Office of the Chief Information Security Officer
Centers for Medicare & Medicaid Services
7500 Security Boulevard
Baltimore, Maryland 21244-1850

OFFICE OF THE CHIEF INFORMATION SECURITY OFFICER

Risk Management, Oversight, And Monitoring

Risk Management Handbook
Volume III
Standard 3.2

CMS Cloud Computing Standard

FINAL
Version 1.00
May 3, 2011

Document Number: CMS-CISO-2011-vIII-std3.2

(This Page Intentionally Blank)

SUMMARY OF CHANGES IN VOL III, STD 3.2, *CMS CLOUD COMPUTING STANDARD* VERSION 1.00

1. Baseline Version.

(This Page Intentionally Blank)

TABLE OF CONTENTS

LIST OF FIGURES

1 INTRODUCTION

Initiatives for using cloud computing in the Federal Government are emerging and evolving at a rapid pace. The Centers for Medicare and Medicaid Services (CMS), Office of the Chief Information Security Officer (OCISO) has developed this security standard to offer clear guidance for the use of cloud computing environments.

In February 2010, the White House launched the *Federal Data Center Consolidation Initiative* (FDCCI) and issued guidance for Federal Chief Information Officer (CIO) Council agencies. The guidance called for agencies to inventory their data center assets, develop consolidation plans throughout fiscal year 2010, and integrate those plans into agency fiscal year 2012 budget submissions. After an 8-month review process, it was determined that the Federal Government was operating and maintaining almost 2,100 data centers (with server utilization rates as low as 7 percent.) The FDCCI review initiated a strategy[1] aimed at reducing Information Technology (IT) infrastructure growth that includes a *Cloud First* policy for services and shrinking the number of data centers by at least 800 by 2015, with goals to:

- Promote the use of Green IT by reducing the overall energy and real estate footprint of government data centers
- *Reduce the cost* of data center hardware, software, and operations
- *Increase* the overall IT *security posture* of the government, and
- *Shift IT investments to more efficient computing platforms and technologies*.

Cloud computing reduces costs by leveraging existing IT infrastructure, increasing server utilization to 60 to 80 percent and provisioning services as needed. Cloud computing also increases efficiency and agility through automation and significantly reduces the administrative burden on internal IT resources. In addition, the need for large upfront capital expenditures and operating expenses is eliminated by purchasing cloud services on demand.

If any CMS cloud computing initiatives do not *consolidate infrastructure* (that is, combine applications and systems onto *less* infrastructure), do not *lower costs*, or do not *meet Federal security requirements*, then they are *not* meeting the intent of the Federal cloud computing initiatives.

No *High* security level data or system shall be recommended for placement into any off-premise (non-government) Cloud Service Provider (CSP)[2], and *Moderate* data or systems shall only be

[1] *The Federal Cloud Computing Strategy* is available at http://www.cio.gov/documents/Federal-Cloud-Computing-Strategy.pdf.

[2] Throughout this document, the term *Cloud Service Provider (CSP)* applies equally to the following types of entity:
1. All internal HHS and OpDiv IT Service organizations
2. Any private or commercial entity (including both non-profit and for-profit organizations) providing or hosting Cloud Computing services or applications
3. Any US Government entity providing or hosting Cloud Computing services or applications on behalf of other US Government Agencies or Departments
4. Academic or Research institutions who provide or host Cloud Computing services or applications
5. State or Local Government entity providing or hosting Cloud Computing services or applications

recommended for placement on *CSP* systems where the host infrastructure have a *FedRAMP Authorization to Operate (ATO)* at the *Moderate* level.

2 DEFINITION OF CLOUD COMPUTING

Cloud computing is a model, as defined[3] by the National Institute of Standards and Technology (NIST), for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned (by a CMS business owner[4]) and released with minimal management effort or service provider interaction. This cloud model promotes availability and is composed of five *Essential Characteristics*, three *Service Models*, and four *Deployment Models*.

2.1 ESSENTIAL CHARACTERISTICS

Cloud computing has several distinct characteristics. The typical cloud computing solution often leverages:

- Massive scale
- Homogeneity
- Virtualization
- Resilient computing
- Low cost software
- Geographic distribution
- Service orientation
- Advanced security technologies

Under the NIST definition, these elements are broken down to the following five basic characteristics.

On-demand self-service. A CMS Business owner can provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service's provider.

Broad network access. Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, laptops, and PDAs).

Resource pooling. The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a sense of location independence in that the customer generally has no control or knowledge over the exact location of the provided

[3] The NIST definition of *cloud computing* can be found at the NIST Cloud Computing workgroup page http://csrc.nist.gov/groups/SNS/cloud-computing/, or at http://csrc.nist.gov/publications/drafts/800-145/Draft-SP-800-145_cloud-definition.pdf.

[4] *Business Owner* is defined as the CMS official (CMS Group director or higher) responsible for the overall procurement, development, integration, modification, or operation and maintenance of an information system.

resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or datacenter). Examples of resources include storage, processing, memory, network bandwidth, and virtual machines.

Rapid elasticity. Capabilities can be rapidly and elastically provisioned, in some cases automatically, to quickly scale out and rapidly released to quickly scale in. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time.

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

2.2 SERVICE MODELS

Software as a Service (SaaS). The capability provided to the consumer is to use the CSP's *applications* running on a cloud infrastructure. The applications are accessible from various client devices through a thin client interface such as a web browser (e.g., web-based email). The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings. Examples may include, but are not limited to:

- Gov-Apps (http://www.apps.gov/, http://apps.usa.gov/, etc.)
- Internet Services
- Blogging/Surveys/Twitter
- Social Networking
- Information/Knowledge Sharing (Wiki)
- Communication (email, Instant Messaging)
- Collaboration (e-meeting)
- Productivity Tools (office)
- Virtual desktop
- Enterprise Resource Planning (ERP)
- Customer Relationship Management (CRM)

Platform as a Service (PaaS). The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages and tools supported by the CSP. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly application hosting environment configurations. Examples may include, but are not limited to:

- Application Development, Data, Workflow, etc.
- Security Services (Single Sign-On, Authentication, etc.)
- Database and Database Management (DBMS)

- Directory Services
- Testing and Developer Tools
- Middleware (Web MQ, WebSphere, etc.)

Infrastructure as a Service (IaaS). The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over; operating systems, storage, deployed applications, and possibly limited control of select networking components (e.g., host firewalls). Examples may include, but are not limited to:

- Mainframes
- Mid-tier Servers
- Storage
- IT Facilities/Hosting Services
- Virtual Machines
- Networking (Networx)

These service models essentially provide some or all of the total IT support necessary to deploy an IT solution. Depending on the scope of the service model selected, business owners will be able to avoid the *details* associated with some portion of that total IT support necessary to deploy a CMS system. Figure 1 illustrates the differences in scope and control between the cloud subscriber and CSP, for each of the service models. In general, the higher the level of support available from a CSP, the more narrow the scope and control the cloud subscriber has over the system. The two lowest layers shown denote the physical elements of a cloud environment, which are under the full control of the cloud provider regardless of the service model. Heating, ventilation, air conditioning (HVAC), power, communications, and other aspects of the physical plant comprise the lowest layer, the facility layer, while computers, network and storage components, and other physical computing infrastructure elements comprise the hardware layer. The remaining layers denote the logical elements of a cloud environment. The virtualized infrastructure layer entails software elements, such as hypervisors, virtual machines, virtual data storage, and supporting middleware components used to realize the infrastructure upon which a computing platform can be established. While virtual machine technology is commonly used at this layer, other means of providing the necessary software abstractions are not precluded. Similarly, the platform architecture layer entails compilers, libraries, utilities, and other software tools and development environments needed to implement applications. The application layer represents deployed software applications targeted towards end-user software clients or other programs, and made available via the cloud.

Figure 1 Scope of Cloud Service Models

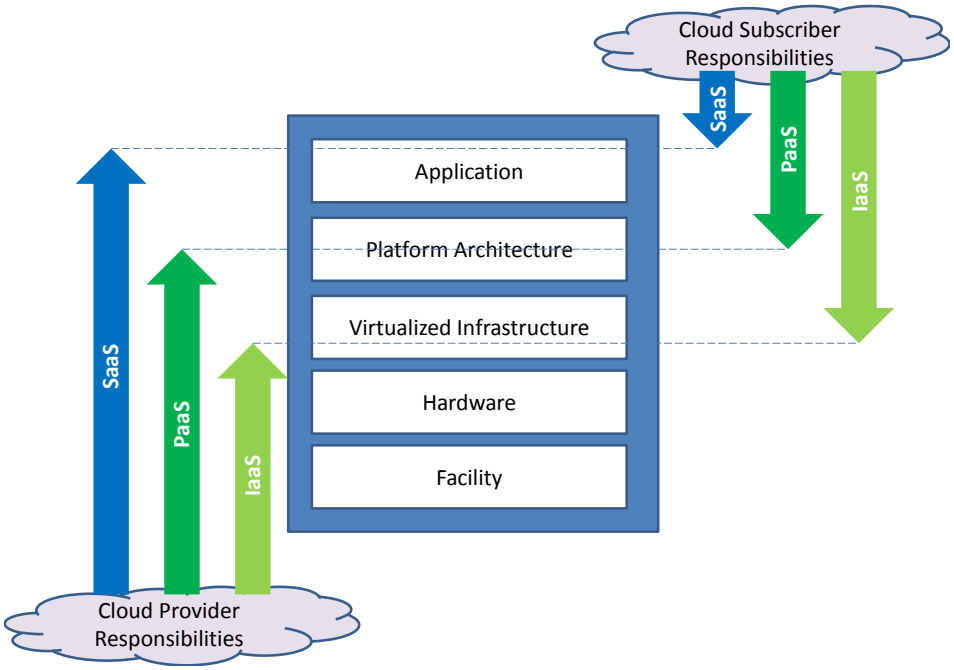

2.3 DEPLOYMENT MODELS

There are three primary cloud deployment models. Each can exhibit the previously listed *characteristics* and *service models.* Their differences lie primarily in the level of access by other CSP customers to those same services and infrastructure.

Private cloud. The cloud infrastructure is operated solely for an organization. It may be managed by the organization or a third party, and may exist on-premise or off-premise.

> *For example, CMS could create a private cloud (either internally or through an external CSP). To enable business owners the capability to utilize this cloud, CMS could create a charge-back pricing schema or a common contract vehicle for the different CMS business owners that use the CMS cloud. This would allow the different CMS organizations to gain access to the IT resources they need, while at the same time allowing CMS to create a sustainable support model for that cloud. The private cloud (including all of the infrastructure, platforms, and support services) would be dedicated for use only by CMS business owners.*

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

> *For example, a consortium of Department of Health and Human Services (HHS) Operating Divisions (OpDivs) or other Government agencies could create a community cloud.*

Payments for use of the cloud might be made via payments through a common contract vehicle. The community cloud (including all of the infrastructure, platforms, and support services) would be dedicated for use only by members of the community.

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

A public cloud is made available to the public for use, and is the most common type of commercially available cloud. Users of a public cloud typically sign up directly with the CSP (usually through a web interface) and make payments based on the provider's pricing schema, and according to the provider's established Terms of Service. Examples include Amazon Web Services, Google applications, SurveyMonkey, or similar cloud-based services.

CMS business owners should ***avoid*** *commercial-grade* public clouds, as there are not enough assurances for security and privacy to meet federal security and privacy requirements. In addition, public cloud CSPs typically *contractually preclude* their customers from access or knowledge of the inner-workings (*proprietary* processes) of their infrastructure, making it impossible to achieve security and privacy compliance.

Hybrid cloud. The cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load-balancing between clouds).

For example, cloud bursting (which occurs when a system or service is too large to be maintained in just a single cloud model) could allow a CMS system to spill over into another (different) cloud during periods of high use (such as an annual enrollment period.) Another example might be for part of a CMS system to live inside one CSP, while the remainder lives in another.

For CMS, Hybrid cloud systems that include a *commercial-grade* public cloud component will likely not be considered for an ATO due to inherent non-compliance with federal security and privacy requirements (see *Public* cloud above.)

Figure 2 Deployment Models

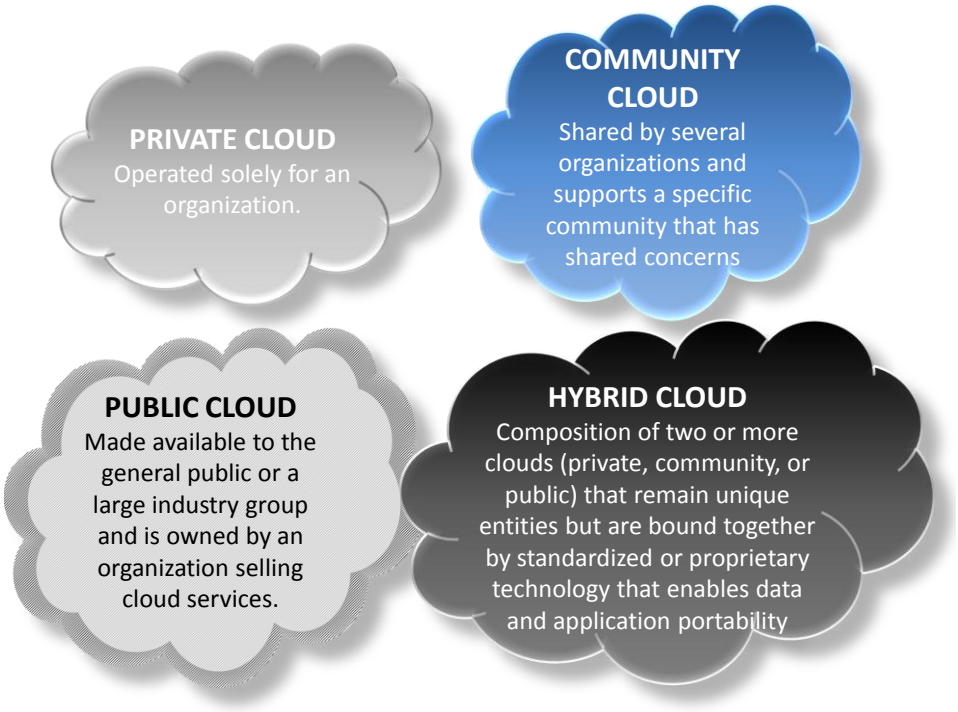

2.4 CLOUD BROKERS

Since there are many cloud service models and providers to choose from, it can be overwhelming for business owners to attempt to navigate through all of the available providers and analyze everything they offer. So, just like insurance brokers, loan brokers, and others, *cloud brokerages*[6] have emerged to seek the best deal for their customer—matching the customer's wants and needs with the best solution available. In cloud brokerage, that may include finding the right infrastructure, the right security solution, the right service levels, and more, at the right price. An effective cloud broker may know what is out there and assist business owners to monitor the current technologies.

There is a growing demand for cloud brokers as intermediaries between end users (such as CMS business owners) and CSPs. From Service Level Agreements (SLAs) with multiple vendors, to compliance and security, a broker handles many cloud related issues for a customer. This approach also enables business owners to switch cloud vendors without worrying about many of the operational details. Some vendor experience in delivering multiple services with stringent SLA requirements, strong enterprise presence, and long lasting relationship with existing government IT departments is a primary requirement when looking for a potential broker for CMS cloud services.

[6] A *cloud broker* is a single point-of-contact for an enterprise for all cloud computing requirements such as service provisioning, service level agreements (SLA) and compliance. A broker sits between the enterprise and multiple cloud service vendors and provides a layer of abstraction.

As such, cloud computing at HHS will utilize the existing governance structure, with the addition of a proposed HHS *Cloud Computing Governance Council (CCGC)* responsible for coordinating cloud computing initiatives department-wide. The HHS CCGC reports to the HHS CIO Council.

The HHS CCGC functions as the HHS point-of-contact for the *Federal Cloud Computing Advisory Council (FCCAC)*, developing HHS-specific cloud computing guidance, recommendations, and policies, taking input from the FCCAC, and applying it to HHS. The HHS CCGC will also review and reach consensus on input from the HHS CIO and Chief Technology Officer (CTO) councils and HHS Enterprise Architecture (EA). The CCGC disseminates this information and artifacts as guidance and/or policy, following the established HHS governance process for approval.

Individual HHS OpDivs have the option to establish internal OpDiv Cloud Computing Governance Councils, or working through their existing governance structures. In either instance, the OpDiv governing bodies will be represented on the HHS CCGC. To date, CMS has *not* established a CMS CCGC.

The HHS CCGC will maintain and update the HHS list of approved cloud providers and their respective rates for standard service offerings. Additionally the CCGC will develop standard RFI templates to be used by the HHS CCGC and the OpDiv CCGC to solicit standard and custom cloud service rates for the HHS approved cloud provider list and OpDiv contract vehicles.

3 CLOUD PROS AND CONS

3.1.1 PROS

3.1.1.1 FAST START-UP

Cloud computing allows business owners to test their business plans very quickly for relatively little start-up cost. Business owners dealing with legislative mandates that appear with little or no time to react may want to consider how to use cloud computing in their business plan. Cloud computing environments, even considering the built-in constraints of the various service models, allow for relatively quick provisioning, prototyping, and actualization of new systems. The cloud computing option can allow developers to (mostly) circumvent the old hardware-requisition process. The benefit of the quick provisioning of new systems allows additional time to be spent on system design and testing. With cloud computing, IT staff can set up a new server; test new services, applications or models; tweak them until they are ready to go live on the production servers; and then tear down the virtual server in the cloud. There are no lengthy delays waiting for a test box to be delivered and configured, and no hefty expenditures on new equipment and host software.

3.1.1.2 FLEXIBLE SCALABILITY

The business of CMS necessitates various *peaks and valleys* in service utilization. Due to the various limited *enrollment* periods, or *update* periods for various Medicare services, systems

may be required to have enormous excess capacity built-in to deal with various sporadic or planned high-use periods. Business owners should consider the variability of the resource utilization of their IT structure to determine if they are a good candidate for cloud services. It may be significantly less costly for business owners to outsource some or all of the services required to deal with peaks and valleys in the CMS system business model.

A well-researched and planned deployment in the cloud (with a well-formed *Service Level Agreement [SLA]*) can be very cost-effective for systems with significant variability during known windows of operation.

3.1.1.3 REDUCED CAPITAL EXPENDITURES

Cloud computing services can allow a business owner to shift funding from *Capital* to *Operations and Maintenance (O&M)*.

With cloud computing, business owners can have zero capital expenditures, despite adding much greater scalability to their system infrastructure. Business owner staff retains control over the computing resources while running applications on the CSP's proven infrastructure. In addition, business owners can acquire the additional capacity they need, when they need it, and correspondingly decommission it when it is no longer required.

3.1.1.4 COLLABORATION

Collaboration can be one of the most important advantages of cloud computing. Multiple users, from all across the nation, can collaborate more easily on documents and projects. Because the information is hosted in the cloud, and not on individual computers, business owners can collaborate with non-CMS stakeholders in a secure CSP environment with nothing more than an Internet connection and some identity management controls.

3.1.1.5 BUSINESS CONTINUITY

An additional benefit of cloud computing is reliable disaster recovery. Even if a server in the cloud fails, the CSP's redundant network can (if effectively designed into the procurement) keep a business owner's applications available. Using virtual servers and automated clustering and redundancy, a CSP can provide a high-level of business continuity with very little additional cost. Through the use of a CSP and its associated SLAs, business owners can procure not only the IT infrastructure necessary to provision and operate their CMS system, but also the required business continuity assurances (through an effective SLA) to maintain their business at the level of continuity that is required. Continuity planning no longer includes a large *capital expenditure,* but instead is procured as a *service*.

3.1.1.6 SECURITY

There *are* security advantages to moving into the cloud. From an *enterprise risk* view, shifting *public* data to an external cloud reduces the exposure of internal CMS *sensitive* data. By moving *public* data into the cloud, business owners can eliminate public portals at the mission-essential core CMS data centers. As a result, the overall exposure of the CMS enterprise is reduced.

3.1.2 CONS

3.1.2.1 HUMAN LIMITATIONS

Exploring cloud-computing models requires an adventuresome spirit and some technical astuteness. Achieving the benefits of cloud computing requires a willingness to stretch and learn, and adapt to new methods and constraints. Without the right mindset, taking on cloud computing can (and will) be very frustrating.

Operating in a cloud environment requires a dramatic change in systems development and design thinking. All of the advantages in cloud computing are achieved by the CSPs though operational efficiencies. The methods used to achieve these efficiencies can be summarized simply as ***mandated enforcement of standardization.*** While this may sound innocuous, the business drivers to standardize means that system developers and designers lose some of the flexibility in design that they may be accustomed to in the traditional CMS systems development environment. The more comprehensive the *cloud service model* (IaaS, PaaS, or SaaS), the more *stringent* the enforcement of standardization, and thus the greater the design constraints. Business owners must consider these constraints before committing to the use of cloud services to ensure their business requirements, and associated security requirements, can be fulfilled within the available service model.

When operating in cloud services models, *customization* cost money—the more profound the customization, the higher the applicable costs. Because typical CMS software development contracts are procured without significant (contractual) regard for long-term maintenance cost of the infrastructure, CMS developers tend to build applications with unique and customized solutions (often marketed as *discriminators* or *differentiators.)* Business owners should be vigilant of developer customizations and monitor system design very closely to ensure developers do not stray from the (cost-effective) service model in which their product must ultimately operate. Developers should be made *contractually* aware of the additional costs of customization that cannot be cost-effectively deployed in the target cloud environment.

3.1.2.2 FOCUS ON LONG-TERM PLANNING

Cloud computing services can allow a business owner to shift funding from *Capital* to *O&M*. However, it should be noted that these savings may *only* be achieved through *economies of scale*. Business owners who purchase a separate CSP for each system development project cannot maximize those savings.

Cloud services are a *service,* and should be treated as such. Buying a new *phone service,* from a *different phone-service provider,* every time you need to add a phone line does not make sense. Neither does procuring additional CSPs for each new system. Business owners should endeavor to consolidate and utilize existing CMS CSP contracts when looking to place systems into a cloud. In addition, contractor proposals for new applications that *include* utilization of a new CSP (vice leveraging existing contracted CSPs) are also not meeting the government objectives of *consolidation* and *cost savings*.

3.1.2.3 BANDWIDTH COSTS

Business owners considering an external cloud service to address storage scalability challenges may want to think again. The network bandwidth needed to deploy and continually access large-scale data repositories in an external cloud can be tremendous. The cloud service cost may be so great that business owners might be better served to buy and host the storage capacity internally (co-located with the application) rather than paying an external cloud provider for it.

Conversely, cloud-hosted applications that require large-scale access to *CMS-hosted* data, or data hosted at *other* CSPs (hybrid clouds) may also require high-bandwidth services. These costs should be carefully quantified when evaluating the use of cloud services.

3.1.2.4 AUTOMATIC SCALING

Cloud services can provide business owners with a simple method to automatically scale their application hosting. However, some problems can arise that may be financially distressing. For example, the ability to automatically scale an application may tend to make developers inattentive the issues that lead to improper resource utilization, leading to increased hosting costs. Worse, inattentiveness by developers to the inherent risks associated with automated scaling can lead to significant incurred costs resulting from low-level exploits of their application. As an example are systems developed without proper controls to identify and deal with unusual activity. Hackers utilizing a simple low-level Distributed Denial of Service (DDoS) attack, which may not take a site down, but will keep the server very busy, may have a significant financial impact on business owners. Since business owners pay CSPs for *usage*, costs can spin wildly out of control with little or no visible change in system performance. Business owners should be aware that CSPs will not be held (contractually) responsible for hosting poorly designed CMS systems. If business owners promote systems to be developed and deployed in the cloud, they must be prepared to deal financially with the potential issues associated with unintended or malicious scaling of services.

3.1.2.5 PERFORMANCE

The use of cloud computing does not necessarily lead to improved application performance, specifically when dealing with network latency. Applications that perform poorly in latency-sensitive conditions are not good candidates for external cloud computing unless the *entire business application* (or at least the latency-sensitive portion) is hosted within the same cloud infrastructure.

3.1.2.6 DATA SECURITY AND PRIVACY

If business owners are looking to achieve and maintain data privacy requirements for the Privacy Act of 1974 (Privacy Act), Health Insurance Portability and Accountability Act (HIPAA), Sarbanes-Oxley Act of 2002 (SOX), E-commerce, and so on, they must be extra vigilant when entering into contracts with CSPs. Business owners need to ensure that the government's rights to data are preserved, and that all federally mandated security and privacy standards are continually met.

3.1.2.7 INCIDENT RESPONSE

When business owners move into the cloud, they are committed to trusting CSP's security model. As a CSP customer, CMS has a reduced capability to respond to audit findings and incidents that involve a CSP. Additionally, CMS may also have a reduced capability to obtain support for forensic investigations. Since much of a CSP offering may include (what the CSP considers) proprietary implementations, CMS may not have the ability to examine those portions of the CSP offering to ensure compliance or security. In fact, many of those restrictions may even be contractually stipulated by the CSP.

Cloud computing alters the scope of *incident response*. Similar to many kinds of outsourcing, business owners give up control over the physical infrastructure, and as a result, the immediate response to incidents suddenly becomes much more complex, as the CSP, possibly with competing priorities, has now become an integral member of the response team. CSP contractual vehicles should include language mandating response times and quality of response (in the SLA) during defined incidents and should aggressively focus on communications between the CMS and CSP incident response teams.

4 CMS CLOUD COMPUTING SECURITY

In all large enterprise organizations, including CMS, it is imperative to use IT assets to their utmost effectiveness and efficiency. The development, maintenance, and secure-operation of IT assets, as well as the protection of sensitive government and citizen information, is the responsibility of any business owner considering the use of cloud computing resources. It should be noted that the requirements of the *Federal Information Security Management Act of 2002* (FISMA) still firmly apply to all cloud computing environments. The cloud portion of a CMS system may receive an Authorization to Operate (ATO) from the General Services Administration (GSA) or the *Federal Risk and Authorization Management Program (FedRAMP)*[7]. However, cloud computing services are nothing more than a General Support System (GSS). And, like any other system hosted on a separate GSS, the CMS Office of the Chief Information Security Officer and the CMS Chief Information Officer ***must*** approve the GSS for operation, as well as approve an ATO for the *remainder* of the CMS systems (Applications) hosted on that GSS.

All systems deployed in a cloud environment must be evaluated for an ATO using the same security standards and requirements as those deployed in traditional environments, and must be maintained *only* within those authorized environments. Directives *encouraging* the use of cloud computing environments do not (and cannot) *waive* any responsibility to meet FISMA requirements, nor any other federal mandates, statutes, or requirements. NIST is developing Special Publication (SP) 800-144 *Guidelines on Security and Privacy in Public Cloud Computing* to assist US Government agencies in making decisions about moving sensitive government data to the cloud. The Department of Health and Human Services (HHS) has also developed the *HHS Cloud Computing Implementation and Governance, Alternative Analysis and Supporting Process* (Version 1.0, dated March 8, 2011) to assist CMS in transitioning

[7] Additional information on the *FedRAMP* program can be found at http://www.FedRAMP.gov.

appropriate systems into a cloud environment. It is essential that decision-makers consider these requirements before committing to apply a specific cloud-computing model to support a CMS mission capability. CMS has the responsibility to ensure that cloud-based solutions are safe and secure. Business owners should carefully consider security needs across a number of dimensions, including but not limited to:

- ***Statutory compliance*** to laws, regulations, and agency requirements
- ***Data characteristics*** to assess which fundamental protections an application's data set requires
- ***Privacy and confidentiality*** to protect against accidental and nefarious access to information
- ***Integrity*** to ensure data is authorized, complete, and accurate
- ***Data controls and access policies*** to determine where data can be stored and who can access physical locations
- ***Governance*** to ensure that CSPs are sufficiently transparent, have adequate security and management controls, and provide the information necessary for CMS to appropriately and independently assess and monitor the efficacy of those controls.

No *High* security level data or system shall be recommended for placement into any off-premise (non-government) CSP, and *Moderate* data or systems shall only be recommended for placement on *CSP* systems where the host infrastructure have a (cloud-specific) *FedRAMP, HHS, or CMS ATO* at the *Moderate* level.

4.1 LAWS AND REGULATIONS

For U.S. federal agencies, the major security and privacy compliance concerns include the Clinger-Cohen Act of 1996, the Office of Management and Budget (OMB) Circular No. A-130, particularly Appendix III, the Privacy Act of 1974, and FISMA. Also of importance are National Archives and Records Administration (NARA) statutes, including the Federal Records Act (44 U.S.C. Chapters 21, 29, 31, 33) and NARA regulations (Title 36 of the Code of Federal Regulations, Chapter XII, Subchapter B).

FISMA requires federal agencies to adequately protect their information and information systems against unauthorized access, use, disclosure, disruption, modification, or destruction. That mandate includes protecting information systems used or operated by an agency or by a contractor of an agency or other organization on behalf of an agency. That is, any external provider handling federal information or operating information systems on behalf of the Federal Government must meet the same security requirements as the source federal agency. The security requirements also apply to external subsystems storing, processing, or transmitting federal information and any services provided by or associated with the subsystem.

Under the Federal Records Act and NARA regulations, agencies are responsible for managing federal records effectively throughout their lifecycle, including records in electronic information systems and in contracted environments. If a contractor holds federal records, the contractor must manage them in accordance with all applicable records management laws and regulations. Managing the records includes secure storage, retrievability, and proper disposition, including transfer of permanently valuable records to NARA in an acceptable format.

Other government and industry-association requirements, such as the HIPAA, may apply to a particular organization. For example, CMS falls under HIPAA standards for private and public health care facilities and applies to both employees and contractors. HIPAA requires both technical and physical safeguards for controlling access to data, which may create compliance issues for some CSPs. In many cases HIPAA requires extensive data and records retention (even longer than NARA) when dealing with *Personal Health Information (PHI)* for both *data* and *metadata* (such as audit logs, access logs, and records of disclosures.) CMS business owners should ensure that these additional requirements are addressed when dealing with CSPs that may have no provisions for these additional requirements in their base services—***especially*** when dealing with *FedRAMP-authorized CSPs* that originate from agency sponsors that do not have (and thus do not *account for*) these additional data-protection requirements.

Electronic discovery involves the identification, collection, processing, analysis, and production of electronic documents in the discovery phase of litigation. Business owners have incentives and obligations to preserve and produce electronic documents, such as complying with audit and regulatory information requests, and with Freedom of Information Act (FOIA) requests. Documents not only include electronic mail, attachments, and other data objects stored on a computer system or storage media, but also any associated metadata, such as dates of object creation or modification, and non-rendered file content (i.e., data that is not explicitly displayed for users). Business owners need to ensure that these discovery capabilities are preserved when utilizing cloud services.

One of the most common compliance issues facing an organization is data location. Use of an in-house computing center allows CMS to structure its computing environment such that they know in detail where data is stored and what safeguards are used to protect the data. In contrast, a characteristic of many cloud-computing services is that detailed information about the *actual* location of data is unavailable or not disclosed to CMS. This situation makes it difficult to ascertain whether sufficient safeguards are in place and whether legal and regulatory compliance requirements are being met. When information crosses *international* borders, the governing legal, privacy, and regulatory regimes can be *very* ambiguous and raise a variety of concerns. Among the concerns to be addressed are: 1) whether the laws in the jurisdiction where the data was collected permit the trans-border flow of sensitive information, 2) whether those laws continue to apply to the data post-transfer, and 3) whether the laws at the destination present additional risks or benefits to CMS. Technical, physical, and administrative safeguards, such as access controls, often apply. For CMS, ***all*** cloud computing services ***must*** be hosted exclusively within U.S. legal jurisdiction (i.e., *physically* within the United States.)

The degree to which CSPs accept liability for exposure of content under their control remains very unclear. In fact, CSPs will seldom step up and *offer* solutions to these complex issues simply because that then places them at some liability if the solution fails. None-the-less, business owners are *always* ultimately accountable for the security and privacy of data held by a cloud provider on their behalf, and must make provisions for these sometimes onerous issues.

4.1.1 PRIVACY

In a November 3, 2010 Memorandum, OMB encouraged agencies to *"seek new approaches for identifying and sharing high-value data responsibly and appropriately."* The purpose of the

Sharing Data While Protecting Privacy Memorandum (M-11-02)[8] is to direct agencies to find solutions that allow data sharing to move forward **in a manner that complies with applicable privacy laws, regulations, and policies**.

In the Memorandum, OMB reminds agencies that while sharing data is encouraged and beneficial, it also must be done in a way that fully protects individual privacy. Hence, agency data sharing must comply with the Privacy Act of 1974 and other applicable privacy laws, regulations, and policies.

For business owners that are considering the use of CSPs for the collection, processing, storage, or transfer of PII or PHI, the following Privacy Impact questions[9] ***must*** be addressed ***in the contract*** with any CSP under consideration:

1. *Does the CSP have a secure environment, federally authorized to at least the standards of confidentiality and integrity from the **Moderate** FIPS-199 level to store records containing PII?*

   If not, the CSP shall not be considered for use with PII or PHI data. The Cloud provider must secure data pursuant to NIST 800-53 requirements.

2. *Does the Cloud provider have the ability to alter **Terms of Service** or contracts without the express written consent of the customer agency?*

   If so, the CSP shall not be considered for use (with PII or PHI data *or any other* federal records data.) The data belongs to the Federal Government. Business owners cannot enter into contracts that may forfeit the Federal Government's exclusivity of ownership.

3. *Will the ownership of data remain under the sole ownership of the Federal Government at all times?*

   If not, the CSP shall not be considered for use (with PII or PHI data *or any other* federal records data.) The data belongs to the Federal Government. Business owners cannot enter into contracts that forfeit the Federal Government's exclusivity of ownership.

4. *Will backup information be returned to the Federal Government in the event the contract is ended or the Cloud provider files for bankruptcy?*

   This item needs to be explicitly addressed in the CSP contract.

5. *Is there a documented process to address the removal and control of agency information upon the termination of the contract between the agency and the cloud provider?*

   If not, the CSP shall not be considered for use with PII or PHI data. This item needs to be explicitly addressed in the CSP contract.

6. *Can the cloud provider utilize any data stored on their systems for any purpose outside agency use?*

[8] M-11-02 is available at http://www.whitehouse.gov/sites/default/files/omb/memoranda/2011/m11-02.pdf
[9] These questions were derived from questions developed by GSA in support of the FedRAMP program. The GSA questions are available at https://sites.google.com/a/fedramp.gov/gsa-iaas/templates/Privacy_Impact_Assessment_Questionnaire.docx?attredirects=0&d=1.

If so, the CSP shall not be considered for use with PII or PHI data. This item needs to be explicitly addressed in the CSP contract.

7. *Does the contract contain language to restrict the sharing of privacy data with any entity not explicitly authorized in the contract?*

   If not, the CSP shall not be considered for use with PII or PHI data. This item needs to be explicitly addressed in the CSP contract.

8. *Does the contract contain language to restrict the storage, transfer, or processing of privacy data to only facilities that fall under the legal jurisdiction of the United States?*

   If not, the CSP shall not be considered for use with PII or PHI data. This item needs to be explicitly addressed in the CSP contract.

9. *What controls are in place to prevent the misuse of data by those having access?*

10. *Does the cloud provider allow for access to data as permitted under current federal law to both authorized federal agencies and individuals wishing to verify their own PII?*

11. *While the data is with the cloud provider, what are the requirements for determining if the data is sufficiently accurate, relevant, timely, and complete to ensure fairness in making determinations?*

12. *Describe what privacy training is provided and who is responsible for protecting the privacy rights of the users of the cloud?*

13. *How does the cloud provider facilitate response to FOIA requests?*

14. *Is there a complete and documented process to report and handle breaches?*

15. *Describe the process that the CSP will use to report, within 1-hour, any potential privacy or security breaches to the agency regardless of whether the breach was intentional or inadvertent.*

16. *Describe the specific redress actions that the agency can take against the cloud provider in the event of a breach.*

4.2 FEDERAL RISK AND AUTHORIZATION MANAGEMENT PROGRAM

4.2.1 FEDRAMP CLOUD SERVICE PROVIDERS

The *Federal Risk and Authorization Management Program* or FedRAMP, is being established (as of the date of this publication) to provide a standard approach to *Assessment and Authorization (A&A)* of cloud computing services and products. FedRAMP allows joint authorizations and continuous security monitoring services for Government and Commercial cloud computing systems *intended for multi-agency use*. Joint authorization of cloud providers results in a common security risk model that can be leveraged across *all* of the Federal Government. The use of a common security risk model ensures that the benefits of cloud-based

technologies are effectively integrated across the various cloud-computing solutions currently proposed within the government. The common risk model will enable the government to *approve CSPs once*, and *use them often*, by ensuring multiple agencies gain the benefit and insight of the FedRAMP's ATO and access to CSP's authorization packages. What this means to CMS business owners is that by using FedRAMP authorized CSPs, they will be assured they already meet the minimum FISMA standards necessary to achieve an ATO for the associated CSP service.

Using FedRAMP-authorized CSPs will allow the CMS business owner to focus their ATO efforts on the remaining portion of the system (not covered under the CSP services.) Because the cloud-based services are vetted and authorized by FedRAMP, CMS does not need to conduct its own risk management program on those cloud services covered under the FedRAMP ATO. This allows CMS business owners to leverage those CSP-provided controls as *Common Controls*[10] in the CMS ATO process. This reduces duplication of effort, the time involved in acquiring services, and overall costs. However, CMS business owners considering the use of cloud services are still encouraged to *further* evaluate the *quality of* FedRAMP services based on their operational needs as well as any additional privacy and security needs that may not be covered in the base FedRAMP authorized services.

The existence of a FedRAMP ATO does not replace the existing CMS A&A process. Rather, it provides a set of inputs to allow for standardized and consistent evaluation of CSP offerings. This simplifies the A&A process for external offerings and shortens the timeframe for granting an Authority to Operate (ATO) for those offerings.

Vendors typically cannot directly request FedRAMP authorization. In order to be evaluated, an agency must sponsor the vendor's service and submit it to FedRAMP for review by a joint authorization board (JAB). In the case of cloud services, the JAB consists of senior executives and technical staff members from the Defense and Homeland Security departments, the General Services Administration, and the sponsoring agency. While FedRAMP is intended to be a government-wide initiative, individual agency involvement is voluntary.

GSA released a draft version of FedRAMP security controls in October 2010 with the intention of issuing the first version by the end of December 2010. However, after reviewing public comments, the federal CIO, GSA, and other officials have decided to take additional time to ensure that critical issues are properly addressed. As such, GSA extended public comments to January 2011. As of the date of this document, FedRAMP is slated for release by the end of the summer 2011.

The HHS Cloud Computing Governance Council (CCGC) will maintain and provide a list of approved CSPs, with coordination with the HHS CISO for Security and Assessment and Authorization (A&A) purposes. When this list is made available, no cloud-hosted service will use a CSP that is not on this list. (Until this list is available, all CSPs will be evaluated as non-FedRAMP providers.)

[10] NIST defines a *common control* as "A security control that is inherited by one or more organizational information systems." They also define security *control inheritance* as "A situation in which an information system or application receives protection from security controls, or portions of security controls (Hybrid controls) that are developed, implemented, assessed, authorized, and monitored by entities other than those responsible for the system or application; entities either internal or external to the organization where the system or application resides."

4.2.2 NON-FEDRAMP CLOUDS AND CLOUD SERVICE PROVIDERS

"We're FISMA compliant", or *"We've completed a SAS 70 audit"*, is one of the first things commercial CSP vendors will tout to potential government customers.

However, being *"FISMA-compliant"* can *only* be judged through FedRAMP, or through some other sponsoring federal agency (such as GSA or CMS) issuing an ATO—there are no commercial equivalents of a federally issued ATO. In addition, simply because a CSP has been granted a federal ATO for one (or more) implementation(s) of a cloud service, does not mean that *ALL* of their cloud services are defacto covered under that ATO. Many CSPs provided a wide variety of services, of which most may not be specifically covered under an active federal ATO. CMS business owners should always ask, *"What specific offering is covered under any advertised FISMA ATO?"*, *"Is that ATO still current?"*, and *"Will the CSP openly disclose to applicable CMS organizations, **all** relevant information regarding those ATOs?"*

Statement on Auditing Standards Number 70: Service Organizations (SAS 70) is a commercial auditing standard that represents that a *service organization* has been through an in-depth audit of *their* (*self-defined*) control objectives. Completing a SAS 70 audit is more of a *self-imposed* exercise by the subject organization. Since the control objectives are tightly defined by the *auditee*, a Type I SAS 70 audit means virtually nothing as a true evaluation of controls present, because it only determines if those controls *exist*. Whereas a Type II will *at least* include the auditor's opinion of the *effectiveness* of those self-defined controls (but again, only toward meeting the auditee-defined control objectives.) Having a SAS 70 audit *conducted* does not necessarily mean that an organization has *"passed"*, or even corrected any deficiencies. The SAS 70 standard was not crafted with cloud computing in mind, but vendors will attempt to use it as a stand-in benchmark.

A better benchmark *may* be ISO 27001, which is a *commercial* information *security* specification published by the *International Organization for Standardization (ISO)*. ISO 27001 is a comprehensive standard that covers many of the operational security aspects of the CSP of which the CMS business owners should be concerned. However, an ISO 27001 certification is only an *indicator* of due diligence, and does not *replace* the first-hand testing requirements of FISMA compliance; nor does it necessarily follow-up on identified deficiencies.

Ultimately, the greatest advantage of FedRAMP-authorized CSPs is that they will already have been contractually bound to adhere to applicable federal laws for data retention, logging, privacy, and security—*and* they will have been tested (and continually monitored) for compliance at the level of service that they provide. This *already-completed* level of testing significantly reduces the amount of testing and documentation necessary for a CMS system to achieve an ATO. Non-FedRAMP authorized CSPs must be ***FULLY*** vetted to ensure that they have been contractually bound to federal standards, and that they have been *fully* tested for compliance. ***CMS business owners seeking to utilize a CSP that is not FedRAMP or GSA authorized (like any other non-authorized GSS) will bear the full burden of required security testing for ALL of the applicable CSP infrastructure, as well as the maintenance cost of the ATO for that CSP.*** This includes utilization of CSPs that have been granted ATOs by *other* federal agencies, but cannot (or *will* not) provide sufficient proof of compliance, testing, and monitoring, to applicable CMS organizations. For this reason, CMS business owners and organizations may benefit from *pooling* their service needs and resources and consolidating on select CSPs (where FedRAMP-

authorized CSP services are not available), and equitably distributing the costs for testing and maintenance of the *CMS* ATOs for those CSPs.

4.3 GENERAL SECURITY REQUIREMENTS

Under FISMA, the security control requirements for cloud services are the same as those of non-cloud CMS General Support System (GSS) and Major Applications (MA). For CMS, those controls are documented in the *CMS Information Security (IS) Acceptable Risk Safeguards (ARS), CMS Minimum Security Requirements (CMSR)* ("the ARS")[11]. Security documentation necessary to achieve and maintain an ATO for systems utilizing cloud services are the same as for non-cloud systems. All security documentation must be uploaded into the CMS FISMA Controls Tracking System (CFACTS). CMS must ensure that security and privacy controls are implemented properly, operate as intended, and meet CMS security and privacy requirements.

The existence of a FedRAMP A&A does not replace the existing CMS Security Authorization process; rather it provides a set of inputs to allow for standardized and consistent evaluation of CSP offerings. The FedRAMP process simplifies the A&A process for external offerings and shortens the timeframe for granting a CMS Authority to Operate (ATO) for those offerings.

It should be noted that, from a security perspective, not all CSPs are the same. Even FedRAMP authorization (as an approved IaaS, PaaS, or SaaS provider) does not mean that the services offered by every FedRAMP authorized CSP are equal. Cloud computing, even when marketed as *"FedRAMP authorized"*, should *not* be viewed as a ***commodity***, but instead as discrete COTs/ GOTs *products*, with different approaches and different levels-of-service. There are significant security discriminators that must be evaluated when shopping for a CSP. Under the FedRAMP program, each CSP is required to provide a complete documented description of *each* NIST security control requirement (listed in NIST SP 800-53), their implemented controls, *and* their current compliance. FedRAMP establishes no *minimum* scope of coverage of the NIST controls—only that their level of coverage is *documented* to the customer. Therefore, some CSPs will provide *more* security control coverage than others will. FedRAMP authorization does *not guarantee compliance*; only that the *compliance status* information must *be provided* to the procuring business owner—with the understanding that the procuring agency is responsible for remediating any remaining gaps in control coverage. The CMS business owner is responsible (and accountable) for the full lifecycle security-compliance of their systems, and should plan for dealing with these remaining gaps.

Business owners should be aware that the overall cost associated with procurement of a CSP is not necessarily directly related to the cost of the contractual procurement of a CSP. In fact, *low-cost* CSPs may actually end up being *more* expensive than moderate or high-cost CSPs—due mainly to differences in their security offerings and total FISMA coverage. The comparative analysis may reveal differences in security control coverage as stark as a comparison of safety features between a *Volvo* and a *Yugo*. Sure, you can get a *great price* on a *Yugo*, but it most definitely will *not* operate at the same safety level as a *Volvo*. While business owners may not need the most robust of business functionalities, there are only two grades of FISMA

[11] The common FISMA government-wide security control requirements are maintained in the NIST SP 800-53. However, the 800-53 requires individual agency input to establish and document a significant amount of *"agency defined"* parameters. CMS accomplishes this final step in the ARS.

compliance: *compliant*, and *non-compliant*. If a CSP is non-compliant with any FISMA requirement (or does not even offer a security service solution for a requirement), it becomes the responsibility of the business owner to make-up for the shortcoming. Business owners should be aware of the potential cost-differential associated with upgrading a *Yugo* to be as safe as a *Volvo.*

4.3.1 CLOUD PROVISIONING SECURITY

The ability to quickly provision services in a cloud environment is perhaps the biggest advantage to the cloud services model. It affords the business owner with the capability to rapidly constitute complex services in a robust environment.

However, the advantage of quick provisioning can quickly be out-weighed by the downside. If this capability is compromised; that is; if some unauthorized entity gains access to this capability, they can quickly provision complex malware tools and hosting capabilities in a CMS provided cloud (*and bill CMS for the capability!*) Or (and perhaps *worse*), they can reset access controls in the provisioning environment to lock out CMS administrators, then quickly and thoroughly *de-provision* an entire CMS system, knocking the system off-line, and *requiring a complete reconstitution* of the system in the cloud (*after* CMS re-establishes control of, and access to, the provisioning environment!)

CMS business owners should aggressively protect the access to this provisioning capability and frequently perform functions to ensure and verify that this capability has not been compromised.

4.3.2 CLOUD INFRASTRUCTURE SECURITY

Another great advantage to cloud computing is the capability to define securely configured master images and push those images out in a rapid manner. Within the CSP, this is efficiently accomplished in a virtual environment. However, the virtual environments of CSPs present new and challenging issues that must be dealt with. The CSP reliance on hypervisors to isolate processes and create application *"sandboxes"* means that, if the virtualization environment is compromised, the scope of the breach can immediately escalate. A breach of the virtualization platform that results in an escape to the hypervisor represents, what some may consider, a *worst-case* security scenario. The virtualization platform (hypervisor/VMM) is software written by human beings, and like everything else humans write, it will likely contain vulnerabilities. Some of these vulnerabilities will result in a breakdown in isolation that the virtualization platform is supposed to enforce. Evil-doers will target this layer with attacks because the benefits of a compromise of this layer are simply too great for them to ignore[12]. Note that these threats may come from a direct attack *on a CMS system*, from an attack on *another customer* within the same CSP cloud, or even ***by*** another customer within the CSP service (which is one reason why CMS does not allow the use of *public* clouds.) While there have been a few *disclosed* attacks on virtualization infrastructure, it is only a matter of time before a *widespread* enterprise breach directly attributed to a hypervisor vulnerability exploitation. Business owners need to ensure that the CSP has extended their vulnerability and configuration management processes to this

[12] There are several known exploits targeting the hypervisor and other virtualization technologies documented at http://www.us-cert.gov.

sensitive layer, just as they should for any sensitive OS. Business owners (and their CSPs) should treat the virtualization platform as *the most sensitive component* of their data center.

NIST SP 800-125, *Guide to Security for Full Virtualization Technologies*, discusses the security concerns associated with full virtualization technologies for server and desktop virtualization, and provides recommendations for addressing these concerns, *in addition to* existing recommended security practices, which remain applicable in most virtual environments.

4.3.3 CLOUD NETWORK AND PERIMETER SECURITY

While enterprise data centers typically have robust perimeter security such as firewalls, network Intrusion Detection Systems (IDS), and Intrusion Prevention Systems (IPS), malware occasionally slips through to compromise the endpoints. Once inside, there is the potential for a spear-phishing attacker to bounce from the compromised CSP administrator endpoint and use credentials from that endpoint to move into the server or cloud infrastructure of another CSP customer to steal data or create mischief.

Most CSPs have robust perimeter security measures to protect their customer's server instances. The CSPs typically have firewalls protecting their customers, but bypassing the CSP perimeter security may be as simple as having a credit card. Potential threats (bad guys) could access the infrastructure simply by renting some CSP services with a stolen credit card and provision their cloud server on the same physical infrastructure as a CMS virtual server. For CMS business owners protecting their CMS data, countering this means the classic *"defense-in-depth"* approach, where business owners need to consider protecting the individual host that might be living in a dynamic, virtualized environment. Business owners need to understand that they are responsible for the security of their CSP's servers and should consider augmenting existing CSP security with host-based security that CMS manages, including firewalls, vulnerability shielding (IDS/IPS), system file integrity, and log inspection. Business owners concerned about whether they might be at risk from infections from other parts of the CSP customer-base should consider requiring additional CSP gateway assessment tools that can determine whether they have been compromised.

4.3.4 CLOUD DATA STORAGE SECURITY

Cloud data-storage offers several advantages to the traditional data-storage model. These advantages may include; automated replication of the data in several regional CSP datacenters, encryption at rest and in transit, and automated data retention. However, these advantages can also be security and privacy headaches. The use of CSPs makes it more difficult for business owners to *really* know where the data is stored, or whether that data is being co-mingled with non-CMS data. Isolation management, that is, ensuring that the data is isolated from everyone except those to whom CMS has given explicit rights, is difficult, especially in a multi-tenant situation (where several CSP customers are sharing IT infrastructure.) The failure of a single storage controller can quickly become a single point of failure, leading to a breach of not only CMS data, but of every other customer's data within the CSP infrastructure.

Encryption, authentication, and authorization are important components of any enterprise security infrastructure—the cloud being no exception. However, if the CSP is encrypting

everyone's data with the same infrastructure and the *same keys*, then if one CSP customer is breached, then everyone is breached. One solution is for business owners to insist on the use of custom encryption keys for each individual CSP client, or even for each CMS instance within the CSP. However, this adds to the overhead management required by the business owners to properly manage and protect the *"master keys."*

4.3.5 FEDERAL RECORDS MANAGEMENT

The National Archive and Records Administration (NARA) recognizes that cloud service and deployment models affect how records[13] may be created, used, and stored in cloud computing environments. As a result, they have issued NARA Bulletin 2010-05[14], *Guidance on Managing Records in Cloud Computing Environments*. NARA has identified several records management challenges with cloud computing environments.

Cloud applications may lack the capability to implement records disposition schedules, including the ability to transfer and permanently delete records or perform other records management functions. Therefore, specific service and deployment models may not meet all of the records management requirements of 36 CFR Part 1236[15] (formerly 36 CFR Part 1234). Examples of these requirements include:

- Maintaining records in a way that maintains their functionality and integrity throughout the record's full lifecycle
- Maintaining links between the records and their metadata
- Transfer of archival records to NARA or deletion of temporary records according to NARA-approved retention schedules.

Depending on the application, CSPs must be made aware of the record retention requirements governing a given body of federal records stored in one or more cloud locations. CMS business owners need to be able to control any proposed deletion of records pursuant to existing authorities, wherever the records may be located in the CSP's cloud. CSPs must also act to ensure that records are accessible to ensure agency responsiveness to discovery, or FOIA/Privacy Act, or other access requests.

CMS business owners are responsible for managing their records in accordance with NARA statues including the Federal Records Act (44 U.S.C. Chapters 21, 29, 31, 33)[16] and NARA regulations (36 CFR Chapter XII Subchapter B). However, NARA recognizes that the differences between cloud deployment models may affect how and by whom (agency/contractor) records management activities can be performed.

---

[13] *Records* or *Federal records* is defined in 44 U.S.C. 3301 as including *"all books, papers, maps, photographs, machine readable materials, or other documentary materials, regardless of physical form or characteristics, made or received by an agency of the United States Government under Federal law or in connection with the transaction of public business and preserved or appropriate for preservation by that agency or its legitimate successor as evidence of the organization, functions, policies, decisions, procedures, operations or other activities of the Government or because of the informational value of the data in them (44 U.S.C. 3301)."* (See CFR 1222.10 for an explanation of this definition at http://www.archives.gov/about/regulations/part-1222.html#1222.10).

[14] NARA Bulletin 2010-05 is available at http://www.archives.gov/records-mgmt/bulletins/2010/2010-05.html.

[15] 36 CFR Part 1236 is available at http://www.archives.gov/about/regulations/part-1236.html.

[16] Applicable Federal records management laws can be found at http://www.archives.gov/about/laws/.

CMS management should follow these NARA guidelines for creating standards and policies for managing CMS' records created, used, or stored in cloud computing environments:

1. Include the CMS records management officer and/or staff in the planning, development, deployment, and use of cloud computing solutions.
2. Define which copy of records will be declared as the agency's record copy and manage these in accordance with 36 CFR Part 1222. Remember, the value of records in the cloud may be greater than the value of any other set because of indexing or other reasons. In such instances, this benefit may require designation of the copies as records.
3. Include instructions for determining if federal records in a cloud environment are covered under an existing records retention schedule.
4. Include instructions on how all records will be captured, managed, retained, made available to authorized users, and retention periods applied. (Note that other statutes like HIPAA may extend some record retention requirements beyond those of NARA.)
5. Include instructions on conducting a records analysis, developing, and submitting records retention schedules to NARA for unscheduled records in a cloud environment. These instructions should include scheduling system documentation, metadata, and related records.
6. Include instructions to periodically test transfers of federal records to other environments, including agency servers, to ensure the records remain portable.
7. Include instructions on how data will be migrated to new formats, operating systems, etc., so that records are readable throughout their entire life cycles. Include in your migration planning provisions for transferring permanent records in the cloud to NARA. An agency choosing to pre-accession[17] its permanent electronic records to NARA is no longer responsible for migration except to meet its business purposes.
8. Resolve portability and accessibility issues through good records management policies and other data governance practices. Data governance typically addresses interoperability of computing systems, portability of data (able to move from one system to another), and information security and access. However, such policies by themselves will not address CMS' compliance with the Federal Records Act and NARA regulations.

Ultimately, CMS maintains responsibility for managing its records whether they reside in a contracted environment or under agency physical custody (see 36 CFR Part 1222.32 (b))[18]. When dealing with a contractor, an agency must include a records management clause in any contract or similar agreement. At a minimum, a records management clause ensures that the federal agency and the contractor are aware of their statutory records management responsibilities.

[17] *Pre-accessioning* is when NARA fully processes (for preservation purposes) permanently valuable electronic records in order to assume physical custody before the records are scheduled to legally become part of the National Archives of the United States. Pre-accession is covered in NATA 2009-03 available at http://www.archives.gov/records-mgmt/bulletins/2009/2009-03.html.

[18] 36 CFR Part 1222.32 (b) is available at http://www.archives.gov/about/regulations/part-1222.html#1222.32.

The following is a NARA-provided general clause that CMS business owners can modify to fit the planned type of service and specific agency records management needs:

> *Use of contractor's site and services may require management of federal records. If the contractor holds federal records, the contractor must manage federal records in accordance with all applicable records management laws and regulations, including but not limited to the Federal Records Act (44 U.S.C. chs. 21, 29, 31, 33), and regulations of the National Archives and Records Administration (NARA) at 36 CFR Chapter XII Subchapter B). Managing the records includes, but is not limited to; secure storage, retrievability, and proper disposition of all federal records including transfer of permanently valuable records to NARA in a format and manner acceptable to NARA at the time of transfer. The agency also remains responsible under the laws and regulations cited above for ensuring that applicable records management laws and regulations are complied with through the life and termination of the contract.*

If a CMS business owner decides to create or join a private or community cloud, they must still meet records management responsibilities. If a CSP ceases to provide services, CMS must continue to meet its records management obligations. Business owners should plan for this contingency.

NARA provides a *Toolkit for Managing Electronic Records* that is available at http://www.archives.gov/records-mgmt/toolkit/pdf/all-nara-tools-by-date.pdf.

4.3.6 IDENTITY MANAGEMENT

Identity, credential, and access management (*"identity management"*) includes both the creation and management of user credentials used for authentication, and attributes used for authorization—to include the enforcement of authorization policy. Identity management also addresses required auditing and reporting for identity-related events. Ultimately, identity management governs access control to information and functionality within the cloud.

OMB Memorandum M-04-04, *E-Authentication Guidance for Federal Agencies* discusses the risk inherent to granting access to an electronic service (web-based, cloud, etc.) It identifies four levels of identity assurance from *"little or no confidence in the asserted identity's validity (anonymous)"* to *"very high confidence in the asserted identity's validity."* To successfully implement a service electronically, CMS must determine the required level of identity assurance for each transaction. This is accomplished through a risk assessment for each transaction. Details of this assessment is covered in the ARS and the CMS *Risk Management Handbook, Volume III, Standard 3.1, Authentication.*

Selection of a cloud environment is driven by the type of user population accessing the cloud resources. For instance, a CSP hosting applications and data accessible to the general public requires very different identity management architecture than one hosting sensitive or proprietary information. A cloud servicing a specific community, such as *healthcare providers*, or CMS *employees* (or *contractors*), requires a more rigorous management of identities.

Under cloud computing, the CMS identity management needs to discuss the following four key challenges:

Identity credential and account provisioning/de-provisioning – The creation and lifecycle management of user accounts and user attributes (where necessary) used for authorization/ access control. This may include issuing credentials, acceptance of externally issued identity credentials, or a mixture of both.

Authentication - The verification of identity credentials (e.g., username/password, digital certificate, security token, biometric, etc.) during a transaction. To meet these needs, cloud applications and services will enter into *trust relationships*[19] with identity and attribute providers for authentication and attribute services. This may include an internal CMS-maintained identity system (such as the Baltimore Enterprise Data Center), or some external identity management provider (such as; the Social Security Administration, Veterans Health Administration; etc.)

Authorization - The management of authorization policy, assessment of a specific access request against that policy, and the enforcement of policy decisions.

Access Privileges - The granting of user privileges affects what a user is authorized to do. The provisioning process includes the evaluation of user attributes and approval of a user's authorized entitlements based on those attributes.

Public cloud services will generally be offered at e-authentication Level 1 or 2. At these levels of assurance, the amount of personal information required to establish and maintain the relationship ranges from very little to none at all. At Level 1, individuals may remain anonymous and can provide non-affiliated identifiers. Account provisioning is based on associating the unique login with user activity ensuring a consistent user experience. Some examples of acceptable Level 1 identity assurance activities include allowing RSS feeds, comments on blogs, and other activities where identity of the correspondent is not necessary. No sensitive information should be exchanged at this level of assurance. At Level 2, cloud services require a rudimentary level of identity verification and credentials meeting SP 800-63 entropy standards. *Security Assertion Markup Language (SAML)*[20] assertions based on UserID/password login are acceptable at this level of assurance and provide a reasonable level of confidence in the claimed identity. Identity e-authentication Level 2 is appropriate for a wide range of CMS business with the public where CMS requires an initial identity assertion (the details of which are verified independently prior to any CMS action). In all cases, the exchange of *Personally Identifiable Information (PII)* must be kept to a minimum, and adequate notice provided of the information exchanged, and its uses. Cloud services requiring *high* or *very high* confidence in

[19] When there are *trust relationships* between cloud services, the authentication mechanism for each service trusts the authentication mechanism for all other trusted services. If a user or application is authenticated by one service, its authentication is accepted by all other domains that trust the authenticating service. Users in a trusted service have access to resources in the trusting service, subject to the access controls that are applied in the trusting domain. *Access to resources* in any discussion of trust relationships always assumes the limitations of access control. Trust relationships allow users and computers to be authenticated (to have their identity verified) by an authentication authority. Access control allows authenticated users to use the resources (files, folders, and virtual containers) that they are authorized to use and prohibits them from using (or even seeing) resources that they are not authorized to use.

[20] Information on *SAML* can be found at http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=security.

the asserted identity (e-authentication Levels 3 or 4), must utilize identity and credentialing solutions commensurate with these levels of assurance, this will usually require a cryptographic credential solution.

Private (government) cloud services, including *community* clouds, have the same requirements as those identified for the *public*. However, where the community is comprised of members of the Federal work force and their contractor support, CMS should utilize the credentials associated with the *Personal Identity Verification (PIV)* cards. These credentials meet M-04-04 e-authentication Level 4, and can be utilized for all government-dedicated cloud activities. Use of the PIV credentials requires the cloud service to process public key infrastructure certificates, to include path discovery and validation.

Clouds serving organizations or individuals *external* to CMS (such as *healthcare providers* and *beneficiaries*) require e-authentication Level 2 or 3, at a minimum (see the CMS *Risk Management Handbook, Volume III, Standard 3.1, Authentication* for specific guidance.) Ideally, the provisioning process will include changes to access entitlements throughout the user lifecycle, up to and including the removal of all access entitlements when the individual no longer requires access.

Authentication of identity credentials may require the CSP to employ mechanisms communicating with the identity provider (which may *not* be the applicable CSP) to verify the validity of the presented identity credential. In other cases, the credential will be presented by the identity provider on behalf of the individual, which negates the requirement for a separate authentication step. The two scenarios to be considered here concern whether the identity credentials are *internally* or *externally* provided. That is, whether the CSP supports *federated* credentials.

In the case of externally provided, or *federated*, credentials, the cloud service recognizes identity credentials from multiple sources certified as meeting the pre-determined level of assurance. Private cloud services may employ trust lists and other related mechanisms to determine access. However, there is a fundamental need to establish the validity of the offered credential, ***before*** access is granted. To do so, three questions must be answered:

- Was the identity credential issued by the claimed identity provider? (*Is the credential **authentic**?*)
- Is the identity credential still in good standing—e.g., not revoked or expired? (*Is the credential **valid**?*)
- Does the identity credential match the claimed identity? (*Does the claimed identity **match** the credential?*)

If the answer to all of these questions is *"Yes"*, then the credential can be considered *verified*.

Assertion-based identity credentials (OpenID, SAML, WSFed, etc.) will be presented to the cloud service by the identity provider. A handshake between the identity provider and the CSP will establish the validity of the identity provider, which then presents the identity assertion. In the case of identity assurance level 1 credentials, this assertion is a unique, anonymous, random identifier used by the CSP to establish a user account and provide a unique experience for the user. At higher levels of identity assurance (Levels 2, 3 or 4), assertions may contain PII and

may be sensitive in nature. In these cases, the identity provider is presenting the credential on behalf of the individual, which answers the three questions above.

Cryptographic-based identity credentials (One-time-passwords, digital certificates) will be presented to the cloud service by the individual without the intervention of the identity provider and must therefore be verified. To do so, the CSP must send a query to the identity provider requesting status. In the case of PKI, this is known as *certificate status checking* and will answer question two above. The answer to question one is determined by establishing the validity of the signature on the certificate itself. The Federal PIV and PIV-Interoperable identity credentials are cryptographic-based; therefore, CSPs whose constituency includes the Federal work force should be able to accept federated cryptographic-based identity credentials.

Locally issued identity credentials, or those ***issued by CMS business owners***, may simplify the credential authentication and access control process, however, they also ***lead to the unmanageable proliferation of user credentials for each individual***. Federation can negate this proliferation. As the Federal Government moves towards federated enterprise architecture, a key requirement for federation is inter-organizational trust. Federation is the recognition and acceptance of identity credentials issued by the Federal government. There are two federation mechanisms in use complementing each other: The *Trust Framework Provider* program and *The Federal PKI*.

The *Trust Framework Provider* program establishes the criteria for identity credential federation at assurance levels 1, 2, and non-PKI level 3 (these would be required for external users such as Medicare providers.) It employs a governance model permitting *industry self-regulation* of identity providers. In this program, the Federal government has established the criteria[21] for industry accreditation bodies (trust frameworks), which in turn certify individual identity providers. The criteria established by the Federal government includes the establishment of federal profiles for the different open source identity technology standards and ensuring the industry accreditation bodies are assessing identity providers against the Federal profiles, applying the principles found in NIST Special Publication 800-63, *Electronic Authentication* Guideline, and applying the appropriate privacy principles. Trust is established from the Federal government through the Trust Framework Providers to the individual Identity Providers and is perpetuated through regular review and audit activities.

The *Federal PKI* deals with the federation of trust for organizations utilizing public key technology and is primarily concerned with assurance levels 3 and 4. The *Federal PKI Policy Authority*[22] governs the peer-to-peer relationships with the *Federal Bridge Certification Authority*, and the hierarchical relationships with the *Federal Common Policy Framework*. Cross certification with either entity ensures membership in the trust federation of the Federal PKI.

A valid, authenticated credential is not guaranteed access to a specific cloud resource. This 'authorization' decision is made by the application owner in the setting of the access control policy. In the public cloud scenario, authorization may be an extension of authentication, with a

[21] As of the date of this publication, only one *framework* has been established—for Assurance Level 1. The *Federal ICAM Trust Framework* can be found at http://openidentityexchange.org/trust-frameworks/us-icam.
[22] More information on the *Federal PKI Policy Authority*, the *Federal Bridge Certification Authority*, and the *Federal Common Policy Framework* can be found at http://www.idmanagement.gov/fpkipa/.

policy decision that *any* valid credential can have access. This is especially true for Level 1 applications. However, at higher levels of assurance, there may be filtering or limiting factors affecting authorization decisions. Authorization may be limited to members of a privileged group, only those individuals whose credentials are issued by a specific identity provider, or only to specific individuals within a privileged group using trust lists. It is up to the *CMS business owner* to make this determination and ensure there is an access control policy to enforce it.

Authorization activities for private (government) cloud activities is more straightforward and will almost always be limited to specific individuals through the use of individual trust lists that tie specific credentials to specific authorized users. In these cases, all authorized users will be pre-provisioned into the system, and *new* users will (most likely) need to complete an out-of-band *on-boarding* or *enrollment* activity.

Closely related to authorization is the concept of access privileges. Once an individual identity has been *authenticated* and *authorized*, further attention may need to be paid to access privileges. This may be based on a set of static permissions or authorizations associated with information contained in the white list or user profile. However, it may also be a dynamic process that queries the identity provider or another authoritative source for *"attributes"* associated with the identity that may then be used to determine access privileges within the cloud.

While not generally a concern for public cloud services, private (government) cloud business owners need to set up trust zones and policy enforcement points or boundaries around information requiring higher levels of assurance. Enforcement points would be useful in defining areas for applying group policies, creating groups, centralizing administration, and having workspaces available to all users within the cloud. The group policies would be applied on a per-user profile. Access would then be limited to personnel with the *need-to-know* and appropriately cleared levels of access. The *Backend Attribute Exchange*[23] technical specification has been defined to assist organizations with the reach-back for attributes and provides a standardized process for requesting and receiving attribute information—allowing approval of attributes relevant to a cloud set up within a trust zone and policy enforcement point.

The CSP must have a system able to enforce or allow CMS-appointed personnel to enforce account management capabilities, such as account lockouts for unsuccessful logon attempts, defined inactivity times, remote access allowances, specific success/failure events, and management of elevated privilege accounts.

All identity credentialing, authentication, authorization, and access control events must be logged and those logs are subject to periodic audit. At a minimum, the CSP must produce logs of all specified success and failure events associated with identity and access management in the cloud environment it manages. These logs must then be archived for a pre-determined amount of time as specified by the ARS. These archived logs must be searchable and or discoverable with CMS-owned or CSP-provided interfaces (command line or graphical user interfaces).

When identifying the security considerations for identity and access management, consider the *most stringent* requirements applicable per OMB M-04-04 (e-authentication) and NIST

[23] The *Backend Attribute Exchange*, published by the *HSPD-12 Working Group* of the *Federal PKI Policy Authority* is available at http://www.idmanagement.gov/awg/documents/BackendArchitectureInterfaceSpec_v100.pdf.

SP800-63; and apply risk-based security controls described in NIST SP 800-53 (as amended) based on the FIPS 199 security categorization as augmented by CMS security requirements.

4.3.7 CONTINUOUS MONITORING

Continuous monitoring is a NIST-required technique to address the security impacts on an information system resulting from changes to the hardware, software, firmware, or operational environment. A well-designed and well-managed continuous monitoring program can effectively transform an otherwise static security control assessment and risk determination process into a dynamic process that provides essential, near real-time security status-related information to organizational officials in order to take appropriate risk mitigation actions and make cost-effective, risk-based decisions regarding the operation of the information system. Continuous monitoring programs provide organizations with an effective mechanism to update Security Plans, Security Assessment Reports, and Plans of Action and Milestones (POA&Ms).

An effective continuous monitoring program includes:

- *Configuration management* and control processes for information systems
- *Security impact analyses* on proposed or actual changes to information systems and environments of operation
- *Assessment of selected security controls* (including system-specific, hybrid, and common controls) based on the defined continuous monitoring strategy
- *Security status reporting* to appropriate officials
- *Active involvement by authorizing officials* in the ongoing management of information system-related security risks

Maintenance of the security Authority To Operate (ATO) will be through continuous monitoring of security controls of the CSP's system and its environment to determine if the security controls at the CSP continue to be effective over time in light of changes that occur in the system and the environment. Through continuous monitoring, security controls and supporting deliverables are updated and submitted to CMS, HHS, and FedRAMP (as applicable) per the schedules in Table 1[24] below. The submitted deliverables provide a current understanding of the security state and risk posture of the information systems. They allow authorizing officials to make credible risk-based decisions regarding the continued operations of the information systems and initiate appropriate responses as needed when changes occur. The deliverable frequencies below are to be considered standards. However, there will be instances, beyond the control of CMS in which deliverables may be required on an ad hoc basis (such as Federal CSIRT data calls.)

Table 1 provides a listing of the deliverables, responsible party, and frequency for completion. The table is organized into:

- **Deliverable** – Detailed description of the reporting artifact. If the artifact is expected in a specific format, that format appears in **bold** text.
- **Frequency** – Frequency under which the artifact should be created and updated.

[24] This table is derived from the *Proposed Security Assessment & Authorization for U.S. Government Cloud Computing* (Draft 0.96, dated 11/2/2010), available at https://info.apps.gov/sites/default/files/Proposed-Security-Assessment-and-Authorization-for-Cloud-Computing.pdf.

- **Responsibility** – Whether FedRAMP or the CSP is responsible for creation and maintenance of the artifact.

Table 1 Continuous Monitoring Deliverables

Deliverable	Frequency	Responsibility	
		FedRAMP (or Sponsor Organization)	Cloud Service Provider
Scan reports of all systems within the boundary for vulnerability (Patch) management. ***(Tool Output Report)***	At *least* Monthly		**X**
Scan for verification of FDCC compliance (USGCB, CIS). ***(SCAP Tool Output)***	At *least* Quarterly		**X**
Incident Response Plan.	Annually		**X**
POAM Remediation ***(Completed POA&M Matrix)***	At *least* Quarterly[25], *preferably* monthly		**X**
Change Control Process	Annually		**X**
Penetration testing ***(Formal plan and results)***	At *least* Annually	**X**	**X**
Independent Verification and Validation (IV&V) of controls	Semi-Annually	**X**	
Scan to verify that boundary has not changed (also that no rogue systems are added after ATO) ***(Tool Output Report)***	At *least* Quarterly		**X**
System configuration management software ***(SCAP Tool Output)***	At *least* Quarterly		**X**
FISMA Reporting data	At *least* Quarterly		**X**
Update Documentation	Annually		**X**

[25] This frequency is *less restrictive* than CMS ARS CA-5 requirements. However, CMS has determined that this is an *acceptable risk* due to the cost impacts of upgrading an approved FedRAMP CSP, which provides only highly-standardized offerings, to *these* additional CMS requirements. If achievable without significant cost impact, business owners should endeavor to meet the ARS CA-5 (monthly) requirements.

Deliverable	Frequency	Responsibility	
Contingency Plan and Test Report	Annually		X
Separation of Duties Matrix	Annually		X
Information Security Awareness and Training Records Results	Annually		X

The Change Control Process is instrumental in ensuring the integrity of the cloud-computing environment. As the system owners as well as other authorizing officials approve changes, they are systematically documented. This documentation is a critical aspect of continuous monitoring since it establishes all of the requirements that led to the need for the change as well as the specific details of the implementation. To ensure that changes to the enterprise do not alter the security posture beyond the parameters set by the FedRAMP Joint Authorization Board (JAB) (or the CMS CISO for non-FedRAMP CSPs), the key documents in the authorization package, which include the *Security Plan*, *Security Assessment Report*, and *Plan Of Actions & Milestones* are updated and formally submitted to FedRAMP (or the sponsor organization) within 30 days of approved modifications.

There are however, changes that are considered routine. These changes can be standard maintenance, addition or deletion of users, the application of standard security patches, or other routine activities. While these changes individually may not have much effect on the overall security posture of the system, in aggregate they can create a formidable security issue. To combat this possibility, these routine changes should be documented as part of the CSP's standard change management process and accounted for via the CSP's internal continuous monitoring plan. Accordingly, these changes must be documented, at a minimum, within the current SSP of the system within 30 days of implementation.

Throughout the *System Development Life Cycle (SDLC),* business owners must be cognizant of changes to the cloud system. Since systems routinely experience changes over time to accommodate new requirements, new technologies, or new risks, they must be routinely analyzed in respect to the security posture. Minor changes typically have little impact to the security posture of a system. These changes can be standard maintenance, adding or deleting users, applying standard security patches, or other routine activities. However, significant changes require an added level of attention and action. NIST defines significant change as *"…a change that is likely to affect the security state of an information system."* Changes, such as installing a new operating system, port modification, new hardware platforms, or changes to the security controls should automatically trigger review for a re-authorization of the system via the FedRAMP (or sponsoring organization) ATO process.

Minor changes must be captured and documented in the SSP of the system within 30 days of implementation. This requirement should be part of the CSP's documented internal continuous monitoring plan. Once the SSP is updated, it must be submitted to FedRAMP or the sponsoring organization, and a record of the change must be maintained internally.

Major or significant changes may require re-authorization via the FedRAMP (or sponsoring organization) process. In order to facilitate a re-authorization, it is the responsibility of both the CSP and the business owner to notify FedRAMP (or sponsoring organization) of the need to make such a significant change. FedRAMP (or the sponsoring organization) will assist and coordinate with all stakeholders the necessary steps to ensure that the change is adequately documented, tested, and approved.

Vulnerability patching is critical. Proprietary operating system vendors (POSV) are constantly providing patches to mitigate vulnerabilities that are discovered. In fact, regularly scheduled monthly patches are published by many POSV to be applied to the appropriate operating system. It is also the case that POSV will, from time-to-time, publish security patches that should be applied on systems as soon as possible due to the serious nature of the vulnerability. Systems running in virtual environments are not exempted from patching. In fact, not only are the operating systems running in a virtual environment to be patched routinely, but oftentimes the virtualization software itself is exposed to vulnerabilities and thus must be patched via either a vendor-based solution or other technical solution.

Open source operating systems require patch and vulnerability management as well. Due to the open nature of these operating systems there needs to be a reliable distribution point for system administrators to safely and securely obtain the required patches. These patches are available at the specific vendor's website.

Database platforms, web platforms and applications, and virtually all other software applications come with their own security issues. It is not only prudent, but also necessary to stay abreast of all of the vulnerabilities that are represented by the IT infrastructure and applications that are in use.

While vulnerability management is indeed a difficult and daunting task, there are proven tools available to assist the system owner and administrator in discovering the vulnerabilities in a timely fashion. These tools must be updated prior to being run. Updates are available at the corresponding vendor's website.

With these issues in mind CMS (and FedRAMP) requires CSPs to provide the following:

- ***Monthly vulnerability scans of all servers.*** Tools used to perform the scan must be provided as well as the version number reflecting the latest update. A formal report must include all vulnerabilities discovered, mitigated or the mitigating strategy. This report should list the vulnerabilities by severity and name[26]. Specificity is crucial to addressing the security posture of the system. All *"High"* level vulnerabilities must be mitigated within thirty days (30) days of discovery. *"Moderate"* level vulnerabilities must be mitigated within ninety (90) days of discovery[27]. It is accepted that, at certain times, the application of certain security patches can cause negative effects on systems. In these situations, it is understood that compensating controls must be used to minimize system performance degradation while

[26] Vendors should utilize the *Common Vulnerabilities and Exposures (CVE)* terminology available through the *National Vulnerability Database* at http://nvd.nist.gov/.

[27] This frequency is *less restrictive* than CMS ARS SI-2 requirements. However, CMS has determined that this is an *acceptable risk* due to the cost impacts of upgrading an approved FedRAMP CSP, which provides only highly standardized offerings, to *these* additional CMS requirements. If achievable without significant cost impact, business owners should endeavor to meet the ARS SI-2 requirements.

serving to mitigate the vulnerability. These *"Workarounds"* must be submitted to FedRAMP and/or the Sponsoring agency for acceptance. All reporting must reflect these activities.

- ***Quarterly FDCC and/or system configuration compliance scans***, with a *Security Content Automation Protocol*[28] *(SCAP)* validated tool, across the entire boundary, which verifies that all servers maintain compliance with the mandated FDCC and/or approved system configuration security settings.
- ***Weekly scans for malicious code***. Internal scans must be performed with the appropriate updated toolset. Monthly reporting is required to be submitted to FedRAMP (or sponsoring organization), where activity is summarized.
- ***All software operating systems and applications are required to be scanned by an appropriate tool to perform a thorough code review to discover malicious code.*** Mandatory reporting to FedRAMP (or sponsoring organization) must include tool used, tool configuration settings, scanning parameters, application scanned (name and version) and the name of the third party performing the scan. Initial report should be included with the SSP as part of the initial authorization package.
- ***Performance of the annual security assessment in accordance with NIST guidelines.*** CSP must perform a security assessment annually or whenever a significant change occurs. This is necessary if there is to be a continuous awareness of the risk and security posture of the system.
- ***Quarterly POA&M remediation reporting***. CSP must provide to FedRAMP a detailed matrix of POA&M activities using the supplied FedRAMP (or sponsoring organization) POA&M Template. This should include milestones met or milestones missed, resources required and validation parameters.
- ***Active Incident Response capabilities*** allow suspect systems to be isolated and inspected for any unapproved or otherwise malicious applications.
- ***Quarterly boundary-wide scans*** are required to be performed on the defined boundary IT system inventory to validate the proper HW and SW configurations as well as search and discover rogue systems attached to the infrastructure. A summary report, inclusive of a detailed network architecture drawing must be provided to FedRAMP.
- ***Change Control Process*** meetings to determine and validate the necessity for suggested changes to HW/SW within the enterprise must be coordinated with FedRAMP to ensure that the JAB is aware of the changes being made to the system.

As part of the authorization process, the CSP system security plan will have documented all of the *"IR"* or Incident Response family of controls. One of these controls (IR-8) requires the development of an Incident Response plan that will cover incident response as documented in the NIST SP 800-61 guidelines. The plan should outline the resources and management support that is needed to effectively maintain and mature an incident response capability. The incident response plan should include these elements:

- Mission
- Strategies and goals
- Senior management approval
- Organizational approach to incident response
- How the incident response team will communicate with the rest of the organization

[28] Information related to the *Security Content Automation Protocol (SCAP)* is available at http://scap.nist.gov.

- Metrics for measuring the incident response capability
- Roadmap for maturing the incident response capability
- How the program fits into the overall organization.

The *incident response program* structure should be discussed within the plan. The response plan must address the possibility that incidents, including privacy breaches and classified spills[29], may impact the cloud and shared cloud customers. In any shared system, communication is the biggest key to success.

Independent Verification and Validation (IV&V) is an integral component to a successful implementation of cloud computing (whether FedRAMP or CMS authorized.) With this in mind, it must be noted that establishing and maintaining an internal CSP expertise of FedRAMP policies, procedures, and processes is going to be required. This expertise will be tasked to perform various IV&V functions with CSP's, sponsoring agencies and commercial entities obtained by CSP's with absolute independence on behalf of FedRAMP. FedRAMP IV&V will be on behalf of the JAB.

As part of the IV&V efforts, FedRAMP (or the sponsoring organization) will periodically perform audits (both scheduled and unscheduled) related strictly to the CSP offering and the established system boundary. This will include, but not be limited to:

- Scheduled annual assessments of the system security documentation
- Verification of testing procedures
- Validation of testing tools and assessments
- Validation of assessment methodologies employed by the CSP and independent assessors
- Verification of the CSP continuous monitoring program
- Validation of CSP risk level determination criteria

Several methods must be employed to accomplish these tasks. In accordance with the new FIMSA requirement and as a matter of implementing industry best practices, FedRAMP (or the sponsoring organization) IV&V may be performed using penetration testing. This testing may be performed with strict adherence to the specific guidelines established by a mutually agreed upon *Rules of Engagement* agreement between IV&V entity and the target stakeholders. Unless otherwise stated in the agreement, all penetration testing will be passive in nature to avoid unintentional consequences. No attempts to exploit vulnerabilities will be allowed unless specified within the *Rules of Engagement* agreement.

4.4 SECURITY IN THE SERVICE MODEL

4.4.1 INFRASTRUCTURE AS A SERVICE (IAAS) SECURITY

Infrastructure as a Service (IaaS) is a cloud service model in which a CMS business owner outsources the ***equipment*** used to support operations, including *storage*, *hardware*, *servers*

[29] *Classified Spills* (also known as contaminations or classified message incidents) occur when classified data is introduced to an unclassified computer system or to a system accredited at a lower classification than the data. For CMS, that might include accidental placement of CMS sensitive data (such as PHI or PII) on a CSP that is not authorized for such data.

and/or *networking* components. The CSP owns the equipment and is responsible for housing, running and maintaining it. IaaS is sometimes also referred to as *Utility* computing, *Hardware as a Service* [HaaS], *pay-per-use*, or *metered services*, and is a service-provisioning model in which a CSP makes computing resources and infrastructure management available to the customer *as needed*, and charges customers (the business owners) for specific usage rather than a flat rate. Like other types of on-demand computing (such as grid computing), this *utility* model seeks to maximize the efficient use of resources and/or minimize associated costs. To make an analogy to other services, IaaS is the IT hardware equivalent of models such as electrical power, which seeks to meet fluctuating customer needs and charge for the resources based on usage rather than on a flat-rate basis.

Since IaaS CSPs are *only providing basic infrastructure*, they will therefore only provide security controls associated the security of those basic components of the total CMS system. It is the responsibility of the CMS business owner to ensure that the remaining required security controls (those associated with middleware platforms, applications, and lifecycle management) be designed, implemented, and maintained to achieve compliance with the entire inventory of security controls requirements.

What this means to a business owner considering the procurement of the services of an IaaS CSP is that a *majority* of security control requirements will likely *not* be addressed by the CSP.

Figure 3 IaaS Security Controls Responsibilities

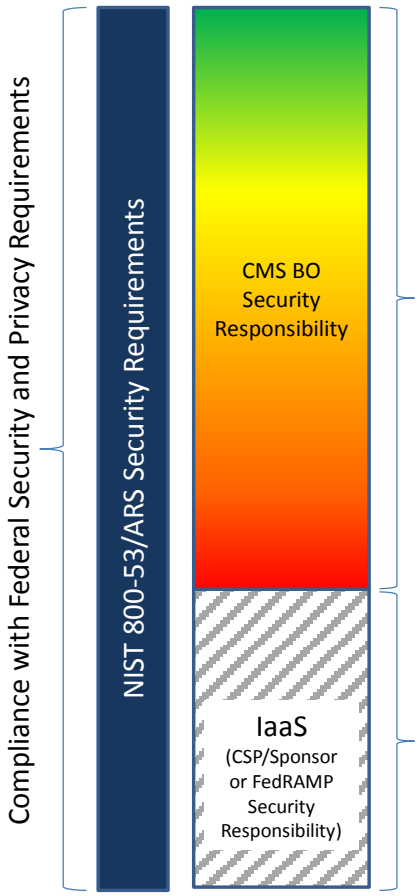

CMS Business Owner (BO) is responsible for implementing and managing security controls that address or include:

- Platforms
  - Middleware
  - Databases
  - Messaging,
  - Some Networking (firewall and *tier* management)
- Applications
  - Content
  - User access controls (authentication)
  - Presentation
  - Application code and code management
  - Application management
- Business Process Management
  - User role management
  - System design and lifecycle controls
  - Risk management
  - System services & acquisition controls

IaaS Cloud Services Provider (CSP) is responsible for implementing and managing security controls that address or include:

- IT Infrastructure, which may include:
  - Facilities
  - Hardware
  - Servers
  - Some Networking

Under the FedRAMP program, these controls will be overseen and managed by GSA (or some other sponsor organization.)

4.4.2 PLATFORM AS A SERVICE (PAAS) SECURITY

Platform as a Service (PaaS) is a cloud service model in which a CMS business owner outsources the ***network, servers, middleware, messaging services, databases, operating systems, and/or storage.*** All of which is usually provided in a virtualized environment, and is used to support operation of hosted applications.

This service delivery model allows the CMS business owner to procure (rent) virtualized servers and associated management services for running existing applications, or for developing and testing new ones. It facilitates deployment of applications without the cost and complexity of buying and managing the underlying hardware and software layers. So, unlike a simple IaaS environment that contains mere hardware, PaaS services include the underlying support software and configurations necessary to host a target application.

In broad terms, this definition means that a CMS business owner is responsible for developing and maintaining an application in the platform of choice—a PaaS provider is responsible for supplying a (pre-configured) platform with enough capabilities to support the business owner application.

The PaaS provider is responsible for its composition; deciding which platform components are present or absent, and how platform components work with each other and the applications they support. However, a PaaS' composition will determine what kinds of applications can be deployed on a platform, since an application cannot operate on a platform that lacks a component the application requires. Thus, PaaS composition can act as both an *enabling* and *limiting* factor, driving the range of applications a platform can support. With PaaS, operating system features can (and likely will) be changed and upgraded frequently (likely without input from the business owner.) That means that business owners should ensure that software developed to operate in a PaaS environment is flexible (generic or standards-based) enough to withstand (forced) upgrades to the PaaS components.

Figure 4 PaaS Security Controls Responsibilities

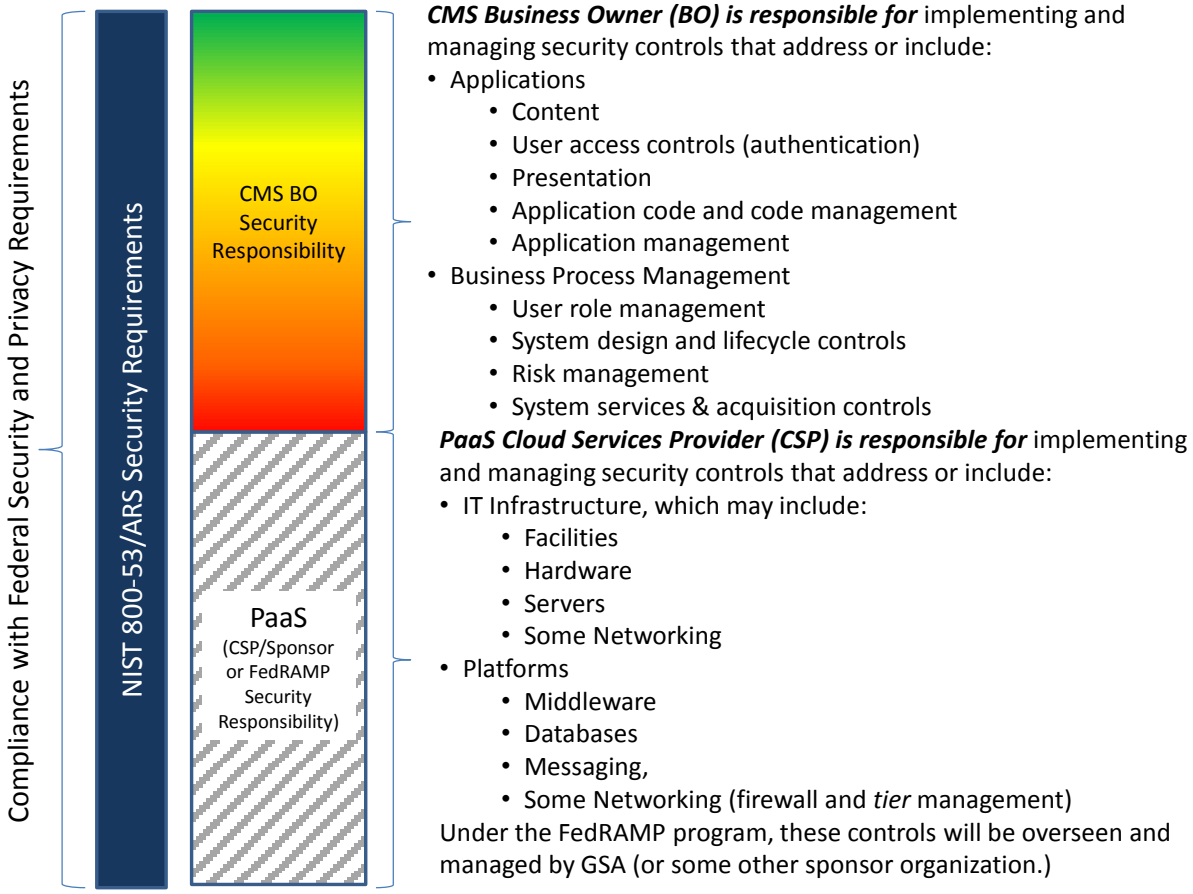

4.4.3 SOFTWARE AS A SERVICE (SAAS) SECURITY

Software as a service (SaaS), sometimes referred to as *software-on-demand*, is software that is deployed over the internet and/or is deployed to run behind a firewall on a local area network. With SaaS, a CSP licenses an application to customers either as a service-on-demand, through a subscription, in a *pay-as-you-go* model. The advantages include:

- Accessible from anywhere with an internet connection

- No local server installation
- Pay-per-use or subscription based payment methods
- Rapid scalability
- System maintenance (backup, updates, security, etc.) often included in service
- Reliability
- Easier administration
- Automatic updates and patch management
- Compatibility: All users will have the same version of software.
- Collaboration

SaaS is software that is developed and hosted by the SaaS CSP that the CMS business owner would access over the Internet (or other network.) Unlike traditional packaged applications that users install on their own computers or servers, the SaaS CSP *owns* the software and runs it on computers in the CSP data center. The CMS business owner does not own the software but effectively rents it, usually for a flat monthly fee per user. Therefore, SaaS is basically subscription-based *commercial off-the-shelf* (COTS) software hosted on a CSP server. SaaS implementations are cheaper because business owners do not have to have to buy additional hardware or infrastructure to make the software work, so there are little (if any) capital expenditures. Business owners generally like SaaS due to the low up-front investment and predictable expense stream. However, depending on the product, the cost advantages of the SaaS model may be a wash after three to five years of monthly fees.

In general, SaaS solutions work best for non-strategic, non-mission-critical processes that are simple, standard, and not highly dependent on, or integrated with, other business functions and systems. SaaS also works well for processes that are being automated for the first time, because there are no legacy processes to replace and thus fewer migration challenges.

Upgrades to SaaS products tend to be almost seamless (although business owners typically cannot *customize* the software) and, unlike with traditional packaged software, *every* user is always on the most recent version of the application. With made-for-SaaS software, upgrades are more effectively pushed out in a series of small changes, rather than large versions, which leads to less user confusion over changes/upgrades, and less overall cost for user training.

However, business owners should have legitimate concerns about keeping their data in a SaaS vendor's systems because they really have little or no direct control over those systems. Business owners should make sure they would always have access to the data and the software (or possibly source code) in the event the CSP goes out of business. It is for this reason that CMS encourages the use of FedRAMP or GSA approved CSPs to ensure those contractual precautions and limitations are available for CMS business owners.

Business owners are accountable for the integrity of the systems used, regardless of who is supplying those systems. If business owners decide to use SaaS, they should make sure they get the same auditing and control requirements from the SaaS vendor that they would get from any third-party provider, including *"safe harbor"* provisions for ensuring data privacy and the ability to audit the CSP's controls.

Figure 5 SaaS Security Controls Responsibilities

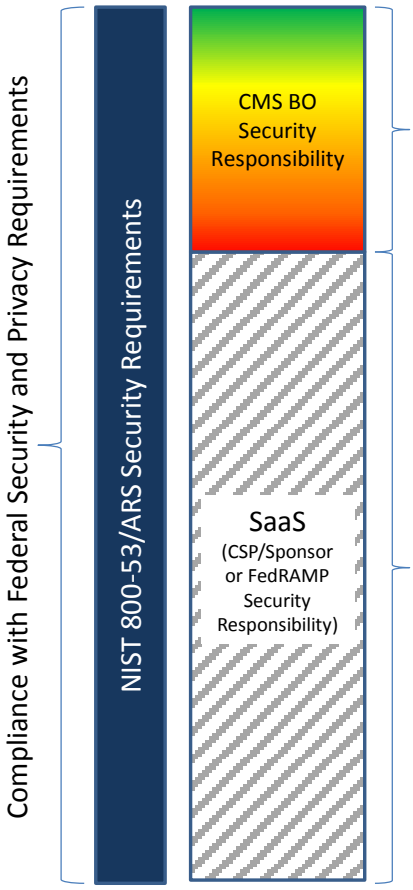

CMS Business Owner (BO) is responsible for implementing and managing security controls that address or include:

- Business Process Management
  - User role management
  - System design and lifecycle controls
  - Risk management
  - System services & acquisition controls

SaaS Cloud Services Provider (CSP) is responsible for implementing and managing security controls that address or include:

- IT Infrastructure, which may include:
  - Facilities
  - Hardware
  - Servers
  - Some Networking
- Platforms
  - Middleware
  - Databases
  - Messaging,
  - Some Networking (firewall and *tier* management)
- Applications
  - Content
  - User access controls (authentication)
  - Presentation
  - Application code and code management
  - Application management

Under the FedRAMP program, these controls will be overseen and managed by GSA (or some other sponsor organization.)

4.5 ACCEPTABLE-USE MATRIX

Cloud-based IT Services carry with them security concerns beyond those associated with traditional IT Service deployment methods. The simple movement of service-delivery mechanisms beyond the comparatively-secure walls of the enterprise datacenter exposes those services to a proportionally-higher level of risk.

The risks associated with Cloud Computing are still being defined and discovered, but it is becoming clear that these risks, like any other IT or business risk, can be mitigated and controlled with a combination of solid governance, relevant and enforceable policies and the implementation of security practices at the earliest stages of application and service development.

Table 2 reflects current HHS guidance[30] based on the intersection of security level (*High*, *Moderate*, or *Low*) with *cloud service and delivery models*. This table is provided to assist

[30] Guidance derived from the *HHS Reference Architecture for Cloud Computing* (Dated August 5/2010) and the *HHS Cloud Computing Implementation and Governance, Alternative Analysis and Supporting Process* (Version 1.0, Dated March 8, 2011)

business owners and system designers in determining the feasibility of hosting a candidate system at a CSP.

Table 2 CMS Cloud Computing Model and Accessibility Matrix

				Deployment Models							
				Private		Community		Public (Gov-Grade)[0]		Hybrid[2]	
		Service Category	Data Security Level	Gov. Provider (GFE)	Commercial Provider	Gov. Provider (GFE)	Commercial Provider	Gov. Provider (GFE	Commercial Provider	Gov. Provider (GFE)	Commercial Provider
Service Models	IaaS	Complex Custom Applications	**Low**	A (Somewhat fit for use)	A (Somewhat fit for use)	A (Somewhat fit for use)	A (Somewhat fit for use)	A (Somewhat fit for use)	A (Somewhat fit for use)	N/A	N/A
			Moderate	C (Somewhat fit for use)	C (Somewhat fit for use) **CMS or FedRAMP ATO Only**	C (Somewhat fit for use)	C (Somewhat fit for use) **CMS or FedRAMP ATO Only**	C (Somewhat fit for use)	C (Somewhat fit for use) **CMS or FedRAMP ATO Only**	N/A	N/A
			High	C (Somewhat fit for use)	TBD (Not fit for use)	C (Somewhat fit for use)	TBD (Not fit for use)	TBD (Not allowed at CMS)	TBD (Not fit for use)	N/A	N/A
	PaaS	Targeted to a particular platform, such as .NET, LAMP, or JAVA	**Low**	A (Fit for use)	A (Fit for use)	A (Fit for use)	A (Fit for use)	A (Somewhat fit for use)	A (Somewhat fit for use)	N/A	N/A
			Moderate	C (Fit for use)	C (Fit for use) **CMS or FedRAMP ATO Only**	C (Fit for use)	C (Fit for use) **CMS or FedRAMP ATO Only**	C (Somewhat fit for use)	C (Somewhat fit for use) **CMS or FedRAMP ATO Only**	N/A	N/A
			High	C (Fit for use)	TBD (Not fit for use)	C (Fit for use)	TBD (Not fit for use)	TBD (Not allowed at CMS)	TBD (Not fit for use)	N/A	N/A
	SaaS	Standardized full-featured Web (thin client) accessible applications	**Low**	A (Somewhat fit for use)	A (Not fit for use)	A (Fit for use)	A (Somewhat fit for use)	A (Not fit for use)	A (Not fit for use)	N/A	N/A
			Moderate	C (Somewhat fit for use)	C (Not fit for use)	C (Fit for use)	C (Somewhat fit for use)	C (Not fit for use)	C (Not fit for use)	N/A	N/A
			High	C (Somewhat fit for use)	TBD (Not fit for use)	C (Fit for use)	TBD (Not fit for use)	TBD (Not fit for use)	TBD (Not fit for use)	N/A	N/A

Legend

Code	Accessibility
TBD	Unknown
N/A	Not Applicable
A	All – Public, Partners, & Gov
B	Gov
C	Partners & Gov

	Recommendation	
1ST	PaaS	Private/Gov & 3rd party CSP
2ND	PaaS	Community/Gov CSP
3RD	SaaS	Public/Gov & 3rd party CSP

Model Preference

1st – Fit for use
2nd – Somewhat fit for use
3rd – Not fit for use

Assumptions: CSPs comply with ALL regulatory and NIST data and data center requirements for government.

Footnotes:

1. CMS business owners should avoid *commercial-grade* public clouds, as there are not enough assurances for security and privacy to meet federal security and privacy requirements. In addition, public cloud CSPs typically contractually-preclude their customers from access or knowledge of the inner-workings of their infrastructure, making it impossible to achieve security and privacy compliance.
2. *Hybrid Clouds* are evaluated by the more restrictive requirements of their components. The most restrictive component cloud factors will be applied.

5 APPROVED

Charles R. Brewer

Digitally signed by Charles R. Brewer
DN: c=US, o=U.S. Government, ou=HHS, ou=CMS, ou=People, cn=Charles R. Brewer, 0.9.2342.19200300.100.1.1=2000389294
Date: 2011.05.03 09:54:24 -04'00'

C. Ryan Brewer
CMS Chief Information Security Officer and
Director, Office of the Chief Information Security Officer

This document will be reviewed periodically, but no less than annually, by the Office of the Chief Information Security Officer (OCISO), and updated *as necessary* to reflect changes in policy or processes. If you have any questions regarding the accuracy, completeness, or content of this document, please contact the OCISO at mailto:ciso@cms.hhs.gov.

(This Page Intentionally Blank)