



**AMD Instinct™ Accelerator  
Qualified Servers  
Q4 2022**

**AMD  
INSTINCT**



# AMD Instinct™ Accelerator Qualified Servers Q4 2022

AMD  
INSTINCT



## AMD Instinct™ Accelerator-Based Servers From Our Partners

Whether your workload is machine learning or high performance computing (HPC), our partners offer servers that can help speed results and return value on your investment. Choose among servers powered by AMD EPYC™ CPUs, the world's fastest server processors,<sup>1</sup> combined with the world's fastest HPC and AI data center accelerators<sup>2</sup>, the AMD Instinct MI200 Series accelerators and other AMD Instinct products.

PARTNER	PRODUCT	RACK UNITS	MAX CPUS	PROCESSOR SERIES	MAX GPUS	AMD INSTINCT MODEL	TYPE	INFINITY FABRIC™ HIVES
AMAX	<a href="#">AceleMax DGS-214A Server</a>	2U	1	EPYC 7002   7003	4	MI210	PCIe®	N/A
	<a href="#">AceleMax DGS-424AI Server</a>	4U	2	EPYC 7003	4	MI250	OAM	4 GPU on board <sup>3</sup>
	<a href="#">AceleMax DGS-428A Server</a>	4U	2	EPYC 7003	8	MI210	PCIe	Dual (4 GPU)
Asus Computer	<a href="#">ESC4000A-E10</a>	2U	1	EPYC 7002   7003	4	MI100	PCIe	N/A
	<a href="#">ESC4000A-E11</a>	2U	1	EPYC 7002   7003	4	MI210	PCIe	Dual (2 GPU)
	<a href="#">ESC8000A-E11</a>	4U	2	EPYC 7003	8	MI100   MI210	PCIe	Dual (4 GPU)
ASRock	<a href="#">1U4G-ROME</a>	1U	1	EPYC 7002	4	MI100	PCIe	N/A
Atos	<a href="#">BullSequana X410-A5 Series 2U1N2S</a>	2U	2	EPYC 7002   7003	8	MI210	PCIe	Quad (2 GPU)
	<a href="#">BullSequana X410-A5 Series 2U1N2S</a>	2U	2	EPYC 7002   7003	4	MI250	OAM	4 GPU on board <sup>3</sup>
	<a href="#">BullSequana X450-A5 Series 2U1N2S</a>	2U	2	EPYC 7002   7003	3	MI210	PCIe	N/A
Cancon	<a href="#">SUMA X7440AQ</a>	4U	2	EPYC 7002	4	MI100	PCIe	N/A
	<a href="#">SUMA X7840HQ</a>	4U	2	EPYC 7002   7003	8	MI100   MI210	PCIe	Dual (4 GPU)
Colfax	<a href="#">CS2460g-EI8 Rackmount Server</a>	2U	2	EPYC 7003	4	MI250	OAM	4 GPU on board <sup>3</sup>
	<a href="#">CS2850g-EK8 Rackmount Server</a>	2U	1	EPYC 7002   7003	8	MI210	PCIe	N/A
	<a href="#">CX4850g-EK8 Rackmount Server</a>	4U	2	EPYC 7002   7003	8	MI100   MI210	PCIe	Dual (4 GPU)
Dell Technologies	<a href="#">PowerEdge R750XA Rack Server</a>	2U	2	3rd Gen Intel® Xeon®	4	MI100   MI210 <sup>4</sup>	PCIe	Dual (2 GPU)
	<a href="#">PowerEdge R7525 Rack Server</a>	2U	2	EPYC 7002   7003	3	MI100   MI210 <sup>4</sup>	PCIe	N/A
	<a href="#">PowerEdge R7515 Rack Server</a>	2U	1	EPYC 7002   7003	1	MI210 <sup>4</sup>	PCIe	N/A

PARTNER	PRODUCT	RACK UNITS	MAX CPUS	PROCESSOR SERIES	MAX GPUS	AMD INSTINCT MODEL	TYPE	INFINITY FABRIC™ HIVES
Exxact Corporation	<a href="#">TensorEX TS2-185671979</a>	2U	1	EPYC 7002   7003	4	MI210	PCIe	N/A
	<a href="#">TensorEX TS2-158632687</a>	2U	2	EPYC 7002   7003	4	MI100   MI210	PCIe	N/A
	<a href="#">TensorEX TS4-173535991</a>	4U	2	EPYC 7002   7003	8	MI100   MI210	PCIe	Dual (4 GPU)
	<a href="#">TensorEX TS2-128058437</a>	2U	2	EPYC 7003	4	MI250	OAM	4 GPU on board <sup>3</sup>
GIGABYTE	<a href="#">G242-Z12 Server</a>	2U	1	EPYC 7003	4	MI210	PCIe	N/A
	<a href="#">G262-Z00 Server</a>	2U	2	EPYC 7002   7003	4	MI250	OAM	4 GPU on board <sup>3</sup>
	<a href="#">G292-Z40 (Microsemi)</a>	2U	2	EPYC 7003	8	MI210	PCIe	Quad (2 GPU)
	<a href="#">G292-Z43 Server</a>	2U	2	EPYC 7003	8	MI100   MI210	PCIe	N/A
	<a href="#">G292-Z44 Server</a>	2U	2	EPYC 7002   7003	8	MI100   MI210	PCIe	Dual (4 GPU)
	<a href="#">G292-Z45 Server</a>	2U	2	EPYC 7003	8	MI100   MI210	PCIe	Quad (2 GPU)
	<a href="#">G482-Z51 Server</a>	4U	2	EPYC 7002   7003	8	MI100	PCIe	Dual (4 GPU)
	<a href="#">G482-Z52 Server</a>	4U	2	EPYC 7002   7003	8	MI100	PCIe	Dual (4 GPU)
	<a href="#">G482-Z53 Server</a>	4U	2	EPYC 7002   7003	8	MI100	PCIe	Dual (4 GPU)
	<a href="#">G482-Z54 Server</a>	4U	2	EPYC 7002   7003	8	MI100   MI210	PCIe	Dual (4 GPU)
	<a href="#">G492-Z51 Server</a>	4U	2	EPYC 7001   7002	8/10	MI100   MI210	PCIe	N/A
	<a href="#">G492-Z52 Server</a>	4U	2	EPYC 7001   7002	8/10	MI100   MI210	PCIe	N/A
	<a href="#">R282-Z93 Server</a>	2U	2	EPYC 7003	3	MI210	PCIe	N/A
	<a href="#">T181-Z70 Server</a>	10U	2	EPYC 7001   7002	4	MI100   MI210	PCIe	N/A
	HPE	<a href="#">ProLiant DL385 Gen10 Plus v2 Server</a>	2U	2	EPYC 7003	3	MI100	PCIe
<a href="#">Cray EX235a Accelerator Blade (water cooled)</a>		2-Node Blade	1 per Node	EPYC 7003	4 per Node	MI250X	OAM	4 GPU on board <sup>3</sup>
<a href="#">Apollo 6500 Gen10 Plus Server</a>		6U	Modular server with nodes as specified below					
<a href="#">ProLiant XL675d Gen10 Plus Server</a>		N/A	2	EPYC 7003	8	MI100   MI210 <sup>4</sup>	PCIe	Dual (4 GPU)
<a href="#">ProLiant XL645d Gen10 Plus Server</a>		N/A	1	EPYC 7003	4	MI100   MI210 <sup>4</sup>	PCIe	4 GPU
Inspur	<a href="#">NF5468A5</a>	4U	2	EPYC 7002	8	MI100   MI210	PCIe	N/A
Koi Computers	<a href="#">G262-Z00 Server</a>	2U	2	EPYC 7002   7003	4	MI250	OAM	4 GPU on board <sup>3</sup>
	<a href="#">G482-Z51 Server</a>	4U	2	EPYC 7002   7003	8	MI100	PCIe	Dual (4 GPU)
	<a href="#">G482-Z52 Server</a>	4U	2	EPYC 7002   7003	8	MI100	PCIe	Dual (4 GPU)
	<a href="#">G292-Z44 Server</a>	2U	2	EPYC 7002   7003	8	MI100   MI210	PCIe	N/A
Lenovo	<a href="#">ThinkSystem SR655 Rack Server</a>	2U	1	EPYC 7003	3	MI210	PCIe	N/A
	<a href="#">ThinkSystem SR665 Rack Server</a>	2U	2	EPYC 7003	3	MI210	PCIe	N/A
	<a href="#">ThinkSystem SR670 v2 Rack Server</a>	2U	2	3rd Gen Xeon	4	MI210	PCIe	4 GPU
Nettrix	<a href="#">X640 G40 Server</a>	4U	2	3rd Gen Xeon	8	MI100   MI210	PCIe	Dual (4 GPU)

PARTNER	PRODUCT	RACK UNITS	MAX CPUS	PROCESSOR SERIES	MAX GPUS	AMD INSTINCT MODEL	TYPE	INFINITY FABRIC™ HIVES
<b>NOR-TECH</b>	<a href="#">HPC Cluster Solutions</a>	Varies	2	EPYC 7002   7003	4   8	MI100   MI210	PCIe	N/A
<b>Penguin Computing</b>	<a href="#">Altus XE2214GT Server</a>	2U	2	EPYC 7002   7003	4	MI100   MI210	PCIe	N/A
	<a href="#">Altus XE4218GT Server</a>	4U	2	EPYC 7002   7003	8	MI100   MI210	PCIe	Dual (4 GPU)
	<a href="#">Altus XO1214GT Server</a>	10U	2	EPYC 7002   7003	4	MI100   MI210	PCIe	N/A
	<a href="#">Relion XE4218GT Server</a>	4U	2	3rd Gen Xeon	8	MI100   MI210	PCIe	Dual (4 GPU)
	<a href="#">7049GP-TRT GPU SuperWorkstation</a>	N/A	2	3rd Gen Xeon	4	MI100	PCIe	N/A
<b>Supermicro</b>	<a href="#">AS-2014CS-TR</a>	2U	1	EPYC 7003	2	MI100   MI210	PCIe	N/A
	<a href="#">AS-2024US-TRT</a>	2U	2	EPYC 7003	2	MI210	PCIe	N/A
	<a href="#">AS-2114GT-DNR (2-Node totals listed)</a>	2U	2	EPYC 7002   7003	6	MI100   MI210	PCIe	N/A
	<a href="#">SBA-4119SG SuperBlade® Server (Maximum 20 1-socket, 1-GPU nodes)</a>	8U	20	EPYC 7002   7003	20	MI100   MI210	PCIe	N/A
	<a href="#">AS-4124GS-TNR</a>	4U	2	EPYC 7002   7003	8/10	MI100   MI210	PCIe	Dual (4 GPU)
	<a href="#">AS-4124GQ-TNMI</a>	4U	2	EPYC 7003	4	MI250	OAM	4 GPU on board <sup>3</sup>
	<a href="#">GPU SuperWorkstation SYS-7049GP-TRT</a>	4U	2	2nd Gen Xeon	4	MI100	PCIe	N/A
	<a href="#">SuperServer 4029GP-TRT</a>	4U	2	2nd Gen Xeon	8	MI100	PCIe	Dual (4 GPU)
	<a href="#">SuperServer 4029GP-TRT2</a>	4U	2	2nd Gen Xeon	8	MI100	PCIe	Dual (4 GPU)
	<a href="#">SYS-740GP-TNRT Rackmountable Workstation</a>	4U	2	3rd Gen Xeon	4	MI210	PCIe	Dual (4GPU)
<b>Symmetric Computing</b>	<a href="#">Ada Departmental GPU Supercomputer</a>	Varies	2	EPYC 7002   7003	N/A	MI100   MI210	PCIe	N/A
<b>TYAN</b>	<a href="#">Transport HX TN83-B8251</a>	2U	2	EPYC 7002   7003	4	MI210	PCIe	N/A
	<a href="#">Transport HX TN83-B8030</a>	N/A	2	EPYC 7002   7003	4	MI210	PCIe	N/A

## For More Information Visit [AMD.com/INSTINCT](https://www.amd.com/instinct)

### FOOTNOTES

- 1: MLN-016B: Results as of 07/06/2021 using SPECrate®2017\_int\_base. The AMD EPYC 7763 scored 854, <http://spec.org/cpu2017/results/res2021q3/cpu2017-20210622-27664.html> which is higher than all other 2P scores published on the SPEC® website.
- 2: World's fastest data center GPU is the AMD Instinct™ MI250X. Calculations conducted by AMD Performance Labs as of Sep 15, 2021, for the AMD Instinct™ MI250X (128GB HBM2e OAM module) accelerator at 1,700 MHz peak boost engine clock resulted in 95.7 TFLOPS peak theoretical double precision (FP64 Matrix), 47.9 TFLOPS peak theoretical double precision (FP64), 95.7 TFLOPS peak theoretical single precision matrix (FP32 Matrix), 47.9 TFLOPS peak theoretical single precision (FP32), 383.0 TFLOPS peak theoretical half precision (FP16), and 383.0 TFLOPS peak theoretical Bfloat16 format precision (BF16) floating-point performance. Calculations conducted by AMD Performance Labs as of Sep 18, 2020 for the AMD Instinct™ MI100 (32GB HBM2 PCIe® card) accelerator at 1,502 MHz peak boost engine clock resulted in 11.54 TFLOPS peak theoretical double precision (FP64), 46.1 TFLOPS peak theoretical single precision matrix (FP32), 23.1 TFLOPS peak theoretical single precision (FP32), 184.6 TFLOPS peak theoretical half precision (FP16) floating-point performance. Published results on the NVIDIA Ampere A100 (80GB) GPU accelerator, boost engine clock of 1410 MHz, resulted in 19.5 TFLOPS peak double precision tensor cores (FP64 Tensor Core), 9.7 TFLOPS peak double precision (FP64), 19.5 TFLOPS peak single precision (FP32), 78 TFLOPS peak half precision (FP16), 312 TFLOPS peak half precision (FP16 Tensor Flow), 39 TFLOPS peak Bfloat 16 (BF16), 312 TFLOPS peak Bfloat16 format precision (BF16 Tensor Flow), theoretical floating-point performance. The TF32 data format is not IEEE compliant and not included in this comparison. <https://www.nvidia.com/content/dam/en-zz/Solutions/Data-Center/nvidia-ampere-architecture-whitepaper.pdf>, page 15, Table 1, MI200-01
- 3: Fully connected CPU and GPU components through AMD Infinity Fabric™ technology with cache coherency.
- 4: Server support coming soon

© 2022 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow, AMD Instinct, ROCm, and EPYC and combinations thereof are trademarks of Advanced Micro Devices, Inc. PCIe® is a registered trademark of PCI-SIG Corporation. SPEC®, SPECrate® and SPEC CPU® are registered trademarks of the Standard Performance Evaluation Corporation. See [www.spec.org](http://www.spec.org) for more information. Other trademarks used in this publication are for identification purposes only and may be trademarks of their respective companies. LE-82101-02 10/22