Cloud Computing

In the Government

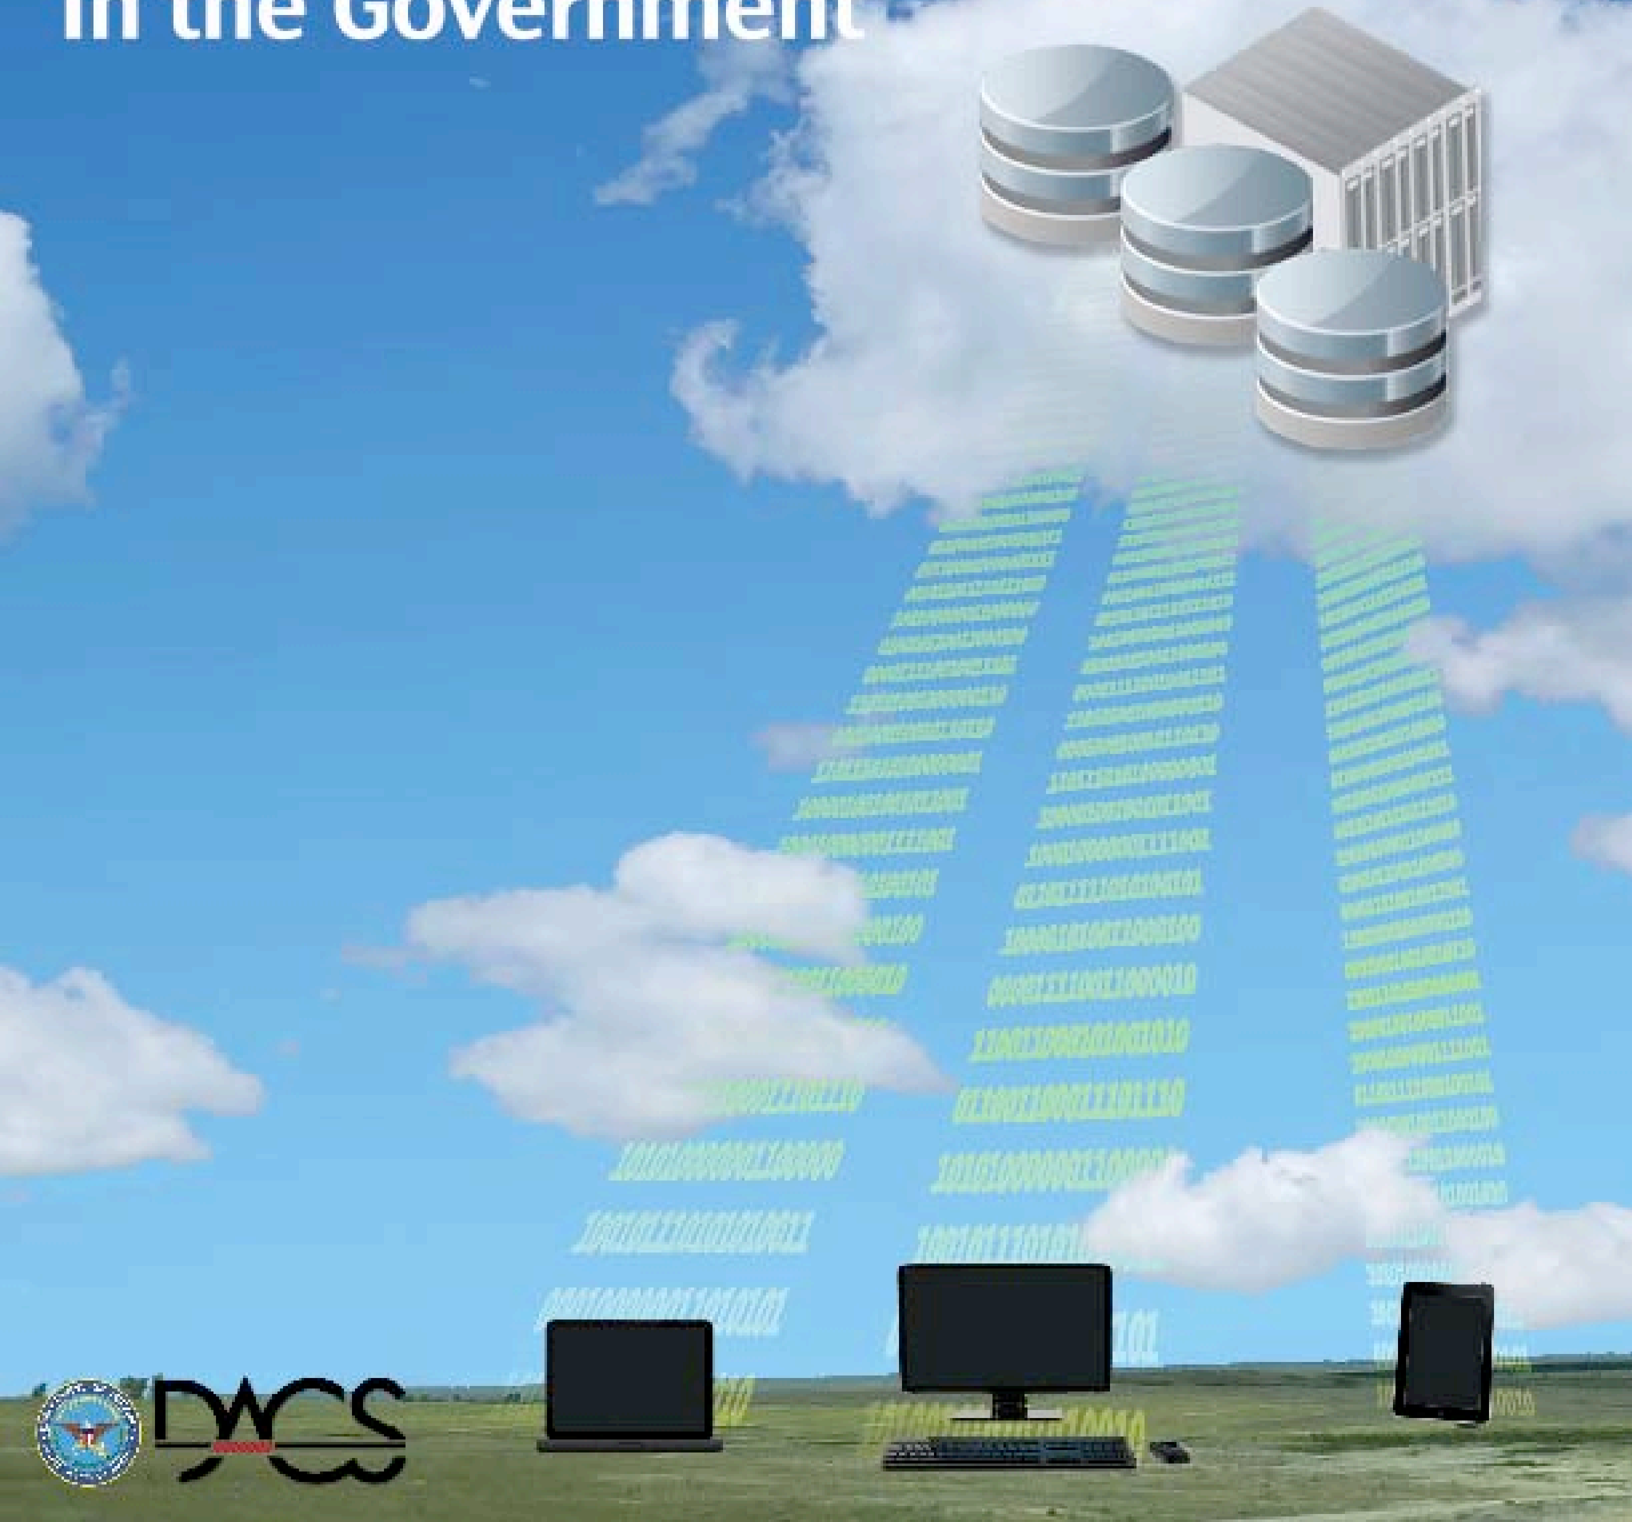

This page intentionally left blank.

Cloud Computing in the Government

A DACS Critical Review and Technology Assessment (CRTA)

DACS Report Number 518136

Contract FA1500-10-D-0010

Prepared for the Defense Technical Information Center

Prepared By

Thomas J. Kwasniewski, DACS Report Lead

EJ Puig, Consultant

Quanterion Solutions Inc.

100 Seymour Road, Suite C102

Utica, NY 13502-1311

This page intentionally left blank.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(23-05-2011)*
08 August 2011

2. REPORT TYPE
Technical

3. DATES COVERED *(From - To)*
N/A

4. TITLE AND SUBTITLE
Cloud Computing in the Government

5a. CONTRACT NUMBER
FA1500-10-D-0010

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER
N/A

6. AUTHOR(S)
EJ Puig and Thomas J. Kwasniewski

5d. PROJECT NUMBER
N/A

5e. TASK NUMBER
N/A

5f. WORK UNIT NUMBER
N/A

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Quanterion Solutions, Inc., 100 Seymour Rd Suite C102
Utica, NY 13502-1311

8. PERFORMING ORGANIZATION REPORT NUMBER

DACS DAN #518136

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Defense Technical Information Center
DTIC/AI
8725 John J. Kingman Rd., STE 0944
Ft. Belvoir, VA 22060

10. SPONSOR/MONITOR'S ACRONYM(S)
DTIC

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT
Approved for Public Release, Distribution Unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT
Cloud computing is currently one of the biggest trends in the information technology (IT) industry. Over the past several years, cloud computing has gained significant momentum to be the cost effective and efficient alternative to managing in-house software applications and hardware systems. The impact that cloud computing has had on business is so great that the United States Government is now looking to the cloud as a means to reorganize its IT infrastructure and to decrease its spending budgets. With top government officials mandating cloud adoption, many agencies already have at least one cloud system online.
This report discusses what cloud computing is and how the government is utilizing this technology. The final section of this report is a case study that is a detailed account of migrating an existing application to the Microsoft Azure cloud environment. The actual implementation of the migration and the issues found are discussed.
The observations and recommendations expressed herein are accurate as of May 19th, 2011.

15. SUBJECT TERMS
Cloud Computing, High Performance Computing, Software Migration, Storage, Grid Computing

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			Thomas McGibbon
U	U	U	UU	130	**19b. TELEPHONE NUMBER** *(include area code)* 315-351-4203

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

This page intentionally left blank.

Table of Contents

List of Figures

List of Tables

Abstract

Cloud computing is currently one of the biggest trends in the information technology (IT) industry. Over the past several years, cloud computing has gained significant momentum to be the cost effective and efficient alternative to managing in-house software applications and hardware systems. The impact that cloud computing has had on business is so great that the United States Government is now looking to the cloud as a means to reorganize its IT infrastructure and to decrease its spending budgets. With top government officials mandating cloud adoption, many agencies already have at least one cloud system online.

This report discusses what cloud computing is and how the government is utilizing this technology. The final section of this report is a case study that is a detailed account of migrating an existing application to the Microsoft Azure cloud environment. The actual implementation of the migration and the issues found are discussed.

The observations and recommendations expressed herein are accurate as of May 19th, 2011.

1 Introduction to the Cloud

Cloud computing has, over the past few years, exploded in popularity. According to Google Trends, search volume indices for cloud computing have increased almost 4-fold since 2007 (Trends). Cloud computing is the biggest trend following commodity clusters in the high performance computing sector (Thomas Sterling, 2009). Although cloud computing is a services-based computing paradigm, it can mean different things to different people. There are many technologies associated with the cloud in order to provide its vast array of services. Depending on specific needs, the cloud and what it means to a given customer can vary immensely. Davies mentions this as the cloud being like real clouds; it can be anything (Davies, Microsofts Azure Skies, 2009). The National Institute of Standards and Technology (NIST) defines cloud computing as "a model for enabling ubiquitous, 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." (Mell & Grance, 2011). A cloud must be an on-demand service, have broad network access, resource pooling capabilities, be highly elastic, and be a measured service (Mell & Grance, 2011). A visual representation of NIST's view of cloud computing can be seen in Figure 1.

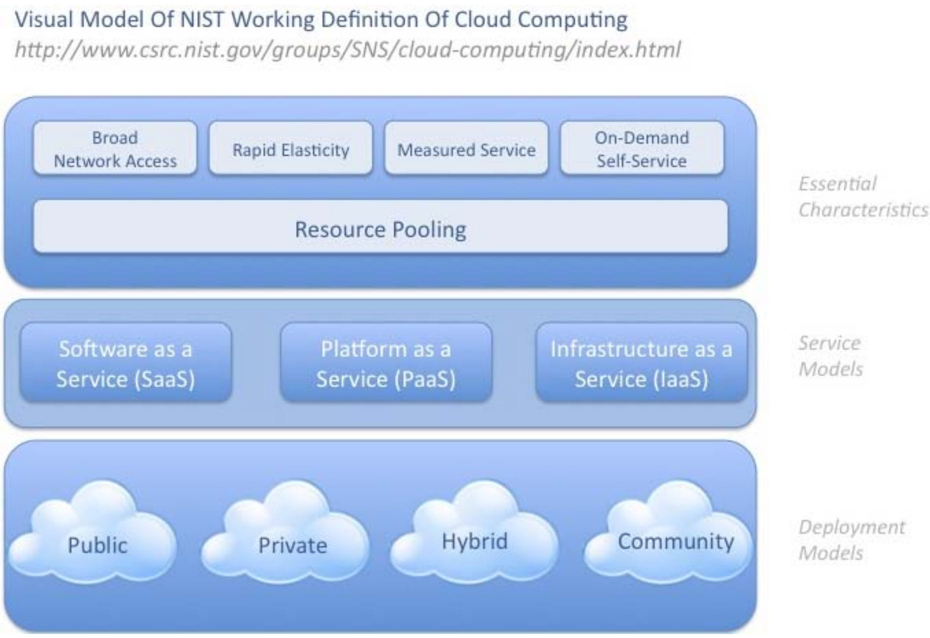

Figure 1: NIST's Cloud Computing Definition

(Cloud Security Alliance, 2011)

In terms of data processing, cloud computing is much more than just distributed or remote computing. A key difference between clusters and clouds is that a cloud without Service Level Agreements (SLAs) is nonspecific and does not guarantee identical properties of the running environment from run to run. Further, clouds favor generic services and users (Thomas Sterling, 2009). Cloud computing is also its own business model for cloud service providers. More accurately, cloud computing is a type of distributed computing utilizing on-demand services over the Internet that charge the customer only for the amount of time that they use the service(s) (Thomas Sterling, 2009). Zhang specifies further that "cloud computing provides elastic computing infrastructure and resources which enable resource-on-demand and pay-as-you-go utility computing models." (Zhang, Schiffman, Gibbs, Kunjithapatham, & Jeong, 2009). Armburst defines utility computing as the service being sold (Armbrust, et al., 2009). Cloud computing allows for the ability to create virtual machine environments that run on top of physical hardware (de Assuncao, di Costanzo, & Buyya, 2009). The user is thus abstracted away from the internal workings of the cloud. Due to this, it makes optimizing code difficult, as the user cannot be certain as to the specifics of the underlying system and hardware. (Thomas Sterling, 2009).

1.1 Advantages

The elasticity of the cloud is a significant characteristic of the cloud computing model. Elasticity in the cloud refers to the ability for up-scaling services due to performance spikes and allows for quick downsizing when usage dissipates instead of decommissioning hardware and/or cutting jobs. This, in effect, gives the customer a real-time supply of resources (Thomas Sterling, 2009).

To process large data sets, a fundamental principle of cloud computing is data locality. Data locality allows the data that a node must compute to be present locally to that node, eliminating the need to fetch the data over the network, which incurs delays. Data locality is a highly efficient way of processing such large data sets (Gu & Grossman, 2009).

Another important aspect of cloud computing is the notion of economies of scale. The beauty of cloud computing is that the cloud service provider only charges customers for the amount of time that they use the cloud. Furthermore, the pricing policy is such that the cost of using more cloud resources over a short period of time is equal to using fewer resources over a longer period of time, given the same task. "A job that might take 40 hours on a 10-node cluster could instead be run on a 400-node cloud cluster in 60 minutes for the same cost." (World, Cloud Computing for Life Sciences, 2009). What this translates to for customers is that it gives them the ability to solve very large and complex problems by harnessing the compute power they need, without the cost associated with over-provisioning.

Cloud computing makes financial sense to cloud service providers. It gives them more of a return on their hardware investments through increased resource use, which increases their hardware utilization. On the customer side, cloud computing reduces the barriers to entry for start-up companies by eliminating the need to invest in costly infrastructure. A primary aim of cloud computing is that cloud services are executed on hardware the customer does not own or operate (Mowbray

& Pearson, 2009). The customer only pays for resources used, not operational infrastructure costs (Thomas Sterling, 2009). For this reason, there is the belief that cloud computing will have a similar impact on large scale computing that foundries have had on the hardware business. Only a few companies have enough capital to own and operate the expensive facilities to manufacture semiconductors (a $3+ billion dollar industry). These foundries are the large companies that produce products for smaller companies that do not have the capital to operate their own plants. This type of relationship allows companies like nVidia to compete in the computer graphics market by transferring their own operational costs and risks to Taiwan Semiconductor Manufacturing Company (the company that manufactures nVidia's chips) (Armbrust, et al., 2009).

By using cloud computing, customers also transfer risk to the cloud service provider; especially the risk of over-provisioning and underutilization (Armbrust, et al., 2009). In other words, what the cloud can offer a customer is the freedom from investing in hardware and software to meet their worst-case needs. They can subscribe to a cloud service provider and pay for the resources that they actually use (Mowbray & Pearson, 2009). This directly affects IT staff. By using the cloud, they are free from the burden of the monotonous tasks that are involved with keeping software and hardware up and running. Transferring this management to a cloud service provider allows IT staff to focus on innovative ideas that address more important issues within their business (Davies, Amylin, Amazon, and the Cloud, 2009). Lastly, using the cloud enables the customer to have access to the latest technologies at a reduced cost, rather than paying a premium and purchasing them outright (Thomas Sterling, 2009).

1.2 Structure

1.2.1 Public vs. Private

In the most general sense, cloud computing is divided into two halves. There are private clouds and public clouds. The differentiation between the two is that private clouds are usually owned by a corporation or university for personal use (Han, Hassan, Yoon, & Huh, 2009). Companies and universities generally use private clouds over public clouds primarily for security reasons (Murray, 2009). In contrast, the public cloud is open for general use, so long as any and all subscription and usage fees for the service are paid. Currently, there are departments within the government, such as the Department of Treasury, which have already taken the leap into the public cloud, bypassing the private cloud altogether. To date, there are many public cloud providers. Some public cloud offerings include IBM, Amazon, Google, and Microsoft.

1.2.2 Other Forms

Although the public/private delineation is the primary way to view cloud computing, there are other types that have emerged. A hybrid cloud is another type of cloud. Hybrid clouds are typically implemented where an organization has their own internal private cloud, as well as a subscription to a public cloud offering. The public cloud subscription augments an organization's private cloud. An example would be storage and backup. Instead of a company running main storage and backup operations, a company can have their primary storage within an on-site private cloud and have this private cloud replicated as backup on a public cloud storage service. Another example would be if a company uses their private cloud for processing and utilization suddenly spikes. When this happens, workloads can be migrated to public clouds to stabilize the system.

There are also community clouds. A community cloud is shared among a very small subset of organizations. An organization might turn to this type of cloud environment when they fear the security issues of the public cloud, but do not have the financial resources to deploy their own private cloud. Sharing essentially a private cloud among two or more organizations keeps security high, while driving down the individual investment.

1.2.3 Layers of the Cloud

Regardless of a cloud being public or private or any of the other types, a cloud is made up of three distinct layers, as shown in Figure 2. These layers are Platform as a Service (PaaS), Software as a Service (SaaS), and Infrastructure as a Service (IaaS). These three layers are general categories, which different services belong to, that cloud service providers make available.

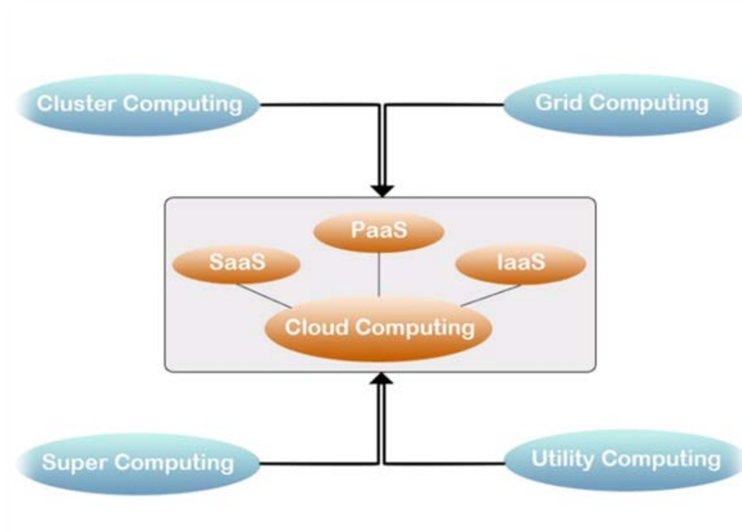

Figure 2: Supporting Technologies of the Cloud Environment

1.2.4 Software as a Service (SaaS)

The SaaS layer of the cloud literally provides software over the Internet. If an application runs in the cloud and it provides direct service to a customer, it runs at the SaaS layer (Lenk, Klems, Nimis, Tai, & Sandholm, 2009). The customer interacts with this software as if it is directly installed on their personal machine. The advantage of SaaS is that it eliminates the installation procedures of software. SaaS also provides management of this software from a central location, without the need of performing multiple configurations at different site locations. This includes updates and patches to SaaS software, as these are performed in a central location which frees users from needing to download and install the updates.

In addition to these advantages, cloud computing allows a customer to deploy such software over the Internet without requiring the customer to provision or build a data center. The cloud service provider handles this burden on their behalf (Armbrust, et al., 2009). Although SaaS might feel new, it is not the first technology to deliver software to a customer via a third party. A service offered at this layer is Google's e-mail service Gmail. Other examples can be seen in **Table 1**.

Company	Service	Description	Layer
Google	Google Docs	"Office" Tools	SaaS
	Gmail	Web based e-mail services	SaaS
Microsoft	Office Live	MS Office Tools Online	SaaS
Salesforce	Salesforce.com	CRM Tols	SaaS

Table 1: Examples of SaaS Services

Types of SaaS

Lenk, et. al., defines two types of applications that run at the SaaS layer. These types of applications are Basic Application Services and Composite Application Services. Essentially, Composite Application Services are applications comprised of Basic Application Services. Any Composite Application Service can be used as a Basic Application Service to develop more complex Composite Application Services. Examples of Basic Application Services include Google Maps and OpenID. An example of a Composite Application Service is MySpace (Lenk, Klems, Nimis, Tai, & Sandholm, 2009).

SaaS Implementation

In the SaaS model, a vendor providing a software service hosts the application and other necessary data such as a database on servers within their own data center. These servers have dedicated support staff. What this accomplishes is that it eliminates a customer's responsibility to purchase and manage the software and servers to host the software themselves. In addition, this allows the customer to extend the lifecycle of their desktop computers significantly, as a remote program places a reduced load on a machine, saving even more money. In short, by hosting applications remotely, a business' budget for IT can be better spent on subscriptions to more software, rather than on hardware to support the software (Chong & Carraro, 2006).

SaaS Maturity Model

SaaS software services are divided up into a four-level maturity model. These four levels are:

1. *Ad Hoc/Custom*
2. *Configurable*
3. *Configurable, Multi-Tenant Efficient*
4. *Scalable, Configurable, Multi-Tenant Efficient*

Level I (or Ad-Hoc/Custom) is quite similar to the Application Service Provider (ASP) model of the past. Each customer has their own instance of an application. Each customer also has their own servers for the software being hosted. Although each customer has their own application and server instance, this customer may actually be a whole company. Employees from this company all connect to the same server and application instance. This is similar to typical Line of Business (LOB) software. Migrating applications to a Level I SaaS model is fairly easy. Little benefit comes from this level, but it does allow such providers to consolidate physical hardware (Chong & Carraro, 2006).

Level II (or the Configurable level) enables the ability to provide multiple customers the same application with the same code base. This is unlike Level I, where the provider had to customize an application for each of its customers. At this level, the customer is provided with customization options to change how the application looks and functions. Even though each customer shares the same code base, each instance of the application remains completely isolated from the others. Moving to Level II allows a SaaS provider to greatly reduce administrative costs of keeping software up to date with upgrades and patches, as they only need to service one code base for all customers. To migrate an application to this level, the application may need to be redesigned to work with configuration metadata.

Metadata allows each customer to customize an application themselves (Chong & Carraro, 2006).

Level III (or the Configurable, Multi-tenant Efficient level) allows a SaaS provider to run an application as a single instance that each customer uses. Configuration metadata is used, as in Level II, to allow customers to change the way the application behaves. To keep customer data separate, authorization and encryption schemes are used. The benefits of this level allow the SaaS provider to consolidate its servers per application running at this level. This results in cost savings for the customer, as utilization of the provider's hardware is greatly increased. The disadvantage to this level is that scalability is limited in the sense that the only way to scale this particular application is to migrate it to a faster server (Chong & Carraro, 2006).

Level IV (or the Scalable, Configurable, Multi-Tenant Efficient level) is essentially Level III on a load-balanced farm, or a cluster of servers. This level allows applications to be easily scaled up and down depending on the demand for the application. Like all lower levels, individual customer data is kept separate and secure from other customers (Chong & Carraro, 2006).

Metadata

Configuration metadata as mentioned earlier allows a customer to change the way an application looks and functions. Configuration metadata relieves a SaaS service provider from customizing specific instances of an application for each customer. By implementing configuration metadata, an application can be provided as a single code-base to all customers and allow those customers to change the interface to the way that they need. Customers can customize an application in the following ways:

User Interface/Branding

This enables a company to use their company's logo and other images as well as font sizes and colors (Chong & Carraro, 2006).

Workflow/Business Rules

This allows a company to change aspects of the work flow of the application to match it with a given company's methodology (Chong & Carraro, 2006).

Extensions to the Data Model

This allows a company to change the way data and information is stored within the application. The result is that the customer is not restricted to doing things predetermined by the designers of the application. They can fit the application's data model into their company how they want. This would include things like changing fields in the tables within the database the application uses (Chong & Carraro, 2006).

Access Control

This allows a company to create user accounts for an application. The company can also set access restrictions on a per-account basis for the application (Chong & Carraro, 2006).

1.2.5 Platform as a Service (PaaS)

Platform as a Service, or PaaS is the second layer of the cloud computing model (CCM). PaaS is a platform provided as a service to application developers. An instance of this layer is a complete application stack on a virtualized computing platform. PaaS, as shown in Figure 3, allows web developers to create and deploy applications without needing to purchase and manage the underlying hardware of the application. Furthermore, developers can use PaaS as a testing suite for their applications. Also, developers can subsequently deploy their application within the same environment. By having development and production happen on the same platform, this makes management of the application significantly easier.

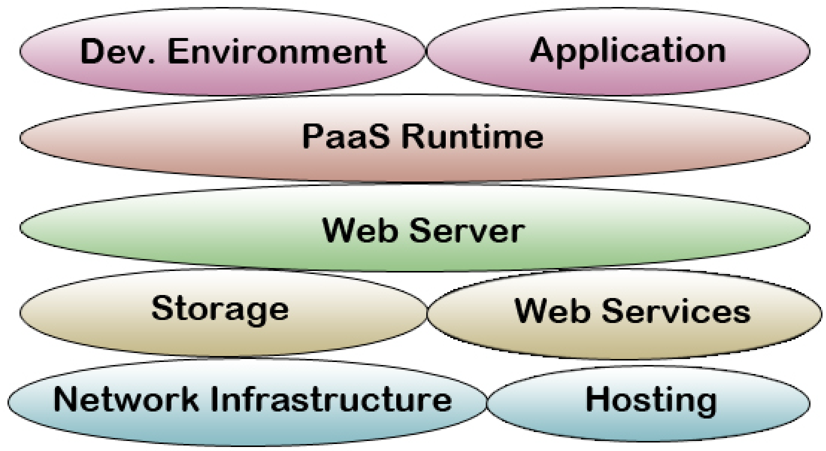

Figure 3: Creating Web Applications in a PaaS Environment

This is quite different from how traditional web applications and web sites have been developed. Currently, the majority of web sites are still developed as static pages. With the PaaS model, the different technologies, and the various development and testing environments that are required to build these websites are consolidated. PaaS allows developers to quickly move through the application development lifecycle. Using a subset of commands, a developer can move from development to testing, and testing to deployment with ease. This allows the company a faster time to market with their product (Hinchcliffe).

Some PaaS services may be language specific. For example, Google App Engine is a PaaS platform that Google provides that uses Python and Java. Others, like Salesforce.com, have proprietary languages ((Architectural Strategies for Cloud Computing), (Hinchcliffe)). These technologies can be seen in **Table 2**. Other examples include Sun's Caroline, Microsoft Azure and the Django web programming framework (Lenk, Klems, Nimis, Tai, & Sandholm, 2009). The developer need not concern themselves with implementing scaling abilities of their application as this, in addition to secure data communication and billing, are implemented automatically by the PaaS provider (Hinchcliffe).

Company	Service	Description	Layer
Google	App Engine	Cloud-based web 2.0 application development environment	PaaS
Microsoft	Azure	Development and runtime environment for MS applications	PaaS
	Live Mesh	Tool to connect to remote devices and resources	PaaS
Salesforce	Force.com	Develop and deploy business applications	PaaS

Table 2: Examples of PaaS Services

1.2.6 Infrastructure as a Service (IaaS)

Finally, the third layer, Infrastructure as a Service (IaaS), is on the lowest level, closest to the physical hardware (Lenk, Klems, Nimis, Tai, & Sandholm, 2009). IaaS is a service such that a customer is provided with hardware including servers and storage, as well as the software needed to use this hardware, e.g., a virtualized operating system. IaaS allows a customer to build their own virtualized system. A customer creates a virtual machine image with all of their required software packages and uploads it to an IaaS provider. The IaaS provider will then run the virtual machine within their cloud. A major benefit of IaaS is that it allows a company to move their applications to the cloud without requiring them to re-architect it. Reasons for not re-architecting an application include not having the time or budget. IaaS gives these companies an easy way to enter into the cloud environment (Chong & Carraro, 2006).

Lenk, et. al., divides IaaS into three layers. The first layer, closest to the hardware, consists of the Physical Resource Set (PRS) and the Virtual Resource Set (VRS) services. Both of these types of services implement an application programming interface (API) to allow for the management of a customer's virtual image. The difference between the two is that PRS is hardware-dependent. This means that a PRS implementation is provider-specific. A VRS implementation is hardware/vendor independent and runs on top of a hypervisor technology such as Xen. Emulab, a network test bed, is an example of a PRS, while Amazon Elastic Compute Cloud (EC2) is a VRS offering. The next layer up is what Lenk, et. al., defines as the Basic Infrastructure services, which include the computational, storage and the network resources. Examples of these are MapReduce, GFS, and OpenFlow, respectively. Finally, the last layer is the Higher Infrastructure Service. Higher Infrastructure Services are built using Basic Infrastructure Services to create services such as Bigtable and Dynamo (Lenk, Klems, Nimis, Tai, & Sandholm, 2009). The concept here is similar to the division of SaaS applications that was discussed earlier.

A major difference between IaaS and PaaS is that, with IaaS, the customer is responsible for managing their IaaS instances. This is exactly the way it would be if they had to manage these systems within their own data center. The cloud service provider's responsibility is to simply keep their data centers up and running. The IaaS provider is not responsible for any down time of any instances that is not a direct result of their data center being down. Amazon's EC2 and their Secure Storage Service (S3) are two of the very many IaaS offerings (Architectural Strategies for Cloud Computing). Examples of IaaS services can be seen in **Table 3**.

Company	Service	Description	Layer
Amazon	Elastic Compute Cloud (EC2)	Virtual resource environment	IaaS
	Dynamo	Structured storage service	IaaS
	Simple Storage Service (S3)	Raw storage service	IaaS
	SimpleDB	Database service	IaaS
	CloudFront	Content delivery network	IaaS
	SQS	Resource messaging service	IaaS
Google	Google Big Table	Storage service	IaaS
	Google File System	File system service	IaaS

Table 3: Examples of IaaS Services

1.3 Origins

Cloud computing resulted from the mounds of data that large companies, like Amazon and Google, have collected over the years. They needed some way to process and analyze this information. Their efforts formed a new data-processing model that fit nicely into their own data analysis and “service-oriented" needs (Thomas Sterling, 2009). Cloud computing, although new, was actually foreshadowed by John McCarthy back in the 1960's. McCarthy, a computer and cognitive scientist believed, that eventually, compute power would be provided as a metered service (An Introduction to Cloud Computing, 2009).

In the history of computing, there were a few specific technologies that paved the way for cloud computing, as shown in Figure 4. Specifically, these technologies are Application Service Providers or ASPs, cluster computing, grid computing, virtualization, and utility computing.

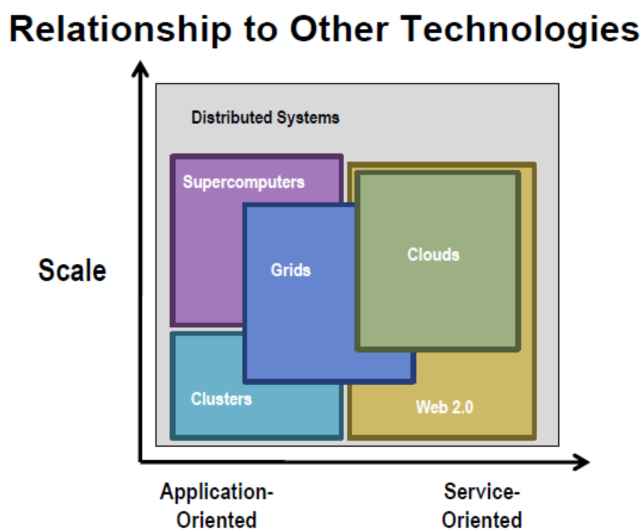

Figure 4: Relationship of HPC Technologies
(Foster, Zahu, Loan, & Lu, 2008)

1.3.1 Application Service Providers

Application Service Providers were the first to deliver applications to a customer over the Internet in the mid-to late-90's. As a third party, ASPs hosted applications that customers needed. This reduced costs to the customer, as it freed them from the necessary management to maintain these applications internally. The problem with ASPs was that every new customer required their own customized version of an application. This meant the ASP would need to set up a new server within their own data center to provide this new customer with their custom application, even if the underlying application was being served to hundreds of other companies. Furthermore, these ASPs did not know enough about the domains in which these applications were developed to make the customizations that customers demanded. On top of this, customization fees incurred for the customer were very high. As a result, the majority of ASPs failed (Difference Between the ASP Model and the SaaS model, 2009).

There are differences between ASPs and SaaS technologies. ASP-oriented applications were client-server applications that were provided over the Internet. SaaS applications are designed specifically to be accessed using a web browser. Another significant difference is that, with SaaS applications, every user shares the same environment for the application. That means everyone shares the same servers, code and configurations. This is known as multi-tenancy, something that ASP lacked. This multi-tenancy of SaaS allows for economies of scale. What this means is, as more customers start to subscribe to the service, costs decrease. SaaS allows the service provider to make simple upgrades to systems without the need to consider special customizations, as one would in the ASP model. Since everyone shares the same environment, the SaaS provider can market their services at a significantly reduced cost that would otherwise be very expensive for one business to purchase outright, or within the ASP model (Difference Between the ASP Model and the SaaS model, 2009).

1.3.2 Virtualization

Virtualization, invented by the Burroughs Corporation, has been around since the 1960's (O'Connor, 2009). Virtualization was originally used to partition mainframes that suffered from underutilization. IBM popularized this technology when they implemented virtualization into their own mainframes, allowing their mainframes to multitask. The motivation for implementing virtualization was to increase utilization, thus maximizing their return on investment (ROI) (VMware History, 2010).

High Performance Computing (HPC) users benefit immensely from virtualization. These benefits include multiple operating systems having the ability to be present on one physical machine at a time. This allows for parallel execution of jobs running in each virtual operating system. This concept can be applied to increase the utilization of servers in data centers. This increases a company's ROI, as well as reducing hardware and maintenance costs. Another benefit regarding virtualization is that each virtualized operating system can be specialized or custom tailored to a specific HPC job (Mergen, Uhlig, Krieger, & Xenidis, 2006).

1.3.3 Cluster Computing

During the 90's, supercomputers were the de facto solution for large computations, as well as web and database servers. Their popularity declined since they were very expensive to purchase and manage. Also, supercomputers were always behind the technology curve, meaning that they were quick to become obsolete. They were also difficult to upgrade without changing the majority of the components within the system. One essentially needed to purchase an entire new supercomputer. Although clusters were nothing new at the time, they started to gain popularity. Clusters are just multiple computers connected together to solve a specific problem. The reason for their popularity at this time was due to them being built using Commercial off the Shelf (COTS) hardware. Using COTS hardware cut costs considerably. Other reasons for their takeover included advancements in networking technologies, and university projects proving the efficiency of clusters built using COTS equipment (Mark Baker, 2000). These projects included Beowulf, Berkeley NOW and HPVM (High Performance Virtual Machines). All three projects exploited COTS hardware. The overall goal of these projects was to provide a cost-effective alternative to supercomputer performance.

Beowulf

The Beowulf project headed by Donald Becker and Thomas Sterling was a break away from large supercomputers. They built the system using COTS components. The Beowulf project was effectively a cost-effective supercomputer alternative (Merkey, 2007).

Berkeley NOW

Berkeley NOW was another academic project. NOW is an acronym for Network of Workstations. Berkeley NOW was a project focused on creating a cluster out of relatively inexpensive hardware. Berkeley NOW believed that, for a better price/performance ratio than a node from a massively parallel processing (MPP) system, they would utilize COTS equipment to implement their cluster. The overall goal of Berkeley NOW was to provide a better cost-to-performance ratio for parallel applications than a typical MPP system, as well as better performance for sequential programs that typically run on a single workstation (Berkeley NOW, 1998).

High Performance Virtual Machines

The HPVM project at the University of San Diego also focused on exploiting the performance of COTS hardware to use in a high performance computing environment. Like a grid, this project involved distributed resources. The difference between this project and others focusing on clusters and grids is that they implemented virtual machines, allowing the system to execute jobs in parallel. The most noteworthy adopter of HPVM is the National Computational Science Alliance (NCSA). The NCSA used HPVM to develop and improve tools and libraries used in scalable parallel computing. As a result of the popularity of HPVMs, it created an increased adoption of the Windows NT system (High Performance Virtual Machines (HPVM), 1999).

Advantages to Cluster Computing

There are several advantages to clusters. The low cost of hardware is a major advantage, since it reduces a business' investment in the system. The cost/performance ratio is much higher with clusters compared to a supercomputer. Inherent in their design, clusters are built using many COTS machines and connecting them together. This allows for incremental upgrade. In contrast to a supercomputer, where essentially the whole system needs to be replaced in order to upgrade the system, a cluster can be upgraded by either connecting more COTS hardware to the system, or by individually replacing one of the nodes within the cluster. Contributing to the low cost of clusters is the open source development of the architecture for a cluster. Open source implies that there is no fee in obtaining the software, further reducing a company's investment. Finally, clusters can be made out of any type of computing hardware. Customers needing a cluster do not have to submit to any one particular brand, as they can make the cluster out of whatever they want. This is not the same for a supercomputer, as these are made by specific vendors. Due to the fact that clusters can be made from any brand of COTS hardware, vendor lock-in is also eliminated (Mark Baker, 2000).

1.3.4 Grid Computing

Grid computing is essentially an extension of cluster computing. One could think of a grid as a number of clusters connected together. However, there are key differences between grids and clusters:

- Clusters generally consist of homogeneous hardware, whereas grids typically are heterogeneous
- Clusters are local to a problem; whereas grids are decentralized
- The motivation to use a grid over a cluster is dependent on the size of the computation one wishes to perform

There are several companies currently using grids. Two well-known grid projects are LHC@Home and the World Community Grid. LHC@Home is a project where volunteers contribute their personal computer's idle time to aid scientists and physicists to develop and exploit particle accelerators, much like CERN's Large Hadron Collider (LHC@Home, 2010). The World Community Grid is another public, voluntary computing grid that is used to benefit humanity. The World Community Grid is exclusively available to public and nonprofit organizations for humanitarian research that otherwise would not be able to be carried out due to the high cost of purchasing a grid for their personal use. The results from executing jobs on the World Community Grid are released to the public domain (World Community Grid, 2010). Grid computing was an initial stepping stone to cloud computing, as it allows users to compute large distributed jobs remotely and at a reduced price, compared to that same user purchasing the infrastructure outright that would perform at the same level as the grid.

1.3.5 Utility Computing

The essence of utility computing is a service provider charging a customer for the amount of time that the customer uses the provider's service over the Internet. The title and concept mimic real utility companies, like water and gas companies. Each of these companies charges their customers for the amount of water or gas that they use. In the realm of cloud computing, potential metrics for charging a customer can include the amount of gigabytes used for storage capacity, or millions of instructions per second (MIPS) executed over the elapsed time that the customer uses a service (Philip, 2004).

Although utility computing and grid computing share similarities, they are, at their roots, very different. Grid computing is a system designed to search for idle compute resources. It is designed to be a "virtual super computer" (Philip, 2004). Utility computing, on the other hand, is a model such that when a system runs out of resources, it subscribes to resources from another location through a usage-based payment method. The utility computing model only reaches out to the service provider when the lack of local resources demand such an action in order to finish a job or meet a deadline (Philip, 2004).

1.4 Utility vs. Grid vs. Cloud Computing

Although utility, grid and cloud computing sound alike, all three have unique differences. Since the emergence of the cloud, grid and utility computing can be viewed as specific subsets of cloud computing. Grids pull resources from specific domains to achieve a single result. Utility computing rents computer resources, including network bandwidth, in an on-demand model. These concepts from utility and grid computing are incorporated into the CCM, creating a new computing paradigm. A major difference between cloud computing vs. utility and grid computing is that the cloud is not restricted to specific networks. The cloud exists on the ubiquitous Internet (Biswas, 2011).

A further distinction between cloud and grid computing is that, even though grids have some virtualization present, it is not up to the same standard as the cloud. Due to virtualization and other characteristics inherent in the cloud design, the cloud is extremely robust and stable. Grid computing, however, has a higher probability of catastrophic failure (Biswas, 2011).

1.5 A Need for High Performance Computing

The final push for the emergence of the cloud was the abundant upsurge in need for High Performance Computing (HPC) resources. The U.S. federal government realized this need beginning with the "Cash for Clunkers" program that began back in June 2009. When dealerships began using the system in July of that year, it only took three days for the system to crash due to higher than estimated usage. To be better prepared for the future, top officials within the government demand the adoption of cloud services which can meet the demand of such increased usage (Kundra, 25 Point Implementation Plan to Reform Federal Information Technology Management, 2010).

2 Cloud Computing in the Government

As already discussed, cloud computing is a natural way for individuals, groups, businesses and organizations to save on IT costs. These savings do not just apply to the private sector. Organizations within the public sector such as the U.S. Federal Government can also benefit from cloud computing. Due to their expensive and inefficient operations, the federal government has declared that federal IT must start a migration to the cloud, which could save the government up to $20 billion dollars annually (Feds Could Save $20 Billion Using Cloud Computing, 2011).

In a memorandum sent to all Chief Information Officers (CIO) within the government, the CIO of the White House, Vivek Kundra, stated that the total number of government-owned data centers grew from 432 in 1998 to more than 1,100 in 2009. Based on agency submissions, as of July 30, 2010 the number of data-centers doubled to 2,094 in one year (OMB Asks Agencies to Review Data Center Targets, 2011). Kundra also reported that, in 2006, federal data centers consumed 6 billion kWh of electricity with the fear of energy consumption doubling by 2011 (Kundra, Federal Data Center Consolidation Initiative, 2010). In addition to these figures, Kundra also stated that federal data centers are running at relatively low utilization rates. Due to these reasons, Kundra created the Federal Data Center Consolidation Initiative (FDCCI).

The FDCCI aims to reduce energy, hardware and software costs, while also decreasing the real estate owned by the government for their data centers (Data Center Consolidation Plan, 2010). One example of the FDCCI being implemented is through the General Dynamics contract with the Defense Intelligence Agency (DIA). In a five year, $40 million dollar contract, General Dynamics will consolidate DIA's seven regional help desks into two enterprise help desks in Washington D.C. and Colorado (DIA awards SITE contract for help desk support, 2011). An overview of the FDCCI can be seen in Figure 5.

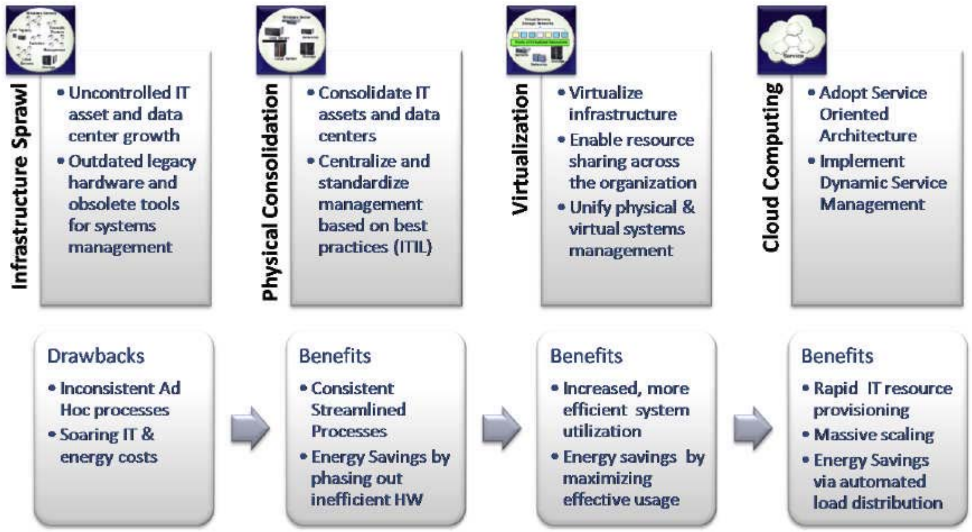

Figure 5: Overview of the FDCCI
(FedSources, 2010)

To lead this effort, Kundra requested every CIO within the government to submit asset inventories and consolidation plans to the CIO Council throughout the year of 2010. The CIO Council is the "...principal interagency forum for improving agency practices related to the design, acquisition, development, modernization, use, sharing, and performance of Federal information resources." (About CIO.gov, 2011). All final consolidation plans were to be submitted and approved by the Office of Management and Budget by December 31st, 2010 (Kundra, Federal Data Center Consolidation Initiative, 2010). Kundra aims to consolidate at least 800 data centers (Walker, 2011).

The FDCCI is a six phase process. The six phases, in order, are Asset Baseline Inventory (I), Application Mapping (II), Analysis and Strategic Decisions (III), Consolidation Design and Transition Plan (IV), Consolidation and Optimization Execution (V), and, finally, Ongoing Optimization Support (VI). Phase I is information gathering of an agency's IT state, including physical hardware, software and usage metrics. Phase II is application profiling. The work in this phase will be the foundation that consolidation and cloud migration plans will be based on. Phase III catalogs major systems. Agencies are also asked to identify solutions and associated transition plans to achieve optimal utilization through either virtualization or cloud computing in order to meet cost savings targets. Phase IV requires agencies to develop technical standards and architectures to use virtualization and cloud computing within their departments. Phase V is the execution of the agency's transition plan. Finally, Phase VI is the ongoing monitoring of the new environment to keep systems optimized (Data Center Consolidation Plan, 2010). A detailed roadmap of these phases can be seen in Figure 6 below.

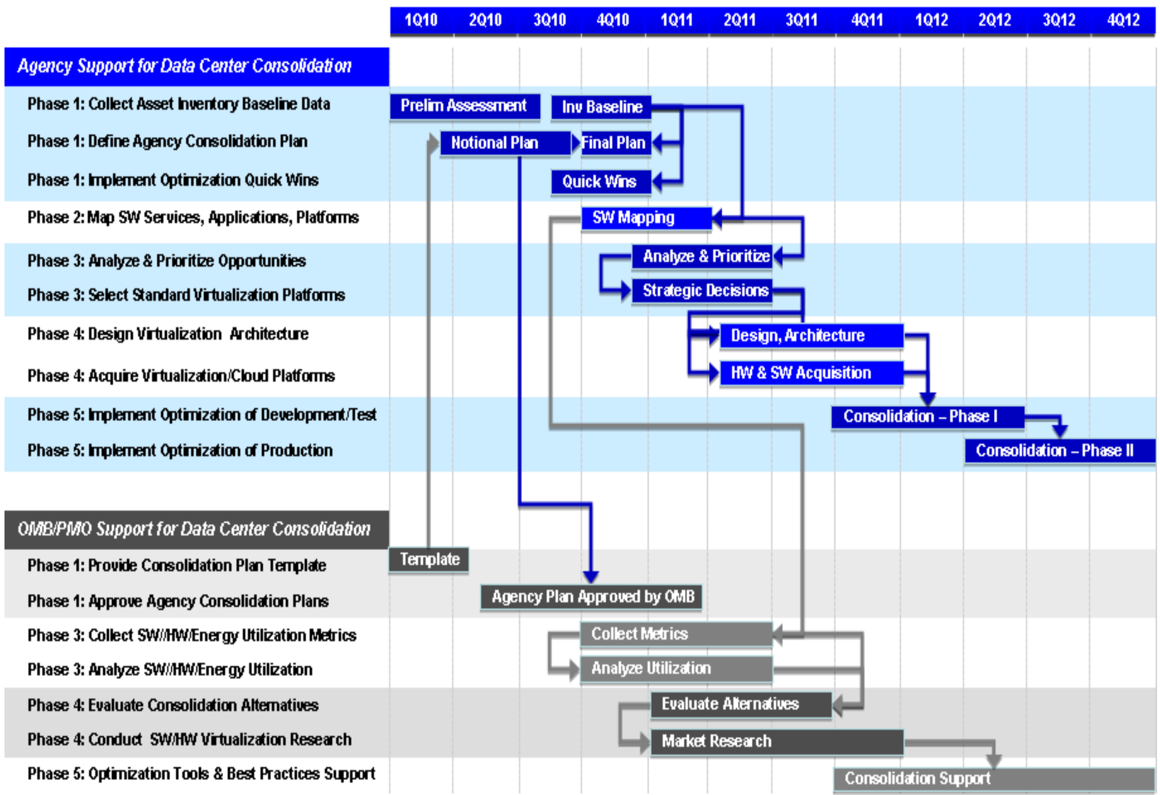

Figure 6: FDCCI phases

(Data Center Consolidation Plan, 2010)

Kundra's FDCCI is a stepping stone, or an initial phase for the government to migrate to the cloud. Kundra's "Cloud First" or Federal Cloud Computing Initiative (FCCI) is part of his "25-Point Implementation to Reform Federal IT" plan. Kundra's strategy to bring the cloud into the government is to use the savings, as shown in Table 4, gained by shutting down 800 data centers from the FDCCI. The savings from this effort will be invested in the cloud initiative. The government is ready to allocate up to $20 billion into cloud spending (Walker, 2011).

Approach	Description	Potential Benefits	Rationale
Decommission	Turn off servers that are not being used or used in-frequently (e.g. dedicated development environments)	• Cost Savings • Energy Efficiency • Frees Floor / Rack Space	• As many as 10-15% of servers may be inactive but still powered on in data centers*
Centralization / Site Consolidation	Move servers/storage to a few selected data centers Consolidate small data centers to larger target centers	• Floor Space Cost Savings • Operational Cost Savings • Increase Rack Utilization • Energy Efficiency	• Approximately 430 Government data centers are categorized as "closets" or small sized data centers (less than 1,000)**
Virtualization	Consolidate several servers onto a single server through virtualization of the OS/Platform	• Floor Space Cost Savings • Increase Rack Utilization • Increase Server Utilization • Energy Efficiency	• Server Utilization is approximately 21% Government wide**
Cloud Computing Alternatives	Move application functions to standard, vendor supported enterprise platforms or services	• Floor Space Cost Savings • Energy Efficiency • Operational Cost Savings • Cap Ex Cost Savings HW/SW • Reduced SW Maintenance • Improved Service Delivery	• Reduce Operational Risk, lower TCO and TCSD • Approximately 40% of Civilian Agency Systems are low-impact FISMA security, and therefore may be low-risk candidates for Cloud Computing solutions

** McKinsey Report: Revolutionizing Data Center Efficiency, July 2008*
*** OMB BDR 09-41 Data Analysis, October, 2009*

Table 4: Expected Benefits of Cloud Adoption
(Data Center Consolidation Plan, 2010)

The cloud adoption timeline is rather aggressive. Kundra requires every government agency to have a single cloud solution up and running by Dec. 2011, and three by June 2012. Savings from adopting the cloud are already being realized by some agencies. The General Services Administration (GSA) has saved $15 million dollars and the US Department of Agriculture (USDA) $20 million dollars by transitioning to the cloud. These are savings from a single cloud solution in each agency (Walker, 2011).

Cloud computing is maturing at a fast pace, though the security policies and standards for government cloud adoption are not. Due to the lack of standardization, many government agencies are taking a cautious step into the cloud before fully committing. Through pilot programs such as migrating e-mail into the cloud, agencies can test these policies and standards to gain a better understanding of cloud computing. This will delay the initial adoption of many cloud endeavors within the government (Yasin, Implementing the cloud-first policy? Start with e-mail, 2010).

To help agencies migrate to the cloud, the government has appointed key groups with varying responsibilities to help expedite the process and to ensure quality cloud adoption principles. The National Institute for Standards and Technology (NIST) has a primary role in the FCCI. Their responsibility is to develop interim standards for cloud computing that relate to high priority security, interoperability, and portability requirements (Technology, 2011).

The Standards Acceleration to Jumpstart Adoption of Cloud Computing (SAJACC) works with the NIST to provide an NIST-hosted portal to exchange candidate information that could potentially be used in standards formation. SAJACC was formed to accelerate standards creation and increase confidence for adopting cloud computing (Standards Acceleration to Jumpstart Adoption of Cloud Computing (SAJACC), 2011).

The federal CIO Council created the Cloud Computing Executive Steering Committee (CCESC) to provide overall management of the FCCI. The committee is led by GSA CIO Casey Coleman and includes IT executives from more than 15 federal agencies (Executive Steering Committee, 2011). The CCESC, in turn, appointed subcommittees. The Cloud Computing Advisory Council is a group of senior IT experts from across the government who are nominated by their department CIOs. They are representatives of best practices, consensus building, and informers of ongoing cloud projects. They also are the primary educator for their respective agency or department for cloud computing within the government (Info.Apps.Gov, 2011). The Communications Working Group manages apps.gov and relays information about the cloud computing initiative (Communications Working Group, 2011). The primary goal of the E-mail Working Group is to promote the adoption of cloud-based e-mail within the government. They are responsible for assuring secure and compliant solutions (E-mail Working Group, 2011). The Operational Excellence Working Group works on data center consolidation efforts, as well as promoting cloud pilot programs (Operational Excellence Working Group, 2011). The Security Working Group is responsible for defining the security needs and solutions of the government within the cloud environment (Security Working Group, 2011). Finally, the Standards Working Group is responsible for addressing portability, interoperability, and manageability issues concerning the government's involvement within the cloud space (Standards Working Group, 2011).

2.1 Cloud Initiatives within the Federal Government

2.1.1 Defense Information Systems Agency (DISA)

The Department of Defense (DoD) is very active with cloud computing. The DoD has declared a Department of Defense Tools Initiative. This initiative aims to maximize the efficiency and quality of software development, as well as reduce duplicate and conflicting software and systems. To meet these goals, the Defense Information Systems Agency (DISA) is developing a suite of tools that can be accessed via their private cloud for software development projects (See Table 5). These products can be found on DISA's Forge.mil website.

Forge.mil refers to a suite of software project tools that are hosted by DISA. These tools are based on CollabNet's TeamForge application life cycle management tool. The goals for the Forge program are to promote collaboration among everyone involved with a given project, including stakeholders, throughout the development lifecycle. Software modularity and standardized test and evaluation criteria are also goals of Forge (Defense Information Systems Agency, 2010). Currently, SoftwareForge and ProjectForge are available. With almost 9,000 users, SoftwareForge enables collaborative software development (Bernhart, 2011). SoftwareForge comes with many features, including software version control, bug tracking, requirements management, and release packaging. SoftwareForge also has wikis, discussion forums, and document repositories (Defense Information Systems Agency, 2010). ProjectForge is similar to SoftwareForge in that it provides the same development and collaborative tools. The difference is that ProjectForge is designed for more restrictive access. Another differentiating factor between the two is that SoftwareForge is free for all valid users. ProjectForge is a service-fee based product. There are currently two types of offerings for ProjectForge, ProjectForge 'On-Demand' and 'Private'. ProjectForge 'On-Demand' is hosted as a multi-tenant application. What this means is that the infrastructure is shared by multiple projects and users who can access their projects concurrently. Permissions that are set and managed by the project owner prevent un-authorized access to projects. ProjectForge 'Private' offers single tenancy, providing the customer with exclusive resources (Defense Information Systems Agency, 2010).

The Forge program also has three tools currently in development. These tools are TestForge, CertificationForge and StandardsForge. TestForge is for testing within the development lifecycle of a software project. TestForge includes functionalities such as defect management; automated unit, functional and regression testing; and a static code analysis. TestForge will also include performance, scalability, reliability, interoperability, operational, Net-Ready Key Performance Parameter (KPP), and Information Assurance (IA) testing. CertificationForge will minimize time and money spent on a project, as each stakeholder of a project won't necessarily have to perform their own testing. Standardization of certification criteria and reporting is a major part of this tool. Every organization and stakeholder involved with a software project will be able to speak the same language, ensuring trust and high visibility. StandardsForge will enforce software development standards throughout the development process. The primary goals for this tool are (Defense Information Systems Agency, 2010):

- to establish a standardized process for development,
- create a reference for special terms and vocabulary for software development,
- define expectations between developer and customer, and
- define the type of software being developed, (e.g., newly developed, modified, reused, bought out, prepared to other standards, or procured off-the shelf).

Given the successes of the project, the future of Forge.mil will support the intelligence community, including the Department of Homeland Security. Currently, however, DISA's plans to support the intelligence community are still in their infancy (Bernhart, 2011).

Another cloud technology that DISA owns is the Global Information Grid (GIG) Content Delivery Service (GCDS). The GIG itself is a 3.7 petabyte system spanning over 14 global datacenters, connected to the Defense network (Brewin, 2011). The GCDS is a "distributed computing platform comprised of globally deployed servers across both the NIPRNET and SIPRNET optimizing the delivery of DoD Web content and Web-based applications." (Defense Information Systems Agency (DISA), 2011). The value of the GCDS is immense. The GCDS was put to the test back in March, 2011, when Japan suffered catastrophic damage from the earthquake and tsunami that affected their country. The United States Navy responded to the natural disaster in several ways, one being the reconnaissance of the Navy Visual News Service (NVNS). They were responsible for sending images and footage to national and international media establishments. The NVNS quickly realized that their personal system was insufficient for the task at hand. Due to the system residing within the United States, NVNS suffered significant networking delays. Before long, the GCDS team provided the Navy with the support that they needed. Integrating the Navy within the GCDS took only six hours, which holds the record for the fastest GCDS integration. Performance was immediately realized. Information dispersion increased dramatically (DISA, 2011). The success of GCDS proves the ability to retrieve and send applications and data in adverse network conditions (Defense Information Systems Agency (DISA), 2011).

Initiative	Description	Layer
Forge	Software Development Tools Suite	PaaS
GCDS	Content delivery, application deployment	PaaS

Table 5: DISA Cloud Initiatives

2.1.2 Army

In response to the FDCCI, the Army is calling upon DISA, as part of their cloud initiatives (see Table 6), to migrate its e-mail services into the cloud. The Army's current e-mail solution is rather expensive. The Army spends $100 dollars per user annually for their email account (Kenyon, Army presses forward with e-mail move, 2011). DISA will host their new enterprise e-mail service within one of the Defense Enterprise Computing Centers. Migration is scheduled to be completed by Sept. 30, 2011. This migration will affect over 1.5 million users, both civilian and some 'secret' accounts. This effort is estimated to save the Army $100 million dollars annually. Another benefit of this migration will allow users to access their accounts from any DoD facility and collaborate with other users worldwide. The Army expects to pay as little as $39 dollars per e-mail account after the migration. The new service will allow 4GB worth of data storage, as opposed to 100MB in the old system. The new system will give extensions to everyone's email address for identification reasons. These extensions will classify active duty military, civilian employees, reservists and contractors (Kenyon, Army presses forward with e-mail move, 2011). After a successful e-mail migration, the Army will begin to implement other cloud solutions such as Enterprise Active Directory (Defense Information Systems Agency, 2010).

The Army Experience Center wanted to upgrade their 10 year old Customer Relationship Management (CRM) Army Recruiting Information Support System (ARISS). With security being top priority, the Army carefully chose a commercially-available SaaS solution. A primary reason why they chose a SaaS solution over upgrading ARISS is that the bids to upgrade ranged from $500,000 to $1,000,000 dollars. The SaaS solution only cost the Army $54,000 to implement. Not only did they save a considerable amount of money, their cloud solution could handle five times the amount of work that a single traditional recruitment center could handle (Kundra, Federal Cloud Computing Strategy, 2011).

With all of the sensor data the Army collects, the Army has designed a solution to keep up with the deluge of information. Army officials have augmented their Distributed Common Ground System-Army (DCGS-A) that collects and shares data and intelligence with cloud computing technologies. By implementing cloud technologies into the DCGS-A system, they can now analyze data as frequently as they desire. This augmented system has been installed and is being tested at Bagram Air Field in Afghanistan. The outlook is for this system is to be installed at edge nodes with a data cache for when bandwidth is limited or non-existent. The cache would automatically update when reconnected to the network. In the future, the system will implement 3G and 4G network capability to increase its bandwidth resources (Corrin, Army pursuing high-tech data sharing at the tactical edge, 2011).

Initiative	Description	Layer
E-mail	Migrate Army E-mail service to DISA Cloud	SaaS
CRM	Migrate ARISS operations to the Cloud	SaaS
DCGS-A	Data collection and dissemination	PaaS

Table 6: Army Cloud Initiatives

2.1.3 Air Force

The United States Air Force (see Table 7) has contracted IBM to architect a secure, self-tuning cloud infrastructure for them in a ten month contract agreement (Humphries, 2010). The Air Force manages about 100 bases and about 700,000 active military personnel globally. The goal of the project is to allow the Air Force to control and monitor the secure information that travels throughout their network. Monitoring will be provided in the form of real time dashboards that deliver status reports on the overall health of the network. From this dashboard, officials can immediately respond to cyber attacks and other network issues. Another aspect of this project is the implementation a self-tuning capability, inherent in its design. to automatically optimize the network in response to dynamic changes (Barrett, 2010). IBM is expected to use their InfoSphere Streams technology for the monitoring services that will enable the Air Force to analyze the massive amounts of data flow within the network (Yasin, Air Force, IBM plan to demonstrate secure cloud computing, 2010).

Initiative	Description	Layer
Self-Tuning Cloud	Secure, Self-Tuning Cloud	IaaS

Table 7: Air Force Cloud Initiative

2.1.4 Navy

The Navy has halted the procurement of new servers and hardware for their data centers. After the FDCCI announcement. Kundra reported that, in the fall of 2010, the DoD had 772 data centers. The Navy plans to consolidate by 25 percent and increase their utilization by 40 percent (Censer, 2011).

2.1.5 Department of Homeland Security

Responding to the FDCCI plan, the Department of Homeland Security (DHS) is reducing their 24 datacenters down to 2 (Perera, 2010). They have future plans to move applications from Citizen and Immigration Services (CIS), Customs and Border Protection, Immigration and Customs Enforcement and United States Visitor and Immigrant Status Indication Technology (U.S. VISIT) into a Justice Department-run facility by the end of the year. Due to security, the Department of Homeland Security (See Table 8) will primarily use a private cloud. This private cloud will be divided into partitions. Each partition will be a "trust zone" in order to deal with the varying security needs of various programs within the DHS (Perera, 2010). There are public-facing agencies within DHS that are able to use the public cloud, such as Federal Emergency Management Agency (FEMA) and CIS. (Yasin, After Data Center Consolidation, Beware Legacy Apps, 2011). These specific agencies are migrating their SharePoint and email services to the cloud (Perera, 2010). By using cloud computing, DHS estimates that it could save $2.4 Billion dollars (Feds Could Save $20 Billion Using Cloud Computing, 2011).

Initiative	Description	Layer
Email/SharePoint	Email and SharePoint Services via Cloud	SaaS

Table 8: Department of Homeland Security Cloud Initiative

2.1.6 Veterans Affairs Department

The US Department of Veterans Affairs (VA) is trying to embrace cloud technologies (Table 9). There is a big push from both doctors and employees as they attempt to use the cloud to make their daily tasks much easier. The main issue with generic services is that they are not secure enough for department use. IT security staff are blocking access to these tools, since there is a history of employees using services without authorization, thereby creating potential security risks (Hoover, VA Employees Using Unauthorized Cloud Services, 2010).

Despite security concerns, the VA is still finding ways to adopt the cloud. In a project called the "Big4", the VA department wants to migrate 600,000 of its employees' email accounts into the cloud. The title of the project refers to the number of data centers that will be required to host the system. The VA has declared that it would consider private cloud solutions, but has also stated that it wants the management of the system entirely outsourced. This means, VA staff will not have administrative access to the applications and systems (Montalbano, VA to Migrate Email to the Cloud, 2011).

Initiative	Description	Layer
Big4	Migrate Email Services to the Cloud	SaaS

Table 9: The VA Department Cloud Initiative

2.1.7 Department of State

The U.S. State Department (see Table 10) is currently decreasing its number of data centers from eleven to two. At the moment, the State Department has a private IaaS cloud called the International Information Programs Content Management System (IIPCMS). This system, managed by Computer Technologies Consultants, hosts over 400 ".gov" websites (Computer Technologies Consultants, 2010). They are hoping to expand their private cloud, to include SaaS, that will run applications such as Microsoft SharePoint. The State Department does not utilize the public cloud due to security and governance issues (Yasin, After Data Center Consolidation, Beware Legacy Apps, 2011).

Initiative	Description	Layer
IIPCMS	Web hosting	IaaS

Table 10: Department of State Cloud Initiative

2.1.8 Department of Treasury

In response to the FDCCI, The Department of Treasury (Table 11) has retooled their website and migrated it to Amazon's EC2 cloud. They also migrated the Treasury's SIGTARP.gov, MyMoney.gov, TIGTA.gov, and IRSOversightBoard.treasury.gov sites to the cloud, as well. The Department of Treasury is the first cabinet-level federal agency to use Amazon to host a Web site (Montalbano, Treasury Sites Jump to Amazon Cloud, 2011). In total, the Department of Treasury estimates that it can save $2.4 billion dollars by utilizing cloud services (Feds Could Save $20 Billion Using Cloud Computing, 2011).

Initiative	Description	Layer
Web Migration	Migrate websites to Amazon’s EC2 Cloud	IaaS

Table 11: Department of Treasury Cloud Initiative

2.1.9 Department of Agriculture

The Department of Agriculture (Table 12) has already saved $20 million by migrating services to the cloud, and expects to increase their savings as more services are migrated (Walker, 2011). One cloud solution implemented within the USDA consolidated their fragmented 21 system, 120,000 employee e-mail service (Foley, 2010). Upgrading their old system would have taken years, as opposed to implementing the cloud solution within a matter of months. This cloud solution immediately removed duplication within the system. They also experienced an immediate savings of $6 million per year (Kundra, Federal Cloud Computing Strategy, 2011).

Initiative	Description	Layer
E-mail	Cloud hosted E-mail	SaaS

Table 12: Department of Agriculture Cloud Initiative

2.1.10 Department of Energy

The Department of Energy (DoE) has been consolidating their data centers (Table 13) for several years. Currently, they have two primary datacenters in Germantown, MD and Albuquerque NM. They have been able to reduce the number of servers by half, from 200 to 100 (Yasin, After Data Center Consolidation, Beware Legacy Apps, 2011).

The DoE's Argonne National Laboratory is working with IBM to create Mira, a 10 petaflop supercomputer. The design is based on Blue Gene. It is scheduled to be operational by 2012. To put the performance of this supercomputer into context, the fastest supercomputer to date is Tianjin National Supercomputer Center's Tianhe-1A system, which is capable of 2.67 petaflops. Mira is a stepping stone for the U.S. endeavor for exascale computing. In addition to Mira, IBM is developing a 20 petaflop supercomputer, Sequoia, for the DoE's Lawrence Livermore National Laboratory, as well as a 10 petaflop machine, Blue Waters, for the National Science Foundation (NSF) (Kenyon, Energy aims to retake supercomputing lead from China, 2011).

Funded by the American Recovery and Reinvestment Act through the Department of Energy, project Magellan is to test the feasibility of cloud computing for scientific study. The project is underway at the Argonne Leadership Computing Facility and the National Energy Research Scientific Computing Center. Mid-range computing hardware will be installed at these locations to create a cloud testbed to examine the effectiveness of cloud computing in a research environment as a cost-and-energy efficient computing alternative (Argonne National Laboratory, 2011).

Initiative	Description	Layer
Mira	10 Petaflop Supercomputer	IaaS
Sequoia	20 Petaflop Supercomputer	IaaS
Magellan	Test cloud within a research environment	N.A.

Table 13: Department of Energy Cloud Initiatives

2.1.11 Department of the Interior

The Department of the Interior's National Business Center currently has four cloud projects (Table 14) underway. NBC Hybrid allows users to seamlessly link NBCFiles to NBCGrid, their IaaS service. NBCApps is their cloud app marketplace. NBCAuth is their security and authentication service that allows users to move seamlessly between different NBC cloud apps without being prompted for their user account information. NBCFiles is their cloud storage service (Department of the Interior National Business Center (NBC), 2011).

Initiative	Description	Layer
NBCGrid	Resource Provisioning	IaaS
NBCApps	Online SaaS Marketplace	SaaS
NBCAuth	Directory and Authentication Service	SaaS
NBCFiles	Cloud Storage	IaaS

Table 14: Department of the Interior Cloud Initiatives

2.1.12 Department of Health

The Department of Health (Table 15) has begun implementing the Electronic Health Records (EHR) system. To help coordinate and manage this task, the DoH has called upon Salesforce.com to provide customer relationship and project management services. The primary use of the Salesforce.com CRM services is to support DoH's Regional Extension Centers in guiding doctors and rural hospitals to understand, use and implement EHR systems (Kundra, Moving to the Cloud, 2010).

Initiative	Description	Layer
CRM	Train doctors to use the EHR systems	SaaS

Table 15: Department of Health Cloud Initiative

2.1.13 General Services Administration

The General Services Administration (GSA), as shown in Table 16, has taken a broad role in the federal government to provide GSA with cloud computing services that increase efficiency, optimize common services and solutions across organizational boundaries as well as enable a transparent, collaborative and participatory government. To carry out this task, the GSA has set up apps.gov. On this website, the GSA provides a large selection of different types of cloud services. These services are primarily SaaS applications. However, the GSA does provide IaaS IT Services, which consist of storage and virtual machine solutions. They also provide a growing list of business applications such as asset and business management SaaS solutions. In addition, the GSA provides productivity applications. These applications consist of tools that people might use on a daily, basis such as collaboration and word processing tools. Lastly, the GSA provides social media apps so that employees and personnel can share information with one another. The Cloud services managed by the GSA are provided by a strict subset of companies who have met government standards and regulations. The vendors currently developing applications for the GSA are

- Apptis Inc. (partnered with Amazon Web Services),
- AT&T, Autonomic Resources (partnered with Carpathia),

- Enomaly,
- Dell,
- CGI Federal Inc.,
- Computer Literacy World (partnered with Electrosoft),
- XO Communications and Secure Networks,
- Computer Technologies Consultants, Inc., (partnered with Softlayer, Inc.),
- Eyak Tech LLC,
- General Dynamics Information Technology (partnered with Carpathia),
- Insight Public Sector partnered with Microsoft,
- Savvis Federal Systems, and
- Verizon Federal Inc (Wali, 2010).

In addition to these initiatives, the GSA announced a request for quotation (RFQ) on May 9th, 2011 for office automation and records management, among other services that will be available through SmartBuy Blanket Purchase Agreements (BPA). It is expected that multiple contracts will be awarded, with a total value of $2.5 billion dollars over the course of five years (Lipowicz, 2011).

The GSA also has its own internal cloud initiatives in the works. The GSA is transitioning from IBM LotusNotes to Google Gmail and Google Apps to cut costs. They expect to reduce their email costs by 50% with the migration over the next five years (Hoover, GSA Picks Google for Email, 2010).

Initiative	Description	Layer
Email	Migrate From IBM LotusNotes to Gmail	SaaS
apps.gov	Online Cloud Marketplace	SaaS & IaaS

Table 16: GSA Cloud Initiatives

2.1.14 National Aeronautics and Space Administration

National Aeronautics and Space Administration's (NASA) Nebula, as shown in Table 17, is under development at the Ames Research Center. Nebula is a cloud environment that will provide all three service levels (SaaS, IaaS, and PaaS). Nebula promotes rapid software development, secure web applications, and code reuse (NASA, 2011). For security, NASA is implementing a virtual Local Area Network (LAN) and data encryption techniques into their Nebula system. This will allow for partitioned data traffic. If any traffic is intercepted, it will be scrambled and impossible to decipher. NASA is also implementing a near real-time security audit within their system. The slightest change within the system, such as an IP request, will trigger a new security audit (Joch, 2011).

Initiative	Description	Layer
Nebula	NASA Cloud Computing Platform	SaaS, PaaS, IaaS

Table 17: NASA Cloud Initiative

2.1.15 North Atlantic Treaty Organization

The North Atlantic Treaty Organization (NATO), as shown in Table 18, is requesting IBM to build an on-premises private cloud system. The goal for this system is to cost-effectively introduce new technologies and consolidate old systems. This is a test for NATO to see if the cloud can develop new solutions for their daily operations, including situational awareness and decision making. A secondary goal for this cloud is to inspire member nations to adopt new technologies of their own to modernize and consolidate their own IT systems (Montalbano, NATO Taps IBM to Build Private Cloud, 2010).

Initiative	Description	Layer
NATO Cloud	IBM to Build Private Cloud for NATO	IaaS

Table 18: NATO Cloud Initiative

2.1.16 National Science Foundation

The National Science Foundation (NSF) has its hands in many diverse research projects relating to cloud computing. The project "Comparative Study of Approaches to Cluster-Based Large Scale Data Analysis" is a partnership between the Massachusetts Instititute of Technology (MIT), Yale and the University of Wisconsin.

The project compares and contrasts MapReduce and parallel database systems for scalable data processing. This project is being hosted on the Cloud Computing Testbed (CCT) operated by researchers at the University of Illinois, Urbana-Champaign (UIUC) (National Science Foundation (NSF), 2011).

The "Hadoop Toolkit for Distributed Text Retrieval" project is taking on the challenge of retrieval of text in a large search space, harnessing the power of Hadoop to try and solve the issues of robustness and scalability. The objective of this project is to modify Hadoop, as it was designed for only batch processing and not real-time search problems. One proposed solution is a distributed in-memory object caching architecture (National Science Foundation (NSF), 2011).

Other projects funded by the NSF are:

- The "Unified Reinforcement Learning Approach for Autonomic Cloud Management" project, which automates the configuration processes of virtualized machines and applications. The goal of this project is to automate resource management in the cloud (National Science Foundation (NSF), 2011).
- The "Commodity Computing in Genomic Research" project aims to develop parallel algorithms to analyze the next generation of sequencing data (National Science Foundation (NSF), 2011).
- The "Data-Intensive Text Processing" project deals with machine translation. The plan is to use network analysis and cross-language information retrieval techniques to create a richer, contextual model for the machine to translate text more accurately (National Science Foundation (NSF), 2011).
- The "Feedback-Controlled Management of Virtualized Resources for Predictable eScience" project aims to accurately report the status of virtualized resources within a cloud framework through feedback control theory (National Science Foundation (NSF), 2011).
- The "Hierarchically-Redundant, Decoupled Storage Project (HaRD)" is developing next generation storage software. The overall objective is to

improve the performance of storage, especially in new use case scenarios such as online photo albums and large-scale data processing. The hope is that this will become an improved file system for MapReduce workloads (National Science Foundation (NSF), 2011).

- The "One Thousand Points of Light" project proposes a cloud proxy network that allows optimized and reliable data-centric operations to be performed at strategic network locations. In this model, proxies may take on several data-centric roles: interacting with cloud services, routing data to each other, caching data for later use, and invoking compute-intensive data operators for intermediate processing. The proposed solution will enable an efficient coupling of cloud services to yield improved end-to-end performance and reliability for newly emerging data-intensive applications (National Science Foundation (NSF), 2011)."
- The "Scaling the Sky with MapReduce/Hadoop" project focuses on developing new algorithms necessary for indexing, accessing and analyzing the petabytes of data associated with astronomical imagery (National Science Foundation (NSF), 2011).
- The "Trustworthy Virtual Cloud Computing" project investigates fundamental research issues leading to new security architectures. This research includes development for new "...security services that enhance the trustworthiness of virtual cloud computing, protection of management infrastructure against malicious workloads, and protection of hosted workloads from potentially malicious management infrastructure." (National Science Foundation (NSF), 2011).
- The "Image Super-Resolution Using Trillions of Examples" project, with the goal being to provide the ability to "infinitely zoom" into a picture using on-line image repositories for reference as to what the picture would look like at the higher resolution (National Science Foundation (NSF), 2011).

- The "Learning Word Relationship Using TupleFlow" project attempts to improve the efficiency of web retrieval results by studying word relationships (National Science Foundation (NSF), 2011).
- The "Where the Ocean Meets the Cloud" project proposes to build a new infrastructure designed to allow massive oceanic simulations and queries on these simulations at interactive speeds (National Science Foundation (NSF), 2011).

2.1.17 National Oceanic and Atmospheric Association

To meet the demands of the FDCCI and the FCCI, the National Oceanic and Atmospheric Association (NOAA), as shown in Table 19, currently has a large internal IT initiative underway. NOAA has two "500-day plans" that they are enacting to overhaul their IT infrastructure. One of their overall goals is to make scientific data easily available to the public. The first 500-day plan is to consolidate all IT. The second is to mature their IT infrastructure to become an effective and efficient IT service. The goals are to secure NOAA's information and IT investments from threats; build robust high performance computing (HPC) capabilities; operate NOAA IT as a customer-focused service provider; increase efficiency and effectiveness through Enterprise-wide IT Solutions; and attract, develop, and retain a skilled IT workforce (NOAA IT Strategic Plan: Executive Summary 2010, 2010).

NOAA has already taken great strides to meet their goals. They have already implemented an initial cyber security ”nerve center” to protect NOAA assets and information. NOAA has also completed the initial phases of the NOAANet program to unify and secure NOAA networks nationwide to reduce costs and to improve performance. In addition, NOAA has completed the initial design for a new Pacific Region Center to Consolidate Pacific NOAA IT and its facilities to a centralized location to achieve organizational efficiency and facilitate collaboration. It is necessary for NOAA to adopt a more efficient and effective means for facilitating higher amounts of environmental data collection and improving the quality of that data, which will allow NOAA to provide forecasts with higher accuracy. Finally, NOAA awarded an initial Unified Communications contract to provide NOAA with a common platform for email, calendar, collaboration and directory tools (NOAA IT Strategic Plan: Executive Summary 2010, 2010).

IT security is an important aspect to the revamping of NOAA's IT, as they are the official voice of the U.S. for severe weather warnings. "Damage to, or loss of, complex, high-cost technology assets such as satellites, ground control systems, and space weather systems due to cyber attacks could take years to repair or replace, at the potential cost of millions of dollars." (NOAA IT Strategic Plan: Executive Summary 2010, 2010).

High Performance Computing will allow NOAA to better forecast the weather by allowing them to use more data in their weather models, increasing their accuracy. Cloud computing will give NOAA the computational resources it needs to provide accurate weather forecasts (NOAA IT Strategic Plan: Executive Summary 2010, 2010). Currently, the Computer Sciences Corporation has a $317 million contract with NOAA to build a new forecast modeling supercomputer. The goal is to have it be one of the top 10 fastest supercomputers in the world (Lais, 2010).

Knowledge retention is highly important to NOAA as baby-boomers are getting ready to retire. NOAA wants to implement a multi-dimensional approach by forging relationships with colleges and universities to provide internship opportunities, while also providing continual education services for current employees. In addition, they want to create wiki pages for employee reference (NOAA IT Strategic Plan: Executive Summary 2010, 2010).

Initiative	Description	Layer
E-mail	Migrate E-mail services to the Cloud	SaaS

Table 19: NOAA Cloud Initiative

2.1.18 Recovery Accountability and Transparency Board

A statement made by Chairman Earl E. Devaney of the Recovery Accountability and Transparency Board (RATB) announced its migration to the cloud (see Table 20). On April 26th 2010, Recovery.gov moved into Amazon's cloud. It estimates that within the first year and a half they will save $750,000 in their effort (Devaney, 2010). The $750,000 in savings is planned to be redirected into its oversight operations of identify fraud, waste and abuse (Recovery.gov Moves to Cloud Computing Infrastructure, 2010)

Initiative	Description	Layer
Web Migration	Migrate website hosting to Amazon's Cloud	IaaS

Table 20: RATB Cloud Initiative

2.2 State Governments

State and local governments are realizing they can benefit from cloud computing just as much as the federal government. In fact, state and local governments have been quicker to adopt cloud computing when compared to the federal government. Microsoft reports that state and local governments are more "agile" over their federal counterparts, making it easier for them to utilize the cloud more quickly. Microsoft has announced 14 new initiatives at the state and local level which include Chicago, IL, Andover, MN, Virginia Beach, VA, and Carlsbad, CA (Montalbano, Local Governments Embrace The Cloud Faster Than Feds, 2011).

Changes are taking place in other state and local governments, as well. New Mexico deployed a private cloud at the end of 2010 to save on operational and infrastructure costs. Los Angeles has partnered with Google to supply its employees with Google's cloud-based SaaS-based Gmail service (Yasin, State IT diet: Consolidation, cloud and shared services, 2011). For the rest of California, the state's government has awarded Computer Sciences Corporation (CSC) a three year, $50 million dollar contract to migrate the state employees' email accounts to either the California Email Service (CES) (based on Microsoft's Business Process Online Suite), or state-hosted C.A. Mail. State agencies favored the Microsoft cloud solution 2-to-1 over migrating to C.A. Mail. Over the next three years, CSC will deal with 130 email systems that California has used to migrate over 105,000 email accounts to CES and over 56,000 accounts to C.A. Mail. California, in this effort, strives to reduce energy consumption and telecommunications hardware by as much as 20% and reduce data center square footage by 50% by July 2011. Also under the plan, California wishes to further reduce energy and hardware costs by 30% by July 2012 (Montalbano, California Agencies Favor Microsoft for Email, 2010). This migration will provide the state with a powerful email service, legal eDiscovery services, and collaborative tools for mobile device users. This migration will eliminate 130 email systems that use three different email services (CSC Wins Cloud Services Contract for California, 2010). Also in California, police officers that guard the Port of Los Angeles have created a collaborative private cloud. This cloud enables them to communicate with any device across any type of IP network, allowing them to share video streams and other data relating to emergencies that they respond to. In the future, public applications based on this technology will be able to perform "real-time, cloud-enabled collaboration across databases, resource centers and even jurisdictions." (McCloskey, 2011).

Minnesota has a five-year plan to consolidate more than 36 data centers down to about two. After the data center consolidation, they plan to move 40% of their critical business operations to a third-party data center service. The rest of their state operations will be transferred to upgraded facilities (Montalbano, Minnesota to Consolidate Datacenters, 2010). Minnesota is also migrating 33,000 employees to Microsoft's Business Productivity Online Suite cloud. In addition to the usual reasons to migrate to the cloud, Minnesota also sees the migration as necessary to alleviate the concerns over the fact that 50% of their employees will be eligible for retirement within the next ten years. According to Minnesota, this is a situation that "creates a huge risk in [administering] government services." (Thibodeau, Minnesota to move e-mail to Microsoft's Cloud, 2010). The state of Minnesota feels that they will be better prepared to deal with retiring employees if they utilize the cloud (Thibodeau, Minnesota to move e-mail to Microsoft's Cloud, 2010).

New York State allows its public, private and chartered schools the ability to use Google Apps Education Edition. Currently, the New York Institute of Technology, New York Teach Centers, Board of Cooperative Educational Services, teachers unions and state professional organizations have adopted the service. The potential customers for New York State could total over 3 million users. Considering New

York and other states, over 8 million customers have subscribed to the service (Claburn, 2010). In another project, New York City gave Microsoft exclusive rights to provide the city with their cloud products. This partnership is estimated to save New York City $50 million over five years (Yasin, State IT diet: Consolidation, cloud and shared services, 2011). During this process, they are working towards centralizing IT management underneath NYC's Department of Technology and Communications (Montalbano, Microsoft Signs NYC to Cloud Computing Deal, 2010). In an unrelated project, a $7.7 million contract with IBM is consolidating 14 of 50 NYC agency datacenters by the end of the year (Yasin, New York City getting data centers under one roof, 2011). The agencies affected in this consolidation include the Finance and Sanitation departments and the City's Chief Medical Examiner. Eventually, all 50 data centers will be consolidated into one 18,000 square foot facility in Brooklyn. The facility and initial cost of hardware will total $11.7 million. Rent for the space will be $2.7 million annually. Under what NYC calls the Citywide IT Infrastructure Services (CITIServ) Program, the consolidation of the 50 data centers NYC will save $100 million by the year 2014. Currently, CITIServ supports 140,000 employees and is saving NYC $200,000 annually. More savings are expected as more agencies are transitioned into the program (Montalbano, NYC Opens Consolidated Data Center, 2011).

Utah has consolidated its 35 datacenters into two, a primary and backup storage pair. Utah was also able to virtualize 75% of their servers. They currently maintain 75 servers. Michigan and Colorado are providing their state, cities and counties with cloud-based services such as email (Yasin, State IT diet: Consolidation, cloud and shared services, 2011). Also, in October of 2010 the state of Wyoming moved its 10,000 employees into Google's cloud, saving the state $1 million annually (Zyskowski, 2011).

3 Mobile Cloud Computing

With the government becoming more efficient by utilizing the cloud, “smart” technologies will start to take a more significant role within the government. A benefit to allowing mobile devices within a business is that, in theory, mobile devices do not have nearly the number of attack vectors as desktop computers. This would lighten up the load of IT management responsibilities for keeping information secure (Vance, Mobile Cloud Computing: 5 Key Trends, 2011). In reality, however, there is the rise in compound threats that are beginning to affect mobile devices. AdaptiveMobile conducted a survey on their network that indicated the presence of compound attacks, which primarily capture login information to user's online bank accounts (Hackers using multiple attack vectors to breach mobile phones, 2011). As smartphone market penetration surges to 40%, with nearly one billion mobile cloud subscriptions by the year 2014, authorization to use mobile devices within the government will be in great demand. (Vance, Mobile Cloud Computing: Is Your Phone Drifting to the Cloud?, 2009). IT professionals must begin to develop policies that will allow mobile device usage within the government, rather than continuing to deny access.

There are many ways to improve IT security policies, such as disallowing data storage and transmission from mobile devices and limiting activity within the corporate system depending on GPS location. Another method is to divide data into access levels and prevent mobile devices access to the most secure data (Vance, Mobile Cloud Computing: 5 Key Trends, 2011). These "common sense" policies have been considered by DISA, which is currently drafting a set of policies that will allow broad-spectrum mobile technologies to use the Global Information Grid without security risks (Corrin, New mobile device policy planned for DOD, 2011). This mobile device adoption stems from DISA's net-centricity movement. One of their goals is to harness the benefits of mobile computing for combat use. Although security risks have been a huge hurdle to overcome, they are slowly adopting a business model of saying "yes" to technologies, rather than falling back on their habit of saying "no" to new technological developments (Corrin, Culture, policies hinder technology adoption at DOD, 2011).

Mobile devices have limited computing resources. Using the cloud to mitigate this issue seems like a natural fit. Berkeley research scientists have come up with a solution. CloneCloud utilizes clouds to provide additional compute resources for compute intensive tasks. The way it works is that the phone communicates with the cloud to orchestrate execution over one or more clones of the mobile device. After the phone calculates battery life and the amount of work to send, the phone then transmits the work to the cloud. To test the system, a researcher created a face recognition application. The application took 100 seconds to run natively on the phone, as opposed to one second on the cloud. Although these are promising results, there are key concerns that need to be addressed for the mobile cloud, which include network bandwidth, availability and security (Perez, 2009).

To address some of the primary security issues with mobile devices and cloud computing, Android v2.2, Google's open-source mobile phone API, incorporates many enterprise security policies that are necessary in order for mobile cloud computing to be viable within the enterprise environment. These security measures include device locking and remote wiping. With the Google Apps Premier and Education Edition App, IT administrators will be able to control mobile access to files and directories without overhead that would interfere with users and their personal files and activities. For this to work, each managed device must have the Google Apps Device Policy application installed (Ingthorsson, Android 2.2 and mobile cloud computing, 2010).

Better web browsers are necessary for the acceleration of mobile cloud computing. Specifically, these new browsers need to make the work being processed remotely look like it is being processed local to the device to deliver a “native installation” feel. Currently, development for cloud apps on mobile devices is cumbersome. There are native browsers on mobile devices that come as part of the underlying operating system. There are also multi-platform browsers that can be downloaded by third party vendors and installed as an application on top of the operating system. This creates the lack of a standard for how an application should function, as these browsers don't all process data in the same way (Ingthorsson, Advanced browsers are key to mobile cloud computing, 2011). The future hope is that mobile web browsers will use standard APIs, such as OneAPI by the GSM Association. Currently, companies such as FeedHenry and RhoMobile are providing tools that allow app development for multi-platform mobile devices without needing expert knowledge of the specific platforms (Ingthorsson, Advanced browsers are key to mobile cloud computing, 2011).

Two pilot programs have been implemented to test the practicality and applicability of mobile phones within government work environments. The Army has released an iPhone app developed by C2 Technologies of Vienna, VA to teach the Army's Patriot Missile crews about the Patriot system in order to use it. The development of this application stems from the Army's Connecting Soldiers to Digital Applications Program, which allows the soldier to train anytime and anywhere (Sideman, 2011). Within the Department of the Interior, the Office of the Secretary is currently executing a pilot program, giving iPads to employees who travel or telework. The benefit of an iPad is that it is one-third of the cost of a government-issued laptop with Internet access. Before jumping into long term contract agreements, officials are holding out, waiting for competition to enter the market to see what they can offer instead, such as the BlackBerry PlayBook (Sternstein, Interior executes test iPads in bid to boost productivity, 2011).

4 Security

Not every program, system and service can be thrown into the cloud. With trade secrets, sensitive documents, identifiable information, and other restricted information that companies internally, and by law, need to protect, security in the cloud is absolutely paramount. These concerns are very real. In one instance in 2007, criminals were able to breach Salesforce.com's systems to steal personal information (Krebs, 2007). In an unrelated event, the VA immediately shut down the unauthorized use of a cloud application hosted on Yahoo! that contained patient data. The breach was discovered after VA IT security employees noticed that doctors and employees were updating a cloud calendar of information containing patient records. In addition, the doctors and staff were also sharing a password which had not been changed for three years. IT security staff promptly removed all patient data from the calendar, changed the password and blocked access to the service. Data from about 900 patients was exposed in the calendar (VA Shuts Down Cloud App After Breach, 2011).

Security is absolutely at the forefront of the obstacles to mass cloud adoption (World, Cloudy Forecast for IBMs Smarter Planet, 2009). Addressing security in the public cloud is of extreme importance for cloud service providers, as the government uses such cloud resources. Customers need to be assured that the level of security deployed in the cloud is at or above the level of security that the federal government mandates.

4.1 Government Standards and Policies

Many government-owned systems and applications contain sensitive data and other protected information. Government systems must pass Federal Information Security Management Act (FISMA) regulations to make certain they perform at or above a certain level of standard to ensure data is kept safe. There are other regulatory laws that protect data, such as Health Insurance Portability and Accountability Act (HIPAA) regulations within the healthcare system. Any agency wishing to migrate to the cloud must ensure that it conforms to any and all regulations when considering a cloud solution (Kundra, Federal Cloud Computing Strategy, 2011). With all new computing paradigms come new security threats. Since the federal government declared the adoption of the cloud, security is of absolute top priority. The government is addressing the issues of security on several fronts.

The Obama administration is requesting to allocate NIST a budget of $100 billion for developing needed policies and standards for cybersecurity and how they relate to cloud computing (Yasin, NIST budget request could more than double cybersecurity spending, 2011). The NIST's involvement in cloud computing is to standardize its secure and effective use within the government. The NIST has started the Cloud Standards Coordination Overview and Contributing Organizations wiki at cloud-standards.org. This wiki is a tool to be used for standards development (National Institute of Standards & Technology (NIST), 2011).

The Federal Risk and Authorization Management Program (FedRAMP) is an assessment process that will be implemented by the CIO Council. The program will analyze security compliance in cloud implementations. FedRAMP is a program to standardize the assessment of cloud solutions and to authorize them as safe to use within the government. FedRAMP will free agencies from performing their own compliance audits. Although FedRAMP will free the bulk of the burden from individual agencies, it is not the final answer to IT Security. Agencies are still ultimately responsible for IT security within their departments. After a FedRAMP approval, agencies will still need to determine if any additional security measures are needed that are appropriate and necessary for their projects. A prime example is browser security. If a browser is hijacked, an attacker can view keystrokes. This will allow them to capture account information to log into cloud environments (Joch, 2011).

Due to the dynamic nature of the cloud, more detailed status reporting is required. An executive order regarding the Federal Information Security

Management Act (FISMA) has been issued stating that supplemental real-time IT status reports are to be submitted to the Office of Management and Budget (OMB) in addition to the annual reports outlined in the original act. These reports now include all cloud service providers (Joch, 2011).

4.2 Cloud Security Alliance

The Cloud Security Alliance (CSA) is a non-profit group that consists of many industry-leading companies and their employees to create a standardized security framework for the cloud. They also strive to educate people through awareness campaigns to promote proper use of security in the cloud. They have a web-based certification program using the Security Guidance for Critical Areas of Focus in Cloud Computing catalog that they released in 2009. The certification is cloud–specific, with topics and content not covered in other certification programs (Jackson, 2011). Companies that are a part of the CSA include Microsoft, Verizon Wireless, Cisco, Dell, HP, Oracle, IronMountain and Lockheed Martin, to name a few (About the Cloud Security Alliance, 2010).

4.3 TrendMicro / VMware Alliance

Due to the expectation that the majority of virtualized servers will be less secure than their physical alternates by 2012, Trend Micro, a provider of anti-virus solutions, and VMware, which specializes in virtualization technologies, have teamed up to provide the first "agent free" anti-virus solution for virtualized datacenters. The standard approach for dealing with anti-virus solutions on virtualized systems has been to install endpoint solutions on each guest. Due to the "hypervisor" in a virtual system, virtual machine guests behave differently from file scans to network requests, compared to their physical counterparts (Trend Micro, 2011).

Treating virtual machine guests in this fashion leads to some key issues. Specifically, instant-on vulnerabilities are security risks when virtual guests are provisioned and de-provisioned rapidly. When these guests are quickly cycling on and off, it is impossible to provide them with a consistent level of security. Also, all guests that remain offline during a system-wide anti-virus update will be vulnerable the next time they are brought online. Only until after they receive the security updates will they be secure. Further, endpoint virus scans, when scheduled to run on guests, put a load on the overall system, thereby degrading performance. With these issues in mind, Trend Micro and VMware have determined that the most efficient way to bring security to a virtualized system is to incorporate it directly into the virtualized infrastructure, utilizing hypervisor introspection. By combining VMware's vShield Endpoint service and Trend Micro's Deep Security Anti-malware product, they argue that it is superior to software-based security services provided on a physical machine (Changing the Game for Anti-Virus in the Virtual Datacenter, 2010).

In addition to this alliance, VMware, has independently bridged the gap between private and public clouds with an update to their vSphere product line. vCloud Connector, a plug-in to vSphere will allow administrators the ability to send locally running VMs and their associated data out to remote cloud vendors to run on their platforms. Although the VM is sent out to a remote location, the administrator can still monitor these VMs in the same way that they monitor their local VMs. This allows an organization to spread out their system load among multiple providers. Currently, the first three providers of this technology are BlueLock, Colt, and Verizon, with more on the way (Schwartz, 2011).

4.4 Privacy Manager

To ensure privacy of sensitive data, Mowbray, et. al., have come up with a client-based "Privacy Manager" that obfuscates and de-obfuscates data. This obfuscation mechanism is driven by a user-created key that is not shared with the cloud service provider. What is sent into the cloud cannot be de-obfuscated until it is returned to the customer. Depending on what and how much obfuscation is used, overall functionality of the cloud services may be limited. This privacy manager will also notify the customer when and how information about them is being used within the cloud as it happens. Through preference settings and different “personas" a customer can take on, the level of privacy a customer can achieve is very flexible (Mowbray & Pearson, 2009).

5 Mass Adoption of the Cloud

Although the cloud is quite inviting, as it offers many incentives for both the cloud service provider and the cloud user, there are improvements that must be made to the overall cloud environment in order for a mass adoption of the cloud to happen.

5.1 Mobile Device Compatibility

Currently, the mobile device industry is a raging market with no end in sight (Vance, Mobile Cloud Computing: Is Your Phone Drifting to the Cloud?, 2009). As the number of mobile device users continues to increase, the demand to use such devices within the government will also rise. In order to efficiently use mobile devices as a replacement for desktop computers, it is paramount that cloud services are provided so that these devices can off-load their processing responsibilities to the cloud for quick results, and to save on battery life. The climbing sales of mobile devices, as well as the push for cloud services within the government, will, at some point, force government IT to adopt mobile users into their security policies.

5.2 Virtualization

As stated earlier, virtualization plays an important role in the success of cloud computing. Virtualization, unfortunately, has a drawback that might inhibit a large target group of customers. The current state of the cloud does not offer many options for specific system types, nor do they guarantee identical system types for every job run. The bigger problem with virtualization is that it prevents users from writing optimized code, especially in the area of capability and capacity computing (Thomas Sterling, 2009). Until more work is done in this area to give users more control over the systems they rent in the cloud, either through better SLA agreements or through metadata, these customers will most likely steer clear of the cloud.

5.3 Costs

Cost savings are attractive to everyone. With the explosion of the amounts of data being collected, it is important that the transmission of this data into the cloud is efficient, incurring minimal cost to the customer. Since networking hardware such as routers are vital to the cloud, they need to be high quality which means that they are very expensive. Networking hardware is so expensive, it makes up the majority of networking costs (Armbrust, et al., 2009). It is important that these costs are not pushed down to the customer.

5.4 Trust

The hesitation surrounding the adoption of the cloud boils down to the issue of trust. The government is not going to risk sensitive data with just any provider. Within the federal government, trust is evaluated essentially through tests. Once cloud services or vendors have passed these tests, they are generally trusted throughout the government. For example, in order for a cloud service provider to be able to provide services on apps.gov, these companies must go through a certification process. These companies must show that their services meet the NIST guidelines. This filtering process is done on behalf of federal agencies wishing to use cloud services from apps.gov. Of course, as mentioned earlier, although services by certain vendors may be available on the apps.gov website, the agency using these services is ultimately responsible for security. The filtering process on apps.gov gives these agencies a good starting point for companies they should consider trusting as they migrate to the cloud. The FedRAMP program is another government initiative designed to authorize the use of certain cloud services and providers among government agencies. Any cloud service or provider that passes the FedRAMP assessment can immediately be trusted, given that it met government standards for FedRAMP approval.

Harris Corporation sees trust as the primary inhibitor to government cloud adoption. Harris Corp. has developed their Harris Trusted Enterprise Cloud. Harris Corp. has had a long-standing relationship with the government, providing them with tech services for many years. With their known reputation, and their efforts to "set the bar high" with their security policies within their environment, they believe that they will be able to woo many hesitant government agencies into the cloud (Yasin, Harris Alliance Targets Government Thirst for Cloud, 2011).

Service providers like Salesforce believe trust should come from setting standards on how to deliver their cloud services to a customer. Salesforce believes in a seven-faceted approach to deliver their services. The first and, arguably, most important standard of their model is complete, end to end, "world-class" security. The security they provide with their services includes physical, network, application, and internal systems security. They also provide a secure strategy to perform data-backup, third party certification of their infrastructure and secure policies for internal operations (Salesforce, 2011).

Transparency comes next. Salesforce believes in giving their customer full disclosure regarding their service performance. This data can be monitored in real-time. In addition, they also provide their customers with availability and transaction performance metrics. Further, Salesforce even provides their customers with maintenance notifications. This data can be viewed on their website at http://trust.salesforce.com (Salesforce, 2011).

Delivering true multi-tenant cloud services is another standard Salesforce set for itself. Single-tenant architectures do not allow for scalability and are not "cloud friendly". Providing multi-tenant services allows Salesforce to provide customers with a highly scalable system where each customer uses the same code base. This allows for performance gains over a single-tenant system. Multi-tenant systems allow Salesforce to manage that single code base for all users, allowing them to effectively manage updates to the system (Salesforce, 2011).

Salesforce believes in a "proof in numbers" approach to indicate how trustworthy a service provider is. Salesforce believes that if a provider has a high number of customers that is generally increasing, that company would be relatively trustworthy. The number of subscribed customers not only indicates a trust level, it also indicates a certain level of proof that the provider's systems are reliable (Salesforce, 2011).

For a service provider to establish themselves as reliable and trustworthy, these companies must provide customers with what they want. Customers overall want maximum performance. In order to provide customers with fast and reliable service, providers must invest in the best equipment in order to meet the demands of the customer. Also, these providers must invest in multiple "substations" situated globally so that customers worldwide have the fastest, most efficient and reliable service possible (Salesforce, 2011).

When dealing with sensitive data, backup services and policies must be implemented to prevent data loss in the event of a disaster. With information that customers inject into the cloud, service providers must ensure that customer data is safe. Salesforce believes the proper way to implement data recovery is to ensure that all data is replicated and dispersed among multiple data centers throughout different parts of the world. Specifically in Salesforce's case, it provides their customers with a 1-to-1 data recovery service in the event of a regional disaster. Their service replicates data in near-real time at the disk level to ensure minimal data loss (Salesforce, 2011).

Last on Salesforce's list of standards for a cloud service provider is to ensure high availability of resources so that customers can use the services they pay for when they want to use them. Service providers must be able to provide a highly reliable power source to their equipment, as well as efficient cooling to keep systems stable and online for use. Having a highly available and robust networking infrastructure enables efficient delivery of services to customers when they request them (Salesforce, 2011).

Retaining trust and confidence in users is now top priority for cloud vendors. On April 21, 2011, Amazon suffered a partial outage in their cloud services. This outage was due to a network configuration error during an upgrade, sending traffic to the wrong network. This error caused many large company websites such as Quora, Foursquare and Reddit to go offline. Amazon's outage lasted for five days before they announced a full recovery (Thibodeau, Amazon Cloud Outage Was Triggered by Network Configuration Error, 2011). Although Amazon was able to recover, this event has left many current subscribers furious. Prospective customers are now rethinking their cloud ambitions. Many are now doubtful about the cloud being superior to corporate IT alternatives. Amazon and other cloud vendors will now have to try and to counter this major blow to the cloud computing industry. Many are now convinced that the cloud is not yet mature enough for businesses. It is speculated that this event will seriously hinder cloud adoption rates. Paul Haugan, CTO of Lynnwood, Washington is scrapping his plans of subscribing to Amazon's cloud services. Amazon's outage convinced him that the cloud is not ready to support his needs (Thibodeau, Amazon Outage Sparks Frustration, 2011).

Although this was a serious event for Amazon, it needs to be kept in perspective. Even though many were affected by the outage, there were a greater number of customers who were not. There is no doubt that Amazon is responsible for the outage. Relying on the cloud vendor for backup and recovery services should only go so far. What customers should learn from this event is that it is essential to have a robust disaster recovery plan separate from the cloud vendor (Thibodeau, Amazon Cloud Outage Was Triggered by Network Configuration Error, 2011).

6 Outlook/Conclusion

Cloud computing has emerged from several different technologies. From seeking better return on investment (ROI), virtual machine technology, distributed computing capabilities, scalability characteristics of clusters, utility based pricing policy, and the desperate need for low cost systems that can process terabytes to petabytes of data, cloud computing comes to life.

The paradigm shifts in business that cloud computing is responsible for are here to stay. The cloud has changed how businesses small and large are run and managed. Cloud computing is also beginning to impact the federal government and how they execute business. Due to the high costs incurred to keep thousands of data centers operational, and the demand for highly efficient and scalable resources, the government is realizing cloud computing is the most cost-effective solution that will provide scalable compute and storage resources. In many ways, cloud computing has become an integral part to a business' success and will become essential in many government operations.

Cloud computing is still very much a developing technology. Like any emerging technology, there will be hiccups, much like what Amazon experienced in late April, 2011. These hiccups shouldn't prevent the cloud from developing further, but should make customers understand that there are risks that need to be considered. Every day, cloud service providers are further developing and improving their cloud technologies. Private development built on top of a vendor's infrastructure is also shaping the cloud landscape. The cloud environment is constantly changing and will continue to grow at a progressive rate as the number of cloud users and developers increase.

The expectation for the future is that cloud computing will evolve to become a vast collection of services that communicate securely, having the ability to jump across various vendor's infrastructures, being able to scale and distribute tasks as needed in a highly efficient manner at minimum cost to the customer (Murray, 2009). This will gain momentum as the number of mobile devices subscribing to the cloud start to emerge in greater force. Over time, cloud computing will mature to a point that will ease concerns and establish a certain level of trust. This sense of trust will, in turn, cause a mass adoption of the cloud. In the meantime, within the government, assessment and authorization programs such as FedRAMP will provide government agencies a first-round filtration of service providers to determine which vendors will be able to meet the government's security needs.

7 Best Practice How-to: Moving Applications to the Cloud

Coming up with reasons to migrate an enterprise application to the cloud is easy. However are these reasons justified within the organization from a business perspective? After justifying the project, what is the next step? Executing the migration process is not necessarily straightforward. There are potentially many caveats which could end up making such a project very expensive and labor intensive, resulting in a disaster. It is important to understand all of the steps of the migration process in order to ensure a smooth migration that meets its projected target and budget.

7.1 Feasibility

7.1.1 Applications Best Suited for the Cloud

Understanding what types of applications to migrate to the cloud is important. Migrating applications to the cloud that cannot benefit from the cloud's characteristics can be a waste of time, effort and investment. Applications that are resource-intensive that perform highly complex mathematical computations are perfect candidates for the cloud. Moving these applications to the cloud would free up local resources for other tasks. From an architecture standpoint, multi-tiered applications are perfect candidates for the cloud. Due to their inherent design, their components are naturally decoupled from one another, allowing for each different component to be migrated at different times. Further, if some of those components are business-critical, requiring them to stay within the company firewall, a multi-tiered design makes it easier to pick and choose which pieces should be migrated, while holding on to those that should stay on-site (Cisco, 2010).

Other types of applications that have the ability of "scaling out", or those that benefit from additional compute resources, would also be a great fit in the cloud. These types of applications can easily benefit from the utility pricing policy of cloud computing when they need more compute resources to complete their tasks (Cisco, 2010). Additional types of applications best suited for cloud migration include those applications with low utilization rates, varying demands for compute resources, the need for progressively more resources over time, the need to be deployed globally, and applications that need testing on various different platforms (Cloud in Steps, 2009).

7.1.2 Business Cases and Market Drivers

Migrating an application to the cloud would increase the performance of the application, but how would the migration benefit the company? Would it increase customer satisfaction? Would it decrease operational costs? There are an abundant number of market drivers to justify a cloud migration. These market drivers can include the prospect of easy upgrades. This benefit of the cloud allows its customers to access the newest technologies and versions of software available at no additional cost. Another market driver could be to reduce the internal labor force at a company, saving on operational costs. A final example of a market driver would be to transfer the responsibilities of data management to the cloud vendor as this task becomes increasingly cumbersome every day (Cloud in Steps, 2009). If a business case cannot be developed in order to justify a particular migration, then further work on the migration project should not continue (Linthicum, 2010).

7.1.3 The Hype Factor

In addition to market drivers, some companies might push the move to the cloud from within, based on concerns of being "left behind" in technology trends. The technology hype cycle represents the typical rise and fall, and subsequent rebound, of a new technology trend. Understanding the hype cycle will allow for a better analysis regarding when, if at all, a company will adopt a new technology. Exactly when a company adopts technology, in this case the cloud, ultimately depends on how conservative they are in terms of risk aversion and their desire to be on the cutting edge of technology.

The cycle has five stages, as in **Figure 7**. The Trigger stage is when a new technology is unveiled with extraordinary claims from proof-of-concept designs. At this stage, these claims are unproven. After this stage is the Peak of Inflated Expectations, which announces many success stories. Early adopters start to emerge, but the majority stay far away. The Trough of Disillusionment is the third stage. This is when technologies fail to meet initial claims. Media coverage typically subsides at this point as well. This is a sink or swim moment for companies providing the technology. They must either improve upon the technology or face failure. The fourth stage, the Slope of Enlightenment, happens when an increased number of uses of the technology, as well as success stories, emerge. Subsequent "generations" of the technology are developed with an increasing number of adopters. The most conservative businesses remain unconvinced. Mainstream adoption of the technology begins at the final stage, called the Plateau of Productivity. Both the technology and the vendors have matured such that the products or services offered have extensive applicability in diverse markets and industries. Those who adopt at this stage do so due to an abundant amount of success stories (Gartner).

7.1.4 Cost Analysis/Financial Assessment

If a business case can be made to justify a migration project, the next step is to perform a cost analysis on the project. Performing a cost analysis by looking at True Cost of Ownership (TCO) and ROI will enable key decision makers to understand how much of a profit or loss such a project would generate. Items to look at when conducting such a cost analysis should include capital expenditures such as physical servers, storage, networking devices and real estate.

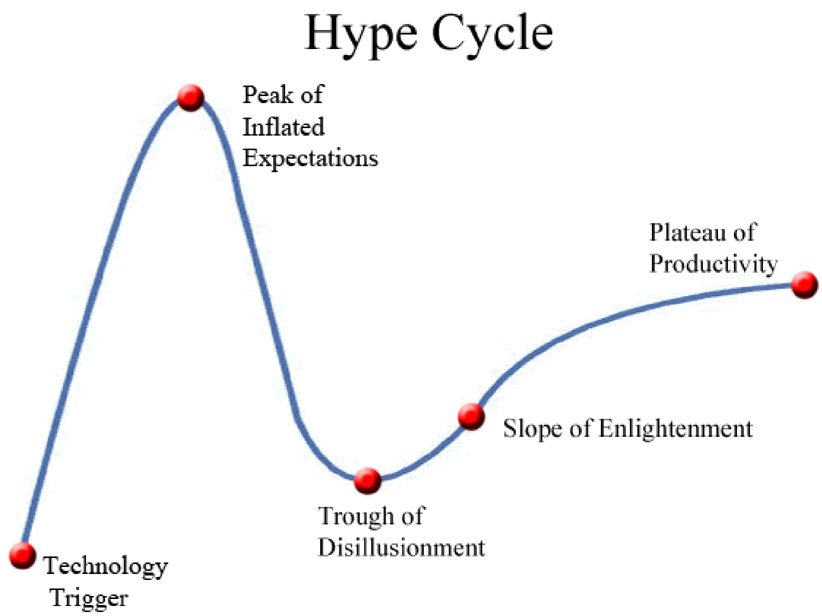

Figure 7: Gartner's Technology Hype Cycle

Operational expenditures, which comprise energy consumption, personnel and staff, raw computing power, bandwidth, support and maintenance should be included. Additionally, overhead costs should be added to the cost analysis. These overhead costs should include all costs that will be allocated during the actual migration, including expert support, migration services from the cloud vendor or third party, and costs concerning governance over the actual migration (SYS-CON Media Inc, 2008). After performing the cost analysis, if the costs of a cloud solution are less than the current system then cloud migration is still a viable option. However, if the current system costs less than a cloud alternative, a move to the cloud might not be the best decision for the company, at least for the time being; especially if current customers are happy with the in-house system (Holland, 2011).

7.2 Understand What the Company Already Has

7.2.1 Business Operation Analysis

After assessing the cost of cloud migration, it is necessary to peer into the business to see how it operates, to ensure there are not any roadblocks in the way. A good strategy is to envision an overall roadmap of the individual steps that make up the different business processes within the company. Each step should include a set of inputs and outputs representing data flow. At this level of resolution, it will make it easy to decide which of these steps are appropriate for cloud migration and which are not. These decisions are based on the data dependencies and decision points of each step (Five Tips for Moving to the Cloud, 2011). Further, what should be the order of the migration? The priority should be dependent on the company's needs. The order in which to migrate should be one where the amount of downtime of services is minimal (Cisco, 2010). For companies with multiple systems they are considering to migrate, an assessment should be made to determine which applications have the lowest utilization rates. These applications are the best candidates to migrate to the cloud, bringing higher ROI to the company through increased utilization (Jong).

For third party software, take extra care when migrating this licensed software to the cloud, as it might violate contract agreements. Straighten out all licensing issues prior to moving forward. There are different options that can be taken. The Bring Your Own License (BYOL) path is a traditional approach to activating software. The license would be granted by the vendor and the customer would activate the product residing in the cloud. Another option is the utility pricing model, where the customer can pay for the product based on usage. The last method is to use the SaaS version of the product and be charged a subscription fee (Varia, 2010).

While assessing which areas of the business process should be migrated, it is important to decrease risk as much as possible when migrating. One issue to consider before proceeding is to contemplate how to retrieve all data injected into a particular cloud in the event of unforeseen circumstances such as the vendor significantly raising their rates, or simply because the company no longer wishes to continue cloud endeavors. Ensure there is a solution in place to back up the data in the cloud in case of such an emergency (Varia, 2010). Within certain industries such as healthcare, policies might prohibit company data from being stored on multi-tenant or shared systems. Even the company itself might impose certain restrictions on how and where data such as intellectual property can be stored. Also, if there are regulations indicating specifically where, in terms of geography, that data can be stored, there needs to be a careful inspection into cloud vendors and how they store customer data. It is important to understand that not all types of data belong in the cloud (Jong). Understanding the level of security an application requires is paramount to reducing risk. The cloud vendor candidate must be able to provide, at the very least, the same level of security that the application implements outside of the cloud (Varia, 2010).

Another very important consideration to keep in mind is how reliant the application is on network resources. If there are strict bandwidth and latency requirements for a particular application, migrating to the public cloud might be a difficult task. It is important to analyze the network performance of cloud vendor candidates. Inspection of the service level agreements (SLA) a vendor can provide in terms of network availability should also be performed (Varia, 2010).

One last consideration should be to determine if any applications require specialized hardware. These particular applications would not fare well in the cloud given that such guarantees of specific hardware may not be met. Check with cloud vendors to see what they can provide in terms of SLA agreements to satisfy the requirements of the application (Jong).

7.2.2 Create Profiles

Once the best application candidates have been selected, a baseline performance profile should be conducted on each of these applications. This profile should be conducted for at least ten to fifteen days per application. Data to capture during this time should include CPU and memory utilization and throughput, storage throughput, latency, and input/output operations per second (IOPS). In addition, network data should also be collected. This data should include network throughput, total number of dropped connections and connections per second. Performing this performance profile will provide a snapshot of how an application behaves and utilizes resources. This is an important step, since cloud vendors have varying sizes of virtual systems to host applications. Selecting the most appropriately-sized system and number of instances will ensure optimal utilization and cost savings (Cisco, 2010).

7.2.3 Define a Migration Strategy

It is necessary at this point to start to develop a standard method of approach to migrate the applications to the cloud. Coming up with such a strategy will allow for a smooth transition and should be applied on a per-application basis. There are three phases to such a strategy: the Enterprise, Operations, and Implementation phases. At the Enterprise phase, employees and departments that rely on the specific application and its data must be notified about the migration. A plan should be developed for these groups so that their daily tasks are not disrupted during the migration. Expert enterprise IT staff should be called upon to uncover these situations and help mitigate such issues (SYS-CON Media Inc, 2008).

The IT staff should also be called upon at the Operations phase of migration planning. The IT staff is needed at this point to deal with the issues of current SLA contracts with current vendors to uncover any issues with a migration. Further, if the application is outsourced to a third party vendor, this company should be included in the migration planning to facilitate a smoother process (SYS-CON Media Inc, 2008).

At the Implementation phase, many issues revolve around transferring the potentially massive amounts of data associated with an application. How will the data be transferred? Specifically, how will the data be stored? A very important aspect that needs to be understood about the cloud is that, in general, data is not stored in the traditional sense of a relational database. Cloud vendors primarily have two types of data storage, BLOB, or Binary Large Object, and non-relational database storage. For cloud compatibility, plan to perform a data model rewrite headed by top IT and software engineers. Vendors do support other types of storage, including relational database tables. These services, however, come at a premium. If the financial budget allows for such services, look into what they offer to see if a data rewrite is actually necessary.

Other questions at this stage should include (Five Tips for Moving to the Cloud, 2011):

- Will the data be transferred all at once?
- Will it be transferred in parts?
- What mechanism is in place to validate the transferred data?
- How will the data be secured during transport/transmission? It is important to plan this process out well in advance to ensure data accuracy and to prevent data loss (SYS-CON Media Inc, 2008). Also, SLA agreements are often overlooked during the actual migration phase.
- What guarantees can the vendor make while the migration phase is underway?
- Who are the people that should be hailed in the event of a problem?
- Can there be guarantees set on the response time of a claim? SLAs are absolutely important during this time for budget, accountability and scheduling reasons.

7.3 Select a Provider

7.3.1 What to Look For In a Vendor

One of the most difficult challenges when dealing with a migration project is actually selecting which cloud vendor to use. It is important to research as many cloud vendors as possible and compare what they offer against the application profiling data that has been collected from a previous step to see if their services can meet the demands of the application.

To make the selection process easier, there are key qualities that a vendor should possess to make them stand out from the competition. Looking for these qualities will allow for an easy filtering process. As an example for cloud vendors supporting PaaS solutions, these vendors should provide SLA agreements for performance and availability. Also, vendors should make guidelines available for maintenance and management of the service, including documentation regarding compatibility for APIs that the platform and the application are to use. For storage services, the vendor should provide some common format that allows its customers to extract/export their data so that they can migrate it elsewhere if they so choose (Cisco, 2010). These qualities are included in the cloud portfolio of all major cloud vendors that have a proven track record, such as Amazon, Google, Salesforce and IBM.

Regarding security, it is critical to realize that multiple customers can be using the same physical hardware concurrently. This creates a situation where vulnerabilities in one customer's system might compromise other customers' systems by breaking into their partitioned virtual machine instance. The vendor should have a security infrastructure in place to prevent such a threat, but customers should inspect for themselves whether or not the safeguards are adequate. Also regarding security, PaaS services will undoubtedly have user account management tools. It is important to identify the similarities and differences of the cloud vendor's management process when compared to the customer's own directory service management tools and processes (Cisco, 2010).

Lastly, the PaaS vendor should provide its customers with monitoring and management tools to control their applications in order to fine tune its performance on the cloud. These tools should come with sufficient documentation to understand how to use them. These tools are especially important when dealing with a system that would be deployed in the cloud across different geographic locations. Access to these different locations needs to be centrally monitored and controlled by the customer for optimization reasons. If the vendor allows for dynamic scaling, they must provide documentation on exactly how the application will scale up and down, as well as how the application or underlying PaaS system will resolve resource contention issues (Cisco, 2010).

7.3.2 Test Drive Vendor's Services

After considering vendor candidates and creating a list of top choices, the next step is to give those vendors' services a test drive. It is essential to become familiar with a vendor's offerings to confirm that a vendor's claims about its services are accurate and that it lives up to expectations. Expert knowledge about a vendor's system is not necessary. Assign a small group of experienced software engineers to become familiar with their APIs, tools, software development kits (SDK's), and security technologies. Once familiar, have them create either a demo piece of software on their cloud that approximates the application to be migrated in terms of major functionalities, or have them migrate an existing application to these vendors' clouds. The proof-of-concept route will provide an approximation of how the real application will behave in the cloud. Using either the proof of concept or migrated application method, comparing performance metrics gathered from vendor candidates against the baseline set of data will determine which cloud vendor can provide the best solution in terms of performance and efficiency. While making the decision of which vendor to choose, consider the amount of potential time and effort it will take to turn the proof-of-concept into a fully functional piece of production software (Varia, 2010).

7.3.3 Ask Lots of Questions - Don't Assume

There are many questions that should be asked before subscribing to a specific vendor. These questions include:

- How much time and effort needs to be put into turning the current application into the migrated cloud version?
- Does the vendor charge extra for scaling services?
- What types of storage will be required for the cloud application? There are many different options for data storage, including raw block storage, and non-relational and relational databases in both licensed and open source formats.
- Does the vendor have the same database services that are currently being used in the in-house version? If not, consideration needs to be given to calculate how difficult or easy it will be to convert to a different service (Burke, 2011).

- With the storage options that are chosen, how will this affect total cost, availability, and the query-ability of this data?
- What will the performance be like?
- Is there a cache that can be utilized to mitigate performance issues? Not all storage options are alike. They all have their pros and cons. It is important to consider each option and how it will affect the overall system performance.
- How long will it take to migrate all of the application data into the vendor's storage service(s) (Varia, 2010)?
- What are the specific data transfer rates?
- Does the vendor have an alternative shipping method where hard drives are shipped directly to the vendor to expedite the process?
- What about data archiving?
- Should an off-site archival service be set up to transfer old data out of the cloud?
- What is the cost of transferring data out of the cloud?
- What is the vendor's process to update its cloud infrastructure software and how will that affect the performance and availability of deployed applications?

These are only some of the many important questions that need to be brought up when deciding what cloud vendor to choose (Holland, 2011)

7.4 Migrate

At this point, a cloud vendor should be selected. The task at hand now is how exactly to migrate the application. This is a delicate issue, since some or all systems being migrated might experience downtime, depending on execution. It is critical to have a game plan in mind to minimize problems.

Prior to migrating anything over to the cloud, an important first step is to make sure that the cloud environment itself is properly "staged", or set up, to accept the application. If this staged environment is a similar platform to the one the in-house application is running on, recompiling the application won't be necessary (Varia, 2010).

There are two general paths that can be taken to actually migrate the application. They are the Forklift and Hybrid migration strategies. The Forklift method is designed for tightly coupled applications. Examples of tightly coupled applications include backup and storage, or archiving systems that can be viewed as an individual component. Essentially the Forklift method is to move the entire application into the cloud all at once (Varia, 2010).

The Hybrid method is great for loosely coupled applications. This method allows parts of the application to be moved to the cloud, leaving the rest behind while maintaining service to all components. This is a low-risk method to migrate an application to the cloud, since this method allows the application to be moved in pieces. Pieces can be moved one at a time and optimized before continuing. It also eliminates the potential of unexpected behaviors if the entire system is picked up and moved into the cloud. It is much easier to isolate a problem when migrating in small increments rather than all at once (Varia, 2010).

7.5 Test Locally

After migrating the application, or components of the application, to the cloud, it is now time to perform rather extensive testing to ensure everything was migrated successfully. The types of tests that should be performed include functional and performance testing of both the application and cloud services such as scaling, as well as security and penetration testing. This process of testing is much the same that a company would do when deploying an in-house application. Of course, the difference here is that what is to be deployed is being hosted remotely. This means that the customer does not have absolute control over the systems as they would in a traditional in-house testing effort (Linthicum, 2010).

Specifically, if the application will scale up and down, it is important to test the scripts that enable such functionality. How quickly do the scripts respond to changing demands? Do they provision/de-provision the right number of instances?

The application itself should go through functional, stress and load tests to make sure it functions properly. Negative tests should also be carried out to test for network outages and disk capacity issues. Finally, security tests should be conducted to ensure the system is rock solid when the application is deployed as a production piece of software (Cloud in Steps, 2009). While in this testing effort, a limited number of "test group" users should be allowed to use the system in order to evaluate its functionality. These test users should be individuals that rely on the system to carry out their daily tasks. These users are essentially experts of the system and would know best what functionalities the system needs. If there are any functionalities missing, now is the point to uncover them and incorporate them into the application before deploying it as a production piece of software (Five Tips for Moving to the Cloud, 2011).

To make the testing effort easier, many larger vendors offer cloud emulation and migration modeling software. As an example, Microsoft includes an Azure cloud emulator that they include with their Azure toolset within their Visual Studio software development integrated development environment (IDE). Not only is such an emulator easier to work with than having to import/export data to the actual cloud, it also allows for the testing effort to be conducted for significantly less money, since cloud resources aren't actually used. These emulators can provide customers with vital information about their cloud application and the underlying system(s), including performance metrics, latency of provisioning more instances, failures and error states, and security, capacities and compatibility (Yunus, 2010).

It is recommended to perform two types of emulations. The first emulation should be run with local test data, and the other with data stored in the cloud. The first emulation will check to make sure the application itself runs without bugs and errors prior to running it on the cloud. The second emulation will test to make sure the connection strings for communication are configured properly.

Performance metrics taken from such an emulator need to be carefully appraised. The reason for this is that these emulators have difficulty inhibiting local machine processes that interfere with the emulator. Due to this, performance metrics gathered from the emulator will under-rate the application's performance.

In addition, the physical hardware that the emulator runs on also affects the performance of the emulation. For instance, if the local machine was a single core system, an emulation using ten nodes will run ten times faster in the cloud.

7.6 Deploy

When switching to a production system, it is recommended to migrate users in groups. This will make it easier to target any problems that might arise from switching user accounts over to the cloud. Migrating in batches will also allow for decision makers to determine if cloud capacity should be increased and to what degree. The batch transition will help minimize system load and make migration management easier (Cisco, 2010).

After user accounts have been moved over, the task at hand is to see what other services the cloud vendor provides that can be utilized to increase performance and efficiency of the application. For instance, if the application did not have elasticity and scaling capabilities prior to the migration, research should be conducted to determine how to implement such technologies into the new application. Other objectives at this stage include double-checking that security is hardened as much as possible. Is the cloud account information secured? Are multiple forms of authentication being used? What different types of user groups are set up for access restrictions? Is data encrypted both in transmission, as well as when the data is stored? What is the recovery strategy? How frequent are the backups taking place? Have the backups been tested to ensure that they will work if they are needed? (Varia, 2010). As part of any contingency plan in the event of a disaster, a service should be activated to notify users of brief down times, as well as systems to come online to capture data trying to be sent to the application. This system will prevent data loss and allow that data to be sent to the application at a later date when it is back online (Five Tips for Moving to the Cloud, 2011).

7.7 Go Live

After going "live", consistent on-going monitoring efforts should be performed. This monitoring is to ensure that the system remains stable, and to make further optimizations that will drive expenses down. Keeping the system stable includes keeping the system up to date with the latest security patches. As far as reducing expenses, it is necessary to understand system usage patterns. If there are particular seasons or geographic locations that don't generate much traffic, scaling back resources by shutting down idle instances will drive utilization rates up and minimize costs. Another way to save on costs is to compress files that are transferred into and out of the cloud, saving on bandwidth fees. Ongoing monitoring of cloud services should also include reviewing usage logs. These logs will be able to determine if there has been unauthorized access to part of the system, as well as serve as a way to prevent over billing of services (Varia, 2010).

Monitoring should also include keeping tabs on the vendor itself. With all of the data collected during the research phase of finding the right vendor to subscribe to, items contained in that data should include all expectations and assumptions of the cloud vendor pertaining to business contracts and responsibilities. This compiled document should be reviewed after six months and one year. Comparing this document with the current state of affairs will reveal whether or not the vendor has lived up to, exceeded, or failed to meet expectations (Five Tips for Moving to the Cloud, 2011).

7.8 Observations

When considering cloud migration, understand that advertisements and claims released by cloud vendors can be misleading. When fully engaged in a cloud migration project, there could very well be times where things don't work the way they are intended, leading to delays and budget increases. Learning from others' experiences with cloud projects is the best way to prevent these types of issues.

One of the more important concepts to keep in mind is that cloud computing is still an emerging technology. With the promise of large profits to be made in the industry, vendors are constantly modifying and changing their technologies. When a company decides to move to the cloud, they must be vigilant with regards to the changing technologies in the cloud. If attention is not given, migration projects could veer off course, since assumptions and expectations that were made early on in the process might not be applicable down the road.

Many urge caution when considering cloud migration. Although it is agreed that businesses do become excited about the cloud and how it can benefit operations, they warn that once it is understood how much effort would be involved in the transition, excitement drops. They also state that enthusiasm is lost when hidden costs such as the risk of losing intellectual property, migration, delays, and provider overhead costs are added into the financial analysis (Dreaming On a Cloud, 2009).

There is even greater concern over existing applications, especially legacy applications, and their compatibility with the cloud. With the expectation for these applications being highly coupled within other corporate systems, intersystem communication and compatibility problems might arise when migrating to the cloud. "Existing applications will have a very tough time in the cloud, considering they were never meant to run free in such a dynamic, open environment." (Cole, Enterprise Applications and the Cloud, 2009). Even Lew Tucker, a CIO with Sun Microsystems, has stated that migrating these applications over to the cloud will be a huge undertaking, as they were not originally developed with open source and SaaS-oriented technologies in mind (Cole, Migration On the Cloud, 2009). Regardless of legacy applications, even current enterprise applications will be faced with significant challenges. "Most applications ported to cloud platforms will need a great deal of rework, testing and redeployment... Most cloud applications will be new applications." (Golden, 2009).

If a company is trying to save on costs by using internal staff to migrate rather than calling upon typically expensive expert services, they might realize too late that they should have called upon the experts from the start. Getting the existing IT staff up to speed on cloud architectures to learn how to implement and monitor these technologies can be difficult. Software developers and other IT staff will need to be trained to redesign existing applications and testing suites. Also, operations personnel will need to be trained to manage systems that are remote to the company. The costs for training and the division of time away from performing their daily tasks might not make this a wise decision (Dreaming On a Cloud, 2009).

From a security standpoint, cloud providers are able to provide most companies with the level of security that they require. The only issue that needs to be kept in mind is that, in public cloud environments, multi-tenant systems are commonplace. This means that applications from multiple customers could run on the same physical system. Consideration needs to be given to the fact that a company can be at risk due to vulnerabilities in some other customer's code that resides on the same system.

7.8.1 Final Thoughts

At this point in time, each vendor's cloud environment is vastly different from another. These differences can include the cost of services, specific capacities of the services, extent of technical support, and migration services. When considering different vendors, be sure to create a spreadsheet or other file that can easily display key differences between each company to make the decision process easier (Neubarth, 2009).

There is no "easy route" to migrate to the cloud. Once in the cloud, it can potentially be just as challenging to migrate away from one vendor to another. For inter-vendor migration, the total amount of effort depends on how different the vendors' environments are. For example, if an application initially resided on Amazon's EC2 IaaS service, it would be relatively easy to transfer it to GoGrid's service, as they are similar architectures. Problems arise when the architectures are dissimilar, spawning new development and testing efforts. Regardless of what type of migration, however, expert advice should be called upon. Keep in mind, however, that expert advice can become expensive (Ivan, 2010).

7.8.2 Typical Results of Cloud Migration

As an example, Eli Lilly's operations were being inhibited by its traditional in-house infrastructure that was constantly expanding. It turned to Google's cloud and saw immediate improvements. Provisioning a new server in Google's cloud now only takes it three minutes, rather than almost eight weeks. Provisioning new collaboration environments saw the same improvement gains. Also, bringing a large Linux cluster online normally took it 12 weeks to provision. With help from Google, it now only takes five minutes (Cearley & Phifer, 2010).

JohnsonDiversey, like other companies, started to realize that their in-house legacy systems were inadequate and their collaborative infrastructure was less than satisfactory. Also, JohnsonDiversey suffered from storage capacity limitations, an issue plaguing industries worldwide. Its solution was to use a multi-faceted approach to cloud computing. By adopting Google apps and Oracle CRM On Demand via Salesforce, it was able to realize great efficiency gains. With help from Google Docs, it can now deploy applications over a weekend, rather than three to four months. JohnsonDiversey has also realized a 20% reduction in bandwidth associated with collaborative functions. By reducing its operating cost of email and collaboration infrastructure by 70%, it was able to see a 100% return on investment in as little as 14 months (Cearley & Phifer, 2010).

8 Case Study: Migrating an Application to the Cloud

This section of the report provides a detailed account of the DACS author's first time efforts to migrate a scientific application to the cloud.

During the initial planning phase of our cloud computing case study, it was decided to focus on providing information to the DACS community that was unique. Having read many reports on cloud computing, a lack of a detailed account of the actual process of migrating a real-world application to the cloud was noticed. The initial thought was that if a suitable application could be found to migrate and port to the cloud, it would provide valuable insight to the reader. Most, if not all, of the cloud case studies found to date dealt with the topic of vendor sales or literature/product descriptions and did not contain any detail of the issues and complications encountered during the process.

Many years of experience in software design, programming, and deployment have taught that caution should be exercised in believing how straightforward a new system will be to implement. Therefore, taking an existing application and moving it to the cloud would provide an excellent first-hand view of the typical migration process. This independent analysis gives our cloud computing report the "value-added" needed to make it stand out as a useful and pragmatic resource for IT decision makers and software engineers.

8.1 Selecting the Application

Finding an existing application that would benefit from the power of the cloud architecture was the next step. A few weeks earlier an associate mentioned the need for a high performance computing resource for a project being worked on. A spreadsheet application was being used to perform Monte Carlo simulations and was reaching its limit in the number of simulations that could feasibly be run and the time required to complete them. At that time, the simulations were taking over a weekend to produce results on a set of 30,000 simulations and seemed to limit out at that number.

This application used Microsoft Excel with Visual Basic for Applications (VBA) code to execute the algorithms used in the scientific model and to randomize the variables used in the Monte Carlo simulations. It was very compute-intensive, not I/O bound and thus, as discussed in Section 7, an excellent candidate for migration to the cloud.

8.2 Choosing the Cloud Platform

Using a spreadsheet program for this application eventually proved an awkward and poor fit for the needs of the project. It was not very configurable, the dataset was embedded, the variables and formulas were scattered on one worksheet, the results displayed back on other sheets, and sorting was a separate user-initiated action. Even if it were possible to push the existing Excel spreadsheet up to the cloud and have it execute via a virtual machine environment and then access it via a user interface, it was not the ideal solution. A standalone application migrated to a platform-as-a-service (PaaS) appeared to be a much more logical and robust approach. This approach would give the most flexibility and capability to adapt to further iterations of the associate's (the customer, in a sense) project. If the model changed or the need for separate models with different algorithms and business logic arose, it would be simple to respond by adding a new module. An application where the user interface, data sets, and domain logic were separated and loosely coupled would be a better approach. The decision was made to rewrite the spreadsheet as a standalone Microsoft Windows application and port to the PaaS cloud.

8.3 Selecting the Vendor

This step and the planning of the Excel spreadsheet rewrite to a standalone application went hand in hand. Driving factors that were immediately apparent:

- The existing algorithms were written in VBA. Customer projects can vary greatly and the complex math and science underlying them can be daunting. For this and future projects, it was important to be able to port the existing formulas as-is, without re-writing them, and in such a way that would not require expert-level understanding of the theory behind them.
- The Monte Carlo simulation randomize functions were also written in VBA
- The precision and format of the results should exactly match the existing spreadsheet to prove the results from the cloud application were correct.

The above factors suggested a port to VB.net as the most prudent solution

- Excel and VB.Net, both being Microsoft products, naturally led to the thought of using the Microsoft cloud product, Windows Azure
- The Windows Azure Home page had a big blue button "Try it free now!"

The above factors were about reducing risk and completing the project in a short period of time with a limited budget.

At this early stage of the project, the idea of migration of an existing application to the cloud and making it into a case study to include in this report was a bit of an experiment. As such, a free cloud trial, including Tools and SDK made sense for this project, as well as for other IT organizations venturing into the cloud for the first time.

A quick tour of the Windows Azure web site confirmed that their platform did support Visual Basic.net along with C#, C++, PHP, Ruby, Python and Java. At this point, due to the factors listed above, the Windows Azure platform and PaaS as the type of cloud platform were clearly the best choices. Final vendor selection would be deferred until after the completion of the Excel spreadsheet transformation into a standalone application.

Throughout the process more research was performed and knowledge gained about Windows Azure as well as other cloud vendors and platforms. Several helpful books for this project included *The Cloud at Your Service* (Rosenberg & Mateos, 2010), *The Cloud Computing Bible* (Sosinsky, 2011), and two books on Windows Azure, *Azure In Action* (Hay & Prince, 2010) and *Programming Windows Azure* (Krishnam, 2010). This reading proved invaluable in helping gain additional insight and knowledge that couldn't be gathered simply by browsing vendor web sites.

8.4 Rewrite the Application for the Cloud

The first step in migrating the application, as shown in Figure 29, was to create a standalone VB.net application that would duplicate the main functionality of the existing Microsoft Excel spreadsheet. Since VB.net was chosen as the language, Microsoft's Integrated Development Environment (IDE) was needed to develop the application. As it turned out, Microsoft was offering a free version of tools called Visual Studio 2010 Express. This appeared to be an ideal starting point since it would not waste money on a toolset if the project had to be abandoned. For other companies considering moving to the cloud or just getting started with the cloud, the idea of a free toolset could be quite enticing. Reading the product description revealed that the functionality was a subset of their professional version and upgrading later would not be a problem. Downloading Visual Studio 2010 Express was straightforward and was up and running in a short period of time.

Subsequently the application needed a sketch of the basic design plan. First, the function of the Excel spreadsheet application had to be understood and then the algorithms and mathematical formulas needed to be reverse engineered. This required some one-on-one discussion with the original application engineer to help better understand this particular model. By utilizing his help, the picture-like Graphical User Interface (GUI) of Excel, and the author’s model, the essence of the task was captured quickly. The logical components of the Excel spreadsheet were divided into building blocks that would need to be created. They were:

- A Graphical User Interface, including:
    - Input parameters for the model
    - Browse controls for input/output data files
    - Function list box

    - Status box to display intermediate and final results
    - Progress bar
    - Run/Stop buttons
- A File I/O section
- An array to hold data set in memory for fast access
- A math section to duplicate formulas
- Control Logic, including the randomize functions for Monte Carlo simulations
- A results and sort routine

The VB.net environment is extremely strong in its graphical user interface (GUI) capability, prompting the choice to start with that aspect of the design first. Reverse engineering the mathematical portions seemed the most challenging part of this phase. The Visual Studio 2010 Express IDE worked flawlessly and the user interface came together quickly. (See Figure 8 below)

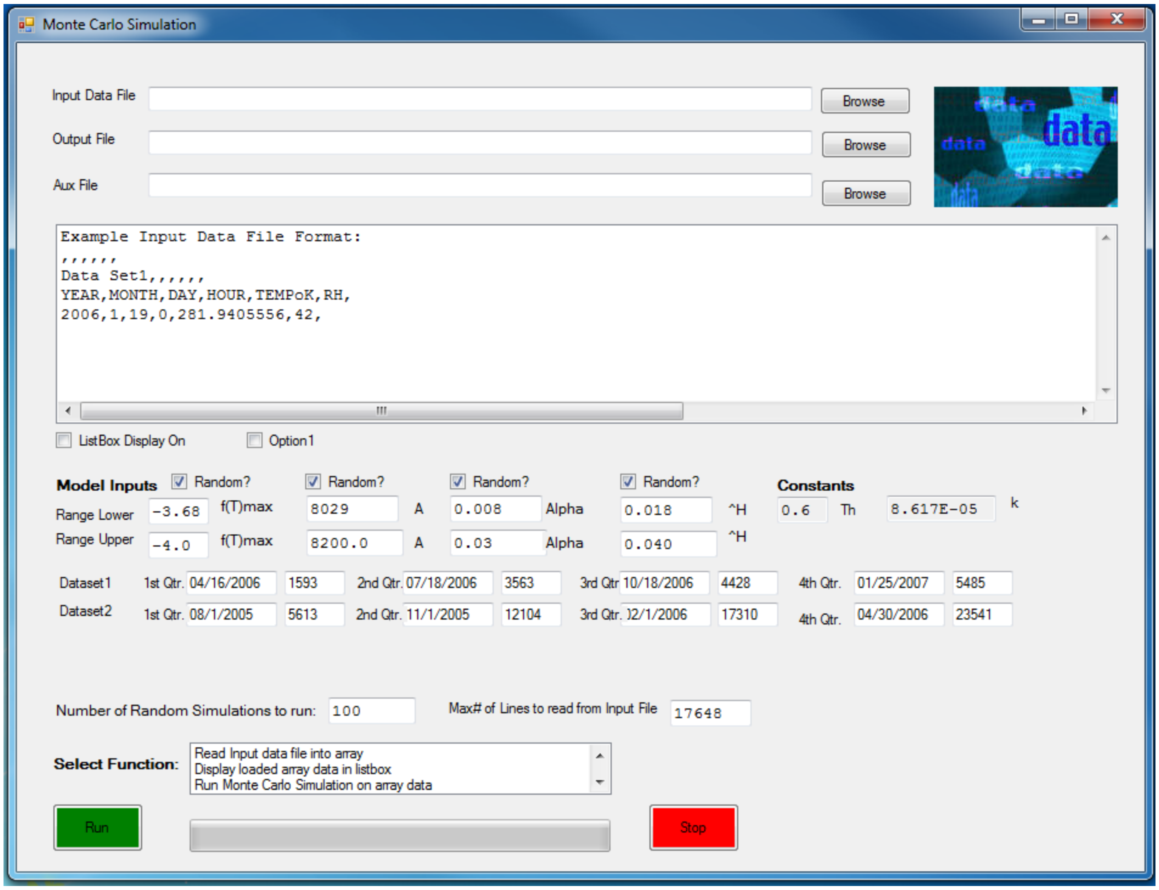

Figure 8: GUI of VB.net App in Visual Studio 2010 Express IDE

Thanks in part to the IntelliSense feature of the IDE coding the File I/O, the general control logic and the randomize functions were easy.

Next it was necessary to tackle the math functions. Since the Excel spreadsheet used VBA, however, porting them to VB.net was almost seamless. A few tricks were needed to intermittently troubleshoot the algorithms in a stepwise fashion so that the code produced the same results as the Excel spreadsheet.

One change to the algorithm was to turn off randomizing of the input parameters in the model, both in Excel and the VB.net, resulting in the answer for one run through the 17,662 rows of data. Again the toolset supported the process with its debugging capability. Stepping through the algorithms with Visual Studio allowed the intermediate results to be compared row-by-row and column-by-column with the Excel results, and the code to be fixed if the calculations were incorrect. The

Excel spreadsheet, as designed by the application engineer, produced eight intermediate results as separate columns within each row, and then contained summation functions at the end. Eventually the program came up with the same correct answers. After plugging in the randomize functions, and testing the new software, the result was a working VB.net application that produced the same results as the original spreadsheet.

8.5 Porting the VB.net App to the Cloud

Preliminary research into Windows Azure stated that it supported VB.net and so it was assumed that it should be ready to migrate. As it turns out, supporting a language is one thing, whereas its host environment is another. The cloud infrastructure is quite different from a desktop PC or server and therefore the existing application needed to be adapted. As in-depth knowledge of Windows Azure was obtained from various white papers such as *Introducing Windows Azure* (Chappell, 2009) and books like *Azure in Action* (Hay & Prince, 2010), a few points of interest surfaced.

One section of (Chappell, 2009) on Windows Azure scenarios seemed to fit the project requirements quite well: "Creating a Parallel Processing Application". This section states that "To interact with the application, the user relies on a single web role instance." Note the term "web role". One question that occurred was how the standalone VB.net application would talk to the user over the internet, as Azure surely had to have some type of interface. It is called an ASP.net app (or another type of web app). In this scenario, the user communicates through a web browser to a web role, which in turn communicates with the worker role (the background processing) via message queues. It turned out that the standalone VB.net application needed to be split into three loosely coupled distinct parts. They would be:

1. A web role to handle the user interface. The ASP.net application would be running under the Internet Information Services (IIS) web server.
2. A worker role with multiple instances. The VB.net module would be running under Windows Server 2008.

3. A class module with the actual business logic. The VB.net class module would be imported by the worker role.

The final architecture showing these parts is shown in Figure 9. The next step was to move the standalone VB.net application to the ASP.net application.

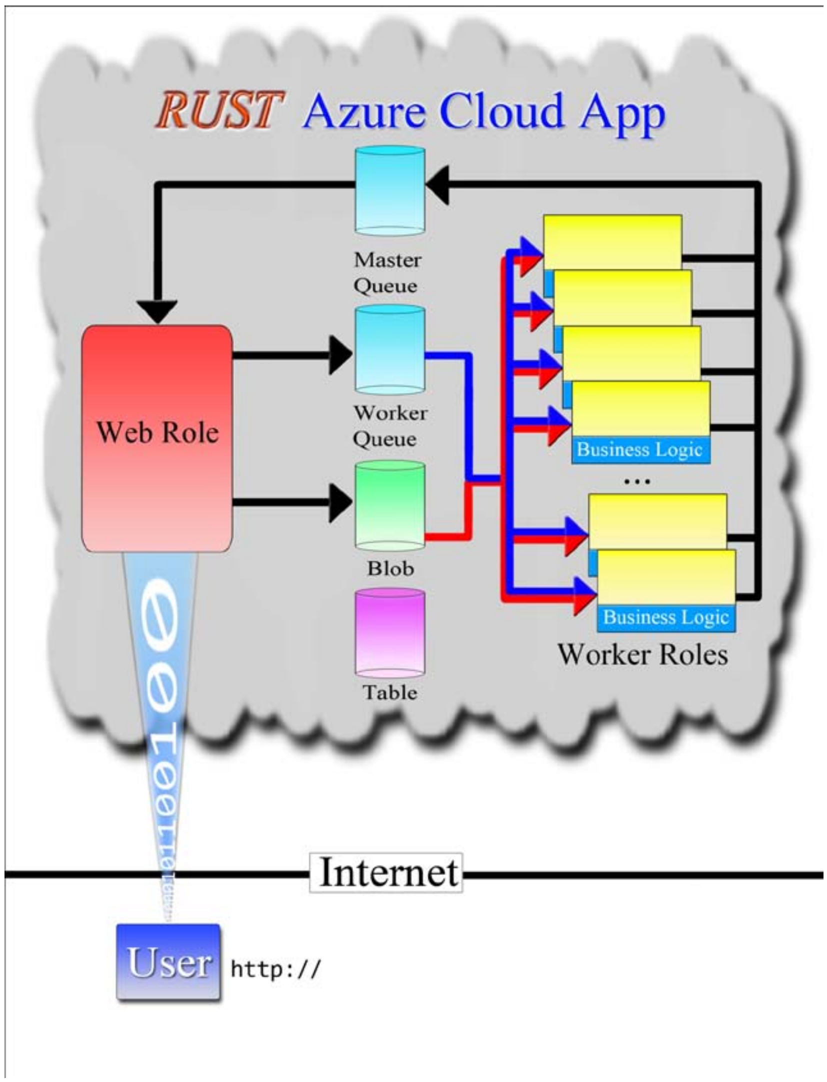

Figure 9: Topography of the Migrated Cloud App

8.6 VB.net to ASP.net

Having read that an existing ASP.net page could be easily added to a Windows Azure application as a web role, it was clear that this step was not time wasted. After completing the VB.net to ASP.net port, the web role GUI would have to be written, and the domain logic would be completed in the form of a VB.net class module, leaving only the worker role/web role inter-communication code. That was the plan at the beginning of this phase, although lessons were learned as it progressed.

The basic idea was to duplicate the existing VB.net GUI functionality in an ASP.net web form. Wherever feasible, the input fields, labels, controls, names etc. had to be kept exactly the same. Also the decision was made to use the vendor's IDE, Visual Studio 2010, throughout the project, and not to code part of the project with a text editor, web design tool, or any other tool. Since the end goal was to develop a working ASP.net front end for a Windows Azure application, the supplied default templates were used initially.

Immediately a problem was encountered, or at the very least a decision needed to be made before getting started. Somewhat confusingly, Visual Studio offers two ways to create an ASP.net application, a web project or a web site. Based on (MacDonald, Mabbutt, & Freeman, 2010), the simpler approach was chosen to use the project-less web site. At first glance, creating a new web site with Visual Studio 2010 produced a template with these main items:

- About.aspx
- Default.aspx
- A master page, site.master
- A style sheet, site.css

Using the tool's "split view" between source and design views revealed a clean style, header, title and, most importantly, a built-in navigation bar. The idea behind the site.master was to give an ASP.net application a consistent look and feel between pages with code inherited from one central source file. Although there was no functional need for an about page or a default page, it was decided to keep them as a starting point for the application separate from the real Monte Carlo simulation code. Once again, the code functionality was separated as much as possible in order to aid in troubleshooting when deploying to the cloud - i.e., the default web page without any cloud specific code should come up first. The built-in navigation code was a significant time saver and, with some minor adjustments to the style sheet, the web site was personalized with a clean look and feel (see Figure 10).

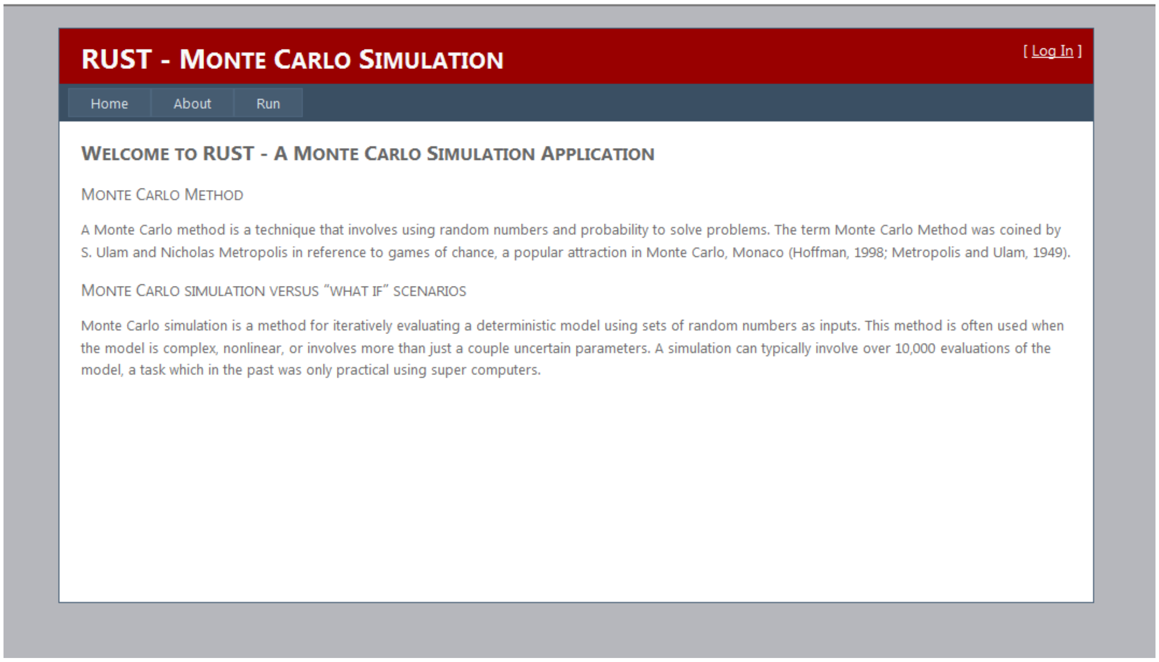

Figure 10: ASP.net Website Home Page

Developing the ASP.net web site in the same language as the standalone model, and using the same IDE used for the standalone VB.net application made the task quite straightforward. Since this was a web application, the method of accessing the data file that the simulation ran against needed to be changed from system I/O to a file upload. However, the built-in server file upload control and example code made this task easy. As stated previously, this migration step involved splitting the domain logic Monte Carlo Simulation into a separate VB.net Class module. It was evident that the Windows Azure application would require using some type of API for the web role/worker role inter-communication, but for this interim step, a simple import statement was used to emulate that functionality.

After completing the above steps, the project was ready to be tested. The Visual Studio 2010 IDE has a built-in web server to test an ASP.net application before actual deployment, and that is what was used to verify that the ASP.net web site would come up and produce the same results as the standalone application. Initial verification was successful (see Figure 11).

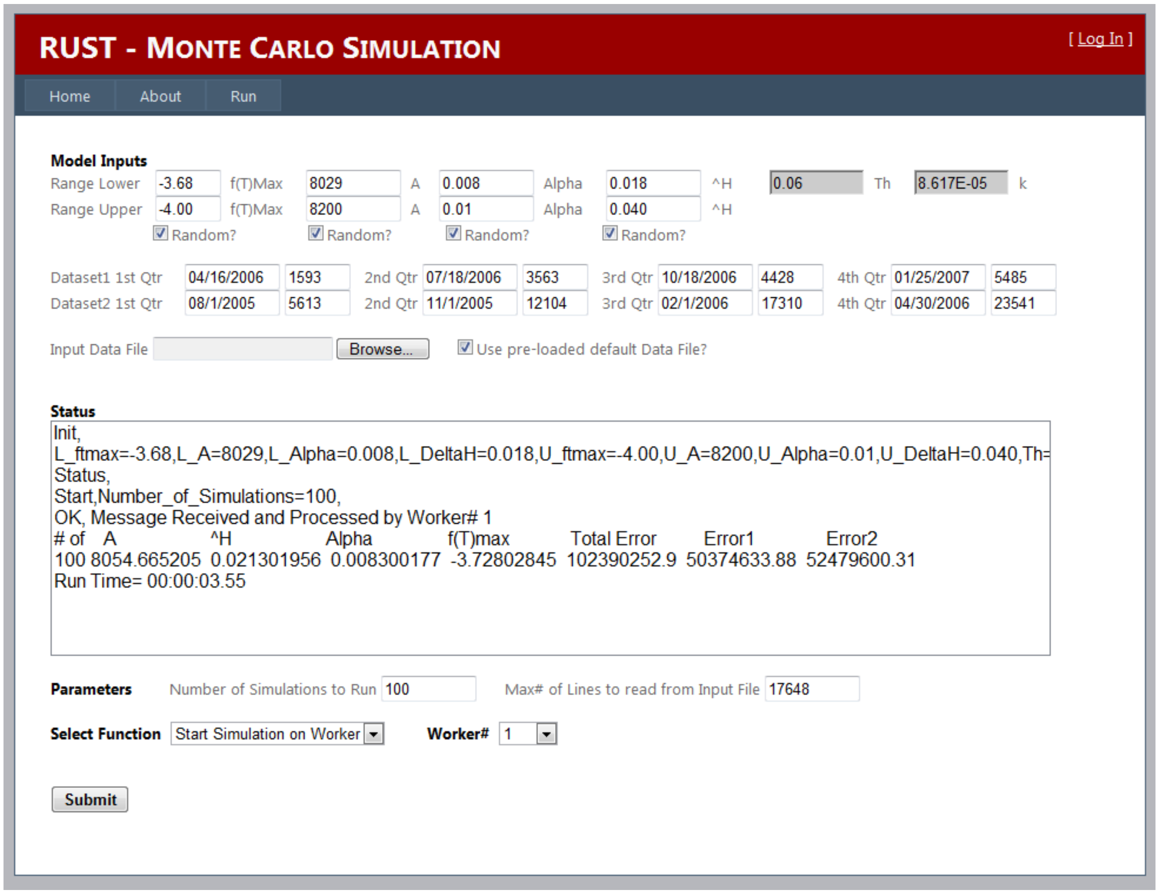

Figure 11: ASP.net Monte Carlo Simulation Tool

Details of anomalies and errors, and their resolution, encountered during this phase of the migration are listed in Section 8.11.

At this point, the project was ready for the actual migration to Windows Azure.

8.7 ASP.net to Windows Azure

8.7.1 Setting up the Windows Azure Development Environment

Finally the project was at the stage where it was ready to begin operating on the Windows Azure platform. At the start of this phase, the free version of the development tools, Visual Basic 2010 Express and Visual Web Developer 2010

Express were still being used but at this point the purchase of Visual Studio 2010 Professional was justifiable.

The Professional version had all of the tools integrated into one IDE, as opposed to the separate tools of the Express version. Interestingly enough, it did not come with the Windows Azure SDK and Windows Azure Tools which extend Visual Studio 2010 "to enable the creation, configuration, building, debugging, running, packaging and deployment of scalable web applications and services on Windows Azure" (Microsoft, 2011). Complications arose from the combination of installing the Professional version to take the place of the Express versions and having to download the tools necessary to develop applications for Windows Azure.

8.7.2 Issues with the Development Platform

Some of the issues encountered with the development platform included:

- Trying to download files from the Microsoft Download Center was unsuccessful during the first attempts. The problem was partially attributed to Internet connectivity, lost connections, etc. and partially to the responsiveness and speed of the host site at different times of the day. Eventually over a period of days and many retries the entire set of tools was downloaded.
- During the March/April 2011 timeframe of this project the Azure SDK changed from Version 1.3 to 1.4 and the installation method changed as well. A much simpler, all-in-one installation of SDK and Tools became available in April 2011. In fact, it changed while an installation of the tools was already in progress.
- After installation it was not intuitive where the tools actually were and how to activate them. For example, there now was an entry under

  Start -> All Programs -> Windows Azure tools For Microsoft Visual Studio 2010 -> Launch Windows Azure Documentation.

  This did not launch any documentation; rather it was just a shortcut to a link back to the Windows Azure web Site. The same was true for the item

Start -> Windows Azure SDK v1.4 -> Windows Azure SDK Documentation. Double clicking on this documentation file caused a re-direction back to a Microsoft Developer Network (MSDN) web page on Windows Azure.

- The above all-in-one installation did not download the Windows Azure Platform Training Kit. This Training Kit later proved to be an invaluable resource. It contained technical content including hands-on labs, presentations, and demonstrations that are designed to help with learning how to use the Windows Azure platform (Microsoft Download Center, 2011).
- Due to operator error, the self-extracting mechanism did not end up extracting and installing the Labs in the Visual Studio 2010 IDE on the two development machines. It seemed as though the installation of the Lab exercises had to be performed in a piecemeal fashion to bring them under Visual Studio. However, they turned out to be terrific working code samples that greatly aided the development effort once they were installed.

8.7.3 Creating the First Windows Azure Project

A major problem arose trying to create the first Windows Azure Project when the Windows Azure Project choice did not appear in the cloud section of the Visual Studio Templates on a workplace PC. However, on a home PC it was present and allowed creation of a Windows Azure Project. In other words, it was supposed to be accessible in both places and was not. This issue was resolved with the help of an Internet search which suggested solutions to fixing the problem.

After resolving the above issue, the next task was to start filling in the Windows Azure Project template with the previously created ASP.net code, including the VB.net Class module. Using code from the hands-on labs of the Windows Azure platform Training Kit, a first cut of the messaging code for the web role/worker role communication was designed and coded. In the Windows Azure platform the web roles and worker roles are loosely coupled and communicate asynchronously with messages stored in storage queues. Basically, the logic of calling a VB.net function or subfunction to initiate an action was replaced with writing a message to a storage queue. After the code for this mechanism was completed, the Azure application was ready to be compiled, built, and then executed in the Azure Emulator on the local machine.

8.7.4 Failures with the First Build

The first build failed to compile successfully, being interrupted by an error stating the path was too long for a part of the project. The IT environment had the *Libraries -> Document* folders mapped to a network drive named by the user that could be automatically backed up at regular intervals. This extended path caused the Visual Studio 2010 path to exceed its maximum limit on the project. The solution was to change the Visual Studio 2010 folder for its application files to a folder directly under the "C:" root directory.

Resolving the above issue allowed the initial running of the web role and a skeleton of the worker role in the Azure emulator. Then, the actual domain logic (Monte Carlo simulation code) was disconnected to enable visualization of the running Azure cloud application. The web role displayed in the browser and the output log of the Azure emulator showed one instance of the worker role up and running.

8.7.5 Debug and Refine the Message Queues

Again using the example code from Labs and various articles on Azure message techniques, the web role/worker role communication logic was designed. All of the examples found to this point, however, used a single message queue with one-way communication from the web role (the master), to the worker role (the slave). Surprisingly, it was even stated that these queues were one-way in nature and that another process should be used to communicate back to the sender (Hay & Prince, 2010).

This one-way communication was not appropriate for this particular application because the web role needed to know when the operations it gave to the worker roles were complete and what the results of those operations were.

A design using two queues was decided upon, one queue for in-bound messages (master queue) and one for out-bound (worker queue). Each message written to a queue had its own unique message ID and each worker role had its own role instance ID. It was now clear that a system could be designed to logically process the items of work. Searching the Internet revealed one simplistic example, written in C#, that used two message queues in this manner (Sawaya, 2010). Incorporating concepts from this example, the design worked with a single role instance. Next, the application was rebuilt with multiple worker role instances and was successfully able to exchange messages in a logical fashion.

8.7.6 Rewrite File Access for Cloud Storage

Now that there was a system in place where the web and worker roles could communicate with each other, the difference between cloud storage and local file access had to be addressed. The Monte Carlo simulation used a comma-delimited text data file consisting of approximately 17,000 lines (512Kb in size) that was loaded into an array for processing. The simulation iterated through the array one row at a time with one set of randomized variables. Successive iterations would yield another random set of variables. It was observed that the worker role could not access the local file system as the standalone application could, hence the need to rewrite this code section for the cloud.

Two options were considered to address this rewrite:

- Embed the data file as a static VB resource
- Upload the file to the web role and then write it to BLOB storage, in which:
    - The web role would send a message to worker role with BLOB ID,
    - The worker role would retrieve the file from BLOB storage.

The second approach would be the most flexible for future implementations that had different data files. Please note that these were only “possibilities”. One Azure publication warned of "gotchas", such as the inability of the Azure emulator to limit local functionality to what was available in the actual cloud environment (Krishnan, 2010). In other words, it was possible to write code to access the local file system or send emails with standard methods and they would function fine while running in the local emulator but they would fail in the Cloud.

Therefore, both approaches were coded and, by default, an attempt was made to load the data file as a VB resource. An operator initiated action could start the second method of uploading the data file to the web role and using BLOB storage to store and retrieve it. The Lab examples were helpful; however they were only manipulating binary files. Searching the MSDN proved to be the answer as the BLOB does support text file storage and by using the appropriate methods and properties, the code was written. After debugging and testing both methods, all could be tied together.

8.7.7 Polishing the Application for the Cloud

As stated previously, this was a scientific application running very time consuming Monte Carlo simulations on a dataset and producing a single final result. Therefore, a typical scenario from an operational viewpoint would be to initialize *n*-worker-roles with different sets of parameters, and then have them start long running simulations. Interim status and final results would be obtained by the web role polling the worker queue for messages and presenting them in a status list box. This status display was accomplished by selecting the Status command from the function list and clicking on the Submit button. However with 10 to 20 worker roles, or more, running, the status update needed to be automated.

A better approach would be to poll the worker message queue at a timed interval and repaint only the status list box. Visual Studio 2010 proved invaluable, as this was accomplished using a combination of three ASP.NET AJAX server controls: the ScriptManager control, the UpdatePanel control, and the Timer control.

This function was made optional to give the user explicit control when needed. After testing, the project was ready to go.

8.7.8 Testing the Completed Azure App in the Local Emulator

Finally, all of the code sections that needed to be changed for the cloud were written and final local testing began. The Windows Azure platform emulator, running under Visual Studio, has the capability of testing locally in two stages. The first is to test an application with local code and local storage; and then the second with local code and actual cloud storage. The first stage was completed without any notable issues.

Testing locally with cloud storage required signing up to obtain a Windows Azure account. Since this was a test and a credit card was required for this part of the project, the costs incurred needed to be thoroughly understood before proceeding.

After reading the various options it was decided to sign up for the *Free* Windows Azure Platform Trial* (see Figure 12). The "free" part referred to the allotted free resources, stating "any monthly usage in excess of the stated amounts will be charged at the standard rates (Windows Azure Platform Fre Trial, 2011)."

The trial service provided the following items:

- Compute
    - 750 hours of an Extra Small Compute Instance
    - 25 hours of a Small Compute Instance
- Storage
    - 20GB
    - 50k Storage transactions
- Data Transfers
    - 20GB in / 20GB out
- AppFabric Access Control transactions - 100k
- AppFabric Service Bus connections - 2
- AppFabric Caching - 128MB cache

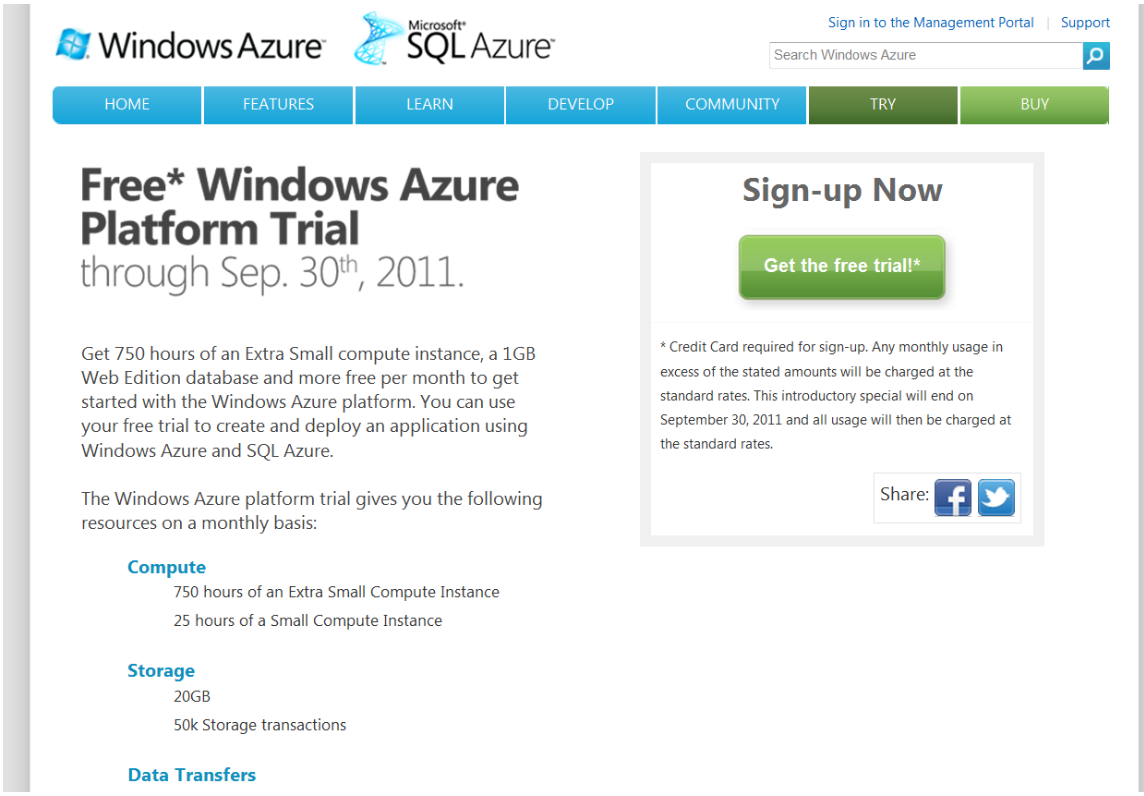

Figure 12: Windows Azure Sign-Up Page

The Extra Small Compute Instances for both the single web role and nine worker roles (for a total of 10 instances) would provide 75 hours of usage. The billing hours refer to time of deployment: even if your web and worker role instances are sitting in an idle loop, you are still getting billed per hour. "If the cube is grey, you're O.K. If the cube is blue, a bill is due." (Hay & Prince, 2010)

After signing up for the Windows Azure account, you are directed to a management portal to obtain storage account credentials. Unfortunately the Windows Azure management portal design had recently changed to a new version with a completely different GUI than the documentation described. After some trial and error, a successful log-in to the older, soon-to-be-discontinued, portal was achieved. This was a somewhat confusing practice for a first-time user, but eventually the storage account setup was completed and the keys needed for the next test phase were obtained.

In order to test code in the local Azure emulator using the cloud storage account, it was necessary to reconfigure both the web role and worker role configuration, specifically the properties on the Settings tab (when using Visual Studio 2010). The *DataConnectionString* and *Diagnostics.ConnectionString* settings needed to be changed from "*UseDevlopmentStorage = True*" to the storage credentials just obtained.

With the re-configuration complete and the application rebuilt and restarted, the Azure application launched and ran. Since the latency of the cloud storage was known to impact performance, the initial concern was only with testing to ensure the application still functioned properly.

Again as a warning, when deploying an application for test purposes, be aware of the charges that you are going to incur.

Do not leave the application deployed or else you will be charged.

Do not think that stopping the application will avoid the charges.

Even when a deployment is in a suspended state, Windows Azure still needs to allocate a virtual machine (VM) for each instance and charge you for it (Miele, 2011).

Details of anomalies and errors, and their resolution, encountered during this phase of the migration are listed in the table in Section 8.11.

8.8 Deploying the Application to the Cloud

The final step in this test project was to actually move the application to the Azure cloud platform. Having done a thorough job of testing the application in the local cloud emulator, confidence was high that the objective was about to be achieved of harnessing the power of the cloud.

As stated above, the project now had a Windows Azure account and a link to the older version of the management portal that matched the "How to deploy an Azure application". The Visual Studio 2010 *Create Service Package Only* option was used to publish the Azure application and upload the Azure application files via the Windows Azure management portal. The application was deployed to the Azure staging area which presented a GUID-type of URL to access the cloud application. A mistake, but also a lesson learned, was that the application was configured for nine worker roles.

The deployment time was directly proportional to the number of instances configured. The average time to deploy the application, with one web role and nine worker roles, was fourteen minutes and ten seconds (14:10). The average time to remove the deployment was one minute and thirty eight seconds (1:38). Until the deployment is solid, it is more expedient to configure the minimum number of roles needed to do a preliminary test of an application.

The moment when the web role and nine worker roles turned from busy to ready, and finally ***green***, was a significant milestone. When the application URL in the Cloud was clicked on, however, a complication arose. A long-spinning browser activity icon was followed by an error stating:
"Server Error - Unknown Error, Cannot display error details from a Remote Server".
The default Web page, which didn't have any Azure specific code, did not appear.

The remoteness of the cloud was evident. The application was known to work fine in the local Azure emulator and, in fact, the ASP.net application was successfully running on an external IIS Web server, so the problem was not immediately apparent. This was a difficult problem, since there wasn't a detailed error message to explain the situation. At the end of the day a decision was made to temporarily undeploy and stop the billing.

After searching the Internet, the conclusion was reached that the problem was really IIS and ASP.net centric. A fix was found that would enable a detailed error message to be produced. The web.config had to be modified to allow a remote server to display a detailed error message.

The next day the application was rebuilt, redeployed, and tried again. After several tries, it was properly configured, finally producing the real error message:

"Default.aspx cannot be found or does not exist".

After researching the problem and searching ASP.net/IIS issues, the answer was found: in the three ASP.net files copied from the ASP.net application to the Azure application, "CodeFile" needed to be changed to "CodeBehind". The root cause was related to how existing, working, ASP.net files are added to an Azure project.

The application worked fine locally, but not in the cloud. This disconnect was a recurring issue discussed previously and an important point of which to be aware.

At last, the web page finally launched. Local testing in the Azure emulator paid off as the application worked properly and as intended. Its performance, as shown in Table 21, was as expected - each worker role subsequently added performed their simulations in the same amount of time. The more workers that were added, the more work that was accomplished.

Table 21: Performance Metrics

Application Type	Number of Simulations	Time in Minutes:Seconds
Excel Spreadsheet	5000	27:29
VB standalone	5000	08:05
ASP.net	5000	02:55
Azure emulator	5000	04:23
Azure Cloud, Extra Small 1 Worker	5000	03:30
Azure Cloud, Extra Small 5 Workers	5 x 5000 = 25,000	03:25
Azure Cloud, Extra Small 9 Workers	9 x 5000 = 45,000	03:15
Azure Cloud, Extra Small 9 Workers	9 x 25000 = 225,000	16:17
Azure Cloud, Extra Small 9 Workers	9 x 50,000 = 450,000	34:52

Details of the anomalies and errors encountered during this phase of the migration, and their resolution, are listed in section 8.11 at the end of this section.

8.9 Conclusions and Lessons Learned from the Migration

8.9.1 The Cloud Really Does Work

As can be seen from Table 21, the Cloud performs as advertised. Using this case study cloud application with 9 worker roles, 225,000 Monte Carlo simulations were performed in a little over 16 minutes, compared to the Excel spreadsheet which took over 20 hours to perform. In test cases, the Windows Azure cloud performed in a linear fashion, with each worker role taking between 3:15 and 3:25 to execute 5000 simulations with random variables. The limit on the trial version was a total of twenty (20) instances, but this can be greatly expanded depending on the subscription plan.

8.9.2 Applications Require a Sizeable Rewrite for the Cloud

Given that almost every migration scenario is different, and this case study only migrated to one cloud platform (a PaaS type), the underlying cloud architecture will dictate that other cloud vendor platforms will require applications to be rewritten to some degree. Golden states, "Once they find out how difficult it is to move an application to an external cloud, their enthusiasm dwindles." (Golden, The Case Against Cloud Computing, 2009).

Figure 13 summarizes the overall process that was taken to migrate the application, once the candidate application was selected and the cloud platform selected.

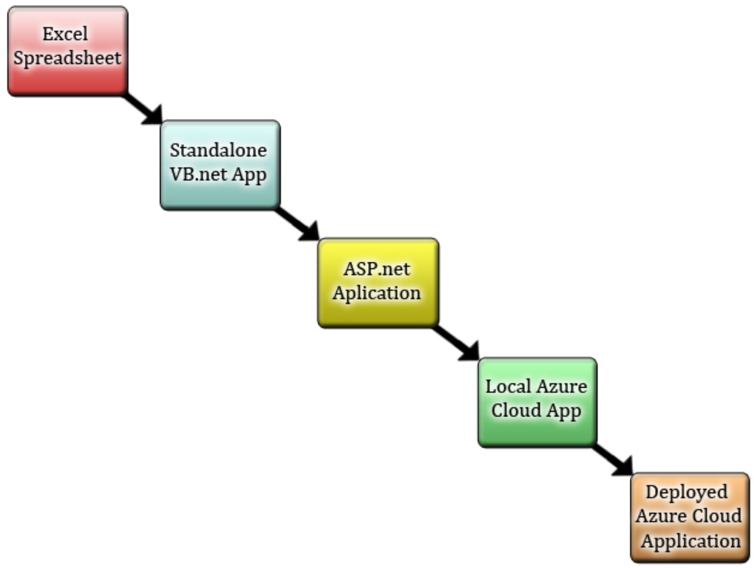

Figure 13: Major Application Migration Development Phases

In this case, the application took one software engineer approximately 28 days, working half-time, for a total of 140 staff-hours.

8.9.3 Understand the Vendor's Billing Policies

Pay only for what you use, or a similar mantra of the cloud vendors is not always accurate. As stated previously, an important point is that the billing hours refer to the time of deployment. Even if the web and worker role instances are sitting in an idle loop, users are still getting billed per hour.

Here is an example where the extra charges may not be readily apparent, stemming from the use of two message queues: "You need another queue and some sort of message correlation to get the response from a worker role. To take this approach with Azure queues, you should take into consideration that every poll on the reception queue will count as a storage transaction and will be billed." (Are queues supposed to be used to send messages back to a web role from worker role, 2011) .

8.9.4 Cloud Platforms are Changing and are Immature

During the timeframe of the migration, from March through May of 2011, the Windows Azure SDK changed from version 1.3 to 1.4. (Now available: Updated Windows Azure SDK and Windows Azure Management Portal, 2011). Also, the Windows Azure management portal changed and that made the deployment section of the books as well as the online documentation obsolete. These changes made it difficult for the author building a Windows Azure project for the first time.

8.10 Complete List of Anomalies/Errors Found

Table 22: List of Anomalies/Errors Found During the Case Study

Stage/Context	Anomaly/Error	Resolution
Excel -> VB.net	None	N/A
VB.net -> ASP.net	Could not load type 'System.ServiceModel.Activation.HttpModule' from assembly 'System.ServiceModel, Version=3.0.0.0, Culture=neutral, PublicKeyToken=b77a5c561934e089'.	This was an IIS error, resolved by running the following from command line: aspnet_regiis.exe /iru
VB.net -> ASP.net	Handler "PageHandlerFactory-Integrated" has a bad module "ManagedPipelineHandler" in its module list.	To repair this problem required running a full silent repair of the .NET Framework 4.0.
VB.net -> ASP.net	ASP.net Server List Box loses the Selected Index when using ASP.NET AJAX server ScriptManager control, the UpdatePanel control, and the Timer control to do a partial page refresh.	
ASP.net -> Azure **See Section 8.8.4**	Build error: The specified path, file name, or both are too long. The fully qualified file name must be less than 260 characters, and the directory name must be less than 248 characters.	Changed the Visual Studio 2010 folder for its application files to a folder directly under the "C:" root directory
ASP.net -> Azure	SetConfigurationSettingPublisher needs to be called before FromConfigurationSetting can be used.	Copied code straight from a vendor Lab exercise, which failed. Needed to add code snippet.
ASP.net -> Azure	Build Errors 1 & 2: Unable to copy file(s) "C:\VS.... to "C:\VS2010\.......	Closed out of Visual Studio 2010 and re-opened or performed a Clean

Stage/Context	Anomaly/Error	Resolution
		Solution before doing a Build or Rebuild.
Local -> Cloud	Azure Account error - P: 84595f40-b921-47d4-b320-e990286e9059 / CC	This was a Timeout/Connectivity issue, resolved by retrying the Windows Azure Portal at a later Date/Time.
Local -> Cloud **See Section 8.9**	Publish Cloud Service fails with error message to the effect: "Cannot accept a .cer file at this time, certificate must be in form of a .csfg file".	Required a workaround: used the Create Service Package Only option and uploaded the Azure application via the Windows Azure Management Portal.
Local -> Cloud **See Section 8.9**	Cannot display the default RUST Web Page in the Cloud, which works in emulator fine: "Server Error - Unknown Error, Cannot display error details from a Remote Server"	Modified the default settings in *Web.config* file to allow the real error message to be displayed. Entry like this was added: <system.web> <customErrors mode="Off" />
Local -> Cloud	Still cannot display the default RUST Web Page in the Cloud, which works in emulator fine: "Server Error - An error occurred during the parsing of a resource required to service this request. Please review the following specific parse error details and modify your source file appropriately, Parse Error default.aspx"	In the three ASP.net files copied from the ASP.net application to Azure, "CodeFile" had to be changed to "CodeBehind". The root cause was in how to add existing, working, ASP.net files to an Azure project. But it worked fine locally just not in the Cloud.
Local -> Cloud	Error when trying to access the Azure account: "Verify your account, As a security precaution complete verification of your account by using a code that will be sent from Windows Live to your mobile"	I was at work with a business phone and did not have a Mobile phone that would accept a text message. Resolved by calling the Windows Azure Support Center.

Bibliography

1. (2010). Retrieved March 14, 2011, from Computer Technologies Consultants: http://usa-ctc.com/infrastructure-as-a-service.html

2. (2011). Retrieved February 23, 2011, from Info.Apps.Gov: http://info.apps.gov/content/cloud-computing-advisory-council

3. *A Breakthrough in Service Delivery for Data Center Workloads â€“ IBM CloudBurst.* manual.

4. *About CIO.gov*. (2011, February 22). Retrieved February 22, 2011, from CIO.gov: http://www.cio.gov

5. *About the Cloud Security Alliance*. (2010). Retrieved March 2010, from Cloud Security Alliance: http://www.cloudsecurityalliance.org/About.html

6. Agency, D. I. (2011). *GIG Content Delivery Service*. Retrieved March 30, 2011, from Defense Information Systems Agency: http://www.disa.mil/gcds/index.html

7. Alliance, C. S. (2009, December). *Security Guidance for Critical Areas of Focus in Cloud Computing v2.1.* Retrieved March 28, 2011, from Cloud Security Alliance: https://cloudsecurityalliance.org/csaguide.pdf

8. *An Introduction to Cloud Computing*. (2009). White paper, from Dialogic.com http://www.dialogic.com/products/docs/whitepapers/12023-cloud-computing-wp.pdf

9. Architectural Strategies for Cloud Computing. (n.d.). Oracle.

10. *Are queues supposed to be used to send messages back to a web role from worker role*. (2011, March 18). Retrieved April 23, 2011, from MSDN: http://social.msdn.microsoft.com/Forums/hu-HU/windowsazuredevelopment/thread/5af2ada3-4757-4570-91ef-8ae563826cd7

11. *Argonne National Laboratory*. (2011). Retrieved March 01, 2011, from CloudBook: http://www.cloudbook.net/directories/gov-clouds/argonne-national-laboratory

12. Armbrust, M., Fox, A., Griffith, R., Joseph, A. D., Katz, R. H., Konwinski, A., et al. (2009). *Above the Clouds: A Berkeley View of Cloud Computing.* techreport.

13. Barrett, L. (2010, February 04). *IBM Lands Air Force Cloud Computing Contract*. Retrieved March 28, 2011, from Datamation: http://itmanagement.earthweb.com/netsys/print.php/3862986

14. *Berkeley NOW*. (1998). Retrieved March 2010, from http://now.cs.berkeley.edu

15. Bernhart, M. (2011, Apr 21). *DISA to Open Forge.mil to Intelligence Community*. Retrieved May 20, 2011, from FierceGovernment: http://www.fiercegovernmentit.com/print/node/15324

16. Biswas, S. (2011, Feb. 01). *Cloud Computing vs Utility Computing vs Grid Computing: Sorting The Differences*. Retrieved March 28, 2011, from Cloud Tweaks: http://www.cloudtweaks.com/2011/02/cloud-computing-vs-utility-computing-vs-grid-computing-sorting-the-differences/

17. Brewin, B. (2011, Jan 03). *DISA Poised to Become Pentagon's Top Cloud Computing Provider*. Retrieved May 20, 2011, from nextgov: http://www.nextgov.com/nextgov/ng_20110103_7911.php

18. Burke, B. (2011, March 30). *Best Practices for Achieving Migration to a Cloud Model*. Retrieved Apr 11, 2011, from DLT Solutions: http://blogs.dlt.com/practices-achieving-migration-cloud-model/

19. Cearley, D., & Phifer, G. (2010, Mar 04). *Case Studies in Cloud Computing.* Retrieved Apr 11, 2011, from Gartner: http://www.gartner.com/it/content/1286700/1286717/march_4_case_studies_in_cloud_computing_dcearley_gphifer.pdf

20. Censer, M. (2011, January 24). *Navy orders a halt to procurement of computer centers*. Retrieved February 23, 2011, from Washington Post: http://www.washingtonpost.com/wp-dyn/content/article/2011/01/21/AR2011012106207.html

21. *Changing the Game for Anti-Virus in the Virtual Datacenter.* (2010, September). Retrieved March 28, 2011, from Trend Micro: http://us.trendmicro.com/imperia/md/content/us/pdf/products/enterprise/datacenter-security/deep-security/wp01_vmware_av_virtualization_100820us.pdf

22. Chappell, D. (2009) *Introducing Windows Azure.* From *www.davidchappell.com*

23. Chong, F., & Carraro, G. (2006). Architecture Strategies for Catching the Long Taill.

24. Cisco. (2010). *Planning the Migration of Enterprise Applications to the Cloud.* San Jose, California: Cisco.

25. Cisco. (2010). *ScanSafe Web Security - Product Data Sheet.* Retrieved March 28, 2011, from Cisco: http://www.scansafe.com/downloads/datasheets/DS_Web_Security.pdf

26. Claburn, T. (2010, October 05). *Google Apps Available to New York Schools*. Retrieved March 28, 2011, from Information Week: http://www.informationweek.com/story/showArticle.jhtml?articleID=227700132

27. *Cloud in Steps.* (2009). Retrieved Apr 11, 2011, from SlideShare: http://www.slideshare.net/imagineashow/migrating-to-cloud-a-step-by-step

28. Cloud Security Alliance. (2011). *Cloud Computing Architectural Framework.* Retrieved Apr 14, 2011, from Cloud Security Alliance: https://wiki.cloudsecurityalliance.org/guidance/images/a/a0/NIST_Visual_Model_of_Cloud_Computing_Definition.jpg

29. Cole, A. (2009, January 29). *Enterprise Applications and the Cloud*. Retrieved May 02, 2011, from IT Business Edge: http://www.itbusinessedge.com/cm/blogs/cole/enterprise-applications-and-the-cloud-can-they-work-together/?cs=30150&decorator=print

30. Cole, A. (2009, Feb 20). *Migration On the Cloud*. Retrieved May 02, 2011, from IT Business Edge: http://www.itbusinessedge.com/cm/blogs/cole/migration-on-the-cloud-more-than-meets-the-eye/?cs=30635&decorator=print

31. *Communications Working Group*. (2011). Retrieved February 23, 2011, from Info.Apps.Gov: http://info.apps.gov/node/7

32. Corrin, A. (2011, Feb. 11). *Army pursuing high-tech data sharing at the tactical edge*. Retrieved Feb. 24, 2011, from Government Computer News: http://gcn.com/Articles/2011/02/09/DOD-cloud-army-dcgs-edge-nodes.aspx?p=1

33. Corrin, A. (2011, March 17). *Culture, policies hinder technology adoption at DOD*. Retrieved March 19, 2011, from Federal Computer Week: http://fcw.com/Articles/2011/03/17/Security-policies-hinder-technology-adoption-at-DOD.aspx?sc_lang=en&p=1

34. Corrin, A. (2011, March 18). *New mobile device policy planned for DOD*. Retrieved March 19, 2011, from DefenseSystems:

http://defensesystems.com/Blogs/Inside-DOD/2011/03/New-DOD-draft-policy-mobile-technology.aspx?sc_lang=en&p=1

35. *CSC Wins Cloud Services Contract for California*. (2010, October 18). Retrieved March 28, 2011, from Business Wire: http://www.businesswire.com/news/home/20101018007262/en/CSC-Wins-Cloud-Services-Contract-California

36. *Data Center Consolidation Plan.* (2010, Feb.). Retrieved March 01, 2011, from CIO.gov: http://cio.gov/documents/Data-Center-Consolidation-Plan.doc

37. Davies, K. (2009). Amylin, Amazon, and the Cloud. 35, 42.

38. Davies, K. (2009). Microsofts Azure Skies. 34.

39. de Assuncao, M. D., di Costanzo, A., & Buyya, R. (2009). Evaluating the cost-benefit of using cloud computing to extend the capacity of clusters. (pp. 141--150). ACM.

40. *Defense Information Systems Agency (DISA)*. (2011). Retrieved March 01, 2011, from CloudBook: http://www.cloudbook.net/directories/gov-clouds/disa-defense-information-systems-agency

41. Defense Information Systems Agency. (2010). *Community: Emerging Techologies*. Retrieved February 2011, from Forge.mil: http://forge.mil/Community-EmergingTechnologies.html#enterpriseTools

42. Defense Information Systems Agency. (2010, December). Enterprise E-mail Partnership Propels Enterprise Services. *The GRID* , pp. 18-19.

43. Defense Information Systems Agency. (2010). *Forge.mil FAQ*. Retrieved February 23, 2011, from Forge.mil: http://forge.mil/Faqs.html

44. Defense Information Systems Agency. (2010). Forge.mil: Transforming the way DoD Innovates IT. United States Government (USG).

45. *Department of the Interior National Business Center (NBC)*. (2011). Retrieved March 01, 2011, from CloudBook: http://www.cloudbook.net/directories/gov-clouds/us-doi-national-business-center--nbc

46. Devaney, E. E. (2010, May 13). *The Cloud*. Retrieved February 24, 2011, from Recovery.gov: http://www.recovery.gov/News/chairman/Pages/may132010.aspx

47. *DIA awards SITE contract for help desk support*. (2011, March 14). Retrieved March 19, 2011, from DefenseSystems:

http://defensesystems.com/Articles/2011/03/14/General-Dynamics-help-desk-support-for-DIA.aspx?sc_lang=en&p=1

48. Difference Between the ASP Model and the SaaS model. (2009).

49. DISA. (2011). *GCDS Brings Navy On-Board in Six Hours to Support Missions in Japan*. Retrieved May 20, 2011, from Disa.mil: http://www.disa.mil/news/stories/gcds_japan_march2011.html

50. *Dreaming On a Cloud*. (2009, March 18). Retrieved April 11, 2011, from ITSkeptic: http://www.itskeptic.org/dreaming-cloud-migration-obstacle-cloud-computing-

51. *E-mail Working Group*. (2011). Retrieved February 23, 2011, from Info.Apps.Gov: http://info.apps.gov/content/e-mail-working-group

52. *Executive Steering Committee*. (2011). Retrieved February 23, 2011, from Info.Apps.Gov: http://info.apps.gov/content/executive-steering-committee

53. *Feds Could Save $20 Billion Using Cloud Computing*. (2011, February 17). Retrieved February 23, 2011, from Federal News Radio: http://www.federalnewsradio.com/index.php?nid=249&sid=2275964

54. FedSources. (2010). *Federal Data Centers Overview.* 8400 Westpark Dr. McLean, VA22102: FedSources.

55. *Five Tips for Moving to the Cloud*. (2011, Feb 02). Retrieved Apr 11, 2011, from TechRepublic: http://www.techrepublic.com/blog/tech-manager/five-tips-for-moving-to-the-cloud/5336

56. Foley, J. (2010, December 08). *USDA Taps Microsoft Cloud for 120,000 Workers*. Retrieved March 28, 2011, from Information Week: http://www.informationweek.com/story/showArticle.jhtml?articleID=228701932

57. Foster, I., Zahu, Y., Loan, R., & Lu, S. (2008). Relationship to Other Technologies. *Cloud Computing and Grid Computing: 360-Degree Compared* . Grid Computing Environments Workshop.

58. Fraser, S., Biddle, R., Jordan, S., Keahey, K., Marcus, B., Maximilien, E. M., et al. (2009). Cloud computing beyond objects: seeding the cloud. (pp. 847--850). ACM.

59. Gartner. (n.d.). *Gartner Hype Cycle.* Retrieved Apr 11, 2011, from Gartner: http://www.gartner.com/technology/research/methodologies/hype-cycle.jsp

60. Gentry, C. (2009). *A Fully Homomorphic Encryption Scheme.* phdthesis.

61. Golden, B. (2009, January 22). *The Case Against Cloud Computing*. Retrieved May 02, 2011, from CIO: http://www.cio.com/article/print/477473

62. Gu, Y., & Grossman, R. L. (2009). Lessons learned from a years worth of benchmarks of large data clouds. (pp. 1--6). ACM.

63. *Hackers using multiple attack vectors to breach mobile phones*. (2011, Feb. 11). Retrieved March 30, 2011, from info-security: http://www.infosecurity-us.com/view/15888/hackers-using-multiple-attack-vectors-to-breach-mobile-phones/

64. *Hackers Using Multiple Attack Vectors to Breach Mobile Phones*. (2011, Feb 11). Retrieved March 28, 2011, from infosecurity: http://www.infosecurity-us.com/view/15888/hackers-using-multiple-attack-vectors-to-breach-mobile-phones/

65. Han, S.-M., Hassan, M. M., Yoon, C.-W., & Huh, E.-N. (2009). Efficient service recommendation system for cloud computing market. (pp. 839--845). ACM.

66. Hay, C., & Prince, B. H. (2010). *Azure in Action.* Greenwich: Manning Publications.

67. High Performance Virtual Machines (HPVM). (1999). http://www-csag.ucsd.edu/projects/hpvm/.

68. Hinchcliffe, D. (n.d.). The Next Evolution in Web Apps: Platform-as-a-Service (PaaS).

69. Holland, R. (2011). *Ten Steps to Successful Cloud Migration.* Cambridge, UK: Eagle Genomics Ltd.

70. Hoover, N. (2010, December 01). *GSA Picks Google for Email*. Retrieved March 28, 2011, from Information Week: http://www.informationweek.com/story/showArticle.jhtml?articleID=228500093

71. Hoover, N. (2010, December 23). *VA Employees Using Unauthorized Cloud Services*. Retrieved March 28, 2011, from Information Week: http://www.informationweek.com/story/showArticle.jhtml?articleID=228900122

72. Humphries, M. (2010, February 05). *U.S. Air Force Wants Secure Cloud Computing*. Retrieved March 28, 2011, from Geek.com: http://www.geek.com/articles/news/u-s-air-force-wants-secure-cloud-computing-selects-ibm-to-prove-it-can-be-done-2010025/

73. Ingthorsson, O. (2011, FEb. 06). *Advanced browsers are key to mobile cloud computing*. Retrieved Mar. 03, 2011, from Cloud Computing Topics: http://cloudcomputingtopics.com/2011/02/advanced-browsers-are-key-to-mobile-cloud-computing/

74. Ingthorsson, O. (2010, November 8). *Android 2.2 and mobile cloud computing*. Retrieved March 03, 2011, from Cloud Computing Topics: http://cloudcomputingtopics.com/2010/11/android-2-2-and-mobile-cloud-computing/

75. Ivan. (2010, Dec 13). *Cloud Vendor Migration Issues*. Retrieved Apr 11, 2011, from Cloud Business Review: http://www.cloudbusinessreview.com/2010/12/13/cloud-vendor-migration-issues.html

76. Jackson, W. (2011, February 15). *What's missing from cloud security*. Retrieved February 24, 2011, from Government Computer News: http://gcn.com/Articles/2011/02/15/RSA-2-Cloud-Security.aspx?admgarea=TC_CLOUD&p=1

77. Joch, A. (2011, February 01). *Four Reasons the Cloud is Safer Than You Thought*. Retrieved February 23, 2011, from Government Computer News: http://gcn.com/Articles/2011/01/31/cloud-security.aspx

78. Jong, A. d. *Cloud Migration Guide.* The Netherlands: Jitscale.

79. Kenyon, H. (2011, February 04). *Army presses forward with e-mail move*. Retrieved February 23, 2011, from DefenseSystems: http://defensesystems.com/Articles/2011/02/04/Army-Begins-Move-to-Enterprise-Email.aspx

80. Kenyon, H. (2011, February 22). *DARPA looks to boost tactical networks with cloud*. Retrieved March 28, 2011, from Defense Systems: http://defensesystems.com/articles/2011/02/22/darpa-communications-programs.aspx

81. Kenyon, H. (2011, February 11). *Energy aims to retake supercomputing lead from China*. Retrieved February 23, 2011, from Government Computer News: http://gcn.com/articles/2011/02/11/energy-supercomputer-to-break-performance-records.aspx

82. Kenyon, H. (2011, March 09). *VA Trying to Shake Off IT Problems*. Retrieved March 28, 2011, from Government Computer News: http://gcn.com/articles/2011/03/09/va-cio-roger-baker-discusses-progress-on-it-issues.aspx

83. Krebs, B. (2007, November). Salesforce.com Acknowledges Data Loss. *The Washington Post* .

84. Krishnan, S. (2010). *Programming Windows Azure.* Sebastopol: O'Reilly Media.

85. Kundra, V. (2010, December 9). 25 Point Implementation Plan to Reform Federal Information Technology Management.

86. Kundra, V. (2011, February 8). Federal Cloud Computing Strategy.

87. Kundra, V. (2010, Feb.). Federal Data Center Consolidation Initiative. Office of Management and Budget.

88. Kundra, V. (2010, Feb.). Memorandum for Chief Information Officers: Federal Data Center Consolidation Initiative. Office of Management and Budget.

89. Kundra, V. (2010, May 13). *Moving to the Cloud*. Retrieved Feb. 22, 2011, from White House: http://www.whitehouse.gov/blog/2010/05/13/moving-cloud

90. Lais, S. (2010, May 24). *CSC makes $317M play for NOAA supercomputer*. Retrieved March 03, 2011, from Washington Technology: http://washingtontechnology.com/articles/2010/05/24/csc-supercomputer-deal.aspx

91. Lenk, A., Klems, M., Nimis, J., Tai, S., & Sandholm, T. (2009). Whats inside the Cloud? An architectural map of the Cloud landscape. (pp. 23--31). IEEE Computer Society.

92. *LHC@Home*. (2010). Retrieved March 2010, from http://lhcathome.cern.ch/

93. Linthicum, D. (2010, Apr 15). *A Step-by-Step Guide for Deploying eCommerce Systems in the Cloud*. Retrieved Apr 11, 2011, from GetElastic: http://www.getelastic.com/depoloying-in-the-cloud/

94. Lipowicz, A. (2011, May 10). *GSA launches $2.5B cloud computing procurement*. Retrieved May 11, 2011, from Washington Technology: http://washingtontechnology.com/articles/2011/05/10/gsa-issues-rfq-for-cloud-computing-options.aspx

95. MacDonald, M., Mabbutt, D., & Freeman, A. (2010). *Pro ASP.NET 4 in VB 2010 .* New York: Apress.

96. Mark Baker, A. A. (2000). Cluster Computing and Applications.

97. McCloskey, P. (2011, March 07). *Cops, Collaboration and the Cloud*. Retrieved March 28, 2011, from Government Computer News: http://gcn.com/Articles/2011/03/07/Editor-Desk-Cops.aspx

98. Mell, P., & Grance, T. (2011). *The NIST Definition of Cloud Computing.* Gaithersburg, MD: National Institute of Standards and Technology.

99. Mergen, M. F., Uhlig, V., Krieger, O., & Xenidis, J. (2006). Virtualization for high-performance computing. *SIGOPS Oper. Syst. Rev. , 40*, 8--11.

100. Merkey, P. (2007). Beowulf History.

101. Microsoft Download Center. (2011, April 18). *Windows Azure Platform Training Kit - April Update* . Retrieved May 03, 2011, from Microsoft Download Center: http://www.microsoft.com/downloads/en/details.aspx?FamilyID=413E88F8-5966-4A83-B309-53B7B77EDF78&displaylang=en

102. Microsoft. (2011, March 16). *Windows Azure SDK and Windows Azure Tools for Microsoft Visual Studio (March 2011)* . Retrieved May 16, 2011, from Microsoft Download Center: http://www.microsoft.com/downloads/en/details.aspx?FamilyID=7a1089b6-4050-4307-86c4-9dadaa5ed018

103. Miele, M. (2011, April 19). *Windows Azure Service Deployment* . Retrieved May 17, 2011, from Technical Notes -- Topics and Examples -- by Michael Miele : http://www.milexsoft.com/2011/04/windows-azure-service-deployment.html

104. Montalbano, E. (2010, December 30). *California Agencies Favor Microsoft for Email*. Retrieved March 28, 2011, from Information Week: http://www.informationweek.com/story/showArticle.jhtml?articleID=228900208

105. Montalbano, E. (2011, Feb. 24). *Local Governments Embrace The Cloud Faster Than Feds*. Retrieved March 18, 2011, from InformationWeek: http://www.informationweek.com/news/government/cloud-saas/showArticle.jhtml?articleID=229219365&cid=RSSfeed_IWK_All

106. Montalbano, E. (2010, October 20). *Microsoft Signs NYC to Cloud Computing Deal*. Retrieved March 28, 2011, from Information Week: http://www.informationweek.com/story/showArticle.jhtml?articleID=227900389

107. Montalbano, E. (2010, November 05). *Minnesota to Consolidate Datacenters*. Retrieved March 28, 2011, from Information Week: http://www.informationweek.com/story/showArticle.jhtml?articleID=228200310

108. Montalbano, E. (2010, December 22). *NATO Taps IBM to Build Private Cloud*. Retrieved March 28, 2011, from Information Week:

http://www.informationweek.com/story/showArticle.jhtml?articleID=228900078

109. Montalbano, E. (2011, March 01). *NYC Opens Consolidated Data Center*. Retrieved March 28, 2011, from Information Week: http://www.informationweek.com/government/cloud/showArticle.jhtml?articleID=229219575

110. Montalbano, E. (2011, January 10). *Treasury Sites Jump to Amazon Cloud*. Retrieved February 23, 2011, from Information Week: http://www.informationweek.com/news/government/cloud-saas/showArticle.jhtml?articleID=229000411

111. Montalbano, E. (2011, Feb 28). *VA to Migrate Email to the Cloud*. Retrieved March 28, 2011, from Information Week: http://www.informationweek.com/story/showArticle.jhtml?articleID=229219490

112. Mowbray, M., & Pearson, S. (2009). A client-based privacy manager for cloud computing. (pp. 1--8). ACM.

113. Murray, P. (2009). Enterprise grade cloud computing. (pp. 1--1). ACM.

114. *NASA*. (2011). Retrieved March 01, 2011, from CloudBook: http://www.cloudbook.net/directories/gov-clouds/nasa-national-aeronautics-and-space-administration

115. *National Institute of Standards & Technology (NIST)*. (2011). Retrieved MArch 01, 2011, from CloudBook: http://www.cloudbook.net/directories/gov-clouds/nist-information-technology-laboratory

116. *National Science Foundation (NSF)*. (2011). Retrieved March 01, 2011, from CloudBook: http://www.cloudbook.net/directories/gov-clouds/national-science-foundation--nsf

117. Neubarth, M. (2009, November 11). *Migrating to the Cloud? Not So Fast*. Retrieved Apr 11, 2011, from CIOZone: http://www.ciozone.com/index2.php?option=com_content&task=view&id=5902&pop=1&page=0&Itemid=572

118. *NOAA IT Strategic Plan: Executive Summary 2010.* (2010). Retrieved March 03, 2011, from CIO.NOAA.GOV: http://www.cio.noaa.gov/FINAL_NOAA_Strategic_IT_Plan_execsum_32410.pdf

119. *NOW AVAILABLE: Updated Windows Azure SDK and Windows Azure Management Portal*. (2011, Mar 09). Retrieved May 18, 2011, from Windows

Azure Team Blog: http://blogs.msdn.com/b/windowsazure/archive/2011/03/09/now-available-updated-windows-azure-sdk-and-windows-azure-management-portal.aspx

120. O'Connor, G. (2009, November 17). *Who Invented Virtualization?* Retrieved March 28, 2011, from PHP Journal: http://php.sys-con.com/node/1175075

121. *OMB Asks Agencies to Review Data Center Targets*. (2011). Retrieved March 28, 2011, from CIO.gov: http://www.cio.gov/pages.cfm/page/OMB-Asks-Agencies-to-Review-Data-Center-Targets

122. *Operational Excellence Working Group*. (2011). Retrieved February 23, 2011, from Info.Apps.Gov: http://info.apps.gov/node/6

123. Perera, D. (2010, October 5). *DHS Embraces public cloud for public websites*. Retrieved March 28, 2011, from Fierce Government: http://www.fiercegovernmentit.com/node/12613/print

124. Perez, S. (2009, May 06). *CloneCloud: The Power of Cloud Computing Comes to Mobile Phones*. Retrieved March 03, 2011, from Read Write Web: http://www.readwriteweb.com/archives/clonecloud_the_power_of_cloud_computing_comes_to_mobile_phones.php

125. Philip, T. (2004). Utility Computing-Identifying the Applicability Domain and Its Boundaries. Zurich, Switzerland: University of Zurich.

126. *Recovery.gov Moves to Cloud Computing Infrastructure.* (2010, May 13). Retrieved February 24, 2011, from Recovery.gov: http://www.recovery.gov/News/mediakit/Documents/Press%20Release05132010.pdf

127. Rosenberg, J., Mateos, A. (2010). *The Cloud at Your Service.* Greenwich: Manning Publications

128. Salesforce. (2011). *The Seven Standards of Cloud Computing Service Delivery*. Retrieved March 28, 2011, from Salesforce.com: http://www.salesforce.com/assets/pdf/datasheets/SevenStandards.pdf

129. Sawaya, G. (2010, May 17). *AzureHelloWorld*. Retrieved April 12, 2011, from AzureHelloWorld-Uv-wiki: http://www.cs.utah.edu/formal_verification/mediawiki/index.php/AzureHelloWorld

130. Schwartz, J. (2011, February 10). *VMware debuts connector for hybrid clouds*. Retrieved February 24, 2011, from Government Computer News:

http://gcn.com/Articles/2011/02/10/ECG-VMware-Debuts-Connector-for-Hybrid-Clouds.aspx?p=1

131. *Security Working Group*. (2011). Retrieved February 23, 2011, from Info.Apps.Gov: http://info.apps.gov/node/5

132. Sideman, A. (2011, Feb. 22). *Readying Patriot missiles for launch? There's an iPhone app for that, too*. Retrieved Feb. 24, 2011, from Government Computer News: http://gcn.com/Articles/2011/02/22/iPhone-teaches-missile-launching.aspx?p=1

133. Sosinsky, B. (2011) *The Cloud Computing Bible*. Wiley Publishing

134. *Standards Acceleration to Jumpstart Adoption of Cloud Computing (SAJACC)*. (2011). Retrieved March 11, 2011, from National Institute of Standards and Technology: http://www.nist.gov/itl/cloud/sajacc.cfm

135. *Standards Working Group*. (2011). Retrieved February 23, 2011, from Info.Apps.Gov: http://info.apps.gov/node/4

136. Sternstein, A. (2011, February 01). *Interior executes test iPads in bid to boost productivity*. Retrieved February 24, 2011, from nextgov: http://www.nextgov.com/nextgov/ng_20110201_9950.php

137. Sternstein, A. (2011, January 28). *White House Considers Relaxing Cloud Security Requirements*. Retrieved March 28, 2011, from nextgov: http://www.nextgov.com/nextgov/ng_20110128_4448.php

138. SYS-CON Media Inc. (2008). *Cloud Application Migration*. Retrieved Apr 11, 2011, from SYS-CON Media Inc: http://cloudcomputing.sys-con.com/node/1458739

139. Technology, N. I. (2011, Feb. 24). *Cloud Computing Collaboration Site*. Retrieved Feb. 24, 2011, from National Institute of Standards and Technology: http://collaborate.nist.gov/twiki-cloud-computing/bin/view/CloudComputing/WebHome

140. Thibodeau, P. (2011, Apr 29). *Amazon Cloud Outage Was Triggered by Network Configuration Error*. Retrieved May 09, 2011, from Computerworld: http://www.computerworld.com/s/article/9216303/Amazon_cloud_outage_was_triggered_by_configuration_error

141. Thibodeau, P. (2011, Apr 22). *Amazon Outage Sparks Frustration*. Retrieved May 09, 2011, from Computerworld: http://www.computerworld.com/s/article/9216098/Amazon_outage_sparks_frustration_doubts_about_cloud

142. Thibodeau, P. (2010, October 11). *Minnesota to move e-mail to Microsoft's Cloud*. Retrieved March 28, 2011, from ComputerWorld: http://www.computerworld.com/s/article/352030/Minnesota_Moves_E_mail_to_Microsoft_s_Cloud

143. Thomas Sterling, D. S. (2009). A High-Performance Computing Forecast: Partly Cloudy. 42-48.

144. Trend Micro. (2011). *Trend Micro and VMware: Partnering to Revolutionize Virtualized Security.* Retrieved March 28, 2011, from Trend Micro: http://us.trendmicro.com/imperia/md/content/us/trendwatch/cloud/revolutionizevirtualizedsecurity.pdf

145. Trends, G. (n.d.). Google Trends Data.

146. *VA Shuts Down Cloud App After Breach*. (2011, January 12). Retrieved March 28, 2011, from FEderal News Radio: http://www.federalnewsradio.com/?nid=249&sid=2231798

147. Vance, J. (2011, January 31). *Mobile Cloud Computing: 5 Key Trends.* Retrieved March 03, 2011, from Datamation: http://itmanagement.earthweb.com/netsys/article.php/3922856

148. Vance, J. (2009, October 02). *Mobile Cloud Computing: Is Your Phone Drifting to the Cloud?* Retrieved MArch 03, 2011, from EarthWeb: IT Management: http://itmanagement.earthweb.com/mowi/article.php/3841611

149. Varia, J. (2010). *Migrating your Existing Applications to the AWS Cloud.* Amzon Web Services.

150. VMware History. (2010).

151. Wali, S. (2010, October 19). *Cloud-Based Infrastructure as a Service Comes to Government*. Retrieved February 23, 2011, from U.S. General Services Administration: http://www.gsa.gov/portal/content/193441

152. Walker, M. B. (2011, February 17). *"CloudFirst" to be Funded by Data Center Consolidation Savings*. Retrieved February 23, 2011, from Fierce Government: http://www.fiercegovernmentit.com/story/kundra-cloud-first-be-funded-data-center-consolidation-savings/2011-02-17?utm_medium=nl&utm_source=internal

153. *Windows Azure Platform Fre Trial*. (2011, May 01). Retrieved May 17, 2011, from Windows Azure: http://www.microsoft.com/windowsazure/free-trial/

154. *World Community Grid*. (2010). Retrieved March 2010, from http://www.worldcommunitygrid.org/

155. World, B.-I. (2009). Cloud Computing for Life Sciences. 23.

156. World, B.-I. (2009). Cloudy Forecast for IBMs Smarter Planet. 36, 42.

157. Yasin, R. (2011, January 14). *After Data Center Consolidation, Beware Legacy Apps.* Retrieved February 23, 2011, from GCN.com: http://www.gcn.com/articles/2011/01/14/afcea-data-center-consolidation.aspx

158. Yasin, R. (2010, February 05). *Air Force, IBM plan to demonstrate secure cloud computing*. Retrieved March 28, 2011, from Defense Systems: http://defensesystems.com/articles/2010/02/04/air-force-ibm-cloud-computing.aspx

159. Yasin, R. (2011, March 16). *Harris Alliance Targets Government Thirst for Cloud*. Retrieved March 19, 2011, from DefenseSystems: http://defensesystems.com/Articles/2011/03/16/Harris-Trusted-Enterprise-Cloud.aspx?sc_lang=en&p=1

160. Yasin, R. (2010, Dec. 17). *Implementing the cloud-first policy? Start with e-mail*. Retrieved Feb. 24, 2011, from Federal Computer Week: http://fcw.com/Articles/2010/12/17/Cloud-First-Policy.aspx?p=1

161. Yasin, R. (2011, Feb. 01). *New York City getting data centers under one roof*. Retrieved Feb. 24, 2011, from Government Computer News: http://gcn.com/Articles/2011/02/01/New-York-IBM-Data-Center-Consolidation.aspx?p=1

162. Yasin, R. (2011, Feb. 16). *NIST budget request could more than double cybersecurity spending*. Retrieved Feb. 24, 2011, from Federal Computer Week: http://fcw.com/Articles/2011/02/16/NIST-cybersecurity-budget.aspx?p=1

163. Yasin, R. (2011, February 04). *State IT diet: Consolidation, cloud and shared services*. Retrieved February 23, 2011, from Government Computer News: http://gcn.com/articles/2011/02/07/feature-cloud-sourcing-side-1.aspx

164. Yunus, M. (2010, June 14). *Understanding Enterprise-to-Cloud Migration and Risks*. Retrieved Apr 11, 2011, from ebizQ: http://www.ebizq.net/topics/cloud_computing/features/12741.html?&pp=1

165. Zhang, X., Schiffman, J., Gibbs, S., Kunjithapatham, A., & Jeong, S. (2009). Securing elastic applications on mobile devices for cloud computing. (pp. 127--134). ACM.

166. Zyskowski, J. (2011, February 09). *Cloud isn't Playing Like We Thought*. Retrieved February 23, 2011, from Federal Computer Week: http://fcw.com/Articles/2011/02/14/HOME-PAGE-Tech-cloud-adoption.aspx