

AMD
INSTINCT



SEPTEMBER 2021 HYPERION USER FORUM

BROCK TAYLOR
Global HPC Solutions Director

HPC WORKLOADS ARE DIVERSIFYING



SCIENTIFIC
RESEARCH



ARTIFICIAL
INTELLIGENCE




MODELING AND
SIMULATION



BIG DATA
ANALYTICS

AMD EPYC™ & AMD INSTINCT™ PROCESSORS ADDRESS THE NEEDS
OF USERS ACROSS MULTIPLE HPC WORKLOADS AND VERTICALS



AMD
EPYC

THE BEST GETS BETTER

200+ WORLD RECORDS AND COUNTING

1001
0101
1010

DATABASES & ANALYTICS

- 30 Relational
- 31 Big Data



Enterprise

- 6 ERP Business Apps
- 46 Java® Based Performance
- 26 Energy Efficiency



HCI/SDI/CLOUD

- 13 Cloud and Virtualization
- 16 Integer Performance




HPC

- 59 High Performance Computing Apps
- 15 Floating Point Performance
- 12 Floating Point Energy Efficiency


REAL WORLD RESULTS WITH AMD




CRASH
SIMULATIONS



COMPUTATIONAL
FLUID DYMANICS




EXPLICIT FINITE
ELEMENT
ANALYSIS



WEATHER
FORECASTING

Open Source



FLUID
DYNAMICS




High Performance → Accelerating Time to Results

AMD EPYC™ Tech Docs and White Papers

www.amd.com/en/processors/server-tech-docs/search

EXAMPLE FROM THE LIBRARY




ANSYS Mechanical “shows exceptional performance” using AMD Optimizing CPU Libraries

AOCL BLIS speeds up ANSYS® Mechanical™

- Measured on 2P server with EPYC 73F3 processor server
- Version 2021 R2 vs 2021 R1
- 1.26x geomean speedup
- Gains as high as 2.12x

AOCL BLIS is a numerical library with optimizations for AMD EPYC™ processors



Source: <https://www.amd.com/system/files/documents/amd-epyc-7003-pb-hpc-ansys-mechanical-aocl-optimized-performance.pdf>

PACKAGING OPTIMIZED APPLICATIONS FOR AMD EPYC™

AMD Toolchain with SPACK

1. AMD Optimizing C/C++ Compiler (AOCC)
2. AMD Optimizing CPU Libraries (AOCL)

Micro Benchmarks/Synthetic

1. HPCG
2. HPL
3. STREAM

SPACK HPC Applications


1. CloverLeaf
2. CP2K
3. GROMACS
4. LAMMPS
5. NAMD
6. Open MPI
7. OpenFOAM
8. WRF

AMD provides one-click Spack recipes for common applications built from source

Spack is an open-source project for packaging scientific applications

- Simplifies the build process
- Abstracts dependencies

More information: <https://developer.amd.com/spack>



AMD INSTINCT™ MI100 GPU

POWERED BY THE ALL-NEW AMD CDNA™ ARCHITECTURE

UP TO
11.5TF First Datacenter GPU to
Surpass 10TF FP64 Barrier

UP TO
4x Faster GPU-to-GPU Performance with
Infinity Architecture vs PCIe® 4.0


UP TO
1.2 TB/s Bandwidth with
32GB HBM2

2X Higher Compute Density
with New Core Design

ALL-NEW Matrix Core Technology
for HPC and AI workloads

NEARLY
7X On AI Workloads with
Mixed Precision and FP16

AMD INFINITY HUB FOR AMD INSTINCT™ ACCELERATORS



<https://www.amd.com/InfinityHub>

Containerized HPC apps and ML frameworks that are ported, optimized and supported on ROCm™ Software Ecosystem

Additional Links to:

- ROCm™ Learning Center
- ROCm™ Docs
- ROCm™ Community Forum

THE AMD INSTINCT™ EDUCATION & RESEARCH (AIER) INITIATIVE



Access to AMD Instinct™ technologies

- ❑ Code development
- ❑ Code validation & certification
- ❑ Benchmark validation



ROCm™ Software and Community

- ❑ Compilers, libraries & management tools
- ❑ Workload containers
- ❑ Community support forums



ROCm™ Learning Center

- ❑ Support DOCs
- ❑ Training and tutorial resources
- ❑ Hackathons

AMD.com/AIER

COMPUTATIONAL SCIENCE STARTS HERE

Be ready for the next generation of supercomputers with resources and initiatives for educators and researchers.

AMD
INSTINCT

AMD
ROCm

DEVELOPER TOOLS FOR AMD

AMD
EPYC

AMD Optimized Tools:

Compilers
Profilers
Debuggers
Math Libraries
ML/AI Libraries
Management Tools
Package Management

AMD
INSTINCT

Including support for additional open-source tools

Developer.AMD.com
AMD.com/ROCm

CASE STUDY: NORTHERN DATA



Northern Data takes HPC to a new level of affordability with AMD technology

“We managed to lower the power consumption about 30 to 40 percent for a comparable workload versus other cloud platforms.”

“The reduction in carbon footprint is especially important to us.”

Northern Data AG is a provider of HPC infrastructure using AMD EPYC™ and AMD Instinct™ Processors

- Datacenters powered by >90% green energy
- Providing cost-effective ML Training services

*Michel Boutouil
General Manager of
Northern Data Software GmbH*

Source <https://www.amd.com/system/files/documents/northern-data-takes-hpc-to-a-new-level-of-cost-and-energy-efficiency-with-amd-technology.pdf>

AMD HPC LEADERSHIP

- Delivering strong HPC & AI roadmap
- World's fastest server processor*
* EPYC 7763
- World's fastest HPC accelerator**
** Instinct MI100
- Broad industry adoption



SEE ENDNOTES: MLN-016B, MI100-03





ENDNOTES

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-057B: Based on SPECrate®2017_int_base on 07/06/2021, a server powered by two 8c AMD EPYC 72F3 CPU has scored 192, <http://spec.org/cpu2017/results/res2021q3/cpu2017-20210621-27508.html> with a per core score of 12.0 which is a higher per core integer base performance score than any currently posted in any SPEC.org publication. SPEC®, SPECrate® and SPEC CPU® are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information.

CDNA-04 - Calculations by AMD Performance Labs as of Oct 5th, 2020 for the AMD Instinct™ MI100 accelerator designed with AMD CDNA 7nm FinFET process technology at 1,200 MHz peak memory clock resulted in 1.2288 TFLOPS peak theoretical memory bandwidth performance. The results calculated for Radeon Instinct™ MI50 GPU designed with “Vega” 7nm FinFET process technology with 1,000 MHz peak memory clock resulted in 1.024 TFLOPS peak theoretical memory bandwidth performance. CDNA-04

MI100-03 - 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 double precision (FP64), 46.1 TFLOPS peak single precision matrix (FP32), 23.1 TFLOPS peak single precision (FP32), 184.6 TFLOPS peak half precision (FP16) peak theoretical, floating-point performance. Published results on the NVidia Ampere A100 (40GB) GPU accelerator resulted in 9.7 TFLOPS peak double precision (FP64), 19.5 TFLOPS peak single precision (FP32), 78 TFLOPS peak half precision (FP16) theoretical, floating-point performance. Server manufacturers may vary configuration offerings yielding different results. MI100-03

MI100-04 - Calculations performed by AMD Performance Labs as of Sep 18, 2020 for the AMD Instinct™ MI100 accelerator at 1,502 MHz peak boost engine clock resulted in 184.57 TFLOPS peak theoretical half precision (FP16) and 46.14 TFLOPS peak theoretical single precision (FP32 Matrix) floating-point performance. The results calculated for Radeon Instinct™ MI50 GPU at 1,725 MHz peak engine clock resulted in 26.5 TFLOPS peak theoretical half precision (FP16) and 13.25 TFLOPS peak theoretical single precision (FP32 Matrix) floating-point performance. Server manufacturers may vary configuration offerings yielding different results. MI100-04

MI100-06 Calculations as of SEP 18th, 2020. AMD Instinct™ MI100 accelerators support PCIe® Gen4 providing up to 64 GB/s peak theoretical transport data bandwidth from CPU to GPU per card. AMD Instinct™ MI100 accelerators include three Infinity Fabric™ links providing up to 276 GB/s peak theoretical GPU to GPU or Peer-to-Peer (P2P) transport rate bandwidth performance per GPU card. Combined with PCIe Gen4 support, this provides an aggregate GPU card I/O peak bandwidth of up to 340 GB/s. Server manufacturers may vary configuration offerings yielding different results. MI100-06

MI100-09 - AMD Instinct™ MI100 accelerators provide 120 compute units and 7,680 stream cores in a 300W accelerator card. Radeon Instinct™ MI50 accelerators provide 60 compute units (CUs) and 3,840 stream cores in a 300W accelerator card. MI100-09

DISCLAIMER AND ATTRIBUTIONS

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions and typographical errors.

The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes.

AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION.

AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY DIRECT, INDIRECT, SPECIAL OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

© 2021 Advanced Micro Devices, Inc. all rights reserved. AMD, the AMD arrow, AMD EPYC, AMD CDNA, AMD Instinct, AMD RDNA, ROCm, and combinations thereof, are trademarks of Advanced Micro Devices, Inc. Other names are for informational purposes only and may be trademarks of their respective owners.

PCIe® is a registered trademark of PCI-SIG Corporation. OpenFOAM® is a registered trademark of OpenCFD Limited. NAMD was developed by the Theoretical Biophysics Group in the Beckman Institute for Advanced Science and Technology at the University of Illinois at Urbana-Champaign. ANSYS® Mechanical™ is a registered trademark of ANSYS, Inc.

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.