Security Standards Council™

Standard: PCI Data Security Standard (PCI DSS)

Version: 2.0

Date: June 2011

Author: Virtualization Special Interest Group
PCI Security Standards Council

Information Supplement:
PCI DSS Virtualization Guidelines

Table of Contents

1 Introduction

Virtualization separates applications, desktops, machines, networks, data and services from their physical constraints. Virtualization is an evolving concept, encompassing a broad range of technologies, tools, and methods, and can bring significant operational benefits to organizations that choose to leverage them. As with any evolving technology, however, the risks also continue to evolve and are often less understood than risks associated with more traditional technologies.

The intent of this Information Supplement is to provide guidance on the use of virtualization in accordance with the Payment Card Industry Data Security Standard (PCI DSS). For the purposes of this paper, all references are made to the PCI DSS version 2.0.

There are four simple principles associated with the use of virtualization in cardholder data environments:

a. If virtualization technologies are used in a cardholder data environment, PCI DSS requirements apply to those virtualization technologies.

b. Virtualization technology introduces new risks that may not be relevant to other technologies, and that must be assessed when adopting virtualization in cardholder data environments.

c. Implementations of virtual technologies can vary greatly, and entities will need to perform a thorough discovery to identify and document the unique characteristics of their particular virtualized implementation, including all interactions with payment transaction processes and payment card data.

d. There is no one-size-fits-all method or solution to configure virtualized environments to meet PCI DSS requirements. Specific controls and procedures will vary for each environment, according to how virtualization is used and implemented.

1.1 Audience

This Information Supplement is intended for merchants and service providers who use or are considering use of virtualization technologies in their cardholder data environment (CDE). This may also be of value for assessors reviewing environments with virtualization as part of a PCI DSS assessment.

> ***Note**: This document presumes a basic level of understanding of virtualization technologies and principles. However, an architectural-level understanding of virtualization technologies is required to assess technical controls in virtualized environments as the nature of these environments, particularly in the areas of process isolation and virtualized networking, can be substantially different from traditional physical environments.*

1.2 Intended Use

This document provides supplemental guidance on the use of virtualization technologies in cardholder data environments and does not replace or supersede PCI DSS requirements. For specific compliance criteria and audit requirements, virtualized environments should be evaluated against the criteria set forth in the PCI DSS.

This document is not intended as an endorsement for any specific technologies, products or services, but rather as recognition that these technologies exist and may influence the security of payment card data.

2 Virtualization Overview

2.1 Virtualization Concepts and Classes

Virtualization refers to the logical abstraction of computing resources from physical constraints. One common abstraction is referred to as a virtual machine, or VM, which takes the content of a physical machine and allows it to operate on different physical hardware and/or along with other virtual machines on the same physical hardware. In addition to VMs, virtualization can be performed on many other computing resources, including operating systems, networks, memory and storage.

The term "workload" is increasingly used to describe the vast array of virtualized resources. For example, a virtual machine is a type of workload. While VMs are the predominant virtualization technology implemented today, there are a number of other workloads to consider, including application, desktop, network, and storage virtualization models. The following types of virtualization are included in the focus of this document.

2.1.1 Operating System

Operating system (OS) virtualization is commonly used to take the resources running in an operating system on a single physical server and separate them into multiple, smaller partitions, such as virtual environments, virtual private servers, guests, zones, etc. In this scenario, all partitions would use the same underlying OS kernel (that is, they would run the same operating system as the base system), but may run different libraries, distributions, etc.

Similarly, application virtualization separates individual instances of an application from the underlying operating system, providing a discrete application "workspace" for each user.

2.1.2 Hardware/Platform

Hardware virtualization is accomplished through hardware partitioning or hypervisor technology. The hypervisor mediates all hardware access for the VMs running on the physical platform. There are two types of hardware virtualization:

Type 1 Hypervisor – A Type 1 hypervisor (also known as "native" or "bare metal") is a piece of software or firmware that runs directly on the hardware and is responsible for coordinating access to hardware resources as well as hosting and managing VMs.

Type 2 Hypervisor – A Type 2 hypervisor (also known as "hosted") runs as an application on an existing operating system. This type of hypervisor emulates the physical resources required by each VM, and is considered just another application as far as the underlying OS is concerned.

2.1.3 Network

Network virtualization distinguishes logical from physical networking. For nearly every type of physical networking component (for example, switches, routers, firewalls, intrusion prevention systems, load balancers, etc.), there is a logical counterpart available as a virtual appliance.

Unlike other standalone hosts (such as a server, workstation, or other system type), network devices operate across the following logical "planes":

- Data plane: Forwards data communications between hosts on the network.
- Control plane: Manages traffic, network and routing information; including communications between network devices related to network topology, state, and routing paths.
- Management plane: Handles direct communications into the device itself for the purpose of device management (for example, configuration, monitoring, and maintenance activities).

2.1.4 Data Storage

Virtualized data storage occurs when multiple physical storage devices on a network are combined and presented as a single storage device. This data consolidation is commonly used in storage area networks (SANs).

One of the benefits of virtualized storage is that the complexity of the storage infrastructure is hidden, out of sight of the user. However, this also presents a significant challenge for entities wishing to document and manage their data stores, as a particular data set may be stored across multiple distributed locations at any one time.

2.1.5 Memory

Memory virtualization is the consolidation of available physical memory from multiple individual systems to create a virtualized "pool" of memory which is then shared among system components.

Similar to virtualized data storage, the consolidation of multiple physical memory resources into a single virtual resource can add levels of complexity when it comes to mapping and documenting data locations.

2.2 Virtual System Components and Scoping Guidance

This section identifies some of the more common virtual abstractions or "virtual system components" that may be present in many virtual environments, and provides high-level scoping guidance for each.

Note that the scoping guidance provided in this section should be considered additional to the underlying principle that PCI DSS applies to all system components, including virtualized components, included in or connected to the cardholder data environment. Determining whether a particular virtual system component is to be considered in scope will depend on the specific technology and how it is implemented in the environment

2.2.1 Hypervisor

The hypervisor is the software or firmware responsible for hosting and managing virtual machines. The hypervisor system component may also include the virtual machine monitor (VMM). The VMM is a software component that implements and manages virtual machine hardware abstraction and can be considered the management function of a hypervisor platform. The VMM manages the system's processor, memory, and other resources to allocate what each virtual machine (also known as a "guest") operating system requires. In some circumstances it provides this functionality in conjunction with hardware virtualization technology.

Scope Guidance: If any virtual component connected to (or hosted on) the hypervisor is in scope for PCI DSS, the hypervisor itself will always be in scope. For additional guidance on the presence of both in-scope and out-of-scope VMs on the same hypervisor, please refer to Section 4.2 Recommendations for Mixed Mode Environments.

Note: *the term "mixed-mode" refers to a virtualization configuration where both in-scope and out-of-scope virtual components are running on the same hypervisor or host.*

2.2.2 Virtual Machine

A Virtual Machine (VM) is a self-contained operating environment that behaves like a separate computer. It is also known as the "Guest", and runs on top of a hypervisor.

Scope Guidance: An entire VM will be in scope if it stores, processes or transmits cardholder data, or if it connects to or provides an entry point into the CDE. If a VM is in scope, both the underlying host system and the hypervisor would also be considered in scope, as they are directly connected to and have a fundamental impact on the functionality and security of the VM.

2.2.3 Virtual Appliance

A virtual appliance can be described as a pre-packaged software image designed to run inside a virtual machine. Virtual appliances are each intended to deliver a specific function, and typically consist of basic operating system components and a single application. Physical network devices such as routers, switches, or firewalls can be virtualized and run as virtual appliances.

A Virtual Security Appliance (VSA)—also known as a Security Virtual Appliance (SVA)—is a virtual appliance consisting of a hardened operating system and a single security application. VSAs are typically assigned a higher level of "trust" than a regular VA, including privileged access to the hypervisor and other resources. In order for the VSA to perform system and network management functions, it usually has increased visibility into the hypervisor and any virtual networks running inside the hypervisor. Some VSA solutions may plug directly into the hypervisor, providing additional security to the platform. Examples of appliances that have virtual implementations include firewalls, IPS/IDS, and anti-virus.

Scope Guidance: Virtual Appliances used to connect or provide services to in-scope system components or networks would be considered in-scope. Any VSA/SVA that could impact the security of the CDE would also be considered in scope.

2.2.4 Virtual Switch or Router

A virtual switch or router is a software component that provides network-level data routing and switching functionality. A virtual switch is often an integral part of a virtualized server platform—for example, as a hypervisor driver, module, or plug-in. A virtual router may be implemented as a distinct virtual appliance or as a component of a physical appliance. Additionally, virtual switches and routers may be used to generate multiple logical network devices from a single physical platform.

Scope Guidance: Networks provisioned on a hypervisor-based virtual switch will be in scope if provisioned with an in-scope component or if they provide services or connect to an in-scope component. Physical devices hosting virtual switches or routers would be considered in scope if any of the hosted components connects to an in-scope network.

2.2.5 Virtual Applications and Desktops

Individual applications and desktop environments can also be virtualized to provide functionality for end users. Virtual applications and desktops are typically installed at a central location and accessed remotely via a remote desktop interface. Virtual desktops can be configured to allow access via multiple device types, including thin clients and mobile devices, and may run using local or remote computing resources. Virtual applications and desktops may be present in point-of-sale, customer service, and other interactions with the payment chain.

Scope Guidance: Virtual applications and desktops will be in scope if they are involved in the processing, storage, or transmission of cardholder data, or provide access to the CDE. If a virtual application or desktop is provisioned on the same physical host or hypervisor as an in-scope component, the virtual application/desktop will also be in scope unless adequate segmentation is in place that isolates all in-scope components from the out-of-scope components. For additional guidance on the presence of both in-scope and out-of-scope components on the same host or hypervisor, please refer to Section 4.2 Recommendations for Mixed Mode Environments.

2.2.6 Cloud Computing

Cloud computing is a rapidly evolving use of virtualization that provides computing resources as a service or utility over public, semi-public, or private infrastructures. Cloud-based service offerings are usually delivered from a "pool" or "cluster" of connected systems and provide service-based access to shared computing resources for multiple users, entities, or tenants.

Scope Guidance: The use of cloud computing presents a number of scoping challenges and considerations. Entities planning to use cloud computing for their PCI DSS environments should first ensure that they thoroughly understand the details of the services being offered, and perform a detailed assessment of the unique risks associated with each service. Additionally, as with any managed service, it is crucial that the hosted entity and provider clearly define and document the responsibilities assigned to each party for maintaining PCI DSS requirements and any other controls that could impact the security of cardholder data.

The cloud provider should clearly identify which PCI DSS requirements, system components, and services are covered by the cloud provider's PCI DSS compliance program. Any aspects of the service not covered by the cloud provider should be identified, and it should be clearly documented in the service agreement that these aspects, system components, and PCI DSS requirements are the responsibility of the hosted entity to manage and assess. The cloud provider should provide sufficient evidence and assurance that all processes and components under their control are PCI DSS compliant.

For additional guidance on the use of cloud environments, please refer to Section 4.3 Recommendations for Cloud Computing Environments.

3 Risks for Virtualized Environments

While virtualization may provide a number functional and operational benefits, moving to a virtual environment doesn't alleviate the risks which existed on the physical systems, and may also introduce new and unique risks. Consequently, there are a number of factors to be considered when implementing virtual technologies, including but not limited to those defined below.

3.1 Vulnerabilities in the Physical Environment Apply in a Virtual Environment

Virtual systems and networks are subject to the same attacks and vulnerabilities that exist in a physical infrastructure. An application that has configuration flaws or is vulnerable to exploits will still have those same flaws and vulnerabilities when installed in a virtual implementation. Similarly, a poorly configured virtual firewall could unwittingly expose internal systems to internet-based attacks in the same way misconfiguration on a physical firewall would do.

Physical threats also apply to virtual implementations; the most securely configured, well-contained logical partitions will still need adequate physical controls for protection of the hardware. For this reason, the physical host system will always remain in scope even where there is logical reduction.

3.2 Hypervisor Creates New Attack Surface

A key risk factor unique to virtual environments is the hypervisor—if this is compromised or not properly configured, all VMs hosted on that hypervisor are potentially at risk. The hypervisor provides a single point of access into the virtual environment and is also potentially a single point of failure. Misconfigured hypervisors could result in a single point of compromise for the security of all hosted components. No matter how securely the individual virtual machines or components may be configured, a compromised hypervisor can override those controls and gain direct access to the virtual systems.

As well as providing a potential entry point to the VMs hosted on it, the hypervisor itself creates a new attack surface that does not exist in the physical world and may be vulnerable to direct attacks. Weaknesses in hypervisor isolation technology, access controls, security hardening, and patching could be identified and exploited, allowing attackers to gain access to individual VMs. Additionally, a hypervisor's default out-of-the-box configuration is often not the most secure; and unless properly configured, even a secure hypervisor can potentially be exploited.

It is critical that access to the hypervisor be restricted according to least privilege and need to know, and that independent monitoring of all activities be enforced. Hypervisors are not created equal, and it is particularly important to choose a solution that supports the required security functions for each environment.

3.3 Increased Complexity of Virtualized Systems and Networks

Virtualized configurations can encompass both systems and networks; for example, VMs may transmit data to each other through the hypervisor, as well as over virtual network connections and through virtual network security appliances such as virtual firewalls. While such configurations may offer significant operational benefits, these additional layers of technology also introduce a considerable amount of complexity that must be carefully managed, and may require additional security controls and complex policy management to ensure appropriate security is applied at each layer. Combined with potential vulnerabilities in virtual operating systems and applications, this increased complexity can also lead to accidental misconfiguration or even entirely new threats that were unforeseen by a system's designers. Because instances of virtual components are often replicated across multiple systems, the presence of such vulnerabilities could result in significant compromise across an entire environment.

3.4 More Than One Function per Physical System

Of particular concern in virtual environments is the possibility that the compromise of one virtual system function could lead to a compromise of other functions on the same physical system. A compromised VM could use virtualization-layer communication mechanisms to launch attacks on other VMs on the same host or even the hypervisor. Virtualization technologies may be able to mitigate some of this risk by enforcing process separation between different functions. Even so, the risk associated with locating multiple functions or components on a single physical system must still be considered. For example, having multiple functions hosted on one physical system increases the possible scope of compromise should an attacker gain physical access to the host system.

See also "Mixing VMs of Different Trust Levels" below for related risk considerations.

3.5 Mixing VMs of Different Trust Levels

One of the challenges when planning a virtualization deployment is to identify appropriate configurations for the variety of workloads to be housed within the particular virtualization technology. The risk of hosting VMs of different trust levels on the same host needs to be carefully assessed. In the virtual context, a VM of lower trust will typically have lesser security controls than VMs of higher trust levels. The lower-trust VM could therefore be easier to compromise, potentially providing a stepping stone to the higher-risk, more sensitive VMs on the same system. Theoretically, hosting VMs of different trust levels on the same hypervisor or host could reduce the overall security for all components to that of the least-protected component (also known as the "security is only as strong as the weakest link" principle).

Due to the increased risks and configuration challenges, the trust and risk level associated with each VM function should be taken into account when considering a virtualized design. Similarly, databases and other systems that store cardholder data require a higher security level than non-

sensitive data stores. The risk of mixing sensitive data with data of lower trust must be carefully assessed.

3.6 Lack of Separation of Duties

It can be particularly challenging to define granular user roles (for example, separation of network administrator from server administrator), and access policies across a distributed, virtualized environment. The risks of failing to properly define roles and access policies are significant because access to the hypervisor can potentially provide broad access to key infrastructure components (including switches, firewalls, payment applications, log-aggregation servers, databases, etc.). Because of the increased accessibility to multiple virtual devices and functions from a single logical location or a user, monitoring and enforcement of appropriate separation of duties is crucial in a virtual environment.

3.7 Dormant Virtual Machines

On many virtualization platforms, VMs can exist in active or dormant states. VMs that are not active (dormant or no longer used) could still house sensitive data such as authentication credentials, encryption keys, or critical configuration information. Inactive VMs containing payment card data can become unknown, unsecured data stores, which are often only rediscovered in the event of a data breach.

Because dormant VMs are not actively used, they can easily be overlooked and inadvertently left out of security procedures. A dormant VM will likely not be updated with the latest security patches, resulting in the system being exposed to known vulnerabilities that the organization thinks have been addressed. Dormant VMs are also unlikely to have up-to-date access policies, and may be excluded from security and monitoring functions, possibly creating an unchecked “back door” to the virtual environment.

Additionally, data in a VM’s memory (which may include, for example, unencrypted PAN) is often captured in its dormant state, resulting in unintentional storage of the data. As this data was in memory when it was captured, it could easily be overlooked and left unprotected, even though it is now “stored” in the dormant VM. Though dormant, inactive VMs represent a viable security threat and therefore must be identified and tracked so appropriate security controls can be applied.

3.8 VM Images and Snapshots

Virtual machine images and snapshots provide a means to quickly deploy or restore virtual systems across multiple hosts within a short period of time. Special attention must be paid to the preparation of VM images and snapshots, as they may capture sensitive data present on the system at the time the image was taken, including contents of active memory. This could result in the inadvertent capture, storage, or even deployment of sensitive information throughout the environment.

Additionally, if images aren't secured and protected from modification, an attacker may gain access and insert vulnerabilities or malicious code into the image. The compromised image could then be deployed throughout the environment, resulting in a rapid compromise of multiple hosts.

3.9 Immaturity of Monitoring Solutions

At the same time that virtualization increases the need for logging and monitoring, it is currently recognized that the tools to monitor the virtual networks, virtual firewalls, virtual compliance systems, etc. are not as mature as their physical counterparts.

Compared to traditional monitoring tools for a physical network, tools for virtual systems may not provide the same level of insight or monitoring within intra-host communications or traffic flowing between VMs on a virtual network. Similarly, specialized tools for monitoring and logging virtual environments may be needed to capture the level of detail required from the multiple components, including hypervisors, management interfaces, virtual machines, host systems, and virtual appliances.

3.10 Information Leakage between Virtual Network Segments

The potential risks of information leakage between logical network segments should be understood when considering network virtualization. Information leakage at the data plane results in sensitive data existing outside of known locations, circumventing the data protection controls that would otherwise apply. Information leakage at the control plane or management plane can be exploited to enable information leakage at the data plane, or to influence network routes and forwarding behavior to bypass network-based security controls. Ideally, virtualization capabilities at all three planes of operation in the network infrastructure should provide controls and features to secure the virtualized infrastructure at a level equivalent to individual physical devices.

3.11 Information Leakage between Virtual Components

Information leakage between virtual components can occur when access to shared resources allows one component to collect information about another component on the same host. For example, an attacker can use a compromised component to gather information about other components running on the same host, and potentially gain enough knowledge for further compromise. In another example, the attacker could obtain access to the underlying operating system memory, resulting in the potential capture of sensitive information from multiple components. A misconfigured hypervisor may also become a conduit for information leakage between hosted virtual components and networks. Isolation of all physical resources (including memory, CPU, network, etc.) is critical to prevent information leakage between VMs and other components or networks on the same host.

4 Recommendations

The controls identified in this section are recommendations and best practices that may assist with meeting PCI DSS requirements in virtual environments.

4.1 General Recommendations

4.1.1 Evaluate risks associated with virtual technologies

Entities should carefully and thoroughly evaluate the risks associated with virtualizing system components before selecting or implementing a virtualization solution. The flow and storage of cardholder data should be accurately documented as part of this risk assessment process to ensure that all risk areas are identified and appropriately mitigated. Virtualization should be deployed with a complete view of its benefits and risks, and a comprehensive, defined set of effective system, application, data, and environmental controls.

Virtualized environments and system components should continue to be included in an annual risk-assessment process. Risk evaluation and management decisions should be fully documented and supported by detailed business and technical evaluations.

4.1.2 Understand impact of virtualization to scope of the CDE

Entities using virtualization to consolidate their environment onto one or more physical hardware platforms may find that, as a result, they now have a complex set of virtual system configurations making it difficult to identify the boundaries or scope of their CDE.

Like physical systems, the scope of PCI DSS across virtual components must be thoroughly verified and documented. The virtual environment should be evaluated using the guidance provided in the “Scope of Assessment for Compliance with PCI DSS Requirements” section of the PCI DSS. If any components running on a single hypervisor are in scope, it is recommended that all components on that hypervisor be considered in-scope as well, including but not limited to virtual machines, virtual appliances, and hypervisor plug-ins. Designing all virtualization components, even those considered out-of-scope, to meet PCI DSS security requirements will not only provide a secure baseline for the virtual environment as a whole, it will also reduce the complexity and risk associated with managing multiple security profiles, and lower the overhead and effort required to maintain and validate compliance of the in-scope components.

4.1.3 Restrict physical access

As identified earlier in this document, hosting multiple components on one physical system could greatly increase the potential impact if an attacker gains physical access to that host system. Physical access controls are therefore particularly important in virtualized environments and should be strengthened as necessary to mitigate the associated risks. When assessing physical controls, consider the potential harm of an unauthorized or malicious individual gaining simultaneous access to all VMs, networks, security devices, applications, and hypervisors that

one physical host could provide. Ensure that all unused physical interfaces are disabled, and that physical or console-level access is restricted and monitored.

4.1.4 Implement defense in depth

In a physical environment, a defense-in-depth approach that encompasses preventive, detective, and responsive controls is a common best practice for securing data and other assets. Logical security controls are typically applied at the network, host, application, and data layer, and physical security controls are implemented to protect media, systems, and facilities from unauthorized physical access. Monitoring the effectiveness of controls and the capacity to respond quickly and effectively to a potential breach are also of paramount importance. A defense-in-depth approach also includes training and educating personnel in the proper use of sensitive assets, the identification of potential security threats, and the appropriate action to be taken in the event of a breach. Additionally, a defense-in-depth environment has well-defined and documented policies, processes, and procedures that are understood and followed by all personnel.

Appropriate security controls should be identified and implemented in a virtualized environment that provide the same level and depth of security as can be achieved in a physical environment. For example, consider how security can be applied to protect each technical layer, including but not limited to the physical device, hypervisor, host platform, guest operating systems, VMs, perimeter network, intra-host network, application, and data layers. Physical controls, documented policies and procedures, and training of personnel should also be a part of a defense-in-depth approach to securing virtual environments.

4.1.5 Isolate security functions

The security functions provided by VMs must be implemented with the same process separation required in the physical world. It is recommended that this requirement be even more stringently enforced in virtualized systems because it significantly complicates the efforts required by an attacker to compromise multiple CDE system components. For example, preventive controls such as a network firewall should never be combined on a single logical host with the payment card data it is configured to protect. Similarly, processes controlling network segmentation and the log-aggregation function that would detect tampering of network segmentation controls should not be mixed. If such security functions are to be hosted on the same hypervisor or host, the level of isolation between security functions should be such that they can be considered as being installed on separate machines.

4.1.6 Enforce least privilege and separation of duties

Accounts and credentials for administrative access to the hypervisor should be carefully controlled, and depending on the level of risk, the use of more restrictive hypervisor access controls is often justified. Entities should consider additional methods for securing administrative access, such as implementing two-factor authentication or establishing dual or split-control of administrative passwords between multiple administrators. Access controls should be assessed for both local and remote access to the hypervisor and management system. Particular attention should be directed to the functions of the individual virtual components, to ensure that appropriate role-based access controls (RBAC) are in place that prevent unnecessary access to resources and enforce separation of duties.

Administrative privileges also need to be appropriately separated. For example, a single-user administrator should not be granted privileged access to firewalls and the monitoring servers for those firewalls. Such broad access could result in undetected tampering and data loss that could have been prevented with properly enforced separation of duties. As a best practice, restrict administrative access by specific VM function, virtual network, hypervisor, hardware, application, and data store.

4.1.7 Evaluate hypervisor technologies

Ensure that the security of the hypervisor has been thoroughly tested prior to deployment and that there is appropriate patch management and other controls to respond to threats and exploits. Identifying and implementing technologies that facilitate strong security practices is critical, as not all hypervisors or VMMs have the functionality to support appropriate security controls.

4.1.8 Harden the hypervisor

Hypervisor platforms should be deployed in a secure manner according to industry-accepted best practices and security guidelines. Careful management of virtual system configurations, patching, and change-control processes is essential to ensure that all hypervisor changes are monitored, authorized, fully tested, and carefully controlled. Due to potential severity of a hypervisor compromise, patches and other mitigating controls should be deployed as soon as possible whenever new security vulnerabilities are discovered, and include immediate testing for the vulnerability to confirm the risk has been addressed.

Because the hypervisor represents a single point of failure, an unauthorized or malicious modification could threaten the integrity of all hosted systems in the environment. The following additional controls are recommended for the hypervisor and any significant management tools.

- Restrict the use of administrative functions to defined endpoint networks and devices, such as specific laptops or desktops that have been approved for such access.
- Require multi-factor authentication for all administrative functions.
- Ensure that all changes are implemented and tested properly. Consider requiring additional management oversight, above and beyond that which is required through the normal change-management process.

- Separate administrative functions such that hypervisor administrators do not have the ability to modify, delete, or disable hypervisor audit logs.
- Send hypervisor logs to physically separate, secured storage as close to real-time as possible.
- Monitor audit logs to identify activities that could indicate a breach in the integrity of segmentation, security controls, or communication channels between workloads.
- Separate duties for administrative functions, such that authentication credentials for the hypervisor do not have access to applications, data, or individual virtual components.
- Before implementing a virtualization solution, verify what security controls the solution supports and how they minimize risk of compromise to the hypervisor.

Note that as the hypervisor and management tools could directly impact the security of virtual components, they should always be considered in scope for PCI DSS.

4.1.9 Harden virtual machines and other components

It's also critical that all individual virtual machines are installed and configured securely and according to industry best practices and security guidelines. The recommendations provided above for hardening the hypervisor are also applicable to VMs and virtual components.

Note that these recommendations may not all be applicable for each type of virtual machine or component. Implementations should be evaluated individually to confirm that the following is considered:

- Disable or remove all unnecessary interfaces, ports, devices and services;
- Securely configure all virtual network interfaces and storage areas;
- Establish limits on VM resource usage;
- Ensure all operating systems and applications running inside the virtual machine are also hardened;
- Send logs to separate, secured storage as close to real-time as possible;
- Validate the integrity of the cryptographic key-management operations;
- Harden individual VM virtual hardware and containers;
- Other security controls as applicable.

Security and hardening requirements may differ depending on the specific services or applications running on each virtual component; consequently, the appropriate security settings will need to be individually determined.

4.1.10 Define appropriate use of management tools

Management tools allow administrators to perform such functions as system back-up, restore, remote connectivity, migration, and configuration changes to virtual systems. Management tools for in-scope components would also be considered in scope, as these tools directly impact the security and functioning of the in-scope components. Access to management tools should be limited to those with a job-related need to have such access. Segregation of roles and responsibilities is recommended for management tool functions, and the use of management tools should be monitored and logged.

4.1.11 Recognize the dynamic nature of VM's

VMs are effectively just "data" that can reside in active states (on a hypervisor) or inactive states (anywhere). Inactive or dormant VMs are effectively stored data sets that could contain sensitive information and virtual device configuration details. An individual with access to a dormant VM could copy and activate it in another location, or they could scan the dormant files for payment card data and other sensitive information. Access to dormant VMs should therefore be restricted, monitored, and carefully controlled.

Inactive VMs that contain payment card data need to be treated with same level of sensitivity and have the same safeguards as any other cardholder data store. Migration paths of inactive VMs should be carefully evaluated since they may bring additional systems into scope. Backups of VMs, active VMs, and inactive VMs should always be protected and securely deleted or secure-wiped when the data is no longer needed. Thorough change-management, monitoring, and alerting processes are essential to ensure that only authorized VM's are added to and removed from the environment, and that all access and actions are recorded.

4.1.12 Evaluate virtualized network security features

Ideally, any deployment of virtualized network infrastructure should include effective security measures at the data plane, control plane, and management plane. This will help minimize the possibility of direct and indirect vulnerabilities cascading from one plane to another and compromising the virtual network devices. While ideal, effective security measures at all three operational planes may not always be possible. In these cases it becomes increasingly important to ensure that the underlying physical components are adequately isolated and secured and do not provide a "path" between virtual network devices. Isolation between virtualized network devices should be such that the virtual systems can be effectively regarded as separate hardware.

Each virtualized device should maintain individual and independent access-controlled configurations. Audit trails for virtual infrastructures should be granular and detailed enough to identify individual access to and activities performed on each specific virtual component. Access controls should enforce least privilege, both individually for each device and across the entire platform.

4.1.13 Clearly define all hosted virtual services

Sometimes shared hosting providers virtualize their offerings, provisioning separate workloads to customers rather than provisioning separate physical systems. Entities considering a hosted virtual service should ensure that the service offering enforces administrative, process, and technical segmentation to isolate each hosted entity's environment from other entities. This isolation should, at a minimum, encompass all PCI DSS controls, including but not limited to segmented authentication, network and access controls, encryption and logging.

Additionally, it is critical to ensure that all details of the service, including responsibilities for maintaining controls which could affect the security or integrity of sensitive data or that could impact the entity's PCI DSS compliance, are clearly defined and documented in a formal agreement.

4.1.14 Understand the technology

Virtualized environments are substantially different from traditional physical environments, and thorough understanding of virtualization technologies is required to effectively evaluate and secure any environment. In the absence of formal virtualization security standards, entities should familiarize themselves with industry-accepted best practices and guidelines for securing virtualized environments. Examples of resources that may provide guidance include publications from:

- The Center for Internet Security (CIS)
- International Organization for Standardization (ISO)
- ISACA (formerly the Information Systems Audit and Control Association)
- National Institute of Standards Technology (NIST)
- SysAdmin Audit Network Security (SANS) Institute

4.2 Recommendations for Mixed-Mode Environments

It is strongly recommended (and a basic security principle) that VMs of different security levels are not hosted on the same hypervisor or physical host; the primary concern being that a VM with lower security requirements will have lesser security controls, and could be used to launch an attack or provide access to more sensitive VMs on the same system.

This principle should also be applied if in-scope and out-of-scope virtual systems are to be located on the same host or hypervisor. As a general rule, any VM or other virtual component that is hosted on the same hardware or hypervisor as an in-scope component would also be in scope for PCI DSS, as both the hypervisor and underlying host provide a connection (either physical, logical, or both) between the virtual components, and it may not be possible to achieve an appropriate level of isolation, or segmentation, between in-scope and out-of-scope components located on the same host or hypervisor.

As stated earlier in this document, any hypervisor or host system that houses an in-scope virtual component would also be in scope for PCI DSS. In order for in-scope and out-of-scope VMs to co-exist on the same host or hypervisor, the VMs must be isolated from each other such that they can effectively be regarded as separate hardware on different network segments with no connectivity to each other. Any system components shared by the VMs, including the hypervisor and underlying host system, must therefore not provide an access path between the VMs.

Even if adequate segmentation between virtual components could be achieved, the resource effort and administrative overhead required to enforce the segmentation and maintain different security levels on each component would likely be more burdensome than applying PCI DSS controls to the system as a whole.

4.2.1 Segmentation in Mixed-Mode Environments

The level of segmentation required for in-scope and out-of-scope systems on the same host must be equivalent to a level of isolation achievable in the physical world; that is, segmentation must ensure that out-of-scope workloads or components cannot be used to access an in-scope component. Unlike separate physical systems, network-based segmentation alone cannot isolate in-scope from out-of-scope components in a virtual environment.

Segmentation of virtual components must also be applied to all virtual communication mechanisms, including the hypervisor and underlying host, as well as any other common or shared component. In virtual environments, out-of-band communications can occur, often via a solution-specific communication mechanism, or through the use of shared resources such as file systems, processors, volatile and non-volatile memory, device drivers, hardware devices, APIs, and so on. Out-of-band communication channels are generally specific to the virtualization technology in use, so a detailed understanding of all underlying mechanisms is critical when planning to segment virtual components. All existing out-of-band channels should be identified and documented—whether they are actively used or not—and appropriate controls implemented to isolate workloads and virtual components. In some cases, physical separation of hardware resources may be required to prevent hardware being used as an access path between virtual components.

It is important to note that out-of-band channels are often required for specific virtual system functions, and it is often not possible to isolate components from these channels without impacting system operations. If it is not feasible for a particular implementation to enforce isolation of in-scope components from out-of-scope components via shared resources or other out-of-band channels, all components accessing the shared resource or out-of-band channel should be considered in scope, as they are effectively connected to the in-scope component.

Process isolation is also an inherent component of segmentation between virtual systems. In a mixed-mode configuration, the hypervisor plays a critical role in enforcing process isolation between the in-scope and out-of-scope systems. It is therefore critical that these controls are functioning properly and that access to hypervisor functions that could affect these controls is strictly controlled and monitored.

As well as isolation of processes and shared resources, virtual storage of cardholder data is a key consideration and is often overlooked when it comes to segmentation of virtual components. Depending on the specific configuration and controls implemented, an entire SAN could potentially be in scope unless it is verified that all in-scope systems and data stores are isolated from all out-of-scope systems and data stores.

4.3 Recommendations for Cloud Computing Environments

Cloud computing provides a broad range of service offerings and deployment models encompassing many different technologies, products and services. Cloud environments may be deployed over a private infrastructure, public infrastructure, or a hybrid of both, as described below.

Private Cloud

A private cloud consists only of system components that are trusted and controlled by the entity. The trusted systems may be located across multiple facilities that are owned by the entity or a third party. Similarly, the systems and components themselves may be the property of the entity, or they could be owned by a service provider and provisioned for dedicated use by a single customer. Irrespective of ownership, systems in private cloud are dedicated to a single entity and computing resources are not shared with any other customer or tenant.

Public Cloud

A public cloud consists of system components that the organization does not own or have any control over. In a public cloud, some part of the underlying systems and infrastructure is always controlled by the cloud service provider. The specific components remaining under the control of the cloud provider will vary according to the type of service—for example, Infrastructure as a Service (IaaS), Platform as a Service (PaaS) and Software as a Service (SaaS). Public cloud services are typically delivered from a “pool” or “cluster” of systems to provide service-based access for multiple customers, or tenants, to shared computing resources. Physical separation between tenants is not practical in a public cloud environment because, by its very nature, all resources are shared by everyone.

Hybrid Cloud

As suggested by the name, a hybrid cloud is a combination of both private and public cloud infrastructures. A hybrid cloud is typically created when an entity interconnects their private cloud with a public cloud or another entity’s private cloud. In a hybrid cloud, ownership and control of data and system components may be divided between three or more separate entities, adding complexity to the task of identifying scope boundaries and defining responsibilities.

Cloud computing also encompasses several types of services, including IaaS, PaaS, and SaaS. Each type of service represents a different assignment of resource management and ownership, which will vary depending on the specific service offering.

For example, an entity subscribing to an IaaS service may retain complete control of, and therefore be responsible for, the ongoing security and maintenance of all operating systems, applications, virtual configurations (including the hypervisor and virtual security appliances), and data. In this scenario, the cloud provider would only be responsible for maintaining the underlying physical network and computing hardware. In an alternative scenario, a SaaS service offering may encompass management of all hardware and software, including virtual components and hypervisor configurations. In this scenario, the entity may only be responsible for protecting their data, and all other security requirements would be implemented and managed by the service provider.

The following diagram provides an example of how an entity's scope and responsibility may vary across different types of cloud service offerings.

Example of how scope and responsibility may differ* by type of cloud service:

Cloud customer responsibility
Cloud service provider responsibility

	Type of Cloud Service		
Area of Responsibility	IAAS	PAAS	SAAS
Data			
Software, user applications			
Operating systems, databases			
Virtual infrastructure (hypervisor, virtual appliances, VMs, virtual networks etc)			
Computer and network hardware (processor, memory, storage, cabling, etc.)			
Data center (physical facility)			

* ***Note:*** *This is an example only. Cloud service offerings should be individually reviewed to determine how responsibilities* between *the cloud provider and cloud customer are assigned.*

In a public cloud environment, the services and computing resources provided by the cloud provider are typically shared across multiple entities, or tenants. This is in contrast to traditional hosting environments where dedicated resources are usually provisioned to each hosted entity or tenant. Unlike traditional hosting environments, where physical isolation between tenants is usually enforced, physical separation between tenants in a cloud environment is not practical because, as mentioned previously, all resources are shared by everyone.

In addition to the challenges of defining scope and assigning responsibilities across a shared infrastructure, the inherent characteristics of many cloud environments present additional barriers to achieving PCI DSS compliance. Some of these characteristics include:

- The distributed architectures of cloud environments add layers of technology and complexity to the environment.
- Public cloud environments are designed to be public-facing, to allow access into the environment from anywhere on the Internet.
- The infrastructure is by nature dynamic, and boundaries between tenant environments can be fluid.
- The hosted entity has limited or no visibility into the underlying infrastructure and related security controls.
- The hosted entity has limited or no oversight or control over cardholder data storage.
- The hosted entity has no knowledge of "who" they are sharing resources with, or the potential risks their hosted neighbors may be introducing to the host system, data stores, or other resources shared across a multi-tenant environment.

In a public cloud environment, additional controls must be implemented to compensate for the inherent risks and lack of visibility into the public cloud architecture. A public cloud environment could, for example, host hostile out-of-scope workloads on the same virtualization infrastructure as a cardholder data environment. More stringent preventive, detective, and corrective controls are required to offset the additional risk that a public cloud, or similar environment, could introduce to an entity's CDE.

These challenges may make it impossible for some cloud-based services to operate in a PCI DSS compliant manner. Consequently, the burden for providing proof of PCI DSS compliance for a cloud-based service falls heavily on the cloud provider, and such proof should be accepted only based on rigorous evidence of adequate controls.

As with all hosted services in scope for PCI DSS, the hosted entity should request sufficient assurance from their cloud provider that the scope of the provider's PCI DSS review is sufficient, and that all controls relevant to the hosted entity's environment have been assessed and determined to be PCI DSS compliant. The cloud provider should be prepared to provide their hosted customers with evidence that clearly indicates what was included in the scope of their PCI DSS assessment as well as what was not in scope; details of controls that were not covered and are therefore the customer's responsibility to cover in their own PCI DSS assessment; details of which PCI DSS requirements were reviewed and considered to be "in place" and "not in place"; and confirmation of when the assessment was conducted.

Any aspects of the cloud-based service not covered by the cloud provider's PCI DSS review should be identified and documented in a written agreement. The hosted entity should be fully aware of any and all aspects of the cloud service, including specific system components and security controls, which are not covered by the provider and are therefore the entity's responsibility to manage and assess.

4.4 Guidance for Assessing Risks in Virtual Environments

Due to the lack of standardization among virtualization technologies and the wide variety of possible implementations, organizations using virtualization need to assess their particular environment, evaluate the associated risks and identify appropriate controls to address that risk.

There are a number of industry-accepted risk assessment methodologies and tools available to help guide the risk assessment process. Whichever process is used, it should include identification of threats and vulnerabilities and result in a clear understanding of the assessed risk to the environment.

Some of the key elements to be considered when performing a risk assessment of virtual environments are provided below.

4.4.1 Define the environment

Before threats and vulnerabilities can be identified and evaluated, an entity must first understand their environment as well as the people, processes, and technologies that comprise or interact with that environment. When defining the environment to be assessed, entities should consider all aspects that have a potential risk impact, regardless of whether they are considered in-scope or out-of-scope for PCI DSS.

Defining the virtual environment should include, at a minimum, the following activities:

- Identification of all components, including hypervisors, workloads, hosts, networks, management consoles and other components;
- Physical site details for each component;
- Description of the primary functions and assigned owners for each component;
- Details of visibility into and between components;
- Identification of traffic flows between different components, between components and hypervisors, and between components and underlying host systems or hardware resources;
- Identification of all intra-host communications and data flows, as well as those between virtual components and other system components;
- Identification of all out-of-band communication channels (whether configured to operate or not) that could allow communications between components;
- Details of all management interfaces and hypervisor access mechanisms, including defined roles and permissions;
- All virtual and physical hardware components such as removable disk drives and USB, parallel, and serial ports.
- Details of the number and types of virtual components on each host, types of segmentation between components and hosts, functions and security levels of all virtualized components, etc.

4.4.2 Identify Threats

This process includes identifying current and potential threats that, if successful or permitted to occur, could result in a loss of confidentiality, integrity, or availability of the environment. Threat considerations should include all scenarios, actions, or events that could result in the deliberate or unintentional circumvention of security controls.

Virtualized environments are typically subject to the same types of threats as traditional environments. However, virtualization itself may provide an additional layer for potential threats to target. An example of potential threats specific to virtualization technologies may include new types of malicious code or logical attacks specifically targeting unique virtual components, such as the hypervisor, or unsecured out-of-band communication channels between shared hardware components.

Having a detailed understanding of the primary function and owner of each component in the environment will help to identify the potential impact of a successful attack or other threat event.

4.4.3 Identify Vulnerabilities

As well as the traditional technical vulnerabilities to be identified (for example, within operating systems, applications, etc.), entities also need to identify vulnerabilities specific to the particular virtualization technologies and configurations implemented in their environment. Additional vulnerabilities may result from the increased complexity introduced by layers of virtualization, the dynamic, shared nature of virtual environments, and the potential lack of visibility into the underlying architecture.

Vulnerabilities are not limited to technical issues. Flaws in operational processes, inadequate training of personnel, lack of control monitoring, and gaps in physical security are examples of additional areas where potential vulnerabilities could exist and be exploited.

4.4.4 Evaluate and Address Risk

The risk assessment should identify whether any additional controls are necessary to secure and protect cardholder data and other sensitive information in a virtualized environment. It should be noted that the implementation of specialized controls may be needed *in addition to* PCI DSS requirements to mitigate potential security issues associated with the use of virtualization technologies.

5 Conclusion

There is no single method for securing virtualized systems. Virtual technologies have many applications and uses, and the security controls appropriate for one implementation may not be suitable for another. As well as the visible functions of a virtual implementation, there are underlying functional and communication services built into virtualization architecture that could provide unknown attack vectors if not properly understood and managed.

As with many evolving technologies, the lack of virtualization industry standards has resulted in a number of vendor-specific best practices and recommendations that may or may not be applicable to a particular environment. Entities need to understand and evaluate their own environments to identify the unique risks virtualization brings, as well as the potential implications to the security of their cardholder data environment.

In a virtual environment, each and every individual component needs to be secured, as the insecurity of one VM or component on a host system could lead to the compromise of other VMs on the same host. Designing all virtualization components, even those considered to be out-of-scope, to meet PCI DSS security requirements will not only provide a secure baseline for the virtual environment as a whole, it will also reduce the complexity and risk associated with managing multiple security profiles, and lower the overhead and effort required to maintain and validate compliance of the in-scope components. For this reason, if any component running on a particular hypervisor or host is in scope for PCI DSS, it is recommended that all components on that hypervisor or host be considered in scope as well.

6 Acknowledgments

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

Alliance Data
Altor Networks
Assurant
AT&T
Bank of America
Capita Group PLC
Cisco
Citrix
Coalfire Systems
ConfigureSoft
DRG
Dufry/Hudson News
Firehost
HP
HyTrust
IGX Global
VMware
LL Bean
Microsoft
Net SPI
Protiviti
Quicktrip
Red Island
Reliant Security
Savvis
SecureState
Southwest Airlines
Speedway
Stanford University
SystemExperts
The Members Group
Trend Micro
Verizon Business

About the PCI Security Standards Council

The mission of the PCI Security Standards Council is to enhance payment account security by driving education and awareness of the PCI Data Security Standard and other standards that increase payment data security.

The PCI Security Standards Council was formed by the major payment card brands American Express, Discover Financial Services, JCB International, MasterCard Worldwide, and Visa Inc. to provide a transparent forum in which all stakeholders can provide input into the ongoing development, enhancement, and dissemination of the PCI Data Security Standard (DSS), PIN Transaction Security (PTS) Requirements, and the Payment Application Data Security Standard (PA-DSS). Merchants, banks, processors, and point-of-sale vendors are encouraged to join as Participating Organizations.

7 Appendix – Virtualization Considerations for PCI DSS

Where virtualization is implemented, all components within the virtual environment will need to be identified and considered in scope for a PCI DSS review, including the individual virtual hosts or devices, guest machines, applications, management interfaces, central management consoles, hypervisors, etc. All intra-host communications and data flows must be identified and documented, as well as those between the virtual component and other system components.

The implementation of a virtualized environment must meet the intent of all PCI DSS requirements, such that the virtualized systems can effectively be regarded as separate hardware. For example, there must be a clear segmentation of functions and segregation of networks with different security levels; segmentation should prevent the sharing of production and test/development environments; the virtual configuration must be secured such that vulnerabilities in one function cannot impact the security of other functions; and attached devices, such as USB/serial devices, should not be accessible by all virtual instances.

Additionally, all virtual management interface protocols should be included in system documentation, and roles and permissions should be defined for managing virtual networks and virtual system components. Virtualization platforms must have the ability to enforce separation of duties and least privilege, to separate virtual network management from virtual server management. Special care is also needed when implementing authentication controls to ensure that users authenticate to the proper virtual system components, and distinguish between the guest VMs (virtual machines) and the hypervisor.

The following section identifies some of the virtualization characteristics described earlier in this document and presents guidance on how these characteristics may be particularly relevant to some PCI DSS control areas. Additional recommendations and best practices are also included for consideration.

> ***Note:*** *This appendix is intended as guidance only. Virtual implementations and configurations will need to be individually evaluated for each particular environment to determine the impact of these considerations to PCI DSS requirements.*
>
> *It is also important to remember that* ***ALL applicable PCI DSS requirements must be evaluated****. The following guidance only identifies some of the potential areas for consideration when virtualization is used.*
>
> *The Virtualization Considerations in this appendix do not replace, supersede, or extend PCI DSS requirements. All best practices and recommendations contained herein are provided as guidance only.*

The following table is divided into two primary columns:

- **PCI DSS Requirements (summarized and abbreviated):** These columns contain summarized extracts of PCI DSS requirements. Note that the full content of the requirements is not provided here—please refer to the *PCI DSS Requirements and Security Assessment Procedures* for all PCI DSS requirements and testing procedures.
- **Virtualization Considerations:** this column identifies characteristics of virtualization technologies that may require additional consideration for PCI DSS Requirements.

PCI DSS Requirements (summarized and abbreviated)		Virtualization Considerations
Requirement 1: Install and maintain a firewall configuration to protect cardholder data	1.1 Establish firewall and router configuration standards.	• Due to the complexity of virtual environments, examination of multiple virtual layers may be needed to ensure that all components and data flows are identified. For example: ○ Virtual firewalls and routers could be embedded within the hypervisor, or could be implemented as virtual network devices or virtual appliances. ○ Similarly, virtual network connections could exist within a host, between hosts, and between a host and the physical network. ○ Inbound and outbound traffic to/from the CDE could include VM-to-VM interactions that never traverse the physical network. ○ Access paths between virtual systems and networks could exist across multiple levels of the virtual infrastructure—for example, between hosts, appliances or hypervisors. • Specialized solutions may be required to monitor and restrict network traffic flowing between virtual systems and networks, including wireless virtual networks. ○ Virtual network configuration changes could have significant impact—for example, a virtual component located on an in-scope network or high security zone could inadvertently be reconfigured or moved to an out-of-scope network or low-security zone. • The assignment of roles and responsibilities may be more complex in virtual environments. For example, a hypervisor administrator account could inadvertently include privileges for administering virtual networks. • Boundaries between trusted and untrusted networks may be dynamic and difficult to define in a virtual shared hosting or cloud-based infrastructure. *(continued on next page)*
	1.2 Build firewall and router configurations that restrict connections between untrusted networks and system components in the CDE. ***Note:*** *An "untrusted network" is any network that is external to the networks belonging to the entity, and/or which is out of the entity's ability to control or manage.*	
	1.3 Prohibit direct public access between the Internet and any system component in the CDE.	

PCI DSS Requirements (summarized and abbreviated)		Virtualization Considerations
	1.4 Install personal firewall software on any mobile and/or employee-owned computers with direct connectivity to the Internet, which are used to access the organization's network.	• The use of remote virtual desktops may inadvertently extend the boundaries of the CDE. ***Additional Best Practices / Recommendations:*** • Do not locate untrusted systems or networks on the same host or hypervisor as systems in the CDE. • Implement physical network segmentation to isolate any systems hosting public-facing or untrusted systems and networks from systems that host virtual components within or connected to the CDE.
Requirement 2: Do not use vendor-supplied defaults for system passwords and other security parameters	2.1 Always change vendor-supplied defaults before installing a system on the network.	• Industry-accepted system-hardening standards may not exist for all implemented virtual technologies. • A virtual component requiring higher security could unintentionally be exposed to additional risk if hosted on the same system or hypervisor as a virtual component of lower security. • Methods for securing insecure services, protocols, or daemons may be needed across multiple virtual layers. • Security parameters and settings may be unique to a particular virtual technology or implementation. • Specialized training may be needed to ensure system administrators and security personnel are knowledgeable in security for virtual technologies. • Non-console administrative access could exist across multiple levels of the virtual architecture—for example: access to hypervisors, management interfaces, and host consoles, as well as to individual VMs, appliances, and other hosted components. • Adequate separation between tenants may not be achievable in a virtual shared-hosting environment or a public cloud environment. ***Additional Best Practices / Recommendations:*** • Do not locate high-security virtual components and low-security virtual components on the same host or hypervisor.
	2.2 Develop configuration standards for all system components. Assure that these standards address all known security vulnerabilities and are consistent with industry-accepted system hardening standards.	
	2.3 Encrypt all non-console administrative access using strong cryptography.	
	2.4 Shared hosting providers must protect each entity's hosted environment and cardholder data. These providers must meet specific requirements as detailed in *Appendix A: Additional PCI DSS Requirements for Shared Hosting Providers.*	

PCI DSS Requirements (summarized and abbreviated)		Virtualization Considerations
Requirement 3: Protect stored cardholder data	3.1 Keep cardholder data storage to a minimum – implement data retention and disposal policies, procedures and processes.	• As well as being present in known locations, cardholder data could exist in archived, off-line or dormant VM images, or be unknowingly moved between virtual systems via dynamic mechanisms such as live migration or storage migration tools. • Sensitive data, such as unencrypted PAN, sensitive authentication data, and cryptographic keys, could be inadvertently captured in active memory and replicated via VM imaging and snapshot functions. • Disk-level encryption could be implemented across multiple virtual layers—for example, on the underlying host, within the VM image, or on a separate network drive that is accessible by the underlying host, hypervisor or VM image. • The use of disk encryption may be subject to specific virtualization-related implementation issues which could render the encryption ineffective. For example, moving or migrating encrypted VM images containing cardholder data to another host, VM image, or removable media could disrupt the effectiveness of the encryption mechanism. • Separating logical access to encrypted file systems from accounts across all virtual layers (including the host system, individual VMs, hypervisor accounts, etc.) adds additional levels of complexity. • Privileged accounts or processes running on the host or hypervisor could inadvertently be granted access to cryptographic keys from within a hosted component. • If cryptographic keys are stored or hosted on the same hypervisor or host as data encrypted with those keys, anyone with access to the hypervisor or host could potentially decrypt the data, rendering the data unprotected. • Specialized tools and processes may be needed to locate and manage cryptographic keys stored on archived, off-line, or relocated images. ***Additional Best Practices / Recommendations:*** • Do not virtualize critical resources used in the generation of cryptographic keys (for example, physical FIPS modules). • If key-management functions are virtualized, do not house virtual components that perform key-management functions or store cryptographic keys on the same hypervisor or host as virtual components that store or access data protected by those keys.
	3.2 Do not store sensitive authentication data after authorization (even if encrypted).	
	3.3 Mask PAN when displayed	
	3.4 Render PAN unreadable anywhere it is stored (including on portable digital media, backup media, and in logs) by using any of the following approaches: • One-way hashes based on strong cryptography (hash must be of the entire PAN) • Truncation (hashing cannot be used to replace the truncated segment of PAN) • Index tokens and pads (pads must be securely stored) • Strong cryptography with associated key-management processes and procedures	
	3.5 Protect keys used to secure cardholder data against disclosure and misuse.	
	3.6 Fully document and implement all key-management processes and procedures for cryptographic keys.	

PCI DSS Requirements (summarized and abbreviated)		Virtualization Considerations
Requirement 4: Encrypt transmission of cardholder data across open, public networks	4.1 Use strong cryptography and security protocols to safeguard sensitive cardholder data over open, public networks.	• Specialized tools may be required to secure sensitive data travelling over virtual networks from unintentional exposure.
	4.2 Never send unprotected PANs by end-user messaging technologies.	
Requirement 5: Use and regularly update anti-virus software or programs	5.1 Deploy anti-virus software on all systems commonly affected by malicious software.	• Multiple anti-virus products may be needed to protect guest operating systems as well as the underlying host operating system (for example, Windows and Linux running on the same host). • Traditional anti-virus mechanisms may interfere with certain virtualization functions. • Traditional anti-virus mechanisms may not provide adequate protection for all virtualization layers.
	5.2 Ensure that all anti-virus mechanisms are current, actively running, and generating audit logs.	

PCI DSS Requirements (summarized and abbreviated)		Virtualization Considerations
Requirement 6: Develop and maintain secure systems and applications	6.1 Ensure that all system components and software are protected from known vulnerabilities by having the latest vendor-supplied security patches installed.	• Specialized tools may be required to deploy and verify patches for virtualized components. • Patching a single host may require coordination of multiple patch schedules to address vulnerabilities across all layers, including the host system, all hosted operating systems and applications, and all virtualization-specific technologies (such as the hypervisor and management console). • Additional patch-management schedules may be needed for dormant or off-line virtual-machine images to ensure they are also protected from known vulnerabilities. • Separation of duties and access controls may need to be enforced across multiple levels—for example, at the hypervisor, individual component, and host level. • Development/test systems and data could be inadvertently moved to production environments, or vice versa, via virtual replication, imaging, or snapshot mechanisms. • Testing of changes to virtualized components may need to consider multiple levels of potential impact. ***Additional Best Practices / Recommendations:*** • Do not locate development/test systems or networks on the same host or hypervisor as production systems or networks.
	6.2 Establish a process to identify and assign a risk ranking to newly discovered security vulnerabilities.	
	6.3 Develop software applications in accordance with PCI DSS, and based on industry best practices.	
	6.4 Follow change control processes and procedures for all changes to system components.	
	6.5 Develop applications based on secure coding guidelines. Prevent common coding vulnerabilities in development processes.	
	6.6 For public-facing web applications, address new threats and vulnerabilities on an ongoing basis.	

PCI DSS Requirements (summarized and abbreviated)		Virtualization Considerations
Requirement 7: Restrict access to cardholder data by business need to know	7.1 Limit access to system components and cardholder data to individuals whose job requires such access.	• Access controls based on need to know and least privilege may need to be implemented across multiple layers to be effective (for example, at the hypervisor, host, management interface and console layer, as well as for individual virtual components, appliances and data stores). • Not all virtualization technologies are able to separate administrative access to the host or hypervisor from administrative access into individual hosted virtual components. This could result in the unauthorized or unnecessary assignment of privileged access within the hosted components. • The use of specialized tools or solutions may therefore be required to ensure effective and granular assignment of privileges across all virtualized layers.
	7.2 Establish an access control system that restricts access based on a user's need to know, and is set to "deny all" unless specifically allowed.	
Requirement 8: Assign a unique ID to each person with computer access	8.1 Assign all users a unique ID before allowing them to access system components or cardholder data	• Unique IDs and secure authentication may be needed across multiple virtual layers as well as for any intermediary technologies used to access virtualized components. • Due to the potential impact of unauthorized hypervisor access, additional authentication controls may be needed—for example, restricting all remote access to the hypervisor to defined source systems, management interfaces, and consoles. • Dormant or off-line virtual components could also contain cardholder data and may also require strong access controls. • Virtual images and snapshots could inadvertently capture passwords in active memory, resulting in unintentional and unprotected storage of the data. • Specialized solutions may be needed to ensure user authentication is applied at the appropriate level, distinguishing between authentication for individual virtual components, data stores, hypervisors, and management systems.
	8.2 In addition to assigning a unique ID, employ at least one of the following: - Something you know - Something you have - Something you are	
	8.3 Incorporate two-factor authentication for remote access to the network.	
	8.4 Render all passwords unreadable during transmission and storage on all system components using strong cryptography.	
	8.5 Ensure proper user identification and authentication management on all system components.	

PCI DSS Requirements (summarized and abbreviated)		Virtualization Considerations
Requirement 9: Restrict physical access to cardholder data	9.1 Use facility entry controls to limit and monitor physical access to systems in the CDE.	• Providing physical access to a single host or hypervisor explicitly grants the equivalent of physical access to all the virtual machines and components running on that host/hypervisor, and potential access to other connected physical systems. • Due to the potential impact of unauthorized physical access, additional authentication and monitoring of physical access may be needed—for example, requiring dual-factor authentication and a supervised escort for all physical access to the data center. • As well as being stored in known locations, cardholder data could also exist on media containing backups of virtual components, or in media containing snapshots, archived, off-line, or dormant VM images.
	9.2 Develop procedures to easily distinguish between onsite personnel and visitors.	
	9.3 Make sure all visitors are authorized.	
	9.4 Maintain a visitor log.	
	9.5 Store media back-ups in a secure location.	
	9.6 Physically secure all media.	
	9.7 Maintain strict control over the internal or external distribution of media.	
	9.8 Ensure management approves any and all media that is moved from a secured area.	
	9.9 Maintain strict control over the storage and accessibility of media.	
	9.10 Destroy media when it is no longer needed for business or legal reasons.	

PCI DSS Requirements (summarized and abbreviated)		Virtualization Considerations
Requirement 10: Track and monitor all access to network resources and cardholder data	10.1 Establish a process for linking all access to system components.	• Logging of activities unique to virtualized environments may be needed to reconstruct the events required by PCI DSS Requirement 10.2. For example, logs from specialized APIs that are used to view virtual process, memory, or offline storage may be needed to identify individual access to cardholder data. • The specific system functions and objects to be logged may differ according to the specific virtualization technology in use. • Audit trails contained within virtual machines are usually accessible to anyone with access to the virtual machine image. • Specialized tools may be required to correlate and review audit log data from within virtualized components and networks. • It may be difficult to capture, correlate, or review logs from a virtual shared hosting or cloud-based environment. ***Additional Best Practices / Recommendations:*** • Do not locate audit logs on the same host or hypervisor as the components generating the audit logs.
	10.2 Implement automated audit trails for all system components.	
	10.3 Record audit trail entries for all system components for each event.	
	10.4 Synchronize all critical system clocks and times.	
	10.5 Secure audit trails so they cannot be altered.	
	10.6 Review logs for all system components at least daily.	
	10.7 Retain audit trail history for at least one year, with a minimum of three months immediately available.	

PCI DSS Requirements (summarized and abbreviated)		Virtualization Considerations
Requirement 11: Regularly test security systems and processes	11.1 Test for the presence of wireless access points and detect unauthorized wireless access points at least quarterly.	• Network scans and testing activities may be needed across multiple virtual layers to ensure coverage of all in-scope components, including all virtual endpoints, hosts, hypervisor interfaces and management consoles. • Additional vulnerability scans may need to be scheduled for dormant/off-line virtual machine images to ensure they are also protected from known vulnerabilities. • Virtualization-specific vulnerabilities may not be detected by traditional vulnerability scanning tools. • Specialized tools may be required to scan and test virtual components and network devices from within virtual systems and networks. • The impact of changes made within a virtualized infrastructure may be complex and scanning schedules may need to be expanded accordingly. • Specialized training on the particular virtualization technologies in use may be needed for resources performing penetration tests of virtualized environments. • Specialized IDS/IPS solutions may be needed to monitor traffic flowing over virtual networks and/or between virtual systems. • Specialized tools may be needed to monitor critical files in virtualized environments. • Controls for monitoring traffic and critical files in the CDE may need to encompass dormant and off-line virtual machine images.
	11.2 Run internal and external network vulnerability scans at least quarterly and after any significant change in the network.	
	11.3 Perform external and internal penetration testing at least once a year and after any significant infrastructure or application upgrade or modification.	
	11.4 Use intrusion-detection systems and/ or intrusion-prevention systems to monitor all traffic at the perimeter of the CDE as well as at critical points inside of the CDE.	
	11.5 Deploy file-integrity monitoring tools to alert personnel to unauthorized modification of critical system files, configuration files, or content files.	

PCI DSS Requirements (summarized and abbreviated)		Virtualization Considerations
Requirement 12: Maintain a policy that addresses information security for all personnel	12.1 Establish, publish, maintain, and disseminate a security policy.	• Specific security policies, usage policies, and operational security procedures may need to be expanded to address unique aspects of virtual environments (for example, technologies implemented, type of infrastructure, deployment models, etc.). • The risk profile of a virtualized environment will be different than for a traditional physical environment. Understanding and assessing risk may require consideration of additional factors unique to a particular environment. • Additional usage policies may be needed to identify proper use of virtualization-based technologies. • Additional user training may be needed to ensure understanding of the security implications and proper use of virtualized technologies. • Details outlining specific controls and assigned responsibilities may need to be included in written agreements with service providers where cardholder data or security controls are under the control of the third party service provider in a virtual environment. • Specialized training may be needed for personnel responsible for monitoring and responding to security incidents in virtualized environments.
	12.2 Develop daily operational security procedures.	
	12.3 Develop usage policies for critical technologies and define proper use of these technologies.	
	12.4 Ensure that the security policy and procedures clearly define information security responsibilities for all personnel.	
	12.5 Assign information security management responsibilities.	
	12.6 Implement a formal security awareness program.	
	12.7 Screen potential personnel.	
	12.8 Maintain and implement policies and procedures to manage service providers.	
	12.9 Implement an incident response plan.	
Requirement A.1: Shared hosting providers must protect the CDE	A.1 Protect each entity's hosted environment and data.	• Adequate resource separation between tenants may not be achievable in a virtual shared hosting environment or a public cloud environment.