MAR | 2011

CLOUD COMPUTING

Making the Cloud Achievable

800.808.4239 | CDWG.com/cloudguide

CDW·G REFERENCE GUIDE

A guide to the latest technology for people who get IT

CDW·G PEOPLE WHO GET IT™

CLOUD COMPUTING REFERENCE GUIDE

IN THIS ISSUE:

DEFINING CLOUD COMPUTING

WHAT THE CLOUD IS AND WHAT IT OFFERS

IT departments are constantly undertaking programs to improve efficiency, optimize productivity, increase agility and streamline costs — all while lowering the total cost of ownership. This quest for constant improvement is a key trait of successful IT management. In an era of constrained budgets and tight staffing, there is an even greater need for organizations to shift resources from mundane operational tasks to activities that will yield greater value.

This drive in the 1980s and early 1990s led many organizations to migrate away from proprietary computing platforms and instead standardize on commodity, x86-based systems. More recently in the 2000s, many organizations consolidated physical server hardware using virtualization technology; shrank data center footprints; embraced web-based computing; and built service-oriented architectures (SOAs) that facilitated the easy reuse of application components.

All of this innovation has paved the way to cloud computing, the highly dynamic, next-generation model for IT.

Many organizations that closely examine how best to meet their efficiency, optimization and agility mandates come to the conclusion that their IT infrastructure needs to be more dynamic. Traditional IT, which binds applications to specific physical hardware, such as servers, storage and network devices, is too rigid to respond to organizational imperatives in a timely fashion.

In such environments, provisioning and testing the infrastructure pieces necessary to support a new application can take months. By breaking the ties between applications and their underlying infrastructure — a key facet of cloud computing — the IT team can slash this process down to days or even hours. The more dynamic the environment, the greater the opportunity for delivering organizational value and enabling a rapid response to changing demands.

It's with good reason that the cloud computing concept has captured the IT industry's attention, becoming one of the most talked-about

technology opportunities of recent time. Because of its dynamic nature, cloud computing offers the promise of increasing IT agility while streamlining operations, improving efficiencies and, potentially, lowering costs.

For these reasons, cloud computing will likely prove one of the most strategic technologies for 2011 and beyond. Interest and use of public cloud services and cloud-related infrastructure purchases are increasing and will continue to do so for years to come.

While there is a great deal of optimistic discussion about cloud computing, it is not for everyone in every situation. IT executives will want to carefully study the cloud computing paradigm and explore what it offers their operations. The sooner the IT department understands the nature of the cloud – and outlines a strategy for its use, where appropriate – the better for the organization overall.

This reference guide will help organizations cut through the chatter so they can discern what the technology truly has to offer. It will introduce the basic concepts of cloud computing and define what makes up a cloud, explore the different service and deployment models, detail the benefits, and outline the foundational technologies.

What Is Cloud Computing?

Many IT managers can relate to being inundated with cloud service and product pitches from industry reps of every ilk, from long-standing application, platform and services developers to newly funded startups. Many have latched on to the "cloud" term when, in fact, their offerings don't meet the criteria for this computing model. For IT buyers, the question of what is and isn't "cloud" can get confusing very quickly.

For example, some sources equate the virtualized data center with cloud infrastructure. The two are not the same, although virtualization is considered a foundational technology for cloud computing. Nor is cloud computing the same as software as a service (SaaS), although a cloud computing platform can support this increasingly popular service model. Cloud computing can refer to externally provisioned services as much as internal IT infrastructure, depending on the technology that the organization uses.

So what is cloud computing?

A widely accepted definition of cloud computing stems from early work done by the National Institute of Standards and Technology (NIST), a U.S. Department of Commerce agency that promotes innovation and industrial competitiveness via measurement science, standards and technology.

Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources that can be rapidly provisioned and released with minimal management effort or service provider interaction.

In other words, cloud computing allows organizations to provide their staff with access to the applications, infrastructure or platforms they need to do their jobs – all via a simple front-end interface, such as a web browser. They might need access to these resources for a few minutes at a time or for a longer period. Depending on the deployment model in use, organizations can pay on a utility basis only for what they use.

If a cloud computing model seems more amorphous than is typical of IT initiatives, that's because it is. It is a further evolution of the strategic principles of virtualization, which rid the organization of "this application runs on this server" designations of legacy IT environments.

In a cloud infrastructure, computing resources are pooled and ready for the taking by any application as needed. These computing resources, typically virtualized, are infinitely scalable on the fly. They can scale upward to meet rising demand and then back down once this application need subsides, creating unprecedented levels of operational efficiency.

Any organization can benefit from this ability to handle unpredictable demand spikes quickly and efficiently. For example, a university could grab infrastructure as needed to handle peak loads during class registration periods, as could a K-12 district doing live streaming of a sporting event.

Likewise, a government agency might employ public cloud infrastructure to support heavy web traffic during crisis periods or when launching new citizen-awareness campaigns.

Computing capacity is not the only pooled resource within the cloud. Other resources include client, storage and network capacity, which enable organizations to acquire disk space and bandwidth as needed.

5 Essential Characteristics

While cloud computing introduces an unprecedented level of fluidity into the traditional rigidness of most IT infrastructures, it doesn't mean anything goes in this new environment. While there are many variations of cloud computing (covered in Chapter 2), five essential characteristics form the foundation for any cloud computing model. These defining characteristics, often intertwined, are as follows:

1. **Rapid elasticity:** Computing resources in a cloud should be capable of rapid provisioning, and at times, apply it automatically. Quick scale-in and scale-out capabilities can alleviate the overprovisioning and underutilization that characterizes traditional IT and service provider environments. Cloud consumers should have a seemingly unlimited pool of resources at their disposal to purchase in any quantity whenever the need arises.
2. **Self-service and on-demand access:** From a desktop and without intervention, IT staff should be able to enter a self-service catalog and acquire server time, network storage or other cloud resources as needed.
3. **Broad network access:** Cloud resources should be available over the network and accessible through standard mechanisms that promote use by almost any user client, be that a desktop computer, notebook, tablet, iPad, Cius or smartphone. Support should be heterogeneous.
4. **Dynamic resource pooling:** Within a typical cloud infrastructure, resources are pooled and available to multiple users at once in a practice known as multitenancy. The pooled resources get dynamically assigned and reassigned according to user demand. In general, cloud users have little to no control — or even knowledge — of where the resources reside. Besides pooled computing processing, a cloud may provide access to banks of storage, memory, network bandwidth and virtual machines.
5. **Measured service:** This computing model allows the monitoring, control and metering of resource usage within the cloud. While this allows for resource use optimization, it also provides transparency into that usage for the cloud provider and users alike.

BUSINESS INTELLIGENCE AND THE CLOUD

To a large degree, cloud computing is all about doing IT smarter and bringing competitive advantage to the organization. Business intelligence (BI) has a similar purpose.

What happens if you join the two?

By virtue of its dynamic and scalable nature, cloud computing can provide a great platform for BI initiatives underway throughout the public sector as well as educational institutions. Rather than devote massive amounts of server, database and storage capacity to periodic numbers crunching, IT can deliver BI as a service, on demand.

When the revenue or budget department wants to analyze its monthly figures — typical BI fare — it heads to the cloud's self-service portal for a BI service. Once the data is tallied, the resources for that service go back into the pool for reuse.

What's more, managing a cloud infrastructure itself requires strong BI capabilities. In fact, any cloud computing project should include a strong IT service management (ITSM) framework that eases the process of creating organizational processes, defining roles and establishing policies for the services delivered via the cloud, and improving services management.

As IT evolves its cloud infrastructure, it can tap into BI services to analyze consumption, performance and utilization trends, for example. With such information in hand, IT will have solid data upon which to make cloud decisions.

The Origin of the Cloud

If some or all of these characteristics sound familiar, they should. The cloud computing paradigm didn't spring up overnight; rather it has evolved over many years.

In the history of modern-era computing, the rise of the commercial Internet in general and the World Wide Web in particular enabled a fundamental shift toward cloud computing.

As web-centric computing took root within IT departments — think of intranets and interorganizational extranets that sprang to life in the late 1990s and early 2000s — a new reality dawned: Operations applications, now web-enabled and pieced together as components under an SOA model, need not be confined to physical resources.

A logical view of IT, enabled by the virtualization of computing, storage and network resources from underlying physical hardware, became not just possible but preferable. This development dovetailed with increasing mobility, which has more and more users demanding remote access to their critical applications whenever they want it, from whatever device they use, wherever they happen to be.

This new-found logical view popularized three principal computing concepts that can be seen in today's cloud architecture:

- Autonomic computing, whereby IT systems are self-managing and self-healing;
- On-demand and utility computing, which describes the provisioning of services on an as-needed basis, similar to how electricity is served;
- Grid computing, a method for pulling together unused computing resources across single or multiple organizations to give super-users ultraprocessing power.

As these concepts show, the IT industry has been inexorably moving toward the automated, dynamic, on-demand IT environment that is cloud computing.

What's different now? Simply put, the maturity level of the underlying technologies. The viability of a cloud architecture is a testament to the maturation of a critical set of technologies, including but not limited to automation, broadband network access, horizontal scaling, provisioning and virtualization.

CHECK OUT

NIST Special Publication 800-53 Revision 1 for guidance on cloud-applicable security controls. It can be found at:

csrc.nist.gov/publications/PubsSPs.html

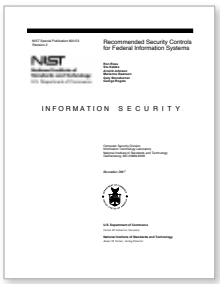

Where is Cloud Computing Going?

As this technology matures, cloud computing is starting to take many shapes. Some IT organizations are building next-generation cloud-based infrastructures for their own use, while others are relying exclusively on Internet-based cloud services to grab computing or storage resources for short periods of time. Still others are embracing a hybrid model that allows bursting from the private into the public cloud.

Likewise, cloud computing is manifesting itself in three service models: IT organizations can purchase software, platforms or infrastructure as cloud services.

For IT professionals, the challenge is to get educated, weigh the pros and cons of the myriad approaches, and determine what features work best for their organization.

The IT group will also need to develop processes to manage the private, public or hybrid cloud services it supports. Understanding the cost and effective evolutionary management of each type of service will prove invaluable to organizations moving to a cloud infrastructure.

A great resource to turn to is the Information Technology Infrastructure Library (ITIL) framework, which can guide organizations on documenting and implementing the key processes and disciplines needed to make a cloud initiative successful.

No doubt, cloud is rapidly evolving, and what it looks like today will differ from what it looks like tomorrow. But rest assured, the cloud will be to this decade and beyond what client/server distributed computing was to the past 20 years, and mainframe computing before that — the go-to template for computing operations. ■

VARIATIONS OF CLOUD COMPUTING

FINDING THE RIGHT OPTION FOR YOUR NEEDS

With the vast array of cloud offerings available on the Internet, testing out a cloud service has become increasingly easy. However, beyond the quick-hit, tactical needs of an individual user or workgroup, using cloud computing on an organizationwide scale is more involved.

One-off test runs with cloud services can serve a purpose: In order to fully comprehend the nature of the technology, getting some hands-on experience with it is important. But planning a cloud strategy that extends well into the future requires a more studied approach.

The goal should be a clear understanding of what the cloud can and cannot do for the organization. With so many opportunities now available, accomplishing this will help an organization clarify what it can expect from the cloud. Developing a roadmap for deployment can then follow.

A well-rounded cloud strategy requires careful consideration of each cloud service type, deployment model and access option weighed against the organization's particular application characteristics.

While cloud computing is quickly becoming a preferred architectural model for IT services, it's important to remember that some applications won't ever be suitable for a cloud environment. At the other extreme, some applications may not belong anywhere else. Where an organization's applications fall along that spectrum can be determined by their relative ubiquity or uniqueness.

If an application is easily replaceable by another of its kind or can run on standardized, commodity infrastructure, then consider it a good cloud candidate. Conversely, if an application requires a high degree of customization, long-term commitment and a specialized platform on which to run, then it isn't cloud-appropriate.

These deciding factors are applicable to any of the three cloud service models that follow: software as a service (SaaS), platform as a service (PaaS) and infrastructure as a service (IaaS).

MAKING THE CLOUD DECISION

Cloud computing is a transformational technology, not only for the IT department but for the entire organization. The IT team will likely want to consult and collaborate with these stakeholders:

- The executive team, including the most senior leader, who will want to be apprised of and sign off on strategic initiatives;
- The CFO and other financial team members, who will need to understand the potential cost reductions, revised cost structures and chargeback programs;
- Privacy, compliance and information security officers, who will need to know how information is handled within the cloud environment;
- Department heads, who need to understand the characteristics of cloud computing and the ramifications for their departments;
- IT managers and staff, who will acquire new roles and responsibilities.

3 Service Models

Software as a service: Often geared toward the end user who needs access through a web browser or other thin client interface, SaaS provides access to applications hosted on a service provider's cloud infrastructure.

An organization can find just about any general office application available via the SaaS model. Customer relationship management, calendaring, e-mail and human resources management are among some of the more common applications delivered as services from the cloud infrastructure.

For IT departments, IT service management, spam filtering, intrusion prevention and other traditional security software are among the application types increasingly available via the SaaS model.

With SaaS, the user organization neither owns the application nor the associated servers, operating systems, storage, network or other IT resources required for its support and delivery. The applications essentially come as-is, with little to no opportunity to tweak user preferences.

Because of this, the commitments between a customer and service provider can be fairly weak, and if one SaaS option doesn't work out, moving to another provider is an option for the organization.

SaaS evolved from application service provider and managed services technologies and is perhaps the best known and most mature type of cloud service available today. Small- and medium-size organizations, limited by staff and budget constraints, often find the most benefit in this model, although organizations of any size will find SaaS useful to some degree.

Note that one common misconception of cloud computing is that SaaS equals cloud — perhaps because of the innumerable SaaS offerings being pitched by cloud companies today. While SaaS can be delivered over a cloud infrastructure, it is not a cloud in and of itself.

Platform as a service: Derived from the SaaS model, PaaS caters to developers' needs. Rather than simply delivering prepackaged applications via the cloud infrastructure, as is the case with SaaS, a PaaS provider offers up the entire computing platform and solutions stack needed for an application.

With PaaS, organizations can deploy acquired or custom applications without incurring the associated upfront provisioning and ongoing maintenance and management costs of the underlying infrastructure. The user organization has application control, the caveat in many cases being that the developers must be comfortable with the PaaS provider's choices for programming languages, interfaces, development tools, database support and the like.

The types of PaaS offerings vary; which option works best for an organization depends on its goals. For example, some cloud companies offer a PaaS and SaaS combo, providing organizations with the ability to customize the packaged application. Although being able to tweak the application may make the SaaS model more attractive, the catch is that when combined with a PaaS offering, SaaS becomes less portable.

The richest PaaS offerings enable an organization to support the entire application development lifecycle.

This means that the PaaS provider not only provides the platform itself but also ensures source code and version control, enables user testing (with rollbacks as needed) and provides change-tracking functions. PaaS can also facilitate collaboration among far-flung developers.

Infrastructure as a service: This service model enables user organizations to forgo deployment of new data center equipment to handle growing operational needs. Rather, an organization obtains needed IT infrastructure from a cloud services provider, often via a self-service catalog.

While a user organization can run applications, databases, operating systems and other software on top of its selected infrastructure, it has no direct control over or access to those machines. The cloud service provider manages the infrastructure, including any scaling up or down as needed.

Infrastructure as a service is similar in concept to a traditional dedicated hosting service, with two major differences: Organizations tap into a shared, highly scalable pool of resources, and they pay for only what's used on a utility basis. In other words, organizations neither have to preorder nor pay for dedicated gear sitting in an outsourced data center.

IT organizations must keep in mind, however, that multitenancy applies across the public cloud infrastructure. They have no way of controlling the types of virtual resources running atop the infrastructure they've provisioned from the cloud.

4 Deployment Models

While matching operational requirements to cloud service type is fairly straightforward, a greater challenge arises when determining which deployment method or methods to use. IT organizations have four choices to choose from.

Public cloud: A public cloud is what most typically comes to mind when people hear the phrase *cloud computing*. With this option, the service provider offers its cloud infrastructure for general use on a self-service, on-demand basis. The target can be individual consumers or small- to medium-size organizations (or larger), depending on the need. Service types vary, as discussed earlier.

Using a public cloud has significant appeal because it requires little to no infrastructure investments while enabling unprecedented levels of scalability. The translation for IT shops: greater efficiencies and increased agility at a relatively low cost. The downside: Sharing resources across a public infrastructure may raise security and regulatory concerns that may nix this option for some organizations. Others worry about locking data into a single vendor's cloud infrastructure.

While public cloud computing allows an organization to avoid many infrastructure expenses, it's important to plan for other associated cost areas tied to a deployment, such as vendor management processes, capacity planning, chargeback systems, incident management and service level agreements (SLAs).

Private cloud: Taking the self-service, on-demand concept in-house, an IT organization can build and maintain a private cloud within its own data center or centers, much as it would any other type of infrastructure. This cloud would be for the exclusive use of the organization's staff or other privileged users.

As an alternative, a private cloud also may run externally at a hosted cloud provider's site. In this case, the provider maintains and manages the cloud infrastructure, which comprises pooled resources dedicated to that single customer's use.

Organizations that are concerned with compliance, privacy, security and data availability are more likely to build a private cloud than use public services. This is also the case for organizations that intend to incorporate legacy infrastructure as part of a cloud environment.

NEW TO CLOUD COMPUTING?

Maybe not. Anyone using an e-mail service such as Gmail, Hotmail or Yahoo! Mail is already operating in the cloud for these e-mail services.

As is the case with a public cloud, a private cloud does have drawbacks, such as capital expenses and infrastructure limitations. While there is greater scalability available within a private cloud than in a legacy environment, it isn't limitless, as it would be in a public arena.

Organizations may want to consider building a private cloud as a step toward an eventual public cloud deployment. This strategy positions the IT department to develop the disciplines and processes needed to move any service to the cloud. An IT services management (ITSM) framework will then be in place when the decision is made to switch to a public cloud.

Community cloud: A community cloud provides an opportunity for multiple organizations with similar needs or like interests to share infrastructure. The community cloud option doesn't provide the full cost benefits of a public cloud, but it can allow organizations to more readily facilitate requirements for higher levels of privacy, security and compliance. Similar to a private cloud, the community version can reside within an organization's data center or at an external site.

Hybrid cloud: Deciding on the right cloud deployment isn't always black and white. Because of this, some organizations opt for a hybrid cloud, mixing and matching among the private, public and community options.

In the most common scenario, an organization would break out of a private cloud by allowing applications to burst into the public cloud to grab additional resources as needed. In an ideal situation, the distinct clouds link via a standard interface that allows for data and application portability between them. Less ideal, but more common, are setups where proprietary technologies bind the different clouds together.

Deciding which cloud deployment method is best depends on a wide array of factors, including cost, control, performance, scalability, security and SLAs, among others. For example, a public cloud's scale-out infrastructure can be a huge boon to some organizations, while others that must comply with stringent privacy and regulatory guidelines might gravitate toward an internal private cloud.

Client Access

Client interface flexibility is one of the great features of cloud computing. As part of a cloud implementation, IT organizations can abstract client resources, such as the operating system, applications, associated data and the client interface, from the underlying physical machine. These resources are pooled within the data center for delivery as a virtual package to the user's device of choice, whether that's a desktop PC, notebook or smartphone.

In essence, being able to deliver a client interface (including any customizations) as part of a virtual container makes the user's experience portable. The look and feel of the interface persists even as the user closes down a desktop and accesses the same application from a smartphone. Not having to change and adjust to a new interface with every device boosts user flexibility and productivity alike.

For IT organizations, such front-end flexibility presents an opportunity to build out a desktop infrastructure that's more cost-efficient and easier to manage. One step that can be taken, for example, is swapping out traditional desktop PCs for thin clients because most processing occurs in the cloud.

As IT centralizes this function, care must be given to the process of application change management. Improper application change processes can potentially affect every user that's reliant on the cloud's services. ■

CLOUD COMPUTING'S VALUE PROPOSITION

BENEFITS ARE FELT AT MULTIPLE LEVELS: END USER, I.T. DEPARTMENT AND ENTERPRISE

Cloud computing, with its promise of highly scalable, dynamic and virtualized resources, helps address some of the biggest challenges that IT organizations face today. It will likely prove a transformational technology, ushering in positive operational changes and untold benefits for many organizations.

For example, cloud computing will ease burdens inside data centers, where data volumes continue to grow exponentially as new application types burden infrastructure and gobble up available bandwidth. Data center consolidation and virtualization help relieve some of the pressure. But absent a firm commitment to cloud computing, many IT teams still remain slowed by rigid silo-based structures within their organization.

Each service request sets off a series of actions to acquire the necessary resources — server, storage, database and network resources — often from different IT groups. It's a slow and cumbersome process that hinders agility, flexibility and innovation.

Many organizations' IT infrastructures remain overly complex. Applications correspond to specific hardware, software to specific operating systems. Should a failure occur somewhere, determining the root cause can be challenging and time-consuming, leading to unacceptable performance degradations, or worse, service interruptions.

Going forward, IT operations need to be far more agile, flexible and capable of responding in real time to an organization's changing needs; today's highly interconnected, global and online-oriented environment demands as much. Cloud computing offers a great opportunity to address many of the problems that IT groups face.

Benefits for the IT Department

The expectations for cloud computing are high, with potential gains for both IT operations and the organization as a whole. What follows is a rundown of cloud computing benefits that

WHAT TO ASK PUBLIC CLOUD SERVICE PROVIDERS

For many organizations, making the move to cloud computing requires a greater leap of faith for public cloud services than for a private cloud infrastructure. Agencies and departments will have questions, and getting the answers from service providers can help assuage concerns. Here are some key questions to cover.

1. Does the service provider help with application integration? How so?
2. How much, if any, customization is allowed to the applications delivered via the cloud?
3. Where will the organization's data reside? (If not by specific data center, a public cloud service provider should at least be able to provide geographic information.)
4. Who will do the migration in and out of the cloud?
5. What kind of dashboard access is offered and can it be customized by user profile?
6. What sort of service level agreements (SLAs) are provided?
7. What are the disaster recovery processes and how will replication of data be done?
8. What kind of help desk services are offered?
9. What if the organization decides that the cloud model doesn't work? Does the provider offer traditional hosting services? How would the transition from one model to the other work?
10. Does the provider have customer references to share?

will improve how the IT department obtains and provisions resources to the organization.

A solution for every need and budget: Cloud computing is available in many shapes, sizes and pricing levels, from the strictly private cloud requiring capital investment in supporting infrastructure to the completely public, pay-as-you-go cloud.

Pricing for public cloud services varies depending on the type of offering – IaaS, PaaS and SaaS (covered in Chapter 2). Organizations can find basic computing and storage resources as needed for reasonable pay-for-use fees.

Increased flexibility: Committing to one style of cloud computing does not preclude the use of another, allowing IT teams newfound flexibility in how they provision resources and deliver services. They can build an internal private cloud, set up a hybrid cloud for bursting into the public arena as needed, use public cloud services where appropriate and even support community clouds for certain user groups.

Better resource utilization: In a traditional IT environment, servers are notoriously underutilized and storage arrays often have inordinate amounts of excess capacity. IT shops know the resource-allocation rules of thumb, and they are rightfully leery of right-sizing their resources for fear of being put in a situation where demand outstrips supply.

Dynamic scalability is an integral feature of the cloud, so the IT team no longer needs to stockpile resources for future needs. However, IT still needs to manage overall capacity within a private cloud. To ensure resources, capacity planning is a must to stay ahead of demand.

Improved efficiency, greater agility: Cloud computing takes the restrictive binding of application to infrastructure out of the IT resource equation, allowing dynamic resource allocation on an as-needed basis. This allows the IT group to operate far more efficiently and with much greater agility than previously possible.

For example, an IT organization working with a public cloud need not build up the needed infrastructure for testing a new application. It can simply tap into an IaaS offering for the resources it requires. Improving efficiency and reducing testing time enables an organization to move new applications and services into production in days – maybe even hours in some circumstances – compared to months.

Infinite scalability: IT departments no longer need to commission servers, storage arrays and the like to handle new or increased application demand. A public cloud infrastructure

gives them an always-available elastic pool of resources from which to draw when needed.

Access to new technology: Public cloud service providers have teams — even entire departments — at work on application upgrades, infrastructure enhancements and protection mechanisms. Their goal is the continuous improvement of the services they deliver. Organizations using those services reap the benefits of these technology advancements without investing the time or budget into their development.

Improved security: Organizations that go either the public or private cloud route can capitalize on up-to-date security architecture. Public cloud users can embed security definitions and policies into their chosen cloud services delivery and management model.

Reduced costs/cost avoidance: IT organizations can realize cost benefits in a number of ways as their infrastructures evolve to embrace the cloud. For example, many organizations have probably already realized reduced capital expenditures from their data center consolidation and virtualization projects. In turn, these cloud-enabling initiatives bring cost relief in the form of lower power and cooling bills.

With a private cloud, organizations can gain savings by leasing cloud infrastructure gear rather than buying it. Whereas buying the servers is a capital expenditure (and will likely lead to utilization of the equipment past its prime to squeeze out total cost of ownership), leasing is an annual expense and the servers can be replaced with fresh technology on a regular schedule.

Being able to replace the equipment in a timely manner with servers that have greater computing capacity ultimately leads to a better TCO via the leasing option.

Organizations that acquire infrastructure from the public cloud also will gain budget benefits as they convert fixed infrastructure costs into variable ones. Breaking the infrastructure into components and automating delivery of those chunks of services can further reduce IT costs.

New cost model: One of the basic tenets of cloud computing is that services are delivered on demand on a pay-as-you-go basis. For organizations that so choose, the ability to measure service usage can serve as the impetus for a new cost model. Whether utilizing a public or private cloud structure, the IT department can monitor, control and meter the use of cloud services, and bill user groups appropriately (if needed).

In other words, as the IT group grows into its role as a service provider, it can morph from a cost center to a revenue unit with profit-and-loss responsibility. End users, in turn, get transparency into their usage and can make smarter budgetary decisions based on that.

Greater visibility for optimized IT services delivery: Cloud computing, with its dynamic service enablement, calls for a strong ITSM component. As the IT organization develops its ITSM capabilities, gaining real- or near-real-time insight into services delivery, it can tap into dynamic cloud resources to boost performance as needed.

Benefits for End Users

While cloud computing offers great benefits to IT department operations, benefits that radiate outward into the larger organization, it also provides improvements directly to end users and their day-to-day operations. What follows is a rundown of those benefits.

Ability to support new application types: As users become more Web 2.0-savvy, IT organizations increasingly must find ways to support rich web applications, multimedia digital content and large-scale data processing — sometimes with little or no warning. Cloud computing, with its elasticity, high levels of scalability and self-service access to resources on demand, gives IT staff the ability to easily meet these changing user requirements.

TRICKLE DOWN BENEFITS

Cloud computing facilitates data center consolidation, freeing up valuable office space and reducing power and cooling expenses.

ENHANCED MOBILITY

Enhanced mobility: The cloud computing model also supports standardized network access to these resource chunks, meaning that the IT department can also meet another rising trend: mobility. Through the self-service portal, users can retrieve cloud services from any desktop computer, notebook, tablet or smartphone. No matter the interface, the experience should be seamless.

Improved collaboration: With the availability of collaboration software as a service, delivered either via a private or public cloud, an organization's users need not go through time-consuming application requests. In an instant, team members can tap into enterprise-class collaborative applications, including calendaring, e-mail, file sharing, instant messaging, social networking and web conferencing.

Strategic focus: By simplifying and automating IT operations and service delivery via the cloud computing model, organizations can alter the typical 70-30 spending ratio, wherein 70 percent of the IT budget is spent on routine maintenance and 30 percent is spent on IT innovation and strategic projects that directly support an organization's mission. With the cloud, spending on innovation can actually rise.

Once fully migrated to the cloud, IT organizations find themselves spending far less time in the data center tending to servers and storage arrays, for example, and far more time listening to users, collaborating with each other and concentrating on strategic initiatives that bring about positive operational and organizational improvements.

Organizations building out private clouds have an opportunity to examine and reset their service management policies. Striving to improve processes and automation when rolling out their cloud offers a way to reduce that 70 percent maintenance spending ratio and put that budget to more strategic use.

Drawbacks to Keep in Mind

The list of cloud computing benefits is long and varied. But no technology, no matter how transformative, comes without a few cautionary notes. Whether or not they represent barriers to entry depends on the organization and its policies. Regardless, IT executives must conduct research, do their due diligence and give each of these potential drawbacks consideration before committing to a cloud initiative.

Additional layers of security: Organizations that procure public cloud services may need to tighten up existing security and/or add additional layers to match the service provider's security measures. This will ensure that assets are protected equally whether inside the public cloud or their own domain.

For example, as IT organizations allow access to cloud-resident applications, they must contemplate how to handle user authentication and identity management. As they no longer will have control over the backend infrastructure, they'll need strong assurances that unauthorized users inside or outside the cloud provider cannot gain access to their data.

In cases where multiple cloud services are in use, IT security professionals might consider establishing a unified access management scheme that allows users single-on access to multiple cloud applications. This will ease the IT security management burden as well as password requirements for end users.

In addition, in cases where sensitive data might be placed in the cloud, encryption of data — while in transit and at rest — should be the norm.

Firewall/proxy issues: Organizations using a hybrid cloud model need to consider how the security technologies being used within their internal clouds will match up to those of the public cloud service provider, as well as how data will flow through the firewall-based perimeter to the external cloud.

In this regard, IT organizations will need the cloud provider of choice to standardize on the same cloud-specific security technologies, such as virtual firewalls, being deployed within the internal private cloud. In some cases, IT organizations may deploy proxy servers that would intercept sensitive data for local delivery rather than via the cloud.

Regulatory compliance: A public cloud provider might offer some insight into what's going on within its network, but it's not going to give an organization complete transparency. Many organizations will need to be able to stipulate where their data resides geographically for reasons of regulatory compliance. They'll also need to verify that only authorized users have access to the data.

Questions worth asking include: Where does the data reside? Who has access to the data — and how is that proven for auditing purposes? What data protection mechanisms and disaster recovery strategies are in place? Will auditors be able to review the provider's overall security practices?

Application integration: Whether using a public or private cloud, IT executives must determine which legacy applications are and are not appropriate for delivery via a cloud model. Once legacy applications eligible for cloud computing are identified, an organization should not assume that the transition will be simple. Many legacy applications, such as those pulling information out of multiple databases, are too rigid to take advantage of the elastic nature of cloud computing.

To reap the full benefits of the cloud's elastic computing model, some applications will need modifications while others may need to be completely re-architected. Failure to think about and plan for the need to adapt applications for use in the cloud could negate the cloud's benefits for those workloads.

Building an application/data architecture is important to ensure alignment with current and future application delivery models. Organizations need to fully understand what the current architecture is and where it is headed in order to plan out a cloud strategy.

Reliance on an Internet connection: The greater the use of cloud services, the more imperative it is that an organization provide high-speed Internet access capable of supporting increased traffic. In addition, when operation applications run in the cloud, users need Internet connectivity from wherever they happen to be. Even though the speed and stability of those connections will be outside the IT department's control, users who have trouble connecting might still lay the blame on the IT organization.

Bandwidth issues: In certain situations, data sets may prove too large for the bandwidth available on some network segments. IT organizations that are considering placing applications with large data sets in a cloud may discover performance degradation issues. There are ways around this, such as moving user clients into the cloud.

Single-provider dependency: No IT organization likes vendor lock-in. Dealing with public cloud service providers is no different. As discussed in Chapter 2, the applications that an organization opts to run in the public cloud should be easy to recreate. Portability will matter if service problems arise and an organization decides to procure cloud service elsewhere. Service providers should be willing to commit to contractual stipulations regarding data availability. ■

PREPARING FOR THE CLOUD

INITIATIVES THAT SUPPORT A TRANSITION TO CLOUD COMPUTING

While there is a great deal of buzz around cloud computing's prospects, many government agencies and educational institutions are still becoming acquainted with this technology. With limited and often constrained IT budgets, organizations' strategic thinking regarding cloud computing is often waylaid by the tactical decisions needed to keep IT resources running day to day.

Arriving at a position to implement a cloud computing initiative is an evolutionary process, one that will require preparatory steps. As laid out, these common steps will benefit the organization, regardless of whether cloud computing is the long-term goal or not. But as IT projects, such as virtualization, are pursued, organizations will want to consider in their planning how these ventures may soon support a cloud computing framework.

Virtualization as a First Step

As mentioned in Chapter 1, virtualization is a foundational technology for cloud computing. Virtualization's cloud value is found in how it abstracts and aggregates data center resources, turning them into logical pools shared among users. The more virtualized the infrastructure, the higher the resource utilization within the shared pool. This allows the workloads to move more fluidly across the data center, and strongly positions the organization for migration to a cloud computing model.

Adopting virtualization in critical IT areas, namely servers, storage, clients and applications, will lay the groundwork for future cloud computing initiatives.

Server virtualization: By most industry estimates, some 30 to 40 percent of server infrastructures are already virtualized, on average. Experts predict that percentage will grow as organizations come to terms with management and security concerns about virtualization.

Virtualization has become a successful data center technology for two primary reasons. First, it enables large-scale consolidation of physical servers. A 20-1 virtual server to physical server ratio is a common rule of thumb for virtualization initiatives, but higher and lower ratios are possible, depending on numerous variables.

Those organizations that have virtualized 75 percent or more of their server infrastructure can look for other virtualization opportunities. They may want to focus on optimizing the virtual server infrastructure while looking at that other 25 percent of the enterprise where virtualization might still be implemented.

The second, but no less significant, reason for server virtualization's popularity is that it can slash IT spending on capital server expenditures and tamp down ongoing operational costs.

That virtualization has become such a critical technology for IT organizations bodes well for cloud computing. Although virtualization isn't the only technology for creating the resource pools characteristic of cloud computing, it promises to be the most commonly used.

Any organization that hasn't looked into server virtualization should do so — not only as a first stepping stone to the

cloud, but also for the host of benefits it brings.

Storage virtualization: Within a virtualized data center, storage is consolidated and its capacity allocated dynamically, just as server resources are. That is done through thin provisioning technology, which allocates disk storage space in a flexible manner among multiple users based on their minimum requirements at any given time, creating a pool of storage from which applications can draw from as needed. Fewer dedicated disks mean better capacity management and increased utilization, among other benefits.

Virtualized client computing: With virtualized client computing, the user's desktop — including operating system, applications and associated data — is separated from the physical machine. Instead, it all runs as a virtualized desktop on a central server.

Users can access their virtualized desktops from any number of devices, including a desktop PC, notebook, smartphone or thin client. Hosting clients centrally smoothes out provisioning, patching and policy enforcement while easing management.

Application virtualization: Regardless of how many applications an organization has, easing deployment and migration processes lays valuable groundwork for a self-service, dynamic cloud computing environment. Application virtualization tools change physical applications into virtual services that run in isolation from each other and the underlying operating system.

With application virtualization, IT managers need not install and support applications on each user's desktop (the same as virtualized client computing). Instead, they can manage the virtual instances from a central console. Furthermore, isolating applications as virtual instances means no two applications can conflict with each other.

IT Governance

Many organizations have studied how to better align IT services with user needs, knowing that doing so can increase agility and improve operational efficiencies. The imperative to adopt strong IT service management and governance becomes even more important under the cloud model, in which services are provisioned dynamically and not tied to specific pieces of the infrastructure.

Organizations that have not embraced IT services management will want to do so as a preparatory step toward the cloud. Consider implementing the ITIL framework for identifying, planning, delivering and supporting IT services to the organization.

Understanding service delivery, for example, will help organizations as they move toward dynamically provisioned services. Implementing a change management database, a core ITIL component, will help ensure that virtual machines are created and deleted in standard fashion and that those changes are documented in a central repository.

Getting started with an IT service catalog component will help organizations determine what set of IT services best meets users' needs and should be provisioned from the cloud's self-service portal, for example.

Cultural Preparation and Acceptance

For many IT professionals, cloud computing and its self-service model can be a daunting prospect. Common concerns include loss of administrative control over service delivery, increased operational workloads and loss of "ownership" of IT resources as computing, storage, database and network resources all get wrapped up and delivered as a service rather than managed separately.

Likewise, many departmental managers may resist the idea of having their workloads drawing from a shared pool of resources. Some may not be happy about paying for IT services they've received "for free" all along, even if a pay-for-use model benefits the organization as a whole.

No doubt, the cloud services delivery model does require a mindset shift, and IT executives need to do some preparation to help their organizations adapt to the new world. ■

UNDERSTANDING VIRTUAL SECURITY

As organizations layer in more and more virtualized resources and move closer to the cloud ideal, they must begin to think about security specifically as it applies to this new environment.

The hypervisor: Maintaining this layer's integrity means establishing basic configuration and vulnerability management strategies. Hardware and software tools are available to help secure the hypervisor.

The network blind spot: As workloads move from one virtual machine (VM) to another in the same physical server, traditional network-based firewalls and intrusion prevention systems (IPSs) can't see this traffic. Placing security controls within the virtual server can provide visibility in case of an inter-VM attack.

Trust zones: One way to address the inter-VM threat is to use virtual security software to create trusted network segments so that VMs of similar trust levels share a common host. Within a trust zone, the IT group is able to enforce policies, as well as monitor, filter and control VM-to-VM traffic.

CHAPTER 5

THE PRIVATE CLOUD

AN INTERNAL RESOURCE THAT OFFERS NUMEROUS BENEFITS

On-demand IT resources, infinite scalability, pay-for-use pricing structures, reduced capital expenditures — the benefits of public cloud computing are hard to ignore. But a "public" approach to cloud computing is not compatible with every organization's situation.

As discussed in Chapter 2, organizations that aren't in a position to pursue public cloud services need not be stuck with legacy infrastructure and outdated cost models. Rather, they can apply the same principles and make similar technology choices within their organization that a service provider would for a public cloud infrastructure.

They can, in effect, build internal private clouds that provide on-demand resources, pay-as-you-go pricing and unprecedented levels of scalability.

In fact, building an internal private cloud might not be such a leap for many organizations — at least on a technology level. That's because their data centers likely have been evolving toward a next-generation ideal for some time already.

Consider that much of today's IT operates on commodity x86 server hardware, and many government agencies and educational institutions have standardized their use of operating systems and software platforms. Many organizations have pursued data center consolidation, with virtualization being the chosen technology for reducing server sprawl (and increasingly applied to data and storage networks as well). Also, organizations are increasingly aligning IT operations with services for improved provisioning. These are among the foundational elements of an internal private cloud.

However, building an internal private cloud requires moving further away from old processes and adopting capabilities such as self-service access to on-demand resources and pay-as-you-go metering. Just because an organization's data center is consolidated and highly virtualized doesn't mean it is operating an internal private cloud.

Is a Private Cloud Right for You?

Many industry watchers have noted that as much as organizations think they're ready for an internal private cloud, few really are. Architecturally, they may be moving in the right direction. But are they ready to add in the capabilities that will enable a shift from a virtual infrastructure into a cloud architecture? For many, the answer is not yet.

Organizations that think they may be ready for an internal private cloud should ask the following questions:

Is the organization prepared to give users the autonomy they'll expect? For the application developer crowd, one of the draws of the public cloud is the ability to start up web applications in minutes. This group of users expects no less in an internal private cloud.

This means readying a self-service portal that will provide quick and easy access to all necessary controls. It also means configuration options for rapid application development and deployment. If developers know they'll need four virtual machines, storage volume and bandwidth, they should be able to grab those resources, sized appropriately and on the fly.

The same ease of use should be available to other end users, who should be able to provision SharePoint or file shares from the self-service portal, too. In this latter case, especially, what's happening in the back end should be completely transparent.

Has the organization standardized its procedures well enough? Because an internal private cloud is so dynamic, it demands standardized operating, deployment and maintenance features that ensure efficiency and consistency. Within an internal private cloud, a workflow engine can ensure that the request/approval process follows standards and procedures.

Organizations that follow ITIL guidelines for IT service management are more likely to be able to answer this question in the affirmative than most others.

How far is the organization willing to take automation? While many organizations have started weeding out manual processes, they're often a long way from fully embracing automation. High levels of automation are important in a private cloud for a number of reasons. For example, the more seamlessly and smoothly workloads can move about the environment, the more efficient and cost-effective the internal private cloud becomes.

Will end users willingly share resources? End users have come to embrace the idea that their data sits on dedicated servers and storage. In an internal private cloud, multitenancy applies — which means accounting and legal applications, for example, run in the same virtual pool as facilities and human resources. Some end users may be uncomfortable with the idea of resource sharing and have concerns that should be addressed during the transition to the cloud.

Does the organization want to begin charging IT usage fees? Although metered usage is part of the formal cloud definition, failure to charge for that use isn't necessarily an internal private cloud deal-breaker. But chargeback does warrant serious consideration.

The cloud's pay-as-you-go nature means organizations can bill or at least track and report on the cost of utilizing IT services. If an organization wants to initiate chargeback, then appropriate metering and tracking software will be part of an internal private cloud's deployment requirements. This process also serves to increase awareness, within separate departments, of the costs associated with delivering IT services.

What Belongs in the Cloud?

Once an organization determines its readiness for an internal private cloud, it must decide exactly what it wants to run in that environment. Not everything belongs there.

As an organization evaluates which of its workloads fit the model, it should look first for applications that have static interfaces and are easy to use. The best candidates shouldn't require massive scale-out. In general, they should run on standardized platforms and commodity hardware. First and foremost, these criteria focus on cost-effectiveness. The more consistency within the internal private cloud, the more cost-effective it will be.

In addition, consider applications that have similar SLA requirements. This, again, speaks to consistency within the internal cloud environment. Supporting a large range of SLAs requires a variegated infrastructure — and this heterogeneity, in turn, drives up deployment and management costs.

Conversely, applications that require high degrees of customization and are continuously targeted for upgrades and improvements aren't suitable for deployment in an internal private cloud. The continuous rate of change to the interfaces would prove too taxing on the dynamically provisioned, self-service model. Moreover, this is often a mission-critical application type, one that supports core operational processes and that is best run on dedicated resources.

Designing Your Cloud

Looking up and down their IT stacks, many organizations will find that they already have the building blocks in place to help them erect an internal private cloud. Some pieces will be more mature than others. Therefore, the focus of design and development efforts will depend on where an IT organization stands on each of the following elements.

Consolidated infrastructure: The more streamlined the IT operation, the easier it is to manage and optimize cloud service delivery and application performance. Consolidate servers, storage and network bandwidth. Respective considerations here include the use of blade server–filled chassis, deployment of storage area

networks (SANs) and migration to 10 Gigabit Ethernet (10 Gig-E) network links.

Dynamic resource pooling: In most internal private clouds, virtualization will be the foundational technology, given its maturity and pervasiveness across many IT operations today. Virtualization is a cloud-enabling technology because it abstracts and aggregates data center resources, turning them into logical pools shared among users. The more virtualized the infrastructure, the higher the resource utilization with the shared pool will be and the more fluidly workloads can move across the data center.

For example, in a highly virtualized data center with shared resource pools, a workload could easily move from virtual machines (VMs) to virtualized storage should the need arise.

Organizations also must be aware that however popular virtualization is, it isn't the only option for creating dynamic resource pools. In some cases, for example, IT professionals might prefer to build their cloud pools using products that enable rapid reprovisioning.

Likewise, if an IT organization already uses or plans to use high-performance computing clusters, it might consider siphoning off excess capacity and creating dynamic resource pools that way.

Resource management: While some organizations may rely on more manual processes today for orchestrating resource assignments based on service requests, the design goal here should be to automate as much as possible. This includes mapping virtual to physical resources. Per service request, resource managers should gather and deploy operating system and application images along with storage and network resources.

Self-service interface: Cloud users should be able to access services from a self-service portal in a manner that meshes with their organizational roles. An end user should be able to pick SharePoint from an IT services catalog without having to also request the back-end resources required for supporting that service.

Meanwhile, application developers could use characteristics such as high performance or high availability to select the appropriate infrastructure for their applications. Ideally, the self-service interface would remain consistent no matter what changes might take place on the back end.

IT service management: The widely used ITIL framework is a good starting point for best practices on creating processes and service policies; building the services catalog; applying capacity, configuration, demand and performance management; monitoring service health; and implementing metering, chargeback and reporting.

In addition, an internal private cloud requires a service governor that would dynamically optimize available resources against service requests based on a range of factors. These could include SLAs, operational policies, scheduled or forecast service demands and the like. In the absence of a service governance tool, IT organizations would want to handle this orchestration manually.

Metered service usage: Strict adherence to the cloud computing definition would have an IT organization present service pricing to users and charge departments for usage. Even if an organization isn't culturally ready to begin billing for IT services, metering service usage and delivering cost estimates could positively impact how end users think about how they might use resources more efficiently and cost effectively.

DON'T OVERLOOK SECURITY

While some IT organizations may be preoccupied with how best to incorporate automation and self-service provisioning, security planning mustn't be left by the wayside. As IT organizations build out their private clouds, they will need to consider the following:

- How will the internal private cloud integrate with the enterprise authentication and authorization scheme?
- How will monitoring and alerting within the cloud be integrated into enterprise systems and be made available for compliance purposes?
- How will the cloud encrypt data while in transit and at rest?
- How will the organization plan for incident response? Forensics becomes tricky the more fluidly resources move around the cloud.

Migrating to the Cloud

Some organizations will have reached high levels of maturity within one or more of these areas, but not across all of them. How an organization approaches its internal private cloud design will depend on the current state of its IT operations.

In some instances — for example, a greenfield opportunity that doesn't involve legacy infrastructure — a quick build-out is possible. Several manufacturers offer "cloud in a box" solutions. These products, which tightly couple software with hardware, typically include a self-service portal, a cost-allocation engine and automated resource management.

But more commonly, IT organizations move to private clouds with legacy infrastructure in place. In such situations a gradual migration is prudent.

Organizations can utilize existing infrastructure, continuously adding virtualization and methodically layering in defining capabilities such as dynamic resource pooling, automated resource management, a self-service interface and usage-based billing as time and resources permit. Gradual migration allows migraters to improve their agility, boost efficiencies and increase operational value each step along the way.

Building a private cloud infrastructure — even slowly — is a considerable challenge that no organization should take lightly. Issues to consider when plotting out a migration include the following:

Legacy applications: Not all applications make sense in a cloud computing model. IT must carefully evaluate and identify which are suitable for the cloud and avoid struggling to try to force-fit those that aren't appropriate.

Legacy infrastructure: Server updates will happen as part of the virtualization process, so organizations will likely have newer hardware migrating into their private cloud infrastructure. Trying to take advantage of older, less flexible hardware that's difficult to manage is a mistake in the dynamic cloud infrastructure.

Scalability: If an organization is looking for hyperscalability, then a private cloud infrastructure might not be the best fit. Private clouds are far more scalable than traditional IT infrastructure, but not as much as a cloud service offered from a public network.

Integration between private and public: An organization should consider how it might use public cloud services in conjunction with its private cloud should this need arise. Will a workload run equally well in either type of environment? Can internal security controls be extended into the public cloud?

Management: IT executives must consider whether their existing application performance, systems monitoring and network management tools are adequate for monitoring and managing the private cloud infrastructure, or whether they'll need more specialized tools.

Budget: As with any major IT project, organizations must carefully examine both capital and operational costs associated with building and managing a private cloud infrastructure, as well as how they'll show a return on investment.

Culture: As mentioned in Chapter 4, many IT staff may be uncomfortable with the self-service, automated nature of a private cloud. So IT executives will need to educate them on the value of embracing the cloud model.

Managing the Cloud

Prior to migrating to a private cloud infrastructure, organizations must determine how they're going to manage them. This is largely determined by the management tools that the organization selects.

A good starting point is to cultivate a holistic, end-to-end view of the IT environment, including the private cloud infrastructure. Cloud management tools present a single view for monitoring and assessing performance of physical and virtual machines as well as multitiered applications and services. These can span traditional physical and virtual environments, or even reach into the public cloud.

In addition, IT organizations that have instituted or are planning to use chargeback mechanisms as part of their private cloud strategies might look for tools that provide real-time usage metering. The more automated this capability, the easier it is to implement.

GOING TO THE LIBRARY

ITIL offers excellent guidance for cloud management and can serve as a go-to resource for planning out a cloud initiative.

Besides understanding management requirements and picking the most appropriate tools for those needs, IT organizations can ease management by simplifying and optimizing their self-service catalogs. A service catalog should provide services well-suited to end-user needs. It should also be built upon interchangeable resources for maximum flexibility.

This requires the IT team to develop a strong understanding of how users will consume the services provided. What's more, the IT organization should simplify the components from which it creates its services — for example, utilizing a standard operating system build as widely as possible.

Organizations with a successful cloud management practice will constantly seek improvement. This means continuous process assessment as well as resource consumption and usage trending analysis.

This approach blends well with one of the primary benefits of a private cloud infrastructure: the ability to adapt quickly to changing requirements. An informed awareness of how the cloud operates coupled with a good understanding of end-user needs will position the cloud as an invaluable resource for the organization. ■

THE HOSTING MANAGED SERVICES ROUTE

CONTRACTING FOR CLOUD SERVICES

Acquiring IT services from a hosting managed services (HMS) provider has been the chosen path to the cloud for many organizations, regardless of size.

In the HMS model, a service provider delivers applications and technology to customers from its own data center infrastructure. HMS customers avoid capital expenditures while lightening their application development and support burdens. HMS providers cater to small- and medium-size organizations and large enterprises alike.

The HMS model is similar to public cloud computing in many ways, but it lacks the cloud's defining characteristics: rapid elasticity, self-service and on-demand access, dynamic resource pooling, broad network access, and metered usage.

Hosted cloud services are available on demand from self-service portals and on a pay-as-you-go basis. The cloud hosting provider dynamically pulls together the resources needed to deliver customer workloads from a shared pool, typically built on a highly virtualized foundation.

Compare this to traditional hosting, in which customers determine how much capacity they'll need up front and pay for that on a contractual basis. In the meantime, inside the data center, the hosting provider dedicates physical infrastructure on a customer-by-customer basis.

It's easy to see how cloud hosting is a natural extension of the traditional managed services model. Cloud services are available from specialty providers that have popped up in the past year or two as well as from Web 2.0 companies and legacy hosting companies.

What Services Are Available?

Organizations can find a full complement of IT services available via the HMS cloud model. Provisioned from within the HMS provider's cloud, the services typically are delivered over the Internet, although dedicated WAN links may be used in certain instances. There are a variety of services available.

Infrastructure: This service offering entails real-time, automated provisioning and scaling of virtual and physical infrastructure such as servers, storage, networks, load balancing and security.

By running web applications in a hosted cloud, for example, an IT organization benefits from the scale-out nature of that infrastructure. Rather than sizing dedicated servers and storage for peak usage — and paying for those unused resources — the organization simply taps into the hosted infrastructure and increases capacity on demand, but only for as long as needed.

In other words, the infrastructure delivered as a service is great for supporting dynamic IT workloads — something that neither traditional internal IT nor hosting managed services do well.

Servers: Created, configured and decommissioned on the fly, servers deployed under this service model typically have operating systems and memory allocation specified by users.

Being able to provision servers from a hosted cloud gives the IT group the flexibility to deliver computing capacity on a project basis, for short periods of time and much more quickly than

ordering a server the old way. Powering up servers on demand works well in both staging and production environments.

Storage: Files or data backups are uploaded and stored on a cloud HMS provider's arrays. Storage capacity is scalable up and down on demand.

A great deal of Web 2.0 data gets stored in the cloud by default, but cloud storage's usefulness goes far beyond that. For example, accommodating high I/O operations per second (IOPS) storage requirements coming from rich media content or the unpredictable growth of an organization's digital archives is another area where cloud storage is useful.

Applications: Common applications are hosted by cloud providers and delivered in the software-as-a-service model, including those for collaboration, customer relationship management (CRM), e-mail and web hosting, among others. While frequently used, these applications aren't necessarily mission-critical, and so are good choices for an off-premises delivery model.

Application development/deployment: In a platform-as-a-service delivery model, the cloud host provides the entire computing platform and solutions stack needed for an application during testing, development and, if desired, deployment.

The HMS Public/Private Cloud Choice

When cloud HMS providers first offered their services, they did so with a very specific customer base in mind — application developers and independent software vendors. But as cloud computing has matured and garnered interest in broader circles, providers have expanded their reach to both individuals and organizations seeking on-demand resources.

This provides more options for an organization as it weighs the public-versus-private cloud decision. Some HMS cloud providers allow organizations to run private clouds on their public infrastructure, giving organizations yet another option to choose from. HMS cloud options include the following:

HMS public cloud: Cloud service providers make virtualized cloud infrastructure available for general consumption on a self-service, on-demand basis. The cloud providers manage multiple tenants from a shared pool of resources, offering high scalability but limited configuration, security and performance options.

HMS private cloud: In this alternative, an organization can run its private cloud on an HMS provider's cloud infrastructure rather than within its own data center. While the provider, rather than the organization's IT group, maintains and manages the infrastructure, this differs from traditional hosting in that self-service provisioning, on-demand access, multitenancy and all other cloud tenets apply.

Most important, it differs from the public cloud in that the HMS cloud provider cordons off the private cloud resources so they're not accessible to other customers. However, the HMS private cloud customer must size its cloud and commit to that set of resources on a long-term basis. After that, HMS cloud providers typically use a pay-as-you-go model to grow and shrink resources on demand.

Comparing Options

Public or private, hosted or not? These are the multifaceted questions being debated within IT organizations. If an organization is particularly concerned with compliance, privacy, security and data availability, it might wish to forgo public cloud services on an institutional basis, or restrict cloud use to applications and nonessential workloads to which those concerns don't apply.

The HMS private cloud offers a convenient option for organizations that don't want to incur the expense of building an internal private cloud but are wary about the openness of a public cloud. HMS private clouds allow customers greater control over their environments than with hosted public cloud services. For example, organizations can make security adjustments, specify infrastructure requirements and fine-tune SLAs.

When compared with the internal private option, the HMS private cloud offers the benefits of infrastructure cost avoidance and places no management burden on internal IT staff. However, when application, infrastructure and security control is nonnegotiable, the internal route is less risky than other private cloud options.

HMS: KEY CONSIDERATIONS

There are many options to choose from for cloud infrastructure hosting: traditional service providers, hosting companies and outsourcers, Web 2.0 enterprises and boutique shops.

Be sure to get a thorough explanation of the business plan of any potential provider. In addition, check out the experience of the management team and the depth of expertise throughout the ranks.

Also, determine whether the cloud hosting provider maintains its own cloud infrastructure end-to-end or collaborates with other outsourcers itself. Many cloud hosting companies supplement their infrastructure. In that case, check out the partners as well.

Choosing an HMS Provider

Whether an IT organization decides to use hosted cloud services or entrust a private cloud to a hosted infrastructure, it needs to query potential providers on a wide range of topics. Here's a checklist of issues that need to be addressed.

Security: It almost goes without saying, but an HMS cloud provider's security practices should weigh heavily in the decision-making process. The following security protocols needed to be a part of any service agreement:

- Assurances, backed by technology implementation, that one tenant can't gain access — either intentionally or by mistake — to another tenant's data on a shared server
- Encryption of data in transit and at rest
- Firewalls at the network perimeter as well as on host servers
- Use of authentication and secure passwords
- Regular reviews and security updates

Servers: Organizations need to know what types of servers are operating in the hosted infrastructure, and why. The provider's replacement procedures for failed or otherwise problematic machines also need to be known.

In addition, organizations will want to ask how the provider handles server redundancy and where its servers are located. An HMS provider won't likely provide specifics, but it should be able to share general geographical and environmental conditions with customers.

Storage: Similar to the server question, an organization needs details on what types of storage the HMS cloud provider uses and why. Information on how quickly storage can be added or subtracted, and at what cost, is also good to know.

Backup and recovery: Any hindrance to accessing data in the cloud is not acceptable. Organizations should dig into hosting candidates' backup procedures, getting details on frequency, location and mean time to recovery.

Monitoring: Organizations should expect continuous monitoring with automated alerts, real-time dashboard visibility into provided services and access to performance statistics, as well as trends analysis.

Service interface: Will organizations have access to the cloud via a web front end or another sort of client interface? Should the provider's back end change, will that be transparent from the user interface perspective?

Cost: Knowing whether an HMS's cloud is the best financial choice means carefully comparing the provider's pricing against internal costs. Give this analysis the attention it deserves.

Support: Look for HMS cloud providers that offer 24x7 support from well-rounded, experienced staff. Depending on application requirements, an organization may require advanced support as well. Will the cloud hosting company provide help in porting data and applications to its cloud?

Negotiating an SLA

Commonly used to measure how a service provider meets quality-of-service designations, SLAs are difficult to define when dealing with a cloud service delivery model. In fact, they remain immature and are often confusing.

Of course, the HMS provider must be accountable for application performance in its cloud infrastructure, but that application also travels across the public Internet on its way to the user. If the end user is at a desk, it travels across the organization's own network too. How does all of this affect a cloud SLA?

If an IT organization is negotiating an SLA with a cloud HMS provider, it should have the agreement reviewed by its legal department. At the very least, organizations should be sure to cover the following when drafting an SLA:

- Speed of initial resource allocation, as well as manual or automated responsiveness to growing and shrinking usage demands
- Granularity of the SLA, as in, does it apply to the infrastructure as a whole or on a machine-by-machine basis?
- Clear accountability for downtime due to scheduled maintenance versus emergency maintenance; avoiding "best effort" language in the SLA
- An exit clause allowing termination of the contract without penalty in the case of recurring incidents
- Definitive rules for what happens when service problems result in refunds instead of service credits, and what the redemption procedures are in each case
- Mandated delivery of monthly reports that analyze performance against agreed-upon metrics, such as server response time, server uptime, storage availability and support staff responsiveness
- Insight into the cloud for monitoring and compliance checks

Migrating to an HMS Cloud Provider

Once cloud choices have been finalized and the SLA is approved, all that's left is moving the organization's data to the provider's infrastructure.

Before the actual migration, IT organizations must test the scalability of that infrastructure as well as its on-demand responsiveness. Promises don't always meet reality. It's better to know this before a migration than after it.

Depending on an IT organization's capabilities and the nature of the cloud services it will be using, the migration process may require some assistance from the HMS provider — or it may require none whatsoever.

IT organizations can expect software-as-a-service deployments to be fairly routine, with applications quickly up and ready for use. Porting data and on-premises applications to a cloud infrastructure is more difficult. As part of a migration plan, an IT organization may need to call on the HMS cloud provider to help optimize the application for use on its infrastructure or to do some code tweaks, for example. ■

This glossary serves as a quick reference to some of the essential terms touched on in this guide. Please note that acronyms are commonly used in the IT field and that variations exist.

GLOSSARY

Application virtualization

This type of client virtualization allows applications to run as virtual services in isolation from each other and the underlying operating system.

Broad network access

Broad network access is an essential cloud characteristic in which capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (such as smartphones, notebooks and PDAs).

Business intelligence (BI)

BI is a broad category of applications and technologies for gathering, storing, analyzing and providing access to data to help users make better operational decisions. BI applications include the activities of decision support systems, querying and reporting, online analytical processing, statistical analysis, forecasting, and data mining.

Cloud computing

Cloud computing enables convenient, on-demand network access to a shared pool of configurable computing resources (networks, servers, storage, applications and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.

Cloud hosting managed service (HMS)

A cloud HMS service provider delivers applications and technology to consumers from its own data center infrastructure.

Cloud storage

In a cloud storage arrangement, files or data backups are uploaded and stored on a cloud HMS provider's arrays. Storage capacity is scalable up and down on demand.

Community cloud

A community cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (similar mission, security requirements, policy or compliance considerations). It may be managed by the organizations or by a third party and may exist on or off premises.

Dynamic resource pooling

With dynamic resource pooling, the provider's computing resources are pooled to serve multiple customers using a multitenant model, with different physical and virtual resources (such as storage, processing or memory) dynamically assigned and reassigned according to organization requirements.

Hybrid cloud

A hybrid cloud infrastructure is composed of two or more clouds (private, community or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (such as cloud bursting for load balancing between clouds).

Information Technology Infrastructure Library (ITIL)

ITIL is a globally recognized collection of best practices for IT service management.

Infrastructure as a Service (IaaS)

IaaS provides customers with the ability to provision processing, storage, networks and other fundamental computing resources where the organization is able to deploy and run arbitrary software, which can include operating systems and applications. The customer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage and deployed applications, and possibly limited control of select networking components (such as host firewalls).

IT service management (ITSM)

ITSM is a discipline for managing IT systems that is philosophically centered on the organization's perspective of IT's contribution to the entity.

Measured service

Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (storage, processing, bandwidth or active user accounts). Resource usage can be monitored, controlled and reported, providing transparency for both the provider and consumer of the utilized service.

National Institute of Standards and Technology (NIST)

NIST is a U.S. Department of Commerce agency that, among other stated responsibilities, promotes effective and secure use of cloud computing within government and industry.

Network virtualization

This form of virtualization is a method for combining the available resources in a network by splitting up the available bandwidth into channels, each of which is independent of the others, and each of which can be assigned (or reassigned) to a particular server or device in real time.

On-demand self-service

This feature is an essential cloud characteristic that allows a customer to unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction by the service provider.

Platform as a Service (PaaS)

PaaS gives the customer the ability to deploy onto the cloud infrastructure applications created using programming languages and tools supported by the provider. The customer does not manage or control the underlying cloud infrastructure, including network, servers, operating systems or storage, but has control over the deployed applications and possibly application hosting environment configurations.

Private cloud

A private cloud infrastructure is operated solely for an organization. It may be managed by the organization or a third party and may exist on or off premises.

Public cloud

A public cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services.

Rapid elasticity

This feature allows capabilities to be rapidly and elastically provisioned, in some cases automatically, to quickly scale out and scale in as needed. To the cloud service consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time.

Server virtualization

This form of virtualization facilitates the masking of server resources, including the number and identity of individual physical servers, processors and operating systems, from server users.

Service level agreement (SLA)

SLAs are commonly used to measure how a service provider meets quality-of-service designations.

Software as a Service (SaaS)

SaaS provides organizations the ability to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through a thin client interface such as a web browser.

Storage virtualization

This form of virtualization allows the pooling of physical storage from multiple network storage devices into what appears to be a single storage device that is managed from a central console. Storage virtualization is commonly used in a storage area network.

Virtualized desktop computing

With this form of client virtualization, the user's desktop – operating system, applications and associated data – is separated from the physical machine. Instead, it runs as a virtualized desktop on a central server. Users can access their virtualized desktops from any number of devices, including a desktop PC, notebook computer, smartphone or thin client.

INDEX

ABOUT THE CONTRIBUTORS

NATHAN COUTINHO is a Solutions Manager for CDW with a focus on virtualization. He has more than 11 years of experience in IT, covering various roles in management, technical sales and consulting. His current responsibilities include evaluating and educating clients about trends and directions in the server, client and storage virtualization spaces.

PAUL SCHAAPMAN is a Solutions Architect for CDW. With more than three decades of experience in IT infrastructure, he has a strong background in virtualization (server and client), server and storage engineering, IT architecture, and IT consulting. Paul was awarded VMware's Virtual Vanguard Award in 2007 for his work on a large virtual infrastructure for the Virginia Farm Bureau.

LOOK INSIDE FOR MORE INFORMATION ON:

- The varieties of cloud deployment options and their benefits
- Determining which cloud option best suits your needs
- Planning a migration to the cloud
- Managing a cloud environment