

## WEATHER MODELING

# WRF® Performance with AMD EPYC™ 7003 Series Processors

MAY 2021

## AMD EPYC™ 7003 Processors


Built on the x86 architecture innovations of the record setting<sup>1</sup> AMD EPYC™ 7002 series processors, the AMD EPYC™ 7003 Series Processors set a new standard for the modern data center. With high frequencies, high core counts, high memory bandwidth and capacity, and up to 32MB of L3 cache per core, the AMD EPYC 7003 series processors enable exceptional HPC performance across multiple industry verticals.

Along with the high memory bandwidth achieved with support for 8 channels DDR4-3200 memory, AMD EPYC 7003 CPUs also synchronize the data fabric clock to match the memory clock speeds, further improving both memory bandwidth and latency. The support for up to 4TB of memory per socket enhances the ability to efficiently utilize very large datasets.

Extra-large caches, reaching up to 256MB per CPU and up to 32MB per core, help to efficiently utilize up to 64 cores per CPU. The large caches drive performance, but also enable exceptional scalability on many key workloads. 128-160 lanes of PCIe® Gen4 offered by EPYC 7003 Series CPUs pushes the ability to efficiently access high-speed network interface cards, high-speed storage, and multiple accelerators.

AMD EPYC 7003 Series Processors are designed to bring faster time-to-value by delivering performance and scalability, while also helping keep your data secure. AMD Infinity Guard<sup>2</sup> helps protect your data both on-prem and in the cloud.

Speed is the new metric for efficiency. The AMD EPYC 7003 Series processors raise the bar once more for workload performance, helping to drive faster time to results for delivering improved business outcomes.



### AMD EPYC 7003 FOR HPC

AMD EPYC 3rd Gen CPUs can deliver the highest per-core performance in the industry by taking advantage of the faster CPU frequencies, lower latency memory and unified cache structure. With AMD EPYC, the HPC community can propel innovations and insights with ground-breaking high-performance computing and advanced security features to deliver excellent results.

### “ZEN 3” CORE & SECURITY

Featuring up to:

- 64 physical cores, 128 threads
- 256MB of L3 cache per CPU
- 32MB of L3 cache per core
- 4TB of DDR4-3200 memory
- 128-160 PCIe® Gen 4 lanes

Industry leading security with Infinity Guard Security

- Secure Boot
- Encrypted memory with SME and SEV

### SCALE OUT AND SCALE UP

Scaling is critical to HPC applications. AMD EPYC 7003 Series processors provide high bandwidth between nodes with support for PCIe Gen 4 enabled network devices and accelerators. Within a node, take advantage of up to 64 cores, 8 memory channels of DDR4-3200, and up to 256MB of L3 cache-per-CPU.

### WRF®

Developed and maintained by the National Center for Atmospheric Research (NCAR), the Weather Research & Forecasting (WRF) model has over 48,000 registered users in over 160 countries. WRF is a flexible and computationally efficient platform for operational forecasting across scales ranging from meters to thousands of kilometers.

## EPYC 7003 Series Architecture Quick Look

The AMD EPYC 7003 Series Processor retains the proven Multi-Chip Module (MCM) Chiplet Architecture of prior successful AMD EPYC server-class processors while making further improvements. One of the most important upgrades is the new “Zen 3” core. The “Zen 3” core is manufactured using a 7nm process and designed to provide a significant instructions per cycle (IPC) uplift over prior generation “Zen 2” cores. Similar to EPYC 7002 Series processors, each core supports Simultaneous Multi-Threading (SMT) allowing up to 2 threads per core. In a typical 2-socket system with 64-core processors, EPYC 7003 Series processors offer up to 128 physical cores per system and up to 256 threads per system.

The L3 cache was also improved in the 3<sup>rd</sup> Gen EPYC processors. EPYC 7003 Series CPUs took the same total L3 cache as the prior generation (up to 256MB/CPU) and created significantly more cache sharing between cores. The 3<sup>rd</sup> Gen EPYC processors now offer a unified 32MB of L3 cache per compute die. With this generation, up to 8 cores per compute die can now share 32MB of unified L3 cache.



Figure 1: EPYC 7003 Processor L3 Cache Layout

The new L3 Cache design can increase the cache hit to miss ratio over the previous design. Improved cache sharing also allows larger blocks to fit directly into the cache whereas previously it would fall into the main memory. Improvements made in the cache fetching and eviction policies manage data more efficiently. All these benefits result in an uplift on HPC workloads in addition to the core and memory improvements.

## EPYC 7003 Series CPU Options and Recommendations by Segment

AMD EPYC 7003 Series CPUs offer 19 different CPU configurations. Below is a table of each CPU with a summary of their features. For driving up per-core performance, pay special attention the 7xF3 processors, which offer the highest cache and frequencies at their respective core-counts of the 7003 series processors. HPC applications come in a wide range of unique characteristics. There is no one-size fits all CPU for the HPC market. Below are the general recommendations. You are encouraged to talk to your AMD sales representative for more detailed guidance and CPU suggestions based on your unique environment and needs.

Model	# CCDs	Cores / Threads	Base Freq (GHz)	Max Boost Freq (Up to GHz) <sup>3</sup>	Default TDP (W)	cTDP (W)	L3 Cache (MB)	NPS	2P/1P
<b>7763</b>	8	64 / 128	2.45	3.50	280W	225-280W	256	1,2,4	2P/1P
<b>7713</b>	8	64 / 128	2.00	3.675	225W	225-240W	256	1,2,4	2P/1P
<b>7713P</b>								1,2,4	1P
<b>7663</b>	8	56 / 112	2.0	3.5	240W	225-240W	256	1,2,4	2P/1P
<b>7643</b>	8	48 / 96	2.3	3.6	225W	225-240W	256	1,2,4	2P/1P
<b>75F3</b>	8	32 / 64	2.95	4.0	280W	225-280W	256	1,2,4	2P/1P
<b>7543</b>	8	32 / 64	2.8	3.7	225W	225-240W	256	1,2,4	2P/1P
<b>7543P</b>								1,2,4	1P
<b>7513</b>	4	32 / 64	2.6	3.65	200W	165-200W	128	1,2,4	2P/1P
<b>74F3</b>	8	24 / 48	3.2	4.0	240W	225-240W	256	1,2,4	2P/1P
<b>7453</b>	4	28 / 56	2.75	3.45	225W	225-240W	64	1,2,4	2P/1P
<b>7443</b>	4	24 / 48	2.85	4.0	200W	165-200W	128	1,2,4	2P/1P
<b>7443P</b>								1,2,4	1P
<b>7413</b>	4	24 / 48	2.65	3.6	180W	165-200W	128	1,2,4	2P/1P
<b>73F3</b>	8	16 / 32	3.5	4.0	240W	225-240W	256	1,2,4	2P/1P
<b>7343</b>	4	16 / 32	3.2	3.9	190W	165-200W	128	1,2,4	2P/1P
<b>7313</b>	4	16 / 32	3.0	3.7	155W	155-180W	128	1,2,4	2P/1P
<b>7313P</b>								1,2,4	1P
<b>72F3</b>	8	8 / 16	3.7	4.1	180W	165-200W	256	1,2,4	2P/1P

Table 1: AMD EPYC 7003 CPU Options

Listed in the table below are a few examples of different HPC market segments, a general sense of the characteristics and sensitivities of applications per segment, and specific AMD EPYC 7003 Series processor recommendations per segment.

Segment	Sensitivity	Example Applications	Suggested Models	Comments
FEA Explicit	Frequency & Cache	LS-DYNA®, Radioss™, Abaqus™, VPS™	75F3   7543 74F3   73F3	Look for CPUs with high frequencies, and large caches. Mid core-counts help increase performance per core to help maximize software investment.
FEA Implicit		Mechanical™, Abaqus Standard™, OptiStruct™		
Molecular Dynamics	Core Count & Frequency	GROMACS, LAMMPS	7763   7713 7663   7643 75F3   7543	Look for CPUs with high core-counts and frequency. These applications scale very well with cores.
Weather	Memory BW & Cache	WRF, IFS	7763   7713 7663   7643 75F3   7543 74F3	Look for CPUs with 256 MB of cache. Large caches help relieve the potential memory bandwidth bottleneck if using high core counts. Look for mid core-count CPUs for per-core licensed codes.
CFD		Fluent®, AcuSolve™		
Oil & Gas		Reveal®, Echos, SAVA		
EDA	Frequency & Cache	VCS®, RedHawk™	73F3   72F3	This market segment is dominated by low core-counts to drive up the frequency and cache per core, helping maximize software investment.

Table 2: HPC Segment Recommendations

## Weather Research Forecast (WRF)

The Weather Research and Forecasting (WRF) Model is a popular mesoscale numerical weather prediction system designed for both atmospheric research and operational forecasting applications. It features two dynamical cores, a data assimilation system, and a software architecture supporting parallel computation and system extensibility. The model serves a wide range of meteorological applications across scales from tens of meters to thousands of kilometers.

For researchers, WRF can produce simulations based on actual atmospheric conditions (i.e., from observations and analyses) or idealized conditions. WRF offers operational forecasting a flexible and computationally efficient platform, while reflecting recent advances in physics, numerics, and data assimilation contributed by developers from the expansive research community. WRF is currently in operational use at National Centers for Environmental Prediction (NCEP) and other national meteorological centers as well as in real-time forecasting configurations at laboratories, universities, and private companies.

## Test Methodology

This document focuses on performance and scaling of the EPYC 7003 Series Processors as well as competitiveness with Intel Xeon Gold 6258R.

Testing was performed on dual-socket EPYC™ 7763-based systems. Each EPYC 7763 processor has 64 cores with a base frequency of 2.45GHz and a boost frequency of up to 3.50GHz. The compute nodes in the cluster are each populated with 1 DIMM per channel of 64-GB, dual-rank, DDR4-3200 DIMMs from Micron®, for a total of 1TB of memory per node. A Mellanox® ConnectX-6 200 Gb/s HDR InfiniBand adapter, utilizing EPYC processors' support for PCIe Gen 4, is also populated on each EPYC processor-based system.

Testing was also run on a dual-socket Intel® Xeon® Gold 6258R-based platforms. The 6258R was selected because it offers the highest frequency of the highest core-count (28c) in the Intel Xeon Gold family of processors. The Intel platforms were populated with 1 DIMM per channel of 64GB, dual-rank, DDR4-2933 DIMMS (768GB total memory), matching the maximum memory speed supported for this processor.

As an additional point of comparison, testing was also done on a cluster of dual socket EPYC™ 7713-based systems. EPYC 7713 processors have 64 cores, but with a lower base frequency of 2.0GHz and a boost frequency of up to 3.675GHz.

This paper compares the performance of the *CONUS 2.5* workload, the weather over the Continental US divided into a grid of 2.5 km<sup>2</sup> blocks. This is a common production workload for WRF.

WRF performance is typically measured in “mean time per step”. To show relative performance, mean time per step was converted to performance (1/ mean time per step), and performance was used as the basis for the comparison (AMD Performance/Intel Performance, higher is better).

Every benchmark was run a minimum of 3 iterations, with the average of the performance results used in the brief. Results of each benchmark were also confirmed to have <1% variability between all runs.

## System Configuration

AMD System Configuration		
CPU	2 x AMD EPYC 7713	2 x AMD EPYC 7763
Frequency: Base   Boost <sup>3</sup>	2.0GHz   3.675GHz	2.45GHz   3.5GHz
Cores	64 Cores/Socket (128 Cores/Node)	
L3 Cache	256MB	
Memory	1TB (16x) Dual-Rank DDR4-3200 64GB DIMMs, 1DPC	
NIC	Mellanox ConnectX-6 HDR 200Gb InfiniBand x16 (OFED-4.5-1.0.1)	
Storage: OS   Data	1 x 256 GB SATA   1 x 1 TB NVMe	
BIOS and Settings	SMT=off, X2APIC=on, IOMMU=off, APBDIS=1, Fixed SOC P-state=0, Determinism=Power, NPS=4, DF C-states=off, PIO, EPIO, TSME=off, PCIe 10-bit tag=on	
OS Settings	Clear caches before every run, NUMA balancing 0, randomize_va_space 0, CC6 Disabled, Governor=Performance	

Table 3: AMD EPYC System Configuration

Competitive System Configuration	
CPU	2 x Intel Xeon Gold 6258R
Frequency: Base   Turbo	2.7GHz   4.0GHz
Cores	28 cores per socket (56 per node)
L3 Cache	38.5MB
Memory	768 GB (12x) Dual-Rank DDR4-2933 64GB DIMMs, 1DPC
NIC	Mellanox ConnectX-6 HDR 200Gb InfiniBand x16 (OFED-4.5-1.0.1)
Storage: OS   Data	1 x 256 GB SATA   1 x 1 TB NVMe
BIOS and Settings	3.3a: Power Management=Extreme Performance, Hyper-threading=Off, SNC=On, ADDDC=Off
OS Settings	clear caches before every run, NUMA balancing 0, randomize_va_space 0

Table 4: Intel System Configuration.

Software	
Solver Version	WRF 2.5
MPI	OpenMPI v4.0.3
OS	RHEL 8.3

Table 5 Software

## WRF Single-Node Performance

Single-node performance testing was performed on multiple AMD EPYC 7003 based systems and Intel® Xeon® Gold 6258R based systems.

The graph below shows relative speedup of third-generation AMD EPYC processors running the WRF CONUS 2.5km model. The EPYC 7713 and 7763 processors deliver outstanding performance, clearly outperforming the Intel Xeon Gold 6258R.




Figure 2: 7003 WRF single-node performance for all systems tested.

All AMD EPYC processors tested outperform the Intel Xeon Gold 6258R baseline in the CONUS 2.5km benchmark. A balance of high memory bandwidth and large caches yields best results:

- The 64-core AMD EPYC 7763 processor outperformed the Intel baseline system by an average of ~1.83, with a base frequency 2.45 GHz and max boost of up to 3.5 GHz. It was the most performant processor tested.
- The 64-core 7713 processor outperformed the Intel baseline system by an average of ~1.73, with a base frequency of 2.0 GHz and a max boost of up to 3.675 GHz. With the same cache 256 MB cache, core count and lower max TDP of 240 W compared to the EPYC 7763, this processor maintained very high performance on this key workload.

## WRF Multi-Node Scaling

The chart below shows the multi-node scaling performance of AMD EPYC 7763-based systems up to 8 compute nodes (1024 cores) running the CONUS 2.5km benchmark model.

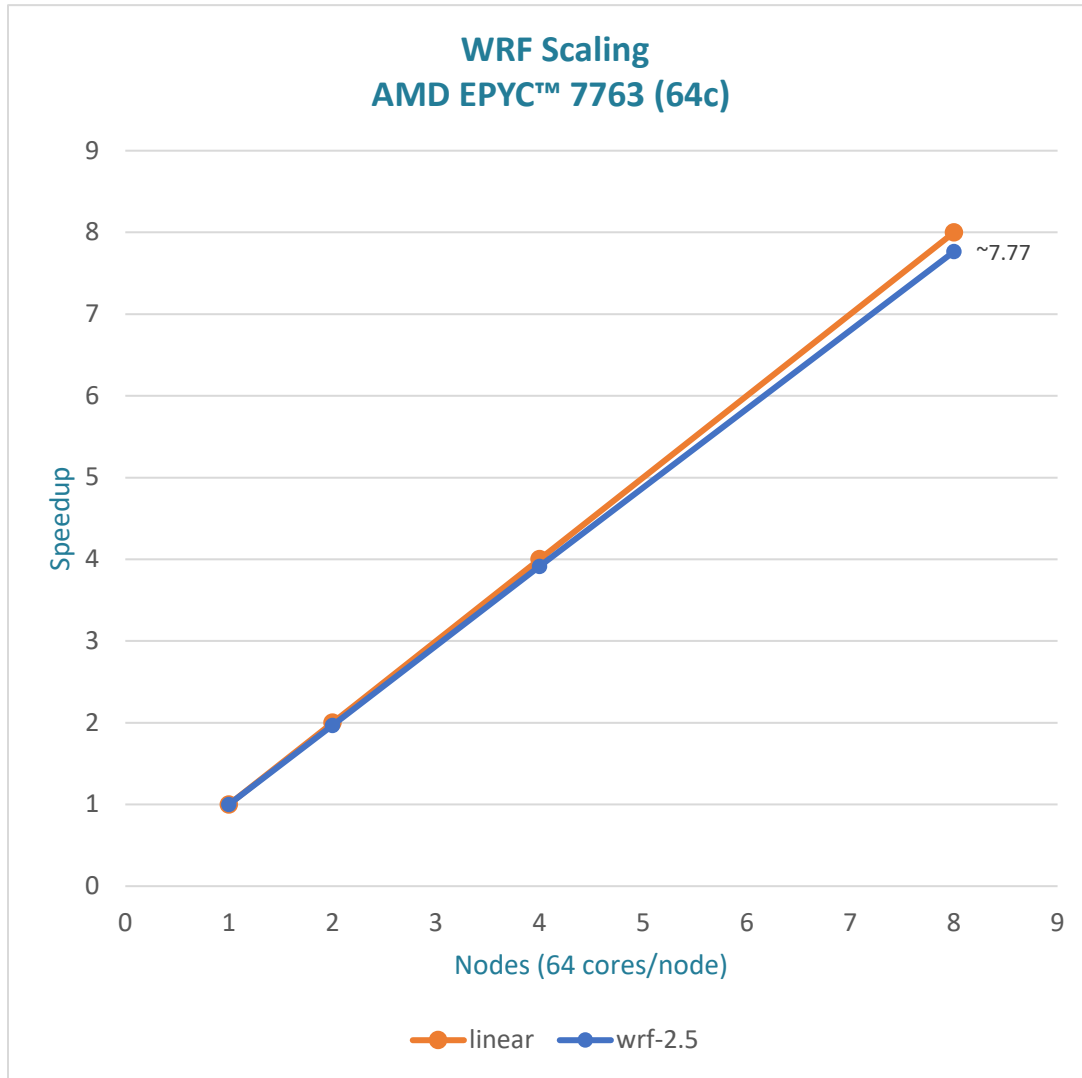


Figure 3: WRF multi-node scaling performance on a core basis for all systems tested.

At 8 nodes, the expectation for scaling would be a speedup of 8. The third-generation EPYC 7763 can achieve a speedup of 7.77 at 8 nodes. This indicates that as the workload is scaled out to multiple systems, the efficiency of each system remains exceptionally high.

## Test Results Summary

The WRF CONUS 2.5km benchmark WRF was run on two different AMD EPYC 7003 series processors and the performance was compared to a baseline Intel® Xeon® Gold® 6258R processor. As a result of these competitive comparisons, single and multi-node scaling results have shown overall performance dominance for all AMD EPYC processors tested as compared to the competition:

- All AMD EPYC processors tested outperformed the Intel baseline system by an average of at least ~1.73x.
- The 64-core AMD EPYC 7763 is the overall performance leader with an average uplift of ~1.83x.
- The 64-core AMD EPYC 7713, with a lower max TDP of 240W, achieves an average per-node uplift of ~1.73x.

## Conclusion

With two high performance cores, a data assimilation system, and a software architecture supporting parallel computation and system extensibility, WRF requires speed and scalability. The 3rd Gen AMD EPYC processors bring excellent power and performance to develop fast, accurate weather modeling, and simulations running on cost effective clustered systems. Whether you are scaling-out to higher node-counts, or driving more density into your clusters, AMD EPYC 7003 series processors give you more levers to tilt the performance balance to your advantage.

### RELATED LINKS

- [AMD EPYC™ Processors](#)
- [AMD EPYC Technical Briefs and Tuning Guides](#)
- [WRF](#) \*

### FOOTNOTES

1. For a complete list of world records see <http://amd.com/worldrecords>. EPYC-22
2. AMD Infinity Guard features vary by EPYC™ Processor generations. Infinity Guard security features must be enabled by server OEMs and/or Cloud Service Providers to operate. Check with your OEM or provider to confirm support of these features. Learn more about Infinity Guard at <https://www.amd.com/en/technologies/infinity-guard> . GD-183
3. For AMD EPYC processors is the maximum frequency achievable by any single core on the processor under normal operating conditions for server systems. EPYC-18

### DISCLAIMERS

\*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.

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.

© 2021 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, EPYC, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Micron and the Micron orbit logo are trademarks of Micron Technology, Inc. NVMe is a trademark of NVM Express, Inc. Mellanox is a registered trademark of Mellanox Technologies. PCIe is a registered trademark of PCI-SIG Corporation. WRF is a registered trademark of the University Corporation for Atmospheric Research. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

### ACKNOWLEDGEMENTS

This document is created with contributions from Alvaro Fernandez.