



Technical Report

PCI DSS 3.2

ONTAP 9

Dan Tulledge, NetApp
November 2018 | TR-4401

Abstract

This technical report is targeted at qualified security assessors as well as storage administrators focused on validating a system against the PCI DSS 3.2 standard. This document provides guidance for meeting the requirements whose controls can be applied to the NetApp® ONTAP® 9 system.

TABLE OF CONTENTS

1	Introduction to PCI DSS	3
2	Build and Maintain a Secure Network	3
2.1	Requirement 1: Install and Maintain a Firewall Configuration to Protect Cardholder Data	3
2.2	Requirement 2: Do Not Use Vendor-Supplied Defaults for System Passwords and Other Security Parameters	5
3	Protect Cardholder Data	8
3.1	Requirement 3: Protect Stored Cardholder Data	8
3.2	Requirement 4: Encrypt Transmission of Cardholder Data Across Open Public Networks	10
4	Maintain a Vulnerability Management Program	11
4.1	Requirement 5: Protect All Systems Against Malware and Regularly Update Antivirus Software or Programs	11
4.2	Requirement 6: Develop and Maintain Secure Systems and Applications	11
5	Implement Strong Access-Control Measures	13
5.1	Requirement 7: Restrict Access to Cardholder Data by Business Need to Know	13
5.2	Requirement 8: Identify and Authenticate Access to System Components	14
5.3	Requirement 9: Restrict Physical Access to Cardholder Data	17
6	Regularly Monitor and Test Networks	18
6.1	Requirement 10: Track and Monitor All Access to Network Resources and Cardholder Data	18
6.2	Requirement 11: Regularly Test Security Systems and Processes	19
6.3	Requirement 12: Maintain a Policy That Addresses Information Security for All Personnel	20
7	Where to Find Additional Information	21
8	Contact Us	22
	Version History	22

1 Introduction to PCI DSS

This technical report provides guidance and information that auditors and system operators will find useful in 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 paper 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.

2 Build and Maintain a Secure Network

2.1 Requirement 1: Install and Maintain a Firewall Configuration to Protect Cardholder Data

Firewalls are devices that control computer traffic allowed between an entity's networks (internal) and untrusted networks (external), as well as traffic into and out of more sensitive areas within an entity's internal trusted networks. The cardholder data environment (CDE) is an example of a more sensitive area within an entity's trusted network.

A firewall examines all network traffic and blocks those transmissions that do not meet the specified security criteria.

All systems must be protected from unauthorized access from untrusted networks, whether entering the system via the internet as e-commerce, employee internet access through desktop browsers, employee email access, dedicated connections such as business-to-business connections, via wireless networks, or via other sources. Often, seemingly insignificant paths to and from untrusted networks can provide unprotected pathways into key systems. Firewalls are a key protection mechanism for any computer network.

Other system components may provide firewall functionality, as long as they meet the minimum requirements for firewalls as defined in Requirement 1. Where other system components are used within the cardholder data environment to provide firewall functionality, these devices must be included within the scope and assessment of Requirement 1.

ONTAP storage systems should be installed behind and protected by an external firewall to conform with the principles of PCI DSS. In addition, ONTAP provides basic firewall functions for controlling management access to services. It is on by default in each node, and therefore protects the entire cluster. The firewall built into ONTAP is not designed to replace a dedicated external firewall, but to provide an additional layer of internal protection by permitting or blocking protocols as needed.

In ONTAP, each storage virtual machine (SVM, formerly known as Vserver) management logical interface (LIF) should have a firewall policy attached to it. Firewall policy entries consist of a protocol type, a firewall action, and a subnet or specific IP address to which the action applies. The firewall policy can be applied 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 have the strictest possible settings; that is, permit only the protocols and specific subnets or IP addresses needed to manage the SVM. For more information, see the [ONTAP 9 Network Management Guide](#).

Evolving Requirements from Previous PCI DSS 3.0 Standard—N/A

Additional Guidance and Clarification from Previous PCI DSS 3.0 Standard

- Added guidance to clarify intent of requirement (1.2.1, 1.3).
- Updated to clarify intent of requirement rather than use of a particular type of technology (1.3.5).

- Increased flexibility by including or equivalent functionality as alternative to personal firewall software. Clarified requirement applies to all portable computing devices that connect to the internet when outside the network and that also access the CDE.

Implications for Data Storage

Best Practice:

Enable firewall services on ONTAP to augment external firewalls.

A secured firewall policy for management interfaces is defined 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 1) 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 Secure Shell (SSH) clients	Allow SSH access from list
telnet	127.0.0.1/32, ::1/128	Deny all

The list in Table 1 provides the most common interfaces and protocols. Others, such as the NetApp FPolicy® component and Key Management Interoperability Protocol (KMIP), might also need to be considered. For more information, refer to the [ONTAP 9 Network Management Guide](#).

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, refer to the [ONTAP 9 Network Management Guide](#).

In addition to the firewall policies just described, the following recommendations enhance security.

The ports used for intracluster traffic should be on a private isolated network in the cluster. Further, the IP subnet used to access the cluster should 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/administrative traffic. This separation can also be performed on a per-SVM basis if desired.

ONTAP uses network services such as Network Time Protocol (NTP), domain name system (DNS), and others. NetApp recommends that these services be provided either by internal network sources or by proxies on the locally trusted network rather than exposing the system to the internet. Again, these services should also be provided on the management IP subnet (VLAN), separate from the data traffic.

The Service Processor (SP) should be in the VLAN for the management plane. The SP enables you to remotely access and administer the storage system and diagnose error conditions. Because it provides an additional point of entry into the system, it is part of the attack surface that must be protected. In keeping with the 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 on the SP and its capabilities, refer to “Managing the IP Addresses That Can Access the SP” in the [System Administration Reference](#).

2.2 Requirement 2: Do Not Use Vendor-Supplied Defaults for System Passwords and Other Security Parameters

Malicious individuals (external and internal to an entity) often use vendor default passwords and other vendor default settings to compromise systems. These passwords and settings are well known by hacker communities and are easily determined through public information.

Cluster administrators administer the overall cluster and have the ability to 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 will be accessed (both control plane and data plane) through the LIF interfaces assigned to that SVM by the cluster administrator.

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

The following command is used 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 some low-level system access and is ordinarily never needed. It is disabled by default, and **should 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. It should be locked, and access to the respective commands should be restricted to accounts with appropriate roles. Create a duplicate account for additional protection and the admin account can also be deleted from the system. This can be accomplished by creating a custom role 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 NetApp 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). Refer to [TR-4220: SNMP Support in Data ONTAP](#) for more information on configuring and using SNMP.

Evolving Requirements from Previous PCI DSS 3.0 Standard

- Removed Secure Sockets Layer (SSL) as an example of a secure technology. Added note that SSL and early TLS are no longer considered to be strong cryptography and cannot be used as a security control after June 30, 2016 (2.2.3 and 2.3).

Additional Guidance and Clarification from Previous PCI DSS 3.0 Standard

- Removed note and testing procedures regarding removal of SSL and early TLS and moved to new Appendix A2 (2.2.3 and 2.3).
- Removed reference to “web-based management,” as requirement already specifies “all non-console administrative access,” which by definition includes any web-based access (2.3).

Implications for Data Storage

ONTAP can be configured to meet these guidelines.

Best Practice:

Configure roles and accounts for ONTAP by using the principle of least privilege.

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 2 for cluster administrator accounts.

Table 2) 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 this table 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.

Similar to the cluster administrator roles in Table 2, there are four default roles for SVM administrators, as shown here.

Table 3) Cluster 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 • 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 3.2. To modify or create new roles, refer to the [Administrator Authentication and RBAC Power 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 -vserver vs
```

The default accounts for cluster administrator and SVM administrator have been 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 keeping with the least-privilege security philosophy.

The ONTAP system also provides web services access to the system to users and applications (using the NetApp Manageability SDK programming interface). NetApp recommends configuring the system to use only HTTPS and disabling HTTP. By default, SSLv3 is enabled and SSLv2 is disabled. NetApp recommends SSLv3, but SSLv2 is provided for backward compatibility.

3 Protect Cardholder Data

3.1 Requirement 3: Protect Stored Cardholder Data

Protection methods such as encryption, truncation, masking, and hashing are critical components of cardholder data protection. If an intruder circumvents other security controls and gains access to encrypted data but without the proper cryptographic keys, the data cannot be read or used by that person. Other effective methods of protecting stored data should also be considered as potential risk mitigation opportunities. For example, methods for minimizing risk include not storing cardholder data unless absolutely necessary, truncating cardholder data if full PAN is not needed, and not sending unprotected PANs using end-user messaging technologies such as email and instant messaging.

Please refer to the PCI DSS and PA-DSS Glossary of Terms, Abbreviations, and Acronyms for definitions of “strong cryptography” and other PCI DSS terms.

Evolving Requirements from Previous PCI DSS 3.0 Standard

- Updated requirement to clarify that any displays of PAN greater than the first six/last four digits of the PAN requires a legitimate business need. Added guidance on common masking scenarios. (3.3)
- New requirement for service providers to maintain a documented description of the cryptographic architecture. *Effective February 1, 2018* (3.5.1)

Additional Guidance and Clarification from Previous PCI DSS 3.0 Standard

- Updated requirement to clarify that any displays of PAN greater than the first six/last four digits of the PAN requires a legitimate business need. Added guidance on common masking scenarios. (3.2.1–3.2.3)
- Clarified in requirement note that additional controls are required if hashed and truncated versions of the same PAN are present in an environment. Added Testing Procedure 3.4.e to assist with validation of the note. Clarified intent of “truncation” in the Guidance column. (3.4)
- Updated testing procedure to clarify the examination of audit logs includes payment application logs. (3.4d)
- Added note to requirement to clarify the requirement applies in addition to all other PCI DSS encryption and key management requirements. (3.4.1)
- Clarified that “HSM” may refer to a “Hardware” or “Host” Security Module. Aligned with language in PCI PTS. (3.5.2)
- Clarified that Testing Procedure 3.6.a only applies if the entity being assessed is a service provider. (3.6)
- Updated testing procedure language to clarify testing involves observation of procedures rather than key-generation method itself, as this should not be observable. Added guidance referring to Glossary definition for “Cryptographic Key Generation.” (3.6.1b)

Implications for Data Storage

PCI DSS 3.2 states that encryption can be performed in one of two ways: as database encryption (either file-level or column-level) or as disk-level encryption. Disk encryption automatically meets the PCI DSS 3.2 requirement to separate security key management from user database management.

NetApp Storage Encryption (NSE) is NetApp’s implementation of FDE, using FIPS-140-2 level 2 validated self-encrypting disks (SEDs) from leading vendors. NetApp Volume Encryption (NVE) encrypts at the volume level, allowing the encryption capability to exist independently of the physical media: solid-state drives (SSDs), NetApp AFF, or even NSE drives.

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:

- Supports the entire suite of storage efficiency technologies from NetApp, including deduplication, compression, and array-based antivirus scanning
- Supports the Gemalto SafeNet KeySecure encryption-key appliance, the Gemalto SafeNet Virtual KeySecure appliance, the IBM SKLM/TLKM key management servers, and the Vormetric Data Security Manager (DSM) to strengthen and simplify long-term key management
- 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

NVE is a software-based, data-at-rest encryption solution available starting with NetApp ONTAP 9.1 management software. Starting with ONTAP 9.2, NVE is FIPS 140-2 compliant. NVE allows ONTAP to encrypt data (using AES 256-bit encryption) per volume for granularity. Starting with NetApp ONTAP 9.3, NVE volume encryption keys can be stored on an external key manager. If the controller and disks are moved without access to the external key manager, the NVE volumes won't be accessible and cannot be decrypted.

NVE:

- 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 the Gemalto SafeNet KeySecure encryption-key appliance, the Gemalto SafeNet Virtual KeySecure appliance, the IBM SKLM/TLKM key management servers, and the Vormetric Data Security Manager (DSM) to strengthen and simplify long-term key management
- 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

Best Practice:

To store payment data, use NetApp Storage Encryption, NetApp Volume Encryption, or both.

3.2 Requirement 4: Encrypt Transmission of Cardholder Data Across Open Public Networks

Sensitive information must be encrypted during transmission over networks that can be easily accessed by malicious individuals. Misconfigured wireless networks and vulnerabilities in legacy encryption and authentication protocols continue to be targets of malicious individuals who exploit these vulnerabilities to gain privileged access to cardholder data environments.

Evolving Requirements from Previous PCI DSS 3.0 Standard

- Removed SSL as an example of a secure technology and added a note to the requirement. See explanation above at 2.2.3. (4.1)
- Updated testing procedure to recognize all versions of SSL as examples of weak encryption. (4.1.1)

Additional Guidance and Clarification from Previous PCI DSS 3.0 Standard

- Removed note and testing procedures regarding removal of SSL/early TLS and moved to new Appendix A2. (4.2)
- Included SMS as an example of end-user messaging technology and added guidance. (4.2)

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 CIFS SMB signing and sealing.

4 Maintain a Vulnerability Management Program

4.1 Requirement 5: Protect All Systems Against Malware and Regularly Update Antivirus Software or Programs

Malicious software, commonly referred to as “malware”—including viruses, worms, and Trojans—enters the network during many business-approved activities, including sending employee email and using the internet, mobile computers, and storage devices, resulting in the exploitation of system vulnerabilities. Antivirus software must be used on all systems commonly affected by malware to protect systems from current and evolving malicious software threats. Additional antimalware solutions can be considered as a supplement to the antivirus software; however, such additional solutions do not replace the need to have antivirus software in place.

Evolving Requirements from Previous PCI DSS 3.0 Standard—N/A

Additional Guidance and Clarification from Previous PCI DSS 3.0 Standard—N/A

Implications for Data Storage

NetApp recommends that antivirus software be run on the computers used to provide card payment services.

For additional security protection, NetApp ONTAP systems can also support integrated antivirus functionality. In collaboration with leading partners, NetApp antivirus scanning integration can detect and prevent the spread of malicious virus code on storage networks.

Antivirus functionality integrated with ONTAP supports off-box integration points with antivirus servers to help guard your business-critical data against known and unknown threats. Multiple servers can be configured to support one or more NetApp storage devices for redundancy and performance.

Best Practice:

Install off-box antivirus servers to work with the ONTAP system to detect and mitigate malware.

4.2 Requirement 6: Develop and Maintain Secure Systems and Applications

Unscrupulous individuals 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 systems must have all appropriate software patches to protect against the exploitation and compromising of cardholder data by malicious individuals and malicious software.

Note: Appropriate software patches are patches that have been evaluated and tested sufficiently to determine that they do not conflict with existing security configurations. You can avoid vulnerabilities in in-house-developed applications by using standard system development processes and secure coding techniques.

Evolving Requirements from Previous PCI DSS 3.0 Standard

- New requirement for change control processes to include verification of PCI DSS requirements impacted by a change. *Effective February 1, 2018* (6.4.6)

Additional Guidance and Clarification from Previous PCI DSS 3.0 Standard

- Added clarification to Guidance column that requirement to patch all software includes payment applications. (6.2)
- Updated requirement to align with testing procedure. (6.4.4)
- Clarified that change control processes are not limited to patches and software modifications. (6.4.5)
- Clarified that training for developers must be up to date and occur at least annually. (6.5)
- Removed Testing Procedure 6.5.b and renumbered remaining testing procedures to accommodate. (6.5a-6.5c)
- Added clarification to testing procedure and Guidance column that if an automated technical solution is configured to alert (rather than block) web-based attacks, there must also be a process to ensure timely response. (6.6)

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's 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 will investigate a suspected vulnerability in our products and confirm 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 will 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 will disclose 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 will credit 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 prior to 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's 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.

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.

- **Security bulletin.** Low- and medium-severity security issues that affect NetApp products.
- **Security notices.** May 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).

See the [NetApp Product Security](#) page for more information.

Best Practice:

Subscribe to NetApp notifications and implement security patches and updates as they are made available. To subscribe, go to [NetApp Security Advisories](#).

5 Implement Strong Access-Control Measures

5.1 Requirement 7: Restrict Access to Cardholder Data by Business Need to Know

To ensure that critical data can be accessed only by authorized personnel, systems and processes must be in place to limit access based on the need to know and according to job responsibilities.

“Need to know” is when access rights are granted to only the least amount of data and the least number of privileges needed to perform a job.

Evolving Requirements from Previous PCI DSS 3.0 Standard—N/A

Additional Guidance and Clarification from Previous PCI DSS 3.0 Standard

- Updated requirement, testing procedures and Guidance column to clarify that one or more access control systems may be used.

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 readonly or AutoSupport user. They 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, AD, or NIS service that provides identity management to users. This allows the restriction of data users to specific volumes.

Implications for NAS Access Through CIFS and NFS

To control file access for NFS and CIFS, 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 take place, export policy rules must be evaluated for NAS access. CIFS shares can leverage export policy rules, but these rules are disabled

by default starting in clustered Data ONTAP 8.2. 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.** CIFS shares use ACL-based share permissions to control whether an authenticated user may 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 (CIFS Workgroup)

Beginning with ONTAP 9, you can configure a CIFS server in a workgroup with CIFS 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: A CIFS 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 CIFS workgroups to maintain your organization's security posture. To validate the CIFS 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.

- External name service servers, such as Active Directory, LDAP, or NIS, can be used to query users and groups for authentication and authorization. For more information on NAS access, see the following technical reports:
- [NFS Best Practice and Implementation Guide](#)
- [Secure Unified Authentication](#)
- [SMB Protocol Best Practices](#)
- [Best Practices Guide for Clustered Data Windows File Services](#)
- [Security Hardening Guide for NetApp ONTAP 9](#)

Best Practice:

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, AD, or NIS to provide authentication access to data users to specific volumes.

5.2 Requirement 8: Identify and Authenticate Access to System Components

Assigning a unique identification (ID) to each person with access ensures that each individual is uniquely accountable for their actions. When such accountability is in place, actions taken on critical data and systems are performed by, and can be traced to, known and authorized users and processes.

The effectiveness of a password is largely determined by the design and implementation of the authentication system--particularly, how frequently password attempts can be made by an attacker, and the security methods to protect user passwords at the point of entry, during transmission, and while in storage.

Note: Note: These requirements are applicable for all accounts, including point-of-sale accounts, with administrative capabilities and all accounts used to view or access cardholder data or to access systems with cardholder data. This includes accounts used by vendors and other third parties (for example, for support or maintenance). These requirements do not apply to accounts used by consumers (e.g., cardholders).

However, Requirements 8.1.1, 8.2, 8.5, 8.2.3 through 8.2.5, and 8.1.6 through 8.1.8 are not intended to apply to user accounts within a point-of-sale payment application that only have access to one card number at a time in order to facilitate a single transaction (such as cashier accounts).

Evolving Requirements from Previous PCI DSS 3.0

- Expanded Requirement 8.3 into subrequirements, to require multifactor authentication for all personnel with nonconsole administrative access, and all personnel with remote access to the CDE.
 - New Requirement 8.3.2 addresses multifactor authentication for all personnel with remote access to the CDE (incorporates former Requirement 8.3).
 - New Requirement 8.3.1 addresses multifactor authentication for all personnel with nonconsole administrative access to the CDE.
 - Requirement 8.3.1 effective February 1, 2018. (8.3, 8.3.1, 8.3.2)

Additional Guidance and Clarification from Previous PCI DSS 3.0

- Clarified that inactive user accounts must be removed/disabled within 90 days. (8.1.4)
- Clarified that testing procedure only applies if the entity being assessed is a service provider, and for nonconsumer customer accounts. (8.1.6.b, 8.2.1.d, 8.2.1.e, 8.2.3.b, 8.2.4.b, 8.2.5.b)
- Clarified that passwords must be changed at least once every 90 days. (8.2.4)
- Clarified this requirement only applies if the entity being assessed is a service provider. (8.5.1)
- Added note to Requirement 8 introduction that the authentication requirements do not apply to accounts used by consumers (e.g. cardholders). (Requirement 8)
- Clarified requirement intended for all third parties with remote access, rather than only vendors. (8.1.4)
- Updated Guidance column to reflect changing industry standards. (8.2.3)
- Clarified correct term is multi-factor authentication rather than two-factor authentication, as two or more factors may be used. (8.3)

Implications for Data Storage

Password policies should be set regarding expiration, number of special characters, and so on. The following parameters can be configured 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 triggers 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 OnCommand® System Manager and OnCommand Unified Manager, and for SSH administrative CLI authentication in ONTAP.

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

Best Practice:

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 OnCommand System Manager ONTAP web user interface or OnCommand Unified Manager web user interface, use Security Assertion Markup Language (SAML) 2.0, where ONTAP OCSM or OCUM is the service provider (SP) 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 be at least eight characters long.
- 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`)
 - 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 triggers 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 6 per PCI-DSS 3.2), the user's 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 the [Administrator Authentication and RBAC Power Guide](#).

For convenience, the login accounts used on the SP can match the general accounts on ONTAP that are used for general administrative access.

In addition to the password protections just described, SP firmware 1.2 and later track failed SSH login attempts from an IP address. If more than 5 repeated login failures are detected from an IP address in any 10-minute period, the SP 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.

5.3 Requirement 9: Restrict Physical Access to Cardholder Data

Any physical access to data or systems that house cardholder data provides the opportunity for individuals to access devices or data and to remove systems or hardcopies, and should be appropriately restricted. For the purposes of Requirement 9, “onsite personnel” refers to full-time and part-time employees, temporary employees, contractors and consultants who are physically present on the entity’s premises. A “visitor” refers to a vendor, guest of any onsite personnel, service workers, or anyone who needs to enter the facility for a short duration, usually not more than one day. “Media” refers to all paper and electronic media containing cardholder data.

Evolving Requirements from Previous PCI DSS 3.0—N/A

Additional Guidance and Clarification from Previous PCI DSS 3.0

- Clarified that the requirement applies to all onsite personnel and visitors. Combined Testing Procedures 9.2.b and 9.2.d to remove redundancy. (9.2)
- Updated testing procedure to clarify both devices and device locations need to be observed. (9.9.1.b)
- Clarified that either video cameras or access controls mechanisms, or both, may be used. (9.1.1)
- Combined testing procedures to clarify that assessor verifies the storage location is reviewed at least annually. (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 3.2 requirements, NetApp recommends NetApp Storage Encryption (NSE) and/or NetApp Volume Encryption (NVE). NSE is NetApp’s implementation of full disk encryption (FDE) using FIPS-140-2 level 2 validated SEDs from leading vendors. NVE is NetApp’s implementation of software encryption of ONTAP volumes for any drive type. NSE and NVE can be combined 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 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 NetApp ONTAP systems. Use NetApp Storage Encryption and NetApp Volume Encryption for additional protection.

6 Regularly Monitor and Test Networks

6.1 Requirement 10: Track and Monitor All Access to Network Resources 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 in all environments allows thorough tracking, alerting, and analysis when something does go wrong. Determining the cause of a compromise is very difficult, if not impossible, without system activity logs.

Evolving Requirements from Previous PCI DSS 3.0

- New requirement for service providers to detect and report on failures of critical security control systems. *Effective February 1, 2018* (10.8, 10.8.1)

Additional Guidance and Clarification from Previous PCI DSS 3.0

- Removed redundant language in guidance column. (10.6)
- Updated requirement to more clearly differentiate intent from Requirement 10.6.2. (10.6.1)
- Renumbered due to addition of new Requirement 10.8.

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 Active IQ (AutoSupport) to determine and isolate root causes for a wide range of issues. Likewise, they also fulfill logging requirements for PCI DSS 3.2.

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

EMS is the ONTAP messaging facility built on the syslog standard. EMS simplifies the management of clusterwide events and how the administrator chooses to be notified. EMS provides a catalogued logging mechanism, and every event has a formal definition. This mechanism allows EMS to provide services such as automatic spam management (for example, message suppression), configurable notifications, assistance with translating low-level data into understandable text, NVRAM 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, ONTAP API (ONTAPI library) calls (such as commands from NetApp manageability tools), and HTTP requests.

Finally, the AutoSupport 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 AutoSupport) 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.

Refer to [Logging in Clustered Data ONTAP](#) for more information on logging. The subject is also covered in the [ONTAP System Administration Reference](#).

The "admin" role is granted access to the SPI web service by default, and the access can be disabled 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.

All logging information should be managed 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 into the system.

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.

6.2 Requirement 11: Regularly Test Security Systems and Processes

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

Evolving Requirements from Previous PCI DSS 3.0

- New requirement for service providers to perform penetration testing on segmentation controls at least every six months. *Effective February 1, 2018* (11.3.4.1)

Additional Guidance and Clarification from Previous PCI DSS 3.0

- Clarified that testing procedure applies where wireless scanning is utilized. (11.1.c)
- Clarified in Guidance Column that a vulnerability scan could be a combination of automated and manual tools, techniques, or other methods. (11.2)
- Removed redundant language from testing procedure. (11.3.2.a)
- Clarified that the intent of the penetration testing is to verify that all out-of-scope systems are segmented (isolated) from systems "in the CDE". (11.3.4)

- Clarified that unauthorized modifications include changes, additions, and deletions of critical system files, etc. (11.5)
- Clarified that all “high risk” vulnerabilities must be addressed in accordance with the entity’s vulnerability ranking (as defined in Requirement 6.1), and verified by rescans. (11.2.1)
- Added Testing Procedure 11.3.4.c to confirm penetration test is performed by a qualified internal resource or qualified external third party. (11.3.4)
- Removed “within the cardholder data environment” from testing procedure for consistency with requirement, as requirement may apply to critical systems located outside the designated CDE. (11.5.a)

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.

Vulnerability scans on ONTAP

In order 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.

For more information, see the [Network Vulnerability Scanner Indicates ONTAP as a FreeBSD System](#) knowledge base article.

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.

6.3 Requirement 12: Maintain a Policy That Addresses Information Security for All Personnel

A strong security policy sets the security tone for the whole entity and informs personnel what is expected of them. All personnel should be aware of the sensitivity of 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 who are “resident” on the entity’s site or otherwise have access to the cardholder data environment.

Evolving Requirements from Previous PCI DSS 3.0

- New requirement for service providers' executive management to establish responsibilities for the protection of cardholder data and a PCI DSS compliance program. *Effective February 1, 2018* (12.4)
- New requirement for service providers to perform reviews at least quarterly, to confirm personnel are following security policies and operational procedures. *Effective February 1, 2018* (12.11, 12.11.1)

Additional Guidance and Clarification from Previous PCI DSS 3.0

- Clarified that the risk assessment process must result in a formal, "documented analysis of risk". (12.2)
- Clarified this requirement only applies if the entity being assessed is a service provider and added related guidance. (12.9)
- Reformatted testing procedure for clarity. (12.3.3)
- Renumbered due to addition of new Requirement 12.4. (12.4.1)
- Clarified intent of security awareness program is to ensure personnel are aware of the cardholder data security policy and procedures. (12.6)
- Clarified that the list of service providers includes a description of the service provided. (12.8.1)
- Added guidance that service provider responsibility will depend on the particular service being provided and the agreement between the two parties. (12.8.2)
- Clarified that review of the incident response plan encompasses all elements listed in Requirement 12.10.1. (12.10.2)

Implications for Data Storage

ONTAP can be configured to conform to and support organizational security policies. These include RBAC, network services (such as NTP), password policies, multifactor authentication for administrative access, data retention, and backup policies.

Best Practice:

Configure ONTAP to conform to the security policy by using roles matched to the responsibility of each user.

7 Where to Find Additional Information

- [ONTAP 9 Network Management Guide](#)
- [System Administration Reference](#)
- [SNMP Support in Data ONTAP](#)
- [Administrator Authentication and RBAC Power Guide](#)
- [FIPS 140-2 validated cryptographic module software](#)
- [NetApp Product Security](#)
- [NetApp Security Advisories](#)
- [NFS Best Practices and Implementation Guide](#)
- [Secure Unified Authentication](#)
- [SMB Protocol Best Practices](#)

- [Logging in Clustered Data ONTAP](#)
- [How to Report Security Issues to NetApp](#)
- [Security Hardening Guide for NetApp ONTAP 9](#)
- [Best Practices Guide for Clustered Data Windows File Services](#)
- [Multifactor Authentication in ONTAP 9.3](#)
- [NetApp Encryption Power Guide](#)
- [Network Vulnerability Scanner Indicates ONTAP as a FreeBSD System](#)
- [PCI-DSS Standards](#)

8 Contact Us

Let us know how we can improve this technical report.

Contact us at docfeedback@netapp.com.

Include TECHNICAL REPORT 4401 in the subject line.

Version History

Version	Date	Document Version History
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: Updates included to address PCI DSS ver 3.2
Version 1.0	May 2015	First 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 © 2015–2018 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

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.