White Paper

Creating and Implementing an Enterprise Cloud Strategy

David Linthicum
Blue Mountain Labs

Blue Mountain Labs
A BICK GROUP COMPANY

Introduction

Cloud computing is about the ability to share IT resources more efficiently. Thus, the core objective of IT, and specifically the CIO, should be to understand the value that this technology can bring to their enterprise or government agency.

The purpose of this paper is to provide a framework of understanding around the use of cloud computing technology, and suggest a process to create and implement your own enterprise cloud strategy for your organization. We will focus on the creation of a cloud computing strategy that defines a sound cloud infrastructure management approach, and selection of the right enabling technology.

Additionally, we'll focus on the vision of cloud computing, or, how to understand the systemic value of this shift in IT, and how to drive more efficiency into your IT infrastructure over time. We'll do this by walking through the process of creating a cloud computing strategic roadmap for your organization, or a plan that will ensure additional strategic IT efficiencies, using cloud computing as a strategic enabling technology.

Let's first understand the real value of cloud computing.

Core to the value of cloud computing is the ability to operate IT with fewer dollars. While we've looked for solutions to this dilemma for years, the reality is that most enterprises have increased IT operations spending without a proportional increase in value.

Cloud computing has the potential to reduce costs while increasing value by allowing the more efficient use of IT resources. Thus fewer staff, hardware, and software resources are required to support the enterprise. The common scenario here is to reduce the internal hardware and software footprint, and at the same time increase capacity and reliability. Later in the paper we'll explain how this occurs.

Another benefit is the CIO's ability to spend more time on new initiatives. Cloud computing promotes innovation, considering that there is little latency between the origin of an idea and deployment of the IT required to put that idea into practice. Today CIOs are more concerned about keeping the existing IT infrastructure in an operational state, and not about working on new projects that can bring more value to the business.

Considering that cloud computing has the ability to do more with less, the core benefit of cloud computing is the ability to increase capacity while at the same time reducing the need to purchase capital equipment (servers, networking equipment, etc.). This allows the enterprise to push that cash back into the core business.

Cloud computing can preserve capital is by upsizing your IT infrastructure on demand, at the user level. In other words, the allocation of IT infrastructure is dynamic in nature, and aligns directly to the needs of the business. Thus, if additional compute or storage resources are required, they are allocated to those users and applications that need them, as they need them, without delays around the procurement of hardware and software, and without having to wait for the system manager to allocate virtual machines, which can take weeks, sometimes months.

This capability is present within the context of cloud computing. Each application is no longer bound to a single physical server. Through the mechanisms of virtualization

and multitenancy, hardware and software users are managed within a limited number of physical servers, and new resources can be dynamically allocated or de-allocated as required.

Like upsizing on-demand, you need to consider what it will take and dollars paid to reduce computing capacity. What does it take to scale down in case you no longer need the computing resources and want to reduce costs as well? Such is the case within many ecommerce systems with capacity requirements that are seasonal. Both upsizing and downsizing on-demand are known as elasticity, or the ability to scale up or down as required by the business but without a significant impact on capital costs.

Another core value of cloud computing is the ability to shift the risk from your enterprise to the cloud computing solution. With internal clouds, this means the core risks around the success of the IT solution are significantly reduced because you avoid the hardware and software costs. With public cloud providers, you shift the capital risks to the provider who is in a much better position to manage those risks. The same core value is present with private clouds as well.

Agility means the ability to change the IT infrastructure faster to adapt to the changing needs of the business, such as market downturns, or the introduction of a key product to capture a changing market. This, of course, provides a strategic advantage and allows the business to have a better chance of long-term survival. These days, many enterprises are plagued by IT infrastructures that are so poorly planned and fragile that they hurt the business by not providing the required degree of agility.

Agility is core to the value of cloud computing. While there are indeed operational and capital cost savings, the ability to quickly adjust our IT infrastructure to meet the changing demands of the business typically provides the greatest benefits over time. However, many times these benefits are not measured. Thus, in order to make a complete business case for cloud computing, CIOs should consider the value of agility, as we will do in this paper.

Cloud Computing Challenges Today

The core challenge that IT and CIOs face today is the fact that only a small percentage of enterprise server capacity is leveraged. While 10% is widely tossed about as an average number, in many enterprises there are hundreds, sometimes thousands of servers that run at 3% to 4% of capacity. This is due to the one-application-one-server concept that has ruled the world of IT for the last 20 years. Excess capacity was never utilized, thus hardware and software assets were not shared.

This lack of utilization is the result of our silo approach to solutions over the last 20 years. A business problem led to a solution within a single application, and a single set of servers to support that application. The result is many servers in the data centers that are underutilized and bound to a single purpose. Thus, you may have 500 or more servers, each running between 5 to 10% of capacity.

Moreover, there is a need for more visibility into heterogeneous virtual environments. As IT struggled to solve the problems around utilization using virtualization technology, the number and types of virtualization solutions have made management of these environments complex and expensive. For instance, enterprises need to monitor oversubscription of available virtual resources they support, or send alerts to administrators that they need to move up to more resources when they approach their limits. In addition, enterprises want notifications when virtual machines are no longer in use and can be decommissioned or re-deployed elsewhere.

The challenge is that management of heterogeneous virtual environments requires you to abstract the underlying complexity from the cloud infrastructure managers. Using this approach, it is much easier to see and manage the underlying virtualized resources via a single user interface and API infrastructure. This means you can allocate and de-allocated virtualized resources as required by the enterprise, as well as load balance and monitor cloud efficiencies.

There is also a need to delegate provisioning without IT losing control. A traditional virtualized environment requires administrator intervention when allocating and de-allocating virtual machines. Today, users would like greater flexibility to configure their environments, such as allocating virtual machines for storage, compute, etc. These types of capabilities lead to the elastic benefits of cloud computing, which is the ability to allocate or de-allocate resources as needed within a multitenant environment, providing efficient access to a single pool of resources.

Enterprises need better capacity planning and cost savings from virtualization to realize the true value of cloud computing. For instance, there is no clear approach or enabling technology to dynamically allocate virtual resources, or what's known in the world of cloud computing as auto provisioning. Thus, the default is to deploy new virtual machines on dedicated servers which results in much lower utilization and a much higher cost model. In addition, humans who make allocation decisions are highly inefficient, error prone, and not dynamic. There needs to be an abstraction and management layer between those who require the resources, and the resources themselves, to assure that we do not give up efficiency for delegation.

There are also compliance issues with cloud computing virtualization that should be addressed, such as logging and reporting (SOX, HIPPA, etc.). In many instances, compliance requires network and storage separation, meaning there are requirements around both the physical and virtual environments. For example, a virtual machine belonging to a research analyst must not be able to sniff the wire and

discover network traffic from an investment banker. In many instances, specific groups are also prohibited from sharing storage hardware, and sometimes compute resources. In other instances there are requirements around where data can be stored, such as restrictions on certain types of data that cannot be transmitted outside of a country, or that must be encrypted in specific ways.

Compliance leads quickly to data security. The core concept here is the ability to insure that sensitive data, both at rest and in flight, has the proper encryption and other security mechanisms protecting it. Also, that the data is not transmitted to places where it won't be protected or compliant, or where it could be vulnerable in other ways. The core notion is to identify the security required for each data set, but also provide management mechanisms to insure that the data security approaches and enabling technology are right for the data under management. There is no one size fits all solution here, thus data security requirements should be considered on a domain-by-domain basis.

Cloud computing technology is often very proprietary in nature, and thus does not provide good portability and heterogeneous management capabilities. Avoiding proprietary solutions should be a core objective of cloud computing. Locking into a particular vendor or provider means that any value gained with private or public cloud computing is quickly lost due to the cost of portability or even interoperability.

Avoiding provider and vendor lock-in means that you must look at open standards as well as how each provider approaches portability. Most technology providers who leverage open systems technology can explain use cases for moving from one system to another, and paths to both portability and interoperability. Again, you need to consider what you're looking to accomplish using this technology in the larger context of the existing technological offerings.

While we've already covered the concept of elasticity earlier in this paper, it's a core consideration when moving to cloud computing. This is the ability to quickly scale up or scale down to the capacity required by the business, to support changing needs. For instance, the ability for an enterprise to scale up to support the added resources required to launch a product without driving waves of hardware and software procurements, or the ability to scale down as the spike in processing requirements passes, thus freeing up those resources for other purposes.

While enterprises like the idea of elasticity, compliance and security issues make cloud computing a tough sell. Thus, we must look at solutions that provide the elasticity required, but also support the security and compliance requirements.

In the next section, we'll suggest some solutions to this problem.

Case Study:

In order to bring the points made in this paper to a more practical vision, let's look at a few case studies around the use of cloud computing within the enterprise.

Global Telecommunications and Cellular Provider

This global organization, like many telecommunications and data communications services providers, is migrating its internal and external service offerings to a managed cloud model. They are doing this to realize the benefits of increased hardware utilization, faster deployment speed, and distributed administration authority. They intend to gain the following values from their cloud computing solution: Manageability, security, portability, and durability. However, while the promise of "the cloud" is the reason they will move to cloud computing, they still need to select the right enabling technology to make the cloud work for them.

The global telecommunications and cellular provider elected to deploy the Abiquo solution to expose managed cloud services to its diverse business units. The idea is that they are able to provide cloud services today, internally, and then they will leverage this deployment as the template for their emerging external managed cloud offering. Starting with Abiquo's Community Edition, they successfully deployed the solution in their US data centers, and they are in the process of migrating to the Enterprise Edition for expanded network and persistent image storage management.

Enterprises need cloud computing management to be successful, including the ability to manage the allocation and de-allocation of virtual machines within the enterprise cloud, as well as the ongoing management of the enterprise cloud to insure that the SLAs are being met and the users are getting access to the resources they require. Without this type of management, the value of their cloud would be significantly diminished considering the amount of technical and administrative effort required to manage the cloud during normal business operations.

Creating your Vision for Cloud Computing

When considering the technical value of cloud computing, we should considering the following core benefits:

- Scalable
- Manageable
- Securable
- Portable
- Durable

Scalable refers to the fact that our cloud computing solution needs to scale up to support the business. While we've talked about the elastic attributes of cloud computing solutions, the concept of scalable means that the cloud computing solution will scale up to support the business processing requirements without a significant increase in hardware and software expenses.

Manageable means that we need to be able to manage our cloud computing solution, including adding, deleting, and managing compute resources across a number of different types of systems without having to move through a great deal of complexity. This is typically accomplished using abstraction layers that can hide the complexities around the underlying virtualization or other system management environments from complex interfaces and APIs. This is an emerging best practice in managing cloud computing environments.

Securable means we can leverage state-of-the-art security best practices, including encryption and identity management, to insure that all data residing within the cloud computing environment meets all security and compliance requirements.

Portable refers to the fact that moving an application to a cloud computing environment does not lock the application and data to that environment, and the application and data are easily portable to other platforms, cloud and non-cloud, to meet the needs of the business. This means your cloud computing solution leverages open standards that will insure investments made around cloud computing technology will be protected as the business needs change, and applications and data move from cloud-to-cloud or from clouds back to traditional systems.

Durable means the cloud computing solution is resilient to internal and external system issues, and provides uptime that meets or surpasses the uptime requirements of the business. This is a critical success factor for cloud computing. Your cloud computing solutions must exceed the up-time record of existing systems to gain trust and value from the existing user community.

Creating a Cloud Roadmap

So, how does one bring cloud computing into their enterprise? We would like to suggest a few basic steps that CIOs and other IT leaders can take to define their cloud computing roadmap. Keep in mind this is not just about solving a tactical problem this year, but moving toward a long-term strategic use of cloud computing that should bring long-term strategic value to the enterprise.

Here are the steps we suggest:

Step 1: Define the New Strategies

Based on the Gartner Group's Key Success Factors for Cloud Computing,[1] you need to address multi-tenancy, public-private cloud usage, roles and responsibilities, governance, globalization, workload policies, standards, integration, leveraging existing systems, and a virtual image repository.

For most in the traditional IT field many of these concepts are new. However, in order to be successful, you need to understand how these key capabilities work and play within your emerging cloud computing strategy.

Multi-tenancy refers to the ability to leverage technology that has many users who simultaneously access the same resources. ***Public or private*** cloud usage means looking at the processing location. Location options include private, within your enterprise, or public, outside of your enterprise.

You must also consider the new ***roles and responsibilities***, in that many things being done now won't be done in the future. Also, there will be many new operational tasks, and you'll need people trained and ready to take them on.

Governance is also an important concept in the world of cloud computing, considering how data, processes, systems, and services are to be managed and controlled. Many cloud systems are made up of many subsystems and services that are tightly coupled. There is a critical need to govern those assets that protect the holistic systems from unauthorized changes that could have a cascading effect on the cloud applications.

Globalization refers to the fact that emerging cloud systems have to exist in the context of many other countries, cultures, regulations, and standards. S***tandards*** include leveraging open environments, which is necessary to avoid vendor lock-in. Globalization also provides the ability to leverage the right ***workload policies*** so we effectively manage the allocation of processing across all virtual machines.

Integration is also a core requirement for building a cloud computing system. We will need to synchronize data and processes between cloud environments, as well as between the cloud and traditional systems. We need to use integration to ***leverage existing systems***, both at the information and behavior levels.

Finally, we need to leverage and manage a ***virtual image repository***, which allows images to be shared and managed effectively.

[1] 2009-2014 Cloud Services Report, Gartner Group

Step 2: Define New Software and Models

When considering cloud computing, there are new software requirements that should be understood when creating the roadmap. They include the addition of cloud application management and cloud infrastructure management, as well as understanding the differences between PaaS, SaaS, IaaS, and existing hypervisors.

Cloud infrastructure management should manage a shared pool of physical computing resources that is controlled and managed by the IT organization, and provides a single point of access/management. This includes the delegation and management of the cloud/virtualized enterprise to create the appropriate resource limits, and provide access to key facilities such as provisioning and governance. Another feature is the automatic workload distribution within the cloud according to policy, including serve-service provisioning and de-provisioning by cloud users, so that the exact needs are met and resources are shared efficiently. Finally, a cloud infrastructure management system should leverage open standards and provide interoperability with features such as automatic conversion between hypervisor types.

The benefits of cloud infrastructure management are clear, and include the significant reduction of workload within IT in support of the cloud, and the ability to support emerging cloud best practices such as auto-provisioning which allows users to leverage the cloud to their exact purposes and thus increase the agility of the infrastructure. Another benefit is the elimination of vendor lock-in, which allows the cloud manager to move processes between hypervisors and platforms, as well as improving utilization through policy-based workload management and reducing reallocation of unused resources. All of this has the side benefit of reducing power consumed, as well as data center space required.

These definitions of SaaS, PaaS, and IaaS come from The National Institute of Standards and Technology (NIST), and are best explained by their definitions published in 2009.

> *"Cloud Software as a Service (SaaS).* The capability provided to the consumer is 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 (e.g., web-based email). The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.
>
> *Cloud Platform as a Service (PaaS)*. The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages and tools supported by the provider. The consumer 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.
>
> *Cloud Infrastructure as a Service (IaaS).* The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over

operating systems, storage, deployed applications, and possibly limited control of select networking components (e.g., host firewalls)."[2]

Step 3: Define Business-Driven Policies

The ability to leverage logical business polices within technology is critical to the success of any cloud computing design and deployment project. This is the link for the business to IT, and thus defines the ability for the technology to carry out the requirements of the business.

In this step we define the business policies we need to consider and create, including the who, what, why, and where around the use of the virtualized resources. Examples of these policies include which groups can access which virtual machines, and what they can accomplish when they access those resources.

This is a very important step considering that the creation and execution of policies define the security and compliance framework for the cloud infrastructure you're creating. Therefore, we suggest that you take the time to define these polices at a logical level, and then define how they will be implemented within the cloud management technology.

Step 4: Create a Deployment Enterprise Strategy

This is the more laborious step in the process. You're taking a strategy and understanding, as defined by the previous steps, and building that out to an actual implementation/deployment strategy. We would suggest you leverage the following check list:

- Access the business.
- Access the culture.
- Access the value.
- Understand your data.
- Understand your services.
- Understand your processes.
- Understand the cloud resources.
- Identify candidate data.
- Identify candidate services.
- Identify candidate processes.
- Create a governance strategy.
- Create a security strategy.
- Create a cloud management strategy.
- Bind candidate services to data and processes.
- Relocate services, processes, and information.
- Implement cloud management.
- Implement security.
- Implement governance.
- Implement operations.

Note that not all of these steps are required; you pick and choose the steps that are most valuable to you. The idea is that we're able to understand the existing "as is" state of enterprise IT, as well as the value. From there we create a migration strategy

[2] *csrc.**nist**.gov/groups/SNS/**cloud-computing/cloud**-def-v15.doc*

for core services, processes, data, and applications, and select only those systems that are right for cloud computing.

Make sure not to neglect the management infrastructure, including operations and cloud management. Core to the success of cloud computing, including obtaining the value, is the ability to properly manage your cloud assets.

The Needs of Cloud Infrastructure Management

Diving a bit deeper into the cloud infrastructure management requirements, there are a few key needs:

- The ability to avoid vendor lock in.
- The ability to provide rapid provisioning.
- The ability to manage virtual sprawl.
- The ability to provide greater efficiency.
- The ability to manage security and compliance.

The ability to avoid vendor lock in refers to the openness and portability of the cloud infrastructure management solution. As already covered in this paper, it allows portability from system to system, hypervisor to hypervisor, without requiring a significant rewrite or other costly processes.

The ability to provide rapid provisioning refers to the cloud infrastructure management solution's ability to allocate virtual machines quickly and directly in line with the requirements of the business. In the past, it could take weeks or months to allocate new virtual machines, and users were typically not allowed to create or eliminate virtual machines for fear they would corrupt the resources. Rapid provisioning allows both the managers and the users to quickly allocate virtual machines as required by the business, allowing the cloud infrastructure to scale up and scale down, as required.

The ability to manage virtual sprawl is the ability to manage the many hypervisors and systems that are typically found within most data centers. Over the years we've built very heterogeneous and complex environments, and each type of virtualization has become a silo within the cloud and enterprise infrastructure. What's needed is a cloud infrastructure management system that's able to span the "virtual sprawl" and provide a common point of management that automatically accounts for differences in images and hypervisors.

The ability to provide greater efficiency is the ability to increase server utilization by providing better management with and between hypervisors. Using a cloud infrastructure management system, all available resources are known and thus can be automatically allocated to particular processes. Thus, server utilization is as efficient as possible.

The ability to manage security and compliance, as discussed earlier in this paper, is the ability for the cloud infrastructure management system and cloud managers to create processes and policies around the use and management of data, both at rest and in flight.

Leveraging Abiquo for Cloud Infrastructure Management

Considering the requirements we've just listed, it's interesting to see how an existing cloud infrastructure management product meets those requirements. Abiquo technology provides the core technology required to operationally manage cloud computing. Abiquo allows the enterprise IT staff to abstract themselves from the underlying complexities of the heterogeneous and complex virtualization systems.

Specifically:

Abiquo is the only cloud infrastructure management solution that supports all major hypervisors, and allows the cloud manager to convert any one hypervisor to any other hypervisor, thus eliminating vendor lock in. Leveraging an intuitive drag-and-drop interface, this system provides the cloud manager with the ability to maximize all available virtualized assets, without considering the vendor implementation.

To manage virtual sprawl, Abiquo separates the physical from the virtual, thus it can delegate provisioning to down to the users, if required. There is no need to wait on the cloud manager to allocate the resources needed by the users. This independence in allocating and de-allocating resources, as required, goes directly to the value of agility and the elasticity of the cloud infrastructure. Additionally, Abiquo can auto-discover all existing virtual machines, and monitor those instances in terms of utilization and other operational behaviors. This includes setting resource limits, and managing billing and chargeback systems.

Abiquo automatically deploys virtualized machines based upon provided business policies. Thus, the cloud manager can set a policy to maximize efficiency, or future tune the cloud infrastructure looking at the tradeoffs of efficiency with compute speed.

In support of security and compliance, additional business polices within the Abiquo technology can specify which security and compliance policies each enterprise group must comply with. This insures that all users and resources behave in accordance with set business policies, and thus compliance issues are automatically avoided.

Call to Action

Cloud computing is not a free lunch. There is much that must be done to be successful in the cloud. What should be done right now is a review of your existing enterprise IT assets and the potential for cloud computing to bring value. Moving forward you'll find the value that cloud computing brings is very specific to your business processes and your existing way of doing IT. That said, most enterprises will find that cloud computing can solve many existing efficiencies issues, and moving to cloud computing will provide clear and measurable value.

Those who leverage virtualization could need some additional planning and perhaps some critical enabling technology to be successful in the cloud. Typically these environments are poorly planned and highly heterogeneous. Thus, a cloud infrastructure management system is required to abstract the complexity from the cloud manager and make more efficient use of the underlying virtualization infrastructure, including the easy allocation and de-allocation of virtual machines, the management of business polices in support of compliance and security, and the ability to avoid technology and vendor lock-in. Most CIOs will find this technology critical to the success of cloud computing within their enterprise.

Abiquo focuses strongly on the needs of the enterprise. Its unique approach to resolving virtualization 1.0 issues make it a strong potential partner for enabling enterprise cloud strategies. Contact www.abiquo.com if you would like to explore this option.

About the Author

David (Dave) S. Linthicum is the founder and CTO of Blue Mountain Labs, an internationally recognized industry expert and thought leader, and the author and coauthor of 13 books on computing, including the best-selling "Enterprise Application Integration" (Addison Wesley). Dave's latest book is "Cloud Computing and SOA Convergence in Your Enterprise, a Step-by-Step Approach." Dave keynotes at many leading technology conferences on cloud computing, SOA, Web 2.0, and enterprise architecture, and has appeared on a number of TV and radio shows as a computing expert.

In his career, Dave has formed or enhanced many of the ideas behind modern distributed computing, including Enterprise Application Integration, B2B Application Integration, and SOA, approaches and technologies in wide use today. For the last 10 years, Dave has focused on the technology and strategies around cloud computing and how to make cloud computing work for the modern enterprise. This includes work with several cloud computing startups.

Dave's industry experience includes tenures as CTO and CEO of several successful software companies, and upper-level management positions in Fortune 100 companies. In addition, he was an associate professor of computer science for eight years and continues to lecture at major technical colleges and universities, including the University of Virginia, Arizona State University, and the University of Wisconsin.

Blue Mountain Labs
11654 Plaza America Drive
#103
Reston, VA 20190

www.bluemountainlabs.com