



NetApp Console local deployment

NetApp Console local deployment

NetApp
August 10, 2026

This PDF was generated from <https://docs.netapp.com/us-en/console-local/whats-new.html> on August 10, 2026. Always check docs.netapp.com for the latest.

Table of Contents

NetApp Console local deployment	1
Release notes	1
What's new with NetApp Console local deployment	1
Known limitations in NetApp Console local deployment	1
Compare NetApp Console local deployment and NetApp Console	1
Get started	4
Learn about NetApp Console local deployment	4
Learn about NetApp Console local deployment identity and access management	6
Learn about fleet management in NetApp Console local deployment	9
Learn about storage classes and policies in NetApp Console local deployment	11
Learn about LLM integration in NetApp Console local deployment	13
Learn about NetApp Backup and Recovery in NetApp Console local deployment	15
Install and set up	15
Quick start for setting up NetApp Console local deployment	15
Install NetApp Console	16
Plan your Console organization	19
Set up your Console organization	24
Configure Console integrations and services	42
Use NetApp Console	49
Log in to NetApp Console local deployment	49
Switch between fleets in NetApp Console local deployment	49
Manage your NetApp Console local deployment user settings	50
Use the NetApp Console assistant in NetApp Console local deployment	52
Manage storage in NetApp Console	54
Learn about fleet management in NetApp Console local deployment	54
Standardize storage behavior	56
Provision storage in NetApp Console local deployment	64
Monitor storage health	65
Storage monitoring in NetApp Console local deployment	65
Monitor with dashboards	66
Monitor fleets using inventory in NetApp Console local deployment	77
Remediate alerts	79
Investigate performance issues	90
Administer and monitor NetApp Console local deployment	100
Learn about administration and monitoring in NetApp Console local deployment	100
Audit NetApp Console local deployment activity	100
Maintain NetApp Console	102
Manage users and access	105
Reference	110
NetApp Console local deployment API	110
Supported LLM providers and models in NetApp Console local deployment	112
ONTAP alerts reference in NetApp Console local deployment	113
Cluster security compliance categories in NetApp Console local deployment	121

Storage VM security compliance categories in NetApp Console local deployment	122
Knowledge and support	123
Register for support in NetApp Console local deployment	123
Get help with NetApp Console local deployment	127

NetApp Console local deployment

Release notes

What's new with NetApp Console local deployment

Learn about the new features and enhancements in the latest release of the NetApp Console local deployment.

Introducing NetApp Console local deployment

[Learn NetApp Console local deployment.](#)

Known limitations in NetApp Console local deployment

Known limitations identify platforms, devices, or functions that are not supported by this release of the product, or that do not interoperate correctly with it. Review these limitations carefully.

These limitations are specific to the setup for the NetApp Console and administration: the agent, the software as a service (SaaS) platform, and more.

Console VM limitations

Possible conflict with IP addresses in the 10.42.0.0/16 and 10.43.0.0/16 ranges

NetApp Console local deployment uses fixed address space for inter-service communication. The IP range 10.42.0.0/16 and 10.43.0.0/16 are used for the internal network. Before deploying Console local deployment, ensure these IP ranges are not assigned to other VMs, DHCP scopes, DNS records, or load balancer VIP pools on the VLANs made available to the Console local deployment.

See Knowledge Base article [Agent IP conflict with existing network](#) for instructions on how to change the IP address of the agent's interfaces.

Shared Linux hosts are not supported

The agent isn't supported on a VM that is shared with other applications. The VM must be dedicated to the agent software.

Third-party agents and extensions

Third-party agents or VM extensions are not supported on the agent VM.

Compare NetApp Console local deployment and NetApp Console



Choose NetApp Console if you manage cloud and on-premises environments and want NetApp to handle platform infrastructure. Choose NetApp Console local deployment if you manage a large on-premises ONTAP estate and need full data sovereignty and policy-driven storage management.

NetApp Console

NetApp hosts and maintains NetApp Console as a SaaS application. You sign up and start managing storage immediately. To manage storage systems, you deploy lightweight Console agents in your cloud or on-premises environments. Agents connect outbound to the NetApp Console platform, and NetApp handles platform authentication, upgrades, and availability automatically.

NetApp Console also supports restricted mode (installed in your own public cloud with limited outbound connectivity) and private mode (installed on-premises or in your cloud with no connectivity to the NetApp Console platform, including air-gapped environments).

[Learn about NetApp Console deployment modes.](#)

NetApp Console supports the broadest range of storage systems, including cloud storage such as Amazon FSx for ONTAP, Cloud Volumes ONTAP, Amazon S3, Google Cloud Storage, E-Series, StorageGRID, and on-premises ONTAP clusters. It provides access to the full suite of NetApp data services, including NetApp Backup and Recovery, NetApp Disaster Recovery, NetApp Ransomware Resilience, NetApp Copy and Sync, and NetApp Data Classification. NetApp Console also provides support contract features such as Digital Advisor, Lifecycle Planning, and Sustainability dashboards.

NetApp Console is best suited for organizations managing hybrid cloud environments that want the most features, less work to keep things running, and flexible licensing.

[Learn about NetApp Console^.](#)

NetApp Console local deployment

NetApp Console local deployment lets you run the NetApp Console control plane entirely within your own on-premises infrastructure, with no management traffic leaving your environment. You deploy the Console as a virtual machine using a virtual appliance image and manage it like any other virtual machine in your vCenter environment. No Console agents are required. The virtual machine communicates directly with your ONTAP clusters.

While NetApp Console private mode also supports air-gapped and on-premises deployments, Console local deployment is designed specifically for enterprise-scale ONTAP fleet management. It adds capabilities that are not available in any NetApp Console deployment mode:

Storage class policies

Define storage standards once as reusable templates that bundle performance, capacity, and tiering requirements. All provisioning uses approved classes, and Console local deployment continuously monitors workloads for compliance and identifies drift.

Fleet management

Organize clusters into folders and fleets that reflect your business structure, then delegate administration at organization, folder, or fleet scope.

Workload Analyzer

Correlate latency, throughput, input/output operations per second, and configuration changes over time to find performance issues before they affect workloads.

Large language model integration

Connect to your own large language model to provision storage using natural language, investigate alerts, and get contextual answers about your storage environment.

Guided and automated remediation

When a workload drifts from its storage class or an alert fires, Console local deployment provides Fix-It recommendations. Use guided steps or one-click automated remediation to restore compliance without manual steps.

Alert delivery via email and webhook

Configure email and webhook notification channels so that alerts reach the right teams when storage conditions change. Webhooks integrate with your existing monitoring and incident management tools.

Console local deployment supports on-premises ONTAP clusters and NetApp Backup and Recovery. It integrates with Active Directory for federated authentication and supports multi-factor authentication for local users.

Console local deployment is best suited for organizations with large on-premises ONTAP environments that require policy-driven provisioning, automated compliance checks, and full data sovereignty.

[Learn about NetApp Console local deployment^.](#)

Feature comparison

Storage management

Feature	NetApp Console	NetApp Console local deployment
On-premises ONTAP clusters	Yes (requires Console agent)	Yes (direct communication from Console virtual machine)
Cloud storage systems (Amazon FSx for ONTAP, Cloud Volumes ONTAP, Amazon S3, and others)	Yes	No
E-Series and StorageGRID	Yes	No
NetApp Backup and Recovery	Yes	Yes
NetApp Disaster Recovery	Yes	No
NetApp Ransomware Resilience	Yes	No
NetApp Copy and Sync	Yes	No
NetApp Data Classification	Yes	No
Digital Advisor, Lifecycle Planning, Sustainability	Yes	No
Storage class policies and compliance monitoring	No	Yes
Guided and automated remediation	No	Yes
Fleet management (folders and fleets)	No	Yes
Workload Analyzer	No	Yes
Large language model integration for AI-assisted operations	No	Yes

Console setup and security

Feature	NetApp Console	NetApp Console local deployment
Deployment model	Software as a service, hosted by NetApp (different deployment modes available)	Self-hosted virtual machine in your own infrastructure
Console agent required	Yes	No
Software upgrades	Automatic, managed by NetApp	Manual
User interface access	NetApp Console application (internet or virtual machine access)	Console virtual machine (local, no internet required)
Data sovereignty	Management traffic goes to the NetApp Console platform	No data or management traffic leaves your environment
Active Directory integration	No	Yes
Identity federation	Yes	No
Multi-factor authentication (local users)	Yes	Yes
Audit logging	Yes	Yes
Licensing	Marketplace subscriptions and bring-your-own-license	Bring-your-own-license

Get started

Learn about NetApp Console local deployment

NetApp Console local deployment lets you host and run the NetApp Console autonomous control plane in your own infrastructure.

You get unified storage management, policy enforcement, fleet-scale automation, and AI-assisted operations. No data, metadata, or management traffic leaves your environment.

Deploy and operate Console in your own infrastructure

NetApp Console local deployment runs in your own infrastructure, so your team controls deployment, connectivity, and daily operations. You deploy the Console agent, maintain the Console virtual machine, and manage the network paths required for monitoring and management traffic. This gives enterprise administrators full control of operational boundaries while keeping management data inside the environment.

- [Install NetApp Console local deployment using an OVA in vCenter.](#)
- [Maintain your vCenter or ESXi host for NetApp Console local deployment.](#)
- [Learn about ports for the on-premises Console agent in NetApp Console local deployment.](#)

Automate storage operations with APIs and service accounts

NetApp Console local deployment supports API-based integrations so your organization can automate

repeatable storage operations. Service accounts let applications authenticate and operate with assigned roles. The same role-based access control and audit controls that apply to interactive users govern those actions. This supports enterprise automation while keeping access scoped and accountable.

- [Add members to your NetApp Console local deployment organization.](#)
- [Learn about the API for NetApp Console IAM](#)

Control access with RBAC and Active Directory integration

NetApp Console local deployment provides zero-trust role-based access control (RBAC) that limits what users can see and do within the fleets and resources they manage. You can integrate with Active Directory so users sign in with their existing corporate credentials, and local users can add multi-factor authentication (MFA) for another layer of verification. Access governance stays aligned with your organization's broader identity and access management practices.

- [Learn about identity and access management in NetApp Console local deployment.](#)
- [Learn about secure access to NetApp Console local deployment.](#)

Organize fleets and delegate storage administration

NetApp Console local deployment scales administration by organizing resources into folders and fleets. You can map fleets to environments, business units, or regions, then delegate administration at organization, folder, or fleet scope to keep ownership clear. This structure helps large storage teams distribute work without losing policy consistency or governance.

- [Learn about folders and fleets in NetApp Console local deployment.](#)
- [Learn about storage fleets in NetApp Console local deployment.](#)

Track and export administrative activity for compliance

Every administrative action generates an audit record. Security and compliance teams can use these records to investigate incidents, show control effectiveness, and satisfy audit requirements. Export audit logs to your syslog server through a webhook for centralized monitoring, longer retention, and analysis. Because NetApp Console local deployment runs in your own environment, log data never leaves your infrastructure.

- [Audit NetApp Console local deployment activity.](#)

Provision storage consistently with storage class policies

NetApp Console local deployment enforces consistent storage behavior through storage classes—templates that bundle performance, capacity, and tiering policies into reusable definitions. Administrators define storage standards once. Administrators and automated workflows use those approved classes for all provisioning activity across your environment. This stops configuration drift, reduces the time junior administrators spend on provisioning decisions, and guarantees that every workload meets your organization's standards immediately. After provisioning, Console local deployment continues to monitor each workload for compliance.

- [Learn about managing storage with NetApp Console local deployment.](#)

Monitor storage class compliance and remediate drift automatically

NetApp Console local deployment monitors every workload against the storage class used to provision it. When a workload drifts—because of a direct cluster configuration change or degrading performance—NetApp Console local deployment flags the violation immediately. Administrators can follow guided remediation steps

or configure automated remediation to resolve common drift conditions without manual intervention. This continuous enforcement reduces audit risk and keeps your environment compliant between scheduled reviews.

Monitor storage health and receive alerts across fleets

NetApp Console local deployment gives you a single place to watch the health and performance of every cluster you manage. Fleet-wide dashboards display clusters and volumes that need attention. Custom dashboards let administrators build monitoring views that match their responsibilities. The Workload Analyzer correlates latency, throughput, resource utilization, and configuration changes over time to help you pinpoint root causes before issues affect workloads.

Configure email and webhook delivery channels so alerts reach the right teams when conditions change. Organization admins set up destinations, and storage admins apply them in alert rules so response paths match your operational model. This helps enterprise teams respond faster to capacity, performance, and protection events.

- [Learn about monitoring storage health in NetApp Console local deployment.](#)
- [Set up notification delivery channels in NetApp Console local deployment.](#)
- [Configure notification rules for alerts in NetApp Console local deployment.](#)

Provision storage and investigate alerts with natural language

NetApp Console local deployment can connect to your own large language model (LLM) to provide AI-assisted storage management. You can describe a workload in plain language to get a provisioning plan, ask the Console to investigate an active alert, or get contextual answers about your storage environment. Every AI action stays within the guardrails of your storage classes and the role-based access controls that apply to your account, and you must confirm sensitive operations before they proceed. Console local deployment sends requests only to the LLM endpoint your administrators configure.

- [Learn about LLM integration in NetApp Console local deployment.](#)

Back up and restore storage with NetApp Backup and Recovery

Beyond provisioning and monitoring, NetApp Console local deployment gives you access to NetApp Backup and Recovery so you can protect your data from the same interface. You can back up and restore your data to support your protection and recovery requirements. Managing data protection alongside your storage keeps governance consistent and reduces the number of tools your teams need to operate.

- [Learn about data services in NetApp Console local deployment.](#)

Learn about NetApp Console local deployment identity and access management

Identity and access management (IAM) in NetApp Console local deployment controls who can sign in, how identities are verified, what resources users can access, and what actions they can perform. In NetApp Console local deployment, IAM includes role-based access control (RBAC), multi-factor authentication (MFA), and directory-backed authentication, as well as the hierarchy and audit model that support delegated administration.

Use this page to understand the NetApp Console local deployment IAM model at a high level. For detailed hierarchy planning examples, [Learn about folders and fleets in NetApp Console local deployment.](#)



You must have the *Super admin*, *Organization admin*, or *Folder or fleet admin* roles to manage IAM in NetApp Console.

What IAM means in NetApp Console local deployment

NetApp Console local deployment IAM combines identity verification, access control, delegation, and auditability:

- *Authentication* determines how users sign in, whether with local credentials or through your directory configuration.
- *Authorization* determines what members can do after they sign in, based on predefined roles and the scope where you assign them.
- *Hierarchy and scoping* determine which folders, fleets, and resources a member can access.
- *Member and credential management* determine how you add users, service accounts, and how you manage MFA and service account secrets.
- *Audit and compliance tracking* records IAM-related activity so you can review who changed access and when.

How authentication works

NetApp Console local deployment supports both local identities and external identity integration.

You can integrate NetApp Console local deployment with Active Directory so users sign in with their existing corporate credentials. This eliminates the need for separate passwords in Console local deployment and lets administrators look up directory users when assigning access.

For local users, MFA adds another verification step to reduce the risk of unauthorized access if credentials are compromised.

To learn more about authentication and secure sign-in options, [Learn about secure access in NetApp Console local deployment](#).

How authorization works

After a user signs in, NetApp Console local deployment uses RBAC to determine what that user can see and do.

Active Directory or LDAP integration handles authentication. NetApp Console local deployment authorization remains separate: administrators assign predefined roles at the organization, folder, or fleet level to control what each member can access.

The same role grants different effective access depending on the scope where you assign it. This model lets you give broad access to central administrators while limiting regional or team-level administrators to the fleets they manage.

Roles that you assign at the organization or folder level are inherited by child scopes. This inheritance model is a key part of how NetApp Console delegates access across folders and fleets.

NetApp designs NetApp Console local deployment roles with least-privilege principles so each role includes only the permissions needed for its tasks.

[Learn about role-based access control and role inheritance in NetApp Console local deployment](#).

How the IAM hierarchy works

NetApp Console local deployment uses a hierarchy to group resources and scope access. The organization is at the top, then folders, then fleets, and then the resources (including possible sub-folders) associated with those fleets.

This hierarchy is what makes delegation possible. For example, an Organization admin can manage access across the entire organization, while a Folder or fleet admin can manage only the resources and members within the part of the hierarchy they own.

NetApp Console IAM is built on three types of components: organizational components that define the hierarchy, resources that are assigned within that hierarchy, and members and roles that control who can access what.

Because roles inherit through this hierarchy, your folder and fleet design directly affects how broadly each access assignment applies.

[Learn how to add fleets to your organization..](#)

Member security and credentials

IAM in NetApp Console local deployment also includes operational controls for securing member access after accounts exist.

For local users, administrators can help users recover access by directing password reset, removing or temporarily disabling MFA, and reviewing local-user MFA behavior. For service accounts, administrators can recreate credentials when secrets are lost or rotated.

These controls are part of IAM because they determine how identities remain secure over time, not just how access is assigned initially.

[Learn how to manage member security.](#)

Audit IAM activity

IAM in NetApp Console local deployment includes the ability to review and export access-related activity.

From the Audit page, you can review actions related to managing your organization, such as adding members, creating fleets, and associating resources. You can also filter audit records by the Tenancy service to focus on IAM-related changes, or export audit logs to an external system.

Auditability is an important IAM principle because it lets you verify who changed access, when the change happened, and whether the change was successful.

[Learn how to audit NetApp Console local deployment activity.](#)

Next steps with IAM in NetApp Console

- [Get started with identity and access in NetApp Console](#)
- [Learn about secure access in NetApp Console local deployment](#)
- [Learn about RBAC in NetApp Console local deployment](#)
- [Monitor or audit Console local deployment activity](#)
- [Learn about the API for NetApp Console IAM](#)

Learn about fleet management in NetApp Console local deployment

Fleet management in NetApp Console local deployment is the practice of organizing storage systems into logical groups so you can apply consistent policies, control access, and monitor health across related clusters. Instead of managing each cluster independently, fleet management lets you treat your fleet as a common pool of resources to support your data storage goals.

As your storage environment grows, managing individual clusters becomes impractical. Fleet management solves this by giving you a structured way to group related systems, delegate responsibilities, and ensure every cluster operates under the same standards.

Why fleet management matters

Fleet management addresses three core challenges:

- **Simplification through abstraction** - Fleet management abstracts away the complexity of managing individual storage infrastructure. Instead of dealing with cluster-by-cluster configuration, you manage your storage estate as a unified whole, reducing operational overhead and decision fatigue.
- **Consistency and scalability** - Without clear organization, different teams make different decisions for similar workloads, leading to fragmented configurations. Fleet management lets you apply standards across dozens of systems at once, ensuring new workloads start from the same baseline across teams and fleets.
- **Governance and control** - As teams grow, you need clear boundaries for who can manage what. Fleet management provides those boundaries through organizational structure and access control, ensuring that storage decisions remain governed and auditable.

How fleets organize your resources

Your organization's resource hierarchy consists of folders and fleets:

For a detailed overview of the hierarchy and access model, see [Learn about folders and fleets in NetApp Console local deployment](#).

- **Organization** - The top level representing your company.
- **Folders** - Help you organize related fleets. For example, you might create a folder for each region or business unit, then create fleets within each folder. Folders can contain other folders or fleets. When you assign a role at the folder level, members inherit that role for all fleets within the folder.
- **Fleets** - Group the storage systems you manage together. You associate storage systems with fleets and assign members to fleets to grant them access.



Folders are an organizational tool and are not visible to members who do not have IAM permissions such as Organization admin, Folder admin, or Fleet admin roles. Members access fleets, not folders.

Common fleet organization patterns

How you organize fleets depends on your operational model:

- **By environment** - Create separate fleets for production, development, and test workloads. This keeps production storage under tighter policies while allowing more flexibility in lower environments.

- **By business unit** - Group clusters that serve a specific department, such as finance, engineering, or HR, into a dedicated fleet. You can then assign storage management responsibilities to team members within that business unit without exposing other fleets to them.
- **By location or data center** - When clusters are distributed across sites, organizing by location lets you monitor site-level health, apply region-specific policies, and track capacity consumption per site.
- **By application tier** - Group clusters that host a specific class of workload, such as databases, file shares, or virtual machines, so you can apply consistent standards tuned to that workload type across the entire fleet.



Use folders to add another level of organization. For example, create a folder for each region and then create environment-based fleets within each regional folder.

How fleet management enables access control

Fleets and folders serve as boundaries for user access and administrative responsibility:

- **Folder-level access** - When you assign a role at the folder level, the member inherits that role for all fleets and resources within the folder. This lets you delegate administrative responsibilities without granting access to the entire organization.
- **Fleet-level access** - When you assign a role at the fleet level, the member has access only to the storage systems within that fleet. This lets you delegate storage management to team members or application owners without exposing unrelated parts of your infrastructure.

Console local deployment provides fleet-level health and performance monitoring so you can see the status of all clusters in a fleet from a single view, identify issues early, and act before they affect workloads.

How storage management works

Storage management is the practice of defining storage standards, applying them consistently, and operating workloads so they stay aligned over time. In Console local deployment, those standards are expressed as storage class policies and storage classes, scoped to fleets, and enforced through provisioning workflows governed by RBAC and audit logging.

Console local deployment connects four concepts into one workflow:

- **Fleets** define the scope where you manage storage. A fleet groups systems so you can control access, apply standards consistently, and manage resources as a unit.
- **Storage class policies** define standards as reusable building blocks (performance, capacity, security, data protection).
- **Storage classes** express workload intent. A storage class is a named template assembled from one or more policies.
- **Provisioning workflows** create storage within guardrails. Different workflows make configuration decisions differently, but all can enforce storage classes and remain governed by RBAC and audit logging.

Provisioning approaches

Console local deployment supports three provisioning approaches, all governed by RBAC, audit logging, and storage class standards:

- **Manual provisioning** - You select the target system and specify configuration details directly. Use this when you need explicit control over system selection and configuration.

- **Automated provisioning** - You provide workload inputs and optionally select a storage class. Console local deployment recommends best-fit placement and applies class-driven settings to reduce manual decisions.
- **AI-assisted (agentic) provisioning** - You describe what you need in natural language. Console local deployment proposes a plan constrained by RBAC and storage classes, then provisions only after you review and approve.

Where to go next

- To understand storage standards, see [Learn about storage classes and policies in NetApp Console local deployment](#).
- To create and manage fleets, see [Manage folders and fleets](#).
- [Provision storage manually](#)
- [Provision storage automatically](#)
- [Provision storage with AI assistance](#)

Learn about storage classes and policies in NetApp Console local deployment

A storage class is a reusable template that defines how storage should behave. When you provision storage using a storage class, NetApp Console local deployment enforces the standards encoded in that class automatically without requiring administrators to make individual configuration decisions for every workload.

Storage classes are built from storage class policies, which define individual dimensions of storage behavior. Together, they let you capture your organization's storage standards once and apply them consistently across fleets.

What storage class policies are

A storage class policy defines one dimension of storage behavior. Policies are reusable building blocks that you combine to create storage classes.

Common policy types include:

- **Performance policies** - Define latency targets, IOPS, or throughput requirements.
- **Capacity policies** - Define provisioning mode, threshold alerts, or growth limits.
- **Security policies** - Define encryption, compliance, or access control requirements.
- **Data protection policies** - Define snapshot schedules, replication targets, or retention periods.

You define policies independently, then combine them when you build a storage class. This lets you reuse the same performance policy across multiple classes, or update a capacity policy in one place and apply that change to every class that uses it.

What storage classes are

A storage class is a named template assembled from one or more policies. It captures your organization's storage standards for a specific type of workload.

For example, you might create a **Production Database** storage class that combines high-performance, high-availability, and strict data protection policies, or a **Development File Share** storage class that combines

moderate performance, flexible capacity, and basic data protection policies.

Storage administrators define storage classes to encode enterprise requirements. All provisioning activity, whether done manually, through guided workflows, or with automation, draws from the library of classes those administrators have approved.

How storage classes enforce standards

When you provision storage, you select a storage class rather than configuring individual settings. Console local deployment uses the policies in that class to identify which clusters in your fleet have the capacity and performance headroom to support the workload, recommend the best placement option, and apply class-driven settings automatically at provisioning time.

After provisioning, Console local deployment continuously evaluates each workload against its storage class standards.

Storage class drift and compliance

When a workload no longer matches its storage class intent, Console local deployment flags it for investigation and remediation.

Issues fall into two categories:

- **Configuration drift:** Storage settings governed by the storage class policy no longer match the intended configuration. This can occur when changes are made directly on the ONTAP cluster or outside standard provisioning workflows.
- **Performance non-compliance:** The workload's runtime behavior no longer meets the storage class performance intent (for example, sustained pressure near a QoS limit or elevated tail latency).

An analysis view explains what changed and why. Guided remediation actions (for example, **Fix it**) may be available to help restore compliance. Some remediations can be applied in place, while others may require aligning the workload to a different storage class to restore the intended behavior.

Where to investigate

You can typically investigate drift and non-compliance from one of these locations (availability depends on your deployment and permissions):

- **Storage classes:** Drift / Compliance
- **Alerts:** Analyze

For step-by-step instructions, see [Remediate storage class drift](#).

End-to-end workflow

The following steps describe the process for using storage classes to standardize and govern storage in your environment:

1. **Create storage class policies** - Policies define the individual dimensions of storage behavior (performance targets, capacity settings, security requirements, data protection schedules). Create policies before building a storage class.
2. **Create a storage class** - Assemble a storage class by combining one policy from each applicable category. You can use a predefined storage class or create a custom one for your organization's standards.
3. **Associate the storage class with a fleet** - After you create a storage class, associate it with the fleet

where it will be used for provisioning and compliance monitoring.

4. **Provision storage using the storage class** - When provisioning a volume or LUN, select a storage class instead of configuring individual settings. Console local deployment uses the class to recommend the best-fit cluster placement, then provisions with the settings defined in the class.
5. **Monitor workloads for drift** - Console local deployment continuously evaluates each workload against the standards encoded in its storage class. If configuration drift or performance non-compliance occurs, it flags the issue and may raise an alert.
6. **Remediate drift to restore compliance** - When drift or performance non-compliance is detected, you can follow guided steps to correct the issue manually, let Console local deployment resolve common drift conditions automatically, or disassociate a workload from its storage class if you need to change its settings.

How storage classes work with fleets

Storage classes are associated with fleets. When you associate a storage class with a fleet, that class becomes available for all provisioning within the fleet. This ensures that every workload provisioned in the fleet follows the same standards, making fleets the primary way to enforce consistent storage behavior across related clusters.

For details on how to organize fleets, see [Learn about fleet management in NetApp Console local deployment](#).

Where to go next

- To understand fleet organization, see [Learn about fleet management in NetApp Console local deployment](#).
- To create policies and storage classes, see [Manage storage classes and policies](#).
- [Provision storage manually](#)
- [Provision storage automatically](#)
- [Provision storage with AI assistance](#)
- To analyze and remediate drift, see [Remediate storage class drift](#)

Learn about LLM integration in NetApp Console local deployment

NetApp Console local deployment connects to a large language model (LLM) for AI-assisted storage management. After setup, users can use natural language to provision storage, investigate alerts, and get storage guidance.



This documentation regarding connecting an LLM to the Console assistant to enable AI features is provided as a technology preview. With this preview offering, NetApp reserves the right to modify offering details, contents, and timeline before General Availability.

AI-assisted management lets users describe requests in plain language while Console local deployment applies your storage policies.

Console local deployment sends LLM requests only to the endpoint you configure.

What you can do with AI-assisted storage management

When you enable LLM integration, Console local deployment provides three capabilities through the Console assistant:

- **Storage provisioning** Describe workload requirements in plain language. Console local deployment recommends and runs a provisioning plan based on your storage classes.
- **Alert response** Ask Console local deployment to investigate an active alert. It analyzes the issue and suggests or runs an action based on assigned permissions.
- **Search and guidance** Ask questions about your storage environment or ONTAP administration. Console local deployment returns contextual answers from documentation and current fleet state.

Choose read-only or read/write access for the Console assistant

Users choose an assistant access mode based on the outcome they want:

- **Read only** mode for guidance and investigation without making infrastructure changes.
- **Read/write** mode for guided execution. Console local deployment shows the proposed action, requests confirmation, and then runs the change.

Understand what controls Console assistant actions

Console local deployment controls Console assistant actions through the same governance model used across the platform:

- Role-based access controls limit what each user can see and do.
- Storage policies constrain provisioning and related actions.
- The Console assistant requires confirmation before it runs storage-changing actions.

Choose an LLM provider path

NetApp Console local deployment supports these LLM provider types:

- **OpenAI** for direct OpenAI model access.
- **OpenAI-compatible** for a compatible endpoint, such as an enterprise gateway or proxy service.
- Anthropic for Anthropic's Claude models.

Model availability depends on the endpoint you configure. Select a model from the list in Console local deployment.

For supported model details, see [Learn about supported LLM providers and models in NetApp Console local deployment](#).

Understand where LLM requests go

Data flow depends on the endpoint path you choose:

- If you configure an OpenAI endpoint, Console local deployment sends prompts and context to the OpenAI endpoint.
- If you configure an OpenAI-compatible endpoint, Console local deployment sends prompts and context to the compatible endpoint that your organization configures.

Protect LLM credentials and control admin access

Protect the integration with these controls:

- Treat the API key as a privileged credential and rotate it according to your organization policy.
- Review provider-side usage and alerts to detect unexpected activity.

Connect your LLM

After you make these decisions, continue with [Connect your LLM and enable the NetApp Console assistant](#).

If connection or runtime checks fail, see [Troubleshoot LLM integration](#).

Learn about NetApp Backup and Recovery in NetApp Console local deployment

NetApp Backup and Recovery is a data service that provides efficient, secure, and cost-effective data protection for all your ONTAP workloads, including volumes, databases, virtual machines, and Kubernetes workloads.

A workload is a logical grouping of resources within a system that is managed as a single unit for protection, governance, detection, and recovery. In Backup and Recovery, a workload is the unit you protect. Its meaning depends on the data source: for Kubernetes a workload is an application, for VM environments it's a virtual machine, and for database environments it's a database.

Support for Backup and Recovery is already built in to all ONTAP systems, so there is no need for additional hardware or media gateways. This makes backup operations simple and cost effective. The NetApp Console simplifies the implementation of any backup strategy, including the full spectrum of 3-2-1 backup variants, without needing multiple resource managers or specialized personnel.



NetApp Backup and Recovery is in preview mode for NetApp Console local deployment. The service is not yet generally available, and the features and functionality are subject to change.

[Learn more about NetApp Backup and Recovery..](#)

Install and set up

Quick start for setting up NetApp Console local deployment

After you install NetApp Console local deployment, set up your organization so that the right teams can securely manage the right storage resources. Use this checklist to plan your structure, organize resources, add members, configure integrations, and enable data services.

Required access roles

Super admin or Organization admin. [Learn about access roles](#).



Install NetApp Console local deployment

- Review the installation workflow to understand the deployment process. [Learn about the installation workflow for NetApp Console local deployment](#).
- Install NetApp Console local deployment in your vCenter. [Install NetApp Console local deployment](#).

2

Plan your organization

- Learn how folders and fleets organize your resources. [Learn about folders and fleets.](#)
- Understand how role-based access control (RBAC) grants members the access they need. [Learn about role-based access control in NetApp Console local deployment.](#)
- Review the identity and access management workflow. [Get started with IAM for fleet and resource management.](#)

3

Set up your organization

- Add your storage systems to NetApp Console local deployment. [Add storage systems.](#)
- Create the folder and fleet hierarchy that matches your business structure. [Create folders and fleets.](#)
- Associate resources with the correct folders and fleets. [Associate resources to folders and fleets.](#)
- Add members and assign their initial roles. [Add members and assign roles.](#)

4

Configure integrations and services

- Integrate with Active Directory so that directory users can access the Console local deployment. [Integrate with Active Directory.](#)
- Connect your large language model (LLM) to enable the Console assistant and AI-assisted management. [Connect your LLM.](#)
- Configure how the Console local deployment delivers notifications. [Configure notification delivery.](#)

5

Set up NetApp Backup and Recovery

- [Obtain a license for NetApp Backup and Recovery.](#) You can skip this step if you do not want access to advanced Backup and Recovery services.
- Add the NetApp Backup and Recovery license to NetApp Console local deployment. [Add the license to NetApp Console local deployment.](#)
- [Grant users access to NetApp Backup and Recovery.](#)

Install NetApp Console

Install NetApp Console local deployment using an OVA in vCenter

You use an OVA to install a NetApp Console local deployment in your VCenter. The OVA download or URL is available through the NetApp Support site.

Prepare to install NetApp Console local deployment

Before installation, make sure your VM host meets the requirements and the NetApp Console local deployment can access the internet and targeted networks. To use NetApp data services such as NetApp Backup and Recovery, create cloud provider credentials for the NetApp Console local deployment to perform actions on your behalf.

Review NetApp Console local deployment host requirements

Make sure your host machine meets installation requirements before installing NetApp Console local deployment.

- CPU: 16 cores or 16 vCPUs
- RAM: 64 GB
- Disk space: 596 GB (thick provisioning recommended)
- vSphere 7.0u3 or higher, with virtual hardware version 19 or higher



Install NetApp Console local deployment in a vCenter environment rather than directly on an ESXi host.

Set up network access for the NetApp Console local deployment

NetApp Console local deployment uses fixed address space for inter-service communication. The IP range 10.42.0.0/16 and 10.43.0.0/16 are used for the internal network. Before deploying Console local deployment, ensure these IP ranges are not assigned to other VMs, DHCP scopes, DNS records, or load balancer VIP pools on the VLANs made available to the Console local deployment.

See Knowledge Base article Agent IP conflict with existing network for instructions on how to change the IP address of the agent’s interfaces.

Install NetApp Console local deployment in your VCenter environment

NetApp supports installing NetApp Console local deployment in your VCenter environment. The OVA file includes a pre-configured VM image that you can deploy in your VMware environment.

Download the OVA or copy the URL

Download the OVA.

1. Download the Console local deployment software from the NetApp Support Site.

Deploy NetApp Console local deployment in your VCenter

Log into your VCenter environment to deploy Console local deployment.

Steps

1. Upload the self-signed certificate to your trusted certificates if your environment requires it. You can replace this certificate after installation.
2. Deploy the OVA from the content library or local system.

From the local system	From the content library
a. Right-click and select Deploy OVF template....	a. Go to your content library and select the Console OVA.
b. Choose the OVA file from the URL or browse to its location, then select Next .	b. Select Actions > New VM from this template

3. Complete the Deploy OVF Template wizard to deploy the Console.

4. Select a name and folder for the VM, then select **Next**.
5. Select a compute resource, then select **Next**.
6. Review the details of the template, then select **Next**.
7. Accept the license agreement, then select **Next**.
8. Choose the type of proxy configuration you want to use: explicit proxy, transparent proxy, or no proxy.
9. Select the datastore where you want to deploy the VM, then select **Next**. Be sure it meets host requirements.
10. Select the network to which you want to connect the VM, then select **Next**. Ensure the network is IPv4 and has outbound internet access to the required endpoints.
11. In the **Customize template** window, complete the following fields:
 - **Proxy information**
 - If you selected explicit proxy, enter the proxy server hostname or IP address and port number, as well as the username, password.
 - If you selected transparent proxy, upload the respective certificate.
 - **Virtual Machine Configuration**
 - **Skip config check:** This check box is unchecked by default which means the NetApp Console local deployment runs a configuration check to validate network access.
 - NetApp recommends leaving this box unchecked so that the installation includes a configuration check of the NetApp Console local deployment. The Configuration check validates that the NetApp Console local deployment has network access to the required endpoints. If the deployment fails because of connectivity issues, you can access the validation report and logs from the NetApp Console local deployment host. In some cases, if you are confident that the NetApp Console local deployment has network access, you can choose to skip the check.
 - **Maintenance password:** Set the password for the `maint` user that allows access to the NetApp Console local deployment maintenance console.
 - **NTP servers:** Specify one or more NTP servers for time synchronization.
 - **Hostname:** Set the hostname for this VM. It must not include the search domain. For example, an FQDN of `console10.searchdomain.company.com` should be entered as `console10`.
 - **Primary DNS:** Specify the primary DNS server to use for name resolution.
 - **Secondary DNS:** Specify the secondary DNS server to use for name resolution.
 - **Search domains:** Specify the search domain name to use when resolving the hostname. For example, if the FQDN is `console10.searchdomain.company.com`, then enter `searchdomain.company.com`.
 - **IPv4 address:** The IP address that is mapped to the hostname.
 - **IPv4 subnet mask:** The subnet mask for the IPv4 address.
 - **IPv4 gateway address:** The gateway address for the IPv4 address.
12. Select **Next**.
13. Review the details in the **Ready to complete** window, select **Finish**.

The vSphere task bar shows the progress as the NetApp Console local deployment is deployed.

14. Power on the VM.

Plan your Console organization

Get started with identity and access in NetApp Console local deployment

In NetApp Console local deployment, set up your folder and fleet hierarchy, add members and starting permissions, and add and place resources in the correct fleets so the right teams can access and manage them.

Required access roles

Super admin, Organization admin, or Folder or fleet admin. [Learn about access roles.](#)

You need the **Organization admin** or **Super admin** role to complete initial setup. After setup, manage resources, member access, and fleet access as your organization changes.

1

Add or discover resources

Add systems to Console local deployment. During initial setup, systems are typically added while the default fleet is selected.

Learn how to create or discover resources:

- [On-premises ONTAP clusters](#)

2

Create your folder and fleet hierarchy

Create additional fleets and folders that match your business hierarchy.

[Learn how to organize your resources with folders and fleets.](#)

3

Associate resources with the correct fleets

Adding or discovering a system in Console local deployment automatically associates the resource with the currently selected fleet. To make a system available to the right teams, associate it with the appropriate fleets in your hierarchy.

- [Learn how to manage resources in folders and fleets.](#)

4

Add members and assign starting permissions

Add members and assign them roles at the organization, folder, or fleet level based on what they manage.

[Learn how to manage members and their permissions.](#)

After initial setup

After initial setup is complete, manage access and resource placement as your organization changes:

- [Manage resources in folders and fleets](#)
- [Manage member access across fleets and folders \(member-centric\)](#)

- [Manage who can access a specific fleet or folder \(fleet-centric\)](#)
- [Delete folders and fleets when no longer needed](#)

Related information

- [Learn about identity and access management in NetApp Console](#)
- [Learn about the API for identity and access](#)

Learn about folders and fleets in NetApp Console local deployment

In NetApp Console local deployment, use folders and storage fleets to organize your NetApp resources and control access according to your business structure—by location, department, or workload type.

Use this page for hierarchy strategy and fleet design patterns. For a complete IAM model of members, roles, and resource scope, [Learn about identity and access management in NetApp Console local deployment](#).

You organize resources in a hierarchy: the organization is at the top, then folders (which can contain other folders or fleets), and then fleets, which contain storage systems. Assign access roles at the organization, folder, or fleet level so users have the right access to resources.



You must have the Organization admin or Folder or fleet admin role to manage folders and fleets in the NetApp Console local deployment.

Organizational components

The organizational components—organization, folders, and fleets—form a hierarchy that determines how resources are grouped and how access is delegated.

Organization

An organization is the top level of the NetApp Console local deployment hierarchy and typically represents your company. Your organization consists of folders, fleets, members, roles, and resources. Resources are associated with fleets, and members are assigned roles at the organization, folder, or fleet level.

Folders

Folders group related fleets to organize them by location, site, or business unit. You can't associate storage systems directly with folders, but assigning a member a role at the folder level gives them access to all fleets in that folder.



Organization admins can add a resource to a folder to delegate the task of associating the resource with fleets to a Folder or fleet admin. [Associate resources with folders and fleets](#).

Fleets

Fleets group storage systems. Assign resources to fleets and grant members access at the fleet level. Resources can belong to multiple fleets. Members with fleet access can manage the resources in that fleet.

For example, you can associate an on-premises ONTAP system with a single fleet or with all fleets in your organization, depending on your needs.

[Learn how to create folders and storage fleets](#).

Resources

A resource is a storage system that the Console local deployment is aware of and that can be assigned to a fleet. Associate a resource with a fleet so that users with access to the fleet can manage the resource.

For example, you might associate an ONTAP cluster with one fleet or with all fleets in your organization. How you associate a resource depends on your organization's needs.

[Learn how to associate resources with fleets.](#)

Members and roles

Members are the users and service accounts in your organization. Roles define what actions they can perform and at what level of the hierarchy—organization, folder, or fleet.

Members

Members of your organization are user accounts or service accounts. A service account is typically used by an application to complete specified tasks without human intervention.

Users with the Organization admin role can add new members to your organization. All users (including service accounts and Active Directory users) must be explicitly added and granted at least one role in order to access Console local deployment.

[Learn how to add members to your organization.](#)

Access roles

The Console local deployment provides access roles that you can assign to the members of your organization.

When you associate a member with a role, you can grant that role for the entire organization, a specific folder, or a specific fleet. The role that you select gives a member permission to the resources in the selected part of the hierarchy.

The NetApp Console local deployment provides granular roles that adhere to the principle of least privilege, which means access roles are designed to give members access to only what they need.

Users can have multiple roles as their duties expand.

Role inheritance

When you assign a role at the organization or folder level, the child folders, fleets, and resources beneath it inherit that role, so your folder and fleet design determines how broadly each assignment applies.

[Learn about role inheritance in NetApp Console local deployment.](#)

Organization design examples

The right organizational structure depends on the size of your organization and how your teams are organized. The following examples show how different organizations can use the folder and fleet hierarchy to manage access effectively.

Small organization strategy

For organizations with fewer than 50 users and centralized storage management, consider a simplified approach using Super admin and Super viewer roles.

Example: ABC Corporation (5-person team)

- **Structure:** Single organization with 3 fleets (Production, Development, Backup)
- **Roles:**
 - 2 senior members: Super admin role for full administrative access
 - 3 team members: Super viewer role for monitoring without modification rights
- **Benefits:** Simplified administration, reduced role complexity, suitable for teams requiring broad access

Multi-regional enterprise strategy

For large organizations with regional operations and specialized teams, implement a hierarchical approach with folders representing geographical or business unit boundaries.

Example: XYZ Corporation (multinational company)

- **Structure:** Organization > Regional folders (North America, Europe, Asia-Pacific) > Fleets per region
- **Platform roles:**
 - 1 Organization admin: Global oversight and policy management
 - 3 Folder or fleet admins: Regional control (one per region)
 - 1 Federation admin: Corporate directory service integration
- **Storage roles by region:**
 - Storage admin: Discover and manage storage systems in assigned regions
 - Storage viewer: Monitor storage resources across regions
 - System health specialist: Manage storage health without system modifications
- **Benefits:** Enhanced security through role segregation, regional autonomy, and compliance with local regulations

Departmental specialization strategy

For organizations with specialized teams requiring specific access, use targeted role assignments based on functional responsibilities.

Example: TechCorp (mid-size technology company)

- **Structure:** Organization > Department folders (IT, Security, Development) > Fleet-specific resources
- **Specialized roles:**
 - Security team: Ransomware resilience admin and Classification viewer roles
 - Backup team: Backup and recovery super admin for comprehensive backup operations
 - Development team: Storage admin for test environment management
 - Compliance team: Operation support analyst for monitoring and support case management
- **Benefits:** Tailored access control, improved operational efficiency, and clear accountability for specialized tasks

Next steps

- [Create folders and storage fleets](#)

- [Associate resources with folders and fleets](#)
- [Manage fleet access assignments](#)
- [Monitor or audit Console local deployment actions](#)

Learn about role-based access control in NetApp Console local deployment

Manage user access to NetApp Console local deployment with role-based access control (RBAC), assigning predefined roles at the organization, folder, or fleet level. Each role grants specific permissions that define what actions users can perform within their assigned scope.

NetApp designs Console local deployment roles with least-privilege, so each role includes only the permissions needed for its tasks. This approach enhances security by limiting access to what each member requires.

After you organize resources into folders and fleets, assign organization members a role or roles for specific folders or fleets, that allow them to perform only their responsibilities.

For example, you can assign a member the Backup and Recovery admin role for a specific fleet level, allowing them to perform Backup and Recovery operations for resources within that fleet, without granting them broader access to the entire organization. This same user can be granted the role for several fleets within your organization.

You can assign members multiple roles for the same scope or different scopes, depending on their responsibilities. For example, a smaller organization might assign the same member both Storage admin and Backup and Recovery admin roles at the organization level, while a larger organization might have different members assigned to each role at the fleet level.

Types of Console organization members

NetApp Console local deployment supports the following types of members:

- *Users*: Individual users can either be local users you add to NetApp Console local deployment who use local credentials or directory users who authenticate through your integrated Active Directory or LDAP service. Both types of users can be assigned roles in NetApp Console local deployment to access resources.
- *Service accounts*: Non-human accounts that applications or services use to interact with NetApp Console local deployment through APIs. You can add service accounts directly to your Console local deployment organization.

[Learn how to add members to your organization.](#)

Predefined roles in NetApp Console local deployment

NetApp Console local deployment includes predefined roles that you can assign to organization members. Each role includes permissions that specify what actions a member can do within their assigned scope (organization, folder, or fleet).

NetApp Console local deployment roles use least-privilege principles that ensure members have only the permissions needed for their tasks, and categorizes roles by the type of access they provide:

- Platform roles: Provide Console local deployment administration permissions
- Data services roles: Provide permissions for managing specific data services, such as Ransomware

Resilience and Backup and Recovery

- Application roles: Provide permissions for managing storage as well as audit Console local deployment events and alerts

You can assign multiple roles to a member based on their responsibilities. For example, you might assign a member both the Storage admin role and the Backup and Recovery admin role for a specific fleet.

To choose access for a member, match the role category to the member's responsibilities, then assign the role at the scope (organization, folder, or fleet) that matches how broadly the member should have that access.

[Learn about how different organizations structure roles and scope in NetApp Console local deployment.](#)

[Learn about the predefined roles that you can assign to members in NetApp Console local deployment.](#)

How role inheritance works

When you assign a role at the organization or folder level, that role is inherited by the child scopes within that part of the hierarchy. This means that organization-level roles apply across the organization, folder-level roles apply to all child folders and fleets in that folder, and fleet-level roles apply only to the resources in that fleet.

Use the scope that matches the access you want the member to keep. For example, assign a role at the folder level when the member needs the same access across several fleets in one region. Assign the role at the fleet level when the member should access only one fleet.

You can't override inherited access at a lower scope. If a member inherits a role from the organization or from a folder, change that assignment at the higher scope where you originally granted it.

[Learn about folders and fleets in NetApp Console local deployment.](#)

[Manage member access.](#)

[Manage fleet access assignments.](#)

Set up your Console organization

Add folders and fleets to your NetApp Console local deployment organization

Create folders and fleets to match your business structure, control access, and organize resources in NetApp Console local deployment.

Required access roles

Super admin, Organization admin, or Folder or fleet admin. [Learn about access roles.](#)

NetApp Console local deployment creates one default fleet when you create an organization. You can add more fleets and group them with folders. [Learn about the resource hierarchy in NetApp Console.](#)

Using folders and fleets to organize resources

Folders and fleets are the two building blocks you use to shape your organization into something that matches how your teams actually work. For example, one folder per region with a production and a development fleet inside each.

- Folders group related fleets and support delegated administration and role inheritance.
- Fleets are where you associate resources for management and how you grant users operational access.

You can create folders and fleets first and add resources and member access later. This lets you plan your

resource hierarchy before adding users and resources in Console local deployment. If your structure is already defined, you can add resources and assign access when you create the fleet. You can update fleets at any time.

For example, put production clusters in a Production fleet and test clusters in a Test fleet, and grant each team access only to the fleet they own.

Folders

Folders organize fleets by business structure, such as business unit, region, or team. You can nest folders to mirror your organization.

Roles assigned at a folder level are inherited by child folders and fleets. You can also use a folder to delegate fleet assignment tasks to a Folder or fleet admin for that part of the hierarchy.



Members work in fleets, not folders. A resource associated only with a folder is not manageable by members until it is associated with a fleet.

For example, a regional enterprise could use folders to organize ONTAP storage by business unit and workload. It might create a top-level folder for North America, subfolders for Production and Development, and fleets for ONTAP clusters that host SAP, VMware, and backup volumes. Storage admins assigned to the North America folder inherit access to all child fleets, while application teams get access only to the fleets they manage.

Fleets

Associate resources with fleets so members can manage them. A fleet can exist directly under the organization or inside a folder.

For example, a healthcare company uses ONTAP storage in separate production and development fleets. The production fleet runs clinical apps and patient record databases, while the development fleet supports testing and validation. This separation protects live data, enforces tighter production access, supports auditability and least-privilege controls, and lets dev teams work without impacting patient-facing workloads.

Users navigate between assigned fleets on the **Systems** page to manage the resources associated with each fleet.

Add a folder or fleet

Add a fleet whenever you need a new boundary for resources and member access, and add a folder when you want to group related fleets and delegate their administration. For example, add a `Production` folder containing a `Production-US-East` fleet and a `Production-EU-West` fleet, then assign a regional lead the Folder or fleet admin role on just their fleet.

You can create up to seven hierarchy levels.

You can add resources and assign member access when you create a fleet or folder, or you can do it later.

Steps

1. Select **Administration > Identity and access**.
2. Select **Organization**.
3. From the **Organization** page, select **Add folder or fleet**.
4. Select **Folder** or **Fleet**.

5. In **Name and location**: Enter a name and choose a location for the folder or fleet.

1. Optional: Expand the **Resources** section to associate resources with the fleet or folder.

- a. Mark the check box next to the resource you want to associate and then select **Associate with the fleet**. If you haven't added systems to Console local deployment yet, you can do this step later.



Members can't access resources in a folder until those resources are assigned to a fleet.

2. Optional: Expand the **Access** section to assign access to members. Select **Add a member** to assign access and a role. By default, the user who creates the fleet has access to it.

[Learn about access roles.](#)

3. Select **Add**.

Rename a folder or fleet

Rename a folder or fleet as needed. Renaming does not affect associated resources or member access.

Steps

1. From the **Organization** page, navigate to a fleet or folder in the table, select **...** and then select **Edit folder** or **Edit fleet**.
2. On the **Edit** page, enter a new name and select **Apply**.

Delete a folder or fleet

Delete folders and fleets you no longer need, such as after team restructuring or fleet completion.

Before you delete a folder or fleet, complete the following checks:

- Verify that the folder or fleet does not contain resources. [Learn how to remove resource associations.](#)
- Review members who still have access to the folder or fleet. [View member access assignments.](#)
- Remove or update member access as needed. [Modify member access assignments.](#)
- If you are deleting a fleet, move resources to the target fleet first. [Learn how to move resources between fleets.](#)

Steps

1. From the **Organization** page, navigate to a fleet or folder in the table, select **...** and then select **Delete**.
2. Confirm that you want to delete the folder or fleet.

Manage fleets and folders in NetApp Console local deployment

After initial setup, you can do the following:

- **Manage resource associations for folders and fleets:** [Learn how to associate resources with fleets and folders.](#)
- **View or update access for one folder or fleet:** [Learn how to grant users access to fleets.](#)
- **Manage one member across multiple folders and fleets:** [Learn how to manage member access.](#)
- **Add additional members and assign starting permissions:** [Learn how to manage members and permissions.](#)

Related information

- [Learn about identity and access in NetApp Console](#)
- [Get started with identity and access in NetApp Console](#)
- [Manage fleet access assignments](#)
- [Learn about the identity and access API](#)

Add storage systems to NetApp Console local deployment

Add storage systems to NetApp Console local deployment to enable monitoring, inventory management, and administration of your storage infrastructure. After a storage system is added, Console local deployment discovers the system and begins collecting information that can be used for monitoring and management tasks.

Before you begin, ensure that you have the required permissions to add storage systems and that the storage system is reachable from Console local deployment.

Steps

1. Select **Storage > Management**.
2. From the Scope drop-down list, select the fleet to which you want to add a storage system.
3. Select **Add system**.
4. Enter the required storage system details, including connection information and credentials.
5. Review the information you entered and select **Add**.

The storage system is added to the selected fleet. Console local deployment begins discovering and monitoring the system. Depending on the environment and network connectivity, it might take several minutes for the system to appear as fully managed.



You can also add a storage system from the **Administration > Identity and access** page by selecting **Add system** and providing the required system information.

Manage resources in folders and fleets in NetApp Console local deployment

Manage resource access in NetApp Console local deployment by associating resources with the correct folder or fleet, which controls which teams can access and manage them.

Required access roles

Super admin, Organization admin, or Folder or fleet admin. [Learn about access roles](#).

Place resources where the right teams can manage them, move them when ownership changes, and remove associations when they are no longer needed.

A *resource* is an entity that Console local deployment recognizes, such as a storage resource or a Backup and Recovery workload.

Console resource types

Console local deployment supports both storage resources and Backup and Recovery workloads. Associate either resource type with a fleet to control access and management.

Storage resources

Storage resources are the most common type of resource in your organization. When you add a storage system to Console local deployment, associate it with a fleet so the appropriate teams can access and manage it. For example, after adding an ONTAP cluster, associate it with the `Production-US-East` fleet.

Backup and Recovery workloads

Backup and Recovery workloads are also considered resources. You can assign member permissions to access Backup and Recovery workloads by associating the workloads with fleets. For example, assign a member access to a Backup and Recovery workload by associating it with a fleet the member can access and granting the appropriate Backup and Recovery role.

View the resources in your organization

View the resources in your organization to find a specific resource and see which fleets or folders it is associated with. If you haven't created any folders or fleets, all resources are associated with the default fleet directly under the organization.

Steps

1. Select **Administration > Identity and access**.
2. Select **Resources**.
3. Select **Advanced Search & Filtering**.
4. Use the available options to find a resource:
 - **Search by resource name:** Enter a text string and select **Add**.
 - **Platform:** Select one or more platforms, such as Amazon Web Services.
 - **Resources:** Select one or more resources, such as Cloud Volumes ONTAP.
 - **Organization, folder, or fleet:** Select the entire organization, a specific folder, or a specific fleet.
5. Select **Search**.

Associate a resource with a folder or fleet

Associate a resource with a fleet or folder so the appropriate teams can manage it. For example, after adding an ONTAP cluster, associate it with the `Production-US-East` fleet so the regional Storage admins for that fleet can manage it.



Users with the Organization admin role can add resources to a folder to delegate fleet association tasks to the Folder or fleet admin for that folder. [Associate a resource with a folder](#).

Steps

1. From the **Resources** page, navigate to a resource in the table, select **...** and then select **Associate to a folder or fleet**.
2. Select a folder or fleet and then select **Accept**.
3. To associate an additional folder or fleet, select **Add folder or fleet** and then select the folder or fleet.

Note that you can only select from the folders and fleets for which you have admin permissions.

4. Select **Associate resources**.
 - If you associated the resource with fleets, members who have permissions for those fleets now have the ability to access the resource from the Console.

- If you associated the resource with a folder, a *Folder or fleet admin* can now access the resource and associate it with a fleet within the folder. [Associate a resource with a folder](#).

View the folders and fleets associated with a resource

Verify which fleets or folders a resource is associated with.



To see which members have access to the resource, review the members assigned to the associated folder or fleet. [Manage fleet access assignments](#).

Steps

1. From the **Resources** page, navigate to a resource in the table, select **...** and then select **View details**.

The following example shows a resource that is associated with one fleet.

Folders (0) Project (1)		Associate to folder or project
Type	Associated folders or projects	
	MyOrganization	
	MyOrganization > Project1	



To see which members have access to the resource, review the associated folder or fleet.

[Manage fleet access assignments](#).
[Manage member access](#).

Remove a resource from a folder or fleet

Remove a resource's association with a folder or fleet to prevent members from managing it there.



To remove a storage system from the entire organization, go to the **Systems** page and remove the system.

Steps

1. From the **Resources** page, navigate to a resource in the table, select **...** and then select **View details**.
2. To remove a resource from a folder or fleet, select next to the folder or fleet.
3. Select **Delete** to remove the association.

Move a resource between fleets

Move a resource from one fleet to another by removing its association with the current fleet and then associating it with the target fleet.



Access to the resource changes as soon as the association changes. If possible, complete both steps in one maintenance window.

Steps

1. From the **Resources** page, navigate to a resource in the table, select **...** and then select **View details**.
2. Select  next to the current fleet and select **Delete**.
3. Select **Add folder or fleet**.
4. Select the target fleet and select **Accept**.
5. Select **Associate resources**.

Related information

- [Learn about identity and access in NetApp Console](#)
- [Get started with identity and access in NetApp Console](#)
- [Manage fleet access assignments after moving resources](#)
- [Learn about the API for identity and access](#)

Add members to your NetApp Console local deployment organization

Add members to your NetApp Console local deployment organization and assign roles to grant access at the organization, folder, or fleet level for users, service accounts, and directory users.

Required access roles

Super admin, Organization admin, or Folder or fleet admin (for folders and fleets that they are administering). [Learn about access roles](#).

Before you begin

- Create the folder and fleet hierarchy that matches your organization. [Learn how to organize your resources with folders and fleets](#).
- Associate resources with the correct fleets before you assign member access. [Learn how to manage resources in folders and fleets](#).
- If you are adding a directory user, integrate your directory service first. [Learn how to integrate with your directory service](#).

Understand how access is granted in NetApp Console local deployment

NetApp Console uses role-based access control (RBAC) to manage permissions. Assign roles to members to provide the required access. Each role defines allowed actions for specific resources.

Note the following about granting access in NetApp Console local deployment:

- You must explicitly add each user to your organization and assign at least one role before that user can access resources.
- A role that you assign at the organization or folder level is inherited by child scopes, so choose the assignment scope carefully to give the member access only where needed. [Learn about role inheritance in NetApp Console local deployment](#).

Add members to your organization

NetApp Console supports three types of members: user accounts, directory users, and service accounts.



You must first configure an email server to use with Console local deployment before you can add users. [Learn how to configure an email server.](#)

Directory users authenticate through your integrated directory service, but you must still add each directory user individually and assign roles in Console local deployment. Create service accounts directly in the Console.

All members must have at least one role explicitly assigned to them in order to access Console local deployment.

When adding a member, choose the organization, folder, or fleet where the member needs access and assign the starting role or roles they need.

Add a user account

You must have the Organization admin or Folder or fleet admin role to add a user to an organization, folder, or fleet so they can access resources.

Steps

1. Select **Administration > Identity and access**.
2. Select **Members**.
3. Select **Add a member**.
4. For **Member Type**, keep **User** selected.
5. For **User's email**, enter the user's email address that is associated with the login that they created.
6. Use the **Select an organization, folder, or fleet** section to choose where the member needs access.

Note the following:

- You can select only the folders and fleets for which you have permissions.
 - When you select an organization or folder, you grant the member permissions to all its contents.
 - You can only assign the **Organization admin** role at the organization level.
7. **Select a category** then select a **Role** that gives the member the starting permissions they need for the organization, folder, or fleet you selected.

[Learn about access roles.](#)

8. To give access to more folders, fleets, or roles, select **Add role**, choose the folder, fleet, or role category, and select a role.
9. Select **Add**.

The Console emails instructions to the user.

Add a service account

Add a service account when you need to automate tasks or connect securely to Console APIs. After you create the account, copy its client ID and client secret for the integration that will use it.

Steps

1. Select **Administration > Identity and access**.

2. Select **Members**.
3. Select **Add a member**.
4. For **Member Type**, select **Service account**.
5. Enter a name for the service account.
6. Use the **Select an organization, folder, or fleet** section to choose where the service account needs access.

Note the following:

- You can only select from the folders and fleets for which you have permissions.
 - Selecting an organization or folder grants the member permissions to all its contents.
 - You can only assign the **Organization admin** role at the organization level.
7. Select a **Category** then select a **Role** that gives the service account the starting permissions it needs for the organization, folder, or fleet you selected.

[Learn about access roles.](#)

8. To give access to more folders, fleets, or roles, select **Add role**, choose the folder, fleet, or role category, and select a role.
9. Download or copy the client ID and client secret.

The Console shows the client secret only once. Copy it securely; you can recreate it later if you lose it.

10. Select **Close**.

Add a directory user to your organization

Add a user from your integrated directory service and grant their first roles. For example, add a new storage engineer from Active Directory and assign the Storage admin role on the `Production-US-East` fleet so they can work in that fleet.

You must first configure your directory service before you can add directory users. [Learn how to integrate with your directory service.](#)

Steps

1. Select **Administration > Identity and access**.
2. Select **Members**.
3. Select **Add a member**.
4. For **Member Type**, select **Directory user**.
5. For **User's email**, enter the email address of the directory user that you want to add. This email address must match the email address associated with the user's login in your directory.
6. Use the **Select an organization, folder, or fleet** section to choose where the directory user needs access.

Note the following:

- You can only select from the folders and fleets for which you have permissions.
- Selecting an organization or folder grants the member permissions to all its contents.
- You can only assign the **Organization admin** role at the organization level.

7. Select a **Category** then select a **Role** that gives the directory user the starting permissions they need for the organization, folder, or fleet you selected.

[Learn about access roles.](#)

8. To give the user additional access to other folders, fleets, or roles, select **Add role**, choose the folder or fleet, role category, and select a role.

Access roles

NetApp Console local deployment access roles

Identity and access management (IAM) in NetApp Console local deployment provides predefined roles that you can assign to the members of your organization across different levels of your resource hierarchy. Before you assign these roles, you should understand the permissions that each role includes. Roles fall into the following categories: platform, application, and data service.

Platform roles

Platform roles grant NetApp Console local deployment administration permissions, including role assignment and user management. The Console local deployment has several platform roles.

Platform role	Responsibilities
Organization admin	Allows a user unrestricted access to all fleets and folders within an organization, add members to any fleet or folder, as well as perform any task and use any data service that does not have an explicit role associated with it. Users with this role manage your organization by creating folders and fleets, assigning roles, adding users, and managing systems if they have the proper credentials.
Folder or fleet admin	Allows a user unrestricted access to assigned fleets and folders. Can add members to folders or fleets they manage, as well as perform any task and use any data service or application on resources within the folder or fleet they are assigned.
Federation admin	Allows a user to create and manage directory service federations with the Console so users can authenticate with their existing directory credentials.
Federation viewer	Allows a user to view existing directory service federations with the Console. Cannot create or manage federations.
Super admin	Gives the user all of admin roles. This role is designed for smaller organizations that may not need to distribute Console responsibilities across multiple users.
Super viewer	Gives the user all viewer roles. This role is designed for smaller organizations that may not need to distribute Console responsibilities across multiple users.

Application roles

The following is a list of roles in the application category. Each role grants specific permissions within its designated scope. Users without the required application or platform role cannot access the respective application.

Application role	Responsibilities
Operation support analyst	Provides access to alerts and monitoring tools such as the Audit page.
Storage admin	Administer storage health and governance functions, discover storage resources, as well as modify and delete existing systems.
Storage viewer	View storage health and governance functions, as well as view previously discovered storage resources. Cannot discover, modify, or delete existing storage systems.
System health specialist	Administer storage and health and governance functions, all permissions of the Storage admin except cannot modify or delete existing systems.

Data service roles

The following is a list of roles in the data service category. Each role grants specific permissions within its designated scope. Users who do not have the required data service role or a platform role will be unable to access the data service.

Data service role	Responsibilities
Backup and recovery super admin	Perform any actions in NetApp Backup and Recovery.
Backup and recovery admin	Perform backups to local snapshots, replicate to secondary storage, and back up to object storage.
Backup and recovery restore admin	Restore workloads in the Backup and Recovery.
Backup and recovery clone admin	Clone applications and data in the Backup and Recovery.
Backup and recovery viewer	View Backup and Recovery information.
Classification viewer	Allows users to view NetApp Data Classification scan results. Users with this role can view compliance information and generate reports for resources that they have permission to access. These users can't enable or disable scanning of volumes, buckets, or database schemas. Data Classification does not have an admin role.
SnapCenter admin	Provides the ability to back up snapshots from on-premises ONTAP clusters using NetApp Backup and Recovery for applications. A member who has this role can complete the following actions: * Complete any action from Backup and Recovery > Applications * Manage all systems in the fleets and folders for which they have permissions * Use all NetApp Console services SnapCenter does not have a viewer role.

Related links

- [Learn about NetApp Console identity and access management](#)
- [Get started with identity and access in NetApp Console](#)
- [Add members and assign roles in a NetApp Console organization](#)
- [Learn about the API for NetApp Console IAM](#)

NetApp Console local deployment platform access roles

Assign platform roles to users to grant permissions to manage the NetApp Console local deployment, assign roles, add users, create fleets, and manage user authentication.

Example for organization roles for a large multi-national organization

XYZ Corporation organizes data storage access by region—North America, Europe, and Asia-Pacific—providing regional control with centralized oversight.

The **Organization admin** in XYZ Corporation's Console local deployment creates an initial organization and separate folders for each region. The **Folder or fleet admin** for each region organizes fleets (with associated resources) within the region's folder.

Regional admins with the **Folder or fleet admin** role actively manage their folders by adding resources and users. These regional admins can also add, remove, or rename folders and fleets they manage. The **Organization admin** inherits permissions for any new resources, maintaining visibility of storage usage across the entire organization.

Within the same organization, one user is assigned the **Federation admin** role to manage the federation of the organization with their corporate IdP. This user can add or remove federated organizations, but cannot manage users or resources within the organization. The **Organization admin** assigns a user the **Federation viewer** role to check federation status and view federated organizations.

The following tables indicate the actions that each Console local deployment platform role can perform.

Organization administration roles

Task	Organization admin	Folder or fleet admin
Create, modify or delete systems from the Console local deployment (add or discover systems)	Yes	Yes
Create folders and fleets, including deleting	Yes	No
Rename existing folders and fleets	Yes	Yes
Assign roles and add users	Yes	Yes
Associate resources with folders and fleets	Yes	Yes
Register for support and submit cases through the Console	Yes	Yes
Use data services that are not associated with an explicit access role	Yes	Yes
View the Audit page and notifications	Yes	Yes

Federation roles

Task	Federation admin	Federation viewer
Create and update directory service integrations	Yes	No
View federations and their details	Yes	Yes

Super admin and viewer roles

The **Super admin** role provides full access to manage Console features, storage, and data services. This role suits those overseeing administration and governance. In contrast, the **Super viewer** role offers read-only access, ideal for auditors or stakeholders who need visibility without making changes.

Organizations should use **Super admin** access sparingly to minimize security risks and align with the principle of least privilege. Most organizations should assign fine-grained roles with only the necessary permissions to reduce risk and improve auditability.

Example for super roles

ABC Corporation has a small team of five that leverages the NetApp Console for data services and storage management. Instead of distributing multiple roles, they assign the **Super admin** role to two senior team members who handle all administrative tasks, including user management and resource configuration. The remaining three team members are assigned the **Super viewer** role, allowing them to monitor storage health and data service status without the ability to modify settings.

Role	Inherited roles
Super admin	• Organization admin • Folder or fleet admin • Backup and recovery super admin • Storage admin
Super viewer	• Organization viewer • Backup viewer • Storage viewer

Operational support analyst access role in NetApp Console local deployment

In NetApp Console local deployment, you can assign the Operational support analyst role to users to provide them access to alerts and monitoring.

Operational support analyst

Task	Can perform
View storage systems	Yes
View the Audit page and notifications	Yes

Task	Can perform
View, download, and configure alerts	Yes

Storage access roles for NetApp Console local deployment

You can assign the following roles to users to provide them access to the storage management features in NetApp Console local deployment. You can assign users an administrative role to manage storage or a viewer role for monitoring.

Administrators can assign storage roles to users for the following storage resources and features:

Storage resources:

- On-premises ONTAP clusters

Console services and features:

- Storage health functions

Example for storage roles in NetApp Console local deployment

XYZ Corporation, a multinational company, has a large team of storage engineers and storage administrators. They allow this team to manage storage assets for their regions while limiting access to core Console local deployment tasks like user management and license management.

Within a team of 12, two users are given the **Storage viewer** role which allows them to monitor the storage resources associated with the Console local deployment fleets they are assigned to. The remaining nine are given the **Storage admin** role which includes the ability to manage software updates, access ONTAP System Manager through the Console local deployment, as well as discover storage resources (add systems). One person on the team is given the **System health specialist** role so they can manage the health of the storage resources in their region, but not modify or delete any systems. This person can also perform software updates on the storage resources for fleets they are assigned.

The organization has two additional users with the **Organization admin** role who can manage all aspects of the Console local deployment, including user management and license management, as well as several users with the **Folder or fleet admin** role who can perform Console local deployment administration tasks for the folders and fleets they are assigned to.

The following table shows the actions each storage role performs.

Feature and action	Storage admin	System health specialist	Storage viewer
Storage Management:			
Discover new resources (add systems)	Yes	Yes	No
View added systems	Yes	Yes	No
Delete systems from the Console	Yes	No	No

Feature and action	Storage admin	System health specialist	Storage viewer
Modify systems	Yes	No	No
System manager access			
May enter credentials	Yes	Yes	No

NetApp Backup and Recovery roles in NetApp Console local deployment

In NetApp Console local deployment, you can assign the following roles to users to provide them access to NetApp Backup and Recovery within the Console. Backup and Recovery roles give you the flexibility to assign users a role specific to the tasks they need to accomplish within your organization. How you assign roles depends on your own business and storage management practices.

The service uses the following roles that are specific to NetApp Backup and Recovery.

- **Backup and recovery super admin:** Perform any actions in NetApp Backup and Recovery.
- **Backup and recovery backup admin:** Perform backups to local snapshots, replicate to secondary storage, and back up to object storage actions in NetApp Backup and Recovery.
- **Backup and recovery restore admin:** Restore workloads using NetApp Backup and Recovery.
- **Backup and recovery clone admin:** Clone applications and data using NetApp Backup and Recovery.
- **Backup and recovery viewer:** View information in NetApp Backup and Recovery, but not perform any actions.

For details about all NetApp Console access roles, see [the Console setup and administration documentation](#).

Roles used for common actions

The following table indicates the actions that each NetApp Backup and Recovery role can perform for all workloads.

Feature and action	Backup and recovery super admin	Backup and recovery backup admin	Backup and recovery restore admin	Backup and recovery clone admin	Backup and recovery viewer
Add, edit, or delete hosts	Yes	No	No	No	No
Install plugins	Yes	No	No	No	No
Add credentials (host, instance, vCenter)	Yes	No	No	No	No
View dashboard and all tabs	Yes	Yes	Yes	Yes	Yes
Start free trial	Yes	No	No	No	No

Feature and action	Backup and recovery super admin	Backup and recovery backup admin	Backup and recovery restore admin	Backup and recovery clone admin	Backup and recovery viewer
Initiate discovery of workloads	No	Yes	Yes	Yes	No
View license information	Yes	Yes	Yes	Yes	Yes
Activate license	Yes	No	No	No	No
View hosts	Yes	Yes	Yes	Yes	Yes
Schedules:					
Activate schedules	Yes	Yes	Yes	Yes	No
Suspend schedules	Yes	Yes	Yes	Yes	No
Policies and protection:					
View protection plans	Yes	Yes	Yes	Yes	Yes
Create, modify, or delete protection plans	Yes	Yes	No	No	No
Restore workloads	Yes	No	Yes	No	No
Create, split, or delete clones	Yes	No	No	Yes	No
Create, modify, or delete policy	Yes	Yes	No	No	No
Reports:					
View reports	Yes	Yes	Yes	Yes	Yes
Create reports	Yes	Yes	Yes	Yes	No
Delete reports	Yes	No	No	No	No
Import from SnapCenter and manage host:					
View imported SnapCenter data	Yes	Yes	Yes	Yes	Yes
Import data from SnapCenter	Yes	Yes	No	No	No
Manage (migrate) host	Yes	Yes	No	No	No

Feature and action	Backup and recovery super admin	Backup and recovery backup admin	Backup and recovery restore admin	Backup and recovery clone admin	Backup and recovery viewer
Configure settings:					
Configure log directory	Yes	Yes	Yes	No	No
Associate or remove instance credentials	Yes	Yes	Yes	No	No
Buckets:					
View buckets	Yes	Yes	Yes	Yes	Yes
Create, edit, or delete bucket	Yes	Yes	No	No	No

Roles used for workload-specific actions

The following table indicates the actions that each NetApp Backup and Recovery role can perform for specific workloads.

Kubernetes workloads

This table indicates the actions that each NetApp Backup and Recovery role can perform for actions specific to Kubernetes workloads.

Feature and action	Backup and recovery super admin	Backup and recovery backup admin	Backup and recovery restore admin	Backup and recovery viewer
View clusters, namespaces, storage classes, and API resources	Yes	Yes	Yes	Yes
Add new Kubernetes clusters	Yes	Yes	No	No
Update cluster configurations	Yes	No	No	No
Remove clusters from management	Yes	No	No	No
View applications	Yes	Yes	Yes	Yes
Create and define new applications	Yes	Yes	No	No
Update application configurations	Yes	Yes	No	No
Remove applications from management	Yes	Yes	No	No

Feature and action	Backup and recovery super admin	Backup and recovery backup admin	Backup and recovery restore admin	Backup and recovery viewer
View protected resources and backup status	Yes	Yes	Yes	Yes
Create backups and protect applications with policies	Yes	Yes	No	No
Unprotect apps and delete backups	Yes	Yes	No	No
View recovery points and resource viewer results	Yes	Yes	Yes	Yes
Restore applications from recovery points	Yes	No	Yes	No
View Kubernetes backup policies	Yes	Yes	Yes	Yes
Create Kubernetes backup policies	Yes	Yes	Yes	No
Update backup policies	Yes	Yes	Yes	No
Delete backup policies	Yes	Yes	Yes	No
View execution hooks and hook sources	Yes	Yes	Yes	Yes
Create execution hooks and hook sources	Yes	Yes	Yes	No
Update execution hooks and hook sources	Yes	Yes	Yes	No
Delete execution hooks and hook sources	Yes	Yes	Yes	No
View execution hook templates	Yes	Yes	Yes	Yes
Create execution hook templates	Yes	Yes	Yes	No
Update execution hook templates	Yes	Yes	Yes	No
Delete execution hook templates	Yes	Yes	Yes	No
View workload summary and analytics dashboards	Yes	Yes	Yes	Yes

Feature and action	Backup and recovery super admin	Backup and recovery backup admin	Backup and recovery restore admin	Backup and recovery viewer
View StorageGRID buckets and storage targets	Yes	Yes	Yes	Yes

Configure Console integrations and services

Integrate NetApp Console local deployment with Active Directory

Integrate NetApp Console local deployment with Active Directory for directory-backed sign-in and user lookup. Use your directory service to authenticate users, then assign access to those users in NetApp Console local deployment.

Required access roles

Super admin or Organization admin. [Learn about access roles.](#)

When you integrate with Active Directory, Console local deployment authenticates users through your existing directory. After you add a directory user, assign Console access roles to control what that user can access.

Active Directory integration also helps you meet compliance requirements by applying your existing controls, audit trails, and policies to Console local deployment access.



- Active Directory integration does not create federated groups, does not synchronize directory-group assignments, and does not automatically grant access. Administrators still add directory users to the organization and assign roles in Console local deployment.
- Only one Active Directory integration is supported at a time. Once an integration is configured, the **Add integration** button is disabled. To switch to a different directory, disable or delete the existing integration first.
- Directory users do not configure or use multi-factor authentication (MFA). Console local deployment supports MFA only for local users. [Learn how to manage member security settings.](#)

Before you begin

Before you integrate with Active Directory, confirm that you have:

- An Active Directory domain and credentials that can query the directory
- The LDAP server address and the base distinguished name (DN) for the directory tree
- Network access from Console local deployment to your LDAP server on port 389 (LDAP) or 636 (LDAPS)
- SSL/TLS certificates, if you plan to use LDAPS

Steps

1. Select **Administration > Identity and access**.
2. Select **Directory services** to open the Federations page.
3. Select **Add integration**.
4. Enter the connection details for your Active Directory server:

- A descriptive name for the integration.
 - The **LDAP host**: hostname or IP address of your LDAP server. For multiple servers, enter the primary.
 - The port: 389 for LDAP or 636 for LDAPS.
 - The **Connection timeout** in milliseconds. The default is 1000 ms.
5. Select **Use SSL/TLS** to use LDAPS. Confirm that your LDAP server supports LDAPS and that the Console can access the required certificates.
 6. Test the connection. If it fails, check the LDAP host, port, network connectivity, and SSL/TLS certificates.
 7. After the test succeeds, enter the directory and user details:
 - **Base DN binding**: Enter the Bind DN for the LDAP/AD account this app will use to connect to the directory, and enter the Bind credential (password) for that account. Console uses these details to bind to LDAP and validate the connection.
 - **User DN**: a user account with permission to query the directory. For example, `CN=ldap-reader,OU=Service Accounts,DC=example,DC=com`.
 8. Select **Add** to save the integration.

After you finish

After you save the integration, add directory users to your organization and assign roles. You must configure and enable at least one Active Directory integration before you can add directory users. [Learn how to add directory users and assign roles](#).

Disable or enable an Active Directory integration

Disable an integration to temporarily suspend directory-backed authentication without deleting the configuration. When you disable a directory service, directory users cannot sign in through the directory.

Steps

1. Select **Administration > Identity and access**.
2. Select **Directory services** to open the Federations page.
3. In the integrations list, locate the integration and select the action menu for that row.
4. Select **Disable** to suspend the integration, or **Enable** to restore it.

Delete an Active Directory integration

Delete an integration to permanently remove the directory service configuration. Before deleting, review any warnings about directory users that are linked to the integration, as their access will be affected.

Steps

1. Select **Administration > Identity and access**.
2. Select **Directory services** to open the Federations page.
3. In the integrations list, locate the integration and select the action menu for that row.
4. Select **Delete** and confirm the deletion.

Connect your LLM and enable the NetApp Console assistant in NetApp Console local deployment

Connect your large language model (LLM) to NetApp Console local deployment to enable the Console assistant and AI-assisted storage management.



This documentation regarding connecting an LLM to the Console assistant to enable AI features is provided as a technology preview. With this preview offering, NetApp reserves the right to modify offering details, contents, and timeline before General Availability.

Required access roles

Super admin or Organization admin. [Learn about access roles.](#)

After setup, users open the Console assistant from the dashboard and choose an access mode:

- In **Read only** mode, the assistant answers questions and provides guidance without making changes.
- In **Read/write** mode, the assistant can act on storage infrastructure. It shows details for review and confirmation before each change. For asynchronous operations, such as volume creation, it creates a job that users can check from the assistant.

To connect your LLM, select a provider path, enter endpoint details and API key, and then select a model in **Administration > AI Tools**. Assistant actions stay within your storage policies and role-based access controls.

Gather what you need before you connect an LLM

Before you connect your LLM, gather the following:

- Your provider type decision: OpenAI or OpenAI-compatible. [Learn about provider choices.](#)
- A valid API key for the provider path you select.
- The endpoint URL for your provider path.
- Your proxy or gateway URL, if your organization requires one.
- Your user name or single sign-on (SSO) ID for the LLM provider account, if required.
- Network access from Console local deployment to the selected endpoint.
- Storage classes and role-based access controls for the assistant actions you plan to allow.

For supported model details, see [Learn about supported LLM providers and models in NetApp Console local deployment.](#)

Connect an LLM to NetApp Console local deployment

Enter provider settings, API key, and model to connect your LLM to Console local deployment.

Steps

1. Log in to the NetApp Console local deployment.
2. Select **Administration > AI Tools**.
3. Select the menu, and then select **Edit**.
4. In **Provider type**, select a provider type.

[Learn about LLM integration in NetApp Console local deployment.](#)

5. If prompted for a provider endpoint, enter the endpoint URL for the provider path you selected.
6. Enter the proxy or gateway URL and credentials.
7. In the **User name** field, enter your user name or SSO ID if your provider path requires it.
8. In the **API key** field, paste the API key from your selected provider path.

9. Select a model from the list.
10. Review the configuration, and then select **Save**.

If save or test fails, verify the provider type, endpoint URL, API key, and selected model, and then try again.

[Troubleshoot your LLM integration.](#)

Verify that LLM integration works

After you save the configuration, verify assistant availability and behavior.

Steps

1. Confirm that the **Console assistant** button appears on the NetApp Console local deployment dashboard.
2. Open the assistant in **Read only** mode and run a basic query.
3. Open the assistant in **Read/write** mode and confirm that storage-changing actions require user confirmation before execution.
4. Verify that the users who need action workflows have the required role-based access controls.

After you complete validation, users can open the Console assistant from the dashboard and describe tasks in plain language. For usage examples and end-user workflows, see [Use the NetApp Console assistant](#).

Protect credentials and monitor LLM usage

- Use a dedicated API key for Console local deployment integration when possible.
- Rotate API keys according to your organization policy.
- Restrict who can edit AI Tools settings to required admin roles only.
- Monitor provider-side API usage and alerts to track abnormal activity or cost spikes.

Troubleshoot LLM integration in NetApp Console local deployment

Use this quick triage flow to diagnose and resolve common LLM integration issues in NetApp Console local deployment.



This documentation regarding connecting an LLM to the Console assistant to enable AI features is provided as a technology preview. With this preview offering, NetApp reserves the right to modify offering details, contents, and timeline before General Availability.

If you have not completed setup, see [Connect your LLM and enable the NetApp Console assistant](#).

Triage LLM integration issues quickly

1. Validate AI Tools configuration in **Administration > AI Tools**.

Confirm provider type, endpoint URL, API key, selected model, and outbound network access.

2. Validate assistant availability on the dashboard.

Confirm the **Console assistant** button appears for expected users and organizations.

3. Validate runtime behavior.

Run a simple read-only request and confirm provider endpoint reachability, model availability, and provider service health.

Troubleshoot by symptom

Symptom	Likely cause	What to check	Next action
Save or test fails in AI Tools	Configuration or connectivity mismatch	Provider type, endpoint URL, API key format, selected model, and outbound access	Correct the value that fails validation and save again
Console assistant button is missing	Unhealthy integration status, org mismatch, or RBAC mismatch	Successful AI Tools save, integration status, current organization, and expected access role	Refresh the session and verify with a known-good admin account
Assistant opens but request fails	Provider or runtime path issue	Endpoint reachability, model availability on that endpoint, provider limits, and temporary provider status	Retry after fixing endpoint/model mismatch or provider-side limit issues
Assistant response is slow or inconsistent	Prompt size, endpoint performance, or network/proxy instability	Short prompt test, provider service health, and network/proxy stability	Use a shorter prompt, try another supported model, and stabilize network or proxy routing

Respond to LLM credential exposure or misuse

If you suspect API key exposure or unauthorized use:

1. Revoke the existing API key in your provider system.
2. Create a new API key.
3. Update the key in **Administration > AI Tools**.
4. Verify successful reconnection with a read-only assistant test.

Set up notification delivery channels in NetApp Console local deployment

In NetApp Console local deployment, you can set up multiple channels for alert notifications, including email and webhooks.

Required access roles

Super admin or Organization admin. [Learn about access roles](#).

By default, Console local deployment displays notifications through its built-in notification system. Configuring additional delivery channels lets you receive notifications in the way that best fits your workflow.

You can send all notifications through the same channels or use different channels for different notification types. For example, you might send urgent storage alerts through email while routing other alert traffic to a third-party monitoring tool through a webhook.

After you set up notification delivery channels, you can configure notification rules and assign them to different

channels. [Learn how to configure notification rules.](#)

Configure an email server

When you configure an email server, Console local deployment sends notifications, alerts, and user invitations through it. Console local deployment supports SMTP email servers, which can be your existing corporate email server or a third-party email service.

Steps

1. Select **Administration > Console**.
2. Select **Configuration** to open the System configurations page.
3. In the **Email server** section, select **Configure**.
4. Enter the connection details for your email server:
 - Add a sender name and a sender email address.
 - The **SMTP host**: hostname or IP address of your SMTP server. For multiple servers, enter the primary.
 - Select the connection protocol from the **Connection security** dropdown list. The port is configured for you and is read-only: 25 for SMTP with STARTTLS, or 465 for SMTPS.

SMTPS (port 465) establishes an encrypted connection immediately and is recommended for security-sensitive environments. STARTTLS (port 25) upgrades from an unencrypted connection and may fall back to unencrypted transmission if the encryption negotiation fails.

5. Select **Test email** to send a test email to a recipient you specify.
6. After the test succeeds, select **Save** to save the configuration.

If the test fails, re-verify the settings you entered and double-check that Console local deployment has network access to the email server.

Update an email server configuration

You can update an existing email server configuration if you want to use a different email server.

Steps

1. Select **Administration > Console**.
2. Select **Configuration** to open the System configurations page.
3. In the **Email server** section, select **Configure**.
4. Select **Clear fields** to clear the existing configuration.
5. Enter the new connection details for your email server.
6. Select **Test email** to send a test email to a recipient you specify.
7. After the test succeeds, select **Save** to save the new configuration.

If the test fails, re-verify the settings you entered and double-check that Console local deployment has network access to the email server.

Remove an email server configuration

You can remove the email server configuration if you no longer want Console local deployment to send emails.

Steps

1. Select **Administration > Console**.
2. Select **Configuration** to open the Systems configuration page.
3. In the **Email server** section, select **Configure**.
4. Select **Clear fields** to clear the existing configuration.
5. Select **Save** to save the cleared configuration, which removes the email server configuration from Console local deployment.

Configure a webhook

Set up webhooks to send notifications from Console local deployment to other tools or services. You can configure multiple webhooks, each pointing to a different endpoint. For example, one for a SIEM and another for an on-call paging service.

After you create a webhook, users with the Storage admin role can assign it to specific alert rules. For example, they can send critical alerts to email while sending all alerts to a third-party monitoring tool through a webhook.



Console local deployment does not enforce access control on webhook endpoints. Any user or system that can reach the endpoint URL receives the alert data. Ensure that access to the receiving application is restricted to authorized users through that application's own access controls.

Before you begin

Confirm that Console local deployment can reach the receiving application over the network, and gather the endpoint URL, HTTP method, and any authentication or certificate details required by that application.

Steps

1. Select **Administration > Console**.
2. Select **Configuration** to open the System configurations page.
3. In the **Webhook integration** section, select **Configure**.
4. Enter the required details for the webhook:
 - A descriptive name for the webhook.
 - The Endpoint URL provided by the receiving application.
 - HTTP method
 - Optional: A self-signed certificate in PEM format.
 - Any required headers or secrets (authentication credentials).
 - The format to use when sending the webhook. JSON and OTEL (OpenTelemetry) are supported.

Use JSON for general-purpose webhooks and custom REST endpoints. Use OTEL for observability platforms that natively consume OpenTelemetry signals, such as Grafana or other OpenTelemetry-compatible collectors.

5. Select **Test webhook** to test the webhook connection and ensure that Console local deployment can successfully send messages to the receiving application.
6. After the test succeeds, select **Save** to create the webhook.

After you save the webhook, it is available for users to select where webhooks are supported, such as Alert delivery options. [Learn how to create alert rules and assign webhooks for alert delivery.](#)

Use NetApp Console

Log in to NetApp Console local deployment

You can sign in to NetApp Console local deployment after your organization administrator has invited you to your Console local deployment organization user account. To login, you use the email address associated with your login.

If you log in but don't see any resources, contact your organization administrator. [Contact your organization administrator.](#)

Steps

1. Open a web browser and go to the IP address where the Console local deployment is installed.
2. On the **Log in** page, enter the email address that's associated with your login.
3. Depending on the authentication method associated with your login, you'll be prompted to enter your credentials.

- A NetApp Console account using your email address and a password



- If MFA is enabled: after entering your password, you'll be prompted to enter a TOTP code (from your authenticator app) or use a recovery code.
- If you are signing in with recovery codes, use them in the order shown in the UI during the recovery-code prompt.

- An Active Directory account using your email address and password

Switch between fleets in NetApp Console local deployment


In NetApp Console local deployment, if you have access to multiple fleets, you can switch between them at any time.



You can't switch to another fleet while you are viewing any of the **Identity and access** pages.

Steps

1. In the top header of the Console local deployment, select **Scope**.
2. Browse through the fleets in your organization, select the fleet that you want, and then select **Switch**.



Manage your NetApp Console local deployment user settings

In NetApp Console local deployment, you can modify your Console local deployment profile including change your password, enable multi-factor authentication (MFA), and see who your Console administrator is.

Change your display name

You can change your Console local deployment display name, which identifies you to other users. You cannot change your username or email address.

Steps

1. Select the profile icon in the upper right corner of the Console local deployment to view the User settings panel.
2. Select the **Edit** icon next to your name.
3. Enter your new display name in the **Name** field.

Configure multi-factor authentication

Configure multi-factor authentication (MFA) to improve security by requiring a second verification method when logging in to NetAppConsole local deployment.

Users who sign in through the NetApp Support Site cannot enable MFA. If this is true for you, you don't see the option to enable MFA in your profile settings.

Do not enable MFA if your user account is used for API access. Multi-factor authentication stops API access when enabled for a user account. Use service accounts for all API access.



You can't configure MFA for your account through the Console local deployment if your account uses Active Directory or another integrated directory service for authentication.

Before you begin

- You must have already downloaded an authentication app, such as Google Authenticator or Microsoft Authenticator, to your device.
- You'll need your password to set up MFA.



If you do not have access to your authentication app or lose your recovery code, contact your Console administrator for help.

Steps

1. Select the profile icon in the upper right corner of the Console to view the User settings panel.
2. Select **Configure** next to the **Multi-Factor Authentication** header.
3. Follow the prompts to set up MFA for your account.
4. When you finish, you'll be prompted to save your recovery codes. Choose to either copy the codes or download a text file containing the codes. Keep these codes somewhere safe. You need the recovery codes if you lose access to your authentication app.



If you need to sign in with a recovery code, use them in the order shown in the UI during the recovery-code prompt.

After you set up MFA, the Console local deployment prompts you to enter a one-time code from your authentication app each time you log in.

Regenerate your MFA recovery code

You can only use a recovery code once. If you lose yours, create a new one.

Steps

1. Select the profile icon in the upper right corner of the the Console local deployment to view the User settings panel.

2. Select  next to the **Multi-Factor Authentication** header.
3. Select **Regenerate recovery code**.
4. Copy the generated recovery codes and save them in a secure location.

Delete your MFA configuration

Deleting MFA removes the second verification step for your next sign-in. To use MFA again, you must configure it again from your user settings.



If you are unable to access your authentication app or recovery code, you will need to contact your Organization administrator to reset your MFA configuration.

Steps

1. Select the profile icon in the upper right corner of the Console local deployment to view the User settings panel.
2. Select  next to the **Multi-Factor Authentication** header.
3. Select **Delete**.

Contact your Organization administrator

If you need to contact your organization administrator, you can send an email to them directly from the Console. The administrator manages user accounts and permissions within your organization.



You must have a default email application configured for your browser to use the **Contact admins** feature.

Steps

1. Select the profile icon in the upper right corner of the Console local deployment to view the User settings panel.
2. Select **Contact admins** to send an email to your organization administrator.
3. Select the email application to use.
4. Finish the email and select **Send**.

Configure dark mode (dark theme)

You can set the Console local deployment to display in dark mode.

Steps

1. Select the profile icon in the upper right corner of the Console local deployment to view the User settings panel.
2. Move the **Dark theme** slider to enable it.

Use the NetApp Console assistant in NetApp Console local deployment

Use the NetApp Console assistant in NetApp Console local deployment to manage storage and get answers with natural language. Describe what you need in plain language, and NetApp Console local deployment acts on it within your storage policies and role-based access.



This documentation regarding connecting an LLM to the Console assistant to enable AI features is provided as a technology preview. With this preview offering, NetApp reserves the right to modify offering details, contents, and timeline before General Availability.

The assistant's actions are limited by your assigned role-based access. [Learn about access roles](#).

See what you can do with the Console assistant

The Console assistant is a natural language interface to NetApp Console local deployment. Ask a question or describe a task, and it returns an answer or a proposed action. The assistant provides these capabilities:

- **Provision storage** Describe a workload in plain language. The assistant recommends a storage class and placement, then creates the volume or LUN after you confirm details. For example, ask it to create a CIFS volume with a specific capacity. The assistant identifies the cluster and storage virtual machine (SVM), fills in required parameters, and provisions storage.
- **Search and get guidance** Ask questions about your storage environment, find Console local deployment features, and get contextual answers from NetApp documentation and current fleet state.
- **Investigate alerts** Ask the assistant to look into an active alert. It analyzes the issue and suggests a fix or, with your confirmation, performs a remediation action.

The assistant sends requests only to the LLM endpoint configured by your administrator in Console local deployment. It requests confirmation before it performs storage-changing actions, and all actions respect role-based access controls assigned to your account.

Confirm that you can use the Console assistant

- An Organization admin or Super admin must enable the Console assistant by connecting a large language model (LLM) to Console local deployment. If you do not see the **Console assistant** button on the dashboard, ask your administrator to enable it. For more information, see [Connect your LLM and enable the NetApp Console assistant](#).
- Make sure you have the role-based access controls required for the actions you want the assistant to perform.

Ask the Console assistant a question or request an action

Open the Console assistant, choose an access mode, and enter a prompt that describes your request.

Steps

1. From the NetApp Console local deployment dashboard, select the **Console assistant** button.
2. Select an access mode:
 - Select **Read only** to ask questions and get guidance without making changes.
 - Select **Read/write** to allow the assistant to act on your storage infrastructure.
3. Enter a prompt that describes your question or task, and then select **Send**.

For example: Create a 10GiB volume for home directories named home, and share it with Everyone set to FULL_CONTROL.

4. Review the assistant's response:
 - For a question, the assistant returns an answer that you can act on.

- For an action, the assistant shows the suggested action, asks for missing details such as permission level, and then asks for confirmation before it proceeds.
- 5. Confirm the action. For asynchronous operations, such as volume creation, the assistant creates a job to complete the request.
- 6. To check on a request, ask the assistant. For example, enter `Let me know when the job completes`, and the assistant returns the current status.

Manage storage in NetApp Console

Learn about fleet management in NetApp Console local deployment

Fleet management in NetApp Console local deployment is the practice of organizing storage systems into logical groups so you can apply consistent policies, control access, and monitor health across related clusters. Instead of managing each cluster independently, fleet management lets you treat your fleet as a common pool of resources to support your data storage goals.

As your storage environment grows, managing individual clusters becomes impractical. Fleet management solves this by giving you a structured way to group related systems, delegate responsibilities, and ensure every cluster operates under the same standards.

Why fleet management matters

Fleet management addresses three core challenges:

- **Simplification through abstraction** - Fleet management abstracts away the complexity of managing individual storage infrastructure. Instead of dealing with cluster-by-cluster configuration, you manage your storage estate as a unified whole, reducing operational overhead and decision fatigue.
- **Consistency and scalability** - Without clear organization, different teams make different decisions for similar workloads, leading to fragmented configurations. Fleet management lets you apply standards across dozens of systems at once, ensuring new workloads start from the same baseline across teams and fleets.
- **Governance and control** - As teams grow, you need clear boundaries for who can manage what. Fleet management provides those boundaries through organizational structure and access control, ensuring that storage decisions remain governed and auditable.

How fleets organize your resources

Your organization's resource hierarchy consists of folders and fleets:

For a detailed overview of the hierarchy and access model, see [Learn about folders and fleets in NetApp Console local deployment](#).

- **Organization** - The top level representing your company.
- **Folders** - Help you organize related fleets. For example, you might create a folder for each region or business unit, then create fleets within each folder. Folders can contain other folders or fleets. When you assign a role at the folder level, members inherit that role for all fleets within the folder.
- **Fleets** - Group the storage systems you manage together. You associate storage systems with fleets and assign members to fleets to grant them access.



Folders are an organizational tool and are not visible to members who do not have IAM permissions such as Organization admin, Folder admin, or Fleet admin roles. Members access fleets, not folders.

Common fleet organization patterns

How you organize fleets depends on your operational model:

- **By environment** - Create separate fleets for production, development, and test workloads. This keeps production storage under tighter policies while allowing more flexibility in lower environments.
- **By business unit** - Group clusters that serve a specific department, such as finance, engineering, or HR, into a dedicated fleet. You can then assign storage management responsibilities to team members within that business unit without exposing other fleets to them.
- **By location or data center** - When clusters are distributed across sites, organizing by location lets you monitor site-level health, apply region-specific policies, and track capacity consumption per site.
- **By application tier** - Group clusters that host a specific class of workload, such as databases, file shares, or virtual machines, so you can apply consistent standards tuned to that workload type across the entire fleet.



Use folders to add another level of organization. For example, create a folder for each region and then create environment-based fleets within each regional folder.

How fleet management enables access control

Fleets and folders serve as boundaries for user access and administrative responsibility:

- **Folder-level access** - When you assign a role at the folder level, the member inherits that role for all fleets and resources within the folder. This lets you delegate administrative responsibilities without granting access to the entire organization.
- **Fleet-level access** - When you assign a role at the fleet level, the member has access only to the storage systems within that fleet. This lets you delegate storage management to team members or application owners without exposing unrelated parts of your infrastructure.

Console local deployment provides fleet-level health and performance monitoring so you can see the status of all clusters in a fleet from a single view, identify issues early, and act before they affect workloads.

How storage management works

Storage management is the practice of defining storage standards, applying them consistently, and operating workloads so they stay aligned over time. In Console local deployment, those standards are expressed as storage class policies and storage classes, scoped to fleets, and enforced through provisioning workflows governed by RBAC and audit logging.

Console local deployment connects four concepts into one workflow:

- **Fleets** define the scope where you manage storage. A fleet groups systems so you can control access, apply standards consistently, and manage resources as a unit.
- **Storage class policies** define standards as reusable building blocks (performance, capacity, security, data protection).
- **Storage classes** express workload intent. A storage class is a named template assembled from one or more policies.

- **Provisioning workflows** create storage within guardrails. Different workflows make configuration decisions differently, but all can enforce storage classes and remain governed by RBAC and audit logging.

Provisioning approaches

Console local deployment supports three provisioning approaches, all governed by RBAC, audit logging, and storage class standards:

- **Manual provisioning** - You select the target system and specify configuration details directly. Use this when you need explicit control over system selection and configuration.
- **Automated provisioning** - You provide workload inputs and optionally select a storage class. Console local deployment recommends best-fit placement and applies class-driven settings to reduce manual decisions.
- **AI-assisted (agentic) provisioning** - You describe what you need in natural language. Console local deployment proposes a plan constrained by RBAC and storage classes, then provisions only after you review and approve.

Where to go next

- To understand storage standards, see [Learn about storage classes and policies in NetApp Console local deployment](#).
- To create and manage fleets, see [Manage folders and fleets](#).
- [Provision storage manually](#)
- [Provision storage automatically](#)
- [Provision storage with AI assistance](#)

Standardize storage behavior

Learn about storage classes and policies in NetApp Console local deployment

A storage class is a reusable template that defines how storage should behave. When you provision storage using a storage class, NetApp Console local deployment enforces the standards encoded in that class automatically without requiring administrators to make individual configuration decisions for every workload.

Storage classes are built from storage class policies, which define individual dimensions of storage behavior. Together, they let you capture your organization's storage standards once and apply them consistently across fleets.

What storage class policies are

A storage class policy defines one dimension of storage behavior. Policies are reusable building blocks that you combine to create storage classes.

Common policy types include:

- **Performance policies** - Define latency targets, IOPS, or throughput requirements.
- **Capacity policies** - Define provisioning mode, threshold alerts, or growth limits.
- **Security policies** - Define encryption, compliance, or access control requirements.

- **Data protection policies** - Define snapshot schedules, replication targets, or retention periods.

You define policies independently, then combine them when you build a storage class. This lets you reuse the same performance policy across multiple classes, or update a capacity policy in one place and apply that change to every class that uses it.

What storage classes are

A storage class is a named template assembled from one or more policies. It captures your organization's storage standards for a specific type of workload.

For example, you might create a **Production Database** storage class that combines high-performance, high-availability, and strict data protection policies, or a **Development File Share** storage class that combines moderate performance, flexible capacity, and basic data protection policies.

Storage administrators define storage classes to encode enterprise requirements. All provisioning activity, whether done manually, through guided workflows, or with automation, draws from the library of classes those administrators have approved.

How storage classes enforce standards

When you provision storage, you select a storage class rather than configuring individual settings. Console local deployment uses the policies in that class to identify which clusters in your fleet have the capacity and performance headroom to support the workload, recommend the best placement option, and apply class-driven settings automatically at provisioning time.

After provisioning, Console local deployment continuously evaluates each workload against its storage class standards.

Storage class drift and compliance

When a workload no longer matches its storage class intent, Console local deployment flags it for investigation and remediation.

Issues fall into two categories:

- **Configuration drift:** Storage settings governed by the storage class policy no longer match the intended configuration. This can occur when changes are made directly on the ONTAP cluster or outside standard provisioning workflows.
- **Performance non-compliance:** The workload's runtime behavior no longer meets the storage class performance intent (for example, sustained pressure near a QoS limit or elevated tail latency).

An analysis view explains what changed and why. Guided remediation actions (for example, **Fix it**) may be available to help restore compliance. Some remediations can be applied in place, while others may require aligning the workload to a different storage class to restore the intended behavior.

Where to investigate

You can typically investigate drift and non-compliance from one of these locations (availability depends on your deployment and permissions):

- **Storage classes:** Drift / Compliance
- **Alerts:** Analyze

For step-by-step instructions, see [Remediate storage class drift](#).

End-to-end workflow

The following steps describe the process for using storage classes to standardize and govern storage in your environment:

1. **Create storage class policies** - Policies define the individual dimensions of storage behavior (performance targets, capacity settings, security requirements, data protection schedules). Create policies before building a storage class.
2. **Create a storage class** - Assemble a storage class by combining one policy from each applicable category. You can use a predefined storage class or create a custom one for your organization's standards.
3. **Associate the storage class with a fleet** - After you create a storage class, associate it with the fleet where it will be used for provisioning and compliance monitoring.
4. **Provision storage using the storage class** - When provisioning a volume or LUN, select a storage class instead of configuring individual settings. Console local deployment uses the class to recommend the best-fit cluster placement, then provisions with the settings defined in the class.
5. **Monitor workloads for drift** - Console local deployment continuously evaluates each workload against the standards encoded in its storage class. If configuration drift or performance non-compliance occurs, it flags the issue and may raise an alert.
6. **Remediate drift to restore compliance** - When drift or performance non-compliance is detected, you can follow guided steps to correct the issue manually, let Console local deployment resolve common drift conditions automatically, or disassociate a workload from its storage class if you need to change its settings.

How storage classes work with fleets

Storage classes are associated with fleets. When you associate a storage class with a fleet, that class becomes available for all provisioning within the fleet. This ensures that every workload provisioned in the fleet follows the same standards, making fleets the primary way to enforce consistent storage behavior across related clusters.

For details on how to organize fleets, see [Learn about fleet management in NetApp Console local deployment](#).

Where to go next

- To understand fleet organization, see [Learn about fleet management in NetApp Console local deployment](#).
- To create policies and storage classes, see [Manage storage classes and policies](#).
- [Provision storage manually](#)
- [Provision storage automatically](#)
- [Provision storage with AI assistance](#)
- To analyze and remediate drift, see [Remediate storage class drift](#)

Understanding storage class policies in NetApp Console local deployment

Storage class policies are the building blocks you use to construct a storage class. Each policy controls a specific aspect of storage behavior. When you combine policies into a storage class, you create a reusable template that the NetApp Console local deployment enforces automatically at provisioning time and monitors continuously throughout a workload's lifecycle.

Policies as building blocks

A storage class is composed of one or more policies, each governing a different dimension of storage behavior. Storage administrators define policies to encode enterprise standards—performance expectations, capacity thresholds, data placement rules—and then assemble those policies into storage class templates. When a workload is provisioned using a storage class, it inherits all of the policies in that class. This design separates the responsibility of defining standards from the act of provisioning, so storage can be provisioned consistently by automated workflows without needing to make individual configuration decisions.

Types of storage class policies

NetApp Console local deployment includes four types of policies that you can use to construct storage classes:

Performance policy

A performance policy sets expected IOPS/TB, peak IOPS/TB, absolute minimum IOPS, and expected latency targets for a workload. The Console local deployment uses these values to recommend clusters with sufficient headroom at provisioning time and to detect IOPS or latency drift after provisioning.

Capacity policy

A capacity policy controls how a volume consumes and manages space. It sets space reservation (thick, thin, or system default), autogrow behavior, FabricPool tiering, and whether to provision NAS workloads as FlexVol or FlexGroup volumes.

Security policy

A security policy sets the encryption, ransomware protection, and FIPS requirements for a workload. Each attribute can be set to a required value or left as "any" to accept the system default.

Data protection policy

A data protection policy sets the number of snapshots retained at each interval to ensure workloads using the policy get a consistent backup schedule.

Default policies

The Console local deployment includes system-defined policies across all four policy types. You can use these policies as-is when you create a storage class, or copy them as a starting point for custom policies tailored to your organization's requirements.

Performance policies

Name	Type	Expected IOPS/TB	Peak IOPS/TB	Absolute min IOPS	Expected latency (ms)	Summary
Extreme for Database Data	System defined	12,288	24,576	2,000	1	Use for transactional database data volumes requiring sustained high IOPS and sub-millisecond latency.
Extreme for Database Logs	System defined	22,528	45,056	4,000	1	Use for database transaction log volumes requiring the highest IOPS floor to prevent write bottlenecks.

Name	Type	Expected IOPS/TB	Peak IOPS/TB	Absolute min IOPS	Expected latency (ms)	Summary
Extreme for Database Shared Data	System defined	16,384	32,768	2,000	1	Use for shared database volumes accessed by multiple hosts requiring high throughput and consistent latency.
Extreme Performance	System defined	6,144	12,288	1,000	1	Use for HPC, AI/ML, and analytics workloads requiring high burst IOPS and predictable low latency.
Performance	System defined	2,048	4,096	500	2	Use for virtualized and enterprise applications requiring reliable performance without extreme IOPS demands.
Value	System defined	128	512	75	17	Use for cost-sensitive workloads such as home directories, file shares, and archives.

Capacity policies

Name	Type	Space reserve	Autogrow mode	Tiering policy	Volume type (NAS)	Summary
default	System defined	System default	System default	none	System default	Use for general workloads where platform default capacity behavior is acceptable.
production	System defined	System default	grow	none	System default	Use for production workloads that must expand automatically to prevent out-of-space failures.

Security policies

Name	Type	Encryption	Ransomware protection	FIPS	Summary
default	System defined	any	any	any	Use for workloads where platform default security behavior is acceptable.
anti-ransomware	System defined	any	on	any	Use for workloads containing sensitive or business-critical data that require active ransomware protection.
production	System defined	enabled	any	any	Use for production workloads where encryption is required for compliance.

Data protection policies

Name	Type	Hourly snapshots	Daily snapshots	Weekly snapshots	Monthly snapshots	Summary
default	System defined	6	2	2	0	Use for standard workloads requiring short-term recovery points without long-term retention overhead.
3month-hourly	System defined	23	6	4	3	Use for regulated or business-critical workloads requiring granular intraday recovery points and three months of historical retention.

Review predefined storage class policies in NetApp Console local deployment

NetApp Console local deployment includes system-defined policies across all four policy types. Use this reference to identify which policy best fits your workload requirements when creating or customizing a storage class.

Performance policies

Extreme Performance

Use for HPC, AI/ML, and analytics workloads requiring high burst IOPS and predictable low latency.

Extreme for Database Data

Use for transactional database data volumes requiring sustained high IOPS and sub-millisecond latency.

Extreme for Database Logs

Use for database transaction log volumes requiring the highest IOPS floor to prevent write bottlenecks.

Extreme for Database Shared Data

Use for shared database volumes accessed by multiple hosts requiring high throughput and consistent latency.

Performance

Use for virtualized and enterprise applications requiring reliable performance without extreme IOPS demands.

Value

Use for cost-sensitive workloads such as home directories, file shares, and archives.

Capacity policies

Autogrow capacity

Use for production workloads that must expand automatically to prevent out-of-space failures.

Security policies

Encryption at-rest

Use for production workloads where encryption is required for compliance.

Ransomware protection

Use for workloads containing sensitive or business-critical data that require active ransomware protection.

Standard security

Use for workloads where platform default security behavior is acceptable.

Data protection policies

default

Use for standard workloads requiring short-term recovery points without long-term retention overhead.

3month-hourly

Use for regulated or business-critical workloads requiring granular intra-day recovery points and three months of historical retention.

30-day recovery

Use for standard workloads requiring short-term recovery points without long-term retention overhead.

90-day recovery

Use for regulated or business-critical workloads requiring granular intra-day recovery points and three months of historical retention.

Create storage classes in NetApp Console local deployment

Create a storage class in the NetApp Console local deployment to standardize how storage is provisioned and governed across your fleets. A storage class is a reusable template that bundles storage class policies—such as performance targets, capacity behavior, security requirements, and data protection schedules—into a single named definition.

When users (or automation) provision a volume or LUN using a storage class, NetApp Console local deployment applies the class's policies automatically. After provisioning, Console local deployment continuously monitors workloads against the storage class to detect drift and can guide or automate remediation to restore compliance.

Before you begin

- Discover at least one ONTAP cluster so Console local deployment has targets for placement recommendations.
- Ensure you have the Organization admin role (or a role that grants storage class management permissions in your environment).
- Create (or identify) the storage class policies you plan to use—performance, capacity, security, and data protection—because storage classes are assembled from policies.

Create a storage class

You must have the Organization admin or Folder or fleet admin role to add a storage class.

Steps

1. Select **Storage > Management**.
2. Select **Storage classes**.
3. Select **Add**.
4. Under **Basic information**, enter the following:
 - **Storage class name**: Enter a unique name for the storage class.
 - **Description**: (Optional) Enter a description of the storage class.
 - **Recommended apps**: (Optional) Enter a list of recommended apps for the storage class.
5. Under **Storage policies**, select one or more policies to associate with the storage class.
6. Select **Add**.

Customize an existing storage class

You must have the Organization admin or Folder or fleet admin role to customize an existing storage class.

Steps

1. Select **Storage > Management**.
2. Select **Storage classes**.
3. For the storage class you want to customize, select ..., and then select **Duplicate**.
4. Under **Basic information**, update the following as required:
 - **Storage class name**: Enter a unique name for the storage class.
 - **Description**: (Optional) Enter a description of the storage class.
 - **Recommended apps**: (Optional) Enter a list of recommended apps for the storage class.
5. Under **Storage policies**, select one or more policies to associate with the storage class.
6. Select **Add**.

Edit a user-defined storage class

You must have the Organization admin or Folder or fleet admin role to edit a user-defined storage class.

Steps

1. Select **Storage > Management**.
2. Select **Storage classes**.
3. For the storage class you want to edit, select ..., and then select **Edit**.
4. Under **Basic information**, edit the following as required:
 - **Storage class name**: Enter a unique name for the storage class.
 - **Description**: (Optional) Enter a description of the storage class.
 - **Recommended apps**: (Optional) Enter a list of recommended apps for the storage class.
5. Under **Storage policies**, select one or more policies to associate with the storage class.
6. Select **Edit**.

Delete a user-defined storage class

You must have the Organization admin or Folder or fleet admin role to delete a user-defined storage class.

Steps

1. Select **Storage > Management**.
2. Select **Storage classes**.
3. For the storage class you want to delete, select ..., and then select **Delete**.
4. Select **Delete**.

Note that this action cannot be undone. If you delete a storage class that is currently in use, any volumes or LUNs provisioned with that class will continue to operate but will no longer be associated with the deleted class.

Provision storage in NetApp Console local deployment

Provision volumes and LUNs in NetApp Console local deployment to make storage resources available for your workloads. During provisioning, you can optionally select a storage class to apply predefined storage policies and organizational standards.

Required access roles

Super admin, Organization admin, Folder or fleet admin, or Storage admin. [Learn about access roles](#).

Provision a volume

Steps

1. Select **Storage > Management**.
2. Select **Fleets**.
3. Select the fleet you want to provision into.
4. Select **Add**.
5. From the menu, select **Volume**.
6. Under **Volume details**:
 - a. Enter a volume name.
 - b. Enter capacity (and choose units if needed).
 - c. (Optional) Select a storage class to apply storage class policies. If you don't select one, the volume is provisioned without storage class policies.
7. Under **System details**:
 - a. Select a system from the eligible systems for the selected fleet.
 - b. Select a storage VM.
8. Under **Host (NAS) details**, choose how hosts access the volume:
 - a. Export by NFS (export policies control which NFS hosts can access the volume)
 - b. Share by SMB/CIFS
9. Select **Add**.

Provision a LUN

Steps

1. Select **Storage > Management**.
2. Select **Fleets**.
3. Select the fleet you want to provision into.
4. Select **Add**.
5. From the menu, select **LUNs**.
6. Under **LUN storage details**:
 - a. Enter a prefix name.
 - b. Enter the number of LUNs to create with that prefix.
 - c. Enter the capacity per LUN (and choose units if needed).
 - d. (Optional) Select a storage class to apply storage class policies. If you don't select one, the LUNs are provisioned without storage class policies.
7. Under **System details**:
 - a. Select a system (if prompted).
 - b. Select a storage VM.
8. Under **Host (SAN) details**, choose how hosts access the volume:
 - a. Choose the host operation system (for example, Windows or Linux) to set recommended defaults for LUN alignment and block size.
 - b. Choose the host mapping group to control which initiators can access the LUNs.
9. Select **Add**.

Monitor storage health

Storage monitoring in NetApp Console local deployment

NetApp Console local deployment helps you monitor storage across fleets with dashboards, alerts, deep-dive analysis, and remediation so you can spot and resolve issues before they affect workloads.

Monitor fleet-wide health with dashboards

Start with dashboards for a quick view of storage health and performance. Use the Home page for at-a-glance metrics, open the dashboards table for predefined and custom dashboards, and move into alerts, investigation, and remediation when you need more detail.

- [View Home page metrics in NetApp Console local deployment](#)
- [Keep dashboards current in NetApp Console local deployment](#)
- [Learn about custom dashboards in NetApp Console local deployment](#)

Investigate performance issues

Use the Workload Analyzer to investigate performance issues on individual volumes. It correlates latency,

throughput, IOPS, and configuration changes so you can pinpoint root causes.

- [Learn about the NetApp Console local deployment Workload Analyzer](#)
- [Investigate high latency on a volume in NetApp Console local deployment](#)
- [Investigate high throughput demand on a volume in NetApp Console local deployment](#)
- [Learn about Workload Analyzer metrics in NetApp Console local deployment](#)

Configure alerts and notifications

Configure alert policies and notification channels so the right people hear about storage conditions that require attention. For example, route a capacity warning to the storage team and a protection alert to the backup admin who can act on it.

[Configure notifications and alert policies in NetApp Console local deployment.](#)
[View storage alerts in NetApp Console local deployment.](#)

Remediate issues

When you detect an issue, remediate it directly from NetApp Console local deployment:

- **Automated remediation** — NetApp Console local deployment applies corrective actions automatically based on predefined rules.
- **Guided remediation** — Follow step-by-step recommendations to resolve issues with full control over each action.

Monitor with dashboards

View Home page metrics in NetApp Console local deployment

Use the Home page dashboard in NetApp Console local deployment as your default starting point for monitoring storage health and performance. It provides high-level metrics for fleet health, alerts, security, capacity usage, latency, throughput, and IOPS.

The Home page dashboard is automatically scoped to only the fleets and systems you are authorized to view. You can also add a custom dashboard to the Home page for role-specific monitoring.

Use the cards as a layered triage workflow. Start with **Fleet health** and **Alerts** to find risk hotspots. Use **System health status** and **Recommended remediations** to identify impacted resources. Use **Capacity usage**, **Latency**, **Throughput**, and **IOPS** to investigate root causes.

Fleet health

The **Fleet health** card shows a high-level status summary for the top five fleets, including system counts, used capacity, security compliance, and critical alerts. Use this card to identify which fleets need immediate attention. Select the fleet name to view a fleet-only dashboard for the selected fleet.

Recommended remediations

The **Recommended remediations** card lists active issues and suggested actions. The card prioritizes remediations based on capacity pressure, offline resources, and protection-related risks.

Alerts

The **Alerts** card summarizes alert volume in the selected time range and groups it by severity. Use this card to understand current incident load and to spot recent changes in critical, warning, or informational alerts.

Security

The **Security** card summarizes compliance and protection metrics, such as system compliance and ransomware-related controls. Use this view to monitor security trends across your resources.

Capacity usage

The **Capacity usage** card displays projected and historical usage trends across selected fleets or systems. Use this card to identify growth patterns and plan for additional capacity.

Latency

The **Latency** card tracks response-time behavior over time across the selected systems. Use this trend to detect emerging performance delays and compare relative latency across resources.

Throughput

The **Throughput** card shows data transfer rates over time for selected systems. Use this card to understand workload intensity and monitor throughput demand changes.

IOPS

The **IOPS** card shows input/output operation rates over time. Use this card to compare activity levels between systems and identify sustained or intermittent I/O demand.

Dashboard card (example metric card)

The lower **Dashboard** area can contain a user-selected custom dashboard such as average latency for a selected production scope. Use these cards for focused monitoring by team, workload, or goal.

System health status

The **System health status** card lists individual systems with their current health state. Use this card to quickly identify unhealthy systems and correlate status changes with alerts and performance signals.

Related information

- [Learn about custom dashboards in NetApp Console local deployment](#)
- [Learn about custom dashboard cards and metrics in NetApp Console local deployment](#)

Use custom dashboards

Learn about custom dashboards in NetApp Console local deployment

Use custom dashboards in NetApp Console local deployment to create focused monitoring views for specific teams, fleets, workloads, and operational goals. Custom dashboards complement the Home page dashboard by adding targeted visibility to broad, always-available monitoring.

How dashboard types work together

Home page dashboard

Ready-to-use monitoring panes for capacity, alerts, performance capacity, licenses, and data protection status. Views are preconfigured and update automatically.

[View Home page metrics in NetApp Console local deployment.](#)

Custom dashboards

Role-specific monitoring views built from predefined cards, including selectable scope, metrics, target resources, and time windows.

Use both together:

- * Start on the Home page for daily baseline monitoring.
- * Use custom dashboards for focused investigation by team, fleet, workload, or operational question.

You can add one custom dashboard to the Home page at a time.

Understand custom dashboards

A custom dashboard is a saved collection of cards that you arrange on a grid. Each card is based on a predefined card template with a metric and chart type. When you add a card, you configure the scope, target objects, aggregation, and time window.

Custom dashboards and cards are private to you. Each card shows data only for the storage resources you are authorized to view.

You can remove the custom dashboard from your Home page and add a different one as monitoring priorities change.

Monitor what matters

Card templates are organized by metric type so you can focus a dashboard on a specific operational area. For example, use capacity cards to watch a volume that is growing faster than expected, or use alert cards to track repeated failures on a busy cluster:

Performance

Latency, IOPS, throughput, utilization, and performance capacity across the fleet, systems, SVMs, volumes, and LUNs.

Capacity

Used capacity, free capacity, capacity used percentage, and snapshot capacity across the fleet, systems, SVMs, volumes, LUNs, and local tiers.

Health

Health status, degraded or critical object counts, failed disks, network interface errors, and hot objects.

Alerts

Critical alert counts, alerts by severity, alerts by impact area, recent alerts, and alert trends over time.

For a complete catalog of predefined cards and metrics, see [Learn about custom dashboard cards and metrics in NetApp Console local deployment.](#)

Resource types

You can scope cards to any level of the ONTAP storage hierarchy: fleet, cluster, system (node), SVM, volume, LUN, and local tier (aggregate). The available resource types depend on the card template you select.

Plan dashboard views

Organize custom dashboards around storage administration goals such as workload performance triage, capacity planning, health risk detection, and alert prioritization.

Related information

- [Get started with dashboards in NetApp Console local deployment](#)
- [View Home page metrics in NetApp Console local deployment](#)
- [Learn about custom dashboard cards and metrics in NetApp Console local deployment](#)

Get started with dashboards in NetApp Console local deployment

In NetApp Console local deployment, follow this workflow to monitor from broad to focused views: review the Home page dashboard, review predefined dashboards, create custom dashboards, and keep dashboards current.

Required access roles

All roles. Dashboards show only the resources you are authorized to view. If you do not see a resource in the list of available resources, you do not have permission to view it.

1

Check baseline health on the Home page

Use the Home page dashboard to review organization-level capacity, alerts, performance, and data protection status.

- [View Home page metrics in NetApp Console local deployment](#)
- [Learn about storage monitoring in NetApp Console local deployment](#)

2

Review other predefined dashboards

Open **Storage > Dashboards** to review predefined dashboards for additional role-specific views.

- [Manage dashboards in NetApp Console local deployment](#)

3

Build custom dashboards for focused monitoring

Create custom dashboards when you need targeted views for selected resources, metrics, and time windows.

- [Create a custom dashboard](#)
- [Customize dashboard cards in NetApp Console local deployment](#)

4

Keep dashboards current

Update dashboards as priorities change by editing, duplicating, or deleting custom dashboards.

- [Manage dashboards in NetApp Console local deployment](#)
- [Learn about custom dashboard cards and metrics in NetApp Console local deployment](#)

Related information

- [Learn about custom dashboards in NetApp Console local deployment](#)
- [View Home page metrics in NetApp Console local deployment](#)

Build dashboards for monitoring goals in NetApp Console local deployment

Create a custom dashboard in NetApp Console local deployment, add cards, and save a monitoring view for daily operations.

If you only need organization-level monitoring that is ready to use, start with predefined Home page dashboards. Create a custom dashboard when you need a role-specific view, resource-level targeting, or a tailored card layout.

Console local deployment dashboard guidelines

- Only the user who creates the dashboard can view it. You cannot share dashboards with other users, even when you add them to the Console Home page.
- You can only view resources that you have permission to view. If you do not see a resource in the list of available resources, you do not have permission to view it.
- Before you begin, identify the metric types and resources you want the dashboard to track. [Learn about custom dashboard cards and metrics in NetApp Console local deployment](#)

Build a dashboard for your monitoring goals

Create a dashboard when you need one place to monitor the metrics that matter to your role, such as fleet health, performance hotspots, and alert trends.

Steps

1. Select **Storage > Dashboards**.
2. Select **Add Dashboard**.

Console local deployment creates a new dashboard with a default name and no description.

3. Select **Add Card**.
4. Select the **Resource type** for the card, such as fleet, system, or volume and then choose resources from the list of available resources.
5. Select **Next** to continue to choose a metric and card type.
6. Choose a metric to display. You can filter the available metrics by type: the **Metric Type: Performance, Capacity, Health, or Alerts** or search for a specific metric by name. The available metrics depend on the resource type you selected.

[Learn about the available metrics.](#)

7. Select a card to include in the dashboard and select **Next**.

You can only select one card at a time. The card preview shows live data for the selected metric and

resource type.

8. Name your card. The name must be unique within the dashboard.
9. Review the card preview, enter a card name and description, and then select **Add**.
10. Repeat these steps for additional cards.
11. Select the pencil icon to edit the dashboard name and description to describe its purpose.
12. Select the action menu to the right of the Add card button and select **Save dashboard**.

The saved dashboard is available under **Storage > Dashboards**.

13. Optionally, add this dashboard to your Home page by selecting **Add to Home page** from the action menu. The dashboard is added to the Home page. Only you can view custom dashboards that you create. It is not shared with others.

Related information

- [View Home page metrics in NetApp Console local deployment](#)
- [Learn about custom dashboards in NetApp Console local deployment](#)
- [Customize dashboard cards in NetApp Console local deployment](#)
- [Manage dashboards in NetApp Console local deployment](#)

Customize cards in NetApp Console local deployment

In NetApp Console local deployment, customize dashboard cards when you need to focus operators on the signals that matter most, such as SLA risk, capacity pressure, or recurring alert noise. Use this topic to shape dashboard views around the operational decisions your team needs to make.

Before you begin

- Create or open a dashboard where you want to place cards.
- Review available templates and metrics in [Learn about custom dashboard cards and metrics in NetApp Console local deployment](#).

Add cards while refining an existing dashboard

Use this option when you are refining an existing dashboard and need to quickly add cards that answer a specific operational question.

Steps

1. Open the dashboard you want to add a card to.
2. Select **Add Card**.
3. Select the **Resource type** for the card, such as fleet, system, or volume and then choose resources from the list of available resources.
4. Select **Next** to continue to choose a metric and card type.
5. Choose a metric to display. You can filter the available metrics by type: the **Metric Type: Performance, Capacity, Health, or Alerts** or search for a specific metric by name. The available metrics depend on the resource type you selected.

[Learn about the available metrics.](#)

6. Select a card to include in the dashboard and select **Next**.

You can only select one card at a time. The card preview shows live data for the selected metric and resource type.

7. Name your card. The name must be unique within the dashboard.
8. Review the card preview, enter a card name and description, and then select **Add**.
9. Repeat these steps for additional cards.
10. Save the dashboard.

Related information

- [View Home page metrics in NetApp Console local deployment](#)
- [Learn about custom dashboards in NetApp Console local deployment](#)
- [Create a custom dashboard](#)
- [Manage custom dashboards](#)
- [Learn about custom dashboard cards and metrics in NetApp Console local deployment](#)

Predefined dashboards and custom dashboard cards reference in NetApp Console local deployment

In NetApp Console local deployment, predefined dashboards and custom dashboard card templates provide monitoring coverage for ONTAP performance, capacity, health, and alert operations.

Predefined dashboards

The following predefined dashboards are available out of the box.

Name	Description
Volumes	Volume performance, capacity, QoS, FabricPool, growth rate, and forecast metrics.
Time Till Full	ONTAP time-till-full predictions for clusters, volumes, and aggregates.
Security	Security compliance, encryption, autonomous ransomware protection, and authentication overview.
SVMs	ONTAP SVM performance, capacity, volume, LIF, protocol (CIFS, FCP, iSCSI, NFS, NVMe/FC), QoS, and latency heatmap monitoring.
Power	ONTAP cluster power consumption, temperature, and fan speed monitoring for nodes, shelves, and aggregates.
Nodes	ONTAP node performance, CPU, network, and backend monitoring.
Network	Network performance metrics including Ethernet, FibreChannel, NVMe/FC, Link Aggregation Groups, and Routes.
LUNs	ONTAP LUN performance, capacity, and efficiency monitoring.

Name	Description
Health	ONTAP cluster health monitoring including errors, warnings, EMS alerts, HA issues, node/disk/shelf health, volumes, licenses, and network ports.
Aggregates	ONTAP aggregate capacity, efficiency ratios, and growth rate monitoring.

Dashboard-level time windows and aggregation

Available time windows are 30 minutes, 1 hour, 24 hours, 48 hours, 7 days, and 30 days. Aggregation options vary by template and include views such as fleet-wide rollup or top-N style results. Time window and aggregation are configured at the dashboard level and apply to all cards in the dashboard.

Dashboard card templates and metrics

Cards are the building blocks of custom dashboards. Each card uses a predefined template with a metric and chart type, then maps that view to specific ONTAP objects and monitoring goals.



Card templates are predefined and not user-created.

Storage operations metric index (at a glance)

Metric types align to common storage administration outcomes, including performance bottlenecks, capacity pressure, health risk, and alert volume.

Metric type	Supported metrics	Storage admin use case	Detailed templates
Performance	Average latency, peak latency, IOPS, throughput (MB/s), utilization, performance capacity	Identify bottlenecks, sustained pressure, and performance headroom	Go to templates
Capacity	Used capacity, free capacity, capacity used (%), snapshot capacity used	Track growth, identify low-free-capacity areas, and monitor snapshot impact	Go to templates
Health	Health status, degraded/critical objects, failed disks, network interface errors, hot objects (by latency)	Detect health regressions and prioritize operational triage	Go to templates
Alerts	Alerts (Critical), alerts by severity, alerts by impact area, recent alerts, alert trend	Monitor active incidents and trend alert volume over time	Go to templates

Supported ONTAP object hierarchy (resource types)

Card data can be represented at multiple ONTAP object levels, from fleet-level visibility to object-level troubleshooting.

Resource type settings map to one or more of the following resource levels:

- Fleet
- Cluster

- System (node)
- SVM
- Volume
- LUN
- Local tier (aggregate)

The available resource levels depend on the template and metric you select.

Chart types for operational analysis

Chart type affects how quickly you can interpret trends, comparisons, distributions, and point-in-time values.

Chart type	Best for
Line chart	Time-series trends, such as latency, IOPS, throughput, utilization, and alert trends over time.
Bar chart	Categorical comparison and distribution, such as alerts by severity or impact area.
Pie or doughnut	Proportional distribution, such as capacity used percentage.
Table	Detailed, multi-column data, such as used capacity, health status, and recent alerts.
Single number (stat)	A single key value at a glance, such as critical alert counts or failed disks.

Performance templates for workload triage

Performance templates focus on latency, throughput, and utilization signals that indicate workload pressure.

Template	Chart type	Description
Average Latency	Line	Average latency across the selected targets over the chosen time window.
Peak Latency	Line	Maximum latency observed across the selected targets over the chosen time window.
IOPS	Line	Sum of I/O operations per second across the selected targets.
Throughput (MB/s)	Line	Sum of data throughput across the selected targets.
Utilization	Line	Average CPU or disk utilization across the selected targets.
Performance Capacity	Table	Headroom analysis comparing current utilization to the optimal point.

Capacity templates for growth and headroom

Capacity templates focus on consumed and available space to support growth planning and risk detection.

Template	Chart type	Description
Used Capacity	Table	Sum of used capacity across the selected targets.

Template	Chart type	Description
Free Capacity	Table	Sum of free or available capacity across the selected targets.
Capacity Used (%)	Pie or doughnut	Average used-to-total ratio across the selected targets.
Snapshot Capacity Used	Table	Sum of snapshot space consumed across the selected targets.

Health templates for risk detection

Health templates focus on object status and failure indicators to support operational triage.

Template	Chart type	Description
Health Status	Table	Latest health state for each selected target object.
Degraded/Critical Objects	Single number	Count of objects whose health state is not OK.
Failed Disks	Single number	Sum of failed disks across the selected systems.
Network Interface Errors	Pie or doughnut	Sum of network interface error counters across the selected targets.
Hot Objects (by Latency)	Table	Resources ranked by peak latency, in descending order.

Alerts templates for incident prioritization

Alerts templates focus on incident volume, severity, and impact area to support monitoring and prioritization.

Template	Chart type	Description
Alerts (Critical)	Single number	Count of critical-severity alerts within the time window.
Alerts by Severity	Bar	Alerts grouped by severity level.
Alerts by Impact Area	Bar	Alerts grouped by impact area, such as capacity, performance, or security.
Recent Alerts	Table	Most recent alerts, sorted by time in descending order.
Alert Trend	Line	Alert count per time bucket over the chosen time window.

Related information

- [View Home page metrics in NetApp Console local deployment](#)
- [Learn about custom dashboards in NetApp Console local deployment](#)
- [Create a custom dashboard](#)
- [Customize dashboard cards in NetApp Console local deployment](#)
- [Manage custom dashboards](#)

Manage custom dashboards

Manage dashboards in NetApp Console local deployment

Manage custom dashboards in NetApp Console local deployment to keep views aligned with operations. You can edit, duplicate, and delete dashboards as teams, fleets, and priorities change.

This task applies only to custom dashboards under **Storage > Dashboards**. Predefined Home page dashboards cannot be edited in the same way.

View dashboards

View custom and predefined dashboards to find the dashboard you want to open, edit, duplicate, or delete.



You can duplicate a predefined dashboard to create a custom version that you can edit.

Steps

1. Select **Storage > Dashboards**.
2. Review the dashboard table to find the dashboard you want to open.
3. Select the dashboard name to open it and view its metrics.
4. If needed, use the available actions to edit, duplicate, or delete custom dashboards.

Update a dashboard as priorities change

Edit a dashboard when your operational priorities change and you need to adjust which cards and metrics are visible.

Steps

1. Select **Storage > Dashboards**.
2. Open the dashboard you want to edit.
3. Modify an existing card on the dashboard by selecting Edit from the actions menu of the card.
4. Add a new card to the dashboard by selecting **Add Card**.

[Customize cards in NetApp Console local deployment.](#)

5. Select **Save changes** to save the dashboard.
 - You can also choose to save your changes as a new dashboard by selecting **Save as new dashboard** from the action menu. This option is useful when you want to keep the original dashboard for other teams or fleets while creating a new version for your current monitoring needs.
 - You can also choose to discard your changes by selecting **Discard changes** from the action menu. This option is useful when you want to revert to the last saved version of the dashboard.

Reuse a custom dashboard for another resource or fleet

Duplicate a dashboard when you want to reuse a proven layout for another team, fleet, or monitoring scenario without rebuilding it from scratch.



You can duplicate a predefined dashboard to create a custom version that you can edit.

Steps

1. Select **Storage > Dashboards**.
2. For the dashboard you want to duplicate, select **Duplicate** from the actions menu.

Console local deployment creates a copy that includes all cards from the original dashboard.

3. Provide a name and description for the duplicated dashboard.
4. Select the duplicate dashboard you just created. Change the metrics, resources, and card types to suit your new monitoring scenario.

[Customize cards in NetApp Console local deployment.](#)

Remove dashboards you no longer use

Delete a dashboard when it no longer supports current operations and you want to keep your dashboard list focused on active use cases.

Steps

1. Select **Storage > Dashboards**.
2. For the dashboard you want to remove, select **Delete**.
3. Confirm the deletion.

Related information

- [View Home page metrics in NetApp Console local deployment](#)
- [Get started with dashboards in NetApp Console local deployment](#)
- [Learn about custom dashboards in NetApp Console local deployment](#)
- [Create a custom dashboard](#)
- [Customize dashboard cards in NetApp Console local deployment](#)

Monitor fleets using inventory in NetApp Console local deployment

In NetApp Console local deployment, use inventory to monitor fleet health and investigate storage resources across your environment. Inventory provides capacity, security, and performance views that help you identify issues, understand resource utilization, and track managed storage systems.

Inventory is primarily an informational and diagnostic workspace. From inventory, you can review systems, volumes, and other storage objects across your fleets and perform common actions such as provisioning storage resources. For advanced storage management operations, NetApp Console local deployment redirects you to System Manager.

Access Inventory

You can access **Inventory** from several areas of Console local deployment, including:

- The Home page
- Fleet overview pages
- Fleet health cards and related inventory views

Depending on where you enter **Inventory**, the displayed scope can be at the organization level or fleet level.

Inventory views

Inventory provides multiple views that help you analyze different aspects of your storage environment.

Capacity

View capacity utilization and identify systems, volumes, and other storage objects that consume storage resources.

Security

Review security-related information and compliance status across managed storage resources.

Performance

Analyze performance-related metrics and identify resources that might require further investigation.

The information displayed varies depending on the selected view and object type.

Investigate storage resources

Use **Inventory** to:

- Monitor fleet health
- Review storage capacity usage
- Track managed storage systems
- Identify volumes and other objects that consume capacity
- Investigate potential security or performance issues
- Locate resources that require administrative attention

Inventory helps you move from high-level visibility into your environment to more detailed investigation of specific storage resources.

Provision storage resources

Inventory includes workflows that allow you to provision storage resources such as volumes and LUNs.

When provisioning resources, you select an existing managed storage system and provide the configuration settings required for the workload.

Inventory and alerts

Alerts are generated for specific storage objects and conditions within your environment.

When you select an alert, Console local deployment might direct you to System Manager for additional investigation or management actions. **Inventory** complements alerts by providing broader visibility into the overall state of your storage environment and managed resources.

Advanced management tasks

Inventory helps you identify resources that require action, but some advanced storage management operations are performed in System Manager.

Examples include:

- Advanced workload management
- Resource relocation and optimization tasks
- Detailed system administration activities

Use **Inventory** to identify and investigate issues, and then use System Manager when deeper management capabilities are required.

Remediate alerts

Learn about alerts in NetApp Console local deployment

NetApp Console local deployment generates alerts that identify health issues on your on-premises ONTAP clusters and for NetApp Backup and Recovery operations.

Console local deployment consolidates alerts from ONTAP clusters and Backup and Recovery operations into a single view. The Alerts page includes a dashboard, an alerts list, and a systems view so you can review alerts across the entire organization or scope them to a specific fleet or system. Each alert identifies the affected system, its severity, and its impact area.

Use alerts in Console local deployment to:

- Detect capacity, connectivity, data protection, performance, and security issues.
- Prioritize alerts by severity and impact area.
- Open an alert in System Manager or Workload Analyzer, then run a guided or automated Fix-It action when the page offers one.
- Route notifications through email to users with specified access roles, or send alert data to an external system through a webhook.
- Export the alerts list to a CSV file for offline analysis.

Alert severity and impact areas

Each alert includes a severity and an impact area:

- **Severity** indicates urgency, from highest to lowest: Critical, Major, Minor, Warning, and Info.
- **Impact area** identifies the affected domain: Capacity, Connectivity, Data protection, Performance, and Security.

Alert sources and scope

- **ONTAP alerts** originate from the ONTAP Event Management System (EMS) on the on-premises clusters that Console local deployment has discovered. [Learn more about the ONTAP EMS and available alerts.](#)
- **NetApp Backup and Recovery alerts** originate from Backup and Recovery actions. [Learn more about NetApp Backup and Recovery alerts.](#)
- Your permissions determine which fleets you can view. Scope the view to the entire organization, a specific fleet, or an individual system. The **Systems health** tab shows per-system status and the **Alerts** tab shows the current alert list for the selected scope.
- The CSV export reflects the current scope, search text, and filters.

Alert notification rules

Create notification rules to determine which alerts generate notifications and who receives them. A notification rule has the following characteristics:

- It matches alerts on the service (such as ONTAP management or Backup and Recovery), severity, impact area, and target system.
- For email delivery, it sends notifications to users with the specified access roles. For webhook delivery, it sends alert data to the configured endpoint—access to that data is controlled by the receiving application, not by Console local deployment.
- It applies to specific systems, not to fleets.
- It can be muted during a planned maintenance window to suppress expected alerts.

Console local deployment includes a default notification rule that is active immediately after installation. The default rule monitors all services and systems for Critical-severity alerts and delivers notifications to the Console Notification Center only—no email or webhook delivery is configured until you set it up. You can view and adjust this rule, and create custom rules to match your environment.

Info-level alerts are visible in the Alerts page but are not delivered by notification unless you explicitly create a rule that includes the Info severity level.

Related information

- [Monitor and investigate alerts](#)
- [Create and manage alert notification rules](#)
- [Learn about ONTAP alerts in NetApp Console local deployment](#)

Monitor and investigate alerts in NetApp Console local deployment

Monitor and investigate alerts in NetApp Console local deployment to keep your storage healthy and minimize disruption.

View active alerts across your entire organization or a specific fleet, prioritize them by severity and impact area, and investigate root causes so you can resolve issues before they escalate. When an alert includes a recommended action or Fix It option, you can move from investigation directly into remediation.

Required access role

All roles except Federation admin and Federation viewer. Users with the Federation admin or Federation viewer role cannot view alerts. [Learn about access roles](#).

For ONTAP alerts, you can access tools such as System Manager and Workload Analyzer directly from the Alerts page to investigate and resolve issues.

For external reporting, you can export the alerts list to a CSV file for offline analysis.



You can also set up notification rules to receive alerts by email or webhook. [Learn how to set up alert notifications](#).

Available alerts

NetApp Console local deployment supports alerts for ONTAP and NetApp Backup and Recovery.

- [Learn about ONTAP alerts in NetApp Console local deployment](#)
- [Learn more about NetApp Backup and Recovery alerts.](#)

View the alerts dashboard

Use the alerts dashboard to get a quick view of current alert activity for the scope you selected. The dashboard summarizes active alerts by severity and impact area, and it includes a chart that shows alert distribution over the past 24 hours so you can spot trends quickly.

You can filter the dashboard to focus on critical alerts or alerts for a specific impact area.

Steps

1. Select **Health > Alerts > Overview**.
2. From the scope selector, choose one of the following:
 - **All** (default) to see alerts across all fleets you have access to
 - A specific fleet to see alerts only for that fleet
3. Filter the dashboard according to the alerts you need to view, including:
 - Severity (Critical, Warning, or Informational)
 - Critical alerts grouped by impact area
 - Impact area (Security, Protection, Availability, Performance, Capacity, or Configuration)

View all alerts

Use the Alerts tab to view the full list of active alerts for the scope you've selected. The list updates to reflect the current organization or fleet scope, and you can search the list for specific alerts by name or description.

Steps

1. Select **Health > Alerts > Overview**.
2. From the scope selector, choose one of the following:
 - **All** (default) to see alerts across all fleets you have access to
 - A specific fleet to see alerts only for that fleet
3. Select the **Alerts** tab to view the full list of active alerts for the selected scope.
4. Optionally, search the list for a specific alert by name or description.

Alerts Systems Health									
Alert (1) Filters applied (1)									
Active x ↓									
Name	Service	Source	Status (1)	Severity	Triggered time	Impact area			
Can't reach Active Directory servers	ONTAP Management	vs0 Storage VM	Active	Critical	June 24, 2026, 8:44 AM	Availability			

View alerts by system

You can view alerts for a specific system within a fleet or your organization.

Steps

1. Select **Health > Alerts > Overview**.

- From the scope selector, choose one of the following:
 - All** (default) to see alerts across all fleets you have access to
 - A specific fleet to see alerts only for that fleet
- Select the **Systems health** tab to view a summary of the alerts associated with the systems within the scope you've selected.
- Select **View alerts** on the row of the respective system to view the full list of active alerts associated with that system.

Investigate an alert

You can investigate an alert to see what triggered it and who received the notification.

When investigating an ONTAP alert, you can also go directly to the System Manager interface for the system that generated the alert to view additional details and take action.

Steps

- Select **Health > Alerts > Overview**.
- From the scope selector, choose one of the following:
 - All** (default) to see alerts across all fleets you have access to
 - A specific fleet to see alerts only for that fleet
- In either tab, select the name of the alert you want to investigate to view its details.

Alerts (3) Filters applied (1)								
Name	Service	Source	Status (1)	Severity	Triggered time	Impact area		
Can't reach Active Directory servers	ONTAP Management	vs0 Storage VM	Active	Critical	June 24, 2026, 12:02 PM	Availability		

- If applicable, select the VM or cluster name to open the System Manager interface for the system that generated the alert.

Can't reach Active Directory servers

Storage VM: **vs0**
Cluster: [C1_sti245-vsimsim-ocvs035e_1781026189](#)

Critical
Severity

Active
Status

Availability
Impact area

12:02 PM Jun 24, 2026
Triggered time

Alert details

Description
All configured AD/LDAP servers for this storage VM are unreachable. Directory sign-in fails until at least one server responds again.

Details
Can't reach any AD/LDAP servers for storage VM vs0; authentication is unavailable.

Related alerts
1 alert

Notifications
0 notification channel

Remediate an alert

When an alert includes a recommended action or Fix It option, use it to move from investigation into remediation without leaving the alert details.

Steps

1. Select **Health > Alerts > Overview**.
2. From the scope selector, choose one of the following:
 - **All** (default) to see alerts across all fleets you have access to
 - A specific fleet to see alerts only for that fleet
3. Select the alert you want to remediate.
4. Review the alert details, then select the recommended action or **Fix It** option if one is available.
5. Follow the guided steps to apply the remediation and then return to the alert details to confirm the alert state.

If the action resolves the issue, the alert no longer appears as active for that scope. If the alert remains active, review the alert details again and continue the investigation.


Analyze an alert

You can analyze an ONTAP alert in Workload Analyzer to understand what triggered it.

When investigating an ONTAP alert, you can also go directly to the System Manager interface for the system that generated the alert to view additional details and take action.

Steps

1. Select **Health > Alerts > Overview**.
2. From the scope selector, choose one of the following:
 - **All** (default) to see alerts across all fleets you have access to
 - A specific fleet to see alerts only for that fleet
3. Select the **Systems health** tab to view the full list of active alerts for the selected scope.
4. In either tab, select the name of the alert you want to investigate to view its details.



Alerts (3) Filters applied (1)							🔍	⬇
Name	Service	Source	Status (1)	Severity	Triggered time	Impact area		
Can't reach Active Directory servers	ONTAP Management	vs0 Storage VM	Active	Critical	June 24, 2026, 12:02 PM	Availability		

5. If applicable, select **Analyze** to open the Workload Analyzer interface for the system that generated the alert.

Volume anti-ransomware monitoring state changed

Analyze

Volume: [TestVol_1782309902536](#) | Cluster: [C1_sti245-vsims-ocvs034c_1782304943](#)

Informational
Severity

Active
Status

Security
Impact area

8:07 AM Jun 24, 2026
Triggered time

Alert details

Description

The anti-ransomware state of a volume changed. Review recent administrative actions and security events to confirm the change is expected.

Details

Volume anti-ransomware state changed.

Related alerts

0 alert

Notifications

0 notification channel

- Enter the time range that covers the period when the alert was triggered, then select **Analyze** to view the analysis results. [Learn about Workload Analyzer.](#)

Export alerts to CSV

Export the alerts list in the NetApp Console local deployment to a CSV file for offline analysis or reporting. The export reflects the current scope (organization or fleet), search text, and filters.

Steps

- Select **Health > Alerts** and then select the **Alerts** tab.
- Set the scope (organization or fleet) and apply any filters or search text you want included in the export.
- Select the download icon to export the alerts list to a CSV file.

Configure notification rules for alerts in NetApp Console local deployment

Configure notification rules in NetApp Console local deployment to control which alerts raise notifications, who receives them, and how they are delivered.

Required access roles

Super admin, Organization admin, Storage admin, Operations support analyst, or any of the admin roles for NetApp Backup and Recovery. [Learn about access roles.](#)

Use notification rules to route the right alerts to the right people through the channels that fit your environment, while reducing noise from alerts that aren't relevant to you. Notification rules complement the Alerts page: you investigate or remediate the alert in the Alerts workflow, then use rules to control how similar alerts are delivered in the future.

A notification rule matches alerts based on conditions you define—such as service, severity, and impact area—and then delivers a notification through the channels you select. The Console local deployment includes default notification rules that you can view and adjust, and you can create your own custom rules. You can also mute rules to temporarily suppress notifications during planned events like maintenance windows.

- [Learn more about the ONTAP EMS and available alerts.](#)

Before you begin

- To deliver notifications by email, your Organization admin must configure an email server in the Console local deployment. To deliver notifications to an external system, your Organization admin must configure a webhook. For the steps, see [Configure notification delivery](#).
- The Console local deployment Notification Center displays notifications by default, so no additional setup is required for in-console notifications.

View notification rules

The notification rules page is the central place to review and manage every rule that applies to your organization. Each rule shows its key attributes so that you can quickly assess and adjust your alerting behavior.

Steps

1. From the left navigation, select **Health > Alerts**.
2. Select **Notification settings**.
3. Review the rules in the list. Each rule shows its active status, name, the services and severity levels it monitors, the roles authorized to receive notifications, the enabled delivery channels, when it was last modified, and any scheduled mute period.
4. To find a specific rule, use the filter and search controls to filter by active status, service, severity, role, channel, or mute status.

Create a notification rule

Create a notification rule to define the conditions that trigger a notification and how that notification is delivered.



You can't set notification rules for fleets. You can set them only for specific systems. In large environments with many systems, consider creating broader rules by severity rather than duplicating rules per system—for example, one Critical rule that covers all systems acts as a catch-all, with additional targeted rules for systems that need different routing or recipients.

Example notification rule for Critical alerts across all systems

Suppose you want to ensure no critical alert goes unnoticed, regardless of which system it originates from. You can create a broad rule that routes all critical alerts to the Console Notification Center for storage admins.

In this example, you create a rule with the following settings:

- **Services:** All
- **Severity:** Critical
- **IAM role:** Storage admin
- **System:** (Leave this blank to apply to all systems)
- **Delivery method:** Console Notification Center

With this rule in place, storage admins see a notification in the Console for every critical alert across all systems. This rule acts as a baseline that you can supplement with more targeted rules.

Example of a targeted notification rule for Storage admin capacity alerts

Suppose your storage admins need to respond immediately to critical capacity issues on a specific production system, but you don't want lower-severity alerts to flood their inboxes. You can create a targeted rule that emails storage admins only when a critical capacity alert fires on that system.

In this example, you create a rule with the following settings:

- **Services:** ONTAP management
- **Severity:** Critical
- **Impact area:** Capacity
- **IAM role:** Storage admin
- **System:** <system_name>
- **Delivery method:** Email

With this rule in place, Console local deployment emails all storage admins for the specified system when a critical capacity alert occurs. Alerts with a lower severity, such as warning or informational, don't generate an email, so the team can focus on the issues that require urgent attention.

Steps

1. From the left navigation, select **Health > Alerts**.
2. Select **Notification settings**, and then select **Add rule**.
3. Define the conditions that the rule matches:
 - **Rule name:** enter a concise, descriptive name. This field is required.
 - **Rule description:** optionally, enter additional context about the rule's purpose.
 - **Services:** select one or more ONTAP or platform services that the rule applies to.
 - **Roles:** select the access roles that users must have to receive notifications when the rule is triggered.
 - **Severity levels:** select which severity levels the rule monitors, such as Critical, Warning, Error, or Information.
 - **Impact area:** select the operational areas to match, such as Capacity or Performance.
 - **Alert name:** select the specific alert types that trigger the rule.
 - **System:** select the system context that the rule applies to.
4. Select the delivery channels for the notification:
 - **Email:** sends alert emails to users who have the selected roles. Requires a configured email server.
 - **Webhook:** sends alert data to an external system. Requires selecting a configured webhook integration.
 - **Console Notification Center:** displays real-time alerts for users who have the selected roles.
5. Optionally, to suppress notifications from this rule during a planned period, such as a maintenance window, set a **Mute notifications** schedule, and choose a recurrence if you want the mute period to repeat. You can add multiple mute windows per rule, which is useful for recurring maintenance cycles such as weekly patching.
6. Select **Create**.

Enable or disable a notification rule

Enable or disable individual notification rules to control which conditions raise notifications. Disable rules that aren't relevant to your environment to reduce noise, and enable rules to start receiving their notifications again.

Steps

1. From the left navigation, select **Health > Alerts**.
2. Select **Notification settings**.
3. Locate the rule you want to change.
4. Toggle the rule's **Active** switch on or off. The change takes effect immediately.

Mute a notification rule

Mute a notification rule to suppress alert delivery during a planned event, such as a maintenance window, without disabling the rule. Alerts continue to appear in the Alerts page during the mute period—only the email, webhook, and in-console notifications are suppressed. When the mute window expires, notifications resume automatically.

You can add multiple mute windows to a single rule and configure each one to recur on a schedule.

Examples of mute windows

- **One-off maintenance window:** You're running a firmware upgrade on a storage system Saturday night and want to suppress the expected alerts during that window. Set a mute window for Saturday 10:00 PM to Sunday 2:00 AM with no recurrence. When the window expires, notifications resume automatically.
- **Recurring weekly patching window:** Your team patches systems every Sunday from midnight to 4:00 AM. Add a mute window with weekly recurrence so notifications are suppressed automatically each week without manual intervention. If a second system has its own patching schedule at a different time, add a second mute window to the same rule with its own start time and recurrence.

Steps

1. From the left navigation, select **Health > Alerts**.
2. Select **Notification rules**.
3. In the rules list, locate the rule you want to mute and select the action menu for that rule.
4. Select **Edit**.
5. In the **Mute notifications** section, set the start and end date and time for the mute window.
6. Optionally, select a recurrence to repeat the window on a schedule.
7. To add additional mute windows, select **Add mute window** and repeat the previous steps.
8. Select **Save**.

Delete a notification rule

Delete a notification rule if you no longer need it. Deleting a rule permanently removes it, so you can't recover it.

Steps

1. From the left navigation, select **Health > Alerts**.
2. Select **Notification settings**.
3. In the rules list, locate the rule you want to delete and select the action menu for that rule.
4. Select **Delete** from the action menu.

Related information

- [Configure notification delivery](#)
- [Learn about Console alerts](#)
- [View alerts](#)

Remediate storage class drift in NetApp Console local deployment

In NetApp Console local deployment, find drift-related alerts, analyze the drift findings, and remediate workloads that no longer match their storage class intent.

Required roles

Storage admin



You can only view and remediate workloads in fleets where you have permission.

Remediate drift

When a workload no longer matches its storage class intent, NetApp Console local deployment raises an alert. Use the alert to analyze the drift and apply the recommended remediation.

You can apply some remediations directly from the alert. For other issues, update the workload configuration or assign a different storage class to restore compliance. The remediation workflow shows the expected impact before you apply changes.

Steps

1. Select **Storage > Storage classes**.

A summary of storage class drift appears in the **Drifts by storage class** area. The complete list appears in the table.

2. In the **Drifts** column, select **View** for the relevant storage class.

The **Alerts > Overview** page opens with filters applied.

3. Select an alert to open the alert details page.
4. Select **Analyze**.
5. Review the findings in **Workload analyzer**.

For configuration drift findings, compare the intended and actual settings to determine what changed.

6. Select the recommended remediation action, such as **Fix it**.
7. Review the proposed changes and apply the remediation.
8. Return to **Storage > Storage classes** and verify that the workload no longer appears as drifted.

Compliance evaluations can take several minutes to reflect recent configuration changes. If the workload is still listed as drifted, return to the alert details page and select **Analyze** to review any remaining findings.

Related information

- [Learn about storage class drift and compliance in NetApp Console local deployment](#)
- [Learn about storage alerts](#)
- [View alerts](#)

Alert remediation actions in NetApp Console local deployment

Use this reference to understand the remediation actions available from **Fix it** in NetApp Console local deployment. For each action, the table describes what the action does and the related alert. Review the action details in the confirmation dialog before you run it.

Remediation actions

Remediation action	Alert name	Alert description	Impact area	Remediation summary
Enable automatic volume growth	Volume full	The volume is running low on space and is nearing full capacity.	Capacity	Automatically increases a volume's size (up to a configured limit) to reduce the risk of running out of space.
Resize volume	Volume full	The volume is running low on space and is nearing full capacity.	Capacity	Increases a volume's size to provide more space immediately.
Bring volume online	Volume offline	The volume is offline and not serving data.	Availability	Brings an offline volume online so it can resume serving data.
Bring aggregate online	Aggregate offline	The aggregate is not online and volumes hosted on it may be unavailable.	Availability	Brings an offline aggregate online to restore access to the volumes it hosts.
Bring LUN online	LUN offline	The LUN is offline and host access is unavailable.	Availability	Brings an offline LUN online so hosts can resume access and I/O.
Unrestrict volume	Volume restricted	The volume is in a restricted state and may not serve I/O normally.	Availability	Returns a restricted volume to an online state so clients can access it again.
Bring NVMe namespace online	NVMe namespace is offline	The NVMe namespace is offline and host access is unavailable.	Availability	Brings an offline NVMe namespace online to restore host access.
Bring intercluster LIF up	Intercluster LIF down	An intercluster LIF is down and cluster-to-cluster communication may be impacted.	Connectivity	Brings an intercluster LIF up to restore cluster-to-cluster connectivity used for replication.

Remediation action	Alert name	Alert description	Impact area	Remediation summary
Re-enable MetroCluster AUSO	MetroCluster automatic unplanned switchover disabled	Automatic unplanned switchover is disabled on this cluster.	Availability	Re-enables automatic unplanned switchover (AUSO) so MetroCluster protection works as intended.
Configure Service Processor network	Service Processor is not configured	The Service Processor (SP) network is not configured.	Manageability	Configures the Service Processor (SP) network settings to enable out-of-band management access.
Assign unassigned disks	Unassigned disks detected	Unassigned disks are present and available capacity may be unused. Assign the disks to a node so they can be used for storage.	Availability	Assigns currently unassigned disks to a node so they can be used for storage.
Root volume space recovery	Node root volume space low	Available space on the node root volume is low and may affect normal system operation.	System health	Frees space on the node root volume by deleting the largest snapshot or the largest core dump file. Deletions are permanent.

Investigate performance issues

Learn about NetApp Console local deployment Workload Analyzer

Use NetApp Console local deployment Workload Analyzer to view a single volume's behavior over time. You can investigate performance degradation, capacity trends, and resource contention issues from the volume itself, from an alert, or from inventory.

How Workload Analyzer identifies root causes

Use the Workload Analyzer to correlate metrics from a single volume across six sections so you can identify root causes quickly. Select a volume and a time range to view events, performance, and capacity together within the same window. This view can help you determine whether storage is the likely source of an application issue or whether another layer, such as networking, is causing the problem.

The analyzer is most useful when you already know which volume is affected. If you open it from an alert or from volume inventory, Console local deployment opens the analyzer with the volume selected so you can focus on the analysis window and the charts.

The analyzer tracks one volume at a time, so use it to investigate a specific issue rather than compare multiple volumes.

Review the capacity section as part of this assessment. When an ONTAP system is more than 85% full, the high utilization can itself cause performance issues, so low available capacity can explain latency that the performance charts alone don't account for.

After you identify a root cause, you can act directly from the analyzer. When automated options are available, Console local deployment provides Fix-It recommendations with guided steps or one-click remediation.



Workload charts for LUNs don't provide the same level of detail as the charts for volumes, so you'll see fewer statistics when you analyze a LUN.

Access Workload Analyzer

You can open the Workload Analyzer in several ways:

- From the **Health** menu, select **Workload Analyzer**, then search for and select the volume you want to investigate in the **Volume** field.
- From the **Storage > Fleets > Inventory** page, navigate to a volume you want to analyze and select **Analyze** from the action menu.
- From the **Alerts > Overview** page, select an alert for the volume you want to investigate, then select **Analyze** from the alert details.

You can also analyze LUNs, but they show fewer statistics than volumes.

Analyze latency, throughput, and storage capacity

Use these six sections to investigate a volume:

Volume selection

Choose the volume and time range to investigate so all charts and events align to the same workload and analysis window.

Event timeline

View current and historical configuration changes, warning alerts, and critical alerts for the selected volume during the analysis window. Use this section to correlate "what changed" with changes in performance or capacity behavior.

Latency analysis

View volume latency over time, either as a total or broken down by delay center—network, data processing, aggregate operations, cluster interconnect, and others. If the volume uses a QoS policy, the analyzer displays the policy's maximum and minimum latency limits as reference lines so you can see how close the workload is to a throttling threshold. Toggle the forecast to project whether latency is trending toward a warning or critical threshold.

Throughput analysis

View IOPS and MB/s over time, with an optional read/write/other breakdown. If the volume uses a QoS policy, the analyzer displays IOPS and throughput limit lines. Toggle the forecast to project whether demand is trending toward QoS limits.

Resource utilization

View how busy the nodes and aggregates serving the volume are during the analysis window. The analyzer plots each resource as an independent line rather than stacking values, so you can identify whether a specific node or aggregate is a contention source. Toggle the forecast to project whether any resource is trending toward a high-utilization threshold.

Capacity analysis

View how physical space and available space in the volume have changed over time. Hover to see a

storage efficiency breakdown including deduplication, compression, and compaction savings. Toggle the forecast to project when the volume will reach warning or critical capacity thresholds. Switch to the snapshot view to see how snapshot space is tracking against the configured snapshot reserve.

Forecast capacity and performance trends

Use the forecast toggle in each analysis section to extend the chart timeline beyond the current time and project future values based on the observed trend. A dashed line distinguishes forecast values from actual data. The analyzer also displays an insight message when the forecast indicates that a threshold breach is likely, along with an estimated time to breach.

Related information

[Investigate high latency on a volume.](#)

[Investigate high throughput demand on a volume.](#)

[Investigate capacity growth on a volume.](#)

[Learn about Workload Analyzer metrics in NetApp Console local deployment.](#)

Investigate high latency on a volume in NetApp Console local deployment

Use Workload Analyzer in NetApp Console local deployment to find the source of a volume's slow response times.

If you open the analyzer from an alert or from volume inventory, the volume is already selected and you can focus on the time range and chart breakdowns.

When application performance degrades, identify which part of the I/O path is causing the slowdown. The latency analysis section breaks down response time by delay center so you can distinguish between network issues, data processing overhead, aggregate contention, and other contributors.

Steps

1. Open Workload Analyzer using one of the following methods:
 - From the **Health** menu, select **Workload Analyzer**, then search for and select the volume you want to investigate in the **Volume** field.
 - From the **Storage > Fleets > Inventory** page, navigate to a volume you want to analyze and select **Analyze** from the action menu.
 - From the **Alerts > Overview** page, select a capacity-related alert for the volume you want to investigate, then select **Analyze** from the alert details.
2. Set the **Time Range** to cover the period when the performance issue occurred, then select **Analyze**.
3. Review the **Event timeline** section to see configuration changes and alerts that line up with the latency increase.

The analyzer plots configuration change events (wrench icon) and alert events (warning and critical icons) along the same time axis as the latency chart, making it easy to see whether a change preceded the degradation.

4. In the **Latency** section, open the **View** dropdown and select **Breakdown by delay center**. The chart shows the contribution of each delay center to the total latency over time. Delay centers include:
 - Read latency
 - Write latency
 - Other latency

5. Hover over a point on the chart where latency is highest to see each delay center's value and its percentage of total latency.

The largest delay center usually points to the resource under contention. When a shared resource cannot keep up with demand, the volumes that use it wait longer for I/O. Reduce the load on that resource, such as by moving workloads or adjusting a QoS limit, to lower latency.

6. Identify the dominant contributor and use the value to determine next steps:
 - High **Read latency** may indicate disk contention. Check the **Resource Utilization** section to confirm aggregate busy percentage.
 - High **Write latency** may indicate NIC saturation or network path issues outside the cluster.
 - High **Other latency** may indicate that inline efficiency settings are consuming more CPU than the workload can tolerate. Consider adjusting efficiency settings or QoS.
 - High **Cloud Latency** may indicate the workload is frequently accessing cold data on the capacity tier. Consider adjusting the tiering policy.
7. If the delay centers don't account for the slowdown, check the **Capacity Analysis** section. When the ONTAP system is more than 85% full, the high utilization can cause performance issues regardless of the delay center breakdown.
8. If latency increased immediately after a change event, use the event details and the related charts together to validate the likely cause:
 - For QoS changes, compare the latency line to the QoS limit lines and review the throughput charts to see whether demand is meeting a policy ceiling.
 - For aggregate or volume moves, compare the latency spike to the node and aggregate utilization values shown for the same timeframe.
 - For tiering policy changes, review the cloud latency contribution to determine whether cold-data access increased after the change.

After you finish

If a configuration or placement issue caused the problem and Console local deployment can resolve it, the alert on the volume may offer a Fix-It option.

[Learn about Workload Analyzer metrics in NetApp Console local deployment.](#)

Investigate high throughput on a volume in NetApp Console local deployment

Use Workload Analyzer in NetApp Console local deployment to examine a volume's IOPS and throughput demand over time.

When a volume's workload is consuming more IOPS or throughput than expected, identify what's driving the demand. The throughput analysis section shows IOPS and MB/s with an optional read, write, and other breakdown so you can identify which type of I/O is driving high demand and whether that demand is trending toward a QoS limit.

If you open the analyzer from an alert or from volume inventory, the volume is already selected and you can focus on the time range and throughput charts.

Steps

1. Open Workload Analyzer using one of the following methods:
 - From the **Health** menu, select **Workload Analyzer**, then search for and select the volume you want to

investigate in the **Volume** field.

- From the **Storage > Fleets > Inventory** page, navigate to a volume you want to analyze and select **Analyze** from the action menu.
- From the **Alerts > Overview** page, select an alert for the volume you want to investigate, then select **Analyze** from the alert details.

2. Set the **Time Range** to cover the period of high demand, then select **Analyze**.
3. Scroll to the **Throughput Analysis** section.
4. Open the **View** dropdown and select **Breakdown (RWO)**.

The charts split into read, write, and other components for both IOPS and MB/s. This lets you determine whether read-heavy access patterns, write-heavy patterns, or a combination is driving the demand.

5. Use the component picker to enable or disable the metrics you want to examine:
 - **Read IOPS** and **Write IOPS** — operations per second by I/O direction
 - **Read MB/s** and **Write MB/s** — throughput in megabytes per second by I/O direction
 - **Other IOPS** and **Other MB/s** — metadata and other non-data operations
 - **QoS IOPS Limit** and **QoS MB/s Limit** — policy ceilings that appear only when the volume uses a QoS policy
 - **Cloud Throughput** — MB/s written to and read from the cloud, shown only for workloads that tier capacity to the cloud through FabricPool
6. Hover over a peak in the chart to see the exact value and percentage breakdown for each component at that point in time.

The hover tooltip also shows the average I/O size (throughput divided by IOPS) for each category, which can indicate whether the workload is performing large sequential I/O or small random I/O.

7. Enable **Show Forecast** to project whether current throughput trends will reach the QoS IOPS or MB/s limit.

If the forecast line is approaching a QoS limit line, ONTAP might throttle the workload soon.

8. If you want to see aggregate demand without the breakdown, open the **View** dropdown and select **Totals**.

The totals view shows a single IOPS line and a single MB/s line, which is useful for a quick overall demand summary.

After you finish

Use the throughput patterns you observe to decide what to do next:

- If read IOPS are very high relative to write IOPS, consider whether read-ahead settings or caching can reduce the load on the volume.
- If the workload is approaching or has reached the QoS IOPS or MB/s limit, use the QoS limit lines and related metric definitions to confirm throttling and decide whether you need to adjust the limit.
- If throughput is high and latency is also elevated, check the **Resource Utilization** section to determine whether the underlying node or aggregate is under contention. Compare throughput peaks, latency spikes, and resource utilization over the same period.
- If the analyzer shows rapid capacity or snapshot growth while demand is increasing, review the capacity and snapshot metrics to understand how that growth might affect the volume.

[Investigate high latency on a volume.](#)

[Learn about Workload Analyzer metrics in NetApp Console local deployment.](#)

Investigate capacity growth on a volume in NetApp Console local deployment

Use Workload Analyzer in NetApp Console local deployment to investigate why a volume is running out of space.

When a volume is filling up or snapshot copies are consuming more space than expected, examine capacity and snapshot trends over time. The capacity analysis section shows how physical and available space change, breaks down storage efficiency savings, and projects when the volume will reach a capacity threshold.

If you open the analyzer from an alert or from volume inventory, the volume is already selected and you can focus on capacity trends and the forecast window.

Steps

1. Open Workload Analyzer using one of the following methods:
 - From the **Health** menu, select **Workload Analyzer**, then search for and select the volume you want to investigate in the **Volume** field.
 - From the **Storage > Fleets > Inventory** page, navigate to a volume you want to analyze and select **Analyze** from the action menu.
 - From the **Alerts > Overview** page, select an alert for the volume you want to investigate, then select **Analyze** from the alert details.
2. Set the **Time Range** to cover the period of capacity growth, then select **Analyze**.
3. Scroll to the **Capacity Analysis** section.
4. Select **Capacity view** from the **View** dropdown.

The chart shows physical capacity as the bottom stacked area and available capacity as the top stacked area. Together they equal the total volume size, so you can see how quickly available space is decreasing.

5. Hover over a point on the chart to see the storage efficiency breakdown, including deduplication, compression, and compaction savings, and the overall storage efficiency ratio.

A declining efficiency ratio while physical capacity rises can indicate that newly written data isn't deduplicating or compressing as well as the existing data.

6. To investigate snapshot consumption, select **Snapshot View** from the **View** dropdown.

The chart shows the snapshot reserve as a horizontal reference line and snapshot used capacity as an area below it. Hover to see the snapshot space used and its percentage of the reserve, the total snapshot count, and the age of the oldest snapshot.

When snapshot consumption exceeds the reserve, snapshots begin consuming space from the volume's data area, which reduces the space available for new writes.

7. If the volume tiers data to the cloud, select **Cloud Tier View** from the **View** dropdown to compare how much capacity is on the local performance tier versus the cloud capacity tier.
8. Enable **Show Forecast** to project when the volume will reach a warning or critical capacity threshold.

The capacity forecast requires at least 10 days of historical data. When fewer than 30 days of capacity remain before the volume is full, the analyzer identifies the storage as almost full. The forecast displays an

insight message with an estimated time to breach.

After you finish

Use the capacity trends you observe to decide what to do next:

- If available capacity is trending toward full, consider increasing the volume size, enabling autogrow, or moving data off the volume.
- If snapshot used capacity is exceeding the reserve, review your snapshot policy and retention to reclaim space.
- If the storage efficiency ratio is dropping, review whether the workload's data type or inline efficiency settings are reducing savings. Use [Learn about Workload Analyzer metrics in NetApp Console local deployment](#) for detailed descriptions of capacity, snapshot, and efficiency metrics.

Workload Analyzer metrics in NetApp Console local deployment

In NetApp Console local deployment, the Workload Analyzer shows metrics across six sections. Each metric description explains how the analyzer renders the metric on the chart and what it indicates about volume behavior.

QoS reference lines in latency charts

If the selected volume uses a QoS policy, you can use two additional reference lines in the metric picker:

Reference line	Description
QoS Limit Max	The maximum latency ceiling defined by the QoS policy. When the latency line approaches or touches this line, ONTAP is throttling the workload, and the QoS limit—rather than a physical resource—is the dominant source of latency.
QoS Limit Min	The guaranteed minimum latency floor defined by the QoS policy. This line indicates the response time floor enforced for the workload. When other workloads use QoS minimums to guarantee their throughput, ONTAP can throttle this workload, which then shows higher latency.

These horizontal lines do not appear in the delay center breakdown on hover. They serve as reference thresholds, not as contributors to latency.

Latency breakdown by I/O type

Use the I/O type breakdown in the **Latency Analysis** section after you select **Breakdown by delay center** from the View dropdown. The chart breaks total latency into three categories based on the direction of the I/O operation. Use these metrics to determine whether a performance issue is affecting reads, writes, or metadata operations.

Latency type	Description	Common causes of elevated values
Read latency	Average time to complete a client read request on the volume, from when the request is received until data is returned. Read requests must traverse the full I/O path from the network protocol layer through the data processing stack to the aggregate where the data resides.	Aggregate or disk contention; cache misses that promote reads to physical disk; cold data retrieved from a cloud capacity tier via FabricPool; cross-node reads when data resides on a different node than the one handling the request.
Write latency	Average time to acknowledge a client write request on the volume. ONTAP acknowledges a write once it is committed to the NVRAM write log; the data is later flushed to the aggregate.	NIC saturation or high round-trip latency between the client and the storage node; synchronous SnapMirror waiting for the destination cluster to confirm receipt before the write can be acknowledged; NVRAM pressure under heavy write loads; CPU overhead from inline deduplication, compression, or compaction running in the write I/O path.
Other latency	Average time to complete non-read, non-write operations on the volume, including file opens, attribute lookups, directory traversals, file creates, and locks. Elevated other latency can affect application responsiveness even when read and write latency appear normal.	CPU pressure from inline efficiency operations competing with metadata processing; high namespace activity from workloads that generate large numbers of file creates, deletes, or directory scans; QoS throttling that constrains total IOPS, leaving fewer IOPS available for metadata operations; volume activation overhead when many volumes are concurrently active.

Throughput components

Use the throughput components in the **Throughput Analysis** section after you select **Breakdown (RWO)** from the View dropdown. The analyzer shows both IOPS and MB/s simultaneously.

Component	Description	Rendered as
Read IOPS	Read operations per second.	Stacked area on the IOPS chart.
Write IOPS	Write operations per second.	Stacked area on the IOPS chart.
Other IOPS	Metadata and other non-data operations per second.	Stacked area on the IOPS chart.
Read MB/s	Read throughput in megabytes per second.	Stacked area on the MB/s chart.
Write MB/s	Write throughput in megabytes per second.	Stacked area on the MB/s chart.
Cloud Throughput	Throughput in megabytes per second between the volume and the cloud capacity tier. This metric appears only for workloads that tier capacity to the cloud through FabricPool.	Stacked area on the MB/s chart.

The Read, Write, and Other components sum to the total IOPS or MB/s value. Hover over the chart to see

each component's value, its percentage of the total, and the average I/O size (MB/s ÷ IOPS) per category.

Resource utilization metrics

Use the **Resource Utilization** section to review resource utilization metrics. The analyzer shows each resource serving the volume as an independent line. Resources do not stack.

Node metrics

Metric	Description
Node utilization	Composite utilization percentage for the node. Hover to see the CPU utilization, network utilization, and available headroom for each node.
CPU utilization	Percentage of node CPU capacity in use. The hover tooltip shows this value.
Network utilization	Percentage of node network capacity in use. The hover tooltip shows this value.
Headroom	Remaining node capacity available for additional workload. The hover tooltip shows this value.

Aggregate metrics

Metric	Description
Aggregate utilization	Disk busy percentage for the aggregate. Hover to see the disk busy value, number of volumes on the aggregate, and available headroom.
Disk busy	Percentage of time the aggregate's disks actively service I/O. The hover tooltip shows this value.
Volume count	Number of volumes currently hosted on the aggregate. The hover tooltip shows this value.
Headroom	Remaining aggregate capacity available for additional workload. The hover tooltip shows this value.

When a node or aggregate utilization line crosses the high-utilization warning threshold, which the chart marks with a dashed reference line, other workloads on that resource may be competing with the selected volume for I/O capacity.

Capacity metrics

Capacity View

Use the **Capacity Analysis** section to review Capacity View metrics after you select **Capacity View** from the View dropdown.

Metric	Description
Physical Capacity	Space consumed on disk after storage efficiency operations. This is the bottom stacked area in the chart. Physical + Available always equals the total volume size.
Available Capacity	Free space remaining in the volume. This is the top stacked area in the chart. It decreases as Physical Capacity grows.

Metric	Description
Storage Efficiency Ratio	Overall efficiency ratio (logical data ÷ physical data). The hover tooltip shows this value. It reflects combined savings from deduplication, compression, and compaction.
Deduplication savings	Space saved by removing duplicate data blocks. The hover tooltip shows this value.
Compression savings	Space saved by compressing data inline. The hover tooltip shows this value.
Compaction savings	Space saved by packing multiple small blocks into a single physical block. The hover tooltip shows this value.

Snapshot View

Use Snapshot View when you want to review snapshot-specific metrics.

Metric	Description	Rendered as
Available snapshot	Pre-allocated space reserved for snapshots, configured as a percentage of the volume size.	Horizontal reference line.
Used snapshot	Actual space consumed by snapshot copies.	Area chart below the reserve line.

Hover over the chart to see the snapshot reserve size, snapshot space used and its percentage of the reserve, total snapshot count, and the age of the oldest snapshot. When snapshot consumption exceeds the reserve, snapshots begin consuming space from the volume's data area.

Forecast behavior

Use the **Show Forecast** toggle in any analysis section when you want to project future values beyond the current time based on the observed trend. The following behaviors apply across all sections:

Aspect	Behavior
Projection window	Projects forward based on the selected time range. For a 2-hour view, the projection window is approximately 1 hour. For a 7-day view, the projection window extends several days.
Visual style	The analyzer renders forecast values as a dashed line. A vertical separator marks the boundary between actual data and projected values.
Threshold alerts	The analyzer displays an insight message when the forecast indicates that the volume will breach a warning or critical threshold, along with an estimated time to breach.
Trend basis	The projection uses the rate of change from the most recent portion of the selected time range. High volatility in recent data reduces forecast confidence.
Minimum data required	Capacity forecasting requires at least 10 days of historical data. When fewer than 30 days of capacity remain before the volume is full, the analyzer identifies the storage as almost full.

Related information

- [Learn about NetApp Console local deployment Workload Analyzer](#)
- [Investigate high latency on a volume](#)
- [Investigate high throughput demand on a volume](#)
- [Investigate capacity growth on a volume](#)

Administer and monitor NetApp Console local deployment

Learn about administration and monitoring in NetApp Console local deployment

After you deploy NetApp Console local deployment, use the administration and monitoring tools to review activity, maintain the VM, and manage member access.

Audit activity

Review user and API activity, track the status of Console operations, download audit logs, and export logs to external tools when you need longer retention or additional analysis.

- [Audit NetApp Console local deployment activity](#)
- [Set up notifications delivery channels in NetApp Console local deployment](#)

Licenses and subscriptions

Use the licensing and subscriptions information to understand service entitlements and subscription status for your Console environment.

[Learn about licenses and subscriptions.](#)

Maintain and troubleshoot the Console local deployment VM

Maintain the VM that hosts Console local deployment, review network and system settings, access diagnostic tools, and troubleshoot issues when the service or agent does not behave as expected.

[Maintain your vCenter or ESXi host for Console local deployment.](#)

Manage users and access

Review member access across your organization, adjust access at the fleet level, and manage security settings such as MFA recovery and service account credentials.

- [Manage member access in NetApp Console](#)
- [View or change fleet access assignments in NetApp Console local deployment](#)
- [Manage member security in NetApp Console local deployment](#)

Audit NetApp Console local deployment activity

You can audit user activity initiated from both the UI and API from the Audit page, and you can monitor the status of operations in NetApp Console local deployment and the user who initiated them.

Required roles

Super admin, Organization admin, Folder and fleet admin, Operations support analyst, Super viewer, Organization viewer, or Folder and fleet viewer. [Learn about access roles](#).

You can view audit activity in the Console local deployment, download a CSV of the audit logs, or configure the audit logs to be sent to your own syslog server by using a webhook.

Audit user activity from the Audit page

Use the Audit page to identify who performed an action or its status.

The Audit page shows the actions that users completed within your organization. For example, you can verify who added a member to an organization or that a fleet was deleted successfully, as well as see other storage management actions that users performed in your organization.

The audit logs show activity for the following services:

- **Tenancy:** actions related to managing your organization, such as adding users, creating fleets, and associating resources.
- **Storage-management:** actions related to managing your storage resources.
- **Federation:** actions related to managing federations, such as creating federations and adding or removing federated organizations.

Steps

1. Select **Administration > Audit**.
2. Use the filters above the table to change which actions display in the table.

For example, you can use the **Service** filter to show actions related to a specific service, or you can use the **User** filter to show actions related to a specific user account.

Filter audit logs for identity and access management actions

To view only actions related to managing your organization—such as adding members, creating fleets, or deleting resources—filter the audit log by the Tenancy service.

Steps

1. Select **Administration > Audit**.
2. Use the filters to narrow the results. Select **Service** and then select **Tenancy**.
3. Use any of the other filters to further refine the results.

For example, use the **User** filter to show IAM actions performed by a specific user account.

Download audit logs from the Audit page

You can download the audit logs from the Audit page to a CSV file. This enables you to keep a record of the actions that users perform in your organization. The downloaded CSV file includes all audit log columns, regardless of the filters or displayed columns on the Audit page.

Steps

1. In the **Audit** page, select the download icon in the upper right corner of the table.

Maintain NetApp Console

Maintain your vCenter or ESXi host for NetApp Console local deployment

In NetApp Console local deployment, you can make changes to your existing VCenter or ESXi host after you deploy the Console local deployment. For example, you can increase the CPU or RAM of the VM instance that hosts the Console local deployment.

Perform these maintenance tasks using the VM web console:

- Increase disk size
- Restart the Console local deployment
- Update static routes
- Update search domains

Limitations

You can only view (not update) information about the IP address, DNS, and gateways through the maintenance console.

Access the VM maintenance console

You can access the maintenance console from the VSphere client.

Steps

1. Open the VSphere client and log in to your VCenter.
2. Select the VM instance that hosts the Console local deployment.
3. Select **Launch Web Console**.
4. Log in to the VM instance using the user name and password that you specified when you created the VM instance. The username is `maint` and the password is the one that you specified when you created the VM instance.

Change the maint user password

You can change the password for the `maint` user.

Steps

1. Open the VSphere client and log in to your VCenter.
2. Select the VM instance that hosts the Console local deployment.
3. Select **Launch Web Console**.
4. Log in to the VM instance using the user name and password that you specified when you created the VM instance. The username is `maint` and the password is the one that you specified when you created the VM instance.
5. Enter `1` to view the `System Configuration` menu.
6. Enter `1` to change the maintenance user password and follow the on-screen prompts.

Increase the CPU or RAM of the VM instance

You can increase the CPU or RAM of the VM instance that hosts Console local deployment.

Edit the VM instance settings in your VCenter or ESXi host, then use the maintenance Console to apply the changes.

Steps in the VSphere client

1. Open the VSphere client and log in to your VCenter.
2. Select the VM instance that hosts the Console local deployment.
3. Right-click the VM instance and select **Edit Settings**.
4. Increase the hard drive space used for /opt or the /var partition.
 - a. Select **Hard Disk 2** to increase the hard drive space used for /opt.
 - b. Select **Hard Disk 3** to increase the hard drive space used for /var.
5. Save your changes.

Steps in the maintenance console

1. Open the VSphere client and log in to your VCenter.
2. Select the VM instance that hosts Console local deployment.
3. Select **Launch Web Console**.
4. Log in to the VM instance using the user name and password that you specified when you created the VM instance. The username is `maint` and the password is the one that you specified when you created the VM instance.
5. Enter `1` to view the ``System Configuration` menu.
6. Enter `2` and follow the on-screen prompts. The console scans for new settings and increases the size of the partitions.

View network settings for the Console local deployment VM

View the network settings for the Console local deployment VM in the VSphere client to confirm or troubleshoot network issues. You can only view (not update) the following network settings: IP address and DNS details.

Steps

1. Open the VSphere client and log in to your VCenter.
2. Select the VM instance that hosts the Console local deployment.
3. Select **Launch Web Console**.
4. Log in to the VM instance using the user name and password that you specified when you created the VM instance. The username is `maint` and the password is the one that you specified when you created the VM instance.
5. Enter `2` to view the `Network Configuration` menu.
6. Enter a number between 1 and 6 to view the corresponding network settings.

Update the static routes for the Console local deployment VM

Add, update, or remove static routes for the Console local deployment VM as needed.

Steps

1. Open the VSphere client and log in to your VCenter.
2. Select the VM instance that hosts the Console local deployment.
3. Select **Launch Web Console**.
4. Log in to the VM instance using the user name and password that you specified when you created the VM instance. The username is `maint` and the password is the one that you specified when you created the VM instance.
5. Enter `2` to view the `Network Configuration` menu.
6. Enter `7` to update static routes and follow the on-screen prompts.
7. Press Enter.
8. Optionally, make additional changes.
9. Enter `9` to commit your changes.

Update domain search settings for the Console local deployment VM

You can update the search domain settings for the Console local deployment VM.

Steps

1. Open the VSphere client and log in to your VCenter.
2. Select the VM instance that hosts the Console local deployment.
3. Select **Launch Web Console**.
4. Log in to the VM instance using the user name and password that you specified when you created the VM instance. The username is `maint` and the password is the one that you specified when you created the VM instance.
5. Enter `2`` to view the `Network Configuration` menu.
6. Enter `8` to update the domain search settings and follow the on-screen prompts.
7. Press Enter.
8. Optionally, make additional changes.
9. Enter `9` to commit your changes.

Access the Console local deployment diagnostic tools

Access diagnostic tools to troubleshoot issues with the Console local deployment. NetApp Support may ask you to do this when troubleshooting issues.

Steps

1. Open the VSphere client and log in to your VCenter.
2. Select the VM instance that hosts the Console local deployment.
3. Select **Launch Web Console**.
4. Log in to the VM instance using the user name and password that you specified when you created the VM instance. The username is `maint` and the password is the one that you specified when you created the VM instance.
5. Enter `3` to view the Support and Diagnostics menu.

6. Enter 1 to access the diagnostic tools and follow the on-screen prompts.

Access the Console local deployment diagnostic tools remotely

You can access diagnostic tools remotely with a tool such as Putty. Enable SSH access to the Console local deployment VM by assigning a one-time password.

SSH access enables advanced terminal features like copy and paste.

Steps

1. Open the VSphere client and log in to your VCenter.
2. Select the VM instance that hosts the Console local deployment.
3. Select **Launch Web Console**.
4. Log in to the VM instance using the user name and password that you specified when you created the VM instance. The username is `maint` and the password is the one that you specified when you created the VM instance.
5. Enter 3 to view the `Support and Diagnostics` menu.
6. Enter 2 to access the diagnostic tools and follow the on-screen prompts to configure a one-time password that expires in 24 hours.
7. Use an SSH tool such as Putty to connect to the Console local deployment VM using the user name `diag` and the one-time password that you configured.

Manage users and access

Manage member access in NetApp Console local deployment

In NetApp Console local deployment, review or change a user's roles across multiple folders or fleets, such as checking everything a storage engineer can access, adding a new role in another region, or removing that user's access when responsibilities change.

Required access roles

Super admin, Organization admin, or Folder or fleet admin (for folders and fleets that they are administering).
[Learn about access roles.](#)

You can view and manage access in two ways depending on your needs. If you want to view the roles that a particular user has across all the fleet in your organization, use the **Members** page.

If you want to confirm who can access a specific fleet, review the members assigned to that fleet instead.
[Manage fleet access assignments.](#)

To add a new member, service account, or directory user and assign their first role, use the onboarding task instead. [Add members and assign roles.](#)

Understand how access is granted in NetApp Console

NetApp Console uses role-based access control (RBAC) to manage member permissions. You can assign predefined roles at the organization, folder, or fleet level, depending on the scope of access that a member needs.

Using role inheritance

A role that you assign at the organization or folder level is inherited by the child scopes beneath it, so change an inherited assignment at the higher scope where you originally granted it.

[Learn about role inheritance in NetApp Console local deployment.](#)

View roles assigned to a member

Confirm exactly which roles a member holds and at which scope before you make a change. For example, check whether a teammate already has the Storage admin role on the `Production-EU-West` fleet.

If you have the *Folder or fleet admin* role, the page displays all members in the organization. However, you can view and manage permissions only for the folders and fleets that you administer. [Learn about Folder or fleet admin actions in NetApp Console local deployment.](#)

1. From the **Members** page, navigate to a member in the table, select **...** and then select **View details**.
2. In the table, expand the respective row for organization, folder, or fleet where you want to view the member's assigned role and select **View** in the **Role** column.

Assign or modify member access

You can adjust a member's access whenever responsibilities change. For example, when a teammate takes on backup duties for a second region, add the Backup and Recovery admin role on that region's fleet without touching their existing storage roles.

If you need to add a new member and assign their first role, see [Add members and assign roles](#).

Add an access role to an existing member

Grant a member an additional role at the organization, folder, or fleet level when their responsibilities expand. For example, give a Storage admin the Backup and recovery admin role at the organization level so they can manage backup and recovery tasks across every fleet without losing their existing storage permissions.

You can assign a member an access role for your organization, folder, or fleet.

Members can have multiple roles within the same fleet and in different fleets. For example, smaller organizations may assign all available access roles to the same user, while larger organizations may have users do more specialized tasks. Alternatively, you could also assign one user the Ransomware resilience admin role at the organization level. In that example, the user would be able to perform Ransomware resilience tasks on all fleets within your organization.

Your access role strategy should align with the way you have organized your NetApp resources.

Steps

1. Select **Administration > Identity and access**.
2. Select **Members**.
3. Select one of the member tabs: **Users**, **Directory users**, or **Service accounts**.
4. Select the actions menu **...** next to the member that you want to assign a role and select **Add a role**.
5. To add a role, complete the steps in the dialog box:
 - **Select an organization, folder, or fleet:** Choose the level of your resource hierarchy that the member should have permissions for.

If you select the organization or a folder, the member will have permissions to everything that resides within the organization or folder.

- **Select a category:** Choose a role category. [Learn about access roles](#).
- **Select a Role:** Choose a role that provides the member with permissions for the resources that are associated with the organization, folder, or fleet that you selected.
- **Add role:** If you want to provide access to additional folders or fleets within your organization, select **Add role**, specify another folder or fleet or role category, and then select a role category and a corresponding role.

6. Select **Add new roles**.

Change a member's assigned role

Replace a member's existing role with a different one when their responsibilities shift. For example, when a Storage viewer on the `Production-US-East` fleet needs the Storage admin role instead.



Each user must have at least one role. You can't remove all roles from a user. If you need to remove all roles, delete the user from your organization.

Steps

1. Select **Administration > Identity and access**.
2. Select **Members**.
3. Select one of the member tabs: **Users**, **Directory users**, or **Service accounts**.
4. From the **Members** page, navigate to a member in the table, select **...** and then select **View details**.
5. In the table, expand the respective row for organization, folder, or fleet where you want to change the member's assigned role and select **View** in the **Role** column to view the roles assigned to this member.
6. You can change an existing role for a member or remove a role.
 - a. To change a member's role, select **Change** next to the role you want to change. You can only change a role to a role within the same role category. For example, you can change from one data service role to another. Confirm the change.
 - b. To unassign a member's role, select  next to the role to unassign the member the respective role. You'll be asked to confirm the removal.

Remove a member from your organization

Remove a member when they leave the organization or change roles in a way that ends their need for Console access. For example, when a contractor's engagement ends or an employee transfers to a team outside your Console organization.



Directory users

Removing a user from your directory prevents that user from authenticating. You should also remove the user from your Console organization to keep the member list accurate and to remove any roles that were assigned in Console local deployment.

Steps

1. Select **Administration > Identity and access**.
2. Select **Members**.

3. Select one of the member tabs: **Users**, **Directory users**, or **Service accounts**.
4. From the **Members** page, navigate to a member in the table, select **...** then select **Delete user**.
5. Confirm that you want to remove the member from your organization.

View or change fleet access assignments in NetApp Console local deployment

Audit or change member access assignments for folders and storage fleets in NetApp Console local deployment. Folder and fleet admins typically work from this view because it shows only the scopes that they administer.

Required access roles

Super admin, Organization admin, or Folder or fleet admin. [Learn about access roles](#).

Alternatively, you can manage all roles for a particular member across multiple folders or fleets. [Manage member access](#).

How folder and fleet access works

Console local deployment uses role-based access control (RBAC). A role that you assign at the folder level applies to all fleets and subfolders within that folder; a role that you assign at the fleet level applies only to that fleet's resources. [Learn about role inheritance in NetApp Console local deployment](#).



You can't change a role at the fleet level if a member inherits it from a folder or from the organization. To change inherited access, change the role at the higher scope.

View members assigned to a folder or fleet

Review exactly who can manage specific folder or fleet. For example, before associating a new production ONTAP cluster with the `Production-US-East` fleet, open the fleet's Access tab to confirm who has access.

Steps

1. Select **Administration > Identity and access**.
2. Select **Organization**.
3. From the Organization page, navigate to a folder or fleet in the table, select **...** and then select **Edit folder** or **Edit fleet**.
4. Select **Access** to view the members who have access to the folder or fleet.

Modify member access from a folder or fleet

Adjust the roles of members who already belong to your organization, scoped to a single folder or fleet. For example, when a teammate moves from a development fleet to a production fleet, promote them from Storage viewer to Storage admin on the production fleet without affecting their access elsewhere.

To add a new member and assign their first role, use the onboarding task instead. [Add members and assign roles](#).

Steps

1. From the Organization page, navigate to a folder or fleet in the table, select **...** and then select **Edit folder** or **Edit fleet**.
2. Select **Access** to view the list of members who have access.

3. Modify member access:

- **Change a member's role:** For any member with a role other than Organization admin, select their existing role and choose a new role. The change applies only at this scope; roles inherited from a higher scope don't appear here.
- **Remove member access at this scope:** For a member who has a role defined directly at this folder or fleet, remove the role to revoke their access at this scope. This doesn't remove the member from your organization or affect roles they have at other scopes.

[Remove a member from your organization.](#)

4. Select **Apply**.

Manage member security in NetApp Console local deployment

Secure user access to your NetApp Console local deployment organization by managing member security settings. You can direct local users to change their own passwords, manage local users' multi-factor authentication (MFA), and recreate service account credentials.

Required access roles

Super admin, Organization admin, or Folder or fleet admin. [Learn about access roles.](#)

Reset user passwords (local users only)

Administrators cannot reset local user passwords directly.

Direct local users to reset their passwords from the Console local deployment login page by selecting **Forgot password?**.



This option is not available for directory users.

Manage a local user's multi-factor authentication (MFA)

If a local user loses access to their MFA device, you can either remove or disable their MFA configuration.



Multi-factor authentication is only available for local users. Directory users cannot enable MFA through Console local deployment.

Use these guidelines to choose the right action:

- Select **Disable** when the user temporarily loses access to their MFA device. Disabling MFA allows one sign-in without MFA for eight hours and then automatically re-enables MFA.
- Select **Remove** when the user must fully reset MFA. After you remove MFA, the user signs in without MFA until they configure MFA again.

If a user has a saved recovery code, they can sign in with it. If a user does not have a recovery code, disable MFA so the user can sign in and regain access.



To manage a local user's multi-factor authentication, you must have an email address in the same domain as the affected user.

Steps

1. Select **Administration > Identity and access**.
2. Select **Members**.
3. From the Members page, navigate to a member in the table, select **...** and then select **Manage multi-factor authentication**.
4. Choose whether to remove or to disable the user's MFA configuration.

After you finish

Review the audit log to confirm who changed MFA access, when they changed it, and whether the action succeeded. [Learn how to audit NetApp Console local deployment activity](#).

Recreate the credentials for a service account

You can create new credentials for a service account if you lose or need to update them.

Creating new credentials deletes the old ones. You cannot use the old credentials.

Steps

1. Select **Administration > Identity and access**.
2. Select **Members**.
3. In the Members table, navigate to a service account, select **...** and then select **Recreate secrets**.
4. Select **Recreate**.
5. Download or copy the client ID and client secret.

The Console local deployment shows the client secret only once. Make sure you copy or download it and store it securely.

Reference

NetApp Console local deployment API

Manage your NetApp Console local deployment organization and fleet IDs

In NetApp Console local deployment, your NetApp Console organization has a name and an ID. You can choose a name for your organization to help identify it. You may also need to retrieve the organization ID for certain integrations.

Required access roles

Super admin or Organization admin. [Learn about access roles](#).


Rename your organization

You can rename your organization.

Steps

1. Select **Administration > Identity and access**.
2. Select **Organization**.

3. From the **Organization** page, navigate to the first row in the table, select **...** and then select **Edit organization**.



4. Enter a new organization name and select **Apply**.

Get the organization ID

The organization ID is used for certain integrations with the Console.

You can view the organization ID from the Organizations page and copy it to the clipboard for your needs.

Steps

1. Select **Administration > Identity and access > Organization**.
2. On the **Organization** page, look for your organization ID in the summary bar and copy it to the clipboard. You can save this for use later or copy it directly to where you need to use it.

Obtain the ID for a fleet


You'll need to obtain the ID for a fleet if you are using the API.

Steps

1. From the **Organization** page, navigate to a fleet in the table and select **...**

The fleet ID displays.

2. To copy the ID, select the copy button.



Related information

- [Learn about identity and access management](#)
- [Get started with identity and access in NetApp Console](#)
- [Learn about the API for identity and access](#)

Supported LLM providers and models in NetApp Console local deployment

NetApp Console local deployment supports OpenAI, OpenAI-compatible, and Anthropic LLM provider types. The models you can select depend on the provider type and endpoint you configure.



This documentation regarding connecting an LLM to the Console assistant to enable AI features is provided as a technology preview. With this preview offering, NetApp reserves the right to modify offering details, contents, and timeline before General Availability.

Choose a model that matches your performance, cost, and governance requirements.

See how provider type affects model availability

The provider type you choose determines which models appear in the Console local deployment model list:

- **OpenAI** — provides models for direct OpenAI integration.
- **OpenAI-compatible** — provides models that your organization's compatible endpoint exposes.
- **Anthropic** — provides models for direct Anthropic integration.

The models you see can vary based on endpoint configuration, proxy or gateway behavior, and organizational policy. Provider types differ by the endpoint you connect to and the models that endpoint exposes, not the transport protocol.

Supported OpenAI models

- GPT-5.4
- GPT-5.5

Supported Anthropic models

- Claude Sonnet 4.6
- Claude Opus 4.6, 4.7 & 4.8

Related information

- [Connect your LLM and enable the NetApp Console assistant.](#)
- [Learn about LLM integration in NetApp Console local deployment.](#)

ONTAP alerts reference in NetApp Console local deployment

This reference lists the ONTAP alerts that NetApp Console local deployment can monitor, organized by impact category.

Severity mapping

The same EMS alert can appear as Critical, Warning, or Info, depending on the ONTAP event that caused it:

- **Critical:** Maps from ONTAP severities `alert`, `emergency`
- **Warning:** Maps from ONTAP severity `error`
- **Info:** Maps from ONTAP severities `notice`, `informational`
- **Other:** Passed through as-is

Dynamic mapping lets a single alert rule emit different severities depending on the runtime context of the EMS event.

The sections below group alerts by impact category and sort each table alphabetically by alert name.

Availability alerts

These alerts can affect system, service, or data availability.

Alert name	Severity	Type	Description
Active Directory Server Connection Down	Critical	EMS	All configured AD/LDAP servers for this SVM are unreachable.
Aggregate isn't online	Critical	Metric	Some aggregates are offline. Volume creation could cause duplicate FSIDs.
Antivirus server busy	Critical, Warning, Info	EMS	The antivirus server is overloaded and can't accept additional scan requests.

Alert name	Severity	Type	Description
AWS credentials not initialized	Critical, Warning, Info	EMS	A module attempted to access AWS IAM role-based credentials before they were initialized.
Cloud tier unreachable	Critical, Warning, Info	EMS	A storage node cannot connect to Cloud Tier object store API. Some data will be inaccessible.
Disk failure	Critical	Metric	A disk has failed, is being sanitized, or is in Maintenance mode.
Disk out of service	Critical, Warning, Info	EMS	A disk was taken out of service due to failure, sanitization, or maintenance.
Disk shelves power supply removed	Critical, Warning, Info	EMS	A power supply unit was removed from the disk shelf.
FC target port commands exceeded	Critical, Warning, Info	EMS	The number of outstanding commands on the physical FC target port exceeds the supported limit.
Giveback of storage pool failed	Critical, Warning, Info	EMS	This event occurs during the relocation of a storage pool (aggregate) as part of a storage failover (SFO) giveback, when the destination node cannot reach the object stores.
HA interconnect down	Critical, Warning, Info	EMS	The high-availability (HA) interconnect is down. Risk of service outage when failover is not available.
InstanceDown	Critical	Metric	The monitored instance is unreachable and may be down or experiencing network issues.
LUN destroyed	Critical, Warning, Info	EMS	A LUN was destroyed and is no longer accessible.
LUN offline	Critical, Warning, Info	EMS	A LUN was taken offline manually.
Main unit fan failed	Critical, Warning, Info	EMS	One or more main unit fans have failed. The system remains operational.
Main unit fan in warning state	Critical, Warning, Info	EMS	One or more main unit cooling fans are in a warning state.
Max sessions per user exceeded	Critical, Warning, Info	EMS	You have exceeded the maximum number of sessions allowed per user over a TCP connection.
Max times open per file exceeded	Critical, Warning, Info	EMS	You have exceeded the maximum number of times that you can open the file over a TCP connection.

Alert name	Severity	Type	Description
MetroCluster automatic unplanned switchover disabled	Critical, Warning, Info	EMS	Automatic unplanned switchover capability has been disabled on this cluster.
MetroCluster monitoring	Critical, Warning, Info	EMS	The system health monitor alert is detected.
NFSv4 store pool exhausted	Critical, Warning, Info	EMS	A NFSv4 store pool has been exhausted.
NetBIOS name conflict	Critical, Warning, Info	EMS	The NetBIOS Name Service has received a negative response to a name registration request, from a remote machine.
No registered scan engine	Critical, Warning, Info	EMS	The antivirus connector notified ONTAP that it does not have a registered scan engine.
No Vscan connection	Critical, Warning, Info	EMS	ONTAP has no Vscan connection to service virus scan requests. This might cause data unavailability if the scan-mandatory option is enabled.
Non-responsive antivirus server	Critical, Warning, Info	EMS	ONTAP detected a non-responsive antivirus server and forcibly closed its Vscan connection.
Nonexistent admin share	Critical, Warning, Info	EMS	Vscan issue: a client has attempted to connect to a nonexistent ONTAP_ADMIN\$ share.
NVRAM battery low	Critical, Warning, Info	EMS	The NVRAM battery capacity is critically low. There might be a potential data loss if the battery runs out of power.
NVMe namespace destroyed	Critical, Warning, Info	EMS	An NVMe namespace was destroyed and is no longer accessible.
NVMe namespace offline	Critical, Warning, Info	EMS	An NVMe namespace was taken offline manually.
NVMe namespace online	Critical, Warning, Info	EMS	An NVMe namespace was brought back online.
NVMe-oF license grace period active	Critical, Warning, Info	EMS	This event occurs on a daily basis when the NVMe over Fabrics (NVMe-oF) protocol is in use and the grace period of the license is active.
NVMe-oF license grace period expired	Critical, Warning, Info	EMS	The NVMe over Fabrics (NVMe-oF) license grace period is over and the NVMe-oF functionality is disabled.

Alert name	Severity	Type	Description
NVMe-oF license grace period start	Critical, Warning, Info	EMS	The NVMe over Fabrics (NVMe-oF) configuration was detected during the upgrade to ONTAP 9.5 software.
Object store host unresolvable	Critical, Warning, Info	EMS	The object store server host name cannot be resolved to an IP address.
Object store intercluster LIF down	Critical, Warning, Info	EMS	The object-store client cannot find an operational LIF to communicate with the object store server.
Object store signature mismatch	Critical, Warning, Info	EMS	The request signature sent to the object store server does not match the signature calculated by the client.
Port Status Down	Critical	EMS	This Ethernet port is down. Traffic on that link may be affected.
REaddir timeout	Critical, Warning, Info	EMS	A REaddir file operation has exceeded the timeout that it is allowed to run in WAFL.
Relocation of storage pool failed	Critical, Warning, Info	EMS	This event occurs during the relocation of an storage pool (aggregate), when the destination node cannot reach the object stores.
SAN "active-active" state changed	Critical, Warning, Info	EMS	The SAN pathing is no longer symmetric.
Service Processor heartbeat missed	Critical, Warning, Info	EMS	ONTAP is not receiving an expected heartbeat signal from the Service Processor (SP).
Service Processor heartbeat stopped	Critical, Warning, Info	EMS	ONTAP is no longer receiving heartbeats from the Service Processor (SP).
Service Processor not configured	Critical, Warning, Info	EMS	This event occurs on a weekly basis, to remind you to configure the Service Processor (SP).
Service Processor offline	Critical, Warning, Info	EMS	The Service Processor (SP) is offline.
SFP in FC target adapter receiving low power	Critical, Warning, Info	EMS	This alert occurs when the power received (RX) by a small form-factor pluggable transceiver (SFP in FC target) is at a level below the defined threshold.
SFP in FC target adapter transmitting low power	Critical, Warning, Info	EMS	This alert occurs when the power transmitted (TX) by a small form-factor pluggable transceiver (SFP in FC target) is at a level below the defined threshold.

Alert name	Severity	Type	Description
Shadow copy failed	Critical, Warning, Info	EMS	A Volume Shadow Copy Service (VSS), a Microsoft Server backup and restore service operation, has failed.
Shelf fan failed	Critical, Warning, Info	EMS	The indicated cooling fan or fan module of the shelf has failed.
SnapMirror lag time is high	Critical	Metric	The SnapMirror relationship is more than one hour behind its expected update schedule.
Storage Failover Interconnect One Or More Links Down	Warning	EMS	One or more HA interconnect links to the partner node are down. Failover redundancy is reduced.
Storage switch power supplies failed	Critical, Warning, Info	EMS	There is a missing power supply in the cluster switch. Redundancy is reduced.
Storage VM stop succeeded	Critical, Warning, Info	EMS	The storage VM was stopped successfully.
SVM Configuration Replication License Invalid	Warning	EMS	The SVM-DR configuration replication license is missing, expired, or not applied
System cannot operate due to main unit fan failure	Critical, Warning, Info	EMS	One or more main unit fans have failed, disrupting system operation. This might lead to a potential data loss.
Too many CIFS authentication	Critical, Warning, Info	EMS	Many authentication negotiations have occurred simultaneously. There are 256 incomplete new session requests from this client.
Unassigned disks	Critical, Warning, Info	EMS	System has unassigned disks - capacity is being wasted.
Volume state offline	Informational	Metric	The volume has been taken offline.
Volume restricted	Critical, Warning, Info	EMS	This event indicates that a flexible volume is made restricted.
Virus detected	Critical, Warning, Info	EMS	A Vscan server reported that a file might be infected with a virus.

Capacity alerts

These alerts indicate storage consumption or capacity pressure.

Alert name	Severity	Type	Description
FabricPool mirror replication resync completed	Critical, Warning, Info	EMS	The FabricPool mirror resync process completed successfully from the primary object store to the mirror object store.
FabricPool space usage limit nearly reached	Critical, Warning, Info	EMS	The total cluster-wide FabricPool space usage of object stores from capacity-licensed providers has nearly reached the licensed limit.
FabricPool space usage limit reached	Critical, Warning, Info	EMS	The total cluster-wide FabricPool space usage of object stores from capacity-licensed providers has reached the license limit.
Node root volume space low	Critical, Warning, Info	EMS	The system has detected that the root volume is dangerously low on space.
QoS monitor memory maxed out	Critical, Warning, Info	EMS	A QoS subsystem's dynamic memory has reached its limit for the current platform hardware.
Volume automatic resizing succeeded	Critical, Warning, Info	EMS	The volume was automatically resized because automatic volume growth is enabled.
Volume full	Critical	Metric	The volume is over 90% full. Free up space or add capacity to prevent write failures.

Configuration alerts

These alerts point to configuration changes or inventory updates.

Alert name	Severity	Type	Description
Disk shelf power supply discovered	Critical, Warning, Info	EMS	A power supply unit was added to the disk shelf.
Volume created	Info	Metric	A volume was created.
Volume deleted	Warning	Metric	A volume was deleted.
Volume modified	Info	Metric	A volume was modified.
WAFL Consistency Check Overdue	Warning	EMS	WAFL consistency checks on this aggregate exceeded the hourly limit.

Performance alerts

These alerts point to latency or node performance issues.

Alert name	Severity	Type	Description
Node NFS latency is high	Critical	Metric	NFS operations on this node are experiencing high latency (>5000 us).

Alert name	Severity	Type	Description
Node panic	Critical, Warning, Info	EMS	The node panicked and was restarted.

Protection alerts

These alerts affect replication, failover, or recovery.

Alert name	Severity	Type	Description
Fanout SnapMirror relationship common snapshot deleted	Critical, Warning, Info	EMS	An older Snapshot copy is deleted as part of a SnapMirror Synchronous resynchronize or update operation.
ONTAP Mediator added	Critical, Warning, Info	EMS	The ONTAP Mediator was successfully added to this cluster.
ONTAP Mediator CA certificate expired	Critical, Warning, Info	EMS	The ONTAP Mediator certificate authority (CA) certificate has expired.
ONTAP Mediator CA certificate expiring	Critical, Warning, Info	EMS	The ONTAP Mediator certificate authority (CA) certificate is due to expire within the next 30 days.
ONTAP Mediator client certificate expired	Critical, Warning, Info	EMS	The ONTAP Mediator client certificate has expired.
ONTAP Mediator client certificate expiring	Critical, Warning, Info	EMS	The ONTAP Mediator client certificate is due to expire within the next 30 days.
ONTAP Mediator not accessible	Critical, Warning, Info	EMS	The ONTAP Mediator is not accessible, either because it has been repurposed or the Mediator package is no longer installed.
ONTAP Mediator removed	Critical, Warning, Info	EMS	The ONTAP Mediator was successfully removed from this cluster.
ONTAP Mediator unreachable	Critical, Warning, Info	EMS	The ONTAP Mediator is unreachable, which means SnapMirror failover is not possible until connectivity is restored.
ONTAP Mediator server certificate expired	Critical, Warning, Info	EMS	The ONTAP Mediator server certificate has expired.
ONTAP Mediator server certificate expiring	Critical, Warning, Info	EMS	The ONTAP Mediator server certificate is due to expire within the next 30 days.
SnapMirror active sync relationship out of sync	Critical, Warning, Info	EMS	The SnapMirror synchronous relationship moved from in-sync to out-of-sync. Data protection is affected.

Alert name	Severity	Type	Description
SnapMirror active sync automatic unplanned failover completed	Critical, Warning, Info	EMS	The SnapMirror Business Continuity (SMBC) automatic unplanned failover operation completed.
SnapMirror active sync automatic unplanned failover failed	Critical, Warning, Info	EMS	The SnapMirror Business Continuity (SMBC) automatic unplanned failover operation failed.
SnapMirror active sync planned failover completed	Critical, Warning, Info	EMS	The SnapMirror Business Continuity (SMBC) planned failover operation completed.
SnapMirror active sync planned failover failed	Critical, Warning, Info	EMS	The SnapMirror Business Continuity (SMBC) planned failover operation failed.
SnapMirror relationship common snapshot failed	Critical, Warning, Info	EMS	There is a failure in creating a common Snapshot copy.
SnapMirror relationship initialization failed	Critical, Warning, Info	EMS	The SnapMirror initialize command fails and no more retries will be attempted.
SnapMirror relationship out of sync	Critical, Warning, Info	EMS	The SnapMirror synchronous relationship moved from in-sync to out-of-sync. Data protection is affected.
SnapMirror relationship resync attempt failed	Critical, Warning, Info	EMS	A SnapMirror synchronization attempt between the source and destination volumes failed.
SnapMirror relationship snapshot is not replicated	Critical, Warning, Info	EMS	The Snapshot copy for SnapMirror Synchronous relationship is not successfully replicated.

Security alerts

These alerts point to threats or unauthorized access.

Alert name	Severity	Type	Description
Ransomware activity detected	Critical, Warning, Info	EMS	To protect the data from the detected ransomware, a Snapshot copy has been taken that can be used to restore original data.
Storage VM anti-ransomware monitoring	Critical, Warning, Info	EMS	The anti-ransomware state of the Storage VM has changed.
Unauthorized user access to admin share	Critical, Warning, Info	EMS	A client tried to connect to the privileged ONTAP_ADMIN\$ share, but the logged-in user isn't part of any active Vscan scanner pool on this SVM.

Alert name	Severity	Type	Description
Volume anti-ransomware monitoring	Critical, Warning, Info	EMS	The anti-ransomware state of a volume is changed.

Cluster security compliance categories in NetApp Console local deployment

Use this reference to review the cluster security compliance categories shown in NetApp Console local deployment, including the recommended value and whether each category affects overall compliance status.

These categories are based on ONTAP security posture checks.

Category	What the category checks	Recommended value	Affects overall compliance
Global FIPS	Whether FIPS 140-2 mode is enabled. When enabled, TLSv1 and SSLv3 are disabled and only TLSv1.1 and TLSv1.2 are allowed.	Enabled	Yes
Telnet	Whether Telnet access is enabled. SSH is the recommended secure remote access method.	Disabled	Yes
Insecure SSH settings	Whether SSH is configured with insecure ciphers, such as ciphers that begin with <code>cbc</code> .	No	Yes
Login banner	Whether a login banner is configured for system access.	Enabled	Yes
Cluster peering	Whether communication between peered clusters is encrypted. Encryption must be enabled on both source and destination clusters.	Encrypted	Yes
Network Time Protocol	Whether one or more NTP servers are configured for the cluster. NetApp recommends at least three NTP servers for resilience.	Configured	Yes
OCSP	Whether Online Certificate Status Protocol validation is enabled for SSL/TLS certificate validation.	Enabled	No
Remote audit logging	Whether log forwarding (syslog) is encrypted.	Encrypted	Yes
AutoSupport HTTPS transport	Whether HTTPS is used as the default transport protocol for AutoSupport messages.	Enabled	Yes
Default admin user	Whether the built-in default admin account is disabled.	Disabled	Yes
SAML users	Whether SAML is configured for single sign-on and MFA-capable authentication.	No	No
Active Directory users	Whether Active Directory authentication is configured for cluster access.	No	No
LDAP users	Whether LDAP authentication is configured for cluster access.	No	No

Category	What the category checks	Recommended value	Affects overall compliance
Certificate users	Whether certificate-based users are configured for cluster login.	No	No
Local users	Whether local users are configured for cluster login.	No	No
Remote shell	Whether RSH is enabled. SSH is preferred for secure remote access.	Disabled	Yes
MD5 in use	Whether ONTAP user accounts still use MD5 hashing instead of stronger hashes such as SHA-512.	No	Yes
Certificate issuer type	The certificate issuer type used by the cluster.	CA-Signed	No

Storage VM security compliance categories in NetApp Console local deployment

Use this reference to review the storage VM (SVM) security compliance categories shown in NetApp Console local deployment, including the recommended value and whether each category affects overall compliance status.

These categories are based on ONTAP security posture checks.

Category	What the category checks	Recommended value	Affects overall compliance
Audit log	Whether SVM audit logging is enabled.	Enabled	Yes
Insecure SSH settings	Whether SSH is configured with insecure ciphers, such as ciphers that begin with <code>cbc</code> .	No	Yes
Login banner	Whether a login banner is configured for users who access SVMs.	Enabled	Yes
LDAP encryption	Whether LDAP encryption is enabled.	Enabled	No
NTLM authentication	Whether NTLM authentication is enabled.	Enabled	No
LDAP payload signing	Whether LDAP payload signing is enabled.	Enabled	No
CHAP settings	Whether CHAP is enabled.	Enabled	No
Kerberos V5	Whether Kerberos V5 authentication is enabled.	Enabled	No
NIS authentication	Whether NIS authentication is configured.	Disabled	No
FPolicy status active	Whether FPolicy is created and active.	Yes	No
SMB encryption enabled	Whether SMB signing and sealing are enabled.	Yes	No
SMB signing enabled	Whether SMB signing is enabled.	Yes	No

Knowledge and support

Register for support in NetApp Console local deployment

NetApp Console local deployment information for this task.

Support registration is required to receive technical support specific to the NetApp Console and its storage solutions and data services. Support registration is also required to enable key workflows for Cloud Volumes ONTAP systems.

Registering for support does not enable NetApp support for a cloud provider file service. For technical support related to a cloud provider file service, its infrastructure, or any solution using the service, refer to "Getting help" in the documentation for that product.

- [Amazon FSx for ONTAP](#)
- [Azure NetApp Files](#)
- [Google Cloud NetApp Volumes](#)

Support registration overview

There are two forms of registration to activate support entitlement:

- Registering your NetApp Console account serial number (your 20 digit 960xxxxxxx serial number located on the Support Resources page in the Console).

This serves as your single support subscription ID for any service within the Console. Each Console account must be registered.

- Registering the Cloud Volumes ONTAP serial numbers associated with a subscription in your cloud provider's marketplace (these are 20 digit 909201xxxxxxx serial numbers).

These serial numbers are commonly referred to as *PAYGO serial numbers* and get generated by the NetApp Console at the time of Cloud Volumes ONTAP deployment.

Registering both types of serial numbers enables capabilities like opening support tickets and automatic case generation. Registration is completed by adding NetApp Support Site (NSS) accounts to the Console as described below.

Register NetApp Console for NetApp support

To register for support and activate support entitlement, one user in your NetApp Console account must associate a NetApp Support Site account with their Console login. How you register for NetApp support depends on whether you already have a NetApp Support Site (NSS) account.

Existing customer with an NSS account

If you're a NetApp customer with an NSS account, you simply need to register for support through the Console.

Steps

1. Select **Administration > Credentials**.
2. Select **User Credentials**.

3. Select **Add NSS credentials** and follow the NetApp Support Site (NSS) authentication prompt.
4. To confirm that the registration process was successful, select the Help icon, and select **Support**.

The **Resources** page should show that your Console account is registered for support.

Note that other Console users will not see this same support registration status if they have not associated a NetApp Support Site account with their login. However, that doesn't mean that your account is not registered for support. As long as one user in the organization has followed these steps, then your account has been registered.

Existing customer but no NSS account

If you're an existing NetApp customer with existing licenses and serial numbers but *no* NSS account, you need to create an NSS account and associate it with your Console login.

Steps


1. Create a NetApp Support Site account by completing the [NetApp Support Site User Registration form](#)
 - a. Be sure to select the appropriate User Level, which is typically **NetApp Customer/End User**.
 - b. Be sure to copy the Console account serial number (960xxxx) used above for the serial number field. This will speed up the account processing.
2. Associate your new NSS account with your Console login by completing the steps under [Existing customer with an NSS account](#).

Brand new to NetApp

If you are brand new to NetApp and you don't have an NSS account, follow each step below.

Steps

1. In the upper right of the Console, select the Help icon, and select **Support**.
2. Locate your account ID serial number from the Support Registration page.



3. Navigate to [NetApp's support registration site](#) and select **I am not a registered NetApp Customer**.
4. Fill out the mandatory fields (those with red asterisks).
5. In the **Product Line** field, select **Cloud Manager** and then select your applicable billing provider.
6. Copy your account serial number from step 2 above, complete the security check, and then confirm that you read NetApp's Global Data Privacy Policy.

An email is immediately sent to the mailbox provided to finalize this secure transaction. Be sure to check your spam folders if the validation email doesn't arrive in few minutes.

7. Confirm the action from within the email.

Confirming submits your request to NetApp and recommends that you create a NetApp Support Site account.

8. Create a NetApp Support Site account by completing the [NetApp Support Site User Registration form](#)
 - a. Be sure to select the appropriate User Level, which is typically **NetApp Customer/End User**.
 - b. Be sure to copy the account serial number (960xxxx) used above for the serial number field. This will speed up processing.

After you finish

NetApp should reach out to you during this process. This is a one-time onboarding exercise for new users.

Once you have your NetApp Support Site account, associate the account with your Console login by completing the steps under [Existing customer with an NSS account](#).

Associate NSS credentials for Cloud Volumes ONTAP support

Associating NetApp Support Site credentials with your Console account is required to enable the following key workflows for Cloud Volumes ONTAP:

- Registering pay-as-you-go Cloud Volumes ONTAP systems for support

Providing your NSS account is required to activate support for your system and to gain access to NetApp technical support resources.

- Deploying Cloud Volumes ONTAP when you bring your own license (BYOL)

Providing your NSS account is required so that the Console can upload your license key and to enable the subscription for the term that you purchased. This includes automatic updates for term renewals.

- Upgrading Cloud Volumes ONTAP software to the latest release


Associating NSS credentials with your NetApp Console account is different than the NSS account that is associated with a Console user login.

These NSS credentials are associated with your specific Console account ID. Users who belong to the Console organization can access these credentials from **Support > NSS Management**.

- If you have a customer-level account, you can add one or more NSS accounts.
- If you have a partner or reseller account, you can add one or more NSS accounts, but they can't be added alongside customer-level accounts.

Steps

1. In the upper right of the Console, select the Help icon, and select **Support**.



2. Select **NSS Management > Add NSS Account**.
3. When you're prompted, select **Continue** to be redirected to a Microsoft login page.

NetApp uses Microsoft Entra ID as the identity provider for authentication services specific to support and licensing.

4. At the login page, provide your NetApp Support Site registered email address and password to perform the authentication process.

These actions enable the Console to use your NSS account for things like license downloads, software upgrade verification, and future support registrations.

Note the following:

- The NSS account must be a customer-level account (not a guest or temp account). You can have multiple customer-level NSS accounts.
- There can be only one NSS account if that account is a partner-level account. If you try to add customer-level NSS accounts and a partner-level account exists, you'll get the following error message:

"The NSS customer type is not allowed for this account as there are already NSS Users of different type."

The same is true if you have pre-existing customer-level NSS accounts and try to add a partner-level account.

- Upon successful login, NetApp will store the NSS user name.

This is a system-generated ID that maps to your email. On the **NSS Management** page, you can display your email from the  menu.

- If you ever need to refresh your login credential tokens, there is also an **Update Credentials** option in the  menu.

Using this option prompts you to log in again. Note that the token for these accounts expire after 90 days. A notification will be posted to alert you of this.

Get help with NetApp Console local deployment

NetApp Console local deployment information for this task.

NetApp provides support for NetApp Console and its cloud services in a variety of ways. Extensive free self-support options are available 24/7, such as knowledge base (KB) articles and a community forum. Your support registration includes remote technical support via web ticketing.

Get support for a cloud provider file service

For technical support related to a cloud provider file service, its infrastructure, or any solution using the service, refer to the documentation for that product.

- [Amazon FSx for ONTAP](#)
- [Azure NetApp Files](#)
- [Google Cloud NetApp Volumes](#)

To receive technical support specific to NetApp and its storage solutions and data services, use the support options described below.

Use self-support options

These options are available for free, 24 hours a day, 7 days a week:

- Documentation

The NetApp Console documentation that you're currently viewing.

- [Knowledge base](#)

Search through the NetApp knowledge base to find helpful articles to troubleshoot issues.

- [Communities](#)

Join the NetApp Console community to follow ongoing discussions or create new ones.

Create a case with NetApp support

In addition to the self-support options above, you can work with a NetApp Support specialist to resolve any issues after you activate support.

Before you get started

- To use the **Create a Case** capability, you must first associate your NetApp Support Site credentials with your Console login. [Learn how to manage credentials associated with your Console login.](#)
- If you're opening a case for an ONTAP system that has a serial number, then your NSS account must be associated with the serial number for that system.

Steps

1. In NetApp Console, select **Help > Support**.
2. On the **Resources** page, choose one of the available options under Technical Support:

- a. Select **Call Us** if you'd like to speak with someone on the phone. You'll be directed to a page on netapp.com that lists the phone numbers that you can call.
- b. Select **Create a Case** to open a ticket with a NetApp Support specialist:
 - **Service:** Select the service that the issue is associated with. For example, **NetApp Console** when specific to a technical support issue with workflows or functionality within the Console.
 - **System:** If applicable to storage, select **Cloud Volumes ONTAP** or **On-Prem** and then the associated working environment.

The list of systems are within scope of the Console organization, and Console agent you have selected in the top banner.

- **Case Priority:** Choose the priority for the case, which can be Low, Medium, High, or Critical.

To learn more details about these priorities, hover your mouse over the information icon next to the field name.

- **Issue Description:** Provide a detailed description of your problem, including any applicable error messages or troubleshooting steps that you performed.
- **Additional Email Addresses:** Enter additional email addresses if you'd like to make someone else aware of this issue.
- **Attachment (Optional):** Upload up to five attachments, one at a time.

Attachments are limited to 25 MB per file. The following file extensions are supported: txt, log, pdf, jpg/jpeg, rtf, doc/docx, xls/xlsx, and csv.

ntapitdemo
NetApp Support Site Account

Service

Select

Working Enviroment

Select

Case Priority

Low - General guidance

Issue Description

Provide detailed description of problem, applicable error messages and troubleshooting steps taken.

Additional Email Addresses (Optional)

Type here

Attachment (Optional)

No files selected

Upload

After you finish

A pop-up will appear with your support case number. A NetApp Support specialist will review your case and get back to you soon.

For a history of your support cases, you can select **Settings > Timeline** and look for actions named "create support case." A button to the far right lets you expand the action to see details.

It's possible that you might encounter the following error message when trying to create a case:

"You are not authorized to Create a Case against the selected service"

This error could mean that the NSS account and the company of record it's associated with is not the same company of record for the NetApp Console account serial number (ie. 960xxxx) or the working environment serial number. You can seek assistance using one of the following options:

- Submit a non-technical case at <https://mysupport.netapp.com/site/help>

Manage your support cases

You can view and manage active and resolved support cases directly from the Console. You can manage the

cases associated with your NSS account and with your company.

Note the following:

- The case management dashboard at the top of the page offers two views:
 - The view on the left shows the total cases opened in the past 3 months by the user NSS account you provided.
 - The view on the right shows the total cases opened in the past 3 months at your company level based on your user NSS account.

The results in the table reflect the cases related to the view that you selected.

- You can add or remove columns of interest and you can filter the contents of columns like Priority and Status. Other columns provide just sorting capabilities.

View the steps below for more details.

- At a per-case level, we offer the ability to update case notes or close a case that is not already in Closed or Pending Closed status.

Steps

1. In the NetApp Console, select **Help > Support**.
2. Select **Case Management** and if you're prompted, add your NSS account to the Console.

The **Case management** page shows open cases related to the NSS account that is associated with your Console user account. This is the same NSS account that appears at the top of the **NSS management** page.

3. Optionally modify the information that displays in the table:
 - Under **Organization's cases**, select **View** to view all cases associated with your company.
 - Modify the date range by choosing an exact date range or by choosing a different time frame.
 - Filter the contents of the columns.
 - Change the columns that appear in the table by selecting  and then choosing the columns that you'd like to display.
4. Manage an existing case by selecting  and selecting one of the available options:
 - **View case**: View full details about a specific case.
 - **Update case notes**: Provide additional details about your problem or select **Upload files** to attach up to a maximum of five files.

Attachments are limited to 25 MB per file. The following file extensions are supported: txt, log, pdf, jpg/jpeg, rtf, doc/docx, xls/xlsx, and csv.

- **Close case**: Provide details about why you're closing the case and select **Close case**.

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.