



Implement file system best practices

Workload Factory for NetApp Console

NetApp

October 07, 2026

This PDF was generated from <https://docs.netapp.com/us-en/data-services-workloads/fsx-storage/configuration-analysis.html> on October 07, 2026. Always check docs.netapp.com for the latest.

Table of Contents

- Implement file system best practices 1
 - Configuration analysis for FSx for ONTAP file systems 1
 - How it works 1
 - Why it matters 1
 - Analysis requirements 1
 - Best practices and recommendations for storage workloads 2
- Implement well-architected FSx for ONTAP file system configurations 5
 - About this task 5
 - Before you begin 5
 - Fix a configuration issue 6
 - Dismiss a configuration analysis 6
 - Reactivate a dismissed configuration analysis 7
- Custom rules 7
 - Learn about custom rules in NetApp Console 8
 - Create a custom rule in NetApp Workload Factory 10
 - Manage custom rules for NetApp Workload Factory 13

Implement file system best practices

Configuration analysis for FSx for ONTAP file systems

NetApp Workload Factory analyzes Amazon FSx for NetApp ONTAP file system configurations regularly to determine if any there are any configuration issues. When issues are found, Workload Factory shows you what the issues are and explains what needs to change to ensure your file system storage achieves peak performance, cost efficiency, and compliance with best practices.

Key capabilities include:

- Daily configuration analysis
- Automatic best practice validations
- Proactive observability
- Insights to action
- AWS Well-Architected Framework advisor

How it works

Workload Factory analyzes your workloads running on supported cloud storage systems deployments daily. The analysis provides well-architected status, insights, and recommendations.

After the daily analysis completes, configurations appear as "optimized" or "not optimized" in the Well-architected dashboard for the deployment. You'll find configuration issues by category and a list of configuration issues and recommendations. You can review the recommendations for configuration issues. Workload Factory can automatically fix some issues; other issues require manual intervention. In this case, Workload Factory provides detailed instructions to help you implement the recommended changes.

You can dismiss the analysis of configurations that do not apply to your environments. This avoids unnecessary alerts and inaccurate optimization results.

You can also create custom rules in plain language to validate your environments against your organization's standards and NetApp best practices. Test rules with a dry run, then schedule them to run across your environments. [Learn more about custom rules.](#)

Why it matters

Workload Factory applies best practices to large storage, database, and VMware environments by combining ongoing assessment with recommendation insights and remediation. Automated fixes reduce human error, ensure uniform management, and preserve performance and reliability across your workload infrastructures.

Analysis requirements

For a complete file system analysis, you must do the following:

- Associate a link. Link connectivity lets Workload Factory analyze all file system configurations like data protection and performance.

[Learn how to associate an existing link or to create and associate a new link.](#)

- Grant *view, planning, and analysis* permissions in your AWS account.

[Learn how to grant permissions to an AWS account.](#)

Best practices and recommendations for storage workloads

The well-architected dashboard assesses storage configurations against ONTAP best practices and the AWS Well-Architected Framework. The assessment also recommends improvements and fixes.

The well-architected analysis categorizes configurations in the following pillars of the framework: *reliability*, *security*, *operational excellence*, *cost optimization*, and *performance efficiency*.

Reliability

Reliability ensures that workloads perform their intended functions correctly and consistently, even when there are disruptions.

- **Volume backups**

Backing up your volumes helps support data retention and compliance needs. Use volume backups to set up automated backups and retention for your data.

- **Snapshot policy**

Schedule local snapshots for efficient backup and quick restores. Snapshots are instant, point-in-time images of your volumes. This recommendation doesn't apply to database environments.

- **Application-consistent snapshots**

For database environments, use application-consistent snapshots with NetApp SnapCenter to take accurate, reliable snapshots of your volume data at a specific moment in time. This keeps your apps running smoothly and your data safe. SnapCenter makes backups easier and helps you restore data quickly and correctly, reducing downtime and protecting your most important workloads.

- **Cross-region replication**

Cross-region replication ensures that your data is replicated to another AWS region, providing enhanced data durability and availability. Set up cross-region replication to support disaster recovery and compliance.

- **SSD capacity headroom**

The SSD storage tier capacity should not exceed 80% utilization. This might impact data reads and writes to your capacity pool storage tier and impact the throughput capacity of your file system. When capacity runs out, data volumes become read-only. Services trying to write new data also fail.

- **Replication policy labels**

The snapshot policy labels of the source volume and the replication policy labels must match to ensure data reliability.

- **File object capacity**

The file capacity threshold should be raised to avoid hitting the volume capacity limit. Low file capacity (inodes) prevents writing additional data to the volume. Workload Factory recommends maintaining file capacity utilization below 80% to allow new file creation.

Security

Security emphasizes protecting data, systems, and assets through risk assessments and mitigation strategies.

- **Ransomware protection (ARP/AI)**

NetApp Autonomous Ransomware Protection with AI (ARP/AI) helps protect your volumes from ransomware threats. Workload Factory recommends enabling ARP/AI for all volumes.

- **Unauthorized iSCSI volume access**

Volumes serving application data using iSCSI should not allow NAS access at the same time. Workload Factory recommends restricting iSCSI volumes from using any other protocol.

Operational excellence

Operational excellence focuses on delivering the most optimal architecture and business value.

- **Automatic capacity management**

Automatic capacity management should be enabled to regularly ensure that the SSD tier doesn't exceed the threshold.

- **Volume capacity headroom**

Keep volume utilization below 80% to prevent potential application read/write impacts. You can increase capacity manually or with volume autogrow.

- **FlexCache write mode**

For optimal performance, set the cache relationship write mode to best suit your workload. Write-around mode provides better performance for read-heavy workloads with small files, whereas write-back mode provides better performance for write-heavy workloads with large files.

- **FlexCache replica size**

Enable volume autosize and scrubbing on cache volumes to maintain optimal size and focus the cache on hot data for peak efficiency.

- **Host-visible capacity reporting**

Enable host-visible capacity reporting to provide better visibility into storage usage at the volume level.

- **Operating System (OS) type**

Confirm that the ONTAP LUN operating system (OS) type value matches the operating system partitioning scheme to achieve I/O alignment. Incorrect configuration might reduce performance.

- **Block device space management**

Configure block device space settings for LUNs used by database host instances to prevent write failures and improve space efficiency on FSx for ONTAP. This configuration applies to the recommended combination of settings for thin-provisioned volumes: space reservation, space allocation, and fractional reserve.

- **Thin provisioning**

Configure thin provisioning for FSx for ONTAP volumes hosting database instances. This approach allows more logical data to be stored than physically available.

Cost optimization

Cost optimization helps you get the most value for your business while keeping costs low.

- **Cold data tiering**

Enable cold data tiering to reduce SSD storage tier utilization. Apply a tiering policy to every volume focused on general data storage. For database volumes, enable cold data tiering where appropriate. FSx for ONTAP continuously scans the data for cold data and moves it to the capacity pool tier without disruption.

- **Storage efficiencies**

Enable storage efficiencies — compaction, compression, and deduplication — to optimize storage utilization and reduce the SSD tier cost.

- **Backup expiration**

Delete any manually created backups for volumes that were deleted, and any backups older than two weeks. These backups don't expire automatically, can increase costs, and add complexity if you no longer need them for SLA or recovery.

- **Snapshot expiration**

Delete old or manually created snapshots you no longer need. Removing them can lower costs and reduce clutter, especially if they aren't required for your SLA or recovery needs.

- **Inactive block devices**

After a block device isn't used for seven days, archive block device data or delete the unused block device to reduce costs.

- **Overprovisioned SSD capacity**

Excess SSD capacity adds unnecessary cost. Reduce the SSD tier while maintaining a 20% free space buffer. The operation may take hours to days, and no other file system operations are supported during that time.

- **Inactive NAS volumes**

The configuration analysis identifies NAS volumes that are not actively used and suggests deleting or archiving the volumes to reduce costs.

Performance efficiency

Performance efficiency helps you ensure that your storage resources are used effectively and that workloads perform optimally.

- **FlexGroup volumes rebalance**

Rebalance your FlexGroup volumes to improve efficiency. After scaling out, capacity and workload activity can become uneven across the volume.

- **FlexVol volumes rebalance**

Rebalance your FlexVol volumes to improve efficiency. After scaling out, capacity and workload activity can become uneven across the volume.

What's next

[Implement well-architected file system configurations](#)

Implement well-architected FSx for ONTAP file system configurations

Use Workload Factory to implement best practices for your FSx for ONTAP file systems. You can easily review the well-architected status, learn about issues with your configurations, and fix issues to improve the architecture of any systems that aren't optimized for reliability, security, efficiency, performance, and cost.

[Learn about the configuration analysis and well-architected status.](#)

About this task

Workload Factory analyzes Amazon FSx for NetApp ONTAP file system deployment configurations daily.

Because requirements for storage configurations vary, you can dismiss the analysis of specific configurations that don't apply to your storage environment.

Link connectivity allows Workload Factory to scan for issues with performance, data protection, and configurations. [Connect to an FSx for ONTAP file system using a link](#) for the most comprehensive analysis of your file system resources.

What is analyzed

Workload Factory analyzes the well-architected status of the following configurations for FSx for ONTAP file systems:

- Reliability: volume backups, snapshot policy, cross-region replication, SSD capacity headroom, replication policy labels, and file object capacity
- Security: Ransomware Protection (ARP/AI) and unauthorized iSCSI volume access
- Operational excellence: automatic capacity management, volume file capacity headroom, FlexCache write mode, FlexCache replica size, host-visible capacity reporting, OS type, block device space management, and thin provisioning
- Cost optimization: cold data tiering, storage efficiencies, backup expiration, snapshot expiration, inactive block devices, overprovisioned SSD capacity, and inactive NAS volumes
- Performance efficiency: FlexGroup volumes rebalance and FlexVol volumes rebalance

Before you begin

- You must [grant operations and remediation permissions](#) in your AWS account.
- [Connect to an FSx for ONTAP file system using a link](#) for the most comprehensive analysis of your file system resources.

Fix a configuration issue

You can fix configuration issues for an FSx for ONTAP file system or for selected volumes in a file system. You can select one or more configurations to fix.

Fix operations can cause instance downtimes or service interruptions. Review each recommendation carefully before fixing a configuration issue.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads** > **Well-architected** > **FSx for ONTAP**.
3. Select **View and fix** for any configuration. Make sure you review the recommendation carefully.

The recommendation explains best practices and potential pitfalls of unoptimized configurations.

4. Select resources and then **Fix**.

Some fix operations are only available for specific operations and not always for more than one resource. For example, you might be able to fix the configuration for a single volume but not for multiple volumes.

5. Review the action items to learn about consequences of this fix such as instance downtimes or service interruptions.
6. Select **Fix** if you want Workload Factory to apply the recommended changes.

Result

The fix process starts. Select the account settings menu and then select **Tracker** to view the status of the operation.

Dismiss a configuration analysis

Stop a configuration analysis indefinitely for one or more resources. You can restart the analysis when needed.

Dismiss to stop configuration analysis for one or more FSx for ONTAP file systems or volumes. You can restart the analysis anytime.

Dismiss a configuration analysis for configurations

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads** > **Well-architected** > **FSx for ONTAP**.
3. Under Configurations, identify the configuration that doesn't apply to your environment.
4. Select  and then select **Dismiss**.
5. In the Dismiss configuration dialog, select **Dismiss** to stop the analysis for the configuration.

Dismiss a configuration analysis for your resources

Steps

1. Log in to the [NetApp Console](#).
- 2.

Select the menu , **Workloads > Well-architected > FSx for ONTAP**.

3. In the FSx for ONTAP well-architected dashboard, under Configurations, select the configuration that doesn't apply to your environment.
4. On the configuration page, select the resource(s) to dismiss from the configuration analysis.
 - For one resource: select  and then select **Dismiss**.
 - For multiple resources: select the resources and then select **Dismiss** next to Bulk action.
5. In the Dismiss resources dialog, select **Dismiss** to stop the analysis.

Result

Workload Factory stops the configuration analysis for selected configurations.

Reactivate a dismissed configuration analysis

Reactivate a dismissed configuration analysis at any time. You can select one or more configurations or resources to reactivate.

Reactivate a configuration analysis for configurations

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Well-architected > FSx for ONTAP**.
3. Select **Dismissed configurations**.
4. Identify the configuration(s) you want to reactivate and select **Reactivate**.

Reactivate a configuration analysis for your resources

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Well-architected > FSx for ONTAP**.
3. In the FSx for ONTAP well-architected dashboard, under Configurations, select the configuration that doesn't apply to your environment.
4. On the configuration page, select **Dismissed configurations**.
5. Select the resource(s) to reactivate.
 - For one resource: select  and then select **Reactivate**.
 - For multiple resources: select the resources and then select **Reactivate** next to Bulk action.

Result

Workload Factory reactivates the configuration analysis and runs a new analysis every day.

Custom rules

Learn about custom rules in NetApp Console

Custom rules let you set, check, and enforce your own architecture standards, along with NetApp best practices, in the Well-architected dashboard.

Overview

The Rule builder lets you create custom best-practice rules using plain language. These rules are checked across all of your resources.

Custom rules help you enforce your organization's standards, risk limits, and operational requirements—things that generic vendor rules might not fully cover.



If your account administrator has turned off AI features, the Rule builder will not be available and you will not be able to create custom rules. [Learn about managing AI features.](#)

How custom rules work

Custom rules use an AI-guided process that turns your plain-language instructions into clear, structured rules. The AI acts as a translator and checker. It suggests a rule, explains how it understood your request, and asks for your approval before the rule is saved or scheduled.

This process ensures that:

- The rule logic is clearly defined and checked for errors
- No free-form code is run
- You can see exactly what will be evaluated before you approve the rule

Limitations

The following limitations apply to custom rules:

- Automatic fixes are not available. You must correct non-compliant resources yourself using the guidance provided.
- Each rule can check only one type of resource. To check multiple resource types, create separate rules.
- The shortest evaluation interval is 1 hour to avoid excessive scanning.
- Rules have read-only access and cannot make changes to your environment.

Rule definition model

Each custom rule consists of the following components:

Rule name

A user-defined name for the rule.

Description

Your rationale for why this rule matters to your organization.

Resource type

The FSx resource to evaluate (file system, volume, or cache). Currently, one resource type per rule.

Scheduler

How often the rule is evaluated (minimum 1 hour, default 24 hours).

Condition logic

The specific criteria to check.

Evaluation scope

Which resources are evaluated (all applicable resources in scope).

Remediation logic

Manual guidance for addressing non-compliant resources.

Rule lifecycle

Custom rules progress through the following states:

Draft

Rule is being created. No evaluation occurs.

Dry-run

One-time evaluation to preview which resources would be affected. Used for validation before scheduling.

Scheduled (Enabled)

Rule is active and runs on the defined schedule.

Disabled

Rule is paused but not deleted. Can be re-enabled.

Deleted

Rule is permanently removed.

Custom rule examples

The following examples illustrate common custom rules:

Production volume protection

"All volumes with tag 'Prod' must have a snapshot policy installed"

This rule automatically creates snapshots to protect production volumes according to company requirements.

Archive volume tiering

"All volumes with tag 'Archive' should have a tiering policy set to 'All'"

This rule lowers costs by automatically moving archived data to cheaper storage.

File system naming standards

"All file systems must follow the naming convention 'fsx-{environment}-{application}'"

This rule validates that file systems adhere to organizational naming conventions for easier management and cost allocation.

What to do next

[Create a custom rule](#)

Create a custom rule in NetApp Workload Factory

Create custom rules by describing them in plain language. The AI rule builder in NetApp Workload Factory turns your request into a validated rule that checks your FSx for ONTAP resources.

About this task

The rule creation process walks you through the steps: defining the rule, reviewing how the AI interpreted it, and checking the results before scheduling regular evaluations. You stay in control at every step. The system suggests a rule based on your prompt, but nothing moves forward until you confirm it.



If your account administrator turned off AI features, you won't be able to create custom rules. Contact your administrator or [learn about AI administration](#).

Limitations

The following limitations apply to custom rules:

- Automatic fixes are not available. You must correct non-compliant resources yourself using the guidance provided.
- Each rule can check only one type of resource. To check multiple resource types, create separate rules.
- The shortest evaluation interval is 1 hour to avoid excessive scanning.
- Rules have read-only access and cannot make changes to your environment.

Custom rule examples

The following examples illustrate common custom rules:

Production volume protection

"All volumes with tag 'Prod' must have a snapshot policy installed"

This rule automatically creates snapshots to protect production volumes according to company requirements.

Archive volume tiering

"All volumes with tag 'Archive' should have a tiering policy set to 'All'"

This rule lowers costs by automatically moving archived data to cheaper storage.

File system naming standards

"All file systems must follow the naming convention 'fsx-{environment}-{application}'"

This rule validates that file systems adhere to organizational naming conventions for easier management and cost allocation.

Before you begin

Before you can create and run custom rules, ensure you have the following:

- You have configured AWS credentials with read-only access. [Add AWS credentials to Workload Factory](#).
- You have ONTAP credentials with viewer permissions. [Add ONTAP credentials to Workload Factory](#).
- You have a link or Console agent configured connectivity to the cloud storage system. [Create and associate a link](#) or [deploy a Console agent](#).
- You understand that custom rule creation uses NetApp's Bedrock AI service.

Create a rule

Use the Rule builder to create custom best-practice rules using plain language. These rules are checked across all of your resources.

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > Well-architected**.
3. From the Well-architected menu, select **Rule builder**.

When you open the rule builder, it shows sample prompts to help you write your custom rule. You can select **Copy prompt** to copy a sample prompt into the AI dialog text box.

4. In the **Create a rule** field, enter your rule description in natural language. Be as specific as possible about:
 - What resource type you're checking (file system, volume, or cache).
 - What condition you're validating.
 - What threshold defines compliance.
 - Why this rule matters (optional but recommended).
5. Select .

The AI analyzes your prompt and generates a rule definition.

Review the AI interpretation

The rule builder displays:

- **Parsed interpretation:** How the AI understood your prompt in plain language.
- **Rule definition:** The structured rule including resource type, condition logic, and evaluation criteria.
- **Reasoning:** Why the AI interpreted your prompt this way.
- **Confidence indicator:** How confident the system is in the interpretation.

Steps

1. Review the interpretation to make sure it matches what you intended.

Check that the rule definition is correct.

If the interpretation is unclear or wrong, the system may ask questions such as:

- * Which resource should be checked?
- * What value or limit defines compliance?

* How often should the rule run?

2. Respond to the AI's clarifying questions.

The AI might ask you to choose from several options to better understand what you need. Answer each question and select **Continue**, or select **Skip** to move past optional questions.

3. After answering all questions, select **Continue** to create your rule.

Test the rule

Before scheduling the rule, run a one-time test to see which resources it will affect.

Steps

1. Select **Validate rule**.

The system evaluates all resources in scope against your rule logic.

2. Review the validation results:

- **Resources evaluated:** Total number of resources checked
- **Compliant:** Resources that meet the rule criteria
- **Non-compliant:** Resources that do not meet the criteria
- **Failed evaluations:** Resources that could not be evaluated (with error details)

3. If some resources failed evaluation, review the error reasons:

- **AWS credential errors:** Read-only permissions may be insufficient for certain resources
- **Connectivity errors:** Link or agent might not be configured for specific file systems
- **ONTAP permission errors:** The viewer user might not have access to certain resources

Save or schedule the rule

After you test your rule, you can save it as a draft or set it to run regularly.

Steps

1. Do one of the following:

- To schedule regular evaluation, set the evaluation frequency (minimum 1 hour, default 24 hours) and select **Schedule**. The rule is enabled and runs on the defined schedule.
- To save without scheduling, select **Save as draft**. The rule is saved but not evaluated until you schedule it.
- To discard the rule, select **Discard** to cancel rule creation and start over.

Result

Your custom rule has been created and is listed in the rule catalog. If you set a schedule, it will check your FSx for ONTAP resources at the specified frequency.

Troubleshooting

No file systems available

Set up your AWS credentials and make sure your environment has at least one FSx for ONTAP file system. The rule builder needs a file system to apply rules.

Low confidence interpretation

More information is needed. Answer the questions or rewrite your prompt with clear details about the resource type, limits, and conditions.

Partial validation failures

Check which resources failed and why. Common causes are missing link or agent settings, or not having enough permissions. Fix the problem, then validate again.

Ambiguous prompt

Be more specific about what you want to check. Include the resource type, the condition, and the threshold. See the prompt examples for help.

Rule targets multiple resource types

Create a separate rule for each resource type. Only one resource type per rule is supported.

What to do next

[Manage custom rules](#)

Manage custom rules for NetApp Workload Factory

The Rules catalog is one place to manage all your custom rules. You can see each rule's details, check its results, and turn it on or off.

View the Rules catalog

The Rules catalog shows all the rules you created, their status, schedule, and latest results.

Steps

1. Log in using one of the [console experiences](#).
2. Select the menu  and then select **Rule builder**.
3. Click **View your rules** to open the Rules catalog.

View rule results

The results show which resources follow your rule and which do not.

Steps

1. In the Rules catalog, click the rule you want to review.
2. Select the **All runs results** tab.
3. Select **View details** to see the results for each run.

If the check failed for some items, read the error message to see why. Common reasons are connection problems, missing permissions, or a link that hasn't been set up.

Edit a custom rule

You can change a rule's logic, schedule, or other settings at any time. When you edit a rule, the rule builder opens with the rule's current settings.

Steps

1. In the Rules catalog, click the rule you want to edit.
2. Click **Edit**.

The rule switches to draft mode and the rule builder opens.

3. Edit the rule prompt or use the conversation to improve the rule.
4. Review the AI's updated interpretation and rule definition.

The system shows what changed from the current rule, including any API calls or logic.

5. Do one of the following depending on the changes made:
 - If minor changes were made, click **Save** to update the rule.
 - If significant changes were made, click **Validate rule** to test the changes, then click **Schedule** to save and schedule the rule, or click **Save as draft** to save without scheduling.
 - To discard the changes, click **Revert to previous version**. The rule returns to its original definition and leaves edit mode.