MITRE

Cybersecurity in the Cloud

The federal landscape for secure cloud services, systems, and solutions

Don Faatz, Mari Spina
1/17/2017

Location (McLean, VA)

Approved for Public Release; Distribution Unlimited. Case Number 17-0247

This page intentionally left blank.

Executive Summary

Deploying data and applications to a cloud computing environment, whether private, community, or public, changes an organization's information technology (IT) security risk profile. This is not to say that risk goes up in the cloud. However, the outsourcing of operational activities, the relinquishment of control over infrastructure components, and the sharing of environments and systems with untrusted entities modifies the threat vector domain and therefore the risks.

All cloud environments utilize new software layers, such as virtualization technologies, within the IT infrastructure. While community and public offerings may employ similar technologies to private environments, the security implications of community and public clouds are more complex. Use of these offerings changes the risk profile because some security responsibility is transferred to the cloud service provider (CSP), and the organization's security perimeter is extended to include the provider's computing resources and personnel. Given these changes, organizations need to understand the risks and appropriate mitigations. This paper seeks to highlight the Federal perspective and considerations for moving from on-premise private to off-premise public/community clouds.

Key to ongoing risk management of cloud-based applications and data is the client's ability to have continuous situational awareness of network status and cyber events. Incident detection and response will be a coordinated effort, requiring data exchange between the partners.

In order to appropriately assess the risks, an understanding of the threat actors and their techniques is needed. The ability to share indicators of compromise and coordinate mitigations will be necessary to prevent and interrupt attacks.

Resiliency and business continuity are key considerations that all application owners must consider. Cloud deployments can offer some advantages in this area. CSPs may offer automated data replication and back-up, and the ability to rapidly restore applications and data that have been compromised. To protect the availability of the data, Federal IT leaders should ensure they have out-of-cloud backups or backups provided by multiple cloud vendors for critical data, ensuring no single point of failure with a given cloud provider. [1]

While vendors have many platform specific security controls unique to their own offerings, encryption provides a consumer-controlled mechanism to protect data moved outside the organization's security perimeter to a public or community cloud. Effective solutions for encryption -in-transit and -at-rest are available. However, processing of data will continue to require decryption for the foreseeable future, opening a window of opportunity for the aggressive attacker. Regardless, organizations should verify that either the CSP's encryption capabilities meet their data protection needs or that additional encryption capabilities can be provided. To support encryption capabilities, a secure key management and restoration capability must be provided. In some cases, legacy applications must be modified to employ or interface with encryption modules.

For applications accessed from a community or public cloud, the organization's Identity and Access Management (IdAM) capabilities will need to be extended to support cloud-deployed applications. The cloud provider may offer IdAM capabilities that can interface to or federate with the organization's capabilities. In this case, the cloud provider capabilities may readily meet

the needs of government organizations. However, if the CSP's capabilities are inadequate, organizational capabilities will need to be replicated or extended into the cloud.

While there are risks with moving capabilities to an external cloud, there also are potential advantages. Cloud providers may be able to better manage infrastructure security concerns such as system configuration and patch management. In addition, economies of scale and homogeneity of infrastructure can give providers advantages in terms of operational cost. They may also have more highly skilled security operations personnel, and a more mature security operations center.

Despite the advantages that the CSP gains through scale and homogeneity, government IT monitoring of community and public cloud-based applications is complicated by the loss of direct control. Organizational security operations and incident response teams may need to develop new approaches to monitoring cloud-deployed systems and learn to combine their data with data from service providers. Government IT teams will need to partner with cloud providers and adjust their detection and response procedures to include this relationship.

Contract Terms and Conditions (T&Cs), Service Level Agreements (SLAs), Organizational Level Agreements (OLAs), and Privacy Level Agreements (PLAs) can be used to manage the relationship between IT organizations and cloud service providers. But actions that move the provider away from promises made in their standard user agreements can increase cost and cause migration delays.

Given the security changes that result from deploying a cloud-based approach, Federal IT leadership should understand the risks and potential mitigations. While private clouds incorporate new technologies into the IT stack that need to be secured, community and public clouds additionally introduce risks due to reduced control and visibility. With these deployment models, the key to secure use of cloud computing is shared understanding of the division of security responsibilities between provider and government client, and the ability to verify that both are meeting their responsibilities. The Federal Risk and Authorization Management Program (FedRAMP) provides information and services to assist government organizations to understand and verify cloud service provider security practices.

Acknowledgments

This paper would not have been possible without the support and contributions of Mindy Rudell, Deirdre Doherty, Brian McKenney, and Prem Ramamurthy

Table of Contents

List of Figures

List of Tables

This page intentionally left blank.

1 Introduction

This paper examines some of the consequences of cloud technologies, shared security responsibilities, and virtual boundaries. It describes issues that organizations planning to use cloud-based computing resources should consider. While it does not offer specific solutions to these challenges, it does provide pointers to cloud service provider guidance and the Federal Risk and Authorization Management Program (FedRAMP) which aid in addressing security challenges. Some of the challenges discussed are unique to cloud computing while others apply to both cloud computing and information technology (IT) outsourcing generally.

From the perspective of information security, cloud computing elicits one of two responses:

- Security issues make cloud computing very risky.
- "Security issues are more perceptual than prohibitive [2]."

Paradoxically, both positions have merit. Along with the potential benefits, this model of computing resource delivery presents Federal IT leaders and security architects with new risks that must be understood and addressed. A better understanding of risks associated with cloud computing can help in identifying appropriate ways to use this IT approach.

The paper begins with a brief description of cloud computing. It considers information security in the clouds from three perspectives—protecting data, protecting systems, and defensive operations.

1.1 Controlling Environments

Moving IT systems to the cloud doesn't necessarily increase risk. But two notable characteristics of public and community cloud computing alter the security posture of an organization's IT program:

- Sharing of security responsibility between the Cloud Service Provider (CSP) and client.
- Expansion of the enterprise security perimeter to the CSP.

Public and community cloud deployment models cede direct control of computing resources to a CSP in exchange for the potential of reduced costs and/or additional capabilities. This transfer of control also transfers some information security responsibilities to the CSP. However, the information owner retains ultimate responsibility and accountability for appropriately protecting the information. The National Institutes of Standards and Technology (NIST) explains this situation in Special Publication 800-53 "Organizations are accountable for the *risk* incurred by use of services provided by external providers … [3]."

Public and community cloud computing also impact the enterprise's computing perimeter, extending it into the cloud of shared resources and making it difficult to define boundaries. Physical boundaries are replaced by virtual boundaries, eliminating the utility of physical separation as risk mitigation and opening the door to attacks from collocated adversaries.

In contrast, consider an organization with poor cybersecurity defense systems and operations. If a move to the cloud results in improved vulnerability mitigations, risk can very well be reduced. Applications and data vary widely in terms of their sensitivity and vulnerability to attack. For this reason, a risk based approach to determining which applications to migrate to the cloud, and which cloud model, cloud service, and CSPs are acceptable, is recommended. The NIST Risk Management Framework [22] provides a structured process for this evaluation.

1.2 Cloud Computing Defined

NIST defines cloud computing as [4]:

> "Cloud computing is a model for enabling ubiquitous, 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 and released with minimal management effort or service provider interaction. This cloud model is composed of five essential characteristics**,** three service models, and four deployment models."

NIST has identified five essential characteristics of cloud computing:

- **On-demand service**: Allows government IT organizations the ability to provision resources as they need them without provider intervention.
- **Broad network access**: Provides access to resources via standard network mechanisms.
- **Resource pooling**: Allows providers to use the same physical resources to provide service simultaneously to different clients.
- **Rapid elasticity**: Allows clients to increase or decrease resources allocated to them possibly without human intervention.
- **Measured service**: Monitors and controls the delivery of resources to consumers ensuring they "get what they pay for" and "pay for what they get."

Broad network access and resource pooling are of particular interest from a security perspective. Both contribute to modifying the enterprise perimeter and potentially increasing the exposure of data and applications.

NIST describes three service models: Software as a Service (SaaS), Platform as a Service (PaaS), and Infrastructure as a Service (IaaS).

- **SaaS** provides the consumer use of applications and information storage in a cloud. Google's Gmail email application and Google Docs office automation applications are examples of SaaS.
- **PaaS** provides the consumer an application-hosting environment for consumer-created or acquired applications. The environment may provide a programming language, a set of application services, data storage, and network connectivity. Google's Google Apps Python-based Web application platform and Microsoft's Azure .Net-based application platform are examples of PaaS.
- **IaaS** provides consumers virtual computing infrastructure, networks, and access to persistent data storage. Consumers can configure the virtual infrastructure and run custom software on top of virtual instances of operating systems such as Microsoft Windows or Linux. Amazon's Elastic Cloud Computing (EC2) service is an example of IaaS.

The cloud service model affects where the line is drawn in transferring security responsibilities to the cloud service provider. With IaaS, responsibility for the security controls within the physical infrastructure and the virtualization layer are transferred to the provider. Responsibility for client operating systems, middleware, and application security remain with the consumer. With PaaS, the provider assumes responsibility for the physical infrastructure, the virtualization layer, and any provided middleware (e.g. .net framework) but responsibility for application security controls remains with the consumer organization. With SaaS, the provider generally assumes all responsibility for security controls offering only application user accounts to the consumers.

Additionally, consumers may be able to leverage application program interfaces (APIs) made available by IaaS and PaaS providers to implement system security. APIs allow consumers an ability to interact with or integrate provider services with deployed systems.

Table 1illustrates how ownership of security controls varies by service model.

Table 1. Ownership for Security Controls Varies by Service Model

	Infrastructure Security	Application Security	Security APIs
SaaS	Cloud Provider	Cloud Provider	N/A
PaaS	Cloud Provider	Cloud Provider and Government Consumer	Potential client use of provider's application security APIs
IaaS	Cloud Provider	Government Consumer	Potential client use of provider's application security APIs

Deployment of the three service models is possible in any of four different deployment models: public cloud, community cloud, private cloud, or hybrid cloud.

- **Public cloud** services are available to the public from a cloud services provider.
- **Community cloud** services are available to a specific community. These services may be provided by a cloud services provider or offered collectively by the community members.
- **Private cloud** services are available only to a single organization. These services may be provided by the organization itself or by a third party.
- **Hybrid cloud** services are a combination of two or more of the other deployment models.

The deployment model determines both degree of control over changes and how the enterprise perimeter is affected. In moving from a traditional model of organization owned and operated resources to a cloud model, public clouds represent the largest change in risk posture. Public and community clouds extend the enterprise perimeter to include the provided services and network paths to those services. A private cloud may not alter the enterprise perimeter. Community clouds vary depending on the number and type of community members. Community clouds with a small number of members with similar characteristics may resemble private clouds, while those with large numbers of diverse members will closely resemble a public cloud.

Although private cloud technologies expose new security issues, they represent less change from a security perspective than community and public clouds. As a result, many security concerns discussed in this paper are relevant for community and public clouds, but may not be directly applicable to private clouds. For more information on private cloud security products, refer to [39].

2 Protecting Data in the Cloud

As data is moved into public and community clouds, physical control over the location of the agency's data is reduced to that which can be contractually specified, such as requiring that data only be stored within national boundaries. Authentication and authorization of identities accessing the data, requires a federated approach. Further, many different tenants employ the CSP's resources concurrently. As a result, there are increased threats associated with colocation of attackers and the malicious insiders. However, use of encryption and Identity and Access Management (IAM) systems represent effective mitigations to threats that would act to compromise confidentiality and integrity.

2.1 Encode the Data at Rest

Encryption provides a form of encoding that renders data indecipherable by anyone not possessing the encryption key. Many cloud providers offer some form of encryption for data stored or transmission within their clouds. For example, the Microsoft Windows Azure PaaS offering makes Cryptographic Service Providers available through the .NET Framework APIs [5]. In Amazon Web Services' IaaS offering, both Elastic Block Storage (virtual disk drives) [23] and Simple Storage Service (S3) [24] provide options to encrypt stored data.

It is important to understand the exact type of protection that is offered by the provider's encryption and how it is administered, before accepting it as adequate. Key generation and management is typically addressed in one of three ways: 1) keys can be generated and managed by the consumer, 2) keys can be generated by the consumer but managed by the CSP, 3) Keys can be generated and managed by the CSP. Each approach has its risk profile implications but the goal of encryption is to limit exposure to adversaries. Today, we are generally able to cost effectively employ encryption-at-rest and encryption-in-transit technologies. However, there is still no practical means of processing encrypted data. This means that there is always a window, although potentially small, where data must be unencrypted to be processed.

The ability to control encryption processes also varies by service delivery model.

- With SaaS, clients rely upon the cryptographic capabilities the CSP and application developers have built into the application. Clients need to verify that either a SaaS application provides appropriate cryptographic capabilities or confidentiality protection is not needed for the data that the application will process.
- With PaaS, clients rely primarily upon the cryptographic capabilities available in the CSP's application hosting environment. A PaaS hosting environment may allow installation of third-party products such as encryption software and public key certificates. If not, clients need to verify that the encryption capabilities of the hosting environment will meet the needs of their applications.
- IaaS offers the most flexibility with respect to encrypting data. Since the CSP supplies a virtual machine running an operating system, the client can install and use third-party encryption software.

For example, the government client organization needs to understand how keys are managed and who has access to the keys. The need for effective key management is not unique to cloud computing but is important any time encryption is used. Keys used in IaaS and PaaS service models should be unique to each client and only accessible by client personnel. In all cases, key management plans need to ensure keys are adequately protected when used and include provisions for key escrow to protect against data loss due to loss of encryption keys. AWS

provides a key management service [25] and dedicated hardware security modules [26] to help consumers manage encryption keys. Additionally, for Federal government clients, encryption and digital signature capabilities need to use cryptographic algorithms approved by the NIST and the implementations need to be Federal Information Processing Standard (FIPS) 140-2 validated.

For added security, key management can be performed by a FIPS 140-2 compliant management system located outside of the CSP in either a 3rd party or on-premises within the agency. Issues associated with the transmission of keys and the effect of associated latency on application operations would then need to be addressed.

2.2 Encode the Data in Transit

Data traversing back and forth between the consumer's security boundary and the cloud provider will need to be protected while in transit. It may cross multiple network providers, including the public Internet. Encrypting data in transit is the primary means of providing this protection. For connections between cloud-deployed applications and users, Transport Layer Security (TLS), as the successor to Secure Sockets Layer (SSL), should be used. For connections between an organization's datacenter and the cloud, or between internal applications and cloud-deployed applications, more persistent encrypted connections such as virtual private networks (VPNs) can be used.

Despite the protection provided by encryption, organizations may not be comfortable with their data traversing the public Internet or multiple network providers. Additionally, public network based communications paths may introduce unacceptable latency in communications. To address these concerns, some cloud providers offer direct connections between an organization and the cloud resources the organization uses. For example, Microsoft's ExpressRoute [30] provides layer 3 connectivity between an organization's network and Azure via a connectivity provider. This provides a dedicated connection in addition to encryption.

It is also important to note that public and community cloud providers will typically provide an array of services to their consumers through Application Program Interfaces (APIs). API servers are quite often accessible via the Internet. CSP's build them like this because they want to provide the broadest access and highest availability possible to their consumer base. Unfortunately, this renders the API servers vulnerable to Internet-based threats[1]. Application developers may assume such communications are private and neglect the need to secure communications to the CSP's API servers.

Encryption-in-Transit services, such as HTTP Secure (HTTPS), can be used to protect access to API server end-points and are typically offered and supported by cloud providers. For example, AWS provides API server end-points that run HTTPS. Directing calls to the AWS HTTPS enabled API servers provides server authentication by the calling entity and establishment of a Transport Layer Security (TLS) encrypted tunnel[2].

2.3 Narrow the Attack Window

Even with encryption, data always will have a window of vulnerability because currently it cannot be processed while it is encrypted[3]. A risk assessment is needed to determine the data and

[2] Making Secure Requests to Amazon Web Servicew, https://aws.amazon.com/articles/1928 .
[3] IBM research [7] has discovered a technique that might eventually allow computation on encrypted data; however, the technique is not likely to be practical for many years [8].

applications for which this exposure is acceptable and are candidates for cloud deployment. The degree of exposure is dependent on the cloud deployment model, the relative sophistication of the cloud provider's security controls, and the client's ability to augment the provider's controls. Assuming equivalent security capabilities, public clouds represent the greatest exposure and private clouds the least. This is a result of a reduced ability to control access to infrastructure and computing resources. When you have little control regarding who can share the systems you use to process data, as is the case with community and public clouds, your exposure to collocated attacker or malicious insider threats is a function of tenant composition[4].

Applications that process information requiring cryptographic protection need to be designed to control exposure. This is accomplished by using techniques such as minimizing the time data is decrypted and clearing storage locations after use. Applications currently used inside an organization may need a security review and enhancement of security deficiencies before deployment to a public or community cloud.

2.4 …. And Throw Away the Key

Having secured the data stored in the cloud to the best degree possible through data handling security controls, the next question to address is, what happens to persistent data when it is deleted? For example, when deleting data in a commercial cloud, there is very little ability for the data owner to verify that the deletion has occurred across the entire system through which it was distributed. Additionally, the data owner has little control or insight into which erasure technique is used, unless specified contractually.

It is important then to understand if and how the cloud provider clears the data when it is deleted or when the client stops using it. For IaaS and PaaS deployments, applications may need to clear persistent storage themselves when deleting data or releasing storage if the provider does not perform this adequately. However, this is not possible for all types of persistent storage. In AWS, for example, applications can overwrite data stored in Elastic Block Storage (EBS) just as they would overwrite data on a physical disk. Simple Storage Service (S3), however, is a write-once, read-many storage service and cannot be overwritten. Hence, consumers are dependent on the Amazon to effectively clear S3. With SaaS, the only options are those provided by the provider. If the provider-supplied data clearing capabilities are not sufficient for a user organization, user-controlled encryption may be sufficient if the encryption keys can be securely deleted.

However, note that key deletion remains an option only until computing capabilities and resources necessary to crack today's cryptography remain unavailable, unobtainable, or overly expensive. On the forefront is Quantum Computing which may render today's cryptography more vulnerable.[5]

2.5 Manage Access

Moving data and applications to a cloud means the Identity and Access Management (IdAM) capability must expand to encompass cloud-based resources. Unlike the private network environment offered by the traditional data center, generation and management of cloud consumer account credentials must be performed using the CSP's IdAM systems. The CSP's IdAM systems typically offer consumers the ability to configure fine-grained controls that implement least privilege, role-based access (RBAC), and multi-factor authentication (MFA)

[4] Section 4.1.1 of this paper discusses malicious insiders.

[5] https://www.washingtonpost.com/world/national-security/nsa-seeks-to-build-quantum-computer-that-could-crack-most-types-of-encryption/2014/01/02/8fff297e-7195-11e3-8def-a33011492df2_story.html

policies. But while consumers can generally only provision and configure these services for their accounts and environments, consumers do not control the CSP's IdAM system. The consumer must therefore trust that the CSP is managing their IdAM system securely.

Most organizations will have a fairly robust enterprise IdAM capability for creating user accounts, providing authentication credentials, and managing user authorizations. To perform these functions with a community or public cloud, either enterprise IdAM capabilities must be integrated with cloud provider IdAM capabilities, or enterprise capabilities must be exposed to the cloud.

Cloud providers offer the ability to federate cloud-based IdAM with enterprise capabilities using standards-based mechanisms such as Security Assertion Markup Language (SAML) tokens. In federated IdAM, cloud-based services depend on existing enterprise IdAM services to authenticate users and authorize the use of services. Decisions from the enterprise services are sent via SAML tokens to and enforced by cloud-based services. Other providers have developed "connectors" that link to client enterprise directories to provide IdAM services.

If a provider does not offer federation or connector capabilities, clients may need to extend portions of their enterprise IdAM functionality into the cloud or manage identities for cloud services directly in the cloud. Care must be taken in making this extension, since enterprise IdAM infrastructure data is sensitive. For example, an organization's active directory environment could be replicated to support cloud-based IdAM, but doing so would expose more information than necessary. Identity, credential, and authorization information also may be exposed during provisioning and use in public or community clouds. A carefully controlled extension of enterprise IdAM will need to be implemented to control and minimize this exposure. Microsoft's Office 365 SaaS offering, for example, provides support for multiple IdAM approaches including federation, replication of some identity data to the cloud, and direct cloud-based management of identities [29].

Another solution for IdAM federation is the use of 3rd party providers that act as Identity Brokers or a Cloud Access & Security Brokers (CASB)[6]. These entities assist in bridging the IdAM divide and are well suited for handling the integration of multiple CSPs.

.

2.6 Lock Up the Root Credential

The first set of credentials received from a CSP to open an account, effectively provide the equivalent of unfettered physical and logical access in a traditional data center. In a traditional data center, personnel such as system administrators and software developers may be given physical access to resources. Logical access controls limit their ability to modify server and application configurations. In cloud computing, the initial credential provided to a consumer is able to create new virtual resources, and modify or delete existing virtual resources. Additionally, this is the only account that generally cannot be restricted in privilege. Some have called this the "god" account because it allows for the provisioning of all CSP services and systems.. Holders of associated account credentials are very powerful and can, if malicious, cause considerable damage.

[6] https://www.skyhighnetworks.com/cloud-university/what-is-cloud-access-security-broker/

AWS provides an IAM Best Practices Guide[7]. The first best practice prescribed by AWS is to lock away account root access. There are ways to do this that range from simple to complicated. Simply guarding and protecting the log-in and password is the simplest, but a much stronger approach can be implemented by enabling Multi-Factor Authentication (MFA) on the account and then diligently controlling access to the one-time password (OTP) device that implements the second factor. In this way, the root account cannot be employed until access to the OTP device is gained. AWS also provides an IAM service that allows the definition of roles for within the virtual environment. Using IAM, developers and system administrator access can be limited to only those actions needed to perform their jobs (least privilege). These least-privilege user roles must be defined ahead of time for all possible actions. Otherwise, the OTP devicewill need to be unlocked regularly to accomplish development and operational tasks.

2.7 Integrate Privacy

In addition to the concerns for protecting sensitive data in the cloud, if personally identifiable information (PII) is stored or processed, it is important to consider the privacy-specific protection requirements that must be met. The American Institute of Certified Public Accountants, Generally Accepted Privacy Principles, defines privacy as, "The rights and obligations of individuals and organizations with respect to the collection, use, retention, and disclosure of personal information [6]." For example, the Health Insurance Portability and Accountability Act (HIPAA) requires the protection and confidential handling of protected health information. Initially, it would seem that data protection should address any difference between the internal processing of PII and cloud-based processing. However, privacy issues can materialize in unexpected ways. The cloud vendor's terms of use might grant the vendor some rights to information stored or processed in their cloud. For example, the AWS Customer Agreement says "We will not access or use Your Content except as necessary to maintain or provide the Service Offerings, or as necessary to comply with the law or a binding order of a governmental body.[8]"

Hence, if PII is stored or processed, it is important to review the privacy protection responsibilities and how they are met by the CSP. Privacy overlays to NIST security have been developed[9] to provide guidance. But if there is any uncertainty, be sure to specify agreements on privacy related roles and responsibilities in the contract with the CSP. To this end, the Cloud Security Alliance (CSA) has developed the Privacy Level Agreement (PLA)[10]. Think of it as a Service Level Agreement (SLA) for privacy.

Additionally, it is important to understand the geographic location where storage and processing of your data will occur. Privacy policies and regulations vary across international boundaries[11]. Jurisdiction over the data tends to reside within the municipality, state, or country in which the data is stored regardless of the reach or operations of the CSP. Understanding the variations and limitations when it comes to data retrieval for forensic or legal purposes can get complicated. Moreover, related regulation continues to change as the industry takes shape and the balance between consumer and citizen rights to privacy are weighed against safety and public interest.

---

[7] http://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html#lock-away-credentials
[8] https://aws.amazon.com/agreement/
[9] https://www.cnss.gov/CNSS/openDoc.cfm?WAYxsitMb7tA394EPT7m+g==

[10] https://cloudsecurityalliance.org/group/privacy-level-agreement/
[11] https://techcrunch.com/2016/08/24/encryption-under-fire-in-europe-as-france-and-germany-call-for-decrypt-law/

The best rule of thumb in this regard is to specify allowable geographic locations for data storage and processing by selection of provider, configuration of resources by geographic region, and in the contract with the CSP. In this way, you can specify your comfort zones knowing full well the level of effort you may need to expend to protect or retrieve your data.

2.8 Key Considerations

The key considerations identified in this section for protecting data in cloud deployments are:

- Encryption and digital signatures are the primary confidentiality and integrity protection for data stored in or transmitted to a public or community cloud.
- While being processed, data may be vulnerable while being processed in a public or community clouds because it must be unencrypted.
- Unless cleared by the consumer, data may remain in persistent storage when the storage service is released.
- Existing internal applications may need analysis and enhancement to operate securely in a public or community cloud.
- Data replication provided by a cloud provider may not be a substitute for backing up critical data to another independent provider or out of the cloud.
- Privacy protection responsibilities should be reviewed if considering moving PII or PHI to the cloud. The contract, SLAs, and PLAs can be effective CSP contracting tools for specifying related data handling requirements.
- Cloud IdAM capabilities vary widely. Integration or federation of provider IdAM with an organization's exiting IdAM capabilities must be considered.
- CSPs IdAM solutions can be leveraged to provide fine grained control for the implementation of least-privilege RBAC over provider managed systems and consumer migrated data.
- The initial credentials used to establish the cloud service account possess root privileges. Special care should be taken to controls and securely manage them.

3 Protecting Systems in the Cloud

Protecting the data moved to the cloud is but one element of an effective security posture. But a move to a public or community cloud is tantamount to extending the organization's perimeter to include the cloud systems and services of the CSP. As such, consideration for the impacts that a perimeter expansion brings is also vitally important.

3.1 Consume Authorized Clouds

FedRAMP[12] was created to reduce the challenges of cloud computing Assessment and Authorization (A&A). It is not cost-effective and efficient for every organization deploying applications to a CSP to define security requirements for the CSP and assess the CSP's success in satisfying the requirements. Because AWS has over a million customers, it is not practical to individually address each customer's need to assess AWS security. FedRAMP provides an approach to assessing a cloud service provider's security once, then sharing that assessment across the federal government.

[12] http://www.fedramp.gov/

The FedRAMP approach begins with defined security control baselines [20] for low, moderate, and high impact cloud services. These baselines are security controls chosen from NIST SP 800-53, with FedRAMP-assigned values for any control parameters. Cloud service providers document how their services satisfy these baseline security controls. The providers' claims are then validated by independent third-party assessment organizations (3PAO). The results of assessment are documented and reviewed by the FedRAMP Joint Authorization Board (JAB). If, in the judgement of board members, a provider has met or, through preparation of a Plan of Actions and Milestone (POA&M), has demonstrated sufficient planning to meet the requirements of the baseline security controls, the provider is granted a FedRAMP Provisional Authority to Operate (P-ATO). Results of the independent assessment are available to cloud consumers to help them understand how a provider secures their services. Departments and agencies use the FedRAMP assessments and the P-ATO as evidence in their assessment and authorization programs.

Agencies need to clearly understand what is addressed by a FedRAMP assessment. The FedRAMP security controls baseline defines the controls a CSP must implement to protect their infrastructure and services. The assessment examines these controls. FedRAMP does not determine the security services that a CSP should offer to clients for use in their applications. NIST defines three categories of security control in SP800-53 [4], common controls, hybrid controls, and system-specific controls. In the context of cloud computing, common controls are security controls implemented by the CSP and common to all consumers. Hybrid controls are controls that are partially implemented by the CSP and partially implemented by the cloud consumer. System-specific controls are implemented by the cloud consumer. An illustration may help clarify these concepts.

NIST SP800-53 security control AC-3 requires an information system to "enforce approved authorization for logical access to information and system resources in accordance with applicable access control policies." Figure 1 illustrates four layers found in a typical IaaS application deployment. In each layer, there are users whose access to information and resources must be controlled to implement AC-3. In the bottom layer, AC-3 is implemented by the CSP as a common control. The CSP both defines the applicable access control policies and enforces approved authorizations. In the layer above this, AC-3 is a hybrid control. The cloud consumer defines the applicable access control policies and the CSP enforces approved authorizations. In the top two layers, AC-3 is a system-specific control with the cloud consumer responsible for both defining access control policies and enforcing approved authorizations.

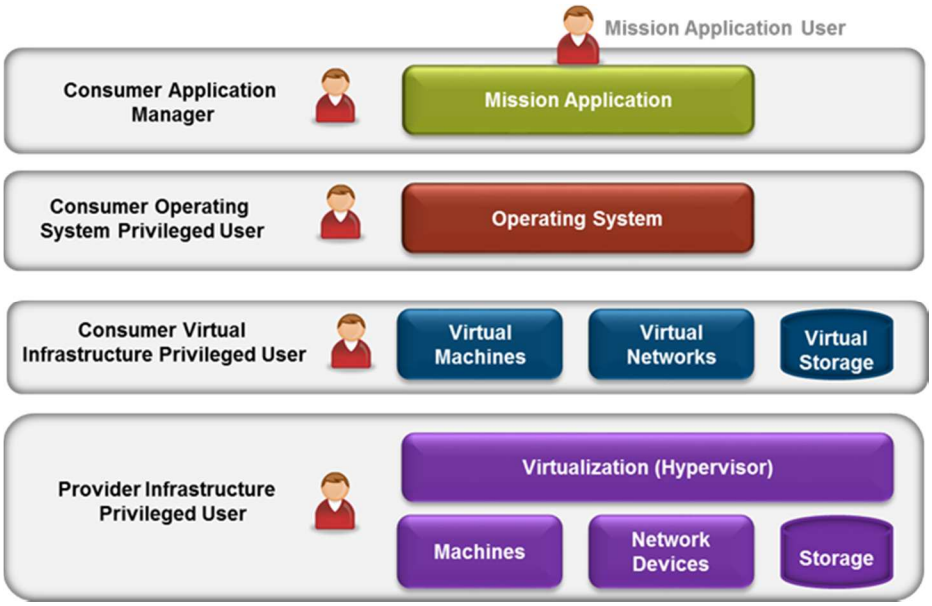

Figure 1. IaaS Service Layers

FedRAMP assessments primarily address common controls. They may also address the CSP-provided portion of some hybrid controls. When it comes to understanding the common and hybrid controls delivered with the cloud service, documentation provided in the CSP's assessment package can help. The Controls Implementation Summary (CIS), provides information on how the CSP has implemented a control. The Customer Responsibility Matrix (CRM) provides information on how the CSP's cloud consumers are expected to employ the security capabilities delivered as a service. Review of both of these documents prior to development is quite useful.

Beyond the initial assessment of a service provider's implementation of baseline security controls, FedRAMP also provides guidelines for continuous monitoring [21]. These guidelines define the frequency of reporting and evidence required for continuous monitoring of the FedRAMP baseline security controls.

In addition to technical information, FedRAMP provides information agencies can use in writing contracts for cloud services. Cloud computing, as noted earlier, involves ceding some security responsibilities to the CSP. A cloud consumer's contract with a CSP determines the services being purchased, including security. Therefore, it is critically important that the contract contain appropriate security clauses. Unfortunately, agency contracting offices often lack experience with cloud computing and security. To help address this, FedRAMP provides templates for standard contract clauses and security control-specific contract clauses [36, 37]. Additionally, the federal CIO Council and the Chief Acquisition Officers Council have jointly published best practices for acquiring information technology as a service [38].

3.2 Use Secure Connectivity

3.2.1 Trusted Connections

To both manage costs and improve security monitoring of network connections between U.S. government departments and agencies and the Internet, the Office of Management and Budget directed the creation and use of Trusted Internet Connections (TICs). All "external connections" carrying network traffic between government systems and external entities, including the Internet, business partners, and state, local and tribal agencies, must transit a TIC which is monitored by the Department of Homeland Security's (DHS's) National Cyber Protection Program (NCPP). The Trusted Internet Connection (TIC) provides a set of security capabilities designed to protect government networks from Internet-sourced attacks. In a three (3) phased initiative defined by National Security Presidential Directive 54/Homeland Security Presidential Directive 23, DHS rolled out TIC standards and policies to allow Federal agencies to deliver TIC solutions as an Access Provider (TICAP). Third Party Commercial entities can do essentially the same under the concept of Managed Trusted Internet Protocol Services (MTIPS). TIC solutions can be built and operated by individual agencies for their own purposes, known as a Single-Service TICAP, or delivered as a broader multi-agency service known as a Multi-Service TICAP.

TIC solutions provide network segmentation, intrusion detection and prevention, and deep packet inspection. The Big Data Cyber Analytic capabilities known as Einstein[13] are most often associated with TIC systems because a standard TIC implementation has the ability to feed raw traffic flows to Einstein for processing and analysis.

When the TIC initiative began in 2005, government organizations did not use cloud service providers as part of their information technology infrastructure and the TIC architecture could be represented very simply as shown

Figure 2.

[13] https://www.dhs.gov/einstein

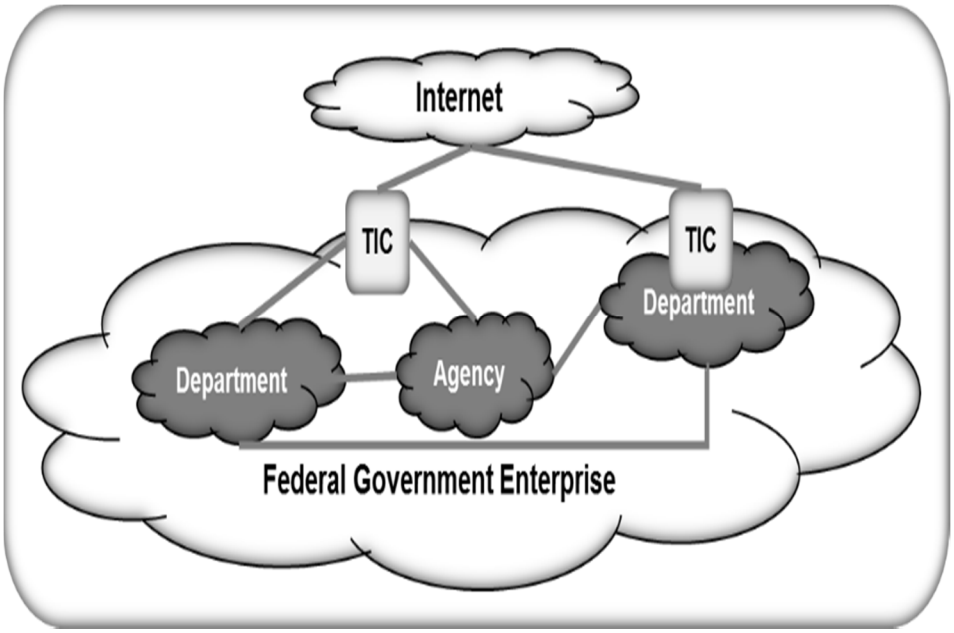

Figure 2. TIC High Level Architecture

The introduction of cloud services into organizations IT infrastructure has significantly complicated this picture. To address the complication, version 2.0 of the TIC reference architecture [31] includes an appendix on cloud considerations. This appendix defines four criteria that determine when interactions between a cloud-deployed application and a federal department or agency must transit a TIC. Additionally, the appendix provides four use cases that describe typical scenarios for interaction between department or agency systems and cloud-provided services. The discussion around each use case provides guidance on how the TIC reference architecture applies to the use case. The guidance requires careful consideration of the D/A's authorization and accreditation boundary, and potential exposure to other tenants and external connections at the CSP.

3.2.2 TIC Ready Providers

DHS and the FedRAMP Program Management Office (PMO) recently released the FedRAMP-TIC Overlay (in draft)[14]. The Overlay provides a selected set of FedRAMP security controls that if achieved, a CSP can be accredited as "TIC Ready". As a result, Federal agencies will be allowed to place private application systems and data currently protected by a TIC into the authorized CSP environment. Users of the systems would no longer be required to authenticate into the Agency's private network to gain access to Agency systems and data. This would allow direct access to systems in the Cloud by the Agency's mobile work force.

FedRAMP-TIC Overlay enables the building of logical network extensions into off-premise commercial cloud environments. Its implied topology acts to reduce the traffic load on the agency's private network and traffic bottle-necking present within the current TIC service structure. A pilot of the FedRAMP-TIC Overlay was conducted with AWS to assess viability of the proposed overlay. As result of this pilot, AWS provides guidance on TIC readiness [33]. If

[14] https://www.fedramp.gov/draft-fedramp-tic-overlay/

the FedRAMP-TIC Overlay is adopted, providers that implement the overlay will offer services that greatly simplify meeting the requirements of the TIC reference architecture, which should lower cost and improve performance.

3.3 Built-In Resiliency

Data and applications stored and processed in a public or community cloud may need to be copied periodically outside the cloud or to another cloud provider with a different IT stack. This will provide backup availability in the event of cloud provider failure or loss of integrity due to a malicious attack. Since CSPs charge for transporting data out of the cloud, the cost of providing an out-of-cloud backup must be considered against the costs of data loss. When planning a solution deployment, it is important to consider how applications and data will be migrated if the provider fails or can no longer meet an agency's needs.

Most, but not all, cloud-based storage includes transparent replication by the CSP. It may be tempting to think of this replication as a fail-safe backup that will protect the client in all circumstances, but it is not. Replication is intended to preserve the "illusion of infinite resource" [1] by ensuring the availability of client data during disruptions to the provider's infrastructure. However, it may not protect the client against significant technical problems or cloud provider business failure. The University of California at Berkley summed it up well, "Even if the company has multiple datacenters in different geographic regions using different network providers, it may have common software infrastructure and accounting systems, or the company may even go out of business. Large customers will be reluctant to migrate to Cloud Computing without a business-continuity strategy for such situations. We believe the best chance for independent software stacks is for them to be provided by different companies, as it has been difficult for one company to justify creating and maintain two stacks in the name of software dependability [1]."

An example of a cloud storage failure is a lightning-related incident at Google's Saint Ghislain, Belgium datacenter in August 2009 [27]. Four lightning strikes to the power grid that supplies the datacenter resulted in the reported permanent loss of 0.000001% of the disk space. Google accepted full responsibility for the loss but reminded customers that, while their datacenters are designed to prevent such incidents, those protections are "no match for Mother Nature." Other large cloud service providers have also experienced data loss incidents at their datacenters. These events help illustrate the need for an appropriate backup strategy for critical data. This strategy may involve replicating data across multiple storage services within a single CSP, replicating to multiple CSPs, or backing up to an on-premises data store. Every solution has costs and risks. For non-critical data, the cost may favor accepting the risk.

Recognizing the value in uptime, CSPs are building in redundancy via implementation of multiple hardware and software stacks, power and communication suppliers, and data center instances. For example, with a geographic region, AWS may have multiple Availability Zones (AZ). Each AZ represents an isolated data center. Consumers can build applications solutions in AWS to take advantage of multiple regions or AZs to manage availability and performance for their users.

Additionally, many CSPs provide auto-scaling capabilities that take advantage of multiple distributed data center resources. For example, AWS provides services that allow applications to automatically scale horizontally. If load increases, new application instances are created. If load decreases, application instances are terminated. If application instances fail, they are replaced with new instances. Special care should be taken when applying auto-scaling capabilities to

applications that depend upon the persistence of transaction or communication session state information. When state persistence is important, migrated applications may require refactoring to take advantage of auto-scaling.

Cyber resiliency is defined as the ability to continue operations and recover quickly in the face of failures or attacks. Commercial cloud service providers such as AWS, Google, and Microsoft design their service to operate 24 hours a day, seven days a week, three hundred and sixty-five days a year with no scheduled down time. Because of the scale of these services, the providers know that failures will occur constantly so their architectures and operating procedures are designed to ensure availability to consumers despite equipment failures. Further, the providers offer consumers services that allow them to create failure-tolerant systems. The SLA is the best place to record CSP uptime commitments.

Netflix uses AWS to deliver content to its customers. In porting its service to AWS from its own data center, Netflix[15] "… found that the best defense against major unexpected failures is to fail often." Netflix developed a tool, Chaos Monkey, that causes failures within their services forcing them to design for failure and thereby making their services more resilient. Combining the resilience provided by a cloud service provider's infrastructure with intelligent use of services offered to consumers can result in very resilient cloud-based applications.

For the government client organization, cloud computing can both simplify and complicate disaster recovery planning. Because most major cloud providers operate several geographically-dispersed data centers, a single natural disaster is unlikely to affect all centers. For example, Amazon EC2 describes its geographic resiliency, "By launching instances in separate Availability Zones, you can protect your applications from failure of a single location. Since mobility of execution and replication of data are core capabilities underlying the resiliency of cloud services, cloud applications remain available [15]." Some level of disaster recovery is inherent in a well-designed, large-scale, cloud computing infrastructure.

That said, circumstances might force a cloud provider to discontinue operations. While not technically a disaster, the impact on a consumer is similar – applications and data need to be reconstituted. To protect themselves, consumers are urged to limit CSP vendor lock-in. Currently, most cloud service offerings are unique to each provider and are not be easily portable. An application built for the Google App Engine platform will not run on Microsoft's Azure Cloud Services platform. Hence, clients may need to develop alternative hosting strategies for applications deployed to the cloud. Software tools such as Terraform, Chef, Puppet, and Ansible provide cross-CSP orchestration capabilities that can help with cross-platform deployments.

For a private cloud, technologies such as virtualization can be employed to help with disaster recovery. Given that virtualized images frequently can be deployed independent of the physical hardware, virtualization provides an inherent continuity of operations capability (i.e., virtualized applications can be easily moved from one data center to another). Of course, this requires data centers in multiple geographic locations and replication of virtual machine images among locations.

Finally, when using cloud services, the possibility that the CSP may go out of business should be considered as a potential risk. To address this risk, an exit plan should be developed prior to contracting with a CSP. Such plans should address continuity of operations and retrieval of data

[15] http://techblog.netflix.com/2012/07/chaos-monkey-released-into-wild.html

stored with the CSP. Processes to continue operations without access to community or public clouds can be built and contract terms and conditions can specify the right to data and the processes and procedures to retrieve it.

3.3.1 Because Resources are Finite

To provide the "illusion of infinite resources," cloud providers must have adequate physical resources. A basic hypothesis is that the peak resource demands of different clients will not occur simultaneously. For a public cloud provider with thousands of clients worldwide, this hypothesis is likely true. However, when constructing community and private clouds, it will be necessary to consider the validity of this assumption. A community cloud built and operated for the Federal Government or Department of Defense could be subject to simultaneous peak usage by all clients, should a significant national emergency occur. Clients of community and private clouds with highly focused availability concerns must consider and plan for this possibility.

3.4 Key Considerations

The key considerations identified in this section for protecting computing and communications infrastructure in cloud deployments are:

- Security controls can be classified as Common (aka "Inheritable"), Hybrid (your CSP gives you the capabilities but you have some operational responsibility), and System-Specific (you get to bring it as part of the system you deploy on top of the CSP' services.) Use of the FedRAMP Customer Responsibility Matrix (CRM) provides a decoder ring.
- Cloud infrastructures typically provide 24/7/365 availability and services to support application resilience.
- Many large-scale cloud providers operate multiple, geographically dispersed, data centers. While this provides resiliency and may improve system availability, it can make tracking the physical location of your data more difficult at retrieval time
- Unique cloud service offerings that are not easily portable can make recovery from provider failure challenging.
- Cloud service providers, through their homogeneous environments and economies of scale, may be able to provide better infrastructure security than many government organizations currently achieve.
- FedRAMP helps consumers understand provider security practices and provides assessments of cloud service security controls that can be leveraged in a department or agency's security assessment and authorization process.
- Service providers offer dedicated network connections from a consumer's location to their virtual resources, eliminating the need to traverse the public Internet.
- NCPP provides important monitoring capabilities and should be planned into the architecture via use of a Trusted Internet Connection. Architectural guidance and examples are available in the TIC 2.0 reference architecture and from some cloud service providers.

4 Operating Defensively in the Cloud

Protecting data moved to the cloud and using the CSP's protection services, such as IdAM, as part of your organization's broader cyber defense solution is the first part of securely using cloud services. Operating within this new environment of outsourced capabilities and shared

responsibility is a new dimension not typically considered in a consumer's private data center.. Becoming effective will take due diligence and practice.

4.1 Understand the Threats

Effectively defending systems requires an understanding of the threats to which the system is exposed. With that knowledge, defenders can take meaningful action at appropriate times. As an example, consider defending property against the impact of a hurricane. Weather forecasts provide good information about the magnitude and timing of a hurricane, allowing people to plan and act. To defend cloud-deployed services, an organization needs to know what threats are being experienced by cloud providers, how those threats relate to their service provider, and what defensive actions are available. The Cloud Security Alliance (CSA) has prepared a report on the top twelve threats to cloud computing in 2016. This report can be used to understand risks associated with cloud computing and take targeted defensive actions.[16]

Additionally, as public clouds become increasingly popular, the migration and consolidation effect that leads to commodity pricing of IT systems and services also creates the high value target. Public clouds become a point at which an attacker can more likely profit. Since public clouds connect to the Internet, they are exposed to Internet-based threats. Beware of the Technology Stack

The hardware and software stack—whether it is commercial off-the-shelf, government off-the-shelf, or proprietary—has an impact on the soundness of the provider's security practices and how readily the government can understand them. For example, both Google and AWS use proprietary hardware and software to implement their clouds [13, 31]. The proprietary cloud infrastructure may be as secure as, or more secure than, the cloud infrastructure constructed of commodity hardware and commercial software. However, there is no standard for comparison. If a cloud vendor is using a proprietary infrastructure, it may be difficult for the government to assess the platform's vulnerabilities, and determine security best practices. Vulnerability scanning is not typically part of the FedRAMP Assessment Package. This makes it difficult to verify CSP vulnerabilities. As a potential mitigation and best practice, the government client should understand the provider's disclosure policy regarding known vulnerabilities, administrative practices, security events, etc. They also should have relevant reporting contractually specified. (Refer to Cloud SLA Considerations for the Government Consumer by Buck and Hanf for more information on contractually specifying this information [14].)

Use of virtual resources simplifies many aspects of software maintenance and configuration management. CSPs create these virtual resources by adding a virtualization layer, known as a hypervisor, to the software stack. Like any software layer, the virtualization layer must be securely configured and maintained. Vulnerabilities and configuration errors in the hypervisor create greater risk than vulnerabilities and configuration errors on individual physical servers because all virtual resources using the virtualization layer's services are affected. Cloud consumers depend on the CSP to configure the hypervisor and perform maintenance to address vulnerabilities.

Vulnerabilities have been uncovered that directly relate to the resource sharing capabilities inherent in virtualization; as an example, consider VMware Transparent Page Sharing (TPS).[17] Unfortunately, the CSP and its consumers can receive great performance benefit from optimized

[16] https://downloads.cloudsecurityalliance.org/assets/research/top-threats/Treacherous-12_Cloud-Computing_Top-Threats.pdf
[17] http://vsphere-land.com/news/why-the-vmware-vsphere-tps-vulnerability-is-a-big-deal.html

resource sharing. In many ways, CSPs may be forced to make a risk-reward decision in this regard that may not be understood or published by the CSP for consumer review. Accordingly, it is important to do the research necessary to understand the CSP's configuration of the virtualization layer to be a party to the risk-reward decision as it pertains to your organization.

The hypervisor itself can also be used in an attack. If the guest operating system on a virtual machine is compromised, that virtual machine can then be used to attack other virtual machines that share the same underlying physical resources. One such technique, known as a side-channel attack, uses a compromised virtual machine to observe the behavior and resource usage of another virtual machine. Through these observations, sensitive information is extracted by the malicious virtual machine from the observed virtual machine. [34]

4.1.1 Recognize the Malicious Insider

Clouds, whether public, community, or private, create an opportunity for a malicious insider with potentially broad access to resources. All three cloud deployment models create a new class of highly privileged insiders—the cloud infrastructure administrators. Operating systems have long had privileged users such as the UNIX root user or the Microsoft Windows administrator. The risk associated with these users often has been managed using a variety of techniques (e.g., limiting the number of platforms on which a person can have privileged access). The cloud approach to providing computing resources may create users with broad privileged access to the entire underlying cloud infrastructure. Given this risk, mitigating controls and access restrictions must be maintained—an unchecked, malicious cloud infrastructure administrator has the potential to inflict significant damage. For public and community clouds, it is important to understand how the vendor reduces the risk posed by cloud administrators. Organizations operating private clouds need to consider what operational and monitoring controls can be used to reduce this risk.

Public and community IaaS clouds significantly increase the number of people who are insiders or "near insiders." Multiple organizations will have virtual machines running on the same physical machine. Administrators of these "neighbor" virtual machines will have privileged access to those virtual machines—an excellent starting point for launching an attack.

Using Amazon's EC2 IaaS offering, [16] demonstrated the ability to map the cloud infrastructure and locate specific target virtual machines. Having located the target, the researchers were able to reliably place a virtual machine that they controlled on the same physical server. This capability enables a variety of virtual-machine-escape or "side channel" attacks to compromise the target. Hence, in multi-tenant IaaS, cloud neighbors are like malicious insiders.

One approach to mitigating this risk is to use a community cloud service with a restricted tenant population. Some CSPs now offer government-only community clouds that provide the similar services to their public clouds but have only government customers. For example, AWS GovCloud and Microsoft Azure Government are community versions of their public cloud services that are restricted to government customers.

Understanding cloud specific threats such as these, makes it easier to develop mitigations. For example, NIST has published a set of Virtual Machine security references (Ref. NIST SP 800-125, -125A, and -125B) to give guidance to cloud consumers and providers engaged in configuring and operationally monitoring for security threat in the virtualization layer of the technology stack. These guides should be referenced when building and deploying virtual private cloud (VPC) systems.

4.2 Monitor Continuously

The challenge of monitoring and defending cloud-based systems depends on the service model and may increase due to shared control of the IT stack. Monitoring and defending systems consists of detecting and responding to inappropriate or unauthorized use of information or computing resources. Much like Microsoft Windows, which has been the dominant desktop operating system and target of choice for malware, large public clouds and community clouds also are high-value targets. Penetrating the substrate of a public or community cloud can provide a foothold from which to attack the applications of all the organizations running on the cloud.

Audit trails from network devices, operating systems, and applications are the first source of information used to monitor systems and detect malicious activity. Some or all of these sources may not be available to a cloud client. With SaaS, all audit trails are collected by the cloud provider. With PaaS, application audit trails may be captured by the client, but operating system and network audit trails are captured by the provider. With IaaS, the government organization may capture audit trails from the virtual network, virtual operating systems, and applications. The provider collects the audit trails for the physical network and the virtualization layer.

While providers cannot make all of the actual physical resource monitoring data available to clients, some providers do make filtered data available. For example, AWS CloudTrail, provides client organizations a log of all AWS service API calls. Coupled with a client's own monitoring of its virtual resources, this can provide a fairly complete picture of the behavior of the client's cloud-based services.

4.3 Maintain Secure Configurations

Protecting software infrastructure in the cloud is an essential activity for maintaining an appropriate security posture. For cloud providers and traditional IT alike, it involves activities such as securely configuring operating systems and network devices, implementing software patches in a timely manner, and tracking the discovery of new vulnerabilities.

The good news in terms of basic infrastructure security such as configuration and patching is that cloud providers may do a better job than what most client organizations currently accomplish. The European Network and Information Security Agency (ENISA) observes, "… security measures are cheaper when implemented on a larger scale. Therefore, the same amount of investment in security buys better protection [10]." Large cloud providers will benefit from these economies of scale.

Cloud providers have an additional benefit—their systems are likely to be homogeneous [11], which is fundamental to delivering commodity resources on demand. Hence, the cloud provider can configure every server identically. Software updates can be deployed rapidly across the provider's infrastructure. As a contrasting example, one large Federal agency has observed that each of its servers is unique. Every server has at least one deviation from defined configuration standards. This heterogeneity adds to the complexity of maintaining infrastructure security.

Homogeneity also has a potential down side. Homogeneity ensures the entire infrastructure has the same vulnerabilities. An attack that exploits an infrastructure vulnerability will affect all systems in a homogeneous cloud. The characteristic that makes routine maintenance easier may increase the impact of a targeted attack.

Although it may be easier for CSPs to maintain infrastructure security, government clients should ensure that they understand the CSP's standards for configuring and maintaining the infrastructure used to deliver cloud services. While some security information is proprietary and sensitive, many CSPs share more information in response to customer needs. For example, Google publishes a white paper providing general information about its security operations and procedures [12]. AWS publishes a series of security operations compliance reports including the ISO/IEC 27001:2013 and SOC 3. To maintain authorization, the FedRAMP requires a monthly reporting by CSPs holding Provisional Authorization to Operate (P-ATO) [21]. The monthly report includes vulnerability scan summaries and notes on configurations that may have changed the security posture of the CSP.

4.4 Respond to Incidents

The government client's incident response team will need to learn the response capabilities offered by the cloud provider, ensure appropriate security SLAs are in place, and develop new response procedures that couple the cloud provider information with its own data. Given the challenge with obtaining provider infrastructure information, a government client's incident response team may need to rethink how it detects some types of malicious activity. For example, an incident response team that provides proactive services such as vulnerability scanning may not be allowed to perform these functions on systems and applications deployed in the cloud. A cloud provider's terms of use may prohibit these activities, as it would be difficult to distinguish legitimate client scanning actions from malicious activities. Standard incident response actions may not be possible in the cloud. For example, a government client's incident response team that proactively deletes known malicious e-mail from users' inboxes may not have this ability in a cloud-based SaaS e-mail system. Given these challenges, it is essential that the appropriate contractual relationship with SLAs be established. However, maintaining close ties with the FedRAMP Program Management Office (PMO) can help government clients to stay abreast of their CSP's security posture and relevant incident response activities.

Additionally, since incident response can become a shared operation between consumer and provider, government clients are turning to Operational Level Agreements (OLAs). Acting in addition to contract terms and conditions and SLAs, OLAs are intended to define operational roles and responsibilities of IT entities interacting procedurally toward common operational objectives. They should specify activities, responsible and support organizations, specific standards and guidelines, and human points of contact. OLA's can go beyond the basic incident reporting specified by the FedRAMP. They can be used to further specify reporting channels and procedure event triggers for cyber defense.

4.5 Key Considerations:

The key considerations identified in this section for monitoring and defending systems in cloud deployments are:

- Large public clouds are high-value targets.
- Incident response teams must develop procedures (with contractual backing) for working with a cloud provider.
- Cloud virtualization technology may create a new class of highly privileged users with broad access to the cloud infrastructure.
- Cloud neighbors pose a similar threat to malicious insiders.

- Security operations centers must understand the technologies used to virtualize processing, storage, and communication and develop effective ways to monitor these technologies.

5 Conclusions

Public and community models of cloud computing cede direct control of computing resources to cloud service providers and extend the enterprise perimeter to include the providers' resources. Therefore, as with any IT outsourcing to a third party, it is essential to have a clear understanding of a provider's security obligations when moving capabilities to the cloud. These obligations, along with reporting and SLAs, should be codified in a contractually binding arrangement. The key to secure use of cloud computing is a clear, shared understanding of the division of security responsibilities between the provider and client, and the ability to verify that both are meeting their responsibilities. The contract between an agency and a CSP is the key to understanding and verifying security responsibilities. FedRAMP helps government cloud consumers do this. The FedRAMP Security Controls [20] matrix defines the security responsibilities of a CSP. Assessment of a CSP's implementation of the security controls, by a third-party assessor, provides the initial verification that the CSP is meeting its security responsibilities. FedRAMP's continuous monitoring strategy [21] helps ensure the CSP continues to meet its security responsibilities over time. FedRAMP contract templates assist agencies to incorporate appropriate security clauses in their cloud service contracts.

Multiple options for cloud usage exist in the marketplace with a variety of shared control and security characteristics. For systems that are moderate risk and provide information for public consumption, a public cloud may be a viable option because the cloud platform can meet system requirements and provide adequate security. In these cases, the scale and homogeneity of resources of the public cloud also may improve infrastructure security posture over its current instantiation. A community cloud may be an option for capabilities that cannot reside in a public cloud. The community cloud can provide some of the system and financial characteristics of a public cloud, while providing enhanced security characteristics. Some CSPs now offer government community clouds that provide the similar services to their public clouds but have only government customers.

In summary, there are challenges to public and community cloud deployment but with careful planning and appropriate mitigations, significant benefits are possible.

5.1 Bibliography

[1] Armbrust, M., et al., "Above the Clouds: A Berkeley View of Cloud Computing," Technical Report Number UCB/EECS-2009-28, University of California at Berkeley, Electrical Engineering and Computer Science, February 2009.

[2] Lockheed Martin Cyber Security Alliance, "Awareness, Trust, and Security to Shape Government Cloud Adoption," http://www.ca.com/Files/IndustryResearch/lm-cyber-security_gov-cloud-adopt_233481.pdf, April 2010.

[3] *National Institutes of Standards and Technology (NIST), Special Publication 800-53 Revision 4 Security and Privacy Controls for Federal Information Systems and Organizations, April 2013.*

[4] Mell, P. and T. Grance, *NIST Special Publication 800-145, The NIST Definition of Cloud Computing, September 2011.*

[5] *Wiggs, J., "Crypto Services and Data Security in Windows Azure," http://msdn.microsoft.com/en-us/magazine/ee291586.aspx, MSDN Magazine, January 2010.*

[6] Mather, T., S. Kumaraswamy, S. Latif, *Cloud Security and Privacy: An Enterprise Perspective on Risks and Compliance (Theory in Practice)*, O'Reilley, September 2009.

[7] Gentry, C., "Fully Homomorphic Encryption Using Ideal Lattices," In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing*, Association of Computing Machinery, June 2009.

[8] Schneier, B., "Homomorphic encryption breakthrough," Schneier on Security, http://www.schneier.com/blog/archives/2009/07/homomorphic_enc.html, July 2009.

[9] Cellan-Jones, R., "The Sidekick Cloud Disaster," BBC News, http://www.bbc.co.uk/blogs/technology/2009/10/the_sidekick_cloud_disaster.html, October 13, 2009.

[10] European Network and Information Security Agency (ENISA), *Cloud Computing: Benefits, Risks, and Recommendations for Information Security*, November 2009.

[11] *Mell, P. and T. Grance, "Effectively and Securely Using the Cloud Computing Paradigm," http://csrc.nist.gov/groups/SNS/cloud-computing/cloud-computing-v26.ppt, October 2009.*

[12] Google, "Google Security Whitepaper Google," https://cloud.google.com/security/whitepaper .

[13] *Shankland, S., "Google Unlocks Once-Secret Server," cnet news, http://news.cnet.com/8301-1001_3-10209580-92.html, April 2009.*

[14] Buck, K., D. Hanf, *Cloud SLA Considerations for the Government Consumer*, 2010.

[15] Amazon.com, "Amazon Elastic Compute Cloud (Amazon EC2)," http://aws.amazon.com/ec2.

[16] Ristepart, T., E. Tromer, H. Sacham, S. Savage, "Hey You Get Off of My Cloud: Exploring Information Leakage in Third-Party Compute Clouds," 2009.

[17] Lemos, R., "Harnessing the Cloud for Hacking," Technology Review, http://www.technologyreview.com/web/24127/?a=f, December 10, 2009.

[18] CIO.gov, http://www.cio.gov/pages.cfm/page/Federal-Risk-and-Authorization-Management-Program-FedRAMP.

[19] Kash, W., "FedRAMP: The Dawn of Approve-Once, Use-Often," Government Computing News, April 30, 2010.

[20] FedRAMP Security Controls Baseline, Version 2.0, https://www.fedramp.gov/files/2015/03/FedRAMP-Rev-4-Baseline-Workbook-FINAL062014.xlsx , June 6, 2014

[21] FedRAMP Continuous Monitoring Strategy and Guide, Version 2.0, https://www.fedramp.gov/files/2015/03/FedRAMP-Continuous-Monitoring-Strategy-Guide-v2.0-3.docx , June6, 2014.
[22] National Institute of Standards and Technology, Special Publication 800-37, *Guide to applying the Risk Management Framework to Federal Information Systems*, http://csrc.nist.gov/publications/nistpubs/800-37-rev1/sp800-37-rev1-final.pdf ,Revision 1, February 2010.
[23] Amazon Web Services, "Amazon EBS Encryption," http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSEncryption.html .
[24] Amazon Web Services, "Protecting Data Using Encryption," http://docs.aws.amazon.com/AmazonS3/latest/dev/UsingEncryption.html .
[25] Amazon Web Services, "What is the AWS Key Management Service ?," http://docs.aws.amazon.com/kms/latest/developerguide/overview.html .
[26] Amazon Web Services, "CloudHSM," https://aws.amazon.com/cloudhsm/ .
[27] Blake, Andrew, *The Washington Times*, "Lightning vs cloud: Google blames data loss at Belgium facility on electrical storm," http://www.washingtontimes.com/news/2015/aug/20/google-blames-data-loss-belgium-facility-electrica/ , August 20, 2015.
[28] The Netflix Tech Blog, "Chaos Monkey Released Into The Wild," http://techblog.netflix.com/2012/07/chaos-monkey-released-into-wild.html , July 30, 2012.
[29] Walther, Henrik, "Exchange Online Identity Models and Authentication Demystified (Part 1)," http://www.msexchange.org/articles-tutorials/office-365/exchange-online/exchange-online-identity-models-and-authentication-demystified-part1.html , November 2015.
[30] McGuire, Cheryl, "ExpressRoute technical overview," https://azure.microsoft.com/en-us/documentation/articles/expressroute-introduction/ , February 2016.
[32] Miller, Rich, "Inside Amazon's Cloud Computing Infrastructure," http://datacenterfrontier.com/inside-amazon-cloud-computing-infrastructure/ , September 2015.
[32] *Trusted Internet Connections (TIC) Reference Architecture Document Version 2.0*, https://www.fedramp.gov/files/2015/04/TIC_Ref_Arch_v2-0_2013.pdf , October 2013.
[33] AWS, "Guidance for Trusted Internet Connection Readiness on AWS," https://d0.awsstatic.com/whitepapers/compliance/Guidance_for_Trusted_Internet_Connection_TIC_Readiness_on_AWS.pdf , February 2016.
[34] Gorka et. al., "Fine-grain Cross-VM Attacks on Xen and VMware are Possible," Worchester Polytechnic Institute, Worchester, MA, 2014.
[35] AWS, "AWS Cloud Trail," https://aws.amazon.com/cloudtrail/ .
[36] FedRAMP Standard Contract Language https://www.fedramp.gov/files/2015/03/FedRAMP_Standard_Contractual_Clauses_062712_0.pdf .
[37] FedRAMP Control Specific Contract Clauses, Version 2.0, https://www.fedramp.gov/files/2015/03/FedRAMP-Control-Specific-Contract-Clauses-v2.1.docx , June 6, 2014.
[38] *Creating Effective Cloud Computing Contracts for the Federal Government*, https://www.fedramp.gov/files/2015/03/Cloud-Best-Practices.pdf , February 24, 2012.
[39] Pizzette and Raines, "Products to Build a Private Cloud," Systems Engineering at MITRE, Cloud Computing Series, June 2010; Release #10-2731.