Systems Engineering at MITRE
CLOUD COMPUTING SERIES

A Decision Process for Applying Cloud Computing in Federal Environments

Geoffrey Raines
Lawrence Pizette

MITRE

March 2010; Revised August 2010

Executive Summary

Commercial industry's use of cloud computing, leveraging economies of scale, commoditizing processing power and storage, and incrementally leasing only the information technology (IT) infrastructure that is needed at a given moment, is presenting a compelling business model to Federal leadership, worthy of further investigation. In considering the application of cloud computing approaches to support Federal IT enterprise portfolios, it is useful to have a structured engineering process, flowing from an initial cloud service concept through to implementation, that clarifies and defines the key decision information required and contemporary cloud options. This paper defines an engineering decision process for applying cloud computing services in a Federal Government context and explores important activities such as:

- Scoping a cloud capability effort
- Determining which cloud services will benefit an organization
- Establishing a business case for cloud services

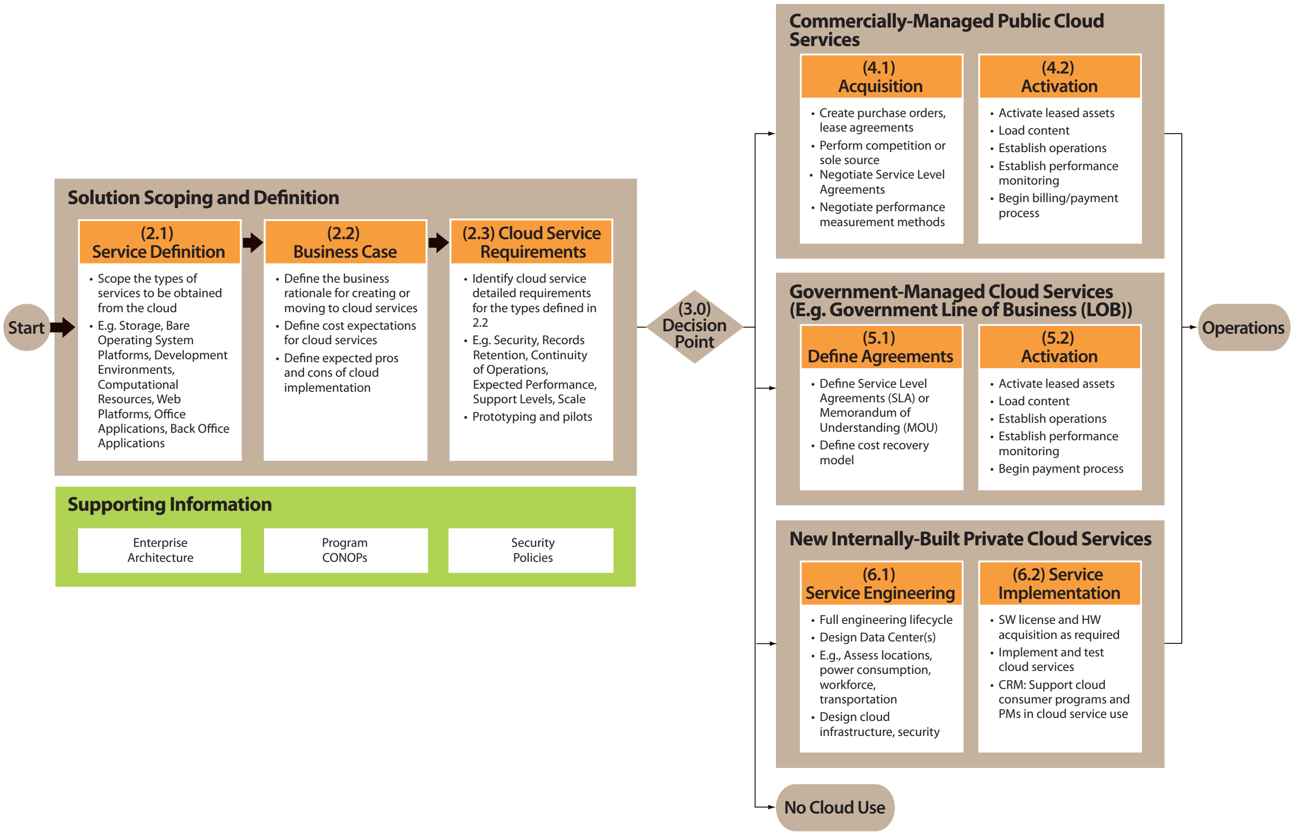
Start
Solution Scoping and Definition
(2.1) Service Definition
• Scope the types of services to be obtained from the cloud
• E.g. Storage, Bare Operating System Platforms, Development Environments, Computational Resources, Web Platforms, Office Applications, Back Office Applications
(2.2) Business Case
• Define the business rationale for creating or moving to cloud services
• Define cost expectations for cloud services
• Define expected pros and cons of cloud implementation
(2.3) Cloud Service Requirements
• Identify cloud service detailed requirements for the types defined in 2.2
• E.g. Security, Records Retention, Continuity of Operations, Expected Performance, Support Levels, Scale
• Prototyping and pilots
Supporting Information
Enterprise Architecture
Program CONOPs
Security Policies
(3.0) Decision Point
Commercially-Managed Public Cloud Services
(4.1) Acquisition
• Create purchase orders, lease agreements
• Perform competition or sole source
• Negotiate Service Level Agreements
• Negotiate performance measurement methods
(4.2) Activation
• Activate leased assets
• Load content
• Establish operations
• Establish performance monitoring
• Begin billing/payment process
Government-Managed Cloud Services (E.g. Government Line of Business (LOB))
(5.1) Define Agreements
• Define Service Level Agreements (SLA) or Memorandum of Understanding (MOU)
• Define cost recovery model
(5.2) Activation
• Activate leased assets
• Load content
• Establish operations
• Establish performance monitoring
• Begin payment process
New Internally-Built Private Cloud Services
(6.1) Service Engineering
• Full engineering lifecycle
• Design Data Center(s)
• E.g., Assess locations, power consumption, workforce, transportation
• Design cloud infrastructure, security
(6.2) Service Implementation
• SW license and HW acquisition as required
• Implement and test cloud services
• CRM: Support cloud consumer programs and PMs in cloud service use
Operations
No Cloud Use

- Defining detailed requirements for cloud services
- Determining when to use internal private clouds or external public clouds
- Assessing when to use community cloud offerings provided by other Government entities
- Understanding when it is appropriate to design and build an internal private cloud.

A Cloud Decision Process—The paper is organized around the inset figure, which depicts an engineering process for the application of cloud computing capabilities in a Federal environment.

Service Definition: The process begins with activities to scope and define the cloud computing effort. Among the many types of possible cloud services in the current marketplace, some subset of services is identified and targeted for use in a Federal organization's portfolio because of an expected benefit. Since "cloud computing" is such a broad term, involving so many possible service types and deployment types, defining exactly which cloud services an organization intends to provide or consume is a fundamental initiating activity. The term cloud computing can be used to refer to a wide range of services that includes data storage, platforms, and applications. As an example, platforms can range from bare operating systems to complete collaboration, development and testing environments, to Web servers, application servers, and databases. It is important to elaborate the organization's service definitions in unambiguous terms. For example, an organization might decide that, "We will establish a means to lease computer storage by the Gigabyte (GB) using a monthly payment such as a purchase card." This bounds the project and service in practical terms, and this definition can then be used to derive design requirements later in the process.

Business Case: Next, a business case is developed that elaborates and clarifies the expected motivation for using cloud capabilities. For example, for many organizations the business case may focus on cost savings, though there could also be compelling motivations in resource agility, new service capabilities, labor savings, and/or leveraging external capital investment. Depending on whether the Federal organization expects to be a cloud service provider, or a service consumer, a number of specific elements are recommended for the business case analysis. For example, a Federal community or private cloud service provider should consider revenue or cost recovery opportunities, capital costs such as facilities, hardware, and software licenses, and expected support staff requirements; along with key service cost drivers such as site redundancy, expected service performance and throughput, and security requirements.

Service Requirements: At this point service requirements are created which will ultimately define the viability of possible commercial and Government-run service provider solutions. The detailed service requirements will include topics such as required security features, required performance, and other service characteristics that will act as differentiators in selecting a solution later in the process. For example, requirements will drive the structure of acquisition documents such as the Statement of Work (SOW) or Statement of Objectives (SOO) if the solution is being leased, or design documents if the solution is built internally. There are several categories of requirements that can be specified for a cloud service, such as cloud service function, service interfaces and standards to be employed, service response times, service failure mechanisms, and support strategies.

Solution Decision: Next, given a defined project scope and service requirements, a solution decision is made for a given cloud capability based on the prior definitional and requirements work. Government organizations, looking to realize the same cloud computing advantages that they see today in the commercial marketplace, have several options for obtaining the benefits of cloud computing, as depicted on the figure's right side. First, the commercial marketplace continues to grow with a wide range of network-based cloud services available for procurement and lease through traditional Federal Acquisition Regulation (FAR) processes, when the commercial service can meet the underlying requirements of the Government. Second, there are several Government run solutions that are being put in place, tailored specifically for Government customers. Third, there is always the option to apply commercial tools to build internal private cloud capabilities completely within Government data center facilities. Within each solution option, common items such as Service Level Agreements (SLAs), contracts, billing, and performance monitoring must be resolved. Finally, the capability moves into its operational phase. Of course, the option to not use any cloud technologies due to an adverse business case remains.

In summary, this document outlines a structured process for scoping a cloud effort, refining potential cloud requirements, developing a business case for a cloud effort, and making a selection among a range of contemporary commercial and Government cloud providers. Fundamentally, this process is focused on understanding and documenting a clear set of service requirements and objectives as an essential early step in selecting a cloud service provider and in deciding to create cloud services for an organization. The engineering process that follows, and the decisions made there, all rest on a robust understanding of the needs of the organization. Cloud computing, when combined with this structured engineering processes, can provide more efficient and agile IT resources for an organization, allowing the commoditization of aspects of IT such as storage and computation, to focus on higher order mission challenges.

Table of Contents

THE BIG PICTURE: In considering the application of cloud computing approaches to support Federal IT enterprise portfolios, it is useful to have a structured engineering process, flowing from an initial cloud service concept through to implementation, that clarifies and defines the key decision information required and contemporary cloud options.

A Decision Process for Applying Cloud Computing in Federal Environments

Geoffrey Raines
Lawrence Pizette

1.0 Introduction

The commercial use of cloud computing, leveraging economies of scale, commoditizing processing power and storage, and incrementally leasing only the information technology (IT) infrastructure that is needed at a given moment, is presenting a compelling business model to Federal leadership, worthy of further investigation. The term "cloud computing" is used to describe a broad Industry movement towards the use of wide area networks, such as the Internet, to enable interaction between IT service providers of many types, and consumers.[1] Commercial service providers are expanding their available cloud offerings to include the entire traditional IT stack from storage, hardware and platforms, to application components, software services, and whole applications. More recently, Government service providers are offering analogous services within Federal networks.

> *"We are moving forward in the direction of cloud computing in the Federal Government, and part of that is to ensure we have to make sure we are leveraging those investments across the Federal Government."*
>
> *– Federal CIO, Vivek Kundra*[2]

The White House's *FY10 Federal Budget* states that, "Cloud-computing will help to optimize the Federal data facility environment and create a platform to provide services to a broader audience of customers… The Federal Government will transform its Information Technology Infrastructure by virtualizing data centers, consolidating data centers and operations, and ultimately adopting a cloud-computing business model."[3] The recently formed *Federal Cloud Computing Advisory Council*, supported by the *CIO Council* under the auspices of the Federal CIO, has initiated efforts to apply cloud computing techniques to large Federal IT challenges, recognizing that cloud computing techniques will be key to producing the underlying scalability and agility in infrastructure required by many Government programs. This ongoing work will also support organizations with the acquisition of commercial cloud capabilities.

While cloud computing concepts offer many potential benefits in the Federal environment, a number of issues remain in its application:

- Commercial offerings may not fully meet Federal needs in areas such as satisfaction of security requirements, Federal record retention rules, and continuity of operations (COOP) requirements
- Federal contracting policies and methods may currently not have the same agility to dynamically engage cloud services as do wholly commercial organizations
- Running cloud capabilities internally, on a smaller scale, may not offer the same benefits as the commercial outsourcing model

Like many large commercial firms, Government organizations will have to decide where cloud computing will fit in their broader IT portfolios, and where the true benefits will be realized. For example, there may be areas where the Government's benefits do not mirror a commercial firm's benefits, given conformance with policy, statutory obligations, or security practices. It is likely that cloud computing, due to inherent cost advantages when implemented on a large scale, will continue to impact the commercial marketplace regarding how commercial IT infrastructure is acquired, maintained, and dynamically scaled. Careful consideration must be given to cloud computing's application in a Federal Government environment.

Document Purpose—This paper defines a decision process for applying cloud computing services in a Federal Government context. As part of this engineering process, the paper will explore activities such as:

- Scoping a cloud capability effort
- Determining which cloud services will benefit an organization
- Establishing a business case for cloud services
- Defining detailed requirements for cloud services
- Determining when to use internal private clouds or external public clouds
- Assessing when to use cloud offerings provided by other Government entities
- Designing and building an internal private cloud

A Cloud Decision Process—Figure 1.0-1 depicts a high-level process for the application of cloud computing capabilities in a Federal environment. The process begins with tasks to scope and define the cloud computing effort. Among the many types of possible cloud services in the current marketplace, some subset of services is identified and targeted for use in a Federal organization's portfolio because of some expected benefit. A business case is developed which elaborates and clarifies the expected motivation for using cloud capabilities. For example, for many organizations the business case may focus on cost savings, though there could also be compelling motivations in resource agility, new

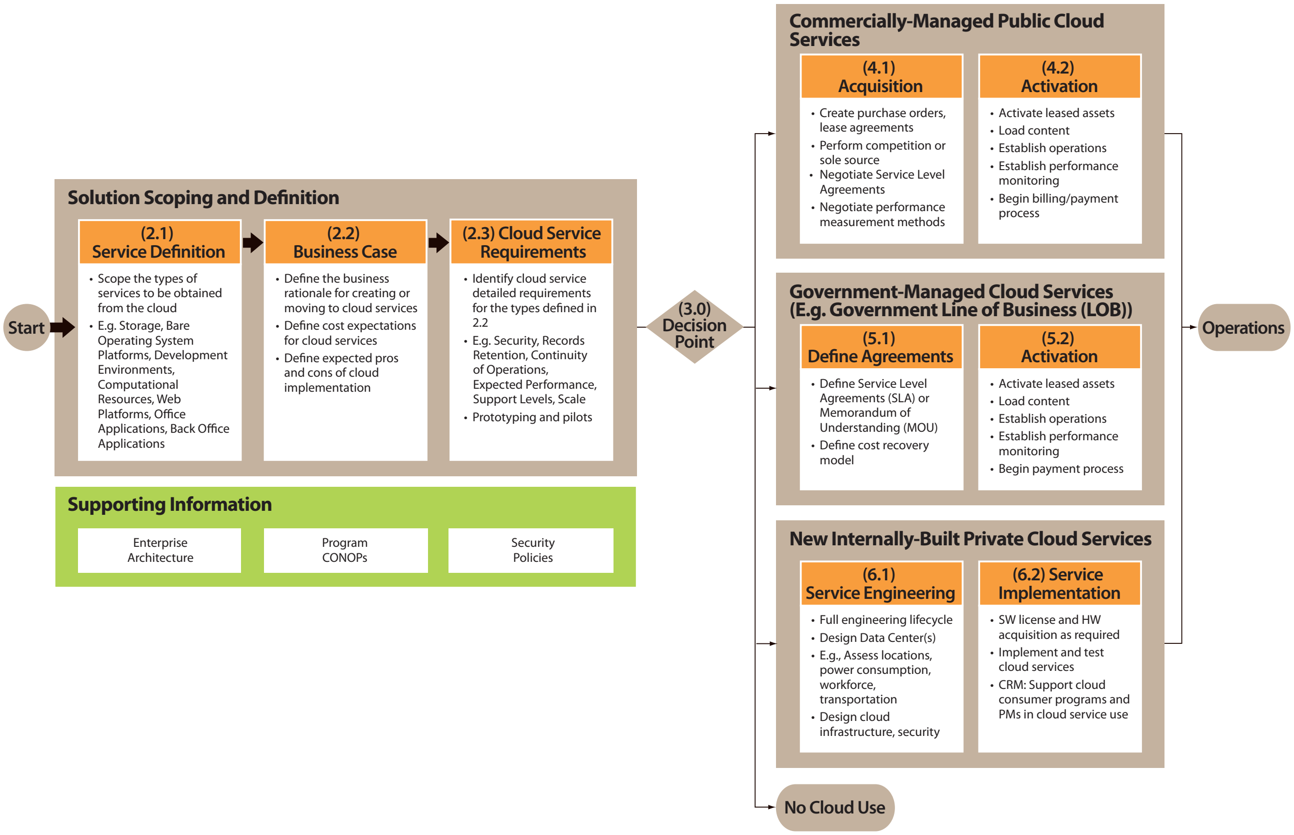
Start
Solution Scoping and Definition
(2.1) Service Definition
• Scope the types of services to be obtained from the cloud
• E.g. Storage, Bare Operating System Platforms, Development Environments, Computational Resources, Web Platforms, Office Applications, Back Office Applications
(2.2) Business Case
• Define the business rationale for creating or moving to cloud services
• Define cost expectations for cloud services
• Define expected pros and cons of cloud implementation
(2.3) Cloud Service Requirements
• Identify cloud service detailed requirements for the types defined in 2.2
• E.g. Security, Records Retention, Continuity of Operations, Expected Performance, Support Levels, Scale
• Prototyping and pilots
Supporting Information
Enterprise Architecture
Program CONOPs
Security Policies
(3.0) Decision Point
Commercially-Managed Public Cloud Services
(4.1) Acquisition
• Create purchase orders, lease agreements
• Perform competition or sole source
• Negotiate Service Level Agreements
• Negotiate performance measurement methods
(4.2) Activation
• Activate leased assets
• Load content
• Establish operations
• Establish performance monitoring
• Begin billing/payment process
Government-Managed Cloud Services (E.g. Government Line of Business (LOB))
(5.1) Define Agreements
• Define Service Level Agreements (SLA) or Memorandum of Understanding (MOU)
• Define cost recovery model
(5.2) Activation
• Activate leased assets
• Load content
• Establish operations
• Establish performance monitoring
• Begin payment process
New Internally-Built Private Cloud Services
(6.1) Service Engineering
• Full engineering lifecycle
• Design Data Center(s)
• E.g., Assess locations, power consumption, workforce, transportation
• Design cloud infrastructure, security
(6.2) Service Implementation
• SW license and HW acquisition as required
• Implement and test cloud services
• CRM: Support cloud consumer programs and PMs in cloud service use
Operations
No Cloud Use

Figure 1.0-1. Leveraging Cloud Computing in a Federal Environment—A General Process for the Application of Cloud Computing

service capabilities, labor savings, and/or leveraging external capital investment. Finally, detailed service requirements are created which will ultimately define the viability of possible commercial and Government-run service providers. The detailed service requirements will include topics such as required security features, required performance, and other service characteristics that will act as differentiators in selecting a solution.

Given a defined project scope and detailed service requirements, a solution decision is made for a given cloud capability based on the prior definitional and requirements work. The solution for a given service, such as cloud-based email, is determined to be satisfied by either commercial offerings, Government offerings, or internally run private cloud services. Within each solution option common items such as service-level agreements (SLAs), contracts, billing, and performance monitoring must be resolved. Finally the capability moves into its operational phase.

Document Organization—The remainder of this document is organized around Figure 1.0-1. Each major element in the figure is labeled with a number that corresponds to a major section in this document. Sections 2.0 and 3.0 describe the steps to scoping and defining a cloud effort, while Sections 4.0 through 6.0 define alternative cloud solutions, each with their own characteristics, benefits, and risks.

2.0 Cloud Computing—Solution Scoping and Definition

This section describes several key activities supporting the successful application of cloud computing approaches. Fundamentally, we want to answer three questions:

- Which cloud computing services do you hope to use? [Section 2.1]
- What is your organizational motivation to move to cloud computing? [Section 2.2]
- For those targeted cloud services, what special requirements do you need to impose on the solution? [Section 2.3]

Accelerators: Where appropriate in the process, we will describe process accelerators that can reduce the time needed to complete a process step in green inset boxes in each section.

2.1 Clarifying Cloud Services

> *"The initial requirements definition and tradeoff phase is rarely performed with sufficient rigor… the importance of spending sufficient time and resources in this initial phase cannot be overemphasized."*
>
> *– Al Grasso, MITRE*

Defining exactly which cloud services an organization intends to provide or consume is a fundamental initiating activity. The term cloud computing can be used to refer to a wide range of services that includes file storage, platforms, and applications. As an example, platforms can range from bare operating systems to complete collaboration, development and testing environments, to Web server application servers, and databases. It is important to elaborate the organization's service definitions in unambiguous terms. For example, an organization might decide that, "We will establish a means to lease computer storage by the GB using a monthly payment such as a purchase card." This bounds the project in practical terms, and this definition can then be used to derive design requirements later in the process.

Accelerators: Taking the time to clearly define and articulate a narrow (as possible) project scope will accelerate the requirements definition and acquisition phases of the project.

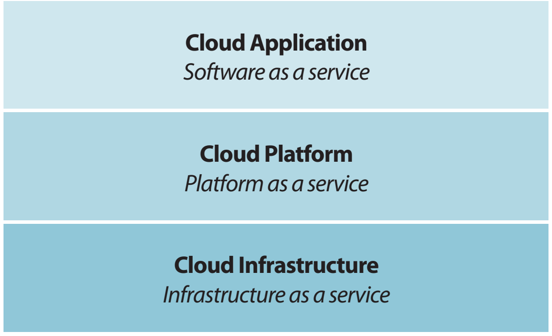

Figure 2.1-1. The NIST Cloud Computing Stack

The commercial cloud marketplace offers a wide range of cloud services that vary in complexity and value. Many analysts describing the cloud marketplace have noticed that it is useful to discuss a 'stack' of services, grouping common services into general categories. The National Institute of Standards and Technology (NIST) has developed one such notional stack for describing cloud computing, with foundational offerings towards the bottom and more complex offerings towards the top.[5] As shown in Figure 2.1-1, the NIST model groups all the services into three general categories.

Cloud Infrastructure: Infrastructure as a Service (IaaS)—At the bottom of the cloud stack, Cloud Infrastructure provides the distributed multi-site physical components to support cloud computing, such as storage and processing resources. This layer allows the infrastructure provider to abstract away details such as which exact hardware an application is using and which data center the application is running in. Advances in server virtualization technologies have made this layer of the stack much more efficient over the past several years, allowing a higher utilization of processing resources than previously practical. Virtual Machine (VM) concepts have also enabled a useful separation of underlying hardware implementation details from the view of developers and the ability to more rapidly scale server resources in response to changing demand.

Cloud Platform: Platform as a Service (PaaS)—Platform offerings provide an infrastructure for developing and operating web-based software applications. Examples include facilities for application design, application development, testing, deployment and hosting as well as application services such as team collaboration, security, application versioning, and application instrumentation.[6] Developer teams frequently work together through their browsers to leverage the virtual cloud platform. Virtual servers run in the cloud can include Web servers, applications servers, and database engines. For some offerings, Application Programming Interfaces (APIs) are provided to predefined network-based functions.

Cloud Applications: Software as a Service (SaaS)—SaaS relies on the cloud for access to what would traditionally have been local desktop software.[7] Advantages of this approach are that the application can be continuously updated by the application provider without issuing and shipping new installation disks. Each time the user logs in to the site, the user will get the latest version of the application. The application provider is also offering a very scalable Web application using a multi-tiered Web architecture, implemented on a considerable infrastructure. Disadvantages include the complete dependence on the underlying network to access the application. When the network is down, the user cannot do any work with the network-based application. In contrast, the desktop version of the software does not require network connectivity for productive work.

In summary, the marketplace offers a wide and ever growing variety of cloud services. Any particular Government program is most likely looking to obtain benefit from some subset of these offerings. Clarifying the scope and definition of the expected cloud services is an input prerequisite of the business case discussed in the next section.

2.2 Developing a Business Case for Cloud Computing

After defining target cloud services, a key activity in scoping and defining a cloud project is thinking through the business case for the use of cloud services. The analysis should focus on the organization's business needs and objectives, clarifying

Accelerators:

- Seek out example business cases from other service-oriented IT efforts of similar scope or purpose to speed the development of the cloud business case.
- Develop standardized reusable business case decision criteria for moving capabilities to a cloud provider.

the expected benefits of the effort, and ensuring that a cloud offering is an appropriate alternative. Depending on the expected type of project, as described in Section 2.1, the business case may be considered from the point of view as a cloud service consumer or a provider.

Items to consider as a Cloud Service Provider— In order to develop a robust business case, a cloud provider must be able to identify its target market, anticipated customers, and their needs. Anticipated customers, such as other Federal agencies and Department of Defense (DoD) programs, should be characterized with accurate requirements to ensure that the cloud offering can provide services with a compelling advantage for its customers. A Federal cloud provider needs to ensure that they can deliver this advantage, while meeting their funding needs and other customer and Government-specific requirements. Example categories to examine are revenue or cost recovery opportunities, costs including hardware/software licensing costs and power consumption, administrative labor, location independence/Continuity of Operations (COOP), performance/throughput, security, and expected performance characteristics such as availability.

Financial considerations for a service provider can include:

- **Revenue or cost recovery opportunities:** A provider should identify the size of their target market and their abilities to recover costs as well as the timing of the cost recovery. It may take time to build the cloud offering and establish incoming revenue flows, funding, or some form of cost recovery.
- **Capital costs:** A provider should examine key capital investments (such as buildings, large hardware purchases) and expenses (such as hardware and software licenses). For a provider, the costs of hardware and software licenses can have an important impact on their ability to economically provide services that meet their customers' needs. Kent Langley explains, "Dollar for dollar the mindful allocation of technical capital can have a dramatic and disruptive effect on the effectiveness of any business and its overall operating capacity."[8] Examples include server purchases and licensing costs for operating systems and virtualization software. Additionally, the fiscal year timing of these expenditures should be examined to ensure the cloud effort will be adequately funded to meet its needs.
- **Support staff:** The ability to administer many environments in a multi-tenant[9] system with a minimum number of support staff can be important to the provider for achieving economies of scale. Many cloud provider technologies reduce cost by effectively expanding the number of systems an administrator can support, using technologies such as virtualization. The administrator cost per machine (or virtual machine) is then lowered.

Service provider requirements considerations can include:

- **Location independence, redundancy, and COOP:** The cloud provider should understand if their users will be leveraging their capabilities for location independence and COOP and, if so, ensure that they can provide this multi-site capability. Costs for providing this capability should be understood, as multiple data centers can affect economies of scale and add complexity to failover scenarios.
- **Performance/throughput:** A cloud provider needs to understand their customers' needs for performance and throughput. These needs can not only affect the quantity and types of hardware and software purchased, but also significantly affect the type of networking capabilities that need to be provided.
- **Security:** Security impacts can be significant for a provider and a customer. Some customers may demand offerings at high classification levels on dedicated hardware and networks, which can significantly impact the business case. Other users may have more flexibility in their security requirements which can allow for a more effective multi-tenant architecture with the associated economies of scale.
- **"Up time":** For a cloud provider, customers' needs for availability and stringent service-level agreements (SLAs) can significantly impact the choices of hardware, software, and support. This can have a large impact on provider costs.

In determining the business case for being a cloud provider, an organization will have to rely upon a forecast for their anticipated usage with a particular offering at a given price. For this offering, the business case should include an estimate of the elasticity of demand. In other words, for a given price, one should project the amount of demand for the service. This will help clarify their usage risk. Costs, such as power, may not be fully known. In order to establish a useful business case, multiple scenarios should be analyzed to understand the implications on cost effectiveness.

***Items to consider as a Cloud Consumer*—**For a cloud consumer, the business case is less complex and needs to be developed based upon organizational goals. The goals may include cost reduction, reduced capital expenditure, ability to scale, location independence or COOP. Meeting mission assurance/security and performance considerations should also be considered in detail. Once these elements are understood, the following attributes can be examined:

Financial considerations for a service consumer business case can include:

- **Costs:** A user of a cloud service can reduce their costs if they can effectively leverage the economies of scale that are available to a large cloud provider. Troy Benohanian writes, "The cloud model disrupts the traditional IT infrastructure by introducing a virtually infinite pool of computing resources that are available on-demand and payable by the hour. Businesses need to figure out which one of their applications currently running inside the firewall can take advantage of this type of virtual IT infrastructure."[10] Additionally, the nature of some costs will change from being capital investment in hardware and infrastructure (CapEx) to a pay-as-you go (OpEx) model with the cloud. The White House notes in the FY10 budget request, "Of the investments that will involve up-front costs to be recouped in outyear savings, cloud-computing is a prime case in point. The Federal Government will transform its Information Technology Infrastructure by virtualizing data centers, consolidating data centers and operations, and ultimately adopting a cloud-computing business model."[11]
- **Support staff:** Cloud computing is essentially an outsourcing model. The consumer should anticipate reduced infrastructure resources support staff.

Service consumer requirements considerations can include:

- **IT agility:** The Department of Interior's National Business Center (NBC) writes, "Agency CIOs and decision makers need to recognize that a large part of the appeal of the Cloud is that it brings technology closer to business users. With Cloud technologies, business users have rapid access to servers, storage, and computing platforms that were once the exclusive domain of IT."[12] Cloud computing approaches put IT agility in the hands of end users and this can be a qualitative benefit in a business case.
- **Location independence/COOP:** A significant advantage of a cloud offering can be the ability to access the cloud from anywhere. This cost of obtaining this capability from the cloud should be compared against the cost of implementing location independence and COOP from other alternatives (such as a program specific investment.)
- **Performance/throughput:** A consumer of cloud services needs to be very aware of the limits of performance and throughput from a cloud offering. SLAs and network capabilities need to be understood in order for the business case to be viable.
- **Security/mission assurance:** In order for a customer to leverage a cloud offering, the consumer organization needs to understand the security and mission assurance characteristics of the cloud and how that cloud offering would affect their business (i.e., what they would do if they lost access to a remote cloud). For example, if a Federal cloud provider has already been through a level of certification and accreditation (C&A) that can facilitate the adoption of the cloud by the consumer, it can lower the program risk of a transition to cloud resources.
- **"Up time":** For a cloud consumer, stringent SLAs that specify characteristics such as "uptime" reliability metrics should be considered essential to the business case.

***Assessing Risks*—**When a Federal organization adopts a cloud approach, they will necessarily have to rely upon another organization to provide capabilities to them. They are trusting another organization to provide capabilities when they most need it.

The consumer may also be opening up a new network attack vector and exposing themselves to other risks. The anticipated benefits may be cost reduction through lower capital expenditures and the ability scale processing as needed. The consumer may get inherent COOP and location independence, which can be high-value capabilities for them.

In essence, the consumer business case needs to effectively weigh the benefits of cost reduction, scalability, location, independence, and other benefits, against the additional exposures or risks added. The exposures may include relying upon another organization, network dependence and performance. In a positive business case, the benefits of cost reduction would be coupled with an acceptable mission assurance posture that would make cloud computing a viable and appropriate choice.

2.3 Cloud Service Requirements

Once the expected cloud services are defined and the motivation for using cloud services is clarified in the business case, service requirements are developed. Of particular interest are requirements which are unique to the Federal Government or to the organization being supported. As Al Grasso writes, "The initial [program] requirements definition and tradeoff phase is rarely performed with sufficient rigor… The importance of spending sufficient time and resources in this initial phase cannot be overemphasized." [13] Capturing these varied aspects of a service's expected operational characteristics is not trivial and is necessary for robust procurement actions or structured solution development.

Accelerators: Set up quick turn-around engineering tasks on existing task order contracting vehicles to identify cloud service requirements. Have these tasks be independent of the main acquisition of the cloud service.

As with most engineering processes, a key driver to a successful solution is a comprehensive set of well-articulated requirements. Requirements will drive the structure of acquisition documents such as the Statement of Work (SOW) or Statement of Objectives (SOO) if the solution is being leased, or design documents if the solution is built internally. There are many types of requirements that can be specified, such as cloud service function, service interfaces and standards to be employed, service response times, service failure mechanisms, and support strategies. It is important to note that these requirements can directly drive cost for the provider and therefore they must be defined with sufficient attention as they generally can not be changed unilaterally once a binding contract is in place without unpleasant consequences. Unlike a custom multiyear system's development effort, leasing a cloud service for many commercial and Government service providers is a 'point of sale' decision, where the user fills in payment and account information in a web browser and selects a predefined service offering for a specified time period. In this context there is advantage to the Government in taking the time to get the requirements as correct as possible up front, and using that information to select services that match the Government's requirements profile.

Examples of general cloud service requirements include:

- **Security requirements:** Federal Government IT programs have a wide range of possible security requirements. For some agencies, unique Federal security requirements can make it difficult for vendors to successfully offer the Government their commercial services. Commercial providers may apply security measures that they feel are sufficient in the marketplace, though they do not completely fulfill Federal needs or policies. Federal requirements can range from compliance with Federal Information Security Management Act (FISMA), compliance with FIPS or agency specific policies, to special Certification and Accreditation (C&A) requirements, or periodic vulnerability scans. Since security requirements can drive the selection of a cloud solution with pass/fail criteria, it is key that a cloud capability program consider all security requirements in depth before proceeding to a decision on general cloud approach.
- **Records management requirements, and records retention:** Federal agency records management programs must comply with regulations promulgated by both National Archives and Records Administration (NARA) and the General Services Administration (GSA) (36 CFR 1220.2)[14] Storing information in the cloud will require some technical mechanism to achieve compliance with records management (RM) laws and policies. Agencies have several legal responsibilities regarding Government records

including “Establishing and maintaining an active, continuing program for the economical and efficient management of the records” (44 U.S.C. 3102) and “establishing safeguards against the removal or loss of records and making requirements and penalties known to agency officials and employees” (44 U.S.C. 3105). The cloud solution has to support analogous record safeguards and retrieval functions, even in the context of a provider termination.

- **Continuity of Operations (COOP):** Depending on the organizational missions supported by the cloud capability, COOP can be a driving solution requirement. The purpose of a COOP capability is to ensure that mission essential functions continue to be available in times of crisis or against a spectrum of threats. Threats can include a wide range of potential emergencies, including localized acts of nature, accidents, and technological and/or attack-related emergencies. COOP requirements may be defined in policy as in *Department of Defense Directive 3020.26, September 8, 2004, Certified Current as of January 1, 2007, USD(P), SUBJECT: Defense Continuity Program (DCP).* The DCP states that “COOP involves plans and capabilities covering the same functional objectives of continuity of Government (COG), must be maintained at a high level of readiness, and be capable of implementation both with and without warning. COOP is not only an integral part of COG and Enduring Constitutional Government (ECG), but is simply ‘good business practice’—part of the Department of Defense’s fundamental mission as a responsible and reliable public institution.”
- **Expected service performance:** The organization’s required performance of the cloud service should be clarified and documented. Cloud service performance expectations could include speed of operation, expected performance incentives or penalties, and methods to accommodate service demand changes.
- **Performance measurement:**[15] Often the contracted cloud service provider cannot be held accountable for the network to which the service is being attached. While the network impacts the consumer-perceived performance of the service, the service provider contractor has no control or management responsibilities for the wide area networking infrastructure, and they may consider the network as a Government-provided capability. In this case, the service provider will rightly require that service performance characteristics, such as response time, be measured at the service provider’s entry point(s) on the Government’s network.
- **Support levels:** Requirements concerning the speed and depth of administrative support must be clarified. Administrative support requirements can cover a number of topics such as the daily hours of prime support, problem escalation times, resolution of recurring problems, and trouble ticket submission methods.
- **Scalability requirements:** The organization should consider requirements concerning the ability of the cloud solution architecture to either grow or shrink over time, with varying levels of processing, storage, or service handling capability.

3.0 Choosing Cloud Providers

With detailed cloud service requirements now understood, the Government organization comes to a decision point regarding which approach to take for a cloud-based service. As shown in Figure 1.0-1, there are three fundamental choices, which are now explored in this paper:

- **Acquire the service commercially, or through a negotiated Government-wide contractual vehicle that accesses a commercial vendor (Section 4.0):** Given the cloud requirements defined in the previous step of the process, each of the three solution options can be considered and evaluated in depth. Acquiring the service commercially requires some understanding of the contemporary commercial marketplace and methodically compares the organization’s detailed requirements against possible commercial cloud offerings. Detailed data gathering and due diligence must be performed regarding external service providers. Unique Government requirements, such as security, may drive an organization away from a wholly commercial solution.
- **Obtain the service through another Government organization (Section 5.0):** There is a growing list of Government-run initiatives and participation in Government cloud computing groups such as the *Federal Cloud Computing Advisory Council* will allow an organization to maintain familiarity with those options.

Cloud Type	Control	Security	Cost	Offerings
Section 4.0 Public commercial cloud offerings	Outsourcing: Least control—services largely driven by commercial requirements	Wide variety in the suitability of service for Government security requirements	Most likely to have best economies of scale	Broad and deep marketplace with a wide variety of offerings—growing rapidly
Section 5.0 Government-run for Government customers—community cloud services	Outsourcing: Control by agreements—MOUs, SLAs, etc.	Cultural understanding and legal compliance with Government standards	More expensive than commercial options	Smaller set of providers with a maturing set of offerings—targeted at Government needs
Section 6.0 Government-built private internal clouds	Internal: Most control over service characteristics—built and managed by organizations IT staff	Built to organizations specifications—can handle high security requirements	Potential to be most expensive	Great flexibility, but can require significant time and capital investment

Figure 3.0-1. Generalizing the Characteristics of Cloud Alternatives

- **Build the cloud service internally and be a service provider (Section 6.0):** The third option, to "build your own" offers a wide range of program risks and costs, though it gives the greatest control to the bill payer.
- A last option is to decide that cloud computing is not beneficial to the organization at this time. For example, the organization may find the expected cost savings are not as large as expected, or that the initial investments to move systems to cloud infrastructure are prohibitive in the context of other budget priorities.

A Government organization must assess these options against their requirements and choose a fundamental path for a cloud solution. Within each path are further considerations and tradeoffs as defined in the next sections.

Figure 3.0-1 generalizes the broad cloud alternatives for Government decision makers in early 2010. Being a generalization, it is safe to say that there are exceptions to the characterizations described in Figure 3.0-1. However, in broad terms, as shown in the figure, building and running a dedicated internal private cloud gives the Government the most control over the details of implementation, technical solutions, staff, and compliance with a variety of laws and policies. Similarly, in general terms, many commercial service providers that we have examined are leveraging much larger economies of scale than the Government programs we have interviewed, driving down unit costs for storage and processing. Finally, there is a healthy and growing marketplace of cloud-based commercial services with hundreds of firms offering cloud services, and a much smaller set of Government providers.

Accelerators: Many agencies possess reports and data intended for Public consumption. The new Data.gov effort is a repository for this type of information. Identifying this Public focused information can be a quick start to hosting information on a cloud computing service.

The ultimate decision on which path to take for a cloud service solution, as shown in Figure 3.0-1, is driven by the characteristics of the organization's applications and data. After performing the due diligence defined in Sections 2.1 through 2.3, which brings the organization a good understanding of the program's scope and requirements, the pivotal solution requirements will be known. For example, if a high security level is a legal requirement, and the Government can only find certified service solutions in Government-run facilities or by constructing their own infrastructure, then the choices become narrower. If the systems in question provide or host information extracts intended for the Public, then a wide range of solutions are available, as long as they prevent information modification and have sufficient availability and bandwidth. In summary, the decision process in this section focuses on the matching of the organization's IT characteristics to the offered service solutions.

Hybrid Approaches—When looking across a portfolio of IT services to be provided to a large organization, a senior leadership team may decide to implement a hybrid model, with some services coming from a cloud provider and some being supported internally. For example, it may be cost effective and compliant with policy to use cloud services for email, while keeping certain mission critical servers on internally operated platforms. Similarly, cloud services such as storage may be used in conjunction with a virtual private network (VPN) to augment internal enterprise storage resources for a time when demand is sufficient. Hybrid approaches can cross the approaches in Sections 4.0 through 6.0 or may cross internal and external service providers.

4.0 Commercially Managed Public Cloud Services

This section discusses the acquisition and use of commercial cloud services, which are generally Internet Protocol (IP) addressable services on the public Internet.

Accelerators:

- When commercial service providers meet your needs, use Government-wide acquisition schedules, such as Apps.gov, to more quickly leverage pre-negotiated services and rates.
- Choose commercial service providers that have already successfully concluded Certification and Accreditation (C&A) activities.
- Look for vendors with solid and verifiable past performance hosting analogous service offerings.

4.1 Acquisition of Cloud Services

Successful acquisitions have always relied on strong requirements management; and when acquiring cloud services, this continues to be true. Fortunately, cloud services have natural network-based interface definition points and performance specifications that can translate into contract requirements. As discussed in Section 2.3, when specifying the requirements for a cloud service, several areas are considered to ensure that the operational service meets the Government's expectations include: interface descriptions, performance measurement methods, service-level agreements (SLAs), data rights, and options for graceful increased and decreased cloud service usage. These topics represent key elements that will provide a foundation for successful consumer/provider relationships. In most cases, when the Government creates contracts with commercial vendors, it is important that these topics are contractually specified.[16]

When acquiring commercial cloud services, Government organizations have several choices. A 'full and open' competition can be held between all vendors who wish to respond. Holding such a competition can be labor intensive as the quantity of bidders can be large. However, for newer services, which are yet to be put a purchase schedule, or to access commercial companies that are not Federally focused, this can be a viable alternative. A more limited competition can occur on an already existing Indefinite Delivery Indefinite Quantity (IDIQ) contract. This option uses pre-qualified vendors who have previously been awarded a blanket contract against which Task Orders (TOs) can be written. With cloud services, some of which are very new, the new startup commercial vendors may not have existed or competed for Federal work when the blanket contract was put in place. In some rarer cases a sole source contract can be awarded to a vendor with a unique service offering, though substantial justification would be needed to explain the rejection of a competition.

GSA Storefront—Currently, a new option for cloud services is being championed by Federal CIO Vivek Kundra. The General Services Administration (GSA) has recently established an online catalog of cloud services. The projected services will be focused on cloud infrastructure and the lower portion of the stack mentioned in Section 2.1, include three fundamental types of cloud offerings – storage, virtual machines, and Web hosting[18]. The establishment of this purchasing catalog will allow for the rapid purchase of a series of fundamental cloud services essential as pre-priced commodities. For example, an organization could lease a quantity of storage measured in GBs for a time period measured in months. [This approach is an example of a storefront catalog, where the Government establishes a predefined purchasing mechanism. Options where the Government sells services to the Government as an approved line of business (LOB) are discussed in Section 5.0] For services in other portions of the IT stack, such as Software as a Service (SaaS), traditional acquisition remains the Government's option.

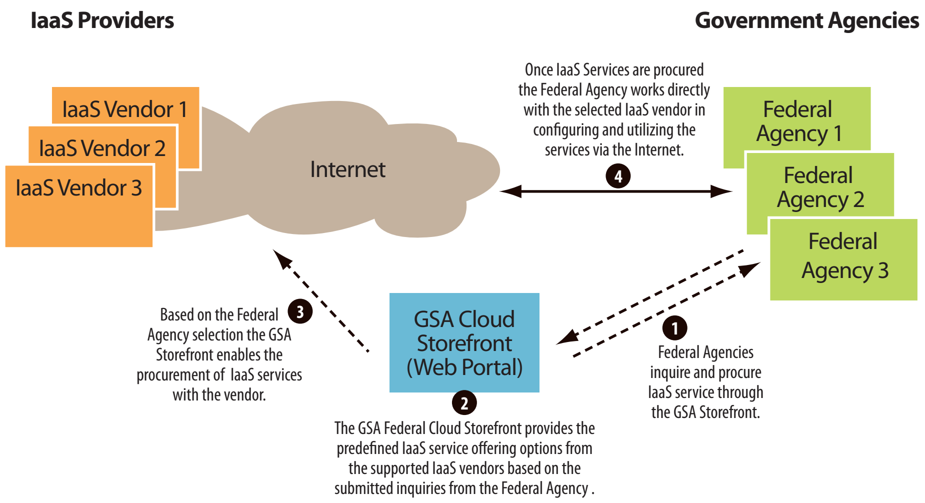
IaaS Providers
Government Agencies
IaaS Vendor 1
IaaS Vendor 2
IaaS Vendor 3
Internet
Once IaaS Services are procured the Federal Agency works directly with the selected IaaS vendor in configuring and utilizing the services via the Internet.
4
Federal Agency 1
Federal Agency 2
Federal Agency 3
3
Based on the Federal Agency selection the GSA Storefront enables the procurement of IaaS services with the vendor.
GSA Cloud Storefront (Web Portal)
1
Federal Agencies inquire and procure IaaS service through the GSA Storefront.
2
The GSA Federal Cloud Storefront provides the predefined IaaS service offering options from the supported IaaS vendors based on the submitted inquiries from the Federal Agency .

Figure 4.1-1. Procuring Cloud Services Through the GSA Federal Cloud Storefront[17]

Acquisition Considerations—Contractual documents for cloud services should define the following key topics as applicable to the cloud services being acquired:

- **Data Center Locations**: Defines characteristics of the location of underlying data centers. For example, a consumer organization could specify only Continental U.S. (CONUS) data centers for service instantiation.
- **Data Center Continuity:** Defines a capability to recovery from failures at particular physical locations and suggests that more than one data center location is used.
- **Data Center Inspection:** Defines the consumer's right to inspect the data center during operations.
- **Demand Elasticity:** Defines an ability to respond to a variety of service demand levels. Since one of the key commercial benefits of cloud services is the ability to scale service demand, this criterion is used to evaluate that feature of a provider.
- **Automated Service Provisioning:** Defines the degree of automation supporting the provisioning of a service. For example, can a consumer organization provision a storage service by providing point of contact information and a funding mechanism such as a credit card? How much staff intervention is required to make this service operational? There are wholly commercial examples where this is completely automated.
- **Provisioning Response Time:** Defines the time it takes to fully provision a service.
- **Service Performance Measurement Tools:** Defines how the performance of cloud services, will be measured. This is central to the operational use of the service-level agreement (SLA).
- **Support:** Defines the mechanism for provider staff support when service issues arise. For example, what is the help desk mechanism?
- **Service Levels:** Defines performance characteristics such as the reliability, up time, or availability of the cloud service. This is a center piece of the SLA. (SLAs are discussed further in Section 5.1)
- **Security Considerations:** Defines security compliance and ongoing C&A compliance approaches.
- **Administration Functions:** Defines how administrative functions such as billing, order tracking, and user support will occur. This effort must include interaction and oversight by the consumer, such that a feedback loop for contract management is created.
- **Interface Definitions:** Defines how technical information about cloud service interfaces will be conveyed to the consumer. This can include Application Programmer Interfaces (APIs) when appropriate.

***Benefits*—**The benefits of utilizing commercial cloud services can include lower capability unit costs, greater burst capability, and faster automated provisioning. Commercial cloud providers can often drive costs down through economies of scale and other cost reduction techniques. For example, they may be able to vertically integrate their supply chain and build their own specialized hardware platforms; locate their data centers in lower cost locations (e.g., electric power, labor, and taxes); minimize labor costs and the number of administrators they need; and reduce overhead and licensing costs through size and volume. By having a larger capacity, they can also provide greater burst capability for their customers. In *A Berkley View of Cloud Computing*, the authors suggest two cases where utility computing burst capability provides an advantage over private clouds: "A first case is when demand for a service varies with time. A second case is when demand is unknown in advance." [19] Lastly, commercial cloud providers often already have an infrastructure established to provide a low-friction automated provisioning of, and payment for, cloud services.

***Risks*—**The Table 1 defines common risks and possible mitigations for using commercially managed public cloud services.

4.2 Activating Cloud Services

After a successful acquisition process, the provisioned cloud services are made operational for the consumer in the cloud. Many commercial public cloud offerings have simplified activation of their services to the point where it is possible to register and access some capabilities in less than an hour. Activating cloud services can include the following general steps:

- Providing identity/account owner
- Establishing a billing mechanism
- Selecting support level
- Selecting and configuring environments
- Loading content
- Testing cloud services

For many commercial offerings, the following steps are interactive through a browser interface. For example, *Amazon.com* provides an on-line interface for registering for their storage and platform services.[21] The consumer can select from on-demand instances or reserved instances.[22] Although Google's offering is different than Amazon's, with Python and Java development environments for the Google AppEngine, both of them have similar Web self-service access.[23]

Table 1. Risks and Mitigations for using Commercially Managed Public Cloud Services

Common Risk/Concern	Mitigation Approach
Imprecise requirements may not adequately support the acquisition process.	Use short-term engineering tasks on existing contract vehicles (e.g., IDIQ task order) to apply subject matter expert (SME) technical experience to the definition and clarification of the service requirements.
Full and open competitions are time consuming and not 'agile' in response to organizational needs.	When leasing commercial cloud services, try first to use pre-negotiated services off of an approved buying schedule, such as the one being established by *Apps.gov*.
Few standards currently exist to ensure portability among cloud service providers, or to ensure common detailed definitions of 'platforms' and 'infrastructure' services.	Currently, there are not many cloud-oriented standards for the consumer to use to mitigate this risk. Continue to look for standards to increase portability among service providers and to reduce lock-in. Consider evolving standards such as the Open Virtualization Format (OVF)[20] which describes standards for virtual machine descriptions. Be cognizant of development languages that are vendor specific, unless program leadership can accept the program risk of being tied to one platform.
The commercial marketplace may not support the whole portfolio of Government IT needs, particularly in areas such as records retention or security.	Consider the organization's entire portfolio of IT needs against the common solution alternatives – commercial market, Government-run, or private internal cloud. Match the service providers' characteristics to the needs of the systems and data under consideration. In the final analysis, only a small portion of an organization's IT portfolio requirements may be outsourced to wholly commercial organizations. For example, the commercial option might support an organization's public facing Web servers or content repositories that are intended for public audiences, and therefore have low security requirements.

Providing Identify/Account Ownership—The first general step in establishing usage on a commercial platform is registering for an account, which in most cases requires a username, password and contact information, such as name, company (or Federal organization), email address and mailing address.

Google[24] and *Amazon.com*[25] integrate this account information with their other services. Through the sign up process, key code information may need to be recorded in order to access the service at a later time.

Establishing Billing For some service providers billing information is provided before services begin, while other offerings will request billing information when volume reaches beyond a free level or introductory period. On publicly available offerings, a credit card or other commonly used payment vehicles can be utilized for billing. Government consumers must often select a payment vehicle that will match the dollar thresholds set in organizational policy. For example, a purchase card might be used for costs up to $2500. Formal payment systems may be used for larger amounts.

Currently service pricing for platforms and storage tends to be very competitive, starting at approximately $0.10 per CPU hour and $0.10 per GB of data transfer.[26] *Salesforce.com* pricing for their SaaS Contact Relationship Management (CRM) system starts $9/month/user for their most basic offering and runs through to $250/month/user for their unlimited offering.[27] Microsoft is entering the market with their Azure offerings with similar costs as Amazon and Google. Microsoft has recently completed a Community Review phase, and going live in January 2010.[28]

Selecting Support Levels During the process of establishing an account, the user may often select a support option. For example, *Amazon.com* offers Silver and Gold levels of support which start at $100 and $400 per month, respectively. The Gold level provides support 365 days per year, 24 hours per day.[29] Similarly, *Force.com* offers a premium level of service for $75/month per user.[30]

Selecting and Configuring Environments After establishing an account, some offerings allow the user to go straight to selecting their operational environment, while some software development activities require the incorporation of local client tools. For development, Microsoft Azure[31] requires client development tools as does Google App Engine.[32] Google App Engine allows for the download of their software development kit (SDK) or an Eclipse plug-in. For monitoring and managing operational environments, client side tools may improve the interaction with the cloud offering; however, most offerings provide a thin-client, Web-based dash board.

For cloud offerings, some mature providers offer a wide-range of cloud environments from which to choose. For example, Amazon EC2 users can select different size instances and a variety of operating system and other environment options shown in Figure 4.2-1 below.

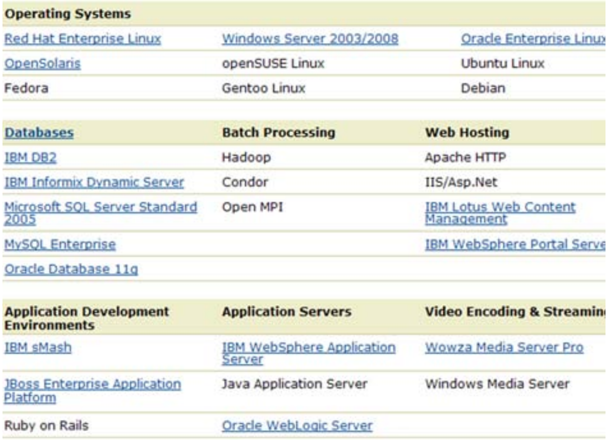

Operating Systems		
Red Hat Enterprise Linux	Windows Server 2003/2008	Oracle Enterprise Linu
OpenSolaris	openSUSE Linux	Ubuntu Linux
Fedora	Gentoo Linux	Debian

Databases	Batch Processing	Web Hosting
IBM DB2	Hadoop	Apache HTTP
IBM Informix Dynamic Server	Condor	IIS/Asp.Net
Microsoft SQL Server Standard 2005	Open MPI	IBM Lotus Web Content Management
MySQL Enterprise		IBM WebSphere Portal Serve
Oracle Database 11g		

Application Development Environments	Application Servers	Video Encoding & Streamin
IBM sMash	IBM WebSphere Application Server	Wowza Media Server Pro
JBoss Enterprise Application Platform	Java Application Server	Windows Media Server
Ruby on Rails	Oracle WebLogic Server	

Figure 4.2-1. Platform Configuration Example: Amazon EC2 Sample Environment Options[33]

When selecting and configuring environments, there may be a variety of security and communications options to choose as well. For example, networking options to establish connectivity with an instance, such as Hyper Text Transfer Protocol (HTTP) and associated ports may need to be selected.

Loading Content—When loading content in cases when platforms or storage is leased, users generally load information to cloud resources via wide area network (WAN), unless very high data volume is needed. For public offerings, WAN throughput is usually charged at a published rate. Transfer costs must be considered. For example, Amazon currently charges between $0.10 and $0.17 per GB for non-load balancing transfers in and out of regional data centers. Google charges between $0.10 and $0.12 per GB for transfers.[35]

While the WAN is sufficient for many cloud based activities, some data transfers are so large that they cannot be reasonably accomplished through WAN connection, due to excessive time requirements or bandwidth cost. For these situations, vendors such as Amazon are offering to transfer data from other media such as disks.[36] In the Berkley white paper, *Above the Clouds: A Berkley View of Cloud Computing*, the authors highlight the importance of transferring large amounts of data: "One opportunity to overcome the high cost of Internet transfers is to ship disks. Jim Gray found that the cheapest way to send a lot of data is to physically send disks or even whole computers via overnight delivery services." [37] As Amazon notes in Steven Lawson's article, "With many enterprise Internet connections, Import/Export will often be faster than online uploads or downloads… The method isn't new: For example, when banks set up new branches and want to have large amounts of information available on site, they typically ship drives because they don't have days to wait for a transfer." [38]

Testing Cloud Services—Operational testing of cloud services builds upon traditional system testing to include attributes of cloud service offerings. The testing program should cover cloud service topics such as the ability to respond to a variety of service demand levels, automated service provisioning, provisioning response times, cloud service interface definitions, security compliance, administrative functions, and service level characteristics such as the reliability, up time, or availability of the cloud service. If applicable, testing from multiple operational locations should be performed to reveal any network-driven issues.

Service Outages—Service outages for commercial cloud providers are often newsworthy events. Service consumers should have a strategy for a degraded network or service outage. Consumer programs should consider the implications to operations if cloud provider services were not available or were degraded, and establish a strategy for continuing mission critical operations under these circumstances. Consider the capabilities that are essential for successful operation during an outage; not all capabilities may be needed during a nonstandard event. Programs should emulate or simulate degraded network conditions to understand the behavior of cloud-dependent systems with reduced or problematic network connectivity. Programs should perform negative testing to understand the implications of service outages on mission critical capabilities and analyze the service implementation architecture for the actual redundancy of service offerings on the network. Consumer programs should incorporate the strategy for continuing operations of mission critical capabilities into system test plans. For example, if an alternate source of information will be utilized in the event of a network outage, test the system with the alternate data source.

5.0 Government Managed Cloud Services

In addition to the commercial options discussed in Section 4.0, a Government consumer may opt to lease cloud services from another Government organization. Currently, in conjunction with an expressed desire from the administration, a number of Government Agencies are organizing cloud offerings to be leased inside the Government. This section focuses on the considerations when obtaining cloud services from another Government organization.

5.1 Defining Cloud Computing Service Agreements

Analogous to the formal contract used in commercial relationships, a key to establishing a cloud computing relationship with between a Government managed provider and Government consumer is the defined terms of service and service level agreement (SLA). These can be codified in a document such as a memorandum of understanding (MOU) between the parties. MOUs for Government entities are similar to contracts discussed in Section 4.0, but less formal, without the legal case law, precedent, and enforceability of contracts. Typical Government-to-Government MOUs will include goals and objectives, responsibilities by party, points of contact,

Accelerators:

- Since security and risk mitigation are likely to be driving factors in selecting a Government-provided offering over a commercial offering, request provider security information early in the decision process.
- Employ a quick proof-of-concept to ensure that the Government cloud service meets your application and data requirements.

period of performance, and signatories. For cloud computing services, they should also include metrics reporting, performance measurement techniques, points of escalation, cost recovery if applicable, and mutual security requirements.

Service Level Agreements**—**SLAs specify many aspects about the delivery of the service that are essential for mission success. Terms for SLAs for cloud services should include characteristics such as:

- Response time
- Hours of operation
- Platform availability
- Expected throughput and utilization ranges
- Maximum permitted down-time
- Performance measurement and reporting requirements
- Performance-based pricing
- Problem resolution thresholds
- Problem escalation and priorities

The SLA should specify the support steps that the consumer can take when the service is failing to meet the terms specified in the agreement. These support steps should include points-of-contact and escalation procedures. The time-to-resolve performance should be specified in the contract based upon the severity of the problem.

As mentioned above, it is important to be precise in the definition of metrics and specify when and where they will be collected. For example, performance is different when measured from the consumer or provider due to the propagation delay of the network. Metrics should measure characteristics under the control of the vendor or they will be unenforceable. Finally, the SLA should describe a mutual management process for the service levels, including periodic reporting requirements and meetings for management assessments.

Network Topology Requirements**—**Network topology, including the implementation of the backbone of the current network, will have an impact on how the services perform as measured from a distributed consumer's perspective. For example, if users are geographically distributed around the country, it is often important to have the service available at multiple locations. In this situation, the MOU should include requirements for a multi-site implementation, and what is often termed "global load balancing", a method for balancing service load across distributed sites. With network-based services, the consumer wants to specify an offering with a minimum number of network hops between the majority of distributed consumers and the provider locations. This can shorten the average time between service requests and service responses.

Multi-site load balancing also positively impacts total system availability. Having multiple sites can provide continuity of operations (COOP) and inherent support for contingency planning, if a given location has a significant outage such as a fire, power failure, or natural disaster. While it is true that temporarily routing consumer requests to a more distant network site is not optimal for the consumer, for many missions operating in a lower performing state is better than having the service being completely unavailable. Forrester's Stan Schatt explains, "Best practices include using these devices [global server load balancers] to enhance server failover, overcoming high availability limitations of a single data center, improving the efficiency and availability of multi-homed links to more than one service provider, and localizing Web content for clients." [39]

Performance Measurement**—**Often the service provider cannot be held accountable for the network that the service is being attached to. While the network impacts the consumer-perceived performance of the service, the service provider contractor has no control or management responsibilities for the wide area networking infrastructure, and they may consider the network as a third-party capability. In this case, the service provider may require that service performance characteristics, such as response time, be measured at the service provider's entry point(s) to the network.

It can be advantageous for a consumer to "instrument" key points on the network to measure performance of service providers. For example, contemporary commercial tools can report back to a centralized data store on service performance, and instrumentation agents can be placed with participating consumers, and at the entry point of the service provider on the network. By gathering data across providers on the performance of pre-planned instrumented service calls throughout typical work periods, service managers can better judge where performance bottlenecks arise. The Government should include requirements for service instrumentation to support broader performance engineering efforts.

Cost—A Government run cloud is likely to provide *less* operational cost savings than a market leading commercial cloud vendor. There are multiple reasons for this. Commercial vendors tend to minimize all factors of cost in a manner the Government cannot match due to policy or security requirements. For example, cloud vendors may employ less costly labor that have not obtained Government security clearances. They may spread their costs over a much larger customer base than the Government enterprise cloud, and they may be able to realize cost savings by not needing to follow Government processes for certification and accreditation, and security. Also, they are more likely to set up data centers in low-cost electric regions or perhaps even outside of CONUS.

When evaluating or defining cloud offerings, the type and timing of payments should be considered – upfront payments and length of time for measuring variable charges should be part of the analysis. For the provider, one-time payments or fixed revenue streams can help to mitigate some of the risk with establishing a cloud offering. Also, longer term units of measure for variable costs may be helpful in allocating computing resources to customers. For example, the *Defense Information Systems Agency (DISA) Rapid Access Computing Environment (RACE)* charges $500 per/month/image in a development and test environment.[40] The monthly time unit is significantly longer than many publicly available commercial offerings. For example, *Google AppEngine* measures and leases computing resources by the hour.[41] Some companies offer multiple options: *Amazon.com Elastic Compute Cloud (EC2)* has a reserved instance and an on-demand instance pricing model. As an example, the reserved pricing model allows for a one-time, non-refundable payment of $227 for one year and $0.03 per hour for a small Linux instance.[42] With no advance payment or commitment, *Amazon.com* also offers $0.10/instance/hour for their small offering to $0.80/instance/hour for their extra large instance. This fine-grained pricing with one-hour time blocks and provisioning in minutes provides for a “pay-as-you-go” offering. It allows the consumer a cost effective means to utilize the cloud in a burst capacity without incurring the long-term expense of mostly unused computing capacity.

Example Cloud Programs—Currently, a number of new cloud provider programs have been initiated by the Government. They include:

- **Defense Information Systems Agency (DISA—Rapid Access Computing Environment (RACE):** DISA offers a cloud service with an MS Windows or Red Hat Linux computing environment.[43] RACE prices are listed in Network World as $500 per month for development and test environments, and $1200 per month for production capabilities.[44] In addition to choosing the operating system, users can optionally select an operating environment that complies with Security Technical Implementation Guides (STIGs). This environment is Defense Enterprise Computing Center (DECC)-compliant and can be directly connected to the DoD Non-secure Internet Protocol Router Network (NIPRnet). According to the *Cloud Computing Journal* RACE provisioning of environments takes 24 hours in development and test environments and 72 hours in production and can be done through a credit card or Military Inter-departmental Purchase Request (MIPR) payment option.[45] Helpdesk support is available 365 days per year, 24 hours per day.

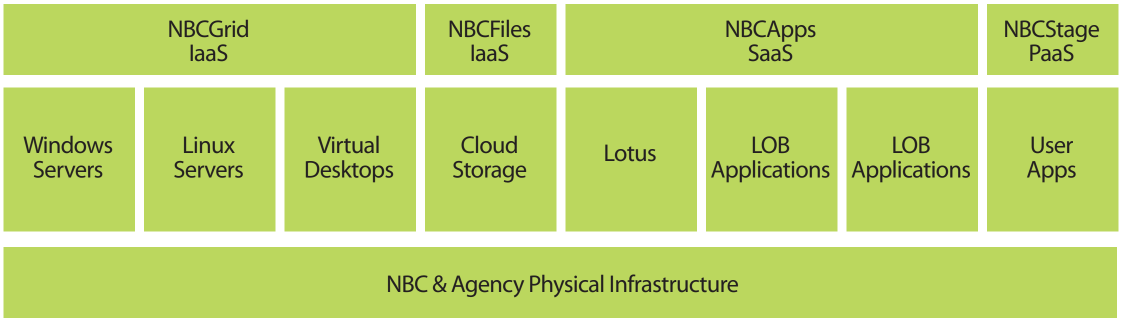

Figure 5.1-1. DOI NBC Cloud Computing Offerings

- **Department of Interior (DOI) National Business Center (NBC):** The Department of the Interior (DOI) National Business Center (NBC) plans to provide multiple offerings to federal agency customers as shown in Figure 5.1-1.[46] The current and planned offerings include: *NBCFiles* as a cloud storage offering; *NBCGrid* as an IaaS; *NBCStage* as PaaS; and *NBCApps* as an SaaS. As an example of their offerings, the *NBCGrid* offering provides "technology agnostic server hosting" and has both metered and pre-paid pricing models. For the end user, *NBCGrid* is intended to facilitate "rapid provisioning (sub 15 minute)."[47] As an example of their SaaS offerings, *NBCApps* will include human resources capabilities and acquisitions capabilities, such as an on-demand version of the Electronic Services Environment (ESE).
- **NASA–Nebula:** NASA's Nebula service will be offering Software as a Service (SaaS), Platform as a Service (PaaS), and Infrastructure as a Service (IaaS).[48] The cloud platform is based upon open source components; the technology stack is shown below in Figure 5.1-2. For example, NASA employs Eucalyptus open source software which provides APIs consistent with Amazon's Elastic Compute Cloud (EC2). This commonality of APIs helped NASA to develop tools for Nebula and can facilitate interoperability of applications for service users.[49] NASA has also considered appropriate network connectivity for its customers: "Nebula is primarily located at the Ames Research Center," writes NASA's Chris Kemp and Joshua McKenty.[50] "The Ames Internet Exchange (AIX) which hosts the cloud, was formerly 'Mae West,' one of the original nodes of the Internet, and is still a major peering location for Tier 1 ISPs."[51] This puts the service very near a central IP exchange point for network traffic, improving end-user performance.

Benefits—Using an existing Government-managed cloud service may save the consuming organization significant time and can lower program risk. For a consumer organization, usage of a Government-managed cloud offering may reduce the time to become operational by avoiding an entire contracting process required when obtaining commercial offerings, or avoiding an systems integration process for those considering building private offerings. Instead, a Government consumer may be able to comfortably leverage a cloud offering with a

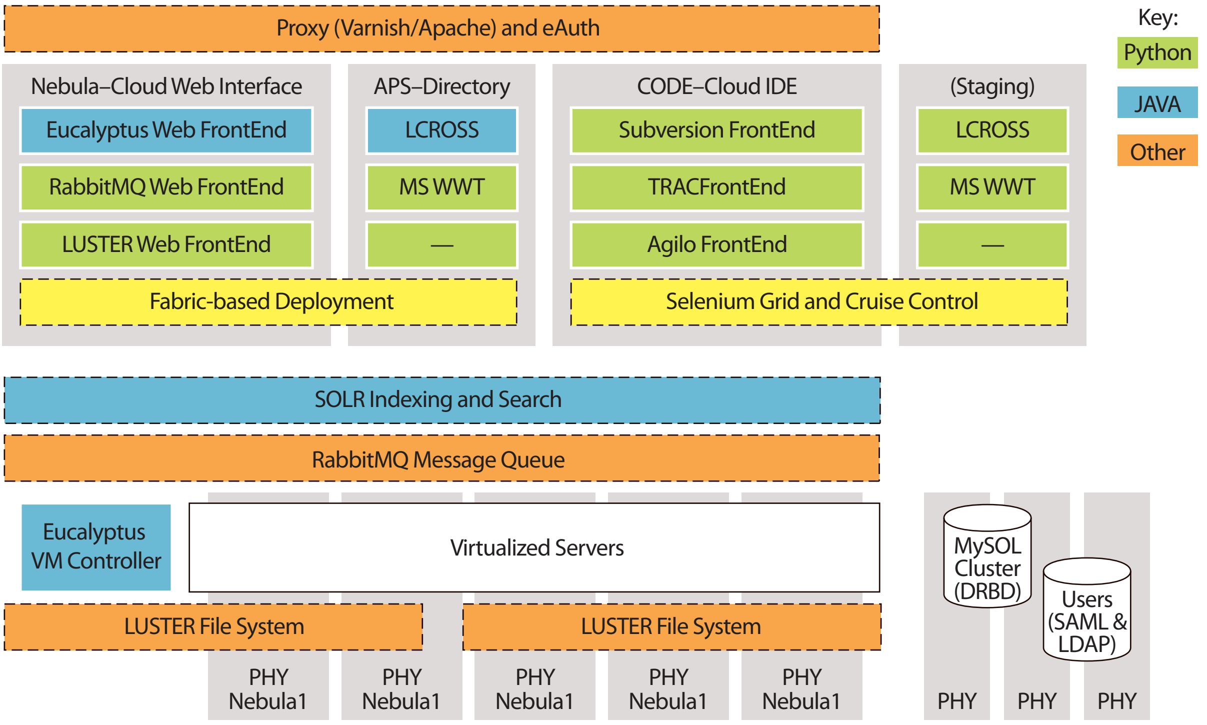

Figure 5.1-2: Nebula Technology Stack[52]

Memorandum of Understanding (MOU) and a Service Level Agreement (SLA). A Government service provider differentiates their offering by providing a customer confidence that they're acquiring services matched to the traditional needs of Government. The Government provider can offer a cloud service that is tailored to the security, records retention, and certification and accreditation needs of a Government customer. This can offer a significant advantage to the Government customer, providing a differentiated product from a potentially less costly commercial offering. For some agencies, unique Federal security requirements can make it difficult for vendors to successfully offer the Government their commercial services. Commercial providers may apply security measures that they feel are sufficient in the marketplace, though they do not completely fulfill Federal needs or policies.

***Risks*—**Table 2 defines common risks and possible mitigations for using Government-managed cloud services.

Table 2. Risks and Mitigations for using Government-managed Cloud Services

Common Risk/Concern	Mitigation Approach
Cost savings for consumers of Government-run cloud services may not be as sizable as commercial alternatives.	Currently, Government-run clouds do not run on the same scale as their commercial counterparts. Therefore, due to smaller economies of scale, the savings realized in the service consumer's business case may be less, and the payback period may be longer. Adjust business cases to accurately reflect expected Government to Government charges, not commercial costs.
Few standards currently exist to ensure portability among cloud service providers, or to ensure common detailed definitions of 'platforms' and 'infrastructure' services.	Currently, there are not many cloud-oriented standards for the consumer to use to mitigate this risk. Continue to look for standards to increase portability among service providers and to reduce lock-in. Consider evolving standards such as the Open Virtualization Format (OVF)[53] which describes standards for virtual machine descriptions. Be cognizant of development languages that are vendor specific, unless program leadership can accept the program risk of being tied to one platform.
Government Service-Level Agreements (SLAs) may not be as clearly defined as their commercial counterparts.	Obtain detailed SLAs from Government service providers and agree on mechanisms for measuring the SLAs. Document a performance feedback process in a Memorandum of Understanding (MOU), and require attendance by service provider personnel able to make commitments for the service provider organization. Define the periodicity of the feedback communications and meetings. Finally, employ a quick proof-of-concept to ensure that the Government cloud service meets implied consumer requirements and expectations.
Government cloud providers may not have the same burst capacity as large scale commercial providers.	Understand the limitations of a Government provider's burst capacity and lead time needed for provider to "scale up" to meet significant increases in usage.

5.2 Activating Government Services

Activating Government-run cloud services, is analogous to commercial-run cloud services, as both are provided by external service providers. Please see section 4.2 for a discussion of these considerations.

6.0 Creating Internal Enterprise Private Cloud Services

Creating an infrastructure for a private cloud of a size that will realize economies of scale analogous to those shown in contemporary commercial offerings, and that will effectively satisfy requirements common to mission critical Government systems, is a major undertaking, requiring substantial capital investment and comprehensive project management and systems engineering. Unlike Sections 4.0 and 5.0 where external organizations provide the cloud, this section discusses the general engineering and requirements considerations for a large-scale effort to provide the infrastructure for a Federal private cloud. This section generally assumes mission critical private clouds with sensitive or classified content.

6.1 Enterprise Cloud Infrastructure Engineering

Some organizations contemplating the creation of large-scale private clouds do not have a ready data center infrastructure sufficient to host a very large scale operation or cannot add this new capability to existing facilities. Expansion of existing facilities can be hampered by power, HVAC, or rack space constraints, or the assessment that the existing facilities are not in a physical location worth large-scale investment.

Accelerators:

- Leverage a proven IT stack (e.g., chipsets, virtualization software) that is known to work together as product integration takes substantial time.
- Anticipate issues when porting legacy applications to virtualized cloud environments. Start porting activities early.

The decision to 'build your own' and not use commercial offerings as described in Section 4.0 has significant implications on project scope and complexity. Project considerations range from the selection of a location with suitable power, transportation, zoning, taxes, workforce, and weather conditions, to design considerations such as network diversity, physical security controls, HVAC and heat dissipation and a failure proof electrical infrastructure. An organization considering the creation of data centers to create a private cloud on a enterprise scale should consider the following:

- **Location Assessments:** The data centers intending to support mission critical IT clouds should avoid areas prone to natural disasters and major population centers to be most survivable. Assessments of the frequency of recurring weather-related events such as hurricanes, floods, and tornados should be completed. Data Center News states, "Companies are shying away from the Gulf region because of the hurricanes; California for mudslides, earthquakes and wildfires; and high-profile metro areas because of the terrorist threat potential."[54] Assuming that a location is weather-neutral, properly zoned, and building codes are generally suitable for a data center, a number of additional considerations drive the choice of a data center location, including:
  - **Power Access Assessment:** Data center planning efforts should quantify the power requirements expected for a data center. A strong relationship with local power providers should be established and expectations for projected consumption should be clarified. Data Center News explains, "The first criteria that come to mind for many data center site selection experts are the price of power and its reliability."[55] With the increase of more dense IT infrastructure, power consumption per square foot of a data center is increasing. Planners should investigate and understand the ability to upgrade or install transmission lines.
  - **Physical Redundancy:** It is likely that any single location in the continental U.S. (CONUS) will be insufficient to support Continuity of Operations (COOP). For many missions more than one data center will need to be built in diverse physical locations. Some locations may be eliminated from consideration because they are too close to existing organizational efforts.
  - **Workforce:** The ability to hire and house technical staff in proximity to the data center

facility should be assessed as well as the ability to staff cleared personnel if applicable.
  - **Transportation Infrastructure:** The transportation system near the project facility should be assessed for its ability to support traffic such as construction equipment, fuel delivery vehicles and staff traffic. Planners should assess whether there are multiple routes to the potential data center, to avoid a single point of failure in the road system, due to accident or disaster, and whether the potential data center is supported by air transportation for urgent requirements.
- **Design Considerations:** Once an optimal location has been established there are a number of major design considerations in implementing the data center, which can also be cost drivers:
  - **Network Diversity:** Due to the fundamental role of the network in cloud computing, the network, and the redundancy of the data center's connections to it, will play a key role in an infrastructure effort. The ability to utilize a diversity of network providers and to route traffic among a group of networked peers is vital to a cloud service provider. A network operations center capable of monitoring and adapting to the health of the network providers is required. For some missions, satellite ground stations may also be required to add to the diversity of communications.
  - **Physical Security:** Considerations in this area include compliance with particular agency security standards and the design and implementation of a wide range of features such as video monitoring, biometric access, perimeter guards, perimeter berms and fences, RF shielding, intrusion detection systems, and fire detection and suppression.
  - **Construction Durability:** Buildings for sensitive mission critical data centers will need to be built to a survivability standard that is beyond standard building construction. Plan for this to be a cost driver in capital expenditures for construction.
  - **Electrical Infrastructure:**[56] The data center must support the use of an alternative source of power for a sustained period of time, in the event of a failure by the primary power supplier. Solutions include banks of generators and large supplies of fuel sufficient to cover the failover time period. Frequent testing of the power alternative will be required once the facility is operational. Planners should account for large scale UPS capabilities to regulate power before it is sent to IT equipment.
  - **HVAC:** Heat dissipation will be a major design and cost driver for the data center. Cloud computing infrastructures are most often built upon highly virtualized blade infrastructures which are continuing to get more dense, and generate much more heat per rack unit (U) than was common even ten years ago. A complex system of heat exchangers, compressors, and air management systems will need to be designed and implemented to achieve sufficient heat reduction. Complex and continuous monitoring will need to be implemented to protect the large investment in computing infrastructure.

***Benefits*—**The chief advantage of building a private cloud is the ability to maintain control over a wide variety of service characteristics, such as the security implementation, continuity of operations (COOP), retention policies, network connectivity and usage policies. The benefits pursued through a private cloud can be tailored to the objectives of the owner's business case. For example, if cost reduction is the primary objective, virtualization and intra-organizational multi-tenancy can be employed to reduce the hardware footprint. The business objectives may be focused on topics such as data center utilization, collaboration, COOP, speed of deployment or location-independent access, perhaps coupled with the need for high security or high mission confidence. In these cases, the flexibility provided by a private cloud implementation may be significant.

***Risks*—**Table 3 defines common risks and possible mitigations for constructing private cloud offerings.

6.2 Cloud Service Implementation

While Section 5.1 discusses establishing the physical infrastructure for the private cloud capability, this section focuses on creating the next layers of the stack which provide services for the organizations end users.

Table 3. Risks and Mitigations for Constructing Private Cloud Offerings

Common Risk/Concern	Mitigation Approach
Cost reductions may take longer to realize as compared to other cloud alternatives.	Look to other comparable Government service provider organizations to understand their integration, development, and operational costs. For example, while the hardware footprint may be significantly reduced as compared to legacy hardware, there may be new costs, such as product and systems integration, virtualization software licensing, application porting, disconnected operations capabilities, and COOP capabilities.
Legacy applications may not easily run in the new cloud environment, increasing costs and schedule time.	Port applications to test cloud environment early to ensure applications can successfully run in the target architecture cloud stack.
When creating cloud services, product integration for a complex series of interdependent components can take longer than expected.	A number of products are required to provide a robust scalable cloud service, and the time required to integrate and test all those products and ensure that they work seamlessly with each other is significant. Look to existing instantiations of Government cloud platforms to understand compatibility.
Acquisition activities in this option are not trivial and may add to the program schedule.	Leveraging a systems integrator or prime contractor may simplify the acquisition process rather than acquiring the hardware and software separately and integrating the capability in-house.

*Commercial Tools—*Government private clouds will most likely rely on commercial tools for their implementation, and therefore Commercial Off The Shelf (COTS) product acquisition will play an additional role in cloud service implementation. Product acquisitions will generally take the form of straightforward hardware and software license procurement, though support agreements and SLAs may need to be negotiated for some items.

As described previously, cloud services cover a wide range of the IT stack and the tools required to implement the cloud services will depend on the program scoping decisions recommended in Section 2.1. Some representative commercial tool examples include:

- **Server Virtualization Example: VMware:**[57] Server virtualization has become a mature option for hosting and managing applications and collections of applications within an enterprise. Currently VMware is a market leader in server virtualization for the data center, though important competitors such as Microsoft Corp., are releasing competing products. The ability to create and manage virtual machines (VMs), including the ability to backup or failover VMs between sites, is a core function of a contemporary cloud capability.
- **Distributed Applications Example: Apache Hadoop[58]/ Google's MapReduce paradigm**:[59] Tools which can take advantage of a large pool of processors and break up a computational problem into a set of smaller jobs will often be needed in a cloud capability supporting Government programs with very large data sets. Hadoop is a Java-based Apache effort that can be licensed without cost and includes a job tracker for dividing up, assigning work, and assembling work results.
- **Distributed Enterprise Server Management Example: 3tera AppLogic:**[60] Tools to construct a scalable enterprise grid from commodity hardware can be important. This example, AppLogic, allows operational VMs to be built from collections of pre-defined machine images, and then connected together into systems of systems, from a graphical interface.

Based on the scope definition of the program as defined in Section 2.1 a data center service architecture will be developed that assembles commercial tools into a comprehensive systems solution using the infrastructure assembled in Section 6.1. This work can be outsourced as a system's integration task, or developed in-house by Government staff.

*On-Demand Capability Testing—*Comprehensive testing and defect remediation is an essential element to establishing trust between Government private cloud providers and consumer programs. As part of the cloud service testing program it is important to establish the performance characteristics of the service as demand increases. For example, if on-demand resource provisioning is available to a consumer, then the private cloud provider should incorporate the testing of on-demand functions into test plans to verify service behavior. In addition, cloud providers should consider testing throughput and scalability beyond the consumer's expected standard processing needs. Testing should identify whether provider services scale appropriately in accordance with the promised service-level agreement (SLA), or whether they gracefully degrade, or break down completely at increased usage levels.

***In-Situ Cloud Testing*—**Testing network-based services that are external to an organization is a challenge. When an organization tests a system that utilizes external services, it is not likely to have complete control over a dedicated test environment that encompasses these external services. The situation is further compounded if the service provider does not have a specific service dedicated for consumers to use just for testing activities, such as systems integration testing. With external services the consuming organization may not be able to effectively conduct volume, throughput, or error testing against an operational service.[61] In contrast, a private cloud infrastructure owned by the Government offers a unique testing option, entitled *In-situ Cloud Testing (ICT).* Since cloud environments are generally built upon virtualization concepts, the concept of VM isolation is well implemented. This allows for very realistic testing environments to be built into the same cloud infrastructure as the operational systems, but in their own virtualized segment of the private cloud. Actual live VMs can be cloned from operational environments into the test environment. Test environments can run at the same classification as the real environment. The testing environment can therefore maintain maximum realism, while at the same time not affecting operational systems or compromising security policy.

***Customer Relationship Management*—**As the creator of a private cloud, the Government is in the role of service provider to other Government organizations who act as consumers. Successfully fulfilling this role implies several key activities:

- **Consumer Program Support:** First the service provider must devote resources provide support to consumer programs and program managers. This support can include SLA reporting, technical support on how to use particular network-based cloud infrastructure services, support with consumer systems testing, support with Certification and Accreditation (C&A) activities, and support with resource scaling and provisioning.
- **Enabling a Feedback Loop:** Connecting with consumers requires a feedback loop to continuously improve the cloud services being offered. Enabling this communication is a key activity of the cloud provider. Example techniques include stakeholder's meetings, online forums, and performance questionnaires.
- **Cloud Service Awareness:** Outreach is a key function of an effective service provider. Services evolve and improve over time and cloud providers must communicate these changes across all consumers. Outreach is also key to expanding the base of consumers and increasing awareness in other parts of the Government organization regarding the services available.

7.0 Conclusion

Government organizations, looking to realize the same cloud computing advantages that they see today in the commercial marketplace, have several options for obtaining the benefits of cloud computing. First, the commercial marketplace continues to grow with a wide range of network-based cloud services available for procurement and lease through traditional Federal Acquisition Regulation (FAR) processes, when the commercial service can meet the underlying requirements of the Government. Second, there are several Government run solutions that are being put in place, tailored specifically for Government customers. And finally, there is always the option to apply commercial tools to build internal private cloud capabilities completely within Government data center facilities.

This document has outlined a structured process for scoping a cloud effort, refining potential cloud requirements, developing a business case for a cloud effort, and making a selection among a range of commercial and Government cloud providers. Fundamentally, this process is focused on understanding and documenting a clear set of service requirements and objectives as an essential early step in selecting a cloud service provider and in deciding to create cloud services for an organization. The engineering process that follows, and the decisions made there, all come back to a robust understanding of the needs of the organization. Cloud computing, when combined with this structured engineering processes, can provide more efficient and agile IT resources for an organization, allowing the commoditization of aspects of IT such as storage and computation, to focus on higher order mission challenges.

Acronyms

Acronym	Definition
API	Application Programmer Interface
CapEx	Capital Expense
CFR	Code of Federal Regulations
CIO	Chief Information Officer
COG	Continuity of Government
CONUS	Continental United States
COOP	Continuity of Operations
COTS	Commercial-Off-The-Shelf
CRM	Customer Relationship Management
DCP	Defense Continuity Program
DISA	Defense Information Systems Agency
DoD	Department of Defense
EC2	Elastic Compute Cloud
FY	Fiscal Year
GB	Gigabyte
GSA	General Services Administration
HTTP	Hyper Text Transfer Protocol
HVAC	Heating, Ventilating and Air Conditioning
IaaS	Infrastructure as a Service
IDIQ	Indefinite Delivery Indefinite Quantity
IP	Internet Protocol
IT	Information Technology
LOB	Line Of Business
MOU	Memorandum of Understanding
NARA	National Archives and Records Administration
NBC	National Business Center
OMB	Office of Management and Budget
OpEx	Operating Expense
PaaS	Platform as a Service
QoS	Quality of Service
RACE	Rapid Access Computing Environment
SaaS	Software as a Service
SDK	Software Development Kit
SLA	Service-Level Agreement
SOA	Service-Oriented Architecture
SOW	Statement of Work
SOO	Statement of Objectives
TO	Task Order
VM	Virtual Machine
VPN	Virtual Private Network
WAN	Wide Area Network

References

[1] Geoff Raines, “Cloud Computing and SOA”, The MITRE Corporation, 2009.

[2] Federal News Radio, “Kundra sees innovation in the cloud”, http://www.federalnewsradio.com/?nid=35&sid=1620108.

[3] FY2010 Federal Budget, “Cross Cutting Programs”, http://www.whitehouse.gov/omb/budget/fy2010/assets/crosscutting.pdf.

[4] Geoff Raines, “Cloud Computing and SOA”, The MITRE Corporation, 2009.

[5] Peter Mell and Tim Grance, “The NIST Definition of Cloud Computing”, http://csrc.nist.gov/groups/SNS/cloud-computing/index.html

[6] Wikipedia, “Platform as a Service”, http://en.wikipedia.org/wiki/Platform_as_a_service.

[7] Robert P. Desisto et al., “Tutorial for Understanding the Relationship Between Cloud Computing and SaaS”, http://www.gartner.com/DisplayDocument?ref=g_search&id=640707.

[8] Kent Langley, “Thoughts on the Business Case for Cloud Computing”, http://www.productionscale.com/home/2009/2/4/thoughts-on-the-business-case-for-cloud-computing.html.

[9] A multi-tenant system allows multiple leasing customers to share the same underlying infrastructure, though they are logically separated.

[10] Troy Benohanian, “The Business Case For Cloud Computing”, http://blog.inm.com/index.php/2009/05/19/the_business_case_for_cloud_computing.

[11] FY2010 Federal Budget, Analytical Perspectives, Cross Cutting Programs, http://www.whitehouse.gov/omb/budget/fy2010/assets/crosscutting.pdf.

[12] National Business Center, Department of Interior, “Federal Cloud Playbook”, August 2009, http://cloud.nbc.gov/PDF/NBC%20Cloud%20White%20Paper%20Final%20(Web%20Res).pdf.

[13] Al Grasso, “Information Technology Acquisition – A Common Sense Approach”, The MITRE Corporation, March 2009, http://www.mitre.org/work/tech_papers/tech_papers_09/atl/.

[14] National Archives, “Frequently Asked Questions about Records Management”, http://www.archives.gov/records-mgmt/faqs/general.html.

[15] Geoff Raines, “Service Oriented Acquisitions”, The MITRE Corporation, April 2009.

[16] Geoff Raines, “Service Oriented Acquisitions”, The MITRE Corporation, April 2009.

[17] GSA, Request For Quote, Statement of Work #17914883, “U.S. Federal Cloud Computing Initiative.”

[18] GSA, “Cloud IT Services”, https://www.apps.gov/cloud/advantage/cloud/category_home.do?BV_SessionID=@@@@0734724691.1267641522@@@@&BV_EngineID=ccciadejkfdmjefcflgcefmdgfhdgji.0&c=IA.

[19] Michael Armbrust, Armando Fox, Rean Griffith, Anthony D. Joseph, Randy Katz, Andy Konwinski, Gunho Lee, David Patterson, Ariel Rabkin, Ion Stoica, and Matei Zaharia, “Above the Clouds: A Berkeley View of Cloud Computing” http://radlab.cs.berkeley.edu/, February 10, 2009.

[20] “Open Virtualization Format”, http://www.vmware.com/appliances/getting-started/learn/ovf.html].

[21] Amazon Elastic Compute Cloud Getting Started Guide (API Version 2009-04-04), http://docs.amazonwebservices.com/AWSEC2/2009-04-04/GettingStartedGuide/.

[22] Amazon Elastic Compute Cloud (Amazon EC2), http://aws.amazon.com/ec2/#pricing.

[23] Google App Engine, Billing and Budgeting Resources,
http://code.google.com/appengine/docs/billing.html.

[24] Google, “Google Accounts”,
https://www.google.com/accounts/NewAccount.

[25] Amazon.com, “Amazon Elastic Compute Cloud Getting Started Guide (API Version 2009-04-04)”,
http://docs.amazonwebservices.com/AWSEC2/2009-04-04/GettingStartedGuide/.

[26] Preethi Dumpala, “Microsoft’s Cloud Pricing In Line With Google’s, Amazon’s”,
http://www.businessinsider.com/microsoft-azure-vs-google-app-engine-vs-amazon-web-services-2009-7.

[27] Salesforce.com, “Salesforce Pricing and Additions”,
http://www.salesforce.com/crm/editions-pricing.jsp.

[28] Windows Azure Platform, “Register for Azure Services”,
http://www.microsoft.com/azure/register.mspx.

[29] Amazon.com, “AWS Premium Support”,
http://aws.amazon.com/premiumsupport/#overview.

[30] Salesforce.com, “Cloud Computing Model and Pricing”,
http://www.salesforce.com/platform/platform-edition/.

[31] Windows Azure Platform, “Register for Azure Services”,
http://www.microsoft.com/azure/register.mspx.

[32] Google.com, “Google App Engine”,
http://code.google.com/appengine/downloads.html#Google_App_Engine_SDK_for_Java.

[33] Amazon.com, “Amazon Elastic Compute Cloud (Amazon EC2) Pricing”,
http://aws.amazon.com/ec2/#pricing].

[34] Amazon.com, “Amazon Elastic Compute Cloud (Amazon EC2)”,
http://aws.amazon.com/ec2/].

[35] Google, “Google App Engine, Billing and Budgeting Resources”,
http://code.google.com/appengine/docs/billing.html.

[36] Amazon.com, “Amazon Elastic Compute Cloud (Amazon EC2) Transferring Large Amounts of Data”,
http://aws.amazon.com/s3/#importexport.

[37] Michael Armbrust et al., “Above the Clouds: A Berkeley View of Cloud Computing”,
http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.pdf.

[38] Steve Lawson, “Amazon S3 lets customers ship big data”,
http://www.macworld.com/article/140739/2009/05/amazon_cloud.html.

[39] Stan Schatt, Giga Information Group, Inc., “Best Practices: Global Server Load Balancing”, 2003.

[40] Network World, October 5, 2009,
http://www.networkworld.com/news/2009/100509-pentagon-cloud-computing.html?page=2.

[41] Google, “Grow Your App Beyond Free Quotas”,
http://googleappengine.blogspot.com/2009/02/new-grow-your-app-beyond-free-quotas.html.

[42] Amazon.com, “Amazon Elastic Compute Cloud (Amazon EC2)”,
http://aws.amazon.com/ec2/.

[43] DISA, “Rapid Access Computing Environment”,
http://www.disa.mil/race/.

[44] Network World, October 5, 2009,
http://www.networkworld.com/news/2009/100509-pentagon-cloud-computing.html?page=2.

[45] Cloud Computing Journal, “DISA Opens DOD Cloud”, October 7, 2009,
http://cloudcomputing.sys-con.com/node/1136332.

[46] National Business Center, Department of Interior, “Federal Cloud Playbook”, August 2009,
http://cloud.nbc.gov/PDF/NBC%20Cloud%20White%20Paper%20Final%20(Web%20Res).pdf.

[47] Cloudbook.net, “National Business Center (NBC) Cloud Computing”,
http://www.cloudbook.net/nbc-gov.

[48] NASA, “NASA's Cloud Computing Platform”,
http://nebula.nasa.gov/.

[49] NASA, “NASA Nebula blog”,
http://nebula.nasa.gov/blog/2009/nov/how-eucalyptus-enables-ec2-compatibility-with-nebu/].

[50] CloudBook.net, “NASA Nebula”,
http://www.cloudbook.net/nebula-gov.

[51] Ibid.

[52] NASA, “Nebula Services”,
http://nebula.nasa.gov/services.

[53] VMware, “Open Virtualization Format”,
http://www.vmware.com/appliances/getting-started/learn/ovf.html.

[54] Matt Stansberry, “Data Center Locations Ranked By Cost”,
http://searchdatacenter.techtarget.com/news/article/0,289142,sid80_gci1204203,00.html].

[55] Ibid.

[56] Generator photo credit: Terremark Worldwide, Inc.

[57] VMware, “Virtualization Products for Cloud Computing”,
https://www.vmware.com/products/product_index.html.

[58] Apache, “Welcome to Apache Hadoop”,
http://hadoop.apache.org/].

[59] Google, “MapReduce: Simplified Data Processing on Large Clusters”,
http://labs.google.com/papers/mapreduce.html].

[60] 3tera, “Cloud Computing Without Compromise”,
http://www.3tera.com/?_kk=3tera%20applogic].

[61] Lawrence Pizette et al., MITRE, 2009.

MITRE

www.mitre.org

Case Number: 10-1070
Document Number: MTR100092

MITRE

www.mitre.org