



Security and data encryption

ONTAP 9

NetApp
September 20, 2024

Table of Contents

- Security and data encryption 1
 - About NetApp ransomware protection 1
 - Autonomous Ransomware Protection 10
 - Virus protection with Vscan 35
 - ONTAP hardening guidelines 75
 - Audit NAS events on SVMs 122
 - Use FPolicy for file monitoring and management on SVMs 168
 - Verify access using security tracing 229
 - Manage encryption with System Manager 241
 - Manage encryption with the CLI 242
 - Enable Zero Trust model 333

Security and data encryption

About NetApp ransomware protection

Ransomware and NetApp's protection portfolio

Ransomware remains one of the most significant threats causing business interruption for organization in 2024. According to the [Sophos State of Ransomware 2024](#), ransomware attacks affected 72% of their surveyed audience. Ransomware attacks have evolved to be more sophisticated and targeted, with threat actors employing advanced techniques like artificial intelligence to maximize their impact and profits.

Organizations must look across their entire security posture from perimeter, network, identity, application, and where the data lives at the storage level and secure these layers. Adopting a data-centric approach to cyber protection at the storage layer is crucial in today's threat landscape. Although no single solution can thwart all attacks, using a portfolio of solutions, including partnerships and third parties, provides a layered defense.

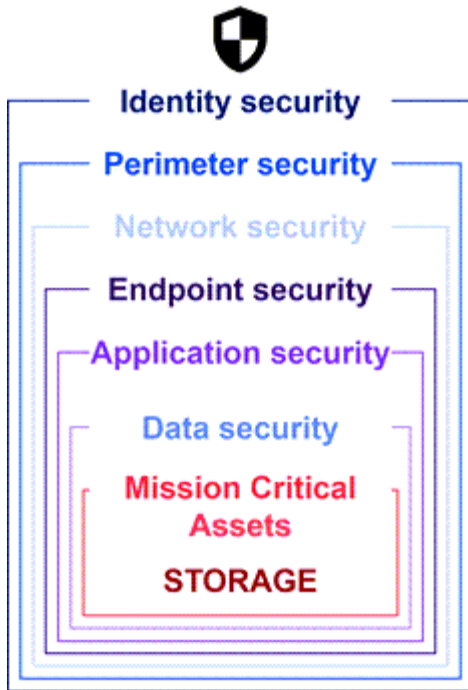
The [NetApp product portfolio](#) provides various effective tools for visibility, detection, and remediation, helping you to spot ransomware early, prevent spread, and recover quickly, if necessary, to avoid costly downtime. Traditional layered defense solutions remain prevalent, as do third parties and partner solutions for visibility and detection. Effective remediation remains a crucial part of the response to any threat. The unique industry approach leveraging immutable NetApp Snapshot technology and SnapLock logical air gap solution is an industry differentiator and the industry best practice for ransomware remediation capabilities.



Beginning in July 2024, content from the technical report *TR-4572: NetApp Ransomware Protection*, which was previously published as a PDF, has been integrated with the rest of the ONTAP product documentation.

Data is the primary target

Cybercriminals increasingly target data directly, recognizing its value. While perimeter, network, and application security are important, they can be bypassed. Focusing on protecting data at its source, the storage layer, provides a critical last line of defense. Gaining access to production data and encrypting or rendering it inaccessible is the objective of ransomware attacks. To get there, attackers must have already pierced existing defenses deployed by organizations today, from perimeter to application security.



Unfortunately, many organizations don't take advantage of security capabilities at the data layer. This is where NetApp ransomware protection portfolio comes in, protecting you at the last line of defense.

The real cost of ransomware

The ransom payment itself is not the largest monetary effect on a business. Although the payment is not insignificant, it pales in comparison to the downtime cost of suffering a ransomware incident.

Ransom payments are just one element of recovery costs when dealing with ransomware events. Excluding any ransoms paid, in 2024 organizations reported a mean cost to recover from a ransomware attack of \$2.73M, an increase of almost \$1M from the \$1.82M reported in 2023 according to the [2024 Sophos State of Ransomware](#) report. For organizations that rely heavily on IT availability, such as e-commerce, equities trading, and health care, costs can be 10 times higher or more.

Cyber insurance costs also continue to rise given the very real likelihood of a ransomware attack on insured companies.












Ransomware protection at the data layer

NetApp understands your security posture is wide and deep across your organization from the perimeter to the where your data lives at the storage layer. Your security stack is complex and should provide security at every level of your technology stack.

Real-time protection at the data layer is even more important and has unique requirements. To be effective, solutions at this layer must offer these critical attributes:

- **Security by design** to minimize chance of successful attack
- **Real-time detection and response** to minimize impact of a successful attack
- **Air-gapped WORM protection** to isolate critical data backups
- **A single control plane** for comprehensive ransomware defense

NetApp can deliver all of this and more.

Secure by Design Data-centric on-box protection	 Immutable backups & snapshots	 Multi-user verification and authentication	 Malicious file blocking	Ransomware Recovery Guarantee No data loss with NetApp Snapshots, guaranteed.
Real-time Detection & Response 99% detection accuracy to minimize attack impact	 AI-powered detection	 Actional intelligence for insider threats		
Air-gapped WORM protection with cyber vaulting Layered approach to further fortify data against ransomware attacks	 Isolated, immutable & indelible WORM snapshots			
Single control plane for comprehensive ransomware defense			BlueXP Ransomware Protection	
 PROTECT Recommends workload protection policies and applies them with one-click.	 DETECT Detects potential attacks on your workload data in near real-time using industry leading AI/ML.	 RESPOND Automatically responds by taking immutable and indelible Snapshots when a potential attack is suspected. Integrates with popular SIEMs.	 RECOVER Rapidly restores workloads with application consistency, through simplified orchestrated recovery.	 GOVERN Implements your ransomware protection strategy and policies, and monitors outcomes.

NetApp's ransomware protection portfolio

NetApp's [built-in ransomware protection](#) delivers real-time, robust, multi-faceted defense for your critical data. At its core, advanced AI-powered detection algorithms continuously monitor data patterns, swiftly identifying potential ransomware threats with 99% accuracy. Reacting quickly to attacks allows our storage to quickly snapshot data and secure the copies ensuring rapid recovery.

To further fortify data, NetApp's [cyber vaulting](#) capability isolates data with a logical air gap. By safeguarding critical data, we ensure rapid business continuity.

NetApp [BlueXP ransomware protection](#) reduces operational burdens with a single control plane to intelligently coordinate and execute an end-to-end workload-centric ransomware defense, so you can identify and protect critical workload data at risk with a single click, accurately and automatically detect and respond to limit the impact of a potential attack, and recover workloads within minutes, not days, safeguarding your valuable workload data and minimizing costly disruption.

As a native, built-in ONTAP solution for protecting unauthorized access to your data, [multi-admin verification \(MAV\)](#) has a robust set of capabilities that ensure that operations such as deleting volumes, creating additional administrative users, or deleting snapshot copies can be executed only after approvals from at least a second designated administrator. This prevents compromised, malicious, or inexperienced administrators from making undesirable changes or deleting data. You can configure as many designated administrator approvers as you want before a snapshot copy can be deleted.



NetApp ONTAP addresses the requirement for web-based [multi-factor authentication \(MFA\)](#) in System Manager and for SSH CLI authentication.

NetApp's ransomware protection offers peace of mind in an ever-evolving threat landscape. Its comprehensive approach not only defends against current ransomware variants but also adapts to emerging threats, providing long-term security for your data infrastructure.

Learn about other protection options

- [Active IQ ransomware protection](#)
- [Cloud Insights Storage Workload Security \(CISWS\)](#)
- [FPolicy](#)
- [SnapLock and tamperproof snapshot copies](#)

Ransomware recovery guarantee

NetApp offers a guarantee to restore Snapshot data if a ransomware attack occurs. Our guarantee: If we can't help you restore your snapshot data, we'll make it right. The guarantee is available on new purchases of AFF A-Series, AFF C-Series, ASA, and FAS systems.

Learn more

- [Recovery guarantee service description](#)
- [Ransomware recovery guarantee blog.](#)

Related information

- [NetApp Support site resources page](#)
- [NetApp product security](#)

SnapLock and tamperproof snapshot copies for ransomware protection

A vital weapon in NetApp's Snap arsenal is SnapLock, which has proven highly effective in safeguarding against ransomware threats. By preventing unauthorized data deletion, SnapLock provides an additional layer of security, ensuring that critical data remains intact and accessible even in the event of malicious attacks.

SnapLock Compliance

SnapLock Compliance (SLC) provides indelible protection for your data. SLC prohibits data from being deleted even when an administrator attempts to re-initialize the array. Unlike other competitive products, SnapLock Compliance is not vulnerable to social engineering hacks through those products' support teams. Data protected by SnapLock Compliance volumes is recoverable until that data has reached its expiration date.

To enable SnapLock, an [ONTAP One](#) license is required.

Learn more

- [Snaplock documentation](#)

Tamperproof Snapshot copies

Tamperproof Snapshot (TPS) copies provide a convenient and fast way to protect data from malicious acts. Unlike SnapLock Compliance, TPS is typically used on primary systems where the user can protect the data for a determined time and left locally for fast recoveries or where data does not need to be replicated off of the primary system. TPS uses SnapLock technologies to prevent the primary snapshot copy from being deleted even by an ONTAP administrator using the same SnapLock retention expiration period. Snapshot copy deletion is prevented even if the volume is not SnapLock enabled, although snapshots do not have the same indelible nature of SnapLock Compliance volumes.

To make snapshot copies tamperproof, an [ONTAP One](#) license is required.

Learn more

- [Lock a snapshot copy for protection against ransomware attacks.](#)

FPolicy file blocking

FPolicy blocks unwanted files from being stored on your enterprise-grade storage appliance. FPolicy also gives you a way to block known ransomware file extensions. A user still has full access permissions to the home folder, but FPolicy doesn't allow a user to store files your administrator marks as blocked. It doesn't matter if those files are MP3 files or known ransomware file extensions.

Block malicious files with FPolicy native mode

NetApp FPolicy native mode (an evolution of the name, File Policy) is a file-extension blocking framework that allows you to block unwanted file extensions from ever entering your environment. It has been part of ONTAP for over a decade and is incredibly useful in helping you protect against ransomware. This Zero Trust engine is valuable because you get extra security measures beyond access control list (ACL) permissions.

In ONTAP System Manager and BlueXP, a list of over 3000 file extensions is available for reference.



Some extensions might be legitimate in your environment and blocking them can lead to unexpected issues. Create your own list that is appropriate for your environment before configuring native FPolicy.

FPolicy native mode is included in all ONTAP licenses.

Learn more

- [Blog: Fighting Ransomware: Part Three — ONTAP FPolicy, another powerful native \(aka free\) tool](#)

Enable user and entity behavior analytics (UEBA) with FPolicy external mode

FPolicy external mode is a file activity notification and control framework that provides visibility of file and user activity. These notifications can be used by an external solution to perform AI-based analytics to detect malicious behavior.

FPolicy external mode can also be configured to wait for approval from the FPolicy server before allowing specific activities to go through. Multiple policies like this can be configured on a cluster, giving you great flexibility.



FPolicy servers must be responsive to FPolicy requests if configured to provide approval; otherwise, storage system performance might be negatively impacted.

FPolicy external mode is included in [all ONTAP licenses](#).

Learn more

- [Blog: Fighting Ransomware: Part Four — UBA and ONTAP with FPolicy external mode.](#)

Cloud Insights Storage Workload Security (CISWS)

Storage Workload Security (SWS) is a feature of NetApp Cloud Insights that greatly enhances the security posture, recoverability, and accountability of an ONTAP

environment. SWS takes a user-centric approach, tracking all file activity from every authenticated user in the environment. It uses advanced analytics to establish normal and seasonal access patterns for every user. These patterns are used to quickly identify suspicious behavior without the need for ransomware signatures.

When SWS detects a potential ransomware, data deletion, or exfiltration attack, it can take automatic actions such as:

- Take a snapshot of the affected volume.
- Block the user account and IP address that is suspected of malicious activity.
- Send an alert to admins.

Because it can take automated action to quickly stop an insider threat as well as track every file activity, SWS makes recovery from a ransomware event much simpler and faster. With advanced auditing and forensics tools built in, users can immediately see what volumes and files were affected by an attack, which user account the attack came from, and what malicious action was performed. Automatic snapshots mitigate the damage and accelerate file restoration.

Total Attack Results

5	0	1,488
Affected Volumes	Deleted Files	Encrypted Files

1,488 Files have been copied, deleted, and potentially encrypted by **1 user account**.

This is potentially a sign of Ransomware Attack.

The extension ".wanna" was added to each file.

Alerts from ONTAP's Autonomous Ransomware Protection (ARP) are also visible in SWS, providing a single interface for customers using both ARP and SWS to protect from ransomware attacks.

Learn more

- [NetApp Cloud Insights](#)

NetApp ONTAP built-in on-box AI-based detection and response

As ransomware threats become more and more sophisticated, so should your defense mechanisms. NetApp's autonomous ransomware protection (ARP) is powered by AI with intelligent anomaly detection that is built in to ONTAP. Turn it on to add another layer of defense to your cyber resiliency.

ARP and ARP/AI are configurable through the ONTAP built-in management interface, System Manager, and enabled on a per-volume basis.

Autonomous Ransomware Protection (ARP)

Autonomous Ransomware Protection (ARP), another native built-in ONTAP solution since 9.10.1, looks at NAS storage volume workload file activity and data entropy to automatically detect potential ransomware. ARP provides administrators with real-time detection, insights, and a data recovery point for unprecedented on-box

potential ransomware detection.

For ONTAP 9.15.1 and earlier versions that support ARP, ARP starts in learning mode to learn typical workload data activity. This can take seven days for most environments. After learning mode is complete, ARP will automatically switch to active mode and start looking for abnormal workload activity that might potentially be ransomware.

If abnormal activity is detected, an automatic snapshot copy is immediately taken, which provides a restoration point as close as possible to the time of attack with minimal infected data. Simultaneously, an automatic alert (configurable) is generated that allows administrators to see the abnormal file activity so that they can determine whether the activity is indeed malicious and take appropriate action.

If the activity is an expected workload, administrators can easily mark it as a false positive. ARP learns this change as normal workload activity and no longer flags it as a potential attack going forward.

To enable ARP, an [ONTAP One](#) license is required.

Learn more

- [Autonomous Ransomware Protection](#)

Autonomous Ransomware Protection/AI (ARP/AI)

Introduced as a tech preview in ONTAP 9.15.1, ARP/AI takes NAS storage systems on-box real-time detection to the next level. The new AI-powered detection technology is trained on over a million files and various known ransomware attacks. In addition to the signals used in ARP, ARP/AI also detects header encryption. The AI power and additional signals allow ARP/AI to deliver better than 99% detection accuracy. This has been validated by SE Labs, an independent test lab that gave ARP/AI its highest AAA rating.

Because training the models continuously happens in the cloud, ARP/AI does not require a learning mode. It is active the moment it is turned on. Continuous training also means that ARP/AI is always validated against new ransomware attack types as they arise. ARP/AI also comes with auto-update capabilities that deliver new parameters to all customers to keep ransomware detection up to date. All other detection, insight, and data recovery point capabilities of ARP are maintained for ARP/AI.

To enable ARP/AI, an [ONTAP One](#) license is required.

Learn more

- [Blog: NetApp's AI-based real-time ransomware detection solution achieves AAA rating](#)

Air-gapped WORM protection with cyber vaulting


NetApp's approach to a cyber vault is a purpose-built reference architecture for a logically air-gapped cyber vault. This approach takes advantage of security hardening and compliance technologies, such as SnapLock, to allow for immutable and indelible snapshots.

Cyber vaulting with SnapLock Compliance and a logical air gap

A growing trend is for attackers to destroy the backup copies and, in some cases, even encrypt them. That is why many in the cybersecurity industry recommend using air gap backups as part of an overall cyber resiliency strategy.

The problem is that traditional air gaps (tape and offline media) can significantly increase restoration time, thus

increasing downtime and the overall associated costs. Even a more modern approach to an air-gap solution can prove problematic. For example, if the backup vault is temporarily opened to receive new backup copies and then disconnects and closes its network connection to primary data to once again be "air gapped", an attacker could take advantage of the temporary opening. During the time the connection is online, an attacker could strike to compromise or destroy the data. This type of configuration also generally adds unwanted complexity. A logical air gap is an excellent substitute for a traditional or modern air gap because it has the same security protection principles while keeping the backup online. With NetApp, you can solve the complexity of tape or disk air gapping with logical air gapping, which can be achieved with immutable snapshot copies and NetApp SnapLock Compliance.



NetApp released the SnapLock feature more than 10 years ago to address the requirements of data compliance, such as the Health Insurance Portability and Accountability Act (HIPAA), Sarbanes-Oxley, and other regulatory data rules. You can also vault primary snapshot copies to SnapLock volumes so that the copies can be committed to WORM, preventing deletion. There are two SnapLock license versions: SnapLock Compliance and SnapLock Enterprise. For ransomware protection, NetApp recommends SnapLock Compliance because you can set a specific retention period during which snapshot copies are locked and cannot be deleted, even by ONTAP administrators or NetApp Support.

Learn more

- [Blog: Layered ransomware protection with NetApp's Cyber Vault solution](#)

Tamperproof snapshot copies

While leveraging SnapLock Compliance as a logical air gap provides the ultimate protection in preventing attackers from deleting your backup copies, it does require you to move the snapshot copies using SnapVault to a secondary SnapLock-enabled volume. As a result, many customers deploy this configuration on secondary storage across the network. This can lead to longer restoration times versus restoring a primary volume Snapshot copy on primary storage.

Beginning in ONTAP 9.12.1, tamperproof snapshot copies provide near SnapLock Compliance level protection for your snapshot copies on primary storage and in primary volumes. There is no need to vault the snapshot copy using SnapVault to a secondary SnapLocked volume. Tamperproof snapshot copies use SnapLock technology to prevent the primary snapshot copy from being deleted, even by a full ONTAP administrator using the same SnapLock retention expiration period. This allows for quicker restore times and the ability for a FlexClone volume to be backed up by a tamperproof, protected snapshot copy, something you cannot do with a traditional SnapLock Compliance vaulted Snapshot copy.

The major difference between SnapLock Compliance and tamperproof snapshot copies is that SnapLock Compliance does not allow the ONTAP array to be initialized and wiped if SnapLock Compliance volumes exist with vaulted Snapshot copies that have not yet reached their expiration date. To make Snapshot copies tamperproof, a SnapLock Compliance license is required.

Learn more

- [Lock a snapshot copy for protection against ransomware attacks](#)

Active IQ ransomware protection

NetApp Active IQ is a digital advisor that simplifies the proactive care and optimization of NetApp storage with actionable intelligence for optimal data management. Fueled by telemetry data from our highly diverse installed base, it uses advanced AI and ML techniques to uncover opportunities to reduce risk and improve the performance and efficiency of your storage environment.

Not only can [NetApp Active IQ](#) help [eliminate security vulnerabilities](#), but it also provides insights and guidance specific to protecting against ransomware. A dedicated wellness card shows the actions needed and the risks addressed, so you can be sure that your systems are meeting those best practices recommendations.



Risks and actions tracked on the Ransomware Defense Wellness page include the following (and much more):

- Volume snapshot copy count is low, decreasing potential ransomware protection.
- FPolicy is not enabled for all storage virtual machines (SVMs) configured for NAS protocols.


To see Active IQ ransomware protection in action, see [NetApp Active IQ](#).

Comprehensive resilience with BlueXP ransomware protection

It is important for ransomware detection to occur as early as possible so that you can prevent the spread and avoid costly downtime. An effective ransomware detection strategy, however, should include more than a single layer of protection. NetApp's ransomware protection takes a comprehensive approach that includes real-time, on-box capabilities extending to data services using BlueXP and an isolated, layered solution for cyber vaulting.

BlueXP ransomware protection

BlueXP is a single control plane to intelligently orchestrate a comprehensive, workload-centric ransomware defense. BlueXP ransomware protection brings together the powerful cyber-resilience features of ONTAP, such as ARP, FPolicy, and tamperproof snapshots, and BlueXP data services, such as BlueXP backup and recovery. It also adds recommendations and guidance with automated workflows to provide an end-to-end defense through a single UI. It operates at the workload level to ensure that the applications that run your business are protected and can be recovered as quickly as possible in case of an attack.



Customer benefits:

- Assisted ransomware preparedness reduces operational overhead and improves efficacy
- AI/ML-powered anomaly detection delivers greater accuracy and faster response to contain risk
- Guided application-consistent restoration allows you to recover workloads more easily and within minutes

BlueXP ransomware protection makes these NIST functions easier to achieve:

- Automatically **discover** and prioritize data in NetApp storage **with a focus on top application-based workloads**.
- **One-click protection** of top-workload data backup, immutable, secure configuration, malicious file blocking, and different security domain.
- **Accurately detect** ransomware as **quickly** as possible using **next-generation AI-based anomaly detection**.
- Automated response and workflows and integration with top **SIEM and XDR solutions**.
- Rapidly restore data using a simplified **orchestrated recovery** to accelerate application uptime.
- Implement your ransomware protection **strategy** and **policies**, and **monitor outcomes**.

Autonomous Ransomware Protection

Autonomous Ransomware Protection overview

Beginning with ONTAP 9.10.1, the Autonomous Ransomware Protection (ARP) feature uses workload analysis in NAS (NFS and SMB) environments to proactively detect and warn about abnormal activity that might indicate a ransomware attack.

When an attack is suspected, ARP also creates new Snapshot copies, in addition to existing protection from scheduled Snapshot copies.

Licenses and enablement

ARP requires a license. ARP is available with the [ONTAP ONE license](#). If you do not have the the ONTAP One license, other licenses are available to use ARP, which differ depending on your version of ONTAP.

ONTAP releases	License
ONTAP 9.11.1 and later	Anti_ransomware
ONTAP 9.10.1	MT_EK_MGMT (Multi-Tenant Key Management)

- If you are upgrading to ONTAP 9.11.1 or later and ARP is already configured on your system, you do not need to purchase the new Anti-ransomware license. For new ARP configurations, the new license is required.
- If you are reverting from ONTAP 9.11.1 or later to ONTAP 9.10.1, and you have enabled ARP with the Anti-ransomware license, you will see a warning message and might need to reconfigure ARP. [Learn about reverting ARP.](#)

You can configure ARP on a per-volume basis using either System Manager or the ONTAP CLI.

ONTAP ransomware protection strategy

An effective ransomware detection strategy should include more than a single layer of protection.

An analogy would be the safety features of a vehicle. You don't rely on a single feature, such as a seatbelt, to completely protect you in an accident. Air bags, anti-lock brakes, and forward-collision warning are all additional safety features that will lead to a much better outcome. Ransomware protection should be viewed in the same way.

While ONTAP includes features like FPolicy, Snapshot copies, SnapLock, and Active IQ Digital Advisor to help protect from ransomware, the following information focuses on the ARP on-box feature with machine learning capabilities.

To learn more about ONTAP's other anti-ransomware features, see [Ransomware and NetApp's protection portfolio.](#)

What ARP detects

ARP is designed to protect against denial-of-service attacks where the attacker withholds data until a ransom is paid. ARP offers real-time ransomware detection based on:

- Identification of the incoming data as encrypted or plaintext.
- Analytics, which detects
 - **Entropy:** an evaluation of the randomness of data in a file
 - **File extension types:** An extension that does not conform to the normal extension type
 - **File IOPS:** A surge in abnormal volume activity with data encryption (beginning in ONTAP 9.11.1)

ARP can detect the spread of most ransomware attacks after only a small number of files are encrypted, take action automatically to protect data, and alert you that a suspected attack is happening.



No ransomware detection or prevention system can completely guarantee safety from a ransomware attack. Although it's possible an attack might go undetected, ARP acts as an important additional layer of defense if anti-virus software has failed to detect an intrusion.

Learning and active modes

ARP has two modes:

- **Learning** (or "dry run" mode)
- **Active** (or "enabled" mode)

When you enable ARP, it runs in *learning mode*. In learning mode, the ONTAP system develops an alert profile based on the analytic areas: entropy, file extension types, and file IOPS. After running ARP in learning mode for enough time to assess workload characteristics, you can switch to active mode and start protecting your data. Once ARP has switched to active mode, ONTAP creates ARP Snapshot copies to protect the data if a threat is detected.

It's recommended you leave ARP in learning mode for 30 days. Beginning with ONTAP 9.13.1, ARP automatically determines the optimal learning period interval and automates the switch, which may occur before 30 days.

In active mode, if a file extension is flagged as abnormal, you should evaluate the alert. You can act on the alert to protect your data or you can mark the alert as a false positive. Marking an alert as a false positive updates the alert profile. For example, if the alert is triggered by a new file extension and you mark the alert as a false positive, you will not receive an alert the next time that file extension is observed. The command `security anti-ransomware volume workload-behavior show` shows file extensions that have been detected in the volume. (If you run this command early in learning mode and it shows an accurate representation of file types, you should not use that data as a basis to move to active mode, as ONTAP is still collecting other metrics.)

Beginning in ONTAP 9.11.1, you can customize the detection parameters for ARP. For more information, see [manage ARP attack detection parameters](#).

Threat assessment and ARP Snapshot copies

In active mode, ARP assesses threat probability based on incoming data measured against learned analytics. A measurement is assigned when ARP detects a threat:

- **Low**: the earliest detection of an abnormality in the volume (for example, a new file extension is observed in the volume).
- **Moderate**: multiple files with the same never-seen-before file extension are observed.
 - In ONTAP 9.10.1, the threshold for escalation to moderate is 100 or more files. Beginning with ONTAP 9.11.1, the file quantity is modifiable; its default value is 20.

In a low threat situation, ONTAP detects an abnormality and creates a Snapshot copy of the volume to create the best recovery point. ONTAP prepends the name of the ARP Snapshot copy with `Anti-ransomware-backup` to make it easily identifiable, for example `Anti_ransomware_backup.2022-12-20_1248`.

The threat escalates to moderate after ONTAP runs an analytics report determining if the abnormality matches a ransomware profile. Threats that remain at the low level are logged and visible in the **Events** section of System Manager. When the attack probability is moderate, ONTAP generates an EMS notification prompting you to assess the threat. ONTAP does not send alerts about low threats, however, beginning with ONTAP 9.14.1, you can [modify alerts settings](#). For more information, see [Respond to abnormal activity](#).

You can view information about a threat, regardless of level, in System Manager's **Events** section or with the `security anti-ransomware volume show` command.

ARP Snapshot copies are retained for a minimum of two days. Beginning with ONTAP 9.11.1, you can modify

the retention settings. For more information, see [Modify options for Snapshot copies](#).

How to recover data in ONTAP after a ransomware attack

When an attack is suspected, the system takes a volume Snapshot copy at that point in time and locks that copy. If the attack is confirmed later, the volume can be restored using the ARP Snapshot copy.

Locked Snapshot copies cannot be deleted by normal means. However, if you decide later to mark the attack as a false positive, the locked copy will be deleted.

With the knowledge of the affected files and the time of attack, it is possible to selectively recover the affected files from various Snapshot copies, rather than simply reverting the whole volume to one of the Snapshot copies.

ARP thus builds on proven ONTAP data protection and disaster recovery technology to respond to ransomware attacks. See the following topics for more information on recovering data.

- [Recover from Snapshot copies \(System Manager\)](#)
- [Restoring files from Snapshot copies \(CLI\)](#)
- [Smart ransomware recovery](#)

Autonomous Ransomware Protection use cases and considerations

Autonomous Ransomware Protection (ARP) is available for NAS workloads beginning with ONTAP 9.10.1. Before deploying ARP, you should be aware of the recommended uses and supported configurations as well as performance implications.

Supported and unsupported configurations

When deciding to use ARP, it's important to ensure that your volume's workload is suited to ARP and that it meets required system configurations.

Suitable workloads

ARP is suited for:

- Databases on NFS storage
- Windows or Linux home directories

Because users could create files with extensions that weren't detected in the learning period, there is a greater possibility of false positives in this workload.

- Images and video

For example, health care records and Electronic Design Automation (EDA) data

Unsuitable workloads

ARP is not suited for:

- Workloads with a high frequency of file create or delete (hundreds of thousands of files in few seconds; for example, test/development workloads).

- ARP's threat detection depends on its ability recognize an unusual surge in file create, rename, or delete activity. If the application itself is the source of the file activity, it cannot be effectively distinguished from ransomware activity.
- Workloads where the application or the host encrypts data.
ARP depends on distinguishing incoming data as encrypted or unencrypted. If the application itself is encrypting the data, then the effectiveness of the feature is reduced. However, the feature can still work based on file activity (delete, overwrite, or create, or a create or rename with a new file extension) and file type.

Supported configurations

ARP is available for NFS and SMB volumes in on-premises ONTAP systems beginning with ONTAP 9.10.1.

Support for other configurations and volume types is available in the following ONTAP versions:

	ONTAP 9.15.1	ONTAP 9.14.1	ONTAP 9.13.1	ONTAP 9.12.1	ONTAP 9.11.1	ONTAP 9.10.1
Volumes protected with SnapMirror asynchronous	✓	✓	✓	✓		
SVMs protected with SnapMirror asynchronous (SVM disaster recovery)	✓	✓	✓	✓		
SVM data mobility (vserver migrate)	✓	✓	✓	✓		
FlexGroup volumes	✓	✓	✓			
Multi-admin verification	✓	✓	✓			

SnapMirror and ARP interoperability

Beginning with ONTAP 9.12.1, ARP is supported on SnapMirror asynchronous destination volumes. ARP is **not** supported with SnapMirror Synchronous.

If a SnapMirror source volume is ARP-enabled, the SnapMirror destination volume automatically acquires the ARP configuration state (learning, enabled, etc), ARP training data, and ARP-created Snapshot of the source volume. No explicit enablement is required.

While the destination volume consists of read-only (RO) Snapshot copies, no ARP processing is done on its data. However, when the SnapMirror destination volume is converted to read-write (RW), ARP is automatically enabled on the RW-converted destination volume. The destination volume does not require any additional learning procedure besides what is already recorded on the source volume.

In ONTAP 9.10.1 and 9.11.1, SnapMirror does not transfer the ARP configuration state, training data, and Snapshot copies from source to destination volumes. Hence when the SnapMirror destination volume is

converted to RW, ARP on the destination volume must be explicitly enabled in learning mode after conversion.

ARP and virtual machines

ARP is supported with virtual machines (VMs). ARP detection behaves differently for changes inside and outside the VM. ARP is not recommended for workloads with high-entropy files inside the VM.

Changes outside the VM

ARP can detect file extension changes on an NFS volume outside of the VM if a new extension enters the volume encrypted or a file extension changes. Detectable file extension changes are:

- .vmx
- .vmxf
- .vmdk
- -flat.vmdk
- .nvram
- .vmem
- .vmsd
- .vmsn
- .vswp
- .vmss
- .log
- -#.log

Changes inside the VM

If the ransomware attack targets the VM and files inside of the VM are altered without making changes outside the VM, ARP detects the threat if the default entropy of the VM is low (for example .txt, .docx, or .mp4 files). Although ARP creates a protective Snapshot in this scenario, it does not generate a threat alert because the file extensions outside of the VM have not been tampered with.

If, by default, the files are high-entropy (for example .gzip or password-protected files), ARP's detection capabilities are limited. ARP can still take proactive Snapshots in this instance; however, no alerts will be triggered if the file extensions have not been tampered with externally.

Unsupported configurations

ARP is not supported in the following system configurations:

- ONTAP S3 environments
- SAN environments

ARP does not support the following volume configurations:

- FlexGroup volumes (in ONTAP 9.10.1 through 9.12.1. Beginning with ONTAP 9.13.1, FlexGroup volumes are supported)
- FlexCache volumes (ARP is supported on origin FlexVol volumes but not on cache volumes)
- Offline volumes

- SAN-only volumes
- SnapLock volumes
- SnapMirror synchronous
- SnapMirror asynchronous (Unsupported only in ONTAP 9.10.1 and 9.11.1. SnapMirror asynchronous is supported beginning with ONTAP 9.12.1. For more information, see [SnapMirror and ARP interoperability](#).)
- Restricted volumes
- Root volumes of storage VMs
- Volumes of stopped storage VMs

ARP performance and frequency considerations

ARP can have a minimal impact on system performance as measured in throughput and peak IOPS. The impact of the ARP feature depends on the specific volume workloads. For common workloads, the following configuration limits are recommended:

Workload characteristics	Recommended volume limit per node	Performance degradation when per-node volume limit is exceeded *
Read-intensive or the data can be compressed.	150	4% of maximum IOPS
Write-intensive and the data cannot be compressed.	60	10% of maximum IOPS

* System performance is not degraded beyond these percentages regardless of the number of volumes added in excess of the recommended limits.

Because ARP analytics run in a prioritized sequence, as the number of protected volumes increases, analytics run on each volume less frequently.

Multi-admin verification with volumes protected with ARP

Beginning with ONTAP 9.13.1, you can enable multi-admin verification (MAV) for additional security with ARP. MAV ensures that at least two or more authenticated administrators are required to turn off ARP, pause ARP, or mark a suspected attack as a false positive on a protected volume. Learn how to [enable MAV for ARP-protected volumes](#).

You need to define administrators for a MAV group and create MAV rules for the `security anti-ransomware volume disable`, `security anti-ransomware volume pause`, and `security anti-ransomware volume attack clear-suspect` ARP commands you want to protect. Each administrator in the MAV group must approve each new rule request and [add the MAV rule again](#) within MAV settings.

Beginning with ONTAP 9.14.1, ARP offers alerts for the creation of an ARP Snapshot and for the observation of a new file extension. Alerts for these events are disabled by default. Alerts can be set at the volume or SVM level. You can create MAV rules at the SVM level using `security anti-ransomware vserver event-log modify` or at the volume level with `security anti-ransomware volume event-log modify`.

Next steps

- [Enable Autonomous Ransomware Protection](#)
- [Enable MAV for ARP-protected volumes](#)

Enable Autonomous Ransomware Protection

Beginning with ONTAP 9.10.1, Autonomous Ransomware Protection (ARP) can be enabled on new or existing volumes. You first enable ARP in learning mode, in which the system analyzes the workload to characterize normal behavior. You can enable ARP on an existing volume, or you can create a new volume and enable ARP from the beginning.

About this task

You should always enable ARP initially in learning (or dry-run) mode. Beginning in active mode can lead to excessive false positive reports.

It's recommended you let ARP run in learning mode for a minimum of 30 days. Beginning with ONTAP 9.13.1, ARP automatically determines the optimal learning period interval and automates the switch, which may occur before 30 days. For more information, see [Learning and active modes](#).



In existing volumes, learning and active modes only apply to newly written data, not to already existing data in the volume. The existing data is not scanned and analyzed, because the characteristics of earlier normal data traffic are assumed based on the new data after the volume is enabled for ARP.

Before you begin

- You must have a storage VM (SVM) enabled for NFS or SMB (or both).
- The [correct license](#) must be installed for your ONTAP version.
- You must have NAS workload with clients configured.
- The volume you want to set ARP on needs to be protected and must have an active [junction path](#).
- The volume must be less than 100% full.
- It's recommended you configure the EMS system to send email notifications, which will include notices of ARP activity. For more information, see [Configure EMS events to send email notifications](#).
- Beginning in ONTAP 9.13.1, it's recommended that you enable multi-admin verification (MAV) so that two or more authenticated user admins are required for Autonomous Ransomware Protection (ARP) configuration. For more information, see [Enable multi-admin verification](#).

Enable ARP

You can enable ARP using System Manager or the ONTAP CLI.

System Manager

Steps

1. Select **Storage > Volumes**, then select the volume you want to protect.
2. In the **Security** tab of the **Volumes** overview, select **Status** to switch from Disabled to Enabled in learning-mode in the **Anti-ransomware** box.
3. When the learning period is over, switch ARP to active mode.



Beginning with ONTAP 9.13.1, ARP automatically determines the optimal learning period interval and automates the switch. You can [disable this setting on the associated storage VM](#) if you want to control the learning mode to active mode switch manually.

- a. Select **Storage > Volumes** and then select the volume that is ready for active mode.
 - b. In the **Security** tab of the **Volumes** overview, select **Switch** to active mode in the Anti-ransomware box.
4. You can verify the ARP state of the volume in the **Anti-ransomware** box.

To display ARP status for all volumes: In the **Volumes** pane, select **Show/Hide**, then ensure that **Anti-ransomware** status is checked.

CLI

The process to enable ARP with the CLI differs if you are enabling it on an existing volume versus a new volume.

Enable ARP on an existing volume

1. Modify an existing volume to enable ransomware protection in learning mode:

```
security anti-ransomware volume dry-run -volume vol_name -vserver svm_name
```

If you're running ONTAP 9.13.1 or later, adaptive learning is enabled so that the change to active state is done automatically. If you do not want this behavior to be automatically enabled, change the setting at the SVM level on all associated volumes:

```
vserver modify svm_name -anti-ransomware-auto-switch-from-learning-to-enabled false
```

2. When the learning period is over, modify the protected volume to switch to active mode if not already done automatically:

```
security anti-ransomware volume enable -volume vol_name -vserver svm_name
```

You can also switch to active mode with the modify volume command:

```
volume modify -volume vol_name -vserver svm_name -anti-ransomware-state active
```

3. Verify the ARP state of the volume.

```
security anti-ransomware volume show
```

Enable ARP on a new volume

1. Create a new volume with anti-ransomware protection enabled before provisioning data.

```
volume create -volume vol_name -vserver svm_name -aggregate aggr_name -size nn -anti-ransomware-state dry-run -junction-path /path_name
```

If you're running ONTAP 9.13.1 or later, adaptive learning is enabled so that the change to active state is done automatically. If you do not want this behavior to be automatically enabled, change the setting at the SVM level on all associated volumes:

```
vserver modify svm_name -anti-ransomware-auto-switch-from-learning-to-enabled false
```

2. When the learning period is over, modify the protected volume to switch to active mode if not already done automatically:

```
security anti-ransomware volume enable -volume vol_name -vserver svm_name
```

You can also switch to active mode with the modify volume command:

```
volume modify -volume vol_name -vserver svm_name -anti-ransomware-state active
```

3. Verify the ARP state of the volume.

```
security anti-ransomware volume show
```

Enable Autonomous Ransomware Protection by default in new volumes

Beginning with ONTAP 9.10.1, you can configure storage VMs (SVMs) such that new volumes are enabled by default for Autonomous Ransomware Protection (ARP) in learning mode.

About this task

By default, new volumes are created with ARP in disabled mode. You can modify this setting in System Manager and with the CLI. Volumes enabled by default are set to ARP in learning (or dry-run) mode.

ARP will only be enabled on volumes created in the SVM after you have changed the setting. ARP will not be enabled on existing volumes. Learn how to [enable ARP in an existing volume](#).

Beginning in ONTAP 9.13.1, adaptive learning has been added to ARP analytics, and the switch from learning mode to active mode is done automatically. For more information, see [Learning and active modes](#).

Before you begin

- The [correct license](#) must be installed for your ONTAP version.
- The volume must be less than 100% full.
- Junction paths must be active.
- Beginning in ONTAP 9.13.1, it's recommended you enable multi-admin verification (MAV) so that two or more authenticated user admins are required for anti-ransomware operations. [Learn more](#).

Switch ARP from learning to active mode

Beginning in ONTAP 9.13.1, adaptive learning has been added to ARP analytics. The switch from learning mode to active mode is done automatically. The autonomous decision by ARP to automatically switch from learning mode to active mode is based on the configuration settings of the following options:

```
-anti-ransomware-auto-switch-minimum-incoming-data-percent  
-anti-ransomware-auto-switch-duration-without-new-file-extension  
-anti-ransomware-auto-switch-minimum-learning-period  
-anti-ransomware-auto-switch-minimum-file-count  
-anti-ransomware-auto-switch-minimum-file-extension
```


After 30 days of learning, a volume is automatically switched to active mode even if one or more of these conditions are not satisfied. That is, if auto-switch is enabled, the volume switches to active mode after a maximum of 30 days. The maximum value of 30 days is fixed and not modifiable.

For more information on ARP configuration options, including default values, see the [ONTAP command reference](#).

Steps

You can use System Manager or the ONTAP CLI to enable ARP by default.

System Manager

1. Select **Storage > Storage VMs** then select the storage VM that contains volumes you want to protect with ARP.
2. Navigate to the **Settings** tab. Under **Security**, locate the **Anti-ransomware** tile then select 
3. Check the box to enable ARP for NAS volumes. Check the additional box to enable ARP on all eligible NAS volumes in the storage VM.



If you have upgraded to ONTAP 9.13.1, the **Switch automatically from learning to active mode after sufficient learning** setting is enabled automatically. This allows ARP to determine the optimal learning period interval and automate the switch to active mode. Turn off the setting if you want to manually transition to active mode.

CLI

1. Modify an existing SVM to enable ARP by default in new volumes:

```
vserver modify -vserver svm_name -anti-ransomware-default-volume-state dry-run
```

At the CLI, you can also create a new SVM with ARP enabled by default for new volumes.

```
vserver create -vserver svm_name -anti-ransomware-default-volume-state dry-run [other parameters as needed]
```

If you upgraded to ONTAP 9.13.1 or later, adaptive learning is enabled so that the change to active state is done automatically. If you do not want this behavior to be automatically enabled, use the following command:

```
vserver modify svm_name -anti-ransomware-auto-switch-from-learning-to-enabled false
```

Pause Autonomous Ransomware Protection to exclude workload events from analysis

If you are expecting unusual workload events, you can temporarily suspend and resume Autonomous Ransomware Protection (ARP) analysis at any time.

Beginning in ONTAP 9.13.1, you can enable multi-admin verification (MAV) so that two or more authenticated user admins are required to pause the ARP. [Learn more](#).

About this task

During an ARP pause, no events are logged nor are any actions for new writes. However, the analytics operation continues for earlier logs in the background.



Do not use the ARP disable function to pause analytics. Doing so disables ARP on the volume and all the existing information around learned workload behavior is lost. This would require a restart of the learning period.

Steps

You can use System Manager or the ONTAP CLI to pause ARP.

System Manager

1. Select **Storage > Volumes** and then select the volume where you want to pause ARP.
2. In the **Security** tab of the Volumes overview, select **Pause anti-ransomware** in the **Anti-ransomware** box.



Beginning with ONTAP 9.13.1, if you are using MAV to protect your ARP settings, the pause operation prompts you to obtain the approval of one or more additional administrators. [Approval must be received from all administrators](#) associated with the MAV approval group or the operation will fail.

CLI

1. Pause ARP on a volume:

```
security anti-ransomware volume pause -vserver svm_name -volume vol_name
```

2. To resume processing, use the `resume` command:

```
security anti-ransomware volume resume -vserver svm_name -volume vol_name
```

3. **If you are using MAV (available with ARP beginning with ONTAP 9.13.1) to protect your ARP settings**, the pause operation prompts you to obtain the approval of one or more additional administrators. Approval must be received from the all administrators associated with the MAV approval group or the operation will fail.

If you are using MAV and an expected pause operation needs additional approvals, each MAV group approver does the following:

- a. Show the request:

```
security multi-admin-verify request show
```

- b. Approve the request:

```
security multi-admin-verify request approve -index[number returned from show request]
```

The response for the last group approver indicates that the volume has been modified and the state of ARP is paused.

If you are using MAV and you are a MAV group approver, you can reject a pause operation request:

```
security multi-admin-verify request veto -index[number returned from show request]
```

Manage Autonomous Ransomware Protection attack detection parameters

Beginning in ONTAP 9.11.1, you can modify the parameters for ransomware detection on a specific Autonomous Ransomware Protection-enabled volume and report a known

surge as normal file activity. Adjusting detection parameters helps improve the accuracy of reporting based on your specific volume workload.

How attack detection works

When Autonomous Ransomware Protection (ARP) is in learning mode, it develops baseline values for volume behaviors. These are entropy, file extensions, and, beginning in ONTAP 9.11.1, IOPS. These baselines are used to evaluate ransomware threats. For more information about these criteria, see [What ARP detects](#).

In ONTAP 9.10.1, ARP issues a warning if it detects both of the following conditions:

- more than 20 files with file extensions not previously observed in the volume
- high entropy data

Beginning in ONTAP 9.11.1, ARP issues a threat warning if *only* one condition is met. For example, if more than 20 files with file extensions that have not previously been observed in the volume are observed within a 24 hour period, ARP will categorize this as a threat *regardless* of observed entropy. (The 24 hour and 20 file values are defaults, which can be modified.)

Beginning in ONTAP 9.14.1, you can configure alerts when ARP observes a new file extension and when ARP creates a Snapshot. For more information, see [Configure ARP alerts](#)

Certain volumes and workloads require different detection parameters. For example, your ARP-enabled volume may host numerous types of file extensions, in which case you may want to modify the threshold count for never-before-seen file extensions to a number greater than the default of 20 or disable warnings based on never-before-seen file extensions. Beginning with ONTAP 9.11.1, you can modify the attack detection parameters so they better fit your specific workloads.

Modify attack detection parameters

Depending on the expected behaviors of your ARP-enabled volume, you may want to modify the attack detection parameters.

Steps

1. View the existing attack detection parameters:

```
security anti-ransomware volume attack-detection-parameters show -vserver  
<svm_name> -volume <volume_name>
```

```

security anti-ransomware volume attack-detection-parameters show
-vserver vs1 -volume voll

Vserver Name : vs1
Volume Name : voll
Is Detection Based on High Entropy Data Rate? : true
Is Detection Based on Never Seen before File Extension? : true
Is Detection Based on File Create Rate? : true
Is Detection Based on File Rename Rate? : true
Is Detection Based on File Delete Rate? : true
Is Detection Relaxing Popular File Extensions? : true
High Entropy Data Surge Notify Percentage : 100
File Create Rate Surge Notify Percentage : 100
File Rename Rate Surge Notify Percentage : 100
File Delete Rate Surge Notify Percentage : 100
Never Seen before File Extensions Count Notify Threshold : 20
Never Seen before File Extensions Duration in Hour : 24

```

2. All of the fields shown are modifiable with boolean or integer values. To modify a field, use the `security anti-ransomware volume attack-detection-parameters modify` command.

For a full list of parameters, see [ONTAP command reference](#).

Report known surges

ARP continues to modify baseline values for detection parameters even in active mode. If you know of surges in your volume activity—either one-time surges or a surge that is characteristic of a new normal—you should report it as safe. Manually reporting these surges as safe helps to improve the accuracy of ARP's threat assessments.

Report a one-time surge

1. If a one-time surge is occurring under known circumstances and you want ARP to report a similar surge in future circumstances, clear the surge from the workload behavior:

```

security anti-ransomware volume workload-behavior clear-surge -vserver
<svm_name> -volume <volume_name>

```

Modify baseline surge

1. If a reported surge should be considered normal application behavior, report the surge as such to modify the baseline surge value.

```

security anti-ransomware volume workload-behavior update-baseline-from-surge
-vserver <svm_name> -volume <volume_name>

```

Configure ARP alerts

Beginning in ONTAP 9.14.1, ARP allows you to specify alerts for two ARP events:

- Observation of new file extension on a volume

- Creation of an ARP Snapshot

Alerts for these two events can be set on individual volumes or for the entire SVM. If you enable alerts for the SVM, the alert settings are inherited only by volumes created after you enable alert. By default, alerts are not enabled on any volume.


Event alerts can be controlled with multi-admin verification. For more information, see [Multi-admin verification with volumes protected with ARP](#).

System Manager

Set alerts for a volume

1. Navigate to **Volumes**. Select the individual volume for which you want to modify settings.
2. Select the **Security** tab then **Event Security Settings**.
3. To receive alerts for **New file extension detected** and **Ransomware snapshot created**, select the dropdown menu under the **Severity** heading. Modify the setting from **Don't generate event** to **Notice**.
4. Select **Save**.

Set alerts for an SVM

1. Navigate to **Storage VM** then select the SVM for which you want to enable settings.
2. Under the **Security** heading, locate the **Anti-ransomware** card. Select  then **Edit Ransomware Event Severity**.
3. To receive alerts for **New file extension detected** and **Ransomware snapshot created**, select the dropdown menu under the **Severity** heading. Modify the setting from **Don't generate event** to **Notice**.
4. Select **Save**.

CLI

Set alerts for a volume

- To set alerts for a new file-extension:

```
security anti-ransomware volume event-log modify -vserver <svm_name> -is
-enabled-on-new-file-extension-seen true
```

- To set alerts for the creation of an ARP Snapshot:

```
security anti-ransomware volume event-log modify -vserver <svm_name> -is
-enabled-on-snapshot-copy-creation true
```

- Confirm your settings with the `anti-ransomware volume event-log show` command.

Set alerts for an SVM

- To set alerts for a new file-extension:

```
security anti-ransomware vserver event-log modify -vserver <svm_name> -is
-enabled-on-new-file-extension-seen true
```

- To set alerts for the creation of an ARP Snapshot:

```
security anti-ransomware vserver event-log modify -vserver <svm_name> -is
-enabled-on-snapshot-copy-creation true
```

- Confirm your settings with the `security anti-ransomware vserver event-log show` command.

Related information

- [Understand Autonomous Ransomware Protection attacks and the Autonomous Ransomware Protection snapshot.](#)

Respond to abnormal activity

When Autonomous Ransomware Protection (ARP) detects abnormal activity in a protected volume, it issues a warning. You should evaluate the notification to determine whether the activity is acceptable (false positive) or whether an attack seems malicious.

About this task

ARP displays a list of suspected files when it detects any combination of high data entropy, abnormal volume activity with data encryption, and unusual file extensions.

When the warning is issued, respond by designating the file activity in one of two ways:

- **False positive**

The identified file type is expected in your workload and can be ignored.

- **Potential ransomware attack**

The identified file type is unexpected in your workload and should be treated as a potential attack.

In both cases, normal monitoring resumes after updating and clearing the notices. ARP records your evaluation to the threat assessment profile, using your choice to monitor subsequent file activities.

In the case of a suspected attack, you must determine whether it is an attack, respond to it if it is, and restore protected data before clearing the notices. [Learn more about how to recover from a ransomware attack.](#)



If you restore an entire volume, there are no notices to clear.

Before you begin

ARP must be running in active mode.

Steps

You can use System Manager or the ONTAP CLI to respond to an abnormal task.

System Manager


1. When you receive an “abnormal activity” notification, follow the link. Alternately, navigate to the **Security** tab of the **Volumes** overview.

Warnings are displayed in the **Overview** pane of the **Events** menu.

2. When a “Detected abnormal volume activity” message is displayed, view the suspect files.

In the **Security** tab, select **View Suspected File Types**.

3. In the **Suspected File Types** dialog box, examine each file type and mark it as either “False Positive” or “Potential Ransomware attack”.

If you selected this value...	Take this action...
False Positive	<p>Select Update and Clear Suspect File Types to record your decision and resume normal ARP monitoring.</p> <div data-bbox="522 835 578 886"></div> <p>Beginning with ONTAP 9.13.1, if you are using MAV to protect your ARP settings, the clear-suspect operation prompts you to obtain the approval of one or more additional administrators. Approval must be received from all administrators associated with the MAV approval group or the operation will fail.</p>
Potential Ransomware Attack	<p>Respond to the attack and restore protected data. Then select Update and Clear Suspect File Types to record your decision and resume normal ARP monitoring.</p> <p>There are no suspect file types to clear if you restored an entire volume.</p>

CLI

1. When you receive a notification of a suspected ransomware attack, verify the time and severity of the attack:

```
security anti-ransomware volume show -vserver svm_name -volume vol_name
```

Sample output:

```
Vserver Name: vs0
Volume Name: voll
State: enabled
Attack Probability: moderate
Attack Timeline: 9/14/2021 01:03:23
Number of Attacks: 1
```

You can also check EMS messages:

```
event log show -message-name callhome.arw.activity.seen
```

2. Generate an attack report and note the output location:

```
security anti-ransomware volume attack generate-report -volume vol_name  
-dest-path file_location/
```

Sample output:

```
Report "report_file_vs0_voll_14-09-2021_01-21-08" available at path  
"vs0:voll/"
```

3. View the report on an admin client system. For example:

```
[root@rhel8 mnt]# cat report_file_vs0_voll_14-09-2021_01-21-08  
  
19 "9/14/2021 01:03:23" test_dir_1/test_file_1.jpg.lckd  
20 "9/14/2021 01:03:46" test_dir_2/test_file_2.jpg.lckd  
21 "9/14/2021 01:03:46" test_dir_3/test_file_3.png.lckd`
```

4. Take one of the following actions based on your evaluation of the file extensions:

◦ False positive

Enter the following command to record your decision, adding the new extension to the list of those allowed, and resume normal anti-ransomware monitoring:

```
anti-ransomware volume attack clear-suspect -vserver svm_name -volume  
vol_name [extension identifiers] -false-positive true
```

Use one of the following parameters to identify the extensions:

```
[-seq-no integer] Sequence number of the file in the suspect list.  
[-extension text, ... ] File extensions  
[-start-time date_time -end-time date_time] Starting and ending times for the range  
of files to be cleared, in the form "MM/DD/YYYY HH:MM:SS".
```

◦ Potential ransomware attack

Respond to the attack and [recover data from the ARP-created backup snapshot](#). After the data is recovered, enter the following command to record your decision and resume normal ARP monitoring:

```
anti-ransomware volume attack clear-suspect -vserver svm_name -volume  
vol_name [extension identifiers] -false-positive false
```

Use one of the following parameters to identify the extensions:

```
[-seq-no integer] Sequence number of the file in the suspect list  
[-extension text, ... ] File extension  
[-start-time date_time -end-time date_time] Starting and ending times for the range  
of files to be cleared, in the form "MM/DD/YYYY HH:MM:SS".
```

There are no suspect file types to clear if you restored an entire volume. The ARP-created backup snapshot will be removed and the attack report will be cleared.

5. If you are using MAV and an expected `clear-suspect` operation needs additional approvals, each

MAV group approver must:

a. Show the request:

```
security multi-admin-verify request show
```

b. Approve the request to resume normal anti-ransomware monitoring:

```
security multi-admin-verify request approve -index[number returned from show request]
```

The response for the last group approver indicates that the volume has been modified and a false positive is recorded.

6. If you are using MAV and you are a MAV group approver, you can also reject a clear-suspect request:

```
security multi-admin-verify request veto -index[number returned from show request]
```

More information

- [KB: Understanding Autonomous Ransomware Protection attacks and the Autonomous Ransomware Protection snapshot.](#)

Restore data after a ransomware attack

Autonomous Ransomware Protection (ARP) creates Snapshot copies named `Anti_ransomware_backup` when it detects a potential ransomware threat. You can use one of these ARP Snapshot copies or another Snapshot copy of your volume to restore data.

About this task

If the volume has SnapMirror relationships, manually replicate all mirror copies of the volume immediately after you restore from a Snapshot copy. Not doing so can result in unusable mirror copies that must be deleted and recreated.

To restore from a Snapshot other than the `Anti_ransomware_backup` Snapshot after a system attack was identified, you must first release the ARP Snapshot.

If no system attack was reported, you must first restore from the `Anti_ransomware_backup` Snapshot copy then complete a subsequent restoration of the volume from the Snapshot copy of your choosing.

Steps

You can use System Manager or the ONTAP CLI to restore your data.

System Manager

Restore after a system attack

1. To restore from the ARP Snapshot, skip to step two. To restore from an earlier Snapshot copy, you must first release the lock on the ARP Snapshot.
 - a. Select **Storage > Volumes**.
 - b. Select **Security** then **View Suspected File Types**
 - c. Mark the files as "False Positive" .
 - d. Select **Update** and **Clear Suspect File Types**
2. Display the Snapshot copies in volumes:


Select **Storage > Volumes**, then select the volume and **Snapshot Copies**.

3. Select  next to the Snapshot copy you want to restore then **Restore**.

Restore if a system attack was not identified

1. Display the Snapshot copies in volumes:

Select **Storage > Volumes**, then select the volume and **Snapshot Copies**.

2. Select  then choose the `Anti_ransomware_backup` Snapshot.
3. Select **Restore**.
4. Return to the **Snapshot Copies** menu, then choose the Snapshot copy you want to use. Select **Restore**.

CLI

Restore after a system attack

1. To restore from the ARP Snapshot copy, skip to step two. To restore data from earlier Snapshot copies, you must release the lock on the ARP Snapshot.



It is only necessary to release the anti-ransomware Snaplock before restoring from earlier Snapshot copies if you are using the `volume snap restore` command as outlined below. If you are restoring data using Flex Clone, Single File Snap Restore or other methods, this is not necessary.

Mark the attack as a "false positive" and "clear suspect":

```
anti-ransomware volume attack clear-suspect -vserver svm_name -volume  
vol_name [extension identifiers] -false-positive true
```

Use one of the following parameters to identify the extensions:

`[-seq-no integer]` Sequence number of the file in the suspect list.

`[-extension text, ...]` File extensions

`[-start-time date_time -end-time date_time]` Starting and ending times for the range of files to be cleared, in the form "MM/DD/YYYY HH:MM:SS".

2. List the Snapshot copies in a volume:

```
volume snapshot show -vserver <SVM> -volume <volume>
```

The following example shows the Snapshot copies in `vol1`:

```
clus1::> volume snapshot show -vserver vs1 -volume vol1
```

Vserver	Volume	Snapshot	State	Size	Total%	Used%
vs1	vol1	hourly.2013-01-25_0005	valid	224KB	0%	0%
		daily.2013-01-25_0010	valid	92KB	0%	0%
		hourly.2013-01-25_0105	valid	228KB	0%	0%
		hourly.2013-01-25_0205	valid	236KB	0%	0%
		hourly.2013-01-25_0305	valid	244KB	0%	0%
		hourly.2013-01-25_0405	valid	244KB	0%	0%
		hourly.2013-01-25_0505	valid	244KB	0%	0%

7 entries were displayed.

3. Restore the contents of a volume from a Snapshot copy:

```
volume snapshot restore -vserver <SVM> -volume <volume> -snapshot  
<snapshot>
```

The following example restores the contents of `vol1`:

```
cluster1::> volume snapshot restore -vserver vs0 -volume vol1  
-snapshot daily.2013-01-25_0010
```

Restore if a system attack was not identified

1. List the Snapshot copies in a volume:

```
volume snapshot show -vserver <SVM> -volume <volume>
```

The following example shows the Snapshot copies in `vol1`:

```
clus1::> volume snapshot show -vserver vs1 -volume voll
```

Vserver	Volume	Snapshot	State	Size	Total%	Used%
vs1	voll	hourly.2013-01-25_0005	valid	224KB	0%	0%
		daily.2013-01-25_0010	valid	92KB	0%	0%
		hourly.2013-01-25_0105	valid	228KB	0%	0%
		hourly.2013-01-25_0205	valid	236KB	0%	0%
		hourly.2013-01-25_0305	valid	244KB	0%	0%
		hourly.2013-01-25_0405	valid	244KB	0%	0%
		hourly.2013-01-25_0505	valid	244KB	0%	0%

```
7 entries were displayed.
```

2. Restore the contents of a volume from a Snapshot copy:

```
volume snapshot restore -vserver <SVM> -volume <volume> -snapshot  
<snapshot>
```

The following example restores the contents of voll:

```
cluster1::> volume snapshot restore -vserver vs0 -volume voll  
-snapshot daily.2013-01-25_0010
```

3. Repeat steps 1 and 2 to restore the volume using the desire Snapshot copy.

Related information

- [KB: Ransomware prevention and recovery in ONTAP](#)

Modify options for automatic Snapshot copies

Beginning with ONTAP 9.11.1, you can use the CLI to control the retention settings for Autonomous Ransomware Protection (ARP) Snapshot copies that are automatically generated in response to suspected ransomware attacks.

Before you begin

You can only modify ARP Snapshots options on a node SVM.

Steps

1. To show all current ARP Snapshot copy settings, enter:

```
vserver options -vserver svm_name arw*
```



The `vserver options` command is a hidden command. To view the man page, enter `man vserver options` at the ONTAP CLI.


- To show selected current ARP Snapshot copy settings, enter:

```
vserver options -vserver svm_name -option-name arw_setting_name
```

- To modify ARP Snapshot copy settings, enter:

```
vserver options -vserver svm_name -option-name arw_setting_name -option-value arw_setting_value
```

The following settings are modifiable:

ARW setting	Description
<code>arw.snap.max.count</code>	Specifies the maximum number of ARP Snapshot copies that can exist in a volume at any given time. Older copies are deleted to ensure that the total number of ARP Snapshot copies are within this specified limit. The <code>-option-value</code> parameter accepts integers between 3 and 8, inclusive. The default value is 6.
<code>arw.snap.create.interval.hours</code>	Specifies the interval <i>in hours</i> between ARP Snapshot copies. A new ARP Snapshot copy is created when an data entropy-based attack is suspected and the most recently created ARP Snapshot copy is older than the specified interval. The <code>-option-value</code> parameter accepts integers between 1 and 48, inclusive. The default value is 4.
<code>arw.snap.normal.retain.interval.hours</code>	Specifies the duration <i>in hours</i> for which an ARP Snapshot copy is retained. When an ARP Snapshot copy reaches the retention threshold, any other ARP Snapshots copy created before it is deleted. No more than one ARP Snapshot copy older than the retention threshold can exist. The <code>-option-value</code> parameter accepts integers between 4 and 96, inclusive. The default value is 48.
<code>arw.snap.max.retain.interval.days</code>	Specifies the maximum duration <i>in days</i> for which an ARP Snapshot copy can be retained. Any ARP Snapshot copy older than this duration is deleted when there is no attack reported on the volume. <div style="border-left: 1px solid #ccc; padding-left: 10px; margin-left: 20px;">  <p>The maximum retention interval for ARP Snapshot copies is ignored if a moderate threat is detected. The ARP Snapshot copy created in response to the threat is retained until you have responded to the threat. Marking a threat as a false positive delete the ARP Snapshot copies on the volume. The <code>-option-value</code> parameter accepts integers between 1 and 365, inclusive. The default value is 5.</p> </div>

ARW setting	Description
<code>arw.snap.create.interval.hours.post.max.count</code>	Specifies the interval <i>in hours</i> between ARP Snapshot copies when the volume already contains the maximum number of ARP Snapshot copies. When the maximum number is reached, an ARP Snapshot copy is deleted to make room for a new copy. The new ARP Snapshot copy creation speed can be reduced to retain the older copy using this option. If the volume already contains the maximum number of ARP Snapshot copies, the interval specified in this option is used for next ARP Snapshot copy creation, instead of <code>arw.snap.create.interval.hours</code> . The <code>-option-value</code> parameter accepts integers between 4 and 48, inclusive. The default value is 8.
<code>arw.surge.snap.interval.days</code>	Specifies the interval <i>in days</i> between ARP Snapshot copies created in response to IO surges. ONTAP creates an ARP Snapshot surge copy when there's a surge in IO traffic and the last created ARP Snapshot copy is older than this specified interval. This option also specifies retention period <i>in day</i> for an ARP surge Snapshot copies. The <code>-option-value</code> parameter accepts integers between 1 and 365, inclusive. The default value is 5.
<code>arw.snap.new.extns.interval.hours</code>	This option specifies the interval <i>in hours</i> between the ARP Snapshot copies created when a new file extension is detected. A new ARP Snapshot copy is created when a new file extension is observed; the previous Snapshot created upon observing a new file extension is older than this specified interval. On a workload that frequently creates new file extensions, this interval helps in controlling the frequency of the ARP Snapshot copies. This option exists independent of <code>arw.snap.create.interval.hours</code> , which specifies the interval for data entropy-based ARP Snapshot copies. The <code>-option-value</code> parameter accepts integers between 24 and 8760. The default value is 48.

Virus protection with Vscan

Antivirus configuration overview

Vscan is an antivirus scanning solution developed by NetApp that allows customers to protect their data from being compromised by viruses or other malicious code.

Vscan performs virus scans when clients access files over SMB. You can configure Vscan to scan on-demand or on a schedule. You can interact with Vscan using the ONTAP command-line interface (CLI) or ONTAP application programming interfaces (APIs).

Related information

[Vscan partner solutions](#)

About NetApp antivirus protection

About NetApp virus scanning

Vscan is an antivirus scanning solution developed by NetApp that allows customers to protect their data from being compromised by viruses or other malicious code. It combines partner-provided antivirus software with ONTAP features to give customers the flexibility they need to manage file scanning.

How virus scanning works

Storage systems offload scanning operations to external servers hosting antivirus software from third-party vendors.

Based on the active scanning mode, ONTAP sends scan requests when clients access files over SMB (on-access) or access files in specific locations, on a schedule or immediately (on-demand).

- You can use *on-access scanning* to check for viruses when clients open, read, rename, or close files over SMB. File operations are suspended until the external server reports the scan status of the file. If the file has already been scanned, ONTAP allows the file operation. Otherwise, it requests a scan from the server.


On-access scanning is not supported for NFS.

- You can use *on-demand scanning* to check files for viruses immediately or on a schedule. We recommend that on-demand scans run only in off-peak hours to avoid overloading existing AV infrastructure, which is normally sized for on-access scanning. The external server updates the scan status of checked files, so that file-access latency is reduced over SMB. If there were file modifications or software version updates, it requests a new file scan from the external server.

You can use on-demand scanning for any path in the SVM namespace, even for volumes that are exported only through NFS.

You typically enable both on-access and on-demand scanning modes on an SVM. In either mode, the antivirus software takes remedial action on infected files based on your software settings.

The ONTAP Antivirus Connector, provided by NetApp and installed on the external server, handles communication between the storage system and the antivirus software.




Virus scanning workflow

You must create a scanner pool and apply a scanner policy before you can enable scanning. You typically enable both on-access and on-demand scanning modes on an SVM.



You must have completed the CIFS configuration.



Next steps

- [Create a scanner pool on a single cluster](#)
- [Apply a scanner policy on a single cluster](#)
- [Create an on-access policy](#)

Antivirus architecture

The NetApp antivirus architecture consists of Vscan server software and associated settings.

Vscan server software

You must install this software on the Vscan server.

- **ONTAP Antivirus Connector**

This is NetApp-provided software that handles scan request and response communication between the SVMs and antivirus software. It can run on a virtual machine, but for best performance use a physical machine. You can download this software from the NetApp Support Site (requires login).

- **Antivirus software**

This is partner-provided software that scans files for viruses or other malicious code. You specify the remedial actions to be taken on infected files when you configure the software.

Vscan software settings

You must configure these software settings on the Vscan server.

- **Scanner pool**

This setting defines the Vscan servers and privileged users that can connect to SVMs. It also defines a scan request timeout period, after which the scan request is sent to an alternative Vscan server if one is available.



You should set the timeout period in the antivirus software on the Vscan server to five seconds less than the scanner-pool scan-request timeout period. This will avoid situations in which file access is delayed or denied altogether because the timeout period on the software is greater than the timeout period for the scan request.

- **Privileged user**

This setting is a domain user account that a Vscan server uses to connect to the SVM. The account must exist in the list of privileged users in the scanner pool.

- **Scanner policy**

This setting determines whether a scanner pool is active. Scanner policies are system-defined, so you cannot create custom scanner policies. Only these three policies are available:

- `Primary` specifies that the scanner pool is active.
- `Secondary` specifies that the scanner pool is active, only when none of the Vscan servers in the primary scanner pool are connected.
- `Idle` specifies that the scanner pool is inactive.

- **On-access policy**

This setting defines the scope of an on-access scan. You can specify the maximum file size to scan, file extensions and paths to include in the scan, and file extensions and paths to exclude from the scan.

By default, only read-write volumes are scanned. You can specify filters that enable scanning of read-only volumes or that restrict scanning to files opened with execute access:

- `scan-ro-volume` enables scanning of read-only volumes.
- `scan-execute-access` restricts scanning to files opened with execute access.



“Execute access” is different from “execute permission.” A given client will have “execute access” on an executable file only if the file was opened with “execute intent.”

You can set the `scan-mandatory` option to off to specify that file access is allowed when no Vscan servers are available for virus scanning. Within on-access mode you can choose from these two mutually-exclusive options:

- **Mandatory:** With this option, Vscan tries to deliver the scan request to the server until the timeout period expires. If the scan request is not accepted by the server, then the client access request is denied.
- **Non-Mandatory:** With this option, Vscan always allows client access, whether or not a Vscan server was available for virus scanning.

• **On-demand task**

This setting defines the scope of an on-demand scan. You can specify the maximum file size to scan, file extensions and paths to include in the scan, and file extensions and paths to exclude from the scan. Files in subdirectories are scanned by default.

You use a cron schedule to specify when the task runs. You can use the `vserver vscan on-demand-task run` command to run the task immediately.

• **Vscan file-operations profile (on-access scanning only)**

The `vscan-fileop-profile` parameter for the `vserver cifs share create` command defines which SMB file operations trigger virus scanning. By default, the parameter is set to `standard`, which is NetApp best practice. You can adjust this parameter as necessary when you create or modify an SMB share:

- `no-scan` specifies that virus scans are never triggered for the share.
- `standard` specifies that virus scans are triggered by open, close, and rename operations.
- `strict` specifies that virus scans are triggered by open, read, close, and rename operations.

The `strict` profile provides enhanced security for situations in which multiple clients access a file simultaneously. If one client closes a file after writing a virus to it, and the same file remains open on a second client, `strict` ensures that a read operation on the second client triggers a scan before the file is closed.

You should be careful to restrict the `strict`` profile to shares containing files that you anticipate will be accessed simultaneously. Since this profile generates more scan requests, it may impact performance.

- `writes-only` specifies that virus scans are triggered only when modified files are closed.

Since `writes-only` generates fewer scan requests, it typically improves performance.

If you use this profile, the scanner must be configured to delete or quarantine unrepairable infected files, so they cannot be accessed. If, for example, a client closes a file after writing a virus to it, and the file is not repaired, deleted, or quarantined, any client that accesses the file `without` writing to it will be infected.



If a client application performs a rename operation, the file is closed with the new name and is not scanned. If such operations pose a security concern in your environment, you should use the `standard` or `strict` profile.

Vscan partner solutions

NetApp collaborates with Trellix, Symantec, Trend Micro, and Sentinel One to deliver industry-leading anti-malware and anti-virus solutions that build upon ONTAP Vscan technology. These solutions help you scan files for malware and remediate any affected files.

As shown in the table below, interoperability details for Trellix, Symantec and Trend Micro are maintained on the NetApp Interoperability Matrix. Interoperability details for Trellix and Symantec can also be found on the partner websites. Interoperability details for Sentinel One and other new partners will be maintained by the partner on their websites.

Partner	Solution documentation	Interoperability details
Trellix (Formerly McAfee)	Trellix Product Documentation	<ul style="list-style-type: none">• NetApp Interoperability Matrix Tool• Supported platforms for Endpoint Security Storage Protection (trellix.com)
Symantec	Symantec Protection Engine 9.0.0	<ul style="list-style-type: none">• NetApp Interoperability Matrix Tool• Support Matrix for Partner Devices Certified with Symantec Protection Engine (SPE) for Network Attached Storage (NAS) 9.x.x• Support Matrix for Partner Devices Certified with Symantec Protection Engine (SPE) for Network Attached Storage (NAS) 8.x (broadcom.com)
Trend Micro	Trend Micro ServerProtect for Storage 6.0 Getting Started Guide	NetApp Interoperability Matrix Tool
Sentinel One	<ul style="list-style-type: none">• SentinelOne Singularity Cloud Data Security• SentinelOne support <p>This link requires a user log-in. You can request access from Sentinel One.</p>	

Partner	Solution documentation	Interoperability details
Deep Instinct	<p>Deep Instinct Prevention for Storage</p> <ul style="list-style-type: none"> • Documentation and Interop <p>This link requires a user log-in. You can request access from Deep Instinct.</p> <ul style="list-style-type: none"> • Data Sheet 	

Vscan server installation and configuration

Vscan server installation and configuration

Set up one or more Vscan servers to ensure that files on your system are scanned for viruses. Follow the instructions provided by your vendor to install and configure the antivirus software on the server.

Follow the instructions in the README file provided by NetApp to install and configure the ONTAP Antivirus Connector. Alternatively, follow the instructions on the [Install ONTAP Antivirus Connector page](#).



For disaster recovery and MetroCluster configurations, you must set up and configure separate Vscan servers for the primary/local and secondary/partner ONTAP clusters.

Antivirus software requirements

- For information about antivirus software requirements, see the vendor documentation.
- For information about the vendors, software, and versions supported by Vscan, see the [Vscan partner solutions](#) page.

ONTAP Antivirus Connector requirements

- You can download the ONTAP Antivirus Connector from the **Software Download** page on the NetApp Support Site. [NetApp Downloads: Software](#)
- For information about the Windows versions supported by the ONTAP Antivirus Connector and interoperability requirements, see [Vscan partner solutions](#).



You can install different versions of Windows servers for different Vscan servers in a cluster.

- .NET 3.0 or later must be installed on the Windows server.
- SMB 2.0 must be enabled on the Windows server.

Install ONTAP Antivirus Connector

Install the ONTAP Antivirus Connector on the Vscan server to enable communication between the system running ONTAP and the Vscan server. When the ONTAP Antivirus Connector is installed, the antivirus software is able to communicate with one or more storage virtual machines (SVMs).

About this task

- See the [Vscan partner solutions](#) page for information about the supported protocols, antivirus vendor software versions, ONTAP versions, interoperability requirements and Windows servers.
- .NET 4.5.1 or later must be installed.
- The ONTAP Antivirus Connector can run on a virtual machine. However, for best performance, NetApp recommends using a dedicated virtual machine for antivirus scanning.
- SMB 2.0 must be enabled on the Windows server on which you are installing and running the ONTAP Antivirus Connector.

Before you begin

- Download the ONTAP Antivirus Connector setup file from the Support Site and save it to a directory on your hard drive.
- Verify that you meet the requirements to install the ONTAP Antivirus Connector.
- Verify that you have administrator privileges to install the Antivirus Connector.

Steps

1. Start the Antivirus Connector installation wizard by running the appropriate setup file.
2. Select *Next*. The Destination Folder dialog box opens.
3. Select *Next* to install the Antivirus Connector to the folder that is listed or select *Change* to install to a different folder.
4. The ONTAP AV Connector Windows Service Credentials dialog box opens.
5. Enter your Windows service credentials or select **Add** to select a user. For an ONTAP system, this user must be a valid domain user and must exist in the scanner pool configuration for the SVM.
6. Select **Next**. The Ready to Install the Program dialog box opens.
7. Select **Install** to begin the installation or select **Back** if you want to make any changes to the settings. A status box opens and charts the progress of the installation, followed by the InstallShield Wizard Completed dialog box.
8. Select the Configure ONTAP LIFs check box if you want to continue with the configuration of ONTAP management or data LIFs.
You must configure at least one ONTAP management or data LIF before this Vscan server can be used.
9. Select the Show the **Windows Installer log** check box if you want to view the installation logs.
10. Select **Finish** to end the installation and to close the InstallShield wizard.
The **Configure ONTAP LIFs** icon is saved on the desktop to configure the ONTAP LIFs.
11. Add an SVM to the Antivirus Connector.
You can add an SVM to the Antivirus Connector by adding either an ONTAP management LIF, which is polled to retrieve the list of data LIFs, or by directly configuring the data LIF or LIFs.
You must also provide the poll information and the ONTAP admin account credentials if the ONTAP management LIF is configured.
 - Verify that the management LIF or the IP address of the SVM is enabled for `management-https`. This is not required when you are only configuring data LIFs.
 - Verify that you have created a user account for the HTTP application and assigned a role which has (at least read-only) access to the `/api/network/ip/interfaces` REST API.
For more information about creating a user, see the [security login role create](#) and [security login create ONTAP man](#) pages.



You can also use the domain user as an account by adding an authentication tunnel SVM for an administrative SVM. For more information, see the [security login domain-tunnel create](#) ONTAP man page or use the `/api/security/accounts` and `/api/security/roles` REST APIs to configure the admin account and role.

Steps

- a. Right-click on the **Configure ONTAP LIFs** icon, which was saved on your desktop when you completed the Antivirus Connector installation, and then select **Run as Administrator**.
- b. In the Configure ONTAP LIFs dialog box, select the preferred configuration type, then perform the following actions:

To create this type of LIF...	Perform these steps...
Data LIF	<ol style="list-style-type: none"> a. Set "role" to "data" b. Set "data protocol" to "cifs" c. Set "firewall policy" to "data" d. Set "service policy" to "default-data-files"
Management LIF	<ol style="list-style-type: none"> a. Set "role*" to "data" b. Set "data protocol" to "none" c. Set "firewall policy" to "mgmt" d. Set "service policy" to "default-management"

Read more about [creating a LIF](#).

After you create a LIF, enter the data or management LIF or IP address of the SVM that you want to add. You can also enter the cluster management LIF. If you specify the cluster management LIF, all SVMs within that cluster that are serving SMB can use the Vscan server.



When Kerberos authentication is required for Vscan servers, each SVM data LIF must have a unique DNS name, and you must register that name as a server principal name (SPN) with the Windows Active Directory. When a unique DNS name is not available for each data LIF or registered as an SPN, the Vscan server uses the NT LAN Manager mechanism for authentication. If you add or modify the DNS names and SPNs after the Vscan server is connected, you must restart the Antivirus Connector service on the Vscan server to apply the changes.

- c. To configure a management LIF, enter the poll duration in seconds. The poll duration is the frequency at which the Antivirus Connector checks for changes to the SVMs or the cluster's LIF configuration. The default poll interval is 60 seconds.
- d. Enter the ONTAP admin account name and password to configure a management LIF.
- e. Click **Test** to check the connectivity and verify the authentication. Authentication is verified only for a management LIF configuration.
- f. Click **Update** to add the LIF to the list of LIFs to poll or to connect to.
- g. Click **Save** to save the connection to the registry.
- h. Click **Export** if you want to export the list of connections to a registry import or registry export file. This is

useful if multiple Vscan servers use the same set of management or data LIFs.

See the [Configure the ONTAP Antivirus Connector page](#) for configuration options.

Configure the ONTAP Antivirus Connector

Configure the ONTAP Antivirus Connector to specify one or more storage virtual machines (SVMs) that you want to connect to by either entering the ONTAP management LIF, poll information, and the ONTAP admin account credentials, or just the data LIF. You can also modify the details of an SVM connection or remove an SVM connection. By default, the ONTAP Antivirus Connector uses REST APIs to retrieve the list of data LIFs if the ONTAP management LIF is configured.

Modify the details of an SVM connection

You can update the details of a storage virtual machine (SVM) connection, which has been added to the Antivirus Connector, by modifying the ONTAP management LIF and the poll information. You cannot update data LIFs after they have been added. To update data LIFs you must first remove them and then add them again with the new LIF or IP address.

Before you begin

Verify that you have created a user account for the HTTP application and assigned a role which has (at least read-only) access to the `/api/network/ip/interfaces` REST API.

For more information about creating a user, see the [security login role create](#) and the [security login create](#) commands.

You can also use the domain user as an account by adding an authentication tunnel SVM for an administrative SVM.

For more information, see the [security login domain-tunnel create](#) ONTAP man page.

Steps

1. Right-click the **Configure ONTAP LIFs** icon, which was saved on your desktop when you completed the Antivirus Connector installation, and then select **Run as Administrator**. The Configure ONTAP LIFs dialog box opens.
2. Select the SVM IP address, and then click **Update**.
3. Update the information, as required.
4. Click **Save** to update the connection details in the registry.
5. Click **Export** if you want to export the list of connections to a registry import or a registry export file. This is useful if multiple Vscan servers use the same set of management or data LIFs.

Remove an SVM connection from the Antivirus Connector

If you no longer require an SVM connection, you can remove it.

Steps

1. Right-click the **Configure ONTAP LIFs** icon, which was saved on your desktop when you completed the Antivirus Connector installation, and then select **Run as Administrator**. The Configure ONTAP LIFs dialog box opens.
2. Select one or more SVM IP addresses, and then click **Remove**.
3. Click **Save** to update the connection details in the registry.

- Click **Export** if you want to export the list of connections to a registry import or registry export file. This is useful if multiple Vscan servers use the same set of management or data LIFs.

Troubleshoot

Before you begin

When you are creating registry values in this procedure, use the right-side pane.

You can enable or disable Antivirus Connector logs for diagnostic purposes. By default, these logs are disabled. For enhanced performance, you should keep the Antivirus Connector logs disabled and only enable them for critical events.

Steps

- Select **Start**, type "regedit" into the search box, and then select `regedit.exe` in the Programs list.
- In **Registry Editor**, locate the following subkey for the ONTAP Antivirus Connector:
`HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\Data ONTAP\Clustered Data ONTAP Antivirus Connector\v1.0`
- Create registry values by providing the type, name, and values shown in the following table:

Type	Name	Values
String	Tracepath	c:\avshim.log

This registry value could be any other valid path.

- Create another registry value by providing the type, name, values, and logging information shown in the following table:

Type	Name	Critical logging	Intermediate logging	Verbose logging
DWORD	Tracelevel	1	2 or 3	4

This enables Antivirus Connector logs that are saved at the path value provided in the TracePath in Step 3.

- Disable Antivirus Connector logs by deleting the registry values you created in Steps 3 and 4.
- Create another registry value of type "MULTI_SZ" with the name "LogRotation" (without quotes). In "LogRotation", provide "logFileSize:1" as an entry for rotation size (where 1 represents 1MB) and in the next line, provide "logFileCount:5" as an entry for rotation limit (5 is the limit).



These values are optional. If they are not provided, default values of 20MB and 10 files are used for the rotation size and rotation limit respectively. Provided integer values do not provide decimal or fraction values. If you provide values higher than the default values, the default values are used instead.

- To disable the user-configured log rotation, delete the registry values you created in Step 6.

Customizable Banner

A custom banner allows you to place a legally binding statement and a system access disclaimer on the *Configure ONTAP LIF API* window.

Step

1. Modify the default banner by updating the contents in the `banner.txt` file in the install directory and then saving the changes.
You must reopen the Configure ONTAP LIF API window to see the changes reflected in the banner.

Enable Extended Ordinance (EO) mode

You can enable and disable Extended Ordinance (EO) mode for secure operation.

Steps

1. Select **Start**, type "regedit" in the search box, and then select `regedit.exe` in the Programs list.
2. In **Registry Editor**, locate the following subkey for ONTAP Antivirus Connector:
`HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\Data ONTAP\Clustered Data ONTAP Antivirus Connector\v1.0`
3. In the right-side pane, create a new registry value of type "DWORD" with the name "EO_Mode" (without quotes) and value "1" (without quotes) to enable EO Mode or value "0" (without quotes) to disable EO Mode.



By default, if the `EO_Mode` registry entry is absent, EO mode is disabled. When you enable EO mode, you must configure both the external syslog server and mutual certificate authentication.

Configure the external syslog server

Before you begin

Take note that when you are creating registry values in this procedure, use the right-side pane.

Steps

1. Select **Start**, type "regedit" in the search box, and then select `regedit.exe` in the Programs list.
2. In **Registry Editor**, create the following subkey for ONTAP Antivirus Connector for syslog configuration:
`HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\Data ONTAP\Clustered Data ONTAP Antivirus Connector\v1.0\syslog`
3. Create a registry value by providing the type, name, and value as shown in the following table:

Type	Name	Value
DWORD	syslog_enabled	1 or 0

Please note that a "1" value enables the syslog and a "0" value disables it.

4. Create another registry value by providing the information as shown in the following table:

Type	Name
REG_SZ	Syslog_host

Provide the syslog host IP address or domain name for the value field.

5. Create another registry value by providing the information as shown in the following table:

Type	Name
REG_SZ	Syslog_port

Provide the port number on which the syslog server is running in the value field.

6. Create another registry value by providing the information as shown in the following table:

Type	Name
REG_SZ	Syslog_protocol

Enter the protocol that is in use on the syslog server, either "tcp" or "udp", in the value field.

7. Create another registry value by providing the information as shown in the following table:

Type	Name	LOG_CRIT	LOG_NOTICE	LOG_INFO	LOG_DEBUG
DWORD	Syslog_level	2	5	6	7

8. Create another registry value by providing the information as shown in the following table:

Type	Name	Value
DWORD	syslog_tls	1 or 0

Please note that a "1" value enables syslog with Transport Layer Security (TLS) and a "0" value disables syslog with TLS.

Ensure a configured external syslog server runs smoothly

- If the key is absent or has a null value:
 - The protocol defaults to "tcp".
 - The port defaults to "514" for plain "tcp/udp" and defaults to "6514" for TLS.
 - The syslog level defaults to 5 (LOG_NOTICE).
- You can confirm that syslog is enabled by verifying that the `syslog_enabled` value is "1". When the `syslog_enabled` value is "1", you should be able to log in to the configured remote server whether or not EO mode is enabled.
- If EO mode is set to "1" and you change the `syslog_enabled` value from "1" to "0", the following applies:
 - You cannot start the service if syslog is not enabled in EO mode.
 - If the system is running in a steady state, a warning appears that says syslog cannot be disabled in EO mode and syslog is forcefully set to "1", which you can see in the registry. If this occurs, you should disable EO mode first and then disable syslog.

- If the syslog server is unable to run successfully when EO mode and syslog are enabled, the service stops running. This might occur for one of the following reasons:
 - An invalid or no syslog_host is configured.
 - An invalid protocol apart from UDP or TCP is configured.
 - A port number is invalid.
- For a TCP or TLS over TCP configuration, if the server is not listening on the IP port, the connection fails and the service shuts down.

Configure X.509 mutual certificate authentication

X.509 certificate based mutual authentication is possible for the Secure Sockets Layer (SSL) communication between the Antivirus Connector and ONTAP in the management path. If EO mode is enabled and the certificate is not found, the AV Connector terminates. Perform the following procedure on the Antivirus Connector:

Steps

1. The Antivirus Connector searches for the Antivirus Connector client certificate and the certificate authority (CA) certificate for the NetApp server in the directory path from where the Antivirus Connector runs the install directory. Copy the certificates into this fixed directory path.
2. Embed the client certificate and its private key in the PKCS12 format and name it "AV_client.P12".
3. Ensure the CA certificate (along with any intermediate signing authority up to the root CA) used to sign the certificate for the NetApp server is in the Privacy Enhanced Mail (PEM) format and named "Ontap_CA.pem". Place it in the Antivirus Connector install directory. On the NetApp ONTAP system, install the CA certificate (along with any intermediate signing authority up to the root CA) used to sign the client certificate for the Antivirus Connector at "ONTAP" as a "client-ca" type certificate.

Configure scanner pools

Configure scanner pools overview

A scanner pool defines the Vscan servers and privileged users that can connect to SVMs. A scanner policy determines whether a scanner pool is active.



If you use an export policy on an SMB server, you must add each Vscan server to the export policy.

Create a scanner pool on a single cluster

A scanner pool defines the Vscan servers and privileged users that can connect to SVMs. You can create a scanner pool for an individual SVM or for all the SVMs in a cluster.

What you'll need

- SVMs and Vscan servers must be in the same domain or in trusted domains.
- For scanner pools defined for an individual SVM, you must have configured ONTAP Antivirus Connector with the SVM management LIF or SVM data LIF.
- For scanner pools defined for all the SVMs in a cluster, you must have configured ONTAP Antivirus Connector with the cluster management LIF.
- The list of privileged users must include the domain user account the Vscan server uses to connect to the

SVM.

- Once the scanner pool is configured, check the connection status to the servers.

Steps

1. Create a scanner pool:

```
vserver vscan scanner-pool create -vserver data_SVM|cluster_admin_SVM -scanner-pool scanner_pool -hostnames Vscan_server_hostnames -privileged-users privileged_users
```

- Specify a data SVM for a pool defined for an individual SVM, and specify a cluster admin SVM for a pool defined for all of the SVMs in a cluster.
- Specify an IP address or FQDN for each Vscan server host name.
- Specify the domain and user name for each privileged user.
For a complete list of options, see the man page for the command.

The following command creates a scanner pool named *SP* on the *vs1* SVM:

```
cluster1::> vserver vscan scanner-pool create -vserver vs1 -scanner-pool SP -hostnames 1.1.1.1,vmwin204-27.fsct.nb -privileged-users cifs\u1,cifs\u2
```

2. Verify that the scanner pool was created:

```
vserver vscan scanner-pool show -vserver data_SVM|cluster_admin_SVM -scanner-pool scanner_pool
```

For a complete list of options, see the man page for the command.

The following command displays the details for the *SP* scanner pool:

```
cluster1::> vserver vscan scanner-pool show -vserver vs1 -scanner-pool SP

Vserver: vs1
Scanner Pool: SP
Applied Policy: idle
Current Status: off
Cluster on Which Policy Is Applied: -
Scanner Pool Config Owner: vserver
List of IPs of Allowed Vscan Servers: 1.1.1.1, 10.72.204.27
List of Host Names of Allowed Vscan Servers: 1.1.1.1, vmwin204-27.fsct.nb
List of Privileged Users: cifs\u1, cifs\u2
```

You can also use the `vserver vscan scanner-pool show` command to view all of the scanner pools on an SVM. For complete command syntax, see the man page for the command.

Create scanner pools in MetroCluster configurations

You must create primary and secondary scanner pools on each cluster in a MetroCluster configuration, corresponding to the primary and secondary SVMs on the cluster.

What you'll need

- SVMs and Vscan servers must be in the same domain or in trusted domains.
- For scanner pools defined for an individual SVM, you must have configured ONTAP Antivirus Connector with the SVM management LIF or SVM data LIF.
- For scanner pools defined for all the SVMs in a cluster, you must have configured ONTAP Antivirus Connector with the cluster management LIF.
- The list of privileged users must include the domain user account the Vscan server uses to connect to the SVM.
- Once the scanner pool is configured, check the connection status to the servers.

About this task

MetroCluster configurations protect data by implementing two physically separate mirrored clusters. Each cluster synchronously replicates the data and SVM configuration of the other. A primary SVM on the local cluster serves data when the cluster is online. A secondary SVM on the local cluster serves data when the remote cluster is offline.

This means that you must create primary and secondary scanner pools on each cluster in a MetroCluster configuration. The secondary pool becomes active when the cluster begins serving data from the secondary SVM. For Disaster Recovery (DR) the configuration is similar to MetroCluster.


This figure shows a typical MetroCluster/DR configuration.



Site A



Site B



Steps

1. Create a scanner pool:

```
vserver vscan scanner-pool create -vserver data_SVM|cluster_admin_SVM -scanner
-pool scanner_pool -hostnames Vscan_server_hostnames -privileged-users
privileged_users
```

- Specify a data SVM for a pool defined for an individual SVM, and specify a cluster admin SVM for a pool defined for all the SVMs in a cluster.
- Specify an IP address or FQDN for each Vscan server host name.
- Specify the domain and user name for each privileged user.



You must create all scanner pools from the cluster containing the primary SVM.

For a complete list of options, see the man page for the command.

The following commands create primary and secondary scanner pools on each cluster in a MetroCluster configuration:

```
cluster1::> vserver vscan scanner-pool create -vserver cifssvm1 -
scanner-pool pool1_for_site1 -hostnames scan1 -privileged-users cifs
\u1,cifs\u2
```

```
cluster1::> vserver vscan scanner-pool create -vserver cifssvm1 -
scanner-pool pool1_for_site2 -hostnames scan1 -privileged-users cifs
\u1,cifs\u2
```

```
cluster1::> vserver vscan scanner-pool create -vserver cifssvm1 -
scanner-pool pool2_for_site1 -hostnames scan2 -privileged-users cifs
\u1,cifs\u2
```

```
cluster1::> vserver vscan scanner-pool create -vserver cifssvm1 -
scanner-pool pool2_for_site2 -hostnames scan2 -privileged-users cifs
\u1,cifs\u2
```

2. Verify that the scanner pools were created:

```
vserver vscan scanner-pool show -vserver data_SVM|cluster_admin_SVM -scanner
-pool scanner_pool
```

For a complete list of options, see the man page for the command.

The following command displays the details for the scanner pool pool1:

```
cluster1::> vserver vscan scanner-pool show -vserver cifssvm1 -scanner
-pool pool1_for_site1

                                Vserver: cifssvm1
                                Scanner Pool: pool1_for_site1
                                Applied Policy: idle
                                Current Status: off
                                Cluster on Which Policy Is Applied: -
                                Scanner Pool Config Owner: vserver
                                List of IPs of Allowed Vscan Servers:
                                List of Host Names of Allowed Vscan Servers: scan1
                                List of Privileged Users: cifs\u1,cifs\u2
```

You can also use the `vserver vscan scanner-pool show` command to view all of the scanner pools on an SVM. For complete command syntax, see the man page for the command.

Apply a scanner policy on a single cluster

A scanner policy determines whether a scanner pool is active. You must activate a scanner pool before the Vscan servers that it defines can connect to an SVM.

About this task

- You can apply only one scanner policy to a scanner pool.
- If you created a scanner pool for all the SVMs in a cluster, you must apply a scanner policy on each SVM individually.

Steps

1. Apply a scanner policy:

```
vserver vscan scanner-pool apply-policy -vserver data_SVM -scanner-pool
scanner_pool -scanner-policy primary|secondary|idle -cluster
cluster_to_apply_policy_on
```

A scanner policy can have one of the following values:

- `Primary` specifies that the scanner pool is active.
- `Secondary` specifies that the scanner pool is active only if none of the Vscan servers in the primary scanner pool are connected.
- `Idle` specifies that the scanner pool is inactive.

The following example shows that the scanner pool named `SP` on the `vs1` SVM is active:

```
cluster1::> vserver vscan scanner-pool apply-policy -vserver vs1
-scanner-pool SP -scanner-policy primary
```

2. Verify that the scanner pool is active:

```
vserver vscan scanner-pool show -vserver data_SVM|cluster_admin_SVM -scanner
-pool scanner_pool
```

For a complete list of options, see the man page for the command.

The following command displays the details for the `SP` scanner pool:

```
cluster1::> vserver vscan scanner-pool show -vserver vs1 -scanner-pool
SP

                Vserver: vs1
                Scanner Pool: SP
                Applied Policy: primary
                Current Status: on
Cluster on Which Policy Is Applied: cluster1
Scanner Pool Config Owner: vserver
List of IPs of Allowed Vscan Servers: 1.1.1.1, 10.72.204.27
List of Host Names of Allowed Vscan Servers: 1.1.1.1, vmwin204-
27.fsct.nb
                List of Privileged Users: cifs\u1, cifs\u2
```


You can use the `vserver vscan scanner-pool show-active` command to view the active scanner pools on an SVM. For the complete command syntax, see the man page for the command.

Apply scanner policies in MetroCluster configurations

A scanner policy determines whether a scanner pool is active. You must apply a scanner policy to the primary and secondary scanner pools on each cluster in a MetroCluster configuration.

About this task

- You can apply only one scanner policy to a scanner pool.
- If you created a scanner pool for all the SVMs in a cluster, you must apply a scanner policy on each SVM individually.
- For disaster recovery and MetroCluster configurations, you must apply a scanner policy to every scanner pool in the local cluster and remote cluster.
- In the policy that you create for the local cluster, you must specify the local cluster in the `cluster` parameter. In the policy that you create for the remote cluster, you must specify the remote cluster in the `cluster` parameter. The remote cluster can then take over virus scanning operations in case of a disaster.

Steps

1. Apply a scanner policy:

```
vserver vscan scanner-pool apply-policy -vserver data_SVM -scanner-pool  
scanner_pool -scanner-policy primary|secondary|idle -cluster  
cluster_to_apply_policy_on
```

A scanner policy can have one of the following values:

- `Primary` specifies that the scanner pool is active.
- `Secondary` specifies that the scanner pool is active only if none of the Vscan servers in the primary scanner pool are connected.
- `Idle` specifies that the scanner pool is inactive.



You must apply all scanner policies from the cluster containing the primary SVM.

The following commands apply scanner policies to the primary and secondary scanner pools on each cluster in a MetroCluster configuration:

```
cluster1::>vserver vscan scanner-pool apply-policy -vserver cifssvm1
-scanner-pool pool1_for_site1 -scanner-policy primary -cluster
cluster1
```

```
cluster1::>vserver vscan scanner-pool apply-policy -vserver cifssvm1
-scanner-pool pool2_for_site1 -scanner-policy secondary -cluster
cluster1
```

```
cluster1::>vserver vscan scanner-pool apply-policy -vserver cifssvm1
-scanner-pool pool2_for_site2 -scanner-policy primary -cluster
cluster2
```

```
cluster1::>vserver vscan scanner-pool apply-policy -vserver cifssvm1
-scanner-pool pool1_for_site2 -scanner-policy secondary -cluster
cluster2
```

2. Verify that the scanner pool is active:

```
vserver vscan scanner-pool show -vserver data_SVM|cluster_admin_SVM -scanner
-pool scanner_pool
```

For a complete list of options, see the man page for the command.

The following command displays the details for the scanner pool pool1:

```
cluster1::> vserver vscan scanner-pool show -vserver cifssvm1 -scanner
-pool pool1_for_site1
```

```

                                Vserver: cifssvm1
                                Scanner Pool: pool1_for_site1
                                Applied Policy: primary
                                Current Status: on
                                Cluster on Which Policy Is Applied: cluster1
                                Scanner Pool Config Owner: vserver
                                List of IPs of Allowed Vscan Servers:
                                List of Host Names of Allowed Vscan Servers: scan1
                                List of Privileged Users: cifs\u1,cifs\u2
```

You can use the `vserver vscan scanner-pool show-active` command to view the active scanner pools on an SVM. For complete command syntax, see the man page for the command.

Commands for managing scanner pools

You can modify and delete scanner pools, and manage privileged users and Vscan servers for a scanner pool. You can also view summary information about the scanner

pool.

If you want to...	Enter the following command...
Modify a scanner pool	<code>vserver vscan scanner-pool modify</code>
Delete a scanner pool	<code>vserver vscan scanner-pool delete</code>
Add privileged users to a scanner pool	<code>vserver vscan scanner-pool privileged-users add</code>
Delete privileged users from a scanner pool	<code>vserver vscan scanner-pool privileged-users remove</code>
Add Vscan servers to a scanner pool	<code>vserver vscan scanner-pool servers add</code>
Delete Vscan servers from a scanner pool	<code>vserver vscan scanner-pool servers remove</code>
View summary and details for a scanner pool	<code>vserver vscan scanner-pool show</code>
View privileged users for a scanner pool	<code>vserver vscan scanner-pool privileged-users show</code>
View Vscan servers for all scanner pools	<code>vserver vscan scanner-pool servers show</code>

For more information about these commands, see the man pages.

Configure on-access scanning

Create an on-access policy

An on-access policy defines the scope of an on-access scan. You can create an on-access policy for an individual SVM or for all the SVMs in a cluster. If you created an on-access policy for all the SVMs in a cluster, you must enable the policy on each SVM individually.

About this task

- You can specify the maximum file size to scan, file extensions and paths to include in the scan, and file extensions and paths to exclude from the scan.
- You can set the `scan-mandatory` option to off to specify that file access is allowed when no Vscan servers are available for virus scanning.
- By default, ONTAP creates an on-access policy named "default_CIFS" and enables it for all the SVMs in a cluster.
- Any file that qualifies for scan exclusion based on the `paths-to-exclude`, `file-ext-to-exclude`, or `max-file-size` parameters is not considered for scanning, even if the `scan-mandatory` option is set to

on. (Check this [troubleshooting](#) section for connectivity issues related to the `scan-mandatory` option.)

- By default, only read-write volumes are scanned. You can specify filters that enable scanning of read-only volumes or that restrict scanning to files opened with execute access.
- Virus scanning is not performed on an SMB share for which the `continuously-available` parameter is set to Yes.
- See the [Antivirus architecture](#) section for details about the *Vscan file-operations profile*.
- You can create a maximum of ten (10) on-access policies per SVM. However, you can enable only one on-access policy at a time.
 - You can exclude a maximum of one hundred (100) paths and file extensions from virus scanning in an on-access policy.
- Some file exclusion recommendations:
 - Consider excluding large files (file size can be specified) from virus scanning because they can result in a slow response or scan request timeouts for CIFS users. The default file size for exclusion is 2GB.
 - Consider excluding file extensions such as `.vhd` and `.tmp` because files with these extensions might not be appropriate for scanning.
 - Consider excluding file paths such as the quarantine directory or paths in which only virtual hard drives or databases are stored.
 - Verify that all exclusions are specified in the same policy, because only one policy can be enabled at a time. NetApp highly recommends having the same set of exclusions specified in the antivirus engine.
- An on-access policy is required for an [on-demand scan](#). To avoid on-access scanning for, you should set `-scan-files-with-no-ext` to false and `-file-ext-to-exclude` to `*` to exclude all extensions.

Steps

1. Create an on-access policy:

```
vserver vscan on-access-policy create -vserver data_SVM|cluster_admin_SVM
-policy-name policy_name -protocol CIFS -max-file-size
max_size_of_files_to_scan -filters [scan-ro-volume,][scan-execute-access]
-file-ext-to-include extensions_of_files_to_include -file-ext-to-exclude
extensions_of_files_to_exclude -scan-files-with-no-ext true|false -paths-to
-exclude paths_of_files_to_exclude -scan-mandatory on|off
```

- Specify a data SVM for a policy defined for an individual SVM, a cluster admin SVM for a policy defined for all the SVMs in a cluster.
- The `-file-ext-to-exclude` setting overrides the `-file-ext-to-include` setting.
- Set `-scan-files-with-no-ext` to true to scan files without extensions.

The following command creates an on-access policy named `Policy1` on the `vs1` SVM:

```
cluster1::> vserver vscan on-access-policy create -vserver vs1 -policy
-name Policy1 -protocol CIFS -filters scan-ro-volume -max-file-size 3GB
-file-ext-to-include "mp*", "tx*" -file-ext-to-exclude "mp3", "txt" -scan
-files-with-no-ext false -paths-to-exclude "\\vol\a b\\", "\\vol\a,b\""
```

2. Verify that the on-access policy has been created: `vserver vscan on-access-policy show -instance data_SVM|cluster_admin_SVM -policy-name name`

For a complete list of options, see the man page for the command.

The following command displays the details for the `Policy1` policy:

```
cluster1::> vserver vscan on-access-policy show -instance vs1 -policy
-name Policy1

                Vserver: vs1
                Policy: Policy1
                Policy Status: off
                Policy Config Owner: vserver
                File-Access Protocol: CIFS
                Filters: scan-ro-volume
                Mandatory Scan: on
                Max File Size Allowed for Scanning: 3GB
                File Paths Not to Scan: \vol\ a b\, \vol\ a,b\
                File Extensions Not to Scan: mp3, txt
                File Extensions to Scan: mp*, tx*
                Scan Files with No Extension: false
```

Enable an on-access policy

An on-access policy defines the scope of an on-access scan. You must enable an on-access policy on an SVM before its files can be scanned.

If you created an on-access policy for all the SVMs in a cluster, you must enable the policy on each SVM individually. You can enable only one on-access policy on an SVM at a time.

Steps

1. Enable an on-access policy:

```
vserver vscan on-access-policy enable -vserver data_SVM -policy-name
policy_name
```

The following command enables an on-access policy named `Policy1` on the `vs1` SVM:

```
cluster1::> vserver vscan on-access-policy enable -vserver vs1 -policy
-name Policy1
```

2. Verify that the on-access policy is enabled:

```
vserver vscan on-access-policy show -instance data_SVM -policy-name
policy_name
```

For a complete list of options, see the man page for the command.

The following command displays the details for the `Policy1` on-access policy:

```
cluster1::> vserver vscan on-access-policy show -instance vs1 -policy
-name Policy1
```

```
                Vserver: vs1
                Policy: Policy1
                Policy Status: on
                Policy Config Owner: vserver
                File-Access Protocol: CIFS
                Filters: scan-ro-volume
                Mandatory Scan: on
Max File Size Allowed for Scanning: 3GB
                File Paths Not to Scan: \vol\a b\, \vol\a,b\
                File Extensions Not to Scan: mp3, txt
                File Extensions to Scan: mp*, tx*
                Scan Files with No Extension: false
```

Modify the Vscan file-operations profile for an SMB share

The *Vscan file-operations profile* for an SMB share defines the operations on the share that can trigger scanning. By default, the parameter is set to `standard`. You can adjust the parameter as necessary when you create or modify an SMB share.

See the [Antivirus architecture](#) section for details about the *Vscan file-operations profile*.



Virus scanning is not performed on an SMB share that has the `continuously-available` parameter set to `Yes`.

Step

1. Modify the value of the Vscan file-operations profile for an SMB share:

```
vserver cifs share modify -vserver data_SVM -share-name share -path share_path
-vscan-fileop-profile no-scan|standard|strict|writes-only
```

For a complete list of options, see the man page for the command.

The following command changes the Vscan file operations profile for an SMB share to `strict`:

```
cluster1::> vserver cifs share modify -vserver vs1 -share-name
SALES_SHARE -path /sales -vscan-fileop-profile strict
```

Commands for managing on-access policies

You can modify, disable, or delete an on-access policy. You can view a summary and details for the policy.

If you want to...	Enter the following command...
Create an on-access policy	<code>vserver vscan on-access-policy create</code>
Modify an on-access policy	<code>vserver vscan on-access-policy modify</code>
Enable an on-access policy	<code>vserver vscan on-access-policy enable</code>
Disable an on-access policy	<code>vserver vscan on-access-policy disable</code>
Delete an on-access policy	<code>vserver vscan on-access-policy delete</code>
View summary and details for an on-access policy	<code>vserver vscan on-access-policy show</code>
Add to the list of paths to exclude	<code>vserver vscan on-access-policy paths-to-exclude add</code>
Delete from the list of paths to exclude	<code>vserver vscan on-access-policy paths-to-exclude remove</code>
View the list of paths to exclude	<code>vserver vscan on-access-policy paths-to-exclude show</code>
Add to the list of file extensions to exclude	<code>vserver vscan on-access-policy file-ext-to-exclude add</code>
Delete from the list of file extensions to exclude	<code>vserver vscan on-access-policy file-ext-to-exclude remove</code>
View the list of file extensions to exclude	<code>vserver vscan on-access-policy file-ext-to-exclude show</code>
Add to the list of file extensions to include	<code>vserver vscan on-access-policy file-ext-to-include add</code>
Delete from the list of file extensions to include	<code>vserver vscan on-access-policy file-ext-to-include remove</code>
View the list of file extensions to include	<code>vserver vscan on-access-policy file-ext-to-include show</code>

For more information about these commands, see the man pages.

Configure on-demand scanning

Configure on-demand scanning overview

You can use on-demand scanning to check files for viruses immediately or on a schedule.

You might want to run scans only in off-peak hours, for example, or you might want to scan very large files that were excluded from an on-access scan. You can use a cron schedule to specify when the task runs.

About this topic

- You can assign a schedule when you create a task.
- Only one task can be scheduled at a time on an SVM.
- On-demand scanning does not support scanning of symbolic links or stream files.



On-demand scanning does not support scanning of symbolic links or stream files.



To create an on-demand task, there must be at least one on-access policy enabled. It can be the default policy or a user created on-access policy.

Create an on-demand task

An on-demand task defines the scope of the on-demand virus scan. You can specify the maximum size of the files to be scanned, the extensions and paths of the files to be included in the scan, and the extensions and paths of the files to be excluded from the scan. Files in subdirectories are scanned by default.

About this task

- A maximum of ten (10) on-demand tasks can exist for each SVM, but only one can be active.
- An on-demand task creates a report, which has information regarding the statistics related to the scans. This report is accessible with a command or by downloading the report file created by the task at the location defined.

Before you begin

- You must have [created an on-access policy](#). The policy can be a default or user-created one. Without the on-access policy, you cannot enable the scan.

Steps

1. Create an on-demand task:

```
vserver vscan on-demand-task create -vserver data_SVM -task-name task_name
-scan-paths paths_of_files_to_scan -report-directory report_directory_path
-report-expiry-time expiration_time_for_report -schedule cron_schedule -max
-file-size max_size_of_files_to_scan -paths-to-exclude paths -file-ext-to
-exclude file_extensions -file-ext-to-include file_extensions -scan-files-with
-no-ext true|false -directory-recursion true|false
```

- The `-file-ext-to-exclude` setting overrides the `-file-ext-to-include` setting.
- Set `-scan-files-with-no-ext` to true to scan files without extensions.

For a complete list of options, see the [command reference](#).

The following command creates an on-demand task named `Task1` on the ``vs1`SVM`:

```
cluster1::> vserver vscan on-demand-task create -vserver vs1 -task
-name Task1 -scan-paths "/vol1/", "/vol2/cifs/" -report-directory
"/report" -schedule daily -max-file-size 5GB -paths-to-exclude
"/vol1/cold-files/" -file-ext-to-include "vmdk?", "mp*" -file-ext-to
-exclude "mp3", "mp4" -scan-files-with-no-ext false
[Job 126]: Vscan On-Demand job is queued. Use the "job show -id 126"
command to view the status.
```



You can use the `job show` command to view the status of the job. You can use the `job pause` and `job resume` commands to pause and restart the job, or the `job stop` command to end the job.

2. Verify that the on-demand task has been created:

```
vserver vscan on-demand-task show -instance data_SVM -task-name task_name
```

For a complete list of options, see the man page for the command.

The following command displays the details for the `Task1` task:

```
cluster1::> vserver vscan on-demand-task show -instance vs1 -task-name
Task1

                Vserver: vs1
                Task Name: Task1
                List of Scan Paths: /vol1/, /vol2/cifs/
                Report Directory Path: /report
                Job Schedule: daily
Max File Size Allowed for Scanning: 5GB
                File Paths Not to Scan: /vol1/cold-files/
                File Extensions Not to Scan: mp3, mp4
                File Extensions to Scan: vmdk?, mp*
Scan Files with No Extension: false
                Request Service Timeout: 5m
                Cross Junction: true
                Directory Recursion: true
                Scan Priority: low
                Report Log Level: info
                Expiration Time for Report: -
```

After you finish

You must enable scanning on the SVM before the task is scheduled to run.

Schedule an on-demand task

You can create a task without assigning a schedule and use the `vserver vscan on-demand-task schedule` command to assign a schedule; or add a schedule while creating the task.

About this task

The schedule assigned with the `vserver vscan on-demand-task schedule` command overrides a schedule already assigned with the `vserver vscan on-demand-task create` command.

Steps

1. Schedule an on-demand task:

```
vserver vscan on-demand-task schedule -vserver data_SVM -task-name task_name
-schedule cron_schedule
```

The following command schedules an on-access task named `Task2` on the `vs2` SVM:

```
cluster1::> vserver vscan on-demand-task schedule -vserver vs2 -task
-name Task2 -schedule daily
[Job 142]: Vscan On-Demand job is queued. Use the "job show -id 142"
command to view the status.
```

To view the status of the job, use the `job show` command. The `job pause` and `job resume` commands, respectively pause and restart the job; the `job stop` command terminates the job.

2. Verify that the on-demand task has been scheduled:

```
vserver vscan on-demand-task show -instance data_SVM -task-name task_name
```

For a complete list of options, see the man page for the command.

The following command displays the details for the `Task 2` task:

```

cluster1::> vserver vscan on-demand-task show -instance vs2 -task-name
Task2

                Vserver: vs2
                Task Name: Task2
                List of Scan Paths: /vol1/, /vol2/cifs/
                Report Directory Path: /report
                Job Schedule: daily
Max File Size Allowed for Scanning: 5GB
                File Paths Not to Scan: /vol1/cold-files/
                File Extensions Not to Scan: mp3, mp4
                File Extensions to Scan: vmdk, mp*
Scan Files with No Extension: false
                Request Service Timeout: 5m
                Cross Junction: true
                Directory Recursion: true
                Scan Priority: low
                Report Log Level: info

```

After you finish

You must enable scanning on the SVM before the task is scheduled to run.

Run an on-demand task immediately

You can run an on-demand task immediately, whether or not you have assigned a schedule.

Before you begin

You must have enabled scanning on the SVM.

Step

1. Run an on-demand task immediately:

```
vserver vscan on-demand-task run -vserver data_SVM -task-name task_name
```

The following command runs an on-access task named `Task1` on the `vs1` SVM:

```

cluster1::> vserver vscan on-demand-task run -vserver vs1 -task-name
Task1
[Job 161]: Vscan On-Demand job is queued. Use the "job show -id 161"
command to view the status.

```



You can use the `job show` command to view the status of the job. You can use the `job pause` and `job resume` commands to pause and restart the job, or the `job stop` command to end the job.

Commands for managing on-demand tasks

You can modify, delete, or unschedule an on-demand task. You can view a summary and details for the task, and manage reports for the task.

If you want to...	Enter the following command...
Create an on-demand task	<code>vserver vscan on-demand-task create</code>
Modify an on-demand task	<code>vserver vscan on-demand-task modify</code>
Delete an on-demand task	<code>vserver vscan on-demand-task delete</code>
Run an on-demand task	<code>vserver vscan on-demand-task run</code>
Schedule an on-demand task	<code>vserver vscan on-demand-task schedule</code>
Unschedule an on-demand task	<code>vserver vscan on-demand-task unschedule</code>
View summary and details for an on-demand task	<code>vserver vscan on-demand-task show</code>
View on-demand reports	<code>vserver vscan on-demand-task report show</code>
Delete on-demand reports	<code>vserver vscan on-demand-task report delete</code>

For more information about these commands, see the man pages.

Best practices for configuring the off-box antivirus functionality in ONTAP

Consider the following recommendations for configuring the off-box functionality in ONTAP.

- Restrict privileged users to virus scanning operations. Normal users should be discouraged from using privileged user credentials. This restriction can be achieved by turning off login rights for privileged users on Active Directory.
- Privileged users are not required to be part of any user group that has a large number of rights in the domain, such as the administrators group or the backup operators group. Privileged users must be validated only by the storage system so that they are allowed to create Vscan server connections and access files for virus scanning.
- Use the computers running Vscan servers only for virus scanning purposes. To discourage general use, disable the Windows terminal services and other remote access provisions on these machines, and grant the right to install new software on these machines only to administrators.
- Dedicate the Vscan servers to virus scanning and do not use them for other operations, such as backups. You might decide to run the Vscan server as a virtual machine (VM). If you run the Vscan server as a VM, make sure that the resources allocated to the VM are not shared and are enough to perform virus

scanning.

- Provide adequate CPU, memory, and disk capacity to the Vscan server to avoid over allocation of resources. Most Vscan servers are designed to use multiple CPU core servers and to distribute the load across the CPUs.
- NetApp recommends using a dedicated network with a private VLAN for the connection from the SVM to the Vscan server so that the scan traffic is not affected by other client network traffic. Create a separate network interface card (NIC) that is dedicated to the antivirus VLAN on the Vscan server and to the data LIF on the SVM. This step simplifies administration and troubleshooting if network issues arise. The antivirus traffic should be segregated using a private network. The antivirus server should be configured to communicate with the domain controller (DC) and ONTAP in one of the following ways:
 - The DC should communicate to the antivirus servers through the private network that is used to segregate the traffic.
 - The DC and antivirus server should communicate through a different network (not the private network mentioned previously), which is not the same as the CIFS client network.
 - To enable Kerberos authentication for antivirus communication, create a DNS entry for the private LIFs and a service principal name on the DC corresponding to the DNS entry created for the private LIF. Use this name when adding a LIF to the Antivirus Connector. The DNS should be able to return a unique name for each private LIF connected to the Antivirus Connector.



If the LIF for Vscan traffic is configured on a different port than the LIF for client traffic, the Vscan LIF might fail over to another node if a port failure occurs. The change makes the Vscan server not reachable from the new node and the scan notifications for file operations on the node fail. Verify that the Vscan server is reachable through at least one LIF on a node so that it can process scan requests for file operations performed on that node.

- Connect the NetApp storage system and the Vscan server by using at least a 1GbE network.
- For an environment with multiple Vscan servers, connect all servers that have similar high-performing network connections. Connecting the Vscan servers improves performance by allowing load sharing.
- For remote sites and branch offices, NetApp recommends using a local Vscan server rather than a remote Vscan server because the former is a perfect candidate for high latency. If cost is a factor, use a laptop or PC for moderate virus protection. You can schedule periodic complete file system scans by sharing the volumes or qtrees and scanning them from any system in the remote site.
- Use multiple Vscan servers to scan the data on the SVM for load-balancing and redundancy purposes. The amount of CIFS workload and resulting antivirus traffic vary per SVM. Monitor CIFS and virus-scanning latency on the storage controller. Monitor the trend of the results over time. If CIFS latency and virus-scanning latency increases due to CPU or application queues on the Vscan servers beyond trend thresholds, CIFS clients might experience long wait times. Add additional Vscan servers to distribute the load.
- Install the latest version of ONTAP Antivirus Connector.
- Keep antivirus engines and definitions up to date. Consult partners for recommendations on how often you should update.
- In a multi-tenancy environment, a scanner pool (pool of Vscan servers) can be shared with multiple SVMs provided that the Vscan servers and the SVMs are part of the same domain or trusted domain.
- The antivirus software policy for infected files should be set to "delete" or "quarantine", which is the default value set by most antivirus vendors. If the "vscan-fileop-profile" is set to "write_only", and if an infected file is found, the file remains in the share and can be opened because opening a file does not trigger a scan. The antivirus scan is triggered only after the file is closed.
- The `scan-engine timeout` value should be lesser than the `scanner-pool request-timeout`

value.

If it is set to a higher value, access to files might be delayed and might eventually time out.

To avoid this, configure the `scan-engine timeout` to 5 seconds less than the `scanner-pool request-timeout` value. Refer to the scan engine vendor's documentation for instructions on how to change the `scan-engine timeout` settings. The `scanner-pool timeout` can be changed by using the following command in advanced mode and by providing the appropriate value for the `request-timeout` parameter:

```
vserver vscan scanner-pool modify.
```

- For an environment that is sized for on-access scanning workloads and requires the use of on-demand scanning, NetApp recommends scheduling the on-demand scan job in off-peak hours to avoid additional loads on the existing antivirus infrastructure.

Learn more about best practices specific to partners at [Vscan partner solutions](#).

Enable virus scanning on an SVM

You must enable virus scanning on an SVM before an on-access or on-demand scan can run.

Steps

1. Enable virus scanning on an SVM:

```
vserver vscan enable -vserver data_SVM
```



You can use the `vserver vscan disable` command to disable virus scanning, if necessary.

The following command enables virus scanning on the `vs1` SVM:

```
cluster1::> vserver vscan enable -vserver vs1
```

2. Verify that virus scanning is enabled on the SVM:

```
vserver vscan show -vserver data_SVM
```

For a complete list of options, see the man page for the command.

The following command displays the Vscan status of the `vs1` SVM:

```
cluster1::> vserver vscan show -vserver vs1
```

```
Vserver: vs1
Vscan Status: on
```

Reset the status of scanned files

Occasionally, you might want to reset the scan status of successfully scanned files on an

SVM by using the `vserver vscan reset` command to discard the cached information for the files. You might want to use this command to restart the virus scanning processing in case of a misconfigured scan, for example.

About this task

After you run the `vserver vscan reset` command, all eligible files will be scanned the next time they are accessed.



This command can affect performance adversely, depending on the number and size of the files to be rescanned.

Before you begin

Advanced privileges are required for this task.

Steps

1. Change to advanced privilege level:

```
set -privilege advanced
```

2. Reset the status of scanned files:

```
vserver vscan reset -vserver data_SVM
```

The following command resets the status of scanned files on the `vs1` SVM:

```
cluster1::> vserver vscan reset -vserver vs1
```

View Vscan event log information

You can use the `vserver vscan show-events` command to view event log information about infected files, updates to Vscan servers, and the like. You can view event information for the cluster or for given nodes, SVMs, or Vscan servers.

Before you begin

Advanced privileges are required to view the Vscan event log.

Steps

1. Change to advanced privilege level:

```
set -privilege advanced
```

2. View Vscan event log information:

```
vserver vscan show-events
```

For a complete list of options, see the man page for the command.

The following command displays event log information for the cluster `cluster1`:

```
cluster1::*> vserver vscan show-events
```

Vserver	Node	Server	Event Type	Event Time
vs1 11:37:38	Cluster-01	192.168.1.1	file-infected	9/5/2014
vs1 11:37:08	Cluster-01	192.168.1.1	scanner-updated	9/5/2014
vs1 11:34:55	Cluster-01	192.168.1.1	scanner-connected	9/5/2014

3 entries were displayed.

Monitor and troubleshoot connectivity issues

Potential connectivity issues involving the scan-mandatory option

You can use the `vserver vscan connection-status show` commands to view information about Vscan server connections that you might find helpful in troubleshooting connectivity issues.

By default, the `scan-mandatory` option for on-access scanning denies file access when a Vscan server connection is not available for scanning. Although this option offers important safety features, it can lead to problems in a few situations.

- Before enabling client access, you must ensure that at least one Vscan server is connected to an SVM on each node that has a LIF. If you need to connect servers to SVMs after enabling client access, you must turn off the `scan-mandatory` option on the SVM to ensure that file access is not denied because a Vscan server connection is not available. You can turn the option back on after the server has been connected.
- If a target LIF hosts all the Vscan server connections for an SVM, the connection between the server and the SVM will be lost if the LIF is migrated. To ensure that file access is not denied because a Vscan server connection is not available, you must turn off the `scan-mandatory` option before migrating the LIF. You can turn the option back on after the LIF has been migrated.

Each SVM should have at least two Vscan servers assigned to it. It is a best practice to connect Vscan servers to the storage system over a different network from the one used for client access.

Commands for viewing Vscan server connection status

You can use the `vserver vscan connection-status show` commands to view summary and detailed information about Vscan server connection status.

If you want to...	Enter the following command...
View a summary of Vscan server connections	<code>vserver vscan connection-status show</code>

If you want to...	Enter the following command...
View details for Vscan server connections	<code>vserver vscan connection-status show-all</code>
View details for connected Vscan servers	<code>vserver vscan connection-status show-connected</code>
View details for available Vscan servers that are not connected	<code>vserver vscan connection-status show-not-connected</code>

For more information about these commands, see the [ONTAP man pages](#).

Troubleshoot virus scanning

For common virus scanning issues, there are possible causes and ways to resolve them. Virus scanning is also known as Vscan.

Issue	How to resolve it
The Vscan servers are not able to connect to the clustered ONTAP storage system.	Check whether the scanner pool configuration specifies the Vscan server IP address. Check also if the allowed privileged users in the scanner pool list are active. To check the scanner pool, run the <code>vserver vscan scanner-pool show</code> command on the storage system command prompt. If the Vscan servers still cannot connect, there might be an issue with the network.
Clients observe high latency.	It is probably time to add more Vscan servers to the scanner pool.
Too many scans are triggered.	Modify the value of the <code>vscan-fileop-profile</code> parameter to restrict the number of file operations monitored for virus scanning.
Some files are not being scanned.	Check the on-access policy. It is possible that the path for these files has been added to the path-exclusion list or that their size exceeds the configured value for exclusions. To check the on-access policy, run the <code>vserver vscan on-access-policy show</code> command on the storage system command prompt.
File access is denied.	Check whether the <code>scan-mandatory</code> setting is specified in the policy configuration. This setting denies data access if no Vscan servers are connected. Modify the setting as needed.

Monitor status and performance activities

You can monitor the critical aspects of the Vscan module, such as the Vscan server connection status, the health of the Vscan servers, and the number of files that have been scanned. This information helps you diagnose issues related to the Vscan server.

View Vscan server connection information

You can view the connection status of Vscan servers to manage the connections that are already in use and the connections that are available for use. Various commands display information about the connection status of Vscan servers.

Command...	Information displayed...
<code>vserver vscan connection-status show</code>	Summary of the connection status
<code>vserver vscan connection-status show-all</code>	Detailed information about the connection status
<code>vserver vscan connection-status show-not-connected</code>	Status of the connections that are available but not connected
<code>vserver vscan connection-status show-connected</code>	Information about the connected Vscan server

For more information about these commands, see the [ONTAP command reference](#).

View Vscan server statistics

You can view Vscan server-specific statistics to monitor performance and diagnose issues related to virus scanning. You must collect a data sample before you can use the `statistics show` command to display the Vscan server statistics.

To complete a data sample, complete the following step:

Step

1. Run the `statistics start` command and the optional `statistics stop` command.

View statistics for Vscan server requests and latencies

You can use ONTAP `offbox_vscan` counters on a per-SVM basis to monitor the rate of Vscan server requests that are dispatched and received per second and the server latencies across all Vscan servers. To view these statistics, complete the following step:

Step

1. Run the `statistics show object offbox_vscan -instance SVM` command with the following counters:

Counter...	Information displayed...
scan_request_dispatched_rate	Number of virus-scanning requests sent from ONTAP to the Vscan servers per second
scan_noti_received_rate	Number of virus-scanning requests received back by ONTAP from the Vscan servers per second
dispatch_latency	Latency within ONTAP to identify an available Vscan server and send the request to that Vscan server
scan_latency	Round-trip latency from ONTAP to the Vscan server, including the time for the scan to run

Example of statistics generated from an ONTAP offbox vscan counter

```

Object: offbox_vscan
Instance: SVM
Start-time: 10/16/2013 10:13:25
End-time: 10/16/2013 10:25:11
Cluster: cluster01
Number of Constituents: 2 (complete_aggregation)
Counter Value
-----
scan_request_dispatched_rate 291
scan_noti_received_rate 292
dispatch_latency 43986us
scan_latency 3433501us
-----

```

View statistics for individual Vscan server requests and latencies

You can use ONTAP `offbox_vscan_server` counters on a per-SVM, per-off-box Vscan server, and per-node basis to monitor the rate of dispatched Vscan server requests and the server latency on each Vscan server individually. To collect this information, complete the following step:

Step

1. Run the `statistics show -object offbox_vscan -instance SVM:servername:nodename` command with the following counters:

Counter...	Information displayed...
scan_request_dispatched_rate	Number of virus-scanning requests sent from ONTAP

scan_latency	Round-trip latency from ONTAP to the Vscan server, including the time for the scan to run to the Vscan servers per second
--------------	---------------------------------------------------------------------------------------------------------------------------

Example of statistics generated from an ONTAP offbox_vscan_server counter

```
Object: offbox_vscan_server
Instance: SVM:vscan_server:node
Start-time: 10/16/2013 10:13:25
End-time: 10/16/2013 10:25:11
Cluster: cluster01
Number of Constituents: 1 (complete_aggregation)
Counter Value
-----
scan_request_dispatched_rate 291
scan_latency 3433830us
-----
```

View statistics for Vscan server utilization

You can also use ONTAP `offbox_vscan_server` counters to collect Vscan server-side utilization statistics. These statistics are tracked on a per-SVM, per-off-box Vscan server, and per-node basis. They include CPU utilization on the Vscan server, queue depth for scanning operations on the Vscan server (both current and maximum), used memory and used network.

These statistics are forwarded by the Antivirus Connector to the statistics counters within ONTAP. They are based on data that is polled every 20 seconds and must be collected multiple times for accuracy; otherwise, the values seen in the statistics reflect only the last polling. CPU utilization and queues are particularly important to monitor and analyze. A high value for an average queue can indicate that the Vscan server has a bottleneck.

To collect utilization statistics for the Vscan server on a per-SVM, per-off-box Vscan server, and per-node basis, complete the following step:

Step

1. Collect utilization statistics for the Vscan server

Run the `statistics show -object offbox_vscan_server -instance SVM:servername:nodename` command with the following `offbox_vscan_server` counters:

Counter...	Information displayed...
<code>scanner_stats_pct_cpu_used</code>	CPU utilization on the Vscan server
<code>scanner_stats_pct_input_queue_avg</code>	Average queue of scan requests on the Vscan server
<code>scanner_stats_pct_input_queue_hiwatermark</code>	Peak queue of scan requests on the Vscan server

scanner_stats_pct_mem_used	Memory used on the Vscan server
scanner_stats_pct_network_used	Network used on the Vscan server

Example of utilization statistics for the Vscan server

```
Object: offbox_vscan_server
Instance: SVM:vscan_server:node
Start-time: 10/16/2013 10:13:25
End-time: 10/16/2013 10:25:11
Cluster: cluster01
Number of Constituents: 1 (complete_aggregation)
Counter Value
```

```
-----
scanner_stats_pct_cpu_used 51
scanner_stats_pct_dropped_requests 0
scanner_stats_pct_input_queue_avg 91
scanner_stats_pct_input_queue_hiwatermark 100
scanner_stats_pct_mem_used 95
scanner_stats_pct_network_used 4
-----
```

ONTAP hardening guidelines

ONTAP security hardening overview

ONTAP provides a set of controls that allow you to harden the ONTAP storage operating system, the industry’s leading data management software. Use the guidance and configuration settings for ONTAP to help your organization meet prescribed security objectives for information system confidentiality, integrity, and availability.

The evolution of the current threat landscape presents an organization with unique challenges for protecting its most valuable assets: data and information. The advanced and dynamic threats and vulnerabilities we face are ever increasing in sophistication. Coupled with an increase in the effectiveness of obfuscation and reconnaissance techniques on the part of potential intruders, system managers must address the security of data and information proactively.



Beginning in July 2024, content from technical reports previously published as PDFs has been integrated with ONTAP product documentation. The ONTAP security documentation now includes content from *TR-4569: Security hardening guide for ONTAP*.

ONTAP image validation

ONTAP provides mechanisms to ensure the ONTAP image is valid at upgrade and at boot time.

Upgrade image validation

Code signing helps verify that ONTAP images installed through nondisruptive image updates or automated nondisruptive image updates, CLIs, or ONTAP APIs are authentically produced by NetApp and have not been tampered with. Upgrade image validation was introduced in ONTAP 9.3.

This feature is a no-touch security enhancement to ONTAP upgrading or reversion. The user is not expected to do anything differently except for optionally verifying the top-level `image.tgz` signature.

Boot-time image validation

Beginning with ONTAP 9.4, Unified Extensible Firmware Interface (UEFI) secure boot is enabled for NetApp AFF A800, AFF A220, FAS2750, and FAS2720 systems and subsequent next-generation systems that employ UEFI BIOS.

During power on, the bootloader validates the whitelist database of secure boot keys with the signature associated with each module that is loaded. After each module is validated and loaded, the boot process continues with the ONTAP initialization. If signature validation fails for any module, the system reboots.



These items apply to ONTAP images and the platform BIOS.

Local storage administrator accounts

Roles, applications, and authentication

ONTAP provides the security-conscious enterprise with the ability to provide granular access to different administrators through different login applications and methods. This helps customers create a data centric zero-trust model.

These are the roles available for admin and storage virtual machine administrators. The login application methods and login authentication methods are specified.

Roles

With role-based access control (RBAC), users have access to only the systems and options required for their job roles and functions. The RBAC solution in ONTAP limits users' administrative access to the level granted for their defined role, which allows administrators to manage users by assigned role. ONTAP provides several predefined roles. Operators and administrators can create, modify, or delete custom access control roles, and they can specify account restrictions for specific roles.

Predefined roles for cluster administrators

This role...	Has this level of access...	To the following commands or command directories
admin	All	All command directories (DEFAULT)

admin-no-fsa (available beginning in ONTAP 9.12.1)	Read/Write	<ul style="list-style-type: none"> • All command directories (DEFAULT) • security login rest-role • security login role
	Read only	<ul style="list-style-type: none"> • security login rest-role create • security login rest-role delete • security login rest-role modify • security login rest-role show • security login role create • security login role create • security login role delete • security login role modify • security login role show • volume activity-tracking • volume analytics
	None	volume file show-disk-usage
autosupport	All	<ul style="list-style-type: none"> • set • system node autosupport
	None	All other command directories (DEFAULT)

backup	All	vserver services ndmp
	Read only	volume
	None	All other command directories (DEFAULT)
readonly	All	<ul style="list-style-type: none"> • security login password <p>For managing own user account local password and key information only</p> <ul style="list-style-type: none"> • set
	None	security
	Read only	All other command directories (DEFAULT)
none	None	All command directories (DEFAULT)



The `autosupport` role is assigned to the predefined `autosupport` account, which is used by AutoSupport OnDemand. ONTAP prevents you from modifying or deleting the `autosupport` account. ONTAP also prevents you from assigning the `autosupport` role to other user accounts.

Predefined roles for storage virtual machine (SVM) administrators

Role name	Capabilities
-----------	--------------

vsadmin	<ul style="list-style-type: none"> • Manage own user account local password and key information • Manage volumes, except volume moves • Manage quotas, qtrees, Snapshot copies, and files • Manage LUNs • Perform SnapLock operations, except privileged delete • Configure protocols: NFS, SMB, iSCSI, FC, FCoE, NVMe/FC and NVMe/TCP • Configure services: DNS, LDAP, and NIS • Monitor jobs • Monitor network connections and network interface • Monitor the health of the SVM
vsadmin-volume	<ul style="list-style-type: none"> • Manage own user account local password and key information • Manage volumes, including volume moves • Manage quotas, qtrees, Snapshot copies, and files • Manage LUNs • Configure protocols: NFS, SMB, iSCSI, FC, FCoE, NVMe/FC and NVMe/TCP • Configure services: DNS, LDAP, and NIS • Monitor network interface • Monitor the health of the SVM
vsadmin-protocol	<ul style="list-style-type: none"> • Manage own user account local password and key information • Configure protocols: NFS, SMB, iSCSI, FC, FCoE, NVMe/FC and NVMe/TCP • Configure services: DNS, LDAP, and NIS • Manage LUNs • Monitor network interface • Monitor the health of the SVM

vsadmin-backup	<ul style="list-style-type: none"> • Manage own user account local password and key information • Manage NDMP operations • Make a restored volume read/write • Manage SnapMirror relationships and Snapshot copies • View volumes and network information
vsadmin-snaplock	<ul style="list-style-type: none"> • Manage own user account local password and key information • Manage volumes, except volume moves • Manage quotas, qtrees, Snapshot copies, and files • Perform SnapLock operations, including privileged delete • Configure protocols: NFS and SMB • Configure services: DNS, LDAP, and NIS • Monitor jobs • Monitor network connections and network interface
vsadmin-readonly	<ul style="list-style-type: none"> • Manage own user account local password and key information • Monitor the health of the SVM • Monitor network interface • View volumes and LUNs • View services and protocols

Application methods

The application method specifies the access type of the login method. Possible values include `console`, `http`, `ontapi`, `rsh`, `snmp`, `service-processor`, `ssh`, and `telnet`.

Setting this parameter to `service-processor` grants the user access to the Service Processor. When this parameter is set to `service-processor`, the `-authentication-method` parameter must be set to `password` because the Service Processor only supports `password` authentication. SVM user accounts cannot access the Service Processor. Therefore, operators and administrators cannot use the `-vserver` parameter when this parameter is set to `service-processor`.

To further restrict access to the `service-processor` use the command `system service-processor ssh add-allowed-addresses`. The command `system service-processor api-service` can be used to update the configurations and certificates.

For security reasons, Telnet and Remote Shell (RSH) are disabled by default because NetApp recommends Secure Shell (SSH) for secure remote access. If there is a requirement or unique need for Telnet or RSH, they

must be enabled.

The `security protocol modify` command modifies the existing cluster-wide configuration of RSH and Telnet. Enable RSH and Telnet in the cluster by setting the `enabled` field to `true`.

Authentication methods

The authentication method parameter specifies the authentication method used for logins.

Authentication method	Description
<code>cert</code>	SSL certificate authentication
<code>community</code>	SNMP community strings
<code>domain</code>	Active Directory authentication
<code>nsswitch</code>	LDAP or NIS authentication
<code>password</code>	Password
<code>publickey</code>	Public key authentication
<code>usm</code>	SNMP user security model



The use of NIS is not recommended due to protocol security weaknesses.

Beginning with ONTAP 9.3, chained two-factor authentication is available for local SSH admin accounts using `publickey` and `password` as the two authentication methods. In addition to the `-authentication-method` field in the `security login` command, a new field named `-second-authentication-method` has been added. Either `publickey` or `password` can be specified as the `-authentication-method` or the `-second-authentication-method`. However, during SSH authentication, the order is always `publickey` with partial authentication, followed by the password prompt for full authentication.

```
[user@host01 ~]$ ssh ontap.netapp.local
Authenticated with partial success.
Password:
cluster1::>
```

Beginning with ONTAP 9.4, `nsswitch` can be used as a second authentication method with `publickey`.

Beginning with ONTAP 9.12.1, FIDO2 can also be used for SSH authentication using a YubiKey hardware authentication device or other FIDO2 compatible devices.

Beginning with ONTAP 9.13.1:

- `domain` accounts can be used as a second authentication method with `publickey`.
- Time-based one-time password (`totp`) is a temporary passcode generated by an algorithm that uses the current time of day as one of its authentication factors for the second authentication method.
- Public key revocation is supported with SSH publickeys as well as certificates which will be checked for expiration/revocation during SSH.

For more information about multifactor authentication (MFA) for ONTAP System Manager, Active IQ Unified

Manager, and SSH, see [TR-4647: Multifactor Authentication in ONTAP 9](#).

Default administrative accounts

The admin account should be restricted because the role of administrator is allowed access using all applications. The diag account allows access to the system shell and should be reserved only for technical support to perform troubleshooting tasks.

There are two default administrative accounts: `admin` and `diag`.

Orphaned accounts are a major security vector that often leads to vulnerabilities, including the escalation of privileges. These are unnecessary and unused accounts that remain in the user account repository. They are primarily default accounts that were never used or for which passwords were never updated or changed. To address this issue, ONTAP supports the removal and renaming of accounts.



ONTAP cannot remove or rename built-in accounts. However, NetApp recommends locking any unneeded built-in accounts with the `lock` command.

Although orphaned accounts are a significant security issue, NetApp strongly recommends testing the effect of removing accounts from the local account repository.

List local accounts

To list the local accounts, run the `security login show` command.

```
cluster1::*> security login show -vserver cluster1

Vserver: cluster1

          Authentication
User/Group Name  Application Method   Role Name   Acct   Is-Nsswitch
                  Locked   Group
-----
admin            console  password  admin     no     no
admin            http     password  admin     no     no
admin            ontapi   password  admin     no     no
admin            service-processor password  admin     no     no
admin            ssh      password  admin     no     no
autosupport      console  password  autosupport no     no
6 entries were displayed.
```

Remove the default admin account

The `admin` account has the role of administrator and is allowed access using all applications.

Steps

1. Create another admin-level account.

To completely remove the default `admin` account, you must first create another admin-level account that uses the `console` login application.



Making these changes might cause some undesired effects. Always test new settings that might affect the security status of the solution on a nonproduction cluster first.

Example:

```
cluster1::*> security login create -user-or-group-name NewAdmin
-application console -authentication-method password -vserver cluster1
```

```
cluster1::*> security login show -vserver cluster1
```

```
Vserver: cluster1
```

		Authentication		Acct	Is-
User/Group Name	Application	Method	Role Name	Locked	Group
-----	-----	-----	-----	-----	
NewAdmin	console	password	admin	no	no
admin	console	password	admin	no	no
admin	http	password	admin	no	no
admin	ontapi	password	admin	no	no
admin	service-processor	password	admin	no	no
admin	ssh	password	admin	no	no
autosupport	console	password	autosupport	no	no

7 entries were displayed.

2. After you create the new admin account, test access to that account with the `NewAdmin` account login. With the `NewAdmin` login, configure the account to have to same login applications as the default or previous admin account (for example, `http`, `ontapi`, `service-processor`, or `ssh`). This step makes sure that access control is maintained.

Example:

```
cluster1::*> security login create -vserver cluster1 -user-or-group-name
NewAdmin -application ssh -authentication-method password
cluster1::*> security login create -vserver cluster1 -user-or-group-name
NewAdmin -application http -authentication-method password
cluster1::*> security login create -vserver cluster1 -user-or-group-name
NewAdmin -application ontapi -authentication-method password
cluster1::*> security login create -vserver cluster1 -user-or-group-name
NewAdmin -application service-processor -authentication-method password
```

3. After all functions have been tested, you can disable the admin account for all applications before removing it from ONTAP. This step serves as a final test to confirm that there are no lingering functions that rely on the previous admin account.

```
cluster1::*> security login lock -vserver cluster1 -user-or-group-name
admin -application *
```

4. To remove the default admin account and all entries for it, run the following command:

```
cluster1::*> security login delete -vserver cluster1 -user-or-group-name
admin -application *
cluster1::*> security login show -vserver cluster1

Vserver: cluster1

                                Authentication                Acct   Is-
Nsswitch
User/Group Name  Application Method   Role Name                Locked Group
-----
-----
NewAdmin         console   password admin                    no     no
NewAdmin         http     password admin                    no     no
NewAdmin         ontapi   password admin                    no     no
NewAdmin         service-processor password admin                no     no
NewAdmin         ssh     password admin                    no     no
autosupport     console   password autosupport        no     no
7 entries were displayed.
```

Set the diagnostic (diag) account password

A diagnostic account named `diag` is provided with your storage system. You can use the `diag` account to perform troubleshooting tasks in the `systemshell`. The `diag` account is the only account that can be used to access the `systemshell` through the `diag` privileged command `systemshell`.



The `systemshell` and the associated `diag` account are intended for low-level diagnostic purposes. Their access requires the diagnostic privilege level and is reserved only to be used with guidance from technical support to perform troubleshooting tasks. Neither the `diag` account nor the `systemshell` is intended for general administrative purposes.

Before you begin

Before accessing the `systemshell`, you must set the `diag` account password by using the `security login password` command. You should use strong password principles and change the `diag` password at regular intervals.

Steps

1. Set the `diag` account user password:

```
cluster1::> set -privilege diag

Warning: These diagnostic commands are for use by NetApp personnel only.
Do you want to continue? \{y|n}: y

cluster1::*> systemshell -node node-01
      (system node systemshell)
diag@node-01's password:

Warning: The system shell provides access to low-level
diagnostic tools that can cause irreparable damage to
the system if not used properly. Use this environment
only when directed to do so by support personnel.

node-01%
```

Multi-admin verification

Beginning with ONTAP 9.11.1, you can use multi-admin verification (MAV) to allow certain operations, such as deleting volumes or Snapshot copies, to be executed only after approvals from designated administrators. This prevents compromised, malicious, or inexperienced administrators from making undesirable changes or deleting data.

Configuring MAV consists of the following:

- [Creating one or more administrator approval groups.](#)
- [Enabling multi-admin verification functionality.](#)
- [Adding or modifying rules.](#)

After initial configuration, only administrators in a MAV approval group (MAV administrators) can modify these elements.

When MAV is enabled, the completion of every protected operation requires three steps:

1. When a user initiates the operation, a [request is generated](#).
2. Before it can be executed, the required number of [MAV administrators must approve](#).
3. After approval, the user completes the operation.

MAV is not intended for use with volumes or workflows that involve heavy automation because each automated task requires approval before the operation can be completed. If you want to use automation and MAV together, NetApp recommends that you use queries for specific MAV operations. For example, you can apply `volume delete` MAV rules only to volumes where automation is not involved, and you can designate those volumes with a particular naming scheme.

For more detailed information about MAV, see the [ONTAP multi-admin verification documentation](#).

Snapshot copy locking

Snapshot copy locking is a SnapLock capability where Snapshot copies are rendered indelible manually or automatically with a retention period on the volume Snapshot policy. The purpose of Snapshot copy locking is to prevent rogue or untrusted administrators from deleting Snapshots on primary or secondary ONTAP system.

Snapshot copy locking was introduced in ONTAP 9.12.1. Snapshot copy locking is also referred to as tamper-proof Snapshot locking. Although it does require the SnapLock license and initialization of the compliance clock, Snapshot copy locking is unrelated to SnapLock Compliance or SnapLock Enterprise. There is no trusted storage administrator, as with SnapLock Enterprise and it does not protect the underlying physical storage infrastructure, as with SnapLock Compliance. This is an improvement over SnapVaulting Snapshot copies to a secondary system. Rapid recovery of locked Snapshots on primary systems can be achieved to restore volumes corrupted by ransomware.

For more details on Snapshot copy locking, see the [ONTAP documentation](#).

Set up certificate-based API access

Instead of user ID and password authentication for REST API or NetApp Manageability SDK API access to ONTAP, certificate-based authentication must be used.



As an alternative to certificate-based authentication for REST API, use [OAuth 2.0 token-based authentication](#).)

You can generate and install a self-signed certificate on ONTAP as described in these steps.

Steps

1. Using OpenSSL, generate a certificate by running the following command:

```
openssl req -x509 -nodes -days 1095 -newkey rsa:2048 -keyout test.key
-out test.pem \> -subj "/C=US/ST=NC/L=RTP/O=NetApp/CN=cert_user"
Generating a 2048 bit RSA private key
.....+++
.....+++
writing new private key to 'test.key'
```

This command generates a public certificate named `test.pem` and a private key named `key.out`. The common name, CN, corresponds to the ONTAP user ID.

2. Install the contents of the public certificate in privacy enhanced mail (pem) format in ONTAP by running the following command and pasting the certificate's contents when prompted:

```
security certificate install -type client-ca -vserver cluster1

Please enter Certificate: Press <Enter> when done
```

3. Enable ONTAP to allow client access through SSL and define the user ID for API access.


```
security ssl modify -vserver cluster1 -client-enabled true
security login create -user-or-group-name cert_user -application ontapi
-authmethod cert -role admin -vserver cluster1
```

In the following example, the user ID `cert_user` is now enabled to use certificate-authenticated API access. A simple Manageability SDK Python script using `cert_user` to display the ONTAP version appears as follows:

```
#!/usr/bin/python

import sys
sys.path.append("/home/admin/netapp-manageability-sdk-9.5/netapp-
manageability-sdk-9.5/lib/python/NetApp")
from NaServer import *

cluster = "cluster1"
transport = "HTTPS"
port = 443
style = "CERTIFICATE"
cert = "test.pem"
key = "test.key"

s = NaServer(cluster, 1, 30)
s.set_transport_type(transport)
s.set_port(port)
s.set_style(style)
s.set_server_cert_verification(0)
s.set_client_cert_and_key(cert, key)

api = NaElement("system-get-version")
output = s.invoke_elem(api)
if (output.results_status() == "failed"):
    r = output.results_reason()
    print("Failed: " + str(r))
    sys.exit(2)

ontap_version = output.child_get_string("version")
print ("V: " + ontap_version)
```

The output of the script displays the ONTAP version.

```
./version.py
```

```
V: NetApp Release 9.5RC1: Sat Nov 10 05:13:42 UTC 2018
```

4. To perform certificate-based authentication with the ONTAP REST API, complete the following steps:
 - a. In ONTAP, define the user ID for http access:

```
security login create -user-or-group-name cert_user -application http  
-authmethod cert -role admin -vserver cluster1
```

- b. On your Linux client, run the following command that produces the ONTAP version as output:

```
curl -k --cert-type PEM --cert ./test.pem --key-type PEM --key  
./test.key -X GET "https://cluster1/api/cluster?fields=version"  
{  
  "version": {  
    "full": "NetApp Release 9.7P1: Thu Feb 27 01:25:24 UTC 2020",  
    "generation": 9,  
    "major": 7,  
    "minor": 0  
  },  
  "_links": {  
    "self": {  
      "href": "/api/cluster"  
    }  
  }  
}
```

More information

- [Certificate based authentication with the NetApp Manageability SDK for ONTAP.](#)

ONTAP OAuth 2.0 token-based authentication for REST API

As an alternative to certificate-based authentication, you can use OAuth 2.0 token-based authentication for REST API.

Beginning with ONTAP 9.14.1, you have the option to control access to your ONTAP clusters using the Open Authorization (OAuth 2.0) framework. You can configure this feature using any of the ONTAP administrative interfaces, including the ONTAP CLI, System Manager, and REST API. However, the OAuth 2.0 authorization and access control decisions can only be applied when a client accesses ONTAP using the REST API.

OAuth 2.0 tokens replace passwords for user account authentication.

For more information about using OAuth 2.0, see the [ONTAP documentation on authentication and authorization using OAuth 2.0.](#)

Login and password parameters

An effective security posture adheres to established organizational policies, guidelines, and any governance or standards that apply to the organization. Examples of these requirements include user name lifetime, password-length requirements, character requirements, and the storage of such accounts. The ONTAP solution provides features and functions to address these security constructs.

New local account features

To support an organization's user account policies, guidelines, or standards, including governance, the following functionality is supported in ONTAP:

- Configuring password policies to enforce a minimum number of digits, lowercase characters, or uppercase characters
- Requiring a delay after a failed login attempt
- Defining the account inactive limit
- Expiring a user account
- Displaying a password expiration warning message
- Notification of an invalid login



Configurable settings are managed by using the security login role config modify command.

SHA-512 support

To enhance password security, ONTAP 9 supports the SHA-2 password hash function and defaults to using SHA-512 for hashing newly created or changed passwords. Operators and administrators can also expire or lock accounts as needed.

Pre-existing ONTAP 9 user accounts with unchanged passwords continue to use the MD5 hash function after the upgrade to ONTAP 9.0 or later. However, NetApp strongly recommends that these user accounts migrate to the more secure SHA-512 solution by having users change their passwords.

The password hash functionality enables you to perform the following tasks:

- Display user accounts that match the specified hash function:

```
cluster1::*> security login show -user-or-group-name NewAdmin -fields
hash-function
vserver  user-or-group-name  application  authentication-method  hash-
function
-----
-----
cluster1 NewAdmin          console     password            sha512
cluster1 NewAdmin          ontapi      password            sha512
cluster1 NewAdmin          ssh         password            sha512
```

- Expire accounts that use a specified hash function (for example, MD5), which forces users to change their

passwords at the next login:

```
cluster1::*> security login expire-password -vserver * -username * -hash
-function md5
```

- Lock accounts with passwords that use the specified hash function.

```
cluster1::*> security login lock -vserver * -username * -hash-function
md5
```

The password hash function is unknown for the internal `autosupport` user in your cluster's administrative SVM. This issue is cosmetic. The hash function is unknown because this internal user does not have a configured password by default.

- To view the password hash function for the `autosupport` user, run the following commands:

```
::> set advanced
::> security login show -user-or-group-name autosupport -instance

                Vserver: cluster1
User Name or Group Name: autosupport
        Application: console
Authentication Method: password
Remote Switch IP Address: -
                Role Name: autosupport
Account Locked: no
        Comment Text: -
Whether Ns-switch Group: no
        Password Hash Function: unknown
Second Authentication Method2: none
```

- To set the password hash function (default: `sha512`), run the following command:

```
::> security login password -username autosupport
```

It does not matter what the password is set to.

```

security login show -user-or-group-name autosupport -instance

                Vserver: cluster1
    User Name or Group Name: autosupport
                Application: console
    Authentication Method: password
Remote Switch IP Address: -
                Role Name: autosupport
    Account Locked: no
    Comment Text: -
    Whether Ns-switch Group: no
    Password Hash Function: sha512
    Second Authentication Method2: none

```

Password parameters

The ONTAP solution supports password parameters that address and support organizational policy requirements and guidelines.

Table 1. Restrictions for management utility user accounts

Attribute	Description	Default	Range
username-minlength	Minimum user name length required	3	3-16
username-alphanum	User name alphanumeric	disabled	Enabled/disabled
passwd-minlength	Minimum password length required	8	3-64
passwd-alphanum	Password alphanumeric	enabled	Enabled/disabled
passwd-min-special-chars	Minimum number of special characters required in the password	0	0-64
passwd-expiry-time	Password expiration time (in days)	Unlimited, which means the passwords never expire	0-unlimited 0 == expire now
require-initial-passwd-update	Require initial password update on first login	Disabled	Enabled/disabled Changes allowed through console or SSH
max-failed-login-attempts	Maximum number of failed attempts	0, do not lock account	-
lockout-duration	Maximum lockout period (in days)	The default is 0, which means the account is locked for one day	-

Attribute	Description	Default	Range
disallowed-reuse	Disallow last N passwords	6	Minimum is 6
change-delay	Delay between password changes (in days)	0	-
delay-after-failed-login	Delay after each failed login attempt (in seconds)	4	-
passwd-min-lowercase-chars	Minimum number of lowercase alphabetic characters required in the password	0, which requires no lowercase characters	0-64
passwd-min-uppercase-chars	Minimum number of uppercase alphabetic characters required	0, which requires no uppercase characters	0-64
passwd-min-digits	Minimum number of digits required in the password	0, which requires no digits	0-64
passwd-expiry-warn-time	Display warning message before password expiration (in days)	Unlimited, which means never warn about password expiration	0, which means warn user about password expiration upon every successful login
account-expiry-time	Account expires in N days	Unlimited, which means the accounts never expire	The account expiration time must be greater than the account inactive limit
account-inactive-limit	Maximum duration of inactivity before account expiration (in days)	Unlimited, which means the inactive accounts never expire	The account inactive limit must be less than the account expiration time

Example

```
cluster1::*> security login role config show -vserver cluster1 -role admin

                                Vserver: cluster1
                                Role Name: admin
                                Minimum Username Length Required: 3
                                    Username Alpha-Numeric: disabled
                                Minimum Password Length Required: 8
                                    Password Alpha-Numeric: enabled
Minimum Number of Special Characters Required in the Password: 0
                                    Password Expires In (Days): unlimited
    Require Initial Password Update on First Login: disabled
        Maximum Number of Failed Attempts: 0
            Maximum Lockout Period (Days): 0
                Disallow Last 'N' Passwords: 6
                    Delay Between Password Changes (Days): 0
                        Delay after Each Failed Login Attempt (Secs): 4
Minimum Number of Lowercase Alphabetic Characters Required in the
Password: 0
Minimum Number of Uppercase Alphabetic Characters Required in the
Password: 0
Minimum Number of Digits Required in the Password: 0
Display Warning Message Days Prior to Password Expiry (Days): unlimited
                                    Account Expires in (Days): unlimited
Maximum Duration of Inactivity before Account Expiration (Days): unlimited
```



Beginning in 9.14.1, there are increased complexity and lockout rules for passwords. This applies only to new installs of ONTAP.

System administration methods

These are important parameters to strengthen ONTAP system administration.

Command-line access

Establishing secure access to systems is a critical part of maintaining a secure solution. The most common command-line access options are SSH, Telnet, and RSH. Of these, SSH is the most secure, industry-standard best practice for remote command-line access. NetApp highly recommends using SSH for command-line access to the ONTAP solution.

SSH configurations

The `security ssh show` command shows the configurations of the SSH key exchange algorithms, ciphers, and MAC algorithms for the cluster and SVMs. The key exchange method uses these algorithms and ciphers to specify how the one-time session keys are generated for encryption and authentication and how server authentication takes place.

```
cluster1::> security ssh show
```

Vserver	Ciphers	Key Exchange Algorithms	MAC Algorithms
nsadhanacluster-2	aes256-ctr, aes192-ctr, aes128-ctr	diffie-helman-group- exchange-sha256, ecdh-sha2-nistp384	hmac-sha2-256 hmac-sha2-512
vs0	aes128-gcm	curve25519-sha256	hmac-sha1
vs1	aes256-ctr, aes192-ctr, aes128-ctr, 3des-cbc, aes128-gcm	diffie-hellman-group- exchange-sha256 ecdh-sha2-nistp384 ecdh-sha2-nistp512	hmac-sha1-96 hmac-sha2-256 hmac-sha2-256- etm hmac-sha2-512

3 entries were displayed.

Login banners

Login banners allow an organization to present any operators, administrators, and even miscreants with terms and conditions of acceptable use, and they indicate who is permitted access to the system. This approach is helpful for establishing expectations for access and use of the system. The `security login banner modify` command modifies the login banner. The login banner is displayed just before the authentication step during the SSH and console device login process. The banner text must be in double quotes (" "), as shown in the following example.

```
cluster1::> security login banner modify -vserver cluster1 -message  
"Authorized users ONLY!"
```

Login banner parameters

Parameter	Description
<code>vserver</code>	Use this parameter to specify the SVM with the modified banner. Use the name of the cluster admin SVM to modify the cluster-level message. The cluster-level message is used as the default for data SVMs that do not have a message defined.
<code>message</code>	<p>This optional parameter can be used to specify a login banner message. If the cluster has a login banner message set, the cluster login banner is used by all data SVMs as well. Setting a data SVM's login banner overrides the display of the cluster login banner. To reset a data SVM login banner to use the cluster login banner, use this parameter with the value "-".</p> <p>If you use this parameter, the login banner cannot contain newlines (also known as ends of lines [EOLs] or line breaks). To enter a login banner message with newlines, do not specify any parameter. You are prompted to enter the message interactively. Messages entered interactively can contain newlines.</p> <p>Non-ASCII characters must use Unicode UTF-8.</p>

Parameter	Description
uri	(ftp http)://(hostname IPv4 Use this parameter to specify the URI from which the login banner is downloaded. The message must not exceed 2048 bytes in length. Non-ASCII characters must be provided as Unicode UTF-8.

Message of the day

The `security login motd modify` command updates the message of the day (MOTD).

There are two categories of MOTD: the cluster-level MOTD and the data SVM-level MOTD. A user logging in to a data SVM's clustershell might see two messages: the cluster-level MOTD followed by the SVM-level MOTD for that SVM.

The cluster administrator can enable or disable the cluster-level MOTD on each SVM individually if needed. If the cluster administrator disables the cluster-level MOTD for an SVM, a user logging in to the SVM does not see the cluster-level message. Only a cluster administrator can enable or disable the cluster-level message.

MOTD Parameter	Description
Vserver	Use this parameter to specify the SVM for which the MOTD is modified. Use the name of the cluster admin SVM to modify the cluster-level message.

MOTD Parameter	Description
message	<p>This optional parameter can be used to specify a message. If you use this parameter, the MOTD cannot contain newlines. If you do not specify any parameter other than the <code>-vserver</code> parameter, you are prompted to enter the message interactively. Messages entered interactively can contain newlines. Non-ASCII characters must be provided as Unicode UTF-8. The message can contain dynamically generated content using the following escape sequences:</p> <ul style="list-style-type: none"> • <code>\</code> - A single backslash character • <code>\b</code> - No output (supported for compatibility with Linux only) • <code>\C</code> - Cluster name • <code>\d</code> - Current date as set on the login node • <code>\t</code> - Current time as set on the login node • <code>\I</code> - Incoming LIF IP address (prints console for a <code>console</code> login) • <code>\l</code> - Login device name (prints console for a <code>console</code> login) • <code>\L</code> - Last login for the user on any node in the cluster • <code>\m</code> - Machine architecture • <code>\n</code> - Node or data SVM name • <code>\N</code> - Name of user logging in • <code>\o</code> - Same as <code>\O</code>. Provided for Linux compatibility. • <code>\O</code> - DNS domain name of the node. Note that the output depends on the network configuration and may be empty. • <code>\r</code> - Software release number • <code>\s</code> - Operating system name • <code>\u</code> - Number of active clustershell sessions on the local node. For the cluster admin: all clustershell users. For the data SVM admin: only active sessions for that data SVM. • <code>\U</code> - Same as <code>\u</code>, but has <code>user</code> or <code>users</code> appended • <code>\v</code> - Effective cluster version string • <code>\W</code> - Active sessions across the cluster for the user logging in (<code>who</code>)

For more information on configuring the Message of the Day in ONTAP, see the [ONTAP documentation on message of the day](#).

CLI session timeout

The default CLI session timeout is 30 minutes. The timeout is important to prevent stale sessions and session piggybacking.

Use the `system timeout show` command to view the current CLI session timeout. To set the timeout value, use the `system timeout modify -timeout <minutes>` command.

Web access with NetApp ONTAP System Manager

If an ONTAP administrator prefers to use a graphical interface instead of the CLI for accessing and managing a cluster, use NetApp ONTAP System Manager. It is included with ONTAP as a web service, enabled by default, and accessible by using a browser. Point the browser to the host name if using DNS or the IPv4 or IPv6 address through `https://cluster-management-LIF`.

If the cluster uses a self-signed digital certificate, the browser might display a warning indicating that the certificate is not trusted. You can either acknowledge the risk to continue access or install a certificate authority (CA) signed digital certificate on the cluster for server authentication.

Beginning with ONTAP 9.3, Security Assertion Markup Language (SAML) authentication is an option for ONTAP System Manager.

SAML authentication for ONTAP System Manager

SAML 2.0 is a widely adopted industry standard that allows any third-party SAML-compliant identity provider (IdP) to perform MFA using mechanisms unique to the IdP of the enterprise's choosing and as a source of single sign-on (SSO).

There are three roles defined in the SAML specification: the principal, the IdP, and the service provider. In the ONTAP implementation, a principal is the cluster administrator gaining access to ONTAP through ONTAP System Manager or NetApp Active IQ Unified Manager. The IdP is third-party IdP software. Beginning with ONTAP 9.3, Microsoft Active Directory Federated Services (ADFS) and the open-source Shibboleth IdP are supported IdPs. Beginning with ONTAP 9.12.1, Cisco DUO is a supported IdP. The service provider is the SAML capability built into ONTAP that is used by ONTAP System Manager or the Active IQ Unified Manager web application.

Unlike the SSH two-factor configuration process, after SAML authentication is activated, ONTAP System Manager or ONTAP Service Processor access requires all existing administrators to authenticate through the SAML IdP. No changes are required to the cluster user accounts. When SAML authentication is enabled, a new authentication method of `saml` is added to existing users with administrator roles for `http` and `ontapi` applications.

After SAML authentication is enabled, additional new accounts requiring SAML IdP access should be defined in ONTAP with the administrator role and the `saml` authentication method for `http` and `ontapi` applications. If SAML authentication is disabled at some point, these new accounts require the `password` authentication method to be defined with the administrator role for `http` and `ontapi` applications and addition of the `console` application for local ONTAP authentication to ONTAP System Manager.

After the SAML IdP is enabled, the IdP performs authentication for ONTAP System Manager access by using methods available to the IdP, such as Lightweight Directory Access Protocol (LDAP), Active Directory (AD), Kerberos, password, and so on. The methods available are unique to the IdP. It is important that the accounts configured in ONTAP have user IDs that map to the IdP authentication methods.

IdPs that have been validated by NetApp are Microsoft ADFS, Cisco DUO, and open-source Shibboleth IdP.

Beginning with ONTAP 9.14.1, Cisco DUO can be used as a second authentication factor for SSH.

For more information about MFA for ONTAP System Manager, Active IQ Unified Manager, and SSH, see [TR-4647: Multifactor Authentication in ONTAP 9](#).

ONTAP System Manager insights

Beginning with ONTAP 9.11.1, ONTAP System Manager provides insights to help cluster administrators

streamline their day-to-day tasks. The security insights are based on the recommendations of this technical report.

Security Insight	Determination
Telnet is enabled	NetApp recommends Secure Shell (SSH) for secure remote access.
Remote Shell (RSH) is enabled	NetApp recommends SSH for secure remote access.
AutoSupport is using an insecure protocol	AutoSupport is not configured to be sent over xref:./ontap-security-hardening/httpS.
Login banner is not configured on the cluster at cluster level	Warning if login banner is not configured for the cluster.
SSH is using insecure ciphers	Warning if SSH uses insecure ciphers.
Too few NTP servers are configured	Warning if the number of NTP servers configured is less than three.
Default admin user not locked	When not using any default administrative accounts (admin or diag) to log in to System Manager, and these accounts are not locked, the recommendation is to lock them.
Ransomware defense - volumes don't have Snapshot policies	No adequate Snapshot policy is attached to one or more volumes.
Ransomware defense - disable Snapshot auto-delete	Snapshot auto-delete is set for one or more volumes.
Volumes are not being monitored for ransomware attacks	Autonomic ransomware protection is supported on several volumes but not yet configured.
SVMs are not configured for autonomic ransomware protection	Autonomic ransomware protection is supported on several SVMs but not yet configured.
Native FPolicy is not configured	FPolicy is not set for NAS SVMs.
Enable autonomic ransomware protection active mode	Several volumes have completed their learning mode and you can switch on active mode
Global FIPS 140-2 compliance is disabled	Global FIPS 140-2 compliance is not enabled.
Cluster is not configured for notifications	Emails, webhooks or SNMP traphosts are not configured to receive notifications.

For more information about ONTAP System Manager insights, see the [ONTAP System Manager insights documentation](#).

ONTAP autonomous ransomware protection

To supplement user behavior analytics for Storage Workload Security, the ONTAP autonomous ransomware protection analyzes volume workloads and entropy to detect ransomware and takes a Snapshot and notifies the administrator when an attack is suspected.

In addition to ransomware detection and prevention using external FPolicy user behavioral analytics (UBA) with NetApp Cloud Insights / Cloud Secure and the NetApp FPolicy partner ecosystem, ONTAP 9.10.1

introduces autonomous ransomware protection. ONTAP autonomous ransomware protection uses a built-in on-box machine learning (ML) capability that looks at volume workload activity plus data entropy to automatically detect ransomware. It monitors for activity that is different from UBA so that it can detect attacks that UBA does not.

For more detailed information about this capability, see [NetApp solutions for ransomware](#) or [ONTAP autonomous ransomware protection documentation](#).

Storage administrative system auditing

Ensure the integrity of event auditing by offloading ONTAP events to a remote syslog server. This server could be a security information event management system such as Splunk.

Send out syslog

Log and audit information is invaluable to an organization from a support and availability standpoint. In addition, the information and details contained in logs (syslog) and audit reports and outputs are generally of a sensitive nature. To maintain security controls and posture, it is imperative that organizations manage log and audit data in a secure manner.

Offloading of syslog information is necessary for limiting the scope or footprint of a breach to a single system or solution. Therefore, NetApp recommends securely offloading syslog information to a secure storage or retention location.

Create a log-forwarding destination

Use the `cluster log-forwarding create` command to create log-forwarding destinations for remote logging.

Parameters

Use the following parameters to configure the `cluster log-forwarding create` command:

- **Destination host.** This name is the host name or IPv4 or IPv6 address of the server to which to forward the logs.

```
-destination <Remote InetAddress>
```

- **Destination port.** This is the port on which the destination server listens.

```
[-port <integer>]
```

- **Log-forwarding protocol.** This protocol is used for sending messages to the destination.

```
[-protocol \{udp-unencrypted|tcp-unencrypted|tcp-encrypted\}]
```

The log-forwarding protocol can use one of the following values:

- `udp-unencrypted`. User Datagram Protocol with no security.
- `tcp-unencrypted`. TCP with no security.
- `tcp-encrypted`. TCP with Transport Layer Security (TLS).
- **Verify destination server identity.** When this parameter is set to `true`, the identity of the log-forwarding destination is verified by validating its certificate. The value can be set to `true` only when the `tcpencrypted` value is selected in the protocol field.

```
[-verify-server \{true|false\}]
```

- **Syslog facility.** This value is the syslog facility to use for the forwarded logs.

```
[-facility <Syslog Facility>]
```

- **Skip the connectivity test.** Normally, the `cluster log-forwarding create` command checks that the destination is reachable by sending an Internet Control Message Protocol (ICMP) ping and fails if it is not reachable. Setting this value to `true` bypasses the ping check so that you can configure the destination when it is unreachable.

```
[-force [true]]
```



NetApp recommends using the `cluster log-forwarding` command to force the connection to a `-tcp-encrypted` type.

Event notification

Securing the information and data leaving a system is vital to maintaining and managing the system's security posture. The events generated by the ONTAP solution provide a wealth of information about what the solution is encountering, the information processed, and more. The vitality of this data highlights the need to manage and migrate it in a secure manner.

The `event notification create` command sends a new notification of a set of events defined by an event filter to one or more notification destinations. The following examples depict the event notification configuration and the `event notification show` command, which displays the configured event notification filters and destinations.

```
cluster1::> event notification create -filter-name filter1 -destinations
email_dest,syslog_dest,snmp-traphost
```

```
cluster1::> event notification show
ID      Filter Name      Destinations
-----
1 filter1 email_dest, syslog_dest, snmp-traphost
```

Storage encryption

To protect sensitive data in the event of a disk that is stolen, returned, or repurposed use hardware-based NetApp Storage Encryption or software-based NetApp Volume Encryption/NetApp Aggregate Encryption. Both mechanisms are FIPS-140-2 validated and when using hardware-based mechanisms with software-based mechanisms, the solution qualifies for Commercial Solutions for Classified (CSfC) Program. It enables enhanced security protection for secret and top-secret data at rest at both the hardware and software layers.

Data-at-rest encryption is important to protect sensitive data in the event of a disk that is stolen, returned, or repurposed.

ONTAP 9 has three Federal Information Processing Standard (FIPS) 140-2-compliant data-at-rest encryption solutions:

- NetApp Storage Encryption (NSE) is a hardware solution that uses self-encrypting drives.
- NetApp Volume Encryption (NVE) is a software solution that enables encryption of any data volume on any drive type where it is enabled with a unique key for each volume.
- NetApp Aggregate Encryption (NAE) is a software solution that enables encryption of any data volume on any drive type where it is enabled with unique keys for each aggregate.

NSE, NVE, and NAE can use either external key management or the onboard key manager (OKM). Use of NSE, NVE, and NAE does not affect ONTAP storage efficiency features. However, NVE volumes are excluded from aggregate deduplication. NAE volumes participate in and benefit from aggregate deduplication.

The OKM provides a self-contained encryption solution for data at rest with NSE, NVE, or NAE.

NVE, NAE, and OKM use the ONTAP CryptoMod. CryptoMod is listed on the CMVP FIPS 140-2 validated modules list. See [FIPS 140-2 Cert# 4144](#).

To begin OKM configuration, use the `security key-manager onboard enable` command. To configure external Key Management Interoperability Protocol (KMIP) key managers, use the `security key-manager external enable` command. Starting with ONTAP 9.6, multitenancy is supported for external key managers. Use the `-vserver <vserver name>` parameter to enable external key management for a specific SVM. Prior to 9.6, the `security key-manager setup` command was used to configure both OKM and external key managers. For onboard key management, this configuration walks the operator or administrator through the passphrase setup and additional parameters for configuring OKM.

A part of the configuration is provided in the following example:

```
cluster1::> security key-manager setup
Welcome to the key manager setup wizard, which will lead you through
the steps to add boot information.
```

Enter the following commands at any time
"help" or "?" if you want to have a question clarified,
"back" if you want to change your answers to previous questions, and
"exit" if you want to quit the key manager setup wizard. Any changes
you made before typing "exit" will be applied.

Restart the key manager setup wizard with "security key-manager setup". To
accept a default
or omit a question, do not enter a value.

Would you like to configure onboard key management? {yes, no} [yes]:
Enter the cluster-wide passphrase for onboard key management. To continue
the configuration, enter the passphrase, otherwise
type "exit":
Re-enter the cluster-wide passphrase:
After configuring onboard key management, save the encrypted configuration
data
in a safe location so that you can use it if you need to perform a manual
recovery
operation. To view the data, use the "security key-manager backup show"
command.

Beginning with ONTAP 9.4, You can use the `-enable-cc-mode true` option with `security key-manager setup` to require that users enter the passphrase after a reboot. For ONTAP 9.6 and later, the command syntax is `security key-manager onboard enable -cc-mode-enabled yes`.

Beginning with ONTAP 9.4, you can use the `secure-purge` feature with advanced privilege to nondisruptively "scrub" data on NVE-enabled volumes. Scrubbing data on an encrypted volume ensures that it cannot be recovered from the physical media. The following command securely purges the deleted files on vol1 on SVM vs1:

```
cluster1::> volume encryption secure-purge start -vserver vs1 -volume vol1
```

Beginning with ONTAP 9.7, NAE and NVE are enabled by default if the VE license is in place, either OKM or external key managers are configured, and NSE is not used. NAE volumes are created by default on NAE aggregates, and NVE volumes are created by default on non-NAE aggregates. You can override this by entering the following command:

```
cluster1::*> options -option-name
encryption.data_at_rest_encryption.disable_by_default true
```


Beginning with ONTAP 9.6, you can use an SVM scope to configure external key management for a data SVM in the cluster. This is best for multitenant environments in which each tenant uses a different SVM (or set of SVMs) to serve data. Only the SVM administrator for a given tenant has access to the keys for that tenant. For more information, see [enable external key management in ONTAP 9.6 and later](#) in the ONTAP documentation.

Beginning in ONTAP 9.11.1, you can configure connectivity to clustered external key management servers by designating primary and secondary key servers on an SVM. For more information, see [configure clustered external key servers](#) in the ONTAP documentation.

Beginning in ONTAP 9.13.1, you can configure external key manager servers in system manager. For more information, see [Manage external key managers](#) in the ONTAP documentation.

Data replication encryption

To supplement data at rest encryption, you can encrypt ONTAP data replication traffic between clusters using TLS 1.2 with a pre-shared key for SnapMirror, SnapVault, or FlexCache.

When replicating data for disaster recovery, caching, or backup, you must protect that data during transport over the wire from one ONTAP cluster to another. Doing so prevents malicious man-in-the-middle attacks against sensitive data while it is in flight.

Beginning with ONTAP 9.6, Cluster Peering Encryption provides TLS 1.2 AES-256 GCM encryption support for ONTAP data replication features such as SnapMirror, SnapVault, and FlexCache. Encryption is setup by way of a pre-shared key (PSK) between two cluster peers.

Customers who use technologies like NSE, NVE, and NAE to protect data at rest can also use end-to-end data encryption by upgrading to ONTAP 9.6 or later to use Cluster Peering Encryption.

Cluster peering encrypts all data between the cluster peers. For example, when using SnapMirror, all peering information as well as all SnapMirror relationships between the source and destination cluster peer are encrypted. You cannot send clear-text data between cluster peers with Cluster Peering Encryption enabled.

Beginning with ONTAP 9.6, new cluster-peer relationships have encryption enabled by default. To enable encryption on cluster peer relationships that were created before ONTAP 9.6, you must upgrade the source and destination cluster to 9.6. In addition, you must use the `cluster peer modify` command to change both the source and destination cluster peers to use Cluster Peering Encryption.

You can convert an existing peer relationship to use Cluster Peering Encryption in ONTAP 9.6 as shown in the following example:

On the Destination Cluster Peer

```
cluster2::> cluster peer modify cluster1 -auth-status-admin use-  
authentication -encryption-protocol-proposed tls-psk
```

When prompted enter a passphrase.

On the Source Cluster Peer

```
cluster1::> cluster peer modify cluster2 -auth-status-admin use-  
authentication -encryption-protocol-proposed tls-psk
```

When prompted enter the same passphrase you created in the previous step.

IPsec data-in-flight encryption

Customers who use data-at-rest encryption technologies such as NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) and Cluster Peering Encryption (CPE) for data replication traffic can now use end-to-end encryption between client and storage across their hybrid multi-cloud data fabric by upgrading to ONTAP 9.8 or later and using IPsec. IPsec provides an alternative to NFS or SMB/CIFS encryption and is the only encryption in flight option for iSCSI traffic.

In some situations, there might be a requirement to protect all client data transported over the wire (or in flight) to the ONTAP SVM. Doing so prevents replay and malicious man-in-the-middle attacks against sensitive data while it is in flight.

Starting with ONTAP 9.8, Internet Protocol Security (IPsec) provides end-to-end encryption support for all IP traffic between a client and an ONTAP SVM. IPsec data encryption for all IP traffic includes NFS, iSCSI, and SMB/CIFS protocols. IPsec provides the only encryption in flight option for iSCSI traffic.

Providing NFS encryption over the wire is one of the main use cases for IPsec. Prior to ONTAP 9.8, NFS over-the-wire encryption required the setup and configuration of Kerberos to utilize krb5p to encrypt NFS data in flight. This is not always simple or easy to accomplish in every customer environment.

Customers who use data-at-rest encryption technologies such as NetApp Storage Encryption (NSE) or NetApp Volume Encryption (NVE) and Cluster Peering Encryption (CPE) for data replication traffic can now use end-to-end encryption between client and storage across their hybrid multi-cloud data fabric by upgrading to ONTAP 9.8 or later and using IPsec.

IPsec is an IETF standard. ONTAP uses IPsec in transport mode. It also leverages the Internet Key Exchange (IKE) protocol version 2, which uses a pre-shared key (PSK) for negotiating key material between the client and ONTAP with either IPv4 or IPv6. By default, IPsec uses Suite-B AES-GCM 256-bit encryption. Suite-B AES-GMAC256 and AES-CBC256 with 256-bit encryption are also supported.

Although the IPsec capability must be enabled on the cluster, it applies to individual SVM IP addresses through the use of a Security Policy Database (SPD) entry. The policy (SPD) entry contains the client IP address (remote IP subnet), SVM IP address (local IP subnet), the encryption cipher suite to use, and the pre-shared secret (PSK) needed to authenticate via IKEv2 and establish the IPsec connection. In addition to the IPsec

policy entry, the client must be configured with the same information (local and remote IP, PSK, and cipher suite) before traffic can flow over the IPsec connection. Beginning with ONTAP 9.10.1, IPsec certificate authentication support is added. This removes IPsec policy limits and enables Windows OS support for IPsec.

If there is a firewall between the client and the SVM IP address, then it must allow the ESP and UDP (port 500 and 4500) protocols, both inbound (ingress) and outbound (egress), for the IKEv2 negotiation to succeed and thus allow IPsec traffic.

For NetApp SnapMirror and cluster peering traffic encryption, cluster peering encryption (CPE) is still recommended over IPsec for secure in-transit over the wire. CPE performs better for these workloads than IPsec. You do not need a license for IPsec, and there are no import or export restrictions.

You can enable IPsec on the cluster and create an SPD entry for a single client and a single SVM IP address as shown in the following example:

```
On the Destination Cluster Peer
```

```
cluster1::> security ipsec config modify -is-enabled true
```

```
cluster1::> security ipsec policy create -vserver vs1 -name test34 -local  
-ip-subnets 192.168.134.34/32 -remote-ip-subnets 192.168.134.44/32
```

```
When prompted enter and confirm the pre shared secret (PSK).
```

FIPS mode and TLS and SSL management

The FIPS 140-2 standard specifies security requirements for cryptographic modules within security systems that protect sensitive information in computer and telecommunication systems. The FIPS 140-2 standard applies *specifically* to the cryptographic module, rather than the product, architecture, data, or ecosystem. The cryptographic module is the specific component (hardware, software, firmware, or a combination of the three) that implements NIST-approved security functions.

Enabling FIPS 140-2 compliance has effects on other systems and communications internal and external to ONTAP 9. NetApp highly recommends testing these settings on a nonproduction system that has console access.

Beginning with ONTAP 9.11.1 and TLS 1.3 support, you can validate FIPS 140-3.



The FIPS configuration applies to ONTAP and the platform BMC.

NetApp ONTAP's FIPS-mode configuration

NetApp ONTAP has a FIPS-mode configuration that instantiates an added level of security to the control plane:

- Beginning in ONTAP 9.11.1 when FIPS 140-2 compliance mode is enabled, TLSv1, TLSv1.1 and SSLv3 are disabled, and only TLSv1.2 and TLSv1.3 remain enabled. It affects other systems and communications that are internal and external to ONTAP 9. If you enable FIPS 140-2 compliance mode and then subsequently disable, TLSv1, TLSv1.1, and SSLv3 remain disabled. Either TLSv1 or TLSv1.3 will remain enabled depending on the previous configuration.

- For versions of ONTAP prior to 9.11.1 when FIPS 140-2 compliance mode is enabled, both TLSv1 and SSLv3 are disabled and only TLSv1.1 and TLSv1.2 remain enabled. ONTAP prevents you from enabling both TLSv1 and SSLv3 when FIPS 140-2 compliance mode is enabled. If you enable FIPS 140-2 compliance mode and then subsequently disable it, TLSv1 and SSLv3 remain disabled, but either TLSv1.2 or both TLSv1.1 and TLSv1.2 are enabled depending on the previous configuration.
- [NetApp Cryptographic Security Module \(NCSM\)](#), which is FIPS 140-2 level 1 validated, provides software-based compliance.



NIST has submitted a FIPS-140-3 standard, and NCSM will have FIPS-140-2 and FIPS-140-3 validations. All FIPS 140-2 validations will move to historical status on September 21, 2026, which is five years after the last day for new certificate submissions.

Enable FIPS-140-2 and FIPS-140-3 compliance mode

Beginning with ONTAP 9, you can enable the FIPS-140-2 and FIPS-140-3 compliance mode for cluster-wide control plane interfaces.

- [Enable FIPS](#)
- [View FIPS status](#)

FIPS enablement and protocols

The `security config modify` command allows you to modify the existing cluster-wide security configuration. If you enable FIPS-compliant mode, the cluster automatically selects only TLS protocols.

- Use the `-supported-protocols` parameter to include or exclude TLS protocols independently from FIPS mode. By default, FIPS mode is disabled, and ONTAP supports the TLSv1.2, TLSv1.1, and TLSv1 protocols.
- For backward compatibility, ONTAP supports adding SSLv3 to the supported-protocols list when FIPS mode is disabled.

FIPS enablement and ciphers

- Use the `-supported-cipher-suites` parameter to configure only the Advanced Encryption Standard (AES) or AES and 3DES.
- You can disable weak ciphers such as RC4 by specifying `!RC4`. By default, the supported cipher setting is `ALL:!LOW:!aNULL:!EXP:!eNULL`. This setting means that all supported cipher suites for the protocols are enabled, except for the ones using 64-bit or 56-bit encryption algorithms with no authentication, no encryption, no exports, and low-encryption cipher suites.
- Select a cipher suite that is available with the corresponding selected protocol. An invalid configuration might cause some functionality to fail to operate properly.
- For the correct cipher string syntax, see the [ciphers page](#) on OpenSSL (published by the OpenSSL software foundation). Beginning with ONTAP 9.9.1 and later releases, you are no longer required to reboot all the nodes manually after modifying the security configuration.

SSH and TLS security hardening

SSH administration of ONTAP 9 requires an OpenSSH client 5.7 or later. SSH clients must negotiate with the Elliptic Curve Digital Signature Algorithm (ECDSA) public key algorithm for the connection to be successful.

To harden TLS security, enable only TLS 1.2 and use cipher suites capable of Perfect Forward Secrecy (PFS).

PFS is a method of key exchange that, when used in combination with encryption protocols like TLS 1.2, helps prevent an attacker from decrypting all network sessions between a client and server.

Enable TLSv1.2 and PFS-capable cipher suites

To enable only TLS 1.2 and PFS-capable cipher suites, use the `security config modify` command from the advanced privilege level.



Before changing the SSL interface configuration, ensure that the client supports the ciphers DHE and ECDHE when connecting to ONTAP to maintain connectivity with ONTAP.

Example

```
cluster1::*> security config modify -interface SSL -supported-protocols
TLSv1.2 -supported-cipher-suites
PSK:DHE:ECDHE:!LOW:!aNULL:!EXP:!eNULL:!3DES:!kDH:!kECDH
```

Confirm `y` for each prompt. For more information on PFS, see this [NetApp blog](#).

Related information

[Federal Information Processing Standard \(FIPS\) Publication 140](#)

Create a CA-signed digital certificate

For many organizations, the self-signed digital certificate for ONTAP web access is not compliant with their InfoSec policies. On production systems, it is a NetApp best practice to install a CA-signed digital certificate for use in authenticating the cluster or SVM as an SSL server.

You can use the `security certificate generate-csr` command to generate a certificate signing request (CSR), and the `security certificate install` command to install the certificate you receive back from the CA.

Steps

1. To create a digital certificate that is signed by the organization's CA, do the following:
 - a. Generate a CSR.
 - b. Follow your organization's procedure to request a digital certificate using the CSR from your organization's CA. For example, using Microsoft Active Directory Certificate Services web interface, go to `<CA_server_name>/certsrv` and request a certificate.
 - c. Install the digital certificate in ONTAP.

Online certificate status protocol

Online Certificate Status Protocol (OCSP) enables ONTAP applications that use TLS communications, such as LDAP or TLS, to receive digital certificate status when OCSP is enabled. The application receives a signed response signifying that the certificate requested is good, revoked, or unknown.

OCSP enables determination of the current status of a digital certificate without requiring certificate revocation lists (CRLs).

By default, OCSP certificate status checking is disabled. It can be turned on with the command `security config ocsd enable -app name`, where the app name can be `autosupport`, `audit_log`, `fabricpool`, `ems`, `kmip`, `ldap_ad`, `ldap_nis_namemap`, or `all`. The command requires advanced privilege level.

SSHv2 management

The `security ssh modify` command replaces the existing configurations of the SSH key exchange algorithms, ciphers, or MAC algorithms for the cluster or an SVM with the configuration settings you specify.



NetApp recommends the following:

- Use passwords for user sessions.
- Use a public key for machine access.

Supported ciphers and key exchanges

Ciphers	Key exchange
aes256-ctr	diffie-hellman-group-exchange-sha256 (SHA-2)
aes192-ctr	diffie-hellman-group-exchange-sha1 (SHA-1)
aes128-ctr	diffie-hellman-group14-sha1 (SHA-1)
aes256-cbc	diffie-hellman-group1-sha1 (SHA-1)
aes192-cbc	-
aes128-cbc	-
aes128-gcm	-
aes256-gcm	-
3des-cbc	-

Supported AES and 3DES symmetric encryptions

ONTAP also supports the following types of AES and 3DES symmetric encryptions (also known as ciphers):

- hmac-sha1
- hmac-sha1-96
- hmac-md5
- hmac-md5-96
- hmac-ripemd160
- umac-64
- umac-64

- umac-128
- hmac-sha2-256
- hmac-sha2-512
- hmac-sha1-etm
- hmac-sha1-96-etm
- hmac-sha2-256-etm
- hmac-sha2-512-etm
- hmac-md5-etm
- hmac-md5-96-etm
- hmac-ripemd160-etm
- umac-64-etm
- umac-128-etm



The SSH management configuration applies to ONTAP and the platform BMC.

NetApp AutoSupport

The AutoSupport feature of ONTAP allows you to proactively monitor the health of your system and automatically send messages and details to NetApp technical support, your organization's internal support team, or a support partner. By default, AutoSupport messages to NetApp technical support are enabled when the storage system is configured for the first time. In addition, AutoSupport begins sending messages to NetApp technical support 24 hours after it is enabled. This 24-hour period is configurable. To leverage the communication to an organization's internal support team, the mail host configuration must be completed.

Only the cluster administrator can perform AutoSupport management (configuration). The SVM administrator has no access to AutoSupport. The AutoSupport feature can be disabled. However, NetApp recommends enabling it because AutoSupport helps speed problem identification and resolution should an issue arise on the storage system. By default, the system collects AutoSupport information and stores it locally even if you disable AutoSupport.

For more details regarding AutoSupport messages, including what is contained in the various messages and where different types of messages are sent, see the [NetApp Active IQ Digital Advisor](#) documentation.

AutoSupport messages contain sensitive data including, but not limited to, the following items:

- Log files
- Context-sensitive data regarding specific subsystems
- Configuration and status data
- Performance data

AutoSupport supports HTTPS, HTTP, and SMTP for transport protocols. Because of the sensitive nature of AutoSupport messages, NetApp strongly recommends using HTTPS as the default transport protocol for sending AutoSupport messages to NetApp support.

In addition, you should leverage the `system node autosupport modify` command to specify the targets of AutoSupport data (for example, NetApp technical support, an organization's internal operations, or partners). This command also allows you to specify what specific AutoSupport details to send (for example, performance data, log files, and so on).

To entirely disable AutoSupport, use the `system node autosupport modify -state disable` command.

Network Time Protocol

Although ONTAP enables you to manually set the time zone, date, and time on the cluster, you must configure the Network Time Protocol (NTP) servers to synchronize the cluster time with at least three external NTP servers.

Problems can occur when the cluster time is inaccurate. Although ONTAP enables you to manually set the time zone, date, and time on the cluster, you must configure the Network Time Protocol (NTP) servers to synchronize the cluster time with external NTP servers.

Beginning with ONTAP 9.5, you can configure your NTP server with symmetric authentication.

You can associate a maximum of 10 external NTP servers by using the `cluster time-service ntp server create` command. For redundancy and quality of time service, you should associate at least three external NTP servers with the cluster.

For details about the configuration of NTP in ONTAP, see [Managing the cluster time \(cluster administrators only\)](#).

NAS file system local accounts (CIFS workgroup)

Workgroup client authentication provides an extra layer of security to the ONTAP solution that is consistent with a traditional domain authentication posture. Use 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.

Starting 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 extra layer of security to the ONTAP solution that is consistent with a traditional domain authentication posture. To configure the CIFS server, use the `vserver cifs create` command. After the CIFS server is created, you can join it to a CIFS domain or join it to a workgroup. To join a workgroup, use the `-workgroup` parameter. Here is an example configuration:

```
cluster1::> vserver cifs create -vserver vs1 -cifs-server CIFSSERVER1
-workgroup Sales
```



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.

NAS file system auditing

NAS file systems occupy an increased footprint in today's threat landscape, audit functions are critical to support visibility.

Security requires validation. ONTAP 9 provides increased auditing events and details across the solution. Because NAS file systems occupy an increased footprint in today's threat landscape, audit functions are critical to support visibility. Because of the improved audit capability in ONTAP 9, CIFS audit details are more plentiful than ever. Key details, including the following, are logged with events created:

- File, folder, and share access
- Files created, modified, or deleted
- Successful file read access
- Failed attempts to read or write files
- Folder permission changes

Create an audit configuration

You must enable CIFS auditing to generate auditing events. Use the `vserver audit create` command to create an audit configuration. By default, the audit log uses a rotation method based on size. You can use a time-based rotation option if specified in the Rotation Parameters field. Additional log audit rotation configuration details include the rotation schedule, the rotation limits, the rotation days of the week, and the rotation size. The following text provides an example configuration depicting an audit configuration using a monthly time-based rotation scheduled for all days of the week at 12:30.

```
cluster1::> vserver audit create -vserver vs1 -destination /audit_log
-rotate-schedule-month all -rotate-schedule-dayofweek all -rotate-schedule
-hour 12 -rotate-schedule-minute 30
```

CIFS audit events

The CIFS audit events are as follows:

- **File share:** Generates an audit event when a CIFS network share is added, modified, or deleted using the related `vserver cifs share` commands.
- **Audit policy change:** Generates an audit event when the audit policy is disabled, enabled, or modified using the related `vserver audit` commands.
- **User account:** Generates an audit event when a local CIFS or UNIX user is created or deleted; a local user account is enabled, disabled, or modified; or a password is reset or changed. This event uses the `vserver cifs users-and-groups local-group` command or the related `vserver services name-service unix-user` command.
- **Security group:** Generates an audit event when a local CIFS or UNIX security group is created or deleted using the `vserver cifs users-and-groups local-group` command or the related `vserver services name-service unix-group` command.

- **Authorization policy change:** Generates an audit event when rights are granted or revoked for a CIFS user or a CIFS group using the `vserver cifs users-and-groups privilege` command.



This functionality is based on the system audit function, which enables an administrator to review what the system is allowing and performing from the perspective of a data user.

Effect of REST APIs on NAS auditing

ONTAP includes the ability for administrator accounts to access and manipulate SMB/CIFS or NFS files using REST APIs. Although REST APIs can only be run by ONTAP administrators, REST API commands do bypass the system NAS audit log. Additionally, file permissions can also be bypassed by ONTAP administrators when using REST APIs. However, the administrator's actions with REST APIs on files are captured in the system command history log.

Create no-access REST API role

You can prevent ONTAP administrators from using REST APIs for file access by creating a REST API role that does not have access to ONTAP volumes via REST. To provision this role, complete the following steps.

Steps

1. Create a new REST role that has no access to storage volumes but has all other REST API access.

```
cluster1::> security login rest-role create nofiles -vserver cluster1
"/api/storage/volumes" -access none
cluster1::> security login rest-role create nofiles -vserver cluster1
"/api" -access all
```

2. Assign the administrator account to the new REST API role you created in the previous step.

```
cluster1::> security login modify -user-or-group-name user1 -application
http -authentication-method password -vserver cluster1 -role nofile
```



If you want to prevent the built-in ONTAP cluster administrator account from using REST APIs for file access, you need to first [create a new administrator account and disable or delete the built-in account](#).

Configure and enable CIFS SMB signing and sealing

You can configure and enable SMB signing that protects the security of the data fabric by making sure that traffic between storage systems and clients is not compromised by replay or man-in-the-middle attacks. SMB signing protects by verifying that SMB messages have valid signatures.

About this task

A common threat vector for file systems and architectures lies in the SMB protocol. To address this vector, the ONTAP 9 solution uses industry-standard SMB signing and sealing. SMB signing protects the security of the data fabric by making sure that traffic between storage systems and clients is not compromised by replay or

man-in-the-middle attacks. It does so by verifying that SMB messages have valid signatures.

Although SMB signing is disabled by default in the interest of performance, NetApp highly recommends that you enable it. In addition, the ONTAP solution supports SMB encryption, which is also known as sealing. This approach enables the secure transport of data on a share-by-share basis. By default, SMB encryption is disabled. However, NetApp recommends that you enable SMB encryption.

LDAP signing and sealing are now supported in SMB 2.0 and later. Signing (protection against tampering) and sealing (encryption) enable secure communication between SVMs and Active Directory servers. Accelerated AES new instructions (Intel AES NI) encryption is now supported in SMB 3.0 and later. Intel AES NI improves on the AES algorithm and accelerates data encryption with supported processor families.

Steps

1. To configure and enable SMB signing, use the `vserver cifs security modify` command and verify that the `-is-signing-required` parameter is set to `true`. See the following example configuration:

```
cluster1::> vserver cifs security modify -vserver vs1 -kerberos-clock
-skew 3 -kerberos-ticket-age 8 -is-signing-required true
```

2. To configure and enable SMB sealing and encryption, use the `vserver cifs security modify` command and verify that the `-is-smb-encryption-required` parameter is set to `true`. See the following example configuration:

```
cluster1::> vserver cifs security modify -vserver vs1 -is-smb-encryption
-required true

cluster1::> vserver cifs security show -vserver vs1 -fields is-smb-
encryption-required
vserver  is-smb-encryption-required
-----
vs1      true
```

NFS securing

Export rules are the functional elements of an export policy. Export rules match client access requests for a volume against specific parameters you configure to determine how to handle the client access requests. An export policy must contain at least one export rule to allow access to clients. If an export policy contains more than one rule, the rules are processed in the order in which they appear in the export policy.

Access control is central to maintaining a secure posture. Therefore, ONTAP uses the export policy feature to limit NFS volume access to clients that match specific parameters. Export policies contain one or more export rules that process each client access request. An export policy is associated with each volume to configure client access to the volume. The result of this process determines whether the client is granted or denied (with a permission-denied message) access to the volume. This process also determines what level of access is provided to the volume.



An export policy with export rules must exist on an SVM for clients to access data. An SVM can contain multiple export policies.

The rule order is dictated by the rule index number. If a rule matches a client, the permissions of that rule are used, and no further rules are processed. If no rules match, the client is denied access.

Export rules determine client access permissions by applying the following criteria:

- The file access protocol used by the client sending the request (for example, NFSv4 or SMB)
- A client identifier (for example, host name or IP address)
- The security type used by the client to authenticate (for example, Kerberos v5, NTLM, or AUTH_SYS)

If a rule specifies multiple criteria, and the client does not match one or more of them, the rule does not apply.

An example export policy contains an export rule with the following parameters:

- `-protocol nfs`
- `-clientmatch 10.1.16.0/255.255.255.0`
- `-rorule any`
- `-rwrule any`

The security type determines which level of access a client receives. The three access levels are read-only, read-write, and superuser (for clients with the user ID 0). Because the access level determined by the security type is evaluated in this order, you must observe the rules listed:

Rules for access-level parameters in export rules

For a client to obtain the following access levels	These access parameters must match the client's security type
Normal user read-only	Read-only (<code>-rorule</code>)
Normal user read-write	Read-only (<code>-rorule</code>) and read-write (<code>-rwrule</code>)
Superuser read-only	Read-only (<code>-rorule</code>) and <code>-superuser</code>
Superuser read-write	Read-only (<code>-rorule</code>) and read-write (<code>-rwrule</code>) and <code>-superuser</code>


The following are valid security types for each of these three access parameters:

- Any
- None
- Never

These security types are not valid for use with the `-superuser` parameter:

- krb5
- ntlm
- sys

Rules for access parameter outcomes

If the client's security type ...	Then ...
Matches a security type specified in the access parameter.	The client receives access for that level with its own user ID.
Does not match a specified security type, but the access parameter includes the option <code>none</code> .	The client receives access for that level and receives the anonymous user with the user ID specified by the <code>-anon</code> parameter.
Does not match a security type specified, and the access parameter does not include the option <code>none</code> .	<p>The client does not receive any access for that level.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;">  <p>This restriction does not apply to the <code>-superuser</code> parameter because this parameter always includes <code>none</code>, even when not specified.</p> </div>

Kerberos 5 and Krb5p

Beginning with ONTAP 9, Kerberos 5 authentication with privacy service (krb5p) is supported. The krb5p authentication mode is secure, and it protects against data tampering and snooping by using checksums to encrypt all traffic between client and server. The ONTAP solution supports 128-bit and 256-bit AES encryption for Kerberos. The privacy service includes verifying the integrity of the received data, authenticating users, and encrypting data before transmission.

The krb5p option is most present in the export policy feature, where it is set as an encryption option. The krb5p authentication method can be used as an authentication parameter, as shown in the following example:

```
cluster1::> vserver export-policy check-access -vserver vs1 -client-ip
10.22.32.42 -volume flex_vol -authentication-method krb5p -protocol nfs3
-access- type read
```

Enable Lightweight Directory Access Protocol signing and sealing

Signing and sealing are supported to enable session security on queries to an LDAP server. This approach provides an alternative to LDAP-over-TLS session security.

Signing confirms the integrity of LDAP payload data using secret key technology. Sealing encrypts the LDAP payload data to avoid transmitting sensitive information in clear text. The session security settings on an SVM correspond to those available on the LDAP server. By default, LDAP signing and sealing are disabled.

Steps

1. To enable this function, run the `vserver cifs security modify` command with the `session-security-for-ad-ldap` parameter.

Options for LDAP security functions:

- **None:** Default, no signing or sealing
- **Sign:** Sign LDAP traffic
- **Seal:** Sign and encrypt LDAP traffic



The sign and seal parameters are cumulative, meaning that if the sign option is used, the outcome is LDAP with signing. However, if the seal option is used, the outcome is both sign and seal. In addition, if a parameter is not specified for this command, the default is none.

The following is an example configuration:

```
cluster1::> vserver cifs security modify -vserver vs1 -kerberos-clock  
-skew 3 -kerberos-ticket-age 8 -session-security-for-ad-ldap seal
```

Create and use a NetApp FPolicy

You can create and use an FPolicy, an infrastructure component of the ONTAP solution, that allows partner applications to monitor and set file access permissions. One of the more powerful applications is Storage Workload Security, a NetApp SaaS application that provides centralized visibility and control of all corporate data access across hybrid cloud environments to ensure security and compliance goals are met.

Access control is a key security concept. Visibility and the ability to respond to file access and file operations are critical for maintaining your security posture. To provide visibility and access control for files, the ONTAP solution uses the NetApp FPolicy feature.

File policies can be set based on file type. FPolicy determines how the storage system handles requests from individual client systems for operations such as create, open, rename, and delete. Beginning with ONTAP 9, the FPolicy file access notification framework is enhanced with filtering controls and resiliency against short network outages.

Steps

1. To leverage the FPolicy feature, you must first create the FPolicy policy with the `vserver fpolicy policy create` command.



In addition, use the `-events` parameter if you use FPolicy for visibility and the collection of events. The additional granularity provided by ONTAP enables filtering and access down to the user name level of control. To control privileges and access with user names, specify the `-privilege-user-name` parameter.

The following text provides an example of FPolicy creation:

```
cluster1::> vserver fpolicy policy create -vserver vs1.example.com
-policy-name vs1_pol -events cserver_evt,v1e1 -engine native -is
-mandatory true -allow-privileged-access no -is-passthrough-read-enabled
false
```

2. After you create the FPolicy policy, you must enable it with the `vserver fpolicy enable` command. This command also sets the priority or sequence of the FPolicy entry.



The FPolicy sequence is important because, if multiple policies have subscribed to the same file access event, the sequence dictates the order in which access is granted or denied.

The following text provides a sample configuration for enabling the FPolicy policy and validating the configuration with the `vserver fpolicy show` command:

```
cluster1::> vserver fpolicy enable -vserver vs2.example.com -policy-name
vs2_pol -sequence-number 5

cluster1::> vserver fpolicy show
Vserver                Policy Name                Sequence  Status
Engine
-----
vs1.example.com        vs1_pol
vs2.example.com        vs2_pol
external
2 entries were displayed.
```

FPolicy enhancements

ONTAP 9 includes the FPolicy enhancements described in the following sections.

Filtering controls

New filters are available for `SetAttr` and for removing notifications on directory activities.

Async resiliency

If an FPolicy server operating in asynchronous mode experiences a network outage, FPolicy notifications generated during the outage are stored on the storage node. When the FPolicy server comes back online, it is alerted of the stored notifications and can fetch them from the storage node. The length of time the notifications can be stored during an outage is configurable up to 10 minutes.

LIF security

A LIF is an IP address or worldwide port name (WWPN) with associated characteristics, such as a role, a home port, a home node, a list of ports to fail over to, and a firewall

policy. You can configure LIFs on ports over which the cluster sends and receives communications over the network. It is critical to understand the security characteristics of each LIF role.

LIF roles

LIF roles can be the following:

- **Data LIF:** A LIF associated with an SVM and used for communicating with clients.
- **Cluster LIF:** A LIF used to carry intracluster traffic between nodes in a cluster.
- **Node management LIF:** A LIF that provides a dedicated IP address for managing a particular node in a cluster.
- **Cluster management LIF:** A LIF that provides a single management interface for the entire cluster.
- **Intercluster LIF:** A LIF used for cross-cluster communication, backup, and replication.

Security characteristics of each LIF role

	Data LIF	Cluster LIF	Node management LIF	Cluster Management LIF	Intercluster LIF
Requires private IP subnet?	No	Yes	No	No	No
Requires secure network?	No	Yes	No	No	Yes
Default firewall policy	Very restrictive	Completely open	Medium	Medium	Very restrictive
Is the firewall customizable?	Yes	No	Yes	Yes	Yes



- Because the cluster LIF is completely open with no configurable firewall policy, it must be on a private IP subnet on a secure isolated network.
- Under no circumstance should any LIF roles be exposed to the internet.

Learn more about securing LIFs, see the [Configure firewall policies for LIFs](#).

Protocol and port security

In addition to performing on-box security operations and functions, the hardening of a solution must also include off-box security mechanisms. Leveraging additional infrastructure devices, such as firewalls, intrusion prevention systems (IPSs), and other security devices, for filtering and limiting access to ONTAP is an effective way to establish and maintain a stringent security posture. This information is a key component for filtering and limiting access to the environment and its resources.

Commonly used protocols and ports

Service	Port/Protocol	Description
SSH	22/TCP	SSH login
telnet	23/TCP	Remote login
Domain	53/TCP	Domain Name Server
HTTP	80/TCP 80/UDP	HTTP
rpcbind	111/TCP 111/UDP	Remote procedure call
NTP	123/UDP	Network Time Protocol
msrpc	135/UDP	Microsoft Remote Procedure Call
Netbios-name	137/TCP 137/UDP	NetBIOS name service
netbios-ssn	139/TCP	NetBIOS service session
SNMP	161/UDP	SNMP
HTTPS	443/TCP	Secure xref:./ontap-security-hardening/http
microsoft-ds	445/TCP	Microsoft directory services
IPsec	500/UDP	Internet Protocol Security
mount	635/UDP	NFS mount
named	953/UDP	Name daemon
NFS	2049/UDP 2049/TCP	NFS server daemon
nrv	2050/TCP	NetApp remote volume protocol
iscsi	3260/TCP	iSCSI target port
Lockd	4045/TCP 4045/UDP	NFS lock daemon
NFS	4046/TCP	NFS mountd protocol
acp-proto	4046/UDP	Accounting protocol
rquotad	4049/UDP	NFS rquotad protocol
krb524	4444/UDP	Kerberos 524
IPsec	4500/UDP	Internet Protocol Security

Service	Port/Protocol	Description
acp	5125/UDP 5133/UDP 5144/TCP	Alternate control port for disk
Mdns	5353/UDP	Multicast DNS
HTTPS	5986/UDP	HTTPS port: listening binary protocol
TELNET	8023/TCP	Node-scope Telnet
HTTPS	8443/TCP	7MTT GUI tool through xref:./ontap-security-hardening/httpS
RSH	8514/TCP	Node-scope RSH
KMIP	9877/TCP	KMIP client port (internal local host only)
ndmp	10000/TCP	NDMP
cifs witness port	40001/TCP	CIFS witness port
TLS	50000/TCP	Transport layer security
Iscsi	65200/TCP	iSCSI port
SSH	65502/TCP	Secure Shell
vsun	65503/TCP	vsun

NetApp internal ports

Port/Protocol	Description
900	NetApp cluster RPC
902	NetApp cluster RPC
904	NetApp cluster RPC
905	NetApp cluster RPC
910	NetApp cluster RPC
911	NetApp cluster RPC
913	NetApp cluster RPC
914	NetApp cluster RPC
915	NetApp cluster RPC
918	NetApp cluster RPC
920	NetApp cluster RPC
921	NetApp cluster RPC
924	NetApp cluster RPC
925	NetApp cluster RPC

Port/Protocol	Description
927	NetApp cluster RPC
928	NetApp cluster RPC
929	NetApp cluster RPC
931	NetApp cluster RPC
932	NetApp cluster RPC
933	NetApp cluster RPC
934	NetApp cluster RPC
935	NetApp cluster RPC
936	NetApp cluster RPC
937	NetApp cluster RPC
939	NetApp cluster RPC
940	NetApp cluster RPC
951	NetApp cluster RPC
954	NetApp cluster RPC
955	NetApp cluster RPC
956	NetApp cluster RPC
958	NetApp cluster RPC
961	NetApp cluster RPC
963	NetApp cluster RPC
964	NetApp cluster RPC
966	NetApp cluster RPC
967	NetApp cluster RPC
7810	NetApp cluster RPC
7811	NetApp cluster RPC
7812	NetApp cluster RPC
7813	NetApp cluster RPC
7814	NetApp cluster RPC
7815	NetApp cluster RPC
7816	NetApp cluster RPC
7817	NetApp cluster RPC
7818	NetApp cluster RPC
7819	NetApp cluster RPC
7820	NetApp cluster RPC

Port/Protocol	Description
7821	NetApp cluster RPC
7822	NetApp cluster RPC
7823	NetApp cluster RPC
7824	NetApp cluster RPC

Audit NAS events on SVMs

SMB and NFS auditing and security tracing

You can use the file access auditing features available for the SMB and NFS protocols with ONTAP, such as native auditing and file policy management using FPolicy.

You should design and implement auditing of SMB and NFS file access events under the following circumstances:

- Basic SMB and NFS protocol file access has been configured.
- You want to create and maintain an auditing configuration using one of the following methods:
 - Native ONTAP functionality
 - External FPolicy servers

Audit NAS events on SVMs

Auditing for NAS events is a security measure that enables you to track and log certain SMB and NFS events on storage virtual machines (SVMs). This helps you track potential security problems and provides evidence of any security breaches. You can also stage and audit Active Directory central access policies to see what the result of implementing them would be.

SMB events

You can audit the following events:

- SMB file and folder access events

You can audit SMB file and folder access events on objects stored on FlexVol volumes belonging to the auditing-enabled SVMs.

- SMB logon and logoff events

You can audit SMB logon and logoff events for SMB servers on SVMs.

- Central access policy staging events

You can audit the effective access of objects on SMB servers using permissions applied through proposed central access policies. Auditing through the staging of central access policies enables you to see what the effects are of central access policies before they are deployed.

Auditing of central access policy staging is set up using Active Directory GPOs; however, the SVM auditing configuration must be configured to audit central access policy staging events.

Although you can enable central access policy staging in the auditing configuration without enabling Dynamic Access Control on the SMB server, central access policy staging events are generated only if Dynamic Access Control is enabled. Dynamic Access Control is enabled through a SMB server option. It is not enabled by default.

NFS events

You can audit file and directory events by utilizing NFSv4 ACL's on objects stored on SVMs.

How auditing works

Basic auditing concepts

To understand auditing in ONTAP, you should be aware of some basic auditing concepts.

- **Staging files**

The intermediate binary files on individual nodes where audit records are stored prior to consolidation and conversion. Staging files are contained in staging volumes.

- **Staging volume**

A dedicated volume created by ONTAP to store staging files. There is one staging volume per aggregate. Staging volumes are shared by all audit-enabled storage virtual machines (SVMs) to store audit records of data access for data volumes in that particular aggregate. Each SVM's audit records are stored in a separate directory within the staging volume.

Cluster administrators can view information about staging volumes, but most other volume operations are not permitted. Only ONTAP can create staging volumes. ONTAP automatically assigns a name to staging volumes. All staging volume names begin with `MDV_aud_` followed by the UUID of the aggregate containing that staging volume (for example: `MDV_aud_1d0131843d4811e296fc123478563412.`)

- **System volumes**

A FlexVol volume that contains special metadata, such as metadata for file services audit logs. The admin SVM owns system volumes, which are visible across the cluster. Staging volumes are a type of system volume.

- **Consolidation task**

A task that gets created when auditing is enabled. This long-running task on each SVM takes the audit records from staging files across the member nodes of the SVM. This task merges the audit records in sorted chronological order, and then converts them to a user-readable event log format specified in the auditing configuration—either the EVTX or XML file format. The converted event logs are stored in the audit event log directory that is specified in the SVM auditing configuration.

How the ONTAP auditing process works

The ONTAP auditing process is different from the Microsoft auditing process. Before you configure auditing, you should understand how the ONTAP auditing process works.

Audit records are initially stored in binary staging files on individual nodes. If auditing is enabled on an SVM, every member node maintains staging files for that SVM. Periodically, they are consolidated and converted to

user-readable event logs, which are stored in the audit event log directory for the SVM.

Process when auditing is enabled on an SVM

Auditing can only be enabled on SVMs. When the storage administrator enables auditing on the SVM, the auditing subsystem checks whether staging volumes are present. A staging volume must exist for each aggregate that contains data volumes owned by the SVM. The auditing subsystem creates any needed staging volumes if they do not exist.

The auditing subsystem also completes other prerequisite tasks before auditing is enabled:

- The auditing subsystem verifies that the log directory path is available and does not contain symlinks.

The log directory must already exist as a path within the SVM's namespace. It is recommended to create a new volume or qtree to hold the audit log files. The auditing subsystem does not assign a default log file location. If the log directory path specified in the auditing configuration is not a valid path, auditing configuration creation fails with the `The specified path "/path" does not exist in the namespace belonging to Vserver "Vserver_name" error.`

Configuration creation fails if the directory exists but contains symlinks.

- Auditing schedules the consolidation task.

After this task is scheduled, auditing is enabled. The SVM auditing configuration and the log files persist across a reboot or if the NFS or SMB servers are stopped or restarted.

Event log consolidation

Log consolidation is a scheduled task that runs on a routine basis until auditing is disabled. When auditing is disabled, the consolidation task verifies that all of the remaining logs are consolidated.

Guaranteed auditing

By default, auditing is guaranteed. ONTAP guarantees that all auditable file access events (as specified by configured audit policy ACLs) are recorded, even if a node is unavailable. A requested file operation cannot be completed until the audit record for that operation is saved to the staging volume on persistent storage. If audit records cannot be committed to the disk in the staging files, either because of insufficient space or because of other issues, client operations are denied.



An administrator, or account user with privilege level access, can bypass the file audit logging operation by using NetApp Manageability SDK or REST APIs. You can determine if any file actions have been taken using NetApp Manageability SDK or REST APIs by reviewing the command history logs stored in the `audit.log` file.

For more information about command history audit logs, see the "Managing audit logging for management activities" section in [System administration](#).

Consolidation process when a node is unavailable

If a node containing volumes belonging to an SVM with auditing enabled is unavailable, the behavior of the auditing consolidation task depends on whether the node's storage failover (SFO) partner (or the HA partner in the case of a two-node cluster) is available:

- If the staging volume is available through the SFO partner, the staging volumes last reported from the node

are scanned, and consolidation proceeds normally.

- If the SFO partner is not available, the task creates a partial log file.

When a node is not reachable, the consolidation task consolidates the audit records from the other available nodes of that SVM. To identify that it is not complete, the task adds the suffix `.partial` to the consolidated file name.

- After the unavailable node is available, the audit records in that node are consolidated with the audit records from the other nodes at that time.
- All audit records are preserved.

Event log rotation

Audit event log files are rotated when they reach a configured threshold log size or on a configured schedule. When an event log file is rotated, the scheduled consolidation task first renames the active converted file to a time-stamped archive file, and then creates a new active converted event log file.

Process when auditing is disabled on the SVM

When auditing is disabled on the SVM, the consolidation task is triggered one final time. All outstanding, recorded audit records are logged in a user-readable format. Existing event logs stored in the event log directory are not deleted when auditing is disabled on the SVM and are available for viewing.

After all existing staging files for that SVM are consolidated, the consolidation task is removed from the schedule. Disabling the auditing configuration for the SVM does not remove the auditing configuration. A storage administrator can reenabling auditing at any time.

The auditing consolidation job, which gets created when auditing is enabled, monitors the consolidation task and re-creates it if the consolidation task exits because of an error. Users cannot delete the auditing consolidation job.

Auditing requirements and considerations

Before you configure and enable auditing on your storage virtual machine (SVM), you need to be aware of certain requirements and considerations.

- The maximum number of auditing-enabled SVMs supported depends on your version of ONTAP:

ONTAP version	Maximum
9.8 and earlier	50
9.9.1 and later	400

- Auditing is not tied to SMB or NFS licensing.

You can configure and enable auditing even if SMB and NFS licenses are not installed on the cluster.

- NFS auditing supports security ACEs (type U).
- For NFS auditing, there is no mapping between mode bits and auditing ACEs.

When converting ACLs to mode bits, auditing ACEs are skipped. When converting mode bits to ACLs, auditing ACEs are not generated.

- The directory specified in the auditing configuration must exist.

If it does not exist, the command to create the auditing configuration fails.

- The directory specified in the auditing configuration must meet the following requirements:
 - The directory must not contain symbolic links.

If the directory specified in the auditing configuration contains symbolic links, the command to create the auditing configuration fails.

- You must specify the directory by using an absolute path.

You should not specify a relative path, for example, `/vs1/././`.

- Auditing is dependent on having available space in the staging volumes.

You must be aware of and have a plan for ensuring that there is sufficient space for the staging volumes in aggregates that contain audited volumes.

- Auditing is dependent on having available space in the volume containing the directory where converted event logs are stored.

You must be aware of and have a plan for ensuring that there is sufficient space in the volumes used to store event logs. You can specify the number of event logs to retain in the auditing directory by using the `-rotate-limit` parameter when creating an auditing configuration, which can help to ensure that there is enough available space for the event logs in the volume.

- Although you can enable central access policy staging in the auditing configuration without enabling Dynamic Access Control on the SMB server, Dynamic Access Control must be enabled to generate central access policy staging events.

Dynamic Access Control is not enabled by default.

Aggregate space considerations when enabling auditing

When an auditing configuration is created and auditing is enabled on at least one storage virtual machine (SVM) in the cluster, the auditing subsystem creates staging volumes on all existing aggregates and on all new aggregates that are created. You need to be aware of certain aggregate space considerations when you enable auditing on the cluster.

Staging volume creation might fail due to non-availability of space in an aggregate. This might happen if you create an auditing configuration and existing aggregates do not have enough space to contain the staging volume.

You should ensure that there is enough space on existing aggregates for the staging volumes before enabling auditing on an SVM.

Limitations for the size of audit records on staging files

The size of an audit record on a staging file cannot be greater than 32 KB.

When large audit records can occur

Large audit records might occur during management auditing in one of the following scenarios:

- Adding or deleting users to or from groups with a large number of users.
- Adding or deleting a file-share access control list (ACL) on a file-share with a large number of file-share users.
- Other scenarios.

Disable management auditing to avoid this issue. To do this, modify the audit configuration and remove the following from the list of audit event types:

- file-share
- user-account
- security-group
- authorization-policy-change

After removal, they will not be audited by the file services auditing subsystem.

The effects of audit records that are too large

- If the size of an audit record is too large (over 32 KB), the audit record is not created and the auditing subsystem generates an event management system (EMS) message similar to the following:

```
File Services Auditing subsystem failed the operation or truncated an audit record because it was greater than max_audit_record_size value. Vserver UUID=%s, event_id=%u, size=%u
```

If auditing is guaranteed, the file operation fails because its audit record cannot be created.

- If the size of the audit record is more than 9,999 bytes, the same EMS message as above is displayed. A partial audit record is created with the larger key value missing.
- If the audit record exceeds 2,000 characters, the following error message shows instead of the actual value:

```
The value of this field was too long to display.
```

What the supported audit event log formats are

Supported file formats for the converted audit event logs are `EVTX` and `XML` file formats.

You can specify the type of file format when you create the auditing configuration. By default, ONTAP converts the binary logs to the `EVTX` file format.

View audit event logs

You can use audit event logs to determine whether you have adequate file security and whether there have been improper file and folder access attempts. You can view and process audit event logs saved in the `EVTX` or `XML` file formats.

- **EVTX file format**

You can open the converted `EVTX` audit event logs as saved files using Microsoft Event Viewer.

There are two options that you can use when viewing event logs using Event Viewer:

- **General view**

Information that is common to all events is displayed for the event record. In this version of ONTAP, the event-specific data for the event record is not displayed. You can use the detailed view to display event-specific data.

- **Detailed view**

A friendly view and an XML view are available. The friendly view and the XML view display both the information that is common to all events and the event-specific data for the event record.

- **XML file format**

You can view and process `XML` audit event logs on third-party applications that support the `XML` file format. `XML` viewing tools can be used to view the audit logs provided you have the `XML` schema and information about definitions for the `XML` fields. For more information about the `XML` schema and definitions, see the [ONTAP Auditing Schema Reference](#).

How active audit logs are viewed using Event Viewer

If the audit consolidation process is running on the cluster, the consolidation process appends new records to the active audit log file for audit-enabled storage virtual machines (SVMs). This active audit log can be accessed and opened over an SMB share in Microsoft Event Viewer.

In addition to viewing existing audit records, Event Viewer has a refresh option that enables you to refresh the content in the console window. Whether the newly appended logs are viewable in Event Viewer depends on whether oplocks are enabled on the share used to access the active audit log.

Oplocks setting on the share	Behavior
Enabled	Event Viewer opens the log that contains events written to it up to that point in time. The refresh operation does not refresh the log with new events appended by the consolidation process.
Disabled	Event Viewer opens the log that contains events written to it up to that point in time. The refresh operation refreshes the log with new events appended by the consolidation process.



This information is applicable only for `EVTX` event logs. `XML` event logs can be viewed through SMB in a browser or through NFS using any `XML` editor or viewer.

SMB events that can be audited

SMB events that can be audited overview

ONTAP can audit certain SMB events, including certain file and folder access events, certain logon and logoff events, and central access policy staging events. Knowing which access events can be audited is helpful when interpreting results from the event logs.

The following additional SMB events can be audited in ONTAP 9.2 and later:

Event ID (EVT/EVTX)	Event	Description	Category
4670	Object permissions were changed	OBJECT ACCESS: Permissions changed.	File Access
4907	Object auditing settings were changed	OBJECT ACCESS: Audit settings changed.	File Access
4913	Object Central Access Policy was changed	OBJECT ACCESS: CAP changed.	File Access

The following SMB events can be audited in ONTAP 9.0 and later:

Event ID (EVT/EVTX)	Event	Description	Category
540/4624	An account was successfully logged on	LOGON/LOGOFF: Network (SMB) logon.	Logon and Logoff
529/4625	An account failed to log on	LOGON/LOGOFF: Unknown user name or bad password.	Logon and Logoff
530/4625	An account failed to log on	LOGON/LOGOFF: Account logon time restriction.	Logon and Logoff
531/4625	An account failed to log on	LOGON/LOGOFF: Account currently disabled.	Logon and Logoff
532/4625	An account failed to log on	LOGON/LOGOFF: User account has expired.	Logon and Logoff
533/4625	An account failed to log on	LOGON/LOGOFF: User cannot log on to this computer.	Logon and Logoff
534/4625	An account failed to log on	LOGON/LOGOFF: User not granted logon type here.	Logon and Logoff
535/4625	An account failed to log on	LOGON/LOGOFF: User's password has expired.	Logon and Logoff

537/4625	An account failed to log on	LOGON/LOGOFF: Logon failed for reasons other than above.	Logon and Logoff
539/4625	An account failed to log on	LOGON/LOGOFF: Account locked out.	Logon and Logoff
538/4634	An account was logged off	LOGON/LOGOFF: Local or network user logoff.	Logon and Logoff
560/4656	Open Object/Create Object	OBJECT ACCESS: Object (file or directory) open.	File Access
563/4659	Open Object with the Intent to Delete	OBJECT ACCESS: A handle to an object (file or directory) was requested with the Intent to Delete.	File Access
564/4660	Delete Object	OBJECT ACCESS: Delete Object (file or directory). ONTAP generates this event when a Windows client attempts to delete the object (file or directory).	File Access
567/4663	Read Object/Write Object/Get Object Attributes/Set Object Attributes	OBJECT ACCESS: Object access attempt (read, write, get attribute, set attribute). Note: For this event, ONTAP audits only the first SMB read and first SMB write operation (success or failure) on an object. This prevents ONTAP from creating excessive log entries when a single client opens an object and performs many successive read or write operations to the same object.	File Access
NA/4664	Hard link	OBJECT ACCESS: An attempt was made to create a hard link.	File Access
NA/4818	Proposed central access policy does not grant the same access permissions as the current central access policy	OBJECT ACCESS: Central Access Policy Staging.	File Access
NA/NA Data ONTAP Event ID 9999	Rename Object	OBJECT ACCESS: Object renamed. This is an ONTAP event. It is not currently supported by Windows as a single event.	File Access

NA/NA Data ONTAP Event ID 9998	Unlink Object	OBJECT ACCESS: Object unlinked. This is an ONTAP event. It is not currently supported by Windows as a single event.	File Access
-----------------------------------	---------------	------------------------------------------------------------------------------------------------------------------------	-------------

Additional information about Event 4656

The `HandleID` tag in the audit XML event contains the handle of the object (file or directory) accessed. The `HandleID` tag for the EVT 4656 event contains different information depending on whether the open event is for creating a new object or for opening an existing object:

- If the open event is an open request to create a new object (file or directory), the `HandleID` tag in the audit XML event shows an empty `HandleID` (for example: `<Data Name="HandleID">00000000000000;00;00000000;00000000</Data>`).

The `HandleID` is empty because the OPEN (for creating a new object) request gets audited before the actual object creation happens and before a handle exists. Subsequent audited events for the same object have the right object handle in the `HandleID` tag.

- If the open event is an open request to open an existing object, the audit event will have the assigned handle of that object in the `HandleID` tag (for example: `<Data Name="HandleID">000000000000401;00;000000ea;00123ed4</Data>`).

Determine what the complete path to the audited object is

The object path printed in the `<ObjectName>` tag for an audit record contains the name of the volume (in parentheses) and the relative path from the root of the containing volume. If you want to determine the complete path of the audited object, including the junction path, there are certain steps you must take.

Steps

1. Determine what the volume name and relative path to audited object is by looking at the `<ObjectName>` tag in the audit event.

In this example, the volume name is "data1" and the relative path to the file is `/dir1/file.txt`:

```
<Data Name="ObjectName"> (data1);/dir1/file.txt </Data>
```

2. Using the volume name determined in the previous step, determine what the junction path is for the volume containing the audited object:

In this example, the volume name is "data1" and the junction path for the volume containing the audited object is `/data/data1`:

```
volume show -junction -volume data1
```

Vserver	Volume	Language	Junction Active	Junction Path	Junction Path Source
vs1	data1	en_US.UTF-8	true	/data/data1	RW_volume

- Determine the full path to the audited object by appending the relative path found in the `<ObjectName>` tag to the junction path for the volume.

In this example, the junction path for the volume:

```
/data/data1/dir1/file.text
```

Considerations when auditing symlinks and hard links

There are certain considerations you must keep in mind when auditing symlinks and hard links.

An audit record contains information about the object being audited including the path to the audited object, which is identified in the `ObjectName` tag. You should be aware of how paths for symlinks and hard links are recorded in the `ObjectName` tag.

Symlinks

A symlink is a file with a separate inode that contains a pointer to the location of a destination object, known as the target. When accessing an object through a symlink, ONTAP automatically interprets the symlink and follows the actual canonical protocol agnostic path to the target object in the volume.

In the following example output, there are two symlinks, both pointing to a file named `target.txt`. One of the symlinks is a relative symlink and one is an absolute symlink. If either of the symlinks are audited, the `ObjectName` tag in the audit event contains the path to the file `target.txt`:

```
[root@host1 audit]# ls -l
total 0
lrwxrwxrwx 1 user1 group1 37 Apr  2 10:09 softlink_fullpath.txt ->
/data/audit/target.txt
lrwxrwxrwx 1 user1 group1 10 Apr  2 09:54 softlink.txt -> target.txt
-rwxrwxrwx 1 user1 group1 16 Apr  2 10:05 target.txt
```

Hard links

A hard link is a directory entry that associates a name with an existing file on a file system. The hard link points to the inode location of the original file. Similar to how ONTAP interprets symlinks, ONTAP interprets the hard link and follows the actual canonical path to the target object in the volume. When access to a hard link object is audited, the audit event records this absolute canonical path in the `ObjectName` tag rather than the hard link path.

Considerations when auditing alternate NTFS data streams

There are certain considerations you must keep in mind when auditing files with NTFS alternate data streams.

The location of an object being audited is recorded in an event record using two tags, the `ObjectName` tag (the path) and the `HandleID` tag (the handle). To properly identify which stream requests are being logged, you must be aware of what ONTAP records in these fields for NTFS alternate data streams:

- EVTX ID: 4656 events (open and create audit events)
 - The path of the alternate data stream is recorded in the `ObjectName` tag.
 - The handle of the alternate data stream is recorded in the `HandleID` tag.
- EVTX ID: 4663 events (all other audit events, such as read, write, getattr, and so on)
 - The path of the base file, not the alternate data stream, is recorded in the `ObjectName` tag.
 - The handle of the alternate data stream is recorded in the `HandleID` tag.

Example

The following example illustrates how to identify EVTX ID: 4663 events for alternate data streams using the `HandleID` tag. Even though the `ObjectName` tag (path) recorded in the read audit event is to the base file path, the `HandleID` tag can be used to identify the event as an audit record for the alternate data stream.

Stream file names take the form `base_file_name:stream_name`. In this example, the `dir1` directory contains a base file with an alternate data stream having the following paths:

```
/dir1/file1.txt  
/dir1/file1.txt:stream1
```



The output in the following event example is truncated as indicated; the output does not display all of the available output tags for the events.

For an EVTX ID 4656 (open audit event), the audit record output for the alternate data stream records the alternate data stream name in the `ObjectName` tag:

```

- <Event>
- <System>
  <Provider Name="Netapp-Security-Auditing" />
  <EventID>4656</EventID>
  <EventName>Open Object</EventName>
  [...]
</System>
- <EventData>
  [...]
  **<Data Name="ObjectType"\>Stream</Data\>
  <Data Name="HandleID"\>00000000000401;00;000001e4;00176767</Data\>
  <Data Name="ObjectName"\>\(data1\);/dir1/file1.txt:stream1</Data\>
  **
  [...]
</EventData>
</Event>
- <Event>

```

For an EVT X ID 4663 (read audit event), the audit record output for the same alternate data stream records the base file name in the `ObjectName` tag; however, the handle in the `HandleID` tag is the alternative data stream's handle and can be used to correlate this event with the alternative data stream:

```

- <Event>
- <System>
  <Provider Name="Netapp-Security-Auditing" />
  <EventID>4663</EventID>
  <EventName>Read Object</EventName>
  [...]
</System>
- <EventData>
  [...]
  **<Data Name="ObjectType"\>Stream</Data\>
  <Data Name="HandleID"\>00000000000401;00;000001e4;00176767</Data\>
  <Data Name="ObjectName"\>\(data1\);/dir1/file1.txt</Data\> **
  [...]
</EventData>
</Event>
- <Event>

```

NFS file and directory access events that can be audited

ONTAP can audit certain NFS file and directory access events. Knowing what access events can be audited is helpful when interpreting results from the converted audit event logs.

You can audit the following NFS file and directory access events:

- READ
- OPEN
- CLOSE
- REaddir
- WRITE
- SETATTR
- CREATE
- LINK
- OPENATTR
- REMOVE
- GETATTR
- VERIFY
- NVERIFY
- RENAME

To reliably audit NFS RENAME events, you should set audit ACEs on directories instead of files because file permissions are not checked for a RENAME operation if the directory permissions are sufficient.

Plan the auditing configuration

Before you configure auditing on storage virtual machines (SVMs), you must understand which configuration options are available and plan the values that you want to set for each option. This information can help you configure the auditing configuration that meets your business needs.

There are certain configuration parameters that are common to all auditing configurations.

Additionally, there are certain parameters that you can use to specify which methods are used when rotating the consolidated and converted audit logs. You can specify one of the three following methods when you configure auditing:

- Rotate logs based on log size
This is the default method used to rotate logs.
- Rotate logs based on a schedule
- Rotate logs based on log size and schedule (whichever event occurs first)



At least one of the methods for log rotation should always be set.

Parameters common to all auditing configurations

There are two required parameters that you must specify when you create the auditing configuration. There are also three optional parameters that you can specify:

Type of information	Option	Required	Include	Your values
<p><i>SVM name</i></p> <p>Name of the SVM on which to create the auditing configuration. The SVM must already exist.</p>	<p><code>-vserver vserver_name</code></p>	<p>Yes</p>	<p>Yes</p>	
<p><i>Log destination path</i></p> <p>Specifies the directory where the converted audit logs are stored, typically a dedicated volume or qtree. The path must already exist in the SVM namespace.</p> <p>The path can be up to 864 characters in length and must have read-write permissions.</p> <p>If the path is not valid, the audit configuration command fails.</p> <p>If the SVM is an SVM disaster recovery source, the log destination path cannot be on the root volume. This is because root volume content is not replicated to the disaster recovery destination.</p> <p>You cannot use a FlexCache volume as a log destination (ONTAP 9.7 and later).</p>	<p><code>-destination text</code></p>	<p>Yes</p>	<p>Yes</p>	

<p><i>Categories of events to audit</i></p> <p>Specifies the categories of events to audit. The following event categories can be audited:</p> <ul style="list-style-type: none"> • File access events (both SMB and NFSv4) • SMB logon and logoff events • Central access policy staging events <p>Central access policy staging events are available beginning with Windows 2012 Active Directory domains.</p> <ul style="list-style-type: none"> • File share category events • Audit policy change events • Local user account management events • Security group management events • Authorization policy change events <p>The default is to audit file access and SMB logon and logoff events.</p> <p>Note: Before you can specify <code>cap-staging</code> as an event category, a SMB server must exist on the SVM. Although you can enable central access policy staging in the auditing configuration without enabling Dynamic Access Control on the SMB server, central access policy staging events are generated only if Dynamic Access Control is enabled. Dynamic Access Control is enabled through a SMB server option. It is not enabled by default.</p>	<pre>-events {file-ops cifs- logon-logoff cap- staging file-share audit-policy-change user-account security- group authorization- policy-change}</pre>	No		
<p><i>Log file output format</i></p> <p>Determines the output format of the audit logs. The output format can be either ONTAP-specific XML or Microsoft Windows EVTX log format. By default, the output format is EVTX.</p>	<pre>-format {xml evtx}</pre>	No		

<p><i>Log files rotation limit</i></p> <p>Determines how many audit log files to retain before rotating the oldest log file out. For example, if you enter a value of 5, the last five log files are retained.</p> <p>A value of 0 indicates that all the log files are retained. The default value is 0.</p>	<p><code>-rotate-limit integer</code></p>	<p>No</p>		
---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-------------------------------------------	-----------	--	--

Parameters used for determining when to rotate audit event logs

Rotate logs based on log size

The default is to rotate audit logs based on size.

- The default log size is 100 MB
- If you want to use the default log rotation method and the default log size, you do not need to configure any specific parameters for log rotation.
- If you want to rotate the audit logs based on a log size alone, use the following command to unset the `-rotate-schedule-minute` parameter: `vserver audit modify -vserver vs0 -destination / -rotate-schedule-minute -`

If you do not want to use the default log size, you can configure the `-rotate-size` parameter to specify a custom log size:

Type of information	Option	Required	Include	Your values
<p><i>Log file size limit</i></p> <p>Determines the audit log file size limit.</p>	<p><code>-rotate-size {integer[KB MB GB TB PB]}</code></p>	<p>No</p>		

Rotate logs based on a schedule

If you choose to rotate the audit logs based on a schedule, you can schedule log rotation by using the time-based rotation parameters in any combination.

- If you use time-based rotation, the `-rotate-schedule-minute` parameter is mandatory.
- All other time-based rotation parameters are optional.
- The rotation schedule is calculated by using all the time-related values.

For example, if you specify only the `-rotate-schedule-minute` parameter, the audit log files are rotated based on the minutes specified on all days of the week, during all hours on all months of the year.

- If you specify only one or two time-based rotation parameters (for example, `-rotate-schedule-month` and `-rotate-schedule-minutes`), the log files are rotated based on the minute values that you specified on all days of the week, during all hours, but only during the specified months.

For example, you can specify that the audit log is to be rotated during the months January, March, and

August on all Mondays, Wednesdays, and Saturdays at 10:30 a.m.

- If you specify values for both `-rotate-schedule-dayofweek` and `-rotate-schedule-day`, they are considered independently.

For example, if you specify `-rotate-schedule-dayofweek` as Friday and `-rotate-schedule-day` as 13, then the audit logs would be rotated on every Friday and on the 13th day of the specified month, not just on every Friday the 13th.

- If you want to rotate the audit logs based on a schedule alone, use the following command to unset the `-rotate-size` parameter: `vserver audit modify -vserver vs0 -destination / -rotate -size -`

You can use the following list of available auditing parameters to determine what values to use for configuring a schedule for audit event log rotations:

Type of information	Option	Required	Include	Your values
<p><i>Log rotation schedule: Month</i></p> <p>Determines the monthly schedule for rotating audit logs.</p> <p>Valid values are January through December, and all. For example, you can specify that the audit log is to be rotated during the months January, March, and August.</p>	<p><code>-rotate-schedule-month</code> <code>chron_month</code></p>	No		
<p><i>Log rotation schedule: Day of week</i></p> <p>Determines the daily (day of week) schedule for rotating audit logs.</p> <p>Valid values are Sunday through Saturday, and all. For example, you can specify that the audit log is to be rotated on Tuesdays and Fridays, or during all the days of a week.</p>	<p><code>-rotate-schedule</code> <code>-dayofweek</code> <code>chron_dayofweek</code></p>	No		
<p><i>Log rotation schedule: Day</i></p> <p>Determines the day of the month schedule for rotating the audit log.</p> <p>Valid values range from 1 through 31. For example, you can specify that the audit log is to be rotated on the 10th and 20th days of a month, or all days of a month.</p>	<p><code>-rotate-schedule-day</code> <code>chron_dayofmonth</code></p>	No		

<p><i>Log rotation schedule: Hour</i></p> <p>Determines the hourly schedule for rotating the audit log.</p> <p>Valid values range from 0 (midnight) to 23 (11:00 p.m.). Specifying <code>all</code> rotates the audit logs every hour. For example, you can specify that the audit log is to be rotated at 6 (6 a.m.) and 18 (6 p.m.).</p>	<pre>-rotate-schedule-hour chron_hour</pre>	<p>No</p>		
<p><i>Log rotation schedule: Minute</i></p> <p>Determines the minute schedule for rotating the audit log.</p> <p>Valid values range from 0 to 59. For example, you can specify that the audit log is to be rotated at the 30th minute.</p>	<pre>-rotate-schedule-minute chron_minute</pre>	<p>Yes, if configuring schedule-based log rotation; otherwise, no.</p>		

Rotate logs based on log size and schedule

You can choose to rotate the log files based on log size and a schedule by setting both the `-rotate-size` parameter and the time-based rotation parameters in any combination. For example: if `-rotate-size` is set to 10 MB and `-rotate-schedule-minute` is set to 15, the log files rotate when the log file size reaches 10 MB or on the 15th minute of every hour (whichever event occurs first).

Create a file and directory auditing configuration on SVMs

Create the auditing configuration

Creating a file and directory auditing configuration on your storage virtual machine (SVM) includes understanding the available configuration options, planning the configuration, and then configuring and enabling the configuration. You can then display information about the auditing configuration to confirm that the resultant configuration is the desired configuration.

Before you can begin auditing file and directory events, you must create an auditing configuration on the storage virtual machine (SVM).

Before you begin

If you plan on creating an auditing configuration for central access policy staging, a SMB server must exist on the SVM.

- Although you can enable central access policy staging in the auditing configuration without enabling Dynamic Access Control on the SMB server, central access policy staging events are generated only if Dynamic Access Control is enabled.



Dynamic Access Control is enabled through a SMB server option. It is not enabled by default.

- If the arguments of a field in a command is invalid, for example, invalid entries for fields, duplicate entries, and non-existent entries, then the command fails before the audit phase.

Such failures do not generate an audit record.

About this task

If the SVM is an SVM disaster recovery source, the destination path cannot be on the root volume.

Step

1. Using the information in the planning worksheet, create the auditing configuration to rotate audit logs based on log size or a schedule:

If you want to rotate audit logs by...	Enter...
Log size	<pre>vserver audit create -vserver vserver_name -destination path -events [{file-ops cifs-logon- logoff cap-staging file-share authorization-policy- change user-account security-group authorization- policy-change}] [-format {xml evtx}] [-rotate-limit integer] [-rotate-size {integer[KB MB GB TB PB]}}</pre>
A schedule	<pre>vserver audit create -vserver vserver_name -destination path -events [{file-ops cifs-logon- logoff cap-staging}] [-format {xml evtx}] [-rotate- limit integer] [-rotate-schedule-month chron_month] [- rotate-schedule-dayofweek chron_dayofweek] [-rotate- schedule-day chron_dayofmonth] [-rotate-schedule-hour chron_hour] -rotate-schedule-minute chron_minute</pre> <div style="display: flex; align-items: center; margin-top: 10px;"> <p>The <code>-rotate-schedule-minute</code> parameter is required if you are configuring time-based audit log rotation.</p> </div>

Examples

The following example creates an auditing configuration that audits file operations and SMB logon and logoff events (the default) using size-based rotation. The log format is `EVTX` (the default). The logs are stored in the `/audit_log` directory. The log file size limit is 200 MB. The logs are rotated when they reach 200 MB in size:

```
cluster1::> vserver audit create -vserver vs1 -destination /audit_log
-rotate-size 200MB
```

The following example creates an auditing configuration that audits file operations and SMB logon and logoff events (the default) using size-based rotation. The log format is `EVTX` (the default). The logs are stored in the `/cifs_event_logs` directory. The log file size limit is 100 MB (the default), and the log rotation limit is 5:

```
cluster1::> vserver audit create -vserver vs1 -destination
/cifs_event_logs -rotate-limit 5
```

The following example creates an auditing configuration that audits file operations, CIFS logon and logoff events, and central access policy staging events using time-based rotation. The log format is `EVTX` (the default). The audit logs are rotated monthly, at 12:30 p.m. on all days of the week. The log rotation limit is 5:

```
cluster1::> vserver audit create -vserver vs1 -destination /audit_log
-events file-ops,cifs-logon-logoff,file-share,audit-policy-change,user-
account,security-group,authorization-policy-change,cap-staging -rotate
-schedule-month all -rotate-schedule-dayofweek all -rotate-schedule-hour
12 -rotate-schedule-minute 30 -rotate-limit 5
```

Related information

- [Enable auditing on the SVM](#)
- [Verify the auditing configuration](#)

Enable auditing on the SVM

After you finish setting up the auditing configuration, you must enable auditing on the storage virtual machine (SVM).

Before you begin

The SVM audit configuration must already exist.

About this task

When an SVM disaster recovery ID discard configuration is first started (after the SnapMirror initialization is complete) and the SVM has an auditing configuration, ONTAP automatically disables the auditing configuration. Auditing is disabled on the read-only SVM to prevent the staging volumes from filling up. You can enable auditing only after the SnapMirror relationship is broken and the SVM is read-write.

Steps

1. Enable auditing on the SVM:

```
vserver audit enable -vserver vserver_name
```

```
vserver audit enable -vserver vs1
```

Related information

- [Create the auditing configuration](#)
- [Verify the auditing configuration](#)

Verify the auditing configuration

After completing the auditing configuration, you should verify that auditing is configured properly and is enabled.

Steps

1. Verify the auditing configuration:

```
vserver audit show -instance -vserver vserver_name
```

The following command displays in list form all auditing configuration information for storage virtual machine (SVM) vs1:

```
vserver audit show -instance -vserver vs1
```

```
          Vserver: vs1
          Auditing state: true
          Log Destination Path: /audit_log
Categories of Events to Audit: file-ops
          Log Format: evtX
          Log File Size Limit: 200MB
          Log Rotation Schedule: Month: -
Log Rotation Schedule: Day of Week: -
          Log Rotation Schedule: Day: -
          Log Rotation Schedule: Hour: -
Log Rotation Schedule: Minute: -
          Rotation Schedules: -
          Log Files Rotation Limit: 0
```

Related information

- [Create the auditing configuration](#)
- [Enable auditing on the SVM](#)

Configure file and folder audit policies

Configure file and folder audit policies

Implementing auditing on file and folder access events is a two-step process. First, you must create and enable an auditing configuration on storage virtual machines (SVMs). Second, you must configure audit policies on the files and folders that you want to monitor. You can configure audit policies to monitor both successful and failed access attempts.

You can configure both SMB and NFS audit policies. SMB and NFS audit policies have different configuration requirements and audit capabilities.

If the appropriate audit policies are configured, ONTAP monitors SMB and NFS access events as specified in

the audit policies only if the SMB or NFS servers are running.

Configure audit policies on NTFS security-style files and directories

Before you can audit file and directory operations, you must configure audit policies on the files and directories for which you want to collect audit information. This is in addition to setting up and enabling the audit configuration. You can configure NTFS audit policies by using the Windows Security tab or by using the ONTAP CLI.

Configuring NTFS audit policies using the Windows Security tab

You can configure NTFS audit policies on files and directories by using the **Windows Security** tab in the Windows Properties window. This is the same method used when configuring audit policies on data residing on a Windows client, which enables you to use the same GUI interface that you are accustomed to using.

Before you begin

Auditing must be configured on the storage virtual machine (SVM) that contains the data to which you are applying system access control lists (SACLs).

About this task

Configuring NTFS audit policies is done by adding entries to NTFS SACLs that are associated with an NTFS security descriptor. The security descriptor is then applied to NTFS files and directories. These tasks are automatically handled by the Windows GUI. The security descriptor can contain discretionary access control lists (DACLS) for applying file and folder access permissions, SACLs for file and folder auditing, or both SACLs and DACLS.

To set NTFS audit policies using the Windows Security tab, complete the following steps on a Windows host:

Steps

1. From the **Tools** menu in Windows Explorer, select **Map network drive**.
2. Complete the **Map Network Drive** box:
 - a. Select a **Drive** letter.
 - b. In the **Folder** box, type the SMB server name that contains the share, holding the data you want to audit and the name of the share.

You can specify the IP address of the data interface for the SMB server instead of the SMB server name.

If your SMB server name is "SMB_SERVER" and your share is named "share1", you should enter \\SMB_SERVER\share1.

- c. Click **Finish**.

The drive you selected is mounted and ready with the Windows Explorer window displaying files and folders contained within the share.

3. Select the file or directory for which you want to enable auditing access.
4. Right-click the file or directory, and then select **Properties**.
5. Select the **Security** tab.
6. Click **Advanced**.

7. Select the **Auditing** tab.
8. Perform the desired actions:

If you want to....	Do the following
Set up auditing for a new user or group	<ol style="list-style-type: none"> a. Click Add. b. In the Enter the object name to select box, type the name of the user or group that you want to add. c. Click OK.
Remove auditing from a user or group	<ol style="list-style-type: none"> a. In the Enter the object name to select box, select the user or group that you want to remove. b. Click Remove. c. Click OK. d. Skip the rest of this procedure.
Change auditing for a user or group	<ol style="list-style-type: none"> a. In the Enter the object name to select box, select the user or group that you want to change. b. Click Edit. c. Click OK.

If you are setting up auditing on a user or group or changing auditing on an existing user or group, the Auditing Entry for <object> box opens.

9. In the **Apply to** box, select how you want to apply this auditing entry.

You can select one of the following:

- **This folder, subfolders and files**
- **This folder and subfolders**
- **This folder only**
- **This folder and files**
- **Subfolders and files only**
- **Subfolders only**
- **Files only**

If you are setting up auditing on a single file, the **Apply to** box is not active. The **Apply to** box setting defaults to **This object only**.



Because auditing takes SVM resources, select only the minimal level that provides the auditing events that meet your security requirements.

10. In the **Access** box, select what you want audited and whether you want to audit successful events, failure events, or both.
 - To audit successful events, select the Success box.
 - To audit failure events, select the Failure box.

Select only the actions that you need to monitor to meet your security requirements. For more information about these auditable events, see your Windows documentation. You can audit the following events:

- **Full control**
- **Traverse folder / execute file**
- **List folder / read data**
- **Read attributes**
- **Read extended attributes**
- **Create files / write data**
- **Create folders / append data**
- **Write attributes**
- **Write extended attributes**
- **Delete subfolders and files**
- **Delete**
- **Read permissions**
- **Change permissions**
- **Take ownership**

11. If you do not want the auditing setting to propagate to subsequent files and folders of the original container, select the **Apply these auditing entries to objects and/or containers within this container only** box.
12. Click **Apply**.
13. After you finish adding, removing, or editing auditing entries, click **OK**.

The Auditing Entry for <object> box closes.

14. In the **Auditing** box, select the inheritance settings for this folder.

Select only the minimal level that provides the auditing events that meet your security requirements. You can choose one of the following:

- Select the Include inheritable auditing entries from this object's parent box.
- Select the Replace all existing inheritable auditing entries on all descendants with inheritable auditing entries from this object box.
- Select both boxes.
- Select neither box.

If you are setting SACLS on a single file, the Replace all existing inheritable auditing entries on all descendants with inheritable auditing entries from this object box is not present in the Auditing box.

15. Click **OK**.

The Auditing box closes.

Configure NTFS audit policies using the ONTAP CLI

You can configure audit policies on files and folders using the ONTAP CLI. This enables you to configure NTFS audit policies without needing to connect to the data using an SMB share on a Windows client.

You can configure NTFS audit policies by using the `vserver security file-directory` command family.

You can only configure NTFS SACLs using the CLI. Configuring NFSv4 SACLs is not supported with this ONTAP command family. See the man pages for more information about using these commands to configure and add NTFS SACLs to files and folders.

Configure auditing for UNIX security style files and directories

You configure auditing for UNIX security style files and directories by adding audit ACEs to NFSv4.x ACLs. This allows you to monitor certain NFS file and directory access events for security purposes.

About this task

For NFSv4.x, both discretionary and system ACEs are stored in the same ACL. They are not stored in separate DACLs and SACLs. Therefore, you must exercise caution when adding audit ACEs to an existing ACL to avoid overwriting and losing an existing ACL. The order in which you add the audit ACEs to an existing ACL does not matter.

Steps

1. Retrieve the existing ACL for the file or directory by using the `nfs4_getfacl` or equivalent command.

For more information about manipulating ACLs, see the man pages of your NFS client.

2. Append the desired audit ACEs.
3. Apply the updated ACL to the file or directory by using the `nfs4_setfacl` or equivalent command.

Display information about audit policies applied to files and directories

Display information about audit policies using the Windows Security tab

You can display information about audit policies that have been applied to files and directories by using the Security tab in the Windows Properties window. This is the same method used for data residing on a Windows server, which enables customers to use the same GUI interface that they are accustomed to using.

About this task

Displaying information about audit policies applied to files and directories enables you to verify that you have the appropriate system access control lists (SACLs) set on specified files and folders.

To display information about SACLs that have been applied to NTFS files and folders, complete the following steps on a Windows host.

Steps

1. From the **Tools** menu in Windows Explorer, select **Map network drive**.
2. Complete the **Map Network Drive** dialog box:
 - a. Select a **Drive** letter.
 - b. In the **Folder** box, type the IP address or SMB server name of the storage virtual machine (SVM) containing the share that holds both the data you would like to audit and the name of the share.

If your SMB server name is “SMB_SERVER” and your share is named “share1”, you should enter \\SMB_SERVER\share1.



You can specify the IP address of the data interface for the SMB server instead of the SMB server name.

c. Click **Finish**.

The drive you selected is mounted and ready with the Windows Explorer window displaying files and folders contained within the share.

3. Select the file or directory for which you display auditing information.
4. Right-click on the file or directory, and select **Properties**.
5. Select the **Security** tab.
6. Click **Advanced**.
7. Select the **Auditing** tab.
8. Click **Continue**.

The Auditing box opens. The **Auditing entries** box displays a summary of users and groups that have SACLs applied to them.

9. In the **Auditing entries** box select the user or group whose SACL entries you want displayed.
10. Click **Edit**.

The Auditing entry for <object> box opens.

11. In the **Access** box, view the current SACLs that are applied to the selected object.
12. Click **Cancel** to close the **Auditing entry for <object>** box.
13. Click **Cancel** to close the **Auditing** box.

Display information about NTFS audit policies on FlexVol volumes using the CLI

You can display information about NTFS audit policies on FlexVol volumes, including what the security styles and effective security styles are, what permissions are applied, and information about system access control lists. You can use the information to validate your security configuration or to troubleshoot auditing issues.

About this task

Displaying information about audit policies applied to files and directories enables you to verify that you have the appropriate system access control lists (SACLs) set on specified files and folders.

You must provide the name of the storage virtual machine (SVM) and the path to the files or folders whose audit information you want to display. You can display the output in summary form or as a detailed list.

- NTFS security-style volumes and qtrees use only NTFS system access control lists (SACLs) for audit policies.
- Files and folders in a mixed security-style volume with NTFS effective security can have NTFS audit policies applied to them.

Mixed security-style volumes and qtrees can contain some files and directories that use UNIX file

permissions, either mode bits or NFSv4 ACLs, and some files and directories that use NTFS file permissions.

- The top level of a mixed security-style volume can have either UNIX or NTFS effective security and might or might not contain NTFS SACLs.
- Because Storage-Level Access Guard security can be configured on a mixed security-style volume or qtree even if the effective security style of the volume root or qtree is UNIX, the output for a volume or qtree path where Storage-Level Access Guard is configured might display both regular file and folder NFSv4 SACLs and Storage-Level Access Guard NTFS SACLs.
- If the path that is entered in the command is to data with NTFS effective security, the output also displays information about Dynamic Access Control ACEs if Dynamic Access Control is configured for the given file or directory path.
- When displaying security information about files and folders with NTFS effective security, UNIX-related output fields contain display-only UNIX file permission information.

NTFS security-style files and folders use only NTFS file permissions and Windows users and groups when determining file access rights.

- ACL output is displayed only for files and folders with NTFS or NFSv4 security.

This field is empty for files and folders using UNIX security that have only mode bit permissions applied (no NFSv4 ACLs).

- The owner and group output fields in the ACL output apply only in the case of NTFS security descriptors.

Step

1. Display file and directory audit policy settings with the desired level of detail:

If you want to display information...	Enter the following command...
In summary form	<pre>vserver security file-directory show -vserver vserver_name -path path</pre>
As a detailed list	<pre>vserver security file-directory show -vserver vserver_name -path path -expand-mask true</pre>

Examples

The following example displays the audit policy information for the path `/corp` in SVM `vs1`. The path has NTFS effective security. The NTFS security descriptor contains both a SUCCESS and a SUCCESS/FAIL SACL entry.

```

cluster::> vserver security file-directory show -vserver vs1 -path /corp
      Vserver: vs1
      File Path: /corp
      File Inode Number: 357
      Security Style: ntfs
      Effective Style: ntfs
      DOS Attributes: 10
      DOS Attributes in Text: ----D---
      Expanded Dos Attributes: -
      Unix User Id: 0
      Unix Group Id: 0
      Unix Mode Bits: 777
      Unix Mode Bits in Text: rwxrwxrwx
      ACLs: NTFS Security Descriptor
      Control:0x8014
      Owner:DOMAIN\Administrator
      Group:BUILTIN\Administrators
      SACL - ACEs
      ALL-DOMAIN\Administrator-0x100081-OI|CI|SA|FA
      SUCCESSFUL-DOMAIN\user1-0x100116-OI|CI|SA
      DACL - ACEs
      ALLOW-BUILTIN\Administrators-0x1f01ff-OI|CI
      ALLOW-BUILTIN\Users-0x1f01ff-OI|CI
      ALLOW-CREATOR OWNER-0x1f01ff-OI|CI
      ALLOW-NT AUTHORITY\SYSTEM-0x1f01ff-OI|CI

```

The following example displays the audit policy information for the path /datavol1 in SVM vs1. The path contains both regular file and folder SACLs and Storage-Level Access Guard SACLs.


```

cluster::> vserver security file-directory show -vserver vs1 -path
/datavol1

        Vserver: vs1
        File Path: /datavol1
File Inode Number: 77
        Security Style: ntfs
Effective Style: ntfs
        DOS Attributes: 10
DOS Attributes in Text: ----D---
Expanded Dos Attributes: -
        Unix User Id: 0
        Unix Group Id: 0
        Unix Mode Bits: 777
Unix Mode Bits in Text: rwxrwxrwx
        ACLs: NTFS Security Descriptor
Control:0xaa14
Owner: BUILTIN\Administrators
Group: BUILTIN\Administrators
SACL - ACEs
        AUDIT-EXAMPLE\marketing-0xf01ff-OI|CI|FA
DACL - ACEs
        ALLOW-EXAMPLE\Domain Admins-0x1f01ff-OI|CI
        ALLOW-EXAMPLE\marketing-0x1200a9-OI|CI

Storage-Level Access Guard security
SACL (Applies to Directories):
        AUDIT-EXAMPLE\Domain Users-0x120089-FA
        AUDIT-EXAMPLE\engineering-0x1f01ff-SA
DACL (Applies to Directories):
        ALLOW-EXAMPLE\Domain Users-0x120089
        ALLOW-EXAMPLE\engineering-0x1f01ff
        ALLOW-NT AUTHORITY\SYSTEM-0x1f01ff
SACL (Applies to Files):
        AUDIT-EXAMPLE\Domain Users-0x120089-FA
        AUDIT-EXAMPLE\engineering-0x1f01ff-SA
DACL (Applies to Files):
        ALLOW-EXAMPLE\Domain Users-0x120089
        ALLOW-EXAMPLE\engineering-0x1f01ff
        ALLOW-NT AUTHORITY\SYSTEM-0x1f01ff

```

Ways to display information about file security and audit policies

You can use the wildcard character (*) to display information about file security and audit policies of all files and directories under a given path or a root volume.

The wildcard character (*) can be used as the last subcomponent of a given directory path below which you want to display information of all files and directories.

If you want to display information of a particular file or directory named as "*", then you need to provide the complete path inside double quotes (" ").

Example

The following command with the wildcard character displays the information about all files and directories below the path /1/ of SVM vs1:

```
cluster::> vserver security file-directory show -vserver vs1 -path /1/*

          Vserver: vs1
          File Path: /1/1
          Security Style: mixed
          Effective Style: ntfs
          DOS Attributes: 10
DOS Attributes in Text: ----D---
Expanded Dos Attributes: -
          Unix User Id: 0
          Unix Group Id: 0
          Unix Mode Bits: 777
Unix Mode Bits in Text: rwxrwxrwx
          ACLs: NTFS Security Descriptor
                Control:0x8514
                Owner: BUILTIN\Administrators
                Group: BUILTIN\Administrators
                DACL - ACEs
                ALLOW-Everyone-0x1f01ff-OI|CI (Inherited)

          Vserver: vs1
          File Path: /1/1/abc
          Security Style: mixed
          Effective Style: ntfs
          DOS Attributes: 10
DOS Attributes in Text: ----D---
Expanded Dos Attributes: -
          Unix User Id: 0
          Unix Group Id: 0
          Unix Mode Bits: 777
Unix Mode Bits in Text: rwxrwxrwx
          ACLs: NTFS Security Descriptor
                Control:0x8404
                Owner: BUILTIN\Administrators
                Group: BUILTIN\Administrators
                DACL - ACEs
                ALLOW-Everyone-0x1f01ff-OI|CI (Inherited)
```

The following command displays the information of a file named as "*" under the path /vol1/a of SVM vs1. The path is enclosed within double quotes (" ").

```
cluster::> vserver security file-directory show -vserver vs1 -path
"/vol1/a/*"

          Vserver: vs1
          File Path: "/vol1/a/*"
          Security Style: mixed
          Effective Style: unix
          DOS Attributes: 10
          DOS Attributes in Text: ----D---
          Expanded Dos Attributes: -
          Unix User Id: 1002
          Unix Group Id: 65533
          Unix Mode Bits: 755
          Unix Mode Bits in Text: rwxr-xr-x
          ACLs: NFSV4 Security Descriptor
                Control:0x8014
                SACL - ACEs
                     AUDIT-EVERYONE@-0x1f01bf-FI|DI|SA|FA
                DACL - ACEs
                     ALLOW-EVERYONE@-0x1f00a9-FI|DI
                     ALLOW-OWNER@-0x1f01ff-FI|DI
                     ALLOW-GROUP@-0x1200a9-IG
```

CLI change events that can be audited

CLI change events that can be audited overview

ONTAP can audit certain CLI change events, including certain SMB-share events, certain audit policy events, certain local security group events, local user group events, and authorization policy events. Understanding which change events can be audited is helpful when interpreting results from the event logs.

You can manage storage virtual machine (SVM) auditing CLI change events by manually rotating the audit logs, enabling or disabling auditing, displaying information about auditing change events, modifying auditing change events, and deleting auditing change events.

As an administrator, if you execute any command to change configuration related to the SMB-share, local user-group, local security-group, authorization-policy, and audit-policy events, a record generates and the corresponding event gets audited:

Auditing Category	Events	Event IDs	Run this command...
-------------------	--------	-----------	---------------------

Mhost Auditing	policy-change	[4719] Audit configuration changed	vserver audit disable enable modify
	file-share	[5142] Network share was added	vserver cifs share create
		[5143] Network share was modified	vserver cifs share modify vserver cifs share create modify delete vserver cifs share add remove
		[5144] Network share deleted	vserver cifs share delete

Auditing

	Rename	and-groups local-user rename
security-group	[4731] Local Security Group created	vserver cifs users-and-groups local-group create vserver services name-service unix-group create
	[4734] Local Security Group deleted	vserver cifs users-and-groups local-group delete vserver services name-service unix-group delete
	[4735] Local Security Group Modified	vserver cifs users-and-groups local-group rename modify vserver services name-service unix-group modify
	[4732] User added to Local Group	vserver cifs users-and-groups local-group add-members vserver services name-service unix-group adduser
	[4733] User Removed from Local Group	vserver cifs users-and-groups local-group remove-members vserver services name-service unix-group deluser
authorization-policy-change	[4704] User Rights Assigned	vserver cifs users-and-groups privilege add-privilege
	[4705] User Rights Removed	vserver cifs users-and-groups privilege remove-privilege reset-privilege

Manage file-share event

When a file-share event is configured for a storage virtual machine (SVM) and an audit is enabled, audit events are generated. The file-share events are generated when the SMB network share is modified using `vserver cifs share` related commands.

The file-share events with the event-ids 5142, 5143, and 5144 are generated when a SMB network share is added, modified, or deleted for the SVM. The SMB network share configuration is modified using the `cifs share access control create|modify|delete` commands.

The following example displays a file-share event with the ID 5143 is generated, when a share object called 'audit_dest' is created:

```
netapp-clus1::*> cifs share create -share-name audit_dest -path
/audit_dest
- System
- Provider
  [ Name] NetApp-Security-Auditing
  [ Guid] {3CB2A168-FE19-4A4E-BDAD-DCF422F13473}
EventID 5142
EventName Share Object Added
...
...
ShareName audit_dest
SharePath /audit_dest
ShareProperties oplocks;browsable;changenotify;show-previous-versions;
SD O:BAG:S-1-5-21-2447422786-1297661003-4197201688-513D:(A;;;FA;;;WD)
```

Manage audit-policy-change event

When an audit-policy-change event is configured for a storage virtual machine (SVM) and an audit is enabled, audit events are generated. The audit-policy-change events are generated when an audit policy is modified using `vserver audit` related commands.

The audit-policy-change event with the event-id 4719 is generated whenever an audit policy is disabled, enabled, or modified and helps to identify when a user attempts to disable auditing to cover the tracks. It is configured by default and requires diagnostic privilege to disable.

The following example displays an audit-policy change event with the ID 4719 generated, when an audit is disabled:

```
netapp-clus1::*> vserver audit disable -vserver vserver_1
- System
  - Provider
    [ Name]  NetApp-Security-Auditing
    [ Guid]  {3CB2A168-FE19-4A4E-BDAD-DCF422F13473}
    EventID 4719
    EventName Audit Disabled
    ...
    ...
    SubjectUserName admin
    SubjectUserSid 65533-1001
    SubjectDomainName ~
    SubjectIP console
    SubjectPort
```

Manage user-account event

When a user-account event is configured for a storage virtual machine (SVM) and an audit is enabled, audit events are generated.

The user-account events with event-ids 4720, 4722, 4724, 4725, 4726, 4738, and 4781 are generated when a local SMB or NFS user is created or deleted from the system, local user account is enabled, disabled or modified, and local SMB user password is reset or changed. The user-account events are generated when a user account is modified using `vserver cifs users-and-groups <local user>` and `vserver services name-service <unix user>` commands.

The following example displays a user account event with the ID 4720 generated, when a local SMB user is created:


```

netapp-clus1::~*> vserver cifs users-and-groups local-user create -user
-name testuser -is-account-disabled false -vserver vserver_1
Enter the password:
Confirm the password:

- System
  - Provider
    [ Name] NetApp-Security-Auditing
    [ Guid] {3CB2A168-FE19-4A4E-BDAD-DCF422F13473}
    EventID 4720
    EventName Local Cifs User Created
    ...
    ...
    TargetUserName testuser
    TargetDomainName NETAPP-CLUS1
    TargetSid S-1-5-21-2447422786-1297661003-4197201688-1003
    TargetType CIFS
    DisplayName testuser
    PasswordLastSet 1472662216
    AccountExpires NO
    PrimaryGroupId 513
    UserAccountControl %%0200
    SidHistory ~
    PrivilegeList ~

```

The following example displays a user account event with the ID 4781 generated, when the local SMB user created in the preceding example is renamed:

```
netapp-clus1::*> vserver cifs users-and-groups local-user rename -user
-name testuser -new-user-name testuser1
- System
- Provider
  [ Name]   NetApp-Security-Auditing
  [ Guid]   {3CB2A168-FE19-4A4E-BDAD-DCF422F13473}
  EventID  4781
  EventName Local Cifs User Renamed
  ...
  ...
  OldTargetUserName testuser
  NewTargetUserName testuser1
  TargetDomainName NETAPP-CLUS1
  TargetSid S-1-5-21-2447422786-1297661003-4197201688-1000
  TargetType CIFS
  SidHistory ~
  PrivilegeList ~
```

Manage security-group event

When a security-group event is configured for a storage virtual machine (SVM) and an audit is enabled, audit events are generated.

The security-group events with event-ids 4731, 4732, 4733, 4734, and 4735 are generated when a local SMB or NFS group is created or deleted from the system, and local user is added or removed from the group. The security-group-events are generated when a user account is modified using `vserver cifs users-and-groups <local-group>` and `vserver services name-service <unix-group>` commands.

The following example displays a security group event with the ID 4731 generated, when a local UNIX security group is created:

```

netapp-clus1::*> vserver services name-service unix-group create -name
testunixgroup -id 20
- System
- Provider
  [ Name]  NetApp-Security-Auditing
  [ Guid]  {3CB2A168-FE19-4A4E-BDAD-DCF422F13473}
  EventID 4731
  EventName Local Unix Security Group Created
  ...
  ...
  SubjectUserName admin
  SubjectUserSid 65533-1001
  SubjectDomainName ~
  SubjectIP console
  SubjectPort
  TargetUserName testunixgroup
  TargetDomainName
  TargetGid 20
  TargetType NFS
  PrivilegeList ~
  GidHistory ~

```

Manage authorization-policy-change event

When authorization-policy-change event is configured for a storage virtual machine (SVM) and an audit is enabled, audit events are generated.

The authorization-policy-change events with the event-ids 4704 and 4705 are generated whenever the authorization rights are granted or revoked for an SMB user and SMB group. The authorization-policy-change events are generated when the authorization rights are assigned or revoked using `vserver cifs users-and-groups privilege` related commands.

The following example displays an authorization policy event with the ID 4704 generated, when the authorization rights for a SMB user group are assigned:

```

netapp-clus1::*> vserver cifs users-and-groups privilege add-privilege
-user-or-group-name testcifslocalgroup -privileges *
- System
- Provider
  [ Name]   NetApp-Security-Auditing
  [ Guid]   {3CB2A168-FE19-4A4E-BDAD-DCF422F13473}
  EventID  4704
  EventName User Right Assigned
  ...
  ...
  TargetUserOrGroupName testcifslocalgroup
  TargetUserOrGroupDomainName NETAPP-CLUS1
  TargetUserOrGroupSid S-1-5-21-2447422786-1297661003-4197201688-1004;
  PrivilegeList
  SeTcbPrivilege;SeBackupPrivilege;SeRestorePrivilege;SeTakeOwnershipPrivile
ge;SeSecurityPrivilege;SeChangeNotifyPrivilege;
  TargetType CIFS

```

Manage auditing configurations

Manually rotate the audit event logs

Before you can view the audit event logs, the logs must be converted to user-readable formats. If you want to view the event logs for a specific storage virtual machine (SVM) before ONTAP automatically rotates the log, you can manually rotate the audit event logs on an SVM.

Step

1. Rotate the audit event logs by using the `vserver audit rotate-log` command.

```
vserver audit rotate-log -vserver vs1
```

The audit event log is saved in the SVM audit event log directory with the format specified by the auditing configuration (XML or EVT), and can be viewed by using the appropriate application.

Enable and disable auditing on SVMs

You can enable or disable auditing on storage virtual machines (SVMs). You might want to temporarily stop file and directory auditing by disabling auditing. You can enable auditing at any time (if an auditing configuration exists).

What you'll need

Before you can enable auditing on the SVM, the SVM's auditing configuration must already exist.

[Create the auditing configuration](#)

About this task

Disabling auditing does not delete the auditing configuration.

Steps

1. Perform the appropriate command:

If you want auditing to be...	Enter the command...
Enabled	<code>vserver audit enable -vserver vserver_name</code>
Disabled	<code>vserver audit disable -vserver vserver_name</code>

2. Verify that auditing is in the desired state:

```
vserver audit show -vserver vserver_name
```

Examples

The following example enables auditing for SVM vs1:

```
cluster1::> vserver audit enable -vserver vs1

cluster1::> vserver audit show -vserver vs1

                Vserver: vs1
                Auditing state: true
                Log Destination Path: /audit_log
                Categories of Events to Audit: file-ops, cifs-logon-logoff
                Log Format: evtX
                Log File Size Limit: 100MB
                Log Rotation Schedule: Month: -
                Log Rotation Schedule: Day of Week: -
                Log Rotation Schedule: Day: -
                Log Rotation Schedule: Hour: -
                Log Rotation Schedule: Minute: -
                Rotation Schedules: -
                Log Files Rotation Limit: 10
```

The following example disables auditing for SVM vs1:

```
cluster1::> vserver audit disable -vserver vs1

                Vserver: vs1
                Auditing state: false
                Log Destination Path: /audit_log
                Categories of Events to Audit: file-ops, cifs-logon-logoff
                Log Format: evtX
                Log File Size Limit: 100MB
                Log Rotation Schedule: Month: -
                Log Rotation Schedule: Day of Week: -
                Log Rotation Schedule: Day: -
                Log Rotation Schedule: Hour: -
                Log Rotation Schedule: Minute: -
                Rotation Schedules: -
                Log Files Rotation Limit: 10
```

Display information about auditing configurations

You can display information about auditing configurations. The information can help you determine whether the configuration is what you want in place for each SVM. The displayed information also enables you to verify whether an auditing configuration is enabled.

About this task

You can display detailed information about auditing configurations on all SVMs or you can customize what information is displayed in the output by specifying optional parameters. If you do not specify any of the optional parameters, the following is displayed:

- SVM name to which the auditing configuration applies
- The audit state, which can be `true` or `false`

If the audit state is `true`, auditing is enabled. If the audit state is `false`, auditing is disabled.

- The categories of events to audit
- The audit log format
- The target directory where the auditing subsystem stores consolidated and converted audit logs

Step

1. Display information about the auditing configuration by using the `vserver audit show` command.

For more information about using the command, see the man pages.

Examples

The following example displays a summary of the auditing configuration for all SVMs:

```
cluster1::> vserver audit show
```

```
Vserver      State  Event Types  Log Format  Target Directory
-----
vs1          false  file-ops     evtX       /audit_log
```


The following example displays, in list form, all auditing configuration information for all SVMs:

```
cluster1::> vserver audit show -instance
```

```
                Vserver: vs1
                Auditing state: true
                Log Destination Path: /audit_log
                Categories of Events to Audit: file-ops
                Log Format: evtX
                Log File Size Limit: 100MB
                Log Rotation Schedule: Month: -
                Log Rotation Schedule: Day of Week: -
                Log Rotation Schedule: Day: -
                Log Rotation Schedule: Hour: -
                Log Rotation Schedule: Minute: -
                Rotation Schedules: -
                Log Files Rotation Limit: 0
```

Commands for modifying auditing configurations

If you want to change an auditing setting, you can modify the current configuration at any time, including modifying the log path destination and log format, modifying the categories of events to audit, how to automatically save log files, and specify the maximum number of log files to save.

If you want to...	Use this command...
Modify the log destination path	<code>vserver audit modify</code> with the <code>-destination</code> parameter
Modify the category of events to audit	<code>vserver audit modify</code> with the <code>-events</code> parameter  To audit central access policy staging events, the Dynamic Access Control (DAC) SMB server option must be enabled on the storage virtual machine (SVM).

Modify the log format	<code>vserver audit modify</code> with the <code>-format</code> parameter
Enabling automatic saves based on internal log file size	<code>vserver audit modify</code> with the <code>-rotate-size</code> parameter
Enabling automatic saves based on a time interval	<code>vserver audit modify</code> with the <code>-rotate-schedule-month</code> , <code>-rotate-schedule-dayofweek</code> , <code>-rotate-schedule-day</code> , <code>-rotate-schedule-hour</code> , and <code>-rotate-schedule-minute</code> parameters
Specifying the maximum number of saved log files	<code>vserver audit modify</code> with the <code>-rotate-limit</code> parameter

Delete an auditing configuration

If you no longer want to audit file and directory events on the storage virtual machine (SVM) and do not want to maintain an auditing configuration on the SVM, you can delete the auditing configuration.

Steps

1. Disable the auditing configuration:

```
vserver audit disable -vserver vserver_name
vserver audit disable -vserver vs1
```

2. Delete the auditing configuration:

```
vserver audit delete -vserver vserver_name
vserver audit delete -vserver vs1
```

Understand the implications of reverting the cluster

If you plan to revert the cluster, you should be aware of the revert process ONTAP follows when there are auditing-enabled storage virtual machines (SVMs) in the cluster. You must take certain actions before reverting.

Reverting to a version of ONTAP that does not support the auditing of SMB logon and logoff events and central access policy staging events

Support for auditing of SMB logon and logoff events and for central access policy staging events starts with clustered Data ONTAP 8.3. If you are reverting to a version of ONTAP that does not support these event types and you have auditing configurations that monitor these event types, you must change the auditing configuration for those audit-enabled SVMs before reverting. You must modify the configuration so that only file-op events are audited.

Troubleshoot auditing and staging volume space issues

Issues can arise when there is insufficient space on either the staging volumes or on the volume containing the audit event logs. If there is insufficient space, new audit records cannot be created, which prevents clients from accessing data, and access requests fail. You should know how to troubleshoot and resolve these volume space issues.

Troubleshoot space issues related to the event log volumes

If volumes containing event log files run out of space, auditing cannot convert log records into log files. This results in client access failures. You must know how to troubleshoot space issues related to event log volumes.

- Storage virtual machine (SVM) and cluster administrators can determine whether there is insufficient volume space by displaying information about volume and aggregate usage and configuration.
- If there is insufficient space in the volumes containing event logs, SVM and cluster administrators can resolve the space issues by either removing some of the event log files or by increasing the size of the volume.



If the aggregate that contains the event log volume is full, then the size of the aggregate must be increased before you can increase the size of the volume. Only a cluster administrator can increase the size of an aggregate.

- The destination path for the event log files can be changed to a directory on another volume by modifying the auditing configuration.



Data access is denied in the following cases:

- The destination directory is deleted.
- The file limit on a volume, which hosts the destination directory, reaches to its maximum level.

Learn more about:

- [How to view information about volumes and increasing volume size.](#)
- [How to view information about aggregates and managing aggregates.](#)

Troubleshoot space issues related to the staging volumes

If any of the volumes containing staging files for your storage virtual machine (SVM) runs out of space, auditing cannot write log records into staging files. This results in client access failures. To troubleshoot this issue, you need to determine whether any of the staging volumes used in the SVM are full by displaying information about volume usage.

If the volume containing the consolidated event log files has sufficient space but there are still client access failures due to insufficient space, then the staging volumes might be out of space. The SVM administrator must contact you to determine whether the staging volumes that contain staging files for the SVM have insufficient space. The auditing subsystem generates an EMS event if auditing events cannot be generated due to insufficient space in a staging volume. The following message is displayed: `No space left on device.` Only you can view information about staging volumes; SVM administrators cannot.

All staging volume names begin with `MDV_aud_` followed by the UUID of the aggregate containing that staging

volume. The following example shows four system volumes on the admin SVM, which were automatically created when a file services auditing configuration was created for a data SVM in the cluster:

```
cluster1::> volume show -vserver cluster1
Vserver   Volume                               Aggregate   State      Type      Size   Available
Used%
-----
cluster1  MDV_aud_1d0131843d4811e296fc123478563412
          aggr0                               online     RW         5GB     4.75GB
5%
cluster1  MDV_aud_8be27f813d7311e296fc123478563412
          root_vs0                             online     RW         5GB     4.75GB
5%
cluster1  MDV_aud_9dc4ad503d7311e296fc123478563412
          aggr1                               online     RW         5GB     4.75GB
5%
cluster1  MDV_aud_a4b887ac3d7311e296fc123478563412
          aggr2                               online     RW         5GB     4.75GB
5%
4 entries were displayed.
```

If there is insufficient space in the staging volumes, you can resolve the space issues by increasing the size of the volume.



If the aggregate that contains the staging volume is full, then the size of the aggregate must be increased before you can increase the size of the volume. Only you can increase the size of an aggregate; SVM administrators cannot.

If one or more aggregates have an available space of less than 2GB (in ONTAP 9.14.1 and earlier) or 5GB (beginning with ONTAP 9.15.1), the SVM audit creation fails. When the SVM audit creation fails, the staging volumes that were created are deleted.

Use FPolicy for file monitoring and management on SVMs

Understand FPolicy

What the two parts of the FPolicy solution are

FPolicy is a file access notification framework that is used to monitor and manage file access events on storage virtual machines (SVMs) through partner solutions. Partner solutions help you address various use cases such as data governance and compliance, ransomware protection, and data mobility.

Partner solutions include both Netapp supported 3rd party Solutions and NetApp products Workload Security and Cloud Data Sense.

There are two parts to an FPolicy solution. The ONTAP FPolicy framework manages activities on the cluster and sends notifications to Partner Application (aka External FPolicy Servers). External FPolicy servers process notifications sent by ONTAP FPolicy to fulfill customer use cases.

The ONTAP framework creates and maintains the FPolicy configuration, monitors file events, and sends notifications to external FPolicy servers. ONTAP FPolicy provides the infrastructure that allows communication between external FPolicy servers and storage virtual machine (SVM) nodes.

The FPolicy framework connects to external FPolicy servers and sends notifications for certain file system events to the FPolicy servers when these events occur as a result of client access. The external FPolicy servers process the notifications and send responses back to the node. What happens as a result of the notification processing depends on the application and whether the communication between the node and the external servers is asynchronous or synchronous.

What synchronous and asynchronous notifications are

FPolicy sends notifications to external FPolicy servers via the FPolicy interface. The notifications are sent either in synchronous or asynchronous mode. The notification mode determines what ONTAP does after sending notifications to FPolicy servers.

- **Asynchronous notifications**

With asynchronous notifications, the node does not wait for a response from the FPolicy server, which enhances overall throughput of the system. This type of notification is suitable for applications where the FPolicy server does not require that any action be taken as a result of notification evaluation. For example, asynchronous notifications are used when the storage virtual machine (SVM) administrator wants to monitor and audit file access activity.

If an FPolicy server operating in asynchronous mode experiences a network outage, FPolicy notifications generated during the outage are stored on the storage node. When the FPolicy server comes back online, it is alerted of the stored notifications and can fetch them from the storage node. The length of time the notifications can be stored during an outage is configurable up to 10 minutes.

Beginning with ONTAP 9.14.1, FPolicy allows you to set up a persistent store to capture file access events for asynchronous non-mandatory policies in the SVM. Persistent stores can help decouple client I/O processing from FPolicy notification processing to reduce client latency. Synchronous (mandatory or non-mandatory) and asynchronous mandatory configurations are not supported.

- **Synchronous notifications**

When configured to run in synchronous mode, the FPolicy server must acknowledge every notification before the client operation is allowed to continue. This type of notification is used when an action is required based on the results of notification evaluation. For example, synchronous notifications are used when the SVM administrator wants to either allow or deny requests based on criteria specified on the external FPolicy server.

Synchronous and asynchronous applications

There are many possible uses for FPolicy applications, both asynchronous and synchronous.

Asynchronous applications are ones where the external FPolicy server does not alter access to files or directories or modify data on the storage virtual machine (SVM). For example:

- File access and audit logging

- Storage resource management

Synchronous applications are ones where data access is altered or data is modified by the external FPolicy server. For example:

- Quota management
- File access blocking
- File archiving and hierarchical storage management
- Encryption and decryption services
- Compression and decompression services

FPolicy persistent stores

Persistent stores can help decouple client I/O processing from FPolicy notification processing to reduce client latency. Beginning with ONTAP 9.14.1, you can set up an FPolicy persistent store to capture file access events for asynchronous non-mandatory policies in the SVM. Synchronous (either mandatory or non-mandatory) and asynchronous mandatory configurations are not supported.

This feature is only available in FPolicy external mode. The partner application you use needs to support this feature. You should work with your partner to ensure this FPolicy configuration is supported.

Beginning with ONTAP 9.15.1, FPolicy persistent store configuration is simplified. The `persistent-store create` command automates volume creation for the SVM and configures the volume with persistent store best practices.

For more information on persistent store best practices, refer to [Requirements, considerations, and best practices for configuring FPolicy](#).

For information on adding persistent stores, refer to [Create persistent stores](#).

FPolicy configuration types

There are two basic FPolicy configuration types. One configuration uses external FPolicy servers to process and act upon notifications. The other configuration does not use external FPolicy servers; instead, it uses the ONTAP internal, native FPolicy server for simple file blocking based on extensions.

- **External FPolicy server configuration**

The notification is sent to the FPolicy server, which screens the request and applies rules to determine whether the node should allow the requested file operation. For synchronous policies, the FPolicy server then sends a response to the node to either allow or block the requested file operation.

- **Native FPolicy server configuration**

The notification is screened internally. The request is allowed or denied based on file extension settings configured in the FPolicy scope.

Note: File extension requests that are denied are not logged.

When to create a native FPolicy configuration

Native FPolicy configurations use the ONTAP internal FPolicy engine to monitor and block file operations based on the file's extension. This solution does not require external FPolicy servers (FPolicy servers). Using a native file blocking configuration is appropriate when this simple solution is all that is needed.

Native file blocking enables you to monitor any file operations that match configured operation and filtering events and then deny access to files with particular extensions. This is the default configuration.

This configuration provides a means to block file access based only on the file's extension. For example, to block files that contain mp3 extensions, you configure a policy to provide notifications for certain operations with target file extensions of mp3. The policy is configured to deny mp3 file requests for operations that generate notifications.

The following applies to native FPolicy configurations:

- The same set of filters and protocols that are supported by FPolicy server-based file screening are also supported for native file blocking.
- Native file blocking and FPolicy server-based file screening applications can be configured at the same time.

To do so, you can configure two separate FPolicy policies for the storage virtual machine (SVM), with one configured for native file blocking and one configured for FPolicy server-based file screening.

- The native file blocking feature only screens files based on the extensions and not on the content of the file.
- In the case of symbolic links, native file blocking uses the file extension of the root file.

Learn more about [FPolicy: Native File Blocking](#).

When to create a configuration that uses external FPolicy servers

FPolicy configurations that use external FPolicy servers to process and manage notifications provide robust solutions for use cases where more than simple file blocking based on file extension is needed.

You should create a configuration that uses external FPolicy servers when you want to do such things as monitor and record file access events, provide quota services, perform file blocking based on criteria other than simple file extensions, provide data migration services using hierarchical storage management applications, or provide a fine-grained set of policies that monitor only a subset of data in the storage virtual machine (SVM).

Roles that cluster components play with FPolicy implementation

The cluster, the contained storage virtual machines (SVMs), and data LIFs all play a role in an FPolicy implementation.

- **cluster**

The cluster contains the FPolicy management framework and maintains and manages information about all FPolicy configurations in the cluster.

- **SVM**

An FPolicy configuration is defined at the SVM level. The scope of the configuration is the SVM, and it only operates on SVM resources. One SVM configuration cannot monitor and send notifications for file access

requests that are made for data residing on another SVM.

FPolicy configurations can be defined on the admin SVM. After configurations are defined on the admin SVM, they can be seen and used in all SVMs.

- **data LIFs**

Connections to the FPolicy servers are made through data LIFs belonging to the SVM with the FPolicy configuration. The data LIFs used for these connections can fail over in the same manner as data LIFs used for normal client access.

How FPolicy works with external FPolicy servers

After FPolicy is configured and enabled on the storage virtual machine (SVM), FPolicy runs on every node on which the SVM participates. FPolicy is responsible for establishing and maintaining connections with external FPolicy servers (FPolicy servers), for notification processing, and for managing notification messages to and from FPolicy servers.

Additionally, as part of connection management, FPolicy has the following responsibilities:

- Ensures that file notification flows through the correct LIF to the FPolicy server.
- Ensures that when multiple FPolicy servers are associated with a policy, load balancing is done when sending notifications to the FPolicy servers.
- Attempts to reestablish the connection when a connection to an FPolicy server is broken.
- Sends the notifications to FPolicy servers over an authenticated session.
- Manages the passthrough-read data connection established by the FPolicy server for servicing client requests when passthrough-read is enabled.

How control channels are used for FPolicy communication

FPolicy initiates a control channel connection to an external FPolicy server from the data LIFs of each node participating on a storage virtual machine (SVM). FPolicy uses control channels for transmitting file notifications; therefore, an FPolicy server might see multiple control channel connections based on SVM topology.

How privileged data access channels are used for synchronous communication

With synchronous use cases, the FPolicy server accesses data residing on the storage virtual machine (SVM) through a privileged data access path. Access through the privileged path exposes the complete file system to the FPolicy server. It can access data files to collect information, to scan files, read files, or write into files.

Because the external FPolicy server can access the entire file system from the root of the SVM through the privileged data channel, the privileged data channel connection must be secure.

How FPolicy connection credentials are used with privileged data access channels

The FPolicy server makes privileged data access connections to cluster nodes by using a specific Windows user credential that is saved with the FPolicy configuration. SMB is the only supported protocol for making a privileged data access channel connection.

If the FPolicy server requires privileged data access, the following conditions must be met:

- A SMB license must be enabled on the cluster.
- The FPolicy server must run under the credentials configured in the FPolicy configuration.

When making a data channel connection, FPolicy uses the credential for the specified Windows user name. Data access is made over the admin share ONTAP_ADMIN\$.

What granting super user credentials for privileged data access means

ONTAP uses the combination of the IP address and the user credential configured in the FPolicy configuration to grant super user credentials to the FPolicy server.

Super user status grants the following privileges when the FPolicy server accesses data:

- Avoid permission checks

The user avoids checks on files and directory access.

- Special locking privileges

ONTAP allows read, write, or modify access to any file regardless of existing locks. If the FPolicy server takes byte range locks on the file, it results in immediate removal of existing locks on the file.

- Bypass any FPolicy checks

Access does not generate any FPolicy notifications.

How FPolicy manages policy processing

There might be multiple FPolicy policies assigned to your storage virtual machine (SVM); each with a different priority. To create an appropriate FPolicy configuration on the SVM, it is important to understand how FPolicy manages policy processing.

Each file access request is initially evaluated to determine which policies are monitoring this event. If it is a monitored event, information about the monitored event along with interested policies is passed to FPolicy where it is evaluated. Each policy is evaluated in order of the assigned priority.

You should consider the following recommendations when configuring policies:

- When you want a policy to always be evaluated before other policies, configure that policy with a higher priority.
- If the success of requested file access operation on a monitored event is a prerequisite for a file request that is evaluated against another policy, give the policy that controls the success or failure of the first file operation a higher priority.

For example, if one policy manages FPolicy file archiving and restore functionality and a second policy manages file access operations on the online file, the policy that manages file restoration must have a higher priority so that the file is restored before the operation managed by the second policy can be allowed.

- If you want all policies that might apply to a file access operation to be evaluated, give synchronous policies a lower priority.

You can reorder policy priorities for existing policies by modifying the policy sequence number. However, to have FPolicy evaluate policies based on the modified priority order, you must disable and reenab

with the modified sequence number.

What the node-to-external FPolicy server communication process is

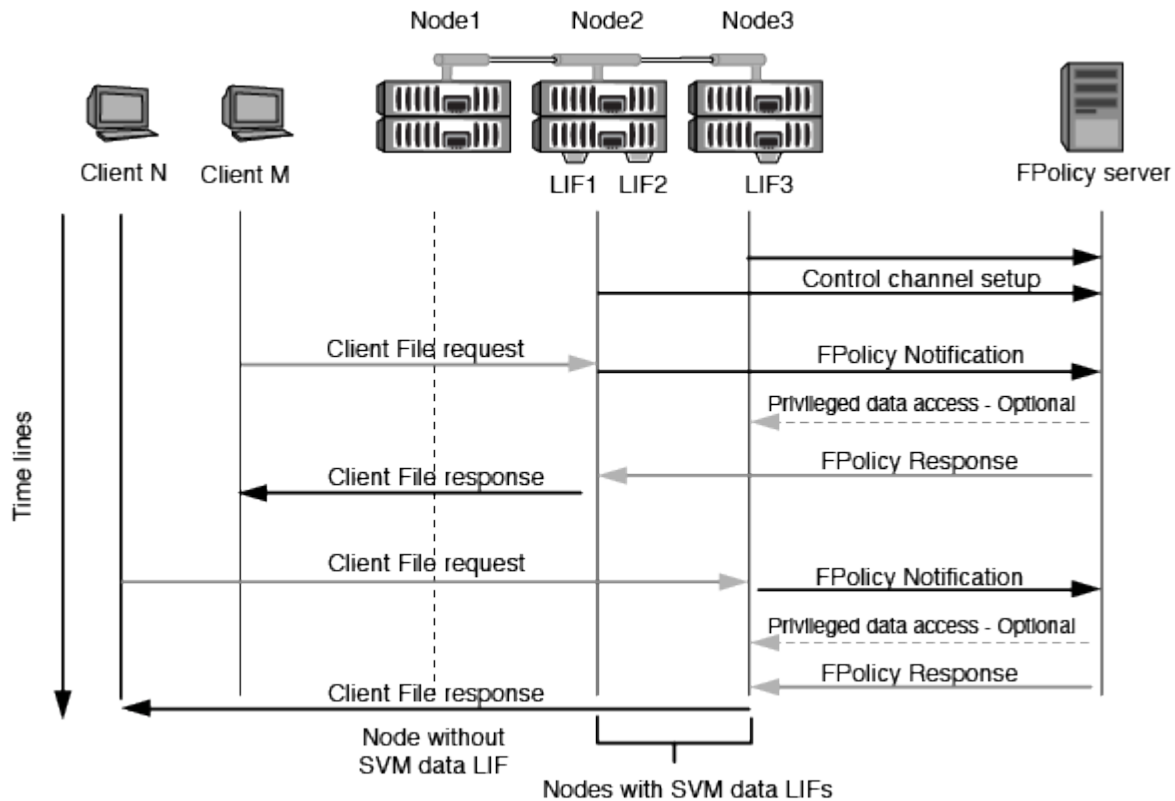
To properly plan your FPolicy configuration, you should understand what the node-to-external FPolicy server communication process is.

Every node that participates on each storage virtual machine (SVM) initiates a connection to an external FPolicy server (FPolicy server) using TCP/IP. Connections to the FPolicy servers are set up using node data LIFs; therefore, a participating node can set up a connection only if the node has an operational data LIF for the SVM.

Each FPolicy process on participating nodes attempts to establish a connection with the FPolicy server when the policy is enabled. It uses the IP address and port of the FPolicy external engine specified in the policy configuration.

The connection establishes a control channel from each of the nodes participating on each SVM to the FPolicy server through the data LIF. In addition, if IPv4 and IPv6 data LIF addresses are present on the same participating node, FPolicy attempts to establish connections for both IPv4 and IPv6. Therefore, in a scenario where the SVM extends over multiple nodes or if both IPv4 and IPv6 addresses are present, the FPolicy server must be ready for multiple control channel setup requests from the cluster after the FPolicy policy is enabled on the SVM.

For example, if a cluster has three nodes—Node1, Node2, and Node3—and SVM data LIFs are spread across only Node2 and Node3, control channels are initiated only from Node2 and Node3, irrespective of the distribution of data volumes. Say that Node2 has two data LIFs—LIF1 and LIF2—that belong to the SVM and that the initial connection is from LIF1. If LIF1 fails, FPolicy attempts to establish a control channel from LIF2.



How FPolicy manages external communication during LIF migration or failover

Data LIFs can be migrated to data ports in the same node or to data ports on a remote node.

When a data LIF fails over or is migrated, a new control channel connection is made to the FPolicy server. FPolicy can then retry SMB and NFS client requests that timed out, with the result that new notifications are sent to the external FPolicy servers. The node rejects FPolicy server responses to original, timed-out SMB and NFS requests.

How FPolicy manages external communication during node failover

If the cluster node that hosts the data ports used for FPolicy communication fails, ONTAP breaks the connection between the FPolicy server and the node.

The impact of cluster failover to the FPolicy server can be mitigated by configuring the failover-policy to migrate the data port used in FPolicy communication to another active node. After the migration is complete, a new connection is established using the new data port.

If the failover-policy is not configured to migrate the data port, the FPolicy server must wait for the failed node to come up. After the node is up, a new connection is initiated from that node with a new Session ID.



The FPolicy server detects broken connections with the keep-alive protocol message. The timeout for purging the session ID is determined when configuring FPolicy. The default keep-alive timeout is two minutes.

How FPolicy services work across SVM namespaces

ONTAP provides a unified storage virtual machine (SVM) namespace. Volumes across the cluster are joined together by junctions to provide a single, logical file system. The FPolicy server is aware of the namespace topology and provides FPolicy services across the namespace.

The namespace is specific to and contained within the SVM; therefore, you can see the namespace only from the SVM context. Namespaces have the following characteristics:

- A single namespace exists in each SVM, with the root of the namespace being the root volume, represented in the namespace as slash (/).
- All other volumes have junction points below the root (/).
- Volume junctions are transparent to clients.
- A single NFS export can provide access to the complete namespace; otherwise, export policies can export specific volumes.
- SMB shares can be created on the volume or on qtrees within the volume, or on any directory within the namespace.
- The namespace architecture is flexible.

Examples of typical namespace architectures are as follows:

- A namespace with a single branch off of the root
- A namespace with multiple branches off of the root
- A namespace with multiple unbranched volumes off of the root

How FPolicy passthrough-read enhances usability for hierarchical storage management

Passthrough-read enables the FPolicy server (functioning as the hierarchical storage management (HSM) server) to provide read access to offline files without having to recall the file from the secondary storage system to the primary storage system.

When an FPolicy server is configured to provide HSM to files residing on a SMB server, policy-based file migration occurs where the files are stored offline on secondary storage and only a stub file remains on primary storage. Even though a stub file appears as a normal file to clients, it is actually a sparse file that is the same size of the original file. The sparse file has the SMB offline bit set and points to the actual file that has been migrated to secondary storage.

Typically when a read request for an offline file is received, the requested content must be recalled back to primary storage and then accessed through primary storage. The need to recall data back to primary storage has several undesirable effects. Among the undesirable effects is the increased latency to client requests caused by the need to recall the content before responding to the request and the increased space consumption needed for recalled files on the primary storage.

FPolicy passthrough-read allows the HSM server (the FPolicy server) to provide read access to migrated, offline files without having to recall the file from the secondary storage system to the primary storage system. Instead of recalling the files back to primary storage, read requests can be serviced directly from secondary storage.



Copy Offload (ODX) is not supported with FPolicy passthrough-read operation.

Passthrough-read enhances usability by providing the following benefits:

- Read requests can be serviced even if the primary storage does not have sufficient space to recall requested data back to primary storage.
- Better capacity and performance management when a surge of data recall might occur, such as if a script or a backup solution needs to access many offline files.
- Read requests for offline files in Snapshot copies can be serviced.

Because Snapshot copies are read-only, the FPolicy server cannot restore the original file if the stub file is located in a Snapshot copy. Using passthrough-read eliminates this problem.

- Policies can be set up that control when read requests are serviced through access to the file on secondary storage and when the offline file should be recalled to primary storage.

For example, a policy can be created on the HSM server that specifies the number of times the offline file can be accessed in a specified period of time before the file is migrated back to primary storage. This type of policy avoids recalling files that are rarely accessed.

How read requests are managed when FPolicy passthrough-read is enabled

You should understand how read requests are managed when FPolicy passthrough-read is enabled so that you can optimally configure connectivity between the storage virtual machine (SVM) and the FPolicy servers.

When FPolicy passthrough-read is enabled and the SVM receives a request for an offline file, FPolicy sends a notification to the FPolicy server (HSM server) through the standard connection channel.

After receiving the notification, the FPolicy server reads the data from the file path sent in the notification and sends the requested data to the SVM through the passthrough-read privileged data connection that is

established between the SVM and the FPolicy server.

After the data is sent, the FPolicy server then responds to the read request as an ALLOW or DENY. Based on whether the read request is allowed or denied, ONTAP either sends the requested information or sends an error message to the client.

Plan the FPolicy configuration

Requirements, considerations, and best practices for configuring FPolicy

Before you create and configure FPolicy configurations on your storage virtual machines (SVMs), you need to be aware of certain requirements, considerations, and best practices for configuring FPolicy.

FPolicy features are configured either through the command line interface (CLI) or through REST APIs.

Requirements for setting up FPolicy

Before you configure and enable FPolicy on your storage virtual machine (SVM), you need to be aware of certain requirements.

- All nodes in the cluster must be running a version of ONTAP that supports FPolicy.
- If you are not using the ONTAP native FPolicy engine, you must have external FPolicy servers (FPolicy servers) installed.
- The FPolicy servers must be installed on a server accessible from the data LIFs of the SVM where FPolicy policies are enabled.



Beginning with ONTAP 9.8, ONTAP provides a client LIF service for outbound FPolicy connections with the addition of the `data-fpolicy-client` service. [Learn more about LIFs and service policies.](#)

- The IP address of the FPolicy server must be configured as a primary or secondary server in the FPolicy policy external engine configuration.
- If the FPolicy servers access data over a privileged data channel, the following additional requirements must be met:
 - SMB must be licensed on the cluster.

Privileged data access is accomplished using SMB connections.
 - A user credential must be configured for accessing files over the privileged data channel.
 - The FPolicy server must run under the credentials configured in the FPolicy configuration.
 - All data LIFs used to communicate with the FPolicy servers must be configured to have `cifs` as one of the allowed protocols.

This includes the LIFs used for passthrough-read connections.

Best practices and recommendations when setting up FPolicy

When setting up FPolicy on storage virtual machines (SVMs), get familiar with the general configuration best practices and recommendations to ensure that your FPolicy configuration provides robust monitoring

performance and results that meet your requirements.

For specific guidelines related to performance, sizing, and configuration, work with your FPolicy partner application.

Persistent stores

Beginning with ONTAP 9.14.1, FPolicy allows you to set up a persistent store to capture file access events for asynchronous non-mandatory policies in the SVM. Persistent stores can help decouple client I/O processing from FPolicy notification processing to reduce client latency. Synchronous (either mandatory or non-mandatory) and asynchronous mandatory configurations are not supported.

- Before using the persistent store functionality, ensure your partner applications support this configuration.
- You need one persistent store for each SVM where FPolicy is enabled.
 - Only one persistent store can be set up on each SVM. This single persistent store needs to be used for all FPolicy configurations on that SVM, even if the policies are from different partners.
- ONTAP 9.15.1 or later:
 - The persistent store, its volume, and its volume configuration is handled automatically when you create the persistent store.
- ONTAP 9.14.1:
 - The persistent store, its volume, and its volume configuration is handled manually.
- Create the persistent store volume on the node with LIFs that expect maximum traffic to be monitored by FPolicy.
 - ONTAP 9.15.1 or later: Volumes are automatically created and configured during persistent store creation.
 - ONTAP 9.14.1: Cluster administrators need to create and configure a volume for the persistent store on each SVM where FPolicy is enabled.
- If the notifications accumulated in the persistent store exceed the size of the volume provisioned, FPolicy starts dropping the incoming notification with appropriate EMS messages.
 - ONTAP 9.15.1 or later: In addition to the `size` parameter, the `autosize-mode` parameter can help the volume grow or shrink in response to the amount of used space.
 - ONTAP 9.14.1: The `size` parameter is configured during volume creation to provide a maximum limit.
- Set the snapshot policy to `none` for the persistent store volume instead of `default`. This is to ensure that there is no accidental restore of the snapshot leading to loss of current events and to prevent possible duplicate event processing.
 - ONTAP 9.15.1 or later: The `snapshot-policy` parameter is automatically configured to `none` during persistent store creation.
 - ONTAP 9.14.1: The `snapshot-policy` parameter is configured to `none` during volume creation.
- Make the persistent store volume inaccessible for external user protocol access (CIFS/NFS) to avoid accidental corruption or deletion of the persisted event records.
 - ONTAP 9.15.1 or later: ONTAP automatically blocks the volume from external user protocol access (CIFS/NFS) during persistent store creation.
 - ONTAP 9.14.1: After enabling FPolicy, unmount the volume in ONTAP to remove the junction path. This makes it inaccessible for external user protocol access (CIFS/NFS).

For more information, refer to [FPolicy persistent stores](#) and [Create persistent stores](#).

Persistent store failover and giveback

The persistent store remains as it was when the last event was received, when there is an unexpected reboot, or FPolicy is disabled and enabled again. After a takeover operation, new events are stored and processed by the partner node. After a giveback operation, the persistent store resumes processing any unprocessed events that might remain from when the node takeover occurred. Live events would be given priority over unprocessed events.

If the persistent store volume moves from one node to another in the same SVM, the notifications that are yet to be processed also move to the new node. You need to re-run the `fpolicy persistent-store create` command on either node after the volume is moved to ensure the pending notifications are delivered to the external server.

Policy configuration

Configuration of the FPolicy external engine, events, and scope for SVMs can improve your overall experience and security.

- Configuration of the FPolicy external engine for SVMs:
 - Providing additional security comes with a performance cost. Enabling Secure Sockets Layer (SSL) communication has a performance effect on accessing shares.
 - The FPolicy external engine should be configured with more than one FPolicy server to provide resiliency and high availability of FPolicy server notification processing.
- Configuration of FPolicy events for SVMs:

Monitoring file operations influences your overall experience. For example, filtering unwanted file operations on the storage side improves your experience. NetApp recommends setting up the following configuration:

- Monitoring the minimum types of file operations and enabling the maximum number of filters without breaking the use case.
 - Using filters for `getattr`, `read`, `write`, `open`, and `close` operations. The SMB and NFS home directory environments have a high percentage of these operations.
- Configuration of FPolicy scope for SVMs:

Restrict the scope of the policies to the relevant storage objects, such as shares, volumes, and exports, instead of enabling them across the entire SVM. NetApp recommends checking the directory extensions. If the `is-file-extension-check-on-directories-enabled` parameter is set to `true`, directory objects are subjected to the same extension checks as regular files.

Network configuration

Network connectivity between the FPolicy server and the controller should be of low latency. NetApp recommends separating FPolicy traffic from client traffic by using a private network.

In addition, you should place external FPolicy servers (FPolicy servers) in close proximity to the cluster with high-bandwidth connectivity to provide minimal latency and high-bandwidth connectivity.



For a scenario in which the LIF for FPolicy traffic is configured on a different port to the LIF for client traffic, the FPolicy LIF might fail over to the other node because of a port failure. As a result, the FPolicy server becomes unreachable from the node which causes the FPolicy notifications for file operations on the node to fail. To avoid this issue, verify that the FPolicy server can be reached through at least one LIF on the node to process FPolicy requests for the file operations performed on that node.

Hardware configuration

You can have the FPolicy server on either a physical server or a virtual server. If the FPolicy server is in a virtual environment, you should allocate dedicated resources (CPU, network, and memory) to the virtual server.

The cluster node-to-FPolicy server ratio should be optimized to ensure that FPolicy servers are not overloaded, which can introduce latencies when the SVM responds to client requests. The optimal ratio depends on the partner application for which the FPolicy server is being used. NetApp recommends working with partners to determine the appropriate value.

Multiple-policy configuration

The FPolicy policy for native blocking has the highest priority, irrespective of the sequence number, and decision-altering policies have a higher priority than others. Policy priority depends on the use case. NetApp recommends working with partners to determine the appropriate priority.

Size considerations

FPolicy performs in-line monitoring of SMB and NFS operations, sends notifications to the external server, and waits for a response, depending on the mode of external engine communication (synchronous or asynchronous). This process affects the performance of SMB and NFS access and CPU resources.

To mitigate any issues, NetApp recommends working with partners to assess and size the environment before enabling FPolicy. Performance is affected by several factors including the number of users, workload characteristics, such as operations per user and data size, network latency, and failure or server slowness.

Monitor performance

FPolicy is a notification-based system. Notifications are sent to an external server for processing and to generate a response back to ONTAP. This round-trip process increases latency for client access.

Monitoring the performance counters on the FPolicy server and in ONTAP gives you the capability to identify bottlenecks in the solution and to tune the parameters as necessary for an optimal solution. For example, an increase in FPolicy latency has a cascading effect on SMB and NFS access latency. Therefore, you should monitor both workload (SMB and NFS) and FPolicy latency. In addition, you can use quality-of-service policies in ONTAP to set up a workload for each volume or SVM that is enabled for FPolicy.

NetApp recommends running the `statistics show -object workload` command to display workload statistics. In addition, you should monitor the following parameters:

- Average, read, and write latencies
- Total number of operations
- Read and write counters

You can monitor the performance of FPolicy subsystems by using the following FPolicy counters.



You must be in diagnostic mode to collect statistics related to FPolicy.

Steps

1. Collect FPolicy counters:

- a. `statistics start -object fpolicy -instance instance_name -sample-id ID`
- b. `statistics start -object fpolicy_policy -instance instance_name -sample-id ID`

2. Display FPolicy counters:

- a. `statistics show -object fpolicy -instance instance_name -sample-id ID`
- b. `statistics show -object fpolicy_server -instance instance_name -sample-id ID`

The `fpolicy` and `fpolicy_server` counters give you information on several performance parameters which are described in the following table.

Counters	Description
“fpolicy” counters	
<code>aborted_requests</code>	Number of screen requests for which processing is aborted on the SVM
<code>event_count</code>	List of events resulting in notification
<code>max_request_latency</code>	Maximum screen requests latency
<code>outstanding_requests</code>	Total number of screen requests in process
<code>processed_requests</code>	Total number of screen requests that went through fpolicy processing on the SVM
<code>request_latency_hist</code>	Histogram of latency for screen requests
<code>requests_dispatched_rate</code>	Number of screen requests dispatched per second
<code>requests_received_rate</code>	Number of screen requests received per second
“fpolicy_server” counters	
<code>max_request_latency</code>	Maximum latency for a screen request
<code>outstanding_requests</code>	Total number of screen requests waiting for response
<code>request_latency</code>	Average latency for screen request
<code>request_latency_hist</code>	Histogram of latency for screen requests
<code>request_sent_rate</code>	Number of screen requests sent to FPolicy server per second
<code>response_received_rate</code>	Number of screen responses received from FPolicy server per second

Manage FPolicy workflow and dependency on other technologies

NetApp recommends disabling an FPolicy policy before making any configuration changes. For example, if you want to add or modify an IP address in the external engine configured for the enabled policy, first disable the policy.

If you configure FPolicy to monitor NetApp FlexCache volumes, NetApp recommends that you not configure FPolicy to monitor read and getattr file operations. Monitoring these operations in ONTAP requires the retrieval of inode-to-path (I2P) data. Because I2P data cannot be retrieved from FlexCache volumes, it must be retrieved from the origin volume. Therefore, monitoring these operations eliminates the performance benefits that FlexCache can provide.

When both FPolicy and an off-box antivirus solution are deployed, the antivirus solution receives notifications first. FPolicy processing starts only after antivirus scanning is complete. It is important that you size antivirus solutions correctly because a slow antivirus scanner can affect overall performance.

Passthrough-read upgrade and revert considerations

There are certain upgrade and revert considerations that you must know about before upgrading to an ONTAP release that supports passthrough-read or before reverting to a release that does not support passthrough-read.

Upgrading

After all nodes are upgraded to a version of ONTAP that supports FPolicy passthrough-read, the cluster is capable of using the passthrough-read functionality; however, passthrough-read is disabled by default on existing FPolicy configurations. To use passthrough-read on existing FPolicy configurations, you must disable the FPolicy policy and modify the configuration, and then reenables the configuration.

Reverting

Before reverting to a version of ONTAP that does not support FPolicy passthrough-read, you must meet the following conditions:

- Disable all the policies using passthrough-read, and then modify the affected configurations so that they do not use passthrough-read.
- Disable FPolicy functionality on the cluster by disabling every FPolicy policy on the cluster.

Before reverting to a version of ONTAP that does not support persistent stores, ensure that none of the FPolicy policies have a configured persistent store. If a persistent store is configured, the revert will fail.

What the steps for setting up an FPolicy configuration are

Before FPolicy can monitor file access, an FPolicy configuration must be created and enabled on the storage virtual machine (SVM) for which FPolicy services are required.

The steps for setting up and enabling an FPolicy configuration on the SVM are as follows:

1. Create an FPolicy external engine.

The FPolicy external engine identifies the external FPolicy servers (FPolicy servers) that are associated with a specific FPolicy configuration. If the internal “native” FPolicy engine is used to create a native file-blocking configuration, you do not need to create an FPolicy external engine.

Beginning with ONTAP 9.15.1, you can use the `protobuf` engine format. When set to `protobuf`, the notification messages are encoded in binary form using Google Protobuf. Before setting the engine format to `protobuf`, ensure that the FPolicy server also supports `protobuf` deserialization. For more information, see [Plan the FPolicy external engine configuration](#)

2. Create an FPolicy event.

An FPolicy event describes what the FPolicy policy should monitor. Events consist of the protocols and file operations to monitor, and can contain a list of filters. Events use filters to narrow the list of monitored events for which the FPolicy external engine must send notifications. Events also specify whether the policy monitors volume operations.

3. Create an FPolicy persistent store (optional).

Beginning with ONTAP 9.14.1, FPolicy allows you to set up [persistent stores](#) to capture file access events for asynchronous non-mandatory policies in the SVM. Synchronous (either mandatory or non-mandatory) and asynchronous mandatory configurations are not supported.

Persistent stores can help decouple client I/O processing from FPolicy notification processing to reduce client latency.

Beginning with ONTAP 9.15.1, FPolicy persistent store configuration is simplified. The `persistent-store-create` command automates volume creation for the SVM and configures the volume for the persistent store.

4. Create an FPolicy policy.

The FPolicy policy is responsible for associating, with the appropriate scope, the set of events that need to be monitored and for which of the monitored events notifications must be sent to the designated FPolicy server (or to the native engine if no FPolicy servers are configured). The policy also defines whether the FPolicy server is allowed privileged access to the data for which it receives notifications. An FPolicy server needs privileged access if the server needs to access the data. Typical use cases where privileged access is needed include file blocking, quota management, and hierarchical storage management. The policy is where you specify whether the configuration for this policy uses an FPolicy server or the internal “native” FPolicy server.

A policy specifies whether screening is mandatory. If screening is mandatory and all FPolicy servers are down or no response is received from the FPolicy servers within a defined timeout period, then file access is denied.

A policy’s boundaries are the SVM. A policy cannot apply to more than one SVM. However, a specific SVM can have multiple FPolicy policies, each with the same or different combination of scope, event, and external server configurations.

5. Configure the policy scope.

The FPolicy scope determines which volumes, shares, or export-policies the policy acts on or excludes from monitoring. A scope also determines which file extensions should be included or excluded from FPolicy monitoring.



Exclude lists take precedence over include lists.

6. Enable the FPolicy policy.

When the policy is enabled, the control channels and, optionally, the privileged data channels are connected. The FPolicy process on the nodes on which the SVM participates begin monitoring file and folder access and, for events that match configured criteria, sends notifications to the FPolicy servers (or to the native engine if no FPolicy servers are configured).



If the policy uses native file blocking, an external engine is not configured or associated with the policy.

Plan the FPolicy external engine configuration

Plan the FPolicy external engine configuration

Before you configure the FPolicy external engine, you must understand what it means to create an external engine and which configuration parameters are available. This information helps you to determine which values to set for each parameter.

Information that is defined when creating the FPolicy external engine

The external engine configuration defines the information that FPolicy needs to make and manage connections to the external FPolicy servers, including the following:

- SVM name
- Engine name
- The IP addresses of the primary and secondary FPolicy servers and the TCP port number to use when making the connection to the FPolicy servers
- Whether the engine type is asynchronous or synchronous
- Whether the engine format is `xml` or `protobuf`

Beginning with ONTAP 9.15.1, you can use the `protobuf` engine format. When set to `protobuf`, the notification messages are encoded in binary form using Google Protobuf. Before setting the engine format to `protobuf`, ensure that the FPolicy server also supports `protobuf` deserialization.

Since the `protobuf` format is supported beginning with ONTAP 9.15.1, you must consider external engine format before reverting to an earlier release of ONTAP. If you revert to an earlier release than ONTAP 9.15.1, work with your FPolicy partner to either:

- Change each engine format from `protobuf` to `xml`
- Delete the engines with an engine format of `protobuf`
- How to authenticate the connection between the node and the FPolicy server

If you choose to configure mutual SSL authentication, then you must also configure parameters that provide SSL certificate information.

- How to manage the connection using various advanced privilege settings

This includes parameters that define such things as timeout values, retry values, keep-alive values, maximum request values, sent and receive buffer size values, and session timeout values.

The `vserver fpolicy policy external-engine create` command is used to create an FPolicy external engine.

What the basic external engine parameters are

You can use the following table of basic FPolicy configuration parameters to help you plan your configuration:

Type of information	Option
<p>SVM</p> <p>Specifies the SVM name that you want to associate with this external engine.</p> <p>Each FPolicy configuration is defined within a single SVM. The external engine, policy event, policy scope, and policy that combine together to create an FPolicy policy configuration must all be associated with the same SVM.</p>	<p><code>-vserver vserver_name</code></p>
<p>Engine name</p> <p>Specifies the name to assign to the external engine configuration. You must specify the external engine name later when you create the FPolicy policy. This associates the external engine with the policy.</p> <p>The name can be up to 256 characters long.</p> <div data-bbox="167 810 220 863" style="border: 1px solid black; border-radius: 50%; width: 30px; height: 30px; display: flex; align-items: center; justify-content: center; margin-left: 20px;"> i </div> <p style="margin-left: 40px;">The name should be up to 200 characters long if configuring the external engine name in a MetroCluster or SVM disaster recovery configuration.</p> <p>The name can contain any combination of the following ASCII-range characters:</p> <ul style="list-style-type: none"> • a through z • A through Z • 0 through 9 • “_”, “-”, and “.” 	<p><code>-engine-name engine_name</code></p>
<p>Primary FPolicy servers</p> <p>Specifies the primary FPolicy servers to which the node sends notifications for a given FPolicy policy. The value is specified as a comma-delimited list of IP addresses.</p> <p>If more than one primary server IP address is specified, every node on which the SVM participates creates a control connection to every specified primary FPolicy server at the time the policy is enabled. If you configure multiple primary FPolicy servers, notifications are sent to the FPolicy servers in a round-robin fashion.</p> <p>If the external engine is used in a MetroCluster or SVM disaster recovery configuration, you should specify the IP addresses of the FPolicy servers at the source site as primary servers. The IP addresses of the FPolicy servers at the destination site should be specified as secondary servers.</p>	<p><code>-primary-servers IP_address,...</code></p>

<p><i>Port number</i></p> <p>Specifies the port number of the FPolicy service.</p>	<p>-port integer</p>
<p><i>Secondary FPolicy servers</i></p> <p>Specifies the secondary FPolicy servers to which to send file access events for a given FPolicy policy. The value is specified as a comma-delimited list of IP addresses.</p> <p>Secondary servers are used only when none of the primary servers are reachable. Connections to secondary servers are established when the policy is enabled, but notifications are sent to secondary servers only if none of the primary servers are reachable. If you configure multiple secondary servers, notifications are sent to the FPolicy servers in a round-robin fashion.</p>	<p>-secondary-servers IP_address,...</p>
<p><i>External engine type</i></p> <p>Specifies whether the external engine operates in synchronous or asynchronous mode. By default, FPolicy operates in synchronous mode.</p> <p>When set to <code>synchronous</code>, file request processing sends a notification to the FPolicy server, but then does not continue until after receiving a response from the FPolicy server. At that point, request flow either continues or processing results in denial, depending on whether the response from the FPolicy server permits the requested action.</p> <p>When set to <code>asynchronous</code>, file request processing sends a notification to the FPolicy server, and then continues.</p>	<p>-extern-engine-type external_engine_type The value for this parameter can be one of the following:</p> <ul style="list-style-type: none"> • synchronous • asynchronous
<p><i>External engine format</i></p> <p>Specify whether the external engine format is xml or protobuf.</p> <p>Beginning with ONTAP 9.15.1, you can use the protobuf engine format. When set to <code>protobuf</code>, the notification messages are encoded in binary form using Google Protobuf. Before setting the engine format to <code>protobuf</code>, ensure that the FPolicy server also supports protobuf deserialization.</p>	<p>- extern-engine-format {protobuf or xml}</p>

<p><i>SSL option for communication with FPolicy server</i></p> <p>Specifies the SSL option for communication with the FPolicy server. This is a required parameter. You can choose one of the options based on the following information:</p> <ul style="list-style-type: none"> • When set to <code>no-auth</code>, no authentication takes place. <p>The communication link is established over TCP.</p> <ul style="list-style-type: none"> • When set to <code>server-auth</code>, the SVM authenticates the FPolicy server using SSL server authentication. • When set to <code>mutual-auth</code>, mutual authentication takes place between the SVM and the FPolicy server; the SVM authenticates the FPolicy server, and the FPolicy server authenticates the SVM. <p>If you choose to configure mutual SSL authentication, then you must also configure the <code>-certificate-common-name</code>, <code>-certificate-serial</code>, and <code>-certificate-ca</code> parameters.</p>	<pre>-ssl-option {no-auth server-auth mutual-auth}</pre>
<p><i>Certificate FQDN or custom common name</i></p> <p>Specifies the certificate name used if SSL authentication between the SVM and the FPolicy server is configured. You can specify the certificate name as an FQDN or as a custom common name.</p> <p>If you specify <code>mutual-auth</code> for the <code>-ssl-option</code> parameter, you must specify a value for the <code>-certificate-common-name</code> parameter.</p>	<pre>-certificate-common -name text</pre>
<p><i>Certificate serial number</i></p> <p>Specifies the serial number of the certificate used for authentication if SSL authentication between the SVM and the FPolicy server is configured.</p> <p>If you specify <code>mutual-auth</code> for the <code>-ssl-option</code> parameter, you must specify a value for the <code>-certificate-serial</code> parameter.</p>	<pre>-certificate-serial text</pre>
<p><i>Certificate authority</i></p> <p>Specifies the CA name of the certificate used for authentication if SSL authentication between the SVM and the FPolicy server is configured.</p> <p>If you specify <code>mutual-auth</code> for the <code>-ssl-option</code> parameter, you must specify a value for the <code>-certificate-ca</code> parameter.</p>	<pre>-certificate-ca text</pre>

What the advanced external engine options are

You can use the following table of advanced FPolicy configuration parameters as you plan whether to customize your configuration with advanced parameters. You use these parameters to modify communication behavior between the cluster nodes and the FPolicy servers:

Type of information	Option
<p><i>Timeout for canceling a request</i></p> <p>Specifies the time interval in hours (h), minutes (m), or seconds (s) that the node waits for a response from the FPolicy server.</p> <p>If the timeout interval passes, the node sends a cancel request to the FPolicy server. The node then sends the notification to an alternate FPolicy server. This timeout helps in handling an FPolicy server that is not responding, which can improve SMB/NFS client response. Also, canceling requests after a timeout period can help in releasing system resources because the notification request is moved from a down/bad FPolicy server to an alternate FPolicy server.</p> <p>The range for this value is 0 through 100. If the value is set to 0, the option is disabled and cancel request messages are not sent to the FPolicy server. The default is 20s.</p>	<p>-reqs-cancel-timeout integer[h m s]</p>
<p><i>Timeout for aborting a request</i></p> <p>Specifies the timeout in hours (h), minutes (m), or seconds (s) for aborting a request.</p> <p>The range for this value is 0 through 200.</p>	<p>-reqs-abort-timeout integer[h m s]</p>
<p><i>Interval for sending status requests</i></p> <p>Specifies the interval in hours (h), minutes (m), or seconds (s) after which a status request is sent to the FPolicy server.</p> <p>The range for this value is 0 through 50. If the value is set to 0, the option is disabled and status request messages are not sent to the FPolicy server. The default is 10s.</p>	<p>-status-req-interval integer[h m s]</p>
<p><i>Maximum outstanding requests on the FPolicy server</i></p> <p>Specifies the maximum number of outstanding requests that can be queued on the FPolicy server.</p> <p>The range for this value is 1 through 10000. The default is 500.</p>	<p>-max-server-reqs integer</p>

<p><i>Timeout for disconnecting a nonresponsive FPolicy server</i></p> <p>Specifies the time interval in hours (h), minutes (m), or seconds (s) after which the connection to the FPolicy server is terminated.</p> <p>The connection is terminated after the timeout period only if the FPolicy server's queue contains the maximum allowed requests and no response is received within the timeout period. The maximum allowed number of requests is either 50 (the default) or the number specified by the <code>max-server-reqs</code> parameter.</p> <p>The range for this value is 1 through 100. The default is 60s.</p>	<pre>-server-progress -timeout integer[h m s]</pre>
<p><i>Interval for sending keep-alive messages to the FPolicy server</i></p> <p>Specifies the time interval in hours (h), minutes (m), or seconds (s) at which keep-alive messages are sent to the FPolicy server.</p> <p>Keep-alive messages detect half-open connections.</p> <p>The range for this value is 10 through 600. If the value is set to 0, the option is disabled and keep-alive messages are prevented from being sent to the FPolicy servers. The default is 120s.</p>	<pre>-keep-alive-interval-integer[h m s]</pre>
<p><i>Maximum reconnect attempts</i></p> <p>Specifies the maximum number of times the SVM attempts to reconnect to the FPolicy server after the connection has been broken.</p> <p>The range for this value is 0 through 20. The default is 5.</p>	<pre>-max-connection-retries integer</pre>
<p><i>Receive buffer size</i></p> <p>Specifies the receive buffer size of the connected socket for the FPolicy server.</p> <p>The default value is set to 256 kilobytes (Kb). When the value is set to 0, the size of the receive buffer is set to a value defined by the system.</p> <p>For example, if the default receive buffer size of the socket is 65536 bytes, by setting the tunable value to 0, the socket buffer size is set to 65536 bytes. You can use any non-default value to set the size (in bytes) of the receive buffer.</p>	<pre>-recv-buffer-size integer</pre>

<p><i>Send buffer size</i></p> <p>Specifies the send buffer size of the connected socket for the FPolicy server.</p> <p>The default value is set to 256 kilobytes (Kb). When the value is set to 0, the size of the send buffer is set to a value defined by the system.</p> <p>For example, if the default send buffer size of the socket is set to 65536 bytes, by setting the tunable value to 0, the socket buffer size is set to 65536 bytes. You can use any non-default value to set the size (in bytes) of the send buffer.</p>	<pre>-send-buffer-size integer</pre>
<p><i>Timeout for purging a session ID during reconnection</i></p> <p>Specifies the interval in hours (h), minutes (m), or seconds (s) after which a new session ID is sent to the FPolicy server during reconnection attempts.</p> <p>If the connection between the storage controller and the FPolicy server is terminated and reconnection is made within the <code>-session-timeout</code> interval, the old session ID is sent to FPolicy server so that it can send responses for old notifications.</p> <p>The default value is set to 10 seconds.</p>	<pre>-session-timeout [integerh][integerm][integer s]</pre>

Additional information about configuring FPolicy external engines to use SSL authenticated connections

You need to know some additional information if you want to configure the FPolicy external engine to use SSL when connecting to FPolicy servers.

SSL server authentication

If you choose to configure the FPolicy external engine for SSL server authentication, before creating the external engine, you must install the public certificate of the certificate authority (CA) that signed the FPolicy server certificate.

Mutual authentication

If you configure FPolicy external engines to use SSL mutual authentication when connecting storage virtual machine (SVM) data LIFs to external FPolicy servers, before creating the external engine, you must install the public certificate of the CA that signed the FPolicy server certificate along with the public certificate and key file for authentication of the SVM. You must not delete this certificate while any FPolicy policies are using the installed certificate.

If the certificate is deleted while FPolicy is using it for mutual authentication when connecting to an external FPolicy server, you cannot reenab a disabled FPolicy policy that uses that certificate. The FPolicy policy cannot be reenabled in this situation even if a new certificate with the same settings is created and installed on the SVM.

If the certificate has been deleted, you need to install a new certificate, create new FPolicy external engines that use the new certificate, and associate the new external engines with the FPolicy policy that you want to reenab by modifying the FPolicy policy.

Install certificates for SSL

The public certificate of the CA that is used to sign the FPolicy server certificate is installed by using the `security certificate install` command with the `-type` parameter set to `client-ca`. The private key and public certificate required for authentication of the SVM is installed by using the `security certificate install` command with the `-type` parameter set to `server`.

Certificates do not replicate in SVM disaster recovery relationships with a non-ID-preserve configuration

Security certificates used for SSL authentication when making connections to FPolicy servers do not replicate to SVM disaster recovery destinations with non-ID-preserve configurations. Although the FPolicy external-engine configuration on the SVM is replicated, security certificates are not replicated. You must manually install the security certificates on the destination.

When you set up the SVM disaster recovery relationship, the value you select for the `-identity-preserve` option of the `snapmirror create` command determines the configuration details that are replicated in the destination SVM.

If you set the `-identity-preserve` option to `true` (ID-preserve), all of the FPolicy configuration details are replicated, including the security certificate information. You must install the security certificates on the destination only if you set the option to `false` (non-ID-preserve).

Restrictions for cluster-scoped FPolicy external engines with MetroCluster and SVM disaster recovery configurations

You can create a cluster-scoped FPolicy external engine by assigning the cluster storage virtual machine (SVM) to the external engine. However, when creating a cluster-scoped external engine in a MetroCluster or SVM disaster recovery configuration, there are certain restrictions when choosing the authentication method that the SVM uses for external communication with the FPolicy server.

There are three authentication options that you can choose when creating external FPolicy servers: no authentication, SSL server authentication, and SSL mutual authentication. Although there are no restrictions when choosing the authentication option if the external FPolicy server is assigned to a data SVM, there are restrictions when creating a cluster-scoped FPolicy external engine:

Configuration	Permitted?
MetroCluster or SVM disaster recovery and a cluster-scoped FPolicy external engine with no authentication (SSL is not configured)	Yes
MetroCluster or SVM disaster recovery and a cluster-scoped FPolicy external engine with SSL server or SSL mutual authentication	No

- If a cluster-scoped FPolicy external engine with SSL authentication exists and you want to create a MetroCluster or SVM disaster recovery configuration, you must modify this external engine to use no authentication or remove the external engine before you can create the MetroCluster or SVM disaster recovery configuration.
- If the MetroCluster or SVM disaster recovery configuration already exists, ONTAP prevents you from creating a cluster-scoped FPolicy external engine with SSL authentication.

Complete the FPolicy external engine configuration worksheet

You can use this worksheet to record the values that you need during the FPolicy external engine configuration process. If a parameter value is required, you need to determine what value to use for those parameters before you configure the external engine.

Information for a basic external engine configuration

You should record whether you want to include each parameter setting in the external engine configuration and then record the value for the parameters that you want to include.

Type of information	Required	Include	Your values
Storage virtual machine (SVM) name	Yes	Yes	
Engine name	Yes	Yes	
Primary FPolicy servers	Yes	Yes	
Port number	Yes	Yes	
Secondary FPolicy servers	No		
External engine type	No		
SSL option for communication with external FPolicy server	Yes	Yes	
Certificate FQDN or custom common name	No		
Certificate serial number	No		
Certificate authority	No		

Information for advanced external engine parameters

To configure an external engine with advanced parameters, you must enter the configuration command while in advanced privilege mode.

Type of information	Required	Include	Your values
Timeout for canceling a request	No		
Timeout for aborting a request	No		
Interval for sending status requests	No		

Maximum outstanding requests on the FPolicy server	No		
Timeout for disconnecting a nonresponsive FPolicy server	No		
Interval for sending keep-alive messages to the FPolicy server	No		
Maximum reconnect attempts	No		
Receive buffer size	No		
Send buffer size	No		
Timeout for purging a session ID during reconnection	No		

Plan the FPolicy event configuration

Plan the FPolicy event configuration overview

Before you configure FPolicy events, you must understand what it means to create an FPolicy event. You must determine which protocols you want the event to monitor, which events to monitor, and which event filters to use. This information helps you plan the values that you want to set.

What it means to create an FPolicy event

Creating the FPolicy event means defining information that the FPolicy process needs to determine what file access operations to monitor and for which of the monitored events notifications should be sent to the external FPolicy server. The FPolicy event configuration defines the following configuration information:

- Storage virtual machine (SVM) name
- Event name
- Which protocols to monitor

FPolicy can monitor SMB, NFSv3, NFSv4, and, beginning in ONTAP 9.15.1, NFSv4.1 file access operations.

- Which file operations to monitor

Not all file operations are valid for each protocol.

- Which file filters to configure

Only certain combinations of file operations and filters are valid. Each protocol has its own set of supported combinations.

- Whether to monitor volume mount and unmount operations


There is a dependency with three of the parameters (`-protocol`, `-file-operations`, `-filters`). The following combinations are valid for the three parameters:



- You can specify the `-protocol` and `-file-operations` parameters.
- You can specify all three of the parameters.
- You can specify none of the parameters.

What the FPolicy event configuration contains

You can use the following list of available FPolicy event configuration parameters to help you plan your configuration:

Type of information	Option
<p>SVM</p> <p>Specifies the SVM name that you want to associate with this FPolicy event.</p> <p>Each FPolicy configuration is defined within a single SVM. The external engine, policy event, policy scope, and policy that combine together to create an FPolicy policy configuration must all be associated with the same SVM.</p>	<p><code>-vserver vserver_name</code></p>
<p>Event name</p> <p>Specifies the name to assign to the FPolicy event. When you create the FPolicy policy you associate the FPolicy event with the policy using the event name.</p> <p>The name can be up to 256 characters long.</p> <div style="border-left: 1px solid #ccc; padding-left: 10px; margin-top: 10px;"> <p> The name should be up to 200 characters long if configuring the event in a MetroCluster or SVM disaster recovery configuration.</p> </div> <p>The name can contain any combination of the following ASCII-range characters:</p> <ul style="list-style-type: none"> • a through z • A through Z • 0 through 9 • " _ ", "-", and "." 	<p><code>-event-name event_name</code></p>

Protocol

Specifies which protocol to configure for the FPolicy event. The list for `-protocol` can include one of the following values:

- `cifs`
- `nfsv3`
- `nfsv4`



If you specify `-protocol`, then you must specify a valid value in the `-file-operations` parameter. As the protocol version changes, the valid values might change.



Beginning in ONTAP 9.15.1, `nfsv4` allows you to capture NFSv4.0 and NFSv4.1 events.

`-protocol protocol`

File operations

Specifies the list of file operations for the FPolicy event.

The event checks the operations specified in this list from all client requests using the protocol specified in the `-protocol` parameter. You can list one or more file operations by using a comma-delimited list. The list for `-file-operations` can include one or more of the following values:

- `close` for file close operations
- `create` for file create operations
- `create-dir` for directory create operations
- `delete` for file delete operations
- `delete_dir` for directory delete operations
- `getattr` for get attribute operations
- `link` for link operations
- `lookup` for lookup operations
- `open` for file open operations
- `read` for file read operations
- `write` for file write operations
- `rename` for file rename operations
- `rename_dir` for directory rename operations
- `setattr` for set attribute operations
- `symlink` for symbolic link operations



If you specify `-file-operations`, then you must specify a valid protocol in the `-protocol` parameter.

```
-file-operations  
file_operations,...
```

Filters

`-filters filter, ...`

Specifies the list of filters for a given file operation for the specified protocol. The values in the `-filters` parameter are used to filter client requests. The list can include one or more of the following:



If you specify the `-filters` parameter, then you must also specify valid values for the `-file-operations` and `-protocol` parameters.

- `monitor-ads` option to filter the client request for alternate data stream.
- `close-with-modification` option to filter the client request for close with modification.
- `close-without-modification` option to filter the client request for close without modification.
- `first-read` option to filter the client request for first read.
- `first-write` option to filter the client request for first write.
- `offline-bit` option to filter the client request for offline bit set.

Setting this filter results in the FPolicy server receiving notification only when offline files are accessed.

- `open-with-delete-intent` option to filter the client request for open with delete intent.

Setting this filter results in the FPolicy server receiving notification only when an attempt is made to open a file with the intent to delete it. This is used by file systems when the `FILE_DELETE_ON_CLOSE` flag is specified.

- `open-with-write-intent` option to filter client request for open with write intent.

Setting this filter results in the FPolicy server receiving notification only when an attempt is made to open a file with the intent to write something in it.

- `write-with-size-change` option to filter the client request for write with size change.
- `setattr-with-owner-change` option to filter the client setattr requests for changing owner of a file or a directory.
- `setattr-with-group-change` option to filter the client setattr requests for changing the group of a file or a directory.
- `setattr-with-sacl-change` option to filter the client setattr requests for changing the SACL on a file or a directory.

This filter is available only for the SMB and NFSv4 protocols.

`setattr-with-dacl-change` option to filter the client setattr requests for changing the DACL on a file or a directory.

This filter is available only for the SMB and NFSv4 protocols.

<p><i>Is volume operation required</i></p> <p>Specifies whether monitoring is required for volume mount and unmount operations. The default is <code>false</code>.</p>	<pre>-volume-operation {true false} -filters filter, ...</pre>
<p><i>FPolicy access denied notifications</i></p> <p>Beginning with ONTAP 9.13.1, users can receive notifications for failed file operations due to lack of permissions. These notifications are valuable for security, ransomware protection, and governance. Notifications will be generated for file operation failed due to lack of permission, which includes:</p> <ul style="list-style-type: none"> • Failures due to NTFS permissions. • Failures due to Unix mode bits. • Failures due to NFSv4 ACLs. 	<pre>-monitor-fileop-failure {true false}</pre>

This option is available only for the SMB protocol.

Supported file operation and filter combinations that FPolicy can monitor for SMB

When you configure your FPolicy event, you need to be aware that only certain combinations of file operations and filters are supported for monitoring SMB file access operations. When this filter is specified, the directory operations are not monitored.

The list of supported file operation and filter combinations for FPolicy monitoring of SMB file access events is provided in the following table:

Supported file operations	Supported filters
close	monitor-ads, offline-bit, close-with-modification, close-without-modification, close-with-read, exclude-directory
create	monitor-ads, offline-bit
create_dir	Currently no filter is supported for this file operation.
delete	monitor-ads, offline-bit
delete_dir	Currently no filter is supported for this file operation.
getattr	offline-bit, exclude-dir
open	monitor-ads, offline-bit, open-with-delete-intent, open-with-write-intent, exclude-dir
read	monitor-ads, offline-bit, first-read
write	monitor-ads, offline-bit, first-write, write-with-size-change

rename	monitor-ads, offline-bit
rename_dir	Currently no filter is supported for this file operation.
setattr	monitor-ads, offline-bit, setattr_with_owner_change, setattr_with_group_change, setattr_with_mode_change, setattr_with_sacl_change, setattr_with_dacl_change, setattr_with_modify_time_change, setattr_with_access_time_change, setattr_with_creation_time_change, setattr_with_size_change, setattr_with_allocation_size_change, exclude_directory

Beginning with ONTAP 9.13.1, users can receive notifications for failed file operations due to lack of permissions. The list of supported access denied file operation and filter combinations for FPolicy monitoring of SMB file access events is provided in the following table:

Supported access denied file operation	Supported filters
open	NA

Supported file operation and filter combinations that FPolicy can monitor for NFSv3

When you configure your FPolicy event, you need to be aware that only certain combinations of file operations and filters are supported for monitoring NFSv3 file access operations.

The list of supported file operation and filter combinations for FPolicy monitoring of NFSv3 file access events is provided in the following table:

Supported file operations	Supported filters
create	offline-bit
create_dir	Currently no filter is supported for this file operation.
delete	offline-bit
delete_dir	Currently no filter is supported for this file operation.
link	offline-bit
lookup	offline-bit, exclude-dir
read	offline-bit, first-read
write	offline-bit, first-write, write-with-size-change

rename	offline-bit
rename_dir	Currently no filter is supported for this file operation.
setattr	offline-bit, setattr_with_owner_change, setattr_with_group_change, setattr_with_mode_change, setattr_with_modify_time_change, setattr_with_access_time_change, setattr_with_size_change, exclude_directory
symlink	offline-bit

Beginning with ONTAP 9.13.1, users can receive notifications for failed file operations due to lack of permissions. The list of supported access denied file operation and filter combinations for FPolicy monitoring of NFSv3 file access events is provided in the following table:

Supported access denied file operation	Supported filters
access	NA
create	NA
create_dir	NA
delete	NA
delete_dir	NA
link	NA
read	NA
rename	NA
rename_dir	NA
setattr	NA
write	NA

Supported file operation and filter combinations that FPolicy can monitor for NFSv4

When you configure your FPolicy event, you need to be aware that only certain combinations of file operations and filters are supported for monitoring NFSv4 file access operations.

Beginning with ONTAP 9.15.1, FPolicy supports the NFSv4.1 protocol.

The list of supported file operation and filter combinations for FPolicy monitoring of NFSv4 or NFSv4.1 file access events is provided in the following table:

Supported file operations	Supported filters
close	offline-bit, exclude-directory
create	offline-bit
create_dir	Currently no filter is supported for this file operation.
delete	offline-bit
delete_dir	Currently no filter is supported for this file operation.
getattr	offline-bit, exclude-directory
link	offline-bit
lookup	offline-bit, exclude-directory
open	offline-bit, exclude-directory
read	offline-bit, first-read
write	offline-bit, first-write, write-with-size-change
rename	offline-bit
rename_dir	Currently no filter is supported for this file operation.
setattr	offline-bit, setattr_with_owner_change, setattr_with_group_change, setattr_with_mode_change, setattr_with_sacl_change, setattr_with_dacl_change, setattr_with_modify_time_change, setattr_with_access_time_change, setattr_with_size_change, exclude_directory
symlink	offline-bit

Beginning with ONTAP 9.13.1, users can receive notifications for failed file operations due to lack of permissions. The list of supported access denied file operation and filter combinations for FPolicy monitoring of NFSv4 or NFSv4.1 file access events is provided in the following table:

Supported access denied file operation	Supported filters
-----------------------------------------------	--------------------------

access	NA
create	NA
create_dir	NA
delete	NA
delete_dir	NA
link	NA
open	NA
read	NA
rename	NA
rename_dir	NA
setattr	NA
write	NA

Complete the FPolicy event configuration worksheet

You can use this worksheet to record the values that you need during the FPolicy event configuration process. If a parameter value is required, you need to determine what value to use for those parameters before you configure the FPolicy event.

You should record whether you want to include each parameter setting in the FPolicy event configuration and then record the value for the parameters that you want to include.

Type of information	Required	Include	Your values
Storage virtual machine (SVM) name	Yes	Yes	
Event name	Yes	Yes	
Protocol	No		
File operations	No		
Filters	No		

Volume operation	No		
Access denied events (support beginning with ONTAP 9.13)	No		

Plan the FPolicy policy configuration

Plan the FPolicy policy configuration overview

Before you configure the FPolicy policy, you must understand which parameters are required when creating the policy as well as why you might want to configure certain optional parameters. This information helps you to determine which values to set for each parameter.

When creating an FPolicy policy you associate the policy with the following:


- The storage virtual machine (SVM)
- One or more FPolicy events
- An FPolicy external engine

You can also configure several optional policy settings.

What the FPolicy policy configuration contains

You can use the following list of available FPolicy policy required and optional parameters to help you plan your configuration:

Type of information	Option	Required	Default
<p><i>SVM name</i></p> <p>Specifies the name of the SVM on which you want to create an FPolicy policy.</p>	<p>-vserver vserver_name</p>	Yes	None

<p>Policy name</p> <p>Specifies the name of the FPolicy policy.</p> <p>The name can be up to 256 characters long.</p> <div style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;">  <p>The name should be up to 200 characters long if configuring the policy in a MetroCluster or SVM disaster recovery configuration.</p> </div> <p>The name can contain any combination of the following ASCII-range characters:</p> <ul style="list-style-type: none"> • a through z • A through Z • 0 through 9 • “_”, “-”, and “.” 	<p><code>-policy-name policy_name</code></p>	<p>Yes</p>	<p>None</p>
<p>Event names</p> <p>Specifies a comma-delimited list of events to associate with the FPolicy policy.</p> <ul style="list-style-type: none"> • You can associate more than one event to a policy. • An event is specific to a protocol. • You can use a single policy to monitor file access events for more than one protocol by creating an event for each protocol that you want the policy to monitor, and then associating the events to the policy. • The events must already exist. 	<p><code>-events event_name, ...</code></p>	<p>Yes</p>	<p>None</p>
<p>Persistent store</p> <p>Beginning with ONTAP 9.14.1, this parameter specifies the persistent store to capture file access events for asynchronous non-mandatory policies in the SVM.</p>	<p><code>-persistent -store persistent_store_name</code></p>	<p>No</p>	<p>None</p>

<p><i>External engine name</i></p> <p>Specifies the name of the external engine to associate with the FPolicy policy.</p> <ul style="list-style-type: none"> • An external engine contains information required by the node to send notifications to an FPolicy server. • You can configure FPolicy to use the ONTAP native external engine for simple file blocking or to use an external engine that is configured to use external FPolicy servers (FPolicy servers) for more sophisticated file blocking and file management. • If you want to use the native external engine, you can either not specify a value for this parameter or you can specify <code>native</code> as the value. • If you want to use FPolicy servers, the configuration for the external engine must already exist. 	<p><code>-engine engine_name</code></p>	<p>Yes (unless the policy uses the internal ONTAP native engine)</p>	<p><code>native</code></p>
<p><i>Is mandatory screening required</i></p> <p>Specifies whether mandatory file access screening is required.</p> <ul style="list-style-type: none"> • The mandatory screening setting determines what action is taken on a file access event in a case when all primary and secondary servers are down or no response is received from the FPolicy servers within a given timeout period. • When set to <code>true</code>, file access events are denied. • When set to <code>false</code>, file access events are allowed. 	<p><code>-is-mandatory {true false}</code></p>	<p>No</p>	<p><code>true</code></p>

<p><i>Allow privileged access</i></p> <p>Specifies whether you want the FPolicy server to have privileged access to the monitored files and folders by using a privileged data connection.</p> <p>If configured, FPolicy servers can access files from the root of the SVM containing the monitored data using the privileged data connection.</p> <p>For privileged data access, SMB must be licensed on the cluster and all the data LIFs used to connect to the FPolicy servers must be configured to have <code>cifs</code> as one of the allowed protocols.</p> <p>If you want to configure the policy to allow privileged access, you must also specify the user name for the account that you want the FPolicy server to use for privileged access.</p>	<pre>-allow -privileged -access {yes no}</pre>	<p>No (unless passthrough-read is enabled)</p>	<p>no</p>
<p><i>Privileged user name</i></p> <p>Specifies the user name of the account the FPolicy servers use for privileged data access.</p> <ul style="list-style-type: none"> • The value for this parameter should use the “domain\user name” format. • If <code>-allow-privileged-access</code> is set to <code>no</code>, any value set for this parameter is ignored. 	<pre>-privileged -user-name user_name</pre>	<p>No (unless privileged access is enabled)</p>	<p>None</p>

<p><i>Allow passthrough-read</i></p> <p>Specifies whether the FPolicy servers can provide passthrough-read services for files that have been archived to secondary storage (offline files) by the FPolicy servers:</p> <ul style="list-style-type: none"> • Passthrough-read is a way to read data for offline files without restoring the data to the primary storage. <p>Passthrough-read reduces response latencies because there is no need to recall files back to primary storage before responding to the read request. Additionally, passthrough-read optimizes storage efficiency by eliminating the need to consume primary storage space with files that are recalled solely to satisfy read requests.</p> <ul style="list-style-type: none"> • When enabled, the FPolicy servers provide the data for the file over a separate privileged data channel opened specifically for passthrough-reads. • If you want to configure passthrough-read, the policy must also be configured to allow privileged access. 	<pre>-is-passthrough -read-enabled {true false}</pre>	<p>No</p>	<p>false</p>
------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-------------------------------------------------------	-----------	--------------

Requirement for FPolicy scope configurations if the FPolicy policy uses the native engine

If you configure the FPolicy policy to use the native engine, there is a specific requirement for how you define the FPolicy scope configured for the policy.

The FPolicy scope defines the boundaries on which the FPolicy policy applies, for example whether the FPolicy applies to specified volumes or shares. There are a number of parameters that further restrict the scope to which the FPolicy policy applies. One of these parameters, `-is-file-extension-check-on-directories-enabled`, specifies whether to check file extensions on directories. The default value is `false`, which means that file extensions on directories are not checked.

When an FPolicy policy that uses the native engine is enabled on a share or volume and the `-is-file-extension-check-on-directories-enabled` parameter is set to `false` for the scope of the policy, directory access is denied. With this configuration, because the file extensions are not checked for directories, any directory operation is denied if it falls under the scope of the policy.

To ensure that directory access succeeds when using the native engine, you must set the `-is-file-extension-check-on-directories-enabled` parameter to `true` when creating the scope.

With this parameter set to `true`, extension checks happen for directory operations and the decision whether to

allow or deny access is taken based on the extensions included or excluded in the FPolicy scope configuration.

Complete the FPolicy policy worksheet

You can use this worksheet to record the values that you need during the FPolicy policy configuration process. You should record whether you want to include each parameter setting in the FPolicy policy configuration and then record the value for the parameters that you want to include.

Type of information	Include	Your values
Storage virtual machine (SVM) name	Yes	
Policy name	Yes	
Event names	Yes	
Persistent store		
External engine name		
Is mandatory screening required?		
Allow privileged access		
Privileged user name		
Is passthrough-read enabled?		

Plan the FPolicy scope configuration

Plan the FPolicy scope configuration overview

Before you configure the FPolicy scope, you must understand what it means to create a scope. You must understand what the scope configuration contains. You also need to understand what the scope rules of precedence are. This information can help you plan the values that you want to set.

What it means to create an FPolicy scope

Creating the FPolicy scope means defining the boundaries on which the FPolicy policy applies. The storage virtual machine (SVM) is the basic boundary. When you create a scope for an FPolicy policy, you must define the FPolicy policy to which it will apply, and you must designate to which SVM you want to apply the scope.

There are a number of parameters that further restrict the scope within the specified SVM. You can restrict the scope by specifying what to include in the scope or by specifying what to exclude from the scope. After you apply a scope to an enabled policy, policy event checks get applied to the scope defined by this command.

Notifications are generated for file access events where matches are found in the “include” options.

Notifications are not generated for file access events where matches are found in the “exclude” options.

The FPolicy scope configuration defines the following configuration information:

- SVM name
- Policy name
- The shares to include or exclude from what gets monitored
- The export policies to include or exclude from what gets monitored
- The volumes to include or exclude from what gets monitored
- The file extensions to include or exclude from what gets monitored
- Whether to do file extension checks on directory objects



There are special considerations for the scope for a cluster FPolicy policy. The cluster FPolicy policy is a policy that the cluster administrator creates for the admin SVM. If the cluster administrator also creates the scope for that cluster FPolicy policy, the SVM administrator cannot create a scope for that same policy. However, if the cluster administrator does not create a scope for the cluster FPolicy policy, then any SVM administrator can create the scope for that cluster policy. If the SVM administrator creates a scope for that cluster FPolicy policy, the cluster administrator cannot subsequently create a cluster scope for that same cluster policy. This is because the cluster administrator cannot override the scope for the same cluster policy.

What the scope rules of precedence are

The following rules of precedence apply to scope configurations:

- When a share is included in the `-shares-to-include` parameter and the parent volume of the share is included in the `-volumes-to-exclude` parameter, `-volumes-to-exclude` has precedence over `-shares-to-include`.
- When an export policy is included in the `-export-policies-to-include` parameter and the parent volume of the export policy is included in the `-volumes-to-exclude` parameter, `-volumes-to-exclude` has precedence over `-export-policies-to-include`.
- An administrator can specify both `-file-extensions-to-include` and `-file-extensions-to-exclude` lists.

The `-file-extensions-to-exclude` parameter is checked before the `-file-extensions-to-include` parameter is checked.

What the FPolicy scope configuration contains

You can use the following list of available FPolicy scope configuration parameters to help you plan your configuration:



When configuring what shares, export policies, volumes, and file extensions to include or exclude from the scope, the include and exclude parameters can include metacharacters such as “?” and “*”. The use of regular expressions is not supported.

Type of information	Option
---------------------	--------

<p>SVM</p> <p>Specifies the SVM name on which you want to create an FPolicy scope.</p> <p>Each FPolicy configuration is defined within a single SVM. The external engine, policy event, policy scope, and policy that combine together to create an FPolicy policy configuration must all be associated with the same SVM.</p>	<p>-vserver vserver_name</p>
<p>Policy name</p> <p>Specifies the name of the FPolicy policy to which you want to attach the scope. The FPolicy policy must already exist.</p>	<p>-policy-name policy_name</p>
<p>Shares to include</p> <p>Specifies a comma-delimited list of shares to monitor for the FPolicy policy to which the scope is applied.</p>	<p>-shares-to-include share_name, ...</p>
<p>Shares to exclude</p> <p>Specifies a comma-delimited list of shares to exclude from monitoring for the FPolicy policy to which the scope is applied.</p>	<p>-shares-to-exclude share_name, ...</p>
<p>Volumes to include Specifies a comma-delimited list of volumes to monitor for the FPolicy policy to which the scope is applied.</p>	<p>-volumes-to-include volume_name, ...</p>
<p>Volumes to exclude</p> <p>Specifies a comma-delimited list of volumes to exclude from monitoring for the FPolicy policy to which the scope is applied.</p>	<p>-volumes-to-exclude volume_name, ...</p>
<p>Export policies to include</p> <p>Specifies a comma-delimited list of export policies to monitor for the FPolicy policy to which the scope is applied.</p>	<p>-export-policies-to-include export_policy_name, ...</p>
<p>Export policies to exclude</p> <p>Specifies a comma-delimited list of export policies to exclude from monitoring for the FPolicy policy to which the scope is applied.</p>	<p>-export-policies-to-exclude export_policy_name, ...</p>
<p>File extensions to include</p> <p>Specifies a comma-delimited list of file extensions to monitor for the FPolicy policy to which the scope is applied.</p>	<p>-file-extensions-to-include file_extensions, ...</p>

<p><i>File extension to exclude</i></p> <p>Specifies a comma-delimited list of file extensions to exclude from monitoring for the FPolicy policy to which the scope is applied.</p>	<pre>-file-extensions-to-exclude file_extensions,...</pre>
<p><i>Is file extension check on directory enabled ?</i></p> <p>Specifies whether the file name extension checks apply to directory objects as well. If this parameter is set to <code>true</code>, the directory objects are subjected to the same extension checks as regular files. If this parameter is set to <code>false</code>, the directory names are not matched for extensions and notifications are sent for directories even if their name extensions do not match.</p> <p>If the FPolicy policy to which the scope is assigned is configured to use the native engine, this parameter must be set to <code>true</code>.</p>	<pre>-is-file-extension-check-on-directories-enabled {true false}</pre>

Complete the FPolicy scope worksheet

You can use this worksheet to record the values that you need during the FPolicy scope configuration process. If a parameter value is required, you need to determine what value to use for those parameters before you configure the FPolicy scope.

You should record whether you want to include each parameter setting in the FPolicy scope configuration and then record the value for the parameters that you want to include.

Type of information	Required	Include	Your values
Storage virtual machine (SVM) name	Yes	Yes	
Policy name	Yes	Yes	
Shares to include	No		
Shares to exclude	No		
Volumes to include	No		
Volumes to exclude	No		
Export policies to include	No		
Export policies to exclude	No		
File extensions to include	No		
File extension to exclude	No		

Is file extension check on directory enabled?	No		
-----------------------------------------------	----	--	--

Create the FPolicy configuration

Create the FPolicy external engine

You must create an external engine to start creating an FPolicy configuration. The external engine defines how FPolicy makes and manages connections to external FPolicy servers. If your configuration uses the internal ONTAP engine (the native external engine) for simple file blocking, you do not need to configure a separate FPolicy external engine and do not need to perform this step.

What you'll need

The [external engine](#) worksheet should be completed.

About this task

If the external engine is used in a MetroCluster configuration, you should specify the IP addresses of the FPolicy servers at the source site as primary servers. The IP addresses of the FPolicy servers at the destination site should be specified as secondary servers.

Steps

1. Create the FPolicy external engine by using the `vserver fpolicy policy external-engine create` command.

The following command creates an external engine on storage virtual machine (SVM) `vs1.example.com`. No authentication is required for external communications with the FPolicy server.

```
vserver fpolicy policy external-engine create -vserver-name vs1.example.com
-engine-name engine1 -primary-servers 10.1.1.2,10.1.1.3 -port 6789 -ssl-option
no-auth
```

2. Verify the FPolicy external engine configuration by using the `vserver fpolicy policy external-engine show` command.

The following command display information about all external engines configured on SVM `vs1.example.com`:

```
vserver fpolicy policy external-engine show -vserver vs1.example.com
```

External Vserver Type	Engine	Primary Servers	Secondary Servers	Port	Engine
vs1.example.com synchronous	engine1	10.1.1.2, 10.1.1.3	-	6789	

The following command displays detailed information about the external engine named “engine1” on SVM vs1.example.com:

```
vserver fpolicy policy external-engine show -vserver vs1.example.com -engine
-name engine1
```

```

Vserver: vs1.example.com
Engine: engine1
Primary FPolicy Servers: 10.1.1.2, 10.1.1.3
Port Number of FPolicy Service: 6789
Secondary FPolicy Servers: -
External Engine Type: synchronous
SSL Option for External Communication: no-auth
FQDN or Custom Common Name: -
Serial Number of Certificate: -
Certificate Authority: -

```

Create the FPolicy event

As part of creating an FPolicy policy configuration, you need to create an FPolicy event. You associate the event with the FPolicy policy when it is created. An event defines which protocol to monitor and which file access events to monitor and filter.

Before you begin

You should complete the FPolicy event [worksheet](#).

Create the FPolicy event

1. Create the FPolicy event by using the `vserver fpolicy policy event create` command.

```
vserver fpolicy policy event create -vserver vs1.example.com -event-name
event1 -protocol cifs -file-operations open,close,read,write
```

2. Verify the FPolicy event configuration by using the `vserver fpolicy policy event show` command.

```
vserver fpolicy policy event show -vserver vs1.example.com
```

Vserver	Event Name	Protocols	File Operations	Filters	Is Volume Operation
vs1.example.com	event1	cifs	open, close, read, write	-	false

Create the FPolicy access denied events

Beginning with ONTAP 9.13.1, users can receive notifications for failed file operations due to lack of permissions. These notifications are valuable for security, ransomware protection, and governance.

1. Create the FPolicy event by using the `vserver fpolicy policy event create` command.

```
vserver fpolicy policy event create -vserver vs1.example.com -event-name event1 -protocol cifs -monitor-fileop-failure true -file-operations open
```

Create FPolicy persistent stores

Persistent stores can help decouple client I/O processing from FPolicy notification processing to reduce client latency. Beginning with ONTAP 9.14.1, FPolicy allows you to set up [persistent stores](#) to capture file access events for asynchronous non-mandatory policies in the SVM. Synchronous (either mandatory or non-mandatory) and asynchronous mandatory configurations are not supported.

Beginning with ONTAP 9.15.1, FPolicy persistent store configuration is simplified. The `persistent-store create` command automates volume creation for the SVM and configures the volume for the persistent store.

There are two ways to create a persistent store, depending on the ONTAP release:

- ONTAP 9.15.1 or later: When you create the persistent store, ONTAP automatically creates and configures its volume at the same time. This simplifies FPolicy persistent store configuration and implements all best practices.
- ONTAP 9.14.1: Manually create and configure a volume and then create a persistent store for the newly created volume.

Only one persistent store can be set up on each SVM. This single persistent store needs to be used for all FPolicy configurations on that SVM, even if the policies are from different partners.

Create a persistent store (ONTAP 9.15.1 or later)

Beginning with ONTAP 9.15.1, use the `fpolicy persistent-store create` command to create the FPolicy persistent store with inline volume creation and configuration. ONTAP automatically blocks the volume from external user protocol access (CIFS/NFS).

Before you begin

- The SVM where you want to create the persistent store must have at least one aggregate.
- You should have access to the aggregates available for the SVM and sufficient permissions to create

volumes.

Steps

1. Create the persistent store, which creates and configures the volume automatically:

```
vserver fpolicy persistent-store create -vserver <vserver> -persistent-store <name> -volume <volume_name> -size <size> -autosize-mode <off|grow|grow_shrink>
```

- The `vserver` parameter is the name of the SVM.
- The `persistent-store` parameter is the name of the persistent store.
- The `volume` parameter is the name of the persistent store volume.



If you want to use an existing, empty volume, use the `volume show` command to find it and specify it in the `volume` parameter.

- The `size` parameter is based on the time duration for which you want to persist the events that are not delivered to the external server (partner application).

For example, if you want 30 minutes of events to persist in a cluster with a 30K notifications per second capacity:

Required Volume Size = 30000 x 30 x 60 x 0.6KB (average notification record size) = 32400000 KB = ~32 GB

To find the approximate notification rate, you can either reach out to your FPolicy partner application or utilize the FPolicy counter `requests_dispatched_rate`.



If you are using an existing volume, the `size` parameter is optional. If you do provide a value for the `size` parameter, it will modify the volume with the size you specify.

- The `autosize-mode` parameter specifies the autosize mode for the volume. The supported autosize modes are:
 - `off` - The volume does not grow or shrink in size in response to the amount of used space.
 - `grow` - The volume automatically grows when used space in the volume is above the grow threshold.
 - `grow_shrink` - The volume grows or shrinks in size in response to the amount of used space.

2. Create the FPolicy policy and add the persistent store name to that policy. For more information, see [Create the FPolicy policy](#).

Create a persistent store (ONTAP 9.14.1)

You can create a volume, and then create a persistent store to use that volume. You can then block the newly created volume from external user protocol access (CIFS/NFS).

Steps

1. Create an empty volume on the SVM that can be provisioned for the persistent store:

```
volume create -vserver <SVM Name> -volume <volume> -state <online> -policy <default> -unix-permissions <777> -size <value> -aggregate <aggregate name>
```

`-snapshot-policy <none>`

It is expected that an administrator user with sufficient RBAC privileges (to create a volume) creates a volume (using the volume cli command or REST API) of the desired size and provide the name of that volume as the `-volume` in the persistent store create CLI command or REST API.

- The `vserver` parameter is the name of the SVM.
- The `volume` parameter is the name of the persistent store volume.
- The `state` parameter should be set to `online` so the volume is available for use.
- The `policy` parameter is set to the FPolicy service policy, if you have one already configured. If not, you can use the `volume modify` command later to add the policy.
- The `unix-permissions` parameter is optional.
- The `size` parameter is based on the time duration for which you want to persist the events that are not delivered to the external server (partner application).

For example, if you want 30 minutes of events to persist in a cluster with a 30K notifications per second capacity:

Required Volume Size = $30000 \times 30 \times 60 \times 0.6\text{KB}$ (average notification record size) = 32400000 KB = ~32 GB

To find the approximate notification rate, you can either reach out to your FPolicy partner application or utilize the FPolicy counter `requests_dispatched_rate`.

- The `aggregate` parameter is needed for FlexVol volumes, otherwise it is not required.
- The `snapshot-policy` parameter must be set to `none`. This ensures that there is no accidental restore of the snapshot leading to loss of current events and prevents possible duplicate event processing.

If you want to use an existing, empty volume, use the `volume show` command to find it and the `volume modify` command to make any needed alterations. Ensure the `policy`, `size`, and `snapshot-policy` parameters are set correctly for the persistent store.

2. Create the persistent store:

```
vserver fpolicy persistent store create -vserver <SVM> -persistent-store <PS_name> -volume <volume>
```

- The `vserver` parameter is the name of the SVM.
- The `persistent-store` parameter is the name of the persistent store.
- The `volume` parameter is the name of the persistent store volume.

3. Create the FPolicy policy and add the persistent store name to that policy. For more information, see [Create the FPolicy policy](#).

Create the FPolicy policy

When you create the FPolicy policy, you associate an external engine and one or more events to the policy. The policy also specifies whether mandatory screening is required,

whether the FPolicy servers have privileged access to data on the storage virtual machine (SVM), and whether passthrough-read for offline files is enabled.

What you'll need

- The FPolicy policy worksheet should be completed.
- If you plan on configuring the policy to use FPolicy servers, the external engine must exist.
- At least one FPolicy event that you plan on associating with the FPolicy policy must exist.
- If you want to configure privileged data access, a SMB server must exist on the SVM.
- To configure a persistent store for a policy, the engine type must be **async** and the policy must be **non-mandatory**.

For more information, see [Create persistent stores](#).

Steps

1. Create the FPolicy policy:

```
vserver fpolicy policy create -vserver-name vserver_name -policy-name
policy_name -engine engine_name -events event_name, [-persistent-store
PS_name] [-is-mandatory {true|false}] [-allow-privileged-access {yes|no}] [-
privileged-user-name domain\user_name] [-is-passthrough-read-enabled
{true|false}]
```

- You can add one or more events to the FPolicy policy.
- By default, mandatory screening is enabled.
- If you want to allow privileged access by setting the `-allow-privileged-access` parameter to `yes`, you must also configure a privileged user name for privileged access.
- If you want to configure passthrough-read by setting the `-is-passthrough-read-enabled` parameter to `true`, you must also configure privileged data access.

The following command creates a policy named “policy1” that has the event named “event1” and the external engine named “engine1” associated with it. This policy uses default values in the policy configuration:

```
vserver fpolicy policy create -vserver vs1.example.com -policy-name policy1
-events event1 -engine engine1
```

The following command creates a policy named “policy2” that has the event named “event2” and the external engine named “engine2” associated with it. This policy is configured to use privileged access using the specified user name. Passthrough-read is enabled:

```
vserver fpolicy policy create -vserver vs1.example.com -policy-name policy2
-events event2 -engine engine2 -allow-privileged-access yes -privileged-
user-name example\archive_acct -is-passthrough-read-enabled true
```

The following command creates a policy named “native1” that has the event named “event3” associated with it. This policy uses the native engine and uses default values in the policy configuration:

```
vserver fpolicy policy create -vserver vs1.example.com -policy-name native1
-events event3 -engine native
```

2. Verify the FPolicy policy configuration by using the `vserver fpolicy policy show` command.

The following command displays information about the three configured FPolicy policies, including the following information:

- The SVM associated with the policy
- The external engine associated with the policy
- The events associated with the policy
- Whether mandatory screening is required
- Whether privileged access is required

```
vserver fpolicy policy show
```

Vserver	Policy Name	Events	Engine	Is Mandatory	Privileged Access
vs1.example.com	policy1	event1	engine1	true	no
vs1.example.com	policy2	event2	engine2	true	yes
vs1.example.com	native1	event3	native	true	no

Create the FPolicy scope

After creating the FPolicy policy, you need to create an FPolicy scope. When creating the scope, you associate the scope with an FPolicy policy. A scope defines the boundaries on which the FPolicy policy applies. Scopes can include or exclude files based on shares, export policies, volumes, and file extensions.

What you'll need

The FPolicy scope worksheet must be completed. The FPolicy policy must exist with an associated external engine (if the policy is configured to use external FPolicy servers) and must have at least one associated FPolicy event.

Steps

1. Create the FPolicy scope by using the `vserver fpolicy policy scope create` command.

```
vserver fpolicy policy scope create -vserver-name vs1.example.com -policy-name policy1 -volumes-to-include datavol1,datavol2
```

2. Verify the FPolicy scope configuration by using the `vserver fpolicy policy scope show` command.

```
vserver fpolicy policy scope show -vserver vs1.example.com -instance
```

```

Vserver: vs1.example.com
Policy: policy1
Shares to Include: -
Shares to Exclude: -
Volumes to Include: datavol1, datavol2
Volumes to Exclude: -
Export Policies to Include: -
Export Policies to Exclude: -
File Extensions to Include: -
File Extensions to Exclude: -

```

Enable the FPolicy policy

After you are through configuring an FPolicy policy configuration, you enable the FPolicy policy. Enabling the policy sets its priority and starts file access monitoring for the policy.

What you'll need

The FPolicy policy must exist with an associated external engine (if the policy is configured to use external FPolicy servers) and must have at least one associated FPolicy event. The FPolicy policy scope must exist and must be assigned to the FPolicy policy.

About this task

The priority is used when multiple policies are enabled on the storage virtual machine (SVM) and more than one policy has subscribed to the same file access event. Policies that use the native engine configuration have a higher priority than policies for any other engine, regardless of the sequence number assigned to them when enabling the policy.



A policy cannot be enabled on the admin SVM.

Steps

1. Enable the FPolicy policy by using the `vserver fpolicy enable` command.

```
vserver fpolicy enable -vserver-name vs1.example.com -policy-name policy1
-sequence-number 1
```

2. Verify that the FPolicy policy is enabled by using the `vserver fpolicy show` command.

```
vserver fpolicy show -vserver vs1.example.com
```

Vserver	Policy Name	Sequence Number	Status	Engine
vs1.example.com	policy1	1	on	engine1

Manage FPolicy configurations

Modify FPolicy configurations

Commands for modifying FPolicy configurations

You can modify FPolicy configurations by modifying the elements that make up the configuration. You can modify external engines, FPolicy events, FPolicy scopes, FPolicy persistent stores, and FPolicy policies. You can also enable or disable FPolicy policies. When you disable the FPolicy policy, file monitoring is discontinued for that policy.

You should disable an FPolicy policy before modifying its configuration.

If you want to modify...	Use this command...
External engines	<code>vserver fpolicy policy external-engine modify</code>
Events	<code>vserver fpolicy policy event modify</code>
Scopes	<code>vserver fpolicy policy scope modify</code>
Persistent store	<code>vserver fpolicy persistent-store modify</code>
Policies	<code>vserver fpolicy policy modify</code>

See the man pages for the commands for more information.

Enable or disable FPolicy policies

You can enable FPolicy policies after the configuration is complete. Enabling the policy sets its priority and starts file access monitoring for the policy. You can disable FPolicy policies if you want to stop file access monitoring for the policy.

What you'll need

Before enabling FPolicy policies, the FPolicy configuration must be completed.

About this task

- The priority is used when multiple policies are enabled on the storage virtual machine (SVM) and more than one policy has subscribed to the same file access event.
- Policies that use the native engine configuration have a higher priority than policies for any other engine, regardless of the sequence number assigned to them when enabling the policy.
- If you want to change the priority of an FPolicy policy, you must disable the policy and then reenabling it using the new sequence number.

Step

1. Perform the appropriate action:

If you want to...	Enter the following command...
-------------------	--------------------------------

Enable an FPolicy policy	<code>vserver fpolicy enable -vserver-name vserver_name -policy-name policy_name -sequence-number integer</code>
Disable an FPolicy policy	<code>vserver fpolicy disable -vserver-name vserver_name -policy-name policy_name</code>

Display information about FPolicy configurations

How the show commands work

It is helpful when displaying information about the FPolicy configuration to understand how the `show` commands work.

A `show` command without additional parameters displays information in a summary form. Additionally, every `show` command has the same two mutually exclusive optional parameters, `-instance` and `-fields`.

When you use the `-instance` parameter with a `show` command, the command output displays detailed information in a list format. In some cases, the detailed output can be lengthy and include more information than you need. You can use the `-fields fieldname[,fieldname...]` parameter to customize the output so that it displays information only for the fields you specify. You can identify which fields that you can specify by entering `?` after the `-fields` parameter.



The output of a `show` command with the `-fields` parameter might display other relevant and necessary fields related to the requested fields.

Every `show` command has one or more optional parameters that filter that output and enable you to narrow the scope of information displayed in command output. You can identify which optional parameters are available for a command by entering `?` after the `show` command.

The `show` command supports UNIX-style patterns and wildcards to enable you to match multiple values in command-parameters arguments. For example, you can use the wildcard operator (`*`), the NOT operator (`!`), the OR operator (`()`), the range operator (`integer...integer`), the less-than operator (`<`), the greater-than operator (`>`), the less-than or equal to operator (`<=`), and the greater-than or equal to operator (`>=`) when specifying values.

For more information about using UNIX-style patterns and wildcards, see the [Using the ONTAP command-line interface](#).

Commands for displaying information about FPolicy configurations

You use the `fpolicy show` commands to display information about the FPolicy configuration, including information about FPolicy external engines, events, scopes, and policies.

If you want to display information about FPolicy...	Use this command...
External engines	<code>vserver fpolicy policy external-engine show</code>

Events	<code>vserver fpolicy policy event show</code>
Scopes	<code>vserver fpolicy policy scope show</code>
Policies	<code>vserver fpolicy policy show</code>

See the man pages for the commands for more information.

Display information about FPolicy policy status

You can display information about the status for FPolicy policies to determine whether a policy is enabled, what external engine it is configured to use, what the sequence number is for the policy, and to which storage virtual machine (SVM) the FPolicy policy is associated.

About this task

If you do not specify any parameters, the command displays the following information:

- SVM name
- Policy name
- Policy sequence number
- Policy status

In addition to displaying information about policy status for FPolicy policies configured on the cluster or a specific SVM, you can use command parameters to filter the command's output by other criteria.

You can specify the `-instance` parameter to display detailed information about listed policies. Alternatively, you can use the `-fields` parameter to display only the indicated fields in the command output, or `-fields ?` to determine what fields you can use.

Step

1. Display filtered information about FPolicy policy status by using the appropriate command:

If you want to display status information about policies...	Enter the command...
On the cluster	<code>vserver fpolicy show</code>
That have the specified status	<code>vserver fpolicy show -status {on off}</code>
On a specified SVM	<code>vserver fpolicy show -vserver vserver_name</code>
With the specified policy name	<code>vserver fpolicy show -policy-name policy_name</code>
That use the specified external engine	<code>vserver fpolicy show -engine engine_name</code>

Example

The following example displays the information about FPolicy policies on the cluster:

```
cluster1::> vserver fpolicy show
```

Vserver	Policy Name	Sequence Number	Status	Engine
FPolicy	cserver_policy	-	off	eng1
vs1.example.com	v1p1	-	off	eng2
vs1.example.com	v1p2	-	off	native
vs1.example.com	v1p3	-	off	native
vs1.example.com	cserver_policy	-	off	eng1
vs2.example.com	v1p1	3	on	native
vs2.example.com	v1p2	1	on	eng3
vs2.example.com	cserver_policy	2	on	eng1

Display information about enabled FPolicy policies

You can display information about enabled FPolicy policies to determine what FPolicy external engine it is configured to use, what the priority is for the policy, and to which storage virtual machine (SVM) the FPolicy policy is associated.

About this task

If you do not specify any parameters, the command displays the following information:

- SVM name
- Policy name
- Policy priority

You can use command parameters to filter the command's output by specified criteria.

Step

1. Display information about enabled FPolicy policies by using the appropriate command:

If you want to display information about enabled policies...	Enter the command...
On the cluster	<code>vserver fpolicy show-enabled</code>
On a specified SVM	<code>vserver fpolicy show-enabled -vserver vserver_name</code>
With the specified policy name	<code>vserver fpolicy show-enabled -policy-name policy_name</code>
With the specified sequence number	<code>vserver fpolicy show-enabled -priority integer</code>

Example

The following example displays the information about enabled FPolicy policies on the cluster:

```
cluster1::> vservers fpolicy show-enabled
Vserver                Policy Name                Priority
-----
vs1.example.com        pol_native                  native
vs1.example.com        pol_native2                 native
vs1.example.com        pol1                        2
vs1.example.com        pol2                        4
```

Manage FPolicy server connections

Connect to external FPolicy servers

To enable file processing, you might need to manually connect to an external FPolicy server if the connection has previously been terminated. A connection is terminated after the server timeout is reached or due to some error. Alternatively, the administrator might manually terminate a connection.

About this task

If a fatal error occurs, the connection to the FPolicy server can be terminated. After resolving the issue that caused the fatal error, you must manually reconnect to the FPolicy server.

Steps

1. Connect to the external FPolicy server by using the `vservers fpolicy engine-connect` command.

For more information about the command, see the man pages.

2. Verify that the external FPolicy server is connected by using the `vservers fpolicy show-engine` command.

For more information about the command, see the man pages.

Disconnect from external FPolicy servers

You might need to manually disconnect from an external FPolicy server. This might be desirable if the FPolicy server has issues with notification request processing or if you need to perform maintenance on the FPolicy server.

Steps

1. Disconnect from the external FPolicy server by using the `vservers fpolicy engine-disconnect` command.

For more information about the command, see the man pages.

2. Verify that the external FPolicy server is disconnected by using the `vservers fpolicy show-engine`

command.

For more information about the command, see the man pages.

Display information about connections to external FPolicy servers

You can display status information about connections to external FPolicy servers (FPolicy servers) for the cluster or for a specified storage virtual machine (SVM). This information can help you determine which FPolicy servers are connected.

About this task

If you do not specify any parameters, the command displays the following information:

- SVM name
- Node name
- FPolicy policy name
- FPolicy server IP address
- FPolicy server status
- FPolicy server type

In addition to displaying information about FPolicy connections on the cluster or a specific SVM, you can use command parameters to filter the command's output by other criteria.

You can specify the `-instance` parameter to display detailed information about listed policies. Alternatively, you can use the `-fields` parameter to display only the indicated fields in the command output. You can enter `?` after the `-fields` parameter to find out which fields you can use.

Step

1. Display filtered information about connection status between the node and the FPolicy server by using the appropriate command:

If you want to display connection status information about FPolicy servers...	Enter...
That you specify	<code>vserver fpolicy show-engine -server IP_address</code>
For a specified SVM	<code>vserver fpolicy show-engine -vserver vserver_name</code>
That are attached with a specified policy	<code>vserver fpolicy show-engine -policy-name policy_name</code>

<p>With the server status that you specify</p>	<pre>vserver fpolicy show-engine -server-status status</pre> <p>The server status can be one of the following:</p> <ul style="list-style-type: none"> • connected • disconnected • connecting • disconnecting
<p>With the specified type</p>	<pre>vserver fpolicy show-engine -server-type type</pre> <p>The FPolicy server type can be one of the following:</p> <ul style="list-style-type: none"> • primary • secondary
<p>That were disconnected with the specified reason</p>	<pre>vserver fpolicy show-engine -disconnect-reason text</pre> <p>Disconnect can be due to multiple reasons. The following are common reasons for disconnect:</p> <ul style="list-style-type: none"> • Disconnect command received from CLI. • Error encountered while parsing notification response from FPolicy server. • FPolicy Handshake failed. • SSL handshake failed. • TCP Connection to FPolicy server failed. • The screen response message received from the FPolicy server is not valid.

Example

This example displays information about external engine connections to FPolicy servers on SVM vs1.example.com:

```
cluster1::> vserver fpolicy show-engine -vserver vs1.example.com
FPolicy
Vserver          Policy      Node      Server      Server-      Server-
-----          -
vs1.example.com policy1    node1     10.1.1.2    connected    primary
vs1.example.com policy1    node1     10.1.1.3    disconnected  primary
vs1.example.com policy1    node2     10.1.1.2    connected    primary
vs1.example.com policy1    node2     10.1.1.3    disconnected  primary
```

This example displays information only about connected FPolicy servers:

```
cluster1::> vserver fpolicy show-engine -fields server -server-status
connected
node          vserver          policy-name      server
-----
node1         vs1.example.com policy1           10.1.1.2
node2         vs1.example.com policy1           10.1.1.2
```

Display information about the FPolicy passthrough-read connection status

You can display information about FPolicy passthrough-read connection status to external FPolicy servers (FPolicy servers) for the cluster or for a specified storage virtual machine (SVM). This information can help you determine which FPolicy servers have passthrough-read data connections and for which FPolicy servers the passthrough-read connection is disconnected.

About this task

If you do not specify any parameter, the command displays the following information:

- SVM name
- FPolicy policy name
- Node name
- FPolicy server IP address
- FPolicy passthrough-read connection status

In addition to displaying information about FPolicy connections on the cluster or a specific SVM, you can use command parameters to filter the command's output by other criteria.

You can specify the `-instance` parameter to display detailed information about listed policies. Alternatively, you can use the `-fields` parameter to display only the indicated fields in the command output. You can enter `?` after the `-fields` parameter to find out which fields you can use.

Step

1. Display filtered information about connection status between the node and the FPolicy server by using the

appropriate command:

If you want to display connection status information about...	Enter the command...
FPolicy passthrough-read connection status for the cluster	<code>vserver fpolicy show-passthrough-read-connection</code>
FPolicy passthrough-read connection status for a specified SVM	<code>vserver fpolicy show-passthrough-read-connection -vserver vserver_name</code>
FPolicy passthrough-read connection status for a specified policy	<code>vserver fpolicy show-passthrough-read-connection -policy-name policy_name</code>
Detailed FPolicy passthrough-read connection status for a specified policy	<code>vserver fpolicy show-passthrough-read-connection -policy-name policy_name -instance</code>
FPolicy passthrough-read connection status for the status that you specify	<code>vserver fpolicy show-passthrough-read-connection -policy-name policy_name -server-status status</code> The server status can be one of the following: <ul style="list-style-type: none">• connected• disconnected

Example

The following command displays information about passthrough-read connections from all FPolicy servers on the cluster:

```
cluster1::> vserver fpolicy show-passthrough-read-connection
```

Vserver	Policy Name	Node	FPolicy Server	Server Status
vs2.example.com	pol_cifs_2	FPolicy-01	2.2.2.2	disconnected
vs1.example.com	pol_cifs_1	FPolicy-01	1.1.1.1	connected

The following command displays detailed information about passthrough-read connections from FPolicy servers configured in the “pol_cifs_1” policy:

```
cluster1::> vserver fpolicy show-passthrough-read-connection -policy-name
pol_cifs_1 -instance

Node: FPolicy-01
Vserver: vs1.example.com
Policy: pol_cifs_1
Server: 1.1.1.1
Session ID of the Control Channel: 8cef052e-2502-11e3-
88d4-123478563412
Server Status: connected
Time Passthrough Read Channel was Connected: 9/24/2013 10:17:45
Time Passthrough Read Channel was Disconnected: -
Reason for Passthrough Read Channel Disconnection: none
```

Verify access using security tracing

How security traces work

You can add permission tracing filters to instruct ONTAP to log information about why the SMB and NFS servers on a storage virtual machine (SVM) allows or denies a client or user's request to perform an operation. This can be useful when you want to verify that your file access security scheme is appropriate or when you want to troubleshoot file access issues.

Security traces allow you to configure a filter that detects client operations over SMB and NFS on the SVM, and trace all access checks matching that filter. You can then view the trace results, which provides a convenient summary of the reason that access was allowed or denied.

When you want to verify the security settings for SMB or NFS access on files and folders on your SVM or if you are faced with an access problem, you can quickly add a filter to turn on permission tracing.

The following list outlines important facts about how security traces works:

- ONTAP applies security traces at the SVM level.
- Each incoming request is screened to see if it matches filtering criteria of any enabled security traces.
- Traces are performed for both file and folder access requests.
- Traces can filter based on the following criteria:
 - Client IP
 - SMB or NFS path
 - Windows name
 - UNIX name
- Requests are screened for *Allowed* and *Denied* access response results.
- Each request matching filtering criteria of enabled traces is recorded in the trace results log.
- The storage administrator can configure a timeout on a filter to automatically disable it.

- If a request matches multiple filters, the results from the filter with the highest index number is recorded.
- The storage administrator can print results from the trace results log to determine why an access request was allowed or denied.

Types of access checks security traces monitor

Access checks for a file or folder are done based on multiple criteria. Security traces monitor operations on all these criteria.

The types of access checks that security traces monitor include the following:

- Volume and qtree security style
- Effective security of the file system containing the files and folders on which operations are requested
- User mapping
- Share-level permissions
- Export-level permissions
- File-level permissions
- Storage-Level Access Guard security

Considerations when creating security traces

You should keep several considerations in mind when you create security traces on storage virtual machines (SVMs). For example, you need to know on which protocols you can create a trace, which security-styles are supported, and what the maximum number of active traces is.

- You can only create security traces on SVMs.
- Each security trace filter entry is SVM specific.

You must specify the SVM on which you want to run the trace.

- You can add permission tracing filters for SMB and NFS requests.
- You must set up the SMB or NFS server on the SVM on which you want to create trace filters.
- You can create security traces for files and folders residing on NTFS, UNIX, and mixed security-style volumes and qtrees.
- You can add a maximum of 10 permission tracing filters per SVM.
- You must specify a filter index number when creating or modifying a filter.

Filters are considered in order of the index number. The criteria in a filter with a higher index number is considered before the criteria with a lower index number. If the request being traced matches criteria in multiple enabled filters, only the filter with the highest index number is triggered.

- After you have created and enabled a security trace filter, you must perform some file or folder requests on a client system to generate activity that the trace filter can capture and log in the trace results log.
- You should add permission tracing filters for file access verification or troubleshooting purposes only.

Adding permission tracing filters has a minor effect on controller performance.

When you are done with verification or troubleshooting activity, you should disable or remove all permission tracing filters. Furthermore, the filtering criteria you select should be as specific as possible so that ONTAP does not send a large number of trace results to the log.

Perform security traces

Perform security traces overview

Performing a security trace involves creating a security trace filter, verifying the filter criteria, generating access requests on an SMB or NFS client that match filter criteria, and viewing the results.

After you are finished using a security filter to capture trace information, you can modify the filter and reuse it, or disable it if you no longer need it. After viewing and analyzing the filter trace results, you can then delete them if they are no longer needed.

Create security trace filters

You can create security trace filters that detect SMB and NFS client operations on storage virtual machines (SVMs) and trace all access checks matching the filter. You can use the results from security traces to validate your configuration or to troubleshoot access issues.


About this task

There are two required parameters for the vserver security trace filter create command:

Required parameters	Description
<code>-vserver vserver_name</code>	<i>SVM name</i> The name of the SVM that contains the files or folders on which you want to apply the security trace filter.
<code>-index index_number</code>	<i>Filter index number</i> The index number you want to apply to the filter. You are limited to a maximum of 10 trace filters per SVM. The allowed values for this parameter are 1 through 10.

A number of optional filter parameters enable you to customize the security trace filter so that you can narrow down the results produced by the security trace:

Filter parameter	Description
<code>-client-ip IP_Address</code>	This filter specifies the IP address from which the user is accessing the SVM.

<p><code>-path path</code></p>	<p>This filter specifies the path on which to apply the permission trace filter. The value for <code>-path</code> can use either of the following formats:</p> <ul style="list-style-type: none"> • The complete path, starting from the root of the share or export • A partial path, relative to the root of the share <p>You must use NFS style directory UNIX-style directory separators in the path value.</p>
<p><code>-windows-name win_user_name</code> or <code>-unix</code> <code>-name ``unix_user_name</code></p>	<p>You can specify either the Windows user name or UNIX user name whose access requests you want to trace. The user name variable is case insensitive. You cannot specify both a Windows user name and a UNIX user name in the same filter.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;">  <p>Even though you can trace SMB and NFS access events, the mapped UNIX user and the mapped UNIX users' groups might be used when performing access checks on mixed or UNIX security-style data.</p> </div>
<p><code>-trace-allow {yes no}</code></p>	<p>Tracing for deny events is always enabled for a security trace filter. You can optionally trace allow events. To trace allow events, you set this parameter to <code>yes</code>.</p>
<p><code>-enabled {enabled disabled}</code></p>	<p>You can enable or disable the security trace filter. By default, the security trace filter is enabled.</p>
<p><code>-time-enabled integer</code></p>	<p>You can specify a timeout for the filter, after which it is disabled.</p>

Steps

1. Create a security trace filter:

```
vserver security trace filter create -vserver vserver_name -index
index_numberfilter_parameters
```

`filter_parameters` is a list of optional filter parameters.

For more information, see the man pages for the command.

2. Verify the security trace filter entry:

```
vserver security trace filter show -vserver vserver_name -index index_number
```

Examples

The following command creates a security trace filter for any user accessing a file with a share path `\\server\share1\dir1\dir2\file.txt` from the IP address 10.10.10.7. The filter uses a complete path for the `-path` option. The client's IP address used to access data is 10.10.10.7. The filter times out after 30 minutes:

```

cluster1::> vserver security trace filter create -vserver vs1 -index 1
-path /dir1/dir2/file.txt -time-enabled 30 -client-ip 10.10.10.7
cluster1::> vserver security trace filter show -index 1
Vserver  Index  Client-IP          Path                Trace-Allow
Windows-Name
-----  -
vs1      1      10.10.10.7       /dir1/dir2/file.txt      no      -

```

The following command creates a security trace filter using a relative path for the `-path` option. The filter traces access for a Windows user named “joe”. Joe is accessing a file with a share path `\\server\share1\dir1\dir2\file.txt`. The filter traces allow and deny events:

```

cluster1::> vserver security trace filter create -vserver vs1 -index 2
-path /dir1/dir2/file.txt -trace-allow yes -windows-name mydomain\joe

cluster1::> vserver security trace filter show -vserver vs1 -index 2
Vserver: vs1
Filter Index: 2
Client IP Address to Match: -
Path: /dir1/dir2/file.txt
Windows User Name: mydomain\joe
UNIX User Name: -
Trace Allow Events: yes
Filter Enabled: enabled
Minutes Filter is Enabled: 60

```

Display information about security trace filters

You can display information about security trace filters configured on your storage virtual machine (SVM). This enables you to see which types of access events each filter traces.

Step

1. Display information about security trace filter entries by using the `vserver security trace filter show` command.

For more information about using this command, see the man pages.

Examples

The following command displays information about all security trace filters on SVM vs1:

```

cluster1::> vserver security trace filter show -vserver vs1
Vserver  Index  Client-IP          Path                Trace-Allow
Windows-Name
-----  -
vs1      1      -                  /dir1/dir2/file.txt  yes      -
vs1      2      -                  /dir3/dir4/          no
mydomain\joe

```

Display security trace results

You can display the security trace results generated for file operations that match security trace filters. You can use the results to validate your file access security configuration or to troubleshoot SMB and NFS file access issues.

What you'll need

An enabled security trace filter must exist and operations must have been performed from an SMB or NFS client that matches the security trace filter to generate security trace results.

About this task

You can display a summary of all security trace results, or you can customize what information is displayed in the output by specifying optional parameters. This can be helpful when the security trace results contain a large number of records.

If you do not specify any of the optional parameters, the following is displayed:

- storage virtual machine (SVM) name
- Node name
- Security trace index number
- Security style
- Path
- Reason
- User name

The user name is displayed depending on how the trace filter is configured:

If the filter is configured...	Then...
With a UNIX user name	The security trace result displays the UNIX user name.
With a Windows user name	The security trace result displays the Windows user name.
Without a user name	The security trace result displays the Windows user name.

You can customize the output by using optional parameters. Some of the optional parameters that you can use

to narrow the results returned in the command output include the following:

Optional parameter	Description
<code>-fields field_name, ...</code>	Displays output on the fields you choose. You can use this parameter either alone or in combination with other optional parameters.
<code>-instance</code>	Displays detailed information about security trace events. Use this parameter with other optional parameters to display detailed information about specific filter results.
<code>-node node_name</code>	Displays information only about events on the specified node.
<code>-vserver vserver_name</code>	Displays information only about events on the specified SVM.
<code>-index integer</code>	Displays information about the events that occurred as a result of the filter corresponding to the specified index number.
<code>-client-ip IP_address</code>	Displays information about the events that occurred as a result of file access from the specified client IP address.
<code>-path path</code>	Displays information about the events that occurred as a result of file access to the specified path.
<code>-user-name user_name</code>	Displays information about the events that occurred as a result of file access by the specified Windows or UNIX user.
<code>-security-style security_style</code>	Displays information about the events that occurred on file systems with the specified security style.

See the man page for information about other optional parameters that you can use with the command.

Step

1. Display security trace filter results by using the `vserver security trace trace-result show` command.

```
vserver security trace trace-result show -user-name domain\user
```

```
Vserver: vs1
```

Node	Index	Filter Details	Reason
node1	3	User:domain\user Security Style:mixed Path:/dir1/dir2/	Access denied by explicit ACE
node1	5	User:domain\user Security Style:unix Path:/dir1/	Access denied by explicit ACE

Modify security trace filters

If you want to change the optional filter parameters used to determine which access events are traced, you can modify existing security trace filters.

About this task

You must identify which security trace filter you want to modify by specifying the storage virtual machine (SVM) name on which the filter is applied and the index number of the filter. You can modify all the optional filter parameters.

Steps

1. Modify a security trace filter:

```
vserver security trace filter modify -vserver vserver_name -index  
index_numberfilter_parameters
```

- `vserver_name` is the name of the SVM on which you want to apply a security trace filter.
- `index_number` is the index number that you want to apply to the filter. The allowed values for this parameter are 1 through 10.
- `filter_parameters` is a list of optional filter parameters.

2. Verify the security trace filter entry:

```
vserver security trace filter show -vserver vserver_name -index index_number
```

Example

The following command modifies the security trace filter with the index number 1. The filter traces events for any user accessing a file with a share path `\\server\share1\dir1\dir2\file.txt` from any IP address. The filter uses a complete path for the `-path` option. The filter traces allow and deny events:

```

cluster1::> vserver security trace filter modify -vserver vs1 -index 1
-path /dir1/dir2/file.txt -trace-allow yes

cluster1::> vserver security trace filter show -vserver vs1 -index 1
          Vserver: vs1
          Filter Index: 1
Client IP Address to Match: -
          Path: /dir1/dir2/file.txt
Windows User Name: -
UNIX User Name: -
Trace Allow Events: yes
Filter Enabled: enabled
Minutes Filter is Enabled: 60

```

Delete security trace filters

When you no longer need a security trace filter entry, you can delete it. Because you can have a maximum of 10 security trace filters per storage virtual machine (SVM), deleting unneeded filters enables you to create new filters if you have reached the maximum.

About this task

To uniquely identify the security trace filter that you want to delete, you must specify the following:

- The name of the SVM to which the trace filter is applied
- The filter index number of the trace filter

Steps

1. Identify the filter index number of the security trace filter entry you want to delete:

```

vserver security trace filter show -vserver vserver_name

vserver security trace filter show -vserver vs1

```

Vserver	Index	Client-IP	Path	Trace-Allow	Windows-Name
vs1	1	-	/dir1/dir2/file.txt	yes	-
vs1	2	-	/dir3/dir4/	no	
mydomain\joe					

2. Using the filter index number information from the previous step, delete the filter entry:

```

vserver security trace filter delete -vserver vserver_name -index index_number

vserver security trace filter delete -vserver vs1 -index 1

```

3. Verify that the security trace filter entry is deleted:

```
vserver security trace filter show -vserver vserver_name
```

```
vserver security trace filter show -vserver vs1
```

Vserver	Index	Client-IP	Path	Trace-Allow
Windows-Name				
-----	-----	-----	-----	-----
vs1	2	-	/dir3/dir4/	no
mydomain\joe				

Delete security trace records

After you finish using a filter trace record to verify file access security or to troubleshoot SMB or NFS client access issues, you can delete the security trace record from the security trace log.

About this task

Before you can delete a security trace record, you must know the record's sequence number.



Each storage virtual machine (SVM) can store a maximum of 128 trace records. If the maximum is reached on the SVM, the oldest trace records are automatically deleted as new ones are added. If you do not want to manually delete trace records on this SVM, you can let ONTAP automatically delete the oldest trace results after the maximum is reached to make room for new results.

Steps

1. Identify the sequence number of the record you want to delete:

```
vserver security trace trace-result show -vserver vserver_name -instance
```

2. Delete the security trace record:

```
vserver security trace trace-result delete -node node_name -vserver  
vserver_name -seqnum integer
```

```
vserver security trace trace-result delete -vserver vs1 -node node1 -seqnum  
999
```

- `-node node_name` is the name of the cluster node on which the permission tracing event that you want to delete occurred.

This is a required parameter.

- `-vserver vserver_name` is the name of the SVM on which the permission tracing event that you want to delete occurred.

This is a required parameter.

- `-seqnum integer` is the sequence number of the log event that you want to delete.

This is a required parameter.

Delete all security trace records

If you do not want to keep any of the existing security trace records, you can delete all of the records on a node with a single command.

Step

1. Delete all security trace records:

```
vserver security trace trace-result delete -node node_name -vserver
vserver_name *
```

- `-node node_name` is the name of the cluster node on which the permission tracing event that you want to delete occurred.
- `-vserver vserver_name` is the name of the storage virtual machine (SVM) on which the permission tracing event that you want to delete occurred.

Interpret security trace results

Security trace results provide the reason that a request was allowed or denied. Output displays the result as a combination of the reason for allowing or denying access and the location within the access checking pathway where access is either allowed or denied. You can use the results to isolate and identify why actions are or are not allowed.

Finding information about the lists of result types and filter details

You can find the lists of result types and filter details that can be included in the security trace results in the man pages for the `vserver security trace trace-result show` command.

Example of output from the `Reason` field in an `Allow` result type

The following is an example of the output from the `Reason` field that appears in the trace results log in an `Allow` result type:

```
Access is allowed because SMB implicit permission grants requested
access while opening existing file or directory.
```

```
Access is allowed because NFS implicit permission grants requested
access while opening existing file or directory.
```

Example of output from the `Reason` field in an `Allow` result type

The following is an example of the output from the `Reason` field that appears in the trace results log in a `Deny`

result type:

```
Access is denied. The requested permissions are not granted by the
ACE while checking for child-delete access on the parent.
```

Example of output from the `Filter details` field

The following is an example of the output from the `Filter details` field in the trace results log, which list the effective security style of the file system containing files and folders that match the filter criteria:

```
Security Style: MIXED and ACL
```

Where to find additional information

After you have successfully tested SMB client access, you can perform advanced SMB configuration or add SAN access. After you have successfully tested NFS client access, you can perform advanced NFS configuration or add SAN access. When protocol access is complete, you should protect the root volume of SVM.

SMB configuration

You can further configure SMB access using the following:

- [SMB management](#)

Describes how to configure and manage file access using the SMB protocol.

- [NetApp Technical Report 4191: Best Practices Guide for Clustered Data ONTAP 8.2 Windows File Services](#)

Provides a brief overview of SMB implementation and other Windows File Services features with recommendations and basic troubleshooting information for ONTAP.

- [NetApp Technical Report 3740: SMB 2 Next-Generation CIFS Protocol in Data ONTAP](#)

Describes SMB 2 features, configuration details, and its implementation in ONTAP.

NFS configuration

You can further configure NFS access using the following:

- [NFS management](#)

Describes how to configure and manage file access using the NFS protocol.

- [NetApp Technical Report 4067: NFS Best Practice and Implementation Guide](#)

Serves as an NFSv3 and NFSv4 operational guide and provides an overview of ONTAP operating system with a focus on NFSv4.

- [NetApp Technical Report 4668: Name Services Best Practices Guide](#)

Provides a comprehensive list of best practices, limits, recommendations, and considerations when configuring LDAP, NIS, DNS, and local user and group files for authentication purposes.

- [NetApp Technical Report 4616: NFS Kerberos in ONTAP with Microsoft Active Directory](#)
- [NetApp Technical Report 4835: How to Configure LDAP in ONTAP](#)
- [NetApp Technical Report 3580: NFSv4 Enhancements and Best Practices Guide Data ONTAP Implementation](#)

Describes the best practices that should be followed while implementing NFSv4 components on AIX, Linux, or Solaris clients attached to systems running ONTAP.

Root volume protection

After configuring protocols on the SVM, you should ensure that its root volume is protected:

- [Data protection](#)

Describes how to create a load-sharing mirror to protect the SVM root volume, which is a NetApp best practice for NAS-enabled SVMs. Also describes how to quickly recover from volume failures or losses by promoting the SVM root volume from a load-sharing mirror.

Manage encryption with System Manager



Encrypt stored data using software-based encryption

Use volume encryption to ensure that volume data cannot be read if the underlying device is repurposed, returned, misplaced, or stolen. Volume encryption does not require special disks; it works with all HDDs and SSDs.

Volume encryption requires a key manager. You can configure the Onboard Key Manager using System Manager. You can also use an external key manager, but you need to first set it up using the ONTAP CLI.

After the key manager is configured, new volumes are encrypted by default.

Steps

1. Click **Cluster > Settings**.
2. Under **Encryption**, click  to configure the Onboard Key Manager for the first time.
3. To encrypt existing volumes, click **Storage > Volumes**.
4. On the desired volume, click  and then click **Edit**.
5. Select **Enable encryption**.



Encrypt stored data using self-encrypting drives

Use disk encryption to ensure that all data in a local tier cannot be read if the underlying device is repurposed, returned, misplaced, or stolen. Disk encryption requires special self-encrypting HDDs or SSDs.

Disk encryption requires a key manager. You can configure the onboard key manager using System Manager. You can also use an external key manager, but you need to first set it up using the ONTAP CLI.

If ONTAP detects self-encrypting disks, it prompts you to configure the onboard key manager when you create the local tier.

Steps

1. Under **Encryption**, click  to configure the onboard key manager.
2. If you see a message that disks need to be rekeyed, click , and then click **Rekey Disks**.

Manage encryption with the CLI

NetApp Encryption overview

NetApp offers both software- and hardware-based encryption technologies for ensuring that data at rest cannot be read if the storage medium is repurposed, returned, misplaced, or stolen.

- Software-based encryption using NetApp Volume Encryption (NVE) supports data encryption one volume at a time
- Hardware-based encryption using NetApp Storage Encryption (NSE) supports full-disk encryption (FDE) of data as it is written.

Configure NetApp Volume Encryption

Configure NetApp Volume Encryption overview

NetApp Volume Encryption (NVE) is a software-based technology for encrypting data at rest one volume at a time. An encryption key accessible only to the storage system ensures that volume data cannot be read if the underlying device is repurposed, returned, misplaced, or stolen.

Understanding NVE

With NVE, both metadata and data (including Snapshot copies) are encrypted. Access to the data is given by a unique XTS-AES-256 key, one per volume. An external key management server or Onboard Key Manager (OKM) serves keys to nodes:

- The external key management server is a third-party system in your storage environment that serves keys to nodes using the Key Management Interoperability Protocol (KMIP). It is a best practice to configure external key management servers on a different storage system from your data.
- The Onboard Key Manager is a built-in tool that serves keys to nodes from the same storage system as your data.

Beginning with ONTAP 9.7, aggregate and volume encryption is enabled by default if you have a volume encryption (VE) license and use an onboard or external key manager. The VE license is included with [ONTAP One](#). Whenever an external or onboard key manager is configured there is a change in how the encryption of data at rest is configured for brand new aggregates and brand new volumes. Brand new aggregates will have NetApp Aggregate Encryption (NAE) enabled by default. Brand new volumes that are not part of an NAE aggregate will have NetApp Volume Encryption (NVE) enabled by default. If a data storage virtual machine

(SVM) is configured with its own key-manager using multi-tenant key management, then the volume created for that SVM is automatically configured with NVE.

You can enable encryption on a new or existing volume. NVE supports the full range of storage efficiency features, including deduplication and compression. Beginning with ONTAP 9.14.1, you can [enable NVE on existing SVM root volumes](#).



If you are using SnapLock, you can enable encryption only on new, empty SnapLock volumes. You cannot enable encryption on an existing SnapLock volume.

You can use NVE on any type of aggregate (HDD, SSD, hybrid, array LUN), with any RAID type, and in any supported ONTAP implementation, including ONTAP Select. You can also use NVE with hardware-based encryption to “double encrypt” data on self-encrypting drives.

When NVE is enabled, the core dump is also encrypted.

Aggregate-level encryption

Ordinarily, every encrypted volume is assigned a unique key. When the volume is deleted, the key is deleted with it.

Beginning with ONTAP 9.6, you can use *NetApp Aggregate Encryption (NAE)* to assign keys to the containing aggregate for the volumes to be encrypted. When an encrypted volume is deleted, the keys for the aggregate are preserved. The keys are deleted if the entire aggregate is deleted.

You must use aggregate-level encryption if you plan to perform inline or background aggregate-level deduplication. Aggregate-level deduplication is otherwise not supported by NVE.

Beginning with ONTAP 9.7, aggregate and volume encryption is enabled by default if you have a volume encryption (VE) license and use an onboard or external key manager.

NVE and NAE volumes can coexist on the same aggregate. Volumes encrypted under aggregate-level encryption are NAE volumes by default. You can override the default when you encrypt the volume.

You can use the `volume move` command to convert an NVE volume to an NAE volume, and vice versa. You can replicate an NAE volume to an NVE volume.

You cannot use `secure purge` commands on an NAE volume.

When to use external key management servers

Although it is less expensive and typically more convenient to use the onboard key manager, you should set up KMIP servers if any of the following are true:

- Your encryption key management solution must comply with Federal Information Processing Standards (FIPS) 140-2 or the OASIS KMIP standard.
- You need a multi-cluster solution, with centralized management of encryption keys.
- Your business requires the added security of storing authentication keys on a system or in a location different from the data.

Scope of external key management

The scope of external key management determines whether key management servers secure all the SVMs in the cluster or selected SVMs only:

- You can use a *cluster scope* to configure external key management for all the SVMs in the cluster. The cluster administrator has access to every key stored on the servers.
- Beginning with ONTAP 9.6, you can use an *SVM scope* to configure external key management for a named SVM in the cluster. That's best for multitenant environments in which each tenant uses a different SVM (or set of SVMs) to serve data. Only the SVM administrator for a given tenant has access to the keys for that tenant.
- Beginning with ONTAP 9.10.1, you can use [Azure Key Vault and Google Cloud KMS](#) to protect NVE keys only for data SVMs. This is available for AWS's KMS beginning in 9.12.0.

You can use both scopes in the same cluster. If key management servers have been configured for an SVM, ONTAP uses only those servers to secure keys. Otherwise, ONTAP secures keys with the key management servers configured for the cluster.

A list of validated external key managers is available in the [NetApp Interoperability Matrix Tool \(IMT\)](#). You can find this list by entering the term "key managers" into the IMT's search feature.

Support details

The following table shows NVE support details:

Resource or feature	Support details
Platforms	AES-NI offload capability required. See the Hardware Universe (HWU) to verify that NVE and NAE are supported for your platform.
Encryption	<p>Beginning with ONTAP 9.7, newly created aggregates and volumes are encrypted by default when you add a volume encryption (VE) license and have an onboard or external key manager configured. If you need to create an unencrypted aggregate, use the following command:</p> <pre>storage aggregate create -encrypt-with-aggr-key false</pre> <p>If you need to create a plain text volume, use the following command:</p> <pre>volume create -encrypt false</pre> <p>Encryption is not enabled by default when:</p> <ul style="list-style-type: none"> • VE license is not installed. • Key manager is not configured. • Platform or software does not support encryption. • Hardware encryption is enabled.
ONTAP	All ONTAP implementations. Support for ONTAP Cloud is available in ONTAP 9.5 and later.
Devices	HDD, SSD, hybrid, array LUN.
RAID	RAID0, RAID4, RAID-DP, RAID-TEC.


Volumes	Data volumes and existing SVM root volumes. You cannot encrypt data on MetroCluster metadata volumes. In versions of ONTAP earlier than 9.14.1, you cannot encrypt data on the SVM root volume with NVE. Beginning with ONTAP 9.14.1, ONTAP supports NVE on SVM root volumes .
Aggregate-level encryption	Beginning with ONTAP 9.6, NVE supports aggregate-level encryption (NAE): <ul style="list-style-type: none"> • You must use aggregate-level encryption if you plan to perform inline or background aggregate-level deduplication. • You cannot rekey an aggregate-level encryption volume. • Secure-purge is not supported on aggregate-level encryption volumes. • In addition to data volumes, NAE supports encryption of SVM root volumes and the MetroCluster metadata volume. NAE does not support encryption of the root volume.
SVM scope	Beginning with ONTAP 9.6, NVE supports SVM scope for external key management only, not for Onboard Key Manager. MetroCluster is supported beginning with ONTAP 9.8.
Storage efficiency	Deduplication, compression, compaction, FlexClone. Clones use the same key as the parent, even after splitting the clone from the parent. You should perform a <code>volume move</code> on a split clone, after which the split clone will have a different key.
Replication	<ul style="list-style-type: none"> • For volume replication, the source and destination volumes can have different encryption settings. Encryption can be configured for the source and unconfigured for the destination, and vice versa. • For SVM replication, the destination volume is automatically encrypted, unless the destination does not contain a node that supports volume encryption, in which case replication succeeds, but the destination volume is not encrypted. • For MetroCluster configurations, each cluster pulls external key management keys from its configured key servers. OKM keys are replicated to the partner site by the configuration replication service.
Compliance	Beginning with ONTAP 9.2, SnapLock is supported in both Compliance and Enterprise modes, for new volumes only. You cannot enable encryption on an existing SnapLock volume.
FlexGroups	Beginning with ONTAP 9.2, FlexGroups are supported. Destination aggregates must be of the same type as source aggregates, either volume-level or aggregate-level. Beginning with ONTAP 9.5, in-place rekey of FlexGroup volumes is supported.
7-Mode transition	Beginning with 7-Mode Transition Tool 3.3, you can use the 7-Mode Transition Tool CLI to perform copy-based transition to NVE-enabled destination volumes on the clustered system.

Related information

[FAQ - NetApp Volume Encryption and NetApp Aggregate Encryption](#)

NetApp Volume Encryption workflow

You must configure key management services before you can enable volume encryption. You can enable encryption on a new volume or on an existing volume.



You must [install the VE license](#) and configure key management services before you can encrypt data with NVE. Before installing the license, you should [determine whether your ONTAP version supports NVE](#).

Configure NVE

Determine whether your cluster version supports NVE

You should determine whether your cluster version supports NVE before you install the license. You can use the `version` command to determine the cluster version.

About this task

The cluster version is the lowest version of ONTAP running on any node in the cluster.

Step

1. Determine whether your cluster version supports NVE:

```
version -v
```

NVE is not supported if the command output displays the text “1Ono-DARE” (for “no Data At Rest Encryption”), or if you are using a platform that is not listed in [Support details](#).

The following command determines whether NVE is supported on `cluster1`.

```
cluster1::> version -v
NetApp Release 9.1.0: Tue May 10 19:30:23 UTC 2016 <1Ono-DARE>
```

The output of `1Ono-DARE` indicates that NVE is not supported on your cluster version.

Install the license

A VE license entitles you to use the feature on all nodes in the cluster. This license is required before you can encrypt data with NVE. It is included with [ONTAP One](#).

Prior to ONTAP One, the VE license was included with the Encryption bundle. The Encryption bundle is no longer offered, but is still valid. Although not currently required, existing customers can choose to [upgrade to ONTAP One](#).

Before you begin

- You must be a cluster administrator to perform this task.
- You must have received the VE license key from your sales representative or have ONTAP One installed.

Steps

1. [Verify that the VE license is installed](#).

The VE license package name is `VE`.

2. If the license is not installed, [use System Manager or the ONTAP CLI to install it](#).

Configure external key management

Configure external key management overview

You can use one or more external key management servers to secure the keys that the cluster uses to access encrypted data. An external key management server is a third-party system in your storage environment that serves keys to nodes using the Key Management Interoperability Protocol (KMIP).



For ONTAP 9.1 and earlier versions, node management LIFs must be assigned to ports that are configured with the node management role before you can use the external key manager.

NetApp Volume Encryption (NVE) supports Onboard Key Manager in ONTAP 9.1 and later. Beginning in ONTAP 9.3, NVE supports external key management (KMIP) and Onboard Key Manager. Beginning in ONTAP 9.10.1, you can use [Azure Key Vault](#) or [Google Cloud Key Manager Service](#) to protect your NVE keys.

Beginning in ONTAP 9.11.1, you can configure multiple external key managers in a cluster. See [Configure clustered key servers](#).

Manage external key managers with System Manager

Beginning with ONTAP 9.7, you can store and manage authentication and encryption keys with the Onboard Key Manager. Beginning with ONTAP 9.13.1, you can also use external key managers to store and manage these keys.

The Onboard Key Manager stores and manages keys in a secure database that is internal to the cluster. Its scope is the cluster. An external key manager stores and manages keys outside the cluster. Its scope can be the cluster or the storage VM. One or more external key managers can be used. The following conditions apply:

- If the Onboard Key Manager is enabled, an external key manager cannot be enabled at the cluster level, but it can be enabled at the storage VM level.
- If an external key manager is enabled at the cluster level, the Onboard Key Manager cannot be enabled.

When using external key managers, you can register up to four primary key servers per storage VM and cluster. Each primary key server can be clustered with up to three secondary key servers.



Configure an external key manager

To add an external key manager for a storage VM, you should add an optional gateway when you configure the network interface for the storage VM. If the storage VM was created without the network route, you will have to create the route explicitly for the external key manager. See [Create a LIF \(network interface\)](#).

Steps



You can configure an external key manager starting from different locations in System Manager.

1. To configure an external key manager, perform one of the following starting steps.

Workflow	Navigation	Starting step
Configure Key Manager	Cluster > Settings	Scroll to the Security section. Under Encryption , select  . Select External Key Manager .
Add local tier	Storage > Tiers	Select + Add Local Tier . Check the check box labeled "Configure Key Manager". Select External Key Manager .
Prepare storage	Dashboard	In the Capacity section, select Prepare Storage . Then, select "Configure Key Manager". Select External Key Manager .
Configure encryption (key manager at storage VM scope only)	Storage > Storage VMs	Select the storage VM. Select the Settings tab. In the Encryption section under Security , select  .

2. To add a primary key server, select **+ Add**, and complete the **IP Address or Host Name** and **Port** fields.
3. Existing installed certificates are listed in the **KMIP Server CA Certificates** and **KMIP Client Certificate**

fields. You can perform any of the following actions:



- Select  to select installed certificates that you want to map to the key manager. (Multiple service CA certificates can be selected, but only one client certificate can be selected.)
 - Select **Add New Certificate** to add a certificate that has not already been installed and map it to the external key manager.
 - Select  next to the certificate name to delete installed certificates that you do not want to map to the external key manager.
4. To add a secondary key server, select **Add** in the **Secondary Key Servers** column, and provide its details.
 5. Select **Save** to complete the configuration.



Edit an existing external key manager

If you have already configured an external key manager, you can modify its settings.

Steps

1. To edit the configuration of an external key manager, perform one of the following starting steps.

Scope	Navigation	Starting step
Cluster scope external key manager	Cluster > Settings	Scroll to the Security section. Under Encryption , select  , then select Edit External Key Manager .
Storage VM scope external key manager	Storage > Storage VMs	Select the storage VM. Select the Settings tab. In the Encryption section under Security , select  , then select Edit External Key Manager .



2. Existing key servers are listed in the **Key Servers** table. You can perform the following operations:
 - Add a new key server by selecting  **Add**.
 - Delete a key server by selecting  at the end of the table cell that contains the name of the key server. The secondary key servers associated with that primary key server are also removed from the configuration.

Delete an external key manager

An external key manager can be deleted if the volumes are unencrypted.

Steps

1. To delete an external key manager, perform one of the following steps.

Scope	Navigation	Starting step
Cluster scope external key manager	Cluster > Settings	Scroll to the Security section. Under Encryption , select  , then select Delete External Key Manager .
Storage VM scope external key manager	Storage > Storage VMs	Select the storage VM. Select the Settings tab. In the Encryption section under Security , select  , then select Delete External Key Manager .

Migrate keys among key managers

When multiple key managers are enabled on a cluster, keys must be migrated from one key manager to another. This process is completed automatically with System Manager.

- If the Onboard Key Manager or an external key manager is enabled at a cluster level, and some volumes are encrypted, then when you configure an external key manager at the storage VM level, the keys must be migrated from the Onboard Key Manager or external key manager at the cluster level to the external key manager at the storage VM level. This process is completed automatically by System Manager.
- If volumes were created without encryption on a storage VM, then keys do not need to be migrated.

Install SSL certificates on the cluster

The cluster and KMIP server use KMIP SSL certificates to verify each other's identity and establish an SSL connection. Before configuring the SSL connection with the KMIP server, you must install the KMIP client SSL certificates for the cluster, and the SSL public certificate for the root certificate authority (CA) of the KMIP server.

About this task

In an HA pair, both nodes must use the same public and private KMIP SSL certificates. If you connect multiple HA pairs to the same KMIP server, all nodes in the HA pairs must use the same public and private KMIP SSL certificates.

Before you begin

- The time must be synchronized on the server creating the certificates, the KMIP server, and the cluster.
- You must have obtained the public SSL KMIP client certificate for the cluster.
- You must have obtained the private key associated with the SSL KMIP client certificate for the cluster.
- The SSL KMIP client certificate must not be password-protected.
- You must have obtained the SSL public certificate for the root certificate authority (CA) of the KMIP server.
- In a MetroCluster environment, you must install the same KMIP SSL certificates on both clusters.



You can install the client and server certificates on the KMIP server before or after installing the certificates on the cluster.

Steps

1. Install the SSL KMIP client certificates for the cluster:

```
security certificate install -vserver admin_svm_name -type client
```

You are prompted to enter the SSL KMIP public and private certificates.

```
cluster1::> security certificate install -vserver cluster1 -type client
```

2. Install the SSL public certificate for the root certificate authority (CA) of the KMIP server:

```
security certificate install -vserver admin_svm_name -type server-ca
```

```
cluster1::> security certificate install -vserver cluster1 -type server-ca
```

Enable external key management in ONTAP 9.6 and later (NVE)

You can use one or more KMIP servers to secure the keys the cluster uses to access encrypted data. Beginning with ONTAP 9.6, you have the option to configure a separate external key manager to secure the keys that a data SVM uses to access encrypted data.

Beginning with ONTAP 9.11.1, you can add up to 3 secondary key servers per primary key server to create a clustered key server. For more information, see [Configure clustered external key servers](#).

About this task

You can connect up to four KMIP servers to a cluster or SVM. A minimum of two servers is recommended for redundancy and disaster recovery.

The scope of external key management determines whether key management servers secure all the SVMs in the cluster or selected SVMs only:

- You can use a *cluster scope* to configure external key management for all the SVMs in the cluster. The cluster administrator has access to every key stored on the servers.
- Beginning with ONTAP 9.6, you can use an *SVM scope* to configure external key management for a data SVM in the cluster. That's best for multitenant environments in which each tenant uses a different SVM (or set of SVMs) to serve data. Only the SVM administrator for a given tenant has access to the keys for that tenant.
- For multitenant environments, install a license for *MT_EK_MGMT* by using the following command:

```
system license add -license-code <MT_EK_MGMT license code>
```

For complete command syntax, see the man page for the command.

You can use both scopes in the same cluster. If key management servers have been configured for an SVM, ONTAP uses only those servers to secure keys. Otherwise, ONTAP secures keys with the key management servers configured for the cluster.

You can configure onboard key management at the cluster scope and external key management at the SVM scope. You can use the `security key-manager key migrate` command to migrate keys from onboard key management at the cluster scope to external key managers at the SVM scope.

Before you begin

- The KMIP SSL client and server certificates must have been installed.
- You must be a cluster or SVM administrator to perform this task.
- If you want to enable external key management for a MetroCluster environment, MetroCluster must be fully configured before enabling external key management.
- In a MetroCluster environment, you must install the same KMIP SSL certificate on both clusters.

Steps

1. Configure key manager connectivity for the cluster:

```
security key-manager external enable -vserver admin_SVM -key-servers  
host_name|IP_address:port,... -client-cert client_certificate -server-ca-cert  
server_CA_certificates
```



- The `security key-manager external enable` command replaces the `security key-manager setup` command. If you run the command at the cluster login prompt, `admin_SVM` defaults to the admin SVM of the current cluster. You must be the cluster administrator to configure cluster scope. You can run the `security key-manager external modify` command to change the external key management configuration.
- In a MetroCluster environment, if you are configuring external key management for the admin SVM, you must repeat the `security key-manager external enable` command on the partner cluster.

The following command enables external key management for `cluster1` with three external key servers. The first key server is specified using its hostname and port, the second is specified using an IP address and the default port, and the third is specified using an IPv6 address and port:

```
cluster1::> security key-manager external enable -vserver cluster1 -key
-servers
ks1.local:15696,10.0.0.10,[fd20:8b1e:b255:814e:32bd:f35c:832c:5a09]:1234
-client-cert AdminVserverClientCert -server-ca-certs
AdminVserverServerCaCert
```

2. Configure a key manager an SVM:

```
security key-manager external enable -vserver SVM -key-servers
host_name|IP_address:port,... -client-cert client_certificate -server-ca-cert
server_CA_certificates
```



- If you run the command at the SVM login prompt, `SVM` defaults to the current SVM. You must be a cluster or SVM administrator to configure SVM scope. You can run the `security key-manager external modify` command to change the external key management configuration.
- In a MetroCluster environment, if you are configuring external key management for a data SVM, you do not have to repeat the `security key-manager external enable` command on the partner cluster.

The following command enables external key management for `svm1` with a single key server listening on the default port 5696:

```
svm11::> security key-manager external enable -vserver svm1 -key-servers
keyserver.svm1.com -client-cert SVM1ClientCert -server-ca-certs
SVM1ServerCaCert
```

3. Repeat the last step for any additional SVMs.



You can also use the `security key-manager external add-servers` command to configure additional SVMs. The `security key-manager external add-servers` command replaces the `security key-manager add` command. For complete command syntax, see the man page.

4. Verify that all configured KMIP servers are connected:

```
security key-manager external show-status -node node_name
```



The `security key-manager external show-status` command replaces the `security key-manager show -status` command. For complete command syntax, see the man page.

```
cluster1::> security key-manager external show-status
```

Node	Vserver	Key Server	Status

node1			
	svm1		
		keyserver.svm1.com:5696	available
	cluster1		
		10.0.0.10:5696	available
		fd20:8b1e:b255:814e:32bd:f35c:832c:5a09:1234	available
		ks1.local:15696	available
node2			
	svm1		
		keyserver.svm1.com:5696	available
	cluster1		
		10.0.0.10:5696	available
		fd20:8b1e:b255:814e:32bd:f35c:832c:5a09:1234	available
		ks1.local:15696	available

8 entries were displayed.

5. Optionally, convert plain text volumes to encrypted volumes.

```
volume encryption conversion start
```

An external key manager must be fully configured before you convert the volumes. In a MetroCluster environment, an external key manager must be configured on both sites.

Enable external key management in ONTAP 9.5 and earlier

You can use one or more KMIP servers to secure the keys the cluster uses to access encrypted data. You can connect up to four KMIP servers to a node. A minimum of two servers is recommended for redundancy and disaster recovery.

About this task

ONTAP configures KMIP server connectivity for all nodes in the cluster.

Before you begin

- The KMIP SSL client and server certificates must have been installed.
- You must be a cluster administrator to perform this task.
- You must configure the MetroCluster environment before you configure an external key manager.
- In a MetroCluster environment, you must install the same KMIP SSL certificate on both clusters.

Steps

1. Configure key manager connectivity for cluster nodes:

```
security key-manager setup
```

The key manager setup starts.



In a MetroCluster environment, you must run this command on both clusters.

2. Enter the appropriate response at each prompt.
3. Add a KMIP server:

```
security key-manager add -address key_management_server_ipaddress
```

```
cluster1::> security key-manager add -address 20.1.1.1
```



In a MetroCluster environment, you must run this command on both clusters.

4. Add an additional KMIP server for redundancy:

```
security key-manager add -address key_management_server_ipaddress
```

```
cluster1::> security key-manager add -address 20.1.1.2
```



In a MetroCluster environment, you must run this command on both clusters.

5. Verify that all configured KMIP servers are connected:

```
security key-manager show -status
```

For complete command syntax, see the man page.


```
cluster1::> security key-manager show -status
```

Node	Port	Registered Key Manager	Status
cluster1-01	5696	20.1.1.1	available
cluster1-01	5696	20.1.1.2	available
cluster1-02	5696	20.1.1.1	available
cluster1-02	5696	20.1.1.2	available

6. Optionally, convert plain text volumes to encrypted volumes.

```
volume encryption conversion start
```

An external key manager must be fully configured before you convert the volumes. In a MetroCluster environment, an external key manager must be configured on both sites.

Manage keys with a cloud provider

Beginning in ONTAP 9.10.1, you can use [Azure Key Vault \(AKV\)](#) and [Google Cloud Platform's Key Management Service \(Cloud KMS\)](#) to protect your ONTAP encryption keys in a cloud-hosted application. Beginning with ONTAP 9.12.0, you can also protect NVE keys with [AWS' KMS](#).

AWS KMS, AKV and Cloud KMS can be used to protect [NetApp Volume Encryption \(NVE\) keys](#) only for data SVMs.

About this task

Key management with a cloud provider can be enabled with the CLI or the ONTAP REST API.

When using a cloud provider to protect your keys, be aware that by default a data SVM LIF is used to communicate with the cloud key management endpoint. A node management network is used to communicate with the cloud provider's authentication services (login.microsoftonline.com for Azure; oauth2.googleapis.com for Cloud KMS). If the cluster network is not configured correctly, the cluster will not properly utilize the key management service.

When utilizing a cloud provider key management service, you should be aware of the following limitations:

- Cloud-provider key management is not available for NetApp Storage Encryption (NSE) and NetApp Aggregate Encryption (NAE). [External KMIPs](#) can be used instead.
- Cloud-provider key management is not available for MetroCluster configurations.
- Cloud-provider key management can only be configured on a data SVM.

Before you begin

- You must have configured the KMS on the appropriate cloud provider.
- The ONTAP cluster's nodes must support NVE.
- [You must have installed the Volume Encryption \(VE\) and multi-tenant Encryption Key Management \(MTEKM\) licenses](#). These licenses are included with [ONTAP One](#).

- You must be a cluster or SVM administrator.
- The data SVM must not include any encrypted volumes or employ a key manager. If the data SVM includes encrypted volumes, you must migrate them before configuring the KMS.

Enable external key management

Enabling external key management depends on the specific key manager you use. Choose the tab of the appropriate key manager and environment.

AWS

Before you begin

- You must create a grant for the AWS KMS key that will be used by the IAM role managing encryption. The IAM role must include a policy that allows the following operations:
 - DescribeKey
 - Encrypt
 - Decrypt
- +
For more information, see AWS documentation for [grants](#).

Enable AWS KMS on an ONTAP SVM

1. Before you begin, obtain both the access key ID and secret key from your AWS KMS.
2. Set the privilege level to advanced:
`set -priv advanced`
3. Enable AWS KMS:
`security key-manager external aws enable -vserver svm_name -region AWS_region -key-id key_ID -encryption-context encryption_context`
4. When prompted, enter the secret key.
5. Confirm the AWS KMS was configured correctly:
`security key-manager external aws show -vserver svm_name`

Azure

Enable Azure Key Vault on an ONTAP SVM

1. Before you begin, you need to obtain the appropriate authentication credentials from your Azure account, either a client secret or certificate.
You must also ensure all nodes in the cluster are healthy. You can check this with the command `cluster show`.
2. Set privileged level to advanced
`set -priv advanced`
3. Enable AKV on the SVM
`security key-manager external azure enable -client-id client_id -tenant-id tenant_id -name -key-id key_id -authentication-method {certificate|client-secret}`
When prompted, enter either the client certificate or client secret from your Azure account.
4. Verify AKV is enabled correctly:
`security key-manager external azure show vserver svm_name`
If the service reachability is not OK, establish the connectivity to the AKV key management service via the data SVM LIF.

Google Cloud

Enable Cloud KMS on an ONTAP SVM

1. Before you begin, obtain the private key for the Google Cloud KMS account key file in a JSON format. This can be found in your GCP account.
You must also ensure all nodes in the cluster are healthy. You can check this with the command `cluster show`.

2. Set privileged level to advanced:

```
set -priv advanced
```

3. Enable Cloud KMS on the SVM

```
security key-manager external gcp enable -vserver svm_name -project-id  
project_id-key-ring-name key_ring_name -key-ring-location key_ring_location  
-key-name key_name
```

When prompted, enter the contents of the JSON file with the Service Account Private Key

4. Verify that Cloud KMS is configured with the correct parameters:

```
security key-manager external gcp show vserver svm_name
```

The status of `kms_wrapped_key_status` will be "UNKNOWN" if no encrypted volumes have been created.

If the service reachability is not OK, establish the connectivity to the GCP key management service via data SVM LIF.

If one or more encrypted volumes is already configured for a data SVM and the corresponding NVE keys are managed by the admin SVM onboard key manager, those keys should be migrated to the external key management service. To do this with the CLI, run the command:

```
security key-manager key migrate -from-Vserver admin_SVM -to-Vserver data_SVM
```

New encrypted volumes cannot be created for the tenant's data SVM until all NVE keys of the data SVM are successfully migrated.

Related information

- [Encrypting volumes with NetApp encryption solutions for Cloud Volumes ONTAP](#)

Enable onboard key management in ONTAP 9.6 and later (NVE)

You can use the Onboard Key Manager to secure the keys that the cluster uses to access encrypted data. You must enable the Onboard Key Manager on each cluster that accesses an encrypted volume or a self-encrypting disk.

About this task

You must run the `security key-manager onboard sync` command each time you add a node to the cluster.

If you have a MetroCluster configuration, you must run the `security key-manager onboard enable` command on the local cluster first, then run the `security key-manager onboard sync` command on the remote cluster, using the same passphrase on each. When you run the `security key-manager onboard enable` command from the local cluster and then synchronize on the remote cluster, you do not need to run the `enable` command again from the remote cluster.

By default, you are not required to enter the key manager passphrase when a node is rebooted. You can use the `cc-mode-enabled=yes` option to require that users enter the passphrase after a reboot.

For NVE, if you set `cc-mode-enabled=yes`, volumes you create with the `volume create` and `volume move start` commands are automatically encrypted. For `volume create`, you need not specify `-encrypt true`. For `volume move start`, you need not specify `-encrypt-destination true`.

When configuring ONTAP data at rest encryption, to meet the requirements for Commercial Solutions for Classified (CSfC) you must use NSE with NVE and ensure the Onboard Key Manager is enabled in Common Criteria mode. Refer to the [CSfC Solution Brief](#) for more information on CSfC.

When the Onboard Key Manager is enabled in Common Criteria mode (`cc-mode-enabled=yes`), system behavior is changed in the following ways:

- The system monitors for consecutive failed cluster passphrase attempts when operating in Common Criteria mode.

If you fail to enter the correct cluster passphrase at boot, encrypted volumes are not mounted. To correct this, you must reboot the node and enter the correct cluster passphrase. Once booted, the system allows up to 5 consecutive attempts to correctly enter the cluster passphrase in a 24-hour period for any command that requires the cluster passphrase as a parameter. If the limit is reached (for example, you have failed to correctly enter the cluster passphrase 5 times in a row) then you must either wait for the 24-hour timeout period to elapse, or you must reboot the node, in order to reset the limit.

- System image updates use the NetApp RSA-3072 code signing certificate together with SHA-384 code signed digests to check the image integrity instead of the usual NetApp RSA-2048 code signing certificate and SHA-256 code signed digests.

The upgrade command verifies that the image contents have not been altered or corrupted by checking various digital signatures. The image update process proceeds to the next step if validation succeeds; otherwise, the image update fails. See the `cluster image man` page for information concerning system updates.

The Onboard Key Manager stores keys in volatile memory. Volatile memory contents are cleared when the system is rebooted or halted. Under normal operating conditions, volatile memory contents will be cleared within 30s when a system is halted.

Before you begin

- You must be a cluster administrator to perform this task.
- You must configure the MetroCluster environment before you configure the Onboard Key Manager.

Steps

1. Start the key manager setup:

```
security key-manager onboard enable -cc-mode-enabled yes|no
```

Set `cc-mode-enabled=yes` to require that users enter the key manager passphrase after a reboot. For NVE, if you set `cc-mode-enabled=yes`, volumes you create with the `volume create` and `volume move start` commands are automatically encrypted. The `-cc-mode-enabled` option is not supported in MetroCluster configurations. The `security key-manager onboard enable` command replaces the `security key-manager setup` command.

The following example starts the key manager setup command on `cluster1` without requiring that the passphrase be entered after every reboot:

```
cluster1::> security key-manager onboard enable
```

```
Enter the cluster-wide passphrase for onboard key management in Vserver  
"cluster1":: <32..256 ASCII characters long text>  
Reenter the cluster-wide passphrase: <32..256 ASCII characters long  
text>
```

2. At the passphrase prompt, enter a passphrase between 32 and 256 characters, or for “cc-mode”, a passphrase between 64 and 256 characters.



If the specified “cc-mode” passphrase is less than 64 characters, there is a five-second delay before the key manager setup operation displays the passphrase prompt again.

3. At the passphrase confirmation prompt, reenter the passphrase.
4. Verify that the authentication keys have been created:

```
security key-manager key query -key-type NSE-AK
```



The `security key-manager key query` command replaces the `security key-manager query key` command. For complete command syntax, see the man page.

The following example verifies that authentication keys have been created for `cluster1`:

```

cluster1::> security key-manager key query -key-type NSE-AK
      Node: node1
      Vserver: cluster1
      Key Manager: onboard
      Key Manager Type: OKM
      Key Manager Policy: -

Key Tag                                Key Type Encryption  Restored
-----
node1                                NSE-AK   AES-256   true

      Key ID:
00000000000000000000200000000000100056178fc6ace6d91472df8a9286daacc00000000
00000000

node1                                NSE-AK   AES-256   true

      Key ID:
00000000000000000000200000000000100df1689a148fdfbf9c2b198ef974d0baa00000000
00000000

2 entries were displayed.

```

5. Optionally, convert plain text volumes to encrypted volumes.

```
volume encryption conversion start
```

The Onboard Key Manager must be fully configured before you convert the volumes. In a MetroCluster environment, the Onboard Key Manager must be configured on both sites.

After you finish

Copy the passphrase to a secure location outside the storage system for future use.

Whenever you configure the Onboard Key Manager passphrase, you should also back up the information manually to a secure location outside the storage system for use in case of a disaster. See [Back up onboard key management information manually](#).

Enable onboard key management in ONTAP 9.5 and earlier (NVE)

You can use the Onboard Key Manager to secure the keys that the cluster uses to access encrypted data. You must enable Onboard Key Manager on each cluster that accesses an encrypted volume or a self-encrypting disk.

About this task

You must run the `security key-manager setup` command each time you add a node to the cluster.

If you have a MetroCluster configuration, review these guidelines:

- In ONTAP 9.5, you must run `security key-manager setup` on the local cluster and `security key-manager setup -sync-metrocluster-config yes` on the remote cluster, using the same passphrase on each.
- Prior to ONTAP 9.5, you must run `security key-manager setup` on the local cluster, wait approximately 20 seconds, and then run `security key-manager setup` on the remote cluster, using the same passphrase on each.

By default, you are not required to enter the key manager passphrase when a node is rebooted. Beginning with ONTAP 9.4, you can use the `-enable-cc-mode yes` option to require that users enter the passphrase after a reboot.

For NVE, if you set `-enable-cc-mode yes`, volumes you create with the `volume create` and `volume move start` commands are automatically encrypted. For `volume create`, you need not specify `-encrypt true`. For `volume move start`, you need not specify `-encrypt-destination true`.



After a failed passphrase attempt, you must reboot the node again.

Before you begin

- If you are using NSE or NVE with an external key management (KMIP) server, you must have deleted the external key manager database.

[Transitioning to onboard key management from external key management](#)

- You must be a cluster administrator to perform this task.
- You must configure the MetroCluster environment before you configure the Onboard Key Manager.

Steps

1. Start the key manager setup:

```
security key-manager setup -enable-cc-mode yes|no
```



Beginning with ONTAP 9.4, you can use the `-enable-cc-mode yes` option to require that users enter the key manager passphrase after a reboot. For NVE, if you set `-enable-cc-mode yes`, volumes you create with the `volume create` and `volume move start` commands are automatically encrypted.

The following example starts setting up the key manager on cluster1 without requiring that the passphrase be entered after every reboot:


```

cluster1::> security key-manager setup
Welcome to the key manager setup wizard, which will lead you through
the steps to add boot information.

...

Would you like to use onboard key-management? {yes, no} [yes]:
Enter the cluster-wide passphrase:    <32..256 ASCII characters long
text>
Reenter the cluster-wide passphrase:  <32..256 ASCII characters long
text>

```

2. Enter `yes` at the prompt to configure onboard key management.
3. At the passphrase prompt, enter a passphrase between 32 and 256 characters, or for “cc-mode”, a passphrase between 64 and 256 characters.



If the specified “cc-mode” passphrase is less than 64 characters, there is a five-second delay before the key manager setup operation displays the passphrase prompt again.

4. At the passphrase confirmation prompt, reenter the passphrase.
5. Verify that keys are configured for all nodes:

```
security key-manager key show
```

For the complete command syntax, see the man page.

```

cluster1::> security key-manager key show

Node: node1
Key Store: onboard
Key ID                                                                                               Used By
-----
-----
0000000000000000000020000000000010059851742AF2703FC91369B7DB47C4722 NSE-AK
000000000000000000002000000000001008C07CC0AF1EF49E0105300EFC83004BF NSE-AK

Node: node2
Key Store: onboard
Key ID                                                                                               Used By
-----
-----
0000000000000000000020000000000010059851742AF2703FC91369B7DB47C4722 NSE-AK
000000000000000000002000000000001008C07CC0AF1EF49E0105300EFC83004BF NSE-AK

```

6. Optionally, convert plain text volumes to encrypted volumes.

```
volume encryption conversion start
```

The Onboard Key Manager must be fully configured before you convert the volumes. In a MetroCluster environment, the Onboard Key Manager must be configured on both sites.

After you finish

Copy the passphrase to a secure location outside the storage system for future use.

Whenever you configure the Onboard Key Manager passphrase, you should also back up the information manually to a secure location outside the storage system for use in case of a disaster. See [Back up onboard key management information manually](#).

Enable onboard key management in newly added nodes

You can use the Onboard Key Manager to secure the keys that the cluster uses to access encrypted data. You must enable Onboard Key Manager on each cluster that accesses an encrypted volume or a self-encrypting disk.

For ONTAP 9.5 and earlier, you must run the `security key-manager setup` command each time you add a node to the cluster.



For ONTAP 9.6 and later, you must run the `security key-manager sync` command each time you add a node to the cluster.

If you add a node to a cluster that has onboard key management configured, you will run this command to refresh the missing keys.

If you have a MetroCluster configuration, review these guidelines:

- Beginning with ONTAP 9.6, you must run `security key-manager onboard enable` on the local cluster first, then run `security key-manager onboard sync` on the remote cluster, using the same passphrase on each.
- In ONTAP 9.5, you must run `security key-manager setup` on the local cluster and `security key-manager setup -sync-metrocluster-config yes` on the remote cluster, using the same passphrase on each.
- Prior to ONTAP 9.5, you must run `security key-manager setup` on the local cluster, wait approximately 20 seconds, and then run `security key-manager setup` on the remote cluster, using the same passphrase on each.

By default, you are not required to enter the key manager passphrase when a node is rebooted. Beginning with ONTAP 9.4, you can use the `-enable-cc-mode yes` option to require that users enter the passphrase after a reboot.

For NVE, if you set `-enable-cc-mode yes`, volumes you create with the `volume create` and `volume move start` commands are automatically encrypted. For `volume create`, you need not specify `-encrypt true`. For `volume move start`, you need not specify `-encrypt-destination true`.



After a failed passphrase attempt, you must reboot the node again.

Encrypt volume data with NVE

Encrypt volume data with NVE overview

Beginning with ONTAP 9.7, aggregate and volume encryption is enabled by default when you have the VE license and onboard or external key management. For ONTAP 9.6 and earlier, you can enable encryption on a new volume or on an existing volume. You must have installed the VE license and enabled key management before you can enable volume encryption. NVE is FIPS-140-2 level 1 compliant.

Enable aggregate-level encryption with VE license

Beginning with ONTAP 9.7, newly created aggregates and volumes are encrypted by default when you have the [VE license](#) and onboard or external key management. Beginning with ONTAP 9.6, you can use aggregate-level encryption to assign keys to the containing aggregate for the volumes to be encrypted.

About this task

You must use aggregate-level encryption if you plan to perform inline or background aggregate-level deduplication. Aggregate-level deduplication is otherwise not supported by NVE.

An aggregate enabled for aggregate-level encryption is called an *NAE aggregate* (for NetApp Aggregate Encryption). All volumes in an NAE aggregate must be encrypted with NAE or NVE encryption. With aggregate-level encryption, volumes you create in the aggregate are encrypted with NAE encryption by default. You can override the default to use NVE encryption instead.

Plain text volumes are not supported in NAE aggregates.

Before you begin

You must be a cluster administrator to perform this task.

Steps

1. Enable or disable aggregate-level encryption:

To...	Use this command...
Create an NAE aggregate with ONTAP 9.7 or later	<pre>storage aggregate create -aggregate aggregate_name -node node_name</pre>
Create an NAE aggregate with ONTAP 9.6	<pre>storage aggregate create -aggregate aggregate_name -node node_name -encrypt-with -aggr-key true</pre>
Convert a non-NAE aggregate to an NAE aggregate	<pre>storage aggregate modify -aggregate aggregate_name -node node_name -encrypt-with -aggr-key true</pre>

Convert an NAE aggregate to a non-NAE aggregate

```
storage aggregate modify -aggregate
aggregate_name -node node_name -encrypt-with
-aggr-key false
```

For complete command syntax, see the man pages.

The following command enables aggregate-level encryption on `aggr1`:

- ONTAP 9.7 or later:

```
cluster1::> storage aggregate create -aggregate aggr1
```

- ONTAP 9.6 or earlier:

```
cluster1::> storage aggregate create -aggregate aggr1 -encrypt-with
-aggr-key true
```

2. Verify that the aggregate is enabled for encryption:

```
storage aggregate show -fields encrypt-with-aggr-key
```

For complete command syntax, see the man page.

The following command verifies that `aggr1` is enabled for encryption:

```
cluster1::> storage aggregate show -fields encrypt-with-aggr-key
aggregate          encrypt-aggr-key
-----
aggr0_vsim4        false
aggr1               true
2 entries were displayed.
```

After you finish

Run the `volume create` command to create the encrypted volumes.

If you are using a KMIP server to store the encryption keys for a node, ONTAP automatically “pushes” an encryption key to the server when you encrypt a volume.

Enable encryption on a new volume

You can use the `volume create` command to enable encryption on a new volume.

About this task

You can encrypt volumes using NetApp Volume Encryption (NVE) and, beginning with ONTAP 9.6, NetApp Aggregate Encryption (NAE). To learn more about NAE and NVE, refer to the [volume encryption overview](#).

The procedure to enable encryption on a new volume in ONTAP varies based on the version of ONTAP you are using and your specific configuration:

- Beginning with ONTAP 9.4, if you enable `cc-mode` when you set up the Onboard Key Manager, volumes you create with the `volume create` command are automatically encrypted, whether or not you specify `-encrypt true`.
- In ONTAP 9.6 and earlier releases, you must use `-encrypt true` with `volume create` commands to enable encryption (provided you did not enable `cc-mode`).
- If you want to create an NAE volume in ONTAP 9.6, you must enable NAE at the aggregate level. Refer to [Enable aggregate-level encryption with the VE license](#) for more details on this task.
- Beginning with ONTAP 9.7, newly created volumes are encrypted by default when you have the [VE license](#) and onboard or external key management. By default, new volumes created in an NAE aggregate will be of type NAE rather than NVE.
 - In ONTAP 9.7 and later releases, if you add `-encrypt true` to the `volume create` command to create a volume in an NAE aggregate, the volume will have NVE encryption instead of NAE. All volumes in an NAE aggregate must be encrypted with either NVE or NAE.



Plaintext volumes are not supported in NAE aggregates.

Steps

1. Create a new volume and specify whether encryption is enabled on the volume. If the new volume is in an NAE aggregate, by default the volume will be an NAE volume:

To create...	Use this command...
An NAE volume	<code>volume create -vserver SVM_name -volume volume_name -aggregate aggregate_name</code>
An NVE volume	<code>volume create -vserver SVM_name -volume volume_name -aggregate aggregate_name -encrypt true +</code> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p>In ONTAP 9.6 and earlier where NAE is not supported, <code>-encrypt true</code> specifies that the volume should be encrypted with NVE. In ONTAP 9.7 and later where volumes are created in NAE aggregates, <code>-encrypt true</code> overrides the default encryption type of NAE to create an NVE volume instead.</p> </div>
A plain text volume	<code>volume create -vserver SVM_name -volume volume_name -aggregate aggregate_name -encrypt false</code>

For complete command syntax, refer to the command reference page for `volume create`.

2. Verify that volumes are enabled for encryption:

```
volume show -is-encrypted true
```

For complete command syntax, see the [ONTAP command reference](#).

Result

If you are using a KMIP server to store the encryption keys for a node, ONTAP automatically "pushes" an encryption key to the server when you encrypt a volume.

Enable encryption on an existing volume

You can use either the `volume move start` or the `volume encryption conversion start` command to enable encryption on an existing volume.

About this task

- Beginning with ONTAP 9.3, you can use the `volume encryption conversion start` command to enable encryption of an existing volume "in place," without having to move the volume to a different location. Alternatively, you can use the `volume move start` command.
- For ONTAP 9.2 and earlier, you can use only the `volume move start` command to enable encryption by moving an existing volume.

Enable encryption on an existing volume with the `volume encryption conversion start` command

Beginning with ONTAP 9.3, you can use the `volume encryption conversion start` command to enable encryption of an existing volume "in place," without having to move the volume to a different location.

After you start a conversion operation, it must be completed. If you encounter a performance issue during the operation, you can run the `volume encryption conversion pause` command to pause the operation, and the `volume encryption conversion resume` command to resume the operation.



You cannot use `volume encryption conversion start` to convert a SnapLock volume.

Steps

1. Enable encryption on an existing volume:

```
volume encryption conversion start -vserver SVM_name -volume volume_name
```

For the entire command syntax, see the man page for the command.

The following command enables encryption on existing volume `vol1`:

```
cluster1::> volume encryption conversion start -vserver vs1 -volume vol1
```

The system creates an encryption key for the volume. The data on the volume is encrypted.

2. Verify the status of the conversion operation:

```
volume encryption conversion show
```

For the entire command syntax, see the man page for the command.

The following command displays the status of the conversion operation:

```
cluster1::> volume encryption conversion show
```

Vserver	Volume	Start Time	Status
vs1	vol1	9/18/2017 17:51:41	Phase 2 of 2 is in progress.

3. When the conversion operation is completed, verify that the volume is enabled for encryption:

```
volume show -is-encrypted true
```

For the entire command syntax, see the man page for the command.

The following command displays the encrypted volumes on `cluster1`:

```
cluster1::> volume show -is-encrypted true
```

Vserver	Volume	Aggregate	State	Type	Size	Available	Used
vs1	vol1	aggr2	online	RW	200GB	160.0GB	20%

Result

If you are using a KMIP server to store the encryption keys for a node, ONTAP automatically “pushes” an encryption key to the server when you encrypt a volume.

Enable encryption on an existing volume with the volume move start command

You can use the `volume move start` command to enable encryption by moving an existing volume. You must use `volume move start` in ONTAP 9.2 and earlier. You can use the same aggregate or a different aggregate.

About this task

- Beginning with ONTAP 9.8, you can use `volume move start` to enable encryption on a SnapLock or FlexGroup volume.
- Beginning with ONTAP 9.4, if you enable “cc-mode” when you set up the Onboard Key Manager, volumes you create with the `volume move start` command are automatically encrypted. You need not specify `-encrypt-destination true`.
- Beginning with ONTAP 9.6, you can use aggregate-level encryption to assign keys to the containing aggregate for the volumes to be moved. A volume encrypted with a unique key is called an *NVE volume* (meaning it uses NetApp Volume Encryption). A volume encrypted with an aggregate-level key is called an *NAE volume* (for NetApp Aggregate Encryption). Plaintext volumes are not supported in NAE aggregates.
- Beginning with ONTAP 9.14.1, you can encrypt an SVM root volume with NVE. For more information, see [Configure NetApp Volume Encryption on an SVM root volume](#).

Before you begin

You must be a cluster administrator to perform this task, or an SVM administrator to whom the cluster administrator has delegated authority.

Delegating authority to run the volume move command

Steps

1. Move an existing volume and specify whether encryption is enabled on the volume:

To convert...	Use this command...
A plaintext volume to an NVE volume	<pre>volume move start -vserver SVM_name -volume volume_name -destination-aggregate aggregate_name -encrypt-destination true</pre>
An NVE or plaintext volume to an NAE volume (assuming aggregate-level encryption is enabled on the destination)	<pre>volume move start -vserver SVM_name -volume volume_name -destination-aggregate aggregate_name -encrypt-with-aggr-key true</pre>
An NAE volume to an NVE volume	<pre>volume move start -vserver SVM_name -volume volume_name -destination-aggregate aggregate_name -encrypt-with-aggr-key false</pre>
An NAE volume to a plaintext volume	<pre>volume move start -vserver SVM_name -volume volume_name -destination-aggregate aggregate_name -encrypt-destination false -encrypt-with-aggr-key false</pre>
An NVE volume to a plaintext volume	<pre>volume move start -vserver SVM_name -volume volume_name -destination-aggregate aggregate_name -encrypt-destination false</pre>

For the entire command syntax, see the man page for the command.

The following command converts a plaintext volume named `vol1` to an NVE volume:

```
cluster1::> volume move start -vserver vs1 -volume vol1 -destination -aggregate aggr2 -encrypt-destination true
```

Assuming aggregate-level encryption is enabled on the destination, the following command converts an NVE or plaintext volume named `vol1` to an NAE volume:

```
cluster1::> volume move start -vserver vs1 -volume vol1 -destination -aggregate aggr2 -encrypt-with-aggr-key true
```

The following command converts an NAE volume named `vol2` to an NVE volume:


```
cluster1::> volume move start -vserver vs1 -volume vol2 -destination
-aggregate aggr2 -encrypt-with-aggr-key false
```

The following command converts an NAE volume named `vol2` to a plaintext volume:

```
cluster1::> volume move start -vserver vs1 -volume vol2 -destination
-aggregate aggr2 -encrypt-destination false -encrypt-with-aggr-key false
```

The following command converts an NVE volume named `vol2` to a plaintext volume:

```
cluster1::> volume move start -vserver vs1 -volume vol2 -destination
-aggregate aggr2 -encrypt-destination false
```

2. View the encryption type of cluster volumes:

```
volume show -fields encryption-type none|volume|aggregate
```

The `encryption-type` field is available in ONTAP 9.6 and later.

For the entire command syntax, see the man page for the command.

The following command displays the encryption type of volumes in `cluster2`:

```
cluster2::> volume show -fields encryption-type

vserver  volume  encryption-type
-----  -
vs1      vol1    none
vs2      vol2    volume
vs3      vol3    aggregate
```

3. Verify that volumes are enabled for encryption:

```
volume show -is-encrypted true
```

For the entire command syntax, see the man page for the command.

The following command displays the encrypted volumes on `cluster2`:

```
cluster2::> volume show -is-encrypted true
```

Vserver	Volume	Aggregate	State	Type	Size	Available	Used
-----	-----	-----	-----	-----	-----	-----	-----
vs1	vol1	aggr2	online	RW	200GB	160.0GB	20%

Result

If you are using a KMIP server to store the encryption keys for a node, ONTAP automatically pushes an encryption key to the server when you encrypt a volume.

Configure NetApp Volume Encryption on an SVM root volume

Beginning with ONTAP 9.14.1, you can enable NetApp Volume Encryption (NVE) on a storage VM (SVM) root volume. With NVE, the root volume is encrypted with a unique key, enabling greater security on the SVM.

About this task

NVE on an SVM root volume can only be enabled after the SVM has been created.

Before you begin

- The SVM root volume must not be on an aggregate encrypted with NetApp Aggregate Encryption (NAE).
- You must have enabled encryption with the Onboard Key Manager or an external key manager.
- You must be running ONTAP 9.14.1 or later.
- To migrate an SVM containing a root volume encrypted with NVE, you must convert the SVM root volume to a plain text volume after the migration completes then re-encrypt the SVM root volume.
 - If the destination aggregate of the SVM migration uses NAE, the root volume inherits NAE by default.
- If the SVM is in an SVM disaster recovery relationship:
 - Encryption settings on a mirrored SVM are not copied to the destination. If you enable NVE on the source or destination, you must separately enable NVE on the mirrored SVM root volume.
 - If all aggregates in the destination cluster use NAE, the SVM root volume will use NAE.

Steps

You can enable NVE on an SVM root volume with the ONTAP CLI or System Manager.

CLI

You can enable NVE on the SVM root volume in-place or by moving the volume between aggregates.

Encrypt the root volume in place

1. Convert the root volume to an encrypted volume:

```
volume encryption conversion start -vserver svm_name -volume volume
```

2. Confirm the encryption succeeded. The `volume show -encryption-type volume` displays a list of all volumes using NVE.

Encrypt the SVM root volume by moving it


1. Initiate a volume move:

```
volume move start -vserver svm_name -volume volume -destination-aggregate aggregate -encrypt-with-aggr-key false -encrypt-destination true
```

For more information about `volume move`, see [Move a volume](#).

2. Confirm the `volume move` operation succeeded with the `volume move show` command. The `volume show -encryption-type volume` displays a list of all volumes using NVE.

System Manager

1. Navigate to **Storage > Volumes**.
2. Next to the name of the SVM root volume you want to encrypt, select  then **Edit**.
3. Under the **Storage and Optimization** heading, select **Enable encryption**.
4. Select **Save**.

Enable node root volume encryption

Beginning with ONTAP 9.8, you can use NetApp Volume Encryption to protect the root volume of your node.



About this task

This procedure applies to the node root volume. It does not apply to SVM root volumes. SVM root volumes can be protected through aggregate-level encryption and, [beginning with ONTAP 9.14.1, NVE](#).

Once root volume encryption begins, it must complete. You cannot pause the operation. Once encryption is complete, you cannot assign a new key to the root volume and you cannot perform a secure-purge operation.

Before you begin

- Your system must be using an HA configuration.
- Your node root volume must already be created.
- Your system must have an onboard key manager or an external key management server using the Key Management Interoperability Protocol (KMIP).

Steps

1. Encrypt the root volume:

```
volume encryption conversion start -vserver SVM_name -volume root_vol_name
```

2. Verify the status of the conversion operation:

```
volume encryption conversion show
```

3. When the conversion operation is complete, verify that the volume is encrypted:

```
volume show -fields
```

The following shows example output for an encrypted volume.

```
::> volume show -vserver xyz -volume vol0 -fields is-encrypted
vserver      volume is-encrypted
-----
xyz          vol0    true
```

Configure NetApp hardware-based encryption

Configure NetApp hardware-based encryption overview

NetApp hardware-based encryption supports full-disk encryption (FDE) of data as it is written. The data cannot be read without an encryption key stored on the firmware. The encryption key, in turn, is accessible only to an authenticated node.

Understanding NetApp hardware-based encryption

A node authenticates itself to a self-encrypting drive using an authentication key retrieved from an external key management server or Onboard Key Manager:

- The external key management server is a third-party system in your storage environment that serves keys to nodes using the Key Management Interoperability Protocol (KMIP). It is a best practice to configure external key management servers on a different storage system from your data.
- The Onboard Key Manager is a built-in tool that serves authentication keys to nodes from the same storage system as your data.

You can use NetApp Volume Encryption with hardware-based encryption to “double encrypt” data on self-encrypting drives.

When self-encrypting drives are enabled, the core dump is also encrypted.



If an HA pair is using encrypting SAS or NVMe drives (SED, NSE, FIPS), you must follow the instructions in the topic [Returning a FIPS drive or SED to unprotected mode](#) for all drives within the HA pair prior to initializing the system (boot options 4 or 9). Failure to do this may result in future data loss if the drives are repurposed.

Supported self-encrypting drive types

Two types of self-encrypting drives are supported:

- Self-encrypting FIPS-certified SAS or NVMe drives are supported on all FAS and AFF systems. These drives, called *FIPS drives*, conform to the requirements of Federal Information Processing Standard Publication 140-2, level 2. The certified capabilities enable protections in addition to encryption, such as preventing denial-of-service attacks on the drive. FIPS drives cannot be mixed with other types of drives on the same node or HA pair.
- Beginning with ONTAP 9.6, self-encrypting NVMe drives that have not undergone FIPS testing are supported on AFF A800, A320, and later systems. These drives, called *SEDs*, offer the same encryption capabilities as FIPS drives, but can be mixed with non-encrypting drives on the same node or HA pair.
- All FIPS validated drives use a firmware cryptographic module that has been through FIPS validation. The FIPS drive cryptographic module does not use any keys that are generated outside of the drive (the authentication passphrase that is input to the drive is used by the drive's firmware cryptographic module to obtain a key encryption key).



Non-encrypting drives are drives that are not SEDs or FIPS drives.



If you are using NSE on a system with a Flash Cache module, you should also enable NVE or NAE. NSE does not encrypt data that resides on the Flash Cache module.

When to use external key management

Although it is less expensive and typically more convenient to use the onboard key manager, you should use external key management if any of the following are true:

- Your organization's policy requires a key management solution that uses a FIPS 140-2 Level 2 (or higher) cryptographic module.
- You need a multi-cluster solution, with centralized management of encryption keys.
- Your business requires the added security of storing authentication keys on a system or in a location different from the data.

Support details

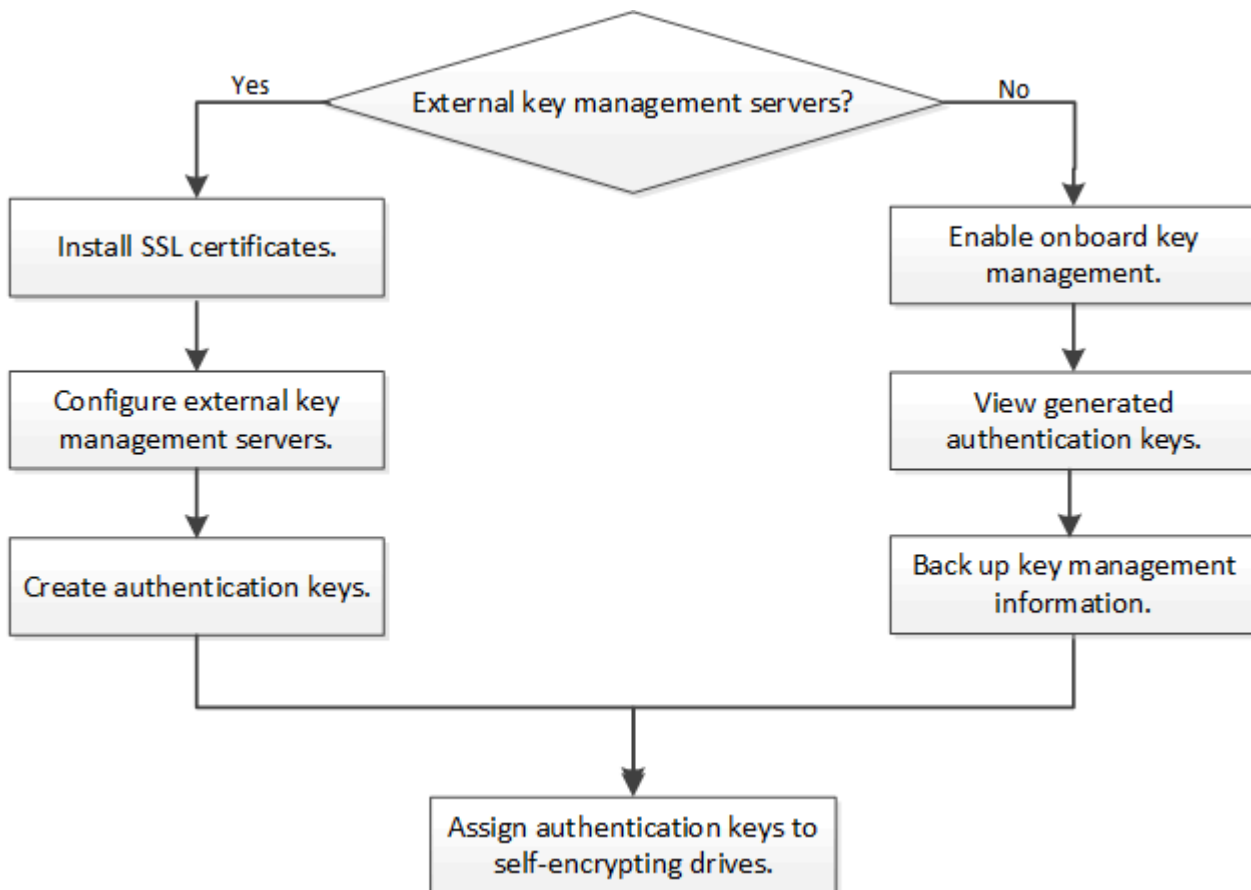
The following table shows important hardware encryption support details. See the Interoperability Matrix for the latest information about supported KMIP servers, storage systems, and disk shelves.

Resource or feature	Support details
Non-homogeneous disk sets	<ul style="list-style-type: none">• FIPS drives cannot be mixed with other types of drives on the same node or HA pair. Conforming HA pairs can coexist with non-conforming HA pairs in the same cluster.• SEDs can be mixed with non-encrypting drives on the same node or HA pair.
Drive type	<ul style="list-style-type: none">• FIPS drives can be SAS or NVMe drives.• SEDs must be NVMe drives.

10 Gb network interfaces	Beginning with ONTAP 9.3, KMIP key management configurations support 10 Gb network interfaces for communications with external key management servers.
Ports for communication with the key management server	Beginning with ONTAP 9.3, you can use any storage controller port for communication with the key management server. Otherwise, you should use port e0M for communication with key management servers. Depending on the storage controller model, certain network interfaces might not be available during the boot process for communication with key management servers.
MetroCluster (MCC)	<ul style="list-style-type: none"> • NVMe drives support MCC. • SAS drives do not support MCC.

Hardware-based encryption workflow

You must configure key management services before the cluster can authenticate itself to the self-encrypting drive. You can use an external key management server or an onboard key manager.



Related information

- [NetApp Hardware Universe](#)
- [NetApp Volume Encryption and NetApp Aggregate Encryption](#)

Configure external key management

Configure external key management overview

You can use one or more external key management servers to secure the keys that the cluster uses to access encrypted data. An external key management server is a third-party system in your storage environment that serves keys to nodes using the Key Management Interoperability Protocol (KMIP).

For ONTAP 9.1 and earlier versions, node management LIFs must be assigned to ports that are configured with the node management role before you can use the external key manager.

NetApp Volume Encryption (NVE) can be implemented with Onboard Key Manager in ONTAP 9.1 and later. In ONTAP 9.3 and later, NVE can be implemented with external key management (KMIP) and Onboard Key Manager. Beginning in ONTAP 9.11.1, you can configure multiple external key managers in a cluster. See [Configure clustered key servers](#).

Collect network information in ONTAP 9.2 and earlier

If you are using ONTAP 9.2 or earlier, you should fill out the network configuration worksheet before enabling external key management.



Beginning with ONTAP 9.3, the system discovers all needed network information automatically.

Item	Notes	Value
Key management network interface name		
Key management network interface IP address	IP address of node management LIF, in IPv4 or IPv6 format	
Key management network interface IPv6 network prefix length	If you are using IPv6, the IPv6 network prefix length	
Key management network interface subnet mask		
Key management network interface gateway IP address		
IPv6 address for the cluster network interface	Required only if you are using IPv6 for the key management network interface	

Port number for each KMIP server	Optional. The port number must be the same for all KMIP servers. If you do not provide a port number, it defaults to port 5696, which is the Internet Assigned Numbers Authority (IANA) assigned port for KMIP.	
Key tag name	Optional. The key tag name is used to identify all keys belonging to a node. The default key tag name is the node name.	

Related information

[NetApp Technical Report 3954: NetApp Storage Encryption Preinstallation Requirements and Procedures for IBM Tivoli Lifetime Key Manager](#)

[NetApp Technical Report 4074: NetApp Storage Encryption Preinstallation Requirements and Procedures for SafeNet KeySecure](#)

Install SSL certificates on the cluster

The cluster and KMIP server use KMIP SSL certificates to verify each other's identity and establish an SSL connection. Before configuring the SSL connection with the KMIP server, you must install the KMIP client SSL certificates for the cluster, and the SSL public certificate for the root certificate authority (CA) of the KMIP server.

About this task

In an HA pair, both nodes must use the same public and private KMIP SSL certificates. If you connect multiple HA pairs to the same KMIP server, all nodes in the HA pairs must use the same public and private KMIP SSL certificates.

Before you begin

- The time must be synchronized on the server creating the certificates, the KMIP server, and the cluster.
- You must have obtained the public SSL KMIP client certificate for the cluster.
- You must have obtained the private key associated with the SSL KMIP client certificate for the cluster.
- The SSL KMIP client certificate must not be password-protected.
- You must have obtained the SSL public certificate for the root certificate authority (CA) of the KMIP server.
- In a MetroCluster environment, you must install the same KMIP SSL certificates on both clusters.



You can install the client and server certificates on the KMIP server before or after installing the certificates on the cluster.

Steps

1. Install the SSL KMIP client certificates for the cluster:

```
security certificate install -vserver admin_svm_name -type client
```

You are prompted to enter the SSL KMIP public and private certificates.


```
cluster1::> security certificate install -vserver cluster1 -type client
```

2. Install the SSL public certificate for the root certificate authority (CA) of the KMIP server:

```
security certificate install -vserver admin_svm_name -type server-ca
```

```
cluster1::> security certificate install -vserver cluster1 -type server-ca
```

Enable external key management in ONTAP 9.6 and later (HW-based)

You can use one or more KMIP servers to secure the keys the cluster uses to access encrypted data. You can connect up to four KMIP servers to a node. A minimum of two servers is recommended for redundancy and disaster recovery.

Beginning in ONTAP 9.11.1, you can add up to 3 secondary key servers per primary key server to create a clustered key server. For more information, see [Configure clustered external key servers](#).

Before you begin

- The KMIP SSL client and server certificates must have been installed.
- You must be a cluster administrator to perform this task.
- You must configure the MetroCluster environment before you configure an external key manager.
- In a MetroCluster environment, you must install the same KMIP SSL certificate on both clusters.

Steps

1. Configure key manager connectivity for the cluster:

```
security key-manager external enable -vserver admin_SVM -key-servers  
host_name|IP_address:port,... -client-cert client_certificate -server-ca-cert  
server_CA_certificates
```



- The `security key-manager external enable` command replaces the `security key-manager setup` command. You can run the `security key-manager external modify` command to change the external key management configuration. For complete command syntax, see the man pages.
- In a MetroCluster environment, if you are configuring external key management for the admin SVM, you must repeat the `security key-manager external enable` command on the partner cluster.

The following command enables external key management for `cluster1` with three external key servers. The first key server is specified using its hostname and port, the second is specified using an IP address and the default port, and the third is specified using an IPv6 address and port:

```
cluster1::> security key-manager external enable -key-servers  
ks1.local:15696,10.0.0.10,[fd20:8b1e:b255:814e:32bd:f35c:832c:5a09]:1234  
-client-cert AdminVserverClientCert -server-ca-certs  
AdminVserverServerCaCert
```

2. Verify that all configured KMIP servers are connected:

```
security key-manager external show-status -node node_name -vserver SVM -key
-server host_name|IP_address:port -key-server-status available|not-
responding|unknown
```



The `security key-manager external show-status` command replaces the `security key-manager show -status` command. For complete command syntax, see the man page.

```
cluster1::> security key-manager external show-status

Node   Vserver   Key Server                                     Status
----   -
node1
  cluster1
    10.0.0.10:5696                             available
    fd20:8b1e:b255:814e:32bd:f35c:832c:5a09:1234 available
    ks1.local:15696                             available
node2
  cluster1
    10.0.0.10:5696                             available
    fd20:8b1e:b255:814e:32bd:f35c:832c:5a09:1234 available
    ks1.local:15696                             available

6 entries were displayed.
```

Enable external key management in ONTAP 9.5 and earlier (HW-based)

You can use one or more KMIP servers to secure the keys the cluster uses to access encrypted data. You can connect up to four KMIP servers to a node. A minimum of two servers is recommended for redundancy and disaster recovery.

About this task

ONTAP configures KMIP server connectivity for all nodes in the cluster.

Before you begin

- The KMIP SSL client and server certificates must have been installed.
- You must be a cluster administrator to perform this task.
- You must configure the MetroCluster environment before you configure an external key manager.
- In a MetroCluster environment, you must install the same KMIP SSL certificate on both clusters.

Steps

1. Configure key manager connectivity for cluster nodes:

```
security key-manager setup
```

The key manager setup starts.



In a MetroCluster environment, you must run this command on both clusters.

2. Enter the appropriate response at each prompt.
3. Add a KMIP server:

```
security key-manager add -address key_management_server_ipaddress
```

```
cluster1::> security key-manager add -address 20.1.1.1
```



In a MetroCluster environment, you must run this command on both clusters.

4. Add an additional KMIP server for redundancy:

```
security key-manager add -address key_management_server_ipaddress
```

```
cluster1::> security key-manager add -address 20.1.1.2
```



In a MetroCluster environment, you must run this command on both clusters.

5. Verify that all configured KMIP servers are connected:

```
security key-manager show -status
```

For complete command syntax, see the man page.

```
cluster1::> security key-manager show -status
```

Node	Port	Registered Key Manager	Status
cluster1-01	5696	20.1.1.1	available
cluster1-01	5696	20.1.1.2	available
cluster1-02	5696	20.1.1.1	available
cluster1-02	5696	20.1.1.2	available

6. Optionally, convert plain text volumes to encrypted volumes.

```
volume encryption conversion start
```

An external key manager must be fully configured before you convert the volumes. In a MetroCluster environment, an external key manager must be configured on both sites.

Configure clustered external key servers

Beginning in ONTAP 9.11.1, you can configure connectivity to clustered external key management servers on an SVM. With clustered key servers, you can designate primary and secondary key servers on a SVM. When registering keys, ONTAP will first attempt to access a primary key server before sequentially attempting to access secondary servers until the operation completes successfully, preventing duplication of keys.

External key servers can be used for NSE, NVE, NAE, and SED keys. An SVM can support up to four primary external KMIP servers. Each primary server can support up to three secondary key servers.

Before you begin

- [KMIP key management must be enabled for the SVM.](#)
- This process only supports key servers that use KMIP. For a list of supported key servers, check the [NetApp Interoperability Matrix Tool](#).
- All nodes in the cluster must be running ONTAP 9.11.1 or later.
- The order of servers list arguments in the `-secondary-key-servers` parameter reflects the access order of the external key management (KMIP) servers.

Create a clustered key server

The configuration procedure depends on whether or not you have configured a primary key server.

Add primary and secondary key servers to an SVM

1. Confirm that no key management has been enabled for the cluster:

```
security key-manager external show -vserver svm_name
```

If the SVM already has the maximum of four primary key servers enabled, you must remove one of the existing primary key servers before adding a new one.

2. Enable the primary key manager:

```
security key-manager external enable -vserver svm_name -key-servers  
server_ip -client-cert client_cert_name -server-ca-certs  
server_ca_cert_names
```

3. Modify the primary key server to add secondary key servers. The `-secondary-key-servers` parameter accepts a comma-separated list of up to three key servers.

```
security key-manager external modify-server -vserver svm_name -key-servers  
primary_key_server -secondary-key-servers list_of_key_servers
```

Add secondary key servers to an existing primary key server

1. Modify the primary key server to add secondary key servers. The `-secondary-key-servers` parameter accepts a comma-separated list of up to three key servers.

```
security key-manager external modify-server -vserver svm_name -key-servers  
primary_key_server -secondary-key-servers list_of_key_servers
```

For more information about secondary key servers, see [Modify secondary key servers](#).

Modify clustered key servers

You can modify external key servers clusters by changing the status (primary or secondary) of particular key

servers, add and removing secondary key servers, or by changing the access order of secondary key servers.

Convert primary and secondary key servers

To convert a primary key server into a secondary key server, you must first remove it from the SVM with the `security key-manager external remove-servers` command.

To convert a secondary key server into a primary key server, you must first remove the secondary key server from its existing primary key server. See [Modify secondary key servers](#). If you convert a secondary key server to a primary server while removing an existing key, attempting to add a new server before completing the removal and conversion can result in the duplication of keys.

Modify secondary key servers

Secondary key servers are managed with the `-secondary-key-servers` parameter of the `security key-manager external modify-server` command. The `-secondary-key-servers` parameter accepts a comma-separated list. The specified order of the secondary key servers in the list determines the access sequence for the secondary key servers. The access order can be modified by running the command `security key-manager external modify-server` with the secondary key servers entered in a different sequence.

To remove a secondary key server, the `-secondary-key-servers` arguments should include the key servers you want to keep while omitting the one to be removed. To remove all secondary key servers, use the argument `-`, signifying none.

For additional information, refer to the `security key-manager external` page in the [ONTAP command reference](#).

Create authentication keys in ONTAP 9.6 and later

You can use the `security key-manager key create` command to create the authentication keys for a node and store them on the configured KMIP servers.

About this task

If your security setup requires you to use different keys for data authentication and FIPS 140-2 authentication, you should create a separate key for each. If that's not the case, you can use the same authentication key for FIPS compliance that you use for data access.

ONTAP creates authentication keys for all nodes in the cluster.

- This command is not supported when Onboard Key Manager is enabled. However, two authentication keys are created automatically when Onboard Key Manager is enabled. The keys can be viewed with the following command:

```
security key-manager key query -key-type NSE-AK
```

- You receive a warning if the configured key management servers are already storing more than 128 authentication keys.
- You can use the `security key-manager key delete` command to delete any unused keys. The `security key-manager key delete` command fails if the given key is currently in use by ONTAP. (You must have privileges greater than “admin” to use this command.)



In a MetroCluster environment, before you delete a key, you must make sure that the key is not in use on the partner cluster. You can use the following commands on the partner cluster to check that the key is not in use:

- `storage encryption disk show -data-key-id key-id`
- `storage encryption disk show -fips-key-id key-id`

Before you begin

You must be a cluster administrator to perform this task.

Steps

1. Create the authentication keys for cluster nodes:

```
security key-manager key create -key-tag passphrase_label -prompt-for-key true|false
```



Setting `prompt-for-key=true` causes the system to prompt the cluster administrator for the passphrase to use when authenticating encrypted drives. Otherwise, the system automatically generates a 32-byte passphrase. The `security key-manager key create` command replaces the `security key-manager create-key` command. For complete command syntax, see the man page.

The following example creates the authentication keys for `cluster1`, automatically generating a 32-byte passphrase:

```
cluster1::> security key-manager key create
Key ID:
000000000000000000002000000000001006268333f870860128fbe17d393e5083b00000000
00000000
```

2. Verify that the authentication keys have been created:

```
security key-manager key query -node node
```



The `security key-manager key query` command replaces the `security key-manager query key` command. For complete command syntax, see the man page. The key ID displayed in the output is an identifier used to refer to the authentication key. It is not the actual authentication key or the data encryption key.

The following example verifies that authentication keys have been created for `cluster1`:

```

cluster1::> security key-manager key query
      Vserver: cluster1
      Key Manager: external
      Node: node1

Key Tag                                Key Type  Restored
-----                                -
node1                                  NSE-AK    yes
      Key ID:
000000000000000000002000000000001000c11b3863f78c2273343d7ec5a67762e00000000
00000000
node1                                  NSE-AK    yes
      Key ID:
000000000000000000002000000000001006f4e2513353a674305872a4c9f3bf79700000000
00000000

      Vserver: cluster1
      Key Manager: external
      Node: node2

Key Tag                                Key Type  Restored
-----                                -
node2                                  NSE-AK    yes
      Key ID:
000000000000000000002000000000001000c11b3863f78c2273343d7ec5a67762e00000000
00000000
node2                                  NSE-AK    yes
      Key ID:
000000000000000000002000000000001006f4e2513353a674305872a4c9f3bf79700000000
00000000

```

Create authentication keys in ONTAP 9.5 and earlier

You can use the `security key-manager create-key` command to create the authentication keys for a node and store them on the configured KMIP servers.

About this task

If your security setup requires you to use different keys for data authentication and FIPS 140-2 authentication, you should create a separate key for each. If that is not the case, you can use the same authentication key for FIPS compliance that you use for data access.

ONTAP creates authentication keys for all nodes in the cluster.

- This command is not supported when onboard key management is enabled.
- You receive a warning if the configured key management servers are already storing more than 128 authentication keys.

You can use the key management server software to delete any unused keys, then run the command again.

Before you begin

You must be a cluster administrator to perform this task.

Steps

1. Create the authentication keys for cluster nodes:

```
security key-manager create-key
```

For complete command syntax, see the man page for the command.



The key ID displayed in the output is an identifier used to refer to the authentication key. It is not the actual authentication key or the data encryption key.

The following example creates the authentication keys for `cluster1`:

```
cluster1::> security key-manager create-key
  (security key-manager create-key)
Verifying requirements...

Node: cluster1-01
Creating authentication key...
Authentication key creation successful.
Key ID: F1CB30AFF1CB30B0010100000000000A68B167F92DD54196297159B5968923C

Node: cluster1-01
Key manager restore operation initialized.
Successfully restored key information.

Node: cluster1-02
Key manager restore operation initialized.
Successfully restored key information.
```

2. Verify that the authentication keys have been created:

```
security key-manager query
```

For complete command syntax, see the man page.

The following example verifies that authentication keys have been created for `cluster1`:


```

cluster1::> security key-manager query

(security key-manager query)

      Node: cluster1-01
    Key Manager: 20.1.1.1
  Server Status: available

Key Tag          Key Type  Restored
-----          -
cluster1-01     NSE-AK   yes
      Key ID:
F1CB30AFF1CB30B0010100000000000A68B167F92DD54196297159B5968923C

      Node: cluster1-02
    Key Manager: 20.1.1.1
  Server Status: available

Key Tag          Key Type  Restored
-----          -
cluster1-02     NSE-AK   yes
      Key ID:
F1CB30AFF1CB30B0010100000000000A68B167F92DD54196297159B5968923C

```

Assign a data authentication key to a FIPS drive or SED (external key management)

You can use the `storage encryption disk modify` command to assign a data authentication key to a FIPS drive or SED. Cluster nodes use this key to lock or unlock encrypted data on the drive.

About this task

A self-encrypting drive is protected from unauthorized access only if its authentication key ID is set to a non-default value. The manufacturer secure ID (MSID), which has key ID 0x0, is the standard default value for SAS drives. For NVMe drives, the standard default value is a null key, represented as a blank key ID. When you assign the key ID to a self-encrypting drive, the system changes its authentication key ID to a non-default value.

This procedure is not disruptive.

Before you begin

You must be a cluster administrator to perform this task.

Steps

1. Assign a data authentication key to a FIPS drive or SED:

```
storage encryption disk modify -disk disk_ID -data-key-id key_ID
```

For complete command syntax, see the man page for the command.



You can use the `security key-manager query -key-type NSE-AK` command to view key IDs.

```
cluster1::> storage encryption disk modify -disk 0.10.* -data-key-id
F1CB30AFF1CB30B0010100000000000A68B167F92DD54196297159B5968923C
```

```
Info: Starting modify on 14 disks.
View the status of the operation by using the
storage encryption disk show-status command.
```

2. Verify that the authentication keys have been assigned:

```
storage encryption disk show
```

For complete command syntax, see the man page.

```
cluster1::> storage encryption disk show
Disk      Mode Data Key ID
-----  ----
-----
0.0.0     data
F1CB30AFF1CB30B0010100000000000A68B167F92DD54196297159B5968923C
0.0.1     data
F1CB30AFF1CB30B0010100000000000A68B167F92DD54196297159B5968923C
[...]
```

Configure onboard key management

Enable onboard key management in ONTAP 9.6 and later

You can use the Onboard Key Manager to authenticate cluster nodes to a FIPS drive or SED. The Onboard Key Manager is a built-in tool that serves authentication keys to nodes from the same storage system as your data. The Onboard Key Manager is FIPS-140-2 level 1 compliant.

You can use the Onboard Key Manager to secure the keys that the cluster uses to access encrypted data. You must enable Onboard Key Manager on each cluster that accesses an encrypted volume or a self-encrypting disk.

About this task

You must run the `security key-manager onboard enable` command each time you add a node to the cluster. In MetroCluster configurations, you must run `security key-manager onboard enable` on the local cluster first, then run `security key-manager onboard sync` on the remote cluster, using the same passphrase on each.

By default, you are not required to enter the key manager passphrase when a node is rebooted. Except in MetroCluster, you can use the `cc-mode-enabled=yes` option to require that users enter the passphrase after a reboot.

When the Onboard Key Manager is enabled in Common Criteria mode (`cc-mode-enabled=yes`), system behavior is changed in the following ways:

- The system monitors for consecutive failed cluster passphrase attempts when operating in Common Criteria mode.

If NetApp Storage Encryption (NSE) is enabled and you fail to enter the correct cluster passphrase at boot, the system cannot authenticate to its drives and automatically reboots. To correct this, you must enter the correct cluster passphrase at the boot prompt. Once booted, the system allows up to 5 consecutive attempts to correctly enter the cluster passphrase in a 24-hour period for any command that requires the cluster passphrase as a parameter. If the limit is reached (for example, you have failed to correctly enter the cluster passphrase 5 times in a row) then you must either wait for the 24-hour timeout period to elapse, or you must reboot the node, in order to reset the limit.

- System image updates use the NetApp RSA-3072 code signing certificate together with SHA-384 code signed digests to check the image integrity instead of the usual NetApp RSA-2048 code signing certificate and SHA-256 code signed digests.

The upgrade command verifies that the image contents have not been altered or corrupted by checking various digital signatures. The image update process proceeds to the next step if validation succeeds; otherwise, the image update fails. See the “cluster image” man page for information concerning system updates.

The Onboard Key Manager stores keys in volatile memory. Volatile memory contents are cleared when the system is rebooted or halted. Under normal operating conditions, volatile memory contents will be cleared within 30s when a system is halted.

Before you begin

- If you are using NSE with an external key management (KMIP) server, you must have deleted the external key manager database.

[Transitioning to onboard key management from external key management](#)

- You must be a cluster administrator to perform this task.
- You must configure the MetroCluster environment before the Onboard Key Manager is configured.

Steps

1. Start the key manager setup command:

```
security key-manager onboard enable -cc-mode-enabled yes|no
```

Set `cc-mode-enabled=yes` to require that users enter the key manager passphrase after a reboot. The `- cc-mode-enabled` option is not supported in MetroCluster configurations. The `security key-manager onboard enable` command replaces the `security key-manager setup` command.

The following example starts the key manager setup command on cluster1 without requiring that the passphrase be entered after every reboot:

```
cluster1::> security key-manager onboard enable
```

```
Enter the cluster-wide passphrase for onboard key management in Vserver  
"cluster1":: <32..256 ASCII characters long text>  
Reenter the cluster-wide passphrase: <32..256 ASCII characters long  
text>
```

2. At the passphrase prompt, enter a passphrase between 32 and 256 characters, or for “cc-mode”, a passphrase between 64 and 256 characters.



If the specified “cc-mode” passphrase is less than 64 characters, there is a five-second delay before the key manager setup operation displays the passphrase prompt again.

3. At the passphrase confirmation prompt, reenter the passphrase.
4. Verify that the authentication keys have been created:

```
security key-manager key query -node node
```



The `security key-manager key query` command replaces the `security key-manager query key` command. For complete command syntax, see the man page.

The following example verifies that authentication keys have been created for cluster1:

```

cluster1::> security key-manager key query
      Vserver: cluster1
      Key Manager: onboard
      Node: node1

Key Tag                                Key Type  Restored
-----                                -
node1                                NSE-AK    yes
      Key ID:
000000000000000000002000000000001000c11b3863f78c2273343d7ec5a67762e00000000
00000000
node1                                NSE-AK    yes
      Key ID:
000000000000000000002000000000001006f4e2513353a674305872a4c9f3bf79700000000
00000000

      Vserver: cluster1
      Key Manager: onboard
      Node: node2

Key Tag                                Key Type  Restored
-----                                -
node1                                NSE-AK    yes
      Key ID:
000000000000000000002000000000001000c11b3863f78c2273343d7ec5a67762e00000000
00000000
node2                                NSE-AK    yes
      Key ID:
000000000000000000002000000000001006f4e2513353a674305872a4c9f3bf79700000000
00000000

```

After you finish

Copy the passphrase to a secure location outside the storage system for future use.

All key management information is automatically backed up to the replicated database (RDB) for the cluster. You should also back up the information manually for use in case of a disaster.

Enable onboard key management in ONTAP 9.5 and earlier

You can use the Onboard Key Manager to authenticate cluster nodes to a FIPS drive or SED. The Onboard Key Manager is a built-in tool that serves authentication keys to nodes from the same storage system as your data. The Onboard Key Manager is FIPS-140-2 level 1 compliant.

You can use the Onboard Key Manager to secure the keys that the cluster uses to access encrypted data. You must enable Onboard Key Manager on each cluster that accesses an encrypted volume or a self-encrypting

disk.

About this task

You must run the `security key-manager setup` command each time you add a node to the cluster.

If you have a MetroCluster configuration, review these guidelines:

- In ONTAP 9.5, you must run `security key-manager setup` on the local cluster and `security key-manager setup -sync-metrocluster-config yes` on the remote cluster, using the same passphrase on each.
- Prior to ONTAP 9.5, you must run `security key-manager setup` on the local cluster, wait approximately 20 seconds, and then run `security key-manager setup` on the remote cluster, using the same passphrase on each.

By default, you are not required to enter the key manager passphrase when a node is rebooted. Beginning with ONTAP 9.4, you can use the `-enable-cc-mode yes` option to require that users enter the passphrase after a reboot.

For NVE, if you set `-enable-cc-mode yes`, volumes you create with the `volume create` and `volume move start` commands are automatically encrypted. For `volume create`, you need not specify `-encrypt true`. For `volume move start`, you need not specify `-encrypt-destination true`.



After a failed passphrase attempt, you must reboot the node again.

Before you begin

- If you are using NSE with an external key management (KMIP) server, you must have deleted the external key manager database.

[Transitioning to onboard key management from external key management](#)

- You must be a cluster administrator to perform this task.
- You must configure the MetroCluster environment before the Onboard Key Manager is configured.

Steps

1. Start the key manager setup:

```
security key-manager setup -enable-cc-mode yes|no
```



Beginning with ONTAP 9.4, you can use the `-enable-cc-mode yes` option to require that users enter the key manager passphrase after a reboot. For NVE, if you set `-enable-cc-mode yes`, volumes you create with the `volume create` and `volume move start` commands are automatically encrypted.

The following example starts setting up the key manager on cluster1 without requiring that the passphrase be entered after every reboot:

After you finish

All key management information is automatically backed up to the replicated database (RDB) for the cluster.

Whenever you configure the Onboard Key Manager passphrase, you should also back up the information manually to a secure location outside the storage system for use in case of a disaster. See [Back up onboard key management information manually](#).

Assign a data authentication key to a FIPS drive or SED (onboard key management)

You can use the `storage encryption disk modify` command to assign a data authentication key to a FIPS drive or SED. Cluster nodes use this key to access data on the drive.

About this task

A self-encrypting drive is protected from unauthorized access only if its authentication key ID is set to a non-default value. The manufacturer secure ID (MSID), which has key ID 0x0, is the standard default value for SAS drives. For NVMe drives, the standard default value is a null key, represented as a blank key ID. When you assign the key ID to a self-encrypting drive, the system changes its authentication key ID to a non-default value.

Before you begin

You must be a cluster administrator to perform this task.

Steps

1. Assign a data authentication key to a FIPS drive or SED:

```
storage encryption disk modify -disk disk_ID -data-key-id key_ID
```

For complete command syntax, see the man page for the command.



You can use the `security key-manager key query -key-type NSE-AK` command to view key IDs.

```
cluster1::> storage encryption disk modify -disk 0.10.* -data-key-id  
0000000000000000020000000000010019215b9738bc7b43d4698c80246db1f4
```

```
Info: Starting modify on 14 disks.  
View the status of the operation by using the  
storage encryption disk show-status command.
```

2. Verify that the authentication keys have been assigned:

```
storage encryption disk show
```

For complete command syntax, see the man page.


```

cluster1::> storage encryption disk show
Disk      Mode Data Key ID
-----  ----
-----
0.0.0    data
0000000000000000000020000000000010019215b9738bc7b43d4698c80246db1f4
0.0.1    data
0000000000000000000020000000000010059851742AF2703FC91369B7DB47C4722
[...]

```

Assign a FIPS 140-2 authentication key to a FIPS drive

You can use the `storage encryption disk modify` command with the `-fips-key-id` option to assign a FIPS 140-2 authentication key to a FIPS drive. Cluster nodes use this key for drive operations other than data access, such as preventing denial-of-service attacks on the drive.

About this task

Your security setup may require you to use different keys for data authentication and FIPS 140-2 authentication. If that is not the case, you can use the same authentication key for FIPS compliance that you use for data access.

This procedure is not disruptive.

Before you begin

The drive firmware must support FIPS 140-2 compliance. The [NetApp Interoperability Matrix Tool](#) contains information about supported drive firmware versions.

Steps

1. You must first ensure you have assigned a data authentication key. This can be done with using an [external key manager](#) or an [onboard key manager](#). Verify the key is assigned with the command `storage encryption disk show`.
2. Assign a FIPS 140-2 authentication key to SEDs:

```

storage encryption disk modify -disk disk_id -fips-key-id
fips_authentication_key_id

```

You can use the `security key-manager query` command to view key IDs.

```

cluster1::> storage encryption disk modify -disk 2.10.* -fips-key-id
6A1E21D8000000000100000000000005A1FB4EE8F62FD6D8AE6754C9019F35A

```

```

Info: Starting modify on 14 disks.
      View the status of the operation by using the
      storage encryption disk show-status command.

```

3. Verify that the authentication key has been assigned:

```
storage encryption disk show -fips
```

For complete command syntax, see the man page.

```
cluster1::> storage encryption disk show -fips
Disk      Mode FIPS-Compliance Key ID
-----  ----
-----
2.10.0    full
6A1E21D8000000000100000000000005A1FB4EE8F62FD6D8AE6754C9019F35A
2.10.1    full
6A1E21D8000000000100000000000005A1FB4EE8F62FD6D8AE6754C9019F35A
[...]
```

Enable cluster-wide FIPS-compliant mode for KMIP server connections

You can use the `security config modify` command with the `-is-fips-enabled` option to enable cluster-wide FIPS-compliant mode for data in flight. Doing so forces the cluster to use OpenSSL in FIPS mode when connecting to KMIP servers.

About this task

When you enable cluster-wide FIPS-compliant mode, the cluster will automatically use only TLS1.2 and FIPS-validated cipher suites. Cluster-wide FIPS-compliant mode is disabled by default.

You must reboot cluster nodes manually after modifying the cluster-wide security configuration.

Before you begin

- The storage controller must be configured in FIPS-compliant mode.
- All KMIP servers must support TLSv1.2. The system requires TLSv1.2 to complete the connection to the KMIP server when cluster-wide FIPS-compliant mode is enabled.

Steps

1. Set the privilege level to advanced:

```
set -privilege advanced
```

2. Verify that TLSv1.2 is supported:

```
security config show -supported-protocols
```

For complete command syntax, see the man page.

```

cluster1::> security config show
          Cluster                                     Cluster
Security
Interface FIPS Mode  Supported Protocols  Supported Ciphers  Config
Ready
-----
-----
SSL        false      TLSv1.2, TLSv1.1, TLSv1  ALL:!LOW:
                                !aNULL:!EXP:
                                !eNULL

```

3. Enable cluster-wide FIPS-compliant mode:

```
security config modify -is-fips-enabled true -interface SSL
```

For complete command syntax, see the man page.

4. Reboot cluster nodes manually.

5. Verify that cluster-wide FIPS-compliant mode is enabled:

```
security config show
```

```

cluster1::> security config show
          Cluster                                     Cluster
Security
Interface FIPS Mode  Supported Protocols  Supported Ciphers  Config
Ready
-----
-----
SSL        true       TLSv1.2, TLSv1.1      ALL:!LOW:
                                !aNULL:!EXP:
                                !eNULL:!RC4

```

Manage NetApp encryption

Unencrypt volume data

You can use the `volume move start` command to move and unencrypt volume data.

Before you begin

You must be a cluster administrator to perform this task. Alternately, you can be an SVM administrator to whom the cluster administrator has delegated authority. For more information, see [Delegate authority to run the volume move command](#).

Steps

1. Move an existing encrypted volume and unencrypt the data on the volume:

```
volume move start -vserver SVM_name -volume volume_name -destination-aggregate
aggregate_name -encrypt-destination false
```

For complete command syntax, see the man page for the command.

The following command moves an existing volume named `vol1` to the destination aggregate `aggr3` and unencrypts the data on the volume:

```
cluster1::> volume move start -vserver vs1 -volume vol1 -destination
-aggregate aggr3 -encrypt-destination false
```

The system deletes the encryption key for the volume. The data on the volume is unencrypted.

2. Verify that the volume is disabled for encryption:

```
volume show -encryption
```

For complete command syntax, see the man page for the command.

The following command displays whether volumes on `cluster1` are encrypted:

```
cluster1::> volume show -encryption
```

Vserver	Volume	Aggregate	State	Encryption State
vs1	vol1	aggr1	online	none

Move an encrypted volume

You can use the `volume move start` command to move an encrypted volume. The moved volume can reside on the same aggregate or a different aggregate.

About this task

The move will fail if the destination node or destination volume does not support volume encryption.

The `-encrypt-destination` option for `volume move start` defaults to `true` for encrypted volumes. The requirement to specify you do not want the destination volume encrypted ensures that you do not inadvertently unencrypt the data on the volume.

Before you begin

You must be a cluster administrator to perform this task. Alternately, you can be an SVM administrator to whom the cluster administrator has delegated authority. For more information, see [delegate authority to run the volume move command](#).

Steps

1. Move an existing encrypted volume and leave the data on the volume encrypted:

```
volume move start -vserver SVM_name -volume volume_name -destination-aggregate aggregate_name
```

For complete command syntax, see the man page for the command.

The following command moves an existing volume named `vol1` to the destination aggregate `aggr3` and leaves the data on the volume encrypted:

```
cluster1::> volume move start -vserver vs1 -volume vol1 -destination
-aggregate aggr3
```

2. Verify that the volume is enabled for encryption:

```
volume show -is-encrypted true
```

For complete command syntax, see the man page for the command.

The following command displays the encrypted volumes on `cluster1`:

```
cluster1::> volume show -is-encrypted true
```

Vserver	Volume	Aggregate	State	Type	Size	Available	Used
vs1	vol1	aggr3	online	RW	200GB	160.0GB	20%

Delegate authority to run the volume move command

You can use the `volume move` command to encrypt an existing volume, move an encrypted volume, or unencrypt a volume. Cluster administrators can run `volume move` command themselves, or they can delegate the authority to run the command to SVM administrators.

About this task

By default, SVM administrators are assigned the `vsadmin` role, which does not include the authority to move volumes. You must assign the `vsadmin-volume` role to SVM administrators to enable them to run the `volume move` command.

Step

1. Delegate authority to run the `volume move` command:

```
security login modify -vserver SVM_name -user-or-group-name user_or_group_name
-application application -authmethod authentication_method -role vsadmin-
volume
```

For complete command syntax, see the man page for the command.

The following command grants the SVM administrator authority to run the `volume move` command.

```
cluster1::>security login modify -vserver engData -user-or-group-name
SVM-admin -application ssh -authmethod domain -role vsadmin-volume
```

Change the encryption key for a volume with the volume encryption rekey start command

It is a security best practice to change the encryption key for a volume periodically. Beginning with ONTAP 9.3, you can use the `volume encryption rekey start` command to change the encryption key.

About this task

Once you start a rekey operation, it must complete. There is no returning to the old key. If you encounter a performance issue during the operation, you can run the `volume encryption rekey pause` command to pause the operation, and the `volume encryption rekey resume` command to resume the operation.

Until the rekey operation finishes, the volume will have two keys. New writes and their corresponding reads will use the new key. Otherwise, reads will use the old key.



You cannot use `volume encryption rekey start` to rekey a SnapLock volume.

Steps

1. Change an encryption key:

```
volume encryption rekey start -vserver SVM_name -volume volume_name
```

The following command changes the encryption key for `vol1` on `SVMvs1`:

```
cluster1::> volume encryption rekey start -vserver vs1 -volume vol1
```

2. Verify the status of the rekey operation:

```
volume encryption rekey show
```

For complete command syntax, see the man page for the command.

The following command displays the status of the rekey operation:

```
cluster1::> volume encryption rekey show
```

Vserver	Volume	Start Time	Status
vs1	vol1	9/18/2017 17:51:41	Phase 2 of 2 is in progress.

3. When the rekey operation is complete, verify that the volume is enabled for encryption:

```
volume show -is-encrypted true
```

For complete command syntax, see the man page for the command.

The following command displays the encrypted volumes on `cluster1`:

```
cluster1::> volume show -is-encrypted true
```

Vserver	Volume	Aggregate	State	Type	Size	Available	Used
vs1	vol1	aggr2	online	RW	200GB	160.0GB	20%

Change the encryption key for a volume with the volume move start command

It is a security best practice to change the encryption key for a volume periodically. You can use the `volume move start` command to change the encryption key. You must use `volume move start` in ONTAP 9.2 and earlier. The moved volume can reside on the same aggregate or a different aggregate.

About this task

You cannot use `volume move start` to rekey a SnapLock or FlexGroup volume.

Before you begin

You must be a cluster administrator to perform this task. Alternately, you can be an SVM administrator to whom the cluster administrator has delegated authority. For more information, see [delegate authority to run the volume move command](#).

Steps

1. Move an existing volume and change the encryption key:

```
volume move start -vserver SVM_name -volume volume_name -destination-aggregate aggregate_name -generate-destination-key true
```

For complete command syntax, see the man page for the command.

The following command moves an existing volume named **vol1** to the destination aggregate **aggr2** and changes the encryption key:

```
cluster1::> volume move start -vserver vs1 -volume vol1 -destination -aggregate aggr2 -generate-destination-key true
```

A new encryption key is created for the volume. The data on the volume remains encrypted.

2. Verify that the volume is enabled for encryption:

```
volume show -is-encrypted true
```

For complete command syntax, see the man page for the command.

The following command displays the encrypted volumes on `cluster1`:

```
cluster1::> volume show -is-encrypted true
```

Vserver	Volume	Aggregate	State	Type	Size	Available	Used
-----	-----	-----	-----	-----	-----	-----	-----
vs1	vol1	aggr2	online	RW	200GB	160.0GB	20%

Rotate authentication keys for NetApp Storage Encryption

You can rotate authentication keys when using NetApp Storage Encryption (NSE).

About this task

Rotating authentication keys in an NSE environment is supported if you are using External Key Manager (KMIP).



Rotating authentication keys in an NSE environment is not supported for Onboard Key Manager (OKM).

Steps

1. Use the `security key-manager create-key` command to generate new authentication keys.

You need to generate new authentication keys before you can change the authentication keys.

2. Use the `storage encryption disk modify -disk * -data-key-id` command to change the authentication keys.

Delete an encrypted volume

You can use the `volume delete` command to delete an encrypted volume.

Before you begin

- You must be a cluster administrator to perform this task. Alternately, you can be an SVM administrator to whom the cluster administrator has delegated authority. For more information, see [delegate authority to run the volume move command](#).
- The volume must be offline.

Step

1. Delete an encrypted volume:

```
volume delete -vserver SVM_name -volume volume_name
```

For complete command syntax, see the man page for the command.

The following command deletes an encrypted volume named `vol1`:

```
cluster1::> volume delete -vserver vs1 -volume vol1
```

Enter `yes` when you are prompted to confirm deletion.

The system deletes the encryption key for the volume after 24 hours.

Use `volume delete` with the `-force true` option to delete a volume and destroy the corresponding encryption key immediately. This command requires advanced privileges. For more information, see the man page.

After you finish

You can use the `volume recovery-queue` command to recover a deleted volume during the retention period after issuing the `volume delete` command:

```
volume recovery-queue SVM_name -volume volume_name
```

[How to use the Volume Recovery feature](#)

Securely purge data on an encrypted volume

Securely purge data on an encrypted volume overview

Beginning with ONTAP 9.4, you can use secure purge to non-disruptively scrub data on NVE-enabled volumes. Scrubbing data on an encrypted volume ensures that it cannot be recovered from the physical media, for example, in cases of “spillage,” where data traces may have been left behind when blocks were overwritten, or for securely deleting a vacating tenant’s data.

Secure purge works only for previously deleted files on NVE-enabled volumes. You cannot scrub an unencrypted volume. You must use KMIP servers to serve keys, not the onboard key manager.

Considerations for using secure purge

- Volumes created in an aggregate enabled for NetApp Aggregate Encryption (NAE) do not support secure purge.
- Secure purge works only for previously deleted files on NVE-enabled volumes.
- You cannot scrub an unencrypted volume.
- You must use KMIP servers to serve keys, not the onboard key manager.

Secure purge functions differently depending upon your version of ONTAP.

ONTAP 9.8 and later

- Secure purge is supported by MetroCluster and FlexGroup.
- If the volume being purged is the source of a SnapMirror relationship, you do not have to break the SnapMirror relationship to perform a secure purge.
- The re-encryption method is different for volumes using SnapMirror data protection versus volumes not using SnapMirror data protection (DP) or those using SnapMirror extended data protection..
 - By default, volumes using SnapMirror data protection (DP) mode re-encrypt data using the volume move re-encryption method.
 - By default, volumes not using SnapMirror data protection or volumes using SnapMirror extended data protection (XDP) mode use the in-place re-encryption method.
 - These defaults can be changed using the `secure purge re-encryption-method [volume-move|in-place-rekey]` command.
- By default, all Snapshot copies in FlexVol volumes are automatically deleted during the secure purge operation. By default, Snapshots in FlexGroup volumes and volumes using SnapMirror data protection are not automatically deleted during the secure purge operation. These defaults can be changed using the `secure purge delete-all-snapshots [true|false]` command.

ONTAP 9.7 and earlier:

- Secure purge does not support the following:
 - FlexClone
 - SnapVault
 - FabricPool
- If the volume being purged is the source of a SnapMirror relationship, you must break the SnapMirror relationship before you can purge the volume.

If there are busy Snapshot copies in the volume, you must release the Snapshot copies before you can purge the volume. For example, you may need to split a FlexClone volume from its parent.

- Successfully invoking the secure-purge feature triggers a volume move that re-encrypts the remaining, unpurged data with a new key.

The moved volume remains on the current aggregate. The old key is automatically destroyed, ensuring that purged data cannot be recovered from the storage media.

Securely purge data on an encrypted volume without a SnapMirror relationship

Beginning with ONTAP 9.4, you can use secure-purge to non-disruptively “scrub” data on NVE-enabled volumes.

About this task

Secure-purge may take from several minutes to many hours to complete, depending on the amount of data in the deleted files. You can use the `volume encryption secure-purge show` command to view the status of the operation. You can use the `volume encryption secure-purge abort` command to terminate the operation.



In order to do a secure purge on a SAN host, you must delete the entire LUN containing the files you want to purge, or you must be able to punch holes in the LUN for the blocks that belong to the files you want to purge. If you cannot delete the LUN or your host operating system does not support punching holes in the LUN, you cannot perform a secure purge.

Before you begin

- You must be a cluster administrator to perform this task.
- Advanced privileges are required for this task.

Steps

1. Delete the files or the LUN you want to securely purge.
 - On a NAS client, delete the files you want to securely purge.
 - On a SAN host, delete the LUN you want to securely purge or punch holes in the LUN for the blocks that belong to the files you want to purge.
2. On the storage system, change to advanced privilege level:

```
set -privilege advanced
```

3. If the files you want to securely purge are in snapshots, delete the snapshots:

```
snapshot delete -vserver SVM_name -volume volume_name -snapshot
```

4. Securely purge the deleted files:

```
volume encryption secure-purge start -vserver SVM_name -volume volume_name
```

The following command securely purges the deleted files on vol1 on SVMvs1:

```
cluster1::> volume encryption secure-purge start -vserver vs1 -volume  
vol1
```

5. Verify the status of the secure-purge operation:

```
volume encryption secure-purge show
```

Securely purge data on an encrypted volume with an SnapMirror asynchronous relationship

Beginning with ONTAP 9.8, you can use a secure purge to non-disruptively “scrub” data on NVE-enabled volumes with an SnapMirror asynchronous relationship.

Before you begin

- You must be a cluster administrator to perform this task.
- Advanced privileges are required for this task.

About this task

Secure-purge may take from several minutes to many hours to complete, depending on the amount of data in the deleted files. You can use the `volume encryption secure-purge show` command to view the status

of the operation. You can use the `volume encryption secure-purge abort` command to terminate the operation.



In order to do a secure purge on a SAN host, you must delete the entire LUN containing the files you want to purge, or you must be able to punch holes in the LUN for the blocks that belong to the files you want to purge. If you cannot delete the LUN or your host operating system does not support punching holes in the LUN, you cannot perform a secure purge.

Steps

1. On the storage system, switch to the advanced privilege level:

```
set -privilege advanced
```

2. Delete the files or the LUN you want to securely purge.
 - On a NAS client, delete the files you want to securely purge.
 - On a SAN host, delete the LUN you want to securely purge or punch holes in the LUN for the blocks that belong to the files you want to purge.
3. Prepare the destination volume in the Asynchronous relationship to be securely purged:

```
volume encryption secure-purge start -vserver SVM_name -volume volume_name  
-prepare true
```

Repeat this step on each volume in your SnapMirror asynchronous relationship.

4. If the files you want to securely purge are in Snapshot copies, delete the Snapshot copies:

```
snapshot delete -vserver SVM_name -volume volume_name -snapshot
```

5. If the files you want to securely purge are in the base Snapshot copies, do the following:

- a. Create a Snapshot copy on the destination volume in the SnapMirror asynchronous relationship:

```
volume snapshot create -snapshot snapshot_name -vserver SVM_name -volume  
volume_name
```

- b. Update SnapMirror to move the base Snapshot copy forward:

```
snapmirror update -source-snapshot snapshot_name -destination-path  
destination_path
```

Repeat this step for each volume in the SnapMirror asynchronous relationship.

- c. Repeat steps (a) and (b) equal to the number of base Snapshot copies plus one.

For example, if you have two base Snapshot copies, you should repeat steps (a) and (b) three times.

- d. Verify that the base Snapshot copy is present:

```
snapshot show -vserver SVM_name -volume volume_name
```

- e. Delete the base Snapshot copy:

```
snapshot delete -vserver svm_name -volume volume_name -snapshot snapshot
```

6. Securely purge the deleted files:

```
volume encryption secure-purge start -vserver svm_name -volume volume_name
```

Repeat this step on each volume in the SnapMirror asynchronous relationship.

The following command securely purges the deleted files on “vol1” on SVM “vs1”:

```
cluster1::> volume encryption secure-purge start -vserver vs1 -volume  
vol1
```

7. Verify the status of the secure purge operation:

```
volume encryption secure-purge show
```

Scrub data on an encrypted volume with a SnapMirror synchronous relationship

Beginning with ONTAP 9.8, you can use a secure purge to non-disruptively "scrub" data on NVE-enabled volumes with a SnapMirror synchronous relationship.

About this task

A secure purge might take from several minutes to many hours to complete, depending on the amount of data in the deleted files. You can use the `volume encryption secure-purge show` command to view the status of the operation. You can use the `volume encryption secure-purge abort` command to terminate the operation.



In order to do a secure purge on a SAN host, you must delete the entire LUN containing the files you want to purge, or you must be able to punch holes in the LUN for the blocks that belong to the files you want to purge. If you cannot delete the LUN or your host operating system does not support punching holes in the LUN, you cannot perform a secure purge.

Before you begin

- You must be a cluster administrator to perform this task.
- Advanced privileges are required for this task.

Steps

1. On the storage system, change to advanced privilege level:

```
set -privilege advanced
```

2. Delete the files or the LUN you want to securely purge.
 - On a NAS client, delete the files you want to securely purge.
 - On a SAN host, delete the LUN you want to securely purge or punch holes in the LUN for the blocks that belong to the files you want to purge.
3. Prepare the destination volume in the Asynchronous relationship to be securely purged:

```
volume encryption secure-purge start -vserver <SVM_name> -volume <volume_name>  
-prepare true
```

Repeat this step for the other volume in your SnapMirror synchronous relationship.

4. If the files you want to securely purge are in Snapshot copies, delete the Snapshot copies:

```
snapshot delete -vserver <SVM_name> -volume <volume_name> -snapshot <snapshot>
```

5. If the secure purge file is in the base or common Snapshot copies, update the SnapMirror to move the common Snapshot copy forward:

```
snapmirror update -source-snapshot <snapshot_name> -destination-path  
<destination_path>
```

There are two common Snapshot copies, so this command must be issued twice.

6. If the secure purge file is in the application-consistent Snapshot copy, delete the Snapshot copy on both volumes in the SnapMirror synchronous relationship:

```
snapshot delete -vserver <SVM_name> -volume <volume_name> -snapshot <snapshot>
```

Perform this step on both volumes.

7. Securely purge the deleted files:

```
volume encryption secure-purge start -vserver <SVM_name> -volume <volume_name>
```

Repeat this step on each volume in the SnapMirror synchronous relationship.

The following command securely purges the deleted files on “vol1” on SVM “vs1”.

```
cluster1::> volume encryption secure-purge start -vserver vs1 -volume  
vol1
```

8. Verify the status of the secure purge operation:

```
volume encryption secure-purge show
```

Change the onboard key management passphrase

It is a security best practice to change the onboard key management passphrase periodically. You should copy the new onboard key management passphrase to a secure location outside the storage system for future use.

Before you begin

- You must be a cluster or SVM administrator to perform this task.
- Advanced privileges are required for this task.

Steps

1. Change to advanced privilege level:

```
set -privilege advanced
```

2. Change the onboard key management passphrase:

For this ONTAP version...	Use this command...
ONTAP 9.6 and later	<code>security key-manager onboard update-passphrase</code>
ONTAP 9.5 and earlier	<code>security key-manager update-passphrase</code>

For complete command syntax, see the man pages.

The following ONTAP 9.6 command lets you change the onboard key management passphrase for `cluster1`:

```
cluster1::> security key-manager onboard update-passphrase
Warning: This command will reconfigure the cluster passphrase for
onboard key management for Vserver "cluster1".
Do you want to continue? {y|n}: y
Enter current passphrase:
Enter new passphrase:
```

3. Enter `y` at the prompt to change the onboard key management passphrase.
4. Enter the current passphrase at the current passphrase prompt.
5. At the new passphrase prompt, enter a passphrase between 32 and 256 characters, or for “cc-mode”, a passphrase between 64 and 256 characters.

If the specified “cc-mode” passphrase is less than 64 characters, there is a five-second delay before the key manager setup operation displays the passphrase prompt again.

6. At the passphrase confirmation prompt, reenter the passphrase.

After you finish

In a MetroCluster environment, you must update the passphrase on the partner cluster:

- In ONTAP 9.5 and earlier, you must run `security key-manager update-passphrase` with the same passphrase on the partner cluster.
- In ONTAP 9.6 and later, you are prompted to run `security key-manager onboard sync` with the same passphrase on the partner cluster.

You should copy the onboard key management passphrase to a secure location outside the storage system for future use.

You should back up key management information manually whenever you change the onboard key management passphrase.

[Backing up onboard key management information manually](#)

Back up onboard key management information manually

You should copy onboard key management information to a secure location outside the storage system whenever you configure the Onboard Key Manager passphrase.

What you'll need

- You must be a cluster administrator to perform this task.
- Advanced privileges are required for this task.

About this task

All key management information is automatically backed up to the replicated database (RDB) for the cluster. You should also back up key management information manually for use in case of a disaster.

Steps

1. Change to advanced privilege level:

```
set -privilege advanced
```

2. Display the key management backup information for the cluster:

For this ONTAP version...	Use this command...
ONTAP 9.6 and later	<code>security key-manager onboard show-backup</code>
ONTAP 9.5 and earlier	<code>security key-manager backup show</code>

For complete command syntax, see the man pages.

+

The following 9.6 command displays the key management backup information for `cluster1`:

+



If you are using NSE on a system with a Flash Cache module, you should also enable NVE or NAE. NSE does not encrypt data that resides on the Flash Cache module.

ONTAP 9.6 and later



If you are running ONTAP 9.8 or later and your root volume is encrypted, follow the procedure for [ONTAP 9.8 or later with encrypted root volume](#).

1. Verify that the key needs to be restored:
`security key-manager key query -node node`
2. Restore the key:
`security key-manager onboard sync`

For complete command syntax, see the man pages.

The following ONTAP 9.6 command synchronize the keys in the onboard key hierarchy:

```
cluster1::> security key-manager onboard sync
```

```
Enter the cluster-wide passphrase for onboard key management in Vserver  
"cluster1":: <32..256 ASCII characters long text>
```

3. At the passphrase prompt, enter the onboard key management passphrase for the cluster.

ONTAP 9.8 or later with encrypted root volume

If you are running ONTAP 9.8 and later, and your root volume is encrypted, you must set an onboard key management recovery passphrase with the boot menu. This process is also necessary if you do a boot media replacement.

1. Boot the node to the boot menu and select option (10) Set onboard key management recovery secrets.
2. Enter `y` to use this option.
3. At the prompt, enter the onboard key management passphrase for the cluster.
4. At the prompt, enter the backup key data.

The node returns to the boot menu.

5. From the boot menu, select option (1) Normal Boot.

ONTAP 9.5 and earlier

1. Verify that the key needs to be restored:
`security key-manager key show`
2. If you are running ONTAP 9.8 and later, and your root volume is encrypted, complete these steps:

If you are running ONTAP 9.6 or 9.7, or if you are running ONTAP 9.8 or later and your root volume is not encrypted, skip this step.

3. Restore the key:

```
security key-manager setup -node node
```

For complete command syntax, see the man pages.

4. At the passphrase prompt, enter the onboard key management passphrase for the cluster.

Restore external key management encryption keys

You can manually restore external key management encryption keys and push them to a different node. You might want to do this if you are restarting a node that was down temporarily when you created the keys for the cluster.

About this task

In ONTAP 9.6 and later, you can use the `security key-manager key query -node node_name` command to verify if your key needs to be restored.

In ONTAP 9.5 and earlier, you can use the `security key-manager key show` command to verify if your key needs to be restored.



If you are using NSE on a system with a Flash Cache module, you should also enable NVE or NAE. NSE does not encrypt data that resides on the Flash Cache module.

Before you begin

You must be a cluster or SVM administrator to perform this task.

Steps

1. If you are running ONTAP 9.8 or later and your root volume is encrypted, do the following:

If you are running ONTAP 9.7 or earlier, or if you are running ONTAP 9.8 or later and your root volume is not encrypted, skip this step.

a. Set the bootargs:

```
setenv kmip.init.ipaddr <ip-address>
```

```
setenv kmip.init.netmask <netmask>
```

```
setenv kmip.init.gateway <gateway>
```

```
setenv kmip.init.interface e0M
```

```
boot_ontap
```

b. Boot the node to the boot menu and select option (11) Configure node for external key management.

c. Follow prompts to enter management certificate.

After all management certificate information is entered, the system returns to the boot menu.

d. From the boot menu, select option (1) Normal Boot.

2. Restore the key:

For this ONTAP version...	Use this command...
ONTAP 9.6 and later	<code>security key-manager external restore -vserver SVM -node node -key-server host_name IP_address:port -key-id key_id -key-tag key_tag</code>
ONTAP 9.5 and earlier	<code>security key-manager restore -node node -address IP_address -key-id key_id -key-tag key_tag</code>



`node` defaults to all nodes. For complete command syntax, see the man pages. This command is not supported when onboard key management is enabled.

The following ONTAP 9.6 command restores external key management authentication keys to all nodes in `cluster1`:

```
cluster1::> security key-manager external restore
```

Replace SSL certificates

All SSL certificates have an expiration date. You must update your certificates before they expire to prevent loss of access to authentication keys.

Before you begin

- You must have obtained the replacement public certificate and private key for the cluster (KMIP client certificate).
- You must have obtained the replacement public certificate for the KMIP server (KMIP server-ca certificate).
- You must be a cluster or SVM administrator to perform this task.
- If you are replacing the KMIP SSL certificates in a MetroCluster environment, you must install the same replacement KMIP SSL certificate on both clusters.



You can install the replacement client and server certificates on the KMIP server before or after installing the certificates on the cluster.

Steps

1. Install the new KMIP server-ca certificate:

```
security certificate install -type server-ca -vserver <>
```

2. Install the new KMIP client certificate:

```
security certificate install -type client -vserver <>
```

3. Update the key manager configuration to use the newly installed certificates:

```
security key-manager external modify -vserver <> -client-cert <> -server-ca
```

-certs <>

If you are running ONTAP 9.6 or later in a MetroCluster environment, and you want to modify the key manager configuration on the admin SVM, you must run the command on both clusters in the configuration.



Updating the key manager configuration to use the newly installed certificates will return an error if the public/private keys of the new client certificate are different from the keys previously installed. See the Knowledge Base article [The new client certificate public or private keys are different from the existing client certificate](#) for instructions on how to override this error.

Replace a FIPS drive or SED

You can replace a FIPS drive or SED the same way you replace an ordinary disk. Make sure to assign new data authentication keys to the replacement drive. For a FIPS drive, you may also want to assign a new FIPS 140-2 authentication key.



If an HA pair is using [encrypting SAS or NVMe drives \(SED, NSE, FIPS\)](#), you must follow the instructions in the topic [Returning a FIPS drive or SED to unprotected mode](#) for all drives within the HA pair prior to initializing the system (boot options 4 or 9). Failure to do this may result in future data loss if the drives are repurposed.

Before you begin

- You must know the key ID for the authentication key used by the drive.
- You must be a cluster administrator to perform this task.

Steps

1. Ensure that the disk has been marked as failed:

```
storage disk show -broken
```

For complete command syntax, see the man page.

```
cluster1::> storage disk show -broken
Original Owner: cluster1-01
Checksum Compatibility: block

Physical
Disk      Outage Reason HA Shelf Bay Chan  Pool  Type  RPM  Usable
Size
-----
0.0.0    admin   failed  0b    1    0    A    Pool0 FCAL  10000 132.8GB
133.9GB
0.0.7    admin   removed 0b    2    6    A    Pool1 FCAL  10000 132.8GB
134.2GB
[...]
```

2. Remove the failed disk and replace it with a new FIPS drive or SED, following the instructions in the hardware guide for your disk shelf model.
3. Assign ownership of the newly replaced disk:

```
storage disk assign -disk disk_name -owner node
```

For complete command syntax, see the man page.

```
cluster1::> storage disk assign -disk 2.1.1 -owner cluster1-01
```

4. Confirm that the new disk has been assigned:

```
storage encryption disk show
```

For complete command syntax, see the man page.

```
cluster1::> storage encryption disk show
Disk      Mode Data Key ID
-----  ----
-----
0.0.0     data
F1CB30AFF1CB30B0010100000000000A68B167F92DD54196297159B5968923C
0.0.1     data
F1CB30AFF1CB30B0010100000000000A68B167F92DD54196297159B5968923C
1.10.0    data
F1CB30AFF1CB30B0010100000000000CF0EFD81EA9F6324EA97B369351C56AC
1.10.1    data
F1CB30AFF1CB30B0010100000000000CF0EFD81EA9F6324EA97B369351C56AC
2.1.1     open 0x0
[...]
```

5. Assign the data authentication keys to the FIPS drive or SED.

[Assigning a data authentication key to a FIPS drive or SED \(external key management\)](#)

6. If necessary, assign a FIPS 140-2 authentication key to the FIPS drive.

[Assigning a FIPS 140-2 authentication key to a FIPS drive](#)

Make data on a FIPS drive or SED inaccessible

Make data on a FIPS drive or SED inaccessible overview

If you want to make data on a FIPS drive or SED permanently inaccessible, but keep the drive's unused space available for new data, you can sanitize the disk. If you want to make data permanently inaccessible and you do not need to reuse the drive, you can destroy it.

- Disk sanitization

When you sanitize a self-encrypting drive, the system changes the disk encryption key to a new random value, resets the power-on lock state to false, and sets the key ID to a default value, either the manufacturer secure ID 0x0 (SAS drives) or a null key (NVMe drives). Doing so renders the data on the disk inaccessible and impossible to retrieve. You can reuse sanitized disks as non-zeroed spare disks.

- Disk destroy

When you destroy a FIPS drive or SED, the system sets the disk encryption key to an unknown random value and locks the disk irreversibly. Doing so renders the disk permanently unusable and the data on it permanently inaccessible.

You can sanitize or destroy individual self-encrypting drives, or all the self-encrypting drives for a node.

Sanitize a FIPS drive or SED

If you want to make data on a FIPS drive or SED permanently inaccessible, and use the drive for new data, you can use the `storage encryption disk sanitize` command to sanitize the drive.

About this task

When you sanitize a self-encrypting drive, the system changes the disk encryption key to a new random value, resets the power-on lock state to false, and sets the key ID to a default value, either the manufacturer secure ID 0x0 (SAS drives) or a null key (NVMe drives). Doing so renders the data on the disk inaccessible and impossible to retrieve. You can reuse sanitized disks as non-zeroed spare disks.

Before you begin

You must be a cluster administrator to perform this task.

Steps

1. Migrate any data that needs to be preserved to an aggregate on another disk.
2. Delete the aggregate on the FIPS drive or SED to be sanitized:

```
storage aggregate delete -aggregate aggregate_name
```

For complete command syntax, see the man page.

```
cluster1::> storage aggregate delete -aggregate aggr1
```

3. Identify the disk ID for the FIPS drive or SED to be sanitized:

```
storage encryption disk show -fields data-key-id,fips-key-id,owner
```

For complete command syntax, see the man page.

```

cluster1::> storage encryption disk show
Disk      Mode Data Key ID
-----  ----
-----
0.0.0    data
F1CB30AFF1CB30B0010100000000000A68B167F92DD54196297159B5968923C
0.0.1    data
F1CB30AFF1CB30B0010100000000000A68B167F92DD54196297159B5968923C
1.10.2   data
F1CB30AFF1CB30B0010100000000000CF0EFD81EA9F6324EA97B369351C56AC
[...]

```

4. If a FIPS drive is running in FIPS-compliance mode, set the FIPS authentication key ID for the node back to the default MSID 0x0:

```
storage encryption disk modify -disk disk_id -fips-key-id 0x0
```

You can use the `security key-manager query` command to view key IDs.

```

cluster1::> storage encryption disk modify -disk 1.10.2 -fips-key-id 0x0

Info: Starting modify on 1 disk.
      View the status of the operation by using the
      storage encryption disk show-status command.

```

5. Sanitize the drive:

```
storage encryption disk sanitize -disk disk_id
```

You can use this command to sanitize hot spare or broken disks only. To sanitize all disks regardless of type, use the `-force-all-state` option. For complete command syntax, see the man page.



ONTAP will prompt you to enter a confirmation phrase before continuing. Enter the phrase exactly as shown on the screen.

```

cluster1::> storage encryption disk sanitize -disk 1.10.2

Warning: This operation will cryptographically sanitize 1 spare or
broken self-encrypting disk on 1 node.
      To continue, enter sanitize disk: sanitize disk

Info: Starting sanitize on 1 disk.
      View the status of the operation using the
      storage encryption disk show-status command.

```


6. Unfail the sanitized disk:

```
storage disk unfail -spare true -disk disk_id
```

7. Check whether the disk has an owner:

```
storage disk show -disk disk_id
```

If the disk does not have an owner, assign one.

```
storage disk assign -owner node -disk disk_id
```

8. Enter the nodeshell for the node that owns the disks you want to sanitize:

```
system node run -node node_name
```

Run the `disk sanitize release` command.

9. Exit the nodeshell. Unfail the disk again:

```
storage disk unfail -spare true -disk disk_id
```

10. Verify that the disk is now a spare and ready to be reused in an aggregate:

```
storage disk show -disk disk_id
```

Destroy a FIPS drive or SED

If you want to make data on a FIPS drive or SED permanently inaccessible and you do not need to reuse the drive, you can use the `storage encryption disk destroy` command to destroy the disk.

About this task

When you destroy a FIPS drive or SED, the system sets the disk encryption key to an unknown random value and locks the drive irreversibly. Doing so renders the disk virtually unusable and the data on it permanently inaccessible. However, you can reset the disk to its factory-configured settings using the physical secure ID (PSID) printed on the disk's label. For more information, see [Returning a FIPS drive or SED to service when authentication keys are lost](#).



You should not destroy a FIPS drive or SED unless you have the Non-Returnable Disk Plus service (NRD Plus). Destroying a disk voids its warranty.

Before you begin

You must be a cluster administrator to perform this task.

Steps

1. Migrate any data that needs to be preserved to an aggregate on another different disk.
2. Delete the aggregate on the FIPS drive or SED to be destroyed:

```
storage aggregate delete -aggregate aggregate_name
```

For complete command syntax, see the man page.

```
cluster1::> storage aggregate delete -aggregate aggr1
```

3. Identify the disk ID for the FIPS drive or SED to be destroyed:

```
storage encryption disk show
```

For complete command syntax, see the man page.

```
cluster1::> storage encryption disk show
Disk      Mode Data Key ID
-----  ----
-----
0.0.0     data
F1CB30AFF1CB30B0010100000000000A68B167F92DD54196297159B5968923C
0.0.1     data
F1CB30AFF1CB30B0010100000000000A68B167F92DD54196297159B5968923C
1.10.2    data
F1CB30AFF1CB30B0010100000000000CF0EFD81EA9F6324EA97B369351C56AC
[...]
```

4. Destroy the disk:

```
storage encryption disk destroy -disk disk_id
```

For complete command syntax, see the man page.



You are prompted to enter a confirmation phrase before continuing. Enter the phrase exactly as shown on the screen.

```
cluster1::> storage encryption disk destroy -disk 1.10.2

Warning: This operation will cryptographically destroy 1 spare or broken
self-encrypting disks on 1 node.
You cannot reuse destroyed disks unless you revert
them to their original state using the PSID value.
To continue, enter
    destroy disk
:destroy disk

Info: Starting destroy on 1 disk.
View the status of the operation by using the
"storage encryption disk show-status" command.
```

Emergency shred data on a FIPS drive or SED


In case of a security emergency, you can instantly prevent access to a FIPS drive or SED, even if power is not available to the storage system or the KMIP server.

Before you begin

- If you are using a KMIP server that has no available power, the KMIP server must be configured with an easily destroyed authentication item (for example, a smart card or USB drive).
- You must be a cluster administrator to perform this task.

Step

1. Perform emergency shredding of data on a FIPS drive or SED:

If...	Then...
<p>Power is available to the storage system and you have time to take the storage system offline gracefully</p>	<ol style="list-style-type: none"> a. If the storage system is configured as an HA pair, disable takeover. b. Take all aggregates offline and delete them. c. Set the privilege level to advanced: <pre>set -privilege advanced</pre> d. If the drive is in FIPS-compliance mode, set the FIPS authentication key ID for the node back to the default MSID: <pre>storage encryption disk modify -disk * -fips-key -id 0x0</pre> e. Halt the storage system. f. Boot into maintenance mode. g. Sanitize or destroy the disks: <ul style="list-style-type: none"> ◦ If you want to make the data on the disks inaccessible and still be able to reuse the disks, sanitize the disks: <pre>disk encrypt sanitize -all</pre> ◦ If you want to make the data on the disks inaccessible and you do not need to save the disks, destroy the disks: <pre>disk encrypt destroy disk_id1 disk_id2 ...</pre> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;">  <p>The <code>disk encrypt sanitize</code> and <code>disk encrypt destroy</code> commands are reserved for maintenance mode only. These commands must be run on each HA node, and are not available for broken disks.</p> </div> h. Repeat these steps for the partner node. This leaves the storage system in a permanently disabled state with all data erased. To use the system again, you must reconfigure it.

<p>Power is available to the storage system and you must shred the data immediately</p>	<p>a. If you want to make the data on the disks inaccessible and still be able to reuse the disks, sanitize the disks:</p> <p>b. If the storage system is configured as an HA pair, disable takeover.</p> <p>c. Set the privilege level to advanced:</p> <pre>set -privilege advanced</pre> <p>d. If the drive is in FIPS-compliance mode, set the FIPS authentication key ID for the node back to the default MSID:</p> <pre>storage encryption disk modify -disk * -fips-key-id 0x0</pre> <p>e. Sanitize the disk:</p> <pre>storage encryption disk sanitize -disk * -force-all-states true</pre>	<p>a. If you want to make the data on the disks inaccessible and you do not need to save the disks, destroy the disks:</p> <p>b. If the storage system is configured as an HA pair, disable takeover.</p> <p>c. Set the privilege level to advanced:</p> <pre>set -privilege advanced</pre> <p>d. Destroy the disks:</p> <pre>storage encryption disk destroy -disk * -force-all-states true</pre>
<p>The storage system panics, leaving the system in a permanently disabled state with all data erased. To use the system again, you must reconfigure it.</p>		
<p>Power is available to the KMIP server but not to the storage system</p>	<p>a. Log in to the KMIP server.</p> <p>b. Destroy all keys associated with the FIPS drives or SEDs that contain the data you want to prevent access to. This prevents access to disk encryption keys by the storage system.</p>	
<p>Power is not available to the KMIP server or the storage system</p>	<p>Destroy the authentication item for the KMIP server (for example, the smart card). This prevents access to disk encryption keys by the storage system.</p>	

For complete command syntax, see the man pages.

Return a FIPS drive or SED to service when authentication keys are lost

The system treats a FIPS drive or SED as broken if you lose the authentication keys for it permanently and cannot retrieve them from the KMIP server. Although you cannot access

or recover the data on the disk, you can take steps to make the SED's unused space available again for data.

Before you begin

You must be a cluster administrator to perform this task.

About this task

You should use this process only if you are certain that the authentication keys for the FIPS drive or SED are permanently lost and that you cannot recover them.

If the disks are partitioned, they must first be unpartitioned before you can start this process.



The command to unpartition a disk is only available at the diag level and should be performed only under NetApp Support supervision. **It is highly recommended that you contact NetApp Support before you proceed.** You can also refer to the Knowledge Base article [How to unpartition a spare drive in ONTAP](#).

Steps

1. Return a FIPS drive or SED to service:

If the SEDS are...	Use these steps...
--------------------	--------------------

Not in FIPS-compliance mode, or in FIPS-compliance mode and the FIPS key is available

- a. Set the privilege level to advanced:
`set -privilege advanced`
- b. Reset the FIPS key to the default manufacture secure ID 0x0:
`storage encryption disk modify -fips-key-id 0x0 -disk disk_id`
- c. Verify the operation succeeded:
`storage encryption disk show-status`
If the operation failed, use the PSID process in this topic.
- d. Sanitize the broken disk:
`storage encryption disk sanitize -disk disk_id`
Verify the operation succeeded with the command `storage encryption disk show-status` before proceeding to the next step.
- e. Unfail the sanitized disk:
`storage disk unfail -spare true -disk disk_id`
- f. Check whether the disk has an owner:
`storage disk show -disk disk_id`

If the disk does not have an owner, assign one.
`storage disk assign -owner node -disk disk_id`
 1. Enter the nodeshell for the node that owns the disks you want to sanitize:

`system node run -node node_name`

Run the disk `sanitize release` command.
- g. Exit the nodeshell. Unfail the disk again:
`storage disk unfail -spare true -disk disk_id`
- h. Verify that the disk is now a spare and ready to be reused in an aggregate:
`storage disk show -disk disk_id`

In FIPS-compliance mode, the FIPS key is not available, and the SEDs have a PSID printed on the label

- a. Obtain the PSID of the disk from the disk label.
- b. Set the privilege level to advanced:
`set -privilege advanced`
- c. Reset the disk to its factory-configured settings:
`storage encryption disk revert-to-original-state -disk disk_id -psid disk_physical_secure_id`
Verify the operation succeeded with the command `storage encryption disk show-status` before proceeding to the next step.
- d. If you are running ONTAP 9.8P5 or earlier, skip to the next step. If you are running ONTAP 9.8P6 or later, unfail the sanitized disk.
`storage disk unfail -disk disk_id`
- e. Check whether the disk has an owner:
`storage disk show -disk disk_id`

If the disk does not have an owner, assign one.
`storage disk assign -owner node -disk disk_id`
 1. Enter the nodeshell for the node that owns the disks you want to sanitize:

`system node run -node node_name`

Run the `disk sanitize release` command.
- f. Exit the nodeshell.. Unfail the disk again:
`storage disk unfail -spare true -disk disk_id`
- g. Verify that the disk is now a spare and ready to be reused in an aggregate:
`storage disk show -disk disk_id`

For complete command syntax, see the [command reference](#).

Return a FIPS drive or SED to unprotected mode

A FIPS drive or SED is protected from unauthorized access only if the authentication key ID for the node is set to a value other than the default. You can return a FIPS drive or SED to unprotected mode by using the `storage encryption disk modify` command to set the key ID to the default.

If an HA pair is using encrypting SAS or NVMe drives (SED, NSE, FIPS), you must follow this process for all drives within the HA pair prior to initializing the system (boot options 4 or 9). Failure to do this may result in future data loss if the drives are repurposed.

Before you begin

You must be a cluster administrator to perform this task.

Steps

1. Set the privilege level to advanced:

```
set -privilege advanced
```

2. If a FIPS drive is running in FIPS-compliance mode, set the FIPS authentication key ID for the node back to the default MSID 0x0:

```
storage encryption disk modify -disk disk_id -fips-key-id 0x0
```

You can use the `security key-manager query` command to view key IDs.

```
cluster1::> storage encryption disk modify -disk 2.10.11 -fips-key-id 0x0
```

```
Info: Starting modify on 14 disks.  
View the status of the operation by using the  
storage encryption disk show-status command.
```

Confirm the operation succeeded with the command:

```
storage encryption disk show-status
```

Repeat the `show-status` command until the numbers in “Disks Begun” and “Disks Done” are the same.

```
cluster1:: storage encryption disk show-status
```

	FIPS	Latest	Start	Execution	Disks
Disks Done	Support	Request	Timestamp	Time (sec)	Begun
Successful					
-----	-----	-----	-----	-----	-----
-----	-----				
cluster1	true	modify	1/18/2022 15:29:38	3	14
5					5

1 entry was displayed.

3. Set the data authentication key ID for the node back to the default MSID 0x0:

```
storage encryption disk modify -disk disk_id -data-key-id 0x0
```

The value of `-data-key-id` should be set to 0x0 whether you are returning a SAS or NVMe drive to unprotected mode.

You can use the `security key-manager query` command to view key IDs.


```
cluster1::> storage encryption disk modify -disk 2.10.11 -data-key-id
0x0
```

```
Info: Starting modify on 14 disks.
      View the status of the operation by using the
      storage encryption disk show-status command.
```

Confirm the operation succeeded with the command:

```
storage encryption disk show-status
```

Repeat the show-status command until the numbers are the same. The operation is complete when the numbers in “disks begun” and “disks done” are the same.

Maintenance mode

Beginning with ONTAP 9.7, you can rekey a FIPS drive from maintenance mode. You should only use maintenance mode if you cannot use the ONTAP CLI instructions in the earlier section.

Steps

1. Set the FIPS authentication key ID for the node back to the default MSID 0x0:

```
disk encrypt rekey_fips 0x0 disklist
```

2. Set the data authentication key ID for the node back to the default MSID 0x0:

```
disk encrypt rekey 0x0 disklist
```

3. Confirm the FIPS authentication key was successfully rekeyed:

```
disk encrypt show_fips
```

4. Confirm data authentication key was successfully rekeyed with:

```
disk encrypt show
```

Your output will likely display either the default MSID 0x0 key ID or the 64-character value held by the key server. The `Locked?` field refers to data-locking.

Disk	FIPS Key ID	Locked?
0a.01.0	0x0	Yes

Remove an external key manager connection

You can disconnect a KMIP server from a node when you no longer need the server. For example, you might disconnect a KMIP server when you are transitioning to volume encryption.

About this task

When you disconnect a KMIP server from one node in an HA pair, the system automatically disconnects the server from all cluster nodes.



If you plan to continue using external key management after disconnecting a KMIP server, make sure another KMIP server is available to serve authentication keys.

Before you begin

You must be a cluster or SVM administrator to perform this task.

Step

1. Disconnect a KMIP server from the current node:

For this ONTAP version...	Use this command...
ONTAP 9.6 and later	<pre>security key-manager external remove-servers -vserver SVM -key-servers host_name IP_address:port,...</pre>
ONTAP 9.5 and earlier	<pre>security key-manager delete -address key_management_server_ipaddress</pre>

In a MetroCluster environment, you must repeat these commands on both clusters for the admin SVM.

For complete command syntax, see the man pages.

The following ONTAP 9.6 command disables the connections to two external key management servers for `cluster1`, the first named `ks1`, listening on the default port 5696, the second with the IP address 10.0.0.20, listening on port 24482:

```
cluster1::> security key-manager external remove-servers -vserver  
cluster-1 -key-servers ks1,10.0.0.20:24482
```

Modify external key management server properties

Beginning with ONTAP 9.6, you can use the `security key-manager external modify-server` command to change the I/O timeout and user name of an external key management server.

Before you begin

- You must be a cluster or SVM administrator to perform this task.
- Advanced privileges are required for this task.
- In a MetroCluster environment, you must repeat these steps on both clusters for the admin SVM.

Steps

1. On the storage system, change to advanced privilege level:

```
set -privilege advanced
```

2. Modify external key manager server properties for the cluster:

```
security key-manager external modify-server -vserver admin_SVM -key-server  
host_name|IP_address:port,... -timeout 1..60 -username user_name
```



The timeout value is expressed in seconds. If you modify the user name, you are prompted to enter a new password. If you run the command at the cluster login prompt, *admin_SVM* defaults to the admin SVM of the current cluster. You must be the cluster administrator to modify external key manager server properties.

The following command changes the timeout value to 45 seconds for the `cluster1` external key management server listening on the default port 5696:

```
cluster1::> security key-manager external modify-server -vserver  
cluster1 -key-server ks1.local -timeout 45
```

3. Modify external key manager server properties for an SVM (NVE only):

```
security key-manager external modify-server -vserver SVM -key-server  
host_name|IP_address:port,... -timeout 1..60 -username user_name
```



The timeout value is expressed in seconds. If you modify the user name, you are prompted to enter a new password. If you run the command at the SVM login prompt, *SVM* defaults to the current SVM. You must be the cluster or SVM administrator to modify external key manager server properties.

The following command changes the username and password of the `svm1` external key management server listening on the default port 5696:

```
svm1::> security key-manager external modify-server -vserver svm11 -key  
-server ks1.local -username svmluser  
Enter the password:  
Reenter the password:
```

4. Repeat the last step for any additional SVMs.

Transition to external key management from onboard key management

If you want to switch to external key management from onboard key management, you must delete the onboard key management configuration before you can enable external key management.

Before you begin

- For hardware-based encryption, you must reset the data keys of all FIPS drives or SEDs to the default value.

[Returning a FIPS drive or SED to unprotected mode](#)

- For software-based encryption, you must unencrypt all volumes.

[Unencrypting volume data](#)

- You must be a cluster administrator to perform this task.

Step

1. Delete the onboard key management configuration for a cluster:

For this ONTAP version...	Use this command...
ONTAP 9.6 and later	<code>security key-manager onboard disable -vserver SVM</code>
ONTAP 9.5 and earlier	<code>security key-manager delete-key-database</code>

For complete command syntax, see the [ONTAP command reference](#).

Transition to onboard key management from external key management

If you want to switch to onboard key management from external key management, you must delete the external key management configuration before you can enable onboard key management.

Before you begin

- For hardware-based encryption, you must reset the data keys of all FIPS drives or SEDs to the default value.

[Returning a FIPS drive or SED to unprotected mode](#)

- You must have deleted all external key manager connections.

[Deleting an external key manager connection](#)

- You must be a cluster administrator to perform this task.

Procedure

The steps to transition your key management depend on the version of ONTAP you are using.

ONTAP 9.6 and later

1. Change to the advanced privilege level:

```
set -privilege advanced
```

2. Use the command:

```
security key-manager external disable -vserver admin_SVM
```



In a MetroCluster environment, you must repeat the command on both clusters for the admin SVM.

ONTAP 9.5 and earlier

Use the command:

```
security key-manager delete-kmip-config
```

What happens when key management servers are not reachable during the boot process

ONTAP takes certain precautions to avoid undesired behavior in the event that a storage system configured for NSE cannot reach any of the specified key management servers during the boot process.

If the storage system is configured for NSE, the SEDs are rekeyed and locked, and the SEDs are powered on, the storage system must retrieve the required authentication keys from the key management servers to authenticate itself to the SEDs before it can access the data.

The storage system attempts to contact the specified key management servers for up to three hours. If the storage system cannot reach any of them after that time, the boot process stops and the storage system halts.

If the storage system successfully contacts any specified key management server, it then attempts to establish an SSL connection for up to 15 minutes. If the storage system cannot establish an SSL connection with any specified key management server, the boot process stops and the storage system halts.

While the storage system attempts to contact and connect to key management servers, it displays detailed information about the failed contact attempts at the CLI. You can interrupt the contact attempts at any time by pressing Ctrl-C.

As a security measure, SEDs allow only a limited number of unauthorized access attempts, after which they disable access to the existing data. If the storage system cannot contact any specified key management servers to obtain the proper authentication keys, it can only attempt to authenticate with the default key which leads to a failed attempt and a panic. If the storage system is configured to automatically reboot in case of a panic, it enters a boot loop which results in continuous failed authentication attempts on the SEDs.

Halting the storage system in these scenarios is by design to prevent the storage system from entering a boot loop and possible unintended data loss as a result of the SEDs locked permanently due to exceeding the safety limit of a certain number of consecutive failed authentication attempts. The limit and the type of lockout protection depends on the manufacturing specifications and type of SED:

SED type	Number of consecutive failed authentication attempts resulting in lockout	Lockout protection type when safety limit is reached
HDD	1024	Permanent. Data cannot be recovered, even when the proper authentication key becomes available again.
X440_PHM2800MCTO 800GB NSE SSDs with firmware revisions NA00 or NA01	5	Temporary. Lockout is only in effect until disk is power-cycled.
X577_PHM2800MCTO 800GB NSE SSDs with firmware revisions NA00 or NA01	5	Temporary. Lockout is only in effect until disk is power-cycled.
X440_PHM2800MCTO 800GB NSE SSDs with higher firmware revisions	1024	Permanent. Data cannot be recovered, even when the proper authentication key becomes available again.
X577_PHM2800MCTO 800GB NSE SSDs with higher firmware revisions	1024	Permanent. Data cannot be recovered, even when the proper authentication key becomes available again.
All other SSD models	1024	Permanent. Data cannot be recovered, even when the proper authentication key becomes available again.

For all SED types, a successful authentication resets the try count to zero.

If you encounter this scenario where the storage system is halted due to failure to reach any specified key management servers, you must first identify and correct the cause for the communication failure before you attempt to continue booting the storage system.

Disable encryption by default

Beginning with ONTAP 9.7, aggregate and volume encryption is enabled by default if you have a volume encryption (VE) license and use an onboard or external key manager. If necessary, you can disable encryption by default for the entire cluster.

Before you begin

You must be a cluster administrator to perform this task, or an SVM administrator to whom the cluster administrator has delegated authority.

Step

1. To disable encryption by default for the entire cluster in ONTAP 9.7 or later, run the following command:

```
options -option-name encryption.data_at_rest_encryption.disable_by_default
```

Enable Zero Trust model

NetApp and Zero Trust

Zero Trust traditionally has been a network-centric approach of architecting micro core and perimeter (MCAP) to protect data, services, applications, or assets with controls known as a segmentation gateway. NetApp ONTAP is taking a data-centric approach to Zero Trust in which the storage management system becomes the segmentation gateway to protect and monitor access of our customer's data. In particular, the FPolicy Zero Trust engine and the FPolicy partner ecosystem becomes a control center to gain a detailed understanding of normal and aberrant data access patterns and identify insider threats.



Beginning in July 2024, content from the technical report *TR-4829: NetApp and Zero Trust: Enabling a data-centric Zero Trust model*, which was previously published as a PDF, has been integrated with the rest of the ONTAP product documentation.

Data is the most important asset your organization has. Insider threats are the cause of 18% of data breaches, according to the 2022 [Verizon Data Breach Investigations Report](#). Organizations can ramp up their vigilance by deploying industry-leading Zero Trust controls around data with NetApp ONTAP data management software.

What Is Zero Trust?

The Zero Trust model was first developed by [John Kindervag[^]] at Forrester Research. It envisions network security from the inside-out rather than from the outside-in. The inside-out Zero Trust approach identifies a microcore and perimeter (MCAP). The MCAP is an interior definition of data, services, applications, and assets to be protected with a comprehensive set of controls. The concept of a secure outer perimeter is obsolete. Entities that are trusted and allowed to successfully authenticate through the perimeter can then make the organization vulnerable to attacks. Insiders, by definition, are already inside the secure perimeter. Employees, contractors, and partners are insiders, and they must be enabled to operate with appropriate controls for performing their roles within your organization's infrastructure.

Zero Trust was mentioned as a technology that offers promise to the DoD in September 2019 [FY19-23 DoD Digital Modernization Strategy](#). It defines Zero Trust as, "A cybersecurity strategy that embeds security throughout the architecture for the purpose of stopping data breaches. This data-centric security model eliminates the idea of trusted or untrusted networks, devices, personas, or processes and shifts to multi-attribute based confidence levels that enable authentication and authorization policies under the concept of least privileged access. Implementing zero trust requires rethinking how we utilize existing infrastructure to implement security by design in a simpler and more efficient way while enabling unimpeded operations."

In August of 2020, the NIST published [Special Pub 800-207 Zero Trust Architecture \(ZTA\)](#). ZTA focuses on protecting resources, not network segments, because the network location is no longer seen as the prime component of the security posture of the resource. Resources are data and computing. ZTA strategies are for enterprise network architects. ZTA introduces some new terminology from the original Forrester concepts. Protection mechanisms called the policy decision point (PDP) and the policy enforcement point (PEP) are analogous to a Forrester segmentation gateway. ZTA introduces four deployment models:

- Device-agent or gateway-based deployment
- Enclave-based deployment (somewhat analogous to the Forrester MCAP)

- Resource portal-based deployment
- Device application sandboxing

For the purposes of this documentation, we use Forrester Research concepts and terminology rather than the NIST ZTA.

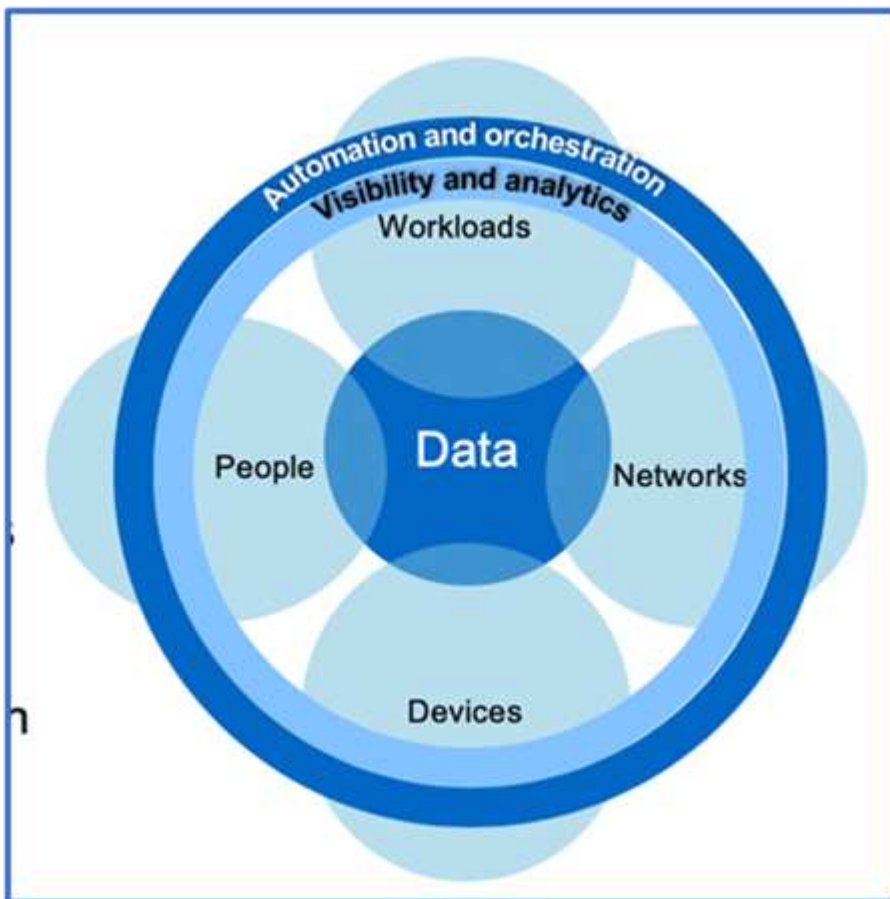
Security resources

For information about reporting vulnerabilities and incidents, NetApp security responses, and customer confidentiality, see the [NetApp security portal](#).

Architect a data-centric approach to Zero Trust with ONTAP

A Zero Trust network is defined by a data-centric approach in which the security controls should be as close to the data as possible. The capabilities of ONTAP, coupled with the NetApp FPolicy partner ecosystem, can provide the necessary controls for the data-centric Zero Trust model.

ONTAP is security-rich data management software from NetApp, and the FPolicy Zero Trust Engine is an industry-leading ONTAP capability that provides a granular, file-based event notification interface. NetApp FPolicy partners can use this interface to provide greater illumination of data access within ONTAP.



Architect a Zero Trust data-centric MCAP

To architect a data-centric Zero Trust MCAP, follow these steps:

1. Identify the location of all organizational data.
2. Classify your data.
3. Securely dispose of data that you no longer require.
4. Understand what roles should have access to the data classifications.
5. Apply the principle of least privilege to enforce access controls.
6. Use multifactor authentication for administrative access and data access.
7. Use encryption for data at rest and data in flight.
8. Monitor and log all access.
9. Alert suspicious access or behaviors.

Identify the location of all organizational data

The FPolicy capability of ONTAP coupled with the NetApp Alliance Partner ecosystem of FPolicy partners lets you identify where your organization's data exists and who has access to it. This is done with user behavioral analytics, which identifies whether data access patterns are valid. More details about user behavioral analytics are discussed in Monitor and log all access. If you do not understand where your data is and who has access to it, user behavioral analytics can provide a baseline to build classification and policy from empirical observations.

Classify your data

In the terminology of the Zero Trust model, classification of data involves identification of toxic data. Toxic data is sensitive data that is not intended to be exposed outside an organization. Disclosure of toxic data could violate regulatory compliance and damage an organization's reputation. In terms of regulatory compliance, toxic data includes cardholder data for the [Payment Card Industry Data Security Standard \(PCI-DSS\)](#), personal data for the EU [General Data Protection Regulation \(GDPR\)](#), or healthcare data for the [Health Insurance Portability and Accountability Act \(HIPAA\)](#). You can use NetApp [BlueXP Classification](#) (formerly known as Cloud Data Sense), an AI-driven toolkit, to automatically scan, analyze, and categorize your data.

Securely dispose of data you no longer require

After classifying your organization's data, you might discover that some of your data is no longer necessary or relevant to the function of your organization. The retention of unnecessary data is a liability, and such data should be deleted. For an advanced mechanism to cryptographically erase data, see the description of secure purge in Data at rest encryption.

Understand what roles should have access to the data classifications and apply the principle of least privilege to enforce access controls

Mapping access to sensitive data and applying the principle of least privilege means giving people in your organization access to only the data required to perform their jobs. This process involves role-based access control ([RBAC](#)), which applies to data access and administrative access.

With ONTAP, a storage virtual machine (SVM) can be used to segment organizational data access by tenants within an ONTAP cluster. RBAC can be applied to data access as well as administrative access to the SVM. RBAC can also be applied at the cluster administrative level.

In addition to RBAC, you can use ONTAP [multi-admin verification \(MAV\)](#) to require one or more administrators to approve commands such as `volume delete` or `volume snapshot delete`. Once MAV is enabled, modifying or disabling MAV requires MAV administrator approval.

Another way of protecting Snapshot copies is with ONTAP [Snapshot copy locking](#). Snapshot copy locking is a SnapLock capability where

Snapshot copies are rendered indelible manually or automatically with a retention period on the volume Snapshot copy policy. Snapshot copy locking is also referred to as tamper-proof Snapshot copy locking. The purpose of Snapshot copy locking is to prevent rogue or untrusted administrators from deleting Snapshot copies on the primary and secondary ONTAP systems. Rapid recovery of locked Snapshot copies on primary systems can be achieved in order to restore volumes corrupted by ransomware.

Use multifactor authentication for administrative access and data access

In addition to cluster administrative RBAC, [multifactor authentication \(MFA\)](#) can be deployed for ONTAP web administrative access and Secure Shell (SSH) command-line access. MFA for administrative access is a requirement for U.S. public sector organizations or those that must follow the PCI-DSS. MFA makes it impossible for an attacker to compromise an account using only a username and password. MFA requires two or more independent factors to authenticate. An example of two-factor authentication is something a user possesses, such as a private key, and something a user knows, such as a password. Administrative web access to ONTAP System Manager or ActiveIQ Unified Manager is enabled by Security Assertion Markup Language (SAML) 2.0. SSH command-line access uses chained two-factor authentication with a public key and password.

You can control user and machine access through APIs with the identity and access management capabilities in ONTAP:

- User:
 - **Authentication and authorization.** Through NAS protocol capabilities for SMB and NFS.
 - **Audit.** Syslog of access and events. Detailed audit logging of CIFS protocol to test authentication and authorization policies. Fine granular FPolicy auditing of detailed NAS access at the file level.
- Device:
 - **Authentication.** Certificate-based authentication for API access.
 - **Authorization.** Default or custom role-based access control (RBAC).
 - **Audit.** Syslog of all actions taken.

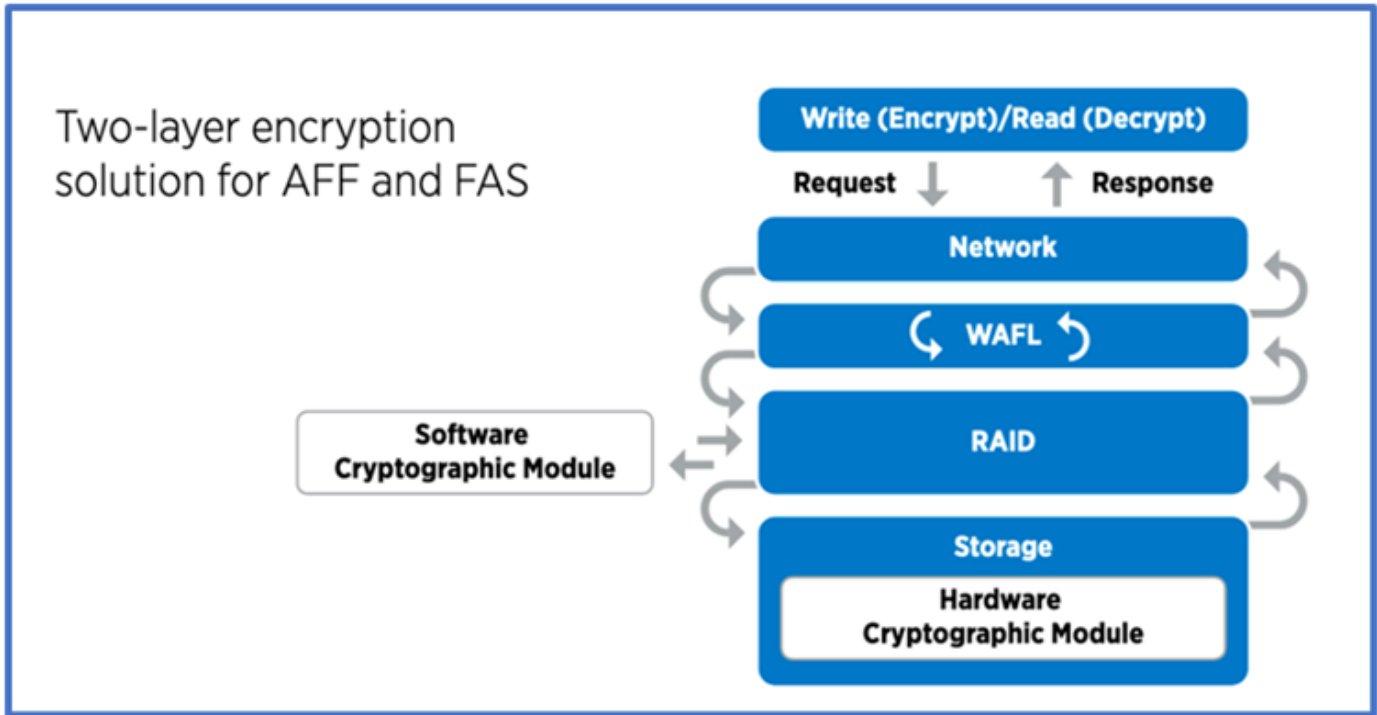
Use encryption for data at rest and data in flight

Data at rest encryption

Each day, there are new requirements for mitigating storage-system risks and infrastructure gaps when an organization repurposes drives, returns defective drives, or upgrades to larger drives by selling or trading them in. As administrators and operators of data, storage engineers are expected to manage and maintain data securely throughout its lifecycle. [NetApp Storage Encryption \(NSE\)](#), [NetApp Volume Encryption \(NVE\)](#), and [NetApp Aggregate Encryption](#) help you encrypt all your data at rest all the time, whether or not it is toxic, and without affecting daily operations. [NSE](#) is an ONTAP hardware [data-at-rest](#) solution that makes use of FIPS 140-2 level 2 validated self-encrypting drives. [NVE and NAE](#) are an ONTAP software [data-at-rest](#) solution that makes use of the [FIPS 140-2 level 1 validated NetApp Cryptographic Module](#). With NVE and NAE, either hard drives or solid-state drives can be used for data-at-rest encryption. Plus, NSE drives can be used to provide a native, layered encryption solution that provides encryption redundancy and additional security. If one layer is breached, then the second layer still secures the data. These capabilities make ONTAP well positioned for [quantum-ready encryption](#).

NVE also provides a capability called [secure purge](#) that cryptographically removes toxic data from data spills when sensitive files are written to a non-classified volume.

Either the [Onboard Key Manager \(OKM\)](#), which is the key manager built into ONTAP, or [approved third-party external key managers](#) can be used with NSE and NVE to securely store keying material.



As seen in the figure above, hardware and software based encryption can be combined. This capability led to the [validation of ONTAP into the NSA's commercial solutions for classified program](#) that allows for storage of top secret data.

Data-in-flight encryption

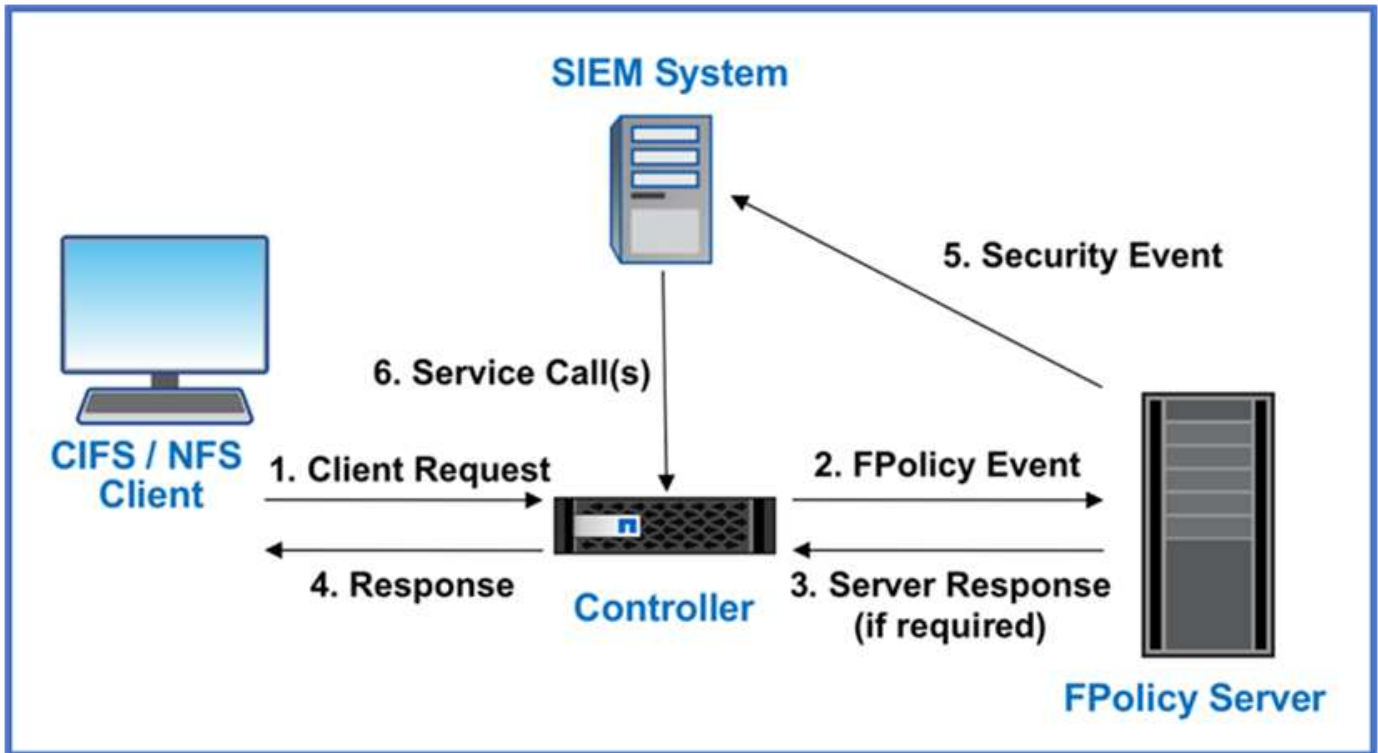
ONTAP data-in-flight encryption protects user data access and control-plane access. User data access can be encrypted by SMB 3.0 encryption for Microsoft CIFS share access or by krb5P for NFS Kerberos 5. User data access can also be encrypted with [IPsec](#) for CIFS, NFS, and iSCSI. Control plane access is encrypted with Transport Layer Security (TLS). ONTAP provides [FIPS](#) compliance mode for control plane access, which enables FIPS-approved algorithms and disables algorithms that are not FIPS approved. Data replication is encrypted with [cluster peer encryption](#). This provides encryption for the ONTAP SnapVault and SnapMirror technologies.

Monitor and log all access

After RBAC policies are in place, you must deploy active monitoring, auditing, and alerting. The FPolicy Zero Trust Engine from NetApp ONTAP, coupled with the [NetApp FPolicy partner ecosystem](#), provides the necessary controls for the data-centric Zero Trust model. NetApp ONTAP is security-rich data management software, and [FPolicy](#) is an industry-leading ONTAP capability that provides a granular file-based event notification interface. NetApp FPolicy partners can use this interface to provide greater illumination of data access within ONTAP. The FPolicy capability of ONTAP, coupled with the NetApp Alliance Partner ecosystem of FPolicy partners, lets you identify where your organization's data exists and who has access to it. This is done with user behavioral analytics, which identifies whether data access patterns are valid. User behavioral analytics can be used to alert for suspicious or aberrant data access that is out of the normal pattern and, if necessary, take actions to deny access.

FPolicy partners are moving beyond user behavioral analytics toward machine learning (ML) and artificial intelligence (AI) for greater event fidelity and fewer, if any, false positives. All events should be logged to a syslog server or to a security information and event management

(SIEM) system that can also employ ML and AI.



NetApp's Storage Workload Security (formerly known as [Cloud Secure](#)) makes use of the FPolicy interface and user behavioral analytics on both cloud and on-premises ONTAP storage systems to give you real-time alerts of malicious user behavior. Storage Workload Security protects organizational data from being misused by malicious or compromised users through advanced machine learning and anomaly detection. Storage Workload Security can identify ransomware attacks or other miscreant behaviors, invoke Snapshot copies and quarantine malicious users. Storage Workload Security also has a forensics capability to view in great detail user and entity activities. Storage Workload Security is a part of NetApp Cloud Insights.

In addition to Storage Workload Security, ONTAP has an onboard ransomware detection capability known as [Autonomous Ransomware Protection](#) (ARP). ARP uses machine learning to determine if abnormal file activity indicates a ransomware attack is underway and invokes a Snapshot copy and alert to administrators. Storage Workload Security integrates with ONTAP to receive ARP events and provides an additional analytics and automatic responses layer.

NetApp security automation and orchestration controls external to ONTAP

Automation allows you to perform a process or procedure with minimal human assistance. Automation enables organizations to scale Zero Trust deployments far beyond manual procedures to defend against miscreant activities that are also automated.

Ansible is an open-source software provisioning, configuration management, and application-deployment tool. It runs on many Unix-like systems, and it can configure both Unix-like systems as well as Microsoft Windows. It includes its own declarative language to describe system configuration. Ansible was written by Michael DeHaan and acquired by Red Hat in 2015. Ansible is agentless, temporarily connecting remotely through SSH or Windows Remote Management (allowing remote PowerShell execution) to perform tasks. NetApp has developed more than [150 Ansible modules for ONTAP software](#), enabling further integration with the Ansible automation framework. Ansible modules for NetApp deliver a set of instructions for how to define the desired state and relay it to the target NetApp environment. Modules are built to support tasks like setting up licensing,

creating aggregates and storage virtual machines, creating volumes, and restoring snapshots to name a few. An Ansible role has been [published on GitHub](#) specific to the NetApp DoD Unified Capabilities (UC) Deployment Guide.

Using the library of available modules, users can easily develop Ansible playbooks and customize them to their own applications and business needs to automate mundane tasks. After a playbook is written, you can run it to execute the specified task, which saves time and improves productivity. NetApp has created and shared sample playbooks that can be used directly or customized for your needs.

Cloud Insights is an infrastructure monitoring tool that gives you visibility into your complete infrastructure. With Cloud Insights, you can monitor, troubleshoot, and optimize all your resources, including your public cloud instances and your private data centers. Cloud Insights can reduce mean time to resolution by 90% and prevent 80% of cloud issues from affecting end users. It can also reduce cloud infrastructure costs by an average of 33% and reduce your exposure to insider threats by protecting your data with actionable intelligence. The Storage Workload Security capability of Cloud Insights enables user behavioral analytics with AI and ML to alert when aberrant user behaviors occur due to an insider threat. For ONTAP, Storage Workload Security makes use of the Zero Trust FPolicy engine.

Zero Trust and hybrid cloud deployments

NetApp is the data authority for the hybrid cloud. NetApp offers a variety of options for extending on-premises data management systems to the hybrid cloud with Amazon Web Services (AWS), Microsoft Azure, Google Cloud Platform (GCP), and other leading cloud providers. NetApp hybrid-cloud solutions support the same Zero Trust security controls that are available with on-premises ONTAP systems and ONTAP Select software-defined storage.

You can easily expand capacity in public clouds without typical capex constraints by using the NetApp Cloud Volumes service, the first enterprise-class, cloud-native file service for AWS and GCP, and Azure NetApp Files for Microsoft Azure. Ideal for data-intensive workloads such as analytics and DevOps, these cloud data services combine elastic, on-demand storage as a service from NetApp with ONTAP data management in a fully managed offering.

For those seeking advanced data services for cloud block or object storage services such as AWS EBS and S3 or Azure storage, Cloud Volumes ONTAP provides data management between your on-premises environment and the public cloud with a single common view. Running in AWS or Azure as an on-demand instance, Cloud Volumes ONTAP provides the storage efficiency, availability, and scalability of ONTAP software. ONTAP enables the movement of data between your on-premises ONTAP systems and AWS or Azure storage environment with NetApp SnapMirror data replication software.

Copyright information

Copyright © 2024 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.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

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.