



Technical Report

PCI DSS 4.0

ONTAP 9

Matt Trudewind, NetApp
September 2022 | TR-4401

Abstract

This technical report is targeted at qualified security assessors and storage administrators focused on validating a system against the PCI DSS 4.0 standard. This document provides guidance for meeting the requirements of the controls that you apply to the NetApp® ONTAP® 9 storage system.

TABLE OF CONTENTS

Introduction to PCI DSS 3

Build and maintain a secure network and systems 3

 Requirement 1: Install and maintain network security controls3

 Requirement 2: Apply secure configurations to all system components7

Protect account data 10

 Requirement 3: Protect stored account data10

 Requirement 4: Protect cardholder data with strong cryptography during transmission over open, public networks.....13

Maintain a vulnerability management program 13

 Requirement 5: Protect all systems and networks from malicious software13

 Requirement 6: Develop and maintain secure systems and software15

Implement strong access control measures..... 17

 Requirement 7: Restrict access to system components and cardholder data by business need to know17

 Requirement 8: Identify users and authenticate access to system components19

 Requirement 9: Restrict physical access to cardholder data23

Regularly monitor and test networks 24

 Requirement 10: Log and monitor all access to system components and cardholder data24

 Requirement 11: Test security of systems and networks regularly26

Maintain an information security policy 28

 Requirement 12: Support information security with organizational policies and programs28

Where to find additional information 30

Contact us 31

Version history..... 31

LIST OF TABLES

Table 1) LIF services.....4

Table 2) Secure management firewall policy entry settings.....6

Table 3) Cluster administrator roles.9

Table 4) SVM administrator roles.....9

Introduction to PCI DSS

This technical report provides guidance and information that auditors and system operators will find useful when applying the Payment Card Industry (PCI) Data Security Standard (DSS) requirements to a storage system that runs the NetApp® ONTAP® 9 system.

ONTAP 9 separates the control plane and management plane functions (used for administration) from the data plane that is accessed by data users. This technical report focuses on administration of the system configuration in the control plane. NetApp expects that user data requirements are met by the applications that have governance over the data and not the storage systems.

This report focuses specifically on providing guidance for the PCI DSS 4.0 standard released in March 2022. PCI DSS 4.0 replaces PCI DSS v3.2.1, which is valid until it is retired on March 31, 2024.. All PCI DSS validations after that date must be validated using PCI DSS 4.0

Build and maintain a secure network and systems

Requirement 1: Install and maintain network security controls

Network security controls (NSCs), such as firewalls and other network security technologies, are network policy enforcement points that typically control network traffic between two or more logical or physical network segments (or subnets) based on predefined policies or rules.

NSCs examine all network traffic entering (ingress) and leaving (egress) a segment and decide, based on the policies defined, whether the network traffic is allowed to pass or whether it is rejected. Typically, NSCs are placed between environments with different security needs or levels of trust. However, in some environments, NSCs control the traffic to individual devices irrespective of trust boundaries. Policy enforcement generally occurs at layer 3 of the OSI model, but data present in higher layers is also frequently used to determine policy decisions.

Traditionally, this function is provided by physical firewalls; however, now this functionality can be provided by virtual devices, cloud access controls, virtualization or container systems, and other software-defined networking technology.

NSCs are used to control traffic within an entity's own networks—for example, between highly sensitive and less sensitive areas—and to protect the entity's resources from exposure to untrusted networks. The cardholder data environment (CDE) is an example of a more sensitive area within an entity's network. Often, seemingly insignificant paths to and from untrusted networks can provide unprotected pathways into sensitive systems. NSCs provide a key protection mechanism for any computer network.

Common examples of untrusted networks include the internet, dedicated connections such as business-to-business communication channels, wireless networks, carrier networks (such as cellular), third-party networks, and other sources outside the entity's ability to control. Untrusted networks also include corporate networks that are considered out-of-scope for PCI DSS, because they are not assessed, and must be treated as untrusted because the existence of security controls is not verified. Although an entity might consider an internal network trusted from an infrastructure perspective, if a network is out of scope for PCI DSS, that network must be considered untrusted for PCI DSS.

ONTAP storage systems must be installed behind and protected by appropriate NSC's, such as an external firewall or other software defined network technology, to conform with the principles of PCI DSS. In addition, ONTAP provides basic NSC and firewall functions for controlling management access to services. It is on by default in each node and protects the entire cluster. Depending on the ONTAP version, NSC's available in ONTAP include service policies and a built in firewall. These are not designed to replace appropriate NSC's, such as a dedicated firewall, but instead provide an additional layer of internal protection by permitting or blocking protocols as required .

Depending on the ONTAP version, each storage virtual machine (SVM, formerly known as Vserver) management LIF can have a firewall policy, a service policy, or a combination of both attached to it. Beginning in ONTAP 9.10.1, firewall service policies are deprecated and are replaced by LIF service policies. In earlier releases, the onboard firewall is managed by using firewall policies. This functionality is accomplished in ONTAP 9.10.1 and later by using a LIF service policy.

Service policies define a collection of network services supported by a LIF. ONTAP provides a set of built-in service policies that you can associate with a LIF. These built-in service policies include management services such as HTTPS, Secure Shell (SSH), DNS, and Lightweight Directory Access Protocol (LDAP).

Firewall policy entries consist of a protocol type, a firewall action, and a subnet or specific IP address to which the action applies. You can apply the firewall policy to each LIF for all traffic going through that interface. Also, the firewall policy can vary for each SVM.

NetApp recommends that an SVM storing payment data has the strictest possible settings; that is, permit only the protocols and specific subnets or IP addresses needed to manage the SVM. For more information on service and firewall policies, see the [ONTAP 9 Network Management Guide](#).

Evolving requirements from previous PCI DSS 3.2.1 standard

- Update to the principal requirement title to reflect the focus on “network security controls.” Replacement of the terms “firewalls” and “routers” with “network security controls” to support a broader range of technologies used to meet the security objectives traditionally met by firewalls.
- Replacement of the requirement for “Description of groups, roles, and responsibilities for management of network components” with a general requirement for roles and responsibilities for Requirement 1.

Additional guidance and clarification from previous PCI DSS 3.2.1 standard

- Refocus of a former null requirement (all content pointed to other requirements) about defining, implementing, and maintaining configuration standards for network security control rule sets (1.2.1).
- Clarification that changes are managed in accordance with the change control process defined at Requirement 6.5.1 (1.2.2).
- Clarification of the intent of reviewing configurations of network security controls at least once every six months (1.2.7).
- Refocus of a former null requirement (all content pointed to other requirements). Clarification that the intent is to implement controls between trusted and untrusted networks (1.4.1).

Implications for data storage

Best practice

Enable firewall and service policies in ONTAP to augment NSC’s such as external firewalls and other software defined network technologies.

A service policy to enable management traffic on a cluster SVM is defined by using the following command:

```
network interface service-policy create -vserver <svm_name> -policy <service_policy_name> -
services <service_name> -allowed-addresses <IP_address/mask, ...>
```

Table 1) LIF services.

Service	Failover limitations	Description
intercluster-core	home-node-only	Core intercluster services

Service	Failover limitations	Description
management-core	home-node-only	Core management services
management-ssh	home-node-only	Services for SSH management access
management-http	home-node-only	Services for HTTP management access
management-https	home-node-only	Services for HTTPS management access
management-autosupport	home-node-only	Services related to posting AutoSupport payloads
management-bgp	home-port-only	Services related to BGP peer interactions
backup-ndmp-control	home-port-only	Services for NDMP backup controls
management-ems	home-port-only	Services for management messaging access
management-ntp-client	home-port-only	Introduced in ONTAP 9.10.1. Services for NTP client access.
management-ntp-server	home-port-only	Introduced in ONTAP 9.11.1. Services for NTP server management access
management-portmap	home-port-only	Services for portmap management
management-rsh-server	home-port-only	Services for rsh server management
management-snmpserver	home-port-only	Services for SNMP server management
management-telnetserver	home-port-only	Services for telnet server management

The list in Table 1 provides the list of services that LIFs can use on a system SVM beginning in ONTAP 9.11.1. You can add services to a custom service policy to define which additional NSC you must apply to the system SVM's. For more information, see [Network management ONTAP](#). As an example of service policy management, you can use the following command to create a new custom service policy named `secure_management` and specify the HTTPS service. This command enables you to specify each service that you want to apply by using the custom policy.

Example: Create service policy 'secure_management'

```
Cluster1::> network interface service-policy create -vserver cluster1 -policy secure-management -
services management-https -allowed-addresses 10.2.0.0/16
```

When completed, you must assign the new policy to a LIF by using the following command:

```
Cluster1::> network interface modify -vserver cluster1 -lif lif1 -service-policy secure-
management
```

Note: You can also specify a service policy during LIF creation.

Note the following changes beginning with ONTAP 9.10.1:

- Firewall service policies are deprecated and are replaced by LIF service policies. In earlier releases, the onboard firewall is managed by using firewall policies. In ONTAP 9.10.1, the onboard firewall is managed by using a LIF service policy.
- All firewall policies are empty and do not open any ports in the underlying firewall. Instead, you must open all ports by using a LIF service policy.
- No action is required after you upgrade to ONTAP 9.10.1 or later to transition from firewall service policies to LIF service policies. The system automatically constructs LIF service policies consistent with the firewall service policies used in in the previous ONTAP release. If you use scripts or other tools that create and manage custom firewall policies, you might need to upgrade those scripts to create custom service policies instead.

For more information, see [Network management ONTAP 9](#). Although service policies are available for commonly used management services, in ONTAP 9.10.1 and earlier, firewall policies are still required. You can define a secured firewall policy for management interfaces by using the following command:

```
system services firewall policy create -vserver <SVM name> -policy <policy-name> -service <protocol_name> -allow-list <ip_address/mask>
```

Note: The first command reference to a new policy creates the policy. Subsequent references add additional entries to that policy.

Table 2) Secure management firewall policy entry settings.

Protocol	-allow-list	Net effect
dns	IP address list for allowed DNS servers	Allow DNS access to list
http	127.0.0.1/32, ::1/128	Deny all
https	IP address list for allowed HTTPS administrators	Allow HTTPS access to list
ndmp	127.0.0.1/32, ::1/128	Deny all
ntp	IP address list for allowed NTP servers	Allow NTP access to list
rsh	127.0.0.1/32, ::1/128	Deny all
snmp	IP address list for allowed SNMP management stations	Allow SNMP access to list
ssh	IP address list for allowed SSH clients	Allow SSH access from list
telnet	127.0.0.1/32, ::1/128	Deny all

Table 2 provides the most common interfaces and protocols. Depending on the ONTAP version, some protocols might be unavailable if they have been replaced by service policies. For more information, see [Network management ONTAP 9](#). As an example of firewall policy management, the following command is used to create a policy named `secure_mgmt`. This command is repeated with appropriate modifications for each entry in Table 1.

Example: Create firewall policy 'secure_mgmt'

```
Cluster1::> system services firewall policy create -vserver cDOT-1 -policy secure_mgmt -service dns -allow-list 10.63.165.0/24, ::1/128
```

When completed, you can verify the policy action entries by using the following command:

```
Cluster1::> system services firewall policy show -policy secure_mgmt
```

The next step applies the firewall policy, created in the previous command, to the interfaces on the cluster and node management interfaces (e0M). Perform this step for the cluster SVM and each cluster node SVM.

For the overall cluster SVM, use the following command as an example:

```
Cluster1::> network interface modify -vserver cDOT-1 -lif cluster_mgmt -firewall-policy secure_mgmt
```

For the user data storage SVMs, use the following commands as examples:

```
Cluster1::> network interface modify -vserver cDOT-1-01 -lif mgmt1 -firewall-policy secure_mgmt
Cluster1::> network interface modify -vserver cDOT-1-02 -lif mgmt1 -firewall-policy secure_mgmt
```

For more information, see [Network management ONTAP 9](#). In addition to the service and firewall policies just described, the following recommendations enhance security:

- The ports used for intracluster traffic must be on a private isolated network in the cluster. The IP subnet used to access the cluster must not be visible on the public internet. Use a private IP subnet (such as 10.10.x.x) that is not accessible outside the secure trusted network. This helps to minimize exposure outside the network. In addition, because ONTAP separates the control plane from the user data plane, it is possible to enhance security further by putting data traffic on a separate VLAN from control and administrative traffic. You can also perform this separation on a per-SVM basis, if desired.
- ONTAP uses network services such as Network Time Protocol (NTP), DNS, and others. NetApp recommends that these services be provided either by internal network sources or by proxies on the trusted network rather than exposing the system to an untrusted network such as the internet. Again, these services must also be provided on the management IP subnet (VLAN), separate from the data traffic.
- The Service Processor (SP) or baseboard management controller (BMC) must be in the VLAN for the management plane. The SP enables you to remotely access and administer the storage system and diagnose error conditions. It is part of the attack surface that must be protected because it provides an additional point of entry into the system,. Along with the service policy and firewall-type functions described earlier, there is a white-list approach for IP addresses in the SP.

In addition, the range of IP addresses allowed to access the SP can be limited. The `system service-processor ssh add-allowed-addresses -allowed-addresses` command permits IP addresses SSH access to the SP.

For detailed information about the SP and its capabilities, see “Managing the IP Addresses That Can Access the SP” in the [Cluster administration for ONTAP 9](#).

Requirement 2: Apply secure configurations to all system components

Malicious individuals, both external and internal to an entity, often use default passwords and other vendor default settings to compromise systems. These passwords and settings are well known and are easily obtained by using public information.

Applying secure configurations to system components reduces the means available to an attacker to compromise the system. Changing default passwords, removing unnecessary software, functions, and accounts, and disabling or removing unnecessary services all help to reduce the potential attack surface.

Cluster administrators administer the overall cluster and can create SVMs. The cluster administrator can then assign resources (aggregates and volumes) to those SVMs. In each SVM, the SVM administrator administers the data storage for that SVM. Because the SVM represents a separate storage machine, each SVM can be set up in a separate network and security domain, managed by the SVM administrator. The SVM is accessed (both control plane and data plane) through the LIF interfaces assigned to that SVM by the cluster administrator.

ONTAP is a hardened appliance (either virtual like NetApp Cloud Volumes ONTAP and ONTAP Select or physical like NetApp AFF and FAS) that does not have any unnecessary services running by default. All services on the ONTAP system are for the purpose of data storage. ONTAP takes a secure by default posture.

When a new system is installed, there is no factory default password for the cluster administrator. During initial setup, you are asked to set the password for these accounts by using the serial port. Use a

password of sufficient length (10 or more characters) and complexity (uppercase/lowercase, special characters) to avoid brute-force guessing attacks.

When cluster administrators create an SVM, they can create a password for the SVM administrator. To require the SVM administrator to change the password immediately, cluster administrators can apply password expiry to the SVM administrator role.

You can display the list of active accounts with the following command:

```
security login show
```

For an initial installation, the following accounts are typically available as the default:

- One cluster admin account with access to:
console, ontapi (NetApp Manageability SDK), http, service-processor, and ssh
- At least one vservers admin account with access to:
ontapi, ssh

You can use the following command to reset a password:

```
cluster1::> security login password -username admin -vservers vs
```

There are two default administrator accounts at the cluster level: admin and diag. The diag account provides low-level system access and is ordinarily never needed. It is disabled by default and must never be used except under the direction of NetApp Support personnel. The admin account is configured with a role (through role-based access control, or RBAC) to have access to all commands necessary to manage the system. You must lock it and restrict access to the respective commands to accounts with appropriate roles. Create a duplicate account for additional protection and you can also delete the admin account from the system. To do this, create a custom role by using the `security login role create` command.

To further enhance security, you can enforce automatic lockout for invalid logins by configuring the role with the `max-failed-attempts` attribute.

The ONTAP system can use SNMP for monitoring and management. For added security, ONTAP supports SNMPv3, which includes authentication and encryption of the SNMP messages. NetApp recommends configuring for both authentication and encryption by using the `security login create` command to create an SNMP user with parameters for authentication and privacy (encryption). For more information about configuring and using SNMP, see [TR-4220: SNMP Support in Data ONTAP](#).

Evolving requirements from previous PCI DSS 3.2.1 standard

- New requirement for roles and responsibilities. This requirement is effective immediately for all v4.0 assessments

Additional guidance and clarification from previous PCI DSS 3.2.1 standard

- Update of the principal requirement title to reflect that the focus is on secure configurations in general, and not just on vendor-supplied defaults.
- Clarification of the intent of the requirement for managing primary functions that require different security levels.
- Clarification that the intent of the requirement is if any insecure services, protocols, or daemons are present.
- Implications for data storage.

You can configure ONTAP to meet these guidelines.

Best practice:

Configure roles and accounts for ONTAP by using the principle of least privilege. The ONTAP operating system is already a hardened appliance that does not have any unnecessary services running by default.

There are two steps to managing admin user accounts. First, roles are established that authorize user capabilities. Second, user accounts are created and assigned to that role. Predefined roles are provided in ONTAP as shown in Table 3 for cluster administrator accounts.

Table 3) Cluster administrator roles.

Role	Default capabilities
Admin	<ul style="list-style-type: none"> All (read, write) access to all command directories
AutoSupport	<ul style="list-style-type: none"> All access to set System mode AutoSupport No other command directories
Backup	<ul style="list-style-type: none"> All access to Vserver services NDMP Read-only access to volumes No access to other command directories
Readonly	<ul style="list-style-type: none"> All access to security login passwords All access to set Read-only access to all other command directories
None	<ul style="list-style-type: none"> No access to any command directories

As Table 3 shows, the admin account is all powerful. Other accounts provide more limited capabilities, and there is a `none` account with no capabilities. These roles provide a starting point for creating custom roles by adding or deleting command directories. For more powerful custom user roles, you can start with the admin role and subtract command directories that are not needed. Conversely, to create a very limited custom role, it might be easier to start with the none role and add the needed command directories. The AutoSupport, Backup, and Readonly roles cover special use cases.

Like the cluster administrator roles, there are four default roles for SVM administrators, as shown in Table 4.

Table 4) SVM administrator roles.

Role	Default capabilities
Vsadmin	<ul style="list-style-type: none"> Manage own administrator account, local password, and public key Manage volumes, quotas, qtrees, NetApp Snapshot™ copies, NetApp FlexCache® files, and files Manage LUNs Configure protocols Configure services Monitor network connections and network interface Monitor the health of an SVM
Vsadmin-volume	<ul style="list-style-type: none"> Manage volumes, quotas, qtrees, FlexCache files, and files Manage LUNs Configure protocols Configure services Monitor network interface

Role	Default capabilities
	<ul style="list-style-type: none"> Monitor the health of an SVM
Vsadmin-protocol	<ul style="list-style-type: none"> Configure protocols Configure services Manage LUNs Monitor network interface Monitor the health of an SVM
Vsadmin-readonly	<ul style="list-style-type: none"> Monitor the health of an SVM Monitor network interface View volumes and LUNs View services and protocols

Notice that the default role list for SVM administrators assumes a separation of duties. For example, the responsibility for managing and configuring data protocol services is separated from the responsibility to manage the storage itself. This approach conforms to the least-privilege principle of common security best practices.

These roles are sufficient for PCI DSS 4.0 To modify or create new roles, see [ONTAP 9 security and data encryption guide](#). After you establish the appropriate roles, you can create user accounts and assign the roles to them. Use `security login create` to create an account and assign a role to it. For example, the following command creates a login that has the user name `monitor`, the application “`ssh`,” the authentication method `password`, and the access-control role `guest` for Vserver `vs`.

```
cluster1::> security login create -username monitor -application ssh -authmethod password -role
guest -vservers vs
```

The default accounts for cluster administrator and SVM administrator are assigned typical roles used by most customers by default. However, as an option, you might decide to lock those accounts and create new accounts with more restricted roles in accordance with the least-privilege security philosophy.

Another option that you can use to further secure the cluster administrator and all other user accounts on the ONTAP system is to use the security feature built into ONTAP 9.11.1 and later called multi-admin verification (MAV). You can use MAV to ensure that certain operations, such as deleting volumes or Snapshot copies, are executed only after approvals from designated administrators. This prevents compromised, malicious, or inexperienced administrators from making undesirable changes or deleting data. To learn more about MAV, see [ONTAP 9 security and data encryption guide](#).

The ONTAP system also provides web services access to the system to users and applications. NetApp recommends configuring the system to use only HTTPS and disabling HTTP. In ONTAP 9.11.1 and later, TLS 1.3 and 1.2 are enabled and TLS 1.1, 1.0, and SSLv3 are disabled by default. NetApp recommends TLS 1.3, but TLS 1.2, 1.1, 1.0 and SSLv3 are provided for backward compatibility.

Protect account data

Requirement 3: Protect stored account data

Protection methods such as encryption, truncation, masking, and hashing are critical components of account data protection. If an intruder circumvents other security controls and gains access to encrypted account data, the data is unreadable without the proper cryptographic keys and is unusable to that intruder. Other effective methods of protecting stored data must also be considered as potential risk-mitigation opportunities. For example, methods for minimizing risk include not storing account data unless necessary, truncating cardholder data if full primary account number (PAN) is not needed, and not

sending unprotected PANs using end-user messaging technologies such as e-mail and instant messaging.

If account data is present in non-persistent memory (for example, RAM, volatile memory), encryption of account data is not required. However, proper controls must be in place to ensure that memory maintains a non-persistent state. Data must be removed from volatile memory after the business purpose (for example, the associated transaction) is complete. If data storage becomes persistent, all applicable PCI DSS requirements apply, including encryption of stored data.

Requirement 3 applies to protection of stored account data unless specifically stated otherwise in an individual requirement.

Evolving requirements from previous PCI DSS 3.2.1 standard

- New requirement for roles and responsibilities. This requirement is effective immediately for all v4.0 assessments.
- New requirement bullet to address stored account data stored before completion of authorization through the implementation of data retention and disposal policies, procedures, and processes. This bullet is a best practice until 31 March 2025. (3.2.1)
- New requirement to encrypt stored account data that is stored electronically before completion of authorization. This requirement is a best practice until 31 March 2025. (3.3.2)
- Clarification that PAN is masked when displayed so that only personnel with a business need can see more than the BIN/last four digits of the PAN (3.4.1).
- New requirement for technical controls to prevent copy and relocation of PAN when using remote-access technologies. Expanded from former Requirement 12.3.10. This requirement is a best practice until 31 March 2025. (3.4.2)
- Removal of pads from the “Index tokens and pads” bullet for rendering PAN unreadable (3.5.1).
- New requirement for keyed cryptographic hashes when hashing is used to render PAN unreadable. This requirement is a best practice until 31 March 2025 (3.5.1.1).
- New requirement that disk-level or partition-level encryption is used only to render PAN unreadable on removable electronic media or, if used on non-removable electronic media, the PAN is also rendered unreadable through a mechanism that meets Requirement 3.5.1. This requirement is a best practice until 31 March 2025 (3.5.1.2).
- New requirement bullet for service providers only to include in the documented description of the cryptographic architecture that the use of the same cryptographic keys in production and test environments is prevented. This bullet is a best practice until 31 March 2025 (3.6.1.1).
- Additional guidance and clarification from previous PCI DSS 3.2.1 standard
- Addition of a requirement to address former testing procedures that any storage of stored account data by issuers is limited to what is needed for a legitimate issuing business need and is secured (3.3.3).

Implications for data storage

PCI DSS 4.0 states that you can perform encryption in one of two ways: as partition-level encryption or as disk-level encryption, and proper key-management practices are required. You can use disk encryption meet the PCI DSS 4.0 requirement.

NetApp Storage Encryption (NSE) is a NetApp implementation of full disk encryption (FDE) , using FIPS-140-2 level 2 validated self-encrypting drives (SEDs) from leading vendors. NetApp Volume Encryption (NVE) encrypts at the volume level, enabling the encryption capability to exist independently of the physical media: SSDs, NetApp AFF, or even NSE drives. NetApp Aggregate Encryption (NAE) encrypts at the aggregate level, allowing all volumes in the aggregate to share encryption and take full advantage of ONTAP storage efficiencies such as deduplication and compression.

NSE is an encryption implementation that provides comprehensive, cost-effective, hardware-based security that is simple to use. This single-source solution can increase overall compliance with industry and government regulations without compromising storage efficiency.

NSE includes the following features

- Supports the entire suite of storage efficiency technologies from NetApp, including deduplication, compression, and array-based antivirus scanning.
- Supports third-party external key management servers provided by NetApp partners like Thales, Entrust, IBM, and more. For a complete list, see the [NetApp Interoperability Matrix Tool](#).
- Helps customers comply with FISMA, HIPAA, PCI, Basel II, SB 1386, and EU Data Protection Directive 95/46/EC and EU General Data Protection Regulations by using FIPS 140-2 validated hardware.
- Complies with the OASIS KMIP standard, offering compatibility with other key managers and encryption devices.
- Supports hyperscaler key management solutions (KMS).

NVE and NAE are software-based, data-at-rest encryption solutions available in ONTAP management software. Both NAE and NVE are FIPS 140-2 compliant. NVE enables ONTAP to encrypt data (using AES 256-bit encryption) per volume for granularity. NAE provides shared AES-256-bit encryption keys across all volumes in the aggregate. You can store NAE aggregates and NVE volume encryption keys on an external key manager. If the controller and disks are moved without access to the external key manager, the NAE aggregates and NVE volumes are inaccessible and cannot be decrypted.

NVE includes the following features:

- Supports the entire suite of storage efficiency technologies from NetApp, including deduplication, compression, and array-based antivirus scanning (NVE aggregate inline deduplication volumes are excluded from this efficiency).
- Supports third party external key management servers provided by NetApp partners like Thales, Entrust, IBM, and more. For a complete list, see the [NetApp Interoperability Matrix Tool](#).
- Helps customers comply with FISMA, HIPAA, PCI, Basel II, SB 1386, and EU Data Protection Directive 95/46/EC and EU General Data Protection Regulations by using [FIPS 140-2 validated cryptographic module software](#). Complies with the OASIS KMIP standard, offering compatibility with other key managers and encryption devices.
- Supports hyperscaler key management solutions (KMS).

NAE includes the following features:

- Supports the entire suite of storage efficiency technologies from NetApp, including aggregate inline deduplication, deduplication, compression, and array-based antivirus scanning
- Supports third-party external key management servers provided by NetApp partners such as Thales, Entrust, IBM, and more. For a complete list, see the [NetApp Interoperability Matrix Tool](#).
- Helps customers comply with FISMA, HIPAA, PCI, Basel II, SB 1386, and EU Data Protection Directive 95/46/EC and EU General Data Protection Regulations by using [FIPS 140-2 validated cryptographic module software](#)
- Complies with the OASIS KMIP standard, offering compatibility with other key managers and encryption devices.
- Supports hyperscaler key management solutions (KMS).

Best practice

To store payment data, use NetApp Storage Encryption, NetApp Aggregate Encryption, NetApp Volume Encryption, or a combination of both..

Requirement 4: Protect cardholder data with strong cryptography during transmission over open, public networks

The use of strong cryptography provides greater assurance in preserving data confidentiality, integrity, and nonrepudiation.

To protect against compromise, you must encrypt PAN during transmission over networks that are easily accessed by malicious individuals, including untrusted and public networks. Misconfigured wireless networks and vulnerabilities in legacy encryption and authentication protocols continue to be targeted by malicious individuals aiming to exploit these vulnerabilities to gain privileged access to CDEs. Any transmissions of cardholder data over an entity's internal network or networks brings that network into scope for PCI DSS because that network stores, processes, or transmits cardholder data. Any such networks must be evaluated and assessed against applicable PCI DSS requirements.

Requirement 4 applies to transmissions of PAN unless specifically called out in an individual requirement.

You can protect PAN transmissions by encrypting the data before it is transmitted, by encrypting the session over which the data is transmitted, or both. Although it is not required that you apply strong cryptography at both the data level and the session level, it is recommended.

Evolving requirements from previous PCI DSS 3.2.1 standard

- New requirement for roles and responsibilities This requirement is effective immediately for all v4.0 assessments (4.1.2).
- New requirement to confirm certificates used for PAN transmissions over open, public networks are valid and not expired or revoked. This requirement is a best practice until 31 March 2025 (4.2.1).
- New requirement to maintain an inventory of trusted keys and certificates. This requirement is a best practice until 31 March 2025.

Additional guidance and clarification from previous PCI DSS 3.2.1 standard

- Update of the principal requirement title to reflect the focus on strong cryptography to protect transmissions of cardholder data.

Implications for data storage

Because of the varied requirements for data encryption, NetApp recommends the use of external VPN encryption for the transmission of both cardholder data and management data across public networks.

Best practice

Use external VPN data encryption to transmit cardholder data. For NAS protocols, use Kerberos 5 authentication with privacy service (krb5p) for NFS and SMB encryption. Internet Protocol security (IPsec) is also available for all IP data traffic. For more information on krg5p, SMB Encryption, and IPsec see [Security hardening guide for NetApp ONTAP 9](#).

Maintain a vulnerability management program

Requirement 5: Protect all systems and networks from malicious software

Malicious software (malware) is software or firmware designed to infiltrate or damage a computer system without the owner's knowledge or consent with the intent of compromising the confidentiality, integrity, or availability of the owner's data, applications, or operating system.

Examples include viruses, worms, Trojans, spyware, ransomware, keyloggers, rootkits, malicious code, scripts, and links.

Malware can enter the network during many business-approved activities, including employee email (for example, through phishing) and use of the internet, mobile computers, and storage devices, resulting in the exploitation of system vulnerabilities.

Using anti-malware solutions that address all types of malware helps to protect systems from current and evolving malware threats.

Evolving requirements from previous PCI DSS 3.2.1 standard

- Replacement of the term “anti-virus” with “anti-malware” throughout to support a broader range of technologies used to meet the security objectives traditionally met by anti-virus software.
- New requirement for roles and responsibilities. This requirement is effective immediately for all v4.0 assessments (5.1.2).
- New requirement to define the frequency of periodic evaluations of system components not at risk for malware in the entity’s targeted risk analysis. This requirement is a best practice until 31 March 2025 (5.2.3.1).
- New requirement to define the frequency of periodic malware scans in the entity’s targeted risk analysis. This requirement is a best practice until 31 March 2025 (5.3.2.1).
- New requirement for a malware solution for removable electronic media. This requirement is a best practice until 31 March 2025 (5.3.3).
- New requirement to detect and protect personnel against phishing attacks. This requirement is a best practice until 31 March 2025 (5.4.1).

Additional guidance and clarification from previous PCI DSS 3.2.1 standard

- Update to the principal requirement title to reflect the focus on protecting all systems and networks from malicious software.
- Clarification of the requirement by changing focus to “system components that are not at risk for malware” (5.2.3).
- Split of one requirement into three to focus each requirement on one area:
 - Keep the malware solution current by using automatic updates (5.3.1).
 - Perform periodic scans and active or real-time scans (with a new option for continuous behavioral analysis) (5.3.2).
 - Generate audit logs by using the malware solution (5.3.4).

Implications for data storage

NetApp recommends that anti malware software is run on the computers used to provide card payment services.

For additional security protection, ONTAP systems can also support protection from malware (such as ransomware) on NAS (NFS and SMB) connected shares. For example, NetApp FPolicy™ in combination with NetApp Cloud Insights, or similar capabilities from our partners, does an excellent job of detecting malware through user behavioral analytics (UBA). It looks for potential malware attacks from the aspect of an individual user’s behavior. Hijacking a single user account is just one avenue a hacker might take when launching a malware attack; malicious actors are constantly evolving their attack techniques.

NetApp Active IQ® and NetApp Active IQ Unified Manager also provide additional layers of detection for ransomware. Active IQ checks ONTAP systems for adherence to NetApp configuration best practices, such as enabling FPolicy. Active IQ Unified Manager generates alerts for abnormal growth of Snapshot copies or storage efficiency loss, which can indicate potential ransomware attacks. In ONTAP 9.10.1 and later, the anti-ransomware on box feature leverages built-in on-box machine learning (ML) that uses volume workload activity and data entropy to automatically detect ransomware. It monitors activity that is different from UBA so that it can detect attacks that UBA does not.

Best practice

Deploy a layered defense approach against malware to protect payment card services with NetApp technologies such as FPolicy, Cloud Insights, on box anti-ransomware protection, and much more. For more detailed information, see [TR-4572 The NetApp solution for ransomware](#).

Requirement 6: Develop and maintain secure systems and software

Actors with bad intentions can use security vulnerabilities to gain privileged access to systems. Many of these vulnerabilities are fixed by vendor provided security patches that must be installed by the entities that manage the systems. All system components must have all appropriate software patches to protect against the exploitation and compromise of account data by malicious individuals and malicious software.

Appropriate software patches are patches that have been evaluated and tested to determine that the patches do not conflict with existing security configurations. For bespoke and custom software, you can avoid numerous vulnerabilities by applying software lifecycle (SLC) processes and secure coding techniques.

Code repositories that store application code, system configurations, or other configuration data that can impact the security of account data or the CDE are in scope for PCI DSS assessments.

For information about the use of PCI SSC-validated software and software vendors and how the use of PCI SSC's software standards might help with meeting the controls in Requirement 6, see page 7 of ["Relationship between PCI DSS and PCI SSC Software Standards"](#).

Note: Requirement 6 applies to all system components, except for section 6.2 for developing software securely, which applies only to bespoke and custom software used on any system component included in or connected to the CDE.

Evolving requirements from previous PCI DSS 3.2.1 standard

- New requirement for roles and responsibilities. This requirement is effective immediately for all v4.0 assessments (6.1.2).
- New requirement to maintain an inventory of bespoke and custom software. This requirement is a best practice until 31 March 2025 (6.3.2).
- New requirement to deploy an automated technical solution for public-facing web applications that continually detects and prevents web-based attacks. This new requirement removes the option in Requirement 6.4.1 to review web applications by using manual or automated application vulnerability assessment tools or methods. This requirement is a best practice until 31 March 2025 (6.4.2).
- New requirement for management of all payment page scripts that are loaded and executed in the consumer's browser. This requirement is a best practice until 31 March 2025 (6.4.3).

Additional guidance and clarification from previous PCI DSS 3.2.1 standard

- Update of the principal requirement title to include "software" rather than "applications." Clarification that Requirement 6 applies to all system components, except for Requirement 6.2, which applies only to bespoke and custom software.
- Replacement of the term "internal and external" with "bespoke and custom" software. Clarification that this requirement applies to software developed for or by the entity for the entity's own use and does not apply to third-party software (6.2.1).
- Move of elements of Requirement 6.5 for training of software developers to align all software development content under Requirement 6.2. Clarification of training requirements for software development personnel (6.2.2).
- Move of requirement for reviewing custom software before release to align all software development content under Requirement 6.2. Split of the requirement to separate general code review practices from those needed if manual code reviews are performed (6.2.3, 6.2.3.1).

- Move of requirements for addressing common coding vulnerabilities to align all software development content under Requirement 6.2. Combining of methods to prevent or mitigate common software attacks into a single requirement and generalization of the language describing each type of attack (6.2.4).
- Addition of a bullet to clarify applicability to vulnerabilities for bespoke and custom and third-party software (6.3.1).
- Removal of the requirement for specific documented procedures and added testing procedures to verify policies and procedures to each related requirement.
- Replacement of the term “development/test and production” with “production and pre-production” environments (6.5.3).
- Replacement of the term “development/test and production” with “production and pre-production” environments. Replacement of the term “separation of duties” and clarification that separation of roles and functions between production and preproduction is intended to provide accountability so that only approved changes are deployed (6.5.4).
- Replacement of the term “testing or development” with “pre- production” environments.
- Clarification that live PANs are not used in preproduction environments except where all applicable PCI DSS requirements are in place (6.5.5).

Implications for data storage

NetApp follows secure development principles throughout our product development lifecycle. NetApp expands and improves on the secure-development programs on a continuing basis. As a part of NetApp standard procedures, we implement secure design principles, developer training, and extensive testing programs.

NetApp follows a multistep process when responding to vulnerabilities and when notifying customers.

- **Vulnerability report received.** NetApp encourages customers and researchers to use PGP-encrypted emails to transmit confidential details to our Vulnerability Response team (PSIRT). NetApp investigates a suspected vulnerability in our products and confirms receipt of the vulnerability report within seven business days.
- **Verification.** After a finder has initiated contact with NetApp regarding a potential vulnerability, NetApp PSIRT engineers verify the vulnerability and provide assessment within the Common Vulnerability Scoring System (CVSS) framework.
- **Resolution development.** NetApp strives to deliver critical fixes and mitigations to the customer base as rapidly as our stringent quality-control standards allow; testing and verification is often a time-intensive process.
- **Notification.** NetApp discloses the minimum amount of information required for customers to assess the impact of a vulnerability in their environment, as well as any steps required to mitigate the threat. NetApp does not intend to provide details that could enable a malicious actor to develop an exploit.
- **Attribution.** NetApp credits external vulnerability discoverers in the advisory if they have provided explicit consent to be identified, and if they provide NetApp the opportunity to remediate and notify our customer base before making the vulnerability public.

To standardize the description of each public vulnerability, NetApp security advisories reference a Common Vulnerabilities and Exposures (CVE) ID. NetApp uses version 3.0 of CVSS to determine vulnerability priority and notification strategy.

NetApp security advisories and notices include the NetApp determined base vulnerability score. We encourage customers who use CVSS for vulnerability classification and management to compute their own temporal and environmental scores to take full advantage of the CVSS metrics.

The following are standard delivery methods for NetApp security information:

- **Security advisory.** Significant security vulnerabilities that directly affect NetApp products and require an upgrade, patch, or direct customer action to remediate.

NetApp Active IQ provides information on potential security risks listed in security advisories. These risks correspond to CVE IDs that are posted by NetApp PSIRT. See the Health Summary section under Security Vulnerability on the [Active IQ](#) landing page and in the security card in Digital Advisor. For more information, see [Active IQ](#).

- **Security bulletin.** Low and medium severity security issues that affect NetApp products.
- **Security notices.** Can be used when a third party makes an unconfirmed public statement about a perceived NetApp product vulnerability, or NetApp products are unofficially implicated in security incidents.
- **Security bug reports.** Information about low-severity security vulnerabilities, available through Bugs Online (requires login).

For more information, see [NetApp Product Security](#).

Best practice

Subscribe to NetApp notifications or use NetApp Active IQ and implement security patches and updates as they are made available. To subscribe, go to [NetApp Security Advisories](#). Find more information about Active IQ on the [Active IQ](#) landing page.

Implement strong access control measures

Requirement 7: Restrict access to system components and cardholder data by business need to know

Unauthorized individuals might gain access to critical data or systems due to ineffective access control rules and definitions. To ensure critical data can only be accessed by authorized personnel, systems and processes must be in place to limit access on a need to know basis and according to job responsibilities.

Access or access rights are created by rules that provide users access to systems, applications, and data, while privileges allow a user to perform a specific action or function in relation to that system, application, or data. For example, a user might have access rights to specific data, but whether they can only read that data, or can also change or delete the data is determined by the user's assigned privileges.

"Need-to-know" refers to providing access to only the least amount of data needed to perform a job. "Least privileges" refers to providing only the minimum level of privileges needed to perform a job.

These requirements apply to user accounts and access for employees, contractors, consultants, and internal and external vendors and other third parties (for example, for providing support or maintenance services). Certain requirements also apply to application and system accounts used by the entity (also called service accounts).

Note: These requirements do not apply to consumers (cardholders).

Evolving requirements from previous PCI DSS 3.2.1 standard

- New requirement for roles and responsibilities. This requirement is effective immediately for all v4.0 assessments (7.1.2).
- New requirement for review of all user accounts and related access privileges. This requirement is a best practice until 31 March 2025 (7.2.4).
- New requirement for assignment and management of all application and system accounts and related access privileges. This requirement is a best practice until 31 March 2025 (7.2.5).

- New requirement for review of all access by application and system accounts and related access privileges. This requirement is a best practice until 31 March 2025 (7.2.5.1).

Additional guidance and clarification from previous PCI DSS 3.2.1 standard

- Update of the principal requirement title to include system components and cardholder data.
- Removal of the requirement for specific documented procedures and added testing procedures to verify policies and procedures to each related requirement (7.2.1, 7.2.2, 7.2.3)
- Clarification that the requirement is about defining an access control model (7.2.1)
- Clarification that the requirement is about approval of required privileges by authorized personnel (7.2.3).

Implications for data storage

The default cluster administrator role provides full access to functions as the “admin” role. In addition, there are several predefined admin roles for the cluster context to limit capabilities for some administrators to read only or AutoSupport user. These default cluster administrator roles include high-level admin roles as well as more limited roles for read-only access.

The predefined roles for cluster administrator and SVM administrators are typically sufficient to provide adequate security. Because of the separation of the data and management planes, user data is usually not directly accessible by administrators. (The primary exception occurs when an administrator creates an additional user account to access user data.) Customer data access is restricted to authorized accounts either locally or, preferably, through an LDAP or Active Directory, that provides identity management to users. This enables you to restrict data users to specific volumes.

Implications for NAS access through SMB and NFS

To control file access for NFS and SMB, standard NAS protocol permissions are used, such as NTFS access control lists (ACLs) and UNIX mode bits. NAS access has four main access considerations.

- **Export policy rules.** Before user authentication can occur, export policy rules must be evaluated for NAS access. SMB shares can leverage export policy rules but these rules are disabled by default in ONTAP 9. NFS exports always base export policy rules and share-level access on a series of factors, including host name/client IP, allowed security type (such as SYS and KRB), and who the user attempting access is.
- **Authentication.** Users must prove that they are who they say they are. This authentication is done through name mapping based on the security style of a file system. If a user requesting access does not map to a valid user, access is denied. Other authentication pieces, such as Kerberos, also might come into play, according to system configuration.
- **Share permissions.** SMB shares use ACL-based share permissions to control whether an authenticated user can access a share. If the user who authenticates is not on the ACL, access to the share is denied. Share-level permissions are different from file-level permissions and are covered in Microsoft documentation.
- **Authorization.** After the user has been authenticated and allowed access to a share, what that user can do at a file or folder level must be verified through the ACLs on the data object. The ACL rights are based on who the user was authenticated as and the security style of the file system.

NAS file system local accounts (SMB workgroup)

Beginning with ONTAP 9, you can configure a SMB server in a workgroup with SMB clients that authenticate to the server by using locally defined users and groups. Workgroup client authentication provides an additional layer of security that is consistent with traditional domain authentication.

Note: An SMB server in workgroup mode supports only Windows NT LAN Manager (NTLM) authentication and does not support Kerberos authentication.

NetApp recommends using the NTLM authentication function with SMB workgroups to maintain your organization's security posture. To validate the SMB security posture, NetApp recommends using the `vserver cifs session show` command to display numerous posture-related details, including IP information, the authentication mechanism, the protocol version, and the authentication type.

You can use external name service servers, such as Active Directory, LDAP, or NIS to query users and groups for authentication and authorization. For more information about NAS access, see the following technical reports:

- TR-4067: NFS in NetApp ONTAP — Best practice and implementation guide
<https://www.netapp.com/us/media/tr-4067.pdf>
- TR-4073: Secure Unified Authentication
<https://www.netapp.com/us/media/tr-4073.pdf>
- TR-4543: SMB Protocol Best Practices
<https://www.netapp.com/us/media/tr-4543.pdf>
- Best Practices Guide for Clustered Data Windows File Services
<https://www.netapp.com/us/media/tr-4191.pdf>
- Security hardening guide for NetApp ONTAP 9
<https://www.netapp.com/us/media/tr-4569.pdf>

Best practices

Use default administration roles to manage authorization for each administrator account.

Assign a unique ID and password to each person with access (for logging purposes). Assign a role to each user account with minimal authority for assigned tasks (least-privilege principle).

Use LDAP, Active Directory, or Network Information Service (NIS) to provide authentication access to data users to specific volumes.

Requirement 8: Identify users and authenticate access to system components

The following are two fundamental principles of identifying and authenticating users:

- Establish the identity of an individual or process on a computer system.
- Prove or verify that the user associated with the identity is who the user claims to be.

Identification of an individual or process on a computer system is conducted by associating an identity with a person or process through an identifier, such as a user, system, or application ID. These IDs (also referred to as accounts) fundamentally establish the identity of an individual or process by assigning unique identification to each person or process to distinguish one user or process from another. When each user or process is uniquely identified, it ensures there is accountability for actions performed by that identity. When accountability is in place, actions taken can be traced to known and authorized users and processes.

The element used to prove or verify the identity is known as the authentication factor. Authentication factors include:

- Something you know, such as a password or passphrase.
- Something you have, such as a token device or smart card.
- Something you are, such as a biometric element.

The ID and the authentication factor together are considered authentication credentials and are used to gain access to the rights and privileges associated with a user, application, system, or service accounts.

These requirements for identity and authentication are based on industry-accepted security principles and best practices to support the payment ecosystem. The NIST Special Publication 800-63, Digital Identity Guidelines provides additional information about acceptable frameworks for digital identity and authentication factors. It is important to note that the NIST Digital Identity Guidelines is intended for US Federal Agencies and should be viewed in its entirety. Many of the concepts and approaches defined in these guidelines are expected to work with each other and not as standalone parameters.

Note: Unless otherwise stated in the requirement, these requirements apply to all accounts on all system components, , including but not limited to the following:

- Point-of-sale accounts
- Accounts with administrative capabilities
- System and application accounts
- All accounts used to view or access cardholder data or to access systems with cardholder data.

Note: This includes accounts used by employees, contractors, consultants, internal and external vendors, and other third parties (for example, for providing support or maintenance services).

Certain requirements are not intended to apply to user accounts that have access to only one card number at a time to facilitate a single transaction (such as IDs used by cashiers on point-of-sale terminals). When items do not apply, they are noted directly within the specific requirement.

Note: These requirements do not apply to accounts used by consumers (cardholders).

Evolving requirements from previous PCI DSS 3.2.1

- New requirement for roles and responsibilities. This requirement is effective immediately for all v4.0 assessments (8.1.2).
- Change of focus of the requirement to allow the use of shared authentication credentials, but only on an exception basis (8.2.2).
- Increase in the number of invalid authentication attempts before locking out a user ID from six to ten attempts (8.3.4).
- New requirement to increase password length from a minimum of seven characters to 12 characters (or if the system does not support 12 characters, a minimum length of eight characters). This requirement is a best practice until 31 March 2025. Clarification that until 31 March 2025, passwords must be a minimum length of at least seven characters in accordance with v3.2.1 Requirement 8.2.3. This requirement applies only if passwords or passphrases are used as an authentication factor to meet Requirement 8.3.1. Addition of a note that this requirement is not intended to apply to user accounts on point-of-sale terminals that have access to only one card number at a time to facilitate a single transaction (8.3.6).
- Addition of the option to determine access to resources automatically by dynamically analyzing the security posture of accounts, instead of changing passwords or passphrases at least once every 90 days (8.3.9).
- New requirement for service providers only—if passwords or passphrases are the only authentication factor for customer user access, then passwords or passphrases are either changed at least once every 90 days or access to resources is automatically determined by dynamically analyzing the security posture of the accounts. This requirement is a best practice until 31 March 2025. Addition of a note that this requirement does not apply to accounts of consumer users accessing their payment card information. Addition of a note that this requirement supersedes Requirement 8.3.10 after it becomes effective, and until that date, service providers can meet either Requirement 8.3.10 or 8.3.10.1 (8.3.10.1).
- New requirement to implement multi-factor authentication (MFA) for all access into the CDE. This requirement is a best practice until 31 March 2025. Added a note to clarify that MFA is required for both types of access specified in Requirements 8.4.2 and 8.4.3; and that applying MFA to one type of access does not replace the need to apply another instance of MFA to the other type of access (8.4.2).

- New requirement for secure implementation of MFA systems. This requirement is a best practice until 31 March 2025 (8.5.1).
- New requirement for management of system or application accounts that can be used for interactive login. This requirement is a best practice until 31 March 2025 (8.6.1).
- New requirement for not hard-coding passwords or passphrases into files or scripts for any application and system accounts that can be used for interactive login. This requirement is a best practice until 31 March 2025 (8.6.2).
- New requirement for protecting passwords or passphrases for application and system accounts against misuse. This requirement is a best practice until 31 March 2025 (8.6.3).

Additional guidance and clarification from previous PCI DSS 3.2.1

- Standardization of the terms “authentication factor” and “authentication credentials.” Removal of the term “non-consumer users” and clarification in the overview that requirements do not apply to accounts used by consumers (cardholders).
- Addition of a note that this requirement is not intended to apply to user accounts on point-of-sale terminals that have access to only one card number at a time to facilitate a single transaction (8.2.1).
- Addition of a note that this requirement is not intended to apply to user accounts on point-of-sale terminals that have access to only one card number at a time to facilitate a single transaction (8.2.2).
- Clarification that this requirement applies only if passwords or passphrases are used as an authentication factor to meet Requirement 8.3.1 (8.3.5).
- Clarification that this requirement applies only if passwords or passphrases are used as an authentication factor to meet Requirement 8.3.1. Addition of a note that this requirement is not intended to apply to user accounts on point-of-sale terminals that have access to only one card number at a time to facilitate a single transaction. Addition of a note that this requirement does not apply to service providers’ customer accounts but does apply to accounts for service provider personnel (8.3.9).

Implications for data storage

You must set password policies regarding expiration, number of special characters, and so on. You can configure the following parameters for each role:

- The required minimum length of a user name
- Whether a mix of alphabetic and numeric characters is required in a user name
- The required minimum length of a password
- Whether a mix of alphabetic and numeric characters is required in a password
- The required number of special characters in a password
- Whether users must change their passwords when logging in to their accounts for the first time
- The number of previous passwords (up to four) that cannot be reused
- The minimum number of days (maximum is 90) that must pass between password changes
- The number of days after which a password expires
- The number of invalid login attempts that trigger the account to be locked automatically
- The number of days for which an account is locked if invalid login attempts reach the allowed maximum

Multifactor administrative access

Beginning with ONTAP 9.3, NetApp is addressing this requirement for administrative web authentication in NetApp ONTAP System Manager and Active IQ Unified Manager, and for SSH administrative CLI authentication in ONTAP.

For more information, see [Multifactor Authentication in ONTAP 9.3](#).

Best practices

For SSH administrative access to ONTAP systems, use a locally administered administrator account with chained primary and secondary authentication methods of `password` and `publickey` or a NIS/LDAP account with chained authentication methods of `nsswitch` and `publickey`.

For the ONTAP System Manager web UI or the Active IQ Unified Manager web UI, use Security Assertion Markup Language (SAML) 2.0, where ONTAP OCSM or OCUM is the service provider role and either Microsoft Active Directory Federation Services (ADFS) or Shibboleth is the identity provider (IdP) role. The authentication factors are configured in the IdP.

Following are the default rules for passwords for ONTAP:

- A password cannot contain the user name.
- A password must contain at least eight characters.
- A password must contain at least one letter and one number.
- A password cannot be the same as the last six passwords.

To enhance user account security, use parameters of the `security login role config modify` command to modify the settings of an access-control role.

- Rule settings for user names:
 - The required minimum length of a user name (`-username-minlength`)
 - Whether a mix of alphabetic and numeric characters is required in a user name (`-username-alphanumeric`)
- Rule settings for passwords:
 - The required minimum length of a password (`-passwd-minlength`). (The required minimum length is 12 characters according to PCI-DSS 4.0.)
 - Whether a mix of alphabetic and numeric characters is required in a password (`-passwd-alphanumeric`)
 - The required number of special characters in a password (`-passwd-min-special-chars`)
 - Whether users must change their passwords when logging in to their accounts for the first time (`-require-initial-passwd-update`)
 - The number of previous passwords that cannot be reused (`-disallowed-reuse`)
 - The minimum number of days that must pass between password changes (`-change-delay`)
 - The number of days after which a password expires (`-passwd-expiry-time`)
- Rule settings about invalid login attempts:
 - The number of invalid login attempts that trigger the account to be locked automatically (`-max-failed-login-attempts`). When the number of a user's invalid login attempts reaches the value specified by this parameter (which has a maximum value of 10 per PCI-DSS 4.0), the user account is locked automatically. The `security login unlock` command unlocks a user account.
 - The number of days for which an account is locked if invalid login attempts reach the allowed maximum (`-lockout-duration`)

You can display the current settings for the rules by using the `security login role config show` command. For information about the `security login role config` commands and the default settings, see the man pages or [Manage failed login attempts](#) in the ONTAP 9 Documentation Center. For convenience, the login accounts used on the service provider can match the general accounts in ONTAP that are used for general administrative access.

In addition to the password protections described previously, service provider firmware 1.2 and later tracks failed SSH login attempts from an IP address. If more than five repeated login failures are detected from an IP address in any 10-minute period, the service provider stops all communication with that IP address for the next 15 minutes. Normal communication resumes after 15 minutes, but if repeated login failures are detected again, communication is again suspended for the next 15 minutes.

Requirement 9: Restrict physical access to cardholder data

Any physical access to cardholder data or systems that store, process, or transmit cardholder data provides the opportunity for individuals to access and remove systems or hardcopies containing cardholder data. You must restrict physical access appropriately to prevent individuals from accessing cardholder data

There are three different areas mentioned in Requirement 9:

- Requirements that specifically refer to sensitive areas apply to those areas only.
- Requirements that specifically refer to the CDE apply to the entire CDE, including any sensitive areas residing within the CDE.
- Requirements that specifically refer to the facility reference the types of controls that might be managed more broadly at the physical boundary of a business premise (such as a building) where CDEs and sensitive areas reside. These controls often exist outside a CDE or sensitive area, for example, a guard desk that identifies, badges, and logs visitors. The term “facility” is used to recognize that these controls might exist at different places within a facility, for instance, at building entry or at an internal entrance to a data center or office space.

Evolving requirements from previous PCI DSS 3.2.1

- New requirement for roles and responsibilities. This requirement is effective immediately for all v4.0 assessments (9.1.2).
- New requirement to define the frequency of periodic POI device inspections based on the entity’s targeted risk analysis. This requirement is a best practice until 31 March 2025 (9.5.1.2.1).

Additional guidance and clarification from previous PCI DSS 3.2.1

- In the overview clarification of the three different areas covered in Requirement 9 (sensitive areas, CDE, and facilities). Whether or not each requirement applies to the CDE, sensitive areas, or facilities is clarified throughout.
- Addition of a requirement to address a former testing procedure bullet to restrict access to consoles in sensitive areas by locking them when not in use (9.2.4).
- Removal of the requirement for procedures to physically secure media (9.5) and merged the procedures into the related requirements. Split of the requirement for storing media backups in a secure location and reviewing the security of the offline backup location at least every 12 months into two requirements (9.4.1, 9.4.1.1, 9.4.1.2).
- Removal of the requirement for procedures for internal and external distribution of media (9.6) and merged the procedures into the related requirements (9.4.2, 9.4.3, 9.4.4).
- Removal of the requirement for procedures for strict control over storage and accessibility of media (9.7) and merged the procedures into the related requirements. Split of requirement for maintaining media inventory logs and conducting media inventories annually into two requirements (9.4.5, 9.4.5.1).
- Removal of the requirement for procedures for media destruction when media is no longer needed (9.8) and merged the procedures into the related requirements. Clarification that options for destroying media when no longer needed includes either destruction of electronic media or rendering cardholder data unrecoverable (9.4.6, 9.4.7).
- Clarification that the focus of the requirement is on “Point-of- interaction (POI) devices that capture payment card data by using direct physical interaction with the payment card form factor.”

Clarification that the requirement applies to deployed POI devices used in card-present transactions (9.5.1).

Implications for data storage

ONTAP systems should be installed in locked rooms and preferably in locked racks. For additional protection to meet or exceed PCI DSS 4.0 requirements, NetApp recommends NSE, NVE, and NAE. NSE is NetApp's implementation of full disk encryption (FDE) using FIPS-140-2 level 2 validated SEDs from leading vendors. NVE and NAE are a NetApp software-based, data-at-rest encryption solution available for any drive type. Both NAE and NVE are FIPS 140-2 compliant. NVE enables ONTAP to encrypt data (using AES 256-bit encryption) per volume for granularity. NAE provides shared AES-256-bit encryption keys across all volumes in the aggregate. You can store NAE aggregates and NVE volume encryption keys on an external key manager. You can combine NSE, NVE, and NAE for two layers of encryption. If physical security is compromised and a disk is physically removed from the system, the encryption on the disk protects the data.

For more information on NSE, NAE, and NVE, see the [NetApp Encryption Power Guide](#).

Best practice

Control physical room access and use enclosed racks with locks to provide physical protection for ONTAP systems. Use NSE and NVE for additional protection.

Regularly monitor and test networks

Requirement 10: Log and monitor all access to system components and cardholder data

Logging mechanisms and the ability to track user activities are critical in preventing, detecting, or minimizing the impact of a data compromise. The presence of logs on all system components and in the CDE enables thorough tracking, alerting, and analysis when something does go wrong. Determining the cause of a compromise is difficult, if not impossible, without system activity logs.

This requirement applies to user activities, including those by employees, contractors, consultants, internal and external vendors, and other third parties (for example, those providing support or maintenance services).

These requirements do not apply to user activity of consumers (cardholders).

Evolving requirements from previous PCI DSS 3.2.1

- New requirement for roles and responsibilities. This requirement is effective immediately for all v4.0 assessments (10.1.2).
- New requirement for the use of automated mechanisms to perform audit log reviews. This requirement is a best practice until 31 March 2025 (10.4.1.1).
- New requirement for a targeted risk analysis to define the frequency of periodic log reviews for all other system components (not defined in Requirement 10.4.1) This requirement is a best practice until 31 March 2025 (10.4.2.1).
- New requirement for all entities to detect, alert, and promptly address failures of critical security control systems. This requirement is a best practice until 31 March 2025. This new requirement applies to all entities - it includes two additional critical security controls not included in Requirement 10.7.1 for service providers (10.7.2).
- New requirement to respond promptly to failures of any critical security controls. For service providers: this is a current PCI DSS v3.2.1 requirement. For all other (non-service provider) entities:

this is a new requirement. This requirement is a best practice (for non-service providers) until 31 March 2025 (10.7.3).

Additional guidance and clarification from previous PCI DSS 3.2.1

- Update to the principal requirement title to reflect focus on audit logs, system components, and cardholder data. Clarification that these requirements do not apply to the user activity of consumers (cardholders). Replacement of the term “Audit trails” with “Audit logs” throughout.

Implications for data storage

NetApp ONTAP systems provide extensive audit logging and monitoring controls. Logs are event-triggered messages that range in severity and are generated by the ONTAP system and recorded in flat text files on the cluster. Logs are the primary resource for administrators, NetApp Support and NetApp Active IQ (AutoSupport) to determine and isolate root causes for a wide range of issues. Likewise, logs also fulfill logging requirements for PCI DSS 4.0.

Several types of logs are provided by ONTAP. The primary types include the event management system (EMS), audit logs, and Active IQ logs.

EMS is the ONTAP messaging facility built on the syslog standard. EMS simplifies the management of cluster wide events and how the administrator chooses to be notified. EMS provides a catalogued logging mechanism, and every event has a formal definition. This mechanism enables EMS to provide services such as automatic spam management (for example, message suppression), configurable notifications, assistance with translating low-level data into understandable text, NVMEM backing of messages, and automatic tagging of messages.

Although EMS captures events, the audit log is used to capture actions. The audit log records the commands sent to the cluster, the user who is sending them, and the success or failure of the command. This information applies to the CLI, REST API calls (such as commands from NetApp manageability tools, or automation), and HTTPS requests.

Finally, the Active IQ logs capture a combination of EMS events, audit log entries, and system state information useful to NetApp Support personnel in diagnosing the health of the system.

Taken together, the three types of logs (EMS, audit, and Active IQ) provide a clear and permanent record of the system for security purposes.

By default, set requests are recorded in `command-history.log` and `mgwd.log`, but get requests are not. To view or modify this setting, perform the `security audit` CLI operations. Regardless of the settings for the security audit commands, set requests are always recorded in the `command-history.log` file.

To access the log files, the Service Processor Infrastructure (SPI) web service is used. The SPI web service is enabled by default, and it can be disabled manually (`vserver services web modify -vserver * -name spi -enabled false`). The SPI web service allows the log files to be downloaded but not altered in the system. The files can also be deleted by an administrative user with sufficiently high authorization.

For more information about logging, see [TR-4303: Logging in Clustered Data ONTAP](#). This subject is also covered in [How ONTAP implements audit logging](#) in the ONTAP 9 Documentation Center.

The admin role is granted access to the SPI web service by default, and you can disable access manually (`services web access delete -vserver cluster_name -name spi -role admin`).

1. Point the web browser to the SPI web service URL in one of the following formats:

```
https://cluster-mgmt-LIF/spi/
```

`cluster-mgmt-LIF` is the name or IP address of the cluster management LIF.

2. When prompted by the browser, enter your user account and password.

After your account is authenticated, the browser displays links to the `/mroot/etc/log/`, `/mroot/etc/crash/`, and `/mroot/etc/mib/` directories of each node in the cluster.

You must manage all logging information as part of a PCI DSS–compliant security policy. The EMS log provides a summary of events that affected the system and might be useful in detecting external attacks (such as DDoS attacks). Likewise, the audit logs are useful for detecting malicious commands entered in the system.

It also might be desirable to send logs to a central location like a remote syslog server so that logs from all systems can be reviewed in one place. ONTAP supports this with the cluster log forwarding capability. For more information see the “Sending out syslog” question in TR-4569: Security hardening guide for NetApp ONTAP 9.

Best practice

Use the audit logging capability in ONTAP to monitor administrative actions. Establish an organizational policy to review the logs on a regular basis. Export logs from ONTAP to an external syslog server for centralized logging and monitoring.

Requirement 11: Test security of systems and networks regularly

Vulnerabilities are discovered continually by malicious individuals and researchers and are introduced by new software. System components, processes, and bespoke and custom software should be tested frequently to ensure security controls continue to reflect a changing environment.

Evolving requirements from previous PCI DSS 3.2.1

- New requirement for roles and responsibilities. This requirement is effective immediately for all v4.0 assessments (11.1.2).
- New requirement to manage all other applicable vulnerabilities (those not ranked as high-risk or critical) found during internal vulnerability scans. This requirement is a best practice until 31 March 2025 (11.3.1.1).
- New requirement to perform internal vulnerability scans by using authenticated scanning. This requirement is a best practice until 31 March 2025 (11.3.1.2).
- New requirement for multitenant service providers to support their customers for external penetration testing. This requirement is a best practice until 31 March 2025 (11.4.7).
- New requirement for service providers to use intrusion-detection and or intrusion-prevention techniques to detect, alert on/prevent, and address covert malware communication channels. This requirement is a best practice until 31 March 2025 (11.5.1.1).
- New requirement to deploy a change -and-tamper- detection mechanism to alert for unauthorized modifications to the HTTP headers and contents of payment pages that are received by the consumer browser. This requirement is a best practice until 31 March 2025 (11.6.1).

Additional guidance and clarification from previous PCI DSS 3.2.1

- A minor update to the principal requirement title.
- Clarification that the intent of the requirement is to manage both authorized and unauthorized wireless access points. A This requirement applies even when a policy exists to prohibit the use of wireless technology (11.2.1).
- Clarification of the following points:(11.4.1):
 - The methodology is defined, documented, and implemented by the entity.
 - Penetration testing results are retained for at least 12 months.
 - The methodology includes a documented approach to assessing and addressing risk posed by exploitable vulnerabilities and security weaknesses found during penetration testing.

- The meaning of testing from inside the network (internal penetration testing) and from outside the network (external penetration testing).
- Clarification that penetration test findings are corrected in accordance with the entity's assessment of the risk posed by the security issue (11.4.4).

Implications for data storage

Routine security validation is an ongoing, dynamic process that should be part of a comprehensive security policy. Many organizations include security scans on a periodic basis (quarterly or more often) using well-known industry tools. These tools operate in different ways and produce different types of results, so they are often useful in combination.

Service providers have additional responsibilities to detect, alert or prevent, and address covert malware communication channels. NetApp ONTAP systems can also support protection from malware (such as ransomware) on NAS (NFS and SMB) connected shares. For example, FPolicy, in combination with NetApp Cloud Insights, or similar capabilities from our partners, does an excellent job of detecting malware through user behavioral analytics (UBA). It looks for potential malware attacks from the aspect of an individual user's behavior. Hijacking a single user account is just one avenue a hacker might take when launching a malware attack; malicious actors are constantly evolving their attack techniques.

Active IQ and Active IQ Unified Manager also provide additional layers of detection for ransomware. Active IQ checks ONTAP systems for adherence to NetApp configuration best practices, such as enabling FPolicy. Active IQ Unified Manager generates alerts for abnormal growth of Snapshot copies or storage efficiency loss, which can indicate potential ransomware attacks. The anti-ransomware on box feature in ONTAP 9.10.1 leverages built-in on-box ML that uses volume workload activity and data entropy to automatically detect ransomware. It monitors activity that is different from UBA so that it can detect attacks that UBA does not.

Vulnerability scans on ONTAP

To understand the results of security scanners, it is important to understand some aspects of how they operate. Very rarely do the scanners perform actual tests of devices for security vulnerabilities. Some security scanners base assumptions about a scanned device's capabilities on release version identifiers found on the device. Those identifiers and the software running on the device might identify a vulnerability that has been remediated, resulting in "false-positive" reports.

For instance, ONTAP and other NetApp products are modified over time as new features are introduced and as suspected security vulnerabilities are identified and remediated. Applicable licenses for open-source components of NetApp products often require that the original release version identifier be used in the code. Therefore, NetApp continually applies fixes to known vulnerabilities, but does not always trigger a version update to be detected by a vulnerability scanner.

Running the vulnerability scan as an authenticated user can help reduce these false positives. PCI DSS 4.0 recommends performing internal vulnerability scans by using authenticated scanning.

For more information, see the NetApp knowledge base article: [Vulnerability Scanner indicates ONTAP as an unsupported Unix version](#).

If you discover a suspected vulnerability (either through a port scanner or otherwise), refer to instructions on the [How to Report Security Issues to NetApp](#) support page. You can also subscribe to security advisories on this page.

Best practice

Establish a policy to run vulnerability scanners against all systems in the data center on a regular basis.

Maintain an information security policy

Requirement 12: Support information security with organizational policies and programs

The organization's overall information security policy sets the tone for the whole entity and informs personnel what is expected of them. All personnel must be aware of the sensitivity of cardholder data and their responsibilities for protecting it.

For the purposes of Requirement 12, "personnel" refers to full-time and part-time employees, temporary employees, contractors, and consultants with security responsibilities for protecting account data or that can impact the security of account data.

Evolving requirements from previous PCI DSS 3.2.1

- Removal of the requirement for a formal organization-wide risk assessment and replacement with specific targeted risk analyses (12.3.1 and 12.3.2).
- Addition of a formal acknowledgment by personnel of their responsibilities (12.1.3).
- Removal of a requirement and addition of new Requirement 3.4.2 for technical controls to prevent copy and relocation of PAN when using remote-access technologies (3.4.2).
- New requirement to perform a targeted risk analysis for any PCI DSS requirement that provides flexibility for how frequently it is performed. This requirement is a best practice until 31 March 2025 (12.3.1).
- New requirement for entities using a customized approach to perform a targeted risk analysis for each PCI DSS requirement that the entity meets with the customized approach. This requirement is effective immediately for all entities undergoing a v4.0 assessment and using a customized approach (12.3.2).
- New requirement to document and review cryptographic cipher suites and protocols in use at least once every 12 months. This requirement is a best practice until 31 March 2025 (12.3.3).
- New requirement to review hardware and software technologies in use at least once every 12 months. This requirement is a best practice until 31 March 2025 (12.3.4).
- New requirement to document and confirm PCI DSS scope at least every 12 months and when there is a significant change to the in-scope environment. This requirement is effective immediately for all v4.0 assessments (12.5.2).
- New requirement for service providers to document and confirm PCI DSS scope at least once every six months and when there is a significant change to the in-scope environment. This requirement is a best practice until 31 March 2025 (12.5.2.1).
- New requirement for service providers for a documented review of the impact to PCI DSS scope and applicability of controls when there are significant changes to organizational structure. This requirement is a best practice until 31 March 2025 (12.5.3).
- New requirement to review and update (as needed) the security awareness program at least once every 12 months. This requirement is a best practice until 31 March 2025 (12.6.2).
- New requirement for security awareness training to include awareness of threats and vulnerabilities that might impact the security of the CDE. This requirement is a best practice until 31 March 2025 (12.6.3.1).
- New requirement for security awareness training to include awareness about the acceptable use of end-user technologies in accordance with Requirement 12.2.1. This requirement is a best practice until 31 March 2025 (12.6.3.2).
- New requirement for service providers to support their customers' requests for information to meet Requirements 12.8.4 and 12.8.5. This requirement is effective immediately for all v4.0 assessments (12.9.2).

- New requirement to perform a targeted risk analysis to define the frequency of periodic training for incident response personnel. This requirement is a best practice until 31 March 2025 (12.10.4.1).
- Merged requirements and update of the security monitoring systems to be monitored and responded to as part of the incident response plan to include the following (12.10.5):
 - Detection of unauthorized wireless access points (former 11.1.2),
 - Change-detection mechanism for critical files (former 11.5.1),
 - New requirement bullet for use of a change- and tamper-detection mechanism for payment pages (relates to new Requirement 11.6.1).

Note: This bullet is a best practice until 31 March 2025.

- New requirement for incident response procedures to be in place and initiated on detection of stored PAN anywhere it is not expected. This requirement is a best practice until 31 March 2025 (12.10.7).

Additional guidance and clarification from previous PCI DSS 3.2.1

- Update of the principal requirement title to reflect that the focus is on organizational policies and programs that support information security.
- Clarification that responsibilities are formally assigned to a chief information security officer or other knowledgeable member of executive management. Merged requirements for formally assigning responsibility for information security (12.1.4).
- Clarification that the intent of the requirement is for acceptable use policies for end-user technologies. Merged and removed requirements to focus on explicit management approval, acceptable uses of technologies, and a list of hardware and software products approved by the company for employee use (12.2.1)
- Clarification that the intent is that all personnel are aware of the entity's information security policy and their role in protecting cardholder data (12.6.1)
- Replacement of the term "Service Provider" with "Third-Party Service Provider (TPSP)". Clarification that the use of a PCI DSS compliant TPSP does not make an entity PCI DSS compliant and does not remove the entity's responsibility for its own PCI DSS compliance (12.8.1 – 12.8.5).
- Replacement of the term "Service Provider" with "Third-Party Service Provider (TPSP)" (12.8.2).
- Replacement of the term "Service Provider" with "Third-Party Service Provider (TPSP)" (12.8.3).
- Replacement of the term "Service Provider" with "Third-Party Service Provider (TPSP)". Clarification that where an entity has an agreement with a TPSP for meeting PCI DSS requirements on behalf of the entity, the entity must work with the TPSP to make sure that the applicable PCI DSS requirements are met. If the TPSP does not meet those applicable PCI DSS requirements, then those requirements are also not in place for the entity (12.8.4).
- Replacement of the term "Service Provider" with "Third-Party Service Provider (TPSP)". Clarification that the information about PCI DSS requirements managed by the TPSP and the entity should include any that are shared between the TPSP and the entity (12.8.5).
- Replacement of the terms "system breach" and "compromise" with "suspected or confirmed security incident (12.10.1)."
- Replacement of the term "alerts" with "suspected or confirmed security incidents (12.10.3)."
- Replacement of the term "system breach" with "suspected or confirmed security incidents (12.10.4)."

Implications for data storage

You can configure ONTAP to conform to and support organizational information security policies. These include RBAC, network services (such as NTP), password policies, MFA for administrative access, data retention, backup policies, and MAV to ensure potentially harmful operations, such as deleting data, have the approval of multiple admins.

Best practice:

Configure ONTAP to conform to the security policy by using roles matched to the responsibility of each user. Use MFA to protect from stolen credentials and use MAV to ensure potentially harmful actions have the approval of multiple admins.

Where to find additional information

- Network management guide ONTAP 9
https://docs.netapp.com/us-en/ontap/pdfs/sidebar/Network_management.pdf
- Cluster administration guide ONTAP 9
https://docs.netapp.com/us-en/ontap/pdfs/sidebar/Cluster_administration.pdf
- SNMP Support in Data ONTAP (NetApp login required)
<https://fieldportal.netapp.com/content/250723>
- Security ONTAP 9
<https://docs.netapp.com/us-en/ontap/pdfs/sidebar/Security.pdf>
- FIPS 140-2 validated cryptographic module software
<https://csrc.nist.gov/projects/cryptographic-module-validation-program/certificate/4144>
- NetApp Product Security
<https://security.netapp.com/>
- NetApp Security Advisories
<https://security.netapp.com/advisory/>
- TR-4067: NFS in NetApp ONTAP best practice and implementation guide
<https://www.netapp.com/us/media/tr-4067.pdf>
- TR-4073: Secure Unified Authentication
<https://www.netapp.com/us/media/tr-4073.pdf>
- TR-4543: SMB Protocol Best Practices
<https://www.netapp.com/us/media/tr-4543.pdf>
- TR-4303: Logging in Clustered Data ONTAP
<https://www.netapp.com/us/media/tr-4303.pdf>
- How to Report Security Issues to NetApp
<https://security.netapp.com/contact/>
- TR-4569: Security hardening guide for NetApp ONTAP 9
<https://www.netapp.com/us/media/tr-4569.pdf>
- TR-4191: Best Practices Guide for Clustered Data ONTAP Windows File Services
<https://www.netapp.com/us/media/tr-4191.pdf>
- TR-4647: Multifactor Authentication in ONTAP 9.3
<https://www.netapp.com/us/media/tr-4647.pdf>
- Security and data encryption ONTAP 9
https://docs.netapp.com/us-en/ontap/pdfs/sidebar/Security_and_data_encryption.pdf
- Vulnerability Scanner indicates ONTAP as an unsupported Unix version
https://kb.netapp.com/Advice_and_Troubleshooting/Data_Storage_Software/ONTAP_OS/Vulnerability_Scanner_indicates_ONTAP_as_an_unsupported_Unix_version
- PCI-DSS Standards
https://www.pcisecuritystandards.org/document_library?category=pcidss&document=pci_dss#agreement

Contact us

Let us know how we can improve this technical report.

Contact us at doccomments@netapp.com.

Include TECHNICAL REPORT 4401 in the subject line.

Version history

Version	Date	Document version history
Version 2.0	September 2022	Matt Trudewind: Updated for PCI DSS 4.0 requirements.
Version 1.3	November 2018	Dan Tulledge: Clarification on NSE drives.
Version 1.2	October 2018	Dan Tulledge: Updated section 2.1 firewall policy -allow-list
Version 1.1	March 2018	Dan Tulledge: Updated PCI DSS version 3.2.
Version 1.0	May 2015	Initial release.

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright information

Copyright © 2015–2022 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

TR-4401-0922