



Planning and Implementing An Optimized Private Cloud

Creating a Cloud Environment That Maximizes Your ROI

Table of contents

3	Why Cloud Computing?
3	Clouds Come in All Shapes and Sizes
5	Where Are You on the Cloud Computing Continuum?
8	The Experts Weigh in
9	Maximizing Your ROI with Moab Cloud Suite
11	The Power Behind Policy-Based Optimization
11	The Advantages of Working with Adaptive Computing

Planning and Implementing An Optimized Private Cloud

Why Cloud Computing?

These days, there is so much hype about cloud computing that it's easy to lose sight of why it's such an important trend. Cloud computing offers IT management and their staffs a tremendous opportunity to solve critical challenges in the data center to better serve business. Today's typical data center is a sprawled, siloed environment with physical and virtual resources that are not utilized to their full capacity or fully aligned with business priorities. What's more, the manual request and delivery processes of today's typical data center can be painfully slow. Those processes simply can't keep up with business demands for IT services and, in many cases, business groups are running out of patience. In fact, many are plunging headlong into public clouds, often without IT input and without considering security issues and other "fuzzy" aspects of cloud computing.

Cloud computing holds the promise for IT management to transform their data centers into highly efficient, largely automated service providers that consistently meet service level guarantees, accommodate self-service requests, allow charging back for services based on actual use, and offer automated reporting to track and improve cloud performance once implemented.

In other words, cloud computing is the right technology for the times. But because the times are changing at breakneck speed, IT decisionmakers have to retain as much flexibility as possible when choosing a cloud computing solution. Special care must be taken not to get locked in to proprietary technologies that limit what you can implement today, and that may be obsolete tomorrow.

Best-of-breed, heterogeneous environments are the wave of the future—environments that allow you to enforce various levels and types of SLAs as you move forward with your cloud strategy, make adjustments, and add different types of users and services. It all comes down to the cliché, "one size does not fit all." In cloud computing, this phrase takes on a whole new meaning.

The potential benefits of cloud computing include:

- ▶ better ability to meet business groups' expectations
- ▶ high availability
- ▶ massive scalability
- ▶ lower costs of hardware, software, electricity, heating and cooling
- ▶ reduced square footage requirements
- ▶ IT staff that is freed from the burden of repetitive manual tasks and can therefore deliver more strategic value
- ▶ less human error
- ▶ better compliance with SLAs
- ▶ better compliance with regulations
- ▶ utility computing capabilities

But the operative word is potential. If you are in the planning stages of your cloud solution, then you are assessing the ways to realize that potential. It's a critical time. You want the most positive vision—the best-case scenario—to become a reality. So, how do you make it happen? Clouds are gathering in data centers the world over. But, unlike Mother Nature's clouds, how do you make yours predictable and consistently attuned to your wants and needs?

Clouds Come in All Shapes and Sizes

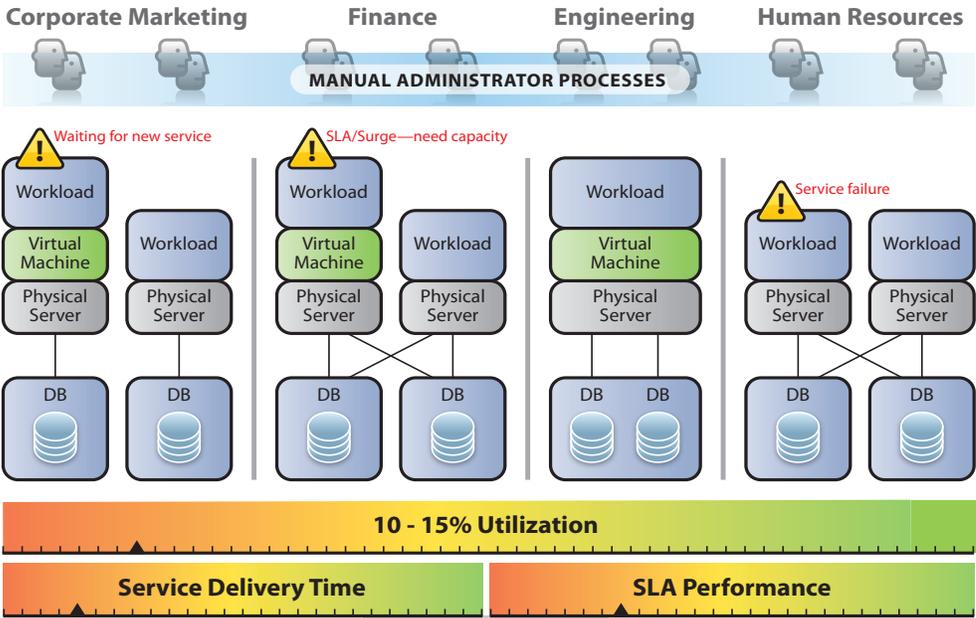
Cloud computing remains a nebulous term and can be defined in all kinds of ways. However, in its ideal form, cloud computing is a highly efficient, cost-effective, and intelligent IT service delivery model. Whether public, private, or hybrid, an optimized cloud solution is customizable to your business and constantly aware of the various physical and virtual resources and workloads in the environment. In addition, the ideal cloud solution offers a level of intelligence and adaptability that can optimize IT resources "on the fly" in order to consistently meet or exceed business requirements.

Within the scope of this paper, we are only covering private cloud strategy considerations because, at least for now, this is where most organizations will reap the most benefits of cloud computing while maintaining optimal control and security. But even within this limited scope, there are wide-ranging options and opportunities for migrating to public clouds or hybrid cloud computing environments when needed.

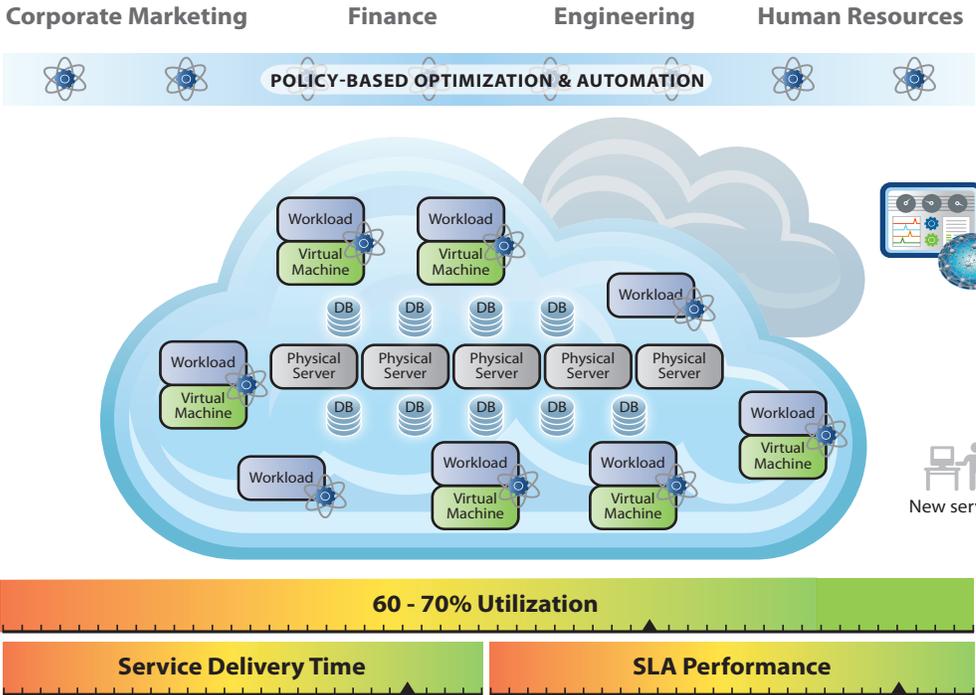
There are three primary cloud computing service models IT departments can offer their businesses:

- ▶ **Infrastructure-as-a-Service (IaaS)** – IT or a cloud service provider owns and administers the backend hardware; the business's IT department takes care of the software and client computing side. Users can request and provision key aspects of infrastructure, including processing, storage, networking, and operating systems. IaaS lets users build their own computing environment on demand without requiring the business to make substantial capital investments in hardware.
- ▶ **Platform-as-a-Service (PaaS)** – an application platform in a cloud. The business's applications, whether in-house-developed or off-the-shelf, are deployed on a cloud provider's or internal private cloud infrastructure. Users can request and provision their own applications using programming languages and tools such as Java, Python and .Net—whatever is supported by the provider. The user does not manage or control the underlying cloud infrastructure, network, servers, operating systems, or storage but retains control over the deployed applications and may also control application hosting environment configurations.

Traditional Data Center



Optimized Cloud / Data Center



Why Cloud Computing?

- ▶ **Application-as-a-Service (AaaS), or Software-as-a-Service (SaaS)** – IT or the cloud service provider supplies the applications and infrastructure on-demand via the cloud. End users simply access the full application service stack, usually from a browser. The user does not manage or control the underlying cloud infrastructure, network, servers, operating systems, storage, or usually even individual application capabilities, with the possible exception of limited, user-specific application configuration settings.

These brief definitions convey only part of the story: service delivery methods. That’s all well and good, but flexibility, scalability and availability are what business professionals expect and IT managers need. So, a cloud solution must support ongoing, continuous optimization to keep services up and running at the highest service levels, including automation of decommissioning and repurposing resources at the appropriate times. This is an important consideration many IT organizations overlook. It’s important to keep in mind that cloud solutions enable a quantum leap in the speed and number of services that can be delivered and still need to be fully managed. You must be prepared to manage all those services day in and day out through their entire lifecycle. That is why an optimized cloud strategy is so vitally important when it comes to making your cloud solution as successful as it can be.

Where Are You on the Cloud Computing Continuum?

Whether you are just starting to feel your way into the cloud or your first cloud implementation is well underway, it makes sense to include IaaS, PaaS and AaaS service models in your long-term plans. Contrary to an all-too-popular trend, which is to dive right in and obtain basic delivery of a static cloud environment, it pays to develop a comprehensive and optimized cloud strategy to ensure your cloud will be able to grow with your needs as you expand the number of services you are delivering and managing through the cloud. Only with an optimized cloud strategy is it possible for IT to successfully support and manage cloud-based services and resources at required service levels and with maximum efficiency—even as the numbers of service requests grow. Moreover, an optimized cloud strategy includes the policy-based automation foundation that can bring together private clouds, public clouds and multiple data centers to provide the overflow capacity that many organizations will want to have in the future.

To some extent, where you are on the cloud computing continuum is a function of how many challenges you’re up against today. The following use case areas represent three of the more common stops along the way toward a dynamic, optimized cloud solution. Each use case adds to and builds on the previous use case area’s capabilities. Consider each set of use case capabilities as you create the optimized cloud strategy for your organization, and as you evaluate cloud solutions.

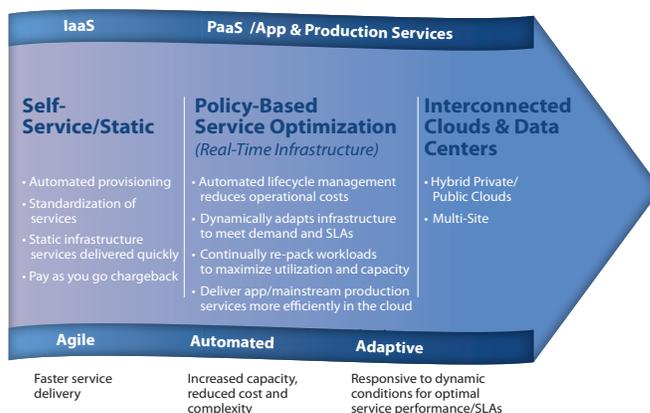


Figure 1. Agile, Automated, Adaptive: Three key stops along the cloud computing continuum.

Use Case Area 1: Agile Service Delivery



Rapid and successful delivery of static IT services and resources to users in minutes for the duration of their self-service request with automated allocation decisions, provisioning, and chargeback based on business policies.

A difficult challenge for IT managers today is that IT services can take weeks or months to deliver, instead of minutes or a day—the quick turnaround times that the business needs to deliver new offerings to customers, respond to market conditions or improve processes. And, of course, business teams and individual users don’t realize (or care) how much IT services cost. Agile Service Delivery, the first stage of a cloud implementation, can resolve these issues. It can deliver business services and resources rapidly while tracking, reporting and charging back the cost of services. Agile Service Delivery can be accomplished by designing and implementing a fast, intelligently automated, self-service, static Infrastructure-as-a-Service delivery model.

Key capabilities for an Agile Service Delivery solution should include:

What every organization should expect from an innovative cloud computing solution are agile delivery and ongoing service optimization and management capabilities to ensure long-term cloud computing success and ROI.

Planning and Implementing An Optimized Private Cloud

 **Self-service, portal-based user creation and management of services**, with catalog-based service template selections that automatically specify available services, the optimal resources needed, and their costs. This includes multi-group management that tailors cloud access, service views and management options to meet the specific needs of multiple different groups, users, or admins using the same enterprise cloud infrastructure.

 **Service catalog templates** that identify the best-practice resource provisioning amounts and processes for each service, as determined by IT. This is vitally important, as lack of best-practice allocation and provisioning leads to service failures, related business disruptions and higher management costs. For example, Virtual Machines (VMs) that are set up on servers without best-practice provisioning can easily end up with not enough memory allocated for the service(s) running on them, and failure ensues. Alternatively, without service definitions in place, many users will overprovision the resources actually needed for their service, leading to wasteful unused capacity. To maximize the value your cloud can offer catalog templates should support the range of virtual machine, physical, and full application services with the option to define custom services for certain users.

 **Optimized service placement** ensures the intelligent matching and allocation of heterogeneous resources via sophisticated policies that model each service request's needs against prioritized current and predicted diverse resource attributes and state. This ensures that the services are allocated the optimal set of physical or virtual resources from the diverse cloud pool—optimized to avoid failure, business disruption and additional management costs and headaches. Policies should be multi-dimensional, taking into consideration current availability and usage levels, historical performance, future reservations or maintenance, and diverse resource attributes. Those attributes should include everything from compute capabilities to active software licenses to ensure that resource allocations are appropriate. This level of policy sophistication can help reduce service failures and costs while optimizing resource utilization and available capacity for other services. It is also key that service placement include policies to pack workload on hypervisors to maximize utilization and capacity as it takes advantage of differences in requested versus actual usage.

 **Automatically initiated provisioning and deprovisioning via workload templates** to rapidly deliver heterogeneous physical and virtual resources using best-practice workflows instead of numerous manual IT

staff tasks. This also removes human errors and compliance inconsistencies from the environment, which are leading causes of instabilities, failures and downtime. Automating these repetitive tasks also ensures the quick decommissioning and repurposing of resources for new services.

 **Chargeback or “showback” for costs of resource use by user, group, and account**, including flexible rates and charging models. This enables IT to re-capture IT resource costs and gain an understanding of IT capacity requirements and costs for different business priorities as the data center transforms to a service provider role. An important requirement is delivery of an estimate of the cost of the service at the time of request, which helps users make better informed decisions on the resources and service levels that will meet their needs and budgets. Ongoing usage cost visibility helps ensure services and resources that are no longer needed get decommissioned so they can be used for other services.

 **Administration dashboard for reporting and management**, including both real-time and trend reporting and tracking of cloud policies, events and activity. This lets IT easily address issues and provides the basis for auditing, performance optimization, planning for capacity, and reporting to the business on value and service delivery.

Agile Service Delivery quickly delivers services and informs IT about current utilization and workload and resource activity. However, its downside is that it leaves services in a static IaaS mode. It doesn't allow you to automatically manage or scale up resources significantly as demand or resource conditions change for a service. The intelligence, dynamism, and control simply aren't there. To reach the full potential of cloud computing, you need capabilities that automate the management of services that are running while adapting your resources to the situation as service demand and conditions change.



Use Case Area 2: Automated Management

In this phase, the cloud environment begins to dynamically adapt through policies based on workload and conditions—self-managing services to maximize utilization and capacity while reducing management cost and complexity.

Planning and Implementing An Optimized Private Cloud

Manual ways of managing services and resources that are currently in place aren't allowing IT staff and management to keep up with IT service requests, especially at the accelerated service-creation rate in cloud environments. In this use case, the objective is to automate ongoing management so both the computing resources and IT staff can better scale to keep services up and running. By automating repetitive manual processes and applying resource, user, group and workload policies and templates—thereby optimizing utilization and capacity of existing physical and virtual resources—IT can effectively meet service delivery expectations and reduce operating expenses. Risk reduction is also a key objective, and it can be accomplished by automating the enforcement of SLAs for cloud service requests with policies that automatically prioritize requests and enforce usage limits. This is key in the shared resource and multi-tenant environment of a cloud.

This level of automated service management eliminates human error, enables resource capacity and utilization maximization, and reduces management complexity and costs. It frees IT staff for new projects and enables them to scale so IT costs are minimized while service to the business improves. And, of course, it helps keep business service users happy, as common resource incidents are responded to immediately and automatically, before users experience any disruption.

Key capabilities for an Automated Service Management solution should include:

 **Automatically initiated live workload migrations to maximize utilization** with VM migration consolidation policies that monitor and pack VM workloads onto underutilized resources to maximize capacity for new services. These automated VM migration policies should eliminate VM sprawl and reduce the amount of IT staff required to manage the growing number of nodes and VMs for IT services.

 **Intelligent power management policies** that reduce power consumption on under-utilized resources, or even power down idle resources and intelligently power them back on in preparation for anticipated workloads. This capability lets you automate decisions as to when to consolidate and migrate live VM workloads from underutilized resources onto more fully utilized resources. As a result, the power usage of underutilized resources can be dramatically reduced.

 **Automated response to common resource incidents and issues** for immediate resolution to eliminate or reduce service disruption and IT staff burden. Your cloud solution should provide you the ability to set custom trigger policy workflows and actions based on specific events or conditions to standardize management and automate repetitive response tasks for the infrastructure.

This also further reduces management costs and headaches for IT staff while ensuring a higher quality of service to users as potential service disruptions are dealt with quickly, often before they impact users.

 **Automate the reservation and running of maintenance jobs across cloud resources** to avoid disruption of regular workloads, reduce management complexity, and increase resource compliance. Resource sets should be able to be identified for maintenance and compliance updating and then have their workloads automatically migrated off to other resources during the maintenance period and automatically migrated back when the maintenance is complete. This ensures that all resources stay in compliance so known issues are avoided before they cause business disruption. Moreover, all manual workload migration tasks are removed for the IT staff, which ensures that new errors don't get introduced to your running services that were initially auto-provisioned from best practices.

Automated Service Management can improve service delivery by leaps and bounds by simply making your complex cloud and data center environment easier and more cost-effective to manage—and much more efficient. It reduces the burden of repetitive tasks and improves compliance while maximizing resource utilization. But it does leave one very important area still left to plan for in your cloud strategy and solution: dynamically adapting and switching out resources to respond to demand and unexpected resource failures. This last step on the cloud continuum is critical to meeting SLAs and ensuring that the services you deliver stay up and running at the performance level the business needs.



Use Case Area 3: Adaptive Cloud Services & Resources

In this step on the cloud continuum, the cloud infrastructure is workload-driven as well as request-driven, dynamically adapting the infrastructure based on workload demand needs and resource conditions, so the cloud is able to self-optimize to meet SLAs. This is a key consideration to enable your cloud to deliver a broader range of services and RO, such as mainstream production services, going beyond just test and development workloads.

Planning and Implementing An Optimized Private Cloud

At its most versatile, cloud infrastructure is adaptable and elastic so it can respond to changing cloud services demand or resource conditions. It should effectively drive dynamic and automatic allocation and optimization of IT resources and make service levels predictable and consistent without IT staff intervention. Adaptive Cloud Services are capable of this level of elasticity because they are based on Real-time Infrastructure (RTI) technology. RTI automation and intelligence provides policy-based optimization and tuning of the runtime environment. What's more, RTI technology reduces risk, ensuring the services the business needs are delivered without interruption and with maximum efficiency.

Adaptive Cloud Services are defined by dynamic service optimization and resource self-healing driven by policies based on designated resource states and performance metrics in order to reduce management cost and complexity and, most importantly, ensure needed service levels are delivered to users and the business. One of the key benefits of cloud computing is the more cost-efficient and responsive use of IT resources to enable the business to deliver more innovations, customer offerings and streamlined business processes faster through IT-based services. This full vision and value is only realized when the policy-based optimization is in place to ensure that the services stay up and running to needed performance levels, which often requires changing the types and amounts of resources available to a service. Where resources are constrained, automated business policies determine how resources are adapted across services to meet business goals. Having Adaptive Cloud Services to complete your optimized cloud strategy and ensure SLAs are met can help reduce the resistance of some business teams in moving their workloads and services to the cloud.

Key capabilities for an Adaptive Cloud Services solution should include:



Automatic live VM workload migration policies to maintain performance that can be set to move workloads from resources as they become overburdened in order to

avoid service performance issues and business disruption. Service workloads should be constantly monitored and automatically migrated to alternative resources when load and performance thresholds are reached, thus ensuring continued optimal service and SLA performance.

Self-healing trigger policies to automatically respond to specific physical or virtual resource pre-failure or failure states—to ensure needed services stay up and running. Constant real-time monitoring and automated re-provisioning and workload migration ensure that failures are healed before they cause a severe service disruption and help ensure SLA uptimes are met.

Automatically enforced resource reservations and usage quotas ensure availability SLAs are met to reduce business risk and ensure satisfaction. These policies can help resolve conflicts between different user, organization and application service requests for limited, shared cloud resources, and help to ensure that service levels are met, especially in multi-tenant situations. Future resource reservations guarantee specific resource availability for a certain time to meet critical group or service SLAs. Soft and hard resource usage limits should be able to be set and enforced for users and new requests prioritized appropriately during periods of high service and resource demand.

Intelligent cloud bursting and data center bursting for dynamic overflow capacity when needed by setting up services on other internal private clouds or data centers or even public clouds via hybrid cloud management as your cloud solution matures. This should include capabilities to automate migrating services back to your private cloud as capacity becomes available.

The Experts Weigh in

According to a recent Gartner Research Report,¹ there are many use cases for emerging private cloud-computing and service optimization solutions, and nearly as many requirements. That's why Gartner did the research that is the basis for the report—to examine the various requirements and help decisionmakers to identify where their data center operations, processes and business needs fit. These, in turn, provide the basis for building their cloud management strategy and can assist in creating an RFP for cloud management tool evaluation.

The Gartner report stressed the emerging customer cloud requirements for integrated capabilities across infrastructure provisioning, configuration, and performance management. Moreover, the report noted that, "As service-to-infrastructure ratios increase, more end users have the potential to be

Look for a policy-based, optimized cloud computing solution that delivers capabilities across all three essential areas—Agile Service Delivery, Automated Service Management, and Adaptive Cloud Services. Together, they enable you to evolve your strategy at your own pace, without management, technology, scale, or performance limitations.

Planning and Implementing An Optimized Private Cloud

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.” In evaluating cloud management platform vendor capabilities in a subsequent report², Gartner evaluated Adaptive Computing’s capabilities with Moab Cloud Suite in this area.

It is this constant and difficult day-to-day service optimization that Moab Cloud Suite provides with easy to use policies. Policies that integrate with your existing monitoring and management tools to automate optimization actions. This dynamic service optimization is a key capability as you look to move business and production applications into your cloud infrastructure now or in the future.

Maximizing Your Cloud ROI with Moab Cloud Suite

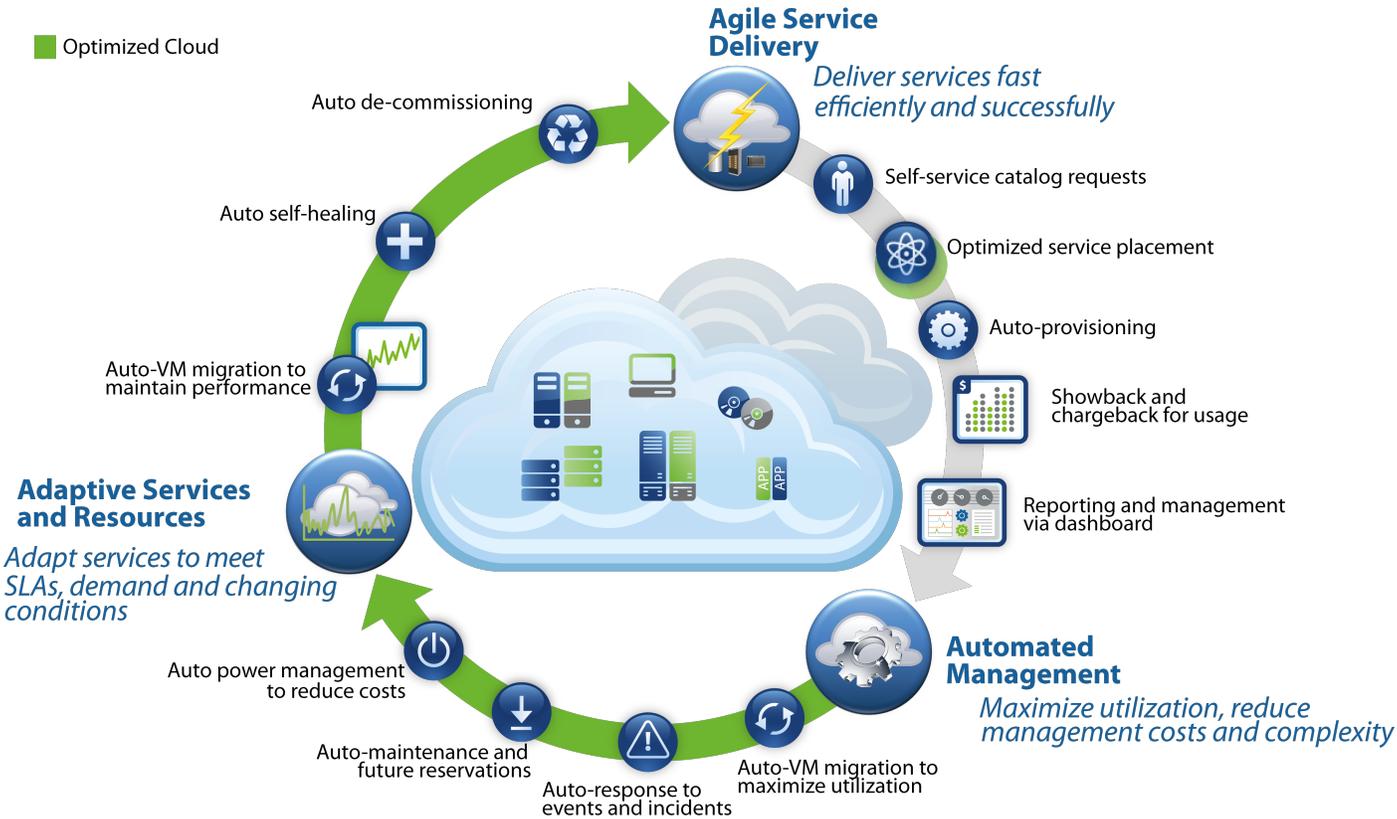
You need to be able to leverage the full benefits of cloud computing for your organization. For that to happen, your cloud solution must be agile, automated and adaptive—unlimited when it comes to evolving and adapting to meet your business needs.

Moab Cloud Suite meets all the criteria to create an ideal private cloud solution for your organization. It is the cloud management platform that accommodates short-term planning that focuses on implementing a static, IaaS cloud strategy in the data center. Moreover, the Moab Cloud Suite policy-based optimization dovetails with long-term planning that calls for delivering more complex and mainstream production applications in your cloud solutions. And, most importantly, Moab Cloud Suite policy engine provides the automation intelligence foundation for implementing a private cloud solution today that has the inherent flexibility to extend to a private/public hybrid solution in the future.

Moab Cloud Suite’s policy-based cloud optimization engine can help you create an agile, automated and adaptive cloud environment that responds faster to business requests and automates across IT processes. Moab Cloud Suite can manage vital tasks such as maximizing resource capacity, recommending the best course of action to resolve issues, and radically reducing the time required to acquire, deliver and optimize fully functional compute platforms and services that are responsive to dynamic business conditions. Moab Cloud Suite intelligently automates and optimizes across the lifecycle of services in the data center or cloud and their resources. It’s the cloud solution that enables:

- ▶ **Self-service catalog requests for services** via an easy-to-use portal to speed request and delivery while reducing IT staff burden and costs. It supports the creation of physical, virtual and complex application services with multi-group management
- ▶ **Optimized service placement** with allocation and hypervisor overcommit policies that intelligently model and allocate the optimal resources to prevent service failure, pack workload and maximize utilization.
- ▶ **Auto provisioning** to quickly deliver complete physical and virtual resources in minutes with service templates that reduce provisioning costs by 60-70%.
- ▶ **Chargeback or showback for service and resource usage** to enable pay-as-you go, including flexible rates and charging models by user and group, with optional usage quota enforcement.
- ▶ **Reporting and management via a customizable administration dashboard** to monitor cloud operations, identify and address issues, and optimize performance with easy to set visual policy controls.
- ▶ **Auto consolidation of VM workloads to maximize utilization and reduce power costs** with an intelligent VM migration policy that eliminates VM sprawl and an auto power management policy that triggers idle servers into hibernate or sleep states and back into idle for incoming workloads.
- ▶ **Automatic maintenance and future resource reservations** automate the tasks to prepare and guarantee availability of resources for maintenance activities and to meet specific user or group SLA requirements. These policies avoid disruption of regular workloads, reduce management complexity and increase the compliance of cloud resources.
- ▶ **Auto live workload migrations to maintain service performance**, with VM migration policies that move VM workloads from resources as they become overburdened to avoid service performance issues.
- ▶ **Auto self-healing and incident response trigger policies** speed service and resource recovery by automating standard processes such as provisioning replacement servers and migrating workloads, diagnosing, and restarting or reprovisioning if a server fails.
- ▶ **Auto de-commissioning of service resources** by initiating de-provisioning workflows and processes at the end of the lifecycle so they can be quickly repurposed for new services.

Optimization & Automation Across the Service Lifecycle



The ideal cloud solution must encompass policy-based optimization. It's a key requirement to ensure that services are delivered quickly and successfully, and that they dynamically stay up and running across their lifecycle so they meet SLAs, with little or no manual intervention.

Planning and Implementing An Optimized Private Cloud

The Power Behind Policy-Based Optimization

One of the key differentiators of Moab Cloud Suite is its ability to intelligently adapt the environment to meet unexpected changes in demand or resource conditions. This is essential for ultimately reducing risk and ensuring that the services that the business needs are delivered without interruption and with maximum efficiency. Something else that other solutions can't match is its unique self-optimizing capability. Moab Cloud Suite accomplishes this with a flexible management abstraction layer that aggregates real-time data across the complex cloud and data center environment to create granular resource models. It combines this with industry-leading multi-dimensional analytic policies that simultaneously factor workload requirements, resource attributes, SLAs and priorities—to drive automated decisions and actions related to service creation, management, adapting and decommissioning. Moab Cloud Suite is specifically designed to integrate with heterogeneous environments of physical and virtual resources, management tools, hypervisors and operating systems, and be able to act as master controller or peer controller, or in a supporting role with other management tools. And, critically important, it is able to look at both your current real-time environment and future scenarios. Thus, it becomes much easier to plan for future needs, including factoring in resource maintenance windows and provisioning and deprovisioning times in order to optimize resource use and energy efficiency.

Not surprising, Moab Cloud Suite is a key component within the world's largest financial institution's private cloud architecture. It's the cloud solution that is letting this global bank consolidate and eliminate redundant and underutilized hardware—driving greater usage of commoditized middleware platforms and increasing efficiency and savings through power reduction and optimization. As a result, the bank is poised to deliver IaaS and PaaS solutions based on application workloads—providing customers with superior service and projected savings to the bank of \$1 billion in traditional IT costs over the next three years.

The Advantages of Working with Adaptive Computing

Cloud computing is a vast new frontier. As such, it represents the potential for tremendous rewards for those who venture in without delay. By the same token, it imposes high risks on anyone who is not careful about how she or he proceeds.

To avoid getting snakebitten, you want a vendor with the experience to guide you on the most direct path to realizing your goals. Adaptive Computing's professional services teams offer proven expertise in delivering control and automated management for massive scale and complex environments. Their knowledge is derived from more than a decade managing infrastructure in the most scale-intensive and complex environments in the world, not to mention working as pioneers in the high-performance computing (HPC) and cloud computing evolutions.

In addition, Adaptive Computing products provide management for the world's largest computing installations and are the preferred policy-based optimization solutions for the leading global data center vendors. Adaptive Computing goes to market with HP, IBM, SGI, Cray, and other leading technology companies, which allows us to deliver dynamic cloud solutions to organizations of all sizes. Together, we enable customers to access the full potential and benefits of elastic cloud capabilities across the entire service lifecycle through Moab real-time Moab optimization policies.

Visit our website or call us to discuss how we can help you design a dynamic cloud strategy for your organization. We'd be glad to demonstrate the capabilities of an optimized cloud, and discuss best practices that have worked consistently for our clients. Best practices that can work for you.

¹ "Configuration and Performance Management Useful for Virtualization, Required for Cloud," Ronni J. Colville, Aneel Lakhani, Gartner Research Report, August 29, 2012.

² "CMP Vendor Landscape," Ronni J. Colville, Donna Scott, Milind Govekar, Aneel Lakhani, Gartner Research Report, September 5, 2012.

Let's talk...Set up a Demonstration...and Test in your Environment

An Adaptive Computing solutions advisor can guide you to the products and services that will best meet your needs and will work with you to set up a live, online demonstration designed specifically for your organization

Contact a solutions advisor by phone or email, or visit our Web site today

North America, Latin America +1 (801) 717.3700
 Europe, Middle East, Africa +44 (0) 1483 243578
 Asia, Pacific, Japan, India +65 6597-7053
 Email: solutions@adaptivecomputing.com
www.adaptivecomputing.com

Corporate Headquarters

1712 S. East Bay Blvd. Suite
 300 Provo, Utah 84606

