

Moving to the Cloud

A white paper produced by the Cloud Computing Use Cases Discussion Group

Final Draft

21 February 2011

Contributors: Salvatore D'Agostino, Joe Armstrong, Rizwan Ahmad, Niranjan Davalbhakta, Raji Gogulapati, Edmund Lau, Eugene Luster, Aurelio A. M. Matsui, Anish Mohammed, David Moskowitz, Mike Nolan, Tom Plunkett, Sakshi Porwal, Amrith Raj Radhakrishnan, Mark B. Sigler, Kamala Sreenivasan, Phillip Stratton, Robert Syputa, Krishna Venkatraman, Michael Versace

Public comments on this document are welcomed and encouraged via the discussion groups referenced at <http://cloudusecases.org>.



This work is licensed under a [Creative Commons Attribution Share Alike 3.0 Unported License](http://creativecommons.org/licenses/by-sa/3.0/).

1 Overview

Cloud computing will change the world of IT as dramatically as anything since the rise of the Web. Before deciding whether to move to the cloud, it is vital to understand the potential and the risks of cloud computing and the organization's requirements for using the cloud.

This paper presents a three-step process for evaluating cloud computing:

1. **Classify Your Information Assets:** Understand the function and value of the organization's applications and data and the risks to the organization if they are lost or compromised.
2. **Determine Your Requirements and Risks:** Define the requirements of the organization and determine if a cloud provider exists that is capable of delivering those requirements while keeping the risks at an acceptable level.
3. **Calculate Your Return on Investment (ROI):** Using the organization's needs, assets, risks and requirements, calculate the cost of moving to the cloud and compare that to your existing costs.

Before discussing the process, there are two important topics to cover: the value propositions of cloud computing and the non-technical considerations that can override any other concerns.

1.1 *The Value Propositions of Cloud Computing*

Before considering moving to the cloud, it is vital to look at the basic value propositions of cloud computing. The NIST definition of cloud computing¹ describes five essential characteristics:

- ◆ **Rapid Elasticity:** Elasticity is defined as the ability to scale resources both up and down as needed.
- ◆ **Measured Service:** Cloud services are controlled and monitored by the cloud provider, and the provider bills the consumer only for what they use.
- ◆ **On-Demand Self-Service:** A consumer can use cloud services as needed without any human interaction with the cloud provider.

¹ You can find the full document on the NIST Cloud Computing page at <http://csrc.nist.gov/groups/SNS/cloud-computing/>. The document states, "This material is public domain although attribution to NIST is requested. It may be freely duplicated and translated." The NIST material in this paper is based on Version 15 of the document, dated 8-19-09. These characteristics are discussed in more detail in the Cloud Computing Use Cases paper.

- ◆ **Ubiquitous Network Access:** The cloud provider's capabilities are available over the network.
- ◆ **Resource Pooling:** Resource pooling allows a cloud provider to share its physical and virtual resources according to consumer demand.

These characteristics provide two significant advantages:

- ◆ **Lower costs:** The ability to add virtual machines, storage and other resources dynamically means consumers can buy hardware based on their normal workloads instead of over-buying to allow for their heaviest workloads. The organization can do the same amount of work with fewer machines. That means lower costs for buying hardware and software, lower costs for keeping machines powered on and cooled, and lower staffing costs because fewer administrators are required.
- ◆ **More responsive organizations:** In many organizations, requisitioning a new machine, database or other resource can take days, weeks, or even months. With cloud computing, those resources can be acquired (and later released) as needed. Even better, that process can be automated so that no human involvement is required.

1.2 *Non-Technical Considerations*

Although most discussions of cloud computing focus on the technologies involved, there are two non-technical considerations that override any other concerns. These should be considered before any decision about cloud computing is made.

1.2.1 **Organizational Challenges**

Cloud computing is changing the world of IT forever. As with any technology trend, an organization adopting cloud computing will encounter changes to the mission, authority, funding and staffing of various departments within the organization. The staff of any department facing a decline in their influence on the organization will almost certainly resist those changes.

Most discussions of cloud computing focus on the technologies that enable it and the value propositions discussed earlier. But without executive-level support, moving to the cloud will not be successful. It is vital that any new cloud-related project be sponsored by a manager who is enlightened and powerful: someone who can both make the right technical decisions and make them happen. Without that support, any wide-ranging attempt to move to the cloud will fail.

1.2.2 Regulatory Concerns

Another non-technical issue is the presence of government and industry regulations. For a variety of reasons, governments around the world are concerned about the use of cloud computing. As an example, many countries have strict privacy laws that prohibit certain data from being stored on a physical machine located outside that country. An organization from such a country storing sensitive data in the cloud must be able to prove that their cloud provider followed those laws.

In addition to government agencies, many trade and industry groups create regulations as well. While those regulations might not be required by law, they represent best practices.

Following these laws and regulations will take precedence over all other requirements. A new law might require an organization to spend its resources changing an application's infrastructure instead of adding features to it. New laws and regulations will be created on an ongoing basis; the CIO's office must be alert for changes to the regulatory landscape as they emerge.

2 Classifying Your Information Assets

Before a thorough investigation of moving to the cloud, it is vital to determine exactly what information assets your organization has. Without an understanding of those assets and their value, assessing the requirements, risks, and benefits of moving to the cloud is impossible.

The information assets of an organization are often more valuable than its physical ones. Those assets can consist of intellectual property, trade secrets, research, financial data, and personal information, among other things. Some of that information is crucial to the success of an organization (or even its continued existence), while other information might be subject to copyright, privacy, or export control restrictions.

When deciding whether to move an asset to the cloud, a vital part of the risk assessment process is classifying that asset. There are three basic parts to this process:

1. **Identification:** The organization must identify the information, where it currently resides, and the policies and regulations in place for storing, accessing, and deleting that information.
2. **Classification:** The organization must classify the information according to its value and the potential damages if the information was lost or accessed inappropriately.

3. Protection: The organization must create a security chain for each class of information.

Once the organization's assets are identified and classified, the security chains should be defined and put into place.

A security chain must protect the organization's information assets at all levels, including physical security, technical security, and procedural and legal steps. Physical security includes measures such as restricting access to data centers, shredding paper documents and destroying tapes and hard drives. Technical security includes everything from the basics of firewalls and access control systems to more advanced techniques such as disabling USB ports. Finally, procedures for handling information assets must be clearly defined and adequately explained to all employees of the organization. In some cases, the procedures may include legal requirements such as laws covering the retention or destruction of data.

Once the organization has classified its information assets and defined the risks and requirements for using them, the decision to move to the cloud will be more straightforward. Moving extremely valuable information to the cloud, especially a non-private cloud, can pose risks that outweigh any benefits of using cloud computing.

In some cases the legal restrictions imposed on certain classes of information will make it impossible to move that information to a non-private cloud. Using a private cloud might still be an option, but a private cloud has risks as well. Moving information to a private cloud might increase the number of the organization's employees who have access to the machines that store and process it. The security chain must be modified to include everyone with access.

2.1 Candidate Applications for Moving to the Cloud

With the benefits of cloud computing in mind, there are several kinds of applications that are good candidates for moving to the cloud:

- ◆ **Pilot Projects:** A cloud pilot project is a good way for an organization to evaluate cloud services to see how useful, reliable and cost-efficient they can be. A pilot project should be a non-critical application that has a limited scope, a short time frame and loosely defined estimates of its ROI. Building a pilot project has the added benefit of giving the organization a chance to learn how to use cloud services.
- ◆ **Variable Workloads:** Some workloads have low requirements the majority of the time, with occasional periods of very high requirements. An organization must buy resources to handle the maximum workload, even though most of the time those resources will be idle. Moving that workload to the cloud can free the organization to buy only the resources to handle

its normal requirements. When the workload peaks, the organization can use cloud computing to provision the resources it needs, then release those resources when the workload returns to normal.

- ◆ **Non-Essential Tasks:** Certain applications and data are essential to an organization's core mission; they typically have much higher requirements for resources and much tighter restrictions on how and where they are used. If there are low-risk applications and data that could be processed off-site, moving them to the cloud would free resources for the rest of the organization.
- ◆ **Data Mining:** Data mining typically requires substantial hardware to find patterns in massive amounts of data. Done in-house, the organization must buy, maintain, power and cool all of that equipment. Moving that task to the cloud can provide substantial savings. The machines required can be virtual machines that run only when needed.
- ◆ **Development and Test:** Development and testing require substantial resources when done on in-house systems. Developers must have the same level of development tools on their machines. Testers must maintain many different machine environments and test applications on all of them. Moving development tools into the cloud ensures that all developers are using the same level of tools, and upgrading the organization to a new version of the tools requires one upgrade in one place. Doing testing in the cloud allows the organization to define a single set of virtual machines for testing; those virtual machines can be started when needed, then shut down when the tests are complete.

3 Determining Your Requirements and Risks

As with any migration, moving to the cloud carries with it some requirements and risks. In most cases, moving to the cloud does not introduce new risks, it merely changes the nature of the existing ones. In addition, the threat posed by each risk varies depending on the type of cloud. Security is always a concern, but security in a non-private cloud involves more variables than security in a private cloud.

Although this paper covers a broad set of common risks, requirements and scenarios, each of those can be affected by the type of cloud being used. For the purposes of this paper, risks and requirements will be discussed in terms of private clouds versus non-private clouds. All of the resources of a private cloud are inside an organization's firewall; all other types of clouds (the public, hybrid and community clouds of the NIST definitions) have at least part of their resources on a shared network.

3.1 Security

Security is consistently mentioned as the most important concern for organizations moving to the cloud. Although the cloud does not introduce any new security threats or issues, it does increase the number of people who have access to the organization's resources. The most significant difference when considering security from a cloud perspective is the organization's loss of control, not any particular technical challenge. With an in-house application, controlling access to sensitive data and applications is crucial. With a cloud-based application, access control is just as important, but the data, infrastructure, platform, or application is under the direct control of the cloud provider.

To adequately secure any system, a number of security controls are necessary. Some of the most common security controls include securing data, storage, networks and endpoints; defining identities and roles and the access control policies for them; and key and certificate management. The services offered by a cloud provider must support all of the security controls the organization needs.

For more information, security is covered in extensive detail in Sections 6 and 7 of the Cloud Computing Use Cases white paper.

3.2 Privacy

Privacy is a concern for any application that deals with sensitive data. Many types of data are subject to privacy laws, copyright protection or export restrictions. An organization's need for privacy often goes beyond the basic controls for data security. It is vital that a cloud provider deliver the added controls needed to protect sensitive data, including the ability for the organization (or government regulators, in extreme cases) to audit the cloud provider to prove that it followed the appropriate procedures.

3.3 Federated Identity / Single Sign-On

As an organization moves applications and data into the cloud, it is likely that the information a user needs will come from different sources, each of which has its own access control mechanisms. Federated identity and single sign-on use an authentication service to vouch that a user with a particular role should be allowed access to a given resource, even if the system controlling that resource has no knowledge of that user.

For more information, Section 6.3 of the Cloud Computing Use Cases paper covers federation, identity management and single sign-on.

3.4 Interoperability and Portability

The rapid provisioning and deprovisioning of cloud computing delivers a great deal of operational flexibility to an organization. That being said, if moving to the

cloud locks the organization to a particular cloud service provider, the organization will be at the mercy of the service level and pricing policies of that provider. With that in mind, portability and interoperability become crucial to providing the freedom to work with multiple cloud providers.

Interoperability is concerned with the ability of systems to communicate. In the world of cloud computing, this means the ability to write code that works with more than one cloud provider simultaneously, regardless of the differences between the providers.² On the other hand, portability is the ability to run components or systems written for one environment in another environment.

As organizations decide whether to move to the cloud, it is important that they consider interoperability and portability. The amount of freedom is likely determined by the type of service used. An application written to use specific services from a particular vendor's Platform as a Service (PaaS) will likely require substantial changes to use similar services from another vendor's PaaS. On the other hand, there are a number of open-source libraries that provide a single, consistent interface to common infrastructure services such as cloud storage. An application written to those interfaces is far more likely to be interoperable and portable.

3.5 Service Level Agreements (SLAs)

An SLA defines the interaction between a cloud service provider and a cloud service consumer. An SLA is the foundation of the consumer's trust in the provider. Among other things, an SLA contains:

- ◆ A set of services the provider will deliver, along with a complete, specific definition of each
- ◆ The responsibilities of the provider and the consumer
- ◆ A set of metrics to determine whether the provider is delivering the service as promised

Depending on the type of cloud service, a provider might need to be certified for certain standards (ISO 27001, for example). Many organizations will also need the ability to monitor and audit the provider to ensure that the terms of the SLA are being met. Finally, the cloud provider must provide transparency, notifying consumers of any outages or problems that occur.

Any organization considering or negotiating an SLA should know its business objectives before agreeing to any terms of service. It is vital that the organization know exactly what it needs as it considers different cloud providers.

²The definitions of interoperability and portability are based on the work at http://www.testingstandards.co.uk/interop_et_al.htm.

For more information, SLAs are covered in extensive detail in Section 8 of the Cloud Computing Use Cases white paper.

3.6 Availability

Availability is a clear requirement for any system, whether it is in the cloud or in the data center down the hall. One risk of cloud computing is that the people responsible for diagnosing a problem and getting the system back online do not work for the organization directly. It is vital that the SLA define the availability the cloud provider will deliver, as well as the recovery procedures in the event of any outages.

Business continuity and disaster recovery are also part of availability. An organization should understand what architecture and technology the cloud provider has in place to recover from system failures, including redundant systems and self-healing infrastructures.

3.7 Performance

Adequate performance is crucial to any successful move to the cloud. If moving to the cloud saves the organization money, yet the performance of applications slows to an unacceptable level, those savings are meaningless.

When moving an application to the cloud, it is important to define the performance the cloud provider must deliver. This is done with Service Level Objectives (SLOs). “Throughput for a request should be less than 3 seconds” and “At least five instances of a virtual machine should be available 99.99999% of the time” are examples of SLOs. The SLOs should be part of the SLA, they should be defined in terms of the organization's objectives, and they should make it clear exactly what performance the cloud provider will deliver.

3.8 Governance

Every organization has policies for deploying, managing, archiving and deleting its applications and data. When moving to the cloud, it is vital that the cloud provider support those policies. As mentioned previously, data is often subject to laws and regulations; the cloud provider's services must keep the organization in compliance, and the provider must be auditable to prove it has done so. The provider's responsibilities for enabling governance should be part of the SLA.

3.9 Testing

The rapid elasticity provided by cloud computing makes it relatively straightforward to test an application as it moves to the cloud. Stress-testing multiple instances of an application under massive loads can be done by starting the application on many virtual machines, then running the test. This is significantly easier and cheaper than building those machines and deploying

them on internal resources. When the test is complete, all of the virtual machines can be shut down.

Applications that use infrastructure services can be tested easily as well. For example, if an application that uses local storage is modified to use cloud storage instead, testing can validate that any operations with the cloud storage service work correctly. Testers should be aware that cloud services often perform much slower than local services. Writing to a disk in the cloud, for example, will take much longer than writing to a disk in the same machine.

Testers should also be aware that many cloud services are massively redundant, meaning that any changes made to a cloud service will be replicated to other machines across the cloud provider's infrastructure. Because that replication takes a certain amount of time and an application has no control over which redundant machine it accesses, testing should account for the fact that an application can access stale data. This might require changes to the applications themselves.

4 Calculating Your ROI

As with deciding to move ahead with any IT project, a thorough analysis of the ROI should be done before deciding to move to the cloud. Here are the things an organization should attempt to quantify:

- ◆ **Hardware Savings:** Moving to the cloud should reduce the organization's need for hardware. In some cases that will mean decommissioning existing machines; in other cases that will mean buying less hardware and software going forward.
- ◆ **Staffing:** With the cloud provider building and maintaining the infrastructure, fewer staff will be needed. The cloud provider's staff will maintain the actual hardware, apply patches to software and handle the day-to-day maintenance of their systems. The savings in staffing should be evaluated according to the type of cloud service being used.
- ◆ **Power and Cooling:** The cost of keeping machines turned on and cooled can be substantial. Having fewer machines in-house will decrease those costs.
- ◆ **Application Changes:** Depending on the type of application, moving to the cloud may require changes to the application itself. For applications that will be hosted in a virtual machine hosted in the cloud, changes might be minimal. On the other hand, applications that will use cloud infrastructure services instead of in-house infrastructure may require substantial changes.
- ◆ **Organizational Efficiency:** The ability to automatically provision and deprovision resources can make an organization much more responsive

and flexible. A more responsive organization has more opportunities to innovate and distinguish itself in the marketplace. This is much harder to quantify; unplugging 20 machines will absolutely lower costs for power and cooling, while flexibility will give an organization the *potential* of better performance.

- ◆ **Governance:** As covered above, using a non-private cloud means that the employees of a cloud provider will be involved in the security chains used to protect the organization's applications and data. Auditing and monitoring the cloud provider's systems will likely be more difficult. The organization should estimate how its policies will be affected and the cost of changing them. Another cost to consider is that a cloud provider might charge a fee for auditing or monitoring its services.
- ◆ **Risks:** A number of risks have been covered in this paper; those risks should be evaluated with the type of cloud service and the type of cloud (private versus non-private) being used.

With these factors in mind, an organization can do a cost-benefit analysis and a risk assessment to determine whether moving to the cloud is worthwhile.

5 Conclusions

There are many benefits to moving applications and data to the cloud, but there are many risks as well. This paper covered the areas organizations should keep in mind as they consider whether the benefits of moving to the cloud are worthwhile. The three-step process is:

1. **Classify Your Information Assets**
2. **Determine Your Risks and Requirements**
3. **Calculate Your ROI**

At every point in the process, the organization should keep its business goals and needs in mind. Moving to a cheaper, more automated system that requires less administration yet fails to provide adequate performance, security, privacy or availability is a disaster.

With a complete understanding of the applications and data and their requirements and risks, an organization can make an objective business decision about the value of cloud computing.