

Solution white paper

Policy-based optimization

Maximize cloud value with HP Cloud Service Automation and Moab Cloud Optimizer



Table of contents

- 3** Executive summary
- 5** Maximizing utilization and capacity to reduce hardware costs
- 6** Preventing service issues with optimized service placement
- 8** Reducing ongoing management costs
- 9** Meet your service performance and availability SLAs
- 12** Summary

Maximize cloud value with policy-based optimization

Executive summary

By now, no one needs to explain the benefits of cloud computing. In fact, you may have already begun planning or implementing a private cloud in your organization in order to satisfy the demand for immediate access to IT services, storage capacity, and computing power. You may also be looking at augmenting that private cloud with the on-demand resources of a public cloud. However, in solving the problem of high service demands, another problem has emerged: how to efficiently manage all those services once delivered in the cloud. In fact, research has shown that two-thirds of customers who have implemented private cloud solutions have actually ended up spending more money than they did before implementation. Without optimization, their automated provisioning and pooled resource utilization savings are quickly devoured by the service demand that cloud unleashes, raising their resource capacity expenses.

Moab Cloud Optimizer for HP Cloud Service Automation

- Maximize
- Optimize
- Reduce
- Ensure

These are some of the reasons why so many customers have turned to HP Cloud Service Automation (CSA). HP Cloud Service Automation lets you deliver private cloud, public cloud, and traditional IT services, and in partnership with Moab Cloud Optimizer provides policy-based optimization. Policy-based optimization maximizes utilization and service performance while reducing management costs. Moab Cloud Optimizer can help guarantee the maximum ROI and success of your HP Cloud Service Automation private cloud, making it capable of not only delivering but also optimizing and managing your development, complex application, and mainstream production services throughout their lifecycle. Moab Cloud Optimizer policies automate the real-world decisions through service tuning and ongoing service and infrastructure management. It takes data center automation to the next level to address the capacity, performance, availability, and management challenges IT teams face in meeting the service needs of their organizations.

Moab Cloud Optimizer for HP Cloud Service Automation provides out-of-the-box policies that let you:

- Maximize utilization of your cloud services by two to three times to help you reduce hardware costs with capacity management policies
- Optimize VM placement to support complex application services with granular service allocation policies
- Reduce maintenance time and costs by up to 50% with policies and management dashboards that automate repetitive tasks
- Ensure that service performance and availability SLAs are met with automated service start, reservation, and migration policies



Right out of the box

Everything works right out of the box. It's like an insurance policy that helps guarantee the optimal ROI and success out of your HP Cloud Service Automation private cloud.

Moab Cloud Optimizer has built-in integration for HP Cloud Service Automation and is available as an optional extension from HP and HP partners. Users experience all the benefits of Moab Cloud Optimizer without even knowing it's there. And for administrators, the new functionality provided by Moab is extremely easy to set up and use. Everything works right out of the box. It's like an insurance policy that helps guarantee the optimal ROI and success out of your HP Cloud Service Automation private cloud. And, it provides the guarantees, ongoing automated management, and optimization needed to support development, complex application, and mainstream production services that more and more customers are moving into their cloud environments.

This white paper will explain how Moab Cloud Optimizer can benefit your organization—whether you're interested in maximizing utilization, service performance and availability SLAs, or management time and cost savings.

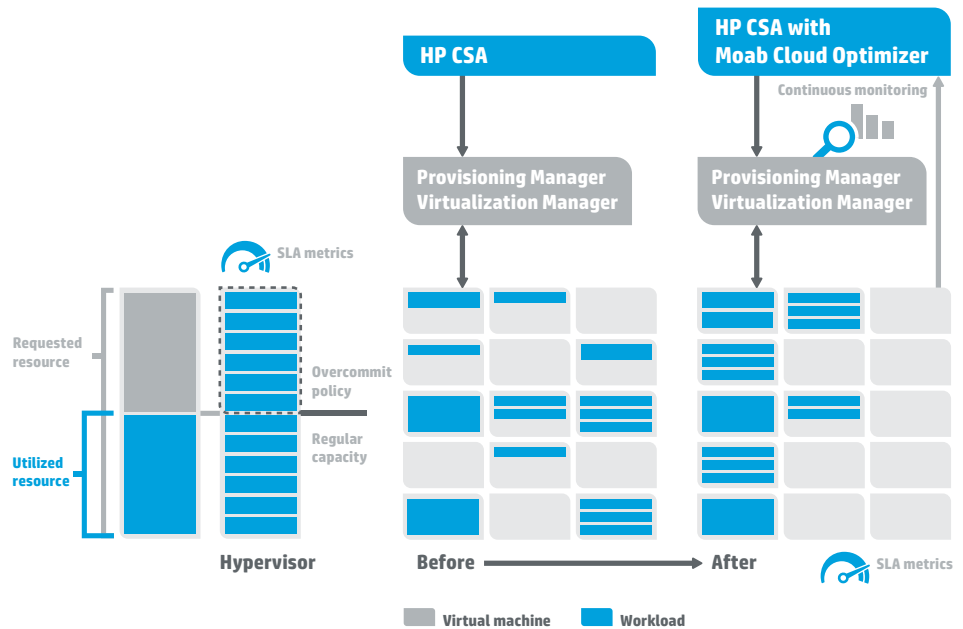
2-3X

Maximum utilization beyond what you experience with virtualization alone

Maximizing utilization and capacity to reduce hardware costs

As you know, many users request virtual machines (VMs) from a service catalog with capabilities that often far exceed their actual needs. In a typical cloud environment, those unused resources are wasted. However, when you add Moab Cloud Optimizer to HP Cloud Service Automation, you can maximize utilization two to three times beyond what you experience with virtualization alone. This allows you to significantly reduce hardware costs and truly maximize resource utilization and available capacity.

That's because Moab Cloud Optimizer uses automatic VM migration and hypervisor overcommit policies to intelligently consolidate live VM workloads onto underutilized resources while maintaining SLA thresholds. Here's how it works:



Moab Cloud Optimizer

- Overcommit capacity optimization
- Automatic migration
- Migration monitoring
- Maximum server utilization

- The Hypervisor Overcommit policy allows hypervisors to be overcommitted at levels the admin sets so they are closer to actual utilization, because users typically request more resources than are actually used. This overcommit capacity optimization is done at first placement and continues ongoing as any policy-based VM migrations occur via the policy.
- The Hypervisor Overcommit policy works in conjunction with the service performance policy with load thresholds set by administrators. These load thresholds enable Moab to identify if any of the workloads are experiencing higher demand that is requiring more of their requested resources, and therefore starting to overburden their hypervisors. If detected, the policy will automatically migrate the affected VM services to alternative hypervisors with more available resources that still meet their requirements. This ensures that service performance is maintained and SLAs are met.
- The Auto VM Migration consolidation policy continually monitors the environment and identifies sprawled VMs that fracture available capacity and could be consolidated to fewer underutilized hypervisors. Migration policy settings let administrators exclude certain VM services and hypervisors that should not be migrated and packed.
- Moab Cloud Optimizer automatically initiates VM migration actions via virtualization management tools to achieve maximum server utilization, packing the workloads for true touchless capacity optimization. The Auto VM Migration consolidation policy ensures that each VM is optimally placed as it migrates so that all service requirements will be met at its new location. This frees up larger blocks of available, usable capacity to run more services, reducing the need to buy more hardware capacity to keep up with peak and growing services demand.

All of this is done automatically without administrator intervention, although full reporting on all migrations and overcommit VM placements is completely visible to administrators in the administrator dashboard. This capability significantly boosts resource utilization and capacity for other services. It improves utilization efficiency while thresholds keep the VM services within SLA comfort zones should the actual usage load increase. Working together, the policies let administrators optimize their capacity automatically without sacrificing anything on performance.

As a result of Moab Cloud Optimizer, more resources are freed up to allow for new services. But the policies also give you the flexibility to easily exclude certain VMs or hypervisors from the migration policy if they shouldn't be considered for consolidation due to service or SLA requirements.

The automatic VM migration policy dashboard makes it easy for you to visually fine tune the automation of consolidation of workloads to prevent VM sprawl and low utilization. Similarly, administrators can easily set the hypervisor overcommit policy out of the box using its policy dashboard by just filling in a couple of fields. All it takes are a few mouse clicks and swipes. Moab Cloud Optimizer includes easy-to-set, multi-factor policy dashboards that let administrators stay in control of the thresholds that tell it when to auto migrate VMs to ensure service performance, when to consolidate to maximize utilization, and how much to overcommit the hypervisors. Essentially, they enable administrators to turn their real-world decision making into automated policies.

Preventing service issues with optimized service placement

Moab Cloud Optimizer can serve other purposes besides consolidation of workloads. It can also help you optimize service placement across diverse resources and data centers, allowing you to better support high service volumes and complex application services. By automating the optimal allocation and reservation across constantly changing resource conditions, it gives you the ability to make the best configuration decisions based on the full scope of each service request. It simultaneously considers multiple dimensions of a service—such as the application, resources, service levels—factoring in current use states, health, load, software licenses, and customizable granular resource attributes that can impact service setup and ongoing performance.

To deliver services to the business quickly and effectively, you should allocate the most optimal set of resources in the cloud pool depending on the incoming service requests. That could mean one resource may be a better fit for a certain request while another resource is better for a different request. Moab intelligently matches workloads to the right resources at a fine-grained level. In applying policies to determine the best set of diverse resources, Moab can consider almost any fine-grained request or resource attribute you can measure, including memory, processor, disk, load, data center, and cost.

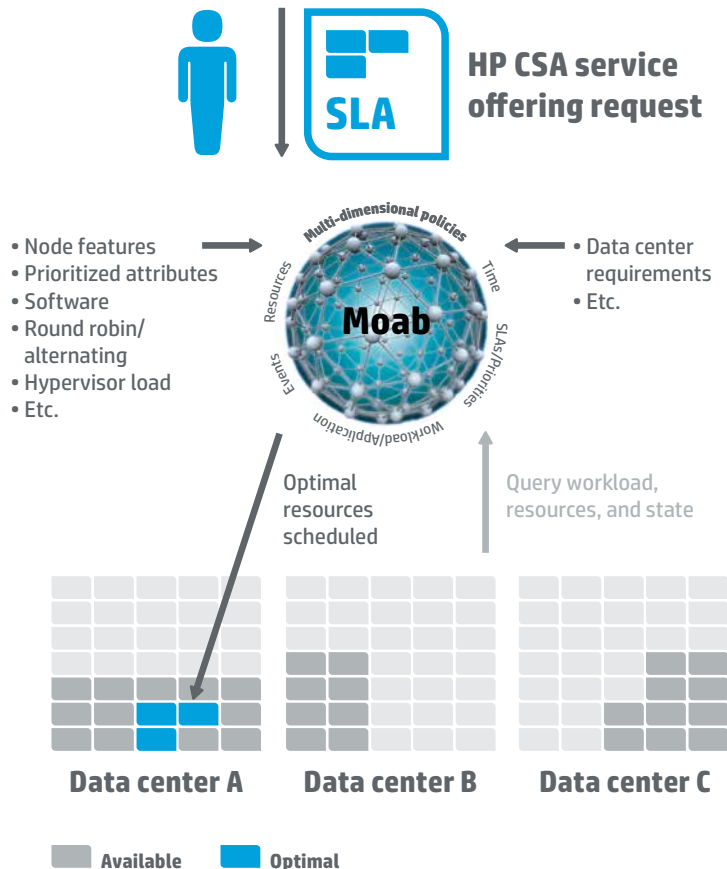
Moab Cloud Optimizer can make fine-tuned allocation placement decisions through various reservation policies, such as:

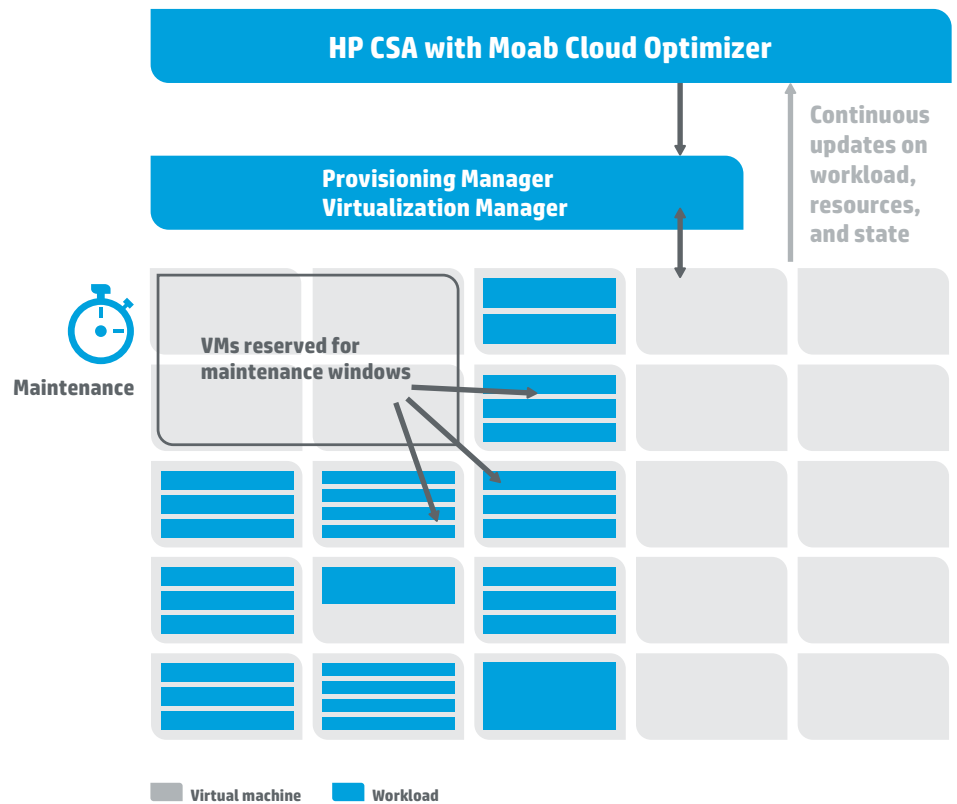
Reservation policies

- Optimize resource utilization
- Optimize workload placement
- Optimize service offerings
- Enhanced service performance

- Moab Node Allocation policies that select the optimal resources for a particular service based on granular, prioritized modeling of the services needs against resource status, health future commitments, and granular capabilities, including memory, processors, load, cost, licenses, I/O bandwidth, and other measurable and customizable attributes. Moab simultaneously prioritizes a range of factors that administrators can set, enabling Moab to filter through all the available resources and pinpoint the optimal set of resources for a service. This ensures a service will stay up and running without requiring administrative support while helping to optimize resource utilization.
- The Moab Node Allocation policy’s ability to optimize placement considering both current workload and future workload reservations is a powerful capability unique to Moab that allows you to optimize workload placement in both resource space and time; now and in the future.
- Moab Cloud Optimizer enables optimized data center placement with a service data center placement policy that can optimize specific HP Cloud Service Automation service offerings to ensure that each service is always placed in a specific data center to meet SLA, security, or other requirements.
- Moab Node Feature Tag policies automate the placement of services on nodes with required features that can enhance service performance. The Node Feature Tag policies can be easily configured by administrators out of the box and added to enhance any HP Cloud Service Automation service offering and optimize the use of any specific resource that can be reported such as high memory, fast network access, type of system (like a development system), special type of processor, or others.

Moab Cloud Optimizer also avoids unhealthy resources by constantly monitoring current health status data and filtering them out so services are not placed on them to avoid potential service issues. By considering the full range of policy points for service placement, you can avoid post-provisioning service failures and optimize the utilization of cloud resources.





50%

Reduction in maintenance cost and time

Reducing ongoing management costs


In addition to helping you make smart placement decisions, all of this automatic allocation and optimization also means you won't have to spend so much time managing your cloud services once they are provisioned. With Moab Cloud Optimizer and HP Cloud Service Automation, you can reduce maintenance costs and time by 50%. Moab maintenance reservation and administrator policy and management dashboards can help you automate repetitive management tasks. This enables your existing IT staff to scale to support the quantum increase in volume of new services that cloud enables and even shift to support new projects.

When making resource selections, Moab Cloud Optimizer can analyze historical data and future resource commitments to get the most utilization out of resources across time. Without Moab, HP Cloud Service Automation can also provide that you will have necessary resources in the future. But Moab has an internal understanding of space and time and keeps a calendar of how long resources are reserved for and used. Armed with this knowledge, Moab ensures resources are available when needed. Moab also allows other requests to proceed as long they expire in time to free up the resources by the time they are required.

Organizations can use this unique capability to plan for system maintenance, such as updating system resources. Usually these tasks must be done with manual tweaks and changes, but with Moab Cloud Optimizer, these tasks can be planned for, resources reserved, and processes automated to ensure no disruption in services or drain on staff time occurs. All other services are placed around the future reservation, and any existing services are automatically migrated off to other resources that meet their needs when the reserved time approaches. This way, running services continue to run uninterrupted and administrators don't waste their time manually migrating services and even potentially placing services in less than optimal locations that may cause issues leading to additional administrative workload.

You can further reduce administrator management time and costs with the administrator policy and management dashboards provided in Moab Cloud Optimizer. They automate more of those repetitive manual service monitoring and management tasks that can limit your staff's ability to scale across more services and to shift from maintenance to new projects that drive the business forward. The range of dashboards in the Moab administrator portal simplifies managing the cloud environment by helping administrators track, diagnose, and take action on issues as well as view utilization to plan for future capacity and set or fine tune policies to optimize cloud performance. The dashboards visualize the cloud environment and sophisticated policies so they can be easily managed out of the box. They also help filter and summarize events, services, and resources to enable quick action. This means administrators spend less time collecting information and identifying issues and more time on optimizing performance and other strategic IT projects.

Admin dashboards accelerate and automate cloud management tasks



Meet your service performance and availability SLAs

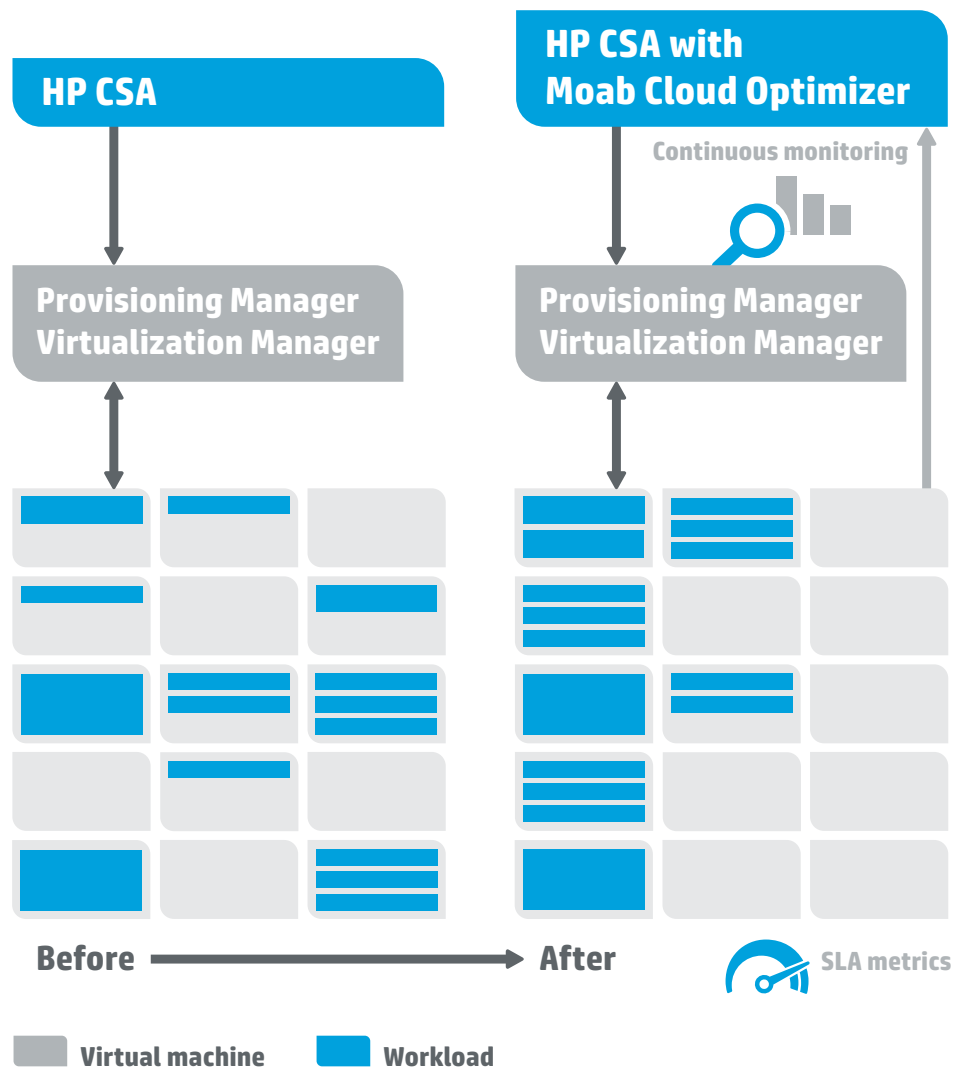
Cloud optimization strategy

Moab Cloud Optimizer and its Automatic VM Migration policies intelligently migrate live VM service workloads from servers as they become overloaded and threaten to exceed SLA thresholds.

Of course, any cloud optimization strategy would be incomplete unless service performance and availability SLAs are consistently met for the supported key test and development or production services. That's why Moab Cloud Optimizer and its Automatic VM Migration policies will always intelligently migrate live VM service workloads from servers as they become overloaded and threaten to exceed SLA thresholds. This helps you adapt your infrastructure as service demand changes, and ensures that you will be able to avoid service performance issues and the risk of any disruption to the business. It works because of the way the Moab Cloud Optimizer policies function:

- The Moab Auto VM Migration Performance policy continually monitors the cloud workload and automatically initiates live VM workload migrations as the servers they are on become overburdened to avoid service issues and maintain service performance levels.

- The multi-factor policy settings let Moab accurately determine when services are at SLA risk based on the IT team’s best practice definition of when a hypervisor is overloaded and may impact service performance. The policy can be based on multiple reported factors that the administrator identifies as most indicative of a hypervisor becoming unable to meet SLAs. This could be either metrics from vCenter or from performance and monitoring tools. The policy metrics are easy to set in the visual policy management dashboard with just a few mouse clicks. In migrating each service, Moab makes sure all the service requirements will be met in the new service location with true touchless optimization.
- This policy is scalable to thousands of nodes across multiple virtual data centers to maximize SLA performance and efficiency instead of just limiting optimization and migration to within small 32 node virtual data centers like other cloud solutions.
- The complementary Moab Migration Exclusion policy enables administrators to identify hypervisors or virtual machines with SLA requirements that Moab excludes in migration policy actions. This policy can also be applied to specific service offerings so that they will not be migrated once running.





Organizations can also use Moab Resource Reservation and Service Start policies to ensure service availability SLAs are consistently met for users. These policies guarantee that the right resources for services will be available at the right time to eliminate disruptions or delays for important activities—like quarter-end financial services or critical product development initiatives—to users and the business. This also helps to reduce management complexity and costs, as administrators don't have to manually remove, then re-add resources to the cloud pool. It's a great tool for overcoming critical team worries about moving services from dedicated resources to the cloud.

The Moab Resource Reservations policy ensures that a certain set of resources will be available now or in the future for a specific user to set up or expand services on. Existing or new services will be automatically migrated off and placed on other resources by Moab when the reserved time is approaching. This ensures priority user workloads always have resource availability.

The Moab Service Start Time policies are integrated right into the HP Cloud Service Automation portal and enable users to make their request for a service to start at a current or future start time and date. They guarantee that those resources will be available and that the service will be auto provisioned at that time, acting as an availability SLA. Without Moab, the requested date would be just that—a request—which may or may not be able to be honored. The start policies also enable Moab to suggest the earliest available start time if capacity is not available for a specific service at a specific request time, to better deliver against SLAs. Moab is able to do this with its unique ability to schedule both for current and future capacity.

Summary

Industry analysts are talking about the types of capabilities that Moab Cloud Optimizer and HP Cloud Service Automation provide as key for the success of organizations as they move from data center automation and virtualization to the cloud, especially as customers push into production-oriented clouds.

Gartner wrote, “As service-to-infrastructure ratios increase, more end users have the potential to be affected. As automation increases the pace of change in the cloud environment, constant optimization is needed to keep the system as a whole from drifting toward poor utilization.” The report also noted, “Cloud operators must understand the impact of changes on performance and service levels, correlate physical to virtual to application and service performance, and ensure that the right resources are allocated with the appropriate capacity needed to meet service levels.”¹

Moab Cloud Optimizer and Cloud Service Automation provide the capabilities and policy automation that help IT teams tackle these challenges.

Engineers from Adaptive Computing and HP have worked together very closely to provide that HP Cloud Service Automation and Moab Cloud Optimizer not only complement each other very well, but are tightly integrated, easy to use, and easy to administer. Moab Cloud Optimizer gives you the policy-based optimization that helps ensure the success of your cloud project across all the bottom line metrics that matter. If you are tasked with making sure your cloud delivers on the value to your business, including fast service delivery as well as ongoing service optimization and management to maximize cost savings and service performance, then you owe it to yourself to take a look at Moab Cloud Optimizer. Contact your HP or HP partner representative today.

**For more information visit
hp.com/go/csa.**

¹ Gartner Research Report, *Configuration and Performance Management Useful for Virtualization, Required for Cloud*, August 2012.

