



Create a new database server

Database workloads

NetApp
January 05, 2026

Table of Contents

Create a new database server	1
Create a Microsoft SQL Server in Workload Factory for Databases	1
About this task	1
Before you begin	2
Step 1: Create a database server	2
Step 2: Enable remote connection on the Microsoft SQL Server	9
Create a PostgreSQL server in NetApp Workload Factory	10
About this task	10
Before you begin	10
Create a PostgreSQL server	10

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:

- 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, 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 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.

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.