



# **NetApp Workload Factory for EDA documentation**

EDA workloads

NetApp  
March 02, 2026

# Table of Contents

- NetApp Workload Factory for EDA documentation ..... 1
- Release notes ..... 2
  - What's new with NetApp Workload Factory for EDA ..... 2
    - 1 March 2026 ..... 2
    - 1 February 2026 ..... 2
    - 4 January 2026 ..... 3
    - 5 October 2025 ..... 4
    - 16 June 2025 ..... 4
    - 4 May 2025 ..... 4
    - 1 December 2024 ..... 4
  - Known limitations of NetApp Workload Factory for EDA ..... 4
    - Operator permissions required ..... 4
- Get started ..... 5
  - Learn about NetApp Workload Factory for EDA ..... 5
  - Quick start for EDA ..... 5
- EDA storage overview ..... 7
  - Use the overview dashboard in EDA workloads ..... 7
    - Overview ..... 7
    - Requirements ..... 7
    - Filter the dashboard ..... 8
    - Cluster health status ..... 8
    - Clusters table ..... 9
- FSx latency analysis ..... 11
  - Monitor volume latency in EDA workloads ..... 11
    - Overview ..... 11
    - Requirements ..... 11
    - Understanding alerts ..... 12
    - Configure latency thresholds ..... 12
    - View latency events ..... 13
    - Understanding basic analysis ..... 14
    - Manage latency configuration ..... 14
    - Best practices ..... 15
- Project usage analysis ..... 16
  - Use the project dashboard ..... 16
    - Setup the dashboard ..... 16
    - View volume details ..... 18
- Use Perforce integration ..... 20
  - Learn about Perforce integration in NetApp Workload Factory for EDA ..... 20
    - What is CI/CD? ..... 20
    - CI/CD features ..... 20
    - Projects and workspaces in EDA ..... 20
    - Automation with Workload Factory Codebox ..... 21
    - Cost ..... 21

|                                                             |    |
|-------------------------------------------------------------|----|
| Licensing                                                   | 21 |
| Regions                                                     | 21 |
| Getting help                                                | 21 |
| EDA requirements                                            | 21 |
| Manage EDA projects                                         | 22 |
| Create a project                                            | 22 |
| View existing projects                                      | 23 |
| Edit a project                                              | 24 |
| View a project's workspaces                                 | 24 |
| Delete a project                                            | 24 |
| Manage versions of NetApp Workload Factory for EDA projects | 25 |
| Create a snapshot of a project                              | 25 |
| Create a clone of a project                                 | 25 |
| Create an EDA workspace                                     | 26 |
| Automate EDA workload tasks with Codebox                    | 26 |
| Integrate EDA with Perforce                                 | 26 |
| Knowledge and support                                       | 28 |
| Register for support for NetApp Workload Factory for EDA    | 28 |
| Support registration overview                               | 28 |
| Register your account for NetApp support                    | 28 |
| Get help with Workload Factory for EDA                      | 30 |
| Get support for FSx for ONTAP                               | 30 |
| Use self-support options                                    | 30 |
| Create a case with NetApp support                           | 30 |
| Manage your support cases (Preview)                         | 33 |
| Legal notices for NetApp Workload Factory for EDA           | 36 |
| Copyright                                                   | 36 |
| Trademarks                                                  | 36 |
| Patents                                                     | 36 |
| Privacy policy                                              | 36 |
| Open source                                                 | 36 |

# NetApp Workload Factory for EDA documentation

# Release notes

## What's new with NetApp Workload Factory for EDA

Learn what's new with the EDA capability of Workload Factory.

### 1 March 2026

#### Overview dashboard for EDA storage management

As an IT administrator managing dozens of FSx for ONTAP file systems, you can use the overview dashboard to answer common operational questions about cluster health, utilization, and volume placement.

The overview dashboard includes:

- **Cluster health status:** Information summarized at the top of the dashboard that highlight performance, capacity, availability and protection, and security events across your clusters.
- **Clusters table:** A searchable, filterable, and paginated table showing SSD capacity, capacity pool, and throughput metrics per cluster, with support for column customization and CSV export.

The previous main dashboard is now accessible in the **Project** tab. EDA tabs are now ordered: Overview, Latency , Project, CI/CD.

[Learn more about the overview dashboard.](#)

#### Automated latency analysis

Latency analysis now includes automated basic analysis. When latency events are detected, the system automatically analyzes ONTAP QoS delay center metrics to identify the source of latency issues.

Basic analysis evaluates QoS policies, disk performance, network issues, and cluster configurations to provide immediate insights in the Details column of the latency events table.

[Learn more about basic latency analysis.](#)

#### New Ask Me bookmark available in Workload Factory

We added a new Ask Me bookmark on every screen in Workload Factory console. This improvement makes it easier and faster for you to access Ask Me whenever you need help. Ask Me is our AI assistant, enabling you to ask questions about your own workload environments, get personalized insights directly from your environment, and refer to previous conversations.

You can open Ask Me from any page by clicking the new bookmark. It launches Ask Me in a side panel without interrupting your current work and offers quick explanations and recommendations related to what you're currently doing.

[Learn more about Ask Me](#)

### 1 February 2026

## Enhanced dashboard filtering with customizable tags

You can now configure up to five custom filters on the EDA dashboard based on your AWS tags. Each custom filter includes a label name, AWS tag key, and a selection type (single or multi-selection).

Multi-selection allows you to select multiple values simultaneously, while single selection restricts you to one value at a time. Custom filters appear in the order you configure them, making it easy to organize your most frequently used filters.

If you don't configure custom filters, the default filters (file system, volume type, and time range) remain available so you can continue to view and interact with your dashboards.

[Learn more about configuring custom filters.](#)

## Volume details view for granular performance analysis

The dashboard now offers two viewing modes: Total view and Volume view. The Total view displays aggregated metrics across all volumes, while the Volume view shows individual performance for the top 10 volumes over time.

In Volume view, interactive hover tooltips provide detailed information for each volume, including volume name, metrics, and time-specific values. When the same volumes appear across multiple components, consistent color coding makes it easier to track specific volumes across different metrics.

[Learn more about viewing volume details.](#)

## Latency analysis for proactive performance monitoring

Latency analysis enables you to monitor volume read and write latency across your FSx for ONTAP file systems. You can configure customizable warning and critical event thresholds to proactively identify performance bottlenecks before they impact your EDA workloads.

The latency events table displays all warning and critical events enabling you to monitor volume performance and identify volumes that require optimization.

This feature requires AWS credentials and is accessible from the Latency menu in the EDA dashboard.

[Learn more about latency analysis.](#)

## 4 January 2026

### NetApp Workload Factory for Builders now NetApp Workload Factory for EDA

Workload Factory for Builders is now Workload Factory for EDA. The name change reflects the focus on electronic design automation (EDA) workloads.

Workload Factory for EDA helps you optimize FSx for ONTAP across multiple file systems. You can optimize for performance and reduce operational costs by automating storage parameters, analyzing performance constraints, and get insights on EDA projects.

Workload Factory for EDA is designed to integrate with your Infrastructure as Code (IaC) frameworks.

### Ask me AI assistant home page integration

The Workload Factory console home page embeds the Ask me AI assistant, enabling you to ask questions about your own storage estate, get personalized insights directly from your environment, and refer to previous

conversations. You can interact with Ask me to understand your workloads, troubleshoot issues, and learn more about Workload Factory — all without leaving the console.

## 5 October 2025

### BlueXP workload factory now NetApp Workload Factory

BlueXP has been renamed and redesigned to better reflect the role it has in managing your data infrastructure. As a result, BlueXP workload factory has been renamed to NetApp Workload Factory.

## 16 June 2025

### Clone support

You can now clone a project in BlueXP Workload Factory for Builders. When you clone a project, Builders creates a new project from a snapshot, with the same configuration as the original. Cloning is useful for quickly creating similar projects or for testing purposes. You can mount the new project clone by following the instructions in Builders.

[Manage versions of BlueXP Workload Factory for Builders projects](#)

## 4 May 2025

### Updated permissions terminology

The Workload Factory user interface and documentation now use "read-only" to refer to read permissions and "read/write" to refer to automate permissions.

## 1 December 2024

### Builders workload initial release

BlueXP Workload Factory for Builders simplifies software version consumption and access, eliminating the need for custom tools or scripts. It enables you to consume software versions as instant clones integrated with Perforce Helix Core as a convenient workspace for your development processes, saving time and resources.

The initial release includes the capability to manage projects and workspaces, and automate actions with Codebox. You can also integrate Builders with Perforce Helix Core, so that you can manage different versions of each project and switch between them quickly.

## Known limitations of NetApp Workload Factory for EDA

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.

### Operator permissions required

NetApp Workload Factory for EDA requires Operator permissions to function correctly.

# Get started

## Learn about NetApp Workload Factory for EDA

Workload Factory for EDA helps you optimize FSx for ONTAP across multiple file systems. You can optimize for performance and reduce operational costs by automating storage parameters, analyzing performance constraints, and get insights on EDA projects.

It is designed to integrate with your Infrastructure as Code (IaC) frameworks.

Workload Factory for EDA provides dashboards and storage automation frameworks and utilities that help you manage multiple FSx for ONTAP file systems.

It offers the following:

- [overview dashboard](#): provides a centralized cluster view to help you monitor FSx for ONTAP cluster health, capacity, and throughput, determine where to place new volumes, and make informed decisions about scaling capacity and throughput.
- [FSx latency analysis](#): proactively monitors volume read and write latency performance with configurable warning and critical event thresholds. Includes automated basic analysis using ONTAP QoS delay center metrics.
- [Project dashboard](#): provides a centralized view of storage consumption across your FSx for ONTAP file systems to help you plan, assess costs, and collect information on storage usage across your projects.
- [CI/CD](#): streamlines and reduces software build time by leveraging FSx for ONTAP volume clone capabilities.

## Quick start for EDA

Get started creating an EDA project. Administrators and team leads can use EDA to administer projects and workspaces for teams of developers.

1

### Log in to Workload Factory

You'll need to [set up an account with Workload Factory](#) and log in using one of the [console experiences](#).

2

### Add credentials and permissions

Choose the permission policies to meet your needs.

If you choose not to grant permissions, you can start using Workload Factory for EDA to copy partially completed code samples.

If you choose to grant permissions, you'll need to add credentials to an account manually that includes selecting workload capabilities, such as EDA and AI, and creating the IAM policies for the required permissions.

[Learn how to add credentials and permissions.](#)

**3****Set up your environment to meet EDA requirements**

You'll need a deployed and discovered FSx for ONTAP file system that contains at least one volume that has been configured as an NFS share.

[Learn more about EDA requirements.](#)

**4****Enable the overview dashboard**

Enable the overview dashboard to monitor the health, capacity, and throughput of your FSx for ONTAP clusters. Accept the consent prompt to allow Workload Factory to begin collecting CloudWatch metrics.

[Learn how to enable the overview dashboard.](#)

**5****Configure the project dashboard**

Configure the project dashboard to monitor system health, performance, and storage usage, enabling you to optimize your EDA environment effectively.

[Learn how to configure the project dashboard.](#)

**6****Configure FSx latency analysis**

Set up latency thresholds to proactively identify and resolve performance issues before they impact simulation run-time. Configure warning and critical alerts, and leverage automated basic analysis.

[Learn how to configure latency monitoring.](#)

**7****Integrate with other services**

Integrate EDA with other services such as a service portal or Perforce Helix Core using the workload Factory REST API.

[Learn how to integrate EDA with Perforce.](#)

# EDA storage overview

## Use the overview dashboard in EDA workloads

As an IT administrator managing EDA workloads across dozens of FSx for ONTAP file systems, you can use the overview dashboard to answer common daily operational questions about the state of your clusters. Use it to determine where to place new volumes or jobs, understand cluster usage at scale, and make informed decisions about scaling capacity and throughput up or down.

### Overview

The overview dashboard provides a centralized cluster view that helps you quickly understand cluster usage across capacity, throughput, and IOPS, and determine where to place new volumes or move existing volumes and SVMs.

The dashboard includes:

- **Cluster health status:** Information summarized at the top of the dashboard that highlights latency, availability and protection, and security events and capacity recommendations across your clusters.
- **Clusters table:** A detailed, searchable table showing usage and performance metrics for each FSx for ONTAP cluster, with support for filtering, sorting, pagination, and CSV export.

### Requirements

Before using the overview dashboard, ensure you meet the following requirements:

#### AWS credentials with view permissions

You must have AWS credentials configured in Workload Factory with at least read (view) permissions for General Storage. Credentials configured with *Basic* mode are not supported.

If you have not configured credentials with view permissions, you are redirected to the AWS credentials setup page when you open the Overview tab.

If you haven't configured AWS credentials, see [Add AWS credentials](#).


#### Activate the EDA dashboard

After credentials with view permissions are confirmed, you must activate the dashboard for Workload Factory to begin collecting CloudWatch metrics for your FSx for ONTAP file systems.



Metrics collection can take some time after you provide consent. The dashboard notifies you while the initial collection is in progress.

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu  and then select **EDA**.
3. Select the **Overview** tab.

4. If no credentials with view permissions are detected, select **Add credentials** and follow the prompts to configure AWS credentials with view permissions. Then return to the **Overview** tab.
5. Review the consent prompt describing the CloudWatch metrics that will be collected for your FSx for ONTAP file systems.
6. Select **Activate** to activate the dashboard and begin metrics collection.

## Result

Workload Factory begins collecting CloudWatch metrics for all FSx for ONTAP file systems associated with your configured AWS credentials. The overview dashboard populates as metrics become available. A notification is displayed if collection is still in progress.

## Filter the dashboard

Use the filters at the top of the overview dashboard to focus on the file systems relevant to your current task. These are Workload Factory-level filters that apply to both the latency and ONTAP events information and the clusters table.

Available filters:

- **Region:** Filter by one or more AWS regions.
- **AWS account:** Filter by one or more AWS accounts associated with your configured credentials.

When you update the filter selections, all information is refreshed to show only the matching file systems and clusters.

## Cluster health status

At the top of the overview dashboard, a summarized view of event activity and health status across your filtered clusters is displayed. This information is only shown when at least one FSx for ONTAP link is associated with your file systems. If no links are available, the information is hidden.

## Latency

The Latency section displays the number of latency events detected across the file systems in scope.

- The count reflects latency events generated by the [FSx latency analysis](#) feature.
- Selecting **Review** navigates to the **Latency** tab.
- You can only view latency information if you have enabled latency monitoring. If you have not configured latency thresholds, select **Configure**. For details on latency monitoring, see [FSx latency analysis](#).

## Capacity

The Capacity section displays the number of EMS messages related to capacity issues, along with the number of file systems affected.

EMS events monitored includes but is not limited to:

- Aggregate nearly full
- Aggregate full
- Volume nearly full

- Volume full
- Snapshot reserve nearly full
- Snapshot reserve full
- Directory size full
- Flex groups full
- Inodes full

For a full list of the EMS events monitored, see [Capacity events](#).

Selecting **Capacity** navigates to the capacity analysis screen in Storage workloads.

### Availability & protection

The Availability & protection section displays the number of EMS messages related to availability and data protection, along with the number of file systems affected.

Events monitored include FlexCache and SnapMirror-related EMS events.

Selecting **Availability & protection** navigates to the availability and protection analysis screen in Storage workloads.

### Security & other

The Security & other section displays the number of EMS messages for events not categorized under Capacity or Availability & protection, along with the number of file systems affected. Events monitored include anti-ransomware protection, NFS authentication failures, and others.

Selecting **Security & other** navigates to the event analysis screen in Storage workloads.

## Clusters table

The clusters table provides a detailed view of each FSx for ONTAP cluster associated with your configured AWS credentials, filtered by the active region and AWS account selections. Data is collected from CloudWatch metrics. The table supports search, column filtering, pagination, column customization, and CSV export.

### Search and filter

- Use the search to look up specific clusters by name, file system ID, region, or other attributes.
- Select any column header to sort the table by that column.
- Use per-column filter controls to narrow results within the table.
- The table is paginated. Use the pagination controls at the bottom of the table to navigate between pages.

The table includes the following columns. Columns shown by default are indicated in the table. You can add or remove non-default columns using the column selector. For throughput columns, hover over the column header for details about the 30-day calculation period and usage formula.

| Column       | Default | Description                                   |
|--------------|---------|-----------------------------------------------|
| Cluster name | Yes     | The name of the FSx for ONTAP cluster.        |
| Region       | Yes     | The AWS region where the cluster is deployed. |

| Column                 | Default | Description                                                                                                                                                 |
|------------------------|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| AWS account            | No      | The AWS account associated with the cluster.                                                                                                                |
| SSD used capacity      | Yes     | The amount of SSD storage currently in use.                                                                                                                 |
| Total SSD capacity     | Yes     | The total provisioned SSD storage capacity.                                                                                                                 |
| SSD usage              | Yes     | The percentage of SSD capacity in use. Use this to identify clusters approaching capacity limits before they impact workload performance.                   |
| Capacity pool          | Yes     | The amount of storage in the capacity pool.                                                                                                                 |
| Average throughput     | Yes     | The average throughput over the last 30 days.                                                                                                               |
| P95 throughput         | No      | The 95th percentile throughput over the last 30 days.                                                                                                       |
| P99 throughput         | Yes     | The 99th percentile throughput over the last 30 days.                                                                                                       |
| Max throughput         | No      | The maximum throughput recorded over the last 30 days.                                                                                                      |
| Throughput usage (P99) | Yes     | P99 throughput relative to the provisioned throughput SKU. Use this to identify clusters where throughput is consistently near or at the provisioned limit. |
| Throughput SKU         | Yes     | The provisioned throughput tier. Compare this against throughput usage to evaluate whether the current SKU is appropriately sized.                          |


### Customize columns

To add or remove columns from the clusters table:

1. Select the column selector icon above the table on the right.
2. Select or deselect the columns you want to show or hide.
3. Select **Apply**.

### Export table to CSV

You can export the currently displayed table data to a CSV file for further analysis or reporting.

1. Apply any filters or column customizations you want reflected in the export.
2. Select the  button above the clusters table.

The CSV file is downloaded and contains all rows currently visible in the table, including only the columns currently shown.

# FSx latency analysis

## Monitor volume latency in EDA workloads

As an IT administrator or DevOps engineer managing EDA workloads, you can use latency analysis to proactively monitor volume performance by tracking read and write latency metrics across your FSx for ONTAP file systems. Configure customizable thresholds for warning and critical events to identify potential performance bottlenecks before they impact simulation run-time and time-to-market. When latency events are detected, automated basic analysis helps identify the root cause.

### Overview

High latency directly impacts simulation run-time and time-to-market for your EDA projects. Unhealthy volumes can cause significant performance degradation, leading to costly production delays. Latency analysis helps you proactively identify, troubleshoot, and remediate operational issues across your entire storage estate before they affect your workloads.

Latency analysis collects and monitors CloudWatch metrics for volume read and write operations. When both latency and IOPS thresholds are breached for all data points within a specified time range, the system generates alerts that appear in the latency events table.

When latency events are detected, the system automatically performs basic analysis using ONTAP QoS delay center metrics to identify the latency source.

This enables you to:

- Identify volumes experiencing performance degradation.
- Distinguish between warning-level and critical-level performance issues.
- Automatically analyze the root cause of latency issues.
- Track latency trends over time to optimize storage configurations.
- Take proactive action before latency impacts workload performance.

### Requirements

To use latency monitoring and analysis features, ensure you meet the following requirements:

#### AWS credentials and permissions

You must add AWS credentials to Workload Factory with read/write permissions. The latency monitoring feature requires access to CloudWatch metrics for all FSx for ONTAP volumes associated with your AWS credentials.

*Basic mode and read-only mode permissions are not supported for latency monitoring.*

If you haven't configured AWS credentials, see [Add AWS credentials](#).

#### FSx for ONTAP file system

You need at least one FSx for ONTAP file system with volumes deployed in your AWS environment. The latency monitoring feature automatically collects metrics for all volumes associated with your configured

AWS credentials.

## Link to FSx for ONTAP

To get insights from basic analysis, you must associate a link with your FSx for ONTAP file system. If no link is already associated, select **Associate link** in EDA, choose whether to create a new link or associate an existing link, and then select **Continue** to automatically go to the link creation page in Storage workloads.

For instructions on creating and associating links, see [Create a link](#).

## Understanding alerts

The latency analysis feature uses CloudWatch alarms to monitor volume performance. Understanding how alerts are triggered helps you configure appropriate thresholds and interpret the results.

### Metrics collected

The system collects the following CloudWatch metrics for each volume:

- **Read latency threshold:** Calculated as  $1000 * m2 / (m1 + 0.000001)$  where  $m1 = \text{DataReadOperations}$  and  $m2 = \text{DataReadOperationTime}$
- **Write latency threshold:** Calculated as  $1000 * m2 / (m1 + 0.000001)$  where  $m1 = \text{DataWriteOperations}$  and  $m2 = \text{DataWriteOperationTime}$

### Alert trigger conditions

An alert is triggered when all of the following conditions are met:

- The latency threshold is exceeded for the operation type (read or write).
- The IOPS threshold is exceeded for the operation type.
- Both conditions persist for all data points within the configured time range.

For example, with default warning thresholds, a read alert triggers only if read latency exceeds 6 ms AND read IOPS exceeds 100 ops/sec for all data points within a 10-minute period.

### Event severity

- **Warning events:** Indicate elevated latency that might need attention.
- **Critical events:** Indicate severe latency that requires immediate investigation.

## Configure latency thresholds

Configuring appropriate latency thresholds enables you to receive timely notifications when volumes are experiencing performance issues. By setting both warning and critical thresholds, you can distinguish between issues that need attention and those requiring immediate action, allowing you to manage your storage estate more effectively and prevent performance problems from impacting production workloads.

You can configure thresholds for both warning and critical events. Each event type includes separate thresholds for read and write operations. The system evaluates these thresholds continuously and generates alerts when conditions are met.




You must set critical event thresholds higher than warning event thresholds to ensure proper alert escalation. If not, you cannot save your configuration.

## About this task

For an alert to trigger, both the latency threshold and the IOPS threshold must be breached for all data points within the specified time range. This dual-condition logic helps reduce false positives by ensuring that high latency is sustained under significant load.

## Steps

1. Log in using one of the [console experiences](#).
2. Select the menu  and then select **EDA**.
3. From the EDA menu, select **Latency**.
4. In the EDA latency configuration page, configure the following thresholds:
  - **Warning events**
    - **Read latency threshold:** Enter the latency threshold in milliseconds. Default: 6 ms.
    - **Read IOPS threshold:** Enter the IOPS threshold in operations per second. Default: 100 ops/sec.
    - **Read time range:** Enter the time range in minutes (5-20). Default: 10 minutes.
    - **Write latency threshold:** Enter the latency threshold in milliseconds. Default: 8 ms.
    - **Write IOPS threshold:** Enter the IOPS threshold in operations per second. Default: 100 ops/sec.
    - **Write time range:** Enter the time range in minutes (5-20). Default: 10 minutes.
  - **Critical events**
    - **Read latency threshold:** Enter the latency threshold in milliseconds. Default: 12 ms.
    - **Read IOPS threshold:** Enter the IOPS threshold in operations per second. Default: 100 ops/sec.
    - **Read time range:** Enter the time range in minutes (5-20). Default: 10 minutes.
    - **Write latency threshold:** Enter the latency threshold in milliseconds. Default: 15 ms.
    - **Write IOPS threshold:** Enter the IOPS threshold in operations per second. Default: 100 ops/sec.
    - **Write time range:** Enter the time range in minutes (5-20). Default: 10 minutes.
5. Select **Apply**.

## Result

Workload Factory begins collecting latency metrics for all FSx for ONTAP volumes associated with your AWS credentials. Metrics are collected at least every 20 minutes. The latency events table displays any volumes that breach your configured thresholds.

## View latency events

As an administrator managing multiple file systems and volumes, the latency events table provides a centralized view of all performance issues requiring your attention. The table displays all warning and critical events detected within the last 72 hours. Each event includes automated basic analysis results in the Details column, helping you quickly identify the root cause of latency issues and prioritize remediation efforts across your estate.

- Only the latest breach for each volume appears in the table. If a volume experiences multiple breaches, only the most recent event is displayed.
- Events are automatically removed after 72 hours.
- The table displays a maximum of 200 events. Older events are removed as new events are added.

## Steps

1. In the **Latency** tab, view the latency events table.
2. Review the information for each event including:
  - **Severity**: Indicates whether the event is Critical or Warning.
  - **Volume name**: The name of the affected volume.
  - **Volume ID**: The ID of the affected volume.
  - **File system**: The FSx for ONTAP file system containing the volume.
  - **Time detected**: When the breach was detected
  - **Median latency**: The median latency value during the breach period.
  - **Details**: Automated basic analysis results identifying the latency source and recommended actions.
3. To sort the table, select any column header. By default, critical events appear first sorted by time, followed by warning events sorted by time.
4. To dismiss one or more events, next to each event select **Dismiss**.
5. To add columns to the table, select the column icon, choose the columns, and select **Apply**.

## Understanding basic analysis

Basic analysis helps you quickly identify the root cause of latency issues without manual investigation. When a latency event is detected, Workload Factory automatically performs basic analysis using ONTAP QoS delay center metrics. The analysis identifies which component is causing the latency and provides actionable guidance in the Details column of the latency events table, enabling you to understand root cause.



There might be slight discrepancies between latency values from ONTAP QoS analysis and CloudWatch data due to different collection methodologies. The basic analysis uses ONTAP data for root cause identification.

## Analysis scenarios

The basic analysis evaluates multiple latency components and provides specific guidance based on the results for each scenario:

- **Flexcache**: Latency per I/O operation for FlexCache operations
- **Capacity pool**: Latency per I/O operation for capacity pool operations
- **QoS min**: Latency per I/O operation for QoS Policy Group Floor
- **QoS max**: Latency per I/O operation for QoS Policy Group Ceiling
- **Disk**: Latency per I/O operation in the Storage subsystem
- **Data**: Latency per I/O operation in the WAFL subsystem file system, which includes tasks such as CPU processing, metadata updates, and cache management
- **Cluster**: Latency per I/O operation across the internally connected nodes in a cluster
- **Other**: Latency per I/O operation on FSx for ONTAP subsystems

## Manage latency configuration

After the initial configuration, you can edit your thresholds.

## Steps

1. In the **Latency** page, select **Edit**.
2. Modify any of the threshold values as needed.



Ensure that critical thresholds remain higher than warning thresholds. The system displays an error if you configure critical thresholds lower than warning thresholds.

3. Select **Apply** to save your changes.

## Best practices

Consider these recommendations when configuring and using latency analysis:

- **Set realistic thresholds:** Configure thresholds based on your workload requirements. Default values provide a starting point but might need adjustment for your specific environment.
- **Start with warning thresholds:** Use warning events to establish baseline performance expectations before fine-tuning critical thresholds.
- **Consider time ranges carefully:** Shorter time ranges (5-10 minutes) detect issues faster but might generate more alerts. Longer time ranges (15-20 minutes) reduce false positives but might delay detection.
- **Monitor trends:** Regularly review the latency events table to identify patterns or recurring issues that might indicate underlying configuration problems.
- **Coordinate IOPS and latency thresholds:** The dual-condition logic means both must be exceeded. Setting very high IOPS thresholds might prevent alerts even when latency is problematic.
- **Review dismissed events:** Periodically review why events were dismissed to identify opportunities for threshold adjustment or infrastructure improvements.

# Project usage analysis

## Use the project dashboard

As an IT administrator or project manager overseeing EDA workloads, you can use the Project tab to observe and analyze project usage across your FSx for ONTAP file systems and volumes. The dashboard helps you plan capacity, assess costs, and make data-driven decisions about resource allocation. You can create up to five custom filters based on your AWS tags, with options for single or multi-selection, to organize and filter the data according to your specific business requirements such as projects, teams, or applications.

The dashboard provides visibility into storage usage patterns, helping you identify resource-intensive projects, track costs, and optimize storage allocation before capacity or performance issues impact your EDA simulations.

The dashboard helps you monitor storage usage for capacity allocated, capacity used, throughput, and IOPS. For proactive latency performance monitoring, see [Monitor volume latency](#).

The CloudWatch metrics collected include:

- Provisioned capacity: Volume-level metric representing the provisioned storage capacity.
- Used capacity: Volume-level metric representing the storage used.
- Average/Maximum throughput: Calculated as the average or maximum of the sum of DataReadBytes and DataWriteBytes over the specified time period.
- Average/Maximum IOPS: Calculated as the average or maximum of the sum of DataReadOperations, DataWriteOperations, and MetadataOperations over the specified time period.

## Setup the dashboard

Setting up custom filters enables you to view metrics organized by your business structure—such as by project, team, or application. This organization helps you track costs accurately, identify which projects consume the most resources, and make informed decisions about resource allocation and budget planning.

To effectively use the dashboard, configure up to five AWS tags on your FSx for ONTAP volumes based on your business requirements, for example tags that represent projects or business units. For details, see [Configure custom filters](#).

AWS tags are metadata for your AWS resources. They help you categorize your AWS resources in different ways, for example, by project, application, or business unit. For more details about tagging, see [What are tags?](#) and [AWS Resource Groups Tagging API Reference](#).

When they are configured, in the **Tag configuration** page provide the AWS tag key names and the corresponding labels to display in your dashboard.

When these tags are applied, Workload Factory begins collecting and displaying the relevant CloudWatch metrics.

Your dashboard becomes a dynamic tool for organizing, tracking costs, and filtering resources based on your organizational needs.

## Configure custom filters



Only the first admin user to access the project dashboard can configure, add, or delete custom filters. All other users, including subsequent admins, have read-only access to the filters.

Custom filters allow you to organize your dashboard data according to how your organization operates. By aligning filters with your project structure, teams, or cost centers, you can quickly access the specific data you need without navigating through irrelevant information. This makes capacity planning, cost tracking, and performance monitoring more efficient.

You can configure up to five custom filters based on your AWS tags. Each custom filter requires three components: a filter label name, an AWS tag key name, and a selection type (single or multi-selection). If you don't configure any custom filters, the default filters (file system, volume type, and time range) remain available so you can still view and interact with your dashboards.

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **EDA**.
3. Select **Project**.

If you have not already configured your dashboard, you are automatically prompted to do so.

4. Select **+ Add filters**.
5. For each custom filter you want to create (up to five), provide the following:
  - **Filter label name:** The display name that appears in the dashboard.
  - **AWS tag key name:** The AWS resource tag key that corresponds to this filter.
  - **Multi-selection:** Choose whether this filter allows single selection or multiple selection. When you select **Multi-selection**, you can select multiple values for this filter simultaneously. Single-selection restricts you to selecting only one value at a time.



Filters appear on the dashboard in the order you configure them. Consider organizing your most frequently used filters first for easier access.

6. Select **Apply**.

You can delete a custom filter by selecting the trash icon next to that filter before applying your changes.

7. To view your changes after applying any tags or filters, select the refresh icon on the dashboard. The dynamic dashboard, configured with the new filters, is displayed in your EDA projects dashboard.
8. To edit the dashboard configuration later, select **Configure**.

## Filter the dashboard

Filtering helps you focus on specific subsets of your storage resources, making it easier to analyze trends, identify issues, and generate reports for stakeholders. By combining multiple filters, you can answer specific business questions like "How much storage is my ASIC project team using?" or "Which volumes in the US-West region have the highest IOPS?"

You can filter the information displayed on the dashboard using a combination of default filters and any custom filters you created.

The following default filters are always available:

- Credentials
- Region
- File system
- Volume type
- Time range

Custom filters you configure appear on the dashboard in addition to these default filters. When using filters:

- **Multi-selection filters** allow you to select multiple values simultaneously to broaden your view. For example, you might select multiple projects to view combined metrics.
- **Single-selection filters** restrict you to selecting only one value at a time, useful when you need to focus on a specific resource or category.

When you have selected your required filters, select the refresh icon to update the dashboard information.

For an explanation of the information shown on a card, select the information icon for that card.

## View volume details

Understanding both aggregated and individual volume performance helps you make informed decisions about capacity planning and performance optimization. The Total view shows overall trends, while the Volume view helps you identify specific volumes that may need attention, optimization, or additional resources.

The dashboard provides two viewing modes to help you analyze your storage metrics: Total view and Volume view. You can switch between these modes using the tabs available on the dashboard.

### Total view

The Total view (default) displays aggregated metrics across all volumes that match your selected filters. This view provides a high-level overview of your overall storage performance, showing combined capacity, IOPS, and throughput metrics. Use this view to understand overall trends, plan for future capacity needs, and report on total resource consumption across projects or teams.

### Volume view

The Volume view displays individual volume performance over time, showing the top 10 volumes for each metric. This view helps you identify specific volumes that are driving resource usage and observe their behavior over the selected time period. Use this view to find performance bottlenecks, identify volumes that may need rightsizing, or pinpoint which specific projects or workloads are consuming the most resources.

To switch to the Volume view, select the **Volume** tab on the dashboard.

### Volume metrics displayed

When you select the Volume view, the dashboard displays the top 10 volumes out of your total number of volumes.

\* **Volume used capacity:** Shows volumes with the highest current used capacity.

\* **IOPS:** Shows volumes with the highest average IOPS during the selected time period.

\* **Throughput:** Shows volumes with the highest average throughput during the selected time period.



The dashboard displays only the top 10 volumes for each metric. If you have more than 10 volumes, some volumes might not be displayed in the detailed view.

When the same volumes appear across **Volume used capacity**, **IOPS**, and **Throughput** metrics, the dashboard uses consistent color coding in the legend to make it easier to track specific volumes across different metrics.

The horizontal axis displays the time range, while a legend shows all volumes (up to 10) represented in the graph.

#### **Interactive volume data**

Interactive hover data provides precise metrics at specific points in time, helping you correlate performance changes with specific events or timeframes. This detailed information is valuable for troubleshooting performance issues, validating optimization efforts, or explaining resource usage patterns to stakeholders.

You can hover over any volume line in the graphs to view detailed information:

#### **Volume used capacity:**

Displays the volume name, used capacity at that point in time, and allocated capacity.

#### **IOPS:**

Displays the volume name, average IOPS for the time range, and maximum IOPS for the time range.

#### **Throughput:**

Displays the volume name, average throughput for the time range, and maximum throughput for the time range.

This interactive data helps you analyze volume performance patterns and identify potential bottlenecks or optimization opportunities.

# Use Perforce integration

## Learn about Perforce integration in NetApp Workload Factory for EDA

Integrating Perforce with CI/CD pipelines enhances the development process by automating builds, tests, and deployments, leading to faster and more reliable software delivery.

Continuous Integration and Continuous Deployment (CI/CD) in EDA is a rapid build environment creation tool for software builders. It enables fast setup of personal development environments, saving time and enabling self-service for developers while empowering DevOps teams to retain control of the infrastructure. Using CI/CD, software developers can quickly create workspaces without needing specialized data storage or understanding of the development infrastructure.

### What is CI/CD?

By using CI/CD, you can streamline the way developers manage and interact with different versions of their software. It works with Perforce Helix Core to instantly clone software versions and create workspaces for development, QA, and CI/CD.

You can easily create a project and assign a volume that represents your software environment and its artifacts. As you update your software, you can take snapshots of the volume, capturing the state of your software at that moment. You can access any software version instantly without resyncing, saving time and resources.

Using the snapshot and clone capabilities of NetApp ONTAP, you can quickly access different versions of your software, so you can develop and release updates faster.

For more information about Workload Factory, refer to the [Workload Factory overview](#).

### CI/CD features

- Create, edit, and remove projects. See [Manage projects](#).
- Create snapshots of defined software versions. See [Manage project versions](#).
- Create, and delete workspaces (based on clones). See [Create a workspace](#).
- Create access policies to control access to a project.
- Analyze the capacity usage of each project.
- Control clone size limits and clone retention for each project.
- Integrate with version control systems such as Perforce. See [Integrate with Perforce](#).

### Projects and workspaces in EDA

You can create a project and assign a volume that represents your software environment and its artifacts. Each time that you create a new version of the software, you need resync the volume data and create a project snapshot to mark the volume state as a known version. The project source volume might get rolling updates and have multiple snapshots to mark multiple versions. You can use each snapshot immediately as an instant clone, a dedicated or shared editable repository available to developers, QA or build processes. A clone in the context of a specific software version is a workspace.

## Automation with Workload Factory Codebox

Workload factory introduces built-in automation with the *Codebox*. The Codebox offers the following automation benefits:

- **Code snippet generation:** Infrastructure-as-Code (IaC) snippets are generated during resource creation, enabling seamless integration with existing orchestration workflows.
- **Infrastructure-as-code co-pilot:** the Codebox is an Infrastructure-as-code (IaC) co-pilot that helps developers and DevOps generate code to execute any operation supported by Workload Factory.
- **Code viewer and automation catalog:** the Codebox provides a code viewer for quick analysis of automation and an automation catalog for quick future re-use.

## Cost

There is no cost for using the CI/CD capability of Workload Factory.

## Licensing

No special licenses are needed from NetApp to use the CI/CD capabilities of Workload Factory.

## Regions

EDA is supported in all commercial regions where FSx for ONTAP is supported. [View supported Amazon regions.](#)

The following AWS regions aren't supported:

- China regions
- GovCloud (US) regions
- Secret Cloud
- Top Secret Cloud

## Getting help

Amazon FSx for NetApp ONTAP is an AWS first-party solution. For support, use the Support Center in your AWS Management Console to open a case. Select "FSx for ONTAP" and the category, then provide the required information.

For general questions about Workload Factory or Workload Factory applications and services, refer to [Get help for EDA for Workload Factory.](#)

## EDA requirements

Ensure that Workload Factory and AWS are set up properly before you use NetApp Workload Factory for EDA. This includes having your AWS log in credentials, a deployed FSx for ONTAP file system, and more.

### Workload factory login and account

You'll need to [set up an account with Workload Factory](#) and log in using one of the [console experiences](#).

## AWS credentials and permissions

You need to add AWS credentials to Workload Factory with read/write permissions, which means you'll be using Workload Factory in *read/write* mode for EDA.

*Basic* mode and *read-only* mode permissions are not supported at this time.



AWS credentials are also required to use the latency monitoring feature, which collects CloudWatch metrics for volume performance analysis. [Learn about latency monitoring.](#)

[Learn how to add AWS credentials to Workload Factory](#)

## FSx for ONTAP file system

You need a minimum of one FSx for ONTAP file system:

- The file system will be used by EDA to store the projects and workspaces that you create.

This FSx for ONTAP file system must use FlexVol volumes. FlexGroup volumes are not supported.

- You'll need to know the AWS region, VPC, and subnet where the AWS FSx for ONTAP file system resides.
- You'll need at least one volume in the filesystem with the following configuration:
  - The volume must be configured as an NFS share.
  - The filesystem must be configured with a link. [Learn more about links.](#)
- You'll need to consider the tag key/value pairs that you want to apply to the AWS resources that are part of this deployment (optional).

[Learn how to deploy and manage FSx for ONTAP file systems](#)

# Manage EDA projects

You can manage EDA projects to control how your code and artifacts are managed for each project in NetApp Workload Factory for EDA.

## Create a project

You can create a new EDA project so that you can leverage the data protection features of your Amazon FSX for NetApp ONTAP filesystem for your code and artifacts.

.Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **EDA**.
3. Select **CI/CD**.
4. Select **Create project**.
5. On the Create project page, provide the following:
  - a. **Project name**: Enter a name for the project.
  - b. **Description**: Enter a description for the project.
  - c. **Filesystem**: Provide the following:

- i. **Credentials:** Select the Amazon AWS credentials to use. EDA uses these credentials to discover FSx for ONTAP filesystems that you can use with this project and to create clones and snapshots of projects.
- ii. **Region:** Select the region that this FSx for ONTAP filesystem resides in.
- iii. **FSx for ONTAP filesystem:** Select an FSx for ONTAP filesystem to use with this project.

You can only select filesystems that are configured with a link. [Learn more about links.](#)

- iv. **Choose a volume:** Select a volume on which to store the project; EDA uses this volume as a software repository.

You can only select volumes that are configured as an NFS share.

- d. **Operation policies:** Provide limits for project clones:

- i. **Maximum retention in days:** Enter the maximum number of days that a clone should be retained. After this number of days, Workload Factory removes the clone.
- ii. **Maximum number of clones per user or group:** Enter the maximum number of clones that can be provisioned for a user or group.
- iii. **Maximum clone size in GiB:** Enter the maximum size in GiB of a project clone.

- e. **Access policies:** Explicitly grant project access to specific users or groups:

- i. **Policy enforcement scope:** Enter single IP addresses or IP address ranges to limit project access to only those IP addresses or ranges.

For example: `172.16.0.0/24`

- ii. **User or group identifiers:** Enter user or group identifiers to limit project access to only those users or groups.

For example: `User1234`

- 6. Select **Create**.

## Result

The project is created, and appears in the list of projects on the Projects page.

## View existing projects

You can view existing projects created in NetApp Workload Factory for EDA by following these steps.

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **EDA**.
3. Select **CI/CD**.
4. Select **Go to Projects page**.
5. View the existing projects listed on the Projects page.

## Edit a project

You can edit the settings of a project at any time.

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **EDA**.
3. Select **CI/CD**.
4. Select **Go to Projects page**.
5. On the Projects page, select **...** for the project you want to edit.
6. Make any needed changes to the project configuration.
7. Select **Save**.

## View a project's workspaces

A clone or snapshot of a project is known as a workspace. When you create a workspace, it is retained for as long as the project's operation policy allows. You can view existing workspaces for a project by following these steps.

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **EDA**.
3. Select **CI/CD**.
4. Select **Go to Projects page**.
5. On the Projects page, choose a project and select **View**.
6. View the status and details of all workspaces for this project.
7. If you see alerts or warnings for a workspace, hover over the alert or warning icon to see the reason.

## Delete a project

You can delete a project when it is no longer needed by following these steps.

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **EDA**.
3. Select **CI/CD**.
4. Select **Go to Projects page**.
5. On the Projects page, select **...** for the project you want to delete.
6. Select **Delete**.
7. In the confirmation dialog, select **Delete**.

### Result

The project is deleted, and any code or artifacts associated with the project are deleted from the volume. Snapshots and clones of the project are retained.

# Manage versions of NetApp Workload Factory for EDA projects

Work with different versions of your EDA projects by creating on-demand snapshots and clones directly from Workload Factory. Snapshots and clones of a project are stored in the filesystem that was associated with the project when it was created. You can also manage snapshots and clones using the [Workload Factory REST API](#).

## Create a snapshot of a project

You can create a snapshot of a project by following these steps.

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **EDA**.
3. Select **CI/CD**.
4. Select **Go to Projects page**.
5. On the Projects page, select **•••** for the project you want to snapshot.
6. In the resulting menu, select **Create a snapshot**.
7. In the **Create snapshot** dialog, choose a name for the snapshot and select **Create**.

## Create a clone of a project

Clone an EDA project from a snapshot by following these steps. When you create a clone, a new editable volume is created to contain the clone.

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **EDA**.
3. Select **CI/CD**.
4. Select **Go to Projects page**.
5. On the Projects page, select **•••** for the project you want to clone.
6. In the resulting menu, select **Create a clone**.
7. In the **Create clone** dialog, do the following:

- a. Enter a name for the clone.

The default name for the clone is the project name with a suffix of the current date and time.

- b. Select a snapshot to use as the base for the clone.
- c. Select **Create**.

### Result

Workload factory creates a new clone of the project, and the clone appears as a new project on the Projects page.

# Create an EDA workspace

A workspace in NetApp Workload Factory for EDA is a Perforce representation of a project at a specific moment in time. Workspaces are created using a project snapshot as a basis. You can create new workspaces within an EDA project. You can create workspaces from the Perforce UI.

## Before you begin

Ensure you have integrated EDA with the Perforce Helix Visual Client. See [Integrate EDA with Perforce](#) for more information.

## Steps

1. Log in to Perforce.
2. In the Perforce menu, select **View > WF**.

The Workload Factory login screen appears within the Perforce UI.

3. Log in using one of the [console experiences](#).
4. Select the menu and then select **EDA**.
5. Select **CI/CD**.
6. Select **Create project** and then select **Create workspace**.
7. On the Create a workspace project page, provide the following:
  - a. Select a snapshot to use as a basis for the workspace.
  - b. Enter a name for the workspace.
  - c. Optionally, enter a user identifier to claim the workspace. This identifier should match the Perforce user ID of the developer that will use this workspace.
8. Select **Create**.

## Result

The workspace is created, and appears in the list of workspaces on the Workspaces page.

# Automate EDA workload tasks with Codebox

You can automate project creation and data protection operations with Codebox. Codebox is an infrastructure as code (IaC) co-pilot that helps you generate code to execute any operations supported by Workload Factory.

Learn more about [Codebox automation](#) and how to use it.

# Integrate EDA with Perforce

Integrate EDA with the Perforce Helix Visual Client (P4V) so that developers can manage your workspaces using the Perforce CLI. This enables developers to quickly switch between projects and workspaces, saving time during development.

## Steps

1. Download the [P4V integration file](#).
2. Open P4V and go to **Tools > Manage Tools > HTML tabs**.
3. Select **Import HTML tabs**.
4. Select the P4V integration XML file and select **Import**.
5. Go to **View > Workload Factory**.

### **Result**

The NetApp Workload Factory for EDA web UI appears as an HTML tab within the P4V client.

### **What's next?**

[Create an Amazon EC2 deployment plan using the migration advisor](#).

# Knowledge and support

## Register for support for NetApp Workload Factory for EDA

Before you can open a support case with NetApp technical support, you need to add a NetApp Support Site account to Workload Factory and then register for support.

Support registration is required to receive technical support specific to NetApp Workload Factory and its storage solutions and services. You must register for support from the NetApp Console, which is a separate web-based console from Workload Factory.

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 Workload Factory documentation for that product.

[Amazon FSx for ONTAP](#)

### Support registration overview

Registering your account ID support subscription (your 20 digit 960xxxxxxx serial number located on the Support Resources page in the NetApp Console) serves as your single support subscription ID. Each NetApp account-level support subscription must be registered.

Registering enables capabilities like opening support tickets and automatic case generation. Registration is completed by adding NetApp Support Site (NSS) accounts to the NetApp Console as described below.

### Register your account for NetApp support

To register for support and activate support entitlement, one user in your account must associate a NetApp Support Site account with their NetApp 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 NetApp Console.

#### Steps

1. In the upper right of the Workload Factory console, select **Help > Support**.

Selecting this option opens the NetApp Console in a new browser tab and loads the Support dashboard.

2. From the NetApp Console menu, select **Administration**, and then select **Credentials**.
3. Select **User Credentials**.
4. Select **Add NSS credentials** and follow the NetApp Support Site (NSS) Authentication prompt.
5. To confirm that the registration process was successful, select the Help icon, and select **Support**.

The **Resources** page should show that your account is registered for support.

|                                                                                                                                   |                                                                                                                                  |
|-----------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
|  9601111222224444455555<br>Account Serial Number |  Registered for Support<br>Support Registration |
|-----------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------|

Note that other NetApp Console users will not see this same support registration status if they have not associated a NetApp Support Site account with their NetApp Console login. However, that doesn't mean that your NetApp account is not registered for support. As long as one user in the account 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 NetApp 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 NetApp 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 NetApp 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 Workload Factory console, select **Help > Support**.

Selecting this option opens the NetApp Console in a new browser tab and loads the Support dashboard.

2. Locate your account ID serial number from the Support Resources page.

|                                                                                                                                   |                                                                                                                                                                                                                                                                                                |
|-----------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  96015585434285107893<br>Account serial number |  Not Registered<br>Add your NetApp Support Site (NSS) <a href="#">credentials</a> to BlueXP<br>Follow these <a href="#">instructions</a> to register for support in case you don't have an NSS account yet. |
|-----------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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 the account 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 NetApp Console login by completing the steps under [Existing customer with an NSS account](#).

## **Get help with Workload Factory for EDA**

NetApp provides support for Workload Factory and its cloud services in a variety of ways. Extensive free self-support options are available 24x7, such as knowledgebase (KB) articles and a community forum. Your support registration includes remote technical support via web ticketing.

### **Get support for FSx for ONTAP**

For technical support related to FSx for ONTAP, its infrastructure, or any solution using the service, refer to "Getting help" in the Workload Factory documentation for that product.

#### [Amazon FSx for ONTAP](#)

To receive technical support specific to Workload Factory and its storage solutions and 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 Workload Factory documentation that you're currently viewing.

- [Knowledge base](#)

Search through the Workload Factory knowledge base to find helpful articles to troubleshoot issues.

- [Communities](#)

Join the Workload Factory 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 register for support. associate your NetApp Support Site credentials with your Workload Factory login. [Learn how to register for support.](#)

## Steps

1. In the upper right of the Workload Factory console, select **Help > Support**.

Selecting this option opens the NetApp Console in a new browser tab and loads the Support dashboard.

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 **Workload Factory**.

- **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 Working Enviroment


Select 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) Upload 

No files selected  

### 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 system serial number. You can seek assistance using one of the following options:

- Use the in-product chat
- Submit a non-technical case at <https://mysupport.netapp.com/site/help>

## Manage your support cases (Preview)

You can view and manage active and resolved support cases directly from the NetApp Console. You can manage the cases associated with your NSS account and with your company.

Case management is available as a Preview. We plan to refine this experience and add enhancements in upcoming releases. Please send us feedback by using the in-product chat.

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 upper right of the Workload Factory console, select **Help > Support**.

Selecting this option opens the NetApp Console a new browser tab and loads the Support dashboard.


2. Select **Case Management** and if you're prompted, add your NSS account to the NetApp Console.

The **Case management** page shows open cases related to the NSS account that is associated with your NetApp 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**.



# Legal notices for NetApp Workload Factory for EDA

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

## Copyright

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

## Trademarks

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

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

## Patents

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

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

## Privacy policy

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

## Open source

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

[NetApp Workload Factory](#)

## 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.