Approved for Public Release; Distribution Unlimited
Case # 10-3208

Systems Engineering at MITRE

CLOUD COMPUTING SERIES

Information Security in the Clouds

Donald Faatz
Lawrence Pizette

MITRE

Executive Summary

Deploying data and applications to a cloud computing environment, whether private, community, or public, changes an organization's information technology (IT) security posture. Private cloud environments utilize new software layers, such as virtualization technologies, within the IT infrastructure. While community and public offerings may employ similar technologies, 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 provider, 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.

While vendors have many platform-specific security controls unique to their own offerings, encryption provides a client-controlled mechanism to protect data moved outside the organization's security perimeter to a public or community cloud. Because encryption replaces physical protection, organizations should verify that either the cloud provider's encryption capabilities meet their data protection needs or additional encryption capabilities can be provided. In addition, as encrypted data cannot be directly processed by most applications, legacy capabilities moved to the cloud may require changes to enable the application to function properly. To protect the data, Federal IT leaders should ensure they have out-of-cloud backups or backups provided by multiple cloud vendors, ensuring no single point of failure with a given cloud provider [1].

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 cloud provider's capabilities are not adequate, 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 cost and timeliness.

Despite the advantages that the cloud provider 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 not have the ability to deploy sensors on the cloud provider's system or collect the data they currently use. Government IT teams will need to partner with cloud providers and adjust their detection and response procedures to include this relationship.

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 needs to be secured, community and public clouds introduce risks due to the lack of 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.

Table of Contents

THE BIG PICTURE: 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.

Information Security in the Clouds

Donald Faatz
Lawrence Pizette

1.0 Introduction

***Control of Environments*—**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 benefits, this new model of computing resource delivery presents Federal IT leaders and security architects with new risks that must be understood. A better understanding of risks associated with cloud computing can help in identifying appropriate ways to use this new IT approach.

Private cloud computing introduces new technologies, such as virtualization and self-service provisioning, that alter the software IT stack. These new technologies must be deployed, configured, and operated in a secure manner. While community and public clouds may share these new technologies as part of their infrastructure, they introduce additional security challenges. Two notable characteristics of public and community cloud computing contribute to a potential increase in risk:

- Loss of direct control creates shared responsibility between provider and client.
- Loss of enterprise security perimeter increases exposure of information.

Public and community models of cloud computing cede direct control of computing resources to a service provider in exchange for reduced costs or additional capabilities. Some responsibility for information security also is transferred to the service provider. However, the information owner retains ultimate responsibility and accountability for appropriately protecting the information. The National Institute 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 break the enterprise's computing perimeter, extending it into the cloud of shared virtual resources and making it difficult to define boundaries. Physical boundaries are replaced by virtual boundaries, eliminating the utility of physical separation as a security control.

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. It does not, however, offer specific solutions as private cloud technology is evolving, and public and community cloud computing is an immature market. Services and architectures are mostly unique to each provider. A solution that works with one provider may be unusable with others.

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

***What is Cloud Computing?*—**NIST defines cloud computing as [4]:

"A model 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 and released with minimal management effort or service provider interaction."

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

- SaaS provides the client 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 client an application-hosting environment for client 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 clients virtual computing infrastructure and access to persistent data storage. Clients 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.

As shown in Figure 1, the service model affects where the line is drawn in transferring security responsibilities to the cloud service provider. With SaaS, responsibility for most or all of the security controls in the application and the infrastructure supporting it are transferred to the provider. With PaaS, the provider assumes responsibility for the underlying virtualization infrastructure, but responsibility for application security controls remains with the client Government organization. However, Government clients may be able to leverage application program interfaces (APIs) made available by the provider for application security. IaaS is similar to PaaS, but without the benefit of provided APIs.

Deployment of the three service models is possible with any of four different deployment models: a public cloud, a community cloud, a private cloud, or a 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 changes and how the enterprise perimeter is affected. 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. In moving from a traditional model of organization-owned and operated resources to a cloud model, public clouds represent the largest change in threat exposure. Private clouds may represent little to no change. 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. Community clouds with large numbers of diverse members will closely resemble a public cloud.

Although private cloud technologies expose new security issues, they represent significantly 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 are not directly applicable to private clouds. For more information on private cloud security products, refer to *Products to Build a Private Cloud* by Pizette and Raines.

NIST also has identified five essential characteristics of cloud computing: on-demand service, broad network access, resource pooling, rapid elasticity, and measured service.

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

Figure 1: Ownership for Security Controls Varies by Service Model

- 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 breaking down the enterprise perimeter and increasing the exposure of data and applications.

Protecting Data in the Cloud

Use of Encryption and Keys—The fundamental purpose of information technology (IT) is to store and process data. Much of this data has some need for confidentiality, integrity, and/or availability protection. Often, aspects of this protection are provided by physical location of the resources storing or processing the data. In traditional IT, those resources are in the organization's computing center.

When moved to a public or community cloud, protections previously provided by physical location of resources no longer apply. The data is stored and processed on the cloud provider's hardware at the provider's data center. Further, several different tenants use storage and transmission resources concurrently. In this environment, encryption and digital signature replace physical location as a means to protect data confidentiality and integrity.

Many cloud providers offer some form of encryption for data stored or transmitted within their clouds. For example, the Microsoft Windows Azure PaaS offering makes Cryptographic Service Providers available through the .NET Framework APIs [5]. 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. For example, the Government client organization needs to understand how keys are managed and who has access to the keys. 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. 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.

Control over encryption varies by service delivery model.

- With SaaS, clients rely upon the cryptographic capabilities the provider has built into the application. Clients need to verify that either an SaaS application provides appropriate cryptographic capabilities or confidentiality protection is not needed for the data that the application will process.
- With PaaS, clients rely upon the cryptographic capabilities available in the provider's application hosting environment. While the PaaS hosting environment may include software development or scripting capabilities, most do not allow installation of third-party products such as encryption software. Therefore, 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 provider supplies a virtual machine running an operating system, the client can install and use third-party encryption software.

Even with encryption, data always will have a window of vulnerability because currently it cannot be processed while it is encrypted.[1] Only data and applications for which this exposure is acceptable are candidates for cloud deployment. The degree of exposure is dependent on the cloud deployment model and the relative sophistication of the cloud provider's security controls. Assuming equivalent security capabilities, public clouds represent the greatest exposure and private clouds the least. The degree of exposure in a community cloud depends on the make-up of the community.

[1] *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 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.

Having secured the data stored in the cloud, the next question to address is, what happens to persistent data when it is deleted? Even if the data is encrypted, it is preferable to have data cleared from persistent storage. For each type of persistent cloud storage used, 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 either does not clear the storage or the clearing mechanism is inadequate. With SaaS, the only options are those provided by the provider.

Personally Identifiable Information—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 whether any privacy-specific protection requirements 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]." 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. For example, the cloud vendor's terms of use might grant the vendor some rights to information stored or processed in their cloud. Hence, if PII is stored or processed, it is important to review the privacy protection responsibilities and how they are met in the cloud.

Data and applications stored and processed in a public or community cloud should be copied periodically outside the cloud or, at a minimum, to another cloud provider with a different IT stack. This will provide backup in the event of cloud provider failure. Most cloud-based storage includes transparent replication by the cloud service provider. 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]."

As an example of a well-publicized cloud failure in September of 2009, users of the T-Mobile Sidekick Smartphone discovered that their address books and other information stored 'in their phones' were missing [9]. In actuality, the Sidekick did not store information in the phone, but instead used a network-based storage provider. Because of a serious technical problem at the storage provider, data for 800,000 users was lost, initially believed irretrievably. As events unfolded, the users' data was unavailable for a few weeks, but eventually was recovered. While this information may not seem as critical as information in Government systems providing essential services to the public or national security, it helps to exemplify the need for a backup strategy. Hopefully, serious technical problems such as those encountered by Sidekick will be rare. However, given the number of cloud providers and the relative maturity of the cloud market, some provider business failures are inevitable.

Identity and Access Management—Moving data and applications to a cloud means the Identity and Access Management (IdAM) capability must expand to encompass cloud-based resources. 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.

Initially, the IdAM capabilities from many cloud providers were very basic–often requiring manual setup using a Web-based interface. However, cloud providers' capabilities are evolving and maturing rapidly. Some providers now offer the ability to federate cloud-based IdAM with enterprise capabilities using standards-based mechanisms such as Security Assertion Markup Language tokens. 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. 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.

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

- Understanding provider security practices and controls is essential for public and community cloud offerings.
- Encryption and digital signatures are the primary confidentiality and integrity protection for data stored or transmitted in a public or community cloud.
- Without appropriate protections, data may be vulnerable while being processed in a public or community cloud.
- Deleted data may remain in persistent storage when the storage is released back to the cloud vendor as a shared, multi-tenant resource.
- Existing internal applications may need analysis and enhancement to operate securely in a public or community cloud.
- Data replication provided by a cloud provider is not a substitute for backing up to another independent provider or out of the cloud.
- Privacy protection responsibilities should be reviewed if considering moving PII to the cloud.
- Cloud IdAM capabilities vary widely. Integration of cloud and enterprise IdAM mechanisms may be challenging.

Protecting Computer and Communications Infrastructure

Software Maintenance and Patching Vulnerabilities—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, ensuring software patches are up-to-date, 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. A potential area for future research would be to employ an instance of a completely different technology stack for the express purpose of validating the integrity of the initial homogeneous infrastructure.

Although it may be easier for cloud providers to maintain infrastructure security, Government clients should ensure that they understand the provider's standards for configuring and maintaining the infrastructure used to deliver cloud services.

While some security information is proprietary and sensitive, many providers are starting to share more information in response to customer needs. For example, Google recently published a white paper providing general information about its security operations and procedures [12].

***The Technology Stack*—**The hardware and software stack—and 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, Google and some other providers use proprietary hardware and software to implement their clouds [13]. 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. There are no commonly accepted secure configurations standards and no public source of vulnerability information for these proprietary infrastructures. 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].)

Similar to the community and public cloud providers described above, Government organizations implementing private cloud solutions may find it easier and faster to maintain secure configurations and timely patching. Unlike physical servers, virtual servers do not have to be configured or patched individually. Instead, the virtual machine images are configured and patched. Measuring compliance also can be simplified by checking the virtual machine images rather than running measurement agents on each virtual server.

***Disaster Recovery*—**In addition to maintaining the currency of software and expeditiously plugging vulnerabilities, cloud computing providers must be able to quickly recover from disaster events. 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. Currently, most cloud service offerings are unique to each provider and may not be easily portable. An application built for the Google Apps platform will not run on Microsoft's Azure platform. Hence, clients may need to develop alternative hosting strategies for applications deployed to the cloud. If dictated by system requirements for uptime and availability, organizations can develop processes to continue operations without access to community or public cloud-based applications.

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).

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

- 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.
- Assessing the security posture of providers is complicated if proprietary hardware or software is used.
- Many large-scale cloud providers operate multiple, geographically dispersed, data centers.
- Unique cloud service offerings that are not easily portable make recovery from provider failure challenging.

Monitoring and Defending Systems in the Cloud

Monitoring and Defending Infrastructure—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. Correlation of events across provider-hosted virtual machines may be difficult, and the ability to place intrusion detection sensors in the virtual machines may be similarly constrained.

To date, most cloud providers have focused on monitoring and defending the physical resources that they control. Unlike their clients, cloud providers have the technical ability to collect audit trail information and place intrusion detection sensors in the infrastructure where their clients cannot. Although they can do this, cloud providers may not be willing or able to share that data. In clouds that host multiple tenants, the provider would need to protect the privacy of all its customers, which complicates the ability to share information. As noted by Buck and Hanf, service-level agreements (SLAs) that specify the exact types of information that will be shared are essential.

Incident Response Team—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 difficulty of 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 email system. Given these challenges, it is essential that the appropriate contractual relationship with SLAs be established.

If the organization is creating a private cloud, there are new challenges that are different from many of the community and public cloud issues. The virtualization layer presents a new attack vector, and many components (e.g., switches, firewalls, intrusion detection devices) within the IT infrastructure may become virtualized. The organization's security operations staff must learn how to safely deploy and administer the virtualization software, and how to configure, monitor, and correlate the data from the new virtual devices.

While cloud computing may make some aspects of incident detection and responses more complex, it has the potential for simplifying some aspects of forensics. When a physical computer is compromised, a forensic analyst's first task is to copy the state of the computer quickly and accurately. Capturing and storing state quickly is a fundamental capability of many IaaS clouds. Instead of needing special-purpose hardware and tools to capture the contents of system memory and copy disks, the forensic analyst uses the inherent capabilities of the virtualization layer in the IaaS cloud. Leveraging this capability will require the incident response team to develop procedures for capturing and using this state information and, in the case of community and public clouds, develop and maintain a working relationship with the cloud provider.

***Malicious Insider*—**Clouds, whether public, community, or private, create an opportunity for a malicious insider. 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, neighbors are similar to malicious insiders.

***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 infrastructure simplifies forensic capture of system state.
- Cloud virtualization technology may create a new class of highly privileged users with broad access to the cloud infrastructure.
- Cloud neighbors pose a threat similar to malicious insiders.

Additional Considerations

This section presents a few issues that are not directly related to protecting data or infrastructure that Government clients should consider in addressing security for cloud deployments.

***Use of Clouds by Adversaries*—**Even organizations that do not use cloud computing may have their IT security posture affected by clouds. All of the advantages that cloud computing provides to legitimate users are potential advantages to cyber adversaries. Cloud computing makes large amounts of computing resource available to adversaries with little more than access to a credit card, which itself may be stolen. To the extent that adversary behavior in the cloud does not differ noticeably from other customers, cloud service providers may have little ability to detect illicit activity. Further, as long as these cyber adversaries pay their bills and do not disrupt other customers, cloud service providers may have little incentive to look for illicit behavior.

In December 2009, *Technology Review* reported the availability of a cloud-based tool for cracking WiFi protected access (WPA) passwords, known as WPA Cracker [17]. This tool reportedly used 400 virtual computers to decrypt passwords in 20 minutes that would have required 5 days on a single physical computer. While clouds are not the only potential resource for adversaries, availability, scale, and ease-of-use may make them a tool of choice.

***Evolving State-of-the-Practice*—**Since cloud computing is a rapidly evolving approach to delivering IT capabilities, cloud provider service offerings are likely to change frequently. In parallel, vulnerabilities in cloud offerings and attack vectors also are evolving rapidly. Hence, organizations deploying applications in a cloud should expect to devote resources to continuously monitoring and understanding these changes and their impact on application security.

***False Illusion of Infinite Resources*—**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.

FedRAMP—To reduce the challenges of cloud computing certification and accreditation(C&A), the Federal Risk and Authorization Management Program (FedRAMP) was recently created. While FedRAMP is currently under development (as of July 2010), its objective is to allow agencies "to save significant time and money by leveraging the FedRAMP authorizations [18]." The ability to leverage capabilities such as FedRAMP to ease adoption of cloud computing offerings is an emerging area in cloud computing that holds significant promise [19].

Key Considerations—The key additional considerations identified in this section for cloud deployments are:

- On-demand, pay-as-you-go, public clouds resources are attractive, useful tools for malicious entities.
- Preserving the security posture of cloud-deployed applications will require constant attention to the rapid change in cloud offerings that may introduce risks.
- Government clients of public and community clouds with high availability requirements must consider and plan for the possibility that a national emergency could cause peak usage of resources.

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, 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.

Multiple options for cloud usage are emerging in the marketplace with a variety of shared control and security characteristics. For systems that are low 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. Finally, for systems used for national security, a public or community cloud may not be a viable option, due to the loss of control and the extended enterprise security perimeter. For these applications, a private cloud may be an effective alternative to harness some of the benefits of cloud computing, while minimizing the risks and changes to security posture.

References

1 Armbrust, M., et al., February 2009, "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.

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

3 National Institutes of Standards and Technology (NIST), June 2009, *Special Publication 800-53 Revision 3 Recommended Security Controls for Federal Information Systems and Organizations.*

4 Mell, P. and T. Grance, October 2009, "The NIST Definition of Cloud Computing," Version 15 csrc.nist.gov/groups/SNS/cloud-computing/cloud-def-v15.doc.

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

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

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

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., October 13, 2009, "The Sidekick Cloud Disaster," *BBC News* http://www.bbc.co.uk/blogs/technology/2009/10/the_sidekick_cloud_disaster.html.

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

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

12 Google, 2010, "Security Whitepaper Google Apps Messaging and Collaboration Products."

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

14 Buck, K., D. Hanf, 2010, "Cloud SLA Considerations for the Government Consumer," *Cloud Computing,* The MITRE Corporation.

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

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

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

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

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

MITRE

www.mitre.org

©2010 The MITRE Corporation
All Rights Reserved
Distribution Unlimited
Case Number: 10-3208

MITRE

www.mitre.org