AMD

Q1 2018 TECH DAY

Mark Papermaster

SVP & CTO, Technology & Engineering

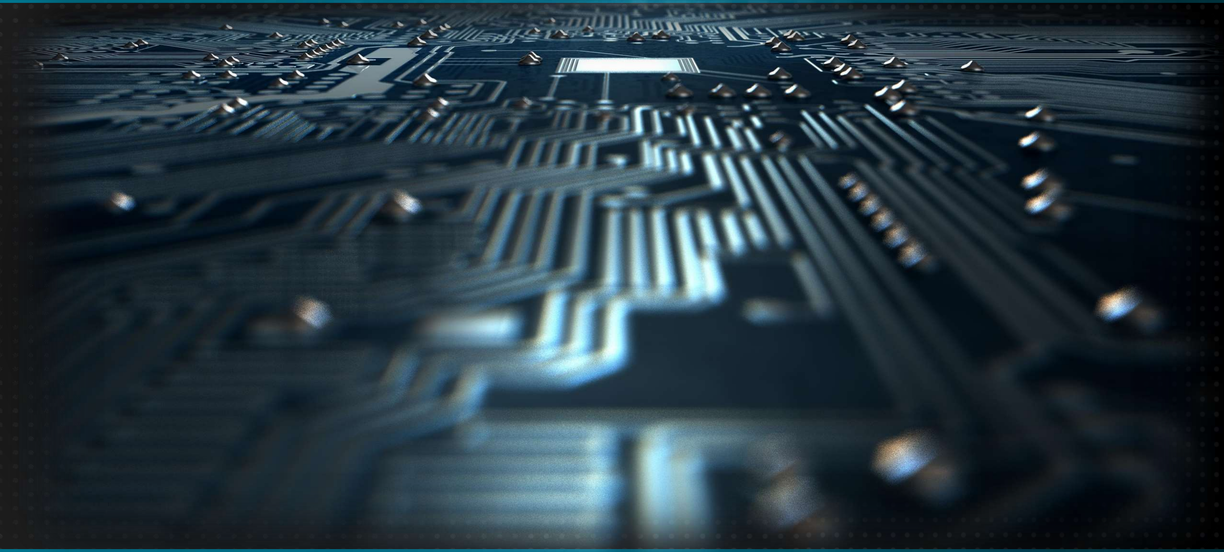

Cautionary Statement

- This presentation contains forward-looking statements concerning Advanced Micro Devices, Inc. (AMD) including but not limited to, AMD's technology roadmap; 12LP technology deployment; AMD's X86 roadmap; AMD's graphics architecture roadmap; AMD's machine learning stack; and the features, functionality, availability, timing, and expected benefits of AMD future products including Radeon™ Instinct Vega 7nm GPUs, 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 Quarterly Report on Form 10-Q for the quarter ended September 30, 2017.

SPECULATIVE EXECUTION: AMD PROCESSOR SECURITY

AMD has a dedicated team with well-defined process for engaging on security issues.

AMD worked closely with the ecosystem to fully understand their findings and evaluate risk across product portfolio.

Due to micro architectural differences and O/S patches being deployed now, we believe there is near zero risk to AMD users.

Google Project Zero (GPZ) Research Title		DETAILS
Variant One	Bounds Check Bypass	Resolved by software/OS updates being made available by vendors/manufacturers.
Variant Two	Branch Target Injection	Differences in AMD architecture mean there is a near zero risk, and vulnerability to Variant 2 has not been demonstrated on AMD processors to-date.
Variant Three	Rogue Data Cache Load	Zero AMD vulnerability or risk because of AMD architecture differences.


CHANGING
THE RULES OF
TECHNOLOGY
DEVELOPMENT

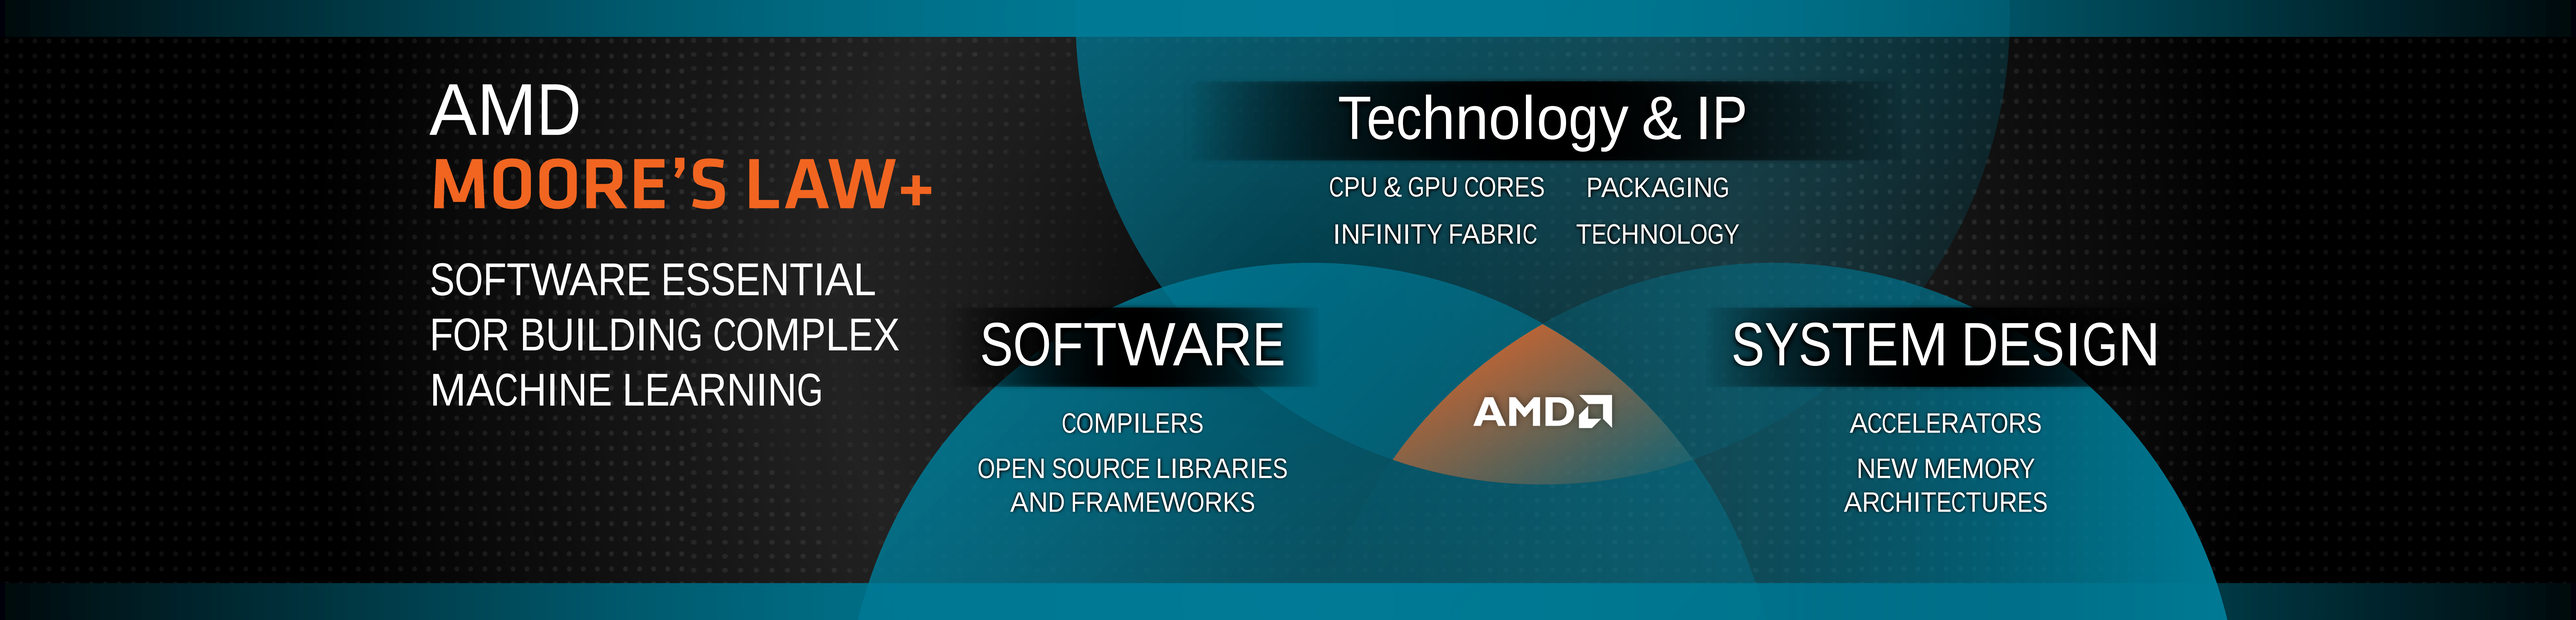
AMD
MOORE'S LAW+
SOFTWARE ESSENTIAL FOR BUILDING COMPLEX MACHINE LEARNING
Technology & IP
CPU & GPU CORES
PACKAGING
INFINITY FABRIC
TECHNOLOGY
SOFTWARE
COMPILERS
OPEN SOURCE LIBRARIES AND FRAMEWORKS
AMD
SYSTEM DESIGN
ACCELERATORS
NEW MEMORY ARCHITECTURES

INTEGRATION:
“ZEN” & “VEGA”

UNITED WITH
INFINITY FABRIC

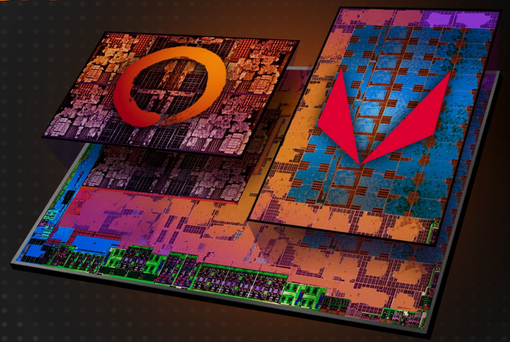

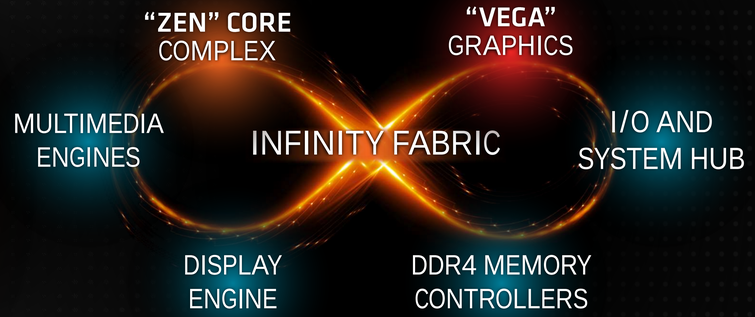

TECHNOLOGY ROADMAP

LEADERSHIP LEVELS THE COMPETITIVE PLAYING FIELD

Roadmap Subject to Change.

12LP Technology Deployment

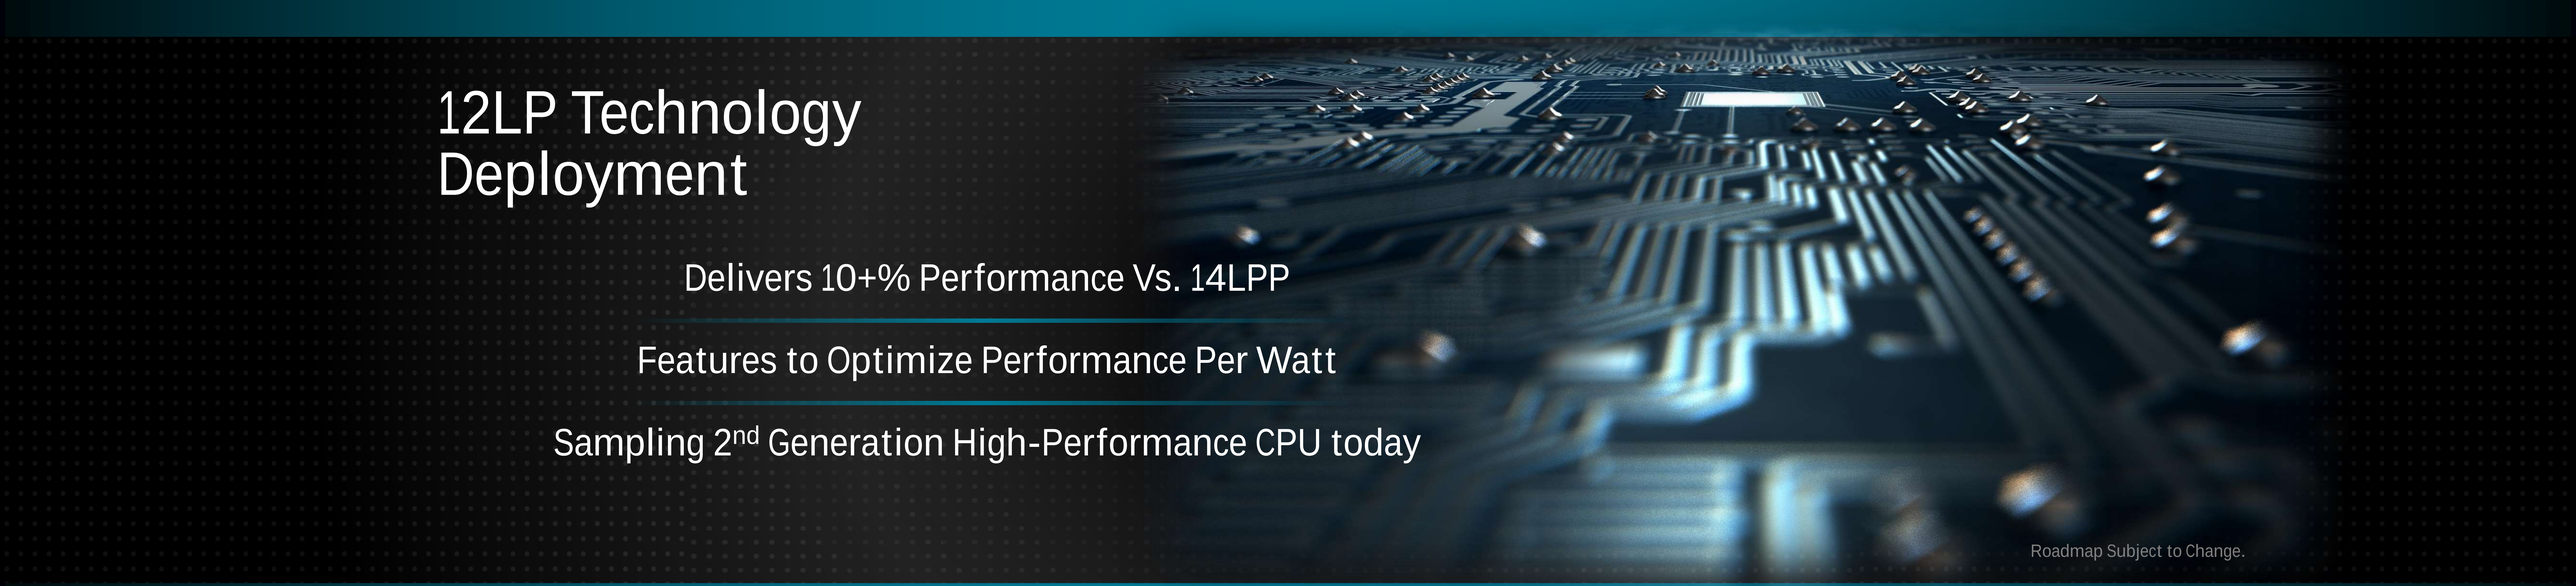

Delivers 10+% Performance Vs. 14LPP

Features to Optimize Performance Per Watt

Sampling 2nd Generation High-Performance CPU today

Roadmap Subject to Change.

“ZEN” CORE | PERFORMANCE LEADERSHIP DELIVERED

HEDT Performance in Mainstream Desktops

World Record Server Performance*

World’s Fastest Processor for Ultrathins**

*Data Reflects EPYC 7601 On SPECrate® 2017_fp_peak and SPECfp®_rate2006.

**See Endnotes for Details.

X86 ROADMAP LEADERSHIP

2017

“ZEN”
14nm

Current Products

“ZEN+”
12nm

Now Sampling

“ZEN 2”
7nm

“ZEN 2” Design is Complete

Improves on “Zen” in Multiple Dimensions

2020

“ZEN 3”
7nm+

On Track

Roadmap Subject to Change.

GRAPHICS ARCHITECTURE ROADMAP

2017

“VEGA”
14nm

“VEGA”
7nm

“NAVI”
7nm

2020

NEXT-GEN
7nm+

Roadmap Subject to Change.

COMPREHENSIVE AMD SOFTWARE INFRASTRUCTURE ENABLES EMERGING MI SOLUTIONS

Complete Software Stack Accelerates Development and Deployment of Custom MI Applications

Open Architecture Invites Community Participation to Develop Fully Optimized Solutions

ANNOUNCING AT CES 2018:

AMD Achieves Production Level Machine Learning Environment

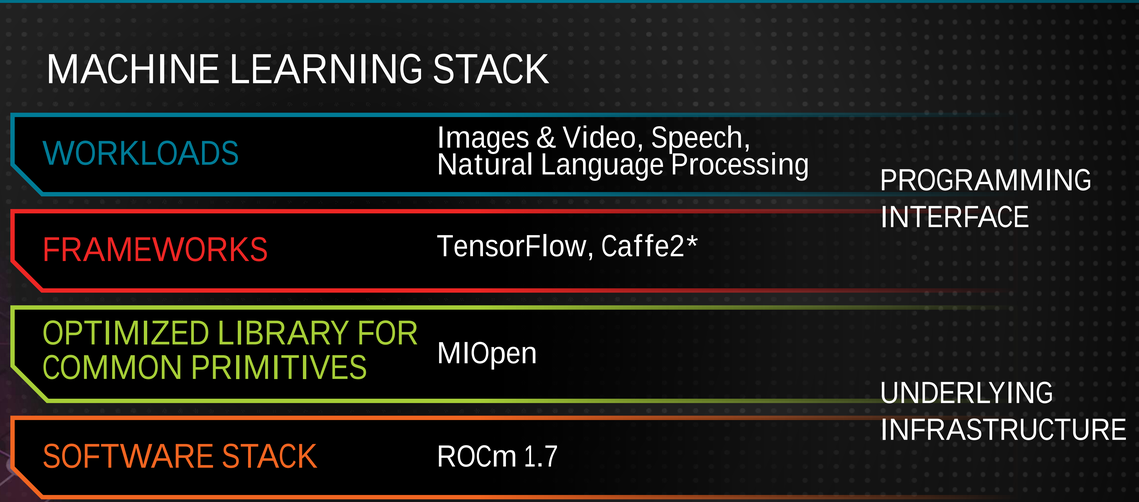

*Coming Soon.

MACHINE INTELLIGENCE SOFTWARE SUPPORT

LIBRARY
UNDERLYING INFRASTRUCTURE

MIOpen

MIOpen, BLAS, FFT, RNG, RCCL, NCCL, cuDNN

SOFWARE STACK
UNDERLYING INFRASTRUCTURE

ROCm ON

OpenCL

ROCm, **HIP, OpenCL™, HCC, CUDA**

AMD MACHINE LEARNING:

Easy to Deploy Porting Capability

CUDA-based application

"HIPify"
Virtually Automatic Conversion

Portable HIP C++

Developer maintains HIP port

AMD

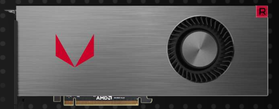

Resulting C++ code runs on NVIDIA or AMD GPUs

NVIDIA

AMD MACHINE LEARNING

THE FIRST TOP-TO-BOTTOM OPEN SOLUTION

AMD Enables Complete Open Architecture Application Development

	AMD / ROCm	NVIDIA / CUDA
Programming Language	Open-Source HIP	Proprietary (CUDA)
Accelerated MI Library	Open-Source MIOpen	Proprietary (cuDNN)
Accelerated Math Libraries	Open-Source rocBLAS, rocRAND, *rocFFT*, *rocSPARSE*	Proprietary (cuBLAS, cuRNG, cuFF, cuSPARSE)
Communication Library	*RCCL*	NCCL
Runtime	Open-Source ROCr	Proprietary
Linux Driver	Open-Source (AMDGPU)	Proprietary
Documented ISA	Open (GCN)	Proprietary

Italics = Under Development.

ENABLING LARGE SCALE ACCELERATED MACHINE LEARNING

CHALLENGE			
Huge working datasets	Security for public cloud & large scale data center deployments	Portability of algorithms & apps across different platform configurations	Large scale data centers require robust platforms
SOLUTION			
Large HBM local memory & coherent access between GPUs in a cluster	Isolate host from the end user with hardware GPU virtualization	Consistent software stack enables different configurations	Same RAS (reliability, availability and serviceability) capabilities for GPU & enterprise-class CPU

INTRODUCING MACHINE LEARNING OPTIMIZED RADEON GRAPHICS

Radeon™ Instinct
"Vega" 7nm

SAMPLING IN 2018

Optimized Enterprise Class Deep Learning Training and Inference with New DL Ops

Enables General Purpose HPC Clusters with New High Speed I/O

Primary Markets: Machine Learning, Automotive, and High Performance Computing

Virtualized Compute with MX-GPU including SR-IOV Support

AMD

Q1 2018 TECH DAY

Thank you

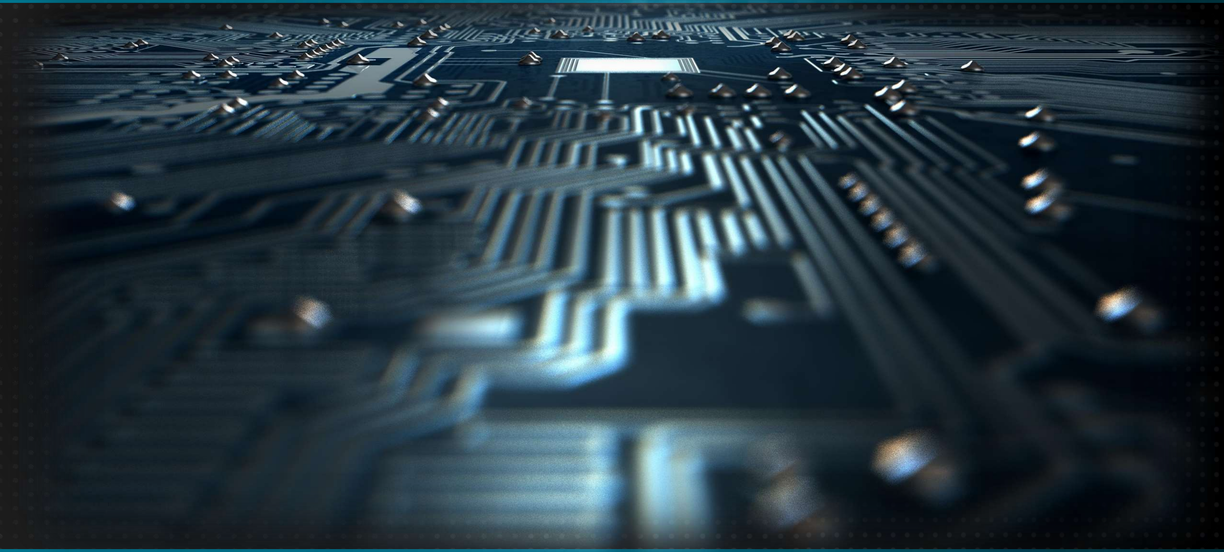

End Notes

Slide 2.

“World Record Server Performance” AMD EPYC 7601-based system scored 2060 on SPECfp®_rate2006, higher than any other 2-socket system score published at www.spec.org as of 5 January 2018.Result available at https://www.spec.org/cpu2006/results/res2017q4/cpu2006-20171114-50603.html. EPYC 7601 scored 2060 on SPECfp_rate2006 in AMD internal testing. 2 x EPYC 7601 CPU in HPE ProLiant DL385 Gen10, SUSE Linux Enterprise Server 12 SP3, x86 Open64 v4.5.2.1 Compiler Suite, 1 TB (16 x 64GB 4Rx4 PC4-2666V-L) memory, 1 x 300 GB 15k RPM SAS. SPEC and SPECfp are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information. NAP-67

AMD EPYC 7601-based system scored 272 on SPECrate®2017_fp_peak, higher than any other 2-socket system score published at www.spec.org as of 5 January 2018. Result available at https://www.spec.org/cpu2017/results/res2017q4/cpu2017-20171114-00845.html EPYC 7601 scored 272 on SPECrate®2017_fp_peak in AMD internal testing. 2 x EPYC 7601 CPU in HPE ProLiant DL385 Gen10, SUSE Linux Enterprise Server 12 SP3, AOCC v1.0.0 compiler, 1 TB (16 x 64GB 4Rx4 PC4-2666V-L) memory, 1 x 300 GB 15k RPM SAS. SPEC and SPECrate are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information. NAP-68

“Processor for ultrathin laptops” defined as 15W nominal processor TDP. Based on testing of the AMD Ryzen™ 7 2700U, AMD Ryzen™ 5 2500U, and Core i7-8550U 15W mobile processors as of 10/6/2017. Performance based on Cinebench R15 nT (“CPU performance”) and 3DMark® TimeSpy (“GPU performance”) in order of AMD Ryzen 7 2700U, AMD Ryzen 5 2500U and Intel 8550U. AMD Ryzen™ 7 2700U: AMD Reference, AMD Ryzen™ 7 2700U with Radeon™ Vega 10 Processor Graphics, 8GB DDR4-2400 RAM, Samsung 850 PRO 512GB SATA SSD, Windows 10 Pro RS2, Graphics driver 23.20.768.9, 26-Sep-201