



# **NetApp Workload Factory for Databases documentation**

## **Database workloads**

NetApp

February 04, 2026

# Table of Contents

NetApp Workload Factory for Databases documentation .....	1
Release notes .....	2
What's new with NetApp Workload Factory for Databases .....	2
02 February 2026 .....	2
04 January 2026 .....	2
18 December 2025 .....	2
27 November 2025 .....	3
02 November 2025 .....	3
06 October 2026 .....	5
1 September 2025 .....	5
12 August 2025 .....	6
04 August 2025 .....	7
30 June 2025 .....	7
03 June 2025 .....	8
04 May 2025 .....	9
04 April 2025 .....	10
03 March 2025 .....	10
03 February 2025 .....	11
06 January 2025 .....	12
01 December 2024 .....	12
3 November 2024 .....	13
29 September 2024 .....	13
1 September 2024 .....	13
4 August 2024 .....	14
7 July 2024 .....	15
Known limitations for NetApp Workload Factory for Databases .....	15
Instance detection support .....	15
AMI image version support .....	15
Custom AMI .....	15
Roll back and retry for failed deployments .....	15
Active Directory and DNS resource rollback .....	15
Always On availability groups configuration support .....	16
Custom encryption keys .....	16
CloudFormation template .....	16
Sandbox support .....	16
Microsoft SQL Server detection and management .....	16
Explore savings .....	16
Multiple FSx for ONTAP file systems .....	16
Optimization limitations .....	16
On-premises savings calculation .....	17
Cross-region replication assessment .....	17
Database host authentication when exploring savings .....	17
Integration with NetApp Backup and Recovery .....	17

Region support . . . . .	17
Registering Oracle on SUSE Linux Enterprise Server 12 . . . . .	17
Get started . . . . .	18
Learn about NetApp Workload Factory for Databases . . . . .	18
What is NetApp Workload Factory for Databases? . . . . .	18
Workload Factory benefits for self-managed databases . . . . .	18
Benefits of Amazon FSx for ONTAP for self-managed databases . . . . .	19
Tools to use NetApp Workload Factory . . . . .	19
Supported configurations . . . . .	20
Integrated AWS services . . . . .	21
Regions . . . . .	21
Getting help . . . . .	21
Quick start for NetApp Workload Factory for Databases . . . . .	21
Use Database workloads . . . . .	23
Explore savings in NetApp Workload Factory for Databases . . . . .	23
About this task . . . . .	23
SQL Server deployment analysis . . . . .	24
Calculator options . . . . .	25
Deploy Microsoft SQL Server on AWS EC2 using FSx for ONTAP . . . . .	31
Create a new database server . . . . .	32
Create a Microsoft SQL Server in Workload Factory for Databases . . . . .	32
Create a PostgreSQL server in NetApp Workload Factory . . . . .	40
Manage resources . . . . .	46
Resource management in NetApp Workload Factory for Databases . . . . .	46
Register resources in NetApp Workload Factory for Databases . . . . .	47
Create a Microsoft SQL database in NetApp Workload Factory for Databases . . . . .	50
Create a sandbox clone in NetApp Workload Factory for Databases . . . . .	52
Automate with Codebox in NetApp Workload Factory for Databases . . . . .	53
Protect Microsoft SQL Server workloads . . . . .	53
About this task . . . . .	53
Before you begin . . . . .	54
Prepare for protection with NetApp Backup and Recovery . . . . .	54
Edit protection for Microsoft SQL Server resources . . . . .	55
Administer and monitor . . . . .	56
Monitor database jobs in Workload Factory for Databases . . . . .	56
About this task . . . . .	56
Monitor jobs . . . . .	56
Implement configuration best practices . . . . .	56
Configuration analysis for database environments in Workload Factory . . . . .	56
Implement well-architected database configurations in Workload Factory . . . . .	58
Analyze error logs in Workload Factory . . . . .	62
About this task . . . . .	62
Before you begin . . . . .	63
Analyze error logs . . . . .	63
Manage clones . . . . .	64

Check the integrity of the data in a sandbox clone .....	64
Revert a database clone in NetApp Workload Factory for Databases .....	65
Refresh a database clone in NetApp Workload Factory for Databases .....	65
Connect a sandbox clone to CI/CD tools .....	66
View connection information of a database clone .....	66
Split a database clone from the source database .....	66
Delete a database clone in NetApp Workload Factory for Databases .....	67
Deregister a resource in NetApp Workload Factory .....	67
<b>Knowledge and support .....</b>	<b>68</b>
<b>Register for support .....</b>	<b>68</b>
Support registration overview .....	68
Register your account for NetApp support .....	68
<b>Get help for NetApp Workload Factory for Databases .....</b>	<b>70</b>
Get support for FSx for ONTAP .....	70
Use self-support options .....	70
Create a case with NetApp support .....	70
Manage your support cases (Preview) .....	73
<b>Legal notices .....</b>	<b>76</b>
<b>Copyright .....</b>	<b>76</b>
<b>Trademarks .....</b>	<b>76</b>
<b>Patents .....</b>	<b>76</b>
<b>Privacy policy .....</b>	<b>76</b>
<b>Open source .....</b>	<b>76</b>

# NetApp Workload Factory for Databases documentation

# Release notes

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

Learn what's new with NetApp Workload Factory for Databases.

### 02 February 2026

#### Improved cost savings for several database hosts on Microsoft SQL Server on-premises

The cost savings analysis for Microsoft SQL Server on-premises storage has been enhanced to provide a consolidated storage solution for several database hosts on one FSx for ONTAP file system. This consolidation optimizes storage costs by reducing the number of file systems required for multiple database hosts, leading to improved cost savings.

[Explore savings for detected hosts on Microsoft SQL Server on-premises](#)

#### Well-architected analysis for Oracle

The well-architected analysis includes four new sub assessments for the storage configuration analysis for Oracle: dNFS enablement, dNFS consistent IP resolution, dNFS configuration file, and dNFS nosharecache. These assessments relate to enabling and setting up Direct NFS (dNFS) for your Oracle environment, which bypasses the host NFS client and performs NFS file operations directly on an NFS server, improving I/O performance and decreasing the load on the host and the storage system because I/O is performed more efficiently.

[Implement well-architected database configurations in Workload Factory](#)

### 04 January 2026

#### 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.

### 18 December 2025

#### Active Directory integration enhancements

Workload Factory for Databases includes three new Active Directory (AD) fields when deploying Microsoft SQL Server using the **Advanced create** option. These enhancements let you to specify Active Directory join preferences and use a managed service account.

The new AD fields are:

- Preferred Domain Controller
- Preferred Organizational Unit Path
- Target Active Directory Group

## 27 November 2025

### Optimize savings in the Databases calculator for Amazon Elastic Block Store (EBS)

Two new calculator features improve the cost savings analysis when running multiple instances with EBS storage so that you can save more by switching to FSx for ONTAP.

- Workload Factory provides a consolidated storage solution for several database hosts on one FSx for ONTAP file system. This consolidation optimizes storage costs by reducing the number of file systems required for multiple database hosts, leading to improved cost savings.
- Workload Factory analyzes your EBS performance usage and then suggests the best and most cost-efficient FSx for ONTAP configuration.

[Explore savings for detected EBS hosts](#)

### Excel report available for the well-architected dashboard

You can download an Excel report of the well-architected dashboard. The report provides the well-architected status for the database resource and recommendations for all resource configurations, including sub-configurations for operating system and ONTAP.

### Error log analyzer available for Oracle databases

The Agentic AI-powered error log analyzer is available for Oracle databases. The feature leverages advanced machine learning algorithms to automatically detect and analyze errors in log files. This tool aims to streamline the troubleshooting process by providing developers with actionable insights and recommendations based on the patterns it identifies in the logs.

[Learn more about the Agentic AI-powered error log analyzer](#)

### Well-architected analysis for Oracle

The well-architected analysis includes two new storage sizing configurations. The analysis assesses and fixes configuration issues related to swap space allocation and file system headroom for existing Oracle database deployments.

[Implement well-architected database configurations in Workload Factory](#)

## 02 November 2025

### Dashboard improvements

The new dashboard is streamlined and better organized to provide a clearer overview of your database resources and the key functionalities of Workload Factory on one screen; one card highlights database resources, one card provides information about the well-architected score and analysis, one card displays error analysis, two cards display monthly cost and potential savings, and one card displays information on sandboxes.

### Well-architected analysis for Oracle

The well-architected analysis includes the following assessments and fixes for Oracle configurations:

- Storage configuration operating system using the NFS protocol: assesses and fixes configuration issues with the NFS configuration for existing Oracle database deployments.
- Storage configuration operating system using the iSCSI protocol: assesses and fixes configuration issues with the iSCSI configuration for existing Oracle database deployments.
- Storage configuration operating system using Automatic Storage Management (ASM): assesses configuration issues with the ASM configuration for existing Oracle database deployments.

## [Implement well-architected database configurations in Workload Factory](#)

### **Permissions changes for Workload Factory for Databases**

Workload Factory for Databases updated the permissions policies to provide more clarity into what it requires for specific actions and granularity for selecting only the permissions you need. When you add credentials, you'll have three permissions options to choose from instead of the previous permissions model which was *read-only* and *read/write*. The new permissions model breaks up the permissions policies as follows:

- *View, planning, and analysis*: view the inventory of database resources, learn about the health of your resources, review the well-architected analysis of your database configurations, get error log analysis, and explore savings
- *Operations and remediation*: perform operational tasks for your database resources and fix issues for database configurations and the underlying FSx for ONTAP file system storage
- *Database host creation*: deploy database hosts and the underlying FSx for ONTAP file system storage according to best practices

When adding credentials, you can select one or more of these permission policies based on the level of access you want to provide to Workload Factory for Databases.

## [Workload Factory permissions reference](#)

### **Well-architected dashboard screen added**

A new *Well-architected* dashboard screen has been added to the Databases menu. From this screen, you can get an aggregated view of the configuration status of your entire database estate.

### **Edit protection for database deployments in NetApp Backup and Recovery**

From the Workload Factory console, you can select to edit protection for a protected resource and then you'll be redirected to Backup and Recovery in the NetApp Console where you can modify the protection policy or schedule.

## [Edit protection for database deployments](#)

### **Error filtering by tags available for the error analysis feature**

You can filter and view Microsoft SQL Server error logs by infrastructure-oriented tags when using the error analysis feature in Workload Factory. This enhancement helps you quickly triage, troubleshoot, and resolve issues.

## [Analyze error logs with tag filtering](#)

## 06 October 2026

### 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.

### Enhancements for Oracle in the Workload Factory console

#### Oracle database resource screen

Each Oracle database has its own dedicated resource screen available from the Databases inventory. The resource screen provides an overview with the database name, status (on or offline), tenancy, and deployment type. Also included are charts with the following data over a three-month period: CPU utilization, latency, IOPS, and throughput. Capacity utilization provides total size for the database, written data size, used solid-state drive capacity, and used capacity pool storage.

From the resource screen, you can view information about the Oracle server (deployment model, OS, edition, version, and more), location (AWS account, region, Availability Zone, and subnet), storage and compute (FSx for ONTAP file system details, database instance type, and associated LUNs and volumes), and connectivity (VPC and access protocol). You can also check the well-architected status of the database configurations and view pluggable databases (PDBs) associated with the database.

#### Support for Oracle pluggable databases

Oracle pluggable databases are viewable in the Databases inventory and from the resource screen of their parent container database. The following information about your PDBs is available: CDB name, host name, protection status, database size, FSx for ONTAP file system, AWS credentials, AWS account, and region.

#### Well-architected analysis for Oracle

The well-architected analysis includes assessments of configuration issues with the storage layout for Oracle databases over NFS or over iSCSI with or without Automatic Storage Management (ASM) and storage configuration issues on the operating system for Oracle over iSCSI LUNs. You can use this information to make informed decisions about your database deployments and ensure they are running efficiently.

### [Implement well-architected database configurations in Workload Factory](#)

#### Databases inventory enhancements

From the instance screen in the Databases inventory, the options to fix, postpone, and dismiss a configuration issue are available for single or multiple instances for Microsoft SQL Server or for single or multiple databases for Oracle.

### Optimize savings based on usage for Amazon Elastic Block Store (EBS)

Workload Factory can analyze your EBS performance usage and then suggest the best and most cost-efficient FSx for ONTAP configuration so that you can save more by switching to FSx for ONTAP.

### [Explore savings for detected storage environments in the Workload Factory console](#)

## 1 September 2025

## Agentic AI-powered error log analyzer

The Agentic AI-powered error log analyzer is a new feature that leverages advanced machine learning algorithms to automatically detect and analyze errors in log files. This tool aims to streamline the troubleshooting process by providing developers with actionable insights and recommendations based on the patterns it identifies in the logs.

[Learn more about the Agentic AI-powered error log analyzer](#)

## Oracle support

Workload factory includes support for Oracle databases. In the workload factory console, you can view your Oracle databases from the inventory, register databases to use advanced features in workload factory, and analyze Oracle databases for alignment with best practices using the well-architected feature. The well-architected analysis determines whether the storage configurations for Oracle databases are optimized. You can use this information to make informed decisions about your database deployments and ensure they are running efficiently.

[Implement well-architected database configurations in workload factory](#)

## Support for Microsoft SQL Server deployments on second-generation FSx for ONTAP file systems

Workload factory supports Microsoft SQL Server deployments on second-generation FSx for ONTAP file systems. This enhancement allows you to leverage the latest features and performance improvements available in the second-generation file systems while managing your SQL Server workloads.

## Windows authentication for SQL Server protection

Authenticating Microsoft SQL Server instances with Windows credentials is embedded in the workflow to prepare Microsoft SQL Server hosts for protection with BlueXP backup and recovery. This used to be a prerequisite step to complete manually. Instead, you'll be prompted to share Windows credentials with administrative access if you haven't registered the hosts in workload factory using Windows credentials.

[Learn how to protect Microsoft SQL Server workloads via the workload factory console.](#)

## Well-architected analysis includes MTU alignment for SQL Server

The well-architected analysis assesses and fixes Maximum Transmission Unit (MTU) misalignment across endpoints for Microsoft SQL Server on FSx for ONTAP storage. Aligning MTU settings helps to optimize network performance and reduce latency for SQL Server workloads.

[Implement well-architected database configurations in workload factory](#)

## 12 August 2025

### BlueXP backup and recovery now supports Microsoft SQL Server workloads

BlueXP backup and recovery enables you to back up, restore, verify, and clone Microsoft SQL Server databases and availability groups. From the workload factory console, you can access and use BlueXP backup and recovery to protect Microsoft SQL Server workloads.

[Learn how to protect Microsoft SQL Server workloads via the workload factory console.](#)

For details about BlueXP backup and recovery, refer to the [Protect Microsoft SQL workloads overview with BlueXP backup and recovery](#).

## 04 August 2025

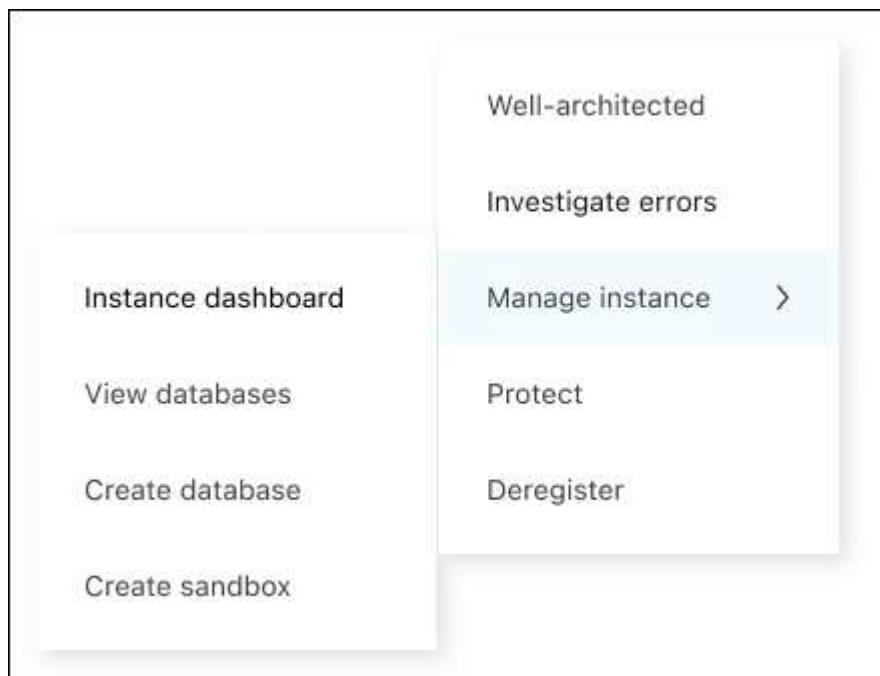
### Well-architected analysis includes high-availability cluster validation

The well-architected analysis now includes validation for high-availability clusters. This validation checks all cluster-related configurations from the server side, including disk availability and configuration on both nodes, Windows cluster configuration, and failover readiness. This ensures that the Windows cluster is properly set up and can successfully failover when needed.

[Implement well-architected database configurations in workload factory](#)

### Multi-level menu available for instances

The workload factory console now includes a multi-level menu for instances. This change provides a more organized and intuitive navigation structure for managing instances. Menu options for instance management include viewing the instance dashboard, viewing databases, creating a database, and creating a sandbox clone.



### New authentication option to explore savings

When the NT Authority\SYSTEM user doesn't have sufficient permissions on the Microsoft SQL Server, you can authenticate with SQL Server credentials or add the missing SQL Server permissions to NT Authority\SYSTEM.

[Explore potential savings for your database environments with Amazon FSx for NetApp ONTAP](#)

## 30 June 2025

### BlueXP workload factory notification service support

The BlueXP workload factory notification service enables workload factory to send notifications to the BlueXP alerts service or to an Amazon SNS topic. Notifications sent to BlueXP alerts appear in the BlueXP alerts panel. 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).

## [Configure BlueXP workload factory notifications](#)

Workload factory provides the following notifications for Databases:

- Well-architected report
- Host deployment

## **Onboarding enhancement for registering instances**

Workload factory for Databases has improved its onboarding process for registering instances running on Amazon FSx for NetApp ONTAP storage. You can now select instances in bulk for registration. Once an instance is registered, you can create and manage database resources within the workload factory console.

## [Instance management](#)

### **Analysis and fix for Microsoft Multipath I/O timeout setting**

The well-architected status for your database instances now includes the analysis and the fix for the Microsoft Multipath I/O (MPIO) timeout setting. Setting the MPIO timeout to 60 seconds ensures FSx for ONTAP storage connectivity and stability during failovers. If the MPIO setting isn't set accordingly, workload factory will provide a fix to set the MPIO timeout value to 60 seconds.

## [Implement well-architected database configurations in workload factory](#)

## **Enhancements to graphics in instance inventory**

From the instance inventory screen, various resource utilization graphs like throughput and IOPS now display 7 days of data so that you can monitor performance of SQL nodes from the workload factory console more efficiently. The performance metrics gathered from the SQL nodes will be saved in Amazon CloudWatch which can be used for Logs Insights or integrating with other analytic services in your environment.

From the Instances and Databases tabs within the inventory, we've enhanced the description and visualization for protection.

## **Support for Windows authentication in workload factory**

Now workload factory supports SQL Server authentication using Windows authenticated users to register instances and benefit from management features.

## [Register instances in workload factory for Databases](#)

## **03 June 2025**

### **PostgreSQL and Oracle detection**

You can now discover the instances that are running PostgreSQL server databases and Oracle database deployments in your AWS account within the workload factory console. Discovered instances will appear in the Databases inventory.

## Updated "Optimization" terminology

Previously called "Optimization", workload factory now uses "well-architected issues" and "well-architected status" to describe the analysis of database configurations and "fix" to describe remediation for opportunities to improve database configurations to meet best practice recommendations.

### [Configuration analysis for database environments in workload factory](#)

## Improved onboarding for instances

Instead of "undetected", "unmanaged", or "managed" term usage for instance management, workload factory now uses "register" for onboarding instances. The new registration process includes authenticating and preparing instances so that you can create, monitor, analyze, and fix resources in your database configurations within the workload factory console. The preparation step in the registration process indicates whether your instances are ready for management.

### [Instance management](#)

## 04 May 2025

### Dashboard enhancements

- Cross-account and cross-region views are available as you navigate between tabs in the BlueXP workload factory console. The new views improve resource management, monitoring, and optimization.
- From the **Potential savings** tile in the dashboard, you'll be able to quickly review what you might save by switching to FSx for ONTAP from Amazon Elastic Block Store or Amazon FSx for Windows File Server.

### Ad hoc scan available for database configurations

BlueXP workload factory for Databases automatically scans managed Microsoft SQL Server instances with FSx for ONTAP storage for potential configuration issues. Now in addition to the daily scan, you can scan at any time.

### Removal of on-premises assessment records

After you've explored savings for a Microsoft SQL server on-premises host, you have the option to remove the on-premises host record from BlueXP workload factory.

## Optimization enhancements

### Clone cleanup

The clone cleanup assessment and remediation identifies and manages old and costly clones. Clones that are older than 60 days can be refreshed or deleted from the BlueXP workload factory console.

### Postpone and dismiss configuration analysis

Some configurations might not apply to your database environments. You now have options to postpone a particular configuration analysis by 30 days or dismiss the analysis.

## 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.

## 04 April 2025

### Optimization enhancements

New optimization assessments, remediations, and the display of multiple resources are available when optimizing your database environments.

### Resiliency assessments

The enhancements include new resiliency assessments to check that data redundancy and disaster recovery capabilities are configured for your database environments.

- FSx for ONTAP backups: analyzes if FSx for ONTAP file systems serving the SQL Server instance's volumes are configured with scheduled FSx for ONTAP backups.
- Cross-region replication: assesses if FSx for ONTAP file systems serving Microsoft SQL Server instances are configured with cross-region replication.

### Compute remediation

Receive Side Scaling (RSS) remediation configures RSS to distribute network processing across multiple processors and ensure efficient load distribution.

### Local snapshot remediation

Local snapshot remediation sets up snapshot policies for volumes for your Microsoft SQL Server instances to keep your database environments resilient in case of data loss.

## Optimize configurations

### Support for multiple resource selection

When optimizing database configurations, you'll now be able to select specific resources instead of all resources.

## Optimize configurations

### Improved Inventory view

The Inventory tab in the workload factory console has been streamlined so that it contains only SQL servers running on Amazon FSx for NetApp ONTAP. Now you'll find SQL servers on-premises and running on Amazon Elastic Block Store and Amazon FSx for Windows File Server in the Explore savings tab.

### Quick create available for PostgreSQL server deployment

You can use this fast deployment option to create a PostgreSQL server with HA configuration and embedded best practices.

## Create a PostgreSQL server in BlueXP workload factory

## 03 March 2025

### PostgreSQL high-availability configuration

You can now deploy a high-availability (HA) configuration for PostgreSQL server.

## [Create a PostgreSQL server](#)

### **Terraform support for PostgreSQL server creation**

You can now use Terraform from the Codebox to deploy PostgreSQL.

- [Create a PostgreSQL database server](#)
- [Use Terraform from Codebox](#)

### **Resiliency assessment for local snapshot schedule**

A new resiliency assessment is available for database workloads. We assess if volumes for your Microsoft SQL Server instances have valid scheduled snapshot policies. Snapshots are point-in time copies of your data and help keep your database environments resilient in case of data loss.

## [Optimize configurations](#)

### **MAXDOP remediation for database workloads**

BlueXP workload factory for Databases now supports remediation for the maximum degree of parallelism (MAXDOP) server configuration. When MAXDOP configuration is suboptimal, you can let BlueXP workload factory optimize the configuration for you.

## [Optimize configurations](#)

### **Email savings analysis report**

When exploring savings for your Amazon Elastic Block Store and FSx for Windows File Server storage environments when compared with FSx for ONTAP, you can now send the recommendation report via email to yourself, team members, and customers.

## **03 February 2025**

### **On-premises database environments cost analysis and migration planning**

BlueXP workload factory for Databases now detects, analyzes, and help you plan for on-premises database migration to Amazon FSx for NetApp ONTAP. You can use the savings calculator to estimate the cost of running your on-premises database environment in the cloud and review recommendations for migrating your on-premises database environment to the cloud.

## [Explore savings for on-premises database environments](#)

### **New optimization assessments for Databases**

The following assessments are now available in BlueXP workload factory for Databases. These assessments are focused on detecting and protecting against potential security vulnerabilities and detecting and alleviating performance bottlenecks.

- **Receive side scaling (RSS) configuration:** checks if the RSS configuration is enabled and if the number of queues is set to the recommended value. The assessment also provides recommendations to optimize the RSS configuration.
- **Maximum Degree of Parallelism (MAXDOP) server configuration:** The assessment checks if MAXDOP is configured correctly and provides recommendations to optimize performance.

- **Microsoft SQL Server patches:** The assessment checks if the latest patches are installed on the SQL Server instances and provides recommendations to install the latest patches.

[Optimize configurations](#)

## 06 January 2025

### Databases Dashboard enhancements

A new design of the Dashboard includes the following graphics and enhancements:

- Host distribution graph shows the number of Microsoft SQL Server hosts and PostgreSQL hosts
- Instances distribution details include the total number of detected instances and the number of managed Microsoft SQL Server and PostgreSQL instances
- Databases distribution details include the total number of databases and the number of managed Microsoft SQL Server and PostgreSQL databases
- Optimization score and statuses for managed and online instances
- Optimization details for storage, compute, and application categories
- Optimization details for Microsoft SQL Server instance configurations like storage sizing, storage layout, ONTAP storage, compute, and applications
- Potential savings for database workloads running on Amazon Elastic Block Store and FSx for Windows File Server storage environments compared to Amazon FSx for NetApp ONTAP storage

### New 'Completed with issues' status in Job monitoring

The job monitoring feature for Databases now provides the new 'Completed with issues' status so that you can learn which sub jobs had issues and what the issues are.

[Monitor your databases](#)

### Assessment and optimization for over-provisioned Microsoft SQL Server licenses

The savings calculator now evaluates whether Enterprise edition is required for your Microsoft SQL Server deployment. If a license is over-provisioned, the calculator recommends downgrading. You'll be able to automatically downgrade the license in Databases by optimizing the application.

- [Explore savings with FSx for ONTAP for your database workloads](#)
- [Optimize your SQL Server workloads](#)

## 01 December 2024

### Continuous optimization adds compute remediation and assessment

Databases now provides insights and recommendations to help you optimize compute resources for Microsoft SQL Server instances. We measure CPU utilization and leverage the AWS Compute Optimizer service to recommend optimal right-sized instance types and notify you of available operating system patches. Optimizing compute resources can help you make informed decisions about instance types, leading to cost savings and efficient resource utilization.

[Optimize compute resource configurations](#)

## PostgreSQL support

You can now deploy and manage standalone PostgreSQL server deployments in Databases.

[Create a PostgreSQL server](#)

## 3 November 2024

### Continuously optimize your Microsoft SQL Server workloads with Databases

BlueXP workload factory introduces continuous guidance and guardrails to ensure ongoing optimization and adherence to best practices for the storage component of your Microsoft SQL Server workloads on Amazon FSx for NetApp ONTAP. This feature continuously scans your Microsoft SQL Server estate offline, providing you with a comprehensive report of insights, opportunities, and recommendations to help you achieve peak performance, cost-efficiency, and compliance.

[Optimize your SQL Server workloads](#)

## Terraform support

You can now use Terraform from the Codebox to deploy Microsoft SQL Server.

- [Create a database server](#)
- [Use Terraform from Codebox](#)

## 29 September 2024

### Explore savings for detected Microsoft SQL servers on FSx for Windows File Server

You can now explore savings of detected Microsoft SQL servers on Amazon EC2 with FSx for Windows File Server storage in the savings calculator. Depending on your SQL server and storage requirements, you might find that FSx for ONTAP storage is the most cost effective for your database workloads.

[Explore savings with FSx for ONTAP for your database workloads](#)

## 1 September 2024

### Explore savings via customization

You can now customize configuration settings for Microsoft SQL server on Amazon EC2 with FSx for Windows File Server and Elastic Block Store storage in the savings calculator. Depending on your storage requirements, you might find that FSx for ONTAP storage is the most cost effective for your database workloads.

[Explore savings with FSx for ONTAP for your database workloads](#)

### Navigate to the savings calculator from the home page

You can now navigate to the savings calculator from the [workload factory console](#) home page. Select from Elastic Block Store and FSx for Windows File Server to get started.

Databases

Detect your existing database estate on AWS, assess potential cost savings by moving to FSx for ONTAP, deploy databases end-to-end with built-in best practices for optimization, and automate thin cloning for CI/CD pipelines.

Deploy database host

Explore savings

Go to Databases inventory

Microsoft SQL Server on EBS

Microsoft SQL Server on FSx for Windows

**4 August 2024**

### Savings calculator enhancements

- Cost estimate descriptions

You can now learn how cost estimates are calculated in the savings calculator. You'll be able to review descriptions of all calculations for your Microsoft SQL Server instances using Amazon Elastic Block Store storage as compared to using Amazon FSx for ONTAP storage.

- Support for Always On availability group

Databases now provides cost savings calculations for the Always On availability group deployment type with Microsoft SQL Server using Amazon Elastic Block Store.

- Optimize SQL server licensing with FSx for ONTAP

The Databases calculator determines whether the SQL license edition you use with Amazon Elastic Block Store storage is optimized for your database workloads. You'll get a recommendation for the optimal SQL license with FSx for ONTAP storage.

- Multiple SQL server instances

Databases now provides cost savings calculations for a configuration hosting multiple Microsoft SQL Server instances using Amazon Elastic Block Store.

- Customize calculator settings

Now you can customize settings for Microsoft SQL Server, Amazon EC2, and Elastic Block Store to explore savings manually. The savings calculator will determine the best configuration based on cost.

[Explore savings with FSx for ONTAP for your database workloads](#)

## Initial release of BlueXP workload factory for Databases

The initial release includes the capability to explore savings with Amazon FSx for NetApp ONTAP as the storage environment for your database workloads, detect, manage, and deploy Microsoft SQL Servers, deploy and clone databases, and monitor these jobs within workload factory.

[Learn about Databases](#)

## Known limitations for NetApp Workload Factory for Databases

Known limitations identify platforms, devices, or functions that are not supported with NetApp Workload Factory for Databases, or that do not interoperate correctly with it. Review these limitations carefully.

### Instance detection support

Workload Factory supports the detection of PostgreSQL instances running only on Amazon Linux operating systems and Oracle instances running only on Red Hat Enterprise Linux.

### AMI image version support

We only show up to four recent versions or revisions of AMI images for each SQL version. We don't allow installation from an AMI image that is older than the latest four versions.

### Custom AMI

Deployment with a custom AMI assumes that install media is present in the image. This is needed for FCI configuration for the following reasons:

- To uninstall and reconfigure to form the FCI cluster
- When selecting a collation other than the collation set in the image for standalone deployment

As part of deployment, Workload Factory doesn't install a custom AMI with multiple Microsoft SQL Server instances. Only the default Microsoft SQL Server instance is selected and configured as part of deployment.

### Roll back and retry for failed deployments

Rollback and retry for failed deployments aren't supported in Workload Factory. You can roll back or retry a failed stack from the CloudFormation console in AWS.

### Active Directory and DNS resource rollback

The following DNS resources aren't removed from Active Directory and DNS when you roll back a test or failed deployment from the CloudFormation console in AWS:

- Each node of the SQL deployment to its management IP address
- Windows Cluster name to the reserved secondary IP addresses of EC2 instances in cases of FCI

- SQL FCI name to the two reserved IP addresses from both EC2 instances in FCI

You'll need to cleanup these up manually or wait for the domain to purge stale entries.

## Always On availability groups configuration support

Managing Always On availability groups configurations isn't supported in Workload Factory.

## Custom encryption keys

Custom encryption keys for FSx for ONTAP aren't listed based on service applicability. You must choose the appropriate key. AWS managed keys are filtered based on applicability to service.

## CloudFormation template

A downloaded or copied Codebox-generated CloudFormation template (YAML file) has a limited retention period of seven days.

## Sandbox support

The maximum number of sandbox clones that can be created from a database server is 90.

## Microsoft SQL Server detection and management

Microsoft SQL Server detection isn't saved. Each time you access Databases in Workload Factory, Microsoft SQL Server detection runs again to identify SQL installations in the region.

## Explore savings

In the Inventory tab, the estimated cost shown for each Microsoft SQL instance is calculated at the FSx for ONTAP file system level and not for volumes hosting the SQL instance.

## Multiple FSx for ONTAP file systems

Workload Factory doesn't support creating or saving configurations for Microsoft SQL Server with multiple FSx for ONTAP file systems. Only one FSx for ONTAP file system configuration deployment is supported.

## Optimization limitations

### Compute rightsizing optimization

When changing to certain instance types, network configuration may be reset which may then cause node connection failure during optimization and result in job failure. Manual intervention may be needed to check and update the DNS setting and iSCSI sessions. Refer to [Amazon Elastic Compute Cloud documentation](#) for more information about EC2 resize limitations.

### ONTAP and OS configuration optimization

Bulk optimization for more than one SQL Server instance is not supported for ONTAP configuration and OS configuration.

## Operating system patch assessment for optimization

The operating system patch assessment that is made for optimization purposes may not work in a private network. The assessment relies on AWS Patch Manager. To understand how to patch Windows EC2 instances in a private network, refer to the AWS Cloud Operations Blog ["How to patch Windows EC2 instances in private subnets Using AWS Systems Manager"](#).

## On-premises savings calculation

Cost and percentage savings calculations of volume size for Microsoft SQL Server on-premises storage isn't taken into consideration.

## Cross-region replication assessment

- The optimization assessment can't determine if cross-region replication (CRR) is being used when the target FSx for ONTAP file system is in a different AWS account.
- Workload Factory assesses only the CRR configuration of the source FSx for ONTAP file system.

## Database host authentication when exploring savings

In some cases with limited permissions, the explore savings page doesn't load data after successful authentication.

## Integration with NetApp Backup and Recovery

After adding hosts to NetApp Backup and Recovery for protection, database discovery sometimes fails.

## Region support

The following AWS regions aren't supported:

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

## Registering Oracle on SUSE Linux Enterprise Server 12

When registering Oracle Database on SUSE Linux Enterprise Server 12, Workload Factory does not install Python dependencies. You must manually configure the required Python version. Workload Factory requires the minimum version of 3.6; however, we recommend version 3.11.

# Get started

## Learn about NetApp Workload Factory for Databases

NetApp Workload Factory for Databases is an end-to-end database deployment and maintenance service with built-in best practices for optimization, automatic thin cloning, and monitoring and resolution capabilities.

### What is NetApp Workload Factory for Databases?

NetApp Workload Factory for Databases detects, assesses, plans, provisions, and moves Microsoft SQL Server data to Amazon FSx for NetApp ONTAP (FSx for ONTAP) deployments optimized to meet your performance and cost expectations while adhering to industry best practices. Throughout the life cycle, NetApp Workload Factory for Databases delivers continuous optimization and management for databases on FSx for ONTAP.

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

### Workload Factory benefits for self-managed databases

Workload Factory offers the following beneficial best practices and automation for self-managed databases.

#### Best practices

- Integrated knowledge from AWS Cloud, Microsoft Windows and SQL servers, and NetApp ONTAP for SQL server deployments on EC2 instances.
- Total cost of ownership optimized deployment.
- End-to-end deployment automation that follows AWS, Microsoft and ONTAP best practices.
- *Quick create* deployment mode helps you avoid potential pitfalls with manual configuration.

#### 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, allowing 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.

#### Workload Factory for Databases features

Workload Factory for Databases offers the following features:

- **Simple, fast deployments:** simplify and streamline your provisioning experience by selecting answers to specification questions and eliminating the time typically required to investigate how to provision and configure your Microsoft SQL Server on AWS.

- **Automated orchestration:** available via the Workload Factory console with *Quick* and *Advanced* create deployment modes, the Chatbot, and AWS CloudFormation via the Codebox.
- **Built-in features:** utilize NetApp, Microsoft, and Amazon best practices and AWS resource selections built in to deployment configuration.
- **Cost estimation:** assess potential cost savings with the Savings calculator that estimates and details storage, compute, SQL licensing, snapshot, and clone itemized costs for existing Microsoft SQL Server deployments with Elastic Block Store and FSx for Windows File Server compared with FSx for ONTAP.
- **Reusable automation templates:** create, re-use, and customize CloudFormation templates from the Workload Factory Codebox for future Microsoft SQL Server deployments in multiple environments.
- **AWS resource discovery and provisioning:** automatically detect Microsoft SQL Servers on FSx for ONTAP, FSx for Windows File Servers, and Elastic Block Store deployed from your AWS account. The Databases Inventory serves as a launching point for exploring costs savings opportunities for servers in other AWS storage systems or a management tool for FSx for ONTAP based servers and instances.
- **Sandbox creation:** create an on-demand, isolated database environment that can be used for testing, integration, diagnostics, and training, without affecting production data.
- **Database creation:** create a user database for your existing Microsoft SQL Servers with *Quick* or *Advanced* create modes to configure your database collation, file names, and sizes. Storage configuration is included.
- **Job Monitoring:** monitor and track database job execution progress and diagnose and troubleshoot issues in case any failure occurs.
- **Continuous optimization:** continuously scans your Microsoft SQL Server estate offline, providing you with a comprehensive report of insights, opportunities, and recommendations to help you achieve operational excellence.

## Benefits of Amazon FSx for ONTAP for self-managed databases

- **Durability, availability, and reliability:** FSx for ONTAP offers multiple features that enhance the durability and availability of self-managed databases hosted on FSx for ONTAP, such as high availability supporting Single and Multiple Availability Zones deployments, application-aware snapshots, improved disaster recovery using replication, and efficient backup.
- **Performance and scalability:** FSx for ONTAP offers performance optimization with high throughput, low latency, high-speed network connections, and scalability with multiple file systems for scaling aggregate performance required for a workload.
- **Data management and efficiency:** FSx for ONTAP offers multiple features that enhance data management and efficiency, such as space efficient thin clones, thin provisioning, compression and deduplication, and tiering infrequently accessed data to the capacity pool.

[Learn about FSx for ONTAP for Workload Factory.](#)

## Tools to use NetApp Workload Factory

You can use NetApp Workload Factory with the following tools:

- **Workload Factory console:** The Workload Factory console provides a visual, holistic view of your applications and projects.
- **NetApp Console:** The NetApp Console provides a hybrid interface experience so that you can use Workload Factory along with other NetApp data services.
- **Ask me:** Use the Ask me AI assistant to ask questions and learn more about Workload Factory without

leaving the Workload Factory console. Access Ask me from the Workload Factory help menu.

- **CloudShell CLI:** Workload Factory includes a CloudShell CLI to manage and operate AWS and NetApp environments across accounts from a single, browser-based CLI. Access CloudShell from the top bar of the Workload Factory console.
- **REST API:** Use the Workload Factory REST APIs to deploy and manage your FSx for ONTAP file systems and other AWS resources.
- **CloudFormation:** Use AWS CloudFormation code to perform the actions you defined in the Workload Factory console to model, provision, and manage AWS and third-party resources from the CloudFormation stack in your AWS account.
- **Terraform NetApp Workload Factory provider:** Use Terraform to build and manage infrastructure workflows generated in the Workload Factory console.

## Supported configurations

Workload Factory supports the following database engines, versions, operating systems, and deployment models according to AWS, NetApp ONTAP, Microsoft SQL Server, Oracle, and PostgreSQL best practices.

Engine	Version	Operating System	Deployment Model
Microsoft SQL Server	SQL Server 2016	Windows Server 2016	FCI, Standalone
Microsoft SQL Server	SQL Server 2019	Windows Server 2016	FCI, Standalone
Microsoft SQL Server	SQL Server 2022	Windows Server 2016	FCI, Standalone
Microsoft SQL Server	SQL Server 2016	Windows Server 2019	FCI, Standalone
Microsoft SQL Server	SQL Server 2019	Windows Server 2019	FCI, Standalone
Microsoft SQL Server	SQL Server 2022	Windows Server 2019	FCI, Standalone
Microsoft SQL Server	SQL Server 2016	Windows Server 2022	FCI, Standalone
Microsoft SQL Server	SQL Server 2019	Windows Server 2022	FCI, Standalone
Microsoft SQL Server	SQL Server 2022	Windows Server 2022	FCI, Standalone
Microsoft SQL Server	SQL Server 2017	Any	FCI, Standalone
Microsoft SQL Server	SQL Server 2016, 2017, 2019, 2022	Any	FCI, Standalone
Oracle	19c	RHEL, SuSe Linux	Standalone (Single/Multi tenant)
Oracle	21c	RHEL, SuSe Linux	Standalone (Single/Multi tenant)
Oracle	19c	RHEL, SuSe Linux	Standalone (Single/Multi tenant)
Oracle	21c	RHEL, SuSe Linux	Standalone (Single/Multi tenant)
Oracle	19c	RHEL, SuSe Linux	Standalone with ASM (Single/Multi tenant)
Oracle	21c	RHEL, SuSe Linux	Standalone with ASM (Single/Multi tenant)

Engine	Version	Operating System	Deployment Model
PostgreSQL	PostgreSQL 15	Amazon Linux 2023 AMI	HA, Standalone instance
PostgreSQL	PostgreSQL 16	Amazon Linux 2023 AMI	HA, Standalone instance

## Integrated AWS services

Databases includes the following integrated AWS services:

- CloudFormation
- Simple Notification Service
- CloudWatch
- Systems Manager
- Secrets Manager

## Regions

Databases 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 questions or technical support issues associated with your FSx for ONTAP file system, infrastructure, or any solution using this service, use the Support Center in your AWS Management Console to open a support case with AWS. Select the “FSx for ONTAP” service and appropriate category. Provide the remaining information required to create your AWS support case.

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

## Quick start for NetApp Workload Factory for Databases

With NetApp Workload Factory for Databases, you can get started immediately in *basic* mode. If you'd like to use Workload Factory to discover hosts, manage resources, and more, you can get started in a few steps.

You must have an AWS account to use Databases.

Follow these steps to get started.

**1**

## Log in to NetApp Workload Factory

You'll need to [set up an account with NetApp 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 Databases 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 Databases and AI, and creating the IAM policies for the required permissions.

**3**

## Discover or deploy resources

With credentials and IAM policies, you can discover existing database resources in the inventory or [deploy a host server](#). The inventory provides a unified interface to manage resources.

**4**

## Explore cost-saving opportunities

When you have databases running on-premises or on AWS with storage on Amazon Elastic Block Store (EBS) or FSx for Windows File Server, you can [use the Explore savings calculator](#) to analyze costs and plan migrations effectively.

**5**

## Implement well-architected database configurations

Workload Factory for Databases regularly analyzes Microsoft SQL Server and Oracle deployments on Amazon FSx for NetApp ONTAP storage from the Well-architected dashboard. To troubleshoot issues from the well-architected dashboard for your database resources, you first need to [register instances](#).

After registering instances, you can view the well-architected status and take action to [implement well-architected database configurations](#) in Workload Factory.

### What's next

When you have registered resources with FSx for ONTAP file system storage in your Databases inventory, you can [create a user database](#) or [clone your host to create a sandbox](#).

# Use Database workloads

## Explore savings in NetApp Workload Factory for Databases

Explore savings in NetApp Workload Factory for Databases for your database workloads by comparing the costs of using Microsoft SQL Server on Amazon Elastic Block Store (EBS), FSx for Windows File Server, and on-premises storage with FSx for ONTAP storage.

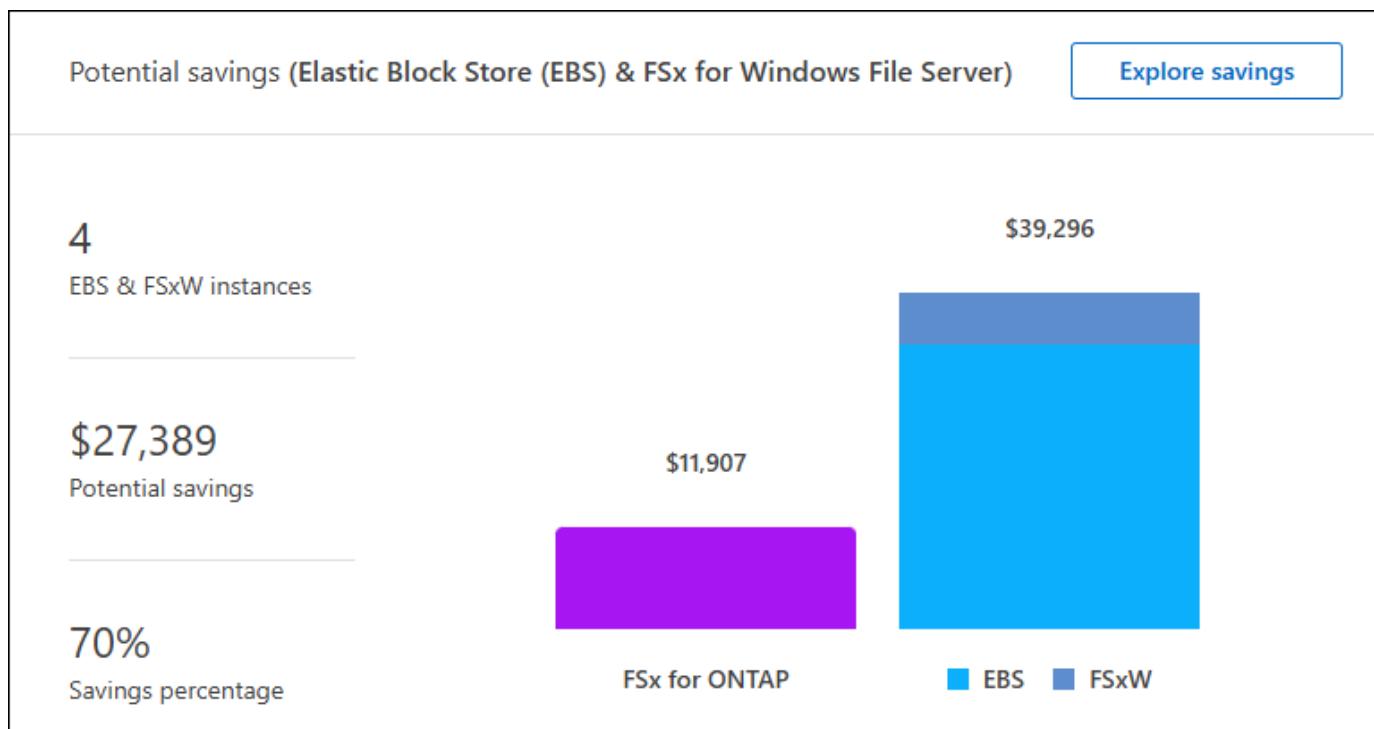
### About this task

Workload Factory has several ways for you to explore savings for your database workloads running on Amazon Elastic Block Store (EBS), FSx for Windows File Server, and on-premises storage - from the Dashboard, from the Inventory tab, and from the Explore savings tab. In all cases, you can use the savings calculator to compare various cost components of running Microsoft SQL Server workloads like storage, compute, SQL license, snapshots, and clones for your database workloads on FSx for ONTAP file systems against Elastic Block Store (EBS), FSx for Windows File Server, and on-premises storage.

If Workload Factory determines that you could save money by running these workloads on an FSx for ONTAP file system, you can deploy Microsoft SQL over FSx for ONTAP directly from the savings calculator in the Workload Factory console. When you have multiple Microsoft SQL Server instances over Elastic Block Store, FSx for Windows File Server, or on-premises storage, we'll recommend an FSx for ONTAP configuration with a single SQL instance.

### Potential savings for all database workloads

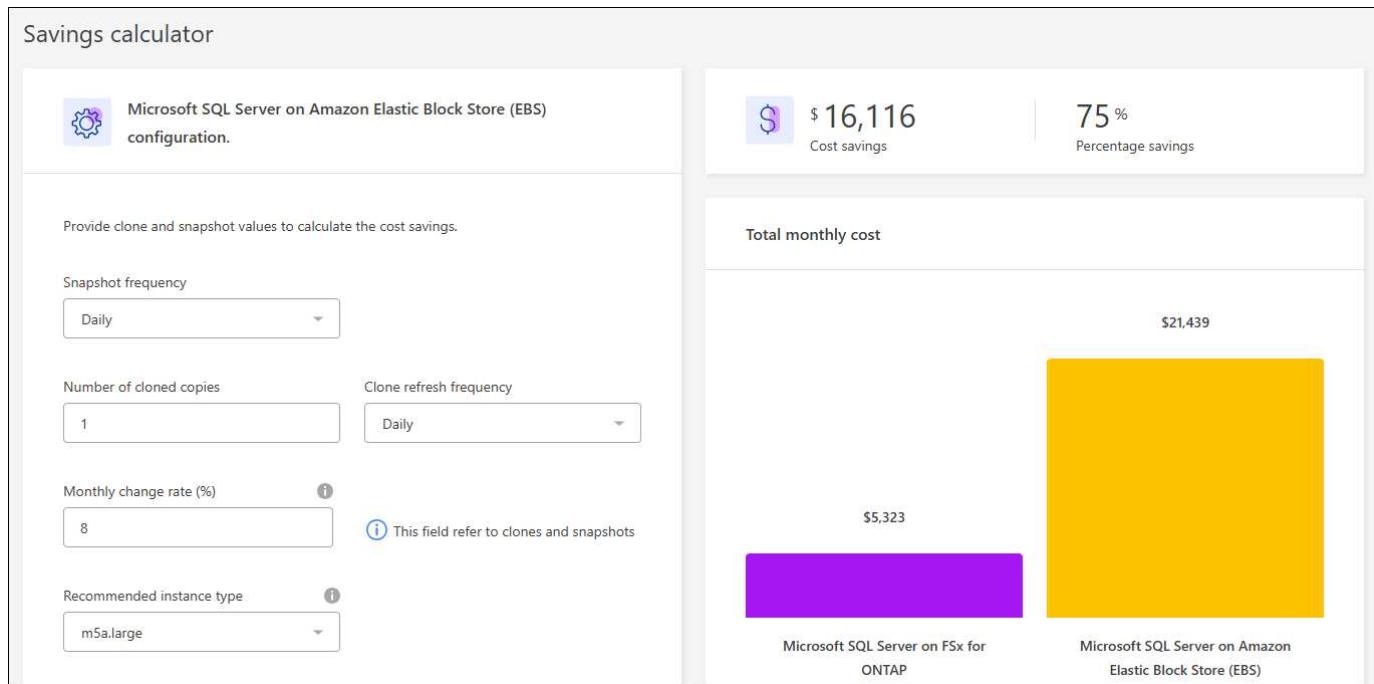
Visit the Databases **Dashboard** in the Workload Factory console to get an overview of potential cost savings for running all of your database workloads on FSx for ONTAP. In the **Potential savings** tile, you can view the number of all database workloads you have on Elastic Block Store and FSx for Windows File Server, the potential cost savings, the savings percentage, and visual representation in the bar graph.



## Savings calculator

You can utilize the savings calculator so you can compare various cost components of running Microsoft SQL Server workloads like storage, compute, SQL license, snapshots, and clones for your database workloads on FSx for ONTAP file systems against Elastic Block Store (EBS), FSx for Windows File Server, and on-premises storage. Depending on your storage requirements, you might find that FSx for ONTAP file systems are the most cost effective for your database workloads.

The calculator displays whether the storage for the database workloads on these Microsoft SQL Servers would cost less if you used an FSx for ONTAP file system. [Learn how to use the calculator.](#)



## SQL Server deployment analysis

The calculator performs a comprehensive analysis of your SQL Server deployment to ensure that the resources and features being utilized are appropriately matched to the SQL Server edition. Here are the key factors and conditions the calculator checks before recommending a downgrade to Standard Edition:

### Deployment model

The calculator evaluates the deployment model and whether Enterprise edition is required.

### Allocated resources

The calculator assesses the conditions of the following license-dependent allocated resources:

- Target Instance vCPUs: The instance has 48 or fewer virtual CPUs.
- Memory Allocation: The instance has 128GB or less of memory.

### Enterprise feature usage

The calculator verifies if any of the following Enterprise features are in use:

- Database-level Enterprise features
- Online index operations
- Resource Governor

- Peer-to-peer or Oracle replication
- R/Python extensions
- Memory-optimized TempDB

If the assessed SQL Server instance doesn't utilize any of the above Enterprise features and meets the resource constraints, the calculator will recommend downgrading the license to Standard Edition. This recommendation is made to help you optimize your SQL Server licensing costs without compromising performance or functionality.

## Calculator options

Two calculator options are available for making the cost comparison between your systems and FSx for ONTAP — customization and detection.

Explore savings via customization: You provide the configuration settings for Microsoft SQL server on Amazon EC2 with EBS or FSx for Windows File Server including the region, deployment model, SQL server edition, monthly data change rate, snapshot frequency, and more.

Explore savings for detected hosts: Workload Factory links to your existing Microsoft SQL servers and pulls in the details to the calculator for automatic comparison. You'll need to grant *view, planning, and analysis* permissions to use this calculator option. You can change the use case, but all other details are automatically determined in the calculation.

Additionally, you can [add AWS credentials](#) to improve the accuracy of the calculator analysis. Select **Calculate savings based on existing resources**. You'll be redirected to the Add credentials page. After you add credentials, select the existing resources to compare with FSx for ONTAP, and select **Explore savings**.

### Explore savings via customization

Follow the steps under the tab for your storage type.

## Amazon Elastic Block Store (EBS)

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Databases**.
3. From Databases, select **Explore savings** then **Microsoft SQL Server on EBS**.
4. In the Savings calculator, provide the following details:
  - a. **Region**: Select a region from the dropdown menu.
  - b. **Deployment model**: Select a deployment model from the dropdown menu.
  - c. **SQL server edition**: Select the SQL server edition from the dropdown menu.
  - d. **Monthly data change rate (%)**: Enter the percentage that clone and snapshot data changes on average per month.
  - e. **Snapshot frequency**: Select a snapshot frequency from the dropdown menu.
  - f. **Number of cloned copies**: Enter the number of cloned copies in the EBS configuration.
  - g. **Monthly SQL BYOL cost (\$)**: Optionally, enter the monthly SQL BYOL cost in dollars.
  - h. Under EC2 specifications, provide the following:
    - **Machine description**: Optionally, enter a name to describe the machine.
    - **Instance type**: Select the EC2 instance type from the dropdown menu.
  - i. Under Volume types, provide the following details for at least one volume type. IOPS and throughput apply to certain disk type volumes.
    - **Number of volumes**
    - **Storage amount per volume (GiB)**
    - **Provisioned IOPS per volume**
    - **Throughput MB/s**
  - j. If you selected the Always On availability deployment model, provide details for **Secondary EC2 specifications** and **Volume types**.

## Amazon FSx for Windows File Server

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Databases**.
3. From Databases, select **Explore savings** then **Microsoft SQL Server on FSx for Windows**.
4. In the Savings calculator, provide the following details:
  - a. **Region**: Select a region from the dropdown menu.
  - b. **Deployment model**: Select a deployment model from the dropdown menu.
  - c. **SQL server edition**: Select the SQL server edition from the dropdown menu.
  - d. **Monthly data change rate (%)**: Enter the percentage that clone and snapshot data changes on average per month.
  - e. **Snapshot frequency**: Select a snapshot frequency from the dropdown menu.
  - f. **Number of cloned copies**: Enter the number of cloned copies in the EBS configuration.

- g. **Monthly SQL BYOL cost (\$)**: Optionally, enter the monthly SQL BYOL cost in dollars.
- h. Under FSx for Windows File Server settings, provide the following:
  - **Deployment type**: Select the deployment type from the dropdown menu.
  - **Storage type**: SSD storage is the supported storage type.
  - **Total storage capacity**: Enter the storage capacity and select the capacity unit for the configuration.
  - **Provisioned SSD IOPS**: Enter the provisioned SSD IOPS for the configuration.
  - **Throughput (MB/s)**: Enter throughput in MB/s.
- i. Under EC2 specifications, select the **Instance type** from the dropdown menu.

After you provide details for your database host configuration, review the calculations and recommendations provided on the page.

Additionally, scroll down to the bottom of the page to view the report by selecting one of the following:

- **Export PDF**
- **Send by email**
- **View the calculations**

To switch to FSx for ONTAP, follow the instructions to [deploy Microsoft SQL Server on AQS EC2 using FSx for ONTAP file systems](#).

## Explore savings for detected hosts

Workload Factory enters the detected Elastic Block Store and FSx for Windows File Server host characteristics so that you can explore savings automatically.

### Before you begin

Complete the following prerequisites before you begin:

- [Grant view, planning, and analysis permissions](#) in your AWS account to detect Elastic Block Store (EBS) and FSx for Windows systems under the **Explore savings** tab and to show the savings calculation in the savings calculator.
- To get instance type recommendations and improve cost accuracy, do the following:
  1. Grant Amazon CloudWatch and AWS Compute Optimizer permissions.
    - a. Sign in to the AWS Management Console and open the IAM service.
    - b. Edit the policy for the IAM role. Copy and add the following Amazon CloudWatch and AWS Compute Optimizer permissions.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "compute-optimizer:GetEnrollmentStatus",
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": "compute-optimizer:PutRecommendationPreferences",
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": "compute-optimizer:GetEffectiveRecommendationPreferences",
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": "compute-optimizer:GetEC2InstanceRecommendations",
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": "autoscaling:DescribeAutoScalingGroups",
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": "autoscaling:DescribeAutoScalingInstances",
      "Resource": "*"
    }
  ]
}
```

## 2. Opt the billable AWS account in to AWS Compute Optimizer.

Follow the steps under the tab for your storage type.

## Amazon Elastic Block Store (EBS)

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Databases**.
3. From the Databases menu, select **Explore savings**.
4. In Explore savings, select the **SQL Server on Elastic Block Store (EBS)** tab.

If Workload Factory detects EBS hosts, you'll be redirected to the Explore savings tab. If Workload Factory doesn't detect EBS hosts, you'll be redirected to the calculator to [explore savings via customization](#).

5. From **Explore savings**, select one or more database hosts running on EBS and then select **Explore savings**.
6. If required, authenticate the database host with SQL Server credentials, Windows credentials, or by adding missing SQL Server permissions.

If the Explore savings page doesn't load data after successful authentication, select the **Inventory** tab to reload the data, and then select the **Explore savings** tab again.

7. In the Savings calculator, optionally, provide the following details on clones and snapshots in your EBS storage for a more accurate cost savings estimate.
  - a. **Snapshot frequency**: Select a snapshot frequency from the menu.
  - b. **Clone refresh frequency**: Select the frequency that clones refresh from the menu.
  - c. **Number of cloned copies**: Enter the number of cloned copies in the EBS configuration.
  - d. **Monthly change rate**: Enter the percentage that clone and snapshot data changes on average per month.
  - e. **Add hosts**: Optionally, select up to five detected EBS hosts to include in the savings calculation.

Workload Factory consolidates multiple SQL Server hosts into a single FSx for ONTAP configuration recommendation to optimize cost savings unless the selected EBS hosts exceed throughput, capacity, or IOPS limits for a single FSx for ONTAP file system.

## Amazon FSx for Windows File Server

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Databases**.
3. From the Databases menu, select **Explore savings**.
4. In Explore savings, select the **SQL Server on FSx for Windows** tab.

If Workload Factory detects FSx for Windows hosts, you'll be redirected to the Explore savings tab. If Workload Factory doesn't detect FSx for Windows hosts, you'll be redirected to the calculator to [explore savings via customization](#).

5. In the Explore savings tab, select **Explore savings** of the database host using FSx for Windows File Server storage.
6. If required, authenticate the database host with SQL Server credentials, Windows credentials, or by

adding missing SQL Server permissions.

If the Explore savings page doesn't load data after successful authentication, select the **Inventory** tab to reload the data, and then select the **Explore savings** tab again.

7. In the Savings calculator, optionally, provide the following details on clones (shadow copies) and snapshots in your FSx for Windows storage for a more accurate cost savings estimate.
  - a. **Snapshot frequency:** Select a snapshot frequency from the menu.  
If FSx for Windows shadow copies are detected, the default value is **Daily**. If shadow copies aren't detected, the default value is **No snapshot frequency**.
  - b. **Clone refresh frequency:** Select the frequency that clones refresh from the menu.
  - c. **Number of cloned copies:** Enter the number of cloned copies in the FSx for Windows configuration.
  - d. **Monthly change rate:** Enter the percentage that clone and snapshot data changes on average per month.

## Microsoft SQL Server on-premises

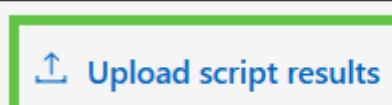
### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Databases**.
3. From the Databases menu, select **Explore savings**.
4. In Explore savings, select the **SQL Server on-premises** tab.
5. From the SQL Server on-premises tab, download the script to assess your on-premises SQL Server environments.
  - a. Download the assessment script. The script is a data collection tool based on PowerShell. It gathers and then uploads SQL Server configuration and performance data to Workload Factory. The migration advisor assesses the data and plans FSx for ONTAP deployment for your SQL Server environment.



[Upload script results](#) [Download script](#)

6. Run the script on the SQL Server host.
7. Upload the script output in the SQL Server on-premises tab in Workload Factory.



[Upload script results](#) [Download script](#)

8. From the SQL Server on-premises tab, select up to five database hosts and then select **Explore savings** to run a cost analysis of the SQL Server on-premises host(s) against FSx for ONTAP.
9. In the Savings calculator, select the region for the on-premises host.
10. Optionally, provide the following details on clones (shadow copies) and snapshots in your on-premises database environment for a more accurate cost savings estimate.

- a. **Snapshot frequency:** Select a snapshot frequency from the menu.  
If FSx for Windows shadow copies are detected, the default value is **Daily**. If shadow copies aren't detected, the default value is **No snapshot frequency**.
- b. **Clone refresh frequency:** Select the frequency that clones refresh from the menu.
- c. **Number of cloned copies:** Enter the number of cloned copies in the on-premises configuration.
- d. **Monthly change rate:** Enter the percentage that clone and snapshot data changes on average per month.

9. For more accurate results, update Compute information and Storage and performance details.

Workload Factory consolidates multiple SQL Server on-premises hosts into a single FSx for ONTAP configuration recommendation to optimize cost savings unless the selected on-premises hosts exceed throughput, capacity, or IOPS limits for a single FSx for ONTAP file system.

After you provide details for your database host configuration, review the calculations and recommendations provided on the page.

Additionally, scroll down to the bottom of the page to view the report by selecting one of the following:

- **Export PDF**
- **Send by email**
- **View the calculations**

To switch to FSx for ONTAP, follow the instructions to [deploy Microsoft SQL Server on AQS EC2 using FSx for ONTAP file systems](#).

#### On-premises host removal

After you've explored savings for a Microsoft SQL server on-premises host, you have the option to remove the on-premises host record from Workload Factory. Select the action menu of the Microsoft SQL Server on-premises host and then select **Delete**.

## Deploy Microsoft SQL Server on AWS EC2 using FSx for ONTAP

If you'd like to switch to FSx for ONTAP to realize cost savings, click **Create** to create the recommended configuration(s) directly from the Create new Microsoft SQL server wizard or click **Save** to save the recommended configuration(s) for later.



Workload Factory doesn't support saving or creating multiple FSx for ONTAP file systems.

#### Deployment methods

With *database host creation permissions*, you can deploy the new Microsoft SQL server on AWS EC2 using FSx for ONTAP directly from Workload Factory. You can also copy the content from the Codebox window and deploy the recommended configuration using one of the Codebox methods.

Without permissions, you can copy the content from the Codebox window and deploy the recommended configuration using one of the Codebox methods.

#### Related information

[Workload Factory permissions reference](#)

# Create a new database server

## Create a Microsoft SQL Server in Workload Factory for Databases

Creating a new Microsoft SQL Server, or database host, in Workload Factory for Databases requires an FSx for ONTAP file system deployment and resources for Active Directory.

### About this task

Before creating a Microsoft SQL Server from Workload Factory, learn about the available storage deployment types for the database host configuration, Microsoft Multi-path I/O configuration, Active Directory deployment, networking details, and the requirements to complete this operation.

After deployment, you'll need to [enable remote connection on the Microsoft SQL Server](#).

### FSx for ONTAP file system deployments

Creating a new Microsoft SQL Server requires an FSx for ONTAP file system as the storage backend. You can use an existing FSx for ONTAP file system or create a new file system. If you select an existing FSx for ONTAP file system as your database server storage backend, we create a new storage VM for the Microsoft SQL workloads.

FSx for ONTAP file systems have two Microsoft SQL Server deployment models: *Failover Cluster Instance (FCI)* or *Standalone*. Different resources are created for the FSx for ONTAP file system depending on the FSx for ONTAP deployment model you select.

- **Failover Cluster Instance (FCI) Microsoft SQL deployment:** A Multiple Availability Zone FSx for NetApp ONTAP file system is deployed when a new FSx for ONTAP file system is selected for FCI deployment. Separate volumes and LUNs are created for data, log, and tempdb files for an FCI deployment. An additional volume and LUN are created for Quorum or witness disk for Windows cluster.
- **Standalone Microsoft SQL deployment:** A Single Availability Zone FSx for ONTAP file system is created when a new Microsoft SQL Server is created. In addition, separate volumes and LUNs are created for data, log, and tempdb files.

### Microsoft Multi-path I/O configuration

Microsoft SQL Server deployment models both require LUN creation using the iSCSI storage protocol. Workload Factory configures Microsoft Multi-path I/O (MPIO) as part of configuring LUNs for SQL Server over FSx for ONTAP. MPIO is configured based on AWS and NetApp best practices.

For more information, refer to [SQL Server High Availability Deployments using Amazon FSx for NetApp ONTAP](#).

### Active Directory

The following occurs for Active Directory (AD) during deployment:

- A new Microsoft SQL service account is created in the domain if you don't provide an existing SQL service account.
- The Windows cluster, node host names, and Microsoft SQL FCI name are added as managed computers to the Microsoft SQL service account.
- The Windows cluster entry is assigned permissions to add computers to the domain.

## User-managed Active Directory security groups

If you select “user-managed Active Directory” during Microsoft SQL Server deployment in Workload Factory, you must provide a security group that allows traffic between the EC2 instances to the directory service for deployment. Workload Factory doesn’t automatically attach the security group for user-managed Active Directory like it does for AWS Managed Microsoft AD.

## Resource rollback

If you decide to rollback your Domain Name System (DNS) resources, the resource records in AD and DNS are not removed automatically. You can remove the records from the DNS server and AD as follows.

- For user-managed AD, first [remove the AD computer](#). Then, connect to the DNS server from DNS manager and [delete the DNS Resource Records](#).
- For AWS Managed Microsoft AD, [install the AD administration tools](#). Next, [remove the AD computer](#). Lastly, connect to the DNS server from DNS manager and [delete the DNS Resource Records](#).

## Before you begin

Ensure you have the following prerequisites before you create a new database host.

### Credentials and permissions

You must [grant database host creation permissions](#) in your AWS account to create a new database host in Workload Factory.

### Active Directory

When connecting to Active Directory, you must have administrative access with permissions to do the following:

- Join the domain
- Create Computer Objects
- Create objects in the default Organization Unit (OU)
- Read all properties
- Make the domain user a local admin on the AD nodes
- Create a Microsoft SQL Server service user in the AD, if it doesn’t exist already

## Step 1: Create a database server

You can use *Quick create* or *Advanced create* deployment modes to complete this task in Workload Factory with *Automate* mode permissions. You can also use the following tools available in the Codebox: REST API, AWS CLI, AWS CloudFormation, and Terraform. [Learn how to use Codebox for automation](#).

 When using Terraform from Codebox, the code you copy or download hides `fsxadmin` and `vsadmin` passwords. You’ll need to re-enter the passwords when you run the code. You’ll need to include the following permissions for the user account in addition to *Automate* mode permissions: `iam:TagRole` and `iam:TagInstanceProfile`. [Learn how to use Terraform from Codebox](#).

During deployment, Workload Factory enables CredSSP for credential delegation to scripts for provisioning SQL. When the CredSSP delegation is blocked for all domain computers with the group policy, deployment fails. Post-deployment, Workload Factory disables CredSSP.

## Quick create



In **Quick create**, FCI is the default deployment model, Windows 2016 is the default Windows version, and SQL 2019 Standard Edition is the default SQL version.

### Steps

1. Log in using one of the [console experiences](#).
2. In the Databases tile, select **Deploy host** and then select **Microsoft SQL Server** from the menu.
3. Select **Quick create**.
4. Under **AWS settings**, provide the following:

- a. **AWS credentials**: Select AWS credentials with automate permissions to deploy the new database host.

AWS credentials with *read/write* permissions let Workload Factory deploy and manage the new database host from your AWS account within Workload Factory.

AWS credentials with *read-only* permissions let Workload Factory generate a CloudFormation template for you to use in the AWS CloudFormation console.

If you don't have AWS credentials associated in Workload Factory and you want to create the new server in Workload Factory, follow **Option 1** to go to the Credentials page. Manually add the required credentials and permissions for *read/write* mode for Database workloads.

If you want to complete the create new server form in Workload Factory so you can download a complete YAML file template for deployment in AWS CloudFormation, follow **Option 2** to ensure you have the required permissions to create the new server within AWS CloudFormation. Manually add the required credentials and permissions for *read* mode for Database workloads.

Optionally, you can download a partially completed YAML file template from the Codebox to create the stack outside Workload Factory without any credentials or permissions. Select **CloudFormation** from the dropdown in the Codebox to download the YAML file.

- b. **Region & VPC**: Select a Region and VPC network.

Ensure deployment subnets are associated with existing interface endpoints and security groups allow access to HTTPS (443) protocol to the selected subnets.

AWS service interface endpoints (SQS, FSx, EC2, CloudWatch, CloudFormation, SSM) and the S3 gateway endpoint are created during deployment if not found.

**VPC DNS attributes** `EnableDnsSupport` and `EnableDnsHostnames` are modified to enable endpoint address resolution if they aren't already set to `true`.

When using a cross-VPC DNS, the security group for endpoints on the other VPC where DNS resides should allow port 443 to deployment subnets. If not, you should provide a DNS resolver from the local VPC when joining a cross-VPC Active Directory. In a multiple replicated Domain Controller environment, if some domain controllers are not reachable from the subnet, you can **Redirect to CloudFormation** and enter `Preferred domain controller` to connect to Active Directory.

- c. **Availability zones**: Select availability zones and subnets according to the Failover Cluster Instance (FCI) deployment model.



FCI deployments are only supported on Multiple Availability Zone (MAZ) FSx for ONTAP configurations.

- i. In the **Cluster configuration - Node 1** field, select the primary availability zone for the MAZ FSx for ONTAP configuration from the **Availability zone** dropdown menu and a subnet from the primary availability zone from the **Subnet** dropdown menu.
- ii. In the **Cluster configuration - Node 2** field, select the secondary availability zone for the MAZ FSx for ONTAP configuration from the **Availability zone** dropdown menu and a subnet from the secondary availability zone from the **Subnet** dropdown menu.

5. Under **Application settings**, enter a user name and password for **Database credentials**.

6. Under **Connectivity**, provide the following:

- a. **Key pair**: Select a key pair.
- b. **Active Directory**:
  - i. In the **Domain name** field, select or enter a name for the domain.
    - A. For AWS-managed Active Directories, domain names appear in the dropdown menu.
    - B. For a user-managed Active Directory, enter a name in the **Search and Add** field, and click **Add**.
  - ii. In the **DNS address** field, enter the DNS IP address for the domain. You can add up to 3 IP addresses.

For AWS-managed Active Directories, the DNS IP address(es) appear in the dropdown menu.
  - iii. In the **User name** field, enter the user name for the Active Directory domain.
  - iv. In the **Password** field, enter a password for the Active Directory domain.

7. Under **Infrastructure settings**, provide the following:

- a. **FSx for ONTAP system**: Create a new FSx for ONTAP file system or use an existing FSx for ONTAP file system.
  - i. **Create new FSx for ONTAP**: Enter user name and password.

A new FSx for ONTAP file system may add 30 minutes or more of installation time.
  - ii. **Select an existing FSx for ONTAP**: Select FSx for ONTAP name from the dropdown menu, and enter a user name and password for the file system.

For existing FSx for ONTAP file systems, ensure the following:
    - The routing group attached to FSx for ONTAP allows routes to the subnets to be used for deployment.
    - The security group allows traffic from the subnets used for deployment, specifically HTTPS (443) and iSCSI (3260) TCP ports.

b. **Data drive size**: Enter the data drive capacity and select the capacity unit.

8. Summary:

- a. **Preview default**: Review the default configurations set by Quick create.
- b. **Estimated cost**: Provides an estimate of charges that you might incur if you deployed the resources shown.

## 9. Click **Create**.

Alternatively, if you want to change any of these default settings now, create the database server with Advanced create.

You can also select **Save configuration** to deploy the host later.

### Advanced create

#### Steps

##### 1. Log in using one of the [console experiences](#).

In the Databases tile, select **Deploy host** and then select **Microsoft SQL Server** from the menu.

##### 2. Select **Advanced create**.

##### 3. For **Deployment model**, select **Failover Cluster Instance** or **Single instance**.

##### 4. Under **AWS settings**, provide the following:

- AWS credentials:** Select AWS credentials with automate permissions to deploy the new database host.

AWS credentials with *read/write* permissions let Workload Factory deploy and manage the new database host from your AWS account within Workload Factory.

AWS credentials with *read-only* permissions let Workload Factory generate a CloudFormation template for you to use in the AWS CloudFormation console.

If you don't have AWS credentials associated in Workload Factory and you want to create the new server in Workload Factory, follow **Option 1** to go to the Credentials page. Manually add the required credentials and permissions for *read/write* mode for Database workloads.

If you want to complete the create new server form in Workload Factory so you can download a complete YAML file template for deployment in AWS CloudFormation, follow **Option 2** to ensure you have the required permissions to create the new server within AWS CloudFormation.

Manually add the required credentials and permissions for *read-only* mode for Database workloads.

Optionally, you can download a partially completed YAML file template from the Codebox to create the stack outside Workload Factory without any credentials or permissions. Select **CloudFormation** from the dropdown in the Codebox to download the YAML file.

- Region & VPC:** Select a Region and VPC network.

Ensure security groups for an existing interface endpoint allow access to HTTPS (443) protocol to the selected subnets.

AWS Service interface endpoints (SQS, FSx, EC2, CloudWatch, Cloud Formation, SSM) and S3 gateway endpoint are created during deployment if not found.

VPC DNS attributes `EnableDnsSupport` and `EnableDnsHostnames` are modified to enable resolve endpoint address resolution if not already set to `true`.

- Availability zones:** Select availability zones and subnets according to the deployment model you selected. Subnets should not share the same route table for high availability.



FCI deployments are only supported on Multiple Availability Zone (MAZ) FSx for ONTAP configurations.

- For single instance deployments:

- In the **Cluster configuration - Node 1** field, select an availability zone from the **Availability zone** from the dropdown menu and a subnet from the **Subnet** dropdown menu.

- For FCI deployments:

- In the **Cluster configuration - Node 1** field, select the primary availability zone for the MAZ FSx for ONTAP configuration from the **Availability zone** dropdown menu and a subnet from the primary availability zone from the **Subnet** dropdown menu.
  - In the **Cluster configuration - Node 2** field, select the secondary availability zone for the MAZ FSx for ONTAP configuration from the **Availability zone** dropdown menu and a subnet from the secondary availability zone from the **Subnet** dropdown menu.

d. **Security group:** Select an existing security group or create a new security group.

Three security groups get attached to the SQL nodes (EC2 instances) during new server deployment.

1. A workload security group is created to allow ports and protocols required for Microsoft SQL and Windows cluster communication on nodes.
2. In case of AWS-managed Active Directory, the security group attached to the directory service gets automatically added to the Microsoft SQL nodes to allow communication with Active Directory.
3. For an existing FSx for ONTAP file system, the security group associated with it is added automatically to the SQL nodes which allows communication to the file system. When a new FSx for ONTAP system is created, a new security group is created for the FSx for ONTAP file system and the same security group also gets attached to SQL nodes.

For a user-managed Active Directory, ensure the security group configured on the AD instance allows traffic from subnets used for deployment. The security group should allow communication to the Active Directory domain controllers from the subnets where EC2 instances for Microsoft SQL are configured.

5. Under **Application settings**, provide the following:

a. Under **SQL Server install type**, select **License included AMI** or **Use custom AMI**.

- i. If you select **License included AMI**, provide the following:

- A. **Operating system:** Select **Windows server 2016**, **Windows server 2019**, or **Windows server 2022**.
  - B. **Database edition:** Select **SQL Server Standard Edition** or **SQL Server Enterprise Edition**.
  - C. **Database version:** Select **SQL Server 2016**, **SQL Server 2019**, or **SQL Server 2022**.
  - D. **SQL Server AMI:** Select a SQL Server AMI from the dropdown menu.

- ii. If you select **Use custom AMI**, select an AMI from the dropdown menu.

b. **SQL Server collation:** Select a collation set for the server.



If the selected collation set isn't compatible for installation, we recommend that you select the default collation "SQL\_Latin1\_General\_CI\_AS".

- c. **Database name:** Enter the database cluster name.
- d. **Database credentials:** Enter a user name and password for a new service account or use existing service account credentials in the Active Directory.

Optional: Select to **Use managed service account** for the SQL Server service account. Use this option if your environment uses MSA (Managed Service Account) or Group Managed Service Accounts (gMSA) where password management is handled by Active Directory.

6. Under **Connectivity**, provide the following:

- a. **Key pair:** Select a key pair to connect securely to your instance.
- b. **Active Directory:** Provide the following Active Directory details:
  - i. In the **Domain name** field, select or enter a name for the domain.
    - A. For AWS-managed Active Directories, domain names appear in the dropdown menu.
    - B. For a user-managed Active Directory, enter a name in the **Search and Add** field, and click **Add**.
  - ii. In the **DNS address** field, enter the DNS IP address for the domain. You can add up to 3 IP addresses.

For AWS-managed Active Directories, the DNS IP address(es) appear in the dropdown menu.
  - iii. In the **User name** field, enter the user name for the Active Directory domain.
  - iv. In the **Password** field, enter a password for the Active Directory domain.
  - v. **Preferred domain controller:** Optionally, enter the preferred domain controller to use for the Active Directory to join.
  - vi. **Preferred organizational unit path:** Optionally, enter the preferred organizational unit (OU) in the Active Directory to join.
  - vii. **Target Active Directory group:** Optionally, enter the target Active Directory group to add the computers to.

7. Under **Infrastructure settings**, provide the following:

- a. **DB Instance type:** Select the database instance type from the dropdown menu.
- b. **FSx for ONTAP system:** Create a new FSx for ONTAP file system or use an existing FSx for ONTAP file system.
  - i. **Create new FSx for ONTAP:** Enter user name and password.

A new FSx for ONTAP file system may add 30 minutes or more of installation time.

- ii. **Select an existing FSx for ONTAP:** Select FSx for ONTAP name from the dropdown menu, and enter a user name and password for the file system.

For existing FSx for ONTAP file systems, ensure the following:

- The routing group attached to FSx for ONTAP allows routes to the subnets to be used for deployment.
- The security group allows traffic from the subnets used for deployment, specifically HTTPS (443) and iSCSI (3260) TCP ports.

- c. **Snapshot policy:** Enabled by default. Snapshots are taken daily and have a 7-day retention period.

The snapshots are assigned to volumes created for SQL workloads.

- d. **Data drive size:** Enter the data drive capacity and select the capacity unit.
- e. **Provisioned IOPS:** Select **Automatic** or **User-provisioned**. If you select **User-provisioned**, enter the IOPS value.
- f. **Throughput capacity:** Select the throughput capacity from the dropdown menu.

In certain regions, you may select 4 GBps throughput capacity. To provision 4 GBps of throughput capacity, your FSx for ONTAP file system must be configured with a minimum of 5,120 GiB of SSD storage capacity and 160,000 SSD IOPS.

- g. **Encryption:** Select a key from your account or a key from another account. You must enter the encryption key ARN from another account.

FSx for ONTAP custom encryption keys aren't listed based on service applicability. Select an appropriate FSx encryption key. Non-FSx encryption keys will cause server creation failure.

AWS-managed keys are filtered based on service applicability.

- h. **Tags:** Optionally, you can add up to 40 tags.

- i. **Simple Notification Service:** Optionally, you can enable the Simple Notification Service (SNS) for this configuration by selecting an SNS topic for Microsoft SQL Server from the dropdown menu.

- i. Enable the Simple Notification Service.
- ii. Select an ARN from the dropdown menu.

- j. **CloudWatch monitoring:** Optionally, you can enable CloudWatch monitoring.

We recommend enabling CloudWatch for debugging in case of failure. The events that appear in the AWS CloudFormation console are high-level and don't specify the root cause. All detailed logs are saved in the C:\cfn\logs folder in the EC2 instances.

In CloudWatch, a log group is created with the name of the stack. A log stream for every validation node and SQL node appear under the log group. CloudWatch shows script progress and provides information to help you understand if and when deployment fails.

- k. **Resource rollback:** This feature isn't currently supported.
8. Summary
  - a. **Estimated cost:** Provides an estimate of charges that you might incur if you deployed the resources shown.
9. Click **Create** to deploy the new database host.

Alternatively, you can save the configuration.

## Step 2: Enable remote connection on the Microsoft SQL Server

After the server deploys, Workload Factory does not enable remote connection on the Microsoft SQL Server. To enable the remote connection, complete the following steps.

### Steps

1. Use computer identity for NTLM by referring to [Network security: Allow Local System to use computer identity for NTLM](#) in Microsoft documentation.
2. Check dynamic port configuration by referring to [A network-related or instance-specific error occurred while establishing a connection to SQL Server](#) in Microsoft documentation.
3. Allow the required client IP or subnet in the security group.

#### What's next

Now you can [create a database in Workload Factory for Databases](#).

## Create a PostgreSQL server in NetApp Workload Factory

Creating a new PostgreSQL server, or database host, in NetApp Workload Factory for Databases requires an FSx for ONTAP file system deployment and resources for Active Directory.

#### About this task

Before creating a PostgreSQL server from Workload Factory, learn about the available storage deployment types for the database host configuration, workload factory modes of operation, and the requirements to complete this operation.

#### FSx for ONTAP file system deployments

Creating a new PostgreSQL server requires an FSx for ONTAP file system as the storage backend. You can use an existing FSx for ONTAP file system or create a new file system. If you select an existing FSx for ONTAP file system as your database server storage backend, we create a new storage VM for the PostgreSQL workloads.

+

FSx for ONTAP file systems have two PostgreSQL server deployment models: *High Availability (HA)* or *single instance*. Different resources are created for the FSx for ONTAP file system depending on the FSx for ONTAP deployment model you select.

- **High Availability (HA) deployment:** A Multiple Availability Zone FSx for NetApp ONTAP file system is deployed when a new FSx for ONTAP file system is selected for HA deployment. Separate volumes and LUNs are created for data, log, and tempdb files for an HA deployment. An additional volume and LUN are created for Quorum or witness disk for Windows cluster. HA deployment configures Streaming replication between the primary and secondary PostgreSQL servers.
- **Single instance deployment:** A Single Availability Zone FSx for ONTAP file system is created when a new PostgreSQL server is created. In addition, separate volumes and LUNs are created for data, log, and tempdb files.

#### Before you begin

You must have [grant database host creation permissions](#) in your AWS account to create a new database host in workload factory.

#### Create a PostgreSQL server

You can use *Quick create* or *Advanced create* deployment modes to complete this task in workload factory with *Automate* mode permissions. You can also use the following tools available in the Codebox: REST API, AWS CLI, AWS CloudFormation, and Terraform. [Learn how to use Codebox for automation](#).

 When using Terraform from Codebox, the code you copy or download hides `fsxadmin` and `vsadmin` passwords. You'll need to re-enter the passwords when you run the code. You'll need to include the following permissions for the user account in addition to *Automate* mode permissions: `iam:TagRole` and `iam:TagInstanceProfile`. [Learn how to use Terraform from Codebox](#).

## Quick create



In **Quick create**, HA is the default deployment model, Windows 2016 is the default Windows version, and SQL 2019 Standard Edition is the default SQL version.

### Steps

1. Log in using one of the [console experiences](#).
2. In the Databases tile, select **Deploy host** and then select **PostgreSQL Server** from the menu.
3. Select **Quick create**.
4. Under **Landing zone**, provide the following:

- a. **AWS credentials**: Select AWS credentials with automate permissions to deploy the new database host.

AWS credentials with *read/write* permissions let workload factory deploy and manage the new database host from your AWS account within workload factory.

AWS credentials with *read-only* permissions let workload factory generate a CloudFormation template for you to use in the AWS CloudFormation console.

If you don't have AWS credentials associated in workload factory and you want to create the new server in workload factory, follow **Option 1** to go to the Credentials page. Manually add the required credentials and permissions for *read/write* mode for Database workloads.

If you want to complete the create new server form in workload factory so you can download a complete YAML file template for deployment in AWS CloudFormation, follow **Option 2** to ensure you have the required permissions to create the new server within AWS CloudFormation. Manually add the required credentials and permissions for *read-only* mode for Database workloads.

Optionally, you can download a partially completed YAML file template from the Codebox to create the stack outside workload factory without any credentials or permissions. Select **CloudFormation** from the dropdown in the Codebox to download the YAML file.

- b. **Region & VPC**: Select a Region and VPC network.

Ensure security groups for an existing interface endpoint allow access to HTTPS (443) protocol to the selected subnets.

AWS service interface endpoints (SQS, FSx, EC2, CloudWatch, CloudFormation, SSM) and the S3 gateway endpoint are created during deployment if not found.

VPC DNS attributes `EnableDnsSupport` and `EnableDnsHostnames` are modified to enable endpoint address resolution if they aren't already set to `true`.

- c. **Availability zones**: Select availability zones and subnets.



HA deployments are only supported on Multiple Availability Zone (MAZ) FSx for ONTAP configurations.

Subnets should not share the same route table for high availability.

- i. In the **Cluster configuration - Node 1** field, select the primary availability zone for the MAZ FSx for ONTAP configuration from the **Availability zone** dropdown menu and a subnet from the primary availability zone from the **Subnet** dropdown menu.
- ii. In the **Cluster configuration - Node 2** field, select the secondary availability zone for the MAZ FSx for ONTAP configuration from the **Availability zone** dropdown menu and a subnet from the secondary availability zone from the **Subnet** dropdown menu.

5. Under **Application settings**, enter a user name and password for **Database credentials**.
6. Under **Connectivity**, select a key pair to connect securely to your instance.
7. Under **Infrastructure settings**, provide the following:
  - a. **FSx for ONTAP system**: Create a new FSx for ONTAP file system or use an existing FSx for ONTAP file system.
    - i. **Create new FSx for ONTAP**: Enter user name and password.
 

A new FSx for ONTAP file system may add 30 minutes or more of installation time.
    - ii. **Select an existing FSx for ONTAP**: Select FSx for ONTAP name from the dropdown menu, and enter a user name and password for the file system.

For existing FSx for ONTAP file systems, ensure the following:

- The routing group attached to FSx for ONTAP allows routes to the subnets to be used for deployment.
- The security group allows traffic from the subnets used for deployment, specifically HTTPS (443) and iSCSI (3260) TCP ports.

- b. **Data drive size**: Enter the data drive capacity and select the capacity unit.

## 8. Summary:

- a. **Preview default**: Review the default configurations set by Quick create.
- b. **Estimated cost**: Provides an estimate of charges that you might incur if you deployed the resources shown.

## 9. Click **Create**.

Alternatively, if you want to change any of these default settings now, create the database server with Advanced create.

You can also select **Save configuration** to deploy the host later.

## Advanced create

### Steps

1. Log in using one of the [console experiences](#).
2. In the Databases tile, select **Deploy host** and then select **PostgreSQL Server** from the menu.
3. Select **Advanced create**.
4. Under **Deployment model**, select **Standalone instance or High availability (HA)**.
5. Under **Landing zone**, provide the following:
  - a. **AWS credentials**: Select AWS credentials with automate permissions to deploy the new database host.

AWS credentials with *automate* permissions let workload factory deploy and manage the new database host from your AWS account within workload factory.

AWS credentials with *read-only* permissions let workload factory generate a CloudFormation template for you to use in the AWS CloudFormation console.

If you don't have AWS credentials associated in workload factory and you want to create the new server in workload factory, follow **Option 1** to go to the Credentials page. Manually add the required credentials and permissions for *read/write* mode for Database workloads.

If you want to complete the create new server form in workload factory so you can download a complete YAML file template for deployment in AWS CloudFormation, follow **Option 2** to ensure you have the required permissions to create the new server within AWS CloudFormation.

Manually add the required credentials and permissions for *read-only* mode for Database workloads.

Optionally, you can download a partially completed YAML file template from the Codebox to create the stack outside workload factory without any credentials or permissions. Select **CloudFormation** from the dropdown in the Codebox to download the YAML file.

b. **Region & VPC:** Select a Region and VPC network.

Ensure security groups for an existing interface endpoint allow access to HTTPS (443) protocol to the selected subnets.

AWS Service interface endpoints (SQS, FSx, EC2, CloudWatch, Cloud Formation, SSM) and S3 gateway endpoint are created during deployment if not found.

VPC DNS attributes `EnableDnsSupport` and `EnableDnsHostnames` are modified to enable resolve endpoint address resolution if not already set to `true`.

c. **Availability zones:** Select availability zones and subnets.

#### For single instance deployments

In the **Cluster configuration - Node 1** field, select an availability zone from the **Availability zone** dropdown menu and a subnet from the **Subnet** dropdown menu.

#### For HA deployments

i. In the **Cluster configuration - Node 1** field, select the primary availability zone for the MAZ FSx for ONTAP configuration from the **Availability zone** dropdown menu and a subnet from the primary availability zone from the **Subnet** dropdown menu.

ii. In the **Cluster configuration - Node 2** field, select the secondary availability zone for the MAZ FSx for ONTAP configuration from the **Availability zone** dropdown menu and a subnet from the secondary availability zone from the **Subnet** dropdown menu.

d. **Security group:** Select an existing security group or create a new security group.

Two security groups get attached to the SQL nodes (EC2 instances) during new server deployment.

1. A workload security group is created to allow ports and protocols required for PostgreSQL.
2. For a new FSx for ONTAP file system, a new security group is created and attached to the SQL node. For an existing FSx for ONTAP file system, the security group associated with it is

added automatically to the PostgreSQL node which allows communication to the file system.

6. Under **Application settings**, provide the following:
  - a. Select the **Operating system** from the dropdown menu.
  - b. Select the **PostgreSQL version** from the dropdown menu.
  - c. **Database server name**: Enter the database cluster name.
  - d. **Database credentials**: Enter a user name and password for a new service account or use existing service account credentials in the Active Directory.
7. Under **Connectivity**, select a key pair to connect securely to your instance.
8. Under **Infrastructure settings**, provide the following:
  - a. **DB Instance type**: Select the database instance type from the dropdown menu.
  - b. **FSx for ONTAP system**: Create a new FSx for ONTAP file system or use an existing FSx for ONTAP file system.
    - i. **Create new FSx for ONTAP**: Enter user name and password.

A new FSx for ONTAP file system may add 30 minutes or more of installation time.
    - ii. **Select an existing FSx for ONTAP**: Select FSx for ONTAP name from the dropdown menu, and enter a user name and password for the file system.

For existing FSx for ONTAP file systems, ensure the following:

      - The routing group attached to FSx for ONTAP allows routes to the subnets to be used for deployment.
      - The security group allows traffic from the subnets used for deployment, specifically HTTPS (443) and iSCSI (3260) TCP ports.
  - c. **Snapshot policy**: Enabled by default. Snapshots are taken daily and have a 7-day retention period.

The snapshots are assigned to volumes created for PostgreSQL workloads.
  - d. **Data drive size**: Enter the data drive capacity and select the capacity unit.
  - e. **Provisioned IOPS**: Select **Automatic** or **User-provisioned**. If you select **User-provisioned**, enter the IOPS value.
  - f. **Throughput capacity**: Select the throughput capacity from the dropdown menu.

In certain regions, you may select 4 GBps throughput capacity. To provision 4 GBps of throughput capacity, your FSx for ONTAP file system must be configured with a minimum of 5,120 GiB of SSD storage capacity and 160,000 SSD IOPS.
  - g. **Encryption**: Select a key from your account or a key from another account. You must enter the encryption key ARN from another account.

FSx for ONTAP custom encryption keys aren't listed based on service applicability. Select an appropriate FSx encryption key. Non-FSx encryption keys will cause server creation failure.

AWS-managed keys are filtered based on service applicability.
  - h. **Tags**: Optionally, you can add up to 40 tags.

- i. **Simple Notification Service:** Optionally, you can enable the Simple Notification Service (SNS) for this configuration by selecting an SNS topic for Microsoft SQL Server from the dropdown menu.
  - i. Enable the Simple Notification Service.
  - ii. Select an ARN from the dropdown menu.
- j. **CloudWatch monitoring:** Optionally, you can enable CloudWatch monitoring.

We recommend enabling CloudWatch for debugging in case of failure. The events that appear in the AWS CloudFormation console are high-level and don't specify the root cause. All detailed logs are saved in the C:\cfn\logs folder in the EC2 instances.

In CloudWatch, a log group is created with the name of the stack. A log stream for every validation node and SQL node appear under the log group. CloudWatch shows script progress and provides information to help you understand if and when deployment fails.

- k. **Resource rollback:** This feature isn't currently supported.
9. Summary
  - a. **Estimated cost:** Provides an estimate of charges that you might incur if you deployed the resources shown.
10. Click **Create** to deploy the new database host.

Alternatively, you can save the configuration.

## What's next

You can manually configure users, remote access, and databases on the deployed PostgreSQL server.

# Manage resources

## Resource management in NetApp Workload Factory for Databases

Managing resources in NetApp Workload Factory for Databases allows you to use advanced features including database and clone creation, resource utilization and monitoring. Additionally, you can analyze the well-architected status of your database configurations and implement configuration best practices to improve performance and lower operational costs. Resource management is only for Microsoft SQL Server and Oracle environments running on FSx for ONTAP file system storage.

You must [register resources](#) to do any of the following management tasks.

Management tasks include:

- Viewing databases from the Inventory
- [Creating a database](#)
- [Creating a database clone \(sandbox\)](#)
- [Implementing well-architected database configurations](#)

## Register resources in NetApp Workload Factory for Databases

Register instances for Microsoft SQL Server and databases for Oracle so that you can monitor instance and database status, resource utilization, protection, and storage performance in NetApp Workload Factory for Databases.

You can register your resources only if they run on FSx for ONTAP file system storage.

### About the task

Registering an instance (SQL Server) or database (Oracle) has three steps - instance or database authentication, FSx for ONTAP authentication, and preparation. Preparation involves making sure that all AWS, NetApp, and PowerShell modules are installed on the instance or database, and that the minimum requirements for Workload Factory for Databases features like [error log analysis](#) or [well-architected review](#) are met.

Workload Factory supports only Microsoft SQL Server instance and Oracle database registration and management. Depending on the AWS account credentials you select in Workload Factory, PostgreSQL hosts might appear in the Inventory. Currently, Workload Factory supports unregistered PostgreSQL instances running only on Amazon Linux operating systems.

### Before you begin

The host for the instance or database must appear in the Inventory. For hosts to appear in the inventory, you must [grant view, planning, and analysis permissions](#) in your AWS account.

#### Registering an instance in a private network

To register an instance (SQL Server) or database (Oracle) in a private network with no external connectivity, the following endpoints need to be available in the VPC with association to the subnets where SQL servers are present. Ensure the interface endpoints allow port 443 in the attached Security Group.

- S3 Gateway/endpoint
- ssm
- ssmessages
- fsx

If you use a proxy server for all outbound connections from EC2 instances, you must allow access to the following domains so that management operations work:

- [.microsoft.com](https://microsoft.com) (SQL Server)
- [.powershellgallery.com](https://powershellgallery.com) (SQL Server)
- [.aws.amazon.com](https://aws.amazon.com)
- [.amazonaws.com](https://amazonaws.com)

### Register a Microsoft SQL Server instance

Registering an instance has three steps - instance authentication, FSx for ONTAP authentication, and preparation to complete missing prerequisites. You can register single or multiple instances.

Workload Factory supports registration for Failover Cluster Instance (FCI) and Standalone deployment for SQL Server.

## Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Databases**.
3. From the Databases menu, select **Inventory**.
4. In the Inventory, select **Microsoft SQL Server** as the engine type.
5. Select the **Instances** tab.
6. Select to register a single instance or multiple instances.
7. To Authenticate instances (step 1), do the following and then select **Next**:
  - a. Select to **Use the same credentials for all instances** or **Manage credentials manually**.
  - b. Authenticate SQL Server and Windows by providing user name and password information.

If instances are authenticated, select **Next**.
8. To Authenticate FSx for ONTAP (step 2), do the following:
  - a. Select to **Use the same credentials for all resources** or **Manage credentials manually**.
  - b. Enter the FSx for ONTAP file system user name and password, and then select **Next**.

If the FSx for ONTAP file system is authenticated, select **Next**.
9. To Prepare (step 3), make sure the instance(s) meets the minimum requirements.

To meet the minimum requirements, the instance must have AWS and NetApp PowerShell modules and PowerShell 7 modules installed, and you must complete the prerequisites for at least one of the capabilities listed under Prerequisite check.

- a. Review the prerequisites in the **Prerequisite check view**.

You must complete all prerequisites for a single capability like **Review well-architected issues and recommendations** to register the instance.

- b. Select **Setup details** for each capability to learn about the capability prerequisites and follow the on-screen instructions to complete any missing prerequisites for a capability.

To have Workload Factory [review and fix well-architected issues](#) for your instances, complete all prerequisites listed under **Review well-architected issues and recommendations** and **Fix well-architected issues** capabilities.

10. When prerequisites are complete, **Register** the instance(s).

## Result

Instance registration initiates. Select the **Job monitoring** tab to track progress.

## Register an Oracle database

Registering an instance has three steps - database authentication, FSx for ONTAP authentication, and preparation to complete missing prerequisites. You can register single or multiple databases.

## Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Databases**.
3. From the Databases menu, select **Inventory**.
4. In the Inventory, select **Oracle** as the engine type.
5. Select the **Databases** tab.
6. Select to register a single database or multiple databases.
7. To Authenticate databases (step 1), do the following:
  - Select to **Use the same credentials for all instances** or **Manage credentials manually**.
  - Authenticate Oracle user and Automatic Storage Management (ASM) grid user (if applicable) by providing user name and password information.If databases are authenticated, select **Next**.
8. To Authenticate FSx for ONTAP (step 2), do the following and then select **Next**:
  - Select to **Use the same credentials for all resources** or **Manage credentials manually**.
  - Enter the FSx for ONTAP file system user name and password.If the FSx for ONTAP file system is authenticated, select **Next**.
9. To Prepare (step 3), make sure the database(s) meets required prerequisites.  
If all required modules are installed and prerequisites are met, select **Next** to register the database. Otherwise, follow these steps.
  - a. Review the prerequisites in the **Prerequisite check view**.  
You must complete all prerequisites for a single capability like **Review well-architected issues and recommendations** to register the database.
  - b. Select **Setup details** for each capability to learn about the capability prerequisites and follow the on-screen instructions to complete any missing prerequisites for a capability.  
To have Workload Factory [review and fix well-architected issues](#) for your databases, complete all prerequisites listed under **Review well-architected issues and recommendations** and **Fix well-architected issues** capabilities.
10. When prerequisites are complete, **Register** the database(s).

## Result

Database registration initiates. Select the **Job monitoring** tab to track progress.

## What's next

After resource registration, you can perform the following tasks.

- View databases from the inventory
- [Create a database](#)
- [Create a database clone \(sandbox\)](#)
- [Implement well-architected database configurations](#)

## Create a Microsoft SQL database in NetApp Workload Factory for Databases

Creating a new Microsoft SQL database enables you to manage the resource within NetApp Workload Factory for Databases.

### About this task

Upon database creation, two new volumes are created in the FSx for ONTAP file system consisting of independent LUNs to host data and log files for the database. The database files in the new database are thin-provisioned and consume only a few MBs of the total size allocated for the new database.

If you want to segregate storage for the database, you can do this by using a *virtual mount point*. The virtual mount point lets you consolidate databases to a few common drives on the host.

Creating a database in workload factory requires *view, planning, and analysis* permissions. Alternatively, you can copy or download a partially completed code template to complete the operation outside workload factory. [Learn about Workload Factory permissions](#) to decide which mode you'd like to use.



Microsoft SQL Servers using SMB protocol don't support database creation.

### Before you begin

Ensure you complete the following prerequisites before you create a new database.

- **Credentials and permissions:** You must have [AWS account credentials and view, planning, and analysis permissions](#) to create a new database in Workload Factory.

Alternatively, you can use the Codebox to copy a template so that you can deploy a database outside of workload factory using REST API. [Learn more about Codebox automation](#).

- **Windows host:** You must have enough drive letters available on the Microsoft SQL Server to create new drives for the new database if you use *Quick create* mode.
- **Microsoft SQL Server:** You must have a managed Microsoft SQL Server in workload factory for Databases to host the new database.
- **AWS Systems Manager:** Ensure the `NT Authority\SYSTEM` user privilege is enabled in the Microsoft SQL host via AWS Systems Manager.

### Create a database

You can use *Quick create* or *Advanced create* deployment modes to complete this task in Workload Factory.

## Quick create

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Databases**.
3. From the Databases menu, select **Inventory**.
4. In the Inventory, select **Microsoft SQL Server** as the database engine type.
5. Select a database server with a managed SQL server instance to create the database in.
6. Click the action menu of the managed instance and then select **Create user database**.
7. On the Create user database page, under Database information, provide the following:
  - a. **Database name**: Enter name for the database.
  - b. **Collation**: Select a collation for the database. The default collation "SQL\_Latin1\_General\_CI\_AS" on Microsoft SQL Server is selected.
8. Under File settings, provide the following:
  - a. **File settings mode**: Select **Quick create**.
  - b. **File names & path**:
    - **Data file name**: Enter the data file name.
    - **Log file name**: Enter the log file name.
  - c. **File sizes**: Enter the data size and log size for the database.
9. Click **Create**.

Alternatively, if you want to change any of these default settings now, change the **File settings mode** to **Advanced create**.

## Advanced create

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Databases**.
3. From the Databases menu, select **Inventory**.
4. In the Inventory, select **Microsoft SQL Server** as the database engine type.
5. Select a database server with a managed SQL server instance to create the database in.
6. Click the action menu of the managed instance and then select **Create user database**.
7. Select **Create user database**.
8. On the Create user database page, under Database information, provide the following:
  - a. **Database name**: Enter name for the database.
  - b. **Collation**: Select the collation for the database. The default collation "SQL\_Latin1\_General\_CI\_AS" on Microsoft SQL Server is selected.
9. Under File settings, provide the following:
  - a. **File settings mode**: Select **Advanced create**.
  - b. **File names & path**:

i. **Data file:** Select a drive letter and enter the data file name.

Optionally, click the box for **Virtual mount point**.

ii. **Log file:** Select a drive letter and enter the log file name.

Optionally, click the box for **Virtual mount point**.

c. **File sizes:** Enter the data size and log size for the database.

10. Click **Create**.

If you created the database host, you can check the job's progress in the **Job monitoring** tab.

## Create a sandbox clone in NetApp Workload Factory for Databases

Creating a sandbox clone of a database in NetApp Workload Factory for Databases lets you use the clone for development, testing, integration, analytics, training, QA, and more without altering the source database.

### About this task

A sandbox clone is created from the most recent snapshot on the source database. It may be cloned in the same Microsoft SQL Server as the source database or cloned in another Microsoft SQL Server as long as they share the same FSx for ONTAP file system.

### Before you begin

Ensure you complete the following prerequisites before you create a sandbox clone.

- **Credentials and permissions:** You must have [AWS account credentials and view, planning, and analysis permissions](#) to create a sandbox clone in Workload Factory.

Alternatively, you can use the Codebox to copy a partially completed template or create a completed template so that you can create the sandbox clone outside of Workload Factory using REST API. [Learn more about Codebox automation](#).

- **Microsoft SQL Server:** You must have a managed Microsoft SQL Server in Workload Factory for Databases to host the new sandbox clone.
- **AWS Systems Manager:** Ensure the NT Authority\SYSTEM user privilege is enabled in the Microsoft SQL host via AWS Systems Manager.
- **Source database:** You need a source database available for the clone.

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Databases**.
3. In Databases, select the **Sandboxes** tab.
4. In the Sandboxes tab, select **Create new sandbox**.
5. On the Create new sandbox page, under Database source, provide the following:
  - a. **Source database host:** Select the source database host.
  - b. **Source database instance:** Select the source database instance.

- c. **Source database:** Select the source database to clone from.
6. Under Database target, provide the following:
  - a. **Target database host:** Select a target database host for the sandbox clone that is in the same VPC and has the same FSx for ONTAP file system as the source host.
  - b. **Target database instance:** Select the target database instance for the sandbox clone.
  - c. **Target database:** Enter a name for the sandbox clone.
7. **Mount:** When cloning a SQL database that has multiple data and/or log files, Workload Factory clones all files under the auto-assigned or defined drive letter.

Select one of the following options:

  - a. **Auto-assign mount point**
  - b. **Define mount point path**

Provide the following to define the mount point path:

  - Enter the drive letter for the data file path.
  - Enter the drive letter for the log file path.

8. **Define tag:** Select a tag to define the sandbox clone.

9. Click **Create**.

To check the job's progress, go to the **Job monitoring** tab.

## Automate with Codebox in NetApp Workload Factory for Databases

You can automate host deployment, database creation, and more with Codebox in NetApp Workload Factory for Databases. 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.

## Protect Microsoft SQL Server workloads

Protect your Microsoft SQL Server applications data using NetApp Backup and Recovery from the Workload Factory console. With this integration, you can achieve the following protection goals: back up workloads with local snapshots on local primary Amazon FSx for NetApp ONTAP (FSx for ONTAP) storage, and replicate workloads to secondary FSx for ONTAP storage.

### About this task

Workload Factory automates discovering resources, validating prerequisites, and configuring and installing the Plug-in for Microsoft SQL Server to prepare for protecting your workloads with NetApp Backup and Recovery. The plug-in is a host-side component of NetApp Software that enables you to protect your Microsoft SQL

## Server workloads

NetApp Backup and Recovery leverages NetApp SnapMirror data replication technology to ensure that all the backups are fully synchronized by creating snapshot copies and transferring them to the backup locations.

For details about protection with Backup and Recovery, refer to the [Protect Microsoft SQL workloads overview with Backup and Recovery](#).

## Before you begin

The following requirements must be met to protect Microsoft SQL Server workloads with Backup and Recovery.

- Ensure that your environment meets [the Backup and Recovery SQL Server requirements](#).
- [Complete NetApp Console requirements](#) including setting up, assigning IAM roles, and installing a Console agent.

If you have organization administrator access to the NetApp account, the backup and recovery super admin role is automatically assigned when you [prepare for protection with NetApp Backup and Recovery](#).

- Set the host resolution on the Connector

To discover databases, you must set host resolution on the Connector. On the hosted device, add the mapping of the IP address to the hostname in the /etc/hosts file.

- [Set up licensing for NetApp Backup and Recovery](#)

## Prepare for protection with NetApp Backup and Recovery

Complete the preparation process to protect your Microsoft SQL Server resources with NetApp Backup and Recovery.

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Databases**.
3. From the Databases menu, select **Inventory**.
4. In the Inventory, select **Microsoft SQL Server** as the engine type.
5. Locate the instance that you want to protect, and then select **Protect** from the menu.
6. If prompted, provide Windows credentials with administrative access.

To use NetApp Backup and Recovery for protection, SQL Server instances must be registered in Workload Factory with Windows credentials.

7. If several Console agents are active and available, select the **Console agent** where you want the workload to be registered and protected.
8. To prepare for data protection, Workload Factory automatically registers your SQL Server resources in Backup and Recovery, configures and installs the Plug-in for Microsoft SQL Server, and discovers resources to meet the prerequisites for protecting your SQL Server instance. Select **Start** to begin the process.
9. After meeting the prerequisites, select **Redirect** to access Backup and Recovery.

## What's next

From Backup and Recovery, create a policy to protect your Microsoft SQL Server instance and databases.

[Learn how to create a policy to protect your Microsoft SQL Server instance and databases.](#)

For related information, refer to the [Backup and Recovery documentation](#) for managing Microsoft SQL Server workloads.

## Edit protection for Microsoft SQL Server resources

You can edit protection for Microsoft SQL Server instances and databases that are already protected in NetApp Backup and Recovery. Editing protection allows you to modify the protection policy or schedule for your protected SQL Server instances.

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Databases**.
3. From the Databases menu, select **Inventory**.
4. In the Inventory, select **Microsoft SQL Server** as the engine type.
5. Select the **Databases** tab.
6. Locate the database to edit protection for, and then select **Edit protection** from the menu.

You'll be redirected to Backup and Recovery in the NetApp Console where you can modify the protection policy or schedule.

# Administer and monitor

## Monitor database jobs in Workload Factory for Databases

Track database jobs and monitor databases within NetApp Workload Factory for Databases for improved visibility and control over database operations.

### About this task

Databases provides job monitoring so you can track job progress, and diagnose and troubleshoot in case any failure occurs. You can filter jobs by type and status, find jobs using the search function, and download the jobs table.

Job monitoring supports up to three levels of monitoring depending on the job. For example, for new database and sandbox clone creation, job monitoring tracks parent jobs and sub-jobs.

#### Job monitoring levels

- Level 1 (parent job): Tracks the host deployment job.
- Level 2 (sub-job): Tracks the sub-jobs related to the host deployment parent job.
- Level 3 (task): Lists the sequence of actions taken on each resource.

#### Job status

The job monitoring feature tracks *in progress*, *completed*, *completed with issues*, and *failed* jobs daily, weekly, bi-weekly, and monthly.

#### Job events retention

Job monitoring events are retained in the user interface for 30 days.

## Monitor jobs

Monitor jobs to track the progress of database operations, and diagnose and troubleshoot in case any failure occurs.

#### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Databases**.
3. From the Databases menu, select **Job monitoring**.
4. In Job monitoring, use the filters or search to narrow job results. You can also download a jobs report.
5. Optionally, select the action menu of the job and click **Go to CloudFormation** to view the job log in the AWS CloudFormation console.

## Implement configuration best practices

### Configuration analysis for database environments in Workload Factory

Workload Factory for Databases analyzes database configurations regularly to determine if there are any issues with Microsoft SQL Server and Oracle deployments on Amazon

FSx for NetApp ONTAP storage. When issues are found, Workload Factory shows you what the issues are and explains what needs to change to ensure your database configurations achieve peak performance, cost efficiency, and compliance with best practices.

Key capabilities include:

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

## Components of the configuration analysis

The configuration analysis includes the following components:

### Well-architected status

The well-architected status refers to the overall health of your database configurations; configurations are categorized as "optimized", "not optimized", or "over-provisioned". The well-architected status of a configuration is referenced from within a Microsoft SQL Server instance or Oracle database in the Workload Factory console.

### Well-architected score

The score includes all currently analyzed configurations and appears as a percentage. A 25% score means that 25% of your database deployments are well-architected. The well-architected score is referenced from the **Well-architected** screen and from within a Microsoft SQL Server instance or Oracle database in the Workload Factory console.

### Configuration categories

The configurations are organized into the following categories: storage, compute, application, resiliency, and cloning. Each category includes specific configuration assessments that are analyzed regularly.

Configuration categories are referenced from the **Well-architected** screen and from within a Microsoft SQL Server instance or Oracle database in the Workload Factory console.

### Analysis scope

The scope of optimization is different depending on which component is being assessed. For example, storage optimization occurs at the SQL instance level whereas compute optimization occurs at the host level.

### Analysis requirements

For a complete database environment analysis, resources must be registered and online.

[Learn how to register resources.](#)

### What's next

[Implement well-architected database configurations](#)

## Implement well-architected database configurations in Workload Factory

Using configuration analysis insights and recommendations, leverage NetApp Workload Factory to implement best practices for your database configurations with Microsoft SQL Server and Oracle. You can easily review the well-architected status, learn about issues with database configurations, and fix configurations that aren't optimized for reliability, security, efficiency, performance, and cost.

You can also dismiss the analysis of specific configurations that don't apply to your environment to avoid unnecessary alerts and inaccurate optimization results.

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

### About this task

Workload Factory analyzes database configurations daily. The daily analysis provides the well-architected status, and insights and recommendations with options to automatically fix configuration issues so that your configurations meet best practices.

You have options to review the recommendations for configuration issues and fix those issues from the Databases inventory within the Workload Factory console.

### What is analyzed

Workload Factory analyzes the well-architected status of the following configurations:

For Microsoft SQL Server instances:

- Storage sizing: includes storage tier, file system headroom, log drive size, and TempDB drive size
- Storage layout: includes user data files placement, log files placement, and TempDB placement
- Storage configuration: includes capacity management, thin provisioning, tiering policy, snapshots, Microsoft Multipath I/O (MPIO) status, and MPIO timeout setting
- Compute: includes rightsizing, operating system patches, and network adapter settings like Receive Side Scaling (RSS), TCP offloading, and MTU alignment
- Applications: includes Microsoft SQL Server licenses, Microsoft SQL Server patch, and MAXDOP settings
- Resiliency: includes local snapshots, FSx for ONTAP backups, cross-region replication (CRR), and Microsoft SQL High Availability.
- Clones: includes options to refresh and delete clones (sandboxes) that were created in or outside of Workload Factory and are older than 60 days

For Oracle databases:

- Storage sizing: includes swap space allocation and file system headroom
- Storage configuration: includes capacity management, thin provisioning, tiering policy, snapshots, storage efficiencies, and operating system configurations for deployments using NFS or iSCSI with or without Automatic Storage Management (ASM) including Microsoft Multipath I/O (MPIO) status and settings, and the following dNFS settings: dNFS enablement, dNFS consistent IP resolution, dNFS configuration file, and dNFS nosharecache
- Storage layout: includes redo log placement, temp tablespace placement, data files placement, archive log placement, control files placement and binaries placement, ASM disk groups LUN count

## Before you begin

- You must [grant \*operations and remediation\* permissions](#) in your AWS account.
- To assess the storage of a Microsoft SQL Server instance or Oracle database, the resource must be registered in Workload Factory and the storage type must be FSx for ONTAP. [Learn how to register resources](#).
- Make sure you review each recommendation carefully before selecting to fix a setting or configuration. For RSS and MAXDOP settings, we suggest that you test the recommended settings to determine performance improvements before making changes to your production environment.



The remediation process may cause instance downtimes or service interruptions. Make sure you review the recommendation carefully before you choose to fix a configuration.

## Fix a configuration issue

Fix configuration issues for SQL Server or Oracle environments running on FSx for ONTAP storage.



The remediation process may cause instance downtimes or service interruptions. Make sure you review the recommendation carefully before you choose to fix a configuration issue.

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Databases**.
3. In the **Inventory**, select the engine type: Microsoft SQL Server or Oracle.
4. Select resources to fix for specific configurations.
  - If you selected Microsoft SQL Server, select the **Instances** tab.
  - If you selected Oracle, select the **Databases** tab.
5. Select **View and fix** to view configuration issues for the instance.
6. On the Well-architected status page for the instance, review the findings from the analysis.

You can filter configuration issues by categories, subcategories, status, severity, and tags.

You can also download a report of the findings by selecting **Export PDF**.

7. Select the dropdown arrow to view the recommendation for any configuration. Recommendations include best practices, potential pitfalls of unoptimized configurations, and important considerations. Make sure you review the recommendation carefully.
8. Select to **View and fix** configuration issues when the option is available.

**Select all** is the default, but you can select specific resources to fix.

- a. For all configurations except clone cleanup, review the recommendation details to learn what will happen if you choose to fix the issue. Some remediation operations may cause instance downtimes or service interruptions.
- b. For clone cleanup, select cloned databases (sandboxes) to refresh or delete.
  - Refreshing a clone synchronizes it with its source database. Refresh is available only for clones created in Workload Factory.

- Deleting a clone removes it permanently, frees up storage space, and reduces costs. You can delete clones created in and outside of Workload Factory.

9. Select **Continue** to fix the configuration issue.

## Result

Workload Factory begins fixing the issue(s). Select the **Job monitoring** tab to view the status of the operation.

## Postpone or dismiss the analysis of database configurations

Postpone or dismiss the analysis of specific database configurations that don't apply to your database environment to avoid unnecessary alerts and inaccurate optimization results. You can reactivate a postponed or dismissed configuration analysis at any time.

The application requirements for database configurations vary. Workload Factory gives you two options for skipping the analysis of specific database configurations so that you can monitor only relevant issues and get an accurate view of the health of relevant configurations. When a specific configuration analysis is postponed or dismissed, the configuration isn't included in the total optimization score.

You can postpone, dismiss, and reactivate the configuration analysis at the configuration level and at the SQL Server instance or at the Oracle database level.

- **Postpone for 30 days:** Postponing the analysis will stop the analysis for 30 days. After 30 days, the analysis will restart automatically.
- **Dismiss:** Dismissing the analysis postpones the analysis indefinitely. You can restart the analysis if needed.

The following instructions describe how to postpone, dismiss, or reactivate an analysis at the configuration level. To complete the following tasks for specific SQL Server instances or Oracle databases, start in the **Dashboard** tab.

## Postpone

Postpone to stop a configuration analysis for 30 days. After 30 days, the analysis will restart automatically.

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Databases**.
3. From the Databases menu, select **Inventory**.
4. In the Inventory, select the database engine type: **Microsoft SQL Server** or **Oracle**.
5. Navigate to the SQL Server instance or Oracle database with the configuration to postpone, select the action menu, and then select **Well-architected**.
6. On the Well-architected analysis page, scroll down to the configuration to postpone, select the action menu, and then select **Dismiss**.
7. In the Dismiss configuration dialog, select **Postpone for 30 days** and then select **Dismiss**.

### Result

The configuration analysis stops for 30 days.

## Dismiss

Dismiss to stop a configuration analysis indefinitely. You can restart the analysis when needed.

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Databases**.
3. From the Databases menu, select **Inventory**.
4. In the Inventory, select the database engine type: **Microsoft SQL Server** or **Oracle**.
5. Navigate to the SQL Server instance or Oracle database with the configuration to postpone, select the action menu, and then select **Well-architected**.
6. On the Well-architected analysis page, scroll down to the configuration to postpone, select the action menu, and then select **Dismiss**.
7. In the Dismiss configuration dialog, select the **Dismiss** option and then select **Dismiss** to confirm dismissal.

### Result

The configuration analysis stops.

## Reactivate

Reactivate a postponed or dismissed configuration analysis at any time.

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Databases**.
3. From the Databases menu, select **Inventory**.
4. In the Inventory, select the database engine type: **Microsoft SQL Server** or **Oracle**.
5. Navigate to the SQL Server instance or Oracle database with the configuration to postpone, select the

action menu, and then select **Well-architected**.

6. On the Well-architected analysis page, select **Dismissed configuration** to display only dismissed configurations.
7. Select **Reactivate** to restart the configuration analysis for the postponed or dismissed configuration.

#### Result

The configuration analysis is reactivated and occurs daily moving forward.

## Analyze error logs in Workload Factory

Use the smart error log analyzer to automatically interpret Microsoft SQL Server and Oracle database error logs so that you can quickly identify and resolve issues. The Agentic AI-based analysis requires Amazon Bedrock integration.

### About this task

Error log analysis and remediation help maintain the health and performance of SQL Server instances and Oracle databases. Interpreting error logs effectively requires careful analysis and expertise. Manual monitoring, error detection, and root cause analysis are time-intensive and prone to errors. These challenges can delay issue resolution, increased downtime, and operational inefficiencies. The smart error log analyzer addresses these challenges with these key benefits:

- Smart grouping: Intelligently consolidates errors by uniqueness, severity, and category, and simplifies the troubleshooting process for faster, more effective resolutions.
- AI-driven investigation: Leverages AI to proactively analyze errors, providing clear, actionable insights to accelerate issue identification without requiring deep expertise.
- Error enrichment: Enhances error logs with external references, offering contextual clarity to improve understanding and decision-making.
- Best-practice remediation: Delivers tailored, remediation recommendations for SQL Server workloads running on FSx for ONTAP, empowering users of all skill levels to resolve issues confidently.

Whenever you use the error log analyzer, you maintain full control over your environment while benefiting from advanced AI analysis.

To use the error log analyzer, you need to activate Amazon Bedrock, select the model Workload Factory uses, create a private endpoint to connect to Amazon Bedrock, add permissions, and create an enterprise license.

#### [Amazon Bedrock pricing](#)

### Data privacy and security

The feature ensures data privacy and security with the following measures:

#### Data sovereignty

Log data and aggregations stay within your AWS account, communicated via private VPC endpoint (Amazon Bedrock), ensuring no public internet exposure.

#### No AI Training

Customer data is not used to train or improve models. Amazon Bedrock processes logs in real time but does

not train on your data. Results are stored in your environment for reference only. For more details, refer to the [Amazon Bedrock data protection documentation](#).

## Before you begin

To use the error log analyzer, you must meet the following prerequisites:

- You must [grant view, planning, and analysis permissions](#) in your AWS account to create a new database host in Workload Factory.
- [Register resources](#) in Workload Factory.
- The following prerequisites also must be met. You will be prompted to complete these prerequisites as part of the steps to analyze log errors.
  - **Amazon Bedrock activation**

Amazon Bedrock is required so that the AI agent running on the SQL Server or Oracle host from Workload Factory can seamlessly connect with Bedrock and fetch AI-based insights for the identified error logs.

- **Networking**

The Amazon Bedrock VPC endpoint ensures private communication of your SQL Server or Oracle host with Amazon Bedrock APIs and eliminates public internet exposure.

Ensure Amazon Bedrock VPC endpoint is associated with the SQL Server or Oracle host's subnet (example: vpce-050cb2f33a1380ffd).

- **AWS IAM permissions**

The following permissions are required for the EC2 instance profile role associated with the SQL Server or Oracle host and for the AWS credentials associated with Workload Factory.

- EC2 instance profile role with "bedrock:InvokeModel" permission

This permission enables the EC2 instance on the corresponding SQL Server or Oracle host to invoke Bedrock models for proactive error investigation and remediation guidance. This profile also ensures secure AI access for tailored insights.

- AWS credentials associated with Workload Factory: "bedrock:GetFoundationModelAvailability" and "bedrock>ListInferenceProfiles" permissions

These permissions verify model availability and configuration in the region of the SQL Server or Oracle host, and ensure reliable, region-specific performance.

- **Oracle user permission**

The following permission grants read access to extended diagnostic alerts, enabling extraction of error/trace details from logs for AI-powered pattern detection and remediation: V\$DIAG\_ALERT\_EXT. This permission is required only for Oracle databases.

## Analyze error logs

Use the Workload Factory console to analyze SQL Server error logs.

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Databases**.
3. From the Databases menu, select **Inventory**.
4. In the Inventory, select **Microsoft SQL Server** or **Oracle** as the database engine type.
5. From the Instances tab, locate the specific SQL Server instance or Oracle database you want to analyze and then select **Investigate errors** from the menu.
6. From the **Error investigation** tab, complete the following prerequisites as described in the console:
  - Amazon Bedrock
  - Networking: Private endpoint for Amazon Bedrock
  - Permissions for EC2 instance profile role
  - Credentials associated with Workload Database Management (wlmdb)
7. When prerequisites are met, select **Investigate now** to use the error log analyzer to gain insights into your SQL Server error logs.

After the scan, errors are displayed in the console, providing a comprehensive view of the issues detected by the Smart error log analyzer.

8. Use filters to refine the displayed errors based on criteria such as severity, time frame, and error code, or by infrastructure-oriented tags such as compute, storage, network, and security.
9. Review the detailed error information, including original error message, AI-based explanation, and suggested remediation steps to resolve the errors.

## Manage clones

### Check the integrity of the data in a sandbox clone

Run an integrity check to determine if sandbox clone data is intact or corrupt in NetApp Workload Factory for Databases.

#### About this task

When you create a sandbox clone from a source database while it is busy, the clone's data may not be in sync with the most recent snapshot of the source database. This operation checks the integrity of all the objects in the sandbox clone to determine if the sandbox clone data is current.

#### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Storage**.
3. From the Databases menu, select **Sandboxes**.
4. In Sandboxes, select the actions menu of the sandbox clone to check integrity for.
5. Select **Run integrity check**.
6. In the Integrity check dialog, click **Integrity check**.
7. Check the status of the integrity check in Sandboxes or in Job monitoring.

If the integrity check fails, we recommend that you do not use the sandbox clone and create a new sandbox clone.

## Revert a database clone in NetApp Workload Factory for Databases

Revert a database (sandbox) clone to its original version at the time of creation in NetApp Workload Factory for Databases.

### About this task

When you clone a database, the clone at creation is a *baseline* clone. The data in the cloned database is the same as the source database at the time of creation. As data in a database clone changes over time, you might want to revert the data back to the baseline when the clone was first created. This operation is called re-baselining a clone. Re-baselining a clone rather than creating a new clone saves space; however, any changes made to the database clone will be deleted.

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Databases**.
3. From the Databases menu, select **Sandboxes**.
4. In Sandboxes, select the action menu of the database clone you want to revert.
5. Select **Re-baseline**.
6. In the Re-baseline dialog, select **Re-baseline**.

## Refresh a database clone in NetApp Workload Factory for Databases

Refresh a database (sandbox) clone in NetApp Workload Factory for Databases so that it is equivalent to the source database at the current moment or a previous point in time.

### About this task

Refreshing a clone updates the clone either to the source database at the current moment or to a snapshot of the source database taken at a previous point in time. Any changes made to the sandbox clone will be deleted.

### Before you begin

A refresh is only possible when the source database is active.

To refresh a database clone from a snapshot, the source database must have at least one snapshot for the operation.

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Databases**.
3. From the Databases menu, select **Sandboxes**.
4. In Sandboxes, select the action menu of the sandbox clone you want to refresh.
5. Select **Refresh**.
6. In the Refresh dialog, select one of the following options:
  - a. **Refresh to current time**
  - b. **Refresh to point in time**

For this option, select the database snapshot from the dropdown menu to refresh to.

7. Click **Refresh**.

## Connect a sandbox clone to CI/CD tools

Connect a sandbox clone to a continuous integration and continuous delivery (CI/CD) pipeline with REST API code to improve software delivery via automation in NetApp Workload Factory for Databases.

### About this task

To deliver a new version of software automatically to your database clone, you should connect to a CI/CD pipeline. Use the REST API code provided from this operation to make the connection.

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Databases**.
3. From the Databases menu, select **Sandboxes**.
4. In Sandboxes, select the actions menu of the sandbox clone to connect to CI/CD tools.
5. Select **Connect to CI/CD tools**.
6. In the CI/CD dialog, copy or download the REST API code needed to connect to CI/CD tools.
7. Click **Close**.

## View connection information of a database clone

View and copy the connection information of a database clone in NetApp Workload Factory for Databases.

### Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Databases**.
3. From the Databases menu, select **Sandboxes**.
4. In Sandboxes, select the action menu of the sandbox clone to view its connection information.
5. Select **Show connection info**.
6. In the Show connection info dialog, copy the connection information if needed.
7. Select **Close**.

## Split a database clone from the source database

In NetApp Workload Factory for Databases, splitting a database clone from its source database creates a new database which will consume a certain amount of storage capacity. The clone gets deleted when the split is complete and the new database appears in the Inventory.

### Before you begin

Consider how much storage capacity is needed for the new database. If needed, [increase file system capacity](#) for the FSx for ONTAP file system before you begin.

## Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Databases**.
3. From the Databases menu, select **Sandboxes**.
4. In Sandboxes, select the action menu of the database clone you want to split.
5. Select **Split**.
6. In the Split dialog, select **Split**.

## Delete a database clone in NetApp Workload Factory for Databases

Delete a sandbox clone in NetApp Workload Factory for Databases when you no longer need it and want to free up storage capacity.

## Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Databases**.
3. From the Databases menu, select **Sandboxes**.
4. In Sandboxes, select the action menu of the sandbox clone you want to delete.
5. Select **Delete**.
6. In the Delete dialog, select **Delete** to confirm deletion.

## Deregister a resource in NetApp Workload Factory

Deregister a resource, such as a Microsoft SQL Server instance or an Oracle database, if you no longer want to deploy or monitor resources from the NetApp Workload Factory console. Deregistering resources also removes the amount of FSx for ONTAP storage capacity used by the instance.

You can register the resource again.

## Steps

1. Log in using one of the [console experiences](#).
2. Select the menu and then select **Databases**.
3. From the Databases menu, select **Inventory**.
4. In the Inventory, select the engine type: **Microsoft SQL Server**, **Oracle**, or **PostgreSQL**.
5. Locate the resource you want to deregister, select the action menu, and then select **Deregister**.

# Knowledge and support

## Register for support

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 960xxxxxxxxx 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.



9601111222224444555555

Account Serial Number



Registered for Support

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.

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 for NetApp Workload Factory for Databases

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 Environment
<input style="width: 150px; height: 30px; border: 1px solid #ccc; border-radius: 5px; padding: 5px; margin-bottom: 5px;" type="text" value="Select"/>	<input style="width: 150px; height: 30px; border: 1px solid #ccc; border-radius: 5px; padding: 5px; margin-bottom: 5px;" type="text" value="Select"/>
<b>Case Priority</b> 	
<input style="width: 300px; height: 30px; border: 1px solid #ccc; border-radius: 5px; padding: 5px; margin-bottom: 5px;" type="text" value="Low - General guidance"/>	
<b>Issue Description</b>	
<p>Provide detailed description of problem, applicable error messages and troubleshooting steps taken.</p>	
<b>Additional Email Addresses (Optional)</b> 	
<input style="width: 300px; height: 30px; border: 1px solid #ccc; border-radius: 5px; padding: 5px; margin-bottom: 5px;" type="text" value="Type here"/>	
<b>Attachment (Optional)</b> 	
<input style="width: 300px; height: 30px; border: 1px solid #ccc; border-radius: 5px; padding: 5px; margin-bottom: 5px;" type="text" value="No files selected"/>	
 <b>Upload</b> 	

## 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.

The screenshot shows a list of cases with the following filters applied:

- Date created: December 22, 2022
- Last updated: December 29, 2022
- Status: Unassigned
- Priority: Critical (P1)
- Priority: High (P2)
- Priority: Medium (P3)
- Priority: Low (P4)

The list includes the following cases:

Last updated	Priority	Status	Notes
December 29, 2022	Critical (P1)	Active	
December 28, 2022	High (P2)	Pending customer	
December 27, 2022	Medium (P3)	Solution proposed	
December 26, 2022	Low (P4)	Pending customer	

- Filter the contents of the columns.

The screenshot shows a list of cases with the following filters applied:

- Last updated: December 29, 2022
- Priority: Critical (P1)
- Status: Active, Pending customer, Solution proposed, Pending closed

The list includes the following cases:

Last updated	Priority	Status	Notes
December 29, 2022	Critical (P1)	Active	
December 28, 2022	High (P2)	Pending customer	
December 27, 2022	Medium (P3)	Solution proposed	
December 26, 2022	Low (P4)	Pending closed	

- Change the columns that appear in the table by selecting and then choosing the columns that you'd like to display.

The screenshot shows a list of cases with the following filters applied:

- Last updated: December 29, 2022
- Priority: Critical (P1)
- Status: Active, Pending customer, Solution proposed, Pending closed

The list includes the following cases:

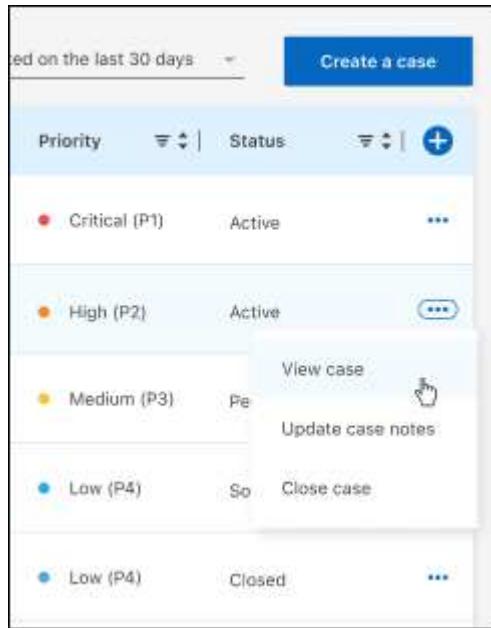
Last updated	Priority	Status	Notes
December 29, 2022	Critical (P1)	Active	Last updated, Priority
December 28, 2022	High (P2)	Pending customer	Cluster name
December 27, 2022	Medium (P3)	Solution proposed	Case owner, Opened by
December 26, 2022	Low (P4)	Pending closed	

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**.



The screenshot shows a list of cases in a management interface. The cases are ordered by priority: Critical (P1), High (P2), Medium (P3), and Low (P4). Each case entry includes its priority color, a status (Active or Pending), and a 'More options' button (indicated by '...'). The 'View case' option for the Medium (P3) case is highlighted with a cursor icon, indicating it is the selected item.

Priority	Status	More options
Critical (P1)	Active	...
High (P2)	Active	...
Medium (P3)	Pending	<b>View case</b> 
Medium (P3)	Pending	Update case notes
Low (P4)	So	Close case
Low (P4)	Closed	...

# Legal notices

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

## Copyright

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

## Trademarks

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

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

## Patents

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

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

## Privacy policy

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

## Open source

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

[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.