

Charting Your Own Cloud Success Story

Research and Insights on Innovative Cloud Management

- 1
Introduction
- 2
From the Gartner Files:
Cool Vendors in Cloud
Management, 2013
- 8
Optimizing Your Private Cloud
Value with Policies
- 11
About Us

Introduction



Private cloud computing is exciting to business leaders for all the obvious reasons: it enables faster innovation and growth as well as a more cost efficient organization. In short, it is about turning IT

from a cost center into a service provider that turns opportunities into realities. Cloud computing helps organizations achieve revenue faster and at a lower cost for each service. Private cloud delivers the business ROI with reduced risk. Whether you are just beginning to consider a private cloud implementation, currently taking advantage of the benefits it has to offer, or looking to expand its value, you should always take the opportunity to plan your cloud strategy to ensure it is optimizing the value it can deliver your organization.

This newsletter features the latest research from Gartner on a select group of innovative cloud management vendors and tools that can help you design and optimize your private cloud strategy. It also presents some key cloud management capabilities to consider in your cloud strategy, based on best practices we have learned from our work with real-world private cloud implementations. Specifically, policy-based optimization will help you achieve greater ROI and agility in your data center. You will learn how to maximize service performance and cost efficiency as well as deployment speed in your private cloud. Using these best practices and Gartner's forward-looking research will help you create a truly optimized private cloud for your organization.

I hope you will find some keys to charting your own cloud success story.

Best Regards,

Rob Clyde
Chief Executive Officer
Adaptive Computing

Cool Vendors in Cloud Management, 2013

Enterprise IT leaders and cloud architects are facing new opportunities to leverage a growing set of offerings from emerging cloud management technology vendors. We examine five such vendors providing cloud management platform and/or cloud migration capabilities.

Key Findings

- Most emerging cloud management vendors will not exist in five years, either having been acquired by larger vendors or succumbing to business loss.
- As enterprises seek to implement public, private and hybrid cloud computing, they increasingly desire cloud management technologies to assist with management, governance and migration.
- Emerging cloud management vendors gain traction because of their innovation, which typically does not exist at large cloud management vendors.

Recommendations

- CIOs, IT infrastructure and operations (I&O) leaders and cloud architects evaluating cloud management platform (CMP) technologies should assess large and emerging vendors. While it is common to assess larger vendors with which the enterprise has a partnership, the smaller vendors must compete on innovation.

- Enterprises and service providers seeking faster, less expensive migration between cloud providers should assess cloud migration technologies. Beware of the complexities of migrating legacy systems to the cloud without reassessment of application architectures (to gain cloud attributes).
- Enterprises and service providers can derive benefits from the technologies of small vendors; however, extreme caution should be taken when the expected ROI is not justified within a two-year horizon, as we expect further market consolidation to occur.
- Because implementing private clouds takes time and is often transformational, plan your cloud computing strategy so that benefits and gains are made every six to nine months, not just at the end of the journey.

Strategic Planning Assumption

By 2016, 60% of Global 2000 enterprises will have invested in CMP technology, up from less than 30% in 2013.

Analysis

This research does not constitute an exhaustive list of vendors in any given technology area, but rather is designed to highlight interesting, new and innovative vendors, products and services. Gartner disclaims all warranties, express or implied, with respect to this research, including any

warranties of merchantability or fitness for a particular purpose.

What You Need to Know

Cloud management vendors and technologies have both proliferated and consolidated over the past year, as a result of serious interest and investment in hybrid cloud computing by enterprise IT organizations, as well as the uptake of the technology by service providers seeking to build new cloud services businesses. The large number of small cloud management providers will increasingly find it difficult to differentiate themselves in a very crowded market, especially with every large infrastructure and management vendor competing against them. However, like most small startups, they invest in best-of-breed capabilities that attract visionary enterprises and service providers to invest in them. We recommend exploiting these innovative capabilities when shorter investment horizons justify the risk, as well as pushing your existing providers for comparable capabilities.

Cool Vendor success often results in acquisition by larger vendors (as was the case for Cloupia, DynamicOps, Gale Technologies and ManageIQ in 2012).

This research focuses on five innovative cloud management vendors, three of which are CMPs (Adaptive Computing, Enstratus [formerly enStratus] and Flexiant), and two of which are cloud migration vendors (HotLink and RiverMeadow Software).

Adaptive Computing

Provo, Utah (www.adaptivecomputing.com)

Analysis by Donna Scott

Why Cool: Adaptive Computing offers a CMP called Moab Cloud Suite, which enables real-time infrastructure through its workload and resource pool optimization capabilities across a heterogeneous physical (Windows, Linux and Unix) and virtual (VMware, Xen and KVM) server infrastructure. This is cool because most CMPs perform provisioning and orchestration, and very few perform dynamic optimization of the environment. The Moab workload management functionality can also be acquired independently and layered on top of other CMPs. Moab offers a policy model that enables workload (and multiworkload service) placement and ongoing optimization decisions to be made automatically, based on a multitude of customizable factors, such as workload proximity (e.g., proximity to key resource attributes to enhance service performance), capacity and performance, quality of service, constraints (in user groups or resource groups), affinities, and budgets (e.g., overcommit allocation policies and those to reduce license management costs). Policies can also be managed and customized with the provided Web services tools, as long as instrumentation is available (e.g., to avoid placing virtual machines [VMs] onto hot servers by assessing the server temperature).

For service provisioning, the Moab mathematical model works in real time to compute a score that matches service request policies to those of available resource policies. The algorithm for initial placement would first identify available capacity across the resource pools, then would eliminate those servers under maintenance and apply the additional rules to develop a prioritized list of preferred placement (the higher the number, the better the fit). All this is done without any coding – just assignment of policies and preferences

at the service definition and resource pool levels. Ongoing service optimization rebalances the loads across resources and data centers if desired, assesses resource performance and thresholds, and determines which VMs should be moved to other servers or resource groups to improve performance, but without violating placement policies. Customers can limit the frequency of automated movement, specify the maximum number of movements or throttle the movements to reduce the performance impacts of the moves.

Adaptive Computing was founded in 2001 as Cluster Resources, focusing on the high-performance computing (HPC) market. In 2010, it entered the cloud computing space, first with Moab and then with its broader CMP, Moab Cloud Suite. Adaptive Computing is primarily financed by its two founders, but took a \$14 million round of funding in 2010 to aid in its growth. The vendor has approximately 120 employees and we estimate \$20 million in annual revenue, with approximately a 50/50 split between the grid product line and the CMP line. It has a go-to-market partnership with HP Software for policy-based optimization. Adaptive Computing prices its products on either a subscription or perpetual basis, and last year it moved from server-capacity-based pricing to VM-based pricing. The suite costs \$500 per VM, with discounting available for volume purchasing. The vendor has added some larger private cloud computing clients in financial services and government, primarily due to its workload management and optimization capabilities, as well as its ability to integrate to other third-party tools.

Challenges: The CMP market is extremely competitive and composed of offerings from large infrastructure providers, such as Citrix, Dell, Microsoft and VMware, and from established IT operations management firms, such as BMC Software, CA Technologies, HP and IBM. Adaptive Computing's challenge is gaining market visibility with its relatively small sales force and with just one industry

partner (HP), which also competes with it. In addition, many startups in the market provide similar CMP capabilities, although most do not have workload optimization functionality. Moreover, while Adaptive Computing has been making its product more user-friendly over the last several releases, understanding the mathematical modeling in order to use it correctly can be challenging and often requires integration to third-party solutions to get at more instrumentation to use in Moab's analysis. In addition, customers report that, while getting better, documentation remains a weakness as does real-time integration with VMware for up-to-date workload optimization.

Who Should Care: CIOs, vice presidents (VPs) and directors of IT, as well as enterprise and infrastructure architects looking to deliver cloud-based, on-demand services that require infrastructure optimization (workload balancing). Service providers may also be interested in this solution, due to its ability to optimize the infrastructure, thus dropping service delivery costs.

Enstratius (formerly enStratus)

Minneapolis, Minnesota
(www.enstratius.com)

Analysis by Milind Govekar

Why Cool: Enstratius has innovated in the CMP area for hybrid and heterogeneous cloud implementation. Its technology has two main components – the Enstratius Cloud Management System and the Enstratius Guest Agent. The Cloud Management System provides three key features – governance, automation and cross-cloud management built on a scalable architecture. It can be implemented as a software as a service (SaaS) deployment or as on-premises software.

Governance includes capabilities such as key management and encryption (including file system encryption using AES-256 encryption) that is kept outside

the cloud provider's infrastructure and application. It also provides role-based and user-defined security that integrates with existing LDAP and Active Directory systems to use current identities and roles. It also provides customized alerting capabilities related to the violation of policies across resources (e.g., server utilization, firewall changes, etc.), including the violation of budgeted costs using various billing codes. Monthly reports can also be generated that aggregate costs by billing codes across various cloud infrastructures. Furthermore, the Cloud Management System can distinctively provide intrusion detection, alerting and authentication using various methods, and it supports single sign-on for multicloud environments.

Automation capabilities include the ability to model the architecture; define policies for applications across multiple cloud zones, regions or services; and then automate the applications' deployment. The Cloud Management System is able to distribute the workload across more than one cloud service using load balancers. Configuration scripts can be defined and orchestrated using the tool, or existing scripts written and maintained in both the Puppet and Chef enterprise product offerings. Scaling can be triggered based on server metrics (CPU usage, memory usage, etc.) or custom metrics (transaction rates, business metrics, etc.) to scale within or across zones, regions or multiple cloud environments. Furthermore the tool is able to set up automated backup and recovery across multiple cloud environments. Although there are CMPs that have some of these capabilities, there are few that have all of them.

Cross-cloud management capabilities in the Cloud Management System particularly enable disaster recovery (DR) capabilities (for example, by using Rackspace for the production environment and providing DR across the Amazon Web Services [AWS] or Microsoft Azure cloud environment). A key feature of the Enstratius Cloud

Management System is that it integrates with multiple cloud providers' APIs. Enstratius supports a wide variety of cloud services providers, including AT&T (Synaptic Storage as a Service), Amazon (AWS), Bluelock, Cloudscaling (Open Cloud System [OCS]), Citrix (CloudStack), CloudSigma, Dell, EMC (Atmos), Eucalyptus Systems, GoGrid, Google (Google Cloud Storage), HP (HP Cloud Services), IBM (SmartCloud), Joyent (Joyent Cloud), OpenStack, Rackspace, Tata Communications (InstaCompute), Terremark, VMware (vSphere and vCloud Express) and Microsoft (Azure).

Enstratius customers like its API support, which enables rapid integration with various environments, and its ease of implementation and installation. The Enstratius Guest Agent can be installed on VMs in the cloud for improved security and automation by establishing a trust relationship between the Cloud Management System and the Guest OS.

The vendor is a privately owned, self-funded company founded in 2007 to deploy and manage enterprise applications in the private, public and hybrid cloud environments. Enstratius employs approximately 50 people and, in addition to a range of small and midsize businesses, it serves 25 large enterprises and service providers.

Challenges: Some of Enstratius' new capabilities are not fully tested in a live, highly scalable production environment (e.g., the virtual LAN [VLAN] isolation capabilities to support virtual private cloud [VPC] environments). Furthermore, some customers have found that configuring Enstratius against cloud platforms where behavior can differ significantly depending on the platform's configuration (such as with a customized VMware vCloud Director installation) can be a challenge. As the number of vendors in the cloud management area grows, enterprises are finding it difficult to differentiate between

their capabilities. We continue to see startup vendors having to work particularly hard at gaining enterprise customers and implementations. They will have to compete, and also cooperate, with some of the larger vendors from the IT operations management and virtual infrastructure areas. They will also compete with other smaller players, such as ServiceMesh, RightScale and others.

Who Should Care: CTOs, cloud architects, VPs of infrastructure, and those in enterprise architecture and DevOps roles. Service providers will also be interested in this solution, due to its secure multitenant support.

Flexiant

London, England (www.flexiant.com)

Analysis by Milind Govekar and Gregor Petri

Why Cool: Flexiant has innovated in the cloud orchestration area, with a solution aiming to assist cloud and hosting service providers in provisioning both simple and more complex cloud services.

Flexiant Cloud Orchestrator is mainly targeted at hosting and service providers; therefore, it is a highly scalable tool able to do policy-based workload placement on various resources in private, hybrid and public cloud environments. It is able to set up servers and users, and to manage rules associated with infrastructure and application resources across a multihypervisor environment. Furthermore, it is able to provide multilevel billing and pricing capabilities for cloud resources used, adding a business management layer to the technical orchestration capabilities. Flexiant Cloud Orchestrator can also set up multilevel VPC environments using templates for firewalls, servers and other components using a designer, and can set up a role-based access environment. This multilevel capability allows for various scenarios involving value-added resellers (VARs) and resellers working with the service provider,

and/or accommodating multilevel customer organizations.

The offering also has an API interface for integrating with, for example, external billing tools and CRM applications. There is an innovative feature with a slightly peculiar (sushi-inspired) name, called Bento Boxes. This functionality (released in November 2012) allows modeling and provisioning of applications, enabling providers to automatically offer and provision more complex services (at a level above individual VMs). While Flexiant's technology is primarily aimed at the service providers, large organizations with a multitenant requirement or with a need for community cloud-type requirements should also consider Flexiant's cloud technology.

Flexiant Cloud Orchestrator is priced per cluster and hypervisor combination annually or per core per month for four different scenarios – Hosting Edition, Hosting Advanced Edition, Service Provider Edition and Service Provider Advanced Edition. The pricing starts from \$150 per cluster and hypervisor combination per year for Hosting Edition and \$1,500 per cluster and hypervisor combination for Service Provider Edition. It is \$8 per core per month for Hosting Edition and \$16 per core per month for Service Provider Edition.

Flexiant originated in 2010 (it started off as a hosting company in 1997, and later in 2009, the hosting company was sold off). It is an investor-funded company, mainly funded by Leopard Rock Capital Partners. Flexiant employs more than 40 employees based in the U.K., Amsterdam and New York. It has approximately 100 paying customers, mainly in EMEA and some in North America.

Challenges: Flexiant's innovative solution will face competition from CMPs that have built-in cloud services orchestration capabilities for a hybrid cloud environment. Furthermore, Flexiant's focus only on hosting and service provider environments excludes it from

the lucrative enterprise market segment. Like every small vendor, it will have to think differently to gain market visibility by developing a good marketing strategy and channel to reach prospects and customers.

Who Should Care: Hosting and service providers with cloud business line responsibility should consider Flexiant, in addition to organizations requiring multiitenancy and community clouds.

HotLink

Santa Clara, California (www.hotlink.com)

Analysis by Aneel Lakhani

Why Cool: HotLink offers a heterogeneous virtualization and cloud management product that plugs into VMware's vCenter management tool. This enables management of non-VMware environments directly from within vCenter and presents non-VMware hosts and VMs plus cloud accounts and instances as if they were vSphere hosts and VMs. HotLink's core platform works via an intermediary data model, transformation engine, and an API layer that enables it to manage, clone, snapshot, provision from template, convert and migrate VMs bidirectionally between any local VMware ESX/vSphere, Hyper-V, KVM and XenServer hypervisors, as well as to and from any CloudStack- or AWS-based public cloud. The core platform works by interacting with each hypervisor directly via an agent and providing this metadata to vCenter. Data and capabilities of each environment are presented in vCenter, and all the management functions of each target environment are made available. For example, if there is a certain capability for a target hypervisor that does not exist for VMware, a control and the respective user interface elements for that capability are created and presented for use. The core platform has been built with the ability to be integrated into other management platforms, and the vendor has just launched a new version of its Hybrid Express product

(see below) that has been integrated with Microsoft System Center Virtual Machine Manager (SCVMM).

On this platform, HotLink has built three products:

- HotLink SuperVISOR: Extends VMware vCenter management to on-premises multihypervisor environments, including Hyper-V, KVM and XenServer
- HotLink SuperVISOR Free Edition: Free version of the full-featured product that works with local non-VMware environments on a maximum of three hosts and for up to 15 powered-on VMs
- HotLink Hybrid Express: Extends VMware vCenter management to AWS Elastic Compute Cloud (EC2) and CloudStack public cloud environments; and Microsoft SCVMM to AWS EC2
- HotLink Hybrid Express Free Edition: Free version of the above product that works with AWS EC2 and supports up to 10 powered-on instances

HotLink was founded in 2010, focused on leveraging the penetration of VMware's hypervisor, as well as the vCenter management tool, and capitalizing on an anticipated demand from enterprise customers to manage non-VMware environments, without the additional overhead of using a different console for each target hypervisor or cloud. The vendor is venture-backed by Foundation Capital and Leapfrog Ventures with a \$10 million first round, and has approximately 25 employees. It has more than a dozen large paying customers and several hundred downloads of its new free edition. Customers are primarily enterprises, but also include educational institutions and the enterprise sides of IT vendors and service providers. The first product was released in September 2011. HotLink prices products based on a base price for the platform

plus a price for each additional hypervisor or cloud platform, plus a price for each managed host or instance, plus annual maintenance and support. Subscription-based licensing is also available. Near-future development plans include integration with additional virtualization and CMPs beyond vCenter.

Challenges: Today, HotLink is offering its product primarily as a means of extending vCenter, with only added support for SCVMM; therefore, it has put itself in both a complementary and competitive position with VMware and Microsoft as each develops or extends its own heterogeneous hypervisor (vCenter Multi-Hypervisor Manager and SCVMM's current VMware management capabilities) and cloud (vCloud Automation Center, formerly DynamicOps and System Center 2012) management capabilities. HotLink is positioning itself as complementary to CMPs by providing fine-grained administration without the need for multiple native consoles, but many of those already offer some form of heterogeneous hypervisor and cloud management, although still requiring multiple administrative consoles for cross-platform administration. This market is both extremely competitive and crowded with vendors both large (such as BMC, Cisco and Microsoft) and small (such as CloudBolt, Emotix and ServiceMesh). Insofar as HotLink is seen as being a VMware extension, its visibility and traction may be dampened by VMware's

development of competitive capabilities. Insofar as HotLink is seen as being in an ambiguous position relative to existing CMP vendors, its visibility and traction may be dampened by companies seeking to promote their own heterogeneous environment management capabilities.

Who Should Care: CIOs, VPs and directors of IT, as well as enterprise and infrastructure architects who need to minimize the operational overhead of running multihypervisor and multicloud environments

RiverMeadow Software

Westford, Massachusetts
(www.rivermeadow.com)

Analysis by Donna Scott

Why Cool: RiverMeadow Software offers a SaaS-based solution for automating the migration of virtual or physical server workloads (images) into and between public, private and hybrid cloud environments. RiverMeadow has targeted service providers as its main prospects, with the use case of migrating existing legacy applications from on-premises to service provider clouds. The value proposition in this use case is to drop the time, costs and risks of migration, especially in professional services labor costs. By using RiverMeadow's cloud migration SaaS

offering, service providers can also achieve higher throughputs and collapse the elapsed time it takes for migrations.

The source environment does not need to be modified or taken out of production, because the SaaS offering does not require agents or need to quiesce servers during migrations. The source and target environments supported are Xen (including Amazon), KVM, ESX and vCloud Director, as well as CloudStack and OpenStack. Source-side support is also available for physical Linux and Windows servers. RiverMeadow also plans to support Hyper-V by year-end 2013. The RiverMeadow SaaS offering copies the server image to its database, extracts platform-specific characteristics (such as drivers, IP addresses and registry entries) and then converts the image to a format suitable for the target environment (adding new drivers and settings as necessary) prior to distributing the image to the target environment. Premigration and postmigration processing is supported to automate operational processes (changing configuration settings or installing service-provider-specific software). Early customer references indicated that RiverMeadow has about a 90% success rate in migrations and has drastically cut down the people time involved in them.

RiverMeadow was founded in 2009, and has approximately 40 employees and a dozen customers. It has received two

rounds of funding (including a second round led by Violin Memory) for a total of \$12.2 million. RiverMeadow has a referral agreement with VMware and OEM reseller agreements with Cisco, Ericsson and Violin Memory. VMware refers the RiverMeadow SaaS offering to service providers as a preferred solution for onboarding to vCloud Director. RiverMeadow has two resellers of enCloud – Violin Memory and VMware. VMware resells RiverMeadow software to service providers for aid with onboarding to vCloud Director. It is also planning to use RiverMeadow for self-service trials for customers considering implementing a private cloud computing environment based on VMware technologies. RiverMeadow prices its software on a per-migration basis, at \$100 to \$300 per migration, depending on the volume of migrations purchased.

Challenges: The cloud migration market is an embryonic one, and while there is significant interest from service providers,

RiverMeadow's success will be dependent on how much is migrated from on-premises to service providers. In addition, one time legacy migrations may not have enough growth opportunity for the vendor, requiring it to expand to other lines of business to achieve growth targets. Hybrid cloud computing and moving workloads between data centers are appealing value propositions to most CIOs; however, they are not easily implementable in the near term. Moreover, there is a lot of capital flocking to many small businesses in this space, some with technologies similar to RiverMeadow's, while others have added focus on an IT service (not a server), as well as the data migration, both of which RiverMeadow does not do. RiverMeadow will have to continue to add value to its migration business to gain the attention of service providers' CTOs and product managers in what will be a crowded marketplace. Lastly, RiverMeadow will have to balance its reselling opportunity with VMware with its other service provider customer

requirements, to be sure it continues to focus on heterogeneous environments, lest it risk business opportunities with non-VMware implementations.

Who Should Care: Large customers that have found themselves with services built in one environment (for example, Amazon) that want to move them to another provider or internally should evaluate RiverMeadow's cloud migration SaaS offering. Service providers seeking self-service and nonintrusive migrations from customers' on-premises environments to their own environments will also be most interested in this vendor.

Source: Gartner Research, G00249819, Donna Scott, Milind Govekar, Aneel Lakhani, Gregor Petri, 10 April 2013

Optimizing Your Private Cloud Value with Policies

Your cloud computing strategy should have goals that include creating innovation, value, and differentiation opportunities for your organization. Incorporating lessons learned from other organizations' established private cloud implementations, and even virtualization and data center automation, helps ensure you address all of the key challenges and capability points to accomplish this. It can also help you identify needs and capabilities you may have not yet uncovered that can be essential to maximizing the value your cloud can deliver your organization.

Leading, innovative organizations have found that policy-based optimization is a powerful way to chart the course to cloud success. This article will share some of the key capabilities that these organizations found most valuable from policy-based optimization, the challenges they helped solve, and the value they helped them achieve for their organizations. These are the capabilities you can plan into your cloud strategy to create a private cloud with not only the speed and agility benefits that you need, but also the additional benefits of greater cost savings, service performance and availability that put you ahead.

Optimized Service Placement Policies Prevent Failures and Boost Performance

Your cloud strategy should incorporate a set of policies that optimize service placement in your cloud. Current cloud customers have found this a key value point for services being delivered in a stable and high performance state. These policies extend beyond the base requirements of the service definition and customize placement to the needs of each service such as data center placement for security, data or latency requirements. It can even consider that certain complex application environments require specialized resources, such as fast network access or I/O, and can prioritize placement on resources with those types of node attributes to boost performance like you would tune a race car before taking it out on the track.

Customers have found these types of policies valuable in optimizing utilization beyond just virtualization as they automate placement across multiple datacenters and nodes based on multiple prioritized current usage factor points the admin can set that make sure resources are utilized to higher levels. These policy capabilities are essential for organizations looking to automate delivery of more complex application service environments and production workloads in their private cloud. The policies combine to go beyond automated provisioning to also ensure that the workload will be able to run optimally and successfully. Organizations should look at adding a combination of data center, node and node feature placement optimization policies for their private cloud to meet this challenge and deliver this added value.

Capacity Management Policies Maximize Utilization and Available Capacity While Reducing Costs


Private cloud customers find that their agile self-service delivery quickly unleashes service demand in their organization. This demand eats through their provisioning cost savings and capacity at an accelerated rate. As a result, demand grows faster than hardware capacity budgets. There is value



in considering this for your organization. You should plan for how to automate capacity management with policies to better utilize the capacity you have to meet this quantum leap in demand and requests.

Capacity management policies help you reclaim the unused, wasted resources that are usually provisioned from the service catalog, but often exceed what is actually used by a service. Hypervisor overcommit policies coupled with SLA performance policies ensure that you make this unused capacity available to new services, maximizing utilization while still ensuring services have the full resources for performance when needed with automated migration.

This policy set should also include an automated VM migration consolidation policy to continuously pack workload and eliminate the VM sprawl and capacity fragmentation customers inevitably see in their cloud environments. This sprawl, often even worse than in just a regular virtualization environment, prevents them from using the capacity they have when and where they need it. This becomes important when you start doing larger, more complex



application environments and production workloads where you may need some of the components located close together. The fractured capacity caused by the VM sprawl may begin to limit your ability to deliver these services optimally within your private cloud.

The combination of the consolidation policies with the hypervisor overcommit policies can maximize utilization two to three times over just virtualization alone. Your capacity can better keep pace with demand and you can reduce initial and ongoing hardware costs by 50%. This can be a significant innovation and cost efficiency differentiator for your organization.

Maintenance Reservation Policies Automate Tasks and Reduce Management Costs by 50%

Customers have also learned that they must find a way to scale their existing staff to handle three to five times the number of services that cloud enables them to deliver when it comes to maintenance tasks. Since they are looking at cloud to reduce their costs, growing the operating costs is not an option, especially since those costs are already a large portion of their IT budget. Customers also want to free up staff and budget for innovation projects and services

that they can offer in the cloud. Policy-based optimization can help with maintenance reservation policies that cut the time and costs admins spend doing repetitive maintenance tasks like updates.

Maintenance reservation policies and administration dashboards should automate reserving resources for maintenance, individually or as rolling updates across the cloud. All of the complexities of migrating workloads to new resources should be automated for the admins, including respecting all optimized placement policies. The policies also ensure that there is no disruption for the users and no new errors or failures introduced. This enables the admins to quickly just do the updates using standard update and configuration automation tools and then remove the reservation to allow workloads to begin using the updated resources again.

Service Performance and Availability Policies Optimize Meeting SLAs

As most customers who have been running private clouds for some time will attest, delivering services quickly ends up being only half the battle. Making sure each service stays performing to SLAs is the other side of cloud success. Ensuring that services are usable to your organization and balancing availability guarantees in the shared resource environment are non-trivial issues. Policy-based optimization can help here by automating the continuous monitoring of resource SLA thresholds and reservations and then automatically adapting the resource infrastructure without admin intervention so that services stay up and running at peak performance and so service availability guarantees are met even in a shared cloud environment.

Automatic VM Migration policies should have flexible, multi-factor SLA threshold settings to ensure that admins can effectively set when hypervisors are reaching overburdened levels



that will start to impact performance. These policies should be out-of-the-box but also open, allowing admins to use APIs to bring in metrics from the management tools of their choice. This prevents locking them into a limited set of just two or three metrics from just their virtualization or cloud vendor that may not meet their SLA needs. These policies should also respect all optimized placement policies to ensure that migrated services are relocated to optimal resources and don't introduce any new points for potential service issues.

Resource and Service Start Reservation policies are two additional key capabilities to add as they ensure service availability will be consistently met for users. These two policies should work together to guarantee the right resources for services will be available at the right time for important activities to the business whether a product development, customer service or financial processing project. The policies automatically migrate off other workload from the resources at the reservation time, prevent other services from being placed on those resources, and can even start the service on the resources in the case of the Service Start Reservations policy. The Service Start Reservation policy can even augment meeting availability SLAs further by suggesting the earliest available start time to the user if capacity is not available for a specific service at a specific requested time.



Policies automate all this availability scheduling and service migration, respective of all other optimization and SLA policies, from a simple request. These policies should be considered for your cloud strategy if you are going to be deploying production workloads in your cloud or servicing multiple groups and lines of business. Both of these will require you to meet the challenge of performance and availability SLAs for services.

Putting Policy-based Optimization to Work for Your Cloud

These four policy capabilities demonstrate some of the customer-proven ways policy-based optimization can maximize the value, innovation and differentiation your cloud can deliver to your business. Putting them to work for your business requires taking the next steps to plan them into your cloud strategy and choosing the right cloud management tool to enable them. Often customers find success with implementing

the optimized service placement and maintenance reservation policies as part of their initial cloud implementations as these are key to the service delivery and management. They then augment their cloud capabilities by implementing the SLA policies and capacity management policies.

These policies become required with increased usage and demand for the cloud services along with expansion to a wider range of production workloads. So, it is critical that in creating or expanding your cloud strategy, you select a cloud management tool that can provide this full set of policy capabilities that are designed not only to give the individual policy value but also to work together. By working together, they can achieve higher cumulative value across cost savings, service level delivery and your choice of infrastructure and management tools.

The Moab Cloud solution family gives you this policy-based cloud optimization. Moab Cloud Optimizer brings this workload optimization to other cloud platforms, such as those from HP, while Moab Cloud Suite delivers policy-based optimization along with a core set of cloud management capabilities. The Moab cloud products also provide these policies out-of-the box with visual policy management dashboards. These dashboards make configuring and managing even the most sophisticated multi-factor policy simple and easy. Moab cloud solutions have been chosen by leading private cloud customers to address the complex challenges cloud presents while maximizing the value it delivers their organization. These solutions can offer your organization the same innovation and value.

Source: Deborah Martin, Product Marketing Director, Adaptive Computing

About Us

Adaptive Computing provides cloud management software. Our patented Moab® technology, provides cloud optimization to improve speed of delivery and capacity utilization, improve SLA levels and flexibility, automate infrastructure maintenance, and deliver policy-based governance for data centers and clouds. Adaptive Computing also provides intelligent workload management software that optimizes high performance computing environments that optimize HPC system and user productivity while ensuring workload SLAs are met. Adaptive Computing currently offers two cloud products to meet customer needs:

- **Moab Cloud Suite:** an intelligent cloud management platform. It provides touchless cloud service optimization across the full cloud service life cycle to speed service delivery time, reduce hardware and operational costs, maximize capacity, and optimize service performance. Moab Cloud Suite uses the patented Moab intelligence engine to automate service requesting, provisioning, optimization, and management based on multi-dimensional policies that mimic real-world decision-making. Moab Cloud Suite enables organizations to create agile, automated and adaptive private clouds from their diverse IT infrastructure.
- **Moab Cloud Optimizer:** Moab Cloud Suite policy-based optimization capabilities are also available to optimize other cloud platforms, such as HP Cloud Service Automation in the Moab® Cloud Optimizer editions. Moab Cloud Optimizer provides out-of-the-box optimization policies to optimize service placement, automate capacity management, automate maintenance tasks, and optimize service performance and availability. You can visually set and monitor its range of cloud policies to continuously maximize cloud ROI, utilization, and service performance while decreasing costs. It enables clouds to support higher capacities and lower costs for workloads and be capable of meeting the SLAs for complex applications and mainstream production services. Moab Cloud Optimizer has built-in integration for and is available as an optional extension to HP Cloud Service Automation from HP and HP Partners.



Contact Us

For more information on how these products can help you create an optimized cloud that maximizes the value delivered to your organization, contact us at solutions@adaptivecomputing.com. You can also access additional information on the web at <http://www.adaptivecomputing.com/products/cloud-products/> or by contacting our regional sales offices nearest you:

Corporate and Americas Office

Adaptive Computing Enterprises Inc.
1712 S. East Bay Blvd., Suite 300
Provo, UT 84606, USA
Telephone: +1 (801) 717-3700
Fax: +1 (801) 717-3738

EMEA Office

Adaptive Computing Limited
3000 Cathedral Hill
Guildford GU2 7YB, UK
Telephone: +44 (0) 1483 243578
Fax: +44 (0) 1483 243666

Asia Pacific Japan Office

Adaptive Computing
9 Battery Road 12th Floor
Straits Trading Building
Singapore 049910
Telephone: +65 6597-7053
Fax: +65 6597-7099



Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.