Enabling Portable Directive-Based Programming at Exascale

Nicholas Malaya and Ron Lieberman
AMD Research
WACCPD, SC20

AMD
AMD EPYC | AMD RADEON INSTINCT | ROCm

Agenda

- Motivation
- Planned offerings for directive programming on AMD GPUs
- Demonstrations of capabilities available now
- Discussions of challenges for Exascale

Motivation: A New Frontier for HPC

https://www.olcf.ornl.gov/wp-content/uploads/2019/05/frontier_specsheet.pdf

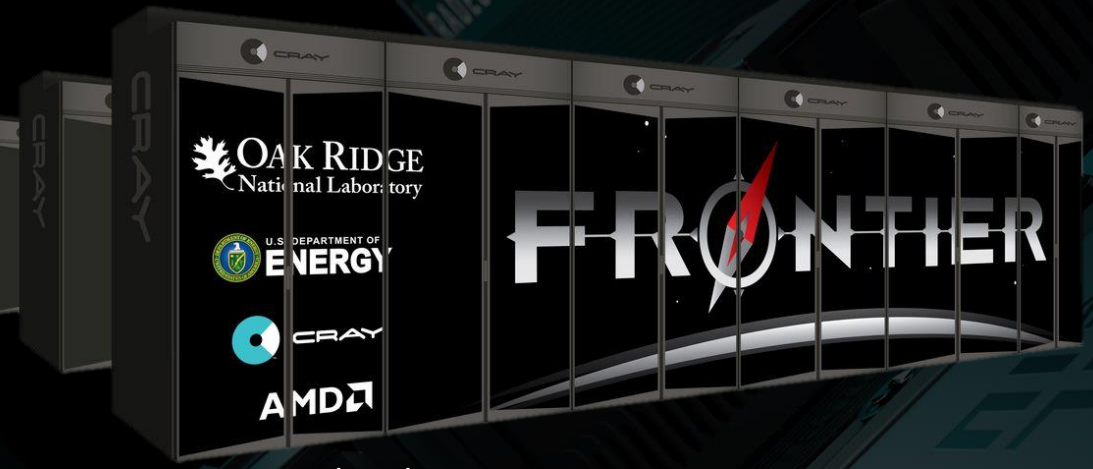

Oak Ridge Frontier Supercomputer
1.5 Exaflops Powering the World's Leading Open Science System

- Custom EPYC™ CPU Optimized for HPC & AI
- HPC-Customized Compute Engines in Radeon Instinct™ GPU
- High Bandwidth, Low Latency CPU & GPU for Coherency
- Open, Portable Software Platform

The Frontier Node at a Glance

- Custom EPYC™ HPC-optimized processor
- Four Radeon Instinct™ accelerators
- Coherent connectivity:
  - Via Infinity Fabric™ interconnect
  - Tightly integrated
  - Unified memory space

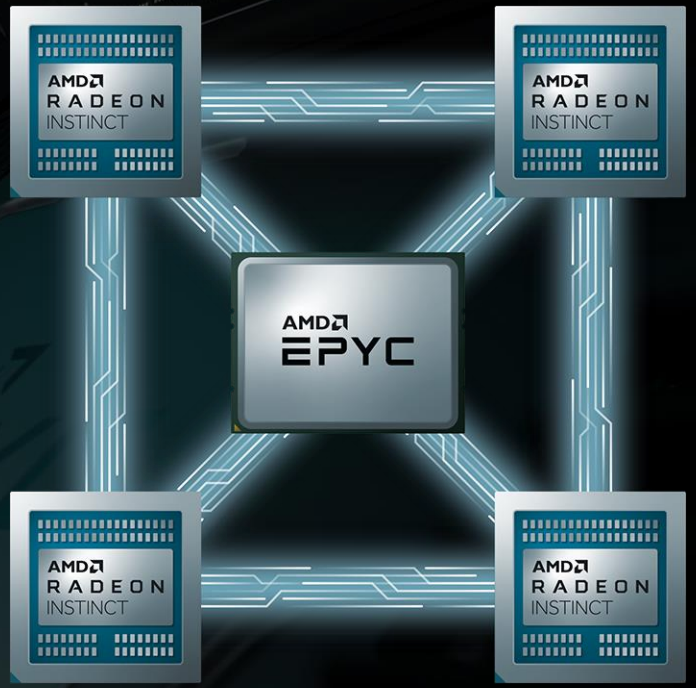

https://www.amd.com/en/products/frontier

Directive Programming on AMD GPUs

OpenMP

- OpenMP support via AOMP
- AOMP: AMD OpenMP Compiler
  - LLVM-based clang driver (LLVM 11) – all the source is open!
  - The last release is AOMP 11.11-1
  - Compiles C/C++ code with OpenMP “target” pragmas
  - Flang - FORTRAN compiler (FORTRAN 2003 standard)
  - Links with libomptarget to produce a binary for offloading to the GPU
- OMPD compliant implementation
  - support for ROC-GDB, Totalview, etc.
- ROCm 3.9 supports AOMP & HIP via single source compiler

What about OpenACC?

- Mentor Graphics
  - Has built in GCC backend supporting the AMDGCN ISA
- GCC
  - OpenACC v2.6 is implemented in GCC and gfortran
  - Mentor released updated compiler May 2020
  - Optimizations and bug fixes ongoing

https://www.mentor.com/embedded-software/sourcery-tools/sourcery-codebench/editions/lite-edition/

What about performance?

- Consider a few representative cases
  - Memory bandwidth bound (STREAM)
  - Abstraction frameworks
  - SPEChpc™ 2021

OpenMP Performance on MI60, ROCm3.9

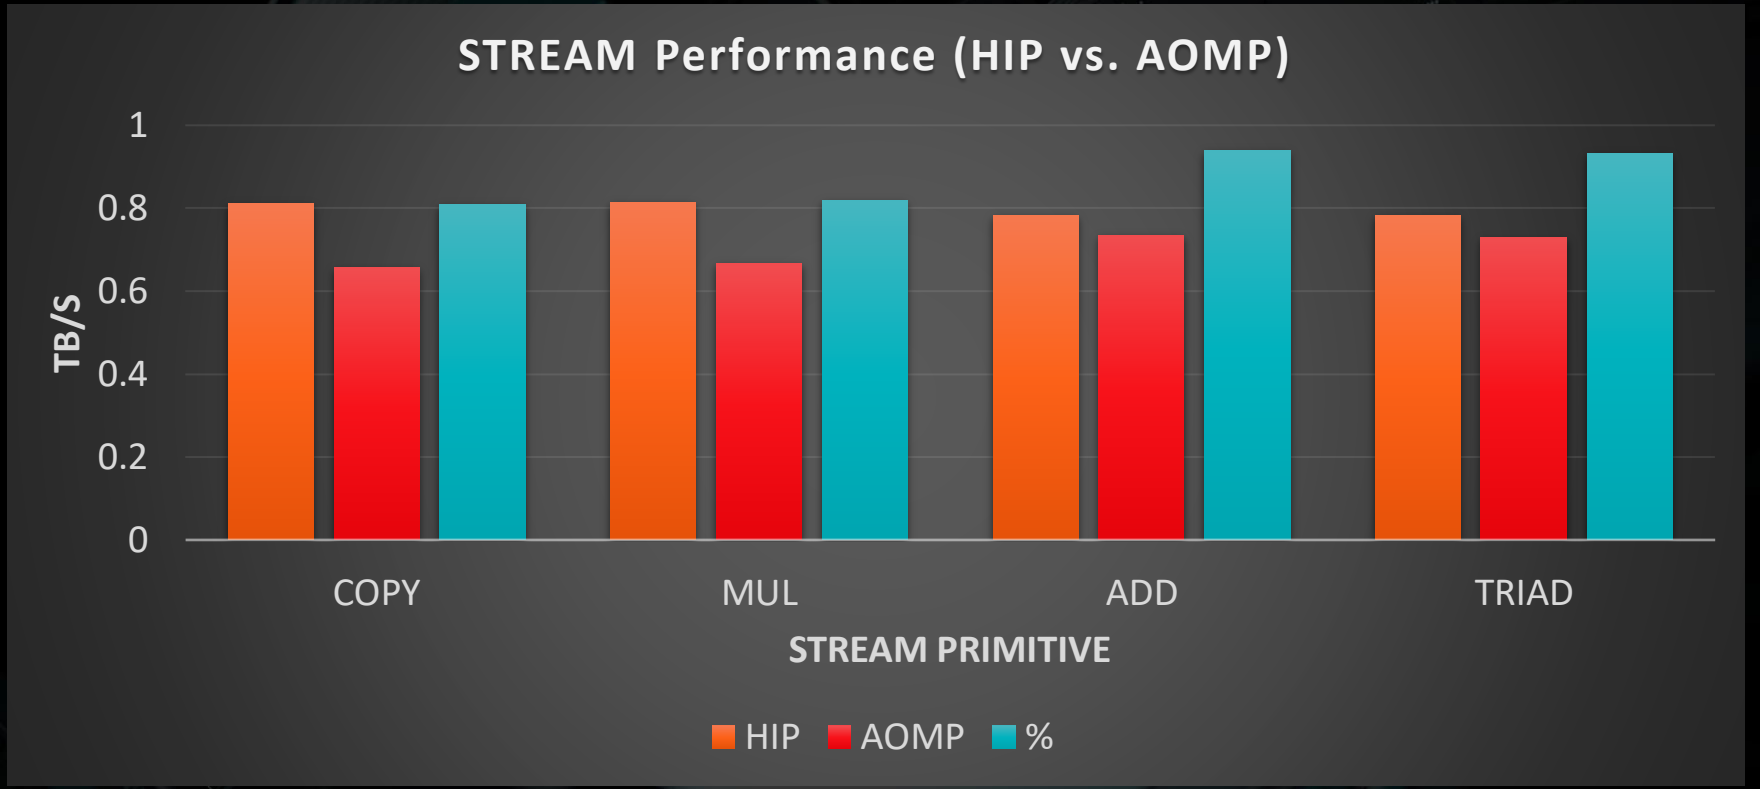

OpenMP Offload: STREAM COPY

```
#ifdef OMP_TARGET_GPU
#pragma omp target teams distribute parallel for simd
#else
#pragma omp parallel for
#endif
  for (int i = 0; i < array_size; i++)
   {
    c[i] = a[i];
    }
```

RAJA Performance on MI60, ROCm3.9

- What about performance via portability layers such as RAJA, Kokkos?

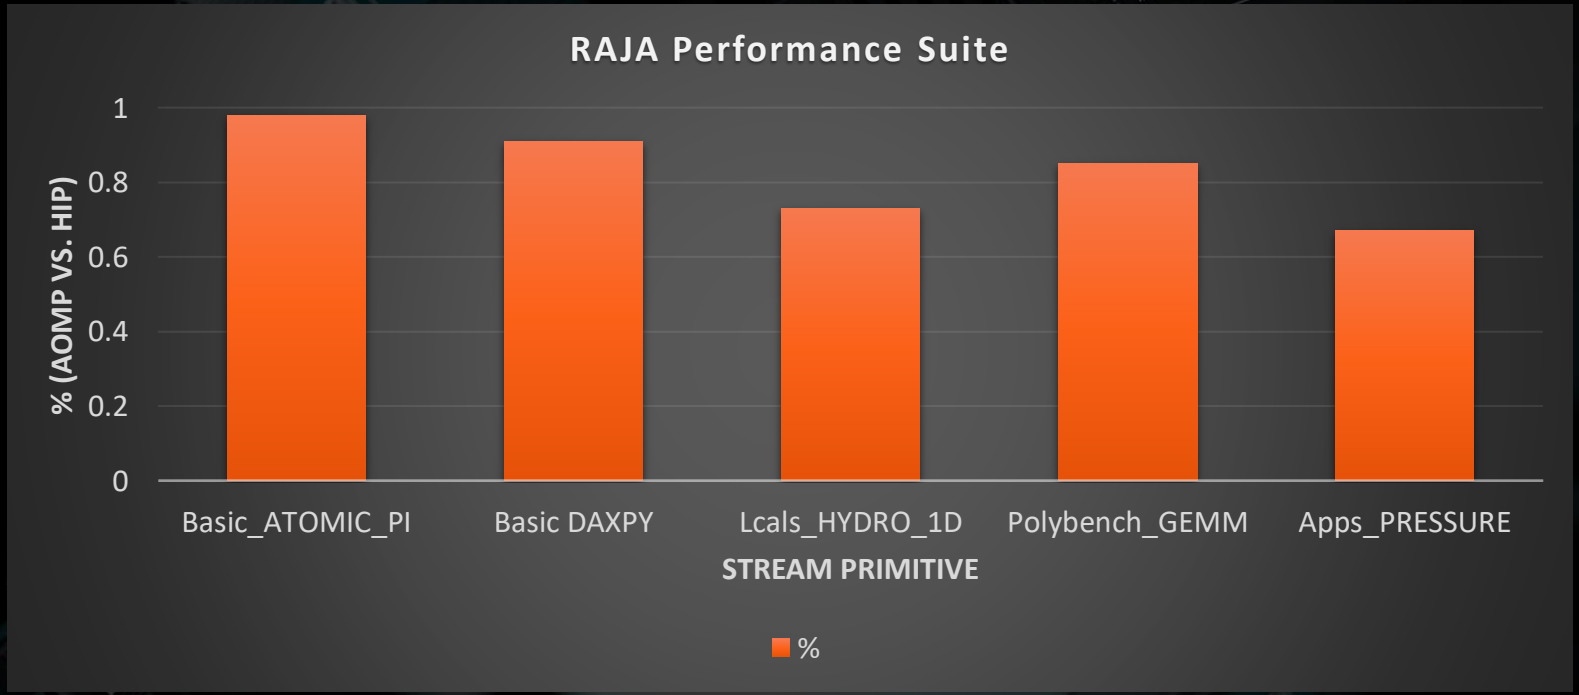

SPEChpc™ 2021

- SPEC HPG® organization working on benchmarks that would be relevant to the Exascale era
- Benchmark was officially announced on Nov 9
- The SPEChpc™ 2021 Benchmark Suites focus on compute intensive parallel performance across one or more nodes
- More information about SPEC HPC® can be obtained from https://www.spec.org

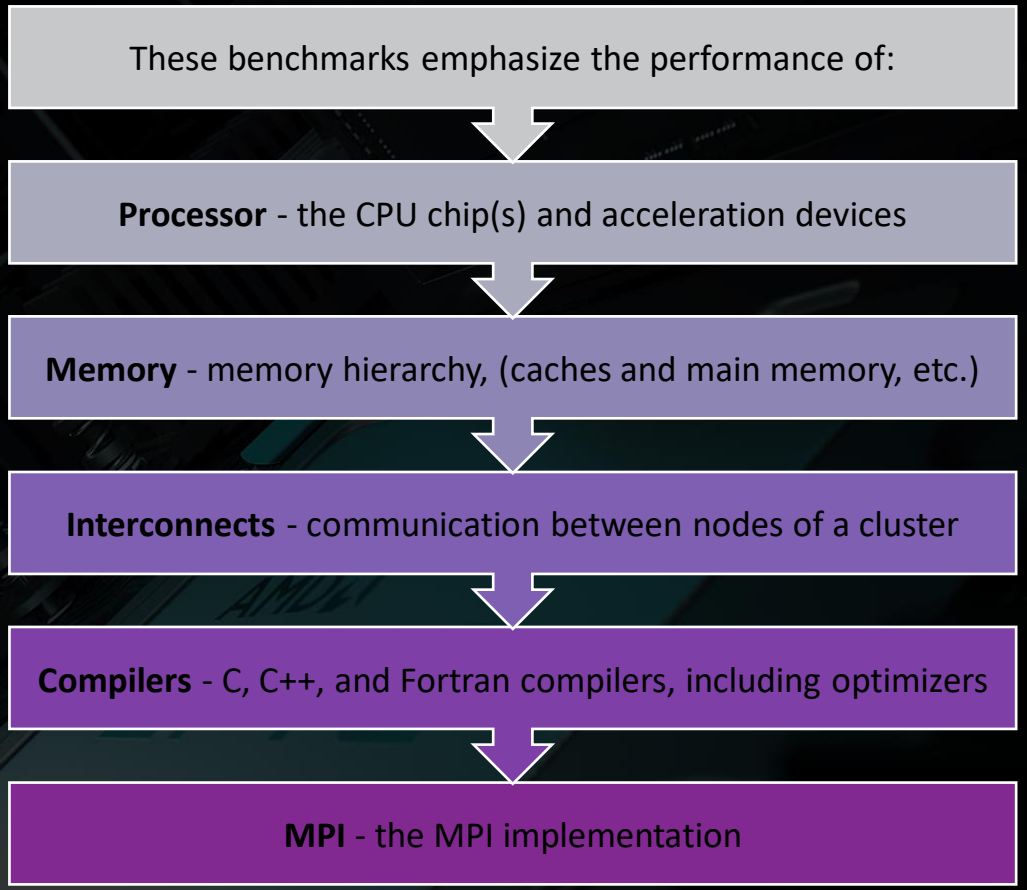

Let's talk Exascale

- Large portion of the FLOPs coming from accelerators
  - Keep work resident on GPU, only move what is needed
  - Map() clause essential: rich APIs for controlling data movement

```
#pragma omp target teams distribute parallel for simd
map(to:a[0,sz]),map(tofrom:c[0,sz])
for (int i = 0; i < sz; i++)
  {
   c[i] = a[i];
   }
```

Let’s talk Exascale

- Large portion of the FLOPs coming from accelerators
- Directives are a powerful mechanism to move more work to the GPU
  - Don’t need to re-write entire subroutines into low-level languages
  - “Premature optimization is the root of all evil”
  - Pass pointers between HIP, BLAS libraries, OMP

Conclusions

- AMD OpenMP compilers enable application work to start *today*
  - Don’t necessarily need AMD hardware to start work
    - MI60s share Instruction Set Architecture (ISA) with future GPUs
  - Portability and standards compliance are good
  - More work to be done to ensure high performance
- Move more work to accelerators
  - Start with directives to quickly move work to an accelerator
  - Think carefully about data placement

Questions?

AMD

Disclaimers 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. Any computer system has risks of security vulnerabilities that cannot be completely prevented or mitigated. 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.

THIS INFORMATION IS PROVIDED 'AS IS." 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 NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY RELIANCE, 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.

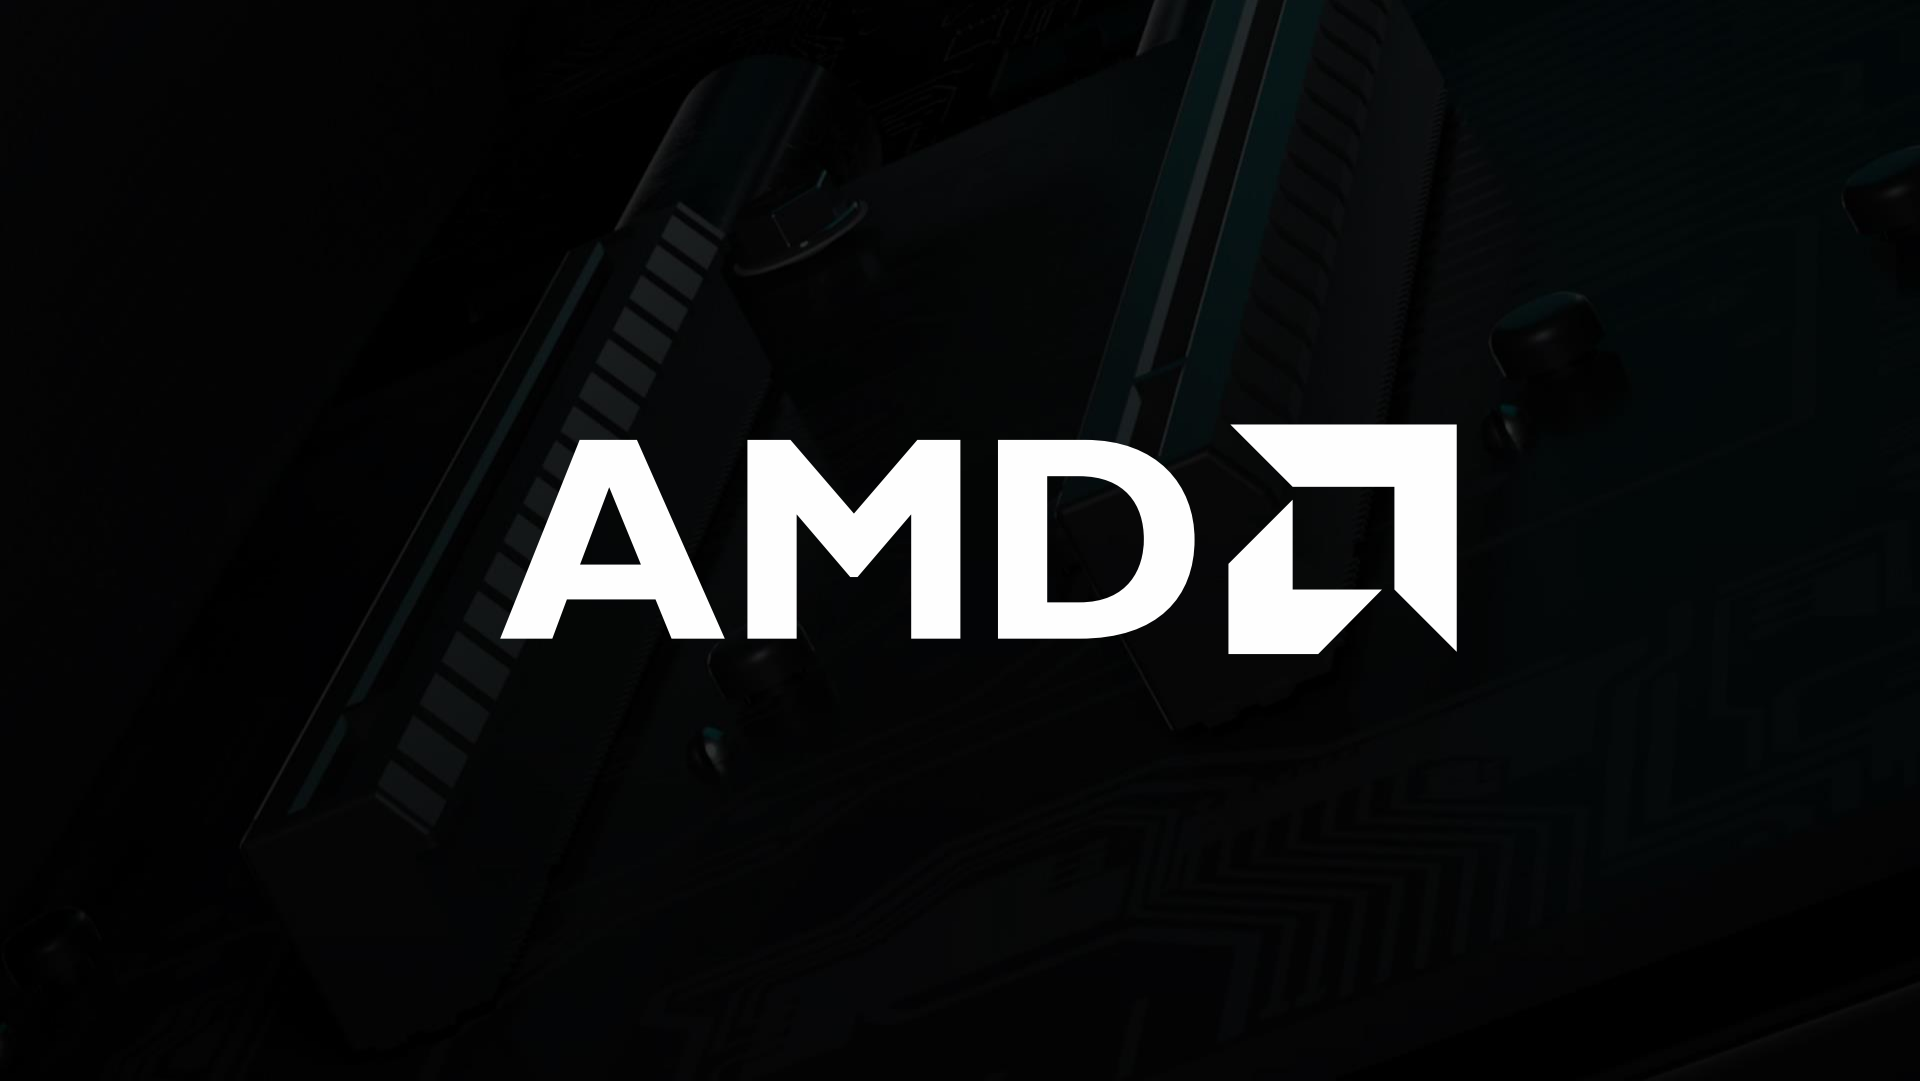
AMD