



Implement configuration best practices

Workload Factory for NetApp Console

NetApp

October 07, 2026

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

Table of Contents

- Implement configuration best practices 1
 - Learn about configuration analysis in NetApp Workload Factory for Databases 1
 - Options for analyzing your database environment. 1
 - Microsoft SQL Server best practices 1
 - Oracle best practices 6
- Implement well-architected database configurations in NetApp Workload Factory 12
 - About this task 12
 - Perform a one-time assessment of your configurations. 13
 - Perform an on-demand assessment 14
 - Before you begin fixing, dismissing, and re-activating configurations 14
 - Fix a configuration issue 15
 - Dismiss the analysis of database configurations 16
 - Reactivate a dismissed configuration analysis 16
 - Related information 17

Implement configuration best practices

Learn about configuration analysis in NetApp Workload Factory for Databases

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 so that your database configurations achieve peak performance, cost efficiency, and compliance with best practices.

Workload Factory provides best practices and recommendations for well-architected database workloads. The well-architected analysis assesses Microsoft SQL Server and Oracle Database configurations and settings related to storage, compute, application, resiliency, and cloning.

Unresolved directive in fsx-databases/optimize-overview.adoc - include::_include/well-architected-workloads.adoc[]

Options for analyzing your database environment

Workload Factory provides the following options for analyzing your database environment:

- **One-time assessment:** Perform a [one-time assessment](#) to find out the well-architected status of your Microsoft SQL Server instances or Oracle databases without storing credentials or registering resources.

For the one-time assessment, you only need to provide credentials once for the Microsoft SQL Server instances or Oracle databases you want to analyze.

- **Continuous assessment:** Workload Factory analyzes your database environment regularly, provides ongoing recommendations for optimizing your configurations, and automatically fixes issues.

For continuous assessment, you need to register your resources, associate a link, and grant permissions in your AWS account.

Microsoft SQL Server best practices

The following best practices and recommendations apply to Microsoft SQL Server workloads.

Storage

Storage tier

Use the primary SSD tier for performance-sensitive SQL workloads.

▼ Details

Use the primary SSD tier for FSx for ONTAP volumes for optimal performance. The capacity pool tier might reduce performance and increase latency. The primary SSD tier is designed for high-performance workloads, while the capacity pool tier is optimized for cost efficiency and might not meet the performance requirements of SQL Server workloads.

File system headroom

Keep enough free capacity above total volume size.

▼ Details

Maintain enough file system capacity above your total volume size. Too little free space causes performance issues; too much wastes storage and increases costs.

Log drive size

Size and monitor log drives to avoid rollbacks and outages.

▼ Details

Size and monitor your SQL Server log drive to prevent transaction rollbacks, database unavailability, data corruption, and performance issues caused by a full log drive.

TempDB drive size

Size TempDB to avoid bottlenecks and timeouts.

▼ Details

Size and monitor TempDB to maintain performance and stability. Too little TempDB space can cause query slowdowns, application timeouts, and system crashes.

Data and log files placement

Separate .mdf and .ldf paths.

▼ Details

Separate data and log files onto different drives to improve parallel I/O, backup flexibility, and restore behavior. We recommend separating data and log LUN paths into different volumes for smaller databases.

TempDB placement

Use a dedicated drive.

▼ Details

Place TempDB on a dedicated drive to avoid I/O contention and improve SQL Server performance and stability.

Thin provisioning

Allocate storage as needed.

▼ Details

Configure thin provisioning for FSx for ONTAP volumes to make better use of storage and reduce costs. Thin provisioning lets you allocate storage as needed instead of reserving all capacity at once.

Autosize and auto-size mode

Grow volumes automatically.

▼ Details

Turn on autosize and set the mode to `grow` so volumes increase automatically before they become full and read-only.

Block device space management

Configure block device space settings for LUNs used by SQL Server.

▼ Details

Configure block device space settings for LUNs used by Microsoft SQL server instances to prevent write failures and improve space efficiency on FSx for ONTAP.

Snapshot copy reserve and autodelete

Keep space for active data.

▼ Details

Set snapshot copy reserve to 0% and enable snapshot autodelete to protect space for active data.

Space management

Expand the volume before deleting snapshot copies.

▼ Details

When space runs low, expand the volume before deleting snapshot copies. Expanding the volume prevents "volume full" errors, reduces write failures, and preserves snapshot recovery points.

Tiering / TCO optimization

Tier snapshot data using the snapshot-only tiering policy.

▼ Details

Use the snapshot-only tiering policy, which moves only snapshot data to the capacity tier while keeping active data on the SSD tier. This approach preserves low-latency performance for SQL workloads while reducing storage costs. Tier snapshot data after a 7-day cooling period.

Storage efficiencies

Use deduplication, compression, and compaction.

▼ Details

Enable deduplication, compression, and compaction to reduce storage use and cost while maintaining performance.

OS type

Match ONTAP LUN OS type to host partitioning.

▼ Details

Match the ONTAP LUN Operating System value to the host partitioning scheme to maintain correct I/O alignment.

MPIO policy, sessions, status, and timeout settings

Configure Multipath I/O (MPIO) policy, sessions, status, and timeout settings.

▼ Details

Configure Multipath I/O (MPIO) for resiliency and throughput. Use five MPIO sessions per target interface for SQL Server on iSCSI LUNs, monitor path status, and set host timeout to 60 seconds.

NTFS allocation unit size

Set to 64K.

▼ Details

Set NTFS allocation unit size to 64K to improve disk efficiency and reduce fragmentation for SQL Server workloads.

Scheduled local snapshots

Disable scheduled snapshots.

▼ Details

Disable scheduled snapshots for FSx for ONTAP volumes used by Microsoft SQL Server. Instead, manage snapshots externally using tools such as NetApp SnapCenter, which create application-consistent backups and help prevent data corruption during restore operations.

Compute

Compute rightsizing

Match EC2 instance size to workload demand.

▼ Details

Rightsize your SQL Server EC2 instance based on workload demand. Increase size when under-provisioned and reduce size when over-provisioned to balance cost and performance.

Operating system patch

Apply current OS patches.

▼ Details

Apply the latest operating system patches to improve security and reliability.

Network adapter settings

Configure Receive Side Scaling (RSS).

▼ Details

Configure Receive Side Scaling (RSS) to distribute network processing across CPUs to improve network performance. Validate settings before production rollout.

MTU alignment

Align Maximum Transmission Unit (MTU) across host and storage paths.

▼ Details

Align EC2 Maximum Transmission Unit (MTU) settings with FSx for ONTAP paths to prevent fragmentation and maintain stable throughput.

Application

License

Review license usage for paid features.

▼ Details

Review SQL Server license usage at the host level. If paid features are not in use, the license is not optimized and may increase cost.

Microsoft SQL Server patch

Apply latest SQL Server patches.

▼ Details

Apply the latest SQL Server patches to improve security and system reliability.

MAXDOP

Use workload-tested values.

▼ Details

Configure MAXDOP to optimize query performance. Values of 4, 8, or 16 typically work well. Test your workload and monitor for parallelism wait types such as CXPACKET.

Resiliency

Application-consistent snapshots

Use application-consistent snapshots with NetApp SnapCenter.

▼ Details

Use application-consistent snapshots with NetApp SnapCenter to take accurate, reliable snapshots of your volume data at a specific moment in time. SnapCenter makes backups easier and helps you restore data quickly, reducing downtime and protecting critical workloads.

Cross-region replication

Protect against regional outages.

▼ Details

Enable cross-region replication (CRR) to improve data availability and disaster recovery in your SQL Server environment. CRR helps protect against regional outages and ensures business continuity by replicating data across different locations.

Backup configuration

Configure FSx for ONTAP backups or AWS Backup.

▼ Details

Configure FSx for ONTAP backups or AWS Backup to meet retention and compliance requirements while avoiding duplicate backup workflows.

Cluster quorum

Configure quorum for failover cluster instances.

▼ Details

Configure quorum correctly for SQL Server Failover Cluster Instances so the cluster can tolerate expected node failures.

Drive letter

Keep drive letters consistent across nodes.

▼ Details

Use the same drive letters on every node in your SQL Server high availability setup. This makes management easier and helps failover work correctly.

Heartbeat settings

Tune heartbeat for stable failover behavior.

▼ Details

Tune heartbeat settings to avoid unnecessary failovers while still detecting real failures quickly.

Shared storage

Make sure both nodes can access shared iSCSI LUNs.

▼ Details

All shared disks (iSCSI LUNs) must be accessible from both nodes in the FCI deployment model so failover can work.

SQL Server service

Validate service account and startup dependencies.

▼ Details

Validate service account permissions, startup type, and dependencies for reliable high availability behavior.

Cloning

Clone cleanup

Remove old clones regularly.

▼ Details

Regularly remove outdated clones to reduce storage cost and operational clutter. Old and unused clones cause high costs.

Oracle best practices

The following best practices and recommendations apply to Oracle workloads.

Storage

File system headroom

Keep enough free capacity above total volume size.

▼ Details

Maintain enough file system capacity above your total volume size. Too little free space causes performance issues; too much wastes storage and increases costs.

Thin provisioning

Allocate storage as needed.

▼ Details

Configure thin provisioning for FSx for ONTAP volumes to make better use of storage and reduce costs. Thin provisioning lets you allocate storage as needed instead of reserving all capacity at once.

Autosize and auto-size mode

Grow volumes automatically.

▼ Details

Turn on autosize and set mode to `grow` so volumes increase automatically before they become full and read-only.

Block device space management

Configure block device space settings for LUNs used by SQL Server.

▼ Details

Configure block device space settings for LUNs used by Oracle databases to prevent write failures and improve space efficiency on FSx for ONTAP.

Snapshot copy reserve and autodelete

Keep space for active data.

▼ Details

Set snapshot copy reserve to 0% and enable snapshot autodelete to protect space for active data.

Space management

Expand the volume before deleting snapshot copies.

▼ Details

When space runs low, expand the volume before deleting snapshot copies. Expanding the volume prevents "volume full" errors, reduces write failures, and preserves snapshot recovery points.

Tiering / TCO optimization

Tier snapshot data using the snapshot-only tiering policy.

▼ Details

Use the snapshot-only tiering policy, which moves only snapshot data to the capacity tier while keeping active data on the SSD tier. This approach preserves low-latency performance for SQL workloads while reducing storage costs. Tier snapshot data after a 7-day cooling period.

Storage efficiencies

Use deduplication, compression, and compaction.

▼ Details

Enable deduplication, compression, and compaction to reduce storage use and cost while maintaining performance.

Swap space

Set swap based on RAM.

▼ Details

Set swap space based on the amount of RAM so the system can handle memory pressure and avoid slowdowns or crashes.

Archive, data, control, redo, and temp placement

Use a dedicated volume for each type of Oracle file to isolate I/O patterns and improve performance.

▼ Details

Place archive logs, data files, control files, redo logs, and temp files on separate drives to isolate I/O patterns and improve resiliency. Keep redundant copies on separate volumes where required.

Oracle binary placement

Use a dedicated volume.

▼ Details

Place Oracle binaries on a dedicated volume to reduce I/O contention. This separation simplifies software

updates and minimizes the risk of accidental modifications or corruption.

Scheduled local snapshots

Disable local scheduled snapshots.

▼ Details

Disable scheduled snapshots for FSx for ONTAP volumes used by Oracle to save space and lower costs. Instead, manage snapshots externally using tools such as NetApp SnapCenter, which create application-consistent backups and help prevent data corruption during restore operations.

NFS rootonly

Disable for dNFS.

▼ Details

Disable the ONTAP `nfs-rootonly` parameter for dNFS. `nfs-rootonly` restricts NFS connections to privileged ports (<1024). Because dNFS processes in NFSv4+ do not run as root and use higher ports, disabling this parameter allows necessary connections.

Binaries export policy

Include superuser and setuid permissions in the dedicated export policy for Oracle binaries.

▼ Details

Make sure that if Oracle binaries are located on an NFS share, the export policy includes superuser and setuid permissions. Superuser (root) access allows NFS clients to map as root, needed for binary execution.

ASM setup

Set up Oracle Automatic Storage Management (ASM) for iSCSI-based storage running on FSx for ONTAP.

▼ Details

Use Oracle Automatic Storage Management (ASM) for iSCSI-based storage running on FSx for ONTAP to optimize performance, simplify storage management, and enhance scalability for Oracle Database deployments.

ASM-only configurations

- **ASM data disk group, ASM logs disk group, ASM FRA disk group, and ASM archive disk group LUNs**

Distribute LUNs across Automatic Storage Management (ASM) disk group resources.

▼ Details

Distribute LUNs across Automatic Storage Management (ASM) data, redo log, and archive log disk groups for optimal performance and redundancy.

- **ASM external redundancy**

Configure Oracle ASM disk groups with EXTERNAL redundancy.

▼ Details

Configure Oracle ASM disk groups with EXTERNAL redundancy for FSx for ONTAP iSCSI LUNs to leverage FSx for ONTAP's built-in high availability, optimize storage efficiency, and reduce costs by avoiding Oracle-level data mirroring.

- **ASM filter driver and ASMLib logical block size alignment**

Configure Oracle ASM Filter Driver (AFD) and ASMLib to use the logical block size of the underlying FSx for ONTAP storage.

▼ **Details**

Configure the Oracle ASM Filter Driver (AFD) and ASMLib in the ASMLib configuration file to use the logical block size of the underlying FSx for ONTAP storage. This ensures that AFD and ASMLib aligns I/O operations with the storage's block size, optimizing performance by minimizing latency and reducing unnecessary I/O overhead.

TCP slot table

Set to 128 for NFS-based Oracle workloads.

▼ **Details**

Configure the kernel parameters for the TCP slot table to 128 for Oracle workloads running on NFS. This configuration enables more concurrent I/O operations, reducing latency in high-throughput scenarios such as Oracle data file access. Kernel parameters govern the maximum number of outstanding NFS requests over TCP, acting like queue depth in Fibre Channel setups. The default Linux values, typically set at 16, do not support optimal database performance.

NFS mount options - database files

Use NFS mount options that optimize performance for database files.

▼ **Details**

Use optimized NFS mount options for database files. Key parameters include `rw` for read/write access, `bg` for background mounting, `hard` for indefinite retries on failures, and `proto=tcp` for reliable delivery. The `rsize` and `wsizes` values of 262144 improve large I/O operations, enhancing throughput, while `nointr` doesn't interrupt long-running tasks. This configuration improves database performance and resilience, particularly in high-throughput environments.

NFS mount options for ADR home

Use NFS mount options that optimize performance for the ADR home directory.

▼ **Details**

Use optimized NFS mount options for the Automatic Diagnostic Repository (ADR) home, or root directory. The ADR home stores logs and traces, which benefit from similar reliability options as core database files but don't require the `nointr` mount option because operations are less interrupt-sensitive.

NFSv4 domain name

Match NFSv4 domain names between the host and NFS server.

▼ **Details**

Match NFSv4 domain names between the host (`/etc/idmapd.conf` or `hostname -d`) and NFS server (`v4-id-domain` in ONTAP).

NFS caching options

Do not disable host caching in standalone deployments.

▼ **Details**

Omit NFS mount options that disable host caching—such as, `actimeo=0` and `noac`—in standalone Oracle deployments to improve efficiency and reliability. These options apply to clustered environments to ensure

cache consistency, but in single-instance environments, they unnecessarily bypass caching and can significantly reduce performance.

dNFS enablement

Enable dNFS.

▼ Details

Enable Direct NFS (dNFS) in your Oracle environment to improve NFS storage performance and reliability.

dNFS consistent IP resolution

Avoid round-robin name resolution.

▼ Details

Avoid the use of Direct NFS (dNFS) with any type of round-robin name resolution, including DNS, DDNS, NIS, or any other method. This includes the DNS load balancing feature available in ONTAP. Consistent IP address resolution maintains database stability and prevents potential crashes or data corruption.

dNFS configuration file

Validate `oranfstab` content.

▼ Details

Check and update the `oranfstab` file content to ensure proper Direct NFS (dNFS) usage. The `oranfstab` file is essential for configuring advanced dNFS features such as multipathing and specific NFS options. Correct configuration improves the efficiency of data access and management.

dNFS no shared cache

Use `nosharecache` for specific multi-mount cases.

▼ Details

Configure the `nosharecache` mount option for environments where Direct NFS (dNFS) is enabled, and a source volume is mounted more than once on a single server with nested NFS mounts. This configuration prevents cache sharing between mounts, ensuring data consistency and optimal performance.

Compute

Operating system patch

Apply current patches.

▼ Details

Apply the latest operating system patches to improve security and reliability.

Transparent huge pages

Disable Transparent Huge Pages (THP).

▼ Details

Disable Transparent Huge Pages (THP) on database hosts running Oracle databases to prevent potential performance degradation and improve memory management.

TCP advanced options

Enable Transmission Control Protocol (TCP) advanced options.

▼ Details

Enable TCP timestamps, SACK, and window scaling for best network performance and reliability.

File system I/O options

`Set filesystemio_options = setall.`

▼ Details

Set `filesystemio_options = setall` for optimal I/O performance. Adjust the system global area (SGA) size if needed when moving away from buffered I/O.

Multiblock read count

Remove the `db_file_multiblock_read_count` parameter.

▼ Details

Remove `db_file_multiblock_read_count` parameter from the text-based FILE `init.ora` or server-managed binary SPFILE (`spfile<SID>.ora`) for Oracle databases to prevent performance issues and allow Oracle to manage this setting automatically.

Oracle Critical Patch Updates

Apply current CPU updates.

▼ Details

Apply the latest Oracle Critical Patch Updates to address security vulnerabilities and maintain the integrity of your database environment. Regularly review and apply patches to protect against potential threats and ensure compliance with security best practices.

Resiliency

Application-consistent snapshots

Use application-consistent snapshots with NetApp SnapCenter.

▼ Details

Use application-consistent snapshots with NetApp SnapCenter to take accurate, reliable snapshots of your volume data at a specific moment in time. SnapCenter makes backups easier and helps you restore data quickly, reducing downtime and protecting critical workloads.

Cross-region replication

Protect against regional outages.

▼ Details

Enable cross-region replication (CRR) to improve data availability and disaster recovery for your Oracle databases. CRR helps protect against regional outages and ensures business continuity by replicating data across different locations.

Backup configuration

Configure FSx for ONTAP backups or AWS Backup.

▼ Details

Configure FSx for ONTAP backups or AWS Backup to meet retention and compliance requirements while avoiding duplicate backup workflows.

Cloning

Clone cleanup

Remove old clones regularly.

▼ Details

Regularly remove outdated clones to reduce storage cost and operational clutter. Old and unused clones cause high costs.

What's next

[Implement well-architected database configurations](#)

Implement well-architected database configurations in NetApp 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 configurations that don't apply to your environment to avoid extra alerts and inaccurate 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 well-architected status, insights, and recommendations.

Many recommendations are available when you grant *view*, *planning*, and *analysis* permissions to your AWS account. If you grant *operations* and *remediation* permissions, all recommendations are available, and you can automatically fix configuration issues to meet best practices.

You can review and fix issues from the Databases inventory or the Well-architected dashboard within the Workload Factory console. Because requirements for database configurations vary, you can dismiss the analysis of specific configurations that don't apply to your database environment. This helps you avoid unnecessary alerts and inaccurate optimization results.

Optionally, you can perform a one-time assessment to review best practices for your database configurations without storing credentials or registering resources.

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

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
- Application: includes Microsoft SQL Server licenses, Microsoft SQL Server patch, and MAXDOP settings
- Resiliency: includes local snapshots, backup configuration, cross-region replication (CRR), Microsoft SQL High Availability, and clone cleanup
- 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
- Compute: includes operating system patches
- Application: includes Oracle Critical Patch Updates
- Resiliency: includes cross-region replication (CRR), application-consistent snapshots, backup configuration, and clone cleanup

The well-architected analysis for Oracle notifies you when a configuration depends on others and lists actions to optimize linked configurations.

Perform a one-time assessment of your configurations

Perform a one-time assessment to find out the well-architected status of your Microsoft SQL Server instances or Oracle databases without storing credentials or registering resources.

You can run the script and upload results multiple times. Workload Factory updates the well-architected status and insights with each upload.

For each configuration, you can view the current value and the recommended value.

Microsoft SQL Server instances

Steps

1. Log in using one of the [console experiences](#).
2. Select the menu  Then select **Databases**.
3. In the **Inventory**, select **SQL Server** as the engine type.

4. Select **One-time assessment** > **Learn about assessment**. Review the information about the one-time assessment.
5. Select **Download script**.
 - a. Copy the script to your Microsoft SQL Server host.
 - b. Run the script in PowerShell. One script collects metadata from all instances.
6. Select **One-time assessment** > **Upload**. Find the file and upload it to the Workload Factory console.
7. View the recommendation for any configuration. Workload Factory includes best practices and potential problems in the recommendations if configurations are not optimized.

Oracle databases

Steps

1. Log in using one of the [console experiences](#).
2. Select the menu  Then select **Databases**.
3. In the **Inventory**, select **Oracle** as the engine type.
4. Select **One-time assessment** > **Learn about assessment**. Review the information about the one-time assessment.
5. Select **Download script**.
 - a. Copy the Python script to your Oracle database host.
 - b. Run the script in the appropriate shell. One script collects metadata from all databases.
6. Select **One-time assessment** > **Upload**. Find the file and upload it to the Workload Factory console.
7. View the recommendation for any configuration. Workload Factory includes best practices and potential problems in the recommendations if configurations are not optimized.

Perform an on-demand assessment

If you grant *operations and remediation* permissions in your AWS account, you can perform an on-demand assessment of your Oracle databases. You don't have to register resources to run this assessment.

Steps

1. Log in using one of the [console experiences](#).
2. Select the menu  Then select **Databases**.
3. In the **Inventory**, select **Oracle** as the engine type.
4. Select **View and fix**.
5. In the Well-architected status tab, select **Analyze now**.
6. View the recommendation for any configuration. Workload Factory includes best practices and potential problems in the recommendations if configurations are not optimized.

Before you begin fixing, dismissing, and re-activating configurations

The following requirements are necessary for continuous assessment of your database environment and automatic fixing of configuration issues. These requirements are not required for the one-time assessment.

- You must [grant operations and remediation permissions](#) in your AWS account.
- You must register the resource in Workload Factory and the storage type must be FSx for ONTAP. [Learn how to register resources](#).
- Review each recommendation before fixing a setting or configuration. For RSS and MAXDOP, test recommended settings to check performance before changing your production environment.

Fix a configuration issue

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



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

Steps

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

Select **Export PDF** to download a report of the findings.

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 see the impact of fixing the issue. Some remediation operations may cause downtime 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.

Dismiss the analysis of database configurations

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 dismissed configuration analysis at any time.

Steps

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

Workload Factory stops the configuration analysis.

You can reactivate the analysis at any time. The configuration is no longer included in the total optimization score.

Reactivate a dismissed configuration analysis

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

Steps

1. Log in to the [NetApp Console](#).
2. Select the menu , **Workloads > 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

Workload Factory reactivates the configuration analysis and runs it every day from now on.

Related information

- [Learn about the configuration analysis for database resources and well-architected status in Workload Factory](#)
- [Connect to an FSx for ONTAP file system using a link](#)
- [Learn about custom rules](#)

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.