



Technical Report

Clustered Data ONTAP Security Guidance Recommendations for Security

Dave Buster, NetApp
April 2015 | TR-4393

Abstract

Clustered Data ONTAP 8.3.0

This document is a brief summary of the security features and settings in clustered Data ONTAP® 8.3.0.

TABLE OF CONTENTS

1	Introduction	4
2	Differences Between 7-Mode and Clustered Data ONTAP Systems	4
2.1	Clustered Data ONTAP Virtualization	4
2.2	Practical Security Implications of Clustered Data ONTAP Architecture	5
2.3	Network Switch Isolation	5
3	Designing a Secure Storage Installation	6
3.1	Security Concepts	6
3.2	Policies and Procedures	6
3.3	Strategy for Clustered Data ONTAP Security Recommendations	7
4	Initial Security Configuration	7
4.1	Role-Based Access Control	7
4.2	Creating Admin User Accounts	10
4.3	Microsoft Active Directory Authentication for Administrator Users	11
4.4	Set Up a local Emergency Administrator Account	12
4.5	Restrict Default Administrative Accounts	12
4.6	Disable Unnecessary Services	12
4.7	Timeouts	13
4.8	Logging	13
4.9	External Ports and VLANs	14
4.10	NTP Protocol	15
4.11	Managing Public Keys	15
4.12	Managing Digital Certificates for Server or Client Authentication	15
4.13	Providing Mutual Authentication	16
5	Management	16
5.1	SSHv2	16
5.2	Telnet	17
5.3	NetApp Manageability SDK and External Software Tools	17
6	Security Validation	18
6.1	Vulnerability Scanners	18
6.2	Port Scanners	18
7	System Monitoring	18
7.1	SNMP	18
7.2	AutoSupport	19
7.2	External Syslog Server	20
8	Hardware Security Considerations	20
8.1	Service Processor	20
9	System Services	21
9.1	Firewall Service	21
9.2	NDMP Service	24
9.3	Web Services	24
10	Summary and Conclusion	25
10.1	Appendix	25
	Checklist of Security Recommendations	25

LIST OF TABLES

Table 1)	Cluster administrator (Cserver) roles	8
Table 2)	SVM administrator (Vserver) roles	9
Table 3)	Installed encryption ciphers and key exchange algorithms	17
Table 4)	Firewall policy protocols	22
Table 5)	Secure management firewall policy entry settings	23

Table 6) Checklist of Security Recommendations25

LIST OF FIGURES

Figure 1) Storage virtual machines.....5
Figure 2) Evolution of security domains.....6

1 Introduction

This document provides a set of practical recommendations to enhance the security of a clustered Data ONTAP system. Note that the protection of user data itself is primarily the responsibility of the appropriate SAN and NAS protocols and configurations. These are secure by default and need little configuration and provide data confidentiality. Therefore, this document focuses on securing the system itself by focusing on administrative interfaces and services. It is intended to reinforce the integrity and availability of a clustered Data ONTAP system in a typical data center environment.

The reader should already be familiar with basic setup and configuration of clustered Data ONTAP. If not, the following documents provide useful details for understanding the information in this guide:

- Clustered Data ONTAP 8.3 System Administration Guide for Cluster Administrators
- Clustered Data ONTAP 8.3 System Administration Guide for Vserver Administrators
- TR-3982: NetApp Clustered Data ONTAP 8.3 and 8.2.x
- TR-4182: Ethernet Storage Best Practices for Clustered Data ONTAP Configurations
- TR-4073: Secure Unified Authentication for NFS

2 Differences Between 7-Mode and Clustered Data ONTAP Systems

2.1 Clustered Data ONTAP Virtualization

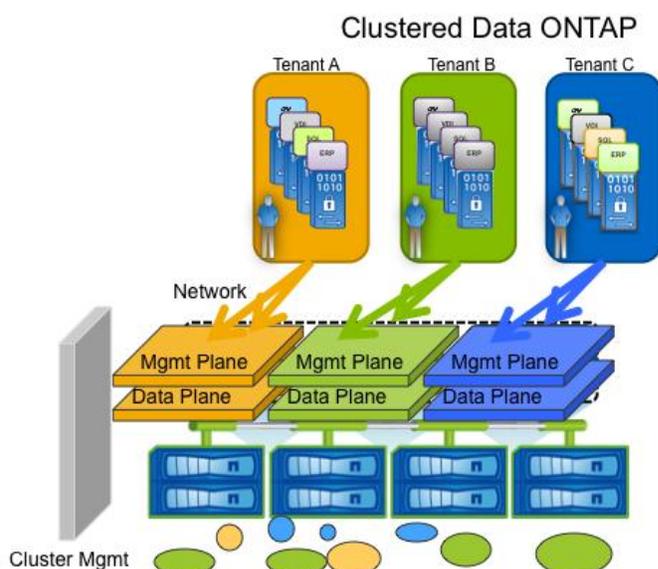
Clustered Data ONTAP uses storage virtual machines (SVMs) to partition the system into separate storage containers. In addition to other duties, a high-level cluster administrator has the ability to create SVMs and then must configure a list of available aggregates and logical interfaces for them. Then, within each SVM, the SVM administrator can operate on the logical aggregate objects.

At the cluster level, a cluster administrator has full privileges to configure clusterwide settings and assign aggregates and logical interfaces. Cluster-level administrators do not have direct storage protocol access (CIFS/NFS). Management of storage protocols only exists within an SVM. If necessary, the cluster administrator can also switch context down to the SVM administrator level to configure settings there as well. However, each SVM administrator can see only those objects assigned to that SVM by the cluster administrator.

Further, within each SVM, the data plane is separate from the management plane. This means that generally, SVM administrators do not have access to datastores, and data users do not have administration rights unless a login account is specifically configured to provide both.

This compartmentalization provides powerful tools for security of clustered Data ONTAP. SVM administrative and authentication domains are kept completely separate from each other. They can therefore be administered by different organizations without risk of interference or contention. More importantly, each SVM exists within its own security domain, with separate authentication, role-based access control, and firewall policy functions.

Figure 1) Storage virtual machines.



2.2 Practical Security Implications of Clusted Data ONTAP Architecture

As stated earlier, clustered Data ONTAP is organized into separate administrative domains by virtue of the SVM-based architecture. From a security standpoint, this creates a desirable compartmentalization of responsibility. However, it is different from 7-Mode systems and must be understood to be managed.

As a result of this architecture, some basic security concepts can be highlighted:

- Unlike the 7-Mode architecture, where vFiler® instances were optional, a minimum of one SVM is required per cluster to be able to administer and use storage resources. This is typically one of the first steps in setting up a new installation. From a security standpoint, both cluster administrator and SVM administrator are always authenticated.
- Because each SVM exists as a separate entity, it carries its own namespace and security domain. Therefore, authentication using LDAP and/or Active Directory® is done on a per-SVM basis, and role-based access controls are managed per SVM. Consequently, the security configuration and management are specific to an SVM and might vary from one SVM to another.
- Within clustered Data ONTAP, management authority can be delegated across various administrator roles. The cluster administrator has authority to create (and delete) SVMs and assign resources such as aggregates, volumes, logical interfaces (LIFs), and so on to each of them. Then, within each SVM, the SVM administrator can manage the necessary hierarchy of volumes, namespaces, directories, and files. This separation of authority can be used to augment security across the system by separating roles on a “need to know” and least privilege management philosophy.

2.3 Network Switch Isolation

The clustered Data ONTAP solution uses an Ethernet switching system as a backplane between nodes. Exposing this switching fabric to other traffic is not supported. Outside (alien) traffic would have an unpredictable effect on performance and would expose the system to unnecessary security risk. Therefore, this infrastructure must be physically isolated to a separate and unique network domain. Physical isolation is often used in military networks to enforce security, and the same principles apply

here. If the switch (rather than just port assignment) is physically isolated, there is no possible way for an incorrectly configured or wired interface to expose the system to security risk.

3 Designing a Secure Storage Installation

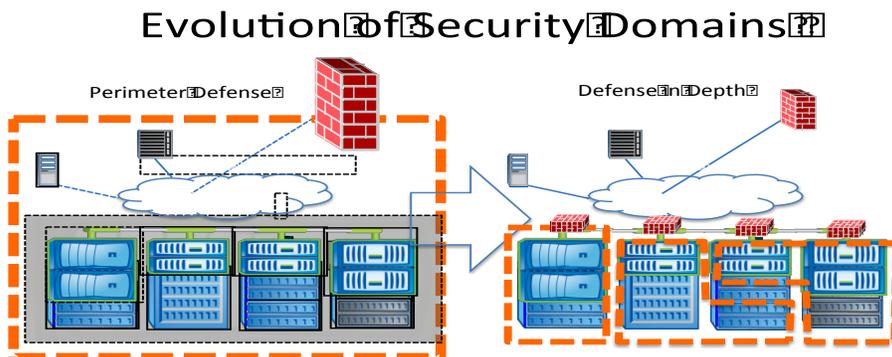
3.1 Security Concepts

Data center security has evolved from a perimeter defense based on firewalls to a defense-in-depth model based on integrated security across systems. Likewise, NetApp® products have been enhanced to provide data confidentiality, integrity, and availability as part of an integrated security domain.

Historically, storage systems have been protected by being in the heart of the data center, typically protected by external firewalls. A storage system existed in a fully trusted network domain and enclosed physical facility. Therefore, the storage system itself was left open for convenience.

Now, however, most organizations adopt a defense-in-depth strategy that relies on layers of defenses to protect critical systems. This reflects a more sophisticated view of security and requires all components in the data center to support an overall security policy and architecture. Therefore, the security functions in clustered Data ONTAP have been enhanced with role-based access controls, firewall policies, and other features on a per-SVM level to support that model.

Figure 2) Evolution of security domains.



Therefore, in addition to securing the storage system itself, it is necessary to practice good security hygiene across the entire IT infrastructure.

3.2 Policies and Procedures

It is assumed that organizations already have security policies in place to protect the system and users. For example, these policies typically address password management (strength, lifecycle) as well as basic concepts such as requiring all administrators (even occasional ones) to have unique IDs and passwords. They also address regular audit log backup and analysis, as well as handling computer forensic evidence upon suspicion of a breach of security.

Clustered Data ONTAP has been architected to be integrated into modern security architectures. However, there is often a trade-off between flexibility and security, and clustered Data ONTAP default configurations are often optimized for flexibility and convenience. In order to conform to security best practices, it is necessary to reconfigure some default settings after initial installation. This document provides guidance and recommendations to do that.

3.3 Strategy for Clustered Data ONTAP Security Recommendations

The steps to configure clustered Data ONTAP for security can be as follows. First, it is necessary to manage the login accounts for administration users. Clustered Data ONTAP provides preconfigured roles to assist with that. Next, it is necessary to selectively enable and disable chosen ciphers and hashes to match the security policies of the organization. Many of the weaker hashes and ciphers are enabled by default for convenience and flexibility, but are recommended to be disabled to enhance security. Finally, some external services need to be securely configured such as network timing, remote access, and administrative web services. These steps are all outlined in this document.

4 Initial Security Configuration

Configuring administrative access to clustered Data ONTAP is a necessary first step to securing the system. Starting with a new installation, it is necessary to configure the cluster-level (Cserver) admin accounts first. Then, after SVMs are created, each SVM (Vserver) administrator can configure additional SVM administrator accounts within that SVM as needed.

Setting up an administrative account (whether cluster-level Cserver or SVM-level Vserver) follows the same process. An account is created using this command:

```
security login create -user-or-group-name doug -application ssh -authmethod password -role admin
```

The necessary parameters include the “Vserver,” “username,” “application (access) method,” and “authentication method.” In addition, a “role” must be specified that correlates to the capabilities of that user. The role is a valuable feature in clustered Data ONTAP for managing permissions of users or groups of users to access and/or change configuration settings in the system. Roles are described in the following section.

4.1 Role-Based Access Control

Role-based access control, or RBAC, offers a powerful paradigm in managing the capabilities of users because the role that is assigned to a user defines what that user can see and do in the system. Standard preconfigured roles are available in clustered Data ONTAP that many customers find useful to use just as they are. In addition, however, it is possible to copy one of the default roles and modify it to suit a particular situation or security policy. This is a powerful capability and contributes to the security and flexibility of the system.

Roles are assigned the capability to use CLI commands either by individually specifying each command or, more commonly, by specifying groups of commands organized into hierarchical directories.

In the sections to follow, we'll examine the standard (default) roles provided with the system and then explore how to modify them as needed.

4.1.1 Standard Roles

For the cluster administrator level, five default Cserver roles are provided with clustered Data ONTAP. These are outlined in Table 1.

Table 1) Cluster administrator (Cserver) 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 SVM services NDMP • Read-only access to volumes • No access to other command directories
Read only	<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 can be seen by the roles in Table 1, the admin account is all powerful. Other accounts provide more limited capabilities. Finally, there is a “none” account with no capabilities. These roles then provide a starting point to create 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 read-only roles cover special use cases.

Similar to the Cserver roles in Table 1, there are four default roles for SVM administrators, as shown in Table 2.

Table 2) SVM administrator (Vserver) roles.

Role	Default Capabilities
Vsadmin	<ul style="list-style-type: none"> • Manage own administrator account local password and public key • Manage volumes, quotas, qtrees, Snapshot® copies, 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 conforms to the least privilege principle of common security best practices.

Also, just like the Cserver roles presented previously, these roles can be copied and modified to create custom roles as needed for each SVM independently.

At this point, many customers use the default roles to establish administrative user accounts. However, as mentioned previously, these roles can be copied (cloned) and then modified as needed to provide exactly the level of capability required for each user. This process is described next.

4.1.2 Modifying Custom Roles

The following command provides the ability to modify the capabilities to access a command directory of an existing role. It can be used, for example, to limit access to certain command sets.

```
Security login role modify
```

```
-vserver xxxx
-role xxxx
-cmddirname xxxx (command /Directory)
-access <Access> (access level) possible options are:
"none", "read-only", "all"
-query <query> Query
```

-Command/Directory specifies the command or command directory to which the role has access. If you want the default setting, use the string "DEFAULT" as the value.

-Query specifies the object that the role is allowed to access. The query must be applicable to the command or directory specified by -cmddirname above. The query object must be enclosed in double quotation marks (" "), and must be a valid field name.

Recommendations for this command are to set command directories to none or read only as far as possible. If write access is required, it can be added as needed.

Taken together, the preceding commands can be used to establish rules (password and so on) and capabilities (command directory) for a role. As a security recommendation, these should be utilized to provide the minimum capabilities required for a role and to enforce organizational policies around passwords and authentication.

In addition to managing the command capabilities of a role, you can also manage other security policies by role. These include password length, complexity, expiration, reuse, lockout, and so on.

For more information about administrator capabilities, see the *Clustered Data ONTAP 8.x System Administration Guide for Cluster Administrators* and the *Clustered Data ONTAP 8.x System Administration Guide for Vserver Administrators*.

4.2 Creating Admin User Accounts

After the roles have been established (either default or custom roles) for each admin account, use the `security login` commands to set up admin user accounts.

The `security login create` and `security login modify` commands create and modify login method for the management users. A login method consists of a user name, an application (access method), and an authentication method. It can optionally include an access-control role name.

```
security login create
-vserver vservice_name
-username user_name
-application application (see choices below)
-authmethod authentication_method (see choices below)
[-role role_name]
[-comment comments]
```

Application

The application or method is determined by a property of the administrator account that is created using the `-application` option; the possibilities are:

- System console (console)
- HTTP and HTTPS (http)
- Data ONTAP API (ontapi)
- Service processor (service-processor)
- SNMP (snmp)
- SSH (ssh)

For more information, see the *Clustered Data ONTAP 8.x System Administration Guide for Cluster Administrators* and the *Clustered Data ONTAP 8.x System Administration Guide for Vserver Administrators*.

Authentication

The authentication method is determined by a property of the administrative account that is created using the `-authmethod` option; the possibilities are:

- User password (password)
- Windows® Active Directory authentication (domain)
- LDAP or NIS authentication (nsswitch): SVM context only
- SSH public key authentication (publickey): ssh application only
- SNMP community strings (community): snmpuser only
- SNMP user-based security model (usm): snmpuser only

In addition, there are two additional authentication methods: cert and public key cert. These provide certificate-based authentication and public key authentication for SSH.

Best practice recommendations for these settings include:

User accounts:

- Use a separate login ID for each person (no sharing)
- Unique login passwords for each person should be kept secret
- Use a sufficiently complicated password and use password policies to enforce complexity

Least privilege principle:

- Allow users access to only what they need to do their job and no more

4.3 Microsoft Active Directory Authentication for Administrator Users

Large organizations with many administrators can benefit from using the same authentication model as for their users (clients). Therefore, they might want to provide a way for administrators to authenticate using LDAP to Active Directory. The process for setting this up is described in the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

4.4 Set Up a local Emergency Administrator Account

This defines the creation of an emergency administrative account to be used during network outages. This account has access using the controller serial console only. Authentication is by local password. Full CLI and security capability roles are granted to this account by it being granted the “admin” role. This account will be defined twice: once for serial console access and once for SP login access.

To create this account, enter the following commands:

1. Create the user.

```
security login create -username E_User1 -application console -authmethod password -role admin -vserver XXXXX -comment "Emergency Console Admin"
```

2. Set the password.

```
security login password -username E_User1
```

You should now log in as the emergency administrative user from the DSC serial console so that you can change the account password.

4.5 Restrict Default Administrative Accounts

There are two default administrative accounts: admin and diag.

Admin Account

The admin account has the role of admin and is allowed access using all applications. Because of the broad scope, many customers choose to lock this account for security reasons and use roles of narrower scope for routine administrative activities. The general methods for creating and modifying accounts have been described in the preceding sections. A detailed procedure description follows.

To accomplish this task, you should first create another account (`security login create`) with the admin role for each application that the new admin_account has permissions to use. Note that the authentication method should not be off-box for the primary admin_account, so that a networking problem does not lock you out of all administrative access.

Recommendation

Delete or lock the default “admin” account.

After you have tested the ability for the new account to access the cluster using the application, you can either delete (`security login delete`) or lock (`security login lock`) the admin account.

Note: The admin console entry cannot be deleted or locked if it is the only account with permissions to that application.

4.6 Disable Unnecessary Services

Any services that are not needed should be disabled. For example, if SNMP is not being used, then that service should be disabled. Likewise, if remote access will never be needed, then those services can be disabled as well. For more information about administrator capabilities, see the Clustered Data ONTAP 8.x System Administration Guide for Cluster Administrators and the Clustered Data ONTAP 8.x System Administration Guide for Vserver Administrators.

4.7 Timeouts

The timeout value specifies how long a CLI session remains idle before being automatically terminated. The CLI timeout value is clusterwide. That is, every node in a cluster uses the same CLI timeout value.

By default, the automatic timeout period of CLI sessions is 30 minutes. This is intended to provide convenience and flexibility during the installation process. However, it should be reduced after initial installation to enhance security. You use the `system timeout` commands (show or modify) to manage the automatic timeout period of CLI sessions.

A timeout period of 30 minutes is useful during initial installation and setup, but might be considered excessive for long-term use. NetApp recommends shortening the timeout value to 5 to 10 minutes.

Note: Commands that have not yet completed will continue to execute after the session times out.

Recommendation

Set the timeout to 10 minutes.

Note: This session timeout will also apply to console sessions.

Configuring Automatic Logout of Idle SSH Connections to SP

You can configure the automatic logout settings so that an SSH connection to the SP is automatically terminated after the connection has been idle for the number of minutes you specify.

Setting changes for automatic logout of idle SP SSH connections take effect only on SSH sessions that start after the changes. Automatic logout does not take effect if you access the SP through the serial console.

You use the `system timeout` commands to manage the automatic timeout period of CLI sessions. The default is 60 minutes. Additional information is in the Clustered Data ONTAP 8.3 System Administration Guide for Cluster Administrators.

4.8 Logging

Audit logging creates a chronological record of management activities. You can specify what types of activities in the management interface are audited.

Data ONTAP enables you to audit both “set” requests and “get” requests from the NetApp Manageability SDK remote API as well as the CLI. You can use the `security audit` commands to manage audit settings. Regardless of the settings for the `security audit` commands, set requests are always recorded in the `command-history.log` file and `mgwd.log` file. If logging all get requests causes the log to be filled too quickly, some customers choose to omit recording of get requests. This might happen, for example, whenever there is continuous polling by an external application.

By default, auditing of set requests is enabled and auditing of get requests is disabled. The `command-history.log` and `mgwd.log` files are rotated when they reach 100MB in size, and their previous 34 copies are preserved (with a maximum total of 35 files, respectively).

4.9 External Ports and VLANs

Ports are either physical ports (NICs) or virtualized ports, such as interface groups or VLANs. A logical interface (LIF) communicates over the network through the port to which it is currently bound. A LIF is essentially an IP address with the following associated characteristics:

- Role
- Home node
- Home port
- Firewall policy
- Failover policy

An Ethernet switch is used to maintain required connectivity for all ports. As a security recommendation, separate VLANs should be used for each of the following types of ports.

VLANs provide logical segmentation of networks by creating separate broadcast domains. A VLAN can span multiple physical network segments. The end stations belonging to a VLAN are related by function or application.

For example, end stations in a VLAN might be grouped by departments, such as engineering and accounting, or by projects, such as release1 and release2. Because physical proximity of the end stations is not essential in a VLAN, you can disperse the end stations geographically and still contain the broadcast domain in a switched network.

If an SVM is configured as a completely separate security domain, then putting it into a separate VLAN is also recommended. This enforces traffic separation and security. If SVMs are in the same security domain, then they can coexist on the same VLAN.

You can create a VLAN for maintaining separate broadcast domains within the same network domain by using the `network port vlan create` command. You cannot create a VLAN from an existing VLAN. You must specify either the `vlan-name` or the `port` and `vlan-id` options when creating a VLAN.

In 8.3, there are new features that can provide even more enhanced partitioning between storage domains. IPspaces provide support for separately owned routing domains with overlapping address spaces. This is common in multi-tenant applications (for example, private 10.10.x.x subnets). Each IP port and each SVM is hosted on one IPspace. This is useful for providing security partitioning between separate tenants. Use the following command to configure an IPspace:

```
network ipspace create -ip-space ipspace_name
```

Likewise, broadcast domains are also configurable now. The broadcast domain defines a layer 2 broadcast domain that is unique (possibly within an IPspace) and owns subnets and is associated with a list of ports and failover groups. Using broadcast domains for each separate VLAN is highly recommended. The following command is used to establish a broadcast domain:

```
network port broadcast-domain create -broadcast-domain broadcast_domain_name -mtu mtu_value[-ip-space ipspace_name] [-ports ports_list]
```

4.10 NTP Protocol

The Network Timing Protocol presents an opportunity for attackers to affect system clocking and therefore interfere with basic system operations. As a good security policy, NTP should use more than 1 server. A maximum of up to 10 NTP servers can be associated with the cluster and can be specified with an IPv4 address, IPv6 address, or fully qualified host name.

Use the `'cluster time-service ntp server'` command to create NTP server references. You can use the `'system services ntp server modify'` command to modify existing server records. You must do this for each member node in the cluster.

The recommended practice is to use multiple servers for timing, but local sources are preferred. This enhances system resiliency and makes it more difficult to launch NTP timing attacks. You will also need to choose a preferred (default) NTP server.

4.11 Managing Public Keys

You can associate, modify, or delete a public key to manage a user's authentication. You use the `security login publickey` commands to manage public keys.

You can manage public keys in the following ways:

- Adding a public key by associating and storing an existing public key in a valid OpenSSH format with a user account. Multiple public keys are allowed for each user account.
- Loading a public key from a universal resource identifier (URI), such as FTP or HTTP, and associating it with a user account. You can also overwrite an existing public key with the one you are loading.
- Displaying information about public keys.
- Modifying a public key that is associated with a specific user.
- Deleting a public key that is associated with a specific user.

To create or modify a public key or load a public key from a URI, your user account must be configured with the `publickey` login method (created by using the `security login create` command with the `-authmethod` parameter set to `publickey`). You use the `security login public key` commands to manage public keys. For information about these commands, see the appropriate `man` pages.

The recommended practice is to obtain a unique public key for each user.

4.12 Managing Digital Certificates for Server or Client Authentication

Use of a digital certificate makes sure that web communications are transmitted in encrypted form, assuring confidentiality and integrity of the data. Data ONTAP enables you to generate, install, and manage a self-signed or certificate authority (CA) signed digital certificate for server authentication. You can generate a certificate signing request, create, install, sign, display, revoke, or delete a digital certificate for server or client authentication.

CA-signed certificates are recommended in a production environment. A CA-signed digital certificate helps prevent man-in-the-middle attacks and provides better security protection than a self-signed digital certificate.

By default, Data ONTAP uses the SHA-2 cryptographic hashing function for signing a CSR or digital certificate, and the SHA1 and MD5 cryptographic hashing functions are also supported. (Note that MD5 is available for backward compatibility, but good security requires that it be disabled, as described elsewhere in this document.) Private keys generated by Data ONTAP are 2048-bit by default. Data

ONTAP also enables you to generate a 512-bit, 1024-bit, or 1536-bit private key. The higher the value, the more secure the key is.

Insecure hashes such as MD5 and SHA1 can be disabled with the following command:

```
security ssh modify -key-exchange-algorithms diffie-hellman-group-exchange-sha256
```

This document applies to clustered Data ONTAP 8.3. In order to view, delete, generate, or install certificates, you must be in the advanced mode in releases before 8.2.0. You can use the command `set advanced` to enter the advanced mode.

Using the following command, you can create, delete, generate, install, and show certificate files: `Security certificate` (add a “?” to see the parameter options).

The generation of a self-signed certificate is performed with the `create` command, and most of the fields have defaults.

For more information, see the section about generating and installing a CA-signed digital certificate in the Data ONTAP 8.3 Cluster-Mode System Administrator’s Guide.

4.13 Providing Mutual Authentication

You can configure the server (which can be the cluster or an SVM) to provide mutual authentication for greater security between the server and a group of admin clients.

When using mutual authentication, also called two-way authentication, both the server and the client present their certificates to each other and validate their respective identities to each other. To configure mutual authentication using a self-signed root CA certificate, you must create a self-signed root CA certificate, enable client authentication, generate and sign a certificate signing request (CSR) for each user, and install the client certificate on the client side. You must also set up user accounts for them to be authenticated by digital certificates.

You can also provide client authentication using a CSR signed by a third-party CA that is installed on the client and installing intermediate certificates of the CA that signed the certificate.

5 Management

5.1 SSHv2

SSH is the primary method to manage the SVM using the SVM administrators using the SVM LIF, and only SSHv2 is currently supported. There are several configuration categories related to SSHv2 use. These are:

- Access lists
- Encryption ciphers and key exchange algorithms
- Session timeout

Individual hosts or subnet IP addresses allowed to make SSH connections are controlled by the firewall policies applied to network interfaces (LIFs) where such connections occur.

5.1.1 SSH Encryption Ciphers and Key Exchange Algorithm Recommendations

Clustered Data ONTAP supports multiple encryption ciphers and key exchange algorithms used by SSH. These can be managed, allowing the use of only those that meet current security policy.

Table 3 shows the installed ciphers and algorithms.

Table 3) Installed encryption ciphers and key exchange algorithms.

Encryption Cipher	Key Exchange Algorithm
aes256-ctr	diffie-hellman-group-exchange-sha256
aes192-ctr	diffie-hellman-group-exchange-sha1
aes128-ctr	diffie-hellman-group14-sha1
aes256-cbc	diffie-hellman-group1-sha1
aes192-cbc	
aes128-cbc	
3des-cbc	

Note: These are applied to each SVM where SSH access might be possible (cluster and data). When the ciphers and algorithms used by the cluster admin SVM are modified, this will become the default for all subsequent data SVMs created.

The following command will display which ciphers and algorithms are in use clusterwide:

```
Security ssh show
```

Note: You can remove unwanted ciphers and algorithms or restore them as needed. However, the HMACs are not configurable.

Security recommendation: Because of known weaknesses with cipher block chaining ciphers, those suffixed by 'cbc' should be disabled and not be used. They are only supported for backward compatibility in older systems.

Enter the following command to modify which ciphers are used by the cluster SVM (and subsequently created SVMs):

```
security ssh modify -vserver st1-IAC-cDOT-1 -key-exchange-algorithms diffie-hellman-group-exchange-sha256 -ciphers aes256-ctr,aes192-ctr,aes128-ctr
```

5.2 Telnet

For security reasons, Telnet is not recommended as a terminal access protocol because it lacks encryption. It is off by default and provided only for backward compatibility on a temporary basis. It will be removed from future clustered Data ONTAP releases.

5.3 NetApp Manageability SDK and External Software Tools

The NetApp Manageability SDKs are XML-based remote procedure calls used primarily to Data ONTAP from NetApp products such as OnCommand® Unified Manager (OCUM), third-party tools, and customer tools. They are also used between other NetApp products.

The transport protocol can be HTTP or HTTPS.

For security, it is recommended to only use HTTPS and to disable HTTP. Some of the NetApp management applications permit HTTPS to fall back to HTTP, but this is not recommended behavior.

More information is in section 9.3 in this document.

6 Security Validation

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 and so are often useful in combination.

6.1 Vulnerability Scanners

Popular scanners include Nessus and Qualys. These scan for well-known vulnerabilities by implementing common exploit signatures and scanning header files in the system for versions of included code that have known vulnerabilities. Note that often a vulnerability scan against a NetApp product might turn up false positives. A vulnerability might be reported simply because of a header scan of the product. However, a reported vulnerability might not be applicable in Data ONTAP because that particular code function is not used or has been mitigated in other ways by the product.

NetApp uses a variety of scanners in internal testing and is constantly running tests. If a vulnerability is detected, a mitigation is published as soon as possible. Typically, this can be a simple configuration change, but if a patch is required, it will be posted as soon as possible on the site. Therefore, if a vulnerability is reported in a customer environment, it should be investigated on the support pages at <http://www.netapp.com/us/legal/vulnerability-handling-response-policy.aspx>.

6.2 Port Scanners

Port scanners interrogate TCP and UDP port numbers for open and responsive ports. The most popular example is NMAP. As a general rule, most ports that are open on a NetApp Data ONTAP system are open for a reason and should be left that way. However, port scanners often report ports open that actually represent no vulnerability. Some of these ports are actively filtered, which means that no “closed” response is sent back to the NMAP scanner.

7 System Monitoring

7.1 SNMP

Enabling SNMP provides a mechanism for monitoring your cluster. Managing SNMP involves enabling SNMP, configuring SNMP users, and configuring SNMP trap hosts for specific events. If you enable SNMP in Data ONTAP, SNMP management systems can query your storage system's SNMP agent for information. The SNMP agent gathers information and forwards it to the SNMP managers. The SNMP agent also generates trap notifications whenever specific events occur. The SNMP agent on the storage system has read-only privileges; that is, it cannot be used for any set operations or for taking a corrective action in response to a trap. In clustered Data ONTAP, SNMP is enabled clusterwide. See TR-4220 for more information on SNMP configuration and usage.

For diagnostic and other network management services, Data ONTAP provides an SNMP agent compatible with SNMP versions v1, v2c, and v3. SNMPv3 offers advanced security by using passphrases and encryption.

SNMPv3 is recommended as more secure than SNMPv1 and SNMPv2c. However, if you want to use SNMPv3, you must configure an SNMPv3 user to run the SNMP utilities from the SNMP manager.

Use the `security login create` command to create an SNMPv3 user. The security level can be authentication, no privacy; authentication, privacy; or no authentication, no privacy.

You are prompted to provide the following information:

- Engine ID (default value is local EngineID)
- Authentication protocol (none, md5, sha)
- Authentication password (eight-character minimum)
- Privacy protocol (none, des)
- Privacy protocol password (passphrase for encryption)

Recommendations

Use `snmpv3` with `auth (sha)` and `priv (des)`.

See the appropriate express guide or TR-4220 for detailed information about SNMP configuration and usage.

7.2 AutoSupport

AutoSupport is a mechanism that proactively monitors the health of your system and, if enabled, automatically provides status information to NetApp technical support, your internal support organization, and a support partner.

On new systems, AutoSupport on demand is enabled by default. It uses secure HTTPS tunnels to send encrypted status information to NetApp Support. However, it is sometimes necessary to use a proxy server to gain access through corporate firewalls. If you choose not to use AutoSupport, it can be disabled at any time using the `system node autosupport modify` command. However, it is recommended to leave AutoSupport enabled so that NetApp can support your system.

Older systems might still use e-mail to communicate this information back to NetApp. Although this method is still supported, moving to the more secure HTTPS model is recommended.

Recommendation

Always use the HTTPS transport when sending AutoSupport messages to NetApp Support:
`system node autosupport modify -node {nodename} -transport https`

Be sure and turn on AOD after switching to HTTPS.

For further information on AutoSupport, see the clustered Data ONTAP 8.3 Administration Guide. See the NetApp Support site for further information concerning MyAutoSupport: <http://now.netapp.com/NOW/knowledge/docs/olio/autosupport/>.

7.2 External Syslog Server

Many customers require the ability to export (send) event messaging to an external syslog server to allow event analysis.

This is accomplished by assigning a syslog server IP address to an “event destination” entry and then routing all event log messages to that event destination.

The following example shows the “built-in” event destinations provided by clustered Data ONTAP:

```
cDOT-1::> event destination show
```

Name	Mail Dest.	SNMP Dest.	Hide Syslog Dest.	Params
allevents	-	-	false	
asup	-	-	false	
criticals	-	-	false	
pager	-	-	false	
traphost	-	-	false	

5 entries were displayed.

The event destination “allevents” will be used and configured to point to an external syslog server.

The following command will set an IP address for an external syslog server (multiple addresses may be given, separated by commas):

Setting IP Address for External Syslog Server

```
event destination modify -name allevents -syslog 10.63.165.225
```

Use **event destination show** to show the result of the preceding command.

Enter the following command to route these messages to the “allevents” destination:

```
event route add-destinations -messagename * -destinations allevents
```

After setting these configurations, there might be some lag time before messages are posted to the external servers.

It is recommended to activate and use an external syslog server to provide the data to help detect attacks on the system. However, just like clustered Data ONTAP, the syslog server should be set up as a secure system following the organization’s security policies.

8 Hardware Security Considerations

All NetApp FAS storage systems provide an out-of-band (OOB) management port for the maintenance and management of the storage system hardware. These network-based ports provide communications using SSH to make sure of confidentiality of the interactive sessions.

8.1 Service Processor

The service processor (SP) CLI commands enable you to remotely access and administer the storage system and diagnose error conditions. Also, the SP extends AutoSupport capabilities by sending alerts and notifications through an AutoSupport message.

In order to access the storage system through the SP interface, an account must have the login service-processor method enabled. The storage system admin role has the service-processor method by default.

SP firmware 1.2 and later will track 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 SP will stop all communication with that IP address for the next 15 minutes. Normal communication will resume after 15 minutes, but if repeated login failures are detected again, communication will again be suspended for the next 15 minutes.

For detailed information on the SP and its capabilities, see the “Managing a Node Remotely by Using the Service Processor” section of the “Data ONTAP 8.3 Cluster-Mode System Administration Guide.”

The recommendation is to configure the SP for remote access in the event of a disruptive event in the system. The login accounts used on the SP can match the general accounts on Data ONTAP used for general administrative access for convenience.

9 System Services

9.1 Firewall Service

Each LIF type has a default role and firewall policy attached to it. Firewall policy entries are composed of a protocol type, a firewall action, and an address list to which the action applies. Note that the firewall service is configured per SVM for each LIF and applies only to the control plane. Note also that a particular protocol (port) might appear open but is ignored for security purposes. For that reason, port scanners might report different results than the configurations for the firewall services presented here.

The network protocols in Table 4 may be managed by a firewall policy.

Table 4) Firewall policy protocols.

Protocol	Actions	Address List Members
DNS	Allow and/or deny	Range of IPv4 and IPv6 addresses
HTTP	Allow and/or deny	Range of IPv4 and IPv6 addresses
HTTPS	Allow and/or deny	Range of IPv4 and IPv6 addresses
NDMP	Allow and/or deny	Range of IPv4 and IPv6 addresses
NTP	Allow and/or deny	Range of IPv4 and IPv6 addresses
RSH	Allow and/or deny	Range of IPv4 and IPv6 addresses
SNMP	Allow and/or deny	Range of IPv4 and IPv6 addresses
SSH	Allow and/or deny	Range of IPv4 and IPv6 addresses
Telnet	Allow and/or deny	Range of IPv4 and IPv6 addresses

The address list for each policy entry is a range of IPv4 and IPv6 networks or an individual IP address associated with the policy. An IP address is shown explicitly, while a range is indicated by IP and subnet mask. For example, the range 192.168.21.50/28 would indicate access to IP addresses from 192.168.21.49 to 192.168.21.63. The IPv4 address 0.0.0.0/0 is a generic entry meaning all IPv4 networks and the addresses contained therein. The IPv6 address ::/0 is a generic entry meaning all IPv6 networks and their related addresses.

Note: Firewall policies may contain a single action entry for allow, deny, or one of each type.

Note: Firewall policies for cluster and intercluster logical interfaces support IPv4 addresses only.

These policies are further broken down by type: cluster, data, intercluster, and mgmt:

- To view all the policies within the cluster:
 - system services firewall policy show
- To view the policy applied on a network interface:
 - network interface show -fields firewall-policy
- To create a new policy:
 - system services firewall policy create-policy *policy_name* –services
 - *service_name* [-allowed-ipslist_of_IP_addresses_and_netmasks]

Following are details on how to create these two policies.

A secured firewall policy for management interfaces is defined using the following commands:

```
system services firewall policy create -policy <policy-name> -service
<protocol_name> -action <allow | deny> -ip-list <IPv4 and/or IPv6
address list>
```

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

Table 5 lists the desired entries for the secured management firewall policy (`secure_mgmt`).

Table 5) Secure management firewall policy entry settings.

Protocol	Action	Access List	Net Effect
DNS	Allow	IP address list for allowed DNS servers	Allow DNS access to list
DNS	Deny	0.0.0.0/0, ::/0	Deny all except preceding allow list
HTTP	Deny	0.0.0.0/0, ::/0	Deny all
HTTPS	Deny	0.0.0.0/0, ::/0	Deny all recommended, but required for applications based on NetApp Manageability SDK management
NDMP	Deny	0.0.0.0/0, ::/0	Deny all
NTP	Allow	IP address list for allowed NTP servers	Allow NTP access to list
NTP	Deny	0.0.0.0/0, ::/0	Deny all except preceding allow list
RSH	Deny	0.0.0.0/0, ::/0	Deny all
SNMP	Deny	0.0.0.0/0, ::/0	Deny all
SSH	Allow	IP address list for allowed SSH clients	Allow SSH access from list
SSH	Deny	0.0.0.0/0, ::/0	Deny all except preceding allow list
Telnet	Deny	0.0.0.0/0, ::/0	Deny all

For this cluster, the following commands are used to create a policy named `secure_mgmt`:

Create Firewall Policy '`secure_mgmt`'

```
Cluster1::> system services firewall policy create -policy secure_mgmt
-service dns -action allow -ip-list 10.63.165.0/24, ::/0
```

Note: The preceding example is too long to display as a single command line in this document. Do not press Enter until the entire command line has been entered.

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

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

This step will apply the firewall policy for secure management, created earlier, to the LIFs on the cluster and node management interfaces (e0M). Perform this step for the cluster SVM and each cluster node SVM.

For the cluster SVM use:

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

For the node SVMs use:

```
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
```

9.2 NDMP Service

There are two default settings that need to be changed: the ndmp account and the clear text password setting. Use the command `system services ndmp show` to list the status and set attributes as follows:

```
System services ndmp modify -node * -clear-text false
```

```
System services ndmp modify -node Cluster07-02 -user-id garlic
```

You will be asked to enter and confirm a password. Use the command `system services ndmp show` to verify the changes.

The recommendation is that NDMP never be in clear text. However, this option remains available for backward compatibility.

9.3 Web Services

Managing Access to Web Services

A web service is an application that users can access by using HTTP or HTTPS. Note that HTTPS is necessary to support applications based on NetApp Manageability SDK and web applications and should therefore be enabled. The cluster administrator can set up the web protocol engine, configure SSL, enable a web service on the cluster, and enable users of a role to access a web service. Use the following command to disable HTTP (and only use HTTPS) for web services:

```
vserver services web modify
```

Data ONTAP supports the following web services:

- **Service processor infrastructure support (spi).** You can enable this service for the nodes or the cluster. Enabling this service makes a node's log and core files available for HTTP or HTTPS access through the cluster's management LIF or any node's management LIF. For security reasons, HTTP is not recommended. HTTPS should always be configured.
- **Data ONTAP APIs (ontapi).** This service enables you to run Data ONTAP APIs to execute administrative functions with a remote program. This service is required for external management tools.

Managing Web Protocol Engine

You can configure the web protocol engine on the cluster to control whether web access is allowed and what SSL versions can be used. You can also display the configuration settings for the web protocol engine.

You can manage the web protocol engine at the cluster level in the following ways:

- Configuring the web protocol engine to control whether remote clients can use HTTP or HTTPS to access web service content (HTTPS is recommended).

- Specifying whether SSLv3 or SSLv2 should be used for secure web access: By default, SSLv3 is enabled, and SSLv2 is disabled. SSLv3 is recommended, but SSLv2 is provided for backward compatibility.
- Displaying the configuration and status of web services: You use the `system services web` commands to manage the web protocol engine at the cluster level. If a firewall is enabled, the firewall policy for the logical interface (LIF) to be used for web services must be set up to allow HTTPS access. Also, SSL for the cluster or SVM that offers the web service must also be enabled, and you must provide a digital certificate for the cluster or SVM.

SSL Services

To check the state of the services, use the command `system services web show`.

To modify the SSL configuration, use the command `system services web modify`.

10 Summary and Conclusion

Secure configuration of clustered Data ONTAP systems is an important consideration. The default configurations and settings have been chosen to be the most convenient and flexible for the majority of customers, especially during installation and setup. However, the recommendations outlined in this document can provide an additional measure of security after the system is put into service. The information in this document provides guidelines for basic secure configuration, but is by no means exhaustive. It exists as a component of a larger security strategy and policy within the data center.

Security should be thought of as a continuum rather than a binary (yes/no) state. It is not possible to say a system is completely secure, but practical steps can be taken to make it reasonably secure. Further, security is a dynamic state and must be constantly evaluated and updated as necessary. This document will continue to be upgraded as NetApp continues to enhance the security of its products. NetApp welcomes any feedback regarding security strategy and positioning of these products.

10.1 Appendix

Checklist of Security Recommendations

Table 6) Checklist of Security Recommendations

Security Recommendation	Location		
Configure separate SVMs for each set of discrete storage users	Section 2.2		
Create/modify roles to limit scope of authority	Section 4.2		
Create separate and secure user accounts for each admin user with appropriate roles	Section 4.2		
Restrict default administrator accounts	Section 4.5		
Change admin password	Section 4.3		
Create reduced scope administrative account	Section 4.3		
Set password and lock the diag account	Section 4.5		
Verify RSH, Telnet are disabled	Section 4.6		
Set CLI session timeouts	Section 4.7		

Security Recommendation	Location		
Set SSH session timeouts	Section 4.7		
Configure audit logging settings	Section 7.3		
Disable IPv6 if not used	Section 4.6		
Configure NTP settings to fixed IP addresses	Section 4.10		
Install digital certificates	Section 4.12		
Set mutual authentication	Section 4.13		
Configure SSH ciphers (disable weak ones)	Section 5.1		
Disable HTTP (use HTTPS) for NetApp Manageability SDK	Section 5.3		
Run external vulnerability scanner	Section 6.1		
Run external port scanner	Section 6.2		
Configure for SNMPv3 only	Section 7.1		
Configure AutoSupport	Section 7.2		
Configure external syslog server	Section 7.3		
Verify and configure firewall	Section 9.1		
Set NDMP service to hash password	Section 9.2		
Set web services to SSLv3	Section 9.3		

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 © 1994–2015 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, Go Further, Faster, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL are trademarks or registered trademarks of NetApp, Inc., in the United States and/or other countries. A current list of NetApp trademarks is available on the Web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.

Cisco and the Cisco logo are trademarks of Cisco in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. TR-4393-0415