SFG | SHARED ASSESSMENTS
The Trusted Source in Third Party Risk Management

BUILDING BEST PRACTICES

Evaluating Cloud Risk for the Enterprise – An Updated Shared Assessments Guide

Revised June 2017

SFG | SHARED ASSESSMENTS

BUILDING BEST PRACTICES:

Evaluating Cloud Risk for the Enterprise – Table of Contents

SFG | SHARED ASSESSMENTS

BUILDING BEST PRACTICES:

Evaluating Cloud Risk for the Enterprise – An Updated Shared Assessments Guide

EXECUTIVE SUMMARY

This Guide is the second iteration of the Evaluating Cloud Risk for the Enterprise, the first one having been published in 2010. In the past seven years we have seen tremendous changes in technology, personnel and business practices. Cloud has now become the de facto industry model for providing a computing service. Mobile has become the most common model for accessing data. Cloud platforms are managing billions of Internet of Things (IoT) devices daily and new exciting developments are evolving, such as microservices, to allow previously unimaginable scalability and efficiencies.[1]

This updated Guide is based on the combined experience of the hundreds of Shared Assessments members and peer organizations across all verticals who have successfully integrated cloud computing into their operations.[2] Many of the Delta Cloud Controls have stood the test of time and some new controls have been added based on industry change and evolving technology.

This Guide fosters successful deployment and monitoring of cloud computing technologies by helping organizations and their risk managers better understand and evaluate the use of cloud computing enterprise-wide. Included are practical recommendations, questions to discuss with cloud providers, and lessons learned for control domains that are cloud-related. The recommendations and guidance in this document may be used in conjunction with the Shared Assessments Program Tools and resources, or may be selectively incorporated into other types of audits or assessments of environments containing cloud elements, such as the AICPA Service Organization Control (SOC) or Statements on Standards for Attestation Engagements (SSAE).

Ultimately, this Cloud Guide targets audiences with varying 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.

Using this Cloud Guide, risk managers can begin to evaluate specific areas of cloud risk, ask the right questions and ensure they get answers they understand.

We hope you enjoy the new Cloud Guide.

Signed: *The Shared Assessments Team*

Through 2018, more than 75% of fully successful implementations are predicted to be delivered with a cloud-native, DevOps-centric service delivery approach.

North Bridge, 2017

By 2020, 92% of all workloads will be processed by cloud data centers and cloud is rapidly becoming the de facto model for providing computer services.

Cisco, 2016

THE CLOUD MODEL

The steady and often invisible movement from models in which data is stored centrally into one in which data is held in distributed systems has, unfortunately, lulled risk and control professionals into a false sense of security. Many have believed there would be ample time to start building much-needed controls for distributed models before they became widely deployed. However, outsourcing and remote work environments have significantly eroded the concept and practicality of perimeter security, and, as a whole, risk management has often lagged in devising the critical alternative controls to protect the increasingly porous enterprise environments.

The implications are profound for both cloud application use and security models: compliance no longer trumps security, and risk concerns don't center only on who is accessing data, but also what data and how. As a result of increased investment in security, 69% of 451 survey respondents plan to move workflows that use regulated data into the cloud – respondents in France lead the pack followed by the US, Germany and the UK.

451 Research, 2016

Cloud deployment is more evolutionary than revolutionary. Yet cloud has been highly disruptive to the IT industry, impacting how companies build and procure IT services.

By 2020:

- 68% of cloud workloads will be in public cloud data centers and 32% will be private cloud data centers.
- Delivery by model will be: 74% total cloud workloads as Software as a Service – SaaS (up from 65% in 2015); 17% Infrastructure as a Service – IaaS (down from 26% in 2015); and 8% Platform as a Service – PaaS (down from 9% in 2015).[3]
- IoT endpoints being estimated to nearly triple by 2020.[4]
- Use of hybrid environments is a significant growth trend, in which onsite resources are combined with hosted private or public cloud computing resources.[5]
- Utilization of hyperscale cloud operators are anticipated to take 47% of all installed data centers by 2020; accounting for 83% of the public cloud server base and 86% of all workloads.[6,7]

Cloud Computing Services: IaaS, PaaS and SaaS

This Guide defines cloud computing as consisting of three distinct service types: Infrastructure as a Service (IaaS); Platform as a Service (PaaS); and Software as a Service (SaaS). As a cloud infrastructure, the IaaS layer can host PaaS and SaaS environments.

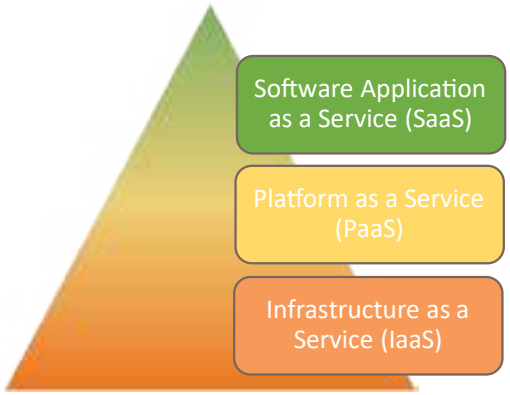

Figure 1: Cloud Computing Service Types

An emerging trend in IaaS is Desktop as a Service (DaaS), where a virtual terminal is implemented within the cloud that replaces internal IT infrastructure on the customer's premises. While not as prevalent as other cloud computing models, DaaS providers have deployed offerings which are slowly being adopted by companies of all sizes. Strategic use of cloud service providers has also given rise to managed service providers (MSPs), which are a type of cloud services broker (CSB). MSPs may own facilities or lease co-location/wholesale data center provider facilities. The MSP provides professional management services to an outsourcer, in effect coordinating all the access points that tie to cloud usage for a particular enterprise. MSPs would typically be able to deliver a management portal, technical support and solution architecture, workload migration services and DevOps automation. They may also offer prebuilt cloud solutions by vertical or use type.[8]

A RISK MANAGEMENT APPROACH FOR CLOUD COMPUTING

The face of third party risk management is changing as supply chain management evolves toward digital and cloud technologies with internal and external collaborations driving adoption of the cloud for platforms and storage solutions. This has led to Business managers and IT specialists frequently asking about the differences between traditional IT outsourcing and cloud models from the perspective of security, auditing and risk management. Risks surrounding cloud model use are depicted in Figure 2 below.[8]

This focus on what is different and new is key, since so many enterprise organizations have been outsourcing successfully using traditional models and have acquired a wealth of corresponding knowledge and experience upon which they can build. 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.

Common and Delta Cloud Controls

To help address these issues, 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 usage. For purposes of this guide, these controls have been grouped into 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 control areas that have particular relevance to cloud environments, and whose cloud audit mechanisms are less well known (e.g., multi-tenancy.)

An enterprise organization evaluating a cloud solution or provider might have a list of 100 control requirements to minimize risk, including concentration of services, IT management processes, data and information security policies protection, patch and antivirus-management, privacy, recovery and resiliency 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 (virtualization, for example) insufficiently examined and the company exposed.

Figure 2: Hierarchy of IT Risks
Cloud Computing Service Types[9]

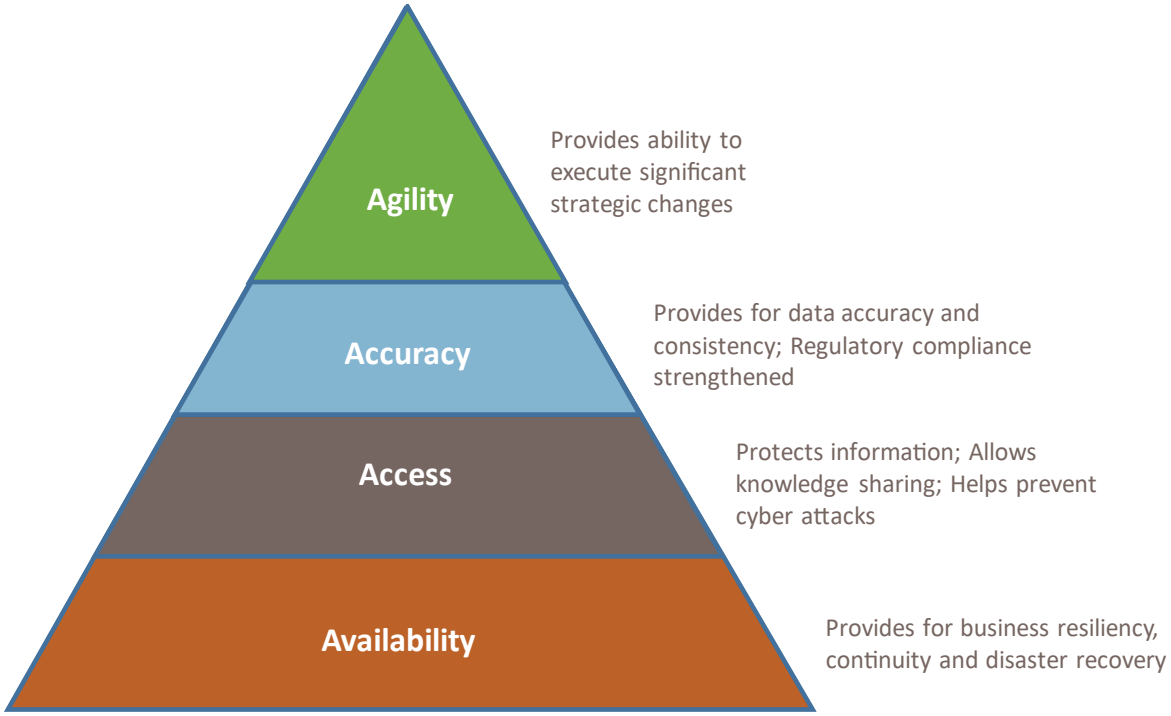

To remedy this, consider applying a risk management approach to cloud engagements. Beginning with an examination of the Common Cloud Controls, use the cloud provider's existing audit reports and certifications to cross reference. Approaching the evaluation this way allows the majority of controls and risks to be analyzed 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. Your team may not be as experienced in evaluating these Delta Cloud Controls and your existing audit programs may not cover them sufficiently. (For example, virtualization is largely ignored or omitted entirely in PCI-DSS, ISO 27001, and HIPAA audits.)

The Delta Cloud Controls section of this Guide includes twelve cloud control domains that include numerous best practices for evaluating cloud provider controls.

Common Cloud Controls

Many of today's enterprise organizations are well versed in evaluating traditional IT controls. An organization should be able to leverage their existing procurement and vendor management processes to complete a significant portion of the evaluation process. When these mature processes and teams are applied to cloud, the effort, time and cost required to evaluate cloud providers can be significantly reduced. One example would be to review the cloud provider's SOC 2, Shared Assessments AUP or ISO 27001 assessment, as these sources can help answer the majority of security questions; and then move to the Delta Cloud Controls (which are not typically contained in the aforementioned reports).

The table below is not intended to be an exhaustive list of Common Cloud Control areas. Instead, it is offered as a starting point to illustrate 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.

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 5
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 5, ISO 27001
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 5, ISO 27001 (Annex A controls), NIST 800-53A (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.	SSAE/SOC 2, PCI DSS, NIST 800-53A (Appendix F procedure catalog, controls portion), Shared Assessments SIG and AUP. Vendor-specific security guidance.

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.

In Common Cloud Controls, commonly used controls are applied to the cloud environment in new ways. With the 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 Controls are divided into twelve domains:

1. Multi-Tenancy
2. Concentration Risk
3. Agile Delivery
4. Virtualization and Containerization
5. Cloud Providers and Locations
6. Legal and Privacy
7. Roles and Responsibilities
8. Identity and Log Management
9. Application Security
10. Vendor Governance and Interdependence
11. Data Retention, Management, Recovery and Destruction
12. E-Discovery and Forensics

As with the Common Cloud Controls, the twelve Delta Cloud Control Domains 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.

TABLE 2: DELTA CLOUD CONTROL AREAS – RISK MANAGEMENT BEST PRACTICES GUIDELINES

Process areas for cloud provider assessment that may impact review of outsourcing risk posture

DOMAIN: 1) MULTI-TENANCY

Significance of the Control Area	Cloud Best Practices
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 severs, 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.	• Document the data-segmentation and separation controls at each of the four main layers: network, physical, system and 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, both across external and internal networks • 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 comingled 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 • 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. Search 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 may be successful if there is weak application data segmentation in place.[10]

TABLE 2: DELTA CLOUD CONTROL AREAS – RISK MANAGEMENT BEST PRACTICES GUIDELINES

DOMAIN: 2) CONCENTRATION RISK

Significance of the Control Area	Cloud Best Practices
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. Concentration risk most commonly occurs in single-vendor relationships, or placing all your eggs in one basket. The loss of that vendor may result in unplanned service outages, disruption of services, damage to your brand and reputation and higher costs. Related to cloud service providers this takes on a different meaning, in that multiple outsourcers may be utilizing the same cloud service provider and/or their services. Therefore, when this provider suffers an outage, many firms may be affected.	• Ensure that you have a clear understanding of the recovery and resiliency capabilities of the Cloud Service provider and that they meet your recovery time objectives. Pay particular attention to their underlying infrastructure e.g. a SaaS provider may own no servers and be deployed wholly on IaaS. • In SaaS, one-off customer product changes are often difficult to execute on. This is due to the fact that all product changes can have a direct impact, good or bad, on all the other customers. Evaluate how open the cloud provider is to customer input in their road-map, e.g., ask what percentage of features over the past year were based on customer requests. • Create a list of your 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).

DOMAIN: 3) AGILE DELIVERY

Significance of the Control Area	Cloud Best Practices
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.	• Decide whether it will be acceptable to your business to receive near continuous (iterative) 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 controls checking 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. • Ensure adequate segregation of duties exists, in that code creators and approvers remain separate in an agile environment.

TABLE 2: DELTA CLOUD CONTROL AREAS – RISK MANAGEMENT BEST PRACTICES GUIDELINES

DOMAIN: 4) VIRTUALIZATION AND CONTAINERIZATION

Significance of the Control Area	Cloud Best Practices
Clients of a Cloud provider (e.g. public cloud) 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. Common mechanisms for providing this separation of data and services are virtualization and containerization.	• 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 Technologies provides a good starting point.[11] • Confirm that the cloud provider has the controls in place to ensure that only authorized snapshots are taken, and their level of classification and storage locations of these snapshots are commensurate in strength with the virtualized production environment. • Review in detail the controls in place around the hypervisor as it manages the virtual environments. Who has administrative access to it? What level of logging is enabled? Are logs reviewed or monitored? Is the hypervisor physical server or network separate from the general system? • Confirm that an application container hardening standard exists and is implemented (if microservices are used). Apart from general OS hardening, this should also include, but not be limited to, specific container hardening requirements such as: create a separate partition for containers; only allow trusted users to control container daemon; audit container daemon; restrict network traffic between containers; configure TLS authentication for the container daemon; consider using a container authorization plugin for granular control; verify that container service file permissions are set e.g., 644 or more restrictive.[12]

DOMAIN: 5) CLOUD PROVIDERS AND LOCATIONS

Significance of the Control Area	Cloud Best Practices
Cloud providers should be able to identify the location or region that client data is stored, processed or transmitted. 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." Clients should know the location or the legal jurisdiction, e.g. US, Germany, etc.	• Request the locations or regions where client data will be stored, processed, accessed or transmitted. Note: In some cases, the PaaS may point to regions only (e.g., US East). • Understand who has access to data. Ask your provider to list all of its vendors that have direct access to your data — in particular, any cloud vendors that will store, process, transmit or have access to your data. • Request that cloud providers contractually alert the client of changes in the aforementioned vendors (if any) 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. • Confirm that the client can define the legal jurisdictions to which client data can be replicated.

DOMAIN: 6) LEGAL AND PRIVACY

Significance of the Control Area	Cloud Best Practices
Before moving a service out of an organization to any third party, a rigorous legal analysis and evaluation needs to be conducted. This is especially important if data will be stored, processed or transmitted in a foreign country. 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.	• Establish who the owner of the data is "Data Controller" and what rights the cloud provider "Data Processor" has to the data. In nearly all cases the client should own the data, and the cloud provider should have no ownership rights to it. • List all locations and service providers that store, process, transmit or access client data and whether these are contractually documented. • Define in the contract the countries or legal jurisdictions 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. One example is that 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 places 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 documented permission before it stores data outside of a specific predefined country.* • Establish whose data privacy policy applies and how the contractual requirements will be implemented. • 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. • Provide contractual assurances that define what data must be encrypted and in what state, e.g., transit or storage.

TABLE 2: DELTA CLOUD CONTROL AREAS – RISK MANAGEMENT BEST PRACTICES GUIDELINES

DOMAIN: 6) LEGAL AND PRIVACY (CONTINUED)

Significance of the Control Area	Cloud Best Practices
Continued from previous page.	• 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.

DOMAIN: 7) ROLES AND RESPONSIBILITIES

Significance of the Control Area	Cloud Best Practices
Internal IT support teams may excel at managing internal servers but may need additional training to manage or provide effective oversight of cloud systems.	• Evaluate how increasing your use of cloud may affect your vendor management skill set requirements; begin this planning early. • Consider re-training staff on vendor management and cloud technologies. Staff will need to fully understand the relationship and technology changes so as 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). • 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.

DOMAIN: 8) IDENTITY AND LOG MANAGEMENT

Significance of the Control Area	Cloud Best Practices
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 two to 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.	• 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 Security Assertion Markup Language (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 if cloud provider's authentication controls support minimum company requirements such as multi-factor password complexity, password expiry, etc. • 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, e.g., SAML, Identity Federation Framework (ID-FF), Web Services (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 cloud dashboards, reports and Application Programming Interfaces (API's) that the cloud provider may make available, so as to ensure that they meet client requirements. Will this provide adequate monitoring capabilities, as cloud providers will typically not expose raw log data to the client for privacy reasons.

TABLE 2: DELTA CLOUD CONTROL AREAS – RISK MANAGEMENT BEST PRACTICES GUIDELINES

DOMAIN: 9) APPLICATION SECURITY

Significance of the Control Area	Cloud Best Practices
Application security is important in both traditional outsourcing models and cloud computing. 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 surface. In an agile model, the CSPs application code changes over very short cycles (measured in weeks). 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.	• 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 polices and processes that quickly become "shelf-ware" 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 securely inspection. Basic checks should include (but are not limited to): having a secure code review before shipping to production; that no credentials appear in source code; that appropriate permissions have been made for the source code etc. • 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): o Planning Planning and requirement phase - application security team should review business requirements and define relevant application security requirements to make sure that the security requirements are an integral part of business requirements. o Design phase - architecture and functional design review from a security perspective to make sure security requirements are designed as defined in the requirements phase. o Coding phase - security team perform security code review to find vulnerabilities in the code through a white box testing methodology. o Testing phase- security teams complete a penetration test to find security vulnerabilities in the code. A combination of black box and white box testing can also be used for a more thorough security test. o Maintenance phase – ongoing security team review and testing the changes and bug fixes.

DOMAIN: 10) VENDOR GOVERNANCE AND INTERDEPENDENCE

Significance of the Control Area	Cloud Best Practices
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. 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.	• Keep it under one service contract (where possible). 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. • 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 process exists to ensure service levels are met and issues are resolved. • Understand the processes interrelationships between the cloud provider and its vendors. 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 narrowly scoped • Cloud provider continuous monitoring and management should be part of your overall information security management system (ISMS) with defined and monitored metrics. • 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.

TABLE 2: DELTA CLOUD CONTROL AREAS – RISK MANAGEMENT BEST PRACTICES GUIDELINES

DOMAIN: 11) DATA RETENTION, MANAGEMENT, RECOVERY AND DESTRUCTION

Significance of the Control Area	Cloud Best Practices
The retention, management and destruction of data in any outsourcing model is critical. 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. Cloud is "always-on" by nature: for this to occur, the cloud provider needs to ensure it has superior operational, change management, disaster recovery and business continuity plans and controls.	• All data should hold an information-classification 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: (a) what is in scope for destruction; (b) whether logical or physical destruction controls will be used; and (c) 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 uptime. • 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 uptime 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 that mirror the provider's primary production environment. 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.

DOMAIN: 12) E-DISCOVERY AND FORENSICS

Significance of the Control Area	Cloud Best Practices
During legal e-discovery, organizations may receive subpoenas requesting relevant data to be retained and shared with a third party. 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.	• 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 comingled with other client data. Providing the disk or back-up tape would expose other clients' data, so the cloud provider would typically refuse 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 table 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.

Practitioners may also wish to refer to the Shared Assessments' companion paper, Assessment of Public Cloud Computing Vendors, which provides additional information for outsourcers on assessment and monitoring of their connections to cloud service providers.

CONCLUSION

We hope that this Guide will provide risk professionals at all levels with a greater understanding of key issues as they plan and implement management programs for their own computing services and other critical enterprise functions in the cloud.

This Guide is based on the combined experience of the hundreds of Shared Assessments members and peer organizations across all verticals who have successfully integrated cloud computing into their operations. Accordingly, it addresses concerns surrounding the significant drivers of increased cloud service use, which include: growing data security requirements; support for greater innovation; the need to improve customer experience; goals for reducing costs in IT service delivery; the need to scale IT resources quickly on demand; and the need to address globalization supply chain issues.

The analyses and guidelines provided here can be used in developing a comprehensive plan for evaluating, risk ranking and cost-effectively selecting cloud providers and solutions, with the goal of enabling successful deployment and integration across departments and lines of business. The recommendations and guidance in this document may be used in conjunction with the Shared Assessments Program Tools, or may be selected and incorporated into other types of audits or assessments of environments containing cloud elements.

Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. This cloud model promotes availability and is composed of five essential characteristics (On-demand self-service, Broad network access, Resource pooling, Rapid elasticity, Measured Service); three service models (Cloud Software as a Service [SaaS], Cloud Platform as a Service [PaaS], Cloud Infrastructure as a Service [IaaS]); and, four deployment models (Private cloud, Community cloud, Public cloud, Hybrid cloud). Key enabling technologies include: (1) fast wide-area networks, (2) powerful, inexpensive server computers, and (3) high-performance virtualization for commodity hardware.

National Institute of Standards in Technology

Definition of Cloud Computing

Additional copies of this guide, as well as Shared Assessments' Assessment of Public Cloud Computing Vendors and other third party risk management resources, may be downloaded at: https://sharedassessments.org/.

1. Microservices is a collection of loosely-combined services. This style is a modification of service-oriented architecture (SOA).
2. For the purposes of this discussion, cloud computing refers to all cloud-related deployment, resources and cloud computing.
3. Cisco Global Cloud Index: Forecast and Methodology. 2015-2020. Cisco. 2016.
4. Cloud Trends in 2017. IDC Opinion Paper. December 2016.
5. Cloud Trends in 2017. IDC Opinion Paper. December 2016.
6. Building Trust in a Cloudy Sky: The state of cloud adoption and security. Intel Security. January 2017.
7. Cisco Global Cloud Index: Forecast and Methodology. 2015-2020. Cisco. 2016.
8. Magic Quadrant for Public Cloud Infrastructure Managed Service Providers, Worldwide. Gartner Research. March 2017.
9. Graphic after: Westerman. 2006. Becker, J.D. & Bailey, E. IT Controls and Governance in Cloud Computing. 2014 AMCIS Proceedings Paper. Twentieth Americas Conference on Information Systems, Savannah. 2014.
10. 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.
11. Guide to Security for Full Virtualization Technologies. Special Publication 800-125. National Institute of Standards and Technology (NIST). January 2011.
12. CIS Docker 1.6 Benchmark. V 1.0.0. Center for Internet Security. 22 April 2015.
13. Networks Drive Digital Ambitions: Investment in digital transformation is being driven by cloud, security and network upgrades. IDG Connect. 2017.
14. Microsoft, Google and IBM Public Cloud Surge is at Expense of Smaller Providers. Synergy Research. February 2017.
15. Magic Quadrant for Cloud Infrastructure as a Service, Worldwide. Gartner Research. May 2015.
16. Magic Quadrant for Public Cloud Infrastructure Managed Service Providers, Worldwide. Gartner Research. March 2017.
17. Industry's Largest Cloud Survey Reveals Cloud Momentum Driving Enterprise to "Re-Orchestrate" Strategy. North Bridge. 2017. [Wikibon Research. July 2016.]
18. Networks Drive Digital Ambitions: Investment in digital transformation is being driven by cloud, security and network upgrades. IDG Connect. 2017.
19. The NIST Definition of Cloud Computing, Special Publication 800-145. September 2011.

Acknowledgments

Project Chair

Niall Browne, SVP Trust & Security, CSO, Domo

Editors

Niall Browne, SVP Trust & Security, CSO, Domo
Marya Roddis, VP Communications, The Santa Fe Group, Shared Assessments Program

Contributors

We'd like to thank the Shared Assessments Steering Committee members who conducted this effort:

Niall Browne, SVP Trust & Security, CSO, Domo
David Hubley, Director, Information Security Risk Management, Capital One
Shawn Malone, Founder & CEO, Security Diligence, LLC

We would also like to acknowledge The Santa Fe Group, Shared Assessments Program subject matter experts and other staff who supported this effort:

Bob Jones, Senior Advisor, The Santa Fe Group, Shared Assessments Program; Principal, RW Jones Associates, LLC; Adjunct, Professor of Economic Crime, Utica College
Katherine Kneeland, Senior Project Manager, The Santa Fe Group, Shared Assessments Program
Marya Roddis, VP Communications, The Santa Fe Group, Shared Assessments Program

---

About the Shared Assessments Program
The Shared Assessments Program has been setting the standard in third party risk management since 2005. Member-driven development of program resources helps organizations to effectively manage the critical components of the third party risk management lifecycle by creating efficiencies and lowering costs for conducting rigorous assessments of controls for cybersecurity, IT, privacy, data security and business resiliency.

Shared Assessments membership and use of the Shared Assessments Program Tools offers companies and their service providers a standardized, more efficient and less costly means for third party risk management programs. Program Tools are kept current with regulations, industry standards and guidelines and the current threat environment; and are adopted globally across a broad range of industries both by service providers and their customers.

The Program Tools follow a two-step "trust, but verify" model for assessing third party risks:

- The trust component is the Standardized Information Gathering (SIG) questionnaire, a holistic tool for risk management assessments, including assessments of cybersecurity, IT, privacy, data security and business resiliency controls.
- The verify portion of the Shared Assessments Program is facilitated by the use of the Shared Assessments (AUP), a holistic tool for standardized onsite risk management assessments that covers the same comprehensive risk areas incorporated into the SIG. The AUP allows an outsourcer to validate the answers provided on the SIG questionnaire. The AUP sets forth the risk control areas to be assessed as part of the onsite assessment, as well as the procedures to be used.
- The Vendor Risk Management Maturity Model (VRMMM) Tool is available for free to any organization that is seeking to evaluate the maturity of their third party risk program with a holistic tool that includes focuses on cybersecurity, IT, privacy, data security and business resiliency controls.

The Shared Assessments Program began specifically addressing cloud computing in 2009 by adding new procedures and cloud-relevant questions into its Program Tools. Enhancements continue to be made to Program Tools to improve their effectiveness, which includes updates that reflect the growing importance of cloud computing across the IT landscape.

Shared Assessments members are national and international organizations of all sizes from all verticals 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 maintains strategic partnerships with global associations, including the member-driven Cloud Security Alliance (CSA), the 70 member and Department of Treasury partnership Financial Services Sector Coordinating Council (FSSCC), and the 18 member Financial and Banking Information Infrastructure Committee (FBIIC). The Santa Fe Group is the strategic partner that manages the Program.

Together, the diverse membership of the Shared Assessments Program works to increase awareness and adoption of the Program Tools across industry sectors and around the globe. Any suggestions or questions may be directed to info@sharedassessments.org.

APPENDIX I: CLOUD COMPUTING – AN OVERVIEW

Cloud Computing Services: IaaS, PaaS and SaaS
This Guide defines cloud computing as consisting of three distinct service types: Infrastructure as a Service (IaaS); Platform as a Service (PaaS); and Software as a Service (SaaS). These three cloud service types can be viewed as a pyramid (see Figure 1). As a cloud infrastructure, the IaaS layer can host PaaS and SaaS environments.

- **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.

  As of this writing, Amazon Web Services (AWS) is maintaining its dominant share of the burgeoning public cloud (IaaS) services market at over 40%, while the three other cloud providers – Microsoft, Google and IBM – are gaining ground, but at the expense of smaller players in the market.[14] In response, many providers are shifting strategies.[15]

- **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. 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 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 SaaS vendor, being responsible for all bug fixes and enhancements to the application, as well as all other application support services.

- **Desktop as a Service (DaaS):** Less commonly used but growing, IaaS can include Desktop as a Service (DaaS) in which no internal IT exists with a virtual terminal being set up through the cloud. Typically, in the past, companies implemented and deployed an internal Virtual Desktop Infrastructure (VDI); however, this is being replaced with cloud providers. DaaS use is typically utilized within companies with BYOD (bring your own device) solutions. This provides a strong control over desktop use, settings and recovery. The DaaS provider monitors the software and hardware and applies all required patches, as well as providing a recovery option for all the virtual terminals involved.

Cloud vs. Hosted Applications (Traditional Model)
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 based on the service type being used. 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.

Cloud vs. Licensed Software Vendors
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.

Managed Service Providers
Strategic use of cloud service providers has given rise to managed service providers (MSP), which are a type of cloud services broker (CSB). MSPs may own facilities or lease co-location/wholesale data center provider facilities. The MSP provides professional management services to an outsourcer, in effect coordinating all the access points that tie to cloud usage for a particular enterprise. All MSPs must be able to deliver a management portal, technical support and solution architecture, workload migration services and DevOps automation. They may also offer prebuilt cloud solutions by vertical or use type.[16]

This market responds directly to the increased focus on business agility and cost reduction that accompanies use

of hyperscale cloud operators. Through 2018, more than 75% of fully successful implementations are predicted to be delivered with a cloud-native, DevOps-centric service delivery approach.[17]

Two trends have increased the need and use of outsourced, online data use, further complicating the nature of data distribution:

- **Increased outsourcing of data, services and critical processes:** Traditionally, data and services resided deep within an organization, protected by the company's own employees and controls. With the advent of widespread outsourcing of data and services, data is now often stored and protected by numerous third and fourth parties, often in multiple, and frequently undisclosed, 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 within the confines of the office environment from criminal compromise. Remote access to company systems and data has allowed information to be stored on workers' local systems, resulting in increasing risk of data loss. Managing access risk is now recognized as a key issue and resources are increasingly being dedicated to data classification, data loss prevention and monitoring and reporting of defined access risk criteria.

One of the core components of cloud is its distributed nature. With distributed models, perimeter security diminishes in effectiveness, leaving significant exposures.

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, including with third and fourth party providers and the associated serviced data. One location may be on encrypted storage arrays within data centers accessed or processed using hardened systems and incorporating use of various data protection controls such as IPS, biometrics and armed guards. After normal business hours, that data may reside on company laptops or smartphones in homes, airports or on the front seat of the car; all locations where these types of controls are of little or no use.

Cloud users should refuse to sacrifice security for convenience. 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. Businesses should expect and demand all the risk management and security controls that traditional on-premise providers have — and more — from their cloud solutions. 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.

Companies have begun to develop strong initiatives in this direction. This is evidenced, in part, by the fact that, in 2016, "60% of current IT investment in UK, Germany, France, the Netherlands and Belgium is focused on digital transformation initiatives designed to support future business expansion and growth rather than keeping the lights on (maintaining applications and services)."[18]

Benefits of Cloud Computing

These qualities allow for building end-to-end visibility across the enterprise's network from a risk management standpoint. Cloud services offer many advantages, including:

- Faster deployment of computing services.
- Lower computing costs in the form of reduced IT infrastructure expense.
- 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 that can provide cost-effective security.
- Ubiquitous access (multiple networks, remote access, mobile devices).
- Improved performance (through the pooling and sharing of hardware resources).
- Improved disaster recovery plans and options.

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.

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

APPENDIX II: ADDITIONAL CLOUD COMPUTING INITIATIVES

The Shared Assessments Program strives to standardize and streamline all third party risk management across the vendor lifecycle by creating efficiencies and lowering costs for conducting rigorous assessments of controls for cybersecurity, IT, privacy, data security and business resiliency.

The following consortia and other organizations are also working to define the controls necessary for cloud's widespread acceptance in the enterprise and therefore improve security and risk management in the cloud. The summaries below provide an overview of some of the initiatives of the following organizations:

- Commission of the European Communities
- Cloud Security Alliance
- European Network and Information Security Agency (ENISA)
- National Institute of Standards and Technology (NIST)
- Regulatory Bodies
- UT System Cyber and Cloud Security Initiative, University of Texas at Austin

Commission of the European Communities

The Commission of the European Communities examines cloud security and adoption issues as part of its follow-up to the public consultation launched in 2009 on the review of the EU's regulatory framework for data protection, the Commission arranged a series of targeted consultation meetings with a number of key stakeholders. The Commission consults with the non-public sector stakeholders on a range of issues that pertain to existing data protection rules, emerging rules (such as the General Data Protection Regulation – GDPR), and identifies problems and possible solutions.

The Directorate-General Information Society published The Future of Cloud Computing: Opportunities for European Cloud Computing Beyond 2010. 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 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 industries, government and individuals worldwide.provide education to help secure all other forms of computing.

Cloud Security Alliance

The Cloud Security Alliance (CSA) is a global not-for-profit organization with a broad geographical distribution. Its mission is to "promote the use of best practices for providing security assurance within Cloud Computing," and provide education to help secure all other forms of computing. The CSA supports a broad spectrum of subject matter expertise, including experts in cloud, security, legal issues, compliance and virtualization that provide leadership and guidance for security best practices. CSA operates a cloud security provider certification program, the Certificate of Cloud Security Knowledge (CCSK) and Certified Cloud Security Professional (CCSP); the CSA Security, Trust & Assurance Registry (STAR), "a three-tiered provider assurance program of self-assessment; 3rd party audit and continuous monitoring program; and conducts comprehensive research in collaboration with industry, higher education and government on a global basis."

European Network and Information Security Agency (ENISA)

The European Union Agency for Network and Information Security (ENISA) is "a centre of expertise for cyber security in Europe." ." The mission of ENISA is to contribute to securing Europe's information society by raising "awareness of network and information security and to develop and promote a culture, of network and information security in society for the benefit of citizens, consumers, enterprises and public sector organizations in the Union." ENISA's strategic objectives are derived from the ENISA regulation, inputs from the Member States and relevant communities, including the private sector. ENISA strategically seeks to: support Europe in facing emerging network and information security challenges; promote network and information security by assisting in development and implementation of policies and laws related to network information security; support information security capacities; and fostering the emerging European network and information security

community. It does so, in part, through educational resources, including its most recent publication, Technical Guidelines for the implementation of minimum security measures for Digital Service Providers and a 2010 report, Cloud Computing Security Risk Assessment, which provides an in-depth and independent analysis that outlines information security benefits and key security risks of cloud computing and makes practical recommendations.

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 provides guidelines on enterprise and government body use of 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. In October 2014, NIST released its special publication 500-293, US Government Cloud Computing Technology Roadmap, Volume II, High-Priority Requirements to Further USG Agency Cloud Computing Adopting.

Regulatory Agencies and International Guidance
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.

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

In recognition of the need for greater IT and cloud security, the US Department of Homeland Security oversees the National Initiative for Cybersecurity Careers and Studies (NICCS). This cloud/virtual environment security training course trains hands-on beginners to advanced techniques with the tools necessary to develop a cloud implementation plan to ensure security.

For US banking organizations, the Federal Financial Institution Examination Council's (FFIEC) 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." 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. The FFIEC guidance was updated in January 2017. In addition, the FFIEC recommends guidance from several organizations including ISACA, IIA, AICPA and the ABA.

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.

UT System Cyber and Cloud Security Initiative, University of Texas at Austin
The University of Texas at Austin "is committed to supporting cybersecurity and identity management education and research." The University of Texas System, in support of its member institutions, is pursuing a strategic initiative focusing on the future needs of cybersecurity organizations and workforce. "The initiative is a convergence of core competencies and capabilities across the UT System's member institutions to build an ecosystem integrating areas of education and training, industry outreach, research, applied technologies, and academic activities mapped to specific economic and educational needs of Texas (public and private sectors). The initiative supports degree programs dedicated to cybersecurity."

APPENDIX III: GLOSSARY

Term	Working Definition
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.
Application Delivery Controllers (ADC)	ADC's are deployed as a management platform to gain visibility between data centers (in-house, private/public cloud) to optimize application performance, security and resource efficiency for both externally-facing web applications and service delivery options. Virtual ADC platforms, in cloud settings, are more flexible than traditional software-based platforms for business applications and protocols.
Asset Classification	The category or type assigned to an asset, derived from the asset classification policy. Asset classifications frequently vary from company to company.
Broad Network Access	Those resources, hosted in cloud network(s) and available from a wide range of devices and locations through online access.
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 ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. This cloud model promotes availability and is composed of five essential characteristics (On-demand self-service, Broad network access, Resource pooling, Rapid elasticity, Measured Service); three service models (Cloud Software as a Service [SaaS], Cloud Platform as a Service [PaaS], Cloud Infrastructure as a Service [IaaS]); and, four deployment models (Private cloud, Community cloud, Public cloud, Hybrid cloud). Key enabling technologies include: (1) fast wide-area networks, (2) powerful, inexpensive server computers, and (3) high-performance virtualization for commodity hardware.[19]
Cloud Deployment Models	• Private cloud operated within the confines of an organization's firewall. Otherwise known as 'enterprise cloud hosting.' • Community cloud - provides a cloud solution to a targeted community of limited users with similar requirements who work collaboratively to govern, manage and secure the solution. • Public cloud – scalable cloud solutions offered to external users to achieve economies of scale and resource sharing in any vertical or jurisdiction where services are shared. There may be no guarantee on where data is stored. • Hybrid cloud – cloud services managed across a blend of external and internal providers. "Cloud bursting" is the term applied when daily needs are met in a private cloud and expanded dynamically as needed into a pre-determined public cloud hybrid use arrangement.
Communication Tree	A document stating when, how and to whom communication must be made if an unexpected event occurs.
Containerization	This provides an alternative to full machine virtualization. Containerization involves encapsulating an application into a separate operating environment for physical or virtual systems to gain the benefits of a virtual machine without the same dependencies.
Controls	The safeguards or countermeasures utilized 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.
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.
Fourth Party	A subcontractor to the outsourcer's third party, regardless of whether the fourth party has a potential materiality/criticality impact on the original outsourcer.
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.

Term	Working Definition
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.
On-Demand Self-Service Cloud Computing	External provision of cloud resources 'on-demand' as needed by an enterprise; the self-service component involves user (customer) access through an online control panel that allows for scaling up of the infrastructure during run time use in a subscription based model with service delivery changes effected by the user that are separate from the actual administration of the assets by the provider.
Outsourcer	The entity delegating a function to a subcontractor through the process of 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 a formally assigned, 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.
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 the most current version of the Shared Assessments Standardized Information Gathering Questionnaire (SIG) and the Agreed Upon Procedures (AUP) procedures. 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.
Rapid Elasticity	The ability to dynamically scale services being provided in direct response to the need of customers for space and other services. This is one of the five fundamental aspects of cloud computing.
Remote Access	The ability to log in to a network from a distant location.

Term	Working Definition
Resource Pooling	Provision of scalable services, such as data storage services and bandwidth services, through the use of software as a service (SaaS) at a meta level so that resource availability can be adjusted to meet unique customer requirements in a manner that is transparent to the user.
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 Governance in IT Context	The entity delegating a function to a subcontractor through the process of outsourcing - the contracting out of a business or technology function — commonly one previously performed in-house — to a third-party provider.
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 as a Service (SecaaS)	The next generation of managed security services dedicated to the delivery, over the Internet, of specialized information-security services.
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.
Shadow IT	IT devices, services and software that are procured outside the user organization's ownership or control.
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 an 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

Term	Working Definition
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.
Virtualization	The creation of a virtual (rather than actual) version of something, such as an operating system, server, storage device or network resource. Virtualization represents a paradigm of data transmission that is multipoint to multipoint in many different physical locations. Virtualization allows organizations to run tens (or even hundreds) of virtual operating systems on the same physical server. Virtualization provides tremendous efficiency of scale; however, can introduce risks and greater difficulty in tracking and protecting data.

P: (505) 466-6434
F: (505) 466-3111
E: info@santa-fe-group.com