# DESIGN LEADERSHIP IMPERATIVE

MARK FUSELIER
SVP, TECHNOLOGY & PRODUCT ENGINEERING

# 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; TAM for data center, PCs, embedded and gaming; and technology trends, innovation and 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.

# COMPUTING POWER & ENERGY USE PROJECTIONS

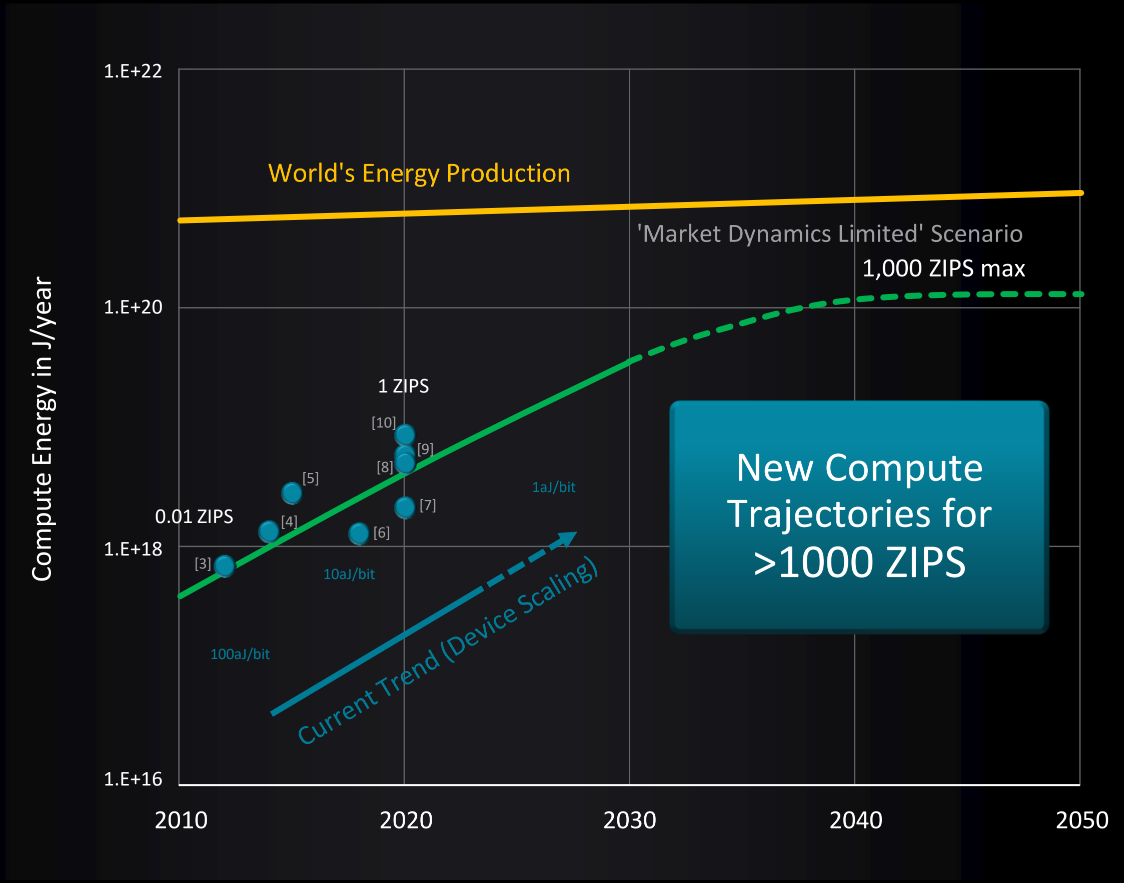


## POWER & COST SCALING IS CRITICAL

# TRENDS IN ENERGY EFFICIENCY

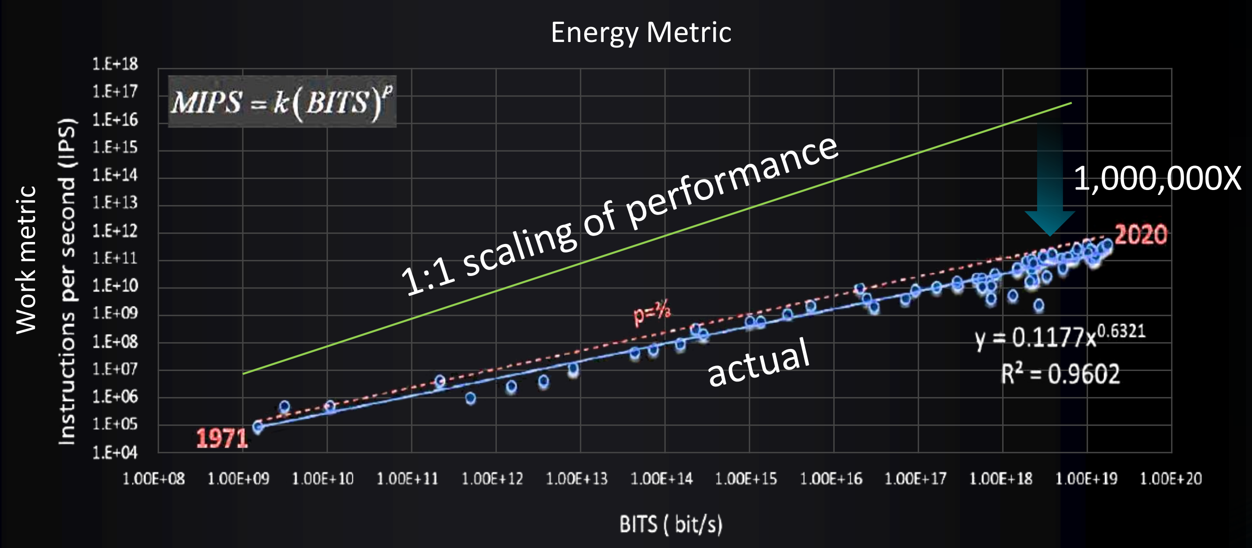


- Demand for computation growth is outpacing the progress realized by Moore's law
- It is now recognized that conventional computing is approaching fundamental limits in energy efficiency

## *MOORE'S LAW SCALING NOT KEEPING PACE*



AMD together we advance_

# COST/MM2 BECOMING PROHIBITIVE

NEW NODE INTRODUCTION RATE IS SLOWING

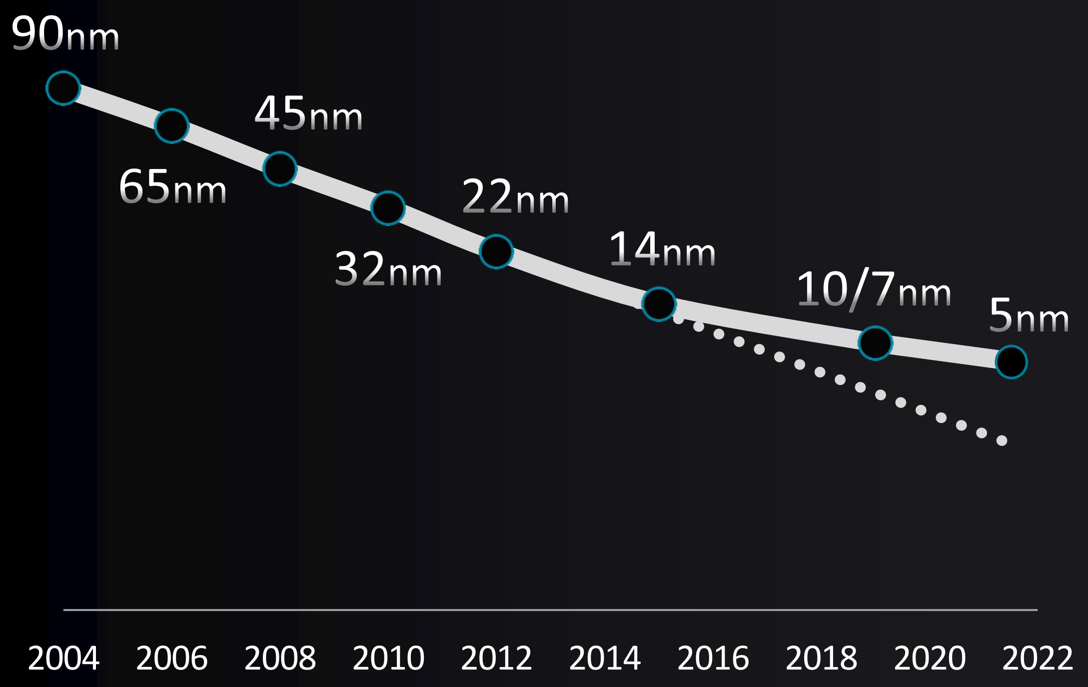


SOURCE: AMD.

WHILE COSTS CONTINUE TO INCREASE

(Cost Per Yielded mm2 for a 250mm2 Die)

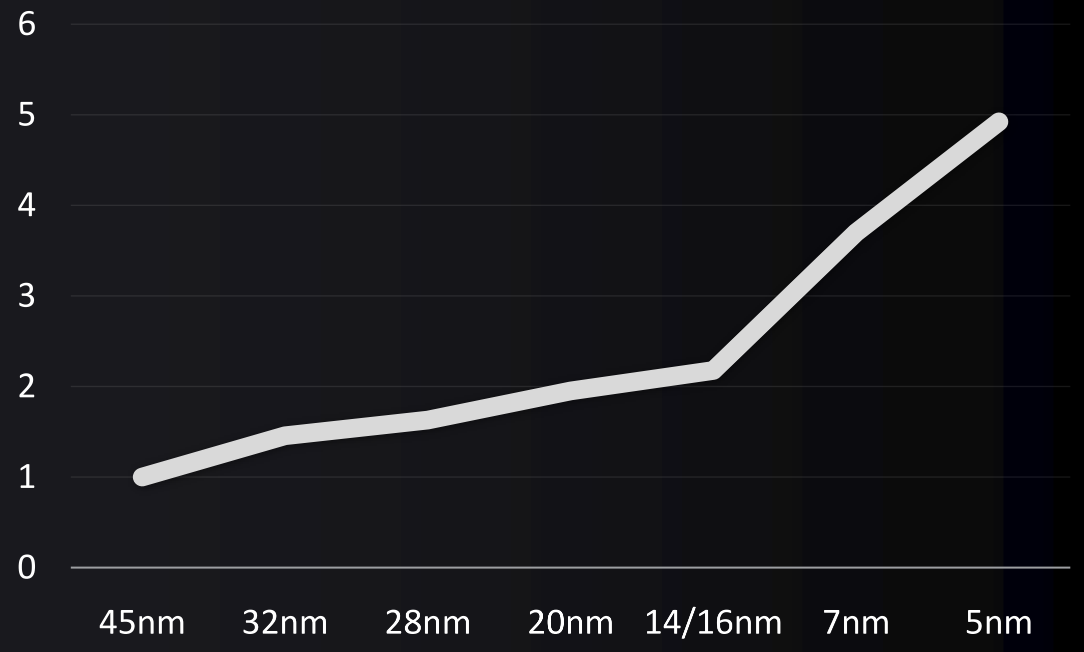


***NEW IMPROVEMENTS REQUIRED:***

***DESIGN & PACKAGING IN ADDITION TO FAB***

# COST SOLUTIONS: CHIPLETS

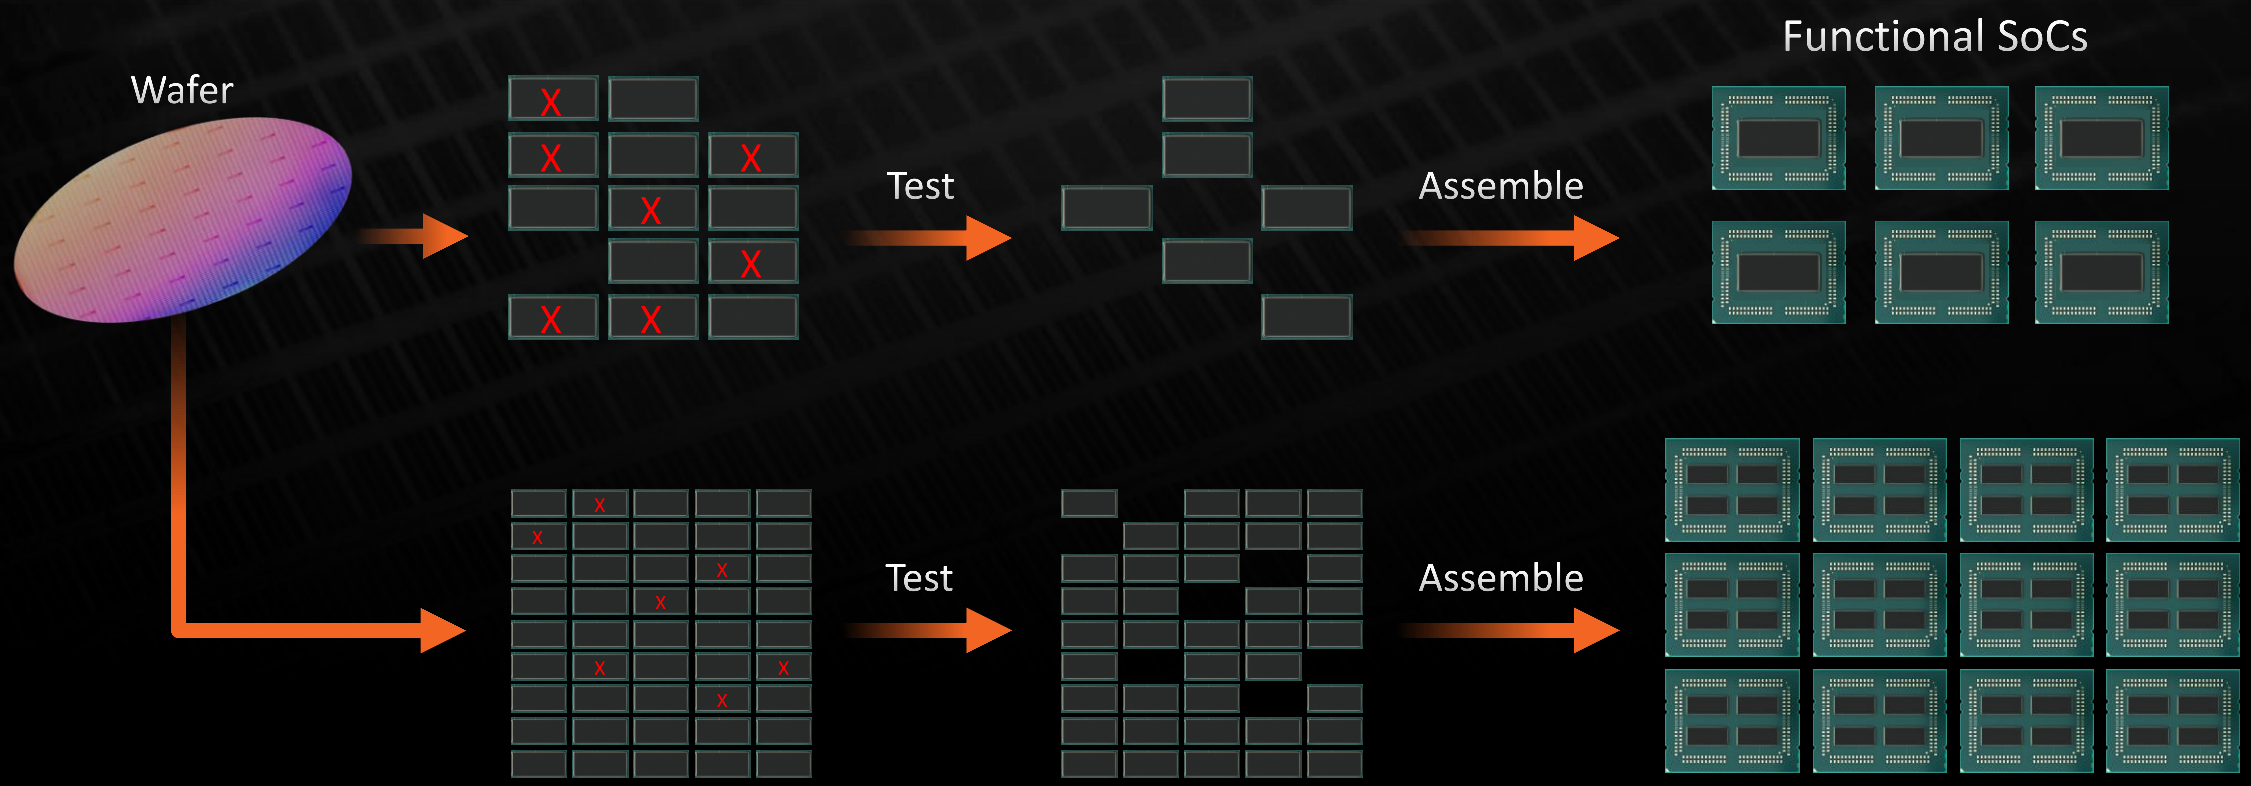


Many More Functional SoCs

Ability to mix and match at a finer grained level

# POWER SOLUTIONS: APPLICATION SPECIFIC COMPUTE w/ MODULAR DESIGN

## APPLICATION-SPECIFIC OPTIMIZATION PROVIDES BETTER PERFORMANCE-PER-WATT

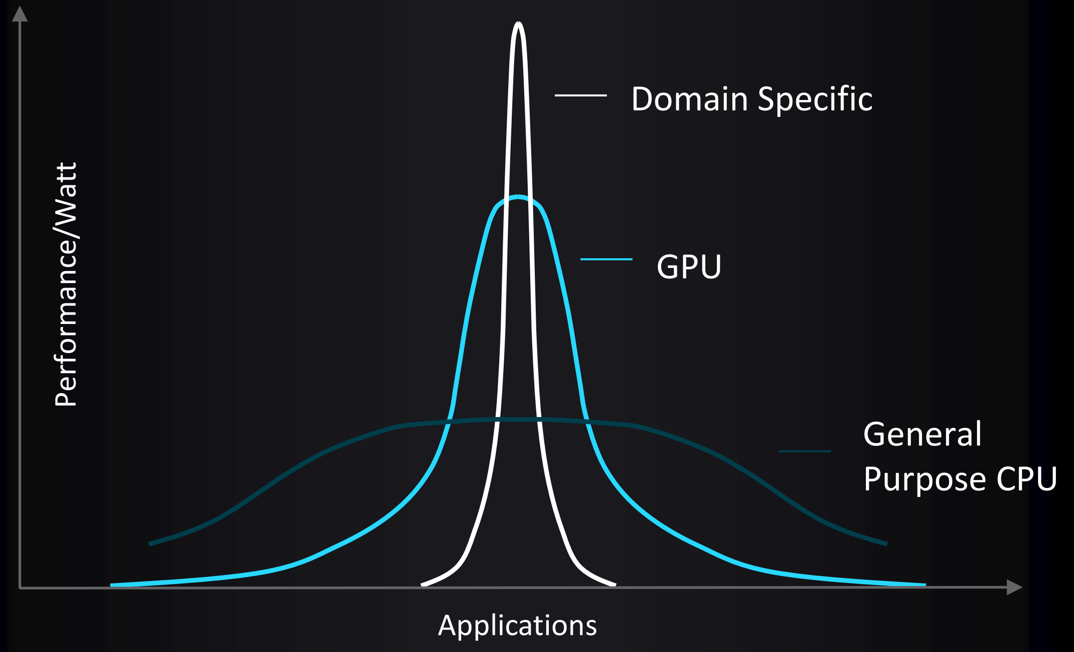


## HOW TO ARCHITECT, DESIGN AND BUILD FUTURE SYSTEMS?

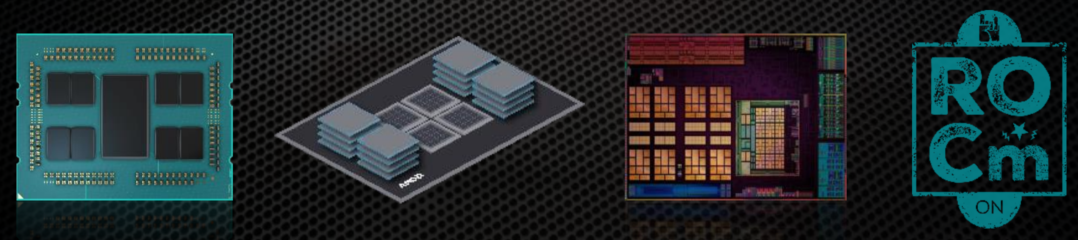


## MODULAR DESIGN

Accelerators are more transistor efficient

Though by definition limited in application

Technology that enables flexible configurations will deliver better perf/transistor at the solution level



AMD
together we advance_

# DESIGN ENGAGEMENT WITH MANUFACTURING PARTNERS

## ENABLING MEANINGFUL INNOVATION


SYSTEM/DESIGN/TECHNOLOGY CO-OPTIMIZATION

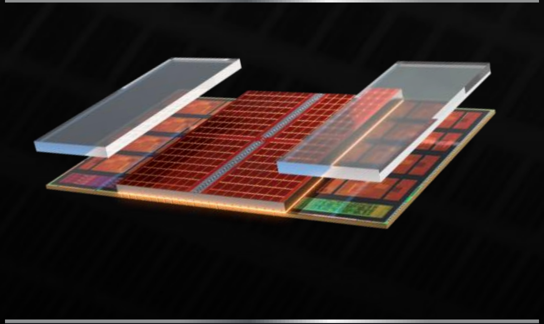

ARCHITECTING WITH ADVANCED PACKAGING

POWER & COST REDUCTION FOCUS

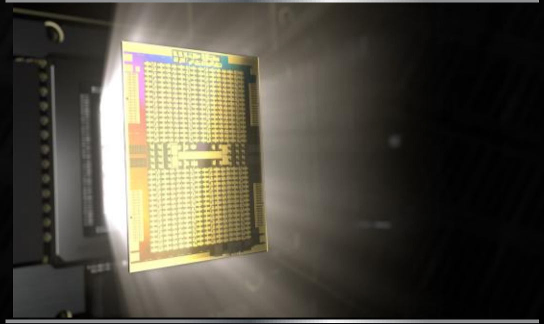

RISK MITIGATE DISRUPTIVE TECHNOLOGIES

**DEEP PARTNERSHIP: DESIGN / FOUNDRY / PACKAGING / SYSTEM**



AMD
together we advance_

AMD

# REFERENCES


1. [Friedrich14] J. Friedrich, et al., “POWER8: A 12-Core Server-Class Processor in 22nm SOI with 7.6Tb/s Off-Chip Bandwidth”, ISSCC, Feb. 2014.
2. [Kurd14] N. Kurd, et al., “Haswell: A Family of IA 22nm Processors”, ISSCC, Feb. 2014.

# ENDNOTES

AMD 3D CHIPLET TECHNOLOGY
Competition 3D architecture picture from SystemPlus. Intel Core i5-L16G7: the first utilization of Intel's Foveros Technology with Package-on-Package configuration in a consumer product. https://www.systemplus.fr/reverse-costing-reports/intel-foveros-3d-packaging-technology/.

3D CHIPLET GAMING DEMO AND PERFORMANCE CHART
Testing by AMD performance labs as of April 28, 2021 based on the average FPS of 32 PC games at 1920x1080 with the High image quality preset using an AMD Ryzen™ 9 5900X processor vs. 12-Core 3D Chiplet Prototype. Results may vary. R5K-078.

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.

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.

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.



AMD
together we advance_

# COPYRIGHT AND DISCLAIMER





AMD
together we advance_