AMD EPYC™ 9005 MICROSOFT® SQL SERVER TUNING GUIDE

PID: 58485 V2.0 | AUGUST 2025

DATE	VERSION	CHANGES
JUNE, 2024	0.1	Initial NDA release
OCTOBER, 2024	1.0	Initial public release
AUGUST, 2025	2.0	Incorporate Feedback from Microsoft

AUDIENCE

This document is intended for a technical audience such as MSSQL® application architects, production deployment, and performance engineering teams who have:

- A background in configuring servers.
- Administrator-level access to both the server management Interface (BMC) and the OS.
- Familiarity with the BMC, OS-specific configuration, monitoring, and troubleshooting tools.

AUTHOR

Bryon Georgson and Sylvester Rajasekaran

MICROSOFT® SQL SERVER TUNING GUIDE CONTENTS

Chapter 1 Introduction 4
1.1 Important Reading 4

Chapter 2 Hardware Configuration 5
2.1 CPU Selection Guidelines 5
2.2 Recommended CPUs 5
2.3 Memory 6
2.4 Network 6
2.5 Storage 7

Chapter 3 BIOS Settings 8

Chapter 4 OS Settings 9
4.1 Windows Configuration 9
4.2 Linux Configuration 9
4.2.1 /etc/sysctl.conf 9
4.2.2 /var/opt/mssql/mssql.conf 9
4.2.3 /etc/rc.local 10
4.2.4 tuned-adm profile 10
4.2.5 Filesystem Mount Options 10

Chapter 5 SQL Server Configuration Settings 11
5.1 SQL Server Soft NUMA Configuration 11
5.2 SQL Server 2022 Instance Trace Flags and Settings 13
5.3 sp_configure 14

Chapter 6 Workload-Specific Considerations 15
6.1 Data Warehouse 15
6.2 Online Transaction Processing 15

Chapter 7 Resources 16

CHAPTER 1

INTRODUCTION

This tuning guide provides detailed descriptions of parameters that can optimize performance on servers built with AMD EPYC™ 9005 Series processors. The default configurations on hardware and BIOS from different OEM vendors may require additional tuning for best performance on specific OS platforms and/or workloads. Thia guide enables additional optimizations on a per platform and workload level by calling out:

- BIOS settings that can impact performance
- Hardware configuration best practices
- Supported OS versions and configurations
- SQL Server performance tuning
- Workload-specific BIOS and OS settings for OLTP and DSS workloads.

1.1 IMPORTANT READING

Please be sure to read the following guides (available from the **AMD Documentation Hub**), which contain important foundational information about 5th Gen AMD EPYC processors:

- AMD EPYC™ 9005 Processor Architecture Overview
- BIOS & Workload Tuning Guide for AMD EPYC™ 9005 Series Processors
- Memory Population Guidelines for AMD EPYC™ 9005 Series Processors

CHAPTER 2

HARDWARE CONFIGURATION

This chapter discusses the hardware components to consider when planning a SQL Server system.

2.1 CPU SELECTION GUIDELINES

Selecting the right processor depends on the following factors:

- **Licensing costs:** Microsoft SQL Server is licensed by the core. Processors with lower core counts and higher frequencies are therefore ideal.
- **Database size, IOPS, and throughput:** Single-socket systems are better for memory latency dependent applications. Dual-socket systems with 8 or more memory channels per socket are better for applications that depend on memory bandwidth.
- **Inter-process data sharing:** It may be better to use AMD EPYC 9005 Series Processors with larger caches when the SQL Server threads either have little data sharing between them or those that do share data can be isolated to a single Core Complex.

2.2 RECOMMENDED CPUS

Selecting the correct CPU is an important part of achieving optimal database application performance. Table 2-1 provides CPU selection guidelines along with the minimum memory requirements for SQL Server running Decision Support Systems (DSS) and Online Transaction Processing (OLTP) applications. You can choose other AMD EPYC processors not listed in Table 2-1 if needed for your specific requirements.

WORKLOAD					
SIZE	**TYPE**	**SIZE/# OF USERS**	**CORES**	**PROCESSOR**	**MEMORY**
SMALL	DSS	• Up to 300GB • Up to 5 concurrent users	16	AMD EPYC 9005 16-core	128GB
MEDIUM	DSS	• 300GB-1TB • Up to 10 concurrent users	32	AMD EPYC 9005 32-core	256GB
LARGE	DSS	• 1TB to 3TB • Up to 20 concurrent users	96	AMD EPYC 9005 96-core	512GB
VERY LARGE	DSS	• >3TB • >20 concurrent users	192	2 x AMD EPYC 9005 96-core	1TB-4TB
SMALL	OLTP	• Up to 100GB • Up to 50 concurrent users	16	AMD EPYC 9005 16-core	256GB
MEDIUM	OLTP	• 100GB-300GB • Up to 100 concurrent users	32	AMD EPYC 9005 32-core	512GB
LARGE	OLTP	• 300GB-1TB • Up to 200 concurrent users	64	AMD EPYC 9005 64-core	1TB-2TB
VERY LARGE	OLTP	• >1TB • >200 concurrent users	96	AMD EPYC 9005 96-core	1TB-4TB
CLOUDS	Mixed	• Private clouds • PaaS • SaaS	192	2 x AMD EPYC 9005 96-core	Size based on unique requirements

Table 2-1: Recommended AMD EPYC CPUs for various SQL Server applications

For cloud deployments, use 8 to 96 vCPU cores with at least 8 GB memory per vCPU instance VM with appropriate storage attached for optimum performance.

2.3 MEMORY

SQL Server systems consume a lot of memory. AMD EPYC 9005 Series Processors include 12 memory channels per CPU socket, which increases both the total addressable memory and memory throughput per socket. Using large pages in the OS system can improve system performance by reducing the amount of system resources to access page table entries. The size of large pages varies from platform to platform. Most Linux distributions enable transparent hugepages by default, and SQL Server will use them if enabled.

2.4 NETWORK

Tuning and configuring an RDBMS network critical for rapidly transferring information in and out of the database. Tuning also ensures that the network does not overtake all system resources. This is less critical on smaller systems that run transactions and queries directly on the local server this will be a less critical component in the tuning process. Review the Windows® Network Tuning Guide for AMD EPYC™ 9005 Series Processors or Linux® Network Tuning Guide for AMD EPYC™ 9005 Series Processors (available from the **AMD Documentation Hub** for OS-specific network tuning information.

Configure the number of RSS queues and set the processor affinity of each network port to its own set of CPUs that are within a Core/Cache Die (CCD). Enable jumbo frames in and OLTP environment for better network throughput. Adjust flow control settings to be identical on the switch and server ports.

On Linux systems, enable Rx Tx buffer size as follows:

```
#To check the pre-set maximums, run the command, example NIC name used here is:
"eth0"
ethtool -g eth0
#command to set both the rx(recieve) and tx (transmit) buffer size to 4 KB.
ethtool -G eth0 rx 4096 tx 4096
#command to check the value is properly configured is: ethtool -g eth0
```

To reduce number of RSS queues to #cores in a CCD:

```
ethtool -l eth0
ethtool -L eth0 <queue type> 16 ethtool -l eth0
```

On Windows systems, use PowerShell cmdlets to set Rx Tx buffer size:

```
# list all the NET adapters Get-NetAdapter
# Using $rssNumaNode socket
Set-NetAdapterRss -Name <network adapter> -Enable 1 -BaseProcessorGroup
<processor group that nic is attached to> -MaxProcessorGroup <processor group
that nic is attached to> -NumberOfReceiveQueues 16 -BaseProcessorNumber 0
-MaxProcessorNumber 15
-MaxProcessors 16 -NumaNode -Profile NUMAStatic
```

2.5 STORAGE

AMD EPYC 9005 Series Processors support PCIe® 5.0 with double the I/O bandwidth of PCIe 4.0. All normal RDBMS I/O layer tuning rules still apply, such as:

- Using RAID devices when possible.
- Separating log and data devices.
- Splitting I/O across the remaining devices, controllers, PCIe bus, etc.

Either mirror database logs or use RAID 10. You may also want to RAID the data, depending on your fault tolerance requirements. Hardware RAID is always more efficient than software RAID.

Storage technology has evolved rapidly over the years. NVMe drives offer significantly faster performance than SATA/SAS. HDDs are cheaper for systems that require substantial amounts of storage, but SSDs are faster, lighter, more durable, and use less energy. Your needs will dictate the best storage for your system.

- Local storage (HDD, SATA/SAS SSDs, NVMe):
  - Log, data, index requirements
  - RAID (either HW or SW)
- Shared Storage (NAS, SAN):
  - Log, data, index requirements
  - RAID (either HW or SW)

CHAPTER 3

BIOS SETTINGS

Table 3-1 describes BIOS options that most impact performance for common RDBMS systems. Please see the **BIOS & Workload Tuning Guide for AMD EPYC™ 9005 Series Processors** (available from the **AMD Documentation Hub** for additional BIOS tuning information.

NAME	VALUE	DESCRIPTION
TSME	Disabled	Transparent Secure Memory Encryption (TSME) provides hardware mem- ory encryption of all data stored on system DIMMs at a cost of 5-7ns of memory latency.
DETERMINISM CONTROL	Enabled	Enables the **Determinism Slider** control.
DETERMINISM SLIDER	Power	Ensures maximum CPU performance by only throttling the CPUs when they reach the cTDP.
TDP CONTROL	Manual	Enable setting a custom TDP.
TDP	OPN Max	Set the TDP in watts to the maximum supported by OPN.
PACKAGE POWER LIMIT CONTROL	Manual	Enable setting a custom Package Power Limit (PPL)
PACKAGE POWER LIMIT	OPN Max	Set the PPL in watts to maximum supported by OPN.
SMT CONTROL	Auto	Enables Simultaneous Multithreading (SMT), which allows two hardware threads per core.
NUMA NODE PER SOCKET (NPS)	1, 2, or 4	Determines the number of NUMA nodes to split the memory channels between. A higher number indicates fewer memory channels per NUMA node, which lowers both memory throughput and latency for a NUMA node. If your configuration has more than 64 logical CPU's (2 times the number of cores when SMT is either **Enabled** or **Auto**) per socket, then set NPS to 2. If your configuration has more than 128 logical CPU cores per socket, then set NPS to 4.
ACPI SRAT LLC AS NUMA DOMAIN	Varies by workload	Dictates whether to report each LLC as NUMA domain to the OS. Allowing processes that use the same data to be scheduled on the set of CPUs that share an L3 cache increases the chance of L3 Cache hits.
LOCAL APIC MODE	X2APIC	In general, interrupt delivery is faster when using x2APIC mode over the legacy xAPIC mode. Windows 2019 and later have support for x2APIC mode.

Table 3-1: Recommended BIOS settings for most RDBMS applications

CHAPTER 4

OS SETTINGS

Microsoft SQL Server can run on both Linux and Windows operating systems. AMD recommends using the latest available OS version. Please see **AMD EPYC™ Processors Minimum Operating System (OS) Versions** for a list of operating systems supported by AMD EPYC 9005 Series Processors. SQL server supports all of the Windows versions listed therein. For Linux support, please cross-reference the AMD EPYC Operating Systems page with the Installation guidance for SQL Server on Linux found at **SQL Server | Microsoft Learn*** to determine OS requirements for both the processors and SQL Server.

4.1 WINDOWS CONFIGURATION

Most of the default Windows OS configuration is already optimized for SQL Server. However, enabling large pages on Windows Server requires modifying the group policy for the database user to enable the **Lock Pages in Memory** setting.

NAME	VALUE	DESCRIPTION
LOCK PAGES IN MEMORY	<SQL Server User Name>	Enable SQL Server user to lock pages in memory through the group policy editor.

Table 4-1: General Windows settings

4.2 LINUX CONFIGURATION

Modify `sysctl.conf`, `mssql.conf`, and `rc.local` as shown in Table 4-2, Table 4-1, and Table 4-4 to improve SQL Server performance on Linux.

4.2.1 /ETC/SYSCTL.CONF

NAME	VALUE	DESCRIPTION
KERNEL.NUMA_BALANCING	0	Disable Linux kernel auto NUMA balancing.
VM.MAX_MAP_COUNT	262144	Increase settings for Virtual Address Space.
VM.SWAPPINESS	1	Sets a strong preference to keeping process memory in physical memory at the expense of filesystem cache.

Table 4-2: General sysctl.conf settings

4.2.2 /VAR/OPT/MSSQL/MSSQL.CONF

NAME	VALUE	DESCRIPTION
MEMORY.MEMORYLIMITMB	85-90% of system memory	Limits the max memory visible to SQL Server.
CONTROL.WRITETHOUGH	1	For NVMe and any drives that support force unit access (FUA). Use `O_DSYNC` for file flag write through requests.
CONTROL.ALTERNATEWRITETHOUGH	0	Controls write, flush behavior takes place.

Table 4-3: Microsoft SQL Server 2022 mssql.conf settings

4.2.3 /ETC/RC.LOCAL

NAME	VALUE	DESCRIPTION
/SYS/BLOCK/<DEVICE>/QUEUE/SCHEDULER	deadline	Modify elevator scheduler for all SQL Server storage devices.
/SYS/BLOCK/<DEVICE>/QUEUE/NR_REQUESTS	1024	Number of concurrent requests to the I/O device.

Table 4-4: /etc/rc.local entries

4.2.4 TUNED-ADM PROFILE

The throughput-performance tuned-adm profile produces optimal performance. This profile sets up system I/O and memory throughput by configuring the CPU governor, kernel scheduler granularity, disk read ahead, swappiness behavior, and dirty cache write back settings. See '/usr/lib/ tuned/throughput-performance/tuned.conf' for profile-specific settings. Please also see the Kernel and CPU settings for high performance section in **Performance Best Practices and Configuration Guidelines for SQL Server on Linux*** for additional information.

4.2.5 FILESYSTEM MOUNT OPTIONS

The EXT4 filesystem handles smaller file sizes (up to 16 TB) than XFS (up to 500 TB). Table 4-5 shows optimal mount options for both filesystems.

NAME	DESCRIPTION
NOATIME	This option disables updating the metadata associated with files in the filesystem with an updated access time. This tracking is redundant because the database logs its own accesses.
NOBARRIER	Disables the filesystem write barrier. Using a write barrier degrades I/O performance by requiring more frequent data flushes.

Table 4-5: XFS filesystem mount options

CHAPTER 5

SQL SERVER CONFIGURATION SETTINGS

SQL Server has many options and settings that can impact database functionality and performance. Test these options in a non-production environment to verify that they meet your requirements.

Note: Some settings alter and disable certain database features that reduce performance, such as checkpointing, logging, and database recoverability. Review and test your settings before implementing them in a production environment.

5.1 SQL SERVER SOFT NUMA CONFIGURATION

By default, soft-NUMA nodes are created automatically whenever the SQL Server engine detects more than eight physical cores per NUMA node or socket at startup. Multi-threaded cores are not differentiated when counting physical cores in a node. SQL Server Database Engine then creates soft- NUMA nodes that ideally contain eight cores but which can go down to five or up to nine logical cores per node.

You can disable soft-NUMA by using the **`ALTER SERVER CONFIGURATION SET SOFTNUMA OFF`** statement and then restarting the database engine. When Soft-NUMA is enabled the SQL Server error log file will display something similar to the following:

```
19 09:21:46.78 Server             Automatic soft-NUMA was enabled because SQL
Server has detected hardware NUMA nodes with greater than 8 physical cores.

2022-01-19 09:21:46.78 Server     Processor affinity turned on: node 0, processor
mask 0x00000000000000ff. Threads will execute on CPUs per affinity settings. This
is an informational message; no user action is required.

2022-01-19 09:21:46.78 Server     Processor affinity turned on: node 1, processor
mask 0x000000000000ff00. Threads will execute on CPUs per affinity settings. This
is an informational message; no user action is required
...

2022-01-19 09:21:46.78 Server     Processor affinity turned on: node 30, processor
mask 0x00ff000000000000. Threads will execute on CPUs per affinity settings. This
is an informational message; no user action is required.

2022-01-19 09:21:46.78 Server     Processor affinity turned on: node 31, processor
mask 0xff00000000000000. Threads will execute on CPUs per affinity settings. This
is an informational message; no user action is required.
```

Note: This example is from a system with two 64-core processors and Simultaneous Multi-Threading enabled.

Things to look for in this output include the first line showing both that Soft-NUMA was enabled and the processor mask for each node. Verify that the mask includes consecutive bits and that they are associated with the CPUs within a core complex. This ensures that the set of processor threads for a given node are all sharing the same last level cache. If the Automatic Soft-NUMA mask does not meet these requirements, then the cross Level 3 cache communication will increase and hurt performance. If this happens, then you must manually configure the Soft-NUMA nodes via the registry.

The following example uses a 2-socket system with dual 64-core AMD EPYC CPUs and SMT enabled, with 128 cores / 256 threads split across 16 Core Complexes running MSSQL-Server 2022. This example will create 16 soft-NUMA nodes and then group the cores and SMT threads to match the Core Complex boundaries.

```
Windows Registry Editor Version 5.00 [HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\
Microsoft SQL Server\160\NodeConfiguration]

[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft SQL Server\160\NodeConfiguration\
Node0] "CPUMask"=hex:00,00,00,00,00,00,ff,ff
"Group"=dword:00000000

[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft SQL Server\160\NodeConfiguration\
Node1] "CPUMask"=hex:00,00,00,00,ff,ff,00,00
"Group"=dword:00000000

[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft SQL Server\160\NodeConfiguration\
Node2] "CPUMask"=hex:00,00,ff,ff,00,00,00,00
"Group"=dword:00000000

[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft SQL Server\160\NodeConfiguration\
Node3] "CPUMask"=hex:ff,ff,00,00,00,00,00,00
"Group"=dword:00000000

[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft SQL Server\160\NodeConfiguration\
Node4] "CPUMask"=hex:00,00,00,00,00,00,ff,ff
"Group"=dword:00000001

[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft SQL Server\160\NodeConfiguration\
Node5] "CPUMask"=hex:00,00,00,00,ff,ff,00,00
"Group"=dword:00000001

[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft SQL Server\160\NodeConfiguration\
Node6] "CPUMask"=hex:00,00,ff,ff,00,00,00,00
"Group"=dword:00000001

[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft SQL Server\160\NodeConfiguration\
Node7] "CPUMask"=hex:ff,ff,00,00,00,00,00,00
"Group"=dword:00000001

[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft SQL Server\160\NodeConfiguration\
Node8] "CPUMask"=hex:00,00,00,00,00,00,ff,ff
"Group"=dword:00000002

[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft SQL Server\160\NodeConfiguration\
Node9] "CPUMask"=hex:00,00,00,00,ff,ff,00,00
"Group"=dword:00000002

[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft SQL Server\160\NodeConfiguration\
Node10] "CPUMask"=hex:00,00,ff,ff,00,00,00,00
"Group"=dword:00000002

[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft SQL Server\160\NodeConfiguration\
Node11] "CPUMask"=hex:ff,ff,00,00,00,00,00,00
"Group"=dword:00000002
```

```
[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft SQL Server\160\NodeConfiguration\
Node12] "CPUMask"=hex:00,00,00,00,00,00,ff,ff
"Group"=dword:00000003

[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft SQL Server\160\NodeConfiguration\
Node13] "CPUMask"=hex:00,00,00,00,ff,ff,00,00
"Group"=dword:00000003

[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft SQL Server\160\NodeConfiguration\
Node14] "CPUMask"=hex:00,00,ff,ff,00,00,00,00
"Group"=dword:00000003

[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Microsoft SQL Server\160\NodeConfiguration\
Node15] "CPUMask"=hex:ff,ff,00,00,00,00,00,00
"Group"=dword:00000003
```

The mask sometimes comes up as alternating between **0x00000000aaaaaaaa** and **0x00000005555555555**. These CPU masks are undesirable. This is especially true when SMT is enabled because this places the sibling threads for a physical core on two separate Soft-NUMA nodes. If your configuration does this by default, then you will get better performance by switching off the Automatic Soft NUMA and using the manual configuration described above to obtain optimal Soft-NUMA alignment to the CCDs.

5.2 SQL SERVER 2022 INSTANCE TRACE FLAGS AND SETTINGS

Microsoft SQL Server 2022 has many trace flags that tweak how the database interacts and behaves. Table 5-1 lists a few trace flags that will help in most environments.

NAME	VALUE	DESCRIPTION
TRACEFLAG	834	Uses large memory pages. This trace flag shouldn't be enabled by default, but it might improve performance on large Windows machines dedicated to a single SQL Server instance. For most SQL Server deployments on Windows, enabling locked pages is the preferred alternative.Use large pages for buffer pool.
TRACEFLAG	652	Disable read ahead. Consider enabling this trace flag for pure OLTP workloads to reduce unnecessary load on the IO subsystem.

Table 5-1: Microsoft SQL Server 2022 trace flags

Follow guideline outlined in **Performance best practices and configuration guidelines for SQL Server on Linux** for the conditions in which trace flag 3979 might need to be enabled.

Consider using this configuration when checkpoints are managed externally (using the CHECKPOINT command) or on large machines if you observe a performance degradation associated with indirect checkpoints and the IO subsystem is capable of handling IO spikes from periodic checkpoints. Set target recovery interval to 0:

```
ALTER DATABASE <db_name> set TARGET_RECOVERY_TIME = 0 seconds;
```

Pre-grow user database files and tempdb files to their maximum intended sizes and turn off auto-growth on db files and tempdb files:

```
ALTER DATABASE <db_name> MODIFY FILE (NAME=<dbfile_name>, SIZE= <file size>,
FILEGROWTH=0
```

Increase the tempdb file count to match the number of cpu threads. A default SQL Server installation will only create 8 tempdb data files, which is not enough for AMD EPYC processors. Remember to keep all tempdb files the same size to help ensure even usage. Further, tempdb should be stored in the fastest/lowest-latency storage available and distributed across multiple storage devices (either by hardware striping technologies or individual file placements):

```
ALTER DATABASE <db_name> ADD FILE (NAME=<dbfile_name>, FILENAME=<filepath\filename,
SIZE=<size of tempdb file>, FILEGROWTH=0);
```

5.3 SP_CONFIGURE

Instance-level SQL Server configuration settings can increase performance across all workloads. Use the **sp_configure** command to set the parameters shown in Table 5-2.

NAME	VALUE	DESCRIPTION
SHOW ADVANCED OPTIONS	1	Enables View/Modify advanced **sp_configure** options.
AUTOMATIC SOFT-NUMA DIS- ABLED	1	Automatic soft-NUMA disabled" recommendation to "SQL Server performance significantly benefits from manual soft-NUMA setting matching the EPYC CPU Core Complex boundaries. Configure Soft-NUMA manually following the steps outlined: **"SQL Server Soft NUMA Con- figuration" on page 11**.
LIGHTWEIGHT POOLING	1	Uses the fiber mode in Windows. This configuration shouldn't be enabled by default, but it might provide a small performance improvement for workloads affected by excessive context switching between user and kernel modes. Not compatible with a number of SQL Server features.
PRIORITY BOOST ENABLED	1	Increases the OS priority for SQL Server threads. This configuration shouldn't be enabled by default because it might cause CPU starvation for essential OS processes under high load, but it might provide an async IO performance improvement on large Windows machines dedicated to a single SQL Server instance.
MAX SERVER MEMORY (MB)	80-90% of system/VM memory	Maximum memory for SQL Server to use.

Table 5-2: Microsoft SQL Server 2022 configuration settings

CHAPTER 6

WORKLOAD-SPECIFIC CONSIDERATIONS

SQL Server databases are generally deployed for one of two primary workloads that each have unique considerations. The Data Warehouse workload is more query oriented, while the Online Transaction Processing workload is more transaction oriented. Some databases run both workloads, which requires either deciding which workload has higher priority or determining the right tuning balance for each of them.

6.1 DATA WAREHOUSE

SQL Server data warehouses consume a high amount of CPU, memory, and I/O resources. SQL Server provides parallel queries to optimize query execution and index operations on systems with more than one processor. SQL Server can perform query or index operations in parallel by using several OS worker threads for both speed and efficiency. In an environment where only 1 query or operation is running at a time (batched queries/ updates etc.), SQL Server can set the max degree of parallelism to twice the number of cores to fully utilize all system threads. In other cases where several query threads are running simultaneously, it may help to reduce the max degree of parallelism such that a single query does not consume all resources and force other queries to run single-threaded.

Setting NPS to 1 normally generates best query throughput performance. However, single socket systems with 32 or more cores show reliable performance with higher NUMA node per socket (NPS) BIOS setting. In those cases, setting this to 2 or 4 may boost performance boost. AMD therefore recommends setting NPS2 for systems with between 32-cores and 48-cores and NPS4 for systems with more than 48 cores. Test this setting with your specific workloads to evaluate performance.

6.2 ONLINE TRANSACTION PROCESSING

OLTP workloads require high transaction log write performance because of their high insert and/or update rates. AMD strongly recommends SSD and NVMe drives for both data areas and redo logs to match CPU performance. The next level of contention issues arises around transactional latencies in memory, the network latencies, or the data disk(s).

CHAPTER 7

RESOURCES

- From the **AMD Documentation Hub**:
  - Microsoft Windows® Tuning Guide for AMD EPYC™ 9005 Series Processors
  - NVMe Tuning Guide for AMD EPYC™ 9005 Series Processors
- **Performance best practices and configuration guidelines for SQL Server on Linux***
- **Soft-NUMA (SQL Server)***

©2025 ADVANCED MICRO DEVICES, INC. ALL RIGHTS RESERVED.

The information contained herein is for informational purposes only and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for particular purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD's products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale.

TRADEMARKS

AMD, the AMD Arrow logo, AMD EPYC, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Microsoft SQL Server is a registered trademark of Microsoft Corporation in the US or other jurisdictions. Other product names and links to external sites used in this publication are for identification purposes only and may be trademarks of their respective companies.

* Links to third party sites are provided for convenience and unless explicitly stated, AMD is not responsible for the contents of such linked sites and no endorsement is implied.

PID: 58485