



Administer Workload Factory

Setup and administration

NetApp

February 02, 2026

This PDF was generated from <https://docs.netapp.com/us-en/workload-setup-admin/log-in.html> on February 02, 2026. Always check docs.netapp.com for the latest.

Table of Contents

Administer Workload Factory	1
Log in to NetApp Workload Factory	1
Manage service accounts	1
Create a service account	2
Delete a service account	3
Build and operate well-architected workloads	3
How it works	3
Why it matters	4
Get started with Workload Factory to detect and correct misconfigurations	4
Best practices and recommendations for storage workloads	4
Best practices and recommendations for database workloads	6
Best practices and recommendations for EVS workloads	10
Related information	11
Configure NetApp Workload Factory notifications	11
Notification types and messages	11
Configure Workload Factory notifications	13
Subscribe to the Amazon SNS topic	14
Filter notifications	14
Automate tasks using Codebox	16
Learn about codebox automation	16
Use Codebox for automation in NetApp Workload Factory	17
Use CloudShell in NetApp Workload Factory	20
About this task	20
CloudShell commands	21
Before you begin	22
Deploy CloudShell	22
Rename a CloudShell session tab	24
Duplicate CloudShell session tab	24
Close CloudShell session tabs	25
Split CloudShell session tabs	25
Update settings for a CloudShell session	25
Remove credentials from NetApp Workload Factory	26

Administer Workload Factory

Log in to NetApp Workload Factory

After you sign up to NetApp Workload Factory, you can log in at any time from the web-based console to start managing your workloads and FSx for ONTAP file systems.

About this task

You can log in to the Workload Factory web-based console using one of the following options:

- Your existing NetApp Support Site (NSS) credentials
- A NetApp cloud login using your email address and a password

Steps

1. Open a web browser and go to the [Workload Factory console](#).
2. On the **Log in** page, enter the email address that's associated with your login.
3. Depending on the authentication method associated with your login, you'll be prompted to enter your credentials:
 - NetApp cloud credentials: Enter your password
 - Federated user: Enter your federated identity credentials
 - NetApp Support Site account: Enter your NetApp Support Site credentials
4. Select **Log in**.

If you have successfully logged in in the past, you'll see the Workload Factory home page and you'll be using the default account.

If this is the first time that you've logged in, you'll be directed to the **Account** page.

- If you are a member of a single account, select **Continue**.
- If you are a member of multiple accounts, select the account and select **Continue**.

Result

You're now logged in and can start using Workload Factory to manage FSx for ONTAP file systems and your workloads.

Manage service accounts

Create service accounts to act as machine users that automate infrastructure operations. You can revoke or change access to service accounts at any time.

About this task

Service accounts are a multi-tenancy functionality provided by NetApp. Account admins create service accounts, control access, and delete service accounts. You can manage service accounts in the NetApp Console or in the NetApp Workload Factory console.

Unlike managing service accounts in the NetApp Console where you can recreate a client secret, Workload Factory supports only creation and deletion of service accounts. If you want to recreate a client secret for a

specific service account in the NetApp Workload Factory console, you'll need to [delete the service account](#), and then [create a new one](#).

Service accounts use a client ID and a secret for authentication rather than a password. Client IDs and secrets are fixed until the account admin decides to change them. To use a service account, you'll need the client ID and secret to generate the access token or you won't gain access. Keep in mind that access tokens are short-lived and can only be used for several hours.

Before you begin

Decide if you want to create a service account in the NetApp console or in the Workload Factory console. There are slight differences. The following instructions describe how to manage service accounts in the Workload Factory console.

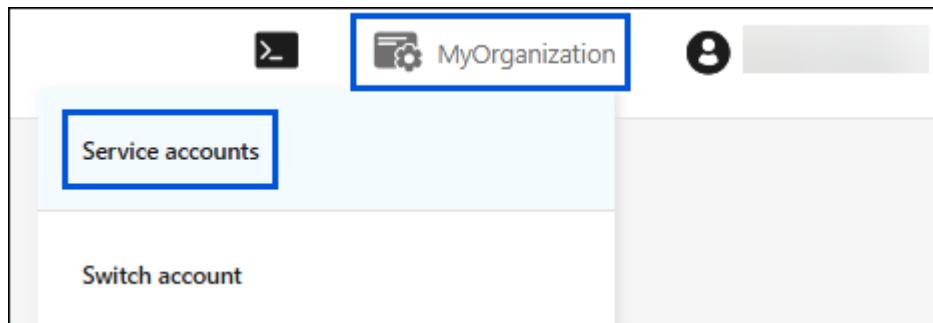
To manage service accounts in the NetApp Console, [learn how identity and access management works](#) and [learn how to add IAM members and manage their permissions](#).

Create a service account

When you create a service account, Workload Factory enables you to copy or download a client ID and client secret for the service account. This key pair is used for authentication with Workload Factory.

Steps

1. In the Workload Factory console, select the **Account** icon, and select **Service accounts**.



2. On the **Service accounts** page, select **Create service account**.
3. In the Create service account dialog, enter a name for the service account in the **Service account name** field.

The **role** is preselected as **account admin**.

4. Select **Continue**.
5. Copy or download the client ID and client secret.

The client secret is visible only once and is not stored anywhere by Workload Factory. Copy or download the secret and store it safely.

6. Optionally, you can get an access token for Auth0 management API by executing a client credentials exchange. The curl example shows how you can take the client ID and secret and use an API to generate the access token which are time-limited. The token provides several hours of access to the NetApp Workload Factory APIs.

7. Select **Close**.

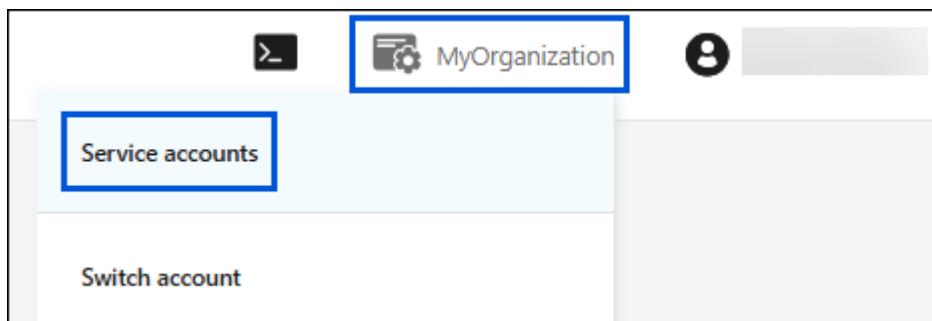
The new service account is created and listed on the Service accounts page.

Delete a service account

Delete a service account if you no longer need to use it.

Steps

1. In the Workload Factory console, select the **Account** icon, and select **Service accounts**.



2. On the **Service accounts** page, select the actions menu and then select **Delete**.
3. In the Delete service account dialog, enter **delete** in the text box.
4. Select **Delete** to confirm deletion.

Build and operate well-architected workloads

Workload Factory, the NetApp management suite for Amazon FSx for NetApp ONTAP, helps you maintain and operate reliable, secure, efficient, and cost-effective storage and database configurations that align with the AWS Well-Architected Framework. Workload Factory provides daily analysis of your storage and database workloads, recommendations, and automatic fixes to promote healthy workload operations. By automating this process, Workload Factory minimizes human error and ensures consistency in workload management.

How it works

Workload Factory analyzes Amazon FSx for NetApp ONTAP file systems, Microsoft SQL Server, and Oracle database deployments daily. The analysis provides well-architected status, insights, and recommendations. You can automatically fix configuration issues to meet best practices and operate efficiently.

After the daily analysis completes, configurations appear as "optimized" or "not optimized" in the Well-architected dashboard for the deployment. You'll find the total optimization score, configuration issues by category, and a list of configuration issues and recommendations. You can review the recommendations for configuration issues. Some issues can be fixed automatically by Workload Factory, while others 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. When you dismiss a specific configuration analysis, Workload Factory does not include the configuration in the total optimization score.

Why it matters

Workload Factory applies best practices to large storage or database environments by combining ongoing assessment with recommendation insights and remediation. Automated fixes reduce human error, ensure uniform management, and maintain performance and reliability. Fixes applied in the Workload Factory console reduce human error and ensure uniform management. Automation ensures configurations are applied correctly and maintained, preserving performance and reliability across your workload infrastructures.

Get started with Workload Factory to detect and correct misconfigurations

To get started with Workload Factory, sign up, add credentials, and establish connectivity to manage AWS resources and optimize workloads using Amazon FSx for NetApp ONTAP.

[Quick start](#)

Best practices and recommendations for storage workloads

Workload Factory assesses storage configurations to provide an in-depth view of ONTAP configuration best practices and for compliance with 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.

- **Schedule FSx for ONTAP backups**

FSx for ONTAP: Backing up your volumes helps support data retention and compliance needs. Use FSx for ONTAP backup to set up automated backups and retention for your data.

- **Schedule local snapshots**

Schedule local snapshots for efficient backup and quick restores. Snapshots are instant, point-in-time images of your volumes.

- **Cross-region replication**

Cross-region replication ensures that your data is replicated to another AWS region, providing enhanced data durability and availability. Workload Factory recommends setting up cross-region replication to help with disaster recovery and compliance.

- **Set up data replication**

To extend data reliability, data can be replicated to an FSx for ONTAP file system in the same region or in another region. Set up data replication to support migration, disaster recovery, and long-term retention across file systems.

- **Increase SSD capacity threshold**

The SSD storage tier capacity should not exceed 80% utilization on an ongoing basis. This might impact data reads and writes to your capacity pool storage tier and impact the throughput capacity of your file

system. Running out of capacity might result in data volumes becoming read-only, and services trying to write new data might fail.

- **Match labels to ensure data reliability**

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

- **Increase file capacity threshold**

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 staying below 80% utilization of the available file capacity on an ongoing basis. Available file capacity is required to create new files in the volume.

Security

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

- **Enable 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 access to volumes**

Volumes serving application data using iSCSI should not allow NAS access in parallel. Workload Factory recommends that volumes accessed via the iSCSI protocol should be restricted to any additional protocols.

Operational excellence

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

- **Enable automatic capacity management**

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

- **Volume capacity utilization threshold**

Workload Factory recommends that volume capacity doesn't exceed 80% utilization on an ongoing basis. This might impact data reads and writes to your application. Volume capacity increases can be manual or automatic using the volume autogrow feature.

- **Volume utilization nearing full**

When a volume is nearing full capacity, Workload Factory recommends taking action to increase the volume capacity to avoid potential application disruptions.

- **Cache relationship write mode**

For optimal performance, Workload Factory recommends the cache relationship write mode that best suits 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.

- **Optimize cache volume size**

Workload Factory recommends enabling volume autosize and scrubbing on cache volumes to maintain optimal size and focus the cache on hot data for peak efficiency.

- **Storage VM logical reporting**

Workload Factory recommends that the default reporting setting is set to logical for a storage VM to provide better visibility into storage usage at the volume level.

Cost optimization

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

- **Optimize TCO by tiering cold data**

Cold data tiering should be enabled to reduce SSD storage tier utilization. Applying a tiering policy to every volume is recommended. FSx for ONTAP scans the data continuously to detect cold data and move it to the capacity storage pool tier without disruption.

- **Enable storage efficiencies**

Storage efficiencies should be enabled - compaction, compression, and deduplication - to optimize storage utilization and reduce the SSD tier cost.

- **Unnecessary snapshot and backup deletion**

Snapshots and backups that are no longer needed should be deleted to reduce costs.

- **Orphaned block devices**

After a block device isn't used for seven days, Workload Factory recommends archiving block device data or deleting the unused block device to reduce costs.

Best practices and recommendations for database workloads

Workload Factory provides a set of best practices and recommendations for operating well-architected database workloads. The well-architected analysis assesses Microsoft SQL Server and Oracle Database configurations and settings related to storage sizing, storage layout, storage configuration, compute, application (SQL Server), and resiliency.

Storage sizing

- **Storage tier**

For best storage performance, create FSx for ONTAP volumes on the primary SSD tier. Using the capacity pool tier might make performance slower and increase latency.

- **File system headroom**

To optimize storage performance, set the file system capacity to 1.35 times the total size of your volumes.

File system headroom percentages are as follows:

- Under-provisioned: < 35%
- Optimized: 35-100%
- Over-provisioned: > 100%

- **Log drive size**

Ensure accurate sizing and regular monitoring of the SQL Server log drive to prevent issues such as transaction rollbacks, database unavailability, data corruption, and performance degradation caused by a full log drive.

Log drive size percentages are as follows:

- Under-provisioned: < 20%
- Optimized: 20-30%
- Over-provisioned: > 30%

- **TempDB drive size**

Ensure accurate sizing and regular monitoring of the SQL Server TempDB to optimize performance and maintain overall stability. Properly configured TempDB prevents performance issues and instability. Insufficient space or high contention can lead to query slowdowns, application timeouts, and system crashes.

TempDB drive size percentages are as follows:

- Under-provisioned: < 10%
- Optimized: 10-20%
- Over-provisioned: > 20%

Storage layout

- **Data files (.mdf) placement**

Separate data and log files onto different drives to improve performance, enable independent backup schedules, and improve restore functionality. For smaller databases, separate data and log LUN paths into different volumes. This separation is required for more than one large database (> 500 GiB).

- **Log files (.ldf) placement**

Separate data and log files onto different drives to improve performance, enable independent backup schedules, and improve restore functionality. For smaller databases, separate data and log LUN paths into different volumes. This separation is required for more than one large database (> 500 GiB).

- **TempDB placement**

Isolate TempDB I/O and avoid I/O contention from other databases by placing TempDB on its own dedicated drive. This optimization improves overall SQL Server performance and stability. Failure to do so can result in significant I/O bottlenecks, slower query performance, and potential system instability.

Storage configuration

- **ONTAP configuration**

Entity	Setting	Recommendation
Volume	<ul style="list-style-type: none"> Thin provisioning (-space-guarantee = none) Autosize on Autosize-mode = grow Fractional reserve = 0% Snapshot copy reserve = 0% Snapshot autodelete (volume/oldest first) Space-mgmt-try-first = volume_grow 	To optimize storage efficiency and cost-effectiveness, configure thin provisioning, autosize, and space management options for your FSx for ONTAP volumes. Without thin provisioning, storage is allocated upfront, leading to inefficient use and higher costs due to over-provisioning; static allocation results in paying for unused capacity, increasing expenses; lack of dynamic allocation hampers scalability and flexibility, impacting performance; and without space reclamation, deleted data occupies space, reducing efficiency.
Volume	<ul style="list-style-type: none"> Tiering-policy = snapshot-only Tiering-minimum-cooling-days = 7 	For optimal database performance and cost efficiency, Workload Factory recommends moving only snapshots to the capacity tier. This strategy ensures high performance while reducing costs. It is especially recommended to tier snapshots that are older than 7 days.
LUN	OS type = windows_2008	ONTAP LUN OS type value should match the operating system partitioning scheme to achieve I/O alignment. Incorrect configuration might result in suboptimal performance.
LUN	Space reservation enabled	When space reservation is enabled, ONTAP reserves enough space in the volume so that writes to those LUNs do not fail because of a lack of disk space.
LUN	Space allocation enabled	This option ensures that FSx for ONTAP notifies the EC2 host when the volume is full and cannot accept writes. This setting also allows FSx for ONTAP to automatically reclaim space when SQL Server on the EC2 host deletes data. If disabled, write failures are possible and space might be inefficiently utilized.

- Windows storage configuration

Entity	Setting	Recommendation
Microsoft Multipath I/O (MPIO)	<ul style="list-style-type: none"> • Status = Enabled • Policy = Round Robin • Number of sessions = 5 	To ensure optimal uptime and data access consistency for Microsoft SQL Server databases on EC2 with underlying LUNs provisioned in FSx for ONTAP, Workload Factory recommends enabling and configuring Multipath I/O (MPIO). MPIO provides multiple paths to FSx for ONTAP, enhancing both resiliency and performance. This best practice protects against potential data loss or downtime by maintaining data access even if a component fails.
Allocation unit size	NTFS allocation unit size = 64K	Set NTFS allocation unit size to 64K to better use disk space, reduce fragmentation, and improve file read/write performance. Failure to configure this properly might lead to inefficient disk usage and degraded performance.

Compute

- **Compute rightsizing**

To ensure optimal performance and cost efficiency for your SQL Server EC2 instance, we recommend rightsizing based on your workload demands. If your current instance is under-provisioned, upgrading will enhance CPU, memory, and I/O capacity. If it is over-provisioned, downgrading will maintain performance while reducing costs.

- **Operating system patch**

Workload Factory recommends applying the latest patches to ensure security, protect SQL Server databases from vulnerabilities, and improve system reliability.

- **Network adapter settings**

Accurate configuration of receive side scaling (RSS) is essential for optimal network performance in Microsoft SQL Server instances. RSS distributes network processing across multiple processors, preventing bottlenecks and enhancing system performance. Workload Factory recommends the following RSS settings:

- Disable TCP Offloading Features: Ensure all TCP offloading features are disabled.
- Number of Receive Queues: Set to 8 if vCPUs > 8. Set to the number of vCPUs if vCPUs ≤ 8.
- RSS Profile: Set to NUMAStatic.
- Base Processor Number: Set to 2.

Following these settings will improve the performance and reliability of your Microsoft SQL Server instances. We suggest that you test the recommended settings to determine performance improvements before making changes to your production environment.

Application (SQL Server)

- **License**

The SQL Server license assessment and recommendation are provided at the host level.

Not optimized: A license is considered "not optimized" when Workload Factory detects that your database infrastructure doesn't use any of the commercial software license features you're paying for. An unoptimized license might result in unnecessary costs.

Optimized: A license is considered "optimized" when the commercial software license for your databases meets your performance requirements.

- **Microsoft SQL Server patch**

Workload Factory recommends applying the latest patches to ensure security, protect SQL Server databases from vulnerabilities, and improve system reliability.

- **MAXDOP**

Set the Maximum Degree of Parallelism (MAXDOP) to optimize query performance by balancing parallel processing. Accurate MAXDOP configuration enhances performance and efficiency. Setting MAXDOP to 4, 8, or 16 generally provides the best results in most use cases. We recommend that you test your workload and monitor for any parallelism-related wait types such as CXPACKET.

Reliability

- **Schedule FSx for ONTAP backups**

Backing up your Microsoft SQL Server volumes is crucial for supporting your data retention and compliance requirements. Use FSx for ONTAP backup to set up automatic backups and retention for your SQL Server data.

- **Schedule local snapshots**

Schedule local snapshots for efficient backup and quick restores. Snapshots are instant, point-in-time images of your volumes.

- **Cross-region replication**

Cross-region replication ensures that your data is replicated to another AWS region, providing enhanced data durability and availability. Workload Factory recommends setting up cross-region replication to help with disaster recovery and compliance.

Best practices and recommendations for EVS workloads

Workload Factory provides best practices and recommendations for operating well-architected Amazon Elastic VMware Service (EVS) workloads. The well-architected analysis assesses EVS configurations to help ensure your VMware environments are optimized for reliability, security, operational excellence, cost optimization, and performance efficiency. From the well-architected status tab in VMware, you'll find insights and recommendations to help implement well-architected best practices for your EVS environments.

The well-architected analysis categorizes configurations in the following pillars of the framework: *reliability* and *security*.

Reliability

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

- **EVS environment resiliency**

Ensure that your EVS cluster nodes are properly distributed across partition placement groups. All nodes should be members of a single partition placement group configured with four or more partitions. Proper partition placement ensures that your EVS cluster nodes are distributed across multiple fault-isolated hardware partitions within an AWS availability zone. Misalignment can result in significant loss of processing power or downtime if a partition fails.

Security

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

- **Cluster node management**

Ensure that your EVS cluster nodes have appropriate EC2 stop and termination protection configured. EVS ESXi nodes should be managed exclusively using vCenter or other VMware-level management tools. Without proper EC2-level protections, nodes could be accidentally stopped or terminated from the EC2 console, which can lead to virtual machine data unavailability or data loss.

Related information

- [Implement well-architected FSx for ONTAP file systems](#)
- [Implement well-architected database workloads](#)
- [Implement well-architected EVS configurations](#)

Configure NetApp Workload Factory notifications

You can configure the NetApp Workload Factory notification service to send notifications as alerts in the NetApp Console or to an Amazon SNS topic. Notifications sent as alerts appear in NetApp Console when you have a Console agent or link deployed. When Workload Factory publishes notifications to an Amazon SNS topic, subscribers to the topic (such as people or other applications) receive the notifications at the endpoints configured for the topic (such as email or SMS messages).

Notification types and messages

Workload Factory sends notifications for the following events:

Event	Description	Notification type	Severity	Workload	Resource type
Some database instances in your account are not well-architected	All Microsoft SQL Server instances in your account have been analyzed for well-architected issues. The description for this event gives the number of well-architected instances and unoptimized instances. Review well-architected status findings and recommendations in the Databases inventory from the Workload Factory console.	Well-architected	Recommendation	Databases	Microsoft SQL Server instance
Microsoft SQL Server/PostgreSQL server deployment succeeded	The deployment of the Microsoft SQL Server or PostgreSQL host succeeded. For more information, go to job monitoring.	Deployment	Success	Databases	FSx for ONTAP, DB host
Microsoft SQL Server/PostgreSQL server deployment failed	The deployment of the Microsoft SQL Server or PostgreSQL host failed. For more information, go to job monitoring.	Deployment	Error	Databases	FSx for ONTAP, DB host
Failed replication relationship creation	A SnapMirror replication relationship creation has failed. For more information, go to Tracker.	Replication	Critical	General storage	FSx for ONTAP
FSx for ONTAP creation failure	An FSx for ONTAP file system creation process has failed. For more information, go to Tracker.	FSx for ONTAP file system action	Critical	General storage	FSx for ONTAP
Automatic SSD capacity or inodes increase success	During a recent automatic capacity management update, the FSx for ONTAP file system either increased SSD capacity or volume inodes successfully. For more information, go to Tracker.	Capacity management	Success	General storage	FSx for ONTAP file
Automatic SSD capacity or inodes increase failure	During a recent automatic capacity management update, the FSx for ONTAP file system failed to increase SSD capacity or volume inodes. For more information, go to Tracker.	Capacity management	Critical	General storage	FSx for ONTAP file systems
FSx for ONTAP issue detected	All FSx for ONTAP file systems have been analyzed for well-architected issues. The scan detected one or more issues. For more information, review the well-architected analysis from the Storage dashboard in the Workload Factory console.	Well-architected analysis	Recommendation	General storage	FSx for ONTAP file systems

Event	Description	Notification type	Severity	Workload	Resource type
Automatic capacity management event for FSx for ONTAP	The SSD performance tier for the FSx for ONTAP file system reached the warning threshold capacity/percentage total.	Capacity management	Warning	General storage	FSx for ONTAP file systems
Automatic inodes management event for FSx for ONTAP	The inode count for the FSx for ONTAP volume reached the warning threshold count/percentage total.	Capacity management	Warning	General storage	FSx for ONTAP file systems

Configure Workload Factory notifications

Configure Workload Factory notifications using the NetApp Console or the Workload Factory console. If you use the NetApp Console, you can configure Workload Factory to send notifications as alerts in the NetApp Console or to an Amazon SNS topic. You can configure notifications from the **Notifications settings** in the NetApp Console.

Before you begin

- You need to configure Amazon SNS and create Amazon SNS topics using either the Amazon SNS console or the AWS CLI.
- Note that Workload Factory supports the **Standard** topic type. This type of topic does not ensure that notifications are sent to subscribers in the order in which they were received, so consider this if you have critical or emergency notifications.

Configure notifications from NetApp Console

Steps

1. Log in to the [NetApp Console](#).
2. From the NetApp Console menu, select **Workloads, Administration**, and then **Notifications setup**.
3. On the Notifications setup page, do the following:
 - a. Optional: Select **Enable NetApp Console notifications** to configure Workload Factory to send notifications in the NetApp Console.
 - b. Select **Enable SNS notifications**.
 - c. Follow the instructions to configure Amazon SNS from the Amazon SNS console.

After you create the topic, copy the topic ARN and enter it in the **SNS topic ARN** field on the **Notifications setup** page.

4. After you verify the configuration by sending a test notification, select **Apply**.

Result

Workload Factory is configured to send notifications to the Amazon SNS topic that you specified.

Configure notifications from Workload Factory console

Steps

1. Log in to the [Workload Factory console](#).
2. From the Workload Factory console menu, select **Workloads, Administration**, and then **Notifications setup**.
3. Select **Enable SNS notifications**.
4. Follow the instructions to configure Amazon SNS from the Amazon SNS console.
5. After you verify the configuration by sending a test notification, select **Apply**.

Result

Workload Factory is configured to send notifications to the Amazon SNS topic that you specified.

Subscribe to the Amazon SNS topic

After you configure Workload Factory to send notifications to a topic, follow the [Instructions](#) in the Amazon SNS documentation to subscribe to the topic so that you can receive notifications from Workload Factory.

Filter notifications

You can reduce unnecessary notification traffic and target specific notification types for specific users by applying filters to the notifications. You can do this using an Amazon SNS policy for SNS notifications, and using the notifications settings in the NetApp Console.

Filter Amazon SNS notifications

When you subscribe to an Amazon SNS topic, you receive all notifications published to that topic by default. If you want to receive only specific notifications from the topic, you can use a filter policy to control which notifications you receive. Filter policies cause Amazon SNS to deliver only the notifications that match the filter policy to the subscriber.

You can filter Amazon SNS notifications by the following criteria:

Description	Filter policy field name	Possible values
Resource type	resourceType	<ul style="list-style-type: none">• DB• Microsoft SQL Server host• PostgreSQL Server host
Workload	workload	WLMDB
Priority	priority	<ul style="list-style-type: none">• Success• Info• Recommendation• Warning• Error• Critical
Notification type	notificationType	<ul style="list-style-type: none">• Deployment• Well-architected

Steps

1. In the Amazon SNS console, edit the subscription details for the SNS topic.
2. In the **Subscription filter policy** area, select to filter by **Message attributes**.
3. Enable the **Subscription filter policy** option.
4. Enter a JSON filter policy in the **JSON editor** box.

For example, the following JSON filter policy accepts notifications from the Microsoft SQL Server resource that are related to the WLMDB workload, have a priority of Success or Error, and provide details on Well-architected status:

```
{  
  "accountId": [  
    "account-a"  
  ],  
  "resourceType": [  
    "Microsoft SQL Server host"  
  ],  
  "workload": [  
    "WLMDB"  
  ],  
  "priority": [  
    "Success",  
    "Error"  
  ],  
  "notificationType": [  
    "Well-architected"  
  ]  
}
```

5. Select **Save changes**.

For other examples of filter policies, refer to [Amazon SNS example filter policies](#).

For further information about creating filter policies, refer to the [Amazon SNS documentation](#).

Filter notifications in the NetApp Console

You can use the NetApp Console notifications settings to filter notifications that you receive in the Console by severity level, such as Critical, Info, or Warning.

For more information about filtering notifications in the Console, refer to the [NetApp Console documentation](#).

Automate tasks using Codebox

Learn about codebox automation

Codebox is an Infrastructure as Code (IaC) co-pilot that helps developers and DevOps generate the code needed to execute any operation supported by NetApp Workload Factory. Codebox is aligned with the Workload Factory permission policies and it sets a clear path for execution readiness as well as providing an automation catalog for quick future reuse.

Codebox capabilities

Codebox provides two key IaC capabilities:

- *Codebox Viewer* shows the IaC that is generated by a specific job flow operation by matching entries and selections from the graphical wizard or from the conversational chat interface. While Codebox Viewer

supports color coding for easy navigation and analysis, it does not allow editing—only copying or saving code to the Automation Catalog.

- *Codebox Automation Catalog* shows all saved IaC jobs, allowing you to easily reference them for future use. Automation catalog jobs are saved as templates and shown in context of the resources that apply to them.

Additionally, when setting up Workload Factory credentials, Codebox dynamically displays the AWS permissions that are needed to create IAM policies. The permissions are provided for each Workload Factory capability that you plan to use (databases, AI, FSx for ONTAP, and so on), and they are customizable. You just copy the permissions from Codebox and then paste them in the AWS Management Console so that Workload Factory has the correct permissions to manage your workloads.

Supported code formats

The supported code formats include:

- Workload Factory REST APIs
- AWS CLI
- AWS CloudFormation
- Terraform

Related information

[Learn how to use Codebox](#).

[Workload Factory REST API documentation](#).

Use Codebox for automation in NetApp Workload Factory

You can use Codebox to generate the code needed to execute any operation supported by NetApp Workload Factory. You can generate code that can be consumed and run using Workload Factory REST APIs, the AWS CLI, and AWS CloudFormation.

Codebox is aligned with the Workload Factory permission policies by populating the appropriate data in the code based on the AWS permissions provided in the Workload Factory account for each user. The code can be used like a template where you can fill in missing information (for example, credentials) or customize certain data before running the code.

How to use Codebox

As you enter values in the Workload Factory UI wizards, you can see the data update in Codebox as you complete each field. When you complete the wizard, but before you select the **Create** button at the bottom of the page, select  to copy in Codebox to capture the code required to build your configuration. For example, this screenshot from creating a new Microsoft SQL Server shows the wizard entries for VPC and availability zones and the equivalent entries in Codebox for a REST API implementation.

Region & VPC us-east-1 | US East (N. Virginia) VPC-1 | 172.30.0.0/20

Availability zones

Select an Availability Zone and subnet for each node. Ensure that each of the selected private subnets have outbound connectivity enabled. The subnet for node 1 must be in the primary Availability Zone for the FSx for ONTAP file system.

Cluster configuration - Node 1: Availability zone Subnet

us-east-1d HCL-CC-1 | 192.168.16.0/24

Cluster configuration - Node 2: Availability zone Subnet

us-east-2d HCL-CC-2 | 192.168.17.0/24

Security group Use an existing security group sg-ad2b38d1

Codebox

Create database

REST API

```
curl --location --request POST https://api.workloads.netapp.com/accounts/acc
--header 'Authorization: Bearer <Token> \
--header 'Content-Type: application/json' \
--data-raw '{
  "networkConfiguration": {
    "vpclId": "vpcl-7d4a2818",
    "vpcCidr": "172.30.0.0/20",
    "availabilityZone1": "us-east-1d",
    "privateSubnet1Id": "subnet-5a37222d",
    "routeTable1Id": "rtb-0de1132a1c54f5e6",
    "availabilityZone2": "us-east-2d",
    "privateSubnet2Id": "subnet-74a1b303",
    "routeTable2Id": "rtb-00d7acd615fac5414",
  },
  "ec2Configuration": {
    "workloadInstanceType": "m5.xlarge",
    "keyPairName": "Key-Pair-1",
  }
}'
```

Copy

With some code formats you can also select the download button to save the code in a file that you can bring to another system. If required, you can edit the code after it has been downloaded so that you can adapt it to other AWS accounts.

Use CloudFormation code from Codebox

You can copy the CloudFormation code generated from Codebox and then launch the Amazon Web Services CloudFormation stack in your AWS account. CloudFormation will perform the actions that you defined in the Workload Factory UI.

The steps to use the CloudFormation code might be different depending on whether you are deploying an FSx for ONTAP file system, creating account credentials, or performing other Workload Factory actions.

Note that the code within a CloudFormation-generated YAML file expires after 7 days for security reasons.

Before you begin

- You'll need to have credentials to log in to your AWS account.
- You'll need to have the following user permissions to use a CloudFormation stack:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "cloudformation>CreateStack",
        "cloudformation>UpdateStack",
        "cloudformation>DeleteStack",
        "cloudformation>DescribeStacks",
        "cloudformation>DescribeStackEvents",
        "cloudformation>DescribeChangeSet",
        "cloudformation>ExecuteChangeSet",
        "cloudformation>ListStacks",
        "cloudformation>ListStackResources",
        "cloudformation>GetTemplate",
        "cloudformation>ValidateTemplate",
        "lambda>InvokeFunction",
        "iam>PassRole",
        "iam>CreateRole",
        "iam>UpdateAssumeRolePolicy",
        "iam>AttachRolePolicy",
        "iam>CreateServiceLinkedRole"
      ],
      "Resource": "*"
    }
  ]
}

```

Steps

1. After you have used the Workload Factory console to define the operation that you want to perform, copy the code in the Codebox.
2. Select **Redirect to CloudFormation** and the Redirect to CloudFormation page is displayed.
3. Open another browser window and log in to the AWS Management Console.
4. Select **Continue** from the Redirect to CloudFormation page.
5. Log in to the AWS account where the code should be run.
6. On the Quick create stack page, under Capabilities, select **I acknowledge that AWS CloudFormation might**
7. Select **Create stack**.
8. Monitor the progress from AWS or from Workload Factory.

Use REST API code from Codebox

You can use the Workload Factory REST APIs generated from Codebox to deploy and manage your FSx for ONTAP file systems and other AWS resources.

You can run the APIs from any host that supports curl and that has internet connectivity.

Note that the authentication tokens are hidden in Codebox, but they are populated when you copy and paste the API call.

Steps

1. After you have used the Workload Factory console to define the operation that you want to perform, copy the API code in the Codebox.
2. Paste the code and run it on your host system.

Use AWS CLI code from Codebox

You can use the Amazon Web Services CLI generated from Codebox to deploy and manage your FSx for ONTAP file systems and other AWS resources.

Steps

1. After you have used the Workload Factory console to define the operation that you want to perform, copy the AWS CLI in the Codebox.
2. Open another browser window and log in to the AWS Management Console.
3. Paste the code and run it.

Use Terraform from Codebox

You can use Terraform to deploy and manage your FSx for ONTAP file systems and other AWS resources.

Before you begin

- You'll need a system where Terraform is installed (Windows/Mac/Linux).
- You'll need to have credentials to log in to your AWS account.

Steps

1. After you have used the Workload Factory console to define the operation that you want to perform, download the Terraform code from the Codebox.
2. Copy the downloaded script archive to the system where Terraform is installed.
3. Extract the zip file and follow the steps in the README.md file.

Use CloudShell in NetApp Workload Factory

Open CloudShell to execute AWS or ONTAP CLI commands from anywhere in the NetApp Workload Factory console.

About this task

CloudShell allows you to execute AWS CLI commands or ONTAP CLI commands in a shell-like environment from within the Workload Factory console. It simulates terminal sessions in the browser, providing terminal features and proxying messages through Workload Factory's backend. It allows you to use the AWS

credentials and ONTAP credentials that you have provided in your NetApp account.

CloudShell features include:

- Multiple CloudShell sessions: deploy multiple CloudShell sessions at one time to issue several sequences of commands in parallel,
- Multiple views: split CloudShell tab sessions so you can view two or more tabs horizontally or vertically at the same time
- Session renaming: rename sessions as needed
- Last session content persistence: re-open the last session if you close it by mistake
- Settings preferences: change the font size and output type
- AI-generated error responses for ONTAP CLI commands
- Autocomplete support: start typing a command and use the **Tab** key to view available options

CloudShell commands

Within the CloudShell GUI interface, you can enter `help` to view available CloudShell commands. After you issue the `help` command, the following reference appears.

Description

NetApp CloudShell is a GUI interface built into NetApp Workload Factory enables you to execute AWS CLI commands or ONTAP CLI commands in a shell-like environment. It simulates terminal sessions in the browser, providing terminal features and proxying messages through the backend in Workload Factory. It enables you to use the AWS credentials and ONTAP credentials that you have provided in your NetApp account.

Available commands

- `clear`
- `help`
- `[--fsx <fsxId>] <ontap-command> [parameters]`
- `aws <aws-command> <aws-sub-command> [parameters]`

Context

Each terminal session runs in a specific context: credentials, region, and optionally FSx for ONTAP file system.

+

All AWS commands execute in the provided context. AWS commands will only succeed if the provided credentials have permissions in the specified region.

+

You can specify ONTAP commands with an optional `fsxId`. If you provide an `fsxId` with an individual ONTAP command, then this ID overrides the ID in the context. If the terminal session doesn't have an FSx for ONTAP file system ID context, then you must provide `fsxId` with each ONTAP command.

+

To update different context specifics, do the following:

- * To change credentials: "using credentials <credentialId>"
- * To change region: "using region <regionCode>"
- * To change FSx for ONTAP file system: "using fsx <fileSystemId>"

Showing Items

- To show available credentials: "show credentials"
- To show available regions: "show regions"
- To show command history: "show history"

Variables

The following are examples of setting and using variables. If a variable value contains spaces, you should set it inside quotes.

+

- * To set a variable: \$<variable> = <value>
- * To use a variable: \$<variable>
- * Example setting a variable: \$svm1 = svm123
- * Example using a variable: --fsx FileSystem-1 volumes show --vserver \$svm1
- * Example setting a variable with string value \$comment1 = "A comment with spaces"

Operators

Shell operators such as pipe |, background execution &, and redirection > aren't supported. Command execution fails if you include these operators.

Before you begin

CloudShell works in the context of your AWS credentials. To use CloudShell, you must provide at least one AWS credential.



CloudShell is available for you to execute any AWS or ONTAP CLI command. However, if you want to work within the context of an FSx for ONTAP file system, make sure you issue the following command: `using fsx <file-system-name>`.

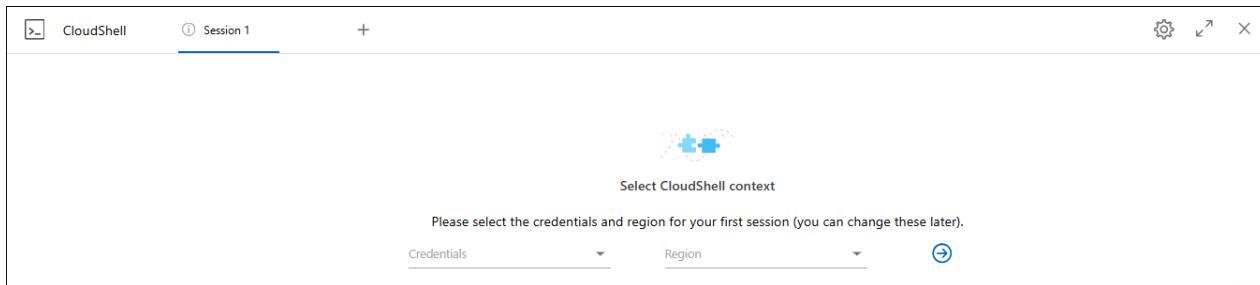
Deploy CloudShell

You can deploy CloudShell from anywhere in the NetApp Workload Factory console. You can also deploy CloudShell from the NetApp Console.

Deploy from Workload Factory console

Steps

1. Log in to the [Workload Factory console](#).
2. From the menu, select **Administration** and then **CloudShell**.
3. In the CloudShell window, select credentials and region for the CloudShell session and then select the arrow to continue.



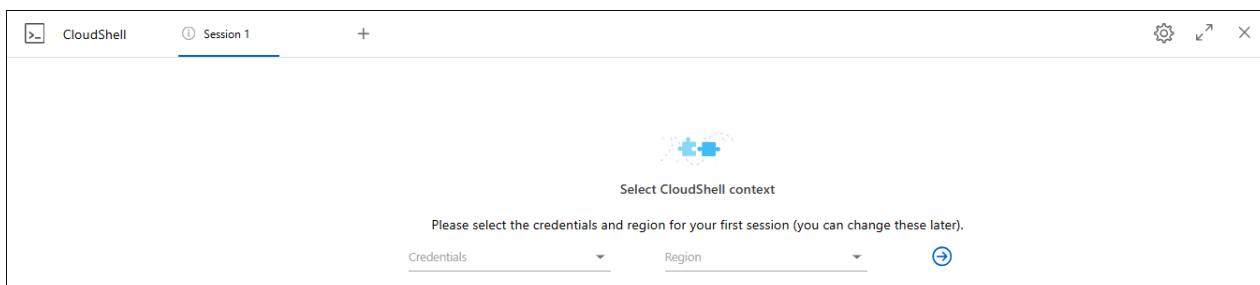
4. Enter `help` to view available [CloudShell commands](#) and instructions or refer to the following CLI reference documents for available commands:
 - [AWS CLI reference](#): For commands related to FSx for ONTAP, select **fsx**.
 - [ONTAP CLI reference](#)
5. Issue commands within the CloudShell session.

If an error occurs after issuing an ONTAP CLI command, select the light bulb icon to get a brief AI-generated error response with a description of the failure, the cause of the failure, and a detailed resolution. Select **Read more** for more details.

Deploy from the NetApp Console

Steps

1. Log in to the [NetApp Console](#).
2. From the menu, select **Workloads** and then **Administration**.
3. From the Administration menu, select **CloudShell**.
4. In the CloudShell window, select credentials and region for the CloudShell session and then select the arrow to continue.



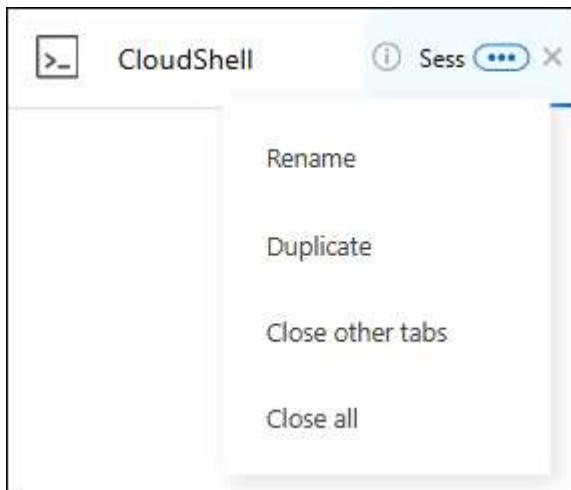
5. Enter `help` to view available CloudShell commands and instructions or refer to the following CLI reference documents for available commands:
 - [AWS CLI reference](#): For commands related to FSx for ONTAP, select **fsx**.

- [ONTAP CLI reference](#)

6. Issue commands within the CloudShell session.

If an error occurs after issuing an ONTAP CLI command, select the light bulb icon to get a brief AI-generated error response with a description of the failure, the cause of the failure, and a detailed resolution. Select **Read more** for more details.

The CloudShell tasks shown in this screenshot can be completed by selecting the actions menu of an open CloudShell session tab. The instructions for each of these tasks follows.



Rename a CloudShell session tab

You can rename a CloudShell session tab to help you identify the session.

Steps

1. Select the actions menu of the CloudShell session tab.
2. Select **Rename**.
3. Enter a new name for the session tab and then click outside the tab name to set the new name.

Result

The new name appears in the CloudShell session tab.

Duplicate CloudShell session tab

You can duplicate a CloudShell session tab to create a new session with the same name, credentials, and region. The code from the original tab isn't duplicated in the duplicated tab.

Steps

1. Select the actions menu of the CloudShell session tab.
2. Select **Duplicate**.

Result

The new tab appears with the same name as the original tab.

Close CloudShell session tabs

You can close CloudShell tabs one at a time, close other tabs you're not working on, or close all tabs at once.

Steps

1. Select the actions menu of the CloudShell session tab.
2. Select one of the following:
 - Select "X" in the CloudShell tab window to close one tab at a time.
 - Select **Close other tabs** to close all other tabs that are open except the one you're working on.
 - Select **Close all tabs** to close all tabs.

Result

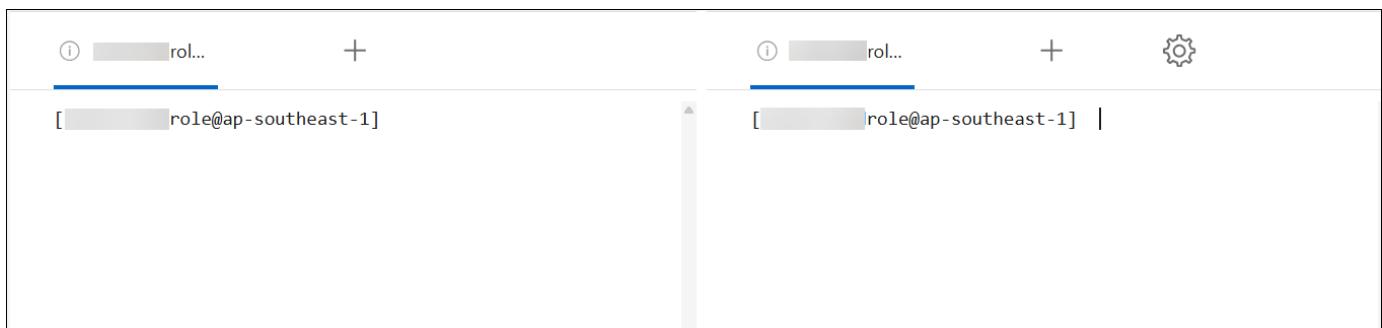
The selected CloudShell session tabs close.

Split CloudShell session tabs

You can split CloudShell session tabs to view two or more tabs at the same time.

Step

Drag and drop CloudShell session tabs to the top, bottom, left, or right of the CloudShell window to split the view.



Update settings for a CloudShell session

You can update font and output type settings for CloudShell sessions.

Steps

1. Deploy a CloudShell session.
2. In the CloudShell tab, select the settings icon.
The settings dialog appears.
3. Update font size and output type as needed.



Enriched output applies to JSON objects and table formatting. All other output appears as plain text.

4. Select **Apply**.

Result

The CloudShell settings are updated.

Remove credentials from NetApp Workload Factory

If you no longer need a set of credentials, you can delete them from Workload Factory. You can only delete credentials that aren't associated with an FSx for ONTAP file system.

Steps

1. Log in using one of the [console experiences](#).
2. From the menu, select **Administration** and then **Credentials**.
3. On the **Credentials** page, do the following:
 - In the Workload Factory console, select the action menu for a set of credentials and then select **Remove**. Select **Remove** to confirm.
 - In the NetApp Console, select the action menu for a set of credentials and then select **Delete**. Select **Delete** to confirm.

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—with 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.