they can accept a job and retain it in an idle state until the MGC directs it to run.


**i** You can turn off job submission on MGM nodes by setting the `DISABLESLAVEJOBSUBMIT` parameter to `TRUE`.

the MGC Moab server obtains full resource information from all three clusters and makes intelligent scheduling decisions and distributes jobs (and data when configured to do so) to the remote clusters. The Moab servers running on clusters MOAB2 and MOAB3 are destinations behaving like a local resource manager. The Moab server running on MOAB1 is a source, loading and using this resource information.

### 23.1.8.B Centralized/Localized Management (Hierarchical)

As with the centralized management model (Moab Grid Control / Moab Grid Member), the hierarchical model allows users to submit jobs to a centralized source server running Moab. However, in the hierarchical model, clusters retain sovereignty, allowing local job scheduling. Thus, if communication between the source and destination clusters is interrupted, the destination cluster(s) can still run jobs locally.

*Image 23-3: Hierarchical Management*




In the hierarchical model, the source Moab server obtains full resource information from all clusters and makes intelligent scheduling decisions across all clusters. As needed, jobs and data are distributed to the remote clusters. Or, if preferred, a destination cluster can also serve as its own source; however, a destination cluster cannot serve as a source to another destination cluster. The centralized management model is recommended for intra-organization grid environments when cluster autonomy and/or local management is necessary.



### 23.1.8.C Localized Management (Peer-to-Peer)

The localized management (peer-to-peer) model allows you to submit jobs on one cluster and schedule the jobs on the other cluster (it currently works with two clusters). For example, a job may be submitted on MOAB1 and run on MOAB2. Jobs can also migrate in the opposite direction (that is, from MOAB2 to MOAB1). The source servers running Moab obtain full resource information from both clusters and make intelligent scheduling decisions across both clusters. Jobs (and data when configured to do so) are migrated to other clusters as needed.

*Image 23-4: Localized Management*



**i** Jobs will not migrate indefinitely. The localized management model limits them to one migration.

This model allows clusters to retain their autonomy while still allowing jobs to run on either cluster. No central location for job submission is needed, and you do not need to submit jobs from different nodes based on resource needs. You can submit a job from any location

and it is either migrated to nodes on the least utilized cluster or the cluster requested in the job submission. This model is recommended for grids in an inter-organization grid environment.

### 23.1.9 Submitting Jobs to the Grid

In any peer-to-peer or grid environment where jobs must be migrated between clusters, use the Moab `msub` command. Once a job has been submitted to Moab using `msub`, Moab identifies potential destinations and migrates the job to the destination cluster.

Using Moab's `msub` job submission command, jobs can be submitted using PBS command file syntax and be run on any cluster using any of the resource managers. For example, a PBS job script can be submitted using `msub` and depending on availability, Moab can translate a subset of the job's directives and execute it on a PBS cluster.

**i** Moab can only stage/migrate jobs between resource managers (in between clusters) that have been submitted using the `msub` command. If jobs are submitted directly to a low-level resource manager, such as PBS, Moab will still be able to schedule them, but only on resources directly managed by the resource manager to which they were submitted.

#### Example

A research lab wants to use spare cycles on its four clusters, each of which is running a local resource manager. In addition to providing better site-wide load balancing, the goal is to also provide some of its users with single point access to all compute resources. Various researchers have made it clear that this new multi-cluster load balancing must not impose any changes on users who are currently using these clusters by submitting jobs locally to each cluster.

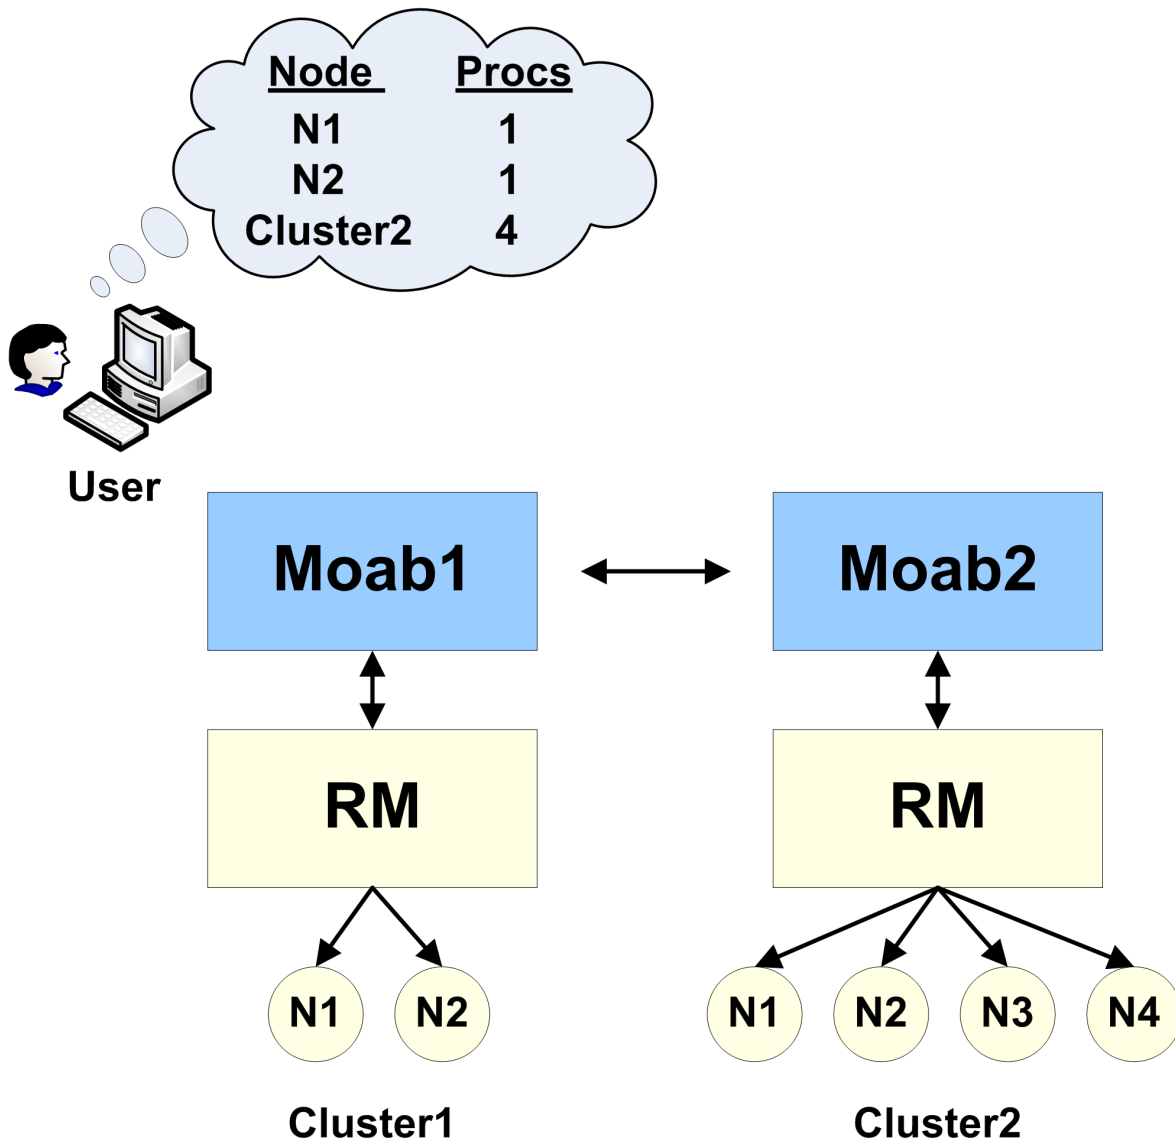
In this example, the scheduler mode of the destination clusters should be set to `NORMAL` rather than `SLAVE`. In `SLAVE` mode, Moab makes no local decisions - it simply follows the directions of remote trusted peers. In `NORMAL` mode, each Moab is fully autonomous, scheduling all local workload and coordinating with remote peers when and how to schedule migrated jobs.

From the perspective of a local cluster user, no new behaviors are seen. Remote jobs are migrated in from time to time, but to the user each job looks as if it were locally submitted. The user continues to submit, view, and manage jobs as before, using existing local jobs scripts.

### 23.1.10 Viewing Jobs and Resources

By default, each destination Moab server will report all compute nodes it finds back to the source Moab server. These reported nodes appear within the source Moab as local nodes each within a partition associated with the resource manager reporting them. If a source resource manager was named `slave1`, all nodes reported by it would be associated with the `slave1` partition. Users and administrators communicating with the source Moab via Moab Cluster Manager, or standard Moab command line tools would be able to view and analyze all reported nodes.

*Image 23-5: Viewing Jobs and Resources*



**i** The grid view will be displayed if either the source or the destination server is configured with grid view.

For job information, the default behavior is to only report to the source Moab information regarding jobs that originated at the source. If information about other jobs is desired, this can be configured as shown in the [Workload Submission and Control](#) section.

---

### Related Topics

- [Resource Control and Access](#)

## 23.2 Grid Configuration Basics

In this topic:

- [23.2.1 Peer Configuration Overview - page 960](#)
- [23.2.2 Initial Configuration - page 960](#)

### 23.2.1 Peer Configuration Overview

In the simplest case, establishing a peer relationship can be accomplished with as few as two configuration lines: one line to indicate how to contact the peer and one line to indicate how to authenticate the server. However, data migration issues, credential mapping, and usage policies must often be addressed in order to make a peer-based grid effective.

To address these issues Moab provides facilities to control how peers inter-operate, enabling full autonomy over both client and server ends of the peer relationship.

### 23.2.2 Initial Configuration

At a minimum, only two parameters must be specified to establish a peer relationship: [RMCFG](#) and [CLIENTCFG\[<X>\]](#). [RMCFG](#) allows a site to specify interface information directing Moab on how to contact and inter-operate with the peer. For peer interfaces, a few guidelines must be followed with the [RMCFG](#) parameter:

- The `TYPE` attribute of the peer must be set to `moab`.
- The `SERVER` attribute must point to the host and user interface port of the remote Moab server.
- The `name` of the resource manager should match the name of the remote peer cluster as specified with the [SCHEDCFG](#) parameter in the peer `moab.cfg`.

```
# moab.cfg on MoabServer01
SCHEDCFG[MoabServer01] MODE=NORMAL SERVER=hpc-01:41111
RMCFG[MoabServer02] TYPE=moab SERVER=hpc-02:40559
...
```

Configuring the `CLIENTCFG` parameter is mandatory. When specifying the `CLIENTCFG` parameter for peers, the following guidelines must be followed:

- The `CLIENTCFG` parameter must be specified in the `moab-private.cfg` file on both peers.
- An `RM:` prefix is required before the peer's name.
- If using default secret key based security, the value of the `KEY` attribute must match the `KEY` value set on the corresponding remote peer.
- The `AUTH` attribute must be set to `admin1` in the `moab-private.cfg` on the destination Moab.

```
# moab-private.cfg on MoabServer01
CLIENTCFG[RM:MoabServer02] KEY=3esfv0=32re2-tdbne
....
```

```
# moab-private.cfg on MoabServer02
CLIENTCFG[RM:MoabServer01] KEY=3esfv0=32re2-tdbne AUTH=admin1
...
```

## 23.3 Centralized Grid Management (Moab Grid Control / Moab Grid Member)

In this topic:

[23.3.1 Moab Grid Control Configuration - page 961](#)

[23.3.2 Moab Grid Member Configuration - page 962](#)

### 23.3.1 Moab Grid Control Configuration

**Note:** Moab Grid Control / Moab Grid Member are also known as MGC/MGM.

The process of setting up the Moab Grid Control Configuration is the same as setting up a [source Moab configuration](#). The Moab Grid Control / Moab Grid Member relationship is configured in each `moab.cfg` on the MGM.

```
# moab.cfg on Master
SCHEDCFG[master] SERVER=master:42559 MODE=NORMAL
...
```

```
# moab-private.cfg on Master
CLIENTCFG[RM:slave1] KEY=3esfv0=32re2-tdbne
...
```

## 23.3.2 Moab Grid Member Configuration

the MGM's relationship with the MGC is determined by the `MODE`. Setting `MODE` to `SLAVE` notifies the MGC to take control of starting jobs on the MGM. the MGC starts the jobs on the MGM. In `SLAVE` mode, jobs can be submitted locally to the MGM, but are not seen or started by the MGC. When a job is submitted locally to the MGM, the job is locked into the cluster and cannot migrate to other clusters.

```
# moab.cfg on Slave
SCHEDCFG[slave1] SERVER=slave1:42559 MODE=SLAVE
...
```

```
# moab-private.cfg on Slave
CLIENTCFG[RM:master] KEY=3esfv0=32re2-tdbne AUTH=admin1
...
```

## 23.4 Hierarchical Grid Management

In this topic:

[23.4.1 Configuring a Peer Server \(Source\) - page 962](#)

[23.4.2 Simple Hierarchical Grid - page 963](#)

### 23.4.1 Configuring a Peer Server (Source)

Peer relationships are enabled by creating and configuring a [resource manager](#) interface using the `RMCFG` parameter. This interface defines how a given Moab will load resource and workload information and enforce its scheduling decisions. In non-peer cases, the `RMCFG` parameter points to a resource manager such as Torque. However, if the `TYPE` attribute is set to `moab`, the `RMCFG` parameter can be used to configure and manage a peer relationship.

## 23.4.2 Simple Hierarchical Grid

The first step to create a new peer relationship is to configure an interface to a destination Moab server. In the following example, cluster C1 is configured to be able to *see* and *use* resources from two other clusters.

```
SCHEDCFG [C1]  MODE=NORMAL  SERVER=head.C1.xyz.com:41111
RMCFG [C2]    TYPE=moab    SERVER=head.C2.xyz.com:40559
RMCFG [C3]    TYPE=moab    SERVER=head.C3.xyz.com:40559
...
```

*C1 allows a global view of the underlying clusters. From C1, jobs can be viewed and modified. C2 and C3 act as separate scheduling entities that can receive jobs from C1. C1 migrates jobs to C2 and C3 based on available resources and policies of C1. Jobs migrated to C2 and C3 are scheduled according to the policies on C2 and C3.*

In this case, one RMCFG parameter is all that is required to configure each peer relationship if standard secret key based authentication is being used and a shared default secret key exists between the source and destination Moabs. However, if peer relationships with multiple clusters are to be established and a per-peer secret key is to be used (highly recommended), then a CLIENTCFG parameter must be specified for the authentication mechanism. Because the secret key must be kept secure, it must be specified in the `moab-private.cfg` file. For the current example, a per-peer secret key could be set up by creating the following `moab-private.cfg` file on the C1 cluster.

```
CLIENTCFG [RM:C2]  KEY=fastclu3t3r
CLIENTCFG [RM:C3]  KEY=14436aaa
```

**i** The key specified can be any alphanumeric value and can be locally generated or made up. The only critical aspect is that the keys specified on each end of the peer relationship match.

Additional information can be found in the [Grid Security](#) section, which provides detailed information on designing, configuring, and troubleshooting peer security.

Continuing with the example, the initial source side configuration is now complete. On the destination clusters, C2 and C3, the first step is to configure authentication. If a shared default secret key exists between all three clusters, then configuration is complete and the clusters are ready to communicate. If per-peer secret keys are used (recommended), then it will be necessary to create matching `moab-private.cfg` files on each of the destination clusters. With this example, the following files would be required on C2 and C3 respectively:

```
CLIENTCFG [RM:C1]  KEY=fastclu3t3r  AUTH=admin1
```

```
CLIENTCFG [RM:C1]  KEY=14436aaa  AUTH=admin1
```

Once peer security is established, a final optional step would be to configure scheduling behavior on the destination clusters. By default, each destination cluster accepts jobs from

each trusted peer. However, it will also be fully autonomous, accepting and scheduling locally submitted jobs and enforcing its own local policies and optimizations. If this is the desired behavior, then configuration is complete.

In the current example, with no destination side scheduling configuration, jobs submitted to cluster C1 can run locally, on cluster C2 or on cluster C3. However, the established configuration does not necessarily enforce a strict master-slave relationship because each destination cluster (C2 and C3) has complete autonomy over how, when, and where it schedules both local and remote jobs. Each cluster can potentially receive jobs that are locally submitted and can also receive jobs from other source Moab servers.

Further, each destination cluster will accept any and all jobs migrated to it from a trusted peer without limitations on who can run, when and where they can run, or how many resources they can use. If this behavior is either too restrictive or not restrictive enough, then destination side configuration will be required.

## 23.5 Localized Grid Management


In this topic:

[23.5.1 Enabling Bi-Directional Job Flow - page 964](#)

[23.5.2 True Peer-to-Peer Grid - page 965](#)

### 23.5.1 Enabling Bi-Directional Job Flow

*Image 23-6: Bi-directional peer-to-peer setup*





For each peer interface, an `RMCFG` parameter is only required for the server (or source side of the interface). If two peers are to share jobs in both directions, the relationship is considered to be bi-directional.

## 23.5.2 True Peer-to-Peer Grid

Previous examples involved grid masters that coordinated the activities of the grid and made it so direct contact between peers was not required. However, if preferred, the master is not required and individual clusters can interface directly with each other in a true peer manner. This configuration is highlighted in the following example:

```
# Cluster A
SCHEDCFG[clusterA] MODE=NORMAL SERVER=clusterA
RMCFG[clusterA]     TYPE=pbs
RMCFG[clusterB]     TYPE=moab SERVER=clusterB:40559
CLIENTCFG[RM:clusterB] AUTH=admin1 KEY=banana16
```

```
# Cluster B
SCHEDCFG[clusterB] MODE=NORMAL SERVER=clusterB
RMCFG[clusterB]     TYPE=pbs
RMCFG[clusterA]     TYPE=moab SERVER=clusterA:40559
CLIENTCFG[RM:clusterA] AUTH=admin1 KEY=banana16
```

**i** If you are using Moab Accounting Manager, the `Start` action is not supported as a non-blocking accounting action in Peer-to-Peer grids. You will need to include `Start` as a blocking action. For example:

```
AMCFG[mam] BLOCKINGACTIONS=Start
```

## 23.6 Resource Control and Access

In this topic:

- 23.6.1 Controlling Resource Information - page 966
  - 23.6.1.A Direct Node View - page 966
  - 23.6.1.B Mapped Node View - page 966
- 23.6.2 Managing Resources with Grid Sandboxes - page 968
  - 23.6.2.A Controlling Access on a Per Cluster Basis - page 969
  - 23.6.2.B Access Control Lists/Granting Access to Local Jobs - page 969

## 23.6.1 Controlling Resource Information

In a Moab peer-to-peer grid, resources can be viewed in one of two models:

- **Direct** - nodes are reported to remote clusters exactly as they appear in the local cluster
- **Mapped** - nodes are reported as individual nodes, but node names are mapped to a unique name when imported into the remote cluster

### 23.6.1.A Direct Node View


Direct node import is the default resource information mode. No additional configuration is required to enable this mode.

### 23.6.1.B Mapped Node View

In this mode, nodes are reported just as they appear locally by the exporting cluster. However, on the importing cluster side, Moab maps the specified node names using the resource manager **object map**. In an object map, node mapping is specified using the `node` keyword as in the following example:

```
SCHEDCFG[gridmaster] MODE=NORMAL
RMCFG[clusterB]      TYPE=moab OMAP=file://$HOME/clusterb.omap.dat
...
node:b_*,*
```

*Image 23-7: Mapped Node View*

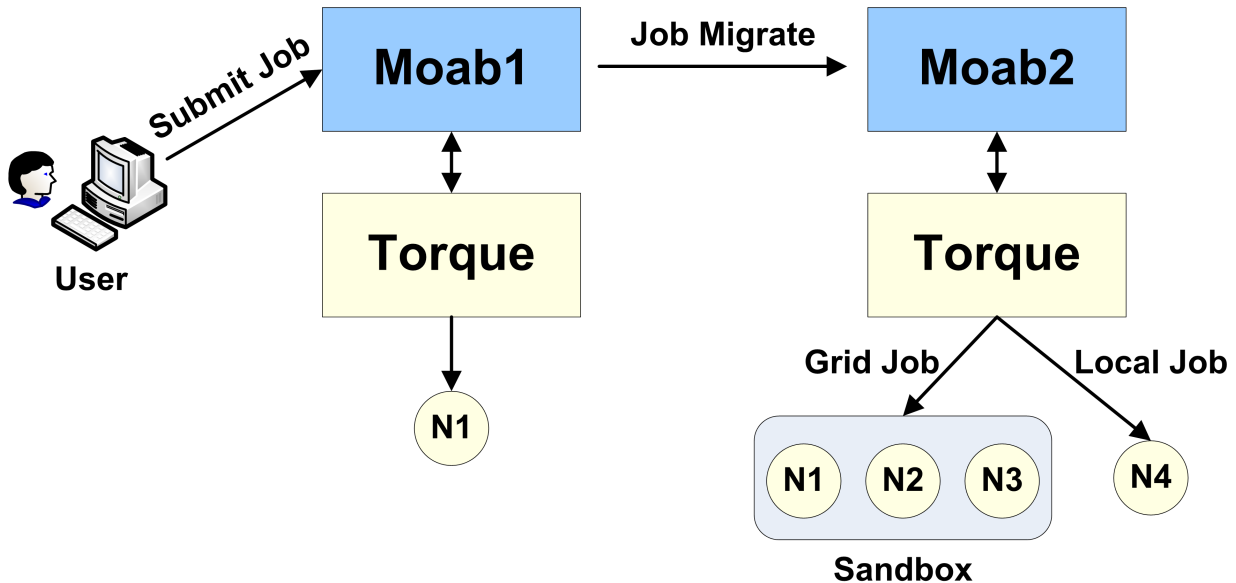


In this example, all nodes reported by `clusterB` have the string `b_` prepended to prevent node name space conflicts with nodes from other clusters. For example, if cluster `clusterB` reported the nodes `node01`, `node02`, and `node03`, cluster `gridmaster` would report them as `b_node01`, `b_node02`, and `b_node03`.

See [object mapping](#) for more information on creating an object map file.

## 23.6.2 Managing Resources with Grid Sandboxes

Image 23-8: Grid Sandbox



A cluster may want to participate in a grid but may desire to dedicate only a set amount of resources to external grid workload or may only want certain peers to have access to particular sets of resources. With Moab, this can be achieved by way of a grid sandbox, which must be configured at the destination cluster. Grid sandboxes can both constrain external resource access and limit which resources are reported to other peers. This allows a cluster to only report a defined subset of its total resources to source peers and restricts peer workload to the sandbox. The sandbox can be set aside for peer use exclusively, or can allow local workload to also run inside of it. Through the use of multiple, possibly overlapping grid sandboxes, a site can fully control resource availability on a per peer basis.

A grid sandbox is created by configuring a [standing reservation](#) on a destination peer and then specifying the `ALLOWGRID` flag on that reservation. This flag tells the Moab destination peer to treat the standing reservation as a grid sandbox, and, by default, only the resources in the sandbox are visible to grid peers. Also, the sandbox only allows workload from other peers to run on the contained resources.

### Example 23-1: Dedicated Grid Sandbox

```
SRCFG[sandbox1] PERIOD=INFINITY HOSTLIST=node01,node02,node03
SRCFG[sandbox1] CLUSTERLIST=ALL FLAGS=ALLOWGRID
...
```

*The standing reservation `sandbox1` creates a grid sandbox, which always exists and contains the nodes `node01`, `node02`, and `node03`. This sandbox will only allow grid workload to run within it by default. This means that the scheduler will not consider the boxed resources for local workload.*

Grid sandboxes inherit all of the same power and flexibility that standing reservations have. See [Managing Reservations](#) for additional information.

**i** The flag `ALLOWGRID` marks the reservation as a grid sandbox and as such, it precludes grid jobs from running anywhere else. However, it does *not* enable access to the reserved resources. The `CLUSTERLIST` attribute in the above example enables access to all remote jobs.

### 23.6.2.A Controlling Access on a Per Cluster Basis

Often clusters may want to control which peers are allowed to use certain sandboxes. For example, Cluster A may have a special contract with Cluster B and will let overflow workload from Cluster B run on 60% of its resources. A third peer in the grid, Cluster C, doesn't have the same contractual agreement, and is only allowed 10% of Cluster A at any given time. Therefore two separate sandboxes must be made to accommodate the different policies.

```
SRCFG[sandbox1] PERIOD=INFINITY HOSTLIST=node01,node02,node03,node04,node05
SRCFG[sandbox1] FLAGS=ALLOWGRID CLUSTERLIST=ClusterB
SRCFG[sandbox2] PERIOD=INFINITY HOSTLIST=node06 FLAGS=ALLOWGRID
SRCFG[sandbox2] CLUSTERLIST=ClusterB,ClusterC,ClusterD USERLIST=ALL
...
```

*This example configuration illustrates how cluster A could set up their sandboxes to follow a more complicated policy. In this policy, `sandbox1` provides exclusive access to nodes 1 through 5 to jobs coming from peer ClusterB by including `CLUSTERLIST=ClusterB` in the definition. Reservation `sandbox2` provides shared access to node 6 to local jobs and to jobs from clusters B, C, and D through use of the `CLUSTERLIST` and `USERLIST` attributes.*

With this setup, the following policies are enforced:

- local jobs can see all nodes and run anywhere except nodes 1 through 5
- jobs from cluster B can see and run only on nodes 1 through 6
- jobs from clusters C and D can see and run only on node 6

As shown in the example above, sandboxes can be shared across multiple peers by listing all sharing peers in the `CLUSTERLIST` attribute (comma-delimited).

### 23.6.2.B Access Control Lists/Granting Access to Local Jobs

It is not always desirable to have the grid sandbox reserve resources for grid consumption, exclusively. Many clusters may want to use the grid sandbox when local workload is high and demand from the grid is relatively low. Clusters may also want to further restrict what kind of grid workload can run in a sandbox. This fine-grained control can be achieved by attaching access control lists (ACLs) to grid sandboxes.

Since sandboxes are basically special standing reservations, the syntax and rules for specifying an ACL is identical to those found in [Managing Reservations](#).

### Example

```
SRCFG[sandbox2] PERIOD=INFINITY HOSTLIST=node04,node05,node06
SRCFG[sandbox2] FLAGS=ALLOWGRID QOSLIST=high GROUPLIST=engineer
...
```

*A cluster decides to dedicate resources to a sandbox, but wishes local workload to also run within it. An additional ACL is then associated with the definition. The reservation `sandbox2` takes advantage of this feature by allowing local jobs running with a QOS of `high`, or under the group `engineer`, to also run on the sandboxed nodes `node04`, `node05`, and `node06`.*

## 23.7 Workload Submission and Control

### Controlling Peer Workload Information

By default, a peer is only responsible for workload that is submitted via that particular peer. This means that when a source peer communicates with destination peers it only receives information about workload it sent to those destination peers. If desired, the destination peers can send information about *all* of its workload: both jobs originating locally and remotely. This is called *local workload exporting*. This might help simplify administration of different clusters by centralizing monitoring and management of jobs at one peer.

To implement local workload exporting, use the `LOCALWORKLOADEXPORT` resource manager flag. For example:

```
RMCFG[ClusterA.INBOUND] FLAGS=LOCALWORKLOADEXPORT # source peer
...
```

*This example shows the configuration on a destination peer (ClusterB) that exports its local and remote workload to the source peer (ClusterA).*

**i** `LOCALWORKLOADEXPORT` does not need to be configured in Moab Grid Control / Moab Grid Member grids.

### Related Topics

- [Job Start Time Estimates](#)

## 23.8 Reservations in the Grid

In some environments, globally-shared resources might need to be managed to guarantee the full environment required by a particular job. Resources such as networks, storage

systems, and license managers can be used only by batch workload but this workload can be distributed among multiple independent clusters. Consequently, the jobs from one cluster can utilize resources required by jobs from another. Without a method of coordinating the needs of the various cluster schedulers, resource reservations will not be respected by other clusters and will be of only limited value.

Using the centralized model, Moab allows the importing and exporting of reservations from one peer server to another. With this capability, a source peer can be set up for the shared resource to act as a clearinghouse for other Moab cluster schedulers. This source peer Moab server reports configured and available resource state and in essence possesses a global view of resource reservations for all clusters for the associated resource.

To allow the destination peer to export reservation information to the source Moab, the `RMCFG` lines for all client resource managers must include the flag `RSVEXPORT`. The source Moab should be configured with a resource manager interface to the destination peer and include both the `RSVEXPORT` and `RSVIMPORT` flags. For the destination peer, `RSVEXPORT` indicates that it should *push* information about newly created reservations to the source Moab, while the `RSVIMPORT` flag indicates that the source Moab server should import and locally enforce reservations detected on the destination peer server.

## 23.9 Grid Usage Policies

In this topic:

- [23.9.1 Grid Usage Policy Overview - page 971](#)
- [23.9.2 Peer Job Resource Limits - page 972](#)
- [23.9.3 Usage Limits via Peer Credentials - page 972](#)
- [23.9.4 Using General Policies in a Grid Environment - page 973](#)
- [23.9.4.A Source Cluster Policies - page 973](#)

### 23.9.1 Grid Usage Policy Overview

Moab allows extensive control over how peers interact. These controls allow the following:

- Limiting which remote users, group, and accounts can utilize local compute resources
- Limiting the total quantity of local resources made available to remote jobs at any given time
- Limiting remote resource access to a specific subset of resources

- Limiting timeframes during which local resources will be made available to remote jobs
- Limiting the types of remote jobs that will be allowed to execute

## 23.9.2 Peer Job Resource Limits

Both source and destination peers can limit the types of jobs they will allow in terms of resources requested, services provided, job duration, applications used, etc using Moab's job template feature. Using this method, one or more job profiles can be created on either the source or destination side, and Moab can be configured to allow or reject jobs based on whether or not the jobs meet the specified job profiles.

When using the `ALLOWJOBLIST` and `REJECTJOBLIST` attributes, the following rules apply:

- All jobs that meet the job templates listed by `ALLOWJOBLIST` are allowed.
- All jobs that do not meet `ALLOWJOBLIST` job templates and which do meet `REJECTJOBLIST` job templates are rejected.
- All jobs that meet no job templates in either list are allowed.

## 23.9.3 Usage Limits via Peer Credentials

With peer interfaces, destination clusters willing to accept remote jobs can [map](#) these jobs onto a select subset of users, accounts, QoSs, and queues. With the ability to lock these jobs into certain credentials comes the ability to apply any arbitrary credential constraints, priority adjustments, and resource limitations normally available within cluster management. Specifically, the following can be accomplished:

- Limit number of active jobs simultaneously allowed
- Limit quantity of allocated compute resources simultaneously allowed
- Adjust job priority
- Control access to specific scheduling features (deadlines, reservations, preemption, etc.)
- Adjust fairshare targets
- Limit resource access



## 23.9.4 Using General Policies in a Grid Environment

While Moab does provide a number of unique grid-based policies for use in a grid environment, the vast majority of available management tools come from the transparent application of cluster policies. Cluster-level policies such as [job prioritization](#), [node allocation](#), [fairshare](#), [usage limits](#), [reservations](#), [preemption](#), and [allocation management](#) all just work and can be applied in a grid in exactly the same manner.

The one key concept to understand is that in a centralized based grid, these policies apply across the entire grid; in a peer-based grid, these policies apply only to local workload and resources.

### 23.9.4.A Source Cluster Policies

In many cases, organizations are interested in treating jobs differently based on their point of origin. This can be accomplished by assigning and/or keying off of a unique credential associated with the remote workload. For example, a site might want to constrain jobs from a remote cluster to only a portion of the total available cluster cycles. This could be accomplished using usage limits, fairshare targets, fairshare caps, reservations, or allocation management based policies.

The examples below show three different approaches for constraining remote resource access.

#### Example 23-2: Constraining Remote Resource Access via Fairshare Caps

```
# define peer relationship and map all incoming jobs to orion account
RMCFG[orion.INBOUND] SET.JOB=orion.set
JOBCFG[orion.set] ACCOUNT=orion
# configure basic fairshare for 7 one day intervals
FSPOLICY DEDICATEDPS
FSINTERVAL 24:00:00
FSDEPTH 7
FSUSERWEIGHT 100
# use fairshare cap to limit jobs from orion to 10% of cycles
ACCOUNTCFG[orion] FSCAP=10%
```

#### Example 23-3: Constraining Remote Resource Access via Fairshare Targets and Preemption

```
# define peer relationship and map all incoming jobs to orion account
RMCFG[orion.INBOUND] SET.JOB=orion.set
JOBCFG[orion.set] ACCOUNT=orion
# local cluster can preempt jobs from orion
USERCFG[DEFAULT] JOBFLAGS=PREEMPTOR
PREEMTPOLICY CANCEL
# configure basic fairshare for 7 one day intervals
FSPOLICY DEDICATEDPS
FSINTERVAL 24:00:00
FSDEPTH 7
FSUSERWEIGHT 100
# decrease priority of remote jobs and force jobs exceeding 10% usage to be
```

```
preemptible
ACCOUNTCFG[orion] FSTARGET=10-
ENABLEFSVIOLATIONPREEMPTION TRUE
```

#### Example 23-4: Constraining Remote Resource Access via Priority and Usage Limits

```
# define peer relationship and map all incoming jobs to orion account RMCFG
[orion.INBOUND] SET.JOB=orion.set
JOBCFG[orion.set] QOS=orion
USERCFG[DEFAULT] QDEF=orion
# local cluster can preempt jobs from orion
USERCFG[DEFAULT] JOBFLAGS=PREEMPTOR
PREEMTPOLICY CANCEL
# adjust remote jobs to have reduced priority
QOSCFG[orion] PRIORITY=-1000
# allow remote jobs to use up to 64 procs without being preemptible and up to 96 as
preemptees
QOSCFG[orion] MAXPROC=64,96
ENABLESPVIOLATIONPREEMPTION TRUE
```

#### Related Topics

- [Grid Sandbox](#) - control grid resource access

## 23.10 Grid Scheduling Policies

In this topic:

- [23.10.1 Peer-to-Peer Resource Affinity Overview - page 974](#)
- [23.10.2 Peer Allocation Policies - page 975](#)
- [23.10.3 Per-partition Scheduling - page 975](#)

### 23.10.1 Peer-to-Peer Resource Affinity Overview

The concept of resource affinity stems from a number of facts:

- Certain compute architectures are able to execute certain compute jobs more effectively than others.
- From a given location, staging jobs to various clusters might require more expensive [allocations](#), more data and network resources, and more use of system services.
- Certain compute resources are owned by external organizations and should be used sparingly.

Regardless of the reason, Moab servers allow the use of peer resource affinity to guide jobs to the clusters that make the best fit according to a number of criteria.

At a high level, this is accomplished by creating a number of job templates and associating the profiles with different peers with varying impacts on estimated execution time and peer affinity.

### 23.10.2 Peer Allocation Policies

A direct way to assign a peer allocation algorithm is with the `PARALLELLOCATIONPOLICY` parameter. Legal values are listed in the following table:

Value	Description
<b>FirstStart</b>	Allocates resources from the eligible peer that can start the job the soonest.
<b>LoadBalance</b>	Allocates resources from the eligible peer with the most available resources; measured in tasks (balances workload distribution across potential peers).
<b>LoadBalanceP</b>	Allocates resources from the eligible peer with the most available resources; measured in percent of configured resources (balances workload distribution across potential peers).
<b>Random</b>	Allocates partitions in a random order each iteration. In general, all the jobs scheduled within the same iteration receive the same randomized list of partitions. This means the randomization happens between iterations and not within the same iteration. One iteration Moab might start with partition X and the next it might start with partition Y.
<b>RoundRobin</b>	Allocates resources from the eligible peer that has been least recently allocated.

**i** The `mdiag -t -v` command can be used to view current calculated partition priority values.

### 23.10.3 Per-partition Scheduling

Per-partition scheduling can be enabled by adding the following lines to `moab.cfg`:

```
PERPARTITIONSCHEDULING TRUE
JOBMIGRATEPOLICY JUSTINTIME
```

To use per-partition scheduling, you must configure fairshare trees where particular users have higher priorities on one partition, and other users have higher priorities on a different partition.

**i** Do not set the `USEANYPARTITIONPRIO` parameter if you use per-partition scheduling. Doing so causes Moab to schedule jobs to the first partition listed, even if nodes from another partition will be available sooner.

## 23.11 Grid Credential Management

In this topic:

[23.11.1 Peer Credential Management Overview - page 976](#)

[23.11.2 Peer Credential Mapping - page 976](#)

[23.11.3 Source and Destination Side Credential Mapping - page 978](#)

[23.11.4 Preventing User Space Collisions - page 978](#)

### 23.11.1 Peer Credential Management Overview

Moab provides a number of credential management features that allow sites to control which local users can utilize remote resources and which remote users can utilize local resources and under what conditions this access is granted.

### 23.11.2 Peer Credential Mapping

If two peers share a common user space (a given user has the same login on both clusters), then there is often no need to enable credential mapping. When users, groups, classes, QoSes, and accounts are not the same from one peer to another, Moab allows a site to specify an Object Map URL. This URL contains simple one to one or expression based mapping for credentials and other objects. Using the `RMCFG` parameter's `OMAP` attribute, a site can tell Moab where to find these mappings. The object map uses the following format:

```
<OBJECTTYPE> : <SOURCE_OBJECTID>, <DESTINATION_OBJECTID>
```

where `<SOURCE_OBJECT>` can be a particular username or an asterisk (\*) that is a wildcard matching all credentials of the specified type, which have not already been matched.

The object map file can be used to translate the following:

Keyword	Objects
<b>account</b>	accounts/projects
<b>class</b>	classes/queues
<b>file</b>	files/directories
<b>group</b>	groups
<b>node</b>	nodes
<b>qos</b>	QoS
<b>user</b>	users

The following `moab.cfg` and `omap.dat` files demonstrate a sample credential mapping.

```
SCHEDCFG[master1] MODE=normal
RMCFG[slave1]    OMAP=file:///opt/moab/omap.dat
...
```

```
user:joe,jsmith
user:steve,sjohnson
group:test,staff
class:batch,serial
user:*,grid
```

In this example, a job that is being migrated from cluster `master1` to the peer `slave1` will have its credentials mapped according to the contents of the `omap.dat` file. In this case, a job submitted by user `joe` on `master1` will be executed under the user account `jsmith` on peer `slave1`. Any credential that is not found in the mapping file will be passed to the peer as submitted. In the case of the user credential, all users other than `joe` and `steve` will be remapped to the user `grid` due to the wildcard matching.

Because the `OMAP` attribute is specified as a URL, multiple methods can be used to obtain the mapping information. In addition to the file protocol shown in the example above, `exec` can be used.

Note that there is no need to use the credential mapping facility to map all credentials. In some cases, a common user space exists but it is used to map all classes/queues on the source side to a single queue on the destination side. Likewise, for utilization tracking purposes, it may be desirable to map all source account credentials to a single cluster-wide account.

### 23.11.3 Source and Destination Side Credential Mapping

Credential mapping can be implemented on the source cluster, destination cluster, or both. A source cluster may want to map all user names for all outgoing jobs to the name `generaluser` for security purposes, and a destination cluster may want to remap all incoming jobs from this particular user to the username `cluster2` and the QoS `grid`.

### 23.11.4 Preventing User Space Collisions


In some cases, a cluster might receive jobs from two independent clusters where grid wide username distinctiveness is not guaranteed. In this case, credential mapping can be used to ensure the uniqueness of each name. With credential mapping files, this can be accomplished using the `<DESTINATION_CREDENTIAL>` wildcard asterisk (\*) character. If specified, this character will be replaced with the exact `<SOURCE_CREDENTIAL>` when generating the destination credential string. For example, consider the following configuration:

```
SCHEDCFG[master1] MODE=normal
RMCFG[slave1] OMAP=file:///opt/moab/omap.dat  FLAGS=client
...
```

```
user:*,cl_*
group:*,*_grid
account:*,temp_*
```

This configuration will remap the usernames of all jobs coming in from the peer `slave1`. The username `john` will be remapped to `cl_john`, the group `staff` will be remapped to `staff_grid` and the account `demo` will be remapped to `temp_demo`.

## 23.12 Grid Data Management

 This method of data staging has been deprecated in Moab Workload Manager 10.0.1 and will be removed from the product in a future release. See [Data Staging Example](#) for information about the new method of staging data.

In this topic:

- 23.12.1 Grid Data Management Overview - page 979
- 23.12.2 Configuring Peer Data Staging - page 979
  - 23.12.2.A Simple Configuration - page 979
  - 23.12.2.B Advanced Configuration - page 980
- 23.12.3 Peer-to-Peer SCP Key Authentication - page 981
  - 23.12.3.A Generate SSH Key on Source Peer - page 981
- 23.12.4 Diagnostics - page 982

## 23.12.1 Grid Data Management Overview

Moab provides a highly generalized data manager interface that can allow both simple and advanced data management services to be used to migrate data amongst peer clusters. Using a flexible script interface, services such as *scp*, *NFS*, and *gridftp* can be used to address data staging needs. This feature enables a Moab peer to push job data to a destination Moab peer.

## 23.12.2 Configuring Peer Data Staging

Moab offers a simple, automatic configuration, as well as advanced configuration options. At a high level, configuring data staging across a peer-to-peer relationship consists of configuring one or more storage managers, associating them with the appropriate peer resource managers, and then specifying data requirements at the local level—when the job is submitted.

To use the data staging features, you must specify the `--with-grid` option at `./configure` time. After properly configuring data staging, you can submit a job to the peer with any user who has SSH keys set up and Moab will automatically or implicitly stage back the standard out and standard error files created by the job. Files can be implicitly staged in or out before a job runs by using the `mstagein` or `mstageout` options of `msub`.

### 23.12.2.A Simple Configuration

Moab automatically does most of the data staging configuration based on a simplified set of parameters (most common defaults) in the configuration file (`moab.cfg`).

Do the following to configure peer data staging:

1. Configure at least two Moab clusters to work in a grid. Refer to information throughout [Moab Workload Manager for Grids](#) for help on configuring Moab clusters to work together as peers in a grid.

2. [Set up SSH keys](#) so that users on the source grid peer can SSH to destination peers without the need for a password.
3. Make necessary changes to the `moab.cfg` file of the source grid peer to activate data staging, which involves creating a new data resource manager definition within Moab. The resource manager provides data staging services to existing peers in the grid. By defining the data resource manager within the `moab.cfg`, Moab automatically sets up all of the necessary data staging auxiliary scripts.

Use the following syntax for defining a data resource manager:

```
RMCFG[<RMName>] TYPE=NATIVE RESOURCETYPE=STORAGE
VARIABLES=DATASPACEUSER=<DataSpaceUser>,DATASPACE=DATASPACEUSER, DATASPACE=DATASPACEUSER, DATASPACE=DATASPACEUSER
SERVER=<DataServer>
```

- `<RMName>`: Name of the RM (defined as a storage RM type by `RESOURCETYPE=STORAGE`).
- `<DataSpaceUser>`: User used to SSH into `<DataServer>` to determine available space in `<DataSpaceDir>`. Moab runs a command similar to the following:
 

```
ssh <DataServer> -l <DataSpaceUser> df <DataSpaceDir>
```
- `<DataSpaceDir>`: Directory where staged data is stored.
- `<DataServer>`: Name of the server where `<DataSpaceDir>` is located.

Define the following URLs:

```
RMCFG[data] CLUSTERQUERYURL=exec://$TOOLSDIR/grid/cluster.query.dstage.pl
RMCFG[data] SYSTEMMODIFYURL=exec://$TOOLSDIR/grid/system.modify.dstage.pl
RMCFG[data] SYSTEMQUERYURL=exec://$TOOLSDIR/grid/system.query.dstage.pl
RMCFG[data] RMINITIALIZEURL=exec://$TOOLSDIR/grid/setup.config.pl
```

4. Associate the data resource manager with a peer resource manager:

```
RMCFG[remote_data] TYPE=NATIVE RESOURCETYPE=STORAGE
VARIABLES=DATASPACEUSER=datauser, DATASPACE=DATASPACEUSER, DATASPACE=DATASPACEUSER, DATASPACE=DATASPACEUSER
RMCFG[remote_cluster] TYPE=MOAB SERVER=clusterhead:42559 DATARM=remote_data
```

5. Restart Moab to finalize changes. You can use the `mschedctl -R` command to cause Moab to automatically restart and load the changes.

When restarting, Moab recognizes the added configuration and runs a Perl script in the Moab tool directory that configures the external scripts (also found in the tools directory) that Moab uses to perform data staging. You can view the data staging configuration by looking at the `config.dstage.pl` file in `$MOABHOMEDIR/etc`.

### 23.12.2.B Advanced Configuration

If you need a more customized data staging setup, contact your account representative.



### 23.12.3 Peer-to-Peer SCP Key Authentication

In order to use `scp` as the data staging protocol, we will need to create SSH keys that allow users to copy files between the two peers, without the need for passwords. For example, if `UserA` is present on the source peer, and his counterpart is `UserB` on the destination peer, then `UserA` will need to create an SSH key and configure `UserB` to allow password-less copying. This will enable `UserA` to copy files to and from the destination peer using Moab's data staging capabilities.

Another common scenario is that several users present on the source peer are mapped to a single user on the destination peer. In this case, each user on the source peer will need to create keys and set them up with the user at the destination peer. Below are steps that can be used to setup SSH keys among two (or more) peers:

**i** These instructions were written for [OpenSSH version 3.6](#) and might not work correctly for older versions.

#### 23.12.3.A Generate SSH Key on Source Peer

As the user who will be submitting jobs on the source peer, run the following command:

```
ssh-keygen -t rsa
```

You will be prompted to give an optional key. Just hit return and ignore this or other settings. When finished, this command will create two files `id_rsa` and `id_rsa.pub` located inside the user's `~/.ssh/` directory.

#### Copy the Public SSH Key to the Destination Peer

Transfer the newly created public key (`id_rsa.pub`) to the destination peer:

```
scp ~/.ssh/id_rsa.pub ${DESTPEERHOST}:~
```

#### Disable Strict SSH Checking on Source Peer (Optional)

By appending the following to your `~/.ssh/config` file you can disable SSH prompts that ask to add new hosts to the 'known hosts file.' (These prompts can often cause problems with data staging functionality.) Note that the `${DESTPEERHOST}` should be the name of the host machine running the destination peer:

```
Host ${DESTPEERHOST}
  CheckHostIP no
  StrictHostKeyChecking no
  BatchMode yes
```

#### Configure Destination Peer User

Now, log in to the destination peer as the destination user and set up the newly created public key to be trusted:

```
ssh ${DESTPEERUSER}@${DESTPEERHOST}
mkdir -p .ssh; chmod 700 .ssh
cat id_rsa.pub >> .ssh/authorized_keys
chmod 600 .ssh/authorized_keys
rm id_rsa.pub
```

If multiple source users map to a single destination user, then repeat the above commands for each source user's SSH public key.

### Configure SSH Daemon on Destination Peer

Some configuration of the SSH daemon may be required on the destination peer. Typically, this is done by editing the `/etc/ssh/sshd_config` file. To verify correct configuration, see that the following attributes are set (not commented):

```
---
RSAAuthentication    yes
PubkeyAuthentication yes
---
```

If configuration changes were required, the SSH daemon will need to be restarted:

```
/etc/init.d/sshd restart
```

### Validate Correct SSH Configuration

If all is properly configured, if you issue the following command source peer it should succeed without requiring a password:

```
scp ${DESTPEERHOST}:/etc/motd /tmp/
```

## 23.12.4 Diagnostics

Verify data staging is properly configured by using the following diagnostic commands:

- **mdiag -R -v**: Displays the status of the storage manager. Verify that you set up the necessary URLs.

```
> mdiag -R -v data
diagnosing resource managers
RM[data]      State: Active  Type: NATIVE  ResourceType: STORAGE
Server:      keche
Timeout:     30000.00 ms
Cluster Query URL:  exec://$TOOLSDIR/grid/cluster.query.dstage.pl
RM Initialize URL:  exec://$TOOLSDIR/grid/setup.config.pl
System Modify URL:  exec://$TOOLSDIR/grid/system.modify.dstage.pl
System Query URL:  exec://$TOOLSDIR/grid/system.query.dstage.pl
Nodes Reported:    1 (scp://keche//tmp/)
Partition:      SHARED
Event Management: (event interface disabled)
```

```

Variables:          DATASPACEUSER=root,DATASPCEDIR=/tmp
RM Languages:      NATIVE
RM Sub-Languages:  -

```

- **checknode -v**: Executing this on the storage node displays the data staging operations associated with the node and its disk usage.

**i** The number of bytes transferred for each file is currently not used.

```

> checknode -v scp://keche//tmp/
node scp://keche//tmp/
State:      Idle (in current state for 00:00:13)
Configured Resources: DISK: 578G
Utilized Resources:  DISK: 316G
Dedicated Resources: ---
  MTBF(longterm):    INFINITY  MTBF(24h):    INFINITY
Active Data Staging Operations:
  job      native.2  complete (1 bytes transferred)
(/home/brian/stage.txt)
  job      native.3  pending (1 bytes)  (/home/brian/stage.txt)
Dedicated Storage Manager Disk Usage: 0 of 592235 MB
Cluster Query URL:  exec://$TOOLSDDIR/grid/cluster.query.dstage.pl
Partition: SHARED Rack/Slot: ---
Flags:      rmdetected
RM[data]:   TYPE=NATIVE
EffNodeAccessPolicy: SHARED
Total Time: 00:12:15 Up: 00:12:15 (100.00%) Active: 00:00:00 (0.00%)
Reservations: ---

```

- **mdiag -n**: Displays the state of the storage node.

```

> mdiag -n
compute node summary
Name                State   Procs   Memory   Opsys
compute1            Idle   4:4     3006:3006  linux
compute2            Down   0:4     3006:3006  linux
scp://keche//tmp/   Idle   0:0     0:0       -
-----            ---    4:8     6012:6012  -----
Total Nodes: 3 (Active: 0 Idle: 2 Down: 1)

```

- **checkjob -v**: Displays the status of the staging request.

**i** The remaining time and size of the file information is currently not used. The information should only be used to see file locations and whether the file has been staged or not.

```

> checkjob -v jobid
...
Stage-In Requirements:
  localhost:/home/brian/stage.txt => keche:/tmp/staged.txt size:0B
status:[NONE] remaining:00:00:01
  Transfer URL: file:///home/brian/stage.txt,ssh://keche/tmp/staged.txt
...

```

To ensure that SCP key authentication is properly configured, the following conditions must be met:

- Moab is running as root.
- You are able to issue the following command as the root user without being prompted for a password:

```
su - <DATASPACEUSER> -c "/usr/bin/ssh <destination host> -l <DATASPACEUSER> 'df -k //tmp/ 2>&1 || echo FAILED'"
```

- You can SSH `<destination host>` without a password.
- The `dataSpaceLocalUser` and `dataSpaceMappedUser` variables in your `/opt/moab/etc/config.dstage.pl` script are set to the same username you assigned through `<DATASPACEUSER>`.

## 23.13 Accounting and Allocation Management

In this topic:

[23.13.1 Peer-to-Peer Accounting Overview - page 984](#)

[23.13.2 Peer-to-Peer Allocation Management - page 985](#)

### 23.13.1 Peer-to-Peer Accounting Overview

When Moab is used to manage resources across multiple clusters, there is a greater need to track and enforce the resource sharing agreements between the resource principals.

The Moab Accounting Manager is an accounting management system that provides usage tracking, charge accounting, and allocation enforcement for resource or service usage in cloud and technical computing environments. It acts like a bank in which credits are deposited into accounts with constraints designating which entities can access the account. As resources or services are utilized, accounts are charged and usage recorded. MAM supports familiar operations such as deposits, withdrawals, transfers, and refunds and provides balance and usage feedback to users, managers, and system administrators. See [5.5 Accounting, Charging, and Allocation Management - page 418](#) for more information.

MAM can be used as a real-time debiting system where jobs are charged at the moment of completion. When used in a multi-site (grid) environment, MAM facilitates trust by allowing lending organizations to manage what the costing rules are for usage of their resources and job submitters to determine how much their job will cost them before they start,

ensuring all parties can agree to the transaction and giving each party a first-hand accounting record.

If the clusters are within a common administrative domain and have a common user space, then a single Moab Accounting Manager will suffice to manage the project allocation and accounting. This works best in Moab Grid Control / Moab Grid Member grids.

### 23.13.2 Peer-to-Peer Allocation Management

The following steps provide an example of setting up the Moab Accounting Manager to manage the allocation and accounting for a multiple cluster grid within a single administrative domain.

First you will need to install Moab Accounting Manager and its database on one or more head nodes. The following is a sample installation. See 'Installing Moab Accounting Manager' in the *Moab Installation and Configuration Guide* for more information.

```
# Install Prerequisites (Perl with suidperl, PostgreSQL, libxml2, ...)
[root] yum install perl perl-suidperl postgresql postgresql-libs postgresql-devel
postgresql-server libxml2 libxml2-devel ncurses-devel readline-devel openssl
# Unpack the tarball
[root] passwd adaptive
[adaptive] mkdir ~/src
[adaptive] cd ~/src
[adaptive] gzip -cd mam-10.0.0.tar.gz | tar xvf -
[adaptive] cd mam-10.0.0
# Install
[adaptive] ./configure
[adaptive] make
[root] make deps
[root] make install
# Configure, create and bootstrap the database
[root] service postgresql initdb
[postgres] echo "host    all             all             192.168.1.1     255.255.255.255
trust" >>var/lib/pgsql/data/pg_hba.conf
[postgres] /usr/local/pgsql/bin/postmaster -i -D /usr/local/pgsql/data >var/log/pgsql
2>&1 &
[postgres] createuser adaptive
[adaptive] createdb mam
[adaptive] psql mam < hpc.sql
# Startup the mam server daemon
[root] cp contrib/init.d/mam/redhat /etc/init.d/mam
[root] chmod +x /etc/init.d/mam
[root] service mam start
```

---

#### Related Topics

- [Grid Credential Management](#)

## 23.14 Grid Security

### Secret Key Based Server Authentication

Secret key based security is required in order for the grid to work. It is enabled in the `moab-private.cfg` file. Configuration of `moab-private.cfg` is covered throughout the grid configuration documentation, as well as in [Appendix E: Security](#).

## 23.15 Grid Diagnostics and Validation

In this topic:

[23.15.1 Peer Management Overview - page 986](#)

[23.15.2 Peer Diagnostic Overview - page 986](#)

### 23.15.1 Peer Management Overview

- Use `mdiag -R` to view interface health and performance/usage statistics.
- Use `mrmctl` to enable/disable peer interfaces.
- Use `mrmctl -m` to dynamically modify/configure peer interfaces.

### 23.15.2 Peer Diagnostic Overview

- Use `mdiag -R` to diagnose general RM interfaces.
- Use `mdiag -S` to diagnose general scheduler health.
- Use `mdiag -R -V job <RMID>` to diagnose peer-to-peer job migration:

```
> mdiag -R -V job peer1
```

- Use `mdiag -R -V data <RMID>` to diagnose peer-to-peer data staging.
- Use `mdiag -R -V cred <RMID>` to diagnose peer-to-peer credential mapping.

## Chapter 24: Data Staging

Sometimes you might need a job to process data that resides at another site. With the proper configuration, you can submit your job with the requirement that it copies data from the external site to yours and, if needed, copy the job's resulting data out to the external site for its owner to use. Data staging is an out-of-band method of moving data without reserving compute nodes or other resources for it.

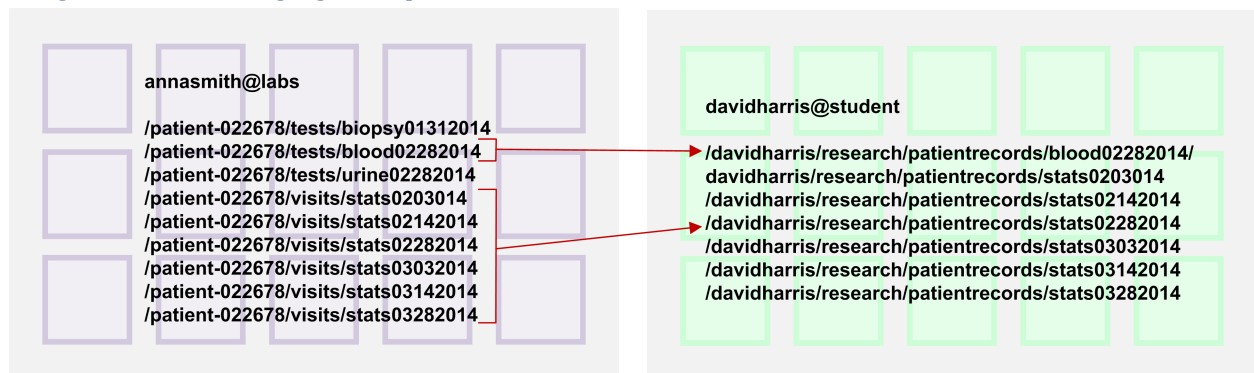
In this section:

- [24.1 Data Staging Example - page 987](#)
- [24.2 Data Staging Tasks - page 988](#)
- [24.3 Data Staging References - page 1012](#)

### 24.1 Data Staging Example

In the example below, which will appear throughout the chapter, a university researcher needs the results of tests done at a hospital to run his job. User  `davidharris`  on the  `student`  server of the university submits a job called  `Moab . 1`  that requires several files stored by user  `annasmith`  on the  `research`  server of the hospital.  `davidharris`  submitted  `Moab . 1`  with certain options in place that instruct  `Moab`  to copy the files to the  `/student/davidharris/research/patientrecords/`  directory on the  `student`  server prior to starting the job.

*Image 24-1: Data Staging Example*



Moab currently supports the following data staging use cases: 1) Staging data to or from a shared file system, 2) Staging data to or from local node storage on a single compute node, and 3) Staging data to or from a shared file system on an unspecified cluster – resolved at job migration – in a grid configuration.

Before you can submit data staging jobs, you must configure certain generic metrics in your partitions, job templates, and the data staging submit filter for data staging scheduling, throttling, and policies.

Moab uses Linux file transfer utilities to stage the data and includes data staging reference scripts that support the `scp` and `rsync` Linux file transfer utilities. The scripts will work for standard installations, but you can customize the script to support data staging to and from an external staging server, the Moab server itself, or a local compute node, depending on your implementation. You can also customize your own script for other file transfer utilities, such as Aspera.

Once you configure your system to support data staging, you can begin creating data staging jobs by attaching the `--stagein`, `--stageinfile`, `--stageinsize`, `--stageout`, `--stageoutfile`, and `--stageoutsize` options to your `msub` commands. See [Staging Data](#) for more information.

---

### Related Topics

- [3.7.30 msub - page 239](#)
- [3.7.30.H Applying the msub Submit Filter - page 259](#)

## 24.2 Data Staging Tasks

In this section:

- [24.2.1 Configuring the SSH Keys for the Data Staging Transfer Script - page 989](#)
- [24.2.2 Configuring Data Staging - page 992](#)
- [24.2.3 Staging Data to or from a Shared File System - page 994](#)
- [24.2.4 Staging Data to or from a Shared File System in a Grid - page 999](#)
- [24.2.5 Staging Data To or From a Compute Node - page 1005](#)
- [24.2.6 Configuring Data Staging with Advanced Options - page 1009](#)

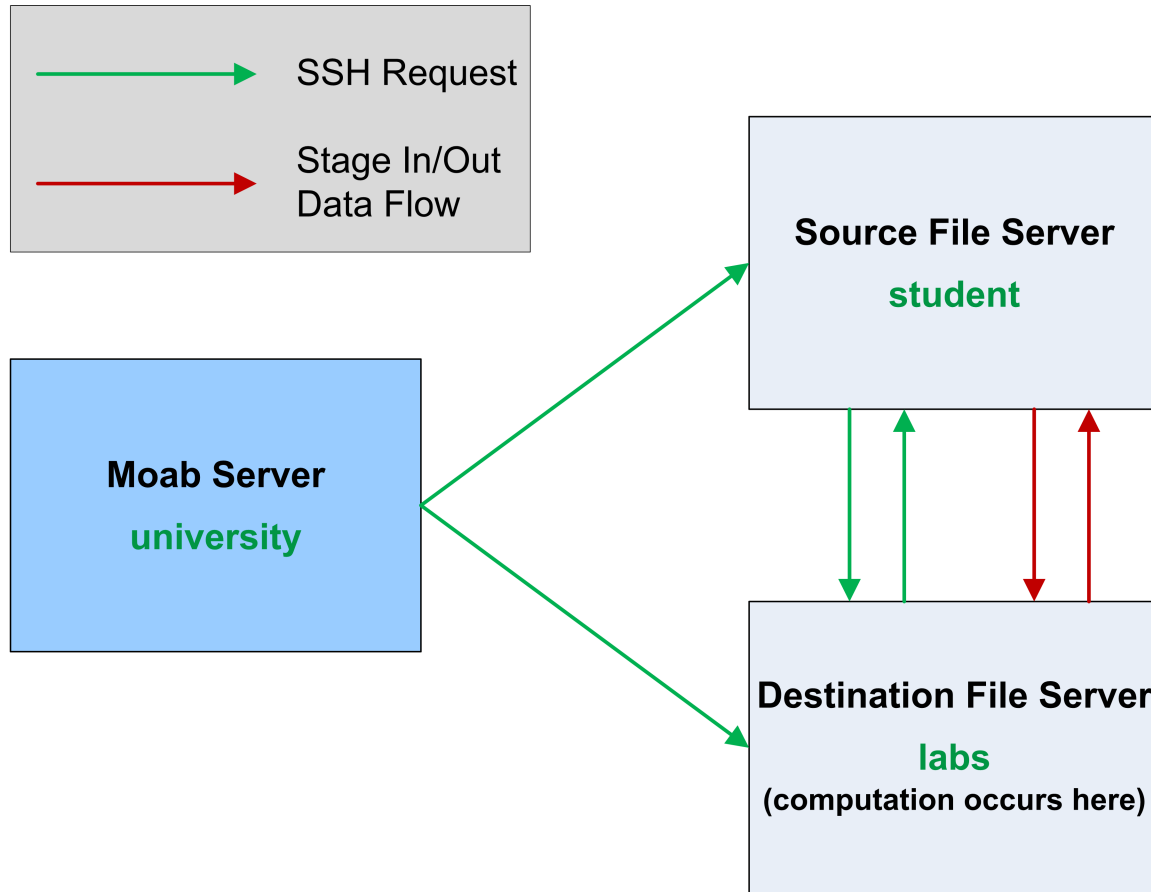


## 24.2.1 Configuring the SSH Keys for the Data Staging Transfer Script

For data staging to work correctly, you must configure SSH keys to allow the data staging scripts to run without passphrases. In the sample data staging server configuration shown in the image below,  `davidharris`  on the  `student`  server stages data from the source server  `student`  to the destination server  `labs` . The computation occurs on the  `labs`  server before Moab stages the output data from  `labs`  back to  `student` .

The image below demonstrates the SSH connections necessary and how you should configure your SSH keys:

*Image 24-2: Data Staging Server Configuration*



For more details on generating keys, see the [ssh-keygen man page](#) and [SSH login without password](#).

### To Configure the SSH Keys for the Data Staging Transfer Script

1. Generate a new SSH key on the Moab server (`university`) if one does not already exist. To do so, run each of the following steps:
  - a. Run `ssh-keygen` to generate a public and private rsa key pair:

```
davidharris@university]$ ssh-keygen
```

- b. Enter the name of the file where you want to store the key, or you can accept the default location:

```
/home/davidharris/.ssh/id_rsa
```

- c. When prompted for a passphrase, leave it blank and press `Enter`. Repeat when prompted to retype passphrase.
2. Install the public key on the source and destination hosts. Note that in this example the source host is `student` and the destination host is `labs`:
  - a. Copy the university public key to `student`. Answer `yes` to continue connecting:

```
[davidharris@university]$ ssh-copy-id -i ~/.ssh/id_rsa.pub student
```

- b. Copy the university public key to `labs`. Answer `yes` to continue connecting:

```
[davidharris@university]$ ssh-copy-id -i ~/.ssh/id_rsa.pub labs
```

**i** The next two steps generate a key-pair for each node. It is acceptable to generate a single key-pair and install it on each node. It does not matter where the key-pair is generated, so long as it is compatible with the SSH client/server.

3. Generate a key pair on the source host (`student`) and install the public key generated to the destination host (`labs`). When prompted for a passphrase, leave it blank and press `Enter`. Repeat when prompted to retype passphrase.

```
[davidharris@student]$ ssh-keygen
[davidharris@student]$ ssh-copy-id -i ~/.ssh/id_rsa.pub labs
```

4. Generate a key pair on the destination host (`labs`) and install the public key generated to the source host (`student`). When prompted for a passphrase, leave it blank and press `Enter`. Repeat when prompted to retype passphrase.

```
[davidharris@labs]$ ssh-keygen
[davidharris@labs]$ ssh-copy-id -i ~/.ssh/id_rsa.pub student
```

5. Ensure that each user who will run data staging jobs has read and write permissions on each source and destination server.
6. Test the configuration. To do so:
  - a. Install the modules required to run the data staging scripts. `python-paramiko` is required for data staging, but `python-mock` is only required if you intend to run the unit test:

```
> yum install python-paramiko python-mock
```

- b. Transfer a file from the source host to the destination host to verify that the keys work for the users configured. To do so, run `/opt/moab/tools/data-staging/ds_move_scp --test=<source>%<destination>` if you use `scp` or `/opt/moab/tools/data-staging/ds_move_rsync --test=<source>%<destination> script` if you use `rsync`. `<source>%<destination>` is configured the same way as the `--stagein` and

--stageout options for msub; for help configuring your source and destination, see [Staging a File or Directory](#).

```
[davidharris@university]$ /opt/moab/tools/data-staging/ds_move_rsync --
test=davidharris@student:/tmp/test%davidharris@labs:/tmp
```

- c. In the same way, transfer a file from the destination host to the source host to verify that the keys work for the users configured:

```
[davidharris@university]$ /opt/moab/tools/data-staging/ds_move_rsync --
test=davidharris@labs:/tmp/test%davidharris@student:/test_processed
```

---

## Related Topics

- [24.1 Data Staging Example - page 987](#)
- [24.2.2 Configuring Data Staging - page 992](#)

## 24.2.2 Configuring Data Staging

You must modify your Moab configuration to enable data staging. In addition to the configuration steps described below, you might also consider customizing the configuration (including the associated scripts) to meet your site's specific needs.

For advanced configuration steps and options, see [Configuring Data Staging with Advanced Options](#).

### To Configure Data Staging

1. Verify that your firewall and network are correctly configured to allow the scripts to operate as designed.
2. If you have not already done so, install the modules required to run the data staging scripts. `python-paramiko` is required for data staging, but `python-mock` is only required if you intend to run the unit test.

```
> yum install python-paramiko python-mock
```

3. If you have not already, follow the instructions found in [Configuring the SSH Keys for the Data Staging Transfer Script](#).
4. Ensure that the data staging scripts are installed on your system. To do so, list the contents of the `/opt/moab/tools/data-staging` directory. You should see the data staging README file, reference scripts, and other related files.

```
> ls -l /opt/moab/tools/data-staging
```

You can copy and modify the reference scripts and configuration files to meet your specific needs. See the README file packaged in the `data-staging` directory for information about modifying these files.

5. Open your `moab.cfg` file for editing and do each of the following tasks:
  - a. Configure the data staging `msub` filter, located in `/opt/moab/tools/data-staging` by default, as a client-side filter. See [Applying the msub Submit Filter](#) for more information.

```
SUBMITFILTER /opt/moab/tools/data-staging/ds_filter
```

The data staging filter checks the `msub` argument syntax to verify that the arguments make sense and are consistent; attempts a dry run connection via SSH and the file transfer utility to ensure that keys exist for the user on the necessary systems; and attempts to determine the size of the data that will be transferred.

You can customize the script to meet your specific needs; the file contains detailed comments illustrating its default behavior to facilitate its modification. If you replace or modify the submit filter, it is your responsibility to ensure that the same functionality described in the paragraph above is present in your filter.

Note that this filter has the `DEFAULT_TEMPLATE` name, which should match the name of the master data staging template in `moab.cfg`. For more information, see [Configuring Data Staging with Advanced Options](#).

- b. Set the data staging bandwidth gmetric (`DATASTAGINGBANDWIDTH_MBITS_PER_SEC`) on each partition associated with an RM to the rate at which its network to be used for data staging transfers data in megabits per second (see [Per-Partition Settings](#) for more information). Moab will use the specified rate and the data staging size specified at job submission (see [Stage in or out File Size](#) for more information) to determine how long staging the data will take and to schedule the job as soon after data staging completes as possible.

#### Example 24-1: Non-grid

```
RMCFG[torque] Type=pbs
PARCFG[torque] GMETRIC[DATASTAGINGBANDWIDTH_MBITS_PER_SEC]=58
```

*Partition torque has a transfer rate of 58 megabits per second. Moab uses the rate when it estimates the time it will take to stage data in and determine when to schedule the job that will use the data.*

#### Example 24-2: Grid

```
RMCFG[m1] type=Moab
PARCFG[m1] GMETRIC[DATASTAGINGBANDWIDTH_MBITS_PER_SEC]=100
```

*Partition m1 has a transfer rate of 100 megabits per second. Moab uses the rate when it estimates the time*

*it will take to stage data in and determine when to schedule the job that will use the data.*

- c. Set the bandwidth generic resource on all nodes to limit the total number of concurrent data staging jobs in your system:

```
NODECFG[GLOBAL] GRES=bandwidth:10
```

*Data staging jobs can use up to 10 units of bandwidth on the system. You can specify the number of units consumed by each data staging job when you configure the data staging job templates.*

- d. Configure moab with 'JOBMIGRATEPOLICY JUSTINTIME'.

**i** DataStaging requires 'JOBMIGRATEPOLICY JUSTINTIME' to ensure the workflow job IDs are not altered upon submission.

- 6. Install the msub client filter on all client submission hosts.

---

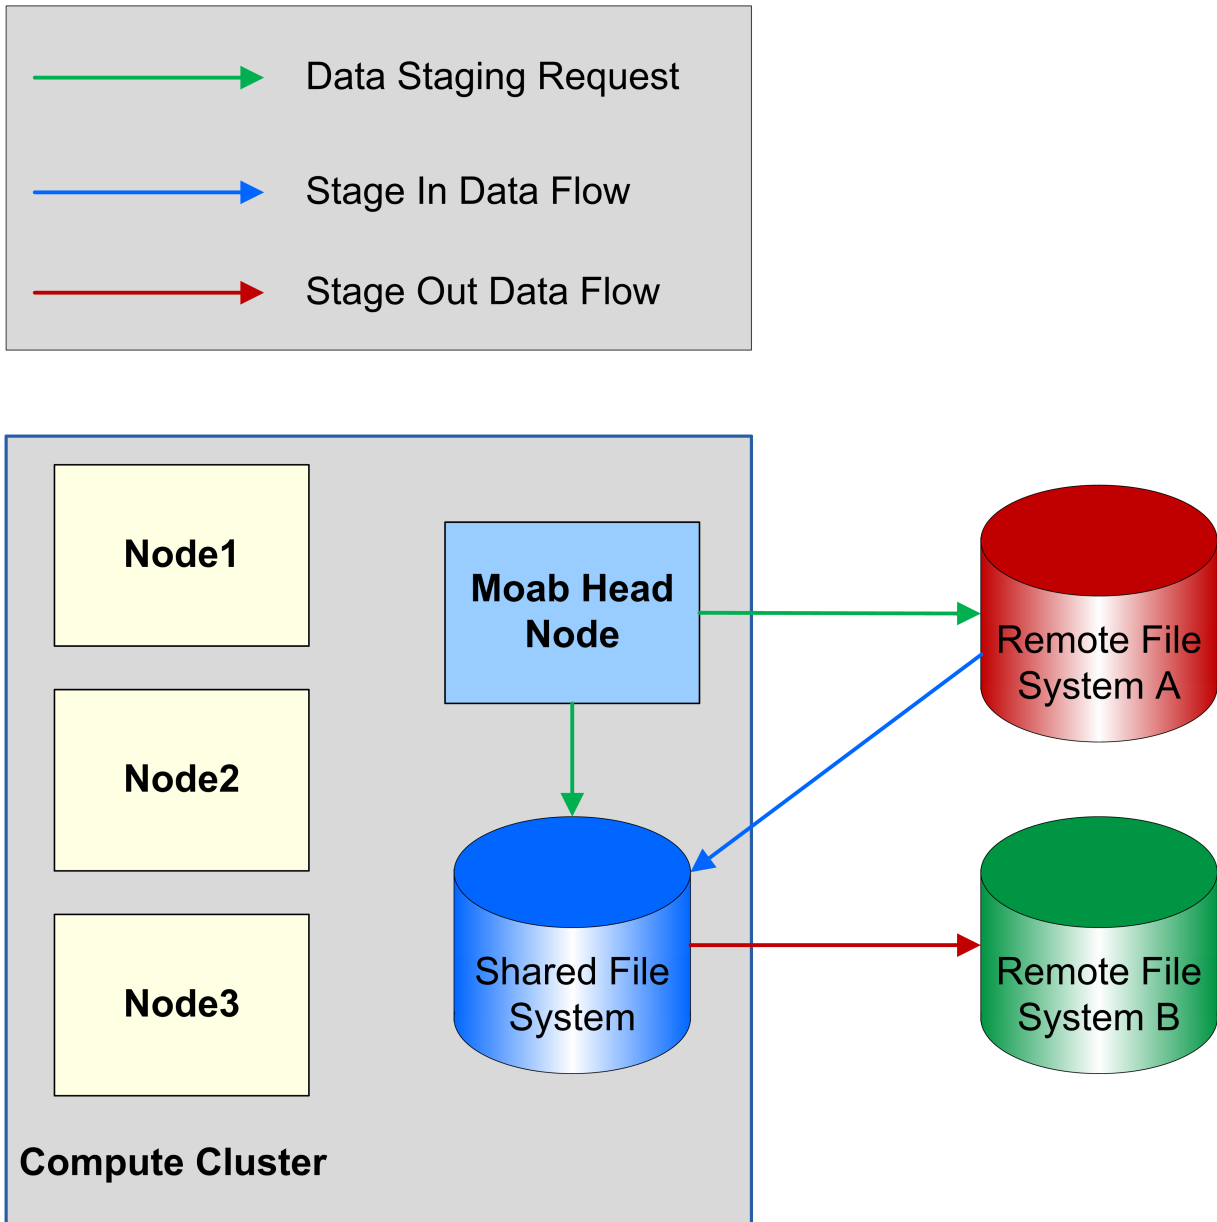
### Related Topics

- [24.1 Data Staging Example - page 987](#)

## 24.2.3 Staging Data to or from a Shared File System

In the most common data staging use case, the cluster utilizes a shared file system between all compute nodes. This type of data staging makes data stored outside of the cluster available to a job that will run on any set of nodes in the cluster. At the time of submission, you must specify where Moab will obtain the data with a username, host name, and path to a file or directory and where on the shared file system Moab will store the data. After the job runs, you can also copy data from the shared file system back to a remote file system.

*Image 24-3: Data Staging To or From a Shared File System*



## To Stage Data to or from a Shared File System

1. If you have not already done so, configure your SSH keys and `moab.cfg` to support data staging. See [Configuring the SSH Keys for the Data Staging Transfer Script](#) and [Configuring Data Staging](#) for more information.
2. Create your job templates for data staging jobs in `moab.cfg`. The templates in the example below create a compute job that stages data in before it starts and stages data

out when it completes. For more information about creating job templates, see [About Job Templates](#).

- a. Create a selectable master template, called `ds` in the example below, that creates a stage in and stage out system job. This name should match the `DEFAULT_TEMPLATE` value in `ds_config.py`. See [Configuring Data Staging with Advanced Options](#) for more information.
- b. For the data staging in job template, called `dsin` in the example below, specify that it will create a data staging job by setting `DATASTAGINGJOB` to `TRUE`. Note that the name of this job template must match the name of the data stage in job template referenced in the master template.
- c. Set the bandwidth `GRES` to the amount of bandwidth a single stage in job should use. This indicates how many of the bandwidth units specified with `NODECFG[GLOBAL]` in [Configuring Data Staging](#) a data staging job with this template should consume.
- d. Add `FLAGS=GRESONLY` to indicate that this data staging job does not require any compute resources.
- e. Create a trigger that executes the `ds_move_scp`, `ds_move_rsync`, or `ds_move_multiplex` script, depending on which file transfer utility you use. Set the `attacherror`, `objectxmlstdin`, and `user` `FLAGS` to attach any trigger `stderr` as a message to the job, pass the job XML to the script, and indicate that the script should run as the job's user, respectively.

**i** If you use the `rsync` protocol, you can configure your data staging jobs to report the actual number of bytes transferred and the total data size to be transferred. To do so, use the `Sets` attribute to `^BYTES_IN.^DATA_SIZE_IN` for stage in jobs and `^BYTES_OUT.^DATA_SIZE_OUT` for stage out jobs. For example, a stage in trigger would look like the following:

```
JOBCFG[dsin]
TRIGGER=EType=start,AType=exec,Action="/opt/moab/tools/data-staging/ds_move_rsync --stagein",Flags=objectxmlstdin:user:attacherror,Sets=^BYTES_IN.^DATA_SIZE_IN
```

A stage out trigger would look like the following:

```
JOBCFG[dsout]
TRIGGER=EType=start,AType=exec,Action="/opt/moab/tools/data-staging/ds_move_rsync --stageout",Flags=objectxmlstdin:user:attacherror,Sets=^BYTES_OUT.^DATA_SIZE_OUT
```

These variables show up as events if you set your `WIKIEVENTS` parameter to `TRUE`.

- f. Create the stage out job, called `dsout` in the example below, by repeating steps 2b - 2e in a new template. In the example below, this template is called `dsout`. Note that



the name of this job template must match the name of the data stage out job template referenced in the data staging master template.

```

JOBCFG[ds]      TEMPLATEDEPEND=AFTEROK:dsin TEMPLATEDEPEND=BEFORE:dsout
SELECT=TRUE

JOBCFG[dsin]    DATASTAGINGSYSJOB=TRUE
JOBCFG[dsin]    GRES=bandwidth:2
JOBCFG[dsin]    FLAGS=GRESONLY
JOBCFG[dsin]    TRIGGER=EType=start,AType=exec,Action="/opt/moab/tools/data-
staging/ds_move_rsync --stagein",Flags=attacherror:objectxmlstdin:user

JOBCFG[dsout]   DATASTAGINGSYSJOB=TRUE
JOBCFG[dsout]   GRES=bandwidth:1
JOBCFG[dsout]   FLAGS=GRESONLY
JOBCFG[dsout]   TRIGGER=EType=start,AType=exec,Action="/opt/moab/tools/data-
staging/ds_move_rsync --stageout",Flags=attacherror:objectxmlstdin:user

```

3. Create the job using `msub`, adding resources and specifying a script as you normally would. Then configure Moab to stage the data for it. To do so:
  - a. At the end of the command, use the `--stagein/--stageout` option and/or `--stageinfile/--stageoutfile` option.
    - The `--stagein/--stageout` option lets you specify a single file or directory to stage in or out. You must set the option equal to `<source>%<destination>`, where `<source>` and `<destination>` are both `[<user>@]<host>: /<path> / [<fileName>]`. See [Staging a File or Directory](#) for format and details.

**i** If the destination partition is down or does not have configured resources, the data staging workflow submission will fail.

```

> msub --stagein=annasmith@labs:/patient-
022678/%davidharris@university:/davidharris/research/patientrecords
<jobScript>

```

*Moab copies the /patient-022678 directory from the hospital's labs server to the university cluster where the job will run prior to job start.*

- The `--stageinfile/--stageoutfile` option lets you specify a file that contains the file and/or directory name(s) to stage in or out. You must set the option equal to `<path>/<fileName>` of the file. The file must contain at least one line with this format: `<source>%<destination>`, where both `<source>` and `<destination>` are `[<user>@]<host>: /<path> [<fileName>]`. See [Staging Multiple Files or Directories](#) for more information.

**i** If the destination partition is down or does not have configured resources, the data staging workflow submission will fail.

```
> msub --stageinfile=/davidharris/research/recordlist <jobScript>
```

*Moab copies all files specified in the /davidharris/research/recordlist file to the cluster where the job will run prior to job start.*

```
/davidharris/research/recordlist:
```

```
annasmith@labs:/patient-022678/tests/blood02282014%davidharris@university:/davidharris/research/patientrecords/blood02282014
annasmith@labs:/patient-022678/visits/stats02032014%davidharris@university:/davidharris/research/patientrecords/stats02032014
annasmith@labs:/patient-022678/visits/stats02142014%davidharris@university:/davidharris/research/patientrecords/stats02142014
annasmith@labs:/patient-022678/visits/stats02282014%davidharris@university:/davidharris/research/patientrecords/stats02282014
annasmith@labs:/patient-022678/visits/stats03032014%davidharris@university:/davidharris/research/patientrecords/stats03032014
annasmith@labs:/patient-022678/visits/stats03142014%davidharris@university:/davidharris/research/patientrecords/stats03142014
annasmith@labs:/patient-022678/visits/stats03282014%davidharris@university:/davidharris/research/patientrecords/stats03282014
```

*Moab copies the seven patient record files from the hospital's labs server to the university cluster where the job will run prior to job start.*

- b. The `--stageinsize/--stageoutsize` option lets you specify the estimated size of the files and/or directories to help Moab more quickly and accurately calculate the amount of time it will take to stage the data and therefore schedule your job correctly. If you are staging data out, then setting `--stageoutsize` is required. If you provide an integer, Moab will assume the number is in megabytes. To change the unit, add another suffix. See [Stage in or out File Size](#) for more information.

```
> msub --stageinfile=/davidharris/research/recordlist --stageinsize=100 <jobScript>
```

*Moab copies the /davidharris/research/recordlist file, which is approximately 100 megabytes, from the biology node to the host where the job will run prior to job start.*

4. To see the status, errors, and other details associated with your data staging job, run `checkjob -v`. See [checkjob](#) for details.

---

## Related Topics

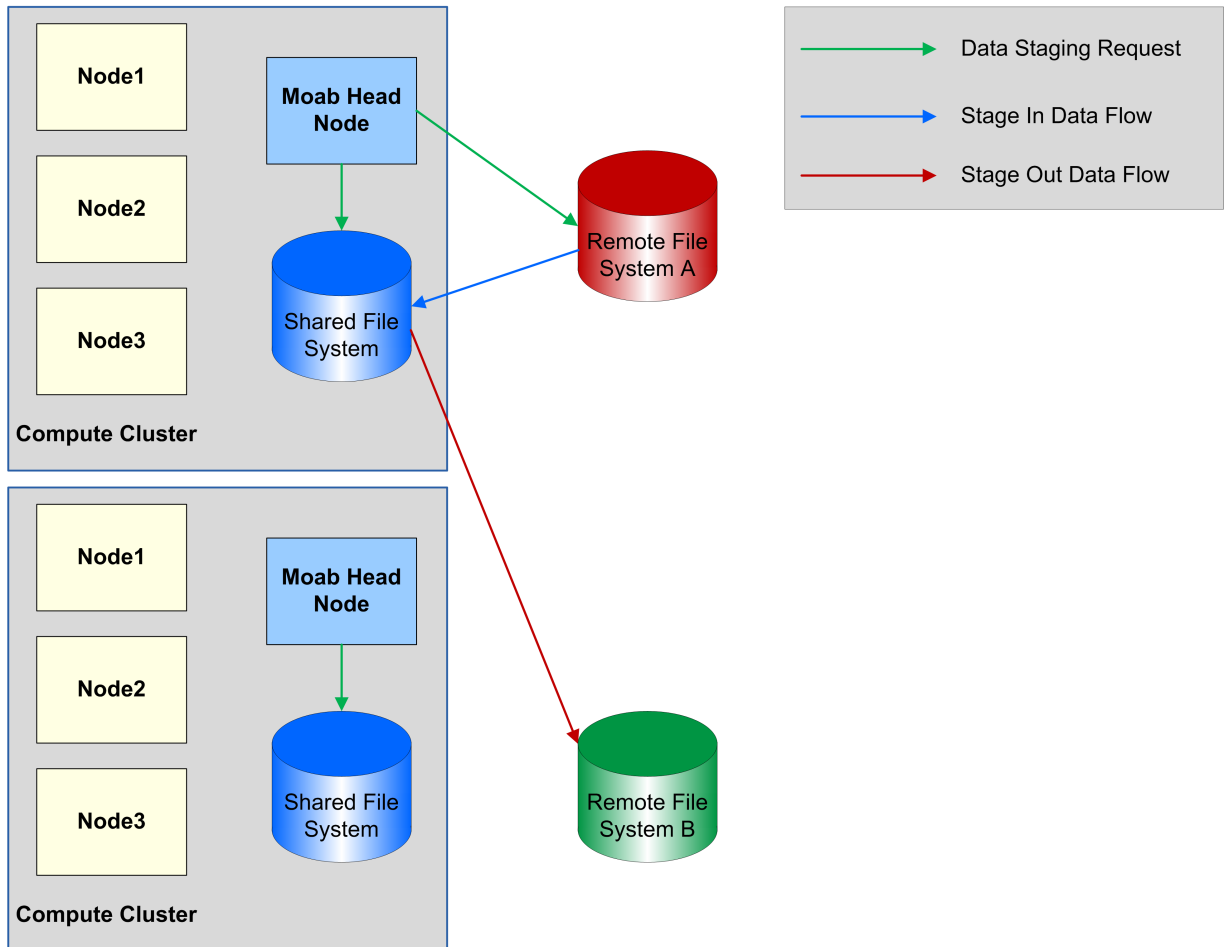
- [24.1 Data Staging Example - page 987](#)
- [24.2.2 Configuring Data Staging - page 992](#)
- [24.2.6 Configuring Data Staging with Advanced Options - page 1009](#)
- [24.2.4 Staging Data to or from a Shared File System in a Grid - page 999](#)

- [24.2.5 Staging Data To or From a Compute Node - page 1005](#)
- 

## 24.2.4 Staging Data to or from a Shared File System in a Grid

You can stage data in an environment where multiple instances of Moab run in a grid configuration. For this type of data staging, each cluster utilizes a shared file system with all compute nodes. This type of data staging will make data available to a job that will run on a set of nodes in one of the clusters in the grid. You must specify where the remote data can be obtained with a username, host name, and path to a file or directory and where on the shared storage Moab will store the data. The remote data source location is known at job submission time, but you must use the `$CLUSTERHOST` placeholder for the host name of the data transfer server on which the job will be scheduled. After the job runs, you can also copy data from the cluster shared file system to a remote file system.

Note that you cannot stage data to or from a local compute node with its own local storage in a grid environment.

*Image 24-4: Data Staging in a Grid*

## To Stage Data to or from a Shared File System in a Grid

1. If you have not already done so, configure your SSH keys and `moab.cfg` to support data staging. See [Configuring the SSH Keys for the Data Staging Transfer Script](#) and [Configuring Data Staging](#) for more information.
2. Create your job templates for data staging jobs in `moab.cfg`. The templates in the example below create a compute job that stages data in before it starts and stages data out when it completes. For more information about creating job templates, see [About Job Templates](#).
  - a. Create a selectable master template, called `ds` in the example below, that creates a stage in and stage out system job. This name should match the `DEFAULT_TEMPLATE` value in `ds_config.py`. For more information, see [Configuring Data Staging with Advanced Options](#).

- b. For the data staging in job template, called `dsin` in the example below, specify that it will create a data staging job by setting `DATASTAGINGJOB` to `TRUE`. Note that the name of this job template must match the name of the data stage in job template referenced in the master template.
- c. Set the staging job template bandwidth `GRES` to the amount of bandwidth a single stage in job should use. This indicates how many of the bandwidth units specified with `NODECFG[GLOBAL]` in [Configuring Data Staging](#) a data staging job with this template should consume.
- d. Set `JOBMIGRATEPOLICY` to `JUSTINTIME`.
- e. Add `FLAGS=GRESONLY` to indicate that this data staging job does not require any compute resources.
- f. Create a trigger that executes the `ds_move_scp`, `ds_move_rsync`, or `ds_move_multiplex` script, depending on which file transfer utility you use. Set the `attacherror`, `objectxmlstdin`, and `user` `FLAGS` to attach any trigger `stderr` as a message to the job, pass the job XML to the script, and indicate that the script should run as the job's user, respectively.

**i** If you use the `rsync` protocol, you can configure your data staging jobs to report the actual number of bytes transferred and the total data size to be transferred. To do so, use the `Sets` attribute to `^BYTES_IN.^DATA_SIZE_IN` for stage in jobs and `^BYTES_OUT.^DATA_SIZE_OUT` for stage out jobs. For example, a stage in trigger would look like the following:

```
JOBCFG[dsin]
TRIGGER=EType=start,AType=exec,Action="/opt/moab/tools/data-staging/ds_
move_rsync --stagein",Flags=objectxmlstdin:user:attacherror,Sets=^BYTES_
IN.^DATA_SIZE_IN
```

A stage out trigger would look like the following:

```
JOBCFG[dsout]
TRIGGER=EType=start,AType=exec,Action="/opt/moab/tools/data-staging/ds_
move_rsync --stageout",Flags=objectxmlstdin:user:attacherror,Sets=^BYTES_
OUT.^DATA_SIZE_OUT
```

These variables show up as events if you set your `WIKIEVENTS` parameter to `TRUE`.

- g. Create the stage out job, called `dsout` in the example below, by repeating steps 2b - 2f in a new template. In the example below, this template is called `dsout`. Note that the name of this job template must match the name of the data stage out job template referenced in the master template.

```
JOBCFG[ds]      TEMPLATEDEPEND=AFTEROK:dsin  TEMPLATEDEPEND=BEFORE:dsout
SELECT=TRUE
```

```

JOBcfg[dsin]    DATASTAGINGSYSJOB=TRUE
JOBcfg[dsin]    GRES=bandwidth:2
JOBcfg[dsin]    FLAGS=GRESONLY
JOBcfg[dsin]    TRIGGER=EType=start,AType=exec,Action="/opt/
moab/tools/data-staging/ds_move_rsync --stagein",Flags=attacherror:objectxmlstdin:user

JOBcfg[dsout]   DATASTAGINGSYSJOB=TRUE
JOBcfg[dsout]   GRES=bandwidth:2
JOBcfg[dsout]   FLAGS=GRESONLY
JOBcfg[dsout]   TRIGGER=EType=start,AType=exec,Action="/opt/
moab/tools/data-staging/ds_move_rsync --stageout",Flags=attacherror:objectxmlstdin:user

```

3. Create the job using *msub*, adding resources and specifying a script as you normally would. Then configure Moab to stage the data for it. To do so:
  - a. At the end of the command, use the `--stagein/--stageout` option and/or `--stageinfile/--stageoutfile` option.
    - The `--stagein/--stageout` option lets you specify a single file or directory to stage in or out. You must set the option equal to `<source>%<destination>`, where `<source>` and `<destination>` are both `[<user>@]<host>: /<path> / [<fileName>]`. See [Staging a File or Directory](#) for format and details.

Note that if you do not know the cluster where the job will run but want the data staged to the same location, you can use the `$CLUSTERHOST` variable in place of a host. If you choose to use the `$CLUSTERHOST` variable, you must first customize the `ds_config.py` file. For more information, see [Configuring the \\$CLUSTERHOST Variable](#) below.

**i** If the destination partition is down or does not have configured resources, the data staging workflow submission will fail.

```
> msub ... --stagein=annasmith@labs:/patient-022678/%\
$CLUSTERHOST:/davidharris/research/patientrecords <jobScript>
```

*Moab copies the /patient-022678 directory from the hospital's labs server to the cluster where the job will run prior to job start.*

- The `--stageinfile/--stageoutfile` option lets you specify a file that contains the file(s) and directory(-ies) to stage in or out. You must set the option equal to `<path>/<fileName>` of the file. The file must contain at least one line with this format: `[<user>@]<host>: /<path> [<fileName>]`. See [Staging Multiple Files or Directories](#) for more information.

**i** If the destination partition is down or does not have configured resources, the data staging workflow submission will fail.

```
> msub ... --stageinfile=/davidharris/research/recordlist <jobScript>
```

*Moab copies all files specified in the /davidharris/research/recordlist file to the cluster where the job will run prior to job start.*

```
/davidharris/research/recordlist:
```

```
annasmith@labs:/patient-
022678/tests/blood02282014%$CLUSTERHOST:/davidharris/research/patientrecords/
blood02282014
annasmith@labs:/patient-
022678/visits/stats02032014%$CLUSTERHOST:/davidharris/research/patientrecords
/stats02032014
annasmith@labs:/patient-
022678/visits/stats02142014%$CLUSTERHOST:/davidharris/research/patientrecords
/stats02142014
annasmith@labs:/patient-
022678/visits/stats02282014%$CLUSTERHOST:/davidharris/research/patientrecords
/stats02282014
annasmith@labs:/patient-
022678/visits/stats03032014%$CLUSTERHOST:/davidharris/research/patientrecords
/stats03032014
annasmith@labs:/patient-
022678/visits/stats03142014%$CLUSTERHOST:/davidharris/research/patientrecords
/stats03142014
annasmith@labs:/patient-
022678/visits/stats03282014%$CLUSTERHOST:/davidharris/research/patientrecords
/stats03282014
```

*Moab copies the seven patient record files from the hospital's labs server to the cluster where the job will run prior to job start.*

- b. The `--stageinsize/--stageoutsize` option lets you specify the estimated size of the files and/or directories to help Moab more quickly and accurately calculate the amount of time it will take to stage the data and therefore schedule your job correctly. If you used the `$CLUSTERHOST` variable to stage in, then setting `--stageinsize` is required. `--stageoutsize` is always required for staging data out. If you provide an integer, Moab will assume the number is in megabytes. To change the unit, add another suffix. See [Stage in or out File Size](#) for more information.

```
> msub ... --stageinfile=/davidharris/research/recordlist --stageinsize=100
<jobScript>
```

*Moab copies the /davidharris/research/recordlist file, which is approximately 100 megabytes, from the biology node to the host where the job will run prior to job start.*

4. To see the status, errors, and other details associated with your data staging job, run `checkjob -v`. See [checkjob](#) for details.

#### 24.2.4.A Configuring the `$CLUSTERHOST` Variable

When you submit a data staging job in a grid environment, you can use the `$CLUSTERHOST` variable instead of specifying a destination if you do not know the cluster where the job will run but want the data staged to the same location. Before the variable

will work correctly, you must first configure it by customizing the `ds_config.py` script to match your unique system.

### Use Case

In a grid with three clusters, you have a partition named `master` where you want all data staged to a host named `gridheadNAS`; a partition named `csdept` where you want all data staged to a host named `fs001.cs.example.edu`; and a partition named `lab` where you want all data staged to a host named `bigfilesystem`.

## To Configure the \$CLUSTERHOST Variable

1. Open the `ds_config.py` file for modification. It is located in `/opt/moab/tools/data-staging/` by default.

```
[moab]$ vi /opt/moab/tools/data-staging/ds_config.py
```

2. Locate the `PARTITION_TO_HOST` parameter:

```
...
PARTITION_TO_HOST = {"partition_1_name": "cluster_1_staging_hostname",
                    "partition_2_name": "cluster_2_staging_hostname",
                    "partition_3_name": "cluster_3_staging_hostname"}
...
```

3. Replace the partition names and associated cluster hostnames with those that match your system. For the use case provided above, you would customize it the following way:

```
...
PARTITION_TO_HOST = {"master": "gridheadNAS",
                    "csdept": "fs001.cs.example.edu",
                    "lab": "bigfilesystem"}
...
```

---

### Related Topics

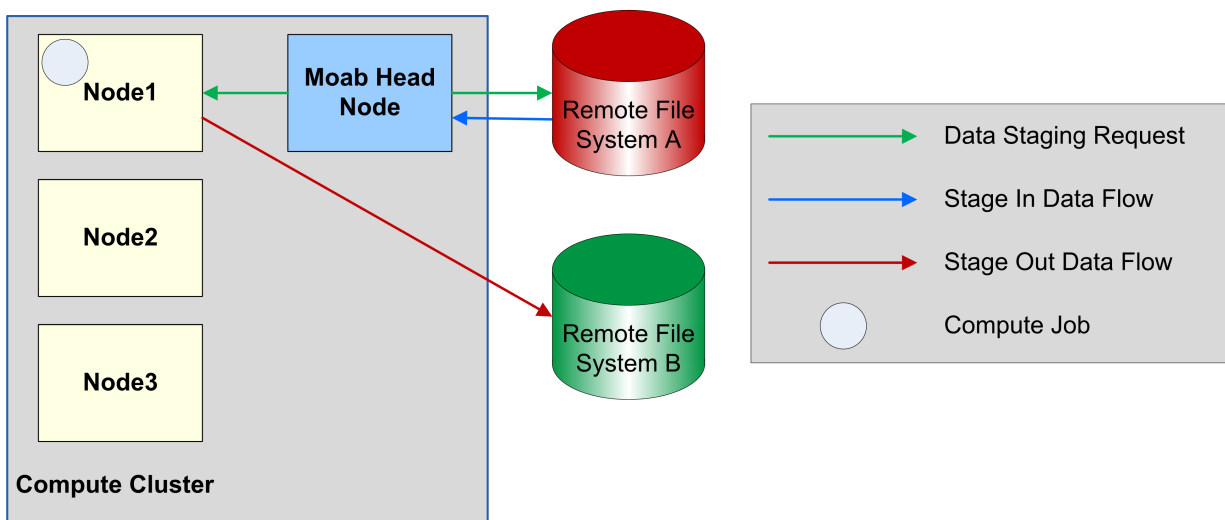
- [24.1 Data Staging Example - page 987](#)
- [24.2.2 Configuring Data Staging - page 992](#)
- [24.2.6 Configuring Data Staging with Advanced Options - page 1009](#)
- [24.2.4 Staging Data to or from a Shared File System in a Grid - page 999](#)



## 24.2.5 Staging Data To or From a Compute Node

You can stage data to or from a local compute node in an environment where each node on the cluster has local storage. This type of data staging will make data stored outside the cluster available to a job that will run on a single node in the cluster. You must specify the username, host name, and path to a file or directory and a location on the compute node where Moab will store the data. You will supply the remote data source location at job submission time, but you must use the `$JOBHOST` placeholder for the name of the compute node. After the job runs, you can also copy data from the local file system to a remote file system.

*Image 24-5: Data Staging To or From a Local Compute Node*



Before staging data to or from a local compute node, follow the procedure in [Configuring Data Staging](#).

### To Stage Data to or from a Local Compute Node

1. If you have not already done so, configure your SSH keys and `moab.cfg` to support data staging. See [Configuring the SSH Keys for the Data Staging Transfer Script](#) and [Configuring Data Staging](#) for more information.
2. Create your job templates for data staging jobs in `moab.cfg`. The templates in the example below create a compute job that stages data in before it starts and stages data out when it completes. For more information about creating job templates, see [About Job Templates](#).
  - a. Create a selectable master template, called `ds` in the example below, that creates a stage in and stage out system job. This name should match the `DEFAULT_TEMPLATE` value in `ds_config.py`. For more information, see [Configuring Data](#)

### Staging with Advanced Options.

- b. For the data staging in job template, called `dsin` in the example below, specify that it will create a data staging job by setting `DATASTAGINGJOB` to `TRUE`. Note that the name of this job template must match the name of the data stage in job template referenced in the master template.
- c. Set the staging job template bandwidth `GRES` to the amount of bandwidth a single stage in job should use. This indicates how many of the bandwidth units specified with `NODECFG[GLOBAL]` in [Configuring Data Staging](#) a data staging job with this template should consume.
- d. For local node data staging it is important that the data staging job has the entire node to itself. To prevent Moab from scheduling another job on the node at the same time as the data staging job, set the `NODEACCESSPOLICY` to `SINGLEJOB` in the staging job template.
- e. Add `INHERITRES=TRUE` to reserve the compute node for the data staging job to prevent other compute jobs from using the node at the same time and creating input, output, and disk conflicts with the data staging job.
- f. Create a trigger that executes the `ds_move_scp`, `ds_move_rsync`, or `ds_move_multiplex` script, depending on which file transfer utility you use. Set the `attacherror`, `objectxmlstdin`, and `user` `FLAGS` to attach any trigger `stderr` as a message to the job, pass the job XML to the script, and indicate that the script should run as the job's user, respectively.

**i** If you use the `rsync` protocol, you can configure your data staging jobs to report the actual number of bytes transferred and the total data size to be transferred. To do so, use the `Sets` attribute to `^BYTES_IN.^DATA_SIZE_IN` for stage in jobs and `^BYTES_OUT.^DATA_SIZE_OUT` for stage out jobs. For example, a stage in trigger would look like the following:

```

JOBCFG[dsin]
TRIGGER=EType=start,AType=exec,Action="/opt/moab/tools/data-staging/ds_
move_rsync --stagein",Flags=objectxmlstdin:user:attacherror,Sets=^BYTES_
IN.^DATA_SIZE_IN

```

A stage out trigger would look like the following:

```

JOBCFG[dsout]
TRIGGER=EType=start,AType=exec,Action="/opt/moab/tools/data-staging/ds_
move_rsync --stageout",Flags=objectxmlstdin:user:attacherror,Sets=^BYTES_
OUT.^DATA_SIZE_OUT

```

These variables show up as events if you set your `WIKIEVENTS` parameter to `TRUE`.

- g. Create the stage out job, called `dsout` in the example below, by repeating steps 2b - 2f in a new template. In the example below, this template is called `dsout`. Note that the name of this job template must match the name of the data stage out job template referenced in the data staging master template.

```

JOBCFG[ds]      TEMPLATEDEPEND=AFTEROK:dsin  TEMPLATEDEPEND=BEFORE:dsout
SELECT=TRUE

JOBCFG[dsin]   DATASTAGINGSYSJOB=TRUE
JOBCFG[dsin]   GRES=bandwidth:2
JOBCFG[dsin]   NODEACCESSPOLICY=SINGLEJOB
JOBCFG[dsin]   INHERITRES=TRUE
JOBCFG[dsin]   TRIGGER=EType=start,AType=exec,Action="/opt/moab/tools/data-
staging/ds_move_rsync --stagein",Flags=attacherror:objectxmlstdin:user

JOBCFG[dsout]  DATASTAGINGSYSJOB=TRUE
JOBCFG[dsout]  GRES=bandwidth:1
JOBCFG[dsout]  NODEACCESSPOLICY=SINGLEJOB
JOBCFG[dsout]  INHERITRES=TRUE
JOBCFG[dsout]  TRIGGER=EType=start,AType=exec,Action="/opt/moab/tools/data-
staging/ds_move_rsync --stageout",Flags=attacherror:objectxmlstdin:user

```

3. Create the job using `msub`, adding resources and specifying a script as you normally would. Then configure Moab to stage the data for it. To do so:

- a. If the compute job does not use all of the node's processors, Moab could schedule another job on the node at the same time. If you did not set `NODEACCESSPOLICY` to `SINGLEJOB` in your `moab.cfg`, set the policy for this job by adding `-l naccesspolicy=singlejob` to your `msub` command.

```
> msub -l naccesspolicy=singlejob... <jobScript>
```

- b. At the end of the command, use the `--stagein/--stageout` option and/or `--stageinfile/--stageoutfile` option.

- The `--stagein/--stageout` option lets you specify a single file or directory to stage in or out. You must set the option equal to `<source>%<destination>`, where `<source>` and `<destination>` are both `[<user>@]<host>:/<path>/[<fileName>]`. See [Staging a File or Directory](#) for format and details.

**i** If the destination partition is down or does not have configured resources, the data staging workflow submission will fail.

If you do not know the host where the job will run but want the data staged to the same location, you can use the `$JOBHOST` variable in place of a host.

```
> msub --stagein=annasmith@labs:/patient-
022678/%$JOBHOST:/davidharris/research/patientrecords <jobScript>
```

*Moab copies the /patient-022678 directory from the hospital's labs server to the node where the job will run prior to job start.*

- The `--stageinfile/--stageoutfile` option lets you specify a file that contains the file and directory name(s) to stage in or out. You must set the option equal to `<path>/<fileName>%` of the file. The file must contain at least one line with this format: `<source>%<destination>`, where `<source>` and `<destination>` are both `[<user>@]<host>:/<path>[/<fileName>]`. See [Staging Multiple Files or Directories](#) for more information.

**i** If the destination partition is down or does not have configured resources, the data staging workflow submission will fail.

```
> msub --stageinfile=/davidharris/research/recordlist <jobScript>
```

*Moab copies all files specified in the /davidharris/research/recordlist file to the host where the job will run prior to job start.*

```
/davidharris/research/recordlist:
```

```
annasmith@labs:/patient-022678/tests/blood02282014%$JOBHOST:/davidharris/research/patientrecords/blood02282014
annasmith@labs:/patient-022678/visits/stats02032014%$JOBHOST:/davidharris/research/patientrecords/stats02032014
annasmith@labs:/patient-022678/visits/stats02142014%$JOBHOST:/davidharris/research/patientrecords/stats02142014
annasmith@labs:/patient-022678/visits/stats02282014%$JOBHOST:/davidharris/research/patientrecords/stats02282014
annasmith@labs:/patient-022678/visits/stats03032014%$JOBHOST:/davidharris/research/patientrecords/stats03032014
annasmith@labs:/patient-022678/visits/stats03142014%$JOBHOST:/davidharris/research/patientrecords/stats03142014
annasmith@labs:/patient-022678/visits/stats03282014%$JOBHOST:/davidharris/research/patientrecords/stats03282014
```

*Moab copies the seven patient record files from the hospital's labs server to the host where the job will run prior to job start.*

- c. The `--stageinsize/--stageoutsize` option lets you specify the estimated size of the files and/or directories to help Moab more quickly and accurately calculate the amount of time it will take to stage the data and therefore schedule your job correctly. If you used the `$JOBHOST` variable to stage in, then setting `--stageinsize` is required. `--stageoutsize` is always required for staging data out. If you provide an integer, Moab will assume the number is in megabytes. To change the unit, add another suffix. See [Stage in or out File Size](#) for more information.

```
> msub --stageinfile=/davidharris/research/recordlist --stageinsize=100
<jobScript>
```

*Moab copies the /davidharris/research/recordlist file, which is approximately 100 megabytes, from the biology node to the host where the job will run prior to job start.*

4. To see the status, errors, and other details associated with your data staging job, run `checkjob -v`. See [checkjob](#) for details.

**i** Your `checkjob` output might include a warning that says "req 1 RM (internal) does not match job destination RM". You can safely ignore this message.

## Related Topics

- [24.1 Data Staging Example - page 987](#)
- [24.2.2 Configuring Data Staging - page 992](#)
- [24.2.6 Configuring Data Staging with Advanced Options - page 1009](#)
- [24.2.5 Staging Data To or From a Compute Node - page 1005](#)
- 

## 24.2.6 Configuring Data Staging with Advanced Options

In this topic:

- [24.2.6.A Using a Different Default Template Name - page 1009](#)
- [24.2.6.B Supporting Multiple File Transfer Script Utilities in a Grid on a Per-Partition Basis - page 1010](#)
- [24.2.6.C Receiving Notification at the Completion of the Data Staging Job - page 1011](#)
- [24.2.6.D Adding a Non-Default Template via `msub` - page 1011](#)
- [24.2.6.E Using `msub` to Return all the Job IDs in the Workflow at Submission Time - page 1012](#)

### 24.2.6.A Using a Different Default Template Name

When you submit a data staging job, a data staging job template is attached to the job automatically. In the reference script configuration, the default template name is `ds`. This is

the template that will be attached to the compute job by the client *msub* filter.

If you would like to change the name of the default template that is automatically attached, you should change the value of `DEFAULT_TEMPLATE` in the `ds_config.py` file installed on all client submit hosts. This name must match the master data staging template name specified in the Moab configuration file.

## To Configure the `DEFAULT_TEMPLATE` Variable

1. Open the `ds_config.py` file for modification. It is located in `/opt/moab/tools/data-staging/` by default.

```
[moab]$ vi /opt/moab/tools/data-staging/ds_config.py
```

2. Locate the `DEFAULT_TEMPLATE` parameter:

```
...
DEFAULT_TEMPLATE = "ds"
...
```

3. Replace the template name with the one specified in the Moab configuration file:

```
ds_config.py
...
DEFAULT_TEMPLATE = "datastaging"
...

moab.cfg
...
JOBCFG[datastaging] TEMPLATEDEPEND=...
```

4. Make these changes on all client submit hosts.

### 24.2.6.B Supporting Multiple File Transfer Script Utilities in a Grid on a Per-Partition Basis

If you want a different transfer script to run based on which partition the job is submitted to, you can configure a multiplexer script that will switch execution to various other scripts based on the partition.

## To Support Multiple File Transfer Script Utilities in a Grid on a Per-Partition Basis

1. Configure the trigger in your job templates in `moab.cfg` to run `ds_move_multiplex` instead of `ds_move_rsync` or `ds_move_scp`.
2. Configure the `PARTITION_TO_SCRIPT` variable in `ds_config.py` to provide a mapping from each partition to the desired script to run.
  - a. Open the `ds_config.py` file for modification. It is located in `/opt/moab/tools/data-staging/` by default.

```
[moab]$ vi /opt/moab/tools/data-staging/ds_config.py
```

- b. Locate the `PARTITION_TO_SCRIPT` parameter:

```
...
PARTITION_TO_SCRIPT =
{"partition_1_name":"/opt/moab/tools/data-staging/ds_move_rsync",
 "partition_2_name":"/opt/moab/tools/data-staging/ds_move_scp",
 "partition_3_name":"/opt/moab/tools/data-staging/ds_move_rsync"}
...
```

- c. Replace the `partition_*_names` with partitions that exist in your configuration. After each partition, specify the script that you want to execute for that partition.

### 24.2.6.C Receiving Notification at the Completion of the Data Staging Job

If you want explicit notification in case of failure of the stage out job, add an additional trigger to the `dsout` job template, which will send email notification to the job's submitter. For more information, see [Using a Trigger to Send Email](#).

```
JOBCFG[dsout]  DATASTAGINGSYSJOB=TRUE
JOBCFG[dsout]  GRES=bandwidth:1
JOBCFG[dsout]  FLAGS=GRESONLY
JOBCFG[dsout]  TRIGGER=EType=start,AType=exec,Action="/opt/moab/tools/data-staging/ds_
move_rsync --stageout",Flags=attacherror:objectxmlstdin:user
JOBCFG[dsout]  TRIGGER=EType=fail,AType=mail,Action="Your (stageout) data staging job
$OID failed."
```

The first trigger listed in the template configuration should be the `exec` trigger. Add the email trigger and any other triggers after the `exec` trigger. You can modify the email trigger to run at completion rather than at failure. You can also add this type of trigger to stage in jobs.

### 24.2.6.D Adding a Non-Default Template via `msub`

You can have multiple data staging template workflows defined in the `moab.cfg`. The submit filter is configured to add only one of them by default. If you want to use one of the other available templates, you can do so by using the `-l template=TEMPLATENAME` option in the `msub` command.

Given the following `moab.cfg`:

```
#Default data staging template:

JOBCFG[ds]      TEMPLATEDEPEND=AFTEROK:dsin TEMPLATEDEPEND=BEFORE:dsout SELECT=TRUE
JOBCFG[dsin]    DATASTAGINGSYSJOB=TRUE
JOBCFG[dsin]    GRES=bandwidth:2
JOBCFG[dsin]    FLAGS=GRESONLY
JOBCFG[dsin]    TRIGGER=EType=start,AType=exec,Action="/opt/moab/tools/data-staging/ds_
move_rsync --stagein",Flags=attacherror:objectxmlstdin:user

JOBCFG[dsout]   DATASTAGINGSYSJOB=TRUE
JOBCFG[dsout]   GRES=bandwidth:1
```

```

JOBCFG[dsout]   FLAGS=GRESONLY
JOBCFG[dsout]   TRIGGER=EType=start,AType=exec,Action="/opt/moab/tools/data-staging/ds_
move_rsync --stageout",Flags=attacherror:objectxmlstdin:user

#experimental data staging template:

JOBCFG[dscustom]   TEMPLATEDEPEND=AFTEROK:dscustomin
TEMPLATEDEPEND=BEFORE:dscustomout SELECT=TRUE
JOBCFG[dscustomin]   DATASTAGINGSYSJOB=TRUE
JOBCFG[dscustomin]   GRES=bandwidth:2
JOBCFG[dscustomin]   FLAGS=GRESONLY
JOBCFG[dscustomin]   TRIGGER=EType=start,AType=exec,Action="/opt/moab/tools/data-
staging/ds_move_custom --stagein",Flags=attacherror:objectxmlstdin:user

JOBCFG[dscustomout]   DATASTAGINGSYSJOB=TRUE
JOBCFG[dscustomout]   GRES=bandwidth:1
JOBCFG[dscustomout]   FLAGS=GRESONLY
JOBCFG[dscustomout]   TRIGGER=EType=start,AType=exec,Action="/opt/moab/tools/data-
staging/ds_move_custom --stageout",Flags=attacherror:objectxmlstdin:user

```

The user could submit a job using the custom data staging template with the following command:

```
[moab]$ msub -l template=dscustom ...
```

### 24.2.6.E Using msub to Return all the Job IDs in the Workflow at Submission Time

By default, *msub* will print the job ID to stdout at the time of submission. If you would like to have *msub* print all of the jobs that are created as part of the data staging workflow template, you can use the *msub --workflowjobids* option to show all the job IDs at submission time:

```

$ echo sleep 60 | msub -l walltime=15 --workflowjobids
MoabA.3.dsin MoabA.3 MoabA.3.dsout

```

This could be useful if you are writing scripts to do your own workflows and you need to programmatically capture the data stage out job name for use in your workflow.

#### Related Topics

- [24.2.2 Configuring Data Staging - page 992](#)

## 24.3 Data Staging References

### 24.3.1 Sample User Job Script

The code below is an example of a job script that a user might use to run a data staging job:



```
#!/bin/bash
#
# Sample data staging job script
#
# stage in directives
#MSUB --stageinsize=1MB
#MSUB --stagein=davidharris@source-server:/tmp/filein.tostage%davidharris@destination-
server:/tmp/filein.staged
#
# stage out directives
#MSUB --stageoutsize=10GB
#MSUB --stageout=davidharris@destination-
server:/tmp/fileout.tostage%davidharris@source-server:/tmp/fileout.staged

# run executable on the destination host using staged data
$HOME/bin/my_compute_executable < /tmp/filein.staged > /tmp/fileout.tostage
```

---

## Related Topics

- [24.1 Data Staging Example - page 987](#)
- [Staging Data - page 255](#)

## Chapter 25: Using NUMA with Moab

Moab works with Torque to support these two Non-Uniform Memory Architecture (NUMA) systems:

- **NUMA-Aware** – Introduced with Torque version 6.0, this configuration supports multi-req jobs and jobs that span hosts. Moab version 9.0 and later is also required.
- **NUMA-Support** – Introduced with Torque version 3.0, this configuration supports only a single instance for pbs\_mom that is treated as if there were multiple nodes running in the cluster. This configuration is only for large-scale SLES systems using SGI Altix and UV hardware.

This chapter serves as a central information repository for the various configuration settings involved when using either NUMA system configuration.



Both systems cannot be configured for use at the same time.

In this chapter:

25.1 Using NUMA-Aware with Moab .....	1014
25.1.1 NUMA Process .....	1015
25.1.2 Installation and Configuration .....	1016
25.1.3 Moab and NUMA Resources .....	1016
25.1.4 Track Dedicated NUMA Resources .....	1017
25.2 Using NUMA-Support with Moab .....	1018

### 25.1 Using NUMA-Aware with Moab

This chapter introduces NUMA-aware scheduling and identifies Moab-, Torque-, and MAM-related functions.



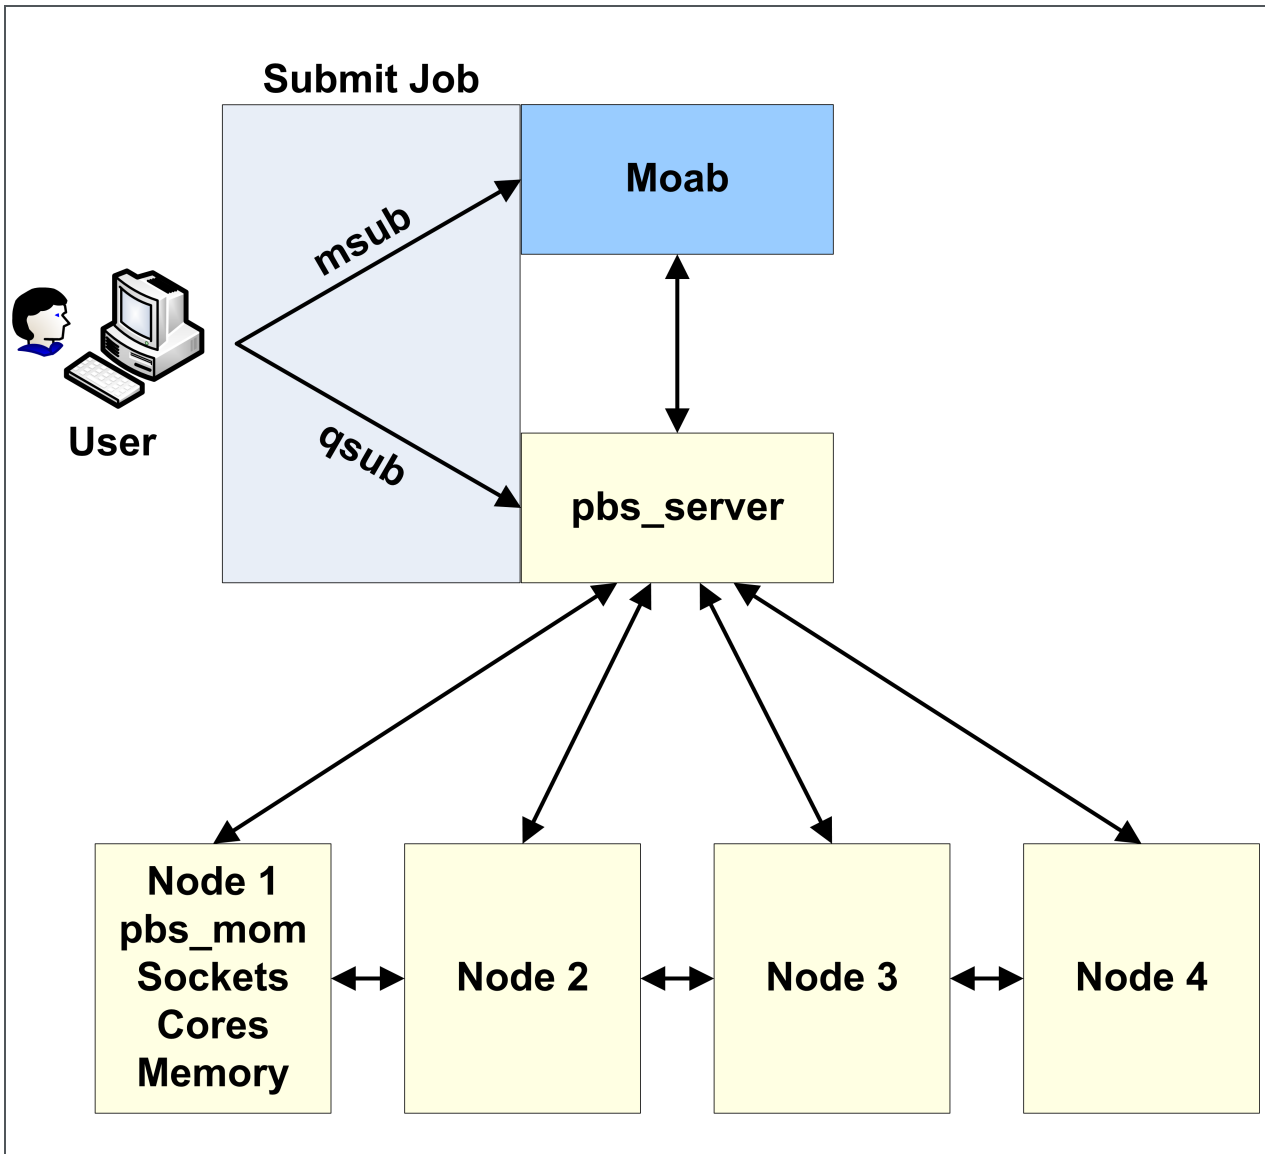
NUMA-aware scheduling is available with Moab Workload Manager 9.0 and later and requires Torque Resource Manager 6.0 and later.

In this topic:

- [25.1.1 NUMA Process - page 1015](#)
- [25.1.2 Installation and Configuration - page 1016](#)
- [25.1.3 Moab and NUMA Resources - page 1016](#)
- [25.1.4 Track Dedicated NUMA Resources - page 1017](#)

## 25.1.1 NUMA Process

The following image provides an example of the NUMA process:



## 25.1.2 Installation and Configuration

Moab does not require any special installation or configuration processes to support NUMA-aware with Torque.

## 25.1.3 Moab and NUMA Resources

Moab uses generic resources to natively understand the concept of 'socket', 'numanode', 'core' and 'thread'. The `msub/qsub -L` syntax for job submissions lets you request placement or allocation of these specific resources.

See '`-L NUMA Resource Request`' in the *Torque Resource Manager Administrator Guide* for more information on using the `-L` syntax.

**i** For the `-L` syntax, submit using `qsub` unless your system requires submissions using `msub`.

- A job requesting a placement of a numanode is requesting exclusive access to the entire numa node and all of its resources including cores, threads, memory, gpus, and mics.
- A job requesting a placement of a socket is requesting exclusive access to the entire socket including numanodes, cores, threads, memory, gpus and mics.

**i** Moab does not require a configuration change to support NUMA-aware scheduling (no new Moab configuration parameters). However, you might need to increase `MAXGRES` to accommodate the additional resources. See [MAXGRES](#).

When using NUMA-aware, the following occurs:

1. Moab imports NUMA resources from Torque and treats them as a special case of generic resources.
2. When a job requests a NUMA resource, such as a socket or numanode, Moab will schedule exclusive access to those resources for the job. If exclusive access to the NUMA resource is not available on a particular node, Moab will look for another node or schedule the job out into the future if no resources are available at the time.

### 25.1.3.A Balanced Resources

Moab assumes that the NUMA resources on a particular node are balanced. This means that each socket has the same amount of resources, including numanodes, cores, threads, memory, and gpus.

Moab supports advanced multi-req resource requests within the same job using the `msub/qsub '-L'` syntax.

In addition, any job that requests NUMA resources will receive a per-task default memory requirement if a memory requirement is *not* specified by the user. .

### 25.1.3.B Job Recommendations

Adaptive Computing provides these recommendations for jobs:

- Jobs that request NUMA resources can share nodes with non-NUMA jobs, but it is *not* recommended. It is recommended that you enforce the separation using policies (queues, reservations, partitions, node-sets, etc.).
- GPU jobs that request cores and processors should *not* share nodes with non-GPU jobs. For example, a job requesting a numanode can be given a numanode with a GPU and therefore block other jobs from consuming that GPU. This can be prevented with proper queue policies and placement.
- A job can run across heterogeneous NUMA resources. For example, a socket on one node can contain more memory than a socket on another node. Use the `NODEALLOCATIONPOLICY` Moab parameter to enable running a job across homogenous NUMA resources.

### 25.1.4 Track Dedicated NUMA Resources

If Moab Accounting Manager is part of your configuration, you can configure MAM to track dedicated NUMA resources ( sockets, numanodes, cores, threads).

As the MAM Admin, run the commands for the individual resources you want to track. The follow example shows the commands for all of the available resources.

```
[mam]$ mam-shell Attribute Create Object=UsageRecord Name=Sockets DataType=Integer
Description="\Number of Sockets Dedicated\"
[mam]$ mam-shell Attribute Create Object=UsageRecord Name=NumaNodes DataType=Integer
Description="\Number of Numa Nodes Dedicated\"
[mam]$ mam-shell Attribute Create Object=UsageRecord Name=Cores DataType=Integer
Description="\Number of Cores Dedicated\"
[mam]$ mam-shell Attribute Create Object=UsageRecord Name=Threads DataType=Integer
Description="\Number of Threads Dedicated\"
```

**i** NUMA resources are only reported to MAM when they are dedicated to the job. As you can specify the placement rules (the NUMA resources that are dedicated), it is *not* recommended to charge for any NUMA resources; use Processors instead.

## 25.2 Using NUMA-Support with Moab

This topic serves as a central information repository for NUMA-support systems. This topic provides basic information and contains links to the various NUMA-aware topics found throughout the documentation.

**i** Support for NUMA-support systems is available only on large-scale SLES systems using SGI Altix and UV hardware and requires Torque 3.0 or later.

### Installation and Configuration

Additional information is provided on configuring Moab for NUMA-support. See [Hardware Integration](#).

## Appendix A: Moab Parameters

See [Initial Moab Configuration](#) for further information about specifying parameters.

**i** If a parameter does not have a set default, the Default value in the table is shown as '--'.

Index: [A](#) [B](#) [C](#) [D](#) [E](#) [F](#) [G](#) [H](#) [I](#) [J](#) [K](#) [L](#) [M](#) [N](#) [O](#) [P](#) [Q](#) [R](#) [S](#) [T](#) [U](#) [V](#) [W](#) [X](#)


### ACCOUNTCFG[<ACCOUNTID>]

<b>Format</b>	List of zero or more space delimited <ATTR>=<VALUE> pairs, where <ATTR> is one of the following: <a href="#">General Credential Flags</a> , CHARGERATE, <a href="#">PRIORITY</a> , <a href="#">ENABLEPROFILING</a> , <a href="#">MEMBERULIST</a> , <a href="#">PLIST</a> , <a href="#">QDEF</a> , <a href="#">QLIST</a> , <a href="#">usage limit</a> , or a <a href="#">fairness usage limit specification</a> ( <a href="#">FSCAP</a> , <a href="#">FSTARGET</a> , and <a href="#">FSWEIGHT</a> ).
<b>Default</b>	---
<b>Description</b>	Specifies account specific attributes. See the <a href="#">account</a> overview for general information and the <a href="#">job flag overview</a> for a description of legal flag values.
<b>Example</b>	<pre>ACCOUNTCFG[projectX] MAXJOB=50 QDEF=highprio</pre> <p><i>Up to 50 jobs submitted under the account ID <code>projectX</code> will be allowed to execute simultaneously and will be assigned the QOS <code>highprio</code> by default.</i></p>

### ACCOUNTINGINTERFACEURL

<b>Format</b>	<URL> where protocol can be one of <code>exec</code> or <code>file</code>
<b>Default</b>	---
<b>Description</b>	The interface to use for real-time export of Moab accounting/auditing information. See <a href="#">Exporting Events in Real-Time</a> for more information.
<b>Example</b>	<pre>ACCOUNTINGINTERFACEURL exec://\$TOOLSDIR/dumpacc.pl</pre>

ACCOUNTWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	1
<b>Description</b>	The priority weight to be applied to the specified account priority. See <a href="#">Credential (CRED) Factor</a> .
<b>Example</b>	<pre>ACCOUNTWEIGHT 100</pre>

ADMIN1, ADMIN2, ADMIN3	
<b>Description</b>	 These parameters are deprecated. Use <a href="#">ADMINCFG</a> .



ADMINCFG[X]	
<b>Format</b>	One or more <ATTR>=<VALUE> pairs, where <ATTR> is one of the following: ENABLEPROXY, USERS, GROUPS, <a href="#">SERVICES</a> , or NAME
<b>Default</b>	---
<b>Description</b>	Allows a site to configure which services and users belong to a particular level of administration. <b>Note:</b> The first user listed in the ADMINCFG [ 1 ] users list is considered to be the primary admin. The option USERS=ALL is allowed. The groups list adds the groups' users as if they were listed individually as USERS. To prevent Moab from assigning a primary user from the first group listed, you must specify a primary user first using the USERS attribute, then list the desired groups.
<b>Example</b>	<pre>ADMINCFG[1] USERS=root, john ADMINCFG[1] GROUPS=admin ADMINCFG[1] SERVICES=ALL ADMINCFG[1] NAME=batchadmin ADMINCFG[3] USERS=bob, carol, smoore ADMINCFG[3] GROUPS=science, math ADMINCFG[3] SERVICES=mjobctl, mcredctl, runjob ADMINCFG[3] NAME=helpdesk</pre> <p><i>Members of the batchadmin admin role and members of the admin group are allowed to run all commands. Members of the helpdesk role and science and math groups are allowed to run mjobctl. They are also able to view and modify credential objects (i.e., users, groups, accounts, etc.). See the <a href="#">security overview</a> for more details.</i></p> <pre>ADMINCFG[4] USERS=ALL SERVICES=checknode</pre> <p><i>All users can execute <a href="#">checknode</a> to get information on any node.</i></p>

AGGREGATENODEACTIONS	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	<p>Consolidates queued node actions into as few actions as possible to reduce communication burden with resource manager. Node actions are queued until the <a href="#">AGGREGATENODEACTIONSTIME</a> setting.</p> <div style="border: 1px solid #0070C0; border-radius: 5px; padding: 5px; margin-top: 10px;"> <p><b>i</b> This might delay some node actions. When reprovisioning, the system job might expire before the provision action occurs; while the action will still occur, the job will not show it.</p> </div>
<b>Example</b>	<pre style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 5px;">AGGREGATENODEACTIONS TRUE</pre> <p><i>Queues node actions together when possible.</i></p>

AGGREGATENODEACTIONSTIME	
<b>Format</b>	<SECONDS>
<b>Default</b>	60
<b>Description</b>	The delay time for the <a href="#">AGGREGATENODEACTIONS</a> parameter to aggregate requests before sending job batches.
<b>Example</b>	<pre style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 5px;">AGGREGATENODEACTIONSTIME 120</pre> <p><i>Sets the AGGREGATENODEACTIONS delay to two minutes.</i></p>

ALLOWMULTIREQNODEUSE	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	By default, Moab does not allow different requirements on the same job to occupy the same node. For example, if a job is submitted with <code>nodes=2:ppn=8+4:fast:ppn=16</code> , it's possible that some of the tasks requested could overlap onto the same node. This parameter instructs Moab to allow overlapping the same node, or not. This parameter also applies to the various <code>-w</code> clauses of an <code>mshow -a</code> command.
<b>Example</b>	<pre>ALLOWMULTIREQNODEUSE TRUE</pre>

ALLOWROOTJOBS	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	Specifies whether batch jobs from the root user (UID=0) are allowed to be executed. <b>Note:</b> The resource manager must also support root jobs.
<b>Example</b>	<pre>ALLOWROOTJOBS TRUE</pre> <i>Jobs from the root user can execute.</i>

ALWAYSEVALUATEALLJOBS	
<b>Format</b>	ALWAYS, FIRSTRSV, or FULLRSVV
<b>Default</b>	FIRSTRSV
<b>Description</b>	<p>Instructs Moab how to handle the scheduling of eligible jobs during the first phase of each scheduling iteration:</p> <ul style="list-style-type: none"> <li>• FIRSTRSV directs Moab to stop considering eligible jobs once a single reservation has been created.</li> <li>• FULLRSV tells Moab to evaluate eligible jobs until reservations have been created for a number of eligible jobs.</li> <li>• ALWAYS directs Moab to always evaluate all eligible jobs.</li> </ul>
<b>Example</b>	<pre>ALWAYSEVALUATEALLJOBS FIRSTRSV</pre>

AMCFG	
<b>Format</b>	One or more key-value pairs as described in <a href="#">AMCFG Parameters and Flags</a> .
<b>Default</b>	---
<b>Description</b>	The interface and policy configuration for the scheduler-accounting manager interface.
<b>Example</b>	<pre>AMCFG [mam] TYPE=MAM STARTFAILUREACTION=HOLD</pre>

APPLICATIONLIST	
<b>Format</b>	Space-delimited list of generic resources.
<b>Default</b>	---
<b>Description</b>	Specifies which generic resources represent actual applications on the cluster/grid. See <a href="#">Managing Consumable Generic Resources</a> for more information.
<b>Example</b>	<pre>NODECFG[node01] GRES=calclab:1,powerhouse:1 RCSOFTWARE=calclab:1,powerhouse:1 NODECFG[node02] GRES=calclab:1,powerhouse:1 RCSOFTWARE=calclab:1,powerhouse:1 APPLICATIONLIST calclab,powerhouse</pre> <p><i>The generic resources <code>calclab</code> and <code>powerhouse</code> will now be recognized and treated as application software.</i></p>

ARRAYJOBPARLOCK	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	If set to <code>TRUE</code> , all subjobs of an array are locked to a single partition. The default behavior when scheduling array subjobs is to span the jobs across partitions when possible. The <a href="#">ARRAYJOBPARLOCK</a> job flag can be used to specify partition locking at submit time. The <a href="#">ARRAYJOBPARSPAN</a> job flag overrides this parameter.
<b>Example</b>	<pre>ARRAYJOBPARLOCK TRUE</pre>

ATTRATTRWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	The priority weight to be applied to jobs with the specified job attribute. See <a href="#">Attribute (ATTR) Factor</a> .
<b>Example</b>	<pre>ATTRATTRWEIGHT 100</pre>

ATTRGRESWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	The priority weight to be applied to jobs requesting the specified <a href="#">generic resource</a> . See <a href="#">Attribute (ATTR) Factor</a> .
<b>Example</b>	<pre>ATTRGRESWEIGHT 200</pre>

ATTRSTATEWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	The priority weight to be applied to jobs with the specified job state. See <a href="#">Attribute (ATTR) Factor</a> .
<b>Example</b>	<pre>ATTRSTATEWEIGHT 200</pre>

ATTRWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	1
<b>Description</b>	The priority component weight to be applied to the ATTR subcomponents. See <a href="#">Attribute (ATTR) Factor</a> .
<b>Example</b>	<pre>ATTRWEIGHT 2 ATTRSTATEWEIGHT 200</pre>

BACKFILLDEPTH	
<b>Format</b>	<INTEGER>
<b>Default</b>	0 (no limit)
<b>Description</b>	The number of idle jobs to evaluate for backfill. The backfill algorithm will evaluate the top <X> priority jobs for scheduling. By default, all jobs are evaluated.
<b>Example</b>	<pre>BACKFILLDEPTH 128</pre> <p><i>Evaluate only the top 128 highest priority idle jobs for consideration for backfill.</i></p>

BACKFILLMETRIC	
<b>Format</b>	One of the following: PROCS, PROCSECONDS, SECONDS, or NODES
<b>Default</b>	PROCS
<b>Description</b>	The criteria used by the backfill algorithm to determine the 'best' jobs to backfill. Only applicable when using the BESTFIT backfill algorithm.
<b>Example</b>	<pre>BACKFILLMETRIC PROCSECONDS</pre>

BACKFILLPOLICY	
<b>Format</b>	One of FIRSTFIT, BESTFIT, or NONE
<b>Default</b>	FIRSTFIT
<b>Description</b>	Specifies which backfill algorithm will be used. See <a href="#">Configuring Backfill</a> for more information.
<b>Example</b>	<pre>BACKFILLPOLICY NONE</pre>

BFCHUNKDURATION	
<b>Format</b>	[ [ [DD:] HH:] MM:] SS
<b>Default</b>	0 (chunking disabled)
<b>Description</b>	The duration during which freed resources will be aggregated for use by larger jobs. Used in conjunction with <a href="#">BFCHUNKSIZE</a> . See <a href="#">Configuring Backfill</a> for more information.
<b>Example</b>	<pre>BFCHUNKDURATION 00:05:00 BFCHUNKSIZE 4</pre> <p><i>Aggregate backfillable resources for up to 5 minutes, making resources available only to jobs of size 4 or larger.</i></p>

BFCHUNKSIZE	
<b>Format</b>	<INTEGER>
<b>Default</b>	0 (chunking disabled)
<b>Description</b>	The minimum job size that can utilize chunked resources. Used in conjunction with <a href="#">BFCHUNKDURATION</a> . See <a href="#">Configuring Backfill</a> for more information.
<b>Example</b>	<pre>BFCHUNKDURATION 00:05:00 BFCHUNKSIZE 4</pre> <p><i>Aggregate backfillable resources for up to 5 minutes, making resources available only to jobs of size 4 or larger.</i></p>



BFMINVIRTUALWALLTIME	
<b>Format</b>	[ [ [DD:] HH:] MM:] SS
<b>Default</b>	---
<b>Description</b>	The minimum job wallclock time for virtual scaling (optimistic-like backfilling). Any job with a wallclock time less than this setting will <i>not</i> be virtually scaled. The value specified relates to a job's original walltime and not its virtually-scaled walltime.
<b>Example</b>	<code>BFMINVIRTUALWALLTIME 00:01:30</code>

BFPRIORITYPOLICY	
<b>Format</b>	One of RANDOM, DURATION, or HWDURATION
<b>Default</b>	---
<b>Description</b>	Specifies policy to use when prioritizing backfill jobs for <a href="#">preemption</a> .
<b>Example</b>	<code>BFPRIORITYPOLICY DURATION</code> <i>Use length of job in determining which backfill job to preempt.</i>

BFVIRTUALWALLTIMECONFLICTPOLICY	
<b>Format</b>	One of the following: PREEMPT
<b>Default</b>	---
<b>Description</b>	Specifies how to handle scheduling conflicts when a virtually scaled job 'expands' to its original wallclock time. This occurs when the job is within one scheduling iteration - <a href="#">RMPOLLINTERVAL</a> - of its virtually scaled wallclock time expiring.
<b>Example</b>	<code>BFVIRTUALWALLTIMECONFLICTPOLICY PREEMPT</code>

BFVIRTUALWALLTIMESCALINGFACTOR	
<b>Format</b>	<DOUBLE>
<b>Default</b>	0 (virtual scaling disabled)
<b>Description</b>	<p>The factor by which eligible jobs' wallclock time is virtually scaled (optimistic-like backfilling).</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> If you do not want scaling, set BFVIRTUALWALLTIMESCALINGFACTOR to '0' (default). Setting to '1' is not recommended as it impacts performance. When set to '1', Moab will exercise the code paths of scaling but no actual scaling will occur.</p> </div>
<b>Example</b>	<pre>BFVIRTUALWALLTIMESCALINGFACTOR .4</pre>

BYPASSCAP	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	The max weighted value allowed from the bypass count subfactor when determining a job's priority (see <a href="#">Priority Factors</a> for more information).
<b>Example</b>	<pre>BYPASSWEIGHT 5000 BYPASSCAP 30000</pre>

BYPASSWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	The weight to be applied to a job's backfill bypass count when determining a job's priority (see <a href="#">Priority Factors</a> for more information).
<b>Example</b>	<pre>BYPASSWEIGHT 5000</pre>

CHECKPOINTDIR	
<b>Format</b>	<STRING>
<b>Default</b>	---
<b>Description</b>	The directory for temporary job checkpoint files (usually of the form <code>jobid.cp</code> ). <b>Note:</b> This is <i>not</i> the directory for Moab's checkpoint file ( <code>.moab.ck</code> ).
<b>Example</b>	<code>CHECKPOINTDIR /tmp/moabcheckpoint</code>

CHECKPOINTEXPIRATIONTIME	
<b>Format</b>	[ [ [DD:] HH:] MM:] SS or UNLIMITED
<b>Default</b>	3,000,000 seconds
<b>Description</b>	Specifies how 'stale' checkpoint data can be before it is ignored and purged.
<b>Example</b>	<code>CHECKPOINTEXPIRATIONTIME 1:00:00:00</code> <i>Expire checkpoint data that has been stale for over 1 day.</i>

CHECKPOINTFILE	
<b>Format</b>	<STRING>
<b>Default</b>	<code>.moab.ck</code>
<b>Description</b>	Name (absolute or relative) of the Moab checkpoint file.
<b>Example</b>	<code>CHECKPOINTFILE /var/adm/moab/.moab.ck</code> <i>Maintain the Moab checkpoint file in the file specified.</i>

CHECKPOINTINTERVAL	
<b>Format</b>	[ [ [DD:] HH:] MM:] SS
<b>Default</b>	00:05:00
<b>Description</b>	<p>Time between automatic Moab checkpoints.</p> <div style="border: 1px solid #0070C0; padding: 5px; margin-top: 10px;"> <p><b>i</b> If <b>RMPOLLINTERVAL</b> does not specify both a minimum and maximum poll time, Moab will ignore <b>CHECKPOINTINTERVAL</b> and checkpoint every iteration.</p> </div>
<b>Example</b>	<pre>CHECKPOINTINTERVAL 00:15:00</pre> <p><i>Moab should checkpoint state information every 15 minutes.</i></p>

CHECKSUSPENDEDJOBPRIORITY	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	Prevents Moab from starting a job on any node containing a suspended job of higher priority.
<b>Example</b>	<pre>CHECKSUSPENDEDJOBPRIORITY FALSE</pre>

CHILDSTDERRCHECK	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	If set to <b>TRUE</b> , child processes Moab executes are considered failed if their standard error stream contains the text 'ERROR'.
<b>Example</b>	<pre>CHILDSTDERRCHECK TRUE</pre>

CLASSCFG[<CLASSID>]	
<b>Format</b>	List of zero or more space delimited <ATTR>=<VALUE> pairs, where <ATTR> is one of the following: General Credential Flags, DEFAULT.ATTR, DEFAULT.DISK, DEFAULT.FEATURES, DEFAULT.GRES, DEFAULT.MEM, DEFAULT.NODE, DEFAULT.NODESET, DEFAULT.PROC, ENABLEPROFILING, EXCL.FEATURES, EXCLUDEUSERLIST, HOSTLIST, IGHOSTLIST, JOBEPILOG, JOBPROLOG, JOBTRIGGER, MAXPROCPERNODE, MAX.NODE, MAX.PROC, MAX.TPN, MAX.WCLIMIT, MIN.NODE, MIN.PROC, MIN.TPN, MIN.WCLIMIT, PARTITION, PRIORITY, PRIORITYF, QDEF, QLIST, REQ.FEATURES, REQUIREDACCOUNTLIST, REQUIREDUSERLIST, RESFAILPOLICY, SYSPRIO, WCOVERRUN, usage limit, or fairshare usage limit specification.
<b>Default</b>	---
<b>Description</b>	Specifies class specific attributes (see <a href="#">Credential Overview</a> for details).
<b>Example</b>	<pre>CLASSCFG[batch] MAXJOB=50 QDEF=highprio</pre> <p><i>Up to 50 jobs submitted to the class batch will be allowed to execute simultaneously and will be assigned the QOS highprio by default.</i></p>
CLASSWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	1
<b>Description</b>	The weight to be applied to the class priority of each job (see <a href="#">Credential (CRED) Factor</a> and <a href="#">credential priority</a> ).
<b>Example</b>	<pre>CLASSWEIGHT 10</pre>

CLIENTCFG[<X>]	
<b>Format</b>	One or more of <ATTR>--<VALUE> pairs, where <X> indicates the specified peer and <ATTR> is one of the following: AUTH, AUTHCMD, AUTHTYPE, HOST, KEY, or DEFAULTSUBMITPARTITION.
<b>Default</b>	---
<b>Description</b>	The shared secret key and authentication method that Moab will use to communicate with the named peer daemon. See <a href="#">Security Overview</a> for more information. <b>Note:</b> The AUTHTYPE and KEY attributes of this parameter can only be specified in the moab-private.cfg config file.
<b>Example</b>	<pre>CLIENTCFG[silverB] KEY=apple7 AUTH=admin1</pre> <p><i>Moab will use the session key apple7 for peer authentication and for encrypting and decrypting messages sent from silverB. Also, client connections from this interface will be authorized at an admin 1 level.</i></p>

CLIENTCONNECTIONTIMEOUT	
<b>Format</b>	<SECONDS>
<b>Default</b>	30
<b>Description</b>	Specifies how long client commands will wait for the initial connection to succeed before giving up and failing.
<b>Example</b>	<pre>CLIENTCONNECTIONTIMEOUT 1</pre> <p><i>Client commands will wait only 1 second for the initial connection. If the client command has not connected within 1 second, it will give up and fail.</i></p>

CLIENTMAXCONNECTIONS	
<b>Format</b>	<INTEGER>
<b>Default</b>	128
<b>Description</b>	Changes the maximum number of connections that can simultaneously connect to Moab. The value can be increased during runtime, but it cannot be decreased. The value cannot be reduced below the default value of 128.
<b>Example</b>	<pre>CLIENTMAXCONNECTIONS 256</pre> <p><i>Doubles the maximum number of connections.</i></p>

CLIENTMAXPRIMARYRETRY	
<b>Format</b>	<INTEGER> or INFINITY
<b>Default</b>	1
<b>Description</b>	Specifies how many times the client command will attempt to retry its connection to the primary server if Moab is not available.
<b>Example</b>	<pre>CLIENTMAXPRIMARYRETRY 5 CLIENTMAXPRIMARYRETRYTIMEOUT 1000</pre> <p><i>The client command will attempt to retry its connection to the primary server 5 times with 1 second intervals before giving up. <b>Note:</b> If INFINITY is specified, Moab will attempt 2,140,000,000 times.</i></p>

CLIENTMAXPRIMARYRETRYTIMEOUT	
<b>Format</b>	<INTEGER> (milliseconds)
<b>Default</b>	2000
<b>Description</b>	Specifies how much time to wait until the client command will attempt to retry its connection to the primary server if Moab is not available.
<b>Example</b>	<pre>CLIENTMAXPRIMARYRETRY 3 CLIENTMAXPRIMARYRETRYTIMEOUT 500</pre> <p><i>The client command will attempt to retry its connection to the primary server 3 times with .5 second intervals before giving up.</i></p>

CLIENTTIMEOUT	
<b>Format</b>	[ [ [DD:]HH:]MM:]SS
<b>Default</b>	00:00:30
<b>Description</b>	Time that Moab client commands will wait for a response from the Moab server. See <a href="#">Client Configuration</a> for more information. <b>Note:</b> Can also be specified as an environment variable.
<b>Example</b>	<pre>CLIENTTIMEOUT 00:15:00</pre> <p><i>Moab clients will wait up to 15 minutes for a response from the server before timing out.</i></p>



CLIENTUIPORT	
<b>Format</b>	<INTEGER>
<b>Default</b>	N/A
<b>Description</b>	<p>Port on which to listen when <code>UIMANAGEMENTPOLICY FORK</code> is specified. This is typically Moab's configured listen port + 1.</p> <div style="border: 1px solid #0070C0; padding: 5px; margin-top: 10px;"> <p><b>i</b> Both <code>CLIENTUIPORT</code> and <code>UIMANAGEMENTPOLICY</code> need to be defined on clients for them to use the backup port when the primary Moab process is busy. <code>UIMANAGEMENTPOLICY</code> should be configured on the server, as well as any client machines.</p> </div>
<b>Example</b>	<pre>UIMANAGEMENTPOLICY FORK CLIENTUIPORT 42560</pre> <p><i>Moab is typically configured to listen on port 42559.</i></p>

CREDDISCOVERY	
<b>Format</b>	TRUE
<b>Default</b>	FALSE
<b>Description</b>	Specifies that Moab should create otherwise unknown credentials when it discovers them in the statistics files.
<b>Example</b>	<pre>CREDDISCOVERY TRUE</pre>

CREDWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	1
<b>Description</b>	The credential component weight associated with the <a href="#">credential priority</a> . See <a href="#">Credential (CRED) Factor</a> for more information.
<b>Example</b>	<pre>CREDWEIGHT 2</pre>

DATASTAGEHOLDTYPE	
<b>Format</b>	Any <a href="#">Job Hold</a> type
<b>Default</b>	DEFER
<b>Description</b>	Specifies what to do if a job's data staging operations fail.
<b>Example</b>	<code>DATASTAGEHOLDTYPE BATCH</code>

DEADLINEPOLICY	
<b>Format</b>	One of CANCEL, HOLD, IGNORE, or RETRY
<b>Default</b>	NONE
<b>Description</b>	Specifies what to do when a requested deadline cannot be reached (see <a href="#">Job Deadlines</a> ).
<b>Example</b>	<code>DEADLINEPOLICY IGNORE</code>

DEFAULTCLASSLIST	
<b>Format</b>	Space-delimited list of one or more <STRING>s.
<b>Default</b>	---
<b>Description</b>	The default classes supported on each node for RM systems that do not provide this information.
<b>Example</b>	<code>DEFAULTCLASSLIST serial parallel</code>

DEFAULTSTARTTIMEQUERY	
<b>Format</b>	Space-delimited list of one or more <STRING>s.
<b>Default</b>	---
<b>Description</b>	Specifies the default classes supported on each node for RM systems that do not provide this information.
<b>Example</b>	<pre>DEFAULTSTARTTIMEQUERY RESERVATION   Moab will estimate job start times using reservations.</pre>

DEFAULTSUBMITPARTITION	
<b>Format</b>	See parameter <a href="#">CLIENTCFG[]</a> .
<b>Default</b>	---
<b>Description</b>	If a user submits a job using <a href="#">msub</a> that does not specify host, feature, or partition constraints, then the <i>msub</i> client will insert the specified default submit partition into the newly submitted job as a hard requirement.
<b>Example</b>	<pre>CLIENTCFG[DEFAULT] DEFAULTSUBMITPARTITION=partition1</pre>

DEFERCOUNT	
<b>Format</b>	<INTEGER>
<b>Default</b>	24
<b>Description</b>	The number of times a job can be deferred before it will be placed in batch hold.
<b>Example</b>	<pre>DEFERCOUNT 12</pre>

DEFERSTARTCOUNT	
<b>Format</b>	<INTEGER>
<b>Default</b>	1
<b>Description</b>	The number of times a job will be allowed to fail in its start attempts before being deferred. <code>JOBRETRYTIME</code> overrides <code>DEFERSTARTCOUNT</code> ; <code>DEFERSTARTCOUNT</code> only begins when the <code>JOBRETRYTIME</code> window elapses. <b>Note:</b> A job's startcount will increase each time a start request is made to the resource manager regardless of whether or not this request succeeded. This means start count increases if job starts fail or if jobs are started and then rejected by the resource manager. (For related information, see <a href="#">Reservation Policies</a> , <a href="#">DEFERTIME</a> , <a href="#">RESERVATIONRETRYTIME</a> , <a href="#">NODEFAILURERESERVETIME</a> , <a href="#">JOBRETRYTIME</a> , and <a href="#">GUARANTEEDPREEMPTION</a> .)
<b>Example</b>	<code>DEFERSTARTCOUNT 3</code>

DEFERTIME	
<b>Format</b>	[ [ [DD:]HH:]MM:]SS
<b>Default</b>	1:00:00
<b>Description</b>	The amount of time a job will be held in the deferred state before being released back to the Idle job queue. <b>Note:</b> A job's defer time will be restarted if Moab is restarted. (For related information, see <a href="#">Reservation Policies</a> , <a href="#">DEFERSTARTCOUNT</a> , <a href="#">RESERVATIONRETRYTIME</a> , <a href="#">NODEFAILURERESERVETIME</a> , <a href="#">JOBRETRYTIME</a> , and <a href="#">GUARANTEEDPREEMPTION</a> .)
<b>Example</b>	<code>DEFERTIME 0:05:00</code>

DELETSTAGEOUTFILES	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	Specifies whether the scheduler should delete explicitly specified stageout files after they are successfully staged. By default, such files are not deleted but are left on the nodes where the job ran.
<b>Example</b>	<pre>DELETSTAGEOUTFILES TRUE Example of an explicit stageout request msub x=MSTAGEOUT:ssh://source_node/tmp/file,file:///results_folder job.cmd</pre> <p><i>With this parameter set to TRUE, /tmp/file on source_node is deleted after it is copied to the specified destination (file:///results_folder). If the parameter is not set, or if it is set to FALSE, /tmp/file remains on source_node after the job terminates.</i></p>

DEPENDFAILUREPOLICY	
<b>Format</b>	HOLD or CANCEL
<b>Default</b>	HOLD
<b>Description</b>	Specifies what happens to a job if its dependencies cannot be fulfilled; that is, what happens when a job depends on another job to complete successfully but the other job fails.
<b>Example</b>	<pre>DEPENDFAILUREPOLICY CANCEL</pre> <p><i>If job A is submitted with depend=afterok:B and job B fails, job A is canceled.</i></p>

DIRECTORYSERVER	
<b>Format</b>	<HOST> [ :<PORT> ]
<b>Default</b>	---
<b>Description</b>	The interface for the directory server.
<b>Example</b>	<code>DIRECTORYSERVER calli3.icluster.org:4702</code>

DISABLEEXCHLIST	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	If the resource manager rejects a job and the value is set to <code>TRUE</code> , then the node is not added to the job's exclude host list.
<b>Example</b>	<code>DISABLEEXCHLIST TRUE</code>

DISABLEINTERACTIVEJOBS	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	<p>Disallows interactive jobs submitted with <code>msub -I</code>.</p> <p><b>Note:</b> It is possible for users to submit interactive jobs directly to a resource manager, which can bypass the <code>DISABLEINTERACTIVEJOBS</code> parameter. However, some resource managers (such as Torque) will check with Moab before allowing an interactive job.</p>
<b>Example</b>	<code>DISABLEINTERACTIVEJOBS TRUE</code>

DISABLEREQUIREDGRESNONE	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	When set to <code>TRUE</code> , this causes Moab to reject msub requests that have a gres of 'none'. <a href="#">ENFORCEGRESACCESS</a> must also be set to <code>TRUE</code> for this feature to work.
<b>Example</b>	<pre>##### moab.cfg ##### ENFORCEGRESACCESS TRUE DISABLEREQUIREDGRESNONE TRUE #####  &gt; msub -A ee -l nodes=1,ttc=5,walltime=600,partition=g02 -l gres=none ERROR: cannot submit job - cannot locate required resource 'none'</pre>

DISABLESAMECREDPREEMPTION	
<b>Format</b>	Comma-delimited list of one or more credentials: ACCT, CLASS, GROUP, QOS, or USER
<b>Default</b>	---
<b>Description</b>	This parameter prevents specified credentials from preempting its own jobs.
<b>Example</b>	<pre>DISABLESAMECREDPREEMPTION QOS,USER</pre>

DISABLESCHEDULING	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	Specifies whether or not the scheduler will schedule jobs. If set to <code>TRUE</code> , Moab will continue to update node and job state but will not start, preempt, or otherwise modify jobs. The command <a href="#">mschedctl -r</a> will clear this parameter and resume normal scheduling.
<b>Example</b>	<pre>DISABLESCHEDULING FALSE</pre>

DISABLESLAVEJOBSUBMIT	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	TRUE
<b>Description</b>	This parameter can be added to the <code>moab.cfg</code> file on an MGM Moab server (in a grid configuration) to prevent users from submitting jobs to the MGC Moab server from the MGM Moab server. Some grid configurations allow the user to submit jobs on the MGM that are migrated to the MGC and submitted from the MGC. Other grid configurations do not allow the jobs to be migrated to the MGC from the MGM, in which case, jobs submitted from the MGM remain idle on the MGM and never run. This parameter will reject the job submissions on the MGM to prevent the submission of jobs that will never run.
<b>Example</b>	<pre>DISABLESLAVEJOBSUBMIT TRUE example (node04 is a slave and node06 is the master) [test@node04 moab-slave]\$ echo sleep 100   msub ERROR: cannot submit job from slave</pre>

DISABLETHRESHOLDTRIGGERS	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	This makes Moab not fire threshold-based triggers, but will log the intended action to the event logs.
<b>Example</b>	<pre>DISABLETHRESHOLDTRIGGERS TRUE</pre>



DISKWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	The priority weight to be applied to the amount of dedicated disk space required per task by a job (in MB).
<b>Example</b>	<pre>RESWEIGHT 10 DISKWEIGHT 100</pre> <p><i>If a job requires 12 tasks and 512 MB per task of dedicated local disk space, Moab will increase the job's priority by <math>10 * 100 * 12 * 512</math></i></p>

DISPLAYFLAGS	
<b>Format</b>	One or more of the following values (space delimited): ACCOUNTCENTRIC, HIDEBLOCKED, HIDECREDS, HIDE DRAINED, NODECENTRIC, USEBLOCKING, USENOBLOCKMSUB
<b>Default</b>	---
<b>Description</b>	<p>By default, no flags (special modifications) are specified. If flags are specified, this controls how Moab client commands display varied information:</p> <ul style="list-style-type: none"> <li>ACCOUNTCENTRIC: Displays account information in <a href="#">showq</a>, rather than group information.</li> <li>HIDEBLOCKED: Prevents <a href="#">showq</a> from listing information about blocked jobs that are not owned by the user if the user is not an admin.</li> <li>HIDECREDS: Users without Moab administrative privileges will not be able to see the credentials of other jobs.</li> <li>HIDE DRAINED: Prevents <a href="#">mdiag -n</a> from displaying nodes in the <b>DRAINED</b> state. An override option of <a href="#">mdiag -n -w nodestate=drained</a> lists only those nodes with a DRAINED state.</li> <li>NODECENTRIC: Displays node allocation information instead of processor allocation information in <a href="#">showq</a>.</li> <li>USEBLOCKING: Disables threading for commands that support it; those commands include <a href="#">showq</a>, <a href="#">mdiag -n</a>, and <a href="#">mdiag -j</a>.</li> <li>USENOBLOCKMSUB: Moab will skip error checking of the msub job submission and queue it up for later processing. The job ID will be returned immediately.</li> </ul>
<b>Example</b>	<code>DISPLAYFLAGS NODECENTRIC</code>

DISPLAYPROXYUSERASUSER	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	If set to TRUE, Moab shows the proxy users instead of the real user on some queries of system jobs that have proxy users. Commands affected include <a href="#">mjobctl -q diag</a> and <a href="#">checkjob</a> .
<b>Example</b>	<code>DISPLAYPROXYUSERASUSER TRUE</code>

DONTCANCELINTERACTIVEHJOBS	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	If set to TRUE, Moab does not cancel interactive jobs that are held.
<b>Example</b>	<code>DONTCANCELINTERACTIVEHJOBS TRUE</code>

DONTENFORCEPEERJOBLIMITS	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	If set to TRUE, only the scheduler that is running the job can cancel the job or enforce other limits.
<b>Example</b>	<code>DONTENFORCEPEERJOBLIMITS TRUE</code>

ENABLEFAILUREFORPURGEDJOB	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	<p>By default, when a job is purged or removed by the Torque resource manager for a walltime violation, the job takes on a state of Completed and a completion code of 0. If TRUE, the job state is set to Removed and has a completion code of 98. ENABLEFAILUREFORPURGEDJOB is for the Torque resource manager only.</p> <div style="border: 1px solid #0070C0; padding: 5px; margin-top: 10px;"> <p><b>i</b> For ENABLEFAILUREFORPURGEDJOB to return Removed job states, you must reset the TORQUE server attribute <code>keep_completed</code> to 0 in <code>qmgr</code>. See 'Queue Attributes' in the <i>Torque Resource Manager Administrator Guide</i> for more information.</p> </div>
<b>Example</b>	<div style="border: 1px dashed #000; padding: 5px; margin-bottom: 5px;"> <p>ENABLEFAILUREFORPURGEDJOB TRUE</p> </div> <div style="border: 1px dashed #000; padding: 5px;"> <p><i>Jobs that are purged or removed by Torque are given a state of Removed and a completion code of 98.</i></p> </div>

ENABLEFSVIOLATIONPREEMPTION	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	<p>If set to TRUE, Moab will allow jobs within the same <a href="#">class/queue</a> to <a href="#">preempt</a> when the preemptee is violating a <a href="#">fairshare</a> target and the preemptor is not.</p>
<b>Example</b>	<div style="border: 1px dashed #000; padding: 5px;"> <p>ENABLEFSVIOLATIONPREEMPTION TRUE</p> </div>

ENABLEHIGHTHROUGHPUT	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	<p>Reduces iteration times by eliminating string error checking during checkpointing, eliminating automatic rack processing, reducing object caching, using vfork rather than fork, reducing RM timeout parameters, and scheduling similar jobs as a chunk rather than individually.</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> If <code>ENABLEHIGHTHROUGHPUT</code> is <code>TRUE</code>, you must set <code>NODEALLOCATIONPOLICY</code> to <code>FIRSTAVAILABLE</code>.</p> </div>

ENABLEJOBARRAYS	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	TRUE
<b>Description</b>	If set to <code>TRUE</code> , job arrays will be enabled.
<b>Example</b>	<code>ENABLEJOBARRAYS TRUE</code>

ENABLENEGJOBPRIORITY	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	If set to <code>TRUE</code> , the scheduler allows job priority value to range from <code>-INFINITY</code> to <code>MMAX_PRIO</code> ; otherwise, job priority values are given a lower bound of '1'. For more information, see <a href="#">REJECTNEGPRIOJOBS</a> .
<b>Example</b>	<code>ENABLENEGJOBPRIORITY TRUE</code> <i>Job priority can range from <code>-INFINITY</code> to <code>MMAX_PRIO</code>.</i>

ENABLENODEADDRLOOKUP	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	If set to <code>TRUE</code> , the scheduler will use the default host name service lookup mechanism (i.e., <code>/etc/hosts</code> , DNS, NIS, etc.) to determine the IP address of the nodes reported by the resource manager. This information is used to correlate information reported by multi-homed hosts.
<b>Example</b>	<pre>ENABLENODEADDRLOOKUP TRUE</pre>

ENABLEPOSUSERPRIORITY	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	If set to <code>TRUE</code> , the scheduler will allow users to specify positive job priority values that will be honored. In other words, users can specify a priority that falls in the range of -1024 to +1023, inclusive. If set to <code>FALSE</code> (the default), user priority values are given an upper bound of '0' when users request a positive priority. See <a href="#">USERPRIORITY</a> .
<b>Example</b>	<pre>ENABLEPOSUSERPRIORITY TRUE</pre> <i>Users can now specify positive job priorities and have them take effect (e.g., <code>msub -p 100 job.cmd</code>).</i>

ENABLESPVIOLATIONPREEMPTION	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	If set to <code>TRUE</code> , Moab will allow jobs within the same <a href="#">class/queue</a> to <a href="#">preempt</a> when the preemptee is violating a <a href="#">soft usage policy</a> and the preemptor is not.
<b>Example</b>	<pre>ENABLESPVIOLATIONPREEMPTION TRUE</pre>

ENFORCEACCOUNTACCESS	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	Specifies whether or not Moab will enforce account access constraints without an <a href="#">accounting manager</a> .
<b>Example</b>	<pre>ENFORCEACCOUNTACCESS TRUE</pre>

ENFORCEGRESACCESS	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	If a user submits a job with a non-existent gres (e.g., in the case of a typo) and ENFORCEGREACCESS is not set in <code>moab.cfg</code> , or is set to FALSE, then the requested gres will be created (but will not exist on any nodes) and the job will be deferred (similar to ENFORCEACCOUNTACCESS).
<b>Example</b>	<pre>ENFORCEGRESACCESS TRUE</pre>

EVENTFILEFORMAT	
<b>Format</b>	One of JSON, WHITESPACE, or WIKI
<b>Default</b>	WIKI
<b>Description</b>	The format to write the event log.
<b>Example</b>	<pre>EVENTFILEFORMAT WIKI</pre>

EVENTSERVER	
<b>Format</b>	<HOST> [ :<PORT> ]
<b>Default</b>	---
<b>Description</b>	The interface for the event server.
<b>Example</b>	<pre>EVENTSERVER calli3.icluster.org:4702</pre>

FEATURENODETYPEHEADER	
<b>Format</b>	<STRING>
<b>Default</b>	---
<b>Description</b>	The header used to specify node type via node features (for example, PBS node attributes).
<b>Example</b>	<pre>FEATURENODETYPEHEADER xnt</pre> <p><i>Moab will interpret all node features with the leading string <code>xnt</code> as a nodetype specification, as used by the accounting manager and other accounting managers, and assign the associated value to the node (for example, <code>xntFast</code>).</i></p>

FEATUREPARTITIONHEADER	
<b>Format</b>	<STRING>
<b>Default</b>	---
<b>Description</b>	The header used to specify node partition via node features (for example, PBS node attributes).
<b>Example</b>	<pre>FEATUREPARTITIONHEADER xpt</pre> <p><i>Moab will interpret all node features with the leading string <code>xpt</code> as a partition specification and assign the associated value to the node (for example, <code>xptFast</code>).</i></p>



FEATUREPROCSPEEDHEADER	
<b>Format</b>	<STRING>
<b>Default</b>	---
<b>Description</b>	The header used to extract node processor speed via node features (i.e., LL features or PBS node attributes). <b>Note:</b> Adding a trailing '\$' character will specify that only features with a trailing number be interpreted. For example, the header 'sp\$' will match 'sp450' but not 'sport'.
<b>Example</b>	<pre>FEATUREPROCSPEEDHEADER xps</pre> <p><i>Moab will interpret all node features with the leading string <code>xps</code> as a processor speed specification and assign the associated value to the node (i.e., <code>xps950</code>).</i></p>

FEATURERACKHEADER	
<b>Format</b>	<STRING>
<b>Default</b>	---
<b>Description</b>	The header used to extract node rack index via node features (i.e., LL features or PBS node attributes). <b>Note:</b> Adding a trailing '\$' character will specify that only features with a trailing number be interpreted. For example, the header 'rack\$' will match 'rack4' but not 'racket'.
<b>Example</b>	<pre>FEATURERACKHEADER rack</pre> <p><i>Moab will interpret all node features with the leading string <code>rack</code> as a rack index specification and assign the associated value to the node (i.e., <code>rack16</code>).</i></p>

FEATURESLOTHEADER	
<b>Format</b>	<STRING>
<b>Default</b>	---
<b>Description</b>	The header used to extract node slot index via node features (i.e., LL features or PBS node attributes). <b>Note:</b> Adding a trailing '\$' character will specify that only features with a trailing number be interpreted. For example, the header 'slot\$' will match 'slot12' but not 'slotted'.
<b>Example</b>	<pre>FEATURESLOTHEADER slot</pre> <p><i>Moab will interpret all node features with the leading string <code>slot</code> as a slot index specification and assign the associated value to the node (i.e., <code>slot16</code>).</i></p>

FEEDBACKPROGRAM	
<b>Format</b>	<STRING>
<b>Default</b>	---
<b>Description</b>	The name of the program to be run at the completion of each job. If not fully qualified, Moab will attempt to locate this program in the 'tools' subdirectory. For more details on how this works and what fields are provided, see <a href="#">User Feedback Overview</a> .
<b>Example</b>	<pre>FEEDBACKPROGRAM /var/moab/fb.pl</pre> <p><i>Moab will run the specified program at the completion of each job.</i></p>

FILEREQUESTISJOBCENTRIC	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	Specifies whether a job's file request is a total request for the job or a per task request.
<b>Example</b>	<pre>FILEREQUESTISJOBCENTRIC TRUE</pre> <p><i>Moab will treat file requests as a total request per job.</i></p>

FILTERCMDFILE	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	TRUE
<b>Description</b>	<p>Running the <a href="#">msub</a> command performs the following operations on the submission script:</p> <ul style="list-style-type: none"> <li>• Replace all comments with spaces (excludes Resource Manager directives)</li> <li>• Strip empty lines</li> <li>• Replace <code>\r</code> with <code>\n</code></li> <li>• Lock job to a PBS resource manager if <code>\$PBS</code> is found in the script</li> </ul> <p>Include the <code>FILTERCMDFILE</code> parameter in the Moab configuration file that resides on the clients (<code>moab.cfg</code> or <code>moab-client.cfg</code>).</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> <code>FILTERCMDFILE</code> must be <code>FALSE</code> for <a href="#">REJECTDOSSCRIPTS</a> to work correctly.</p> </div>
<b>Example</b>	<pre>FILTERCMDFILE FALSE</pre> <p><i>Running the <code>msub</code> command does not perform the actions detailed earlier.</i></p>

FORCENODEREPROVISION	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	When set to TRUE, this config option causes Moab to reprovision a node, even if it is to the same operating system (in essence rewriting the OS).
<b>Example</b>	<pre>FORCENODEREPROVISION TRUE</pre>

FORCERSVSUBTYPE	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	Specifies that admin reservations must have a <a href="#">subtype</a> associated with them.
<b>Example</b>	<pre>FORCERSVSUBTYPE TRUE</pre> <i>Moab will require all admin reservations to include a subtype.</i>

FREETIMELOOKAHEADDURATION	
<b>Format</b>	[ [ [DD:]HH:]MM:]SS
<b>Default</b>	2 Months
<b>Description</b>	Specifies how far ahead Moab will look when calculating free time on a node.
<b>Example</b>	<pre>FREETIMELOOKAHEADDURATION 7:00:00:00</pre> <i>Moab will look 1 week ahead when it calculates free time on a node.</i>

FSACCOUNTWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	1000
<b>Description</b>	The weight assigned to the account subcomponent of the fairshare component of priority.
<b>Example</b>	<code>FSACCOUNTWEIGHT 10</code>

FSCAP	
<b>Format</b>	<DOUBLE>
<b>Default</b>	0 (NO CAP)
<b>Description</b>	The maximum allowed absolute value for a job's total preweighted fairshare component.
<b>Example</b>	<code>FSCAP 10.0</code> <i>Moab will bind a job's preweighted fairshare component by the range +/- 10.0.</i>

FSCLASSWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	1000
<b>Description</b>	The weight assigned to the class subcomponent of the fairshare component of priority.
<b>Example</b>	<code>FSCLASSWEIGHT 10</code>

FSDECAY	
<b>Format</b>	<DOUBLE>
<b>Default</b>	1.0
<b>Description</b>	Specifies decay rate applied to past fairshare interval when computing effective fairshare usage. Values can be in the range of 0.01 to 1.0. A smaller value causes more rapid decay causing <i>aged</i> usage to contribute less to the overall effective fairshare usage. A value of 1.0 indicates that no decay will occur and all fairshare intervals will be weighted equally when determining effective fairshare usage. See <a href="#">Fairshare Overview</a> .
<b>Example</b>	<pre> FSPOLICY    DEDICATEDPS FSDECAY     0.8 FSDEPTH     8           </pre> <p><i>Moab will apply a decay rate of 0.8 to all fairshare windows.</i></p>

FSDEPTH	
<b>Format</b>	<INTEGER>
<b>Default</b>	8
<b>Description</b>	<b>Note:</b> The number of available fairshare windows is bounded by the MAX_FSDEPTH value (32 in Moab). See <a href="#">Fairshare Overview</a> .
<b>Example</b>	<pre> FSDEPTH 12           </pre>

FSDISABLEIFNEGATIVE	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	Disable the fairshare component of priority if it is negative for a particular job.
<b>Example</b>	<pre> FSDISABLEIFNEGATIVE TRUE           </pre> <p>Moab will disable any fairshare components that are negative when determining a job's priority.</p>

FSENABLECAPRIORITY	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	Fairshare priority will increase to target and stop.
<b>Example</b>	<code>FSENABLECAPRIORITY TRUE</code>

FSGROUPWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	1000
<b>Description</b>	The weight assigned to the group subcomponent of the fairshare component of priority.
<b>Example</b>	<code>FSGROUPWEIGHT 4</code>

FSINTERVAL	
<b>Format</b>	[ [ [DD:]HH:]MM:]SS
<b>Default</b>	12:00:00
<b>Description</b>	The length of each fairshare <a href="#">window</a> .
<b>Example</b>	<code>FSINTERVAL 12:00:00</code> <i>Track fairshare usage in 12 hour blocks.</i>

FSJPUWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	The fairshare weight assigned to jobs per user.
<b>Example</b>	<code>FSJPUWEIGHT 10</code>

FSMOSTSPECIFICLIMIT	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	When checking policy usage limits in a fairshare tree, if the most specific policy limit is passed then do not check the same policy again at higher levels in the tree. For example, if a user has a MaxProc policy limit then do not check the MaxProc policy limit on the account for this same user.
<b>Example</b>	<code>FSMOSTSPECIFICLIMIT TRUE</code>

FSPOLICY	
<b>Format</b>	<POLICY> [*] See <a href="#">5.3.1.A FSPOLICY - Specifying the Metric of Consumption - page 403</a> for valid values.
<b>Default</b>	[NONE]
<b>Description</b>	The unit of tracking fairshare usage.
<b>Example</b>	<code>FSPOLICY DEDICATEDPES</code> <i>Moab will track fairshare usage by dedicated processor-equivalent seconds.</i>



FSPPUWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	The fairshare weight assigned to processors per user.
<b>Example</b>	<code>FSPPUWEIGHT 10</code>

FSPSPUWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	The fairshare weight assigned to processor-seconds per user.
<b>Example</b>	<code>FSPSPUWEIGHT 10</code>

FSQOSWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	1000
<b>Description</b>	The priority weight assigned to the QOS fairshare subcomponent.
<b>Example</b>	<code>FSQOSWEIGHT 16</code>

FSTARGETISABSOLUTE	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	Specifies whether Moab should base fairshare targets off of delivered cycles or up/available cycles.
<b>Example</b>	<pre>FSTARGETISABSOLUTE TRUE</pre>

FSTREE	
<b>Format</b>	List of zero or more space delimited <ATTR>=<VALUE> pairs, where <ATTR> is one of the following: <a href="#">SHARES</a> or <a href="#">MEMBERLIST</a>
<b>Default</b>	---
<b>Description</b>	The share tree distribution for job fairshare prioritization (see <a href="#">Hierarchical Share Tree Overview</a> ).
<b>Example</b>	<pre>FSTREE[geo] SHARES=16 MEMBERLIST=geo103,geo313,geo422</pre>

FSTREEACLPOLICY	
<b>Format</b>	OFF, PARENT, or FULL
<b>Default</b>	FULL
<b>Description</b>	Specifies how Moab should interpret credential membership when building the FSTREE (see <a href="#">Hierarchical Share Tree Overview</a> ).
<b>Example</b>	<pre>FSTREEACLPOLICY PARENT</pre> <i>Credentials will be given access to their parent node when applicable.</i>

FSTREEISREQUIRED	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	Specifies whether a job must have an applicable node in the partition's FSTREE in order to execute within that partition (see <a href="#">Hierarchical Share Tree Overview</a> ).
<b>Example</b>	<pre>FSTREEISREQUIRED TRUE</pre> <p><i>Jobs must have an applicable node in the FSTREE in order to execute.</i></p>

FSTREEUSERISREQUIRED	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	Specifies whether the user must be given explicit access to a branch in the FSTREE (see <a href="#">Hierarchical Share Tree Overview</a> ).
<b>Example</b>	<pre>FSTREEUSERISREQUIRED TRUE</pre> <p><i>Users must be given explicit access to FSTREE nodes in order to gain access to the FSTREE.</i></p>

FSUSERWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	1000
<b>Description</b>	The priority weight assigned to the user fairshare subfactor.
<b>Example</b>	<pre>FSUSERWEIGHT 8</pre>

FSWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	1
<b>Description</b>	The priority weight assigned to the summation of the fairshare subfactors (see <a href="#">Priority Factor</a> and <a href="#">Fairshare</a> overviews).
<b>Example</b>	<pre>FSWEIGHT 500</pre>

GEVENTCFG[<GEVENT>]	
<b>Format</b>	List of zero or more space delimited <ATTR>=<VALUE> pairs. See <a href="#">10.9.1 Configuring Generic Events - page 617</a> for details on values you can assign to each attribute.
<b>Default</b>	---
<b>Description</b>	Specifies how the scheduler should behave when various cluster events are detected. See the <a href="#">10.9 Enabling Generic Events - page 616</a> for more information.
<b>Example</b>	<pre>GEVENTCFG[hitemp] ACTION=avoid,record,notify REARM=00:10:00 GEVENT[nodeerror] SEVERITY=3</pre> <p><i>If a hitemp event is detected, Moab adjusts the node allocation policy to minimize the allocation of the node. Moab also sends emails to cluster administrators and reports the event in the Moab event log.</i></p>

GRES CFG[<GRES>]	
<b>Format</b>	List of zero or more space delimited <ATTR>=<VALUE> pairs, where <ATTR> is one of the following: FEATUREGRES, LICENSE, INVERTTASKCOUNT, PRIVATE, TYPE, or STARTDELAY.
<b>Default</b>	---
<b>Description</b>	<p>Specifies associations of generic resources into resource groups:</p> <ul style="list-style-type: none"> <li>• When LICENSE is set to TRUE, Moab will pass this generic resource to the accounting manager in the Licenses property rather than the Resources property.</li> <li>• When PRIVATE is set to TRUE, Moab puts the requested generic resource on a separate job request. By default, a private request is a request with 1 task with X number of generic resources per task.</li> <li>• When INVERTTASKCOUNT and PRIVATE are set to TRUE, Moab makes the generic resource's private request a request with X number of tasks with 1 generic resource per task.</li> </ul> <p>See <a href="#">12.6 Managing Consumable Generic Resources</a> for more information.</p>
<b>Example</b>	<pre>GRES CFG[scsi1] TYPE=fastio GRES CFG[scsi2] TYPE=fastio GRES CFG[scsi3] TYPE=fastio</pre> <p><i>The generic resources scsi1, scsi2, and scsi3 are all associated with the generic resource type fastio.</i></p>

GRES TOJOBATTR	
<b>Format</b>	Comma-delimited list of generic resources
<b>Default</b>	---
<b>Description</b>	The list of generic resources will also be interpreted as JOB features. See <a href="#">Managing Reservations</a> .
<b>Example</b>	<pre>GRES TOJOBATTR matlab,ccs</pre> <p><i>Jobs that request the generic resources matlab or ccs will have a corresponding job attribute assigned to them.</i></p>


GROUPCFG[<GROUPID>]	
<b>Format</b>	List of zero or more space delimited <ATTR>=<VALUE> pairs, where <ATTR> is one of the following: <a href="#">General Credential Flags</a> , PRIORITY, <a href="#">ENABLEPROFILING</a> , QLIST, QDEF, PLIST, FLAGS, <a href="#">usage limits</a> , or a <a href="#">fairshare usage limit</a> specification.
<b>Default</b>	---
<b>Description</b>	Specifies group specific attributes. See the <a href="#">flag overview</a> for a description of legal flag values.
<b>Example</b>	<pre>GROUPCFG[staff] MAXJOB=50 QDEF=highprio</pre> <p><i>Up to 50 jobs submitted by members of the group <code>staff</code> will be allowed to execute simultaneously and will be assigned the QOS <code>highprio</code> by default.</i></p>

GROUPWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	1
<b>Description</b>	The priority weight assigned to the specified group priority (see <a href="#">Credential (CRED) Factor</a> ).
<b>Example</b>	<pre>GROUPWEIGHT 20</pre>

GUARANTEEDPREEMPTION	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	<p>Causes Moab to lock <a href="#">PREEMPTOR</a> jobs until <a href="#">JOBRETRYTIME</a> expires (essentially, waiting for the <a href="#">PREEMPTEE</a> jobs to finish).</p> <p>It may take some time for the <a href="#">PREEMPTEE</a> jobs to clear out. During that time, the <a href="#">PREEMPTOR</a> job might want to look elsewhere to run, which would be disruptive as it might preempt another set of jobs. If you want to prevent this, it is recommended that you set <code>GUARANTEEDPREEMPTION</code> to <code>TRUE</code>.</p> <p>For related information, see <a href="#">About preemption</a>, <a href="#">Reservation Policies</a>, <a href="#">DEFERSTARTCOUNT</a>, <a href="#">DEFERTIME</a>, <a href="#">RESERVATIONRETRYTIME</a>, <a href="#">NODEFAILURERESERVETIME</a>, and <a href="#">JOBRETRYTIME</a>.</p>
<b>Example</b>	<code>GUARANTEEDPREEMPTION TRUE</code>

HALOCKCHECKTIME	
<b>Format</b>	[[[DD:]HH:]MM:]SS
<b>Default</b>	9
<b>Description</b>	Specifies how frequently the secondary server checks the timestamp on the lock file. See <a href="#">High Availability Overview</a> for more info.
<b>Example</b>	<code>HALOCKCHECKTIME 00:00:15</code> <p><i>The Moab fallback server will check the health of the Moab primary server every 15 seconds.</i></p>

HALOCKUPDATETIME	
<b>Format</b>	[[[DD:]HH:]MM:]SS
<b>Default</b>	3
<b>Description</b>	Specifies how frequently the primary server checks the timestamp on the lock file. See <a href="#">High Availability Overview</a> for more info.
<b>Example</b>	<pre>HALOCKUPDATETIME 00:00:03</pre> <p><i>The Moab primary server will check the timestamp of the lock file every 3 seconds.</i></p>

IDCFG	
<b>Format</b>	One or more of the following attribute/value pairs: BLOCKEDCREDLIST, CREATECRED, REFRESHPERIOD, REQUIREUSERLIST, RESETCREDLIST, or SERVER. See <a href="#">Identity Managers</a> for additional information.
<b>Default</b>	---
<b>Description</b>	<p>This parameter enables the <a href="#">identity manager</a> interface allowing credential, policy, and usage information to be shared with an external information service.</p> <div style="border: 1px solid #0070C0; border-radius: 5px; padding: 5px; margin-top: 10px;"> <p> Only one identity manager can be configured at a time.</p> </div>
<b>Example</b>	<pre>IDCFG[info] SERVER=exec:///usr/local/bin/dbquery.pl REFRESHPERIOD=30:00</pre> <p><i>Moab will refresh credential info every half hour using the STDOUT of the specified script.</i></p>



IGNOREMDATASTAGING	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	When set to <code>TRUE</code> , Moab will ignore any resource manager specific data staging on a job and assume the resource manager is processing the request. Currently, this only applies to PBS.
<b>Example</b>	<pre>IGNORERMDATASTAGING TRUE</pre>

IGNORECLASSES	
<b>Format</b>	[!]<CLASS> [, <CLASS>] . . .
<b>Default</b>	---
<b>Description</b>	By default, if using the Torque resource manager, jobs from all listed classes are ignored and not scheduled, tracked, or otherwise processed by Moab. If the <b>not</b> (i.e., '!') character is specified, only jobs from listed classes are processed. See the <a href="#">Moab Side-by-Side</a> for more information.
<b>Example</b>	<pre>IGNORECLASSES dque, batch</pre> <i>Moab will ignore jobs from classes dque and batch.</i>

IGNOREJOBS	
<b>Format</b>	[!]<JOBID> [, <JOBID>] . . .
<b>Default</b>	---
<b>Description</b>	By default, listed jobs are ignored and not scheduled, tracked, or otherwise processed by Moab. If the <b>not</b> (i.e., '!') character is specified, only listed jobs are processed. See the <a href="#">Moab Side-by-Side Analysis</a> for more information.
<b>Example</b>	<pre>IGNOREJOBS !14221, 14223</pre> <i>Moab will ignore jobs all classes except 14221 and 14223.</i>

IGNORENODES	
<b>Format</b>	[ ! ] <NODE> [ , <NODE> ] . . .
<b>Default</b>	---
<b>Description</b>	By default, all listed nodes are ignored and not scheduled, tracked, or otherwise processed by Moab. If the <b>not</b> (i.e., '!') character is specified, only listed nodes are processed. See the <a href="#">Moab Side-by-Side Analysis</a> for more information.
<b>Example</b>	<pre>IGNORENODES !host3,host4</pre> <p><i>Moab will only process nodes host3 and host4.</i></p>


IGNOREPREEMPTTEEPRIORITY	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	By default, preemptor jobs can only <b>preempt</b> preemptee jobs if the preemptor has a higher job <b>priority</b> than the preemptee. When this parameter is set to true, the priority constraint is removed allowing any preemptor to preempt any preemptees once it reaches the top of the eligible job queue.
<b>Example</b>	<pre>IGNOREPREEMPTTEEPRIORITY TRUE</pre> <p><i>A preemptor job can preempt any preemptee jobs when it is at the top of the eligible job queue.</i></p>

IGNOREUSERS	
<b>Format</b>	[!]<USERNAME> [, <USERNAME>] . . .
<b>Default</b>	---
<b>Description</b>	By default, if using the Torque resource manager, jobs from all listed users are ignored and not scheduled, tracked, or otherwise processed by Moab. If the <b>not</b> (i.e., '!') character is specified, only jobs from listed users are processed. See the <a href="#">Moab Side-by-Side Analysis</a> for more information.
<b>Example</b>	<pre>IGNOREUSERS testuser1,annapolis</pre> <p><i>Moab will ignore jobs from users testuser1 and annapolis.</i></p>

#INCLUDE	
<b>Format</b>	<STRING>
<b>Default</b>	---
<b>Description</b>	Specifies another file that contains more configuration parameters. If <STRING> is not an absolute path, Moab will search its home directory for the file.
<b>Example</b>	<pre>#INCLUDE moab.acct</pre> <p><i>Moab will process the parameters in moab.acct, as well as moab.cfg</i></p>

INSIGHTENDPOINT	
<b>Format</b>	<hostname>[:<port>]
<b>Default</b>	---
<b>Description</b>	Enables Moab Workload Manager to connect to Moab Insight. <hostname> is the server where Insight is located. <hostname> is required, <port> is optional.

### INSTANTSTAGE

<b>Description</b>	 This parameter is deprecated and may be removed in a future release. Use <a href="#">JOBMIGRATEPOLICY</a> .
--------------------	---

### INVALIDFSTREEMSG

<b>Format</b>	"<STRING>"
<b>Default</b>	"no valid fstree node found"
<b>Description</b>	The error message that should be attached to jobs that cannot run because of a fairshare tree configuration violation.
<b>Example</b>	<pre>INVALIDFSTREEMSG "account is invalid for requested partition"</pre>

### JOBACTIONONNODEFAILURE

<b>Format</b>	<a href="#">CANCEL</a> , <a href="#">FAIL</a> , <a href="#">HOLD</a> , <a href="#">IGNORE</a> , <a href="#">NOTIFY</a> , or <a href="#">REQUEUE</a>
<b>Default</b>	---
<b>Description</b>	<p>By default, Moab does not report information when a node allocated to an active job has failed (state is <a href="#">down</a>).</p> <p>Use this parameter to specify the action to take if Moab detects that a node allocated to an active job has failed. Moab only reports this information via diagnostic commands. If this parameter is set, Moab will cancel or requeue the active job. See <a href="#">4.4.6 Allocated Resource Failure Policy for Jobs - page 373</a> for more information.</p> <p><b>Note:</b> The <a href="#">HOLD</a> value is only applicable when using <a href="#">checkpointing</a>.</p>
<b>Example</b>	<pre>JOBACTIONONNODEFAILURE REQUEUE</pre> <p><i>Moab will requeue active jobs that have allocated nodes that have failed during the execution of the job.</i></p>

JOBAGGREGATIONTIME	
<b>Format</b>	[ [ [DD:] HH:] MM:] SS
<b>Default</b>	0
<b>Description</b>	The minimum amount of time the scheduler should wait after receiving a job event until it should process that event. This parameter allows sites with bursty job submissions to process job events in groups decreasing total job scheduling cycles and allowing the scheduler to make more intelligent choices by aggregating job submissions and choosing between the jobs. See <a href="#">Appendix I: Considerations for Large Clusters - page 1390</a> for more information.
<b>Example</b>	<pre>JOBAGGREGATIONTIME 00:00:04 RMPOLLINTERVAL    30,30</pre> <p><i>Moab will wait 4 seconds between scheduling cycles when job events have been received and will wait 30 seconds between scheduling cycles otherwise.</i></p>

JOBCFG	
<b>Format</b>	<p>&lt;ATTR&gt;=&lt;VAL&gt; where &lt;ATTR&gt; is one of <a href="#">ACCOUNT</a>, <a href="#">CLASS</a>, <a href="#">CPUCLOCK</a>, <a href="#">CPULIMIT</a>, <a href="#">DESCRIPTION</a>, <a href="#">DPROCS</a>, <a href="#">ENV</a>, <a href="#">EXEC</a>, <a href="#">FLAGS</a>, <a href="#">GNAME</a>, <a href="#">GRES</a>, <a href="#">GROUP</a>, <a href="#">MEM</a>, <a href="#">NODEACCESSPOLICY</a>, <a href="#">NODES</a>, <a href="#">NODESET</a>, <a href="#">PARTITION</a>, <a href="#">PREF</a>, <a href="#">PRIORITY</a>, <a href="#">PRIORITYF</a>, <a href="#">QOS</a>, <a href="#">RARCH</a>, <a href="#">RFEATURES</a>, <a href="#">RM</a>, <a href="#">ROPSYS</a>, <a href="#">SELECT</a>, <a href="#">SYSTEMJOBTYPE</a>, <a href="#">TASKS</a>, <a href="#">TASKPERNODE</a>, <a href="#">TEMPLATEDDEPEND</a>, <a href="#">UNAME</a>, <a href="#">USER</a>, <a href="#">VARIABLE</a>, <a href="#">WCLIMIT</a></p>
<b>Default</b>	---
<b>Description</b>	<p>Specifies attributes for jobs that satisfy the specified profile. The <code>SELECT</code> attribute allows users to specify the job template by using <code>msub -l template=</code>. The <code>JOBCFG</code> parameter supports the following attributes:</p> <p><code>ACCOUNT</code>, <code>CLASS</code>, <code>CPUCLOCK</code>, <code>CPULIMIT</code>, <code>DESCRIPTION</code>, <code>DPROCS</code>, <code>ENV</code>, <code>EXEC</code>, <code>FLAGS</code>, <code>GNAME</code>, <code>GRES</code>, <code>GROUP</code>, <code>MEM</code>, <code>NODEACCESSPOLICY</code>, <code>NODES</code>, <code>NODESET</code>, <code>PARTITION</code>, <code>PREF</code>, <code>PRIORITY</code>, <code>PRIORITYF</code>, <code>QOS</code>, <code>RARCH</code>, <code>RFEATURES</code>, <code>RM</code>, <code>ROPSYS</code>, <code>SELECT</code>, <code>SYSTEMJOBTYPE</code>, <code>TASKS</code>, <code>TASKPERNODE</code>, <code>TEMPLATEDDEPEND</code>, <code>UNAME</code>, <code>USER</code>, <code>VARIABLE</code>, <code>WCLIMIT</code></p> <p>It also supports the following Wiki attributes:</p> <p><code>ARGS</code>, <code>DMEM</code>, <code>DDISK</code>, <code>DWAP</code>, <code>ERROR</code>, <code>EXEC</code>, <code>EXITCODE</code>, <code>GATTR</code>, <code>GEVENT</code>, <code>IWD</code>, <code>JNAME</code>, <code>NAME</code>, <code>PARTITIONMASK</code>, <code>PRIORITYF</code>, <code>RDISK</code>, <code>RSWAP</code>, <code>RAGRES</code>, <code>RCGRES</code>, <code>TASKPERNODE</code>, <code>TRIGGER</code>, <code>VARIABLE</code>, <code>NULL</code></p> <p><b>Note:</b> The <i>index</i> to the <code>JOBCFG</code> parameter can either be an admin-chosen <a href="#">job template</a> name or the exact name of job reported by one or more <a href="#">workload queries</a>. See <a href="#">Wiki Attributes</a> and <a href="#">Job Template Extensions</a>.</p>
<b>Example</b>	<div style="border: 1px dashed gray; padding: 10px; border-radius: 10px;"> <pre>JOBCFG[sql] RFEATURES=sqlnode QOS=service</pre> </div> <div style="border: 1px dashed gray; padding: 10px; border-radius: 10px; margin-top: 10px;"> <p><i>When the <code>sql</code> job is detected, it will have the specified default qos and node feature attributes set.</i></p> </div>

JOBPURGETIME	
<b>Format</b>	[ [ [DD:] HH:] MM:] SS
<b>Default</b>	00:05:00
<b>Description</b>	The amount of time Moab will preserve detailed information about a completed job (see <a href="#">showq -c</a> and <a href="#">checkjob</a> ).
<b>Example</b>	<pre>JOBPURGETIME 02:00:00</pre> <p><i>Moab will maintain detailed job information for 2 hours after a job has completed.</i></p>

JOBTRUNCATENLCP	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	TRUE
<b>Description</b>	Specifies whether Moab will store only the first node of the node list for a completed job in the checkpoint file.
<b>Example</b>	<pre>JOBTRUNCATENLCP TRUE</pre> <p><i>JOBTRUNCATENLCP reduces the amount of memory Moab uses to store completed job information.</i></p>

JOBEXTENDSTARTWALLTIME	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	<p>Extends the job walltime when Moab starts the job up to the lesser of the maximum or the next reservation (rounded down to the nearest minute).</p> <div style="border: 1px solid #004a7c; border-radius: 10px; padding: 10px; background-color: #e6f2ff; margin-top: 10px;"> <p><b>i</b> JOBEXTENDSTARTWALLTIME TRUE and JOBEXTENDDURATION cannot be configured together. If they are in the same moab.cfg or are both active, then the JOBEXTENDDURATION will not be honored.</p> </div>
<b>Example</b>	<div style="border: 1px dashed #ccc; border-radius: 10px; padding: 5px; margin-bottom: 10px;"> <pre>JOBEXTENDSTARTWALLTIME TRUE</pre> </div> <p>Submit job with a minimum wallclock limit and a walltime; for example:</p> <div style="border: 1px dashed #ccc; border-radius: 10px; padding: 5px; margin-bottom: 10px;"> <pre>echo sleep 500   msub -A ee -l nodes=5,minwclimit=5:00,walltime=30:00,partition=g02</pre> </div> <div style="border: 1px dashed #ccc; border-radius: 10px; padding: 10px; background-color: #fff9e6; margin-top: 10px;"> <p><i>At job start, Moab recognizes the nodes assigned to the specified job and extends the walltime for the job (one time at job start) up to the lesser of the maximum walltime requested or the least amount of time available for any of the nodes until the next reservation on that node.</i></p> </div>




JOBFAILRETRYCOUNT	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	<p>The number of times a job is requeued and restarted by Moab if the job fails (if the job itself returns a non-zero exit code). Some types of jobs might succeed if automatically retried several times in short succession. This parameter was created with these types of jobs in mind. Note that the job in question must also be restartable (the job needs to have the 'RESTARTABLE' flag set on it) and the RM managing the job must support requeuing and starting completed jobs.</p> <p>If a job fails too many times, and reaches the number of retries given by JobFailRetryCount, then a UserHold is placed on the job and a message is attached to it signifying that the job has a 'restart count violation.'</p>
<b>Example</b>	<pre>JOBFAILRETRYCOUNT 7</pre> <p><i>Any job with a RESTARTABLE flag is requeued, if it fails, up to 7 times before a UserHold is placed on it.</i></p>

JOBIDWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	---
<b>Description</b>	The weight to be applied to the job's ID. See <a href="#">Attribute (ATTR) Factor</a> .
<b>Example</b>	<pre>JOBIDWEIGHT -1</pre> <p><i>Later jobs' priority will be negatively affected.</i></p>

JOBMATCHCFG	
<b>Format</b>	<ATTR>=<VAL> where <ATTR> is one of <a href="#">JMIN</a> , <a href="#">JMAX</a> , <a href="#">JDEF</a> , <a href="#">JSET</a> , or <a href="#">JSTAT</a>
<b>Default</b>	---
<b>Description</b>	The job templates that must be matched and that will be applied in the case of a match.
<b>Example</b>	<pre>JOBMATCHCFG[sql] JMIN=interactive JSTAT=istat</pre>

JOBMAXHOLDTIME	
<b>Format</b>	[ [ [DD:]HH:]MM:]SS
<b>Default</b>	0 (meaning, no max hold time)
<b>Description</b>	The amount of time a job can be held before it is canceled automatically.
<b>Example</b>	<pre>JOBMAXHOLDTIME 02:00:00</pre> <i>Moab will keep jobs in any HOLD state for 2 hours before canceling them.</i>

JOBMAXNODECOUNT	
<b>Format</b>	<INTEGER>
<b>Default</b>	1024
<b>Description</b>	The maximum number of nodes that can be allocated to a job. After changing this parameter, Moab must be restarted. <b>Note:</b> This value cannot exceed either <b>M_MAX_NODE</b> or <b>M_MAX_TASK_PER_JOB</b> . If larger values are required, these values must also be increased. Moab must be restarted before changes to this command will take effect. The command <a href="#">mdiag -S</a> will indicate if any job node count overflows have occurred. See <a href="#">Consideration for Large Clusters</a> .   Moab only reads in this setting when starting up (or restarting).
<b>Example</b>	<pre>JOBMAXNODECOUNT 4000</pre>

JOBMAXOVERRUN	
<b>Format</b>	[[ [DD:]HH:]MM:]SS,[[ [DD:]HH:]MM:]SS
<b>Default</b>	(no soft limit), 10 minutes (hard limit)
<b>Description</b>	<p>Soft and hard limit of the amount of time Moab will allow a job to exceed its wallclock limit before it first sends a mail to the primary admin (soft limit) and then terminates the job (hard limit). See <a href="#">WCVIOLATIONACTION</a> or <a href="#">Usage-based Limits</a>.</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> If you run Moab with the Torque resource manager, you must set the <code>\$ignwalltime</code> parameter to <code>true</code> in the <code>/var/spool/torque/mom_priv/config</code> file; otherwise the <code>pbs_mom</code> will kill any job that exceeds its walltime. See '<code>\$ignwalltime</code>' in the <i>Torque Resource Manager Administrator Guide</i> for more information.</p> </div>
<b>Example</b>	<pre>JOBMAXOVERRUN 15:00,1:00:00</pre> <p><i>Jobs may exceed their wallclock limit by up to 1 hour, but Moab will send an email to the primary administrator when a job exceeds its walltime by 15 minutes.</i></p>

JOBMAXPREEMPTCOUNT	
<b>Format</b>	<INTEGER>
<b>Default</b>	0 (No Limit)
<b>Description</b>	Maximum number of times a job can be preempted before it is no longer preemptible.
<b>Example</b>	<pre>JOBMAXPREEMPTCOUNT 5</pre> <p><i>Any job may be preempted up to 5 times, after which it is no longer preemptible.</i></p>

JOBMAXPREEMPTPERITERATION	
<b>Format</b>	<INTEGER>
<b>Default</b>	0 (No Limit)
<b>Description</b>	Maximum number of jobs allowed to be preempted per iteration.
<b>Example</b>	<code>JOBMAXPREEMPTPERITERATION 10</code>

JOBMAXSTARTPERITERATION	
<b>Format</b>	<INTEGER>
<b>Default</b>	0 (No Limit)
<b>Description</b>	Maximum number of jobs allowed to start per <a href="#">iteration</a> .
<b>Example</b>	<code>JOBMAXSTARTPERITERATION 10</code>

JOBMAXSTARTTIME	
<b>Format</b>	[ [ [DD:] HH:] MM:] SS
<b>Default</b>	-1 (NO LIMIT)
<b>Description</b>	length of time a job is allowed to remain in a 'starting' state. If a 'started' job does not transition to a running state within this amount of time, Moab will cancel the job, believing a system failure has occurred.
<b>Example</b>	<code>JOBMAXSTARTTIME 2:00:00</code> <i>Jobs may attempt to start for up to 2 hours before being canceled by the scheduler</i>

JOBMAXTASKCOUNT	
<b>Format</b>	<INTEGER>
<b>Default</b>	32768
<b>Description</b>	The total number of tasks allowed per job.
<b>Example</b>	<pre>JOBMAXTASKCOUNT 226000</pre>

JOBMIGRATEPOLICY	
<b>Format</b>	One of the following: IMMEDIATE, JUSTINTIME, or AUTO
<b>Default</b>	AUTO
<b>Description</b>	Upon using the <a href="#">msub</a> command to submit a job, you can allocate the job to immediately ( <b>IMMEDIATE</b> ) migrate to the resource manager, or you can instruct Moab to only migrate the job to the resource manager when it is ready to run ( <b>JUSTINTIME</b> ). Specifying <b>AUTO</b> allows MOAB to determine on a per-job basis whether to use <b>IMMEDIATE</b> or <b>JUSTINTIME</b> .
<b>Example</b>	<pre>JOBMIGRATEPOLICY JUSTINTIME</pre>

JOBNAMEWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	---
<b>Description</b>	The weight to be applied to the job's name if the Name contains an integer. See <a href="#">Attribute (ATTR) Factor</a> .
<b>Example</b>	<pre>JOBNAMEWEIGHT 1</pre>

JOBNODEMATCHPOLICY	
<b>Format</b>	AUTO, EXACTNODE, or EXACTPROC
<b>Default</b>	AUTO
<b>Description</b>	<p>Specifies additional constraints on how compute nodes are to be selected:</p> <ul style="list-style-type: none"> <li>• AUTO overrides the JOBNODEMATCHPOLICY (packs the jobs on any node).</li> <li>• EXACTNODE indicates that Moab should select as many nodes as requested even if it could pack multiple tasks onto the same node.</li> <li>• EXACTPROC indicates that Moab should select only nodes with exactly the number of processors configured as are requested per node even if nodes with excess processors are available.</li> </ul>
<b>Example</b>	<pre>JOBNODEMATCHPOLICY EXACTNODE</pre> <p><i>In a PBS/Native job with resource specification nodes=&lt;x&gt;:ppn=&lt;y&gt;, Moab will allocate exactly &lt;y&gt; task on each of &lt;x&gt; distinct nodes.</i></p>

JOBPREEMPTMAXACTIVETIME	
<b>Format</b>	[ [ [DD:]HH:]MM:]SS
<b>Default</b>	0
<b>Description</b>	The amount of time when a job can be eligible for preemption. See <a href="#">Job Preemption</a> .
<b>Example</b>	<pre>JOBPREEMPTMAXACTIVETIME 00:05:00</pre> <p><i>A job is preemptable for the first 5 minutes of its run time.</i></p>

JOBPREEMPTMINACTIVETIME	
<b>Format</b>	[ [ [DD:]HH:]MM:]SS
<b>Default</b>	0
<b>Description</b>	The minimum amount of time a job must be active before being considered eligible for preemption. See <a href="#">Job Preemption</a> .
<b>Example</b>	<pre>JOBPREEMPTMINACTIVETIME 00:05:00</pre> <p><i>A job must execute for 5 minutes before Moab will consider it eligible for preemption.</i></p>

<b>JOBPRIOACCRUALPOLICY</b>	
<b>Format</b>	ACCRUE or RESET
<b>Default</b>	ACCRUE
<b>Description</b>	<p>Specifies how Moab should track the dynamic aspects of a job's priority. ACCRUE indicates that the job will accrue queuetime based priority from the time it is submitted unless it violates any of the policies not specified in <a href="#">JOBPRIOEXCEPTIONS</a>. RESET indicates that it will accrue priority from the time it is submitted unless it violates any of the <a href="#">JOBPRIOEXCEPTIONS</a>. However, with RESET, if the job does violate <a href="#">JOBPRIOEXCEPTIONS</a> then its queuetime based priority will be reset to 0.</p> <div style="border: 1px solid #4a7c9c; border-radius: 10px; padding: 10px; margin: 10px 0;"> <p><b>i</b> JOBPRIOACCRUALPOLICY is a global parameter, but can be configured to work only in QOSCFG:</p> <pre style="border: 1px dashed #ccc; padding: 5px; display: inline-block;">QOSCFG[arrays] JOBPRIOACCRUALPOLICY=ACCRUE</pre> </div> <p>The following old JOBPRIOACCRUALPOLICY values have been <b>deprecated</b> and should be adjusted to the following values:</p> <ul style="list-style-type: none"> <li>• QUEUEPOLICY = ACCRUE and JOBPRIOEXCEPTIONS SOFTPOLICY, HARDPOLICY</li> <li>• QUEUEPOLICYRESET = RESET and JOBPRIOEXCEPTIONS SOFTPOLICY, HARDPOLICY</li> <li>• ALWAYS = ACCRUE and JOBPRIOEXCEPTIONS ALL</li> <li>• FULLPOLICY = ACCRUE and JOBPRIOEXCEPTIONS NONE</li> <li>• FULLPOLICYRESET = RESET and JOBPRIOEXCEPTIONS NONE</li> </ul>
<b>Example</b>	<pre style="border: 1px dashed #ccc; padding: 5px; display: inline-block;">JOBPRIOACCRUALPOLICY RESET</pre> <div style="border: 1px dashed #ccc; border-radius: 10px; padding: 10px; margin: 10px 0;"> <p><i>Moab will adjust the job's dynamic priority subcomponents (i.e., QUEUETIME, XFACTOR, and TARGETQUEUETIME, etc. each iteration that the job does not violate any JOBPRIOEXCEPTIONS, if it is found in violation, its queuetime will be reset to 0).</i></p> </div>



JOBPRIOEXCEPTIONS	
<b>Format</b>	Comma-delimited list of any of the following: DEFER, DEPENDS, SOFTPOLICY, HARDPOLICY, IDLEPOLICY, USERHOLD, BATCHHOLD, and SYSTEMHOLD (ALL or NONE can also be specified on their own)
<b>Default</b>	NONE
<b>Description</b>	<p>Specifies exceptions for calculating a job's dynamic priority (QUEUE TIME, XFACTOR, TARGETQUEUE TIME). Normally, when a job violates a policy, is placed on hold, or has an unsatisfied dependency, it will not accrue priority. Exceptions can be configured to allow a job to accrue priority in spite of any of these violations. With DEPENDS a job will increase in priority even if there exists an unsatisfied dependency. With SOFTPOLICY, HARDPOLICY, or IDLEPOLICY a job can accrue priority despite violating a specific limit. With DEFER, USERHOLD, BATCHHOLD, or SYSTEMHOLD a job can accrue priority despite being on hold.</p> <div style="border: 1px solid #0056b3; padding: 5px; margin-top: 10px;"> <p><b>i</b> JOBPRIOEXCEPTIONS is a global parameter, but can be configured to work only in QOSCFG:</p> <pre style="border: 1px dashed #ccc; padding: 5px; display: inline-block;">QOSCFG[arrays] JOBPRIOEXCEPTIONS=IDLEPOLICY</pre> </div>
<b>Example</b>	<pre style="border: 1px dashed #ccc; padding: 5px; display: inline-block;">JOBPRIOEXCEPTIONS BATCHHOLD,SYSTEMHOLD,DEPENDS</pre> <p style="border: 1px dashed #ccc; padding: 5px; display: inline-block; margin-top: 5px;"><i>Jobs will accrue priority in spite of batchholds, systemholds, or unsatisfied dependencies.</i></p>
JOBPRIOF	
<b>Format</b>	<ATTRIBUTE> [<VALUE>]=<PRIORITY> where <ATTRIBUTE> is one of ATTR, GRES or STATE
<b>Default</b>	---
<b>Description</b>	Specifies attribute priority weights for jobs with specific attributes, generic resource requests, or states. State values must be one of the standard Moab <a href="#">job states</a> . See <a href="#">Attribute-Based Job Prioritization</a> .
<b>Example</b>	<pre style="border: 1px dashed #ccc; padding: 5px; display: inline-block;">JOBPRIOF          STATE[Running]=100  STATE[Suspended]=1000  ATTR[PREEMPTEE]=200 GRES[biocalc]=5 ATTRATTRWEIGHT   1 ATTRSTATEWEIGHT  1</pre> <p style="border: 1px dashed #ccc; padding: 5px; display: inline-block; margin-top: 5px;"><i>Moab will adjust the job's dynamic priority subcomponents.</i></p>

JOBPURGETIME	
<b>Format</b>	[ [ [DD:] HH:] MM:] SS
<b>Default</b>	30
<b>Description</b>	The amount of time Moab will keep a job record that is no longer reported by the resource manager. Useful when using a resource manager that <i>drops</i> information about a job due to internal failures. See <a href="#">JOBPCPURGETIME</a> . Set to 0 to purge immediately if the resource manager does not report the job.
<b>Example</b>	<pre>JOBPURGETIME 00:05:00</pre> <p><i>Moab will maintain a job record for 5 minutes after the last update regarding that object received from the resource manager.</i></p>

JOBREJECTPOLICY	
<b>Format</b>	One or more of CANCEL, HOLD, IGNORE, MAIL, or RETRY
<b>Default</b>	HOLD
<b>Description</b>	The action to take when the scheduler determines that a job can never run. CANCEL issues a call to the resource manager to cancel the job. HOLD places a <i>batch</i> hold on the job preventing the job from being further evaluated until released by an administrator. <b>Note:</b> Administrators can dynamically alter job attributes and possibly <i>fix</i> the job with <a href="#">mjobctl -m</a> . With IGNORE, the scheduler will allow the job to exist within the resource manager queue but will neither process it nor report it. MAIL will send email to both the admin and the user when rejected jobs are detected. If RETRY is set, then Moab will allow the job to remain idle and will only attempt to start the job when the policy violation is resolved. Any combination of attributes can be specified. See <a href="#">QOSREJECTPOLICY</a> .
<b>Example</b>	<pre>JOBREJECTPOLICY MAIL, CANCEL</pre>

JOBREMOVEENVVARLIST	
<b>Format</b>	Comma-delimited list of strings
<b>Default</b>	---
<b>Description</b>	<p>Moab will remove the specified environment variables from the job's environment before migrating the job to its destination resource manager. This is useful when jobs submit themselves from one cluster to another with the full environment.</p> <div style="border: 1px solid #0070C0; border-radius: 10px; padding: 10px; background-color: #E6F2FF;"> <p><b>i</b> This parameter is currently only supported with Torque resource managers.</p> </div>
<b>Example</b>	<pre>JOBREMOVEENVVARLIST PBS_SERVER,TZ</pre> <p><i>Moab will remove the environment variables PBS_SERVER and TZ before submitting jobs.</i></p>

JOBRETRYTIME	
<b>Format</b>	[ [ [DD:]HH:]MM:]SS
<b>Default</b>	00:00:60
<b>Description</b>	<p>Period of time Moab will continue to attempt to start a job that has failed to start due to transient failures or that has successfully started and was then rejected by the resource manager due to transient failures. (For related information, see <a href="#">Reservation Policies</a>, <a href="#">DEFERSTARTCOUNT</a>, <a href="#">DEFERTIME</a>, <a href="#">RESERVATIONRETRYTIME</a>, <a href="#">NODEFAILURERESERVETIME</a>, and <a href="#">GUARANTEEDPREEMPTION</a>.)</p>
<b>Example</b>	<pre>JOBRETRYTIME 00:05:00</pre> <p><i>Moab will try for up to 5 minutes to restart jobs if the job start has failed due to transient errors.</i></p>

LIMITEDJOBBCP	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	TRUE
<b>Description</b>	By default, Moab does not update the checkpoint file or the cache for a job unless it has undergone a qualifying change (such as a change of state). Setting LIMITEDJOBBCP to FALSE will cause Moab to update the checkpoint file and the cache for all jobs each iteration, even without a qualifying change. For clusters that routinely run large numbers of jobs (e.g., more than 15,000), setting this parameter to FALSE might adversely affect iteration times.
<b>Example</b>	<pre>LIMITEDJOBBCP FALSE</pre> <p><i>Moab will update the checkpoint file and cache for all jobs each iteration.</i></p>

LIMITEDNODECP	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	Specifies whether there should be limited node checkpointing (see <a href="#">Consideration for Large Clusters</a> ).
<b>Example</b>	<pre>LIMITEDNODECP TRUE</pre> <p><i>Moab will only maintain scheduler checkpoint information for nodes with explicitly modified job attributes (some minor node performance and usage statistics may be lost).</i></p>

LOADALLJOB	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	Specifies whether Moab should load, during startup, all non-completed jobs in the checkpoint files regardless of whether or not their corresponding resource managers are active. For example, this allows source peers to continue showing remote jobs in the queue based on checkpointed info, even though the destination peer is offline.
<b>Example</b>	<pre>LOADALLJOB TRUE</pre> <p><i>Moab will load, at startup, all non-completed jobs from all checkpoint files.</i></p>

LOCKFILE	
<b>Format</b>	<STRING>
<b>Default</b>	---
<b>Description</b>	The path for the lock (pid) file used by Moab.
<b>Example</b>	<pre>LOCKFILE /var/spool/moab/lock</pre>

LOGDIR	
<b>Format</b>	<STRING>
<b>Default</b>	log
<b>Description</b>	The directory where log files will be maintained. If specified as a relative path, LOGDIR will be relative to \$ (MOABHOMEDIR). See <a href="#">Logging Overview</a> for more information.
<b>Example</b>	<pre>LOGDIR /var/spool/moab</pre> <p><i>Moab will record its log files directly into the /var/spool/moab directory</i></p>


LOGFACILITY	
<b>Format</b>	Colon delimited list of one or more of the following: CORE, SCHED, SOCK, UI, LL, CONFIG, STAT, SIM, STRUCT, FS, CKPT, BANK, RM, PBS, WIKI, ALL
<b>Default</b>	ALL
<b>Description</b>	Specifies which types of events to log (see <a href="#">Logging Overview</a> ).
<b>Example</b>	<pre>LOGFACILITY RM:PBS</pre> <p><i>Moab will log only events involving general resource manager or PBS interface activities.</i></p>

LOGFILE	
<b>Format</b>	<STRING>
<b>Default</b>	moab.log
<b>Description</b>	Name of the Moab log file. This file is maintained in the directory pointed to by <LOGDIR> unless <LOGFILE> is an absolute path (see <a href="#">Logging Overview</a> ).
<b>Example</b>	<pre>LOGFILE moab.test.log</pre> <p><i>Log information will be written to the file moab.test.log located in the directory pointed to by the LOGDIR parameter.</i></p>

LOGFILEMAXSIZE	
<b>Format</b>	<INTEGER>
<b>Default</b>	10000000
<b>Description</b>	Maximum allowed size (in bytes) of the log file before it will be rolled (see <a href="#">Logging Overview</a> ).
<b>Example</b>	<pre>LOGFILEMAXSIZE 50000000</pre> <p><i>Log files will be rolled when they reach 50 MB in size</i></p>

LOGFILEROLLDEPTH	
<b>Format</b>	<INTEGER>
<b>Default</b>	3
<b>Description</b>	Number of old log files to maintain (i.e., when full, <code>moab.log</code> will be renamed <code>moab.log.1</code> , <code>moab.log.1</code> will be renamed <code>moab.log.2</code> , ...). See <a href="#">Logging Overview</a> .
<b>Example</b>	<pre>LOGFILEROLLDEPTH 5</pre> <p><i>Moab will maintain and roll the last 5 log files.</i></p>

LOGLEVEL	
<b>Format</b>	<INTEGER> (0-9)
<b>Default</b>	0
<b>Description</b>	The verbosity of Moab logging where 9 is the most verbose ( <b>Note:</b> each logging level is approximately an order of magnitude more verbose than the previous level). See <a href="#">Logging Overview</a> .
<b>Example</b>	<pre>LOGLEVEL 4</pre> <p><i>Moab will write all Moab log messages with a threshold of 4 or lower to the <code>moab.log</code> file.</i></p>

LOGLEVELOVERRIDE	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	<p>When this parameter is on, if someone runs a command with <code>--loglevel=&lt;x&gt;</code>, that loglevel, if higher than the current loglevel, is used on the scheduler side for the duration of the command. All logs produced during that time are put into a separate log file (this creates a 'gap' in the normal logs). This can be very useful for debugging, but it is recommend that this be used only when diagnosing a specific problem so that users can't affect performance by submitting multiple <code>--loglevel</code> commands.</p> <div style="border: 1px solid #0070C0; border-radius: 10px; padding: 5px; margin-top: 10px;"> <p> This parameter does not work with threaded commands (such as <code>showq</code>, <code>mdiag -n</code>, and <code>mdiag -j</code>).</p> </div>
<b>Example</b>	LOGLEVELOVERRIDE TRUE

LOGPERMISSIONS	
<b>Format</b>	<INTEGER>
<b>Default</b>	644
<b>Description</b>	The octal number that represents read, write, and execute permissions.
<b>Example</b>	LOGPERMISSIONS 600
	<p><i>Allows the file owner to read and write permissions, but denies rights to the group and others.</i></p>



LOGROLLACTION	
<b>Format</b>	<STRING>
<b>Default</b>	---
<b>Description</b>	Specifies a script to run when the logs roll. The script is run as a trigger and can be viewed using <a href="#">mdiag -T</a> . For example, a script can be specified that always moves the first rolled log file, <code>moab.log.1</code> , to an archive directory for longer term storage.
<b>Example</b>	<pre>LOGROLLACTION /usr/local/tools/logroll.pl</pre>


MAILFROMADDR	
<b>Format</b>	<EMAILADDRESS>
<b>Default</b>	---
<b>Description</b>	Sets the FROM address for all emails sent from Moab. Used in conjunction with MAILPROGRAM.
<b>Example</b>	<pre>MAILFROMADDR it@yourdomain.com</pre>


MAILPROGRAM	
<b>Format</b>	[<Full_Path_To_Mail_Command>   DEFAULT   NONE ] [@<DEFAULTMAILDOMAIN>]
<b>Default</b>	NONE
<b>Description</b>	<p>If set to NONE, no mail is sent. If set to DEFAULT, Moab sends mail via the system's default mail program (usually <code>/usr/bin/sendmail</code>). If set to the local path of a mail program, Moab uses the specified mail program to send mail.</p> <p>By default, Moab mail notification is disabled. To enable, you must set MAILPROGRAM to DEFAULT or specify some other locally available mail program. If the <i>default mail domain</i> is set, emails will be routed to this domain unless a per-user domain is specified using the EMAILADDRESS attribute of the USERCFG parameter. If neither of these values is set, Moab uses '@localhost' as the mail domain. See <a href="#">Notify Admins</a>.</p> <p>For jobs, the email address used on the <code>msub -M</code> option overrides all other user email addresses. Additionally, administrators are notified in the case of job violations.</p>
<b>Example</b>	<pre>MAILPROGRAM DEFAULT</pre> <p><i>Moab sends mail via the system's default mail program, /usr/bin/sendmail.</i></p> <pre>MAILPROGRAM /usr/local/bin/sendmail@mydomain.com</pre> <p><i>Moab sends mail via the mail program located at /usr/local/bin/sendmail with default mail domain @mydomain.com</i></p>



MAXGRES	
<b>Format</b>	<INTEGER>
<b>Default</b>	512
<b>Description</b>	<p>Specifies how many generic resources Moab should manage.</p> <div style="border: 1px solid #0056b3; border-radius: 5px; padding: 5px; margin-bottom: 5px;"> <p><b>i</b> In Moab 9.0, four new internal generic resources were added to support NUMA. You might need to increase MAXGRES to accommodate the additional resources.</p> </div> <div style="border: 1px solid #0056b3; border-radius: 5px; padding: 5px;"> <p><b>i</b> Moab only reads in this setting when starting up (or restarting).</p> </div>
<b>Example</b>	<code>MAXGRES 1024</code>

MAXGMETRIC	
<b>Format</b>	<INTEGER>
<b>Default</b>	10
<b>Description</b>	<p>Specifies how many generic metrics Moab should manage.</p> <div style="border: 1px solid #0056b3; border-radius: 5px; padding: 5px;"> <p><b>i</b> Moab only reads in this setting when starting up (or restarting).</p> </div>
<b>Example</b>	<code>MAXGMETRIC 20</code>

MAXGREENSTANDBYPOOLSIZE	
<b>Format</b>	<INTEGER>
<b>Default</b>	-1
<b>Description</b>	Set the size of the standby pool. This is the number of idle nodes that will be powered on and idle. As the workload changes, Moab turns nodes on or off to try to meet this goal. A value of -1 disables the standby pool. See <a href="#">Enabling Green Computing</a> .
<b>Example</b>	<pre>MAXGREENSTANDBYPOOLSIZE 5</pre> <p><i>Moab keeps up to 5 idle nodes powered on to be kept on standby.</i></p>

MAXJOB	
<b>Format</b>	<INTEGER>
<b>Default</b>	51200
<b>Description</b>	<p>The maximum quantity of jobs for which Moab should allocate memory used for tracking jobs. If Moab is tracking the maximum quantity of jobs specified by this parameter, it rejects subsequent jobs submitted by any user since it has no memory left with which to track newly submitted jobs.</p> <p>If a user submitted a job with the <i>msub</i> command, this rejection behavior requires the user to resubmit the job at a later time after other jobs have completed, which frees memory in which Moab can place later-submitted jobs. If a user submitted a job with the Torque <i>qsub</i> command, Torque will automatically resubmit the job to Moab until Moab accepts it.</p> <p>The <a href="#">mdiag -S</a> command indicates if any job overflows have occurred.</p> <p>If this parameter's value is changed, it does not go into effect until Moab restarts. Moab reads the parameter only on initial startup and uses its value to allocate the memory it uses to track jobs.</p> <div style="border: 1px solid #0070C0; border-radius: 5px; padding: 5px; margin-top: 10px;">  Moab only reads in this setting when starting up (or restarting). </div>
<b>Example</b>	<pre>MAXJOB 75000</pre>

MAXNODE	
<b>Format</b>	<INTEGER>
<b>Default</b>	5120
<b>Description</b>	<p>The maximum number of compute nodes supported.</p> <p> Moab only reads in this setting when starting up (or restarting).</p>
<b>Example</b>	<pre>MAXNODE 10000</pre>

MAXRSVPERNODE	
<b>Format</b>	<INTEGER>
<b>Default</b>	64
<b>Description</b>	<p>The maximum number of reservations on a node.</p> <p>For large SMP systems (&gt;512 processors/node), Adaptive Computing advises adjusting the value to approximately twice the average sum of admin, standing, and job reservations present.</p> <p>A second number, led by a comma, can also be specified to set a maximum number of reservations for nodes that are part of the SHARED partition.</p> <p>The maximum possible value of MAXRSVPERNODE is 8192 for a global node and 4096 for any other node.</p> <p>Moab must be restarted for any changes to this parameter to take effect. The command <code>mdiag -S</code> indicates whether any node reservation overflows have occurred. See <a href="#">Considerations for Large Clusters</a>.</p> <div style="border: 1px solid #d9534f; background-color: #f8d7da; padding: 5px; margin: 5px 0;"> <p> Do not lower the MAXRSVPERNODE value while there are active jobs in the queue. This can lead to queue instability and certain jobs could become stuck or disconnected from the system.</p> </div> <div style="border: 1px solid #4a7c9c; background-color: #d9edf7; padding: 5px; margin: 5px 0;"> <p> Moab only reads in this setting when starting up (or restarting).</p> </div>
<b>Example</b>	<div style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 5px;"> <pre>MAXRSVPERNODE 64</pre> <p><i>64 is the maximum number of reservations on a single node.</i></p> </div> <div style="border: 1px dashed #ccc; padding: 5px;"> <pre>MAXRSVPERNODE 100,7000</pre> <p><i>100 is the maximum number of reservations on a single node, and 7000 is the maximum number of reservations for global nodes.</i></p> </div>


MEMREFRESHINTERVAL	
<b>Format</b>	[ [ [DD:]HH:]MM:] :SS   job:<COUNT>
<b>Default</b>	---
<b>Description</b>	The time interval or total job query count at which Moab will perform garbage collection to free memory associated with resource manager APIs that possess memory leaks (i.e., Loadleveler, etc.).
<b>Example</b>	<pre># free memory associated with leaky RM API MEMREFRESHINTERVAL 24:00:00</pre> <p><i>Moab will perform garbage collection once every 24 hours.</i></p>

MEMWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	The coefficient to be multiplied by a job's MEM (dedicated memory in MB) factor. See <a href="#">Resource Priority Overview</a> .
<b>Example</b>	<pre>RESWEIGHT 10 MEMWEIGHT 1000</pre> <p><i>Each job's priority will be increased by <math>10 * 1000 * \langle \text{request memory} \rangle</math>.</i></p>

MESSAGEQUEUEADDRESS	
<b>Format</b>	The IP address of the machine on which Moab is generating events.
<b>Default</b>	* (all)
<b>Description</b>	When a user subscribes to the events Moab provides and delivers via zeroMQ, s/he must do so by specifying <code>tcp://&lt;ipAddress&gt;:&lt;port&gt;</code> . MESSAGEQUEUEADDRESS specifies the <code>&lt;ipAddress&gt;</code> , which must match the IP address of the machine on which Moab is installed. To specify the port, see <a href="#">MESSAGEQUEUEPORT</a> .
<b>Example</b>	<pre>MESSAGEQUEUEADDRESS 10.1.0.10</pre> <p><i>To subscribe to Moab events, users must use <code>tcp://10.1.0.10:&lt;port&gt;</code>.</i></p>

MESSAGEQUEUEPORT	
<b>Format</b>	The port of the machine on which Moab is generating events.
<b>Default</b>	5563
<b>Description</b>	When a user subscribes to the events Moab provides and delivers via zeroMQ, s/he must do so by specifying <code>tcp://&lt;ipAddress&gt;:&lt;port&gt;</code> . MESSAGEQUEUEPORT specifies the <code>&lt;port&gt;</code> , which must match the port of the machine on which Moab is installed. To specify the IP address, see <a href="#">MESSAGEQUEUEADDRESS</a> .
<b>Example</b>	<pre>MESSAGEQUEUEPORT 1010</pre> <p><i>To subscribe to Moab events, users must use <code>tcp://&lt;ipAddress&gt;:1010</code>.</i></p>



MESSAGEQUEUESECRETKEY	
<b>Format</b>	<STRING>
<b>Default</b>	---
<b>Description</b>	<p>Causes Moab to encrypt the events delivered via zeroMQ using the Advanced Encryption Standard (AES) algorithm. Must be a <a href="#">Base64</a>-encoded, 128-bit (16-byte) key. Messages will be encrypted using AES in CBC mode where inputs are padded with PKCS5 padding. The initialization vector is calculated by using an MD5 hash of the key specified in MESSAGEQUEUESECRETKEY.</p> <div style="border: 1px solid #0070C0; border-radius: 5px; padding: 5px; margin-top: 10px;"> <p> MESSAGEQUEUESECRETKEY can only be specified in the moab-private.cfg file.</p> </div>
<b>Example</b>	<pre>MESSAGEQUEUESECRETKEY 1r6RvfqJa6voezy5wAx0hw==</pre>

MINADMINSTIME	
<b>Format</b>	<INTEGER>
<b>Default</b>	60 seconds
<b>Description</b>	The minimum time a job will be <a href="#">suspended</a> if suspended by an administrator or by a scheduler policy.
<b>Example</b>	<pre>MINADMINSTIME 00:10:00</pre> <div style="border: 1px dashed #000; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p><i>Each job suspended by administrators or policies will stay in the suspended state for at least 10 minutes.</i></p> </div>

MINPRIORITYJOBRSVSIZE	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	The minimum total job size, in processors, for a job to receive a priority reservation. Jobs smaller than this value will still be started during normal and backfill scheduling, but will not be eligible for priority reservations.
<b>Example</b>	<pre>MINPRIORITYJOBRSVSIZE 4</pre> <p><i>Any job requesting less than four processors will not receive a priority reservation.</i></p>

MISSINGDEPENDENCYACTION	
<b>Format</b>	CANCEL, HOLD, or RUN
<b>Default</b>	HOLD
<b>Description</b>	Controls what Moab does with a dependent job when its dependency job cannot be found when Moab evaluates the dependent job for scheduling. This only affects jobs whose dependent job cannot be found.
<b>Example</b>	<pre>MISSINGDEPENDENCYACTION CANCEL</pre> <p><i>Any job that has a dependent job that cannot be found is canceled.</i></p>

MONGOREPLICASETNAME	
<b>Format</b>	<name>
<b>Default</b>	---
<b>Description</b>	If MONGOSERVER is a comma separated list of HostAndPort strings (i.e., replicaset), MONGOREPLICASETNAME must be set to the name used when defining the replica set within MongoDB.
<b>Example</b>	<pre>rs.initiate(   {     _id: "myReplSet",     version: 1,     members: [       { _id: 0, host : "mongodb0.example.net:27017" },       { _id: 1, host : "mongodb1.example.net:27017" },       { _id: 2, host : "mongodb2.example.net:27017" }     ]   } )</pre> <pre>MONGOSERVER mongodb0.example.net:27017,mongodb1.example.net:27017,mongodb2.example.net:27017 MONGOREPLICASETNAME myReplSet MONGOSSLMODE enabled MONGOSSLCAFILE /etc/ssl/mongodb-cert.crt</pre>

MONGOSERVER	
<b>Format</b>	<server>[:<port>]
<b>Default</b>	---
<b>Description</b>	The MongoDB server DNS or IP and optional port number. The port number defaults to the MongoDB default (27017) when not given.
<b>Example</b>	<pre>MONGOSERVER localhost:27017</pre>

MONGOSSLCAFILE	
<b>Format</b>	<file>
<b>Default</b>	---
<b>Description</b>	<file> is a file containing the public cert located in the net.ssl.PEMKeyFile defined in the mongod.conf (i.e., /etc/mongod.conf) file on the MongoDB host. MONGOSSLCAFILE is ignored when MONGOSSLMODE is either undefined or set to disabled.
<b>Example</b>	<pre>MONGOSERVER localhost:27017 MONGOSSLMODE enabled MONGOSSLCAFILE /etc/ssl/mongodb-cert.crt</pre>

MONGOSSLMODE	
<b>Format</b>	enabled or disabled
<b>Default</b>	disabled
<b>Description</b>	Enables or disables encryption of MongoDB network traffic. MONGOSSLCAFILE is required when MONGOSSLMODE is set to 'enabled'. 'preferred' can be added if the Mongo driver supports it as an option.
<b>Example</b>	<pre>MONGOSERVER localhost:27017 MONGOSSLMODE enabled MONGOSSLCAFILE /etc/ssl/mongodb-cert.crt</pre>


MSUBQUERYINTERVAL	
<b>Format</b>	<INTEGER>
<b>Default</b>	5 seconds
<b>Description</b>	<p>The length of the interval (in seconds) between job queries when using <code>msub -K</code>. Jobs submitted with the <code>-K</code> option query the scheduler every <code>MSUBQUERYINTERVAL</code> seconds until the job is completed.</p> <p><code>MSUBQUERYINTERVAL</code> can exist as an environment variable. Any value in <code>moab.cfg</code> overrides the environment variable.</p> <p><b>Note:</b> If <code>MSUBQUERYINTERVAL</code> is set to 0, the <code>-K</code> option will be disabled. Jobs will still submit correctly, but the client will not continue to check on the job.</p>
<b>Example</b>	<pre>MSUBQUERYINTERVAL 60</pre> <p><i>If a user uses the <code>msub -K</code> command, the client remains open and queries the server every 60 seconds until the job completes.</i></p>

NODEACCESSPOLICY	
<b>Format</b>	One of the following: <a href="#">SHARED</a> , <a href="#">SHAREDONLY</a> , <a href="#">SINGLEACCOUNT</a> , <a href="#">SINGLECLASS</a> , <a href="#">SINGLEGROUP</a> , <a href="#">SINGLEJOB</a> , <a href="#">SINGLETASK</a> , <a href="#">SINGLEUSER</a> , or <a href="#">UNIQUEUSER</a>
<b>Default</b>	<b>SHARED</b>
<b>Description</b>	Specifies how node resources will be shared by various tasks (see the <a href="#">Node Access Overview</a> for more information).
<b>Example</b>	<pre>NODEACCESSPOLICY SINGLEUSER</pre> <p>Moab will allow resources on a node to be used by more than one job provided that the jobs are all owned by the same user.</p>

NODEAFFINITYPOLICY	
<b>Format</b>	POSITIVE or DEFAULT
<b>Default</b>	<b>DEFAULT</b>
<b>Description</b>	When multiple reservations are on the same node and a job has access to some with a positive affinity <i>and</i> to others with a negative affinity, then the last reservation's affinity wins (by default). When NODEAFFINITYPOLICY is set to POSITIVE <i>and</i> a job has any positive affinity on the node, then the positive affinity will have precedent over any negative affinity.
<b>Example</b>	<pre>NODEAFFINITYPOLICY POSITIVE</pre> <p>If a job has any positive affinity to a node, it will take precedent over any negative affinity.</p>

NODEALLOCATIONPOLICY	
<b>Format</b>	One of the following: <a href="#">CONTIGUOUS</a> , <a href="#">CPULOAD</a> , <a href="#">FIRSTAVAILABLE</a> , <a href="#">LASTAVAILABLE</a> , <a href="#">MAXBALANCE</a> , <a href="#">MINRESOURCE</a> , <a href="#">PRIORITY</a> , or <a href="#">PLUGIN</a>
<b>Default</b>	LASTAVAILABLE
<b>Description</b>	Specifies how Moab should allocate available resources to jobs. See <a href="#">Node Allocation Overview</a> for more information.  <div style="border: 1px solid #0070C0; border-radius: 5px; padding: 5px; background-color: #E6F2FF;"> <p><b>i</b> If <a href="#">ENABLEHIGHTHROUGHPUT</a> is TRUE, you must set <a href="#">NODEALLOCATIONPOLICY</a> to <a href="#">FIRSTAVAILABLE</a>.</p> </div>
<b>Example</b>	<pre>NODEALLOCATIONPOLICY MINRESOURCE</pre> <p><i>Moab will apply the node allocation policy MINRESOURCE to all jobs by default.</i></p>

NODEALLOCRESFAILUREPOLICY	
<b>Format</b>	One of the following: CANCEL, HOLD, IGNORE, MIGRATE, NOTIFY, or REQUEUE
<b>Default</b>	NONE
<b>Description</b>	Specifies how Moab should handle active jobs that experience node failures during execution. See the <a href="#">RESFAILPOLICY</a> resource manager extension or the <a href="#">Node Availability Overview</a> .
<b>Example</b>	<pre>NODEALLOCRESFAILUREPOLICY REQUEUE</pre> <p><i>Moab will requeue jobs that have allocated nodes fail during execution.</i></p>

NODEAVAILABILITYPOLICY	
<b>Format</b>	<p>&lt;POLICY&gt; [:&lt;RESOURCETYPE&gt;] ...</p> <p>where &lt;POLICY&gt; is one of COMBINED, DEDICATED, or UTILIZED and &lt;RESOURCETYPE&gt; is one of PROC, MEM, SWAP, or DISK</p>
<b>Default</b>	COMBINED
<b>Description</b>	<p>Specifies how available node resources are reported. Moab uses the following calculations to determine the amount of available resources:</p> <p>Dedicated (use what Moab has scheduled to be used):  Available = Configured - Dedicated</p> <p>Utilized (use what the resource manager is reporting is being used):  Available = Configured - Utilized</p> <p>Combined (use the larger of dedicated and utilized):  Available = Configured - (MAX (Dedicated, Utilized))</p> <p>Moab marks a node as busy when it has no available processors, so NODEAVAILABILITYPOLICY, by affecting how many processors are reported as available, also affects node state. See <a href="#">Node Availability Policies</a> for more information.</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p> Beginning with the 8.1.2 release, you can also set NODEAVAILABILITYPOLICY at NODECFG. See <a href="#">NODEAVAILABILITYPOLICY - page 593</a> for instructions on setting this at the local level.</p> </div>
<b>Example</b>	<div style="border: 1px dashed black; padding: 5px; margin-bottom: 5px;"> <p>NODEAVAILABILITYPOLICY DEDICATED:PROC COMBINED:MEM</p> </div> <div style="border: 1px dashed black; padding: 5px;"> <p><i>Moab will ignore resource utilization information in locating available processors for jobs but will use both dedicated and utilized memory information in determining memory availability.</i></p> </div>



NODEBUSYSTATEDELAYTIME	
<b>Format</b>	[ [ [DD:]HH:]MM:]SS
<b>Default</b>	0:01:00 (one minute)
<b>Description</b>	Length of time Moab will assume busy nodes will remain unavailable for scheduling if a system reservation is not explicitly created for the node.
<b>Example</b>	<pre>NODEBUSYSTATEDELAYTIME 0:30:00</pre> <p><i>Moab will assume busy nodes are not available for scheduling for at least 30 minutes from the current time. Therefore, these nodes will never be allocated to starting jobs. Also, these nodes will only be available for reservations starting more than 30 minutes in the future.</i></p>

NODECATCREDLIST	
<b>Format</b>	<p>&lt;LABEL&gt;=&lt;NODECAT&gt; [, &lt;NODECAT&gt;] ... [ &lt;LABEL&gt;=&lt;NODECAT&gt; [, &lt;NODECAT&gt;] ... ] ...</p> <p>where &lt;LABEL&gt; is any string and &lt;NODECAT&gt; is one of the defined node categories.</p>
<b>Default</b>	---
<b>Description</b>	If specified, Moab will generate node category groupings and each iteration will assign usage of matching resources to pseudo-credentials with a name matching the specified label. See the <a href="#">Node Categorization</a> section for more information.
<b>Example</b>	<pre>NODECATCREDLIST down=BatchFailure,HardwareFailure,NetworkFailure idle=Idle</pre> <p><i>Moab will create a down user, group, account, class, and QoS and will associate BatchFailure, HardwareFailure, and NetworkFailure resources with these credentials. Additionally, Moab will assign all Idle resources to matching idle credentials.</i></p>

NODECFG[X]	
<b>Format</b>	List of space delimited <ATTR>=<VALUE> pairs, where <ATTR> is one of the following:  ACCESS, ARCH, CHARGERATE, COMMENT, ENABLEPROFILING, FEATURES, GRES, MAXJOB, MAXJOBPERUSER, MAXLOAD, MAXPE, MAXPEPERJOB, MAXPROC, NODEAVAILABILITYPOLICY, NODEINDEX, NODETYPE, OS, OSLIST, PARTITION, POWERPOLICY, PREEMPTMAXCPULOAD, PREEMPTMINMEMAVAIL, PREEMTPOLICY, PRIORITY, PRIORITYF, PROCSPEED, PROVRM, RACK, RADISK, RCDISK, RCMEM, RCPROC, RCSWAP, SIZE, SLOT, SPEED, TRIGGER, VARIABLE
<b>Default</b>	---
<b>Description</b>	Specifies node-specific attributes for the node indicated in the array field. See the <a href="#">General Node Administration Overview</a> for more information.
<b>Example</b>	<pre>NODECFG[nodeA] MAXJOB=2 SPEED=1.2</pre> <p><i>Moab will only allow 2 simultaneous jobs to run on node nodeA and will assign a relative machine speed of 1.2 to this node.</i></p>

NODEDOWNSTATEDELAYTIME	
<b>Format</b>	[ [ [DD:]HH:]MM:]SS
<b>Default</b>	-1 (never)
<b>Description</b>	Length of time Moab will assume <a href="#">down</a> , <a href="#">drained</a> (offline), or corrupt nodes will remain unavailable for scheduling if a system reservation is not explicitly created for the node. The default specification of '-1' causes Moab to never create job reservations on down nodes. See <a href="#">Node Availability</a> for more information.
<b>Example</b>	<pre>NODEDOWNSTATEDELAYTIME 0:30:00</pre> <p><i>Moab will assume down, drained, and corrupt nodes are not available for scheduling for at least 30 minutes from the current time. Therefore, these nodes will never be allocated to starting jobs. Also, these nodes will only be available for reservations starting more than 30 minutes in the future.</i></p>


NODEDOWNTIME	
<b>Format</b>	[ [ [DD:] HH:] MM:] SS
<b>Default</b>	---
<b>Description</b>	The maximum time a previously reported node remains unreported by a resource manager before the node is considered to be in the <a href="#">down</a> state. This can happen if communication with a resource manager or a peer server is lost for more than the specified length of time, or if there is communication with the resource manager but it fails to report the node status.
<b>Example</b>	<pre>NODEDOWNTIME 10:00</pre> <p><i>If Moab loses communication with the resource manager for more than 10 minutes, it sets the state of all nodes belonging to that resource manager to DOWN.</i></p>

NODEDRAINSTATEDELAYTIME	
<b>Format</b>	[ [ [DD:] HH:] MM:] SS
<b>Default</b>	3:00:00 (three hours)
<b>Description</b>	Length of time Moab will assume <a href="#">drained</a> nodes will remain unavailable for scheduling if a system reservation is not explicitly created for the node. Specifying '-1' will cause Moab to never create job reservations on drained nodes. See <a href="#">Node Availability</a> for more information.
<b>Example</b>	<pre>NODEDRAINSTATEDELAYTIME 0:30:00</pre> <p><i>Moab will assume down, drained, and corrupt nodes are not available for scheduling for at least 30 minutes from the current time. Therefore, these nodes will never be allocated to starting jobs. Also, these nodes will only be available for reservations starting more than 30 minutes in the future.</i></p>

NODEFAILURERESERVETIME	
<b>Format</b>	[ [ [DD:]HH:]MM:]SS
<b>Default</b>	0:05:00
<b>Description</b>	Duration of reservation Moab will place on any node where it detects a failure from the resource manager (0 indicates no reservation will be placed on the node). See <a href="#">Node Availability</a> for more information. See also <a href="#">RMCFG[] NODEFAILURERSVPROFILE</a> . (For related information, see <a href="#">Reservation Policies</a> , <a href="#">DEFERSTARTCOUNT</a> , <a href="#">DEFERTIME</a> , <a href="#">RESERVATIONRETRYTIME</a> , <a href="#">JOBRETRYTIME</a> , and <a href="#">GUARANTEEDPREEMPTION</a> .)
<b>Example</b>	<pre>NODEFAILURERESERVETIME 10:00</pre> <p><i>Moab will reserve failed nodes for 10 minutes.</i></p>

NODEIDFORMAT	
<b>Format</b>	<STRING>
<b>Default</b>	*\$N*
<b>Description</b>	Specifies how a node ID can be processed to extract possible node, rack, slot, and cluster index information. The value of the parameter can include the markers \$C (cluster index), \$N (node index), \$R (rack index), or \$S (slot index) separated by *(asterisk - representing any number of non-numeric characters) or other characters to indicate this encoding. See <a href="#">Node Selection</a> for more information on use of node, rack, and slot indices.
<b>Example</b>	<pre>NODEIDFORMAT *\$R*\$S</pre> <p><i>Moab will extract rack and slot information from the cluster node IDs (i.e., tg-13s08).</i></p>

NODEIDLEPOWERTHRESHOLD	
<b>Format</b>	<INTEGER>
<b>Default</b>	60 seconds
<b>Description</b>	Specifies how long to allow a node to be idle before performing a power action. Increasing the idle duration prevents power on/off thrashing.
<b>Example</b>	<pre>NODEIDLEPOWERTHRESHOLD 300</pre> <p><i>Moab will wait 5 minutes before performing a power action on a node that has become idle.</i></p>

NODEIDLEPURGETIME	
<b>Format</b>	<SECONDS>
<b>Default</b>	0
<b>Description</b>	<p>When dynamic nodes are created in Moab, they are usually created with a request ID. Dynamic nodes created with a request ID are eligible to be scheduled for purging using NODEIDLEPURGETIME.</p> <p>NODEIDLEPURGETIME is the amount of time for all nodes with the same request ID to be idle before Moab begins firing the node end trigger for each iteration.</p> <p>If one or more of the nodes with the same request ID becomes non-idle, Moab stops firing the node end trigger for all of the nodes with the same request ID until the NODEIDLEPURGETIME is once again met.</p> <div style="border: 1px solid #0070C0; border-radius: 5px; padding: 5px; margin-top: 10px;">  A value of 0 disables this feature.         </div>
<b>Example</b>	<pre>NODEIDLEPURGETIME 300</pre> <p><i>Moab will begin purging groups of dynamic nodes with the same request ID when all nodes with the same request ID have been idle for 300 seconds.</i></p> <p><i>Here is an example of how to create a dynamic node with a request ID:</i></p> <pre>qmgr -c "create node elastic_node01 np=16,TTL=2022-6-16T17:17:8Z,requestid=unique_identifierXYZ"</pre>

NODEMAXLOAD	
<b>Format</b>	<DOUBLE>
<b>Default</b>	0.0
<b>Description</b>	<p>The maximum CPU load on an idle or running node. If the node's load reaches or exceeds this value, Moab will mark the node <b>busy</b>.</p> <p>You can also set the MAXLOAD at NODECFG. However, setting NODECFG MAXLOAD to -1 unsets this parameter setting. See <a href="#">10.4.1.D MAXLOAD - page 600</a> for instructions on setting this at the local level.</p>
<b>Example</b>	<pre>NODEMAXLOAD 0.75</pre> <p><i>Moab will adjust the state of all idle and running nodes with a load &gt;= .75 to the state busy.</i></p>

NODEMEMOVERCOMMITFACTOR	
<b>Format</b>	<DOUBLE>
<b>Default</b>	---
<b>Description</b>	<p>The parameter overcommits available and configured memory and swap on a node by the specified factor (for example: mem/swap * factor). Used to show that the node has more mem and swap than it really does. Only works for PBS RM types.</p>
<b>Example</b>	<pre>NODEMEMOVERCOMMITFACTOR .5</pre> <p><i>Moab will overcommit the memory and swap of the node by a factor of 0.5.</i></p>

NODESETATTRIBUTE	
<b>Format</b>	FEATURE or VARATTR
<b>Default</b>	---
<b>Description</b>	The type of node attribute by which node set boundaries will be established. See <a href="#">Node Set Overview</a> .
<b>Example</b>	<pre>NODESETPOLICY    ONEOF NODESETATTRIBUTE FEATURE</pre> <p><i>Moab will create node sets containing nodes with common features.</i></p>

NODESETDELAY	
<b>Format</b>	[ [ [DD:]HH:]MM:]SS
<b>Default</b>	0:00:00
<b>Description</b>	<p>Causes Moab to attempt to span a job evenly across nodesets unless doing so delays the job beyond the requested NODESETDELAY.</p> <p><b>i</b> Must use with <a href="#">NODESETPLUS</a> set to SPANEVENLY; if you do not want to use SPANEVENLY, use <a href="#">NODESETISOPTIONAL</a> instead of NODESETDELAY.</p>
<b>Example</b>	<pre>NODESETPLUS SPANEVENLY NODESETDELAY    5:00</pre> <p><i>Moab tries to span the job evenly across nodesets unless doing so delays the job by 5 minutes.</i></p>

NODESETISOPTIONAL	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	TRUE
<b>Description</b>	Specifies whether or not Moab will start a job if a requested node set cannot be satisfied. See <a href="#">Node Set Overview</a> .
<b>Example</b>	<pre>NODESETISOPTIONAL TRUE</pre> <p><i>Moab will not block a job from running if its node set cannot be satisfied.</i></p>

NODESETLIST	
<b>Format</b>	<ATTR>[{ :,  }<ATTR>] ...
<b>Default</b>	---
<b>Description</b>	The list of node attribute values that will be considered for establishing node sets. See <a href="#">Node Set Overview</a> .
<b>Example</b>	<pre>NODESETPOLICY      ONEOF NODESETATTRIBUTE   FEATURE NODESETLIST        switchA,switchB</pre> <p><i>Moab will allocate nodes to jobs either using only nodes with the switchA feature or using only nodes with the switchB feature.</i></p>



NODESETPLUS	
<b>Format</b>	DELAY or SPANEVENLY
<b>Default</b>	---
<b>Description</b>	<p>Specifies how Moab distributes jobs among nodesets. See <a href="#">Node Set Overview</a>.</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> Neither SPANEVENLY nor DELAY values of the <b>NODESETPLUS</b> parameter will work with multi-req jobs or preemption.</p> </div>
<b>Example</b>	<div style="border: 1px dashed gray; padding: 5px; margin-bottom: 10px;"> <p>NODESETPLUS SPANEVENLY</p> <p><i>Moab attempts to fit all jobs on a single nodeset or to span them evenly across a number of nodesets, unless doing so would delay a job beyond the requested <b>NODESETDELAY</b>.</i></p> </div> <div style="border: 1px dashed gray; padding: 5px;"> <p>NODESETPLUS DELAY</p> <p><i>Moab attempts to schedule the job within a nodeset for the configured <b>NODESETDELAY</b>. If Moab cannot find space for the job to start within <b>NODESETDELAY</b> (Moab considers future workload to determine if space will open up in time and might create a future reservation), then Moab schedules the job and ignores the nodeset requirement.</i></p> </div>
NODESETPOLICY	
<b>Format</b>	ANYOF, FIRSTOF, or ONEOF
<b>Default</b>	---
<b>Description</b>	<p>Specifies how nodes will be allocated to the job from the various node set generated. See <a href="#">Node Set Overview</a>.</p>
<b>Example</b>	<div style="border: 1px dashed gray; padding: 5px; margin-bottom: 10px;"> <p>NODESETPOLICY ONEOF          NODESETATTRIBUTE NETWORK</p> </div> <div style="border: 1px dashed gray; padding: 5px;"> <p><i>Moab will create node sets containing nodes with common network interfaces.</i></p> </div>

NODESETPRIORITYTYPE	
<b>Format</b>	One of AFFINITY, BESTFIT, FIRSTFIT, WORSTFIT, or MINLOSS
<b>Default</b>	FIRSTFIT
<b>Description</b>	Specifies how resource sets will be selected when more than one feasible resource can be found. See <a href="#">Node Set Overview</a> .
<b>Example</b>	<pre>NODESETPRIORITYTYPE BESTFIT NODESETATTRIBUTE    PROCSPEED</pre> <p><i>Moab will select the resource set that most closely matches the set of resources requested.</i></p>

NODETOJOBATTRMAP	
<b>Format</b>	Comma-delimited list of node features
<b>Default</b>	---
<b>Description</b>	<p>Job requesting the listed node features will be assigned a corresponding job attribute. These job attributes can be used to enable <a href="#">reservation access</a>, adjust <a href="#">job priority</a> or enable other capabilities.</p> <p><b>i</b> Feature/attribute matching is case-sensitive. In particular, keep in mind that PREEMPTEE and INTERACTIVE require uppercase.</p>
<b>Example</b>	<pre>NODETOJOBATTRMAP fast,big</pre> <p><i>Jobs requesting node feature <code>fast</code> or <code>big</code> (for instance, with <b>-l feature=fast</b> or <b>-W x=feature:big</b>) will be assigned a corresponding job attribute.</i></p>

NODEUNTRACKEDRESDELAYTIME	
<b>Format</b>	[ [ [DD:] HH:] MM:] SS
<b>Default</b>	0:00:00
<b>Description</b>	<p>Length of time Moab will assume untracked generic resources will remain unavailable for scheduling if a system reservation is not explicitly created for the node.</p> <p>If NODEUNTRACKEDRESDELAYTIME is enabled and there is an untracked resource preventing a job from running, then the job remains in the idle queue instead of being deferred.</p>
<b>Example</b>	<pre>NODEUNTRACKEDRESDELAYTIME 0:30:00</pre> <p><i>Moab will assume untracked generic resources are not available for scheduling for at least 30 minutes from the current time. Therefore, these nodes will never be allocated to starting jobs. Also, these nodes will only be available for reservations starting more than 30 minutes in the future.</i></p>

NODEWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	<p>The weight that will be applied to a job's requested node count before this value is added to the job's cumulative priority. <b>Notes:</b> This weight currently only applies when a nodecount is specified by the user job; if the job only specifies tasks or processors, no node factor will be applied to the job's total priority. In order for NODEWEIGHT to function properly, JOBNODEMATCHPOLICY should be set to EXACTNODE. These will be rectified in future versions.</p>
<b>Example</b>	<pre>NODEWEIGHT 1000</pre>

NOLOCALUSERENV	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	If TRUE, specifies that a user's UserID, GroupID, and HomeDirectory are available on the Moab server host.
<b>Example</b>	<pre>NOLOCALUSERENV TRUE</pre>

NOJOBHOLDNORESOURCES	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	If TRUE, Moab does not place a hold on jobs that don't have feasible resources. For example, suppose there are 20 processors available for ClassA and 50 processors for the entire system. If a job requests 21 or more processors from ClassA, or 51 or more processors from the entire system, Moab idles the job (instead of putting a hold on it) until the resources become available.
<b>Example</b>	<pre>NOJOBHOLDNORESOURCES TRUE</pre>

NOTIFICATIONPROGRAM	
<b>Format</b>	<STRING>
<b>Default</b>	---
<b>Description</b>	The name of the program to handle all notification call-outs.
<b>Example</b>	<pre>NOTIFICATIONPROGRAM tools/notifyme.pl</pre>

NOWAITPREEMPTION	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	---
<b>Description</b>	Generally when a job is trying to preempt another, it just waits for the original jobs that it chose to preempt to end. If this parameter is on, the preemptor will continue trying to preempt jobs until it can get in.
<b>Example</b>	<pre>NOWAITPREEMPTION  TRUE</pre>

OSCREDLOOKUP	
<b>Format</b>	BESTEFFORT or NEVER
<b>Default</b>	BESTEFFORT
<b>Description</b>	<p>When set to <code>BESTEFFORT</code>, Moab will create a 'moab user' and try to associate it what a 'system user' (if possible).</p> <p>When set to <code>NEVER</code>, this disables all Moab OS credential lookups, including UID, GID, user to group mappings, and any other OS specific information.</p> <p>Setting <code>OSCREDLOOKUP</code> by itself does not allow job submission; additional configuration is required. When submitting jobs from user accounts that do not exist on the head node (where Moab and Torque are running), you must also set the <a href="#">PROXYJOBSUBMISSION</a> flag in addition to specifying configuration settings in the resource manager configuration file. See the example that follows for information on required resource manager settings.</p>
<b>Example</b>	<pre>OSCREDLOOKUP NEVER RMCFG[] FLAGS=PROXYJOBSUBMISSION</pre> <p>To allow job submission, in the Torque configuration file (torque.cfg):</p> <pre>VALIDATEPATH FALSE</pre> <p>Run the following qmgr directive:</p> <pre>set server disable_server_id_check = True</pre> <p>Restart both Moab and pbs_server.</p>

PARALLOCATIONPOLICY	
<b>Format</b>	One of BestFit, BestFitP, FirstStart, LoadBalance, LoadBalanceP, Random, or RoundRobin
<b>Default</b>	FirstStart
<b>Description</b>	The approach to use to allocate resources when more than one eligible partition can be found. See <a href="#">Grid Scheduling Policies</a> for more information.
<b>Example</b>	<pre>PARALLOCATIONPOLICY LOADBALANCE</pre> <p><i>New jobs will be started on the most lightly allocated partition.</i></p>

PARCFG	
<b>Format</b>	<a href="#">NODEPOWEROFFDURATION</a> , <a href="#">NODEPOWERONDURATION</a> , <a href="#">NODEALLOCATIONPOLICY</a> or one or more key-value pairs as described in the <a href="#">Partition Overview</a>
<b>Default</b>	---
<b>Description</b>	The attributes, policies, and constraints for the given partition.
<b>Example</b>	<pre>PARCFG[oldcluster] MAX.WCLIMIT=12:00:00</pre> <p><i>Moab will not allow jobs to run on the oldcluster partition, which has a wallclock limit in excess of 12 hours.</i></p>

PBSACCOUNTINGDIR	
<b>Format</b>	<PATH>
<b>Default</b>	---
<b>Description</b>	When specified, Moab will write out job events in standard PBS/ Torque trace-job format to the specified directory using the standard PBS/TORQUE log file naming convention. See 'Using tracejob to Locate Job Failures' in the <i>Torque Administrator Guide</i> for more information.
<b>Example</b>	<pre>PBSACCOUNTINGDIR /var/spool/torque/sched_logs/</pre> <p><i>Job events will be written to the specified directory (can be consumed by PBS's tracejob command).</i></p>

PERPARTITIONSCHEDULING	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	<p>By default, Moab's scheduling routine schedules each job on each partition using the following algorithm:</p> <pre style="margin-left: 20px;"> <i>prioritize</i> <i>foreach (job)</i>     <i>find the partition on which that job should run</i>     <i>schedule job</i>                     </pre> <p>In this model, a job's priority is the same on each partition as it uses a single global priority. Because a job's priority is the same on every partition, Moab prioritizes the queue once and then schedules the prioritized queue across all partitions.</p> <p>When PERPARTITIONSCHEDULING TRUE is set, the following algorithm is used:</p> <pre style="margin-left: 20px;"> <i>foreach (partition)</i>     <i>prioritize</i>     <i>foreach (job)</i>         <i>schedule job</i>                     </pre> <p>In this case, each partition can have a unique priority configuration and Moab will re-prioritize the jobs for each partition on the system. Each job is prioritized and scheduled on each partition. See <a href="#">PARCFG</a> for more information. Also, note that Moab will order the partitions as they are discovered in the moab.cfg file. Partitions should be explicitly ordered via PARCFG in the moab.cfg file.</p>
<b>Example</b>	<div style="border: 1px dashed gray; padding: 10px; margin-bottom: 10px;"> <pre style="margin: 0;"> PERPARTITIONSCHEDULING TRUE PARCFG[p1] CONFIGFILE=/opt/moab/etc/p1.cfg PARCFG[p2] CONFIGFILE=/opt/moab/etc/p2.cfg                     </pre> </div> <div style="border: 1px dashed gray; padding: 10px; margin-bottom: 10px;"> <p><i>Rather than prioritizing the job queue once, Moab prioritizes the job queue for each partition, p1 and p2 respectively, and schedules each partition in turn using the policies located in their respective configuration files. See <a href="#">Per-Partition Settings</a> for more information.</i></p> </div>



PEWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	The coefficient to be multiplied by a job's PE (processor equivalent) priority factor.
<b>Example</b>	<pre>RESWEIGHT 10 PEWEIGHT 100</pre> <p><i>Each job's priority will be increased by <math>10 * 100 * \text{its PE factor}</math>.</i></p>

PREEMPTIONALGORITHM	
<b>Format</b>	PREEMPTORCENTRIC or PREEMPTTEECENTRIC
<b>Default</b>	PREEMPTTEECENTRIC
<b>Description</b>	PREEMPTTEECENTRIC specifies Moab will use a custom scheduling policy that ignores many policies such as JOBNODEMATCHPOLICY, NODEALLOCATIONPOLICY, NODEACCESSPOLICY. This custom scheduling policy will ensure the fewest and least important (by priority) preemptees are disturbed by the preemptor. PREEMPTORCENTRIC specifies Moab uses the normal scheduling policy and obeys all configured policies.
<b>Example</b>	<pre>PREEMPTIONALGORITHM PREEMPTORCENTRIC</pre> <p><i>Moab schedules the jobs as if the preemptees were not active and results in optimal placement for the preemptor.</i></p>

PREEMTPOLICY	
<b>Format</b>	One of the following: <b>CANCEL</b> , <b>REQUEUE</b> , <b>SUSPEND</b> , or <b>CHECKPOINT</b>
<b>Default</b>	REQUEUE
<b>Description</b>	<p>Specifies how preemptable jobs will be <b>preempted</b>:</p> <ul style="list-style-type: none"> <li>• If this policy is set to <b>REQUEUE</b>, preemptible jobs should be marked as <b>RESTARTABLE</b>.</li> <li>• If this policy is set to <b>SUSPEND</b>, preemptible jobs should be marked as <b>SUSPENDABLE</b>.</li> </ul> <div style="border: 1px solid #004a87; padding: 5px; margin-top: 10px;"> <p><b>i</b> Moab uses preemption escalation to preempt resources if the specified preemption facility is not applicable. This means if the policy is set to <b>SUSPEND</b> and the job is not <b>SUSPENDABLE</b>, Moab might attempt to requeue or even cancel the job.</p> </div>
<b>Example</b>	<pre>PREEMTPOLICY CHECKPOINT</pre> <p><i>Jobs that are to be preempted will be checkpointed and restarted at a later time.</i></p>

PREEMTPRIOJOBSELECTWEIGHT	
<b>Format</b>	<DOUBLE>
<b>Default</b>	256.0
<b>Description</b>	<p>Determines which jobs to preempt based on size or priority. The higher the value, the more emphasis is placed on the priority of the job, causing the lower priority jobs to be preempted first. The lower the value, the more emphasis is placed on the size of the job, causing the smaller jobs to be preempted first. If set to 0, job priority will be ignored, job size will take precedence and the smallest jobs will be preempted.</p> <p>The special setting of -1 places the emphasis solely on resource utilization. This means that jobs will be preempted in a manner that keeps the resource utilization at the highest level, regardless of job priority or size.</p>
<b>Example</b>	<pre>PREEMTPRIOJOBSELECTWEIGHT 220.5</pre>

PREEMPTRTIMEWEIGHT	
<b>Format</b>	<DOUBLE>
<b>Default</b>	0
<b>Description</b>	If set to anything other than 0, a job's remaining time is added into the calculation of which jobs will be preempted. If a positive weight is specified, jobs with a longer remaining time are favored. If a negative weight is specified, jobs with a shorter remaining time are favored.
<b>Example</b>	<pre>PREEMPTRTIMEWEIGHT 1.5</pre>

PREEMPTSEARCHDEPTH	
<b>Format</b>	<INTEGER>
<b>Default</b>	unlimited
<b>Description</b>	Specifies how many preemptible jobs will be evaluated as potential targets for serial job preemptors. See <a href="#">Preemption Overview</a> for more information.
<b>Example</b>	<pre>PREEMPTSEARCHDEPTH 8</pre> <i>Serial job preemptors will only consider the first 8 feasible preemptee jobs when determining the best action to take.</i>

PRIORITYTARGETDURATION	
<b>Format</b>	[ [ [DD:]HH:]MM:]SS
<b>Default</b>	---
<b>Description</b>	The <i>ideal</i> job duration that will maximize the value of the <a href="#">WALLTIMEWEIGHT</a> priority factor. If specified, this factor will be calculated as the distance from the ideal. Consequently, in most cases, the associated subcomponent weight should be set to a negative value.
<b>Example</b>	<pre>WALLTIMEWEIGHT -2500 PRIORITYTARGETDURATION 1:00:00</pre>

PRIORITYTARGETPROCCOUNT	
<b>Format</b>	<INTEGER>{+ - %}
<b>Default</b>	---
<b>Description</b>	The ideal job requested proc count that will maximize the value of the <a href="#">PROCWEIGHT</a> priority factor. If specified, this factor will be calculated as the distance from the ideal (proc count - ideal = coefficient of PROCWEIGHT). Consequently, in most cases, the associated subcomponent weight should be set to a negative value.
<b>Example</b>	<pre>PROCWEIGHT -1000 PRIORITYTARGETPROCCOUNT 64</pre>

PROCWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	The coefficient to be multiplied by a job's requested processor count priority factor.
<b>Example</b>	<pre>PROCWEIGHT 2500</pre>

PROFILECOUNT	
<b>Format</b>	<INTEGER>
<b>Default</b>	600
<b>Description</b>	<p>The number of statistical profiles to maintain.</p> <p>PROFILECOUNT must be set high enough that at least one day of statistics is maintained. The statistics time window can be determined by measuring PROFILEDURATION * PROFILECOUNT. If PROFILEDURATION is one hour then PROFILECOUNT must be at least 24 so 24 hours worth of statistics are maintained. If PROFILEDURATION is 30:00 then PROFILECOUNT must be set to at least 48. If PROFILECOUNT is not high enough for at least one day of statistics, Moab adjusts it automatically.</p>
<b>Example</b>	<code>PROFILECOUNT 300</code>

PROFILEDURATION	
<b>Format</b>	[ [ [DD:]HH:]MM:]SS
<b>Default</b>	00:30:00
<b>Description</b>	<p>The duration of each statistical profile. The duration cannot be more than 24 hours, and any specified duration must be a factor of 24. For example, factors of 1/4, 1/2, 1, 2, 3, 4, 6, 8, 12, and 24 are acceptable durations.</p>
<b>Example</b>	<code>PROFILEDURATION 24:00:00</code>

PURGETIME	
<b>Format</b>	[ [ [DD:] HH:] MM:] SS
<b>Default</b>	0
<b>Description</b>	The amount of time Moab will keep a job or node record for an object no longer reported by the resource manager. Useful when using a resource manager that 'drops' information about a node or job due to internal failures. <b>Note:</b> This parameter is superseded by <a href="#">JOBPURGETIME</a> .
<b>Example</b>	<pre>PURGETIME 00:05:00</pre> <p><i>Moab will maintain a job or node record for 5 minutes after the last update regarding that object received from the resource manager.</i></p>

PUSHCACHETOWEBSERVICE	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	Specifies whether or not you want to send cache objects (nodes, jobs, services, etc.) to Moab Web Services.
<b>Example</b>	<pre>PUSHCACHETOWEBSERVICE TRUE</pre>

QOSCFG[<QOSID>]	
<b>Format</b>	List of zero or more space delimited <ATTR>=<VALUE> pairs, where <ATTR> is one of the following: <a href="#">General Credential Flags</a> , ACLBLTHRESHOLD, ACLQTTTHRESHOLD, ACLXFTHRESHOLD, <a href="#">ENABLEPROFILING</a> , FSTARGET, <a href="#">JOBPRIOACCRUALPOLICY</a> , <a href="#">JOBPRIOEXCEPTIONS</a> , <a href="#">MEMBERULIST</a> , <a href="#">PLIST</a> , <a href="#">PREEMPTTEES</a> , <a href="#">PREEMPTMAXTIME</a> , <a href="#">PREEMPTMINTIME</a> , <a href="#">PREEMPTQTTTHRESHOLD</a> , <a href="#">PREEMPTXFTHRESHOLD</a> , <a href="#">PRIORITY</a> , <a href="#">QFLAGS</a> , <a href="#">QTTARGET</a> , <a href="#">QTWEIGHT</a> , <a href="#">REQRID</a> , <a href="#">RSVQTTTHRESHOLD</a> , <a href="#">RSVXFTHRESHOLD</a> , <a href="#">XFTARGET</a> , <a href="#">XFWEIGHT</a> , <a href="#">usage limit</a> .
<b>Default</b>	---
<b>Description</b>	Specifies QOS specific attributes. See the <a href="#">flag overview</a> for a description of legal flag values. See the <a href="#">QOS Overview</a> section for further details.
<b>Example</b>	<pre>QOSCFG[commercial] PRIORITY=1000 MAXJOB=4 MAXPROC=80</pre> <p><i>Moab will increase the priority of jobs using QOS commercial, and will allow up to 4 simultaneous QOS commercial jobs with up to 80 total allocated processors.</i></p>

QOSDEFAULTORDER	
<b>Format</b>	Comma-delimited list of QOS names.
<b>Default</b>	---
<b>Description</b>	Sets a global QOS default order for all QOSs, which overrides any specific default QOS. If the order is defined as b, a, c and a user has access to c, a and submits a job without requesting a specific QOS, the job is assigned a as the default QOS.
<b>Example</b>	<pre>QOSDEFAULTORDER b,a,c</pre> <p><i>If the job does not have a QOS specified, it is assigned a QOS from the QOSDEFAULTORDER list (if the user has access to one of them).</i></p>

QOSISOPTIONAL	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	An entity's default QoS will be the first QoS specified in the QLIST parameter. When this parameter is set to TRUE the default QoS for the associated credential (user, account, class, etc.) will not be automatically set to the first QoS specified in the QLIST.
<b>Example</b>	<pre>QOSISOPTIONAL TRUE USERCFG[bob] QLIST=high,low</pre> <p><i>Moab will set the QOSList for user bob to high and low but will not set the QDEF. Should bob decide to submit to a particular QOS he will have to do so manually.</i></p>

QOSREJECTPOLICY	
<b>Format</b>	One or more of CANCEL, HOLD, IGNORE, or MAIL
<b>Default</b>	HOLD
<b>Description</b>	<p>The action to take when Moab determines that a job cannot access a requested QoS. CANCEL issues a call to the resource manager to cancel the job. HOLD places a <i>batch</i> hold on the job preventing the job from being further evaluated until released by an administrator. (<b>Note:</b> Administrators can dynamically alter job attributes and possibly <i>fix</i> the job with <code>mjobctl -m</code>.) With IGNORE, Moab will ignore the QoS request and schedule the job using the default QoS for that job. MAIL will send email to both the admin and the user when QoS request violations are detected. Most combinations of attributes can be specified; however, if both MAIL and IGNORE are specified, Moab will not implement MAIL. Similarly, while CANCEL and HOLD are mutually exclusive, CANCEL will supersede HOLD if both are specified.</p> <p>Also see <a href="#">JOBREJECTPOLICY</a>.</p>
<b>Example</b>	<pre>QOSREJECTPOLICY MAIL, CANCEL</pre>




QOSWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	1
<b>Description</b>	The weight to be applied to the qos priority of each job (see <a href="#">Credential (CRED) Factor</a> ).
<b>Example</b>	<pre>QOSWEIGHT 10</pre>

QUEUETIMECAP	
<b>Format</b>	<DOUBLE>
<b>Default</b>	0 (NO CAP)
<b>Description</b>	The maximum allowed absolute preweighted queue time priority factor.
<b>Example</b>	<pre>QUEUETIMECAP 10000 QUEUETIMEWEIGHT 10</pre> <p><i>A job that has been queued for 40 minutes will have its queue time priority factor calculated as 'Priority = QUEUETIMEWEIGHT * MIN (10000,40)'.</i></p>

QUEUETIMEWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	1
<b>Description</b>	Specifies multiplier applied to a job's queue time (in minutes) to determine the job's queue time priority factor.
<b>Example</b>	<pre>QUEUETIMEWEIGHT 20</pre> <p><i>A job that has been queued for 4:20:00 will have a queue time priority factor of 20 * 260.</i></p>

REALTIMEDBOBJECTS	
<b>Format</b>	Comma-delimited list of one or more of the following: JOB, NODE, RSV (reservation), TRIG (trigger), VC (virtual container). You can also specify ALL or NONE.
<b>Default</b>	ALL
<b>Description</b>	Specifies which objects Moab will store in the unixodbc database.
<b>Example</b>	<pre>REALTIMEDBOBJECTS JOB,RSV,TRIG</pre> <p><i>Moab stores jobs, reservations, and triggers in the uxodbc database. It will no longer record real time information about nodes and VCs.</i></p>

RECORDEVENTLIST	
<b>Format</b>	One or more comma (',') or plus ('+') separated events of <a href="#">GEVENT</a> , ALLSCHEDCOMMAND, AMCREATE, AMDELETE, AMEND, AMPAUSE, AMQUOTE, AMRESUME, AMSTART, AMUPDATE, JOBCANCEL, JOBCHECKPOINT, JOBEND, JOBFAILURE, JOBHOLD, JOBMIGRATE, JOBMODIFY, JOBPREEMPT, JOBREJECT, JOBRESUME, JOBSTART, JOBSUBMIT, JOBVARSET, JOBVARUNSET, NODEADD, NODEDELETE, NODEDESTROY, NODEDOWN, NODEFAILURE, NODEMODIFY, NODEPOWEROFF, NODEPOWERON, NODEPROVISION, NODEUP, NOTE, QOSVIOLATION, RMDOWN, RMPOLLEND, RMPOLLSTART, RMUP, RSVCANCEL, RSVCREATE, RSVEND, RSVMODIFY, RSVSTART, SCHEDCOMMAND, SCHEDCYCLEEND, SCHEDCYCLESTART, SCHEDFAILURE, SCHEDMODIFY, SCHEDPAUSE, SCHEDRECYCLE, SCHEDRESUME, SCHEDSTART, SCHEDSTOP, TRIGEND, TRIGFAILURE, TRIGSTART, TRIDTHRESHOLD, or ALL
<b>Default</b>	JOBSTART, JOBCANCEL, JOBEND, JOBFAILURE, SCHEDPAUSE, SCHEDSTART, SCHEDSTOP, TRIGEND, TRIGFAILURE, TRIGSTART
<b>Description</b>	Specifies which events should be recorded in the appropriate event file found in Moab's <code>stats/</code> directory. These events are recorded for both local and remotely staged jobs (see <a href="#">Event Log Overview</a> ). <b>Note:</b> If a plus character is included in the list, the specified events will be added to the default list; otherwise, the specified list will replace the default list.
<b>Example</b>	<pre>RECORDEVENTLIST JOBSTART,JOBCANCEL,JOBEND</pre> <p><i>When a local and/or remote job starts, is canceled, or ends, the respective event will be recorded.</i></p>
REJECTDOSSCRIPTS	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	TRUE
<b>Description</b>	Moab rejects DOS-formatted scripts submitted with the <code>msub</code> command. This is useful if you use Slurm as your resource manager, since it does not handle DOS scripts well. For REJECTDOSSCRIPTS to work correctly, <a href="#">FILTERCMDFILE</a> must be FALSE. Otherwise, Moab modifies the script instead of rejecting it, leading to job errors.
	<div style="border: 1px solid black; border-radius: 10px; padding: 5px; display: inline-block;">  This parameter is deprecated and may be removed in a future release. </div>

### REJECTDOSSCRIPTS


<b>Example</b>	<pre>REJECTDOSSCRIPTS FALSE</pre> <p><i>Moab does not reject DOS-formatted scripts submitted with <code>msub</code>.</i></p>
----------------	--

### REJECTINFEASIBLEJOBS

<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	If zero feasible nodes are found for a job among all the nodes on the cluster and all the resource managers are reporting 'Active', the scheduler rejects the job. See <a href="#">JOBREJECTPOLICY</a> for more information.
<b>Example</b>	<pre>REJECTINFEASIBLEJOBS TRUE JOBREJECTPOLICY      MAIL, CANCEL</pre> <p><i>Any job with zero feasible nodes for execution will be rejected.</i></p>

### REJECTNEGPRIOJOBS

<b>Format</b>	<BOOLEAN>
<b>Default</b>	TRUE
<b>Description</b>	If enabled, the scheduler will refuse to start any job with a negative priority. See <a href="#">Job Priority Overview</a> and <a href="#">ENABLENEGJOBPRIORITY</a> for more information.
<b>Example</b>	<pre>ENABLENEGJOBPRIORITY TRUE REJECTNEGPRIOJOBS    TRUE</pre> <p><i>Any job with a priority less than zero will be rejected.</i></p>

REMAPCLASS	
<b>Format</b>	<ClassID>
<b>Default</b>	---
<b>Description</b>	<p>Specifies which class/queue will be remapped based on the processors, nodes, and node features requested and the resource limits of each <a href="#">class</a>. See <a href="#">Remap Class Overview</a> for more information.</p> <div style="border: 1px solid #0070C0; border-radius: 5px; padding: 5px; margin-top: 10px;"> <p> In order to use REMAPCLASS, you must specify a DEFAULTCLASS.</p> </div>
<b>Example</b>	<pre>RMCFG[internal]  DEFAULTCLASS=batch REMAPCLASS      batch CLASSCFG[small]  MAX.PROC=2 CLASSCFG[medium] MAX.PROC=16 CLASSCFG[large]  MAX.PROC=1024</pre> <div style="border: 1px dashed #000; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p><i>Class batch will be remapped based on the number of processors requested.</i></p> </div>

REMAPCLASSLIST	
<b>Format</b>	Comma-delimited list of class names
<b>Default</b>	---
<b>Description</b>	<p>The order in which classes will be searched when attempting to <a href="#">remap</a> a class. Only classes included in the list will be searched and Moab will select the first class with matches. <b>Note:</b> If no REMAPCLASSLIST is specified, Moab will search all classes and will search them in the order they are discovered. See <a href="#">Remap Class Overview</a> for more information.</p>
<b>Example</b>	<pre>RMCFG[internal]  DEFAULTCLASS=batch REMAPCLASS      batch REMAPCLASSLIST  short,medium,long</pre> <div style="border: 1px dashed #000; border-radius: 10px; padding: 10px; margin-top: 10px;"> <p><i>Class batch will be re-mapped to one of the listed classes.</i></p> </div>

REMOTEFAILTRANSIENT	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	Only applicable to Moab configurations with multiple resource managers able to run jobs (such as in a grid environment). When Moab attempts to migrate a job to one of these resource managers, a remote failure may occur. For example, a destination peer in a grid that has an error accepting a job results in a remote error, and the job is rejected. <code>REMOTEFAILTRANSIENT</code> controls how Moab reacts to remote errors. By default, Moab considers such an error permanent and does not try to migrate the same job to that resource manager again. If <code>REMOTEFAILTRANSIENT</code> is set to <code>TRUE</code> , then Moab considers such an error as transient and will not exclude the erring resource manager in future migration attempts.
<b>Example</b>	<pre>REMOTEFAILTRANSIENT  TRUE</pre>

REMOVETRIGOUTPUTFILES	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	When Moab launches external trigger actions, the standard output and error of those trigger actions are redirected to files located in Moab's spool directory. By default, these files are cleaned every 24 hours. (Files older than 24 hours are removed.) If, however, you want to have Moab immediately remove the spool files after they are no longer needed, set <code>RemoveTrigOutputFiles</code> to <code>TRUE</code> .
<b>Example</b>	<pre>REMOVETRIGOUTPUTFILES  TRUE</pre>

RESCAP	
<b>Format</b>	<DOUBLE>
<b>Default</b>	0 (NO CAP)
<b>Description</b>	The maximum allowed absolute preweighted job resource priority factor.
<b>Example</b>	<pre>RESCAP 1000</pre> <p><i>The total resource priority factor component of a job will be bound by +/- 1000</i></p>

RESERVATIONDEPTH[X]	
<b>Format</b>	<INTEGER>
<b>Default</b>	1
<b>Description</b>	The number of priority reservations that are allowed in the associated reservation bucket. <b>Note:</b> The array index, X, is the bucket label and can be any string up to 64 characters. This label should be synchronized with the <a href="#">RESERVATIONQOSLIST</a> parameter. See <a href="#">Reservation Policies</a> .
<b>Example</b>	<pre>RESERVATIONDEPTH[bigmem] 4 RESERVATIONQOSLIST[bigmem] special,fast,joshua</pre> <p><i>Jobs with QOSs of special, fast, or joshua can have a cumulative total of up to 4 priority reservations.</i></p>

RESERVATIONPOLICY	
<b>Format</b>	One of the following: CURRENTHIGHEST, HIGHEST, NEVER
<b>Default</b>	CURRENTHIGHEST
<b>Description</b>	Specifies how Moab reservations will be handled. See also <a href="#">RESERVATIONDEPTH</a> . See <a href="#">Reservation Policies</a> .
<b>Example</b>	<pre>RESERVATIONPOLICY          CURRENTHIGHEST RESERVATIONDEPTH[DEFAULT] 2</pre> <p><i>Moab will maintain reservations for only the 2 currently highest priority jobs.</i></p>

RESERVATIONQOSLIST[X]	
<b>Format</b>	One or more QOS values or [ALL]
<b>Default</b>	[ALL]
<b>Description</b>	Specifies which QOS credentials have access to the associated reservation bucket. <b>Note:</b> The array index, X, is the bucket label and can be any string up to 64 characters. This label should be synchronized with the <a href="#">RESERVATIONDEPTH</a> parameter. See <a href="#">Reservation Policies</a> .
<b>Example</b>	<pre>RESERVATIONDEPTH[big]      4 RESERVATIONQOSLIST[big]   hi,low,med</pre> <p><i>Jobs with QOSs of hi, low, or med can have a cumulative total of up to 4 priority reservations.</i></p>



RESERVATIONRETRYTIME	
<b>Format</b>	[ [ [DD:]HH:]MM:]SS
<b>Default</b>	60 seconds
<b>Description</b>	Period of time Moab will continue to attempt to allocate resources to start a job after the time resources should be made available. This parameter takes into account resource manager node state race conditions, nodes with residual high load, network glitches, etc. (For related information, see <a href="#">Reservation Policies</a> , <a href="#">DEFERSTARTCOUNT</a> , <a href="#">DEFERTIME</a> , <a href="#">NODEFAILURERESERVETIME</a> , <a href="#">JOBRETRYTIME</a> , and <a href="#">GUARANTEEDPREEMPTION</a> .)
<b>Example</b>	<pre>RESERVATIONRETRYTIME 00:05:00</pre> <p><i>Moab will try for up to 5 minutes to maintain immediate reservations if the reservations are blocked due to node state, network, or batch system based race conditions.</i></p>

RESOURCELIMITMULTIPLIER[<PARID>]	
<b>Format</b>	<p>&lt;RESOURCE&gt;:&lt;MULTIPLIER&gt; [, . . . ]</p> <p>Where &lt;RESOURCE&gt; is one of the following:            NODE, PROC, JOBPROC, MEM, JOBMEM, SWAP, DISK, or WALLTIME</p>
<b>Default</b>	1.0
<b>Description</b>	If set to less than one, then the hard limit will be the specified limit and the soft limit will be the specified limit multiplied by the multiplier. If set to a value greater than one, then the specified limit will be the soft limit and the hard limit will be the specified limit multiplied by the multiplier. See <a href="#">Usage-based Limits</a> .
<b>Example</b>	<pre>RESOURCELIMITMULTIPLIER PROC:1.1, MEM:2.0</pre> <p><i>Sets hard limit for PROC at 1.1 times the PROC soft limit, and the hard limit of MEM to 2.0 times the MEM soft limit.</i></p>

RESOURCELIMITPOLICY	
<b>Format</b>	<p>&lt;RESOURCE&gt;: [&lt;SPOLICY&gt;, ]&lt;HPOLICY&gt; :                      [&lt;SACTION&gt;, ]&lt;HACTION&gt; [:                      [&lt;SVIOLATIONTIME&gt;, ]&lt;HVIOLATIONTIME&gt;] . . .</p> <p>Where RESOURCE is one of CPU, DISK, JOBMEM, JOBPROC, MEM, MINJOBPROC, NETWORK, PROC, SWAP, or WALLTIME                      where *POLICY is one of ALWAYS, EXTENDEDVIOLATION, or BLOCKEDWORKLOADONLY                      and where *ACTION is one of CANCEL, CHECKPOINT, NOTIFY, REQUEUE, SIGNAL, or SUSPEND.</p>
<b>Default</b>	No limit enforcement.
<b>Description</b>	Specifies how the scheduler should handle jobs that utilize more resources than they request. See <a href="#">Usage-based Limits</a> .
<b>Example</b>	<pre>RESOURCELIMITPOLICY MEM:ALWAYS,BLOCKEDWORKLOADONLY:REQUEUE,CANCEL</pre> <p><i>Moab will cancel all jobs that exceed their requested memory limits.</i></p>

RESTARTINTERVAL	
<b>Format</b>	[ [ [DD:]HH:]MM:]SS
<b>Default</b>	---
<b>Description</b>	Causes Moab daemon to recycle/restart when the given interval of time has transpired.
<b>Example</b>	<pre>RESTARTINTERVAL 20:00:00</pre> <p><i>Moab daemon will automatically restart every 20 hours.</i></p>

RESOURCEQUERYDEPTH	
<b>Format</b>	<INTEGER>
<b>Default</b>	3
<b>Description</b>	Maximum number of options that will be returned in response to an <a href="#">mshow -a</a> resource query.
<b>Example</b>	<pre>RESOURCEQUERYDEPTH 1</pre> <p><i>The mshow -a command will return at most 1 valid collection of resources.</i></p>

RESWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	1
<b>Description</b>	All resource priority components are multiplied by this value before being added to the total job priority. See <a href="#">Job Prioritization</a> .
<b>Example</b>	<pre>RESWEIGHT 5 MEMWEIGHT 10 PROCWEIGHT 100 SWAPWEIGHT 0 RESCAP 2000</pre> <p><i>The job priority resource factor will be calculated as <math>\text{MIN}(2000, 5 * (10 * \text{JobMemory} + 100 * \text{JobProc}))</math>.</i></p>

RMCFG	
<b>Format</b>	One or more key-value pairs as described in the <a href="#">Resource Manager Configuration Overview</a>
<b>Default</b>	---
<b>Description</b>	The interface and policy configuration for the scheduler-resource manager interface. Described in detail in the <a href="#">Resource Manager Configuration Overview</a> .
<b>Example</b>	<pre>RMCFG[Torque3] TYPE=PBS</pre> <p><i>The PBS server will be used for resource management.</i></p>

RMMSGIGNORE	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	<p>Specifies whether or not Moab should adjust node state based on generic resource manager failure messages. See 'Compute Node Health Check' in the <i>Torque Resource Manager Administrator Guide</i> for more information.</p> <div style="border: 1px solid #0070C0; border-radius: 5px; padding: 5px; margin-top: 10px;"> <p><b>i</b> For green or ONDEMAND computing, RMMSGIGNORE must be set to TRUE to prevent Moab from powering off a down node.</p> </div>
<b>Example</b>	<pre>RMMSGIGNORE TRUE</pre> <p><i>Moab will load and report resource manager failure messages but will not adjust node state as a result of them.</i></p>

RMPOLLINTERVAL	
<b>Format</b>	[ <MINPOLLTIME> , ] <MAXPOLLTIME> where poll time is specified as [ [ [DD:] HH:] MM:] SS
<b>Default</b>	0, 30
<b>Description</b>	<p>The interval between RM polls. The poll interval will be no less than MINPOLLTIME and no more than MAXPOLLTIME. If you specify a single value, Moab interprets the value as the MAXPOLLTIME with a MINPOLLTIME of 0.</p> <div style="border: 1px solid #005596; padding: 5px; margin-top: 10px;"> <p><b>i</b> If you use Torque as your resource manager, prevent communication errors by giving tcp_timeout at least twice the value of the Moab RMPOLLINTERVAL.</p> </div>
<b>Example</b>	<pre>RMPOLLINTERVAL 30,45</pre> <div style="border: 1px dashed #ccc; padding: 10px; margin-top: 10px;"> <p><i>Moab will refresh its resource manager information between a minimum of 30 seconds and a maximum of 45 seconds. <b>Note:</b> This parameter specifies the default global poll interval for all resource managers.</i></p> </div>

RMRETRYTIMECAP	
<b>Format</b>	[ [ [DD:] HH:] MM:] SS
<b>Default</b>	1:00:00
<b>Description</b>	Moab attempts to contact RMs that are in state 'corrupt' (not down). If the attempt is unsuccessful, Moab tries again later. If the second attempt is unsuccessful, Moab increases the gap (the gap grows exponentially) between communication attempts. RMRETRYTIMECAP puts a cap on the length between connection attempts.
<b>Example</b>	<pre>RMRETRYTIMECAP 24:00:00</pre> <div style="border: 1px dashed #ccc; padding: 10px; margin-top: 10px;"> <p><i>Moab stops increasing the gap between connection attempts once the retry gap reaches 24 hours.</i></p> </div>



RSVLIMITPOLICY	
<b>Format</b>	HARD or SOFT
<b>Default</b>	---
<b>Description</b>	Specifies what limits should be enforced when creating reservations.
<b>Example</b>	<pre>RSVLIMITPOLICY HARD</pre> <p><i>Moab will limit reservation creation based on the HARD limits configured.</i></p>

RSVNODEALLOCATIONPOLICY	
<b>Format</b>	One of the following: FIRSTAVAILABLE, LASTAVAILABLE, MINRESOURCE, CPULOAD, CONTIGUOUS, MAXBALANCE, or PRIORITY
<b>Default</b>	LASTAVAILABLE
<b>Description</b>	Specifies how Moab should allocate available resources to reservations.
<b>Example</b>	<pre>RSVNODEALLOCATIONPOLICY MINRESOURCE</pre> <p><i>Moab will apply the node allocation policy MINRESOURCE to all reservations by default.</i></p>


RSVNODEALLOCATIONPRIORITYF	
<b>Format</b>	User specified algorithm
<b>Default</b>	---
<b>Description</b>	When <a href="#">RSVNODEALLOCATIONPOLICY</a> is set to PRIORITY, this parameter allows you to specify your own priority algorithm. The priority functions available are the same as the <a href="#">node priority functions</a> .
<b>Example</b>	<pre>RSVNODEALLOCATIONPOLICY PRIORITY RSVNODEALLOCATIONPRIORITYF 'SPEED + .01 * AMEM - 10 * JOBCOUNT'</pre>


RSVPROFILE[X]	
<b>Format</b>	<p>One or more of the following:</p> <p><i>Allowed:</i></p> <p><a href="#">TRIGGERACL</a> (<a href="#">ACCOUNTLIST</a>, <a href="#">CLASSLIST</a>, <a href="#">GROUPLIST</a>, <a href="#">MAXTIME</a>, <a href="#">QOSLIST</a>, <a href="#">USERLIST</a>)</p> <p>HostExp ( <a href="#">HOSTLIST</a> )</p> <p>Features (<a href="#">NODEFEATURES</a>)</p> <p><a href="#">FLAGS</a></p> <p><a href="#">TASKCOUNT</a></p> <p><a href="#">RSVACCESSLIST</a></p> <p><b>Note:</b> Lists of more than one ACL value cannot be whitespace delimited. Such lists must be delimited with a comma, pipe, or colon.</p> <p><i>Not allowed:</i></p> <p><a href="#">ACCESS</a></p> <p><a href="#">CHARGEACCOUNT</a></p> <p><a href="#">DAYS</a></p> <p><a href="#">DEPTH</a></p> <p><a href="#">ENDTIME</a></p> <p><a href="#">OWNER</a></p> <p><a href="#">PARTITION</a></p> <p><a href="#">PERIOD</a></p> <p><a href="#">PRIORITY</a></p> <p><a href="#">RESOURCES</a></p> <p><a href="#">STARTTIME</a></p> <p><a href="#">TPN</a></p>
<b>Default</b>	---
<b>Description</b>	Specifies attributes of a reservation profile using syntax similar to that for specifying a standing reservation. See <a href="#">Using Reservation Profiles</a> for details.
<b>Example</b>	<pre style="border: 1px dashed gray; padding: 10px;">RSVPROFILE[fast] USERLIST=john,steve RSVPROFILE[fast] QOSLIST=high,low RSVPROFILE[fast] TRIGGER=ETYPE=start,OFFSET=5:00,ATYPE=exec, ACTION="/opt/moab/rp.pl"</pre> <p style="border: 1px dashed gray; padding: 5px; margin-top: 10px;"><i>Moab will create a reservation profile including trigger and ACL information.</i></p>


RSVSEARCHALGO	
<b>Format</b>	LONG or WIDE
<b>Default</b>	NONE
<b>Description</b>	<p>When Moab is determining when and where a job can run, it either searches for the most resources (WIDE) or the longest range of resources (LONG). In almost all cases, searching for the longest range is ideal and returns the soonest starttime. In some rare cases, however, a particular job may need to search for the most resources. In those cases sites can configure this parameter to prevent the starvation of large jobs that fail to hold onto their reservation starttimes. See the <a href="#">WIDERSVSEARCHALGO</a> job flag.</p> <p>If this parameter is not set, it will be displayed in <code>mschedctl -l</code> as NONE but the algorithm that is used will be LONG.</p>
<b>Example</b>	<pre>RSVSEARCHALGO WIDE</pre>

SCHEDCFG	
<b>Format</b>	<p>List of zero or more space delimited &lt;ATTR&gt;=&lt;VALUE&gt; pairs, where &lt;ATTR&gt; is one of the following:  <a href="#">FBSEVER</a>, <a href="#">FLAGS</a>, <a href="#">MAXJOBID</a>, <a href="#">MINJOBID</a>, <a href="#">HTTPSERVERPORT</a>, <a href="#">MODE</a>, <a href="#">RECOVERYACTION</a>, <a href="#">SERVER</a>, <a href="#">TRIGGER</a>, or <a href="#">USEDATABASE</a></p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p> MAXRECORDEDCJOBID is deprecated.</p> </div>
<b>Default</b>	---
<b>Description</b>	<p>Specifies scheduler policy and interface configuration.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> The <code>SERVER</code> attribute can also be set using the environment variable <code>\$MOABSERVER</code>. Using this variable allows you to quickly change the Moab server that client commands will connect to.</p> <pre style="border: 1px dashed #ccc; padding: 5px; margin-top: 5px;">&gt; export MOABSERVER=cluster2:12221</pre> </div>
<b>Example</b>	<pre>SCHEDCFG[zylem3] SERVER=geronimo.scc.com:3422 MODE=NORMAL</pre> <div style="border: 1px dashed #ccc; padding: 10px; margin-top: 10px;"> <p><i>Moab will execute in NORMAL mode on the host geronimo.scc.com.</i></p> </div>



SERVERCSALGO	
<b>Format</b>	HMAC64 HMACSHA2
<b>Default</b>	HMAC64
<b>Description</b>	<p>Sets the algorithm used for message digests and message authentication codes:</p> <ul style="list-style-type: none"> <li>• HMAC64: the default (SHA-1)</li> <li>• HMACSHA2: more secure (SHA-512)</li> </ul> <div style="border: 1px solid #d9534f; padding: 5px; margin-top: 10px;"> <p> If you are using Moab Web Services, then you must set the MWS configuration parameter <code>moab.messageDigestAlgorithm</code> to match the value of <code>SERVERCSALGO</code>. See 'moab.messageDigestAlgorithm' in the <i>Moab Web Services Reference Guide</i> for more details.</p> </div>
<b>Example</b>	<div style="border: 1px dashed #ccc; padding: 5px; margin-bottom: 5px;"> <pre>SERVERCSALGO HMACSHA2</pre> </div> <div style="border: 1px dashed #ccc; padding: 5px;"> <p><i>Moab will use SHA-512 for message digests and message authentication codes.</i></p> </div>

SERVERHOST	
<b>Description</b>	<div style="border: 1px solid #4f81bd; padding: 5px;"> <p> This parameter is deprecated and may be removed in a future release. See <a href="#">SCHEDCFG</a> for replacement parameter.</p> </div>

SERVERMODE	
<b>Description</b>	<div style="border: 1px solid #4f81bd; padding: 5px;"> <p> This parameter is deprecated and may be removed in a future release. See <a href="#">SCHEDCFG</a> for replacement parameter.</p> </div>

SERVERNAME	
<b>Format</b>	<STRING>
<b>Default</b>	<SERVERHOST>
<b>Description</b>	The name the scheduler will use to refer to itself in communication with peer daemons. See <a href="#">SCHEDCFG</a> for replacement parameter.
<b>Example</b>	<code>SERVERNAME moabA</code>

SERVERPORT	
<b>Format</b>	<INTEGER> (range: 1-64000)
<b>Default</b>	40559
<b>Description</b>	Port on which moab will open its user interface socket. See <a href="#">SCHEDCFG</a> for replacement parameter.
<b>Example</b>	<code>SERVERPORT 30003</code> <i>Moab will listen for client socket connections on port 30003.</i>

SERVERSUBMITFILTER	
<b>Format</b>	<PATH>
<b>Default</b>	---
<b>Description</b>	The location of a global job submit filter script. When you configure a global job submit filter, Moab executes it on the head node and uses it to filter every job submission it receives. See <a href="#">Server-Based Submit Filter</a> for more information about job submit filters.
<b>Example</b>	<code>SERVERSUBMITFILTER /opt/moab/scripts/globalfilter.pl</code> <i>Moab uses /opt/moab/scripts/globalfilter.pl to filter every job submitted to Moab.</i>

SERVICEWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	1
<b>Description</b>	The service component weight associated with the service factors. See <a href="#">Service (SERV) Factor</a> for more information.
<b>Example</b>	<pre>SERVICEWEIGHT 2</pre>

SHOWMIGRATEDJOBSASIDLE	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	By default, migrated jobs in the grid will show as blocked. This is to prevent jobs from counting against the idle policies of multiple clusters rather than just the cluster to which the job was migrated.
<b>Example</b>	<pre>SHOWMIGRATEDJOBSASIDLE TRUE</pre> <i>When set to TRUE, migrated jobs will show as idle and will count against the idle policies of the cluster showing the job as migrated.</i>

SPOOLDIR	
<b>Format</b>	<STRING>
<b>Default</b>	---
<b>Description</b>	The directory for temporary spool files created by Moab while submitting a job to the RM.
<b>Example</b>	<pre>SPOOLDIR /tmp/moab/spool</pre>

SPOOLDIRKEEPTIME	
<b>Format</b>	<INTEGER> (seconds) or [[[DD:]HH:]MM:]SS
<b>Default</b>	---
<b>Description</b>	The interval to delete spool files and other temporary files that have been left in the spool directory. If not set, Moab will remove the spool files after a year.
<b>Example</b>	<pre>SPOOLDIRKEEPTIME 4:00:00</pre>




SPVIOLATIONWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	The weight to be applied to a job that violates <a href="#">soft</a> usage limit policies (see <a href="#">Service (SERVICE) Component</a> ).
<b>Example</b>	<pre>SPVIOLATIONWEIGHT 5000</pre>




SRCFG[X]	
<b>Format</b>	<p>One or more of the following &lt;ATTR&gt;=&lt;VALUE&gt; pairs            ACCESS, <a href="#">ACCOUNTLIST</a>, <a href="#">CHARGE</a>, <a href="#">CHARGEACCOUNT</a>, <a href="#">CHARGEUSER</a>,            CLASSLIST, <a href="#">CLUSTERLIST</a>, <a href="#">COMMENT</a>, <a href="#">DAYS</a>, DEPTH, <a href="#">DISABLE</a>, <a href="#">ENDTIME</a>,  <a href="#">FLAGS</a>, GROUPLIST, <a href="#">HOSTLIST</a>, <a href="#">JOBATTRLIST</a>, MAXTIME, NODEFEATURES,  <a href="#">OWNER</a>, PARTITION, <a href="#">PERIOD</a>, <a href="#">PROFILE</a>, PRIORITY, QOSLIST,  <a href="#">REQUIREDACCTLIST</a>, <a href="#">REQUIREDTPN</a>, <a href="#">REQUIREDUSERLIST</a>, <a href="#">RESOURCES</a>,  <a href="#">ROLLBACKOFFSET</a>, <a href="#">RSVACCESSLIST</a>, <a href="#">RSVGROUP</a>, <a href="#">STARTTIME</a>, <a href="#">TASKCOUNT</a>,            TIMELIMIT, TPN, TRIGGER, or <a href="#">USERLIST</a></p> <p><b>Note:</b> HOSTLIST and ACL list values must be comma-delimited. For example:            HOSTLIST=nodeA, nodeB</p>
<b>Default</b>	---
<b>Description</b>	Specifies attributes of a standing reservation. See <a href="#">Managing Reservations</a> for details.
<b>Example</b>	<pre>SRCFG[fast] STARTTIME=9:00:00 ENDTIME=15:00:00 SRCFG[fast] HOSTLIST=node0[1-4]\$ SRCFG[fast] QOSLIST=high,low</pre> <p><i>Moab will create a standing reservation running from 9:00 AM to 3:00 PM on nodes 1 through 4 accessible by jobs with QOS high or low.</i></p>

STARTCOUNTCAP	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	The max weighted value allowed from the startcount subfactor when determining a job's priority (see <a href="#">Priority Factors</a> for more information).
<b>Example</b>	<pre>STARTCOUNTWEIGHT 5000 STARTCOUNTCAP 30000</pre>



STARTCOUNTWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	The weight to be applied to a job's startcount when determining a job's priority (see <a href="#">Priority Factors</a> for more information).
<b>Example</b>	<pre>STARTCOUNTWEIGHT 5000</pre>



STATDIR	
<b>Format</b>	<STRING>
<b>Default</b>	stats
<b>Description</b>	The directory where Moab statistics will be maintained.
<b>Example</b>	<pre>STATDIR /var/adm/moab/stats</pre>



STATPROCMAX	
<b>Format</b>	<INTEGER>
<b>Default</b>	1
<b>Description</b>	<p>The maximum number of processors requested by jobs to be displayed in matrix outputs (as displayed by the <a href="#">showstats -f</a> command).</p> <div style="border: 1px solid red; padding: 5px; margin: 5px 0;">  <b>Caution:</b> Altering this setting will reset all recorded statistical data related to it. Do not change this parameter via <code>mschedctl -m</code> (or <code>changeparam</code>).         </div> <div style="border: 1px solid blue; padding: 5px; margin: 5px 0;">  <b>Moab only reads in this setting when starting up (or restarting).</b> </div>
<b>Example</b>	<pre style="border: 1px dashed gray; padding: 5px;">STATPROCMAX      256 STATPROCSTEPcount 4 STATPROCSTEPsize 4</pre> <div style="border: 1px dashed gray; padding: 5px; margin: 5px 0;"> <p><i>Each matrix output will display data in rows for jobs requesting between 4 and 256 processors.</i></p> </div> <div style="border: 1px solid blue; padding: 5px; margin: 5px 0;">  <b>A NONE in services will still allow users to run <a href="#">showq</a> and <a href="#">checkjob</a> on their own jobs.</b> </div>



STATPROCMIN	
<b>Format</b>	<INTEGER>
<b>Default</b>	1
<b>Description</b>	<p>The minimum number of processors requested by jobs to be displayed in matrix outputs (as displayed by the <a href="#">showstats -f</a> command).</p> <div style="border: 1px solid red; padding: 5px; margin: 5px 0;">  <b>Caution:</b> Altering this setting will reset all recorded statistical data related to it. Do not change this parameter via <code>mschedctl -m</code> (or <code>changeparam</code>).         </div> <div style="border: 1px solid blue; padding: 5px; margin: 5px 0;">  Moab only reads in this setting when starting up (or restarting).         </div>
<b>Example</b>	<div style="border: 1px dashed gray; padding: 10px; margin-bottom: 10px;"> <pre>STATPROCMIN      4 STATPROCSTEPcount 4 STATPROCSTEPsize 4</pre> </div> <div style="border: 1px dashed gray; padding: 10px; margin-bottom: 10px;"> <p><i>Each matrix output will display data in rows for jobs requesting between 4 and 256 processors.</i></p> </div> <div style="border: 1px solid blue; padding: 5px;">  A <b>NONE</b> in services will still allow users to run <a href="#">showq</a> and <a href="#">checkjob</a> on their own jobs.         </div>







STATPROCSTEPCOUNT	
<b>Format</b>	<INTEGER>
<b>Default</b>	5
<b>Description</b>	<p>The number of rows of processors requested by jobs to be displayed in matrix outputs (as displayed by the <a href="#">showstats -f</a> command).</p> <div style="border: 1px solid red; padding: 5px; margin: 5px 0;">  <b>Caution:</b> Altering this setting will reset all recorded statistical data related to it. Do not change this parameter via <code>mschedctl -m</code> (or <code>changeparam</code>).         </div> <div style="border: 1px solid blue; padding: 5px; margin: 5px 0;">  Moab only reads in this setting when starting up (or restarting).         </div>
<b>Example</b>	<pre>STATPROCMIN      4 STATPROCSTEPCOUNT 4 STATPROCSTEPsize 4</pre> <div style="border: 1px dashed gray; padding: 5px; margin-top: 10px;"> <p><i>Each matrix output will display data in rows for jobs requesting between 4 and 256 processors.</i></p> </div>

STATPROCSTEPsize	
<b>Format</b>	<INTEGER>
<b>Default</b>	4
<b>Description</b>	<p>The processor count multiplier for rows of processors requested by jobs to be displayed in matrix outputs (as displayed by the <a href="#">showstats -f</a> command).</p> <div style="border: 1px solid red; padding: 5px; margin: 5px 0;">  <b>Caution:</b> Altering this setting will reset all recorded statistical data related to it. Do not change this parameter via <code>mschedctl -m</code> (or <code>changeparam</code>).         </div> <div style="border: 1px solid blue; padding: 5px; margin: 5px 0;">  Moab only reads in this setting when starting up (or restarting).         </div>
<b>Example</b>	<pre>STATPROCMIN      4 STATPROCSTEPCOUNT 4 STATPROCSTEPsize 4</pre> <div style="border: 1px dashed gray; padding: 5px; margin-top: 10px;"> <p><i>Each matrix output will display data in rows for jobs requesting between 4 and 256 processors.</i></p> </div>

STATTIMEMAX	
<b>Format</b>	[ [DD:] HH:] MM:] SS
<b>Default</b>	00:15:00
<b>Description</b>	<p>The maximum amount of time requested by jobs to be displayed in matrix outputs (as displayed by the <a href="#">showstats -f</a> command).</p> <div style="border: 1px solid #f08080; padding: 5px; margin: 5px 0;">  <b>Caution:</b> Altering this setting will reset all recorded statistical data related to it. Do not change this parameter via <code>mschedctl -m</code> (or <code>changeparam</code>).         </div> <div style="border: 1px solid #4682b4; padding: 5px; margin: 5px 0;">  <b>Moab only reads in this setting when starting up (or restarting).</b> </div>
<b>Example</b>	<pre>STATTIMEMAX      02:08:00 STATTIMESTEPCOUNT 4 STATTIMESTEPSIZE 4</pre> <div style="border: 1px dashed #ccc; padding: 5px; margin-top: 10px;"> <p><i>Each matrix output will display data in columns for jobs requesting between 2 and 128 minutes.</i></p> </div>

STATTIMEMIN	
<b>Format</b>	[ [DD:] HH:] MM:] SS
<b>Default</b>	00:15:00
<b>Description</b>	<p>The minimum amount of time requested by jobs to be displayed in matrix outputs (as displayed by the <a href="#">showstats -f</a> command).</p> <div style="border: 1px solid #f08080; padding: 5px; margin: 5px 0;">  <b>Caution:</b> Altering this setting will reset all recorded statistical data related to it. Do not change this parameter via <code>mschedctl -m</code> (or <code>changeparam</code>).         </div> <div style="border: 1px solid #4682b4; padding: 5px; margin: 5px 0;">  <b>Moab only reads in this setting when starting up (or restarting).</b> </div>
<b>Example</b>	<pre>STATTIMEMIN      00:02:00 STATTIMESTEPCOUNT 4 STATTIMESTEPSIZE 4</pre> <div style="border: 1px dashed #ccc; padding: 5px; margin-top: 10px;"> <p><i>Each matrix output will display data in columns for jobs requesting between 2 and 128 minutes.</i></p> </div>

STATTIMESTEPCOUNT	
<b>Format</b>	<INTEGER>
<b>Default</b>	6
<b>Description</b>	<p>The number of columns of time requested by jobs to be displayed in matrix outputs (as displayed by the <a href="#">showstats -f</a> command).</p> <div style="border: 1px solid #d9534f; padding: 5px; margin: 5px 0;"> <p> <b>Caution:</b> Altering this setting will reset all recorded statistical data related to it. Do not change this parameter via <code>mschedctl -m</code> (or <code>changeparam</code>).</p> </div> <div style="border: 1px solid #4f81bd; padding: 5px; margin: 5px 0;"> <p> Moab only reads in this setting when starting up (or restarting).</p> </div>
<b>Example</b>	<pre>STATTIMEMIN      00:02:00 STATTIMESTEPCOUNT 4 STATTIMESTEPSIZE 4</pre> <div style="border: 1px dashed #ccc; padding: 5px; margin-top: 10px;"> <p><i>Each matrix output will display data in columns for jobs requesting between 2 and 128 minutes.</i></p> </div>

STATTIMESTEPSIZE	
<b>Format</b>	<INTEGER>
<b>Default</b>	4
<b>Description</b>	<p>The time multiplier for columns of time requested by jobs to be displayed in matrix outputs (as displayed by the <a href="#">showstats -f</a> command).</p> <div style="border: 1px solid #d9534f; padding: 5px; margin: 5px 0;"> <p> <b>Caution:</b> Altering this setting will reset all recorded statistical data related to it. Do not change this parameter via <code>mschedctl -m</code> (or <code>changeparam</code>).</p> </div> <div style="border: 1px solid #4f81bd; padding: 5px; margin: 5px 0;"> <p> Moab only reads in this setting when starting up (or restarting).</p> </div>
<b>Example</b>	<pre>STATTIMEMIN      00:02:00 STATTIMESTEPCOUNT 4 STATTIMESTEPSIZE 4</pre> <div style="border: 1px dashed #ccc; padding: 5px; margin-top: 10px;"> <p><i>Each matrix output will display data in columns for jobs requesting between 2 and 128 minutes.</i></p> </div>

STOPITERATION	
<b>Format</b>	<INTEGER>
<b>Default</b>	-1 (don't stop)
<b>Description</b>	Specifies which scheduling iteration Moab will stop and wait for a command to resume scheduling.
<b>Example</b>	<pre>STOPITERATION 10</pre> <p><i>Moab should stop after iteration 10 of scheduling and wait for administrator commands.</i></p>

STOREJOBSUBMISSION	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	---
<b>Description</b>	<p>When set to TRUE, specifies that Moab will save a job's submit arguments and script to \$MOABHOMEDIR/stats/jobarchive/jobNumber.</p> <div style="border: 1px solid #0070C0; padding: 5px; margin-bottom: 5px;"> <p><b>i</b> STOREJOBSUBMISSION currently does not work with jobs submitted using the Torque qsub command. Instead use Torque job logging. For more information, see 'Job Logging' in the <i>Torque Administrator Guide</i>.</p> </div> <div style="border: 1px solid #0070C0; padding: 5px;"> <p><b>i</b> Moab does not manage any of the files it creates in the stats directory. Therefore, cluster administrators should keep this fact in mind when enabling STOREJOBSUBMISSION, and implement appropriate archival and/or pruning tasks to avoid overuse of disk space.</p> </div>
<b>Example</b>	<pre>STOREJOBSUBMISSION TRUE</pre>

STRICTPROTOCOLCHECK	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	Specifies how Moab reacts to differences in XML protocols when communicating with other Moab peers. If set to <code>TRUE</code> , Moab will reject any communication that does not strictly conform to the expected protocol. If set to <code>FALSE</code> (the default), Moab will not reject XML that has extra or unknown attributes.
<b>Example</b>	<pre>STRICTPROTOCOLCHECK TRUE</pre> <p><i>Moab will reject any XML communication that does not strictly conform to the expected protocol definition.</i></p>

STRICTSCHEDULING	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	If Moab fails to get a reservation or an allocation for a priority job, then it will stop the scheduling cycle for that iteration and not attempt to start any other jobs, even through backfill.
<b>Example</b>	<pre>STRICTSCHEDULING TRUE</pre> <p><i>Moab will stop the scheduling cycle for one iteration if it fails to get a reservation or allocation for a priority job. It will not attempt to start any other jobs until the next iteration.</i></p>

SUBMITENVFILELOCATION	
<b>Format</b>	FILE or PIPE
<b>Default</b>	---
<b>Description</b>	<p>If set to FILE, these behaviors are expected:</p> <ul style="list-style-type: none"> <li>• The environment file is owned by a user with 600 permissions.</li> <li>• Moab writes the environment variables ('\0' delimited) to a random file in Moab's spool directory.</li> <li>• Moab adds the <code>--export-file=&lt;path_to_file&gt;</code> on the sbatch command line.</li> <li>• Moab deletes the file after the job completes.</li> </ul> <p>If set to PIPE, these behaviors are expected:</p> <ul style="list-style-type: none"> <li>• Moab creates a pipe and passes the read end of the pipe's file descriptor to sbatch.</li> <li>• Moab's parent process writes the environment ('\0' delimited) into the write end of the pipe.</li> </ul> <p>Adaptive Computing recommends that you configure this parameter for a more secure environment.</p>
<b>Example</b>	<code>SUBMITENVFILELOCATION PIPE</code>

SUBMITFILTER	
<b>Format</b>	<STRING>
<b>Default</b>	---
<b>Description</b>	The directory of a given <a href="#">submit filter script</a> .
<b>Example</b>	<code>SUBMITFILTER /home/submitfilter/filter.pl</code>

SUBMITHOSTS	
<b>Format</b>	space delimited list of host names
<b>Default</b>	---
<b>Description</b>	If specified, SUBMITHOSTS specifies an explicit list of hosts where jobs can be submitted.
<b>Example</b>	<pre>SUBMITHOSTS hostA hostB</pre>

SUSPENDRESOURCES[<PARID>]	
<b>Format</b>	<RESOURCE> [ , ... ] Where <RESOURCE> is one of the following: NODE, PROC, MEM, SWAP, DISK
<b>Default</b>	---
<b>Description</b>	List of resources to dedicate while a job is suspended (available in Moab version 4.5.1 and higher).
<b>Example</b>	<pre>SUSPENDRESOURCES [base] MEM, SWAP, DISK</pre> <i>While a job is suspended in partition base, the memory, swap and disk for that job will remain dedicated to the job.</i>

SYSCFG	
<b>Format</b>	List of zero or more space delimited <ATTR>=<VALUE> pairs, where <ATTR> is one of the following: PRIORITY, FSTARGET, QLIST, QDEF, PLIST, FLAGS, or a <a href="#">fairness policy</a> specification.
<b>Default</b>	---
<b>Description</b>	Specifies system-wide default attributes. See the <a href="#">Attribute/Flag Overview</a> for more information.
<b>Example</b>	<pre>SYSCFG PLIST=Partition1 QDEF=highprio</pre> <p><i>By default, all jobs will have access to partition Partition1 and will use the QOS highprio.</i></p>

SWAPWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	The priority weight assigned to the virtual memory request of a job.
<b>Example</b>	<pre>SWAPWEIGHT 10</pre>

SYSTEMMAXPROCPERJOB	
<b>Format</b>	<INTEGER>
<b>Default</b>	-1 (NO LIMIT)
<b>Description</b>	The maximum number of processors that can be requested by any single job.
<b>Example</b>	<pre>SYSTEMMAXPROCPERJOB 256</pre> <p><i>Moab will reject jobs requesting more than 256 processors.</i></p>



SYSTEMMAXPROCSECONDPERJOB	
<b>Format</b>	<INTEGER>
<b>Default</b>	-1 (NO LIMIT)
<b>Description</b>	The maximum number of proc-seconds that can be requested by any single job.
<b>Example</b>	<pre>SYSTEMMAXJOBPROCSECONDPERJOB 86400</pre> <p><i>Moab will reject jobs requesting more than 86400 proc seconds (i.e., 64 processors * 30 minutes will be rejected, while a 2 processor * 12 hour job will be allowed to run).</i></p>

SYSTEMMAXJOBWALLTIME	
<b>Format</b>	[ [ [DD:] HH:] MM:] SS
<b>Default</b>	-1 (NO LIMIT)
<b>Description</b>	The maximum amount of wallclock time that can be requested by any single job.
<b>Example</b>	<pre>SYSTEMMAXJOBWALLTIME 1:00:00:00</pre> <p><i>Moab will reject jobs requesting more than 1 day of walltime.</i></p>

TARGETQUEUETIMEWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	The weight assigned to the time remaining until the queuetime is reached.
<b>Example</b>	<pre>TARGETQUEUETIMEWEIGHT 10</pre>

TARGETWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	1
<b>Description</b>	The weight to be applied to a job's queuetime and expansion factor target components (see <a href="#">Job Prioritization</a> ).
<b>Example</b>	<pre>TARGETWEIGHT 1000</pre>

TARGETXFACTORWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	The weight assigned to the distance to the target expansion factor.
<b>Example</b>	<pre>TARGETXFACTORWEIGHT 10</pre>

TASKDISTRIBUTIONPOLICY	
<b>Format</b>	<div style="border: 1px solid #004a7c; border-radius: 10px; padding: 10px; background-color: #e6f2ff;"> <p><b>i</b> This parameter is deprecated and may be removed in a future release. (Regardless of how <code>TASKDISTRIBUTIONPOLICY</code> is configured, Moab always packs tasks by filling each node before moving to the next one in the list, in the order provided by the resource manager).</p> </div>

THREADPOOLSIZE	
<b>Format</b>	<INTEGER>
<b>Default</b>	2X number of core processors (MAX: 64)
<b>Description</b>	Governs the number of threads used when processing job scheduling. Scalability and performance might improve with multi-threading; to throttle, limit the number of threads used.
<b>Example</b>	<code>THREADPOOLSIZE 10</code>

TOOLS DIR	
<b>Format</b>	<STRING>
<b>Default</b>	tools
<b>Description</b>	The directory where Moab tools will be maintained (commonly used in conjunction with <a href="#">Native Resource Managers</a> , and <a href="#">Triggers</a> ).
<b>Example</b>	<code>TOOLS DIR /var/adm/moab/tools</code>

TRACKSUSPENDEDJOBUSAGE	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	Track the memory usage of suspended jobs on the nodes on which they are suspended and factor the memory usage into the scheduling of idle jobs.
<b>Example</b>	<code>TRACKSUSPENDEDJOBUSAGE TRUE</code> <i>Moab will track the memory usage of suspended jobs.</i>

TRAPFUNCTION	
<b>Format</b>	<STRING>
<b>Default</b>	---
<b>Description</b>	The functions to be trapped.
<b>Example</b>	<code>TRAPFUNCTION UpdateNodeUtilization GetNodeSResTime</code>

TRAPJOB	
<b>Format</b>	<STRING>
<b>Default</b>	---
<b>Description</b>	The jobs to be trapped.
<b>Example</b>	<code>TRAPJOB pros23.0023.0</code>

TRAPNODE	
<b>Format</b>	<STRING>
<b>Default</b>	---
<b>Description</b>	The nodes to be trapped.
<b>Example</b>	<code>TRAPNODE node001 node004 node005</code>

TRAPRES	
<b>Format</b>	<STRING>
<b>Default</b>	---
<b>Description</b>	The reservations to be trapped.
<b>Example</b>	<code>TRAPRES interactive.0.1</code>

TRIGCHECKTIME	
<b>Format</b>	<INTEGER> (milliseconds)
<b>Default</b>	2000
<b>Description</b>	Each scheduling iteration, Moab will have a period of time where it handles commands and other UI requests. This time period is controlled by <a href="#">RMPOLLINTERVAL</a> . During this time period, known as the UI phase, Moab will periodically evaluate triggers. Usually this only takes a fraction of a second, but if the number of triggers are large it could take up substantially more time (up to several seconds). While Moab is evaluating triggers, it doesn't respond to UI commands. This makes Moab feel sluggish and unresponsive. To remedy this, use the parameter <code>TRIGCHECKTIME</code> . This parameter tells Moab to only spend up to X milliseconds processing triggers during the UI phase. After X milliseconds has gone by, Moab will pause the evaluating of triggers, handle any pending UI events, and then restart the trigger evaluations where it last left off.
<b>Example</b>	<code>TRIGCHECKTIME 4000</code>

TRIGEVALLIMIT	
<b>Format</b>	<INTEGER>
<b>Default</b>	1
<b>Description</b>	Each scheduling iteration, Moab will have a period of time where it handles commands and other UI requests. This time period is controlled by <a href="#">RMPOLLINTERVAL</a> . During this time period, known as the UI phase, Moab will periodically evaluate triggers. The number of times Moab evaluates all triggers in the system is controlled by the TRIGEVALLIMIT parameter. By default, this is set to 1. This means that Moab will evaluate all triggers at most once during the UI phase. Moab will not leave the UI phase and start other scheduling tasks until ALL triggers are evaluated at least one time. If TrigEvalLimit is set to 5, then Moab will wait until all triggers are evaluated five times.
<b>Example</b>	<pre>TRIGEVALLIMIT 3</pre>

UIMANAGEMENTPOLICY	
<b>Format</b>	One of FORK or NONE
<b>Default</b>	NONE
<b>Description</b>	<p>When set with FORK, and with <a href="#">CLIENTUIPORT</a> specified, Moab creates a new process to handle specific command requests in order to reduce command processing time. Currently, these commands are supported:</p> <ul style="list-style-type: none"> <li>• <a href="#">checkjob</a></li> <li>• <a href="#">showbf</a></li> <li>• <a href="#">showres</a></li> <li>• <a href="#">showstart</a></li> </ul> <div style="border: 1px solid #0056b3; border-radius: 5px; padding: 5px; margin-top: 10px;"> <p><b>i</b> See <a href="#">Considerations for Large Clusters</a> for additional information on reducing command time (also known as low latency).</p> </div> <div style="border: 1px solid #0056b3; border-radius: 5px; padding: 5px; margin-top: 10px;"> <p><b>i</b> This parameter should be configured on the server, as well as any client machines.</p> </div>
<b>Example</b>	<pre>UIMANAGEMENTPOLICY FORK CLIENTUIPORT 42560</pre>

UJOBWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	Weight assigned by jobs per user. -1 will reduce priority by number of active jobs owned by user.
<b>Example</b>	<pre>UJOBWEIGHT 10</pre>

UMASK	
<b>Format</b>	<INTEGER>
<b>Default</b>	0022 (octal) (produces 0644 permissions)
<b>Description</b>	The file permission mask to use when creating new fairshare, stats, and event files. See the <code>umask</code> man page for more details.
<b>Example</b>	<pre>UMASK 0127</pre> <i>Create statistics and event files that are 'read-write' by owner and 'read' by group only.</i>

UNSUPPORTEDDEPENDENCIES	
<b>Format</b>	Comma-delimited string
<b>Default</b>	---
<b>Description</b>	Specifies <a href="#">dependencies</a> that are not supported and should not be accepted by job submissions. A maximum of 30 dependencies is supported.
<b>Example</b>	<pre># moab.cfg UNSUPPORTEDDEPENDENCIES before,beforeok,beforenotok,on &gt; msub -l depend=before:105 cmd.sh ERROR: cannot submit job - error in extension string</pre>

UPROCWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	Weight assigned by processors per user. -1 will reduce priority by number of active procs owned by user.
<b>Example</b>	<code>UPROCWEIGHT 10</code>

USAGECONSUMEDWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	The weight assigned to per job processor second consumption.
<b>Example</b>	<code>USAGECONSUMEDWEIGHT 10</code>

USAGEEXECUTIONTIMEWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	The priority weight assigned to the total job execution time (measured in seconds since job start). See <a href="#">Preemption Overview</a> .
<b>Example</b>	<code>USAGEEXECUTIONTIMEWEIGHT 10</code>



USAGEPERCENTWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	The weight assigned to total requested resources consumed.
<b>Example</b>	<code>USAGEPERCENTWEIGHT 5</code>

USAGEREMAININGWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	The weight assigned to remaining usage.
<b>Example</b>	<code>USAGEREMAININGWEIGHT 10</code>

USAGEWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	1
<b>Description</b>	The weight assigned to the percent and total job usage subfactors.
<b>Example</b>	<code>USAGEWEIGHT 100</code>

USEANYPARTITIONPRIO	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	<p>The FSTREE data from the first feasible FSTREE will be used when determining a job's start priority, rather than having no FSTREE data considered.</p> <div style="border: 1px solid #0070C0; padding: 5px; margin-top: 10px;"> <p><b>i</b> Do not set USEANYPARTITIONPRIO if you use per-partition scheduling. Doing so causes to schedule jobs to the first partition listed, even if nodes from another partition will be available sooner.</p> </div>
<b>Example</b>	USEANYPARTITIONPRIO TRUE

USECPRSVNODELIST	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	TRUE
<b>Description</b>	Specifies whether Moab should use the checkpointed reservation node list when rebuilding reservations on startup. If this is not used then Moab will use the reservation's specified host expression during rebuilding.
<b>Example</b>	USECPRSVNODELIST FALSE

USEDATABASE	
<b>Format</b>	INTERNAL
<b>Default</b>	-
<b>Description</b>	Specifies whether Moab should store profile statistics, checkpoint information, and event information in an integrated database. See <a href="#">Layout of Scheduler Components with Integrated Database</a> for more information.
<b>Example</b>	USEDATABASE INTERNAL


USEJOBREGEX	
<b>Format</b>	BOOLEAN
<b>Default</b>	FALSE
<b>Description</b>	Specifies whether <i>mjobctl</i> supports regular expressions.
<b>Example</b>	<pre>USEJOBREGEX TRUE  [user@linux]\$ mjobctl -c 8[1-3]  job '81' cancelled job '82' cancelled job '83' cancelled</pre>

USEMOABCTIME	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	<p>When Moab finds new jobs on the resource manager, it creates a job inside of Moab for each job in the resource manager. By default, when Moab creates a new job, it uses the time the job was submitted to the resource manager to calculate how long the job has been in the queue (Moab processing time - job creation in resource manager), which is then used in determining the job's priority.</p> <p>In a system where more jobs are submitted to a resource manager than Moab can handle in one iteration, there is the possibility of jobs running out of order. For example, two jobs are both submitted at time 5. The first submitted job is processed first at time 6. So the first job's effective queue duration would be 1 (6-5). On the next iteration, the second job is processed at time 8. So the second job's effective queue duration would be 3 (8-5), indicating that it has been in the queue longer than the other job. Since the later job has a higher effective queue duration it will get a higher priority and could be scheduled to run before earlier submitted jobs.</p> <p>Setting USEMOABCTIME to TRUE tells Moab to use the creation time of the job in Moab rather than the creation time in the resource manager. This corrects the possible problem of having later submitted jobs having higher priorities and starting before earlier submitted jobs.</p>
<b>Example</b>	<pre>USEMOABCTIME TRUE</pre>

USEMOABJOBID	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	<p>Specifies whether to return the Moab job ID when running 'msub', or the resource manager's job ID if it is available.</p> <div style="border: 1px solid #0070C0; padding: 5px; margin-top: 10px;"> <p><b>i</b> USEMOABJOBID can also be set at the job level. The job level setting overrides this (global) setting in moabcfg. See <a href="#">USEMOABJOBID</a> for more information.</p> </div>
<b>Example</b>	<pre>USEMOABJOBID TRUE</pre>

USERCFG[<USERID>]	
<b>Format</b>	<p>List of zero or more space delimited &lt;ATTR&gt;=&lt;VALUE&gt; pairs, where &lt;ATTR&gt; is one of the following:</p> <p><a href="#">General Credential Flags</a>, CDEF, DEFAULT.TPN, DEFAULT.WCLIMIT, EMAILADDRESS, <a href="#">ENABLEPROFILING</a>, FSCAP, FSTARGET, <a href="#">JOBFLAGS</a>, <a href="#">MAX.ARRAYSUBJOBS</a>, MAX.WCLIMIT, NOEMAIL, OVERRUN, PLIST, <a href="#">PRIORITY</a>, PRIVILEGES, QLIST, QDEF, or a <a href="#">usage limit</a>.</p>
<b>Default</b>	---
<b>Description</b>	<p>Specifies user specific attributes. For general user attribute information, See the <a href="#">Credential Overview</a>. For a description of legal flag values, see <a href="#">flag overview</a>.</p>
<b>Example</b>	<pre>USERCFG[john] MAXJOB=50 QDEF=highprio USERCFG[john] EMAILADDRESS=john@company.com</pre> <div style="border: 1px dashed gray; padding: 5px; margin-top: 10px;"> <p><i>Up to 50 jobs submitted under the user ID john will be allowed to execute simultaneously and will be assigned the QOS highprio.</i></p> </div>

USERPRIOCAP	
<b>Format</b>	<INTEGER>
<b>Default</b>	---
<b>Description</b>	The priority cap to be applied to the user specified job priority factor. Under Moab, only negative user priorities can be specified. See <a href="#">Credential (Service) Factor</a> .
<b>Example</b>	<pre>USERPRIOWEIGHT 10 USERPRIOCAP -10000</pre>

USERPRIOWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	1
<b>Description</b>	<p>The weight to be applied to the user specified job priority. Under Moab, only negative user priorities can be specified. If this weight is set, users can reduce the priority of some of their jobs to allow other jobs to run earlier. See <a href="#">Credential (Service) Factor</a> and <a href="#">User Selectable Prioritization</a>.</p> <div style="border: 1px solid #005596; padding: 5px; margin-top: 10px;"> <p> If Viewpoint is part of your configuration, this value must be at least 1. Otherwise, Moab will not take into consideration any user priority information specified for a job that was created using Viewpoint.</p> </div>
<b>Example</b>	<pre>USERPRIOWEIGHT 10</pre>

USERWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	1
<b>Description</b>	The weight to be applied to the user <a href="#">priority</a> of each job. See <a href="#">Credential (CRED) Factor</a> .
<b>Example</b>	<pre>USERWEIGHT 10</pre>

USESYSLOG	
<b>Format</b>	<BOOLEAN> [<FACILITY>]
<b>Default</b>	FALSE:daemon
<b>Description</b>	Specifies whether or not the scheduler will report key events to the system syslog facility. If the <FACILITY> is specified, Moab will report events to this syslog facility. See <a href="#">Logging Facilities</a> for more information.
<b>Example</b>	<pre>USESYSLOG TRUE:local3</pre> <p><i>Moab will report key events, commands, and failures to syslog using the local3 facility.</i></p>

USESYSTEMQUEUE TIME	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	Specifies whether or not job prioritization should be based on the time the job has been eligible to run (i.e., idle and meets all fairness policies [TRUE] or the time the job has been idle [FALSE]). See <a href="#">Priority Factors</a> for more info. <b>Note:</b> This parameter has been superseded by the <a href="#">JOBPRIOACCRUALPOLICY</a> parameter.
<b>Example</b>	<pre>USESYSTEMQUEUE TIME FALSE</pre> <p><i>The queuetime and expansion factor components of a job's priority will be calculated based on the length of time the job has been in the idle state.</i></p>

USEUSERHASH	
<b>Format</b>	<BOOLEAN>
<b>Default</b>	FALSE
<b>Description</b>	Enables searching of the user buffer using the user hash key instead of doing sequential searches of the user buffer.
<b>Example</b>	<pre>USEUSERHASH TRUE</pre>

WALLTIMECAP	
<b>Format</b>	<DOUBLE>
<b>Default</b>	0 (NO CAP)
<b>Description</b>	The maximum total preweighted absolute contribution to job priority that can be contributed by the walltime component. This value is specified as an absolute priority value, not as a percent.
<b>Example</b>	<pre>WALLTIMECAP 10000</pre> <i>Moab will bound a job's preweighted walltime priority component within the range +/- 10000.</i>

WALLTIMEWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	The priority weight to be applied to the amount of walltime requested by a job (in seconds) (see <a href="#">Resource (RES) Factor</a> ).
<b>Example</b>	<pre>RESWEIGHT 10 WALLTIMEWEIGHT 100</pre> <i>Increase the priority of longer duration jobs.</i>


WCACCURACYCAP	
<b>Format</b>	<DOUBLE>
<b>Default</b>	0 (NO CAP)
<b>Description</b>	The maximum total preweighted absolute contribution to job priority that can be contributed by the wallclock accuracy component. This value is specified as an absolute priority value, not as a percent.
<b>Example</b>	<pre>WCACCURACYCAP 10000</pre> <p><i>Moab will bound a job's preweighted wallclock accuracy priority component within the range +/- 10000.</i></p>

WCACCURACYWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	The priority weight to be applied to the job's historical user wallclock accuracy (range 0.0 to 1.0) (see <a href="#">Fairshare (FS) Factor</a> ).
<b>Example</b>	<pre>FSWEIGHT 10 WCACCURACYWEIGHT 100</pre> <p><i>Favor jobs with good wallclock accuracies by giving them a priority increase.</i></p>



WCVIOLATIONACTION	
<b>Format</b>	One of CANCEL or PREEMPT
<b>Default</b>	CANCEL
<b>Description</b>	The action to take when a job exceeds its wallclock limit. If set to CANCEL, the job will be terminated. If set to PREEMPT, the action defined by <a href="#">PREEMTPOLICY</a> parameter will be taken. See <a href="#">JOBMAXOVERRUN</a> or <a href="#">Usage-based limits</a> .
<b>Example</b>	<pre>WCVIOLATIONACTION  PREEMPT PREEMTPOLICY       REQUEUE</pre> <p><i>Moab will requeue jobs that exceed their wallclock limit.</i></p>

WEBSERVICESURL	
<b>Format</b>	<URL>
<b>Default</b>	---
<b>Description</b>	If specified, Moab sends data to Moab Web Services (MWS) to be stored in a database. This allows Moab to spend more cycles on scheduling instead of database interaction. The sending occurs via HTTP PUT.
<b>Example</b>	<pre>WEBSERVICESURL https://mws-staging.ac:8080/mws/rm/moab/dump</pre> <p><i>Moab sends data that needs to be stored in a database to the specified URL.</i></p>

WIKIEVENTS	
<b>Format</b>	<BOOLEAN>
	<div style="border: 1px solid #0070C0; border-radius: 10px; padding: 5px; display: inline-block;">  This parameter is deprecated and may be removed in a future release. </div>

XFACTORCAP	
<b>Format</b>	<DOUBLE>
<b>Default</b>	0 (NO CAP)
<b>Description</b>	The maximum total preweighted absolute contribution to job priority that can be contributed by the expansion factor component. This value is specified as an absolute priority value, not as a percent.
<b>Example</b>	<pre>XFACTORCAP 10000</pre> <p><i>Moab will bound a job's preweighted XFactor priority component within the range +/- 10000.</i></p>

XFACTORWEIGHT	
<b>Format</b>	<INTEGER>
<b>Default</b>	0
<b>Description</b>	The weight to be applied to a job's minimum expansion factor before it is added to the job's cumulative priority.
<b>Example</b>	<pre>XFACTORWEIGHT 1000</pre> <p><i>Moab will multiply a job's XFactor value by 1000 and then add this value to its total priority.</i></p>

XFMINWCLIMIT	
<b>Format</b>	[ [ [DD:] HH:] MM:] SS
<b>Default</b>	-1 (NO LIMIT)
<b>Description</b>	The minimum job wallclock limit that will be considered in job expansion factor priority calculations.
<b>Example</b>	<pre>XFMINWCLIMIT 0:01:00</pre> <p><i>Jobs requesting less than 1 minute of wallclock time will be treated as if their wallclock limit was set to 1 minute when determining expansion factor for priority calculations.</i></p>

## Appendix B: Multi-OS Provisioning

In this section:

- [B.1 xCAT Plug-in Configuration Parameters - page 1185](#)
- [B.2 Configuration Validation - page 1192](#)
- [B.3 Deploying Images with Torque - page 1193](#)
- [B.4 Installing Moab on the Management Node - page 1193](#)
- [B.5 Integrating MSM and xCAT - page 1194](#)
- [B.6 Moab Configuration File Example - page 1195](#)
- [B.7 MSM Configuration - page 1196](#)
- [B.8 MSM Installation - page 1196](#)
- [B.9 Troubleshooting - page 1197](#)
- [B.10 Verifying the Installation - page 1197](#)
- [B.11 xCAT Configuration Requirements - page 1200](#)

Moab can dynamically provision compute machines to requested operating systems and power off compute machines when not in use. Moab can intelligently control xCAT and use its advanced system configuration mechanisms to adapt systems to current workload requirements. Moab communicates with xCAT using the Moab Service Manager (MSM). MSM is a translation utility that resides between Moab and xCAT and acts as aggregator and interpreter. The Moab Workload Manager will query MSM, which in turn queries xCAT, about system resources, configurations, images, and metrics. After learning about these resources from MSM, Moab then makes intelligent decisions about the best way to maximize system utilization.

In this model Moab gathers system information from two resource managers. The first is Torque, which handles the workload on the system; the second is MSM, which relays information gathered by xCAT. By leveraging these software packages, Moab intelligently adapts clusters to deliver on-site goals.

This document assumes that xCAT has been installed and configured. It describes the process of getting MSM and xCAT communicating, and it offers troubleshooting guidance for basic integration. This document offers a description for how to get Moab communicating with MSM and the final steps in verifying a complete software stack.

## B.1 xCAT Plug-in Configuration Parameters

Plugin parameters that begin with an underscore character are specific to the xCAT plug-in; others are common to all plug-ins and can either be set in the `RMCFG [msm]` for all plug-ins, or per plug-in in the `APPCFG [<plugin_name>]`.

<code>_CQxCATSessions</code>	<code>_ModifyTorque</code>	<code>_UseStates</code>
<code>_DoNodeStat</code>	<code>_NodeRange</code>	<code>_VerifyRPower</code>
<code>_DORVitals</code>	<code>_NoRollbackOnError</code>	<code>_xCATHost</code>
<code>_DoxCATStats</code>	<code>_PowerString</code>	Description
<code>_ESXStore</code>	<code>_QueueRPower</code>	LogLevel
<code>_FeatureGroups</code>	<code>_ReportNETADDR</code>	Module
<code>_HVxCATPasswdKey</code>	<code>_RPowerQueueAge</code>	PollInterval
<code>_ImagesTabName</code>	<code>_RPowerQueueSize</code>	TimeOut
<code>_LockDir</code>	<code>_RPowerTimeOut</code>	
<code>_MaskOSWhenOff</code>	<code>_UseOpIDs</code>	

<code>_CQxCATSessions</code>	
<b>Format</b>	Positive integer > 1
<b>Default</b>	10
<b>Description</b>	MSM will divide the node list generated by <code>nodeIs</code> into this many groups and simultaneously query xCAT for each group. The value may need tuning for large installations, higher values will cause the time to complete a single cluster query to go down, but cause a higher load on the xCAT headnode.

<code>_DoNodeStat</code>	
<b>Format</b>	0 or 1
<b>Default</b>	1
<b>Description</b>	If set to 0, MSM will not call <code>nodestat</code> to generated a substate. This can be used to speed up the time it takes to query xCAT, and you do not need the substate visible to Moab.

<code>_DORVitals</code>	
<b>Format</b>	0 or 1

<b>_DORVitals</b>	
<b>Default</b>	0
<b>Description</b>	When set to 1, MSM will poll rvitals power and led status (see the xCAT rvitals man page). This only works with IBM BMCs currently. In order to use this, xCAT should respond without error to the <code>rvitals &lt;noderange&gt; watts</code> and <code>rvitals &lt;noderange&gt; leds</code> commands. Status is reported as <code>GMETRTIC[watts]</code> and <code>GMETRIC[leds]</code> . See also the <a href="#">_PowerString</a> configuration parameter.

<b>_DoxCATStats</b>	
<b>Format</b>	0 or 1
<b>Default</b>	0
<b>Description</b>	If Set to 1, MSM will track performance statistics about calls to xCAT, and the performance of higher level operations. The information is available via the script <code>\$MSMHOMEDIR/contrib/xcat/xcatstats.pl</code> . This parameter is useful for tuning the <code>POLLINTERVAL</code> and <a href="#">_CQxCATSessions</a> configuration parameters.

<b>_ESXStore</b>	
<b>Format</b>	Mountable NFS Path
<b>Default</b>	---
<b>Description</b>	Location of ESX stores.

<b>_FeatureGroups</b>	
<b>Format</b>	Comma-delimited string of xCAT group names.
<b>Default</b>	---
<b>Description</b>	MSM builds the OSLIST for a node as the intersection of <code>_FEATUREGROUPS</code> , features specified in <code>x_msm</code> for that image, and the nodes group membership. The value 'torque' is special, and indicates that the image uses Torque, and the node should be added/removed from Torque during provisioning when used in conjunction with the <a href="#">_ModifyTorque</a> parameter.

<b>_HVxCATPasswdKey</b>	
<b>Format</b>	key value in the xCAT passwd table
<b>Default</b>	vmware
<b>Description</b>	This is where MSM gets the user/password to communicate with ESX hypervisors.

<b>_ImagesTabName</b>	
<b>Format</b>	Existing xCAT table that contains your image definitions.
<b>Default</b>	x_msm
<b>Description</b>	This table specifies the images that can be presented to Moab in a node's OSLIST. The xCAT schema for this table is defined in <code>\$MSMHOMEDIR/contrib/xcat/MSM.pm</code> , which needs to be copied to the <code>\$XCATROOT/lib/perl/xCAT_schema</code> directory.

<b>_LockDir</b>	
<b>Format</b>	Existing path on MSM host
<b>Default</b>	<code>\$MSMHOMEDIR/lck</code>
<b>Description</b>	This is a path to where MSM maintains lock files to control concurrency with some Xen and KVM operations.

<b>_MaskOSWhenOff</b>	
<b>Format</b>	0 or 1
<b>Default</b>	0
<b>Description</b>	When set, this parameter will cause MSM to report <code>OS=None</code> for nodes that are powered off. This may be useful when mixing stateless and stateful images, forcing Moab to request provisioning instead of just powering on a node.

<b>_ModifyTorque</b>	
<b>Format</b>	0 or 1
<b>Default</b>	0
<b>Description</b>	When set, this parameter will cause MSM to add and removes nodes from Torque as required by provisioning. See the <a href="#">_FeatureGroups</a> parameter as well.

<b>_NodeRange</b>	
<b>Format</b>	Any valid noderange (see the xCAT noderange man page).
<b>Default</b>	All
<b>Description</b>	When MSM queries xCAT this is the noderange it will use. At sites where xCAT manages other hardware that Moab is not intended to control, it is important to change this.

<b>_NoRollbackOnError</b>	
<b>Format</b>	0 or 1
<b>Default</b>	0
<b>Description</b>	When an error occurs and rollback is activated (as it is by default), rollback causes a reversion to the previous successful request. <code>_NoRollbackOnError</code> is useful for debugging to determine the xCAT state if no rollback occurred. If set to 1 and an error occurs between MSM and xCAT when creating a node, assigning a name (DNS) to a node, or assigning an IP address (DHCP) to a node, then no rollback occurs.

<b>_PowerString</b>	
<b>Format</b>	single quote delimited string
<b>Default</b>	'AC Avg Power'
<b>Description</b>	Only meaningful when used with <code>_DORVitals=1</code> . Some BMCs return multiple responses to the rvitals command, or use slightly different text to describe the power metrics. Use this parameter to control what is reported to Moab. You



<b>_PowerString</b>	
	can use <code>'\$MSMLIBDIR/contrib/xcat/dump.xcat.cmd.plrvitals &lt;node_name&gt; power'</code> and examine the output to determine what the appropriate value of this string is.

<b>_QueueRPower</b>	
<b>Format</b>	0 or 1
<b>Default</b>	0
<b>Description</b>	When set, this parameter will cause MSM to aggregate rpower requests to xCAT into batches. The timing and size of these batches is controlled with the <a href="#">_RPowerQueueAge</a> and <a href="#">_RPowerQueueSize</a> parameters. <b>Note:</b> This can significantly reduce load on the xCAT headnode, but will cause the power commands to take longer, and MSM shutdown to take longer.

<b>_ReportNETADDR</b>	
<b>Format</b>	0 or 1
<b>Default</b>	0
<b>Description</b>	When set, this parameter will cause MSM to report NETADDR=<hosts.ip from xCAT>.

<b>_RPowerQueueAge</b>	
<b>Format</b>	Positive integer values
<b>Default</b>	30
<b>Description</b>	Only meaningful when used with <a href="#">_QueueRPower</a> . MSM will send any pending rpower requests when the oldest request in the queue exceeds this value (seconds).

<b>_RPowerQueueSize</b>	
<b>Format</b>	Positive integer values
<b>Default</b>	200
<b>Description</b>	Only meaningful when used with <a href="#">_QueueRPower</a> . MSM will send any pending rpower requests when the queue depth exceeds this value.

<b>_RPowerTimeOut</b>	
<b>Format</b>	Positive integer values
<b>Default</b>	60
<b>Description</b>	Only meaningful when used with <a href="#">_VerifyRPower</a> . If nodes do not report the expected power state in this amount of time, a GEVENT will be produced on the node (or system job).

<b>_UseOpIDs</b>	
<b>Format</b>	0 or 1
<b>Default</b>	0
<b>Description</b>	When set, this parameter will cause errors to be reported as GEVENTs on the provided system job, instead of a node (Moab 5.4 only, with appropriate Moab CFG)

<b>_UseStates</b>	
<b>Format</b>	Valid xCAT chain.currstate values (see the xCAT chain man page)
<b>Default</b>	boot,netboot,install
<b>Description</b>	Nodes that do not have one of these values in the xCAT chain.currstate field will reported with STATE=Updating. Use this configuration parameter to prevent Moab from scheduling nodes that are updating firmware, etc.

<b>_VerifyRPower</b>	
<b>Format</b>	0 or 1
<b>Default</b>	0
<b>Description</b>	If set, MSM will attempt to confirm that rpower requests were successful by polling the power state with rpower stat until the node reports the expected state, or <a href="#">_RPowerTimeOut</a> is reached. <b>Note:</b> This can create significant load on the xCAT headnode.

<b>_xCATHost</b>	
<b>Format</b>	<xcat_headnode>:<xcatd_port>
<b>Default</b>	localhost:3001
<b>Description</b>	Use to configure MSM to communicate with xCAT on another host.

<b>Description</b>	
<b>Format</b>	Double quoted string containing brief description of plugin.
<b>Default</b>	---
<b>Description</b>	This information is not visible in Moab, but shows up in <code>msmctl -a</code> .

<b>LogLevel</b>	
<b>Format</b>	1-9
<b>Default</b>	5
<b>Description</b>	Used to control the verbosity of logging, 1 being the lowest (least information logged) and 9 being the highest (most information logged). For initial setup and testing, 8 is recommended, then lowering to 3 (only errors logged) for normal operation. Use 9 for debugging, or when submitting a log file for support.

Module	
<b>Format</b>	Moab::MSM::App::xCAT
<b>Default</b>	---
<b>Description</b>	Name of the plugin module to load.

PollInterval	
<b>Format</b>	Integer > 0
<b>Default</b>	60
<b>Description</b>	<p>MSM will query xCAT every <code>POLLINTERVAL</code> seconds to update general node status. This number will likely require tuning for each specific system. In general, to develop this number, you should pick a fraction of the total nodes MSM will be managing (<code>1/_CQXCATSESSIONS</code>), and time how long it takes run <code>nodestat</code>, <code>rpower stat</code>, and optionally <code>rvitals</code> on these nodes, and add ~15%.</p> <p>Increasing the <code>POLLINTERVAL</code> will lower the overall load on the xCAT headnode, but decrease the responsiveness to provisioning and power operations.</p>

TimeOut	
<b>Format</b>	Integer value > <code>POLLINTERVAL</code>
<b>Default</b>	300
<b>Description</b>	<p>This parameter controls how long MSM will wait for child processed to complete (all xCAT commands are run in child processes). After <code>TIMEOUT</code> seconds, if a child has not returned it will be killed, and an error reported for the operation.</p>

## B.2 Configuration Validation

Set up environment to manually call MSM commands:

```
# substitute appropriate value(s) for path(s)
export MSMHOMEDIR=/opt/moab/tools/msm
export MSMLIBDIR=/opt/moab/tools/msm
export PATH=$PATH:/$MSMLIBDIR/contrib:$MSMLIBDIR/bin
```

Verify that MSM starts without errors:

```
> msmd
```

Verify that the expected nodes are listed, without errors, using the value of `_NODERANGE` from `msm.cfg`:

```
> nodels <_NODERANGE>
```

Verify that the expected nodes, are listed in the cluster query output from MSM:

```
> cluster.query.pl
```

Provision all nodes through MSM for the first time (pick an image name from `x_msm`):

```
> for i in `nodels <_NODERANGE>`; do node.modify.pl $i --set os=<image_name>;done
```

Verify the nodes correctly provision and that the correct OS is reported (which might take some time after the provisioning requests are made):

```
> cluster.query.pl
```

## B.3 Deploying Images with Torque

When using MSM + xCAT to deploy images with Torque, there are some special configuration considerations. Most of these also apply to other workload resource managers.

Note that while the MSM xCAT plugin contains support for manipulating Torque directly, this is not an ideal solution. If you are using a version of xCAT that supports prescripts, it is more appropriate to write prescripts that manipulate Torque based on the state of the xCAT tables. This approach is also applicable to other workload resource managers, while the xCAT plugin only deals with Torque.

Several use cases and configuration choices are discussed in what follows.

Each image should be configured to report its image name through Torque. In the `Torquepbs_mommom_config` file the `opsys` value should mirror the name of the image. See 'Appendix C: Node Manager (MOM) Configuration' in the *Torque Administrator Guide* for more information.

## B.4 Installing Moab on the Management Node

Moab is the intelligence engine that coordinates the capabilities of xCAT and Torque to dynamically provision compute nodes to the requested operating system. Moab also

schedules workload on the system and powers off idle nodes. Download and install Moab.

## B.5 Integrating MSM and xCAT

Copy the `x_msm` table schema to the xCAT schema directory:

```
> cp $MSMHOMEDIR/contrib/xcat/MSM.pm $XCATROOT/lib/perl/xCAT_schema
```

Restart `xcatd` and check the `x_msm` table is correctly created:

```
> service xcatd restart
```

```
> tabdump x_msm
```

Prepare xCAT images and ensure they provision correctly (see [xCAT documentation](#)).

Populate the `x_msm` table with your image definitions:

```
> tabedit x_msm

#flavorname,arch,profile,os,nodeset,features,vmoslist,hvtype,hvgroupname,vmgroupname,comments,disable
"compute","x86_64","compute","centosX.X","netboot","torque",,,,,,
"science","x86","compute","scientific_linux","netboot","torque",,,,,,
```

- **flavorname** - A user specified name for the image and settings; also an xCAT group name, nodes are added to this group when provisioned
- **arch** - Architecture as used by xCAT
- **profile** - Profile as used by xCAT
- **os** - Operating system as used by xCAT
- **nodeset** - One of netboot|install|statelite
- **features** - Names of xCAT groups that identify special hardware features ('torque' and 'paravirt' are special cases)
- **vmoslist** - Note: Not used. List of flavorname's this image can host as VMs (hypervisor images only)
- **hvtype** - Note: Not used. One of esx|xen|kvm (hypervisor images only)
- **hvgroupname** - Note: Not used. Name of xCAT group nodes will be added to when provisioned to this image
- **vmgroupname** - Note: Not used. Name of xCAT group VMs will be added to when hosted on a hypervisor of this image

- **comments** - User specified comments
- **disable** - Flag to temporarily disable use of this image

Ensure all xCAT group names in the x\_msm table exist in the xCAT nodegroup table:

```
> tabedit nodegroup
```

Edit as necessary to simulate the following example:

```
#groupname, grouptype, members, wherevals, comments, disable
"compute",,,,,,
"esxi4",,,,,,
"esxhv",,,,,,
"esxvmgmt",,,,,,
```

After making any necessary edits, run the following command:

```
> nodels compute,esxi4,esxhv,esxvmgmt
# should complete without error, ok if doesn't return anything
```

## B.6 Moab Configuration File Example

Moab stores its configuration in the `moab.cfg` file: `/opt/moab/etc/moab.cfg`. A sample configuration file, set up and optimized for Adaptive Computing follows:

```
SCHEDCFG[Moab]          SERVER=gpc-sched:42559
ADMINCFG[1]             USERS=root,egan
LOGLEVEL                7

# How often (in seconds) to refresh information from Torque and MSM
RMPOLLINTERVAL         60,60
RESERVATIONDEPTH       10
DEFERTIME               0
TOOLS DIR              /opt/moab/tools

#####
# Torque and MSM configuration                                     #
#####
RMCFG[torque]          TYPE=PBS
RMCFG[msm]             TYPE=NATIVE:msm FLAGS=autosync,NOCREATERESOURCE RESOURCETYPE=PROV
RMCFG[msm]             TIMEOUT=60
RMCFG[msm]             PROVDURATION=10:00
AGGREGATENODEACTIONS   TRUE

#####
# ON DEMAND PROVISIONING SETUP                                   #
#####
QOSCFG[od]             QFLAGS=PROVISION
USERCFG[DEFAULT]       QLIST=od
NODEALLOCATIONPOLICY    PRIORITY
NODECFG[DEFAULT]       PRIORITYF=1000*OS+1000*POWER
NODEAVAILABILITYPOLICY DEDICATED
CLASSCFG[DEFAULT]     DEFAULT.OS=scinetcompute
```

```
#####
# GREEN POLICIES #
#####
NODECFG[DEFAULT]      POWERPOLICY=ONDEMAND
PARCFG[ALL]           NODEPOWEROFFDURATION=20:00
NODEIDLEPOWERTHRESHOLD 600
# END Example moab.cfg
```

## B.7 MSM Configuration

Edit `$MSMHOMEDIR/msm.cfg` and configure the xCAT plug-in. Below is a generic example for use with Torque without virtualization. See the section on configuration parameters for a complete list of parameters and descriptions.

```
# MSM configuration options
RMCFG[msm]      PORT=24603
RMCFG[msm]      POLLINTERVAL=45
RMCFG[msm]      LOGFILE=/opt/moab/log/msm.log
RMCFG[msm]      LOGLEVEL=8
RMCFG[msm]      DEFAULTNODEAPP=xcat

# xCAT plugin specific options
APPCFG[xcat]    DESCRIPTION="xCAT plugin"
APPCFG[xcat]    MODULE=Moab::MSM::App::xCAT
APPCFG[xcat]    LOGLEVEL=3
APPCFG[xcat]    POLLINTERVAL=45
APPCFG[xcat]    TIMEOUT=3600
APPCFG[xcat]    _USEOPIDS=0
APPCFG[xcat]    _NODERANGE=moab,esxcompute
APPCFG[xcat]    _USESTATES=boot,netboot,install
APPCFG[xcat]    _LIMITCLUSTERQUERY=1
APPCFG[xcat]    _RPOWERTIMEOUT=120
APPCFG[xcat]    _DONODESTAT=1
APPCFG[xcat]    _REPORTNETADDR=1
APPCFG[xcat]    _CQXCATSESSIONS=4
```

## B.8 MSM Installation

- Determine the installation directory (usually `/opt/moab/tools/msm`)
- Untar the MSM tarball into the specified directory (making it the MSM home directory, or `$MSMHOMEDIR`)
- Verify the required Perl modules and version are available:

```
> perl -e 'use Storable 2.18'
> perl -MXML::Simple -e 'exit'
> perl -MProc::Daemon -e 'exit'
> perl -MDBD::SQLite -e 'exit'
```



## B.9 Troubleshooting

- **msmctl -a does not report the xCAT plugin** - Check the log file (path specified in `msm.cfg`) for error messages. A common cause is missing Perl modules (Storable, DBD::SQLite, xCAT::Client).
- **cluster.query.pl does not report any nodes** - Check that the xCAT command `nodeIs<noderange>`, where `<noderange>` is the value configured for `_NODERANGE` in `msm.cfg`, outputs the nodes expected.
- **cluster.query.pl does not report OS** - MSM must provision a node to recognize what the current operating system is. It is not sufficient to look up the values in the `nodetype` table because MSM has no way of recognizing whether `nodeset` and `rpower` were run with the current values in the `nodetype` table.
- **cluster.query.pl does not report OSLIST, or does not report the expected OSLIST for a node** - Check that the node belongs to the appropriate groups, particularly any listed in the `features` field of the `x_msm` table for the missing image name.

## B.10 Verifying the Installation

When Moab starts it immediately communicates with its configured resource managers. In this case Moab communicates with Torque to get compute node and job queue information. It then communicates with MSM to determine the state of the nodes according to xCAT. It aggregates this information and processes the jobs discovered from Torque.

When a job is submitted, Moab determines whether nodes need to be provisioned to a particular operating system to satisfy the requirements of the job. If any nodes need to be provisioned Moab performs this action by creating a provisioning system job (a job that is internal to Moab). This system job communicates with xCAT to provision the nodes and remain active while the nodes are provisioning. Once the system job has provisioned the nodes it informs the user's job that the nodes are ready at which time the user's job starts running on the newly provisioned nodes.

When a node has been idle for a specified amount of time (see [NODEIDLEPOWERTHRESHOLD](#)), Moab creates a power-off system job. This job communicates with xCAT to power off the nodes and remains active in the job queue until the nodes have powered off. Then the system job informs Moab that the nodes are powered off but are still available to run jobs. The power off system job then exits.

To verify correct communication between Moab and MSM run the `mdiag -R -v msm` command:

```

$ mdiag -R -v msm
diagnosing resource managers
RM[msm]      State: Active Type: NATIVE:MSM ResourceType: PROV
Timeout:    30000.00 ms
Cluster Query URL: $HOME/tools/msm/contrib/cluster.query.xcat.pl
Workload Query URL: exec://$TOOLSDIR/msm/contrib/workload.query.pl
Job Start URL:  exec://$TOOLSDIR/msm/contrib/job.start.pl
Job Cancel URL:  exec://$TOOLSDIR/msm/contrib/job.modify.pl
Job Migrate URL: exec://$TOOLSDIR/msm/contrib/job.migrate.pl
Job Submit URL:  exec://$TOOLSDIR/msm/contrib/job.submit.pl
Node Modify URL: exec://$TOOLSDIR/msm/contrib/node.modify.pl
Node Power URL:  exec://$TOOLSDIR/msm/contrib/node.power.pl
RM Start URL:   exec://$TOOLSDIR/msm/bin/msmd
RM Stop URL:    exec://$TOOLSDIR/msm/bin/msmctl?-k
System Modify URL: exec://$TOOLSDIR/msm/contrib/node.modify.pl
Environment:
MSMHOMEDIR=/home/wightman/test/scinet/tools//msm;MSMLIBDIR=/home/wightman/test/scinet/
tools//msm
Objects Reported: Nodes=10 (0 procs) Jobs=0
Flags:            autosync
Partition:        SHARED
Event Management: (event interface disabled)
RM Performance:   AvgTime=0.10s MaxTime=0.25s (38 samples)
RM Languages:     NATIVE
RM Sub-Languages: -

```

To verify nodes are configured to provision use the `checknode -v` command. Each node will have a list of available operating systems.

```

$ checknode n01
node n01
State:      Idle (in current state for 00:00:00)
Configured Resources: PROCS: 4 MEM: 1024G SWAP: 4096M DISK: 1024G
Utilized Resources: ---
Dedicated Resources: ---
Generic Metrics:  watts=25.00,temp=40.00
Power Policy:    Green (global policy) Selected Power State: Off
Power State:     Off
Power:          Off
MTBF(longterm): INFINITY MTBF(24h): INFINITY
Opsys:          compute Arch: ---
OS Option:      compute
OS Option:      computea
OS Option:      gpfscompute
OS Option:      gpfscomputea
Speed:          1.00 CPUload: 0.000
Flags:          rmdetected
RM[msm]:        TYPE=NATIVE:MSM ATTRO=POWER
EffNodeAccessPolicy: SINGLEJOB
Total Time: 00:02:30 Up: 00:02:19 (92.67%) Active: 00:00:11 (7.33%)

```

To verify nodes are configured for Green power management, run the `mdiag -G` command. Each node will show its power state.

```

$ mdiag -G
NOTE: power management enabled for all nodes
Partition ALL: power management enabled
Partition NodeList:
Partition local: power management enabled
Partition NodeList:

```

```

node n01 is in state Idle, power state On (green powerpolicy enabled)
node n02 is in state Idle, power state On (green powerpolicy enabled)
node n03 is in state Idle, power state On (green powerpolicy enabled)
node n04 is in state Idle, power state On (green powerpolicy enabled)
node n05 is in state Idle, power state On (green powerpolicy enabled)
node n06 is in state Idle, power state On (green powerpolicy enabled)
node n07 is in state Idle, power state On (green powerpolicy enabled)
node n08 is in state Idle, power state On (green powerpolicy enabled)
node n09 is in state Idle, power state On (green powerpolicy enabled)
node n10 is in state Idle, power state On (green powerpolicy enabled)
Partition SHARED: power management enabled

```

To submit a job that dynamically provisions compute nodes, run the `msub -l os=<image>` command:

```

$ msub -l os=computea job.sh
yuby.3
$ showq
active jobs-----
JOBID            USERNAME        STATE PROCS    REMAINING      STARTTIME
provision-4      root            Running      8      00:01:00  Fri Jun 19 09:12:56
1 active job
                8 of 40 processors in use by local jobs (20.00%)
                2 of 10 nodes active          (20.00%)

eligible jobs-----
JOBID            USERNAME        STATE PROCS    WCLIMIT        QUEUE TIME
yuby.3           wightman        Idle      8      00:10:00  Fri Jun 19 09:12:55
1 eligible job

blocked jobs-----
JOBID            USERNAME        STATE PROCS    WCLIMIT        QUEUE TIME

0 blocked jobs
Total jobs: 2

```

Notice that Moab created a provisioning system job named `provision-4` to provision the nodes. When `provision-4` detects that the nodes are correctly provisioned to the requested OS, the submitted job `yuby.3` runs:

```

$ showq
active jobs-----
JOBID            USERNAME        STATE PROCS    REMAINING      STARTTIME
yuby.3           wightman        Running      8      00:08:49  Fri Jun 19 09:13:29
1 active job
                8 of 40 processors in use by local jobs (20.00%)
                2 of 10 nodes active          (20.00%)

eligible jobs-----
JOBID            USERNAME        STATE PROCS    WCLIMIT        QUEUE TIME

0 eligible jobs

blocked jobs-----
JOBID            USERNAME        STATE PROCS    WCLIMIT        QUEUE TIME

0 blocked jobs
Total job: 1

```

The `checkjob` command shows information about the provisioning job, as well as the submitted job. If any errors occur, run the `checkjob -v <jobid>` command to diagnose failures.

## B.11 xCAT Configuration Requirements

Observe the following xCAT configuration requirements before installing MSM:

- Configure xCAT normally for your site.
  - Test the following commands to verify proper function:
    - **rpower**
    - **nodeset**
    - **makedhcp**
    - **makedns**
    - **nodestat**
    - **rvitals**
  - If MSM will run on a different machine than the one on which xCAT runs, install the xCAT client packages on that machine, and test the previously listed commands on that machine as well.
  - Configure and test all stateful/stateless images you intend to use.
- Configure xCAT to use either PostgreSQL or MySQL. Note that the default of SQLite might not function properly when MSM drives xCAT.
  - PostgreSQL: See [Setting Up PostgreSQL as the xCAT DB](#) for more information.
  - MySQL: See [Setting Up MySQL as the xCAT DB](#) for more information.

**i** You must have a valid Moab license file (`moab.lic`) with provisioning and green enabled. For information on acquiring an evaluation license, please contact [info@adaptivecomputing.com](mailto:info@adaptivecomputing.com).

## Appendix C: Moab Event Dictionary

**i** See [Logging Overview](#) for more information about Moab logging.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x1000005</b>	USER	system.moab	INFO	MWM_TESTING_INFO	Testing with argument1: %s. and argument2: %s and argument3: %s and argument4: %s	Internal error for testing diagnostics.
<b>0x1000065</b>	USER	domain.life cycle	INFO	MWM_JOB_CANCEL	Job %s was canceled. %s	The job was canceled.
<b>0x1000066</b>	USER	domain.life cycle	INFO	MWM_JOB_END_SUCCESSFUL	Job %s finished successfully at %s.	The job finished successfully.
<b>0x1000068</b>	USER	domain.life cycle	INFO	MWM_JOB_USER_HOLD	Job %s had a user hold applied.	A user hold was applied to the job.
<b>0x1000069</b>	USER	domain.life cycle	INFO	MWM_JOB_SYSTEM_HOLD	Job %s had a system hold applied.	A system hold was applied to the job.
<b>0x100006a</b>	USER	domain.life cycle	INFO	MWM_JOB_BATCH_HOLD	Job %s had a batch hold applied.	A batch hold was applied to the job.
<b>0x100006b</b>	USER	domain.life cycle	INFO	MWM_JOB_DEFER_HOLD	Job %s had a defer hold applied.	A defer hold was applied to the job.
<b>0x100006c</b>	USER	domain.life cycle	INFO	MWM_JOB_MODIFY	Job %s was modified. %s	One of the attributes of

Cod e	Es cal ation Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
						the job was modified either via a user initiated action or an automated action.
<b>0x100006d</b>	US ER	doma in.life cycle	IN F O	MWM_JOB_REJECT	Job %s was rejected. %s	The job was rejected for some reason.
<b>0x100006e</b>	US ER	doma in.life cycle	IN F O	MWM_JOB_RELEASE	Job %s was released.	Any holds placed on the job have been released, and the job is not prevented from running due to any hold action. The job might still not be able to run due to other considerations.
<b>0x100006f</b>	US ER	doma in.life cycle	IN F O	MWM_JOB_START	Job %s was started. %s	The job was started on its designated node[s].
<b>0x1000070</b>	US ER	doma in.life cycle	IN F O	MWM_JOB_SUBMIT	Job %s was submitted. %s	The job has been submitted to Moab and is being evaluated and processed.

Cod e	Es cal ation Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
<b>0x1000071</b>	US ER	doma in.life cycle	IN F O	MWM_JOB_CREATED	Job %s was created.	The job has been created and will be queued for execution.
<b>0x1000072</b>	US ER	doma in.life cycle	IN F O	MWM_JOB_REQUEUE	Job %s was requeued. %s	The job has been requeued so it can be executed again.
<b>0x1000073</b>	US ER	doma in.life cycle	IN F O	MWM_JOB_CANCEL_CLEANUP_STARTED	Job %s is being cleaned up due to cancel request.	The job has been issued a cancel request and is being cleaned up.
<b>0x1000074</b>	US ER	doma in.life cycle	IN F O	MWM_JOB_CLEANUP_STARTED	Job %s is being cleaned up.	The job has ended and is being cleaned up.
<b>0x1000075</b>	US ER	doma in.life cycle	IN F O	MWM_JOB_DEFERRED	Job %s has been deferred.	The job has been deferred.
<b>0x1000076</b>	US ER	doma in.life cycle	IN F O	MWM_JOB_RENAME	Job %s has been renamed to %s.	The job has been renamed.
<b>0x10000c9</b>	US ER	doma in.life cycle	IN F O	MWM_NODE_EVAC_VMS	Evacuating VMs off node %s.	Evacuating VMs off the node.
<b>0x100012c</b>	US ER	doma in.life cycle	IN F O	MWM_RSV_CREATE	Reservation %s was created. %s	The reservation has been created and is stored in the

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
						system.
<b>0x100012d</b>	USER	domain.life cycle	INFO	MWM_RSV_START	Reservation %s has started.	The reservation has started.
<b>0x100012e</b>	USER	domain.life cycle	INFO	MWM_RSV_END	Reservation %s has ended.	The reservation has ended.
<b>0x1000190</b>	USER	system.moab	INFO	MWM_SCHED_COMMAND	The following scheduler command was submitted: %s	External commands are submitted to Moab in a variety of ways. This event documents the command line and possibly other information associated with the command. These commands typically have the ability to change behavior/state within Moab. Commands that are typically queries are not included.
<b>0x1000</b>	USER	system.mo	INFO	MWM_SCHED_	A scheduler iteration is beginning. %s	Moab periodically



Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>192</b>		ab	0	CYCLE_START		checks through submitted jobs and makes decisions regarding which jobs are scheduled. One of these iterations is beginning now.
<b>0x1000193</b>	USER	system.moab	INFO	MWM_SCHED_CYCLE_END	A scheduler iteration is ending. %s	Moab periodically checks through submitted jobs and makes decisions regarding which jobs are scheduled. One of these iterations is ending now.
<b>0x1000194</b>	USER	system.moab	INFO	MWM_SCHED_PAUSE	The scheduler has been paused. %s	The Moab scheduler has been administratively paused. New jobs can be submitted and existing jobs will continue to run, but no new jobs will be scheduled as long as Moab is paused.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x1000195</b>	USER	system.moab	INFO	MWM_SCHED_RECYCLE	The scheduler has been recycled. %s	The Moab scheduler has been administratively recycled. The process will cleanly exit and save the state data. It will then restart, read in the data, and resume scheduling.
<b>0x1000196</b>	USER	system.moab	INFO	MWM_SCHED_RESUME	The scheduler has been resumed.	The Moab scheduler has been administratively resumed. A new scheduling iteration will begin immediately and continue regularly.
<b>0x1000197</b>	USER	system.moab	INFO	MWM_SCHED_START	The scheduler has started.	The Moab scheduler has started.
<b>0x1000198</b>	USER	system.moab	INFO	MWM_SCHED_STOP	The scheduler has stopped. %s	The Moab scheduler has stopped.
<b>0x10001f4</b>	USER	domain.lifecycle	INFO	MWM_TRIG_CREATE	Trigger %s has been created.	The named trigger has been created and is now

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
						recognized in the Moab system.
<b>0x10001f5</b>	USER	domain.life cycle	INFO	MWM_TRIG_START	Trigger %s has started.	The named trigger has started its action.
<b>0x10001f6</b>	USER	domain.life cycle	INFO	MWM_TRIG_END	Trigger %s has ended. %s	The named trigger has finished its action.
<b>0x10001f8</b>	USER	domain.life cycle	INFO	MWM_TRIG_THRESHOLD	Trigger %s threshold event: %s	A trigger threshold has been encountered. Additional details regarding the threshold might be included in the text.
<b>0x1000258</b>	USER	domain.life cycle	INFO	MWM_VM_SUBMIT	VM %s has been submitted.	The named VM has been submitted and is now recognized in the Moab system.
<b>0x1000259</b>	USER	domain.life cycle	INFO	MWM_VM_DESTROY	VM %s has been terminated.	The named VM has finished its lifecycle and is now removed and added to the completed

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
						table.
<b>0x100025a</b>	USER	domain.life cycle	INFO	MWM_VM_CANCEL	VM %s has been canceled.	The named VM has been canceled.
<b>0x100025b</b>	USER	domain.life cycle	INFO	MWM_VM_END	VM %s has been ended.	The named VM has been canceled because it has exceeded its allocated walltime.
<b>0x100025c</b>	USER	domain.life cycle	INFO	MWM_VM_MIGRATE_START	VM %s migration has started. (%s)	The named VM has started its migration. Additional information might be provided regarding source and destination nodes.
<b>0x100025d</b>	USER	domain.life cycle	INFO	MWM_VM_MIGRATE_END	VM %s migration has finished. (%s)	The named VM has finished its migration. Additional information might be provided regarding source and destination nodes.
<b>0x1000</b>	USER	domain.life	INFO	MWM_VM_MANUAL_	VM %s migration started manually. (%s)	The named VM

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
25f		cycle	0	MIGRATE_START		migration has been started manually. Additional information might be provided regarding source and destination nodes.
0x1000260	USER	domain.life cycle	INFO	MWM_VM_READY	VM %s is ready.	The named VM is ready. It has been linked to an internal tracking job.
0x1002711	USER	system.moab	INFO	MWM_PARAMETER_SET_TO_VALUE_INFO	Parameter '%s' is set to '%s'.	A parameter was set to a specified value. This is usually accomplished via a configuration file.
0x1002741	USER	system.moab	INFO	MWM_SOCKET_EXCEPTION	Exception detected in select for socket %s.	The select() system call indicated an exception for this socket.
0x1002742	USER	system.moab	INFO	MWM_SOCKET_EXCEPTION_REASON	Exception identified as '%s' in select for socket %s.	The select() system call indicated an exception for this socket. It has been identified with

Cod e	Es calati on Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
						an error ID by getsockopt().
<b>0x1002748</b>	US ER	system.moab	INFO	MWM_MOAB_STARTED_ON_CORRECT_HOST	Server started on host '%s' %s.	Moab is started on either the primary or fallback server.
<b>0x1002762</b>	US ER	system.moab	INFO	MWM_CONFIG_LINE_SUCCESSFUL	Configuration line '%s' successfully processed.	The line in the configuration file was processed without error.
<b>0x1002935</b>	US ER	system.moab	INFO	MWM_ACTIVE_JOB_REMOVED_FROM_QUEUE	Active %s job %s has been removed from the queue, default to successful completion.	The job was removed from the indicated resource manager while it was still active. By default, it is assumed to complete successfully unless more information is available (i.e., ENABLEFAILU REFORPURGE DJOB).
<b>0x1002936</b>	US ER	system.moab	INFO	MWM_INACTIVE_JOB_REMOVED_FROM_QUEUE	Inactive %s job %s has been removed from the queue, default to status 'canceled'.	The job was removed from the indicated resource manager while it was still

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
						active. By default, it is given status 'canceled' unless more information is available (i.e., ENABLEFAILU REFORPURGE DJOB).
<b>0x1002937</b>	USER	system.moab	INFO	MWM_RM_DOWN_SKIPPING_WORK	RM %s state is %s, skipping %s.	The specified resource manager is not in a good state. Certain actions might be skipped while it is in this state.
<b>0x100296a</b>	USER	system.moab	INFO	MWM_CANNOT_RESUME_JOB	Cannot resume job '%s' (%s).	Check the PBS server log to see reason for failure.
<b>0x100296b</b>	USER	system.moab	INFO	MWM_CANNOT_LOCATE_RESOURCE	Cannot locate %s '%s'.	Unable to find the resource specified.
<b>0x100296c</b>	USER	system.moab	INFO	MWM_CANNOT_SET_JOB_CLASS	Cannot set class on job '%s' to '%s' (%s).	The job could not be modified.
<b>0x100296d</b>	USER	system.moab	INFO	MWM_NATIVE_ACTION_MISSING	%s action not specified for native interface. %s.	The native interface allows custom actions to be

Cod e	Es calati on Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
						specified. No value was specified for this action.
<b>0x100296e</b>	US ER	system.moab	INFO	MWM_COMMAND_SENT	Command sent to server.	A command was sent.
<b>0x100296f</b>	US ER	system.moab	INFO	MWM_JOB_MAXPREMPT	JOBMAXPREEMPTPERITERATION reached: %s of %s.	The maximum value was reached.
<b>0x1002970</b>	US ER	system.moab	INFO	MWM_JOB_CHANGED_STATES	Job '%s' changed states from '%s' to '%s'.	The state changed.
<b>0x1002971</b>	US ER	system.moab	INFO	MWM_JOB_RELEASING_RESERVATION	Job '%s' was requeued/rejected. Releasing reservation.	The job no longer holds the reservation.
<b>0x1002972</b>	US ER	system.moab	INFO	MWM_NODE_CHANGED_STATES	Node '%s' changed states from '%s' to '%s'.	The node state changed.
<b>0x1002973</b>	US ER	system.moab	INFO	MWM_JOB_ACTION_SUCCESSFUL	Job '%s' successfully %s.	The job action completed.
<b>0x1002974</b>	US ER	system.moab	INFO	MWM_ALLOC_TEMP_MEMORY	Cannot allocate temp memory for %s completed jobs.	The system might be low on memory.
<b>0x1002</b>	US ER	system.moab	INFO	MWM_ACTION_	Action '%s' launched with message '%s'. PID = '%s'	Scheduler action is about



Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
975		ab	O	LAUNCHED		to be executed.
0x1002976	USER	system.moab	INFO	MWM_JOB_ADJUSTMENT	Adjusting allocated %s to %s for job '%s'.	The value is being changed.
0x1002977	USER	system.moab	INFO	MWM_ALL_JOBS_LOADED	All jobs loaded.	The jobs have been loaded.
0x1002978	USER	system.moab	INFO	MWM_ALL_NODES_LOADED	All located non-native nodes loaded (%s).	The nodes have been loaded.
0x1002979	USER	system.moab	INFO	MWM_BACKFILL_POLICY_DISABLED	Backfill policy disabled.	The policy was disabled.
0x100297a	USER	system.moab	INFO	MWM_JOB_LOAD	Cannot load job '%s'.	The job failed to load.
0x100297b	USER	system.moab	INFO	MWM_CANNOT_CREATE_RSV	Cannot create reservation.	The request to create the given reservation has failed.
0x100297c	USER	system.moab	INFO	MWM_MODIFY_PARTITION	Cannot modify partition of running job '%s'.	Must wait until job completes.
0x100297f	USER	system.moab	INFO	MWM_CANNOT_CREATE_RSV_IN_PARTITION	Cannot create reservation for job '%s' in partition '%s'.	Job cannot be run on requested partition.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
0x1002982	USER	system.moab	INFO	MWM_CLUSTER_QUERY_GETDATA	Cluster query getdata failed for native interface.	The resource manager might be down or unresponsive.
0x1002a0b	USER	system.moab	INFO	MWM_SENDING_CLIENT_COMMAND	Sending %s command: '%s'.	The specified command is being sent to the server.
0x1002a0e	USER	system.moab	INFO	MWM_SCHED_SHUTDOWN_REQUEST	The scheduler has received a user shutdown request.	The Moab scheduler has received a request to shut down. It will be processed as soon as possible.
0x1002a0f	USER	system.moab	INFO	MWM_SCHED_RECYCLE_REQUEST	The scheduler has received a user recycle request.	The Moab scheduler has received a request to recycle. It will be processed as soon as possible.
0x1002a10	USER	system.moab	INFO	MWM_SCHED_PAUSE_DESCRIPTION	Scheduling will be disabled, cluster information will continue to be updated.	This is a description of what happens when the scheduler is paused.
0x1002a11	USER	system.moab	INFO	MWM_SCHED_STOP_TIMESTAMP	Scheduling will stop in %s at iteration %s.	This provides a log message of when the scheduler will stop.

Cod e	Es cal ation Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
<b>0x1002a12</b>	US ER	system.moab	INFO	MWM_SCHED_RESUME_TIMESTAMP	Scheduling will resume in %s seconds.	This provides a log message of when the scheduler will resume.
<b>0x1002a13</b>	US ER	system.moab	INFO	MWM_SCHED_RESTART_TIME_REACHED	Scheduler restart time reached (scheduler will restart).	The configured restart time was reached. (RESTARTINTERVAL or MEMREFRESHINTERVAL.
<b>0x1002a14</b>	US ER	system.moab	INFO	MWM_SCHED_COMPLETE_SCHEDULING	Scheduling complete. Sleeping for %s seconds.	The scheduling portion of the iteration is complete. Additional jobs will not be scheduled until the next iteration.
<b>0x1002a17</b>	US ER	system.moab	INFO	MWM_ABOUT_TO_EXEC	About to exec() '%s'.	The process is about to be executed.
<b>0x1002a18</b>	US ER	system.moab	INFO	MWM_JOB_ARRAY_CANCEL_POLICY	Sub-job %s exit code %s canceled job array %s with policy %s.	A job within an array job finished and, depending on its exit code and the policy in place, the entire array job might cancel.
<b>0x1</b>	US	system	IN	MWM_	Reservation completion for	The

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>002a19</b>	ER	m.moab	FO	RESERVATION_COMPLETION_DELAYED	job '%s' delayed from %s to %s.	reservation end time is later than initially expected for this job.
<b>0x1002a1b</b>	USER	system.moab	INFO	MWM_VM_ORPHANED	VM '%s' successfully orphaned.	The VM has been separated from its tracking job.
<b>0x1002a1c</b>	USER	system.moab	INFO	MWM_VM_REPORTED_DESTROYED	VM '%s' reported destroyed via RM - removing VM.	The VM is no longer available from the resource manager, so it is being removed from the scheduler.
<b>0x1002a1d</b>	USER	system.moab	INFO	MWM_VM_STALE_REPORT	VM '%s' has not been reported in %s seconds.	The VM is no longer being reported from the resource manager. No action is currently being taken.
<b>0x1002a1e</b>	USER	system.moab	INFO	MWM_WIKI_KEYWORD_NOT_HANDLED	Wiki keyword '%s'(%s) not handled.	The keyword was not recognized, so it will be ignored.
<b>0x1002a20</b>	USER	system.moab	INFO	MWM_ROLLING_LOGFILE	Rolling logfile '%s' to '%s'.	The old logfile will be closed and logging will resume in

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
						the new file.
<b>0x1002a23</b>	USER	system.moab	INFO	MWM_NODE_LOCATED	Nodes located for job %s: %s of %s required (%s feasible).	List of nodes located for a specific job.
<b>0x1002a2d</b>	USER	system.moab	INFO	MWM_JOB_PAL_SET	Partition access list set to value: %s.	The partition access list (PAL) is set.
<b>0x1002a2e</b>	USER	system.moab	INFO	MWM_JOB_PREEMPTED_BY_JOB	Job %s preempted job %s - added idle resources (T: %s; N: %s; P: %s)/remaining (T: %s; N: %s; P: %s).	Job was preempted by another job.
<b>0x1002a2f</b>	USER	system.moab	INFO	MWM_JOB_CAN_START_WITHOUT_PREEMPTION	Job %s would start in %s without preemption (PC: %s).	Job can start without preemption.
<b>0x1002a32</b>	USER	system.moab	INFO	MWM_PARTITION_RESOURCES	Partition %s nodes/procs available %s/%s (%s jobs examined).	General partition information.
<b>0x1002a33</b>	USER	system.moab	INFO	MWM_RSV_OPERATION	Performing '%s' operation on reservation expression '%s' (%s matches).	This is operation is caused by a mrsvctl command.
<b>0x1002a34</b>	USER	system.moab	INFO	MWM_PREEMPTING_JOBS	Preempting jobs to allow job %s to start - required resources T: %s; N: %s; P: %s.	Preempting jobs.
<b>0x1</b>	USER	system	INFO	MWM_	Preempt usage tracking	Environment

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>002a35</b>	ER	m.moab	FO	MOABTRACKSUSPEND	enabled (env).	variable MOABTRACKSUSPEND set.
<b>0x1002a36</b>	USER	system.moab	INFO	MWM_JOB_MAX_PREEMPTEE_LIMIT	Single job max preemptee limit (%s) reached.	Max requirements exceeded on job.
<b>0x1002a37</b>	USER	system.moab	INFO	MWM_QUEUES_DETECTED	Queues detected: %s.	Resource manager found queues on cluster query update.
<b>0x1002a38</b>	USER	system.moab	INFO	MWM_JOB_START_TIME_CHANGED	Start time changed from %s to %s on job %s.	The job's start time was changed via the resource manager.
<b>0x1002a39</b>	USER	system.moab	INFO	MWM_STORING_CHECKPOINT_INFO	Storing object to checkpoint.	The object's state is being checkpointed.
<b>0x1002a3a</b>	USER	system.moab	INFO	MWM_PBS_DATA_UP_TO_DATE	PBS raw data already up to date.	The resource manager is already updated.
<b>0x1002a3b</b>	USER	system.moab	INFO	MWM_PBS_DATA_UPDATED	PBS data updated for iteration %s.	The resource manager is now updated.
<b>0x1002a3c</b>	USER	system.moab	INFO	MWM_STARTED_MESSAGE_QUEUE	Started message queue thread.	The message queue is now operational.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x1002a3d</b>	USER	system.moab	INFO	MWM_JOBS_SELECTED_IN_PARTITION	Total jobs selected in partition %s: %s/%s.	Identifies the selected jobs in a partition.
<b>0x1002a3e</b>	USER	system.moab	INFO	MWM_TASKS_LOCATED_FOR_JOB	Tasks located for job %s: %s of %s required (%s feasible).	Identifies the tasks available for a job.
<b>0x1002a3f</b>	USER	system.moab	INFO	MWM_CLIENT_REQUEST	Client requesting command '%s'.	Client requested command.
<b>0x1002a40</b>	USER	system.moab	INFO	MWM_REQUEST_TO_CANCEL_JOB	Request to cancel job '%s' sent, but could not confirm cancellation (pending response).	Client did not get a confirmation as expected.
<b>0x1002a41</b>	USER	system.moab	INFO	MWM_RESERVATION_NOT_ALLOWED_FOR_JOB	Reservation not allowed for job %s in %s.	Reservation not allowed in specified condition.
<b>0x1002a42</b>	USER	system.moab	INFO	MWM_RESERVED_JOB_STARTED	Reserved job '%s' started.	Reserved job started.
<b>0x1002a43</b>	USER	system.moab	INFO	MWM_RESOURCE_AVAILABLE_AFTER_SCHEDULING	Resources available after scheduling: N: %s P: %s.	Resources available after scheduling.
<b>0x1</b>	USER	system	INFO	MWM_	Restoring job '%s' from	Restoring job

Cod e	Es calati on Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
<b>002 a44</b>	ER	m.mo ab	F O	RESTORIN G_ DEFERRED_ JOB	deferred state.	from deferred state.
<b>0x1 002 a45</b>	US ER	syste m.mo ab	IN F O	MWM_RM_ DUPLICAT E_QUERY	RM %s already has a pending query - skipping get data query.	Duplicate queries cannot be performed simultaneously.
<b>0x1 002 a46</b>	US ER	syste m.mo ab	IN F O	MWM_RM_ PEER_ COMMAND	Sending peer server command to %s:%s (Cmd: %s, Requestor: %s, Key: %s...).	A command has been sent to a peer Moab grid server.
<b>0x1 002 a47</b>	US ER	syste m.mo ab	IN F O	MWM_SET_ ATTRIBUT E_ON_NODE	Setting %s on node %s to %s.	A command has been sent to a peer Moab grid server.
<b>0x1 002 a48</b>	US ER	syste m.mo ab	IN F O	MWM_SET_ ATTRIBUT E_ON_JOB	Setting %s on job %s to %s (%s).	A command has been sent to a peer Moab grid server.
<b>0x1 002 a49</b>	US ER	syste m.mo ab	IN F O	MWM_ INVALID_ STORAGE_ DATA	Storage data from MWS RM (%s) is not a valid object.	Invalid object from Storage data.
<b>0x1 00a 73c</b>	AD MI N	syste m.mo ab	IN F O	MWM_ CANNOT_ STAT_FILE_ INFO	Cannot stat file '%s', errno: %s (%s).	The stat() system call failed. This is not always significant as it is sometimes used to test the existence



Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
						of a file that may or may not be there. Use the errno and associated message to determine possible causes.
<b>0x100a743</b>	ADMIN	system.moab	INFO	MWM_FAILED_SELECT	Select for socket %s failed, errno: %s (%s).	The select() system call failed. Use the errno and associated message to determine possible causes.
<b>0x100a744</b>	ADMIN	system.moab	INFO	MWM_SELECT_TIMEOUT	Select for socket %s timed out after %s seconds with no valid descriptors.	The select() system call timed out. This may or may not be an error. Check MTU.
<b>0x100a75d</b>	ADMIN	system.moab	INFO	MWM_CONFIG_VALUE_OUT_OF_RANGE	Configuration parameter '%s' has an invalid value '%s'. Range is limited by %s.	Check the line in the configuration file for the attribute.
<b>0x100a767</b>	ADMIN	system.moab	INFO	MWM_NO_MONGOSERVER_SPECIFIED	Moab is built to use Mongo, but no MONGOSERVER is specified.	Cannot connect to the Mongo server since the MONGOSERVER parameter

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
						was unspecified. Add MONGOSERVER parameter to moab.cfg.
<b>0x100a76d</b>	ADMIN	system.moab	INFO	MWM_REMOVING_OBJECT_FROM_MONGO	Removing object '%s' from Mongo DB '%s'.	The object is being removed from the database.
<b>0x100a789</b>	ADMIN	system.moab	INFO	MWM_UNABLE_TO_ALLOCATE_NODES_FOR_RSV	Cannot allocate nodes for reservation '%s'. (%s)	Cannot allocate a node list that matches the requirements for this reservation. This might not be serious since multiple passes can occur.
<b>0x100a8ab</b>	ADMIN	system.moab	INFO	MWM_VM_EXCEED_TTL	VM '%s' has reached TTL (%s). Must be removed manually.	The given VM has reached its time to live.
<b>0x100a8e1</b>	ADMIN	system.moab	INFO	MWM_PLUGIN_LOADED_SUCCESS	Successfully loaded NodeAllocation plugin '%s' for partition '%s'.	A NodeAllocation plugin was loaded without error.
<b>0x100a918</b>	ADMIN	system.moab	INFO	MWM_JOB_NO_QUEUE_TIME	No QueueTime has been specified for job.	Configure the job with a queue time.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x100a922</b>	ADMIN	system.moab	INFO	MWM_FIND_PEER	Cannot find client peer for job %s (Name: %s).	The resource manager cannot be located.
<b>0x100a95b</b>	ADMIN	system.moab	INFO	MWM_RSV_FULL	Full reservation '%s' reserved %s procs in partition '%s' to start in %s at (%s).	The full reservation has been reserved.
<b>0x100a97d</b>	ADMIN	system.moab	INFO	MWM_RSV_PREREQ_JOB	Cannot create reservation for prerequisite job '%s'.	Could not obtain a reservation for this job.
<b>0x100a97e</b>	ADMIN	system.moab	INFO	MWM_ANNOTATE_JOB	Cannot annotate job '%s' with message '%s'.	Unable to modify the job with the annotation.
<b>0x100a980</b>	ADMIN	system.moab	INFO	MWM_UPDATE_JOB	Cannot update job '%s'.	The update on the job from XML failed.
<b>0x100a981</b>	ADMIN	system.moab	INFO	MWM_REMAP_CLASS	Cannot remap class for RM job '%s' (%s).	Unable to modify the job with the new class.
<b>0x100a983</b>	ADMIN	system.moab	INFO	MWM_COMPLETING_JOB	Completing job '%s'.	The job finished.
<b>0x100a984</b>	ADMIN	system.moab	INFO	MWM_JOBS_DETECTED	There were %s %s jobs detected on RM '%s'.	The resource manager reported these jobs.
<b>0x1</b>	AD	system	IN	MWM_	Cannot suspend job '%s'	Check the PBS

Cod e	Es calati on Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
<b>00a 985</b>	MI N	m.mo ab	F O	SUSPEND_ JOB	(%s).	server log to see reason of failure.
<b>0x1 00a 986</b>	AD MI N	syste m.mo ab	IN F O	MWM_ STALE_ PARTITION	Attempting to remove stale partition for completed job '%s'.	About to perform the stated operation.
<b>0x1 00a 987</b>	AD MI N	syste m.mo ab	IN F O	MWM_ STALE_ PARTITION_ SUCCESS	Successfully removed stale partition for completed job '%s'.	Successfully performed the stated operation.
<b>0x1 00a 988</b>	AD MI N	syste m.mo ab	IN F O	MWM_ CANCEL_ NOQUEUE_ JOB	Canceling No-queue job '%s'.	About to perform the stated operation.
<b>0x1 00a 989</b>	AD MI N	syste m.mo ab	IN F O	MWM_ SIGNAL_ JOB	Cannot signal job '%s' (%s).	The resource manager did not respond to the signal request.
<b>0x1 00a 98a</b>	AD MI N	syste m.mo ab	IN F O	MWM_ RSV_ JOB_ CREDS	Cannot set up reservation job credentials.	The user, account, or group credentials might not be valid.
<b>0x1 00a 98b</b>	AD MI N	syste m.mo ab	IN F O	MWM_ CP_ CORRUPT_ NODE_ LINE	Corrupt node line detected (%s).	The line does not contain the correct syntax for a checkpoint.
<b>0x1</b>	AD	syste	IN	MWM_	Evaluating reservation '%s'.	About to

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>00a98c</b>	MIN	m.moab	FO	EVALUATING_RSV		perform the stated operation.
<b>0x100a98d</b>	ADMIN	system.moab	INFO	MWM_EXPIRING_CHECKPOINT_DATA	Expiring checkpoint data for %s '%s'. Not updated in %s.	The object's checkpoint data has expired.
<b>0x100a98e</b>	ADMIN	system.moab	INFO	MWM_JOB_PREVIOUSLY_REMOVED	Job '%s' was previously removed.	The job has already been removed.
<b>0x100a98f</b>	ADMIN	system.moab	INFO	MWM_JOB_STARTED_BY_USER	Job '%s' was started by user '%s'.	The job started.
<b>0x100a990</b>	ADMIN	system.moab	INFO	MWM_JOB_NOT_STARTED_BY_USER	Job '%s' could not be started by user '%s' (%s).	The job could not be started.
<b>0x100a991</b>	ADMIN	system.moab	INFO	MWM_RM_JOB_NOT_STARTED	Job '%s' could not be started with %s RM '%s' (%s).	The job could not be started.
<b>0x100a992</b>	ADMIN	system.moab	INFO	MWM_JOB_CANCELED_EXTERNALLY	Job '%s' appears to have been canceled externally.	The job was canceled.
<b>0x100a993</b>	ADMIN	system.moab	INFO	MWM_JOB_COMPLETED_SINGLE_ITERATION	Job '%s' appears to have been started and completed in a single iteration.	The job completed.
<b>0x100a994</b>	ADMIN	system.moab	INFO	MWM_JOB_PROCESSING_G_	Job processing completed.	The jobs have been processed.

Cod e	Es calati on Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
				COMPLETE D		
<b>0x100a995</b>	AD MI N	syste m.mo ab	IN F O	MWM_ PROCESSIN G_JOB	Processing job '%s' in state '%s'.	Processing a single job.
<b>0x100a996</b>	AD MI N	syste m.mo ab	IN F O	MWM_JOB_ SUSPENDED	Job '%s' suspended through %s RM.	The resource manager suspended the job.
<b>0x100a997</b>	AD MI N	syste m.mo ab	IN F O	MWM_JOB_ RESUMED	Job '%s' resumed through %s RM.	The resource manager resumed the job.
<b>0x100a998</b>	AD MI N	syste m.mo ab	IN F O	MWM_JOB_ FEASIBILIT Y_CHECK_ DISABLED	Job feasibility check disabled (env).	This feature has been disabled.
<b>0x100a999</b>	AD MI N	syste m.mo ab	IN F O	MWM_JOB_ USAGE_ SENT	Job usage sent for job '%s'.	The usage sent as feedback to user.
<b>0x100a99a</b>	AD MI N	syste m.mo ab	IN F O	MWM_ LOADING_ JOBS	Loading %s job(s).	The jobs are about to be loaded.
<b>0x100a99b</b>	AD MI N	syste m.mo ab	IN F O	MWM_ LOADING_ NODE_ RECORDS	Loading %s node record(s).	The node records are about to be loaded.
<b>0x100a99c</b>	AD MI N	syste m.mo ab	IN F O	MWM_ LOADED_ WORKLOA D_BUFFER	Loaded %s workload buffer (%s bytes), processing jobs.	The workload buffer was loaded.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x100a99d</b>	ADMIN	system.moab	INFO	MWM_JOB_REJECTED_INFINITE_WALLTIME	Job '%s' rejected (requested infinite walltime).	Jobs must have a walltime limit.
<b>0x100a99e</b>	ADMIN	system.moab	INFO	MWM_JOB_REJECTED_PARTITION	Job '%s' rejected in partition %s (exceeds maximum task size: %s > %s).	Adjust JOBMAXTASKCOUNT in the configuration file.
<b>0x100a99f</b>	ADMIN	system.moab	INFO	MWM_JOB_ALREADY_EXISTS	Job '%s' already exists but is not a duplicate.	The ID of the job matched a completed job.
<b>0x100a9a0</b>	ADMIN	system.moab	INFO	MWM_JOB_ALREADY_BATCH_HOLD	Job '%s' is already on batch hold.	Trying to place a job on hold that is already in that state.
<b>0x100a9a1</b>	ADMIN	system.moab	INFO	MWM_JOB_REQUESTS_RSV	Job '%s' requests reservation '%s' (not deferring).	The job is requesting the reservation.
<b>0x100a9a2</b>	ADMIN	system.moab	INFO	MWM_RM_CONNECTION_FAILED	Connection to RM '%s' failed. Not deferring job '%s' (Reason: %s).	Refer to the reason message.
<b>0x100a9a3</b>	ADMIN	system.moab	INFO	MWM_DEFER_DISABLED	Defer disabled.	The job cannot be deferred.
<b>0x100a9a4</b>	ADMIN	system.moab	INFO	MWM_MWS_CLUSTER_QUERY	Cluster query retrieval failed for MWS RM '%s'.	The resource manager did not respond to the request.
<b>0x100a</b>	ADMIN	system.moab	INFO	MWM_JOB_INVALID_PARTITION	Job '%s' specifies an invalid partition.	The job must reference a

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
9a5	N	ab	O	PARTITION		valid partition.
0x100a9a6	ADMIN	system.moa	INFO	MWM_JOB_INVALID_QOS	Cannot set QoS on job '%s' to '%s' - invalid QoS.	The job must use a valid QoS.
0x100a9a7	ADMIN	system.moa	INFO	MWM_JOB_INVALID_ACCOUNT	Cannot set account on job '%s' to '%s' - invalid account (%s).	The job must use a valid account.
0x100a9a8	ADMIN	system.moa	INFO	MWM_CHECKING_IDLE_JOB	Checking idle job '%s' (priority: %s) partition %s.	Checking the job.
0x100a9a9	ADMIN	system.moa	INFO	MWM_CHECKING_SUSPENDED_JOB	Checking suspended job '%s' (priority: %s).	Checking the job.
0x100a9aa	ADMIN	system.moa	INFO	MWM_CHECKPOINT_TEST_ENABLED	Checkpoint test enabled (env).	The feature has been enabled with an environment variable.
0x100a9ab	ADMIN	system.moa	INFO	MWM_ADD_NODE_FAILED	Could not add node because MNodeAdd failed.	The node could not be added to the object.
0x100a9ac	ADMIN	system.moa	INFO	MWM_ATTEMPTING_RESERVATION	Attempting reservation of %s procs in %s for %s.	The scheduler will try to make the reservation.
0x100a	ADMIN	system.moa	INFO	MWM_FAIRSHARE_	Fairshare rolled to interval %s.	The interval has changed.



Cod e	Es calation Level	Topi c	S e verity	Event Name	Message Template	Comment
<b>9ad</b>	N	ab	O	INTERVAL		
<b>0x100a9ae</b>	ADMIN	system.moab	INFO	MWM_INVALID_ARCHITECTURE	Invalid architecture.	The architecture is not a valid value.
<b>0x100a9af</b>	ADMIN	system.moab	INFO	MWM_INVALID_PSEUDOJOB	Invalid pseudo-job.	The pseudo-job is not a valid value.
<b>0x100a9b0</b>	ADMIN	system.moab	INFO	MWM_HOLD_TYPE	Hold type '%s' selected.	The given hold type was specified.
<b>0x100a9b1</b>	ADMIN	system.moab	INFO	MWM_MESSAGE_SENT	Message sent to server.	The message was sent.
<b>0x100a9b2</b>	ADMIN	system.moab	INFO	MWM_JOB_LOCATED	Located job '%s' in partition '%s' reserved to start %s.	The specified job has been located.
<b>0x100a9b3</b>	ADMIN	system.moab	INFO	MWM_TOTAL_JOBS_DETECTED	Total jobs detected: %s.	Number of counted jobs returned from the workload query.
<b>0x100a9b4</b>	ADMIN	system.moab	INFO	MWM_NO_WORKLOAD_DETECTED	No workload reported by any RM.	No jobs were reported across all the resource manager queries.
<b>0x100a9b5</b>	ADMIN	system.moab	INFO	MWM_LOADING_JOB	Loading job '%s' in state '%s' (%s bytes).	The job is being loaded.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x100a9b6</b>	ADMIN	system.moab	INFO	MWM_JOB_START_REJECTED	Local constraints rejected the starting of job '%s'.	The job cannot start.
<b>0x100a9b7</b>	ADMIN	system.moab	INFO	MWM_INVALID_STAT_TYPE	Invalid stat type '%s' requested.	Not a valid value.
<b>0x100a9b8</b>	ADMIN	system.moab	INFO	MWM_ORPHAN_PARTITION	Creating temporary job to process orphan partition '%s' for job '%s'.	The job was not found in active or completed job tables.
<b>0x100a9b9</b>	ADMIN	system.moab	INFO	MWM_DISABLE_ACTION_PROGRAM	Disabling action program '%s'.	An invalid action program was requested.
<b>0x100a9ba</b>	ADMIN	system.moab	INFO	MWM_DISABLE_JOB_FB_PROGRAM	Disabling job feedback program '%s' (%s).	An invalid job feedback program was requested. See documentation for <a href="#">FEEDBACKPROGRAM</a> .
<b>0x100a9bb</b>	ADMIN	system.moab	INFO	MWM_CP_RESTART_STATE_IGNORED	Checkpoint restart state '%s' ignored.	The restart state specified is being ignored.
<b>0x100a9bc</b>	ADMIN	system.moab	INFO	MWM_CP_RESTART_STATE_SUCCESS	Starting scheduler with checkpoint restart state '%s'.	The restart state specified is being used.
<b>0x1</b>	AD	system	IN	MWM_	Destroying node '%s'.	The specified

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>00a9bd</b>	MIN	m.moab	FO	DESTROYING_NODE		node is being destroyed.
<b>0x100a9be</b>	ADMIN	system.moab	INFO	MWM_IGNOREING_NODE	Ignoring node '%s'.	The specified node is being ignored.
<b>0x100a9c0</b>	ADMIN	system.moab	INFO	MWM_CANNOT_ADJUST_JOB_HOLDS	Cannot adjust holds on remote peer for job '%s' (%s).	Unable to modify the job.
<b>0x100a9c1</b>	ADMIN	system.moab	INFO	MWM_JOB_CANNOT_CREATE_RESERVATION	Cannot create reservation for job '%s' (previously reserved to start in %s)).	Failed to create reservation for job.
<b>0x100a9c2</b>	ADMIN	system.moab	INFO	MWM_TRIGGER_LOAD_OUTPUT	Cannot load output data for trigger '%s' (File: %s).	The file might not exist or might be inaccessible.
<b>0x100a9c3</b>	ADMIN	system.moab	INFO	MWM_PBS_SERVER_CONNECT	Connected to PBS server %s:%s on sd %s.	Connection established.
<b>0x100a9c4</b>	ADMIN	system.moab	INFO	MWM_NO_JOB_DATA	No job data was sent by %s RM.	The data sent by the resource manager did not contain job information.
<b>0x100a9c5</b>	ADMIN	system.moab	INFO	MWM_JOB_RESUMED_WITH_PROCS	Job '%s' resumed on %s processors.	The resource manager resumed the job.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
0x100a9c6	ADMIN	system.moab	INFO	MWM_JOB_SINGALED	Job '%s' successfully signaled (action: %s, signal: %s).	The job responded to the signal request.
0x100a9c7	ADMIN	system.moab	INFO	MWM_JOB_CANCELED_RM	Job '%s' canceled through %s RM.	The job was canceled.
0x100a9c8	ADMIN	system.moab	INFO	MWM_JOB_ASSIGNED_DEFAULT_GROUP	Job '%s' assigned default group '%s'.	The job was modified.
0x100a9c9	ADMIN	system.moab	INFO	MWM_FILE_EXECUTE_PERMISSION	File '%s' does not have user execute permission (st_mode = %s).	The permissions must be modified.
0x100a9ca	ADMIN	system.moab	INFO	MWM_INSUFFICIENT_PREEMPT_JOBS	Inadequate preempt jobs (%s) located for %s job (P: %s of %s, N: %s of %s).	Not enough jobs could be preempted.
0x100a9cb	ADMIN	system.moab	INFO	MWM_READ_STAT_INDEX	Cannot read stat index for location %s:%s:%s.	The checkpoint did not have the stat information.
0x100a9cc	ADMIN	system.moab	INFO	MWM_BACKFULL_JOB_PREEMPT	Backfill job '%s' no longer preemptible (%s > %s) in partition '%s'.	The job cannot be preempted.
0x100a9cd	ADMIN	system.moab	INFO	MWM_STARTTIME	Cannot obtain desired starttime (%s != %s).	The job cannot be adjusted to the given start

Cod e	Es calati on Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
				UNAVAILAB LE		time.
<b>0x100a9ce</b>	AD MI N	system.mo ab	INFO	MWM_STARTTIME_ADJUSTED	Timeframe for reservation %s adjusted forward by %s seconds.	The reservation has been adjusted to the given start time.
<b>0x100a9cf</b>	AD MI N	system.mo ab	INFO	MWM_RESERVATION_ROLLBACK	Time: %s RollbackOffset: %s RsvStartTime: %s RsvDuration %s.	The reservation is being considered for rollback.
<b>0x100a9d0</b>	AD MI N	system.mo ab	INFO	MWM_RESERVATION_NOT_REQUIRED	Reservation '%s' not required for specified period.	The reservation is not required for this time period.
<b>0x100a9d1</b>	AD MI N	system.mo ab	INFO	MWM_RM_INTERFACE_RECOVERED	The interface for RM '%s' has been recovered.	A previously corrupt interface is now working.
<b>0x100a9d2</b>	AD MI N	system.mo ab	INFO	MWM_NTR_JOB_FOUND	Found an NTR (next to run) job - stopping idle job scheduling.	The job will now be run.
<b>0x100a9d3</b>	AD MI N	system.mo ab	INFO	MWM_GRES_KEYWORD_NO_VALUE	GRes keyword '%s' passed in with no value.	A value must be specified.
<b>0x100a9d4</b>	AD MI N	system.mo ab	INFO	MWM_JOB_EXTENSION_STRING	Job '%s' has invalid extension string - '%s'.	The system is unable to process the

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
						string.
<b>0x100a9d5</b>	ADMIN	system.moab	INFO	MWM_JOB_PROCESS_FAILURE	Job '%s' is invalid. It cannot be processed (%s).	There was an error loading the job. It will be rejected.
<b>0x100a9d6</b>	ADMIN	system.moab	INFO	MWM_JOB_MODIFIED_RM	Job '%s' has been modified through %s RM.	The job was modified.
<b>0x100a9d7</b>	ADMIN	system.moab	INFO	MWM_IDLE_BACKLOG_SIZE	Idle backlog: %s seconds (%s hours).	The idle backlog status is given.
<b>0x100a9d8</b>	ADMIN	system.moab	INFO	MWM_SET_RESOURCES	Inadequate resources found in any set (%s < %s).	None of the node sets have the resources needed.
<b>0x100a9d9</b>	ADMIN	system.moab	INFO	MWM_JOB_PREEMPTING_JOB	Job %s preempting job %s (statemtime: %s) (preempted this iteration: %s).	One job preempted another.
<b>0x100a9da</b>	ADMIN	system.moab	INFO	MWM_UPDATE_SCHEDULE_STATS	Iteration: %s; scheduling time: %s seconds.	Normal statistics update.
<b>0x100a9db</b>	ADMIN	system.moab	INFO	MWM_JOB_STARTED_RM	Job '%s' started through %s RM on %s procs.	The job has started.
<b>0x100a9dc</b>	ADMIN	system.moab	INFO	MWM_JOB_DELAY	Job delay: %s; reservation retry time: %s (StateDelayNC: %s; JobRsvDelayNC: %s).	The job has been delayed.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x100a9dd</b>	ADMIN	system.moab	INFO	MWM_JOB_COMPLETE	Job '%s' completed. X: %s; T: %s; PS: %s; A: %s (RM: %s/%s).	The job completed.
<b>0x100a9de</b>	ADMIN	system.moab	INFO	MWM_JOB_RESERVED_TASKS	Job '%s' reserved %s tasks (partition %s) to start in %s on %s (WC: %s).	The job has reserved the tasks.
<b>0x100a9df</b>	ADMIN	system.moab	INFO	MWM_EVENT_INTERFACE_ENABLED	Event interface enabled for wiki RM %s on port %s.	The interface is now functional.
<b>0x100a9e0</b>	ADMIN	system.moab	INFO	MWM_RM_RESOURCE_S_DETECTED	There were %s %s resources detected on RM '%s'.	The given resources were found.
<b>0x100a9e3</b>	ADMIN	system.moab	INFO	MWM_EXTENDING_RESERVATION	Extending reservation by %s seconds (trigger still active).	The reservation is being extended.
<b>0x100a9e4</b>	ADMIN	system.moab	INFO	MWM_EXTENDING_RESERVATION_OVERRUN_JOB	Extending reservation for overrun job '%s' by %s seconds.	The reservation is being extended.
<b>0x100a9e5</b>	ADMIN	system.moab	INFO	MWM_JOB_LOCATED_BESTFIT	Located bestfit job '%s' (size: %s; duration: %s).	Backfill found a job that best fits the available resources.
<b>0x1</b>	AD	system	IN	MWM_JOB_	The best partition for job	Backfill found

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
00a9e6	MIN	m.moab	FO	BEST_PARTITION	'%s' is '%s'.	a job that best fits the available resources.
0x100a9e7	ADMIN	system.moab	INFO	MWM_CPA_PARTITION_DESTROY	Destroying CPA partition '%s' for job '%s' with cookie %s (%s).	The partition is being destroyed.
0x100a9e8	ADMIN	system.moab	INFO	MWM_RESOURCE_S_LOCATED	Located resources for %s tasks (%s) in best partition '%s' for job '%s' at time offset %s.	The listed resources have been located.
0x100a9e9	ADMIN	system.moab	INFO	MWM_MINIMUM_EFFICIENCY_REACHED	Minimum efficiency reached (%s percent) on iteration %s.	The threshold has been reached.
0x100a9ea	ADMIN	system.moab	INFO	MWM_JOB_START_PARTITION	Cannot start job '%s' in partition '%s' (scheduler mode: %s).	The job could not be started.
0x100a9eb	ADMIN	system.moab	INFO	MWM_JOB_FEASIBLE_NODES	Inadequate feasible nodes found for job '%s':%s in partition '%s' (%s < %s).	The job could not be scheduled.
0x100a9ec	ADMIN	system.moab	INFO	MWM_JOB_LOADED	Job '%s' loaded: TC=%s UGC=%s,%s,%s WC=%s ST=%s %s %s.	The job was loaded.
0x100a9ed	ADMIN	system.moab	INFO	MWM_JOB_RSV_CREATE	Cannot create new reservation for job %s (shape[%s] %s).	Check the reservation time, nodes, and account.
0x100a9ee	ADMIN	system.moab	INFO	MWM_JOB_LOCATE_NODES	Cannot locate nodes for job '%s' req[%s] (%s additional needed).	Not enough nodes are available to



Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
						run the job.
<b>0x100a9ef</b>	ADMIN	system.moab	INFO	MWM_JOB_START_RM_DISABLED	Cannot start job '%s' since RM '%s' is disabled.	Not enough nodes are available to run the job.
<b>0x100a9f0</b>	ADMIN	system.moab	INFO	MWM_JOB_START_RESERVE_TIME	Cannot start job '%s' reserve time in %s.	The time to schedule has already arrived.
<b>0x100a9f1</b>	ADMIN	system.moab	INFO	MWM_ERROR_IN_EXE_STDERR	Error detected in '%s' due to presence of the word 'ERROR' in stderr (%s).	The executable failed.
<b>0x100a9f2</b>	ADMIN	system.moab	INFO	MWM_ERROR_IN_STDERR	Error detected due to presence of the word 'ERROR' in stderr.	The child process failed.
<b>0x100a9f3</b>	ADMIN	system.moab	INFO	MWM_CHECKJOB_STATE	Job '%s' State: %s Expected State: %s QueueTime: %s.	The job is in the listed state. The expected state might not be the same.
<b>0x100a9f4</b>	ADMIN	system.moab	INFO	MWM_JOB_NODELIST	Cannot obtain nodelist for job '%s':%s in range %s.	The nodes are not available.
<b>0x100a9f5</b>	ADMIN	system.moab	INFO	MWM_JOB_RESUME	Job '%s' cannot be resumed since allocated nodes are not available (node '%s' state '%s').	The resource manager resumed the job.
<b>0x100a9f6</b>	ADMIN	system.moab	INFO	MWM_CLEARING_EXPIRED_	Clearing expired %s reservation '%s' on iteration %s (start: %s end:	The reservation has expired.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
				RESERVATION	%s).	
<b>0x100a9f7</b>	ADMIN	system.moa	INFO	MWM_CPA_RETRY	CPA retry detected - will re-attempt partition creation in 2 seconds.	The partition may be created.
<b>0x100a9f9</b>	ADMIN	system.moa	INFO	MWM_JOBS_STARTED	There were %s %s jobs started in partition '%s' on iteration %s.	The jobs were started.
<b>0x100a9fa</b>	ADMIN	system.moa	INFO	MWM_TASKS_ALLOCATED	There were %s of %s tasks allocated for job '%s':%s.	The tasks were allocated.
<b>0x100a9fb</b>	ADMIN	system.moa	INFO	MWM_CLASSES_DETECTED	There were %s %s classes/queues detected on RM '%s'.	The classes were detected.
<b>0x100a9fd</b>	ADMIN	system.moa	INFO	MWM_JOB_DELAYED_RSV	Delayed reservation detected for reserved job '%s' (%s seconds) attempting squeeze.	Attempting to fit the job into the reservation.
<b>0x100a9fe</b>	ADMIN	system.moa	INFO	MWM_DUMPING_RESERVATIONS	Dumping reservations on iteration %s.	All the reservations will be dumped to the log.
<b>0x100a9ff</b>	ADMIN	system.moa	INFO	MWM_ALLOCPARTITION_MISSING	ALLOCPARTITION missing from completed job '%s' - restoring variable with value '%s'.	The value is being substituted.
<b>0x100aa00</b>	ADMIN	system.moa	INFO	MWM_RECEIVED_NODELIST	Received nodelist through %s RM.	The nodelist was received.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x100a01</b>	ADMIN	system.moab	INFO	MWM_SERVICE_REQUEST_FROM_HOST	Received service request from host '%s'.	The request was received.
<b>0x100a02</b>	ADMIN	system.moab	INFO	MWM_RECEIVED_WORKLOAD	Received workload info through %s RM '%s' (%s bytes).	The workload was received.
<b>0x100a03</b>	ADMIN	system.moab	INFO	MWM_RSV_REMOVED_FROM_CACHE	Removing reservation '%s' from cache.	The cached reservation is being removed.
<b>0x100a04</b>	ADMIN	system.moab	INFO	MWM_RECOVER_READ_SOCKET	RECOVER: attempting to read socket connection.	The recovery function is attempting to communicate via sockets.
<b>0x100a05</b>	ADMIN	system.moab	INFO	MWM_GREEDY_BACKFILL	Improved list found by greedy backfill in %s searches (utility: %s; processors available: %s).	The object is being removed from the database.
<b>0x100a06</b>	ADMIN	system.moab	INFO	MWM_RESERVATION_NAME_AND_GROUP	Name='%s' RsvGroup='%s'.	The object is being removed from the database.
<b>0x100a15</b>	ADMIN	system.moab	INFO	MWM_TRIGGERS_DISABLED	Triggers disabled. %s.	Triggers are disabled. This message indicates when this flag is being set, and when an action

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
						is being skipped because the flag is set.
<b>0x100a16</b>	ADMIN	system.moab	INFO	MWM_USER_NOT_AUTHORIZED	User %s is not authorized to %s.	This user does not have permissions to accomplish the listed task.
<b>0x100a1f</b>	ADMIN	system.moab	INFO	MWM_SUCCESSFULLY_OPENED_SOCKET	Opened service socket on port %s.	A socket was successfully opened listening on the remote port.
<b>0x100a21</b>	ADMIN	system.moab	INFO	MWM_NO_CHECKPOINT_INFO	No checkpoint information available for '%s'.	Checkpoint information was not available.
<b>0x100a22</b>	ADMIN	system.moab	INFO	MWM_NODE_INDEX_TABLE_ENABLED	Node index table enabled.	Enabled by environment variable: MOABUSENO DEINDEX.
<b>0x100a24</b>	ADMIN	system.moab	INFO	MWM_UNKNOWN_NODE_SLOT	Node slot not yet set on node '%s'.	Delaying setting rack until slot is known.
<b>0x100a25</b>	ADMIN	system.moab	INFO	MWM_NODE_STATUS	Node '%s' status: state='%s' rsvlist='%s' joblist='%s'.	General node status.
<b>0x1</b>	AD	system	IN	MWM_NO_	No jobs in queue.	There were no

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>00a a26</b>	MIN	m.moab	FO	JOBS_IN_QUEUE		jobs in the scheduler queue indicating the scheduler has nothing to process.
<b>0x1 00a a27</b>	ADMIN	system.moab	INFO	MWM_NO_NODE_DATA	No node data sent by %s RM.	The resource manager did not receive any node data in cluster query.
<b>0x1 00a a28</b>	ADMIN	system.moab	INFO	MWM_NO_PREEMPTIBLE_RESOURCES	No preemptible resources found for job %s (tc: %s; class: '%s'; qos: %s; priority: %s; partition %s).	Indicates the scheduler could not find any jobs for preemption.
<b>0x1 00a a29</b>	ADMIN	system.moab	INFO	MWM_NO_PRIORITY_RESERVATION_CREATED_FOR_POLICY	No priority reservations created for policy '%s' for job %s.	Job reservation for a specific policy was unable to be created.
<b>0x1 00a a2a</b>	ADMIN	system.moab	INFO	MWM_NO_QUEUES_DETECTED	No queues detected.	Resource manager attempted to obtain queue information. Check resource manager for configured queues.
<b>0x1 00a a2b</b>	ADMIN	system.moab	INFO	MWM_NOT_ADDING_RM	Not adding RM '%s'.	The partition is not adding

Cod e	Es calati on Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
						the specified resource manager. This situation is most common in grid configurations where resource manager names are similar.
<b>0x100aa30</b>	ADMI N	system.moab	INFO	MWM_ORPHAN_PARTITION_REPORTED_FOR_JOB	Orphan partition %s reported for job %s. %s.	The resource manager reported the partition as orphaned.
<b>0x100aa31</b>	ADMI N	system.moab	INFO	MWM_PARAMETER_CANNOT_BE_CHANGED	Parameter '%s' cannot be changed while Moab is running.	Configuration file must be changed, and Moab must be restarted.
<b>0x100aa4c</b>	ADMI N	system.moab	INFO	MWM_NODE_REMOVED	Dynamic node '%s' is being removed. RequestID = '%s', TTL = %s, Reason = '%s'	A dynamic node was removed due to one of the following reasons: TTL expiration, mnodectl deletion, node idle purge time, not reported in most recent cluster query.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x100a4d</b>	ADMIN	system.moab	INFO	MWM_NODE_ADDED	Node '%s' was newly reported in the last cluster query. RequestID = '%s', TTL = %s	A node was newly reported in the last cluster query.
<b>0x100a4e</b>	ADMIN	system.moab	INFO	MWM_NODE_IGNORED_TTL	Ignored dynamic node '%s' reported in cluster query with expired TTL. RequestID = '%s', TTL = %s	A dynamic node was reported in the last cluster query with an expired TTL - ignoring.
<b>0x100a4f</b>	ADMIN	system.moab	INFO	MWM_NODE_IDLE_PURGE_TIME_EXCEEDED	dynamic node '%s' exceeded node idle purge time. RequestID = '%s'	dynamic node exceeded node idle purge time
<b>0x100a50</b>	ADMIN	system.moab	INFO	MWM_NODE_UNREMOVED	Dynamic node '%s' is being re-added	A dynamic node was re-added
<b>0x100a51</b>	ADMIN	system.moab	INFO	MWM_NODE_TTL_MODIFIED	Node '%s' TTL modified. RequestID = '%s', Old TTL = %s, New TTL = %s	The node TTL has been modified.
<b>0x100a52</b>	ADMIN	system.moab	INFO	MWM_NODE_REQUESTID_MODIFIED	Node '%s' RequestID modified. Old RequestID = '%s', New RequestID = '%s'	The node RequestID has been modified.
<b>0x100a63</b>	ADMIN	system.moab	INFO	MWM_NO_MONGOREPLICASET_SPECIFIED	A set (replicaset) of Mongo servers (comma separated list) was given in MONGOSERVER, but MONGOREPLICASETNAME	Cannot connect to the Mongo server. A set

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
					is undefined.	(replicaset) of Mongo servers (comma separated list) was defined by MONGOSERVER, but MONGOREPLICASETNAME is undefined. Add MONGOREPLICASETNAME parameter to moab.cfg.
<b>0x100c3e8</b>	INTERNAL	telemetry.timing	INFO	MWM_PERFORMANCE_TIMER	File:%s,Function:%s,Line:%s,Pid:%s,Duration:%s	Microsecond performance timer.
<b>0x100c3e9</b>	INTERNAL	telemetry.timing	INFO	MWM_FUNCTION_START	File:%s,Function:%s,Line:%s,Pid:%s,TimeIn:%s	Function start time.
<b>0x100c3ea</b>	INTERNAL	telemetry.timing	INFO	MWM_FUNCTION_END	File:%s,Function:%s,Line:%s,Pid:%s,TimeOut:%s,Duration:%s	Function end time.
<b>0x100c3eb</b>	INTERNAL	telemetry.output	INFO	MWM_FUNCTION_OUTPUT	File:%s,Function:%s,Line:%s,Pid:%s	Function output.
<b>0x100c3f2</b>	INTERNAL	telemetry.cpu	INFO	MWM_CPU_MONITOR	Description:%s,Pid:%s,Ticks:%s,Percent:%s	CPU Utilization.



Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x100c3f3</b>	INTERNAL	telemetry.stats	INFO	MWM_JOB_INFORMATION	Phase:%s,Count:%s	Job Process Information.
<b>0x100c3fc</b>	INTERNAL	telemetry.stats	INFO	MWM_SCHEDULE_ITERATION	Schedule Phase:%s,Milliseconds:%s	Last Schedule Iteration.
<b>0x100e72f</b>	INTERNAL	system.moab	INFO	MWM_SOCKET_REMOTE_DISCONNECT	Reading from a socket failed. It appears the client disconnected, errno: %s (%s).	The recv() system call failed. Use the errno and associated message to determine possible causes.
<b>0x100e7f0</b>	INTERNAL	system.moab	INFO	MWM_VM_LINKED_TO_NEW_TRACKING_JOB	Setting VMTracking job for VM '%s' to job '%s'.	A VM is associated with a tracking job.
<b>0x10001f7</b>	USER	domain.life cycle	WARN	MWM_TRIGGER_FAILURE	Trigger %s has failed.	The named trigger has finished its action, but it returned with a failure status.
<b>0x100025e</b>	USER	domain.life cycle	WARN	MWM_VM_MIGRATE_END_ERROR	VM %s migration has finished with an error: (%s)	The named VM has finished its migration. There was a problem during the migration.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
						Additional information might be provided regarding the error specifics.
<b>0x11000261</b>	USER	domain.life cycle	WARN	MWM_VM_MIGRATE_SUBMIT	Failed to submit VM migration job for VM %s. (%s)	Failed to submit VM Migration job.
<b>0x11000262</b>	USER	domain.life cycle	WARN	MWM_VM_NO_FEASIBLE_NODES	Failed to find a feasible node/hypervisor on which to run VM %s. Check setup job %s for details.	The named VM has been submitted, but no node/hypervisor meets all requirements.
<b>0x11000384</b>	USER	system.moab	WARN	MWM_VM_LICENSE_ERROR	There is an error with the Moab license: (%s)	There was a licensing error. Additional information might be provided regarding the error specifics.
<b>0x1100284a</b>	USER	system.moab	WARN	MWM_BAD_COMMAND_LINE_FLAG	Unexpected flag detected: '%s'.	The command line syntax that was received contains an invalid flag. Check the documentation and retry.
<b>0x1</b>	USER	system	WARN	MWM_	Ignoring incorrect	Valid range is

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>1002858</b>	ER	m.moab	ARN	NODESETMAXUSAGE_FAILURE	NODESETMAXUSAGE value '%s'.	from 0.0 to 1.0 inclusive.
<b>0x11002883</b>	USER	system.moab	WARN	MWM_CANNOT_DESTROY_STATIC_RM	Trigger cannot destroy static RM.	A trigger cannot destroy a static resource manager. Refer to trigger 'destroy'.
<b>0x11002966</b>	USER	system.moab	WARN	MWM_RM_JOB_SUBMIT_FAILURE	RM %s job submit failed: %s.	Error while submitting the job to the resource manager.
<b>0x11002a0c</b>	USER	system.moab	WARN	MWM_CANNOT_MODIFY_RM_JOB	Cannot modify %s for RM job %s - '%s'.	The listed attribute of the job could not be changed.
<b>0x11004004</b>	POWER_USER	system.moab	WARN	MWM_TESTING_WARNING	Testing with argument1: %s. and argument2: %s and argument3: %s	Internal error for testing diagnostics.
<b>0x11008199</b>	ADMIN	system.moab	WARN	MWM_SCHED_SET_COUNTER	%s cannot be set lower than it's current value. %s < %s	A counter cannot be set lower than it's current value.
<b>0x1100a713</b>	ADMIN	system.moab	WARN	MWM_CANNOT_LOAD_FILE	Cannot load %s file %s - %s.	Failed to load a file into Moab. Make sure it exists and that permissions

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
						are correct.
<b>0x1100a71a</b>	ADMIN	system.moab	WARNING	MWM_FAILED_TO_WAIT_FOR_CHILD	Failed to wait for child, pid: %s, errno: %s (%s).	The wait() system call failed. Use the errno and associated message to determine possible causes.
<b>0x1100a71d</b>	ADMIN	system.moab	WARNING	MWM_CANNOT_CHMOD_FILE	Failure changing permissions of file: '%s' to mode:'%s', errno: %s (%s).	The chmod() system call failed. Use the errno and associated message to determine possible causes.
<b>0x1100a71e</b>	ADMIN	system.moab	WARNING	MWM_CANNOT_OPEN_FILE_WARNING	Cannot open %s file '%s', errno: %s (%s).	The fopen() system call failed. Use the errno and associated message to determine possible causes.
<b>0x1100a722</b>	ADMIN	system.moab	WARNING	MWM_CANNOT_WRITE_FILE_WARNING	Cannot write to file '%s', errno: %s (%s).	The fwrite() system call failed. Use the errno and associated message to determine possible

Cod e	Es cal ation Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
						causes.
<b>0x1 100 a72 3</b>	AD MI N	syste m.mo ab	W A R N	MWM_ CANNOT_ CLOSE_ FILE_ DESCRIPTO R	Cannot close file descriptor %s, errno: %s (%s).	The close() system call failed. Use the errno and associated message to determine possible causes.
<b>0x1 100 a72 4</b>	AD MI N	syste m.mo ab	W A R N	MWM_ CANNOT_ RENAME_ FILE	Failure renaming file '%s' to '%s', errno: %s (%s).	The rename() system call failed. Use the errno and associated message to determine possible causes.
<b>0x1 100 a72 6</b>	AD MI N	syste m.mo ab	W A R N	MWM_ CANNOT_ BIND_TO_ PORT	Cannot bind to port %s, errno: %s (%s).	The bind() system call failed. Use the errno and associated message to determine possible causes.
<b>0x1 100 a72 a</b>	AD MI N	syste m.mo ab	W A R N	MWM_ CANNOT_ SEND_TO_ SOCKET	Cannot send %s byte packet, errno: %s (%s).	The send() system call failed. Use the errno and associated message to determine possible

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
						causes.
<b>0x1100a72c</b>	ADMIN	system.moab	WARNING	MWM_FAILED_GETSOCKOPT_WARNING	Cannot get socket %s option, errno: %s (%s).	The getsockopt() system call failed. Use the errno and associated message to determine possible causes.
<b>0x1100a72d</b>	ADMIN	system.moab	WARNING	MWM_FAILED_SETSOCKOPT_WARNING	Cannot set socket %s option, errno: %s (%s).	The setsockopt() system call failed. Use the errno and associated message to determine possible causes.
<b>0x1100a738</b>	ADMIN	system.moab	WARNING	MWM_CANNOT_SET_UMASK	Failure setting umask on file '%s', errno: %s (%s).	The umask() system call failed. Use the errno and associated message to determine possible causes.
<b>0x1100a73a</b>	ADMIN	system.moab	WARNING	MWM_FAILED_FCNTL_WARNING	Cannot set %s option on file descriptor, errno: %s (%s).	The fcntl() system call failed. Use the errno and associated message to determine

Cod e	Es calati on Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
						possible causes.
<b>0x1100a73b</b>	ADMI N	system.moab	WARN	MWM_CANNOT_STAT_FILE_WARNING	Cannot get stats on file '%s', errno: %s (%s).	The stat() system call failed. Use the errno and associated message to determine possible causes.
<b>0x1100a73e</b>	ADMI N	system.moab	WARN	MWM_GET_HOSTNAME_CLIENT	Cannot get hostname of the client, errno: %s (%s).	The getnameinfo() system call failed. Use the errno and associated message to determine possible causes.
<b>0x1100a74d</b>	ADMI N	system.moab	WARN	MWM_CONFIG_FILE_NOT_FOUND_WARNING	Cannot locate configuration file '%s' in '%s'.	Check for the existence of this file.
<b>0x1100a753</b>	ADMI N	system.moab	WARN	MWM_RM_CONFIG_INVALID_VALUE	Invalid %s value '%s' specified for RM '%s'.	Check the line in the configuration file for the parameter.
<b>0x1100a754</b>	ADMI N	system.moab	WARN	MWM_RM_CONFIG_PROCESS_ATTR	Failed to process attribute '%s' for resource manager '%s'.	Check the line in the configuration file for the

Cod e	Es cal ation Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
						parameter.
<b>0x1100a755</b>	AD MI N	syste m.mo ab	W A R N	MWM_RM_ CONFIG_ ATTR	RM attribute '%s' not handled.	Check the line in the configuration file for the parameter.
<b>0x1100a756</b>	AD MI N	syste m.mo ab	W A R N	MWM_RM_ CONFIG_ TIMEOUT	Resource manager '%s' has a timeout of less than 50 ms.	Check the line in the configuration file for the parameter.
<b>0x1100a758</b>	AD MI N	syste m.mo ab	W A R N	MWM_ CONFIG_ PARAM_ DEFAULT_ VALUE	Configuration parameter '%s[%s]' was not assigned a value. Using default.	Check the line in the configuration file to see if this behavior is desired.
<b>0x1100a759</b>	AD MI N	syste m.mo ab	W A R N	MWM_ CONFIG_ PARAM_ INTEGER_ DEFAULT_ VALUE	Configuration parameter '%s[%s]' has a value '%s' that is not an integer. Using default.	Check the line in the configuration file for the integer value.
<b>0x1100a75a</b>	AD MI N	syste m.mo ab	W A R N	MWM_ CONFIG_ PARAM_ DOUBLE_ DEFAULT_ VALUE	Configuration parameter '%s[%s]' has a value '%s' that is not a double. Using default.	Check the line in the configuration file for the double value.
<b>0x1100a75b</b>	AD MI N	syste m.mo ab	W A R N	MWM_ CONFIG_ PARAM_ NULL_ VALUE	Configuration parameter '%s[%s]' has a NULL value.	Check the line in the configuration file.



Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x1100a75c</b>	ADMIN	system.moab	WARN	MWM_CONFIG_PARAM_INVALID	Configuration parameter '%s' has an invalid value.	Check the line in the configuration file for the attribute.
<b>0x1100a75e</b>	ADMIN	system.moab	WARN	MWM_CONFIG_PARAM_UNKNOWN	Configuration parameter '%s[%s]' is not defined.	Check the line in the configuration file for the undefined parameter.
<b>0x1100a75f</b>	ADMIN	system.moab	WARN	MWM_CONFIG_ATTR_EXTRACTION	Configuration parameter '%s[%s]' attribute value '%s' cannot be extracted.	Check the line in the configuration file for the attribute.
<b>0x1100a760</b>	ADMIN	system.moab	WARN	MWM_INVALID_CONFIG_LINE	Cannot process line '%s'.	Check the line syntax against the documentation.
<b>0x1100a763</b>	ADMIN	system.moab	WARN	MWM_UNKNOWN_ADMINCFG_PARAMETER	Unknown ADMINCFG parameter '%s'.	Check the syntax in the configuration file.
<b>0x1100a764</b>	ADMIN	system.moab	WARN	MWM_UNKNOWN_MID_ATTR	Unknown identity attribute '%s'.	Check the MIDCFG lines in the configuration file.
<b>0x1100a76</b>	ADMIN	system.moab	WARN	MWM_UNKNOWN_AM_ATTR	Unknown account manager attribute '%s'.	Check the AMCFG lines in

Cod e	Es calati on Level	Topi c	S e ver it y	Event Name	Message Template	Comment
5			N			the configuration file.
0x1100a766	ADMIN	system.moab	WARN	MWM_UNKNOWN_ATTRIBUTE_SPECIFIED	Unknown attribute '%s' specified for %s %s.	An error occurred while parsing the configuration for the listed object. The specified attribute is unknown or invalid.
0x1100a769	ADMIN	system.moab	WARN	MWM_MONGOSERVER_CONNECTION_FAILURE	Unable to connect to Mongo server '%s' (%s).	The program will continue to try and connect in the background.
0x1100a76e	ADMIN	system.moab	WARN	MWM_EVENT_QUERY_ODBC	Event querying is only supported with ODBC.	Check the USEDATABASE option.
0x1100a770	ADMIN	system.moab	WARN	MWM_DB_CONNECT	Cannot connect to DB-- falling back to file and memory-based storage (%s).	Verify that the database is running.
0x1100a772	ADMIN	system.moab	WARN	MWM_DATABASE_STATS	Unable to retrieve statistics from the database.	Verify that the database is running.
0x1100a774	ADMIN	system.moab	WARN	MWM_SERVER_CONNECTION	The system was unable to connect to the server %s:%s - attempting fallback server %s.	Make sure the server's address is correct and it

Cod e	Es cal ati on Le vel	Topi c	S e v er ity	Event Name	Message Template	Comment
				N_FAILED_ TRYING_ FALLBACK		is running.
<b>0x1 100 a77 5</b>	AD MI N	syste m.mo ab	W A R N	MWM_ PRIMARY_ SERVER_ FAILED_ TRYING_ BACKUP	The system was unable to connect to the server %s (%s:%s) - trying backup server (%s:%s).	Make sure the server's address is correct and it is running.
<b>0x1 100 a77 8</b>	AD MI N	syste m.mo ab	W A R N	MWM_ INVALID_ REDUCE_ CLIENTMAX CONNECTIO NS	Reducing CLIENTMAXCONNECTIONS to %s from %s not allowed during runtime.	Decreasing the value of CLIENTMAXCONNECTIONS cannot be done during runtime.
<b>0x1 100 a78 5</b>	AD MI N	syste m.mo ab	W A R N	MWM_ UNABLE_ TO_START_ JOB	Cannot start job %s. (%s)	The job failed to start.
<b>0x1 100 a78 8</b>	AD MI N	syste m.mo ab	W A R N	MWM_ UNABLE_ TO_ ALLOCATE_ NODES_ FOR_JOB	Cannot allocate nodes for job %s. (%s)	Cannot allocate a node list that matches the requirements for this job. This might not be serious since multiple passes may occur.
<b>0x1 100 a7d 2</b>	AD MI N	syste m.mo ab	W A R N	MWM_ UNABLE_ LOAD_PBS_ JOB	Cannot load PBS job '%s'.	Could not load a job discovered from a PBS

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
						resource manager into Moab.
<b>0x1100a7e3</b>	ADMIN	system.moab	WARNING	MWM_TRUNCATING_ATTRIBUTE_FOR_CLASS	Truncating %s for class: %s (rm reports: %s; Moab enforces: %s).	The resource manager reports a certain value for a class, but Moab has been instructed to keep it within certain limits. The value will be truncated to keep it within the limits.
<b>0x1100a7e4</b>	ADMIN	system.moab	WARNING	MWM_UNEXPECTED_JOB_STATE	Unexpected job state '%s' detected for job %s.	The listed job was found to be in a state that was not expected. This may or may not be an error condition.
<b>0x1100a821</b>	ADMIN	system.moab	WARNING	MWM_CLOCK_SKEW_DETECTED	Clock skew detected (%s time for job %s in %s).	A reported time associated with the job appears to be wrong. This could be because of a lack of synchronization between system clocks

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
						on all nodes.
<b>0x1100a83e</b>	ADMIN	system.moab	WARN	MWM_INVALID_WIKI_ATTRIBUTE	Encountered invalid wiki attribute while reading '%s'.	Check the syntax of the attribute.
<b>0x1100a83f</b>	ADMIN	system.moab	WARN	MWM_DUPLICATE_WIKI_ATTRIBUTE	Wiki attribute '%s' is already set.	Check for duplicate instances of the attribute.
<b>0x1100a840</b>	ADMIN	system.moab	WARN	MWM_VM_UNSUPPORTED_WIKI_ATTRIBUTE	Wiki attribute '%s' is unsupported for VM creation.	Remove the attribute.
<b>0x1100a843</b>	ADMIN	system.moab	WARN	MWM_ADD_NODE_FAILURE	Cannot add node '%s' to global node table. Index is already used.	Cannot have two nodes with the same name.
<b>0x1100a844</b>	ADMIN	system.moab	WARN	MWM_HT_ADD_NODE_FAILURE	Cannot add node '%s' to hash table. Index is already used.	Cannot have two nodes with the same name.
<b>0x1100a846</b>	ADMIN	system.moab	WARN	MWM_VM_MIGRATION_FAILURE	Cannot migrate VM '%s'.	The VM might not be eligible for migration.
<b>0x1100a84c</b>	ADMIN	system.moab	WARN	MWM_UNEXPECTED_SUBCOMMAND_RECEIVED	Unexpected subcommand '%s' received.	The communication from a Moab client includes an unknown subcommand.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x1100a84d</b>	ADMIN	system.moab	WARN	MWM_UNABLE_TO_REGISTER_JOB_AM	Unable to register job %s with accounting manager for job %s. reason: '%s' message: '%s'.	The accounting manager was unable to register the listed job for a certain action. An optional reason and/or message might be given to assist in diagnosis.
<b>0x1100a851</b>	ADMIN	system.moab	WARN	MWM_IMPROPER_VM_MIGRATION_DECISION	The migration decision for the VM was not properly set up.	The information indicating the destination node is missing.
<b>0x1100a85b</b>	ADMIN	system.moab	WARN	MWM_JOB_UPDATE_NULL_STARTTIME	Start time is NULL for job update.	Specify a start time that is greater than zero.
<b>0x1100a85c</b>	ADMIN	system.moab	WARN	MWM_JOB_UPDATE_NULL_DISPATCHTIME	Dispatch time is NULL for job update.	Specify a dispatch time that is greater than zero.
<b>0x1100a862</b>	ADMIN	system.moab	WARN	MWM_UNABLE_TO_REGISTER_RESERVATION_AM	Unable to register reservation %s with accounting manager for %s processors for reservation %s.	The accounting manager was unable to register the listed reservation for a certain action.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x1100a863</b>	ADMIN	system.moab	WARN	MWM_DEPRECATED_PARAMETER_VALUE	Deprecated value '%s' specified for parameter '%s'. %s	The listed value is no longer valid for this parameter. A hint might be provided with the message. Check the most recent documentation for the software version.
<b>0x1100a864</b>	ADMIN	system.moab	WARN	MWM_INVALID_FILE_ATTRIBUTES_WARNING	Invalid value '%s' specified for %s (%s).	Checking a file to see whether it exists, is executable, etc, has produced unexpected results.
<b>0x1100a865</b>	ADMIN	system.moab	WARN	MWM_CANNOT_SET_JOB_ATTRIBUTE_VIA_TEMPLATE_WARNING	Cannot set %s %s via template %s.	Failed to set the listed attribute to the listed value for a specified job template.
<b>0x1100a866</b>	ADMIN	system.moab	WARN	MWM_KILL_PROCESS_FAILURE	Unable to kill process %s.	The system tried to kill the given process and failed.
<b>0x1100a86a</b>	ADMIN	system.moab	WARN	MWM_NODELIST_STRING_BUFFER	Insufficient buffer space to convert a node list into a string.	The buffer must be larger to hold all the nodes.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x1100a86b</b>	ADMIN	system.moab	WARN	MWM_MAX_NODES_EXCEEDED	The maximum number of nodes associated with a reservation has been exceeded.	The number of nodes must be reduced.
<b>0x1100a871</b>	ADMIN	system.moab	WARN	MWM_VC_ALREADY_ADDED	The virtual container '%s' is already an ancestor of VC '%s'.	Cannot create a circular chain, must maintain a hierarchical structure.
<b>0x1100a872</b>	ADMIN	system.moab	WARN	MWM_VC_REMOVAL_FAILURE	The virtual container '%s' cannot be removed.	This is an internal error.
<b>0x1100a874</b>	ADMIN	system.moab	WARN	MWM_RESERVATION_JOB_NOT_FOUND	Unable to find the job for reservation '%s'.	The host job for the reservation is NULL.
<b>0x1100a875</b>	ADMIN	system.moab	WARN	MWM_SINGLE_USE_RESERVATION_DESTRUCTION	Unable to destroy a single-use reservation.	This is an internal error.
<b>0x1100a878</b>	ADMIN	system.moab	WARN	MWM_UNEXPECTED_JOB_SUBMISSION_POLICY	The system encountered an unexpected job submission policy (%s).	The job submission policy did not match a defined policy.
<b>0x1100a87</b>	ADMIN	system.moab	WARN	MWM_SIMULATION_JOB_	Unable to simulate workload by creating job records (1000 attempts).	The system might be low on memory.



Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>9</b>			N	RECORDS		
<b>0x1100a880</b>	ADMIN	system.moa	WARN	MWM_JOB_TRANSITION_XML_MESSAGE	Unable to add messages to job '%s' transition XML.	The system might be low on memory.
<b>0x1100a884</b>	ADMIN	system.moa	WARN	MWM_PBS_API_STALE	PBS API is stale - re-initializing.	Re-initializing the PBS environment.
<b>0x1100a885</b>	ADMIN	system.moa	WARN	MWM_UNABLE_TO_GET_PBS_QUEUE_INFO	Cannot process PBS queue info for RM %s (node %s) - no data available.	Unable to get any information on the PBS queues. Make sure that there was at least a queue set up in PBS.
<b>0x1100a886</b>	ADMIN	system.moa	WARN	MWM_UNABLE_TO_SET_JOB_ATTRIBUTE	Cannot set job '%s' attribute '%s:%s' to '%s' (rc: %s; '%s').	There was a problem while changing the job attribute and the error status was displayed in rc.
<b>0x1100a887</b>	ADMIN	system.moa	WARN	MWM_UNABLE_TO_CONNECT_PBS_SCHEDULER	Cannot connect to PBS event/scheduler port %s.	Ensure the PBS scheduler is running and listening on the specified port.
<b>0x1100a88</b>	ADMIN	system.moa	WARN	MWM_NODE_UNUSABLE_	Idle node %s is unusable (inadequate disk space in /var).	Ensure that the node has sufficient disk

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>8</b>			N	NO_DISK		space.
<b>0x1100a889</b>	ADMIN	system.moab	WARN	MWM_NODE_UNUSABLE_BAD_STATE	Node '%s' is unusable in state 'NONE'.	The node has become unusable because of its state being NONE.
<b>0x1100a88a</b>	ADMIN	system.moab	WARN	MWM_UNABLE_TO_FIND_USERS_GID	Cannot locate OS GID information for user '%s' - ignoring user.	Moab was unable to find the GID of this user. Make sure that this user has a GID.
<b>0x1100a88b</b>	ADMIN	system.moab	WARN	MWM_UNABLE_TO_FIND_USERS_UID	Cannot locate OS information for user '%s' - ignoring user - %s.	Moab was unable to find the user on the system. Make sure that this user exists.
<b>0x1100a88c</b>	ADMIN	system.moab	WARN	MWM_UNABLE_TO_FIND_GID_LIST	Cannot locate OS group list information for user '%s' - ignoring user.	Moab was unable to find the group list for this user.
<b>0x1100a88d</b>	ADMIN	system.moab	WARN	MWM_TIMEOUT	Command '%s' timed out, or wait failed after %s seconds.	Increasing the TIMEOUT settings in moab.cfg might help.
<b>0x1100a88e</b>	ADMIN	system.moab	WARN	MWM_INSUFFICIENT_POLICIES	Insufficient policies specified; hpolicy=%s,spolicy=%s.	Please revise your policies along with their actions.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x1100a88f</b>	ADMIN	system.moab	WARN	MWM_NO_STDOUT	Request succeeded with no stdout but stderr='%s'.	Typically there will also be stdout when there is stderr. Depending on the request this might be the intended result.
<b>0x1100a890</b>	ADMIN	system.moab	WARN	MWM_CANNOT_CONNECT_WIKI	Cannot connect to Wiki event port %s.	Failure to connect to the Wiki event port.
<b>0x1100a891</b>	ADMIN	system.moab	WARN	MWM_MISSING_COLON_STR	Colon delimiter not located in %s wiki string '%s..' in %s.	Check that the string contains the right format.
<b>0x1100a892</b>	ADMIN	system.moab	WARN	MWM_ADD_DEPENDENCY_FAIL	Failed to add dependencies to job %s's submission.	There was a problem in adding the dependencies to the job.
<b>0x1100a893</b>	ADMIN	system.moab	WARN	MWM_INVALID_PRIORITY_FUNCTION	Invalid priority function '%s' on job '%s'.	The priority function applied to the job was invalid.
<b>0x1100a894</b>	ADMIN	system.moab	WARN	MWM_MISSING_JOB_REQ	Invalid job '%s'; no requirements.	The job was invalid because it was missing the requirements.
<b>0x1</b>	ADMIN	system.moab	W	MWM_	Invalid job %s; no requirement at index %s.	The job was

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
100a895	N	ab	ARN	MISSING_JOB_REQ_AT_INDEX		invalid because an index was missing requirements.
0x1100a896	ADMIN	system.moab	WARN	MWM_INVALID_WIKI_STRING_MISSING_EQUAL	Malformed wiki string '%s' - no '='.	The wiki string was missing an equal sign '='.
0x1100a897	ADMIN	system.moab	WARN	MWM_INVALID_EMPTY_WIKI_STRING	Malformed wiki string '%s' - EOF.	The wiki string was empty.
0x1100a898	ADMIN	system.moab	WARN	MWM_INCORRECT_STAGE_LOCATION	stage-data source location is being incorrectly reported via wiki '%s' != '%s'.	The stage data source location was incorrectly reported in wiki.
0x1100a899	ADMIN	system.moab	WARN	MWM_INVALID_VM_OBJECT_ID	VM '%s' is not a valid object, ignoring.	The VM does not have a valid object ID.
0x1100a89a	ADMIN	system.moab	WARN	MWM_INVALID_JSON_CLUSTER	Could not parse JSON cluster query data from MWS RM (%s): %s.	The JSON construct(s) for the cluster might contain some invalid syntax.
0x1100a89b	ADMIN	system.moab	WARN	MWM_INVALID_JSON_WORKLOAD	Could not parse JSON workload query data from MWS RM (%s): %s.	The JSON construct(s) for the workload might contain some invalid

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
						syntax.
<b>0x1100a89c</b>	ADMIN	system.moab	WARN	MWM_MISSING_REQUIRED_PROPERTIES	JSON cluster query data from MWS RM (%s) does not contain required properties (%s, %s, %s).	The JSON constructs for the cluster query data are missing the required properties.
<b>0x1100a89d</b>	ADMIN	system.moab	WARN	MWM_INVALID_JSON_CLUSTER_OBJECT	JSON cluster query data from MWS RM (%s) is not a valid object.	Review the JSON construct for the cluster query data to ensure its syntax is correct.
<b>0x1100a89e</b>	ADMIN	system.moab	WARN	MWM_INVALID_JSON_WORKLOAD_OBJECT	JSON workload query data from MWS RM (%s) is not a valid object.	Review the JSON construct for the workload query data to ensure its syntax is correct.
<b>0x1100a89f</b>	ADMIN	system.moab	WARN	MWM_EMPTY_RESPONSE	Empty %s response from RM (%s).	The response from the resource manager query was empty.
<b>0x1100a8a0</b>	ADMIN	system.moab	WARN	MWM_INVALID_NODE_DATA	Nodes data from MWS RM (%s) is not a valid object.	The response from the Moab Web Services resource manager query was empty.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x1100a8a1</b>	ADMIN	system.moab	WARN	MWM_CREATE_RESERVATION_FAIL	Cannot create requested reservation (%s).	The request to create the given reservation has failed.
<b>0x1100a8a2</b>	ADMIN	system.moab	WARN	MWM_LOCATE_RSVPROFILE_FAIL	Cannot locate RSVPROFILE '%s'.	Moab failed to find the given RSVPROFILE. Confirm that the file exists.
<b>0x1100a8a3</b>	ADMIN	system.moab	WARN	MWM_LOCATE_RSV_PARENT_FAIL	Cannot locate parent '%s' for reservation '%s'.	Moab failed to locate the parent of the given reservation.
<b>0x1100a8a4</b>	ADMIN	system.moab	WARN	MWM_LOCATE_COMMAND_FAIL	Cannot locate command '%s'.	Moab failed to locate the given command. Confirm that the command exists.
<b>0x1100a8a5</b>	ADMIN	system.moab	WARN	MWM_UNSUPPORTED_SCHED_CMD	Received unexpected sched command '%s'.	Received an unexpected mschedctl command. Confirm that the used option is supported.
<b>0x1100a8a6</b>	ADMIN	system.moab	WARN	MWM_UNSUPPORTED_EVENT	Unsupported event '%s' from RM '%s'.	The given event is not supported by the given resource

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
						manager.
<b>0x1100a8a7</b>	ADMIN	system.moab	WARN	MWM_VM_MIGRATE_FAIL	VM %s should migrate from node %s but cannot locate valid destination - %s (policy).	Attempt to migrate the given VM from the given node failed. Check that the destination is valid.
<b>0x1100a8a8</b>	ADMIN	system.moab	WARN	MWM_JOB_START_FAIL	Start of system job %s failed; no action specified.	Failed to start a job because there was no action specified.
<b>0x1100a8a9</b>	ADMIN	system.moab	WARN	MWM_VMTRACKING_JOB_FAIL	VM '%s' reported a system job failure on VMTracking job '%s'.	The given VM reported it failed on the given VMTracking job.
<b>0x1100a8aa</b>	ADMIN	system.moab	WARN	MWM_VMTRACKING_EXCEEDED_WALTIME	VM '%s' exceeded its allocated walltime. VMTracking job '%s' (pointing to job '%s').	The given VM has exceeded its allocated walltime on the associated VM tracking job.
<b>0x1100a8ac</b>	ADMIN	system.moab	WARN	MWM_UNKNOWN_POWER_STATE	No RM can report node '%s' power state for system job '%s'.	No resource manager can report the power state for the given nodes on the given job.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x1100a8ad</b>	ADMIN	system.moab	WARN	MWM_ADD_GLOBAL_NODE_FAIL	Cannot add global node '%s'.	Failed to add the given global node.
<b>0x1100a8ae</b>	ADMIN	system.moab	WARN	MWM_UNKNOWN_CLIENT	Client ID '%s' is unknown.	Moab failed to recognize the name/ID of the given client.
<b>0x1100a8af</b>	ADMIN	system.moab	WARN	MWM_JOB_DEBIT_ACCOUNT	Unable to charge funds for job.	The account manager failed to debit the account for the job.
<b>0x1100a8b0</b>	ADMIN	system.moab	WARN	MWM_JOB_RESERVE_ACCOUNT	Unable to reserve funds for job (Reason: %s).	The account manager failed to reserve funds on the account for the job.
<b>0x1100a8b1</b>	ADMIN	system.moab	WARN	MWM_CANCEL_LEIN	Unable to cancel lien for instance '%s' (Reason: %s).	The account manager failed to release the lien.
<b>0x1100a8b2</b>	ADMIN	system.moab	WARN	MWM_RESERVATION_RESERVE_ACCOUNT	Unable to reserve funds for reservation (Reason: %s).	The account manager failed to reserve funds on the account for the reservation.
<b>0x1100a8b3</b>	ADMIN	system.moab	WARN	MWM_JOB_TASK_DISTRIBUTION	The system cannot distribute the tasks allocated for a job.	Check the tasks specified in the job.



Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x1100a8b4</b>	ADMIN	system.moab	WARN	MWM_JOB_DEFAULT_CLASS	Job cannot run with default class '%s'.	Check the limits set on the class.
<b>0x1100a8b5</b>	ADMIN	system.moab	WARN	MWM_RM_START_JOB	Cannot start job through a resource manager.	The resource manager might not be set to run the job.
<b>0x1100a8b6</b>	ADMIN	system.moab	WARN	MWM_JOB_DEPENDENCY_UPDATE	Cannot find job '%s' to update dependency '%s' for job '%s'.	The dependency job for the specified job is missing.
<b>0x1100a8b7</b>	ADMIN	system.moab	WARN	MWM_EXPIRED_CHECKPOINT	The checkpoint has expired.	Items within the checkpoint might no longer be valid.
<b>0x1100a8b8</b>	ADMIN	system.moab	WARN	MWM_BAD_CHECKPOINT_LINE	The system encountered an incorrectly formed checkpoint line for key '%s'.	All lines must end with a NEWLINE character.
<b>0x1100a8b9</b>	ADMIN	system.moab	WARN	MWM_CONVERT_XML_FROM_STRING	XML data cannot be obtained from an XML string ('%s').	There was an error converting from a string that should contain XML into internal XML data structures.
<b>0x1100a8ba</b>	ADMIN	system.moab	WARN	MWM_CONVERT_XML_TO_	An XML string cannot be constructed from XML data.	There was an error

Cod e	Es calati on Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
<b>a</b>			N	STRING		converting from internal XML data structures into an XML string representation.
<b>0x1 100 a8b b</b>	AD MI N	system.mo ab	W A R N	MWM_SOCKET_OPERATION	Cannot %s message on sd %s within %s second timeout.	There is a communication error with sockets.
<b>0x1 100 a8b c</b>	AD MI N	system.mo ab	W A R N	MWM_COMPLETE_JOB_RECORD	Could not create job record for completed job %s - %s.	The system might be low on memory.
<b>0x1 100 a8b d</b>	AD MI N	system.mo ab	W A R N	MWM_CREATE_JOB_DEPENDENCY	Could not create template job dependency %s - %s.	The system might be low on memory.
<b>0x1 100 a8b e</b>	AD MI N	system.mo ab	W A R N	MWM_CANNOT_FIND_SMP_NODE_BY_FEATURE	Could not find SMP node by feature '%s'.	The feature did not match any of the SMP nodes.
<b>0x1 100 a8bf</b>	AD MI N	system.mo ab	W A R N	MWM_CHECKPOINT_PROCESS_COMPLETED_JOB	Could not process completed job from checkpoint.	Examine the checkpoint entry for the job.
<b>0x1 100 a8c 0</b>	AD MI N	system.mo ab	W A R N	MWM_CANNOT_FIND_SMP_NODE_IN_	Could not find SMP node in partition '%s'.	The feature did not match any of the SMP nodes in the

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
				PARTITION		partition.
<b>0x1100a8c1</b>	ADMIN	system.moab	WARN	MWM_JOB_REPORTED_BY_RM_NOT_OWNER	Job '%s' is being reported by RM '%s' but is owned by RM '%s'.	The resource manager reporting does not own the job.
<b>0x1100a8c2</b>	ADMIN	system.moab	WARN	MWM_PEER_RM_UNKNOWN_LANGUAGE	Peer RM '%s' reported unknown language: '%s'.	The language does not match a known format.
<b>0x1100a8c3</b>	ADMIN	system.moab	WARN	MWM_PEER_RM_UNKNOWN_SUBLANGUAGE	Peer RM '%s' reported unknown sub-language: '%s'.	The language does not match a known format.
<b>0x1100a8c5</b>	ADMIN	system.moab	WARN	MWM_JOB_ARCH_VALUE	Job '%s' does not have a valid arch (architecture) value '%s'.	Check the specified value for the architecture.
<b>0x1100a8c6</b>	ADMIN	system.moab	WARN	MWM_INVALID_HOST_REQ	Job '%s' does not have a valid host requirement '%s'.	Check the specified value for the requirement.
<b>0x1100a8c7</b>	ADMIN	system.moab	WARN	MWM_JOB_OPYSYS_VALUE	Job '%s' does not have a valid operating system value '%s'.	Check the specified value for the operating system.
<b>0x1100a8c8</b>	ADMIN	system.moab	WARN	MWM_UNSUPPORTED_REQ	Resource requirement '%s' not supported.	The requirement specified is unsupported.

Cod e	Es calati on Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
<b>0x1100a8c9</b>	ADMI N	system.moab	WAR N	MWM_NO_TASKS	Job loaded in active state with no tasks allocated.	Jobs must have at least one task.
<b>0x1100a8ca</b>	ADMI N	system.moab	WAR N	MWM_SINGLE_ITERATION_JOB_COMPLETION	Scheduler cannot handle job completion in a single iteration.	The job must not start and complete while the scheduler is sleeping.
<b>0x1100a8cb</b>	ADMI N	system.moab	WAR N	MWM_MAX_TASKS_EXCEEDED	The number of tasks associated with a job has exceeded the maximum (%s).	The number of tasks must be reduced or the scheduler must be configured to accept more tasks.
<b>0x1100a8cc</b>	ADMI N	system.moab	WAR N	MWM_NODE_OUT_OF_RANGE	Job '%s' node index (%s) at task list index (%s) is out of range.	This is an internal limit.
<b>0x1100a8cd</b>	ADMI N	system.moab	WAR N	MWM_NODE_NULL	Job '%s' node index (%s) at task list index (%s) is NULL.	This is an internal limit.
<b>0x1100a8ce</b>	ADMI N	system.moab	WAR N	MWM_NODESET_CONSTRAINTS	Nodeset constraints prevent use of task for job '%s':%s at %s.	The specified nodeset cannot run the task.
<b>0x1100a8cf</b>	ADMI N	system.moab	WAR N	MWM_DEFAULT_WALLTIME	Job assigned default walltime limit (%s).	Unlimited or no walltime limit specified.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x1100a8d0</b>	ADMIN	system.moab	WARN	MWM_PARTITION_ACCESS	Job cannot access requested partitions (%s).	The partition access list disallows the job.
<b>0x1100a8d1</b>	ADMIN	system.moab	WARN	MWM_UNABLE_TO_ALLOCATE_TASKS_FOR_JOB	Cannot allocate tasks for job at %s.	The system might be low on memory.
<b>0x1100a8d2</b>	ADMIN	system.moab	WARN	MWM_IGNORING_PARTIAL_RANGE	Ignoring partial time range since full range previously located.	The system will use the full range instead.
<b>0x1100a8d3</b>	ADMIN	system.moab	WARN	MWM_DESTINATION_RM	Cannot locate a valid destination resource manager for job.	The submitted job could not be sent to a resource manager.
<b>0x1100a8d4</b>	ADMIN	system.moab	WARN	MWM_JOB_CREDENTIALS	Cannot authenticate the submitted job (Reason: %s).	The user for the job is not a member of a group or account with access.
<b>0x1100a8d5</b>	ADMIN	system.moab	WARN	MWM_SMPNODE_BY_FEATURE	Could not find SMPNode by feature %s (%s).	None of the nodes has the feature specified.
<b>0x1100a8d6</b>	ADMIN	system.moab	WARN	MWM_SMPNODE_BY_PARTITION	Could not find SMPNode in partition %s.	The SMPNode specified is not in the given partition.

Cod e	Es calati on Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
<b>0x1 100 a8d 7</b>	AD MI N	syste m.mo ab	W A R N	MWM_RSV_ ATTR_TO_ STRING	Reservation '%s' attribute '%s' could not be converted to a string.	There is no string conversion routine for that attribute type.
<b>0x1 100 a8d 8</b>	AD MI N	syste m.mo ab	W A R N	MWM_AM_ STATUS	Account manager sent failure message - %s.	Check status message.
<b>0x1 100 a8d 9</b>	AD MI N	syste m.mo ab	W A R N	MWM_AM_ FAILURE	Native accounting manager call '%s' failed using input XML '%s'.	Check XML syntax.
<b>0x1 100 a8d a</b>	AD MI N	syste m.mo ab	W A R N	MWM_AM_ INSUFFICIE NT_FUNDS	Account manager - Insufficient funds '%s'.	Validate that the user has access to account.
<b>0x1 100 a8d b</b>	AD MI N	syste m.mo ab	W A R N	MWM_ MIGRATE_ JOB	Unable to migrate job '%s' to RM '%s' (%s).	Check the error message.
<b>0x1 100 a8d c</b>	AD MI N	syste m.mo ab	W A R N	MWM_ RESERVE_ PRIORITY_ JOB	Unable to reserve priority job.	Check the error message.
<b>0x1 100 a8d d</b>	AD MI N	syste m.mo ab	W A R N	MWM_ SYNC_JOB	Job '%s' not synchronized to start with job '%s'.	The two jobs must start at the same time.
<b>0x1 100</b>	AD MI	syste m.mo	W A	MWM_ SYNC_JOB_	Job '%s' could not start. Requeuing any	The other jobs should be back

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>a8de</b>	N	ab	RN	REQUEUE	synchronized jobs.	on the queue.
<b>0x1100a8df</b>	ADMIN	system.moa	WARN	MWM_UNHANDLED_PLUGIN_EXCEPTION	A node allocation plugin '%s' encountered an unhandled exception '%s'.	Consult the documentation for the plugin.
<b>0x1100a8e0</b>	ADMIN	system.moa	WARN	MWM_PLUGIN_LOADED_FAILURE	Error loading node allocation plugin '%s' for partition '%s' %s.	A NodeAllocation plugin was not loaded because of an error. Default node allocation will be used.
<b>0x1100a8e2</b>	ADMIN	system.moa	WARN	MWM_UNKNOWN_JOB_DEPENDENCY	Unknown job dependency '%s' on job.	The job is trying to use a dependency that is unknown.
<b>0x1100a8e3</b>	ADMIN	system.moa	WARN	MWM_UNSUPPORTED_JOB_DEPENDENCY	Unknown job dependency type '%s' on job.	The job is trying to use a dependency type that is unsupported.
<b>0x1100a8e4</b>	ADMIN	system.moa	WARN	MWM_MISSING_JOB_DEPENDENCY	Cannot find dependency job. MissingDependencyAction is '%s'.	Check for the existence of the job dependency.
<b>0x1100a8e5</b>	ADMIN	system.moa	WARN	MWM_PARTITION_REP_NODE	Corrupt partition representative node.	Check the representative node for the partition.

Cod e	Es calati on Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
<b>0x1100a8e6</b>	ADMI N	system.moab	WARN	MWM_PARTITION_ATTRIBUTE	Partition attribute '%s' is not configurable.	Consult the documentation to see which attributes can be configured.
<b>0x1100a8e7</b>	ADMI N	system.moab	WARN	MWM_UNABLE_TO_LOCATE_NODE	Unable to locate specified nodes for job.	Could not find a node in the job's node list.
<b>0x1100a8e8</b>	ADMI N	system.moab	WARN	MWM_NODES_MISSING_FROM_FEASIBLE_LIST	Specified node(s) not found in feasible hostlist.	Could not find a node.
<b>0x1100a8e9</b>	ADMI N	system.moab	WARN	MWM_HOSTLIST_HAS_TOO_FEW_TASKS	A hostlist has too few tasks available for job '%s': '%s' (%s < %s).	More nodes are needed to satisfy the task requirements.
<b>0x1100a8ea</b>	ADMI N	system.moab	WARN	MWM_TASKS_REMAINING	A hostlist was unable to handle all tasks (%s remain).	More nodes are needed to satisfy the task requirements.
<b>0x1100a8eb</b>	ADMI N	system.moab	WARN	MWM_NODE_DOWN	Unable to detect node '%s' for '%s' seconds. Marking it down or removing it.	Make sure the node is up and running.
<b>0x1100a8ec</b>	ADMI N	system.moab	WARN	MWM_NODE_RESET_EMPTY	Unable to reset node. Node list empty.	Must specify a valid node to reset.



Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x1100a8ed</b>	ADMIN	system.moab	WARN	MWM_NODE_RESET_URL	Unable to reset node. NODEPOWERURL not specified.	Must specify a valid URL for the node to reset.
<b>0x1100a8ee</b>	ADMIN	system.moab	WARN	MWM_AM_DOWN	The account manager is not currently running.	Check the status of the account manager.
<b>0x1100a8ef</b>	ADMIN	system.moab	WARN	MWM_KEYBOARD_ACTIVITY_PREEMPT_JOB	Keyboard activity on node prevented job preemption.	Jobs can be preempted only if the keyboard is idle.
<b>0x1100a8f0</b>	ADMIN	system.moab	WARN	MWM_KEYBOARD_ACTIVITY_SET_NODE_STATE	Keyboard activity on node prevented setting the node state to '%s'.	Node states can be changed only if the keyboard is idle.
<b>0x1100a8f1</b>	ADMIN	system.moab	WARN	MWM_JOB_MISMATCHED_TIMES	Fixing job '%s' with invalid '%s' times (%s - %s).	Check the times for the specified job.
<b>0x1100a8f2</b>	ADMIN	system.moab	WARN	MWM_JOB_OPSYS	Cannot add operating system '%s' to job.	Check the type of operating system specified.
<b>0x1100a8f3</b>	ADMIN	system.moab	WARN	MWM_JOB_ARCH	Cannot add architecture '%s' to job.	Check the type of architecture specified.
<b>0x1100</b>	ADMIN	system.moab	WARN	MWM_LOCATE_AM	Cannot locate the account manager '%s'.	Check the account

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>a8f4</b>	N	ab	R N			manager command option syntax.
<b>0x1 100 a8f5</b>	ADMIN	system.moab	WARN	MWM_LOCATE_RM	Cannot locate the resource manager '%s'.	Check the resource manager command option syntax.
<b>0x1 100 a8f6</b>	ADMIN	system.moab	WARN	MWM_LOCATE_RMID	Cannot locate the resource manager ID '%s'.	Check the ID command option syntax.
<b>0x1 100 a8f7</b>	ADMIN	system.moab	WARN	MWM_LOCATE_PARTITION	Cannot locate the partition '%s'.	Check the partition command option syntax.
<b>0x1 100 a8f8</b>	ADMIN	system.moab	WARN	MWM_RM_QUEUE_MODIFY	Command to modify RM queue failed on resource manager %s - '%s'.	Queue may be configured to reject modify requests.
<b>0x1 100 a8f9</b>	ADMIN	system.moab	WARN	MWM_RM_QUEUE_CREATE	Command to create RM queue failed on resource manager %s - '%s'.	System may be configured to reject queue creation requests.
<b>0x1 100 a8fa</b>	ADMIN	system.moab	WARN	MWM_RM_QUEUE_CREATE_MISSING_ARGS	Command to create RM queue failed - arguments missing.	The user must supply the needed arguments to the command.
<b>0x1 100</b>	ADMIN	system.moab	WARN	MWM_STATIC_RM_	An attempt was made to destroy a static resource	Static resource managers

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>a8fb</b>	N	ab	R N	DESTRUCTION	manager.	cannot be destroyed.
<b>0x1 100 a8fc</b>	ADMIN	system.moab	WARN	MWM_ADD_SYSTEM_USER	Unable to create a new user '%s' in the system.	The system might be low on memory.
<b>0x1 100 a8fd</b>	ADMIN	system.moab	WARN	MWM_ADD_PARTITION	The system was unable to create partition '%s'.	The system might be low on memory.
<b>0x1 100 a8fe</b>	ADMIN	system.moab	WARN	MWM_CORE_LIMIT	System core limit set to %s (complete core files might not be generated).	Expand the system core limit to ensure the complete core dump can be saved.
<b>0x1 100 a8ff</b>	ADMIN	system.moab	WARN	MWM_KEY_FILE_PERMISSIONS	The .moab.key file exists, but the file permissions prevent access (%s).	Check the ownership permissions on the file.
<b>0x1 100 a900</b>	ADMIN	system.moab	WARN	MWM_STATS_PERIOD_TYPE	The system could not process stats for period type %s.	'Day' is the only period type currently supported.
<b>0x1 100 a901</b>	ADMIN	system.moab	WARN	MWM_STATS_BUFFER_SIZE	The system could not process stats for period type %s (buffer too small).	The buffer allocated was too small to hold the data.
<b>0x1 100 a902</b>	ADMIN	system.moab	WARN	MWM_STATS_FILE	The system could not create the stats file '%s'.	Check the path and user permissions on the directory.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
0x1100a903	ADMIN	system.moab	WARN	MWM_JOB_NO_ACCOUNT	No account specified for job '%s'.	Check the job for an account specification.
0x1100a904	ADMIN	system.moab	WARN	MWM_SET_JOBATTR_FAIL	Cannot set attribute '%s' to value '%s' on jobmatch '%s'.	Failed to set the given attribute to the given value on the given job.
0x1100a905	ADMIN	system.moab	WARN	MWM_JOBATTR_NOT_SUPPORTED	JobAttr not supported. '%s'.	The given attribute is not a supported job attribute.
0x1100a906	ADMIN	system.moab	WARN	MWM_INVALID_TRIGGER_DEFINITION	Invalid trigger definition: %s.	The given trigger is invalid. Check that the given trigger has been defined.
0x1100a907	ADMIN	system.moab	WARN	MWM_ATTRIBUTE_NOT_HANDLED	System attribute '%s' not handled.	Check that the given attribute was spelled correctly.
0x1100a908	ADMIN	system.moab	WARN	MWM_QOS_IN_PARAM_NOT_FOUND	Cannot locate QOS '%s' for parameter %s.	Make sure that the given QOS exists.
0x1100a909	ADMIN	system.moab	WARN	MWM_INVALID_PROFILEDURATION_VAL	Invalid PROFILEDURATION specified, modified internally to %s (see documentation).	The entered PROFILEDURATION value is invalid. Moab uses the given value instead.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x1100a90a</b>	ADMIN	system.moab	WARN	MWM_NO_DATA_STAGING_PATH	No path in data staging specification '%s' (bad format).	Verify the data staging path is specified.
<b>0x1100a90b</b>	ADMIN	system.moab	WARN	MWM_UNSUPPORTED_RM_DATA_STAGING	Cannot stage-out stdout/stderr (unsupported RM type '%s').	Failed to stage-out stdout/stderr because the given resource manager does not support such feature.
<b>0x1100a90c</b>	ADMIN	system.moab	WARN	MWM_RM_DATA_ON_NON_EXISTING_JOB	Storage RM '%s' reporting data operation for non-existent job '%s'.	The given resource manager is reporting data operation on the non-existing job.
<b>0x1100a90d</b>	ADMIN	system.moab	WARN	MWM_INVALID_RM_DATA_STAGE	Data stage for RM '%s' not possible as it has no nodelist.	Check CLUSTERQUE RYURL to ensure it at least has a nodelist.
<b>0x1100a90e</b>	ADMIN	system.moab	WARN	MWM_DATA_STAGE_IN_FAIL	Data stage in failed for job '%s' file '%s' (%s).	Failed to complete the data staging in operation for the job on the given file due to error(s).
<b>0x1100a90f</b>	ADMIN	system.moab	WARN	MWM_UNABLE_TO_	Cannot remove data staging block for job '%s'.	Failed to remove the data staging

Cod e	Es calati on Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
			N	REMOVE_ DATA_ STAGE		block for the given job.
<b>0x1 100 a91 0</b>	AD MI N	system.mo ab	W A R N	MWM_ DATA_ STAGE_ OUT_FAIL	Data stage out failed for job '%s' file '%s' (%s).	Failed to complete the data staging out operation for the job on the given file due to error (s).
<b>0x1 100 a91 1</b>	AD MI N	system.mo ab	W A R N	MWM_JOB_ INVALID_ OPSYS	Job '%s' cannot request OS '%s').	The requested operating system is not available for the job.
<b>0x1 100 a91 2</b>	AD MI N	system.mo ab	W A R N	MWM_ NODE_ BUFFER_ OVERFLOW	Node buffer is full (check license and MAXNODE parameter).	Try increasing the node buffer size (MAXNODE).
<b>0x1 100 a91 3</b>	AD MI N	system.mo ab	W A R N	MWM_ REMOVE_ NODE_ WITH_ RESERVATI ON	Unable to remove node '%s' because of reservation references.	Remove the reservations from the node.
<b>0x1 100 a91 4</b>	AD MI N	system.mo ab	W A R N	MWM_JOB_ PURGE_RM_ INACTIVE	Unable to purge job '%s' because the resource manager '%s' is inactive.	Check the status of the resource manager.
<b>0x1 100 a91 5</b>	AD MI N	system.mo ab	W A R N	MWM_JOB_ PURGE_RM_ NO_ RESOURCES	Unable to purge job '%s' because the resource manager '%s' is reporting no resources.	Check the status of the resource manager.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x1100a916</b>	ADMIN	system.moab	WARN	MWM_JOB_NOT_DETECTED	Job '%s' in state '%s' no longer detected (Last Detected %s > PurgeTime %s).	The job might have been purged in the meantime.
<b>0x1100a917</b>	ADMIN	system.moab	WARN	MWM_BACKFILL_DEPTH_REACHED	The backfill depth (BFDEPTH) has been reached so no more jobs will be backfilled this iteration.	Wait for the next iteration or increase the depth.
<b>0x1100a91c</b>	ADMIN	system.moab	WARN	MWM_CHECKPOINT_CREATE_RSV_FROM_XML	Unable to create a reservation from checkpoint XML.	The system might be low on memory.
<b>0x1100a91d</b>	ADMIN	system.moab	WARN	MWM_JOB_CACHE_REMOVAL	Failed to remove job %s (ID = %s) from the cache.	The job was missing from the system hash table.
<b>0x1100a91e</b>	ADMIN	system.moab	WARN	MWM_INVALID_CRED_VALUE	Invalid credential value '%s'.	Check the syntax in the configuration file.
<b>0x1100a91f</b>	ADMIN	system.moab	WARN	MWM_INVALID_CRED_ATTR	Invalid credential attribute '%s'.	Check the syntax in the configuration file.
<b>0x1100a920</b>	ADMIN	system.moab	WARN	MWM_MAX_JOBS_EXCEEDED	The maximum number of jobs has been exceeded.	Increase the value of the MAXJOB setting.
<b>0x1100</b>	ADMIN	system.moab	WARN	MWM_JOB_FAILED_	A job failed while processing PBS resources.	May not have been able to

Cod e	Es calati on Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
<b>a921</b>	N	ab	R N	PROCESSING_PBS		locate host or vnode.
<b>0x1100a923</b>	ADMI N	system.mo ab	WARN	MWM_AM_INSUFFICIENT_BALANCE	Insufficient balance in primary account '%s' to run job '%s' (attempting fallback credentials).	Validate that the user has access to account.
<b>0x1100a924</b>	ADMI N	system.mo ab	WARN	MWM_AM_JOB_SUBMIT_VALIDATION	Job submission validation failed for job '%s' -- taking action '%s'.	Validate that the job has access.
<b>0x1100a925</b>	ADMI N	system.mo ab	WARN	MWM_TOO_MANY_NODE_SETS	The maximum number of node sets has been exceeded.	This is a configurable setting.
<b>0x1100a926</b>	ADMI N	system.mo ab	WARN	MWM_CLASS_SET_LIST_INVALID	The specified class set list is invalid '%s'.	Check the documentation for valid classes.
<b>0x1100a928</b>	ADMI N	system.mo ab	WARN	MWM_RM_START	Cannot start resource manager '%s' (Reason: %s).	The resource manager might not be available.
<b>0x1100a929</b>	ADMI N	system.mo ab	WARN	MWM_STDOUT_FAIL	Request succeeded with no stdout. stderr= '%s'.	The standard out might not have been specified.
<b>0x1100a92a</b>	ADMI N	system.mo ab	WARN	MWM_UNSUPPORTED_REQ_ATTR	Unsupported req attribute '%s'.	The attribute is not one that can be set.



Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x1100a92b</b>	ADMIN	system.moab	WARN	MWM_UNUPPORTED_GENERAL_ATTR	Unsupported general attribute '%s'.	The attribute is not one that can be set.
<b>0x1100a92e</b>	ADMIN	system.moab	WARN	MWM_CREATE_ACCOUNT	Unable to create account '%s' on the account manager.	Verify that the account manager is running.
<b>0x1100a92f</b>	ADMIN	system.moab	WARN	MWM_QUERY_ACCOUNT	Unable to query account '%s' on the account manager.	Verify that the account name is correct.
<b>0x1100a930</b>	ADMIN	system.moab	WARN	MWM_ACCOUNT_ADD_USER	Unable to add user '%s' to account '%s' on the account manager.	Verify that the account name is correct.
<b>0x1100a931</b>	ADMIN	system.moab	WARN	MWM_ACCOUNT_DEPOSIT	Unable to deposit '%s' credits to account '%s' on the account manager.	Verify that the account name is correct.
<b>0x1100a932</b>	ADMIN	system.moab	WARN	MWM_ALLOCATE_REQ	Unable to allocate requirement '%s' using NAllocPolicy '%s' (%s).	The system might be low on memory.
<b>0x1100a933</b>	ADMIN	system.moab	WARN	MWM_VALID_STAT_DATA	Unable to generate valid statistic data for external query.	The system might be low on memory.
<b>0x1100a93</b>	ADMIN	system.moab	WARN	MWM_JOB_QOS_REQUEST	Job '%s' cannot request QOS '%s'.	The requested QOS is not available for

Cod e	Es calati on Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
<b>4</b>			N			the job.
<b>0x1100a938</b>	ADMI N	system.moab	WARN	MWM_GRES_OVERFLOW	GRES overflow.	Unable to add another GRES.
<b>0x1100a939</b>	ADMI N	system.moab	WARN	MWM_RM_NO_RESOURCES	The resource manager '%s' is reporting no resources.	Check the nodes on the resource manager.
<b>0x1100a93d</b>	ADMI N	system.moab	WARN	MWM_EMPTY_FILE	File '%s' is empty.	Check the file specified.
<b>0x1100a93f</b>	ADMI N	system.moab	WARN	MWM_PREEMPT_NONACTIVE_JOB	Cannot preempt non-active job '%s' (state: '%s' estate: '%s').	The job must currently be active to preempt it.
<b>0x1100a940</b>	ADMI N	system.moab	WARN	MWM_REQUEUE_NONSTARTABLE_JOB	Cannot requeue non-startable job '%s' (canceling instead).	The job could not be requeued.
<b>0x1100a941</b>	ADMI N	system.moab	WARN	MWM_NODE_RACK_VALUE	Invalid rack value '%s' specified for node %s (must be digit).	Check the value of the rack parameter.
<b>0x1100a942</b>	ADMI N	system.moab	WARN	MWM_ENCODE_JOB_MESSAGE	Cannot encode job message.	Check the value of the rack parameter.
<b>0x1100</b>	ADMI N	system.moab	WARN	MWM_INVALID_	Invalid PROFILECOUNT specified, modified	The

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>a943</b>	N	ab	RN	PROFILECOUNT_VAL	internally to %s (see documentation).	PROFILECOUNT value is invalid. Moab uses the default value instead.
<b>0x1100a944</b>	ADMIN	system.moab	WARN	MWM_FAIRSHARE_FILE	Cannot load fairshare file '%s' for slot %s.	Check for the existence of the fairshare file in the file system.
<b>0x1100a945</b>	ADMIN	system.moab	WARN	MWM_REQ_ATTR_ALREADY_SET	Requirement attribute %s '%s' is already set.	Check the attribute setting in the configuration file.
<b>0x1100a946</b>	ADMIN	system.moab	WARN	MWM_ZERO_START_TIME	StartTime set to zero for reservation on job '%s'.	Check the start time for the specified job.
<b>0x1100a947</b>	ADMIN	system.moab	WARN	MWM_EXISTING_RESERVATION	Reservation created for reserved job '%s' (existing reservation '%s' deleted).	Only one reservation can exist at a time for the reserved job.
<b>0x1100a948</b>	ADMIN	system.moab	WARN	MWM_CANNOT_PARSE_REQ_LINE	Cannot parse requirement line for job '%s'.	The syntax of the requirement line is incorrect.
<b>0x1100a949</b>	ADMIN	system.moab	WARN	MWM_UNKNOWN_RESOURCE_TYPE	Unknown resource type '%s' for job '%s'.	Check the documentation for valid resource types.

Cod e	Es cal ation Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
<b>0x1100a94a</b>	ADMI N	system.moab	WARN	MWM_UNKNOWN_TRANSACTION_ATTR	Unknown transaction attribute '%s'.	Check the documentation for valid transaction attributes.
<b>0x1100a94c</b>	ADMI N	system.moab	WARN	MWM_ID_MANAGER_DOWN	The identity manager is down.	Check the status of the identity manager.
<b>0x1100a94d</b>	ADMI N	system.moab	WARN	MWM_PROCESS_ID_LINE	Unable to process the ID line '%s'.	Check the syntax of the attribute/value pairs.
<b>0x1100a94e</b>	ADMI N	system.moab	WARN	MWM_JOB_ID_MISSING	Unable to locate the job ID for a job submitted to the resource manager.	Check the job being submitted`.
<b>0x1100a950</b>	ADMI N	system.moab	WARN	MWM_EMPTY_NODELIST	The nodelist is empty for reservation '%s'.	Reservations should include a node list.
<b>0x1100a951</b>	ADMI N	system.moab	WARN	MWM_FAIRSHARE_PAL	Fairshare does not allow specified PAL (%s).	The fairshare algorithm is reverting to the original PAL.
<b>0x1100a952</b>	ADMI N	system.moab	WARN	MWM_JOB_START_TIME	Cannot find earliest start time for job '%s'.	Resources needed to run the job may never be available.
<b>0x1</b>	AD	system	W	MWM_	Cannot modify node '%s'	The node

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>100a953</b>	MIN	m.moab	ARN	NODE_MODIFY	Error(%s).	could not be modified.
<b>0x1100a954</b>	ADMIN	system.moab	WARN	MWM_TRIGGER_RSV_CREATE	Unable to create a trigger reservation.	Check the reservation time, nodes, and account.
<b>0x1100a955</b>	ADMIN	system.moab	WARN	MWM_MISSING_TRIGGER	Trigger '%s' with PID '%s' does not exist--completing!	The process might have already completed.
<b>0x1100a956</b>	ADMIN	system.moab	WARN	MWM_EMPTY_HOSTLIST	The hostlist is empty for reservation.	Reservations should include a hostlist.
<b>0x1100a957</b>	ADMIN	system.moab	WARN	MWM_RSV_POLICY_VIOLATION	Unable to create requested reservation due to a policy violation (%s).	Reservations must conform to existing policies.
<b>0x1100a958</b>	ADMIN	system.moab	WARN	MWM_RSV_CREATE_FAILURE	Unable to create requested reservation at time %s (%s).	Resources are unavailable at requested time.
<b>0x1100a959</b>	ADMIN	system.moab	WARN	MWM_RSV_OWNER	Cannot process owner '%s' for standing reservation '%s' (%s).	Consult the error message.
<b>0x1100a95a</b>	ADMIN	system.moab	WARN	MWM_RSV_PARTIAL	Partial reservation %s reserved %s of %s procs in partition '%s' to start in %s at (%s) %s.	Entire reservation could not be filled.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x1100a95c</b>	ADMIN	system.moab	WARN	MWM_RSV_NEGATIVE_JOBCOUNT	Reservation %s jobcount is %s, should not decrement less than 0.	JobCount cannot be negative.
<b>0x1100a95d</b>	ADMIN	system.moab	WARN	MWM_EMPTY_REQ_NODELIST	Req node list empty for job %s:%s in state %s (job nodelist copied to req nodelist).	Job should include a req node list.
<b>0x1100a95e</b>	ADMIN	system.moab	WARN	MWM_TASK_ALLOCATION_INFO	Cannot locate task allocation info for job %s:%s in state %s.	Job should include a task list.
<b>0x1100a95f</b>	ADMIN	system.moab	WARN	MWM_INVALID_THREADPOOL_SIZE	Invalid ThreadPoolSize '%s' (must be a non-negative integer no larger than %s).	Check the size for a valid value.
<b>0x1100a960</b>	ADMIN	system.moab	WARN	MWM_INVALID_QUEUE_TIME	Job '%s' has invalid system queue time (SQ: %s > ST: %s).	Check the job queue time value.
<b>0x1100a961</b>	ADMIN	system.moab	WARN	MWM_NO_WCLIMIT	Job '%s' has no WCLimit specified.	Check the job for the correct value.
<b>0x1100a962</b>	ADMIN	system.moab	WARN	MWM_AM_INVALID_PROTOCOL	Invalid protocol '%s' specified for account manager '%s'.	Communication with the account manager must be over a supported protocol.
<b>0x1</b>	AD	system	W	MWM_NO_	No external power interface	Cannot set the

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>100a963</b>	MIN	m.moab	ARN	POWER_INTERFACE	- cannot set power state '%s' on node '%s%s%s'.	power state on the node without a power interface.
<b>0x1100a964</b>	ADMIN	system.moab	WARN	MWM_JOB_STARTED_ON_ANOTHER_RM	Job '%s' started externally: (rc: %s; errmsg: '%s'; Tasklist: '%s').	Two or more resource managers are running side-by-side and the job is already running on one of them.
<b>0x1100a965</b>	ADMIN	system.moab	WARN	MWM_COMMAND_FAILED_CHILD_PROCESS	Job submit request failed with child process status code='%s', stderr='%s', stdout='%s', EMsg='%s'.	Review the status code and error message for further information.
<b>0x1100a967</b>	ADMIN	system.moab	WARN	MWM_AM_REGISTER_JOB	Unable to register job creation with account manager for job '%s', reason: '%s'.	Check the status of the account manager.
<b>0x1100a968</b>	ADMIN	system.moab	WARN	MWM_AM_DEPRECATED_PARAMETER	Use of the '%s' parameter is deprecated. %s	Check the documentation for the new parameter syntax.
<b>0x1100a969</b>	ADMIN	system.moab	WARN	MWM_AM_INVALID_ACTION	Invalid action '%s' specified in '%s' for account manager '%s'.	Check the documentation for valid actions for the account manager.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x1100a9bf</b>	ADMIN	system.moab	WARN	MWM_ERASING_JOB	Erasing job '%s' by address.	The specified job could not be found by name. The entire job table was searched to find the matching job.
<b>0x1100a9e1</b>	ADMIN	system.moab	WARN	MWM_JOB_INVALID_TASK_LAYOUT	Job '%s' has invalid task layout (TPN:%s * N:%s != T:%s).	The task layout does not compute.
<b>0x1100a9e2</b>	ADMIN	system.moab	WARN	MWM_JOB_ACCESS_QOS	Job '%s' does not have access to QOS '%s' (QAL:%s).	The QoS is not accessible from the job.
<b>0x1100a9f8</b>	ADMIN	system.moab	WARN	MWM_DUPLICATE_SYSTEMJID	Duplicate SystemJID '%s' [JState: %s] found from RM '%s'.	The SystemJID must be unique.
<b>0x1100a9fc</b>	ADMIN	system.moab	WARN	MWM_CANNOT_PING_RM	Cannot ping RM '%s' because a file was not specified.	A file path to a valid file is needed.
<b>0x1100aa07</b>	ADMIN	system.moab	WARN	MWM_VM_CONTAINER_NODE	Cannot find or add container node '%s' for VM '%s'.	The node could not be found.
<b>0x1100aa4b</b>	ADMIN	system.moab	WARN	MWM_DATASTAGING_DYNAMIC	Failed to calculate dynamic walltime for data staging system job '%s'.	Unlimited walltime will be used.



Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
				WALLTIME_CALCULATION_FAILURE		
<b>0x1100aa53</b>	ADMIN	system.moab	WARN	MWM_PBS_JOB_FULL_REPORT_TIME_MISMATCH	PBS server (%s) is configured with a job_full_report_time (%s) less than Moab's RMPollInterval (%s). This could lead to incomplete job information.	Moab/TORQUE configuration mismatch.
<b>0x1100aa54</b>	ADMIN	system.moab	WARN	MWM_BACKLOGCOMPLETIONTIME_NEEDS_ENABLEPROFILING	BacklogCompletionTime cannot be calculated; EnableProfiling must be enabled on the QOS to gather stats.	BacklogCompletionTime needs EnableProfiling enabled.
<b>0x1100aa62</b>	ADMIN	system.moab	WARN	MWM_NATIVE_RM_OVERRIDE	RM '%s', of type '%s' is overriding default '%s' operation with configured 'native' call.	A non-native RM that is configured with native calls can override the default functions.
<b>0x1100e710</b>	INTERNAL	system.moab	WARN	MWM_INVALID_ARG_VALUE	Invalid arguments passed to this function.	One or more arguments passed to this function were not valid. This is an internal error logged for informational purposes.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x1100e72b</b>	INTERNAL	system.moab	WARN	MWM_SEND_SENT_NO_DATA	No data was sent to the socket when it should have been.	The send() system call reported no data was sent when data should have been sent.
<b>0x1100e72e</b>	INTERNAL	system.moab	WARN	MWM_SOCKET_BLOCKED_UNEXPECTEDLY	Read operations on the socket were blocked when it should have been available.	A socket operation reported that the operation was blocked. Previous information indicated that this operation should have been available.
<b>0x1100e73f</b>	INTERNAL	system.moab	WARN	MWM_MUTEX_LOCK	Cannot lock mutex semaphore using pthread_mutex_lock().	This is an operating system call problem.
<b>0x1100e740</b>	INTERNAL	system.moab	WARN	MWM_MUTEX_UNLOCK	Cannot unlock mutex semaphore using pthread_mutex_unlock().	This is an operating system call problem.
<b>0x1100e771</b>	INTERNAL	system.moab	WARN	MWM_THREAD_DB_INIT	Thread %s attempting to re-initialize database info struct.	Internal error condition.
<b>0x1100e84b</b>	INTERNAL	system.moab	WARN	MWM_CORRUPT_COMMAND_RECEIVED	Corrupt command '%s' received.	The communication packet received from a Moab client

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
						command is malformed.
<b>0x1100e919</b>	INTERNAL	system.moab	WARN	MWM_CHECKPOINT_NO_XML	The checkpoint data does not contain XML.	Internal error.
<b>0x1100e91a</b>	INTERNAL	system.moab	WARN	MWM_CHECKPOINT_INVALID_XML	The checkpoint data does not contain valid XML (%s).	Internal error.
<b>0x1100e91b</b>	INTERNAL	system.moab	WARN	MWM_CHECKPOINT_UPDATE_RSV_FROM_XML	Unable to update a reservation from checkpoint XML.	Internal error.
<b>0x1100e927</b>	INTERNAL	system.moab	WARN	MWM_JOB_ATTR_TO_STRING	Job attribute '%s' not yet translated to string value.	Internal warning.
<b>0x1100e92c</b>	INTERNAL	system.moab	WARN	MWM_MISSING_STATUS_CODE	The status code was missing from the S3 response.	This is an internal error.
<b>0x1100e92d</b>	INTERNAL	system.moab	WARN	MWM_MISSING_STATUS_VALUE	The status value was missing from the S3 response.	This is an internal error.
<b>0x1100e93a</b>	INTERNAL	system.moab	WARN	MWM_SIMULATION_NO_JOBS	No jobs loaded in simulation.	Internal simulation error.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x1100e93b</b>	INTERNAL	system.moab	WARN	MWM_SIMULATION_JOB_DETECTED_TRACEBUFFER	Job '%s' previously detected in tracefile (MJobTraceBuffer[%s]/JC: %s; IT: %s).	Internal simulation error.
<b>0x1100e93c</b>	INTERNAL	system.moab	WARN	MWM_SIMULATION_JOB_DETECTED	Job '%s' previously detected in tracefile (Job/JC: %s; IT: %s).	Internal simulation error.
<b>0x1100e93e</b>	INTERNAL	system.moab	WARN	MWM_READ_COMMAND_OUTPUT	Cannot read output of command '%s'.	This is an internal communications error.
<b>0x1100e94b</b>	INTERNAL	system.moab	WARN	MWM_THREAD_TIMEOUT	Thread %s killed (%s micro-second time out reached).	This is an internal issue.
<b>0x1100e94f</b>	INTERNAL	system.moab	WARN	MWM_INVALID_XML_RM	Invalid XML data for resource manager '%s'.	Check the XML syntax.
<b>0x21000067</b>	USER	domain.life cycle	ERROR	MWM_JOB_END_FAILED	Job %s failed at %s. %s	The job finished unsuccessfully.
<b>0x210000ca</b>	USER	domain.life cycle	ERROR	MWM_NODE_EVAC_VMS_ERROR	Error evacuating VMs off node %s. %s	There was an error while attempting to evacuate the VMs off the node.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x21002882</b>	USER	system.moab	ERROR	MWM_NODE_MODIFY_FAILURE	Cannot modify node state of '%s' Error(%s).	The node state could not be modified.
<b>0x21002a1a</b>	USER	system.moab	ERROR	MWM_DEPRECATED_RM_FEATURE	RM flag SUBMITJOBSASROOT not supported with this version, %s. Must be >= 2.4.8.	The resource manager version should be updated to get support for this feature.
<b>0x21008003</b>	ADMIN	system.moab	ERROR	MWM_TESTING_ERROR	Testing with argument1: %s. and argument2: %s.	Internal error for testing diagnostics.
<b>0x21008263</b>	ADMIN	domain.lifecycle	ERROR	MWM_VC_SCHEDULE_FAILURE	Failed to schedule virtual container '%s'.	This is an internal error.
<b>0x21008388</b>	ADMIN	system.moab	ERROR	MWM_INVALID_LICENSE_USE_ERROR	Use of %s requires license with %s enabled	Requested feature must be licensed. Please contact your sales representative at Adaptive Computing for assistance.
<b>0x2100a718</b>	ADMIN	system.moab	ERROR	MWM_CANNOT_FORK_ERROR	Cannot fork the process, errno: %s (%s).	The fork() system call failed. Use the errno and associated

Cod e	Es calati on Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
						message to determine possible causes.
<b>0x2100a719</b>	ADMI N	system.moab	ERROR	MWM_CANNOT_EXEC_PROGRAM	Cannot exec action '%s', errno: %s (%s).	The exec() system call failed to execute the command. This might be because the command does not exist or the permissions do not allow it to be run. Use the errno and associated message to determine possible causes.
<b>0x2100a71b</b>	ADMI N	system.moab	ERROR	MWM_CANNOT_CHOWN_FILE	Failure changing ownership of file: '%s' to uid:'%s', gid:'%s', errno: %s (%s).	The chown() system call failed. Use the errno and associated message to determine possible causes.
<b>0x2100a71f</b>	ADMI N	system.moab	ERROR	MWM_PIPE_READ_FAILED	Failed to read pipe on command '%s', errno: %s (%s).	The fread() system call failed. Use the errno and associated message to

Cod e	Es calati on Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
						determine possible causes.
<b>0x2 100 a72 0</b>	AD MI N	system.mo ab	E R R O R	MWM_ CANNOT_ READ_FILE	Cannot read file '%s', errno: %s (%s).	The fread() system call failed. Use the errno and associated message to determine possible causes.
<b>0x2 100 a72 1</b>	AD MI N	system.mo ab	E R R O R	MWM_ CANNOT_ WRITE_TO_ FILE	Failure writing to file, errno: %s (%s).	The write() system call failed. Use the errno and associated message to determine possible causes.
<b>0x2 100 a72 5</b>	AD MI N	system.mo ab	E R R O R	MWM_ CANNOT_ GET_ HOSTNAME	Cannot get hostname '%s', errno: %s (%s).	The gethostname() system call failed. Use the errno and associated message to determine possible causes.
<b>0x2 100 a72 7</b>	AD MI N	system.mo ab	E R R O R	MWM_ CANNOT_ CREATE_ SOCKET	Failure creating a socket, errno: %s (%s).	The socket() system call failed. Use the errno and associated

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
						message to determine possible causes.
<b>0x2100a728</b>	ADMIN	system.moab	ERROR	MWM_CANNOT_CONNECT_TO_HOST	Failure connecting to server '%s' on port %s, errno: %s (%s).	The connect() system call failed. Use the errno and associated message to determine possible causes.
<b>0x2100a730</b>	ADMIN	system.moab	ERROR	MWM_EPOCH_FAIL	Epoch Fail, time: '%s' cannot be converted to an epoch time, errno: %s (%s).	The mktime() system call failed. Use the errno and associated message to determine possible causes.
<b>0x2100a731</b>	ADMIN	system.moab	ERROR	MWM_MEMORY_ALLOCATION_FAILURE_MALLOC	Failure allocating memory (malloc), allocating '%s' bytes, errno: %s (%s).	The malloc() system call failed. Use the errno and associated message to determine possible causes.
<b>0x2100a732</b>	ADMIN	system.moab	ERROR	MWM_MEMORY_ALLOCATION_FAILURE_CALLOC	Failure allocating memory (calloc), allocating '%s' elements of size '%s' bytes, errno: %s (%s) in file %s:%s.	The calloc() system call failed. Use the errno and associated



Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
						message to determine possible causes.
<b>0x2100a733</b>	ADMIN	system.moab	ERROR	MWM_MEMORY_ALLOCATION_FAILURE_REALLOC	Failure allocating memory (realloc), allocating '%s' bytes, errno: %s (%s).	The realloc() system call failed. Use the errno and associated message to determine possible causes.
<b>0x2100a734</b>	ADMIN	system.moab	ERROR	MWM_CANNOT_DUPLICATE_STRING	Failure duplicating string, errno: %s (%s).	The strdup() system call failed. Use the errno and associated message to determine possible causes.
<b>0x2100a735</b>	ADMIN	system.moab	ERROR	MWM_CANNOT_CHANGE_PROCESS_GROUP	Failure changing process group, errno: %s (%s).	The setpgrp() system call failed. Use the errno and associated message to determine possible causes.
<b>0x2100a736</b>	ADMIN	system.moab	ERROR	MWM_CANNOT_CREATE_THREAD	Failure creating thread: '%s', errno: %s (%s).	The pthread_create() system call failed. Use the errno and

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
						associated message to determine possible causes.
<b>0x2100a737</b>	ADMIN	system.moab	ERROR	MWM_CANNOT_TRUNCATE_FILE	Failure truncating a file '%s', errno: %s (%s).	The truncate() system call failed. Use the errno and associated message to determine possible causes.
<b>0x2100a739</b>	ADMIN	system.moab	ERROR	MWM_PIPE_OPEN_FAILED	Failed to open pipe on command '%s', errno: %s (%s)	The popen() system call failed. Use the errno and associated message to determine possible causes.
<b>0x2100a73d</b>	ADMIN	system.moab	ERROR	MWM_CHANGE_DIR_FAILURE	OS call to change directory to '%s' failed errno: %s (%s).	The chdir() system call failed. Use the errno and associated message to determine possible causes.
<b>0x2100a74a</b>	ADMIN	system.moab	ERROR	MWM_CANNOT_LOCK_MOAB_PID_	Cannot lock the PID file '%s'. Is Moab already running?	Moab tries to ensure that only one instance of

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
			R	FILE		itself is running. In the default configuration it will exit if it cannot obtain a lock.
<b>0x2100a74e</b>	ADMIN	system.moab	ERROR	MWM_CONFIG_FILE_NOT_FOUND_ERROR	Cannot locate configuration file in any predetermined location.	Moab cannot find the configuration file. Verify that it is present and installed in a proper location.
<b>0x2100a757</b>	ADMIN	system.moab	ERROR	MWM_MWS_RM_CONFIGURATION	The resource manager with Moab Web Services (%s) does not have a base URL, username, and password configured.	Correctly configure the Moab Web Services resource manager.
<b>0x2100a761</b>	ADMIN	system.moab	ERROR	MWM_STRICT_INVALID_CONFIG_LINE	Error processing line #%s: %s - (%s).	Check the line number in the configuration file.
<b>0x2100a768</b>	ADMIN	system.moab	ERROR	MWM_MONGOSERVER_INITIALIZATION_FAILED	Failed to initialize connection to Mongo server '%s' RelicatSetName: '%s' SSLMode: '%s' SSL CA File: '%s'.	Failed to initialize connection to the configured MONGOSERVER. Check the following: (1) network connection to Mongo server;

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
						and (2) check MONGouser and MONGOPASSWORD parameters in moab-private.cfg.
<b>0x2100a76a</b>	ADMIN	system.moab	ERROR	MWM_MONGOSERVER_AUTHENTICATION_FAILURE	Failed to authenticate to Mongo server (%s).	Check user credentials.
<b>0x2100a76b</b>	ADMIN	system.moab	ERROR	MWM_MONGOSERVER_WRITE_FAILURE	Unable to write out transition object '%s'.	The BSON information is invalid or missing.
<b>0x2100a76c</b>	ADMIN	system.moab	ERROR	MWM_MONGOSERVER_DOWN	The Mongo server is down.	Check the status of the server.
<b>0x2100a76f</b>	ADMIN	system.moab	ERROR	MWM_DB_CHECKPOINT_OBJECT	Unable to checkpoint object to the database (%s).	Make sure the database is running.
<b>0x2100a773</b>	ADMIN	system.moab	ERROR	MWM_BACKUP_SERVER_CONNECTION_FAILED	The system was unable to connect to the backup server %s (%s:%s).	Make sure the backup server's address is correct.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
0x2100a776	ADMIN	system.moab	ERROR	MWM_CONNECTIO N_REFUSED	Connection to the server was refused (%s).	Primary server refused and no fallback server available.
0x2100a777	ADMIN	system.moab	ERROR	MWM_UNABLE_ TO_ CONNECT	Cannot send request to %s:%s (%s may not be running).	Unable to connect to the scheduler program.
0x2100a779	ADMIN	system.moab	ERROR	MWM_CLIENT_ MAX_ CONNECTIO NS_ REACHED	Cannot accept connection number %s (transaction number %s) from '%s' (limit reached).	May need to increase the CLIENTMAXC ONNECTIONS configuration setting.
0x2100a77a	ADMIN	system.moab	ERROR	MWM_SERVER_ CONNECTIO N_FAILED	The system was unable to connect to the server %s:%s - %s.	Make sure the server's address is correct and it is running.
0x2100a77c	ADMIN	system.moab	ERROR	MWM_COMMUNIC ATION_ ERROR	Communication error %s:%s (%s).	General error trying to communicate with the host.
0x2100a77d	ADMIN	system.moab	ERROR	MWM_CANNOT_ PARSE_ SERVER_ RESPONSE_ STATUS	Cannot parse server response (status).	The response sent from the server is malformed.
0x2	ADMIN	system.moab	E	MWM_	Cannot parse server response (data).	The response

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
100a77e	N	ab	RROR	CANNOT_PARSE_SERVER_RESPONSE_DATA		sent from the server is malformed.
0x2100a77f	ADMIN	system.moab	ERROR	MWM_INVALID_FS_TARGET	Invalid type specified for FSTarget.	Fairshare target type is invalid.
0x2100a780	ADMIN	system.moab	ERROR	MWM_COULD_NOT_ADD_FS_TREE_NODE	Could not add fstree node %s.	Unable to add a node to the fairshare configuration tree.
0x2100a781	ADMIN	system.moab	ERROR	MWM_CANNOT_ADD_MANAGER_TO_FS_TREE	Could not add manager %s to fstree.	Unable to add a manager to the fairshare tree.
0x2100a782	ADMIN	system.moab	ERROR	MWM_CRED_MANAGER_OVERFLOW	CredManager overflow, manager %s not added.	Credential Manager could not add another manager.
0x2100a783	ADMIN	system.moab	ERROR	MWM_CRED_MANAGER_OVERFLOW_CHILD	CredManager overflow while adding managers to child in fstree.	Fairshare tree configuration problem.
0x2100a786	ADMIN	system.moab	ERROR	MWM_UNABLE_TO_SELECT_TASKS_FOR_	Cannot select tasks for job %s. (%s)	Cannot select a node list that matches the

Cod e	Es cal ation Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
			R	JOB		requirements for this job. This might not be serious since multiple passes may occur.
<b>0x2100a787</b>	ADMI N	system.moab	ERROR	MWM_CANNOT_GET_TASK_ON_RESERVATION	Cannot get tasks on (ERR: %s/no reservation/iteration %s).	Cannot select tasks that meet the requirements.
<b>0x2100a78a</b>	ADMI N	system.moab	ERROR	MWM_BEST_VAL_ACHIEVED_BUT_SCHEDULE_EMPTY	BestVal %s achieved but schedule is empty.	Best value has been set, but the schedule is empty.
<b>0x2100a78b</b>	ADMI N	system.moab	ERROR	MWM_JOB_SCHEDULING_FAILURE_NO_RESERVATION	Scheduling failure %s (policy violation/no reservation) iteration: %s. (%s)	The job was not scheduled because no reservations are available.
<b>0x2100a78c</b>	ADMI N	system.moab	ERROR	MWM_UNSUPPORTED_SERVICE	Service '%s' (%s) not supported.	A request for an unsupported service was sent.
<b>0x2100a78d</b>	ADMI N	system.moab	ERROR	MWM_INVALID_CLASS_HOST_EXPRESSION	Invalid class host expression received (%s) : %s.	Failed to expand the class's host pattern to a list.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
				N		
<b>0x2100a78e</b>	ADMIN	system.moa	ERROR	MWM_TOO_MANY_COALLOCATION_REQUESTS	Too many co-allocation requests (%s > %s).	Too many co-allocation requests were received.
<b>0x2100a78f</b>	ADMIN	system.moa	ERROR	MWM_INVALID_JOBID_COUNTER	Min Job ID '%s' must be less than Max Job ID '%s'.	Invalid job ID was encountered.
<b>0x2100a790</b>	ADMIN	system.moa	ERROR	MWM_PARAMETER_NOT_HANDLED	Parameter[%s] '%s' not handled.	The specified parameter was not handled due to an unknown format.
<b>0x2100a791</b>	ADMIN	system.moa	ERROR	MWM_CIRCULAR_JOB_DEPENDENCY	Job cannot be dependent on itself.	The job is trying to use itself as a dependency, which creates a circular dependency and is invalid.
<b>0x2100a792</b>	ADMIN	system.moa	ERROR	MWM_CANNOT_CREATE_AM	Cannot create AM %s.	Could not create account manager object.
<b>0x2100a793</b>	ADMIN	system.moa	ERROR	MWM_INVALID_FLUSH_INTERVAL	%s for AM %s.	An invalid flush interval has been entered.



Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
			R			
<b>0x2100a794</b>	ADMIN	system.moab	ERROR	MWM_FAILED_SERVER_AUTH	Unable to authenticate server.	The server could not be authenticated.
<b>0x2100a795</b>	ADMIN	system.moab	ERROR	MWM_NO_QUOTE	No quote output provided in response.	No quote output provided in response.
<b>0x2100a796</b>	ADMIN	system.moab	ERROR	MWM_UNABLE_TO_PARSE_XML	Unable to parse XML (%s): %s.	Unable to parse XML data.
<b>0x2100a797</b>	ADMIN	system.moab	ERROR	MWM_INVALID_QUOTE	Invalid quote amount (%s).	Quote is invalid.
<b>0x2100a798</b>	ADMIN	system.moab	ERROR	MWM_RECURRING_COST	Unable to determine recurring cost.	Unable to determine recurring cost.
<b>0x2100a799</b>	ADMIN	system.moab	ERROR	MWM_AVAILABLE_PORT_NOT_FOUND	Cannot locate an available port for listening.	After trying to bind to a large number of ports, none were found to be available. Check network socket status

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
						for saturation.
<b>0x2100a79a</b>	ADMIN	system.moab	ERROR	MWM_CANNOT_RESOLVE_IP_FROM_HOSTNAME	Cannot resolve IP address from hostname '%s', getaddrinfo() rc: %s (%s).	There is a failure matching an IP address to a hostname. Check DNS, /etc/hosts or applicable nameservice.
<b>0x2100a79b</b>	ADMIN	system.moab	ERROR	MWM_UNKNOWN_CHECKPOINT_TYPE	Unexpected checkpoint type, %s.	Unknown checkpoint type while reading from the file.
<b>0x2100a79c</b>	ADMIN	system.moab	ERROR	MWM_CHECKPOINT_FILE_LINE_NOT_HANDLED	Line '%s' not handled in checkPoint file '%s'.	Please contact Adaptive Computing for assistance.
<b>0x2100a79d</b>	ADMIN	system.moab	ERROR	MWM_CANNOT_ADD_DEFAULT_GROUP	Cannot add default group.	Default group cannot be added.
<b>0x2100a79e</b>	ADMIN	system.moab	ERROR	MWM_CANNOT_ADD_GROUP	Cannot add group %s.	Group cannot be added.
<b>0x2100a79f</b>	ADMIN	system.moab	ERROR	MWM_ACCOUNT_NOT_ACCESSIBL	Account '%s' is not accessible by job '%s'.	The job is not authorized to run under the listed account.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
			R	E_BY_JOB		
0x2100a7a0	ADMIN	system.mom.ab	ERROR	MWM_CANNOT_DETERMINE_DEFAULT_ACCOUNT	Unable to determine default account for job '%s', user '%s'.	There is not a default account type for this job.
0x2100a7a1	ADMIN	system.mom.ab	ERROR	MWM_CANNOT_CREATE_RESERVATION	Cannot create reservation for job '%s'.	Failed to create reservation for job.
0x2100a7a2	ADMIN	system.mom.ab	ERROR	MWM_INVALID_NODELIST_BAD_TASKCOUNT	Invalid nodelist for job %s:%s (inadequate taskcount, %s < %s).	Invalid nodelist due to inadequate task count.
0x2100a7a3	ADMIN	system.mom.ab	ERROR	MWM_INVALID_NODELIST_BAD_NODECOUNT	Invalid nodelist for job %s:%s (inadequate nodecount, %s < %s).	Invalid nodelist due to inadequate node count.
0x2100a7a4	ADMIN	system.mom.ab	ERROR	MWM_INVALID_ALLOCATION_POLICY	Invalid allocation policy (%s).	Invalid allocation policy.
0x2100a7a5	ADMIN	system.mom.ab	ERROR	MWM_NO_MEMORY_FOR_ALLOCPARTITION	Cannot set ALLOCPARTITION variable for job %s (no memory).	No memory remaining to create job variable.

Cod e	Es cal ation Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
				VARIABLE		
<b>0x2 100 a7a 6</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ BASIL_ RSVID_NOT_ FOUND	Cannot locate BASIL RSVID (job 'ALLOCPARTITION' variable) that was just created.	Cannot locate BASIL reservation ID stored in the ALLOCPARTITION variable.
<b>0x2 100 a7a 7</b>	AD MI N	syste m.mo ab	E R R O R	MWM_JOB_ ADD_CLASS_ ATTR	Cannot add class for job %s (Class: %s).	Unable to add a class requirement attribute to a job.
<b>0x2 100 a7a 8</b>	AD MI N	syste m.mo ab	E R R O R	MWM_JOB_ ADD_DRM_ ATTR	Cannot set destination RM for job %s (RM: %s).	Unable to add a destination resource manager attribute to a job.
<b>0x2 100 a7a 9</b>	AD MI N	syste m.mo ab	E R R O R	MWM_JOB_ FLAGS_ INVALID_ SOURCE	Attempting to set job flags from invalid format.	Job flags must be created using documented formats.
<b>0x2 100 a7a a</b>	AD MI N	syste m.mo ab	E R R O R	MWM_SET_ SIZE_ON_ NONEXISTE NT_REQ	Requirement must be created before size is set.	Unable to set the size of an unallocated requirement.
<b>0x2 100 a7a b</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ADD_ GROUP_TO_ JOB_ FAILURE	Cannot add group for job %s (Group: %s).	Unable to set a group attribute on a job.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x2100a7ac</b>	ADMIN	system.moab	ERROR	MWM_NULL_JOB_NAME	Cannot add an empty name as an alternate name attribute for job %s.	No value specified. Make sure the alternate job name has a value.
<b>0x2100a7ad</b>	ADMIN	system.moab	ERROR	MWM_SPACES_IN_JOB_NAME	Attempted to set a job name (%s) with space(s) for job %s.	A job name with space(s) was specified. Job names cannot contain embedded spaces.
<b>0x2100a7ae</b>	ADMIN	system.moab	ERROR	MWM_ADD_QOS_TO_JOB_FAILURE	Cannot add QOS for job %s (QOS: %s).	Unable to set a QOS attribute on a job.
<b>0x2100a7af</b>	ADMIN	system.moab	ERROR	MWM_ADD_SRM_TO_JOB_FAILURE	Cannot add Submit RM for job %s (RM: %s).	Unable to find the entered name as an available resource manager.
<b>0x2100a7b0</b>	ADMIN	system.moab	ERROR	MWM_ADD_VARIABLE_TO_JOB_FAILURE	Cannot set variable for job %s (no variable name specified).	Only variables with names can be added as a job attribute.
<b>0x2100a7b1</b>	ADMIN	system.moab	ERROR	MWM_ADD_USER_TO_JOB_FAILURE	Cannot add user for job %s (User: %s).	Unable to set a user attribute on a job.

Cod e	Es calati on Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
<b>0x2 100 a7b 2</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ADD_ NODE_TO_ JOB_ FAILURE	Cannot add node for job %s (Node: %s).	Unable to set a node attribute on a job.
<b>0x2 100 a7b 3</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ADD_ ACCOUNT_ TO_JOB_ FAILURE	Cannot add account for job %s (Name: %s).	Failed to add account to the job.
<b>0x2 100 a7b 5</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ INVALID_ TIME_ STRING	Invalid format for time specification: '%s'.	A string that describes a time cannot be parsed because the format is wrong, or the values are out of range.
<b>0x2 100 a7b 6</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ CANNOT_ FIND_ ARRAY_JOB	Cannot find array job at index %s for job '%s'.	Array job is missing.
<b>0x2 100 a7b 7</b>	AD MI N	syste m.mo ab	E R R O R	MWM_JOB_ BUFFER_ FULL	Job buffer is full (ignoring job '%s').	Ignoring job since job buffer is full. Try increasing the value specified for the MAXJOB parameter.
<b>0x2 100 a7b 8</b>	AD MI N	syste m.mo ab	E R R O	MWM_ CANNOT_ FIND_ MASTER_	Cannot find master job (%s) for job '%s'; job array slot limits may not be enforced.	Cannot find the master job that is associated

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
			R	JOB		with a job array.
<b>0x2100a7b9</b>	ADMIN	system.moab	ERROR	MWM_INVALID_ACTION_STRING	The action string (%s) is invalid.	The format of the action string is '<operation type>:<operation ID>:<operation action>' Example: job:145+146+147:cancel where 145,146 and 147 are job IDs.
<b>0x2100a7ba</b>	ADMIN	system.moab	ERROR	MWM_INVALID_OBJECT_TYPE	The object type %s is invalid.	The format of the action string is '<operation type>:<operation ID>:<operation action>' Example: job:145+146+147:cancel where 145,146 and 147 are job IDs.
<b>0x2100a7bb</b>	ADMIN	system.moab	ERROR	MWM_JOB_NOT_FOUND	Unable to locate job %s.	The named job was not located in the system.
<b>0x2</b>	AD	system	E	MWM_JOB_	Completed trigger action is	The job is not

Cod e	Es cal ation Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
<b>100 a7b c</b>	MI N	m.mo ab	R R O R	IN_BAD_STATE_FOR_COMPLETE	specified for job %s but it is in an invalid state.	a system job and is not allowed to be started by the resource manager.
<b>0x2 100 a7b d</b>	AD MI N	syste m.mo ab	E R R O R	MWM_JOB_CANNOT_BE_HELD	Job %s cannot be put into hold state.	The resource manager cannot hold the job, usually because the job is not in a state that can be held.
<b>0x2 100 a7b e</b>	AD MI N	syste m.mo ab	E R R O R	MWM_CANNOT_SET_TRIGVAR	Cannot set trigger variable on job %s.	The trigger variables on a job cannot be set.
<b>0x2 100 a7bf</b>	AD MI N	syste m.mo ab	E R R O R	MWM_CANNOT_SET_REQATTR	Cannot set request attribute variable on job %s.	The request attribute variables on a job cannot be set.
<b>0x2 100 a7c 0</b>	AD MI N	syste m.mo ab	E R R O R	MWM_CANNOT_ADJUST_GRES	Cannot adjust generic resources for job %s.	The generic resources of the job could not be modified.
<b>0x2 100 a7c 1</b>	AD MI N	syste m.mo ab	E R R O R	MWM_INVALID_GRES_VALUE	Invalid value '%s' for GRes '%s' %s.	The value being set on the generic resource is not valid.



Cod e	Es cal ation Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
<b>0x2 100 a7c 2</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ CANNOT_ MODIFY_ ATTRIBUTE	Attribute %s cannot be modified for job %s.	The job's attribute could not be modified.
<b>0x2 100 a7c 3</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ COULD_ NOT_SEND_ SIGNAL	Signal %s could not be sent to job %s.	The resource manager was unable to send the signal to the job.
<b>0x2 100 a7c 4</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ COULD_ NOT_ START_JOB	Could not start job %s in %s.	The resource manager was unable to start the job.
<b>0x2 100 a7c 5</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ UNABLE_ TO_ REQUEUE_ JOB	Cannot requeue job %s.	The job could not be requeued.
<b>0x2 100 a7c 6</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ UNHANDLE D_ACTION	The action %s was not handled.	The action was undefined in this function.
<b>0x2 100 a7c 7</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ UNRECOGN IZED_ ATTRIBUTE	The attribute %s is not recognized.	The attribute is not in the lookup table.
<b>0x2 100 a7c 8</b>	AD MI N	syste m.mo ab	E R R O	MWM_ UNRECOGN IZED_JOB_ ACTION	The job action %s is not recognized.	The job action is not in the lookup table.

Cod e	Es cal ation Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
			R			
<b>0x2 100 a7c 9</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ UNABLE_ TO_ CANCEL_JOB	Job %s could not be canceled.	The job could not be canceled.
<b>0x2 100 a7c a</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ UNABLE_ TO_HOLD_ JOB	Job %s could not be held in.	The job was unable to be put into a hold state.
<b>0x2 100 a7c b</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ INVALID_ PBS_ SBINDIR	Invalid SBINDIR specified (%s).	Check paths for the directory containing pbs_iff.
<b>0x2 100 a7cc</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ UNABLE_ TO_ CONNECT_ PBS_SRVR	Cannot connect to PBS server '%s'; rc: %s (pbs_errno=%s, '%s').	Make sure the pbs_server process is running.
<b>0x2 100 a7c d</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ UNABLE_ GET_SRVR_ INFO	Cannot get server info: %s.	Make sure that the pbs_server process is running.
<b>0x2 100 a7c e</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ UNABLE_ LOAD_ SRVR_INFO	Cannot load PBS server info: %s.	Make sure that the pbs_server process is running.
<b>0x2</b>	AD	syste	E	MWM_	Cannot load PBS cluster	Make sure that

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>100a7cf</b>	MIN	m.moab	RROR	UNABLE_LOAD_PBS_CLUSTER	info: %s.	the pbs_server process is running.
<b>0x2100a7d0</b>	ADMIN	system.moab	RROR	MWM_UNABLE_LOAD_PBS_WORKLOAD	Cannot load PBS workload info: %s.	Make sure that the pbs_server process is running.
<b>0x2100a7d1</b>	ADMIN	system.moab	RROR	MWM_UNABLE_LOAD_PBS_QUEUE	Cannot load PBS queue info: %s.	Make sure the path to the queue configuration is accessible by Moab.
<b>0x2100a7d3</b>	ADMIN	system.moab	RROR	MWM_UNABLE_PROCESS_NODE_INFO	Cannot process node info.	Make sure the resource manager is running.
<b>0x2100a7d4</b>	ADMIN	system.moab	RROR	MWM_NODE_BUFFER_FULL	Node buffer is full (ignoring node '%s').	Try increasing the node buffer.
<b>0x2100a7d5</b>	ADMIN	system.moab	RROR	MWM_JOB_CANNOT_START	Job '%s' cannot be started: (cannot generate Tasklist).	Check the PBS server log to see reason of failure.
<b>0x2100a7d6</b>	ADMIN	system.moab	RROR	MWM_JOB_CANNOT_START_TASK	Job '%s' cannot be started: (empty Tasklist).	Check the PBS server log to see reason of failure.

Cod e	Es calati on Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
			R	EMPTY		
<b>0x2 100 a7d 7</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ UNABLE_ SET_NODE_ COUNT	Cannot set nodecount for job '%s' - %s.	Check the PBS server log to see reason of failure.
<b>0x2 100 a7d 8</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ UNABLE_ SET_ WALLTIME	Cannot set walltime for job '%s' - %s.	Check the PBS server log to see reason of failure.
<b>0x2 100 a7d 9</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ UNABLE_ SET_ TASKLIST	Cannot set Tasklist for job '%s' - %s.	Check the PBS server log to see reason of failure.
<b>0x2 100 a7d a</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ UNABLE_ TO_START_ JOB_RC	Job '%s' cannot be started: (rc: %s; errmsg: '%s'; Tasklist: '%s').	Check the PBS server log to see reason of failure.
<b>0x2 100 a7d b</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ UNABLE_ TO_SIGNAL_ JOB	%s' cannot be signaled: %s.	Check the PBS server log to see reason of failure.
<b>0x2 100 a7d c</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ UNABLE_ TO_ SUSPEND_ JOB	Job '%s' cannot be suspended: %s.	Check the PBS server log to see reason of failure.
<b>0x2</b>	AD	syste	E	MWM_	Job '%s' cannot be	Check the PBS

Cod e	Es cal ation Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
<b>100a7dd</b>	MI N	m.mo ab	R R O R	UNABLE_ TO_ RESUME_ JOB	resumed: %s.	server log to see reason of failure.
<b>0x2100a7de</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ UNABLE_ TO_FIND_ RESOURCE	Failed to find/add %s generic resource.	Failure to find/add GPUs/MICs to the global GRES/MIC slots.
<b>0x2100a7df</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ UNABLE_ TO_SET_ CREDENTIALS	Cannot authenticate job '%s' (U: %s; G: %s; A: '%s').	Could not set the credentials on the job.
<b>0x2100a7e0</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ UNABLE_ TO_ QUEUE	PBS job '%s' cannot be requeued (rc: %s; '%s').	Check the PBS server log to see reason of failure.
<b>0x2100a7e1</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ UNABLE_ TO_ CHECKPOINT	PBS job '%s' cannot be checkpointed (rc: %s; '%s').	Check the PBS server log to see reason of failure.
<b>0x2100a7e2</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ UNABLE_ TO_ RELEASE	PBS job '%s' cannot be released from hold (rc: %s; '%s').	Check the PBS server log to see reason of failure.
<b>0x2100a7e5</b>	AD MI N	syste m.mo ab	E R R O	MWM_ UNABLE_ TO_FIND_ ACCOUNT	Cannot find account for job %s (Name: %s).	Make sure the account exists.

Cod e	Es calati on Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
			R			
<b>0x2100a7e6</b>	ADMI N	system.moab	ERROR	MWM_INVALID_ARGUMENT	Command '%s' args not handled.	An unsupported argument was used.
<b>0x2100a7e7</b>	ADMI N	system.moab	ERROR	MWM_INVALID_LOGDIR	LogDir '%s' is invalid.	Make sure that the path to the logs directory exists.
<b>0x2100a7e8</b>	ADMI N	system.moab	ERROR	MWM_INVALID_SPOOLDIR	SpoolDir '%s' is invalid.	Make sure that the path to the spool directory exists.
<b>0x2100a7e9</b>	ADMI N	system.moab	ERROR	MWM_INVALID_STATDIR	StatDir '%s' is invalid.	Make sure that the path to the stat directory exists.
<b>0x2100a7ea</b>	ADMI N	system.moab	ERROR	MWM_INVALID_TOOLS DIR	ToolsDir '%s' is invalid.	Make sure that the path to the tools directory exists.
<b>0x2100a7eb</b>	ADMI N	system.moab	ERROR	MWM_CANNOT_CREATE_DAT_FILE	Cannot create/modify dat file: '%s'.	Moab encountered an error creating the dat file.
<b>0x2</b>	AD	system	E	MWM_	The '%s' feature is not	Moab can be

Cod e	Es calati on Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
<b>100a7ec</b>	MIN	m.moab	RROR	FEATURE_NOT_AVAILABLE_IN_BUILD	available in the build of Moab.	configured with various features. The listed feature is not available in the binary being run.
<b>0x2100a7ed</b>	ADMIN	system.moab	ERROR	MWM_FEATURE_NOT_AVAILABLE_WITH_LICENSE	The '%s' feature is not enabled with the current Moab license.	Moab can be licensed with various features. The listed feature is not available with the current license. Contact Adaptive Computing for more information.
<b>0x2100a7ee</b>	ADMIN	system.moab	ERROR	MWM_RESOURCE_LIMIT_EXCEEDED	The maximum number of '%s' (%s) has been reached.	Moab has certain resources that are limited. This error occurs when you have reached or exceeded those limits. Contact Adaptive Computing for more information.
<b>0x2100a7f1</b>	ADMIN	system.moab	ERROR	MWM_CANNOT_CREATE_VM_	Failed to create migration job for VM %s.	The migration job was not created. Check MIGRATETEM

Cod e	Es cal ation Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
			R	MIGRATIO N_JOB		PLATE on workflow and its trigger.
<b>0x2 100 a7f2</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ CANNOT_ OPEN_ EXTENSIO N_ INTERFACE	Cannot open extension interface socket on port %s.	There was a failure opening the HTTP extension service. This feature will not work until the problem is corrected.
<b>0x2 100 a7f3</b>	AD MI N	syste m.mo ab	E R R O R	MWM_JOB_ USER_ AUTHENTIC ATION	The system was unable to connect the given user to job %s (User: %s, Group: %s).	Check the credentials of the given user and/or group.
<b>0x2 100 a7f4</b>	AD MI N	syste m.mo ab	E R R O R	MWM_JOB_ AUTHENTIC ATION	The system was unable to authenticate the user connected with job %s (User: %s, Group: %s, Account %s) - %s.	Check the credentials of the given user and/or group.
<b>0x2 100 a7f5</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ SEND_ DATA_ FAILED	The system was unable to send data to the server %s (%s:%s).	Make sure the server's address is correct and that the server is running.
<b>0x2 100 a7f6</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ RECEIVE_ DATA_ FAILED	The system was unable to receive data from the server %s (%s:%s).	Make sure the server's address is correct and that the server is running.



Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x2100a7f7</b>	ADMIN	system.moab	ERROR	MWM_JOB_OVERLAP	Job '%s' overlaps an existing job.	Check the job being created for overlap.
<b>0x2100a7f8</b>	ADMIN	system.moab	ERROR	MWM_JOB_CREATION	The system was unable to create job '%s'	Verify that the job being created is correctly specified.
<b>0x2100a7f9</b>	ADMIN	system.moab	ERROR	MWM_MISSING_STATUS_ELEMENT	The status element was missing from the S3 response.	This is an internal error.
<b>0x2100a7fa</b>	ADMIN	system.moab	ERROR	MWM_VC_WORKFLOW_JOB	Virtual container '%s' was marked as workflow, but could not find job that created it.	This is an internal error.
<b>0x2100a7fb</b>	ADMIN	system.moab	ERROR	MWM_VC_COMBINE_JOBS	Failed to combine jobs in virtual container '%s'.	This is an internal error.
<b>0x2100a7fc</b>	ADMIN	system.moab	ERROR	MWM_VC_SCHEDULE_TIME_FAILURE	Failed to schedule virtual container '%s' for requested time.	This is an internal error.
<b>0x2100a7fd</b>	ADMIN	system.moab	ERROR	MWM_VC_RESERVATION_FAILURE	Failed to find a reservation for virtual container '%s'.	This is an internal error.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
			R			
<b>0x2100a7fe</b>	ADMIN	system.moab	ERROR	MWM_VC_RESERVATION_CREATE_FAILURE	Failed to create a reservation for jobs in virtual container '%s'.	This is an internal error.
<b>0x2100a7ff</b>	ADMIN	system.moab	ERROR	MWM_VC_RESOURCE_FAILURE	Requested resources are not available at any time for virtual container '%s'.	This is an internal error.
<b>0x2100a800</b>	ADMIN	system.moab	ERROR	MWM_NONEXISTING_JOB_USER	Job template %s requests non-existent user %s.	Make sure the user exists.
<b>0x2100a801</b>	ADMIN	system.moab	ERROR	MWM_NONEXISTING_JOB_GROUP	Job template %s requests non-existent group %s.	Make sure the group exists.
<b>0x2100a802</b>	ADMIN	system.moab	ERROR	MWM_NONEXISTING_JOB_QOS	Job template %s requests non-existent QoS %s.	Make sure the QoS exists.
<b>0x2100a803</b>	ADMIN	system.moab	ERROR	MWM_UNABLE_CREATE_CLASS	Unable to create class %s for job template %s.	Make sure the class exists.
<b>0x2</b>	AD	system	E	MWM_	Job template %s requests	Make sure the

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
100a804	MIN	m.moab	ERROR	NONEXISTING_JOB_ACCOUNT	non-existent account %s.	account exists.
0x2100a805	ADMIN	system.moab	ERROR	MWM_INVALID_WALLTIME_SPECIFIED	Invalid walltime specification '%s'.	Make sure the format for walltime is correct.
0x2100a806	ADMIN	system.moab	ERROR	MWM_UNABLE_PARSE_WIKI_STR	Cannot parse wiki string for job '%s'.	Make sure the format for wiki string is correct.
0x2100a807	ADMIN	system.moab	ERROR	MWM_MISSING_STATS_XML_ELEMENT	%s is not a valid template job stat child element.	Make sure there is a stats element in the XML.
0x2100a808	ADMIN	system.moab	ERROR	MWM_NULL_NODE_POINTER	Node pointer is NULL and cannot be used to find SMP node.	Node pointer is NULL and cannot be used to find SMP node by node.
0x2100a809	ADMIN	system.moab	ERROR	MWM_PINDEX_OUT_OF_RANGE	PIndex is less than -1 which is out of range.	PIndex must be greater than or equal to -1 to find a node by partition.
0x2100a80a	ADMIN	system.moab	ERROR	MWM_FEATURE_OUT_OF_RANGE	Feature is less than -1 which is out of range.	Feature must be greater than or equal

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
			R			to -1 to find a node by feature.
<b>0x2100a80b</b>	ADMIN	system.moab	ERROR	MWM_INCORRECT_ARG	Incorrect argument in %s: %s, %s,%s.	Name must point to a valid string, Feature must be greater than or equal to -1, and N must point to a valid node.
<b>0x2100a80c</b>	ADMIN	system.moab	ERROR	MWM_NODE_ALLOCATION_ERROR	Failed to allocate a node named %s.	Call to MUMalloc failed, system is probably low on memory.
<b>0x2100a80d</b>	ADMIN	system.moab	ERROR	MWM_FAILED_TO_APPEND_MSMPNODE	Failed to append smpnode %s to MSMPNodes.	The call to append the node to the array list failed, probably due to a low memory condition.
<b>0x2100a80e</b>	ADMIN	system.moab	ERROR	MWM_NULL_SMPNODE_POINTER	Cannot initialize node because pointer is NULL.	Call to MSMPNodeInitiaize must have a valid pointer to a valid node.
<b>0x2100</b>	ADMIN	system.moab	ERROR	MWM_NULL_	Cannot reset node because pointer is NULL.	Call to MSMPNodeRes

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>a80f</b>	N	ab	R O R	SMPNODE_ POINTER_ IN_RESET		etStats must have a valid pointer to a valid node.
<b>0x2 100 a81 0</b>	AD MI N	system.mo ab	E R R O R	MWM_ RESET_ NODE_ FAILED	Call to MSMPNodeResetStats failed.	Call to MSMPNodeResetStats failed. The most likely cause is passing a NULL pointer to SMPNode.
<b>0x2 100 a81 1</b>	AD MI N	system.mo ab	E R R O R	MWM_ FREE_ NODE_ FAILED	Call to free MSMPNodes failed.	Call to free MSMPNodes failed, most likely due to corrupted memory.
<b>0x2 100 a81 2</b>	AD MI N	system.mo ab	E R R O R	MWM_ NULL_ NODE_IN_ UPDATE	Node pointer in %s cannot be NULL.	Node pointer cannot be NULL when trying to update node.
<b>0x2 100 a81 3</b>	AD MI N	system.mo ab	E R R O R	MWM_ CANNOT_ FIND_NODE	Unable to find SMP node with node %s.	Unable to find SMP node by node.
<b>0x2 100 a81 4</b>	AD MI N	system.mo ab	E R R O R	MWM_ EMPTY_ NODE_LIST	Updating node from list with empty node list.	Updating node from node list must not be called with an empty node list.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
0x2100a815	ADMIN	system.moab	ERROR	MWM_BAD_ARG_IN_FEASIBLE_JOB	Incorrect argument to function %s: %s, %s.	A parameter in the function was incorrect.
0x2100a816	ADMIN	system.moab	ERROR	MWM_CANNOT_FIND_INDEX_IN_LIST_FOR_FEATURE	Could not find index into NodeSetList for node feature %s.	Could not find index into NodeSetList for node feature.
0x2100a817	ADMIN	system.moab	ERROR	MWM_CANNOT_CREATE_SR	Could not create standing reservation: %s.	Failed to create the named standing reservation.
0x2100a818	ADMIN	system.moab	ERROR	MWM_UNEXPECTED_STATISTICS_TYPE	Unexpected statistics type: %s.	Number is not a member of MMStatTypeEnum enumeration.
0x2100a819	ADMIN	system.moab	ERROR	MWM_CANNOT_PROCESS_VM_ATTRIBUTE	Cannot process VM attribute %s for VM %s.	Either AttrName or NodeName is not found in string.
0x2100a81a	ADMIN	system.moab	ERROR	MWM_CANNOT_FIND_NODE_FOR_VM	Cannot find node %s for VM %s.	The node does not exist or cannot be found.
0x2100a81	ADMIN	system.moab	ERROR	MWM_CANNOT_LOAD_JOB	Cannot load job %s (state: %s).	There was an error creating

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>b</b>			R O R			a job in Moab that was reported by the resource manager.
<b>0x2100a81c</b>	ADMIN	system.moab	E R R O R	MWM_CANNOT_CREATE_CHECKPOINT_FILE_ENTRY	Cannot create checkpoint file entry.	There was an error writing a checkpoint file entry for the associated objects.
<b>0x2100a81d</b>	ADMIN	system.moab	E R R O R	MWM_CANNOT_CREATE_OBJECT_FROM_CHECKPOINT_FILE	Cannot create object from checkpoint file entry.	There was an error reading a checkpoint file entry for the associated objects.
<b>0x2100a81e</b>	ADMIN	system.moab	E R R O R	MWM_TASKLIST_TOO_LARGE	The tasklist for job '%s' is too large (size = %s, growth = %s).	The system has a fixed maximum size for the task map for each job.
<b>0x2100a81f</b>	ADMIN	system.moab	E R R O R	MWM_TASKLIST_MISSING	The tasklist for job '%s' is missing.	The system requires that each job has at least one task assigned.
<b>0x2100a820</b>	ADMIN	system.moab	E R R O R	MWM_TASK_DISTRIBUTION_UNKNOWN	The system encountered an unknown type of task distribution (%s).	This is an internal error.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
0x2100a822	ADMIN	system.moab	ERROR	MWM_INCOMPATIBLE_CHARGE_POLICY	Periodic charging disabled due to incompatible job charge policy (%s).	The job charge policy is undefined.
0x2100a823	ADMIN	system.moab	ERROR	MWM_INCOMPLETE_JOB_TEMPLATE_ACTION	The job template '%s' has an incomplete action specification.	Job templates must fully specify the action to be performed.
0x2100a824	ADMIN	system.moab	ERROR	MWM_INCOMPLETE_JOB_TEMPLATE_GENERIC	The job template '%s' has an incomplete generic system job specification.	Job templates must fully specify the generic system job.
0x2100a825	ADMIN	system.moab	ERROR	MWM_DUPLICATE_JOB_TEMPLATE_VMID	The job template '%s' has a job '%s' that requests an existing VMID.	Virtual machine IDs cannot be shared across job templates.
0x2100a826	ADMIN	system.moab	ERROR	MWM_UNKNOWN_JOB_TEMPLATE_VMID	The requested VMID '%s' could not be found or already has a tracking job.	Virtual machine IDs can only be assigned to a single job.
0x2100a827	ADMIN	system.moab	ERROR	MWM_UNABLE_TO_MODIFY_JOB	The job '%s' on account '%s' cannot be modified in the resource manager.	The job previously submitted to the resource manager cannot be modified.
0x2	ADMIN	system.moab	E	MWM_	The system failed to generate a workflow virtual	This is an internal error.



Cod e	Es cal ation Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
<b>100a828</b>	N	ab	R R O R	WORKFLOW_VC_FAILURE	container for job '%s'.	
<b>0x2100a829</b>	AD MI N	syste m.mo ab	E R R O R	MWM_CREATE_JOB_TEMPLATE_FAILURE	The system failed to create job template '%s'.	The job could not be created or one of its attributes could not be set.
<b>0x2100a82a</b>	AD MI N	syste m.mo ab	E R R O R	MWM_VC_NOT_FOUND	The system could not find the virtual container for job '%s'.	This is an internal error.
<b>0x2100a82b</b>	AD MI N	syste m.mo ab	E R R O R	MWM_JOB_MIGRATION_FAILED	The system failed to migrate a remote job (%s).	Make sure the resource manager has not been disabled.
<b>0x2100a82c</b>	AD MI N	syste m.mo ab	E R R O R	MWM_JOB_START_XML_FAILURE	The system could not generate the command line needed to start job: '%s'.	The proper command line could not be derived from the XML structure.
<b>0x2100a82d</b>	AD MI N	syste m.mo ab	E R R O R	MWM_JOB_START_FAILURE_RESPONSE	The system could not start job - Reason: '%s'.	The system was unable to start the job for the specified reason.
<b>0x2</b>	AD MI	syste m.mo	E	MWM_JOB_	The system could not cancel job - Reason: '%s'.	The system

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>100a82e</b>	N	ab	RROR	CANCEL_FAILURE_RESPONSE		was unable to cancel the job for the specified reason.
<b>0x2100a82f</b>	ADMIN	system.moab	RROR	MWM_JOB_SIGNAL_FAILURE_RESPONSE	The system could not signal job - Reason: '%s'.	The system was unable to signal the job for the specified reason.
<b>0x2100a830</b>	ADMIN	system.moab	RROR	MWM_JOB_MODIFY_FAILURE_RESPONSE	The system could not modify job - Reason: '%s'.	The system was unable to modify the job for the specified reason.
<b>0x2100a831</b>	ADMIN	system.moab	RROR	MWM_JOB_REQUEUE_FAILURE_RESPONSE	The system could not requeue job - Reason: '%s'.	The system was unable to requeue the job for the specified reason.
<b>0x2100a832</b>	ADMIN	system.moab	RROR	MWM_SEND_EVENT_FAILURE	The system could send event '%s' to resource manager '%s' (%s).	The system was unable to send the event.
<b>0x2100a833</b>	ADMIN	system.moab	RROR	MWM_UNEXPECTED_BACKFILL_POLICY	The system encountered an unexpected backfill policy '%s' (using '%s' instead).	The backfill policy did not match a defined policy.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
0x2100a834	ADMIN	system.moab	ERROR	MWM_NODE_LIST_ALLOCATION	The system was unable to allocate a node list for job '%s' in partition '%s'.	The system might be low on memory.
0x2100a835	ADMIN	system.moab	ERROR	MWM_BAD_NODE_IN_NODELIST	The reservation nodelist for job '%s' has an invalid node at index %s.	Check the nodes specified for the reservation.
0x2100a836	ADMIN	system.moab	ERROR	MWM_RESERVATION_SPANS_PARTITIONS	The reservation request for job '%s' spans partitions (node %s partition %s).	Reservations that span partitions must have the COALLOC flag set.
0x2100a837	ADMIN	system.moab	ERROR	MWM_ADJUST_JOB_RESERVATION_FAILURE	The system failed to adjust job '%s' reservation on node %s.	This is an internal error.
0x2100a838	ADMIN	system.moab	ERROR	MWM_OBJECT_TYPE_INVALID	The object type specified (%s) is not valid.	A valid object type must be specified.
0x2100a839	ADMIN	system.moab	ERROR	MWM_MISSING_OBJECT_ID	The object ID is missing.	A valid object ID must be specified.
0x2	ADMIN	system.moab	E	MWM_MISSING_	The action is missing.	A valid action must be

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>100a83a</b>	N	ab	R R O R	ACTION		specified.
<b>0x2100a83b</b>	ADMIN	system.moab	E R R O R	MWM_PIPE_BUFFER_FAILED	The system could not open a bi-directional pipe.	A valid action must be specified.
<b>0x2100a83c</b>	ADMIN	system.moab	E R R O R	MWM_STD_OUT_FAILED	Failed to load stdout file '%s'.	Check the file name and path.
<b>0x2100a83d</b>	ADMIN	system.moab	E R R O R	MWM_STD_ERR_FAILED	Failed to load stderr file '%s'.	Check the file name and path.
<b>0x2100a841</b>	ADMIN	system.moab	E R R O R	MWM_CREATE_NODE_FAILURE	Unable to create node '%s' (check license and MAXNODE parameter).	check license and MAXNODE parameter.
<b>0x2100a842</b>	ADMIN	system.moab	E R R O R	MWM_PARTITION_CREATE_FAILURE	The system was unable to create a shared partition for the global node.	The system might be low on memory.
<b>0x2100a845</b>	ADMIN	system.moab	E R R O R	MWM_HT_FIND_NODE_FAILURE	Cannot find node '%s' in hash table.	A node by the given name might not have been created.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
0x2100a847	ADMIN	system.moab	ERROR	MWM_HT_FIND_VM_FAILURE	Cannot find VM '%s' in hash table.	A VM with the given name might not have been created.
0x2100a848	ADMIN	system.moab	ERROR	MWM_COMMAND_FAILED	Command '%s' failed. StatusCode: %s; Response: '%s'.	Check the command syntax and parameters.
0x2100a849	ADMIN	system.moab	ERROR	MWM_HASH_TABLE_INITIALIZATION	There was an unexpected hash table initialization error.	The hash table for jobs to delete never initialized correctly.
0x2100a84e	ADMIN	system.moab	ERROR	MWM_UNABLE_TO_AUTHENTICATE_JOB	Unable to authenticate job %s when UID or GID is empty (UID=%s, GID=%s).	Either the UID or the GID field is empty.
0x2100a84f	ADMIN	system.moab	ERROR	MWM_MISSING_JOB_TASKCOUNT	Job does not have a taskcount specified.	Each job must have an associated taskcount.
0x2100a850	ADMIN	system.moab	ERROR	MWM_FAILED_EXCLUDE_NODELIST	The system failed to add an exclude nodelist to a submission.	The job exclude hostlist could not be converted into a string.
0x2100a85	ADMIN	system.moab	ERROR	MWM_CANNOT_	Failed to submit migration job for VM %s.	Check MIGRATETEM

Cod e	Es cal ation Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
<b>2</b>			R O R	SUBMIT_V M_ MIGRATIO N_JOB		PLATE on workflow and its trigger.
<b>0x2 100 a85 3</b>	AD MI N	syste m.mo ab	E R R O R	MWM_WEB_ SERVICES_ WRITE_ FAILURE	Error %s encountered while trying to write to web services.	Encountered problem trying to put HTTP data to web server.
<b>0x2 100 a85 4</b>	AD MI N	syste m.mo ab	E R R O R	MWM_WEB_ SERVICES_ URL_ MISSING	Missing URL in call to web services.	Web services must have a valid destination URL.
<b>0x2 100 a85 5</b>	AD MI N	syste m.mo ab	E R R O R	MWM_RM_ PARTITION_ CREATE_ FAILURE	The system was unable to create a partition for RM '%s'.	The system might be low on memory.
<b>0x2 100 a85 6</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ PARSE_MPP_ NODES_ FAILURE	The system failed to parse the MPP nodes value '%s'.	Check the MPP names.
<b>0x2 100 a85 7</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ FIND_MPP_ NODES_ FAILURE	The system failed to find node '%s' in the MPP nodes value '%s'.	Check the MPP names.
<b>0x2 100 a85 9</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ NODE_SET_ TYPE_ INVALID	The node set type specified (%s) is not valid.	Check the NODESETLIST option.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x2100a85a</b>	ADMIN	system.moab	ERROR	MWM_GRES_ADD_FAILURE	Unable to add the GRESTOJOBATTRMAP '%s'.	The limit has been reached.
<b>0x2100a85d</b>	ADMIN	system.moab	ERROR	MWM_NOT_MWS_RM	The resource manager is not Moab Web Services.	Make sure the resource manager has Moab Web Services.
<b>0x2100a85e</b>	ADMIN	system.moab	ERROR	MWM_MWS_RM_CURL_CONNECTION	The system could not initialize a cURL connection to the MWS RM.	The cURL command to connect to the resource manager has failed.
<b>0x2100a85f</b>	ADMIN	system.moab	ERROR	MWM_MWS_RM_CURL_CONNECTION_EXPANDED	Could not connect to MWS RM (%s) at '%s%s' as '%s', response code: %s; cURL error: %s (%s); MWS response: '%s'.	The connection has failed.
<b>0x2100a860</b>	ADMIN	system.moab	ERROR	MWM_MWS_RM_JSON_CLUSTER_QUERY_EMPTY	JSON cluster query data from MWS RM (%s) is null or empty.	The query must contain valid JSON data.
<b>0x2100a861</b>	ADMIN	system.moab	ERROR	MWM_MWS_RM_JSON_WORKLOAD_QUERY_EMPTY	JSON workload query data from MWS RM (%s) is null or empty.	The query must contain valid JSON data.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
0x2100a867	ADMIN	system.moab	ERROR	MWM_JOB_TRANSITION_FAILURE	Unable to transition a job.	The job was missing requirements.
0x2100a868	ADMIN	system.moab	ERROR	MWM_SET_JOB_VARIABLE	Unable to set a job pref variable.	The system is probably low on memory.
0x2100a869	ADMIN	system.moab	ERROR	MWM_ARRAY_EXPANSION	Unable to expand the size of an array.	The system is probably low on memory.
0x2100a86c	ADMIN	system.moab	ERROR	MWM_VC_FIND_FAILURE	The system could not find the virtual container '%s'.	Check the name of the VC.
0x2100a86d	ADMIN	system.moab	ERROR	MWM_VC_USER_CREDENTIALS	User '%s' does not have access to virtual container '%s'.	Check the rights granted to the VC.
0x2100a86e	ADMIN	system.moab	ERROR	MWM_VC_BEING_DELETED	Virtual container '%s' is being deleted; cannot add jobs to it.	Only add jobs to VCs that are not being deleted.
0x2100a86f	ADMIN	system.moab	ERROR	MWM_PARTITION_STATUS	Unable to query the status of a partition - %s.	Check to make sure the resource manager is



Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
			R			running.
<b>0x2100a870</b>	ADMIN	system.moab	ERROR	MWM_FIND_JOB_TEMPLATE	The system failed to find job template '%s'.	Check the template name for the given job.
<b>0x2100a873</b>	ADMIN	system.moab	ERROR	MWM_PROCESS_EVENT	Unable to process the generic event.	During processing, unable to get a description of the event.
<b>0x2100a876</b>	ADMIN	system.moab	ERROR	MWM_JOB_CAN_NEVER_RUN	Unable to allocate tasks for job at any time.	Job tasks must match available resources.
<b>0x2100a877</b>	ADMIN	system.moab	ERROR	MWM_NODE_NOT_IN_PARTITION	Node is not associated with any partition.	Node must be in a partition.
<b>0x2100a87a</b>	ADMIN	system.moab	ERROR	MWM_NODE_COUNT_EXCEEDS_LICENSE	The number of nodes '%s' exceeds the current license limit '%s'.	A different license is needed to use more nodes.
<b>0x2100a87b</b>	ADMIN	system.moab	ERROR	MWM_KILL_FAILURE	OS call to kill process (PID: %s) %s failed).	This is an operating system error.
<b>0x2</b>	AD	system	E	MWM_	Job does not have any	Each job must

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
100a87c	MIN	m.moab	RROR	MISSING_JOB_REQUIREMENTS	requirements specified.	have requirements attached.
0x2100a87d	ADMIN	system.moab	ERROR	MWM_JOB_MISSING_DISPATCH_TIME	Job loaded in alloc state '%s' with no dispatch time.	The job must have a dispatch time.
0x2100a87e	ADMIN	system.moab	ERROR	MWM_UNEXPECTED_OBJECT_TYPE	The object type '%s' was not expected in this operation.	Verify that a valid object type is given.
0x2100a87f	ADMIN	system.moab	ERROR	MWM_JOB_TRANSITION_XML	Unable to create XML element from job transition object.	The system might be low on memory.
0x2100a881	ADMIN	system.moab	ERROR	MWM_VM_CREATE_RESERVATION	Cannot create reservation for VM '%s'.	Failed to create reservation for the given VM.
0x2100aa08	ADMIN	system.moab	ERROR	MWM_VM_FIELD_VALUE	VM '%s' has an invalid '%s%s%s' field value.	The field value for the VM is invalid.
0x2100aa09	ADMIN	system.moab	ERROR	MWM_NODE_FIELD_VALUE	Node '%s' has an invalid '%s%s%s' field value.	The field value for the node is invalid.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x2100aa0a</b>	ADMIN	system.moab	ERROR	MWM_JOB_FIELD_VALUE	Job '%s' has an invalid '%s%s%s' field value.	The field value for the job is invalid.
<b>0x2100aa4a</b>	ADMIN	system.moab	ERROR	MWM_INCOMPLETE_JOB_TEMPLATE_DATASTAGING	The job template '%s' has an incomplete data staging system job specification.	Job templates must fully specify the data staging system job.
<b>0x2100aa55</b>	ADMIN	system.moab	ERROR	MWM_CANNOT_ADD_DEFAULT_USER	Cannot add default user.	Default user cannot be added.
<b>0x2100aa56</b>	ADMIN	system.moab	ERROR	MWM_CANNOT_ADD_DEFAULT_ACCOUNT	Cannot add default account.	Default account cannot be added.
<b>0x2100aa57</b>	ADMIN	system.moab	ERROR	MWM_CANNOT_ADD_DEFAULT_CLASS	Cannot add default class.	Default class cannot be added.
<b>0x2100aa58</b>	ADMIN	system.moab	ERROR	MWM_CANNOT_CREATE_CP_FILE	Cannot create checkpoint file. Cannot open %s file '%s', errno: %s (%s).	The fopen() system call failed. Use the errno and associated message to determine possible causes.

Cod e	Es calati on Level	Topi c	S e v e r i t y	Event Name	Message Template	Comment
<b>0x2 100 aa5 9</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ PAST_PROC_ SEC_LIMIT_ DAILY	Maximum daily processor seconds reached - Elastic Computing disabled. Total: %s, Max: %s	The daily threshold has been reached - Elastic Computing disabled.
<b>0x2 100 aa5 a</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ PAST_PROC_ SEC_LIMIT_ MONTHLY	Maximum monthly processor seconds reached - Elastic Computing disabled. Total: %s, Max: %s	The monthly threshold has been reached - Elastic Computing disabled.
<b>0x2 100 aa5 b</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ PAST_PROC_ SEC_LIMIT_ QUARTERL Y	Maximum quarterly processor seconds reached - Elastic Computing disabled. Total: %s, Max: %s	The quarterly threshold has been reached - Elastic Computing disabled.
<b>0x2 100 aa5 c</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ PAST_PROC_ SEC_LIMIT_ YEARLY	Maximum yearly processor seconds reached - Elastic Computing disabled. Total: %s, Max: %s	The yearly threshold has been reached - Elastic Computing disabled.
<b>0x2 100 aa5 d</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ PAST_QOS_ PROC_SEC_ LIMIT_ DAILY	Maximum daily processor seconds reached for QOS %s - Elastic Computing disabled. Total: %s, Max: %s	The daily QOS threshold has been reached - Elastic Computing disabled.
<b>0x2 100 aa5 e</b>	AD MI N	syste m.mo ab	E R R O R	MWM_ PAST_QOS_ PROC_SEC_ LIMIT_ MONTHLY	Maximum monthly processor seconds reached for QOS %s - Elastic Computing disabled. Total: %s, Max: %s	The monthly QOS threshold has been reached - Elastic

Cod e	Es cal ation Level	Topi c	S e v e r i t y	Event Name	Message Template	Comment
						Computing disabled.
<b>0x2100aa5f</b>	ADMIN	system.moa	ERROR	MWM_PAST_QOS_PROC_SEC_LIMIT_QUARTERLY	Maximum quarterly processor seconds reached for QOS %s - Elastic Computing disabled. Total: %s, Max: %s	The quarterly QOS threshold has been reached - Elastic Computing disabled.
<b>0x2100aa60</b>	ADMIN	system.moa	ERROR	MWM_PAST_QOS_PROC_SEC_LIMIT_YEARLY	Maximum yearly processor seconds reached for QOS %s - Elastic Computing disabled. Total: %s, Max: %s	The yearly QOS threshold has been reached - Elastic Computing disabled.
<b>0x2100aa61</b>	ADMIN	system.moa	ERROR	MWM_NODE_CANNOT_LICENSE	Node '%s' does not report socket/gpu/mic information. It cannot be licensed.	An RM that reports socket/gpu/mic information for each node must be used with this license type.
<b>0x2100c001</b>	INTERNAL	system.moa	ERROR	MWM_NOT_IMPLEMENTED	Function %s has not been implemented yet.	This error is used when we've stubbed out code but do not expect it to be called in production environments. It's not helpful except for internal diagnostics.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x2100e729</b>	INTERNAL	system.moab	ERROR	MWM_CANNOT_SEND_TO_SOCKET_DETAILED	Cannot send %s of %s bytes to socket descriptor %s - errno: %s (%s).	The send() system call failed. Socket is blocked (select() indicated socket was available--check MTU).
<b>0x2100e77b</b>	INTERNAL	system.moab	ERROR	MWM_CLIENT_COUNT_NEGATIVE	Client count fell below zero on socket %s.	This is an internal error. The number of client connections should always be zero or greater.
<b>0x2100e784</b>	INTERNAL	system.moab	ERROR	MWM_HOSTLIST_MISSING	A hostlist was specified but now it is NULL/EMPTY.	The job claims to have a specified hostlist, but at the current point in processing no list can be found. This is most likely an internal problem.
<b>0x2100e7b4</b>	INTERNAL	system.moab	ERROR	MWM_REQATTR_UNSUPPORTED_OPERATION	Operation (%s) not supported on required attributes (reqattrs).	See documentation for supported operators allows on required attributes (reqattrs).

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x2100e7ef</b>	INTERNAL	system.moab	ERROR	MWM_VM_NOT_LINKED_TO_TRACKING_JOB	VM '%s' not linked to VMTracking job '%s' (linked to job '%s').	A VM must be associated with a tracking job.
<b>0x31002a2c</b>	USER	system.moab	ALERT	MWM_NO_TASKS_FOUND_ON_JOB	No tasks found for job '%s'.	Check job submission arguments for desired requirements.
<b>0x31008385</b>	ADMIN	system.moab	FATAL	MWM_EXPIRED_LICENSE	%s License has expired.	A license file was found but it has expired. Please contact your sales representative at Adaptive Computing for assistance.
<b>0x31008386</b>	ADMIN	system.moab	FATAL	MWM_EVALUATION_EXPIRED	%s evaluation period has expired.	The evaluation period has expired. Please contact your sales representative at Adaptive Computing for assistance.
<b>0x31008387</b>	ADMIN	system.moab	FATAL	MWM_UNEXPECTED_LICENSE_ERROR	Moab will now exit. Unexpected error while reading license: %s	Moab was unable to verify that the license file was valid. Please contact your sales representative

Cod e	Es calati on Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
						at Adaptive Computing for assistance.
<b>0x3 100 a71 2</b>	AD MI N	syste m.mo ab	F A T A L	MWM_ UNABLE_ TO_ ALLOCATE_ MEMORY	Unable to allocate memory.	One or more calls to allocate memory failed.
<b>0x3 100 a71 4</b>	AD MI N	syste m.mo ab	F A T A L	MWM_ CANNOT_ RESTORE_ UID	Cannot restore EUID to '%s' for server, errno: %s (%s).	The setuid() system call failed. There was a failure resetting the UID of the process. This might be because the process is running as a different user. Use the errno and associated message to determine possible causes.
<b>0x3 100 a71 5</b>	AD MI N	syste m.mo ab	F A T A L	MWM_ CANNOT_ CHANGE_ UID	Cannot change UID to user '%s' (UID: %s) errno: %s (%s).	The setuid() system call failed. Use the errno and associated message to determine possible causes.



Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x3100a716</b>	ADMIN	system.moab	FATAL	MWM_CANNOT_RESTORE_GID	Cannot restore GID to '%s' for server, errno: %s (%s).	The setgid() system call failed. There was a failure resetting the GID of the process. This might be because the process is running in a different group. Use the errno and associated message to determine possible causes.
<b>0x3100a717</b>	ADMIN	system.moab	FATAL	MWM_CANNOT_FORK_INTO_BACKGROUND	Cannot fork the process into the background, errno: %s (%s).	The fork() system call failed. Moab must do this to daemonize unless run with the '-d' flag. This is usually due to low system resources.
<b>0x3100a71c</b>	ADMIN	system.moab	FATAL	MWM_CANNOT_CHANGE_FILE_OWNERSHIP_FATAL	Cannot change ownership of %s file to uid:%s gid:%s errno: %s (%s).	The fchown() system call failed. Use the errno and associated message to determine possible causes.

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
<b>0x3100a745</b>	ADMIN	system.moab	FATAL	MWM_CANNOT_GET_SERVER_HOSTNAME	Cannot determine hostname and attribute '%s' of parameter %s is not specified.	Moab failed to obtain system host name or ip address information from the operating system.
<b>0x3100a746</b>	ADMIN	system.moab	FATAL	MWM_HA_MOAB_NOT_STARTED_ON_CORRECT_HOSTS	The server must be started on host '%s' or on alternate '%s' (currently on '%s').	Moab must be started on either the primary or alternate host for high availability.
<b>0x3100a747</b>	ADMIN	system.moab	FATAL	MWM_MOAB_NOT_STARTED_ON_CORRECT_HOST	The server must be started on host '%s' (currently on '%s').	Moab must be started on specified host as identified by the SCHEDCFG parameter.
<b>0x3100a749</b>	ADMIN	system.moab	FATAL	MWM_MOAB_ALREADY_RUNNING	Moab is already running. Cannot open user interface socket on port %s.	Cannot open user interface socket, which is most likely caused by Moab already running.
<b>0x3100a74b</b>	ADMIN	system.moab	FATAL	MWM_CANNOT_LOCATE_FULL_PATH	Cannot locate the full path for '%s'.	Check the path to make sure the Moab executable is in it. Restart manually to work around this problem

Code	Escalation Level	Topic	Severity	Event Name	Message Template	Comment
						temporarily.
<b>0x3100a74c</b>	ADMIN	system.moab	FATAL	MWM_CANNOT_RESTART_SCHEDULER	Exec failed when attempting to restart the scheduler '%s' rc: %s.	Check permissions on this executable to correct and restart manually to work around.
<b>0x3100a750</b>	ADMIN	system.moab	FATAL	MWM_CANNOT_CONNECT_TO_DB_WITH_STRICT_CONFIG_CHECK_ON	StrictConfigCheck ON and cannot connect to DB-- please check DB engine and configuration (%s).	Moab was unable to connect to the database and with strict configuration on Moab must exit.
<b>0x3100a751</b>	ADMIN	system.moab	FATAL	MWM_USER_NOT_AUTHORIZED_TO_RUN_THIS_PROGRAM	The user '%s' (UID: %s) is not authorized to run this program.	The user has insufficient privileges to run the program.
<b>0x3100a752</b>	ADMIN	system.moab	FATAL	MWM_PROBLEMS_WITH_KEY_FILE	Problems with key file.	Key file does not exist or ownership of key file is invalid.
<b>0x3100aa0d</b>	ADMIN	system.moab	FATAL	MWM_STRICT_CONFIG_EXIT	Exiting because of strict configuration check.	Moab is configured to exit if there are any errors in configuration files or file/directory

Cod e	Es cal ation Le vel	Topi c	S e v e r i t y	Event Name	Message Template	Comment
						layout. One of these errors has occurred.
<b>0x3100c002</b>	INTE RN AL	system.mo ab	FATAL	MWM_TESTING_FATAL	Testing with single argument: %s.	Internal error for testing diagnostics.
<b>0x3100e74f</b>	INTE RN AL	system.mo ab	FATAL	MWM_CORRUPT_CHECKPOINT_FILE	Unable to read the checkpoint file.	Please contact Adaptive Computing for assistance.

## Appendix D: Adjusting Default Limits

Moab is distributed in a configuration capable of supporting multiple architectures and systems ranging from a few processors to several thousand processors. However, in spite of its flexibility, for performance reasons, it still contains a number of default object limits parameters and static structures defined in header files. These limits constrain such things as the maximum number of jobs, reservations, and nodes that Moab can handle and are set to values that provide a reasonable compromise between capability and memory consumption for most sites. However, many site administrators want to increase some of these settings to extend functionality, or decrease them to save consumed memory. The most common parameters are listed in the table below. Parameters listed in the Moab configuration file (`moab.cfg`) can be modified by restarting Moab. To change parameters listed in `moab.h`, please contact technical support.

Moab currently possesses hooks to allow sites to create local algorithms for handling site specific needs in several areas. The `contrib` directory contains a number of sample local algorithms for various purposes. The `MLocal.c` module incorporates the algorithm of interest into the main code. The following scheduling areas are currently handled via the `MLocal.c` hooks:

- Local Job Attributes
- Local Node Allocation Policies
- Local Job Priorities
- Local Fairness Policies

### CLIENTMAXCONNECTIONS

<b>Location</b>	<code>moab.cfg</code> (dynamic parameter)
<b>Default</b>	128
<b>Max tested</b>	---
<b>Description</b>	Maximum number of connections that can simultaneously connect to Moab. See <a href="#">CLIENTMAXCONNECTIONS</a> for additional information on this parameter.

### JOBMAXNODECOUNT

<b>Location</b>	<code>moab.cfg</code> (dynamic parameter)
-----------------	---

JOBMAXNODECOUNT	
<b>Default</b>	1024
<b>Max tested</b>	8192
<b>Description</b>	Maximum number of compute nodes that can be allocated to a job. After changing this parameter, Moab must be restarted for changes to take effect. The value cannot exceed that of the <a href="#">MAXNODE</a> parameter (specified in <code>moab.cfg</code> ). If you specify a value higher than the limit set for the <code>MAXNODE</code> parameter, the value will match <code>MAXNODE</code> . <code>JOBMAXNODECOUNT</code> can also be specified within configure using <code>--with-maxjobsize=&lt;NODECOUNT&gt;</code> . See <a href="#">JOBMAXNODECOUNT</a> for additional information on this parameter.

JOBMAXTASKCOUNT	
<b>Location</b>	<code>moab.cfg</code> (dynamic parameter)
<b>Default</b>	32768
<b>Max tested</b>	250000
<b>Description</b>	Total number of tasks allowed per job. See <a href="#">JOBMAXTASKCOUNT</a> for additional information on this parameter.

MAXGRES	
<b>Location</b>	<code>moab.cfg</code> (dynamic parameter)
<b>Default</b>	512
<b>Max tested</b>	---
<b>Description</b>	Total number of distinct generic resources that can be managed. See <a href="#">MAXGRES</a> for additional information on this parameter.

MAXJOB	
<b>Location</b>	<code>moab.cfg</code> (dynamic parameter)

MAXJOB	
<b>Default</b>	51200
<b>Max tested</b>	500,000
<b>Description</b>	Maximum number of jobs that can be evaluated simultaneously. (Can also be specified within <b>configure</b> using <b>--with-maxjobs=&lt;JOBCOUNT&gt;</b> .) See <a href="#">MAXJOB</a> for additional information on this parameter.

MAXNODE	
<b>Location</b>	moab.cfg (dynamic parameter)
<b>Default</b>	5120
<b>Max tested</b>	160000
<b>Description</b>	Maximum number of compute nodes supported. See <a href="#">MAXNODE</a> for additional information on this parameter.

MAXRSVPERNODE	
<b>Location</b>	moab.cfg (dynamic parameter)
<b>Default</b>	64
<b>Max tested</b>	1024
<b>Description</b>	Maximum number of reservations a node can simultaneously support. See <a href="#">MAXRSVPERNODE</a> for additional information on this parameter.

MMAX_ATTR	
<b>Location</b>	moab.h
<b>Default</b>	128
<b>Max tested</b>	512

### MMAX\_ATTR

<b>Description</b>	Total number of distinct node attributes (PBS node attributes/LL node features) that can be tracked.
--------------------	--

### MMAX\_CLASS

<b>Location</b>	moab.h
<b>Default</b>	24
<b>Max tested</b>	64
<b>Description</b>	Total number of distinct job classes/queues available.

### MMAX\_FSDEPTH

<b>Location</b>	moab.h
<b>Default</b>	24
<b>Max tested</b>	32
<b>Description</b>	Number of active fairshare windows.

### MMAX\_PAR

<b>Location</b>	moab.h
<b>Default</b>	32
<b>Max tested</b>	32
<b>Description</b>	Maximum number of partitions supported.

### MMAX\_QOS

<b>Location</b>	moab.h
-----------------	--------



MMAX_QOS	
<b>Default</b>	128
<b>Max tested</b>	128
<b>Description</b>	Total number of distinct QoS objects available to jobs.

MMAX_RACK	
<b>Location</b>	moab.h
<b>Default</b>	200
<b>Max tested</b>	200
<b>Description</b>	Total number of distinct rack objects available within cluster.

MMAX_RANGE	
<b>Location</b>	moab.h
<b>Default</b>	2048
<b>Max tested</b>	2048
<b>Description</b>	<p>Total number of distinct timeframes evaluated.</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> This is proportional to the size of the cluster and the number of simultaneously active jobs in the cluster. (Can be specified within <code>./configure</code> using <code>--with-maxrange=&lt;RANGECOUNT&gt;</code>.) Increasing this value will not increase the size of total memory consumed by Moab but may result in minor slowdowns in the evaluation and optimization of reservations.</p> </div>

MMAX_REQ_PER_JOB	
<b>Location</b>	moab.h
<b>Default</b>	5

MMAX_REQ_PER_JOB	
<b>Max tested</b>	64
<b>Description</b>	Total number of unique requirement structures a job can have. Limits the number of <code>-w</code> clauses in the <code>mshow -a</code> command. It also limits the number of <code>-l nodes=X+Y+Z</code> a normal HPC job can have.

---

### Related Topics

- [Appendix I: Considerations for Large Clusters](#)

## Appendix E: Security

Moab provides role and host based authorization, encryption, and DES, HMAC, and MD5 based authentication. The following sections describe these features in more detail and tell how to control access to sensitive configuration information.

In this section:

- [E.1 Authentication \(Interface Security\) - page 1359](#)
- [E.2 Authorization - page 1364](#)
- [E.3 Host Security for Compute Resources - page 1367](#)
- [E.4 Securing Sensitive Configuration Information - page 1368](#)

### E.1 Authentication (Interface Security)

Moab supports password-challenge, DES, HMAC, and MD5 based authentication. Authentication protocols can be specified on a per interface basis allowing independent realms of trust with per realm secret keys and even per realm authentication protocols.

In this topic:

- [E.1.1 Mauth Authentication - page 1359](#)
- [E.1.2 Munge Authentication - page 1362](#)
- [E.1.3 Server Response Control - page 1363](#)
- [E.1.4 Checksum Algorithm for Client Authentication - page 1363](#)
- [E.1.5 Interface Development Notes - page 1363](#)

#### E.1.1 Mauth Authentication

Mauth is a tool provided with Moab that provides client authentication services. With mauth enabled, each client request is packaged with the client ID, a timestamp, and an encrypted key of the entire request generated using the shared secret key.

This tool is enabled by providing a secret key. A random key is selected when the Moab `./configure` script is run and can be regenerated at any time by rerunning `./configure` and rebuilding Moab. If desired, this random key can be overridden by specifying a new key in the protected `.moab.key` file as in the example below:



Moab must be shut down before setting a new secret key. Use the `service moab stop` or `mschedctl -k` commands to shut down Moab.

```
> vi /opt/moab/etc/.moab.key
(insert key)
> cat /opt/moab/etc/.moab.key
XXXXXXXXXX
# secure file by setting owner read-only permissions
> chmod 400 /opt/moab/etc/.moab.key
# verify file is owned by root and permissions allow only root to read file
> ls -l /opt/moab/etc/.moab.key
-r----- 1 root root 15 2007-04-05 03:47 /opt/moab/etc/.moab.key
```

**i** Be aware of the following:

- All directories in the path containing `.moab.key` must be owned by the root or primary Moab user. It must not be writable by 'other' in its permissions.
- The `.moab.key` file will need to be on each host that is authorized to run Moab client commands.
- The `.moab.key` file must reside in the same directory as the `moab.cfg` file.
- By default:
  - The `.moab.key` file will be owned by the user root and its contents will be read by the mauth tool, which provides client authorization services. If desired, the ownership of this file can be changed so long as this file is readable by the Moab server and the mauth tool. This can be accomplished if the Moab [primary administrator](#), the owner of mauth, and the owner of `.moab.key` are the same.
  - It is up to the individual cluster administrators to determine whether to use the `.moab.key` file. For sites with source code, the use of `.moab.key` can be mandated by using `./configure --with-keyfile`.
  - mauth is located in the `install bin` directory. If an alternate name or alternate file location is desired, this can be specified by setting the AUTHCMD attribute of the [CLIENTCFG](#) parameter within the `moab.cfg` file as in the following example.

```
CLIENTCFG AUTHCMD=/opt/sbin/mauth
```

### E.1.1.A Configuring Peer-Specific Secret Keys

Peer-specific secret keys can be specified using the [CLIENTCFG](#) parameter. This key information must be kept secret and consequently can only be specified in the `moab-private.cfg` file. With regard to security, there are two key attributes that can be set.

(Other resource managers or clients such as Moab Accounting Manager or a Wiki interface can also use the attributes to configure their authentication algorithms. The default, unless otherwise stated, is always DES. These attributes are listed in the table below:

AUTH	
<b>Format</b>	One of ADMIN1, ADMIN2, or ADMIN3
<b>Default</b>	---
<b>Description</b>	The level of control/information available to requests coming from this source/peer.
<b>Example</b>	<pre>CLIENTCFG[RM:clusterB] AUTH=admin1 KEY=14335443</pre>

AUTHTYPE	
<b>Format</b>	One of DES, HMAC, HMAC64, HMACSHA2, or MD5.
<b>Default</b>	DES
<b>Description</b>	The encryption algorithm to use when generating the message checksum.
<b>Example</b>	<pre>CLIENTCFG[AM:mam] AUTHTYPE=HMAC64</pre>

HOST	
<b>Format</b>	<STRING>
<b>Default</b>	---
<b>Description</b>	The hostname of the remote peer. Peer requests coming from this host will be authenticated using the specified mechanism. This parameter is optional.
<b>Example</b>	<pre>CLIENTCFG[RM:clusterA] HOST=orx.pb13.com KEY=banana6</pre>

KEY	
<b>Format</b>	<STRING>
<b>Default</b>	---
<b>Description</b>	The shared secret key to be used to generate the message checksum.
<b>Example</b>	<code>CLIENTCFG[RM:clusterA] KEY=banana6</code>

The `CLIENTCFG` parameter takes a string index indicating which peer service will use the specified attributes. In most cases, this string is simply the defined name of the peer service. However, for the special cases of resource and accounting managers, the peer name should be prepended with the prefix `RM:` or `AM:` respectively, as in `CLIENTCFG[AM:mam]` or `CLIENTCFG[RM:devcluster]`.

**i** The first character of any secret key can be viewed by trusted administrators using specific diagnostic commands to analyze Moab interfaces. If needed, increase the length of the secret keys to maintain the desired security level.

## E.1.2 Munge Authentication

Moab also integrates with `MUNGE`, an open source authentication service created by Lawrence Livermore National Laboratory (<https://dun.github.io/munge/>). `MUNGE` works with Moab to authenticate user credentials being passed between the Moab client and the Moab server or from Moab server to Moab server.

To set up `MUNGE` in a cluster or grid, download and install `MUNGE` on every node in the cluster or grid by following the installation steps found at <https://dun.github.io/munge/>. The `MUNGE` secret key must reside on each node in the cluster or grid. Before starting the Moab daemon, the `MUNGE` daemon must be running on all nodes.

To enable Moab to use `MUNGE` for authentication purposes, specify the `MUNGE` executable path in the `moab.cfg` file using `CLIENTCFG` and `AUTHCMD` as in the following example. The `MUNGE` executable path must reside in each client's `moab.cfg` file as well.

```
CLIENTCFG AUTHCMD=/usr/bin/munge
```

**i** Moab requires that the `MUNGE` and `UNMUNGE` executable names be 'munge' and 'unmunge' respectively. It also assumes that the `UNMUNGE` executable resides in the same directory as the `MUNGE` executable.

### E.1.2.A Configuring Munge Command Options

Moab also integrates with MUNGE command line options. For example, to set up Moab to use a specific socket that was created when the MUNGE daemon was started, use `CLIENTCFG` and `AUTHCMDOPTIONS` to specify the newly created socket. The `AUTHCMDOPTIONS` attribute, like `AUTHCMD`, must also reside in the client's `moab.cfg` file.

```
CLIENTCFG    AUTHCMD=/usr/bin/munge
CLIENTCFG    AUTHCMDOPTIONS="-S /var/run/munge/munge.socket.2"
```

### E.1.3 Server Response Control

If a request is received that is corrupt or cannot be authenticated, Moab will report some limited information to the client indicating the source of the failure, such as 'bad key,' 'malformed header,' and so forth. In the case of highly secure environments, or to minimize the impact of sniffing or denial of service attacks, Moab can be configured to simply drop invalid requests. This is accomplished by adding the `DROPBADREQUEST` attribute to the `CLIENTCFG` parameter in the `moab-private.cfg` file as in the following example:

```
CLIENTCFG [DEFAULT] DROPBADREQUEST=TRUE
```

### E.1.4 Checksum Algorithm for Client Authentication

The `SERVERCSALGO` parameter lets you choose the algorithm used for message digests and message authentication codes:

- HMAC64: the default (SHA-1)
- HMACSHA2: more secure (SHA-512)



If you are using Moab Web Services, then you must set the MWS configuration parameter `moab.messageDigestAlgorithm` to match the value of `SERVERCSALGO`. See 'moab.messageDigestAlgorithm' in the *Moab Web Services Reference Guide* for more information.

### E.1.5 Interface Development Notes

Sample checksum generation algorithm code can be found in the [Socket Protocol Description](#) document.

## E.2 Authorization

### E.2.1 Role Based Authorization Security Configuration

Moab provides access control mechanisms to limit how the scheduling environment is managed. The primary means of accomplishing this is through limiting the users and hosts that are trusted and have access to privileged commands and data.

With regard to users, Moab breaks access into three distinct levels.

#### Level 1 Moab Admin (Administrator Access)

Level 1 Moab administrators have global access to information and unlimited control over scheduling operations. By default, they are allowed to control scheduler configuration, policies, jobs, reservations, and all scheduling functions. They are also granted access to all available statistics and state information. Level 1 administrators are specified using the [ADMINCFG\[1\]](#) parameter.

#### Level 2 Moab Admin (Operator Access)

Level 2 Moab administrators are specified using the [ADMINCFG\[2\]](#) parameter. By default, the users listed under this parameter are allowed to change all job attributes and are granted access to all informational Moab commands.

#### Level 3 Moab Admin (Help Desk Access)

Level 3 administrators are specified via the [ADMINCFG\[3\]](#) parameter. By default, they are allowed access to all informational Moab commands. They cannot change scheduler or job attributes.

#### E.2.1.A Configuring Role Based Access

Moab allows site specific tuning of exactly which functions are available to each administrator level. Moab also provides two additional administrator levels ([ADMINCFG\[4\]](#) and [ADMINCFG\[5\]](#)) that can be used for site specific needs.

**i** [ADMINCFG\[5\]](#) is different from other administrator levels because, when set, all commands are authorized at that level by default. Furthermore, if a service is set at level 5, all other services are disallowed.

To configure Moab role based access, use the [ADMINCFG](#) parameter:

```
ADMINCFG[1]    USERS=root, john SERVICES=ALL NAME=admin
```



```

ADMINCFG[3]   USERS=joe,mary   SERVICES=mdiag,mrsvctl,mcredctl NAME=power
ADMINCFG[5]   USERS=joy,blake  SERVICES=NONE   NAME=users
...

```

**i** A NONE in services will still allow users to run [showq](#) and [checkjob](#) on their own jobs.

To determine the role of system users and what commands they can run, use the [mcredctl -q role user:<USERID>](#) command.

Using the SERVICES attribute of the ADMINCFG parameter, access to an arbitrary selection of services can be enabled on a per administrator-level basis. Possible services include the following:

Service	Description
<a href="#">changeparam</a>	Change any scheduling policy or parameter. (This command is deprecated. Use <a href="#">mschedctl -m</a> instead).
<a href="#">checkjob</a>	View detailed information for any job.
<a href="#">checknode</a>	View detailed information for any node.
<a href="#">mbal</a>	Perform real-time load-balancing of interactive commands.
<a href="#">mcredctl</a>	View and modify credential attributes.
<a href="#">mdiag</a>	Provide diagnostic reports for resources, workload, and scheduling.
<a href="#">mjobctl</a>	Modify, control, and view jobs.
<a href="#">mnodectl</a>	Modify, control, and view nodes.
<a href="#">mrmctl</a>	Modify, control, and view resource managers.
<a href="#">mrsvctl</a>	Modify, control, and view reservations.
<a href="#">mschedctl</a>	Modify, control, and view scheduler behavior.
<a href="#">mshow</a>	View existing configuration and predicted resource availability.
<a href="#">showstats</a>	View all scheduler and credential statistics.

Service	Description
<b>releaseres</b>	Release all reservations. (This command is deprecated. Use <code>mrsvctl -r</code> instead).
<b>runjob</b>	Immediately execute any job (see <code>mjobctl -x</code> ).
<b>setqos</b>	Set QoS on any job. (This command is deprecated. Use <code>mjobctl -m</code> instead).
<b>setres</b>	Create any reservation. (This command is deprecated. Use <code>mrsvctl -c</code> instead).
<b>setspri</b>	Set system priority on any job. (This command is deprecated. Use <code>mjobctl -p</code> instead).
<b>showconfig</b>	Show all scheduler configuration parameters. (This command is deprecated. Use <code>mschedctl -l</code> instead).
<b>showres</b>	Show detailed information for any reservation.
<b>showstate</b>	Show detailed information for all jobs, including their locations, and display job error messages, if any.

### E.2.1.B Account and Class/Queue Admins

While the `ADMINCFG` parameter allows organizations to provide controlled access to scheduling objects, it does not allow for distribution along organizational boundaries. For example, a site may set up a level 3 administrator to be able to view statistics, diagnose jobs, and modify job priorities; it does not provide a way to differentiate one type of job from another. If a site administrator wanted to allow control based on the queue or account associated with a job, they would best accomplish this using the credential `MANAGERS` attribute.

A credential manager allows a user to be trusted to administer workload and policies for an associated subgroup of jobs. For example, in the configuration below, a number of queue and account managers are configured:

```
CLASSCFG[orion] MANAGERS=johns
CLASSCFG[xray]  MANAGERS=steve2
CLASSCFG[gamma] MANAGERS=steve2, jpw
ACCOUNTCFG[bio] MANAGERS=charles
```

By default, the specified managers can do anything to a job that the actual job owner could do. By default, this would include the ability to view cumulative and per job statistics, see

job details, modify job priorities and holds, cancel and preempt jobs, and otherwise adjust policies and constraints within the associated credential.

## E.3 Host Security for Compute Resources

Host level security can vary widely from one site to another with everything from pure on-your-honor based clusters to complete encrypted VLAN based network security and government approved per job scrubbing procedures being used. The following documentation describes some best practices in use throughout the industry.

### E.3.1 Minimal Host Security Enforcement

For minimal host security, no additional configuration is required.

### E.3.2 Medium Host Security Enforcement

- Login Access
  - PAM — Enable/disable access by modifying `/etc/security/access.conf`.
- Processes
  - Kill all processes associated with job user (dedicated).
  - Kill all processes associated with job session (dedicated/shared). Use `ps -ju <USER>` or `ps -js <SESSID>`.
- IPC (Inter-Process Communication)
  - Remove shared memory, semaphores, and message queues (use `ipcs/ipcrm`).
  - Remove named pipes.
- Network/Global File System Access
  - Explicitly unmount user home and global file systems.
- Local Temporary File Systems
  - Where possible, mount local file systems read-only.
  - Clear `/tmp`, `/scratch` and other publicly available local file systems.
  - Remove user files with `shred`; `shred` is a Linux command that first overwrites files completely before removing them, preventing remnant data from surviving on the hard drive.

### E.3.3 Strict Host Security Enforcement

- VLAN creation
- Host rebuild
  - U.S. Dept. of Energy Disk/File Sanitization ([Clearing, Sanitizing, and Destroying Disks](#))
  - U.S. Dept. of Defense Scrubbing [Software](#) (DOD-5520)

## E.4 Securing Sensitive Configuration Information

The `moab.cfg` file may include sensitive configuration information, such as user or group fairshare targets that determine job priority and scheduling for individual users or groups. Sensitive configuration information can be moved to a separate file in an access-controlled directory and included in `moab.cfg` using an `#INCLUDE` directive. For example, the following commands create a directory that requires root permissions to read or execute, and a `.cfg` file that can be used for sensitive configuration information:

```
# mkdir -m 500 /opt/moab/etc/secure
# echo "ARRAYJOBPARLOCK TRUE" > /opt/moab/etc/secure/moab.secure.cfg
```

Adding the following line to `moab.cfg` will cause Moab to use the contents of the protected `.cfg` file:

```
#INCLUDE secure/moab.secure.cfg
```

## Appendix F: Initial Moab Testing

Moab has been designed with a number of key features that allow testing to occur in a *no risk* environment. These features allow you to safely run Moab in test mode even with another scheduler running whether it be an earlier version of Moab or another scheduler altogether. In test mode, Moab collects real-time job and node information from your resource managers and acts as if it were scheduling live. However, its ability to actually affect jobs (that is, start, modify, cancel, charge, and so forth) is disabled.

### F.0.1 Scheduler Modes

This topic describes the test modes Moab offers to provide a minimal configuration for verifying such things as proper configuration and operation.

Central to Moab testing is the `MODE` attribute of the `SCHEDCFG` parameter. This parameter attribute allows administrators to determine how Moab will run. The possible values for `MODE` are `NORMAL`, `MONITOR`, `INTERACTIVE`, `TEST`, `SINGLESTEP`, and `SLAVE`. For example, to request monitor mode operation, include the following in the `moab.cfg` file:

```
SCHEDCFG MODE=MONITOR
```

### F.0.2 Normal Mode

If initial evaluation is complete or not required, you can place the scheduler directly into *production* by setting the `MODE` attribute of the `SCHEDCFG` parameter to `NORMAL` and (re)starting the scheduler.

### F.0.3 Monitor Mode (or Test Mode)

Monitor mode allows evaluation of new Moab releases, configurations, and policies in a risk-free manner. In monitor mode, the scheduler connects to the resource manager(s) and obtains live resource and workload information. Using the policies specified in the `moab.cfg` file, the monitor-mode Moab behaves identical to a live or normal-mode Moab except the ability to start, cancel, or modify jobs is disabled. In addition, allocation management does not occur in monitor mode. This allows safe diagnosis of the scheduling state and behavior using the various diagnostic client commands. Further, the log output can also be evaluated to see if any unexpected situations have arisen. At any point, the scheduler can be dynamically changed from monitor to normal mode to begin *live* scheduling.

To set up Moab in monitor mode, do the following:

```
> vi moab.cfg
    (change the MODE attribute of the SCHEDCFG parameter from NORMAL to MONITOR)
> moab
```

Remember that Moab running in monitor mode will not interfere with your production scheduler.

### F.0.3.A Running Multiple Moab Instances Simultaneously

If running multiple instances of Moab, whether in simulation, normal, or monitor mode, make certain that each instance resides in a different home directory to prevent conflicts with configuration, log, and statistics files. Before starting each additional Moab, set the `MOABHOMEDIR` environment variable in the execution environment to point to the local home directory. Also, each instance of Moab should run using a different [port](#) to avoid conflicts.

**i** If running multiple versions of Moab, not just different Moab modes or configurations, set the `$PATH` variable to point to the appropriate Moab binaries.

To *point* Moab client commands (such as [showq](#)) to the proper Moab server, use the appropriate command line [arguments](#) or set the environment variable `MOABHOMEDIR` in the client execution environment as in the following example:

```
# point moab clients/server to new configuration
> export MOABHOMEDIR=/opt/moab-monitor
# set path to new binaries (optional)
> export PATH=/opt/moab-monitor/bin:/opt/moab-monitor/sbin:$PATH
# start Moab server
> moab
# query Moab server
> showq
```

**i** `moabd` is a safe and recommended method of starting Moab if things are not installed in their default locations.

### F.0.4 Interactive Mode

Interactive mode allows for evaluation of new versions and configurations in a manner different from monitor mode. Instead of disabling all resource and job control functions, Moab sends the desired change request to the screen and asks for permission to complete it. For example, before starting a job, Moab might post something like the following to the screen:

```
Command: start job 1139.ncsa.edu on node list test013,test017,test018,test021
Accept: (y/n) [default: n]?
```

The administrator must specifically accept each command request after verifying that it correctly meets desired site policies. Moab will then execute the specified command. This mode is useful in validating scheduler behavior and can be used until configuration is appropriately tuned and all parties are comfortable with the scheduler's performance. In most cases, sites will want to set the scheduling mode to normal after verifying correct behavior.

## Appendix G: Integrating Other Resources with Moab

Moab can interface with most popular resource managers, many cluster services, and numerous general protocols. The following links provide additional information.

In this section:

- [G.1 Compute Resource Managers - page 1372](#)
- [G.2 Hardware Integration - page 1376](#)
- [G.3 Torque/PBS Integration Guide - RM Access Control - page 1380](#)
- [G.4 Torque/PBS Config - Default Queue Settings - page 1381](#)
- [G.5 Provisioning Resource Managers - page 1382](#)

### G.1 Compute Resource Managers

#### G.1.1 Moab-Torque Integration Guide

In this topic:

- [G.1.1.A Integration Steps - page 1372](#)
- [G.1.1.B Torque/Moab Considerations - page 1375](#)
- [G.1.1.C Troubleshooting - page 1376](#)

##### G.1.1.A Integration Steps

#### Install Torque

- Installing Torque Resource Manager

 Keep track of the PBS target directory, `$PBSTARGDIR`



## Install Moab

- Untar the Moab distribution file.
- Change the directory to the `moab-10.0.1` directory.
- Run `./configure`.
- Specify the PBS target directory (`$PBSTARGDIR` from step 2.1) when queried by `./configure`.

Moab interfaces to PBS by utilizing a few PBS libraries and include files. If you have a non-standard PBS installation, you may need to modify `Makefile` and change `PBSIP` and `PBSLP` values and references as necessary for your local site configuration.

The `./configure` script automatically sets up Moab so that the user running `configure` will become the default Primary Moab Administrator (`$MOABADMIN`). This can be changed by modifying the `ADMINCFG[1] USERS=<USERNAME>` line in the Moab configuration file (`moab.cfg`). The primary administrator is the first user listed in the `USERS` attribute and is the ID under which the Moab daemon runs.

Some Tru64 and IRIX systems have a local `libnet` library that conflicts with PBS's `libnet` library. To resolve this, try setting `PBSLIB` to `'${PBSLIBDIR}/libnet.a -lpbs'` in the Moab `Makefile`.

Moab is 64-bit compatible. If PBS/Torque is running in 64-bit mode, Moab likewise needs to be built in this manner to use the PBS scheduling API (i.e., for IRIX compilers, add `-64` to `OSCCFLAGS` and `OSLDLFLAGS` variables in the `Makefile`).

When starting both Torque and Moab it is best to have a small delay between starting the servers. In general (and especially for very fast or very large systems) this is recommended startup procedure:

- Start Torque.
- Start Moab with scheduling paused (`moab -P`) to give it a chance to load everything in the checkpoint file and to sync with Torque.
- Unpause Moab with `mschedctl -r`.

## Configure Torque

### General Configuration for All Versions of Torque

- Make `$MOABADMIN` a PBS admin.
  - By default, Moab only communicates with the `pbs_server` daemons and the `$MOABADMIN` should be authorized to talk to this daemon (see [suggestions](#) for more information).

- (OPTIONAL) Set default PBS queue, nodecount, and walltime attributes. (See [suggestions](#) for more information.)
- (OPTIONAL - Torque Only) Configure Torque to report completed job information by setting the `qmgrkeep_completed` parameter:>

```
> qmgr -c 'set server keep_completed = 300'
```

**i** PBS nodes can be configured as *time shared* or *space shared* according to local needs. In almost all cases, space shared nodes provide the desired behavior.

**i** PBS/Torque supports the concept of virtual nodes. Using this feature, Moab can individually schedule processors on SMP nodes. The online Torque documentation describes how to set up the `$PBS_HOME/server_priv/nodes` file to enable this capability. (For example, `<NODENAME>np=<VIRTUAL NODE COUNT>`)

### Version-Specific Configuration for Torque

Do not start the `pbs_sched` daemon. This is the default scheduler for Torque; Moab provides this service.

**i** Moab uses PBS's scheduling port to obtain real-time event information from PBS regarding job and node transitions. Leaving the default `qmgr` setting of `set server scheduling=True` allows Moab to receive and process this real-time information.

### Configure Moab

By default, Moab automatically interfaces with Torque/PBS when it is installed. Consequently, in most cases, the following steps are not required:

- Specify PBS as the primary resource manager by setting `RMCFG[base]TYPE=PBS` in the Moab configuration file (`moab.cfg`).

If a non-standard PBS installation/configuration is being used, additional Moab parameters may be required to enable the Moab/PBS interface as in the line `RMCFG[base]HOST=$PBSSERVERHOSTPORT=$PBSSERVERPORT`. See the [Resource Manager Overview](#) for more information.

**i** Moab's user interface port is set using the `SCHEDCFG` parameter and is used for user-scheduler communication. This port must be different from the PBS scheduler port used for resource manager-scheduler communication.

### G.1.1.B Torque/Moab Considerations

The default meaning of a node for Torque and Moab are not the same. By default, a node is a host in Torque. The node can have one or more execution slots (procs) allocated to it in the `TORQUE_HOME/server_priv/nodes` file. However, the number of nodes recognized by Torque is equivalent to the number of node entries in the `TORQUE_HOME/server_priv/nodes` file. A node specification from `qsub` such as `-1 nodes=2:ppn=2` will direct Torque to allocate to execution slots on two separate nodes.

Moab is more liberal in its interpretations of a node. To Moab, the `qsub` request above would be interpreted to mean allocate four tasks with at least two tasks on a node. Where Torque would require two nodes for the request, Moab will place all four tasks on the name node (host) if four execution slots are available.

If a cluster has four nodes with eight processors each, Torque still sees only four nodes. Moab sees 32 nodes. However, if a user made a `qsub` request with `-1 nodes=10`, Torque would reject the request because there are only four nodes available. To enable Torque to accommodate Moab's more liberal node interpretation, the server parameter `available_resources.nodect` can be set as a server parameter in Torque. The value of `available_resources.nodect` should equal at least the number of execution slots in the cluster.

For our example, cluster `available_resources.nodect` should be 32. With this parameter set, the user can now make a request such as `-1 nodes=8:ppn=2`. In this example, the user is still limited to a maximum node request of 32.

With `available_resources.nodect` set in Torque, Moab can be directed to honor the default Torque behavior by setting `JOBNODEMATCHPOLICY` to `EXACTNODE`.

### PBS Features Not Supported by Moab

Moab supports basic scheduling of all PBS node specifications.

### Moab Features Not Supported by PBS

PBS does not support the concept of a job QoS or other extended scheduling features by default.

### G.1.1.C Troubleshooting

On TRU64 systems, the PBS `libpbs` library does not properly export a number of symbols required by Moab. This can be worked around by modifying the Moab `Makefile` to link the PBS `rm.o` object file directly into Moab.

## G.2 Hardware Integration

### G.2.1 Moab-NUMA-Support Integration Guide



This topic is for NUMA-support systems on large-scale SLES systems using SGI Altix and UV hardware only and requires Torque 3.0 or later.

Scheduling a shared-memory NUMA type system (not the same as a modern SMP-based individual compute node, which cannot share memory between compute nodes) requires some special configuration. Additionally, Moab can use [NODESETs](#) to guarantee feasibility of large memory jobs and to enforce node allocation based on the system's interconnect network topology.

In this topic:

[G.2.1.A Configuration - page 1376](#)

[G.2.1.B Job Submission - page 1380](#)

#### G.2.1.A Configuration

### To Integrate Moab and NUMA

1. Configure Moab to schedule large memory jobs. Because Moab creates a partition for each resource manager by default, you must configure the cluster controlled by the resource manager to be a shared-memory system to support jobs spanning multiple nodes/blades. To do so, use the [PARCFG](#) parameter:

```
RMCFG[sys-uv] TYPE=Torque
PARCFG[sys-uv] FLAGS=SharedMem
```

*Cluster `sys-uv` is now configured as a shared-memory system to Moab.*

2. Configure [NODESETs](#) as shown below:

```

NODESETISOPTIONAL FALSE
NODESETATTRIBUTE FEATURE
NODESETPOLICY ONEOF
NODESETPRIORITYTYPE FIRSTFIT

```

*The NODESET parameters tell Moab that performing node allocation using node sets is required, that the node set name is a feature name assigned to compute nodes, that a job must fit within the available nodes of one node set, and that Moab should use the first node set that contains sufficient available nodes to satisfy the job's request.*

3. To configure Moab to perform topology-aware node allocation using node sets, you must create a node set definition for each set of nodes that has the same number of maximum network 'hops' from any node to every other node within the node set. For an example, see the following sample scenario:

#### Use Case

The SGI UV 1000 has a two-socket blade with a physical organization of 16 blades within a blade chassis (SGI term is Intra-Rack Unit or IRU), two blade chassis (IRUs) within a rack, and up to four racks within a single UV system. The UV 1000 interconnect network has a topology that requires zero hops between the two sockets on the same physical blade, one hop between an even-odd blade pair (e.g., blades 0 and 1, 2 and 3, etc.), two hops between all even-numbered or all odd-numbered blades within an IRU, three hops maximum between all blades within an IRU, four hops maximum between all even-numbered blades or all odd-numbered blades within a UV system, and five hops maximum between all blades within a UV system.

- a. Define topology-aware node definitions to parallel the compute nodes reachable within a specific hop count. For the UV 1000, this means the sockets of each blade will belong to six separate node set definitions (i.e., one each for 0, 1, 2, 3, 4, and 5 hops).
- b. Define multiple node sets for different nodes reachable in a specific hop count based on the context of where they are in the network topology; that is, you must create a separate and distinct node set definition for each pair of blades reachable with one hop, for each IRU for its nodes reachable in three hops, etc.
- c. Moab node sets are usually defined as compute node features; that is, each node set defined to Moab should appear as a 'feature' name on one or more compute nodes. Which node set/feature names appear on each compute node depends on where the compute node is in the interconnect network topology.

Since the SGI UV operating system identifies each blade socket as a separate NUMA node, each NUMA node within a UV system is traditionally an individual compute node to Moab (although Torque has the ability to redefine a compute node definition

by grouping OS NUMA nodes, which some UV installations do to define a blade as a compute node).

For the sake of illustration, this example assumes each OS NUMA node, which is a UV blade socket, is also a compute node in Moab. This means each compute node (blade socket) will have six feature names assigned, where each feature name must reflect both the compute node's location in the network topology and the hop count the name represents. A feature name is constructed by using the same root name for a hop count and a number for the topology location at the hop-count level.

For example, the root feature name 'blade' represents the zero-hop count and the numbers '0', '1', etc, represent the consecutively numbered blades throughout the entire UV system, which yields feature names of 'blade0' for the first blade in the system, 'blade1' for the second blade, etc, to 'blade127' for the last blade in a fully populated 4-rack UV system. To illustrate further, the root feature name 'iru' represents the 3-hops count and the numbers '0' through '7' represent the eight IRUs within a full 4-rack UV system.

- d. For each compute node, configure the correct feature name for each of the hop counts possible and its location within the topology at the hop-count level (e.g., blade (0 hops), blade pair (1 hop), odd- or even-numbered nodes within an IRU (2 hops), IRU (3 hops), odd- or even-numbered nodes within the UV (4 hops), and UV system (5 hops)). The following example illustrates the feature names assigned to the compute nodes for an SGI UV 1000 system using the following root feature names:

- blade (0 hops)
- pair (1 hop)
- eiru (2 hops for even-numbered blades within an IRU)
- oiru (2 hops for odd-numbered blades within an IRU)
- iru (3 hops)
- esys (4 hops for even-numbered blades within a UV system)
- osys (4 hops for odd-numbered blades within a UV system)
- sys (5 hops)

Note that nodes 0 and 1 are not given any feature names. This is because the operating system instance for the UV system runs on the first blade and in order to not adversely affect OS performance, no jobs should run on the same compute resources as the operating system; therefore, these nodes have no node set feature names and therefore will never be chosen to run jobs. In addition, some of the first feature names at a specific hop count-level are omitted (such as pair0) since it makes no sense to define them when the first blade is a substantial part of the nodes making up a node set.

The node name of a UV system has the same name as the UV system's host name plus the NUMA node's relative socket number.

```

/var/spool/torque/server_priv/nodes:
sys-uv0
sys-uv1
sys-uv2  blade1          oiru0 iru0 osys sys
sys-uv3  blade1          oiru0 iru0 osys sys
sys-uv4  blade2    pair1    eiru0 iru0 esys sys
sys-uv5  blade2    pair1    eiru0 iru0 esys sys
sys-uv6  blade3    pair1    oiru0 iru0 osys sys
sys-uv7  blade3    pair1    oiru0 iru0 osys sys
sys-uv8  blade4    pair2    eiru0 iru0 esys sys
sys-uv9  blade4    pair2    eiru0 iru0 esys sys
sys-uv10 blade5    pair2    oiru0 iru0 osys sys
sys-uv11 blade5    pair2    oiru0 iru0 osys sys
sys-uv12 blade6    pair3    eiru0 iru0 esys sys
sys-uv13 blade6    pair3    eiru0 iru0 esys sys
sys-uv14 blade7    pair3    oiru0 iru0 osys sys
sys-uv15 blade7    pair3    oiru0 iru0 osys sys
sys-uv16 blade8    pair4    eiru0 iru0 esys sys
sys-uv17 blade8    pair4    eiru0 iru0 esys sys
sys-uv18 blade9    pair4    oiru0 iru0 osys sys
sys-uv19 blade9    pair4    oiru0 iru0 osys sys
sys-uv20 blade10   pair5    eiru0 iru0 esys sys
sys-uv21 blade10   pair5    eiru0 iru0 esys sys
sys-uv22 blade11   pair5    oiru0 iru0 osys sys
sys-uv23 blade11   pair5    oiru0 iru0 osys sys
sys-uv24 blade12   pair6    eiru0 iru0 esys sys
sys-uv25 blade12   pair6    eiru0 iru0 esys sys
sys-uv26 blade13   pair6    oiru0 iru0 osys sys
sys-uv27 blade13   pair6    oiru0 iru0 osys sys
sys-uv28 blade14   pair7    eiru0 iru0 esys sys
sys-uv29 blade14   pair7    eiru0 iru0 esys sys
sys-uv30 blade15   pair7    oiru0 iru0 osys sys
sys-uv31 blade15   pair7    oiru0 iru0 osys sys
sys-uv32 blade16   pair8    eiru1 iru1 esys sys
sys-uv33 blade16   pair8    eiru1 iru1 esys sys
sys-uv34 blade17   pair9    oiru1 iru1 osys sys
sys-uv35 blade17   pair9    oiru1 iru1 osys sys
...
sys-uv62 blade31   pair15   oiru1 iru1 osys sys
sys-uv63 blade31   pair15   oiru1 iru1 osys sys
sys-uv64 blade32   pair16   eiru2 iru2 esys sys
sys-uv65 blade32   pair16   eiru2 iru2 esys sys
...
sys-uv126 blade63   pair31   oiru3 iru3 osys sys
sys-uv127 blade63   pair31   oiru3 iru3 osys sys
sys-uv128 blade64   pair32   eiru4 iru4 esys sys
sys-uv129 blade64   pair32   eiru4 iru4 esys sys
...
sys-uv190 blade95   pair47   oiru5 iru5 osys sys
sys-uv191 blade95   pair47   oiru5 iru5 osys sys
sys-uv192 blade96   pair48   eiru6 iru6 esys sys
sys-uv193 blade96   pair48   eiru6 iru6 esys sys
...
sys-uv252 blade126 pair63   eiru7 iru7 esys sys
sys-uv253 blade126 pair63   eiru7 iru7 esys sys
sys-uv254 blade127 pair63   oiru7 iru7 osys sys
sys-uv255 blade127 pair63   oiru7 iru7 osys sys

```

4. Define the order that Moab should check node sets for available nodes. Since the `NODESETPRIORITYTYPE` has a value of `FIRSTFIT`, the node sets must be ordered from smallest to largest so Moab will always choose the node set with the fewest nodes

required to satisfy the job's request. This means listing all blades, blade pairs, even and odd IRUs, IRUs, even and odd system, and system, respectively.

```
moab.cfg:
NODESETLIST
blade1,blade2,blade3,...,blade127,pair1,pair2,pair3,...,pair63,eiru0,oiru0,eiru1,oiru1,
...,eiru7,oiru7,iru0,iru1,...,iru7,esys,osys,sys
```

5. Configure Moab to use the `PRIORITYNODEALLOCATIONPOLICY`. This allocation policy causes Moab to allocate enough nodes to fulfill a job's processor and memory requirement.

```
NODEALLOCATIONPOLICY PRIORITY
```

6. Set `NODEACCESSPOLICY` to `SINGLEJOB` to ensure that Moab will schedule large memory requests correctly and efficiently. This is necessary even when a job uses only the memory of a NUMA node.

```
NODEACCESSPOLICY SINGLEJOB
```

The policy `SINGLEJOB` tells Moab not to allow jobs to share NUMA resources (cores and memory), which for a shared-memory system is very important for fast job execution. For example, if Moab scheduled a job to use the cores of a NUMA node where memory is used by another job, both jobs would execute slowly (up to 10 times more slowly).

### G.2.1.B Job Submission

Jobs can request processors and memory using the `-l nodes=<number of cpus>` and `-l mem=<amount of memory>` syntaxes. You should not have `JOBNODEMATCHPOLICYEXACTNODE` configured on a NUMA system. You must use the `sharedmem` job flag on submission to force the job to run only on a `sharedmem` partition or cluster and to indicate that the job can span multiple nodes. For example:

```
qsub -l nodes=3,mem=640sgb,flags=sharedmem
```

## G.3 Torque/PBS Integration Guide - RM Access Control

### G.3.1 Server Configuration

Using the PBS `qmgr` command, add the Moab administrator as both a *manager* and *operator*:



```
> qmgr
Qmgr: set server managers += <MOABADMIN>@*.<YOURDOMAIN>
Qmgr: set server operators += <MOABADMIN>@*.<YOURDOMAIN>
Qmgr: quit
```

For example:

```
> qmgr
Qmgr: set server managers += staff@*.ucsd.edu
Qmgr: set operators += staff@*.ucsd.edu
Qmgr: quit
```

**i** If desired, the Moab administrator can be enabled as a manager and operator only on the host on which Moab is running by replacing "`*.<YOURDOMAIN>`" with "`<MOABSERVERHOSTNAME>`".

## G.3.2 (Optional) MOM Configuration

If direct Moab to `pbs_mom` communication is required, the `mom_priv/config` file on each compute node where `pbs_mom` runs should be set as in the following example:

```
$restricted *.<YOURDOMAIN>
$clienthost <MOABSERVERHOSTNAME>
```

**i** For security purposes, sites may want to run Moab under a non-root user ID. If so, and Moab-`pbs_mom` communication is required, the `mom_priv/config` files must be world-readable and contain the line `'$restricted *.<YOURDOMAIN>'`. (i.e., `'$restricted *.uconn.edu'`).

## G.4 Torque/PBS Config - Default Queue Settings

### G.4.1 Default Queue

To set the default queue (the queue used by jobs if a queue is not explicitly specified by the user), issue the following:

```
>> qmgr
Qmgr: set system default_queue = <QUEUENAME>
Qmgr: quit
```

## G.4.2 Queue Default Node and Walltime Attributes

To set a default of one node and 15 minutes of walltime for a particular queue, issue the following:

```
> qmgr
Qmgr: set queue <QUEUENAME> resources_default.nodect = 1
Qmgr: set queue <QUEUENAME> resources_default.walltime = 00:15:00
Qmgr: quit
```

## G.4.3 System-Wide Default Node and Walltime Attributes

To set system wide defaults, set the following:

```
> qmgr
Qmgr: set server resources_default.nodect = 1
Qmgr: set server resources_default.walltime = 00:15:00
Qmgr: quit
```

# G.5 Provisioning Resource Managers

## G.5.1 Validating an xCAT Installation for Use with Moab

In this topic:

- [G.5.1.A Introduction to Validating xCAT Configuration - page 1382](#)
- [G.5.1.B Verifying Node List - page 1383](#)
- [G.5.1.C Reporting Node Status - page 1383](#)
- [G.5.1.D Verifying Hardware Management Configuration - page 1383](#)
- [G.5.1.E Verifying Provisioning Images - page 1384](#)

### G.5.1.A Introduction to Validating xCAT Configuration

This document describes a series of steps to validate xCAT configuration prior to configuring Moab to manage hardware via xCAT. It is assumed the reader is familiar with xCAT and the xCAT configuration on the target site. This document does not provide xCAT configuration documentation or troubleshooting information; refer to the [xCAT documentation](#) for such information.

### G.5.1.B Verifying Node List

Verify that all nodes that Moab will manage are known to xCAT with the xCAT `nodels` command. Ensure that all expected (and no unexpected) nodes are listed. You may find it useful to create new group names to identify Moab-managed nodes.

```
[root@h0 moab]# nodels hyper,compute
h1
h2
h3
h4
h5
h7
kvmn1
kvmn10
kvmn2
kvmn3
kvmn4
kvmn5
kvmn6
kvmn7
kvmn8
[root@h0 moab]#
```

### G.5.1.C Reporting Node Status

Verify that all nodes report their status correctly using the xCAT `nodestat` command. Ensure that all nodes show the correct status (`sshd`, `installing`, `noping`, and so forth); there should not be any timeouts or error messages.

```
[root@h0 moab]# nodestat hyper,compute |sort
h1: pbs,sshd
h2: pbs,sshd
h3: pbs,sshd
h4: pbs,sshd
h5: pbs,sshd
h7: noping
kvmn10: noping
kvmn1: pbs,sshd
kvmn2: pbs,sshd
kvmn3: pbs,sshd
kvmn4: pbs,sshd
kvmn5: pbs,sshd
kvmn6: pbs,sshd
kvmn7: pbs,sshd
kvmn8: noping
kvmn9: noping
[root@h0 moab]#
```

### G.5.1.D Verifying Hardware Management Configuration

Verify that all nodes that Moab will manage have hardware management interfaces correctly configured using the xCAT `nodels` and `rpower` commands. After each of the `rpower` commands, verify the requested state was achieved with `rpower stat`.

```
[root@h0 moab]# nodels h1,kvmm1 nodehm.mgt nodehm.power
h1: nodehm.power: ilo
h1: nodehm.mgt: ilo
kvmm1: nodehm.power: kvm
kvmm1: nodehm.mgt: kvm
[root@h0 moab]# rpower h1,kvmm1 off
h1: off
kvmm1: off
[root@h0 moab]# rpower h1,kvmm1 stat
h1: off
kvmm1: off
[root@h0 moab]# rpower h1,kvmm1 boot
h1: on reset
kvmm1: on reset
[root@h0 moab]# rpower h1,kvmm1 stat
h1: on
kvmm1: on
[root@h0 moab]#
```

### G.5.1.E Verifying Provisioning Images

Verify that all operating system images that Moab uses are configured correctly in xCAT. For stateful images, test that all combinations of operating system, architecture, and profile install correctly.

```
[root@h0 moab]# rinstall -o centosX.X -a x86_64 -p hyper h1
h1: install centosX.X-x86_64-hyper
h1: on reset
[root@n100 ~]# sleep 15 && nodestat n05
n05: ping install centosX.X-x86_64-hyper
[root@h0 moab]#
```

For stateless images, test that nodes are able to network boot the images:

```
[root@h0 moab]# nodech h5 nodetype.os=centosX.X nodetype.arch=x86_64
nodetype.profile=hyper
[root@h0 moab]# nodeset h5 netboot
h5: netboot centosX.X-x86_64-hyper
[root@h0 moab]# rpower h5 boot
h5: on reset
[root@h0 moab]# sleep 60 && nodestat h5
h5: pbs, sshd
[root@h0 moab]#
```

#### Related Topics

- [Native Resource Manager Overview](#)
- [Resource Provisioning](#)

## Appendix H: Interfacing with Moab (APIs)

Moab provides numerous interfaces allowing it to monitor and manage most services and resources. It also possesses flexible interfaces to allow it to interact with peer services and applications as both a broker and an information service. This appendix is designed to provide a general overview and links to more detailed interface documentation.

Moab interfaces to systems providing various services and using various protocols. This appendix is designed to assist users who want to enable Moab in new environments using one of the existing interfaces. It does not cover the steps required to create a new interface.

In this section:

- [H.1 Accounting Interfaces - page 1385](#)
- [H.2 Grid Interfaces - page 1385](#)
- [H.3 Identity and Credential Management Interfaces - page 1386](#)
- [H.4 Job Submission and Management Interface - page 1386](#)
- [H.5 Query and Control APIs - page 1387](#)
- [H.6 Resource Management Interfaces - page 1389](#)

### H.1 Accounting Interfaces

Moab accounting interfaces allow Moab to export local utilization statistics, events, and accounting information to site specific scripts.

### H.2 Grid Interfaces

Moab provides interfaces to allow interaction with various grid brokers and services. These interfaces allow Moab to provide services, as well as utilize services.

#### H.2.1 Services Utilized

- Information Services (import and utilize information service data in making scheduling decisions)
- Job Migration
- Data Migration

- Credential Mapping
- Security and Delegation

See [Chapter 23: Moab Workload Manager for Grids - page 948](#) for more information on utilized services.

## H.2.2 Services Provided

- Information Services (provide resource, workload, and credential information)
- Job Migration
- Data Migration
- Credential Mapping

See [23.1 Grid Basics - page 950](#) for more information on provided services.

## H.3 Identity and Credential Management Interfaces

Moab's identity and credential management interfaces allow Moab to exchange credential and user configuration, access, policy, and usage information:

- [Identity Manager](#)
- [Accounting Manager](#)
- [Moab Workload Manager for Grids](#)

## H.4 Job Submission and Management Interface

Moab's Job Submission and Management Interface provides interfaces to query resource availability, submit, modify, and manage jobs, and query the status of active and completed jobs:

- Resource Availability Query
  - Determine quantity, state, and configuration of configured resources (idle, busy, and down nodes)
  - Determine quantity and configuration of all available resources (idle nodes)
  - Determine resources available subject now and in the future for potential job

- Determine best target cluster destination for potential job
- Determine largest/longest job that could start immediately
- Determine estimated start time for potential job
- Determine earliest guaranteed start time for potential job
- Reserve Resources
  - Reserve specific resources for desired time frame
- Submit Job ([XML format](#))
  - Submit job to specific cluster
  - Submit job to global job queue
- Manage Job
  - Hold job
  - Adjust job priority
  - Modify job executable, args, data requirements, job dependencies, duration, hostcount, or other attributes
  - Suspend/resume job
  - Checkpoint/requeue job
  - Cancel job
  - Migrate job
  - Adjust job quality of service (QoS)
- Query Job
  - Determine job state, utilization, or output results for idle, active, or completed job
  - Determine estimated start time
  - Determine guaranteed start time

## H.5 Query and Control APIs

The Moab Cluster and Grid Suites provides a (Moab) workload manager server that supports a broad array of client services. These services can be directly accessed via Moab client commands. The Query and Control APIs allow external portals and services to obtain information about compute resources, workload, and usage statistics.

## H.5.1 CLI (Command Line Interface) XML API

All Moab client commands can report results in XML format to allow the information to be easily integrated into peer services, portals, databases, and other applications. To request that a client command report its output in XML, specify the `--format=xml` flag as in the following example:

```
> showq --format=xml
<Data>
<Object>queue</Object>
<cluster LocalActiveNodes="1" LocalAllocProcs="1" LocalIdleNodes="0"
LocalIdleProcs="3" LocalUpNodes="1"
  LocalUpProcs="4" RemoteActiveNodes="0" RemoteAllocProcs="0" RemoteIdleNodes="0"
RemoteIdleProcs="0"
  RemoteUpNodes="0" RemoteUpProcs="0" time="1128451812"></cluster>
<queue count="1" option="active">
<job AWDuration="11672" EEDuration="1128451812" Group="[DEFAULT]" JobID="Moab.2"
MasterHost="cw2" PAL="2"
  QOS="bug3" ReqAWDuration="54000" ReqNodes="1" ReqProcs="1" RsvStartTime="1128451812"
RunPriority="0"
  StartPriority="1" StartTime="1128451812" StatPSDed="11886.580000"
StatPSUtl="11886.580000" State="Running"
  SubmissionTime="1128451812" SuspendDuration="0" User="smith"></job>
</queue>
<queue count="1" option="eligible">
<job EEDuration="1128451812" Group="jacksond" JobID="customer.35" QOS="bug"
ReqAWDuration="3600"
  ReqProcs="1" StartPriority="1" StartTime="0" State="Idle"
SubmissionTime="1128451812" SuspendDuration="0"
  User="johnson"></job>
<queue><queue count="0" option="blocked"></queue>
</Data>
```

### H.5.1.A Common Query/Control Services

- jobs
  - query status - [mdiag -j](#) (XML details)
  - submit - [msub](#) (XML format)
  - cancel - [mjobctl -c](#)
- nodes
  - query status - [mdiag -n](#) (XML details)
  - create resource reservation - [mrsvctl -c](#)
  - destroy resource reservation - [mrsvctl -r](#)



## H.6 Resource Management Interfaces

Moab can monitor, schedule, and control services and resources using multiple protocols. These protocols include the following:

- [script/flat file](#)
- [Resource Manager Specific Interfaces](#) - Torque

Using the resource manager interfaces, Moab can do the following:

- monitor resources (compute host, network, storage, and software license based resources)
  - load configuration, architecture, and feature information
  - load state, utilization, and workload information
  - load policy and ownership information
- manage resources
  - dynamically reconfigure and reprovision resource hardware (processors, memory, etc.)
  - dynamically reconfigure and reprovision resource software (operating system, application software, filesystem mounts, etc.)
  - dynamically reconfigure and reprovision resource security (VPNs, VLANs, host security, etc.)
- monitor workload (batch jobs, interactive jobs, persistent services, dynamic services, distributed services)
  - load state, resource requirement, and required environment information
  - load user, group, and credential information
  - load utilization, resource allocation, and policy information
- manage workload
  - migrate jobs from one resource to another (intra-cluster and inter-cluster)
  - modify jobs for translation and optimization purposes
  - suspend, resume, checkpoint, restart, and cancel jobs
- query cluster policies and configuration

# Appendix I: Considerations for Large Clusters

There are several key considerations in getting a batch system to scale.

In this section:

- [I.1 Handling Large Jobs - page 1390](#)
- [I.2 Handling Large Numbers of Jobs - page 1390](#)
- [I.3 Handling Large Numbers of Nodes - page 1395](#)
- [I.4 Handling Large SMP Systems - page 1396](#)
- [I.5 Resource Manager Scaling - page 1396](#)
- [I.6 Server Sizing - page 1397](#)

---

## Related Topics

- [Appendix D: Adjusting Default Limits](#)

## I.1 Handling Large Jobs

For large jobs, additional parameters beyond those specified for [large node](#) systems might be required. These include settings for the maximum number of [tasks per job](#), and the maximum number of [nodes per job](#).

## I.2 Handling Large Numbers of Jobs

In this topic:

- [I.2.1 Set a Minimum RMPOLLINTERVAL - page 1391](#)
- [I.2.2 Reduce Command Processing Time - page 1391](#)
- [I.2.3 Minimize Job Processing Time - page 1392](#)
- [I.2.4 Load All Non-Completed Jobs at Startup - page 1393](#)
- [I.2.5 Reducing Job Start Time - page 1393](#)
- [I.2.6 Reducing Job Reservation Creation Time - page 1393](#)
- [I.2.7 Optimizing Backfill Time - page 1394](#)

- [I.2.8 Constraining Moab Logging - LOGLEVEL - page 1394](#)
- [I.2.9 Preemption - page 1394](#)
- [I.2.10 Handling Transient Resource Manager Failures - page 1394](#)
- [I.2.11 Constrain the Number of Jobs Preempted Per Iteration - page 1394](#)
- [I.2.12 Scheduler Settings - page 1395](#)
- [I.2.13 Configure Torque for Large Job Numbers - page 1395](#)

## I.2.1 Set a Minimum RMPOLLINTERVAL

With event driven resource managers like Torque, each time a job is submitted the resource manager notifies the scheduler. In an attempt to minimize response time, the scheduler starts a new scheduling cycle to determine if the newly submitted job can run. In systems with large numbers of jobs submitted at once, this might not result in the desired behavior for two reasons. First, by scheduling at every job submission Moab schedules newly submitted jobs onto available resources in a first come, first served basis rather than evaluating the entire group of new jobs at once and optimizing the placement accordingly. Second, by launching a scheduling iteration for every job submitted, Moab places a heavy load on the resource manager. For example, if a user were to submit 1000 new jobs simultaneously, for each job submitted, the resource manager contacts the scheduler, the scheduler starts a new iteration, and in this iteration, the scheduler contacts the resource manager requesting updated information on all jobs and resources available.

Setting a minimum RMPOLLINTERVAL causes the scheduler to not process jobs as quickly as they are submitted, but rather to wait a minimum amount of time to allow more jobs be submitted and to process these new jobs in groups.

```
RMPOLLINTERVAL 30,60
```

*If the system is busy, schedule every 30 seconds. If it is not busy, schedule every 60 seconds.*

## I.2.2 Reduce Command Processing Time

If your system's scheduling cycle regularly takes longer than the CLIENTTIMEOUT value, you can configure Moab to fork a copy of itself that will respond to certain information-only client commands ([checkjob](#), [showbf](#), [showres](#), and [showstart](#)). This enables you to run intense diagnostic commands while Moab is in the middle of its scheduling process.

When you set UIMANAGEMENTPOLICYFORKCLIENTUIPORT<port number> on the server side, Moab forks a copy of itself that will listen for client commands on a separate port. For example, systems that run client commands, such as submit hosts, can set CLIENTUIPORT41560. This will allow the clients to run commands such as *checkjob*, *showbf*, *showres* and *showstart* and have cached information returned from the

previous scheduling iteration. Moab prints a disclaimer at the top of each command that was populated by the forked process stating that the information may be an iteration behind.

**i** See [CLIENTTIMEOUT](#), [CLIENTUIPORT](#), and [UIMANAGEMENTPOLICY](#) for parameter information.

*Example I-1: Sample configuration:*

```
UIMANAGEMENTPOLICY    FORK
CLIENTUIPORT          41560
```

*Moab forks a copy of itself on port 41560, where it will watch for checkjob, showbf, showres, and showstart commands until the main scheduling process completes.*

*Example I-2: Sample command output:*

```
$ checkjob 34
-----
NOTE: The following information has been cached by the remote server
and may be slightly out of date.
-----

job 34

State: Idle
Creds: user:wightman group:company class:batch
WallTime: 00:00:00 of 00:01:00
SubmitTime: Thu May 22 14:17:06
(Time Queued Total: 00:00:18 Eligible: 00:00:18)

TemplateSets: DEFAULT
Total Requested Tasks: 1

Req[0] TaskCount: 1 Partition: ALL

SystemID: scale
SystemJID: 34

IWD: $HOME/test/scale
SubmitDir: $HOME/test/scale
Executable: sleep 60
```

## I.2.3 Minimize Job Processing Time

Use the [ENABLEHIGHTHROUGHPUT](#) parameter. By default, Moab processes all job attributes, filters, remap classes, job arrays, and other information when a job is submitted. This requires full access to the Moab configuration and significantly increases the processing time Moab needs when jobs are submitted. By setting [ENABLEHIGHTHROUGHPUT](#) to `TRUE`, Moab stores the job information in an internal

queue and returns the job ID immediately. The internal queue is processed when Moab begins its next scheduling iteration. This enables Moab to process hundreds of jobs per second rather than 20-30 per second. Because the jobs are processed in a separate queue after the job has been returned, it is recommended that `MAILPROGRAM` be configured. Moab will send an email to the user if a job is rejected.

Because the job is not fully processed, some attributes might change after the job has been submitted. For example, when a job class is remapped, the new class is not reflected until Moab begins its next scheduling iteration. Additionally, job arrays are not instantiated until Moab begins its next scheduling cycle.

**i** If `ENABLEHIGHTHROUGHPUT` is `TRUE`, you must set `NODEALLOCATIONPOLICY` to `FIRSTAVAILABLE`.

## I.2.4 Load All Non-Completed Jobs at Startup

Use the `LOADALLJOB` parameter. By default, Moab loads non-complete jobs for active resource managers only. By setting `LOADALLJOB` to `TRUE`, Moab will load all non-complete jobs from all checkpoint files at startup, regardless of whether their corresponding resource manager is active.

## I.2.5 Reducing Job Start Time

Use the `ASYNCSTART` parameter. By default, Moab will launch one job at a time and verify that each job successfully started before launching a subsequent job. For organizations with large numbers of very short jobs (less than 2 minutes in duration), the delay associated with confirming successful job start can lead to productivity losses. If tens or hundreds of jobs must be started per minute, and especially if the workload is composed primarily of serial jobs, then the resource manager `ASYNCSTART` flag can be set. When set, Moab will launch jobs optimistically and confirm success or failure of the job start on the subsequent scheduling iteration. Also consider adding the `ASYNCDELETE` flag if users frequently cancel jobs.

## I.2.6 Reducing Job Reservation Creation Time

Use the `RMCFGJOBRSVRECREATE` attribute. By default, Moab destroys and re-creates job reservations each time a resource manager updates any aspect of a job. Historically, this stems from the fact that certain resource managers would inadvertently or intentionally migrate job tasks from originally requested nodes to other nodes. To maintain synchronization, Moab would re-create reservations each iteration therefore incorporating these changes. On most modern resource managers, these changes never occur, but the

effort required to handle this case grows with the size of the cluster and the size of the queue. Consequently, on very large systems with thousands of nodes and thousands of jobs, a noticeable delay is present. By setting `JOBRSVRECREATE` to `FALSE` on resource managers that do not exhibit this behavior, significant time savings per iteration can be obtained.

## I.2.7 Optimizing Backfill Time

Use the `OPTIMIZEDBACKFILL` flag. Speeds up backfill when a system reservation is in use.

## I.2.8 Constraining Moab Logging - LOGLEVEL

Use the `LOGLEVEL` parameter. When running on large systems, setting `LOGLEVEL` to 0 or 1 is normal and recommended. Only increase `LOGLEVEL` above 0 or 1 if you have been instructed to do so by Moab support.

## I.2.9 Preemption

When preemption is enabled Moab can take considerably more time scheduling jobs for every scheduling iteration. Preemption increases the number of options available to Moab and therefore takes more time for Moab to optimally place jobs. If you are running a large cluster or have more than the usual amount of jobs (>10000), consider disabling preemption. If disabling preemption is not possible, consider limiting its scope to only a small subset of jobs (as both preemptors and preemptees).

## I.2.10 Handling Transient Resource Manager Failures

Use the `RMCFGMAXITERATIONFAILURECOUNT` attribute.

## I.2.11 Constrain the Number of Jobs Preempted Per Iteration

Use the `JOBMAXPREEMPTPERITERATION` parameter.

**i** For very large job count systems, configuration options controlling the maximum supported limits might need to be adjusted, including the maximum number of reservations and the maximum number of supported evaluation ranges.

## I.2.12 Scheduler Settings

If using Moab, there are a number of parameters that can be set on the scheduler, which may improve Torque performance. In an environment containing a large number of short-running jobs, the `JOBAGGREGATIONTIME` parameter can be set to reduce the number of workload and resource queries performed by the scheduler when an event based interface is enabled. Setting `JOBAGGREGATIONTIME` instructs the scheduler to ignore events coming from the resource manager and perform scheduling at regular intervals, rather than around resource manager events. If the `pbs_server` daemon is heavily loaded and PBS API timeout errors (i.e., 'Premature end of message') are reported within the scheduler, the `TIMEOUT` attribute of the `RMCFG` parameter can be set with a value of between 30 and 90 seconds.

## I.2.13 Configure Torque for Large Job Numbers

Torque's `use_job_subdirs` server parameter enables Torque to handle large numbers of jobs more efficiently. For more information, see 'use\_jobs\_subdirs' in the *Torque Resource Manager Administrator Guide*.

## I.3 Handling Large Numbers of Nodes

For very large clusters ( $\geq 10,000$  processors) default scheduling behavior might not scale as desired. To address this, the following parameters should be considered:

Parameter	Recommended Settings
<b>RMPOLLINTERVAL</b>	In large node environments with large and long jobs, scheduling overhead can be minimized by increasing <code>RMPOLLINTERVAL</code> above its default setting. If an event-driven resource management interface is available, values of two minutes or higher can be used. Scheduling overhead can be determined by looking at the scheduling load reported by <code>mdiag -S</code> .
<b>LIMITEDNODECP</b>	Startup/shutdown time can be minimized by disabling full node state checkpointing that includes some statistics covering node availability.
<b>SCHEDCFG FLAGS="FASTRSVSTARTUP"</b>	When you have reservations on a large number of nodes, it can take Moab a long time to recreate them on startup. Setting the <code>FASTRSVSTARTUP</code> scheduler flag greatly reduces startup time.

\* For clusters where the number of nodes or processors exceeds 50,000, the maximum stack size for the shell in which Moab is started might need to be increased (as Moab might

crash if the stack size is too small). On most UNIX/Linux based systems, the command `ulimit -s unlimited` can be used to increase the stack size limit before starting Moab. This can be placed in your Moab startup script.

**i** See [Appendix D](#) for further information on default and supported object limits.

Avoid adding large numbers of **NODECFG** lines in the `moab.cfg` or `moab.d/*.cfg` files to keep the Moab boot time low.

For example, adding a configuration line to define features for each node in a large cluster (such as `NODECFG[x] Features+=green,purple`) can greatly increase the Moab boot time. If Moab processes 15 node configuration lines per second for a 50,000-node system, it could add approximately 55 minutes of node configuration processing to the Moab boot time.

In this case, it is better to define the node features in the resource manager configuration.

## I.4 Handling Large SMP Systems

For large-way SMP systems (> 512 processors/node) Moab defaults might need adjustment.

Parameter	Recommended Settings
<b>MAXRSVPERNODE</b>	By default, Moab does not expect more than 64 jobs per node to be running or have future reservations. Increasing this parameter to a value larger than the expected maximum number of jobs per node is advised.

## I.5 Resource Manager Scaling

### Proper Resource Manager Configuration

- Torque
  - See Large Cluster Considerations in the *Torque Resource Manager Administrator Guide*.



## I.6 Server Sizing

See 'Hardware Requirements' in the *Moab HPC Suite Installation and Configuration Guide* for recommendations.

## Appendix J: Configuring Moab as a Service

Scripts that follow can be used to start up Moab services automatically upon a reboot. To enable a service script, copy the script to `/etc/rc.d/init.d/S97moab`, edit the file to make needed localization changes (adjust binary paths, execution user, etc.), and add links to the `rc3.d` and `rc5.d` directories as in the example that follows:

```
> cp mwm.service /etc/rc.d/init.d/S97moab
> vi /etc/rc.d/init.d/S97moab
   (make needed localizations)
> ln -s /etc/rc.d/init.d/S97moab /etc/rc.d/rc3.d
> ln -s /etc/rc.d/init.d/S97moab /etc/rc.d/rc5.d
```

In this section:

[J.1 Moab Grid Scheduler Service Script - page 1398](#)

[J.2 Moab Workload Manager Service Scripts - page 1399](#)

### J.1 Moab Grid Scheduler Service Script

#### Sample Script

```
#!/bin/sh
#
# moab          This shell script takes care of starting and stopping
#              the Moab Workload Manager.
#
# Source function library.
if [ -f /etc/init.d/functions ] ; then
    . /etc/init.d/functions
elif [ -f /etc/rc.d/init.d/functions ] ; then
    . /etc/rc.d/init.d/functions
else
    exit 0
fi

# Read in the command arguments
case "$1" in
    start)
        # Next start Moab Workload Manager...
        echo -n "Starting Moab Workload Manager: "
        daemon su -l -s /bin/sh root -c "/opt/moab/sbin/moab"
        echo
        ;;
    stop)
        # Stop Moab
        echo -n "Shutting down Moab Workload Manager: "
        killproc moab
```

```

        echo
        ;;

restart)
    $0 stop
    $0 start
    ;;

*)
    echo $"Usage: moab {start|stop|restart}"
    exit 1
esac

exit 0

```

## J.2 Moab Workload Manager Service Scripts

- [Moab Workload Manager Script](#)
- [Moab Workload Manager + Torque Script](#)

### Moab Workload Manager Script

```

#!/bin/sh
#
# moab          This shell script takes care of starting and stopping
#              the Moab Workload Manager.
#
# Source function library.
if [ -f /etc/init.d/functions ] ; then
    . /etc/init.d/functions
elif [ -f /etc/rc.d/init.d/functions ] ; then
    . /etc/rc.d/init.d/functions
else
    exit 0
fi

# Read in the command arguments
case "$1" in
start)
    # Next start Moab Workload Manager...
    echo -n $"Starting Moab Workload Manager: "
    daemon su -l -s /bin/sh root -c "/opt/moab/sbin/moab"
    echo
    ;;

stop)
    # Stop Moab
    echo -n $"Shutting down Moab Workload Manager: "
    killproc moab
    echo
    ;;

restart)
    $0 stop
    $0 start

```

```

        ;;
    *)
        echo $"Usage: moab {start|stop|restart}"
        exit 1
    esac
exit 0

```

## Moab Workload Manager + Torque Script

```

#!/bin/sh
#
# moab          This shell script takes care of starting and stopping
#              the TORQUE resource manager and Moab Workload Manager.
#
# description: TORQUE is a scalable resource manager which manages jobs in
#              cluster environments. Moab is a cluster scheduler which uses
#              TORQUE to schedule jobs on that cluster.
#
# Source function library.
if [ -f /etc/init.d/functions ] ; then
    . /etc/init.d/functions
elif [ -f /etc/rc.d/init.d/functions ] ; then
    . /etc/rc.d/init.d/functions
else
    exit 0
fi

# Read in the command arguments
case "$1" in
    start)
        # Start TORQUE services first...

        echo -n $"Starting TORQUE services: "
        daemon /usr/local/torque/sbin/pbs_mom
        daemon /usr/local/torque/sbin/pbs_server
        echo

        # Next start Moab Workload Manager...
        echo -n $"Starting Moab Workload Manager: "
        daemon su -l -s /bin/sh root -c "/opt/moab/sbin/moab"
        echo
        ;;

    stop)
        # Stop Moab first...

        echo -n $"Shutting down Moab Workload Manager: "
        killproc moab
        echo

        echo -n $"Shutting down TORQUE services: "
        killproc pbs_server
        echo
        killproc pbs_mom
        echo
        ;;

    restart)
        $0 stop

```

```
    $0 start
    ;;
*)
    echo $"Usage: moab {start|stop|restart}"
    exit 1
esac

exit 0
```

## Appendix K: Migrating from Maui 3.2

This guide is intended to help facilitate migrating from Maui to Moab. If you do not have Moab yet, you can download a [free evaluation version](#). At a high level, migrating from Maui 3.2 to Moab involves minimal effort. In fact, Moab fully supports all Maui parameters and commands. Migration can consist of nothing more than renaming `maui.cfg` to `moab.cfg` and launching Moab using the `moab` command. With this migration, the biggest single issue is becoming aware of all the new facilities and capabilities available within Moab. Beyond this, migration consists of a few minor issues that might require attention such as some [statistics and priorities](#).

Another approach of migrating from Maui to Moab is to configure Moab in Monitor mode and run it beside Maui. Maui will continue to perform the scheduling and control workload. Moab will simply monitor the cluster environment using the policies configured in `moab.cfg`. Moab will not have the ability to affect workload, providing a safe and risk-free environment to evaluate Moab without affecting your production environment. You can also have Moab capture resource and workload trace files and allow Moab to simulate what it would have done if it controlled workload. When you feel comfortable with and want to run Moab live on your cluster, all you need to do is change the mode to NORMAL, stop Maui, and restart Moab. Current jobs will remain running and Moab will take over control of scheduling.

As with any migration, we suggest that you back up important files such as the following: `maui.cfg`, `maui.log` and `maui.ck`.

In this section:

- [K.1 Migrating from Maui to Moab - page 1402](#)
- [K.2 Other Notes - page 1403](#)
- [K.3 Running Maui and Moab Side-By-Side - page 1404](#)

### K.1 Migrating from Maui to Moab

1. Install Moab Workload Manager.
2. Copy your `maui.cfg` file to the `MOABHOMEDIR/etc (/opt/moab/etc)` and rename it `moab.cfg`.
3. Stop Maui.
4. Start Moab.

5. If applicable, re-apply those configurations found in the [Statistics and Checkpointing](#) section that need adjustment after migration, as well as any parameters in `moab.cfg` that point to a Maui file like `maui.log`.

## K.2 Other Notes

This section provides more information on the minor differences between Maui and Moab, and the changes you might need to make.

In this topic:

[K.2.1 File Naming - page 1403](#)

[K.2.2 Statistics and Checkpointing - page 1403](#)

[K.2.3 Verify Configuration File Compatibility - page 1404](#)

[K.2.4 Environment Variables - page 1404](#)

### K.2.1 File Naming

Moab uses slightly different naming than Maui. The following table displays these changes:

File	Maui	Moab
<b>executable</b>	<code>maui</code>	<code>moab</code>
<b>logs</b>	<code>maui.log</code>	<code>moab.log</code>
<b>configuration file</b>	<code>maui.cfg</code>	<code>moab.cfg</code>

### K.2.2 Statistics and Checkpointing

Moab supports Maui version 3.2 or higher workload statistics, allowing it to process historical statistics based on these statistics files. No changes are required to use these statistics.

Moab does not support the Maui 3.2 checkpointing format. Because of this, state information checkpointed under Maui will not be available at the time of the migration. The loss of this information will have the following impact:

- Admin reservations, if any, will need to be re-created.
- Processed credential and scheduler statistics (displayed by `showstats`) will be lost.
- Admin job system priority configured by the `setspri` command and QoS assignments configured by the `setqos` command, if any, will be lost.

### K.2.3 Verify Configuration File Compatibility

The command `mdiag -C` will perform diagnostics on your new configuration file and might prove helpful in identifying any issues.

### K.2.4 Environment Variables

Scheduler environment variables are supported under Moab with obvious naming changes. Sample environment variables follow:

Maui	Moab
MAUIHOMEDIR	MOABHOMEDIR
MAUIDEBUG	MOABDEBUG
MAUICRASHVARIBALE	MOABCRASHVARIABLE
MAUIENABLELOGBUFFERING	MOABENABLELOGBUFFERING
MAUIRECOVERYACTION	MOABRECOVERYACTION
MAUI-COMMANDS-PATH	MOAB-COMMANDS-PATH
MAUIENABLELOGBUFFERING	MOABENABLELOGBUFFERING

## K.3 Running Maui and Moab Side-By-Side

1. Install Moab Workload Manager on your cluster. (Installation steps will differ slightly from a typical Moab installation.)



- a. Run `./configure`.
  - b. Run `make`.
  - c. You will need to set your `MOABHOMEDIR` environment variable to the location where you built Moab by typing `export MOABHOMDIR=[make directory]`.
2. To have Moab use all the same policies as Maui, copy `maui.cfg` to the `MOABHOMEDIR/etc` and rename it `moab.cfg`.
    - You can also start your `moab.cfg` file from scratch. Just use the `moab.cfg` already in the `MOABHOMEDIR/etc`.
  3. Make sure that the port in `moab.cfg` is different than the port used in `maui.cfg`.
  4. In the `moab.cfg` file, add the parameter, `SERVERMODE=MONITOR`.
    - If you used the `moab.cfg` from scratch, on the `SCHEDCFG` line add `MODE=MONITOR`.
  5. You will need to either put the Moab commands in your environment path (located in `MOABHOMEDIR/bin`) or run the commands from their location if you still want to use the Maui commands in your environment path.
  6. Run Moab Workload Manager using the `moab` command located in `MOABHOMEDIR/bin`.

## Appendix L: Node Allocation Plug-in Developer Kit

Each time Moab schedules a job, it must choose the nodes on which the job will run. Moab uses the Node Allocation policy to select the available nodes to be used. Because there are so many different systems and cluster topologies, you now have the ability to create and use a node allocation plugin for allocating nodes based on your cluster's interconnect topology.

The plugin policy allows you to write your own algorithm to choose which nodes will be used. This algorithm is contained in a shared library that Moab loads at run time.

To obtain the Plug-in Developer Kit (PDK) with the header file and example code, contact your sales representative.

In this section:

[L.1 Moab Configuration - page 1406](#)

[L.2 Writing the Plug-In - page 1408](#)

### L.1 Moab Configuration

The actual loading of a plug-in is accomplished by specifying the plug-in in the Moab configuration file, `moab.cfg`.

In this topic:

[L.1.1 Moab.cfg - page 1406](#)

[L.1.2 Syntax Rules - page 1407](#)

[L.1.3 Troubleshooting - page 1408](#)

#### L.1.1 Moab.cfg

We recommend that you store all Moab plug-ins in the `$MOABHOMEDIR/lib` directory (e.g., `/opt/moab/lib`) as shared libraries (`*.so`). The name of the actual plug-in shared library file is up to the plug-in developer, which means you must give the correct name in the `moab.cfg` file to form the absolute plug-in filename.

If a plug-in's specified shared library filename starts with a forward slash (/), it is an absolute file path name and Moab simply uses it without alteration. For example, if a plug-in's specified shared library filename is `/opt/moab/plugins/plugin.so`, Moab will use it as the absolute plug-in file path name.

If a plug-in's specified shared library filename does not start with a forward slash (/), it is a plug-in name and Moab forms the plug-in's absolute path name by concatenating the Moab home directory, `/lib/lib`, the specified plug-in name, and `.so` to obtain the absolute path name. For example, if the `$MOABHOMEDIR` environment variable contains `/opt/moab` and the plug-in name is `plugin`, Moab will create `/opt/moab/lib/libplugin.so` and use it as the absolute plug-in file path name.

## L.1.2 Syntax Rules

In order for Moab to use a plug-in for the Node Allocation policy, instead of a built-in Moab policy, you must configure the policy in the `moab.cfg` file with the value `'PLUGIN:'` followed by the plug-in's shared library file name. The examples below assume the environment variable `$MOABHOMEDIR` has a value of `/opt/moab`. Note the use of relative and absolute plug-in shared library file path names in the parameter value and how they affect Moab's construction of the full path name.

Partition	Plug-in Name	moab.cfg Parameter	Moab-derived Full Path Name
<b>global</b>	<code>plugin.so</code>	<code>NODEALLOCATIONPOLICY PLUGIN:plugin.so</code>	<code>/opt/moab/lib /libplugin.so</code>
<b>global</b>	<code>/usr/local /plugins /plugin.so</code>	<code>NODEALLOCATIONPOLICY PLUGIN:/usr/local /plugins/plugin.so</code>	<code>/usr/local/plugins /plugin.so</code>
<b>abc</b>	<code>plugin.so</code>	<code>PARCFG[abc] NODEALLOCATIONPOLICY =PLUGIN:plugin.so</code>	<code>/opt/moab/lib /libplugin.so</code>
<b>xyz</b>	<code>/usr/local /plugins /plugin.so</code>	<code>PARCFG[xyz] NODEALLOCATIONPOLICY =PLUGIN:/usr/local /plugins/plugin.so</code>	<code>/usr/local/plugins /plugin.so</code>

## L.1.3 Troubleshooting

There are several commands that can be used to confirm that the plug-in Node Allocation Policy was loaded properly.

### mschedctl -l

`mschedctl -l` is used to print out Moab's in memory configurations. If the plug-in policy, with its full path, doesn't show for the configured partition then Moab failed to load the partition. Note that when the `NODEALLOCATIONPOLICY` is configured globally, it is configured on the 'ALL' partition.

```
$ mschedctl -l -v|grep ^NODEALLOCATIONPOLICY
NODEALLOCATIONPOLICY[ALL] PLUGIN:/opt/moab/lib/libfirstavailable.so
NODEALLOCATIONPOLICY[a] PLUGIN:/opt/moab/lib/liblastavailable.so
NODEALLOCATIONPOLICY[b] CONTIGUOUS
NODEALLOCATIONPOLICY[c] PLUGIN:/opt/moab/lib/libfirstavailable.so
NODEALLOCATIONPOLICY[d] [NONE]
```

### mdiag -C

`mdiag -C` is used to validate the `moab.cfg` configuration. With a plug-in node allocation policy, Moab will validate that it can successfully load the plug-in and that all of the required symbols are present.

```
$ mdiag -C
...
INFO: line #35 is valid: 'NODEALLOCATIONPOLICY PLUGIN:firstavailable'
INFO: line #36 is valid: 'PARCFG[a]NODEALLOCATIONPOLICY=PLUGIN:lastavailable'
INFO: line #37 is valid: 'PARCFG[b]NODEALLOCATIONPOLICY=CONTIGUOUS'
INFO: line #38 is valid: 'PARCFG[d]NODEALLOCATIONPOLICY=PLUGIN:firstavailable'
```

## L.2 Writing the Plug-In

In this topic:

[L.2.1 Node Allocation Plug-in - page 1408](#)

[L.2.2 API and Data Structures - page 1409](#)

### L.2.1 Node Allocation Plug-in

A plug-in is a shared library that has specific functions and variables that will be called directly from Moab. The plug-in conforms to a C language API. The API is specified through an include file: `moab-plugin.h`. This file must be included in the plug-in code. The

include file provides function definitions, structures and variables that will be used when communicating with Moab.

When you write the plug-in, you need to ensure that the plug-in code is robust. If the plug-in crashes, Moab will crash. You will need to handle your own memory appropriately. If the plug-in has memory leaks, Moab will have similar issues. If you want to maintain logs, the plug-in will need to be responsible for its own logging.

## L.2.2 API and Data Structures

The Application Programmer Interface (API) for the Moab Node Allocation plug-in consists of three data items and three entry points that must be supplied to Moab by the plug-in.

Plug-in Supplied Data	Description
<b>const char *PLUGIN_NAME = "Node Allocation plugin 1.1";</b>	This character pointer is used by Moab when logging information regarding the operation of the plug-in.
<b>const char *PLUGIN_TYPE = PLUGIN_TYPE_ NAME_ NODEALLOCATION;</b>	This character pointer is used by Moab to verify the type of plug-in. The value of this data is supplied by the <code>moab-plugin.h</code> source file. The plug-in must set this as shown so that Moab does not attempt to use a plug-in incorrectly. Moab uses this to determine whether the plug-in API type is correct and to allow Moab to correctly communicate with the plug-in.
<b>const char *PLUGIN_VERSION = PLUGIN_API_ VERSION;</b>	This character pointer is used by Moab to verify the API version number. The value of this data is supplied by the <code>moab-plugin.h</code> source file. The plug-in must set this as shown so that the correct version of the <code>moab-plugin.h</code> is supplied to Moab. Moab uses this to determine whether the API version is correct and to allow Moab to correctly communicate with the plug-in.

Load Time API	Description
<b>initialize()</b>	<pre>int initialize(const char *name, void **data_handle)</pre> <p>The plug-in must supply an <code>initialize()</code> entry point. This entry point is called for each use instance of the plug-in. For example, if the plug-in is used on two different partitions, the <code>initialize()</code> entry point will be called once for each partition.</p> <ul style="list-style-type: none"> <li>Name — The name is the unique identifier that is used to distinguish multiple instances of the plug-in and for logging. When configured</li> </ul>

Load Time API	Description
	<p>globally, the name “ALL” will be given.</p> <ul style="list-style-type: none"> <li>Data handle — The data_handle points to a location where the plug-in should store a pointer to any internal data needed by the plug-in between calls to the API. The actual format and structure of the data is up to the plug-in. Moab will supply this pointer back to the plug-in each time a plug-in entry point is called. This data can provide context for the plug-in usage instance.</li> </ul>
<b>Return codes</b>	<p>The initialize() entry point should return one of two return statuses as defined in moab-plugin.h:</p> <pre>#define  PLUGIN_RC_SUCCESS  0 #define  PLUGIN_RC_FAILURE  1</pre>
<b>Gathering node info</b>	<p>The initialize() entry point must gather any information about system nodes, their topology, interconnection, and configuration that it needs to make correct node allocations. Since Moab does not know what information the plug-in might need, the plug-in must gather this information itself.</p>
<b>Memory considerations</b>	<p>The plug-in may allocate memory for temporary or persistent data as needed, but <i>must</i> de-allocate or return the memory when finished. Not returning memory can result in memory leaks and unstable operation on the part of Moab.</p>
<b>Multiple access</b>	<p>A given loaded plug-in can be used by more than one partition. This means that the plug-in must maintain its internal data in such a way that calls to the plug-in for the separate partitions do not conflict. It is recommended that internal data be allocated and a pointer to the data be kept in the data_handle described above as opposed to using global or static variables. Any global or static data will be shared between possible multiple instances of the plug-in.</p>

Runtime API	Description
<b>node_allocate</b>	<pre>int node_allocate (     void          *data_handle,     const char    *job_name,     int           container_count,     nalloc_container_t container[])</pre> <p>The plug-in must provide a node_allocate() entry point. This entry point is called each time Moab needs to determine where (on what nodes) a job will eventually run. Note that this entry point can be called many times before the job is actually scheduled to run.</p>

Runtime API	Description																																																																																
	<ul style="list-style-type: none"> <li>Data structures — Moab uses C data structures to pass information and lists of nodes to the plug-in and receive them back from the plug-in. See <code>moab-plugin.h</code> for the definitions of these structures and for information on how they relate to one another.</li> </ul>																																																																																
<p><b>Operations</b></p>	<p>A node allocation request consists of one or more requirements. Each of these requirements is provided within a “container” structure. The container has information regarding the requirement to be met, the count and list of all nodes that are available to meet the requirement and a place to return the list of nodes that the plug-in has chosen to use for the job.</p> <table border="1" data-bbox="383 667 1411 1604"> <thead> <tr> <th>Command</th> <th>Moab Job Task Count</th> <th>Job Node Count</th> <th>Job Tasks Per Node</th> <th>No of CFG Procs</th> <th>No of AVL Procs</th> <th>Plug-in Node Mapped TC</th> <th>requirement -&gt;taskcount</th> <th>return_node_count</th> </tr> </thead> <tbody> <tr> <td colspan="9" style="text-align: center;"><b>Non-ExactNode</b></td> </tr> <tr> <td>-l nodes=12</td> <td>12</td> <td>0</td> <td>0</td> <td>8</td> <td>8</td> <td>8</td> <td>12</td> <td>2</td> </tr> <tr> <td>-l nodes=12:ppn=2</td> <td>24</td> <td>0</td> <td>2</td> <td>8</td> <td>8</td> <td>8</td> <td>24</td> <td>3</td> </tr> <tr> <td colspan="9" style="text-align: center;"><b>ExactNode</b></td> </tr> <tr> <td>-l nodes=4</td> <td>4</td> <td>4</td> <td>0</td> <td>8</td> <td>8</td> <td>1</td> <td>4</td> <td>4</td> </tr> <tr> <td>-l nodes=4:ppn=2</td> <td>8</td> <td>4</td> <td>2</td> <td>8</td> <td>8</td> <td>2</td> <td>8</td> <td>4</td> </tr> <tr> <td>-l nodes=12</td> <td>12</td> <td>0</td> <td>0</td> <td>8</td> <td>6</td> <td>6</td> <td>12</td> <td>2</td> </tr> </tbody> </table> <p>The duty of the plug-in is to use the information that it has previously gathered (during the initialization) to select from the available nodes those that will best fulfill the requirements.</p> <p>The basic algorithm is to consume all the taskcount and memory on each node</p>									Command	Moab Job Task Count	Job Node Count	Job Tasks Per Node	No of CFG Procs	No of AVL Procs	Plug-in Node Mapped TC	requirement ->taskcount	return_node_count	<b>Non-ExactNode</b>									-l nodes=12	12	0	0	8	8	8	12	2	-l nodes=12:ppn=2	24	0	2	8	8	8	24	3	<b>ExactNode</b>									-l nodes=4	4	4	0	8	8	1	4	4	-l nodes=4:ppn=2	8	4	2	8	8	2	8	4	-l nodes=12	12	0	0	8	6	6	12	2
Command	Moab Job Task Count	Job Node Count	Job Tasks Per Node	No of CFG Procs	No of AVL Procs	Plug-in Node Mapped TC	requirement ->taskcount	return_node_count																																																																									
<b>Non-ExactNode</b>																																																																																	
-l nodes=12	12	0	0	8	8	8	12	2																																																																									
-l nodes=12:ppn=2	24	0	2	8	8	8	24	3																																																																									
<b>ExactNode</b>																																																																																	
-l nodes=4	4	4	0	8	8	1	4	4																																																																									
-l nodes=4:ppn=2	8	4	2	8	8	2	8	4																																																																									
-l nodes=12	12	0	0	8	6	6	12	2																																																																									

Runtime API	Description
	<p>until the consumed task count is greater than or equal to the container's task_count and memory requirements.</p> <p>A job's taskcount is calculated differently based on the JOBNODEMATCHPOLICY parameter. By default, it isn't defined and <code>-1 nodes=#</code> actually requests the number of tasks without respect to the number of nodes. In this case, the plug-in should consume all the tasks of each chosen node until the taskcount is greater and/or equal to the container's taskcount requirement. The plug-in is for node allocation and not task placement.</p> <p>When the JOBNODEMATCHPOLICY EXACTNODE is configured, then <code>-1 nodes=#</code> means the job wants # of nodes with 1 task per node. In this case, the nodes passed to the plug-in will have a taskcount that is mapped down to what the job can only use on that node. Each node's taskcount should be consumed on each node until the summed amount is equal to the container's requirement taskcount requirement.</p>
<b>Errors and return codes</b>	<p>The plug-in may internally log any errors encountered and must return a success or error status as defined in <code>moab-plugin.h</code>:</p> <pre>#define PLUGIN_RC_SUCCESS 0 #define PLUGIN_RC_FAILURE 1</pre>
<b>Multiple access safe</b>	The <code>node_allocate()</code> entry point must support multiple access as described above.

Unload Time API	Description
<b>finish()</b>	<pre>void finish(void *data_handle)</pre> <p>The plug-in must supply a <code>finish()</code> entry point. This entry point is called when Moab is preparing to disable and/or unload an instance of the plug-in.</p>
<b>Memory/resource cleanup</b>	The plug-in must de-allocate and free up any resources acquired either during the <code>initialize()</code> entry point or during any calls to the <code>node_allocate()</code> entry point. When the last entry point returns, there should be no allocated memory or other resources still in use by the plug-in instance.
<b>Multiple access safe</b>	The <code>finish()</code> entry point must support multiple access as described above.





## Appendix M: Scalable Systems Software Specification

In this section:

- [M.1 Scalable Systems Software Job Object Specification - page 1414](#)
- [M.2 Scalable Systems Software Resource Management and Accounting Protocol \(SSSRMAP\) Message Format - page 1450](#)
- [M.3 Scalable Systems Software Node Object Specification - page 1476](#)
- [M.4 Scalable Systems Software Resource Management and Accounting Protocol \(SSSRMAP\) Wire Protocol - page 1486](#)

### M.1 Scalable Systems Software Job Object Specification

SSS Job Object Specification  
Draft Release Version 3.1.0  
26 April 2011

Scott Jackson, PNNL  
David Jackson, Ames Lab  
Brett Bode, Ames Lab

#### Status of this Memo

This document describes the job object to be used by Scalable Systems Software compliant components. It is envisioned for this specification to be used in conjunction with the SSSRMAP protocol with the job object passed in the Data field of Requests and Responses. Queries can be issued to a job-cognizant component in the form of modified XPATH expressions to the Get field to extract specific information from the job object as described in the SSSRMAP protocol.

#### Abstract

This document describes the syntax and structure of the SSS job object. A job model is described that is flexible enough to support the specification of very simple jobs, as well as jobs with elaborate and complex specification requirements in a way that avoids complex

structures and syntax when it is not needed. The basic assumption is that a solitary job specification should be usable for all phases of the job lifecycle and can be used at submission, queuing, staging, reservations, quotations, execution, charging, accounting, etc. This job specification provides support for multi-step jobs, as well as jobs with disparate task descriptions. It accounts for operational requirements in a grid or meta-scheduled environment where the job is executed by multiple hosts in different administrative domains that support different resource management systems.

## Table of Contents

- [1.0 Introduction](#)
  - [1.1 Goals](#)
  - [1.2 Non-Goals](#)
  - [1.3 Examples](#)
    - [1.3.1 Very Simple Example](#)
    - [1.3.2 Moderate Example](#)
    - [1.3.3 Elaborate Example](#)
- [2.0 Conventions used in this document](#)
  - [2.1 Keywords](#)
  - [2.2 Table Column Interpretations](#)
  - [2.3 Element Syntax Cardinality](#)
- [3.0 The Job Model](#)
- [4.0 JobGroup Element](#)
  - [4.1 JobGroup Properties](#)
    - [4.1.1 Simple JobGroup Properties](#)
    - [4.1.2 Job](#)
    - [4.1.3 JobDefaults](#)
  - [4.2 JobGroup Reference](#)
- [5.0 Job and JobDefaults Element](#)

- 5.1 Job Properties
  - 5.1.1 Simple Job Properties
  - 5.1.2 Feature Element
  - 5.1.3 OutputFile Element
  - 5.1.4 ErrorFile Element
  - 5.1.5 InputFile Element
  - 5.1.6 NotificationList Element
  - 5.1.7 ResourceLimit Element
  - 5.1.8 Credentials
  - 5.1.9 Environment Element
    - 5.1.9.1 Variable Element
  - 5.1.10 Node Element
  - 5.1.11 TaskDistribution Element
  - 5.1.12 Dependency Element
  - 5.1.13 Consumable Resources
  - 5.1.14 Resource Element
  - 5.1.15 Extension Element
  - 5.1.16 TaskGroup
  - 5.1.17 TaskGroupDefaults
- 5.2 Job Reference
- 6.0 TaskGroup and TaskGroupDefaults Element
  - 6.1 TaskGroup Properties
    - 6.1.1 Simple TaskGroup Properties
    - 6.1.2 Task
    - 6.1.3 TaskDefaults
  - 6.2 TaskGroup Reference
- 7.0 Task and TaskDefaults Element

- 7.1 Task Properties
  - 7.1.1 Simple Task Properties
- 7.2 Task Reference
- 8.0 Property Categories
  - 8.1 Requested Element
  - 8.2 Delivered Element
- 9.0 AwarenessPolicy Attribute
- 10.0 References
- 11.0 Units of Measure Abbreviations

## 1.0 Introduction

This specification proposes a standard XML representation for a job object for use by the various components in the SSS Resource Management System. This object will be used in multiple contexts and by multiple components. It is anticipated that this object will be passed via the Data Element of SSSRMAP Requests and Responses.

## 1.1 Goals

There are several goals motivating the design of this representation.

The representation needs to be inherently flexible. We recognize we will not be able to exhaustively include the ever-changing job properties and capabilities that constantly arise.

The representation should use the same job object at all stages of that job's lifecycle. This object will be used at job submission, queuing, scheduling, charging and accounting, therefore it might need to distinguish between requested and delivered properties.

The design must account for the properties and structure required to function in a meta or grid environment. It needs to include the capability to support local mapping of properties, global namespaces, etc.

The equivalent of multi-step jobs must be supported. Each step (job) can have multiple logical task descriptions.

Many potential users of the specification will not be prepared to implement the complex portions or fine-granularity that others need. There needs to be a way to allow the more complicated structure to be added as needed while leaving more straightforward cases simple.

There needs to be guidance for how to understand a given job object when higher order features are not supported by an implementation, and which parts are required, recommended and optional for implementers to implement.

It needs to support composite resources.

It should include the ability to specify preferences or fuzzy requirements.

## 1.2 Non-Goals

Namespace considerations and naming conventions for most property values are outside of the scope of this document.

## 1.3 Examples

### *Example M-1: Very Simple Example*

This example shows a simple job object that captures the requirements of a simple job:

```
<Job>
  <Id>PBS.1234.0</Id>
  <State>Idle</State>
  <User>scottmo</User>
  <Executable>/bin/hostname</Executable>
  <Processors>16</Processors>
  <Duration>3600</Duration>
</Job>
```

### *Example M-2: Moderate Example*

This example shows a moderately complex job object that uses features such as required versus delivered properties:

```
<Job>
  <Id>PBS.1234.0</Id>
  <Name>Heavy Water</Name>
  <Project>nwchemdev</Project>
  <User>peterk</User>
  <Application>NWChem</Application>
  <Executable>/usr/local/nwchem/bin/nwchem</Executable>
  <Arguments>-input basis.in</Arguments>
  <InitialWorkingDirectory>/home/peterk</InitialWorkingDirectory>
  <Machine>Colony</Machine>
  <QualityOfService>BottomFeeder</QualityOfService>
  <Queue>batch_normal</Queue>
  <State>Completed</State>
  <StartTime>1051557713</StartTime>
  <EndTime>1051558868</EndTime>
  <Charge>25410</Charge>
  <Requested>
    <Processors op="GE">12</Processors>
    <Memory op="GE" units="GB">2</Memory>
  </Requested>
  <Duration>3600</Duration>
```

```

</Requested>
<Delivered>
  <Processors>16</Processors>
  <Memory metric="Average" units="GB">1.89</Memory>
  <Duration>1155</Duration>
</Delivered>
<Environment>
  <Variable name="PATH">/usr/bin:/home/peterk</Variable>
</Environment>
</Job>

```

**Example M-3: Elaborate Example**

This example uses a job group to encapsulate a multi-step job. It shows this protocol's ability to characterize complex job processing capabilities. A component that processes this message is free to retain only that part of the information that it requires. Superfluous information can be ignored by the component or filtered out (by XSLT for example).

```

<JobGroup>
  <Id>workflow1</Id>
  <State>Active</State>
  <Name>ShuttleTakeoff</Name>
  <JobDefaults>
    <StagedTime>1051557859</StagedTime>
    <SubmitHost>asteroid.lbl.gov</SubmitHost>
    <SubmitTime>1051556734</SubmitTime>
    <Project>GrandChallenge18</Project>
    <GlobalUser>C=US,O=LBL,CN=Keith Jackson</GlobalUser>
    <User>keith</User>
    <Environment>
      <Variable name="LD_LIBRARY_PATH">/usr/lib</Variable>
      <Variable name="PATH">/usr/bin:/bin</Variable>
    </Environment>
  </JobDefaults>
  <Job>
    <Id>fr15n05.1234.0</Id>
    <Name>Launch Vector Initialization</Name>
    <Executable>/usr/local/gridphys/bin/lvcalc</Executable>
    <Queue>batch</Queue>
    <State>Completed</State>
    <Machine>SMP2.emsl.pnl.gov</Machine>
    <StartTime>1051557713</StartTime>
    <EndTime>1051558868</EndTime>
    <Quote>https://www.pnl.gov/SMP2#654321</Quote>
    <Charge units="USD">12.75</Charge>
    <Requested>
      <Duration>3600</Duration>
      <Processors>2</Processors>
      <Memory>1024</Memory>
    </Requested>
    <Delivered>
      <Duration>1155</Duration>
      <Processors consumptionRate="0.78">2</Processors>
      <Memory metric="Max">975</Memory>
    </Delivered>
    <TaskGroup>
      <TaskCount>2</TaskCount>
      <TaskDistribution type="TasksPerNode">1</TaskDistribution>
      <Task>
        <Node>node1</Node>

```

```

    <Process>99353</Process>
  </Task>
</Task>
  <Node>node12</Node>
  <Process>80209</Process>
</Task>
</TaskGroup>
</Job>
<Job>
  <Id>fr15n05.1234.1</Id>
  <Name>3-Phase Ascension</Name>
  <Queue>batch_normal</Queue>
  <State>Idle</State>
  <Machine>Colony.emsl.pnl.gov</Machine>
  <Priority>1032847</Priority>
  <Hold>System</Hold>
  <StatusMessage>Insufficient funds to start job</StatusMessage>
  <Requested>
    <Duration>43200</Duration>
  </Requested>
  <TaskGroup>
    <TaskCount>1</TaskCount>
    <Name>Master</Name>
    <Executable>/usr/local/bin/stage-coordinator</Executable>
    <Memory>2048</Memory>
    <Resource name="License" type="ESSL2">1</Resource>
    <Feature>Jumbo-Frame</Feature>
  </TaskGroup>
  <TaskGroup>
    <Name>Slave</Name>
    <TaskDistribution type="Rule">RoundRobin</TaskDistribution>
    <Executable>/usr/local/bin/stage-slave</Executable>
    <NodeCount>4</NodeCount>
    <Requested>
      <Processors group="-1">12</Processors>
      <Processors conj="Or" group="1">16</Processors>
      <Memory>512</Memory>
      <Node aggregation="Pattern">fr15n.*</Node>
    </Requested>
  </TaskGroup>
</Job>
</JobGroup>

```

## 2.0 Conventions Used in this Document

### 2.1 Keywords

The keywords **MUST**, **MUST NOT**, **REQUIRED**, **SHALL**, **SHALL NOT**, **SHOULD**, **RECOMMENDED**, **MAY**, and **OPTIONAL** in this document are to be interpreted as described in [RFC2119](#).

### 2.2 Table Column Interpretations

The columns of the property tables in this document have the following meanings:



Element Name	Name of the XML element (xsd:element) see [DATATYPES]
<b>Type</b>	Data type defined by xsd (XML Schema Definition) as: <ul style="list-style-type: none"> <li>• String — xsd:string (a finite length sequence of printable characters)</li> <li>• Integer — xsd:integer (a signed finite length sequence of decimal digits)</li> <li>• Float — xsd:float (single-precision 32-bit floating point)</li> <li>• Boolean — xsd:boolean (consists of the literals 'true' or 'false')</li> <li>• DateTime — xsd:int (a 32-bit unsigned long in GMT seconds since the EPOCH)</li> <li>• Duration — xsd:int (a 32-bit unsigned long measured in seconds)</li> </ul>
<b>Description</b>	Brief description of the meaning of the property
<b>Appearance</b>	An indication of whether the given property must appear in the parent element. It assumes the following meanings: <ul style="list-style-type: none"> <li>• MUST — This property is REQUIRED when the parent is specified</li> <li>• SHOULD — This property is RECOMMENDED when the parent is specified.</li> <li>• MAY — This property is OPTIONAL when the parent is specified.</li> </ul>
<b>Compliance</b>	An indication of the relative importance of supporting the given property: <ul style="list-style-type: none"> <li>• MUST — A compliant implementation MUST support this property.</li> <li>• SHOULD — A compliant implementation SHOULD support this property.</li> <li>• MAY — A compliant implementation MAY support this property.</li> </ul>
<b>Categories</b>	Some properties can be categorized into one of several categories. Letters in this column indicate that the given property can be classified in the following property categories. <ul style="list-style-type: none"> <li>• R — This property can be encompassed in a Requested element.</li> <li>• D — This property can be encompassed in a Delivered element.</li> </ul>

### 2.3 Element Syntax Cardinality

Selected elements in the element syntax sections use regular expression wildcards with the following meanings:

Wildcard	Description
*	Zero or more occurrences

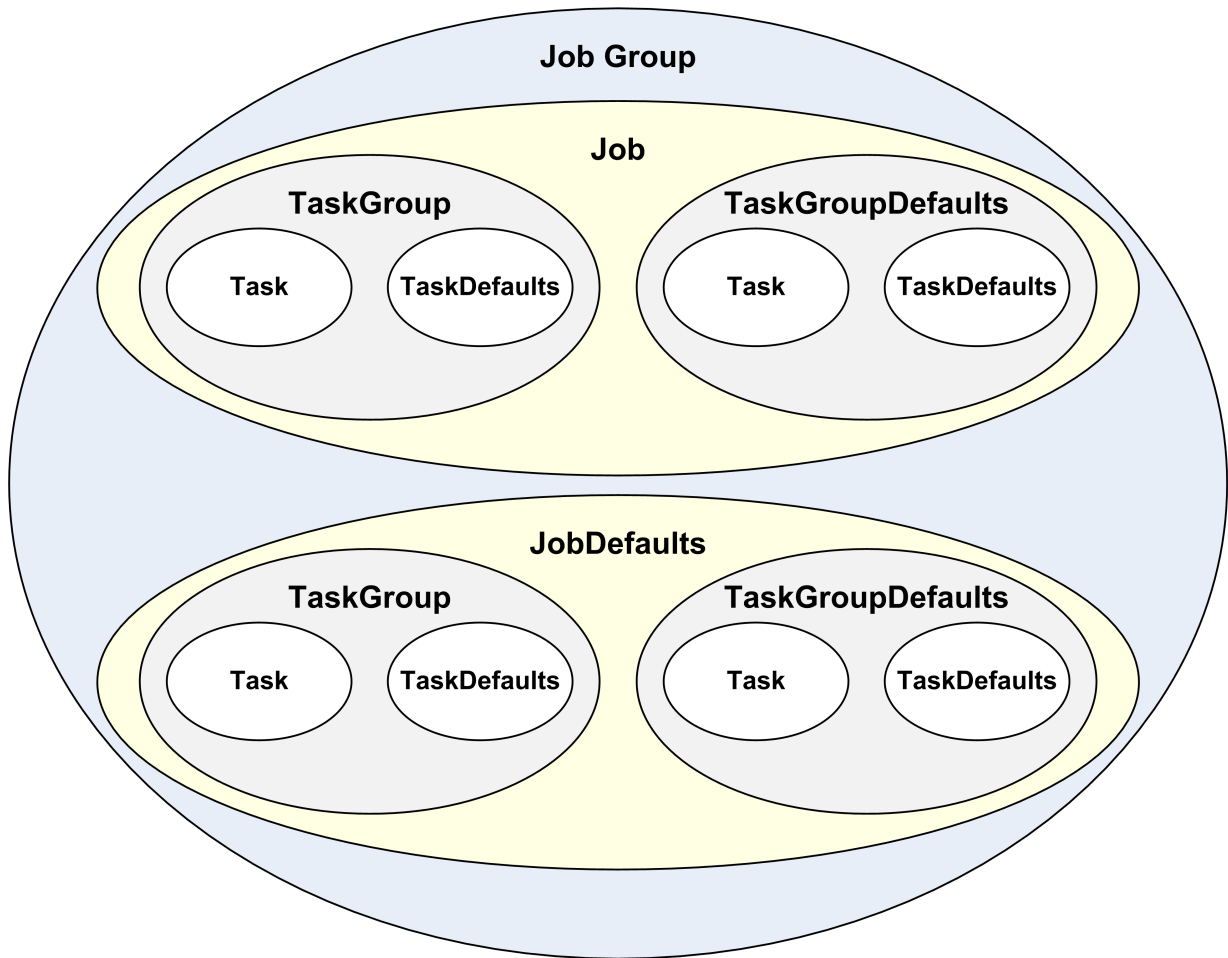
Wildcard	Description
+	One or more occurrences
?	Zero or one occurrences

The absence of one of these symbols implies exactly one occurrence.

### 3.0 The Job Model

The primary object within the job model is a job. A job can be thought of as a single schedulable entity and will be the object normally seen in job queues.

*Image M-1: JobGroup contains Job and JobDefaults, which contain TaskGroup and TaskGroupDefaults*



Jobs with dependencies on other jobs can be submitted in a job group. Jobs within a job group form a DAG (directed acyclic graph) where the nodes are jobs and the edges represent dependencies on the status of previous jobs. A job group will consist of at least one job. A job group can optionally specify job defaults, which are a set of job properties to be assumed by all jobs within the job group unless overridden within the job.

A job can consist of multiple tasks, which are the finest grained work unit and represent an endpoint for executing a given process instance. For example, a job that requests 3 nodes and 4 processors will have 4 tasks, two on one node and one on each of two nodes. Tasks can be grouped into task groups, which are logical aggregations of tasks and their common properties. Submit filters, prologs, epilogs, notification scripts, etc. run once only for each job. Whereas task groups function as logical descriptions of tasks and their properties, they also describe the number of such tasks and the nodes that they run on. As an example, a master task group (consisting of a single task) might ask for a node with a MATLAB license, 2GB of memory and an Internet connected network adapter while a slave task group (consisting of 12 tasks) could be targeted for nodes with more CPU bandwidth -- all within the same job and utilizing a common MPI ring. Tasks (and therefore taskgroups) can have different executables or environments, specify different consumable resources or node properties. A job, therefore, can specify one or more task group. A job that does not specify an explicit task group is considered as having a single implicit task group. A job can optionally specify task group defaults, which are a set of task group properties to be assumed by all task groups within the job unless overridden within a task group.

A task group can specify one or more tasks. A task group that does not specify an explicit task is considered as having a single implicit task. A task group can optionally specify task defaults, which are a set of task properties to be assumed by all tasks within the task group unless overridden within a task.

## 4.0 JobGroup Element

A JobGroup is an optional element that aggregates one or more interdependent jobs. Some resource managers support the submission of job groups (multi-step jobs) and queries on the status of an entire job group.

- A compliant implementation MAY support this element.
- A JobGroup MUST specify one or more JobGroup Properties.
- A JobGroup MUST contain one or more Jobs.
- A JobGroup MAY contain zero or more JobsDefaults.

The following illustrates this element's syntax:

```
<JobGroup>
  <!-- JobGroup Properties -->+
  <Job/>+
```

```
<JobDefaults/>?
</JobGroup>
```

## 4.1 JobGroup Properties

JobGroup Properties are properties that apply to the job group as a whole. These include the job group ID, jobs and job defaults, and other simple optional job properties.

### Simple JobGroup Properties

Simple (unstructured) job group properties are enumerated in the table below:

*Table M-1: Simple JobGroup Properties*

Element Name	Type	Description	Appearance	Compliance
<b>CreationTime</b>	DateTime	Date and time that the job group was instantiated	MAY	MAY
<b>Description</b>	String	Description of the job group	MAY	MAY
<b>Id</b>	String	Job group identifier	MUST	MUST
<b>Name</b>	String	Name of the job group	MAY	SHOULD
<b>State</b>	String	State of the job group as a whole. Valid states can include NotQueued, Unstarted, Active, and Completed.	MAY	SHOULD

### Job

A job group **MUST** contain one or more jobs.

See the next section for element details.

### JobDefaults

A job group **MAY** contain zero or one job defaults.

See the next section for element details.

## 4.2 JobGroup Reference

When a simple reference to a predefined job group is needed in an encapsulating element, a JobGroup element is used with the text content being the job group ID:

```
<JobGroup> workflow1</JobGroup>
```

## 5.0 Job and JobDefaults Element

The Job and JobDefaults elements are of the same structure. A Job element encapsulates a job and can be expressed as a standalone object. A JobDefaults element can only appear within a JobGroup and represents the defaults to be taken by all jobs within the job group. Job properties in Job elements override any properties found in a sibling JobDefaults element.

- A compliant implementation **MUST** support the Job element.
- A compliant implementation **MAY** support the JobDefaults element only if it supports the JobGroup element.
- A job **MUST** specify one or more Job Properties.
- One or more TaskGroup elements **MAY** appear at this level.
- Zero or one TaskGroupDefaults elements **MAY** appear at this level.

The following illustrates this element's syntax:

```
<Job>
  <!-- Job Properties -->+
  <TaskGroup/>*
  <TaskGroupDefaults/>?
</Job>
```

## 5.1 Job Properties

Job Properties apply to a particular job or as default properties to all jobs. They include the job ID, job credentials, task groups, task group defaults, and other simple optional properties.

### Simple Job Properties

Simple (unstructured) job properties are enumerated in the table below:

**Table M-2: Simple Job Properties**

Element Name	Type	Description	Appearance	Compliance	Categories
<b>Application</b>	String	Type of application such as Gaussian or Nwchem	MAY	MAY	
<b>Architecture</b>	String	Type architecture for the nodes on which this job must run	MAY	MAY	RD
<b>Arguments</b>	String	The arguments for the executable	MAY	SHOULD	
<b>Charge</b>	Float	The amount charged for the job	MAY	SHOULD	
<b>Checkpointable</b>	Boolean	Can this job be checkpointed?	MAY	MAY	
<b>CpuDuration</b>	Duration	Number of cpu seconds used by the job	MAY	SHOULD	
<b>DeadlineTime</b>	DateTime	Date and time that a job must end by	MAY	MAY	
<b>EligibleTime</b>	DateTime	Date and time that a job must start after	MAY	MAY	

Element Name	Type	Description	Appearance	Compliance	Categories
<b>EndTime</b>	DateTime	Date and time that a job ended (independent of success or failure)	MAY	MUST	
<b>Executable</b>	String	Executable. This can be an absolute or relative path or a URI.*	MAY	MUST	
<b>ExitCode</b>	Integer	Exit code for the job	MAY	SHOULD	
<b>GlobalJob</b>	String	Globally unique job identifier (possibly in the form of a URI)	MAY	SHOULD	
<b>Hold</b>	String	Hold(s) on the job. There may be multiple instances of this element if there is more than one ld on the job	MAY	SHOULD	
<b>InitialWorking-Directory</b>	String	Initial working directory	MAY	SHOULD	
<b>Interactive</b>	Boolean	Is this an interactive job?	MAY	SHOULD	

Element Name	Type	Description	Appearance	Compliance	Categories
<b>Id</b>	String	A local job identifier assigned to the job by the local resource manager	MUST	MUST	
<b>Name</b>	String	Name of the job	MAY	SHOULD	
<b>State</b>	String	State of the job. Valid states can include Idle, Hold, Running, Suspended, or Completed	MAY	MUST	
<b>Type</b>	String	Type of job. Meaning of this extension property is context specific.	MAY	MAY	
<b>Machine</b>	String	Name of the system or cluster that runs the job	MAY	MUST	RD
<b>Network</b>	String	Type of network adapter required by the job	MAY	MAY	RD
<b>NodeCount</b>	Integer	Number of nodes used by the job	MAY	MUST	RD



Element Name	Type	Description	Appearance	Compliance	Categories
<b>OperatingSystem</b>	String	Operating System required by the job	MAY	MAY	RD
<b>Partition</b>	String	Name of the partition where the job should run	MAY	MAY	RD
<b>Priority</b>	Integer	Current queue priority (or rank) for the job	MAY	SHOULD	
<b>QualityOfService</b>	String	Name of the Quality of Service (QoS)	MAY	SHOULD	RD
<b>Queue</b>	String	Name of the Queue (or class) that the job runs in	MAY	SHOULD	RD
<b>Quote</b>	String	Identifier for a guaranteed charge rate quote obtained by the job	MAY	MAY	
<b>Reservation</b>	String	Identifier for a reservation used by the job	MAY	MAY	RD
<b>ReservationTime</b>	DateTime	Date and time that a reservation	MAY	MAY	

Element Name	Type	Description	Appearance	Compliance	Categories
		was placed for the job			
<b>ResourceManager Type</b>	String	Type of resource manager required to run this job	MAY	MAY	RD
<b>Restartable</b>	Boolean	Can this job be restarted?	MAY	MAY	
<b>Shell</b>	String	Specified the shell necessary to interpret the job script	MAY	MAY	
<b>StagedTime</b>	DateTime	Date and time that a job was staged to the local resource management system	MAY	MAY	
<b>StartCount</b>	Integer	Number of times the scheduler tried to start the job	MAY	MAY	
<b>StartTime</b>	DateTime	Date and time that the job started	MAY	MUST	
<b>StatusMessage</b>	String	Natural language message that can be used to provide detail on	MAY	SHOULD	

Element Name	Type	Description	Appearance	Compliance	Categories
		why a job failed, isn't running, etc.			
<b>SubmitTime</b>	DateTime	Date and time that a job was submitted	MAY	SHOULD	
<b>SubmitHost</b>	String	FQDN of host where the job was submitted from	MAY	SHOULD	
<b>Suspendable</b>	Boolean	Can this job be suspended?	MAY	MAY	
<b>SuspendDuration</b>	Integer	Number of seconds the job was in the Suspended state	MAY	MAY	
<b>TimeCategory</b>	String	This allows the specification of shifts like PrimeTime for charging purposes	MAY	MAY	
<b>Duration</b>	Duration	Number of seconds in the Running state	SHOULD	MUST	RD

\* The Executable can be a script or a binary executable. If it is already on the target system it can be referenced by an absolute or relative pathname (relative to InitialWorkingDirectory). If it is passed with the job in a File object (see SSSRMAP), it can be referenced by an absolute or relative URI. An absolute URI would specify a URL where

the file can be downloaded (like with `wget`). A relative URI is specified by preceding an identifier by a pound sign, as in:

```
<Executable>#Script</Executable>
```

It will be found in a `File` object included along with the `Job` object with the `Script` as an identifier, as in:

```
<File id="Script">echo hello world</File>
```

## Feature Element

The `Feature` element connotes an arbitrary named feature of a node:

- A compliant implementation **SHOULD** support this element.
- This element **MAY** appear zero or one times within a given set of `Job Properties`.
- This element is of type `String`.
- This element **MAY** have an `aggregation` attribute of type `String` that provides a way to indicate multiple values with a single expression. A compliant implementation **MAY** support the `aggregation` attribute if the `Feature` element is supported. Possible values for this attribute include:
  - `List` — a comma-separated list of features
  - `Pattern` — a regular expression (perl5) matching desired features
- If an `aggregation` attribute is specified with the value of `List`, this element **MAY** also have a `delimiter` attribute of type `String` that indicates what delimiter is used to separate list elements. The default list delimiter is a comma.
- This element **MAY** be categorized as a requested or delivered property by being encompassed by the appropriate element.

The following is an example of a feature element:

```
<Feature aggregation="List">feature1,feature2</Feature>
```

## OutputFile Element

The `OutputFile` element specifies the name of the file to which the output stream (`stdout`) from the job will be written:

- This element's character content is the name of the file. If this element is omitted or it is empty, then an appropriate output file is auto-determined by the queuing system.
- This element **MAY** have a `redirectList` attribute, which is a comma-separated list of output redirection attributes of type `String`. A compliant implementation **SHOULD** support this attribute if `OutputFile` is supported. Possible values for this

attribute include:

- Append — opens the output file for append
- Close — closes and discards the output stream
- Flush — output is written to output file as it is generated
- Keep — leave the output file on the execution host
- Merge — merges the output stream into the error stream

Note that when using the `redirectList` attributes, the cumulative effect of the `ErrorFile` and `OutputFile` directives may be order dependent.

The following is an example of an `OutputFile` element:

```
<OutputFile redirectList="Append">~/myjob.out</OutputFile>
```

## ErrorFile Element

The `ErrorFile` element specifies the name of the file to which the error stream (`stderr`) from the job will be written:

- This element's character content is the name of the file. If this element is omitted or it is empty, then an appropriate error file is auto-determined by the queuing system.
- This element MAY have a `redirectList` attribute, which is a comma-separated list of error redirection attributes of type `String`. A compliant implementation SHOULD support this attribute if `ErrorFile` is supported. Possible values for this attribute include:
  - Close — closes and discards the error stream
  - Append — opens the error file for append
  - Flush — output is written to output file as it is generated
  - Keep — leave the output file on the execution host
  - Merge — merges the error stream into the output stream

Note that when using the `redirectList` attributes, the cumulative effect of the `ErrorFile` and `OutputFile` directives may be order dependent.

The following is an example of an `ErrorFile` element:

```
<ErrorFile redirectList="Merge"></ErrorFile>
```

## InputFile Element

The `InputFile` element specifies the name of the file from which the input stream (`stdin`) for the job will be read:

- This element's character content is the name of the file. If this element is omitted or it is empty, then an appropriate input file is auto-determined by the queuing system.
- This element MAY have a `redirectList` attribute, which is a comma-separated list of input attributes of type `String`. A compliant implementation SHOULD support this attribute if `InputFile` is supported. Possible values for this attribute include:
  - `Close` — closes and discards the input stream

The following is an example of an `InputFile` element:

```
<InputFile redirectList="Close"></InputFile>
```

## NotificationList Element

The `NotificationList` element specifies the job-related events or conditions for which a notification will be sent.

- This element's character content is a comma-separated list of events or conditions for which a notification should be sent. Possible values for the elements of this list include:
  - `JobStart` — send a notification when the job starts
  - `JobEnd` — send a notification when the job ends
  - `All` — send notifications for all notifiable events
  - `None` — do not send notifications for any events
- This element MAY have a `uri` attribute of type `String`, which indicates where the notification is to be sent. A compliant implementation MAY support this attribute if `NotificationList` is supported. The `uri` is in the format: `[scheme://]authority` with the scheme being `smtp` and the authority being an email address by default.

The following is an example of a `NotificationList` element:

```
<NotificationList uri="smith@business.com">JobStart,JobEnd</NotificationList>
```

## ResourceLimitElement

The `ResourceLimit` element represents a resource limit with its name and value:

- This element MUST have a `name` attribute of type `String`. A compliant implementation MUST support the `name` attribute if `ResourceLimit` is supported.
- This element MAY have a `type` attribute of type `String` that can have the values `Hard` or `Soft`. If the limit is enforced by the operating system, a hard limit is one that cannot be increased once it is set while a soft limit can be increased up to the

value of the hard limit. If the type attribute is omitted, both the soft and hard limits are set.

- This element’s character content is the resource limit’s value.

Some typical names include:

Name	Description
<b>CoreFileSize</b>	Maximum core file size
<b>CpuTime</b>	CPU time in seconds
<b>DataSegSize</b>	Maximum data size
<b>FileSize</b>	Maximum file size
<b>MaxMemorySize</b>	Maximum resident set size
<b>MaxProcesses</b>	Maximum number of processes
<b>MaxSwap</b>	Virtual memory limit
<b>MaxMemLock</b>	Maximum locked-in-memory address space
<b>MaxProcessors</b>	Maximum processors
<b>MaxMemory</b>	Maximum memory
<b>MaxDisk</b>	Maximum disk space
<b>MaxNetwork</b>	Maximum network bandwidth
<b>MaxFileIO</b>	Maximum file i/o
<b>OpenFiles</b>	Maximum number of open files
<b>Stacksize</b>	Maximum stack size

The following is an example of a ResourceLimit element:

```
<ResourceLimit name="CPUTime">1000000</ResourceLimit>
```

## Credentials

Credentials are a special group of job properties that characterize an authenticated token or ID. They can be categorized in both requested and delivered forms.

Credential job properties are enumerated in the table below:

*Table M-3: Credential Job Properties*

Element Name	Type	Description	Appearance	Compliance	Categories
<b>Project</b>	String	Name of the Project or Charge Account	MAY	SHOULD	RD
<b>GlobalUser</b>	String	Globally unique user identifier. This could be an X.509 DN for example	MAY	SHOULD	RD
<b>Group</b>	String	Name of the local group ID	MAY	MAY	RD
<b>User</b>	String	Name of the local user ID for the job	MAY	MUST	RD

## Environment Element

The Environment element encapsulates environment variables:

- This element MAY have an export attribute of type Boolean, which if set to `True`, indicates that all environment variables in the context of the job submission process should be exported in the job's execution environment.
- A compliant implementation SHOULD support this element.
- An Environment element MAY appear zero or one times within a given set of Job (or TaskGroup) Properties.
- An Environment element MAY contain one or more Variable elements.

The following illustrates this element's syntax:

```
<Environment>
  <Variable/>+
</Environment>
```

## Variable Element

The Variable element represents an environment variable with its name and value.



This element **MUST** have a `name` attribute of type `String`. A compliant implementation **MUST** support the `name` attribute if `Variable` is supported. This element's character content is the environment variable's value.

The following is an example of a `Variable` element:

```
<Variable name="PATH">/usr/bin:/home/sssdemo</Variable>
```

## Node Element

The `Node` element represents a node:

- A compliant implementation **SHOULD** support this element.
- This element **MAY** appear zero or one times within a given set of Job Properties.
- This element is of type `String`.
- This element **MAY** have an `aggregation` attribute of type `String` that provides a way to indicate multiple values with a single expression. A compliant implementation **MAY** support the `aggregation` attribute if the `Feature` element is supported. Possible values for this attribute include:
  - `List` - a comma-separated list of features
  - `Pattern` - a regular expression (perl5) matching desired features
  - `Range` - a range of nodes of the form: `<prefix>[5-23,77]`
- If an `aggregation` attribute is specified with the value of `List`, this element **MAY** also have a `delimiter` attribute of type `String` that indicates what delimiter is used to separate list elements. The default list delimiter is a comma.
- This element **MAY** have a `count` attribute of type `Integer` that indicates the instance count of the specified node(s).
- This element **MAY** be categorized as a requested or delivered property by being encompassed by the appropriate element.

The following is an example of a `Node` element:

```
<Node aggregation="Pattern">node[1-5]</Node>
```

## TaskDistribution Element

The `TaskDistribution` element describes how tasks are to be mapped to nodes. This mapping can be expressed as a rule name, a task per node ratio or an arbitrary geometry.

- A compliant implementation **SHOULD** support this element.
- This element **MAY** appear zero or one times in a given set of Job (or `TaskGroup`) Properties.

- This element is of type String.
- This element MAY have a `type` attribute of type String that provides a hint as to the type of mapping guidance provided. It can have values including `Rule`, `TasksPerNode`, `ProcessorsPerTask` or `Geometry`. A compliant implementation MAY support the `type` attribute if the `TaskDistribution` element is supported.
- It is possible to use `Processors`, `NodeCount` and `TaskCount` elements to specify a set of mutually contradictory task parameters. When this occurs, components are responsible for resolving conflicting requirements.

The following are three examples of a `TaskDistribution` element:

```
<TaskDistribution type="TasksPerNode">2</TaskDistribution>
<TaskDistribution type="Rule">RoundRobin</TaskDistribution>
<TaskDistribution type="Geometry">{1,4}{2}{3,5}</TaskDistribution>
```

## Dependency Element

The `Dependency` element allows a job's execution to depend on the status of other jobs. In a job group (multi-step job), some jobs may delay execution until the failure or success of other jobs creating in general a Directed Acyclic Graph relationship between the jobs. This element's content is of type String and represents the job that the current job is dependent upon. Since a job can have two or more dependencies, this element can appear more than once in a given job scope. A compliant implementation SHOULD support this element if job groups are supported:

- A compliant implementation SHOULD support this element.
- This element MAY appear zero or more times in a given set of Job (or TaskGroup) Properties.
- This element is of type String and contains the `JobId` that the current job is dependent upon.
- This element MAY have a `condition` attribute of type String that indicates the basis for determining when the current job executes in relation to the specified job. A compliant implementation MUST support this attribute if this element is supported. Possible values for this attribute include:
  - `OnSuccess` this job should run after the referenced job only if it completes successfully (this is the default if the `type` attribute is omitted)
  - `OnFailure` this job should run after the referenced job only if it fails
  - `OnExit` this job should run after the referenced job exits
- If the `condition` attribute is equal to `OnExit`, this element MAY have a `code` attribute of type Integer that indicates the exit code that will trigger this job to run. If

the code attribute is omitted, then the current job should run after the referenced job for any exit status.

- This element MAY have a `designator` attribute of type String that indicates that indicates the property of the job that identifies it as the dependent job. A compliant implementation MAY support this attribute if this element is supported.

Possible values for this attribute include:

- `JobId` the job this job is dependent upon is specified by `JobId` (this is the default if the `designator` attribute is omitted)
- `JobName` the job(s) this job is dependent upon are specified by `JobName`

The following is an example of a `Dependency` element:

```
<Dependency condition="OnSuccess" designator="JobId">PBS.1234.0</Dependency>
```

## Consumable Resources

Consumable Resources are a special group of properties that can have additional attributes and can be used in multiple contexts. In general a consumable resource is a resource that can be consumed in a measurable quantity.

- A consumable resource MAY have a `context` attribute of type String that indicates the sense in which the resource is used. A compliant implementation MAY support this attribute. Possible values for this attribute include:
  - `Configured` — run this task only on nodes having the specified configured resources
  - `Available` — run this task only on nodes having the specified available resources. (this is the default if the `context` attribute is omitted)
  - `Used` — the task used the indicated resources (this is analogous to being including in a `Delivered` block)
  - `Dedicated` — the indicated amount of the resource should be dedicated to the task
- A consumable resource MAY have a `units` attribute that is of type String that specifies the units by which it is being measured. If this attribute is omitted, a default unit is implied. A compliant implementation MAY support this attribute if the element is supported.
- A consumable resource MAY have a `metric` attribute that is of type String that specifies the type of measurement being described. For example, the measurement can be a `Total`, an `Average`, a `Min` or a `Max`. A compliant implementation MAY support this attribute if the element is supported.

- A consumable resource MAY have a `duration` attribute of type `Duration` that indicates the amount of time for which that resource was used. This need only be specified if the resource was used for a different amount of time than the duration for the job. A compliant implementation MAY support this attribute if the element is supported.
- A consumable resource MAY have a `consumptionRate` attribute of type `Float` that indicates the average percentage that a resource was used over its duration. For example, an overbooked SMP running 100 jobs across 32 processors might want to scale the usage and charge by the average fraction of processor usage actually delivered. A compliant implementation MAY support this attribute if the element is supported.
- A consumable resource MAY have a `dynamic` attribute of type `Boolean` that indicates whether the resource allocated for this job should be allowed to grow or shrink dynamically. For example, if `processors` is specified with `dynamic` equal to `True`, the job can be dynamically allocated more processors as they become available. The growth bounds can be indicated via the `op` attribute, which is inherited when a consumable resource element is encapsulated within a *Requested* element. A compliant implementation MAY support this attribute if the element is supported.

Simple consumable resources are listed in the table below:

*Table M-4: Simple Consumable Resources*

Element Name	Type	Description	Appearance	Compliance	Categories
<b>Disk</b>	Float	Amount of disk	MAY	SHOULD	RD
<b>Memory</b>	Float	Amount of memory	MAY	SHOULD	RD
<b>Network</b>	Float	Amount of network	MAY	MAY	RD
<b>Processors</b>	Integer	Number of processors	MAY	MUST	RD
<b>Swap</b>	Float	Amount of virtual memory	MAY	MAY	RD

The following are two examples for specifying a consumable resource:

```
<Memory metric="Max" units="GB">483</Memory>
<Processors duration="1234" consumptionRate="0.63">4</Processors>
```

## Resource Element

In addition to the consumable resources enumerated in the above table, an extensible consumable resource is defined by the Resource element:

- A compliant implementation SHOULD support this element.
- This element MAY appear zero or more times within a given set of job (or task group) properties.
- Like the other consumable resources, this property MAY be categorized as a requested or delivered property by being encompassed in the appropriate element.
- This element is of type Float.
- This element shares the *same properties and attributes as the other consumable resources* but it requires an additional `name` (and optional `type`) attribute to describe it.
- It MUST have a `name` attribute of type String that indicates the type of consumable resource being measured. A compliant implementation MUST support this attribute if the element is supported.
- It MAY have a `type` attribute of type String that distinguishes it within a general resource class. A compliant implementation SHOULD support this attribute if the element is supported.

The following are two examples for specifying a Resource element:

```
<Resource name="License" type="MATLAB">1</Resource>
<Resource name="Telescope" type="Zoom2000" duration="750" metric="KX">10</Resource>
```

## Extension Element

The Extension element provides a means to pass extensible properties with the job object.

Some applications might find it easier to use a named extension property than discover and handle elements they do not understand or anticipate by name:

- A compliant implementation MAY support this element.
- This element MUST have a `name` attribute of type String that gives the extension property's name. A compliant implementation MUST support this attribute if this element is supported.
- This element MAY have a `type` attribute of type String that characterizes the context within which the property should be understood. A compliant implementation SHOULD support this attribute if this element is supported.
- This element's character content, which is of type String, is the extension property's value.

The following is an example of an Extension element:

```
<Extension type="Scheduler" name="Restartable">true</Extension>
```

## TaskGroup

A job MAY specify one or more task groups.

See the next section for element details.

## TaskGroupDefaults

A job MAY specify zero or more task group defaults.

See the next section for element details.

## 5.2 Job Reference

When a simple reference to a predefined job is needed in an encapsulating element, a Job element is used with the text content being the job ID:

```
<Job> job123</Job>
```

## 6.0 TaskGroup and TaskGroupDefaults Element

The `TaskGroup` and `TaskGroupDefaults` elements have the same structure. A `TaskGroup` element aggregates tasks. A `TaskGroupDefaults` element can only appear within a Job (or JobDefaults) and represents the defaults to be taken by all task groups within the job. Task group properties in `TaskGroup` elements override any properties found in a sibling `TaskGroupDefaults` element.

- A compliant implementation MAY support the `TaskGroup` element.
- A compliant implementation MAY support the `TaskGroupDefaults` element.
- A task group MUST specify one or more TaskGroup Properties.
- One or more Task elements MAY appear at this level.
- Zero or one TaskDefaults elements MAY appear at this level.

The following illustrates this element's syntax:

```
<TaskGroup>
  <!-- TaskGroup Properties -->+
  <!-- Job Properties -->*
  <Task>+
  <TaskDefaults>?
</TaskGroup>
```

## 6.1 TaskGroup Properties

TaskGroup Properties apply to a particular task group or as default properties to encompassed task groups. These properties include the task group ID, its tasks, task defaults, and other simple task group properties.

### Simple TaskGroup Properties

Simple (unstructured) task group properties are enumerated in the table below:

*Table M-5: Simple TaskGroup Properties*

Element Name	Type	Description	Appearance	Compliance
<b>TaskCount</b>	Integer	Number of tasks in this taskgroup	MAY	MUST
<b>Id</b>	String	A task group identifier unique within the job	MAY	MAY
<b>Name</b>	String	A task group name (such as Master)	MAY	SHOULD

### Task

A task group MAY specify zero or more tasks.

See the next section for element details.

### TaskDefaults

A task group MAY specify zero or more task defaults.

See the next section for element details.

## 6.2 TaskGroup Reference

When a simple reference to a predefined task group is needed in an encapsulating element, a TaskGroup element is used with the text content being the task group ID:

```
<TaskGroup> tgl</TaskGroup>
```

## 7.0 Task and TaskDefaults Element

The `Task` and `TaskDefaults` elements have the same structure. A `Task` element contains information specific to a task (like the process ID or the host it ran on). A `TaskDefaults` element can only appear within a `TaskGroup` (or `TaskGroupDefaults`) element and represents the defaults for all tasks within the task group. Task properties in `Task` elements override any properties found in a sibling `TaskDefaults` element.

- A compliant implementation MAY support the `TaskGroup` element.
- A compliant implementation MAY support the `TaskGroupDefaults` element.
- A task group MUST specify one or more `TaskGroup Properties`.
- One or more `Task` elements MAY appear at this level.
- Zero or one `TaskDefaults` elements MAY appear at this level.

The following illustrates this element's syntax:

```
<Task>
  <!-- Task Properties -->+
  <!-- Job Properties -->*
</Task>
```

## 7.1 Task Properties

Task Properties are properties that apply to a particular task or as default properties to encompassed tasks. These properties include the task ID and other task properties.

### Simple Task Properties

Simple (unstructured) task properties are enumerated in the table below:

*Table M-6: Simple Task Properties*

Element Name	Type	Description	Appearance	Compliance
<b>Node</b>	String	Name of the node this task ran on	MAY	MUST
<b>Session</b>	Integer	Session ID for the task group or job	MAY	MAY
<b>Id</b>	String	A task identifier unique within the taskgroup	MAY	MAY



## 7.2 Task Reference

When a simple reference to a predefined task is needed in an encapsulating element, a `Task` element is used with the text content being the task ID:

```
<Task>1</Task>
```

## 8.0 Property Categories

Certain properties need to be classified as being in a particular category. This is done when it is necessary to distinguish between a property that is requested and a property that was delivered. When no such distinction is necessary, it is recommended that the property not be enveloped in one of these elements. In general, a property should be enveloped in a category element only if it is expected that the property will need to be attributed to more than one property category, or if it needs to make use of some of the special attributes inherited from the category.

### 8.1 Requested Element

A requested property reflects properties as they were requested. A disparity might occur between the requested value and the value delivered if a preference was expressed, if multiple options were specified, or if ranges or pattern matching was specified.

- A compliant implementation SHOULD support this element.

The following illustrates the syntax of this element:

```
<Requested>
  <!-- Requested Properties -->+
</Requested>
```

The following describes the attributes and elements for the example above:

```
/Requested
```

This element is used to encapsulate requested properties:

```
/Requested/<Requested Property>
```

Requested properties appear at this level.

Requested Properties inherit some additional attributes:

- A requested property MAY have an `op` attribute of type `String` that indicates a conditional operation on the value. A compliant implementation SHOULD support this attribute. Valid values for the `op` attribute include `EQ` meaning equals (which is the default), `NE` meaning not equal, `LT` meaning less than, `GT` meaning greater than, `LE`

meaning less than or equal to, GE meaning greater than or equal to, Match that implies the value is a pattern to be matched.

- A requested property MAY have a `conj` attribute of type String that indicates a conjunctive relationship with the previous element. A compliant implementation MAY support this attribute. Valid values for the `conj` attribute include And (which is the default), Or, Nand meaning and not, and Nor meaning or not.
- A requested property MAY have a `group` attribute of type Integer that indicates expression grouping and operator precedence much like parenthetical groupings. A compliant implementation MAY support this attribute. A positive grouping indicates the number of nested expressions being opened with the property while a negative grouping indicates the number of nested expressions being closed with the property.
- A requested property MAY have a `preference` attribute of type Integer that indicates a preference for the property along with a weight (the weights are taken as a ratio to the sum of all weights in the same group). A compliant implementation MAY support this attribute. If a group of positive valued preference alternatives are specified, at least one of the preferences must be satisfied for the job to run. If a group of negative valued preferences are specified, the preferences will try to be met according to their weights but the job will still run even if it can't satisfy any of the preferred properties. (Weight ranking can be removed by making all weights the same value (1 or -1 for example).
- A requested property MAY have a `performanceFactor` attribute of type Float that provides a hint to the scheduler of what performance tradeoffs to make in terms of resources and start time. A compliant implementation MAY support this attribute.

The following are four examples of using Requested Properties:

```

<Requested>
  <Processors op="GE">8</Processors>
  <Processors op="LE">16</Processors>
  <Duration>3600</Duration>
</Requested>
<Requested>
  <NodeCount>1</NodeCount>
  <Node aggregation="Pattern">fr15.*</Node>
</Requested>
<Requested>
  <User group="1">scottmo</User>
  <Account group="-1">mscfops</Account>
  <User conj="Or" group="1">amy</User>
  <Account group="-1">chemistry</Account>
</Requested>
<Requested>
  <Memory preference="2">1024</Memory>
  <Memory preference="1">512</Memory>
</Requested>

```

## 8.2 Delivered Element

A delivered property reflects properties as they were actually utilized, realized or consumed. It reflects the actual amounts or values that are used, as opposed to a limit, choice or pattern as may be the case with a requested property.

- A compliant implementation SHOULD support this element.

The following illustrates the syntax of this element:

```
<Delivered>
  <!-- Delivered Properties -->+
</Delivered>
```

The following describes the attributes and elements for the example above:

```
/Delivered
```

This element is used to encapsulate delivered properties:

```
/Delivered/<Delivered Property>
```

Delivered properties appear at this level.

Delivered Properties inherit some additional attributes:

- A delivered property MAY have a `group` attribute of type Integer that indicates expression grouping and operator precedence much like parenthetical groupings. A compliant implementation MAY support this attribute. A positive grouping indicates the number of nested expressions being opened with the property while a negative grouping indicates the number of nested expressions being closed with the property. The purpose of this attribute would be to logically group delivered properties if they were used in certain aggregations (like a job that spanned machines).

The following are the same four examples distinguishing the delivered amounts and values:

```
<Delivered>
  <Processors>12</Processors>
  <Duration>1234</Duration>
</Delivered>
<Delivered>
  <Node>fr15n03</Node>
</Delivered>
<Delivered>
  <User>scottmo</User>
  <Account>mscfops</Account>
</Delivered>
<Delivered>
  <Memory>1024</Memory>
</Delivered>
```

## 9.0 AwarenessPolicy Attribute

A word or two should be said about compatibility mechanisms. With all the leeway in the specification with regard to implementing various portions of the specification, problems might arise if an implementation simply ignores a portion of a job specification that is critical to the job function in certain contexts. Given this situation, it might be desirable in some circumstances for jobs to be rejected by sites that fail to fully support that job's element or attributes. At other times, it might be desirable for a job to run, using a best-effort approach to supporting unimplemented features. Consequently, we define an `awarenessPolicy` attribute that can be added as an optional attribute to the Job element or any other containment or property element to indicate how the property (or the default action for the elements that the containment element encloses) must react when the implementation does not understand an element or attribute.

An awareness policy of `Reject` will cause the server to return a failure if it receives a client request where it does not support an associated element name or attribute name or value. It is reasonable for an implementation to ignore (not even look for) an element or attribute that would not be critical to its function as long as ignoring this attribute or element would not cause an incorrect result. However, any element or attribute that was present that would be expected to be handled in a manner that the implementation does not support must result in a failure.

An awareness policy of `Warn` will accept the misunderstood element or attribute and continue to process the job object on a best effort basis. However a warning **MUST** be sent (if possible) to the requestor enumerating the elements and attributes that are not understood.

An awareness policy of `Ignore` will accept the unsupported element or attribute and continue to process the job object on a best effort basis. The action could be to simply ignore the attribute.

- This name of this attribute is `awarenessPolicy`.
- This attribute is of type `String`.
- This attribute can have values of `Reject`, `Warn` or `Ignore`.
- A compliant implementation **MAY** support this attribute.
- An implementation that does not support an attribute **MUST** reject any job object that contains elements or attributes that it does not support. Furthermore, it **SHOULD** return a message to the requestor with an indication of the element or attribute name it did not understand.
- This attribute **MAY** be present in a property or containment element.
- If an implementation does support the attribute, but it is absent, the default value of `Reject` is implied.

- Individual elements in the job object may override the containing object’s awareness policy default by including this attribute. For example, a job might specify an awarenessPolicy of Reject at its root (the Job element) but may want to allow a particular subset of elements or attributes to be ignored if not understood. Conversely, a job with a default awarenessPolicy of Ignore might want to classify a subset of its optional elements as Reject if they are indispensable to its correct interpretation. An implementation can opt to check or not check for this attribute at any level it wants but must assume a Reject policy for any elements it does not check.

## 10.0 References

### ISO 8601

ISO (International Organization for Standardization). Representations of dates and times, 1988-06-15. <https://www.iso.ch/markete/8601.pdf>

### DATATYPES

XML Schema Part 2: Datatypes. Recommendation, 02 MAY 2001. <https://www.w3.org/TR/xmlschema-2/>

## 11.0 Units of Measure Abbreviations

Abbreviation	Definition	Quantity
<b>B</b>	byte	1 byte
<b>KB</b>	Kilobyte	2 <sup>10</sup> bytes
<b>MB</b>	Megabyte	2 <sup>20</sup> bytes
<b>GB</b>	Gigabyte	2 <sup>30</sup> bytes
<b>TB</b>	Terabyte	2 <sup>40</sup> bytes
<b>PB</b>	Petabyte	2 <sup>50</sup> bytes
<b>EB</b>	Exabyte	2 <sup>60</sup> bytes

Abbreviation	Definition	Quantity
<b>ZB</b>	Zettabyte	2 <sup>70</sup> bytes
<b>YB</b>	Yottabyte	2 <sup>80</sup> bytes
<b>NB</b>	Nonabyte	2 <sup>90</sup> bytes
<b>DB</b>	Doggabyte	2 <sup>100</sup> bytes

## M.2 Scalable Systems Software Resource Management and Accounting Protocol (SSSRMAP) Message Format

Resource Management Interface Specs  
Release v. 3.0.4  
18 JUL 2005

Scott Jackson  
Brett Bode  
David Jackson  
Kevin Walker

### Status of this Memo

This is a specification defining an XML message format used between Scalable Systems Software components. It is intended that this specification will continue to evolve as these interfaces are implemented and thoroughly tested by time and experience.

### Abstract

This document is a specification describing a message format for the interaction of resource management and accounting software components developed as part of the Scalable Systems Software Center. The SSSRMAP Message Format defines a request-response syntax supporting both functional and object-oriented messages. The protocol is specified in XML Schema Definition. The message elements defined in this specification are intended to be framed within the Envelope and Body elements defined in the SSSRMAP Wire Protocol specification document.

## Table of Contents

- 1.0 Introduction
- 2.0 Conventions Used in this Document
  - 2.1 Keywords
  - 2.2 XML Case Conventions
  - 2.3 Schema Definitions
- 3.0 Encoding
  - 3.1 Schema Header and Namespaces
  - 3.2 Element Descriptions
    - 3.2.1 The Request Element
    - 3.2.2 The Object Element
    - 3.2.3 The Get Element
    - 3.2.4 The Set Element
    - 3.2.5 The Where Element
    - 3.2.6 The Option Element
    - 3.2.7 The Data Element
    - 3.2.8 The File Element
    - 3.2.9 The Count Element
    - 3.2.10 The Response Element
    - 3.2.11 The Status Element
    - 3.2.12 The Value Element
    - 3.2.13 The Code Element
    - 3.2.14 The Message Element
  - 3.3 Modified XPATH Expressions
    - 3.3.1 Sample Modified XPATH expressions
  - 3.4 Examples
    - 3.4.1 Sample Requests
    - 3.4.2 Sample Responses

- [4.0 Error Reporting](#)
- [5.0 References](#)

## 1.0 Introduction

A major objective of the Scalable Systems Software [SSS] Center is to create a scalable and modular infrastructure for resource management and accounting on terascale clusters including resource scheduling, grid-scheduling, node daemon support, comprehensive usage accounting and user interfaces emphasizing portability to terascale vendor operating systems. Existing resource management and accounting components feature disparate APIs (Application Programming Interfaces) requiring various forms of application coding to interact with other components.

This document proposes a common message format expressed in an XML request-response syntax to be considered as the foundation of a standard for communications between and among resource management and accounting software components. In this document this standard is expressed in two levels of generality. The features of the core SSSRMAP protocol common to all resource management and accounting components in general are described in the main body of this document. The aspects of the syntax specific to individual components are described in component-specific binding documents.

## 2.0 Conventions Used in this Document

### 2.1 Keywords

The keywords 'MUST', 'MUST NOT', 'REQUIRED', 'SHALL', 'SHALL NOT', 'SHOULD', 'RECOMMENDED', 'MAY', and 'OPTIONAL' in this document are to be interpreted as described in RFC2119 [RFC2119].

### 2.2 XML Case Conventions

In order to enforce a consistent capitalization and naming convention across all SSSRMAP specifications 'Upper Camel Case' (UCC) and 'Lower Camel Case' (LCC) Capitalization styles shall be used. UCC style capitalizes the first character of each word and compounds the name. LCC style capitalizes the first character of each word except the first word. [XML\_CONV][FED\_XML]

1. SSSRMAP XML Schema and XML instance documents SHALL use the following conventions:
  - Element names SHALL be in UCC convention (example: <UpperCamelCaseElement/>).



- Attribute names SHALL be in LCC convention (example: <UpperCamelCaseElement lowerCamelCaseAttribute="Whatever"/>).
2. General rules for all names are:
- Acronyms SHOULD be avoided, but in cases where they are used, the capitalization SHALL remain (example: XMLSignature).
  - Underscores (`_`), periods (`.`) and dashes (`-`) MUST NOT be used (example: use `JobId` instead of `JOB.ID`, `Job_ID` or `job-id`).

## 2.3 Schema Definitions

SSSRMAP Schema Definitions appear like this

In case of disagreement between the schema file and this specification, the schema file takes precedence.

## 3.0 Encoding

Encoding tells how a message is represented when exchanged. SSSRMAP data exchange messages SHALL be defined in terms of XML schema [XML\_SCHEMA].

## 3.1 Schema Header and Namespaces

The header of the schema definition is as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<schema
  xmlns="https://www.w3.org/2001/XMLSchema"
  xmlns:sssrmap="https://scidac.org/ScalableSystems/SSSRMAP"
  targetNamespace="https://www.scidac.org/ScalableSystems/SSSRMAP"
  elementFormDefault="qualified">
```

## 3.2 Element Descriptions

The following subsections describe the elements that make up SSSRMAP messages. SSSRMAP messages are transmitted in the `Body` and `Envelope` elements as described in the SSSRMAP Wire Protocol specification [WIRE\_PROTOCOL].

### The Request Element

The `Request` element specifies an individual request. An object-oriented request will have at least one `Object` element while a functional request will not have one. Depending on context, the `Request` element MAY contain one or more `Get` elements or one or more `Set` elements and any number of `Where` elements. `Option`, `Data`, `File` or `Count`

elements may also be included. If a component supports it, chunking may be requested where large response data is possible. Setting the chunking attribute to 'True' requests that the server break a large response into multiple chunks (each with their own envelope) so they can be processed in separate pieces.

Only an `action` attribute is required. All other attributes are optional.

Attribute	Description
<b>action</b>	Specifies the action or function to be performed
<b>actor</b>	The authenticated user sending the request
<b>id</b>	Uniquely maps the request to the appropriate response
<b>chunking</b>	Requests that segmentation be used for large response data if set to 'True'
<b>chunkSize</b>	Requests that the segmentation size be no larger than the specified amount

```
<complexType name="RequestType">
  <choice minOccurs="0" maxOccurs="unbounded">
    <element ref="sssrmmap:Object" minOccurs="0" maxOccurs="unbounded"/>
    <element ref="sssrmmap:Option" minOccurs="0" maxOccurs="unbounded"/>
    <choice minOccurs="0" maxOccurs="1">
      <element ref="sssrmmap:Get" minOccurs="1" maxOccurs="unbounded"/>
      <element ref="sssrmmap:Set" minOccurs="1" maxOccurs="unbounded"/>
    </choice>
    <element ref="sssrmmap:Where" minOccurs="0" maxOccurs="unbounded"/>
    <element ref="sssrmmap:Data" minOccurs="0" maxOccurs="unbounded"/>
    <element ref="sssrmmap:Count" minOccurs="0" maxOccurs="1"/>
    <any namespace="##other" minOccurs="0" maxOccurs="unbounded"/>
  </choice>
  <attribute name="action" type="string" use="required"/>
  <attribute name="actor" type="string" use="required"/>
  <attribute name="id" type="string" use="optional"/>
  <attribute name="chunking" type="sssrmmap:BoolType" use="optional"/>
  <attribute name="chunkSize" type="positiveInteger" use="optional"/>
</complexType>

<element name="Request" type="sssrmmap:RequestType"/>
```

## The Object Element

The `Object` element is used in an object-oriented request to specify the object receiving the action. It is possible to have multiple `Object` elements in a request if an implementation supports multi-object queries.

The object class name is specified as text content. All attributes are optional.

- `join` – the type of join to be performed with the preceding object
  - A `join` attribute of 'Inner' specifies an inner join. This is the default.
  - A `join` attribute of 'FullOuter' specifies a full outer join.
  - A `join` attribute of 'LeftOuter' specifies a left outer join.
  - A `join` attribute of 'RightOuter' specifies a right outer join.
  - A `join` attribute of 'Cross' specifies a cross join.
  - A `join` attribute of 'Union' specifies a union join.

```

<complexType name="ObjectType">
  <simpleContent>
    <extension base="string">
      <attribute name="join" type="string" use="optional"/>
    </extension>
  </simpleContent>
</complexType>

<element name="Object" type="sssrmap:ObjectType"/>

```

## The Get Element

The `Get` element is used to indicate the data fields to be returned in a query. `Get` is typically used within requests with `action="query"`. Multiple `Get` elements cause the fields to be returned in the order specified. If no `Get` elements are specified, the query will return a default set of fields.

Only a `name` attribute is required. All other attributes are optional.

Attribute	Description
<b>name</b>	The name of the data field to be returned. This MUST be of the form of a 'Modified XPATH expression' as described in a later section.
<b>op</b>	<p>The operator to be used to aggregate or perform an operation on the returned values:</p> <ul style="list-style-type: none"> <li>• An <i>op</i> attribute of 'Sort' specifies an ascending sort operation</li> <li>• An <i>op</i> attribute of 'Tros' specifies a descending sort operation</li> <li>• An <i>op</i> attribute of 'Sum' returns the sum (only valid for numeric values)</li> <li>• An <i>op</i> attribute of 'Max' returns the maximum value</li> <li>• An <i>op</i> attribute of 'Min' returns the minimum value</li> <li>• An <i>op</i> attribute of 'Count' returns the number of values</li> <li>• An <i>op</i> attribute of 'Average' returns the average of the values</li> <li>• An <i>op</i> attribute of 'GroupBy' signifies that aggregates are grouped by this field</li> </ul>

Attribute	Description
<b>object</b>	Specifies the object for which you want the named attribute in a multi-object query.
<b>units</b>	The units in which to return the value (if applicable)

```

<complexType name="GetType">
  <attribute name="name" type="string" use="required"/>
  <attribute name="object" type="string" use="optional"/>
  <attribute name="op" type="sssrmap:GetOperatorType" use="optional"/>
  <attribute name="units" type="string" use="optional"/>
</complexType>

<element name="Get" type="sssrmap:GetType"/>

<simpleType name="GetOperatorType">
  <restriction base="string">
    <enumeration value="Sort"/>
    <enumeration value="Tros"/>
    <enumeration value="Count"/>
    <enumeration value="Sum"/>
    <enumeration value="Max"/>
    <enumeration value="Min"/>
    <enumeration value="Average"/>
    <enumeration value="GroupBy"/>
  </restriction>
</simpleType>

```

## The Set Element

The `Set` element is used to specify the object data fields to be assigned values. `Set` is typically used within requests with `action="Create"` or `action="Modify"`. The use of `Get` or `Set` elements within a request is mutually exclusive.

The assignment value (to which the field is being changed) is specified as the text content. A `Set` element without a value can be used as an assertion flag. Only the `name` attribute is required. All other attributes are optional.

Attribute	Description
<b>name</b>	The name of the field being assigned a value. This MUST be of the form of a 'Modified XPATH expression' as described in a later section.
<b>op</b>	The operator to be used in assigning a new value to the name. If an <code>op</code> attribute is not specified and a value is specified, the specified value will be assigned to the named field ('assign'). <ul style="list-style-type: none"> <li>An <code>op</code> attribute of 'Assign' assigns value to the named field</li> <li>An <code>op</code> attribute of 'Inc' increments the named field by the value</li> <li>An <code>op</code> attribute of 'Dec' decrements the named field by the value</li> </ul>

Attribute	Description
<b>units</b>	The units corresponding to the value being set

```

<complexType name="SetType">
  <simpleContent>
    <extension base="string">
      <attribute name="name" type="string" use="required"/>
      <attribute name="op" type="sssrmmap:SetOperatorType" use="optional"/>
      <attribute name="units" type="string" use="optional"/>
    </extension>
  </simpleContent>
</complexType>

<element name="Set" type="sssrmmap:SetType"/>

<simpleType name="SetOperatorType">
  <restriction base="string">
    <enumeration value="Assign"/>
    <enumeration value="Inc"/>
    <enumeration value="Dec"/>
  </restriction>
</simpleType>

```

### The Where Element

A `Request` element can contain one or more `Where` elements that specify the search conditions for which objects the action is to be performed on.

The condition value (against which the field is tested) is specified as the text content. A `Where` element without a value can be used as a truth test. Only the `name` attribute is required. All other attributes are optional.

Attribute	Description
<b>name</b>	The name of the data field to be tested. This MUST be of the form of a 'Modified XPATH expression' as described in a later section.
<b>op</b>	<p>The operator to be used to test the name against the value. If an <code>op</code> attribute is not specified and a value is specified, the field will be tested whether it is equal to the value ('EQ'):</p> <ul style="list-style-type: none"> <li>• An <code>op</code> attribute of 'EQ' specifies an equality comparison</li> <li>• An <code>op</code> attribute of 'LT' specifies a 'less than' comparison</li> <li>• An <code>op</code> attribute of 'GT' specifies a 'greater than' comparison</li> <li>• An <code>op</code> attribute of 'LE' specifies a 'less than or equal to' test</li> <li>• An <code>op</code> attribute of 'GE' specifies a 'greater than or equal to' test</li> <li>• An <code>op</code> attribute of 'NE' specifies a 'not equal to' test</li> <li>• An <code>op</code> attribute of 'Match' specifies a regular expression matching comparison</li> </ul>

Attribute	Description
<b>conj</b>	Indicates whether this test is to be ANDed or ORed with the immediately preceding where condition: <ul style="list-style-type: none"> <li>• A <code>conj</code> attribute of 'And' specifies an 'and' conjunction</li> <li>• A <code>conj</code> attribute of 'Or' specifies an 'or' condition</li> <li>• A <code>conj</code> attribute of 'AndNot' specifies an 'and not' conjunction</li> <li>• A <code>conj</code> attribute of 'OrNot' specifies an 'or not' condition</li> </ul>
<b>group</b>	Indicates an increase or decrease of parentheses grouping depth: <ul style="list-style-type: none"> <li>• A positive number indicates the number of left parentheses to precede the condition [i.e., <code>group="2"</code> represents <code>"((condition)"</code>].</li> <li>• A negative number indicates the number of right parentheses to follow the condition [i.e., <code>group="-2"</code> represents <code>"condition))"</code>].</li> </ul>
<b>object</b>	Specifies the object for the first operand in a multi-object query.
<b>subject</b>	Specifies the object for the second operand in a multi-object query.
<b>units</b>	Indicates the units to be used in the value comparison

```

<complexType name="WhereType">
  <simpleContent>
    <extension base="string">
      <attribute name="name" type="string" use="required"/>
      <attribute name="op" type="sssrmap:OperatorType" use="optional"/>
      <attribute name="conj" type="sssrmap:ConjunctionType" use="optional"/>
      <attribute name="group" type="integer" use="optional"/>
      <attribute name="units" type="string" use="optional"/>
    </extension>
  </simpleContent>
</complexType>

<element name="Where" type="sssrmap:WhereType"/>

<simpleType name="WhereOperatorType">
  <restriction base="string">
    <enumeration value="EQ"/>
    <enumeration value="GT"/>
    <enumeration value="LT"/>
    <enumeration value="GE"/>
    <enumeration value="LE"/>
    <enumeration value="NE"/>
    <enumeration value="Match"/>
  </restriction>
</simpleType>

```

## The Option Element

The `Option` element is used to indicate processing options for the command. An option might be used to indicate that command usage or special formatting is desired, or that the

command is to be invoked with particular options.

The option value is specified as the text content. An `Option` element without a value can be used as an assertion flag. Only the `name` attribute is required. All other attributes are optional.

Attribute	Description
<b>name</b>	The name of the field being assigned a value.
<b>op</b>	The operator to be used to disassert the option: <ul style="list-style-type: none"> <li>An <code>op</code> attribute of 'Not' specifies that the option is not asserted</li> </ul>
<b>conj</b>	Indicates whether this test is to be ANDed or ORed with the immediately preceding where condition: <ul style="list-style-type: none"> <li>A <code>conj</code> attribute of 'And' specifies an 'and' conjunction</li> <li>A <code>conj</code> attribute of 'Or' specifies an 'or' condition</li> <li>A <code>conj</code> attribute of 'AndNot' specifies an 'and not' conjunction</li> <li>A <code>conj</code> attribute of 'OrNot' specifies an 'or not' condition</li> </ul>

```

<complexType name="OptionType">
  <simpleContent>
    <extension base="string">
      <attribute name="name" type="string" use="required"/>
      <attribute name="op" type="sssrmmap:OptionOperatorType" use="optional"/>
      <attribute name="conj" type="sssrmmap:ConjunctionType" use="optional"/>
    </extension>
  </simpleContent>
</complexType>

<element name="Option" type="sssrmmap:OptionType"/>

<simpleType name="OptionOperatorType">
  <restriction base="string">
    <enumeration value="Not"/>
  </restriction>
</simpleType>

```

## The Data Element

A `Request` or `Response` element can have one or more `Data` elements that allow the supplying of context-specific data. A request might pass in a structured object via a `Data` element to be acted upon. Typically a query will result in a response with the data encapsulated within a `Data` element.

The following attributes are optional:

Attribute	Description
<b>name</b>	Object name describing the contents of the data.
<b>type</b>	Describing the form in which the data is represented: <ul style="list-style-type: none"> <li>• A <code>type</code> attribute of 'XML' indicates the data has internal xml structure and can be recursively parsed by an XML parser</li> <li>• A <code>type</code> attribute of 'Binary' indicates the data is an opaque dataset consisting of binary data</li> <li>• A <code>type</code> attribute of 'String' indicates the data is an ASCII string</li> <li>• A <code>type</code> attribute of 'Int' indicates the data is an integer</li> <li>• A <code>type</code> attribute of 'Text' indicates the data is in formatted human-readable text</li> <li>• A <code>type</code> attribute of 'HTML' indicates the data is represented in HTML</li> </ul>

```

<complexType name="DataType">
  <sequence>
    <any namespace="##any" processContents="lax" minOccurs="0" maxOccurs="unbounded"/>
  </sequence>
  <attribute name="name" type="string" use="optional"/>
  <attribute ref="ssrmap:Type" use="optional"/>
</complexType>

<element name="data" type="ssrmap:DataType"/>

```

## The File Element

A `Request` or `Response` element can have one or more `File` elements of type `String` that allow the inclusion of files. The files can be either text or binary and can be referenced by objects inside the `Data` element. A file can be compressed using the gzip algorithm [ZIP]. A binary file or a compressed file must be base64 encoded as defined in XML Digital Signatures (<https://www.w3.org/2000/09/xmlsig#base64>). Metadata describing the modes and properties of the resulting file are passed as parameters. The text or base64 encoded file data forms the string content of the `File` element.

The following attributes are optional:

Attribute	Description
<b>id</b>	Specifies an identifier that allows the file to be referenced from within another object. If more than one <code>File</code> elements are specified, this attribute is REQUIRED in each of them.
<b>name</b>	Specifies the name to give the file upon creation on the target system. This can be an absolute or relative pathname (relative to the <code>InitialWorkingDirectory</code> ).



Attribute	Description
<b>owner</b>	Indicates what owner the file should be changed to. By default, it will be changed to the UserId that the authenticated actor maps to on the target system. Note that this function should succeed only if the requestor has the privileges to do so (i.e., authenticated as root).
<b>group</b>	Indicates what group the file should be changed to. By default, it will be set to the primary groupid of the UserId that the authenticated actor maps to on the target system. Note that this function should succeed only if the requestor has the proper privileges.
<b>mode</b>	Indicates the permissions the file should possess. By default, it will be set according to the default umask for the UserId that the authenticated actor maps to on the target system. Note that this function should not set permissions for the file that exceed the privileges for the actor. These permissions can be specified using either an octal number or symbolic operations (as accepted by the GNU chmod(1) command).
<b>compressed</b>	Indicates whether the file has been compressed: <ul style="list-style-type: none"> <li>• A <code>compressed</code> attribute of 'True' indicates the file has been compressed.</li> <li>• A <code>compressed</code> attribute of 'False' indicates the file has not been compressed. This is the default.</li> </ul>
<b>encoded</b>	Indicates whether the file has been base64 encoded: <ul style="list-style-type: none"> <li>• An <code>encoded</code> attribute of 'True' indicates the file has been encoded.</li> <li>• An <code>encoded</code> attribute of 'False' indicates the file has not been encoded. This is the default.</li> </ul>

```

<complexType name="FileType">
  <sequence>
    <any namespace="##any" processContents="lax" minOccurs="0" maxOccurs="unbounded"/>
  </sequence>
  <attribute name="name" type="string" use="optional"/>
  <attribute name="owner" type="string" use="optional"/>
  <attribute name="group" type="string" use="optional"/>
  <attribute name="mode" type="string" use="optional"/>
  <attribute name="compressed" type="boolean" use="optional"/>
  <attribute name="encoded" type="boolean" use="optional"/>
</complexType>
<element name="file" type="sssrmmap:FileType"/>

```

## The Count Element

A single `Count` element can be included within a `Request` or `Response` and is context-specific. This can be used to represent the number of objects acted upon or returned.

```

<element name="Count" type="positiveInteger"/>

```

## The Response Element

The `Response` element specifies an individual response. It **MUST** contain a `Status` element. It **MAY** also contain `Count` and any number of `Data` or `File` elements. If chunking has been requested and is supported by the server, a large response can be broken up into multiple chunks (each with their own envelope). The `chunkNum` attribute can be used to indicate which chunk the current one is. The `chunkMax` attribute can be used to determine when all the chunks have been received (all chunks have been received if `chunkNum=chunkMax` or `chunkMax=0`).

It **MAY** have any of the following attributes:

Attribute	Description
<b>id</b>	Uniquely maps the response to the corresponding request
<b>chunkNum</b>	Integer indicating the current chunk number [1 is implied when this attribute is missing or blank]
<b>chunkMax</b>	Integer indicating the number of chunks expected [-1 means unknown but more chunks to follow; 0 means unknown but this is the last chunk; 0 is implied if this attribute is missing or blank]

```
<complexType name="ResponseType">
  <choice minOccurs="0" maxOccurs="unbounded">
    <element ref="sssrmap:Status" minOccurs="1" maxOccurs="1"/>
    <element ref="sssrmap:Count" minOccurs="0" maxOccurs="1"/>
    <element ref="sssrmap:Data" minOccurs="0" maxOccurs="unbounded"/>
    <element ref="sssrmap:File" minOccurs="0" maxOccurs="unbounded"/>
    <any minOccurs="0" maxOccurs="unbounded" namespace="##other"/>
  </choice>
  <attribute name="object" type="string" use="optional"/>
  <attribute name="action" type="string" use="optional"/>
  <attribute name="id" type="string" use="optional"/>
  <attribute name="chunkNum" type="integer" use="optional"/>
  <attribute name="chunkMax" type="integer" use="optional"/>
</complexType>

<element name="Response" type="sssrmap:ResponseType"/>
```

## The Status Element

A `Response` element **MUST** contain a single `Status` element that indicates whether the reply represents a success, warning or failure. This element is composed of the child elements `Value`, `Code` and `Message`. Of these, `Value` and `Code` are required, and `Message` is optional.

```
<complexType name="StatusType">
  <choice minOccurs="1" maxOccurs="unbounded">
    <element ref="sssrmap:Value" minOccurs="1" maxOccurs="1"/>
    <element ref="sssrmap:Code" minOccurs="1" maxOccurs="1"/>
    <element ref="sssrmap:Message" minOccurs="0" maxOccurs="1"/>
    <any minOccurs="0" maxOccurs="unbounded" namespace="##other"/>
  </choice>
</complexType>
```

```

</complexType>
<element name="Status" type="sssmap:StatusType"/>

```

## The Value Element

The `Value` element is of type `String` and **MUST** have a value of 'Success', 'Warning' or 'Failure':

```

<simpleType name="StatusValueType">
  <restriction base="string">
    <enumeration value="Success"/>
    <enumeration value="Warning"/>
    <enumeration value="Failure"/>
  </restriction>
</simpleType>
<element name="Value" type="sssmap:StatusValueType"/>

```

## The Code Element

A `Response` element must contain a single `Code` element that specifies the 3-digit status code for the response. Refer to the next section on Error Reporting for a description and listing of supported status codes.

```

<simpleType name="CodeType">
  <restriction base="string">
    <pattern value="[0-9]{3}"/>
  </restriction>
</simpleType>
<element name="Code" type="sssmap:CodeType"/>

```

## The Message Element

A `Response` element can contain a single `Message` element that is context specific to the success or failure response. The message should be an error message if status is false. If present for a successful response, it can be used as a human readable message for a user interface.

```

<element name="Message" type="string"/>

```

## 3.3 Modified XPATH Expressions

The `name` attribute used within the `Get`, `Set` and `Where` Elements **MUST** have the form of a modified XPATH expression as defined in this section. Usually this will just be the simple name of the object property. Some complex objects, such as the SSS Job Object and the SSS Node Object, however, are represented in a structured way with nested elements. In order to define a consistent and flexible way to access and manipulate these objects, as well as keeping the flat XML objects simple and straightforward, SSSRMAP specifies that a 'Modified XPATH' syntax be used.

In essence, 'Modified XPATH' is defined to be an XPATH [XPATH] expression with the exception that the '/' can be omitted from the beginning of the expression when a document search is desired. Therefore, on the server side, a standard XPATH routine can be used by prepending '/' to any expression that does not begin with a '/'.

The response data should always include all of the structure of the queried object necessary to place the requested data in its proper context.

See the XPATH specification for a full description of XPATH. The XPath 1.0 Recommendation is <https://www.w3.org/TR/1999/REC-xpath-19991116>. The latest version of XPath 1.0 is available at <https://www.w3.org/TR/xpath>.

## Sample Modified XPATH Expressions

Consider the following hypothetical object(s) (which might be returned within a Data element):

```
<Job>
  <JobId>PBS.1234.0</JobId>
  <Requested>
    <Memory op="GE">512</Memory>
    <Processors>2</Processors>
    <WallDuration>P3600S</WallDuration>
  </Requested>
  <Utilized>
    <Memory metric="Average">488</Memory>
    <WallDuration>P1441S</WallDuration>
  </Utilized>
</Job>
```

To get everything above for this job you would not need a Get element:

```
<Request action="Query">
  <Object>Job</Object>
  <Where name="JobId">PBS.1234.0</Where>
</Request>
```

If you used `<Get name="JobId" />` you would get back:

```
<Job>
  <JobId>PBS.1234.0</JobId>
</Job>
```

If you used `<Get name="Memory" />` (or `name="/Job/*/Memory"`) you would get:

```
<Job>
  <Requested>
    <Memory op="GE">512</Memory>
  </Requested>
  <Utilized>
    <Memory metric="Average">488</Memory>
  </Utilized>
</Job>
```

If you used `<Get name="Requested/Memory" />` (or `name="/Job/Requested/Memory"`) you would get:

```

<Job>
  <Requested>
    <Memory op="GE">512</Memory>
  </Requested>
</Job>

```

If you used `<Get name="Memory[@metric='Average']"/>` (or `name="Memory[@metric]"`) you would get:

```

<Job>
  <Utilized>
    <Memory metric="Average">488</Memory>
  </Utilized>
</Job>

```

## 3.4 Examples

### Sample Requests

Requesting a list of nodes with a certain configured memory threshold (batch format):

```

<Request action="Query" id="1">
  <Object>Node</Object>
  <Get name="Name" />
  <Get name="Configured/Memory" />
  <Where name="Configured/Memory" op="GE" units="MB">512</Where>
</Request>

```

Activating a couple of users:

```

<Request action="Modify">
  <Object>User</Object>
  <Set name="Active">True</Set>
  <Where name="Name">scott</Where>
  <Where name="Name" conj="Or">brett</Where>
</Request>

```

Submitting a simple job:

```

<Request action="Submit">
  <Object>Job</Object>
  <Data>
    <Job>
      <User>xdp</User>
      <Account>youraccount</Account>
      <Command>myprogram</Command>
      <InitialWorkingDirectory>/usr/home/scl/xdp</InitialWorkingDirectory>
      <RequestedNodes>4</RequestedNodes>
      <RequestedWCTime>100</RequestedWCTime>
    </Job>
  </Data>
</Request>

```

### Sample Responses

A response to the available memory nodes query (batch format):

```

<Response id="1">
  <Status>
    <Value>Success</Value>
    <Code>000</Code>
  </Status>
  <Count>2</Count>
  <Data>
    <Node>
      <Name>fr01n01</Name>
      <Configured>
        <Memory>512</Memory>
      </Configured>
    </Node>
    <Node>
      <Name>fr12n04</Name>
      <Configured>
        <Memory>1024</Memory>
      </Configured>
    </Node>
  </Data>
</Response>

```

Two users successfully activated:

```

<Response>
  <Status>
    <Code>000</Code>
    <Message>Two users were successfully modified</Message>
  </Status>
  <Count>2</Count>
</Response>

```

A failed job submission:

```

<Response>
  <Status>
    <Value>Failure</Value>
    <Code>711</Code>
    <Message>Invalid account specified. The job was not submitted.</Message>
  </Status>
</Response>

```

## 4.0 Error Reporting

SSSRMAP requests will return a status and a 3-digit response code to signify success or failure conditions. When a request is successful, a corresponding response is returned with the status element set to Success and the code element set to '000'. When a request results in an error detected by the server, a response is returned with the status element set to Failure and a 3-digit error code in the code element. An optional human-readable message can also be included in a failure response providing context-specific detail about the failure. The default message language is US English. (The status flag makes it easy to signal success or failure and allows the receiving peer some freedom in the amount of parsing it wants to do on failure [BXXP]).

## Success Codes

Code	Response Text in US English
<b>0xx</b>	Request was successful
<b>000</b>	General Success
<b>010</b>	Help/usage reply
<b>020</b>	Status reply
<b>030</b>	Subscription successful
<b>035</b>	Notification successful (Ack)
<b>040</b>	Registration successful
<b>050-079</b>	Component-defined
<b>080-099</b>	Application-defined

## Warning Codes

Code	Response Text in US English
<b>1xx</b>	Request was successful but includes a warning
<b>100</b>	General warning (examine message for details)
<b>102</b>	Check result (Did what you asked but may not have been what you intended -- or information is suspect)
<b>110</b>	Wire Protocol or Network warning
<b>112</b>	Redirect
<b>114</b>	Protocol warning (something was wrong with the protocol, but best effort guesses were applied to fulfill the request)

Code	Response Text in US English
120	Message Format warning
122	Incomplete specification (request missing some essential information -- best effort guess applied)
124	Format warning (something was wrong with the format but best effort guesses were applied to fulfill the request)
130	Security warning
132	Insecure request
134	Insufficient privileges (Response was sanitized or reduced in scope due to lack of privileges)
140	Content or action warning
142	No content (The server has processed the request but there is no data to be returned)
144	No action taken (nothing acted upon -- i.e., deletion request did not match any objects)
146	Partial content
148	Partial action taken
150-179	Component-defined
180-199	Application-defined

## Wire Protocol Codes

Code	Response Text in US English
2xx	A problem occurred in the wire protocol or network



Code	Response Text in US English
<b>200</b>	General wire protocol or network error
<b>210</b>	Network failure
<b>212</b>	Cannot resolve host name
<b>214</b>	Cannot resolve service port
<b>216</b>	Cannot create socket
<b>218</b>	Cannot bind socket
<b>220</b>	Connection failure
<b>222</b>	Cannot connect
<b>224</b>	Cannot send data
<b>226</b>	Cannot receive data
<b>230</b>	Connection rejected
<b>232</b>	Timed out
<b>234</b>	Too busy
<b>236</b>	Message too large
<b>240</b>	Framing failure
<b>242</b>	Malformed framing protocol
<b>244</b>	Invalid payload size
<b>246</b>	Unexpected end of file
<b>250-279</b>	Component-defined
<b>280-299</b>	Application-defined

## Message Format Codes

Code	Response Text in US English
<b>3xx</b>	A problem occurred in the message format
<b>300</b>	General message format error
<b>302</b>	Malformed XML document
<b>304</b>	Validation error (XML Schema)
<b>306</b>	Namespace error
<b>308</b>	Invalid message type (Something other than Request or Response in Body)
<b>310</b>	General syntax error in request
<b>311</b>	Object incorrectly (or not) specified
<b>312</b>	Action incorrectly (or not) specified
<b>313</b>	Invalid Action
<b>314</b>	Missing required element or attribute
<b>315</b>	Invalid Object (or Object-Action combination)
<b>316</b>	Invalid element or attribute name
<b>317</b>	Illegal value for element or attribute
<b>318</b>	Illegal combination
<b>319</b>	Malformed Data
<b>320</b>	General syntax error in response
<b>321</b>	Status incorrectly (or not) specified
<b>322</b>	Code incorrectly (or not) specified

Code	Response Text in US English
324	Missing required element or attribute
326	Invalid element or attribute name
327	Illegal value for element or attribute
328	Illegal combination
329	Malformed Data
340	Pipelining failure
342	Request identifier is not unique
344	Multiple messages not supported
346	Mixed messages not supported (both requests and responses in same batch)
348	Request/response count mismatch
350-379	Component-defined
380-399	Application-defined

## Security Codes

Code	Response Text in US English
4xx	A security requirement was not fulfilled
400	General security error
410	Negotiation failure
412	Not understood
414	Not supported

Code	Response Text in US English
416	Not accepted
420	Authentication failure
422	Signature failed at client
424	Authentication failed at server
426	Signature failed at server
428	Authentication failed at client
430	Encryption failure
432	Encryption failed at client
434	Decryption failed at server
436	Encryption failed at server
438	Decryption failed at client
440	Authorization failure
442	Authorization failed at client
444	Authorization failed at server
450-479	Component-defined
480-499	Application-defined

## Event Management Codes

Code	Response Text in US English
5xx	Failure conditions in event messaging

Code	Response Text in US English
<b>500</b>	General Event Management failure
<b>510</b>	Subscription failed
<b>520</b>	Notification failed
<b>550-579</b>	Component-defined
<b>580-599</b>	Application-defined

## Reserved Codes

Code	Response Text in US English
<b>6xx</b>	Reserved for future use

## Server Application Codes

Code	Response Text in US English
<b>7xx</b>	A server-side application-specific error occurred
<b>700</b>	General failure
<b>710</b>	Not supported
<b>712</b>	Not understood
<b>720</b>	Internal error
<b>730</b>	Resource unavailable (insufficient resources -- software, hardware or a service I rely upon is down)
<b>740</b>	Business logic
<b>750-</b>	Component-defined

Code	Response Text in US English
<b>779</b>	
<b>780-799</b>	Application-defined

## Client Application Codes

Code	Response Text in US English
<b>8xx</b>	A client-side application-specific error occurred
<b>800</b>	General failure
<b>810</b>	Not supported
<b>812</b>	Not understood
<b>820</b>	Internal error
<b>830</b>	Resource unavailable
<b>840</b>	Business logic
<b>850-879</b>	Component-defined
<b>880-899</b>	Application-defined

## Miscellaneous Codes

Code	Response Text in US English
<b>9xx</b>	Miscellaneous failures
<b>999</b>	Unknown failure

## 5.0 References

- [BEEP] M. Rose, “The Blocks Extensible Exchange Protocol Core”, [RFC 3080](#), March 2001.
- [FED\_XML] “[U.S. Federal XML Guidelines](#)”.
- [HMAC] H. Krawczyk, M. Bellare, R. Canetti, “HMAC, Keyed-Hashing for Message Authentication”, [RFC 2104](#), February 1997.
- [HTTP] “Hypertext Transfer Protocol – HTTP/1.1”, [RFC 2616](#), June 1999.
- [RFC2119] S. Bradner, “Key Words for Use in RFCs to Indicate Requirement Levels”, [RFC 2119](#), March 1997.
- [RFC3117] M. Rose, “On the Design of Application Protocols”, [Informational RFC 3117](#), November 2001.
- [SHA-1] U.S. Department of Commerce/National Institute of Standards and Technology, “[Secure Hash Standard](#)”, FIPS PUB 180-1.
- [SSS] “Scalable Systems Software”, <https://www.scidac.org/ScalableSystems>
- [WIRE\_PROTOCOL] S. Jackson, B. Bode, D. Jackson, K. Walker, “Systems Software Resource Management and Accounting Protocol (SSSRMAP) Wire Protocol”, [SSS Resource Management and Accounting Documents](#), January 2004.
- [XML] Bray, T., et al, “[Extensible Markup Language \(XML\) 1.0 \(Second Edition\)](#)”, 6 October 2000.
- [XML\_CONV] “[I-X and <I-N-CA> XML Conventions](#)”.
- [XML\_DSIG] D. Eastlake, J. Reagle Jr., D. Solo, “[XML Signature Syntax and Processing](#)”, W3C Recommendation, 12 February 2002.
- [XML\_ENC] T. Imamura, B. Dillaway, E. Smon, “[XML Encryption Syntax and Processing](#)”, W3C Candidate Recommendation, 4 March 2002.
- [XML\_SCHEMA] D. Beech, M. Maloney, N. Mendelshohn, “[XML Schema Part 1: Structures Working Draft](#)”, April 2000.
- [XPath 1.0] J. Clark, S. DeRose, “[XML Path Language \(XPath\) Version 1.0](#)”, 16 November 1999.
- [XRP] E. Brunner-Williams, A. Damaraju, N. Zhang, “[Extensible Registry Protocol \(XRP\)](#)”, Internet Draft, expired August 2001.
- [ZIP] J. Gailly, M. Adler, “The gzip home page”, <https://www.gzip.org/>

## M.3 Scalable Systems Software Node Object Specification

SSS Node Object Specification  
Release Version 3.1.0  
26 April 2011

Scott Jackson, PNNL  
David Jackson, Ames Lab  
Brett Bode, Ames Lab

### Status of this Memo

This is a specification of the node object to be used by Scalable Systems Software compliant components. It is envisioned for this specification to be used in conjunction with the SSSRMAP protocol with the node object passed in the Data field of Requests and Responses. Queries can be issued to a node-cognizant component in the form of modified XPATH expressions to the Get field to extract specific information from the node object as described in the SSSRMAP protocol.

### Abstract

This document describes the syntax and structure of the SSS node object. This node model takes into account various node property categories such as whether it represents a configured, available or utilized property.

### Table of Contents

- [1.0 Introduction](#)
  - [1.1 Goals](#)
  - [1.2 Examples](#)
    - [1.2.1 Simple Example](#)
    - [1.2.2 Elaborate Example](#)
- [2.0 Conventions Used in this Document](#)
  - [2.1 Keywords](#)
  - [2.2 Table Column Interpretations](#)



- 2.3 Element Syntax Cardinality
- 3.0 The Node Model
- 4.0 Node Element
  - 4.1 Uncategorized Node Properties
    - 4.1.1 Simple Node Properties
    - 4.1.2 Extension Element
  - 4.2 Property Categories
    - 4.2.1 Configured Element
    - 4.2.2 Available Element
    - 4.2.3 Utilized Element
  - 4.3 Categorized Node Properties
    - 4.3.1 Consumable Resources
    - 4.3.2 Resource Element
- 5.0 Units of Measure Abbreviations

## 1.0 Introduction

This specification proposes a standard XML representation for a node object for use by the various components in the SSS Resource Management System. This object will be used in multiple contexts and by multiple components. It is anticipated that this object will be passed via the Data Element of SSSRMAP Requests and Responses.

## 1.1 Goals

There are several goals motivating the design of this representation.

It needs to be inherently flexible. We recognize we will not be able to exhaustively include the ever-changing node properties and capabilities that constantly arise.

The same node object should be used at all stages of its lifecycle. This object needs to distinguish between configured, available and utilized properties of a node.

Its design takes into account the properties and structure required to function in a meta or grid environment. It should eventually include the capability of resolving namespace and locality issues, though the earliest versions will ignore this requirement.

One should not have to make multiple queries to obtain a single piece of information (i.e., there should not be two mutually exclusive ways to represent a node resource).

It needs to support resource metric, as well as unit specifications.

## 1.2 Examples

### Simple Example

This example shows a simple expression of the Node object:

```
<Node>
  <Id>Node64</Id>
  <Configured>
    <Processors>2</Processors>
    <Memory>512</Memory>
  </Configured>
</Node>
```

### Elaborate Example

This example shows a more elaborate Node object:

```
<Node>
  <Id>64</Id>
  <Name>Netpipe2</Name>
  <Feature>BigMem</Feature>
  <Feature>NetOC12</Feature>
  <Opsys>AIX</Opsys>
  <Arch>Power4</Arch>
  <Configured>
    <Processors>16</Processors>
    <Memory units="MB">512</Memory>
    <Swap>512</Swap>
  </Configured>
  <Available>
    <Processors>7</Processors>
    <Memory metric="Instantaneous">143</Memory>
  </Available>
  <Utilized>
    <Processors wallDuration="3576">8</Processors>
    <Memory metric="Average" wallDuration="3576">400</Memory>
  </Utilized>
</Node>
```

## 2.0 Conventions Used in this Document

### 2.1 Keywords

The keywords **MUST**, **MUST NOT**, **REQUIRED**, **SHALL**, **SHALL NOT**, **SHOULD**, **RECOMMENDED**, **MAY**, and **OPTIONAL** in this document are to be interpreted as described in [RFC2119](#).

## 2.2 Table Column Interpretations

In the property tables, the columns are interpreted to have the following meanings:

Property	Description
<b>Element Name</b>	Name of the XML element (xsd:element)
<b>Type</b>	<p>Data type defined by xsd (XML Schema Definition) as:</p> <ul style="list-style-type: none"> <li>• String — xsd:string(a finite length sequence of printable characters)</li> <li>• Integer — xsd:integer(a signed finite length sequence of decimal digits)</li> <li>• Float — xsd:float (single-precision 32-bit floating point)</li> <li>• Boolean — xsd:boolean (consists of the literals “true” or “false”)</li> <li>• DateTime — xsd:dateTime (discreet time values are represented in ISO 8601 extended format CCYY-MM-DDThh:mm:ss where CC represents the century, YY the year, MM the month and DD the day. The letter T is the date/time separator and hh, mm, ss represent hour, minute and second respectively. This representation may be immediately followed by a Z to indicate Coordinated Universal Time (UTC) or, to indicate the time zone (i.e., the difference between the local time and Coordinated Universal Time), immediately followed by a sign, + or -, followed by the difference from UTC.</li> <li>• Duration — xsd:duration (a duration of time is represented in ISO 8601 extended format PnYnMnDTnHnMnS, where nY represents the number of years, nM the number of months, nD the number of days, T is the date/time separator, nH the number of hours, nM the number of minutes and nS the number of seconds. The number of seconds can include decimal digits to arbitrary precision.)</li> </ul>
<b>Description</b>	Brief description of the meaning of the property
<b>Appearance</b>	<p>Indicates whether the given property has to appear within the parent element. It assumes the following meanings:</p> <ul style="list-style-type: none"> <li>• MUST — This property is REQUIRED when the parent is specified.</li> <li>• SHOULD — A compliant implementation SHOULD support this property.</li> <li>• MAY — A compliant implementation MAY support this property.</li> </ul>
<b>Compliance</b>	<p>Indicates whether a compliant implementation has to support the given property:</p> <ul style="list-style-type: none"> <li>• MUST — A compliant implementation MUST support this property.</li> <li>• SHOULD — A compliant implementation SHOULD support this property.</li> <li>• MAY — A compliant implementation MAY support this property.</li> </ul>

Property	Description
<b>Categories</b>	<p>Some properties can be categorized into one of several categories. Letters in this column indicate that the given property can be classified in the following property categories:</p> <ul style="list-style-type: none"> <li>• C — This property can be encompassed in a Configured element.</li> <li>• A — This property can be encompassed in an Available element.</li> <li>• U — This property can be encompassed in a Utilized element.</li> </ul>

## 2.3 Element Syntax Cardinality

The cardinality of elements in the element syntax sections can make use of regular expression wildcards with the following meanings:

Wildcard	Description
*	Zero or more occurrences
+	One or more occurrences
?	Zero or one occurrences

The absence of one of these symbols implies one and only one occurrence.

## 3.0 The Node Model

The primary element within the node model is a node. One can speak of some node properties as being a configured, available or utilized property of the node.

## 4.0 Node Element

The Node element is the root element of a node object and is used to encapsulate a node:

- A node object **MUST** have exactly one Node element.
- A compliant implementation **MUST** support this element.
- A node **MUST** specify one or more Node Properties.

## 4.1 Uncategorized Node Properties

Uncategorized Node Properties are properties that apply to the node as a whole and do not need to be distinguished between being configured, available or utilized. These include the node ID and other optional node properties.

### Simple Node Properties

Simple (unstructured) node properties are enumerated in the table below:

*Table M-7: Simple Node Properties*

Element Name	Type	Description	Appearance	Compliance
<b>Id</b>	String	Node identifier	MUST	MUST
<b>Name</b>	String	Node name or pattern	MAY	MAY
<b>OpSys</b>	String	Operating System	MAY	SHOULD
<b>Arch</b>	String	Architecture	MAY	SHOULD
<b>Description</b>	String	Description of the node	MAY	MAY
<b>State</b>	String	State of the node. Valid states can include Offline, Configured, Unknown, Idle, and Busy.	SHOULD	MUST
<b>Features</b>	String	Arbitrary named features of the node (comma-delimited string)	MAY	SHOULD

### Extension Element

The `Extension` element provides a means to pass extensible properties with the node object. Some applications might find it easier to deal with a named extension property than discover and handle elements for which they do not understand or anticipate by name.

- A compliant implementation **MAY** support this element.
- This element **MUST** have a name attribute that is of type `String` and represents the name of the extension property. A compliant implementation **MUST** support this attribute if this element is supported.

- This element MAY have a type attribute that is of type String and provides a hint about the context within which the property should be understood. A compliant implementation SHOULD support this attribute if this element is supported.
- The character content of this element is of type String and is the value of the extension property.

The following is an example of an `Extension` element:

```
<Extension type="Chemistry" name="Software">NWChem</Extension>
```

## 4.2 Property Categories

Certain node properties (particularly consumable resources) need to be classified as being in a particular category. This is done when it is necessary to distinguish between a property that is configured versus a property that is available or utilized. For example, a node might be configured with 16 processors. At a particular time, 8 might be utilized, 7 might be available and 1 disabled. When a node property must be categorized to be understood properly, the property MUST be enveloped within the appropriate Property Category Element.

### Configured Element

A configured node property reflects resources pertaining to the node that could in principle be used though they might not be available at this time. This information could be used to determine if a job could ever conceivably run on a given node.

- A compliant implementation MUST support this element.

The following is an example of using Configured Properties:

```
<Configured>
  <Processors>16</Processors>
  <Memory units="MB">512</Memory>
</Configured>
```

### Available Element

An available node property refers to a resource that is currently available for use.

- A compliant implementation SHOULD support this element.

The following is an example of specifying available properties:

```
<Available>
  <Processors>7</Processors>
  <Memory units="MB">256</Memory>
</Available>
```

## Utilized Element

A utilized node property reflects resources that are currently utilized.

- A compliant implementation SHOULD support this element.

The following is an example of specifying utilized properties:

```
<Utilized>
  <Processors>8</Processors>
  <Memory metric="Average">207</Memory>
</Utilized>
```

## 4.3 Categorized Node Properties

### Consumable Resources

Consumable Resources are a special group of node properties that can have additional attributes and can be used in multiple categories. In general a consumable resource is a resource that can be consumed in a measurable quantity.

- A consumable resource MUST be categorized as being a configured, available or utilized node property by being a child element of a Configured, Available or Utilized element respectively.
- A consumable resource MAY have a units attribute that is of type String that specifies the units by which it is being measured. If this attribute is omitted, a default unit is implied. A compliant implementation MAY support this attribute if the element is supported.
- A consumable resource MAY have a metric attribute that is of type String that specifies the type of measurement being described. For example, the measurement can be a Total, an Average, a Min or a Max. A compliant implementation MAY support this attribute if the element is supported.
- A consumable resource MAY have a wallDuration attribute of type Duration that indicates the amount of time for which that resource was used. This need only be specified if the resource was used for a different amount of time than the wallDuration for the step. A compliant implementation MAY support this attribute if the element is supported.
- A consumable resource MAY have a consumptionRate attribute of type Float that indicates the average percentage that a resource was used over its wallDuration. For example, an overbooked SMP running 100 jobs across 32 processors might want to scale the usage and charge by the average fraction of processor usage actually delivered. A compliant implementation MAY support this attribute if the element is supported.

Simple consumable resources are listed in the table below:

**Table M-8: Consumable Resource Node Properties**

Element Name	Type	Description	Appearance	Compliance	Categories
<b>Processors</b>	Integer	Number of processors	MAY	MUST	CAU
<b>Memory</b>	Float	Amount of memory	MAY	SHOULD	CAU
<b>Disk</b>	Float	Amount of disk	MAY	SHOULD	CAU
<b>Swap</b>	Float	Amount of virtual memory	MAY	MAY	CAU
<b>Network</b>	Float	Amount of network	MAY	MAY	CAU

The following are two examples for specifying a consumable resource:

```
<Memory metric="Max" units="GB">483</Memory>
<Processors wallDuration="1234" consumptionRate="0.63">4</Processors>
```

## Resource Element

In addition to the consumable resources enumerated in the above table, an extensible consumable resource is defined by the Resource element:

- A compliant implementation SHOULD support this element.
- This element MAY appear zero or more times within a given set of node properties.
- Like the other consumable resources, this property MUST be categorized as a configured, available or utilized property by being encompassed in the appropriate elements.
- This element is of type Float.
- It shares the other same properties and attributes as the other consumable resources but it requires an additional name (and optional type) attribute to describe it.
- This element MUST have a name attribute of type String that indicates the type of consumable resource being measured. A compliant implementation MUST support this attribute if the element is supported.



- This element MAY have a type attribute of type String that distinguishes it within a general resource class. A compliant implementation SHOULD support this attribute if the element is supported.

The following are two examples for specifying a Resource element:

```
<Resource name="License" type="MATLAB">1</Resource>
<Resource name="Telescope" type="Zoom2000" wallDuration="750"
metric="KX">10</Resource>
```

## 4.4 Node Reference

When a simple reference to a predefined node is needed in an encapsulating element, a Node element is used with the text content being the node ID:

```
<Node>node1</Node>
```

- This element MAY have an aggregation attribute of type String that provides a way to indicate multiple values with a single expression. A compliant implementation MAY support the aggregation attribute if the Feature element is supported. Possible values for this attribute include:
  - List a comma-separated list of features
  - Pattern a regular expression (perl5) matching desired features
  - Range a range of nodes of the form: <prefix>[5-23, 77]
- If an aggregation attribute is specified with the value of List, this element MAY also have a delimiter attribute of type String that indicates what delimiter is used to separate list elements. The default list delimiter is a comma.
- This element MAY have a count attribute of type Integer that indicates the instance count of the specified node(s).

The following is another example of a Node element:

```
<Node aggregation="Pattern">node[1-5]</Node>
```

## 5.0 Units of Measure Abbreviations

Abbreviation	Definition	Quantity
<b>B</b>	byte	1 byte
<b>KB</b>	Kilobyte	2 <sup>10</sup> bytes

Abbreviation	Definition	Quantity
<b>MB</b>	Megabyte	2 <sup>20</sup> bytes
<b>GB</b>	Gigabyte	2 <sup>30</sup> bytes
<b>TB</b>	Terabyte	2 <sup>40</sup> bytes
<b>PB</b>	Petabyte	2 <sup>50</sup> bytes
<b>EB</b>	Exabyte	2 <sup>60</sup> bytes
<b>ZB</b>	Aettabyte	2 <sup>70</sup> bytes
<b>YB</b>	Yottabyte	2 <sup>80</sup> bytes
<b>NB</b>	Nonabyte	2 <sup>90</sup> bytes
<b>DB</b>	Doggabyte	2 <sup>100</sup> bytes

## M.4 Scalable Systems Software Resource Management and Accounting Protocol (SSSRMAP) Wire Protocol

Resource Management Interface Specs  
Release v. 3.0.3  
13 May 2004

Scott Jackson  
Brett Bode  
David Jackson  
Kevin Walker

### Status of this Memo

This is a specification defining a wire level protocol used between Scalable Systems Software components. It is intended that this specification will continue to evolve as these interfaces are implemented and thoroughly tested by time and experience.

## Abstract

This document is a specification describing a connection-oriented XML-based application layer client-server protocol for the interaction of resource management and accounting software components developed as part of the Scalable Systems Software Center. The SSSRMAP Wire Protocol defines a framing protocol that includes provisions for security. The protocol is specified in XML Schema Definition and rides on the HTTP protocol.

## Table of Contents

- [1.0 Introduction](#)
- [2.0 Conventions Used in this Document](#)
  - [2.1 Keywords](#)
  - [2.2 XML Case Conventions](#)
  - [2.3 Schema Definitions](#)
- [3.0 Encoding](#)
  - [3.1 Schema Header and Namespaces](#)
  - [3.2 The Envelope Element](#)
  - [3.3 The Body Element](#)
- [4.0 Transport Layer](#)
- [5.0 Framing](#)
  - [5.1 Message Header Requirements](#)
  - [5.2 Message Chunk Format](#)
  - [5.3 Reply Header Requirements](#)
  - [5.4 Reply Chunk Format](#)
  - [5.5 Message and Reply Tail Requirements and Multiple Chunks](#)
  - [5.6 Examples](#)
    - [5.6.1 Sample SSSRMAP Message Embedded in HTTP Request](#)
    - [5.6.2 Sample SSSRMAP Reply Embedded in HTTP Response](#)
- [6.0 Asynchrony](#)
- [7.0 Security](#)

- 7.1 Security Token
  - 7.1.1 The SecurityToken Element
  - 7.1.2 Security Token Types
    - 7.1.2.1 Symmetric Key
    - 7.1.2.2 Asymmetric Key
    - 7.1.2.3 Password
    - 7.1.2.4 Cleartext
    - 7.1.2.5 Kerberos
    - 7.1.2.6 GSI (X.509)
  - 7.1.3 Example
- 7.2 Authentication
  - 7.2.1 The Signature Element
  - 7.2.2 The DigestValue Element
  - 7.2.3 The SignatureValue Element
  - 7.2.4 Signature Example
- 7.3 Confidentiality
  - 7.3.1 The EncryptedData Element
  - 7.3.2 The EncryptedKey Element
  - 7.3.3 The CipherValue Element
  - 7.3.4 Encryption Example
- 8.0 References

## 1.0 Introduction

A major objective of the Scalable Systems Software [SSS] Center is to create a scalable and modular infrastructure for resource management and accounting on terascale clusters including resource scheduling, grid-scheduling, node daemon support, comprehensive usage accounting and user interfaces emphasizing portability to terascale vendor operating systems. Existing resource management and accounting components feature disparate APIs (Application Programming Interfaces) requiring various forms of application coding to interact with other components.

This document proposes a wire level protocol expressed in an XML envelope to be considered as the foundation of a standard for communications between and among

resource management and accounting software components. Individual components additionally need to define the particular XML binding necessary to represent the message format for communicating with the component.

## 2.0 Conventions Used in this Document

### 2.1 Keywords

The keywords **MUST**, **MUST NOT**, **REQUIRED**, **SHALL**, **SHALL NOT**, **SHOULD**, **RECOMMENDED**, **MAY**, and **OPTIONAL** in this document are to be interpreted as described in [RFC2119](#).

### 2.2 XML Case Conventions

In order to enforce a consistent capitalization and naming convention across all SSSRMAP specifications 'Upper Camel Case' (UCC) and 'Lower Camel Case' (LCC) Capitalization styles shall be used. UCC style capitalizes the first character of each word and compounds the name. LCC style capitalizes the first character of each word except the first word. [XML\_CONV][FED\_XML]

1. SSSRMAP XML Schema and XML instance documents **SHALL** use the following conventions:
  - Element names **SHALL** be in UCC convention (example: `<UpperCamelCaseElement/>`).
  - Attribute names **SHALL** be in LCC convention (example: `<UpperCamelCaseElement lowerCamelCaseAttribute="Whatever"/>`).
2. General rules for all names are:
  - Acronyms **SHOULD** be avoided, but in cases where they are used, the capitalization **SHALL** remain (example: XMLSignature).
  - Underscores (`_`), periods (`.`) and dashes (`-`) **MUST NOT** be used (example: use JobId instead of JOB.ID, Job\_ID or job-id).

### 2.3 Schema Definitions

SSSRMAP Schema Definitions appear like this

In case of disagreement between the schema file and this specification, the schema file takes precedence.

## 3.0 Encoding

Encoding tells how a message is represented when exchanged. SSSRMAP data exchange messages SHALL be defined in terms of XML schema [XML\_SCHEMA].

## 3.1 Schema Header and Namespaces

The header of the schema definition is as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<schema
  xmlns="https://www.w3.org/201/XMLSchema"
  xmlns:sssrmmap="https://www.scidac.org/ScalableSystems/SSSRMAP"
  targetNamespace="https://www.scidac.org/ScalableSystems/SSSRMAP"
  elementFormDefault="qualified">
```

## 3.2 The Envelope Element

SSSRMAP messages and replies are encapsulated in the `Envelope` element. There are two possibilities for the contents of this element. If the contents are unencrypted, this element **MUST** contain a `Body` element and **MAY** contain a `Signature` element (refer to the section on [Security](#)). If the contents are encrypted, this element **MUST** contain exactly one `EncryptedData` element (refer to the section on [Security](#)). The `Envelope` element **MAY** contain namespace and other xsd-specific information necessary to validate the document against the schema. In addition, it **MAY** have any of the following attributes, which might serve as processing clues to the parser:

Attribute	Description
<b>type</b>	A message type providing a hint as to the body contents such as 'Request' or 'Notification'
<b>component</b>	A component type such as 'QueueManager' or 'LocalScheduler'
<b>name</b>	A component name such as 'OpenPBS' or 'Maui'
<b>version</b>	A component version such as '2.2p12' or '3.2.2'

```
<complexType name=EnvelopeType">
  <choice minOccurs="1" maxOccurs="1">
    <choice minOccurs="1" maxOccurs="2">
      <element ref="sssrmmap:Signature" minOccurs="0" maxOccurs="1"/>
      <element ref="sssrmmap:Body" minOccurs="1" maxOccurs="1"/>
    </choice>
    <element ref="sssrmmap:EncryptedData" minOccurs="1" maxOccurs="1"/>
  </choice>
  <attribute name="type" type="string" use="optional"/>
```

```

<attribute name="component" type="string" use="optional"/>
<attribute name="name" type="string" use="optional"/>
<attribute name="version" type="string" use="optional"/>
</complexType>

<element name="Envelope" type="ssrmap:EnvelopeType"/>

```

### 3.3 The Body Element

- SSSRMAP messages and replies are encapsulated in the `Body` element. This element **MUST** contain exactly one `Request` or `Response` element.

```

<complexType name="BodyType">
  <choice minOccurs="1" maxOccurs="1">
    <element ref="ssrmap:Request" minOccurs="0" maxOccurs="1"/>
    <element ref="ssrmap:Response" minOccurs="0" maxOccurs="1"/>
    <any minOccurs="0" maxOccurs="1" namespace="##other"/>
  </choice>
</complexType>

<element name="Body" type="ssrmap:BodyType"/>

```

### 4.0 Transport Layer

This protocol will be built over the connection-oriented reliable transport layer TCP/IP. Support for other transport layers could also be considered, but native support for TCP/IP can be found on most terascale clusters and automatically handles issues such as reliability and connection fullness for the application developer implementing the SSSRMAP protocol.

### 5.0 Framing

Framing specifies how the beginning and ending of each message is delimited. Given that the encoding will be expressed as one or more XML documents, clients and servers need to know when an XML document has been fully read in order to be parsed and acted upon.

SSSRMAP uses the HTTP 1.1 [HTTP] protocol for framing. HTTP uses a byte-counting mechanism to delimit the message segments. HTTP chunked encoding is used. This allows for optional support for batched messages, large message segmentation and persistent connections.

### 5.1 Message Header Requirements

The HTTP request line (first line of the HTTP request header) begins with POST and is followed by a URI and the version of the HTTP protocol that the client understands. It is suggested for this protocol that the URI consist of a single slash, followed by the protocol name in uppercase (i.e., /SSSRMAP), though this field is not checked and could be empty, a single slash or any URI.

The Content-Type must be specified as `test/xml`. Charset can be optionally specified and defaults to US-ASCII. It is recommended that charset be specified as 'utf-8' for maximum interoperability.

The Transfer-Encoding must be specified as `chunked`. The Content-Length must NOT be specified as the chunk size is specified in the message chunk.

Other properties such as User-Agent, Host and Date are strictly optional.

## 5.2 Message Chunk Format

A message chunk consists of a chunk size in hexadecimal format (whose value is the number of bytes in the XML message not including the chunk size and delimiter) delimited by a CR/LF “\r\n” and followed by the message payload in XML that consists of a single XML document having a root element of `Envelope`.

## 5.3 Reply Header Requirements

The HTTP response line (first line of the HTTP response header) begins with HTTP and a version number, followed by a numeric code and a message indicating what sort of response is made. These response codes and messages indicate the status of the entire response and are as defined by the HTTP standard. The most common response is 200 OK, indicating that the message was received and an appropriate response is being returned.

The Content-Type must be specified as `text/xml`. Charset can be optionally specified and defaults to US-ASCII. It is recommended that charset be specified as 'utf-8' for maximum interoperability.

The Transfer-Encoding MUST be specified as `chunked`. The Content-Length must NOT be specified.

Other properties such as Server, Host and Date are strictly optional.

## 5.4 Reply Chunk Format

A reply chunk consists of a chunk size in hexadecimal format (whose value is the number of bytes in the XML reply not including the chunk size and delimiter) delimited by a CR/LF “\r\n” and followed by the reply payload in XML that consists of a single XML document having a root element of `Envelope`.

## 5.5 Message and Reply Tail Requirements and Multiple Chunks

This specification only requires that single chunks be supported. A server can optionally be configured to handle requests with persistent connections (multiple chunks). It will be the



responsibility of clients to know whether a particular server supports this additional functionality. After all chunks have been sent, a connection is terminated by sending a zero followed by a carriage return-linefeed combination (0\r\n) and closing the connection.

## 5.6 Examples

### Sample SSSRMAP Message Embedded in HTTP Request

```
POST /SSSRMAP HTTP/1.1\r\n
Content-Type: text/xml; charset="utf-8"\r\n
Transfer-Encoding: chunked\r\n
\r\n
9A\r\n
<Envelope .../>
0\r\n
```

### Sample SSSRMAP Reply Embedded in HTTP Response

```
HTTP/1.1 200 OK\r\n
Content-Type: text/xml; charset="utf-8"\r\n
Transfer-Encoding: chunked\r\n
\r\n
2B4\r\n
<Envelope .../>
0\r\n
```

## 6.0 Asynchrony

Asynchrony (or multiplexing) allows for the handling of independent exchanges over the same connection. A widely-implemented approach is to allow pipelining (or boxcarring) by aggregating requests or responses within the body of the message or via persistent connections and chunking in HTTP 1.1. Pipelining helps reduce network latency by allowing a client to make multiple requests of a server, but requires the requests to be processed serially [RFC3117]. Parallelism could be employed to further reduce server latency by allowing multiple requests to be processed in parallel by multi-threaded applications.

Segmentation can become necessary if the messages are larger than the available window. With support for segmentation, the octet-counting requirement that you need to know the length of the whole message before sending it can be relegated to the segment level – and you can start sending segments before the whole message is available. Segmentation is facilitated via 'chunking' in HTTP 1.1.

The current SSSRMAP strategy supports only a single request or response within the Body element. A server can optionally support persistent connections from a client via HTTP chunking. Segmentation of large responses is also optionally supported via HTTP chunking.

Later versions of the protocol could allow pipelined requests and responses in a single Body element.

## 7.0 Security

SSSRMAP security features include capabilities for integrity, authentication, confidentiality, and non-repudiation. The absence or presence of the various security features depend upon the type of security token used and the protection methods you choose to specify in the request.

For compatibility reasons, SSSRMAP specifies six supported security token types. Extensibility features are included allowing an implementation to use alternate security algorithms and security tokens. It is also possible for an implementation to ignore security features if it is deemed nonessential for the component. However, it is highly RECOMMENDED that an implementation support at least the default security token type in both authentication and encryption.

### 7.1 Security Token

A security token can be included in either the Signature block, and/or in the EncryptedData block (both described later) as an implicit or explicit cryptographic key. If this element is omitted, the security token is assumed to be a secret key shared between the client and the server.

#### The SecurityToken Element

This element is of type String. If the security token conveys an explicit key, this element's content is the value of the key. If the key is natively expressed in a binary form, it must be converted to base64 encoding as defined in XML Digital Signatures (<https://www.w3.org/2000/09/xmldsig#base64>). If the type is not specified, it is assumed to be of type 'Symmetric'.

It can have any of the following optional attributes:

Attribute	Description
<b>type</b>	<p>The type of security token (described subsequently):</p> <ul style="list-style-type: none"> <li>A <code>type</code> attribute of 'Symmetric' specifies a shared secret key between the client and server. This is the default.</li> <li>A <code>type</code> attribute of 'Asymmetric' specifies the use of public private key pairs between the client and server.</li> <li>A <code>type</code> attribute of 'Password' encrypts and authenticates with a user password known to both the client and server.</li> </ul>

Attribute	Description
	<ul style="list-style-type: none"> <li>A <code>type</code> attribute of 'Cleartext' allows the passing of a cleartext username and password and depends on the use of a secure transport (such as SSL or IPsec).</li> <li>A <code>type</code> attribute of 'Kerberos5' specifies a kerberos token.</li> <li>A <code>type</code> attribute of 'X509v3' specifies an X.509 certificate.</li> </ul>
<b>name</b>	The name of the security token that serves as an identifier for the actor making the request (useful when the key is a password, or when the key value is implicit as when a public key is named but not included).

```

<complexType name="SecurityTokenType" mixed="true">
  <simpleContent>
    <extension base="string">
      <attribute name="type" type="string" use="optional">
        <attribute name="name" type="string" use="optional">
      </extension>
    </simpleContent>
  </complexType>

<element name="SecurityToken" type="sssrmmap:SecurityTokenType"/>

```

## Security Token Types

SSSRMAP defines six standard security token types:

### Symmetric Key

The default security token specifies the use of a shared secret key. The secret key is up to 128-bits long and known by both client and server. When using a symmetric key as a security token, it is not necessary to specify the `type` attribute with value 'Symmetric' because this is assumed when the attribute is absent. The `name` attribute should be specified indicating the actor issuing the request. If the user provides a password to be sent to the server for authentication, then the password is encrypted with the secret key using a default `method="kw-tripledes"` (XML ENCRYPTION <https://www.w3.org/TR/xmlenc-core1/>), base64 encoded and included as the string content of the `SecurityToken` element. If the client authenticated the user, then the `SecurityToken` element is empty. The same symmetric key is used in both authentication and encryption.

### Asymmetric Key

Public and private key pairs can be used to provide non-repudiation of the client (or server). The client and the server must each have their own asymmetric key pairs. This mode is indicated by specifying the `type` attribute as 'Asymmetric'. The `name` attribute should be specified indicating the actor issuing the request. If the user provides a password to be sent to the server for authentication, then the password is encrypted with the server's public key using a default `method="rsa-1_5"` (XML ENCRYPTION <https://www.w3.org/TR/xmlenc-core1/>), base64 encoded and included as the string

content of the *SecurityToken* element. If the client authenticated the user, then the *SecurityToken* element is empty. The sender's private key is used in authentication (signing) while the recipient's public key is used for encryption.

### Password

This mode allows for a username password combination to be used under the assumption that the server also knows the password for the user. This security token type is indicated by specifying a value of 'Password' for the *type* attribute. The password itself is used as the cryptographic key for authentication and encryption. The *name* attribute contains the user name of the actor making the request. The *SecurityToken* element itself is empty.

### Cleartext

This security mode is equivalent to passing the username and password in the clear and depends upon the use of a secure transport (such as SSL or IPsec). The purpose of including this security token type is to enable authentication to occur from web browsers over SSL or over internal LANs who use IPsec to encrypt all traffic. The password (or a hash of the password like in /etc/passwd) would have to be known by the server for authentication to occur. In this mode, neither encryption nor signing of the hash is performed at the application layer. This mode is indicated by specifying a value of 'Cleartext' for the *type* attribute. The *name* attribute contains the user name of the actor making the request and the string content of the *SecurityToken* element is the unencrypted plaintext password.

### Kerberos

The use of a Kerberos version 5 token is indicated by specifying 'Kerberos5' in the *type* attribute. The *name* attribute is used to contain the kerberos user ID of the actor making the request. The *SecurityToken* element contains two sub elements. The *Authenticator* element contains the authenticator encoded in base64. A *Ticket* element contains the service-granting ticket, also base64 encoded.

### GSI (X.509)

The Grid Security Infrastructure (GSI), which is based on public key encryption, X.509 certificates, and the Secure Sockets Layer (SSL) communication protocol can be indicated by specifying a *type* attribute of 'X509v3'. The *name* attribute contains the userid used that the actor was mapped to in the local system. The string content of the *SecurityToken* element is the GSI authentication message including the X.509 identity of the sender encoded in base64.

### Example

```
<SecurityToken type="Asymmetric" name="scottmo">
MIIEZzCCA9CggAwIBAgIQEmtJZc0rqrKh5i...
</SecurityToken>
```

## 7.2 Authentication

Authentication entails how the peers at each end of the connection are identified and verified. Authentication is optional in an SSSRMAP message or reply. SSSRMAP uses a digital signature scheme for authentication that borrows from concepts in XML Digital Signatures [XML\_DSIG]. In addition to authentication, the use of digital signatures also ensures integrity of the message, protecting exchanges from third-party modification.

When authentication is used, a `Signature` element is prepended as the first element within the `Envelope` element. All of the security modes will create a digest of the data for integrity checking and store this in base64 encoding in a `DigestValue` element as a child of the `Signature` element. The digital signature is created by encrypting the hash with the appropriate security token and storing this value in a `SignatureValue` element as a child of the `Signature` element. The security token itself is included as a child of the `Security` element within a `SecurityToken` element.

There are a number of procedural practices that must be followed in order to standardize this approach. The digest (or hash) is created over the contents of the `Envelope` element (not including the `Element` tag or its attributes). This might be over one or more `Request` or `Notify` elements (or `Response` or `Ack` elements) and necessarily excludes the `Signature` Element itself. (Note that any data encryption is performed after the creation of the digital signature and any decryption is performed before authenticating so the `EncryptedData` element will not interfere with this process. Therefore, the signature is always based on the (hashed but) unencrypted data). For the purposes of generating the digest over the same value, it is assumed that the data is first canonicalized to remove extraneous whitespace, comments, etc according to the XML Digital Signature algorithm (<https://www.w3.org/TR/2001/REC-xml-c14n-20010315>) and a transform is applied to remove namespace information. As a rule, any binary values are always transformed into their base64 encoded values when represented in XML.

### The Signature Element

The `Signature` element **MUST** contain a `DigestValue` element that is used for integrity checking. It **MUST** also contain a `SecurityToken` element that is used to indicate the security mode and token type, and to verify the signature. It **MUST** contain a `SignatureValue` element that contains the base64 encrypted value of the signature wrought on the hash **UNLESS** the security token type indicates `Cleartext` mode where a signature would be of no value with the encryption key being sent in the clear -- in this case we use the password itself for authentication).

```
<complexType name="SignatureType">
  <choice minOccurs="2" maxOccurs="3">
    <element ref="sssrmmap:DigestValue" minOccurs="1" maxOccurs="1"/>
    <element ref="sssrmmap:SignatureValue" minOccurs="1" maxOccurs="1"/>
    <element ref="sssrmmap:SecurityToken" minOccurs="0" maxOccurs="1"/>
  </choice>
</complexType>

<element name="Signature" type="sssrmmap:SignatureType"/>
```

## The DigestValue Element

The `DigestValue` element contains the cryptographic digest of the message data. As described above, the hash is generated over the `Body` element. The data to be hashed must first be canonicalized and appropriately transformed before generating the digest since typically an application will read in the XML document into an internal binary form, then marshal (or serialize) the data into a string, which is passed as input to the hash algorithm. Different implementations marshal the data differently so it is necessary to convert this to a well-defined format before generating the digest or the clients will generate different digest values for the same XML. The SHA-1 [SHA-1] message digest algorithm (<https://www.w3.org/2000/09/xmlsig#sha1>) SHALL be used as the default method for generating the digest. A `method` attribute is defined as an extensibility option in case an implementation wants to be able to specify alternate message digest algorithms.

It MAY have a `method` attribute:

Attribute	Description
<b>method</b>	<p>The message digest algorithm.</p> <ul style="list-style-type: none"> <li>A <code>method</code> attribute of 'sha1' specifies the SHA-1 message digest algorithm. This is the default and is implied if this attribute is omitted.</li> </ul>

```
<complexType name="DigestValueType">
  <simpleContent>
    <extension base="string">
      <attribute name="method" type="string" use="optional"/>
    </extension>
  </simpleContent>
</complexType>

<element name="DigestValue" type="sssrmap:DigestValueType"/>
```

## The SignatureValue Element

The `SignatureValue` element contains the digital signature that serves the authentication (and potentially non-repudiation) function. The string content of the `SignatureValue` element is a base64 encoding of the encrypted digest value. The HMAC algorithm [HMAC] based on the SHA1 message digest (<https://www.w3.org/2000/09/xmlsig#hmac-sha1>) SHALL be used as the default message authentication code algorithm for user identification and message integrity. A `method` attribute is defined as an extensibility option in case an implementation wants to be able to specify alternate digital signature algorithms.

It MAY have a `method` attribute:

Attribute	Description
<b>method</b>	The digest signature algorithm.

Attribute	Description
	<ul style="list-style-type: none"> <li>A method attribute of 'hmac-sha1' specifies the HMAC SHA-1 digital signature algorithm. This is the default and is implied if this attribute is omitted.</li> </ul>

```

<complexType name="SignatureValueType">
  <simpleContent>
    <extension base="string">
      <attribute name="method" type="string" use="optional"/>
    </extension>
  </simpleContent>
</complexType>

<element name="SignatureValue" type="sssrmap:SignatureValueType"/>

```

## Signature Example

### Pre-authentication:

```

<Envelope>
  <Body>
    <Request action="Query" actor="kenneth">
      <Object>User</Object>
      <Get name="EmailAddress"></Get>
      <Where name="Name">scott</Where>
    </Request>
  </Body>
</Envelope>

```

### Post-authentication:

```

<Envelope>
  <Signature>
    <DigestValue>
      LyLsF0Pi4wPU...
    </DigestValue>
    <SignatureValue>
      DJbchm5gK...
    </SignatureValue>
    <SecurityToken type="Asymmetric" name="kenneth">
      MIIEZzCCA9CggAwIBAgIQEmtJZc0rqKh5i...
    </SecurityToken>
  </Signature>
  <Body>
    <Request action="Query" actor="kenneth">
      <Object>User</Object>
      <Get name="EmailAddress"></Get>
      <Where name="Name">scottmo</Where>
    </Request>
  </Body>
</Envelope>

```

## 7.3 Confidentiality

Confidentiality involves encrypting the sensitive data in the message, protecting exchanges against third-party interception and modification. Confidentiality is optional in an SSSRMAP message or reply. When confidentiality is required, SSSRMAP sessions use block cipher encryption with concepts borrowed from the emerging XML Encryption [XML\_ENC] standard.

When confidentiality is used, encryption is performed over all child elements of the `Envelope` element (i.e., on the message data, as well as any signature). The encrypted data is not signed -- rather the signature is encrypted. This data is replaced in-place within the envelope with an `EncryptedData` element. The data is first compressed using the gzip algorithm [ZIP]. Instead of encrypting this compressed data with the security token directly, a 192-bit random session key is generated by the sender and used to perform symmetric encryption on the compressed data. This key is itself encrypted with the security token and included with the encrypted data as the value of the `EncryptedKey` element as a child of the `EncryptedData` element. The ciphertext resulting from the data being encrypted with the session key is passed as the value of a `CipherValue` element (also a child of the `EncryptedData` element). As in the case with authentication, the security token itself is included as a child of the `Security` element within a `SecurityToken` element.

### The EncryptedData Element

When SSSRMAP confidentiality is required, the `EncryptedData` element MUST appear as the only child element in the `Envelope` element. It directly replaces the contents of these elements including the data and any digital signature. It MUST contain an `EncryptedKey` element that is used to encrypt the data. It MUST contain a `CipherValue` element that holds the base64 encoded ciphertext. It MAY also contain a `SecurityToken` element that is used to indicate the security mode and token type. If the `SecurityToken` element is omitted, a Symmetric key token type is assumed. Confidentiality is not used when a security token type of 'Cleartext' is specified since it would be pointless to encrypt the data with the encryption key in the clear.

```
<complexType name="EncryptionDataType">
  <choice minOccurs="0" maxOccurs="1">
    <element ref="sssrmmap:EncryptedKey" minOccurs="1" maxOccurs="1"/>
    <element ref="sssrmmap:CipherValue" minOccurs="1" maxOccurs="1"/>
    <element ref="sssrmmap:SecurityToken" minOccurs="1" maxOccurs="1"/>
  </choice>
</complexType>

<element name="EncryptedData" type="sssrmmap:EncryptionDataType"/>
```

### The EncryptedKey Element

The `EncryptedKey` element is a random session key encrypted with the security token. This approach is used for a couple of reasons. In the case where public key encryption is



used, asymmetric encryption is much slower than symmetric encryption and it makes sense to use a symmetric key for encryption and pass along it along by encrypting it with the recipient’s public key. It is also useful in that the security token, which does not change very often (compared to the session key, which changes for every connection) is used on a very small sampling of data (the session key), whereas if it was used to encrypt the whole message an attacker could more effectively exploit an attack against the ciphertext. The CMS Triple DES Key Wrap algorithm 'kw-tripledes' SHALL be used as the default method for key encryption. The session key is encrypted using the security token, base64 encoded and specified as the string content of the `EncryptedKey` element. A `method` attribute is defined as an extensibility option in case an implementation wants to be able to specify alternate key encryption algorithms.

It is REQUIRED that an implementation use a cryptographically secure Pseudo-Random number generator. It is RECOMMENDED that the session key be cryptographically generated (such as cyclic encryption, DES OFB, ANSI X9.17 PRNG, SHA1PRNG, or ANSI X12.17 (used by PGP)).

It MAY have a `method` attribute:

Attribute	Description
<b>method</b>	<p>The key encryption algorithm.</p> <ul style="list-style-type: none"> <li>A <code>method</code> attribute of 'kw-tripledes' specifies the CMS Triple DES Key Wrap algorithm. This algorithm is specified by the XML Encryption [XML_ENC] URI <a href="https://www.w3.org/TR/xmlenc-core1/">https://www.w3.org/TR/xmlenc-core1/</a>. It involves two Triple DES encryptions, a random and known Initialization Vector (IV) and a CMS key checksum. A 192-bit key encryption key is generated from the security token, lengthened as necessary by zero-padding. No additional padding is performed in the encryptions. This is the default and is implied if this attribute is omitted.</li> </ul>

```

<complexType name="EncryptedKeyType">
  <simpleContent>
    <extension base="string">
      <attribute name="method" type="string" use="optional"/>
    </extension>
  </simpleContent>
</complexType>

<element name="EncryptedKey" type="sssrmmap:EncryptedKeyType"/>

```

## The CipherValue Element

The `CipherValue` element contains the message (and possibly signature) data encrypted with the random session key. The ciphertext is compressed using the gzip algorithm [ZIP], encrypted by the designated method, base64 encoded and included as the string content of the `CipherValue` element. The Triple DES algorithm with Cipher Block Chaining (CBC) feedback mode SHALL be used as the default method for encryption. A `method` attribute is

defined as an extensibility option in case an implementation wants to be able to specify alternate data encryption algorithms.

It MAY have a method attribute:

Attribute	Description
<b>method</b>	<p>The data encryption algorithm.</p> <ul style="list-style-type: none"> <li>A <code>method</code> attribute of 'tripleDES-cbc' specifies the Triple DES algorithm with Cipher Block Chaining (CBC) feedback mode. This algorithm is specified by the XML Encryption [XML_ENC] URI identifier <a href="https://www.w3.org/TR/xmlenc-core1/">https://www.w3.org/TR/xmlenc-core1/</a>. It specifies the use of a 192-bit encryption key and a 64-bit Initialization Vector (IV). Of the key bits, the first 64 are used in the first DES operation, the second 64 bits in the middle DES operation, and the third 64 bits in the last DES operation. The plaintext is first padded to a multiple of the block size (8 octets) using the padding scheme described in [XMLENC] for Block Encryption Algorithms (Padding per PKCS #5 will suffice for this). The resulting cipher text is prefixed by the IV. This is the default and is implied if this attribute is omitted.</li> </ul>

```
<complexType name="CipherValueType">
  <simpleContent>
    <extension base="string">
      <attribute name="method" type="string" use="optional"/>
    </extension>
  </simpleContent>
</complexType>

<element name="CipherValue" type="sssrmmap:CipherValueType"/>
```

## Encryption Example

In this example, a simple request is demonstrated without a digital signature for the sake of emphasizing the encryption plaintext replacement.

Pre-encryption:

```
<Envelope>
  <Body>
    <Response>
      <Status>true</Status>
      <Code>000</Code>
      <Count>1</Count>
      <Data>
        <User>
          <EmailAddress>Scott.Jackson@pnl.gov</EmailAddress>
        </User>
      </Data>
    </Response>
  </Body>
</Envelope>
```

Post-encryption:

```

<Envelope>
  <EncryptedData>
    <EncryptedKey>
      NAKe9iQofYhyOfiHZ29kkEFVJ30CAwEAAaMSM...
    </EncryptedKey>
    <CipherValue>
      mPCadVfOMx1NzDaKMHNgFkR9upTW4kgBxyPW...
    </CipherValue>
    <SecurityToken type="Asymmetric" name="kenneth">
      MIEZzCCA9CggAwIBAgIQEmtJZc0rqrKh5i...
    </SecurityToken>
  </EncryptedData>
</Envelope>

```

## 8.0 References

- [RFC2119] S. Bradner, "Key Words for Use in RFCs to Indicate Requirement Levels", [RFC 2119](#), March 1997.
- [BEEP] M. Rose, "The Blocks Extensible Exchange Protocol Core", [RFC 3080](#), March 2001.
- [HMAC] H. Krawczyk, M. Bellare, R. Canetti, "HMAC, Keyed-Hashing for Message Authentication", [RFC 2104](#), February 1997.
- [SHA-1] U.S. Department of Commerce/National Institute of Standards and Technology, "[Secure Hash Standard](#)", FIPS PUB 180-1.
- [SSS] "Scalable Systems Software", <https://www.scidac.org/ScalableSystems>
- [HTTP] "Hypertext Transfer Protocol – HTTP/1.1", [RFC 2616](#), June 1999.
- [XML\_CONV] "[I-X and <I-N-CA> XML Conventions](#)".
- [FED\_XML] "[U.S. Federal XML Guidelines](#)".
- [RFC3117] M. Rose, "On the Design of Application Protocols", [Informational RFC 3117](#), November 2001.
- [XML\_DSIG] D. Eastlake, J. Reagle Jr., D. Solo, "[XML Signature Syntax and Processing](#)", W3C Recommendation, 12 February 2002.
- [XML\_ENC] T. Imamura, B. Dillaway, E. Smon, "[XML Encryption Syntax and Processing](#)", W3C Candidate Recommendation, 4 March 2002.
- [XRP] E. Brunner-Williams, A. Damaraju, N. Zhang, "[Extensible Registry Protocol \(XRP\)](#)", Internet Draft, expired August 2001.
- [XML] Bray, T., et al, "[Extensible Markup Language \(XML\) 1.0 \(Second Edition\)](#)", 6 October 2000.
- [XML\_SCHEMA] D. Beech, M. Maloney, N. Mendelsohn, "[XML Schema Part 1: Structures Working Draft](#)", April 2000.
- [ZIP] J. Gailly, M. Adler, "The gzip home page", <https://www.gzip.org/>



## Appendix N: Moab Resource Manager Language Interface Overview

The Moab RM Language (formerly called WIKI) is the language that some resource managers use to communicate with Moab, specifically a native RM. Generally each line represents a single resource or workload in Moab. The line contains the name of the resource or workload followed by a set of `<attr>=<val>` pairs. Although the Moab RM language follows the same data format for all RMs, each RM type receives and returns it differently. For instructions and examples on using Moab RM language with a native RM, see and [Managing Resources Directly with the Native Interface](#) respectively.

In this section:

- [N.1 Moab RM Language Socket Protocol Description - page 1505](#)
- [N.2 Moab Resource Manager Language Data Format - page 1509](#)

### N.1 Moab RM Language Socket Protocol Description

In this topic:

- [N.1.1 RM Language Overview - page 1505](#)
- [N.1.2 Checksum Algorithm \('C' Version\) - page 1506](#)
- [N.1.3 Header Creation \(PERL Code\) - page 1508](#)
- [N.1.4 Header Processing \(PERL Code\) - page 1508](#)

#### N.1.1 RM Language Overview

Moab RM language is formerly known as WIKI. The Moab scheduler uses a simple protocol for socket connections to the user client and the resource manager as described below:

```
<SIZE><CHAR>
```

```
CK=<CKSUM><WS>TS=<TIMESTAMP><WS>AUTH=<AUTH><WS>DT=<DATA>
```

Attribute	Description
<code>&lt;SIZE&gt;</code>	8 character decimal ASCII representation of the size of the packet following

Attribute	Description
	'<SIZE><CHAR>' Leading zeroes must be used to pad this value to 8 characters if necessary.
<CHAR>	A single ASCII character
<CKSUM>	A 16 character hexadecimal ASCII DES-based checksum calculated using the algorithm below* and <SEED> selected and kept secret by the site admins. The checksum is performed on the line from TS= to the end of the message including <DATA>.
<WS>	a series of white space characters consisting of either tabs and/or space characters.
<TIMESTAMP>	ASCII representation of epoch time
<AUTH>	Identifier of user requesting service (i.e., USERNAME)
<DT>	Data to be sent

An example header follows:

```
00001057 CK=cdf6d7a7ad45026f TS=922401962 AUTH=sched DT=<DATA>
```

where <DATA> is replaced by actual message data.

## N.1.2 Checksum Algorithm ('C' Version)

### *Checksum Algorithm ('C' Version)*

```
#define MAX_CKSUM_ITERATION 4

int GetChecksum(
    char *Buf,
    int BufSize,
    char *Checksum,
    char *CSKey) /* Note: pass in secret key */
{
    unsigned int crc;
    unsigned int lword;
    unsigned int irword;
    int index;
    unsigned int Seed;
    Seed = (unsigned int)strtoul(CSKey, NULL, 0);
    crc = 0;
    for (index = 0; index < BufSize; index++)
    {
```

```

crc = (unsigned int)DoCRC((unsigned short)crc,Buf[index]);
}
lword = crc;
irword = Seed;
PSDES(&lword,&irword);
sprintf(Checksum,"%08x%08x",
lword,
irword);
return(SUCCESS);
}

unsigned short DoCRC(
unsigned short crc,
unsigned char onech)
{
int index;
unsigned int ans;
ans = (crc ^ onech << 8);
for (index = 0;index < 8;index++)
{
if (ans & 0x8000)
ans = (ans <<= 1) ^ 4129;
else
ans <<= 1;
}
return((unsigned short)ans);
}

int PSDES(
unsigned int *lword,
unsigned int *irword)
{
int index;
unsigned int ia;
unsigned int ib;
unsigned int iswap;
unsigned int itmph;
unsigned int itmpl;
static unsigned int c1[MAX_CKSUM_ITERATION] = {
0xcba4e531, 0x537158eb, 0x145cdc3c, 0x0d3fdeb2 };
static unsigned int c2[MAX_CKSUM_ITERATION] = {
0x12be4590, 0xab54ce58, 0x6954c7a6, 0x15a2ca46 };
itmph = 0;
itmpl = 0;
for (index = 0;index < MAX_CKSUM_ITERATION;index++)
{
iswap = *irword;
ia = iswap ^ c1[index];
itmpl = ia & 0xffff;
itmph = ia >> 16;
ib = (itmpl * itmpl) + ~(itmph*itmph);
ia = (ib >> 16) | ((ib & 0xffff) << 16);
*irword = (*lword) ^ ((ia ^ c2[index]) + (itmpl * itmph));
*lword = iswap;
}
return(SUCCESS);
}

```

## N.1.3 Header Creation (PERL Code)

(taken from PNNL's QBank client code)

### Header Creation (PERL Code)

```
#####
#
# subroutine wiki($COMMAND)
#
# Sends command to Moab server and returns the parsed result and status
#
#####
sub wiki
{
  my($COMMAND,$REQUEST,$result);
  my($sockaddr,$hostname);
  my($name,$aliases,$proto,$port,$type,$len,$thisaddr);
  my($thisport,$thatport,$response,$result);
  $COMMAND = shift;
  #
  # Establish socket connection
  #
  $sockaddr = 'S n a4 x8';
  chop ($hostname = `hostname`);
  ($name,$aliases,$proto)=getprotobyname('tcp');
  ($name,$aliases,$type,$len,$thisaddr)=gethostbyname($hostname);
  ($name,$aliases,$type,$len,$thataddr)=gethostbyname($BANKHOST);
  $thisport=pack($sockaddr, &AF_INET,0,$thisaddr);
  $thatport=pack($sockaddr, &AF_INET,$BANKPORT,$thataddr);
  socket(S, &PF_INET,&SOCK_STREAM,$proto) || die "cannot create socket\n";
  bind(S,$thisport) || die "cannot bind socket\n";
  connect(S,$thatport) || die "cannot connect socket\n";
  select(S); $| = 1; # Turn on autoflushing
  select(stdout); $| = 1; # Select STDOUT as default output
  #
  # Build and send command
  #
  $REQUEST="COMMAND=$COMMAND AUTH=$AUTH";
  chomp($CHECKSUM = `QSUM "$REQUEST"`);
  $REQUEST .= " CHECKSUM=$CHECKSUM";
  my $command=pack "a8 a1 A*",sprintf("%08d",length($REQUEST))," ",$REQUEST;
  print S "$command"; # Send Command to server
  @REPLY=();
  while ( ) { push(@REPLY,$_); } # Listen for Reply
  
  $STATUS=grep(/STATUSCODE=(\d*)/,&&$1,@REPLY); # STATUSCODE stored in $STATUS
  grep(s/.*RESULT=//,@REPLY); # Parse out the RESULT
  return @REPLY;
  
}
```

## N.1.4 Header Processing (PERL Code)

### Header Processing (PERL Code)

```

sysread(NS,$length,8); # Read length string
sysread(NS,$delimiter,1); # Read delimiter byte
$DEBUG && print STDERR "length=[$length]\tdelimiter=[$delimiter]\n";
while($length) {
```



```

$DEBUG && print STDERR "Awaiting $length bytes -- ".`date`;
$length==sysread(NS,$request,$length); # Read request
sleep 1;
}
%REQUEST=();
chomp($request);
foreach (@REQUEST=&shellwords($request)) # Parse arguments into array
{
($key,$value)=split(/=/,$_);
$REQUEST{$key}=$value unless defined $REQUEST{$key};
}
$request =~ s/\s+CHECKSUM=.*/; # Strip off the checksum
print STDERR "REQUEST=$request\n";
chomp($checksum=`QSUM "$request"`);
$me=$REQUEST{AUTH};
$command=$REQUEST{COMMAND};
if (!grep($command eq $_,@VALIDCMDS))
{ $REPLY = "STATUSCODE=0 RESULT=$command is not a valid command\n";}
elsif ($checksum ne $REQUEST{CHECKSUM})
{ $REPLY = "STATUSCODE=0 RESULT=Invalid Checksum\n";}
else
{ $REPLY = do $command(@REQUEST); }
$len=sprintf("%08d",length($REPLY)-1);
$delim=' ';
$DEBUG && print STDERR "REPLY=${len}${delim}$REPLY\n";
$buf="$len"."$delim"."$REPLY";
syswrite(NS,$buf,length($buf));
close NS;

```

## N.2 Moab Resource Manager Language Data Format

In this topic:


- [N.2.1 Query Resources Data Format - page 1509](#)
- [N.2.2 Query Workload Data Format - page 1514](#)

### N.2.1 Query Resources Data Format

#### *Query Resources Data Format*

Name	Format	Default	Description
<b>ADISK</b>	<INTEGER>	0	Available local disk on node (in MB)
<b>AFS</b>	<fs id="X" size="X" io="Y" rcount="X" wcount="X"	0	Available filesystem state

Name	Format	Default	Description
	ocount="X"></fs>[...]		
<b>AMEMORY</b>	<INTEGER>	0	Available/free RAM on node (in MB)
<b>APROC</b>	<INTEGER>	1	Available processors on node
<b>ARCH</b>	<STRING>	---	Compute architecture of node
<b>ARES</b>	One or more comma-delimited <NAME>:<VALUE> pairs (i.e., MATLAB:6,COMPILER:100 )	---	Arbitrary consumable resources currently available on the node
<b>ASWAP</b>	<INTEGER>	0	Available swap on node (in MB)
<b>CCLASS</b>	One or more bracket enclosed <NAME>:<COUNT> pairs (i.e., [batch:5] [sge:3])	---	Run classes supported by node. Typically, one class is 'consumed' per task. Therefore, an 8 processor node may have 8 instances of each class it supports present, ie [batch:8] [interactive:8]
<b>CDISK</b>	<INTEGER>	0	Configured local disk on node (in MB)
<b>CFS</b>	<STRING>	0	Configured filesystem state
<b>CMEMORY</b>	<INTEGER>	0	Configured RAM on node (in MB)
<b>CONTAINERNODE</b>	<STRING>	---	The physical machine that is hosting the virtual machine. Only valid on VMs.

Name	Format	Default	Description
<b>CPROC</b>	<INTEGER>	1	Configured processors on node
<b>CPULOAD</b>	<DOUBLE>	0.0	One minute BSD load average
<b>CPUSPEED</b>	<INTEGER>	---	The node's processor speed in MHz
<b>CRES</b>	One or more comma-delimited <NAME>:<VALUE> pairs (i.e., MATLAB:6,COMPILER:100 )	---	Arbitrary consumable resources supported and tracked on the node, ie software licenses or tape drives
<b>CSWAP</b>	<INTEGER>	0	Configured swap on node (in MB)
<b>FEATURE</b>	One or more colon delimited <STRING>'s (i.e., WIDE:HSM)  Punctuation and escapes are not allowed (.,;\t\n, etc.).	---	Generic attributes, often describing hardware or software features, associated with the node
<b>GEVENT</b>	GEVENT [ <EVENTNAME>]=<STRING>	---	<b>Generic event</b> occurrence and context data
<b>GMETRIC</b>	GMETRIC [ <METRICNAME> ]=<DOUBLE>	---	Current value of <b>generic metric</b> (i.e., 'GMETRIC [temp]=103.5').
<b>IDLETIME</b>	<INTEGER>	---	Number of seconds since last detected keyboard or mouse activity (often used with desktop harvesting)
<b>MAXTASK</b>	<INTEGER>	<CPROC>	Maximum number of tasks allowed on the node at any given time

Name	Format	Default	Description
<b>NETADDR</b>	<STRING>	---	The IP address of the machine
<b>NODEINDEX</b>	<INTEGER>	---	The <a href="#">node's index</a>
<b>OS</b>	<STRING>	---	Operating system running on node
<b>OSLIST</b>	One or more comma-delimited <STRING>'s with quotes if the string has spaces (i.e., "SAS7 AS3 Core Baseline Build v0.1.0", "RedHat AS3-U5Development Build v0.2").	---	Operating systems accepted by node
<b>OTHER</b>	<ATTR>=<VALUE> [ , <ATTR>=<VALUE> ] . . .	---	Opaque node attributes assigned to node
<b>PARTITION</b>	<STRING>	DEFAULT	Partition to which node belongs
<b>POWER</b>	<BOOLEAN>		Whether the machine is on or off
<b>PRIORITY</b>	<INTEGER>	---	<a href="#">Node allocation priority</a>
<b>RACK</b>	<INTEGER>	0	Rack location of the node
<b>SLOT</b>	<INTEGER>	0	Slot location of the node
<b>STATE*</b>	One of the following: Idle, Running, Busy, Unknown, Drained, Draining, or Down	Down	State of the node
<b>UPDATETIME*</b>	<EPOCHTIME>	0	Time node information was last updated

Name	Format	Default	Description
<b>VARATTR</b>	<code>&lt;ATTR1&gt;=&lt;VAL1&gt;</code> <code>[=&lt;displayName1&gt;]</code> <code>[+&lt;ATTR2&gt;=&lt;VAL2&gt;</code> <code>[=&lt;displayName2&gt;]]...</code>	---	<p>Plus-delimited (+) list of <code>&lt;ATTR&gt;=&lt;VAL&gt;</code> <code>[=&lt;displayName&gt;]</code> pairs that jobs can request. You can replace any of the equals signs with colons if desired.</p> <p>Specifying a display name allows you to choose a name that will be displayed in the Mongo database instead of the unique ID (the <code>&lt;VALUE&gt;</code>).</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>i</b> If you give two different attributes the same value and one of them also has a display name specified, both attributes will appear with the same display name.</p> </div>
<b>VARIABLE</b>	<code>&lt;ATTR&gt;=&lt;VAL&gt;</code>	---	Generic variables to be associated with node
<b>VMOSLIST</b>	<code>&lt;STRING&gt;</code>	---	Comma-delimited list (,) of supported virtual machine operating systems for this node
<b>XRES</b>	<p>One or more comma-delimited <code>&lt;NAME&gt;:&lt;VALUE&gt;</code> pairs (i.e.,</p> <pre>MATLAB:6,COMPILER:100 )</pre>	---	Amount of external usage of a particular generic resource

\* indicates required field

Node states have the following definitions:

State	Description
<b>Busy</b>	Node is running some jobs and will not accept additional jobs.
<b>Down</b>	Resource Manager problems have been detected. Node is incapable of running jobs.
<b>Draining</b>	Node is responding but will not accept new jobs.
<b>Idle</b>	Node is ready to run jobs but currently is not running any.
<b>Running</b>	Node is running some jobs and will accept additional jobs.
<b>Unknown</b>	Node is capable of running jobs but the scheduler will need to determine if the node state is actually Idle, Running, or Busy.

## N.2.2 Query Workload Data Format

### *Query Workload Data Format*

Name	Format	Default	Description
<b>ACCOUNT</b>	<STRING>	---	AccountID associated with job
<b>ARGS</b>	<STRING>	---	Job command-line arguments
<b>COMMENT</b>	<STRING>	0	Job <a href="#">resource manager extension</a> arguments including qos, dependencies, reservation constraints, etc.
<b>COMPLETETIME*</b>	<EPOCHTIME>	0	Time job completed execution
<b>DDISK</b>	<INTEGER>	0	Quantity of local disk space (in MB) that must be dedicated to each task of the job
<b>DGRES</b>	name:value [, name:value]	---	Dedicated generic resources per task
<b>DPROCS</b>	<INTEGER>	1	Number of processors dedicated to the job

Name	Format	Default	Description
<b>DSWAP</b>	<INTEGER>	0	Quantity of virtual memory (swap, in MB) that must be dedicated to each task of the job
<b>ENDDATE</b>	<EPOCHTIME>	[ANY]	Time by which job must complete
<b>ENV</b>	<STRING>	---	Job environment variables
<b>ERROR</b>	<STRING>	---	File to contain STDERR
<b>EVENT</b>	<EVENT>	---	Event or exception experienced by job
<b>EXEC</b>	<STRING>	---	Job executable command
<b>EXITCODE</b>	<INTEGER>	---	Job exit code
<b>FLAGS</b>	<STRING>	---	Job flags
<b>GEOMETRY</b>	<STRING>	---	String describing task geometry required by job
<b>GNAME*</b>	<STRING>	---	GroupID under which job will run
<b>HOSTLIST</b>	Comma or colon delimited list of hostnames - suffix the hostlist with a carat (^) to mean superset; suffix with an asterisk (*) to mean subset; otherwise, the hostlist is interpreted as an exact set	[ANY]	List of required hosts on which job must run (see <a href="#">TASKLIST</a> )  A subset means the specified hostlist is used first to select hosts for the job. If the job requires more hosts than are in the hostlist, they will be obtained from elsewhere if possible. If the job does not require all of the jobs in the hostlist, it will use only the ones it needs.  A superset means the hostlist is the <i>only</i> source of hosts that should be considered for running the job. If the job can't find the necessary resources in the hosts in this list it should <i>not</i> run. No other hosts should be considered in allocating the job.
<b>INPUT</b>	<STRING>	---	File containing STDIN

Name	Format	Default	Description
<b>IWD</b>	<STRING>	---	Job's initial working directory
<b>NAME</b>	<STRING>	---	User specified name of job
<b>NODES</b>	<INTEGER>	1	Number of nodes required by job (see <a href="#">Node Definition</a> for more info)
<b>OUTPUT</b>	<STRING>	---	File to contain STDOUT
<b>PARTITIONMASK</b>	One or more colon delimited <STRING> <sub>s</sub>	[ANY]	List of partitions where job can run
<b>PREF</b>	Colon delimited list of <STRING> <sub>s</sub>	---	List of preferred node features or variables. See <a href="#">PREF</a> for more information.
<b>PRIORITY</b>	<INTEGER>	---	System priority (absolute or relative - use '+' and '-' to specify relative)
<b>QOS</b>	<INTEGER>	0	Quality of service requested
<b>QUEUETIME*</b>	<EPOCHTIME>	0	Time job was submitted to resource manager
<b>RARCH</b>	<STRING>	---	Architecture required by job
<b>RCLASS</b>	List of bracket enclosed <STRING> : <INTEGER> pairs	---	List of <CLASSNAME> : <COUNT> pairs indicating type and number of class instances required per task. (i.e., [batch:1] or [batch:2] [tape:1])
<b>RDISK</b>	<INTEGER>	0	Local disk space (in MB) required to be configured on nodes allocated to the job
<b>RDISKCMP</b>	One of >=, >, ==, <, or <=	>=	Local disk comparison (i.e., node must have > 2048 MB local disk)
<b>REJCODE</b>	<INTEGER>	0	Reason job was rejected



Name	Format	Default	Description
<b>REJCOUNT</b>	<INTEGER>	0	Number of times job was rejected
<b>REJMESSAGE</b>	<STRING>	---	Text description of reason job was rejected
<b>REQRSV</b>	<STRING>	---	Name of reservation where job must run
<b>RESACCESS</b>	<STRING>	---	List of reservations where job can run
<b>RFEATURES</b>	Colon delimited list <STRING>'s	---	List of features required on nodes
<b>RMEM</b>	<INTEGER>	0	Real memory (RAM, in MB) required to be configured on nodes allocated to the job
<b>RMEMCMP</b>	One of '>=', '>', '==', '<', or '<='	>=	Real memory comparison (i.e., node must have >= 512MB RAM)
<b>ROPSYS</b>	<STRING>	---	Operating system required by job
<b>RSOFTWARE</b>	<RESTYPE> [{:}<COUNT>] [@<TIMEFRAME>]	---	Software required by job
<b>RSWAP</b>	<INTEGER>	0	Virtual memory (swap, in MB) required to be configured on nodes allocated to the job
<b>RSWAPCMP</b>	One of '>=', '>', '==', '<', or '<='	>=	Virtual memory comparison (i.e., node must have ==4096 MB virtual memory)
<b>SID</b>	<STRING>	---	System ID (global job system owner)
<b>STARTDATE</b>	<EPOCHTIME>	0	Earliest time job should be allowed to start
<b>STARTTIME*</b>	<EPOCHTIME>	0	Time job was started by the resource manager

Name	Format	Default	Description
<b>STATE*</b>	One of the following: Idle, Running, Hold, Suspended, Completed, or Removed	Idle	State of job
<b>SUSPENDTIME</b>	<INTEGER>	0	Number of seconds job has been suspended
<b>TASKLIST</b>	One or more comma-delimited <STRING>'s	---	List of <b>allocated</b> tasks, or in other words, comma-delimited list of node ID's associated with each active task of job (i.e., cl01, cl02, cl01, cl02, cl03) The tasklist is initially selected by the scheduler at the time the StartJob command is issued. The resource manager is then responsible for starting the job on these nodes and maintaining this task distribution information throughout the life of the job (see <a href="#">HOSTLIST</a> ).
<b>TASKS*</b>	<INTEGER>	1	Number of tasks required by job (see <a href="#">Task Definition</a> for more info)
<b>TASKPERNODE</b>	<INTEGER>	0	Exact number of tasks required per node
<b>UNAME*</b>	<STRING>	---	UserID under which job will run
<b>UPDATETIME*</b>	<EPOCHTIME>	0	Time job was last updated
<b>WCLIMIT*</b>	[[HH:]MM:]SS	864000	Walltime required by job

\* indicates required field

Job states have the following definitions:

State	Definition
<b>Completed</b>	Job has completed

State	Definition
<b>Hold</b>	Job is in the queue but is not allowed to run
<b>Idle</b>	Job is ready to run
<b>Removed</b>	Job has been canceled or otherwise terminated externally
<b>Running</b>	Job is currently executing
<b>Suspended</b>	job has started but execution has temporarily been suspended

**i** Completed and canceled jobs should be maintained by the resource manager for a brief time, perhaps 1 to 5 minutes, before being purged. This provides the scheduler time to obtain all final job state information for scheduler statistics.

---

### Related Topics

- [Managing Resources Directly with the Native Interface](#)

## Appendix O: SCHEDCFG Flags

Flag	Description
<b>AGGREGATENODEFEATURES</b>	<p>AGGREGATENODEFEATURES causes Moab to aggregate features reported by the different RMs. For example, if you have two RMs reporting different features for the same node, Moab will add both features together (instead of one being overwritten by the other).</p> <p>In order to set features manually, you can use <code>mnodectl -m features</code> (for details, see <a href="#">mnodectl</a>).</p>
<b>ALLOWCREDENTIALSWITHSPACES</b>	<p>ALLOWCREDENTIALSWITHSPACES lets Moab ignore POSIX standards and allows groups, users, and accounts with spaces in their names.</p>
<b>ALLOWINFINITEJOBS</b>	<p>ALLOWINFINITEJOBS allows infinite wallclock times to be accepted. Previously, jobs with infinite job times were allowed by default.</p>
<b>ALLOWMULTICOMPUTE</b>	<p>ALLOWMULTICOMPUTE tells Moab how to resolve conflicting information from different resource managers. If ALLOWMULTICOMPUTE is specified, Moab will use the STATE and OS information from the resource manager that reports the node as online.</p>
<b>ALLOWPERJOBNODESETISOPTIONAL</b>	<p>ALLOWPERJOBNODESETISOPTIONAL specifies whether Moab will read the <a href="#">NODESETISOPTIONAL</a> resource manager extension on an individual job or use the global setting.</p>
<b>CANCELFAILEDDEPENDENCYJOBS</b>	<p>Automatically cancels dependency jobs that will never run because of an unmet requirement. For example, if you ran a job with both an afterok and afternotok job attached to it and that job was successful, the afterok job would run, leaving the afternotok job idle in the queue. If you set CANCELFAILEDDEPENDENCYJOBS, Moab will cancel the job with the failed dependency and remove it from the queue. For more information about job dependencies, see <a href="#">Job Dependencies</a>.</p>


Flag	Description
	<div style="border: 1px solid #ccc; padding: 5px; background-color: #e6f2ff;">  If you want to cancel all jobs that a specified <code>&lt;job_id&gt;</code> depends on, use <code>mjobctl -c flags=follow-dependency &lt;job_id&gt;</code> instead.         </div>
<b>CHECKCIRCULARDEPENDENCIES</b>	<p>If you regularly submit job dependencies based on job names (and not job IDs) it is possible to accidentally create a circular dependency where jobs end up blocked. Moab can recursively check for circular dependencies when jobs are submitted by enabling this flag. Note that this check can be extensive depending on the workload submitted.</p>
<b>DISABLEPERJOBNODESETS</b>	<p>Disables a job's ability to override the system specified node set. See <a href="#">13.3 Resource Manager Extensions</a> for more information.</p>
<b>DISABLEPARTIALNODERESERVATIONS</b>	<p>Blocks partial node reservations.</p>
<b>ENABLEDYNAMICNODES</b>	<p>Enables the ability to automatically remove nodes from Moab that are no longer reported by the resource manager.</p>
<b>ENABLEJOBTRIGGERSONRSV</b>	<p>Enable job start triggers based on the job's reservation and not the actual start of the job (allows for negative offset job start triggers).</p>
<b>ENABLEMOABJOBENV</b>	<p>Puts the Moab job variables on every job.</p>
<b>ENFORCERESERVEDNODES</b>	<p>Without this flag Moab tries to optimize the reservation for a job before it starts, meaning a job can start on nodes that weren't part of its reservation. With this flag Moab tries to start jobs only on the nodes that were reserved.</p>
<b>ENFORCESAMENODESET</b>	<p>The same nodeset is not enforced across job requirements by default, rather each requirement is scheduled separately and the nodesets are determined on a per-req basis. To have Moab enforce the same nodeset across all job requirements set this flag.</p>

Flag	Description
<b>EXTENDEDGROUPSUPPORT</b>	Allows Moab to consider a user's secondary Linux groups when dealing with reservation <a href="#">ACLs</a> .
<b>FASTGROUPLOOKUP</b>	Moab will use the system call <code>getgrouplist</code> to gather group information. This can significantly improve performance on some LDAP systems.
<b>FASTRVSTARTUP</b>	<p>Speeds up start time if there are existing reservations.</p> <div data-bbox="797 653 1414 810" style="border: 1px solid #0070C0; padding: 5px; margin-bottom: 5px;"> <p><b>i</b> When you set the <b>FASTRVSTARTUP</b> flag, Moab will also set the <b>DISABLEPARTIALNODERESERVATIONS</b> flag.</p> </div> <div data-bbox="797 825 1414 915" style="border: 1px solid #0070C0; padding: 5px; margin-bottom: 5px;"> <p><b>i</b> <b>FASTRVSTARTUP</b> is incompatible with partial node reservations.</p> </div> <div data-bbox="797 930 1414 1020" style="border: 1px solid #0070C0; padding: 5px;"> <p><b>i</b> <b>FASTRVSTARTUP</b> maintains the resource manager-reported node order at all times.</p> </div> <p>On very large systems, if there is a reservation in the checkpoint file on all the nodes, it would take a really long time for Moab to start up. For every node in the reservation, Moab checks every other node. With this flag, Moab just uses the nodelist that was checkpointed to create the reservation. It speeds up the startup process because it doesn't have to check every node. Where Moab would take 8 - 10 minutes to start up with an 18,000 node reservation without the flag, Moab can start up in 2-3 minutes with the flag.</p> <p>With the flag you will see one difference in <a href="#">checknode</a>. A reservation that uses all the procs on a node initially shows that all the procs are blocked. Without the flag, and as jobs fill on the node, the blocked resources will be configured - dedicated (ex. 5/6). With the flag, the blocked resources will always be what the reservation is blocking and won't change when jobs fill on the node.</p> <p>Without flag: Reservations: brian.1x1 User -00:12:52 -&gt; INFINITY ( INFINITY)</p>

Flag	Description
	<p>Blocked Resources@-00:00:02 Procs: 5/6 (83.33%) Mem: 0/5000 (0.00%)                      Blocked Resources@00:04:58 Procs: 6/6 (100.00%) Mem: 0/5000 (0.00%)                      m.2x1 Job:Running -00:00:02 -&gt; 00:04:58 (00:05:00)                      Jobs: m.2</p> <p>With flag:                      Reservations:                      brian.1x1 User -00:00:15 -&gt; INFINITY (INFINITY)                      Blocked Resources@-00:00:02 Procs: 6/6 (100.00%) Mem: 0/5000 (0.00%)                      Blocked Resources@00:04:58 Procs: 6/6 (100.00%) Mem: 0/5000 (0.00%)                      m.1x1 Job:Running -00:00:02 -&gt; 00:04:58 (00:05:00)                      Jobs: m.1</p>
<b>FILELOCKHA</b>	<p>This is a <a href="#">High Availability</a> feature. FILELOCKHA prevents scheduling conflicts between multiple Moab servers.</p>
<b>FREECOMPLETEDJOBSSUBMITSTRING</b>	<p>Moab frees the job submit string for completed jobs, decreasing the amount of memory needed during operation. This is useful in environments with large job scripts that can create a large memory footprint.</p>
<b>IGNOREPIDFILELOCK</b>	<p>Moab frees the job submit string and environment for completed jobs, decreasing the amount of memory needed during operation. This is useful in environments with large job scripts and environments that can create a large memory footprint.</p>
<b>INTERACTIVWCACCURACY</b>	<p>Moab will assume all interactive jobs are 100% accurate with respect to walltime. This is useful for sites that don't enforce walltimes for interactive jobs but don't want users punished for inaccurate interactive job walltimes.</p>
<b>JOBSUSERSVWALLTIME</b>	<p>Allows jobs submitted without a walltime request</p>

Flag	Description
	<p>or default walltime received from a class or queue but with an <code>ADVRES:reservation</code> to inherit their walltime limit from the reservation instead of the Moab default. The job walltime limit is then the remaining time of the reservation to which the job was submitted.</p>
<b>NOCLASSUPDATE</b>	<p>While running against Torque, Moab will not update classes when it refreshes each iteration. Moab loads the classes at startup, but does not refresh them until the next time it is restarted.</p>
<b>NORMALIZETASKDEFINITIONS</b>	<p>Instructs Moab to normalize all tasks that it receives via an <code>mshow -a</code> command. Moab normalizes the task definition to one processor and then changes the tasks requested to the number of processors requested. For example, when the following is received by Moab:</p> <pre data-bbox="797 940 1414 982" style="border: 1px dashed gray; padding: 5px;">mshow -a -w mintasks=1@procs:4+mem:4096</pre> <p>It is changed to this:</p> <pre data-bbox="797 1052 1414 1094" style="border: 1px dashed gray; padding: 5px;">mshow -a -w mintasks=4@procs:1+,mem:1024,tpn=4</pre>
<b>OPTIMIZEDBACKFILL</b>	<p>Turns on an optimization within the <code>FIRSTFIT</code> backfill algorithm that checks whether there is a system-wide reservation blocking most jobs from running. This flag speeds up backfill scheduling (in the case where there is a system-wide reservation blocking most users) by checking access to the reservation sooner rather than later. This flag will be the default in future versions of Moab.</p>
<b>PRIORITYPOLICYBLOCKING</b>	<p>By default, a job that violates a policy is placed into the blocked queue. Jobs with a lower priority, but that do not violate the policy, will run. This can lead to situations where small jobs starve out larger, higher priority jobs.</p> <p>When you set the <code>PRIORITYPOLICYBLOCKING</code> flag, Moab allows the job that violates the policy to continue consuming the policy slots while it remains blocked. With the policy slots consumed, the smaller, lower priority jobs will not run. The</p>



Flag	Description
	<p>higher priority job will continue to consume the policy slots until it has consumed enough to actually run.</p> <p>Note that because the blocked job consumes policy slots, this will inevitably lead to lower system utilization.</p>
<b>SETDEFAULTHOSTLISTEXACTSET</b>	<p>By default, all Hostlist requests will be interpreted as an exactset request. See <a href="#">HOSTLIST</a> for more information.</p>
<b>SHOWCOMPLETEDDEPENDENCIES</b>	<p>Continues showing dependencies on a job even after the dependencies have been satisfied.</p>
<b>SHOWREQUESTEDPROCS</b>	<p>Shows requested processors regardless of NodeAccessPolicy in <a href="#">showq</a>. When <code>SINGLEJOB NODEACCESSPOLICY</code> is used and the job requests one processor, <i>showq</i> displays the job with one processor.</p>
<b>SHOWUSERJOBONLY</b>	<p>Causes Moab, when a non-admin user runs <i>showq</i>, to return only that user's jobs. If an administrator runs <i>showq</i> when this flag is set, Moab returns the jobs of all users; no restrictions are placed on administrators.</p>
<b>STRICTSPOOLDIRPERMISSIONS</b>	<p>Enforces at least a 511 permission on the Moab spool directory.</p>
<b>SUSPENDEDJOBNODEBFINELIGIBLE</b>	<p>Turns off back fill on nodes where a job was suspended due to preemption.</p> <div data-bbox="797 1423 1414 1583" style="border: 1px solid red; padding: 5px; margin-top: 10px;">  When enabled, <code>CHECKSUSPENDEDJOBPRIORITY</code> must be set to True. See <a href="#">CHECKSUSPENDEDJOBPRIORITY</a>.                 </div>
<b>UNMIGRATEONDEFER</b>	<p>Forces Moab to unmigrate a job in a grid if it enters a deferred state.</p>
<b>USERMCOMPLETEDJOBSTATS</b>	<p>Calculate job statistics using the RM reported metrics rather than Moab's internal metrics.</p>

