
AMD
ACCELERATED COMPUTING
AMD

CAUTIONARY STATEMENT

CAUTIONARY STATEMENT

This presentation contains forward-looking statements concerning Advanced Micro Devices, Inc. (AMD) such as the features, functionality, performance, availability, timing and expected benefits of AMD products and product roadmaps, which are made pursuant to the Safe Harbor provisions of the Private Securities Litigation Reform Act of 1995. Forward-looking statements are commonly identified by words such as "would," "may," "expects," "believes," "plans," "intends," "projects" and other terms with similar meaning. Investors are cautioned that the forward-looking statements in this presentation are based on current beliefs, assumptions and expectations, speak only as of the date of this presentation and involve risks and uncertainties that could cause actual results to differ materially from current expectations. Such statements are subject to certain known and unknown risks and uncertainties, many of which are difficult to predict and generally beyond AMD's control, that could cause actual results and other future events to differ materially from those expressed in, or implied or projected by, the forward-looking information and statements. Investors are urged to review in detail the risks and uncertainties in AMD's Securities and Exchange Commission filings, including but not limited to AMD's most recent reports on Forms 10-K and 10-Q.

AMD does not assume, and hereby disclaims, any obligation to update forward-looking statements made in this presentation, except as may be required by law.

WORLD'S LARGEST HYPERSCALERS

RUN ON AMD EPYC™

EPYC PROCESSORS HAVE BEEN DESIGNED INTO DATA CENTERS BY TEN OF

THE WORLD'S LARGEST HYPERSCALE COMPANIES

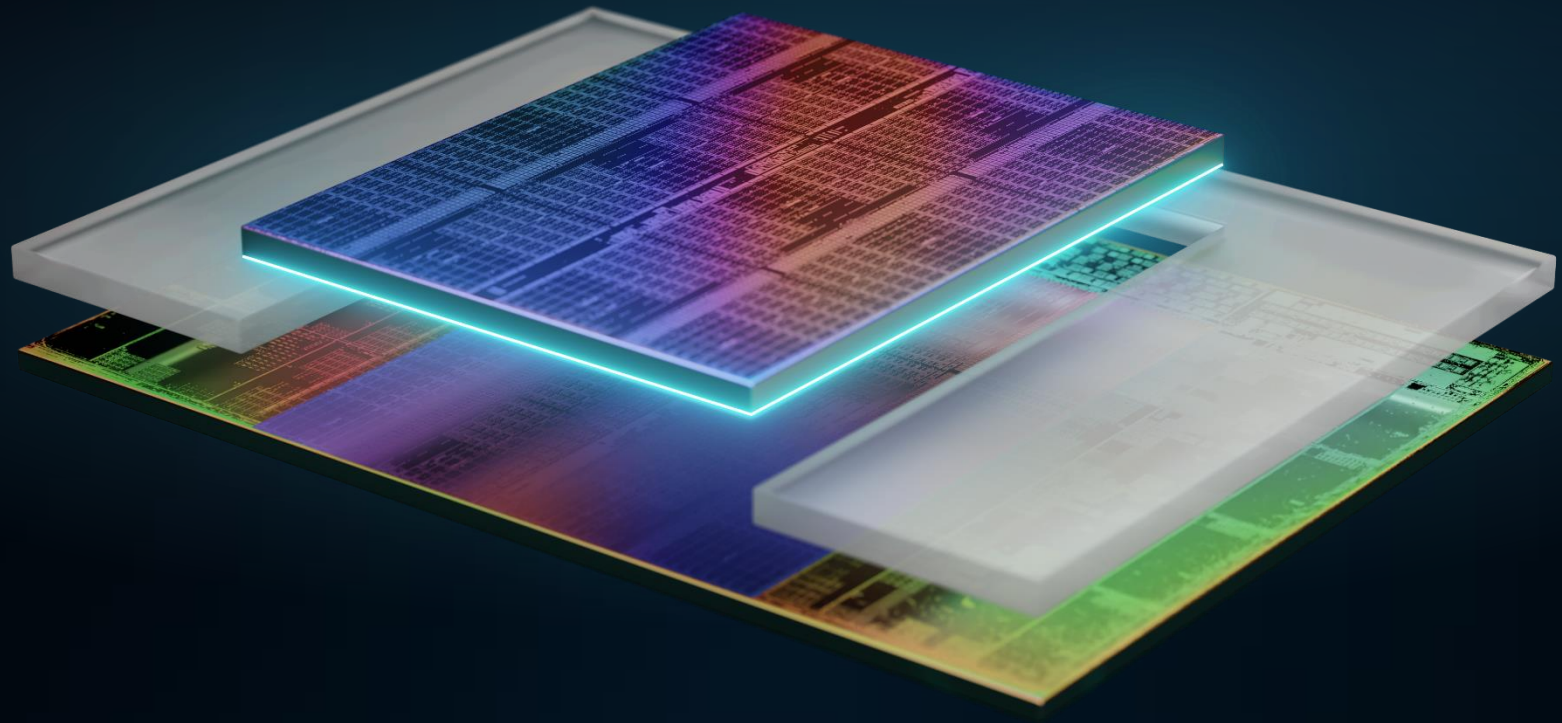

WORLD'S FIRST SERVER CPU WITH 3D CHIPLET TECHNOLOGY

3rd GEN AMD EPYC™ WITH 3D V-CACHE

3X L3 CACHE*	UP TO **804**MB TOTAL CACHE PER SOCKET	UP TO **64** "ZEN 3" CORES	**SP3** SOCKET COMPATIBLE

*3rd Gen AMD EPYC™ Processors with 3D V-Cache have 768MB of L3 Cache, while 3rd Gen EPYC processors without 3D V-Cache have 256MB

CHALLENGING DATA CENTER WORKLOADS

DEMAND COMPUTING INNOVATION

GENERAL PURPOSE
COMPUTING

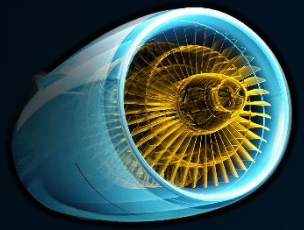

TECHNICAL
COMPUTING

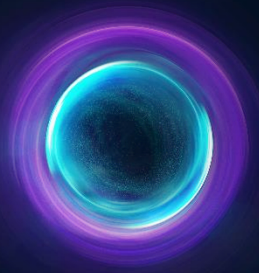

ACCELERATED
COMPUTING

CLOUD NATIVE
COMPUTING

DESIGNED FOR
TECHNICAL COMPUTING

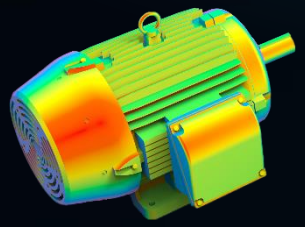

FINITE ELEMENT ANALYSIS

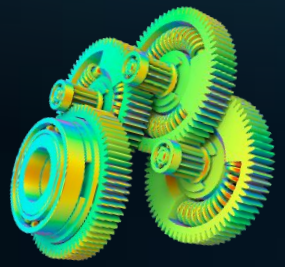

STRUCTURAL ANALYSIS

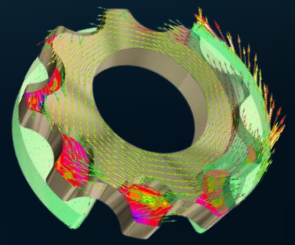

COMPUTATIONAL FLUID DYMANICS

ELECTRONIC DESIGN AUTOMATION

TODAY'S LEADER FOR TECHNICAL COMPUTING WORKLOADS

PERFORMANCE UPLIFT PER CORE

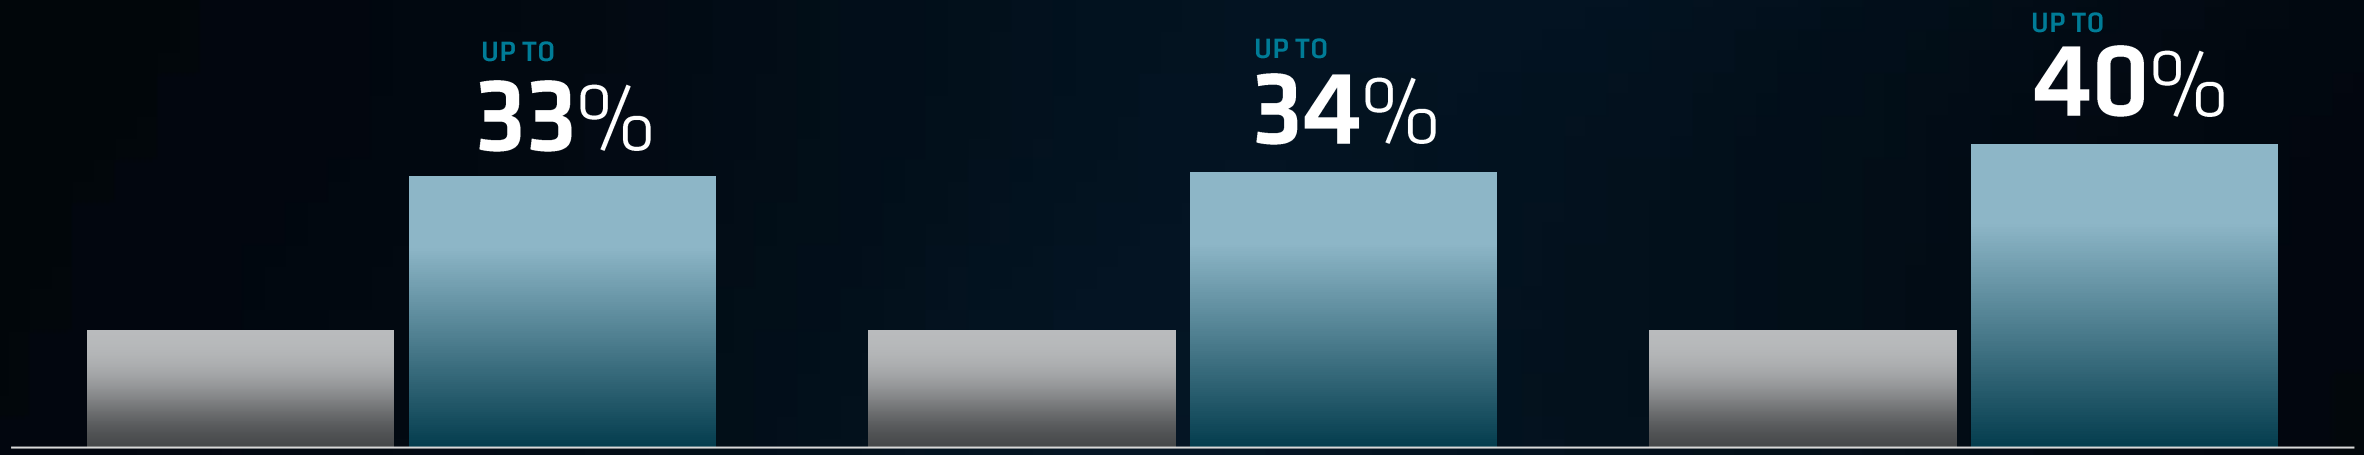

Ansys Mechanical®
Finite Element Analysis
mech-cg1

Altair Radioss®
Structural Analysis
rad-neon

Ansys CFX®
Fluid Dynamics
CFX_10

2P XEON® 8362
32 CORES

2P EPYC™ 75F3
32 CORES

See endnotes MLN-130A, MLN-075A and MLN-080B

AMD

SYNOPSYS® VCS® FUNCTIONAL VERIFICATION PERFORMANCE

AMD EPYC™ "Milan"
16-Core Without 3D V-Cache

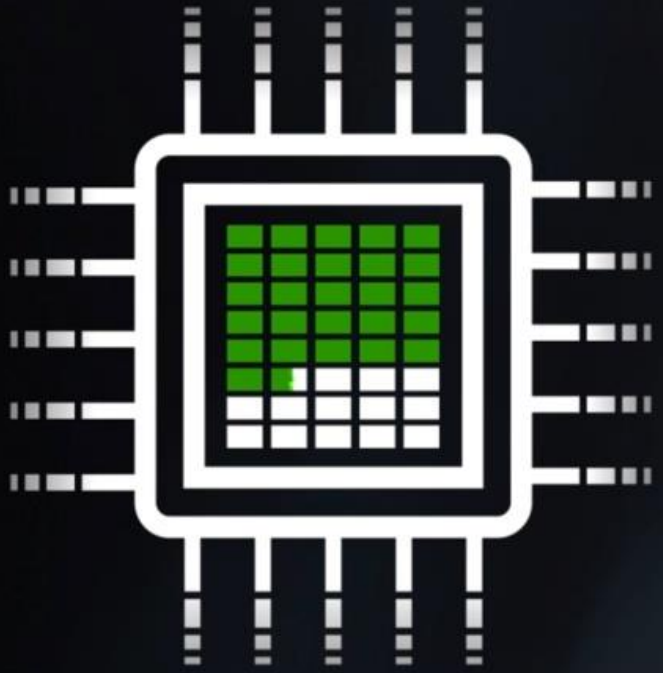

JOBS COMPLETED: **24.40**

AMD EPYC™ "Milan-X"
16-Core With 3D V-Cache

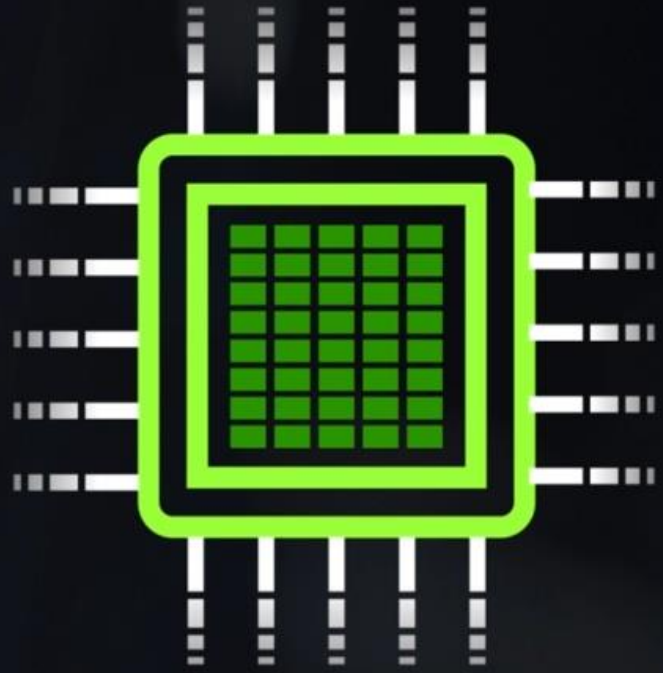

JOBS COMPLETED: **40.55**

Results may vary. See endnotes MLNX-001R

24.4
JOBS/HOUR

3rd GEN AMD EPYC™ 16-CORE
WITHOUT 3D V-CACHE

66%
FASTER RTL VERIFICATION

SYNOPSYS® VCS®

40.6
JOBS/HOUR

3rd GEN AMD EPYC™ 16-CORE
WITH 3D V-CACHE

Results may vary. See endnote MLNX-001R

"MILAN-X"

ISV ECOSYSTEM

AUTOMOTIVE

CHEMICAL ENGINEERING

FINANCE

ENERGY

GEOSCIENCES

LIFE SCIENCES

MANUFACTURING

ALTAIR

Ansys

cadence

SIEMENS

SYNOPSYS®

CRASH SIMULATIONS

FLUID DYNAMICS

FORECASTING

STRUCTURAL ANALYSIS

EXPLORATION

SIMULATION

DESIGN VERIFICATION

FASTER | **BEST SOLUTION VALUE** | **EFFICIENT**

BROADER DATA CENTER BENEFITS OF AMD 3D V-CACHE

ENTERPRISE APPLICATIONS

Data Mining

Risk Analysis

Anomaly Detection

MEDIA AND ENTERTAINMENT

Live Broadcasting

Visual Effects

Realtime Rendering

ARTIFICIAL INTELLIGENCE

Recommendation Engines

Natural Language Processing

Image Recognition

CHALLENGING DATA CENTER WORKLOADS

DEMAND ACCELERATED COMPUTING INNOVATION

GENERAL PURPOSE
COMPUTING

TECHNICAL
COMPUTING

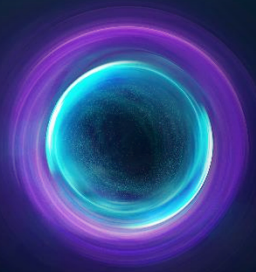

ACCELERATED
COMPUTING

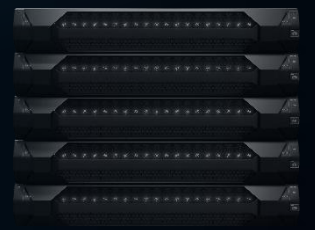

CLOUD NATIVE
COMPUTING

COMPUTE GPU ARCHITECTURE ROADMAP

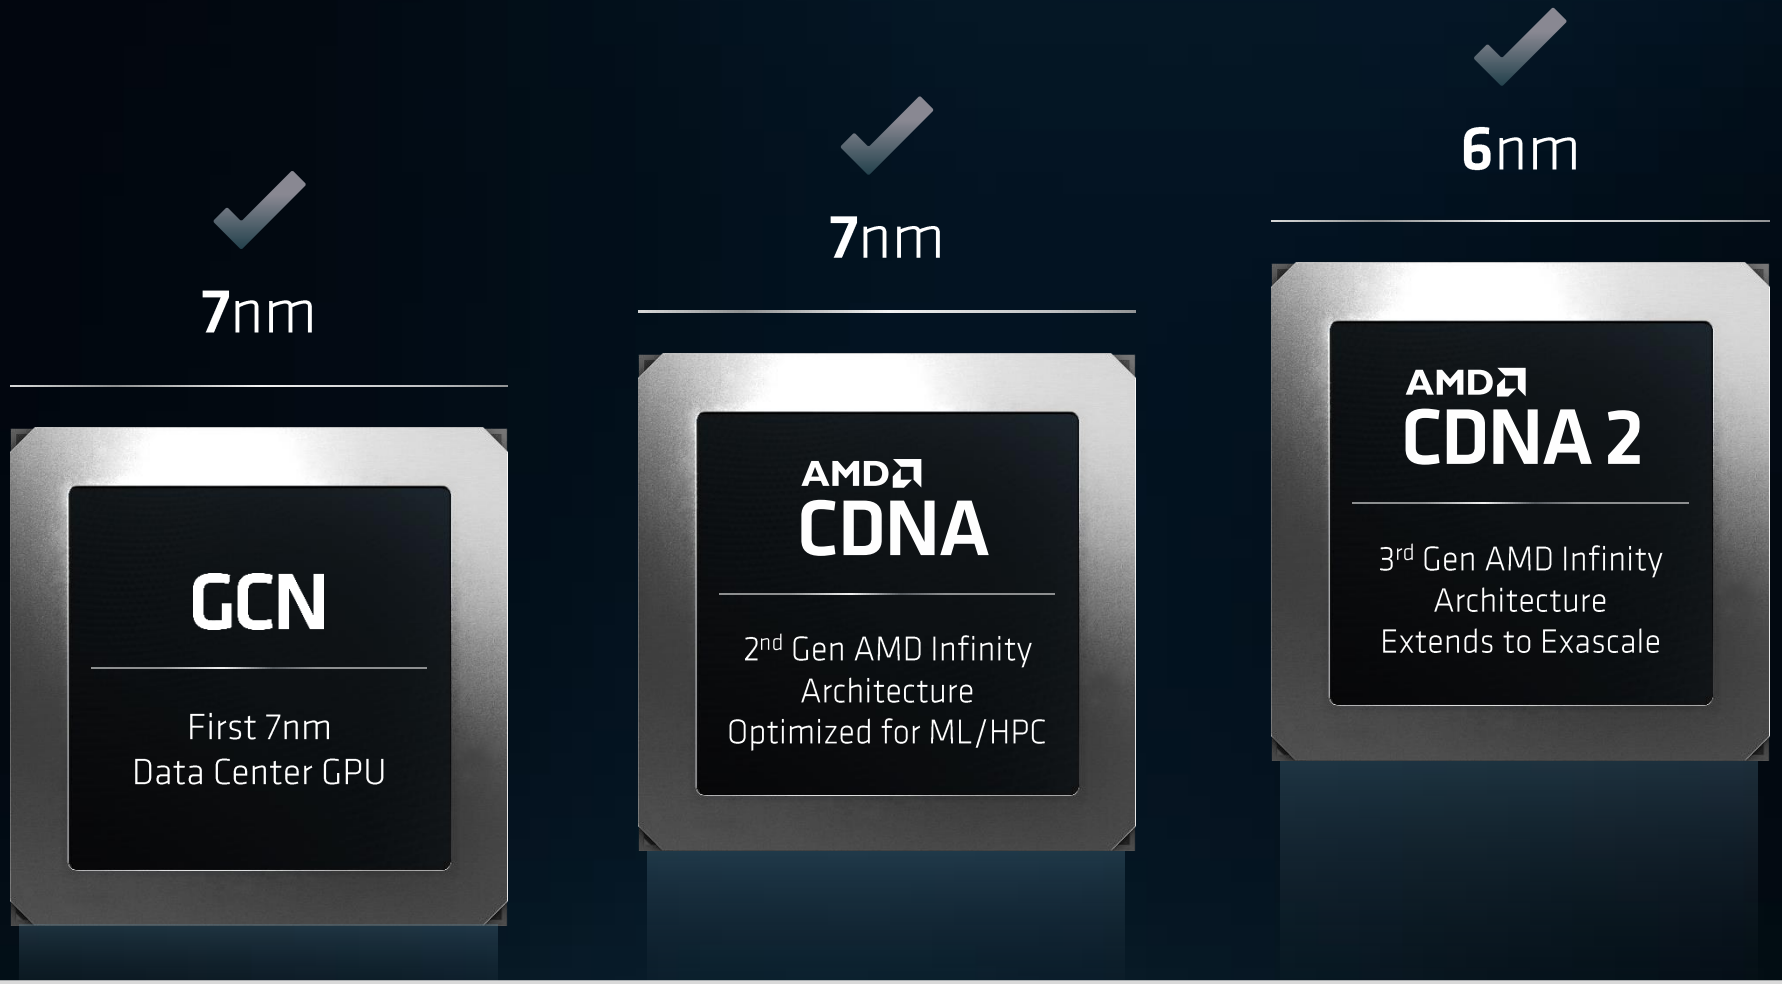

ROADMAP SUBJECT TO CHANGE

AMD INSTINCT™ MI200 SERIES

WORLD'S FIRST EXASCALE-CLASS GPU

LEADERSHIP **HPC**	LEADERSHIP **AI**	LEADERSHIP **SCIENCE**
Up to 4.9X Faster than Competition	Up to 20% Faster than Competition	Fueling Exascale Discoveries with ROCm™ Open Ecosystem

See endnotes MI200-01, MI200-02

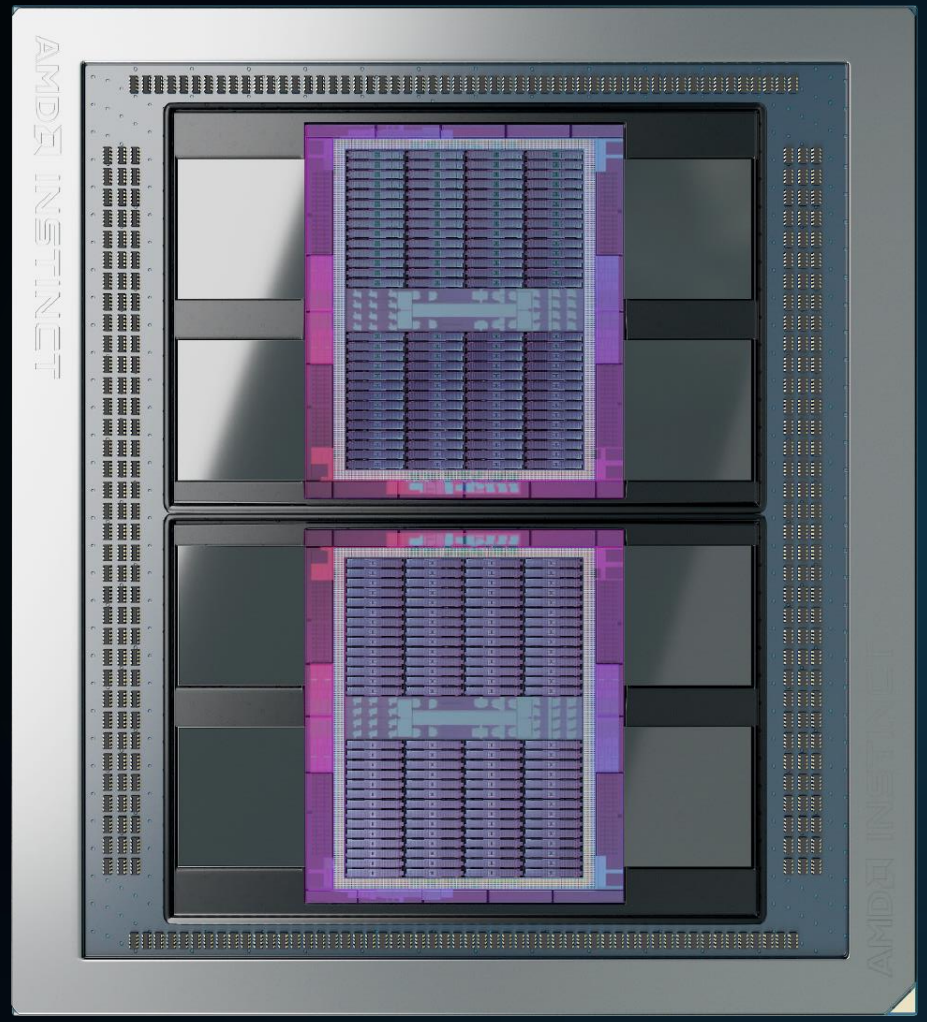

AMD INSTINCT™ MI200 SERIES

WORLD'S MOST ADVANCED DATA CENTER ACCELERATOR

UP TO 58B
Transistors in 6nm

UP TO 220
Compute Units

UP TO 880
2nd Gen Matrix Cores

UP TO 128
GB HBM2E @ 3.2 TB/s

See endnotes MI200-07, MI200-31

AMD INSTINCT™ MI200 SERIES

AMD INSTINCT™

MI200 OAM SERIES

MI250, MI250X

AMD INSTINCT™

MI200 PCIe® SERIES

MI210

Coming Soon

AMD

AMD INSTINCT™ MI200 SERIES

AMD CDNA 2

WORKLOAD-OPTIMIZED COMPUTE ARCHITECTURE

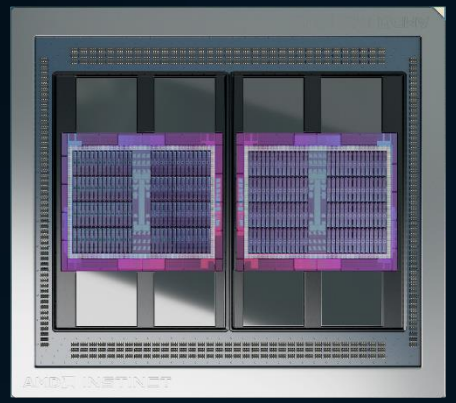

FIRST MULTI-DIE GPU

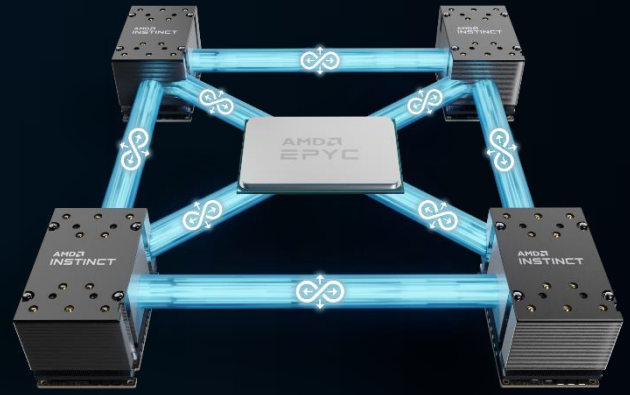

3rd GEN AMD INFINITY ARCHITECTURE

AMD

AMD INSTINCT™ MI200 SERIES

KEY INNOVATIONS

TWO AMD CDNA™2 CHIPLETS

ULTRA HIGH BANDWIDTH CHIPLET INTERCONNECT

COHERENT CPU-TO-GPU INTERCONNECT

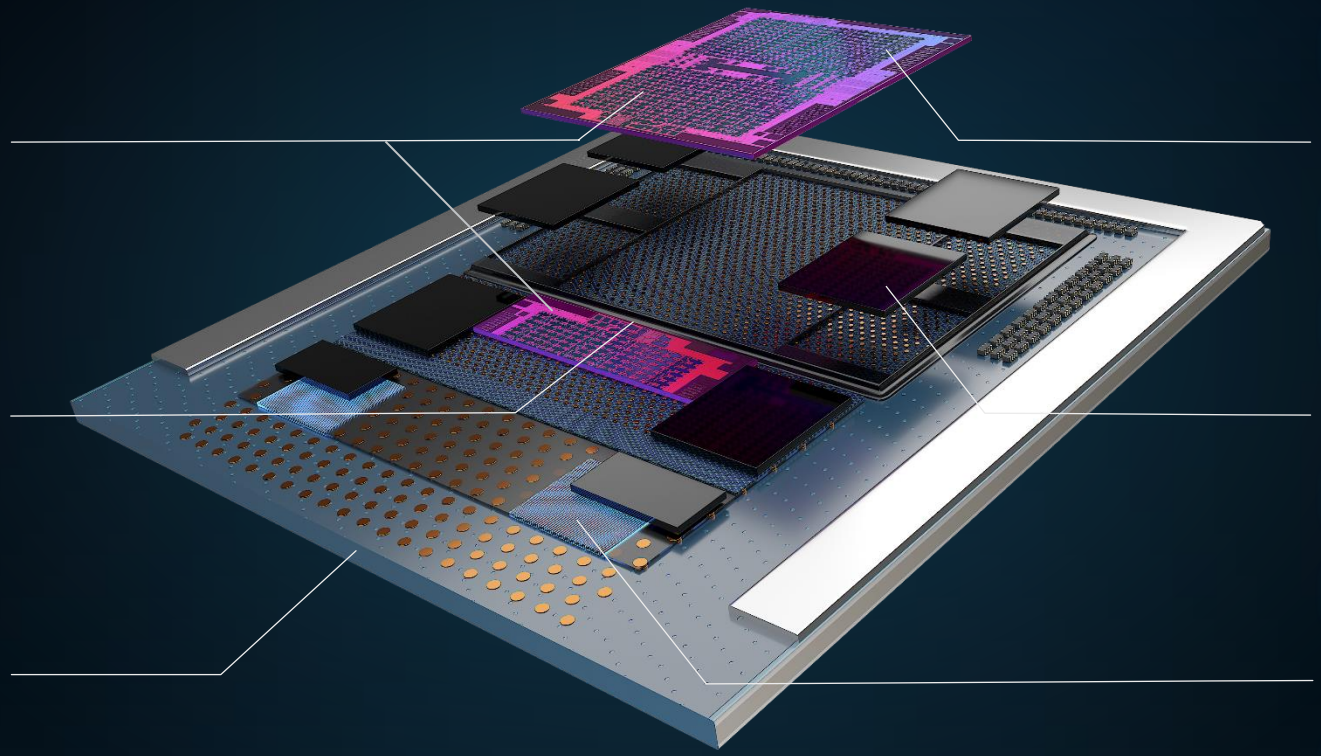

2nd GEN MATRIX CORES FOR HPC & AI

EIGHT STACKS OF HBM2E

2.5D ELEVATED FANOUT BRIDGE (EFB)

AMD INSTINCT™ **MI200 OAM SERIES**

SHATTERING PERFORMANCE BARRIERS IN HPC & AI

PEAK PERFORMANCE	A100	MI200*	INSTINCT ADVANTAGE
FP64 VECTOR	9.7 TF	47.9 TF	**4.9X**
FP32 VECTOR	19.5 TF	47.9 TF	**2.5X**
FP64 MATRIX	19.5 TF	95.7 TF	**4.9X**
FP32 MATRIX	N/A	95.7 TF	N/A
FP16, BF16 MATRIX	312 TF	383 TF	**1.2X**
MEMORY SIZE	80 GB	128 GB	**1.6X**
MEMORY BANDWIDTH	2.0 TB/s	3.2 TB/s	**1.6X**

*MI250X. See endnotes MI200-01, MI200-07.
Note: the A100 TF32 data format is not IEEE FP32 compliant, so not included in this comparison.

DELIVERING PERFORMANCE FOR HPC

FASTEST HPC APPLICATION PERFORMANCE ACROSS A RANGE OF DOMAINS

HPC BENCHMARKS

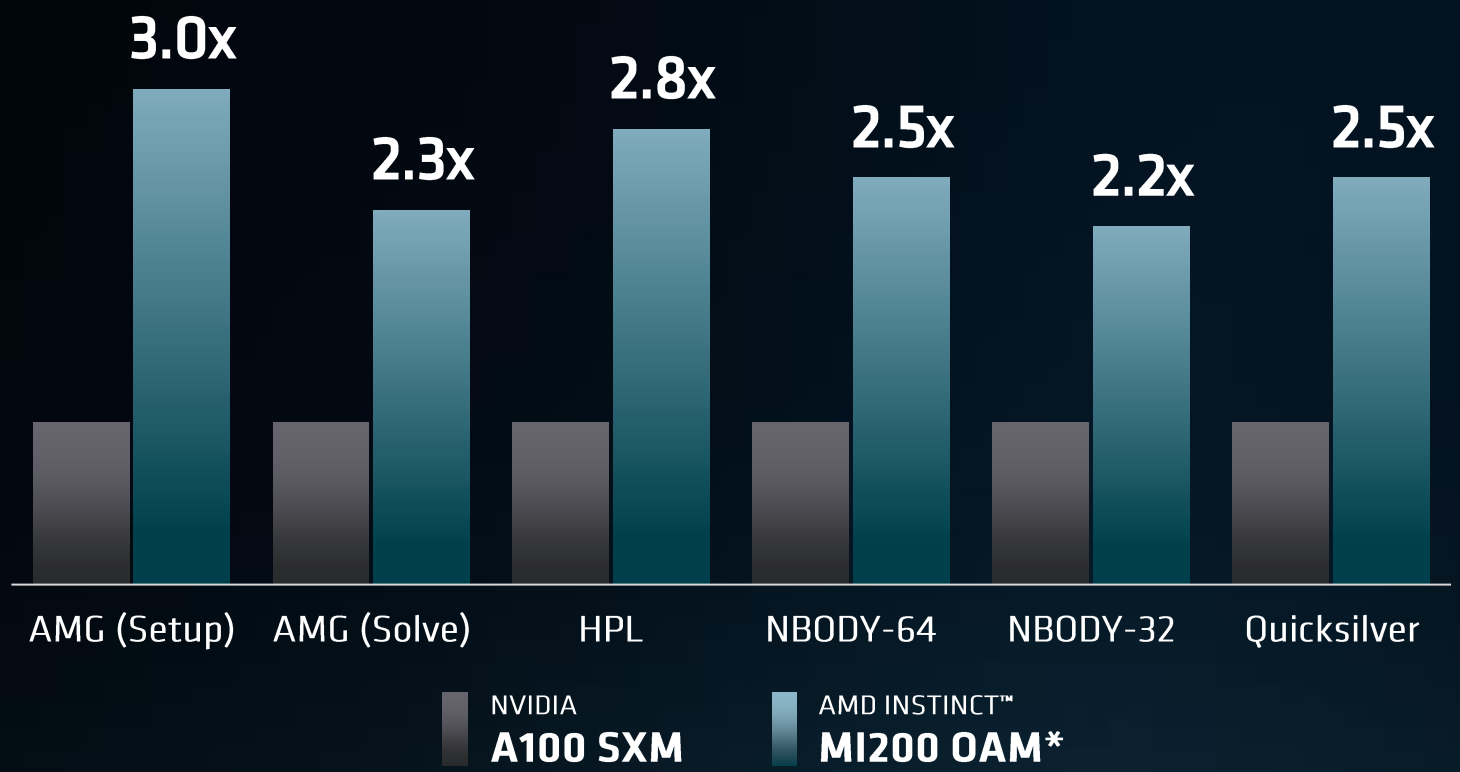

HPC APPLICATIONS

APPLICATION	MI200* ADVANTAGE OVER A100
OpenMM	2.4X
LAMMPS	2.2X
HACC	1.9X
LSMS	1.6X
MILC	1.4X

*MI250X. Results may vary. See endnotes MI200-15A, MI200-16A, MI200-17A, MI200-19A, MI200-20A, MI200-21A, MI200-22A, MI200-23A, MI200-24A, MI200-25A, MI200-26A

LEADERSHIP PERFORMANCE IN SCIENCE

SOLVING GLOBAL CHALLENGES, ACCELERATING TIME TO DISCOVERY

8.9m
ATOM STEPS/S

NVIDIA
A100

2.2X
FASTER SIMULATION THROUGHPUT
LAMMPS ReaxFF

19.5m
ATOM STEPS/S

AMD INSTINCT™
MI200*

Result may vary. *MI250X. See endnote MI200-15A

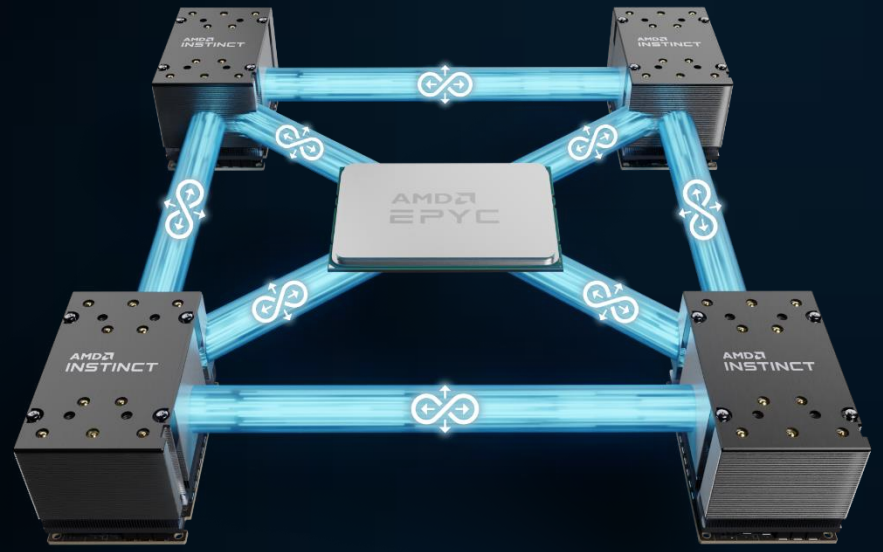

3rd GEN AMD INFINITY ARCHITECTURE

ENABLING UNIFIED COMPUTE AT EXASCALE

REMOVING BARRIERS	PROGRAMMER PRODUCTIVITY	ACCESSIBLE PERFORMANCE
Exceptional System Bandwidth & Performance	CPU & GPU Memory Coherency	Designed to make it easier to accelerate legacy codes on GPUs

CONVERGENCE OF HPC & AI

ENABLING ENHANCED SCIENTIFIC DISCOVERY

“GENOA”

DESIGNED AS THE HIGHEST PERFORMANCE PROCESSOR FOR GENERAL PURPOSE COMPUTING

- UP TO 96 “ZEN 4” CORES IN 5NM
- DDR5
- PCIe® 5.0 | CXL
  WITH BREAKTHROUGH MEMORY EXPANSION
- ENHANCED SECURITY FEATURES

CHALLENGING DATA CENTER WORKLOADS

DEMAND ACCELERATED COMPUTING INNOVATION

GENERAL PURPOSE
COMPUTING

TECHNICAL
COMPUTING

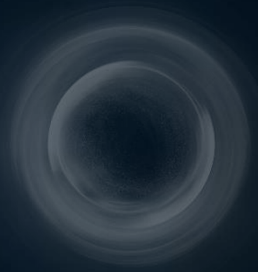

ACCELERATED
COMPUTING

CLOUD NATIVE
COMPUTING

AMD

DESIGNED FOR CLOUD NATIVE COMPUTING LEADERSHIP

"ZEN 4c" IN 5NM

- "ZEN 4" CORE OPTIMIZED FOR SCALE-OUT PERFORMANCE
- SIGNIFICANTLY IMPROVED POWER EFFICIENCY
- DENSITY-OPTIMIZED CACHE HIERARCHY

“BERGAMO”

DESIGNED AS THE HIGHEST PERFORMANCE PROCESSOR FOR CLOUD NATIVE COMPUTING

- UP TO 128 “ZEN 4c” CORES IN 5NM
- SAME “ZEN 4” ISA
- BREAKTHOUGH PERFORMANCE AND POWER EFFICIENCY
- SAME SOCKET AND PLATFORM AS “GENOA”

AMD EPYC™ CPU LEADERSHIP ROADMAP

ROADMAP SUBJECT TO CHANGE

LEADING EVERYWHERE IN THE DATA CENTER

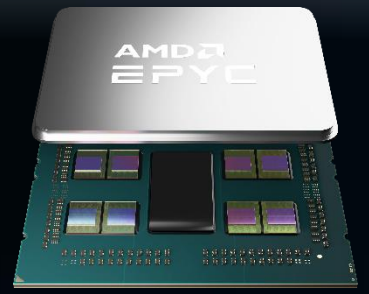

"MILAN-X"

TECHNICAL COMPUTING

Highest
Performance/Core

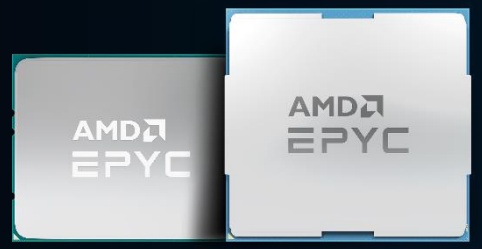

"MILAN" ⊕ "GENOA"

GENERAL PURPOSE COMPUTING

Highest
Performance Across Workloads

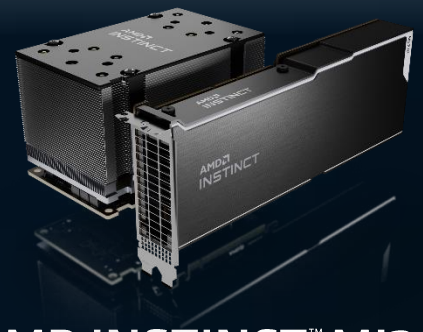

AMD INSTINCT™ MI200

ACCELERATED COMPUTING

Highest
Performance HPC GPU

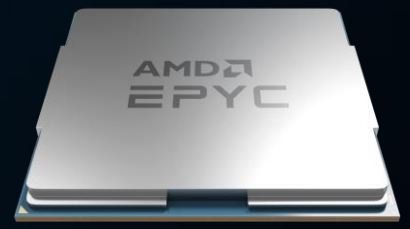

"BERGAMO"

CLOUD-NATIVE COMPUTING

Highest
Thread Density

WITH UNMATCHED SYSTEM ARCHITECTURE TO POWER THE EXASCALE ERA FROM SUPERCOMPUTERS TO CLOUD

See endnote MI200-01, MLN-016B, MLNX-026

AMD

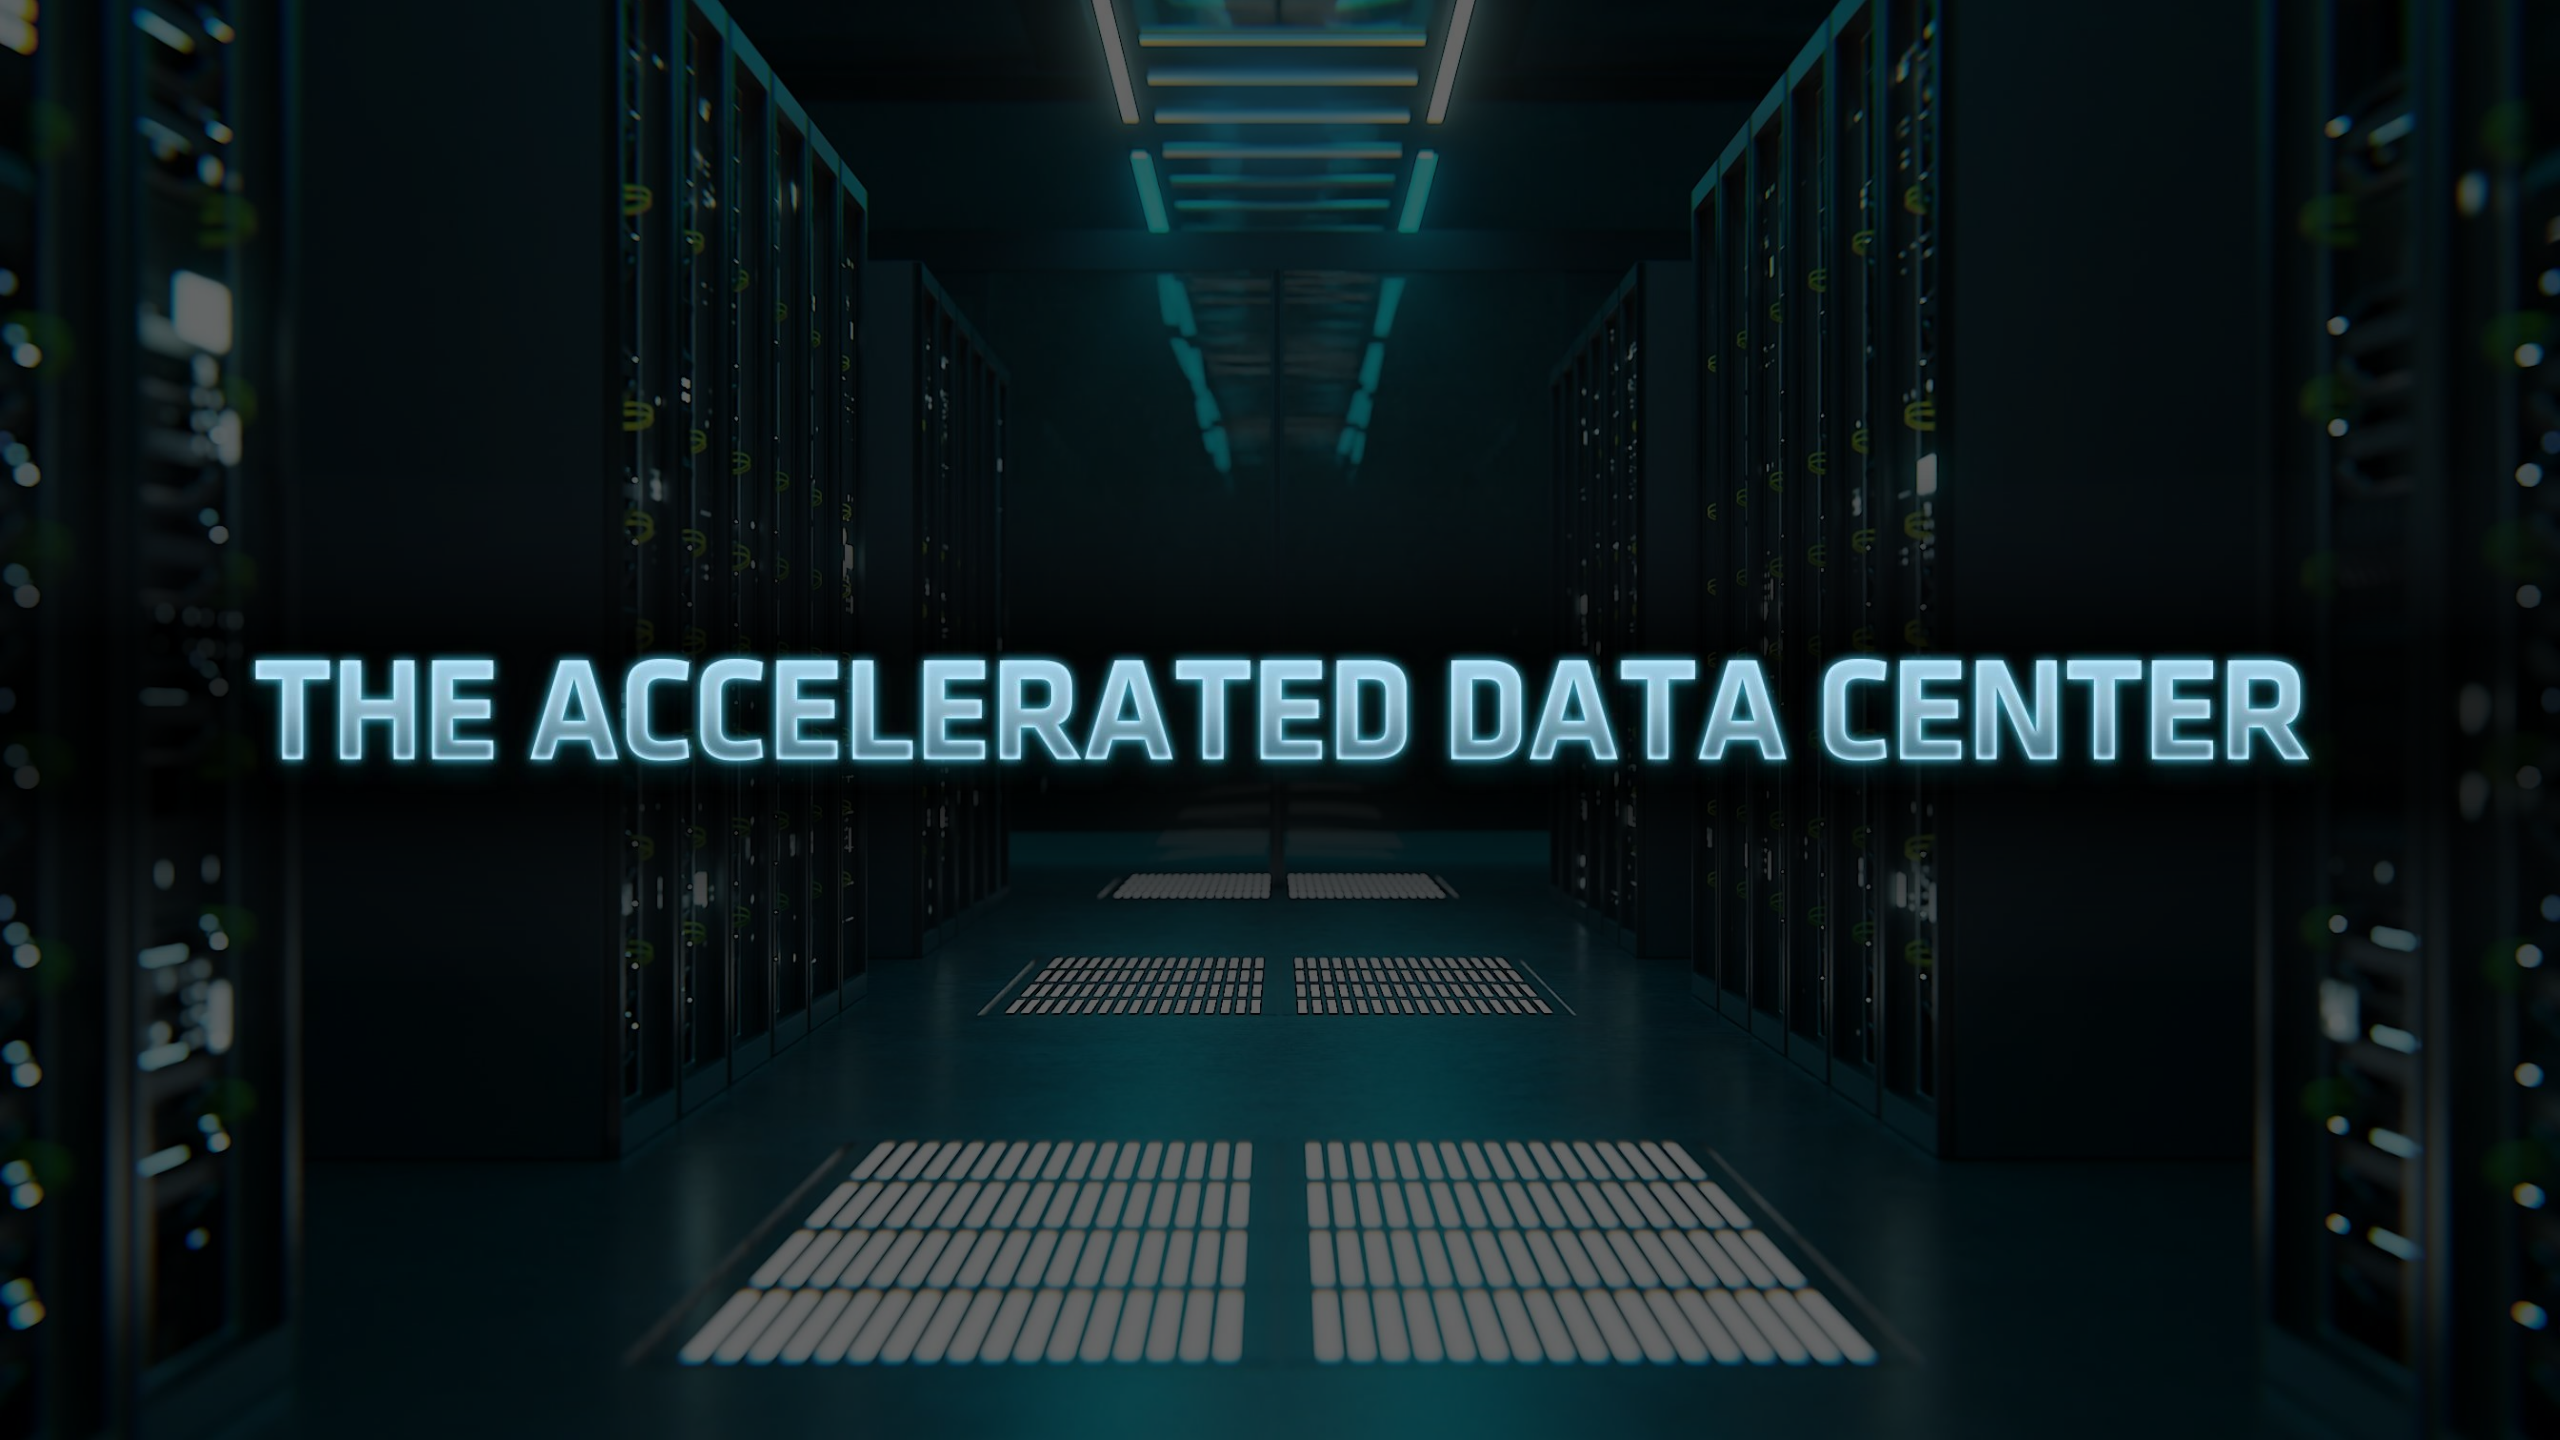
THE ACCELERATED DATA CENTER


AMD

Endnotes

MI200-01 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.

MI200-02 Calculations conducted by AMD Performance Labs as of Sep 15, 2021, for the AMD Instinct™ MI250X accelerator (128GB HBM2e OAM module) at 1,700 MHz peak boost engine clock resulted in 95.7 TFLOPS peak double precision matrix (FP64 Matrix) theoretical, floating-point performance. Published results on the NVidia Ampere A100 (80GB) GPU accelerator resulted in 19.5 TFLOPS peak double precision (FP64 Tensor Core) theoretical, floating-point performance. Results found at https://www.nvidia.com/content/dam/en-zz/Solutions/Data-Center/nvidia-ampere-architecture-whitepaper.pdf, page 15, Table 1.

MI200-07 Calculations conducted by AMD Performance Labs as of Sep 21, 2021, for the AMD Instinct™ MI250X and MI250 (128GB HBM2e) OAM accelerators designed with AMD CDNA™ 2 6nm FinFet process technology at 1,600 MHz peak memory clock resulted in 3.2768 TFLOPS peak theoretical memory bandwidth performance. MI250/MI250X memory bus interface is 4,096 bits times 2 die and memory data rate is 3.20 Gbps for total memory bandwidth of 3.2768 TB/s ((3.20 Gbps*(4,096 bits*2))/8). The highest published results on the NVidia Ampere A100 (80GB) SXM GPU accelerator resulted in 2.039 TB/s GPU memory bandwidth performance. https://www.nvidia.com/content/dam/en-zz/Solutions/Data-Center/a100/pdf/nvidia-a100-datasheet-us-nvidia-1758950-r4-web.pdf

MI200-15A Testing Conducted by AMD performance lab as of 10/7/2021, on a single socket Optimized 3rd Gen EPYC CPU powered server, with 4x AMD Instinct™ MI250X OAM (128 GB HBM2e) 560W GPUs with AMD Infinity Fabric™ technology, using LAMMPS ReaxFF/C, patch_2Jul2021 plus AMD optimizations to LAMMPS and Kokkos that are not yet available upstream resulted in a median score of 4x MI250X = 19,482,180.48 ATOM-Time Steps/sVs.Dual AMD EPYC 7742@2.25GHz CPUs with 4x NVIDIA A100 SXM 80GB (400W) using LAMMPS classical molecular dynamics package ReaxFF/C, patch_10Feb2021 resulted in a published score of 8,850,000 (8.85E+06) ATOM-Time Steps/s. https://developer.nvidia.com/hpc-application-performance 19,482,180.48/8,850,000=2.20x (220%) the/1.2x (120%) faster. Container details found at:https://ngc.nvidia.com/catalog/containers/hpc:lammps Information on LAMMPS: https://www.lammps.org/index.html Server manufacturers may vary configurations, yielding different results. Performance may vary based on use of latest drivers and optimizations

Endnotes

MI200-16A Testing Conducted by AMD performance lab as of 10/18/2021, on a single socket Optimized 3rd Gen AMD EPYC™ CPU powered server with 1x AMD Instinct™ MI250X OAM (128 GB HBM2e) 560W GPU with AMD Infinity Fabric™ technology, using HACC, plus AMD optimizations to HACC that are not yet available upstream resulted in a median score of 1x MI250X = 4,400,000 (4.40E+06) Particles/s Vs. Testing Conducted by AMD performance lab as of 10/18/2021, on Nvidia DGX dual socket AMD EPYC 7742@2.25GHz CPU server with 1x NVIDIA A100 SXM 80GB (400W), using HACC resulted in a median score of 1x A100 = 2,290,000 (2.29E+06) Particles/s. Information on HACC: https://asc.llnl.gov/coral-2-benchmarks/gpu-versions-and-other-supplementary-material https://asc.llnl.gov/sites/asc/files/2020-09/coral-hacc-benchmark-summary-v1.7.pdf Server manufacturers may vary configurations, yielding different results. Performance may vary based on use of latest drivers and optimization

MI200-17A Testing conducted by AMD performance lab as of 10/13/2021, on a single socket Optimized 3rd Gen AMD EPYC™ CPU server with 1x AMD Instinct™ MI250X OAM (128 GB HBM2e) 560W GPU with AMD Infinity Fabric™ technology, using LSMS, plus AMD optimizations to LSMS that are yet available upstream resulted in a median score of 1x MI250X = 3,950,000,000 (3.95E+09) Atom Interactions/s Vs. Testing conducted by AMD performance lab as of 9/27/2021, on Nvidia DGX dual socket AMD EPYC 7742@2.25GHz CPU server with 1x NVIDIA A100 SXM 80GB (400W), using LSMS resulted in a median score of 2,440,000,000 (2.44E+09) Atom Interactions/s. Information on LSMS: https://github.com/mstsuite/lsms, Information on GFortran: https://gcc.gnu.org/fortran/, Information on GCC Compiler: https://gcc.gnu.org/ Server manufacturers may vary configurations, yielding different results. Performance may vary based on use of latest drivers and optimizations

MI200-19A Testing Conducted by AMD performance lab as of 10/1/2021, on a single socket Optimized 3rd Gen AMD EPYC™ CPU server with 4x AMD Instinct™ MI250X OAM (128 GB HBM2e) 560W GPUs with AMD Infinity Fabric™ technology running AMG (Set up) FOM, resulting in a median score of 4x MI250X = 16,773,660,000 FOM_Setup / Sec (Setup Phase Time) Vs. Testing Conducted by AMD performance lab as of 10/1/2021, on Nvidia DGX dual socket AMD EPYC 7742@2.25GHz CPU server with 4x NVIDIA A100 SXM 80GB (400W) running AMG (Set up) FOM, resulting in a median score of 4x A100 = 5,507,144,000 FOM_Setup / Sec (Setup Phase Time). Information on AMG_Setup: https://asc.llnl.gov/coral-2-benchmarks , https://asc.llnl.gov/sites/asc/files/2020-09/AMG_Summary_v1_7.pdf , Server manufacturers may vary configurations, yielding different results. Performance may vary based on use of latest drivers and optimizations

MI200-20A Testing Conducted by AMD performance lab as of 10/1/2021, on a single socket Optimized 3rd Gen AMD EPYC™ CPU server, with 4x AMD Instinct™ MI250X OAM (128 GB HBM2e) 560W GPUs with AMD Infinity Fabric™ technology using AMG (Solve) FOM resulting in a median score of 4x MI250X = 73,318,380,000 FOM_Solve / Sec (Solve Phase Time) Vs. Testing Conducted by AMD performance lab as of 10/1/2021, on Nvidia DGX dual socket AMD EPYC 7742@2.25GHz CPU server with 4x NVIDIA A100 SXM 80GB (400W), using AMG (Solve) FOM resulting in a median score of 4x A100 = 31,476,470,000 FOM_Solve / Sec (Solve Phase Time). Information on AMG_Solve: https://asc.llnl.gov/coral-2-benchmarks, https://asc.llnl.gov/sites/asc/files/2020-09/AMG_Summary_v1_7.pdf Server manufacturers may vary configurations, yielding different results. Performance may vary based on use of latest drivers and optimizations

MI200-21A Testing Conducted by AMD performance lab as of 9/22/2021, on a single socket Optimized 3rd Gen AMD EPYC™ CPU server with 1x AMD Instinct™ MI250X OAM (128 GB HBM2e) 560W GPU with AMD Infinity Fabric™ technology using Nvidia Nbody 32 CUDA sample version 11.2.152 converted to HIP plus AMD optimizations to Nbody 32 that are not yet available upstream resulting in a median score of 2.3x MI250X = 31.72 Particles (Body-to-Body) Interactions/s Vs. Testing Conducted by AMD performance lab as of 9/22/2021, on Nvidia DGX dual socket AMD EPYC 7742@2.25GHz CPU server with 1x NVIDIA A100 SXM 80GB (400W) using Nbody 32 sample code version 11.2.152 resulting in a median score of 14.12 Particles (Body-to-Body) Interactions/s. Information on Nbody 32: https://developer.download.nvidia.com/compute/DevZone/C/html_x64/Physically-Based_Simulation.html, https://github.com/AMD-HPC/nbody-nvidia. Server manufacturers may vary configurations, yielding different results. Performance may vary based on use of latest drivers and optimizations

Endnotes

MI200-22A Testing Conducted by AMD performance lab as of 9/22/2021, on a single socket Optimized 3rd Gen AMD EPYC™ CPU server with 1x AMD Instinct™ MI250 X OAM GPU (128GB HBM2e) with AMD Infinity Fabric™ technology, using Nbody 64 CUDA Sample version 11.2.152 converted to HIP. Nvidia Nbody 64 samples code version 11.2.152, plus AMD optimizations to Nbody 64 that are not yet available upstream resulted in a median score of 19.245 Particles (Body-to-Body) Interactions/s. Vs. Testing Conducted by AMD performance lab as of 9/22/2021, on Nvidia DGX dual socket AMD EPYC 7742@2.25GHz CPU server with 1x NVIDIA A100 SXM 80GB (400W) using benchmark Nvidia Nbody 64 sample code version 11.2.152 resulting in a median score of 7.631 Particles (Body-to-Body) Interactions/s. Information on Nbody 64: https://developer.download.nvidia.com/compute/DevZone/C/html_x64/Physically-Based_Simulation.html, https://github.com/AMD-HPC/nbody-nvidia. Server manufacturers may vary configurations, yielding different results. Performance may vary based on use of latest drivers and optimizations.

MI200-23A Testing Conducted by AMD performance lab as of 10/6/2021, on a single socket Optimized 3rd Gen AMD EPYC™ CPU server with 1x AMD Instinct™ MI250X OAM (128 GB HBM2e) 560W GPU with AMD Infinity Fabric™ technology using Quicksilver - LLNL-CODE-684037 converted to HIP, plus AMD optimizations to Quicksilver that are on AMD Github branch resulted in a median score of 214,000,000 Segments/s Vs. Testing Conducted by AMD performance lab as of 9/22/2021, on Nvidia DGX dual socket AMD EPYC 7742@2.25GHz CPU server with 1x NVIDIA A100 SXM 80GB (400W) using Quicksilver - LLNL-CODE-684037 run with CUDA code version 11.2.152 resulted in a median score of 85,500,000 Segments/s. Information on Quicksilver: AMD branch based on LLNL version for this testing: https://github.com/moes1/Quicksilver/tree/AMD-HIP , LLNL version: https://github.com/LLNL/Quicksilver & Quicksilver info sheet: https://hpc.llnl.gov/sites/default/files/Quicksilver_CTS.pdf. Note: A proxy app for the Monte Carlo Transport Code, Mercury. LLNL-CODE-684037. Server manufacturers may vary configurations, yielding different results. Performance may vary based on use of latest drivers and optimizations

MI200-24A Testing Conducted by AMD performance lab as of 10/12/2021, on a single socket Optimized 3rd Gen AMD EPYC™ CPU server with 1x AMD Instinct™ MI250X OAM (128 GB HBM2e) 560W GPU with AMD Infinity Fabric™ technology using benchmark OpenMM_amoebagk v7.6.0, (converted to HIP) and run at double precision (8 simulations*10,000 steps) plus AMD optimizations to OpenMM_amoebagk that are not yet upstream resulted in a median score of 387.0 seconds or 223.2558 NS/Day Vs. Nvidia DGX dual socket AMD EPYC 7742@2.25GHz CPU server with 1x NVIDIA A100 SXM 80GB (400W) using benchmark OpenMM_amoebagk v7.6.0, run at double precision (8 simulations*10,000 steps) with CUDA code version 11.4 resulted in a median score of 921.0 seconds or 93.8111 NS/Day. Information on OpenMM: https://openmm.org/ Server manufacturers may vary configurations, yielding different results. Performance may vary based on use of latest drivers and optimizations

MI200-25A Testing Conducted by AMD performance lab as of 9/30/2021, on a single socket Optimized 3rd Gen AMD EPYC™ CPU server with 1x AMD Instinct™ MI250X OAM (128 GB HBM2e) 560W GPUs with AMD Infinity Fabric™ technology using MILC benchmark version 7.8.1 developer version MILC_QCD on Github, Apex Medium test module, plus AMD optimizations to MILC that are not yet available upstream resulted in a median score 1,604.567 Total Time (Seconds). Vs. Dual AMD EPYC 7742@2.25GHz CPUs with 1x NVIDIA A100 SXM 80GB (400W) using MILC benchmark version develop_c30ed15e (quda0.8-patch4Oct2017), Apex Medium test module, resulted in a published score of 2,262 Total Time (Seconds). https://developer.nvidia.com/hpc-application-performance Nvidia MILC Container details found at: https://ngc.nvidia.com/catalog/containers/hpc:milc Information on MILC: https://web.physics.utah.edu/~detar/milc/ MILC Manual Server manufacturers may vary configurations, yielding different results. Performance may vary based on use of latest drivers and optimizations

Endnotes

MI200-26A Testing Conducted by AMD performance lab as of 10/14/2021, on a single socket Optimized 3rd Gen AMD EPYC™ CPU server, with 1x AMD Instinct™ MI250X OAM (128 GB HBM2e) 560W GPU with AMD Infinity Fabric™ technology using benchmark HPL v2.3, plus AMD optimizations to HPL that are not yet upstream resulted in a median score of 42.26 TFLOPS Vs. Nvidia DGX dual socket AMD EPYC 7742@2.25GHz CPU server with 1x NVIDIA A100 SXM 80GB (400W) using benchmark HPL Nvidia container image 21.4-HPL resulting in a median score of 15.33 TFLOPS. Information on HPL: https://www.netlib.org/benchmark/hpl/ Nvidia HPL Container Detail: https://ngc.nvidia.com/catalog/containers/nvidia:hpc-benchmarks Server manufacturers may vary configurations, yielding different results. Performance may vary based on use of latest drivers and optimizations

MI200-31 As of October 20th, 2021, the AMD Instinct™ MI200 series accelerators are the “Most advanced server accelerators (GPUs) for data center,” defined as the only server accelerators to use the advanced 6nm manufacturing technology on a server. AMD on 6nm for AMD Instinct MI200 series server accelerators. Nvidia on 7nm for Nvidia Ampere A100 GPU. https://developer.nvidia.com/blog/nvidia-ampere-architecture-in-depth/ MI200-31

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. SPEC®, SPECrate® and SPEC CPU® are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information.

MLN-075A Altair™ Radioss™ comparison based on AMD internal testing as of 09/27/2021 measuring the time to run the neon, t10m, and venbatt test case simulations using a server with 2x AMD EPYC 75F3 versus 2x Intel Xeon Platinum 8362. Neon crash impact is the max result test case. Results may vary.

MLN-080B ANSYS® CFX® 2021.1 comparison based on AMD internal testing as of 09/27/2021 measuring the average time to run the Release 14.0 test case simulations (converted to jobs/day - higher is better) using a server with 2x AMD EPYC 75F3 utilizing 1TB (16x 64 GB DDR4-3200) versus 2x Intel Xeon Platinum 8380 utilizing 1TB (16x 64 GB DDR4-3200). Results may vary.

MLN-130A ANSYS® Mechanical® 2021 R2 comparison based on AMD internal testing as of 09/27/2021 measuring the average of all Release 2019 R2 test case simulations using a server with 2x AMD EPYC 75F3 versus 2x Intel Xeon Platinum 8380. Steady state thermal analysis of a power supply module 5.3M (cg1) is max result. Results may vary.

MLNX-001R: EDA RTL Simulation comparison based on AMD internal testing completed on 9/20/2021 measuring the average time to complete a test case simulation. comparing: 1x 16C 3rd Gen EPYC CPU with AMD 3D V-Cache Technology versus 1x 16C AMD EPYC™ 73F3 on the same AMD “Daytona” reference platform. Results may vary based on factors including silicon version, hardware and software configuration and driver versions.

MLNX-026 Estimated SPECrate®2017_fp_base comparison based on AMD internal testing and best performing systems published at www.spec.org as of 10/28/2021. Configurations: 2x 32C AMD EPYC CPU with AMD V-Cache Technology versus 2x 32C Intel Xeon Platinum 8362 for an estimated 1.18x the performance/Core. The AMD EPYC CPU score is estimated because SPECrate®2017_fp_base was run on pre-production hardware. Customer systems, planned for 1H’22, are expected to be similar. SPEC®, SPEC CPU®, and SPECrate® are registered trademarks of the Standard Performance Evaluation Corporation. see www.spec.org for more information.