

AMD EPYC™ 9005

Microsoft® Windows® Server Tuning Guide



together we advance_data center computing

PID: 58471
v1.0
October 2024

© 2024 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, Windows, and Azure are registered trademarks of Microsoft Corporation in the US and other countries. PCIe is a registered trademark of PCI-SIG Corporation. 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.

DATE	VERSION	CHANGES
June, 2024	0.1	Initial NDA release
October, 2024	1.0	Initial public release

AUDIENCE

This document is intended for a technical audience with a background of configuring Microsoft® Windows® servers.



MICROSOFT® WINDOWS® SERVER

TUNING GUIDE

CONTENTS

CHAPTER 1 - INTRODUCTION	1
1.1 - Operating Systems	1
1.2 - Important Reading	2
CHAPTER 2 - SYSTEM CONFIGURATION BEST PRACTICES	3
2.1 - Processors	3
2.2 - Memory	3
2.3 - I/O	3
2.3.1 - Storage	4
2.3.2 - Network Interface Cards (NICs)	4
CHAPTER 3 - BIOS SETTINGS BEST PRACTICES	7
3.1 - IOD Settings	7
3.2 - NUMA Settings	7
3.3 - Memory Settings	8
3.4 - Power Settings	8
3.5 - Core Setting	8
3.6 - I/O Setting	8
3.7 - Collaborative Processor Performance Control (CPPC)	8
3.8 - NUMA	9
3.9 - x2APIC	10
CHAPTER 4 - WINDOWS TUNINGS AND TOOLS	11
4.1 - OS updates	11
4.2 - Processor Affinity	11
4.2.1 - Task Manager	11
4.2.1.1 - PowerShell	12
4.3 - Managing Microsoft Hyper V	13

CHAPTER 5 - PERFORMANCE MONITORING TOOLS - - - - - 15

5.1 - Microsoft Windows Tools and Resources15

5.2 - AMD Tools and Resources15

5.3 - Other Documents16



CHAPTER 1: INTRODUCTION

This tuning guide describes various parameters that can optimize performance of servers powered by AMD EPYC™ 9005 Series Processors that are running Microsoft® Windows® Server Operating Systems.

The default OEM hardware and BIOS configurations normally provide the best performance across all operating systems and for general workloads; however, additional tuning can improve performance of specific workloads.

This tuning guide discusses the following topics:

- Supported OS versions
- System configuration best practices
- BIOS settings that may impact performance
- OS information and commands that relate to optimization
- Information on further resources to assist with performance and analysis

1.1 - OPERATING SYSTEMS

AMD recommends using the latest available Operating System version. Please see [AMD EPYC™ Processors Minimum Operating System \(OS\) Versions](#) for detailed OS version information.

5th Gen AMD EPYC processors support the following version of Microsoft Windows Server:

- **Windows Server 2025:** This version is in preview as of the time of publishing this tuning guide.
- **Windows Server 2022:** Released in September of 2021. This is the most current Windows Server recommended release.
- **Windows Server 2019:** Includes x2APIC support, as described in [“x2APIC” on page 10](#).

Windows Server 2022 is the current Long Term Servicing Channel (LTSC) Windows Server release. Microsoft may release monthly Windows Cumulative Updates, but full official LTSC Windows Server releases occur only every 3-4 years. AMD's Windows Server support is based upon these LTSC releases i.e., Windows Server 2022, Windows Server 2019, etc. Please see [Windows Server Servicing Channels](#)* for more information.

Note: Microsoft officially introduced the future release of Windows Server 2025 (WS25) in early January, 2024. Microsoft has released [publications](#) about WS25 and you can currently download the Windows Server 2025 Preview build (26100) from their [Evaluation](#)* site. Microsoft has not officially announced a release date for WS25, but AMD is currently validating this newer Windows Server operating system with AMD EPYC 9005 processors.*

You can obtain the latest Windows Server versions from any one of the following links:

- Evaluation Channels:
 - [Windows Server 2019 Evaluation](#)*
 - [Windows Server 2022 Evaluation](#)*
 - [Windows Server 2025 Evaluation](#)*
- [Visual Studio subscription](#) (formerly MSDN or Microsoft Developer Network)*.
- [Microsoft Partner Network](#) (MPN)*.
- [Volume Licensing Service Center](#) (VLSC), if you have valid Software Assurance*

Note: You should install Windows Cumulative Updates as Microsoft releases them.

AMD EPYC 9005 Series Processors require using the latest Windows Server Full Media Refresh releases from respective licensing channels. For more details, latest requirements, and latest known issues and limitations of Windows operating system running on AMD EPYC 9005 based systems, please see the [Windows Server support and installation instructions for the AMD EPYC 9004 Series server processors](#)*.

If you use the Windows Preinstallation Environment (WinPE) to deploy AMD EPYC 9005-based servers (such as for applying Windows Server images and applications from a network share, provisioning the system [e.g., its hard drive], adding a custom shell or GUI to automate deployment, etc.) then you will need an updated WinPE image that includes the latest support for AMD EPYC 9005 Series Processors.

AMD recommends using the ADK 10.1.26100.1 (May 2024) release for full support of AMD EPYC 9005 Series Processors under a WinPE environment. This release supports the Windows Server 2025 Preview build and can be downloaded directly from Microsoft: [Download the ADK 10.1.26100.1 \(May 2024\)](#)*. In addition to this ADK release, you'll also be able to download the [Windows PE add-on 10.1.26100.1 \(May 2024\)](#)*.

If your Windows provisioning environment cannot be upgraded to utilize the latest WinPE release, then AMD recommends using the Windows Server 2022 WinPE image and customizing it by adding the latest Windows Server 2022 Cumulative Update package available from the [Update Catalog](#)*. You can search the official Microsoft Web site for specific information about WinPE and how to update an image, such as [Create Bootable Windows PE Media](#)*, but the high-level steps are:

1. [Download](#)* and install the Assessment and Deployment Kit (ADK) & the “WinPE add-on for the ADK” for Windows Server 2022.
2. Download the latest Windows Server 2022 Cumulative Update package and all prerequisite packages, if any, as required by the Cumulative Update package (See details in its release page) from the [2022 update site](#)*.
3. Add the prerequisite packages followed by the Cumulative Update package to the Windows Server 2022 WinPE image.

1.2 - 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: SYSTEM CONFIGURATION BEST PRACTICES

2.1 - PROCESSORS

AMD provides a variety of EPYC 9005 Series Processors to support a full range of workloads and environments. Processors with higher core counts will provide more computational resources. Processors that run at higher frequencies also will increase performance. Processors with access to larger L3 Cache may also show increased performance. Systems configured with dual 9005 EPYC processors may also provide greater computing power. Please see [AMD EPYC Processors](#) for more details on available 5th Gen AMD EPYC Orderable Part Numbers (OPNs) and models.

2.2 - MEMORY

AMD EPYC 9005 Series Processors support up to 6TB of 6000 MT/sec RAM (1 DIMM slot per channel) or up to 12 TB of 4000 MT/sec RAM (2 DIMM slots per channel) per processor. In general, it is important to have enough memory for corresponding workloads in order to avoid degraded performance caused by excessive paging.

AMD recommends populating all twelve memory channels per CPU socket with an equal number of equal-capacity DIMMs. If your system provides 2 DIMM slots per channel, then populate each channel with at least a single DIMM before populating the second DIMM slot on a given channel.

Please see the latest version of [Memory Population Guidelines for AMD Family 1Ah Models 00h-0Fh and Models 10h-1Fh Socket SP5 Processors](#) - (login required; please review the latest version if multiple versions are present) for additional memory population guidelines.

2.3 - I/O

I/O intensive workloads may perform better if they execute on the cores and memory that correspond to the same quadrant or NUMA node as the PCIe device's I/O hub.

You can obtain more detailed device by locating the device within Windows Device Manager, right-clicking it, and then selecting **Properties** to display the device **Properties** page. Select the **Details** tab and then click **Property Field** to display a list of device-specific settings and their corresponding values.

3rd Gen and prior AMD EPYC processors included the **Preferred I/O** and **Relaxed Ordering** settings that helped optimize network and disk I/O performance. 4th Gen AMD EPYC processors (9xx4 models) and newer include architectural enhancements that deliver optimal network and disk I/O performance by default without the need for either of these features.

2.3.1 - Storage

Systems powered by AMD EPYC processors support a variety of high-performance devices. For example, NVMe SSDs provide both power savings and performance advantages for I/O queues, interrupt processing, etc. compared to older storage technologies. Always provide enough appropriately-distributed disks to avoid bottlenecks. Placing the paging file on its own disk will avoid contention from other processes attempting to access a shared disk.

2.3.2 - Network Interface Cards (NICs)

Place network intensive workloads on the cores of the socket which the NIC also connects to. Many modern day NICs support Receive Side Scaling (RSS), which can distribute network workloads across multiple logical processors. This support allows you to determine the nearest processor-PCIe slot distance:

1. [Start Windows PowerShell*](#).
2. Execute the `Get-NetAdapterHardwareInfo` cmdlet to obtain the connection name for each NIC.

```
PS C:\Users\Administrator> Get-NetAdapterHardwareInfo
```

Name	Segment	Bus	Device	Function	Slot	NumaNode	PciLinkSpeed	PciLinkWidth	Version
Ethernet 2	0	33	0	0	238	9	8.0 GT/s	8	1.1
Ethernet	0	33	0	1	238	9	8.0 GT/s	8	1.1

Figure 2-1: `Get-NetAdapterHardwareInfo` cmdlet and output

3. Execute the `Get-NetAdapterRss -Name <Connection Name>` cmdlet to get the `RssProcessorArray` elements that each show the distance between the processor cores and the PCIe slot where the NIC is currently plugged in.

AMD recommends placing workloads on the cores with the nearest distance, with 0 being optimal because those processor cores belong to the NUMA node that is closest to the NIC.

`RssProcessorArray` returns information formatted as `A : B / C` where

- A = Processor Group
- B = Logical Processor ID within this group/node
- C = Distance between the PCIe slot the NIC is plugged in to and the Logical Processor

AMD recommends using cores with the nearest distance (ideally C = 0, which implies that the processor cores are within the closest NUMA node to the NIC). Please see the latest version of the *Windows® Network Tuning Guide for AMD EPYC™ 9005 Series Processors* (available from the [AMD Documentation Hub](#)) for additional NIC tuning information.

[illegible]

Figure 2-2: Sample `Get -NetAdapterRss` output

THIS PAGE INTENTIONALLY LEFT BLANK.



CHAPTER 3: BIOS SETTINGS BEST PRACTICES

Various BIOS settings can impact Windows Server performance. Default OEM BIOS settings provide general high performance for their servers, but specific workloads and contexts benefit from additional tuning. See the *BIOS & Workload Tuning Guide for AMD EPYC™ 9005 Series Processors* (available from AMD EPYC Tuning Guides) for additional tuning guidelines.

If present, the following BIOS settings can impact both performance and other factors, such as power consumption.

3.1 - IOD SETTINGS

- **xGMI Link Max Speed:** Increasing speed between SoC communication links (which increases power to SoC which may have the unintended consequence of preventing core frequency boost).
- **xGMI Link Width:** Widening the communication link between sockets.
- **SoC P-state:** Forcing the Data Fabric power state to highest performing state.
- **Data Fabric C-states:** Preventing the I/O Die from going into low power state.

3.2 - NUMA SETTINGS

- **LLC as NUMA:** Use the ACPI SRAT table to define the NUMA domain based on CCX boundary, so number of NUMA domains is equal to the number of Last Level Caches or CCXs. Users who desire smaller NUMA Nodes, or to work around potential OS limitations of representing multiple CCXs within the same NUMA Node, may consider enabling LLC as NUMA.
- **NUMA Nodes per Socket (NPS):** This setting also relates to memory interleaving of the memory channels per socket. For example, NPS1 implies a single NUMA domain with all cores within the 9005 processor (socket) and all the corresponding memory in one NUMA domain. Memory is interleaved across the socket's twelve memory channels. All PCIe devices on the socket belong to this single NUMA domain. NPS2 i.e., 2 nodes per socket, interleaves memory across six channels. NPSx setting may be set independently of the LLC as NUMA Domain BIOS setting. With Hyper-V and 2P-64c\SMT enable part, NPSx = 2 or 4 should be used if LLC as NUMA is not enabled.

Setting **LLC as NUMA** to **Enabled** yields NUMA nodes equal to the number of CCDs or L3 caches in the system, such as 16 NUMA nodes on a 2P 64-core EPYC 9005 -based server. Concurrently, you can also use the **NUMA Nodes per Socket** (NPSx) setting to define the memory interleave. However, disabling **LLC as NUMA** means that **NUMA Nodes per Socket** (NPSx) will both provide the memory interleave info and dictate the number of NUMA nodes.

Some NUMA settings may not be available for all AMD EPYC 9005 Series Processor models. For more details, please see the latest version of [Socket SP5/SP6 Platform NUMA Topology for AMD Family 1Ah Models 00h-0Fh and Models 10h-1Fh](#) (login required).

3.3 - MEMORY SETTINGS

Memory Clock Speed: The memory and I/O Die each have a corresponding clock and lower latencies attained as the fabric clock (FCLK) now runs coupled with Memory Clock, to further improve memory latency.

3.4 - POWER SETTINGS

- **Power Determinism:** May increase power to die, to maximize core performance. See the latest version of [Power/Performance Determinism](#) for additional information.
- **CPPC:** Collaborative Processor Performance Control allows OS control over processor boost, as described in [“Collaborative Processor Performance Control \(CPPC\)” on page 8](#).

3.5 - CORE SETTING

SMT: Allowing for 2 execution threads per core. Turning on SMT will provide further gains. However, because there are resources within the core that are shared 2x performance is usually not expected.

3.6 - I/O SETTING

Local APIC Mode: Depending upon Windows Server OS, x2APIC may be suggested. See [“x2APIC” on page 10](#), below.

3.7 - COLLABORATIVE PROCESSOR PERFORMANCE CONTROL (CPPC)

Collaborative Processor Performance Control (CPPC) is defined in the Advanced Configuration and Power Interface (ACPI) specification and provides a mechanism for the OS to potentially request varying performance levels from processors. CPPC replaces traditional discrete frequency P-state requests.

Windows Server provides three power plans:

- Balanced (default).
- High Performance
- Power Saver

The default Windows Server power plan is **Balanced**. To obtain best performance, you can change the power plan to **High Performance**. Each of the three power plans has a separate GUID, for High Performance it is: `8c5e7fda-e8bf-4a96-9a85-a6e23a8c635c`.

For **High Performance** Mode with CPPC, use the Windows `Powercfg, exe` tool using a Window Command window running as Administrator:

- To show the current Power plan/scheme type: `Powercfg /L`

- To change to **High Performance**: `Powercfg /s 8c5e7fda-e8bf-4a96-9a85-a6e23a8c635c`
- To make the power scheme active: `Powercfg /setactive scheme_current`

Additionally, CPPC supports an autonomous mode, which AMD recommends, where AMD's System Management Unit (SMU) hardware promptly selects processor performance levels best suited to the current workload and power limits.

```
Powercfg /setacvalueindex scheme_current sub_processor PERFAUTONOMOUS 1
Powercfg /setacvalueindex scheme_current sub_processor PERFEPP 0
Powercfg /setacvalueindex scheme_current sub_processor PERFBOOSTMODE 4
Powercfg /setactive scheme_current
```

3.8 - NUMA

Using the **NUMA Nodes Per Socket** (NPSx) and **LLC as NUMA** BIOS settings mentioned prior (NUMA related BIOS settings), you can abstract the AMD EPYC 9005 processor into a variety of NUMA configurations. In Windows, systems with fewer than 64 logical processors always have a single Processor Group. Windows will try and assign processors that are closest to each other into the same Processor Group.

There may potentially be multiple NUMA Nodes within that Processor Group. The Windows scheduler sees each processor group as a single entity. The maximum size of a Windows Processor Group is currently 64 logical processors.

Prior to Windows Server 2022, the OS scheduler only assigns application threads of a process to a single processor group. Regardless of NPS settings, applications must be multi-group aware to take advantage of all the processors. Otherwise, their affinity will be to a single processor group. Starting with Windows Server 2022, applications are no longer constrained by default to a single processor group. The OS scheduler assigns thread across all processor groups by default without requiring applications to be multi-group aware. Please see [Processor Groups](#)* for more information about Processor Groups in Windows.

Windows historically supports a maximum of 64 logical processors per NUMA Node. When there were more than 64 logical processors, Windows would internally split the node into multiple soft NUMA nodes. Furthermore, all that node's memory is associated with the first soft NUMA node. Starting with Windows Server 2022 Microsoft introduced: `KeQueryNodeActiveAffinity2` API, which allows a NUMA node to span across more than a single processor group

Note: In Windows Server 2019, Task Manager displays soft NUMA nodes internally created by the OS. If there are more than 64 logical processors per NUMA node, then the number of NUMA nodes shown in Task Manager will not match the actual number of NUMA nodes configured in the system. Windows Server 2022 fixes this Task Manager limitation.

Please see [NUMA Architecture](#)* for more information about NUMA in Windows.

Be aware of BIOS settings which allow for more than 64 logical processors per node. For example, a 2 socket, 128-core AMD EPYC 9005-based system with SMT enabled and with NPSx equal to 0 (or 1), would have 512 (or 256) threads per NUMA Node respectively. Settings which result in >64 threads per NUMA node are not supported by Windows Server 2019. Users who desire smaller NUMA Nodes, or wish to work around potential OS limitations of representing multiple CCXs within the same NUMA Node, may consider enabling **LLC as NUMA**.

Please see the NUMA Topology section of the *AMD EPYC™ 9005 Series Architecture Overview* (available from the [AMD Documentation Hub](#)) for additional information.

3.9 - x2APIC

On systems with more than 255 logical processors such as dual-socket servers using the AMD EPYC 95x5, that is 2P 64c with SMT enabled, and which are running Windows Server 2019; x2APIC mode must be used to take advantage of all the cores. x2APIC also provides a potentially more efficient mechanism for interrupt delivery and is necessary for more than 255 logical processor support.

Server Operating System	WS2019 and Newer
Support for Interrupt Mechanism	x2APIC
Support for Max. OPN Core Count / Logical Processor	2 x 192c 768 LPs

Table 3-1: x2APIC support by Windows Server version

Please see [“Operating Systems” on page 1](#) for more details about supported Windows operating systems and builds.



CHAPTER 4: WINDOWS TUNINGS AND TOOLS

This chapter provides settings and links to tools that may help you analyze and optimize Windows Server performance on systems powered by AMD EPYC 9005 Series Processors.

4.1 - OS UPDATES

Microsoft regularly releases OS security and functionality patches. AMD recommends running Windows Update and installing the latest OS patches.

4.2 - PROCESSOR AFFINITY

You can set processor affinity in Windows using either:

- Task Manager
- Windows PowerShell

4.2.1 - Task Manager

To set Windows processor affinity using Task Manager running as Administrator:

1. Select the app you want to affinitize.
2. Right-click the desired app, and then select **Details** to show the app processes

3. Right-click the desired process, and then select **Set Affinity**.




Figure 4-1: Setting affinity using Windows Task Manager

4.2.1.1 - PowerShell

To set Windows processor affinity using PowerShell:

1. Open PowerShell.
2. Obtain affinity by entering `Get-Process appname | Select-Object ProcessorAffinity`.
3. Set affinity by entering `$Process = Get-Process app; $Process.ProcessorAffinity=mask- value`.

For example, running `get-process` shows all currently-running processes:

```
PS C:\> Get-Process win32calc | Select-Object ProcessorAffinity
```

`ProcessorAffinity` will return a bitmask representing the processors that the threads in the associated process can run on. So, to run a process on first CCX (assuming SMT is off):

```
PS C:\> $Process = Get-Process win32calc; $Process.ProcessorAffinity=15
PS C:\> Get-Process win32calc | Select-Object ProcessorAffinity
```

Calculator will now only run on the first 4 logical processors.




Figure 4-2: Sample Calculator output

4.3 - MANAGING MICROSOFT HYPER V

Hyper-V is either a Type 1 hypervisor or bare-metal hypervisor, meaning that it runs directly on the underlying hardware and not on the OS. Hyper-V adds a level of protection by virtualizing system resources to guest VMs. The Microsoft OS portfolio uses Hyper-V for virtualization in Windows Client, Windows Server, Azure Stack HCI, and Azure.

Hyper-V provides isolated guest operating system environments running on a single server platform. Each isolated partition receives its own resources. Hyper-V has a parent (or management) Root Partition, which is a virtual machine partition with unique access and increased privileges. This root partition creates the requested isolated child partitions. AMD recommends not running other applications on the Root Partition.

Enabling Hyper-V on AMD EPYC 9005-based servers requires the following:

- **Windows Server 2025:** At the time of this publication, Windows Server 2025 Preview build (version 26100) is available for evaluation purposes and can be downloaded directly from Microsoft.
- **Windows Server 2022:** If the OS build is prior to 20348.859, then install the latest Windows Cumulative Update before enabling Hyper-V.
- **Windows Server 2019:** If the OS build is prior to 17763.3532, then install the latest Windows Cumulative Update before enabling Hyper-V.

Use the following cmdlet to enable Hyper-V on Windows Server via PowerShell:

```
Install-WindowsFeature -Name Hyper-V -IncludeManagementTools -Restart
```

[Hyper-V Minimum Root configuration \(Miniroot\)](#)* allows you to specify the maximum number of logical processors available to the root partition. Set the root partition number of logical processors to N by opening an elevated Command Prompt window, executing the following `bcdedit` command, and then rebooting for the command to take effect:

```
C:\ > bcdedit /set hypervisorrootproc N
```

Please see:

- [Hyper-V Host CPU Resource Management](#)* for an overview of Hyper-V Miniroot.
- [Windows Server 2019 Hyper-V host behavior running in the Minroot configuration](#)* for behaviors and known issues specific to Windows Server 2019 Hyper-V Miniroot.

`Cpugroups.exe` is a tool from the Hyper-V Team that allows you to allocate processing resources at a much more granular level. See [Virtual Machine Resource Controls](#)* for information about Microsoft VM CPU tools.

[Hyper-V Integration Services](#)* provide enlightened guest OS drivers for Hyper-V and will improve performance of the guest OS while it is running on Hyper-V.

For Hyper-V scalability and capacity planning, please consult the current maximum configuration limits supported by Hyper-V. For example, Windows Server 2019 Hyper-V supports a maximum of 320 logical processors to the root partition. The root partition may also not support all 320 logical processors depending on the NUMA configuration. In this example, a dual-socket system powered by 96-core AMD EPYC processors with SMT enabled and NPS=4 running Windows Server 2019 Hyper-V will have only 320 out of the 384 available logical processors supported by the root partition. In the root partition, Task Manager shows **320** in the **Host logical processors** field and **384** in the **Logical processors** field.




Figure 4-3: Task Manager view showing 320 available logical processors

Be aware that the family of AMD EPYC 9005 Series Processors will include models with physical more than 96 and up to 192 processor cores. A 96-core AMD EPYC 9005 processor with Simultaneous Multithreading (SMT) enabled will produce 384 total processor threads. This same processor in a dual 2P server will produce 768 total threads. Windows Server 2019 limits this configuration because Hyper-V only supports up to 512 total logical processors/threads on the host and will therefore not utilize all 768 threads. In such scenarios, AMD recommends using Windows Server 2022 or newer with support for up to at least 1024 total threads on the Hyper-V host.

5th Gen AMD EPYC processors support a maximum of 384 threads per processor socket, meaning that the number of threads per NUMA node will increase compared to prior AMD EPYC processors. Windows Hyper-V has a limit of 256 logical processor threads per NUMA node. Lower NUMA node granularities, such as NPS1 or NPS0, could cause the theoretical total number of threads in a NUMA node could easily exceed this 256 limitation. For example, a 5th Gen AMD EPYC processor with 384 total threads (physical cores + SMT) set to NPS1 will not have all 384 threads started due to the Hyper-V NUMA limitation and only 256 threads will be made available. AMD recommends working around this limitation by using a more granular NPS setting in cases where total thread counts will exceed the 256 thread per Hyper-V NUMA limitation. In the previous example, using NPS2 instead of NPS1 will allow Hyper-V to start 192 threads per NUMA node. At the time of publication, AMD is working with Microsoft to address this Hyper-V limitation in current and future Windows Server releases.

Windows Server 2025 (Preview) will continue support both Gen1 and Gen2 Virtual Machines (VMs). However, the WS25 Hyper-V Manager UI will create Gen2 VMs by default. If you decide to create a Gen1 VM with the Hyper-V Manager, then Gen1 VM creation will fail on host servers that enable greater than 64 total logical processors because a Gen1 VM is limited to 64 vCPUs and WS25 Hyper-V will create the VM with more than 64 vCPUs by default when available. If this occurs, you can instead create a Gen1 VM using the Powershell `New-VM` cmdlet with the `-Generation 1` parameter to create a VM with one vCPU by default. You can then use the Hyper-V Manager to configure the Gen1 VM.

Additional details are available for specific Windows Server Releases at [Plan for Hyper-V scalability in Windows Server*](#). Microsoft continues to update these maximum limits in response to business needs. Please continue to consult this article on an ongoing basis.



CHAPTER 5: PERFORMANCE MONITORING TOOLS

5.1 - MICROSOFT WINDOWS TOOLS AND RESOURCES

Windows includes the following resources:

- **Windows Resource Monitor:** Monitors CPU, memory, disk, and network usage in real time.
- **Windows Performance Monitor:** View OS\Hyper-V performance counters in either real time or saved from the Performance recorder. This feature is part of Windows Admin Center and can thus track multiple servers.

Microsoft also provides various kits, such as:

- **Windows Assessment and Deployment Kit (ADK):** Includes a variety of useful performance tools, such as:
- **Windows Performance Recorder (WPR) / xPerf:** Recording utility based on Event Tracing for Windows (ETW).
- **Windows Performance Analyzer:** Creates graphs and data tables of recorded events.
- **Windows Software Development Kit (SDK):** Includes the powerful Windows Kernel Debugger (WinDbg) that can investigate OS functionality and view data structures.

Please also see the following Microsoft documents for additional information:

- [Windows NUMA Support and APIs*](#)
- [Powercfg configuration options*](#)

5.2 - AMD TOOLS AND RESOURCES

- [uProf for Windows:](#) Profiling tool that can monitor system metrics and performance counters.

Other AMD EPYC resources are available from:

- <https://developer.amd.com/resources/epyc-resources/>
- www.amd.com/epyc-tech-docs
- www.amd.com/epyc-tuning-guides

5.3 - OTHER DOCUMENTS

[ACPI Specification v6.5*](#)

THIS PAGE INTENTIONALLY LEFT BLANK.

Microsoft® Windows® Server Tuning Guide for AMD EPYC™ 9005 Processors

PID: 58471

Chris Karamatas and David Tanaka

