NetApp

Administer and monitor

Database workloads

NetApp
February 04, 2026

Table of Contents

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.

> (i) 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.

> (i) 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

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 datahase 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**.