cloud security alliance℠

SECURITY GUIDANCE FOR CRITICAL AREAS OF FOCUS IN CLOUD COMPUTING V3.0

INTRODUCTION

The guidance provided herein is the third version of the Cloud Security Alliance document, **"Security Guidance for Critical Areas of Focus in Cloud Computing,"** which was originally released in April 2009. The permanent archive locations for these documents are:

http://www.cloudsecurityalliance.org/guidance/csaguide.v3.0.pdf (this document)

http://www.cloudsecurityalliance.org/guidance/csaguide.v2.1.pdf (version 2 guidance)

http://www.cloudsecurityalliance.org/guidance/csaguide.v1.0.pdf (version 1 guidance)

In a departure from the second version of our guidance, each domain was assigned its own editor and peer reviewed by industry experts. The structure and numbering of the domains align with industry standards and best practices. We encourage the adoption of this guidance as a good operating practice in strategic management of cloud services. These white papers and their release schedule are located at:

http://www.cloudsecurityalliance.org/guidance/

In another change from the second version, there are some updated domain names. We have these changes: **Domain 3: Legal Issues: Contracts and Electronic Discovery** and **Domain 5: Information Management and Data Security.** We now have added another domain, which is **Domain 14: Security as a Service**

TABLE OF CONTENTS

FOREWORD

Welcome to the third version of the Cloud Security Alliance's "Security Guidance for Critical Areas of Focus in Cloud Computing." As cloud computing begins to mature, managing the opportunities and security challenges becomes crucial to business development. We humbly hope to provide you with both guidance and inspiration to support your business needs while managing new risks.

The Cloud Security Alliance has delivered actionable, best practices based on previous versions of this guidance. As we continue to deliver tools to enable businesses to transition to cloud services while mitigating risk, this guidance will act as the compass for our future direction. In v3.0, you will find a collection of facts and opinions gathered from over seventy industry experts worldwide. We have compiled this information from a range of activities, including international chapters, partnerships, new research, and conference events geared towards furthering our mission. You can follow our activities at www.cloudsecurityalliance.org.

The path to secure cloud computing is surely a long one, requiring the participation of a broad set of stakeholders on a global basis. However, we should happily recognize the progress we are seeing: new cloud security solutions are regularly appearing, enterprises are using our guidance to engage with cloud providers, and a healthy public dialogue over compliance and trust issues has erupted around the world. The most important victory we have achieved is that security professionals are vigorously engaged in securing the future, rather than simply protecting the present.

Please stay engaged on this topic and continue to work with us to complete this important mission.

Best Regards,

Jerry Archer

Dave Cullinane

Nils Puhlmann

Alan Boehme

Paul Kurtz

Jim Reavis

The Cloud Security Alliance Board of Directors

ACKNOWLEDGMENTS

Domain Authors/Contributors

Domain 1: Chris Hoff, Paul Simmonds

Domain 2: Marlin Pohlman, Becky Swain, Laura Posey, Bhavesh Bhagat

Domain 3: Francoise Gilbert, Pamela Jones Harbour, David Kessler, Sue Ross, Thomas Trappler

Domain 4: Marlin Pohlman, Said Tabet

Domain 5: Rich Mogull, Jesus Luna

Domain 6: Aradhna Chetal, Balaji Ramamoorthy, Jim Peterson, Joe Wallace, Michele Drgon, Tushar Bhavsar

Domain 7: Randolph Barr, Ram Kumar, Michael Machado, Marlin Pohlman

Domain 8: Liam Lynch

Domain 9: Michael Panico, Bernd Grobauer, Carlo Espiritu, Kathleen Moriarty, Lee Newcombe, Dominik Birk, Jeff Reed

Domain 10: Aradhna Chetal, Balaji Ramamoorthy, John Kinsella, Josey V. George, Sundararajan N., Devesh Bhatt, Tushar Bhavsar

Domain 11: Liam Lynch

Domain 12: Paul Simmonds, Andrew Yeomans, Ian Dobson, John Arnold, Adrian Secombe, Peter Johnson, Shane Tully, Balaji Ramamorthy, Subra Kumaraswamy, Rajiv Mishra, Ulrich Lang, Jens Laundrup, Yvonne Wilson

Domain 13: Dave Asprey, Richard Zhao, Kanchanna Ramasamy Balraj, Abhik Chaudhuri, Melvin M. Rodriguez

Domain 14: Jens Laundrup, Marlin Pohlman, Kevin Fielder

Peer Reviewers

Valmiki Mukherjee, Bernd Jaeger, Ulrich Lang, Hassan Takabi, Pw Carey, Xavier Guerin, Troy D. Casey, James Beadel, Anton Chuvakin, Tushar Jain, M S Prasad, Damir Savanovic, Eiji Sasahara, Chad Woolf, Stefan Pettersson, M S Prasad, Nrupak Shah, Kimberley Laris, Henry St. Andre, Jim Peterson, Ariel Litvin, Tatsuya Kamimura, George Ferguson, Andrew Hay, Danielito Vizcayno,

K.S. Abhiraj, Liam Lynch, Michael Marks, JP Morgenthal, Amol Godbole, Damu Kuttikrishnan, Rajiv Mishra, Dennis F. Poindexter, Neil Fryer, Andrea Bilobrk, Balaji Ramamoorthy, Damir Savanovic

Editorial Team

Archie Reed: Domains 3, 8, 9

Chris Rezek: Domains 2, 4, 5, 7, 13, 14

Paul Simmonds: Domains 1, 6, 10, 11, 12

CSA Staff

Technical Writer/Editor: Amy L. Van Antwerp

Graphic Designer: Kendall Scoboria

Research Director: J.R. Santos

LETTER FROM THE EDITORS

Over the past three years, the Cloud Security Alliance has attracted around 120 corporate members and has a broad remit to address all aspects of cloud security, including compliance, global security-related legislation and regulation, identity management, and the challenge of monitoring and auditing security across a cloud-based IT supply chain. CSA is becoming the focal point for security standards globally, aligning multiple, disparate government policies on cloud security and putting forward standards for ratification by international standards bodies.

CSA sees itself as a cloud security standards incubator, so its research projects use rapid development techniques to produce fast results. To this end, the CSA Guidance editorial team is proud to present the third version of its flagship "Security Guidance for Critical Areas of Focus in Cloud Computing." This work is a set of best security practices CSA has put together for 14 domains involved in governing or operating the cloud (Cloud Architecture, Governance and Enterprise Risk Management, Legal: Contracts and Electronic Discovery, Compliance and Audit, Information Management and Data Security, Portability and Interoperability, Traditional Security, Business Continuity and Disaster Recovery, Data Center Operations, Incident Response, Notification and Remediation, Application Security, Encryption and Key Management, Identity and Access Management, Virtualization, and Security as a Service).

CSA guidance in its third edition seeks to establish a stable, secure baseline for cloud operations. This effort provides a practical, actionable road map to managers wanting to adopt the cloud paradigm safely and securely. Domains have been rewritten to emphasize security, stability, and privacy, ensuring corporate privacy in a multi-tenant environment.

Over the past two years, version 2.1 of the guidance has served as the foundation for research in multiple areas of cloud security. Deliverables now in use from the TCI Architecture to the GRC Stack were inspired by previous versions of the guidance, and it is our hope that this version will be no different. The guidance serves as a high level primer for chief executives, consumers, and implementers wishing to adopt cloud services as an alternative or supplement to traditional infrastructure. However, the guidance is designed with innovation in mind. Those with an entrepreneurial mindset should read this work with an eye toward the inferred services and approaches many of the authors have included in the domain creation. Investors and corporate decision makers will also find this work of interest, as it serves as a roadmap for innovation and development already in place in companies throughout the world. Security practitioners and educators will find elements of this book both authoritative and thought provoking, and as the industry evolves, the value the authors have included should prove influential and timely.

In the third edition, the guidance assumes a structural maturity in parallel with multinational cloud standards development in both structure and content. Version 3.0 extends the content included in previous versions with practical recommendations and requirements that can be measured and audited. Please note that different interpretations of the term "requirements" exist, which we use throughout the document. Our guidance does not represent a statutory obligation, but "requirements" was chosen to represent guidance appropriate for virtually all use cases we could envision, and also aligns our guidance with similar well-accepted documents. CSA industry expert authors have endeavored to present a working product that is measured and balanced between the interests of cloud providers and tenants. Controls focus on the preservation of tenant data ownership integrity while embracing the concept of a shared physical infrastructure. Guidance Version 3.0 incorporates the highly dynamic nature of cloud computing, industry learning curve, and new developments within other research projects such as Cloud Controls Matrix, Consensus Assessments Initiative, Trusted Cloud Initiative, and GRC Stack Initiative and ties in the various CSA activities into one comprehensive C-level best practice. The Security Guidance v3.0 will serve as the gateway to emerging standards being

developed in the world's standards organization and is designed to serve as an executive-level primer to any organization seeking a secure, stable transition to hosting their business operations in the cloud.

On behalf of the Cloud Security Alliance, we would like to thank each and every volunteer for their time and effort in the development and editing of this new release of our flagship guidance document. While we believe this is our best, most widely reviewed work to date, the topic is still evolving and although our foremost intent is to guide, we also intend to inspire the readers to become involved in improving and commenting on the direction those composing the body of work have outlined. We humbly and respectfully submit this effort to the industry and await the most important component of any dialog, your opinion. We are eager to hear your feedback regarding this updated guidance. If you found this guidance helpful or would like to see it improved, please consider joining the Cloud Security Alliance as a member or contributor.

Best Regards,

Paul Simmonds Chris Rezek Archie Reed

Security Guidance v3.0 Editors

AN EDITORIAL NOTE ON RISK

Throughout this Guidance we make extensive recommendations on reducing your risk when adopting cloud computing, but not all the recommendations are necessary or even realistic for all cloud deployments. As we compiled information from the different working groups during the editorial process, we quickly realized there simply wasn't enough space to provide fully nuanced recommendations for all possible risk scenarios. Just as a critical application might be too important to move to a public cloud provider, there might be little or no reason to apply extensive security controls to low-value data migrating to cloud-based storage.

With so many different cloud deployment options — including the SPI service models (SPI refers to Software as a Service, Platform as a Service, or Infrastructure as a Service, explained in depth in Domain 1); public vs. private deployments, internal vs. external hosting, and various hybrid permutations — no list of security controls can cover all circumstances. As with any security area, organizations should adopt a risk-based approach to moving to the cloud and selecting security options. The following is a simple framework to help evaluate initial cloud risks and inform security decisions.

This process is not a full risk assessment framework, nor a methodology for determining all your security requirements. It's a quick method for evaluating your tolerance for moving an asset to various cloud computing models.

Identify the Asset for the Cloud Deployment

At the simplest, assets supported by the cloud fall into two general categories:

1. Data
2. Applications/Functions/Processes

We are either moving information into the cloud, or transactions/processing (from partial functions all the way up to full applications).

With cloud computing our data and applications don't need to reside in the same location, and we can choose to shift only parts of functions to the cloud. For example, we can host our application and data in our own data center, while still outsourcing a portion of its functionality to the cloud through a Platform as a Service.

The first step in evaluating risk for the cloud is to determine exactly what data or function is being considered for the cloud. This should include potential uses of the asset once it moves to the cloud to account for scope creep. Data and transaction volumes are often higher than expected.

Evaluate the Asset

The next step is to determine how important the data or function is to the organization. You don't need to perform a detailed valuation exercise unless your organization has a process for that, but you do need at least a rough assessment of how sensitive an asset is, and how important an application/function/process is.

For each asset, ask the following questions:

1. How would we be harmed if the asset became widely public and widely distributed?
2. How would we be harmed if an employee of our cloud provider accessed the asset?
3. How would we be harmed if the process or function were manipulated by an outsider?
4. How would we be harmed if the process or function failed to provide expected results?
5. How would we be harmed if the information/data were unexpectedly changed?
6. How would we be harmed if the asset were unavailable for a period of time?

Essentially we are assessing confidentiality, integrity, and availability requirements for the asset; and how the risk changes if all or part of the asset is handled in the cloud. It's very similar to assessing a potential outsourcing project, except that with cloud computing we have a wider array of deployment options, including internal models.

Map the Asset to Potential Cloud Deployment Models

Now we should have an understanding of the asset's importance. Our next step is to determine which deployment models we are comfortable with. Before we start looking at potential providers, we should know if we can accept the risks implicit to the various deployment models: private, public, community, or hybrid; and hosting scenarios: internal, external, or combined.

For the asset, determine if you are willing to accept the following options:

1. Public.
2. Private, internal/on-premises.
3. Private, external (including dedicated or shared infrastructure).
4. Community; taking into account the hosting location, potential service provider, and identification of other community members.
5. Hybrid. To effectively evaluate a potential hybrid deployment, you must have in mind at least a rough architecture of where components, functions, and data will reside.

At this stage you should have a good idea of your comfort level for transitioning to the cloud, and which deployment models and locations fit your security and risk requirements.

Evaluate Potential Cloud Service Models and Providers

In this step focus on the degree of control you'll have at each SPI tier to implement any required risk management. If you are evaluating a specific offering, at this point you might switch to a fuller risk assessment.

Your focus will be on the degree of control you have to implement risk mitigations in the different SPI tiers. If you already have specific requirements (e.g., for handling of regulated data) you can include them in the evaluation.

Map Out the Potential Data Flow

If you are evaluating a specific deployment option, map out the data flow between your organization, the cloud service, and any customers/other nodes. While most of these steps have been high-level, before making a final decision it's absolutely essential to understand whether, and how, data can move in and out of the cloud.

If you have yet to decide on a particular offering, you'll want to sketch out the rough data flow for any options on your acceptable list. This is to insure that as you make final decisions, you'll be able to identify risk exposure points.

Conclusions

You should now understand the importance of what you are considering moving to the cloud, your risk tolerance (at least at a high level), and which combinations of deployment and service models are acceptable. You should also have a good idea of potential exposure points for sensitive information and operations.

These together should give you sufficient context to evaluate any other security controls in this Guidance. For low-value assets you don't need the same level of security controls and can skip many of the recommendations — such as on-site inspections, discoverability, and complex encryption schemes. A high-value regulated asset might entail audit and data retention requirements. For another high-value asset not subject to regulatory restrictions, you might focus more on technical security controls.

Due to our limited space, as well as the depth and breadth of material to cover, this document contains extensive lists of security recommendations. Not all cloud deployments need every possible security and risk control. Spending a little time up front evaluating your risk tolerance and potential exposures will provide the context you need to pick and choose the best options for your organization and deployment.

SECTION I // CLOUD ARCHITECTURE

DOMAIN 1 //
CLOUD COMPUTING ARCHITECTURAL FRAMEWORK

This domain, the Cloud Computing Architectural Framework, provides a conceptual framework for the rest of the Cloud Security Alliance's guidance. The contents of this domain focus on a description of cloud computing that is specifically tailored to the unique perspective of IT network and security professionals.

The final section of this domain provides a brief introduction to each of the other domains.

Understanding the architectural framework described in this domain is an important first step in understanding the remainder of the Cloud Security Alliance guidance. The framework defines many of the concepts and terms used throughout the other domains.

Overview. The following three sections define this architectural perspective in terms of:

- The terminology used throughout the guidance, to provide a consistent lexicon
- The architectural requirements and challenges for securing cloud applications and services
- A reference model that describes a taxonomy of cloud services and architectures

1.1 What Is Cloud Computing?

Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services). Cloud computing is a disruptive technology that has the potential to enhance collaboration, agility, scaling, and availability, and provides the opportunities for cost reduction through optimized and efficient computing. The cloud model envisages a world where components can be rapidly orchestrated, provisioned, implemented and decommissioned, and scaled up or down to provide an on-demand utility-like model of allocation and consumption.

From an architectural perspective, there is much confusion surrounding how cloud is both similar to and different from existing models of computing and how these similarities and differences impact the organizational, operational, and technological approaches to network and information security practices. There is a thin line between conventional computing and cloud computing. However, cloud computing will impact the organizational, operational, and technological approaches to data security, network security, and information security good practice.

There are many definitions today that attempt to address cloud from the perspective of academicians, architects, engineers, developers, managers, and consumers. This document focuses on a definition that is specifically tailored to the unique perspectives of IT network and security professionals.

1.2 What Comprises Cloud Computing?

This version of the Cloud Security Alliance's Guidance features definitions that are based on published work of the scientists at the U.S. National Institute of Standards and Technology (**NIST**)[1] and their efforts around defining cloud computing.

NIST's publication is generally well accepted, and the Guidance aligns with the NIST Working Definition of Cloud Computing (NIST 800-145 as of this writing) to bring coherence and consensus around a common language to focus on use cases rather than semantic nuances.

It is important to note that this guide is intended to be broadly usable and applicable to organizations globally. While NIST is a U.S. government organization, the selection of this reference model should not be interpreted to suggest the exclusion of other points of view or geographies.

NIST defines cloud computing by describing five essential characteristics, three cloud service models, and four cloud deployment models. They are summarized in visual form in Figure 1 and explained in detail below.

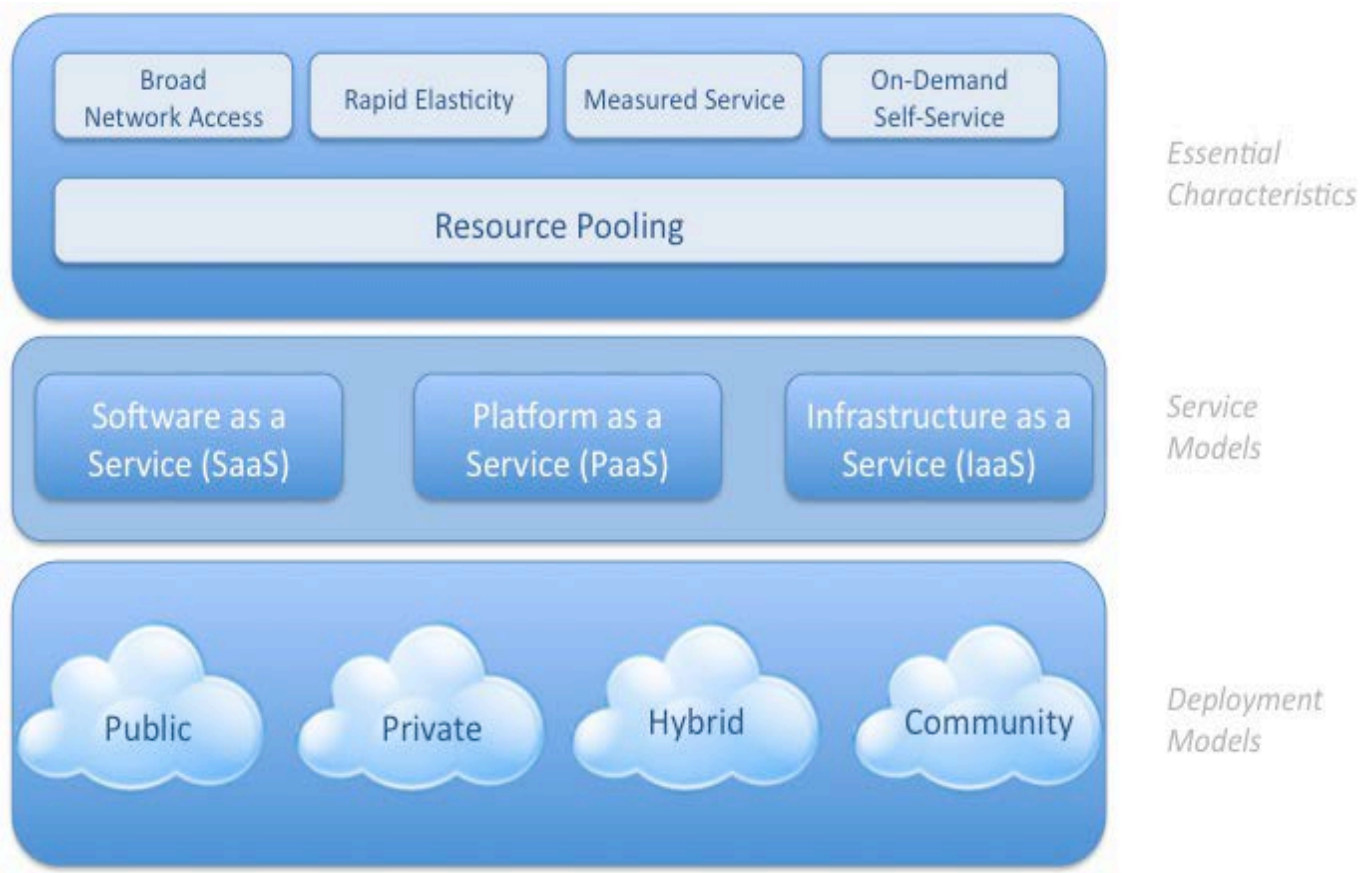

Figure 1—NIST Visual Model of Cloud Computing Definition[2]

[1] **NIST** - National Institute of Standards and Technology

1.3 The Characteristics of Cloud Computing

It is important to recognize that cloud services are often but not always utilized in conjunction with, and enabled by, virtualization technologies. There is no requirement, however, that ties the abstraction of resources to virtualization technologies, and in many offerings virtualization by hypervisor or operating system container is not utilized.

Further, it should be noted that multi-tenancy is not called out as an essential cloud characteristic by NIST but is often discussed as such. Although not an essential characteristic of cloud computing in the NIST model, CSA has identified multi-tenancy as an important element of cloud.

1.4 Multi-Tenancy

For this document multi tenancy is considered an important element, and the following section will outline the CSA's understanding/definition as an important element of cloud computing.

Multi-tenancy in its simplest form implies use of same resources or application by multiple consumers that may belong to same organization or different organization. The impact of multi-tenancy is visibility of residual data or trace of operations by other user or tenant.

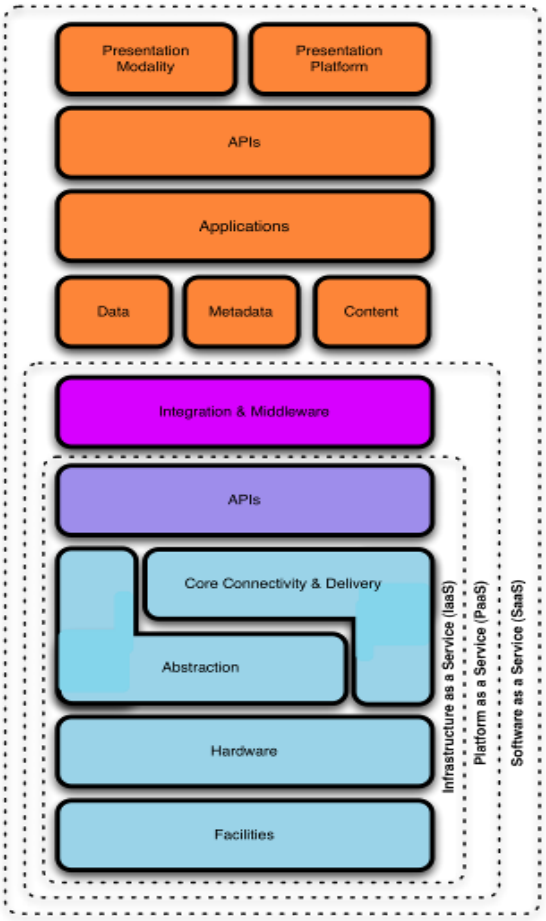

Figure 2—Multi-Tenancy

Multi-tenancy in cloud service models implies a need for policy-driven enforcement, segmentation, isolation, governance, service levels, and chargeback/billing models for different consumer constituencies.

Consumers may choose to utilize a public cloud providers' service offering on an individual user basis or, in the instance

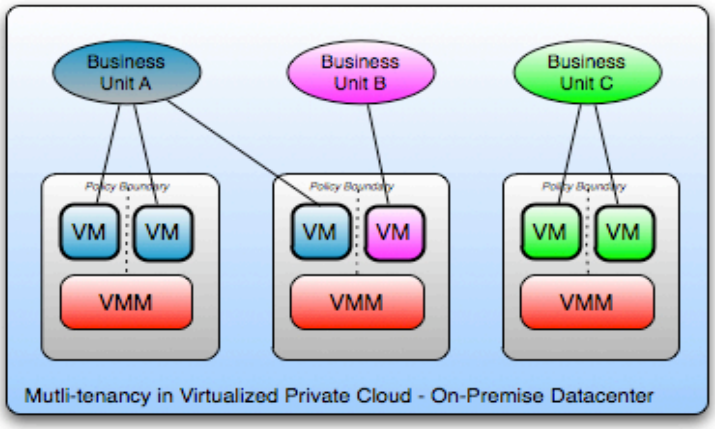

Private Cloud of Company XYZ with 3 business units, each with different security, SLA, governance and chargeback policies on shared infrastructure

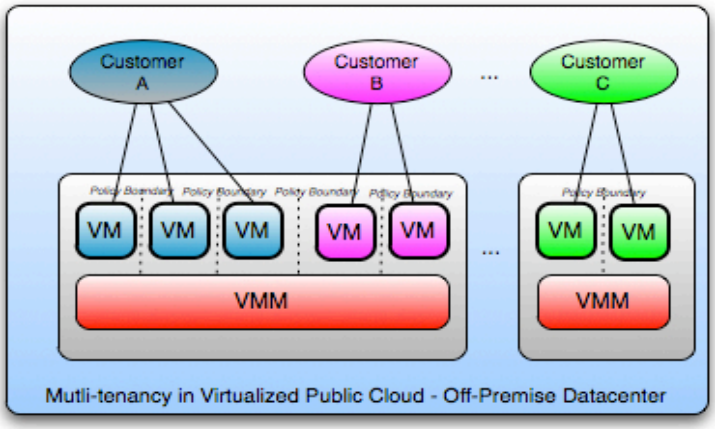

Public Cloud Provider with 3 business customers, each with different security, SLA, governance and billing policies on shared infrastructure

of private cloud hosting, an organization may segment users as different business units sharing a common infrastructure.

From a provider's perspective, multi-tenancy suggests an architectural and design approach to enable economies of scale, availability, management, segmentation, isolation, and operational efficiency. These services leverage shared infrastructure, data, metadata, services, and applications across many different consumers.

Multi-tenancy can also take on different definitions depending upon the cloud service model of the provider; inasmuch as it may entail enabling the capabilities described above at the infrastructure, database, or application levels. An example would be the difference between an Infrastructure-as-a-Service (**IaaS**)[2], Software-as-a-Service (**SaaS**)[3], and (**PaaS**)[4] multi-tenant implementation.

Cloud deployment models place different importance on multi-tenancy. However, even in the case of a private cloud, a single organization may have a multitude of third party consultants and contractors, as well as a desire for a high degree of logical separation between business units. Thus, multi-tenancy concerns should always be considered.

1.5 Cloud Reference Model

Understanding the relationships and dependencies between cloud computing models is critical to understanding cloud computing security risks. IaaS is the foundation of all cloud services, with PaaS building upon IaaS, and SaaS in turn building upon PaaS as described in the Cloud Reference Model diagram. In this way, just as capabilities are inherited, so are information security issues and risk. It is important to note that commercial cloud providers may not neatly fit into the layered service models. Nevertheless, the reference model is important for relating real-world services to an architectural framework and understanding that the resources and services require security analysis.

IaaS includes the entire infrastructure resource stack from the facilities to the hardware platforms that reside in them. It 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**[5], which allows management and other forms of interaction with the infrastructure by consumers.

***Infrastructure as a Service* (IaaS)**, delivers computer infrastructure (typically a platform virtualization environment) as a service, along with raw storage and networking. Rather than purchasing servers, software, data-center space, or network equipment, clients instead buy those resources as a fully outsourced service.

***Software as a service* (SaaS)**, sometimes referred to as "on-demand software," is a software delivery model in which software and its associated data are hosted centrally (typically in the (Internet) cloud) and are typically accessed by users using a thin client, normally using a web browser over the Internet.

***Platform as a service* (PaaS)**, is the delivery of a computing platform and solution stack as a service. PaaS offerings facilitate deployment of applications without the cost and complexity of buying and managing the underlying hardware and software and provisioning hosting capabilities. This provides all of the facilities required to support the complete life cycle of building and delivering web applications and services entirely available from the Internet.

[2] **IaaS** - Infrastructure as a Service
[3] **SaaS** - Software as a Service
[4] **PaaS** - Platform as a Service
[5] **API** - Application Programming Interface

PaaS sits on top of IaaS and adds an additional layer of integration with application development frameworks, middleware capabilities, and functions such as database, messaging, and queuing. These services allow developers to build applications on the platform with programming languages and tools that are supported by the stack.

SaaS in turn is built upon the underlying IaaS and PaaS stacks and provides a self-contained operating environment that is used to deliver the entire user experience, including the content, its presentation, the application(s), and management capabilities.

It should therefore be clear that there are significant trade-offs to each model in terms of integrated features, complexity versus openness (extensibility), and security. Generally, SaaS provides the most integrated functionality built directly into the offering, with the least consumer extensibility, and a relatively high level of integrated security (at least the provider bears a responsibility for security).

PaaS is intended to enable developers to build their own applications on top of the platform. As a result, it tends to be more extensible than SaaS, at the expense of customer-ready features. This tradeoff extends to security features and capabilities, where the built-in capabilities are less complete, but there is more flexibility to layer on additional security.

IaaS provides few if any application-like features, but enormous extensibility. This generally means less integrated security capabilities and functionality beyond protecting the infrastructure itself. This model requires that operating systems, applications, and content be managed and secured by the cloud consumer.

The key takeaway for security architecture is that the lower down the stack the cloud service provider stops, the more security capabilities and management consumers are responsible for implementing and managing themselves.

Service levels, security, governance, compliance, and liability expectations of the service and provider are contractually stipulated, managed to, and enforced, when a service level agreement (**SLA's)**[6], is offered to the consumer. There are two types of SLA's, negotiable and non-negotiable. In the absence of an SLA, the consumer administers all aspects of the cloud under its control. When a non-negotiable SLA is offered, the provider administers those portions stipulated in the agreement. In the case of PaaS or IaaS, it is usually the responsibility of the consumer's system administrators to effectively manage the residual services specified in the SLA, with some offset expected by the provider for securing the underlying platform and infrastructure components to ensure basic service availability and security. It should be clear in all cases that one can assign/transfer responsibility but not necessarily accountability.

Narrowing the scope or specific capabilities and functionality within each of the cloud delivery models, 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'.

While a broader review of the growing set of cloud computing solutions is outside the scope of this document, the OpenCrowd Cloud Solutions taxonomy in the figure below provides an excellent starting point, however the CSA does not specifically endorse any of the solutions or companies shown below. It provides the below diagram to demonstrate the diversity of offerings available today.

[6] **SLA** - Service Level Agreement

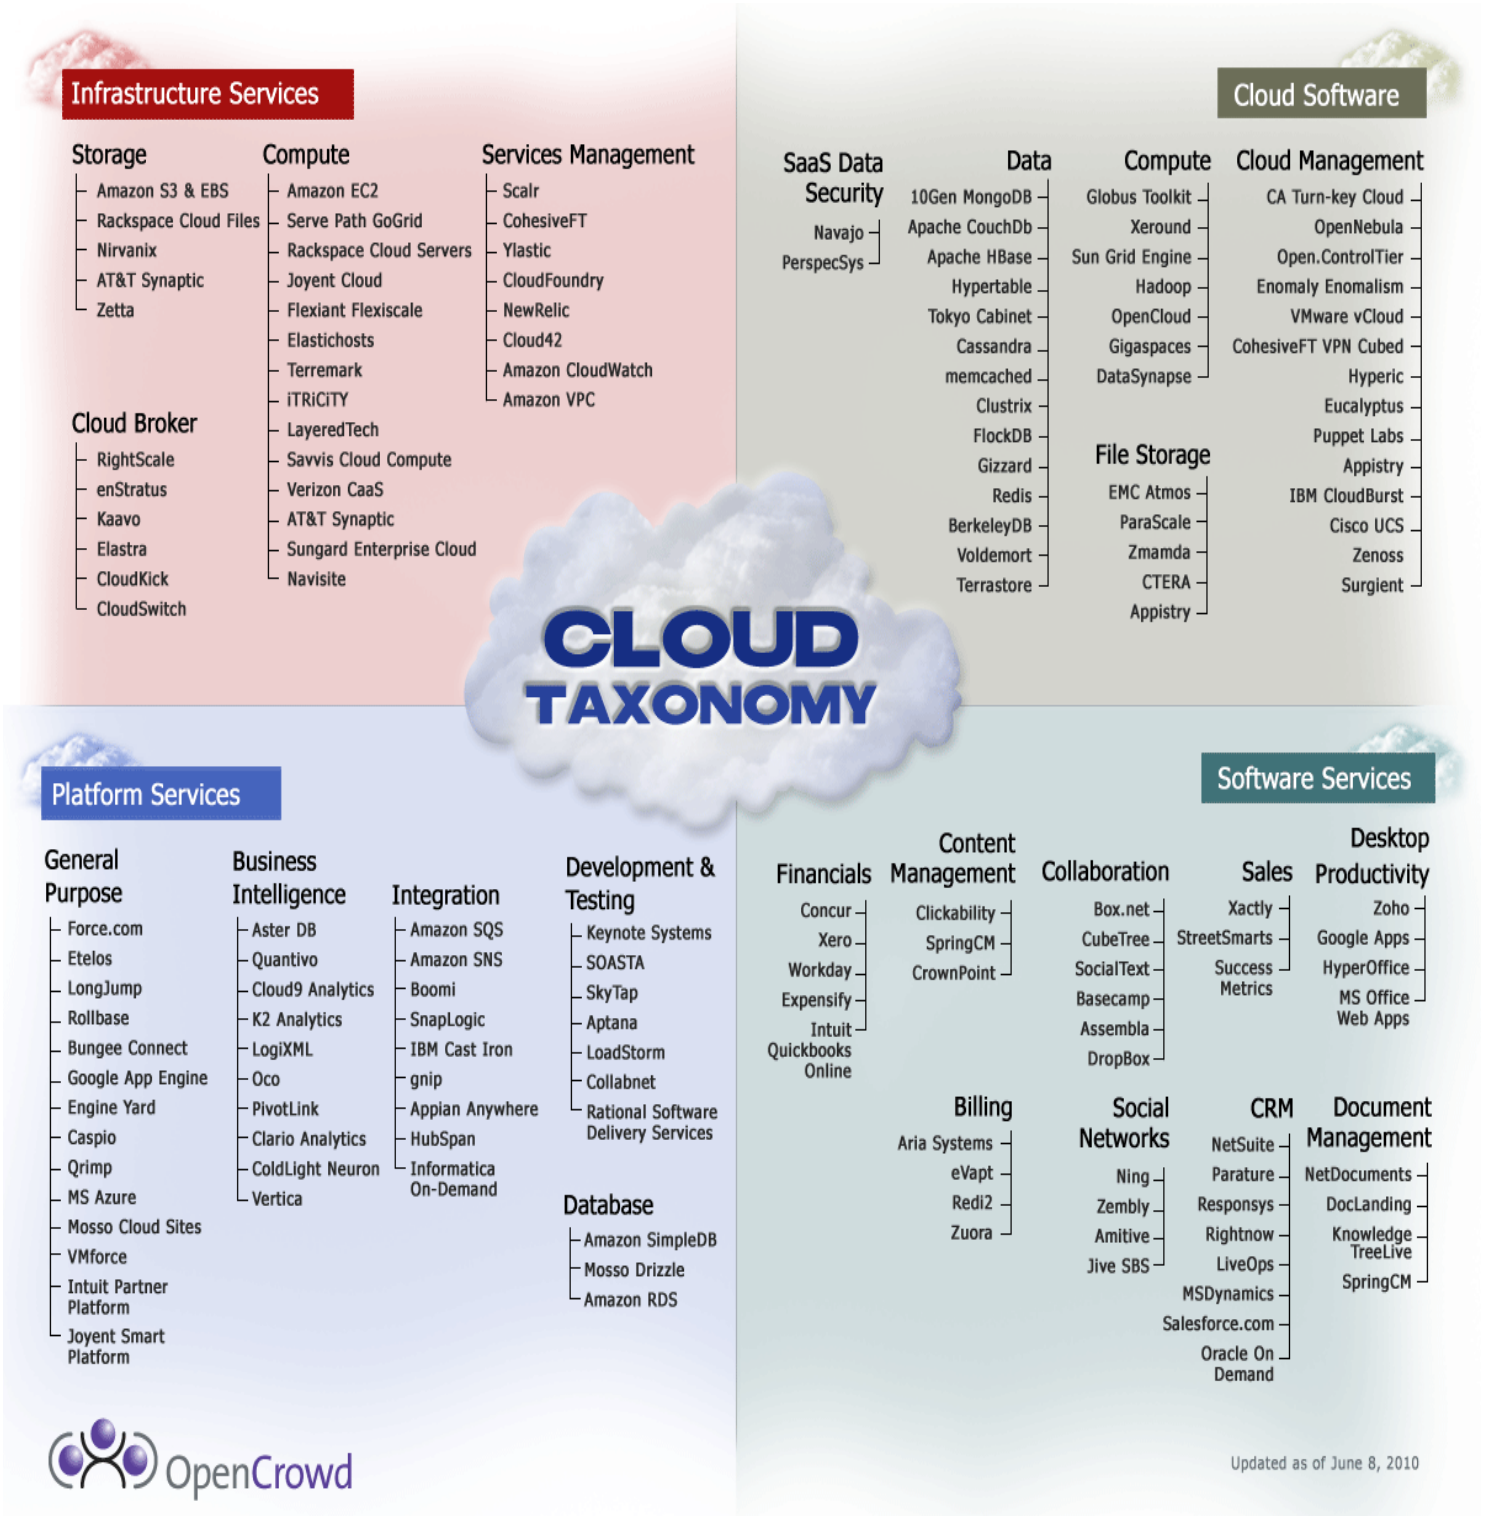

Figure 3 – OpenCrowd Taxonomy[7]

For an excellent overview of the many cloud computing use cases, the Cloud Computing Use Case Group produced a collaborative work to describe and define common cases and demonstrate the benefits of cloud, with their goal being to "...bring together cloud consumers and cloud vendors to define common use cases for cloud computing...and highlight

[7] http://www.opencrowd.com/assets/images/views/views_cloud-tax-lrg.png

the capabilities and requirements that need to be standardized in a cloud environment to ensure interoperability, ease of integration, and portability."

1.5.1 Cloud Security Reference Model

The cloud security reference model addresses the relationships of these classes and places them in context with their relevant security controls and concerns. For organizations and individuals grappling with cloud computing for the first time, it is important to note the following to avoid potential pitfalls and confusion:

- The notion of *how* cloud services are deployed is often used interchangeably with *where* they are provided, which can lead to confusion. Public or private clouds may be described as external or internal, which may not be accurate in all situations.
- The manner in which cloud services are consumed is often described relative to the location of an organization's management or security perimeter (usually defined by the presence of a known demarc). While it is still important to understand where security boundaries lie in terms of cloud computing, the notion of a well-demarcated perimeter is an anachronistic concept for most organizations.
- The re-perimeterization and the erosion of trust boundaries already happening in the enterprise is amplified and accelerated by cloud computing. Ubiquitous connectivity, the amorphous nature of information interchange, and the ineffectiveness of traditional static security controls which cannot deal with the dynamic nature of cloud services, all require new thinking with regard to cloud computing. The Jericho Forum[8] has produced a considerable amount of material on the re- perimeterization of enterprise networks, including many case studies.

The deployment and consumption modalities of cloud should be thought of not only within the context of 'internal' versus 'external' as they relate to the physical location of assets, resources, and information; but also by whom they are being consumed; and who is responsible for their governance, security, and compliance with policies and standards.

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

- The types of assets, resources, and information being managed
- Who manages them and how
- Which controls are selected and how they are integrated
- Compliance issues

For example, a **LAMP**[9] stack deployed on Amazon's AWS EC2 would be classified as a public, off-premise, third-party managed IaaS solution, even if the instances and applications/data contained within them were managed by the consumer or a third party. A custom application stack serving multiple business units, deployed on Eucalyptus under a

[8] http://www.jerichoforum.org

[9] LAMP-Linux (operating system), Apache HTTP Server, MySQL (database software) and Perl/PHP/Python, the principal components to build a viable general purpose web server

corporation's control, management, and ownership, could be described as a private, on-premise, self-managed SaaS solution. Both examples utilize the elastic scaling and self-service capabilities of cloud.

The following table summarizes these points:

Table 1—Cloud Computing Deployment Models

	Infrastructure 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
Private/ Community	Organization Or Third Party Provider	Organization Third Party Provider	On-Premise Off-Premise	Trusted
Hybrid	Both Organization & Third Party Provider	Both Organization & Third Party Provider	Both On-Premise & Off-Premise	Trusted & Untrusted

[1] *Management includes: governance, 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 and speaks to ownership versus control*
[4] *Trusted consumers of service are those who are considered part of an organization's legal/contractual/ policy 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.*

Another way of visualizing how combinations of cloud service models, deployment models, physical locations of resources, and attribution of management and ownership, is the Jericho Forum's Cloud Cube Model[10], shown in the figure to the right:

The Cloud Cube Model illustrates the many permutations available in cloud offerings today and presents four criteria/dimensions in order to differentiate cloud "formations" from one another and the manner of their provision, in order to understand how cloud computing affects the way in which security might be approached.

The Cloud Cube Model also highlights the challenges of understanding and mapping cloud models to control frameworks and standards such as ISO/IEC 27002, which provides "...a series of guidelines and general principles for initiating,

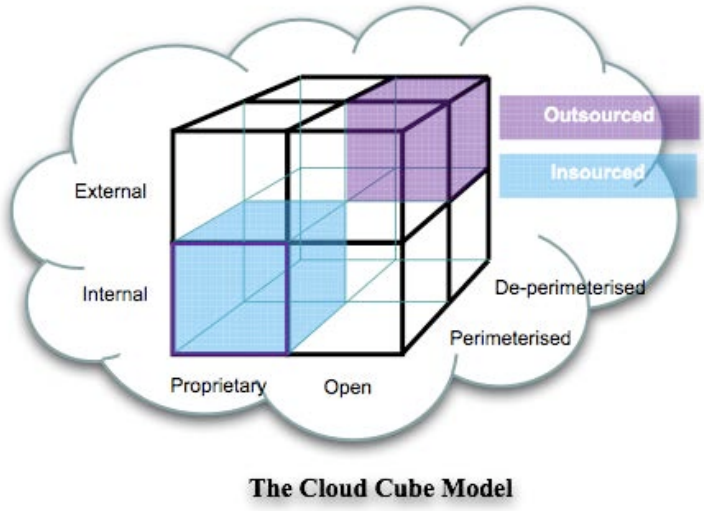

Figure 4—Jericho Cloud Cube Model

[10] http://www.opengroup.org/jericho/cloud_cube_model_v1.0.pdf

implementing, maintaining, and improving information security management within an organization."

The ISO/IEC 27002, section 6.2, "External Parties" control objective states: "…the security of the organization's information and information processing facilities should not be reduced by the introduction of external party products or services…"

As such, the differences in methods and responsibility for securing the three cloud service models mean that consumers of cloud services are faced with a challenging endeavor. Unless cloud providers can readily disclose their security controls and the extent to which they are implemented to the consumer and the consumer knows which controls are needed to maintain the security of their information, there is tremendous potential for misguided risk management decisions and detrimental outcomes.

First, one classifies a cloud service against the cloud architecture model. Then it is possible to map its security architecture as well as business, regulatory, and other compliance requirements against it as a gap-analysis exercise. The result determines the general "security" posture of a service and how it relates to an asset's assurance and protection requirements.

The figure below shows an example of how a cloud service mapping can be compared against a catalogue of compensating controls to determine which controls exist and which do not — as provided by the consumer, the cloud service provider, or a third party. This can in turn be compared to a compliance framework or set of requirements such as PCI DSS, as shown.

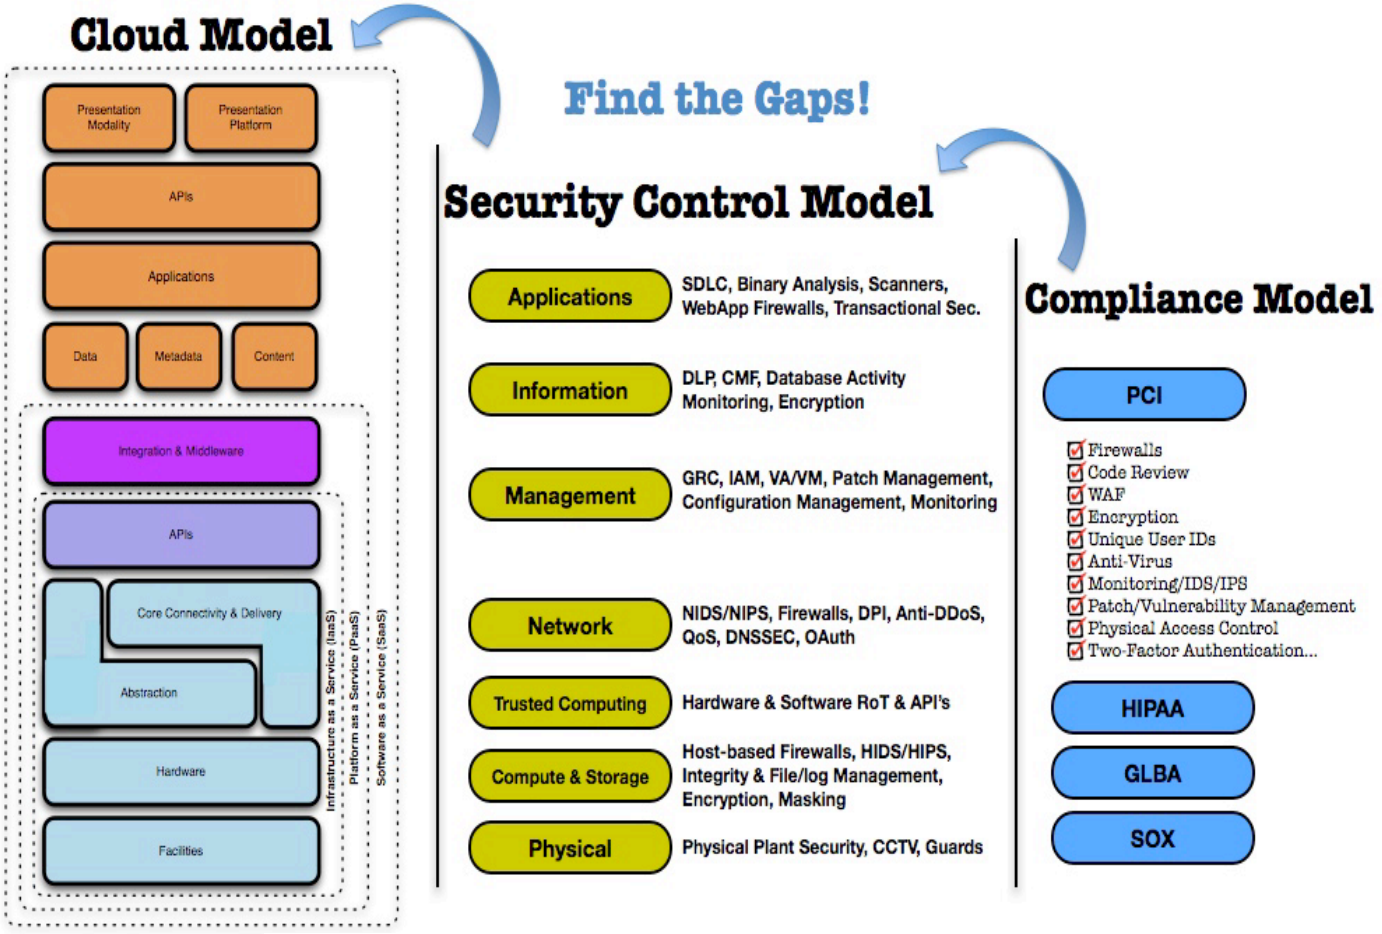

Figure 5—Mapping the Cloud Model to the Security Control & Compliance

Once this gap analysis is complete, per 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. This, in turn, helps to determine how the gaps and ultimately risks should be addressed: accepted, transferred, or mitigated.

It is important to note that the use of cloud computing as an operational model does not inherently provide for or prevent achieving compliance. The ability to comply with any requirement is a direct result of the service and deployment model utilized and the design, deployment, and management of the resources in scope.

For an excellent overview of control frameworks which provides good illustrations of the generic control framework alluded to above, see the Open Security Architecture Group's[11] "landscape" of security architecture patterns documentation, or the always useful and recently updated NIST 800-53 revision 3 Recommended Security Controls for Federal Information Systems and Organizations security control catalogue.

1.5.2 What Is Security for Cloud Computing?

Security controls in cloud computing are, for the most part, no different than security controls in any IT environment. However, because of the cloud service models employed, the operational models, and the technologies used to enable cloud services, cloud computing may present different risks to an organization than traditional IT solutions.

An organization's security posture is characterized by the maturity, effectiveness, and completeness of the risk-adjusted security controls implemented. These controls are implemented in one or more layers ranging from the facilities (physical security), to the network infrastructure (network security), to the IT systems (system security), all the way to the information and applications (application security). Additionally, controls are implemented at the people and process levels, such as separation of duties and change management, respectively.

As described earlier in this document, the security responsibilities of both the provider and the consumer greatly differ between cloud service models. Amazon's AWS EC2 infrastructure as a service offering, as an example, includes vendor responsibility for security up to the hypervisor, meaning they can only address security controls such as physical security, environmental security, and virtualization security. The consumer, in turn, is responsible for security controls that relate to the IT system (instance) including the operating system, applications, and data.

The inverse is true for Salesforce.com's customer resource management (CRM) SaaS offering. Because Salesforce.com provides the entire "stack," the provider is not only responsible for the physical and environmental security controls, but it must also address the security controls on the infrastructure, the applications, and the data. This alleviates much of the consumer's direct operational responsibility.

There is currently no way for a naive consumer of cloud services to simply understand what exactly he/she is responsible for [though reading this guidance document should help], but there are efforts underway by the CSA and other bodies to define standards around cloud audit.

One of the attractions of cloud computing is the cost efficiencies afforded by economies of scale, reuse, and standardization. To bring these efficiencies to bear, cloud providers have to provide services that are flexible enough to serve the largest customer base possible, maximizing their addressable market. Unfortunately, integrating security into these solutions is often perceived as making them more rigid.

[11] www.opensecurityarchitecture.org

This rigidity often manifests in the inability to gain parity in security control deployment in cloud environments compared to traditional IT. This stems mostly from the abstraction of infrastructure, and the lack of visibility and capability to integrate many familiar security controls, especially at the network layer.

The figure below illustrates these issues: in SaaS environments the security controls and their scope are negotiated into the contracts for service; service levels, privacy, and compliance are all issues to be dealt with legally in contracts. In an IaaS offering, while the responsibility for securing the underlying infrastructure and abstraction layers belongs to the provider, the remainder of the stack is the consumer's responsibility. PaaS offers a balance somewhere in between, where securing the platform falls onto the provider, but both securing the applications developed against the platform and developing them securely, belong to the consumer.

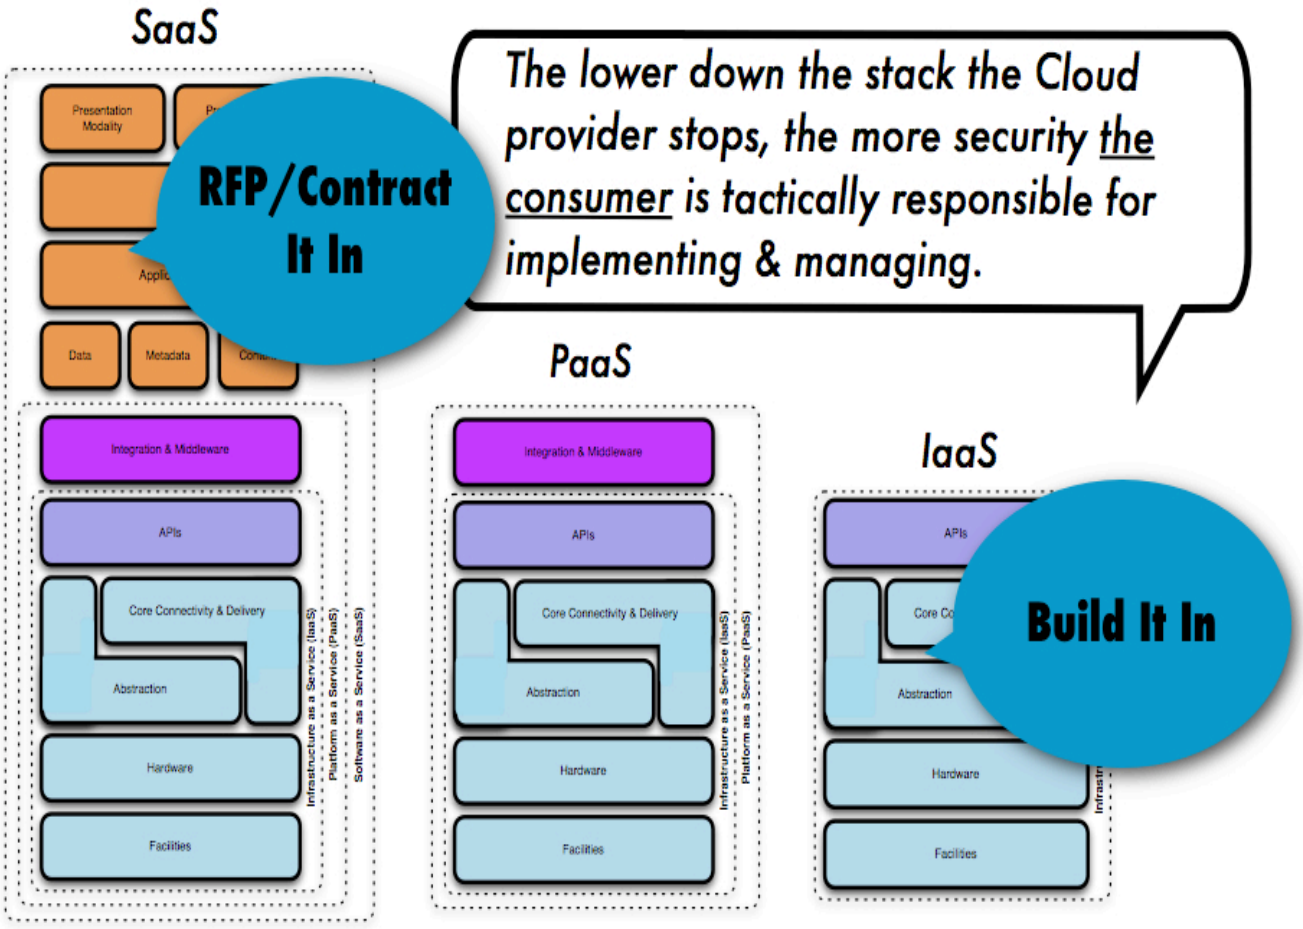

Figure 6—How Security Gets Integrated

Understanding the impact of these differences between service models and how they are deployed is critical to managing the risk posture of an organization.

1.5.3 Beyond Architecture: The Areas of Critical Focus

The thirteen other domains which comprise the remainder of the CSA guidance highlight areas of concern for cloud computing and are tuned to address both the strategic and tactical security 'pain points' within a cloud environment and can be applied to any combination of cloud service and deployment model.

The domains are divided into two broad categories: governance and operations. The governance domains are broad and address strategic and policy issues within a cloud computing environment, while the operational domains focus on more tactical security concerns and implementation within the architecture.

Table 2a— Governance Domains

DOMAIN	GUIDANCE DEALING WITH…
Governance and Enterprise Risk Management	The ability of an organization to govern and measure enterprise risk introduced by cloud computing. Items such as legal precedence for agreement breaches, ability of user organizations to adequately assess risk of a cloud provider, responsibility to protect sensitive data when both user and provider may be at fault, and how international boundaries may affect these issues.
Legal Issues: Contracts and Electronic Discovery	Potential legal issues when using cloud computing. Issues touched on in this section include protection requirements for information and computer systems, security breach disclosure laws, regulatory requirements, privacy requirements, international laws, etc.
Compliance and Audit	Maintaining and proving compliance when using cloud computing. Issues dealing with evaluating how cloud computing affects compliance with internal security policies, as well as various compliance requirements (regulatory, legislative, and otherwise) are discussed here. This domain includes some direction on proving compliance during an audit.
Information Management and Data Security	Managing data that is placed in the cloud. Items surrounding the identification and control of data in the cloud, as well as compensating controls that can be used to deal with the loss of physical control when moving data to the cloud, are discussed here. Other items, such as who is responsible for data confidentiality, integrity, and availability are mentioned.
Portability and Interoperability	The ability to move data/services from one provider to another, or bring it entirely back in-house. Together with issues surrounding interoperability between providers.

Table 2b - Operational Domains

DOMAIN	GUIDANCE DEALING WITH…
Traditional Security, Business Continuity and Disaster Recovery	How cloud computing affects the operational processes and procedures currently used to implement security, business continuity, and disaster recovery. The focus is to discuss and examine possible risks of cloud computing, in hopes of increasing dialogue and debate on the overwhelming demand for better enterprise risk management models. Further, the section touches on helping people to identify where cloud computing may assist in diminishing certain security risks, or entails increases in other areas.
Data Center Operations	How to evaluate a provider's data center architecture and operations. This is primarily focused on helping users identify common data center characteristics that could be detrimental to on-going services, as well as characteristics that are fundamental to long-term stability.
Incident Response, Notification and Remediation	Proper and adequate incident detection, response, notification, and remediation. This attempts to address items that should be in place at both provider and user levels to enable proper incident handling and forensics. This domain will help you understand the complexities the cloud brings to your current incident-handling program.
Application Security	Securing application software that is running on or being developed in the cloud. This includes items such as whether it's appropriate to migrate or design an application to run in the cloud, and if so, what type of cloud platform is most appropriate (SaaS, PaaS, or IaaS).
Encryption and Key Management	Identifying proper encryption usage and scalable key management. This section is not prescriptive, but is more informational in discussing *why* they are needed and identifying issues that arise in use, both for protecting access to resources as well as for protecting data.
Identity and Access Management	Managing identities and leveraging directory services to provide access control. The focus is on issues encountered when extending an organization's identity into the cloud. This section provides insight into assessing an organization's readiness to conduct cloud-based Identity, Entitlement, and Access Management (IdEA).
Virtualization	The use of virtualization technology in cloud computing. The domain addresses items such as risks associated with multi-tenancy, VM isolation, VM co-residence, hypervisor vulnerabilities, etc. This domain focuses on the security issues surrounding system/hardware virtualization, rather than a more general survey of all forms of virtualization.

Security as a Service	Providing third party facilitated security assurance, incident management, compliance attestation, and identity and access oversight. Security as a service is the delegation of detection, remediation, and governance of security infrastructure to a trusted third party with the proper tools and expertise. Users of this service gain the benefit of dedicated expertise and cutting edge technology in the fight to secure and harden sensitive business operations.

1.6 Cloud Deployment Models

Regardless of the service model utilized (SaaS, PaaS, or IaaS), there are four deployment models for cloud services with derivative variations that address specific requirements.

It is important to note that there are derivative cloud deployment models emerging due to the maturation of market offerings and customer demand. An example of such is virtual private clouds — a way of utilizing public cloud infrastructure in a private or semi-private manner and interconnecting these resources to the internal resources of a consumer's datacenter, usually via virtual private network (VPN) connectivity.

The architectural mind-set used when designing "solutions" has clear implications on the future flexibility, security, and mobility of the resultant solution, as well as its collaborative capabilities. As a rule of thumb, perimeterized solutions are less effective than de-perimeterized solutions in each of the four areas. Careful consideration should also be given to the choice between proprietary and open solutions for similar reasons.

Deployment models

- *Public Cloud*. The cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services.
- *Private Cloud*. The cloud infrastructure is operated solely for a single organization. It may be managed by the organization or by a third party and may be located on-premise or off-premise.
- *Community Cloud.* The cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (e.g., mission, security requirements, policy, or compliance considerations). It may be managed by the organizations or by a third party and may be located on-premise or off-premise.
- *Hybrid Cloud.* The cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load-balancing between clouds).

1.7 Recommendations

Cloud service delivery is divided among three archetypal models and various derivative combinations. The three fundamental classifications are often referred to as the "SPI Model," where "SPI" refers to Software, Platform or Infrastructure (as a Service), respectively.

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

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

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

The NIST model and this document do not directly address the emerging service model definitions associated with cloud service brokers, those providers that offer intermediation, monitoring, transformation/portability, governance, provisioning, and integration services and negotiate relationships between various cloud providers and consumers.

In the short term, as innovation drives rapid solution development, consumers and providers of cloud services will enjoy varied methods of interacting with cloud services in the form of developing API's and interfaces and so cloud service brokers will emerge as an important component in the overall cloud ecosystem.

Cloud service brokers will abstract these possibly incompatible capabilities and interfaces on behalf of consumers to provide proxy in advance of the arrival of common, open, and standardized ways of solving the problem longer term with a semantic capability that allows the consumer a fluidity and agility in being able to take advantage of the model that works best for their particular needs.

It is also important to note the emergence of many efforts centered on the development of both open and proprietary API's, which seek to enable things such as management, security, and interoperability for cloud. Some of these efforts include the Open Cloud Computing Interface Working Group, Amazon EC2 API, VMware's DMTF-submitted vCloud API, Sun's Open Cloud API, Rackspace API, and GoGrid's API, to name just a few. Open, standard API's will play a key role in cloud portability and interoperability as well as common container formats such as the DMTF's Open Virtualization Format (OVF).

While there are many working groups, drafts, and published specifications under consideration at this time, it is natural that consolidation will take effect as market forces, consumer demand, and economics pare down this landscape to a more manageable and interoperable set of players.

1.8 Requirements

Cloud services exhibit five essential characteristics that demonstrate their relation to, and differences from, traditional computing approaches.

- ✓ **On-demand self-service.** A consumer can unilaterally provision computing capabilities such as server time and network storage as needed automatically without requiring human interaction with a service provider.
- ✓ **Broad network access.** Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, laptops, and **PDA's**)[12] as well as other traditional or cloud-based software services.
- ✓ **Resource pooling.** The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a degree of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources, but may be able to specify location at a higher level of abstraction (e.g., country, state, or datacenter). Examples of resources include storage, processing, memory, network bandwidth, and virtual machines. Even private clouds tend to pool resources between different parts of the same organization.
- ✓ **Rapid elasticity.** Capabilities can be rapidly and elastically provisioned — in some cases automatically — to quickly scale out; and rapidly released to quickly scale in. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time.
- ✓ **Measured service.** Cloud systems automatically control and optimize resource usage by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, or active user accounts). Resource usage can be monitored, controlled, and reported — providing transparency for both the provider and consumer of the service.

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 frameworks, and management frameworks; and in turn to compliance standards.

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, it is impossible to address more detailed issues rationally. This architectural overview, along with the thirteen other areas of critical focus, will provide the reader with a solid foundation for assessing, operationalizing, managing, and governing security in cloud computing environments.

[12] **PDA** - Personal Digital Assistant

REFERENCES

[1] NIST definition of Cloud.
NIST 500-292 “NIST Cloud Computing Reference Architecture”

[2] NIST definitions and API homepages.
www.cloud-standards.org

[3] Jericho Forum Cloud Cube Model.
www.opengroup.org/jericho/cloud_cube_model_v1.0.pdf

SECTION II // GOVERNING IN THE CLOUD

DOMAIN 2 //
GOVERNANCE & ENTERPRISE RISK MANAGEMENT

The fundamental issues of governance and enterprise risk management in cloud computing concern the identification and implementation of the appropriate organizational structures, processes, and controls to maintain effective information security governance, risk management, and compliance. Organizations should also assure reasonable information security across the information supply chain, encompassing providers and customers of cloud computing services and their supporting third party vendors, in any cloud deployment model.

An effective governance and enterprise risk management cloud computing program flows from well-developed information security governance processes as part of the organization's overall corporate governance obligations of due care. Well-developed information security governance processes result in information security management programs that are scalable with the business, repeatable across the organization, measurable, sustainable, defensible, continually improving, and cost-effective on an ongoing basis.

For many cloud deployments, a major element of governance will be the agreement between provider and customer. For custom environments, detailed care can be taken, and negotiated, for each provision. For larger scale customer or providers, there will be a decision whether to trade off attention to detail vs. scalability of effort. Attention can be prioritized base on criticality or value at risk for the particular workload (e.g., up-time and availability may be more important for email than for HR systems). As the space continues to mature, projects like Cloud Audit or STAR will provide more standardized governance methods, therefore greater scalability.

Overview. This domain addresses:

- Governance
- Enterprise Risk Management

This section maps to the Cloud Control Matrix Controls DG-01, IS-02 and the use of GRC-XML and CloudAudit to establish solvency.

2.1 Corporate Governance

Corporate governance is the set of processes, technologies, customs, policies, laws, and institutions affecting the way an enterprise is directed, administered or controlled. Corporate governance also includes the relationship among the many stakeholders involved and the goals of the company involved. Good governance is based on the acceptance of the rights of shareholders, as the true owners of the corporation, and the role of senior management as trustees. There are many models of corporate governance; however, all follow five basic principles:

- Auditing supply chains
- Board and management structure and process
- Corporate responsibility and compliance
- Financial transparency and information disclosure

- Ownership structure and exercise of control rights

A key factor in a customer decision to engage a corporation is the confidence that expectations will be met. For cloud services, the interdependencies among multiple services make it more difficult for a customer to sort out the responsible party. If that results in less confidence in a particular vendor, then further engagement with that vendor is less likely. If this becomes a systemic feature, the loss of confidence in one actor will rollover to others, and the market failure will increase the likelihood of both external action and alternative participants.

Stakeholders should carefully consider the monitoring mechanisms that are appropriate and necessary for the company's consistent performance and growth.

2.2 Enterprise Risk Management

Enterprise risk management (ERM) is rooted in the commitment by every organization to provide value for its stakeholders. All businesses face uncertainty and one of management's challenges is to determine how one organization can measure, manage, and mitigate that uncertainty. Uncertainty presents both opportunity and risk with potential to increase or decrease the value of the organization and its strategies.

Information risk management is the process of identifying and understanding exposure to risk and capability of managing it, aligned with the risk appetite and tolerance of the data owner. Hence, it is the primary means of decision support for IT resources dedicated to delivering the confidentiality, integrity, and availability of information assets.

This section maps to the Cloud Control Matrix Controls DG-08 and the use of ISO31000, ISF and ISACA guidelines to establish solvency.

Enterprise risk management in business includes the methods and processes used by organizations to manage risks and seize opportunities related to the achievement of their objectives. In a cloud environment, management selects a risk response strategy for specific risks identified and analyzed, which may include:

- Avoidance—exiting the activities giving rise to risk
- Reduction—taking action to reduce the likelihood or impact related to the risk
- Share or insure—transferring or sharing a portion of the risk to finance it
- Accept—no action is taken due to a cost/benefit decision

Risk management is naturally a balancing process with the goal not necessarily to minimize uncertainty or variation, but rather the goal of maximizing value in line with risk appetite and strategy.

There are many variables, values, and risk in any cloud opportunity or program that affect the decision whether a cloud service should be adopted from a risk or business value standpoint. Each enterprise has to weigh those variables to decide whether the cloud is an appropriate solution.

Cloud computing offers enterprises many possible benefits, some of these benefits include:

- Optimized resource utilization

- Cost savings for cloud computing tenants
- Transitioning of capital expenses
- (CAPEX) to operating expenses (OPEX)
- Dynamic scalability of IT power for clients
- Shortened life cycle development of new applications or deployments
- Shortened time requirements for new business implementation

Customers should view cloud services and security as supply chain security issues. This means examining and assessing the provider's supply chain (service provider relationships and dependencies) to the extent possible. This also means examining the provider's own third party management. Assessment of third party service providers should specifically target the provider's incident management, business continuity and disaster recovery policies, and processes and procedures; and should include review of co-location and back-up facilities. This should include review of the provider's internal assessments of conformance to its own policies and procedures and assessment of the provider's metrics to provide reasonable information regarding the performance and effectiveness of its controls in these areas. Incident information can be specified in contracts, SLAs, or other joint agreements, and can be communicated automatically or periodically, directly into reporting systems or delivered to key personnel. The level of attention and scrutiny should be connected to the value at risk – if the third party will not directly access enterprise data, then the level of risk drops significantly and vice versa.

Customers should review the risk management processes and governance of their providers, and ensure that practices are consistent and aligned.

2.3 Permissions

Permissions

- Adopt an established risk framework for monitoring and measuring corporate risk
- Adopt metrics to measure risk management performance (e.g., Security Content Automation Protocol (**SCAP**) [13], Cybersecurity Information Exchange Framework (**CYBEX**)[14], or **GRC-XML**[15]).
- Adopt a risk centric viewpoint of corporate governance with senior management taking the role of trustee for both the shareholders and the stakeholders in the supply chain.
- Adopt a framework from legal perspective to account for differences across jurisdictions.

[13] **SCAP** - Security Content Automation Protocol
[14] **CYBEX** - Cybersecurity Information Exchange Framework
[15] **GRC-XML** - technology standards to enable and enhance the sharing of information between various technologies that support GRC efforts

2.4 Recommendations

- Reinvest the cost savings obtained by cloud computing services into increased scrutiny of the security capabilities of the provider, application of security controls, and ongoing detailed assessments and audits to ensure requirements are continuously met.
- User organizations should include review of specific information security governance structure and processes, as well as specific security controls, as part of their due diligence for prospective provider organizations. The provider's security governance processes and capabilities should be assessed for sufficiency, maturity, and consistency with the user's information security management processes. The provider's information security controls should be demonstrably risk-based and clearly support these management processes.
- Collaborative governance structures and processes between customers and providers should be identified as necessary, both as part of the design and development of service delivery, and as service risk assessment and risk management protocols, and then incorporated into service agreements.
- Security departments should be engaged during the establishment of Service Level Agreements (**SLA's**) [16] and contractual obligations to ensure that security requirements are contractually enforceable.
- Metrics and standards for measuring performance and effectiveness of information security management should be established prior to moving into the cloud. At a minimum, organizations should understand and document their current metrics and how they will change when operations are moved into the cloud and where a provider may use different (potentially incompatible) metrics.
- Due to the lack of physical control over infrastructure by customers, in many Cloud Computing deployments, SLA's, contract requirements, and provider documentation play a larger role in risk management than with traditional, enterprise owned infrastructure.
- Due to the on-demand provisioning and multi-tenant aspects of cloud computing, traditional forms of audit and assessment may not be available or may be modified. For example, some providers restrict vulnerability assessments and penetration testing, while others limit availability of audit logs and activity monitoring. If these are required per your internal policies, you may need to seek alternative assessment options, specific contractual exceptions, or an alternative provider better aligned with your risk management requirements.
- If the services provided in the cloud are essential to corporate operations a risk management approach should include identification and valuation of assets, identification and analysis of threats and vulnerabilities and their potential impact on assets (risk and incident scenarios), analysis of the likelihoods of events/scenarios, management-approved risk acceptance levels and criteria, and the development of risk treatment plans with multiple options (control, avoid, transfer, accept). The outcomes of risk treatment plans should be incorporated into service agreements.
- Risk assessment approaches between provider and user should be consistent with consistency in impact analysis criteria and definition of likelihood. The user and provider should jointly develop risk scenarios for the cloud service; this should be intrinsic to the provider's design of service for the user, and to the user's assessment of cloud service risk.

[16] **SLA** - Service Level Agreement

- Due to the evolving nature of cloud and its providers, care should be taken to include vendor risk, e.g., business survivability of providers, portability of data and applications, and interoperability of services.
- Asset inventories should account for assets supporting cloud services and under the control of the provider. Asset classification and valuation schemes should be consistent between user and provider.
- The service, and not just the vendor, should be the subject of risk assessment. The use of cloud services, and the particular service and deployment models to be utilized, should be consistent with the risk management objectives of the organization, as well as with its business objectives.
- Cloud Computing service customers and providers should develop robust information security governance, regardless of the service or deployment model. Information security governance should be collaboration between customers and providers to achieve agreed-upon goals that support the business mission and information security program. Governance should include periodic review, and the service model may adjust the defined roles and responsibilities in collaborative information security governance and risk management (based on the respective scope of control for user and provider), while the deployment model may define accountability and expectations (based on risk assessment).
- Customers of cloud services should ask whether their own management has defined risk tolerances with respect to cloud services and accepted any residual risk of utilizing cloud services.
- Where a provider cannot demonstrate comprehensive and effective risk management processes in association with its services, customers should carefully evaluate use of the vendor as well as the user's own abilities to compensate for the potential risk management gaps.
- Organizations should define risk metrics for engaging providers based on business and technical exposures. These metrics could include the type of data covered, the variety of user types relating to the information, and the vendors and other counterparties involved.

2.5 Requirements

- ✓ Provide transparency to stakeholders and shareholders demonstrating fiscal solvency and organizational transparency.
- ✓ Respect the interdependency of the risks inherent in the cloud supply chain and communicate the corporate risk posture and readiness to consumers and dependant parties.
- ✓ Inspect and account for risks inherited from other members of the cloud supply chain and take active measures to mitigate and contain risks through operational resiliency.

DOMAIN 3 //
LEGAL ISSUES: CONTRACTS AND ELECTRONIC DISCOVERY

This domain highlights some of the legal aspects raised by cloud computing. It provides general background on legal issues that can be raised by moving data to the cloud, some issues for consideration in a cloud services agreement, and the special issues presented by electronic discovery under Western litigation.

This domain provides an overview of selected issues and it is not a substitute for obtaining legal advice.

Overview. This domain will address the following topics:

- Summary of specific legal issues raised by moving data to the cloud
- Considerations for a cloud services agreement
- Special issues raised by e-discovery

3.1 Legal Issues

In many countries throughout the world, numerous laws, regulations, and other mandates require public and private organizations to protect the privacy of personal data and the security of information and computer systems. For example, in the Asia Pacific region, Japan, Australia, New Zealand, and many others have adopted data protection laws that require the data controller to adopt reasonable technical, physical, and administrative measures in order to protect personal data from loss, misuse, or alteration, based on the Privacy and Security Guidelines of the Organization for Economic Cooperation and Development (**OECD**) [17], and the Asia Pacific Economic Cooperation's (**APEC**)[18] Privacy Framework.

In Europe, the European Economic Area (**EEA**)[19] Member States have enacted data protection laws that follow the principles set forth in the 1995 European Union (EU) Data Protection Directive[20] and the 2002 ePrivacy Directive (as amended in 2009). These laws include a security component, and the obligation to provide adequate security must be passed down to subcontractors. Other countries that have close ties with the EEA, such as Morocco and Tunisia in Africa, Israel and Dubai in the Middle East have also adopted similar laws that follow the same principles.

North, Central, and South America countries are also adopting data protection laws at a rapid pace. Each of these laws includes a security requirement and places on the data custodian the burden of ensuring the protection and security of personal data wherever the data are located, and especially when transferring to a third party. For example, in addition to the data protection laws of Canada, Argentina and Colombia, which have been in existence for several years, Mexico, Uruguay, and Peru have recently passed data protection laws that are inspired mainly from the European model and may include references to the APEC Privacy Framework as well.

[17] **OECD** - Organization for Economic Cooperation and Development
[18] **APEC** - Asia Pacific Economic Cooperation
[19] **EEA** - European Economic Area
[20] EU Directive 95/46/EC

In Japan, the Personal Information Protection Act requires the private sectors to protect personal information and data securely. In the healthcare industry, profession-specific laws, such as the Medical Practitioners' Act, the Act on Public Health Nurses, Midwives and Nurses, and the Pharmacist Act, require registered health professionals for confidentiality of patient information.

Organizations that do business in the United States may be subject to one or more data protection laws. The laws hold organizations responsible for the acts of their subcontractors. For example, the security and privacy rules under the Gramm-Leach-Bliley Act (**GLBA**) [21] or the Health Insurance Portability and Accountability Act of 1996 (**HIPAA**) require that organizations compel their subcontractors, in written contracts, to use reasonable security measures and comply with data privacy provisions. Government agencies, such as the Federal Trade Commission (FTC) or the State Attorneys General have consistently held organizations liable for the activities of their subcontractors. The Payment Card Industry (PCI) Data Security Standards (DSS), which apply to credit card data anywhere in the world, including data processed by subcontractors has similar requirements.

The following sections provide examples of legal issues that may arise in connection with the transfer of personal data to the cloud or the processing of personal data in the cloud.

Table 1 — Obligatory Predicates

ISSUE	DESCRIPTION
U.S. Federal Laws	Numerous federal laws and their related regulations, such as GLBA, HIPAA, Children's Online Privacy Protection Act of 1998 ("COPPA"), together with orders issued by the FTC, require companies to adopt specific privacy and security measures when processing data, to require similar precautions in their contracts with the third party service provider.
U.S. State Laws	Numerous state laws also create an obligation on companies to provide adequate security for personal data and to require their service providers to do the same. 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 and state laws. Companies that are subject to PCI DSS or ISO 27001 must both comply with specified standards and pass onto their subcontractors the same obligation to meet the standard to which they are subject.
International Regulations	Many countries have adopted data protection laws that follow the European Union model, the OECD model or the APEC 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 when third parties process the data. The data controller is required to ensure that any third party

[21] **GLBA** - Gramm-Leach-Billey Act

	processing personal data on its behalf takes adequate technical and organizational security measures to safeguard the data.
Contractual Obligations	Even if a specific activity is not regulated, companies may have a contractual obligation to protect the personal information of their clients, contacts or employees, to ensure that the data are not used for secondary uses, and are not disclosed to third parties. This obligation may stem, for example, from the Terms and Conditions and Privacy Statement that a company post on its website. Alternately, the company may have entered into contracts (such as service agreements) with its customers, in which it has made specific commitments to protect the data (personal data or company data), limit their use, ensure their security, use encryption, etc. The organization must ensure that, when data in its custody are hosted in the cloud, it will have the continued ability to meet the promises and commitments that it made in its privacy notice(s) or other contracts. For example, the company may have agreed to make only specific uses of the data. Data in the cloud must be used only for the purposes for which they were collected. 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 also allow these access, modification and deletion rights to be exercised to the same extent as it would in a non-cloud relationship.
Prohibition against cross border transfers	Many laws, throughout the world, prohibit or restrict the transfer of information out of the country. In most cases, the transfer is permitted only if the country to which the data are transferred offers an adequate protection of personal information and privacy rights. The purpose of this adequacy requirement is to ensure that the individual data subjects whose data are transferred across borders will be able to enjoy, in the new country where their data were transferred, privacy rights and privacy protections that are similar to, and not less than, those that were afforded to them before the transfer. Thus, it is important for a cloud user 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. Depending on the country, the requirements for ensuring this adequate protection may be complex and stringent. In some cases, it may be necessary to obtain prior permission of the local Data Protection Commissioner.

3.2 Contract Considerations

When data is transferred to a cloud, the responsibility for protecting and securing the data typically remains with the collector or custodian of that data, even if in some circumstances, this responsibility may be shared with others. When it relies on a third party to host or process its data, the custodian of the data remains liable for any loss, damage, or

misuse of the data. It is prudent, and may be legally required, that the data custodian and the cloud provider enter into a written (legal) agreement that clearly defines the roles, expectations of the parties, and allocates between them the many responsibilities that are attached to the data at stake.

The laws, regulations, standards and the related best practices discussed above, also require data custodians 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).

3.2.1 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 relinquishment of control over the company data is restricted by law or creates serious security concerns.

In addition, the company should—and in some cases may be legally required to—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.

3.2.2 Contract

The parties must enter into a written contract. Depending on the nature of the services, the contract may commonly 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.

If issues are not addressed in the contract, the cloud service customer should consider alternate means of achieving the goal, an alternate provider, or not sending the data to the cloud. For example, if the cloud service customer wishes to send HIPAA-covered information to the cloud, the customer will need to find a cloud service provider that will sign a HIPAA business associate agreement or else not send that data to the cloud.

Below are brief descriptions of some cloud-specific issues. In addition, the attached checklist provides a comprehensive (but not exhaustive) list of issues to consider when reviewing a cloud services contract.

3.2.3 Monitoring, Testing and Updating

The cloud environment is not static. It evolves, and the parties must adapt. Periodic monitoring, testing, and evaluation of the services are recommended, in order to ensure that the required privacy and security measures are being used, and the processes and policies are being followed.

In addition, the legal, regulatory, and technical landscape is likely to change at a rapid pace. New security threats, new laws, new compliance requirements must be addressed promptly. The parties must keep abreast of the legal and other requirements and ensure that the operations remain compliant with applicable laws, and that the security measures in place keep evolving as new technologies and new laws emerge.

Cloud Audit and Cloud Trust Protocol are two mechanisms to automate monitoring and testing of cloud supply chains. In addition, the ITU-T is working on an X.1500 Cloud Auditing specification referred to as CYBEX.

3.3 Special Issues Raised by E-Discovery

This section addresses the unique requirements of litigation in the United States. U.S. litigants rely heavily on documents when arguing their case. One of the particularities of the American judicial system – in great contrast to most other countries – is that a US litigant must provide its adversary with ALL documents that pertain to the case. It must not only provide the documents that are favorable to its case, but also the documents that are favorable to the other litigant.

In recent years, there have been numerous scandals where litigants were accused to have voluntarily deleted, lost, or modified important evidence that was detrimental to their case. As a result, the rules of procedures have been changed to clarify the obligations of the parties, especially in the case of electronically stored information or "ESI."

Since the cloud will become the repository of most ESI that is needed in a litigation or investigation, cloud service providers and their clients must carefully plan how they will be able to identify all documents that pertain to a case in order to be able to fulfill the stringent requirements imposed by the E-Discovery provisions of the Federal Rules of Civil Procedure, and the State equivalents to these laws.

In this regard, the cloud service client and provider need to consider the following issues in matters when a client is subject to a discovery request and potentially relevant data exists with the cloud provider.

3.3.1 Possession, Custody, and Control

In most jurisdictions in the United States, a party's obligation to produce relevant information is limited to documents and data within its possession, custody or control. Hosting relevant data at a third-party, even a cloud provider, generally does not obviate a party's obligation to produce information as it may have a legal right to access or obtain the data. However, not all data hosted by a cloud provider may be in the control of a client (e.g., disaster recovery systems, certain metadata created and maintained by the cloud provider to operate its environment). Distinguishing the data that is and is not available to the client may be in the interest of the client and provider. The obligations of the cloud service provider as cloud data handler with regard to the production of information in response to legal process is an issue left to each jurisdiction to resolve.

3.3.2 Relevant Cloud Applications and Environment

In certain litigations and investigations, the actual cloud application or environment could itself be relevant to resolving the dispute in the litigation or investigation. In these circumstances, the application and environment will likely be outside the control of the client and require a subpoena or other discovery process on the provider directly.

3.3.3 Searchability and E-Discovery Tools

Because of the cloud environment, a client may not be able to apply or use e-discovery tools that it uses in its own environment. Moreover, a client may not have the ability or administrative rights to search or access all of the data hosted in the cloud. For example, where a client could access multiple employees' e-mail accounts on its own server at once, it may not have this ability with e-mail accounts hosted in the cloud. As such, clients need to account for the potential additional time, and expense, that this limited access will cause.

3.3.4 Preservation

Generally speaking, in the United States, a party is obligated to undertake reasonable steps to prevent the destruction or modification of data or information in its possession, custody, or control that it knows, or reasonably should know, is relevant to a pending or reasonably anticipated litigation or government investigation. Depending on the cloud service and deployment model that a client is using, preservation in the cloud can be very similar to preservation in other IT infrastructures, or it can be significantly more complex.

In the European Union, information preservation is governed under Directive 2006/24/EC of the European Parliament and of the Council of 15 March 2006. Japan, South Korea, and Singapore have similar data protection initiatives. Within South America, Brazil and Argentina have the Azeredo Bill, and the Argentina Data Retention Law 2004, Law No. 25.873, 6 February 2004, respectively.

3.3.4.1 Costs and Storage

Preservation can require that large volumes of data be retained for extended periods. What are the ramifications of this under the service level agreement ("SLA")? What happens if the preservation requirements outlast the terms of the SLA? If the client preserves the data in place, who pays for the extended storage and at what cost? Does the client have the storage capacity under its SLA? Can the client effectively download the data in a forensically sound manner so it can preserve it off-line or near-line?

3.3.4.2 Scope of Preservation

Absent good cause or a specific need, a requesting party is only entitled to data that is hosted in the cloud that contains relevant information, not all the data in the cloud or in the application. However, if the client does not have the ability to preserve relevant information or data in a granular way, it may be required to over-preserve in order to effect reasonable preservation, depending on the litigation or investigation.

3.3.4.3 Dynamic and Shared Storage

The burden of preserving data in the cloud may be relatively modest if the client has space to hold it in place, the data is relatively static, and the people with access are limited and know to preserve it. However, in a cloud environment that programmatically modifies or purges data, or one where the data is shared with people unaware of the need to preserve, preservation can be more difficult. After a client determines that such data is relevant and needs to be preserved, the client may need to work with the provider to determine a reasonable way to preserve such data.

3.3.5 Collection

Because of the potential lack of administrative control a client has over its data in the cloud, collection from the cloud can be more difficult, more time-consuming, and more expensive than from behind a client's firewall. In particular, a client may not have the same level of visibility across its cloud data, and it may have more difficulty comparing the data it has collected with the data in the cloud to determine that export was reasonably complete and accurate.

3.3.5.1 Access and Bandwidth

In most cases, a client's access to its data in the cloud will be determined by its SLA. This may limit its ability to collect large volumes of data quickly and in a forensically sound manner (i.e., with all reasonably relevant metadata preserved). Clients and cloud providers are well served to consider this issue early and establish a protocol (and a cost) for extraordinary access in the case of litigation and investigations to allow for collection. Absent these agreements, clients should consider the extra time and cost implicated by collection in the cloud when making representations to requesting parties and courts.

3.3.5.2 Functionality

Related to access and bandwidth, but different. Clients' right of access may provide them access to a full range of data, but not provide them the degree of functionality that would best assist them in a given situation. By way of example, a client may have access to three years of retail transactional data, but may only be able to download data two weeks at a time due to functionality constraints. Moreover, a client may not have full view into all the metadata that actually exists, but rather only a more limited degree of metadata.

3.3.5.3 Forensics

Bit-by-bit imaging of a cloud data source is generally difficult or impossible. For obvious security reasons, providers are reluctant to allow access to their hardware, particularly in a multi-tenant environment where a client could gain access to other clients' data. Even in a private cloud, forensics may be extremely difficult, and clients may need to notify opposing counsel or the courts of these limitations. Luckily, forensics is rarely warranted in cloud computing, not because it is cloud computing, but because it is usually a structured data hierarchy or virtualization that does not lend itself to forensic analysis.

3.3.5.4 Reasonable Integrity

A client subject to a discovery request should undertake reasonable steps to validate that its collection from its cloud provider is complete and accurate, especially where ordinary business procedures are unavailable and litigation-specific measures are being used to obtain the information. This process is separate and apart from verifying, that the data stored in the cloud is accurate, authenticated, or admissible.

3.3.5.5 Not Reasonably Accessible

Because of differences in how a client's data is stored and the client's access rights and privileges, not all of a client's data in the cloud may be equally accessible. The client (and the provider) should analyze requests for information and the pertinent data structure for relevance, materiality, proportionality and accessibility.

3.3.6 Direct Access

Outside of the cloud environment, a requesting party's direct access to a responding party's IT environment is not favored. In the cloud environment, it is even less favored and may be impossible for the same reasons as forensics. Importantly, a client may not be able to provide direct access because the hardware and facilities are outside its possession, custody or control, and a requesting party would need to subpoena or negotiate directly with the provider.

3.3.7 Native Production

Cloud service providers often store data in highly proprietary systems and applications in the cloud that clients do not control. Production of data in this native format may be useless to requesting parties, as they will not be able to understand the information produced. In these circumstances, it may be best for all concerned – requesting party, producing party, and provider – that the relevant information be exported using standard reporting or exporting protocols that exist within the cloud environment.

3.3.8 Authentication

Authentication in this context refers to forensic authentication of data that is admitted into evidence. This should not be confused with user authentication, which is a component of Identity Management. Storing data in the cloud does not affect the analysis for authentication of the data to determine if it should be admitted into evidence. The question is whether the document is what it purports to be. An e-mail is no more or less authentic because it was stored behind a company's firewall or was stored in the cloud. The question is whether it was stored with integrity and the court can trust that it has not been altered since it was sent or received.

3.3.9 Admissibility and Credibility

Absent other evidence, such as tampering or hacking, documents should not be considered more or less admissible or credible merely because they were created or stored in the cloud.

3.3.10 Cooperation between Provider and Client in e-Discovery

It is in the best interests of both providers and clients to consider the complications caused by discovery at the beginning of their relationship and to account for it in their SLAs. Providers may want to consider designing their cloud offerings to include discovery services to attract clients ("Discovery by Design"). In any event, clients and providers should consider including an agreement to reasonably cooperate with each other in the event of discovery requests against either.

3.3.11 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.).

REFERENCES

International Treaties and Agreements

[1] OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data (1980).

[2] OECD Guidelines for the Security of Information Systems and Networks (2002).

[3] OECD Recommendation on Cross-border Cooperation in the Enforcement of Laws Protecting Privacy.

Publications

[4] GILBERT, Francoise. © 2009-2011. Global Privacy & Security. Aspen Publishing / Wolters Kluwer (2 volumes).

[5] GILBERT, Francoise. 2011. Cloud Service Providers Can Be Both Data Processors and Data Controllers (BNA Privacy & Security Law Report 10 PVLR 266 (2011). Journal of Internet Law, Volume 15, Number 2, page 3.

[6] POWER, Michael E. AND TROPE, Roland L. 2005. Sailing in Dangerous Waters: A Director's Guide to Data Governance. American Bar Association.

[7] SMEDINGHOFF, Thomas. 2008 Information Security Law: Emerging Standard for Corporate Compliance (ITGP).

Websites

[8] Cloud computing definitions and business models:
http://p2pfoundation.net/Cloud_ComputingDefinition (technical aspects, business models)

[9] 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 //
COMPLIANCE AND AUDIT MANAGEMENT

Organizations face new challenges as they migrate from traditional data centers to the cloud. Delivering, measuring, and communicating compliance with a multitude of regulations across multiple jurisdictions is one of the largest challenges. Customers and providers alike need to understand and appreciate the differences and implications on existing compliance and audit standards, processes, and practices. The distributed and virtualized nature of cloud requires significant framework adjustment from approaches based on definite and physical instantiations of information and processes.

Cloud has the potential to improve transparency and assurance, through its more centralized and consolidated management platforms. Moreover, the outsourced solutions from cloud providers reduce the scale-dependency of compliance. With providers able to deliver first-day compliant solutions, new firms (for-profit and non-profit) would be able to enter markets and take actions that would have been cost-prohibitive in a pre-cloud era. Governments and other organizations previously reluctant to outsource IT operations due to issues of security and compliance may be more ready to adopt a cloud model, where compliance can be partly addressed through contractual delegation.

In addition to providers and customers, regulators and auditors are also adjusting to the new world of cloud computing. Few existing regulations were written to account for virtualized environments or cloud deployments. A cloud consumer can be challenged to show auditors that the organization is in compliance. Understanding the interaction of cloud computing and the regulatory environment is a key component of any cloud strategy. Cloud customers must consider and understand the following:

- Regulatory implications for using a particular cloud service or providers, giving particular attention to any cross-border or multi-jurisdictional issues when applicable
- Assignment of compliance responsibilities between the provider and customer, including indirect providers (i.e., the cloud provider of your cloud provider)
- Provider capabilities for demonstrating compliance, including document generation, evidence production, and process compliance, in a timely manner
- Relationships between customer, providers and auditors (both the customer's and provider's) to ensure required (and appropriately restricted) access and alignment with governance requirements

Overview. This domain will address the following topics:

- Compliance
- Audit

4.1 Compliance

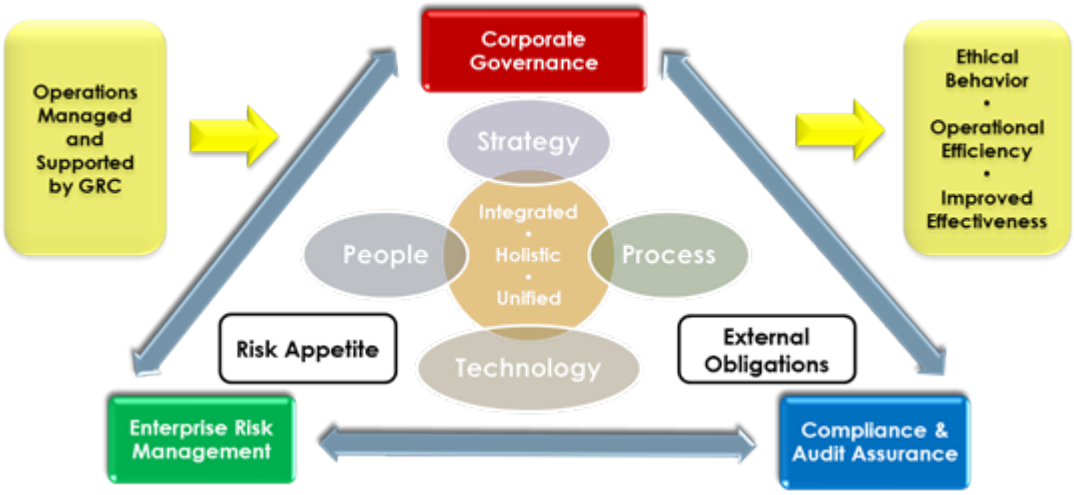

Figure 1—GRC Value Ecosystem

- **Corporate Governance:** the balance of control between stakeholders, directors and managers of an organization providing consistent management, cohesive application of policies, guidance and controls, and enabling effective decision-making
- **Enterprise Risk Management:** methods and processes (framework) used by organizations to balance decision-making based on identifying particular events or circumstances relevant to the organization's objectives (risks and opportunities), assessing them in terms of likelihood and magnitude of impact, determining a response strategy, and monitoring progress to protect and create value for their stakeholders
- **Compliance and Audit Assurance:** awareness and adherence to corporate obligations (e.g., corporate social responsibility, ethics, applicable laws, regulations, contracts, strategies and policies) by assessing the state of compliance, assessing the risks and potential costs of non-compliance against the costs to achieve compliance, and hence prioritize, fund, and initiate any corrective actions deemed necessary

Information technology in the cloud is increasingly subject to a plethora of policies and regulations. All stakeholders expect organizations to proactively comply with regulatory guidelines and requirements across multiple jurisdictions. IT governance is a necessity to deliver against these requirements and all organizations need a strategy to deliver.

Governance includes the processes and policies that enable the smooth execution of organizational objectives within the constraints of the external environment. Governance requires compliance activities to ensure that operations are fully aligned with those processes and policies. In this sense, compliance is focused on aligning with external requirements (e.g., law, regulation, industry standards) while governance is focused on aligning with internal requirements (e.g., board decisions, corporate policy).

Compliance can be defined as the awareness and adherence to obligations (e.g., corporate social responsibility, applicable laws, ethical guidelines), including the assessment and prioritization of corrective actions deemed necessary and appropriate. In some environments, particularly those highly regulated, the transparency aspect can even be dominant with reporting requirements getting more attention than compliance itself. In the best circumstances, compliance is not an inhibitor of organizational effectiveness, but a complement to internally determined policies.

Regulations typically have strong implications for information technology and its governance, particularly in terms of monitoring, management, protection, and disclosure). IT governance is a supporting element in overall corporate governance, enterprise risk management, compliance, and audit/assurance.

Cloud can be an enabling technology for governance and compliance, centralizing control and transparency through its management platforms, particularly for internally management cloud. By leveraging cloud services, sub-scale organizations can achieve the same level of compliance as much larger and highly resources entities. Security and assurance services are one way third-parties can play a role in compliance assessment and communication.

Any compliance approach will need to include participation across the organization, including IT. The role of external providers needs to be carefully considered, and responsibility for including them in governance, indirectly or directly, should be explicitly assigned within the customer organization.

In addition, the following represent a number of cloud security standards that are in development within ISO/IEC and ITU-T:

- ISO/IEC 27017: Cloud Computing Security and Privacy Management System-Security Controls
- ISO/IEC 27036-x: Multipart standard for the information security of supplier relationship management that is planned to include a part relevant to the cloud supply chain
- ITU-T X.ccsec: Security guideline for cloud computing in telecommunication area
- ITU-T X.srfcts: Security requirements and framework of cloud-based telecommunication service environment (X.srfcts)

ITU-T X.sfcse: Security functional requirements for Software as a Service (SaaS) application environment

4.2 Audit

Proper organizational governance naturally includes audit and assurance. Audit must be independently conducted and should be robustly designed to reflect best practice, appropriate resources, and tested protocols and standards.

Both internal and external audit and controls have legitimate roles to play for cloud, for both the customer and provider. Greater transparency may be best during initial stages of cloud introduction, to increase stakeholder comfort levels. An audit is one method to provide assurance that operational risk management activities are thoroughly tested and reviewed.

An audit plan should be adopted and supported by the most senior governing elements of the organization (e.g., the board and management). Regular and independent audits of critical systems and controls, including the accompanying audit trail and documentation will support improvements in efficiency and reliability.

Many organizations use a maturity model (e.g., CMM, PTQM) as a framework for analyzing process effectiveness. In some cases, a more statistical approach to risk management is adopted (e.g., Basel and Solvency accords for financial services) and as the field matures more specialized models for risk can be adopted as appropriate for the function or line of business.

For cloud, these practices will need to be revised and enhanced. Just as with previous models of information technology, audit will need to take advantage of the potential of cloud, as well as increase scope and scale to manage its novel aspects.

4.3 Recommendations

When engaging a provider, involve the appropriate legal, procurement, and contracts teams within the customer organization. The standard terms of services may not address compliance needs, and would need to be negotiated.

Specialized compliance requirements for highly regulated industries (e.g., finance, health care) should be considered when using a cloud service. Organizations who understand their current requirements should consider the impact of a distributed IT model, including the impact of cloud providers operating in diverse geographic locations and different legal jurisdictions.

Determine how existing compliance requirements will be impacted by the use of cloud services, for each workload (i.e., set of applications and data), in particular as they relate to information security. As with any outsourced solution, organizations need to understand which of their cloud partners are and should be processing regulated information. Examples of impacted policies and procedures include activity reporting, logging, data retention, incident response, controls testing, and privacy policies.

Understand the contractual responsibilities of each party. The baseline expectations will vary by deployment model with the customer having more control and responsibility in an IaaS model, and the provider having the dominant role for SaaS solutions. Particularly important is chained requirements and obligations – not just the customer to their direct cloud provider, but between the end customer and the provider's cloud provider.

Compliance with laws and industry regulation and its requirement (i.e. laws, technical, legal, compliance, risk, and security) is critical and must address during requirements identification stage. Any information processed, transmitted, stored, or viewed that is identified as Personal Identifiable Information (**PII**)[22] or private information faces a plethora of compliance regulation worldwide that may vary country or state. Since cloud was designed to be geographically diverse and scalable, solution data may be stored, processed, transmitted, or retrieved from many locations or multiple data centers of CSP. Some regulatory requirements specify controls that are difficult or impossible to achieve in certain cloud service types (e.g., geographic requirements may be inconsistent with distribute storage). Customers and providers must agree how to collect, store, and share compliance evidence (e.g., audit logs, activity reports, system configurations).

- Prefer auditors that are "cloud aware" that will be familiar with the assurance challenges (and advantages) of virtualization and cloud.
- Request cloud Provider's SSAE 16 SOC2 or ISAE 3402 Type 2 report. These will provide a recognizable starting point of reference for auditors and assessors.
- Contracts should provide for third-party review of SLA metrics and compliance (e.g., by a mutually-selected mediator).

[22] **PII** - Personal Identifiable Information

4.3 Requirements

- ✓ A right to audit clause gives customers the ability to audit the cloud provider, which supports traceability and transparency in the frequently evolving environments of cloud computing and regulation. Use a normative specification in the right to audit to ensure mutual understanding of expectations. In time, this right should be supplanted by third-party certifications (e.g., driven by ISO/IEC 27001/27017).
- ✓ A right to transparency clause with specified access rights can provide customers in highly regulated industries (including those in which non-compliance can be grounds for criminal prosecution) with required information. The agreement should distinguish between automated/direct access to information (e.g., logs, reports) and 'pushed' information (e.g., system architectures, audit reports).
- ✓ Providers should review, update, and publish their information security documents and GRC processes regularly (or as required). These should include vulnerability analysis and related remediation decisions and activities.
- ✓ Third-party auditors should be mutually disclosed or selected in advance, jointly by provider and customer.
- ✓ All parties should agree to use a common certification assurance framework (e.g., from ISO, COBIT) for IT governance and security controls.

DOMAIN 5 //
INFORMATION MANAGEMENT AND DATA SECURITY

The primary goal of information security is to protect the fundamental data that powers our systems and applications. As companies transition to cloud computing, the traditional methods of securing data are challenged by cloud-based architectures. Elasticity, multi-tenancy, new physical and logical architectures, and abstracted controls require new data security strategies. In many cloud deployments, users even transfer data to external — or even public — environments in ways that would have been unthinkable only a few years ago.

Managing information in the era of cloud computing is a daunting challenge that affects all organizations; even those that aren't seemingly actively engaged in cloud-based projects. It begins with managing internal data and cloud migrations and extends to securing information in diffuse, cross-organization applications and services. Information management and data security in the cloud era demand both new strategies and technical architectures. Fortunately not only do users have the tools and techniques needed, but the cloud transition even creates opportunities to better secure data in our traditional infrastructure.

The authors recommend using a Data Security Lifecycle (explored below) for evaluating and defining cloud data security strategy. This should be layered with clear information governance policies, and then enforced by key technologies such as encryption and specialized monitoring tools.

Overview. This domain includes three sections:

- Section 1 provides background material on cloud information (storage) architectures.
- Section 2 includes best practices for information management, including the Data Security Lifecycle.
- Section 3 details specific data security controls, and when to use them.

5.1 Cloud Information Architectures

Cloud information architectures are as diverse as the cloud architectures themselves. While this section can't possibly cover all potential permutations, there are certain consistent architectures within most cloud services.

5.1.1 Infrastructure as a Service

IaaS, for public or private cloud, generally includes the following storage options:

- **Raw storage.** This includes the physical media where data is stored. May be mapped for direct access in certain private cloud configurations.
- **Volume storage.** This includes volumes attached to IaaS instances, typically as a *virtual hard drive*. Volumes often use *data dispersion* to support resiliency and security.

- **Object storage.** Object storage is sometimes referred to as file storage. Rather than a virtual hard drive, object storage is more like a file share accessed via **API's**[23] or web interface.
- **Content Delivery Network.** Content is stored in object storage, which is then distributed to multiple geographically distributed nodes to improve Internet consumption speeds.

5.1.2 Platform as a Service

PaaS both provides and relies on a very wide range of storage options.

PaaS may provide:

- **Database as a Service.** A multitenant database architecture that is directly consumable as a service. Users consume the database via APIs or direct **SQL**[24] calls, depending on the offering. Each customer's data is segregated and isolated from other tenants. Databases may be relational, flat, or any other common structure.
- **Hadoop/MapReduce/Big Data as a Service.** *Big Data* is data whose large scale, broad distribution, heterogeneity, and currency/timeliness require the use of new technical architectures and analytics. *Hadoop* and other *Big Data* applications may be offered as a cloud platform. Data is typically stored in *Object Storage* or another distributed file system. Data typically needs to be close to the processing environment, and may be moved temporally as needed for processing.
- **Application storage.** Application storage includes any storage options built into a PaaS application platform and consumable via API's that doesn't fall into other storage categories.

PaaS may consume:

- **Databases.** Information and content may be directly stored in the database (as text or binary objects) or as files referenced by the database. The database itself may be a collection of IaaS instances sharing common back-end storage.
- **Object/File Storage.** Files or other data are stored in object storage, but only accessed via the PaaS API.
- **Volume Storage.** Data may be stored in IaaS volumes attached to instances dedicated to providing the PaaS service.
- **Other.** These are the most common storage models, but this is a dynamic area and other options may be available.

5.1.3 Software as a Service

As with PaaS, SaaS uses a very wide range of storage and consumption models. SaaS storage is always accessed via a web-based user interface or client/server application. If the storage is accessible via API then it's considered PaaS. Many SaaS providers also offer these PaaS APIs.

[23] **API** - Application Program Interface
[24] **SQL** - Structural Query Language is programming language designed for managing data

SaaS may provide:

- **Information Storage and Management.** Data is entered into the system via the web interface and stored within the SaaS application (usually a back-end database). Some SaaS services offer data set upload options, or PaaS API's.
- **Content/File Storage.** File-based content is stored within the SaaS application (e.g., reports, image files, documents) and made accessible via the web-based user interface.

SaaS may consume:

- **Databases.** Like PaaS, a large number of SaaS services rely on database back-ends, even for file storage.
- **Object/File Storage.** Files or other data are stored in object storage, but only accessed via the SaaS application.
- **Volume Storage.** Data may be stored in IaaS volumes attached to instances dedicated to providing the SaaS service.

5.2 Data (Information) Dispersion

Data (Information) Dispersion is a technique that is commonly used to improve data security, but without the use of encryption mechanisms. These sorts of algorithms (**IDA**[25] for short) are capable of providing high availability and assurance for data stored in the cloud, by means of data fragmentation, and are common in many cloud platforms. In a fragmentation scheme, a file *f* is split into *n* fragments; all of these are signed and distributed to *n* remote servers. The user then can reconstruct *f* by accessing *m* arbitrarily chosen fragments. The fragmentation mechanism can also be used for storing long-lived data in the cloud with high assurance.

When fragmentation is used along with encryption, data security is enhanced: an adversary has to compromise *m* cloud nodes in order to retrieve *m* fragments of the file *f*, and then has to break the encryption mechanism being used.

5.3 Information Management

Before we can discuss specific data security controls, we need a model to understand and manage our information. *Information management* includes the processes and policies for both understanding how your information is used, and governing that usage. In the *data security* section, specific technical controls and recommendations are discussed to monitor and enforce this governance.

5.4 The Data Security Lifecycle

Although *Information Lifecycle Management* is a fairly mature field, it doesn't map well to the needs of security professionals. The Data Security Lifecycle is different from Information Lifecycle Management, reflecting the different needs of the security audience. This is a summary of the lifecycle, and a complete version is available at http://www.securosis.com/blog/data-security-lifecycle-2.0

[25] **IDA** - Intrusion Detection Algorithms

The lifecycle includes six phases from creation to destruction. Although it is shown as a linear progression, once created, data can bounce between phases without restriction, and may not pass through all stages (for example, not all data is eventually destroyed).

1. **Create.** Creation is the generation of new digital content, or the alteration/updating/modifying of existing content.
2. **Store.** Storing is the act committing the digital data to some sort of storage repository and typically occurs nearly simultaneously with creation.
3. **Use.** Data is viewed, processed, or otherwise used in some sort of activity, not including modification.
4. **Share.** Information is made accessible to others, such as between users, to customers, and to partners.
5. **Archive.** Data leaves active use and enters long-term storage.
6. **Destroy.** Data is permanently destroyed using physical or digital means (*e.g.,* cryptoshredding).

Figure 1—Data Lifecycle

5.4.1 Locations and Access

The lifecycle represents the phases information passes through but doesn't address its location or how it is accessed.

Locations

This can be illustrated by thinking of the lifecycle not as a single, linear operation, but as a series of smaller lifecycles running in different operating environments. At nearly any phase data can move into, out of, and between these environments.

Due to all the potential regulatory, contractual, and other jurisdictional issues it is extremely important to understand both the logical and physical locations of data.

Access

When users know where the data lives and how it moves, they need to know who is accessing it and how. There are two factors here:

1. Who accesses the data?

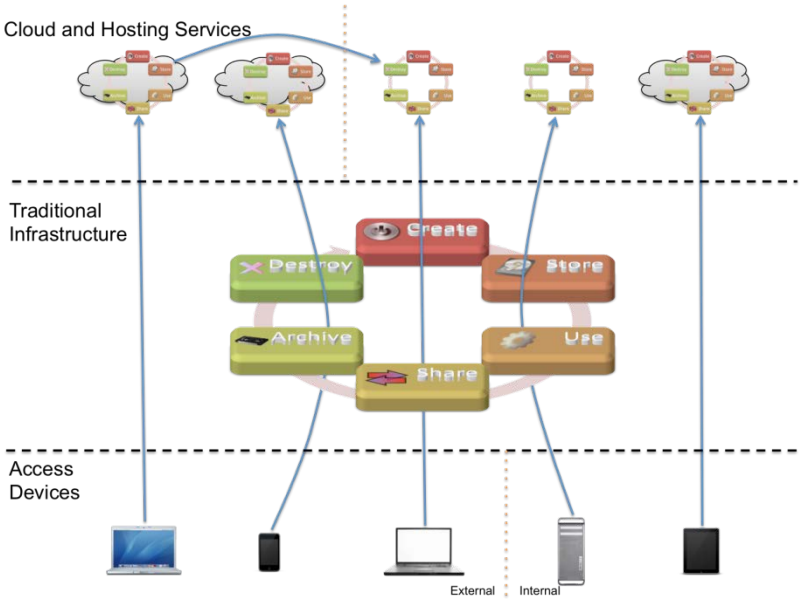

Figure 2—Cloud Access Devices

2. How can they access it (device & channel)?

Data today is accessed using a variety of different devices. These devices have different security characteristics and may use different applications or clients.

5.4.2 Functions, Actors, and Controls

The next step identifies the *functions* that can be performed with the data, by a given actor (person or system) and a particular location.

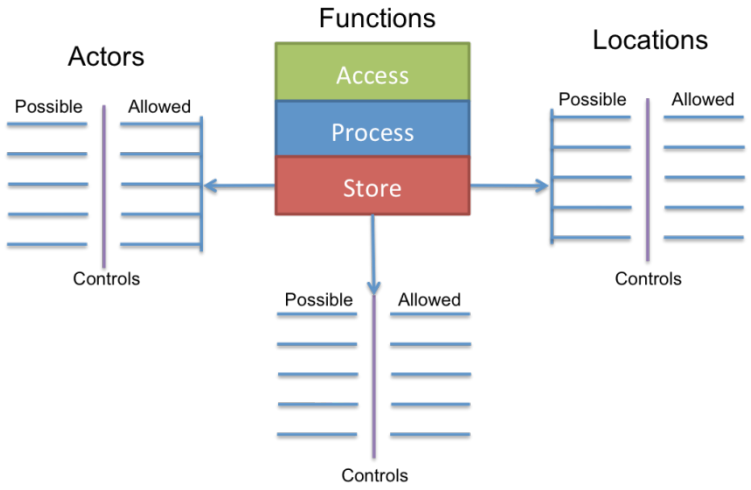

Figure 3—Functions vs. Controls

Functions

There are three things we can do with a given datum:

- **Access.** View/access the data, including creating, copying, file transfers, dissemination, and other exchanges of information.
- **Process.** Perform a transaction on the data: update it; use it in a business processing transaction, etc.
- **Store.** Hold the data (in a file, database, etc.).

The table below shows which functions map to which phases of the lifecycle:

Table 1—Information Lifecycle Phases

	Create	Store	Use	Share	Archive	Destroy
Access	X	X	X	X	X	X
Process	X		X			
Store		X			X	

An *actor* (person, application, or system/process, as opposed to the access device) performs each function in a *location*.

Controls

A *control* restricts a list of *possible* actions down to *allowed* actions. The table below shows one way to list the possibilities, which the user then maps to controls.

Table 2—Possible and Allowed Controls

Function		Actor		Location	
Possible	Allowed	Possible	Allowed	Possible	Allowed

5.5 Information Governance

Information governance includes the policies and procedures for managing information usage. It includes the following key features:

- **Information Classification.** High-level descriptions of important information categories. Unlike with *data classification* the goal isn't to label every piece of data in the organization, but rather to define high-level categories like "regulated" and "trade secret" to determine which security controls may apply.
- **Information Management Policies.** Policies to define what activities are allowed for different information types.
- **Location and Jurisdictional Polices.** Where data may be geographically located, which also has important legal and regulatory ramifications.
- **Authorizations.** Define which types of employees/users are allowed to access which types of information.
- **Ownership.** Who is ultimately responsible for the information.
- **Custodianship.** Who is responsible for managing the information, at the bequest of the owner.

5.6 Data Security

Data security includes the specific controls and technologies used to enforce information governance. This has been broken out into three sections to cover detection (and prevention) of data migrating to the cloud, protecting data in transit to the cloud and between different providers/environments, and protecting data once it's within the cloud.

5.6.1 Detecting and Preventing Data Migrations to the Cloud

A common challenge organizations face with the cloud is managing data. Many organizations report individuals or business units moving often sensitive data to cloud services without the approval or even notification of IT or security.

Aside from traditional data security controls (like access controls or encryption), there are two other steps to help manage unapproved data moving to cloud services:

1. Monitor for large internal data migrations with Database Activity Monitoring (**DAM**) [26]and File Activity Monitoring (**FAM**)[27].
2. Monitor for data moving to the cloud with URL filters and Data Loss Prevention.

[26] **DAM** - Database Activity Monitoring
[27] **FAM** - File Activity Monitoring

Internal Data Migrations

Before data can move to the cloud it needs to be pulled from its existing repository. Database Activity Monitoring can detect when an administrator or other user pulls a large data set or replicates a database, which could indicate a migration.

File Activity Monitoring provides similar protection for file repositories, such as file shares.

Movement to the Cloud

A combination of URL filtering (web content security gateways) and Data Loss Prevention (DLP) can detect data moving from the enterprise into the cloud.

URL filtering allows you to monitor (and prevent) users connecting to cloud services. Since the administrative interfaces for these services typically use different addresses than the consumption side, the user can distinguish between someone accessing an administrative console versus a user accessing an application already hosted with the provider.

Look for a tool that offers a cloud services list and keeps it up to date, as opposed to one that requires creating a custom category, and the user managing the destination addresses.

For greater granularity, use Data Loss Prevention. DLP tools look at the actual data/content being transmitted, not just the destination. Thus the user can generate alerts (or block) based on the classification of the data. For example, the user can allow corporate private data to go to an approved cloud service but block the same content from migrating to an unapproved service.

The insertion point of the DLP solution can determine how successfully data leakage can be detected. For example, availability of cloud solutions to various users (e.g., employees, vendors, customers) outside of the corporate network environment avoids or nullifies any DLP solutions if they are inserted at the corporate boundary.

5.6.2 Protecting Data Moving To (And Within) the Cloud

In both public and private cloud deployments, and throughout the different service models, it's important to protect data in transit. This includes:

- Data moving from traditional infrastructure to cloud providers, including public/private, internal/external and other permutations.
- Data moving between cloud providers.
- Data moving between instances (or other components) within a given cloud.

There are three options (or order of preference):

1. **Client/Application Encryption.** Data is encrypted on the endpoint or server before being sent across the network or is already stored in a suitable encrypted format. This includes local client (agent-based) encryption (e.g., for stored files) or encryption integrated in applications.

2. **Link/Network Encryption.** Standard network encryption techniques including SSL, VPNs, and SSH. Can be hardware or software. End to end is preferable but may not be viable in all architectures.

3. **Proxy-Based Encryption.** Data is transmitted to a proxy appliance or server, which encrypts before sending further on the network. Often a preferred option for integrating into legacy applications but is not generally recommended.

5.6.3 Protecting Data in the Cloud

With such a wide range of options and technologies available in cloud computing, there is no way to cover all possible security options. The following are some of the more useful technologies and best practices for securing data within various cloud models.

5.6.3.1 Content Discovery

Content discovery includes the tools and processes to identify sensitive information in storage. It allows the organization to define policies based on information type, structure, or classification and then scans stored data using advanced content analysis techniques to identify locations and policy violations.

Content discovery is normally a feature of Data Loss Prevention tools; for databases, it is sometimes available in Database Activity Monitoring products. Scanning can be via accessing file shares or a local agent running on an operating system. The tool must be "cloud aware" and capable of working within your cloud environment (e.g., able to scan object storage). Content discovery may also be available as a managed service.

5.6.3.2 IaaS Encryption

5.6.3.2.1 Volume Storage Encryption

Volume encryption protects from the following risks:

- Protects volumes from snapshot cloning/exposure
- Protects volumes from being explored by the cloud provider (and private cloud admins)
- Protects volumes from being exposed by physical loss of drives (more for compliance than a real-world security issue)

IaaS volumes can be encrypted using three methods:

- **Instance-managed encryption.** The encryption engine runs within the instance, and the key is stored in the volume but protected by a passphrase or keypair.
- **Externally managed encryption.** The encryption engine runs in the instance, but the keys are managed externally and issued to the instance on request.
- **Proxy encryption.** In this model you connect the volume to a special instance or appliance/software, and then connect your instance to the encryption instance. The proxy handles all crypto operations and may keep keys either onboard or external.

5.6.3.2.2 Object Storage Encryption

Object storage encryption protects from many of the same risks as volume storage. Since object storage is more often exposed to public networks, it also allows the user to implement *Virtual Private Storage*. Like a VPN, a **VPS**[28] allows use of a public shared infrastructure while still protecting data, since only those with the encryption keys can read the data even if it is otherwise exposed.

- **File/Folder encryption and Enterprise Digital Rights Management.** Use standard file/folder encryption tools or EDRM to encrypt the data before placing in object storage.
- **Client/Application encryption.** When object storage is used as the back-end for an application (including mobile applications), encrypt the data using an encryption engine embedded in the application or client.
- **Proxy encryption.** Data passes through an encryption proxy before being sent to object storage.

5.6.3.3 PaaS Encryption

Since PaaS is so diverse, the following list may not cover all potential options:

- **Client/application encryption.** Data is encrypted in the PaaS application or the client accessing the platform.
- **Database encryption.** Data is encrypted in the database using encryption built in and supported by the database platform.
- **Proxy encryption.** Data passes through an encryption proxy before being sent to the platform.
- **Other.** Additional options may include API's built into the platform, external encryption services, and other variations.

5.3.4.4 SaaS Encryption

SaaS providers may use any of the options previously discussed. It is recommended to use per-customer keys when possible to better enforce multi-tenancy isolation. The following options are for SaaS consumers:

- **Provider-managed encryption.** Data is encrypted in the SaaS application and generally managed by the provider.
- **Proxy encryption.** Data passes through an encryption proxy before being sent to the SaaS application.

Encryption operations should use whatever encryption method is most appropriate, which may include shared keys or public/private keypairs and an extensive **PKI/PKO** [29](Public Key Infrastructure/Operations) structure. Please see Domain 11 for more information on encryption and key management.

[28] **VPS** - Virtual Private Storage
[29] **PKI/PKO** - Public Key Infrastructure/Operations

5.3.5 Data Loss Prevention

Data Loss Prevention (DLP) is defined as:

Products that, based on central policies, identify, monitor, and protect data at rest, in motion, and in use, through deep content analysis.

DLP can provide options for how data found violation of policy is to be handled. Data can be blocked (stopping a workflow) or allowed to proceed after remediation by encryption using methods such as DRM, ZIP, or OpenPGP.

DLP is typically used for content discovery and to monitor data in motion using the following options:

- **Dedicated appliance/server.** Standard hardware placed at a network chokepoint between the cloud environment and the rest of the network/Internet or within different cloud segments.
- **Virtual appliance**
- **Endpoint agent**
- **Hypervisor-agent**. The DLP agent is embedded or accessed at the hypervisor level, as opposed to running in the instance.
- **DLP SaaS.** DLP is integrated into a cloud service (e.g., hosted email) or offered as a standalone service (typically content discovery).

5.3.6 Database and File Activity Monitoring

Database Activity Monitoring (DAM) is defined as:

Database Activity Monitors capture and record, at a minimum, all Structured Query Language (SQL) activity in real time or near real time, including database administrator activity, across multiple database platforms; and can generate alerts on policy violations.

DAM supports near real time monitoring of database activity and alerts based on policy violations, such as SQL injection attacks or an administrator replicating the database without approval. DAM tools for cloud environments are typically agent-based connecting to a central collection server (which is typically virtualized). It is used with dedicated database instances for a single customer, although in the future may be available for PaaS.

File Activity Monitoring (FAM) is defined as:

Products that monitor and record all activity within designated file repositories at the user level, and generate alerts on policy violations.

FAM for cloud requires use of an endpoint agent or placing a physical appliance between the cloud storage and the cloud consumers.

5.3.7 Application Security

A large percentage of data exposures are the result of attacks at the application layer, particularly for web applications. Please see Domain 10 for more information on application security.

5.3.8 Privacy Preserving Storage

Almost all cloud-based storage systems require some authentication of participants (cloud user and/or CSP) to establish trust relations, either for only one endpoint of communication or for both. Although cryptographic certificates can offer sufficient security for many of these purposes, they do not typically cater to privacy because they are bound to the identity of a real person (cloud user). Any usage of such a certificate exposes the identity of the holder to the party requesting authentication. There are many scenarios (e.g., storage of Electronic Health Records) where the use of such certificates unnecessarily reveals the identity of their holder.

Over the past 10-15 years, a number of technologies have been developed to build systems in a way that they can be trusted, like normal cryptographic certificates, while at the same time protecting the privacy of their holder (i.e., hiding the real holder's identity). Such attribute-based credentials are issued just like ordinary cryptographic credentials (e.g., X.509 credentials) using a digital (secret) signature key. However, attribute-based credentials (ABCs) allow their holder to transform them into a new credential that contains only a subset of the attributes contained in the original credential. Still, these transformed credentials can be verified just like ordinary cryptographic credentials (using the public verification key of the issuer) and offer the same strong security.

5.3.9 Digital Rights Management (DRM)

At its core, Digital Rights Management encrypts content, and then applies a series of *rights*. Rights can be as simple as preventing copying, or as complex as specifying group or user-based restrictions on activities like cutting and pasting, emailing, changing the content, etc. Any application or system that works with DRM protected data must be able to interpret and implement the rights, which typically also means integrating with the key management system.

There are two broad categories of Digital Rights Management:

- **Consumer DRM** is used to protect broadly distributed content like audio, video, and electronic books destined for a mass audience. There are a variety of different technologies and standards, and the emphasis is on one-way distribution.
- **Enterprise DRM** is used to protect the content of an organization internally and with business partners. The emphasis is on more complex rights, policies, and integration within business environments and particularly with the corporate Directory Service.

Enterprise DRM can secure content stored in the cloud well but requires deep infrastructure integration. It's most useful for document based content management and distribution. Consumer DRM offers good protection for distributing content to customers but does not have a good track record with most technologies being cracked at some point.

5.4 Recommendations

- Understand the cloud storage architecture in use, which will help determine security risk and potential controls.
- Choose storage with data dispersion when available.
- Use the Data Security Lifecycle to identify security exposures and determine the most appropriate controls.
- Monitor key internal databases and file repositories with DAM and FAM to identify large data migrations, which could indicate data migrating to the cloud.
- Monitor employee Internet access with URL filtering and/or DLP tools to identify sensitive data moving to the cloud. Select tools that include predefined categories for cloud services. Consider using filtering to block unapproved activity.
- Encrypt all sensitive data moving to or within the cloud at the network layer, or at nodes before network transmission. This includes all service and deployment models.
- When using any data encryption, pay particular attention to key management (see Domain 11).
- Use content discovery to scan cloud storage and identify exposed sensitive data.
- Encrypt sensitive volumes in IaaS to limit exposure due to snapshots or unapproved administrator access. The specific technique will vary depending on operational needs.
- Encrypt sensitive data in object storage, usually with file/folder or client/agent encryption.
- Encrypt sensitive data in PaaS applications and storage. Application-level encryption is often the preferred option, especially since few cloud databases support native encryption.
- When using application encryption, keys should be stored external to the application whenever possible.
- If encryption is needed for SaaS, try to identify a provider that offers native encryption. Use proxy encryption if that isn't available and /or trust levels must be assured.
- Use DLP to identify sensitive data leaking from cloud deployments. It is typically only available for IaaS, and may not be viable for all public cloud providers.
- Monitor sensitive databases with DAM and generate alerts on security policy violations. Use a cloud-aware tool.
- Consider privacy preserving storage when offering infrastructure or applications where normal access could reveal sensitive user information.
- Remember that most large data security breaches are the result of poor application security.
- Cloud providers should not only follow these practices, but expose data security tools and options to their customers.

- Removal of data from a cloud vendor either due to expiry of contract or any other reason should be covered in detail while setting up the SLA. This should cover deletion of user accounts, migration or deletion of data from primary / redundant storage, transfer of keys, etc.

5.5 Requirements

- ✓ Use the Data Security Lifecycle to identify security exposures and determine the most appropriate controls.
- ✓ Due to all the potential regulatory, contractual, and other jurisdictional issues it is extremely important to understand both the logical and physical locations of data.
- ✓ Monitor employee Internet access with URL filtering and/or DLP tools to identify sensitive data moving to the cloud.
- ✓ Encrypt all sensitive data moving to or within the cloud at the network layer, or at nodes before network transmission.
- ✓ Encrypt sensitive volumes in IaaS to limit exposure due to snapshots or unapproved administrator access.
- ✓ Encrypt sensitive data in PaaS applications and storage.

REFERENCES

[1] RABIN, M. O. 1989. Efficient Dispersal of Information for Security, Load Balancing, and Fault Tolerance. J. ACM, 36(2), 335–348.

[2] SECUROSIS. 2011. The Data Security Lifecycle. http://www.securosis.com/blog/data-security-lifecycle-2.0

[3] SECUROSIS. 2011. Understanding and Selecting a Data Loss Prevention Solution. http://www.securosis.com/research/publication/report-data-loss-prevention-whitepaper

[4] SECUROSIS. 2008. Understanding and Selecting a Database Activity Monitoring solution. http://www.securosis.com/research/publication/report-selecting-a-database-activity-monitoring-solution/

[5] CHAUM, D. L. Feb. 1981. Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms. Communications of the ACM, 24 (2), 84-90.

DOMAIN 6 //
INTEROPERABILITY AND PORTABILITY

The advent of cloud computing offers unprecedented scalability to an organization's IT processing and administrative capability unlike those available in "traditional" in-house infrastructures. Almost instantaneously, additional capacity can be added, moved, or removed in response to dynamically changing processing needs. A new application support system can be initiated to meet increased demand in a matter of hours rather than weeks. Should demand fall back, the additional capacity can be shut down just as quickly with no surplus hardware now sitting idled. Gaining the benefits of this more elastic environment requires both interoperability and portability to be the design goals of any cloud-implemented system, from IaaS through to SaaS.

At one end of the scale, Interoperability and Portability allows you to scale a service across multiple disparate providers on a global scale and have that system operate and appear as one system. At the other end, Interoperability and Portability allows the easy movement of data and applications from one platform to another, or from one service provider to another.

Portability and interoperability are not considerations unique to cloud environments and their related security aspects are not new concepts brought about by cloud computing. However, the open and often shared processing environments that exist within the cloud bring a need for even greater precautions than are required for traditional processing models. Multi-tenancy means data and applications reside with data and applications of other companies and that access to confidential data (intended or unintended) is possible through shared platforms, shared storage, and shared networks.

This section defines the critical consideration which should be addressed when designing for portability and interoperability.

Overview. The following sections define Interoperability and Portability in terms of:

- An introduction to Interoperability
- Recommendations to ensure Interoperability
- An introduction to Portability
- Recommendations for Portability

6.1 An Introduction to Interoperability

Interoperability is the requirement for the components of a cloud eco-system to work together to achieve their intended result. In a cloud computing eco-system the components may well come from different sources, both cloud and traditional, public and private cloud implementations (known as hybrid-cloud). Interoperability mandates that those

components should be replaceable by new or different components from different providers and continue to work, as should the exchange of data between systems.

Businesses, over time, will make decisions that lead to the desire to change providers. Reasons for this desired change include:

- An unacceptable increase in cost at contract renewal time
- The ability to get the same service at a cheaper price
- A provider ceases business operations
- A provider suddenly closes one or more services being used without acceptable migration plans
- Unacceptable decrease in service quality, such as a failure to meet key performance requirements or achieve service level agreements (**SLA's**)[30]
- A business dispute between cloud customer and provider

A lack of interoperability (and also portability) can lead to being locked to a particular cloud service provider.

The degree to which interoperability can be achieved or maintained when considering a cloud project often will depend on the degree to which a cloud provider uses open, or published, architectures and standard protocols and standard **API's**[31]. Though many vendors of "open" and "standards based" cloud provision provide propriety hooks and extensions (e.g. Eucalyptus) and enhancements that can impede both interoperability and portability.

6.2 An Introduction to Portability

Portability defines the ease of ability to which application components are moved and reused elsewhere regardless of provider, platform, OS, infrastructure, location, storage, the format of data, or API's.

Portability and interoperability must be considered whether the cloud migration is to public, private, or hybrid cloud deployment solutions. They are important elements to consider for service model selection regardless of whether a migration strategy is to Software as a Service (SaaS), Platform as a Service (PaaS), or Infrastructure as a Service (IaaS).

Portability is a key aspect to consider when selecting cloud providers since it can both help prevent vendor lock-in and deliver business benefits by allowing identical cloud deployments to occur in different cloud provider solutions, either for the purposes of disaster recovery or for the global deployment of a distributed single solution.

Achieving portability for a cloud service is generally reliant on the two services operating in the same architectural octant of the Cloud Cube, as defined in Domain One. Where services operate in different octants, then porting a service usually means migrating the service back "in-house" before re-outsourcing it to an alternative cloud service.

[30] **SLA** - Service Level Agreement
[31] **API** - Application Program Interface

Failure to appropriately address portability and interoperability in a cloud migration may result in failure to achieve the desired benefits of moving to the cloud and can result in costly problems or project delays due to factors that should be avoided such as:

- Application, vendor, or provider lock-in – choice of a particular cloud solution may restrict the ability to move to another cloud offering or to another vendor
- Processing incompatibility and conflicts causing disruption of service – provider, platform, or application differences may expose incompatibilities that cause applications to malfunction within a different cloud infrastructure
- Unexpected application re-engineering or business process change – moving to a new cloud provider can introduce a need to rework how a process functions or require coding changes to retain original behaviors
- Costly data migration or data conversion — lack of interoperable and portable formats may lead to unplanned data changes when moving to a new provider
- Retraining or retooling new application or management software
- Loss of data or application security – different security policy or control, key management or data protection between providers may open undiscovered security gaps when moving to a new provider or platform

Moving services to the cloud is a form of outsourcing; the golden rule of outsourcing is "understand up-front and plan for how to exit the contract". Portability (and to an extent interoperability) should therefore be a key criterion of any organizations strategy to move into cloud services, allowing for a viable exit strategy to be developed.

6.3 Recommendations

6.3.1 Interoperability Recommendations

Hardware – Physical Computer Hardware

The hardware will inevitably vary or change over time and from provider to provider leaving unavoidable interoperability gaps if direct hardware access is required.

- Whenever possible, use virtualization to remove many hardware level concerns, remembering that virtualization doesn't necessarily remove all hardware concerns, especially on current systems.
- If hardware must be directly addressed, it is important to ensure that the same or better physical and administrative security controls exist when moving from one provider to another.

Physical Network Devices

The network devices including security devices will be different from service providers to service providers along with its API and configuration process.

- To maintain interoperability the Network physical hardware and network & security abstraction should be in virtual domain. As far as possible API's should have the same functionally.

Virtualization

While virtualization can help to remove concerns about physical hardware, distinct differences exist between common hypervisors (such as ZEN, VMware and others).

- Using open virtualization formats such as OVF to help ensure interoperability.
- Document and understand which specific virtualization hooks are used no matter the format. It still may not work on another hypervisor.

Frameworks

Different platform providers offer different cloud application frameworks and differences do exist between them that affect interoperability.

- Investigate the API's to determine where differences lie and plan for any changes necessary that may be required to application processing when moving to a new provider.
- Use open and published API's to ensure the broadest support for interoperability between components and to facilitate migrating applications and data should changing a service provider become necessary.
- Applications in the cloud often interoperate over the Internet and outages can be anticipated to occur. Determine how failure in one component (or a slow response) will impact others and avoid stateful dependencies that may risk system data integrity when a remote component fails.

Storage

Storage requirements will vary for different types of data. Structured data will most often require a database system, or require application specific formats. Unstructured data will typically follow any of a number of common application formats used by Word Processors, Spreadsheets and Presentation Managers. Here the concern should be to move data stored with one service to another seamlessly.

- Store unstructured data in an established portable format.
- Assess the need for encryption for the data in transit.
- Check for compatible database systems and assess conversion requirements if needed.

Security

Data and applications in the cloud reside on systems the user doesn't own and likely has only limited control over. A number of important items to consider for interoperable security include:

- Use SAML or WS-Security for authentication so the controls can be interoperable with other standards-based systems. See domain 12 for more detail.
- Encrypting data before it is placed into the cloud will ensure that it cannot be accessed inappropriately within cloud environments. See domain 11 for more detail on encryption.

- When encryption keys are in use, investigate how and where keys are stored to ensure access to existing encrypted data is retained. See domain 11 for more detail on key management.
- Understand your responsibilities and liabilities should a compromise occur due to unanticipated "gaps" in protection methods offered by your service provider.
- Log file information should be handled with the same level of security as all other data moving to the cloud. Ensure that log files are interoperable to ensure continuity of log analysis pre-and post move as well as compatibility with whatever log management system is in use.
- When completing a move ensure that all data, logs, and other information is deleted from the original system.

6.3.2 Portability Recommendations

There are a number of issues standing in the path of moving to the cloud. Portability considerations and recommendations that impact moving to the cloud include;

- **Service Level**. SLA's will differ across providers, and there is a need to understand how this may affect your ability to change providers.
- **Different architectures**. Systems in the cloud may reside on disparate platform architectures. It is important to be aware of how these will limit portability by understanding service and platform dependencies, which may include API's, hypervisors, application logic, and other restrictions.
- **Security integration**. Cloud systems introduce unique portability concerns for maintaining security, including:
- Authentication and identity mechanisms for user or process access to systems now must operate across all components of a cloud system. Using open standards for Identity such as SAML will help to ensure portability. Developing internal IAM system to support SAML assertions and internal system to accept SAML will aid future portability of system to the cloud.
- Encryption keys should be escrowed locally, and when possible maintained locally
- Metadata is an aspect of digital information that is often and easily overlooked as (typically) metadata is not directly visible when working with files and documents. Metadata becomes an important consideration in the cloud, because metadata moves with the document. When moving files and their metadata to new cloud environments ensure copies of file metadata are securely removed to prevent this information from remaining behind and opening a possible opportunity for compromise.

6.3.3 Recommendations for Different Cloud Models

There are a number of generic risks and recommendations that are common to all cloud models.

- When substituting cloud providers it is normal to expect resistance from the legacy cloud provider. This must be planned for in the contractual process as outlined in Domain 3, in your Business Continuity Program as outlined in Domain 7, and as a part of your overall governance in Domain 2.

- Understand the size of data sets hosted at a cloud provider. The sheer size of data may cause an interruption of service during a transition, or a longer transition period than anticipated. Many customers have found that using a courier to ship hard drives is faster than electronic transmission for large data sets.
- Document the security architecture and configuration of individual component security controls so they can be used to support internal audits, as well as to facilitate migration to new providers and aid the validation of the new environment.

Infrastructure as a Service (IaaS)

Where the responsibility of the cloud provider is to provide basic computing utilities such as storage, computing power, etc., the cloud customer is responsible for a majority of application design tasks related to interoperability. The cloud provider should provide standardized hardware and computing resources that can interact with various disparate systems with minimal efforts. The cloud provider should strictly adhere to industry standards to maintain interoperability. The provider should be able to support complex scenarios such as cloud brokerage, cloud bursting, hybrid clouds, multi- cloud federation, etc.

- Understand how virtual machine images can be captured and ported to new cloud providers and who may use different virtualization technologies. Example: Distributed Management Task Force (DMTF) Open Virtualization format (OVF).
- Identify and eliminate (or at least document) any provider-specific extensions to the virtual machine environment.
- Understand what practices are in place to make sure appropriate de-provisioning of VM images occurs after an application is ported from the cloud provider.
- Understand the practices used for decommissioning of disks and storage devices.
- Understand hardware/platform based dependencies that need to be identified before migration of the application/data.
- Ask for access to system logs, traces, and access and billing records from the legacy cloud provider.
- Identify options to resume or extend service with the legacy cloud provider in part or in whole if new service proves to be inferior.
- Determine if there are any management-level functions, interfaces, or API's being used that are incompatible with or unimplemented by the new provider.
- Understand costs involved for moving data to and from a cloud provider
- Determine what means are supported for moving data as efficiently to the cloud as possible through using standard capabilities such as data compression.
- Understand what security is provided and who maintains access to encryption keys.

Platform as a Service (PaaS)

The cloud provider is responsible to provide a platform on which the consumers can build their systems. They provide with a runtime environment and an integrated application stack. It allows developers to quickly develop and deploy custom applications on the offered platforms without the need to build the infrastructure. The cloud provider provides the entire infrastructure and its maintenance to its consumers.

- When possible, use platform components with a standard syntax, open API's, and open standards, e.g. Open Cloud Computing Interface (**OCCI**)[32]
- Understand what tools are available for secure data transfer, backup, and restore.
- Understand and document application components and modules specific to the PaaS provider and develop application architecture with layers of abstraction to minimize direct access to proprietary modules.
- Understand how base services like monitoring, logging, and auditing would transfer over to a new vendor.
- Understand what protections are provided for data placed into the cloud and data generated and maintained in the cloud.
- Understand control functions provided by the legacy cloud provider and how they would translate to the new provider.
- When migrating to a new platform, understand the impacts on performance and availability of the application and how these impacts will be measured.
- Understand how testing will be completed prior to and after migration to verify that the services or applications are operating correctly. Ensure that both provider and user responsibilities for testing are well known and documented.

Software as a Service (SaaS)

The cloud provider provides application capabilities over the cloud, and the client just manages his/her operations and the information flowing in and out of the system. The client needs a browser, and majority of the administrative at all the levels rests with the provider.

- Perform regular data extractions and backups to a format that is usable without the SaaS provider.
- Understand whether metadata can be preserved and migrated.
- If needed use data escrow services.
- Understand that any custom tools will have to be redeveloped, or the new vendor must provide those tools, or commit to port (and support) these tools.
- Review and audit to ensure the consistency of control effectiveness across old and new providers.

[32] **OCCI** - Open Cloud Computing Interface

- Ensure backups and other copies of logs, access records, and any other pertinent information which may be required for legal and compliance reasons can be migrated.
- Understand the management, monitoring, and reporting interfaces and their integration between environments.
- Test and evaluate all applications before migration, and dual run if feasible prior to cut-over.

Private Cloud

Private cloud is when the consumer runs a cloud environment / service within their enterprise or uses private cloud offering from the cloud providers (typically extending the internal network into a service providers hosting centre).

- Ensure interoperability exists between common hypervisors such as KVM, VMware, Xen.
- Ensure standard API's are used for management functions such as users and their privilege management, VM image management, Virtual Machine management, Virtual Network management, Service management, Storage management, Infrastructure management, Information Management, etc.

Public Cloud

Interoperability in public cloud means exposing most common cloud interfaces. They may be vendor specific or open specifications and interfaces such as OCCI, libcloud, etc.

- Ensure that the cloud providers expose common and/or open interfaces to access all cloud functions in their service offering.

Hybrid Cloud

In this scenario the consumer's local private infrastructure should have the capability to work with external cloud providers. A common scenario is "cloud bursting", where an enterprise shares the load with external cloud providers to meet peak demands.

- Ensure that the cloud providers expose common and/or open interfaces to access all cloud functions in their service offering.
- Ensure the ability to federate with different cloud providers to enable higher levels of scalability.

REFERENCES

[1] http://msdn.microsoft.com/en-us/library/cc836393.aspx

[2] http://blogs.msdn.com/b/eugeniop/archive/2010/01/12/adfs-wif-on-amazon-ec2.aspx

[3] http://download.microsoft.com/download/6/C/2/6C2DBA25-C4D3-474B-8977-E7D296FBFE71/EC2-Windows%20SSO%20v1%200--Chappell.pdf

[4] http://www.zimbio.com/SC+Magazine/articles/6P3njtcIjmR/Federation+2+0+identity+ecosystem

[5] http://www.datacenterknowledge.com/archives/2009/07/27/cloud-brokers-the-next-big-opportunity/

[6] http://blogs.oracle.com/identity/entry/cloud_computing_identity_and_access

[7] http://www.opengroup.org/jericho/cloud_cube_model_v1.0.pdf

[8] http://www.burtongroup.com

[9] http://www.pkware.com/appnote

[10] http://www.apps.ietf.org/rfc/rfc4880.html

SECTION III // OPERATING IN THE CLOUD

DOMAIN 7 //
TRADITIONAL SECURITY, BUSINESS CONTINUITY, & DISASTER RECOVERY

With the emergence of cloud computing as a preferred technology for outsourcing IT operations, the security issues inherent in the hosting model have assumed greater significance and criticality. Inherent in the concept of cloud computing are the risks associated with entrusting confidential and sensitive data to third parties or pure-play cloud service providers (**CSP**)[33].

The evolution of cloud services has enabled business entities to do more with less: fewer resources and better operating efficiency. This has many tangible benefits for business, yet there are inherent security risks that must be evaluated, addressed, and resolved before businesses will have confidence in securely outsourcing their IT requirements to cloud service providers.

This section maps to Cloud Control Matrix Domains IS-01 and IS-02 as well as ISO/IEC 27002 Clause 9.

One purpose of this domain is to assist cloud service users to share a common understanding of traditional security (physical security) with cloud service. Traditional security can be defined as the measures taken to ensure the safety and material existence of data and personnel against theft, espionage, sabotage, or harm. In the context of cloud information security, this is about information, products, and people.

Proper information security deploys many different layers to achieve its goal. This is referred to as "layered security" or "defense in depth." When implementing security measures, managers should acknowledge that no measure is one hundred percent secure. Information security uses the depth of its layers to achieve a combined level of security. A weakness in any one of these layers can cause security to break. Physical protection is the initial step in a layered approach to cloud information security. If it is nonexistent, implemented incorrectly, weak, exercised inconsistently, treated as a project (fire-n-forget), or properly reviewed and maintained, the best logical security measures will not make up for the physical security weakness, and security overall can fail.

An effective traditional security program flows from a well-developed series of risk assessments, vulnerability analysis, BCP/DR policies, processes, and procedures that are reviewed and tested on a regular basis. Well-developed physical security programs will result in physical security that is scalable with the business, repeatable across the organization, measurable, sustainable, defensible, continually improving, and cost-effective on an ongoing basis.

Overview: Some of the security risks associated with cloud computing are unique, and it is in this context the business continuity, disaster recovery, and traditional security environments of a cloud service provider need to be assessed thoroughly (e.g., using standard industry guidelines such as **TOGAF**[34], **SABSA**[35], **ITIL**[36], **COSO**[37], or **COBIT**[38]). This domain addresses:

[33] **CSP** - Cloud Service Provider
[34] **TOGAF** - The Open Group Architecture Framework
[35] **SABSA** - Sherwood Applied Business Security Architecture

- Establishing a Physical Security Function
- Human Resources Physical Security
- Business Continuity
- Disaster Recovery

This section maps to Cloud Control Matrix Domains FS-01, FS-02, FS-03, and FS-04 as well as ISO/IEC 27002 Clause 9.

7.1 Establishing a Traditional Security Function

Outdated security for IT equipment, network technology, and telecommunications is often overlooked in many organizations. This has resulted in many organizations installing computer equipment, networks, and gateways in buildings that did not have proper physical facilities designed to secure the assets or maintain availability.

To establish proper physical security for IT equipment, network technology, and telecommunications assets in a cloud environment, it is important that responsibilities be assigned to personnel who are appropriately placed in a cloud provider's organization. An individual in a management position within a cloud provider is responsible for managing planning, implementation, and maintenance of relevant plans and procedures. Personnel responsible for physical security need to be trained and have their performance evaluated. In establishing a physical security function within a cloud environment, the following must be considered:

- The security needs for the equipment and services being protected
- The human resources that are in place for physical security
- How legacy physical security efforts have been managed and staffed prior to transition to cloud
- Financial resources available for these efforts

Physical security can be as simple as adding a locked door or as elaborate as implementing multiple layers of barriers and armed security guards. Proper physical security uses the concept of layered defense in appropriate combinations for managing risk by deterring and delaying physical security threats. Physical threats to infrastructure, personnel, and systems are not limited to intrusions. To mitigate these risks, combinations of both active and passive defense are deployed, to include having measures such as:

- Obstacles to deter and delay events, incidents, and attacks
- Detection systems to monitor security and environmental conditions
- Security response designed to repel, apprehend, or discourage attackers

Physical security normally takes one of several forms in design and implementation:

- Environmental design

[36] **ITIL** - Information Technology Infrastructure Library
[37] **COSO** - Committee of Sponsoring Organizations
[38] **COBIT** - Control Objectives for Information and Related Technology

- Mechanical, electronic, and procedural controls
- Detection, response, and recovery procedures
- Personnel identification, authentication, authorization, and access control
- Policies and procedures, including training of staff

7.1.1 Evaluation of Traditional Physical Security

When evaluating the traditional security of a CSP, consumers consider the aspects of the infrastructure as a service/physical presence of the foundational data center provider. These include the physical location of the facility and the documentation of critical risk and recovery factors.

7.1.1.1 Physical Location of the CSP Facility

Consumers should conduct a critical evaluation of the data center's physical location. If they are dependent on a cloud supply chain, it is important to understand the cloud infrastructure on which they depend.

The following are suggestions in evaluating the physical location of the facility:

- Check if the location of the facility falls under any active seismic zone and the risks of seismic activity
- The facility should not be located in a geographic region which is prone to: flooding, landslides, or other natural disasters
- The facility should not be in an area with high crime, political or social unrest
- Check the accessibility of the facility's location (and frequency of inaccessibility)

7.1.1.2 Documentation Review

The documentation supporting recovery operations is critical in evaluating the readiness of the hosting company to recover from a catastrophic event. The following sets of documentation should be inspected prior to engagement of a physical data center provider:

- Risk Analysis
- Risk Assessments
- Vulnerability Assessments
- Business Continuity Plans
- Disaster Recovery Plans
- Physical and Environmental Security Policy
- User Account Termination Procedures
- Contingency Plan, including test protocols

- Incident Reporting and Response Plan, including test protocols
- Emergency Response Plan
- Facility Layout – emergency exits, positioning of CCTV cameras, secure entry points
- Fire Exit Route Map and Fire Order Instructions
- Emergency Evacuation Plan and Procedures
- Crisis Communication Procedures
- Emergency Contact Numbers
- User Facility Access Review/Audit Records
- Security Awareness Training documentation, presentation, handouts, etc.
- Security Awareness Attendance Records
- Succession Planning for key executives
- Technical Documents – electrical wiring diagrams, BMS, UPS, AHU details
- Maintenance Schedule of Electrical, Generator, and CCTV
- Emergency fuel service providers contracts
- List of Authorized Personnel allowed entry inside facility
- Security Staff profiles – bio and background information
- Background Check Reports of Security Staff (must be performed every year)
- Annual Maintenance Contracts for key equipment and devices (focus on **SLA's**[39] for equipment/devices downtime and restoration)

When inspecting the documents, there are areas of critical focus that the purchaser of cloud services should focus on to ensure that his/her risk is mitigated. The following advice may prove critical in securing a cloud consumer's business interest when transitioning to cloud:

- Check whether all the documents are up to date and current. These documents must be reviewed by the CSP at least once a year. The revision dates and sign off by management must be included and validated as proof of them being reviewed internally.
- Further, the policy and procedure documents (that are suitable for employee viewing) should be made available through a common Intranet site where authorized employees of the CSP can access them anytime for reference. Adequate care must be taken by the security team to ensure the uploaded documents are the latest versions duly approved by management.

[39] **SLA** - Service Level Agreement

- All policies and procedures will be effective only when employees are aware of them. To this end, check whether a CSP has security awareness program in place. At the minimum, the CSP should ensure that employees are given adequate security awareness training at least once each year and receive sign off from them. Also, new employees joining the organization shall undergo a security orientation session as part of the induction program where key policies and procedures are to be covered with formally signed attendance records maintained and available for review at any time. To make the program effective, senior staff from the security team must conduct the session.

7.1.1.3 Compliance with International/Industry Standards on Security

Ensure that the CSP is compliant with global security standards like ISO 27001 ISMS or other industry-standards such as TOGAF, SABSA, ITIL, COSO, or COBIT. These activities will prove invaluable in assessing the CSP's level of security and its maturity.

- Verify the compliance certificate and its validity.
- Look for verifiable evidence of resource allocation, such as budget and manpower to sustain the compliance program.
- Verify internal audit reports and evidence of remedial actions for the findings.

7.1.1.4 Visual Walk-Through Inspection of the CSP's facility

Area Coverage

Data Center Perimeter Security should be evaluated when determining what areas require physical coverage. The following are high-risk areas that should be secured:

- Administrative areas
- Reception
- Parking Area
- Storage Area
- Fire Exits
- CCTV Command Center
- Air Handling Unit (AHU) Room
- Locker Room
- UPS Room
- Generator Room
- Fuel Storage Tanks

Signage

Look for the following signage that must be displayed prominently in the facility at appropriate places:

- Fire Escape Route Maps and Emergency Exits
- Fire Order Instructions
- Fire Safety Signages
- Security Posters and instructions
- Anti-tailgating Posters
- Temperature/Humidity-related information
- Warning and Instructional Signage
- Emergency Contact Numbers
- Escalation Chart

7.1.2 Security Infrastructure

Perimeter security is important as it serves as the first line of protection against intruders and unwanted visitors. The principles of perimeter security has undergone sea change with technological advancements. The Four D's of Perimeter Security consists of Deter, Detect, Delay and Deny phases for intruders wanting access to the facility.

The following qualities are preferential when selecting a physical infrastructure provider. Depending on the design and function of the cloud service provider, the following list should be closely adhered to in the selection process. Due care should be taken to ensure the physical infrastructure is adequate for the facility's size and nature and scale of operations. Security controls must be strategically positioned and conform to acceptable quality standards consistent with prevalent norms and best practices.

- Secure Entry Points – Access control systems (proximity cards/biometric access)
- Access Control System linked with fire control panel for emergency release
- Motion-sensing alarms, thermal tracking devices, glass-breakage detection
- Fire safety equipment – wet riser, hydrants, hoses, smoke detectors and water sprinklers
- Fire extinguishers
- Fire exits (must not be locked or blocked)
- Panic Bars in fire exit doors
- Alarm sirens and lights
- CCTV Cameras and DVR server (including backup timelines)

- Door closures and time-delay door alarms
- Gas-based fire suppressants inside Data Centers
- Paper Shredders near printers
- Degaussing devices or disk shredders
- Emergency Response Team Kit (ERT Kit)
- Two-way Radio devices (Walkie-talkie handsets) for security staff
- Duress Alarms underneath security desk and vantage (concealed) points
- Door Frame Metal Detectors at entrance and Hand-held Metal Detectors (if needed)
- Fire-proof Safe to safe keep important documents/media

7.2 Human Resources Physical Security

This section maps to Cloud Control Matrix Domains IS-15, FS-05, FS-06, FS-07 and FS-08 as well as ISO/IEC 27002 Clause 9.

The purpose of the human resources physical control is to minimize the risk of the personnel closest to the data disrupting operations and compromising the cloud. A knowledgeable actor with physical access to a console can bypass most logical protective measures by simply rebooting the system or accessing a system that is already turned on with root or administrator access. A wiring closet can provide hidden access to a network or a means to sabotage existing networks. Consider the following measures:

- Roles and responsibilities (e.g., through a RACI-style matrix)
- Background verification and screening agreements
- Employment agreement (e.g., NDA's)
- Employment termination
- Awareness and training of company policies (i.e., Code or Business Conduct)

Roles and responsibilities are part of a cloud environment, in which people and processes, along with technology, are integrated to sustain tenant security on a consistent basis. Segregation of duties, requires at least two persons with separate job responsibilities to complete a transaction or process end-to-end. Avoidance of conflict of interest is essential to the protection of cloud consumers and measures should be implemented to monitor or avoid this risk. Segregation of duties originated in accounting and financial management; its benefits extend to other risk mitigation needs, such as physical security, availability, and system protection. Segregation of duties is implemented via eliminating high-risk role combinations, e.g., not having the same person who approves a purchase order also able to facilitate payment. The principle is applied to role division in cloud development and operations, as well as a software development life cycle. An example common to cloud software development would be the separation of those who develop applications from the staff who operate those systems. Ensure there are no unauthorized backdoors remaining

in the final delivered product. Ensure different personnel manage different critical infrastructure components. Additionally, granting staff the least amount of access privilege required for them to perform their duties will further reduce but not eliminate risk. The segregation of duties and least privilege/access are principles that support a cloud provider's goal to protect and leverage the organization's information assets. A cloud security management program requires the assignment of key roles and responsibilities that may be held by individuals or groups. These roles and responsibilities must be formally defined in the organization's information security policy framework and formally reviewed and approved by senior management in line with their fiduciary GRC (Governance Risk and Compliance) duties and responsibilities.

Additionally, the development of effective HR security must include employment and confidentiality agreements, background checks (when legally permitted), and legally sound hiring and termination practices. Additional measures to consider, if they are applied across all areas of the organization, include formalized job descriptions, appropriate training, security clearances, job rotation, and mandatory vacations for staff in sensitive or high risk roles.

7.3 Assessing CSP Security

Some of the security risks associated with cloud computing are unique, partly due to an extended data centric chain of custody, and it is in this context the business continuity, disaster recovery, and traditional security environments of a cloud service provider need to be assessed thoroughly and in reference to industry standards.

Traditional or Physical Security of the cloud computing service provider's facility is important and needs to be thoroughly assessed from various parameters. This is an area of highest similarity – the security requirements of a cloud and non-cloud data center are fairly similar.

A holistic view and understanding of the "people, process, technology" model or philosophy of the CSP would immensely help in evaluating the maturity of the CSP and flag open issues with their approach towards security which must be resolved, approved, and closed before proceeding.

Organizational maturity and experience contributes a great deal to the effective handling of physical security programs and any contingencies that may arise. Invariably, there is a strong human element involved in the effective administration of physical security programs. The level of management support and the caliber of the security leadership are significant factors in protecting company assets with management support being critical.

Physical security is generally the first line of defense against unauthorized as well as authorized access to an organization's physical assets and the physical theft of records, trade secrets, industrial espionage, and fraud.

7.3.1 Procedures

Cloud service providers should ensure that the following documents are made available for inspection on demand by clients:

- Background Checks (once yearly) by third party vendors
- Non-Disclosure Agreements
- Implement "need to know" and "need to have" policies for information sharing

- Separation of duties
- User Access Administration
- Defined Job Description (Role and Responsibilities)
- Role-based Access Control System
- User Access Reviews

7.3.2 Security Guard Personnel

Where human monitoring and intervention are necessary, physical security staff comprised of guards, supervisors and officers should be posted (on 24/7 basis) at the CSP's facility.

Among other things, the Site and Post instructions should include the following:

- Checking employee, contract staff, and visitor credentials and use of the sign-in log
- Issuing and recovering visitor badges
- Curbing tail-gating by employees
- Handling visitors and movement within the facility
- Handling security-relevant phone calls
- Monitoring intrusion, fire alarm systems and dispatch personnel to respond to alarms
- Controlling movement of materials into and out of the building and enforcing property pass regulations
- Enforcing rules and regulations established for the building
- Patrolling inside facility
- CCTV monitoring
- Key control and management
- Executing emergency response procedures
- Escalating security-related issues to security manager
- Accepting and dispatching mail
- Escorting unattended business visitors inside the office

7.3.4 Environmental Security

The CSP's facilities should protect both personnel and assets by implementing controls that will protect the environment from environmental hazards. These controls may include but are not limited to: temperature and humidity controls, smoke detectors and fire suppression systems.

7.3.4.1 Environmental Controls

- The data center should be equipped with specific environmental support equipment according to published internal standards, local and/or regional rules or laws including an emergency/uninterruptible power supply.
- Equipment/devices required for environmental controls must be protected to reduce risks from environmental threats and hazards and to reduce the risk of unauthorized access to information.

7.3.4.2 Equipment Location and Protection

The following controls must be considered for systems classified as containing Restricted or Confidential information:

- Equipment is located in a physically secure location to minimize unnecessary access.
- Environmental conditions such as humidity that could adversely affect the operation of computer systems are monitored.
- Security staff shall take into account the potential impact of a disaster happening in nearby premises, e.g., a fire in a neighboring building, water leaking from the roof or in floors below ground level, or an explosion in the street.
- Methods for thoroughly destroying and disposing of discarded media (e.g., disk drives)

7.3.4.3 Equipment Maintenance

To ensure continued availability and integrity, equipment is properly maintained with equipment maintenance controls, including:

- Maintaining equipment in accordance with the supplier's recommended service intervals and specifications
- Permitting only authorized maintenance personnel to carry out repairs and service equipment
- Maintaining records of suspected or actual faults and all preventive and corrective maintenance.
- Using appropriate controls when sending equipment off premises for maintenance. Examples of appropriate controls include proper packaging and sealing of containers, storage in safe and secure places, and clear and complete shipping and tracking instructions.
- Maintaining appropriate policies and procedures for asset control, including records retention for all hardware, firmware, and software encompassing traceability, accountability, and ownership

A thorough review of the CSP's facility would enable the prospective client to understand and evaluate the maturity and experience of the security program. Generally, with the focus on IT security, physical security gets limited attention. However, with the range of threat scenarios prevalent today it is imperative that the physical security receives the

attention it deserves, especially, in an environment where the clients' data may be co-resident with a number of other clients (including competitors), physical security assumes greater significance. Physical Security is one of many interconnected lines of defense against intruders and corporate saboteurs who may want access to a CSP's facility for nefarious purposes.

7.4 Business Continuity

Traditionally, the three tenets of information security are confidentiality, integrity, and availability. Business Continuity deals with the continuity component of those three requirements. The transition to a Cloud Service Provider includes an assessment of the uptime the provider contractually commits to. However, this Service Level Agreement (SLA) may not be enough to satisfy the customer. Consideration should be made to the potential impact should a significant outage occur. Based on recent high profile service disruptions into third party provisioned services, the authors would suggest that maintaining continuity of service is a critical dependency on the business to maintain operations.

The following guidelines should be considered with regard to maintaining the continuity of a given service. Although many of these guidelines will likely apply for internally provisioned services as they would for third party provisioned services (e.g. Cloud), these guidelines are written with the pretext that the responsibility rests with the third party.

7.5 Disaster Recovery

One of the most interesting aspects of cloud storage for IT is how it can be leveraged for backup and disaster recovery (DR). Cloud backup and DR services are targeted at reducing the cost of infrastructure, applications, and overall business processes. Cloud backup and DR must aim to make reliable data protection affordable and easy to manage. The challenges to cloud storage, cloud backup, and DR in particular involve mobility, information transfer to and from the cloud, availability, assuring optimal business continuity, scalability and metered payment. Cloud disaster recovery solutions are built on the foundation of three fundamentals: a fully virtualized storage infrastructure, a scalable file system and a compelling self-service disaster recovery application that responds to customers' urgent business requirements.

Customers transitioning disaster recovery operations to the cloud should review the existence of the following organizations or teams within the service provider's disaster recovery program:

- Emergency Response Team (ERT)
- Crisis Management Team
- Incident response team

The composition of the above teams should be reviewed in detail along with crisis communication procedure.

7.5.1 Restoration Priorities

Review the service providers documented restoration plan: This plan should include details on the priorities regarding restoration sequencing. This should correlate directly with the SLA, as contractually committed, with regards to the

services acquired by the customer and the criticality of the service. The Restoration plan should incorporate and quantify the **RPO**[40] (Recovery Point Objective) and **RTO**[41] (Recovery Time Objective) for services.

Detail the Information security controls that are considered and implemented during the restoration process, which should include as an example:

- Clearances of staff involved during the restoration process
- Physical security controls implemented at alternate site
- Specified dependencies relevant to the restoration process (suppliers and outsource partners)
- Minimum separation for the location of the secondary site if the primary site is made unavailable

7.6 Permissions

- Ensure proper facility design.
- Adopt integrated physical and logical security systems that reinforce one another.
- Establish service level agreements that require the inheritance of employment security obligations and responsibilities by later levels of the supply chain.

7.7 Recommendations

7.7.1 Policy Recommendations

- Cloud providers should consider adopting as a security baseline the most stringent requirements of any customer, such that systems, facilities, and procedures are at a system high level. To the extent these security practices do not negatively impact the customer experience, stringent security practices should prove to be cost effective and quantified by reducing risk to personnel, revenue, reputation, and shareholder value.
- Alternately, providers may target a set of users with lower security requirements, or offer a baseline level to all customers with the potential to up-sell and implement additional measures for those who value them. In the latter case, it should be recognized that some customers will be interested only in providers that deliver uniformly high security. This balancing act includes systems, facilities, and documented procedures.
- Providers should have robust compartmentalization of job duties, perform background checks, require and enforce non-disclosure agreements for employees, and restrict employee knowledge of customers to a least privilege, need to know basis.

[40] **RPO** - Recovery Point Objective
[41] **RTO** - Recovery Time Objective

7.7.2 Transparency Recommendations

- Transparency regarding the security posture of the CSP should be required. Onsite visit to the CSP's facility or data center will help in performing an on-the-spot assessment and gaining a clear understanding of the different security measures that have been put in place. However, due to the on-demand provisioning and multi-tenant aspects of cloud computing, traditional forms of audit and assessment may not be available, or may be modified (e.g., shared access to a third-party inspection).
- To enhance effectiveness of the onsite assessment, the visit to the CSP facility or data center should be carried out unannounced (if need be with the CSP being informed about a broad time window rather than specific times). This will enable to have real assessment on the ground on a normal business day instead of giving an opportunity to the CSP to 'keep up appearances' during a client or third-party visit.
- When direct examination is pursued, the assessment team should comprise at least two members or more with specialists drawn from IT, Information Security, Business Continuity, Physical Security, and Management (e.g., department heads or data owners) functions.
- Customers should request and acquire business continuity planning and disaster recovery documentation prior to visit, including relevant certifications (e.g., based on ISO, **ITIL**[42] standards), and audit reports and test protocols.

7.7.3 Human Resources Recommendations

- Consumers should check to see if the CSP deploys competent security personnel for its physical security function. A dedicated security manager is highly recommended to provide the necessary leadership and drive to the physical security program. Leading industry certifications such as **CISA**[43], **CISSP**[44], **CISM**[45], ITIL, or **CPP**[46] (from **ASIS**[47]) would be helpful in validating the incumbent's knowledge and skills in physical security.
- Consumers should request a thorough review of the reporting structure of the security manager. This will help in determining whether the position has been given due significance and responsibilities. The security manager should report to a functional superior and his/her GRC Committee if one exists. They should not report to Facilities or IT. It would be better if this position reports to the CEO through another chain (e.g., through the CRO or head counsel) in terms of independence and objectivity of the position.

7.7.4 Business Continuity Recommendations

- The customer should review the contract of third party commitments to maintain continuity of the provisioned service. However, the customer should strongly consider further analysis. Typically the customer acts as the Data Controller and where personal data is held, there are likely to be specific regulatory requirements to

[42] **ITIL** - Information Technology Infrastructure Library
[43] **CISA -** Certified Information Security Auditor
[44] **CISSP** - Certified Information System Security Professional
[45] **CISM** - Certified Information Security Manager
[46] **CPP** - Certified Privacy Professional
[47] **ASIS** - American Society for Industrial Security

ensure appropriate controls are employed. Such requirements apply even in the event that a third party data processor is utilized.

- The customer should review the third party Business Continuity processes and any particular certification. For example, the CSP may adhere and certify against BS 25999, the British Standard for Business Continuity Management (BCM). The customer may wish to review the scope of the certification and documented details of the assessment.
- The customer should conduct an onsite assessment of the CSP facility to confirm and verify the asserted controls used to maintain the continuity of the service. It may not be entirely necessary to conduct this unannounced assessment of the CSP facility for the sole purpose of verifying specific BCP controls, as typically such controls are only likely to be engaged when a disaster/event were to occur.
- The customer should ensure that he/she receives confirmation of any BCP/DR tests undertaken by the CSP. While many of the recommendations already mentioned focus on documented assertions that the service will maintain continuity, the true test of these is in the event of a significant incident. Without awaiting an actual disaster to occur, the customer should stress the importance of getting formal confirmation of BCP/DR tests, and whether the tests satisfied the SLAs contractually committed.

7.7.5 Disaster Recovery Recommendations

- Cloud customers should not depend on a singular provider of services and should have a disaster recovery plan in place that facilitates migration or failover should a supplier fail.
- IaaS providers should have contractual agreements with multiple platform providers and have the tools in place to rapidly restore systems in the event of loss.
- Data validation should be an automated or user initiated validation protocol that allows the customer to check their data at any time to ensure the data's integrity.
- Incremental backups should frequently update a replica of all protected systems or snapshots at intervals set by the user for each system, so the consumer determines the settings according to recovery point objectives.
- Full site, system, disk, and file recovery should be accessible via a user-driven, self-service portal that allows the user the flexibility to choose which file disk or system they want to recover.
- The cloud provider should implement fast SLA-based data recovery.
- The SLA should be negotiated up front, and the customer should pay for the SLA required to ensure that there is no conflict of interest. No data, no file or system disk, should take more than 30 minutes to recover.
- WAN optimization between the customer and the physical site should be in place so that the cloud enables full data mobility at reduced bandwidth, storage utilization, and cost.

7.8 Requirements

- ✓ All parties must ensure proper structural design for physical security.
- ✓ All supply chain participants must respect the interdependency of deterrent, detective, and authentication solutions.
- ✓ End consumers must inspect, account for, and fix personnel risks inherited from other members of the cloud supply chain. They must also design and implement active measures to mitigate and contain personnel risks through proper separation of duties and least privilege access.

DOMAIN 8 //
DATA CENTER OPERATIONS

In order for Cloud Computing to evolve, the provider must advance the enterprise data center beyond simply using virtualization to manage server assets. In order to enable business agility, green technology, provider openness, increasingly unique ideas in power generation and data center construction and management, the data center has to morph for long-term cloud success.

The "Next Generation Data Center", a term that has been around for several years, has grown into data center operations that includes business intelligence adaptation within the data center, understanding the applications running in the data center, and the requirement of hosting large scale analytical clusters are evolving as well. The data center is not a standalone entity but an entity that needs to be as agile as the application and also be connected to other data centers so that latency is managed as well as security.

Overview. This domain will address the following topics:

- Physical security considerations as related in the CCM
- Automated data center use case mapping
- The new data center? Cloud computing at home
- Cloud infrastructure dissemination and the data center

CCM considerations and how new ideas in cloud data center affect each other

8.1 Data Center Operations

New concepts in this section:

- **Cloud Application Mission**. The industry or application mission housed within the data center. For example, a health care or e-commerce application mission.
- **Data Center Dissemination**. Cloud infrastructures that operate together but are in physically separate physical locations.

Service based automation and predictive analytics to enable service-based automation have been long represented by Information Technology Service Management[48] **(ITSM)** using Information Technology Infrastructure Library[49] **(ITIL)** standards for data center evolution. Different types of applications housed by data centers require automation. Those who operate the data center benefit greatly by understanding what is running inside it and how the data center as a whole needs to respond to varying use.

[48] **ITSM** - Information Technology Service Management
[49] **ITIL** - Information Technology Infrastructure Library

The Cloud Security Alliance's Cloud Controls Matrix has a number of physical requirements based upon different standards and regulatory requirements. The physical security domain in this version of guidance and the Cloud Controls Matrix should be read by the data center professional to get an understanding of requirements inside and outside the data center. For reference, the following table illustrates data center controls needed based upon the mission of the applications housed within the data center. The table is not all-inclusive but provides some examples cross-referencing a Cloud Control Matrix control and specification to an application type or mission.

Table 1— Application Mission by Control

APPLICATION MISSION	CONTROL	SPECIFICATION
Health Care (HIPAA[50])	Facility Security -Security Policy	Policies and procedures shall be established for maintaining a safe and secure working environment in offices, rooms, facilities and secure areas.
Card Processing/Payment (PCI[51])	Facility Security - User Access	Physical access to information assets and functions by users and support personnel shall be restricted.
Power Generation (NERC CIP[52])	Facility Security - Controlled Access Points	Physical security perimeters (fences, walls, barriers, guards, gates, electronic surveillance, physical authentication mechanisms, reception desks and security patrols) shall be implemented to safeguard sensitive data and information systems.

The list above is not meant to be exhaustive in this chapter. The reader can look at the matrix and based upon the standards the organization wants to abide by or which regulations the organization must adhere to can be seen there.

An application running in the data center that contains regulated information (governed under an information security or application security standard) will be audited. The result of the physical audit findings undertaken by the data center operator can then be published to the customers of the data center operator or included in an application query infrastructure such as that provided by Cloud Audit.

In past versions of the Guidance, the reader was instructed to conduct their own audits. For many data center operators or cloud providers this might not be physically possible. In multi-tenant environments the operator or provider cannot normally accommodate visits by every customer to conduct an audit. The customer should require the operator or provider to provide independent audit results.

This idea brings in service automation. By automating reporting, logging, and the publication of audit results the data center operator can provide their customer with evidence that, based upon the application mission, the data center

[50] **HIPAA** - Healthcare Information Portability and Protection Act
[51] **PCI** - Payment Card Industry. Specifically PCI DSS, which is Data Security Standard
[52] **NERC CIP** - North American Electric Reliability Corporation Critical Infrastructure Protection

specific controls are in place and satisfactory. Cloud Audit, Cloud Trust Protocol, and CYBEX (X.1500) can automate the publication of audit findings through a common accessible interface.

Further automation in the data center relies on the library that contains the assets being housed with the data center. By understanding how the assets in the library use resources in the data center, the operations management can predict which tenants are using resources. If the data center uses concepts such as **PoD's**[53] and virtual data center **VMDC**[54] then the data center is as agile as it can be promoting the cloud or virtualized business quickly.

8.1.1 New and Emerging Models

Recently (Summer 2011) there was more news about home-based cloud platforms. In these types of infrastructures modeled after **SETI@home**[55], a cloud is based on the compute assets of volunteers exposing their home/office computers to support other applications. The data centers in these cases are the homes of each of the volunteers. These types of clouds would work well as community-based application hosting environments, but not regulated environments where standards are audited. For example, if a cloud is hosted on 100,000 home computers there would be no way to audit a data center that is effectively cut up into 100,000 pieces and scattered across a large geographical area. This type of infrastructure would host a community based set of applications based upon interest (book club for example) or a residential web site.

The cloud is increasingly being viewed as a commodity or as a utility. There are efforts in the industry to create Security as a Service or create broker infrastructures for identity, interoperability, and business continuity amongst other reasons. The application then is being pulled apart and placed into specialized physical environments that focus on specific needs of an organization or the applications they run.

Data center dissemination takes the application and places it across many other specialized data centers that house and manage specific needs. By disseminating the application across physical boundaries the application is less burdened in the cloud but harder to control and manage.

8.2 Permissions

- Dissemination of data center collaboration. Data center automation having to span multiple physical unrelated data centers will need software to orchestrate what the data center needs for logging and report generation during audits.
- Home based clouds where the data center is personal. Auditing for standards and compliance are near impossible in home based clouds. Regulated environments and standards based environments will have difficulty with home-based clouds based on the controls needed. There may be aspects to an application where some part of the application can be disseminated to home-based infrastructure.

[53] **PoD** - Point of Delivery. A rack-able aggregated set of power, compute, storage access, and network components contained in a single unit

[54] **VMDC** - Virtual Multi-tenant Data Center. A concept using modular, easily rack-able components to quickly expand a data center such as PoD's

[55] **SETI@home** - http://setiathome.berkeley.edu/

8.3 Recommendations

- Organizations building cloud data centers should incorporate management processes, practices, and software to understand and react to technology running inside the data center.
- Organizations buying cloud services should ensure that the provider has adopted service management processes and practices to run their data centers and have adopted racking techniques that ensure agile and highly available resources inside the data center.
- Understand the mission of what is running in the data center. Given the controls in the Cloud Control Matrix the data center being built or purchased must conform to physical and asset security requirements.
- Data center locations are important. If technology and application components are spread across data centers, then there will be latency between the data centers.
- Organizations buying cloud services must clearly understand and document which parties are responsible for meeting compliance requirements, and the roles they and their cloud provider when assessing compliance.

8.4 Requirements

The Cloud Security Alliance has many sources of information to help with the construction or remodeling of data centers for the cloud. The controls matrix highlights requirements across a very broad set of security standards and regulations. Cloud Audit and other projects within the CSA also can help with construction and management of data centers and the technology running within them.

- ✓ Fully understand Control Matrix requirements based upon what is going to run in the data center. Use a common denominator that satisfies most application missions.
- ✓ Use IT service management techniques to ensure availability, security, and asset delivery and management.
- ✓ If the data center is owned by a provider, audit against a regulatory and security standard template and publish results to the customer.

DOMAIN 9 //
INCIDENT RESPONSE

Incident Response (IR) is one of the cornerstones of information security management. Even the most diligent planning, implementation, and execution of preventive security controls cannot completely eliminate the possibility of an attack on the information assets. One of the central questions for organizations moving into the cloud must therefore be: what must be done to enable efficient and effective handling of security incidents that involve resources in the cloud?

Cloud computing does not necessitate a new conceptual framework for Incident Response; rather it requires that the organization appropriately maps its extant IR programs, processes, and tools to the specific operating environment it embraces. This is consistent with the guidance found throughout this document; a gap analysis of the controls that encompass organizations' IR function should be carried out in a similar fashion.

This domain seeks to identify those gaps pertinent to IR that are created by the unique characteristics of cloud computing. Security professionals may use this as a reference when developing response plans and conducting other activities during the preparation phase of the IR lifecycle. To understand the challenges cloud computing poses to incident handling, we must examine, which challenges the special characteristics of cloud computing and the various deployment and service models pose for incident handling.

This domain is organized in accord with the commonly accepted Incident Response Lifecycle as described in the National Institute of Standards and Technology Computer Security Incident Handling Guide (NIST 800-61) [1]. After establishing the characteristics of cloud computing that impact IR most directly, each subsequent section addresses a phase of the lifecycle and explores the potential considerations for responders.

Overview. This domain will address the following topics:

- Cloud computing impact on Incident Response
- Incident Response Lifecycle
- Forensic accountability

9.1 Cloud Computing Characteristics that Impact Incident Response

Although cloud computing brings change on many levels, certain characteristics [2] of cloud computing bear more direct challenges to IR activities than others [3].

First, the on demand self-service nature of cloud computing environments means that a cloud customer may find it hard or even impossible to receive the required co-operation from their cloud service provider (**CSP**)[56] when handling a security incident. Depending on the service and deployment models used, interaction with the IR function at the CSP

[56] **CSP** - Cloud Service Provider

will vary. Indeed, the extent to which security incident detection, analysis, containment, and recovery capabilities have been engineered into the service offering are key questions for provider and customer to address.

Second, the resource pooling practiced by cloud services, in addition to the rapid elasticity offered by cloud infrastructures, may dramatically complicate the IR process, especially the forensic activities carried out as part of the incident analysis. Forensics has to be carried out in a highly dynamic environment, which challenges basic forensic necessities [4] such as establishing the scope of an incident, the collection and attribution of data, preserving the semantic integrity of that data, and maintaining the stability of evidence overall. These problems are exacerbated when cloud customers attempt to carry out forensic activities, since they operate in a non-transparent environment (which underscores the necessity of support by the cloud provider as mentioned above).

Third, resource pooling as practiced by cloud services causes privacy concerns for co-tenants regarding the collection and analysis of telemetry and artifacts associated with an incident (e.g. logging, netflow data, memory, machine images, and storage, etc.) without compromising the privacy of co-tenants. This is a technical challenge that must be addressed primarily by the provider. It is up to the cloud customers to ensure that their cloud service provider has appropriate collection and data separation steps and can provide the requisite incident-handling support.

Fourth, despite not being described as an essential cloud characteristic, cloud computing may lead to data crossing geographic or jurisdictional boundaries without the explicit knowledge of this fact by the cloud customer. The ensuing legal and regulatory implications may adversely affect the incident handling process by placing limitations on what may or may not be done and/or prescribing what must or must not be done during an incident across all phases of the lifecycle [5]. It is advisable that an organization includes representatives from its legal department on the Incident Response team to provide guidance on these issues.

Cloud computing also presents opportunities for incident responders. Cloud continuous monitoring systems can reduce the time it takes to undertake an incident handling exercise or deliver an enhanced response to an incident. Virtualization technologies, and the elasticity inherent in cloud computing platforms, may allow for more efficient and effective containment and recovery, often with less service interruption than might typically be experienced with more traditional data center technologies. Also, investigation of incidents may be easier in some respects, as virtual machines can easily be moved into lab environments where runtime analysis can be conducted and forensic images taken and examined.

9.2 The Cloud Architecture Security Model as a Reference

To a great extent, deployment and service models dictate the division of labor when it comes to IR in the cloud ecosystem. Using the architectural framework and security controls review advocated in Domain 1 (see Cloud Reference Model Figure 1.5.2a) can be valuable in identifying what technical and process components are owned by which organization and at which level of the “stack.”

Cloud service models (IaaS, PaaS, SaaS) differ appreciably in the amount of visibility and control a customer has to the underlying IT systems and other infrastructure that deliver the computing environment. This has implications for all phases of Incident Response as it does with all other domains in this guidance document.

For instance, in a SaaS solution, response activities will likely reside almost entirely with the CSP, whereas in IaaS, a greater degree of responsibility and capability for detecting and responding to security incidents may reside with the

customer. However, even in IaaS there are significant dependencies on the CSP. Data from physical hosts, network devices, shared services, security devices like firewalls and any management backplane systems must be delivered by the CSP. Some providers are already provisioning the capability to deliver this telemetry to their customers and managed security service providers are advertising cloud-based solutions to receive and process this data.

Given the complexities, the Security Control Model described in Domain 1 (Figure 1.5.1c), and the activities an organization performs to map security controls to your particular cloud deployment should inform IR planning and vice versa. Traditionally, controls for IR have concerned themselves more narrowly with higher-level organizational requirements; however, security professionals must take a more holistic view in order to be truly effective. Those responsible for IR should be fully integrated into the selection, purchase, and deployment of any technical security control that may directly, or even indirectly, affect response. At a minimum, this process may help in mapping of roles/responsibilities during each phase of the IR lifecycle.

Cloud deployment models (public, private, hybrid, community) are also considerations when reviewing IR capabilities in a cloud deployment; the ease of gaining access to IR data varies for each deployment model. It should be self-evident that the same continuum of control/responsibility exists here as well. In this domain, the primary concern is with the more public end of the continuum. The authors assume that the more private the cloud, the more control the user will have to develop the appropriate security controls or have those controls delivered by a provider to the user's satisfaction.

9.3 Incident Response Lifecycle Examined

NIST 800-61 [1] describes the following main stages of the IR lifecycle: preparation; detection & analysis; containment, eradication & recovery. This section examines the specific challenges of cloud computing for these stages and provides recommendations as to how these challenges can be met.

9.3.1 Preparation

Preparation may be the most important phase in the Incident Response Lifecycle when information assets are deployed to the cloud. Identifying the challenges (and opportunities) for IR should be a formal project undertaken by information security professionals within the cloud customer's organization prior to migration to the cloud. If the level of IR expertise within the organization is deemed insufficient, experienced external parties should be consulted. This exercise should be undertaken during every refresh of the enterprise Incident Response Plan.

In each lifecycle phase discussed below, the questions raised and suggestions provided can serve to inform the customer's planning process. Integrating the concepts discussed into a formally documented plan should serve to drive the right activities to remediate any gaps and take advantage of any opportunities.

Preparation begins with a clear understanding and full accounting of where the customer's data resides in motion and at rest. Given that the customer's information assets may traverse organizational, and likely, geographic boundaries necessitates threat modeling on both the physical and logical planes. Data Flow diagrams that map to physical assets, and map organizational, network, and jurisdictional boundaries may serve to highlight any dependencies that could arise during a response.

Since multiple organizations are now involved, Service Level Agreements (**SLA**)[57] and contracts between the parties now become the primary means of communicating and enforcing expectations for responsibilities in each phase of the IR lifecycle. It is advisable to share IR plans with the other parties and to precisely define and clarify shared or unclear terminology. When possible, any ambiguities should be cleared up in advance of an incident.

It is unreasonable to expect CSP's to create separate IR plans for each customer. However, the existence of some (or all) of the following points in a contract/SLA should give the customer organization some confidence that its provider has done some advanced planning for Incident Response:

- Points of Contact, communication channels, and availability of IR teams for each party
- Incident definitions and notification criteria, both from provider to customer as well as to any external parties
- CSP's support to customer for incident detection (e.g., available event data, notification about suspicious events, etc.)
- Definition of roles/responsibilities during a security incident, explicitly specifying support for incident handling provided by the CSP (e.g., forensic support via collection of incident data/artifacts, participation/support in incident analysis, etc.)
- Specification of regular IR testing carried out by the parties to the contract and whether results will be shared
- Scope of post-mortem activities (e.g, root cause analysis, IR report, integration of lessons learned into security management, etc.)
- Clear identification of responsibilities around IR between provider and consumer as part of SLA

Once the roles and responsibilities have been determined, the customer can now properly resource, train, and equip its Incident Responders to handle the tasks that they will have direct responsibility for. For example, if a customer-controlled application resides in a PaaS model and the cloud provider has agreed to provide (or allow retrieval of) platform-specific logging, having the technologies/tools and personnel available to receive, process, and analyze those types of logs is an obvious need. For IaaS and PaaS, aptitude with virtualization and the means to conduct forensics and other investigation on virtual machines will be integral to any response effort. A decision about whether the particular expertise required is organic to the customer organization or is outsourced to a Third Party is something to be determined during the preparation phase. Please note that outsourcing then prompts another set of contracts/**NDA's**[58] to manage.

Between all involved parties, communication channels must be prepared. Parties should consider the means by which sensitive information is transmitted between parties to ensure that out-of-band channels are available and that encryption schemes are used to ensure integrity and authenticity of information. Communication during IR can be facilitated by utilizing existing standards for the purpose of sharing indicators of compromise or to actively engage another party in an investigation. For example, the Incident Object Description Exchange Format (**IODEF**)[59] [6] as well as the associated Real-time Inter-network Defense (**RID**)[60] standard [7,8] were developed in the Internet Engineering Task

[57] **SLA** - Service Level Agreements
[58] **NDA** - Non-Disclosure Agreement
[59] **IODEF** - Incident Object Description Exchange Format
[60] **RID** - Real-time Inter-network Defense

Force (**IETF**)[61] and are also incorporated in the International Telecommunication Union's (**ITU**)[62] Cybersecurity Exchange (**CYBEX**)[63] project. IODEF provides a standard XML schema used to describe an incident, RID describes a standard method to communicate the incident information between entities, which includes, at least, a CSP and tenant.

The most important part of preparing for an incident is testing the plan. Tests should be thorough and mobilize all the parties who are likely going to be involved during a true incident. It is unlikely that a CSP has resources to participate in tests with each of its customers; the customer should therefore consider role-playing as a means to identify which tasking or requests for information are likely to be directed at the CSP. This information should be used to inform future discussions with the provider while in the preparation phase. Another possibility is for the customer to volunteer to participate in any testing the CSP may have planned.

9.3.2 Detection and Analysis

Timely detection of security incidents and successful subsequent analysis of the incident (what has happened, how did it happen, which resources are affected, etc.) depend on the availability of the relevant data and the ability to correctly interpret that data. As outlined above, cloud computing provides challenges in both cases. Firstly, availability of data to a large extent depends on what the cloud provider supplies to the customer and may be limited by the highly dynamic nature of cloud computing. Secondly, analysis is complicated by the fact that the analysis at least partly concerns non-transparent, provider-owned infrastructure, of which the customer usually has little knowledge and – again – by the dynamic nature of cloud computing, through which the interpretation of data becomes hard, sometimes even impossible.

Putting aside the technical challenges of incident analysis for a moment, the question on how a digital investigation in the cloud should be conducted in order to maximize the probative value (i.e. credibility) of the evidence at the time of writing remains largely unanswered. Hence, until legal cases involving cloud incidents have become more common place and commonly accepted best practice guidelines exist, analysis results for cloud security incidents incur the risk of not standing up in court.

Until standards, methods, and tools for detecting and analyzing security incidents have caught up with the technical developments introduced by cloud computing, incident detection and analysis will remain especially challenging for cloud environments. Cloud customers must rise to this challenge by making sure that they have access to (1) the data sources and information that are relevant for incident detection/analysis as well as (2) appropriate forensic support for incident analysis in the cloud environment(s) they are using.

9.3.3 Data Sources

As in any hosted IT service integration, the IR team will need to determine the appropriate logging required to adequately detect anomalous events and identify malicious activity that would affect their assets. It is imperative for the customer organization to conduct an assessment of what logs (and other data) are available, how they are collected and processed, and finally, how and when they may be delivered by the CSP.

[61] **IETF** - Internet Engineering Task Force
[62] **ITU** - International Telecommunication Union's
[63] **CYBEX** - Cybersecurity Exchange

The main data source for the detection and subsequent analysis of incidents on the customer side, is logging information. The following issues regarding logging information must be taken into consideration:

- **Which information should be logged?** Examples of log types that may be relevant are audit logs (e.g. network, system, application, and cloud administration roles and accesses, backup and restore activities, maintenance access, and change management activity), error logs (e.g., kernel messages of hypervisors, operating systems, applications, etc.), security-specific logs (e.g., IDS logs, firewall logs, etc.), performance logs, etc. Where existing logging information is not sufficient, additional log sources have to be negotiated/added.
- **Is the logged information consistent and complete?** A prime example of inconsistent source logging information is failure to synchronize clocks of log sources. Similarly, incomplete information regarding the time zone in which the timestamps in logs are recorded makes it impossible to accurately interpret the collected data during analysis.
- **Is the cloud's dynamic nature adequately reflected in the logged information?** The dynamic behavior of cloud environments also is a frequent reason for inconsistent and/or incomplete logging. For example, as new cloud resources (VM's, etc.) are brought online to service demand, the log information produced by the new resource instance will need to be added to the stream of log data. Another likely problem is the failure to make dynamic changes in the environment explicit in the log information. For example, consider the case that web service requests to a certain PaaS component are logged, but may be serviced dynamically by one of various instances of this service. Incomplete information regarding the question, which instance served which request, may then make proper analysis hard or impossible, e.g., if the root cause of an incident is a single compromised instance.
- **Are overall legal requirements met?** Privacy issues regarding co-tenants, regulation regarding log data in general and personally identifiable information in particular, etc., may place limitations on the collection, storage, and usage of collected log data. These regulatory issues must be understood and addressed for each jurisdiction where the company's data is processed or stored.
- **What log retention patterns are required?** Legal and compliance requirements will direct specific log retention patterns. Cloud customers should understand and define any extended log retention patterns to meet their need over time to support their requirements for incident analysis/forensics.
- **Are the logs tamper-resistant?** Ensuring that logs are in tamper resistant stores is critical for accurate legal and forensic analysis. Consider the use of write-once devices, separation of servers used to storage logs from application servers and access controls to servers storing logs as critical aspects of this requirement.
- **In which format is logging information communicated?** Normalization of logging data is a considerable challenge. The use of open formats (such as the emerging CEE [9]) may ease processing at the customer side.

The cloud provider can only detect some incidents because such incidents occur within the infrastructure owned by the cloud provider. It is important to note that the SLA's must be such that the cloud provider informs cloud customers in a timely and reliable manner to allow for agreed IR to occur. For other incidents (perhaps even detectable by the customer) the cloud provider may be in a better position for detection. Cloud customers should select cloud providers that optimally assist in the detection of incidents by correlating and filtering available log data.

The amount of data produced from the cloud deployment may be considerable. It may be necessary to investigate cloud provider options regarding log filtering options from within the cloud service before it is sent to the customer to reduce network and customer internal processing impacts. Additional considerations include the level of analysis or correlation performed by the CSP and the cloud tenant to identify possible incidents prior to forensics. If analysis is performed at the CSP, the escalation and hand-off points for the incident investigation must be determined.

9.3.4 Forensic and Other Investigative Support for Incident Analysis

Although still immature, efforts are already underway within the forensic community to develop the tools and protocols to collect and examine forensic artifacts derived especially from virtualized environments; also, forensic support required for PaaS and SaaS environments is subject to ongoing research.

It is important that the customer understands the forensic requirements for conducting incident analysis, researches to what extent CSP's meet these requirements, chooses a CSP accordingly, and addresses any remaining gaps. The amount of potential evidence available to a cloud customer strongly diverges between the different cloud service and deployment models.

For IaaS services, customers can execute forensic investigations of their own virtual instances but will not be able to investigate network components controlled by the CSP. Furthermore, standard forensic activities such as the investigation of network traffic in general, access to snapshots of memory, or the creation of a hard disk image require investigative support to be provided by the CSP. Also advanced forensic techniques enabled by virtualization such as generating snapshots of virtual machine states or VM introspection on live systems require forensic support by the CSP.

With PaaS and SaaS security incidents that have their root cause in the underlying infrastructure, the cloud customer is almost completely reliant on analysis support of the CSP, and as mentioned previously, roles and responsibilities in IR must be agreed upon in the SLAs. With PaaS, the customer organization will be responsible for any application layer code that is deployed to the cloud. Sufficient application logging is required for incident analysis in scenarios where the root cause lies within the application (e.g., a flaw in the application code). In this case, support by the CSP could take the form of facilitating application log generation, secure storage, and secure access via a read-only API [10]. SaaS providers that generate extensive customer-specific application logs and provide secure storage as well as analysis facilities will ease the IR burden on the customer. This may reduce application level incidents considerably.

Providers that use their management backplane/systems to scope an incident and identify the parts of a system that have been under attack or are under attack, and provide that data to the cloud customer, will greatly enhance the response in all service models.

To prepare for incident analysis in a given cloud environment, the customer's IR team should familiarize themselves with information tools the cloud vendor provides to assist the operations and IR processes of their customer. Knowledge base articles, FAQs, incident diagnosis matrices, etc. can help fill the experience gap a cloud customer will have with regard to the cloud infrastructure and its operating norms. This information may, for example, assist the IR team in discriminating operational issues from true security events and incidents.

9.3.5 Containment, Eradication, and Recovery

As with the other phases of Incident Response, close coordination with all stakeholders is required to ensure that strategies developed to contain, eradicate, and recover from an incident are effective, efficient, and take into consideration all legal and privacy implications. The strategies must be also consistent with business goals and seek to minimize disruption to service. This is considerably more challenging when multiple organizations are involved in the response, as is the case with cloud computing.

Options for this phase will differ depending upon the deployment and service model, and also the layer of the stack at which the attack was targeted. There may be multiple strategies that can be employed, possibly by different entities equipped with different technological solutions. If at all possible, thought exercises should be conducted in the preparation phase to anticipate these scenarios and a conflict resolution process identified. Customers must also consider how their provider will handle incidents affecting the provider itself or affecting other tenants on a shared platform in addition to incidents that are directly targeted at their own organization.

Consumers of IaaS are primarily responsible for the containment, eradication, and recovery from incidents. Cloud deployments may have some benefits here. For example, isolating impacted images without destroying evidence can be achieved by pausing these images. As discussed in the introduction, the relative ease with which nodes can be shut down and new instances brought up may help to minimize service interruption when a code fix needs to be deployed. If there are issues with a particular IaaS cloud, then the customer may have the option of moving the service on to another cloud, especially if they have implemented one of the meta-cloud management solutions.

The situation is more complicated for SaaS and PaaS deployments. Consumers may have little technical ability to contain a Software or Platform as a Service incident other than closing down user access and inspecting/cleaning their data as hosted within the service prior to a later re-opening. But especially for SaaS, even these basic activities may be difficult or impossible without adequate support by the CSP, such as fine-grained access control mechanisms and direct access to customer data (rather than via the web interface).

In all service models, providers may be able to assist with certain categories of attack, such as a Denial of Service (**DoS**)[64]. For example, smaller enterprises may benefit from the economies of scale, which allow for more expensive mitigation technologies, such as DoS protection, to be extended to their sites. As for the previous phases, the extent to which facilities at the provider will be made available to the customer to assist in responding to an attack should be identified in the preparation phase. In addition, the conditions under which the provider is obligated to provide assistance to responding to an attack should be contractually defined.

The SLA's and the IR plan should be flexible enough to accommodate a "Lessons Learned" activity after the recovery. A detailed Incident Report based on the IR activities is to be written and shared with impacted parties, i.e., between the cloud customer, CSP and other affected/involved organizations. The Incident Report should include the timeline of the incident, analysis of the root cause or vulnerability, actions taken to mitigate problems and restore service, and recommendations for long-term corrective action.

Corrective actions are likely to be a blend of customer-specific and provider supported, and the provider's Incident Response team should provide a section with their perspective of the incident and proposed resolution. After an initial review of the Incident Report by the customer and CSP, joint discussions should be held to develop and approve a remediation plan.

9.4 Recommendations

- Cloud customers must understand how the CSP defines events of interest versus security incidents and what events/incidents the cloud-service provider reports to the cloud customer in which way. Event information that is supplied using an open standard can facilitate the processing of these reports at the customer side.

[64] **DoS** - Denial of Service

- Cloud customers must set up proper communication paths with the CSP that can be utilized in the event of an incident. Existing open standards can facilitate incident communication.
- Cloud customers must understand the CSP's support for incident analysis, particularly the nature (content and format) of data the CSP will supply for analysis purposes and the level of interaction with the CSP's incident response team. In particular, it must be evaluated whether the available data for incident analysis satisfies legal requirements on forensic investigations that may be relevant to the cloud customer.
- Especially in case of IaaS, cloud customers should favor CSP's that leverage the opportunities virtualization offers for forensic analysis and incident recovery such as access/roll-back to snapshots of virtual environments, virtual-machine introspection, etc.
- Cloud customers should favor CSP's that leverage hardware assisted virtualization and hardened hypervisors with forensic analytic capabilities.
- For each cloud service, cloud customers should identify the most relevant incident classes and prepare strategies for the incident containment, eradication, and recovery incidents; it must be assured that each cloud provider can deliver the necessary assistance to execute those strategies.
- Cloud customers should obtain and review a CSP's history for incident response. A CSP can provide industry recommendations from existing customers about its IRP.

9.5 Requirements

- ✓ For each cloud-service provider that is used, the approach to detecting and handling incidents involving resources hosted at that provider must be planned and described in the enterprise incident response plan.
- ✓ The SLA of each cloud-service provider that is used must guarantee the support for incident handling required for effective execution of the enterprise incident response plan for each stage of the incident handling process: detection, analysis, containment, eradication, and recovery.
- ✓ Testing will be conducted at least annually. Customers should seek to integrate their testing procedures with that of their provider (and other partners) to the greatest extent possible. Ideally, a team (comprising Customer and CSP members) should carry out various health check tests for an incident response plan, and accordingly, suggestions should be implemented into a new version of incident response plan.

REFERENCES

[1] GRANCE, T., KENT, K., and KIM, B. Computer Security Incident Handling Guide. NIST Special Publication 800-61.

[2] MELL, P. and GRANCE, T. The NIST Definition of Cloud Computing, NIST Special Publication 800-145.

[3] GROBAUER, B. and SCHRECK, T. October 2010. Towards Incident Handling in the Cloud: Challenges and Approaches. In Proceedings of the Third ACM Cloud Computing Security Workshop (CCSW), Chicago, Illinois.

[4] WOLTHUSEN, S. 2009. Overcast: Forensic Discovery in Cloud Environments. In Proceedings of the Fifth International Conference on IT Security Incident Management and IT Forensics.

[5] REED, J. 2011. Following Incidents into the Cloud. SANS Reading Room

[6] DANYLIW, R., et al. 2007. The Incident Object Description Exchange Format, IETF Internet Draft RFC 5070.

[7] MORIARTY, K. 2010. Real-time Inter-network Defense, IETF Internet Draft RFC 6045.

[8] MORIARTY, K., and TRAMMELL, B. 2010. Transport of Real-time Inter-network Defense (RID) Messages, IETF Internet Draft RFC 6046.

[9] FITZGERALD, E., et al. 2010. Common Event Expression (CEE) Overview. Report of the CEE Editorial Board.

[10] BIRK, D. and WEGENER, C. 2011. Technical Issues of Forensic Investigations in Cloud Computing Environments In Proceedings of 6th International Workshop on Systematic Approaches to Digital Forensic Engineering (IEEE/SADFE), Oakland, CA, USA.

DOMAIN 10 //
APPLICATION SECURITY

Cloud environments, particularly public cloud environments, by virtue of their flexibility and openness challenge many fundamental assumptions about application security. Some of these assumptions are well understood, however, many are not. This section is intended to provide guidance on how cloud computing influences security over the lifetime of an application, from design to operations to ultimate decommissioning. This guidance is for all stakeholders (including application designers, security professionals, operations personnel, and technical management) on how to best mitigate risk and manage assurance when designing Cloud Computing applications.

Cloud Computing is a particular challenge for applications across the layers of Software as a Service (SaaS), Platform as a Service (PaaS) and Infrastructure as a Service (IaaS). Cloud-based software applications require a design rigor similar to an application connecting to the raw Internet—the security must be provided by the application without any assumptions being made about the external environment. However, the threats that applications are going to be exposed to in a cloud environment will be more than those experienced in a traditional data center. This creates the need for rigorous practices that must be followed when developing or migrating applications to the cloud.

Overview. This Application Security domain has been organized into the following areas of focus:

- Secure **SDLC**[65] (General practices for secure Software Development Life Cycle and nuances specific to the Cloud)
- Authentication, Authorization, Compliance –Application Security Architecture in the Cloud
- Identity and the consumption of identity as it relates to Cloud Application Security
- Entitlement processes and risk-based access management as it relates to cloud encryption in cloud-based applications
- Application authorization management (policy authoring/update, enforcement)
- Application Penetration Testing for the Cloud (general practices and nuances specific to cloud-based Applications)
- Monitoring Applications in the Cloud
- Application authentication, compliance, and risk management and the repercussions of multi-tenancy and shared infrastructure
- The difference between avoiding malicious software and providing application security

[65] **SDLC** - Software Development Life Cycle

10.1 Secure SDLC (Software Development Life Cycle)

A Secure Software Development Life Cycle (SSDLC) (also referred by a few as Secure Development Life Cycle (SDLC)) has assumed increased importance when migrating and deploying applications in the cloud. Organizations should ensure that the best practices of application security, identity management, data management, and privacy are integral to their development programs and throughout the lifecycle of the application.

Developing for a cloud environment is different than the traditional hosting environments in the following areas:

- The control over physical security is substantially reduced in public cloud scenarios.
- The potential incompatibility between vendors when services (for example, Storage) are migrated from one vendor to another.
- Protection of data through the lifecycle must be considered. This includes transit, processing and storage.
- The combinations of web services in the cloud environment can potentially cause security vulnerabilities to be present.
- The ability to access logs, especially in a shared public cloud, is more difficult and should be specified as a part of the service level agreement.
- Fail-over for data and data security in the cloud has to be more detailed and layered than traditional environments.
- Assuring (and demonstrating evidence for) compliance to relevant industry and government regulations is typically more difficult within a cloud environment.

In implementing a SSDLC, organizations must adopt best practices for development, either by having a good blend of processes, tools, and technologies of their own or adopting one of the maturity models such as these:

- Building Security In Maturity Model (BSIMM2)
- Software Assurance Maturity Model (SAMM)
- Systems Security Engineering Capability Maturity Model (SSE-CMM)

10.1.1 Application Security Assurance Program

Organizations should have an application security assurance program that ensures for the applications that are being migrated and/or developed and maintained in a cloud environment the following:

- With adequate executive support, goals and metrics are defined, implemented and tracked.
- A security and a privacy policy for applications in the cloud has been established to meet the legal and regulatory compliance requirements that are aligned with the business needs and regulatory obligations of the organization.

- Adequate capability and capacity for security assurance is available within the organization to architect, design, develop, test and deploy secure applications by on-boarding, training suitable resources in a timely manner.
- Security and privacy risk evaluations are performed for all applications to ensure the requirements are correctly defined.
- Processes for ensuring security and privacy requirements for the development and maintenance process in the cloud are defined and implemented.
- Configuration and change management must be auditable and verifiable
- Physical security risk evaluations for the application and the data are performed, and the access to all cloud infrastructure components is adequate to meet those requirements.
- Formal coding best practices, considering the strengths and weaknesses of language used should be followed during the development phase.
- Privacy and security measures must be auditable and verifiable.

10.1.2 Verification and Validation

10.1.2.1 Design Review

Some functions are more security sensitive than others and may not be viable candidates for running in the cloud and should be considered when the specific application design and requirements are specified.

The following principles should be followed in order to develop a secure design for the application. Where these principles cannot be met by a cloud architecture, they should be remediated by appropriate technical and/or compensating controls. Failure to do this brings into question the viability of a cloud deployment.

- **Least privilege**. This principle maintains that an individual, process, or other type of entity should be given the minimum privileges and resources for the minimum period of time required to complete a task. In many cases, least privilege can only be implemented effectively using fine-grained, contextual application authorization management with security policy automation[66] mechanisms.
- **Segregation of duties**. This is a control policy according to which no person should be given responsibility for, or access to, more than one related function.
- **Defense in depth**. This is the application of multiple layers of protection wherein a subsequent layer will provide protection if a previous layer is breached.
- **Fail safe**. If a cloud system fails it should fail to a state in which the security of the system and its data are not compromised. For example, to ensure the system defaults to a state in which a user or process is denied access to the system.
- **Economy of mechanism**. This promotes simple and comprehensible design and implementation of protection mechanisms, so that unintended access paths do not exist or can be readily identified and eliminated.

[66] www.policyautomation.org

- **Complete mediation**. This is where every request by an entity[67] to access an object in a computer system must be authorized.
- **Open design**. This is an open-access cloud system design that has been evaluated and peer-reviewed by a community of experts resulting in a more secure design.
- **Least common mechanism**. This states that a minimum number of mechanisms (especially protection mechanisms) should be common across multiple applications, minimizing one application's ability to corrupt or subvert another application.
- **Weakest link**. It is important to identify the weakest mechanisms in the security chain and layers of defense and improve them, so that risks to the system are mitigated to an acceptable level.

10.1.3 Construction

10.1.3.1 Code Review

It is recommended to define and follow secure software development at the organization level. The guidelines spelled out in the Fundamental Practices for Secure Software Development from SAFECode[68], CERT (**SEI**)[69] or ISO Standards could be followed.

Dynamic code analysis examines the code as it executes in a running cloud application, with the tester tracing the external interfaces in the source code to the corresponding interactions in the executing code, so that any vulnerabilities or anomalies that arise in the executing interfaces are simultaneously located in the source code, where they can then be fixed.

Unlike static analysis, dynamic analysis enables the tester to exercise the software in ways that expose vulnerabilities introduced by interactions with users and changes in the configuration or behavior of environment components.

Some of the best practices for writing a secure code and reviewing are listed below:

- The minimum necessary information should be included in cloud server code. Comments should be stripped from operational code, and names and other personal information should be avoided.
- Utilize source code analysis tools to check for typical programming errors such as Buffer Overflows, Format String Attacks, Race Conditions, etc.
- Verify and validate all inputs, user, computer and inter-system. Content injection and several other attacks are possible when the cloud infrastructure takes any input and applies the content of that input into commands or SQL statements.
- When using object code (binaries), for example, where 3rd party libraries are being used, utilize a testing service capable of static vulnerability testing on object code.

10.1.3.2 Security Testing

[67] An entity could be a user, code, a device, an organization or agent
[68] http://www.safecode.org/
[69] https://www.cert.org/secure-coding/

Penetration test is a security testing methodology that gives the tester insight into the strength of the target's network security by simulating an attack from a malicious source. The process involves an active analysis of the cloud system for any potential vulnerability that may result from poor or improper system configuration, known and/or unknown hardware or software flaws, or operational weaknesses in process or technical countermeasures. This analysis is carried out from the position of a potential attacker, and can involve active exploitation of security vulnerabilities.

The type of cloud model has a huge impact on the penetration testing or in deciding if penetration test is possible. Generally, Platform as a Service (PaaS) and Infrastructure as a Service (IaaS) clouds are likely to permit penetration testing. However, Software as a Service (SaaS) providers are not likely to allow customers to penetration test their applications and infrastructure, with the exception of third parties performing the cloud providers' own penetration tests for compliance or security best practices.

Penetration testing is typically carried out within a "black box" scenario, that is, with no prior knowledge of the infrastructure to be tested. At its simplest level, the penetration test involves three phases:

1. **Preparation**. This is where a formal contract is executed containing non-disclosure of the client's data and legal protection for the tester. At a minimum, it lists the IP addresses to be tested.
2. **Execution**. In this phase the penetration test is executed, with the tester looking for potential vulnerabilities.
3. **Delivery**. The results of the evaluation are communicated to the tester's contact in the organization, and corrective action is advised.

Whether the penetration test is a full knowledge (white box) test, a partial knowledge (gray box) test, or a zero knowledge (black box) test, after the report and results are obtained, mitigation techniques have to be applied to reduce the risk of compromise to an acceptable or tolerable level. The test should have the widest possible scope to address vulnerabilities and corresponding risks to such areas as applications, remote access systems and other related IT assets.

10.1.3.3 Interoperability Testing

Interoperability testing evaluates whether a cloud application can exchange data (interoperate) with other components or applications. Interoperability testing activities determine the capability of applications to exchange data via a common set of exchange formats, to read and write the same file formats, and to communicate using the same protocols.

A major goal of interoperability testing is to detect interoperability problems between cloud software applications before these applications are put into operation. Interoperability testing requires the majority of the application to be completed before testing can occur.

As well as testing for interoperability, these tests should confirm that all data exchanges, protocols and interfaces used are using secure transfers of information.

Interoperability testing typically takes one of three approaches:

1. **Testing all pairs**. This is often conducted by a third-party independent group of testers who are knowledgeable about the interoperability characteristics across software products and between software vendors.
2. **Testing some of the combinations**. This involves testing only part of the combinations and assuming the untested combinations will also interoperate.

3. **Testing against a reference implementation**. This establishes a reference implementation, e.g., using the accepted standard, and testing all products against this reference.

10.1.4 Quantitative Improvement

10.1.4.1 Metrics

Any application security assurance program should collect metrics, which can be analyzed and used to report the status of secure development on a periodic basis. The metrics collection and reporting could be enhanced as any application security program attains more maturity.

Some of the metrics recommended are:

- Percentage of applications and data assets in the cloud evaluated for risk classification in the past quarter and/or year
- Costs of the Application Security Assurance program in a quarter and/or in a year in a project/program for cloud-based applications
- Estimates of past loss due to security issues, if any, in the applications being developed and/or deployed in the cloud

10.1.4.2 Use of Automated SDLC Security Technology Tools and Features

People-centric SDLC activities (processes, training, and testing) are necessary but often not sufficient or viable for good application security. Where feasible, automated tools should be used to construct secure applications and automatically build security into applications.

Such tools that automatically generate technical security features are often tied into development and integration/orchestration tools. For example, technical authorization policy rules can be automatically generated (at development/integration/mash-up time) from security requirement specifications by tools that analyze applications and their interactions[70].

Similarly, some automated testing can be done at the development/integration stage, and information assurance evidence can be generated.

For cloud, this can be done at the subscriber end during development or mash-up (especially for IaaS), or the cloud provider can provide the technology (the subscriber can configure if necessary), especially in the cloud application platform for PaaS.

For SaaS, it is likely that most security automation will be built-in, configured, and operated by the cloud provider.

[70] This scientific field is called “model-driven security”, see www.modeldrivensecurity.org

10.2 Authentication, Authorization, and Compliance – Application Security Architecture in the Cloud

10.2.1 Cloud Services/Applications Development and Business Challenges

There are new potential risks associated with access to sensitive data and systems. A clear understanding of the following security risks within the application and business environment is critical for addressing the full scope of security and privacy issues in reference to the Cloud services/applications:

- **Lack of control**. This is where cloud subscribers typically lack control over cloud security policies and controls.
- **Lack of visibility**. This is where cloud subscribers typically lack visibility into cloud security policy enforcement and controls effectiveness.
- **Lack of manageability**. This is where cloud subscribers are often not sufficiently able to manage cloud application security, especially access and audit policies.
- **Loss of governance**. This is where the organization may not have direct control of the infrastructure; here trust in the provider (sometimes a naive trust) and its own ability to provide proper security is paramount.
- **Compliance risk**. This is where the cloud provider impacts the organization's ability to comply with regulations, privacy expectations, and industry standards, because data and systems may exist outside the organization's direct control.
- **Isolation failure**. This is where multi-tenancy and resource sharing are defining characteristics of the cloud. Thus it is entirely likely for competing companies to be using the same cloud services, in effect, running their workloads shoulder-to-shoulder. Keeping memory, storage, and network access isolated is essential.
- **Data protection**. This is where the organization relinquishes direct control over data; it relies on the provider to keep that data secure, and when it is deleted, the provider should ensure (or be able to prove) that it is permanently destroyed.
- **Management interfaces and access configuration**. Cloud applications are accessed and managed through the Internet, and potentially involve complex and control requirements. The risk associated with a security breach is therefore increased and proper access authorization must be carefully considered.

10.2.2 Technical Risks and Solutions

Most Cloud service providers include some form of Identity, Entitlement, and Access Management (**IdEA**)[71] in the cloud service's design. Often user authentication and authorization is delegated to the customer's user management system using a federation standard.

[71] **IdEA** - Identity, Entitlement, and Access Management

Support for Identity, Entitlement and Access Management impacts the customer in that integration in constrained by the credential passing mechanism. Infrastructure such as billing and metering that depend on identity management also present integration and migration risks.

Support for Identity, Entitlement, and Access management has integration implications for the customer. These implications include securely passing credentials and attributes, provisioning additional users, etc. Business operations within the cloud service provider are also affected; these operations include billing and accounting resource utilization. As a result, it is important to consider Identity, Entitlement, and Access management integration as an integral part of the design.

The application's IdEA capabilities (or lack thereof), such as an application's ability to accept a **SAML**[72] assertion, will impact the cloud service governance, integration, and user experience, thus understanding the IdEA requirements of the particular cloud application is a critical part of the requirements definition.

Typical IdEA requirements in a cloud application design include:

- Understanding how the cloud application will provision accounts to users, power users, and administrators – triggers for these could be links to internal HR systems or cloud-based HR platforms
- Provisioning of cloud services for service-to-service integration, e.g., internal applications to cloud-based services
- The ability to accept claim and assertions (identifiers and attributes) from a variety of sources, and entities based on federation standards (e.g., SAML, WS FED, etc.)
- The ability to make risk-based entitlement decisions about access to (and within) the cloud application, based on the identity and attributes of all the entities (users, devices, code, organization, agents) in the chain.
- A rich, risk-based entitlement language leading to access management (authoring/distribution/update, etc.) for protected resources (i.e., what is allowed for each resource)
- Support for internal security and regulatory-policy compliance requirements, such as claims-based authentication, or at a minimum role-based access control
- User activity monitoring, logging, and reporting dictated by internal policies and regulatory compliance, such as SOX, PCI, and HIPAA.

A variety of Identity providers or Service providers may generate tokens such as SAML, **OpenID**[73], or **OAuth**[74] tokens for session caching allowing a pass-through sign-on capability. Applications to be deployed in cloud should have capability to integrate with these claims/assertion services and Applications/services should be designed to support the open standards for Federation, i.e. SAML, OAuth, OpenID.

[72] **SAML** - Security Assertion Markup Language, an XML-based OASIS open standard for exchanging authentication and authorization data between security domains

[73] **OpenID** - an open standard permitting users to be authenticated in a decentralized manner

[74] **OAuth** - Open Authorization, an open standard for authorization, allowing users to share their private resources with tokens instead of credentials

The Entitlement management process will require the ability for defining, managing, and accessing the access control rules for the cloud-based applications through a centralized interface. Such an interface/service could itself be hosted on the cloud or internally and can leverage standards such as **XACML**[75]. The main challenge here is manageability: With increasing security policy and compliance complexity, IT complexity, and IT agility, the task of translating security policies into security implementation gets more time-consuming, repetitive, expensive, and error-prone and easily can amount to the bulk of security costs for end-user organizations as traditional users are managed into and out of access control lists for role-based access control, while expensive engines process these lists to ensure segregation-of-duties have not been breached.

Instead, defining a set of rules into an entitlement layer, fed by the claims, (assertions,) and attributes of the entities in the transaction significantly simplifies and enhances the control an organization has over its applications leading to the end subscriber organizations (and cloud providers) lowering their cost and improving policy implementation accuracy.

Table 1— Simple Entitlement Matrix for a Cloud HR Application

Claim / Attribute	Corporate HR Managers Access	User Corporate Access	Corporate HR Managers Home Access (Corp. Laptop)	User Home Access (Own Device)
ID: Organization Id	Valid	Valid	Valid	No
ID: User Identifier	Valid	Valid	Valid	Valid
ID: Device	Valid	Valid	Valid	No
Attrib: Device is clean	Valid	Valid	Valid	Unknown
Attrib: Device is patched	Valid	Valid	Valid	Unknown
Attrib: Device IP (is on corp. net. ?)	Valid	Valid	No	No
Attrib: User is HR manager	Valid	No	Valid	No
Access Result	Read/write access to all HR accounts	Read/write access to users HR account only	Read/write access to users HR account only	Read-only access to users HR account only

To integrate application security controls, data security and privacy protection, the services should use auditable industry standards, e.g. ISAE 3402/SSAE 16 (replaces SAS 70), PCI, HIPAA and ISO 27002. Each one comes with controls in a variety of categories that govern operation of a cloud provider's data center as well as the applications that can be hosted in such an environment.

It is important to evaluate the different security claims and make a sound decision on which standards apply for the applications and services being hosted in a cloud environment. A thorough analysis based on requirements should be done to identify service level objectives upfront to avoid major code changes to application code, deployment, and support tools for both the customers as well as the cloud provider organizations.

[75] **XACML**- eXtensible Access Control Markup Language, an OASIS standard

10.2.3 Compliance Building Blocks

Irrespective of which standards are used, achieving compliance to run an application in a cloud has some basic building blocks and the foundation of all standards is provided by the cloud provider's physical infrastructure. Infrastructure controls include things like protecting the facility from natural disasters, assuring reliable electrical power in the event of outages, and backing up data in the event of a hardware failure. They also include controls governing the cloud provider's processes and policies such as system administrative auditing, access and authorization to access the data center, and methods used for internal security reviews and how they are performed and reported.

The next layer on top of the infrastructure controls is a collection of application controls. Multiple levels of security are required, such as the transport layer that must be secure; when data leaves the data center, it must be encrypted with encryption keys under enterprise control. Some applications may need message layer security, digital signing, and other added security features in order to be compliant with some standards for storing or transmitting Personally identifiable information in order to meet privacy requirements. All such application controls for the service/application to be moved to Cloud should be identified during the design phase so they can be appropriately integrated into the architecture design and developed as per requirements. Notable standards are PCI –DSS, SOX, ISAE 3402/SSAE 16,HIPAA, and other privacy standards.

10.3 Identity, Entitlement, & Access Management for Cloud Application Security

Traditional in house enterprise applications could be protected with traditional edge security controls like firewalls, proxies, etc. This could very well meet the risk level and security requirements of the enterprise as the applications are running on trusted networks, trusted hardware, etc. The enterprise could also leverage their enterprise directory infrastructure for authenticating its users to these applications and maintain all access decisions within the applications. The perimeter for the enterprise is well defined in this case.

When the user moves these applications to the cloud, all these traditional controls are no longer effective enough to protect as these applications are running on un-trusted networks (de-parameterization). Applications could be residing with other co-tenants of same service provider (resource pooling) and could be accessed from anywhere through any type of device. This changes the very nature of security requirements for the cloud applications. As per www.rationalsurvivability.com cloud anatomy is referred as the following:

Figure 1—Cloud Anatomy

To the above referenced structure the user can now add the ways he/she can access these applications. This anatomy can then be viewed as:

Figure 2—Cloud Delivery Components

From the anatomy above, you can clearly see that your application is a window to your data, and the new perimeter is the content (data) and context by which the user tries to access that data. This makes applying security controls to the cloud applications critical. The context of accessing the data becomes very important and needs a rich collection of identifiers and attributes with which to make access decisions. With consumerization of IT, enterprises are now faced with the reality of "Bring Your Own Device" (BYOD). So device identification and the attributes of that device also become an important factor for determining the access control.

Identity should not just be viewed as a reference for authenticating the entity but also gathers more information about the user for making access decisions. Identity also includes the identities of the devices that applications run on (VM image identity), privileged users that manage that VM image (could be both enterprise users as well as service provider users), identities for other applications and services that application needs to interact with, identities of administrative users to manage the application, and external identities outside of the enterprise that need access to the application like B2B, B2C, etc. Also note that access decisions will be based on attributes that are not identity-related, and policy authoring/management tools need to support such non-identity attributes (see "Authorization management & policy automation" below).

In this section we will look into how Identity, Entitlement, and Access management affects the cloud application security. IdEA can be divided broadly into five main components:

1. Authentication
2. Authorization
3. Administration
4. Audit & Compliance
5. Policy

10.3.1 Authentication

Authentication refers to establishing/asserting the identity to the application. This is usually done in two phases. The first phase is disambiguating the identity and the second phase is validating the credential already provided to the user. Some of the main drivers for authentication to cloud applications are device independence, common and simple UI, and single protocol universal across the devices. Also, many service providers expose their services in form of **API's**[76], and these API's are designed for accepting tokens rather than the passwords.

In a regular enterprise application, the authentication is done against the enterprise user store (Active Directory or LDAP), and the authenticating credential is typically userID/Password. For cloud-based application, authenticating using the enterprise credentials gets trickier. Some enterprises establish VPN tunnel from the service provider to the enterprise network so that they can authenticate against the enterprise user directory. Even though this solution might work, enterprise should take into consideration latency issues, connectivity issues, and BCP/DR planning etc., and this solution should not be used or designed into new cloud applications. Enterprises should plan for using open standards like SAML and WS-Federation.

There is also increase in use of enterprise applications by partners and customers of the enterprise. This is also true for cloud applications. These users rarely want to maintain separate identities for their 3rd party access (but today often have no choice). So enterprises should plan for "Bring Your Own Identity" (BYOI), and the cloud application needs to be designed to consume Identity and attributes from multiple organizations.

Since the cloud applications are accessible widely through various devices, authenticating with simple userID/password should be deprecated as a solution. Enterprises should plan for using stronger authentication. Consumers should consider strong authentication for the original identity confirmation and determine the type of credential that meets their risk requirement (RSA token, OTP over SMS or phone, Smartcard/PKI, Biometrics etc.). This then will enable identifiers and attributes with a strong level of authentication to be passed to the cloud application and better risk decisions to be made about access management by the entitlement layer.

Enterprises should plan for using risk-based authentication for their cloud applications. This type of authentication is based on device identifier, geolocation, ISP, heuristic information, etc. Cloud application should not only perform authentication during the initial connection but should also perform risk-based authentication based on the transactions being performed within the application.

Cloud applications should also leverage convergence of standards where applicable, such as SAML and OAuth. As mentioned earlier in this section, cloud service API's are designed to accept tokens and not passwords, so a user trying to access cloud services from their mobile device first has to authenticate to their Identity Provider (today, probably their enterprise), and a SAML assertion is generated and passed on to cloud service provider. Upon successful validation of the SAML assertion, an OAuth token is generated and passed on to the mobile device. The mobile device then passes on these tokens to access cloud services **REST**[77] based API's.

[76] **API** - Application Program Interface
[77] **REST -** Representational state transfer, a style of software architecture for distributed hypermedia systems

10.3.2 Authorization and Access Control

Authorization in broadest terms refers to enforcing the rules by which access is granted to the resources. The Entitlement process implements business policies that in turn translate to access into enterprise resources. For cloud-based applications, authorization should not only be performed based on the content but also by the context.

For user-centric authorization model, the user is the Policy Decision Point (**PDP**)[78]. The user determines the access for their resources, and the service provider acts as Policy Enforcement Point (**PEP**)[79]. OAuth is widely used for this model, and User Managed Access (UMA) is also an emerging standard in this space.

For an enterprise-centric authorization model, the enterprise is the PDP or Policy Access Point (**PAP**)[80], and the service provider acts as PEP. In some cases, enterprises implement cloud security gateways for PEP. The enterprise customer should consider use of XACML and centralized policy management.

Cloud applications could be leveraging multiple types of services. Some services could be legacy applications exposed as web services utilizing middleware, or the web services could be native cloud web services. The diversity of the delivery supply chain, although abstracted by the web service interface, may complicate the governance process. Design time governance includes defining the services, developing the services, registering the services, and implementing policy requirement for accessing these services. Runtime governance includes discovering the services, implementing security restrictions for calling the services, enforcing security restrictions for accessing the service, and auditing all access. Use open standards like W3C **WS**[81]-policy for defining security and management policy assertions, WS-security for enforcing access restrictions, WS-trust for implementing Secure Token Service (**STS**)[82] to validate and issue tokens, and exchange token formats, etc.

There are different types of authorization models namely Role-based, Rule-based, Attribute-based access, Claims-based, and Authorization-based access control (such as **ZBAC)**[83]. Enterprises that already own Web Access Management (**WAM)**[84] solution should leverage these solutions to seamlessly protect cloud applications as well. Most of WAM products support Rule and Role-based access controls.

Application architects and designers should plan to migrate to Rule-based using claims and attributes as the source for those rules via the Entitlement process described above, and depreciate other legacy solutions.

When using attribute-based access control, the Identity Provider (**IdP**)[85] passes attributes to the Cloud Service Provider for enforcement. Identity Providers should ensure:

- Attributes attached to the identity need not strictly refer to the user identity such as first name, last name, email address, etc. It could also include IP address, location information, group affiliation, phone number, etc.
- Care should be taken for sharing attributes that directly identify the user as it raises privacy issues.

---

[78] **PDP** - Policy Decision Point
[79] **PEP** - Policy Enforcement Point
[80] **PAP** - Policy Access Point
[81] **WS** - Web Service
[82] **STS** - Secure Token Service
[83] Described in publications by Alan Karp, HP Labs
[84] **WAM** - Web Access Management
[85] **IdP -** Identity Provider

- Enterprises should also plan for the attribute complexity for making access control decisions. They should know which attribute provider to contact for a particular attribute-based on authoritativeness. There are attribute aggregators that enterprise can leverage. This could either complicate or simplify the trust. Enterprises should take into account conflict resolution complexities, handling incomplete data, etc.
- Enterprises should also take into account attribute extensibility like validation (verifiability), terms of use, date, etc.
- Enterprises should take into consideration privacy, attribute release policies, and consent. Examples include EU privacy directives, State and local laws, etc. Location of the IdP, CSP, and user (jurisdictional issues) should also be factored into this decision.
- Only minimal information required for the access control should be released.
- Enterprises should ensure attributes that are not identity-centric are also supported.
- Enterprises should ensure that access policies and entitlement policies are manageable in addition to being technically enforceable. Potential solutions include the use of policy automation technologies (maybe tied into the PaaS application mash-up tools).

The main goal of Claims-based access control is controlled sharing of the information. The claims are based on the context of the transaction. When planning to use claims-based authorization, an enterprise should consider:

- Usage of meaningful claims (verified email address instead of just email address)
- Type, surety, freshness, and quality of the claim (if the claim is cached outside of claims provider then the freshness of the claim is lost).
- Appropriate authority of the claims based on the context, e.g., a telecom company having authority to verify your phone number, an email provider having authority to verify your email address, etc.
- Utilization of claim brokers where possible as they could be used for abstraction from various claims providers, e.g., they could create a package of claims at desired confidence levels and create a central point for user permission
- Minimal release of claim as required by the transaction

The cloud application could also be a mash-up of other cloud applications running on the same or different service providers. The enterprise should plan for how the users are authenticated seamlessly across all these cloud applications and how the users' profiles such as group association, entitlements, roles, etc. are shared across these cloud applications for granular access controls. Enterprises are recommended to use open standards for this use case (SAML, OAuth, XACML, etc.).

10.3.3 Administration/Management

Identity Management (**IDM**)[86] within the enterprise is mainly focused on managing users (provisioning) and managing access policies (for enterprise applications). IDM is a very important component of IdEA for not only providing timely access to the users but also timely revocation of access when the user leaves or timely management of access when the user moves to a different role. Within the enterprise, identity management is usually tightly integrated and is directly connected to the data stores (users, policies, etc.); in most deployments, it is also heavily customized. Due to the distributed nature of cloud applications applying the same principle becomes a non- starter, as IDM might not have direct access to the data stores on the service provider. Moreover, there are no standard API's for provisioning. Many service providers have not adopted Service Provisioning Mark-up Language (**SPML**)[87].

IDM in the context of cloud computing should not only manage the user identities. It should extend this to manage cloud application/services identities, access control policies for these cloud applications/services, privileged identities for the applications/services, etc.

Current federated provisioning is implemented with proprietary API's exposed by the service provider. The PUSH model that is followed by the enterprise IDM will not work with cloud applications as it might overload the service providers.

The new emerging standard is Simple Cloud Identity Management (**SCIM)**[88] with the main goal of this standard to make the management of identities cheaper with easier and faster implementation. The secondary goal is to ease the migration of user identities into and out of the cloud. SCIM is simple because it uses well defined core schema, cloud-friendly because it uses RESTful API as supported by many cloud service providers, and supports identity management because it works with existing protocols such as SAML, OpenID connect etc. Based on these facts (at the time of writing), SCIM might get adopted as an industry standard for identity provisioning.

Some of the challenges that enterprises consider for identity management are:

- How to sync changes about identities/access between enterprise -> cloud, cloud -> cloud, cloud -> enterprise
- How to de-provision identities and access across the enterprise and cloud
- How to author/update/manage access policies in a manageable, scalable, low-maintenance, low-cost way.

The current solution for many enterprises is the adoption of a hybrid IDM solution that spans both enterprise and cloud.

Access policy management is a major application security challenge and often requires maximum security automation as a solution: Security policy automation is particularly important for cloud computing because cloud users will demand support for regulatory compliance policy management from cloud providers, but will at the same time judge the financial ROI by the same measures as they do for cloud computing in general, i.e., by how much it cuts their up-front capital expenditure and their in-house manual maintenance cost.

[86] **IDM -** Identity Management
[87] **SPML** - Service Provisioning Mark-up Language
[88] **SCIM** - Simple Cloud Identity Management

10.3.4 Audit/Compliance

Enterprises using cloud services should answer three fundamental questions:

1. What cloud resources does a user have access to?
2. What cloud resources does a user actually access?
3. Which access policy rules were used as a basis for a decision?

With current cloud deployments, enterprise customers have very limited visibility into cloud service providers for audit data. An enterprise needs access to this data not only for meeting the business driven compliance but also to meet industry regulations and also deal with fraud disputes.

Currently the IDM market is moving towards Identity and Access Governance (**IAG**)[89] market. Enterprises should also consider use of SIEM (Security Incident & Event Management) tools to correlate cloud application access log data and your policy data to generate policy compliance reports as well as use of auditable industry standards such as ISAE 3402/SSAE 16, HIPPA, DSS PCI, ISO27002, etc.

General IdEA considerations for cloud application security are:

- Identity, Entitlement, and Access management should not be an afterthought; rather, it should be integrated into an application's SDLC starting with the requirements gathering.
- During the design phase consider how to control access to the application using a Claims-based access whenever possible.
- Consider using tools like SAPM (Shared Account Password Management) for managing highly privileged accounts within the application. This allows for segregation of duties and least privilege.
- If the enterprise already has web access management tools, ensure that those tools can be extended into a cloud environment, i.e., by adding a SAML capability.
- Cloud applications might need to leverage services offered by service providers such as logging, database connectivity, etc. Most service providers expose these as web services or API. Access to these services could be controlled by OAuth tokens. Thus cloud applications should take into consideration supporting various token types like OAuth, API keys, etc.
- Ensure that you follow an agile development process and that the application is built into modularized components. This allows the application to leverage new emerging standards in the future like Mozilla's browserID, Microsoft's U-Prove, and Kantara Initiative's UMA (User Managed Access).

Be aware of the threats for cloud applications, which include:

- **Spoofing.** Assuming the identity of another user
- **Tampering.** Modifying the data on transit

[89] **IAG** - Identity and Access Governance

- **Repudiation.** Denying the origin of transaction (request or response)
- **Information disclosure.** Unauthorized disclosure of data
- **Denial of Service.** Affecting the availability
- **Elevation of Privilege.** Assuming the role or entitlement

These threats could be addressed by IdEA as follows:

- **Spoofing.** Authentication (strong authentication)
- **Tampering.** Digital Signature or Hash (As used in SAML assertions)
- **Repudiation.** Digital signature (as used in SAML assertions), audit logging
- **Information disclosure.** SSL, encryption (Strictly not IdEA specific)
- **Denial of Service.** Security Gateways (Web Services security gateways)
- **Elevation of Privilege.** Authorization (OAuth)

10.3.5 Policy Management

Access policy management (often called "authorization management" when done entitlement-centric) is the process of specifying and maintaining access policies to resources in access policies, based on attributes including caller-related identities and related attributes (e.g. caller authentication), context attributes (e.g. environment/business/IT related), and target-related attributes (e.g. throttling or QoS access policies)[90].

Entitlement management forms part of authorization and access management, which additionally includes authoring and maintaining policies for attributes that are not identity-related but are required (in addition to identity and its attributes) to make a meaningful access decision.

Entitlement/Authorization also takes attributes into account that are not related to an identity, e.g.:

- General state of the IT landscape, business/business process, interconnection of IT systems or business processes, or environment, etc. (e.g. crisis level, emergency situation)
- Other decisions made by other entities (e.g. approvals, prior decisions)
- Attributes related to the protected target resource (e.g. QoS or throttling policies)

Typically the authorization management, decisioning, and enforcement process is performed in one of three places:

1. Using a central/external Policy Enforcement point / Policy Server / Policy-as-a-Service
2. Embedded as part of the Cloud application
3. Using an Identity-aaS or Persona-aaS (an entities Persona is its Identity with selected attributes).

[90] Lang, U. "Access Policies for Middleware", PhD thesis, University of Cambridge Computer Laboratory, 2003

10.3.5.1 Cloud Issues vs. Policy Management

Authorization/entitlement management for cloud faces several issues[91].

Firstly, entitlement management for cloud has the specific issue that cloud subscribers often do not have sufficient control over technical access policy decision-making and enforcement in the cloud infrastructure. Most cloud providers today do not offer subscriber-configurable policy enforcement points (e.g. based on the OASIS XACML standard), and cloud providers naturally cannot pre-configure subscriber-specific policies for subscribers (because they are subscriber-specific).

Secondly, an entitlement management complexity for interconnected clouds (mash-ups) is that access needs to be controlled for the interconnected cloud mash-ups, and not only for each individual cloud node. This means that policies need to be authored for service chains and delegation across the interconnected cloud mash-up in mind.

10.3.5.2 Authorization Management Best Practice

- Establish whether an identity/entitlement-centric perspective is the best way for your organization to author and maintain access policies; in many cases a protected-resource-centric perspective may be easier to author and maintain because the goal is often to protect resources, and policies are often distributed to the protected end-systems for automatic policy enforcement (e.g. in entitlement/authorization management systems). In those cases identity is merely one attribute in an access policy that is written with the goal of enforcement at the protected end-system in mind.
- Make sure that policies are specified in a manageable form. This includes specifying policies that are generic, specified at a sufficiently high level of abstraction, and expressed close to the understanding of the relevant organization/business/humans. Mechanisms and tools are available to generate the detailed technical access policy rules from such a manageable form (e.g. using model-driven security policy automation).

10.3.5.3 Architectures for Interfacing to Access Policy Providers

Policy Decision/Enforcement Points (PEP's/PDP's) using standard protocols (e.g. XACML) or proprietary protocols (e.g. direct web service or other middleware calls) can access policy servers (which contain the rules for an interconnected cloud mash-ups). The architecture is usually one (server) to many (PDP/PEP's) if the policy covers a single trust domain (e.g. an enterprise intranet). However, in more large-scale deployments, there can be several federated policy servers that service many different PDP/PEP's. Several access management products now support authorization management rules (e.g. in XACML) that can be used to express entitlements for identities. In addition, several authorization management products are available that can be used to author authorization policies from a more target-resource-centric perspective.

10.3.5.4 Provisioning of Access Policies

In addition to identity + attribute provisioning, access policies need to be provisioned (see above "10.3.5.3 Architectures for Interfacing to Access Policy Providers"). Moreover, non-identity attributes need to be provisioned, e.g., from directory services or other attribute sources. Both need to be provisioned to the PDP/PEP's, and timeliness and correctness play a critical role.

[91] Details: Lang U, Schreiner R, Analysis of recommended cloud security controls to validate OpenPMF, Information Security Technical Report (2011), doi:10.1016/j.istr.2011.08.001

10.3.5.5 Managing Access Policies for Cloud

Making authoring and maintaining access policies manageable is a major challenge; there are typically simply too many technical rules to manage, used policy languages and attributes that do not match the understanding of human administrators, technical rules that need to be updated frequently to remain correct after each time systems change (e.g. for agile cloud mash-ups), and it is hard to establish that the level of confidence/assurance of the technical policy enforcement matches the intent of the human administrator. As a consequence, it is critical to carefully plan the tools and processes to make this access policy authoring/updating process manageable through automation.

Current solutions include automated approaches to turn high-level security policies into (low-level) technical access rules, including:

- Model-driven security[92], the tool-supported process of modeling security requirements at a high level of abstraction, and using other information sources available about the system (produced by other stakeholders). These inputs, which are expressed in Domain Specific Languages (DSL), are then transformed into enforceable security rules with as little human intervention as possible. It also includes the run-time security management (e.g. entitlements / authorizations), i.e., run-time enforcement of the policy on the protected IT systems, dynamic policy updates, and the monitoring of policy violations.
- Clustering technical access rules into similar groups to reduce the complexity
- Visual attempts to make technical policies easier to understand

10.3.5.6 Authorization in the Cloud Best Practice

- Carefully consider whether a protected-resource-centric perspective to authoring access policies may be more suitable for your environment than an identity-centric perspective.
- Ensure manageability of access policies, especially for dynamically changing cloud mash-ups. This includes consistency of policy authoring, policy distribution, enforcement, and update. Consider the use of automated tools and approaches (e.g. model-driven security) to generate the technical access rules needed for policy enforcement.
- Designate clear responsibilities for policy management and policy auditing.
- Ensure your cloud provider offers authorization management PEP's/PDP's that can be configured with the subscriber-specific authorization policy, and that your policy server can interface correctly with the policy selected.
- Consider the use of "policy-as-a-service" as the policy server if you need a central policy server for a cloud mash-up.

Current best practices for selecting authorization services:

- The most important authorization management services feature is cloud subscriber policy manageability, because managing access policies is the biggest challenge around authorization.

[92] NIST IR 7628

- Services should allow for as-automatic-as-possible technical policy generation (and update!) from human-intuitive, generic security policy requirements.
- If politically viable for your organization, and if available for you, "policy-as-a-service" should be considered as an option of outsourcing the policy authoring and updating. Most likely this will be acceptable within community clouds where the "policy-as-a-service" is offered to a closed community.
- Ensure services have an import and/or export function into standards such as OASIS XACML.
- Ensure services can interface with PEP/PDP's installed in the cloud infrastructure and with Policy Monitoring Points for incident monitoring/auditing.

10.4 Application Penetration Testing for the Cloud

A penetration test involves the process of evaluating the residual vulnerabilities present in the application or system layer that can be potentially exploited by an external or internal hacker with malicious intent. The test would typically involve active analysis of the surfaces of the application or system as a "black box" and attempts to identify typical vulnerabilities that can be prevalent as a result of bad programming or hardening practices.

Open Web Application Security Project (**OWASP)**[93] in its OWASP Testing Guide V3.0 recommends nine types of Active Security Testing categories as follows:

1. Configuration Management Testing
2. Business Logic Testing
3. Authentication Testing
4. Authorization testing
5. Session Management Testing
6. Data Validation Testing
7. Denial of Service Testing
8. Web Services Testing
9. Ajax Testing (RIA Security Testing)

The above categories of Security testing will be equally applicable for an application that is going to be deployed on the Cloud as the nature of Application Vulnerabilities from a technical perspective is not going to change. However, depending upon the type of Cloud Deployment Model additional threats vectors (that would have not come into the equation for a non-cloud deployment) could be induced.

An example of such a threat vector in a SAAS deployment would be induced by multi-tenancy when the same application run time is being used to service multiple tenants and their segregated data.

[93] OWASP - Open Web Application Security Project, www.owasp.org

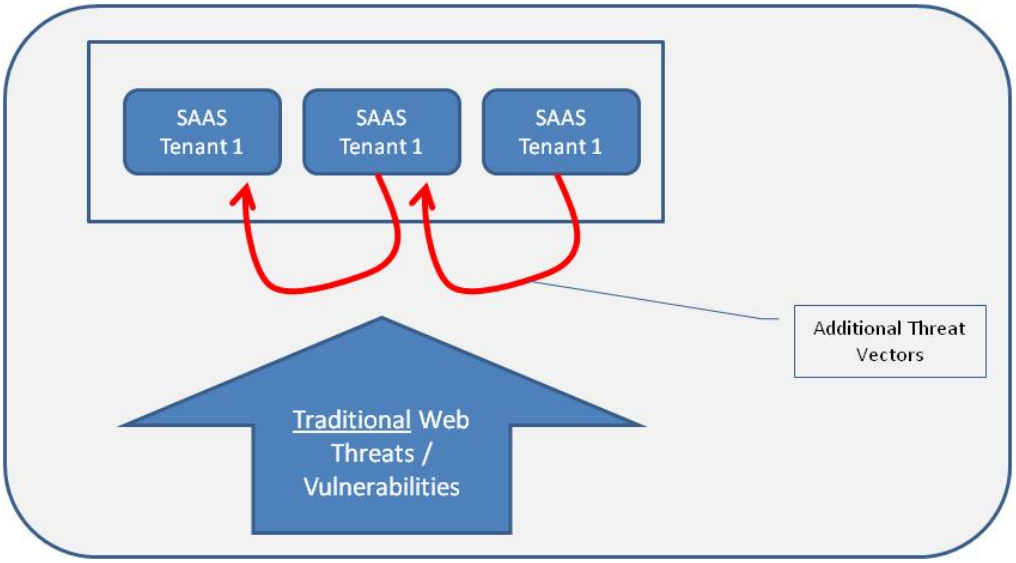

Figure 3—Threat Vector Inheritance

Additional classes of testing will need to be developed and included to address threats that arise as a result of the deployment model of the Application in the Cloud. An illustration of the same is provided in the table below.

Table 2— Threat Vector Inheritance

CLOUD MODEL ON APPLICATION IS BEING DEPLOYED	ADDITIONAL THREATS INDUCERS	EXAMPLES OF THREATS	TRADITIONAL SECURITY TESTING CATEGORIES STILL RELEVANT	ADDITIONAL TESTING CATEGORIES
SAAS	Multi-tenancy at an Application Level	A different tenant using the same SAAS infrastructure gains access to another tenants data through the web layer vulnerabilities (a privilege escalation)	▪ Configuration Management Testing ▪ Business Logic Testing ▪ Authentication Testing ▪ Authorization testing ▪ Session Management Testing ▪ Data Validation Testing ▪ Denial of Service Testing ▪ Web Services Testing ▪ Ajax Testing (RIA Security Testing)	▪ Multi-Tenancy Testing (an extension of privilege escalation)
PAAS	Multi-tenancy at a Platform	Same as above	Same as above	Same as above

	Level			
IAAS	Multi-tenancy at an Infrastructure Level	Deficiencies in virtualization security (improper implementation of VM zoning, segregation leading to inter VM attacks across multiple IAAS tenants)	▪ Traditional Infrastructure Vulnerability Assessment (need to "define" this)	▪ Inter VM Security / Vulnerability Testing

10.5 Monitoring Applications in the Cloud

As with other aspects of cloud security, what and how one monitors a cloud-based system varies with the type of cloud under consideration. What it means to monitor applications in the cloud and how to monitor different types of cloud applications are explained in detail below.

10.5.1 Application Monitoring in the Cloud: Give and Take

For this document, we are limiting "monitoring" to focus on application security monitoring. In particular, the following categories of metrics should be addressed:

1. **Log monitoring.** It is not just to archive logs for compliance purposes. Understand the potential output that could be sent to these logs, and monitor for actionable events. An application logging errors is of zero use unless a process exists to detect and respond to those errors.

2. **Performance monitoring**. This plays a large factor in shared computing. A significant change in the performance of one application could be the symptom of another customer using more than their fair share of a limited resource (e.g., CPU, memory, SAN storage), or it could be the symptom of malicious activity either with the application being monitored or with another application in the shared infrastructure.

3. **Monitoring for malicious use**. This is a blend of auditing and monitoring required to be successful. The enterprise must understand what happens when a malicious user attempts to gain access, or use permissions that they do not have. Audit logs must log failed (and successful) login attempts. Do data-validation functions log anything? If an application experiences a significant increase in traffic load, is an alert created anywhere?

4. **Monitoring for compromise**. Here the key is how quickly and efficiently an organization responds to the compromise. Depending on the complexity of the application, determining compromise might be relatively easy (e.g., "User A is logged in twice") or may require more effort (e.g., developing heuristic algorithms to monitor data usage). This is a good example of an item that, when addressed earlier in the SDLC, can be easier to manage.

5. **Monitoring for policy violations (especially access control)**. It is also important to monitor and audit how a Policy Decision Point came to a decision, i.e., which policy rules were applied to make a specific access decision. This is in line with a general policy-driven monitoring approach that avoids the typical monitoring problems of false-positives and incident overload.

These are some of the key concepts behind log monitoring – the "take" side of the equation. Equally important, the developer of an application is responsible for the "give" side: His application must provide a solid logging subsystem to allow the monitoring system to efficiently do its job:

1. **Easily parsable**. Logs should be written in a format that can be easily parsed by a separate system. A good example would be using a well-known and accepted format such as XML. A bad example would be writing log entries of non-delineated, multi-line text output.

2. **Easily readable**. Unless writing logs in a binary format that will, without a doubt, never be directly read by a human, a log entry should be understandable by a person with a technical background, familiar with the application.

3. **Well documented**. It is not enough to just write logs to a file. Error codes need to be documented and should be unique. If a particular log entry has a known path to resolution, document the resolution, or provide a reference to it.

10.5.2 Monitoring Applications in Different Cloud Types

With an IAAS-based application, monitoring the application is almost "normal," compared to "legacy" applications deployed in non-shared environments. The customer needs to monitor issues with the shared infrastructure or with attempted unauthorized access to an application by a malicious co-tenant.

Monitoring applications deployed in platform-clouds requires additional work. Unless the platform provider also provides a monitoring solution capable of monitoring the deployed application, the customer has two choices: Either write additional application logic to perform the monitoring tasks within the platform or send logs to a remote monitoring system, be that the customer's in-house monitoring system, or a third party service.

As SAAS-based applications provide the least flexibility, it should not come as a surprise that monitoring the security of these types of applications is the most difficult. Before using a SAAS product, customers must have a thorough understanding of:

- How does the provider monitor their applications?
- What type of audit, log, or alert information will the provider send to the customer? Does the customer have the ability to select what information they will receive?
- How will the provider transmit this information to the customer? (Twitter? Email? Custom API?)

One final point when considering application security monitoring in the cloud: While providers (or 3rd party cloud monitoring services) may have built a monitoring system to monitor a customer's applications, those monitoring systems are monitoring hundreds, if not thousands, of customers. The provider, as a business, wants this monitoring system to work "well enough." If the customer has the resources, running his/her own monitoring system that monitors just his/her applications will almost always be more responsive and more informative than that of a cloud provider.

10.6 Recommendations

10.6.1 Security Assurance Recommendations

- Functional and regulatory security and privacy requirements are defined that meet the needs of cloud development and/or deployment.
- A detailed assessment of the attack vectors and risks in the cloud environment are understood, and the mitigations strategies are integrated into the requirements.
- An impact assessment for all risks and attack vectors is undertaken, and documented, together with the potential for loss or damage from each scenario.
- Security and privacy requirements and efforts should be prioritized on likelihood and impact.

10.6.2 Risk Analysis Recommendations

- Risk analysis of the applications for security and privacy (confidentiality, integrity and availability) are undertaken, and threat models should be built and maintained.
- Risks from the perspective of development and deployment in the cloud should be analyzed and related threat models maintained.
- Attack vectors and impact analysis specific to cloud architectures should be catalogued and maintained.
- Traceability between security assurance features and all identified risks / threats should be maintained.

10.6.3 Architecture Recommendations

- Secure software architecture frameworks should be developed and maintained.
- Cloud computing architecture patterns that explicitly mitigate threats (for example, from "Open Security Architecture"[94] or TOGAF/SABSA[95]) should be used.
- Reusable building blocks in the application architecture are available for mitigating commonly known security and breach scenarios.
- Cloud-specific secure data architectures should be used to enhance the chosen security architectural framework, which will address cloud specific issues and threats, such as:
  - The monitoring of dynamic database servers
  - Understanding where the database is exactly hosted at any point in time

[94] www.opensecurityarchitecture.org
[95] www.opengroup.org

- Centrally logging all activity, across disparate (potentially global) systems to provide a holistic view of the application and flag suspicious events
- Define where encryption must be used (see Domain 12)
- Provide adequate segregation of duties within the system, the data, and all privileged activities by third parties, capable of being monitored by staff of the organization that owns the data

10.6.3 Penetration Testing Applications on Cloud Recommendations

- Carry out regular Web Application Penetration Testing to check for OWASP Top 10 vulnerabilities
- Categorize vulnerabilities based on criticality / Impact and have a process for remediation
- Carry out manual tests from a multi-tenancy perspective to validate that privileges cannot be escalated or data segregated based on lack of session enforcement.
- For applications being migrated to an IAAS or PAAS environment, a security assessment needs to be carried out to ensure that the underlying security controls such as VM zoning and segregation, virtualization security, etc. has been effectively put in place and does not pose a risk to the application ecosystem.

REFERENCES

[1] The Building Security In Maturity Model. http://bsimm.com/

[2] OpenSAMM – Software Assurance Maturity Model. http://www.opensamm.org/

[3] DAVIS, NOOPUR. Secure Software Development Life Cycle Processes. Software Engineering Institute

[4] SP-011: Cloud Computing Pattern. http://www.opensecurityarchitecture.org/cms/en/library/patternlandscape/251-pattern-cloud-computing

[5] KRUTZ, RONALD L. and VINES, RUSSEL DEAN. 2010. Cloud Security- A Comprehensive Guide to Secure Cloud Computing. Wiley Publishing, Inc., Indianapolis, IN.

[6] SARNA, DAVID E.Y. 2010. Implementing and Developing Cloud Computing Applications. Auerbach Publications.

[7] BELK, MARK, COLES, MATT, et al. 2011. Fundamental Practices for Secure Software Development: A Guide to the Most Effective Secure Development Practices in Use Today, 2nd EDITION. Software Assurance Forum for Excellence in Code. http://www.safecode.org/publications/SAFECode_Dev_Practices0211.pdf

[8] RITTINGHOUSE, JOHN W. and RANSOME, JAMES F. 2009. "Cloud Security Challenges" in Cloud Computing: Implementation, Management, and Security. Auerbach Publications. http://www.infosectoday.com/Articles/Cloud_Security_Challenges.htm

[9] Guidelines on Security and Privacy in Public Cloud Computing. Computer Security Division Information Technology Laboratory. 2011. National Institute of Standards and Technology - Gaithersburg, MD 20899-8930 http://csrc.nist.gov/publications/drafts/800-144/Draft-SP-800-144_cloud-computing.pdf

[10] Homomorphic Encryption. Making Cloud Computing More Secure. http://www.technologyreview.in/computing/37197/

[11] Cloud Data Protection. Best Practices. http://www.ciphercloud.com/blog/?cat=10

DOMAIN 11 //
ENCRYPTION AND KEY MANAGEMENT

For a security professional, it is obvious that if an organization needs to store data and doesn't trust those who can access or use the data, then the data must be encrypted. Inside an on-premise data center where the organization controls all assets, data is encrypted because some regulations say the data *must* be encrypted (PCI DSS for example).

In the cloud, where there are multiple tenants and administrators working for someone else it would seem obvious that much more data would need to be encrypted. If that is the case, how do those processes work and how does the organization manage their keys? Encrypting everything increases complexity. On the other hand, is it even necessary to encrypt these volumes of data if it causes business process complexity amongst other issues? Is there another way to reduce the need to encrypt data and subsequently manage the keys? This chapter looks at these issues.

Overview. This domain will address the following topics:

- Introduction to Encryption
- Alternative approaches to Encryption
- Cryptography in cloud deployments
- Encryption in Cloud Databases
- Key management in the cloud
- Storage and safe-guard of keys

To encrypt or not to encrypt? That is the question. If yes, how do I manage the keys? If no, are risks too high?

11.1 Introduction to Encryption

Data classified as confidential for reasons of regulatory compliance or corporate secrecy must be protected. As confidential information that is currently managed within internal systems increasingly moves to the cloud, it must be protected with the same diligence. Moving data to the cloud does not remove any requirements for confidentiality and data protection. The loss of control of data outside the secured corporate perimeter (de-perimeterization) increases the complexity of protecting data and increases the risk of compromise.

There are a number of factors to consider regarding data encryption in the cloud, which include:

- Protecting data through encryption as it moves to the cloud requires more than just ensuring that a secure transfer channel (i.e. TLS) is used. Encrypting the transfer of data to the cloud does not ensure the data is protected in the cloud. Once data arrives in the cloud, it should remain protected both at rest and in use.
- For unstructured files that must be protected when stored or shared in the cloud use data-centric encryption, or encryption embedded into the file format whenever practical to apply protection directly to files.

- Understand how all encryption / decryption keys will be managed for the entire lifecycle of the data. Whenever possible avoid any reliance on cloud providers to protect and appropriately use the keys that protect your critical information.
- Avoid opportunities for lapses in the employee safeguards of others, or of regional laws that may provide undesired, but mandated access to your encrypted files. If only you have the keys, only you can access your files.
- Do not forget to protect files that are often overlooked, but which frequently include sensitive information. Log files and metadata can be avenues for data leakage.
- Encrypt using sufficiently durable encryption strengths (such as AES-256) that comply with the same corporate and regulatory mandates used for encrypting internally maintained files. Use open, validated formats and avoid proprietary encryption formats wherever possible.

11.2 Alternative Approaches to Encryption

There are good reasons to look at alternate approaches to encrypting data in the cloud. For many organizations sending data into the cloud is equivalent to transferring custodial relationship.

For those organizations that have issues with sending unsecured data outside their organization there are alternatives:

- **Tokenization**. This is where public cloud service can be integrated/paired with a private cloud that stores sensitive data. The data sent to the public cloud is altered and would contain a reference to the data residing in the private cloud.
- **Data Anonymization**. This is where (for example) Personally Identifiable Information (**PII**)[96] and Sensitive Personal Information (**SPI**)[97] are stripped before processing.
- **Utilizing cloud database controls**. This is where the access controls built into the database are deemed to provide adequate levels of segregation.

As a rule, good data management practices are essential before moving data into the cloud, to understand whether all or just some of the data need to be encrypted, protected by an alternative method, or not protected at all.

When evaluating what to protect through encryption of other alternative methods there are risks of data sharing[98] that can be broken down into two primary categories: disclosure and misuse, under the following areas:

- **Accidental public disclosure**. Making information or data readily available to the general public via publication or posting on the web.
- **Accidental or malicious disclosure**. The act of making information or data available to a third party(s) as a result of inadequate data protection.

[96] **PII** - Personally Identifiable Information
[97] **SPI** - Sensitive Personal Information
[98] http://www.caida.org/data/sharing/

- **Compelled disclosure to third parties**. The obligations of having to respond to subpoenas requesting data disclosure in lawsuits.
- **Government disclosure**. The release of data to government entities, either by law, or by court order (such as the Patriot Act).
- **Misuse of user or network profiles**. The ability to analyze and data mine to derive sensitive information from seemingly benign traffic data, and thereby reveal user behaviors, associations, preferences or interests.
- **Inference misuse**. Being able to synthesize first-order or second-order identifiers to draw inferences about a person's behavior or identity.
- **Re-identification and de-anonymizing misuse**. Having access to enough "anonymized" information to be able to infer the original subject.

11.3 Cryptography in Cloud Deployments

There are two complementary concepts used in the encryption section, they are:

- **Content Aware Encryption**. Used in Data Leak Prevention, content aware software understands a data type or format and encrypts based upon policy settings. For example a credit card number is encrypted in an email being sent to law enforcement.
- **Format Preserving Encryption.** Encryption that preserves format is a result that encrypts a message and produces a result like the input message. For example, a 16-digit credit card number is a 16-digit number after encryption, a telephone number would look like a telephone number, and an English word would look like an English word.

The ability to encrypt from the enterprise to the cloud without user intervention is to the preferred way to make data safe. Content aware software can be leveraged for public cloud encryption if the software can be configured to be protocol aware as well and encrypt fields in a REST http transaction to a public cloud application. The Data Leak Prevention (**DLP**)[99] use case today is met by products that can enforce data protection leaving the enterprise, usually by email, and encrypt data before the transaction leaves the enterprise. This principle can be used in cloud data protection; however, the DLP product may generate alerts. A content aware service would need to detect, encrypt, and log but not alert.

Format preserving encryption takes content aware a step further by being sensitive to the data needing encryption and maintaining the data format and type. For example, with conventional encryption, a credit card being encrypted would render a **cipher-text**[100] that would no longer be a 16-digit number. Format preserving encryption would generate a cipher text value that is 16 digits in addition to being encrypted.

By also preserving the data type and format the service providing encryption can then easily change values in line over a wide variety of protocols. The key challenge to format preserving encryption is in encrypting large clear text values such

[99] Data Leak Prevention (**DLP**) products have an enforcement mode that detects data leaving a secured device or the enterprise and encrypts it.

[100] **Cipher text** - The result of an encryption operation. The input is known as clear text.

as an email stored in the cloud. Bulk scale encryption is normally how text values are encrypted using block **ciphers**[101]. In the format preserving case, each word would be encrypted into a character string of the same length, which takes time. The result, however, would be cipher-text data values that can be stored in fields of the same data-type as the original plain text.

Encryption in cloud applications poses some issues for business applications that the application architecture needs to address. These are:

- If data in the application is needed to search for records or objects, then an encrypted **primary key**[102] would make that difficult.
- If the cloud application set contains batch jobs or other types of processes that work on sensitive data, particularly PII and SPI data, and those processes are moved to the cloud, that situation will complicate key management.

An application that needs to find records or objects in a database may choose to develop another way to store a unique value such as tokenization. Tokens are often used in credit card environments to ensure the credit card number is minimally accessed in applications. A unique token generated from the value can be used to develop a new primary key that the application can use without exposing sensitive data in a public cloud.

As will be discussed in section 11.4 below, where possible, keys should not be stored in the cloud and must be maintained by the enterprise or a trusted key management service provider.

Processes that need to operate on clear text data and run in the cloud with other business applications and data must have access to keys or a service in order to perform their functions. See section 11.4 for more details on key management in the cloud.

11.3.1 Encryption in Cloud Databases

The first thing to consider is whether it is necessary to encrypt the data. All databases provide the ability to restrict access to data. If properly implemented, that may be enough to protect confidentiality.

Other reasons that may require the encryption to protect data stored in the database are:

- To hide it from those with privileged access to the database (Database Administrators (**DBA's**)[103], for example)
- To comply with legal statutes (such as California's SB1386 law)
- To store it in a schema for which the data owner cannot control the account credentials accessing the data (using shared accounts, for example)

When using a cloud database and particularly SaaS solution employing a database, the database ability to function correctly may be compromised unless it can operate on the encrypted data, necessitating the database or cloud application to have access to the keys.

---

[101] **Ciphers** - Algorithm based software/hardware that perform encryption/decryption and signing/verifying
[102] **Primary key** - A database column/field/attribute that is used to uniquely identify records in a database
[103] **DBA** - Database Administrator

Data encryption comes at the price of complexity and performance, and there are effective alternatives to encryption:

- **Use object security**. Use SQL grant and revoke statements to restrict which accounts can access the data. The accounts to which you grant access must be controlled to ensure that you are only allowing access to authorized users.
- **Store a secure hash**. Rather than storing the data directly, store a hash of the data. This allows your program to prove that the holder has the correct value without actually storing it.

11.4 Key Management

One of the more difficult processes in public cloud computing is key management. The multi-tenant model of public cloud solutions causes key management issues for processes running there.

The easiest use cases are those that have applications running in the public cloud and keys that encrypt data going to the public cloud from the enterprise are used within the enterprise only. As described in section one, there are encryption engines that can encrypt data on the way out and decrypt data on the way back in. An application using cryptographic keys gets complicated when other processes, such as batch processes, need access to keys to decrypt data and those processes reside in the public cloud.

Enterprise users of encryption need to have keys of their own so that a single shared key is not used across the enterprise. The easiest way to accomplish such specific keys is a cryptographic engine for each user or **entity**[104] to have keys assigned (and managed) based on the entities identity. In this way, anything that is encrypted specifically for an entity is maintained for that entity. If an entity needs to share access to data in a group setting then group level keys can be associated with the application that maintains group access, and entities within that group can share the keys. The keys should be maintained within the enterprise as discussed earlier in this section.

Where data is stored in a public cloud environment, there are problems when exiting that environment to be able to prove that all data (especially PII or SPI data, or data subject to regulatory assurance regimes) has been deleted from the public cloud environment, including all other media, such as back-up tapes. Maintaining local key management allows such assurance by revoking (or just deleting/losing) the key from the key management system, thus assuring that any data remaining in the public cloud cannot be decrypted.

11.4.1 Storage and Safe-Guarding of Keys

Encrypting data has little value if both providers as well as users of cloud services do not vigorously enforce the processes around key management.

On the provider side, a lack of **SOD**[105] (Segregation of Duties) around access to key servers and servers having encrypted data should be a cause for concern, as well as DBA's having access to individual keys for databases, or the architecture of the database service reliant on a single key.

[104] **Entity** - For the purpose of identity, could be a user, code, a device, an organization or agent
[105] **SOD** - Segregation of Duties

Controls around protecting the keys themselves, by using **KEK**[106] (Key Encrypting Keys) and generation of encryption keys in-memory, and only storing the encrypted key of key servers are all valid architectural solutions that should be considered when architecting any solution.

Keys managed on the client side that protect keys on devices that are not themselves secure (such as mobile devices) or devices which do not have the same level of controls as the encrypting system itself should be a cause for concern.

11.5 Recommendations

General Recommendations

- Use best practice key management practices when using any form of encryption/decryption product.
- Use off-the-shelf technology where possible to get the best practices from a credible source.
- Use best practice key management practices and obtain technology and products for encryption, decryption, signing, and verifying from credible sources.
- It is highly recommended that organizations maintain their own keys or use a trusted cryptographic service from a source that currently maintains such a service.
- If an organization needs to run analytics or other processes using data stored in the cloud then the organization should develop on top of a platform such as Hadoop and have that data derived from the cloud source. Such development platforms, including Hadoop, have their own set of security issues but those are beyond the scope of this chapter.
- Key scoping can be maintained at the individual or group level.
- Group access can be managed with off-the-shelf technology such as DRM systems and other software running on the desktop/laptop that encrypts disks, folders, and email messages.

Recommendations - Encryption within Databases

- Use standard algorithms. Do not invent/use proprietary scrambling techniques. Proprietary encryption algorithms are unproven and easily broken.
- Avoid old insecure encryption standards such as Data Encryption Standard (**DES**)[107].
- Use object security. You should still use basic object security (SQL grant and revoke statements) to prevent access to even the encrypted data.
- Do not encrypt primary keys or indexed columns. If you encrypt a primary key, you will have to encrypt all referencing foreign keys. If you encrypt an indexed column, you may end up with slow queries when trying to use the encrypted value.
- Use a columnar approach to encryption (since big data system uses this).

[106] **KEK** - Key Encrypting Keys
[107] **DES** - Data Encryption Standard

11.6 Requirements

- ✓ In order to maintain best practices and pass audits the organization should manage their keys in the custody of their own enterprise or that of a credible service from a cryptographic service provider.
- ✓ Keys used in existing encryption technology such as DRM and disk encryption products should be managed by central, internal to the enterprise, key storage technology. Hardware Security Modules (HSM) should be used to store keys as well as process cryptographic operations such as encryption/decryption, signing and verifying.
- ✓ Enterprise users should go through a registration process to enable cryptographic operations and other processes in the enterprise, such as Content Aware or Format Preserving systems can access encryption/decryption keys as needed.
- ✓ Deploy technology integrated into corporate systems based on the identity of all components in the processing chain to make entitlement decisions.
- ✓ Manage keys used by the cryptographic processes using binding cryptographic operations.
- ✓ Use existing systems such as **E-DRM**[108] or DLP if possible.
- ✓ Binding cryptographic operations and key management to corporate identity systems will provide the organization with the most flexible integration and uses technology that the organization already knows works and has been audited and or reviewed.

[108] **E-DRM** - Enterprise Digital Rights Management. A process that protects content such as internal corporate communications or copyrighted material.

DOMAIN 12 //
IDENTITY, ENTITLEMENT, & ACCESS MANAGEMENT

The concepts behind Identity, Entitlement, and Access Management used in traditional computing require fundamental changes in thinking when implementing a cloud environment, particularly splitting it into three discrete functions, Identity, Entitlement, and Authorization/Access Management (IdEA).

For most organizations, implementing a traditional application means implementing a server, possibly in a **DMZ**[109], and in most cases tied into a Directory Service (**DS**) [110] (such as Microsoft's Active Directory,Novell's eDirectory or Open LDAP) for user authentication. In some cases it means implementing an application or using a web-delivered service using its own stand-alone authentication system, much to the annoyance of the users who then have to remember sets of credentials (or worse, reuse credentials from other, perhaps more trusted, domains).

In contrast, a well implemented cloud service or application-identity should be consumed from a variety of external sources together along with the associated attributes (remembering that an identity applies not only to **Users**[111], but also Devices, **Code**[112], Organizations and Agents which all have identity and attributes). Leveraging all the multiple identities and attributes involved in a transaction enables the cloud system to make better holistic risk-based decisions (defined by the **entitlement process**[113] and implemented by the authorization & access management components) about granular access to the system, processes, and data within the cloud system / application.

This process of using multiple sources of Identity and their related attributes is critical when a cloud application is likely to be Internet-facing, and is also likely to be one of the main hurdles for organizations wanting to use "true" cloud services and instead opt to implement virtualization technologies in their own DMZ connected to their own internal DS.

This **de-perimeterized**[114] approach to identity, entitlement, and access management provides a more flexible and secure approach but also can be implemented equally well inside the corporate boundary (or perimeter).

Overview. The following sections cover the key aspects of Identity, Entitlement, and Access Management in a cloud environment:

- Introduction to Identity in a cloud environment
- Identity architecture for the Cloud
- Identity Federation

[109] **DMZ** - DeMilitarized Zone

[110] **DS** or "Directory Service" is used through this section as an abbreviation for a generic corporate directory service, used for username and password login.

[111] Typically humans; for a wider definition and expansion refer to www.opengroup.org/jericho/Jericho%20Forum%20Identity%20Commandments%20v1.0.pdf

[112] Code includes all forms of code, up to including applications and self-protecting data.

[113] "Entitlement" is the process of mapping privileges (e.g., access to an application or its data) to identities and the related attributes.

[114] De-perimterization is a term coined by the Jericho Forum® (www.jerichoforum.org)

- Provisioning and governance of Identity and Attributes
- Authorization and Access Management
- Architectures for interfacing to Identity and Attribute providers
- Level of trust with Identity and Attributes
- Provisioning of accounts on cloud systems
- Application Design for Identity
- Identity and Data Protection

12.1 Terminology Used in this Document

The language used around identity is confused, with some terms having diametrically opposite means to different people. To avoid confusion while reading this domain, some of the identity terms used within this domain are defined below:

- **Identity**. The means by which an *Entity* can consistently and comprehensively be identified as unique.
- **Identifier**. The means by which an *Identity* can cryptographically asserted, usually using public-key technology.
- **Entity**. Discrete types that will have *Identity*; these are to Users, Devices, Code, Organizations and Agents.
- **Entitlement**. The process of mapping privileges (e.g., access to an application or its data) to *Identities* and the related *Attributes*.
- **Reduced Sign-on (RSO)**. The use of an account and/or credential synchronization tool to minimize the number of credentials (usually username and password) a user has to remember; most of these solutions result in some form of security compromise.
- **Single Sign On (SSO)**. The ability to pass *Identity* and *Attributes* to a cloud service, securely, using secure standards such as **SAML**[115] and **OAuth**[116].
- **Federation**. The connection of one *Identity* repository to another.
- **Persona**. *Identity* plus the particular *Attributes* that provide context to the environment the *Entity* is operating within. A *Persona* may be an aggregation of an individual *Identity* together with an Organizational *Identity* and Organization *Attributes* (e.g. a corporate *Persona*, Fred Smith as CEO of ACME Corp., or a Personal Computer belonging to ACME Corp.).
- **Attributes**. Facets of an *Identity*

[115] **SAML**- Security Assertion Markup Language, an XML-based OASIS open standard for exchanging authentication and authorization data between security domains

[116] **OAuth**-Open Authorization, an open standard for authorization, allowing users to share their private resources with tokens instead of credentials

12.2 Introduction to Identity in a Cloud Environment

An identity eco-system faces scaling problems (think of the move from a small village where everyone knows everyone else, to a large town or city). As the industry expands identity systems from single computers into global enterprises and then into cloud deployment models, the ability to identify all the entities involved in a transaction become significantly more difficult.

However, with cloud, the use of *Identity* for all *Entities* in the transaction value-chain, and the move to risk-based decisions, cannot only mitigate the risk but also potentially improve security.

The following key points need to be considered when implementing a cloud based solution that needs to use identity information:

- The strength with which an *Identity* can be asserted will feed into the risk calculation when interacting with that *Identity* (examples include Anonymous; self-assert; validated by a known reputable organization with a strong assertion of organizational *Identity*).
- The *Attributes* of a *Persona*, like *Identity*, will have a strength with which an *Attribute* can be asserted that feed into the risk calculation when interacting with that *Persona*. Assertion strength ranges from self-asserted to validated by a known reputable organization (with a strong assertion of organizational *Identity*).
- *Identity* and *Attributes* will need to be consumed from multiple sources, thus cloud solutions / architectures will need the ability to consume multiple disparate sources of *Identity* and *Attributes*.
- There will be instances when a transient *Identity* is sufficient (enough information about an *Entity* to deem them unique).
- There will be instances where pseudo-anonymity is desirable (such as voting).

12.3 Identity Architecture for the Cloud

The move from a traditional architecture of a perimeterized organization with traditional server based applications in internal computer centers affords little flexibility to an organization. The move to cloud-based architectures allows greater flexibility, whether deployed internally within the organizational boundaries (a private cloud) or external public clouds (SaaS, PaaS or IaaS).

The table on the following page shows how identity needs to vary between traditional implementation and cloud implementation, dependent on the type of cloud architecture implemented.

Table 1—Identity Architecture Assertions

ARCHITECTURE TYPE	TRADITIONAL ARCHITECTURE	CLOUD ARCHITECTURE
Internal / Perimeterized	Connected to the internal DS Identities must be maintained within the DS to be used by the application, potentially using reduced sign-on solutions.	Ability to accept multiple sources of *Identity* and *Attributes*.
Internal / De-perimeterized	Need to tightly control and connect to organizational services using VPN tunnels at back end. Not a recommended architecture.	Use assertions to provide *Identity* and *Attributes* to access cloud services.
External / Perimeterized	External hosting means extending perimeter to the provider of the server. *Identity* is extended into an environment the consumer does not manage, often putting a replica of the DS into that environment for performance.	Use assertions to provide *Identity* and *Attributes* to access cloud services.
External / De-perimeterized	External hosting means extending internal *Identity* into a foreign environment, but a back-end leaded line or VPN. *Identity* is extended into an environment the consumer does not own or manage, often replicating DS into that environment for performance	Use assertions to provide *Identity* and *Attributes* to access cloud services.

Whereas in a traditional "**IAM**"[117] architecture, often all the components are stand-alone as part of a single server, a cloud architecture is potentially more complex taking *Identity* and *Attributes* from a number of sources and making authorization / access management decisions via a set of Entitlement Rules defined by the Entitlement Process.

[117] **IAM**-*Identity* and Access Management

In Figure 1, *Identity* and *Attributes* are sourced from (potentially) multiple sources and feed into an authorization/access management layer that translates the entitlement rules into access.

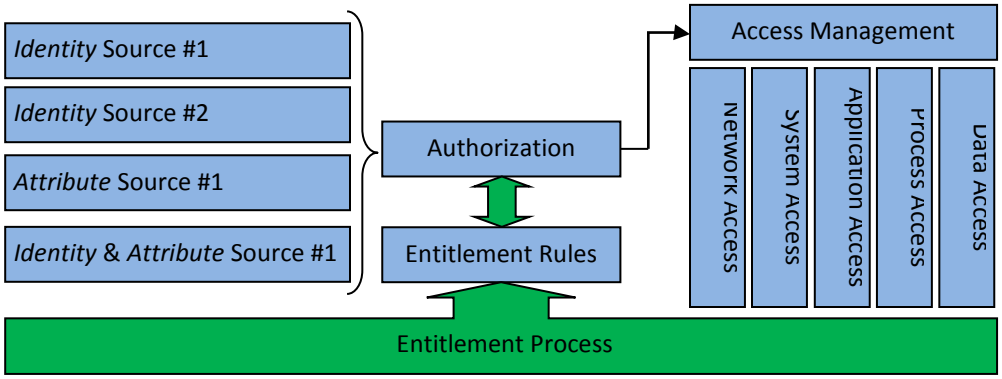

Figure 1: Generic Identity, Entitlement & Access Management System

Access Management should (depending on the business / security requirements, and the type of cloud model, IaaS, PaaS or SaaS being deployed) govern access to the;

- **Network layer.** Without meeting the entitlement rules it may not even be possible to "see" (i.e. Ping or route) to the cloud system. The entitlement rules may also direct access to particular interfaces.
- **System layer.** The entitlement rules may define the protocols that are permitted to access and modify systems, such as terminal server vs. web.
- **Application layer.** The entitlement rules may map *Identity* and/or *Attributes* to functionality provided by a specific application, such a being presented with a reduced set of menus or options.
- **Process layer.** The entitlement rules can be used to define the processes (or functions) that can be run within an application. Entitlement may also define that enhanced functions (such as transferring money out of the eco-system) need additional verification (which may be obtained directly or derived in the background).
- **Data layer.** The entitlement rules may limit access to areas of the data and file structure or even individual files or fields within files (e.g., in a database). At a more advanced level, entitlement could be used to auto-redact documents, such that two users accessing identical documents would view different content (e.g., constructing a specific dynamic view of a database table).

The entitlement process starts with the customer to turn business requirements and security requirements into a set of entitlement rules. This process will define the identities and *Attribute*s required to be able to evaluate the rules. These rules in turn drive the authorization/ access system.

12.4 Identity Federation

Conceptually speaking, federation is the interconnection of disparate Directories Services. Some organizations opt for a federation gateway, (a "Bridge" or "Federation Hub") in order to externalize their federation implementation, where the

federation and the rules by which *Identity* is managed within that "bridge" is governed by a set of rules, usually a legal contract, thus allowing other partners in this bridge a defined level of trust in identities not directly issued by themselves.

Technologically speaking, federation is the use of SAML to offer portability to disparate and independent security domains with some organizations extending their DS environment via a gateway product that will handle SAML assertions. Other organizations will consume native SAML assertions from an identity service.

In both these types of federation architectures, it is essential to understand the provenance of the *Identity* and *Attributes* that are being asserted.

Federation standards are used widely for SaaS deployment models for both identity federation and access control. There are no similar standards for PaaS or IaaS deployment models. Cloud Consumers leveraging IaaS deployment models should take into consideration how they manage the lifecycle of identities (shared accounts, named accounts, privileged accounts etc.). Enterprises that leverage the Privileged *Identity* Management (PIM) tools for Super User Management (SUPM) and Shared Account Password Management(SAPM) should investigate extending these tools to support cloud deployments. Enterprise or Cloud Consumers must have a well-defined policy for HPA (Highly Privileged Access).

12.5 Provisioning and Governance of Identity and Attributes

When talking about provisioning, typically we think about user provisioning, but to make rich, risk-based decisions, the cloud system / application needs *Identity* and *Attributes* from all entities involved in the transaction and potentially other *Attributes* from other systems / processes.

Some examples of *Identity* and *Attributes* are as follows (not an exhaustive list):

- User Assertions: User *Identifier* (The public part of a Public/Private key pair)
- User Name (User Name should be just another *Attribute* of *Identity*)
- Credential strength/trust
- Location Assertions; IP-Address, Geo-location, GPS, Cellular Service Location
- Organization *Identity* (*Identifier* – crypto) and Organization Assertions
- Device *Identity* (*Identifier* – crypto) and Device Assertions; Functionality Required, Functionality Offered, Sandbox capability, Secure container, Cleanliness of device
- Code *Identity* (*Identifier* – crypto) and Code Assertions
- Training record / compliance, etc.

The master source of *Identity* and the *Attributes* of an *Identity* (which may be from a different source) need to be identified in the design of the entitlement process.

As a rule, the cloud service or application itself should avoid being the master source for *Identity* (exceptions may be a cloud based HR service, or a cloud *Identity*-as-a-Service offering). However, during the transition to cloud services (not best practice) the cloud service / application may need to hold identities or operate a mixed-mode model.

All *Attributes* should be linked to an *Identity*, as without the associated *Identifier* and level of trust with that *Identifier* the *Attributes* have no provenance. While this may at first sight be counterintuitive, the strength in the entitlement process lies in defining those *Attributes* necessary to make the rules work the way the business requires them to and then identifying the authoritative source (or as close as possible) to provide those *Attributes* (with the related *Entity Identifier*). Examples would include:

- Security threat level: Organizational, Government, or Outsourced provider *Identity*
- Approvals or prior decisions made by other Entities: *Entity Identity*
- QoS or throttling policies related to a protected target resource; System *Identity*

12.6 The Entitlement Process

The entitlement process starts with the customer to turn business requirements and security requirements into a set of rules that will govern authorization and access to the various aspects of the cloud system. This process will then define the identities and *Attribute*s that are required to properly evaluate the entitled rules. The entitlement process and the derived rules should not only drive the authorization and access management of a cloud system, they can also specify a degree of negotiation/entitlement at all layers of the cloud infrastructure, e.g., to allow protocols and interfaces at the network and/or system layer.

The entitlement process should be embedded into any business requirements document and also the technical requirements document; it should also feature as an integral part of the cloud vendor's provisioning / "customer on-boarding" process.

The entitlement process does not stop once the cloud service is up and running, but the entitlement rules and the subsequent rules that drive authorization and access must be the subject of regular review. The entitlement process must then be audited by the business "system-owner" against the business requirement. Any audit must include the threat and risk assessment and any regulatory requirements.

Current solutions include automated approaches to turn high-level security policies into (low-level) technical access rules, including:

- **Model-driven security**[118], a tool-supported process of modeling security requirements at a high level of abstraction and using other information sources available about the system (produced by other stakeholders)
- Clustering technical access rules into similar groups to reduce the complexity
- Visual attempts to make technical policies easier to understand

The entitlement process should define those *Entities*, *Identities*, and *Attributes* that are required to make meaningful authorization and access decisions. It should also define those *Attributes* that are fixed within the process, or those that

[118] www.modeldrivensecurity.org

have a temporal (change over time) aspect to them, and therefore either the time interval at which they must be revalidated, or the trigger within the process, will force revalidation.

Where *Identity* and *Attributes* that need to be sourced from outside the business's control are defined in the entitlement process, the Organizational *Identity* (*Identifier*) of that provider (*Entity*) must be on-boarded as well, and thus (at some point in time) off-boarded.

Typically the entitlement rules are interpreted in one of three places:

1. Using a central/external Policy Enforcement point / Policy Server / Policy-as-a-Service
2. Embedded as part of the Cloud application
3. Using an Identity-aaS (IDaaS)

12.7 Authorization and Access Management

Authorization and Access Management is the process by which the entitlement rules are translated (via the Authorization layer) into Access Management rules.

In most cloud based systems, the Authorization layer is likely to be a "Policy Decision Point" (**PDP)**[119] or the point that evaluates and issues authorization decisions, and the Access Management layer, the "Policy Enforcement Point" (**PEP)**[120]**,** the point that enforces the PDP's decision.

The PDP and PEP will be part of an authorization eco-system that uses **XACML**[121] (eXtensible Access Control Markup Language) as a declarative access control policy language implemented in XML.

A PEP could be as simple as an IF (conditional) statement in the cloud application or as advanced as an agent running on an application server or a filter in an XML-gateway that intercepts access requests, gathers necessary data (*Attributes*) to be able to evaluate the Entitlement Rules, and then makes and implements these decisions.

This is not to mandate the use of XACML, PDP's, and PEP's in a cloud environment, as the functionality could potentially be implemented in other ways (probably in a closed or proprietary eco-system).

PDP's can be implemented outside of the cloud environment, possibly within the customer's environment. This can potentially have a number of advantages such as interfacing to an internal DS and/or the ability to integrate logs about the decision made directly into an internal logging system.

12.8 Architectures for Interfacing to Identity and Attribute Providers

There are three basic architectures for interfacing to *Identity* and *Attribute* providers:

[119] **PDP** - Policy Decision Point
[120] **PEP** - Policy Enforcement Point
[121] **XACML** - eXtensible Access Control Markup Language

1. A "hub-and-spoke" model where *Identity* and *Attributes* are centrally managed (coordinated) by the hub, which then interacts with the cloud service(s) or cloud application(s)
2. The free-form model where the cloud service and/or application can be configured to accept *Identities* and *Attributes* from multiple sources
3. The hybrid solution, where the components are distributed, potentially using other cloud services.

Each model has its merits, and the choice will be based on the number of factors, including:

- Where the customers for the service have their identity
- The capability of the cloud service chosen
- The capability of the enterprise to provide assertion-based *Identity* and *Attributes*.

12.8.1 Hub and Spoke Model

The "hub and spoke" approach typically allows the cloud service to interface directly with the organization for its *Identity* and *Attribute* information, ideally in the form of standards-based assertion protocols, such as OAuth & SAML.

The organization's internal systems are responsible for keeping track of users, other entities and the *Attributes*. This is most like a traditional IAM system, and thus probably the easiest to transition to for cloud solutions being implemented by organizations, as most DS or LDAP systems can have a SAML capability "bolted-on".

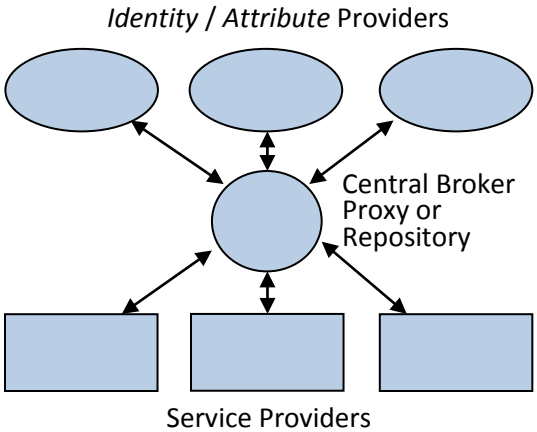

Figure 2— "Hub & Spoke" Model

It is likely in this model that the entitlement process might also be handled within the organization through the use of a Policy Enforcement Point and Policy Server and communicated via XACML (though XACML isn't that widely used for this yet). Figure 2 illustrates the hub-and-spoke approach.

One benefit of this approach is that maintaining a Policy Enforcement Point within the organization allows the integration of audit logs to be maintained within the organization and even correlated with other disparate audit trails (outside of the cloud environment, or from other cloud environments) to get the complete picture required. Examples of this include Segregation of Duties analysis and satisfaction of regulatory requirements.

The hub-and-spoke approach is likely to be used when a high degree of control is required over all "users" with a central enrollment process. This is more likely in organizations that are subject to heavy regulation. The hub-and-spoke should also lessen the dependency on the *Identity*/*Attribute* providers, as *Attributes* are often stored (duplicated) within the central hub.

This is also the model used when an organization subscribes to a Bridge or "Identity Hub."

12.8.2 Free Form Model

In the "free-form" model, the cloud service / application is responsible for maintaining the sources of *Identity* and *Attributes*. This solution is more suited for a public facing solution or a solution with a large number of disparate partners.

The free form approach has the advantage that it is easier to setup, at least for current federation protocols (such as SAML) but relies on a good implementation of the entitlement model to allow it to scale to a large amount of "users."

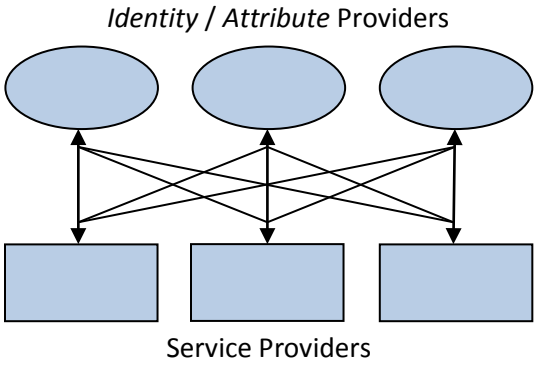

Figure 3—"Free Form" Model

One approach is to setup a point-to-point federated trust relationship (using protocols such as SAML and OAuth) between the service and *Attribute*/*Identity* providers, but this approach needs an efficient process to on-board and off-board those providers.

The free-form model provides challenges to provisioning "users", as the environment of new entities connecting is likely to be more ad-hoc. Again, careful design of the Entitlement Process will help to alleviate this problem. Figure 3 above illustrates the point-to-point approach.

12.8.3 Hybrid Model

The Hybrid model is (by definition) a mix of both the hub & spoke and free-form model. For example, the entitlement rules may be held inside the organization and pushed to a PDP, which in itself is a cloud service, and then those decisions are delivered to multiple disparate PEP's that are part of separate cloud services. In more large-scale deployments, there can be several federated policy servers that service many different PDP/PEP's. The hybrid model will also be found in organizations that mix traditional and/or legacy computing with a (public and/or private) cloud environment.

The hybrid model may offer an efficient use of distributed resources, but it risks becoming complex with the attendant scope for security loopholes to creep in. It also makes long-term maintenance more problematic (the reasoning behind simple rules is easy to understand when all who implemented them are long gone).

The hybrid model will also have issues of logging decisions and actions taken with the potential need to bring all logs back into a single place in a common format to allow a holistic view to be taken.

The potential complexity of the hybrid model stresses the need to be able to use visualization tools to develop, maintain, and audit the translation of the Entitlement Rules into actual access control.

12.9 Level of Trust with Identity and Attributes

Identity and *Attributes* come with many levels of trust, both in the various identities being used in a transaction and with the *Attributes* attached to those identities. Traditionally this lack of trust has led to organizations having to maintain identities for anyone who needs access to their systems, which can be (in some cases) for hundreds of thousands of people who they do not employ or directly manage.

Some organizations (Military/Aerospace, Pharmaceutical, etc.) that need to collaborate with a pre-agreed level of trust use a "Bridge" or "Federation Hub" (see section 12.4), where trusted identities also have trusted *Attributes* associated with them.

During the entitlement process it is essential to understand not only the *Attributes* required, but also the source of those *Attributes*, the organization that will provide them, and the strength (level of trust) with which they can be asserted.

To accept *Attributes* from an external organization with any defined level of trust will require an on-boarding process for that organization, and the *Identity* (*Identifier*) of the organization that will be asserting those *Attributes*.

As a rule, the aim should be to source *Identity* and *Attributes* from the master / authoritative source of those *Attributes* with all *Attributes* having a known *Identity* asserting them, as the level of trust that can be placed in the *Attribute* cannot exceed the level of trust that can be placed in the *Identity* asserting the *Attribute*.

Where *Attributes* are being uniquely generated within the cloud system itself, then a governance process must be in place to ensure that all *Attributes* are accurate and have appropriate lifecycle management.

12.10 Provisioning of Accounts on Cloud Systems

Where it is necessary to provision an "account" on cloud systems (typically for a user, but it could be for any *Entity* type) there are challenges when provisioning (and de-provisioning) these accounts, as the normal "push" model used within organizations is generally not a viable solution for a cloud implementation.

At the time of writing, there are no widely used or de-facto provisioning standards; **SPML**[122] (Service Provisioning Markup Language) has not been widely adopted by the cloud providers, and **SCIM** [123](Simple Cloud *Identity* Management) is only starting to emerge as a potential standard.

The key to provisioning entities on a cloud system is to understand the complete lifecycle management of an account, from creation, management, and eventually de-commissioning (including deletion and/or archiving) across all the systems that both provide and consume the *Identity* and *Attributes*.

There are some key issues that need to be addressed with sources of *Identity* and *Attributes* when it comes to provisioning:

- The link to Human Resources (or the authoritative source of person-user information) is problematic as HR is often only the master source for staff on regular payroll.
- There are usually no authoritative information sources for partner information and their devices.
- The ability to provision other entities (particularly organizations and devices) does not exist in most organizations.
- Public *Identity* services generally only provide self-asserted *Identity* and only about people; it does not extend to the other *Entity* types.

[122] **SPML** - Service Provisioning Markup Language
[123] **SCIM** - Simple Cloud *Identity* Management

- De-provisioning needs to extend to all entities, thus most organizations do not have the ability to off-board another organization when the contract finishes or revoke code from operating on systems when it is found to be faulty or obsolete.

These issues and the immaturity of provisioning standards stress the need for good planning and a holistic approach to how *Identity*, *Attributes*, accounts, and lifecycle management of all *Entity*-types will operate in the cloud eco-system being developed.

12.11 Identity-as-a-Service

Cloud Identity as a Service **(IDaaS)**[124] is a broad term that covers the management of any part of the *Identity*, Entitlement, and Authorization/Access Management in the cloud service.

This encompasses service for software, platform, or infrastructure services, and for both public and private clouds. Hybrid solutions are also possible, whereby identities can still be managed internally within an organization, while other components such as authentication are externalized through a Service Oriented Architecture (**SOA**)[125]. This effectively creates a Platform as a Service (PaaS) layer to facilitate a cloud-based IAM solution.

For more information refer to the section covering Identity-as-a-Service in Domain 14 – "Security-as-a-Service".

12.12 Compliance & Audit

The outcome of the entitlement rules may well need to be logged together with the decisions made by the entitlement rules / authorization process for compliance or security reasons. Compliance and audit is integrally tied to *Identity*. Without proper *Identity* management, there is no way to assure regulatory compliance. Auditing also requires proper *Identity* management, and the use of log files is of little value without a working *Identity* system.

12.13 Application Design for Identity

This section applies just to application design as it applies to Identity and should be read in conjunction with Domain 10 – Application Security.

Designing cloud based systems or applications necessitates a change in mindset when it comes to *Identity* as *Identity* and *Attribute* information will be consumed by the cloud service or application, needing to be held for at least the duration of the transaction, and probably some facets maintained longer, but because the cloud environment may likely not be a part of an organization's physical or logical jurisdiction, and may even be in a different legal jurisdiction, the service and application design may need to be substantially different from the design practices used in traditional client server in a DMZ owned and managed by the organization.

[124] **IDaaS** - Cloud *Identity* as a Service
[125] **SOA** - Service Oriented Architecture

The design goal should be to minimize the need for *Identity* and *Attributes*. When possible, start from the principle that identification is not required while understanding the threshold where there will be a need to switch from basic "on-the-fly" account provisioning to an "identified" user account. Examples include:

- Unique sessions can be established using other *Attributes*, e.g., the IP address of the connecting device (understanding that IP addresses can be spoofed) or a unique session cookie.
- In many cases *Attribute*-based entitlement alone will be adequate with no need for user information or an actual *Identity*; don't assume a *Persona* is needed to tie to a session or even an account.
- When encountering a new *Entity* for the first time (say authenticating with a SAML assertion) then create a basic account on-the-fly. [Note that this approach requires thought about de-provisioning.]
- Use *Attribute* derivation whenever possible, (e.g. don't ask for date of birth, instead query "if over 18" [if DoB > (today – 18 years)].

When generating any unique accounts, decide whether the system will consume an external unique *Identifier* from the *Entity* or whether the system will need to generate its own unique *Identifier* (such as a customer reference number).

There must be careful thought put into cloud systems that maintain user accounts. There must be careful design thought put into how the cloud user accounts will be synchronized with existing user accounts in other systems (either internal or other cloud systems), particularly around the integration with a "joiners and leavers" process, and the changes in access required when people move internally. Designing a cloud system to scale (think of 100,000 users with an unconnected username and/or unsynchronized password) requires the need to avoid forcing a common help desk process, including manual or semi-automated synchronization scripts, password strength validation processes, password reset processes, password resets after a compromise, etc. all due to a lack of initial design thought about consuming external identities.

Avoid trying to extend an internal DS into the cloud service and/or replicating the organization's DS over the Internet (generally very insecure) or via a back-channel (leased line or VPN) as this exposes an organization's entire DS into an environment the organization does not control. Also be wary of the promise of RSO (reduced-sign-on) products as RSO generally works by compromising on-log-in security internally, more so when trying to extend RSO to a cloud environment.

As a rule, cloud services, and thus cloud applications, should accept the standard SSO federation formats such as SAML and OAuth (or even the less widely accepted WS-Federation).

When designing an application to consume *Identity* and *Attributes*, remember that *Identity* encompasses all Entities, and that the application security should, where possible, be part of a holistic approach that includes all layers; the Network layer; the System layer; the Application layer; the Process layer; and the Data layer (as detailed in section 12.3). An application could (for example) offer two methods of connecting a full, rich connection using Web/AJAX/Java or an Citrix style "screen-scrape" connection with the type of connection permitted determined by the Entitlement Rules (defined in the Entitlement process).

12.14 Identity and Data Protection

Holding aspects of an *Identity* that comprises Personal Identifiable Information (**PII**)[126], and particularly information classified as Sensitive Personal Information (**SPI**)[127], is an issue for all organizations. Cloud services managed or located outside of the organization will need specialist advice to ensure all applicable laws and regulations are being adhered to.

When considering which laws or jurisdictions might apply, the following (non-exhaustive) list should be considered:

- All the countries of the data subjects
- The country in which the organization operates
- Countries in which the organization has legal entities
- Countries in which the organization lists on the stock exchange or issues shares
- The country or countries where the cloud services are physically located
- The relevant legislation, regulations, and also pseudo-regulation (such as PCI-DSS)

12.15 Consumerization and the Identity Challenge

Interacting with consumers and/or consumer devices brings a number of challenges and opportunities in cloud-based services and applications. The ability of the consumer and consumer devices to interface directly to an Internet-facing cloud service strips away a layer of network complexity but introduces a series of security challenges which can potentially be mitigated using *Identity*.

However, in the consumer space, standards for device and user *Identity* are fragmented and (by definition) will rarely have the same level of conformance and standardization that can be achieved in a corporate environment.

Unfortunately, most consumer devices and consumers themselves have no easy or standard way to enroll themselves or their devices into an authentication system providing strong authentication, and thus, authorization without strong *Identity* is difficult. Even when users have an existing strong authentication method (for example with their bank) for one account, this can almost never be reused with another account/provider. This has resulted in a situation where *Identity* for the consumer has already passed the limits of scalability. Over 61 percent of people use the same password whenever they can[128]; this results in every additional registration or authentication causing the loss of potential customers.

Solving this problem with seamless access to applications will facilitate business, and clear separation between *Identity* and authorization will facilitate additional uses, for example allowing one individual to delegate use of their *Persona* linked to a specific credit card on behalf of another's transactions.

12.16 Identity Service Providers

Consuming information about *Identity* from an external service brings its own issues. Levels of trust in the providing organization and validation of *Attributes* are just two examples. Most current proposals or actual offerings for

[126] **PII** - Personal Identifiable Information
[127] **SPI** - Sensitive Personal Information
[128] http://www.guardian.co.uk/technology/2008/jan/17/security.banks

comprehensive and consistent *Identity* frameworks are extrapolations of single or group players' needs by those with little or no understanding of other communities' needs.

Nearly all open/public consumable *Identity* services deal only with user verification. Those that offer personal information (*Attributes* as well as *Identity*) do so using *Attributes* that are either self-asserted or not from authoritative sources.

Examples of sources of *Identity* and *Attributes* are as follows:

- National Government
  - United States, NSTIC – strategy & aspiration only
  - German "EID card," Austrian "Citizen Card," Estonian "ID Card," Finland "Citizen Certificate," Hong Kong "Smart ID Card," Malaysian "MyCad"
- Public – integration via API's
  - Facebook
  - Amazon
  - Google
  - Microsoft Passport (Windows Live ID)
  - OpenID providers (Various)
  - Twitter
- Bridges[129] or "Hubs"
  - Research / Education Bridge (REBCA[130]), serving the US higher education sector
  - Federal PKI Architecture (Federal Bridge) serving all US federal agencies.
  - CertiPath/Transglobal Secure Collaboration Program[131], serving the aerospace and defense industries
  - SAFE-BioPharma Association[132], serving the biopharmaceutical and healthcare industry
- *Identity* Service offerings
  - Check / validate my postal code and address (various)
  - Experian / Equifax
  - 3D card verification (Visa/MasterCard)

[129] www.the4bf.com
[130] www.hebca.org
[131] www.certipath.com / www.tscp.org/
[132] www.safe-biopharma.org/

- eBay / PayPal / X.Commerce

12.17 Recommendations

12.17.1 Federation Recommendations

- Consumers, Implementers, and Providers should agree on the context and definition of "federation" being used.
- Implementers should understand what trust relationship and transitive trust exist and the need for bi-direction trust relationships.
- Implementers should, where possible, use federation based on open standards such as SAML and OAuth.
- If using a "Bridge" or "Federation Hub", then Implementers should understand the nature and relationship of the trusts that exist between different members of the club. Understand what it could mean to your entitlement rules if there is another member signed up to the cloud or federating to another bridge.
- Implementers should understand that public *Identity* providers such as Facebook, Yahoo, or Google provide a source of (low grade, typically self-asserted) *Identity* with no guarantees that they will not federate to other providers in the future.
- Implementers should deprecate examples of bad solution design solutions to get *Identity* from a DS linked into the access management system of a cloud solution. Such examples include in-band VPN's and out-of-band leased-lines.

12.17.2 Provisioning and Governance Recommendations

- All *Attributes* should be sourced from as close to the authoritative / master source as possible.
- As a rule, the cloud service or application itself should avoid being the master source for *Identity*.
- The cloud service or application itself should only be the master source for *Attributes* it directly controls.
- All *Attributes* consumed should have a known level of trust.
- All *Attributes* consumed should be linked to an *Identity*.
- The *Identifier* of a defined *Entity* should sign all *Attributes* consumed.
- Each *Attribute* should have a lifecycle that is fit-for-purpose.
- Each *Identity* (and related *Identifier*) should have a lifecycle that is fit-for-purpose.

12.17.3 Entitlement Recommendations

- All parties in the entitlement (definition) process should be clearly identified.

- There should be clear designated responsibilities for entitlement rule approval and sign-off.
- A process to manage changes to the Entitlement Rules should be clearly defined.
- A frequency (or trigger) for auditing the Entitlement Rules should be clearly defined.
- The entitlement process should focus on producing Entitlement Rules that are simple and minimal and designed using the principle of least privilege.
- The entitlement process should focus on producing entitlement rules that minimize the exposure of *Identity* or avoid needing to consume *Identity* altogether.
- *Attributes* that are temporal (such as geolocation) need real-time *Attribute* checking through a lifetime of transaction to revalidate the entitlement rules.
- Entitlement rules should be triggered by a process (or attempt to initiate a process, such as money transfer out of environment). In some environments, best practice would be for the entitlement rules to disable such functions. In others, best practice would be to require additional *Identity* or *Attributes* at the point of execution to ensure the *Entity* is entitled to perform the process.
- Implementers should ensure bi-directional trust to ensure the optimal secure relationship for the transaction. The entitlement process should define this.
- The design of the entitlement rules should include delegation[133] of access by a secondary *Entity* to some, but not necessarily all, information the primary *Entity* can access.
- The design of entitlement should include the seizing of access (including legal seizure), although the designer of the entitlement rules will need to take into account the jurisdiction of the system, organization, and entities involved. Legal advice must be taken prior to implementing any seizure of access.
- Where practical, management interfaces, tools, or other visualization technologies should be used to help in management of Entitlement and to help ensure the interpretation meets the business and/or regulatory (e.g. SOX Segregation-of-Duties) requirements.

12.17.4 Authorization and Access Recommendations

- Implementers should ensure services have an import and/or export function into standards such as OASIS XACML.
- When using a PDP in a cloud environment, implementers should understand how authorization decision logging would be extracted and/or integrated into an organization logging for a holistic picture of access.
- Implementers should ensure existing (legacy) services can interface with PEP/PDP's.
- Implementers should ensure any PDP is adequate to properly translate the entitlement rules defined in the entitlement process.

[133] Delegation is newly supported in XACML 3.0

- Implementers should consider the use of "policy-as-a-service" as the policy server if there is a need for a central policy server (for example for cloud mashups).

12.17.5 Architecture Recommendations

- Implementers should ensure any cloud provider offers authorization management PEPs/PDP's that can be configured with entitlement rules.
- Implementers should ensure that all components of the *Identity*, Entitlement, and Authorization / Access Management (IdEA) work correctly together.
- Implementers should ensure that Policy Decision/Enforcement Points (PEP's/PDP's) use standard protocols (e.g. XACML) and avoid (or depreciate) proprietary protocols (e.g. direct web service or other middleware calls).
- Implementers should ensure any strong authentication service is OAuth compliant. With an OAuth-compliant solution, organizations can avoid becoming locked into one vendor's authentication credentials.
- Cloud services and applications should support the capability to consume authentication from authoritative sources using SAML.
- Implementers should ensure services have an import and/or export function into standards such as OASIS XACML.
- Implementers should ensure services can interface with PEP/PDPs installed in the cloud infrastructure and with Policy Monitoring Points for incident monitoring/auditing.
- Implementers should ensure that logging of authorization decision and access actually granted can be logged in a common format using standard secure protocols.

12.17.6 Entitlement Recommendations

- Implementers should ensure that each *Identity* and *Attribute* defined in the entitlement process matches the level of trust that is needed (or is acceptable) in both the *Identity*/*Attribute* itself and also matches the source that provides it.
- Implementers should ensure all sources of *Identity* / *Attributes* provide organizational *Identity*.
- Implementers should ensure that *Attributes* are validated at master / source whenever possible, or as close as possible.
- Implementers should ensure *Attribute* use correctly leads to the right conclusion. (Your context may be different to the originator of the *Attribute*)
- Implementers should ensure that the *Identity* / *Attribute* source has both the standards of data quality and a governance mechanism that meets your needs.

- Consumers should be aware that reputational trust can be an important source of trust. Through the entitlement definition, consumers should be aware of small value transitions, leading to an increase in transactional trust, which may be defrauded on a subsequent large transaction.

12.17.7 Provisioning Recommendations

- Providers should understand whether SPML or SCIM could be a viable option for provisioning.
- Implementers should follow the rule of least privilege when provisioning an account. In the case of entities such as computing devices, a link to organizational asset registries is desirable.
- Most systems and applications have a one-to-one relationship between the user and access and no concept of delegation.
- Implementers should ensure that provisioning and de-provisioning are not limited to user identities. Architectures must include authorization for all *Entity* types.
- Implementers should ensure provisioning and de-provisioning are done in real time.
- Providers should ensure the maintenance of both *Identity* and *Attributes* are critical if entitlement is to be accurate.

12.17.8 Identity Compliance & Audit Recommendations

- Implementers should ensure that the applicable logs from the entitlement rules / authorization process are capable of being made available.
- Implementers should ensure where logs need to be integrated into a wider (possibly remote) system (e.g. for wider fraud detection or segregation of duties analysis) to ensure the availability, timeliness, format, and transmission security of the logs is adequate.
- When logging access decisions, implementers should group the *Attributes* together with the entitlement logic used at the time of the decision, and the outcome should be recorded.
- All cloud participants should remember that *Attributes* with a temporal component might need to be revalidated, and hence re-logged, during the lifetime of the session.
- When logging PII or SPI then, whenever possible, implementers should use *Attribute* derivation to minimize the exposure of PII or SPI in the logs.
- Consumers should be aware that logs containing PII or SPI will be subject to data protection legislation.

12.17.9 Application Design Recommendations

- Implementers should use ITU X.805 / 3-layer definition of User, System, and Management layers to ensure segregation.

- Implementers should minimize the need for *Identity* and *Attributes* in an application design.
- When possible, design cloud systems to consume *Identity* and *Attributes* from external sources.
- Implementers should ensure the cloud system supports standard SSO federation formats such as SAML and OAuth.
- Implementers should take a holistic approach to security, using *Identity* and *Attributes* across all the layers of the system.
- Implementers should remember that mutual authentication is critical at all levels, and even more important in cloud environments, just as the cloud environment needs entities and other systems to authenticate who they are, so the cloud system needs to be able to authenticate in return.

12.17.10 Data Protection Recommendations

- Implementers should minimize the use and storage of PII or SPI. This should be done in the design phase of the entitlement process to ensure only Identities and *Attributes* essential to the process are used.
- The implementer should consider the following technologies to minimize exposure of PII or SPI:
  - Encryption
  - Tokenization
  - Homomorphic Encryption[134]

  Refer to Domain 11 "Encryption & Key Management" for more information.
- Implementers should consider using best practice approaches to protecting SPI such as using a dual-key approach, one held by the subject (or keyed against their log-in), and one by the system for use with processing.
- Implementers should understand how administrator access to PII and SPI might be restricted or stopped.
- Implementers should understand how a "**Subject Access Request**[135]" can be dealt with in the legal timeframe mandated especially when the data may be held on a cloud system not owned / managed by the organization that received the request.
- If there is a need to share PII or SPI, consumers should understand how the approval of the subject of the PII/SPI will be obtained.
- Implementers should reduce PII/SPI being stored, especially when not the authoritative source, and only reference those attributed from the authoritative source rather than store (and maintain) them.
- Implementers should understand the processes by which the maintenance of PII/SPI (whether *Identity* or *Attributes*) will be handled in a timely manner.

[134] At the time of release, Homomorphic Encryption is currently in the early stages of product implementation.
[135] A "Subject Access Request" is the legal right in some countries to request any PII or SPI held about yourself

12.17.11 Identity Implementation Recommendations

- Implementers should start from the principle of *Identity* re-use rather than the unique enrollment of new users and/or devices.
- Consumers should understand where existing sources of *Identity* can provide sufficient levels of trust and be re-used.
- Providers should understand what *Attributes* about the user and devices can be asserted to a sufficient level of trust for the transaction being undertaken.
- When appropriate, consumers should allow low risk transactions to take place using low grade level of authentication. Only escalate the *Identity* required when the transaction value / risk increases.
- Providers should provide a critical assessment of the *Identity* and *Attributes* being required during the Entitlement Process when considering consumers and consumer devices.
- Providers should understand what technologies can be used with consumer devices to increase assurance levels, especially technologies than can be used in the background.
- Consumers should understand where the management of consumer devices will not be performed and the level of assurance this provides; this could range from no assurance to good assurance.
- Consumers should understand where a level of assurance and legal liability resides should an issue arise with a transaction from a consumer device.

12.18 Requirements

- ✓ Implementers must design the common service layers to act independently to enable the removal of application silos without sacrificing existing information security policies and procedures.
- ✓ All cloud participants must respect the integrity of the supply chain and respect existing IAM practices in place. Elements such as privacy, integrity, and audit ability must be respected. *Identity* integrity and audit must be preserved when moving data off-site and/or decoupling the pillars of the solution into web service architecture.

DOMAIN 13 //
VIRTUALIZATION

Virtualization is one of the key elements of Infrastructure as a Service (IaaS) cloud offerings and private clouds, and it is increasingly used in portions of the back-end of Platform as a Service (PaaS) and SaaS (Software as a Service) providers as well. Virtualization is also, naturally, a key technology for virtual desktops, which are delivered from private or public clouds.

The benefits of virtualization are well known, including multi-tenancy, better server utilization, and data center consolidation. Cloud providers can achieve higher density, which translates to better margins, and enterprises can use virtualization to shrink capital expenditure on server hardware as well as increase operational efficiency.

However, virtualization brings with it all the security concerns of the operating system running as a guest, together with new security concerns about the hypervisor layer, as well as new virtualization specific threats, inter-VM (Virtual Machine) attacks and blind spots, performance concerns arising from CPU and memory used for security, and operational complexity from "VM sprawl" as a security inhibitor. New problems like instant-on gaps, data comingling, the difficulty of encrypting virtual machine images, and residual data destruction are coming into focus.

Overview. While there are several forms of virtualization, by far the most common is the virtualized operating system, and that is the focus for this domain. This domain covers these virtualization-related security issues:

- Virtual machine guest hardening
- Hypervisor security
- Inter-VM attacks and blind spots
- Performance concerns
- Operational complexity from VM sprawl
- Instant-on gaps
- Virtual machine encryption
- Data comingling
- Virtual machine data destruction
- Virtual machine image tampering
- In-motion virtual machines

Virtualization brings with it all the security concerns of the guest operating system, along with new virtualization-specific threats.

13.1 Hypervisor Architecture Concerns

13.1.1 VM Guest Hardening

Proper hardening and protection of a virtual machine instance, including firewall (inbound/outbound), HIPS, web application protection, antivirus, file integrity monitoring, and log monitoring can be delivered via software in each guest or using an inline virtual machine combined with hypervisor-based **API's**[136].

13.1.2 Hypervisor Security

The hypervisor needs to be locked down and hardened using best practices. The primary concerns for enterprises and virtualization users should be the proper management of configuration and operations as well as physical security of the server hosting the hypervisor.

13.1.3 Inter-VM Attacks and Blind Spots

Virtualization has a large impact on network security. Virtual machines may communicate with each other over a hardware backplane, rather than a network. As a result, standard network-based security controls are blind to this traffic and cannot perform monitoring or in-line blocking. In-line virtual appliances help to solve this problem; another approach to this issue is hardware-assisted virtualization, which requires API-level integration with hypervisors and virtualization management frameworks. Migration of virtual machines is also a concern. An attack scenario could be the migration of a malicious VM in a trusted zone, and with traditional network-based security controls, its misbehavior will not be detected. Installing a full set of security tools on each individual virtual machine is another approach to add a layer of protection.

13.1.4 Performance Concerns

Installing security software designed for physical servers onto a virtualized server can result in severe degradation in performance, as some security tasks like antivirus scanning are CPU-intensive. The shared environment in virtualized servers leads to resource contention. Especially with virtual desktops or high-density environments, security software needs to be virtualization-aware or it needs to perform security functions on a single virtual machine to support other virtual machines.

13.1.5 Operational Complexity from VM Sprawl

The ease with which VM's can be provisioned has led to an increase in the number of requests for VM's in typical enterprises. This creates a larger attack surface and increases the odds of a misconfiguration or operator error opening a security hole. Policy-based management and use of a virtualization management framework is critical.

[136] **API** - Application Program Interface

13.1.6 Instant-On Gaps

The ease with which a virtual machine can be stopped or started, combined with the speed at which threats change, creates a situation where a virtual machine can be securely configured when it is turned off, but by the time it is started again, threats have evolved, leaving the machine vulnerable. Best practices include network-based security and "virtual patching" that inspects traffic for known attacks before it can get to a newly provisioned or newly started VM. It is also possible to enforce **NAC**[137] (Network Access Control)-like capabilities to isolate stale VM's until their rules and pattern files are updated and a scan has been run.

13.1.7 VM Encryption

Virtual machine images are vulnerable to theft or modification when they are dormant or running. The solution to this problem is to encrypt virtual machine images at all times, but there are performance concerns at this time. For high security or regulated environments, the performance cost is worth it. Encryption must be combined with administrative controls, DLP, and audit trails to prevent a snapshot of a running VM from "escaping into the wild," which would give the attacker access to the data in the VM snapshot.

13.1.8 Data Comingling

There is concern that different classes of data (or VM's hosting different classes of data) may be intermixed on the same physical machine. In **PCI**[138] terms, we refer to this as a mixed-mode deployment. We recommend using a combination of VLANs, firewalls, and **IDS/IPS**[139] to ensure VM isolation as a mechanism for supporting mixed mode deployments. We also recommend using data categorization and policy-based management (e.g., DLP) to prevent this. In cloud computing environments, all tenants in the multi-tenant virtual environment could potentially share the lowest common denominator of security.

13.1.9 VM Data Destruction

When a VM is moved from one physical server to another, enterprises need assurances that no bits are left behind on the disk that could be recovered by another user or when the disk is de-provisioned. Zeroing memory/storage or encryption of all data are solutions to this problem. For encryption keys should be stored on a policy-based key-server away from the virtual environment. In addition, if a VM is migrated while it is running, it may be at risk itself during the migration if encryption, or proper wiping, is not used.

13.1.10 VM Image Tampering

Pre-configured virtual appliances and machine images may be misconfigured or may have been tampered with before you start them.

---

[137] **NAC** - Network Access Control
[138] **PCI** - Payment Card Industry
[139] **IDS** - Intrusion Detection Systems; **IPS**- Intrusion Prevention Systems

13.1.11 In-Motion VM

The unique ability to move virtual machines from one physical server to another creates a complexity for audits and security monitoring. In many cases, virtual machines can be relocated to another physical server (regardless of geographic location) without creating an alert or track-able audit trail.

13.2 Recommendations

- Customers should identify which types of virtualization the cloud provider uses, if any.
- Implementers should consider a zoned approach with production environments separate from test/development and highly sensitive data/workloads.
- Implementers should consider performance when testing and installing virtual machine security tools, as performance varies widely. Virtualization-aware server and network security tools are also important to consider.
- Customer should evaluate, negotiate, and refine the licensing agreements with major vendors in virtualized environments.
- Implementers should secure each virtualized OS by using hardening software in each guest instance or use an inline virtual machine combined with hypervisor-based **API's**[140].
- Virtualized operating systems should be augmented by built-in security measures, leveraging third party security technology to provide layered security controls and reduce dependency on the platform provider alone.
- Implementers should ensure that secure by default configurations follow or exceed available industry baselines.
- Implementers should encrypt virtual machine images when not in use.
- Implementers should explore the efficacy and feasibility of segregating VM's and creating security zones by type of usage (e.g., desktop vs. server), production stage (e.g., development, production, and testing), and sensitivity of data on separate physical hardware components such as servers, storage, etc.
- Implementers should make sure that the security vulnerability assessment tools or services cover the virtualization technologies used.
- Implementers should consider implementing data automated discovery and labeling solutions (e.g., DLP) organization-wide to augment the data classification and control between virtual machines and environments.
- Implementers should consider patching virtual machine images at rest or protect newly spun-up virtual machines until they can be patched.
- Implementers should understand which security controls are in place external to the VM's to protect administrative interfaces (web-based, API's, etc.) exposed to the customers.

[140] **API** - Application Program Interface

13.3 Requirements

- ✓ VM-specific security mechanisms embedded in hypervisor API's must be utilized to provide granular monitoring of traffic crossing VM backplanes, which will be opaque to traditional network security controls.
- ✓ Implementers must update the security policy to reflect the new coming security challenges of virtualization.
- ✓ implementers must encrypt data accessed by virtual machines using policy-based key servers that store the keys separately from the virtual machine and the data.
- ✓ Customers must be aware of multi-tenancy situations with your VM's where regulatory concerns may warrant segregation.
- ✓ Users must validate the pedigree and integrity of any VM image or template originating from any third party, or better yet, create your own VM instances.
- ✓ Virtualized operating systems must include firewall (inbound/outbound), Host Intrusion Prevention System(**HIPS**)[141], Network Intrusion Prevention System (**NIPS**)[142], web application protection, antivirus, file integrity monitoring, and log monitoring, etc. Security countermeasures can be delivered via software in each guest virtual instance or by using an inline virtual machine combined with hypervisor-based API's.
- ✓ Providers must clean any backup and failover systems when deleting and wiping the VM images.
- ✓ Providers must have a reporting mechanism in place that provides evidence of isolation and raises alerts if there is a breach of isolation.

[141] **HIPS** - Host Intrusion Prevention System
[142] **NIPS** - Network Intrusion Prevention System

DOMAIN 14 //
SECURITY AS A SERVICE

Cloud Computing represents one of the most significant shifts in information technology the industry has experienced. Reaching the point where computing functions as a utility has great potential, promising expansive innovations. One such innovation is the centralization of security resources. The security industry has recognized the benefits of a standardized security framework for both the providers and consumers. In the context of a cloud service level agreement between providers and consumers, a standardized security framework takes the form of a document that specifies which security services are provided how and where. With the maturation of security offerings based on standard frameworks, cloud consumers have recognized the need to centralize computing resources for providers and consumers. One of the milestones of the maturity of cloud as a platform for business operations is the adoption of Security as a Service (SecaaS) on a global scale and the recognition of how security can be enhanced. The worldwide implementation of security as an outsourced commodity will eventually minimize the disparate variances and security voids.

SecaaS is looking at Enterprise security from the cloud – this is what differentiates it from most of the other work / research on cloud security. Predominantly cloud security discussions have focused on how to migrate to the Cloud and how to ensure Confidentiality, Integrity, Availability and Location are maintained when using the Cloud. SecaaS looks from the other side to secure systems and data in the cloud as well as hybrid and traditional enterprise networks via cloud-based services. These systems may be in the cloud or more traditionally hosted within the customer's premises. An example of this might be the hosted spam and AV filtering.

Overview. This domain will address the following topics:

- The Ubiquity of Security as a Service in the Marketplace
- Concerns when Implementing Security As a Service
- Advantages of Implementing Security As a Service
- The Diversity of Services that can be categorized as Security As A Service

This document corresponds to the Security as a Service publication as well as the CSA Cloud Control Matrix controls.

14.1 Ubiquity of Security as a Service

Customers are both excited and nervous at the prospects of cloud computing. They are excited by the opportunities to reduce capital costs and excited for a chance to divest infrastructure management and focus on core competencies. Most of all, they are excited by the agility offered by the on-demand provisioning of computing resources and the ability to align information technology with business strategies and needs more readily. However, customers are also very concerned about the security risks of cloud computing and the loss of direct control over the security of systems for which they are accountable. Vendors have attempted to satisfy this demand for security by offering security services in a cloud platform, but because these services take many forms and lack transparency regarding deployed security

controls, they have caused market confusion and complicated the selection process. This has led to limited adoption of cloud-based security services thus far. Security as a Service is experiencing an exponential growth with Gartner predicting that cloud-based security service usage will more than triple in many segments by 2013.

Numerous security vendors are now leveraging cloud-based models to deliver security solutions. This shift has occurred for a variety of reasons including greater economies of scale and streamlined delivery mechanisms. Consumers are increasingly faced with evaluating security solutions that do not run on premises. Consumers need to understand the unique, diverse, and pervasive nature of cloud delivered security offerings so that they are in a position to evaluate the offerings and to understand if the offerings will meet their needs.

14.2 Concerns When Implementing Security as a Service

Despite the impressive array of benefits provided by cloud security services such as dynamic scalability, virtually unlimited resources, and greater economies of scale that exist with lower or no cost of ownership, there are concerns about security in the cloud environment. Some security concerns are around compliance, multi-tenancy, and vendor lock-in. While these are being cited as inhibitors to the migration of security into the cloud, these same concerns exist with traditional data centers.

Security in the cloud environment is often based on the concern that lack of visibility into security controls implemented means systems are not locked down as well as they are in traditional data centers and that the personnel lack the proper credentials and background checks. Security as a Service providers recognize the fragility of the relationship and often go to extreme lengths to ensure that their environment is locked down as much as possible. They often run background checks on their personnel that rival even the toughest government background checks, and they run them often. Physical and personnel security is one of the highest priorities of a Security as a Service provider.

Compliance has been raised as a concern given the global regulatory environment. Security as a Service providers have also recognized this and have gone to great efforts to demonstrate their ability to not only meet but exceed these requirements or to ensure that it is integrated into a client's network. Security as a Service providers should be cognizant of the geographical and regional regulations that affect the services and their consumers, and this can be built into the offerings and service implementations. The most prudent Security as a Service providers often enlist mediation and legal services to preemptively resolve the regulatory needs of the consumer with the regional regulatory requirements of a jurisdiction. When deploying Security as a Service in a highly regulated industry or environment, agreement on the metrics defining the service level required to achieve regulatory objectives should be negotiated in parallel with the SLA documents defining service.

As with any cloud service, multi-tenancy presents concerns of data leakage between virtual instances. While customers are concerned about this, the Security as a Service providers are also highly concerned in light of the litigious nature of modern business. As a result, a mature offering may take significant precautions to ensure data is highly compartmentalized and any data that is shared is anonymized to protect the identity and source. This applies equally to the data being monitored by the SecaaS provider and to the data held by them such as log and audit data from the client's systems (both cloud and non-cloud) that they monitor.

Another approach to the litigious nature of multi-tenant environments is increased analytics coupled with semantic processing. Resource descriptors and applied jurimetrics, a process through which legal reasoning is interpreted as high-

level concepts and expressed in a machine-readable format, may be employed proactively to resolve any legal ambiguity regarding a shared resource.

When utilizing a Security as a Service vendor, an enterprise places some, many or all security logging, compliance, and reporting into the custody of a provider that might sometimes have proprietary standards. In the event the enterprise seeks a new provider, they must concern themselves with an orderly transition and somehow find a way for the existing data and log files to be translated correctly and in a forensically sound manner.

It is important to note that other than multi-tenancy, each of these concerns is not "cloud unique" but are problems faced by both in-house models and outsourcing models. For this reason, non-proprietary unified security controls, such as those proposed by the Cloud Security Alliance Cloud Control Matrix, are needed to help enterprises and vendors benefit from the Security as a Service environment.

14.3 Advantages When Implementing Security as a Service

The potential strategic benefits of leveraging centralized security services are well understood by technical experts who witness the daily efficiencies gained. Just as cloud computing offers many advantages to both providers and consumers, Cloud Security as a Service offers many significant benefits due to a number of factors, including aggregation of knowledge, broad actionable intelligence, and having a full complement of security professionals on hand at all times, to name a few. Companies that are actively involved in the centralization and standardization of security best practices typically gain significant medium and long-term cost savings and competitive benefits over their rivals in the market due to the efficiencies gained. Security delivered as a service enables the users of security services to measure each vendor by a singular security standard thus better understanding what they are getting.

14.3.1 Competitive Advantages

Companies that employ third party security service providers gain a competitive edge over their peers due to early access to information helpful in understanding the risk proposition of a given IT strategy. Furthermore, through the use of a centralized security infrastructure, consumers are better able to stem the inclusion of undesirable content. Companies making use of a third party to report on regulatory compliance and measure obligatory predicates—the inherited legal and contractual obligations connected to identities and data—might result in the avoidance of costly litigation and fines that their competitors are vulnerable to. Once holistic security services are adopted and implemented, providers reap the competitive benefits of being able to assure their clients that they meet security best practice. Clients making use of these services have the advantage of being able to point to security providers as a part of their compliance framework and to third party assurance providers for proof of the achievement of service level agreement obligations.

14.3.2 Improved Vendor Client Relationship

There are many clear-cut benefits of security as a service. Transparency provided by a third party assurance service enables customers to understand exactly what they are getting, enabling easier comparison of vendor services and holding vendors to clear and agreed standards. Migration services enable the migration of data and services from one vendor to another. By leveraging migration services, consumers and providers are better enabled to exert market

pressures on their tertiary suppliers, enhancing the value for the enterprises that consume the services and securing the supply chain.

14.4 Diversity of Existing Security as a Service Offerings

Security as a Service is more than an outsourcing model for security management; it is an essential component in secure business resiliency and continuity. As a business resiliency control, Security as a Service offers a number of benefits. Due to the elastic model of services delivered via the cloud, customers need only pay for the amount they require, such as the number of workstations to be protected and not for the supporting infrastructure and staffing to support the various security services. A security focused provider offers greater security expertise than is typically available within an organization. Finally, outsourcing administrative tasks, such as log management, can save time and money, allowing an organization to devote more resources to its core competencies.

Gartner predicts that cloud-based security controls for messaging applications such as anti-malware and anti-spam programs will generate 60 percent of the revenue in that industry sector by 2013.

The areas of Cloud Security as a Service that most likely will interest consumers and security professionals are:

- Identity Services and Access Management Services
- Data Loss Prevention (DLP)
- Web Security
- Email Security
- Security Assessments
- Intrusion Management, Detection, and Prevention (IDS/IPS)
- Security Information and Event Management (SIEM)
- Encryption
- Business Continuity and Disaster Recovery
- Network Security

14.4.1 Identity, Entitlement, and Access Management Services

Identity-as-a-service is a generic term that covers one or many of the services that may comprise an identity eco-system, such as Policy Enforcement Points (PEP-as-a-service), Policy Decision Points (PDP-as-a-service), Policy Access Points (PAP-as-a-service), services that provide Entities with Identity, services that provide attributes, and services that provide reputation.

All these Identity services can be provided as a single stand-alone service, as a mix-and-match service from multiple providers, or today most probably a hybrid solution of public and private, traditional IAM, and cloud services.

These Identity services should provide controls for identities, access, and privileges management. Identity services need to include people, processes, and systems that are used to manage access to enterprise resources by assuring the identity of an entity is verified, then granting the correct level of access based on this assured identity. Audit logs of activities such as successful and failed authentication and access attempts should be managed by the application / solution or the SIEM service. Identity, Entitlement, and Access Management services are a Protective and Preventative technical control.

14.4.2 Data Loss Prevention

Monitoring, protecting, and demonstrating protection of data at rest, in motion, and in use both in the cloud and on premises, Data Loss Prevention (DLP) services offer protection of data usually by running as a client on desktops / servers and enforcing policies around what actions are authorized for particular data content. Where these differ from broad rules like 'no ftp' or 'no uploads to web sites' is the level to which they understand data, e.g., the user can specify no documents with numbers that look like credit cards can be emailed; anything saved to USB storage is automatically encrypted and can only be unencrypted on another office owned machine with a correctly installed DLP client; and only clients with functioning DLP software can open files from the file server. Within the cloud, DLP services may be offered as something that is provided as part of the build such that all servers built for that client get the DLP software installed with an agreed set of rules deployed. In addition, DLP may leverage central ID- or cloud brokers to enforce usage profiles. The ability to leverage a service to monitor and control data flows from an enterprise to the various tiers in the cloud service supply chain may be used as a preventative control for transborder transport, and subsequent loss, of regulated data such as **PII**[143]. This DLP offering is a preventative technical control.

14.4.3 Web Security

Web Security is real-time protection offered either on premise through software / appliance installation or via the Cloud by proxying or redirecting web traffic to the cloud provider. This provides an added layer of protection on top of other protection such as anti-malware software to prevent malware from entering the enterprise via activities such as web browsing. Policy rules around types of web access and the time frames when this is allowed can also be enforced via these technologies. Application authorization management can be used to provide an extra level of granular and contextual security enforcement for web applications. Web Security is a protective, detective, and reactive technical control.

14.4.4 Email Security

Email Security should provide control over inbound and outbound email, protecting the organization from phishing, malicious attachments, enforcing corporate polices such as acceptable use and spam prevention, and providing business continuity options. In addition, the solution should allow for policy-based encryption of emails as well as integrating with various email server solutions. Digital signatures enabling identification and non-repudiation are also features of many email security solutions. The Email Security offering is a protective, detective, and reactive technical control.

[143] **PII**-Personally Identifiable Information

14.4.5 Security Assessment

Security assessments are third party or customer-driven audits of cloud services or assessments of on premises systems via cloud provided solutions based on industry standards. Traditional security assessments for infrastructure, applications and compliance audits are well defined and supported by multiple standards such as NIST, ISO, and **CIS**[144]. A relatively mature toolset exists, and a number of tools have been implemented using the SecaaS delivery model. In the SecaaS delivery model, subscribers get the typical benefits of this cloud-computing variant—elasticity, negligible setup time, low administration overhead, and pay per use with low initial investments.

While not the focus of this effort, additional challenges arise when these tools are used to audit cloud environments. Multiple organizations, including the CSA have been working on the guidelines to help organizations understand the additional challenges:

- Virtualization awareness of the tool, frequently necessary for IaaS platform auditing
- Support for common web frameworks in PaaS applications
- Compliance Controls for Iaas, PaaS, and Saas platforms
- Automated incident and breach notification tools for maintenance of cloud supply chain integrity
- Standardized questionnaires for XaaS environments, that help address:
  - What should be tested in a cloud environment?
  - How does one assure data isolation in a multi-tenant environment?
  - What should appear in a typical infrastructure vulnerability report?
  - Is it acceptable to use results provided by cloud provider?

14.4.6 Intrusion Detection/Prevention (IDS/IPS)

Intrusion Detection/Prevention systems monitor behavior patterns using rule-based, heuristic, or behavioral models to detect anomalies in activity that present risks to the enterprise. Network IDS/IPS have become widely used over the past decade because of the capability to provide a granular view of what is happening within an enterprise network. The IDS/IPS monitors network traffic and compares the activity to a baseline via rule-based engine or statistical analysis. IDS is typically deployed in a passive mode to passively monitor sensitive segments of a client's network whereas the IPS is configured to play an active role in the defense of the clients network. In a traditional infrastructure, this could include De-Militarized Zones (**DMZ**'s)[145] segmented by firewalls or routers where corporate Web servers are locate or monitoring connections to an internal database. Within the cloud, IDS systems often focus on virtual infrastructure and cross-hypervisor activity where coordinated attacks can disrupt multiple tenants and create system chaos. Intrusion Detection Systems are detective technical controls, and Intrusion Prevention Systems are detective, protective, and reactive technical controls.

[144] **CIS**-Center for Internet Security
[145] **DMZ**-De-Militarized Zone

14.4.7 Security Information & Event Management (SIEM)

Security Information and Event Management (SIEM) systems aggregate (via push or pull mechanisms) log and event data from virtual and real networks, applications, and systems. This information is then correlated and analyzed to provide real time reporting and alerting on information or events that may require intervention or other types of responses. The logs are typically collected and archived in a manner that prevents tampering to enable their use as evidence in any investigations or historical reporting. The SIEM Security as a Service offering is a Detective technical control but can be configured to be a protective and reactive technical control.

14.4.8 Encryption

Encryption is the process of obfuscating/ encoding data using cryptographic algorithms, the product of which is encrypted data (referred to as cipher-text). Only the intended recipient or system that is in possession of the correct key can decode (un-encrypt) the cipher-text. Encryption for obfuscation systems typically consist of one or more algorithms that are computationally difficult (or infeasible) to break one or more keys, and the systems, processes, and procedures to manage encryption, decryption, and keys. Each part is effectively useless without the other, e.g., the best algorithm is easy to "crack" if an attacker can access the keys due to weak processes.

In the case of one-way cryptographic functions, a digest or hash is created instead. One-way cryptographic functions include hashing, digital signatures, certificate generation and renewal, and key exchanges. These systems typically consist of one or more algorithms that are easily replicated but very resistant to forgery, along with the processes and procedures to manage them. Encryption when outsourced to a Security as a Service provider is classified as a protective and detective technical control.

14.4.9 Business Continuity and Disaster Recovery

Business Continuity and Disaster Recovery are the measures designed and implemented to ensure operational resiliency in the event of any service interruptions. They provide flexible and reliable failover and DR solutions for required services in the event of a service interruption, whether natural or man-made. For example, in the event of a disaster scenario at one location, machines at different locations may protect applications in that location. This Security as a Service offering is a reactive, protective, and detective technical control.

14.4.10 Network Security

Network Security consists of security services that restrict or allocate access and that distribute, monitor, log, and protect the underlying resource services.

Architecturally, network security provides services that address security controls at the network in aggregate or those controls specifically addressed at the individual network of each underlying resource. In cloud / virtual environments and hybrid environments, network security is likely to be provided by virtual devices alongside traditional physical devices. Tight integration with the hypervisor to ensure full visibility of all traffic on the virtual network layer is key. These Network Security offerings include detective, protective, and reactive technical controls.

14.5 Permissions

- Implementers may employ pattern recognition of user activities.
- Implementers may employ secure legal mediation of security metrics for **SLA**[146] expectation management
- Implementers may employ provide trusted channels for penetration testing.

14.6 Recommendations

- Implementers should ensure secure communication channels between tenant and consumer.
- Providers should supply automated secure and continuous notification throughout the supply chain on a need-to-know basis.
- Providers should supply secured logging of internal operations for service level agreement compliance.
- Consumers should request addition of third party audit and SLA mediation services.
- All parties should enable Continuous Monitoring of all interfaces through standardized security interfaces such as SCAP (NIST), CYBEX (ITU-T), or RID & IODEF (IETF).

14.6 Requirements

14.6.1 Identity as a Service Requirements

- ✓ Providers of IaaS must provide cloud customers provisioning/de-provisioning of accounts (of both cloud & on-premise applications and resources).
- ✓ Providers of IaaS must provide cloud customers authentication (multiple forms and factors).
- ✓ Providers of IaaS must provide cloud customers identity life cycle management.
- ✓ Providers of IaaS must provide cloud customers directory services.
- ✓ Providers of IaaS must provide cloud customers directory synchronization (multi-lateral as required).
- ✓ Providers of IaaS must provide cloud customers federated SSO.
- ✓ Providers of IaaS must provide cloud customers web SSO (granular access enforcement & session management - different from federated SSO).

[146] **SLA**-Service Level Agreement

- ✓ Providers of IaaS must provide privileged session monitoring.
- ✓ Providers of IaaS must provide granular access management.
- ✓ Providers of IaaS must provide tamper-proof storage of audit records (including an option for non-repudiation).
- ✓ Providers of IaaS must provide policy management (incl. authorization management, role management, compliance policy management).
- ✓ Providers of IaaS must provide cloud customers authorization (both user and application/system).
- ✓ Providers of IaaS must provide cloud customers authorization token management and provisioning.
- ✓ Providers of IaaS must provide cloud customers user profile and entitlement management (both user and application/system).
- ✓ Providers of IaaS must provide cloud customers support for policy and regulatory compliance monitoring and/or reporting.
- ✓ Providers of IaaS must provide cloud customers federated provisioning of cloud applications.
- ✓ Providers of IaaS must provide privileged user and password management (including administrative, shared, system and application accounts).
- ✓ Providers of IaaS must provide cloud customers Role-Based Access Control (RBAC) (where supported by the underlying system/service).
- ✓ Providers of IaaS must provide cloud customers optionally support DLP integration.
- ✓ Providers of IaaS must provide cloud customers optionally support granular activity auditing broken down by individual.
- ✓ Providers of IaaS must provide cloud customers segregation of duties based on identity entitlement.
- ✓ Providers of IaaS must provide cloud customers compliance-centric reporting.
- ✓ Providers of IaaS must provide cloud customers centralized policy management.
- ✓ Providers of IaaS must provide cloud customers usable management interfaces.
- ✓ Providers of IaaS must provide cloud customers unified access control & audit.

- ✓ Providers of IaaS must provide cloud customers Interoperability and heterogeneity among various providers.
- ✓ Providers of IaaS must provide cloud customers scalability.

14.6.2 DLP SecaaS Requirements

- ✓ Providers of DLP must provide cloud customers with data labeling and classification.
- ✓ Providers of DLP must provide cloud customers with identification of Sensitive Data.
- ✓ Providers of DLP must provide cloud customers with predefined policies for major regulatory statues.
- ✓ Providers of DLP must provide cloud customers with context detection heuristics.
- ✓ Providers of DLP must provide cloud customers with structured data matching (data-at-rest).
- ✓ Providers of DLP must provide cloud customers with SQL regular expression detection.
- ✓ Providers of DLP must provide cloud customers with traffic spanning (data-in-motion) detection.
- ✓ Providers of DLP must provide cloud customers with Real Time User Awareness.
- ✓ Providers of DLP must provide cloud customers with security level assignment.
- ✓ Providers of DLP must provide cloud customers with custom attribute lookup.
- ✓ Providers of DLP must provide cloud customers with automated incident response.
- ✓ Providers of DLP must provide cloud customers with signing of data.
- ✓ Providers of DLP must provide cloud customers with cryptographic data protection and access control.
- ✓ Providers of DLP must provide cloud customers with machine-readable policy language.

14.6.3 Web Services SecaaS Requirements

- ✓ Providers of Web Services SecaaS must provide must provide cloud customers with web monitoring and filtering.
- ✓ Providers of Web Services SecaaS must provide must provide cloud customers with Malware, Spyware, and Bot Network analyzer and blocking.
- ✓ Providers of Web Services SecaaS must provide must provide cloud customers with phishing site blocker.
- ✓ Providers of Web Services SecaaS must provide must provide cloud customers with instant messaging scanning.
- ✓ Providers of Web Services SecaaS must provide must provide cloud customers with email security.

- ✓ Providers of Web Services SecaaS must provide must provide cloud customers with bandwidth management / traffic control.
- ✓ Providers of Web Services SecaaS must provide must provide cloud customers with Data Loss Prevention.
- ✓ Providers of Web Services SecaaS must provide must provide cloud customers with fraud prevention.
- ✓ Providers of Web Services SecaaS must provide must provide cloud customers with Web Access Control.
- ✓ Providers of Web Services SecaaS must provide must provide cloud customers with backup.
- ✓ Providers of Web Services SecaaS must provide must provide cloud customers with SSL (decryption / hand off).
- ✓ Providers of Web Services SecaaS must provide must provide cloud customers with usage policy enforcement.
- ✓ Providers of Web Services SecaaS must provide must provide cloud customers with vulnerability management.
- ✓ Providers of Web Services SecaaS must provide must provide cloud customers with web intelligence reporting.

14.6.4 Email SecaaS Requirements

- ✓ Providers of Email Security SecaaS must provide cloud customers with accurate filtering to block spam and phishing.
- ✓ Providers of Email Security SecaaS must provide cloud customers with deep protection against viruses and spyware before they enter the enterprise perimeter.
- ✓ Providers of Email Security SecaaS must provide cloud customers with flexible policies to define granular mail flow and encryption.
- ✓ Providers of Email Security SecaaS must provide cloud customers with rich, interactive reports and correlate real-time reporting.
- ✓ Providers of Email Security SecaaS must provide cloud customers with deep content scanning to enforce policies.
- ✓ Providers of Email Security SecaaS must provide cloud customers with the option to encrypt some / all emails based on policy.
- ✓ Providers of Email Security SecaaS must provide cloud customers with integration capability to various email server solutions.

14.6.5 Security Assessment SecaaS Requirements

- ✓ Providers of Security Assessment SecaaS must provide cloud customers with detailed governance processes and metrics (Implementers should define and document and process by which policies are set and decision making is executed).

- ✓ Providers of Security Assessments and Governance offerings should implement an automated solution for notifying members of their immediate supply chain in the event of breach or security incident.
- ✓ Providers of Security Assessment SecaaS must provide cloud customers with proper risk management (Implementers should define and document and process for ensuring that important business processes and behaviors remain within the tolerances associated with those policies and decisions).
- ✓ Providers of Security Assessment SecaaS must provide cloud customers with details of compliance (Implementers should define and document process-of-adherence to policies and decisions).
- ✓ Providers of Security Assessment SecaaS must provide cloud customers with policies that can be derived from internal directives, procedures, and requirements or external laws, regulations, standards and agreements.
- ✓ Providers of Security Assessment SecaaS must provide cloud customers with technical compliance audits (automated auditing of configuration settings in devices, operating systems, databases, and applications).
- ✓ Providers of Security Assessment SecaaS must provide cloud customers with application security assessments (automated auditing of custom applications).
- ✓ Providers of an assessment and governance service offering must provide cloud customers with vulnerability assessments—automated probing of network devices, computers, and applications for known vulnerabilities and configuration issues.
- ✓ Providers of Security Assessment SecaaS must provide cloud customers with penetration testing (exploitation of vulnerabilities and configuration issues to gain access to an environment, network or computer, typically requiring manual assistance)
- ✓ Providers of Security Assessment SecaaS must provide cloud customers with a security rating.

14.6.6 Intrusion Detection SecaaS Requirements

- ✓ Providers of Intrusion Detection SecaaS must provide cloud customers with identification of intrusions and policy violations.
- ✓ Providers of Intrusion Detection SecaaS must provide cloud customers with automatic or manual remediation actions.
- ✓ Providers of Intrusion Detection SecaaS must provide cloud customers with Coverage for Workloads, Virtualization Layer (VMM/Hypervisor) Management Plane
- ✓ Providers of Intrusion Detection SecaaS must provide cloud customers with deep packet inspection using one or more of the following techniques: statistical, behavioral, signature, heuristic.
- ✓ Providers of Intrusion Detection SecaaS must provide cloud customers with system call monitoring.
- ✓ Providers of Intrusion Detection SecaaS must provide cloud customers with system/application log inspection.

- ✓ Providers of Intrusion Detection SecaaS must provide cloud customers with integrity monitoring OS (files, registry, ports, processes, installed software, etc.)
- ✓ Providers of Intrusion Detection SecaaS must provide cloud customers with integrity monitoring VMM/Hypervisor.
- ✓ Providers of Intrusion Detection SecaaS must provide cloud customers with VM Image Repository Monitoring.

14.6.7 SIEM SecaaS Requirements

- ✓ Providers of SIEM SecaaS must provide cloud customers with real time log /event collection, de-duplication, normalization, aggregation and visualization.
- ✓ Providers of SIEM SecaaS must provide cloud customers with forensics support.
- ✓ Providers of SIEM SecaaS must provide cloud customers with compliance reporting and support.
- ✓ Providers of SIEM SecaaS must provide cloud customers with IR support.
- ✓ Providers of SIEM SecaaS must provide cloud customers with anomaly detection not limited to email.
- ✓ Providers of SIEM SecaaS must provide cloud customers with detailed reporting.
- ✓ Providers of SIEM SecaaS must provide cloud customers with flexible data retention periods and flexible policy management

14.6.8 Encryption SecaaS Requirements

- ✓ Providers of Encryption SecaaS must provide cloud customers with protection of data in transit.
- ✓ Providers of Encryption SecaaS must provide cloud customers with protection of data at rest.
- ✓ Providers of Encryption SecaaS must provide cloud customers with key and policy management.
- ✓ Providers of Encryption SecaaS must provide cloud customers with protection of cached data.

14.6.9 Business Continuity and Disaster Recovery Requirements

- ✓ Providers of Business Continuity & Disaster Recovery SecaaS must provide cloud customers with flexible infrastructure.
- ✓ Providers of Business Continuity & Disaster Recovery SecaaS must provide cloud customers with secure backup.
- ✓ Providers of Business Continuity & Disaster Recovery SecaaS must provide cloud customers with monitored operations.
- ✓ Providers of Business Continuity & Disaster Recovery SecaaS must provide cloud customers with third party service connectivity.

- ✓ Providers of Business Continuity & Disaster Recovery SecaaS must provide cloud customers with replicated infrastructure component.
- ✓ Providers of Business Continuity & Disaster Recovery SecaaS must provide cloud customers with replicated data (core / critical systems).
- ✓ Providers of Business Continuity & Disaster Recovery SecaaS must provide cloud customers with data and/or application recovery.
- ✓ Providers of Business Continuity & Disaster Recovery SecaaS must provide cloud customers with alternate sites of operation.
- ✓ Providers of Business Continuity & Disaster Recovery SecaaS must provide cloud customers with tested and measured processes and operations to ensure operational resiliency.
- ✓ Providers of Business Continuity & Disaster Recovery SecaaS must provide cloud customers with geographically distributed data centers / infrastructure.
- ✓ Providers of Business Continuity & Disaster Recovery SecaaS must provide cloud customers with Network survivability.

14.6.10 Network Security SecaaS Requirements

- ✓ Providers of Network Security SecaaS must provide cloud customers with details of data threats.
- ✓ Providers of Network Security SecaaS must provide cloud customers with details of access control threats.
- ✓ Providers of Network Security SecaaS must provide cloud customers with access and authentication controls.
- ✓ Providers of Network Security SecaaS must provide cloud customers with security gateways (firewalls, WAF, SOA/API).
- ✓ Providers of Network Security SecaaS must provide cloud customers with security products (IDS/IPS, Server Tier Firewall, File Integrity Monitoring, DLP, Anti-Virus, Anti-Spam).
- ✓ Providers of Network Security SecaaS must provide cloud customers with security monitoring and incident response.
- ✓ Providers of Network Security SecaaS must provide cloud customers with DoS protection/mitigation.
- ✓ Providers of Network Security SecaaS must provide cloud customers with Secure “base services” like DNSSEC, NTP, OAuth, SNMP, management network segmentation, and security.
- ✓ Providers of Network Security SecaaS must provide cloud customers with traffic / netflow monitoring.
- ✓ Providers of Network Security SecaaS must provide cloud customers integration with Hypervisor layer.

REFERENCES

[1] Security and Economic Benefits of Standardization for Security as a Service. September 2011 Proceedings. United Nations ITU-T.

[2] Gartner. Gartner Says Security Delivered as a Cloud-Based Service Will More Than Triple in Many Segments by 2013. July 15, 2008. http://www.gartner.com/it/page.jsp?id=722307