



Security Guidance
for
Critical Areas of Focus
in
Cloud Computing

Prepared by the
Cloud Security Alliance
April 2009

Table of Contents

Foreword.....	3
Acknowledgments.....	4
Introduction.....	5
Executive Summary and Key Guidance	6
Section I. Cloud Architecture	14
Domain 1: Cloud Computing Architectural Framework	15
Section II. Governing in the Cloud.....	25
Domain 2: Governance and Enterprise Risk Management.....	26
Domain 3: Legal	30
Domain 4: Electronic Discovery.....	41
Domain 5: Compliance and Audit	44
Domain 6: Information Lifecycle Management	48
Domain 7: Portability and Interoperability	51
Section III. Operating in the Cloud.....	54
Domain 8: Traditional Security, Business Continuity and Disaster Recovery	55
Domain 9: Data Center Operations.....	59
Domain 10: Incident Response, Notification and Remediation.....	62
Domain 11: Application Security	65
Domain 12: Encryption and Key Management	72
Domain 13: Identity and Access Management	74
Domain 14: Storage	77
Domain 15: Virtualization	79
Appendix A. Contact Information	83

Foreword

Welcome to the Cloud Security Alliance's initial report, "Security Guidance for Critical Areas of Focus in Cloud Computing". From our first organizing meeting in Silicon Valley in early December of 2008, we have moved rapidly to garner industry support and have reached out to a multitude of subject matter experts to develop this report. We look forward to your participation in reviewing this document and providing your feedback at public venues and in our working groups.

My role as a Chief Information Security Officer is dual-purposed as it pertains to this subject. I am responsible for security assurance as both a consumer and provider of cloud computing services. Depending upon which hat I am wearing, I will have stronger affinity with any particular guidance recommendation in this document. However, if you accept the proposition that cloud computing allows for the realization of economic efficiencies in computing, I strongly feel that the most cost effective way to secure the cloud is to do it right the first time. Implementing a high standard of security benefits both providers and consumers alike.

The very nature of how businesses use information technology is being transformed by the 'on-demand' cloud computing model. It is imperative that information security leaders are engaged at this early stage to help assure that the rapid adoption of cloud computing builds in information security best practices without impeding the business. I am proud to be a co-founder of this important initiative.

Best,

A handwritten signature in blue ink that reads "Dave Cullinane". The signature is written in a cursive style and is contained within a thin black rectangular border.

Dave Cullinane, CPP, CISSP
CISO & Vice President eBay MP Global Information Security

Acknowledgments

The Cloud Security Alliance and its initial work product have been made possible through the tireless contribution of many volunteers we would like to thank here.

Alliance Co-Founders

Nils Puhlmann, Qualys
Jim Reavis, Cloud Security Alliance

Editor

Jim Reavis, Cloud Security Alliance

Industry Advisors

Jerry Archer, Intuit
Alan Boehme, ING
Larry Brock, DuPont
Dave Cullinane, eBay

Paul Kurtz, Good Harbor Consulting
Izak Mutlu, Salesforce.com
Nils Puhlmann, Qualys
Lynne Terwoerds, Barclays

Primary Authors

Jeff Bardin, Treadstone 71
Jon Callas, PGP
Shawn Chaput, Privity
Pam Fusco, UAT
Francoise Gilbert, IT Law Group
Christofer Hoff, Rational Survivability
Dennis Hurst, HP
Subra Kumaraswamy, Sun
Liam Lynch, eBay

Scott Matsumoto, Cigital
Brian O'Higgins, Third Brigade
Jean Pawluk, Visa
George Reese, enStratus
Jeff Reich, FUDless
Jeffrey Ritter, Waters Edge Consulting
Jeff Spivey, RiskIQ
John Viega, McAfee

Contributing Reviewers

Phil Agcaoili, Dell
Todd Barbee, New Dominion Bank
Girish Bhat, SAVVIS
Glenn Brunette, Sun
Jake Brunetto, Intuit
Sean Catlett, Barclays
Anton Chuvakin, Qualys
Joshua Davis, Qualcomm
Dr Ken Fauth, CPP
Jeff Forristal, Zscaler
Robert Fly, Salesforce.com
Edward Haletky, AstroArch Consulting
Jim Hietala, The Open Group

Michael Johnson, Security GRC2
Shail Khiyara, Cloud Computing
Mark Leary, Northrop Grumman
Tim Mather, RSA Security
Dave Morrow, Secure Business Operations
Josh Pennell, IOActive
Ben Rothke, BT
Stephen Sengam, Fox/Newscorp
Ward Spangenberg, IOActive
Michael Sutton, Zscaler
Dave Tyson, eBay
Dov Yoran, MetroSITE Group
Josh Zachry, Rackspace

We would like to thank the CEOs of our founding charter companies for their encouragement and contributions to this endeavor:

Jay Chaudhry
Zscaler, Inc.

Philippe Courtot
Qualys, Inc.

Phil Dunkelberger
PGP Corporation

We would also like to acknowledge Peter M. Mell of NIST for his excellent work on Cloud Computing Definitions

Introduction

We are continuously bombarded with news of information technology's next big thing, a disruptive trend in computing with far reaching implications. Many of these trends are no more than a marketer's dream - hype sells technology and it becomes difficult to separate real change from an incremental upgrade. Cloud Computing is having its moment in the sun, as the concept of utilizing computing as an on-demand subscription creates operating and economic efficiencies. Some deride the cloud as nothing new and in many respects they are correct. Henry Ford's Model T was not a new invention, but the revolution that ensued cannot be denied. We believe Cloud Computing to be a very important trend that in many ways is beginning to fulfill the early promise of the Internet and will create unanticipated change in business with its ubiquitous adoption. Phase one of the Internet was connectivity, with Cloud Computing we are leveraging that connectivity to optimize the utility of computing.

While we do see Cloud Computing as being a major change coming to every business, as information security practitioners, we recognize that there are verities which must not change: good governance, managing risks and common sense. Cloud Computing is an unstoppable force and we encourage security practitioners to lead and help accelerate its secure adoption aided by common sense, rather than standing on the sidelines and letting the business move forward without us.

Some evangelists of cloud computing encourage us to focus on the model as a black box, the seamless presentation of your information on demand. Pay no attention to how it works: resources are dynamically allocated, loads are balanced in real time and data is archived automatically. Our message to the security practitioner is that in these early days of cloud computing, you must look under the hood of your cloud providers and you must do so using the broadest precepts of your profession in order to properly assure that the service engagements meet and exceed the security requirements of your organization.

The Cloud Security Alliance is a grassroots effort to facilitate the mission to create and apply best practices to secure cloud computing. Incorporated as a not-for-profit organization, our efforts will seek to provide a voice for security practitioners. However, recognizing that a secure cloud is a shared responsibility, we will be inclusive of all organizations and points of view to fulfill this mission.

What follows is our initial report, outlining areas of concern and guidance for organizations adopting cloud computing. The intention is to provide security practitioners with a comprehensive roadmap for being proactive in developing positive and secure relationships with cloud providers. Much of this guidance is also quite relevant to the cloud provider to improve the quality and security of their service offerings. As with any initial foray, there will certainly be guidance that we could improve upon. We will quite likely modify the number of domains and change the focus of some areas of concern. We seek your help to improve this guidance to make version 2.0 of this document an even better asset to the security practitioner and cloud provider. We will be kicking off numerous online activities and in-person regional events to share our findings and connect with experts to increase our knowledge base. Here is how you can get involved:

- Visit our website to find out how you can help: www.cloudsecurityalliance.org
- Join our LinkedIn group to collaborate with us: www.linkedin.com/groups?gid=1864210

Executive Summary of Key Guidance

Cloud security requires companies and stakeholders to navigate a portfolio of domains. This executive summary provides a quick guide to the initial guidance that the expert contributors have developed within each domain. Every attempt has been made to focus on areas of concern that are either unique to cloud computing, or are greatly exacerbated by the model. This executive summary is not a substitute for reading all of the domain documents and understanding the assumptions and reasoning behind the guidance.

Domain 1: Cloud Computing Architectural Framework

There are five Principal Characteristics of Cloud Computing:

1. Abstraction of Infrastructure
2. Resource Democratization
3. Services Oriented Architecture
4. Elasticity/Dynamism of Resources
5. Utility model of Consumption & Allocation

There are three Cloud Service Delivery Models:

1. Infrastructure as a Service (IaaS)
2. Platform as a Service (PaaS)
3. Software as a Service (SaaS)

The three Cloud Service Delivery Models, their interrelationship and logical boundaries can be expressed in the Cloud Reference Model included in this domain.

There are four Cloud Service Deployment and Consumption Modalities

1. Private
2. Public
3. Managed
4. Hybrid

It is critical to be aware of the tradeoffs between extensibility (openness) and security responsibility within the three Cloud Service Delivery Models

- SaaS (Software as a Service): least extensibility and greatest amount of security responsibility taken on by the cloud provider
- IaaS (Infrastructure as a Service): greatest extensibility and least amount of security responsibility taken on by the cloud provider
- PaaS (Platform as a Service): lies somewhere in the middle, with extensibility and security features which must be leveraged by the customer

Domain 2: Governance and Enterprise Risk Management

- √ A portion of the cost savings obtained by cloud computing services must be invested into the increased scrutiny of the security capabilities of the provider and ongoing detailed audits to ensure requirements are continuously met.

Security Guidance for Critical Areas of Focus in Cloud Computing

- √ The Domain 1 Principals of Cloud Computing which make it very flexible and affordable create a relationship dynamism, which must be mitigated by ongoing risk management.
- √ Providers should have regular third party risk assessments and these should be made available to customers.
- √ Require listings of all third party relationships of the cloud provider.
- √ Understand financial viability of cloud provider.
- √ Understand the cloud provider's key risk and performance indicators and how these can be monitored and measured from a customer perspective.
- √ Request a divulgence of all policies, procedures and processes comprising the cloud provider's Information Security Management System (ISMS).
- √ The onus is on the customer to perform extensive due diligence of a cloud provider for usage in mission critical business functions or for hosting regulated personally identifiable information. At this point in time, customers should consider Private and Hybrid Cloud models for these types of business needs, unless rigorous due diligence determines a Public Cloud is acceptable.
- √ Contracts are not your only governance tool but should encompass the broad due diligence required of a cloud provider.

Domain 3: Legal

- √ Contracts are the key legal enforcement mechanism and must be negotiable to reflect any organization's unique needs and the dynamic nature of cloud computing.
- √ Plan for both an expected and unexpected termination of the relationship in the contract negotiations and for an orderly return or secure disposal of your assets.
- √ Understand that there will be conflicts between the laws the cloud provider must comply with and those governing the cloud customer. Perform due diligence to find out what those conflicts are.
- √ Gain a clear expectation of the cloud provider's response to legal requests for information.
- √ Understand any secondary uses of data by the cloud provider and develop contract language to prohibit them if necessary.
- √ Identify potential for cross-border data transfers and develop contract language to prohibit them if necessary.
- √ Develop service level agreement monitoring language in the contracts.

Domain 4: Electronic Discovery

- √ Cloud providers have become custodians of primary data assets for which customers have legal responsibilities to preserve and make available in legal proceedings (electronic discovery), even if the customer is not in direct possession or control.
- √ Cloud Computing challenges the presumption that organizations have control over data for which they are legally responsible. Electronic discovery has become an essential function to which a cloud provider is indispensable; if neglected by customers and providers, the adverse legal risks are substantial.
- √ Customers and cloud providers must have a mutual understanding of each others' roles and responsibilities related to Electronic Discovery, including such activities as litigation hold, discovery searches, who provides expert testimony, etc.
- √ Cloud providers are advised to assure their information security systems are responsive to customer requirements to preserve data as authentic and reliable, including both primary information and metadata, log files and other related information.
- √ Going forward, the Records and Information Management (RIM) domain of knowledge must be adapted to Cloud Computing.

Domain 5: Compliance and Audit

- √ Classify data and systems to understand compliance requirements
- √ Understand data locations, in particular the copies of data that are made and how they are controlled.
- √ Maintain a right to audit on demand as your regulatory mandates and business needs may change rapidly.
- √ Perform external risk assessments, including a Privacy Impact Assessment.
- √ While SAS 70 Type II audits and ISO 27001 certifications can indicate widely varying levels of security competency, in the aggregate they are better than no certifications whatsoever.
- √ It is critical to examine the scope of SAS 70 Type II audits and ISO 27001 certifications. Going forward, we advocate greater uniformity in comprehensive certification scoping. This will lead to increased security assurance for the customer and a decrease in ad hoc audits, an expensive drag on cloud provider productivity.

Domain 6: Information Lifecycle Management

- √ Understand the logical segregation of information and protective controls implemented.

Security Guidance for Critical Areas of Focus in Cloud Computing

- √ Understand the privacy restrictions inherent in data entrusted to your company, this information potentially will not be allowed to be held by a cloud provider without very specific partner designations.
- √ Understand cloud provider policies and processes for data retention and destruction and how they compare with internal organizational policy. Be aware that data retention assurance may be easier for the cloud provider to demonstrate, but data destruction may be very difficult.
- √ Negotiate penalties payable by the cloud provider for data breaches to ensure this is being taken seriously. If practical, customer should seek to recover all breach costs as part of their provider contract. If impractical, customer should explore other risk transference vehicles, such as insurance, to recover breach costs.
- √ Perform regular backup and recovery tests to assure that logical segregation and controls are effective.
- √ Assure that cloud provider personnel controls are in place to provide a logical segregation of duties.

Domain 7: Portability and Interoperability

- √ It is important to understand Domain 1 Architectural Framework and in particular the distinction between SaaS, PaaS and IaaS as a prerequisite to accurately assess portability and interoperability risks.
- √ For Software as a Service (SaaS), perform regular data extractions and backups to a format that is usable and not proprietary to the SaaS provider.
- √ For Infrastructure as a Service (IaaS), deploy applications in runtime in a way that is abstracted from the machine image. Backups should also be machine independent.
- √ For Platform as a Service (PaaS), careful application development techniques should be followed to minimize potential lock-in for the customer. The PaaS provider marketplace is evolving quickly and there are a variety of development environments available, which will vary greatly in their adherence to standards and compatibility with other PaaS providers. The onus is on the customer to have portability as a key design goal and an architecture that supports the necessary abstraction layers to make this goal achievable.
- √ Understand who the competitors are to your cloud providers and what their capabilities are to assist in migration.
- √ Advocate open standards, particularly as it relates to application development.

Domain 8: Traditional Security, Business Continuity and Disaster Recovery

- √ Centralization of data means the risk of insider threats from within the cloud provider is a significant concern.
- √ Cloud providers should adopt as a security baseline the most stringent requirements of any customer.
- √ Providers should have robust compartmentalization of job duties and limit knowledge of customers to that which is absolutely needed to perform job duties.
- √ Customers should perform onsite inspections of cloud provider facilities whenever possible.
- √ Customers should inspect cloud provider disaster recovery and business continuity plans.
- √ Customers should identify physical interdependencies in provider infrastructure.

Domain 9: Data Center Operations

- √ Understand how your cloud provider has implemented Domain 1's "Five Principal Characteristics of Cloud Computing" and how that technology architecture and infrastructure impacts their ability to meet service level agreements.
- √ While the technology architecture infrastructure of cloud providers will differ, they must all be able to demonstrate comprehensive compartmentalization of systems, networks, management, provisioning and personnel.
- √ If feasible, discover the cloud provider's other clients to assess the impact their business fluctuations may have on your customer experience with the cloud provider.
- √ Understand how resource democratization occurs within your cloud provider to best predict the likelihood of system availability and performance during your business fluctuations.
- √ For IaaS and PaaS, clearly understand the cloud provider's patch management policies and procedures and how this impacts the applications you have developed for these environments. This understanding should be reflected in contract language.
- √ As in Domain 8, review business continuity and disaster recovery plans from an IT perspective and how it relates to people and processes. Cloud provider's technology architecture may use new and unproven methods for failover. Customer's own business continuity plans should also address impacts and limitations of Cloud computing.
- √ Test cloud provider's customer service function regularly to determine their level of mastery in supporting the services.

Domain 10: Incident Response, Notification and Remediation

- √ Any data classified as private for the purpose of data breach regulations should always be encrypted to reduce the consequences of a breach incident. Customer should stipulate encryption requirements (algorithm, key length and key management at a minimum) contractually.
- √ Security Operation Centers (SOC) often assume a single governance model related to incident response, which is inappropriate to multi-tenant cloud providers
- √ Cloud providers need application layer logging frameworks to provide granular narrowing of incidents to a specific customer.
- √ Cloud providers should construct a registry of application owners by application interface (URL, SOA service, etc.)
- √ Application level firewalls, proxies and other application logging tools are key capabilities currently available to assist in responding to incidents in multi-tenant environments.
- √ Cloud providers and customers need defined processes for collaboration to determine appropriate remediation during incidents to minimize service downtime.

Domain 11: Application Security

- √ IaaS, PaaS and SaaS create differing trust boundaries for the software development lifecycle, which must be accounted for during the development, testing and production deployment of applications.
- √ For IaaS, a key success factor is the presence of trusted virtual machine images. The best alternative is the ability to provide your own virtual machine image conforming to internal policies.
- √ The best practices available to harden host systems within DMZs should be applied to virtual machines. Limiting services available to only those needed to support the application stack is appropriate.
- √ Securing inter-host communications must be the rule, there can be no assumption of a secure channel between hosts, whether existing in a common data center or even on the same hardware platform.
- √ Managing and protecting application “secret keys” is critical.
- √ Attention should be paid to consider how malicious actors will react to new cloud application architectures that obscure application components from their scrutiny. Hackers are likely to attack visible code, including but not limited to code running in the user context. They will also likely attack infrastructure and perform extensive black box testing.

Domain 12: Encryption and Key Management

- √ From a risk management perspective, unencrypted data existent in the cloud may be considered “lost” by the customer.
- √ Application providers who are not controlling backend systems should assure that data is encrypted when being stored on the backend.
- √ Use encryption to separate data holding from data usage.
- √ Segregate the key management from the cloud provider hosting the data, creating a chain of separation. This protects both the cloud provider and customer from conflict when being compelled to provide data due to a legal mandate and can potentially solve some problems discussed in Domain3: Legal and Domain 4: Electronic Discovery.
- √ When stipulating encryption in contract language, assure that the encryption is adhering to existing industry or government standards, as applicable.

Domain 13: Identity and Access Management

- √ The key critical success factor to managing identities at cloud providers is to have a robust federated identity management architecture and strategy internal to the organization.
- √ Insist upon standards enabling federation: primarily SAML, WS-Federation and Liberty ID-FF federation
- √ Validate that cloud provider either support strong authentication natively or via delegation and support robust password policies that meet and exceed cloud customer internal policies.
- √ Understand that the current state of granular application authorization on the part of cloud providers is non-existent or proprietary.
- √ Consider implementing Single Sign-on (SSO) for internal applications and leveraging this architecture for cloud applications.
- √ Using cloud-based “Identity as a Service” providers may be a useful tool for outsourcing some identity management capabilities and facilitating federated identity management with cloud providers. For example, they may be useful for abstracting and managing complexities such as differing versions of SAML, etc. Be aware that they become a critical new cloud provider for your organization and must be vetted with this broad guidance document.

Domain 14: Storage

- √ Understand the storage architecture and abstraction layers to verify that the storage subsystem does not span domain trust boundaries
- √ Ascertain if knowing storage geographical location is possible.
- √ Understand what controls are used during storage provisioning to partition multiple customers.
- √ Understand the cloud provider's data search capabilities.
- √ Understand cloud provider storage retirement processes. Related to Domain 6: Information Lifecycle Management, data destruction is extremely difficult in a multi-tenant environment and cloud provider should be utilizing strong storage encryption that renders data unreadable when storage is recycled, disposed of, or accessed by any means outside of authorized applications.
- √ Understand circumstances under which storage can be seized by a third party or government entity.
- √ Understand how encryption is managed on multi-tenant storage. Is there a single key for all customers, one key per customer, or multiple keys per customer?
- √ Can the cloud provider support long term archiving, will the data be available several years later and will the decryption and associated technologies still be useable?

Domain 15: Virtualization

- √ Virtualized operating systems should be augmented by third party security technology to provide layered security controls and reduce dependency on the platform provider alone.
- √ The simplicity of invoking new machine instances from a VM platform creates a risk that insecure machine images can be created. Secure by default configuration needs to be assured by following or exceeding available industry baselines.
- √ Virtualization also contains many security advantages such as creating isolated environments and better defined memory space, which can minimize application instability and simplify recovery.
- √ VM-specific security mechanisms embedded in hypervisor APIs need to be utilized to provide granular monitoring of traffic crossing VM backplanes, which will be opaque to traditional network security controls.
- √ Administrative access and control of virtualized operating systems is crucial and should include strong authentication integrated with enterprise identity management, as well as tamper proof logging and integrity monitoring tools.

Section I. Cloud Architecture

Domain 1: Cloud Computing Architectural Framework

Contributors: *Christofer Hoff*

Problem Statement

Cloud Computing (“Cloud”) is a catch-all term that describes the evolutionary development of many existing technologies and approaches to computing that at its most basic, separates application and information resources from the underlying infrastructure and mechanisms used to deliver them with the addition of elastic scale and the utility model of allocation. Cloud computing enhances collaboration, agility, scale, availability and provides the potential for cost reduction through optimized and efficient computing.

More specifically, Cloud describes the use of a collection of distributed services, applications, information and infrastructure comprised of pools of compute, network, information and storage resources. These components can be rapidly orchestrated, provisioned, implemented and decommissioned using an on-demand utility-like model of allocation and consumption. Cloud services are most often, but not always, utilized in conjunction with and enabled by virtualization technologies to provide dynamic integration, provisioning, orchestration, mobility and scale.

While the very definition of Cloud suggests the decoupling of resources from the physical affinity to and location of the infrastructure that delivers them, many descriptions of Cloud go to one extreme or another by either exaggerating or artificially limiting the many attributes of Cloud. This is often purposely done in an attempt to inflate or marginalize its scope. Some examples include the suggestions that for a service to be Cloud-based, that the Internet must be used as a transport, a web browser must be used as an access modality or that the resources are always shared in a multi-tenant environment outside of the “perimeter.” What is missing in these definitions is context.

From an architectural perspective given this abstracted evolution of technology, there is much confusion surrounding how Cloud is both similar and differs from existing models and how these similarities and differences might impact the organizational, operational and technological approaches to Cloud adoption as it relates to traditional network and information security practices. There are those who say Cloud is a novel sea-change and technical revolution while others suggest it is a natural evolution and coalescence of technology, economy and culture. The truth is somewhere in between.

There are many models available today which attempt to address Cloud from the perspective of academicians, architects, engineers, developers, managers and even consumers. We will focus on a model and methodology that is specifically tailored to the unique perspectives of IT network and security professionals.

The keys to understanding how Cloud architecture impacts security architecture are a common and concise lexicon coupled with a consistent taxonomy of offerings by which Cloud services and architecture can be deconstructed, mapped to a model of compensating security and operational controls, risk assessment and management frameworks and in turn, compliance standards.

Setting the Context: Cloud Computing Defined

Understanding how Cloud Computing architecture impacts security architecture requires an understanding of Cloud's principal characteristics, the manner in which cloud providers deliver and deploy services, how they are consumed and ultimately how they need to be safeguarded.

The scope of this area of focus is not to define the specific security benefits or challenges presented by Cloud Computing as these are covered in depth in the other 14 domains of concern:

- Governance and Enterprise Risk Management
- Legal
- Electronic Discovery
- Compliance and Audit
- Information Lifecycle Management
- Portability and Interoperability
- Traditional Security, Business Continuity and Disaster Recovery
- Data Center Operations
- Incident Response, Notification, Remediation
- Application Security
- Encryption and Key Management
- Identity and Access Management
- Storage
- Virtualization

We will discuss the various approaches and derivative offerings of Cloud and how they impact security from an architectural perspective using an in-process model developed as a community effort associated with the Cloud Security Alliance.

Principal Characteristics of Cloud Computing

Cloud services are based upon five principal characteristics that demonstrate their relation to, and differences from, traditional computing approaches:

1. Abstraction of Infrastructure

The compute, network and storage infrastructure resources are abstracted from the application and information resources as a function of service delivery. Where and by what physical resource that data is processed, transmitted and stored on becomes largely opaque from the perspective of an application or services' ability to deliver it. Infrastructure resources are generally pooled in order to deliver service regardless of the tenancy model employed – shared or dedicated. This abstraction is generally provided by means of high levels of virtualization at the chipset and operating system levels or enabled at the higher levels by heavily customized file systems, operating systems or communication protocols.

2. Resource Democratization

The abstraction of infrastructure yields the notion of resource democratization – whether infrastructure, applications, or information – and provides the capability for pooled resources to be made available and accessible to anyone or anything authorized to utilize them using standardized methods for doing so.

3. Services Oriented Architecture

As the abstraction of infrastructure from application and information yields well-defined

Security Guidance for Critical Areas of Focus in Cloud Computing

and loosely-coupled resource democratization, the notion of utilizing these components in whole or part, alone or with integration, provides a services oriented architecture where resources may be accessed and utilized in a standard way. In this model, the focus is on the delivery of service and not the management of infrastructure.

4. Elasticity/Dynamism

The on-demand model of Cloud provisioning coupled with high levels of automation, virtualization, and ubiquitous, reliable and high-speed connectivity provides for the capability to rapidly expand or contract resource allocation to service definition and requirements using a self-service model that scales to as-needed capacity. Since resources are pooled, better utilization and service levels can be achieved.

5. Utility Model of Consumption & Allocation

The abstracted, democratized, service-oriented and elastic nature of Cloud combined with tight automation, orchestration, provisioning and self-service then allows for dynamic allocation of resources based on any number of governing input parameters. Given the visibility at an atomic level, the consumption of resources can then be used to provide an “all-you-can-eat” but “pay-by-the-bite” metered utility-cost and usage model. This facilitates greater cost efficiencies and scale as well as manageable and predictive costs.

Cloud Service Delivery Models

Three archetypal models and the derivative combinations thereof generally describe cloud service delivery. The three individual models are often referred to as the “SPI Model,” where “SPI” refers to Software, Platform and Infrastructure (as a service) respectively and are defined thusly¹:

1. Software as a Service (SaaS)

The capability provided to the consumer is to use the provider’s applications running on a cloud infrastructure and accessible from various client devices through a thin client interface such as a Web browser (e.g., web-based email). The consumer does not manage or control the underlying cloud infrastructure, network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

2. Platform as a Service (PaaS)

The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created applications using programming languages and tools supported by the provider (e.g., java, python, .Net). The consumer does not manage or control the underlying cloud infrastructure, network, servers, operating systems, or storage, but the consumer has control over the deployed applications and possibly application hosting environment configurations.

3. Infrastructure as a Service (IaaS)

The capability provided to the consumer is to rent processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, deployed applications, and possibly select networking components (e.g., firewalls, load balancers).

¹ Credit: Peter M. Mell, NIST

Security Guidance for Critical Areas of Focus in Cloud Computing

Understanding the relationship and dependencies between these models is critical. IaaS is the foundation of all Cloud services with PaaS building upon IaaS, and SaaS – in turn – building upon PaaS. We will cover this in more detail later in the document.

The OpenCrowd Cloud Solutions Taxonomy² shown in Figure 1 provides an excellent reference that demonstrates the swelling ranks of solutions available today in each of the models above.

Narrowing the scope or specific capabilities and functionality within each of the *aaS offerings or employing the functional coupling of services and capabilities across them may yield derivative classifications. For example “Storage as a Service” is a specific sub-offering within the IaaS “family,” “Database as a Service” may be seen as a derivative of PaaS, etc.

Each of these models yields significant trade-offs in the areas of integrated features, openness (extensibility) and security. We will address these later in the document.

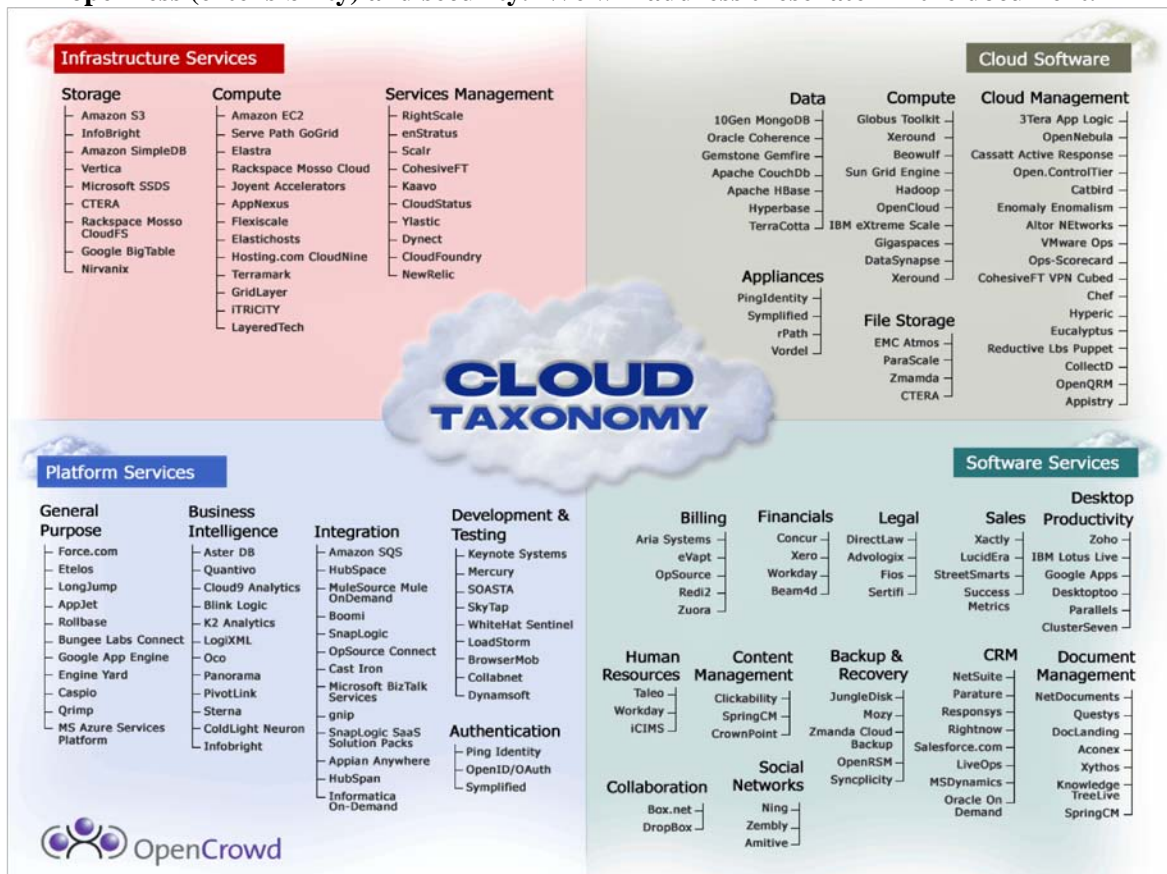


Figure 1 - The OpenCrowd Cloud Taxonomy

Cloud Service Deployment and Consumption Modalities

Regardless of the delivery model utilized (SaaS, PaaS, IaaS,) there are four primary ways in which Cloud services are deployed and are characterized:

² The OpenCrowd Taxonomy – <http://www.opencrowd.com/views/cloud.php/2>

Security Guidance for Critical Areas of Focus in Cloud Computing

1. **Private**

Private Clouds are provided by an organization or their designated service provider and offer a single-tenant (dedicated) operating environment with all the benefits and functionality of elasticity and the accountability/utility model of Cloud.

The physical infrastructure may be owned by and/or physically located in the organization's datacenters (on-premise) or that of a designated service provider (off-premise) with an extension of management and security control planes controlled by the organization or designated service provider respectively.

The consumers of the service are considered "trusted." Trusted consumers of service are those who are considered part of an organization's legal/contractual umbrella including employees, contractors, & business partners. Untrusted consumers are those that may be authorized to consume some/all services but are not logical extensions of the organization.

2. **Public**

Public Clouds are provided by a designated service provider and may offer either a single-tenant (dedicated) or multi-tenant (shared) operating environment with all the benefits and functionality of elasticity and the accountability/utility model of Cloud.

The physical infrastructure is generally owned by and managed by the designated service provider and located within the provider's datacenters (off-premise.)

Consumers of Public Cloud services are considered to be untrusted.

3. **Managed³**

Managed Clouds are provided by a designated service provider and may offer either a single-tenant (dedicated) or multi-tenant (shared) operating environment with all the benefits and functionality of elasticity and the accountability/utility model of Cloud.

The physical infrastructure is owned by and/or physically located in the organization's datacenters with an extension of management and security control planes controlled by the designated service provider.

Consumers of Managed Clouds may be trusted or untrusted.

4. **Hybrid**

Hybrid Clouds are a combination of public and private cloud offerings that allow for transitive information exchange and possibly application compatibility and portability across disparate Cloud service offerings and providers utilizing standard or proprietary methodologies regardless of ownership or location. This model provides for an extension of management and security control planes

Consumers of Hybrid Clouds may be trusted or untrusted.

³ This is a relatively new term I am introducing in order to clarify and distinguish between deployment modalities of service

Security Guidance for Critical Areas of Focus in Cloud Computing

The difficulty in using a single label to describe an entire service/offering is that it actually attempts to describe the following elements:

- Who manages it
- Who owns it
- Where it's located
- Who has access to it
- How it's accessed

The notion of Public, Private, Managed and Hybrid when describing Cloud services really denotes the attribution of management and the availability of service to specific consumers of the service.

It is important to note that the characterizations that describe *how* Cloud services are deployed are often used interchangeably with the notion of *where* they are provided; as such, you may often see public and private clouds referred to as “external” or “internal” clouds. This can be very confusing.

The manner in which Cloud services are offered and ultimately consumed is then often described relative to the location of the asset/resource/service owner's management or security “perimeter” which is usually defined by the presence of a “firewall.”

While it *is* important to understand where within the context of an enforceable security boundary an asset lives, the problem with interchanging or substituting these definitions is that the notion of a well-demarcated perimeter separating the “outside” from the “inside” is an anachronistic concept.

It is clear that the impact of the re-perimeterization and the erosion of trust boundaries we have seen in the enterprise is amplified and accelerated due to Cloud. This is thanks to ubiquitous connectivity provided to devices, the amorphous nature of information interchange, and the ineffectiveness of traditional static security controls which cannot deal with the dynamic nature of Cloud services and the mobility and velocity at which Cloud services operate.

Thus the deployment and consumption modalities of Cloud should be thought of not only within the construct of “internal” or “external” as it relates to asset/resource/service physical location, but also by whom they are being consumed and who is responsible for their governance, security and compliance to policies and standards.

This is not to suggest that the on- or off-premise location of an asset/resource/information does not affect the security and risk posture of an organization, because it does, but it also depends upon the following:

- The types of application/information/services being managed
- Who manages them and how
- How controls are integrated
- Regulatory issues

Table 1 illustrates the summarization of these points:

Security Guidance for Critical Areas of Focus in Cloud Computing

Table 1 - Cloud Computing Service Deployment

	Managed By ¹	Infrastructure Owned By ²	Infrastructure Located ³	Accessible and Consumed By ⁴
Public	Third Party Provider	Third Party Provider	Off-Premise	Untrusted
Managed	Third Party Provider	Third Party Provider	On-Premise	Trusted or Untrusted
Private				Trusted
Hybrid	Both Organization & Third Party Provider	Both Organization & Third Party Provider	Both On-Premise & Off-Premise	Trusted & Untrusted

¹ Management includes: operations, security, compliance, etc...

² Infrastructure implies physical infrastructure such as facilities, compute, network & storage equipment

³ Infrastructure Location is both physical and relative to an Organization's management umbrella

⁴ Trusted consumers of service are those who are considered part of an organization's legal/contractual umbrella including employees, contractors, & business partners. Untrusted consumers are those that may be authorized to consume some/all services but are not logical extensions of the organization.

As an example, one could classify a service as IaaS/Public/External (Amazon's AWS/EC2 offering is a good example) as well as SaaS/Managed/Internal (an internally-hosted, but third party-managed custom SaaS stack using Eucalyptus, as an example.)

Thus when assessing the impact a particular Cloud service may have on one's security posture and overall security architecture, it is necessary to classify the asset/resource/service within the context of not only its location but also its criticality and business impact as it relates to management and security. This means that an appropriate level of risk assessment is performed prior to entrusting it to the vagaries of "The Cloud."

Which Cloud service deployment and consumption model is used depends upon the nature of the service and the requirements that govern it. As we demonstrate later in this document, there are significant trade-offs in each of the models in terms of integrated features, extensibility, cost, administrative involvement and security.

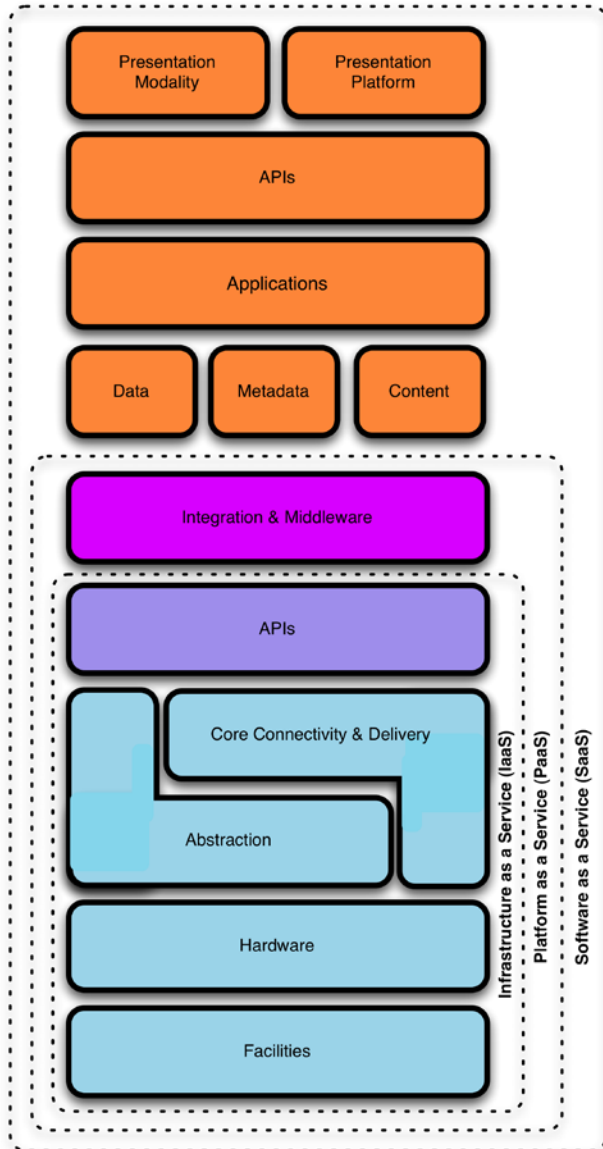
It is therefore important to be able to classify a Cloud service quickly and accurately and compare it to a reference model that is familiar to an IT networking or security professional.

Reference models such as that shown in Figure 2 allows one to visualize the boundaries of *aaS definitions, how and where a particular Cloud service fits, and also how the discrete *aaS models align and interact with one another. This is presented in an OSI-like layered structure with which security and network professionals should be familiar.

Security Guidance for Critical Areas of Focus in Cloud Computing

Considering each of the *aaS models as a self-contained “solution stack” of integrated functionality with IaaS providing the foundation, it becomes clear that the other two models – PaaS and SaaS – in turn build upon it.

Figure 2 - Cloud Reference Model



Each of the abstract layers in the reference model represents elements which when combined, comprise the services offerings in each class.

IaaS includes the entire infrastructure resource stack from the facilities to the hardware platforms that reside in them. Further, IaaS incorporates the capability to abstract resources (or not) as well as deliver physical and logical connectivity to those resources. Ultimately, IaaS provides a set of API's which allows for management and other forms of interaction with the infrastructure by the consumer of the service.

Amazon's AWS Elastic Compute Cloud (EC2) is a good example of an IaaS offering.

PaaS sits atop IaaS and adds an additional layer of integration with application development frameworks, middleware capabilities and functions such as database, messaging, and queuing that allows developers to build applications which are coupled to the platform and whose programming languages and tools are supported by the stack. Google's AppEngine is a good example of PaaS.

SaaS in turn is built upon the underlying IaaS and PaaS stacks and provides a self-contained operating environment used to deliver the entire user experience including the content, how it is presented,

the application(s) and management capabilities.

SalesForce.com is a good example of SaaS.

Security Guidance for Critical Areas of Focus in Cloud Computing

It should therefore be clear that there are significant trade-offs in each of the models in terms of features, openness (extensibility) and security.

Figure 3 demonstrates the interplay and trade-offs between the three *aaS models:

- Generally, SaaS provides a large amount of integrated features built directly into the offering with the least amount of extensibility and a generally a relatively high level of security (or at least a responsibility for security on the part of the provider).
- PaaS generally offers less integrated features since it is designed to enable developers to build their own applications on top of the platform and is therefore more extensible than SaaS by nature, but due to this balance trades off on security features and capabilities.
- IaaS provides few, if any, application-like features, provides for enormous extensibility but generally less security capabilities and functionality beyond protecting the infrastructure itself since it expects operating systems, applications and content to be managed and secured by the consumer.




Figure 3 - Trade-off's Across *aaS Offerings

The key takeaway from a security architecture perspective in comparing these models is that the lower down the stack the Cloud service provider stops, the more security capabilities and management the consumer is responsible for implementing and managing themselves.

This is critical because once a Cloud service can be classified and referenced against the model, mapping the security architecture, business and regulatory or other compliance requirements against it becomes a gap-analysis exercise to determine the general “security” posture of a service and how it relates to the assurance and protection requirements of an asset.

Figure 4 below shows an example of how mapping a Cloud service can be compared to a catalog of compensating controls to determine what existing controls exist and which do not, as provided by either the consumer, the Cloud service provider or another third party.

Security Guidance for Critical Areas of Focus in Cloud Computing

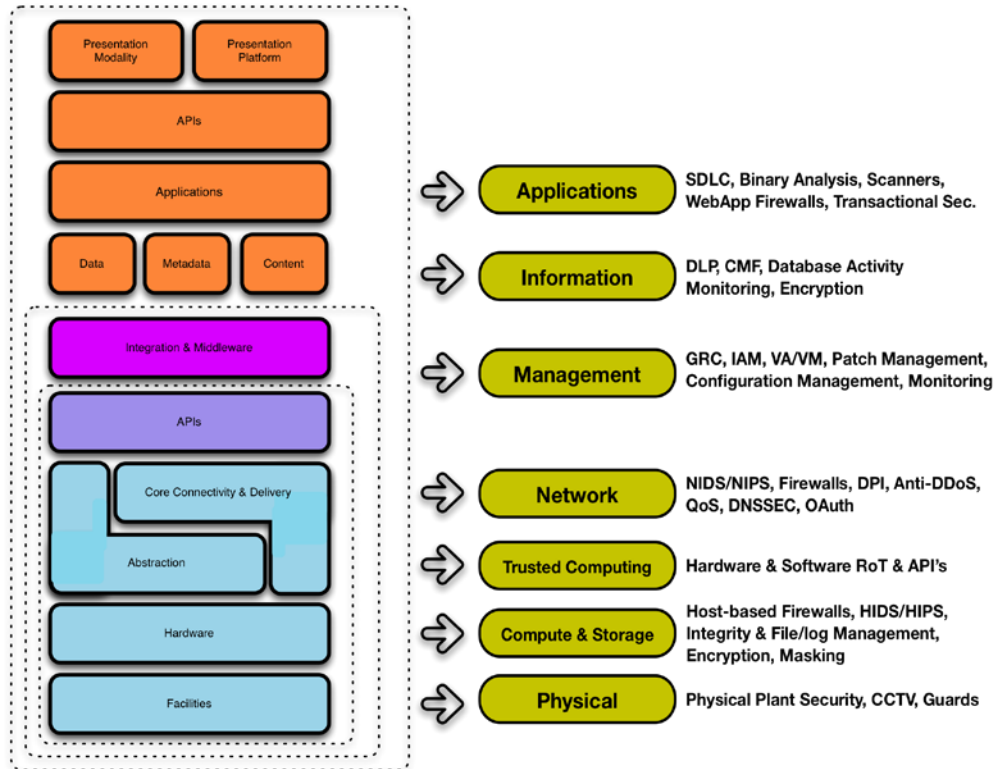


Figure 4 - Mapping the Cloud Model to the Security Model

Once this gap analysis is complete as governed by the requirements of any regulatory or other compliance mandates, it becomes much easier to determine what needs to be done in order to feed back into a risk assessment framework to determine how the gaps and ultimately how the risk should be addressed: accept, transfer or mitigate.

Conclusion

Understanding how architecture, technology, process and human capital requirements change or remain the same when deploying Cloud Computing services is critical. Without a clear understanding of the higher-level architectural implications of Cloud services, it is impossible to address more detailed issues in a rational way.

The keys to understanding how Cloud architecture impacts security architecture are a common and concise lexicon coupled with a consistent taxonomy of offerings by which Cloud services and architecture can be deconstructed, mapped to a model of compensating security and operational controls, risk assessment and management frameworks and in turn, compliance standards.

Section II. Governing in the Cloud

Domain 2: Governance and Enterprise Risk Management

Contributors: *Jeff Bardin*

Problem Statement

The ability to govern and measure enterprise risk within a company owned data center is difficult and surprisingly still in the stages of maturation in most organizations. Cloud computing brings new unknowns to governance and enterprise risk.

Online agreements and contracts of this type may be untested in a court of law and consumers have yet to taste an extended outage of services that they may someday determine to need on a 24x7 basis. Questions still remain on the ability of user organizations to assess the risk of the provider through onsite assessments. Governance of the provider's facilities and services could leave the user without recourse or recompense, placing the risk of use squarely on the shoulders of the user.

The storage and use of information that is considered sensitive by nature may be allowed but it could be unclear as to who is responsible in the event of a breach. If both the code authored by the user and the service delivered by the provider is flawed, who is responsible? Current statutes cover the majority of the United States but how are the laws of foreign countries, especially the European Union to be interpreted in the event of disputes? Many questions remain with respect to Cloud Governance and Enterprise Risk. We will attempt to cover some of the issues, define some solutions and raise a call for more collaboration and action.

Issues

- Many organizations seek the speed of deployment and lower cost of Cloud Services to support the delivery of mission critical services. This shift needs to be accompanied by an investment of a portion of the savings into increased regulatory scrutiny of risk, privacy and security of Cloud Services. Assessing the risk of a Cloud Service provider is a point in time effort that must be continued with ongoing risk management to effectively mitigate Cloud risks.
- Many of the leading Cloud Service Providers accept no responsibility for the data being stored in their infrastructure thereby not accepting any transference of risk. In addition, service level agreements are included in the online contracts defining the services being rendered and are potentially non-negotiable. If you want the service, you accept the online terms of use with conditional clauses and privacy statements that are subject to change without notice.
- Service Level Agreements focus upon availability of services and may not explain service quality, resolution times, critical success factors, key performance indicators, or offer any recourse to the user. The on time delivery of any new or remediated services may not be defined or mentioned.
- Financial viability of the cloud service provider is a critical issue and should be assessed in initial due diligence, and on an ongoing basis.

Security Guidance for Critical Areas of Focus in Cloud Computing

- Provider site certifications such as SAS 70, WebTrust® and SysTrust®, Service Capability & Performance (SCP) or ISO27001 can be directed as desired by the provider and are a point in time certification if there is any such certification.
- Third party support to the provider is not divulged or defined which could lead to additional governance issues and decrease the user risk due to ill conceived procedures or those that are lacking.
- Accurate metrics may not be available for the service(s) in use thereby limiting the available reporting and bringing pricing and service into question.
- Communication to service users is ad-hoc and at the convenience of the provider.
- Service incident response, service recovery and overall site resiliency could lead to extended outages of mission critical service components.
- Governance over regulatory and statutory audits as well as industry standard assessments could be left to the user without a view into the processes, procedures, and practices of the provider in the areas of access, identity management, and segregation of duties non-inclusively leaving control risks as an unknown quantity.
- The governance of data stored in Cloud Services is at risk due to the lack of clarity on data recovery, data backups, offsite storage, residual data due to virtual provisioning and data removal upon contract termination as well as the physical location of the data and the laws of that location (state / country).

Guidance

Organizations considering using Cloud Services should perform in depth due diligence prior to the execution of any service Terms of Service, Service Level Agreements or use. This due diligence should also assess the alignment, or misalignment of risks known at present and abilities of partners to work within and contribute to the customer's enterprise risk management program for the duration of the engagement. Here are some recommendations until such time as pertinent security and privacy concerns are addressed:

- Careful and comprehensive due diligence is required before deciding to store or transact any information in a Public Cloud that qualifies as personally identifiable information based upon regulations and statutes.
 - Examine creating a Private (Virtual) Cloud or a Hybrid Cloud that provides the appropriate level of controls while maintaining risk at an acceptable level.
 - Review what type of provider you require such as software (SaaS), infrastructure (IaaS) or platform (PaaS).
 - For PCI related transactions, use existing payment processors that are certified.
- Careful and comprehensive due diligence is required before deciding to use Public Cloud Services for mission critical components of your business unless you can manage customer expectations and draft an appropriate contract.
 - Examine creating a Private Cloud or a Hybrid Cloud that provides the appropriate level of controls while maintaining risk at an acceptable level.
 - Review what type of provider you require such as software (SaaS), infrastructure (IaaS) or platform (PaaS).

Security Guidance for Critical Areas of Focus in Cloud Computing

- Search for available Vertical Cloud service providers.
- For PCI related transactions, use existing payment processors who passed their point-in-time PCI assessment.
- Review what type of provider you prefer such as software, infrastructure or platform.
- Gain clarity on how pricing is truly performed with respect to bandwidth and CPU utilization in a shared environment. Compare usage as measured by the cloud service provider with your own log data, to ensure accuracy.
- Request clear documentation on how the facility and services are assessed for risk and audited for control weaknesses, the frequency of assessments and how control weaknesses are mitigated in a timely manner. Ask the service provider if they make the results of risk assessments available to their customers.
- Require the definition of what the provider considers to be critical success factors, key performance indicators and how they measure them relative to IT Service Management (Service Support and Service Delivery).
- Require a listing of all provider third party vendors, their third party vendors, their roles and responsibilities to the provider and their interfaces to your services.
- Request divulgence of incident response, recovery, and resiliency procedures for any and all sites and associated services.
- Request a review of all documented policies, procedures and processes associated with the site and associated services assessing the level of risk associated with the service.
- Require the provider to deliver a comprehensive list of the regulations and statutes that govern the site and associated services and how compliance with these items is executed.
- Perform full contract or terms of use due diligence to determine roles, responsibilities, and accountability. Ensure legal counsel review including an assessment of the enforceability of local contract provisions and laws in foreign or out-of-state jurisdictions.
 - Determine whether due diligence requirements encompass all material aspects of the cloud provider relationship, such as the provider's financial condition, reputation (e.g., reference checks), controls, key personnel, disaster recovery plans and tests, insurance, communications capabilities and use of subcontractors.
 - Scope of services;
 - Performance standards;
 - Rapid provisioning – de-provisioning;
 - Methods of multi-tenancy and resource sharing;
 - Pricing;
 - Controls;
 - Financial and control reporting;
 - Right to audit;
 - Ownership of data and programs;
 - Procedures for address a Legal Hold;
 - Confidentiality and security;
 - Regulatory compliance;
 - Indemnification;
 - Limitation of liability;
 - Dispute resolution;
 - Contract duration;
 - Restrictions on, or prior approval for, subcontractors;

Security Guidance for Critical Areas of Focus in Cloud Computing

- Termination and assignment, including timely return of data in a machine-readable format;
 - Insurance coverage;
 - Prevailing jurisdiction (where applicable);
 - Choice of Law (foreign outsourcing arrangements);
 - Regulatory access to data and information necessary for supervision; and
 - Business Continuity Planning.
- Review any material subcontractor relationships identified by the cloud provider or in the outsourcing contracts.

Call to Action – Collaboration

The relative maturity of Cloud Services will lead to history repeating itself with respect to security issues. Consumers, Businesses, Cloud Service Providers, and Information Security and Assurance professionals need to collaborate to shine a light on the potential issues and solutions listed above and to discover those not yet identified. The Cloud Security Alliance calls for collaboration in setting standard terms and requirements that drive governance and enterprise risk issues to a mature and acceptable state that allows for negotiation. Cloud Services is an inevitable evolution of the information society with a need for clearly defined governance and well thought out enterprise risk strategies appropriate for each different cloud offering. The CSA is working to address these issues so organizations can take full advantage of the agility, rich service options, flexible pricing and cost savings of Cloud Services.

References

3Tera. (2009, March 30). *3Tera AppLogic VPDC Service Level Agreement (SLA)*. Retrieved April 1, 2009, from 3Tera: <http://www.3tera.com/Terms/cloud-computing-sla.php>

Amazon. (2009, March 12). *AWS Customer Agreement*. Retrieved March 15, 2009, from Amazon Web Services Customer Agreement: <http://aws.amazon.com/agreement/>

enomaly. (2009, April 1). *enomaly Elastic Computing*. Retrieved April 1, 2009, from enomaly: <http://www.enomaly.com/Terms-of-Use.431.0.html>

FFIEC. (n.d.). *Outsourcing Technology Services Frame Set*. Retrieved March 21, 2009, from Booklet: Outsourcing Technology Services: http://www.ffiec.gov/ffiecinfobase/html_pages/outsourc_book_frame.htm

Microsoft. (2008, November). *About - Terms of Use | Azure Services Platform*. Retrieved March 14, 2009, from Azure Services Platform: <http://portal.ex.azure.microsoft.com/tou.aspx>

Nicholas A. Benvenuto, D. B. (2005, August 24). *Outsourcing - A Risk Management Perspective*. Rolling Meadows, Illinois, USA.

OpSource Marketing. (2009, April 1). *OpSource Brochure*. Retrieved April 1, 2009, from OpSource: <http://www.opsourcenet.com/datasheets/OpSourceBrochureDS.pdf>

OpSource. (2009, April 1). *OpSource AUP*. Retrieved April 1, 2009, from OpSource: <http://www.opsourcenet.com/uap.php>

ParaScale. (2009, April 1). *ParaScale Cloud Storage Software*. Retrieved April 1, 2009, from ParaScale: <http://www.parascale.com/index.php/home>

Domain 3: Legal

Contributors: *Francoise Gilbert*

© 2009 Francoise Gilbert

Problem Statement

Numerous United States laws require public and private organizations to protect the security of their information and computer systems.⁴ These laws may cover only specific markets; for example, the financial services or the health care industry. They may cover only specific regions; for example, a State security breach disclosure law is intended to protect the residents of that State. They may apply anywhere in the country; for example the unfair and deceptive practices provisions of the FTC Act. Any organization that does business in the United States is subject to one or more of these laws. These obligations to protect the security of personal or other data in its custody apply no matter where the information assets are located, including in the cloud.

In many instances, these obligations rest on the senior management. Under the Sarbanes Oxley Act, for example, responsibility for the company's financial data lies with the company Chief Executive Officer and Chief Financial Officer. The Gramm Leach Bliley Act (GLBA) Security Safeguards place responsibility for security with the Board of Directors of financial institutions. The Health Information Portability and Accountability Act (HIPAA) Security Safeguards require covered entities to designate a security official to be responsible for compliance with the security rule. The consent decrees issued by the Federal Trade Commission require the designation of an individual to coordinate, and be accountable for, the company's information security program.⁵ Numerous cases have also held that by virtue of their fiduciary obligations to a company, corporate directors have a duty of care, which includes responsibility for the security of the company's information systems; this responsibility may extend to safeguarding the integrity of all stored data.⁶

Where a third party – such as a cloud service provider – processes data on behalf of an organization, the same requirements as described above apply to mandate that the organization impose appropriate security measures on its subcontractors, service providers, or outsourcers. For example, the security and privacy rules under GLBA or HIPAA require that organizations require, in written contracts, their subcontractors to use reasonable security measures and comply with data privacy provisions. Abroad, similar concepts apply in many countries. For example, the Data Protection Laws of all Member States of the European Union include a security component that must be passed down to subcontractors. The laws hold organizations responsible for the acts of their subcontractors. Government agencies, such as the Federal Trade Commission or the State

⁴ Many foreign data protection laws create the same obligation. See, for example, the data protection laws that implement the principles in the 1995 Data Protection Directive of the European Union.

⁵ See, e.g., Federal Trade Commission, *In the Matter of Genica Corporation, Consent Order*, Feb. 5, 2009; <http://www.ftc.gov/os/caselist/0823113/090125genicaagree.pdf>

⁶ Example: *Caremark International Inc. Derivative Litigation* (698 International 959 (Del. Ch. 1996)); *Bell v. Michigan Council*, 2005 Mich App. Lexis 353 (Feb. 2005).

Security Guidance for Critical Areas of Focus in Cloud Computing

Attorneys General in the United States, and the Data Protection Agencies in the European Union, have consistently held organizations liable for the activities of their subcontractors.

In the cloud computing environment, important data, files, and records are entrusted to a third party. The company relinquishes most controls over these data. One or several cloud service providers might hold, store, or process a substantial part, or all of the company's data in their servers. Several vendors might provide services associated with the operation of the cloud. Each service provider holds one of the keys that allow access to the data. If the hosting or application service provider fails, the data might not be accessible. If the cloud spam filter fails, access to the data may be compromised, and the cloud service client may come under attack or the flow of its operations might be interrupted.

Continued Compliance Obligations

Organizations are custodians of the personal and other data entrusted to them. They have the same obligations to protect the integrity, security, and confidentiality of this data in-cloud and off-cloud. Data in the custody of cloud service providers must receive the same protection as when in the hands of their original owner or custodian. This duty of care is translated into different activities at different stages of the relationship between an organization and the cloud service provider: pre-contract due diligence, contract term negotiation, post contract monitoring, and termination and transition.

Due Diligence

Before entering into a cloud computing arrangement, a company should evaluate its own practices, needs, and restrictions, in order to identify the legal barriers and compliance requirements, associated with a proposed cloud computing transaction. For example, it should determine whether its business model allows for the use of cloud computing services, and under which conditions. The nature of its business might be such that any relinquishing control over the company data is restricted by law or creates serious security concerns.

In addition, the company should conduct due diligence of the proposed cloud service provider, in order to determine whether the offering will allow the company to fulfill its continued obligation to protect its assets. For example, the company may want to determine the technical or financial capabilities of the cloud service provider. How secure are its operations? Is it equipped to handle business interruption and disaster recovery?

Contract

After the parties have agreed in principle to the proposed arrangement, they must enter into a written agreement. Depending on the nature of the services, the contract may be in the form of a click-wrap agreement, which is not negotiated; or the parties may negotiate a more complex written document that is tailored to the specific situation. If a click-wrap agreement is the only agreement available, the cloud service client should balance the risks from foregoing negotiations against the actual benefits, financial savings, and ease of use promised by the cloud service provider. If the parties can negotiate a contract, they should ensure that the provisions of this contract address the needs and obligations of the parties both during the term of the contract and upon termination. Detailed, comprehensive provisions addressing the unique needs and risks of operating in a cloud environment should be negotiated.

Monitoring, Testing and Updating

Throughout the life of the relationship, both parties should ensure the performance of the contract according to its terms. The cloud service client should conduct periodic monitoring, testing, and evaluation of the services provided in order to verify that the required privacy and security measures are being used, and the processes and policies are being followed. The cloud service provider may have to change its practices and procedures in order to address new threats. Both the cloud service provider and cloud service client may need to adapt to new compliance requirements, if new laws are passed, or regulations are enacted during the term of the contract.

Termination and Transition

The termination of the cloud service contract may occur through the natural expiration of the term, or after a breach or default by one of the parties, or as the consequence of financial difficulties, through bankruptcy or reorganization proceedings. It is a stage where data are the most at risk because the cloud service client and provider are distracted. Nevertheless, the legal compliance requirements are not relaxed. It is therefore even more important for the parties to ensure the proper – and secure – winding down of the relationship in order to limit the risk of loss or alteration of the data. The cloud environment may create unique risks or enhanced exposure. The technology used – i.e., a distributed computing environment – may make it difficult to locate the data. Further, the parties are likely to be located in different jurisdictions, each with a different legal regime.

Selected Legal Issues

The following sections provide examples of legal issues that may arise in connection with cloud computing arrangements.

Understanding legal restrictions

There are numerous legal restrictions to entering into a cloud computing arrangement. These restrictions may come from federal laws or state laws, and their related regulations. They may stem from standards, or from preexisting contracts. They may result from foreign laws, and many other sources.

Federal Laws

Numerous federal laws and their related regulations, together with the orders issued by the FTC, require companies to adopt specific privacy and security measures when entering into a contract with the third party service provider.

State Laws

Numerous state laws also create an obligation on companies when entrusting certain categories of data to third parties. See, for example, the restrictions on the use of social security numbers or driver license numbers. State laws that address information security issues generally require at a minimum that the company have a written contract with the service provider, with reasonable security measures. See for example the extensive requirements under the Massachusetts Security Regulations.

Standards

Standards such as PCI DSS or ISO 27001 also create a domino effect similar to that of federal state laws. Companies that are subject to PCI DSS or ISO 27001 must pass onto their subcontractors the same obligation to meet the standard to which they are subject.

Corporate Duty of Care

Numerous laws or common law place upon company management a duty and an obligation to safeguard the company's assets. The proposed cloud computing arrangement must provide sufficient reasonable security safeguards that are commensurate with the sensitivity of the materials that will be entrusted to the cloud service provider.

Foreign Laws

Many foreign countries have adopted data protection laws that follow the European Union model. Under these laws, the data controller (typically the entity that has the primary relationship with an individual) remains responsible for the collection and processing of personal data even where the data are processed by third parties. The data controller is required to ensure that any third party processing personal data on its behalf takes adequate technical and organizational security measures to safeguard the data. This includes conducting due diligence before entering into the contract.

Understanding where the data will be located

Location of the data

When a business entrusts its data to a third party, it is vulnerable. Its data are sitting in someone's computer, and in someone else's facility. Many things can go wrong. The cloud service provider may go out of business or may decide to hold the data hostage if there is a dispute. It is important for a company to understand in which country its data will be hosted, because the location of the data directly affects the choice of the law that will govern the data. If the data reside in China, it is likely that Chinese law will govern access to the servers where the data are hosted. If the client demands access to its data would Chinese law apply since the data are stored in China? Further, Chinese law may permit the Chinese Government to have unlimited access to the data stored in its territory whereas there might be stricter restrictions to access by the United States Government to data stored in the United States.

Location of the service provider

The cloud service client may also wish to ensure that the cloud service provider with which it contracted actually performs the services so that the client may keep a direct relationship with, and control over, the custodian of its data. To this end, the client may want to contractually limit the cloud service provider's ability to subcontract its obligations.

Combination or commingling of the data

The client may want to ensure that its data will be stored separately from the data of other clients. If its data are combined or commingled with those of other clients, these data may be more vulnerable. For example, viruses might be transmitted. If another client is the victim of a hack attack, the attack might affect the availability or integrity of the data of other companies located in the same environment. If data are commingled or combined, the client might have problems getting its data back.

Technical and logistical restrictions

It might be difficult or impossible for the cloud service provider to assure the client of the location of its data at all times because of the inherent nature of cloud computing, unless specific servers are dedicated to servicing a particular client. The cloud service provider may also need to address logistics. It may be important for the cloud service provider to have the ability to subcontract hosting or other services to third parties. This flexibility might be necessary in order to limit financial exposures, and ensure the ability to reconfigure a cloud network as needed.

Ensuring privacy protection

Privacy obligations

Companies have a legal obligation to protect the privacy of their clients or employees, to ensure that the data are not used for secondary uses, and are not disclosed to third parties. In a cloud environment, each organization must ensure that it will continue to be able to fulfill the legal requirements to which it is subject (e.g., through HIPAA, GLBA) and to meet the promises and commitments it made in its privacy notices. It must also ensure that individuals' choices about their information are respected in the cloud environment. For example, individuals may have agreed only to specific uses of their information. Data in the cloud must be used only for the purposes for which they were collected.

Data must not be transferred out of the cloud without compliance with the applicable privacy law or privacy notice. See for example the restrictions under HIPAA. Outward transfers of data may only occur when authorized by law, or as provided for in the privacy notice, or with the consent of the individual data subjects.

If the privacy notice allows individual data subjects to have access to their personal data, and to have this information modified or deleted, the cloud service provider must allow this access, modification and deletion rights to be exercised to the same extent, within the same delay, as it would in an off-cloud relationship.

Cloud service provider's privacy obligations

The cloud computing service provider may prefer to limit the extent to which the client can retain control over its data, in order to keep all of its clients aligned within the same structure or business model. Giving a client control over its data (such as location of the data, or access to the data) is costly, reduces efficiencies, and limits the cloud service provider's freedom to move the data as convenient for its operations.

On the other hand, cloud service providers that target organizations in certain markets – e.g., the healthcare industry – will have to be prepared to agree to comply with the legal restrictions that govern these markets. For example, HIPAA (as modified by the HITECH Act) would preclude cloud service providers from receiving electronic prohibited health information unless appropriate security measures are in place.

Coping with foreign laws

Prohibition against cross border transfers

A global company that wishes to take advantage of cloud services will want to ensure that this use does not jeopardize its subsidiaries, clients, business partners and others which may be subject to foreign laws with different restrictions than those in effect in the United States. The US company will want to know where the personal data of its employees, clients and others will be located, so that it can address the specific restrictions that foreign data protection laws may impose. For example, a German subsidiary may not oppose the use of a cloud service provider in Argentina, but it will object to the transfer of its data to Turkey, Mexico, or the United States. Knowing where the cloud service provider will host the data is a prerequisite to implementing the required measures to ensure compliance with local laws that restrict the cross border flow of data.

Security Guidance for Critical Areas of Focus in Cloud Computing

Cloud computing users and service providers will need a clear understanding of the complex restrictions and requirements created under the data protection laws of the European Union member states and of several other non-EU countries with similar laws. Cumbersome restrictions hamper the transfer of data outside of these countries. Their laws require data controllers (who originally collected the data) to inform individuals that their data will be processed abroad, and to obtain their consent to the transfer. In addition, the data controller and the recipient of the data may have to enter into special contracts that must be approved by the local Data Protection Authority. For data that are sent to the United States, the US recipient of the data may self-certify of its data protection procedures by registering with the U.S. Department of Commerce. Additional contracts limiting the use or reuse of the data may be required. Permission from local Data Protection Authorities may be needed.

Cloud service provider's compliance obligations

Cloud service providers should plan strategically where their servers (or the servers of their own service providers or associates) will be located. The choice of the location of servers will directly affect jurisdictional issues and choice of law. Similarly, they should evaluate carefully which market they want to approach and where their clients are located. Refraining from doing business in or with countries with restrictive laws might help cloud computing service providers to reduce the risk of running into unexpected problems.

Secondary use of the data

Secondary uses of the data

As the custodian of the personal data of its employees or clients, and of the company's other intellectual property assets, a company that uses cloud computing services should ensure that it retains ownership of its data. If it is concerned about the confidentiality, security, or privacy of its databases, or if it wishes to control who has access to its data, it will want to understand what the cloud service provider could do with the data in its custody. In most cases, it will want to limit the scope of what the cloud service provider is permitted to do with the company's data.

Controlling who has access to the metadata associated with its data or with the uses of its data may be important. A company that holds sensitive personal data, company trade secrets, or other valuable information may wish to limit access to, or use of the traffic information associated with this data by the cloud service provider. For example, who looked at what information, and when or what queries or searches were run may have great value. The cloud service provider may want the ability to mine the company's data or metadata for secondary uses, such as for marketing or market research purposes. Numerous cloud service providers offer free access to their services or their applications with the view to mine the data in their custody in order to offer advertising services.

In other cases, an organization or an individual may not mind the potential intrusion in their affairs if they determine that the financial benefit and ease of access to their information through the cloud outweighs the potential that third parties may access their files, pictures, or correspondence.

Bankruptcy issues

In today's environment, businesses also need to plan for the possibility that their contractor might go out of business, be reorganized or file for bankruptcy. Each such event raises the issue of ownership of assets. A business should plan as appropriate to ensure that in the event of the cloud service provider's demise, it will be able to retrieve its data from the service provider bankruptcy trustee. On the other hand, a cloud service provider may wish to ensure that it has the necessary

authority to dispose of the data of a client who is not paying for its services or has filed for protection from its creditors.

Complying with information security laws

Legal requirements

Companies must ensure a reasonable level of security at all times, in order to protect their intellectual property and other assets, and the personal data of their employees, clients, and contacts. This obligation stems from numerous laws, regulations, standards, cases, and best practices. Further, the general duty of care that is imposed on company management creates another mandate for corporate leaders. The obligation to maintain a reasonable level of security may also result from contracts. If it fails to provide an appropriate level of security, a company may be exposed to private suits and class actions, as well as to enforcement actions by enforcement agencies.

The laws that address security issues (e.g., GLBA, HIPAA, Data Protection Laws in Europe, etc.), and the security standards (e.g. PCI DSS, ISO 27001) require that businesses enter into services agreements with their subcontractors and service providers and mandate the use of adequate technical, physical and organizational measures in order to safeguard the data entrusted to them. These laws, regulations, standards and the related best practices, also require businesses to ensure that these obligations will be fulfilled by conducting due diligence (before execution of the contract) or security audits (during performance of the contract).

In a cloud computing environment, numerous security issues may need to be addressed. They will have to be reflected in specific provisions of the service agreement that clarify the respective commitments of the cloud service provider and the client. For example, the client may want the cloud computing services agreement to provide for protections from other clients of the cloud service provider, such as through the use of “silos” that will ensure that one client’s data are segregated from another client’s data. The client may also want to specify protections from the cloud service provider’s employees, such as physical access restrictions, or limitations to what information may be used. Technical measures may be specified, such as the use of strong passwords, encryption, or firewalls.

The client may also wish to negotiate the right to conduct periodic audits of the security measures, in order to verify that the cloud service provider complies with the requirements to which the client is subject. Thorough due diligence and security audits allow the identification of problems and issues, and thereby reduce risk and potential damages. They may also prove useful in a lawsuit to counter a negligence claim by showing that the parties had exercised due care.

Cloud service provider’s viewpoint

The cloud service provider may attempt to retain the freedom to define the applicable security measures in order to retain control, limit expenses, and ensure flexibility. Customized security requirements are costly.

Nevertheless, cloud service providers should be mindful of the numerous legal constraints to which they are subject. In addition to obligations that may stem from contracts, cloud service providers are subject to numerous laws and regulations. In the past few years, an increasing number of laws have begun targeting cloud service providers as well as their clients. These laws place directly on the cloud service provider requirements that are similar to those that are placed on the client side. For example, the recent HITECH Act (which addresses healthcare issues)

Security Guidance for Critical Areas of Focus in Cloud Computing

creates specific obligations and liabilities for cloud service providers. The State security breach disclosure laws place on cloud service providers specific obligations to disclose security breaches that affect their clients' data.

Handling security breaches

More than 44 US States have adopted security breach disclosure laws that require the custodian of specified categories of personal data to notify individuals whose data might have been compromised in a breach of security. Both the cloud service provider and the client need to ensure that the other party will promptly disclose the existence of a breach. The client may have to inform its own customers or employees of the occurrence of the breach. The cloud service provider may also have to inform its other clients that the breach occurred.

Since cloud computing creates an environment where much of the resources are shared, both parties have responsibilities to promptly disclose a breach of security that affects its operation. This disclosure is necessary to prevent a ripple effect and more infections within or outside the cloud. Both parties may need to be prepared to address a security breach. They should have in place a security incident response plan to address the security breach thoroughly and expeditiously.

Ensuring business continuity

It is not clear whether the cloud environment creates more risks of business interruption than the traditional computing environment. The fact that numerous platforms are interconnected with each other may reduce the risk of a fatal outage. At the same time, the structure of the cloud may also make it more attractive to hackers or terrorists who might increase the strength and frequency of their attacks.

All users of cloud computing services need to ensure the continuity of their operation and uninterrupted access to their data. Any disruption in the operations will be critical. Disruptions or interruptions always affect the security of the systems and networks, and they may affect the viability of the company's business. Several laws and regulations create a legal obligation to maintain their applications running. For example, ensuring business continuity and disaster recovery is required by the HIPAA Security Safeguards. A hospital that provides technology or medical information database services to the physicians on its staff must provide continued access to critical information in order to ensure proper patient care.

The outage may affect only the cloud service provider's operation, or it may affect the Internet service. Most businesses will want to ensure that the cloud service provider has in place proper business continuity and disaster recovery capabilities because business continuity is essential to protect the viability of the business, and in some cases because of compliance requirements.

Addressing an outage of the Internet service might be more problematic, and raises additional legal issues. Since the company has a continuing obligation to ensure the protection and availability of its data, it may opt not to use a cloud environment for data that are critical to the company and that it must be able to access on a 24x7x365 basis.

Most reputable cloud service providers have measures for disaster recovery and business continuity. However, these measures may not always be sufficiently extensive, efficient, or

Security Guidance for Critical Areas of Focus in Cloud Computing

sophisticated. A cloud service provider may need to evaluate its clients' needs in order to provide the adequate level of assurance of the continued operations.

Responding to litigation requests

If there is a civil suit in which the cloud service client is a party, or if there is an investigation by a government agency, the cloud service provider is likely to receive a request for access to the information that it holds as the hosting entity.

Electronic Discovery

The cloud service client may also be involved in litigation. In this regard, it is likely that it will have to produce evidence that might be located in the cloud. It is very important that the cloud service client work in advance with its service provider to identify how the parties will work and cooperate to address compliance with the E-Discovery provisions of the Federal Rules of Civil Procedure and the State equivalents to these laws.

Response to a subpoena or search warrant

The cloud service provider is likely to receive from third parties a request to provide information, in the form of a Subpoena, a Warrant, or Court Order in which access to the client data is requested. The client may want to have the ability to fight the request for access in order to protect the confidentiality or secrecy of the data sought. To this end, the cloud services agreement should require the cloud service provider to notify the company that a subpoena was received, and give the company time to fight the request for access.

The cloud service provider might be tempted to reply to the request by opening its facilities and providing the requestors with whatever information is identified in the access request. Before doing so, the cloud service provider should ensure that the request is in good order, and uses the appropriate legal method. The cloud service provider should carefully analyze the request before disclosing information in its custody. Complex laws apply depending on the specific nature of the information, its location, etc. For example, different rules apply for requesting access to the content of an email depending on whether or not the email has been opened, and how long the email has been stored. Different rules apply if the information requested is the content of the email, or only the transactional data about the email (e.g., when sent, to whom, etc.).

Monitoring the cloud service activities

During the term of the relationship, the cloud service client must supervise the activities of the cloud service provider, in order to ensure that adequate security measures are being used, and that they comply with the contract. Monitoring the cloud service providers' performance, testing the vulnerabilities of the system, requesting changes if new threats and vulnerabilities are uncovered are part of the company's legal obligations. See for example the requirements under the Security Safeguards under HIPAA or GLBA, or those in the orders issued by the Federal Trade Commission or the State Attorneys General.

To this end the cloud services agreement must provide for the company's ability to conduct this monitoring and these tests. In addition, the company management must plan for and organize these tests, and this monitoring.

While it is in its interest to ensure the security of its client's data, the cloud service provider will want to push back on monitoring and testing by the client. These audits and tests are disruptive

Security Guidance for Critical Areas of Focus in Cloud Computing

and cost significant management time. The cloud service provider may offer instead to provide the company with a certificate from an independent third party with a description of the results of the audit.

Ensuring a smooth termination

Numerous events may lead to the termination of the relationship. The contract for cloud services may expire at the end of its term and not be renewed. The contract may also be terminated for cause or without cause according to its terms. A default or a material breach by one of the parties, or financial difficulties and bankruptcy proceedings might also cause the termination of the relationship.

In all instances, the company remains responsible for the data it entrusted to the cloud service provider, and it must retrieve this data, or ensure its destruction if they are no longer needed. This requirement may stem from various data protection laws and regulations, or from a common law duty of care. In addition, the company may be responsible to its own customers for ensuring that the customer's information is always available. See, for example, the obligations under some data protection laws.

Too frequently, companies have great problems recovering their data. The cloud service provider may keep the data hostage as a way to generate needed funds. It may have commingled the company's data with the data of other clients to save space or for technical reasons, and may be unable to disentangle the different components.

The parties should ensure that the service agreement anticipates these problems, defines proper procedures in the event of termination, and identifies work-arounds to address disputes. For example, one alternative would be to require that the cloud service provider periodically deliver to the company copies of the data in its custody. This may not be a viable alternative if the structure of the databases do not allow for the segregation of data.

The cloud service provider may have trouble providing the company with copies of the data, for example if the data are commingled with other data. If the company requests periodic copies of its databases, the cloud service provider may need to charge a fee in order to compensate for the administrative time in preparing the copies.

References

Books

Francoise Gilbert, *Global Privacy & Security* (Aspen Publishing 2009).

E. Michael Power & Roland L. Trope, *Sailing in Dangerous Waters: A Director's Guide to Data Governance* (American Bar Association, 2005).

Thomas Smedinghoff, *Information Security Law: Emerging Standard for Corporate Compliance* (ITGP 2008).

Websites

Cloud computing definitions and business models:
http://p2pfoundation.net/Cloud_Computing
Definition, technical aspects, business models

Security Guidance for Critical Areas of Focus in Cloud Computing

Cloud Computing Incidents Database:

http://wiki.cloudcommunity.org/wiki/CloudComputing:Incidents_Database

Records and monitors verifiable, noteworthy events that affect cloud computing providers, such as outages, security issues, and breaches

Domain 4: Electronic Discovery

Contributor: *Jeffrey Ritter, Esq.*

Problem Statement

Legal systems in both developed and developing economies generally presume that a company, as a legal entity, will possess and control the records and information assets that may serve as evidence in legal proceedings in which a company may be involved. Further, there are various important affirmative legal duties for a company to preserve and produce those records and information assets in those legal proceedings, including regulatory reporting (tax records, environmental discharge reports), compliance audits, internal investigations and, of course, civil litigation.

In the 21st century, cloud computing business models (a) challenge the presumption that a company possesses, or even controls, all of the digital business information for which the law imposes duties to preserve and produce, and (b) potentially jeopardize a company's ability to preserve and produce required records and information. As a result, companies face substantial barriers to implementing cloud computing solutions if, as a result, their compliance capabilities and legal profiles are compromised.

The rules of evidence in both civil and common law systems also emphasize the importance that business records and information be authentic and reliable as evidence. It is vital to assuring the long-term success of cloud computing that information security management be a strong feature of the varied service agreements, in order that the records and information in the custody and control of the service provider align to the standards for authenticity and reliability required by their customers and the surrounding legal environments.

Essential Legal Rules

During the last five years, increased legal focus on digital information as evidence has initiated a far more comprehensive awareness of the importance of records and information management ("RIM") programs and services as an integral part of corporate governance. The overall effectiveness of legal systems, both in developed and developing economies, demands a heavy reliance on corporations, partnerships and other business types, as legal entities, to create, manage and control the records and information that document events, transactions and activities of potential legal significance. As a result, in order to be in business, companies face varied affirmative legal duties to preserve and produce their records in legal proceedings, including regulatory reporting (tax records, environmental discharge reports), compliance audits, internal investigations and, of course, civil litigation.

As the first decade in the 21st century nears completion, the business community has seen a rapid transformation toward the creation and management of all business records in digital form. Generally, the laws and regulations, as well as judicial rules, have adapted quickly to this digital shift; however, companies have been far less agile in understanding the impact of digital computing on their legal obligations and developing, in turn, appropriate RIM programs and services. Today, the costs confronted by companies in electronic discovery activities (usually for civil litigation) are one of the largest uncontrolled costs in business, largely as a result of historic

Security Guidance for Critical Areas of Focus in Cloud Computing

failures to invest appropriately in RIM programs and services for digital assets required by law to be maintained and accessible.

The Impact of Cloud Computing on Records and Information Management

Cloud computing business models routinely involve:

- The transfer to a service provider of operational control of the security environment in which business information is created, processed, and managed, and
- A shift to the service provider of the primary responsibility to maintain and store for its customer all of the related digital assets (inbound information, processed information and operational data) for at least the duration of the business relationship.

But doing so creates significant legal and compliance concerns.

In transferring both operational control and primary storage responsibilities, cloud computing contradicts (a) the presumption that a company possesses and controls the digital business information for which the law imposes duties to preserve and produce, and (b) jeopardizes a company's ability to preserve and produce required records and information. Routinely, these issues are not addressed when defining the cloud computing business models, developing and offering service proposals (and related terms and conditions) or initial negotiations.

As a result, those charged with legal governance and compliance will often oppose moving forward with cloud computing solutions that otherwise offer compelling business advantages, creating real barriers to the solutions. When the issues are raised, delays in negotiations occur, pricing schedules are adjusted (usually to the detriment of the customer's budget), and the resulting contract is more complex.

Perhaps worse, the client's legal and compliance managers will proceed to contract signing without an awareness of the overall impact of these issues and allow the solutions to be installed without considering their implications for future compliance. The results are often ticking time-bombs, creating potentially adverse future situations that are neither favorable to the client nor the service provider. When the client has a legal duty to produce information in the control of the service provider, when the contract is silent, the cooperation (and costs) of the service provider can be quite random and unpredictable.

The Rules of Evidence

The rules of evidence that govern in both civil law countries (such as France or Italy) and common law countries (such as the United States or Australia) introduce further challenges to various cloud computing business models. These rules emphasize that any business records offered as evidence must be shown to be authentic and reliable; courts now are realizing that information security management services are critical to understand in order to make decisions on whether digital information may be accepted as evidence.

As cloud providers take responsibility for, and execute, the management and storage of digital information, it is vital that they employ information security services that will assure the records under their control will, in fact, serve the critical function of serving as evidence for the truth on behalf of the client and their stakeholders (employees, customers, management, investors).

Security Guidance for Critical Areas of Focus in Cloud Computing

Failing to do so places in jeopardy one of the most useful functions of digital information, and undermines the value and cost-effectiveness of cloud computing services.

Many enterprises will find themselves in a situation in which they will have to maintain full copies of Cloud-resident data to ensure forensic integrity, availability, performance, “search ability”, and/or preservation. The cost of maintaining a full copy of such data may outweigh the benefits of implementing a Cloud-based solution to begin with. Such a situation should be carefully considered by enterprises with high levels of civil litigation activity regarding intellectual property disputes, financial fraud, and other disputes centered around sensitive data.

Electronic Discovery Services in the Cloud

For many companies, the risks are further increased since the service providers for electronic discovery (forensic services, collection, processing, review, production) and law firms are increasingly using the Cloud to make the collected electronically stored information (ESI) available to all of the parties. These business practices create several risk areas relating to information security:

- The security controls employed to transfer ESI from the corporate systems (or their cloud service providers) to the e-discovery vendor or law firm;
- The security services provided to protect the collected ESI from improper access, alteration or destruction while held in the custody of the e-discovery vendor or law firm; and
- The security services employed to allow authorized users (including adverse counsel) online access through the cloud to the collected ESI held by the e-discovery vendor or law firm.

Conclusion

Cloud computing challenges one of the most important functions of digital information—serving as reliable evidence for proving the truth. Information security services can be enormously useful at defining, and addressing, the numerous legal and business issues that arise when digital information with legal significance is shifted from inside the enterprise into the custody and management of service providers in the cloud.

References:

Federal Rules of Civil Procedure

Defending Electronic Mail as Evidence—The Critical E-Discovery Questions (available at www.cqdiscovery.com)

Contracting for Information Security: Model Contract Terms (published by the Internet Security Alliance and available at www.cqdiscovery.com)

Contracting for Certified Information Security: Model Contract Terms and Analysis (published by the Internet Security Alliance and available at www.cqdiscovery.com)

Domain 5: Compliance and Audit

Contributors: *Shawn R. Chaput*

Problem Statement

With cloud computing resources as a viable and cost effective means to outsource entire systems, maintaining compliance with your security policy and the various regulatory and legislative requirements to which your company is privy can become even more difficult to demonstrate. Worse, the cost of auditing that compliance is likely to increase without proper planning. With that in mind, it's imperative to consider all of your requirements and options prior to progressing with cloud computing plans.

Issues

There are many issues to assess when considering the use of cloud computing resources. These topics need to be kept in mind when evaluating your regulatory obligations and how they may impact the option of computing in the cloud.

Firstly, typical cloud computing environments have unclear boundaries as to where the processing or storage of data physically occurs or resides. While this is a fundamental element of cloud computing, it raises concerns with respect to many privacy legislations. For instance, the European Union's Data Protection Directive (EU DPD) has restrictions around the movement and access of specific types of data across political borders. Regional breach notification laws become difficult to conceptualize as well. Understanding where a cloud security breach occurred may not be as simple as it would otherwise be in a typical enterprise managed solution and thus the potential notification obligations become unclear..

Secondly, similar to the first point, without understanding how a cloud computing vendor stores your data, multiple copies of your data may exist. Again, depending on the type of data, this could prove troublesome with respect to regulatory compliance. Further, many laws have been interpreted to proscribe that sensitive data such as Personally Identifiable Information (PII) should be segregated from other types of data to be adequately protected. Using the cloud and an array of virtualized systems, it can become increasingly difficult to demonstrate appropriate segregation in a dynamic and shared environment.

Most importantly, the ability to audit all of these requirements in a cloud environment can be difficult and significantly more expensive. Further, these aren't the only audit challenges you'll encounter with cloud computing. To fully understand the infrastructure and relative security associated with a cloud computing solution, one might suggest an external audit, but without having a formal adopted framework for such activities, the validity of these assessments could be suspect. As with most audits, isolating the scope becomes a fundamental task to understanding the environment. Assuming a third party audit, driven by the provider of cloud based services, the scope of the audit may not meet the needs of your company or regulators. The validity of these reports can then become irrelevant.

Other audit challenges exist as well. Typical drivers of cost, such as usage and Service Level Agreements (SLA)s need to be considered for audit purposes as well. A means of formalizing the metrics indicative of meeting the obligations and other criteria need to be accounted for and auditable; none of which, at present exist in any standardized approach.

Security Guidance for Critical Areas of Focus in Cloud Computing

Using cloud computing resources and maintaining compliance may not be simple to do but it certainly is possible. Depending on the maturity of your organization's security program, utilizing cloud providers and remaining compliant, may be an achievable task or present significant challenges.

Guidance

Know Your Legal Obligations

Paramount to all other task and guidance and regardless of whether you intend to use cloud computing resources, your organization must understand necessary legal requirements. This isn't a trivial task and rarely are the requirements easily identified for companies in a meaningful and practical manner, but the activity must be conducted. The regulatory landscape is typically dictated by the industry in which you reside. Depending on where your organization operates, you are likely subject to a lengthy collection of legislation which governs how you treat specific types of data and it is your obligation to understand it and remain compliant. Without understanding your obligations, an organization cannot formulate its data processing requirements. It is recommended that you engage your internal auditors, external auditors and legal council to ensure nothing gets left out.

Classify / Label your Data & Systems

In order to adequately protect your data, your company must classify it. Considering the regulatory and legislative requirements discussed earlier, your organization needs to classify its data to isolate that which requires the most stringent protection from the public or otherwise less sensitive data. The data and systems must also be clearly labeled and the processes surrounding the handling of the data formalized. At this point, your organization can decide to only consider cloud computing resources for data and systems not classified at a certain level which would be subject to burdensome regulatory requirements.

External Risk Assessment

A third party risk assessment with respect to the systems and data being considered for cloud resources should be conducted to ensure all risks are identified and accounted for. This should include a Privacy Impact Assessment (PIA) as well as other typical Threat Risk Assessments (TRA). Impacts to other internal systems, such as Data Leakage Protection (DLP) systems should also be considered. Be prepared to discover extensive risks with costly remediation strategies in order to consider cloud computing for regulated data.

Do Your Diligence / External Reports

At a minimum, you need to understand the security of the organization hosting your cloud computing resources and what they're prepared to offer. If you have very stringent security requirements, you may want to mandate that your cloud provider be certified to ISO/IEC 27001:2005 annually. It's also likely that your organization will need to improve your processes and operational security maturity to manage your cloud provider to that level of security. It is important to utilize the risk assessment and data classification exercises previously mentioned to provide the amount of security required ensuring the appropriate confidentiality, integrity and availability of your data and systems without over spending. Assuming ISO/IEC 27001:2005 certification is too costly or not available within the class of service you seek, the assurance statement most likely to be available is the Statement on Auditing Standards (SAS)70 Type II. Work these requirements into the contract requirements and ensure you see a previous certificate of compliance prior to agreement.

Security Guidance for Critical Areas of Focus in Cloud Computing

Similarly, your company should demand the results of external security scans and penetration tests on a regular basis due to the unique attack surfaces associated with cloud computing.

The value of certifications such as ISO/IEC 27001:2005 or audit statements like SAS 70 are the source of significant debate among security professionals. Skeptics will point out that through the scoping process, an organization can exclude critical systems and process from scrutiny and present an unrealistic picture of organizational security. This is a legitimate issue, and our recommendation is that domain experts develop standards relating to scoping of these and other certifications, so that over time broad scoping will be expected by the customer. An ISO certification that has been based upon a comprehensive security program is an outcome that customers must demand. In the end, this will benefit the cloud provider as well, as a certifiably robust security program will pay for itself in reduced requests for audit.

Understand Where the Data Will Be

If your company is considering using cloud computing resources for regulated data, it may become imperative to understand where the data will be processed and stored under any and all situations. Of course, this task is far from simple for all parties, including cloud computing providers. However, with respect to legislative compliance surrounding where data can and cannot be transmitted or stored, the cloud computing provider will need to be able to demonstrate assurance that the data will be where they say it is and only there.

This applies to third parties and other outsourcers used by the cloud computing provider. If the provider has reciprocal arrangements or other types of potential outsourcing of the resources, strict attention to how this data is managed, handled and located must precipitate through to that third party arrangement.

If the potential provider you've engaged cannot do this, investigate others. As this requirement becomes more prevalent, it's likely the option will likely become more available. Remember, if that assurance cannot be provided, some of your data and processing cannot use public cloud computing resources as defined in Domain 1 without exception. Private clouds may be the appropriate option in this case.

Other Considerations

Aside from having your typical comprehensive "right to audit" clauses in contracts, further consideration should be given to the content of the agreement with respect to its "auditability". Specific care should be given to the process of defining the necessary SLAs and security Service Level Objectives (SLO)s which necessitate the collection of documentation outlining expectations. Make sure these documents are available for both parties and that supporting documentation is created to demonstrate achievement. This will likely require the development of standardization with how cloud computing operations function. For instance, a standardized approach to how data is stored, processed or accessed and controlled would be necessary. These standards would be the primary means against which the SLAs and SLOs are measured. This can also provide the foundation for further performance metrics and other requirements, such as how and when the infrastructure gets scaled up and out.

But do not forget, the security of the cloud computing environment isn't mutually exclusive of your organizations internal policies, procedures, standards, guidelines and processes. As part of any audit of the environment your organization may commission, your internal organization's operations should be included. Outsourcing the processing and storage of a system doesn't make it inherently more secure as the technical security is only as strong as your company's weakest process.

Security Guidance for Critical Areas of Focus in Cloud Computing

References

Wood, Lamont, “Cloud Computing and Compliance: Be Careful Up There”, ITWorld, January 30, 2009

Mather, Tim, “Cloud Computing is on the Up, but what are the Security Issues?”, Secure Computing Magazine (UK), March 2, 2009.

Cloud Computing: Bill of Rights, http://wiki.cloudcomputing.org/wiki/CloudComputing:Bill_of_Rights

Raywood, Dan, “Data Privacy Clarification Could Lead to Greater Confidence in Cloud Computing”, Secure Computing Magazine (UK), March 9, 2009.

“Auditing the Cloud”, Grid Gurus, http://gridgurus.typepad.com/grid_gurus/2008/10/auditing-the-cl.html, October 20, 2008

Roiter, Neil, “How to Secure Cloud Computing”, Information Security Magazine, http://searchSecurity.techtarget.com/magazineFeature/0,296894,sid14_gci1349670,00.html, March 2009.

Domain 6: Information Lifecycle Management

Contributors: *Jeffrey N. Reich*

Problem Statement

Every organization has data. Management of data as information should be maintained throughout the life of the data. Placing data “in the cloud” as part of services removes many of the direct physical controls that a data owner has over data and does not necessarily replace those physical controls with an equal level of other compensating controls. Since logical controls depend on physical controls as a foundation, the entire data control structure may be at risk.

Issues and Guidance

Historically, information, in the form of computerized data, has been best protected with a combination of physical and logical controls under the management of the data owner. Cloud Computing, focusing on services, introduces a level of abstraction that masks much of the physical infrastructure. As a result, most of the physical controls are no longer under the management of the data owner. Every organization should have a defined Information Lifecycle for its data.

The lifecycle begins when data are created. This starts when data structures are defined and codified. Data Classification should be put in place then and continued throughout the lifecycle. Appropriate data classification must take into consideration the sensitivity of the data and who should be able to access or control the data under given circumstances. When data are placed into the cloud, some of the considerations to be added to the lifecycle are:

- Specific identification of all controls being used – What levels of physical and logical controls can be verified to be in place at all times without exception? The user must always assume that without validation of physical and logical access controls, the controls are not in place.
- Validating data integrity – The probability of unwanted data disclosure drops dramatically if the data are encrypted. When you encrypt data that reside in the cloud, the cloud provider may not be able to assert data integrity since they probably do not have the encryption key. If data are corrupted, how can one determine the cause?
- Identifying and controlling personnel with access – The data owner is responsible for determining who should access data, when and under what conditions. Cloud providers should offer contractual language that warrants the denial of access to data to its personnel and other customers. All cloud provider employees with access to the physical storage location should be bonded and insured.

Security Guidance for Critical Areas of Focus in Cloud Computing

- Terms of service may not adequately address control and disclosure of sensitive information such as trade secrets, etc. – This should be quite straightforward and should not be ignored. Based on the risk tolerance of the cloud customer, the cloud customer should determine the value of the data and potentially include that as part of a service penalty, along with associated costs, in the event of a breach. Individual negotiations should determine what parts, if any, of this are included. Customers lacking negotiating leverage should explore insurance and other risk transference strategies to recover full breach costs.
- Privacy requirements – After the user determines what privacy and disclosure requirements exist; steps must be taken to ensure that the cloud provider can support those requirements. In particular, some regulations and standards require that unless the cloud provider is designated as an affiliate or partner, the data cannot be made available to them under any circumstances.
- Backup and recovery – Many cloud providers offer data backup functions that allow data to be commingled with data from other cloud customers. Customers should perform tests often to validate the continued logical segregation of data.
- Data Retention and Destruction – Data retention and destruction schedules are the responsibility of the data owner. Data retention practices should be relatively straightforward and easily verified. Data destruction processes can be much more convoluted in the cloud. Many legacy data destruction practices depend on multiple overwrites or physical destruction of the media. Sharing resources in the cloud means that the customer cannot always confirm that media used for their data have not been reused for other purposes before all traces of deleted data are gone. This includes all slack in the data structures and on the media.
- Responding to hold orders, subpoenas, discovery requirements and other compelled disclosures – When a customer is compelled to disclose information, contamination must not exist. Not only does the customer need to ensure that all data are intact and disclosed properly; the customer must ensure that no other data are affected. More importantly, when a cloud provider has custody of customer data, the provider may not be required to and may not inform the customer that the provider is compelled to disclose the customer's data.
- Cross-boundary issues – This is an issue that will grow in complexity as cloud computing is more widely used. Many jurisdictions apply their own privacy, retention and disclosure laws on any data that physically reside within their jurisdiction. One hypothetical example could include having data that originated in Jurisdiction A be stored within the boundaries of Jurisdiction B. If jurisdiction A had no specific code or regulation controlling use of the data, the organization using the data would create their business plans according to that. If Jurisdiction B had regulations requiring that data of that sort be classified and not leave the boundaries of Jurisdiction B without government approval, then as soon as those data entered the boundaries of Jurisdiction B, the organization would be compelled to follow those regulations even if they never knew in advance that the data were to reside there. [*see Domain 3: Legal, Cross Border Transfers*]

Security Guidance for Critical Areas of Focus in Cloud Computing

- Viability and continued existence of cloud providers – Perhaps the most uncomfortable portion of dealing with service providers of any sort is addressing how to act should the service provider no longer exist. Similar to the issues concerning availability during hardware failures and data destruction, the confirmed existence of data housed exclusively with one provider comes into question. Customers should consider having an alternate location that is out-of-band from the cloud provider that can provide the needed services and data. *[Provider viability is also discussed in see Domain 3: Legal, Cross Border Transfers and Domain 7: Portability and Interoperability]*

Domain 7: Portability and Interoperability

Contributors: *John Viega*
George Reese

Problem Statement

Businesses using the cloud should be prepared for the worst. Yes, large cloud providers can make it easier to handle general availability issues well. But what happens when the cloud provider isn't good enough? What are the considerations for moving applications across clouds? Or should businesses perhaps consider building a cross-cloud solution in the first place?

Issues and Guidance

Large cloud providers often can offer geographic redundancy in the cloud, hopefully enabling high availability on a single provider. Nonetheless, it's highly advisable to do basic business continuity planning, to help minimize the damage should a worst-case scenario transpire. For example, some companies may someday suddenly find themselves with an urgent need to switch cloud providers, for one of many reasons, including:

- An unacceptable increase in cost at contract renewal time.
- The provider suddenly goes out of business.
- The provider suddenly discontinues one or more of the cloud services being used by the customer, without migration plans.
- The provider's service quality is degraded for any number of reasons.

While larger providers may not be much of a risk for going out of business without first attempting to restructure, it's still good to be prepared for the unexpected. Even a large, profitable company can increase their prices, or may decide that, for whatever reason, the economics of its cloud business are not worth the cost of keeping the services running. Some simple architectural considerations can help minimize the damage should these kinds of scenarios occur.

First, the ability to recover from a disaster of this nature depends on the type of cloud service:

Software as a Service (SaaS), where a cloud provider owns the application and hosts a customer's data. In such a scenario, the primary considerations are:

- Companies may need to migrate their data to a new application. To that end, they should make sure that they have easy access to their data in a format that is documented, so that they can process it, should the need arise.
- Companies should keep regular backups of their data, outside the service provider, in order to avoid losing data altogether, such as through vendor incompetence, a sudden bankruptcy or a dispute with the vendor.
- Companies should consider best-of-breed providers—those whose competitors have existing capabilities to migrate that provider's data, should the need arise.

Security Guidance for Critical Areas of Focus in Cloud Computing

- If a migration does need to take place, there will generally be additional expense in training people on a new system.

Infrastructure as a Service (IaaS). When using this model, the primary considerations are:

- Ensure that applications get deployed on top of the virtual machine image at runtime, limiting cloud-specific items to an abstraction layer in the machine image. This will minimize the level of effort should a need to switch providers arise.
- Backups should be kept in a cloud-independent format, and independent of the machine image.
- Copies of backups should be moved out of the cloud regularly. Note that cloud providers typically do not provide services to do this.

Platform as a Service (PaaS), where companies can build their own applications on top of a specific execution platform that is hosted in the cloud. Here, the considerations are:

- If an application development architecture is not employed to create an abstraction layer around the PaaS provider, migration can easily require significant re-writes, perhaps including changes in programming languages.
- It is more difficult to predict whether the provider will be able to meet a company's ongoing needs, because there of limited transparency into the operations, and reduced control over the execution environment. PaaS providers in general, as of this writing, have limited disclosure about security practices and policies. Customers of PaaS providers cannot assume that the development platform is lacking significant vulnerabilities that would expose their data to other applications.
- Data also needs to be backed up off-cloud.

We recommend that companies scrutinize their PaaS provider carefully when considering this as part of their IT infrastructure to ascertain level of provider transparency into security practices and policies. Also, careful application development techniques should be followed to minimize potential lock-in for the customer. The PaaS provider marketplace is evolving quickly and there are a variety of development environments available, which will vary greatly in their adherence to standards and compatibility with other PaaS providers. The onus is on the customer to have portability as a key design goal and an architecture that supports the necessary abstraction layers to make this goal achievable.

Cross-cloud applications

Sometimes there may be reasons to run an application across multiple clouds, when various parts of the application are hosted in different cloud environments. For instance, it may be useful to embed a 3rd-party SaaS offering into a larger application, where the rest of the application uses hosted infrastructure. But, note that, in such a scenario, the reliability of the system will probably only be as good as the weakest link. For instance, if the SaaS provider is inaccessible, it may render the entire application useless.

Another cross-cloud strategy is to use multiple providers to implement redundancy, by building applications that span cloud providers. While we are advocates of building applications that can be moved easily across cloud providers, we see no benefit in cross-provider redundancy. There

Security Guidance for Critical Areas of Focus in Cloud Computing

are several competent providers that already offer geographic redundancy, where the provider can automatically manage availability problems by starting up machine instances in foreign geographies. In a cross-provider scenario, such management becomes the responsibility of the company using cloud services.

Conclusion

We can certainly hope that cloud providers will provide high levels of service at a reasonable price, for the long haul. Nonetheless, it is valuable to perform business continuity planning to hedge against the worst case scenarios. If such planning is done early in the life of a project, it should not cost much to gain assurance of these issues. Primarily, one should make sure that data is easily portable between providers by abstracting the differences. And, in the infrastructure model, one should ensure that the application itself is cloud agnostic.

Section III. Operating in the Cloud

Domain 8: Traditional Security, Business Continuity and Disaster Recovery

Contributors: Jeff Spivey, Ken Fauth, Michael Johnson, Todd Barbee, Pam Fusco, Dave Tyson, Mark Leary, Ben Rothke and Dave Morrow.

Problem Statement

Risk is a component of opportunity and is very much connected to innovation, discovery and growth. Our discussion on risk in the cloud will thus be focused on opportunity rather than the downside of risk.

Cloud computing is new technology and subject to history's lessons of value gained through trust and innovation. It is our hope that the following discussion and examination of possible risks will increase dialogue and debate on the overwhelming demand for better enterprise risk management models changing the paradigm to all risks, not just security risks.

These models will be holistic, hyperadaptive and nimble as they manage emerging technology risks on the backdrop of traditional risks. Traditional security, business continuity and disaster recovery are evolving, while in parallel the Cloud continues to grow significantly. These simultaneous changes create the immediate need for an ontology, standard models and processes, collaboration among stakeholders, open and frequent discussion of the cloud providers and customers.

As an example, applying traditional physical security models from an external perspective looking in only continues the failure of prior management of security risks. Our mandate is to aggressively anticipate and respond to fast paced, new technology risks and soon-to-be new solutions.

Our challenge is to collaborate on risk identification, recognize interdependencies, integrate and leverage resources in a dynamic and forceful way. Even with Cloud Services, the need for traditional security, business continuity and disaster recovery is a fundamental protection requirement and it will not go away. Cloud Services and its accompanying infrastructure will assist in some ways to diminish certain security issues, but it may increase others and will never eliminate the need for security. While major shifts in business and technology continue, traditional security principles should remain.

Issues

- The silo approach to the insider threat continues to be one of the largest vulnerabilities. The rogue vendor, employee or visitor with current and in-depth knowledge of building or system vulnerabilities and possible secretive exploitation over time can be difficult to detect.
- Interdependencies of critical infrastructure (power, water, etc.), geography (earthquake, hurricanes, etc.), will vary significantly at cloud provider facilities.
- Incomplete and evolving convergence of physical security and IT is creating new and unknown vulnerabilities. Surveillance, alarms, telecommunications, and access control

Security Guidance for Critical Areas of Focus in Cloud Computing

can now all reside on the IT backbone; a convergence that introduce new risks and single points of failure that can have dramatic consequences .

- Outsourcing of some or all functions creates varied standards, processes, accountability, stability, etc. Providers will have a variety of customers who, in turn, will have varied regulatory mandates and standards, each of which must be complied with. While outsourcing various functions can significantly reduce infrastructure and staff costs, it can also cause confusion between provider and customer regarding responsibilities, accountability and redress for failure to meet the required standards.
- Access control should be a united perimeter for the physical and logical access. The silo approach found most prevalent ignores information of physical presence before a logical breach of any site, whether it is accomplished in an overt or covert, physical or cyber manner may lead to the potential compromise of someone's information, and will at least force expenditure of investigative time for a damage assessment, down time, repair or replacement of equipment or structure and impact the reputation of the cloud provider, the cloud's customer, and the customer's customer.
- Due to the rapid build out of cloud computing, weak, missing or overly specific security policies will create ineffective management of traditional security risks along with business continuity and disaster recovery. Additionally, due to incomplete policy/procedures or poor training, staff, vendors, visitors, etc. do not understand their individual roles in the overall security and continuity plans. When an event occurs, lack of preparation makes the effects of the event or breach more severe.
- Outsourcing cloud center security responsibilities can create inconsistency in security application of policy and procedures.

Guidance

- Providers should:
 - Examine the highest level of requirements by potential clients and provide that as the baseline level of security for their enterprise offerings.
 - Confirm the traditional security, business continuity, and disaster recovery issues in the various international compliance standards, compliance with international and country law, industry standards
 - Contractually specify and agree upon which party is responsible for ensuring compliance with relevant parts of the standards in question. When responsibility falls through the cracks, security incidents occur.
 - Use an employee screening process with periodic re-checks and substance abuse checking with penalties to act as a deterrent. Integrate Human Resource, Security, etc. policies to assure the early detection and process for dealing with the rogue employees, vendors and others creating the insider threat.
 - Compartmentalize jobs and access to lessen the chances of revealing information that can assist in unwanted activities. Compartmentalization also hinders

Security Guidance for Critical Areas of Focus in Cloud Computing

unauthorized activity by limiting knowledge and access to areas to only those having a *need to know* what is present or occurring. Access should be only to those who require it to do their jobs and not granted solely due to status or position in the organization

- Institute surveillance of the both the physical and cyber sides of the data center(s) serving the cloud
 - Implement physical access controls to assure that everyone entering the cloud center is authorized. Have strong visitor procedures with enforcement of all security rules. Regular review of access control systems is vital to maintain appropriate levels of access and promptly reflect changes in personnel due to hiring and departures.
 - If outsourcing security responsibilities at a cloud center, assure contractual third party adherence to policies/procedures and their involvement in your Enterprise Risk Management (ERM) model of risk management. Additionally, a monitoring program should be established, including an early warning process if the outsourcer falls out of compliance.
 - Work to develop a security culture throughout the company to ensure employees continuously understand their duties and responsibilities to protect company and customer data
 - Perform frequent risk assessments by both internal and external staff. Risk assessments performed by internal staff become somewhat routine over time. The security of cloud centers is a continually evolving process with new technologies and new vulnerabilities to be addressed. The outside view is always a good idea, especially in fast moving and new services.
 - Establish processes to frequently review potential infrastructure issues including: fire protection; working back-up generators or battery power in the event of an electrical outage, HVAC testing and other essential building operations.
 - Invest in traditional physical security applications include: fences, locks, controlled entry/exit point and other emergency points, interior movement restrictions and visual monitoring. The physical restraining of movement denies access to a targeted area. Creating physical barriers (bollards, fences, solid walls, floors and ceilings, locks, doors, cameras) to both individuals and vehicles protects the area from destructive devices or actions.
 - Respond immediately to a breach or intruder. Make sure everyone involved in the response knows their roles and the overall response process. To physically respond to an alert can include apprehending an intruder, providing medical help, saving individuals or equipment/materials.
- Customers of cloud computing should:
 - Clearly understand what data they are putting into the cloud, the security requirements for that data, and make sure the provider is also aware of the nature of the data. Customers should include the security requirements of their data and that of their customers as criteria in conducting their enterprise risk assessments and contracts.

Security Guidance for Critical Areas of Focus in Cloud Computing

- If possible, inspect the vendor's physical and personnel security measures, disaster recovery and business continuity plans. Obtain specific written commitments from the vendor on recovery objectives for critical capabilities.
 - Require cloud vendors to take contractual responsibility for security. Many current Terms of Service are written to absolve the vendor of any responsibility for the security or availability of data they house
 - Identify interdependencies in the vendor's infrastructure to include: Site selection risk include natural (earthquake fault line, tornado, floods), man-made (chemical plant, airport, adjacent high risk targets) and stand-off distances (clear zones). Type of building, single or multi-tenant use, and location within structure. The securing of feeder lines for redundancy of electricity, water for cooling and communication lines
 - Integrate cloud computing security, business continuity and disaster recovery into the customer's own policy and procedures. These would reflect the holistic and comprehensive Enterprise Risk Management policies for the entire company.
- Regulators and standards organizations should:
 - Update and modify their standards to account for the differences in architecture and operations of cloud environments
 - Build flexibility into standards and regulatory structures to allow for changes in technology and business practices. Specifying technologies to counter specific risks can become a hindrance when the environment changes.

Domain 9: Data Center Operations

Contributors: *Jim Reavis*
Pam Fusco
Josh Zachry

Problem Statement

The anticipated explosion in cloud computing providers is accompanied by similar growth in data centers to fuel the on-demand computing services. Cloud providers include both large, well known technology bellwethers, as well as thousands of cloud startups and emerging growth companies. These many cloud providers use a wide variety of enabling technology: off-the-shelf virtualization, open source software and many proprietary technologies to optimize resource sharing. Some may be Software as a Service (SaaS) providers that are vertically integrated with Infrastructure as a Service (IaaS) or Platform as a Service (PaaS) providers. This drive to create maximum efficiencies in cloud computing may lead to new and unproven cloud data center IT infrastructure and uneven data center management practices, which need to be vetted by the cloud customer.

Issues

To the casual observer, it may seem that the future of cloud computing lies in the hands of a small number of large and well-known technology companies. This could not be further from the truth, as history shows us that IT innovation comes from the most unexpected corners, and cloud providers will be no different. Whether you are engaging with a cloud provider with a highly recognized brand name or a new market entrant, a savvy customer needs to be aware of the cloud-centric data center operations issues relevant to the delivery of their service.

A key underlying premise of the economic model driving cloud computing is that sharing resources creates efficiencies. Internal data centers must be architected for periods of peak business activity, which could vary greatly from average activity and leave the data center idle a great majority of the time. The economic model dictates that cloud providers will seek to utilize resources, both human and technological to the maximum degree possible to gain a competitive advantage and maximize operating profit margins.

Maximizing the sharing of resources inside the cloud data center does not absolutely run counter to the need to compartmentalize and segment those resources for the purpose of security assurance. However, it does not guarantee compartmentalization either, and it remains a significant open question and recurrent theme of many domains within this whitepaper. As is obvious from our Domain 1 Cloud Computing Architectural Framework, broad definitions of cloud computing are necessary to depict the many different approaches to building a cloud service. There is not a standard cloud data center operating system or set of APIs. In these early stages, this is a good thing from the perspective of fostering innovation. In many cases, providers are adapting traditional technologies to the cloud purpose, but at some point a wave of innovation in best practices, storage, networking, security, application development and many other technologies will sweep away the incumbents. Until that happens, a cloud customer must be willing to invest a greater amount of time to understand how their cloud provider is managing their data center.

Guidance

It is critical to understand how your cloud provider has implemented Domain 1's "Five Principal Characteristics of Cloud Computing", and how that technology architecture and infrastructure impacts their ability to meet service level agreements. Your provider's technology architecture could be a combination of IT products and other cloud services, IaaS storage, for example. The OpenCrowd Taxonomy in Domain 1 is a useful reference for some of these components, but you should also dig deeper to understand the cloud provider's selections on products as varied as switches to firewalls to uninterruptible power supplies. Understanding their selections can allow you to research vulnerabilities, mean time between failure (MTBF), service availability, and assemble that data into a system-wide assessment of risks.

How does the cloud provider's compartmentalization work? While the technology architecture and infrastructure of cloud providers will differ, they must all be able to demonstrate comprehensive compartmentalization of systems, networks, management, provisioning and personnel. It is important to assure that the controls segregating each layer of the infrastructure are harmonized and do not cancel each other out. For example, is the storage compartmentalization easily bypassed by management tools or poor key management?

Understand how resource democratization and dynamism occurs within your cloud provider to best predict the likelihood of system availability and performance during your business fluctuations. The theory of cloud computing to date exceeds its practice, and many customers have incorrect assumptions about the automation involved. As you reach capacity in your provisioned resource, will the additional required resources be made available seamlessly?

Related to the above point, if feasible, discover the cloud provider's other clients to assess the impact their business fluctuations may have on your customer experience with the cloud provider. Your data and applications are at some data center(s) and you have neighbors. It is common sense to understand who your neighbors are, and if they have the potential to consume inordinate amounts of resources and impact your performance metrics. Also, is your neighbor a high profile target likely to get attacked?

For IaaS and PaaS, clearly understand the cloud provider's patch management policies and procedures and how this impacts the applications you have developed for these environments. You should certainly expect more restrictions in cloud provider flexibility than with your own systems, however there are many important details related to notification, rollbacks, regression testing and services windows that may be critical to you. When patching and potentially rebooting a shared resource, do all customers of that resource get impacted? This understanding should be reflected in contact language.

As in Domain 8, review business continuity and disaster recovery plans from an IT perspective. The cloud provider's technology architecture may use new and unproven methods for failover. The customer's own business continuity plans should also be integrated with and address the limitations of the cloud provider's plan.

Logging practices of cloud providers need to be carefully scrutinized. What are the standard logging practices for the cloud provider? What is stored? How long is it stored? How is it segmented? How are different roles, such as administrative, user and audit, supported by the logging system?

Security Guidance for Critical Areas of Focus in Cloud Computing

Customers should test cloud provider's customer service function regularly to determine their level of mastery in supporting the services.

One good indicator for understanding the cloud provider's commitment to the operational quality of their data center is the presence of staging facilities to pilot changes and new products, both on the part of the provider as well as for use by the customer. Asking if the provider has a laboratory and getting a walkthrough or tangible quantification of lab capabilities is recommended.

Domain 10: Incident Response, Notification, and Remediation

Contributor: *Liam Lynch*

Problem Statement

Cloud computing has quickly become a way to not only manage agile development of Internet applications but also an inexpensive way to host applications as well. Across the fabric of cloud instances there are subtle nuances of what a cloud means but for the purposes of incident response the focus is fairly simple. Applications provide value by managing data and the data is owned by users of the application and therefore much of that data can fall into a standard data classification: confidential, restricted, customer non public, etc. There are other types of incidents that can affect an application in the cloud, which relate to data access, but stand alone as potentially serious for a user, and they are the OWASP Top 10 security vulnerabilities.

The problem to focus on will be that applications deployed to cloud fabrics will not always be designed with security and that will mean that the results will be vulnerable applications, which are going to cause incidents.

Privacy experts have much to contribute to cloud security and they also have requirements in the incident response area regarding notification and remediation. Large cloud providers' privacy councils may have further requirements regarding data access by employees to application data that is not governed by user agreements and privacy policy. Application data not managed by a cloud provider's own application should have access restricted to employees in the case where the infrastructure is owned and maintained by the cloud provider.

In order to remediate privacy concerns further, which encompass prevention of internal disclosure, emerging best practices within cloud providers suggest that all data falling into privacy data classification be automatically encrypted in foreign application data stores. In this way, accidental or malicious access to private data can be avoided.

Lastly, the complexities of a large scale cloud provider providing SaaS, PaaS and IaaS capabilities create significant incident response issues. The provider is faced with hosting hundreds of thousands of application instances, some of which the provider maintains, some of which are customer maintained. From an incident monitoring perspective, foreign applications widen the operations aspect in the security operations center (SOC) functionality. A SOC today monitors for issues from firewalls and intrusion detection platforms, however open cloud application deployment can pose issues if all hundreds of thousands of applications have to reside behind a firewall in order to be monitored.

Issues and Guidance

An incident response strategy for cloud providers will have to be responsible for identification and notification, and highlight options for remediation of criminal access to application data. In addition, application data management and access has different meaning and regulatory requirements depending on where the user of the application lives (country). An incident may exist for a user in Germany where a user in another country may not. This complication makes incident identification a non trivial exercise.

Security Guidance for Critical Areas of Focus in Cloud Computing

There are ways to lessen the complexity however. To reduce the possibility of a data breach for example, data that falls into a common data classification as private or non public can be encrypted while it is stored and only available decrypted to the application that has been created to manage the data. One type of incident however can make encrypted data remain open to a breach. If an application has an SQL injection vulnerability, then the application will access its data and the system will decrypt the data and then the application will serve the data back to the requester. In the SQL injection case the requester is a criminal harvesting data. This case is an incident.

To further complicate incident discovery the hosting entity will have employees or contractors to manage the systems the applications run on and data is stored in. These personnel should not be able to access data owned by other entities and attempts to break into the system and or successfully obtain data are also an incident. The possibility of insider threat is an incident response responsibility as well.

The Cloud Computing Community¹ incident database has captured the following types of incidents over that past 18 months:

- Malware infection.
- Data breach.
- Man-In-The-Middle discovery.
- Session hijacking.
- User impersonation.

With the onset of cloud computing to a larger scale and availability to a larger audience the types and occurrences of incidents will increase. A criminal who commits fraud, for example, will simply follow the money and look for weak applications.

There is a need and a market for ways to detect that an application in a cloud is vulnerable and under attack. Conventionally, and in addition to firewalls and Intrusion Detection Systems, application level firewalls or hybrid cloud configurations can be used to detect vulnerabilities and those events can be sent to the operations center (SOC) of the cloud host. Application level firewalls can also be used in front of the database management systems to get a full path vulnerability detection framework in place. In the case of a secure-sensitive cloud application, transactions to the applications can be processed by high speed proxies before traffic reaches the application.

The application framework can also provide components that provide protection against OWASP vulnerabilities. At sophisticated cloud providers, there are frameworks to protect against cross site scripting, cross site request forgery, SQL injection, and authentication weaknesses amongst others. There is no guarantee however that every application developed in a cloud would use the mitigation framework unless the frameworks were built into the platform.

On the users' side there will be the effects of malware and robotic application use as the rise of infections reach critical mass. We need to be able to detect infected users and robotic application use for malicious consequences. The incident response process should be able to detect and notify an application owner about these situations. Targeted malware is of a specific concern as instances of this type of malware are rising constantly.

Security Guidance for Critical Areas of Focus in Cloud Computing

Application level logging is a detection tool used widely in large enterprises and can be used to analyze traffic and report incidents as well. Given that cloud applications will be developed by a wide variety of software developers, application logging will be inconsistent unless a standard application logging framework exists and logs transactions independent of the application developer. The detection of issues can be uniformly applied to all applications and then plugged into an incident process at the SOC. Alternately, cloud security providers, while they proxy transactions, can be detecting issues across the malicious spectrum and those services can plug into the incident process at the SOC.

Notification, however, is not as simple as current incident processes are today within a SOC or other types of monitoring tools. The assumption is that what is being monitored belongs to a single governance and incident response model and that notification and remediation is handled by one group even though the application space may be varied. In the case of hundreds of thousands of application owners, the notification process is more complex and likely that notification won't work the same way. A registry of all application owners by their application interface (URL, SOA service, etc) and who to contact in case of an incident has to be built, which may not scale to the current SOC/SIM models in use today.

With the eventuality that clouds will host hundreds of thousands of applications and that those applications store and manage data for users, the scope of regulatory control will also increase. If there were a data breach from an application in the cloud and the FTC² in the US were to investigate, the hosting provider would have to know which of the applications was breached, who is responsible for that application, and if another application was involved. Take for example a Platform as a Service (PaaS) provider that also provides access to consumer data as part of the offering. A theoretical breach may have been perpetrated against a 3rd party cloud application but most of their data may have come from the PaaS provider systems through one of its SOA services. Who will the FTC investigate? Best practices from a privacy perspective would dictate that differing applications' data be kept separate and unavailable to other applications including the PaaS provider. In that way a breach will focus on one entity and not several. Unfortunately, the data line will blur quickly and an incident process will have to be very specific about what will happen and also know the dependencies so the incident notification can be delivered to all parties in a line of data custody.

Finally, remediation of incidents will be dependent upon what specific data can be surfaced about the incident and the ability to plug that data into data management tools for prosecution if required. Application owners will have to be a part of the process so that remediation can take place, and options other than shutting an application down will be what an application owner would want. The reality of an incident however is that an application shutdown is normally the first act taken when a breach or other malicious use is detected, and final remediation is completed when the issue is fully diagnosed.

References

- ¹ Cloud Computing Community Incident Database: http://wiki.cloudcommunity.org/wiki/CloudComputing:Incidents_Database
- ² FTC questions cloud-computing security. Cnet news. http://news.cnet.com/8301-13578_3-10198577-38.html

Domain 11: Application Security

Contributors: *Dennis Hurst*
Scott Matsumoto

Problem Statement

Application software running on or being developed for cloud computing platforms presents different security challenges depending on the delivery model of the that particular platform. The first question a CISO must answer is whether it's appropriate to migrate or design an application to run on a cloud computing platform and the second question is what type of cloud platform is most appropriate. For application security, the answer to each of these questions has two: what security controls must the application provide over and above the controls inherent in the cloud platform and how must an enterprise's secure development lifecycle change to accommodate cloud computing? Both answers must be continually re-evaluated as the application is maintained and enhanced over time.

Issues and Guidance

Many enterprise security programs have an application security program to address the unique security risks in this realm. Designing and building applications targeted for deployment on a cloud platform will require that existing application security programs re-evaluate current practices and standards. The changes to an enterprise's current application security practices and standards need to address the subtle differences of the cloud platforms. Some of these differences come from the multi-tenant environment of cloud platforms, the lack of direct control over the environment, and access to data by the cloud platform vendor. These differences must be addressed by an application through a set of application level controls and through the service agreement with the cloud vendor.

For a CISO, addressing cloud application security is a function of extending an enterprise's existing application security policy, standards and tools to a cloud platform. The level and nature of the necessary extensions depend on the delivery model of the cloud services defined by the Domain1 Cloud Computing Architectural Framework. Each of the main delivery models and its impact on application security is described below.

Infrastructure as a Service (IaaS) Delivery Model Application Security

In an Infrastructure as a Service (IaaS) cloud platform, the cloud vendor provides a set of virtualized components such as virtual machines, raw storage and other components that can be used to construct and run an application. The most basic component is a virtual machine and the virtual OS where the application resides. See figure 1.

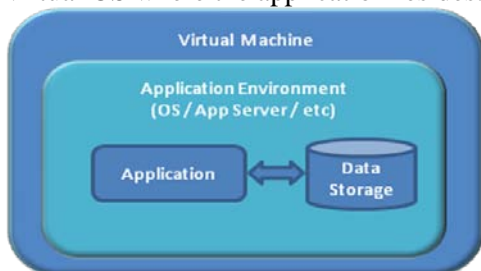


Figure 1 - Virtual Machine of an IaaS

In IaaS environments, the local data storage is not persisted across machine restarts, so most applications use some form of external, persistent storage. The IaaS environments provide additional components for persistent storage, but that storage is always remote. See figure 2.




Figure 2 - Cloud-based persistent storage

IaaS Impact on Application Security Architecture

The architecture for IaaS hosted applications has a high resemblance to legacy web applications, namely a web-based, n-Tier distributed architecture. For distributed applications, running in an enterprise there are many controls put in place to secure the host and the network connecting the distributed hosts. Comparable controls do not exist by default in an IaaS platform and must be added through configuration or as application-level controls.

Trusting the Virtual Machine Image

IaaS providers make a vast number of virtual machine images available to their customers. Some of these virtual machine images are provided by the IaaS provider itself, but some are provided by other customers. When a virtual image from the IaaS provider is used it should undergo the same level of security verification and hardening for hosts within the enterprise. The best alternative is to provide one's own image that conforms to the same security policies as internal trusted hosts. An alternative is to use virtual images from a trusted third party. One example of a trusted third party is a service provider that provides value added services above the infrastructure components provided by the IaaS provider.

Hardening Hosts

IaaS platforms provide the ability to block and filter traffic based on IP address and port, but these facilities are not equivalent to the network security controls in most enterprises. Hosts running within an IaaS infrastructure are akin to hosts running in the DMZ of your enterprise's network. All of the same precautions used to harden hosts running in the DMZ should be applied to the virtual images.

A best practice for cloud-based applications is to build custom operating system and application platform images that have only the capabilities necessary to support the application stack. Limiting the capabilities of the underlying application stack not only limits the overall attack surface of the host, but also greatly reduces the number of patches needed to keep that application stack secure.

Securing Inter-host Communication

Most enterprise applications do not worry about security communication between hosts of a distributed application, so long as traffic does not traverse an untrusted network. A cloud-based

Security Guidance for Critical Areas of Focus in Cloud Computing

application must design in explicit controls to prevent the disclosure of sensitive information between hosts. The application must take on the responsibility for securing the communication in a cloud-based application, because the hosts are running in a shared infrastructure with other companies. Also, the administrators that maintain the data center running the hosts and network should not be afforded the same level of trust as administrators of an internal data center.

Securing such communication depends on the type of communication. For synchronous communication, such as point-to-point network connections, channel level security is sufficient. For asynchronous communication such as using a message queue-based mechanism, message-based security is needed to protect the sensitive information while the data is in transit.

Managing Application Keys

IaaS platforms use a “secret key” to identify a valid account. The account key must be passed on all of the calls to make use of the services provided by the IaaS provider, such as the calls to connect and communicate between application nodes. Most application security programs have standards and best practices for handling key material, but these standards and practices will need some modification to for application keys.

Additional Requirements for Handling of Sensitive Information

Applications running on an IaaS platform must ensure that sensitive information not leak during processing. All of the precautions for handling sensitive information for enterprise applications apply to IaaS hosted applications. Additional filtering and masking is needed for handling operation and exception logging; especially when debugging information is logged because the storage for this information is shared and managed by an outside party.

IaaS Platform Impact on the Software Development Lifecycle

A fundamental aspect of application security is how security is integrated into the development lifecycle. This concept has been articulated in many different formats such as the Secure Development Lifecycle (Microsoft), various sections of the Payment Card Industry (PCI) Data Security Standard and other sources. All of the security issues related to application security still apply when applications move to a cloud platform, however a number of new issues arise. One key issue occurs when the development life cycle crosses a trust boundary from an internal or “trusted” environment into the cloud.

Applications running on an IaaS platform have a different trust relationship between the development environment and the deployment environment than traditional enterprise applications. In a traditional enterprise application, all of the environments are within the enterprise as is shown in figure 3. Within an enterprise, this trust is created by the secure host and secure networks provided as part of the enterprise’s computing infrastructure.




Figure 3 – SDLC Trust Model for Internal Application

When an application runs on an IaaS platform, the application's production environment and some parts of the test environment run with different trust assumptions than the development environment. Figure 4 shows the different environments for development, test and production.




Figure 4 - SDLC Trust Model for IaaS Hosted Application

When an application is moved to an external environment, trust must be codified between the customer and the vendor through the Service Level Agreement (SLA) provided by the vendor. Application security must be represented as a clearly articulated set of actions and guarantees within the SLA, such as providing documentation of security measures taken by the vendor and allowing for reasonable security testing by the customer related to ongoing activities such as logging, audit reports and other activities.

Regaining issues of trust between internal and external environments is similar to the problem of operating an application at a Managed Service Provider (MSP). The difference between operating an application in the cloud is the limited duration of persistent data on the cloud resources versus the physical resources at an MSP.

Platform as a Service (PaaS) Delivery Model Application Security

Platform as a Service (PaaS) providers deliver not only the runtime environment for the application, but also an integrated application stack. A PaaS provides additional application building blocks. These additional application building blocks layer on top of services provided by IaaS platforms. For example, an IaaS provides a message queue for asynchronous messaging whereas a PaaS provides an Enterprise Service Bus (ESB) that provides both the asynchronous messaging as well as services such as message routing.

The Domain 1 Cloud Reference Model describes these initial capabilities as the Integration and Middleware layer. The relevant layers are shown the excerpt of the Cloud Reference Model in figure 5.




Figure 5 - Cloud Reference Model: Application Capabilities Provided by PaaS

PaaS Impact on Application Security Architecture

PaaS platforms also provide the programming environment to access and utilize the additional application building blocks. This programming environment has a visible impact on the application architecture. One such impact is that of the constraints on what services the application can request of the operating system. For example, a PaaS environment may limit access to well defined parts of the file system. These restrictions are put in place by the PaaS to allow the PaaS to better manage its multi-tenant environment.

Even though the PaaS platform's application building blocks are similar to their enterprise counterparts, for example both have ESBs, the multi-tenant nature of the cloud computing environment means that the application's assumption about trust must be re-evaluated. Just like IaaS environments where the network is multi-tenanted, an ESB within a PaaS environment will be shared. Securing the messages on the ESB becomes the responsibility of the application because controls such as segmenting ESBs based on data classification are not available in PaaS environments.

Managing Application Keys

Just as in IaaS platforms, PaaS platforms require an application key on all API calls to the platform itself and for calls to services within the PaaS environment from the hosted application. The application key must be maintained and secured along with all other credentials required by the application.

Additional Requirements for Handling of Sensitive Information

PaaS platforms have the same requirements for application level handling of sensitive information as the ones from IaaS platforms.

PaaS Platform Impact on the Software Development Lifecycle

Developing applications for a PaaS platform can add risk associated with the software development lifecycle. This risk comes from the lack of secure design and coding patterns, technology specific application security standards and application security assurance tools for software built on this platform. These cornerstones of a secure development lifecycle must be updated for the specific PaaS environment.

Each enterprise looking to extend its current secure development lifecycle will have to develop this knowledge and tools. Web-based, n-Tier applications have a rich body of knowledge about common types of vulnerabilities and their mitigation. Similar knowledge for PaaS environments must still be developed.

Software as a Service (SaaS) Delivery Model Application Security

Software as a Service (SaaS) provides the same management of infrastructure and programming environment and layers in specific application capabilities. The application's capabilities provide end-user functions as well as become part of the programming platform. The application's capabilities can be extended by adding custom code extensions. External applications can exchange data through the APIs the SaaS platform provides. Figure 6 shows these integration points relative to appropriate layers of the Domain 1 Cloud Reference Model.

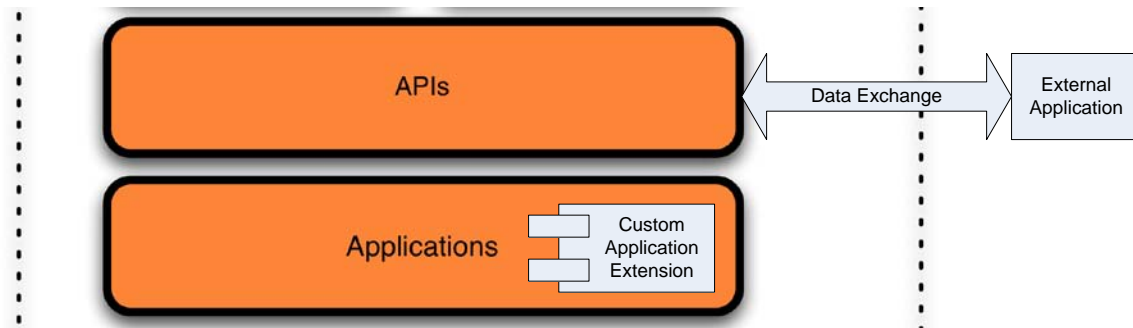


Figure 6 - SaaS platform customization

SaaS Impact on Application Security Architecture

SaaS platforms inherit all of the same security architecture concerns and mitigations as PaaS and IaaS environments. The application security architecture for any custom code extensions is the same as for the application itself. Data exchanged through the SaaS platform's external APIs are subject to existing security policy and standards for any type of external data exchange.

SaaS Platform Impact on the Software Development Lifecycle

Like PaaS platforms, SaaS platforms represent a new programming environment and existing secure design and coding patterns, technology-specific standards and application security assurance tools must be developed and adopted by the organization. In addition to these concerns for the software development lifecycle (SDLC) within the organization, an enterprise must be equally concerned about the SDLC of the SaaS platform vendor. This concern is true for all of the other cloud delivery models, but it is especially true since the application is now shared between the SaaS vendor and the enterprise.




Figure 7 - SDLC Trust Boundaries with SaaS Vendor

An enterprise must have a way to trust that the vendor's development lifecycle is as secure as its own. Appropriate due diligence should be given to ensuring in the SLA, that the maturity of SaaS vendor's SDLC through either internal or external verification (audit).

Final Thoughts: How will the malicious actors react?

As application development practices and security hardening evolve within the different cloud delivery models, it is useful to consider what the reaction will be from malicious actors. To the extent that major application components are not exposed via SOA or in the user presentation, the hacker will be unable to examine and attempt to reverse engineer these components. We should be able to predict that the malicious actor will ruthlessly examine available active code, such as JavaScript, Flash and others. They will also seek to attack infrastructure that is standardized,

Security Guidance for Critical Areas of Focus in Cloud Computing

where they can leverage a body of vulnerability research knowledge. We can also expect hackers to focus on extensive black-box testing strategies. It will be important for the application security professional to stay abreast of the latest tools and techniques hackers develop specifically to attack cloud providers.

References

Amazon Elastic Compute Cloud Developer Guide, <http://docs.amazonwebservices.com/AWSEC2/2009-03-01/DeveloperGuide/>
Amazon Simple Storage Service Developer Guide, <http://docs.amazonwebservices.com/AmazonS3/2006-03-01/>
Amazon SimpleDB Developer Guide, <http://docs.amazonwebservices.com/AmazonSimpleDB/2007-11-07/DeveloperGuide/>
Amazon Simple Queue Service Developer Guide, <http://docs.amazonwebservices.com/AWSSimpleQueueService/2008-01-01/SQSDeveloperGuide/>
Azure Services Platform, <http://msdn.microsoft.com/en-us/library/dd163896.aspx>
Windows Azure SDK, <http://msdn.microsoft.com/en-us/library/dd179367.aspx>
Python Runtime Environment, <http://code.google.com/appengine/docs/>
OWASP Top Ten Project, http://www.owasp.org/index.php/Category:OWASP_Top_Ten_Project
Force.com Web Services API Developer's Guide, <http://www.salesforce.com/us/developer/docs/api/index.htm>
The Force.com Workbook, <http://wiki.developerforce.com/index.php/Forcedotcomworkbook>
Building Security In Maturity Model, <http://www.bsi-mm.com/>

Domain 12: Encryption and Key Management

Contributors: *Jon Callas*

Problem Statement

Cloud computing changes the way we think about computing by removing the specifics of location from its resources. Cloud computing can be thought of as radical deperimeterization; similarly to the way that we think about the network itself as an abstract cloud of network links, cloud computing abstracts all computing and networking resources. However, in divorcing components from location, this creates security issues that result from this lack of any perimeter. In such a world, there is only one way to secure the computing resources: strong encryption and scalable key management.

Issues and Guidance

Strong encryption with key management is the core mechanism that cloud computing systems must use to protect data. Encrypted data is intrinsically protected; if someone has the data without its corresponding keys, they cannot use the data at all. This principle is reflected in international laws and regulations, particularly privacy and data loss requirements. Safe harbor provisions in laws and regulations consider lost encrypted data as not lost at all. The encryption provides resource protection while key management mediates access to these resources.

Encryption not only protects the data while it is in transit, but while it is at rest. The radical deperimeterization that cloud computing provides blurs the very distinction between data that is at rest or in motion. Consequently, cloud computing security requires the pervasive use of data-in-transit mechanisms as well as data-at-rest mechanisms that are used in concert with each other. In fact, unencrypted data in the cloud could be considered to be lost by its mere unprotected presence in the cloud. In any case where the application provider does not control the application back end, they must assume that all communications and all storage may be visible to arbitrary outsiders, even though that may occur only through an error in some other provider. They must encrypt all data in motion and at rest, even when other services might also be protecting that data.

Cloud customers and providers alike must guard against data loss and data breach and must exploit the advantage that encrypted data and services have in law and regulation. Customers want their providers to encrypt their data so that the customers do not have to report incidents that the providers have to the customer's customers. Likewise, the provider wants to encrypt its customers' data for the same reason. This forms a secure information supply chain since most customers are also providers of goods and services and most providers are also customers of goods and services.

Similarly, customers and providers alike have needs for data integrity systems that only encryption and key management can provide. Customers need to know that the data that they are storing in a service is not altered through error, neglect, nor malice. The providers themselves want to be able to demonstrate to their customers that the data in the cloud service is in fact the same data that the customer placed into the .

Providers also want to construct a key management scheme that keys each customer's information separate. This creates an automatic separation of customers' data from each other

Security Guidance for Critical Areas of Focus in Cloud Computing

even in the cases where all customers share a common repository. Here, encryption is a mechanism to create a data minimization system while preserving economies of scale.

The key management needs of cloud computing services allow us to bring into information technology a core principle from finance, manufacturing, and government, the principle of separation of roles. That principle, that no single entity should be able to do everything, has traditionally been hard to bring into information technology. Information is fluid and those who hold it can do anything with it, unless there are confidentiality and integrity checks on it that only encryption can provide. In a traditional information management architecture, we finesse this problem by creating trusted holders of the data. The cloud breaks this down, but simultaneously recreates the ability to use encryption for separation of roles. Using encrypted data requires both the encrypted data and the keys that encrypt that data. By separating the data and the keys, we can create a chain of separation as well as a chain of custody with two or three parties involved at each step.

The “cyber supply chain” creates a web of cloud data services that connect to each other — data here, computes there, visualization somewhere else — also creates a new threat that does not exist elsewhere, but its solution also exists through encryption. That threat comes from an external party that tries to get an organization or person’s data by requesting the cloud provider. That request can be a simple request, or it can come with discovery or subpoena. Providers can protect themselves by separating the role between data holding and data use. Customers can protect themselves with the same mechanisms.

Finally, there should be no misunderstanding of what constitutes robust encryption. When specifying the use of encryption on the part of the cloud provider contractually, it should include existing industry or government standards, as applicable.

Cloud computing provides new economies for information technologies as well as new challenges. Properly constructed, it be more secure than traditional IT infrastructure, but that construction requires strong encryption and scalable key management.

Domain 13: Identity and Access Management

Contributors: *Subra Kumaraswamy*
Jim Reavis

Problem Statement

Managing identities and leveraging directory services to solve business problems remains one of the greatest challenges facing IT today. While an enterprise can leverage several cloud computing services without integrating identity management, in the long run extending an organization's identity into the cloud is a necessary precursor towards strategic use of on-demand computing services to run your business. Supporting today's aggressive adoption by the business of an admittedly immature cloud ecosystem requires an honest assessment of an organization's readiness to conduct cloud-based Identity and Access Management (IAM), as well as understanding the capabilities of that organization's cloud computing providers.

Issues and Guidance

The litany of issues that organizations have with their identity and access management is well known: a single employee having hundreds of individual user accounts, former employees having system access long after having been terminated and users having inappropriately elevated privileges. A single large enterprise may have hundreds or even thousands of line of business software applications, most of which are not integrated into a single sign-on (SSO) solution. It is a reasonable presumption that both small businesses and large enterprises will by necessity engage with several cloud computing providers that provide the optimal combination of IT services. It therefore stands to reason that an organization should ideally be leveraging an optimized internal IAM strategy and practice into the cloud. The following represents our guidance:

The most important factor for an enterprise to have success in effectively managing identities in the cloud is the presence of a robust federated identity management capability within the organization: architecture & systems, user and access lifecycle management, audit and compliance capabilities. Viral adoption of cloud services driven by business units that don't leverage your own federated identity management infrastructure and processes risk repeating the mistakes - that caused you to implement enterprise identity management solutions in the first place. In addition, you may inherit the risk of sensitive internal identity information, inappropriately hosted at several cloud providers.

Standards support for achieving IdM federation with your cloud providers is crucial. The current capabilities of cloud computing providers to integrate with your Identity Provider (IdP) will vary widely, based on the particular provider and the class of cloud service (SaaS, IaaS, PaaS) they are providing. Do not be surprised to find that many cloud providers are immature in support of the top federation standards: SAML, WS-Federation and Liberty ID-FF federation standards. It appears as though SAML is emerging as the leading standard that enables single sign-on (SSO). As it pertains to IaaS and PaaS, you will generally need to build in this integration yourself. It is critical to ask your cloud provider about their support of federation standards early on to ensure alignment with your standards, and you should not be shy about pushing the cloud providers to broadly support industry standards for federation. Failure to do so will lead to unplanned increases in integration costs or incompatibility with your internal IAM architecture.

Security Guidance for Critical Areas of Focus in Cloud Computing

You should understand the cloud provider's support for user management processes including user provisioning, de-provisioning and overall lifecycle management of users and access in the cloud in an automated way. These factors will impact the efficiency of the overall federated identity management system and will be essential to answer audit questions.

Authentication capabilities are another important area to clarify on the part of your cloud provider ahead of time. You also need to perform due diligence to assure that the cloud provider's password policies and strong authentication capabilities meet or exceed your own policies and requirements. It is not uncommon for cloud provider to delegate authentication to your corporate identity provider using federation standard such as SAML. In that model you will have the flexibility to enforce the appropriate authentication strength (e.g. one time password) mandated by your information protection and data classification policies and standards.

Cloud customers should be aware that granular application authorization is immature at this point. Where it does exist, it is usually implemented in a proprietary fashion specific to the cloud provider. As a long term strategy, customers should be advocating for greater support of XACML-compliant entitlement management on the part of cloud providers, even if XACML has not been implemented internally. XACML provides a standardized language and method of access control and policy enforcement across all applications that enforce a common authorization standard. At the very least, CISO's should be thinking about authorization standards and avoid any temptation to customize solution based on providers capability.

The reality is that you may need to make tactical investments today to augment what a given cloud provider is capable of delivering out of the box to sufficiently integrate with your own IdM solution. Again, the better your federated strategy and support of standards is, the better chance you will have of minimizing additional investments. A good strategy towards the maturation of your own IdM in order to make it “cloud friendly” is to start enabling SSO within your own enterprise applications, for your existing user base of employees, partners and contractors. This internal SSO framework can be leveraged and the IAM processes, practice extended into the cloud.

One of the investments you may consider is an Identity as a Service solution to bridge between cloud providers or even outsource some Identity Mgt functions. Identity as a Service may abstract some of the complexity by way of unique SaaS/PaaS provider support for SSO by way of federation. For example, some providers support SAML 1.1, while others support SAML 2.0 only. By leveraging Identity as a Service, you can outsource the integration issues to the service provider while maintaining consistent directory synchronization between your enterprise and identity service provider directory. Be aware that this service represents an additional cloud provider to your enterprise and they will need to be vetted with the guidance from these 15 domains as well.

Summary

- A strong internal Federated Identity Management strategy is the foundation for cloud IAM.
- Federated standards support on the part of both the cloud customer and cloud provider is critical: SAML, WS-Federation and Liberty ID-FF federation.
- Provider's password and strong authentication capabilities should meet or exceed internal requirements.

Security Guidance for Critical Areas of Focus in Cloud Computing

- Be careful of provider portability issues when adopting application authorization schemas proprietary to a cloud provider. Advocate XACML-compliant entitlement management.
- Implementing internal single sign-on (SSO) for enterprise applications can be leveraged to simplify cloud provider engagements and implementation.
- Burgeoning Identity as a Service providers may be appropriate solutions to simplify IdM complexities, understand that this becomes another cloud provider relationship to vet and manage.

References

Upcoming publication “Cloud Security and Privacy” – Tim Mather, Subra Kumaraswamy, Shahid Latif

Domain 14: Storage

Contributors: *Jean Pawluk*
 Liam Lynch
 Jim Reavis

Problem Statement

The requirement for secure, accessible data storage has been increasing at an exponential rate for short term and ongoing use, both inside the enterprise and within the cloud. In this rush to save time and money, the security risks of storage in the cloud are often minimized or ignored. Yet the overall security of cloud data storage, the persistence of data and the reliability of the information deserves careful consideration.

Organizations considering the cost savings and agility for time to market must still consider the basic tenants of information security. They are Confidentiality, Integrity, and Availability. Data storage falls squarely into these 3 system qualities of an information security program. Your privacy group and legal advisors, as a close business partner will have additional requirements especially in the case where data is shared between public and private clouds.

Do you know how your storage provider plans to ensure that your data is still reliable and available when your business needs it? Will you and your customers be able to trust the promises that all stored private and confidential information is protected? Does the geographical location of where the data being processed and stored matter? Is your data is being securely stored and kept separate from the other residents in the data storage farm?

Issues and Guidance

There are a growing number of Infrastructure as a Service (IaaS) storage services available, the OpenCrowd Taxonomy in Domain 1 lists nine and there are certainly many others. Additionally, Software as a Service (SaaS) and Platform as a Service (PaaS) provider offerings are driving the demand for large amounts of cloud-based storage which is opaque to the customer. While these different forms of cloud computing present storage to the customer in different ways, in keeping with our overarching theme that customers “look under the hood”, there are several common areas of concern to consider.

Customers need to consider the reality of what happens with stored data and if the actions represented in the user interface (files being deleted, records being updated) are in alignment with what a different interface would represent (administrative program or forensics tool). What really happens with your data when it is deleted at the SaaS/PaaS/IaaS level? Is it truly destroyed or are copies retained? How can the data destruction be proven? As pointed out in Domain 6: Information Lifecycle Management, data destruction is extremely difficult in a multi-tenant environment. We believe it is fundamental that the customer understand cloud provider storage retirement processes. We recommend that the cloud provider should be utilizing strong storage encryption that renders data unreadable when storage is recycled, disposed of, or accessed by any means outside of authorized applications.

Customers need to have some level of understanding of the storage technology and architecture. Our concerns do not weigh heavily on the storage hardware, but rather on how that storage is organized via software. In a Service Oriented Architecture (SOA), a layer of abstraction is

Security Guidance for Critical Areas of Focus in Cloud Computing

introduced, which has huge benefits. However, do layers of abstraction potentially hide a storage subsystem that in effect has multiple ingress/egress points which may be insecure? Could a grid storage system or a technology such as Hadoop potentially cross boundaries between different clouds or other trust boundaries? Applications that are built to access data need to be careful if the data is spread across private and public clouds. Given appropriate SLAs and policy enforcement at the data level an application can be built to access data in both types of clouds. Due care should be built into the data access subsystem to not copy private cloud data to the public cloud for processing efficiency alone.

In other domains in this whitepaper, such as Domain 3 Legal and Domain 5 Compliance and Audit, the importance of understanding the location of data for regulatory purposes is underscored. We agree that this is a question that should always be asked. Some cloud providers are able to host data within specific countries or regions. Because of the dynamic nature of the cloud, we would also want to know if it is possible to be alerted of any change in location via our SLA. With geography as a backdrop, it is an important due diligence activity to understand what the ramifications are of storage located under the purview of any foreign governing entity. Would the governing entity be able to seize the storage or would the cloud provider be compelled to submit encryption keys?

As in Domain 9 Data Center Operations, we are concerned about assuring compartmentalization during the provisioning and normal operations of storage servicing multiple clients. We recommend that the customer understand how encryption is managed on multi-tenant storage. Is there a single key for all customers, one key per customer, or multiple keys per customer? While there are many application level dependencies that limit our ability to provide blanket guidance, when possible we recommend granular encryption policies supported by multiple keys per customer to support compartmentalization and to limit the scope of breaches.

In general, as part of due diligence we recommend gaining an overall understanding of your cloud provider's storage management capabilities. What is the nature of their data search and analysis capabilities? What sorts of data migration systems are in place? How are data archiving and backup functions performed? Is your data transferred to offline and portable storage and how is that storage managed? How is long term archiving managed? Will the current encryption technology and keys be maintained to allow access to data many years later? Knowing as much as possible about the storage security capabilities within the cloud provider enables the customer to create alignment with its data lifecycle management strategy and data classification policies.

Domain 15: Virtualization

Contributor: *Brian O'Higgins*

Problem Statement:

Cloud computing incorporates highly virtualized environments, in addition to varying levels of service provided with Software as a Service (SaaS), Infrastructure as a Service (IaaS), or Platform as a Service (PaaS) business offerings. The security implications are markedly different than those common in the world of physical data centers. At lesser service levels offered by a cloud provider, a lesser amount of security is integrated into the service. Some cloud providers deliver raw compute resources, leaving application security entirely up to its customers. Others deliver an application that runs in a secure environment managed by the SaaS provider. At the most basic level, visibility, control and audit capabilities underlie the questions of virtualization security in cloud computing environments.

Issues

Top new risks and security considerations introduced by virtualization in cloud environments include the following:

- New technology, old vulnerabilities remain, and new ones arise.
- Loss of security by default.
- Commingling challenges integrity.
- Jurisdiction, control and regulatory issues.
- New management challenges also impact security.

Virtualization is the underlying technology enabling cloud computing and transforming the face of the modern datacenter. Whenever new technology is introduced, new opportunities are created for things to go wrong. Virtualized environments use the same operating systems, enterprise and web applications as physical environments. The ability for malware to remotely exploit vulnerabilities in these systems and applications remains a primary threat to virtualized environments and is an even greater risk in instance-based cloud computing environments where patch management responsibilities lie with the system lessee. Virtual environments also introduce new risk due to the dynamic nature of virtual machines (VMs). Because VMs can quickly be reverted to previous instances, and easily moved between physical servers, it is difficult to achieve and maintain consistent security. Also don't forget the fundamental rule that as complexity goes up, security goes down. This is the situation with cloud services, and when combined with the rush to deploy and limited experience, security problems are all but guaranteed.

Loss of 'security by default' will be more apparent in the early days of cloud services and virtualization until experience and best practices catch up. 'Security-by-default' is the implicit security existing in day-to-day operations. Consider a typical large organization IT environment. There are numerous subject matter experts or groups with particular domain expertise, and all play a role in securely deploying new application servers, for example. This might involve network specialists to install and configure NICs, platform experts to set up and test the server,

Security Guidance for Critical Areas of Focus in Cloud Computing

operations people to configure and test a particular application. Firewalling is heavily used to establish security zones for particular applications, and network zoning plays a large role in the security architecture. Physical security typically plays a large role as well, when you consider the doors, locks, and other access controls that are usually in place around the server room.

In the virtualized world many of the mechanisms that provide security-by-default are missing, and special attention must be paid to replace them. If the virtualization platform is a VMware-based environment, for example, and someone from the operations group logs on and dispatches a new Windows server and application, can you be certain the guest OS is locked down appropriately? What about enforcement of zones? Were any issues related to commingling of data introduced?

Core virtualization technology itself introduces new attack surfaces with the hypervisor and other management components, but more important is the severe impact virtualization has on network security. Virtual machines now communicate over a hardware backplane, not a network. As such, standard network security controls are now blinded to this traffic and they cannot do their job of monitoring or other in-line blocking functions. These controls need to take a new form to run in the virtual environment. An interim step in moving from old network-style security to full virtualization security is to deploy network appliances as virtual appliances. However, care must be taken as the network context is lost as VMs move freely from one physical box to another. In fact, this new mobility for server images can be compared to laptops in the enterprise. At one point PCs could rely on being protected by a secure perimeter, but that model was broken when the endpoints could easily move outside the secure perimeter.

Commingling of data that arises with centralized services and repositories is a concern. A centralized database as provided by a cloud computing service in theory should improve security over data distributed over a vast number and mix of endpoints. However, this is an example of centralizing risk, and a breach can have significant consequences.

Another issue is likely commingling of VMs of different sensitivities and security. In cloud computing environments, the lowest common denominator of security will be shared by all tenants in the multitenant virtual environment unless a new security architecture can be achieved that does not ‘wire in’ any network dependency for protection.

Regulatory compliance and legal issues have been substantially covered in other domains, but it should be mentioned again as a key issue of concern specifically to virtualization, as this is often the mechanism that obscures data location, a key regulatory concern.

New management challenges for virtualization result in challenges for security and compliance. The on-demand nature of cloud computing allows IT organizations to increase automation to save costs. New virtual machines can be configured and spun up with a few mouse clicks, in contrast to the “old days” when it took weeks to order a new server, get it delivered, configured, tested, and made operational. Policy and controls must be in place so only appropriate individuals can dispatch new VMs. Each VM needs to be configured to an appropriate locked-down state. Cloning a misconfigured VM augments the risk exponentially. Backup strategies often take advantage of archiving VMs on a regular basis, and this leads to VM sprawl. These images all need to be managed, on-line and off-line. VMs need to be fully patched, and this can be challenging if an off-line unpatched image might come back on-line at any point.

Guidance

To mitigate against these new risks, these priority areas need to be addressed, among others:

- Augmenting platform security with third-party security mechanisms
- The use of inherent security advantages afforded by virtualization technology
- The use of VM-based security mechanisms
- The control and monitoring of administrative access to servers and applications
- Configuring secure images and following risk-based security standards, such as PCI-DSS

Security is a concern for virtualization platform providers and steps are being taken to incorporate security mechanisms into these platforms. While it seems tempting to rely on a particular platform provider alone for security, it is not enough. In time, it will provide a baseline that continually improves, but to be more secure, additional controls will be necessary. While perfect security is impossible, you are actually secure if you are just a bit better than the effort an enemy will use against you. For practical purposes, this means be more secure than your neighbor, as attackers go after the easy targets first. Organizations can achieve a superior level of security by focusing time and attention on the issue.

Virtualization itself delivers some inherent security advantages through the properties of isolation. Moving an application, or set of applications to a particular VM provides a generally isolated environment. As well as potentially reducing data commingling issues, other advantages include increased redundancy, failover, and application longevity, for example by minimizing application instability and allowing old applications to run on current hardware. Fault isolation can be at the hardware level, system recovery can be made simpler (re-start a VM), and in general, cleaner images are maintained as fresh copies can be rapidly dispatched to flush problems. The defined memory space for each VM also helps in resisting buffer overflows. VM-based security is now necessary. Physical network security controls do not work when the network taps they require are not available. Virtual appliance gateway versions of these network controls can provide some monitoring functions, but cannot effectively block traffic that runs between VMs on the same physical box. New hypervisor APIs such as VMsafe enable greater visibility into the network-based activity on the hardware backplane in local deployments, but these mechanisms will be ineffective for end-user organizations needing to gain visibility of their cloud computing resources. Cloud service providers will leverage this level of introspection when available for the applicable virtualization platform, and when performance allows. Dedicated virtualization security technology will mix network and host-based security techniques to achieve adequate protection of the dynamic virtual infrastructure. VM-based controls such as firewall, application control, IDS/IPS, log inspection, integrity monitoring and web application protection are necessary to provide comprehensive threat protection for virtual environments.

The dynamic nature of the virtual infrastructure increases the potential for exposure and risk that must be mitigated. Other mechanisms become mandatory when dealing with dynamic growth and management requirements. For example, using strong authentication, most likely combined with identity management software, especially for system administrators to control creation, configuration and cloning of virtual machines. It is crucial to monitor the integrity of the system

Security Guidance for Critical Areas of Focus in Cloud Computing

and the access to administrative functions. To have the greatest visibility and protection, log inspection and integrity monitoring capabilities must be deployed at the virtual machine level.

Lock down a virtual image appropriately to create a ‘gold standard’, and clone away from this base. Examples of configuration checklists can be found at the Center for Internet Security (CIS), <http://www.cisecurity.org/benchmarks.html>, and on the Information Assurance Support Environment (designed to support the US Department of Defense Community) <http://iase.disa.mil/stigs/checklist/>. Follow risk-based security standards when possible. Rather than forcing long checklists, these newer standards focus on maintaining a secure environment by assessing your particular situation and doing what is necessary to comply. PCI-DSS is a great example, and is required when you are dealing with credit card and associated payment transactions. Encryption of end user data sent to the cloud also becomes necessary in most cases.

Conclusion

Virtualization is the enabling technology for cloud computing. Organizations not leveraging cloud computing today are likely looking to cloud computing for tomorrow. Datacenters that have consolidated physical servers to multiple virtual machine instances on virtualized servers can take immediate steps to increase security in their virtualized environment, as well as prepare these virtual machines for the migration to cloud environments when appropriate. Questions of visibility, control and audit capabilities need to be addressed in the cloud services environment. CISOs engaging in cloud computing need to be able to answer the following:

- Will I still have the same security policy control over my applications and services?
- Can I prove to my organization and my customers that I am still secure and meeting my SLAs?
- Am I still compliant, and can I prove it to my auditors?
- How do I solve old security problems, like vulnerabilities, patch management, and system integrity in this new environment?

Appendix A. Contact Information

Cloud Security Alliance

www.cloudsecurityalliance.org

info@cloudsecurityalliance.org

LinkedIn Group

www.linkedin.com/groups?gid=1864210

Twitter

www.twitter.com/cloudsa

#cloudsa

Facebook

<http://www.facebook.com/group.php?gid=57805091701>