Security Standards Council®

Standard: PCI Data Security Standard (PCI DSS)

Version: 2.0

Date: February 2013

Author: Cloud Special Interest Group
PCI Security Standards Council

Information Supplement:

PCI DSS Cloud Computing Guidelines

Table of Contents

1 Executive Summary

Cloud computing is a form of distributed computing that is yet to be standardized[1]. There are a number of factors to be considered when migrating to cloud services, and organizations need to clearly understand their needs before they can determine if and how they will be met by a particular solution or provider. As cloud computing is still an evolving technology, evaluations of risks and benefits may change as the technology becomes more established and its implications become better understood.

Cloud security is a shared responsibility between the cloud service provider (CSP) and its clients. If payment card data is stored, processed or transmitted in a cloud environment, PCI DSS will apply to that environment, and will typically involve validation of both the CSP's infrastructure and the client's usage of that environment. The allocation of responsibility between client and provider for managing security controls does not exempt a client from the responsibly of ensuring that their cardholder data is properly secured according to applicable PCI DSS requirements.

It's important to note that all cloud services are not created equal. Clear policies and procedures should be agreed between client and cloud provider for all security requirements, and responsibilities for operation, management and reporting should be clearly defined and understood for each requirement.

1.1 Intended Use

This document provides guidance on the use of cloud technologies and considerations for maintaining PCI DSS controls in cloud environments. This guidance builds on that provided in the PCI DSS Virtualization Guidelines and is intended for organizations using, or thinking of using, providing, or assessing cloud technologies as part of a cardholder data environment (CDE).

This document is structured as follows:

- **Executive Summary –** Includes a brief summary of some key points and provides context for the remainder of the document.
- **Cloud Overview –** Describes the deployment and service models discussed throughout this document.
- **Cloud Provider/ Cloud Customer Relationships –** Discusses how roles and responsibilities may differ across different cloud service and deployment models
- **PCI DSS Considerations –** Provides guidance and examples to help determine responsibilities for individual PCI DSS requirements, and includes segmentation and scoping considerations.
- **PCI DSS Compliance Challenges –** Describes some of the challenges associated with validating PCI DSS compliance in a cloud environment.
- **Additional Security Considerations –** Explores a number of business and technical security considerations for the use of cloud technologies.
- **Conclusion –** Presents recommendations for starting discussions about cloud services.

[1] *NIST Guidelines on Security and Privacy in Public Cloud Computing* (SP SP800-144)

The following appendices are included to provide additional guidance:

- **Appendix A: PCI DSS Responsibilities for different Service Models –** Presents additional considerations to help determine PCI DSS responsibilities across different cloud service models.
- **Appendix B: Sample Inventory –** Presents a sample system inventory for cloud computing environments.
- **Appendix C: PCI DSS Responsibility Matrix –** Presents a sample matrix for documenting how PCI DSS responsibilities are assigned between cloud provider and client.
- **Appendix D: PCI DSS Implementation Considerations –** Suggests a starting set of questions that may help in determining how PCI DSS requirements can be met in a particular cloud environment.

This document is intended to provide an initial point of discussion for cloud providers and clients, and does not delve into specific technical configurations. This document does not endorse the use of any specific technologies, products, or services.

The information in this document is intended as supplemental guidance and does not supersede, replace or extend PCI DSS requirements. For the purposes of this document, all references made are to PCI DSS version 2.0.

1.2 Audience

The information in this document is intended for merchants, service providers, assessors and other entities looking for guidance on the use of cloud computing in the context of PCI DSS. For example:

- **Merchants –** The security and PCI DSS considerations are applicable to all types of cloud environments, and may be useful to merchants managing their own cloud infrastructure as well as those looking to engage with a third party. Guidance for working with third-party cloud providers and PCI DSS compliance challenges may also be useful.
- **Cloud service providers –** The security and PCI DSS considerations may provide useful information for CSPs to assist their understanding of the PCI DSS requirements, and may also help CSPs to better understand their clients' PCI DSS needs. Guidance on CSP/client relationships and PCI DSS compliance challenges may also be useful for providers.
- **Assessors –** The security and PCI DSS considerations may help assessors to understand what they might need to know about an environment in order to be able to determine whether a PCI DSS requirement has been met.

1.3 Terminology

The following terms are used throughout this document:

- **CSP – Cloud Service Provider**. The CSP, or cloud provider, is the entity providing the cloud service. The CSP acquires and manages the infrastructure required for providing the services, runs the cloud software that provides the services, and delivers the cloud services through network access.[2]
- **Cloud customer or client –** The entity subscribing to a service provided by a cloud provider. May include merchants, service providers, payment processors, and other entities utilizing cloud services. May also be referred to as a cloud tenant.

[2] *NIST Cloud Computing Reference Architecture* (SP 500-292)

2 Cloud Overview

Cloud computing provides a model for enabling on-demand network access to a shared pool of computing resources (for example: networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or cloud provider interaction.[3]

Cloud computing can be used to provide clients with access to the latest technologies without a costly investment in hardware and software. Due to the economies of scale associated with the delivery of cloud services, CSPs can often provide access to a greater range of technologies and security resources than the client might otherwise have access to. Client organizations without a depth of technically-skilled personnel may also wish to leverage the skills and knowledge provided by CSP personnel to securely manage their cloud operations.

Cloud computing therefore holds significant potential to help organizations reduce IT complexity and costs, while increasing agility. Cloud computing is also seen as a means to accommodate business requirements for high availability and redundancy, including business continuity and disaster recovery.

2.1 Deployment and Service Models

Deployment models are defined to distinguish between different models of ownership and distribution of the resources used to deliver cloud services to different customers. Cloud environments may be deployed over a private infrastructure, public infrastructure, or a combination of both. The most common deployment models, as defined by NIST, include:

- **Private cloud** – The cloud infrastructure is operated solely for a single organization (client). It may be managed by the organization itself or a third-party provider, and may be on-premise or off-premise. However, it must be solely dedicated for the use of one entity.
- **Community cloud –** The cloud infrastructure is shared by several organizations and supports a specific community with shared requirements or concerns (for example, business model, security requirements, policy, or compliance considerations). It may be managed by the organizations or a third party, and may be on-premise or off-premise.
- **Public cloud –** The cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services. Public cloud infrastructure exists on the premises of the cloud provider.
- **Hybrid cloud –** The cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by technology to enable portability. Hybrid clouds are often used for redundancy or load-balancing purposes—for example, applications within a private cloud could be configured to utilize computing resources from a public cloud as needed during peak capacity times (sometimes called “cloud-bursting”).

With respect to understanding roles and responsibilities, this paper is largely focused on public cloud scenarios. However, many of the concepts discussed remain applicable to the other deployment models.

[3] *The NIST Definition of Cloud Computing* (SP 800-145)

Service models identify different control options for the cloud customer and cloud provider. For example, SaaS customers simply use the applications and services provided by the CSP, where IaaS customers maintain control of their own environment hosted on the CSP's underlying infrastructure.

The three most commonly used service models are described as follows[4]:

Software as a Service (SaaS) – Capability for clients to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through either a thin client interface, such as a web browser, or a program interface.

Platform as a Service (PaaS) – Capability for clients to deploy their applications (created or acquired) onto the cloud infrastructure, using programming languages, libraries, services, and tools supported by the provider.

Infrastructure as a Service (IaaS) – Capability for clients to utilize the provider's processing, storage, networks, and other fundamental computing resources to deploy and run operating systems, applications and other software on a cloud infrastructure.

The main difference between service levels relates to how control is shared between client and CSP, which in turn impacts the level of responsibility for both parties. It should be noted that, other than in a truly private cloud (on-premise) scenario, the client rarely has any control over hardware, and it is the degree to which virtual components, applications and software are managed by the different parties that differentiates the service models. As a general rule, SaaS provides clients with the least amount of control, whereas IaaS offers the most control for the client.

It's important to note that these descriptions for deployment and service models, although widely accepted by the industry, may not be universally followed by cloud providers or reflect actual cloud environments. For example, a CSP might be selling a "private cloud" service that does not meet the intent of "private" as it is described above. Similarly, the details of what is and what is not included in a particular service will probably vary between CSPs, even if they each identify their service by the same term (IaaS, PaaS, or SaaS).

The level of security responsibility across the cloud service models generally migrates towards the client as the client moves from a SaaS model (least client responsibility) to an IaaS model (most client responsibility). The greatest level of responsibility for the CSP to maintain security and operational controls is present in the SaaS service model.

Figure 1 on the following page shows how control is typically shared between the CSP and client across different service models.

[4] Adapted from *The NIST Definition of Cloud Computing* (SP 800-145)

Figure 1: Level of control/responsibility for client and CSP across different service models

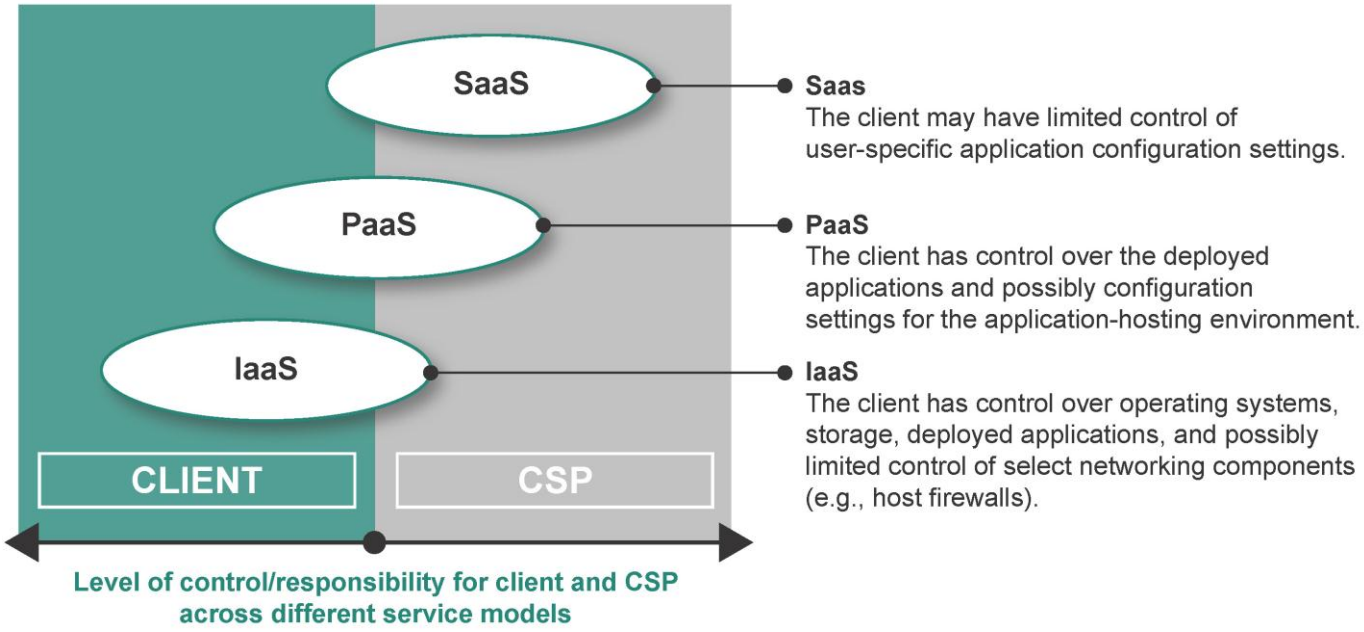

While clients may be attracted to the SaaS and PaaS models due to the resource savings and reduced responsibility for administering the cloud environment, they should be aware that these models also correspond to a greater loss of control of the environment housing their sensitive data. Contractual agreements and ongoing due diligence become especially critical where control is outsourced, to ensure that the required security measures are being met and maintained by the CSP for the duration of the agreement.

3 Cloud Provider / Cloud Customer Relationships

3.1 Understanding Roles and Responsibilities

The lines of accountability and responsibility will be different for each service and deployment model. Clear policies and procedures should be agreed upon between client and cloud provider for all security requirements, and clear responsibilities for operation, management and reporting need to be defined for each requirement.

3.2 Roles and Responsibilities for Different Deployments Models

The entity performing the role of CSP will vary according to the type of deployment model. For example, the CSP role may be assigned entirely to an external third party (as in a public cloud), or the role may be undertaken by an internal department or business function (as in an on-premise private cloud). Similarly, the role of CSP may be assigned to more than one entity in a community or hybrid cloud scenario.

To understand how responsibilities are assigned in a particular deployment model, consider the following:

- **Private cloud –** Where a private cloud is managed on-premise, the CSP role may be undertaken within the client organization. For example, the IT department could take on the role of CSP with various operational departments as its clients. In this scenario, the client organization retains full control of their environment and its security and compliance.

  Dedicated, private clouds may also be provisioned off-premise by a third-party CSP. In this case, the delineation of responsibility will also depend on the particular service model, as described in Section 3.3, "Responsibilities for Different Service Models."

- **Community cloud –** The CSP could be one of the client organizations within the community or a separate third party. The delineation of responsibility follows the particular service model implemented.

- **Public cloud –** The CSP is a third party that is an organizationally-separate entity to its clients. The cloud is deployed within a CSP's environment and responsibility is delineated according to the particular service model, as defined by the CSP.

- **Hybrid cloud –** The CSP role may be assigned to both internal and third-party entities for different elements of the overall cloud infrastructure. Responsibility will be assigned based on the combination of deployment models and service models implemented.

The responsibility for implementation, operation, and management of security controls will be shared differently within each of the cloud models, and needs to be clearly understood by both the client and CSP. The client also needs to understand the level of oversight or visibility they will have into security functions that are outside their control. If these security responsibilities are not properly assigned, communicated, and understood, insecure configurations or vulnerabilities could go unnoticed and unaddressed, resulting in potential exploit and data loss or other compromise.

3.3 Responsibilities for Different Service Models

In all deployment models, and particularly in public cloud environments, it is important for all parties to understand the specific elements of the service model used and its associated risks. Any cloud deployment model that is not truly private (on-premise) is by nature a shared responsibility model, where a portion of responsibility for the cloud service falls under the realm of the CSP, and a portion of responsibility also falls to each client. The level of responsibility that falls to the CSP or the client is determined by the cloud service model being utilized—that is, IaaS, PaaS, or SaaS. Clear delineation of responsibilities should be established as a prerequisite to any cloud service implementation to provide a baseline for the cloud operation.

Figure 2 on the following page illustrates how control of the different technical layers is often shared across different service models. For illustration purposes, different layers of the cloud stack are described as follows:

Layer	Description
Application Program Interface (API) or Graphical User Interface (GUI)	The interface used by the client or their customers to interact with the application. The current most common API is RESTful HTTP or HTTPS. The current most common GUI is an HTTP or HTTPS based Web site.
Application	The actual application being used by one or more clients or their customers.
Solution stack	This is the programming language used to build and deploy applications. Some examples include .NET, Python, Ruby, Perl, etc.
Operating systems (OS)	In a virtualized environment, the OS runs within each VM. Alternatively, if there is no underlying hypervisor present, the operating system runs directly on the storage hardware.
Virtual machine (VM)	The virtual container assigned for client use.
Virtual network infrastructure	For communications within and between virtual machines
Hypervisor	When virtualization is used to manage resources, the hypervisor is responsible for allocating resources to each virtual machine. It may also be leveraged for implementing security.
Processing and memory	The physical hardware that supplies CPU time and physical memory.
Data Storage	The physical hardware used for file storage.
Network	This can be a physical or virtual network. It is responsible for carrying communications between systems and possibly the Internet.
Physical facility	The actual physical building where the cloud systems are located.

Appendix B illustrates a sample inventory for cloud computing systems, as guidance for how CSPs and their customers can document the different layers of the cloud environment.

Figure 2: Example of how control may be assigned between CSP and clients across different service models.

Cloud Layer	Service Models		
	IaaS	PaaS	SaaS
Data	Client	Client	Client
Interfaces (APIs, GUIs)	Client	Client	CSP
Applications	Client	Client	CSP
Solution Stack (Programming languages)	Client	Client	CSP
Operating Systems (OS)	Client	CSP	CSP
Virtual Machines	Client	CSP	CSP
Virtual network infrastructure	Client	CSP	CSP
Hypervisors	CSP	CSP	CSP
Processing and Memory	CSP	CSP	CSP
Data Storage (hard drives, removable disks, backups, etc.)	CSP	CSP	CSP
Network (interfaces and devices, communications infrastructure)	CSP	CSP	CSP
Physical facilities / data centers	CSP	CSP	CSP

Note: *This table provides an example of how responsibilities might be assigned according to common descriptions of the different service models. However, it's important to note that the technology layers and their corresponding lines of responsibility may be different for each CSP, even if they use the same terminology to describe their service, and the individual service offerings may or may not align with the responsibly assignments indicated above.*

Some CSPs offer multiple "options" for their services—for example, a CSP may have one IaaS offering that includes a client-controlled hypervisor and a separate IaaS offering with no client access to the hypervisor. It's imperative that clients and CSPs clearly document and understand where the boundaries are in their particular relationship rather than assuming that any particular responsibility model applies to them.

Even where a client does not have control over a particular layer, they may still have some responsibility for the configurations or settings that the CSP maintains on their behalf. For example, a client may need to define firewall rules and review firewall rule-sets for those firewalls applicable to the protection of their environment, even though the CSP actually configures and manages the firewalls. Similarly, clients may be responsible for approving and reviewing user access permissions to their data resources, while the CSP configures the access according to client needs.

The allocation of responsibility for managing security controls does not exempt a client from the responsibility of ensuring that their cardholder data is properly secured.

3.4 Nested Service-Provider Relationships

Nested service-provider relationships are not uncommon in cloud scenarios, as CSPs sometimes rely on other third-party companies to deliver their services. For examples, some CSPs use third-party storage providers as part of their cloud service offering, while some might partner with other CSPs for redundancy or fail-over as part of their cloud-delivery strategy.

Identifying all third-party relationships that the CSP has in place is important in order to understand the potential ramifications to a client's environment. The existence of multiple nested relationships—for example, where there is a chain of vendors and/or other providers required for delivery of a cloud service—will also add complexity to both the CSP's and the client's PCI DSS assessment process.

4 PCI DSS Considerations

4.1 Understanding PCI DSS Responsibilities

The responsibilities delineated between the client and the CSP for managing PCI DSS controls are influenced by a number of variables, including but not limited to:

- The purpose for which the client is using the cloud service.
- The scope of PCI DSS requirements that the client is outsourcing to the CSP.
- The services and system components that the CSP has validated within its own operations.
- The service option that the client has selected to engage the CSP (IaaS, PaaS or SaaS).
- The scope of any additional services the CSP is providing to proactively manage the client's compliance (for example, additional managed security services).

The client needs to clearly understand the scope of responsibility that the CSP is accepting for each PCI DSS requirement, and which services and system components are validated for each requirement. For example, PCI DSS Requirements 6.1 and 6.2 address the need for vulnerabilities to be identified, ranked according to risk, and deployed in a timely manner. If not properly defined, a client could assume that the CSP is managing this process for the entire cloud environment, whereas the CSP could be managing vulnerabilities for their underlying infrastructure only, and assuming that the client is managing vulnerabilities for operating systems and applications.

4.2 PCI DSS Responsibilities for Different Service Models

As a general rule, the more aspects of a client's operations that the CSP manages, the more responsibility the CSP has for maintaining PCI DSS controls. However, outsourcing maintenance of controls is not the same as outsourcing responsibility for the data overall. Cloud customers should not make assumptions about any service, and should clearly spell out in contracts, memorandums of understanding, and/or SLAs exactly which party is responsible for securing which system components and processes.

Figure 3 on the following page provides an example of how responsibilities for PCI DSS requirements may be shared between clients and CSPs across the three service models. There will of course be exceptions and variations across each individual service, and this table is provided as a guideline for clients and CSPs to help plan discussions and negotiations.

Responsibilities have been identified as follows:

- **Client –** Generally each client will retain responsibility for maintaining and verifying the requirement.
- **CSP –** Generally the CSP will maintain and verify the requirement for their clients.
- **Both –** Generally responsibility is "shared" between the client and the CSP. This may be due to the requirement applying to elements present in both the client environment and the CSP-managed environment, or because both parties need to be involved in the management of a particular control.

***Appendix A** includes additional considerations for determining how PCI DSS responsibilities may be assigned for each service model.*

Figure 3: Example of how PCI DSS responsibilities may be shared between clients and CSPs.

	Client
	CSP
	BOTH Client and CSP

PCI DSS Requirement	Example responsibility assignment for management of controls		
	IaaS	**PaaS**	**SaaS**
1: Install and maintain a firewall configuration to protect cardholder data	Both	Both	CSP
2: Do not use vendor-supplied defaults for system passwords and other security parameters	Both	Both	CSP
3: Protect stored cardholder data	Both	Both	CSP
4: Encrypt transmission of cardholder data across open, public networks	Client	Both	CSP
5: Use and regularly update anti-virus software or programs	Client	Both	CSP
6: Develop and maintain secure systems and applications	Both	Both	Both
7: Restrict access to cardholder data by business need to know	Both	Both	Both
8: Assign a unique ID to each person with computer access	Both	Both	Both
9: Restrict physical access to cardholder data	CSP	CSP	CSP
10: Track and monitor all access to network resources and cardholder data	Both	Both	CSP
11: Regularly test security systems and processes	Both	Both	CSP
12: Maintain a policy that addresses information security for all personnel	Both	Both	Both
PCI DSS Appendix A: Additional PCI DSS Requirements for Shared Hosting Providers	CSP	CSP	CSP

***Note:** The sample responsibilities illustrated in this table do not include consideration for any activities or operations performed outside of a hypothetical cloud service offering. This table provides an example of how PCI DSS responsibilities might be assigned for different service models. However, each CSP ultimately defines their own service, and particular service offerings may or may not be consistent with those illustrated above. Clients and CSPs should clearly document their responsibilities as applicable to their particular agreement.*

The concept of “shared” or “joint” responsibility can be a particular tricky path to navigate. While some services and functions will be relatively straightforward to scope and establish boundaries, many services and functions will overlap if not clearly demarcated at the outset of the service relationship.

Where the CSP maintains responsibility for PCI DSS controls, the client is still responsible for monitoring the CSP's ongoing compliance for all applicable requirements. CSPs should be able to provide their clients with ongoing assurance that requirements are being met, and where the CSP is managing requirements on behalf of the client, they should have mechanisms in place to provide the customer with the applicable records—for example, audit logs showing all access to client data.

Clients are still required to validate their compliance in accordance with payment brand programs.

Appendix C *illustrates a sample PCI DSS Responsibly Matrix, as guidance for how CSPs and their customers can document PCI DSS responsibility assignments.*

Appendix D *includes Implementation Considerations for PCI DSS Requirements.*

4.3 Security as a Service (SecaaS)

Security as a Service, or SecaaS, is sometimes used to describe the delivery of security services using a SaaS-based delivery model. SecaaS solutions not directly involved in storing, processing, or transmitting CHD may still be an integral part of the security of the CDE. As an example, a SaaS-based anti-malware solution may be used to update anti-malware signatures on the client's systems via a cloud-delivery model. In this example, the SecaaS offering is delivering a PCI DSS control to the client's environment, and the SecaaS functionality will need to be reviewed to verify that it is meeting the applicable requirements.

4.4 Segmentation Considerations

Outside of a cloud environment, individual client environments would normally be physically, organizationally, and administratively separate from each other. Clients utilizing a public or otherwise shared cloud must rely on the CSP to ensure that their environment is adequately isolated from the other client environments.

In addition to enforcing separation between client environments, segmentation may also be desired within a client's environment to isolate their CDE components from non-CDE components in order to reduce their own PCI DSS scope.

Segmentation on a cloud-computing infrastructure must provide an equivalent level of isolation as that achievable through physical network separation. Mechanisms to ensure appropriate isolation may be required at the network, operating system, and application layers; and most importantly, there should be guaranteed isolation of data that is stored. Client environments must be isolated from each other such that they can be considered separately managed entities with no connectivity between them. Any systems or components shared by the client environments, including the hypervisor and underlying systems, must not provide an access path between environments. Any shared infrastructure used to house an in-scope client environment would be in scope for that client's PCI DSS assessment.

A segmented cloud environment exists when the CSP enforces isolation between client environments. Examples of how segmentation may be provided in shared cloud environments include, but are not limited to:

- Traditional Application Service Provider (ASP) model, where physically separate servers are provided for each client's cardholder data environment.
- Virtualized servers that are individually dedicated to a particular client, including any virtualized disks such as SAN, NAS or virtual database servers.

- Environments where clients run their applications in separate logical partitions using separate database management system images and do not share disk storage or other resources.

The PCI DSS assessor must validate the effectiveness of the segmentation to ensure it provides adequate isolation. If adequate segmentation is provided between clients, the client environment and the CSP-managed environment and processes would be in scope for a client's PCI DSS assessment. If adequate segmentation is not in place or cannot be verified, the entire cloud environment would be in-scope for any one client's assessment. Examples of "non-segmented" cloud environments include but are not limited to:

- Environments where organizations use the same application image on the same server and are only separated by the access control system of the operating system or the application.
- Environments where organizations use different images of an application on the same server and are only separated by the access control system of the operating system or the application.
- Environments where organizations' data is stored in the same instance of the database management system's data store.

Without adequate segmentation, all clients of the shared infrastructure, as well as the CSP, would need to be verified as being PCI DSS compliant in order for any one client to be assured of the compliance of the environment. This will likely make compliance validation unachievable for the CSP or any of their clients.

4.4.1 Segmentation Challenges

Segmentation in traditional hosted environments can be applied via separate physical servers and security measures applied using known methods. The difference in a cloud environment is that there are common shared layers (such as hypervisors and virtual infrastructure layers), which can present a single point of entry (or attack) for all systems above or below those shared layers. The security applied to these layers is therefore critical not only to the security of the individual environments they support, but also to ensure that segmentation is enforced between different client environments.

Once any layer of the cloud architecture is shared by CDE and non-CDE environments, segmentation becomes increasingly complex. This complexity is not limited to shared hypervisors; all layers of the infrastructure that could provide an entry point to a CDE must be included when verifying segmentation.

In a private cloud environment, one approach that may help reduce the complexity of segmentation efforts could be to locate all CDE virtual components on a dedicated "CDE hypervisor," and ensure all non-CDE virtual components are located on separate hypervisors, adequately segmented from the CDE hypervisor.

The need for adequate segmentation of client environments in a public or shared cloud is underscored by the principle that the other client environments running on the same infrastructure are to be considered untrusted networks. The client has no way of confirming whether other client environments are securely configured, patched appropriately to protect against attack, or that they are not already compromised or even designed to be malicious. This is particularly relevant where a CSP offers IaaS and PaaS services, as the individual clients have greater control and management of their environments.

4.4.2 Segmentation Responsibilities

Ultimately, the CSP needs to take ownership of the segmentation between clients and verify it is effective and provides adequate isolation between individual client environments, between client environments and the CSP's own environment, and between client environments and other untrusted environments (such as the Internet). Applicable PCI DSS controls for the segmentation functions would also be the CSP's responsibility (for example, firewall rules, audit logging, documentation, reviews, etc.). The client is responsible for the proper configuration of any segmentation controls implemented within their own environment (for example, using virtual firewalls to separate in-scope VMs from out-of-scope VMs), and for ensuring that effective isolation is maintained between in-scope and out-of-scope components.

Clients wishing to implement segmentation within their cloud environment also need to consider how the CSP's environment and processes may impact the effectiveness of the segmentation. For example, CSP systems could be providing connectivity between the client's own VMs that is not visible to the client. Clients should also consider how the CSP manages offline or dormant VMs, and whether in-scope and out-of-scope VMs could potentially be stored together by the CSP without active segmentation controls.

4.4.3 Segmentation Technologies

Traditional network segmentation technologies consist of hardware devices such as firewalls, switches, routers, and so forth. These physical components could be used to separate VMs hosted on the same or multiple hypervisors similar to the manner in which systems could be segmented in a "physical" network. This would require hypervisors with multiple network interfaces and PCI DSS compliant configurations for the various types of network hardware. Additionally, virtual counterparts of firewalls, switches and routers now exist and can be incorporated into a virtual environment.

As mentioned above, a key consideration is how secure the "common layers" (such as hypervisors and shared physical components) are, and whether they represent a potential attack surface between zones or clients. The answer is that yes, they do; however the associated risks are still not well understood.

Examples of controls to be considered when evaluating segmentation options include, but are not limited to:

- Physical firewalls and network segmentation at the infrastructure level
- Firewalls at the hypervisor and VM level
- VLAN tagging or zoning in addition to firewalls
- Intrusion-prevention systems at the hypervisor and/or VM level to detect and block unwanted traffic
- Data-loss-prevention tools at the hypervisor and/or VM level
- Controls to prevent out-of-band communications occurring via the underlying infrastructure
- Isolation of shared processes and resources from client environments
- Segmented data stores for each client
- Strong, two-factor authentication
- Separation of duties and administrative oversight
- Continuous logging and monitoring of perimeter traffic, and real-time response

4.5 Scoping Considerations

Merchant or other organizations looking to store, process, or transmit payment card data in a cloud environment should clearly understand the impact that extending their CDE into the cloud will have on their PCI DSS scope. For example, in a private-cloud deployment, an organization could either implement adequate segmentation to isolate in-scope systems from other systems and services, or they could consider their private cloud to be wholly in scope for PCI DSS. In a public cloud, the client organization and CSP will need to work closely together to define and verify scope boundaries, as both parties will have systems and services in scope.

> ***Appendix D*** *includes Implementation Considerations for PCI DSS Requirements.*

Recommendations for minimizing and simplifying PCI DSS scope in a cloud environment include:

- Don't store, process or transmit payment card data in the cloud. This is the most effective way to keep a cloud environment out of scope, as PCI DSS controls are not required if there is no payment card data to protect.
- Implement a dedicated physical infrastructure that is used only for the in-scope cloud environment. The scoping process will be simplified if all in-scope operations are limited to a known, defined set of physical and virtual system components that are managed independently from other components. Once defined, the client will be reliant on the CSP's ability to ensure scope boundaries are maintained—for example, by ensuring that all segmentation controls are operating effectively and that any new components connected to the in-scope environment are immediately brought into scope and protected accordingly.
- Minimize reliance on third-party CSPs for protecting payment card data. The more security controls the CSP is responsible for, the greater the scope of the CDE will potentially be, thereby increasing the complexity involved in defining and maintaining CDE boundaries.

Ensuring that clear-text account data is never accessible in the cloud may also assist to reduce the number of PCI DSS requirements applicable to the cloud environment. As an example, let's say the client performs all encryption and decryption operations and all key-management functions[5] in their own data center and uses a third-party cloud only to store or transmit encrypted data. In this scenario, clear-text data would never exist in the cloud environment—not even temporarily or in memory. Additionally, the cloud environment would never have access to cryptographic keys or key-management processes.

It should be noted that the encrypted data is still in scope for PCI DSS (generally for the entity that controls or manages the encrypted data and/or the cryptographic keys[6]) to ensure that applicable controls are in place. However, by keeping all encryption/decryption and key-management operations isolated from the cloud, the number of PCI DSS requirements that the CSP is required to maintain may be reduced, as these requirements will instead be applicable to the client's own environment and personnel. The CSP will still be in scope for any PCI DSS requirements it manages on behalf of the client—for example, access controls managed by the CSP will need to be verified to ensure that only authorized persons (as determined by the client) have access to the encrypted data, and that access is not granted to unauthorized persons.

[5] In accordance with PCI DSS Requirements
[6] Refer to FAQ "*Is encrypted cardholder data in scope for PCI DSS?"* on PCI SSC website for additional guidance.

Alternatively, if clear-text account data is present (for example, in memory) in the cloud environment, or the ability to retrieve account data exists (for example, if decryption keys and encrypted data are present), all applicable PCI DSS requirements would apply to that environment.

4.5.1 Scoping Examples for Different Deployment Models

For private cloud environments, segmentation efforts are focused on isolating CDE components from non-CDE components to reduce the number of systems in scope for PCI DSS. In public or shared cloud environments, segmentation between clients is critical for the security of the entire client environment, and is additional to any segmentation managed by the client within their environment for the purposes of scoping.

A number of simple scoping examples are presented here to provide guidance.

Scenario	Environment description	PCI DSS scoping guidance
Case 1: Private Cloud hosted and controlled by entity seeking PCI DSS compliance, with segmentation.	• All CDE VMs are hosted on a single, dedicated hypervisor; non-CDE VMs are hosted on a separate hypervisor(s). • Validated segmentation of CDE systems from non-CDE systems using a combination of physical and logical controls	The CDE hypervisor and VMs, and all cloud components that are not segmented are in scope (segmentation must be validated as providing effective isolation)
Case 2: Private Cloud hosted and controlled by entity seeking PCI DSS compliance, no segmentation.	• All VMs are hosted on one or more hypervisors; some VMs are considered part of the CDE and some are not. • No segmentation of CDE systems from non-CDE systems.	The entire cloud environment and all connected systems are in scope and considered part of the CDE (similar to a flat network).

Scenario	Environment description	PCI DSS scoping guidance
Case 3: Third-party CSP hosting a "PCI DSS compliant" public cloud supporting multiple clients, with validated segmentation for client environments.	• VMs may be on one or multiple hypervisors, all hypervisors and VMs are configured by CSP to support PCI DSS requirements. • Multiple clients hosted on each hypervisor. • Validated segmentation of client environments using a combination of physical and logical controls.	The CSP is responsible for compliance of all elements of the cloud service provided. Each client's scope would include their own environment (for example, VMs, applications etc.) and any other elements not managed by the CSP. Segmentation must be validated as providing effective isolation between clients as part of the CSP's validation, and may require additional validation as part of each client's validation.
Case 4: Third-party CSP hosting a "PCI DSS compliant" public cloud supporting multiple clients, no client segmentation.	• VMs may be on one or multiple hypervisors, all hypervisors configured by CSP to support PCI DSS requirements. • Multiple clients hosted on each hypervisor, VM configuration managed by each client. • Segmentation between client environments is not verified.	Entire cloud service and all client environments are in scope. Note that validating PCI DSS compliance may be intractable and infeasible as every client environment would need to be included in the assessment.

5 PCI DSS Compliance Challenges

Storing, processing, or transmitting cardholder data in the cloud brings that cloud environment into scope for PCI DSS, and it may be particularly challenging to validate PCI DSS compliance in a distributed, dynamic infrastructure such as a public or other shared cloud. For example, it can be difficult to identify which system components are in scope for a particular service, or identify who is responsible for particular PCI DSS controls. Some of the technical controls and auditing processes traditionally used to attain a measurable level of assurance in static environments (for example, in-house data storage servers) are not designed for rapidly-changing cloud environments and processes (for example, cloud bursting, continual deployment and retirement of virtual machines, dynamic IP addressing, and so on). Additionally, clients and assessors often can't "see and touch" CDE systems as they would in a traditional environment (for example, by visiting the data center).

The distributed architectures of cloud environments add layers of technology and complexity that challenge traditional assessment methods. For example, how does an assessor determine an appropriate sample size for a dynamic cloud environment in which systems can appear and disappear in minutes?

Examples of compliance challenges include but are not limited to:

- Clients may have little or no visibility into the CSP's underlying infrastructure and the related security controls.
- Clients may have limited or no oversight or control over cardholder data storage. Organizations might not know where cardholder data is physically stored, or the location(s) can regularly change. For redundancy or high availability reasons, data could be stored in multiple locations at any given time.
- Some virtual components do not have the same level of access control, logging, and monitoring as their physical counterparts.
- Perimeter boundaries between client environments can be fluid.
- Public cloud environments are usually designed to allow access from anywhere on the Internet.
- It can be challenging to verify who has access to cardholder data processed, transmitted, or stored in the cloud environment.
- It can be challenging to collect, correlate, and/or archive all of the logs necessary to meet applicable PCI DSS requirements.
- Organizations using data-discovery tools to identify cardholder data in their environments, and to ensure that such data is not stored in unexpected places, may find that running such tools in a cloud environment can be difficult and result in incomplete results. It can be challenging for organizations to verify that cardholder card data has not "leaked" into the cloud.
- Many large providers might not support right-to-audit for their clients. Clients should discuss their needs with the provider to determine how the CSP can provide assurance that required controls are in place.

These challenges will impact a number of factors related to how PCI DSS compliance is managed, including how segmentation is implemented, how PCI DSS assessments are scoped, how individual PCI DSS requirements are validated, and which party will perform particular validation activities.

At a high level, CSPs can be identified as those that have been validated as meeting a particular level of PCI DSS compliance and those that have not. The recommended practice for clients with PCI DSS considerations is to work with CSPs whose services have been independently validated as being PCI DSS compliant.

5.1 What does "I am PCI compliant" mean?

Much stock is placed in the statement "I am PCI compliant", but what does this actually mean for the different parties involved?

Use of a PCI DSS compliant CSP does not result in PCI DSS compliance for the clients. The client must still ensure they are using the service in a compliant manner, and is also ultimately responsible for the security of their CHD—outsourcing daily management of a subset of PCI DSS requirements does not remove the client's responsibility to ensure CHD is properly secured and that PCI DSS controls are met. The client therefore must work with the CSP to ensure that evidence is provided to verify that PCI DSS controls are maintained on an ongoing basis—an Attestation of Compliance (AOC) reflects a single point in time only; compliance requires ongoing monitoring and validation that controls are in place and working effectively.

Even where a cloud service is validated for certain PCI DSS requirements, this validation does not automatically transfer to the client environments within that cloud service. For example, a CSP's validation may have included use of up-to-date anti-virus software on the CSP's systems; however, this validation might not extend to the individual client OS or VMs (such as in an IaaS service). Additionally, the client must still maintain compliance for all of their own operations—for example, by ensuring anti-virus is installed and updated on all client-side systems used to connect into the cloud environment.

Similarly, a client's PCI DSS compliance does not result in any claim of compliance for the CSP, even if the client's validation included elements of the service managed by the CSP.

Regarding the applicability of one party's compliance to the other, consider the following:

a) If a CSP is compliant, this does not mean that their clients are.

b) If a CSP's clients are compliant, this does not mean that the CSP is.

c) If a CSP and the client are compliant, this does not mean that any other clients are.

The CSP should ensure that any service offered as being "PCI compliant" is accompanied by a clear and unambiguous explanation, supported by appropriate evidence, of which aspects of the service have been validated as compliant and which have not.

5.2 Verifying Scope of Validated Services and Components

Clients should first verify that the service they are using is the one that has been validated. CSPs that have validated PCI DSS compliance may be included on a list published by a payment card brand or they may not; either way, the client will need to obtain details of the CSP's compliance validation in order to determine whether the service they are using is wholly covered.

Considerations for the client may include:

- How long has the CSP been PCI DSS compliant? When was their last validation?
- What specific services and PCI DSS requirements were included in the validation?
- What specific facilities and system components were included in the validation?
- Are there any system components that the CSP relies on for delivery of the service that were not included in the PCI DSS validation?
- How does the CSP ensure that clients using the PCI DSS compliant service cannot introduce non-compliant components to the environment or bypass any PCI DSS controls?

CSPs should provide their clients with evidence that clearly identifies what was included in the scope of their PCI DSS assessment, as well as the specific PCI DSS requirements that the environment was assessed against, and the date of the assessment. All aspects of the cloud service ***not*** covered by the CSP's PCI DSS assessment should also be identified and documented, as these will need to be validated either by the client or the CSP in order for a client's assessment to be completed. The client must have a detailed understanding of any security requirements that are not covered by the provider and are therefore the client's responsibility to implement, manage, and validate as part of their own PCI DSS compliance.

CSPs that have undergone an independent PCI DSS assessment to validate their compliance will have the results summarized in an Attestation of Compliance (AOC) and detailed in a Report on Compliance (ROC). The Executive Summary and Scope of Work sections of the ROC should detail the scope of the assessment including the specific components, facilities, and services that were assessed.

5.3 Verifying PCI DSS Controls Managed by the Cloud Provider

As with all hosted services in scope for PCI DSS, the client organization should request sufficient evidence and assurance from their CSP that all in-scope processes and components under the CSP's control are PCI DSS compliant. This verification may be completed by the client's assessor (such as a QSA or ISA) as part of client's PCI DSS assessment. If the CSP has already undergone a PCI DSS assessment that was performed by another assessor, the client's assessor will need to verify that the CSP's validation is current, that the assessment covered all services provided to or used by the client, and that all applicable requirements were found to be in place for the environments and systems in scope.

CSPs that have undergone PCI DSS compliance assessment and validation should be able to provide their clients with the following:

- Proof of compliance documentation (such as the AOC and applicable sections from the ROC), including date of compliance assessment
- Documented evidence of system components and services that were included in the PCI DSS assessment
- Documented evidence of system components and services that were excluded from the PCI DSS assessment, as applicable to the service
- Appropriate contract language, if applicable

CSPs that have not undergone a PCI DSS compliance assessment will need to be included in their client's assessment. The CSP will need to agree to provide the client's assessor with access to their environment in order for the client to complete their assessment. The client's assessor may require onsite access and detailed information from the CSP, including but not limited to:

- Access to systems, facilities, and appropriate personnel for on-site reviews, interviews, physical walk-throughs, etc.
- Policies and procedures, process documentation, configuration standards, training records, incident response plans, etc.
- Evidence (such as configurations, screen shots, process reviews, etc.) to show that all applicable PCI DSS requirements are being met for the in-scope system components
- Appropriate contract language, if applicable

The client and CSP will need to agree upon which assessment activities can be performed by the client and which testing is the responsibility of the CSP. For example, in an IaaS/PaaS service, the client may wish to test within their own environment and whatever else they can access, such as the boundaries between themselves and other clients, or between themselves and the CSP's systems. However, if such testing is not permitted by the CSP, the client will have to rely on the CSP to perform and validate these requirements. In SaaS environments, the client will have limited or no visibility or permission to perform testing, and will generally be reliant on the CSP for all testing and validation. Defined testing activities and their associated controls and permissions should be detailed in the SLA.

The CSP also needs to be able to provide clients with specific details as applicable to the ongoing maintenance of PCI DSS compliance. For example, depending on the service provided, the CSP may need to produce copies of log files, patch update records, or firewall rule-sets that specifically apply to an individual client's environment.

CSPs wishing to provide a PCI DSS compliant service may wish to consider isolating the PCI DSS compliant services from their non-PCI compliant services. This may help to simplify the compliance validation process for both the CSP and for their individual clients. It may also help the CSP to standardize the PCI DSS compliant services being provided to their clients.

6 Additional Security Considerations

While the use of cloud services can provide an attractive opportunity for organizations of all sizes to outsource and utilize centrally-managed security resources, organizations should also be aware of the risks and challenges associated with a particular cloud choice before moving their sensitive data or services into the cloud environment. This section explores some of these additional security considerations.

6.1 Governance, Risk and Compliance

One of the primary challenges with cloud environments is that governance, compliance, and risk management are typically shared between the client and CSP. This shared delineation of responsibilities emphasizes the importance of a strong governance and risk-management structure. Without a clear governance strategy, the client may be unaware of issues arising from use of the cloud service, and the CSP may be unaware of issues within the client environment that could impact their service provision. A strategy for shared governance and communication should be established between client and CSP to enable clear communication of all aspects of the relationship from operational performance to security risk management and issue resolution. Reporting and monitoring mechanisms should be made available to client organizations to provide assurance that effective governance is applied by the CSP.

6.1.1 Risk Management

Consistent with a risk-management approach for in-house services, outsourced cloud services should be assessed against an organization's risk criteria with the intent of identifying critical assets, analyzing potential vulnerabilities and threats to those assets, and developing an appropriate risk-mitigation strategy (see PCI DSS Requirement 12.1.2). Lack of physical control of infrastructure, as occurs when the environment is outsourced to a third-party CSP, renders a thorough risk-management process all the more important.

In traditional environments, the physical location of sensitive data can be restricted to dedicated systems, facilitating the identification and implementation of effective risk-mitigation controls. However, the advent of new technologies requires a reevaluation of traditional risk strategies. For example, data in cloud environments is no longer tied to a physical system or location, reducing the effectiveness of traditional security mechanisms to protect data from risk. Traditional security approaches that build security controls "around" sensitive data may therefore need to evolve to address this new risk environment.

Similarly, traditional forms of risk assessment might not take into consideration particular cloud characteristics, such as a pay-as-you go model or multi-tenancy (described in Section 6.5.7), and may therefore require new or modified procedures.

6.1.2 Due Diligence

A CSP that stores, processes, or transmits cardholder data on behalf of a client, provides a security service for the protection of a client's cardholder data, or could otherwise impact the security of a client's cardholder data, would be considered a third-party service provider of the client. As with all service providers, clients should follow a thorough due-diligence process (see PCI DSS Requirement 12.8) prior

to engagement of the CSP. The specific due-diligence process and goals will vary for each client organization, however common objectives typically include:

1. Confirming the provider has a history of sound work practices and ethical behavior and is legitimately performing the services the client believes them to be
2. Verifying that the provider is compatible with the client's business image and risk profile
3. Identifying potential risks or circumstances associated with the provider that may impact the client's operations or business
4. Identifying elements of the service that need to be clarified, and that need to be included in contracts or service agreements

Due diligence is not simply reading the provider's marketing material or relying on a provider's claims of "PCI compliance" or secure operations. Clients should be sufficiently assured that they are engaging with a provider that can meet their security and operational needs before undertaking any such engagements. The scope of the due-diligence exercise should consider, at a minimum, the topics discussed throughout this document, as applicable to a client's particular requirements.

6.1.3 Service Level Agreements (SLAs)

The use of cloud services includes the deployment of a defined service model and should always be underwritten by comprehensive service level agreements (SLAs). The secure delivery of any cloud service is dependent on the CSP's personnel, processes, and technologies, while the secure usage of cloud services remains the responsibility of the client.

Typically, cloud-hosting agreements are concerned with "up-time" and high availability, with little or no mention or assurance of security. However, the client is ultimately responsible for ensuring the service they're using meets their security requirements and compliance obligations.

SLAs and other written agreements between the CSP and client should clearly identify the delineation of responsibilities between parties, including responsibilities for implementing and managing different security controls. These SLAs and agreements should be established as a prerequisite to any cloud service implementation. PCI DSS compliance validation and testing activities (with the associated controls, permissions, and schedules) should also be clearly detailed in the SLA.

Failure to develop and agree upon appropriate SLAs may result in issues for the client if the cloud service does not meet the needs and demands of their business. SLAs should be established and agreed as part of any contract and service negotiations. Performance, availability, integrity, and confidentiality should be considered and SLAs agreed for each service managed and/or operated by the CSP. Written agreements should also cover activities and assurances to be provided by both parties upon termination of the service provision.

6.1.4 Business Continuity Plans and Disaster Recovery

Organizational requirements for business-continuity plans (BCP) and disaster recovery (DR) apply to the client's outsourced environments as they do for client-managed facilities. Clients should consider whether the CSP's continuity and recovery procedures are sufficient to meet the client's organizational requirements, and the PCI DSS scope of the cloud service should include any fail-over sites and systems

that might be used to store, process, or transmit cardholder data in BCP or DR situation, The ability to perform tests of the BCP and DR capabilities and/or to observe results of the CSP's testing should also be considered.

6.1.5 Human resources

Management of the CSP's human resources is largely out of the control of the client. The client's due-diligence processes should include an understanding of the CSP's human resources and employee engagement practices, as inappropriately or under-qualified personnel may expose data to unnecessary risk. PCI DSS Requirement 12.7 provides a basis for assessing the recruitment process at the CSP.

6.2 Facilities and Physical Security

Cloud services are only "cloud" in concept. In reality, cloud services involve physical resources located at the CSP environment which are accessed remotely from the client's environment. Similarly to other third-party providers, CSPs of public and shared clouds provide services to multiple clients whose data and virtual components co-exist in the same physical location and on the same physical systems as other clients. Poor physical security controls at a CSP facility may expose many clients' data to unnecessary risk, and poor environmental controls may impact the performance and integrity of the service provision.

In a private cloud, the physical location of all components is known and can be verified. When using a public cloud, different elements of the environment, such as VMs, hypervisors, virtual network devices, etc., could be frequently relocated according to the CSP's load-balancing strategy. Verifying that appropriate physical security is in place can be challenging in an environment where data and infrastructure can be in multiple different locations at different times. A client should seek assurance that their physical security requirements are consistently applied across all potential locations.

6.3 Data sovereignty and Legal considerations

Depending on the deployment and service model adopted, and due to the dynamic nature of cloud operations, it may not be known where particular information actually resides. This may result in concerns over data ownership and potential conflicts between domestic or international legal and regulatory requirements. For example, the CSP's infrastructure may result in data traversing or being stored in politically or economically unstable countries.

Understanding the legal jurisdictions that apply to data in different countries or regions can be a challenge for the client organization. For example, clients subject to regional laws restricting cross-border flows of data will need to verify all locations and flows of their data to ensure their cloud service is compliant with their legal obligations.

Other legal considerations include requirements for electronic discovery, evidence preservation and integrity, and data custody. CSPs should have documented processes for responding to legal requests for seizure of records, including data/audit logs belonging to the CSP and their clients. Clients should understand the ramifications of such laws in the countries where their data exists, as well as the processes that their CSP will engage in.

6.4 Data Security Considerations

Further to the data-sovereignty considerations mentioned above, public-cloud providers often have multiple data storage systems located in multiple data centers, which may often be in multiple countries or regions. Consequently, the client may not know the location of their data, or the data may exist in one or more of several locations at any particular time. Additionally, a client may have little or no visibility into the controls protecting their stored data. This can make validation of data security and access controls for a specific data set particularly challenging.

It is recommended that data-security needs are evaluated for all types of information being migrated to a cloud environment, not only cardholder data. For example, operational data, security policies and procedures, system configurations and build standards, log files, audit reports, authentication credentials, cryptographic keys, incident response plans, and employee contact details are just some of the types of data with different security requirements that may need to be considered. If data security processes are not clearly defined and documented, the data may be unintentionally exposed or subject to unnecessary risk that could result in loss or inappropriate disclosure.

6.4.1 Data Acquisition

The client will ultimately determine how and when the cardholder data is acquired in the cloud environment. End-to-end processes and data flows must be documented across both client and cloud provider networks, so that it is clearly understood where cardholder data is located and how it is traversing the infrastructure (see PCI DSS Requirement 1.1.2). This will also help the client and CSP to identify where each entity acquires and relinquishes cardholder data throughout the process.

6.4.2 Data storage and persistence

In addition to the known range of intended storage locations, data may also be present in other CSP systems used for maintenance of the cloud infrastructure, such as VM images, backups, monitoring logs, and so on. Cardholder data stored in memory could also be written to disk for recovery or high availability purposes (for example, in the case of virtual machine suspension or snapshot). Such stored data may easily be “forgotten” and so not protected by data security controls. All potential capture points should be identified and managed as necessary to prevent unintended or unsecured storage or transmission of sensitive data. Specialized tools and processes may be needed to locate and manage data stored on archived, off-line, or relocated images.

Potential hypervisor access to data in memory should also be taken into consideration, to ensure that client-defined access controls are not unintentionally bypassed by CSP administrator personnel.

6.4.3 Data lifecycle management

For all cloud models, clear requirements for data retention, storage and secure disposal should form an integral part of the engagement process to ensure that sensitive data is:

- Retained for as long as needed,
- Not retained any longer than needed,
- Stored only in appropriate and secured locations,

- Accessible only to those with a business need, and
- Handled in accordance with the client's security policy

(See PCI DSS Requirements 3, 7, and 10.7)

Because all environments outside the client-controlled environment could potentially be untrusted, cloud services should support the secure transmission of cardholder data throughout the cloud infrastructure, between the client and cloud environments, between client environments, and between the cloud infrastructure and other public networks. It is recommended that sensitive data be encrypted for all transmissions through any cloud environment that is not entirely private and/or controlled by the client. Cloud environments outside of the client-controlled environment should be treated as "open" or "public" networks (see PCI DSS Requirement 4.1).

In a distributed cloud environment, verifying that all instances of cardholder data have been securely deleted in accordance with the client's data-retention policy is subject to the same challenges identified above for validating data security and access controls. Disposal of cardholder data must be conducted using secure methods in accordance with PCI DSS requirements, and all locations of cardholder data from within both the client and CSP environments need to be included. The disposal method should ensure that data is not recoverable upon completion of the disposal process.

6.4.4 Data Classification

Data classification and the management of data according to its classification will vary from organization to organization. A defined data-classification system can help organizations identify data that is sensitive or confidential, and data with specific security needs. This in turn allows organizations to assign appropriate protection mechanisms based on the security needs of different data types, and helps to prevent sensitive data from being inadvertently mishandled or treated as non-sensitive.

Organizations should ensure that their particular data security needs can be met by the cloud service before migrating that data into the cloud environment. Considerations should include how storing data types with different levels of sensitivity in the same virtual environment may impact the protection levels required for each data type. Cardholder data, user credentials and passwords, and cryptographic keys are examples of sensitive data that must be protected according to their individual needs.

6.4.5 Data Encryption and Cryptographic Key Management

In a public-cloud environment, one client's data is typically stored with data belonging to multiple other clients. This makes a public cloud an attractive target for attackers, as the potential gain may be greater than that to be attained from attacking a number of organizations individually. Strong data-level encryption should be enforced on all sensitive or potentially sensitive data stored in a public cloud. Because compromise of a CSP could result in unauthorized access to multiple data stores, it is recommended that cryptographic keys used to encrypt/decrypt sensitive data be stored and managed independently from the cloud service where the data is located. At a minimum, key-management servers should be located in a separate network segment and protected with separate access credentials from the VMs that are using the keys and the data encrypted with them.

Only defined, authorized key custodians should have access to cryptographic keys. Because access to both keys and encrypted data provides the ability to decrypt the data, clients will need to verify who has access to cryptographic keys, who has access to the encrypted data, and who has access to both. If a client shares encryption keys with the CSP, or engages the CSP as a key custodian, details of CSP access permissions and processes will also need to be reviewed and verified.

This consideration is particularly critical if cryptographic keys are stored or hosted by a third-party CSP that also hosts the encrypted data. If CSP personnel have access to a client's keys and the client's encrypted data, the client may have unintentionally granted the CSP ability to decrypt their sensitive data.

Any data that is decrypted in the cloud may be inadvertently captured in clear text in process memory or VMs via cloud maintenance functions (such as snapshots, backups, monitoring tools, etc.). To avoid this risk, clients may choose to keep all encryption/decryption operations and key management on their own premises, and use a public cloud only for storage of the encrypted data. Applicable controls must be applied to the encryption, decryption, and key-management processes to ensure data can only be retrieved (decrypted) by those who are authorized with a defined business need.

CSPs providing cryptographic-key management services for their clients should ensure that secret or private keys are not shared among clients—each client should be provided with a unique key or set of keys. Secure channels for access to the cloud environment also require proper key management. If the CSP uses images or cloning for protecting VMs, it is strongly recommended that keys not be cloned as part of the VM image—each clone or VM instance should have its own key(s)[7].

6.4.6 Decommissioning and Disposal

In addition to data disposal, resource decommissioning requirements should be defined to support clients' future decisions to migrate to a new CSP, decommission their cloud resources, or move out of a cloud environment altogether. The CSP should provide data-disposal mechanisms that provide assurance to the client that all data has been securely removed and deleted from the cloud environment. Procedures for "termination of service" should be clearly defined and documented

Clients may choose to ensure that all data is encrypted with strong cryptography to reduce the risk to any residual data left behind on CSP systems. However, clients should be aware that leaving potentially unknown quantities of encrypted data on CSP systems after their agreement has been terminated is likely to be a violation of their data-retention policy.

6.5 Technical Security Considerations

Technical security considerations for cloud environments generally include all those that apply to virtualization technologies, as well as those directly related to the different deployment and service models.

6.5.1 Evolving Security Technologies

As mentioned above, virtualization security considerations will also apply to cloud environments. There are many industry resources—including the *PCI DSS Virtualization Guidelines,* available from the PCI

[7] *The need for unique host keys has been well documented in a number of vendor technical papers*.

SSC website—that discuss security considerations for the use of virtual technologies. Some of these considerations include:

- It is difficult to maintain up-to-date, secure configurations on virtual machines when they are being activated and deactivated in rapid cycles —virtual machines that are dormant for any period of time may be improperly secured or introduce security vulnerabilities when activated.
- Security and monitoring solutions for virtual networks are still evolving and are not as mature as those available for physical networks.

Additionally, traditional security software and security device functions often do not scale well to a cloud environment. For example:

- Management of VM-to-VM traffic that does not pass through traditional network-based security controls may require the use of additional host-based security controls to monitor and control the traffic.
- Traditional agent-based software security solutions that are not designed for virtualized environments may cause operational issues. For example, software agents, such as those often used for anti-virus, each use a small percentage of memory and processing resources; this can result in a large overhead when multiple agents are installed on multiple VMs on the same host.
- Scheduled scans or updates occurring simultaneously across multiple VMs may result in an extreme load on the underlying system and reduce overall performance of all hosted VMs.

6.5.2 Identity and Access Management

Individual user identification and authentication for both CSP and client personnel is essential for access control and accountability (see PCI DSS Requirements 7 and 8). Shared credentials (such as user accounts and passwords) should not be used in the CSP environment—for example, for system administration and maintenance—nor should generic or shared accounts be assigned to or used by clients.

The use of a single client credential that covers multiple cloud services for that client is also a potential concern. For example, let's say a CSP issues a client with a user account and password that has administrator privilege in one environment and user-level privilege for a separate, unrelated cloud service. Compromise of the client's user-level account in the second environment could therefore result in the attacker gaining administrator-level access to the first environment. Client accounts and passwords should be unique for each service, and any account with elevated privilege (such as administrator) should be restricted for a specific service or function, and not used for activities or access that do not require such privilege.

6.5.3 Logging and Audit Trails

CSPs should be able to segregate log data applicable for each client and provide it to each respective client for analysis without exposing log data from other clients. Additionally, the ability to maintain an accurate and complete audit trail may require logs from all levels of the infrastructure, requiring involvement from both the CSP and the client. For example, the CSP could manage system-level, operating-system, and hypervisor logs, while the client configures logging for their own VMs and

applications. In this scenario, the ability to associate various log files into meaningful events would require correlation of client-controlled logs and those controlled by the CSP.

6.5.4 Hypervisor Access and Introspection

In large cloud environments it can be difficult to keep track of which hypervisors are running which VMs, as VMs can be dynamically assigned across a pool of hypervisors based upon load balancing needs. Hypervisor configuration and access is particularly important as it provides a single point of entry to all its VMs, and can potentially be used to gain access to sensitive data and resources on separate VMs.

An additional consideration is the degree to which the hypervisor is used to deliver security functionality to the VMs. For example, a simple, hardened hypervisor may be very secure but offer limited security capabilities, whereas a more complex, security-capable hypervisor with improved functionality could potentially present a greater risk if compromised[8].

Functionality that allows the hypervisor to control and monitor individual VM activity from outside the VMs is known as introspection[9]. Hypervisor introspection expands the functionality of the hypervisor to allow a deeper analysis of the data being processed by the VM, and typically includes visibility into stored data files as well as monitoring of network traffic, memory and program execution, and other elements of the VM.

Depending on the particular technology implemented, introspection can provide the CSP with a level of real-time auditing of VM activity that may otherwise be unattainable. This can help the CSP to monitor for and detect suspicious activity within and between VMs. Additionally, introspection may facilitate cloud-efficient implementations of traditional security controls—for example, hypervisor-managed security functions such as malware protection, access controls, firewalling and intrusion detection between VMs.

Two potential challenges with introspection are that it can bypass role-based access controls and it can be used without leaving a forensic audit trail within the VM itself. For example, to view a data file, a user typically authenticates to the VM, resulting in an authentication audit trail and ensuing the user's access is controlled according to their defined permissions. If file-access logging is enabled in the VM and the user views a file, the access is recorded to show what was accessed by whom and when.

With introspection, files can be accessed from within the privileged state of the hypervisor. As no authentication to the VM itself is required, file access leaves no audit trail on the VM and the VM contains no evidence that the file was accessed. In this example, the access would need to be logged via the introspection tool itself, which would typically not be in the client's control. While this may be less of an issue within a private cloud environment, it is an important consideration for clients of public cloud services.

Additionally, since introspection is designed to have full visibility into each VM, it may be difficult to restrict such access to only specific files or programs in memory. Any personnel (for example, CSP employees or possibly other hosted clients) with access to the introspection function could potentially have access to data and processes on any VM running on that hypervisor. Introspection access must therefore be carefully managed, controlled and monitored to ensure that role-based access and segregation of duties

[8] *Don't Bloat the Hypervisor! What to know about Introspection* (Webinar), Tim Mather 2011.
[9] *A Virtual Machine Introspection Based Architecture for Intrusion Detection,* Tal Garfinkel, Mendel Rosenblum 2002.

is maintained. For example, the ability to configure introspection auditing should not be available to personnel with the ability to access hosted VMs via the introspection tool.

CSPs using introspection should be able to provide their clients with all applicable introspection logs for that client's environment including, but not limited to, authentication details, disk and memory access requests, and API calls. All introspection activity should be mapped to the individual user account performing the activity, and logs should be reviewed on a continual basis to ensure the integrity and confidentiality of client data has been maintained.

Where introspection is used by a third-party CSP, the clients may wish to consider implementing data-level security controls (such as strong cryptography with all key storage and encryption/decryption operations external to the cloud service) to avoid exposing sensitive data to the enhanced monitoring features that introspection provides.

6.5.5 Security of Interfaces and APIs

Application programming interfaces (APIs) and other software interfaces are an integral component of cloud computing, supporting interoperability and rapid delivery of cloud services. APIs can be configured to provide access to a variety of functions, allowing clients and CSPs to interact and manage their interactions within the cloud service.

As web services and APIs are by nature publicly accessible, their security is critical to the security of the resources they provide access to. If not properly developed, managed, and secured, these interfaces can be exploited or compromised, resulting in unexpected behavior and potentially unauthorized access. For example, a poorly-coded API could result in weak authentication protocols, poor access controls, and limited auditing capability. Such weaknesses could lead to the exposure of authentication credentials and other sensitive data. If the APIs are not properly secured, they could also be exploited or altered by an attacker to redirect data flows or alter application behavior.

APIs and other public interfaces should be designed to prevent both accidental misuse and malicious attempts to bypass security policy. Strong authentication and access controls, strong cryptography, and real-time monitoring are examples of controls that should be in place to protect these interfaces.

6.5.6 Security of Client Systems

Client systems used to access the cloud environment should not be overlooked, as they could potentially become a weak point in a client's cloud security strategy. Client organizations need to ensure their systems and internal processes do not provide unauthorized access to the cloud environment. For example, if a client workstation or other device is compromised, an attacker may be able to use legitimate credentials and an authorized channel to gain access to the cloud environment from the compromised client system. The client will therefore need to ensure their client-side devices are appropriately secured and protected from unauthorized physical and logical access. Client-side systems used to access cardholder data in the cloud would also be in scope for all applicable PCI DSS requirements.

6.5.7 Multi-tenancy

In a multi-tenant cloud environment, client organizations generally have no knowledge of the other clients with whom they share resources (for example, virtual infrastructure, data stores, etc.), or how other clients are securing (or not securing) their environments that access the shared resources.

Whether unsavory clients can pose a risk to other clients using the same provider will largely depend on the controls the CSP has in place to segregate clients from one another, and to monitor and detect suspicious activity on the shared infrastructure and between client environments. Before engaging with a CSP, clients should consider how the CSP verifies that their clients are who they say they are, and how the CSP detects potentially suspicious behavior once the clients are onboard. Clients should also ask the CSP what controls they have in place to ensure that the security posture of one client cannot affect the security posture of another client.

6.6 Incident Response and Investigation

Clients need to know when an issue, incident, or breach has occurred and the impact to their environment and/or to their data. Issues, incidents, and data breaches should be communicated by the CSP in a timely manner. Clients should consider whether their CSP requires all clients to immediately notify the CSP of potential breaches in their environments, allowing the CSP to respond more quickly to contain the breach and minimize its impact to other clients.

Definitions of what constitutes a breach or incident requiring notification between client and cloud provider should be agreed. Notification processes and timelines should be included in SLAs, and incident response plans should include notification requirements. The potential for client data to be captured by third parties during a breach investigation should also be clearly understood.

Investigating potential breaches in cloud environments brings additional challenges. For example, compromised VM instances may be deactivated before anyone is aware that a breach occurred. It may be nearly impossible to properly investigate a breach when the source of the breach is no longer in use or even exists.

7 Conclusion

In addition to the business and risk considerations, the implementation of security controls in a cloud environment may require specialized technical knowledge and skills. It is therefore crucial that prior to migrating payment card operations into a cloud environment, an organization engages their technical, legal, due diligence, information security, and compliance teams to work together to define the client's needs and evaluate potential cloud service offerings against those needs.

Regarding third-party or public clouds, clients should consider that while they can outsource the day-to-day operational management of the data environment, they retain responsibility for the data they put in the cloud. Clients are encouraged to "shop around" until they find a CSP who can provide the level of security and assurance they require. Potential clients are encouraged to:

- UNDERSTAND your risk and security requirements first.
- CHOOSE a deployment model that aligns with your security and risk needs.
- EVALUATE different service options.
- KNOW what you want from your CSP.
- COMPARE providers and service offerings.
- ASK questions of the CSP and verify the responses, for example:
    - What does each service consist of exactly, and how is the service delivered?
    - What does the service provide with respect to security maintenance, PCI DSS compliance, segmentation, assurance, and what is the client responsible for?
    - How will the CSP provide ongoing evidence that security controls continue to be in place and are kept up to date?
    - What will the CSP commit to in writing?
    - Are other parties involved in the service delivery, security, or support?
- DOCUMENT everything with your provider in written agreements—for example, SLAs / Terms of Service contracts, etc.
- REQUEST written assurances that security controls will be in place and maintained.
- REVIEW the service and written agreements periodically to identify if anything has changed.

CSPs are encouraged to work with their clients to understand their security and compliance needs. Both parties should be willing to maintain open communication and monitoring to avoid any misunderstandings or gaps in security responsibilities.

Appendix A: Sample PCI DSS Responsibilities for Different Service Models

This Appendix expands on Figure 3 (in Section 4), and provides examples of how responsibilities for PCI DSS requirements may be shared between clients and CSPs across the three service models. There will of course be exceptions and variations across each individual service, and this table is provided as a guideline for clients and CSPs to help plan discussions and negotiations.

The descriptions in this table are intended to reflect the CSP's responsibilities with respect to the services it provides, and do not consider the CSP's responsibilities for its internal infrastructure and operations not directly involved in providing services to their clients. Similarly, client responsibilities do not include consideration for the client systems used to access the cloud service, or for any client systems in scope for PCI DSS that are outside of the cloud service.

PCI DSS Requirements	Common Considerations	Example responsibility assignment for management of controls		
		IaaS	PaaS	SaaS
Requirement 1: *Install and maintain a firewall configuration to protect cardholder data.*	**IaaS:** Typically, network security is a shared responsibility: the client is responsible for securing networks within and between their own environments, while the CSP provides network security at the cloud perimeter and between the CSP's clients. The CSP manages firewalls on the CSP-managed network and any infrastructure firewalls not visible to the cloud customer. Any firewalls above the infrastructure layer may be the responsibility of the cloud customer. The CSP-managed firewalls could also be shared by multiple cloud customers. **PaaS:** Firewalls above the infrastructure layer may be the responsibility of the cloud customer or CSP. The cloud customer could be directly responsible for implementing and managing firewalls on the provided platform, and/or they may define firewall configurations, which the CSP then implements for the customer's environment. The CSP-managed firewalls could also be shared with other cloud customers. **SaaS:** The network is wholly owned and managed by the CSP, and consequently, all firewall functions are typically managed by the CSP. **In all scenarios,** the client may still need to define, approve and/or periodically review the services, protocols, and ports permitted into their environment, even if the CSP is managing the firewalls in question.	**Client and CSP**	**Client and CSP**	**CSP**

PCI DSS Requirements	Common Considerations	Example responsibility assignment for management of controls		
		IaaS	PaaS	SaaS
Requirement 2: *Do not use vendor-supplied defaults for system passwords and other security parameters*	**IaaS:** Secure configuration of OS and applications is typically responsibility of the client while secure configuration of underlying devices is the responsibility of CSP. There may also be virtual devices that customer is responsible for maintaining. **PaaS:** The OS is often controlled by CSP but some services may include a level of client access to OS—both parties will need to clarify which entity is applying secure configuration and hardening at the OS level. Applications and software above the OS are likely to be controlled by the client. Secure configuration of network devices will be managed by the CSP. **SaaS:** The CSP typically manages configuration and hardening of all devices, OS and applications.	**Client and CSP**	**Client and CSP**	**CSP**
Requirement 3: *Protect stored cardholder data*	**IaaS and PaaS:** The client is generally responsible for the manner in which information is secured (such as the use of encryption mechanisms) and in what format—for example, flat files, databases entries etc. Physical locations of the information stores might be unknown to the client, and storage locations may need to be identified. Data retention is defined by the client; however the CSP controls the actual storage areas. The use of controls to prevent unintended or additional retention (for example, via snapshots, backups, etc.) also needs to be considered. **SaaS:** Typically controlled and managed by the CSP as part of the predefined service. The CSP may also define the retention periods. Clients may have very little to no control over how or where their data, including CHD, is stored.	**Client and CSP**	**Client and CSP**	**CSP**

PCI DSS Requirements	Common Considerations	Example responsibility assignment for management of controls		
		IaaS	PaaS	SaaS
Requirement 4: *Encrypt transmission of cardholder data across open, public networks*	**IaaS and PaaS:** Mechanisms for transmission are typically controlled by the client while the underlying technology is managed by the CSP; however, this will depend on the technologies in use. Controls to prevent unintended transmission of data outside of client environment are generally maintained by the CSP, depending on the particular service. The client should be aware of how data is transmitted between components in order to ensure data is encrypted for all transmissions over non-private channels. This may include transmissions within the client's own environment (for example, between client VMs). **SaaS:** The CSP retains full control over transmission mechanisms. The client has little to no control over how or where data is transmitted within the cloud environment. The client is responsible for ensuring clear-text data is not passed to the CSP for transmission to public networks or untrusted environments (such as other cloud clients).	**Client**	**Client and CSP**	**CSP**
Requirement 5: *Use and regularly update anti-virus software or programs*	**IaaS:** Protection of the OS and client VMs is typically the responsibility of client. Anti-virus updates apply to the host OS as well as any VMs in the client environment running their own OS. There may also be virtual devices that the client is responsible for keeping up to date. Anti-malware protection for underlying devices/infrastructure remains the responsibility of the CSP. **PaaS:** Generally managed by whoever controls the OS; some PaaS services include client responsibility for OS maintenance. Anti-virus updates will apply to the underlying OS as well as any VMs in the client environment running their own OS. **SaaS:** The CSP typically manages the security and anti-virus for the environment.	**Client**	**Client and CSP**	**CSP**

PCI DSS Requirements	Common Considerations	Example responsibility assignment for management of controls		
		IaaS	PaaS	SaaS
Requirement 6: *Develop and maintain secure systems and applications*	**IaaS:** Patching and maintenance of OS and applications are typically the responsibility of the client, while patching and maintenance of underlying devices remains the responsibility of the CSP. There may also be virtual devices that the client is responsible for maintaining. Secure coding is typically the client's responsibility (they may either use their own applications or choose secure commercial applications). **PaaS:** Patching and maintenance of underlying devices remains the responsibility of the CSP. OS patching and maintenance may also be controlled by CSP; however, some PaaS services include client responsibility for OS maintenance—entities will need to determine which party is responsible for applying patches/updates. If the CSP provides patching, the client should verify that patches are deployed in a timely manner. Patching of applications is typically managed by the client, depending on the service and agreements. Secure coding of application is the responsibility of whoever develops/controls the applications, which may be either the client or the CSP, and may vary for different applications. **SaaS:** The client may control or manage the APIs or they may share responsibility with the CSP. The CSP typically manages patching and updates of all devices, OS, and applications, and is also responsible for secure coding of software; however, the client should verify that patches are deployed in a timely manner.	**Client and CSP**	**Client and CSP**	**Client and CSP**
Requirement 7: *Restrict access to cardholder data by business need to know*	**IaaS and PaaS:** Generally the client is responsible for defining access to their data files. Physical location of the information stores might be unknown to the client and may need to be identified. The CSP controls the physical storage areas, and CSP-managed access controls are often cumulative to the client-defined controls. The use of controls to prevent unintended access to data (for example, to data captured via snapshots, backups, etc.) should also be considered. **SaaS:** The client defines data access needs for their own personnel; however, access to data is ultimately controlled by CSP.	**Client and CSP**	**Client and CSP**	**Client and CSP**

PCI DSS Requirements	Common Considerations	Example responsibility assignment for management of controls		
		IaaS	PaaS	SaaS
Requirement 8: *Assign a unique ID to each person with computer access*	**IaaS and PaaS:** The client is responsible for ensuring all accounts under their control use unique IDs and strong authentication. The CSP is responsible for ensuring strong authentication is used for the underlying infrastructure. Compared to the IaaS model, the CSP retains significant administrative access rights In SaaS and PaaS models. **SaaS:** The CSP has ultimate control of accounts at all levels. Depending on the particular service, the client may have the ability to create user-level accounts within the application or service, or they may be assigned user accounts that the CSP maintain on their behalf. The client is responsible for ensuring all the accounts they use have strong passwords.	**Client and CSP**	**Client and CSP**	**Client and CSP**
Requirement 9: *Restrict physical access to cardholder data*	**All service models:** Generally managed by the CSP for all service models. The client rarely has any physical access to cloud systems; and the CSP might not permit onsite visits or client audits. This will depend on the particular CSP as well as the distribution of data across different locations; the clients may not know which location houses their data.	**CSP**	**CSP**	**CSP**
Requirement 10: *Track and monitor all access to network resources and cardholder data*	**IaaS and PaaS:** The CSP typically manages monitoring and logging for underlying devices and infrastructure, including hypervisors, while the client is responsible for monitoring and logging within their own virtual environments. The ability to associate various log files in order to reconstruct events may require correlation between client-controlled logs and those controlled by the CSP. Some monitoring activities may be built into the service agreement for the CSP to manage on behalf of clients. Details of what data will be captured and what will be made available to the client will need to be defined. **SaaS:** The client typically relies on the CSP for all monitoring and logging, but may have limited application-level logging such as user logon/logoff, account management, and basic reporting.	**Client and CSP**	**Client and CSP**	**CSP**

PCI DSS Requirements	Common Considerations	Example responsibility assignment for management of controls		
		IaaS	PaaS	SaaS
Requirement 11: *Regularly test security systems and processes*	**IaaS and PaaS:** Testing is generally managed by whoever has control of the particular aspect of the environment. However, CSPs may prohibit client testing, in which case clients may need to rely on the CSP. If the CSP is performing scans, the client needs to verify which instances/VMs are covered. IDS/IPS may not be provided by the CSP. Generally the client can use FIM to monitor their own virtual environments (including data, applications, and logs), while monitoring of system/device files is managed by the CSP. **SaaS:** The client doesn't have visibility or permission to perform scans and typically relies on the CSP for all scans, testing, and monitoring.	**Client and CSP**	**Client and CSP**	**CSP**
Requirement 12: *Maintain a policy that addresses information security for all personnel*	**All service models:** While the CSP and client may define agreed-upon procedures (for example, in the SLA), each party maintains their own security policies and internal procedures. Defined roles and responsibilities, training, and personnel security requirements are the responsibility of each party for their respective personnel. Clients should ensure that the CSP policies and procedures are appropriate for the client's risk and security needs. Incident response in particular requires awareness and coordination between both parties.	**Client and CSP**	**Client and CSP**	**Client and CSP**
Appendix A: *Additional PCI DSS Requirements for Shared Hosting Providers*	The requirements for shared hosting providers to ensure separation between clients apply to third-party provided cloud services.	**CSP**	**CSP**	**CSP**

Appendix B: Sample Inventory

This appendix provides an example inventory for components used in cloud environments. Use of an inventory can help to identify the types of components involved in delivery of the service and responsibly for securing them. This example is not intended to be applicable to any particular scenario, and is intended to provide a starting point for scoping discussions between clients and CSPs.

When preparing an inventory, consider the following:

- The type of information collected should be relevant for the client's business needs as well as the CSP's
- The level of detail collected should be appropriate for both parties to reach a clear understanding of the components involved, their use, and who manages/secures them.

Type/Layer	Component Description/Purpose	Type of Component	Number of components	Implementation Notes	Responsibility for securing component
Note: Actual layers will vary depending on structure of CSP service offerings	*For example: Firewall, OS, application, web server, hypervisor, router, database, etc.*	*For example: Is component physical, logical or virtual? Static or dynamic?*	*Number of components used in relation to client's service*	*Defined usage, location, etc., as applicable*	*For example: CSP only, client only, or shared*
Data					
Interfaces – APIs/GUIs					
Applications					
Programming stack					
Operating Systems					
Virtual Machines					
Virtual networking					
Hypervisors					

Type/Layer	Component Description/Purpose	Type of Component	Number of components	Implementation Notes	Responsibility for securing component
Note: Actual layers will vary depending on structure of CSP service offerings	*For example: Firewall, OS, application, web server, hypervisor, router, database, etc.*	*For example: Is component physical, logical or virtual? Static or dynamic?*	*Number of components used in relation to client's service*	*Defined usage, location, etc., as applicable*	*For example: CSP only, client only, or shared*
Processing/Memory					
Data Storage					
Network devices					
Physical facilities					

Note: *This is intended as a general example only. It may be necessary to reorganize the different technology layers or define additional component characteristics as applicable to a particular environment. Additionally, entities may wish to identify responsibilities for each system component in greater detail than provided for here.*

Appendix C: Sample PCI DSS Responsibility Matrix

A PCI DSS responsibility matrix may help to clarify and confirm how responsibilities for maintaining PCI DSS requirements are shared between the client and CSP. Responsibilities should always be defined in written agreements.

Considerations for each PCI DSS requirement include:

- Does the CSP perform/manage/maintain the required control?
- How is the control implemented, and what are the supporting processes? (E.g., process for patch updates would include details of testing, scheduling, approvals, etc.)
- What layers of the cloud architecture are covered by the CSP for the requirement? What layers of the architecture are not covered by the CSP and are specifically the responsibility of the client?
- How will the CSP provide ongoing assurance and/or evidence to the client that controls are met? (For example, periodic reports, real-time notifications, results of testing, etc.)

PCI DSS Requirement	**Responsibility** (CSP only, client only, or shared)	**Specific coverage/ scope of client responsibly**	**Specific coverage/ scope of CSP responsibility**	**How and when CSP will provide evidence of compliance to Client**
1.1 Establish firewall and router configuration standards that include the following:				
1.1.1 A formal process for approving and testing all network connections and changes to the firewall and router configurations				
1.1.1 A formal process for approving and testing all network connections and changes to the firewall and router configurations				
1.1.2 Current network diagram with all connections to cardholder data, including any wireless networks				

PCI DSS Requirement	Responsibility (CSP only, client only, or shared)	Specific coverage/ scope of client responsibly	Specific coverage/ scope of CSP responsibility	How and when CSP will provide evidence of compliance to Client
1.1.3 Requirements for a firewall at each Internet connection and between any demilitarized zone (DMZ) and the internal network zone				
1.1.4 Description of groups, roles, and responsibilities for logical management of network components				
1.2 Build firewall and router configurations that restrict connections between untrusted networks and any system components in the cardholder data environment. *Note: An "untrusted network" is any network that is external to the networks belonging to the entity under review, and/or which is out of the entity's ability to control or manage.*				
1.2.1 Restrict inbound and outbound traffic to that which is necessary for the cardholder data environment.				
1.2.2 Secure and synchronize router configuration files.				
…And so on.				

Note: *This is intended as an example only to provide a starting point for discussions between clients and CSPs. It is not intended as a requirement or an extension of PCI DSS compliance responsibilities. However, it may provide a useful tool to help to clarify responsibilities in agreements between clients and providers.*

Appendix D: PCI DSS Implementation Considerations

The questions in this appendix are intended as suggestions to help start conversations between clients and CSPs in order to understand the characteristics of a particular cloud environment, which may in turn help determine if and how PCI DSS requirements can be met in that environment. These questions alone will not determine whether or not applicable PCI DSS requirements can be met, however, they may be a useful addition to questions directly related to specific PCI DSS requirements.

Information in this table incorporates guidance from the following sources:

- CSA Consensus Assessments Initiative Questionnaire
- ENISA Information Assurance Requirements

Please also refer to the PCI DSS Virtualization Guidelines for additional PCI DSS considerations for virtualization technologies.

PCI DSS Requirement	Considerations for Cloud Environments

Build and Maintain a Secure Network
 	
1: Install and maintain a firewall configuration to protect cardholder data
 	
2: Do not use vendor-supplied defaults for system passwords and other security parameters
	▪ How is separation between tenants assured?
	▪ How are boundaries enforced between trusted (internal to the client) networks and untrusted networks (such as CSP, other client, or public-facing networks)?
	▪ Are physical or virtual firewalls used?
	▪ Who manages and audits firewall configurations?
	▪ How are changes to firewall and network configurations tracked and managed?
	▪ What technologies are used in the provision of the cloud service—e.g., hardware, software, virtual technologies?
	  o Is there a current list of all hardware and software components in the environment?
	  o Can the actual components used by a particular client be identified?
	▪ How are configuration standards assured on different components of the infrastructure?
	  o Are API interfaces standardized?
	  o What is the process for provisioning new components?
	  o Are virtual images hardened before being enabled?
	  o Are hardened images protected from unauthorized access?
	▪ How are systems with high security classifications segregated from systems with low security classifications?
	▪ How are shared resources (such as processing, memory, and storage) managed to ensure they cannot be manipulated—for example by overloading—in order to gain access to other client environments or data?

PCI DSS Requirement	Considerations for Cloud Environments
Protect Cardholder Data *3: Protect stored cardholder data* *4: Encrypt transmission of cardholder data across open, public networks*	▪ Where are the "known" data storage locations? Where are data centers located? ▪ Which legal jurisdiction(s) applies to client data? ▪ Does the CSP have any business, legal or regulatory requirements that could impact retention of client data? ▪ How is access to client data restricted to only that client's users and applications? ▪ How are VM images, snapshots, and backups managed to prevent unnecessary capture of sensitive data? ▪ How is data securely deleted from memory and stored images? Will data remnants exist in terminated VMs? ▪ If cryptographic keys are provided by CSP, are unique keys generated for each client? ▪ Where are encryption/decryption processes being performed? Who controls each process? ▪ Where are cryptographic keys stored, and who controls the keys? Are data-encryption keys stored and managed separately from the data they protect? ▪ Where is encrypted data stored, and who has access to the keys and encrypted data? ▪ How is security and access defined for the virtualized resources used for generation of cryptographic keys? ▪ What process is followed in the event of a suspected key compromise? ▪ Is all client data securely purged from all CSP systems upon termination of the agreement? ▪ How are communications secured between client and other environments? How are communications secured within the cloud itself? ▪ Are APIs configured to enforce strong cryptography and authentication? ▪ Is mutual authentication implemented between CSP and client systems?
Maintain a Vulnerability Management Program *5: Use and regularly update anti-virus software or programs* *6: Develop and maintain secure systems and applications*	▪ Are VMs protected from within the VM or from the hypervisor? ▪ How are VM images (including inactive and replicated VMs) ensured to have up-to-date anti-malware and patches before they are enabled for use? ▪ How are patches managed (e.g., prioritized, tested, approved, and deployed), for both underlying CSP systems and provisioned client environments? ○ What is the process for each layer of the cloud service—e.g., physical network devices, host operating systems, hypervisors, virtualized components (including VMs, virtual network devices), applications, etc.? ▪ How are APIs and web services protected from vulnerabilities? ▪ Are standardized interfaces and coding languages used? ▪ How are development/test systems and data prevented from being inadvertently migrated into production environments, and vice versa (e.g., through virtual replication, imaging in or snapshot mechanisms)?

PCI DSS Requirement	Considerations for Cloud Environments

Implement Strong Access Control Measures
 	
7: Restrict access to cardholder data by business need to know
 	
8: Assign a unique ID to each person with computer access
 	
9: Restrict physical access to cardholder data
	▪ How is user authentication applied at different levels?
	▪ How are layers of access controls managed to ensure the aggregate access is not more than intended?
	▪ Which CSP personnel have ability to access client data?
	▪ How are CSP privilege assignments reviewed and monitored?
	▪ How is segregation of duties maintained (for example, between administrative and auditing functions)?
	  ○ Is administrative access to systems or hypervisor separate from access to client VMs and data stores?
	  ○ Are separate credentials used for different security functions?
	▪ How are least-privilege and need-to-know determined for CSP personnel?
	▪ How are credentials de-provisioned?
	  ○ Does de-provisioning apply across all geographically distributed locations?
	  ○ Could de-provisioned credentials be retained in offline images?
	▪ Is remote access for CSP personnel permitted from untrusted networks?
	▪ Are controls in place to prevent the capture of passwords in active memory, and to ensure that virtualized images do not contain authentication credentials?
	▪ Is two-factor authentication required for client access?
	▪ Does the CSP use any shared passwords (e.g., for maintenance)?
	▪ Does the CSP maintain direct ownership and control over all data storage systems and facilities?
	▪ Who has physical access to data centers and systems?
	▪ How are data-storage systems protected from physical or direct console access?
	▪ How are backups of VMs and data secured?
	▪ How is physical media inventoried, secured, monitored, and tracked?
	▪ Is media re-used? How is data permanently removed from end-of-life or reusable media?

PCI DSS Requirement	Considerations for Cloud Environments

Regularly Monitor and Test Networks
 	
10: Track and monitor all access to network resources and cardholder data
 	
11: Regularly test security systems and processes
	▪ How are activities traced back to individual client personnel or individual CSP personnel?
	▪ Can the specific system components used by a client at a particular time be identified?
	▪ What types of events are recorded in audit logs?
	▪ How are audit logs correlated between client environments (such as a VM image) and CSP infrastructure (such as the hypervisor or underlying system)?
	▪ How are audit logs monitored and reviewed?
	▪ How are clocks synchronized between virtual instances and underlying systems/hardware?
	▪ How is testing for wireless technologies performed and managed?
	▪ How are all variations of VM images (including inactive VMs) scanned for vulnerabilities?
	▪ What defenses are in place to protect against 'internal' attacks (originating from CSP's or other client network) and "external" attacks (originating from the Internet or other public network)?
	▪ Is penetration testing performed across different layers of the environment (e.g., between VMs and the CSP's management network, or between clients on shared infrastructure)?
	▪ How is security testing managed for CSP infrastructure vs. client environments?
	  ○ What testing are clients allowed to perform on their internet-facing systems?
	  ○ How are clients prevented from performing penetration testing on other clients' environments?

PCI DSS Requirement	Considerations for Cloud Environments

Maintain an Information Security Policy
 	
12: Maintain a policy that addresses information security for all personnel
	▪ How does the CSP identify potential risks?
	▪ Are clients notified upon changes to the CSP's security and/or privacy policies?
	▪ Does the CSP have mechanisms in place to ensure secure operational procedures are followed?
	▪ How does the CSP screen personnel?
	  o Are different levels of screening used for different roles or regions?
	  o Does screening cover all personnel with physical access to data centers at all locations?
	▪ Does the CSP outsource any aspect of the cloud service to other providers (e.g., data storage, security services, etc.)?
	▪ What measures are taken to ensure that the CSP's security policies are maintained by their third-party providers?
	▪ What processes are in place to detect, assess, escalate, and respond to potential breaches?
	  o What mechanisms are in place for clients to report a suspected breach?
	  o What criteria are used to define whether an "incident" or a "breach" has actually occurred?
	  o What notifications are provided and when?
	  o How would a breach at one client impact other clients on the same infrastructure?
	  o How is evidence collected, managed, and shared?
	▪ What happens to client data in the event of a breach to CSP systems?
	▪ Can a client's data be collected as part of another client's (or CSP's) breach investigation (either by authorities or third party investigators)?
	▪ Are disaster-recovery processes, systems and facilities implemented with the same security controls as production environments?

Appendix A: *Additional PCI DSS Requirements for Shared Hosting Providers*	▪ How is isolation maintained across different layers, including between virtual machines, physical machines, networks, storage systems (e.g., storage area networks), management networks and support systems? ▪ What controls are in place to prevent data leakage between clients, and between client and CSP? ▪ Are resource-isolation mechanisms in place?

Acknowledgements

PCI SSC would like to acknowledge the contribution of the Cloud Special Interest Group (SIG) in the preparation of this document. The Cloud SIG consists of representatives from the following organizations:

Accuvant, Inc.
Acertigo AG
Anitian Enterprise Security
Assurant, Inc.
AT&T Consulting Inc.
Bank of America – Global Information Security – PCI Adherence
BrightLine CPAs & Associates Inc.
BT Plc.
Capita Plc.
Citi
CloudPassage
Coalfire
Computer Services, Inc.
Comsec Consulting Ltd.
Crowe Horwath LLP
DataPipe
Deloitte New Zealand
Diebold, Inc.
Direct Insite Corp.
Dollar Thrifty Automotive Group, Inc.
Equinox Payments, LLC
FishNet Security, Inc.
Fort Consult A/S (Denmark)
Habif, Arogeti & Wynne LLP
HP Information Security UK Ltd
Hytrust, Inc
Information Risk Management Plc.
Integralis Ltd
International Cards Processing Services
Internet Security Auditors
IOActive, Inc.
IQ Information Quality
Kilrush Consultancy Ltd.
Kingston Smith Consulting LLP
Knowit Secure AB
KPMG, LLP
Layered Tech
Liaison Technologies Inc.
McGladrey LLP
Market America, Inc.
Merchant Link
Microsoft
Nationwide Building Society
PathMaker Group
PayPros
PricewaterhouseCoopers LLP
Privity Systems Inc.
Protiviti
Retalix
RightScale
Royal Bank of Scotland
SecurityMetrics, Inc.
Sense of Security Pty Ltd.
Specsavers
Structured Communications Systems, Inc.
Tesco Plc.
Thales e-Security
Trend Micro Inc.
TrustNet
Trustwave
The UK Cards Association
University of North Carolina at Chapel Hill
Vanguard Integrity Professionals
Venda
VendorSafe Technologies, LLC.
VeriFone UK & Ireland Services Ltd.
Verizon Enterprise Solutions
Verizon Wireless
WEX Inc.
Woolworths Limited

References

This document draws the following additional sources of reference. These sources are recommended as additional guidance on securing cloud-computing environments.

Source[10]	Reference
Cloud Security Alliance (CSA) https://cloudsecurityalliance.org/	▪ *Security Guidance for Critical Areas of Focus in Cloud Computing v3.0* ▪ *How to do PCI DSS in the Cloud* (Courseware) ▪ *Top Threats to Cloud Computing V1.0* ▪ *Consensus Assessments Initiative Questionnaire v1.1* ▪ *Hypervisor vs. Host Based Security*
European Network and Information Security Agency (ENISA) http://www.enisa.europa.eu/	▪ *Cloud Computing – Benefits, risks and recommendations for information security*
National Institute of Standards and Technology (NIST) http://csrc.nist.gov/publications/	▪ *The NIST Definition of Cloud Computing (S P 800-145)* ▪ *Cloud Computing Synopsis and Recommendations (SP 800-146)* ▪ *Guidelines on Security and Privacy in Public Cloud Computing (SP 800-144)* ▪ *NIST Cloud Computing Reference Architecture (SP 500-292)*
Information Commissioners Office (ICO) http://www.ico.gov.uk/	▪ *Guidance on the use of cloud computing*
ISACA http://www.isaca.org/	▪ *Cloud Computing Management Audit/Assurance Program* ▪ *Meeting PCI DSS When Using a Cloud Service Provider*
ISC2 International Information Systems Security Certification Consortium, Inc., https://www.isc2.org/	▪ *Security in the Skies – Cloud computing security concerns, threats, and controls*
The Open Web Application Security Project (OWASP) https://www.owasp.org	▪ OWASP Cloud-10 Project

Cloud Computing Security Research Library
http://searchcloudsecurity.com
	▪ Robert Zigweid Webinars:
	  ○ *Collision Course: PCI Data and the Cloud*
	  ○ *Security Role Play: Merchants and Cloud Service Providers*
	  ○ *Tools to Assess PCI Compliance in Cloud*
	▪ *Technical Guide on Compliance and Cloud Security*

PCI SSC https://www.pcisecuritystandards.org	▪ *PCI DSS Virtualization Guidelines*
Webinar / various sources	▪ *Don't Bloat the Hypervisor. What to know about Introspection* (Webinar - https://www.brighttalk.com/webcast/288/24370)

[10] Links to third party websites are subject to change

About the PCI Security Standards Council

The PCI Security Standards Council is an open global forum that is responsible for the development, management, education, and awareness of the PCI Data Security Standard (PCI DSS) and other standards that increase payment data security. Founded in 2006 by the major payment card brands American Express, Discover Financial Services, JCB International, MasterCard Worldwide, and Visa Inc., the Council has over 600 Participating Organizations representing merchants, banks, processors, and vendors worldwide. To learn more about playing a part in securing payment card data globally, please visit: pcisecuritystandards.org.