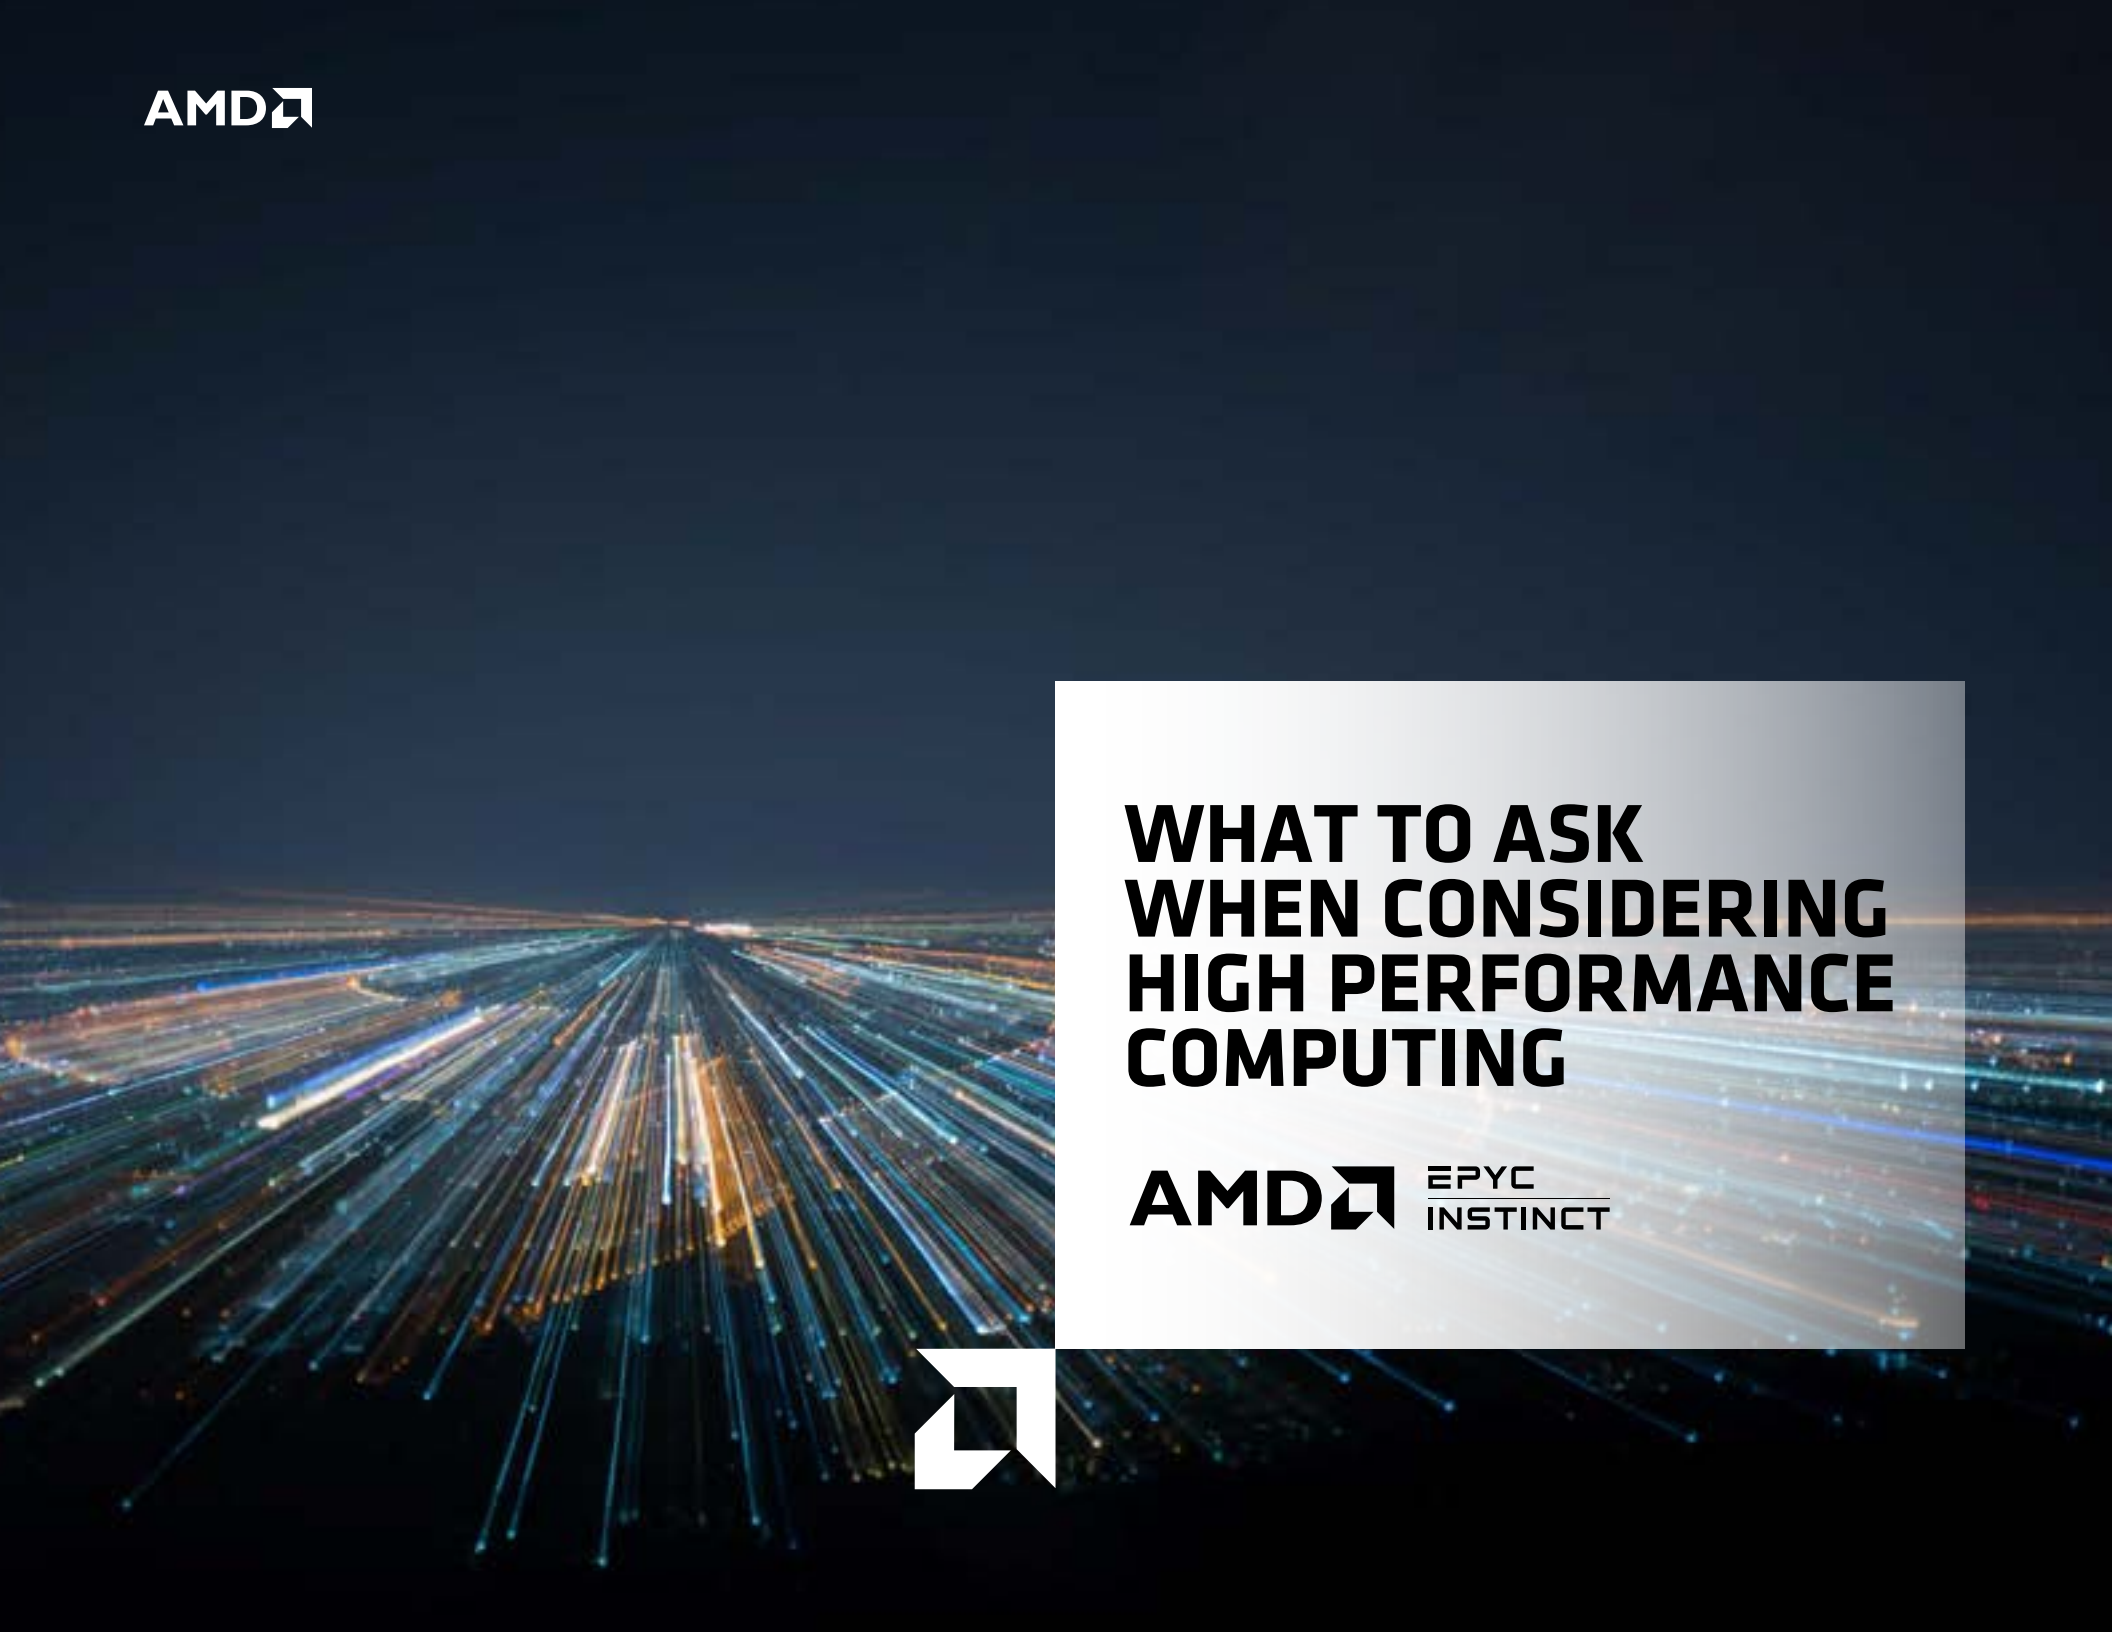
AMD
WHAT TO ASK
WHEN CONSIDERING
HIGH PERFORMANCE
COMPUTING
AMD
EPYC
INSTINCT

# TABLE OF CONTENTS

From manufacturing to healthcare, high performance computing (HPC) is being used to solve the world's most complex problems, as well as everyday challenges. In this eBook, we provide a high-level overview of things to consider when you're looking at HPC for your organization.



# THINGS TO CONSIDER WHEN BUILDING YOUR HPC SYSTEM

When exploring high performance computing (HPC) for your organization, there are five key areas of consideration.

- WORKLOAD
- COMPUTE
- PERFORMANCE TOOLS
- STORAGE
- DEPLOYMENT

# WORKLOAD REQUIREMENTS

Using the right tool for the job is as true in the world of high performance computing as it is on a job site. Ensuring your HPC system is optimized for your workloads helps ensure you have the speed you need to get fast results.

## SYSTEM CHARACTERISTICS

**What are the most important system characteristics for the desired set of workloads? Is your application CPU-bound?**

–

**Does it require a lot of memory or does it spend a lot of time communicating?**

Understanding these characteristics can help you select system components optimized for your workloads.

## DATA

**How does the size and type of data storage fit the simulation needs?**

–

**What provides the data and how does it get to the system?**

## HOW AMD CAN HELP

3rd Gen AMD EPYC™ processors can deliver up to 1.4 times faster performance for HPC workloads.[1] If you're ready for accelerators, combine them with AMD Instinct™ accelerators for additional compute capability.

# COMPUTE POWER

Understanding your workload can help you select processors with the optimal core count, frequency, and memory.

### CORES AND LICENSING

**How many cores should your CPUs have?**

### FREQUENCY

**What frequency do you need?**

### MEMORY AND CACHE

**What is the ideal memory bandwidth and cache size?**

### HOW AMD CAN HELP

2nd and 3rd Gen AMD EPYC™ processors scale from 8 to 64 cores and support PCIe® 4.0 to enable the bandwidth that can help improve efficiency and speed. See below for examples of different HPC market segments, as well as general characteristics and sensitivities of segment applications.

| Segment | Sensitivity | Example Applications | Comments |
|---|---|---|---|
| FEA Explicit | Frequency and Cache | LS-DYNA®, Radioss™, Abaqus™, VPS™ | 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 and Frequency | GROMACS, LAMMPS | Look for CPUs with high core counts and frequency. These applications scale very well with cores. |
| Weather | Memory BW and Cache | WRF, IFS | |
| CFD | | Fluent®, AcuSolve™ | 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. |
| Oil and Gas | | Reveal®, Echos, SAVA | |
| EDA | Frequency and Cache | VCS®, RedHawk™ | This market segment is dominated by low core counts to drive up the frequency and cache per core, helping maximize software investment. |

→ LEARN MORE

# PERFORMANCE TOOLS

Performance tools play an important supporting role in high performance computing. Help ensure your workloads can take advantage of system functionality by using software and tools that are designed for your systems. This can also provide an additional performance boost.

## COMPILERS AND LIBRARIES

**Which compilers, profilers, and/or libraries can you use to optimize your system?**

### HOW AMD CAN HELP

AMD Developer Central offers tools and resources to help optimize applications on AMD hardware, including the AMD Optimizing C/C++ Compiler (AOCC), the AMD Optimizing CPU Libraries (AOCL), and AMD ROCm™ for accelerators.

# STORAGE

Depending on the desired workloads and scale of the system, storage options for HPC comes in different forms.

### SHARED STORAGE

**What kind of shared storage do you need? Parallel file systems, NFS?**

### SCRATCH SPACE

**Do you need local or network scratch storage?**

### BACKUP AND RECOVERY

**How will you back up data and recover, if necessary?**

### HOW AMD CAN HELP

Servers based on 2nd and 3rd Gen AMD EPYC™ SoC can take advantage of the 128 lanes of PCIe® Gen 4 to offer dense storage and I/O configurations.

# DEPLOYMENT

Understanding how you'll power and cool your system will influence how you size a cluster. Also, decide how you'll manage the cluster, whether you'll build your own software, or use commercial or open-source applications. Identify the benefits and which solution best fits your deployment.

### POWER AND COOLING

**Do you have power and cooling to accommodate your HPC system?**

### CLUSTER MANAGEMENT

**What tools will you use to run and maintain the cluster?**

### CLOUD POTENTIAL

**Is HPC in the cloud a viable option?**

### HOW AMD CAN HELP

AMD EPYC™ processors can be used in both air- and liquid-cooled environments, offering great data center flexibility and efficiency. See how your data center could reduce greenhouse gas (GHG) emissions and lower TCO with AMD EPYC™ processors. Try the AMD EPYC™ Bare Metal and Greenhouse Gas Emissions TCO Estimation tool.

–

With AMD Infinity Guard, the advanced security features built into every AMD EPYC™ processor, you can take advantage of HPC in the cloud with confidence.

# AMD EPYC™ PROCESSORS AND AMD INSTINCT™ ACCELERATORS FOR HPC

From high frequencies to dense core counts to large caches, AMD EPYC™ processors offer the flexibility and performance that HPC workloads demand. Combine them with AMD Instinct™ MI200 accelerators, featuring the world's fastest HPC and AI GPUs, and give your system an even bigger leap in performance.[2] Find out how AMD can help you with HPC.

→ LEARN MORE ABOUT AMD EPYC™ PROCESSORS FOR HPC

→ LEARN MORE ABOUT AMD INSTINCT™ ACCELERATORS FOR HPC

# HOW HPC SUCCEEDS WITH AMD

### HPC INFRASTRUCTURE PROVIDER

Northern Data delivered on its promise of more affordable AI and ML and got closer to its goal of climate-neutral computing, with AMD EPYC™ processors and AMD Instinct™ GPUs.

→ Read more

### COMPUTER-AIDED DESIGN AND DEVELOPMENT

Promarin used AMD EPYC™ processors to help transform development of its marine propellers, increasing efficiency and making them more environmentally friendly.

→ Read more

### COVID-19 HPC RESEARCH

The AMD COVID-19 HPC Fund helps research institutions accelerate our understanding of infectious diseases, with the help of AMD EPYC™ processors and AMD Instinct™ GPUs.

→ Read more

### GENERAL PURPOSE, MULTIUSER HPC SYSTEM

The University of Florida is expanding human knowledge using supercomputers powered by AMD EPYC™ processor-based Lenovo servers.

→ Read more

### CAD, CFD, THERMAL AND SHOCK ENGINEERING SERVICES

Ten Tech, LLC increased customer satisfaction by accelerating simulation and modeling to deliver products faster, with AMD EPYC™ processors.

→ Read more

### GENERAL PURPOSE, MULTIUSER HPC SYSTEM

The San Diego Supercomputer Center delivers HPC resources to tens of thousands of academic- and industry-based researchers across the United States, with AMD EPYC™ processors.

→ Read more

### PARTICLE PHYSICS RESEARCH

The Dutch National Institute for Subatomic Physics (Nikhef) developed a deeper understanding of subatomic particles that serve as the building blocks of the universe, with AMD EPYC™ processors

→ Read more

### ANIMATION AND VISUAL EFFECTS STUDIO

Jellyfish Pictures gained a competitive advantage, taking on more projects and expanding its global remote workforce with AMD EPYC™ processors

→ Read more

# RESOURCES

## AMD OPTIMIZING COMPILERS AND LIBRARIES

Learn more about the compilers and libraries AMD offers: AMD Optimizing C/C++ Compiler (AOCC), the AMD Optimizing CPU Libraries (AOCL), and AMD ROCm™ for accelerators.

## AMD HPC SOLUTION BRIEFS

Before making a solution decision, reference AMD Solution Briefs to learn how we work with your applications and workloads.

[1] SPECrate®2017_fp_base comparison based on best performing systems published at www.spec.org as of 07/06/2021. Configurations: 2x AMD EPYC 7763 (651 SPECrate®2017_fp_base, http://spec.org/cpu2017/results/res2021q1/cpu2017-20210219-24944.html, $15780 1Ku price total, 560W total TDP) versus 2x Intel Xeon Platinum 8380 (489 SPECrate®2017_fp_base, http://spec.org/cpu2017/results/res2021q2/cpu2017-20210521-26361.html, $16198 1Ku price total, 540W total TDP) for 1.33x the performance at 1.37x the score per total CPU $; 0.83x the performance/Core; 1.28x the performance/Watt. Top EPYC 7002 Series result: 2x AMD EPYC 7H12 (543 SPECrate®2017_fp_base, http://spec.org/cpu2017/results/res2020q3/cpu2017-20200707-23397.html). AMD 1Ku pricing and Intel ARK.intel.com specifications and pricing as of 4/6/21. SPEC®, SPEC CPU®, and SPECrate® are registered trademarks of the Standard Performance Evaluation Corporation. see www.spec.org for more information. MLN-086B

[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

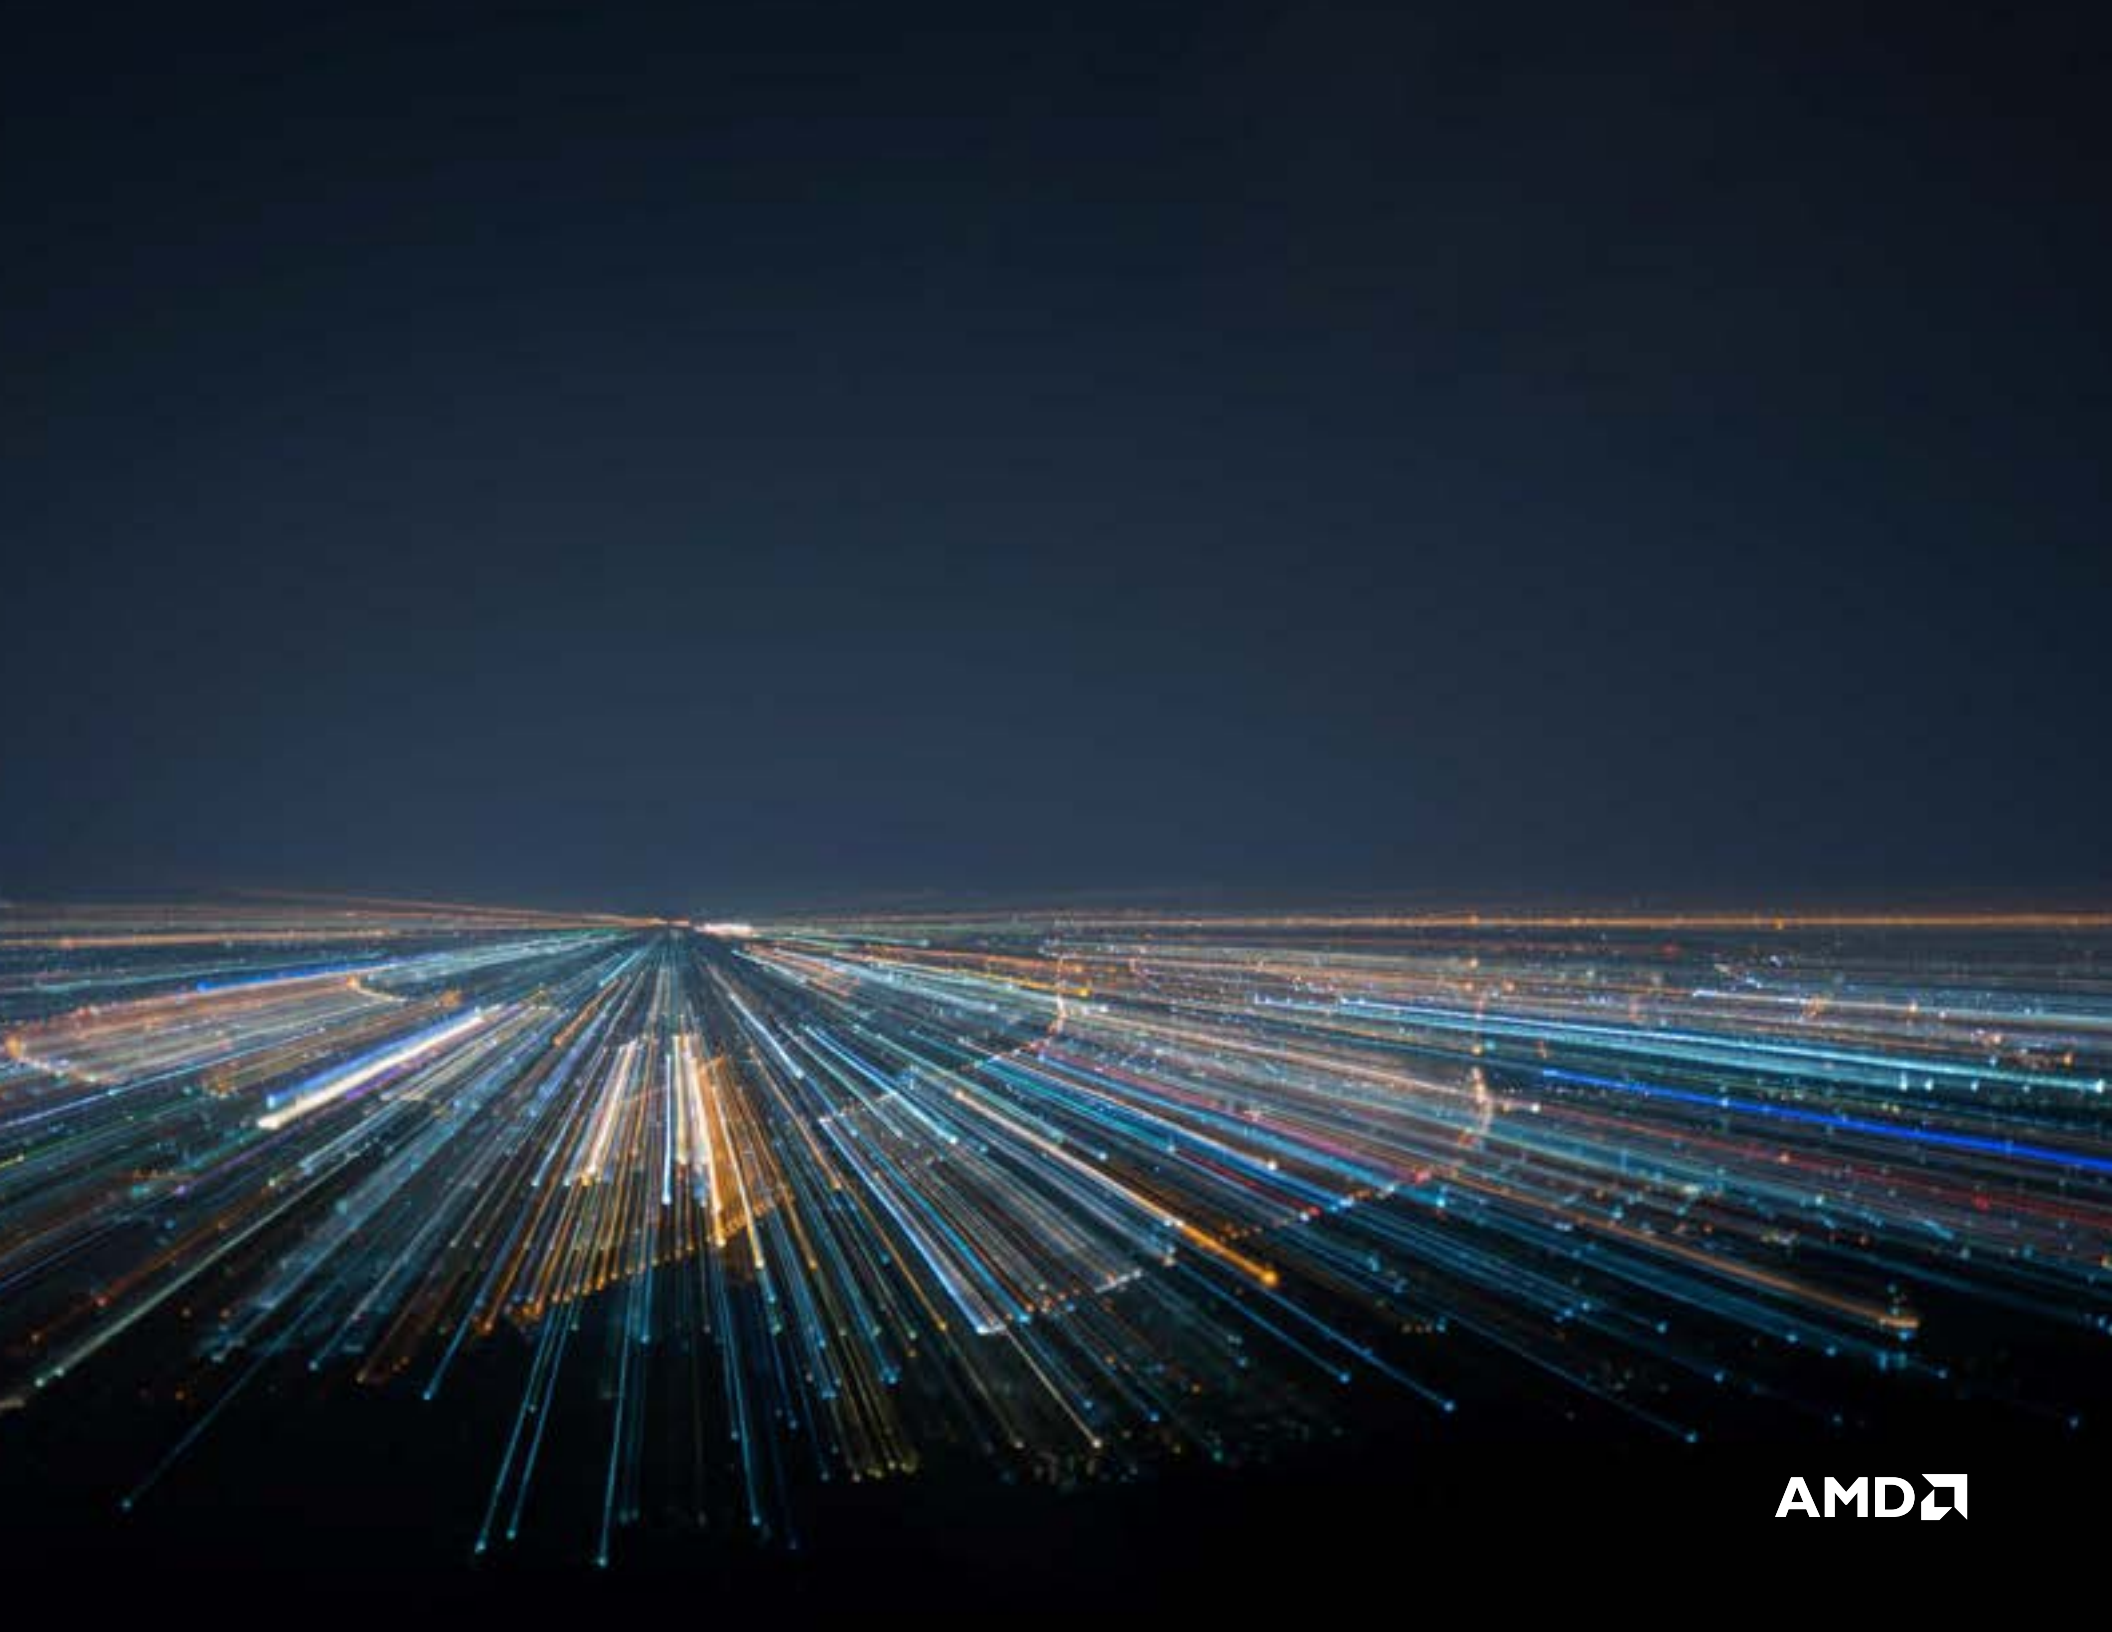
AMD