



Evaluating Cloud Risk for the Enterprise:
A Shared Assessments Guide

October 2010

Published By

THE **SANTA FE** GROUP

©2010 The Shared Assessments Program. All Rights Reserved.

Table of Contents

About the Shared Assessments Program	4
Acknowledgments	6
Foreword	7
Introduction	8
Cloud Computing: An Overview	11
A Risk Management Approach: Common and Delta Controls	15
Cloud Computing Case Study	40
Glossary	43
Appendix: Additional Cloud Computing Initiatives	48

©Shared Assessments 2010

Complete and accurate documents created under the Shared Assessments Program may be downloaded from the official Shared Assessments Program website at www.sharedassessments.org.

While retaining copyrights, the Shared Assessments Program makes specific documents available to the public for the purpose of conducting self-assessments and third-party security assessments. Licenses for other uses are available from the Shared Assessments Program. Individuals and organizations should review the terms of use prior to downloading, copying, using or modifying Shared Assessment Program documents.

This notice must be included on any copy of the Shared Assessments Program documents, excluding assessors' AUP reports.

The Shared Assessments Program is administered by The Santa Fe Group (www.santa-fe-group.com). Questions about this document and the Program should be directed to:

Joyce Crawshaw
Director Client Relations
The Santa Fe Group
505-466-6434
sharedassessments@santa-fe-group.com

About the Shared Assessments Program

The vendor assessment control evaluation process has long been inefficient and costly. Many organizations that assess their technology service providers produce and distribute their own proprietary questionnaire to each of their service providers. The volume of diverse client questionnaires to which service providers must respond puts a significant strain on their resources. The disparity of information requested from questionnaire to questionnaire can cause delays for all parties.

Shared Assessments was created by leading financial institutions, the Big Four accounting firms and leading service providers to inject standardization, consistency, speed, efficiency and cost savings into the service provider vendor assessment process. Through [membership](#) in the Shared Assessments Program and use of the Shared Assessments tools (the [Standardized Information Gathering Questionnaire](#), or “SIG” and [Agreed Upon Procedures](#) or “AUP”), Shared Assessments strives to eliminate redundancies and create efficiencies, giving all parties a faster, more efficient and less costly means of conducting rigorous and comprehensive security, privacy and business continuity assessments.

To promote adoption, the Shared Assessments Program makes its standards available for download at the Shared Assessments website, www.sharedassessments.org. These documents are reviewed annually by Shared Assessments members and updated for consistency with evolving security, privacy and business continuity standards.

A Global Community

Shared Assessments [members](#) are national and international organizations of all sizes that understand the importance of comprehensive standards for managing risk. Members include representatives from a range of industries: financial institutions, healthcare organizations, retailers, and telecommunications companies. Membership also includes service providers, consulting companies and assessment firms of all sizes. All of these companies are committed to being best-in-class members of a global community of risk management experts who understand the value of implementing efficient and effective standard assessment practices.

In addition to its members, the Shared Assessments Program has strategic alliances with global associations including the [National Association of Software and Services Companies](#) (NASSCOM) and the [Securities Industry and Financial Markets Association](#) (SIFMA). Shared Assessments also continues its affiliation with [BITS](#). [The Santa Fe Group](#) manages the Program. Together, Shared Assessments' diverse membership works to increase awareness and adoption of the Program tools across industry sectors and around the globe.

Shared Assessments and Cloud Computing

With the publication of version 5.0 of the SIG and AUP in 2009, the Shared Assessments Program began specifically addressing cloud computing by adding six new procedures to its on-site assessment tool (the AUP) and inserting cloud-relevant questions into several sections of the Shared Assessments questionnaire (the SIG). Enhancements continue to be made to the SIG and AUP to improve their effectiveness, including updates that reflect the growing importance of cloud computing across the IT landscape.

About The Santa Fe Group

[The Santa Fe Group](#) is a strategic consulting firm that specializes in business strategy, payments strategies, risk management, emerging technologies and innovation. Drawing on a national network of consultants and executives with specialized expertise, The Santa Fe Group provides strategic consulting, senior executive briefings, research studies, educational and training programs, publications, seminars and other services. The Santa Fe Group manages the [Shared Assessments Program](#), including facilitating its Member Forum and working groups and managing the Shared Assessments standards. The Santa Fe Group's Chairman and CEO is Catherine A. Allen, an award-winning financial services visionary and the founding CEO of financial services industry consortium BITS.

Acknowledgments

Project Chair

Niall Browne, CISO & VP Information Security, LiveOps Inc.

Editors

Niall Browne, CISO & VP Information Security, LiveOps Inc.

Susanna Space, Vice President, Communications and Business Development, The Santa Fe Group

Contributors

Georges Ataya, ISACA

Brian G. Barnier, ValueBridge Advisors, USA

Paul A. Bateman, Goldman Sachs

Lisa Picard, Fishnet Security

Daniel Burks, US Bank

French Caldwell, Gartner, Inc.

Chris Carcich, CISSP CISM CISM MCSE

John A. DiMaria, Certified Six Sigma BB, HISP, eFortresses Inc.

Clint Harris, CISSP, AT&T Consulting Services, Inc.

Donna Hiers, US Bank

Eddie Holt, KPMG LLP

Aileen Johnson, Target

Randy Kiriara, Target

Daniel Kramer, Target

Mark Lundin, KPMG LLP

Andrew MacLeod, ISACA

Adrian Mikeliunas, CISSP, CISA, PCI-QSA, AT&T Consulting Solutions, Inc.

Ray Murphy, Navy Federal Credit Union

Pritesh Parekh, MS, MBA, CISSP, CISM, CISA, CEH, Yodlee, Inc.

Kevin Scott, Deluxe Corporation

Sue Suhling, Target

Becky Swain, CIPP, CIPP/IT, CISSP, CISA, Cisco; Co-Founder, Co-Chair Cloud

Security Alliance Controls Matrix

Donald Williams, Churchill & Harriman

Foreword

This guide was published to help businesses and individuals understand and evaluate the use of cloud computing in large enterprises. The authors' overarching goal is to raise awareness of cloud computing risks in order to better enable Enterprise businesses to successfully deploy cloud computing technologies.

This guide approaches various aspects of cloud computing environments from a risk perspective. While many aspects of cloud computing resemble traditional hosting environments, new technologies often present unique and unknown risks that must be considered before and during migration to these environments. To address this, the discussion of control considerations is broken into two categories:

1. **Common Controls:** These are mature control areas associated with traditional IT services environments that are also applicable to cloud-based services, and whose audit mechanisms are considered mature.
2. **Delta Controls:** These are higher-risk control areas that have particular relevance to cloud environments, and whose cloud audit mechanisms are less mature.

Also included in this Guide are practical recommendations, questions to discuss with cloud providers, and lessons learned for each of the Common and Delta Cloud Control areas.

The Shared Assessments Program will convert the cloud control areas, recommendations and best practices set out here into formal control objectives that comply with formal audit standards. These controls will be incorporated into the 2010 and 2011 Shared Assessments standards (the AUP and SIG)¹.

The recommendations and guidance in this document may also be selected and incorporated into other types of audits or assessments of environments containing cloud elements, such as the Statements on Auditing Standards (SAS) 70 or the upcoming Statements on Standards for Attestation Engagements (SSAE) 16.

Ultimately, this guide targets audiences with a variety of levels of cloud expertise and knowledge. Cloud users may choose to read the document from start to finish or read relevant sections, using it as a reference tool.

We hope you find this guide helpful. Any suggestions or questions should be directed to the Shared Assessments Program Cloud Computing Working Group at cloud@sharedassessments.org.

¹ The Shared Assessments Agreed Upon Procedures and Standardized Information Gathering Questionnaire. For more information about these documents, please visit <http://sharedassessments.org/about/programtools.html>.

Introduction

By 2012, 20 percent of businesses will own no IT assets. Several interrelated trends are driving the movement toward decreased IT hardware assets, such as virtualization, cloud-enabled services, and employees running personal desktops and notebook systems on corporate networks.

The need for computing hardware, either in a data center or on an employee's desk, will not go away. However, if the ownership of hardware shifts to third parties, then there will be major shifts throughout every facet of the IT hardware industry. For example, enterprise IT budgets will either be shrunk or reallocated to more strategic projects; enterprise IT staff will either be reduced or reskilled to meet new requirements, and/or hardware distribution will have to change radically to meet the requirements of the new IT hardware buying points.

— Gartner, January 2010²

Cloud computing is one of the most talked about trends in the IT industry. Referred to as "transformational" and "a game-changer" in analyst articles and news stories, cloud represents a market-changing shift with widespread business impact.

While many papers and articles have explored cloud computing for smaller organizations or departments within larger businesses, this Guide was created with a different audience in mind. The goal of this Guide is to help individuals understand cloud from the perspective of the entire enterprise, with the goal of enabling successful deployment across departments and lines of business.

To this end, this Guide offers:

- An overview of the characteristics of cloud computing
- A risk management approach to evaluating cloud computing, including an exploration of Common and Delta Cloud Computing Controls with detailed practical recommendations for each control area
- A case study outlining the evaluation and implementation of cloud by one of the largest organizations in the US
- A list of industry leaders and ongoing initiatives relating to cloud

This Guide does not advocate for any position on cloud computing. Instead, it describes cloud computing services, identifies key issues related to cloud, and offers analyses that can be used in developing a comprehensive plan for evaluating, risk ranking and cost-effectively selecting cloud providers and solutions.

² Gartner Highlights Key Predictions for IT Organizations and Users in 2010 and Beyond, January 13, 2010. See <http://www.gartner.com/it/page.jsp?id=1278413>

While the term "cloud computing" is relatively new, one of the core components of cloud is its distributed nature. (Cloud has alternately been called "distributed computing.") Prior to the advent of cloud environments, data had become progressively more distributed, moving from a model in which it is stored and accessed in a central location to a much less centralized model.

Over the past ten years, two trends have emerged that have further increased data distribution:

- **Increased outsourcing of data, services and processes:** Traditionally, data and services have resided deep within an organization, protected by the company's interest in keeping its data safe and secure. With the advent of widespread outsourcing of data and services, data is now often stored and protected by numerous third parties, often in multiple locations.
- **Remote access for workers:** For decades, the vast majority of staff and contractors drove to the office, accessed their company's critical data from inside its walls, and then went home, leaving the company's critical assets and data protected from criminal compromise. With the advent of laptops, VPNs, Blackberrys, and other devices, most workers now have the opportunity to work effectively from home, whether temporarily or on a full-time basis. Remote access to company systems and data has allowed information to be stored on workers' local systems, resulting in increasing risk of data loss due to loss and theft.

The growing adoption of cloud computing has created a pressing need for further analysis and investigation of the controls in distributed environments. Rather than fitting into contained silos protecting the so-called "four walls" of the building, today's controls need to focus on all of the locations where data resides. One location may be within encrypted storage arrays in data centers, where the data is protected by biometrics, IPS, armed guards and hardened systems, at five o'clock that data may also sit on company laptops in homes, in airports and on the front seat of the car — all locations where perimeter controls are of little or no use.

Outsourcing and remote work environments have significantly eroded the concept and practicality of perimeter security, and risk management as a whole has often lagged in devising the critical alternative controls to protect the increasingly porous enterprise environments. With distributed models, perimeter security diminishes in effectiveness, leaving significant exposures. Unfortunately, the steady and often invisible movement from centralized to distributed systems over the past ten years has lulled risk and control professionals into a false sense of security. Many have believed there would be plenty of time to start building much-needed controls for distributed models before they became widely deployed.

Cloud computing has exposed data's distributed nature. Today, it is impossible to deny that location-centric and perimeter-based data controls are of limited value. Business and IT managers must shift their focus to "data-centric" controls: controls that are focused on protecting the data itself, rather than any one location. It is these controls, in the enterprise context, that are the focus of this Guide.

Cloud users should not have to sacrifice security for convenience. Clients should expect and demand all of the risk management and security controls of traditional on-premise providers — and more — from their cloud solutions. Cloud providers are uniquely positioned to build environments and corresponding security controls from the ground up. Still, the onus is on client companies to ensure that proper due diligence is completed. Using this Guide as a starting point, companies can begin to evaluate specific areas of cloud risk, ask the right questions, and ensure they get answers they understand. If a cloud provider cannot adequately respond to specific information requests, such as the exact location of data and the corresponding controls, enterprise users should consider selecting a provider that can.

Cloud Computing: An Overview

Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. This cloud model promotes availability and is composed of five essential characteristics, three service models, and four deployment models.

—National Institute of Standards in Technology
Definition of Cloud Computing³

Computing is always evolving. Since the 1970s, trend has followed trend, from mainframes to PCs to laptops to Blackberrys and iPads. Today, cloud is ubiquitous in discussions of IT and its future, and offers a unique blend of both old and new elements of computing, making for a compelling and powerful concept.

Cloud computing is more evolutionary than revolutionary; its underlying components have been in existence for some time. Yet cloud is highly disruptive to the IT industry, affecting all levels of IT management and vendors of all kinds. Whether cloud will prove to be a sustainable IT consumption model that delivers superior business value in comparison to previous computing models remains to be seen.

Characteristics of Cloud Computing

What makes cloud computing unique? Cloud services generally have the following characteristics that set them apart from other technology providers:

- Users often don't own, house or control the computing assets. Instead, computers and storage are housed in external data centers.⁴
- Service is delivered on a pay-per-use (utility) or subscription model.
- Resources and services are often virtual and shared by multiple parties.
- Services are delivered via the Internet.

These qualities allow cloud to offer unprecedented options for software utilization and flexibility.

³ See <http://csrc.nist.gov/groups/SNS/cloud-computing/>

⁴One exception to this is an in-house private cloud.

Benefits of Cloud Computing

Cloud services offer many advantages, including:

- Speed, including faster deployment of software services
- Lower computing costs in the form of reduced IT infrastructure expense (including hardware and software maintenance costs); companies pay only for what they use
- Reduced dependence on internal IT resources
- Online applications that facilitate collaboration
- Accessibility by low-end devices of computational or storage-intensive applications
- Scalability through on-demand computing
- Ubiquitous access (multiple networks, remote access, mobile devices)
- Improved performance (through the pooling and sharing of hardware resources)
- Cost-effective security via economies of scale (multiple clients share the cost of enterprise security controls)

In addition, companies can realize process and cost efficiencies when IT services that are essential to the delivery of cloud services — such as system administration, data backup, security and hardware/software maintenance — are shifted to the cloud provider.

Cloud Computing Services: IaaS, PaaS and SaaS

There is no single definition of cloud computing. Instead, cloud consists of many different types of services. This Guide defines cloud computing as consisting of three distinct service types:

- **Infrastructure as a Service (IaaS):** IaaS vendors offer turnkey data-center infrastructure to customers. IaaS is a provision model in which an organization outsources the equipment used to support operations, including storage, hardware, servers and networking components. Customers typically develop their own applications but do not necessarily want to provide and manage the computing infrastructure required to run them. An IaaS vendor often provides these services on a capacity-based payment stream.

IaaS evolved from application hosting outsourcing. What makes IaaS different from traditional hosting is its "multi-tenant" nature: multiple customers share certain aspects of the cloud infrastructure. The phrase "utility computing" is often associated with IaaS. Telecommunications vendors, for example, may be in a strong position to offer IaaS due to their traditional hosting services and network strengths.

Hardware, telecommunications and outsourcing vendors are rapidly moving into the IaaS market. These vendors see economies of scale that can be exploited by building massive data centers to serve multiple customers with the need for scalable, on-demand computing resources. Considering current customer investments in their own data centers, there is enormous potential for leading vendors offering cloud to grow in this space. Many of these vendors have core competencies in operating large data centers, allowing them to enter the

market with a high degree of credibility for their perceived stability, availability, protection and recoverability.

- **Platform as a Service (PaaS):** Most commonly used by application developers, PaaS vendors offer hardware and software infrastructure for the development of business applications. If a customer does not want to acquire and manage development tools such as programming languages, databases and related infrastructure, a PaaS vendor can provide them on an as-needed basis. PaaS is increasingly being used as a marketplace for applications by developers such as Google and Salesforce.com. This significantly reduces capital costs and can speed the development of business applications.

With PaaS, the customer develops its own application using the PaaS cloud rather than its own onsite development environment. Once developed, the application is typically run "from the cloud" and made available for use by the customer via the Internet and a web browser.

- **Software as a Service (SaaS):** Considered the most mature cloud computing service, SaaS refers to a business application delivered over the Internet in which users interact with the application through a web browser. SaaS applications are designed with a significant degree of network and device independence. SaaS is most commonly used by individuals, small- to mid-sized businesses and departments within larger enterprises.

The SaaS vendor provides the business application in a complete, ready-to-run state, with the application residing on computing infrastructure that is either owned or managed by the SaaS vendor or outsourced to a third-party vendor in a hosted or IaaS model.

The business application is developed and maintained by the SaaS vendor, which is responsible for all bug fixes and enhancements to the application, as well as all services related to the underlying hardware and software infrastructure supporting the application.

These three cloud service types can be viewed as a pyramid (Figure 1): IaaS is the lowest level of cloud service and forms the base layer; PaaS is the middle; and SaaS is the top of the pyramid. As a cloud infrastructure, the IaaS layer can host PaaS and SaaS environments.

Figure 1: Three Cloud Computing Service Types



Cloud vs. Hosted Applications (Traditional Model)

Hosted application services have been available in the IT marketplace for many years. While cloud and hosted applications have much in common, it is important to understand how cloud is different.

Cloud and hosted applications are similar in that they are both forms of outsourcing. However, cloud can bundle a software product with an ongoing service, while a hosted application is most often a pure service in which the customer typically provides the application. With a hosted application, the hosting vendor and the customer supplying the application often share responsibility for security. With cloud, the vendor is responsible for most of the security controls and incident preparation. Infrastructure may be shared among unrelated customers of a hosted application provider; cloud environments offer a higher degree of sharing with the potential for multiple customers to use one cloud solution.

As with any vendor model, an organization can outsource the responsibility for the service, but not the associated risk or accountability.

Cloud vs. Licensed Software Vendors

Cloud delivery models are attracting the attention of leading vendors across all segments of the IT industry. Many vendors of traditionally licensed business software are attracted to cloud because it offers a way to extend of their current business model, with the potential for greater sales, profitability and customer longevity. Cloud providers also view on-demand technology as advantageous because it meets customer demand for speed-to-market and efficiency, allowing clients to outsource responsibilities related to application administration and maintenance.

For all of these reasons, cloud can represent a substantial risk to vendors with a market-dominating presence in licensed software. These vendors are adopting cloud more slowly in order to avoid cannibalizing their entrenched products. This may prove to be a competitive disadvantage as other vendors with less to lose aggressively enter this market.

Take, for example, Google and Microsoft. Google is aggressively developing cloud applications, in many cases making them available for free. Meanwhile, Microsoft, with its enormous base of licensed-software customers, is approaching the cloud market more cautiously. Google may not have a core competency in licensed software, but the company is exploiting its knowledge of the Internet and datacenters to launch web-based applications in a cloud model.

A Risk Management Approach: Common and Delta Controls

Business managers and IT specialists frequently ask about the differences between traditional outsourcing and cloud models from the perspective of security, auditing and risk management.

This seemingly innocuous line of questioning is critical, since so many enterprise organizations have been outsourcing successfully using traditional models and have acquired a wealth of corresponding knowledge and experience. Companies are naturally eager to leverage their prior knowledge, expertise and lessons learned with well-used traditional services models, rather than having to "reinvent the wheel" in order to evaluate the cloud model and cloud providers.

To help answer these questions, companies first need to separate the traditional controls that are also present in cloud models from those controls that are considered particularly relevant to cloud models. For purposes of this guide, these controls have been grouped in to two categories:

1. **Common Cloud Controls:** These are mature control areas associated with traditional IT services environments that are also applicable to cloud-based services, and whose audit mechanisms are considered mature.
2. **Delta Cloud Controls:** These are higher-risk control areas that have particular relevance to cloud environments, and whose cloud audit mechanisms are less mature.

An enterprise organization evaluating a cloud solution or provider might have a list of 100 controls to examine, from IT management processes, information security polices and risk management, to antivirus, recovery and capacity management. Each of these areas presents a different level of risk. Since resources are always finite, spending an equal amount of time examining each of the 100 controls without regard to their importance or risk is likely to leave the higher risk control areas (visualization, for example) insufficiently examined, and the company exposed.

To remedy this, consider applying a risk management approach to cloud engagements. Beginning with an examination of the Common Cloud Controls, use the cloud provider's existing audit reports and certifications. Approaching the evaluation this way allows the majority of controls and risks to be checked much more efficiently, while maintaining rigor by using mature testing methods. It also helps avoid unnecessary duplication of effort.

Next, move on to the higher risk and newer Delta Cloud Controls that have particular significance to the cloud. Your team may not be as experienced in evaluating these controls, and your existing audit programs may not cover them sufficiently. (For example, virtualization is largely ignored or omitted entirely in PCI-DSS, ISO 27002, and HIPAA audits.)

The Delta Cloud Controls section of this Guide provides an overview of 12 cloud computing Delta Control areas. These areas include numerous recommendations for examining and evaluating these cloud controls.

Common Cloud Controls

Many of today’s enterprise organizations are well versed in evaluating traditional IT controls. Much of the auditing guidance that has existed for many years may also be applied to cloud environments. Traditional outsourcing models and cloud providers can each be compliant with the same standards and practices. If a cloud provider can demonstrate its compliance with an existing guidance, for example the Shared Assessments AUP, SAS 70, ISO 27001 or PCI-DSS, then a significant portion of the assessment for those areas may be completed using standard methodologies, often with little additional cost.

Common Cloud Control areas are divided into two categories:

1. [Frequently Used Information System Controls](#)
2. [Mature Relationship, Procurement and Vendor Management Processes](#)

1. Frequently Used Information Systems Controls

Frequently used risk frameworks, management processes and controls are often agnostic to delivery environment. They are used to evaluate the operational or service delivery risk of an IT environment (operational stability, availability, protection and recovery). The widely used guidance documents cited in the table below illustrate that many assessment models were in use long before cloud emerged.

Table 1: Guidance Types, Characteristics and Examples

Guidance Type	Characteristics	Examples
Frameworks	Illustrates how multiple guidance areas (sometimes called “domains”) relate to each other and contain multiple levels of depth. Often include capability models, RACI tables, process models and some level of controls.	ISACA COBIT 4.1, ISACA Risk IT Based on COBIT
Management Processes	Illustrates how management processes are used to implement capabilities to achieve objectives. Typically includes input-output tables, goals and objectives tables and some level of controls.	ISACA COBIT 4.1, ISACA Risk IT Based on COBIT, ISO 27001-2005 (main contents)

Level 1 and 2 Controls	Illustrates how business objectives are related to IT objectives. Typically includes controls related to the processes (to provide assurance in a security, procurement or change-management process), but not specific to an individual environment, hardware device, software or facility.	ISACA COBIT 4.1, ISACA Risk IT Based on COBIT, ISO 27001-2005 (Annex A controls), NIST 800-53A Rev1 (Appendix F procedure catalog, controls portion), Shared Assessments AUP (controls portion)
Level 3 and 4 Controls	Illustrates more granular aspects of process control and/or controls specific to an individual environment, hardware device, software or facility. For example, patching a specific server or managing availability in a particular website configuration.	PCI DSS, NIST 800-53A Rev1 (Appendix F procedure catalog, controls portion), Shared Assessments SIG and AUP. Vendor-specific security guidance.

Considerations

- In an assessment, the frameworks, management process and controls must be utilized together. The failures that many experience with the checklist approach to controls (including those with certification) are well known; checklists have significant weaknesses when used outside of highly defined and rarely changing processes that are clearly scoped in advance. Superficial use of frameworks has also led to failure. Instead of relying on quick solutions, risk management must be a living process, continually improving risk governance, risk evaluation and risk response, including preparedness and controls.
- There are significant content similarities among control guidance documents. For this reason, many mappings exist to illustrate the overlap from guidance to guidance. One of the better established and globally accepted control frameworks is ISACA's COBIT 4.1. It is mapped to more guidance documents than most and is considered by many to be the "Rosetta Stone" of IT-related guidance. The Shared Assessments Standardized Information Gathering questionnaire (SIG) and Agreed Upon Procedures (AUP) also map to a number of other guidance documents, including COBIT, PCI-DSS, ISO 27001 and others.
- The Cloud Cube Model (Figure 1) illustrates the different types of cloud deployments. The table below illustrates some of the *environment types*. Each combination of deployments and environments carries its own risk considerations, including threats, ability to respond and jurisdictional requirements. Each environment combination is distinct. For example, the "private cloud/off-shore" combination is protected from mixed use, but carries risks because the delivery elements are under the jurisdiction of the laws of another country.

Table 2: Environment Types

	Owned by Your Corporate Family		Owned by a Third Party or Joint Venture	
	Traditional Internal	“In-Sourced” /Private Cloud	Private Cloud	Shared/Public Cloud
On-shore				
Off-shore				

- Understanding a) the differences in the cloud environments and b) the adequacy of your current controls creates an opportunity to address IT-related business risks in cloud environments more easily. Begin by assessing risk in a specific type of cloud delivery environment (see Table 2), as well as the people, technology and processes that will be used to deliver the service.
- To tailor the assessment to a new delivery environment, look closely at *what is changing*. Changes introduce different risks (threats, exposures, vulnerabilities, frequency and impact) and thus demand new or improved management practices and controls to bring the risk within acceptable limits.
- Table 3 below may be used as an aid in change analysis.

Table 3: Change Analysis Table

Aspect Subject to Change	Considerations	Change Type	Common Control
Physical Location	Political and country risk, jurisdiction requirements, time zone, latency		
People	Skill, length of service, ongoing training and continuous improvement approach		
Business Process	New interactions for customers or internal users		
Technology:			
Application	Range of user access, potential for user error		
Middleware	A connection point between building blocks		
Data	Range of user access		
Servers	Configurations and connections		
Storage	Data mixing, locations		
Networking	Data mixing, locations of transit		
IT Process Management	Key to stability, availability, protection,		

and System Software	recoverability		
Facilities	Environmental vulnerabilities, facility condition and vulnerabilities		
IT Management Processes	Measured with maturity models such as Shared Assessments, ISO 27001 and COBIT		
Assurance Processes	Measured with Shared Assessments, COBIT, IT Assurance Guide and others		

- To get the most out of this table:
 - **Talk with your cloud provider.** Work with the provider to understand exactly how the service is being delivered and how it will be different from your current environment. While there are many reasons to be cautious, there are likely to be pleasant surprises, too. Cloud providers have the volume and scale to refine capabilities (process, training and tools) to a degree that would be impossible for most individual enterprises.
 - **Mark changes and similarities in each row of the table.** This will help you identify changes relating to:
 - **Information processing:** Management process and location of processing hardware, software and people
 - **Information flow:** Management process and location (physical and geographic) of data networks (dedicated and shared)
 - **Note the composition of the new environment.** Note whether it is: a) composed of standard and stable logical building blocks in a standard and stable configuration with proven management and assurance processes (Is it different from yours but reliable?); b) composed of standard logical building blocks but assembled in ways that create risks at the new connections, or is the environment not yet fully stabilized (for example, because of a high proportion of new staff); or c) composed of new building blocks, new design and/or new management and assurance processes. In conducting this analysis, be sure to evaluate at each layer of Table 3, since business processes can rest on many different underlying elements.
 - **Increase review depth according to the degree of difference.** This does not necessarily mean examining a large number of new controls. Once the environment is identified clearly, simply apply your own familiar and known controls and assurance procedures. (See the Delta Cloud Controls section for more on this.)

- **Focus any new controls on new and not-yet-stabilized points of connection.** While these controls may be "new" at the level of COBIT, Shared Assessments and other level one and two controls, your organization may have used them in the past to capture or parse information. (See the Delta Cloud Controls section for more on this.)

2. Mature Relationship, Procurement and Vendor Management Processes

Most enterprise organizations have considerable knowledge and experience in the area of traditional outsourcing models, both from a process and people perspective. For the most part, this knowledge and experience can be transferred to cloud models.

In many ways, cloud procurement is similar to the acquisition of a traditional software product. In both situations, the client must define business requirements, including functional and non-functional requirements. Both are dependent on a scalable due-diligence process to assess vendor viability, including a review of vendor financial stability and the ability of the vendor to support the product adequately.

When these mature process and teams are applied to cloud, the effort, time and cost required to evaluate cloud providers can be significantly reduced.

Considerations

- Organizations should be able to use their existing mature procurement and vendor management process to start the evaluation process.
- Ensure that the business/IT evaluation team (including vendor management, process improvement, risk management, quality control, business continuity, project management, security, compliance, internal control, audit and others) has undergone sufficient training in cloud computing and knows how to effectively evaluate cloud offerings.
- Non-functional requirements need to be identified, reviewed with the vendor and incorporated into the contract.
- Information sensitivity must be determined and appropriate security requirements agreed upon to ensure the information is protected in a manner that is commensurate with the sensitivity and importance of the data contained in the cloud.
- The contract must include terms and conditions that allow the organization to conduct a periodic assessment (for performance, risk, compliance and other purposes), such as the Shared Assessments AUP or SIG, COBIT, or other standards to determine vendor compliance with organizational standards and policies.

- Consider inserting a risk, controls and preparedness addendum specifying key policies that the cloud provider must implement.

The above is not intended to be an exhaustive list of Common Cloud Control areas. Instead, it is offered as a starting point in illustrating the significant security inspection and knowledge overlap between cloud and traditional outsourcing models. These overlaps may be leveraged to significantly reduce the time, effort and cost involved in evaluating a cloud environment.

Delta Cloud Controls

Once an organization completes an evaluation of Common Cloud Controls, the Delta Cloud Controls should be next. These control areas have the highest significance and risk in the cloud, and less industry knowledge exists to evaluate them in the cloud model.

For the business leader, these controls have substantive implications in the decision to use a cloud service provider, and should be included in conversations with management in evaluating cloud service proposals. In the Common Cloud Controls section, commonly used controls were applied to the cloud environment in new ways. With Delta Cloud Controls, new control areas are required to address the use of new technologies, significantly new service models, or nuances in how these controls apply to cloud.

Delta Cloud Control areas are divided into twelve categories:

1. [Multi-Tenant Platforms](#)
2. [Multi-Client Prioritization](#)
3. [Agile Delivery](#)
4. [Virtualization](#)
5. [Data Location, Cloud Layers and Cloud Providers](#)
6. [Cloud Management: Roles and Division of Responsibilities](#)
7. [Contracts, Data Privacy and Jurisdictional Issues](#)
8. [Identity and Log Management](#)
9. [Web Application Security](#)
10. [Cloud Vendor Interdependence and Governance](#)
11. [Data Retention, Management, Recovery and Destruction Cycles](#)
12. [E-Discovery and Forensics](#)

As with the Common Cloud Controls, the twelve Delta Cloud Control areas are intended not as an exhaustive list, but rather as a means of highlighting the primary areas of significance between cloud and traditional hosting environments.

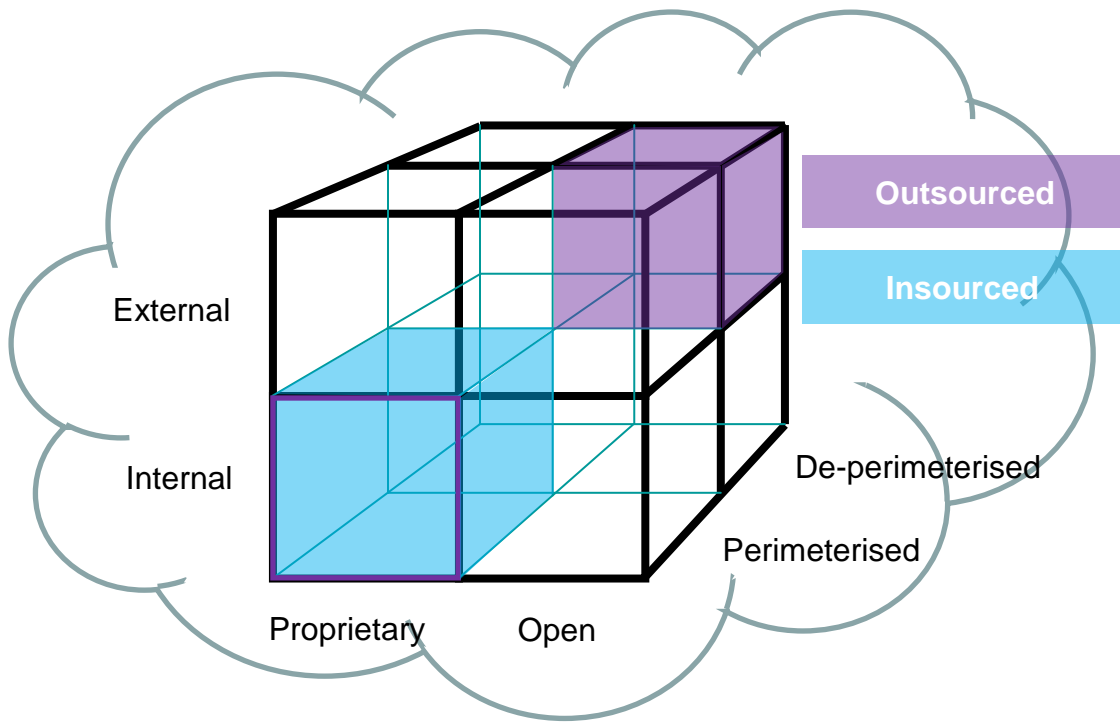
1. Multi-Tenant Platforms

One of the fundamental characteristics of cloud environments is the shared infrastructure upon which the services run. Hundreds or even thousands of clients may be using the same physical fabric at any given time. Data typically transverses and often resides on the same physical infrastructure, which creates obvious data-separation and data-leakage concerns. Today's standard industry audit controls focus primarily on the physical and logical segmentation of servers, lacking depth in inspecting the key areas of data segmentation and separation. These need to be incorporated into risk, security and audit programs, so that the data segmentation and separation controls required by cloud can be evaluated.

Considerations

- Document the data-segmentation and separation controls at each of the four main layers: (1) network, (2) physical, (3) system and (4) application.
- Evaluate each of the above controls at each layer, as well as the number and type of controls at each layer. For example, cloud data separation controls are typically weaker at the physical layer (as there is often no physical separation), requiring controls on the other three layers to be far stronger.
- Pay particular attention to the application controls, since this is the layer where the majority of critical cloud controls will reside. A cloud solution that appears to have few or weaker controls at this layer in relation to network/physical/system could be cause for concern.
- Request the details of the number, skill set and strength of the cloud application security team. With cloud, critical security controls have moved up the stack from the network and systems layers to the application layers. The provider must be able to demonstrate that it has the necessary application security skill set in-house to protect client data.
- Ascertain whether client data will be encrypted at storage and in network transmissions.
- Determine whether each client is provided with a unique encryption key or encryption keys are shared. Unique client keys are a strong control that can render co-mingled data unreadable in the database by another client. This unique encryption key control helps protect data from being readable in the event that it is inadvertently leaked from one client to another, as the other client will not have access to the decryption key to view the leaked data.
- Investigate whether software or hardware keys are used and if they meet any industry standards, for example, FIPS 140 2-3.
- Investigate whether and how the application provides service and data segmentation among clients. The cloud provider may be able to demonstrate that client data is meta-tagged; see Figure 1 below.

Figure 2: Cloud Cube Model⁵



- Evaluate how the permissioning model prevents client A from seeing client B's data.
- Request permission to carry out a penetration test of the cloud platform.⁶ Look for characteristics in the page or site that uniquely identify the client site, for example, the URL may read "Site ID=1." Modify these parameters (for example, change the URL to read "ID-2") to see if you can access another client's site or data. If you can, they can just as easily see yours. (This test is successful in a surprising number of instances due to weak application data segmentation.) Trying this in a test environment helps avoid the risk of inadvertently viewing other clients' confidential data.

⁵ Used with permission from the Jericho Forum. See www.opengroup.org/jericho/cloud_cube_model_v1.0.pdf

⁶ Ensure that this penetration test is carried out on a non-production environment with test data to avoid any risk of exposing other clients' data on the cloud platform.

2. Multi-Client Prioritization

In a traditional outsourcing environment, servers are often dedicated to specific clients. Clients have a high degree of control, with requested changes — for example, adding a new feature, changing a landing page or changing the logging level — typically affecting only that client.

With cloud's shared servers and infrastructure, one client's changes can have an adverse impact on the other clients sharing the infrastructure. For this reason, cloud providers are naturally cautious about making specific changes or customization requests for individual clients. The result is a shift in the fundamental "one-to-one" client/provider relationship to a "one-to-many" model, in which there is one provider and many clients to consider for each change, however minute.

Considerations

- Evaluate whether a reduction in the level of control and turnaround time for unique changes is acceptable to your business. This as a good litmus test for making the decision to move to a cloud environment. (Of course, the benefits of cloud should be weighed here too.)
- Evaluate how often unique change requests are likely to be made. Daily? Weekly? Monthly? Consider whether the cloud provider could realistically meet those requirements and the associated cost.
- Create a list of expected business requests, from adding a new feature to fixing a bug to shutting down the site in the event of a compromise. Ensure that agreements as to which changes are permitted and the associated timelines and costs are included in service-level agreements (SLAs).

3. Agile Delivery

One of the foundations of cloud is its agile nature, which is inherent in its roots in innovation and rapid change. In IT, "agile development" refers to a group of software development methodologies that are based on iterative development. Requirements and solutions are quickly shaped through collaboration among cross-functional teams.

Cloud product delivery cycles (inception to delivery) often occur within days or weeks instead of the annual or semiannual major releases typical of more traditional environments. This reduced delivery time means less time to complete a risk evaluation of operational stability, availability, protection/security and recovery, as well as less time for deployment and release management. For this reason, security programs that examine long release cycles are of little use in cloud environments. For example, if a provider is completing two-week delivery releases and it has ten engineering agile teams that each release ten features per software release (two weeks), 100 product features will be delivered every two weeks.

This mass volume of feature changes demands thorough risk evaluation. Unless the risk management teams (including security, continuity, recovery, change, release, facilities and all other risk areas) can move at the same speed or faster than the development teams, security and risk assessments will quickly fall behind.

When this happens, business risk grows. In the scenario above, the risk evaluation and response teams must increase their work or the change rate must decrease. The risk manager must make the situation clear to executives when reporting the risk, including the level of risk that falls outside acceptable limits and the company's ability to fully understand and respond to the risk. Businesses typically will not slow down to facilitate a slow security or risk management evaluation process.

Considerations

- Decide whether it will be acceptable to receive continuous (iterative or "drip") releases.
- Request detailed information on how the provider ensures agile risk management, including all elements of risk management (not only release or security). Risk management capabilities can degrade quickly in a fast-paced environment where there is less time to inspect and evaluate the risks presented by changes.
- Optimize the risk management processes, tools and service levels to allow for rapid and meaningful risk and controls evaluation for iterative and agile projects.
- Determine whether the cloud provider uses manual or automated controls checking, and how often it is completed. The answer can help determine whether the cloud provider's risk control checks are appropriate for cloud's rapid release cycles. For example, a cloud provider that completes a manual code review monthly with a biweekly release cycle would be a red flag. Daily automated code reviews that are rapid and scalable would indicate a better controls evaluation program.

4. Virtualization

Cloud clients often share a common physical infrastructure in which one client's data is stored, processed and transmitted on the same shared physical fabric (such as RAM or a hard disk) as other clients' data. In cloud computing, the majority of logical separation controls are not physical (i.e., separate servers). Instead, separation is enforced through logical system and application controls designed to help ensure data segmentation and integrity across the platform. One common mechanism for providing this separation of data and services is "virtualization." Virtualization is the creation of a virtual (rather than actual) version of something, such as an operating system, server, storage device or network resource. When referring to data in transmission, the notion of point-to-point no longer applies. Virtualization represents a new paradigm: multi-point to multi-point in many different physical locations.

Virtualization allows organizations to run tens (or even hundreds) of virtual operating systems on the same physical server. The result is tremendous efficiency of scale. However, virtualization can introduce risks, such as data scoping and greater difficulty in tracking and protecting data.

Virtualization's ability to allow companies to create snapshots of the environment and data can present a security issue, since images and snapshots may contain sensitive data, such as passwords and personal data. There are typically far fewer controls in place to prevent the copying of a virtual image or snapshot than there are to prevent the copying of server data to a backup tape. Due to the apparent lack of controls, virtual image snapshots are often copied to insecure locations, such as administrator desktops. Numerous unauthorized and unprotected copies can exist, increasing the likelihood of client data exposure.

While virtualization has been available for years, it is only with cloud computing that it has seen widespread use. Most companies are not nearly as knowledgeable about protecting and auditing the security of virtual environments as they are in protecting traditional systems, such as routers and servers.

Considerations

- Request copies of the cloud provider's virtualization-hardening guides and policies, and complete a gap assessment against industry controls. The National Institute of Standards and Technology's Guide to Security for Full Virtualization Technologies⁷ provides a good starting point.
- Confirm that the cloud provider has the controls in place to ensure that only authorized snapshots are taken, and that these snapshots' level of classification and storage location are commensurate in strength with the production virtualization environment.
- Review in detail the controls in place around the hypervisor as it manages the virtual environments. Who has administrative access to it? What kind of logging is enabled? Is the hypervisor physical server or network separate from the general system?

5. Data Location, Cloud Layers and Cloud Providers

Cloud providers should be able to identify the specific location of client data. Historically, when cloud providers were asked to explain exactly where client data was located, they tended to respond with the ambiguous statement that "it's in the cloud." This is no longer an acceptable answer.

⁷ See <http://csrc.nist.gov/publications/drafts/800-125/Draft-SP800-125.pdf>

While most of today's cloud providers will offer detailed information about where data is stored and the cloud layers upon which the data sits, some still cannot or will not answer more detailed questions. These providers tend to be multinational organizations operating across dozens of data centers in numerous countries; the distributed nature of their technology and/or architecture may make it difficult for them to provide a clear answer. These situations notwithstanding, cloud providers should be able to tell their clients exactly where their data is, and which of the provider's vendors have access to it. If this information isn't available, it's up to the client to decide whether their risk tolerances will allow them to use the provider.

Considerations

- Request the locations where client data will be stored, processed, accessed or transmitted, including country and system types, and incorporating network and data diagrams.
- Understand who has access to your data. Because of cloud's portability and low cost of entry, many cloud providers use and operate on other cloud providers' SaaS, PaaS and IaaS platforms. Ask your provider to list all of its vendors — in particular any cloud vendors — that will store, process, transmit or have access to company data.
- Make cloud providers contractually obligated to alert the client of changes in vendors and material infrastructure. Review how these changes and notifications would be incorporated into the vendor risk management process; notifications that are not acted upon are of little use.
- Consider requesting that the cloud vendor complete the Shared Assessments Target Data Tracker.⁸ The Target Data Tracker helps companies address three critical questions that clients should ask cloud providers prior to beginning a control evaluation:
 - What data of mine do you have?
 - Where are all the locations that my data is stored, processed, transmitted and accessed?
 - Do you share my data with other third parties?

Service provider audits often leave clients with a good sense of the strengths or weakness of the provider's controls, but without clarity about what target data is being managed, where it is located, and whether the data is sent to other dependent service providers. For example, the client may not know about the cloud provider's supporting storage or transport services, disaster recovery/business continuity locations or international contractors. Consequently, an audit may focus on the wrong data types or locations, or completely fail to evaluate all the environments where the data is stored. When the cloud provider can clearly answer the questions listed above without retreating to the statement that "it's somewhere in the cloud," then the client is far better equipped to examine the cloud provider, the locations and the data.

⁸ See www.sharedassessments.org/download/files.html

6. Cloud Management: Roles and Division of Responsibilities

If Gartner's prediction materializes,⁹ where will company assets be and who will manage them? Take, for example, the migration from Microsoft Office to Google Docs, or the transition of email systems to third-party cloud providers. In-house IT management teams are not necessarily the best candidates for managing the new cloud-based services. Internal IT support teams may excel at managing internal Exchange servers but lack the skill sets necessary to manage cloud systems.

Moving assets into the cloud may require significant realignment of client support departments. Roles and the division of responsibilities often shift significantly when an organization begins using cloud services. For this reason, organizations must clearly define roles for managing cloud vendor relationships and service delivery.

Considerations

- Evaluate how increasing your use of cloud may affect your vendor management skill set requirements. (Begin this planning early.)
- Make the most efficient use of staff responsible for internal assets if internal systems and services are moved to a cloud environment.
- Consider re-training staff on vendor management and cloud technologies. They will need to fully understand the relationship and technology aspects to be effective in managing cloud vendors.
- Define and document who is responsible for, accountable for and informed of all aspects of the service (for example, legal, vendor management, change management, business owners and problem management).
- Create a RACI (Responsible, Accountable, Consulted, Informed) matrix (Figure 3) that includes the client and the cloud provider to enhance accountability between the two organizations. A RACI matrix is especially useful in clarifying roles and responsibilities in cross-functional/departmental projects and managing processes in cloud environments. Share the matrix with the cloud provider.
- Create a communication tree and share it with internal teams and the cloud provider.

⁹ Gartner has predicted that 20 percent of businesses will own no IT assets by 2012. See "Gartner Highlights Key Predictions for IT Organizations and Users in 2010 and Beyond," January 13, 2010; <http://www.gartner.com/it/page.jsp?id=1278413>

Figure 3: RACI Matrix¹⁰

Code	Name	Project Sponsor	Business Analyst	Project Manager	Technical Architect	Applications Development
Stage A	Manage Sales					
Stage B	Assess Job					
Stage C	Initiate Project					
CD4	Security Governance (draft)	C	C	A	I	I
C10	Functional Requirements	A	R	I	C	I
C11	Business Acceptance Criteria	A	R	I	C	I
Stage D	Design Solution					

7. Contracts, Data Privacy and Jurisdictional Issues

Before moving a service out of an organization to any third party, a rigorous legal analysis and evaluation should be conducted. This is especially important if data will be stored, processed or transmitted in a foreign country.

Cloud is no different. In a cloud relationship, a number of issues stand out, including the daisy-chain (or point-to-multi-point) cloud service provider model, the co-mingling of data at the physical layer, and the often ambiguous location of client data.

A client may outsource a service, but it cannot outsource its risk and compliance obligations. Contractual relationships must be well defined, including establishing a good understanding of who the “control owner” is and the associated legal roles and responsibilities, which should be agreed on by all parties.

Considerations

- Establish who the owner of the data is and what rights the cloud provider has to the data. In nearly all cases the client should own the data, and the cloud provider should have no rights to it.
- List all locations and service providers that store, process, transmit or access client data and whether these are contractually documented.

¹⁰ Source: http://en.wikipedia.org/wiki/File:RACI_Matrix.png

- Define in the contract the countries where company data will be stored.
- Determine whether any foreign country where the data resides has a propensity to take possession of IT assets or block access to key data needed for business operations. These events could result in loss of business revenue and potential penalties for legal violations in the company's home jurisdiction.
- Investigate thoroughly any conflict in countries' data privacy and legal requirements. For example, a data privacy conflict could arise if the client and cloud provider are located in the US and the provider has multiple datacenters in the US, but also has a datacenter in Germany for disaster recovery and resilience. The US could mandate certain data be deleted (due to a US data privacy law breach) while German law may require that the data be retained (as evidence in a subsequent legal case). In this scenario, the conflict of laws between jurisdictions puts the data at risk.
- Ensure that client data only resides in one jurisdiction (where permissible) as this requirement can significantly negate jurisdictional complications. In this case, ensure that the cloud provider requests permission before it stores data outside of a specific pre-defined country
- Establish whose data privacy policy applies and how the contractual requirements will be implemented. In the majority of cases the client's data privacy policy should take precedence over the cloud provider's.
- Provide contractual assurances so that applications and data will be resilient in the event of planned or unplanned disruptions or outages, with business continuity and disaster recovery planning and backup and redundancy mechanisms in place. SLAs should define financial penalties in the event of a business disruption.
- Provide contractual assurances that define what data must be encrypted and in what state, e.g., transit or storage.
- Contractually require that the cloud provider notify the client of any breach within a specific period. It is important that (a) your company is notified of "suspected" as well as "actual" breaches; (b) the notification period is within hours (not days) of the breach; and (c) the breach notification stopwatch starts when the breach is "discovered" rather than when the investigation is completed. (An investigation can take months to complete.)
- Ensure that contractual and financial terms protect the client from a data breach by the cloud provider.

8. Identity and Log Management

Ideally, a staff member or client end user should not need an additional username and password to access data or services that are managed by the cloud provider. Similarly, having two

authentication databases — one for the client and another for the cloud provider, with a username, password, and permissions for each — is neither manageable, scalable nor secure. (Enormous effort would be required for on-boarding, completing periodic password changes, changing access rights and removing users across the two systems.)

Unified identity management is an essential component of cloud, from a business, usability and security perspective. Businesses using cloud may be presented with the challenge of integrating their existing identity management solutions with that of the cloud provider. If this integration cannot be achieved, then the client may have to allow the cloud provider permission to access its authentication environment or vice versa, neither of which is ideal from a security perspective. This disjointed method may pose risk in the form of improper or unapproved entitlements. The provider may also lack an effective mechanism for allowing the client to perform periodic user entitlement reviews required for standards or regulatory compliance.

Significant progress has been made in this area in the past three years with the advent of Identity-as-a-Service (IDaaS) providers, which provide open, federated standards such as SAML and OpenID to permit transparent user single sign on (SSO) among cloud environments.

Log management, i.e., who has access to the logs, is another management issue that can be contentious and unless agreed upon in advance. A cloud provider will rarely provide raw logs to the client when requested, as the logs may contain other clients' data. Providing logs to one client could expose other clients' data.

Considerations

- Determine whether your identity management solution can integrate with the cloud provider's and the costs associated with integration.
- If your organization does not support identity federation standards such as SAML or OpenID, consider adding this functionally now to help prevent costly individual integrations. Conducting ample due diligence on this at the start of the engagement is highly recommended; supporting multiple non-integrated authentication systems can be prohibitively expensive.
- Determine whether the cloud provider's identity management solution allows for organizational control in managing identities. (Some frameworks allow users to control their own identities.)
- Determine what protocol (SAML, ID-FF, WS-Federation, etc.) should be used for communication among identity management solutions. Solutions that use different protocols may not be able to communicate to support activities such as provisioning, access management, identity management and activity/security monitoring.

- Determine who will manage the identities. Will management be client or cloud-provider-based? If cloud-provider-based, discuss workflow considerations and SLAs with the provider.
- Evaluate whether the provider's authentication, access control, accountability and logging will satisfy your organization's regulatory and legal requirements.
- Agree on who will be responsible for adding and removing users (for example, terminations) and establish a corresponding SLA.
- Agree on the availability of entitlement lists. Will the provider allow periodic entitlement reviews?
- Evaluate how user actions and system events will be audited and monitored, and from where. If the cloud provider is supplying the solution, determine whether or not your IT organization will have access to it or the logs.
- Review the functionality and usefulness of dashboard, reports and application programming interfaces (APIs) that the cloud provider will expose to ensure that they meet client requirements. Will this provide adequate monitoring capabilities? Cloud providers will typically not expose raw log data to the client; clients usually have to rely on what the cloud provider tells them and the reports and dashboards they provide.

9. Web Application Security

Application security is important in both traditional outsourcing models and cloud computing. However, with cloud the importance of application security becomes absolutely critical. Cloud is typically an open environment, and cloud providers are exposing an increasing number of web interfaces and APIs to the Internet — far more than traditional closed on-premise solutions, significantly increasing the application attack exposure.

Cloud providers run applications, and these applications require code. In an agile model, the code changes every two weeks. (Standard software delivery releases for agile cloud is two weeks.) Unless agile security software development processes, code-review and penetration-test programs are in place and moving at the same pace as the two-week software delivery releases, vulnerabilities will increase. Significant effort is required to build and maintain an adequate level of application experience and maturity to achieve true cloud security.

For this reason, cloud providers particularly must excel in application security, and must be able to demonstrate that they have the application security team, knowledge and processes to protect client data in the cloud.

Considerations

- Evaluate the depth of the provider's application security team. Are they in-house or part-time consultants? How many are on the team? What is their level of experience?
Companies should devote time to examining this area, since a cloud provider may have in-depth application security policies and processes that quickly become "shelfware" unless a strong application security team is in place that can move at the same speed (or more quickly) than cloud and its software development cycles.
- Evaluate whether the cloud provider uses application-layer firewalls. In a cloud environment, application firewalls are essential. Since the application has broader exposure, the attack space increases and standard network firewalls and access controls are not sufficient to protect the applications.
- Review the sanity checklist pre- and post-deployment to assure the cloud provider utilizes sufficient application security inspection. Basic checks should include (but are not limited to): having a secure code review before shipping to production; that no password appears in clear-text format; that appropriate permissions have been made for the source code; and that no password appears in database connection strings.
- Ensure sufficient hardening procedures exist for web and application servers. Often confused with OS requirements, hardening procedures lock down the web and application servers. This is essential; an application that is not hardened can quickly be compromised in the cloud.
- Review the security-development programs, code-review cycles, and penetration-test cycles to determine whether they can keep up and are sufficient to secure the code in the software delivery releases.
- Ensure that an application security remediation program is defined, and that it includes fixing the vulnerabilities that are found based on priority. All vulnerabilities should be prioritized and must be fixed and patched within SLAs agreed upon by the client and the cloud provider.
- Ensure that application security is integrated at every phase of software development life cycle (SDLC). These phases should include (but are not limited to):
 - During the planning and requirement phase, the application security team should review the business requirement and define security requirements of the application. The objective of this phase from security perspective is to make sure that the security requirements are an integral part of business requirements.
 - During the design phase, architecture and functional design must be reviewed from a security perspective. The goal of the design review is to make sure security

requirements are designed as defined in the requirements phase. This is usually a manual process that is specific to the application and the client's internal processes.

- During the coding phase, the security team must perform security code review. The objective of security code review is to find vulnerabilities in the code through a white box testing methodology. There are a number of commercial and open source tools available depending on the coding language and the application security needs.
- During the testing phase, the security teams must complete a penetration test. The objective of penetration testing is to find security vulnerabilities in the code through black box testing. A combination of black box and white box testing can also be used for a more thorough security test. There are several open-source tools and commercial tools available in the market that can automate the penetration-testing process.
- During the maintenance phase, the security team continues reviewing and testing the changes and bug fixes made to the application.
- To ensure that data is sufficiently protected at time of de-commissioning or the end of the contract, for example, application security controls are not removed too early in the decommissioning process, placing the data at risk.

10. Cloud Vendor Interdependence and Governance

A cloud provider may choose one vendor for the hardware platform, a second for the software platform, a third for backups, and a fourth for disaster recovery. This "multi-vendor" environment is becoming increasingly common in organizations and in particular with cloud providers. Integrating the various back-end and front-end components of a managed solution is no small task, regardless of whether it is in the cloud or on-premise.

Cloud interdependence can result in a lack of clarity about where client data resides, what controls apply, and most importantly, who is legally responsible for protecting the data.

Considerations

- Keep it under one service contract. Make sure the cloud provider's (and its vendors') support service delivery model for your organization is covered under one maintenance contract.
- Require contract provisions that track the data's physical location(s). Contract arrangements with different vendors are likely to become problematic if customers are required to engage each cloud provider separately. The challenge is to ensure the benefit of deploying a cloud solution is not outweighed by the complexity of doing business in the cloud. The cloud provider should provide a single point of contact, a single contract and a

single point of accountability to look to when things go wrong. Be wary of cloud providers that blame their vendors for an incident, and vice versa, when the client is looking for, and needs accountability.

- Ensure consistent quality. When different vendors have varying service levels and poorly defined incident escalation points, problems can easily go unnoticed and unreported. Make sure a documented holistic mechanism exists to ensure service levels are met and issues are resolved.
- Understand the processes interrelationships between the cloud provider and its vendors. When data is dispersed across vendors, it can be hard to get a “big picture” view of the provider’s IT operation, the vendor interrelationships and how they interact with each other. It is critical to understand all the vendors’ roles that are in scope so that your security evaluation process is not fragmented or isolated, or to too small a scope for the operation.
- Follow processes that ensure holistic information and IT security. Improve procedures, formalizing processes such as escalation points and SLAs. Ensuring that everyone is contractually adhering to procedures and policies helps eliminate ineffective and redundant procedures, resulting in streamlined operations.
- Make monitoring the cloud provider’s environment part of your regular security forum responsibilities. It's still your data and your processes. Cloud provider monitoring and management should be part of your overall information security management system (ISMS) with defined and monitored metrics. Using the Shared Assessments AUP, SIG or a SAS 70 Type II as a baseline, identify the gaps and create a road map for improvement to drive your cloud program. This will help ensure transparency and prevent fragmentation, allowing for a proactive (rather than reactive) evaluation environment.
- Ensure that sufficient governance and risk management oversight exists within the client's organization to be able to effectively manage and monitor the relationship with the cloud provider and its vendors. Guard against an "out of sight, out of mind" mentality: it's still your data and your service even if it is hosted or directly managed by the cloud provider.

11. Data Retention, Management, Recovery and Destruction Cycles

The retention, management and destruction of data in any outsourcing model is critical, whether in a traditional outsourcing model or cloud. When a client moves any services or data into the cloud they must have confidence that the provider will protect its data from loss or compromise.

Data protection controls need to apply to all phases of the data cycle, from copying data to the cloud provider, to day-to-day management, to removing or destroying the data at the end of the contract. Cloud providers must be able to clearly demonstrate to the client that they are capable of carrying out these critical data management processes effectively in the cloud.

Cloud is "always-on" by nature: clients can get to the data from wherever they want, whenever they want. For this to occur, the cloud provider needs to ensure it has superior operational, change management, disaster recovery and business continuity plans and controls.

Considerations

- All data should hold an information-classification tag to ensure that the commensurate data-management controls can be applied to the data based on its business function and data classification.
- Provisions should be made for data retention and deletion, with the contract following the client's destruction policy. Particular care should be taken, with discussions of how cloud data destruction will happen without impacting other clients' data, including shared databases and offline storage; file or backup (tape) storage may contain multiple clients' data.
- Definitions of certificates of destruction, if any, should include 1) what is in scope for destruction, 2) whether logical or physical destruction controls will be used and 3) how those controls will be applied.
- Ensure that a contractual agreement exists between the client and cloud provider for defining data retention requirements potentially bound by legal electronic discovery (also known as a "legal hold").
- Conduct proper due diligence on how client data kept on offline media, including backup tapes, will be destroyed. The cloud provider's agreement to destroy all data at the end of a contract is often far harder to execute than it might seem. Take, for example, a multi-tenant cloud provider that backs up all data onto a single backup tape (or a series of tapes). At the end of the contract the provider can easily remove the data from online databases and stores. However removal from backup tapes is far more complicated (and in some cases not feasible) since each backup tape may contain data from multiple clients. Physical destruction of the tape would destroy other clients' data along with it. Logical destruction would be costly and complicated, too: the data would have to be restored to an electronic medium, the specific client data found and deleted, and the other clients' data copied back onto the backup tape. This cumbersome process is neither feasible nor cost effective in a majority of cases.
- Ensure that change-management and incident-response procedures are in compliance with general standards, including those of the Information Technology Infrastructure Library (ITIL).

- Request the provider's service-level availability, for example, five nines up-time or only one nine up-time.
- Evaluate the cloud provider's definition of up-time, to ensure that any interruption that would impact your business is covered under this definition. A narrow scope for up-time often results in over-hyped numbers that do not match cloud provider availability.
- Ensure that financial penalties apply for not meeting up-time SLAs. (SLAs are largely ineffective without these.)
- Decide who will be responsible for monitoring up-time and releasing the up-time statistics: the client, the cloud provider or both.
- Review disaster recovery and business continuity plans and procedures and ensure that they match your company's business requirements.
- Ensure that the disaster recovery and business continuity locations have adequate levels of security for the provider's primary production environment. Cloud provider fail-over sites may be an on-demand solution with significantly reduced security controls, resulting in client data being put at risk at the time of fail-over.

12. E-Discovery and Forensics

During e-discovery, organizations may receive subpoenas requesting relevant data to be retained and shared with a third party as part of a legal action. Even if the relevant data is stored in-house by the client, e-discovery can still be challenging, costly and time consuming.

Particular challenges arise when the relevant data is being stored by a cloud provider.

Considerations

- Discuss possible e-discovery scenarios with your internal legal, IT and business teams to determine which are most relevant. Then request the cloud provider's e-discovery and forensics procedures and compare these with the scenarios.
- Determine whether providing copies of the data is sufficient, or whether "the original data on the original hard disk" is required. This can be a significant point of contention if the original hard disk or back-up tape is required, as client data may be co-mingled with other client data. Providing the disk or back-up tape would expose other clients' data, so the cloud provider refuses to do so, because of its own data privacy and legal requirements. Conduct an analysis of potential e-discovery scenarios and discuss these issues with the cloud provider.
- Ensure that an e-discovery forensics capability and associated costs and timelines are detailed in the contract.

The above is not intended an exhaustive list of Delta Cloud Controls. Instead, these areas and their corresponding considerations may be used to highlight critical risk management areas in cloud environments. Enterprise organizations may use these guidelines as a starting point for building effective risk management and controls evaluation programs that include cloud computing.

Cloud Computing Case Study

Project Overview

Business leaders at a large, US-based enterprise (referred to here as “Ace Corporation”) needed a way to exchange information with the company’s trading partners. Initially, the business leaders used in-house applications, but each one proved inadequate in scalability and functionality. So, the business department decided it would try something new: web-based collaborative software running on SaaS.

A vice president in the business department contacted Ace Corporation’s technology sourcing department to discuss her department’s business needs. (The sourcing department was a busy place, since Ace worked with tens of thousands of vendors.) She outlined the kind of technology she thought would be required and learned about the steps that would have to be taken in order to source the collaborative software her department wanted. Her department had sourced software products in the past, but they had not yet used any cloud computing services.

Following the call to the technology sourcing team, the vice president and her colleagues documented their functional and non-functional business requirements. This initial scoping process was similar to sourcing any functional software evaluation. They then met several times with the centralized technology sourcing team, whose members educated the business team about software industry trends, including what they should consider in evaluating and selecting a SaaS solution.

Similarities to Traditional Procurement

The SaaS acquisition process is similar in many ways to traditional software product acquisition. In both cases, the client must define business requirements, including functional and non-functional requirements. Both are dependent on a scalable due diligence process to assess vendor viability, including reviewing the vendor’s financial stability and its ability to adequately support the product. Also, packaged software and SaaS solutions both require careful evaluation of software functionality and negotiation of cost.

Sourcing Process

Ace Corporation’s technology sourcing managers recognized the differences between sourcing traditional software applications and SaaS services. Together, they worked to enhance the company’s organizational technology sourcing process to support the effective evaluation, negotiation and selection of SaaS. Inherent in their process was a recognition of the risks involved in the operational services that are offered as a part of any SaaS solution.

To support the execution of this enhanced process, Ace’s managers forged a strong working relationship with key partners in the company’s IT departments that would be instrumental in

executing an effective SaaS sourcing process. These teams worked together to help business teams determine SaaS-specific non-functional requirements (operational and security) that would need to be communicated, negotiated and ultimately documented in the resulting contract with a SaaS vendor.

Ace Corporation's software sourcing lifecycle consisted of these documented phases:

- Procurement Request
- Sourcing Strategy
- Negotiate and Select
- Contract Initiation
- Project Closure

SaaS-specific enhancements were made during the procurement request, sourcing strategy and negotiate and select phases of the software sourcing lifecycle.

Procurement Request

Working with the technology sourcing team, the business department decided it would consider both onsite deployed software and SaaS solutions. The technology sourcing team engaged the information security team early on in the process to work with the business department to review and categorize the type of information that might reside in the SaaS solution. The business department needed to be able to store the confidential information within the SaaS platform — a requirement that had direct bearing on the security controls that would be used to evaluate the proposed SaaS solution. The business team also worked with an IT infrastructure team to evaluate SaaS-specific non-functional requirements as well as the desired service levels — criteria that would not be necessary to consider with in-house traditional software products.

The information security team liaised with the business department to determine the nature of the information that would be contained in the selected application. The project team identified the information that users would enter into the system and established a preliminary information classification system. The categories included “highly confidential (regulated),” “confidential,” “internal” and “public.”

Sourcing Strategy

The technology sourcing team used the information documented in the procurement request phase to create an RFP for the chosen vendors. Team members also established a method for evaluating SaaS and onsite deployed software in anticipation of the varied SaaS and in-house service models offered in the vendor proposals.

The project functional and non-functional requirements were communicated to the vendors that were selected to participate in the evaluation. Requirements relevant only to SaaS offerings were

identified, including security and infrastructure non-functional requirements and SaaS-specific service levels.

Negotiate and Select

The technology sourcing team managed the RFP process. Their work included receiving numerous vendor proposals and presentations and attending meetings with prospective vendors.

The team executed the steps below for each SaaS vendor being considered:

- Negotiation of requirements and service-level gaps
- Documentation of agreed upon product or service changes with delivery timeframes and penalties for missed dates
- Client acceptance of risk related to gaps that cannot be closed
- Definition and negotiation of service levels, including critical service levels and key measurements, such as penalties (needed definition of reporting requirements to facilitate vendor adherence to critical service levels)
- Finalizing terms of the security agreement, which were consistent with previously identified SaaS-specific security requirements communicated to vendors
- Mitigation for missed critical service levels specified in the contract
- Agreement on use of defined templates for penalties related to missed critical service levels
- Negotiation of costs, including caps on price increases
- Finalizing vendor risk profile and confirming the data risk tier for the selected SaaS vendor

In partnership with the business team, the sourcing team eventually decided on a SaaS vendor. The vendor that was selected was considered to be best positioned to meet the project requirements because it offered:

- A SaaS on-demand cost model
- Frequent feature additions and enhancements via semi-monthly software releases
- A scalable platform that allowed for future expected growth
- Open web communication interfaces that provided seamless information exchanges with Ace Corporation's trading partners
- Third-party ownership and responsibility for the direct management of the application and service
- Adherence to all of Ace Corporation's requirements and service-level requirements

The sourcing team negotiated a contract and service levels with this SaaS vendor. Today, Ace Corporation's business department leaders are able to exchange information with their key trading partners via a scalable and highly functional SaaS-based collaborative software service.

Glossary

Acceptable Use Policy: Part of the information security framework that defines what users are and are not allowed to do with the IT systems of the organization. It should contain a subset of the information security policy and refer users to the full security policy when relevant. It should also clearly define the sanctions applied if a user violates the policy.

Acknowledgement of Acceptable Use: A written attestation from a user of an information system indicating the user's acceptance and willingness to comply with the relevant information systems control policies.

Asset Classification: The category or type assigned to an asset, derived from the asset classification policy. Asset classifications frequently vary from company to company.

Business Continuity Plan (BCP): A process that defines exactly how, for which applications, and for how long a business plans to continue functioning after a disruptive event. The business continuity plan is usually an overarching plan that includes both operational and technology-related tasks.

Business Impact Analysis (BIA): This term is applicable across Technology Risk Management, in both information security and business continuity planning domains. An impact analysis results in the differentiation between critical and non-critical business functions. A function may be considered critical if there is an unacceptable impact to stakeholders from damage to the function. The perception of the acceptability of disruption may be modified by the cost of establishing and maintaining appropriate business or technical recovery solutions. A function may also be considered critical if dictated by law.

Business Process: An end-to-end service made available to internal or external parties that usually corresponds to standard service products that the service provider offers to clients.

Cloud Computing (NIST Definition): A model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. This cloud model promotes availability and is composed of five essential characteristics, three service models, and four deployment models.¹¹

Communication Tree: A document stating when, how and to whom communication must be made if an unexpected event occurs.

Controls: Safeguards or countermeasures to avoid, counteract or minimize risks. Controls may prevent risk from occurring, detect that risk has occurred or limit the negative impact of a risk once it has occurred.

¹¹ More information at <http://csrc.nist.gov/groups/SNS/cloud-computing/>

Facility: A structure or building, or multiple structures or buildings, in which operations are conducted for the services provided. These operations include handling, processing and storage of information, data or systems, as well as personnel that support the operations.

Risk Governance in IT Context: The role of Risk Governance in the IT context is to assure that investments in IT generate business value and mitigate the risks, for example security and privacy risks, that are associated with IT. Risk Governance is aided by the use of common control frameworks and standards and usually involves various risk management activities such as threat/risk assessments and response/remediation planning.

Infrastructure as a Service (IaaS): A provision model in which an organization outsources the equipment used to support operations, including storage, hardware, servers and networking components. The service provider owns the equipment and is responsible for housing, running and maintaining it.

Non-Public Information: Any personally identifiable or company proprietary information that is not publicly available. Non-Public Information includes but is not limited to: certain company proprietary information, such as internal policies and memoranda; and personal information such as an individual's name, address or telephone number. It also includes information requiring higher levels of protection according to the company's security policy, such as company proprietary trade secrets or personal information that bundles an individual's name, address or telephone number with a Social Security number, driver's license number, account number, credit or debit card number, personal identification number, health information, religious opinions or a user ID or password.

Non-Public Personal Information (NPPI): Any personally identifiable financial information that is not publicly available. Non-Public Personal Information includes but is not limited to name, address, city, state, Zip code, telephone number, Social Security number, credit card number, bank account number and financial history.

Notice Consent Language: Any Data Subject consent language in a Privacy Notice to be accepted by a Data Subject (expressly or by implication). The language may relate to consent to the entire Privacy Notice or to particular uses of the Target Privacy Data where a Data Subject's non-consent to this use of the Target Privacy Data results in a Data Subject rejecting the Privacy Notice. Examples of uses include cross-border transfer of Target Privacy Data, special use of the Target Privacy Data, or special local regulatory requirements.

Outsourcing: The contracting out of a business or technology function — commonly one previously performed in-house — to a third-party provider.

Owner: An individual or entity that has approved management responsibility for controlling the production, development, maintenance, use and security of the assets. Ownership is not an indication of property rights to the asset.

Ownership: A formally assigned responsibility for a given asset.

Permission: Any Data Subject permission (opt in or opt out) required to use or share Target Privacy Data that can be easily switched on and off, including for the following purposes: marketing; affiliate sharing; product use; promotions; newsletters; tailoring services to the Data Subject's particular requirements; behavioral and purchasing patterns; social networking; and professional networking, excluding Notice Consent Language.

Platform as a Service (PaaS): A model whereby the customer rents hardware, operating systems, storage and network capacity over the Internet for use in running existing applications or developing and testing new ones.

Privacy Applicable Law: Relevant laws, enactments, regulations, binding industry codes, regulatory permits and licenses that are in effect and address the protection, handling and privacy of Target Privacy Data, selected as being in scope by the service provider or client.

Privacy Inventory Flow: The most current Target Privacy Data inventory/list and flow by Data Subject Category that has been approved by management of the organization. A privacy inventory flow identifies the ownership of the Target Privacy Data, its sources, collection methods, storage locations, uses (by whom, where and for what purpose), sharing within the service provider and among its third parties, trans-border flows and adequacy mechanisms chosen to ensure the protection of such Target Privacy Data, security, retention and deletion schedules and mechanisms.

Privacy Notice: Notice given to Data Subjects on the collection, use, storage, sharing, transfer, retention and destruction of their Target Privacy Data in accordance with Privacy Applicable Law and organization policy.

Privacy Policy: An organization's internal policy adopted for the lifecycle of the Target Privacy Data.

Protected Target Data: Target Data or any other data that requires a higher level of protection or special treatment due to its sensitivity under: Security Applicable Law; company security policy; and/or as identified in the Scope Definition of Protected Target Data of the Shared Assessments Standardized Information Gathering Questionnaire (SIG) and Agreed Upon Procedures (AUP). This may include: Target Data, such as name, address or telephone number in conjunction with Social Security number, driver's license number, account number, credit or debit card number, personal identification number, user ID or password; an individual's health information; company trade secrets or certain confidential information. For data that falls under the definitions of both Target Data and Protected Target Data, (for example, credit card details).

Protected Target Privacy Data: Any Target Privacy Data that requires a higher level of protection or special treatment under Privacy Applicable Law due to its sensitivity, e.g., encryption. This includes EU "sensitive personal data" (health, religion, criminal records, trade union membership, sexual orientation and race). In the US, Protected Target Privacy Data includes name, address or telephone number in conjunction with Social Security number, driver's license number, account number, credit or debit card number, personal identification number, or user ID or password

Publicly Accessible: In networking terms, able to accept a connection originating from the public domain, e.g., the Internet.

Remote Access: The ability to log in to a network from a distant location.

Residual Risk Rating Scoring Method: A calculation of the risk that remains after security controls have been applied.

RACI (Responsible, Accountable, Consulted, Informed) Matrix: A common model used to define roles and responsibilities for members of cross-functional initiatives. The matrix allows members to easily understand which groups are responsible and accountable for activities and which must be consulted or informed.

Risk Management: Management of business outcomes through consideration of threats, exposures and vulnerabilities that might put objectives at risk. Some common methods used to help manage risk include assessing probability and impact of threat, assessing inherent and residual risk, and then prioritizing treatment of risk according to objectives defined through a business impact analysis and the organization's risk appetite.

Risk Prioritization Scoring Method: A systematic approach that quantifies risk in terms of loss potential, then sequences individual risks to determine the order in which compensating controls should be implemented.

Secure Perimeter: A space fully enclosed by walls that surround the immediate perimeter and that extend from floor to ceiling (beyond raised floors and ceilings), which is contained, and whose points of entry are secured.

Secure Workspace: An environment from where people work from their desks with the purpose of accessing, editing or inputting Target Data on a computer, telephone or physical media, e.g., a business process outsourcing or call center environment.

Secure Workspace Perimeter: A space fully enclosed by walls that surround the workspace that is contained and whose points of entry and exit are secured.

Security Applicable Law: Applicable laws, enactments, regulations, binding industry codes, regulatory permits and licenses that are in effect that address the protection, handling and security of Target Data and Protected Target Data and that are determined to be in scope by the service provider or client at the scoping of the engagement.

Security Policy: A published document or set of documents defining requirements for one or more aspects of information security.

Sensitive Information: Also known as "Target Data," any customer data stored at the service provider's facility. This data may be stored in the form of physical media, digital media or any other storage medium.

Server: A computer that makes services — such as access to data files, programs and peripheral devices — available to workstations on a network.

Service Provider: An organization that provides outsourced services, such as data processing, business operations, applications, systems or staffing.

Software as a Service (SaaS): A model of software deployment whereby a provider licenses an application to customers for use as a service on demand.

Target Data: A client’s Non-Public Personal Information (NPPI), Protected Health Information (PHI), Personal Information (PI) or Non-Public Information that is stored, transmitted or processed by the service provider. Target Data may also include any data selected as being in scope by the Service Provider or Client at the scoping of the engagement. Any reference to Target Data includes Protected Target Data, where applicable.

Target Privacy Data: Any information relating to a Data Subject, who can be identified, directly or indirectly, by that information and in particular by reference to an identification number or to one or more factors specific to his or her physical, physiological, mental, economic, cultural or social identity. Examples of Target Privacy Data include name, address, telephone number or email address. Target Privacy Data may exist in any media or format. Any reference to Target Privacy Data includes Protected Target Privacy Data, where applicable.

Target System: Computer hardware and software in scope for the engagement that contains Target Data.

Third Party: All entities or persons that work on behalf of the organization but are not its employees, including consultants, contingent workers, clients, business partners, service providers, subcontractors, vendors, suppliers, affiliates and any other person or entity that accesses Target Privacy Data.

Threat Impact Calculation Method: A systematic method of determining the loss potential of a particular threat that is based on the value of assets affected.

Threat Probability Calculation Method: A systematic method of determining the potential for a particular threat to occur, based on the likelihood of the occurrence collected from internal staff, past records and official security records.

Threats x Vulnerability x Asset Value = Total Risk (Threats x Vulnerability x Asset Value) x Controls Gap = Residual Risk

Virtual Private Network (VPN): A communication tunnel running through a shared network, such as the Internet, which uses encryption and other security mechanisms to ensure the data cannot be intercepted and that the data senders and receivers are authenticated.

Appendix:

Additional Cloud Computing Initiatives

The Shared Assessments AUP (formal audit procedures) and SIG (risk questionnaire) offer detailed standards for cloud use related to security, privacy and business continuity. As with any standard, though, organizations should never rely on one source alone. A number of consortia and other organizations have worked to define the controls necessary for cloud's widespread acceptance in the enterprise. The summaries below provide an overview of the initiatives of:

- Cloud Security Alliance
- Commission of the European Communities
- European Network and Information Security Agency
- ISACA
- National Institute of Standards and Technology
- The Open Group / Jericho Forum
- US Regulatory Agencies

Cloud Security Alliance

The Cloud Security Alliance (CSA) is a global not-for-profit organization with a broad geographical distribution with 50 corporate members, 12 nonprofit affiliations and more than 9,000 members, with 300 new members joining each week. The CSA supports a broad spectrum of subject matter expertise, including experts in cloud, security, legal issues, compliance and virtualization. All provide leadership and guidance for security best practices.

The group's mission is "To promote the use of best practices for providing security assurance in Cloud Computing, and provide education on the uses of Cloud Computing to help secure all other forms of computing."

The CSA leads a number of projects, including:

- **Security Guidance for Critical Areas of Focus in Cloud Computing.** The CSA's flagship research project, the Security Guide provides a broad catalog of best practices for cloud with 13 domains to address both broad governance and specific operational issues. Version 2.1 was released in December 2009¹² and Version 3 is slated for release in mid-2011.

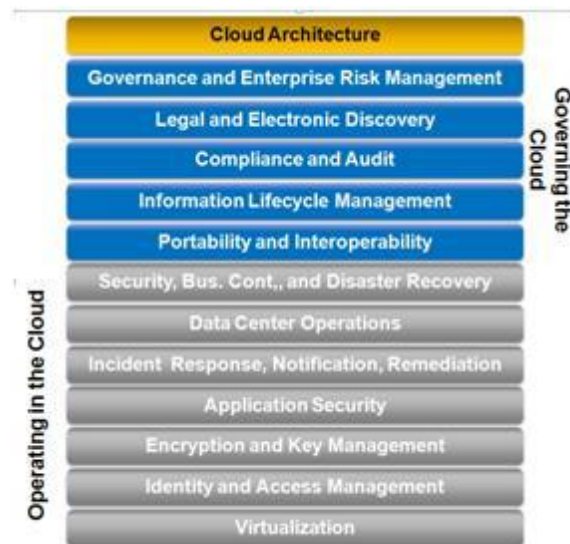
In 2010 the Security Guide group focused on translations, wiki workspaces and whitepapers, including:

- Identity & Access Management Whitepaper (April 2010)

¹² See <http://www.cloudsecurityalliance.org/guidance.html>

- Application Security Whitepaper (July 2010)

Figure 4: Cloud Architecture Diagram¹³



- **Certificate of Cloud Security Knowledge (CCSK).** The first user certification program for secure cloud computing, CCSK is designed to ensure that a broad range of professionals with responsibilities related to cloud computing have a demonstrated awareness of the security threats and best practices for securing the cloud.¹⁴
- **Controls Matrix (CM).** The CM was designed to provide fundamental security principles to guide cloud vendors and assist prospective cloud customers in assessing the overall security risk of a cloud provider. The CM provides a controls framework and offers readers a detailed understanding of security concepts and principles that align to industry-accepted security standards, regulations, and controls frameworks. Version 2 of the CM is planned for release at the November 2010 Cloud Security Congress in Orlando.
- **Consensus Assessments Initiative.** This project offers research tools and processes to perform consistent measurements of cloud providers using a lightweight "common criteria." This group will release tools and process at RSA Europe in London in October 2010 and at the Cloud Security Congress in Orlando in November.

¹³ Used with permission from the Cloud Security Alliance. See <http://www.cloudsecurityalliance.org/csaguide.pdf>

¹⁴ See press release and FAQs at http://www.cloudsecurityalliance.org/ccsk_faq.html

- **Cloud Metrics.** This initiative is building security control metrics for baseline capabilities of cloud service providers.
- **Trusted Cloud Initiative (TCI).** This project works with cloud providers to develop industry-recommended secure and interoperable identity, access and compliance management configurations and practices. In 2010, the group will develop reference models, education, certification criteria and a cloud provider self-certification toolset. The toolset will be vendor-neutral.¹⁵
- **Top Threats to Cloud Computing.** This group's goal is to help organizations make educated risk management decisions about cloud adoption. Its work is intended to complement that of the Security Guidance for Critical Areas in Cloud Computing (bullet one above). The group publishes a Top Threats Report and updates it twice yearly to reflect expert consensus on probable threats¹⁶

The CSA's work serves as a foundation for other compliance-related research projects.

Figure 5: Security Integration¹⁷

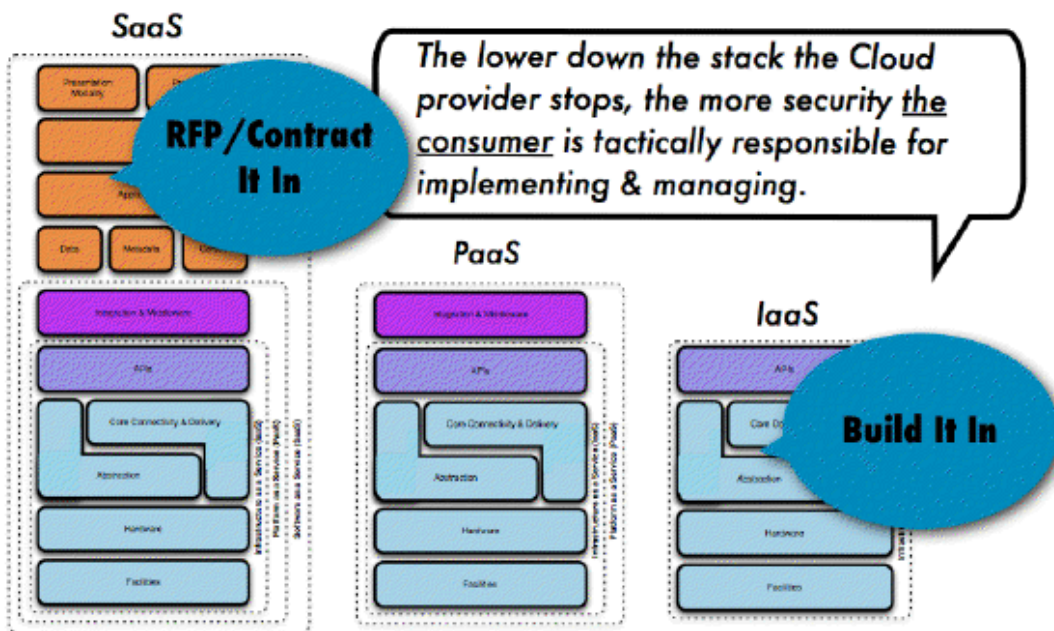


Figure 7 - How Security Gets Integrated

¹⁵ For more information, see <http://www.trusted-cloud.com/>

¹⁶ See <http://www.cloudsecurityalliance.org/topthreats/csathreats.v1.0.pdf>

¹⁷ Used with permission from the Cloud Security Alliance. See <http://www.cloudsecurityalliance.org/csaguide.pdf>

Commission of the European Communities

In Europe, the Directorate-General Information Society published *The Future of Cloud Computing: Opportunities for European Cloud Computing Beyond 2010*.¹ This document outlines the European Commission's position on the subject of cloud computing. Keeping in line with the Directorate-General's mission of supporting innovation and developing information and communication technologies (ICTs), the document examines different aspects of "opportunities" presented by cloud computing and how Europe will leverage them. The idea is to create a base for a regulatory environment in which cloud computing services can be developed, not only for the benefit of European citizens, but in order to join international efforts to address legal and technical issues aimed at enabling cloud computing on a "global" scope.

The Directorate-General's recommendations are based on cloud computing's technology and the non-functional and economic aspects of cloud. The publication also details "gaps and open areas," which they see in the different available cloud computing implementations, and what they have to offer. Legal and regulatory environments are structured very differently across Europe as well as globally. This creates a difficult environment in which to foster international or "global"-enabled clouds because it requires a very complex examination of the privacy laws and rights for each jurisdiction in which the data resides or is processed.

The Directorate-General concludes that while cloud computing can bring together infinite amounts of data that can be used to solve many local and global problems as well as promote "green" computing worldwide, many areas still need to be researched to enable cloud computing to become an effective solution that benefits industry, government and individuals worldwide.

European Network and Information Security Agency

Supported by a group of subject matter experts representing industry, academia and governmental organizations, the European Network and Information Security Agency (ENISA) has conducted a risk assessment on cloud computing business models and technologies. The result is the "Cloud Computing Security Risk Assessment,"¹⁸ an in-depth and independent analysis that outlines information security benefits and key security risks of cloud computing and makes practical recommendations.

ISACA

As an independent, nonprofit, global membership association, ISACA (www.isaca.org) engages in the development, adoption and use of globally accepted, industry-leading knowledge and practices for information systems. ISACA helps its members achieve individual and organizational success, resulting in greater trust in, and value from, information systems. Its members and certification holders are qualified and skilled professionals who make a difference.

¹⁸ [Cloud Computing Security Risk Assessment.pdf](#)

ISACA members receive a subscription to *ISACA Journal*, an authoritative, peer-reviewed publication that reports on timely topics for business and information technology professionals.

National Institute of Standards and Technology

The National Institute of Standards and Technology (NIST) promotes the effective and secure use of cloud computing technology in government and industry by providing technical guidance and promoting standards.

NIST's website (www.nist.gov) defines cloud computing and offers guidance.¹⁹ NIST's definition serves as a foundation for upcoming publications on cloud models, architectures, and deployment strategies. Computer scientists at NIST developed the cloud draft definition in collaboration with industry and government; it is expected it to evolve over time as cloud computing and cloud technology mature.

The NIST Guide refers to organizations that are interested in authorizing an information system such as a cloud environment. Organizations in these settings are collectively responsible and accountable for the information system, jointly accepting the IT-related business risks.

The Open Group / Jericho Forum

The Jericho Forum is an international group of organizations working together to define and promote solutions related to de-perimeterisation. Before it was formally established in 2003, the Jericho Forum existed as a loose affiliation of corporate chief information security officers.

One of the group's early publications was a position paper entitled the Jericho Forum Commandments.²⁰ The Commandments are guidelines for survival in a de-perimeterised world.

The Jericho Forum began in the UK, and nearly all of its original members worked for corporations in positions with global responsibilities. Today, the group has participation from users in Europe, North America and Asia Pacific; vendors in Europe and North America; and academic institutions in Europe and Asia Pacific.

US Regulatory Agencies

In addition to the groups listed above, US regulatory bodies have issued substantial guidance on third-party service providers that are generally applicable to cloud environments. In using this guidance, regulators routinely caution that it should not be applied in a "check-box" fashion.

¹⁹ For more information on NIST's work, see [Agenda and Presentations from the May 20, 2010 Cloud Computing Forum & Workshop; NIST Definition of Cloud Computing v15; Presentation on Effectively and Securely Using the Cloud Computing Paradigm v26](#)

²⁰ See http://www.opengroup.org/jericho/commandments_v1.2.pdf

Rather, the criteria listed in the guidance should be viewed as indicators of a broader and healthy operational risk management capability.

For US banking organizations, the Federal Financial Institution Examination Council's IT Examination guidance is the primary practical reference. Two of the eleven booklets are particularly applicable to cloud environments: "Outsourcing Technology Services" and "Supervision of Technology Service Providers."²¹ The other booklets (which cover operations, business continuity and other topics) are applicable to financial institutions and their service providers, and can be used to provide additional detail in client evaluations of service providers. In addition, the FFIEC recommends guidance from several organizations including ISACA, IIA, AICPA and the ABA. Of these, ISACA offers the most relevant guidance through these publications:

- Control Objectives for Basel II²² (for use in international engagements)
- COBIT Mapping: Mapping FFIEC to COBIT 4.1²³
- Risk IT Based on COBIT²⁴ (for cross-industry use, was designed to be integrated into financial institutions' operational risk functions)

For insurance in the US, the National Association of Insurance Commissioners now uses COBIT for IT Examination.²⁵

In many other countries, ISACA guidance is officially required, recommended or the de-facto criteria. Check your national regulator's website or ask your examiner before your examination begins for more information.

²¹ See http://www.ffiec.gov/ffiecinfobase/html_pages/it_01.html

²² See <http://www.isaca.org/Knowledge-Center/Research/ResearchDeliverables/Pages/IT-Control-Objectives-for-Basel-II-The-Importance-of-Governance-and-Risk-Management-for-Compliance.aspx>

²³ See <http://www.isaca.org/Knowledge-Center/Research/ResearchDeliverables/Pages/COBIT-Mapping-Mapping-FFIEC-With-COBIT-41.aspx>

²⁴ See www.isaca.org/riskit

²⁵ http://www.naic.org/committees_e_examover_it_exam.htm