



NetApp Technical Reports

NetApp Technical Reports

NetApp
June 09, 2026

This PDF was generated from <https://docs.netapp.com/us-en/ontap-technical-reports/index.html> on June 09, 2026. Always check docs.netapp.com for the latest.

Table of Contents

- NetApp Technical Reports 1
- ONTAP and application and database technical reports 2
 - Microsoft SQL Server 2
 - MySQL 2
 - Oracle 2
 - PostgreSQL 4
 - SAP HANA 4
 - Epic 4
- Business continuity technical reports 5
 - SnapMirror active sync (formerly SM-BC) 5
 - MetroCluster 5
- ONTAP data protection and disaster recovery technical reports 6
 - SnapMirror 6
 - Application and infrastructure with SnapMirror 6
 - ONTAP cyber vault 6
- ONTAP FlexCache and FlexGroup volume technical reports 8
 - FlexCache 8
 - FlexCache write-back 8
 - FlexGroup volumes 8
- ONTAP NAS technical reports 9
 - NFS 9
 - SMB 9
 - Multiprotocol 9
 - ONTAP S3 9
 - Name services 9
 - NAS security 10
- ONTAP networking technical reports 11
- ONTAP SAN technical reports 12
- Security 13
 - ONTAP security technical reports 13
 - ONTAP cyber vault 13
 - Ransomware 13
 - Zero Trust 13
 - Multifactor authentication 13
 - Multitenancy 14
 - Standards 14
 - Attribute-based access control 14
- NetApp solution for ransomware 14
 - Ransomware and NetApp’s protection portfolio 14
 - SnapLock and tamperproof snapshots for ransomware protection 17
 - FPolicy file blocking 18
 - Data Infrastructure Insights Storage Workload Security 19
 - NetApp ONTAP built-in on-box AI-based detection and response 19

Air-gapped WORM protection with cyber vaulting in ONTAP	21
Digital Advisor ransomware protection	22
Comprehensive resilience with NetApp ransomware protection	22
NetApp and Zero Trust	23
NetApp and Zero Trust	24
Architect a data-centric approach to Zero Trust with ONTAP	25
NetApp security automation and orchestration controls external to ONTAP	29
Zero Trust and hybrid cloud deployments	30
Attribute-based access control	30
Attribute-based access control with ONTAP	30
Approaches to attribute-based access control (ABAC) in ONTAP	31
Security hardening	44
ONTAP security hardening guides	44
Hardening guides	44
ONTAP security hardening guidelines	44
ONTAP security hardening overview	44
ONTAP image validation	45
Local storage administrator accounts	45
System administration methods	60
ONTAP autonomous ransomware protection	66
Storage administrative system auditing	66
Storage encryption in ONTAP	68
Data replication encryption	70
IPsec data-in-flight encryption	71
FIPS mode and TLS and SSL management in ONTAP	72
Create a CA-signed digital certificate	74
Online certificate status protocol	75
SShv2 management	75
NetApp AutoSupport	76
Network Time Protocol	77
NAS file system local accounts (CIFS workgroup)	77
NAS file system auditing	78
Configure and enable CIFS SMB signing and sealing	80
NFS securing	81
Enable Lightweight Directory Access Protocol signing and sealing	83
Create and use a NetApp FPolicy	83
Security characteristics of LIF roles in ONTAP	85
Protocol and port security	86
Storage Systems	90
AFX	90
NetApp AFX overview: Learn about NetApp AFX	90
What's new in NetApp AFX	94
How NetApp AFX architecture differs from unified ONTAP	94
Performance	112
Management tools	119

Networking, security, and operations	122
Hardware details	124
Maximums and limits	128
Where to find additional information	129
ONTAP SnapCenter technical reports	130
SnapCenter for Oracle	130
SnapCenter for Microsoft SQL Server	130
SnapCenter for Microsoft Exchange Server	130
SnapCenter for SAP HANA	130
SnapCenter hardening guide	131
ONTAP tiering technical reports	132
ONTAP virtualization technical reports	133
Legal notices	135
Copyright	135
Trademarks	135
Patents	135
Privacy policy	135
Open source	135
ONTAP	135
ONTAP Mediator for MetroCluster IP configurations	136

NetApp Technical Reports

ONTAP and application and database technical reports

ONTAP is the foundation for data management and data protection for many enterprise application and database technologies. The following technical reports provide guidance on NetApp recommended practices and implementation procedures for Microsoft SQL Server, MySQL, Oracle, PostgreSQL, SAP HANA and Epic.

Microsoft SQL Server

SQL Server is the foundation of Microsoft's data platform, delivering mission-critical performance with in-memory technologies and faster insights on any data, whether on the premises or in the cloud.

[Best practice for Microsoft SQL Server with ONTAP](#)

Learn how storage administrators and database administrators can successfully deploy Microsoft SQL Server on ONTAP storage.



This documentation replaces the previously published technical report *TR-4590: Best practice guide for Microsoft SQL Server with ONTAP*.

[TR-4976: Virtualized Microsoft SQL Server performance on NetApp AFF A-Series and C-Series systems](#)

Learn about Microsoft SQL Server performance characteristics using a NetApp AFF A-Series and C-Series systems as well as guidance on how to select the right system based on workload.

[TR-4714: Best practices for Microsoft SQL Server using SnapCenter](#)

Learn now to successfully deploy Microsoft SQL Server on ONTAP storage using SnapCenter technology for data protection.

MySQL

This document describes the configuration requirements and provides guidance on tuning and storage configuration for deploying MySQL on ONTAP.

[MySQL database on NetApp ONTAP best practices](#)

MySQL and its variants, including MariaDB and Percona, are widely used for many enterprise applications. These applications range from global social networking sites and massive ecommerce systems to SMB hosting systems containing thousands of database instances. Learn about the configuration requirements and guidance on tuning and storage configuration for deploying MySQL on ONTAP.



This documentation replaces the previously published technical report *TR-4722: MySQL database on NetApp ONTAP best practices*.

Oracle

ONTAP is designed for Oracle databases. For decades, ONTAP has been optimized for the unique demands of relational database I/O and multiple ONTAP features were created specifically to service the needs of Oracle databases and even at the request of Oracle Inc. itself.

[Oracle databases on ONTAP](#)

Learn about the recommended practices that enable storage administrators and database administrators to successfully deploy Oracle on ONTAP storage.

[Oracle data protection with ONTAP](#)

Learn about the recommended practices that enable storage administrators and database administrators to successfully backup, recover, replicate and provide disaster recovery to Oracle on ONTAP storage.

[Oracle disaster recovery with ONTAP](#)

Learn about the recommended practices, test procedures, and other considerations for operating Oracle databases on a MetroCluster and SnapMirror Business Continuity.

[Migration of Oracle databases to ONTAP storage systems](#)

Learn about the overall considerations for planning a migration strategy, the three different levels in which data movement takes place, and details some of the various procedures available.



Documentation linked above replaces these previously published technical reports *TR-3633: Oracle databases on ONTAP*; *TR-4591: Oracle data protection: backup, recovery, replication*; *TR-4592: Oracle on MetroCluster*; and *TR-4534: Migration of Oracle databases to NetApp storage systems*

[TR-4969: Oracle database performance on AFF A-Series and C-Series](#)

ONTAP is a powerful data-management platform with native capabilities that include inline compression, nondisruptive hardware upgrades, and the ability to import a LUN from a foreign storage array. Up to 24 nodes can be clustered together, simultaneously serving data through Network File System (NFS), Server Message Block (SMB), iSCSI, Fibre Channel (FC), and Nonvolatile Memory Express (NVMe) protocols. In addition, Snapshot technology is the basis for creating tens of thousands of online backups and fully operational database clones. In addition to the rich feature set of ONTAP, there are a wide variety of user requirements, including database size, performance requirements, and data protection needs. Learn about bare metal database performance using AFF storage systems, including both the A-Series and C-Series, and it covers both maximums and the practical difference between the two AFF options.

[TR-4971: Virtualized Oracle database performance on AFF A-Series and C-Series](#)

ONTAP is a powerful data-management platform with native capabilities that include inline compression, nondisruptive hardware upgrades, and the ability to import a LUN from a foreign storage array. Up to 24 nodes can be clustered together, simultaneously serving data through Network File System (NFS), Server Message Block (SMB), iSCSI, Fibre Channel (FC), and Nonvolatile Memory Express (NVMe) protocols. In addition, Snapshot technology is the basis for creating tens of thousands of online backups and fully operational database clones. In addition to the rich feature set of ONTAP, there are a wide variety of user requirements, including database size, performance requirements, and data protection needs. Learn about virtualized database performance using AFF storage systems, including both the A-Series and C-Series, and it covers both maximums and the practical difference between the two AFF options.

[TR-4695: Database storage tiering with FabricPool](#)

Learn about the benefits and configuration options of FabricPool with various databases, including the Oracle relational database management system (RDBMS).

[TR-4899: Oracle database transparent application failover with SnapMirror active sync](#)

SnapMirror active sync (formerly SM-BC) and Oracle Real Application Cluster (RAC) can provide transparent application failover (TAF) and continuity in the face of site outages and true disasters. Learn about the configuration guidance and recommended practices of a AFF storage array with SnapMirror Active Sync as the storage component of Oracle RAC.

[TR-4876: Oracle Multitenancy with ONTAP solution and deployment best practices](#)

Learn about the solution recommended practices on how to provision, manage, and protect Oracle Multitenant

databases by using ONTAP storage to maximize the benefits of both Oracle Multitenant databases and the features of ONTAP software.

PostgreSQL

PostgreSQL comes with variants that include PostgreSQL, PostgreSQL Plus, and EDB Postgres Advanced Server (EPAS). PostgreSQL is typically deployed as the back-end database for multitier applications. NetApp ONTAP is an excellent choice for running PostgreSQL databases due to its reliability, high performance and efficient data management capabilities.

[PostgreSQL database on ONTAP best practices](#)

PostgreSQL comes with variants that include PostgreSQL, PostgreSQL Plus, and EDB Postgres Advanced Server (EPAS). PostgreSQL is typically deployed as the back-end database for multi-tier applications. It is supported by common middleware packages (such as PHP, Java, Python, Tcl/Tk, ODBC, and JDBC) and has historically been a popular choice for open-source database management systems. Learn about the configuration requirements and guidance on tuning and storage configuration for deploying PostgreSQL on ONTAP.



This documentation replaces the previously published technical report *TR-4770: PostgreSQL database on ONTAP best practices*.

SAP HANA

[SAP HANA database solutions on ONTAP](#)

Best practices for configuring, managing and automating SAP solutions can be found on the NetApp SAP Solutions page.

Epic

[Epic on ONTAP best practices](#)

A guide to understanding the best practices for deploying Epic on premises and in the cloud while meeting configuration standards for proper deployment on ONTAP.



This documentation replaces the previously published technical report *TR-3923: NetApp best practices for Epic*.

Business continuity technical reports

NetApp offers a wide range of solutions that rationalize where applications and data live to improve performance cost effectively. Data protection, replication, and continuous availability: ONTAP data management can simplify data protection with set-it-and-forget-it policy management, while delivering business continuity with MetroCluster and SnapMirror active sync.



These technical reports expand on the [ONTAP SnapMirror active sync](#) and [ONTAP MetroCluster](#) product documentation.

SnapMirror active sync (formerly SM-BC)

[TR-4878: SnapMirror active sync](#)

SnapMirror active sync is a continuously available storage solution with application-level granularity, available for ONTAP running on AFF or All SAN Array (ASA) storage systems, to meet the RPO 0 and RTO 0 needs of the most critical business applications.

MetroCluster

[TR-4705: NetApp MetroCluster solution architecture and design](#)

This document describes high-level architecture and design concepts for MetroCluster features in ONTAP.

MetroCluster IP

[TR-4689: NetApp MetroCluster IP](#)

MetroCluster is a continuously available storage solution for ONTAP running on FAS and AFF systems. MetroCluster IP is the latest evolution that uses an Ethernet-based back-end storage fabric. MetroCluster IP provides a highly redundant configuration to meet the needs of the most critical business applications. MetroCluster IP is included in ONTAP and provides NAS and SAN connectivity for clients and servers that use ONTAP storage.

MetroCluster FC

[TR-4375: NetApp MetroCluster FC](#)

MetroCluster provides continuous data availability across geographically separated data centers for mission-critical applications. Learn about MetroCluster FC recommended practices, design decisions, and supported configurations.

ONTAP data protection and disaster recovery technical reports

SnapMirror is a cost-effective, easy-to-use unified replication solution across the data fabric. It replicates data at high speeds over LAN or WAN. You get high data availability and fast data replication for your business-critical applications, such as Microsoft Exchange, Microsoft SQL Server, and Oracle, in both virtual and traditional environments. When you replicate data to one or more ONTAP storage systems and continually update the secondary data, your data is kept current and is available whenever you need it. No external replication servers are required.



These technical reports expand on the [ONTAP Data protection and disaster recovery](#) product documentation.

SnapMirror

SnapMirror Asynchronous

[TR-4015: SnapMirror Asynchronous configuration and best practices](#)

Learn about recommended practices for configuring SnapMirror Asynchronous (SM-A) replication of volumes, consistency groups, and storage virtual machines (SVM disaster recovery).

[TR-4678: Data protection and backup ONTAP FlexGroup volumes](#)

Learn about recommended data protection and backup for FlexGroup volumes. Topics include Snapshot copies, SnapMirror, and other data protection and backup solutions.

SnapMirror Synchronous

[TR-4733: SnapMirror Synchronous configuration and best practices](#)

Learn about recommended practices for configuring SnapMirror Synchronous (SM-S) replication.

SnapMirror Three-Data-Center DR

[TR-4832: Three-Data-Center disaster recovery using NetApp SnapMirror for ONTAP 9.7](#)

Learn about a three-data-center disaster recovery configuration using ONTAP SnapMirror technology for replication.

Application and infrastructure with SnapMirror

[TR-4900: VMware Site Recovery Manager with ONTAP](#)

ONTAP has been a leading storage solution for VMware vSphere environments since its introduction into the modern data center in 2002, and it continues to add innovative capabilities to simplify management while reducing costs. Learn about recommended the ONTAP solution for VMware Site Recovery Manager (SRM), VMware's industry leading disaster recovery (DR) software, including the latest product information and recommended practices to streamline deployment, reduce risk, and simplify ongoing management.

ONTAP cyber vault

[ONTAP cyber vault](#)

NetApp's ONTAP based cyber vault provides organizations with a comprehensive and flexible solution for

protecting their most critical data assets. By leveraging logical air-gapping with robust hardening methodologies, ONTAP enables you to create secure, isolated storage environments that are resilient against evolving cyber threats. With ONTAP, you can ensure the confidentiality, integrity, and availability of your data while maintaining the agility and efficiency of your storage infrastructure.

ONTAP FlexCache and FlexGroup volume technical reports

NetApp NAS solutions simplify data management and help you keep pace with growth while optimizing costs. ONTAP NAS solutions give you nondisruptive operations, proven efficiency, and seamless scalability within a unified architecture. Powered by ONTAP, scale-out NAS leverages the massive ONTAP ecosystem, with a significant innovation lead and vision for aggressive future innovation.



These technical reports expand on the [ONTAP FlexCache volume](#) and [ONTAP FlexGroup volume](#) product documentation.

FlexCache

[TR-4743: FlexCache in ONTAP](#)

FlexCache is a caching technology that creates sparse, writable replicas of volumes on the same or different ONTAP clusters. It can bring data and files closer to the user for faster throughput with a smaller footprint. Learn how FlexCache can be used, the recommended practices, limits, and considerations for design and implementation.

FlexCache write-back

[FlexCache write-back](#)

Introduced in ONTAP 9.15.1, FlexCache write-back is an alternate mode of operation for writing at a cache. Write-back allows the write to be committed to stable storage at the cache and acknowledged to the client without waiting for the data to make it to the origin. The data is asynchronously flushed back to the origin. The result is a globally distributed file system that enables writes to perform at near-local speeds for specific workloads and environments, offering significant performance benefits.

FlexGroup volumes

[TR-4571: NetApp ONTAP FlexGroup volumes best practices and implementation guide](#)

Learn about FlexGroup volumes, recommended practices and implementation tips. FlexGroup volumes are an evolution of ONTAP scale-out NAS containers, that blends nearly infinite capacity with predictable, low latency performance in metadata-heavy workloads.

[TR-4678: Data protection and backup of FlexGroup volumes](#)

Learn about data protection and backup for FlexGroup volumes including Snapshot copies, SnapMirror, and other data protection and backup solutions.

ONTAP NAS technical reports

NetApp NAS solutions simplify data management and help you keep pace with growth while optimizing costs. ONTAP NAS solutions provide for nondisruptive operations, efficiency, and seamless scalability within a unified architecture. Powered by NetApp ONTAP, scale-out NAS leverages the massive ONTAP ecosystem, with a significant innovation lead and vision for aggressive future innovation.



These technical reports expand on the [ONTAP NAS storage management](#) and [ONTAP S3 storage management](#) product documentation.

NFS

[TR-4067: NFS in ONTAP best practice and implementation guide](#)

Learn about basic concepts, support information, configuration tips, and recommended practices for NFS in ONTAP.

[TR-4962: NFSv4.2 extended attributes](#)

Learn about enabling and using NFSv4.2 extended attributes in ONTAP 9.12.1 and later.

SMB

[TR-4740: SMB 3.0 multichannel](#)

Microsoft introduced Multichannel in the SMB 3.0 protocol with the goal of improving the SMB3 protocol by addressing the performance and reliability limitations of SMB1 and SMB2. Learn about the Multichannel feature in ONTAP, including its capabilities, recommended practices, and performance test results.

Multiprotocol

[TR-4887: Multiprotocol NAS in ONTAP overview and best practices](#)

Learn how multiprotocol NAS access works in ONTAP and the recommended practices for multiprotocol environments.

ONTAP S3

[TR-4814: S3 in ONTAP best practices](#)

Learn about recommended practices for using the Amazon Simple Storage Service (S3) with ONTAP software as well as capabilities and configurations for using ONTAP as an object store with native S3 applications or as a tiering destination for FabricPool.

Name services

[TR-4523: DNS load balancing in ONTAP](#)

Learn how to configure ONTAP for use with DNS load balancing methodologies including DNS in ONTAP, various configuration methods, and recommended practices.

[TR-4668: Name services best practices guide](#)

Learn about recommended practices, limits, and considerations when implementing network-attached storage

(NAS) solutions such as CIFS/SMB and NFS in ONTAP.

[TR-4835: How to configure LDAP in ONTAP multiprotocol NAS identity management](#)

Learn how to configure Lightweight Directory Access Protocol (LDAP) identity management in ONTAP for multiprotocol NAS.

NAS security

[TR-4616: NFS Kerberos in ONTAP](#)

Learn about NFS Kerberos in ONTAP including configuration steps with Active Directory and Red Hat Enterprise Linux (RHEL) clients.

ONTAP networking technical reports

ONTAP provides a variety of different networking capabilities and configurations to meet the most demanding scale out applications. Using the networking capabilities and features, companies can create reliable and secure access to their data.



These technical reports expand on the [ONTAP network management](#) product documentation.

[TR-4949: BGP/VIP with ONTAP in the data center](#)

Learn how to quickly deploy a basic BGP configuration in ONTAP.

ONTAP SAN technical reports

ONTAP SAN storage delivers a simplified SAN experience that provides high availability for your organization's mission-critical databases and other SAN workloads. With best-in-class data services integration with Oracle, SAP, and Microsoft SQL Server databases, plus VMware and other leading hypervisors, ONTAP SAN delivers accelerated time-to-value for enterprise database applications.



These technical reports expand on the [ONTAP SAN storage management](#) product documentation.

[TR-4080: Best practices for modern SAN in ONTAP](#)

Learn about block protocols in ONTAP as well as recommendations practices.

[TR-4684: Implementing and configuring modern SANs with NVMe over Fabrics \(NVMe-oF\)](#)

Learn how to implement and configure NVMe over Fabrics transports (NVMe over Fibre Channel and NVMe over TCP). Topics include design, implementation, configuration, management guidelines and recommended practices to build highly available, high-performance modern SAN solutions using NVMe protocols and transports.

[TR-4968: NetApp All-SAN Array data availability and integrity](#)

Learn how the various data protection and data integrity features of a All SAN array systems work to achieve maximum application uptime plus recommended practices for designing, implementing, and managing a SAN network.

[Modern SAN Cloud-Connected Flash Solution](#)

This NetApp Verified Architecture has been jointly designed and verified by NetApp, VMware, andBroadcom. It uses the latest Brocade, Emulex, and VMware vSphere technology solutions along with NetApp all-flash storage, which sets a new standard for enterprise SAN storage and data protection that will drive superior business value.

Security

ONTAP security technical reports

ONTAP continues to evolve, with security as an integral part of the solution. The latest releases of ONTAP contain many new security features that are invaluable for your organization to protect its data across your hybrid cloud, prevent ransomware attacks, and adhere to industry recommended practices. These new features also support your organization's move toward a Zero Trust model.



These technical reports expand on the [ONTAP security and data encryption](#) product documentation.

ONTAP cyber vault

[ONTAP cyber vault](#)

NetApp's ONTAP based cyber vault provides organizations with a comprehensive and flexible solution for protecting their most critical data assets. By leveraging logical air-gapping with robust hardening methodologies, ONTAP enables you to create secure, isolated storage environments that are resilient against evolving cyber threats. With ONTAP, you can ensure the confidentiality, integrity, and availability of your data while maintaining the agility and efficiency of your storage infrastructure.

Ransomware

[TR-4572: The NetApp solution for ransomware](#)

Learn how ransomware has evolved; and how to identify attacks, prevent the spread, and recover as quickly as possible using the NetApp solution for ransomware. The guidance and solutions provided in this document are designed to help organizations have cyber-resilient solutions while meeting their prescribed security objectives for information system confidentiality, integrity, and availability.

[TR-4526: Compliant WORM storage using NetApp SnapLock](#)

Many businesses rely on some use of write once, read many (WORM) data storage to meet regulatory compliance requirements or simply to add another layer to their data protection strategy. Learn how to integrate SnapLock, the WORM solution in ONTAP, into environments that require WORM data storage.

Zero Trust

[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. ONTAP takes 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.

Multifactor authentication

[TR-4647: Multifactor authentication in ONTAP best practices and implementation guide](#)

Learn about ONTAP's multifactor authentication capability for administrative access using System Manager, Active IQ Unified Manager and ONTAP secure shell (SSH) CLI authentication.

[TR-4717: ONTAP SSH authentication with a common access card](#)

Learn how to configure and test third-party SSH clients, in conjunction with ActivClient software, to authenticate an ONTAP storage administrator via the public key stored on a common access card (CAC) when it is configured in ONTAP.

Multitenancy

[TR-4160: Secure multitenancy in ONTAP](#)

Learn how to implement secure multitenancy using storage VMs in ONTAP, including design considerations and recommended practices.

Standards

[TR-4401: PCI-DSS 4.0 and ONTAP](#)

Learn how to validate a system against the PCI DSS 4.0 standard and meet the requirements of the controls that you apply to a NetApp ONTAP system.

Attribute-based access control

[Attribute-based access control with ONTAP](#)

Learn how to configure NFSv4.2 security labels and extended attributes (xattrs) to support role-based access control (RBAC) and attribute-based access control (ABAC), an authorization strategy that defines permissions based on user, resource, and environmental attributes.

NetApp solution for ransomware

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.

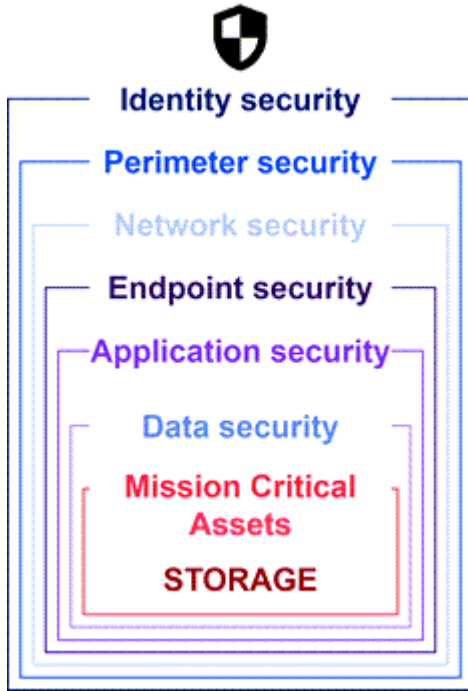
The [NetApp product portfolio](#) provides various effective tools for visibility, detection, and remediation, helping you to spot ransomware early, prevent spread, and recover quickly, if necessary, to avoid costly downtime. Traditional layered defense solutions remain prevalent, as do third parties 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, is available on docs.netapp.com.

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.

<p>Secure by Design Data-centric on-box protection</p>	 Immutable backups & snapshots	 Multi-user verification and authentication	 Malicious file blocking	<p>Ransomware Recovery Guarantee</p> <p>No data loss with NetApp Snapshots, guaranteed.</p>											
<p>Real-time Detection & Response 99% detection accuracy to minimize attack impact</p>	 AI-powered detection	 Actional intelligence for insider threats													
<p>Air-gapped WORM protection with cyber vaulting Layered approach to further fortify data against ransomware attacks</p>	 Isolated, immutable & indelible WORM snapshots														
<p>Single control plane for comprehensive ransomware defense</p> <p style="text-align: right;">BlueXP Ransomware Protection</p> <table border="0" style="width: 100%; text-align: center;"> <tr> <td data-bbox="194 913 251 955"></td> <td data-bbox="373 913 422 955"></td> <td data-bbox="568 913 617 955"></td> <td data-bbox="763 913 820 955"></td> <td data-bbox="966 913 1023 955"></td> <td data-bbox="1161 913 1218 955"></td> </tr> <tr> <td>PROTECT Recommends workload protection policies and applies them with one-click.</td> <td>DETECT Detects potential attacks on your workload data in near real-time using industry leading AI/ML.</td> <td>RESPOND Automatically responds by taking immutable and indelible Snapshots when a potential attack is suspected. Integrates with popular SIEMs.</td> <td>RECOVER Rapidly restores workloads with application consistency, through simplified orchestrated recovery.</td> <td>GOVERN Implements your ransomware protection strategy and policies, and monitors outcomes.</td> <td></td> </tr> </table>											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.
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 [NetApp 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 snapshots 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 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

- [Digital Advisor ransomware protection](#)
- [Data Infrastructure Insights Storage Workload Security](#)
- [FPolicy](#)
- [SnapLock and tamperproof snapshots](#)

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 snapshots 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 snapshots

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 from being deleted even by an ONTAP administrator using the same SnapLock retention expiration period. Snapshot 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 snapshots tamperproof, an [ONTAP One](#) license is required.

Learn more

- [Lock a snapshot 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 the NetApp Console, 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.](#)

Data Infrastructure Insights Storage Workload Security

Storage Workload Security (SWS) is a feature of NetApp Data Infrastructure 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, or data deletion, 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 Data Infrastructure 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 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 in ONTAP

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 snapshots 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 snapshots 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 snapshots are locked and cannot be deleted, even by ONTAP administrators or NetApp Support.

Learn more

- [Blog: ONTAP cyber vault overview](#)

Tamperproof snapshots

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 snapshots 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 on primary storage.

Beginning in ONTAP 9.12.1, tamperproof snapshots provide near SnapLock Compliance level protection for your snapshots on primary storage and in primary volumes. There is no need to vault the snapshot using SnapVault to a secondary SnapLocked volume. Tamperproof snapshots use SnapLock technology to prevent the primary snapshot 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, something you cannot do with a traditional SnapLock Compliance vaulted snapshot.

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

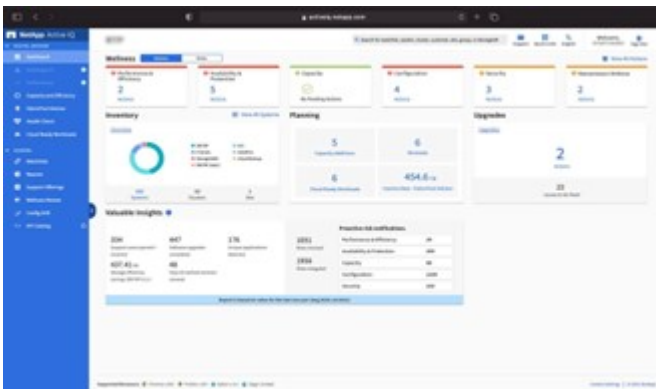
Learn more

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

Digital Advisor ransomware protection

Digital Advisor powered by Active IQ 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 Digital Advisor](#) 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 count is low, decreasing potential ransomware protection.
- FPolicy is not enabled for all storage virtual machines (SVMs) configured for NAS protocols.

To see ransomware protection in action, see [Digital Advisor](#).

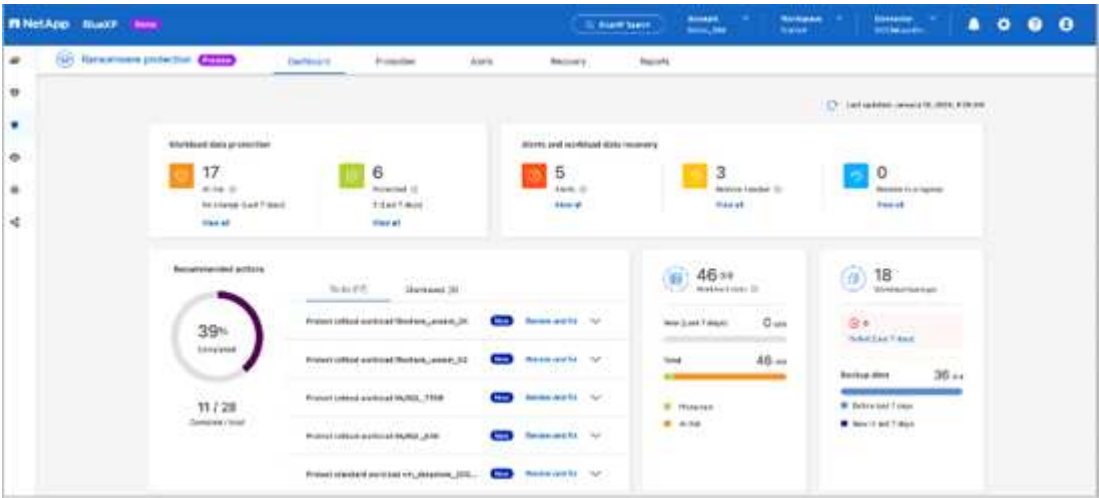
Comprehensive resilience with NetApp ransomware protection

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 the NetApp Console and an isolated, layered solution for cyber vaulting.

NetApp ransomware protection

The NetApp Console is a single control plane to intelligently orchestrate a comprehensive, workload-centric ransomware defense. NetApp ransomware protection brings together the powerful cyber-resilience features of ONTAP, such as ARP, FPolicy, and tamperproof snapshots, and NetApp data services, such as NetApp 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

NetApp 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**.

NetApp and Zero Trust

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, is available on docs.netapp.com.

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 use 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.

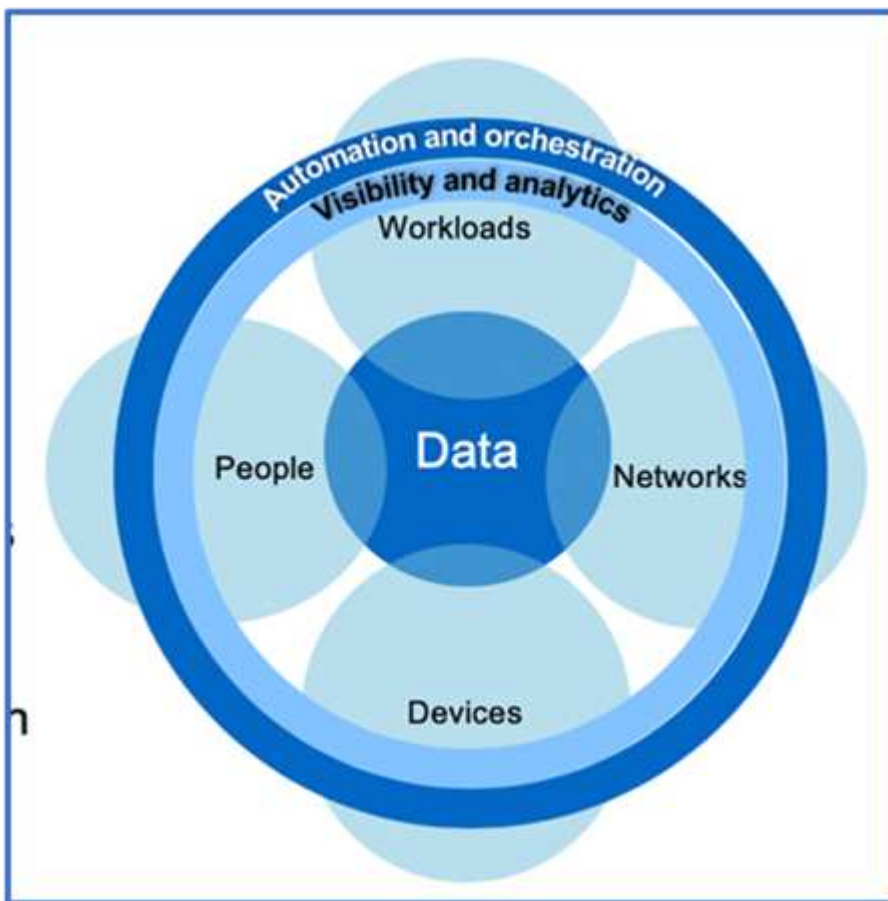
Security resources

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 [NetApp Data 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 snapshots is with ONTAP [snapshot locking](#). Snapshot locking is a SnapLock capability where snapshots are rendered indelible manually or automatically with a retention period on the volume snapshot policy. Snapshot locking is also referred to as tamper-proof snapshot locking. The purpose of snapshot locking is to prevent rogue or untrusted administrators from deleting snapshots on the primary and secondary ONTAP

systems. Rapid recovery of locked snapshots 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.

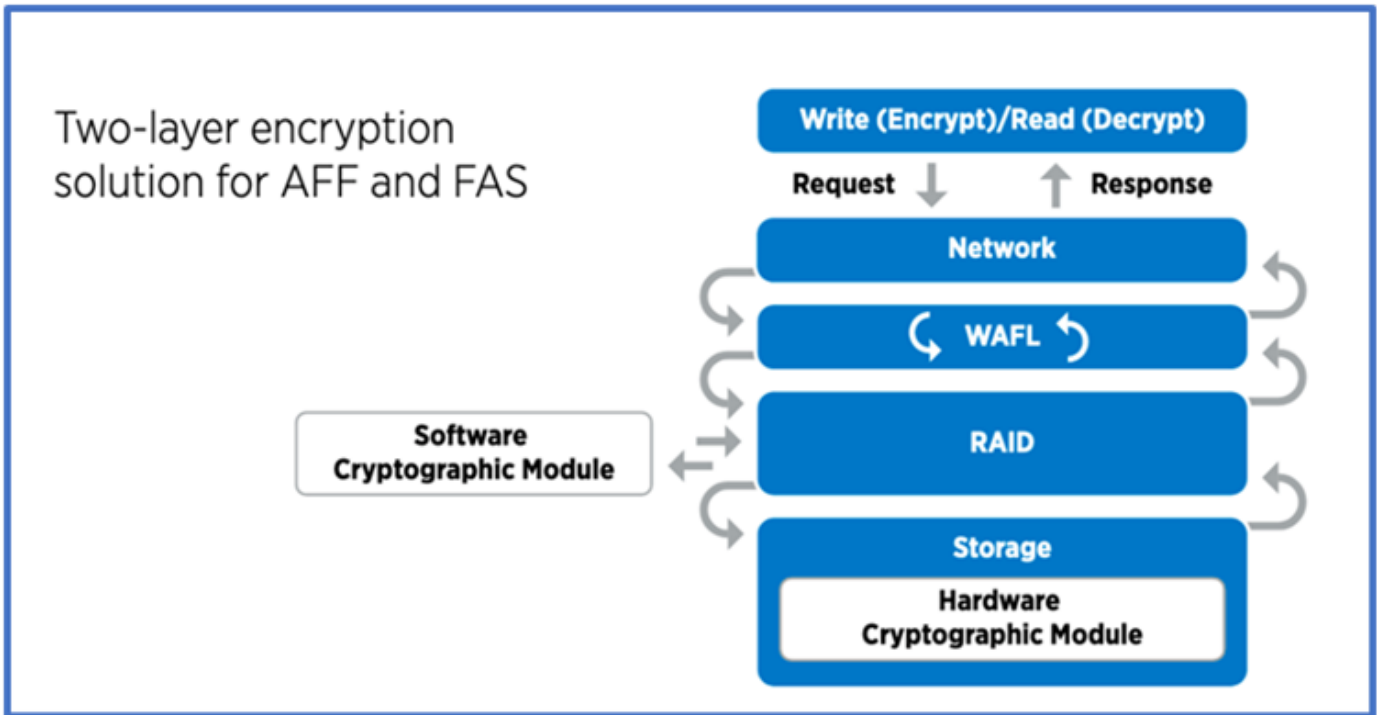
Use encryption for data at rest and data in flight

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 [DII Storage Workload Security](#) 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 snapshots 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 Data Infrastructure 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 and alert to administrators. Storage Workload Security integrates with ONTAP to receive ARP events and provides an additional analytics and automatic responses layer.

Learn more about the commands described in this procedure in the [ONTAP command reference](#).

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.

By 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.

Data Infrastructure Insights is an infrastructure monitoring tool that gives you visibility into your complete infrastructure. With Data Infrastructure Insights, you can monitor, troubleshoot, and optimize all your resources, including your public cloud instances and your private data centers. Data Infrastructure 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 Data Infrastructure 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, 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 enterprise-class, cloud-native file services for AWS (FSxN), Google Cloud (GCNV), 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.

ONTAP enables the movement of data between your on-premises ONTAP systems and AWS, Google Cloud, or Azure storage environment with NetApp SnapMirror data replication software.

Attribute-based access control

Attribute-based access control with ONTAP

Beginning with 9.12.1, you can configure ONTAP with NFSv4.2 security labels and extended attributes (xattrs) to support role-based access control (RBAC) with attributes and attribute-based access control (ABAC).

ABAC is an authorization strategy that defines permissions based on user attributes, resource attributes, and environmental conditions. The integration of ONTAP with NFS v4.2 security labels and xattrs complies with NIST standards for ABAC solutions, as set forth in NIST Special Publication 800-162.

You can use NFS v4.2 security labels and xattrs to assign files user-defined attributes and labels. ONTAP can integrate with ABAC-oriented identity and access management software to enforce granular file and folder access control policies based on these attributes and labels.

Related information

- [Approaches to ABAC with ONTAP](#)
- [NFS in NetApp ONTAP: Best practice and implementation guide](#)

Approaches to attribute-based access control (ABAC) in ONTAP

ONTAP provides several approaches you can use to achieve file-level attribute-based access control (ABAC), including NFS v4.2 security labels and extended attributes (xattrs) using NFS.

NFS v4.2 security labels

Beginning with ONTAP 9.9.1, the NFS v4.2 feature called Labeled NFS is supported.

NFS v4.2 security labels are a way to manage granular file and folder access by using SELinux labels and Mandatory Access Control (MAC). These MAC labels are stored with files and folders, and they work in conjunction with UNIX permissions and NFS v4.x ACLs.

Support for NFS v4.2 security labels means that ONTAP now recognizes and understands the NFS client's SELinux label settings. NFS v4.2 security labels are covered in RFC-7204.

Use cases for NFS v4.2 security labels include the following:

- MAC labeling of virtual machine (VM) images
- Data security classification for the public sector (secret, top secret, and other classifications)
- Security compliance
- Diskless Linux

Enable NFS v4.2 security labels

You can enable or disable NFS v4.2 security labels with the following command (advanced privilege required):

```
vserver nfs modify -vserver <svm_name> -v4.2-seclabel <disabled|enabled>
```

Learn more about `vserver nfs modify` in the [ONTAP command reference](#).

Enforcement modes for NFS v4.2 security labels

Beginning with ONTAP 9.9.1, ONTAP supports the following enforcement modes:

- **Limited Server Mode:** ONTAP cannot enforce the labels but can store and transmit them.



The ability to change MAC labels is up to the client to enforce.

- **Guest Mode:** If the client is not labeled NFS-aware (v4.1 or lower), MAC labels are not transmitted.



ONTAP does not currently support Full Mode (storing and enforcing MAC labels).

NFS v4.2 security labels examples

The following example configuration demonstrates concepts using Red Hat Enterprise Linux release 9.3 (Plow).

The user `jrsmith`, created based on John R. Smith's credentials, has the following account privileges:

- Username = `jrsmith`
- Privileges = `uid=1112(jrsmith) gid=1112(jrsmith) groups=1112(jrsmith) context=user_u:user_r:user_t:s0`

There are two roles: the admin account that is a privileged user and user `jrsmith` as described in the following MLS privileges table:

Users	Role	Type	Levels
<code>admins</code>	<code>sysadm_r</code>	<code>sysadm_t</code>	<code>t:s0</code>
<code>jrsmith</code>	<code>user_r</code>	<code>user_t</code>	<code>t:s1 - t:s4</code>

In this example environment, user `jrsmith` has access to files at the levels `s0` to `s3`. We can enhance the existing security classifications, as outlined below, to ensure that administrators do not have access to user-specific data.

- `s0` = privilege admin user data
- `s0` = unclassified data
- `s1` = confidential
- `s2` = secret data
- `s3` = top secret data

NFS v4.2 security labels example with MCS

In addition to Multi-Level Security (MLS), another capability called Multi-Category Security (MCS) allows you to define categories such as projects.

NFS security label	Value
<code>entitySecurityMark</code>	<code>t:s01 = UNCLASSIFIED</code>

Extended attributes (xattrs)

Beginning with ONTAP 9.12.1, ONTAP supports xattrs. xattrs allow metadata to be associated with files and directories beyond what is provided by the system, such as access control lists (ACLs) or user-defined attributes.

To implement xattrs, you can use `setfattr` and `getfattr` command-line utilities in Linux. These tools provide a powerful way to manage additional metadata for files and directories. They should be used with care, as improper use can lead to unexpected behavior or security issues. Always refer to the `setfattr` and `getfattr` man pages or other reliable documentation for detailed usage instructions.

When xattrs is enabled on an ONTAP filesystem, users can set, modify, and retrieve arbitrary attributes on files. These attributes can be used to store additional information about the file that is not captured by the standard set of file attributes, such as access control information.

There are several requirements and limits for using xattrs in ONTAP:

- Red Hat Enterprise Linux 8.4 or later
- Ubuntu 22.04 or later
- Each file can have up to 128 xattrs
- Xattr keys are limited to 255 bytes
- The combined key or value size is 1,729 bytes per xattr
- Directories and files can have xattrs
- To set and retrieve xattrs, `w` or write mode bits must be enabled for the user and group

Xattrs are utilized within the user namespace and do not carry any intrinsic significance to ONTAP itself. Instead, their practical applications are determined and managed exclusively by the client-side application that interacts with the file system.

Xattr use case examples:

- Recording the name of the application responsible for creating a file
- Maintaining a reference to the email message from which a file was obtained
- Establishing a categorization framework for organizing file objects
- Labeling files with the URL of their original download source

Commands for managing xattrs

- `setfattr` sets an extended attribute of a file or directory:

```
setfattr -n <attribute_name> -v <attribute_value> <file or directory name>
```

Sample command:

```
setfattr -n user.comment -v test example.txt
```

- `getfattr` retrieves the value of a specific extended attribute or lists all extended attributes of a file or directory:

Specific attribute:

```
getfattr -n <attribute_name> <file or directory name>
```

All attributes:

```
getfattr <file or directory name>
```

Sample command:

```
getfattr -n user.comment example.txt
```

Xattr key value pair examples

The following table shows two xattr key value pair examples:

xattr	Value
user.digitalIdentifier	CN=John Smith jrsmith, OU=Finance, OU=U.S.ACME, O=US, C=US
user.countryOfAffiliations	USA

User permissions with ACE for xattrs

An access control entry (ACE) is a component within an ACL that defines the access rights or permissions granted to an individual user or a group of users for a specific resource, such as a file or directory. Each ACE specifies the type of access allowed or denied and is associated with a particular security principal (user or group identity).

Access Control Entry (ACE) required for xattrs

- Retrieve xattr: The permissions required for a user to read the extended attributes of a file or directory. The "R" signifies that read permission is necessary.
- Set xattrs: The permissions needed to modify or set the extended attributes. "a," "w," and "T" represent different examples of permissions, such as append, write, and a specific permission related to xattrs.
- Files: Users need append, write, and potentially a special permission related to xattrs to set extended attributes.
- Directories: A specific permission "T" is required to set extended attributes.

File type	Retrieve xattr	Set xattrs
File	R	a,w,T
Directory	R	T

Integration with ABAC identity and access control software

To fully harness the capabilities of ABAC, ONTAP can integrate with an ABAC-oriented identity and access management software.

In an ABAC system, the Policy Enforcement Point (PEP) and Policy Decision Point (PDP) play crucial roles. The PEP is responsible for enforcing access control policies, while the PDP makes the decision on whether to grant or deny access based on the policies.

In a practical setting, an organization would employ a blend of NFS security labels and xattrs. These are used to represent a variety of metadata, including classification, security, application, and content, which are all instrumental in making ABAC decisions. xattrs, for instance, can be used to store the resource attributes that the PDP uses for its decision-making process. An attribute could be defined to represent the classification level

of a file (for example, "Unclassified", "Confidential", "Secret", or "Top Secret"). The PDP could then utilize this attribute to enforce a policy that restricts users to access only files that have a classification level equal to or lower than their clearance level.



This content assumes that the customer's identity, authentication, and access services include at minimum a PEP and a PDP that act as intermediaries for access to the file system.

Example process flow for ABAC

1. User presents credentials (for example, PKI, OAuth, SAML) to system access to PEP and gets results from PDP.

The PEP's role is to intercept the user's access request and forward it to the PDP.

2. The PDP then evaluates this request against the established ABAC policies.

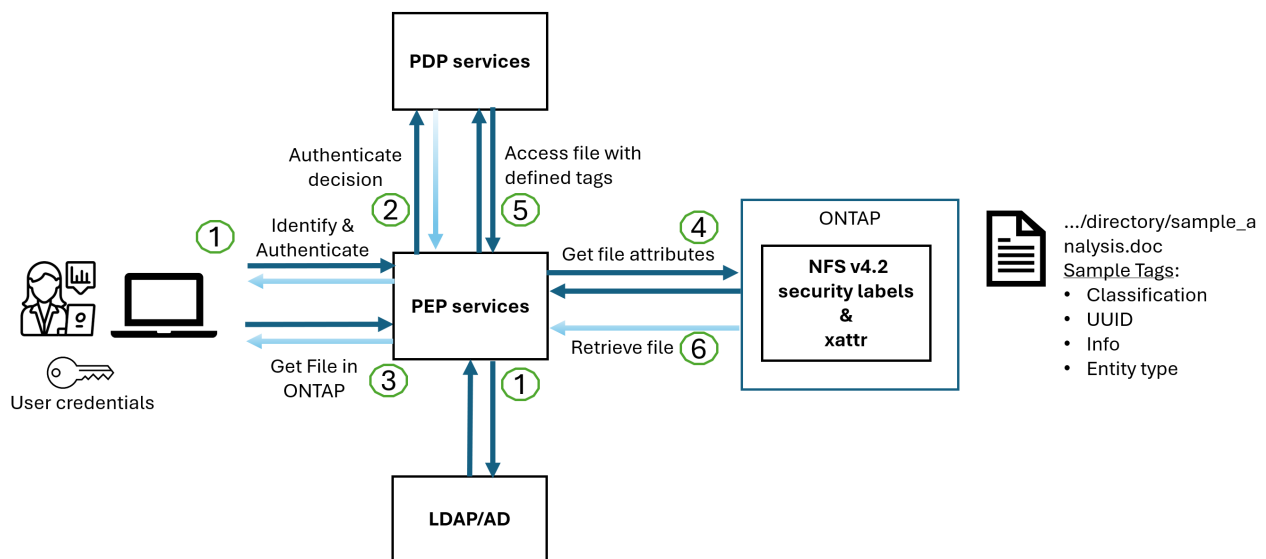
These policies consider various attributes related to the user, the resource in question, and the surrounding environment. Based on these policies, the PDP makes an access decision to either allow or deny and then communicates this decision back to the PEP.

PDP provides policy to PEP to enforce. The PEP then enforces this decision, either granting or denying the user's access request as per the PDP's decision.

3. After a successful request, the user requests a file stored in ONTAP (AFF, AFF-C, for example).
4. If the request is successful, PEP gets fine-grain access control tags from document.
5. PEP requests policy for user based on that user's certs.
6. PEP makes a decision based on policy and tags if the user has access to the file and lets the user retrieve the file.



The actual access might be done using tokens.



ONTAP cloning and SnapMirror

ONTAP's cloning and SnapMirror technologies are designed to provide efficient and reliable data replication

and cloning capabilities, ensuring that all aspects of file data, including xattrs, are preserved and transferred along with the file. xattrs are critical as they store additional metadata associated with a file, such as security labels, access control information, and user-defined data, which are essential for maintaining the file's context and integrity.

When a volume is cloned using ONTAP's FlexClone technology, an exact writable replica of the volume is created. This cloning process is instantaneous and space-efficient, and it includes all file data and metadata, ensuring that xattrs are fully replicated. Similarly, SnapMirror ensures that data is mirrored to a secondary system with full fidelity. This includes xattrs, which are crucial for applications that rely on this metadata to function correctly.

By including xattrs in both cloning and replication operations, NetApp ONTAP ensures that the complete dataset, with all its characteristics, is available and consistent across primary and secondary storage systems. This comprehensive approach to data management is vital for organizations that require consistent data protection, quick recovery, and adherence to compliance and regulatory standards. It also simplifies the management of data across different environments, whether on-premises or in the cloud, providing users with the confidence that their data is complete and unaltered during these processes.



NFS v4.2 security labels have the caveats defined in [NFS v4.2 security labels](#).

Auditing changes to labels

Auditing changes to xattrs or NFS security labels is a critical aspect of file system management and security. Standard file system auditing tools enable the monitoring and logging of all changes to a file system, including modifications to xattrs and security labels.

In Linux environments, the `auditd` daemon is commonly used to establish auditing for file system events. It allows administrators to configure rules to watch for specific system calls related to xattr changes, such as `setxattr`, `lsetxattr`, and `fsetxattr` for setting attributes and `removexattr`, `lremovexattr`, and `fremovexattr` for removing attributes.

ONTAP FPolicy extends these capabilities by providing a robust framework for real-time monitoring and control of file operations. FPolicy can be configured to support various xattr events, offering granular control over file operations and the ability to enforce comprehensive data management policies.

For users utilizing xattrs, especially in NFS v3 and NFS v4 environments, only certain combinations of file operations and filters are supported for monitoring. The list of supported file operation and filter combinations for FPolicy monitoring of NFS v3 and NFS v4 file access events is detailed below:

Supported file operations	Supported filters
<code>setattr</code>	<code>offline-bit</code> , <code>setattr_with_owner_change</code> , <code>setattr_with_group_change</code> , <code>setattr_with_mode_change</code> , <code>setattr_with_modify_time_change</code> , <code>setattr_with_access_time_change</code> , <code>setattr_with_size_change</code> , <code>exclude_directory</code>

Example of an auditd log snippet for a setattr operation:

```
type=SYSCALL msg=audit(1713451401.168:106964): arch=c000003e syscall=188
success=yes exit=0 a0=7fac252f0590 a1=7fac251d4750 a2=7fac252e50a0 a3=25
items=1 ppid=247417 pid=247563 auid=1112 uid=1112 gid=1112 euid=1112
suid=1112 fsuid=1112 egid=1112 sgid=1112 fsgid=1112 tty=pts0 ses=141
comm="python3" exe="/usr/bin/python3.9"
subj=unconfined_u:unconfined_r:unconfined_t:s0-s0:c0.c1023
key="*set-xattr*"ARCH=x86_64 SYSCALL=**setxattr** AUID="jrsmith"
UID="jrsmith" GID="jrsmith" EUID="jrsmith" SUID="jrsmith"
FSUID="jrsmith" EGID="jrsmith" SGID="jrsmith" FSGID="jrsmith"
```

Enabling [ONTAP FPolicy](#) for users working with xattrs provides a layer of visibility and control that is essential for maintaining the integrity and security of the file system. By leveraging FPolicy's advanced monitoring capabilities, organizations can ensure that all changes to xattrs are tracked, audited, and aligned with their security and compliance standards. This proactive approach to file system management is why enabling ONTAP FPolicy is highly recommended for any organization looking to enhance its data governance and protection strategies.

Examples of controlling access to data

The following example entry for data stored in John R. Smith's PKI cert shows how NetApp's approach can be applied to a file and provide fine-grained access control.



These examples are for illustrative purposes, and it is the customer's responsibility to determine the metadata associated with NFS v4.2 security labels and xattrs. Details on updating and label retention are omitted for simplicity.

Example PKI cert values

Key	Value
entitySecurityMark	t:s01 = UNCLASSIFIED

Key	Value
Info	<pre> { "commonName": { "value": "Smith John R jrsmith" }, "emailAddresses": [{ "value": "jrsmith@dod.mil" }], "employeeId": { "value": "00000387835" }, "firstName": { "value": "John" }, "lastName": { "value": "Smith" }, "telephoneNumber": { "value": "938/260-9537" }, "uid": { "value": "jrsmith" } } </pre>
specification	"DoD"
uuid	b4111349-7875-4115-ad30-0928565f2e15
adminOrganization	<pre> { "value": "DoD" } </pre>

Key	Value
briefings	<pre>[{ "value": "ABC1000" }, { "value": "DEF1001" }, { "value": "EFG2000" }]</pre>
citizenshipStatus	<pre>{ "value": "US" }</pre>
clearances	<pre>[{ "value": "TS" }, { "value": "S" }, { "value": "C" }, { "value": "U" }]</pre>
countryOfAffiliations	<pre>[{ "value": "USA" }]</pre>

Key	Value
digitalIdentifier	<pre>{ "classification": "UNCLASSIFIED", "value": "cn=smith john r jrsmith, ou=dod, o=u.s. government, c=us" }</pre>
dissemTos	<pre>{ "value": "DoD" }</pre>
dutyOrganization	<pre>{ "value": "DoD" }</pre>
entityType	<pre>{ "value": "GOV" }</pre>
fineAccessControls	<pre>[{ "value": "SI" }, { "value": "TK" }, { "value": "NSYS" }]</pre>

These PKI entitlements show John R. Smith's access details, including access by data type and attribution.

In scenarios where IC-TDF metadata is stored separately from the file, NetApp advocates for an additional layer of fine-grained access control. This involves storing access control information at both the directory level and in association with each file. As an example, consider the following tags linked to a file:

- NFS v4.2 security labels: Utilized for making security decisions

- xattrs: Provide supplementary information pertinent to the file and the organizational program requirements

The following key-value pairs are examples of metadata that could be stored as xattrs and offer detailed information about the file's creator and associated security classifications. This metadata can be leveraged by client applications to make informed access decisions and to organize files according to organizational standards and requirements.

Example of xattr key-value pairs

Key	Value
user.uuid	"761d2e3c-e778-4ee4-997b-3bb9a6a1d3fa"
user.entitySecurityMark	"UNCLASSIFIED"
user.specification	"INFO"

Key	Value
user.Info	<pre> { "commonName": { "value": "Smith John R jrsmith" }, "currentOrganization": { "value": "TUV33" }, "displayName": { "value": "John Smith" }, "emailAddresses": ["jrsmith@example.org"], "employeeId": { "value": "00000405732" }, "firstName": { "value": "John" }, "lastName": { "value": "Smith" }, "managers": [{ "value": "" }], "organizations": [{ "value": "TUV33" }, { "value": "WXY44" }], "personalTitle": { "value": "" }, "secureTelephoneNumber": { "value": "506-7718" }, "telephoneNumber": { "value": "264/160-7187" }, "title": { "value": "Software Engineer" }, }, </pre>

Key	Value
user.geo_point	[-78.7941, 35.7956]

Related information

```
}  
}
```

- [NFS in NetApp ONTAP: Best practice and implementation guide](#)
- [ONTAP command reference](#)
- Request for comments (RFC)
 - [RFC 7204: Requirements for Labeled NFS](#)
 - [RFC 2203: RPCSEC_GSS Protocol Specification](#)
 - [RFC 3530: Network File System \(NFS\) Version 4 Protocol](#)

Security hardening

ONTAP security hardening guides

These technical reports provide guidance on how to harden NetApp ONTAP as well as other NetApp products.



These technical reports expand on the [ONTAP security and data encryption](#) product documentation.

Hardening guides

[TR-4569: Security hardening guide for NetApp ONTAP](#)

Learn how to configure NetApp ONTAP to help organizations meet prescribed security objectives for information system confidentiality, integrity, and availability.

[Security hardening guide for ONTAP tools for VMware vSphere](#)

Learn how to configure ONTAP tools for VMware vSphere to help organizations meet prescribed security objectives for information system confidentiality, integrity, and availability.

[TR-4957: Security hardening guide for NetApp SnapCenter](#)

Learn how to configure NetApp SnapCenter software to help organizations meet prescribed security objectives for information system confidentiality, integrity, and availability.

[TR-4963: Security hardening guide: NetApp Backup and Recovery for Applications](#)

Learn how to configure NetApp Cloud Backup for Applications to help organizations meet prescribed security objectives for information system confidentiality, integrity, and availability.

[TR-4943: Security hardening guide for NetApp Active IQ Unified Manager](#)

Learn how to configure NetApp Active IQ Unified Manager to help organizations meet prescribed security objectives for information system confidentiality, integrity, and availability.

[TR-4945: Security hardening guide for NetApp Manageability SDK](#)

Learn how to configure NetApp Manageability SDK (NMSDK) to help organizations meet prescribed security objectives for information system confidentiality, integrity, and availability.

[Security hardening guide for MetroCluster Tiebreaker host and database](#)

Learn how to configure the NetApp MetroCluster Tiebreaker host and database to help organizations meet prescribed security objectives for information system confidentiality, integrity, and availability.

ONTAP security 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 the technical report *TR-4569: Security hardening guide for ONTAP*, which was previously published as a PDF, is available on docs.netapp.com.

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

ONTAP 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 with 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, snapshots, 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, except volume moves • Manage quotas, qtrees, snapshots, 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 snapshots • 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, snapshots, 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.



You cannot remove or rename built-in accounts. If an administrator removes the account, upon reboot, the built-in account will be recreated. **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

User/Group Name      Application      Authentication Method      Role Name      Acct Locked      Is-Nsswitch 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.
```

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 snapshots, 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 locking

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

Snapshot locking was introduced in ONTAP 9.12.1. Snapshot locking is also referred to as tamper-proof snapshot locking. Although it does require the SnapLock license and initialization of the compliance clock, snapshot 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 snapshots to a secondary system. Rapid recovery of locked snapshots on primary systems can be achieved to restore volumes corrupted by ransomware.

For more details, see the [snapshot locking 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.

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

All passwords must be distinct from the user name.

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	-
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

Attribute	Description	Default	Range
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

```

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

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	Autonomous ransomware protection is supported on several volumes but not yet configured.
SVMs are not configured for autonomous ransomware protection	Autonomous ransomware protection is supported on several SVMs but not yet configured.
Native FPolicy is not configured	FPolicy is not set for NAS SVMs.
Enable autonomous 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](#).

System Manager session timeout


You can change the System Manager session inactivity timeout. The default inactivity timeout is 30 minutes. A timeout is important to prevent stale sessions and session piggybacking.



If SAML is configured, the inactivity timeout is controlled by settings on the IdP.

Steps

1. Select **Cluster > Settings**.

2. In **UI settings**, select  .
3. In the **Inactivity timeout** box, type a minutes value between 2 and 180 or enter "0" to disable the timeout.
4. Select **Save**.

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 Data Infrastructure Insights Storage Workload Security 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 in ONTAP

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 with 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 with 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 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 in 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.

Starting in ONTAP 9.15.1, cluster peering encryption provides TLS 1.3 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 use 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).
```

Related information

[Prepare to use IP security on the ONTAP network](#)

FIPS mode and TLS and SSL management in ONTAP

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 with 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 TLSv1.2 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 the TLSv1.3 (beginning with ONTAP 9.11.1) and TLSv1.2 protocols are enabled.
- Previous ONTAP releases had the following TLS protocols enabled by default:
 - TLSv1.1 (disabled by default beginning with ONTAP 9.12.1)
 - TLSv1 (disabled by default beginning with ONTAP 9.8)
- 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 ocsp 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 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 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::> vservers cifs create -vservers vs1 -cifs-server CIFS_SERVER1
-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 `vservers 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 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, 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 `vservers 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::> vservers audit create -vservers 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 `vservers 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.



The `/api/storage/volumes` REST API is used for more than just file access. It is used by System Manager, and other GUI interfaces to create, view, and modify volumes.

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>)

For a client to obtain the following access levels	These access parameters must match the client's security type
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-left: 1px solid #ccc; padding-left: 10px; margin-left: 20px;">  <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.

Security characteristics of LIF roles in ONTAP

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.
- LIF roles should never be exposed to the internet.

To learn more about securing LIFs, see [Configure firewall policies for LIFs](#). This page also provides details

about LIF service policies beginning with ONTAP 9.10.1.

To learn more about how to create a new service policy, see the `network interface service-policy create` command in the [Command Reference](#).

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/TCP	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

Service	Port/Protocol	Description
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
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

Port/Protocol	Description
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
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

Port/Protocol	Description
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
7821	NetApp cluster RPC
7822	NetApp cluster RPC
7823	NetApp cluster RPC
7824	NetApp cluster RPC

Storage Systems

AFX

NetApp AFX overview: Learn about NetApp AFX

NetApp AFX is still ONTAP — it's simply a different way to leverage the benefits of ONTAP. NetApp AFX delivers disaggregated architecture for NAS and object workloads while still providing the feature-rich ONTAP software you know and love.

An Evolution of Innovation: NetApp ONTAP

NetApp ONTAP was created in 1992 as a new way to serve NFS workloads to multiple clients, while revolutionizing performance, data resiliency, point in time copies, and more. It originally only supported NFSv2, but as demand grew for newer NFS versions, other data protocols, more data protection features, NetApp ONTAP had to evolve.

Below is an abbreviated timeline of some of the major ONTAP evolutions that have taken place over the past 30+ years.

The evolution of NetApp ONTAP

Decade	Features
1990s	NFSv2, NFSv3, CIFS, Snapshots, WAFL filesystem, AutoSupport, SnapMirror
2000s	SnapVault, SnapLock, NFSv4, Block protocols, FlexVols, FlexClone, GX/Scale out, Deduplication, file clones
2010s	Clustered ONTAP, inline storage efficiencies, NFSv4.1/pNFS, NVMe, RAID-TEC, FlexGroup volumes, volume encryption, AFF, Cloud Volumes ONTAP, QoS, FabricPool, Azure NetApp Files (ANF), Google Cloud NetApp Volumes (GCNV), All SAN Array (ASA)
2020s	SnapMirror Business Continuity, FlexCache, ONTAP S3, IPsec, Autonomous Ransomware Protection, SnapMirror in and out of ANF and GCNV, Amazon FSxN, ASAr2, NetApp AFX you are here

NetApp ONTAP personalities

For years, ONTAP ran as a single unified platform, combining file, block, and object into a single system. This positioned ONTAP as the Swiss army knife of the datacenter – a platform that could do anything you asked it to.

However, as applications evolved and datacenter requirements changed with a series of shifts in the IT industry, demand for platforms that could do specific things grew. As a result, there are now a few different ways to consume ONTAP.

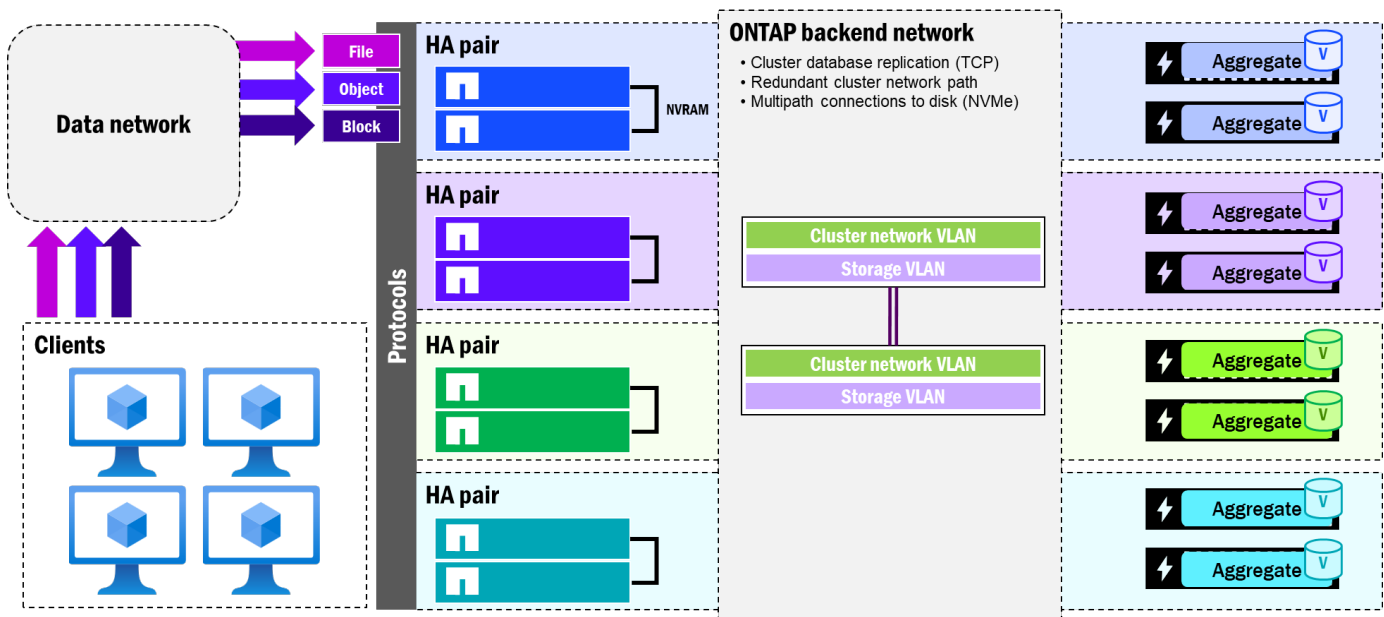
NetApp ONTAP personalities

ONTAP Personality	Description
Unified ONTAP	Same ONTAP you have always known, still being actively developed and improved.
All SAN Arrays	ONTAP that supports only iSCSI, FCP and NVMe over FC/TCP and provides active/active controller functionality, as well as disaggregated concepts in ASAr2.
Cloud-resident ONTAP	ONTAP that runs in the cloud; either on bare metal systems sitting in a cloud datacenter or virtualized ONTAP instances.
Disaggregated ONTAP/AFX	ONTAP using disaggregated architecture for high performance NAS and object workloads.

Unified ONTAP architecture overview

The following diagram shows the general architecture of unified ONTAP, with a description of how it all fits together below it.

General ONTAP architecture



Some key aspects of the ONTAP architecture:

- File, object, and block support
- Protocols are served via a front-end client data network
- Multiple independent nodes can be clustered together
- Each node can provide independent floating IP addresses for data and management use cases
- Clustered nodes connect to a NetApp-provided backend switch over a cluster VLAN
- Nodes are presented as HA pairs to provide resiliency in case of hardware or power outages

- HA pairs have directly connected NVRAM cards that replicate to protect writes
- Each node is assigned at least one aggregate and owns a subset of the total number of disks
- In failovers, a node's disks get reassigned to the HA partner (and only the HA partner)
- Disk shelves are usually directly attached to nodes via multipath cabling, but higher end systems introduce the concept of storage networks on the same backend cluster switch
- Volumes (FlexVols and FlexGroup volumes) provide the entry points to the storage for data access

What is NetApp AFX?

One thing to keep in mind here is that NetApp AFX is still ONTAP.

It's simply a different way to leverage the benefits of ONTAP. The same image you would install for your Unified ONTAP or All SAN Array systems is the same image used with NetApp AFX. The codebase is identical, but the way the systems boot up determines what code path is followed, how the backend storage is presented, and what features and protocols are supported.

NetApp AFX delivers disaggregated architecture for NAS and object workloads while still providing the feature-rich ONTAP software you know and love. Disaggregated ONTAP refers to a storage architecture where every NetApp controller node sees the same capacity by way of redundant network switches and a high speed, low latency network. This approach allows the controller nodes and storage capacity to scale independently of each other to more closely meet the needs of various high-performance workloads. In other words, when you need additional performance, simply add controller nodes. If you need additional capacity, add shelf enclosures. This gives storage administrators the flexibility to present a more cost-effective storage solution to their end users.

Because no controller node independently owns disks, there are also no physical aggregates with their own capacity and performance constraints. Instead, the capacity is presented as a shared model that all nodes can interact with, and ONTAP can automatically manage.

Disaggregated ONTAP

Compute nodes



High Speed, Low Latency Network

Capacity

Key terms and concepts

Below are terms directly related to AFX. For ONTAP-specific terminology, see the product documentation.

[ONTAP product documentation.](#)

Disaggregated ONTAP

Refers to the new architecture powered by ONTAP that offers the ability to scale compute and capacity independently of one another. The term "Disaggregated ONTAP" is not a product name but a means to differentiate between unified ONTAP and NetApp AFX architectures.

NetApp AFX

NetApp AFX is the official product name for the disaggregated ONTAP architecture and was announced at NetApp Insight 2025.

Compute nodes

A compute node in NetApp AFX refers to the storage controller node (and is often used interchangeably in docs). These nodes have no on-board disk and are intended to be fully modular to enable the independent scale that disaggregated ONTAP architecture provides.

Storage Availability Zone

A Storage Availability Zone (SAZ) is the single pool of capacity in a NetApp AFX cluster, where all disks are shared across all nodes. The SAZ enables functionality such as shared capacity, enhanced performance, global deduplication, and more.

What's new in NetApp AFX

This section covers the latest updates for NetApp AFX and will be updated for each new release. Check back here periodically for new information, but also check out the version release notes.

What's new in the latest release of NetApp ONTAP for AFX?

Latest release:

ONTAP 9.19.1RC1 (as of May 2026)

New features in NetApp AFX:

- Global deduplication
- Dynamic storage efficiency
- Advanced readahead
- 32 node support
- 32PB support
- 512 constituent volumes per FlexGroup volume

Version history

Version	Date	Document version history
Version 1.0	June 2026	Initial release

How NetApp AFX architecture differs from unified ONTAP

NetApp AFX introduces significant architectural differences from unified ONTAP in how storage is presented, how nodes interact with disks, and how capacity is managed.

Previously, we showed a general picture of how unified ONTAP architecture delivers file, object, and block data storage via directly connected HA pairs that own their own sets of disks and present physical capacity via aggregates of disks. In this section, we'll discuss some of the major differences between unified ONTAP and NetApp AFX architectures in further detail.

How to tell if a system is running NetApp AFX

The main way to see if your system is running NetApp AFX is to run the following command:

```
AFX::> node show -fields personality
node           personality
-----
afx-01         AFX
afx-02         AFX
```

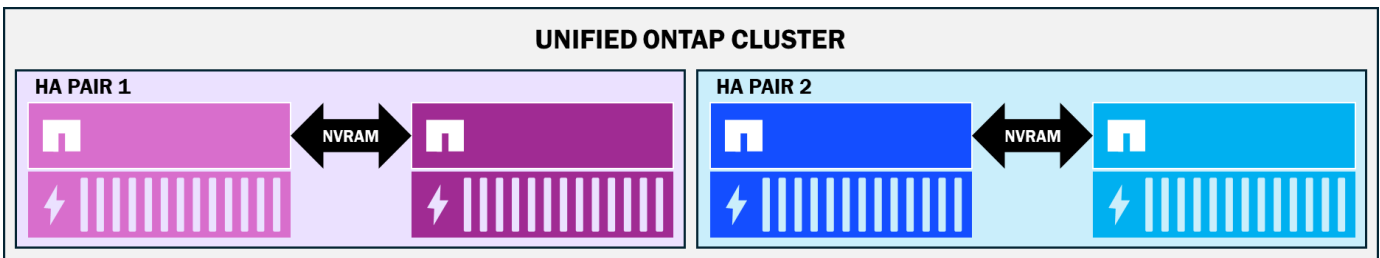
Another clue is the new Storage Availability Zone, but that is also a concept available to NetApp All-SAN

Arrays (ASA). You can see your capacity via that command.

```
AFX::> storage availability-zone show
Availability Zone Name: storage_availability_zone_0
Availability Zone UUID: 545cb59f-32e9-11f1-a2f5-
d039eabdd925
Total Size: 69.59TB
Physical Used: 837.1GB
Physical Used Percent: 1%
Available: 68.77TB
Metadata Used: 837.1GB
Log and Recovery Metadata: 834.6GB
Delayed Frees: 2.50GB
Physical User Data Without Snapshot Copies: 17.24MB
Logical User Data Without Snapshot Copies: 17.24MB
Efficiency Ratio Without Snapshot Copies: 1.00:1
Space Full Threshold Percent: 98%
Space Nearly Full Threshold Percent: 95%
```


Node-to-disk relationships

In the unified ONTAP architecture, reads and writes get directed to a specific subset of disks. So even if you have 24 shelves of disks in a 24-node cluster (one shelf per node), at any given time each node can only directly access one shelf of disks, which limits the capacity and performance available in the cluster.



In addition, because NVRAM is directly connected between HA pairs, nodes must physically reside next to one another and are more tightly coupled as failover targets. For instance, when one node fails over to its partner node, the only disks it physically has access to are the disks in the HA pair domain.

Unified ONTAP cluster during HA failover




In NetApp AFX, there are some major shifts in how disks are presented to compute nodes.

All disks are visible to all storage nodes—no disk ownership

In NetApp AFX, nodes and shelves are all connected to the same backend switch, which makes it possible for ONTAP to extend the overall visibility domain for disks to the entire stack. As a result, no nodes own any specific disks. Rather, all disks participate in a single capacity pool called a Storage Availability Zone, which provides simpler capacity management and increased performance potential (more available disks means more available performance).

NetApp AFX Storage Availability Zone

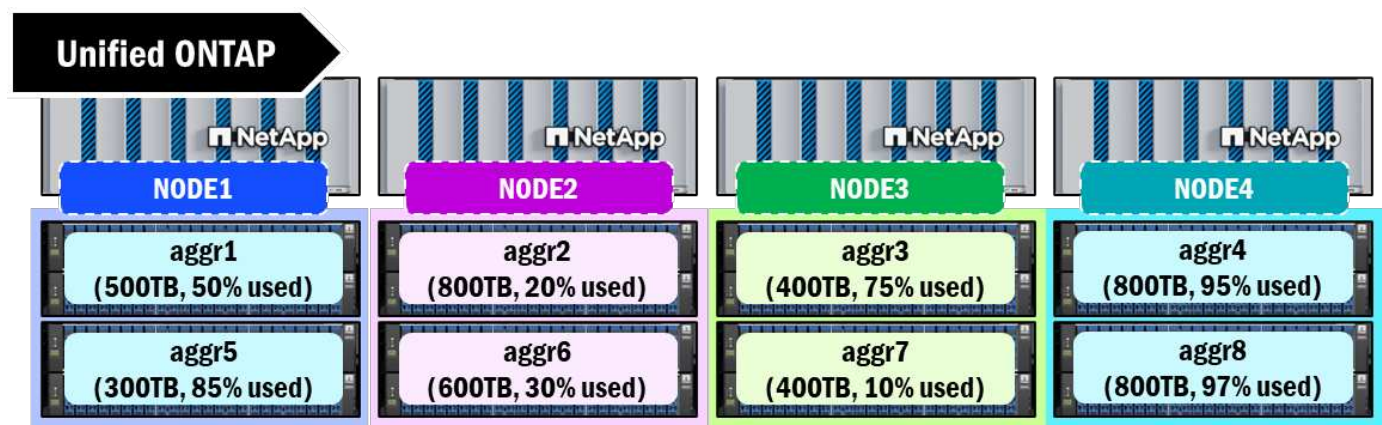


No more physical aggregates

Unified ONTAP collects disks into RAID groups and then combines them into a capacity construct known as an aggregate. This aggregate is how physical capacity is presented to the storage and is the boundary of space available for creating volumes to serve data to end users. Every node must have at least one aggregate assigned to it and these aggregates have a current limit of 800TB. Once that limit is reached, there is no more space available for additional writes.

Physical aggregates can also present some capacity management challenges, as storage admins will sometimes need to manually shuffle volumes around to keep a balance of capacity across cluster nodes. These challenges can also become magnified when leveraging a scale-out volume architecture (such as a FlexGroup volume). Aggregates can also vary in size, disk counts, disk types, etc., which can also creates some performance differences as you traverse nodes.


Aggregates in unified ONTAP



NetApp AFX takes the concept of a physical aggregate and virtualizes it, makes it ONTAP-managed, and then

moves physical capacity management from a per-node methodology to per-cluster via the new Storage Availability Zone. This single pool of capacity provides a "what you see is what you get" approach to space management.

NetApp AFX Storage Availability Zone



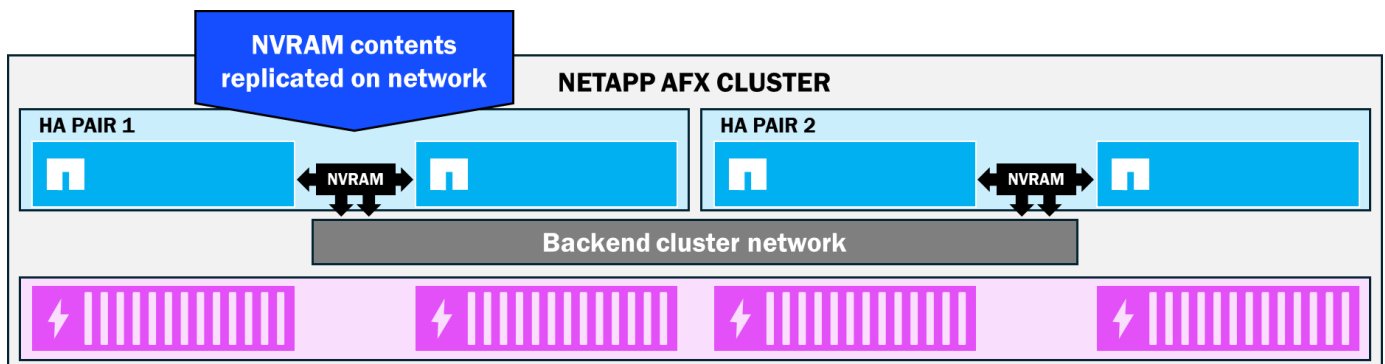
NVRAM moved from direct connection to switched replication

ONTAP uses NVRAM as a staging to protect incoming writes to a cluster. Each node in an ONTAP cluster has a battery-backed NVRAM card. When a write is sent to a volume from a client, it is stored in NVRAM first. NVRAM contents are then flushed to disk when the NVRAM is filled or when a 10s timer expires (whichever comes first). This is known as a consistency point.

The NVRAM contents are also constantly replicating between HA pairs, which further helps protect data consistency, because in the event of a node failure, the NVRAM contents will be preserved on the surviving node and be committed to disk.

In unified ONTAP clusters, the NVRAM cards between HA pairs are connected directly to one another. NetApp AFX moves NVRAM replication to the backend cluster network. As a result, HA partner nodes don't have such a strict distance requirement for nodes. Instead, HA pairs can be separated up to the maximum distance of ethernet.

NetApp AFX NVRAM replication

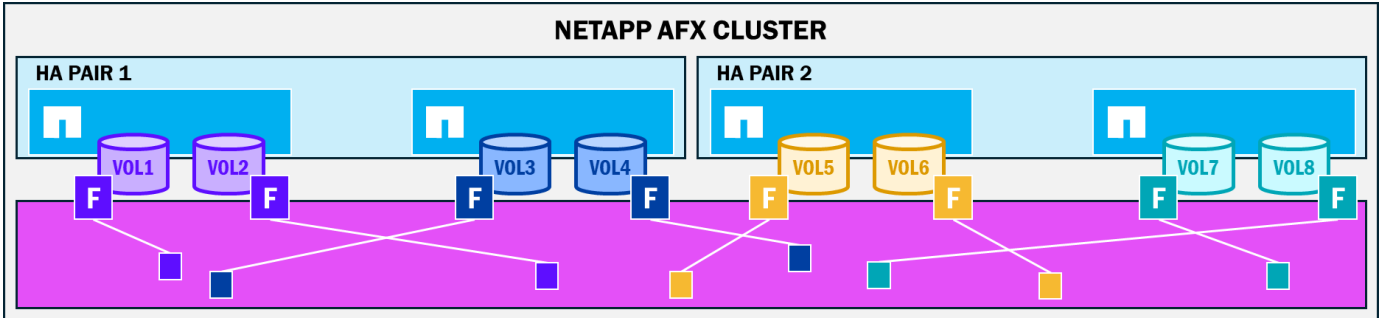


Data written to any (and all) disks in the availability zone

NetApp AFX removes the concept of disk ownership and moves the physical aggregate construct to a virtualized approach managed by ONTAP, where the capacity purchased for the cluster is all available to the nodes attached to the cluster. With AFX, all nodes have the ability to write to any and all disks in the Storage Availability Zone, regardless of what the node:volume ownership is. Nodes still have a concept of volume

ownership since writes still have a path through the NVRAM, but that data can land anywhere in the available capacity. This means that a higher number of disks can participate in a single workload, which provides performance benefits.


How data lands in a Storage Availability Zone



Independent scale of capacity and compute nodes

With hardware resources decoupled in the NetApp AFX architecture, nodes no longer require associated disks to be added alongside one another. When a cluster is running low on performance-related resources such as RAM, CPU, or network throughput, only storage nodes need to be added to the cluster and can leverage the existing Storage Availability Zone. Conversely, if the requirement is capacity, only shelves would need to be added to the mix. This flexibility ensures that you only buy the resources you will need, thus avoiding overprovisioning.


NetApp AFX – Independent scale



Linear scaling of node performance

As nodes are added to an AFX cluster, more CPU, RAM, and network resources are introduced into the workload. As these resources are incorporated into the environment, performance increases are linear in nature. The graphic below shows how that performance would increase as nodes are added.

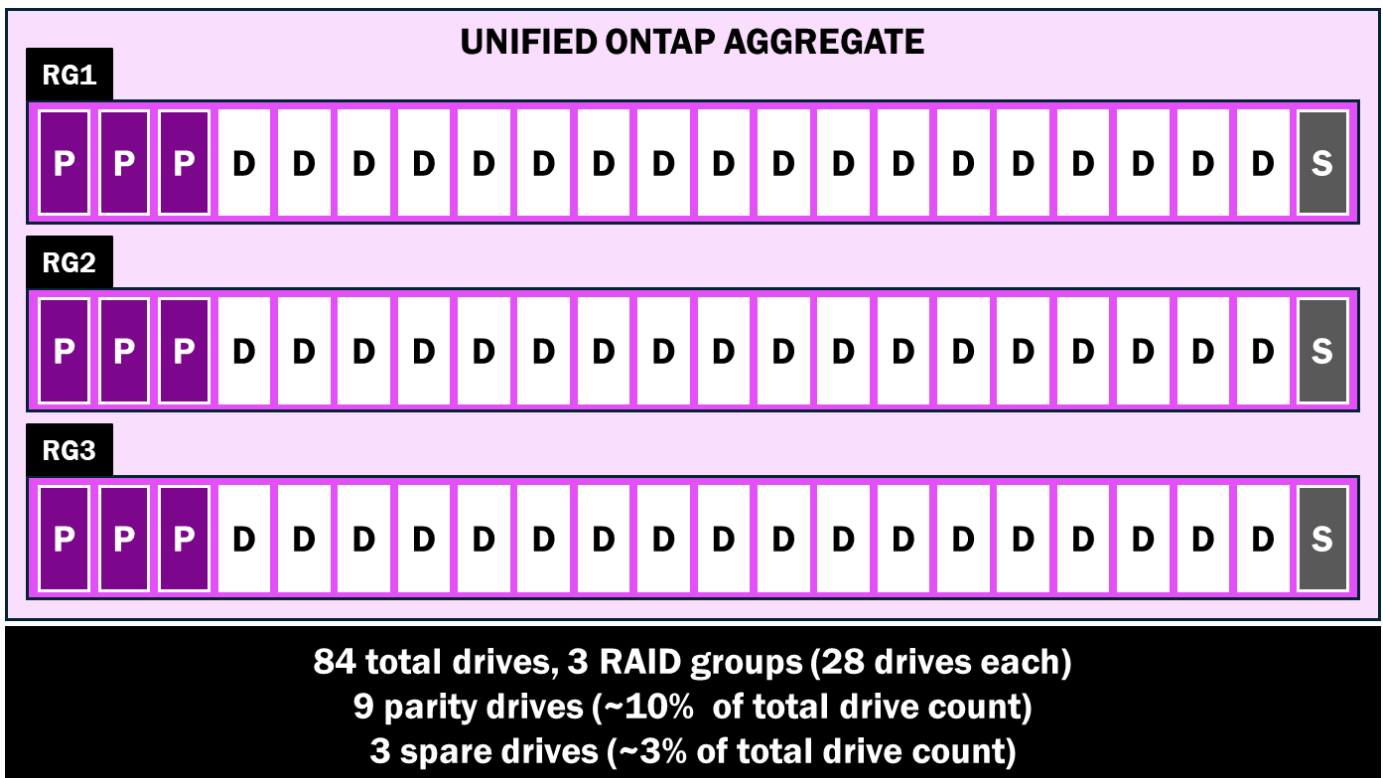
Linear performance increases with NetApp AFX node additions



Larger RAID groups, fewer parity drives

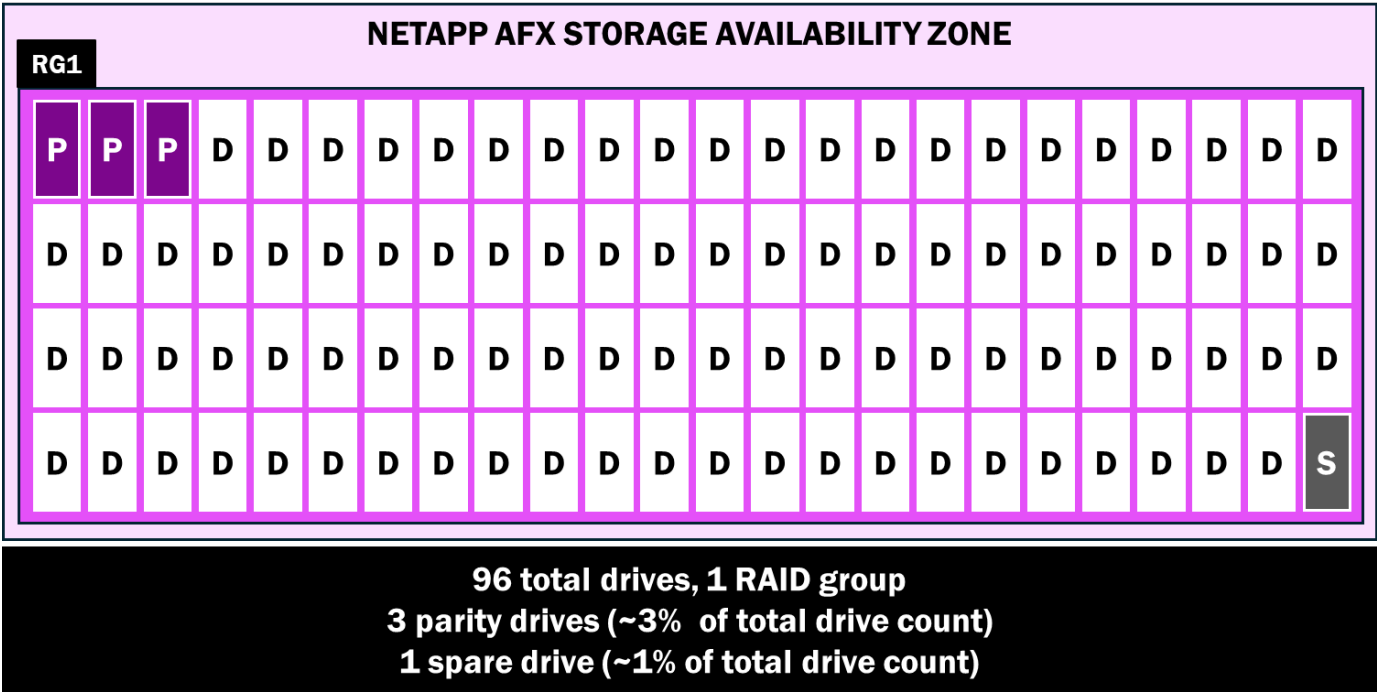
ONTAP provides a blend of data protection and performance for disks via RAID groups – specifically RAID-TEC, which offers triple parity protection in the event of disk failures. RAID-TEC can survive up to three simultaneous drive failures in a RAID group. In unified ONTAP, RAID groups have a maximum disk count of 28, where 3 drives are spent on parity and 1 drive is reserved as a spare. As a result, 24 of the 28 drives are used for data operations/RAID stripes.

Unified ONTAP RAID groups



NetApp AFX still leverages RAID-TEC, but increases the RAID group size to 96 drives while only requiring 3 parity drives and 1 spare. Larger RAID groups deliver greater overall performance, while drive failure exposure is minimized by a combination of low failure rates for SSD, more evenly distributed operations across a larger set of drives, as well as improvements to rebuilds of data drives from parity in NetApp AFX.

NetApp AFX Storage Availability Zone RAID group



The following table approximates the amount of usable raw capacity for 84 disks in unified ONTAP and NetApp AFX with varying drive sizes.

Approximate raw capacity comparison, 84 drives – Unified ONTAP and NetApp AFX

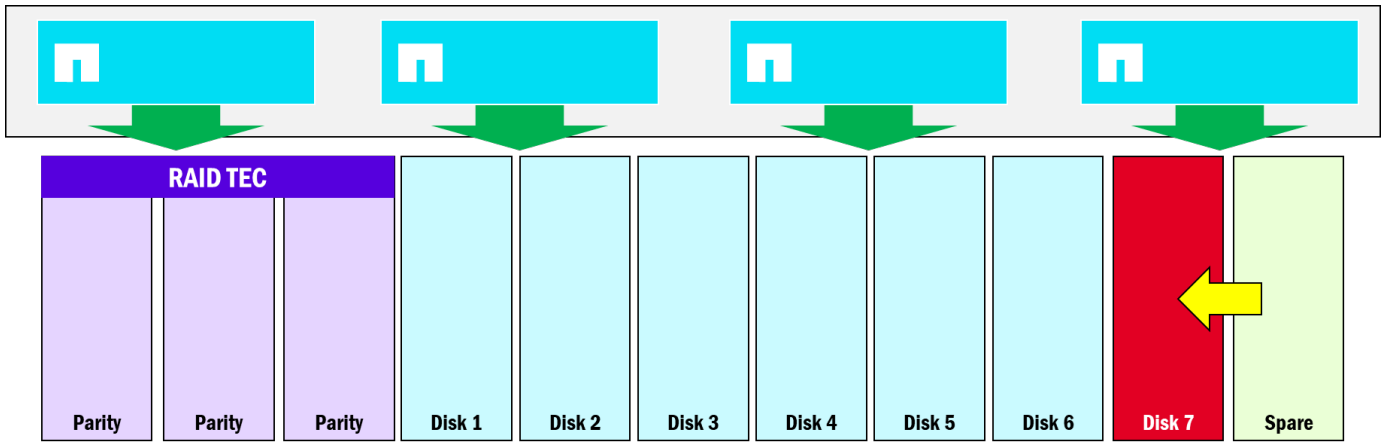
Drive size	Approximate raw capacity (Unified)	Approximate raw capacity (AFX)
7.6TB	~547.2TB	~608TB (+60.8TB)
15.3TB	~1101.6TB	~1224TB (+122.4TB)
30.6TB	~2203.2TB	~2448TB (+244.7TB)
60.1TB	~4327.2TB	~4808TB (+480.8TB)

Faster disk failure rebuild times

In unified ONTAP, each node owns a subset of disks in the storage stack. This means that that node only writes to those disks, but also that disk rebuilds are only handled by a single node in the event of a disk failure.

NetApp AFX eschews the need for disk ownership. As a result, all drives can be written to from a single node if needed. That also means that when a drive needs to be rebuilt from parity, all nodes in the cluster participate, so drive rebuilds can occur more quickly than if a single node had to do it alone.

Disk rebuilds in NetApp AFX




Deduplication domains

Deduplication allows a storage system to find duplicate blocks in its file system and then create pointers to a single block to reduce the total amount of used capacity. In unified ONTAP, deduplication follows a specific boundary for what blocks can be reduced. Those boundaries depend on the type of deduplication in use. In general:

- Volume-based deduplication → Volume boundary
- Cross-volume deduplication → Aggregate boundary

Unified ONTAP deduplication domains



The table below shows the capacity behaviors for duplicated data in different scenarios in unified ONTAP. As file copies span nodes and aggregates (and thus, deduplication domains), space savings get reduced.


Deduplication behaviors in different scenarios for identical 10GB files – unified ONTAP

Scenario	Space used
Four copies of same 10GB file, same volume (volume deduplication)	10GB
Four copies of same 10GB file, different volumes, same aggregate (cross-volume deduplication enabled)	10GB
Four copies of same 10GB file, 4 different volumes, 4 different aggregates (cross-volume deduplication enabled)	40GB

Because NetApp AFX removes physical aggregates and moves capacity management to the new Storage Availability Zone, the deduplication domain boundaries also change. In AFX, the deduplication domain is at the volume-level (like unified ONTAP) and the node (rather than the aggregate) prior to 9.19.1.

Starting in ONTAP 9.19.1, AFX supports a global deduplication domain at the Storage Availability Zone level, so all duplicated blocks in the cluster storage pool are treated the same.

NetApp AFX – Global deduplication domain (ONTAP 9.19.1)



The table below shows the capacity behaviors for duplicated data in different scenarios in NetApp AFX.

Deduplication behaviors in different scenarios for identical 10GB files – NetApp AFX

Scenario	Space used
Four copies of same 10GB file, same volume (volume deduplication)	10GB (9.18.1) 10GB (9.19.1)
Four copies of same 10GB file, different volumes, same node (cross-volume deduplication enabled)	10GB (9.18.1) 10GB (9.19.1)
Four copies of same 10GB file, 4 different volumes, 4 different nodes (cross-volume deduplication enabled)	40GB (9.18.1) 10GB (9.19.1)

Features that have been removed/are unsupported

NetApp AFX is designed for high-performance NAS and object workloads – particularly (but not exclusive to) those in the AI training and inference space. With the design of NetApp AFX, some decisions were made to disable some of the features in ONTAP.

- Because of the focus on high-performance NAS and object, block workloads have been removed from the NetApp AFX solution. There is no support for FCP, iSCSI or NVME data protocols and there are no plans to add block protocols.
- Disaggregated is synonymous to de-aggregated, which means that aggregates (at least as a physical storage administration concept) have been removed. Removal of the physical aggregate not only simplifies the management of capacity in ONTAP, but it also provides the mechanism to allow for a single pool of capacity.
- Aggregates being removed means that aggregate-specific features are also removed. Metrocluster, for instance, leverages aggregate-level mirroring for its site failover capabilities. As such, Metrocluster is also removed from NetApp AFX. Site failover functionality will instead be provided by the new SnapMirror Active-Sync for NAS feature offered in ONTAP 9.19.1GA.
- The cold data tiering feature called FabricPool is also currently unavailable to NetApp AFX since it is also

aggregate-specific.

- Copy-based volume moves are also no longer necessary in NetApp AFX, due to the new capacity architecture. For more information, see [Zero-copy volume moves](#).
- Feature removal also means some CLI/GUI/REST API changes, so any commands or API calls for features no longer supported will also be removed.
- ZAPI is currently unavailable to NetApp AFX.
- NFS copy offload for virtualization (FlexGroup volumes with Granular Data Distribution only)

ONTAP management changes

In general, NetApp AFX management doesn't change the mechanisms used to manage a cluster. Admins can still leverage the CLI, GUI, and REST APIs to log into and configure a cluster. But NetApp AFX did present an opportunity to improve upon some of how storage management operations are conducted.

Simpler capacity management

The NetApp AFX Storage Availability Zone reduces the management endpoints from a node and aggregate-based approach to a single pool of capacity available to the entire cluster. As volumes grow and shrink, ONTAP automatically borrows and returns capacity to and from the Storage Availability Zone.

Because of this, storage administrators no longer need to worry about locating and managing available free space across up to 24 nodes and potentially hundreds of aggregates. Instead, there is only one place where capacity is managed and viewed.

For instance, in the CLI of unified ONTAP, if you wanted to see total physical capacity information for a cluster, you would use "aggregate show-space", which would then print out every aggregate entry. In NetApp AFX, you have "cluster space show", which will show only the single Storage Availability Zone.

Side by side comparison of capacity CLI commands in unified ONTAP and NetApp AFX

```
unified::> aggr show-space
Aggregate : aggr1

Feature
-----
Volume Footprints      250.2TB      50%
Aggregate Metadata     31.06MB      0%
Snapshot Reserve       8.41GB       5%
Total Used              1.2TB        95%

Total Physical Used    225.2TB      50%
Total Provisioned Space 450.2TB      90%
.....
8 entries were displayed.
```

Unified ONTAP

```
AFX::> cluster space show
Availability Zone Name: storage


Total Size: 3PB
Physical Used: 2PB
Physical Used Percent: 75%
Available: 1PB
Metadata Used: 1.71TB
Log and Recovery Metadata: 1.71TB
Delayed Frees: 1.22GB
Physical User Data Without Snapshot Copies: 3.33MB
Logical User Data Without Snapshot Copies: 3.33MB
Efficiency Ratio Without Snapshot Copies: 1.00:1
Space Full Threshold Percent: 98%
Space Nearly Full Threshold Percent: 95%

1 entry was displayed.
```

AFX

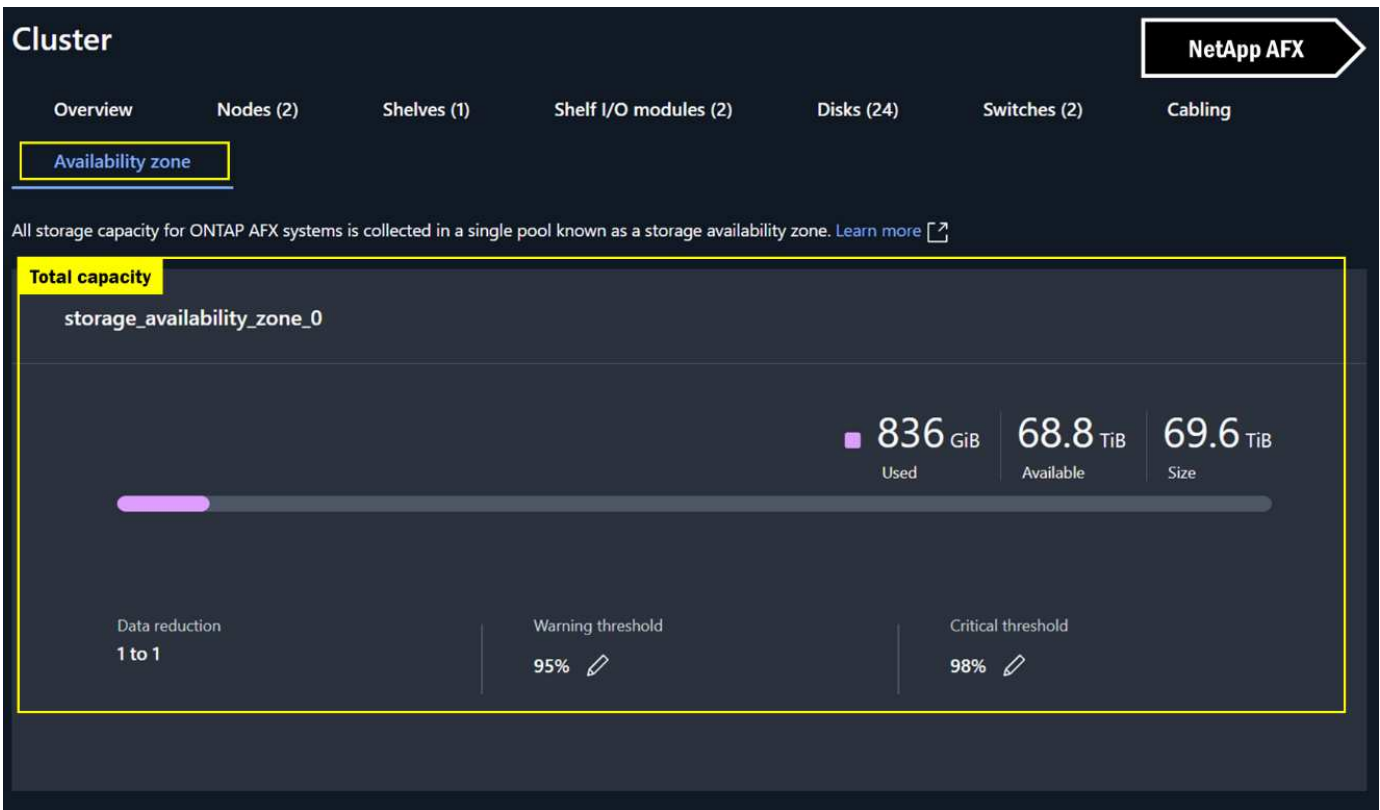
In the Unified ONTAP System Manager GUI, tiers are used to show capacity. In fact, the GUI does attempt to show holistic capacity for the cluster by adding up totals, but it still will show the overall usage on a per-aggregate basis.

System Manager capacity views – Unified ONTAP



In NetApp AFX System Manager, the view is virtually the same for cluster space, but because there are no aggregates, there's no extra math to be done. The capacity you see is the capacity you get.

System Manager capacity views – NetApp AFX




FlexGroup volume management improvements

A FlexGroup volume consists of multiple underlying FlexVol constituent volumes created across multiple nodes and aggregates in the cluster and presented as a single large namespace to NAS clients. FlexGroup volumes provide performance, scale, load balancing, and file count benefits to high performance workloads. However, because they coordinate across nodes and aggregates, they occasionally run into some physical limitations when capacity starts to fill, as the independent file systems provided by aggregates also have independent capacity usage and limits. For instance, if an aggregate with FlexGroup volume constituents starts to fill up before other aggregates in the cluster, then the entire FlexGroup itself could be subject to capacity or performance issues.

As a result, storage administrators may find themselves worrying too much about underlying FlexGroup infrastructure and become less focused on maintaining other aspects of the environment.


FlexGroup volume layout - Unified ONTAP aggregates



NetApp AFX presents capacity in a single Storage Availability Zone, which more closely mirrors the way FlexGroup volumes are intended to work. Instead of multiple constituent volumes across multiple disparate aggregates of potentially varying sizes, all volumes reside in the same pool of capacity, which greatly simplifies the overall management overhead of using a FlexGroup volume.

In addition, AFX enables Advanced Capacity Balancing by default for FlexGroup volumes, which helps better distribute larger files in the volume. Now, FlexGroup volume constituents become less of a management concept and instead do their work quietly in the background.

FlexGroup volume layout - NetApp AFX



Automated storage administration tasks

With the Storage Availability Zone in NetApp AFX, all capacity is shared across all nodes. While nodes still own volumes, ONTAP manages the capacity usage of each node automatically by borrowing and releasing capacity based on what each node needs at any given time. That means storage administrators no longer need to be concerned with how best to balance usable space.

Additionally, RAID group management is automated by ONTAP, where newly added disks will be added to existing or new RAID groups without admin intervention. ONTAP also manages volume moves across nodes without needing to copy data.

Zero-copy volume moves


Unified ONTAP provides a way to nondisruptively move volumes across nodes or aggregates as a way to manage performance and capacity usage across the cluster.

When a volume move is started, the following happens:

- A new empty volume is created on the specified destination aggregate
- Volume metadata (such as storage efficiency information, file handles, etc.) are replicated to the new destination volume
- Volume data is replicated to the destination volume over the backend cluster network via SnapMirror technology—destination aggregate needs to have available free space for the move, or the move job will fail
- Volume replication happens again to ensure both volumes are consistent with any data changes
- A cutover process is initiated to bring the origin volume offline and promote the destination volume as the new origin volume for clients
- Client IO experiences a brief pause during cutover, but no remounts are required

In NetApp AFX, the Storage Availability Zone presents all capacity to all nodes, and all nodes can write to any disk in that pool. Once the data is placed, it stays where it lands – even if the volume is moved. That means no data copy is required. The volume move process is identical to unified ONTAP, minus the need to replicate data via SnapMirror. No extra capacity required.

Zero copy volume moves in NetApp AFX



Having lightweight volume mobility allows AFX to automate away many of the administration tasks without the performance or capacity constraints, and these volume moves are used in a few new features provided by NetApp AFX, as described in topics seen below.

HA failover behavior


In unified ONTAP, nodes own disks and aggregates, where data is served via volumes. Writes are performed using a local node's NVRAM to flush to the disks the node owns. When a node is rebooted or fails, ONTAP will trigger a takeover of the failed node's resources, where disk and aggregate ownership is transferred to the partner node. Network interfaces are also failed over to ports in the IP space, and since NVRAM contents are constantly replicating across the HA pair, the node will flush the NVRAM contents to commit the failed node's writes to disks. After that, the surviving node will own the failed node's aggregates and volumes until giveback of the node occurs. That means that all traffic to those volumes – as well as the volumes already owned by the surviving node – will be processed on a single node until the failover issue is resolved.

As part of the initial unified ONTAP cluster deployment, it is recommended to plan ahead for failovers to help

avoid a single node overloading its partner. That in itself presents a challenge, as it is difficult to predict what volumes might be performance bullies, but features such as nondisruptive volume move and volume quality of service policies can aid in mitigation.


The images below show how unified ONTAP clusters can incur uneven performance balance across nodes, as well as how a failover can create performance degradation in some cases.

Unified ONTAP – potential imbalances in node utilization



When an HA pair's nodes become imbalanced with volume count and performance utilization, node failovers will impact overall performance, since the surviving node will now own all of the failed node's volumes. Meanwhile, other nodes in the cluster may have room to take on additional work.

Unified ONTAP – Failover impact on node utilization



In the above, when an HA partner has to take on additional work, it can potentially become overloaded and impact performance for all volumes on that node. Volume moves can help alleviate the situation, but those require copies across nodes (which requires available free space), and the time that takes may exceed the time it takes for the nodes to fail back. Additionally, if you relocate a volume, it will not fail back to the original node. Instead, it will remain on the node you moved it to.


With NetApp AFX, node failovers take on some different behaviors.

- Since nodes don't own disks and there are no physical aggregates, a node failover won't require transfer of those resources. Instead, only the network interfaces and volume ownership transfer to other nodes.
- NVRAM commits still happen, but over the HA network instead of a direct connection.
- Once volumes conduct the initial failover to the partner node, AFX will re-distribute the volumes across other surviving nodes in the cluster. This is made possible by zero-copy volumes moves.
- When the node is recovered, the volumes will move back to the original node.

NetApp AFX already maintains performance balance across nodes in the cluster to keep a relatively even

utilization, so when a failover happens and volumes are rebalanced, the node utilization should be roughly the same across the cluster.

NetApp AFX - Volume rebalance after failover



Node additions and removals

Both unified ONTAP and NetApp AFX allow for nodes to be added and removed from the cluster. However, due to some of the architectural differences, the process for node additions and removals differ a bit.

Node addition/removal in unified ONTAP

We've already learned that unified ONTAP has a direct node to disk ownership and that all nodes must have some disks and at least one aggregate attached to them. With that in mind, the following holds true for additions and removals.

- Node additions in unified ONTAP don't require any extra steps, but to provide balanced performance across all nodes (including the new nodes), volumes would need to be moved to the new nodes. This requires upfront analysis of the existing volumes and their workloads, decisions about which volumes to move, and then the actual volume moves taking place, which, again, would require a copy of that data across the backend cluster network.
- Node removals in unified ONTAP would require a manual evacuation of existing volumes on the node, which means you must identify which nodes can host which volumes to maintain even performance, and you must have enough free capacity to provide a place for those volumes to move. If free capacity is a challenge, additional volume moves may be needed to shuffle workloads around a bit in the cluster. Node removals also HA pair removals, so the work involved is doubled. Because nodes own disks, an entire disk reinitialization would also be required for those nodes. Each of these things add time and effort to what should be a relatively simple task.

Node addition/removal in NetApp AFX

We've also learned that NetApp AFX doesn't leverage the standard node to disk ownership and does not use physical aggregates to present capacity to the cluster. Because of that, node additions and removals behave a bit differently.

- Node additions in NetApp AFX will not require the same up-front volume analysis, nor will it require administrative intervention to ensure that each node has an even balance of volumes. Instead, ONTAP automatically balances volume counts across newly added nodes to maintain relatively even performance profiles. ONTAP automatically will move volumes across nodes without copying anything, reducing the time, capacity, and effort required for adding nodes to a cluster.
- Node removals in NetApp AFX also don't require much – if any – manual intervention. When a node is

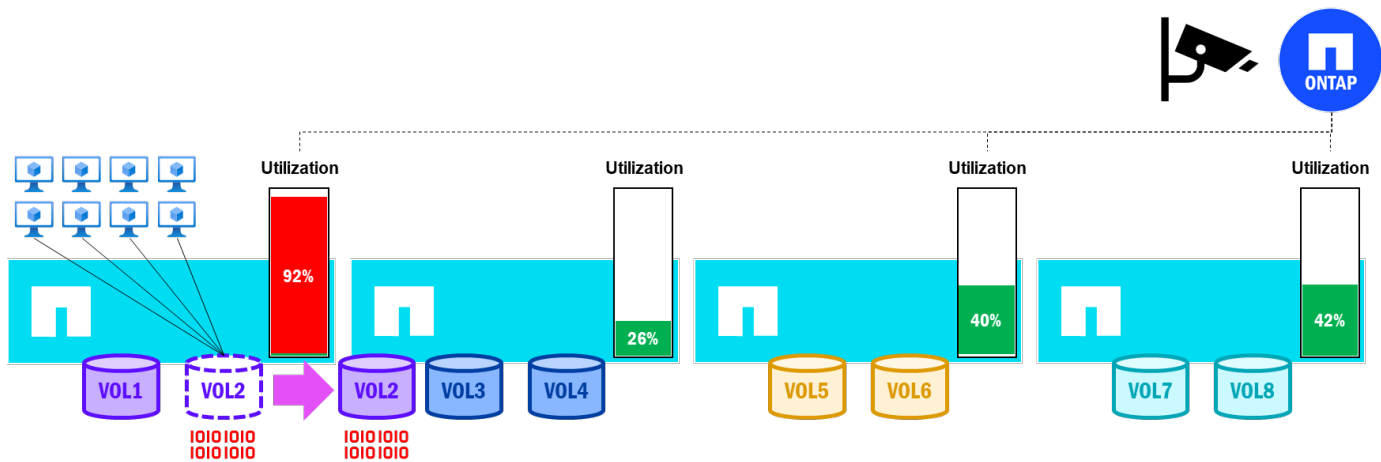
tagged for removal, ONTAP automatically moves volumes across nodes (again, without copying) to evacuate the nodes being removed. And because there are no disks owned by the nodes, there is no need to reinitialize disks after removing nodes. This makes nodes in AFX modular in nature and easy to scale up or down.

Performance-driven volume moves


NetApp AFX's zero-copy volume move functionality means that it can rebalance volumes as needed without copying data, which allows it to perform quickly and without needing added capacity. This means that volume moves can become a greater part of the automated load balancing available to ONTAP clusters. Now that it costs relatively nothing to move a volume, ONTAP can leverage this valuable tool to incorporate features like performance-driven load balancing of volumes.

In NetApp AFX running ONTAP 9.18.1 and later, node, HA pair, and volume utilization is being constantly monitored, while performance data is collected and analyzed. If a node's utilization falls outside of the defined thresholds, then ONTAP will automatically select a volume to move to a less utilized node in an effort to keep balanced performance across the cluster.

Performance-driven volume moves in NetApp AFX – high utilization triggers a volume move



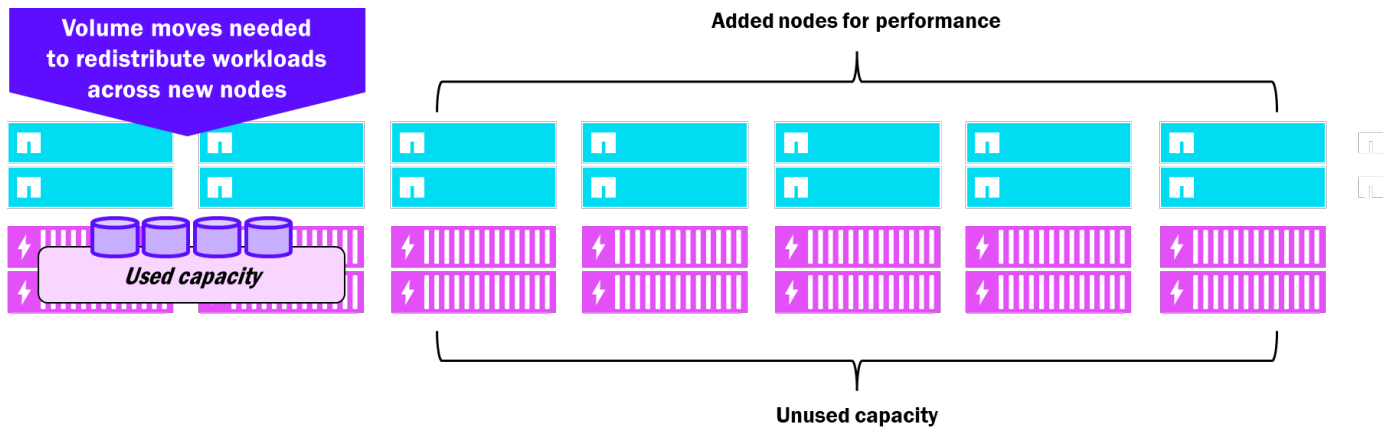
Performance-driven volume moves in NetApp AFX – Balanced node utilization after volume move



Cluster scale and expansion

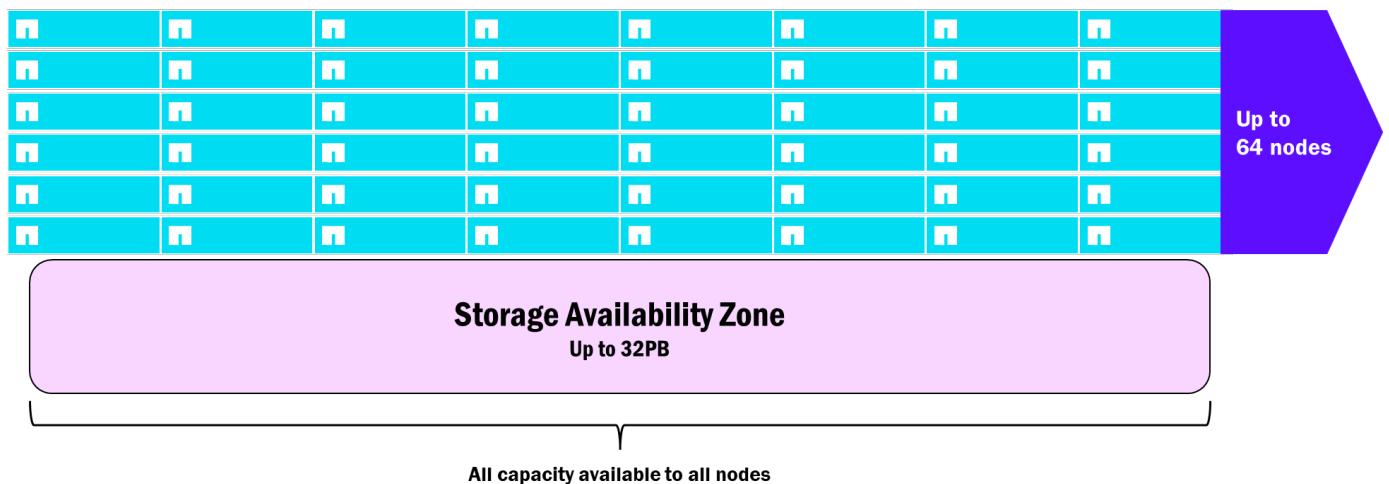
Unified ONTAP clusters support up to 24 nodes and each node that is added must also be added with disks (both for system functionality and data services). Disk shelves can be added to the cluster, but they are always connected to a single HA pair and owned only by a single node, even if the cluster is 24 nodes in size. This means that capacity is added to a cluster even when only performance is required, and that performance increase is mostly relegated to a specific set of disks owned by the new nodes. As a result, you may end up with extra capacity that you don't necessarily need.

Unified ONTAP – added scale considerations



NetApp AFX supports larger scale for clusters. Starting in 9.19.1, AFX clusters can reach 32 nodes in a single cluster. And because all nodes can see and access all disks, they can all share the performance and capacity (up to 32PB as of ONTAP 9.19.1) of those drives so that there is never any stranded resource. Volume moves don't require copies, so ONTAP is able to automatically move volumes to newly added nodes to ensure evenly distributed node utilization, while capacity is evenly distributed via the Storage Availability Zone.

NetApp AFX – added scale considerations



Root volume changes

In NetApp ONTAP, each node is assigned a root volume, which is used for system-specific files and functions, such as log files, boot images, core files, cluster databases, and more.

In unified ONTAP, those root volumes lived on physical root aggregates. To reduce the amount of capacity the root aggregates used, they were created across data drive partitions via Advanced Disk Partitioning (ADP).

NetApp AFX removes physical aggregates from the equation, and as a result, removes the need for root aggregates and ADP to be used. Root volumes are still a concept, but they now live in virtualized areas of the capacity pool and do not require extra configuration. Additionally, root volume functionality changes. Boot images and replicated cluster databases are moved off the storage stack and into an on-board boot media found on each AFX node. Now, if access to the storage stack is lost, nodes can still boot and maintain cluster eligibility, which eases the complexity of troubleshooting.

Onboard boot media

NetApp AFX nodes leverage on-board boot media, which is an NVMe attached M.2 device roughly 3.8TB in size. These boot devices contain boot image files and replicated databases that are separate from the storage enclosures, which provides extra redundancy in the case of disk access issues. If the boot media fails, the node will be taken over by its HA partner and the boot media can be replaced. Once replaced, a new ONTAP image would be loaded to the device by a storage administrator and ONTAP will automatically rebuild the cluster database to restore full functionality.

Performance

NetApp AFX was built with performance and scale in mind – specifically geared towards workloads that require high read and write throughput and can provide simple, linear scale.

Per-node performance

Each NetApp AFX storage node provides a specific amount of throughput for reads and writes. As nodes are added to the cluster, they linearly increase that performance, as covered in the "Linear scaling of node performance" section in this document.

Currently, the node types are [AFX 1K](#) and provide throughput for reads and writes at roughly the amounts below. As newer hardware is available to NetApp AFX, these limits may change.

NOTE: Maximum performance was reached using multiple clients reading and writing multiple files as shown in the "Benchmark results" section below.

Per-node performance estimates

Node type	Read performance max	Write performance max
AFX 1K	~35GB/s	~10GB/s



For the most up-to-date performance estimates, check with your NetApp sales team.

Per-shelf performance

Each shelf contains high performance shelf modules with 16 x 100GB Ethernet ports that leverage RoCEv2 communication for high bandwidth storage interaction with the compute nodes in the cluster. Like any physical resource, these shelves have maximums that can be achieved – particularly since NetApp AFX can present multiple nodes pointing to the same set of disks. The following table shows the estimated maximum read and write performance for a single shelf for TLC and QLC drives. For more information on TLC and QLC differences, see [TLC vs. QLC](#).

Per-shelf performance estimates

Shelf module type	Read performance max	Write performance max
NSM 140	140GB/s (TLC and QLC)	70GB/s TLC 35GB/s QLC



For the most up-to-date performance estimates, check with your NetApp sales team.

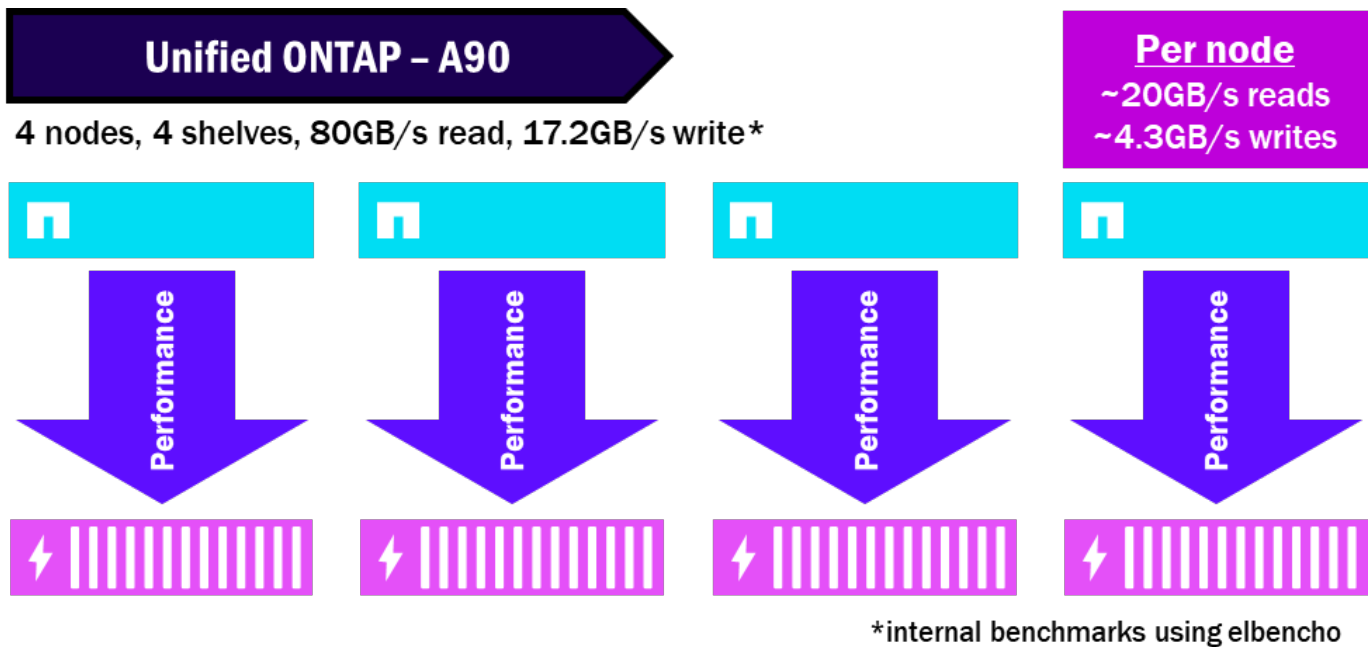
Performance density

Decoupling storage nodes from shelves in the disaggregated ONTAP architecture allows for more nodes to push traffic to fewer shelves, which helps reduce the overall datacenter footprint needed to get maximum performance with only the capacity you need.

This concept of "performance density" enables storage administrators to get the most out of the hardware they have while never having to overprovision their storage environment.

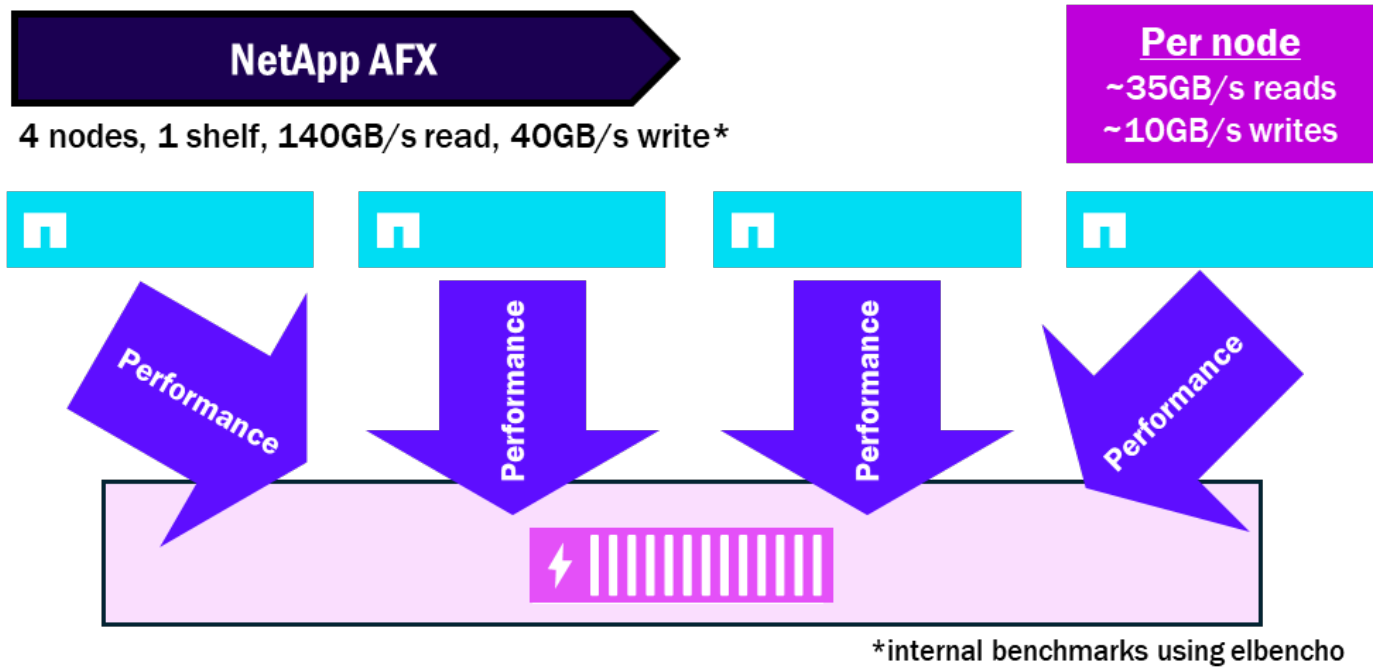
For instance, in a unified ONTAP cluster, since each node has its own set of disks, the performance is directed only to the disks owned by the node, and since only one node can access one set of disks, it can't necessarily saturate the available disks and achieve its maximum performance.

Unified ONTAP – How performance is divided up



NetApp AFX pools all the disks into a single Storage Availability Zone, so all nodes can leverage all the disks. And because the disks and nodes are decoupled, you won't need as many shelves to get the same performance. This condenses down the performance and maximizes the shelf's maximum performance potential.

NetApp AFX – Performance density



Node to shelf ratios

Unified ONTAP nodes require at least one set of disks per node and can have multiple shelves attached to a single node. As a result, there can be performance bottlenecks at the single node that may not be able to saturate its own disks.

NetApp AFX presents all disk shelves to all nodes. Each shelf contains modules with 16 x 100GB RoCE-capable interfaces to increase the total amount of performance allowed per shelf. Because of this, you can saturate a single shelf with multiple nodes that will be reading and writing to the same set of disks.

As of ONTAP 9.19.1, the node:shelf saturation ratio is around 4:1.

Benchmark results

The following section covers benchmark results using a NetApp AFX cluster with the following configuration parameters.

- 4 nodes, 4 data interfaces
- 2 shelves (7.6TB drives)
- ONTAP 9.19.1
- NFSv4.2 (pNFS, session trunking)
- FlexGroup volume
- [EIBencho](#) benchmark
- Writes: `elbencho --hosts=x.x.x.[y-z] -d -w -b 1M -t 80 --iodepth 1 --direct -s 600g /fio_vol1/`
- Reads: `elbencho --hosts=x.x.x.[y-z] -r -b 256k -t 80 --lat --iodepth 2 --direct -s 600g --infloop /fio_vol1/`
- 4 Cisco C240 M8 servers, 2 port * 200GbE CX-7 cards, 80 threads
- NFS mount options: `rw,vers=4.2,rsiz=1048576,wsiz=1048576,trunkdiscovery,proto=tcp`

The configuration above achieved very near the maximum reads available to the 4-node cluster (~134GB/s) and was right at the maximum writes allowed per node (40GB/s).

NetApp AFX – EIBencho read performance, 4 nodes



NetApp AFX – EIBencho write performance, 4 nodes



Aggressive read-ahead

In media streaming workloads, a 4K movie is often split into tens of thousands of files, each typically 50 MB to 250 MB in size. Each file represents a frame, and the application reads an entire frame in a single request. To maintain a smooth, uninterrupted stream with no visible buffering, these frame reads must complete without drops.

ONTAP provides a volume-level option (`-aggressive-readahead-mode`) to optimize these workloads. Starting in ONTAP 9.19.1, a new `cross_file_sequential_read` mode for aggressive readahead has been introduced on AFX to accelerate workloads with predictable I/O patterns across similar file types (for example, media rendering and streaming).

`cross_file_sequential_read` predicts the next file to be read based on its name and begins readahead on those files before the client issues the read call. The prediction logic assumes that all files in a directory follow a naming pattern with a monotonically increasing numeric suffix (eg. `file1`, `file2`, `file3`). All files in the directory must follow this pattern, using either decimal or hexadecimal numbering. Filenames can be up to 255 characters long. The logic is extensionagnostic and generates the next set of filenames in the current directory based solely on the current filename. If a filename previously generated using base10 numbering does not exist in the directory, the names are regenerated using hexadecimal numbering. If none of the generated filenames exist, no prefetch is issued for that set. Prefetching resumes when the next client read is issued.

With these options enabled, [frametest](#) performance benchmarks were able to read 30,000 4K frames at 30 frames per second with 30 clients (NFSv3 and SMB3) and 34 clients (NFSv4.1), without a single dropped frame.

While cross-file sequential read is primarily designed for media workloads, other readheavy workloads with predictable access patterns and file names - such as AI training and inference - can also benefit.

Considerations and caveats

- Shared buffer cache – Aggressive read ahead uses the same buffer cache as other volumes on the node. Enabling it may affect read performance for other volumes on that node.
- Underlying storage performance – If files cannot be read quickly enough (for example, on HDD-based FAS systems), cached data may be evicted before the client read occurs, negating the benefits of read ahead.
- Access pattern requirements – If the workload's read pattern is not sequential, or if files in a directory are not named in an increasing sequential order, the `cross_file_sequential_read` aggressive read ahead mode will not provide meaningful benefits.

NFSv4.x performance enhancements

NFS version 3 has been a gold standard for NFS applications for decades – starting all the way back in 1995, when it was first officially released. Its blend of performance and resiliency has made it difficult to consider a move to newer NFS versions for good reason.

However, NFSv3 is not without its limitations. Statelessness of the protocol, while great for performance and for minimizing disruptions on storage failover, is not so great for data consistency and lock management. An NFS server doesn't really keep track of lock states, so if there is a failure, the NFS server may or may not release the locks, and the NFS client may not know whether a file is locked or not.

Security for NFSv3 is also a bit lacking. The protocol requires multiple open firewall ports to function properly and numeric IDs are sent in plaintext over the wire. Furthermore, NFS does not have robust ACL support, and does not include native file and folder auditing. As a result of these limitations, NFSv4 was created in 2003 via [link:https://datatracker.ietf.org/doc/html/rfc3530](https://datatracker.ietf.org/doc/html/rfc3530) [RFC-3530^] (obsoleted in 2015 by [link:https://datatracker.ietf.org/doc/html/rfc7530](https://datatracker.ietf.org/doc/html/rfc7530) [RFC-7530^]).

While NFSv4.x has been around for over 20 years, there still hasn't been widespread adoption of it for a few reasons.

- Complexity of identity management: Many environments don't have a name service infrastructure in place to properly take advantage of the name string and Kerberos security requirements in NFSv4.x.
- Need for newer NFS clients: This concern is less pressing in today's modern NFS environments, as the further we get from the initial release date of NFSv4. Almost all currently used OS's include NFS clients with full NFSv4 support, but there are still legacy systems that may not have the necessary NFSv4.x packages. In fact, some applications still require use of older NFS versions.
- "If it ain't broke, don't fix it" mentality: Enterprise IT organizations are notoriously conservative in adoption of newer technologies – even ones that have been around for over 20 years. And if the current NFS version is working fine, why change?
- Performance concerns: Performance of a stateful protocol like NFSv4.x has lagged behind stateless NFSv3 for much of the past 20 years. In the past, the performance impact often outweighed the NFSv4.x benefits.

NFSv4.x improvements in ONTAP 9.18.1 using AFX

Some architectural changes to ONTAP have provided a much needed performance boost to NFS in general, and have made some serious inroads to improving NFSv4.x performance in general.

Below is a high-level summary of some of those changes.

Sequential read enhancement: NFSv4.1 30% better than NFSv3

ONTAP 9.18.1 introduces support for multipath IO with NFSv4.1. Rather than processing reads from the WAFL file system, MPIO shifts read operations into a network domain to be served in a multipath-safe manner. This approach reduces context switches, providing greater overall parallelism in sequential read traffic, as well as reducing the overhead from buffer management by bypassing WAFL.

Random read enhancement for FlexGroup volumes: NFSv4.1 within 7% of NFSv3

FlexGroup volumes are volumes that take many underlying constituent volumes and present them as a single unified namespace. In AFX, FlexGroup volumes have Advanced Capacity Balancing enabled by default, which will write files larger than 10GB across multiple constituent volumes as multipart files. Because of the remote location of these file parts, random reads traditionally have had a modest perf disadvantage with NFSv4.x (around 18% less than NFSv3). ONTAP 9.18.1 introduces support for cached IO for multipart reads with NFSv4.x to help address this.

NOTE: This change does not apply to FlexVol volumes.

Sequential writes: +10% improvement from previous releases

An improvement of how we replicate NVLOG data used for HA failover functionality increased overall sequential write performance for NetApp AFX systems.

Metadata operations: Within 15% performance of NFSv3 for EDA benchmarks


NFSv4.1 traditionally serializes all OPEN and CLOSE operations, with a cluster node processing them one at a time before they can be sent from the network to WAFL. ONTAP 9.18.1 introduces Concurrent Open Close (COC), which eliminates network serialization by changing how race conditions are resolved, which removes OPEN/CLOSE bottlenecks seen in prior releases.

All of these changes – along with the architecture changes brought about in AFX – have made it possible to improve the overall NFSv4.1 performance in ONTAP 9.18.1.

Sequential IO results

One of the areas where some modest performance improvements were seen was with sequential IO (ie, IO that is predictable and issued consecutively). In standard performance tests using fio, AFX running ONTAP 9.18.1 improved sequential read performance by nearly 30% and sequential write performance by 10%.

NetApp AFX – NFSv4.1 sequential IO performance in ONTAP 9.18.1



	AFX 9.17.1	AFX 9.18.1
Seq. reads	220GB/s	283GB/s
Seq. writes	70.6GB/s	77.7GB/s

Metadata-heavy workload results


Even more impressive are the improvements with one of NFSv4.x's top performance pain points – metadata. These are random IOs, usually in the 4K range, used to manage file owners and attributes, creating and listing files, and so on. Because of the statefulness of NFSv4.x, these types of operations tend to cost more in CPU and latency, which in turn reduces overall possible performance.

With the changes in AFX ONTAP 9.18.1, NFSv4.x performance for these types of workloads has improved substantially and has closed the gap on NFSv3 performance (within 15%).

Our performance engineering teams compared performance of standard AI image, EDA, and software build

benchmarks and found massive gains from the previous ONTAP release.

NetApp AFX – NFSv4.1 metadata IO performance in ONTAP 9.18.1



	9.17.1	9.18.1	Delta
AI Image processing	209 KIOPS	239 KIOPS	+16%
Software dev	600 KIOPS	950 KIOPS	+58%
EDA	480 KIOPS	881 KIOPS	+84%

Management tools

While there are some fairly large architectural differences in how storage is presented in NetApp AFX, the feature set and the way ONTAP is managed is virtually unchanged. This is entirely by design. ONTAP is ONTAP and there should be little to no learning curve when possible. And in this case, NetApp AFX is still running ONTAP.

Management – CLI

In NetApp AFX, the CLI is nearly identical to what unified ONTAP's CLI offers. Commands are still run at the cluster level for the most part, while there are still node-level commands that can be issued. There are still top-level command directories and sub-commands, as well as tab completion of commands. In addition, all of the same CLI shortcuts work (for example, filtering content using `-fields`).

The only real differences in the NetApp AFX CLI have to do with what was added and what was removed (as covered in the "" section of this document). If a feature (such as aggregates or Metrocluster) have been removed, then the corresponding CLI command will no longer be available.

In addition, where new features and functionality have been added to NetApp AFX, new commands will also be available. For instance, the new [NetApp AI Data Engine \(AIDE\)](#) add-on interacts directly with a NetApp AFX cluster over the backend cluster network. As a result, the `dcn` and `data-engine` command directories are added to the NetApp AFX CLI.

The following shows the top-level command directories available at admin privilege in a NetApp AFX cluster. Newly added commands in bold.

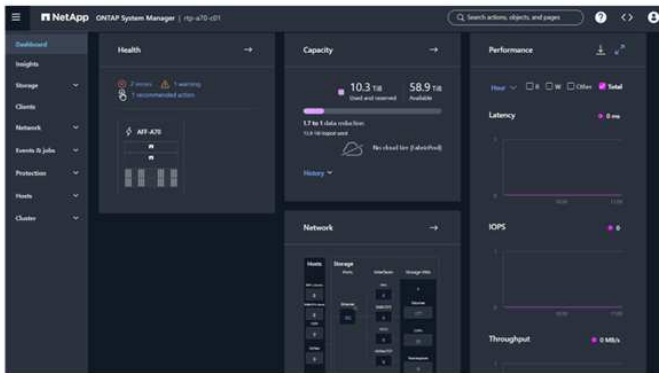
```
AFX::>
```

```
cluster      data-engine  dcn          event        exit
history      job          man          metrocluster network
qos          redo        rows        run          security
set          snaplock    snapmirror   statistics   statistics-v1
storage      system      top          up           volume
vserver
```

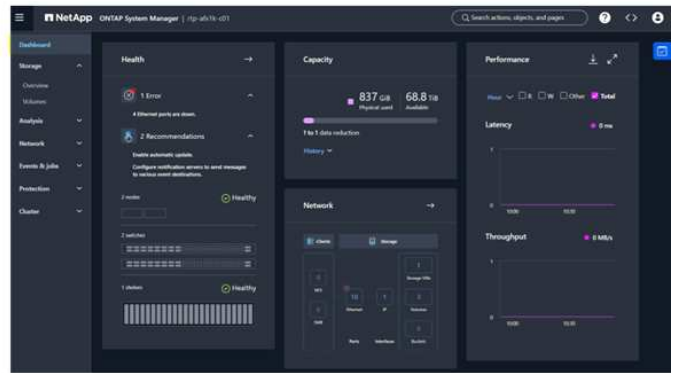
Management – GUI

System Manager is still the GUI interface to interact with NetApp AFX clusters for management of resources, and when you first log in to it, you likely won't be able to tell at first glance that you aren't still logged into a unified ONTAP system.

Side-by-side look at unified ONTAP and NetApp AFX System Manager landing pages




System Manager Dashboard view
Unified ONTAP



System Manager Dashboard view
NetApp AFX


As you can see from the above, there aren't many obvious differences in System Manager for unified ONTAP and NetApp AFX. There are a few tells, however.

How NetApp AFX differs in System Manager



As you drill down into menus and dashboards, most everything remains the same. Volumes still show used and total capacity. Network interfaces still show ports and IP addresses. You can still configure protection policies for SnapMirror. Hardware views are still available. But NetApp AFX also enhances the views a bit by offering a new cabling view, where you can drill down and see where all the cables are connected throughout the storage stack.

NetApp AFX cabling view



REST APIs

NetApp AFX also makes a concerted effort to retain a majority of the ONTAP REST API calls, which means if you've created a suite of automation tools based on REST API, they should still be able to be used in most cases. The main exceptions include aggregates, Metrocluster, SAN, and some performance counters. Complete REST API documentation can be found within the NetApp AFX cluster's System Manager by navigating to https://clus_mgmt_ip/docs/api.

Off-box ONTAP management tools

NetApp AFX provides some support for off-box ONTAP management tools, such as:

- NetApp System Manager
- NetApp Console
- Grafana Harvest (25.11.0 and later)
- NetApp Trident (25.10 and later)

NetApp AFX does not currently support:

- NetApp ActiveIQ Unified Manager

Networking, security, and operations

NetApp AFX supports the same networking stack, security features, data protection technologies, nondisruptive operations, and volume types as unified ONTAP, with some adjustments unique to the disaggregated architecture.

Networking

The networking stack in NetApp AFX is identical to unified ONTAP.

- Data LIFs are still used to present network addresses to internal and external services.
- Each node has its own set of physical and virtual ports.
- VLANs, ifgroups, and BGP are all still supported.
- LIFs can still failover between physical nodes and ports in the cluster.
- IPspaces/broadcast domains are still configured the same.
- Each SVM can have its own dedicated data network.
- Management networks can be segmented from data networks.
- Client data networks don't change outside of the new addition of 400GB networking support.
- Backend cluster switches are still configured via a "golden configuration" file provided by NetApp.

Some key network differences include:

- Backend cluster network ports now support only 100GB connections to the cluster switches.
- Because the cluster switches are 400GB capable, the backend node connections use 4 x 100GB breakout cables to reduce the number of ports used on the switches.
- NVRAM is now mirrored between HA pairs across the backend cluster network via a new HA VLAN configuration on the switches.

- A new DCN network is added by default for the AI Data Engine. These IP addresses are automatically generated and can be changed as needed.

Security

NetApp AFX runs ONTAP, which means it uses the same security as ONTAP. All of the cryptomods are identical, which means that security certifications will be identical once the certification processes are completed. NetApp AFX also leverages the same support for security ciphers as unified ONTAP.

In addition, NetApp AFX supports many of the security features provided by unified ONTAP, including (but not limited to):

- Autonomous Ransomware Protection
- Secure Multi-tenancy
- Encryption at-rest (volume encryption) and in-flight (TLS 1.3)
- Self-encrypting drives (SEDs)
- NFS and SMB Kerberos authentication and encryption
- Multi-admin verification
- SnapLock

For information on which certifications unified ONTAP has received (as well as other security hardening information), see:

- [NetApp Product Security](#)
- [ONTAP security hardening overview](#)

Snapshots and Data Protection

NetApp AFX leverages the same snapshot and replication technologies as unified ONTAP, with no major changes to the way these features work. In fact, AFX can replicate to and from unified ONTAP systems with the same [rules and configurations](#) you are familiar with.

The only exception in AFX for replication has to do with FlexGroup volumes replicating to a unified ONTAP system. In that case, the destination unified ONTAP system must be running ONTAP 9.16.1 or later to provide Advanced Capacity Balancing support.

Nondisruptive operations

ONTAP provides nondisruptive operations, such as volume moves, upgrades, cluster maintenance, storage failovers, and more. NetApp AFX provides the same nondisruptive operations, with some adjustments.

- Volume moves are still nondisruptive, but no longer require a copy.
- Storage failovers are still nondisruptive, but after initial failover, volumes rebalance across all surviving nodes in the cluster.
- LIF migrations are identical.
- Hardware maintenance and upgrades are identical.

Volume types

Unified ONTAP supports several different volume types, such as:

- FlexVols
- FlexGroup volumes
- FlexCache
- FlexClone
- Object buckets

NetApp AFX provides full support for each of these volume types, as well as full interoperability for FlexCache volumes with unified ONTAP systems.

For more information on how FlexGroup volumes benefit from AFX architecture, see [FlexGroup volume management improvements](#).

Hardware details

The following section covers details about the NetApp AFX cluster hardware. For the most recent information about NetApp AFX hardware, see <https://hwu.netapp.com>.

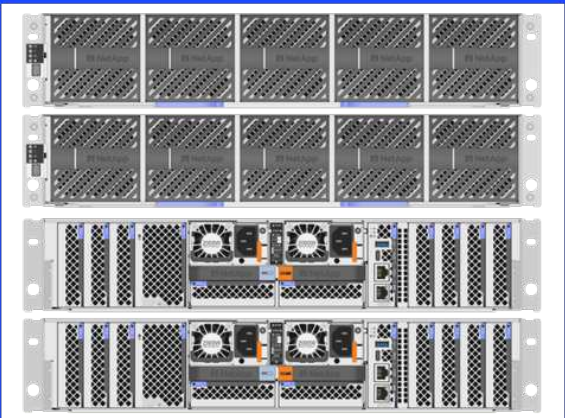
Supported hardware

For the latest information on supported hardware for NetApp AFX, check <https://hwu.netapp.com>.

Nodes

NetApp AFX nodes are based on the AFF A1K model of nodes provided for unified ONTAP clusters. These nodes have no on-board disk for storage and are intended to be modular to easily add and remove nodes based on performance requirements. Each node uses 2U of rack space and increases performance linearly as they are added to an AFX cluster.

NetApp AFX 1K node details



The image shows a NetApp AFX 1K node, a 2U server rack unit. It features a front panel with multiple drive bays and a central control area with two fans. The node is shown from a front-three-quarter perspective.

AFX 1K

Processor/Memory (per HA pair)

- CPU: 208 cores, 52 cores per CPU, 4x1.7 Ghz
- RAM: 2048GB
- NVRAM: 128GB
- On-board NVMe 3.8TB boot device

I/O cards

- HA/cluster/storage (100GB only; 4x100GB breakout)
- Up to 400GB client data network
- 18x IO expansion per HA pair

Hardware slots

NetApp AFX 1K nodes use the following slot assignments.

- Slot 1 is dedicated to HA replication

- Slot 7 is dedicated to cluster replication
- Slots 10 and 11 are dedicated to storage shelf communication

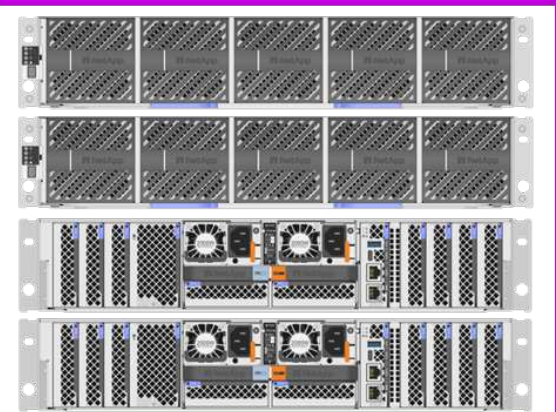
Shelves

NetApp AFX shelves use the same enclosures that the AFF systems leverage. The main difference in shelves in AFX is in the module used. The NSM140 modules provide enhanced performance capabilities and help make disaggregated ONTAP possible. A few key considerations:

- Only fully populated shelves supported.
- Shelves are automatically detected by NetApp AFX when plugged in.
- No shelf removal support currently.

NetApp AFX shelf enclosure details

AFX shelf enclosure



Enclosure

- 2U chassis
- 24 NVMe drives

NSM modules

- 2x NSM140 modules
- 16x 100GB ports (4x100GB breakout)
- RoCE connected
- Only switch connected
- Faster processor, more memory


Switches

NetApp AFX still makes use of backend cluster switches for intracluster communication, such as cluster database replication, remote data operations, storage operations, and NVRAM mirroring. Functionally, the only differences between unified ONTAP and NetApp AFX cluster switches are:

- 400GbE support
- New HA VLAN
- For more information, see the [Cisco switches data sheet](#).

NetApp AFX cluster switches

AFX storage switch




Switch details

- Cisco 9332 (32 ports, 1U) or 9364 (64 ports, 2U)
- 400-Gigabit QSFP-DD ports (32 or 64)
- 4 x 100GB breakout support
- MACsec encryption support
- 120MB memory buffer
- 32GB system memory
- 6 CPU cores

Cluster switch connectivity

NetApp AFX heavily leverages backend cluster switches for many of its major architectural concepts. For instance, cluster interfaces, storage adapters, storage shelves, and NVRAM cards all plug into the cluster switches. Currently, all of these interfaces support 100GB communication only, while the switches support 400GB. As a result, the 100GB interfaces plug into the 400GB interfaces on the switch by way of 4 x 100GB breakout cables. This approach reduces the number of ports used on the switches. For instance, 16 x 100GB storage shelf module ports will use just 4 ports on the switches, while the 8 total ports on the nodes use 2 switch ports.

NetApp AFX switch cabling



Disk types and sizes

NetApp AFX currently only supports NVMe attached SSDs in the following sizes:

- 7.6TB
- 15.3TB
- 30.1TB
- 60.6TB

TLC vs. QLC

NetApp AFX can leverage both TLC and QLC flash types. The 7.6 and 15.3TB drives are TLC models, while drives larger than 30.1TB will be QLC. Regardless of media type, the NVIDIA SuperPod certification performance standards can be met.

All drives used with NetApp AFX are performance-rated drives, and both QLC and TLC will perform nearly identically for read traffic. QLC lags a bit behind write performance as compared to TLC and can experience slightly more wear-leveling with write-heavy workloads.

For performance numbers, see [Per-shelf performance](#).

When choosing the drive type in use, consider what workloads the storage will be hosting and the

performance/capacity density tradeoffs involved.

Maximums and limits

The following section covers the maximums and minimums of a NetApp AFX cluster as compared to unified ONTAP.

For the most recent limits, see [Hardware Universe](#).

NetApp AFX maximums and minimums

Limit	Unified ONTAP	NetApp AFX
Volume count (cluster)	<ul style="list-style-type: none"> 30,000 (platform dependent) 	<ul style="list-style-type: none"> 30,000
Aggregate size	<ul style="list-style-type: none"> 800TB 	<ul style="list-style-type: none"> N/A
Node count	<ul style="list-style-type: none"> 24 	<ul style="list-style-type: none"> 8 (9.17.1) 32 (9.19.1RC1)
Total capacity	<ul style="list-style-type: none"> Node and drive dependent. See hwu.netapp.com for details. 	<ul style="list-style-type: none"> 2PB (9.17.1) 3PB (9.18.1RC1) 16PB (9.18.1GA release) 20PB (9.19.1RC0) 32PB (9.19.1RC1)
Supported shelf count	<ul style="list-style-type: none"> 8 per HA pair 192 per cluster 	<ul style="list-style-type: none"> 12 per cluster (9.18.1GA) 17 per cluster (9.19.1RC0) 25 per cluster (9.19.1RC1) (no HA pair limitation)
Total supported drive count	<ul style="list-style-type: none"> 240 per HA pair 2880 (24 nodes) 	<ul style="list-style-type: none"> 288 per cluster (based on shelf count limits; no HA pair limitation) 408 per cluster (9.19.1RC0) 600 per cluster (9.19.1RC1)
Volume size	<ul style="list-style-type: none"> 300TB (FlexVol) 60PB (FlexGroup) 	<ul style="list-style-type: none"> 300TB (FlexVol) 60PB (FlexGroup)*
Maximum files per volume	<ul style="list-style-type: none"> 2 billion (FlexVol) 400 billion (FlexGroup) 	<ul style="list-style-type: none"> 2 billion (FlexVol) 400 billion (FlexGroup)

Limit	Unified ONTAP	NetApp AFX
Maximum files per directory (default 320MB maxdirsize value)	<ul style="list-style-type: none"> • ~4 million (FlexVol) • ~2 million (FlexGroup) 	<ul style="list-style-type: none"> • ~4 million (FlexVol) • ~2 million (FlexGroup)
SnapMirror concurrent transfers	<ul style="list-style-type: none"> • 250 	<ul style="list-style-type: none"> • 250
Number of qtrees	<ul style="list-style-type: none"> • 4096 	<ul style="list-style-type: none"> • 4096
Maximum TCP connections per node	<ul style="list-style-type: none"> • 100,000 	<ul style="list-style-type: none"> • 100,000
Maximum locks per node	<ul style="list-style-type: none"> • 3 million 	<ul style="list-style-type: none"> • 3 million
Maximum data interface count (cluster)	<ul style="list-style-type: none"> • 4096 	<ul style="list-style-type: none"> • 4096
Maximum constituent count – FlexGroup	<ul style="list-style-type: none"> • 200 	<ul style="list-style-type: none"> • 200 • 512 in 9.19.1 RC1
Maximum number of intercluster LIFs	<ul style="list-style-type: none"> • 8 	<ul style="list-style-type: none"> • 8
Maximum number of IP spaces	<ul style="list-style-type: none"> • 512 	<ul style="list-style-type: none"> • 512
Maximum number of FlexCache per origin	<ul style="list-style-type: none"> • 100 	<ul style="list-style-type: none"> • 100
Max number of FlexCache (node)	<ul style="list-style-type: none"> • 400 	<ul style="list-style-type: none"> • 400

Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

- Hardware universe
<https://hwu.netapp.com/>
- NetApp AFX product page
<https://www.netapp.com/afx/>
- NetApp AFX: Scalable, secure AI data infrastructure
<https://www.netapp.com/blog/afx-scalable-secure-ai-data-infrastructure/>
- Deciding between NFSv3 or NFSv4.x? The choice is getting clearer...
<https://community.netapp.com/t5/Tech-ONTAP-Blogs/Deciding-between-NFSv3-or-NFSv4-x-The-choice-is-getting-clearer>
- NetApp AFX Data Sheet
<https://www.netapp.com/media/142853-ds-3466-netapp-afx-datasheet.pdf>

ONTAP SnapCenter technical reports

SnapCenter provides a, unified platform for application-consistent data protection and clone management. SnapCenter simplifies backup, restore, and clone lifecycle management with application-integrated workflows. Leveraging storage-based data management, SnapCenter enables increased performance and availability and reduced testing and development times.



These technical reports expand on the [SnapCenter](#) product documentation.

SnapCenter for Oracle

[TR-4700: SnapCenter Plug-In for Oracle database best practices](#)

NetApp SnapCenter is a unified, scalable platform for Oracle-consistent data protection that automates complex operations with centralized control and oversight. Learn about the recommended practices for deploying Oracle databases with SnapCenter.

[TR-4964: Oracle Database backup, restore and clone with SnapCenter Services](#)

Learn how to set up SnapCenter Services to backup, restore, and clone Oracle databases deployed to Amazon FSx for ONTAP storage and EC2 compute instances. Although it is much easier to set up and use, SnapCenter Services deliver key functionalities available through the SnapCenter interface.

SnapCenter for Microsoft SQL Server

[TR-4714: Best practices for Microsoft SQL Server using NetApp SnapCenter](#)

Learn how to successfully deploy Microsoft SQL Server on NetApp storage using SnapCenter for data protection.

SnapCenter for Microsoft Exchange Server

[TR-4681: Best practices for Microsoft Exchange Server using NetApp SnapCenter](#)

Learn how to successfully deploy Microsoft Exchange Server on NetApp storage using SnapCenter for data protection.

SnapCenter for SAP HANA

[TR-4614: SAP HANA backup and recovery with SnapCenter](#)

SnapCenter is a unified, scalable platform for application-consistent data protection for SAP HANA and other databases. SnapCenter provides centralized control and oversight, while delegating the ability for users to manage application-specific backup, restore, and clone jobs. With SnapCenter, database and storage administrators learn a single tool to manage backup, restore, and cloning operations for a variety of applications and databases.

[TR-4926: SAP HANA on Amazon FSx for NetApp ONTAP - Backup and recovery with SnapCenter](#)

Learn about the recommended practices for SAP HANA data protection on Amazon FSx for NetApp ONTAP and SnapCenter. Topics include SnapCenter concepts, configuration recommendations, and operation workflows, including configuration, backup operations, and restore and recovery operations.

[TR-4667: Automating SAP HANA System copy and clone operations with SnapCenter](#)

SnapCenter storage cloning and the option to flexibly define pre-cloning and post-cloning operations allows SAP Basis administrators to accelerate and automate SAP system copy, clone, or refresh operations. Learn now the option to choose any SnapCenter Snapshot backup at any primary or secondary storage allows you to address your most important use cases, including logical corruption, disaster recovery testing, or the refresh of an SAP QA system.

[TR-4719: SAP HANA system replication backup and recovery with SnapCenter](#)

Learn how SnapCenter technology and the SAP HANA plug-in can be used for backup and recovery in an SAP HANA System Replication environment.

[TR-4667: Automating SAP HANA system copy and clone operations with SnapCenter](#)

The ability to create application-consistent NetApp Snapshot backups on the storage layer is the foundation for the system copy and system clone operations. Storage-based Snapshot backups are created by using the NetApp SnapCenter Plug-In for SAP HANA and interfaces provided by the SAP HANA database. SnapCenter registers Snapshot backups in the SAP HANA backup catalog so that the backups can be used for restore and recovery as well as for cloning operations.

SnapCenter hardening guide

[TR-4957: Security hardening guide for NetApp SnapCenter](#)

Learn how to configure SnapCenter to help organizations meet prescribed security objectives for information system confidentiality, integrity, and availability.

ONTAP tiering technical reports

With the FabricPool data tiering solution, an enterprise's overall user experience of flash systems improves while avoiding the pain of rearchitecting applications for storage efficiency. FabricPool reduces the storage footprint and the associated costs of a system's environment. Active data remains on high-performance SSDs. Inactive data is tiered to low-cost object storage while preserving storage efficiencies.



These technical reports expand on the [ONTAP FabricPool](#) product documentation.

[TR-4598: FabricPool best practices](#)

Learn about the capabilities, requirements, implementation, and recommended practices for FabricPool.

[TR-4826: NetApp FabricPool with StorageGRID recommendation guide](#)

Learn about the recommended practices for deploying and sizing StorageGRID as a capacity tier for the ONTAP component FabricPool. This document also covers core capabilities, requirements, implementation, and recommended practices when using StorageGRID.

[TR-4695: Database storage tiering with NetApp FabricPool](#)

Learn about the benefits and configuration options of FabricPool with various databases, including the Oracle relational database management system (RDBMS).

ONTAP virtualization technical reports

NetApp virtualization solutions help deliver maximum value from your servers. With a responsive virtual server infrastructure built on groundbreaking, high-performance ONTAP flash systems, you gain the ability to access your data much faster. Your granular virtual infrastructure scales without disruption to multiple petabytes of data, delivering the performance you need for shared access to multiple workloads. ONTAP helps streamline and reduce the complexity of your virtual server infrastructure deployment with key partnerships, deployment guidance, application integration, and superior design. ONTAP provides many recommended practices and solutions for a robust virtualization environment both on-premises and in the cloud.

These technical reports expand on the [ONTAP tools for VMware vSphere](#) product documentation.

[TR-4597: VMware vSphere for ONTAP](#)

ONTAP has been a leading storage solution for VMware vSphere environments for almost two decades and continues to add innovative capabilities to simplify management while reducing costs. This document introduces the ONTAP solution for vSphere, including the latest product information and recommended practices, to streamline deployment, reduce risk, and simplify management.

[TR-4400: VMware vSphere Virtual Volumes \(vVols\) with NetApp ONTAP](#)

ONTAP has been a leading storage solution for VMware vSphere environments for over two decades and continues to add innovative capabilities to simplify management while reducing costs. This document covers ONTAP capabilities for VMware vSphere Virtual Volumes (vVols), including the latest product information and use cases along with recommended practices and other information to streamline deployment and reduce errors.

[TR-4900: VMware Site Recovery Manager with NetApp ONTAP](#)

ONTAP has been a leading storage solution for VMware vSphere environments since its introduction into the modern data center in 2002, and it continues to add innovative capabilities to simplify management while reducing costs. This document introduces the ONTAP solution for VMware Site Recovery Manager (SRM), VMware's industry leading disaster recovery (DR) software, including the latest product information and recommended practices to streamline deployment, reduce risk, and simplify ongoing management.

[Introduction to automation for ONTAP and vSphere](#)

Automation has been an integral part of managing VMware environments since the first days of VMware ESX. The ability to deploy infrastructure as code and extend practices to private cloud operations helps to alleviate concerns surrounding scale, flexibility, self-provisioning, and efficiency. This document introduces the ONTAP solution for automating ONTAP and VMware vSphere environment.

[WP-7353: ONTAP tools for VMware vSphere - Product security](#)

This document describes the techniques and technology used to secure ONTAP tools for VMware vSphere 9.X from both existing and emerging threats in product environments.

[WP-7355: SnapCenter plug-in VMware vSphere - Product security](#)

This document describes the techniques and technology used to secure the NetApp SnapCenter Plug-in for VMware vSphere 4.X from both existing and emerging threats in product environments.

[TR-4568: NetApp deployment guidelines and storage best practices for Windows Server](#)

Microsoft Windows Server is an enterprise-class operating system that covers networking, security, virtualization, cloud, virtual desktop infrastructure, access protection, information protection, web services,

application platform infrastructure, and much more. This document focuses on Microsoft Windows, with a particularly heavy emphasis on Hyper-V virtualization technology, including the latest product information and recommended practices, to streamline deployment, reduce risk, and simplify management.

Legal notices

Legal notices provide access to copyright statements, trademarks, patents, and more.

Copyright

<https://www.netapp.com/company/legal/copyright/>

Trademarks

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<https://www.netapp.com/company/legal/trademarks/>

Patents

A current list of NetApp owned patents can be found at:

<https://www.netapp.com/pdf.html?item=/media/11887-patentspage.pdf>

Privacy policy

<https://www.netapp.com/company/legal/privacy-policy/>

Open source

Notice files provide information about third-party copyright and licenses used in NetApp software.

ONTAP

[Notice for ONTAP 9.16.1](#)
[Notice for ONTAP 9.16.0](#)
[Notice for ONTAP 9.15.1](#)
[Notice for ONTAP 9.15.0](#)
[Notice for ONTAP 9.14.1](#)
[Notice for ONTAP 9.14.0](#)
[Notice for ONTAP 9.13.1](#)
[Notice for ONTAP 9.12.1](#)
[Notice for ONTAP 9.12.0](#)
[Notice for ONTAP 9.11.1](#)
[Notice for ONTAP 9.10.1](#)
[Notice for ONTAP 9.10.0](#)
[Notice for ONTAP 9.9.1](#)
[Notice for ONTAP 9.8](#)
[Notice for ONTAP 9.7](#)
[Notice for ONTAP 9.6](#)
[Notice for ONTAP 9.5](#)
[Notice for ONTAP 9.4](#)
[Notice for ONTAP 9.3](#)

[Notice for ONTAP 9.2](#)
[Notice for ONTAP 9.1](#)

ONTAP Mediator for MetroCluster IP configurations

[9.9.1 Notice for ONTAP Mediator for MetroCluster IP configurations](#)
[9.8 Notice for ONTAP Mediator for MetroCluster IP configurations](#)
[9.7 Notice for ONTAP Mediator for MetroCluster IP configurations](#)

Copyright information

Copyright © 2026 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.