

# CAUTIONARY STATEMENT

This presentation contains forward-looking statements concerning Advanced Micro Devices, Inc. (AMD) including, but not limited to AMD's strategy and focus, expected datacenter total addressable market; expected benefits from EPYC™ based AWS R5, M5 and T3 instances, the features, functionality, availability, timing, deployment, and expected benefits of AMD future products and technologies, including "Rome" and AMD Radeon Instinct™ MI60; and the benefits and expectations of 7nm process technology, 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 29, 2018.

# AMD

## NEXT HORIZON

DR. LISA SU  
President and CEO

# OUR JOURNEY

---

High-Performance Technologies

Great Products

Ambitious Goals

Undaunted Determination

# OUR STRATEGY AND FOCUS

## GRAPHICS



Gaming

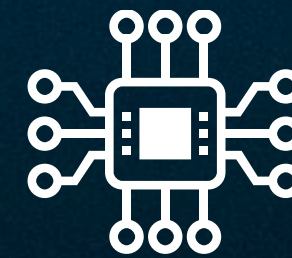


Compute & AI



Virtual & Augmented Reality

## COMPUTE

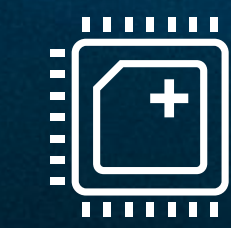


Client Systems



Infrastructure & Cloud

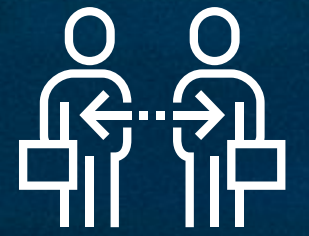
## SOLUTIONS



Semi-Custom



Vertical Platforms



Partnerships

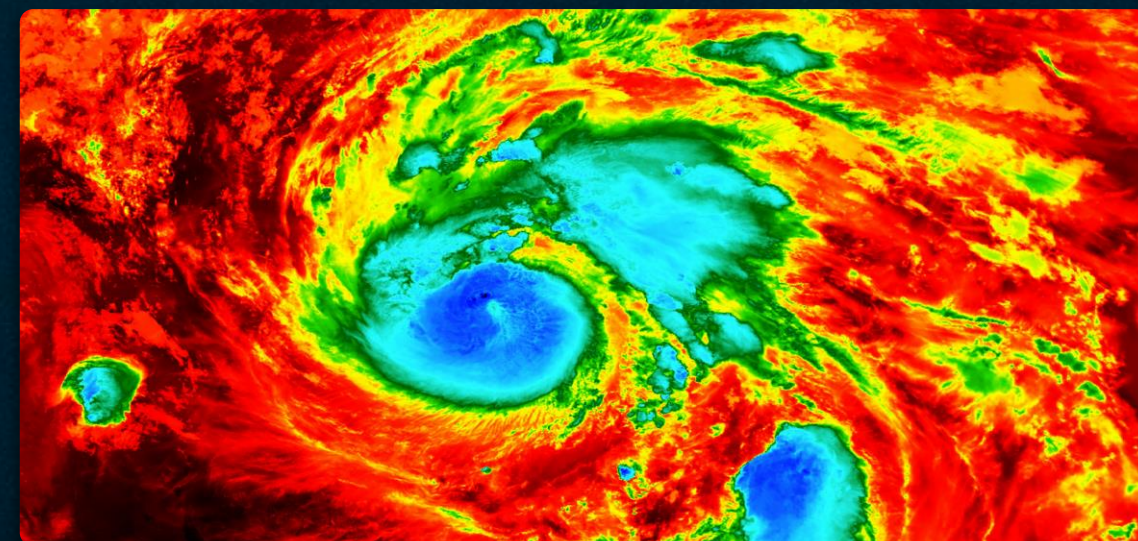


THE DATACENTER HOLDS  
ENORMOUS POTENTIAL  
TO IMPROVE OUR WORLD

# THE WORLD'S TOUGHEST PROBLEMS DEMAND NEW THINKING



LARGE SCALE SIMULATIONS



CLIMATE CHANGE



EDUCATION



ENERGY SOLUTIONS



COMPUTATIONAL BIOLOGY



DISEASE PREVENTION

# THE WORKLOADS OF THE FUTURE REQUIRE INCREDIBLE AMOUNTS OF COMPUTE POWER



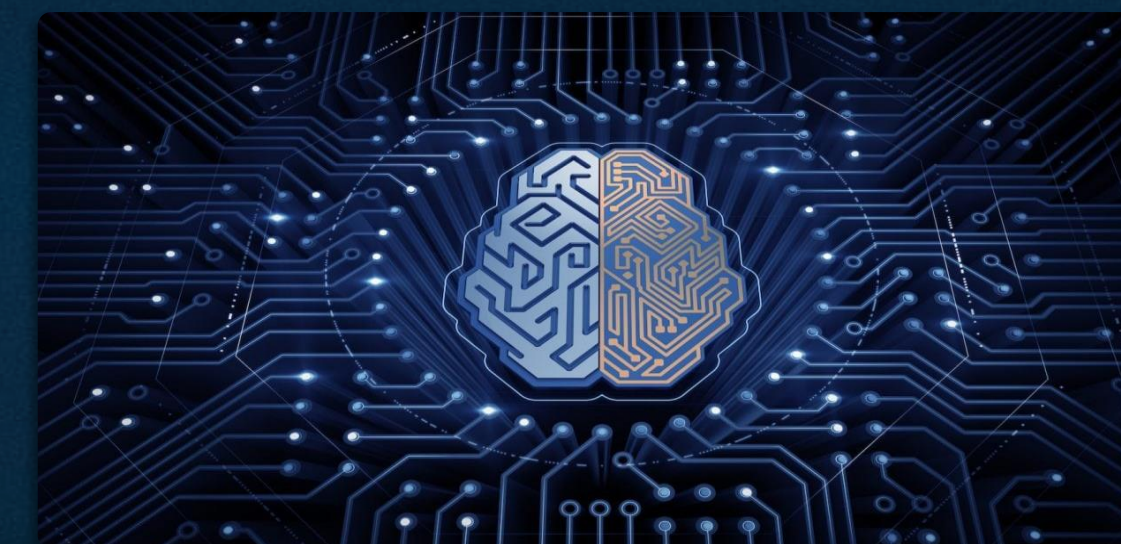
HIGH PERFORMANCE  
COMPUTING



CLOUD, HYPERSCALE  
AND VIRTUALIZATION



MACHINE  
INTELLIGENCE



IMMERSIVE AND INSTINCTIVE  
COMPUTING



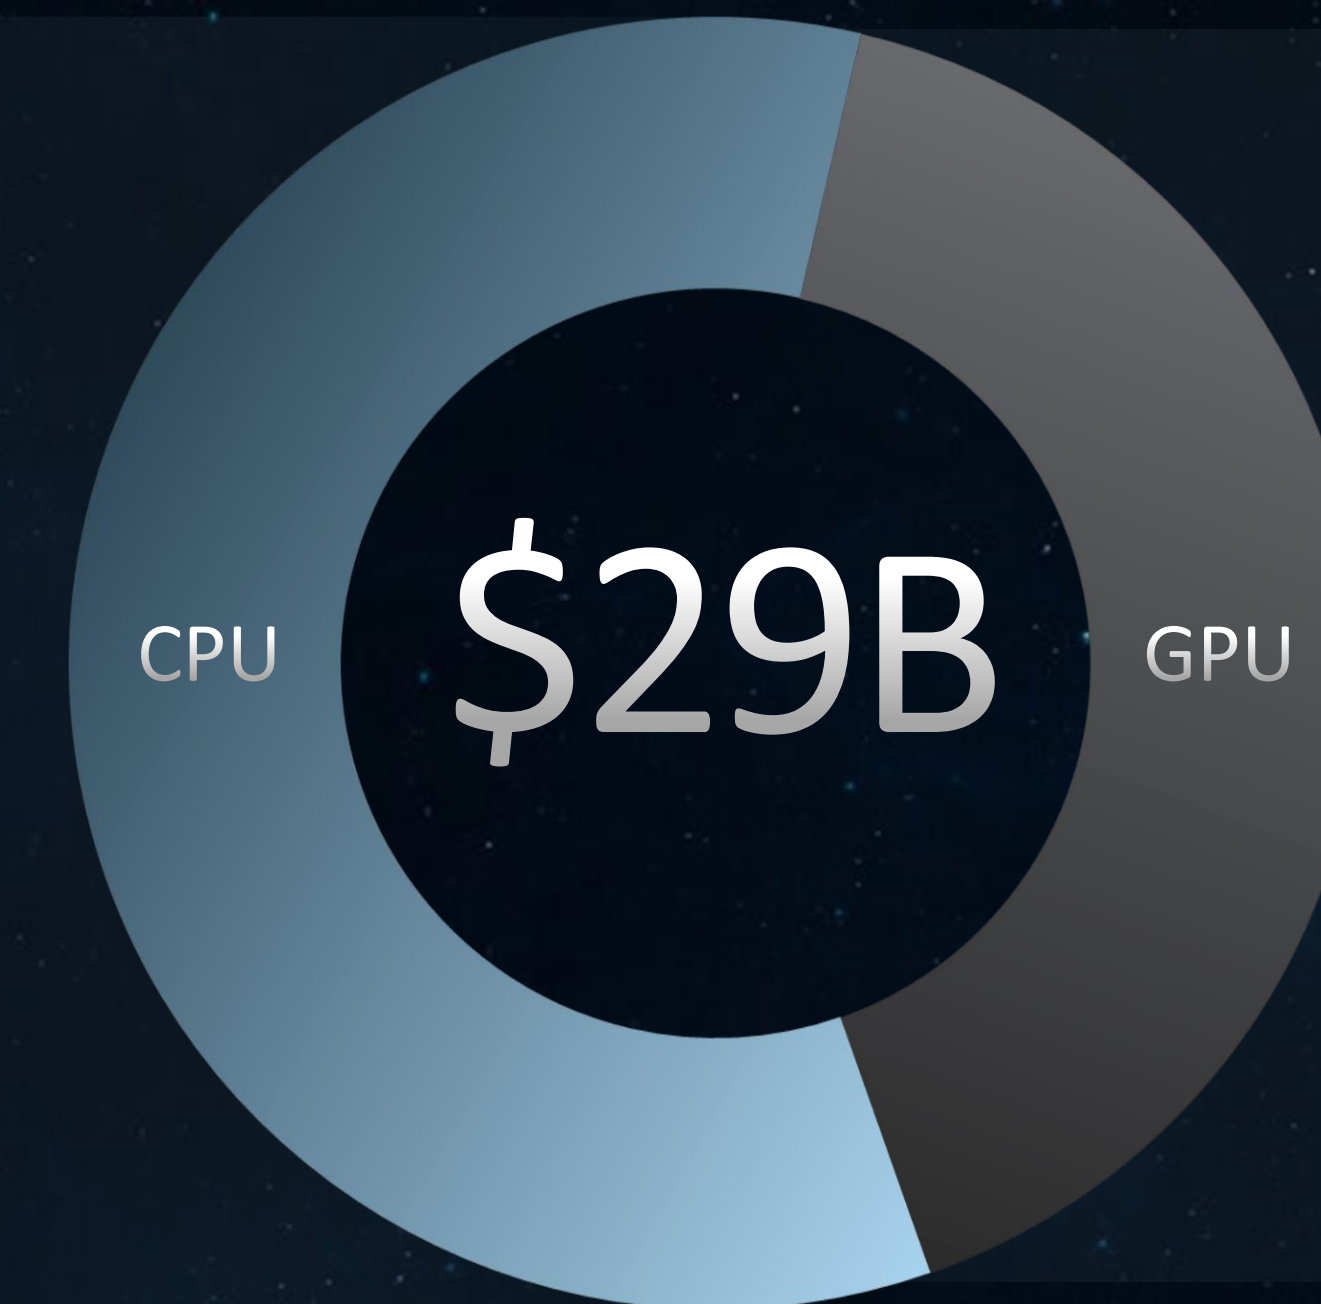
BIG DATA  
ANALYTICS



SOFTWARE-DEFINED  
STORAGE

...AND REPRESENTS AN INCREDIBLE  
OPPORTUNITY

AMD  
EPYC



AMD  
RADEON  
INSTINCT

DATACENTER TAM  
by 2021

AMD Estimate For 2021

# MODERN DATACENTERS REQUIRE NEW THINKING

---

All-New Design Approach  
to CPU and GPU Architecture

Leadership Interconnects  
for System Performance

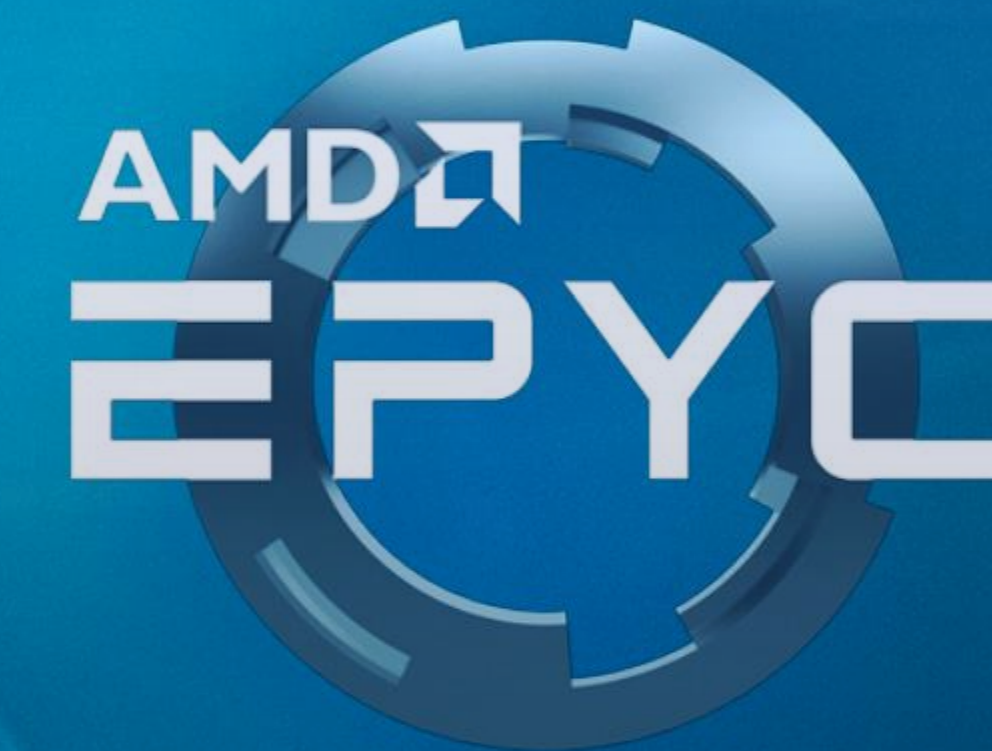
Aggressive  
Long-Term Roadmaps

IN 2017, AMD EPYC™ DELIVERED  
**A NEW ERA IN THE DATACENTER**

Performance Leadership at  
Every Competitive Price Point

More Cores, More Memory  
Bandwidth, More I/O

Unmatched  
Total Cost of Ownership



# AMD RADEON INSTINCT™ MI25 BROUGHT CHOICE TO THE CLOUD

Designed for Compute  
and Cloud Workloads

~25 Teraflops  
of Performance

Optimized MIOpen Libraries  
for Deep Learning

AMD  
RADEON  
INSTINCT

ROCm ENABLED THE FIRST COMPLETE  
**OPEN SOURCE PLATFORM  
FOR GPU COMPUTE**

Unlocked GPU Power To  
Accelerate Computational Tasks

Optimized for HPC and  
Ultra-Scale Class Computing

Enabling Innovation,  
Collaboration, and Efficiency



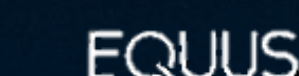
## AMD DATACENTER MOMENTUM

# THE INDUSTRY RESPONDED

## OEM/ODM

Hewlett Packard  
EnterpriseDELL EMCCISCOLenovoSUPERMICRASUSTYANwlstronGIGABYTEInventecSugonH3CQCT

## SYSTEM INTEGRATORS/PARTNERS

CIARAAMAXCRAYinspuratípa technologiesSILICON  
MECHANICSINTERNATIONAL  
COMPUTER CONCEPTSBOXXKOI  
COMPUTERS  
SOLUTIONS TODAY WITH  
TOMORROW'S TECHNOLOGYPENGUIN  
COMPUTINGEQUUSAtosEXXACTAspen SystemsMicroway

## AMD CLOUD MOMENTUM

# THE INDUSTRY RESPONDED

### CPU CLOUD DEPLOYMENTS

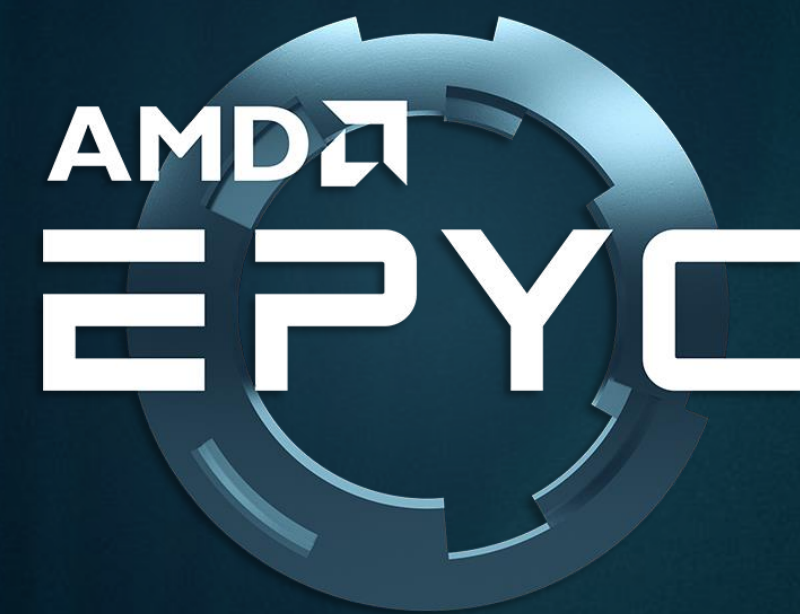


### GPU CLOUD DEPLOYMENTS



# ANNOUNCING TODAY

---



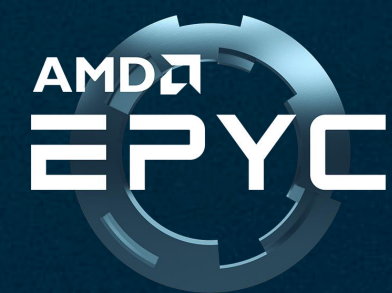
# Matt Garman

Vice President, Compute Services



# AMD EPYC™ NOW POWERING THREE NEW INSTANCES FOR AMAZON ELASTIC COMPUTE CLOUD (EC2)

Exceptional Performance-per-Dollar for General Purpose and Memory Optimized Workloads



AVAILABLE TODAY



ANNOUNCING AVAILABILITY  
OF AMD EPYC™ CPU BASED  
R5a, M5a AND  
T3a INSTANCES

Features AMD EPYC™ Processors' Industry-leading  
Core Density And Memory Bandwidth

Exceptional Performance-per-Dollar for General  
Purpose and Memory Optimized Workloads

Seamless Application Migration



The background of the entire image is a view of Earth from space. The left side shows the dark, silhouetted continents of North and South America, with city lights glowing as a network of golden-yellow dots. The right side shows the bright blue curve of the Earth's horizon against the blackness of space, which is filled with distant stars and a few nebulae. A bright light source, likely the sun, is positioned just behind the horizon line, creating a lens flare effect.

TODAY, WE REVEAL  
**THE NEXT HORIZON**

# EXTENDING OUR LEADERSHIP DESIGN APPROACH TO DATACENTER CPUs

---

New “Zen 2”  
High-Performance Core

Modular System Design Extends Compute,  
Memory Bandwidth and I/O Leadership

7nm Process Technology for  
Performance, Density and Power

# DESIGNING INNOVATIVE GPUs FOR THE CLOUD

High-Performance  
Compute Engines

Extensive Mixed Precision Ops for  
Optimum Deep Learning Performance

Leadership Bandwidth and  
Interconnect Technology



INTRODUCING TODAY

# AMD RADEON INSTINCT™ MI60

World's First 7nm GPU

World's Fastest  
HPC PCIe® Capable Accelerator

Industry's Only  
Hardware-Virtualized GPU



PREVIEWING TODAY

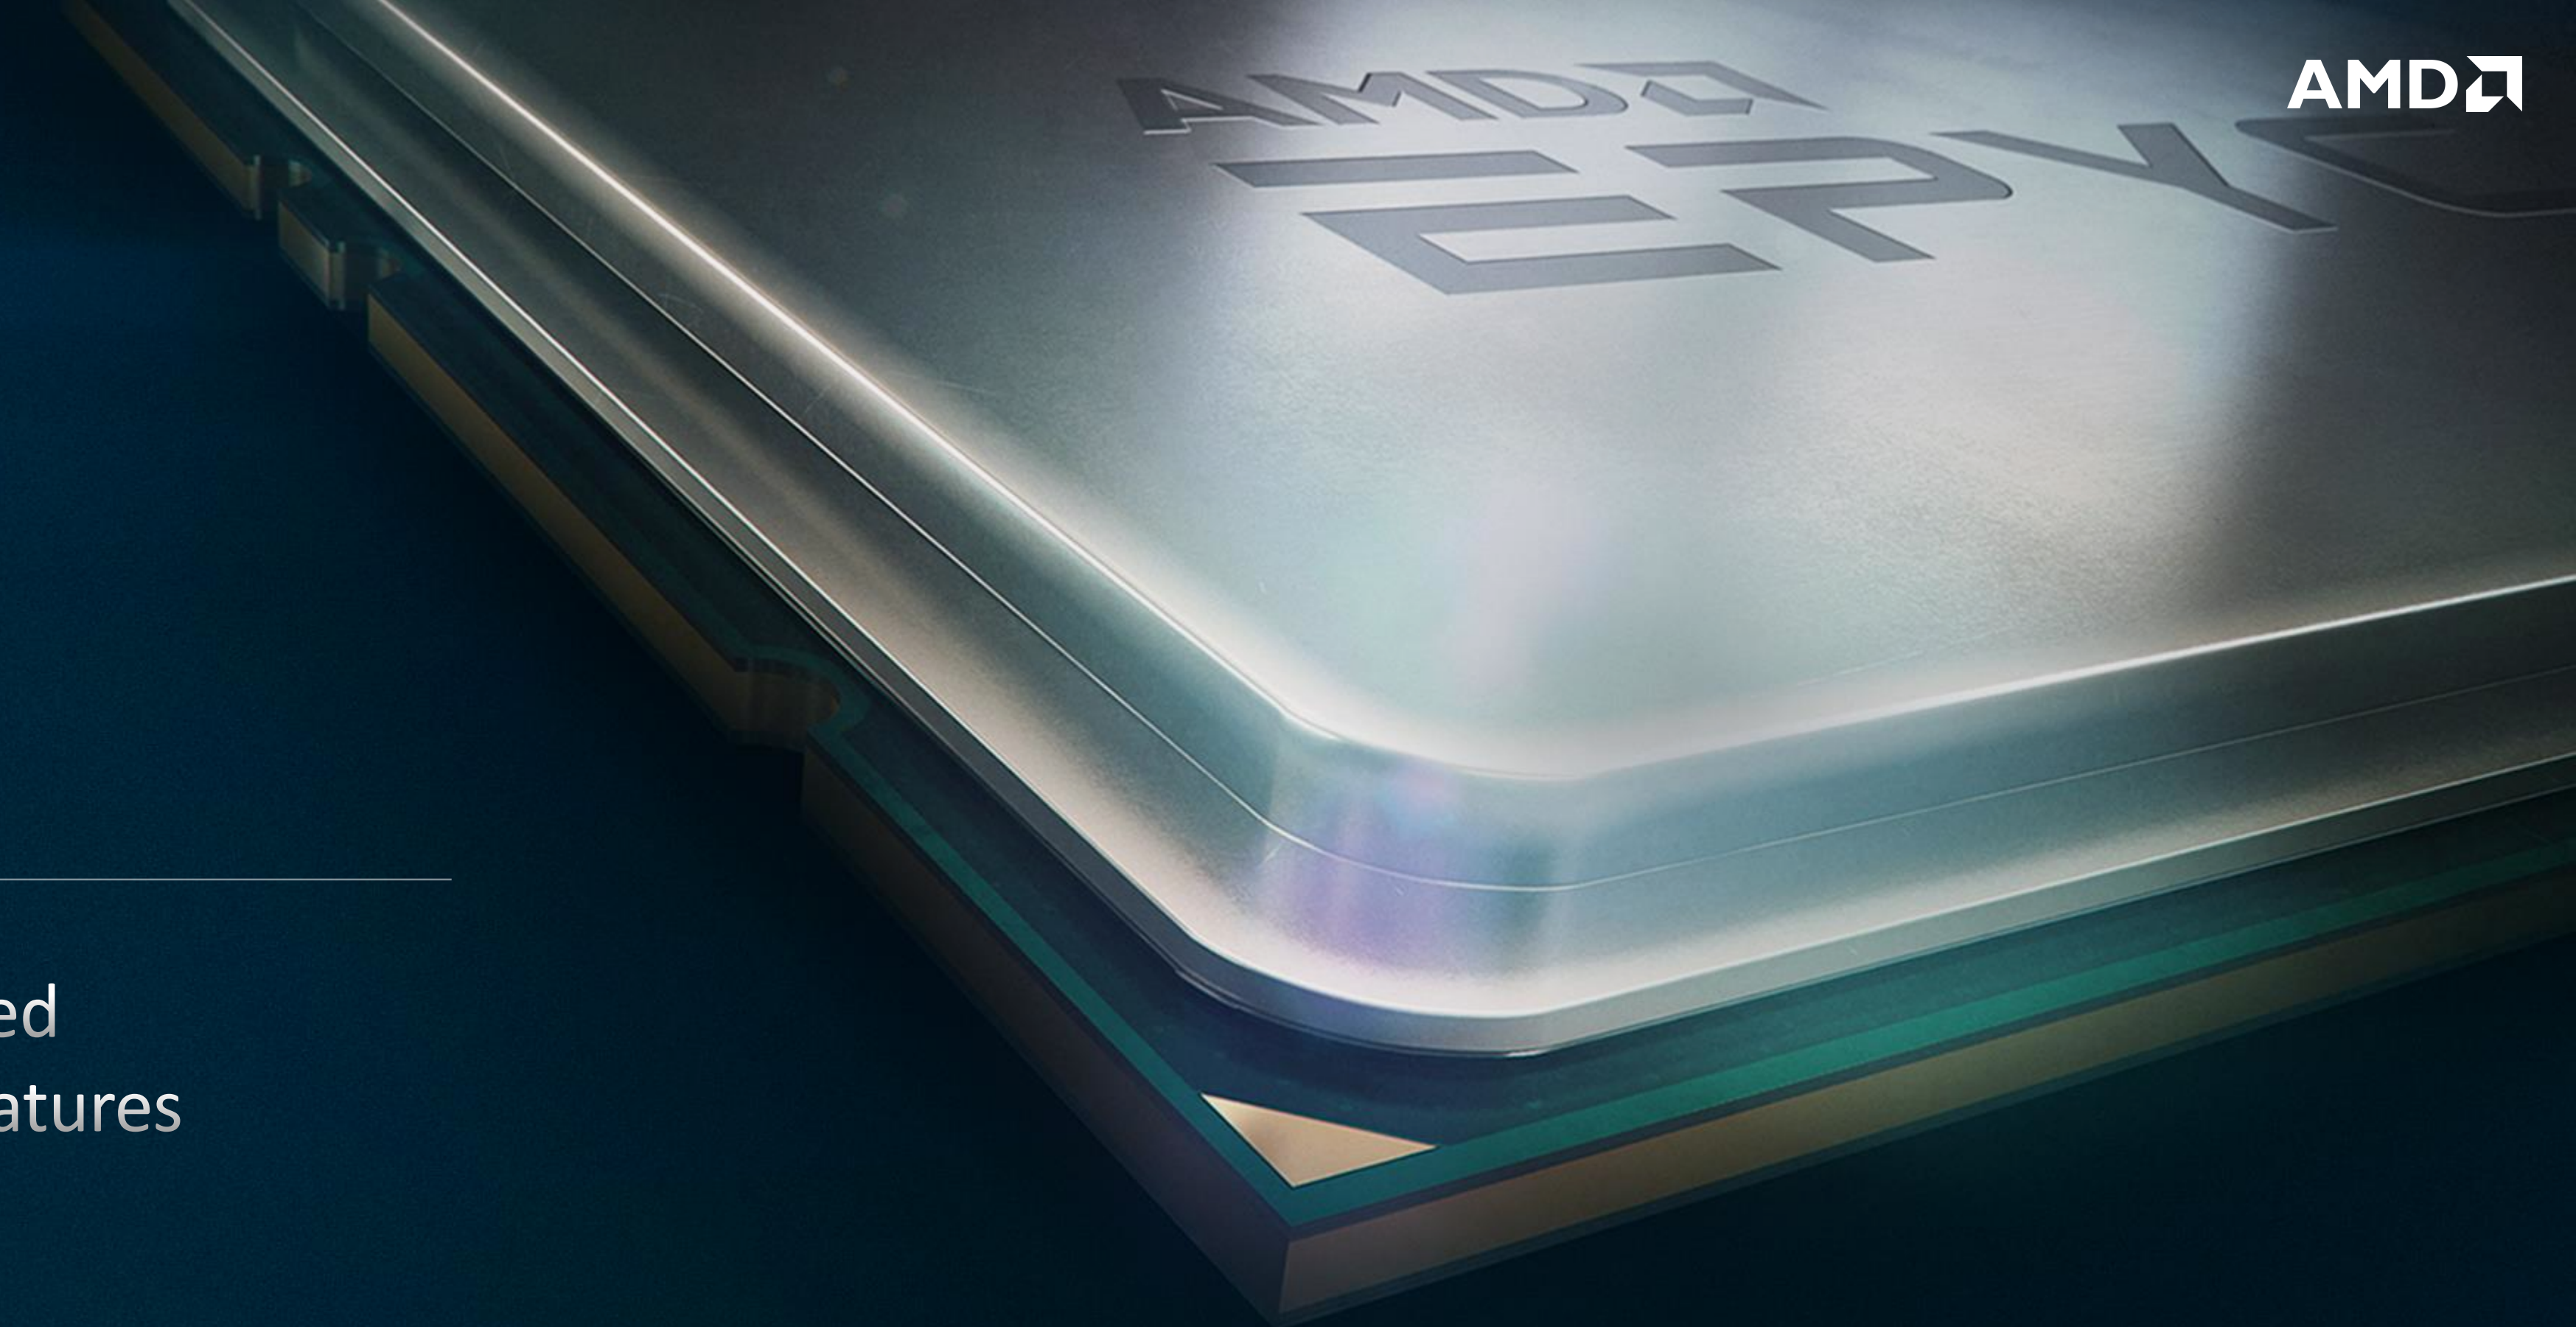
# “ROME”

## AMD NEXT GENERATION EPYC™ CPU

World's First  
7nm Datacenter CPU

Increased IPC and Significant  
Total Performance Uplift

Advanced  
Security Features





# TOTAL DATACENTER COMMITMENT

# ENDNOTES

Slide 8

**Claim:** Datacenter TAM will be \$29M by 2021.

**Substantiation:** AMD internal estimate for 2021. Forward looking statement – please see cautionary statement.

# ENDNOTES

## Slide 10

**Claim:** Performance leadership at every competitive price point.

**Substantiation:** NAP-87 – Estimates based on SPECrate®2017\_int\_base using the GCC-02 v7.2 compiler. AMD-based system scored 196 in tests conducted in AMD labs using an “Ethanol” reference platform configured with 2 x AMD EPYC 7601 SOC’s, 512GB memory (16 x 32GB 2R DDR4 2666MHz), Ubuntu 17.04, BIOS 1002E. Intel-based Supermicro SYS-1029U-TRTP server scored 169.8 in tests conducted in AMD labs configured with 2 x Xeon 8160 CPU’s, 768GB memory (24 x 32GB 2R DDR4 2666MHz), SLES 12 SP3 4.4.92-6.18-default kernel, BIOS set to Extreme performance setting.

NAP-88 – Estimates based on SPECrate®2017\_int\_base using the GCC-02 v7.2 compiler. AMD-based system scored 149 in tests conducted in AMD labs using an “Ethanol” reference platform configured with 2 x AMD EPYC 7401 SOC’s, 512GB memory (16 x 32GB 2R DDR4 2666MHz), Ubuntu 17.04, BIOS 1002E. Intel-based Supermicro SYS-1029U-TRTP server scored 118.1 in tests conducted in AMD labs configured with 2 x Xeon 6130 CPU’s, 768GB memory (24 x 32GB 2R DDR4 2666MHz), SLES 12 SP3 4.4.92-6.18-default kernel, BIOS set to Extreme performance setting.

NAP-89 – Estimates based on SPECrate®2017\_int\_base using the GCC-02 v7.2 compiler. AMD-based system scored 123 in tests conducted in AMD labs using an “Ethanol” reference platform configured with 2 x AMD EPYC 7351 SOC’s, 512GB memory (16 x 32GB 2R DDR4 2666MHz), Ubuntu 17.04, BIOS 1002E. Intel-based Supermicro SYS-1029U-TRTP server scored 86.2 in tests conducted in AMD labs configured with 2 x Xeon 5118 CPU’s, 768GB memory (24 x 32GB 2R DDR4 2666MHz running at 2400), SLES 12 SP3 4.4.92-6.18-default kernel, BIOS set to Extreme performance setting.

NAP-90 – Estimates based on SPECrate®2017\_int\_base using the GCC-02 v7.2 compiler. AMD-based system scored 113 in tests conducted in AMD labs using an “Ethanol” reference platform configured with 2 x AMD EPYC 7351 SOC’s, 512GB memory (16 x 32GB 2R DDR4 2666MHz), Ubuntu 17.04, BIOS 1002E. Intel-based Supermicro SYS-1029U-TRTP server scored 78.7 in tests conducted in AMD labs configured with 2 x Xeon 4116 CPU’s, 768GB memory (24 x 32GB 2R DDR4 2666MHz running at 2400), SLES 12 SP3 4.4.92-6.18-default kernel, BIOS set to Extreme performance setting.

NAP-98 – Based on SPECrate®2017\_int\_peak results published on [www.spec.org](http://www.spec.org) as of April 2018. AMD-based system scored 310 on a Supermicro A+ Server 2123BT-HNC0R configured with 2 x AMD EPYC 7601 SOC’s (\$4200 each at AMD 1ku pricing), 1TB memory (16 x 64GB DDR4 2666MHz), SUSE 12 SP3, Supermicro BIOS 1.0b, using the AOCC 1.0 compiler. Intel-based system scored 309 on a Cisco UCS C220 M5 server configured with 2 x 8180M CPU’s (\$13,011 each per ark.intel.com), 384GB memory (24\*16GB 2R DDR4 2666MHz), SLES 12 SP2, BIOS v3.2.1d, using the ICC 18.0.0.128 compiler.

NAP-99 – Based on SPECrate®2017\_fp\_peak results published on [www.spec.org](http://www.spec.org) as of April 2018. AMD-based system scored 279 on a Supermicro A+ Server 4023S-TRT configured with 2 x AMD EPYC 7601 SOC’s (\$4200 each at AMD 1ku pricing), 1TB memory (16 x 64GB DDR4 2666MHz), SUSE 12 SP3, Supermicro BIOS 1.0b, using the AOCC 1.0 compiler. Intel-based system scored 250 on a Cisco UCS C220 M5 server configured with 2 x 8180M CPU’s (\$13,011 each per ark.intel.com), 384GB memory (24\*16GB 2R DDR4 2666MHz), SLES 12 SP2, BIOS v3.2.1d, using the ICC 18.0.0.128 compiler

# ENDNOTES

## Slide 10

**Claim:** More cores, more memory bandwidth, more I/O.

**Substantiation:** NAP-42 – AMD EPYC™ 7601 processor supports up to 8 channels of DDR4-2667, versus the Xeon Platinum 8180 processor at 6 channels of DDR4-2667.

NAP-43 – AMD EPYC 7601 processor includes up to 32 CPU cores versus the Xeon Platinum 8180 processor with 28 CPU cores.

NAP-44 – A single AMD EPYC™ 7601 processor offers up to 2TB/processor (x 2 = 4TB), versus a single Xeon Platinum 8180 processor at 768Gb/processor (x 2 = 1.54TB).

NAP-56 – AMD EPYC™ processor supports up to 128 PCIe® Gen 3 I/O lanes (in both 1 and 2-socket configuration), versus the Intel® Xeon® SP Series processor supporting a maximum of 48 lanes PCIe® Gen 3 per CPU, plus 20 lanes in the chipset (max of 68 lanes on 1 socket and 116 lanes on 2 socket).

**Claim:** Unmatched total cost of ownership.

**Substantiation:** Dell EMC PowerEdge R7415 AMD EPYC VMware vSAN Mixed Workloads Performance: [https://www.demartek.com/Reports\\_Free/Demartek\\_Dell\\_AMD\\_vSAN\\_Mixed\\_Workloads\\_Performance\\_Summary\\_Report\\_2018-04.pdf](https://www.demartek.com/Reports_Free/Demartek_Dell_AMD_vSAN_Mixed_Workloads_Performance_Summary_Report_2018-04.pdf).

Results not independently verified by AMD.

# ENDNOTES

Slide 17

Claim: Exceptional Performance-per-Dollar for General Purpose and Memory Optimized Workloads

**Substantiation:** Source: Amazon Web Services Press Release “AWS Introduces New Amazon EC2 Instances Featuring AMD EYPC Processors at LINK (actual cost savings with AMD may vary)

Slide 18

**Claim:** Exceptional Performance-per-Dollar for General Purpose and Memory Optimized Workloads

**Substantiation:** Source: Amazon Web Services Press Release “AWS Introduces New Amazon EC2 Instances Featuring AMD EYPC Processors at LINK (actual cost savings with AMD may vary)

**Claim:** EPYC delivers seamless application migration for Amazon Elastic Compute Cloud (EC2).

**Substantiation:** Source: Amazon Web Services Press Release “AWS Introduces New Amazon EC2 Instances Featuring AMD EYPC Processors at LINK (actual cost savings with AMD may vary)

# ENDNOTES

## Slide 20

**Claim:** Modular system design extends compute, memory bandwidth and I/O leadership.

**Substantiation:** NAP-42 – AMD EPYC™ 7601 processor supports up to 8 channels of DDR4-2667, versus the Xeon Platinum 8180 processor at 6 channels of DDR4-2667.

NAP-43 – AMD EPYC 7601 processor includes up to 32 CPU cores versus the Xeon Platinum 8180 processor with 28 CPU cores.

NAP-44 – A single AMD EPYC™ 7601 processor offers up to 2TB/processor (x 2 = 4TB), versus a single Xeon Platinum 8180 processor at 768Gb/processor (x 2 = 1.54TB).

NAP-56 – AMD EPYC™ processor supports up to 128 PCIe® Gen 3 I/O lanes (in both 1 and 2-socket configuration), versus the Intel® Xeon® SP Series processor supporting a maximum of 48 lanes PCIe® Gen 3 per CPU, plus 20 lanes in the chipset (max of 68 lanes on 1 socket and 116 lanes on 2 socket). Based on Zen2 design parameters versus Zen1 and currently shipping products – core count increase from 32 to up to 64 per socket. Memory bandwidth with “Zen 2” design parameters including increased memory speed across eight memory channels, I/O leadership extending to PCIeGen4”

## Slide 22

**Claim:** Radeon Instinct MI60 has the world’s fastest HPC PCIe capable accelerator.

**Substantiation:** Calculated on Oct 22, 2018, the Radeon Instinct MI60 GPU resulted in 7.4 TFLOPS peak theoretical double precision floating-point (FP64) performance. AMD TFLOPS calculations conducted with the following equation: FLOPS calculations are performed by taking the engine clock from the highest DPM state and multiplying it by xx CUs per GPU. Then, multiplying that number by xx stream processors, which exist in each CU. Then, that number is multiplied by 1/2 FLOPS per clock for FP64. TFLOP calculations for MI60 can be found at

<https://www.amd.com/en/products/professional-graphics/instinct-mi60> External results on the NVidia Tesla V100 (16GB card) GPU accelerator resulted in 7 TFLOPS peak double precision (FP64) floating-point performance. Results found at:

<https://images.nvidia.com/content/technologies/volta/pdf/437317-Volta-V100-DS-NV-US-WEB.pdf> AMD has not independently tested or verified external/third party results/data and bears no responsibility for any errors or omissions therein. RIV-3

# ENDNOTES

## Slide 23

**Claim:** “Rome” is the world’s first 7nm datacenter CPU.

**Substantiation:** Based on AMD’s internal competitive assessment of all currently shipping or publicly disclosed X86 Server CPU products as of 9/6/2018

**Claim:** “Rome” delivers increased IPC and significant total performance uplift.

**Estimated increase in instructions per cycle (IPC) is based on AMD internal testing for “Zen 2” across microbenchmarks, compared to prior “Zen 1” generation CPU using combined floating point and integer benchmarks.**

# DISCLAIMER & ATTRIBUTION

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.

Vega, Polaris and Zen are codenames for AMD architecture and are not product names.

## **ATTRIBUTION**

© 2018 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, EPYC, Ryzen, Radeon, and combinations thereof are trademarks of Advanced Micro Devices, Inc. in the United States and/or other jurisdictions.

Microsoft and DirectX are registered trademarks of Microsoft Corporation in the US and other jurisdictions. Xbox is a registered trademark of Microsoft Corporation.

“PlayStation”, the “PS” Family logo, and “PSP” are registered trademarks of Sony Interactive Entertainment Inc.

Other names are for informational purposes only and may be trademarks of their respective owners.

# CAUTIONARY STATEMENT

This presentation contains forward-looking statements concerning Advanced Micro Devices, Inc. (AMD) including, but not limited to the features, functionality, availability, timing, deployment, and expected benefits of AMD future products, including “Zen 2”, “Zen 3” and “Zen 4”; the benefits and expectations of 7nm process technology; and AMD’s technology 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 Quarterly Report on Form 10-Q for the quarter ended September 29, 2018.

The image features a dark, starry space background with a bright sun or star rising over the horizon of a planet, likely Earth, showing blue oceans and white clouds. The AMD logo is rendered in a metallic, blue, 3D style. Below it, the words "NEXT HORIZON" are in a white, bold, sans-serif font.

# AMD

## NEXT HORIZON

MARK PAPERMASTER  
CTO and SVP, Technology & Engineering

# OUR HIGH-PERFORMANCE JOURNEY CONTINUES

---

“Zen” Delivered  
as Promised

“Zen 2” High-Performance  
Leadership

Strong CPU Roadmap  
for the Future



# DELIVERING AS PROMISED



# ZEN IN THE MARKET



# 7nm PROCESS NODE

# LEADING EDGE 7nm PROCESS TECHNOLOGY

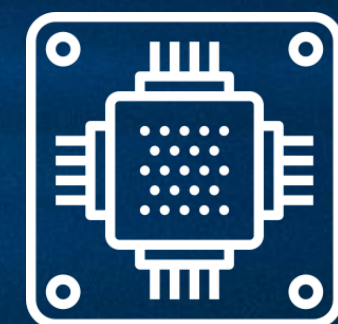
Major Node,  
Significant Investment

Faster, Smaller, Lower  
Power Transistors

Multiple Products in  
Development

Deep Partnerships with TSMC  
and Design Automation Vendors

# 7nm DRIVES SIGNIFICANT COMPUTE EFFICIENCY



DENSITY 2x

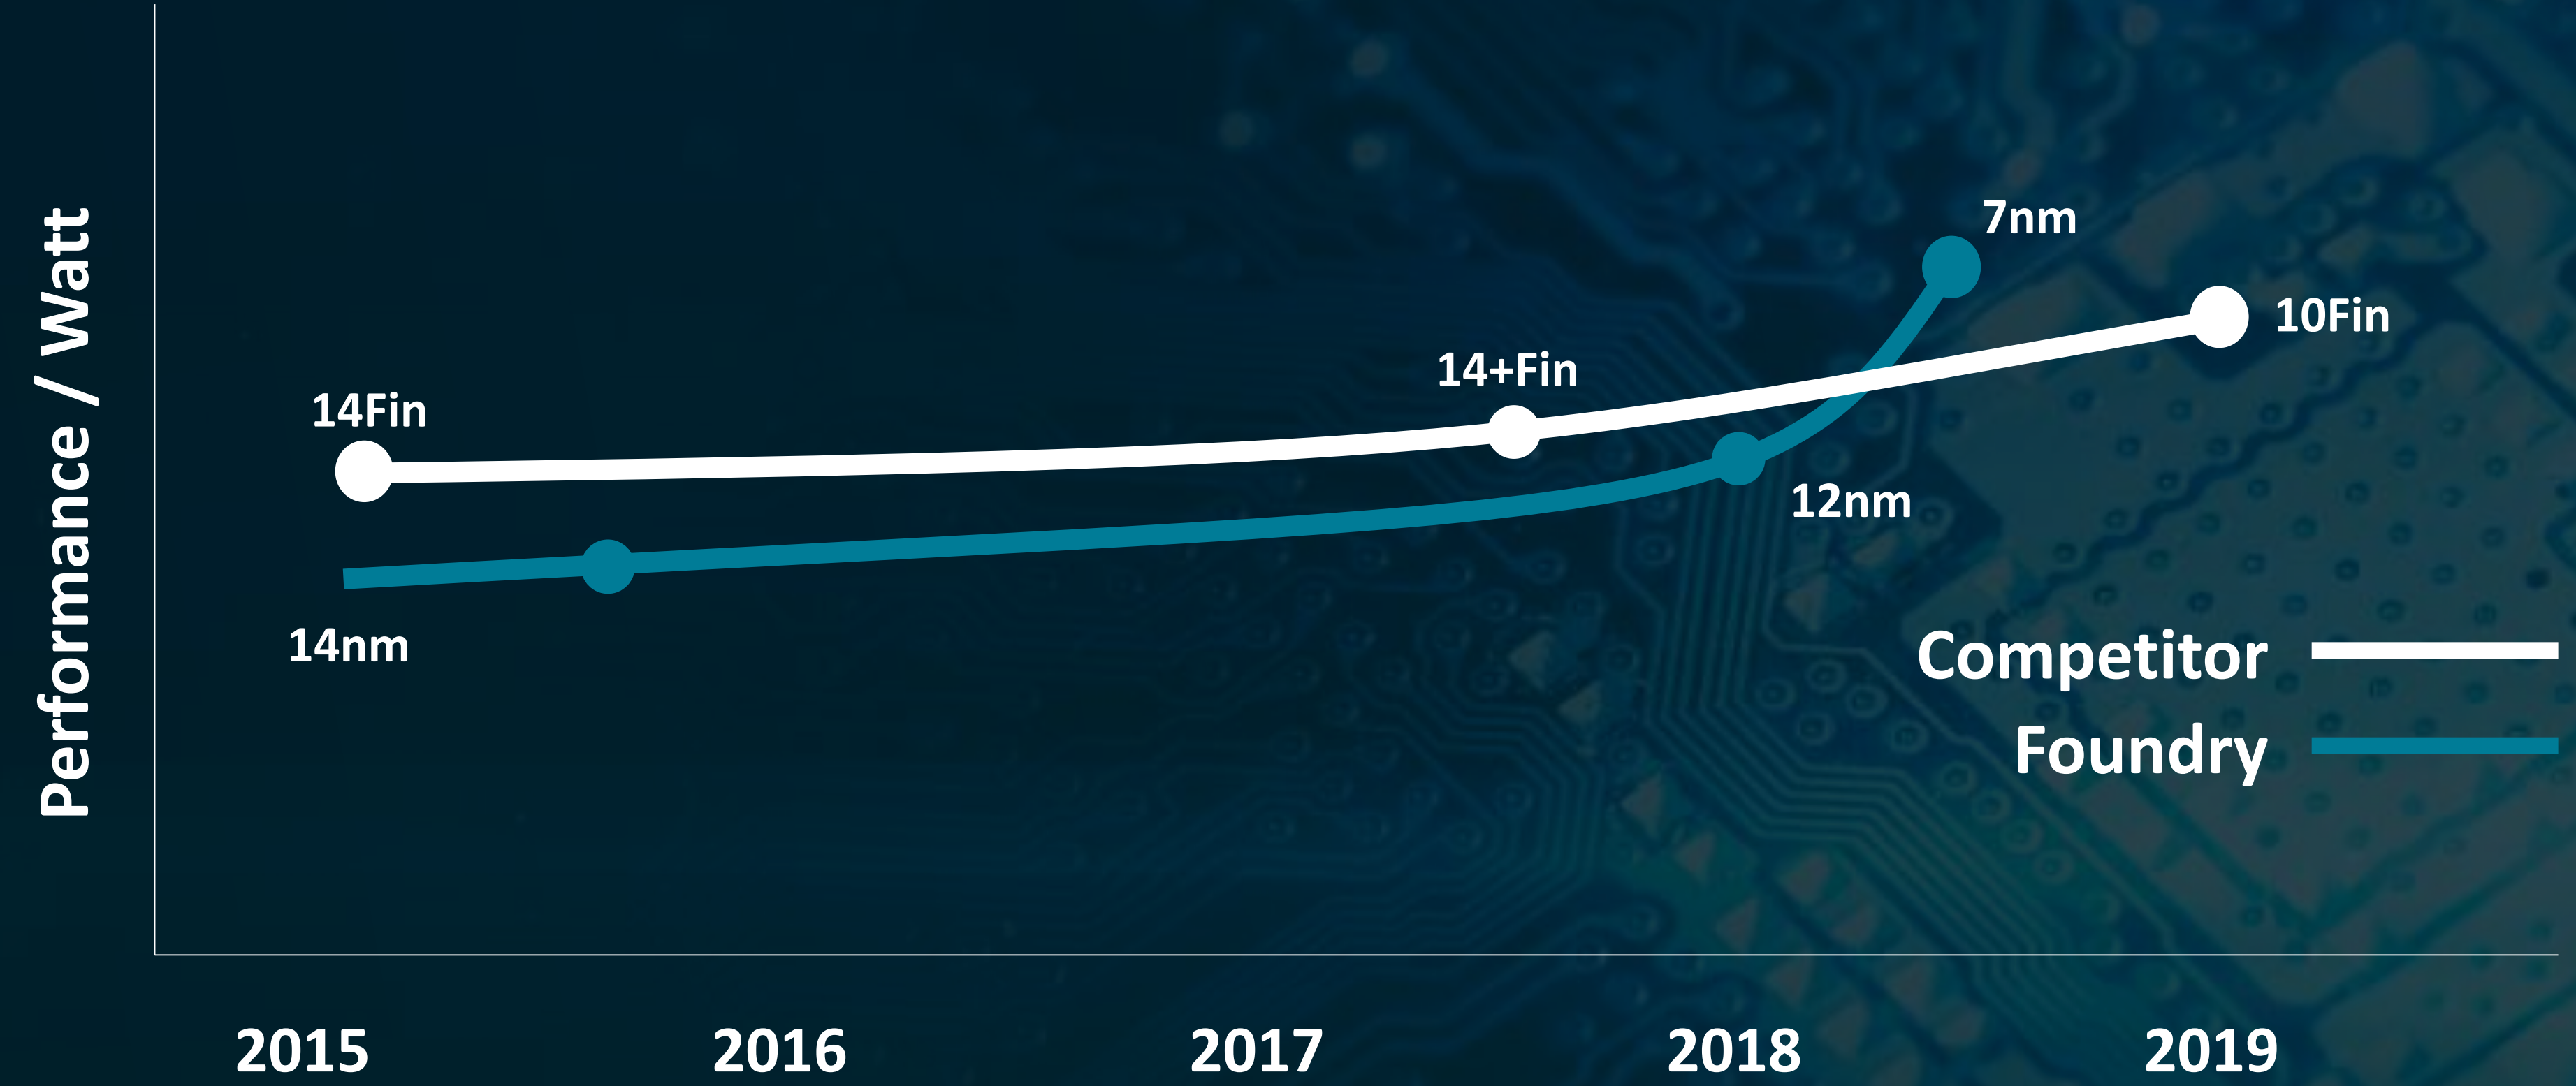


POWER 0.5x  
(Same Performance)



PERFORMANCE >1.25x  
(Same Power)

# 7nm LEVELS COMPETITIVE PLAYING FIELD



# TSMC

Dr. Y.J. Mii, Senior Vice President Research & Development & Technology Development

—





# DESIGN

World's First High-Performance x86  
7nm CPU

CPU Core Execution  
Enhancements

Further Security  
Enhancements

Modular Design for Configurability  
and Ease of Manufacture

# DELIVERS UP TO 2x THROUGHPUT

Improved Execution  
Pipeline

Doubled Floating Point  
and Load Store

Doubled Core  
Density

Half the Energy  
Per Operation



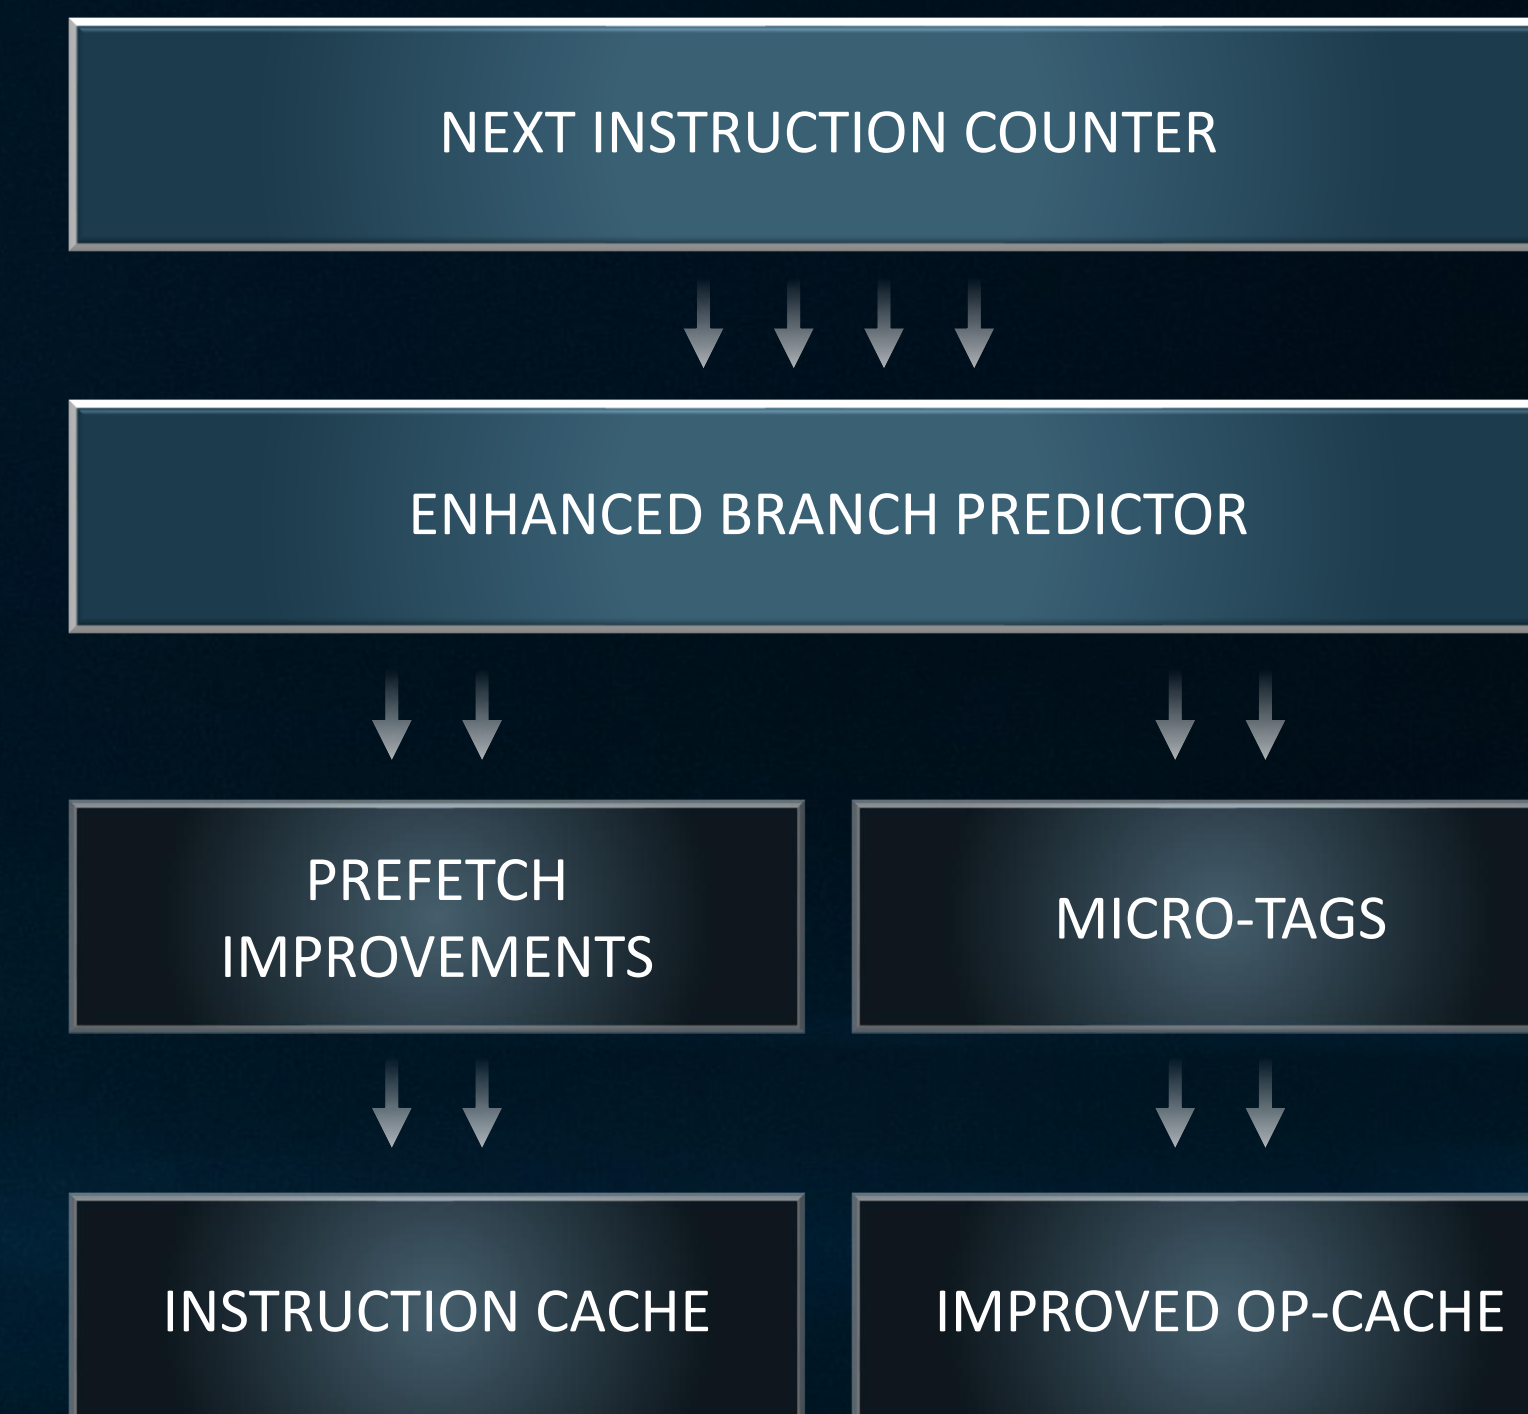
# NEW FRONT END ADVANCES

Improved Branch  
Predictor

Better Instruction  
Pre-Fetching

Re-Optimized  
Instruction Cache

Larger  
Op Cache



# FLOATING POINT ADVANCES

Doubled Floating  
Point Width to  
256-Bit

Doubled  
Load / Store  
Bandwidth

Increased  
Dispatch / Retire  
Bandwidth

Maintained High  
Throughput for All  
Modes



# STRONG SECURITY GETS STRONGER

---

“Zen” Delivered Industry  
Leading Memory Encryption  
with Increased Flexibility

“Zen” Software  
Mitigations Robust

Hardware Enhanced Spectre  
Mitigations with “Zen 2”

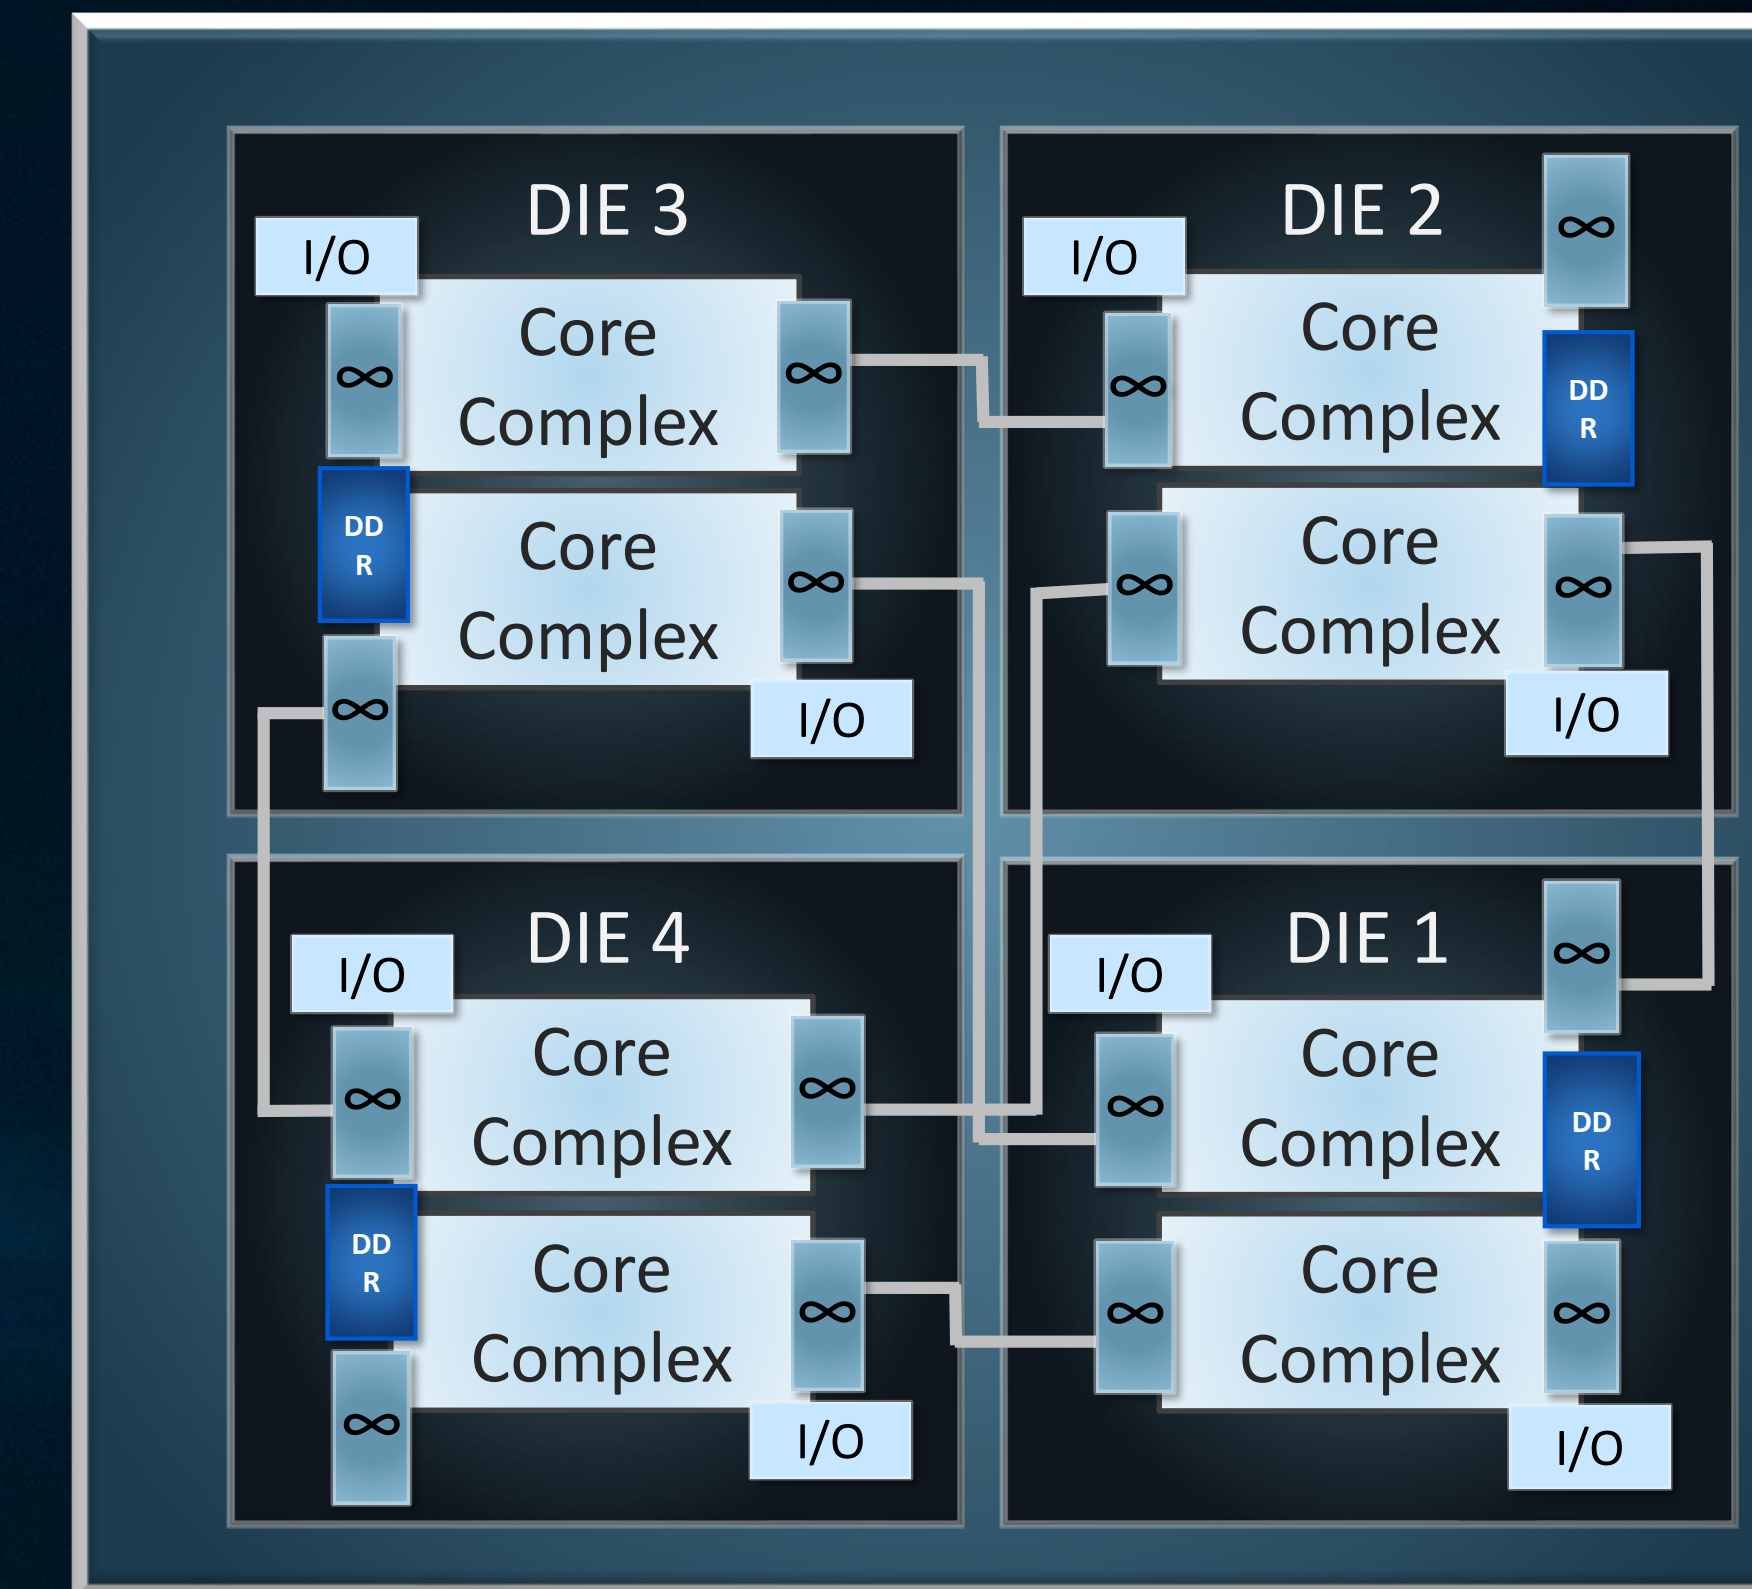
# “ZEN”

## INTRODUCED OUR MULTI-CHIP APPROACH

Inherent Yield  
Advantage

Increased Peak  
Compute

Highly  
Configurable





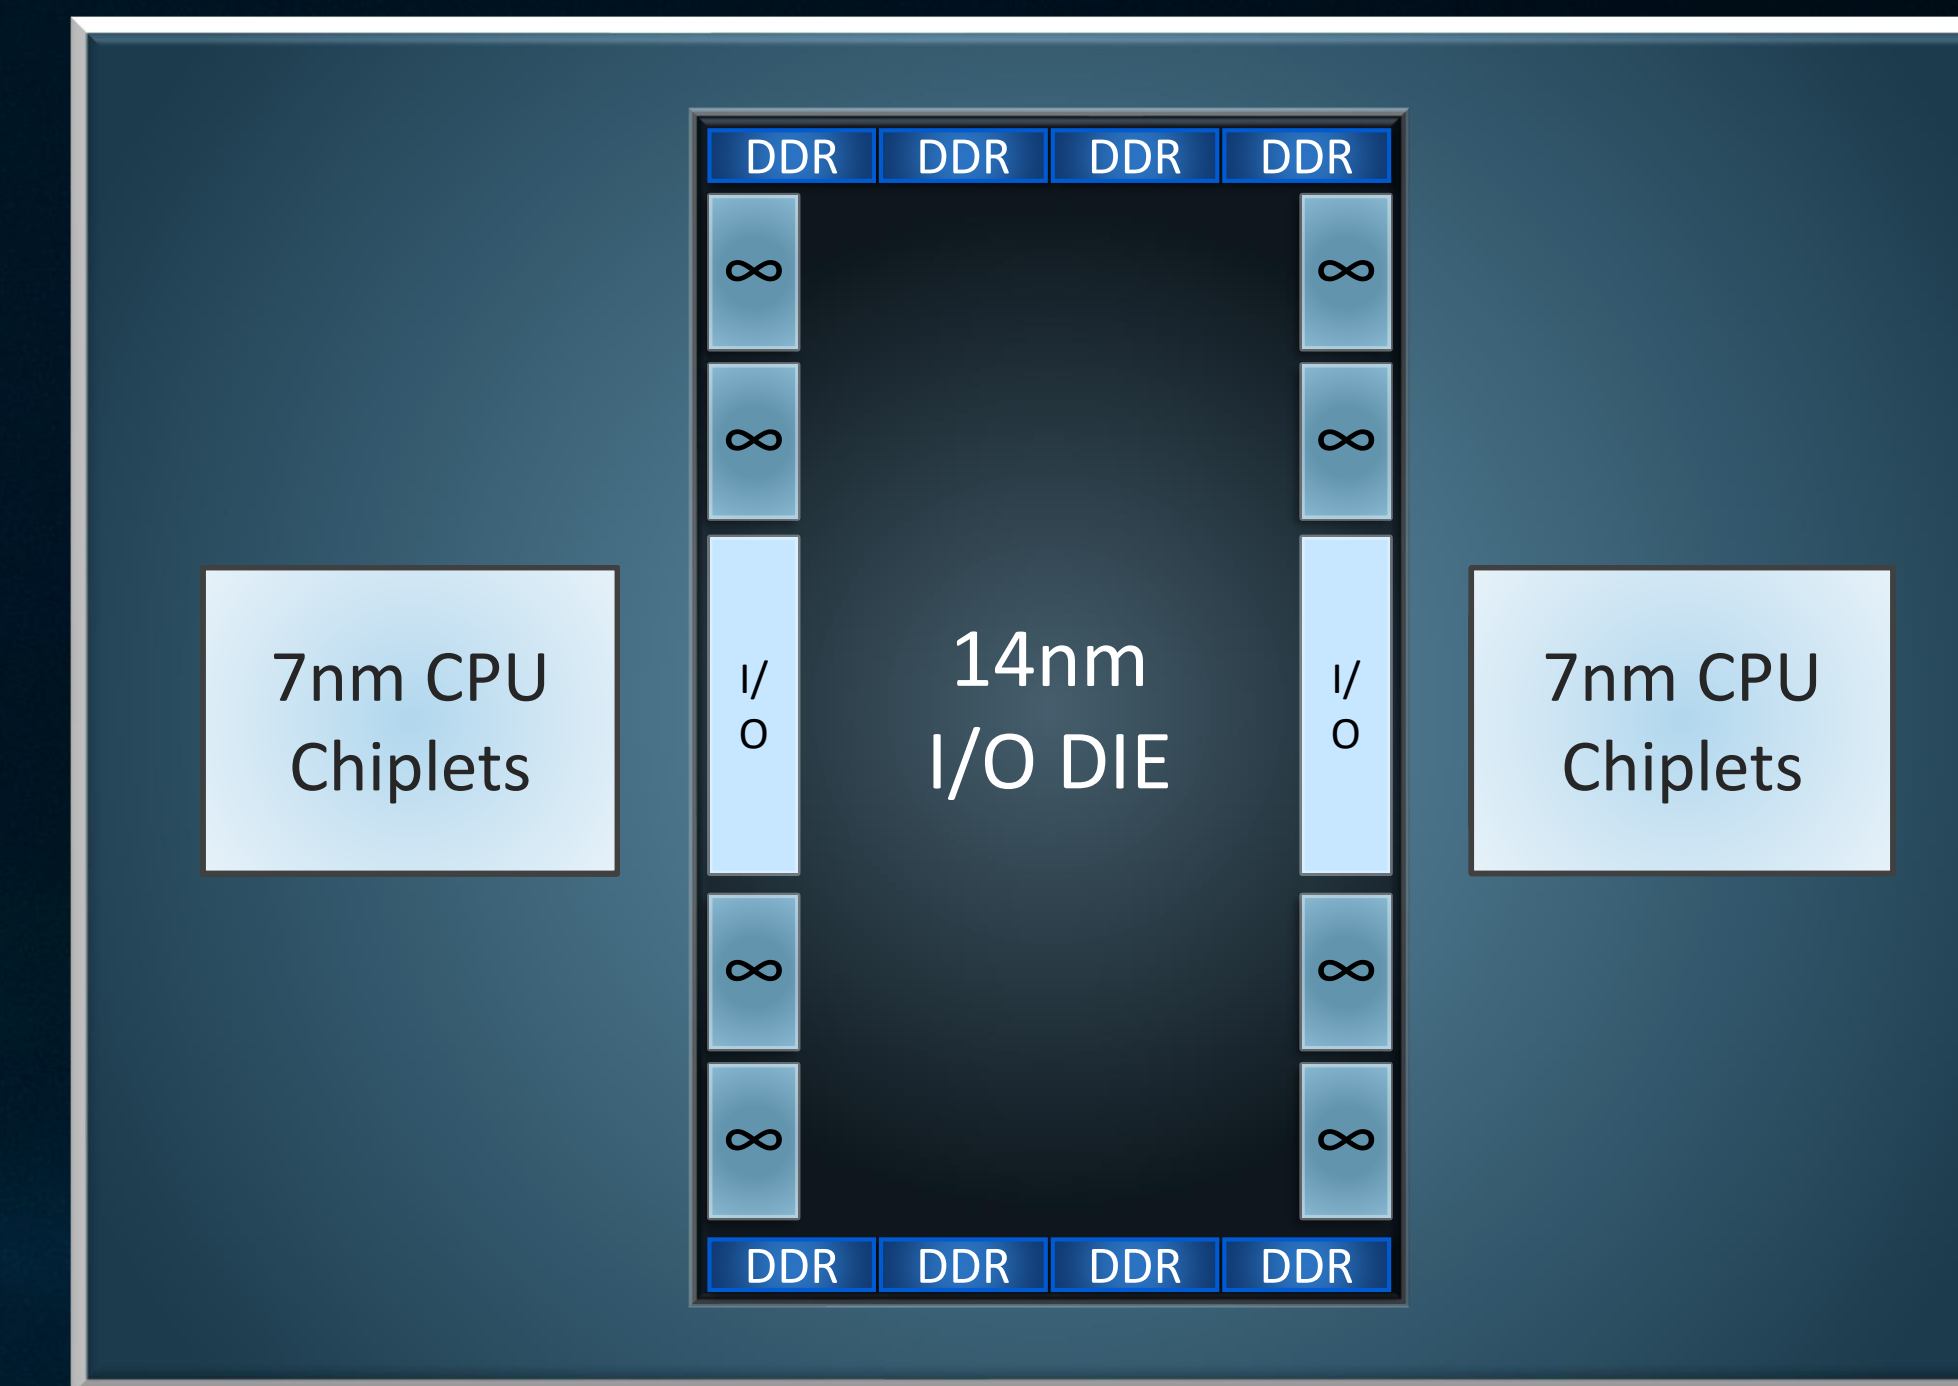
# REVOLUTIONARY CHIPLLET DESIGN

Each IP in its  
Optimal Technology

Infinity Fabric  
Enabled Modularity

Optimized I/O Die  
Improves Latency  
and Power

7nm Tech for  
CPU Performance  
and Power



“Zen 2” Based EPYC Processors

# LEADING-EDGE PERFORMANCE

---

Double Throughput at  
Same Power

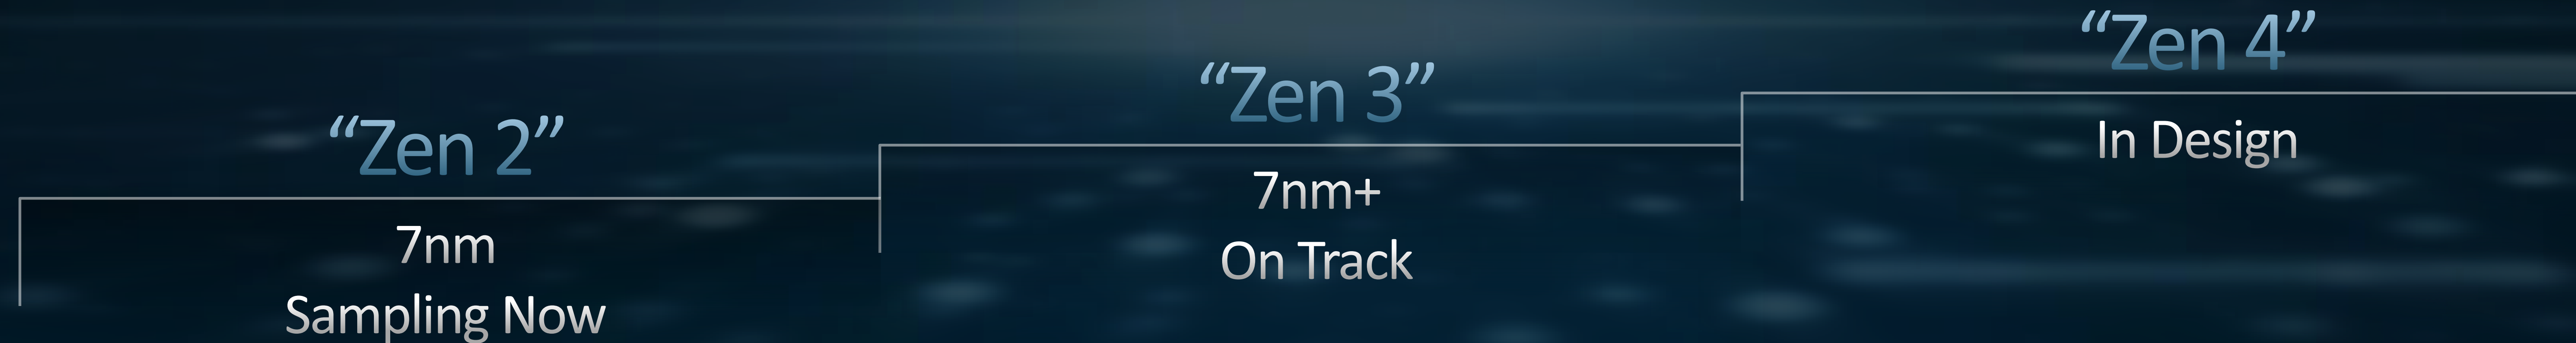
Higher Instructions  
Per Clock

Revolutionary Chiplet  
System Design

Enhanced  
Security



# HIGH-PERFORMANCE MOMENTUM



AMD

NEXT HORIZON

# DISCLAIMERS

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.

## **ATTRIBUTION**

© 2018 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, EPYC, Radeon, and combinations thereof are trademarks of Advanced Micro Devices, Inc. in the United States and/or other jurisdictions. "Rome", "Zen", "Zen 2", and "Naples" are AMD codenames and not intended to be product names.

# CAUTIONARY STATEMENT

This presentation contains forward-looking statements concerning Advanced Micro Devices, Inc. (AMD) including, but not limited to AMD's positioning in the datacenter market; expected TAM for AMD GPUs; the features, functionality, availability, timing, deployment, and expected benefits of AMD future products, including AMD Radeon Instinct MI60; the benefits and expectations of 7nm process technology; and AMD's technology 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 Quarterly Report on Form 10-Q for the quarter ended September 29, 2018.



# AMD

## NEXT HORIZON

The logo features the AMD brand name in a metallic, 3D blue font. Below it, the words "NEXT HORIZON" are written in a white, bold, sans-serif font. The background is a dark space scene with a bright sun or star rising over the horizon of a planet, creating a lens flare effect.

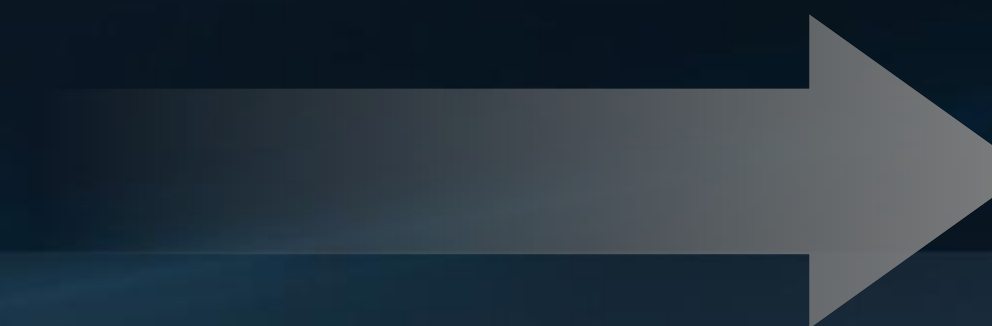
DAVID WANG

Senior Vice President of Engineering,  
Radeon Technologies Group

# AMD RADEON™ GRAPHICS MOVING TO THE DATACENTER



Consumer and Professional



# ADDRESSING EVOLVING DATACENTER GPU USE CASES

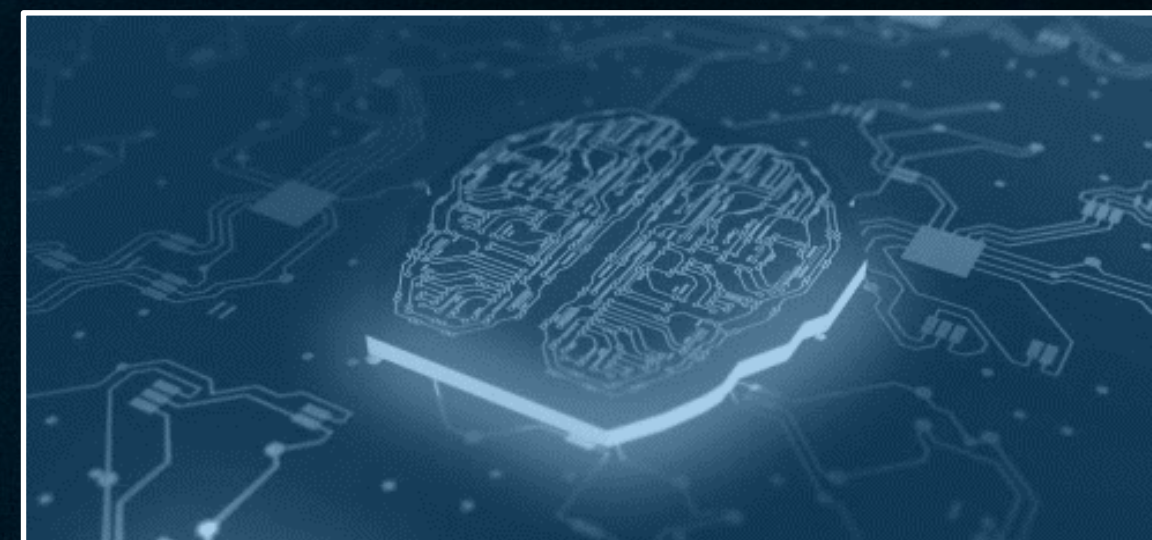
\$12 BILLION COMBINED TAM BY 2021



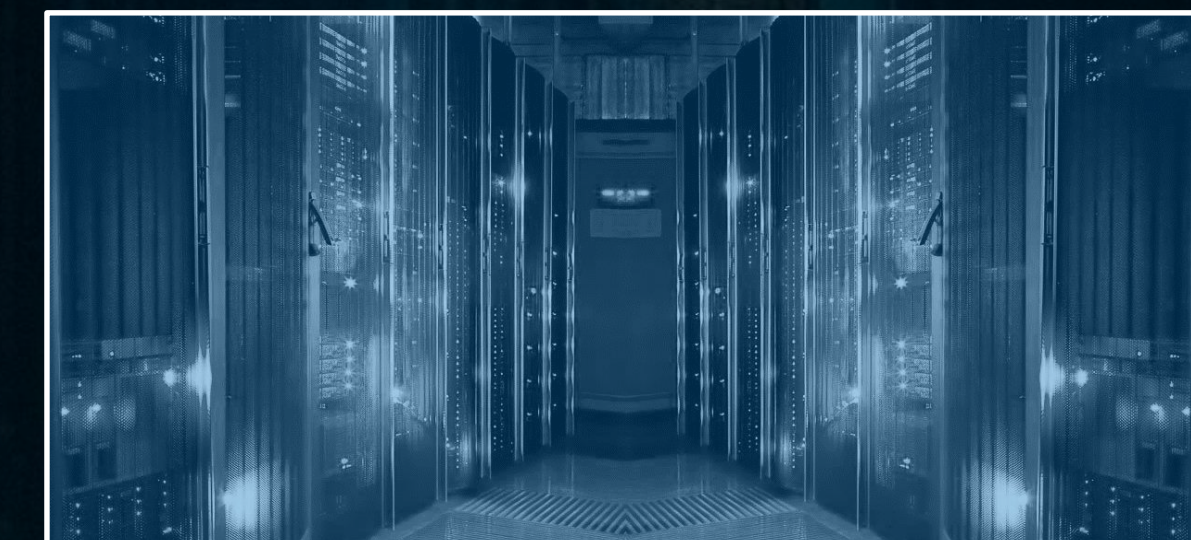
Cloud Gaming



Virtual Desktop & Workstation



Machine Learning



High Performance Computing

DESIGNED FOR THE CLOUD

# AMD RADEON INSTINCT™ MI60

World's First  
7nm GPU

High-Performance,  
Flexible Vega  
Architecture

Leading Edge  
Memory and  
Scalability

Industry's Only  
Hardware-Virtualized  
GPU



# WORLD'S FIRST 7nm GPUs

UP TO

2x MORE  
DENSITY

13.2B TRANSISTORS, 331 mm<sup>2</sup>

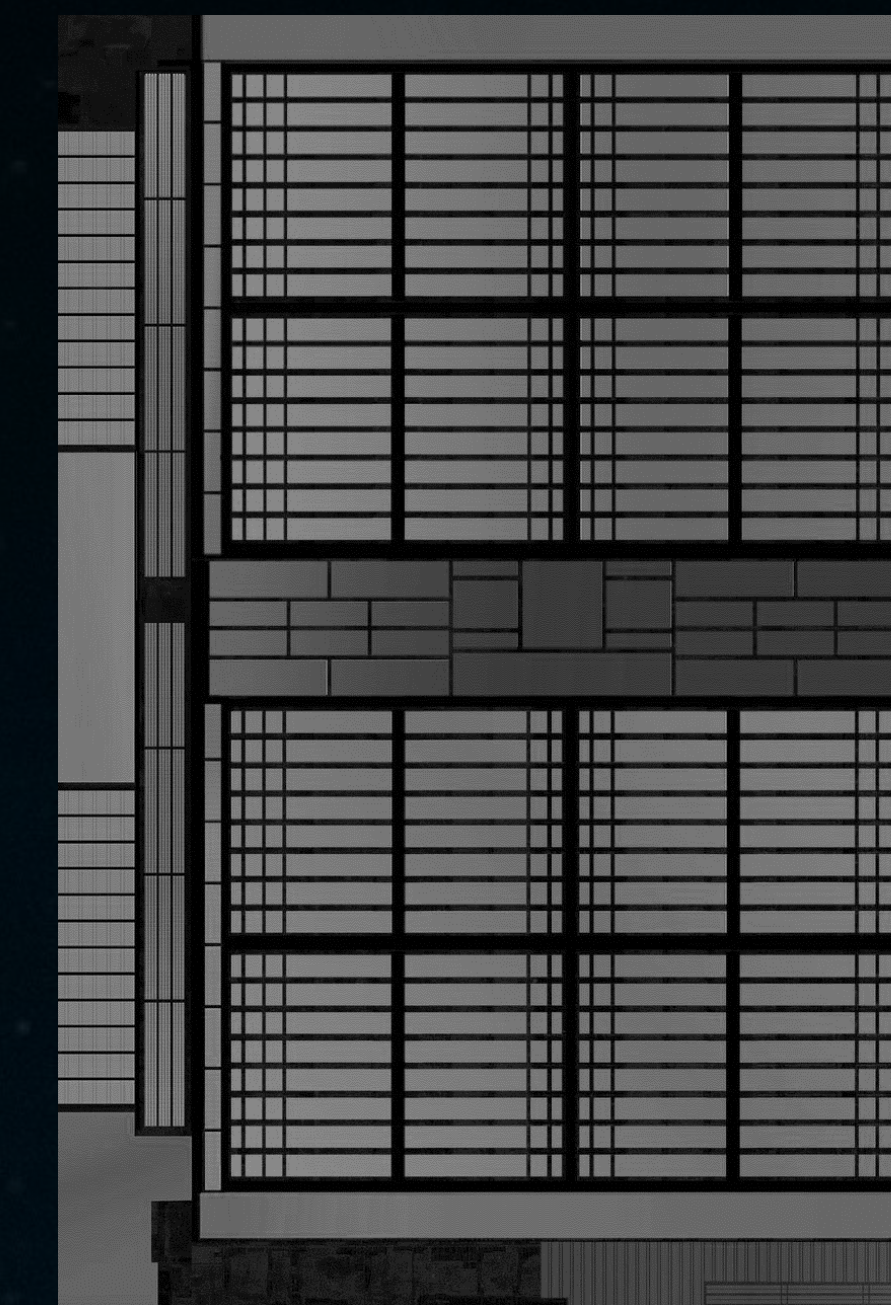
>1.25x HIGHER  
PERFORMANCE

AT THE SAME POWER

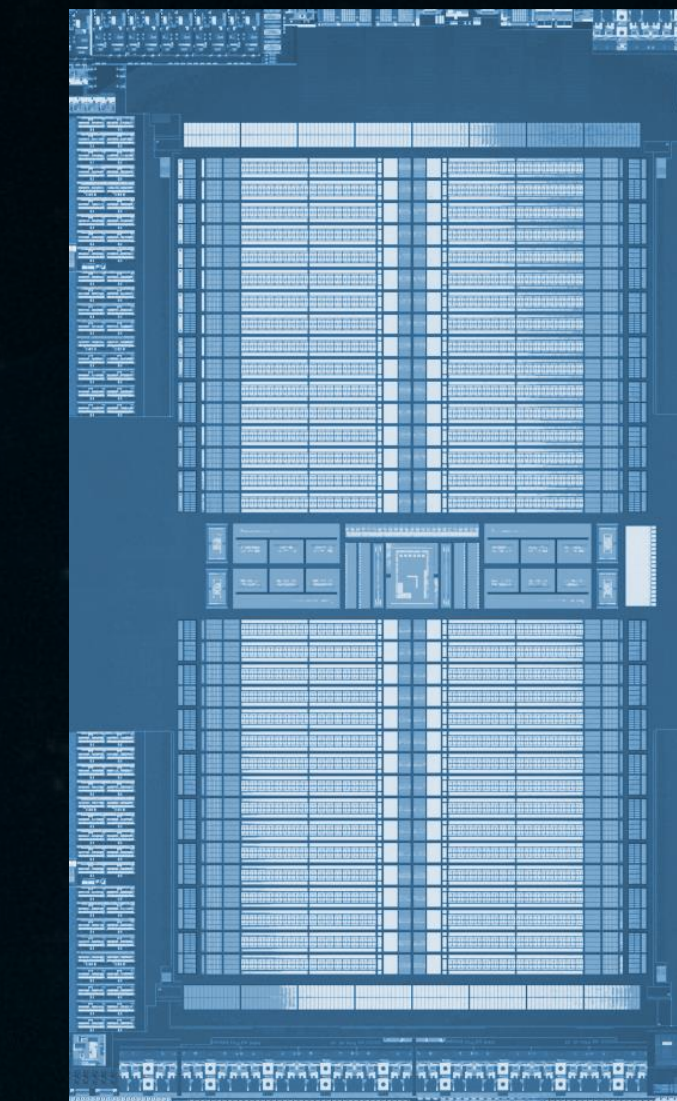
50% LOWER  
POWER

AT THE SAME FREQUENCY

Higher Performance per Watt



MI25



MI60

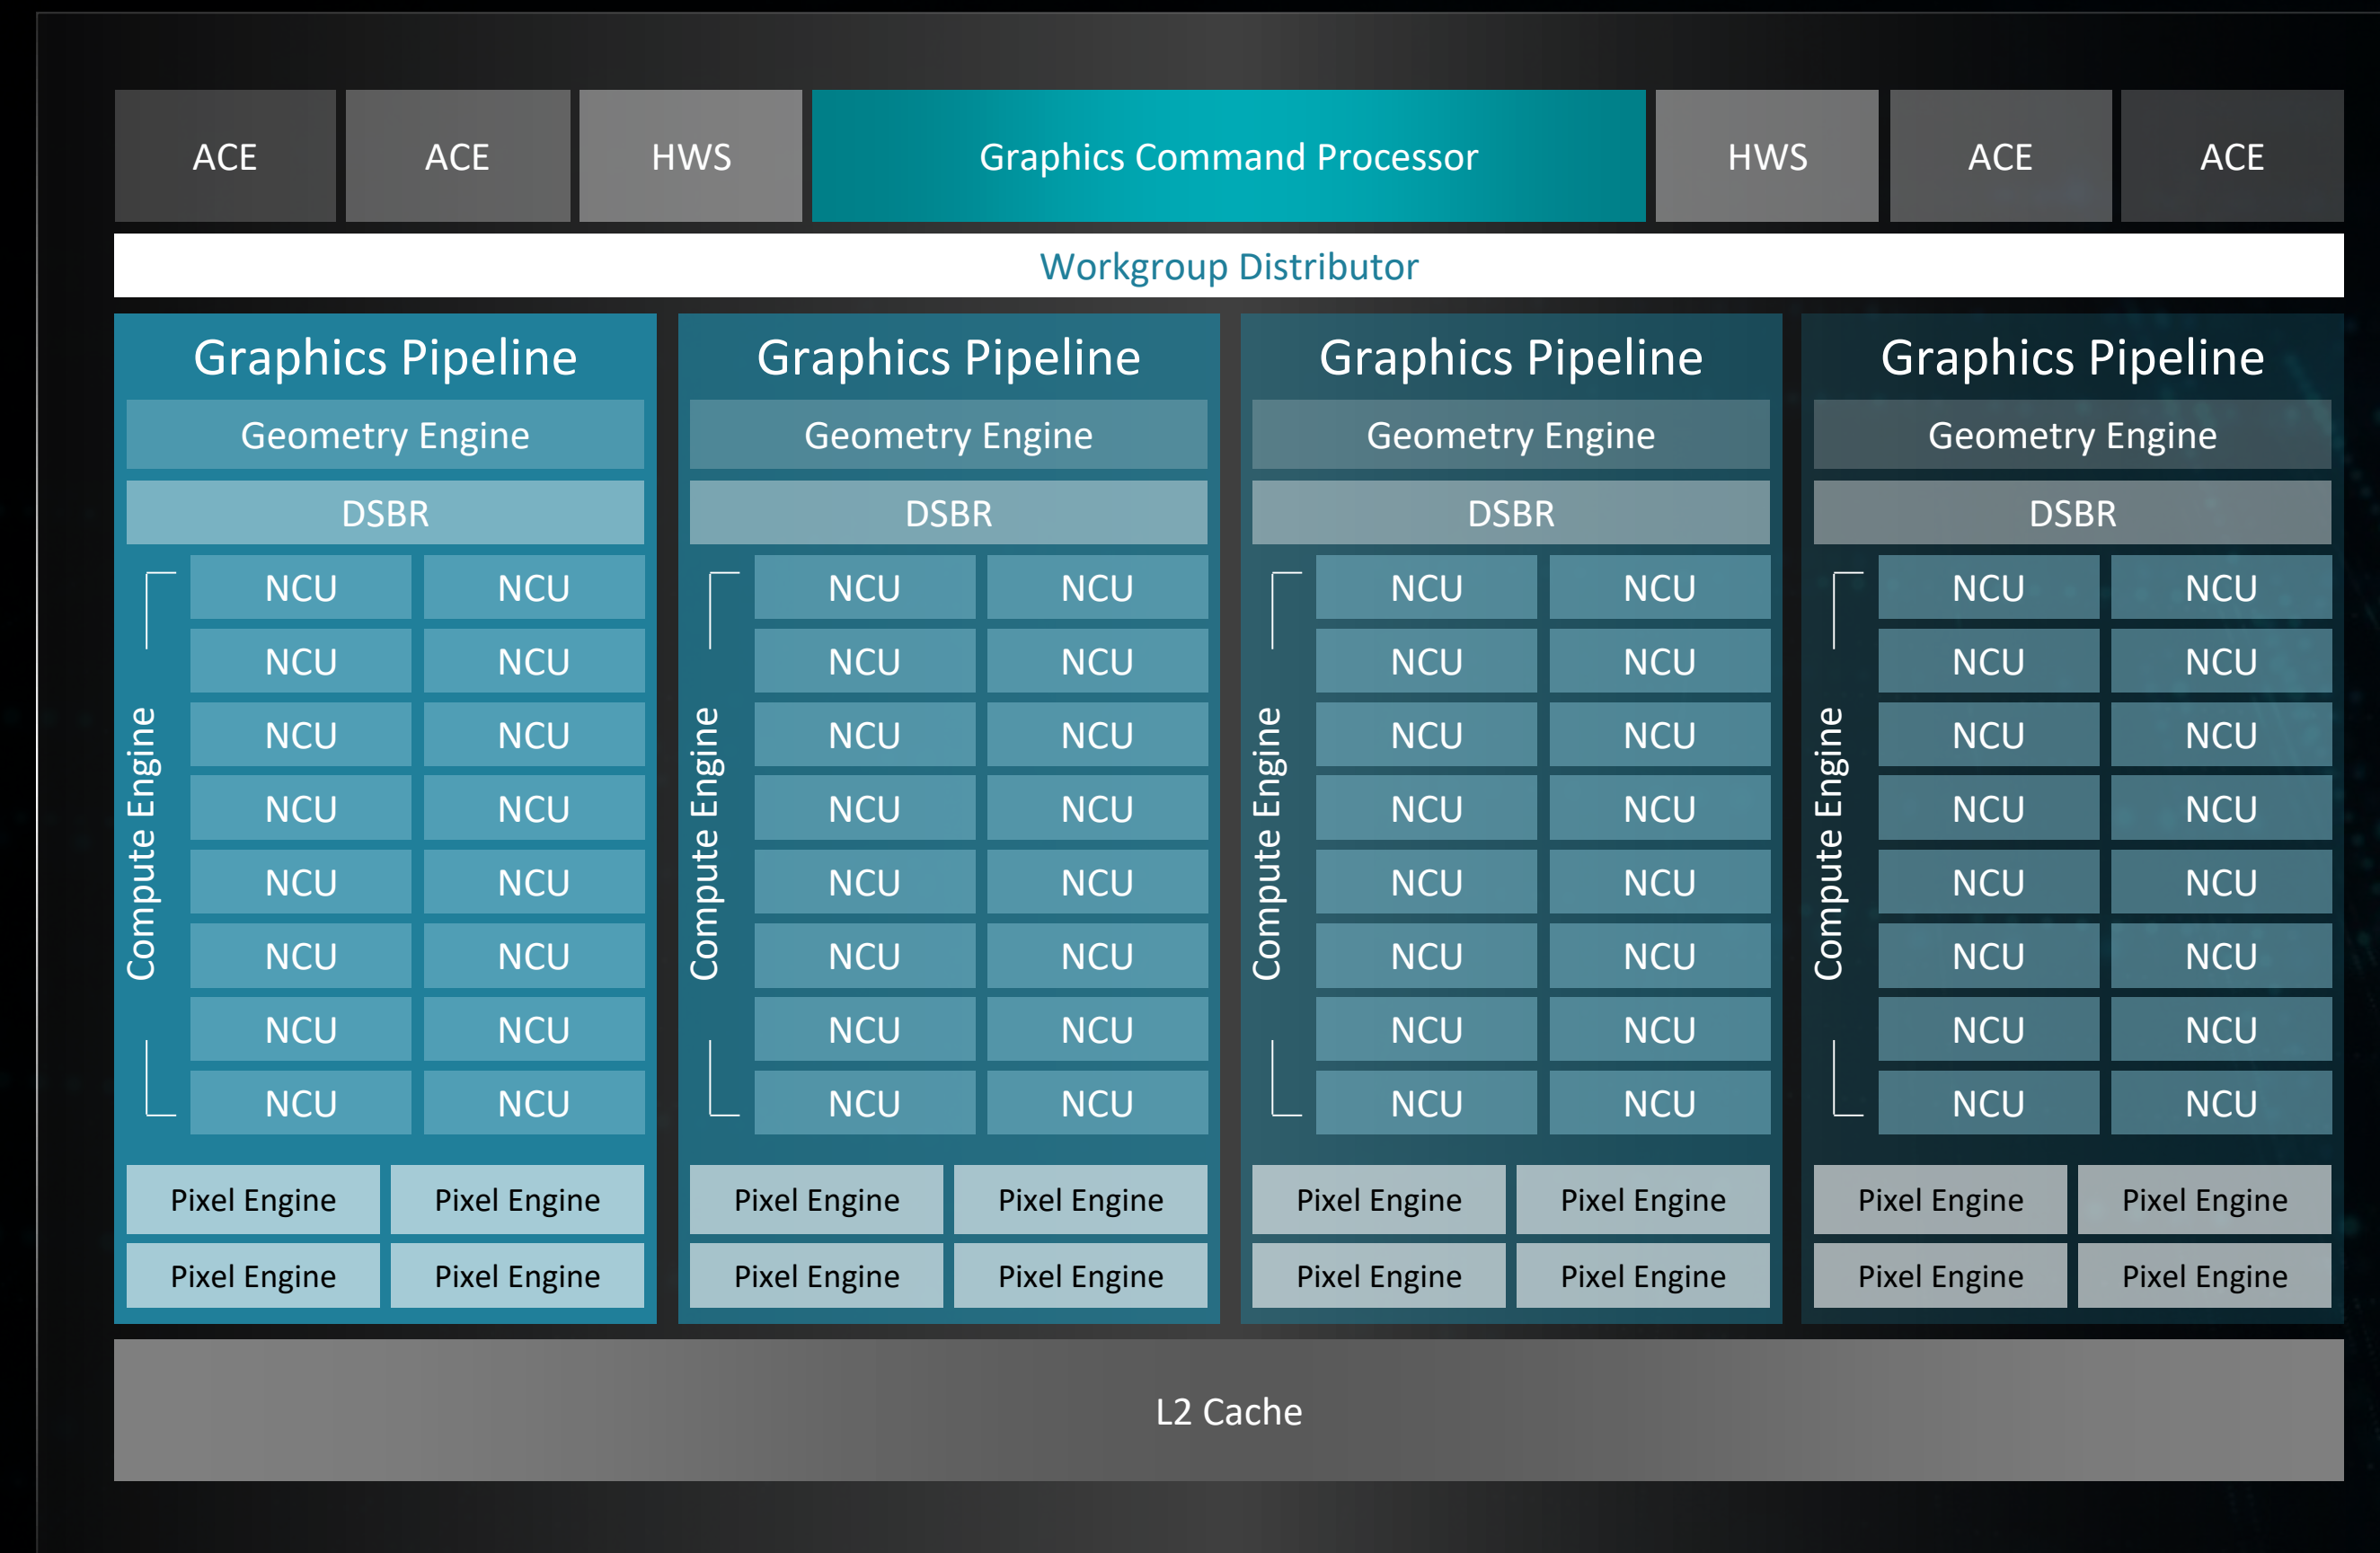
# HIGH-PERFORMANCE, FLEXIBLE VEGA ARCHITECTURE

MI60: World's Fastest  
FP64 and FP32 PCIe  
Capable GPU

Machine Learning  
Operations for  
Training and Inference

Flexible Architecture  
for Different  
Workloads

End-to-End  
ECC  
Protection

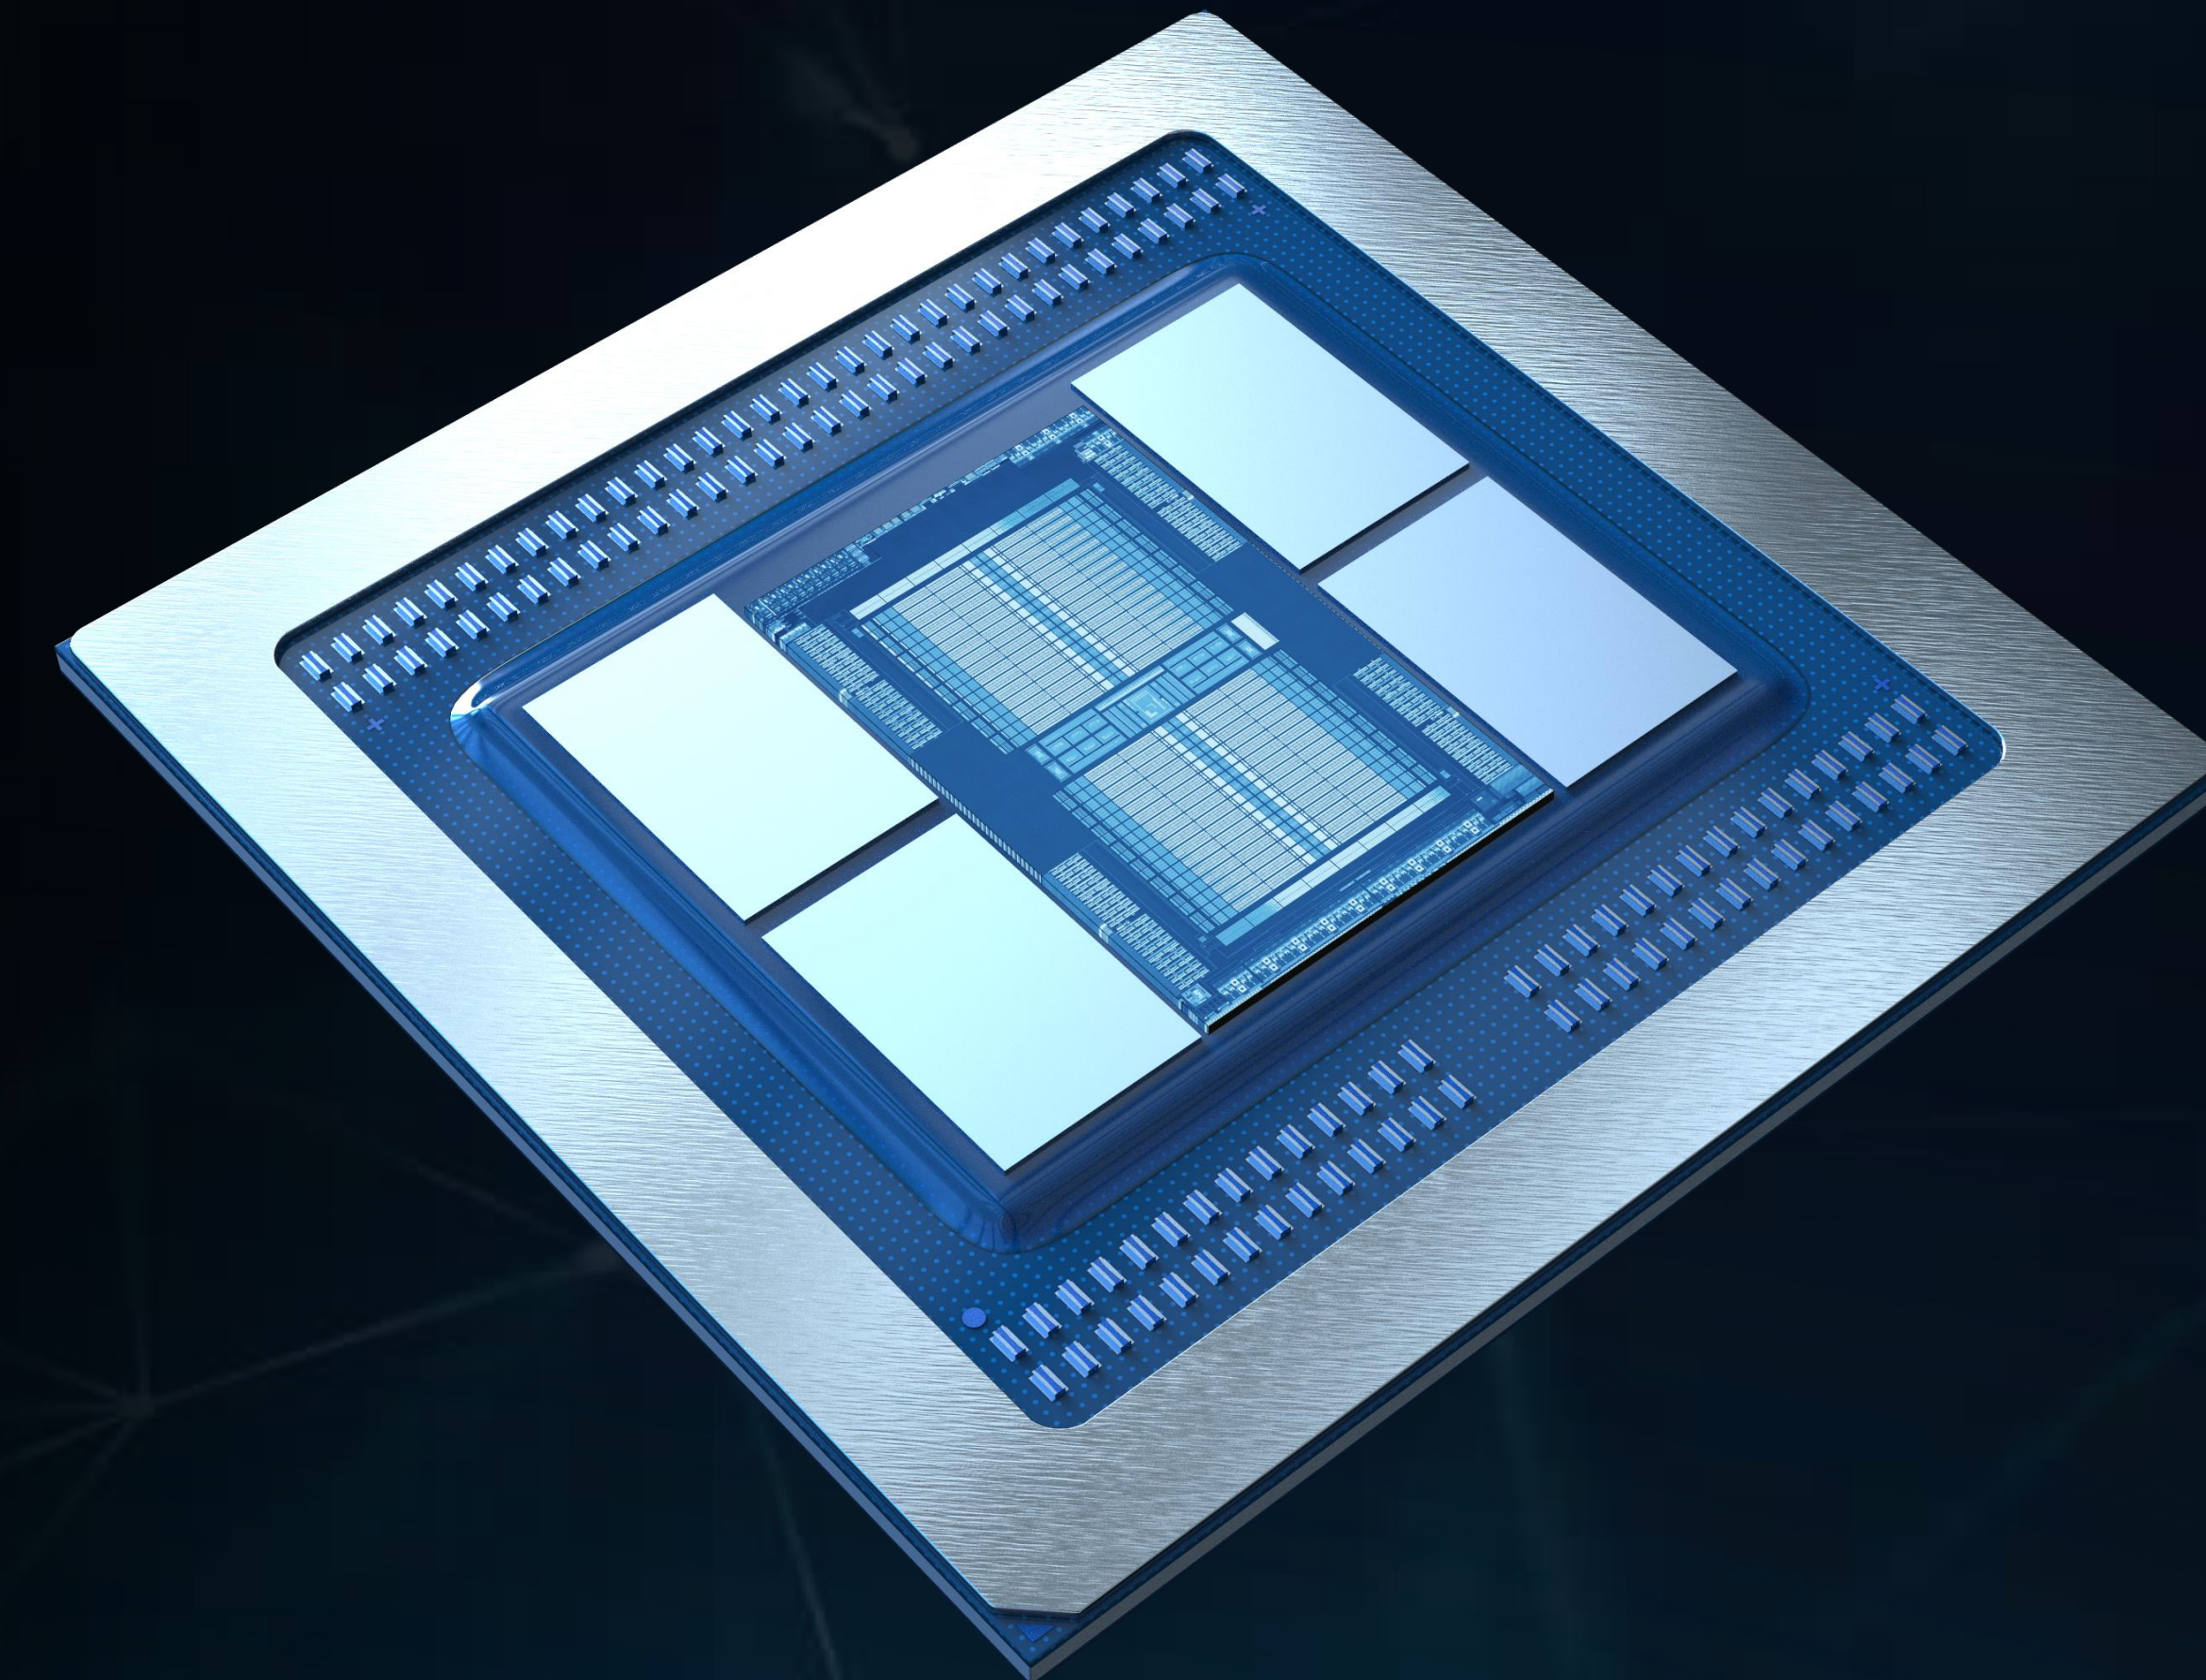


# HIGHLY EFFICIENT MEMORY SUBSYSTEM WITH 1 TB/S MEMORY BANDWIDTH

32 GB of High Bandwidth  
Memory (HBM2)

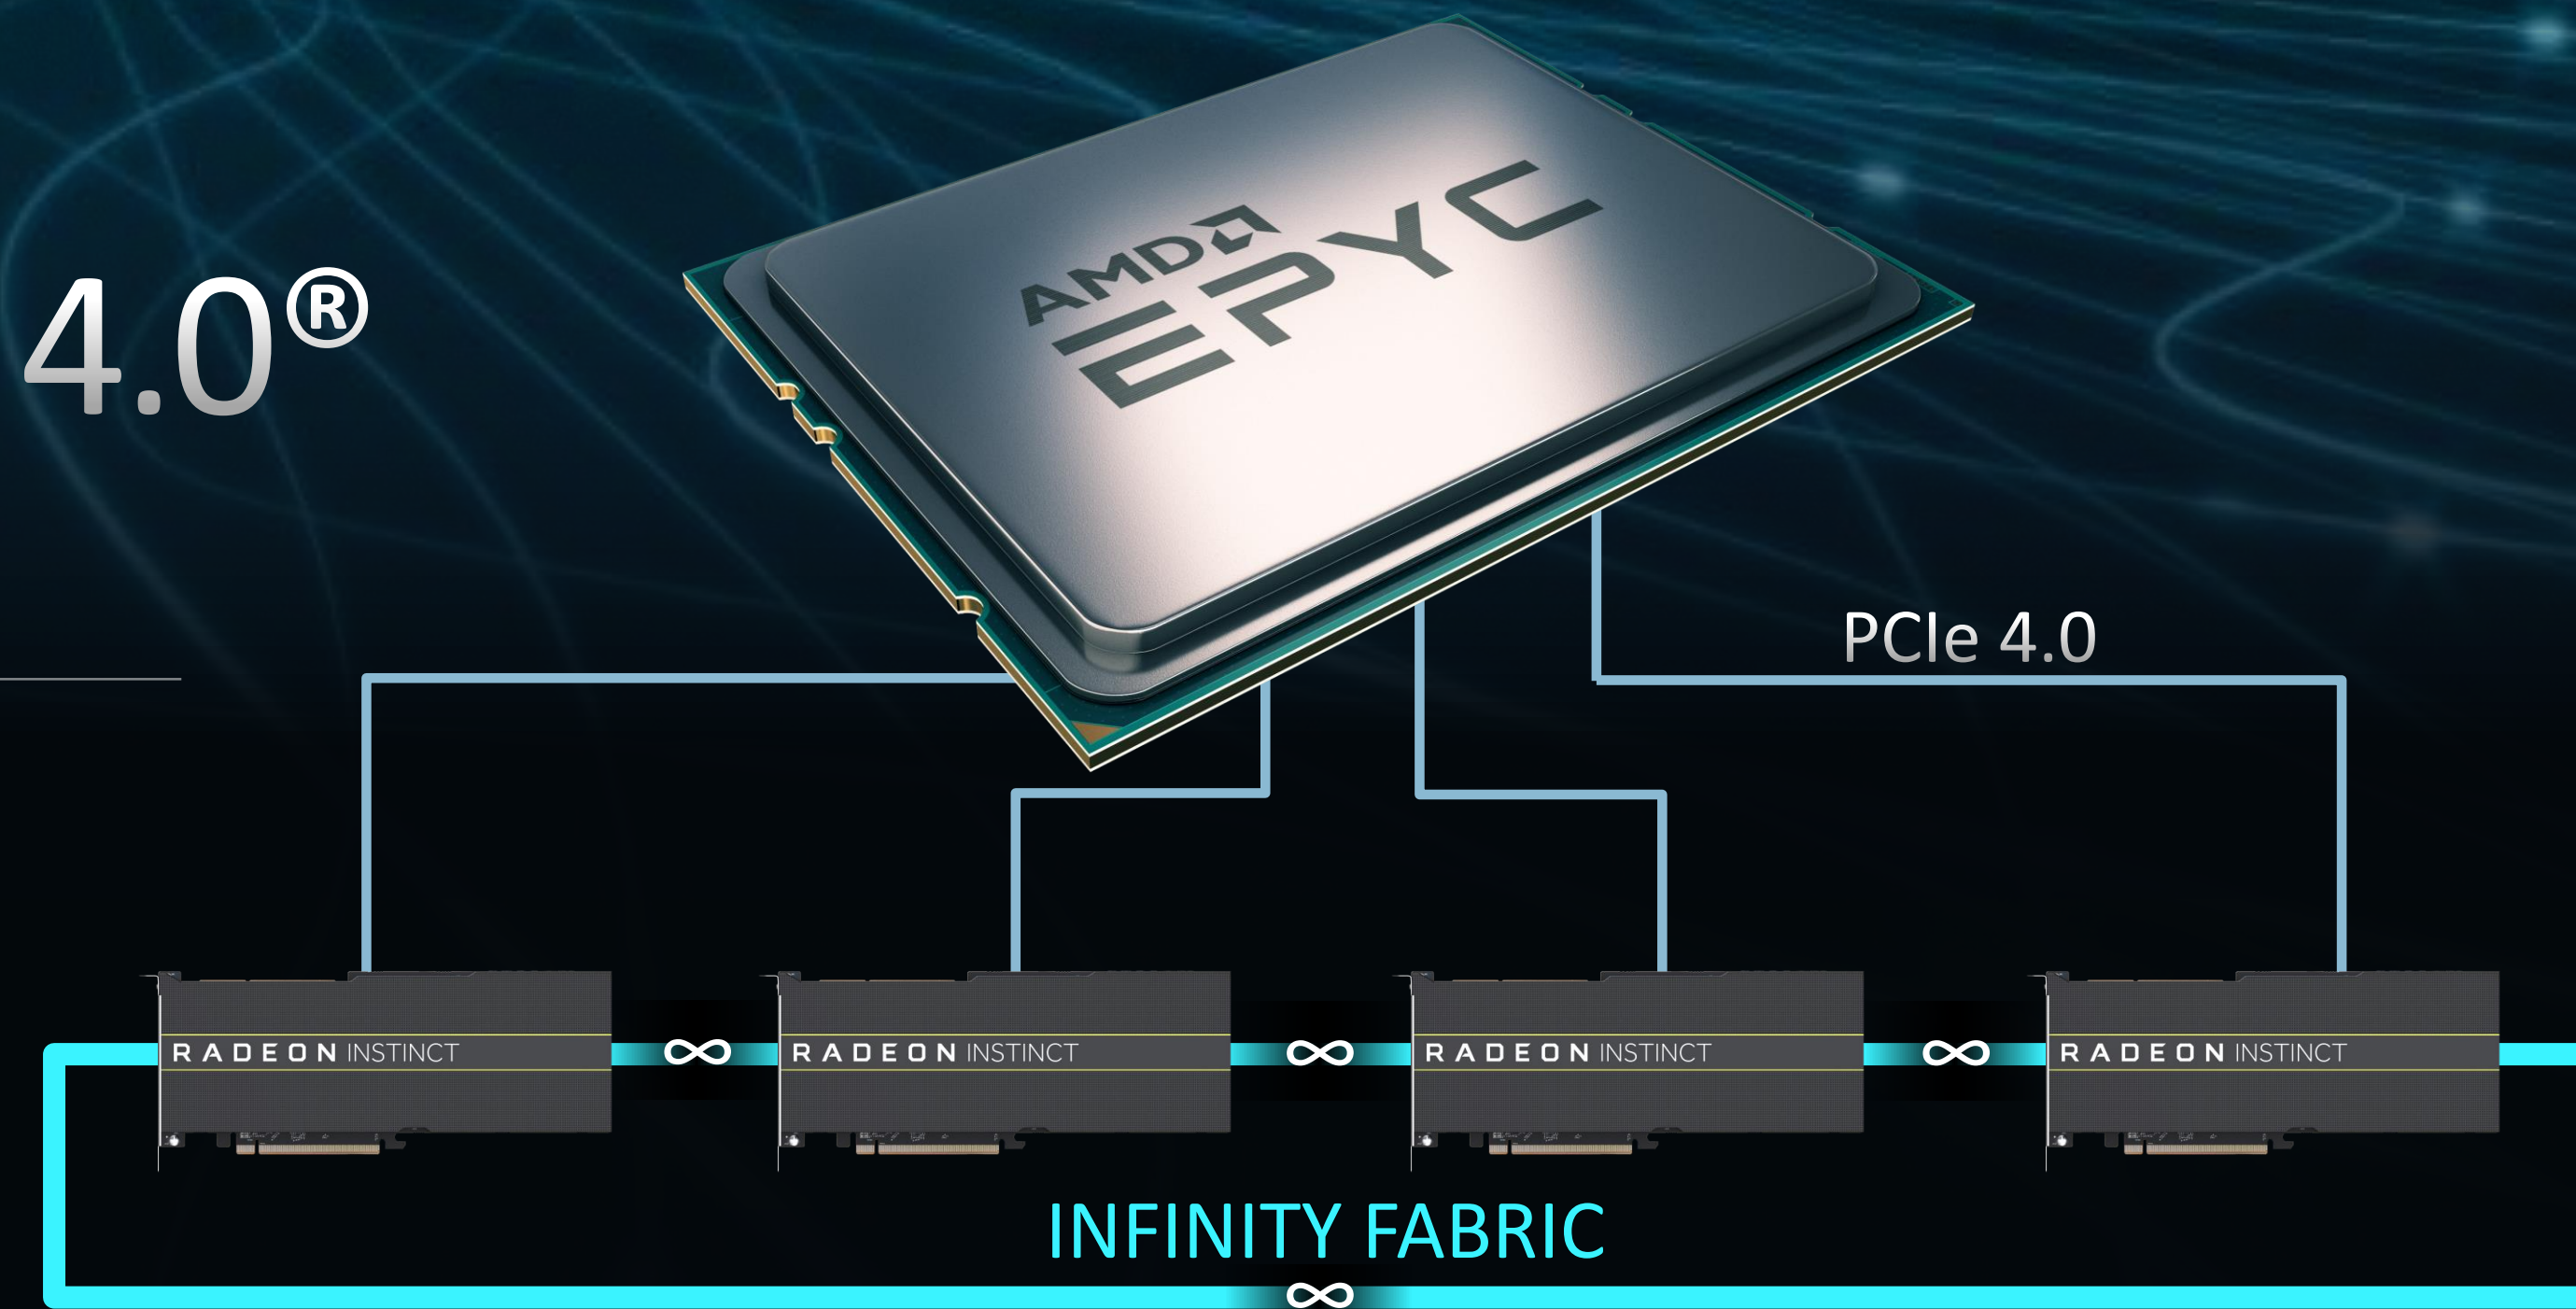
Built-in Error Detection  
and Correction with ECC

Support Huge Datasets with Enterprise Class Reliability



# INDUSTRY-FIRST PCIe 4.0<sup>®</sup> CAPABLE GPU

64 GB/s Bi-Directional  
CPU-to-GPU Bandwidth



Leading Edge Connectivity

# INFINITY FABRIC LINKS

100 GB/s per Link for  
GPU-to-GPU Bandwidth

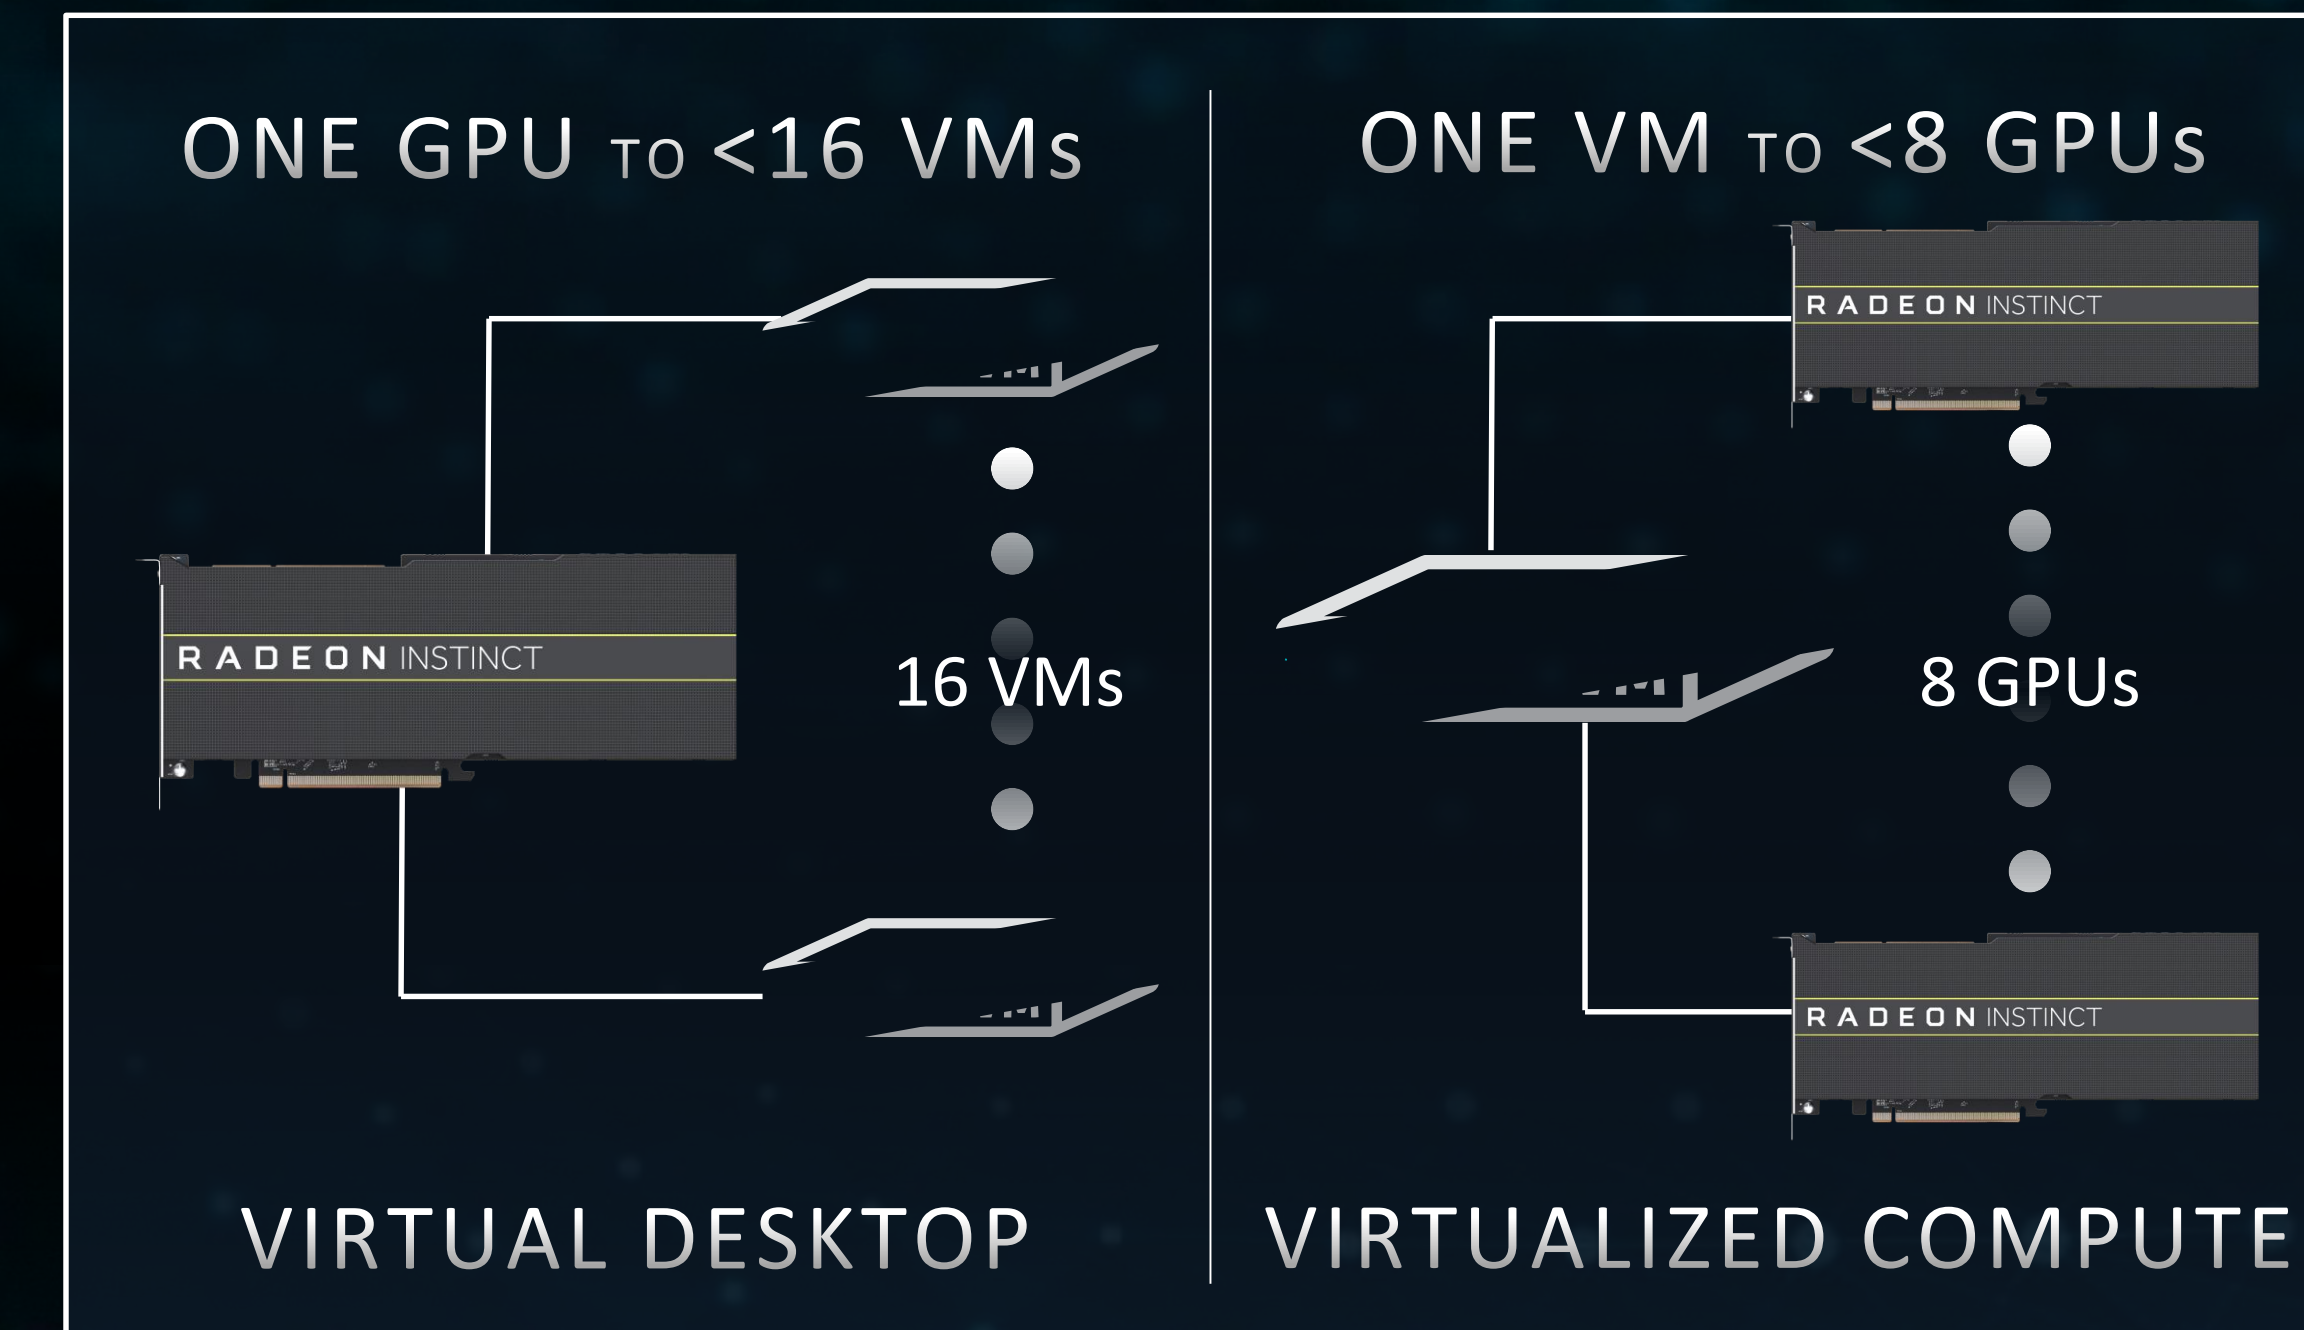
# THIRD-GENERATION HARDWARE-VIRTUALIZED GPU

Deploy without  
Software Overhead

Data Isolation Helps  
with Security

Drive Many  
Virtual Machines  
per GPU

Scale Up Performance  
with Multiple GPUs





# OPEN SOURCE FOUNDATION FOR MACHINE LEARNING

## Announcing ROCm 2.0

Latest Machine  
Learning  
Frameworks

Optimized  
Math  
Libraries

Dockers and  
Kubernetes  
Support

Up-Streamed for  
Linux Kernel  
Distributions

Applications

Machine Learning Apps

Frameworks

TensorFlow

Caffe 2

PyTorch

MXnet

Middleware  
and Libraries

MIOpen

BLAS,FFT,RNG

RCCL

Eigen

ROCm

OpenMP

HIP

OpenCL™

Python

Fully Open Source ROCm Platform

Devices

GPU

CPU

APU

DLA



EMBRACED BY THE

# OPEN SOURCE COMMUNITY

Extends ROCm Access to  
**MILLIONS OF LINUX KERNEL USERS**

**Caffe2**

**PYTORCH**

**tvm**



TensorFlow



PaddlePaddle



**GitHub**



docker

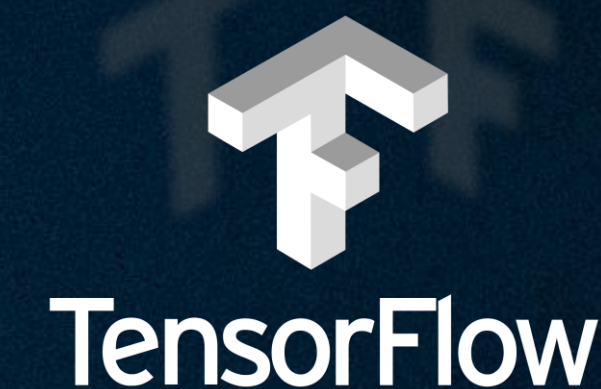
**Numba**



EMBRACED BY

INDUSTRY LEADERS

“Google believes that open source is good for everyone. We've seen how helpful it can be to open source machine learning technology, and we're glad to see AMD embracing it. With the ROCm open software platform, TensorFlow users will benefit from GPU acceleration and a more robust open source machine learning ecosystem.”

**Rajat Monga**

Engineering Director

TENSORFLOW, GOOGLE



EMBRACED BY

INDUSTRY LEADERS



Zhenyu Hou

Executive Director, System Technologies

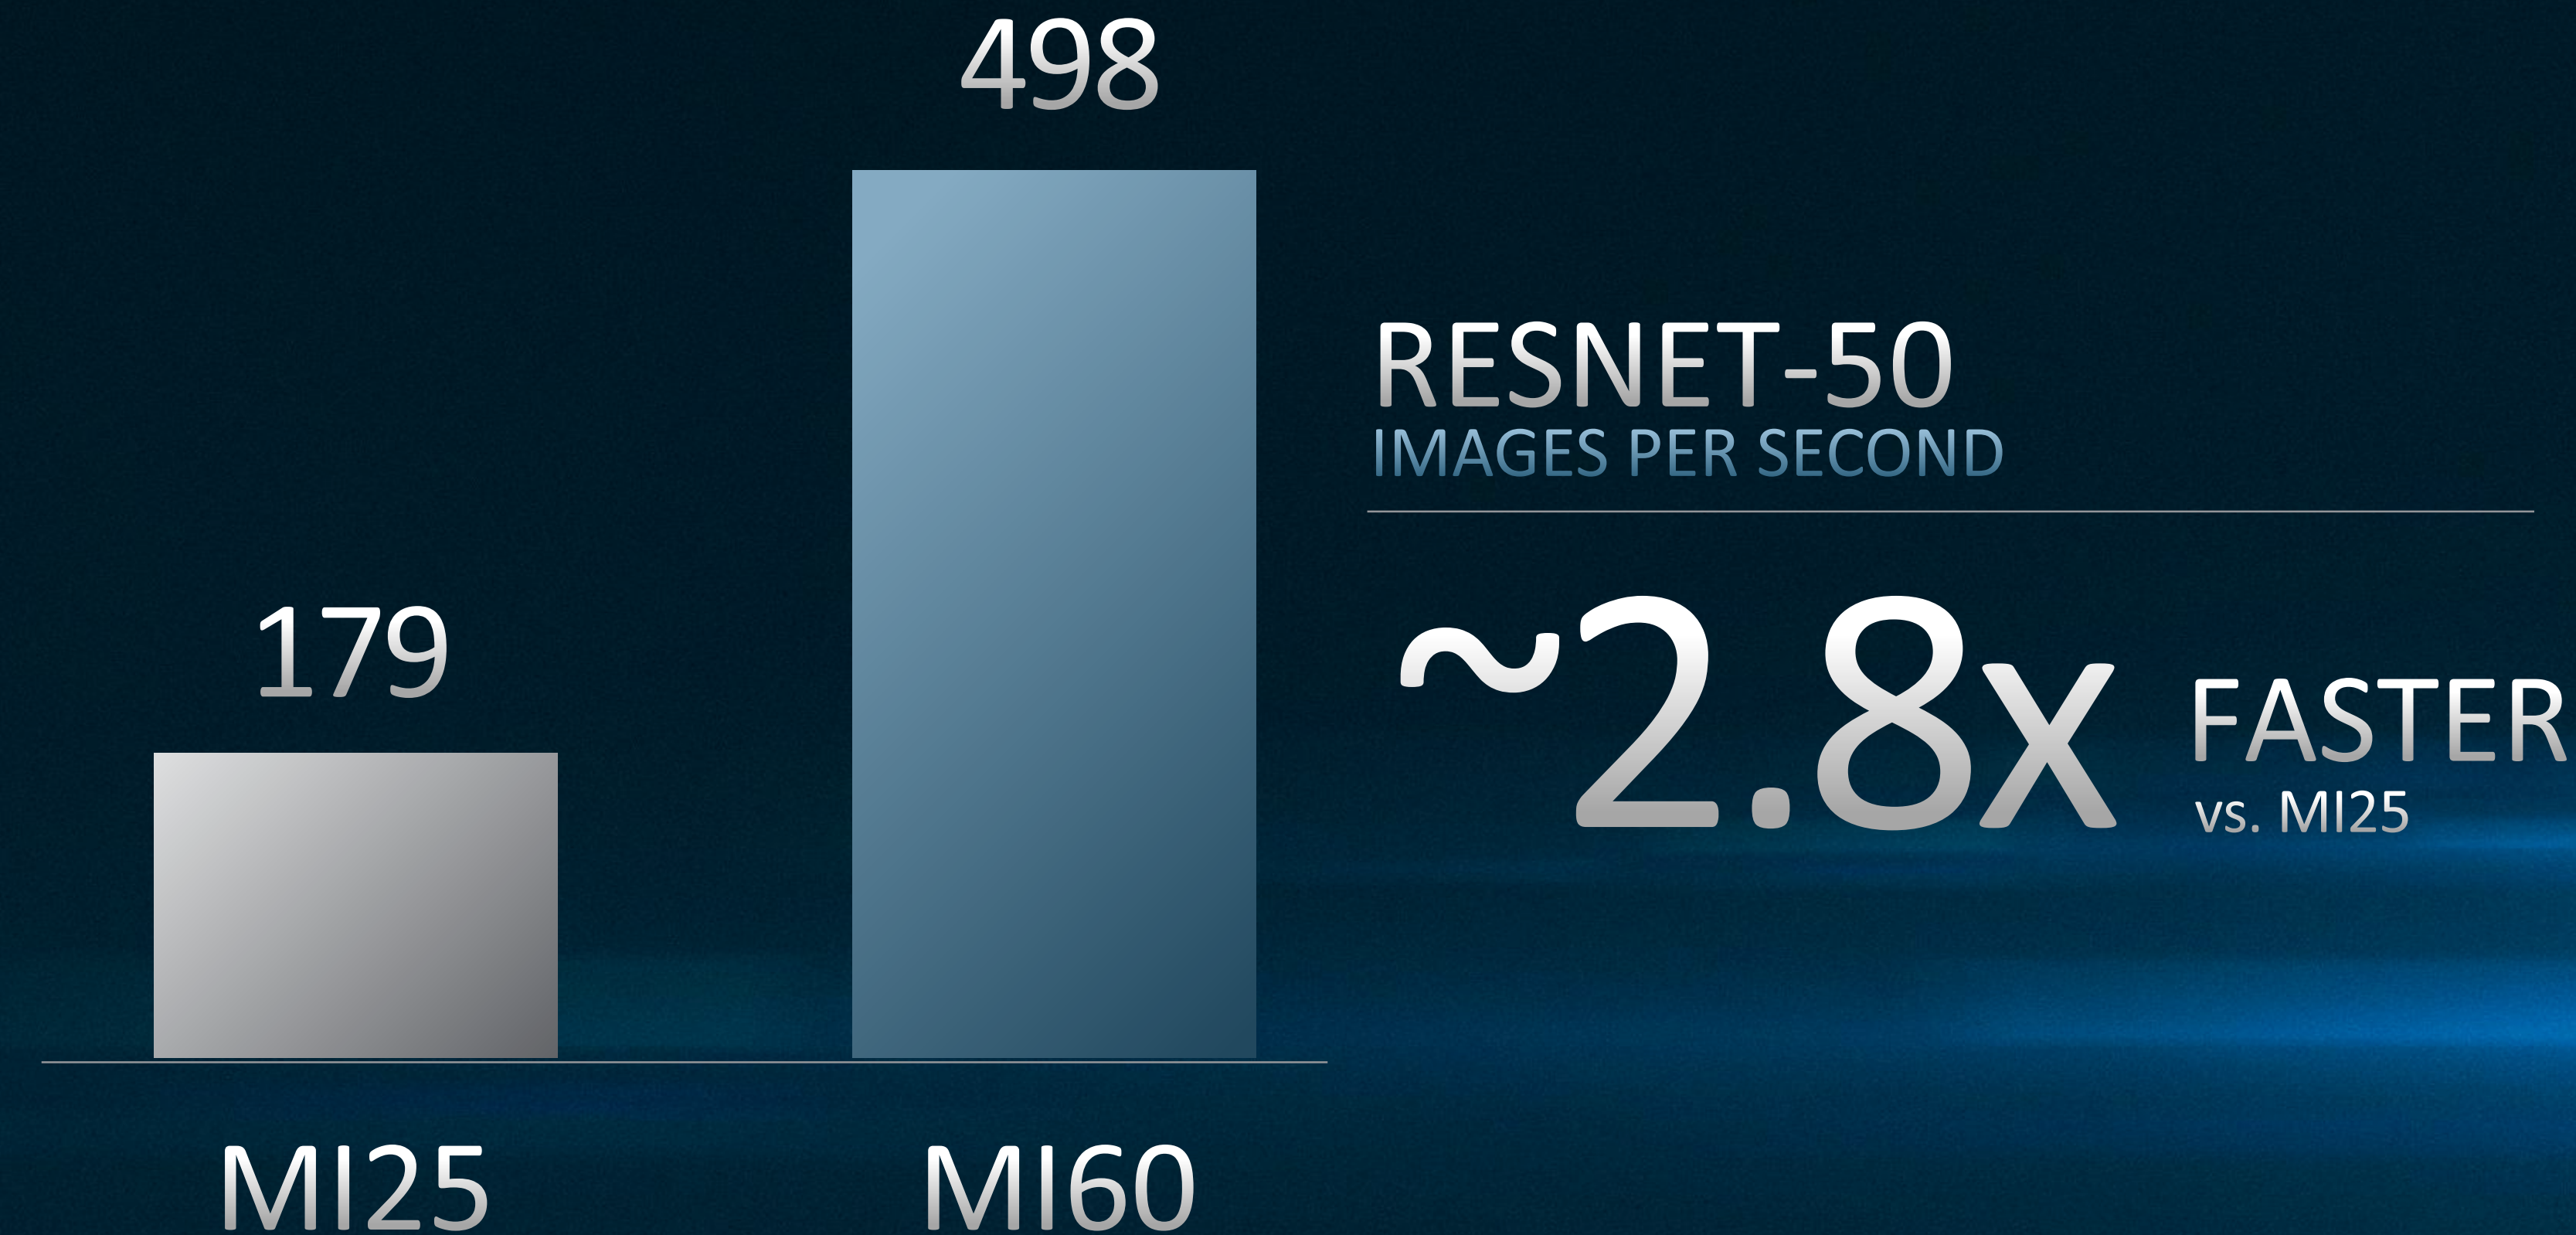
# AMD RADEON INSTINCT™ MI60

Generational Improvement



# AMD RADEON INSTINCT™ MI60

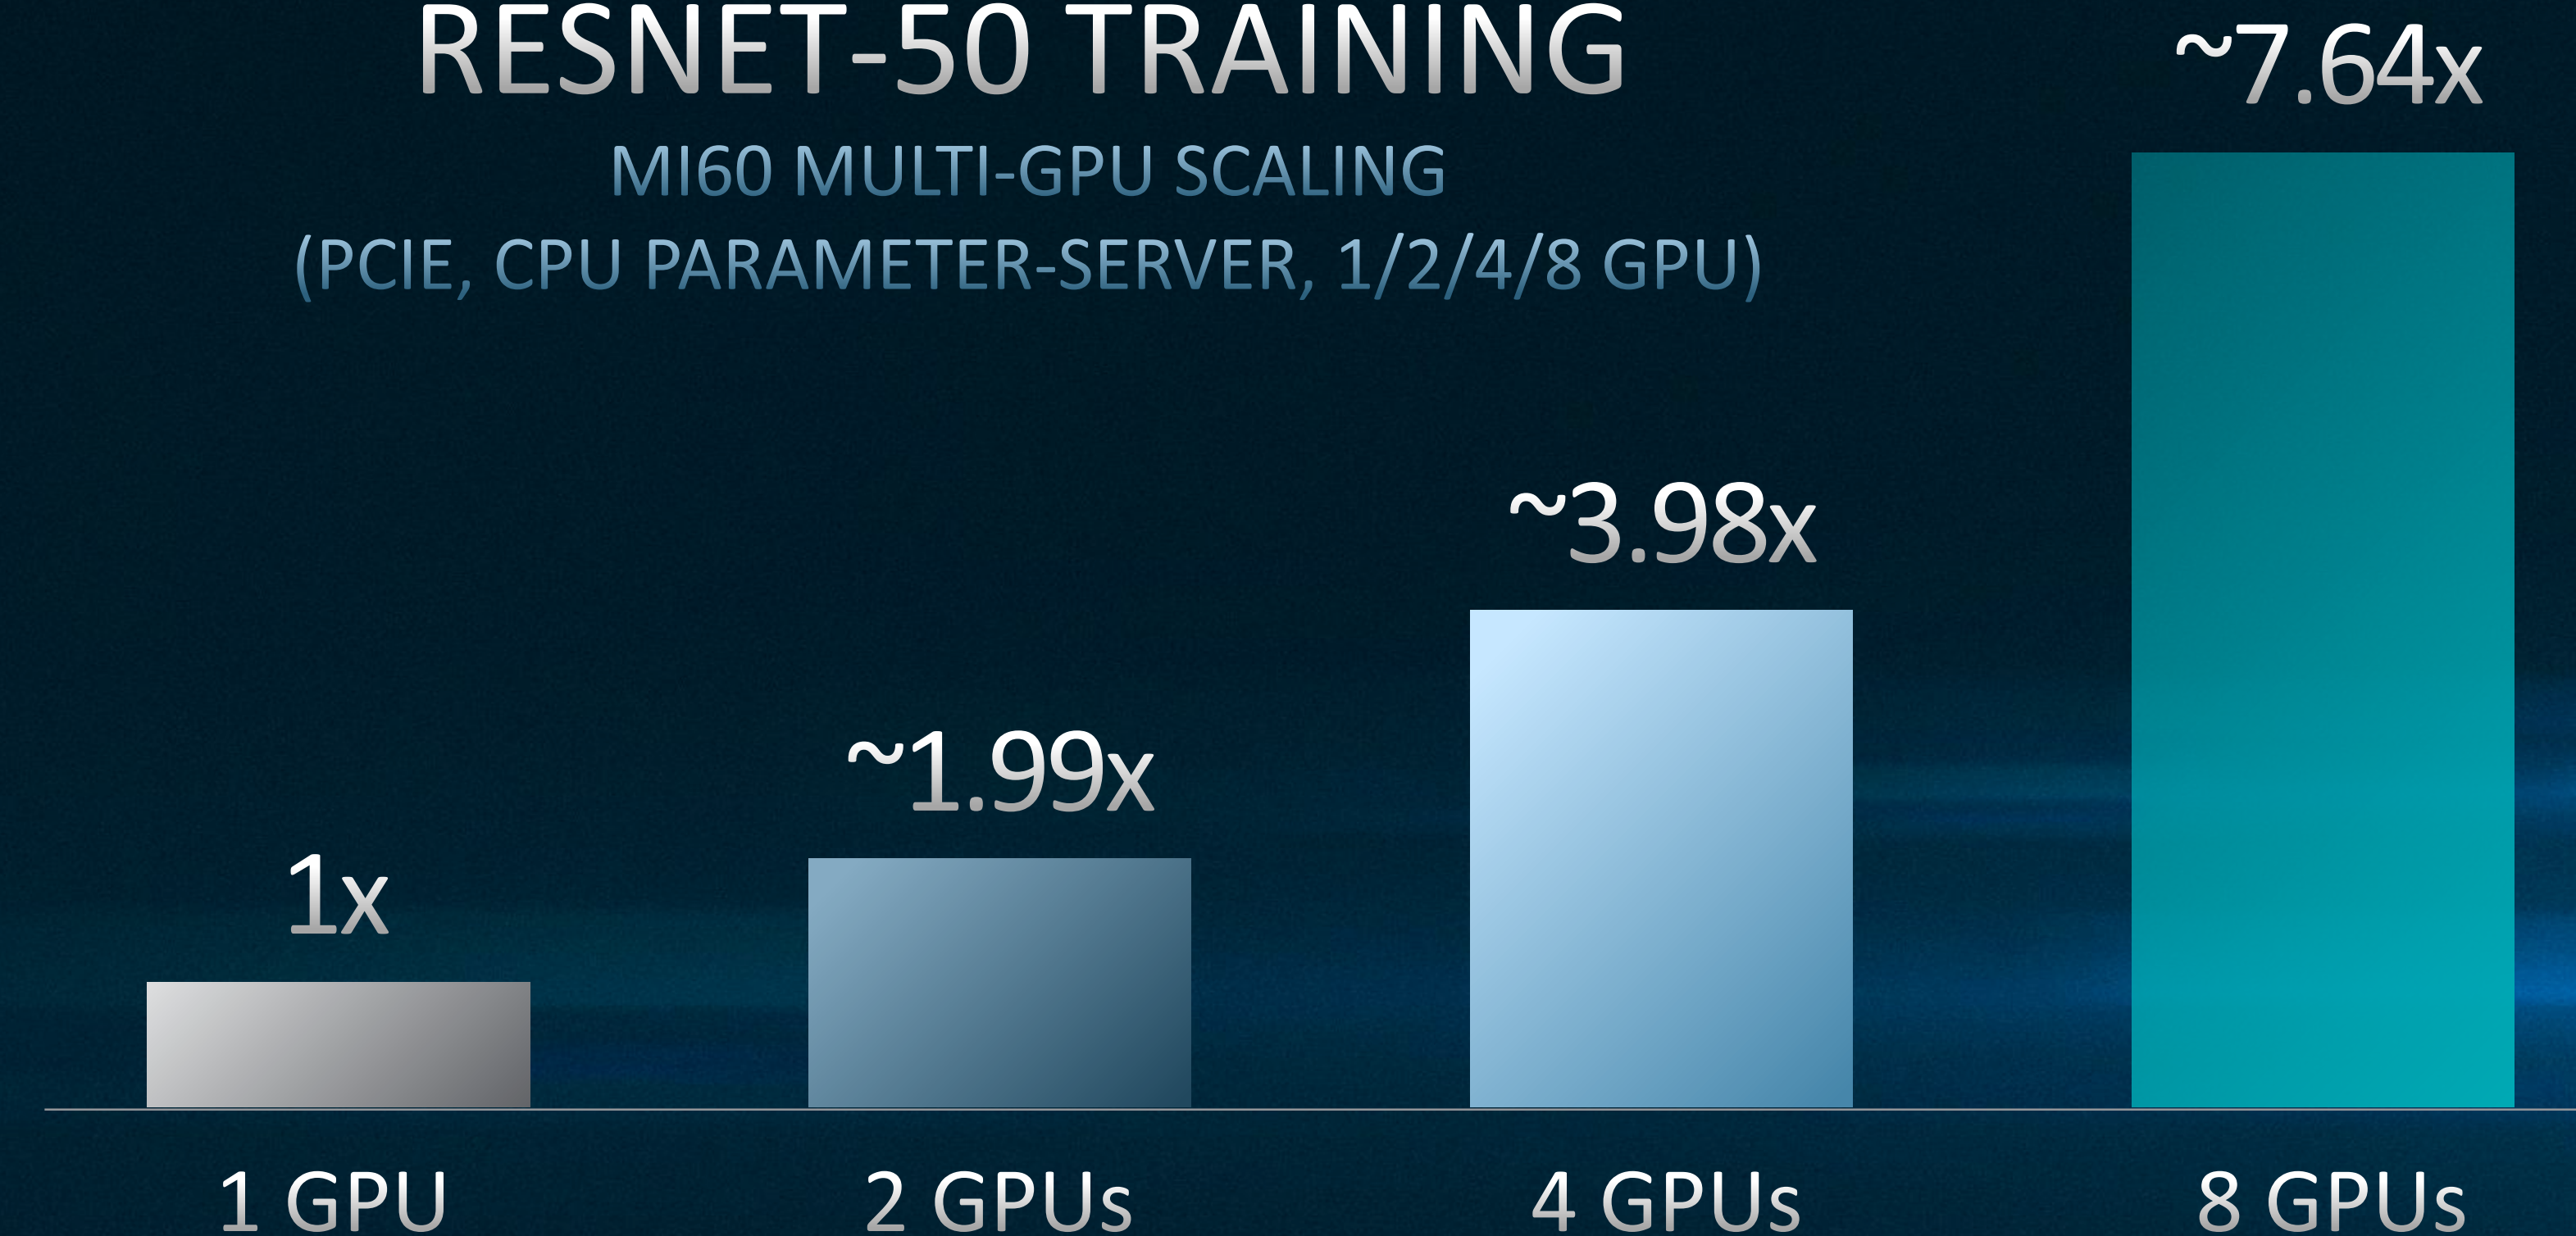
Generational Improvement



# AMD RADEON INSTINCT™ MI60

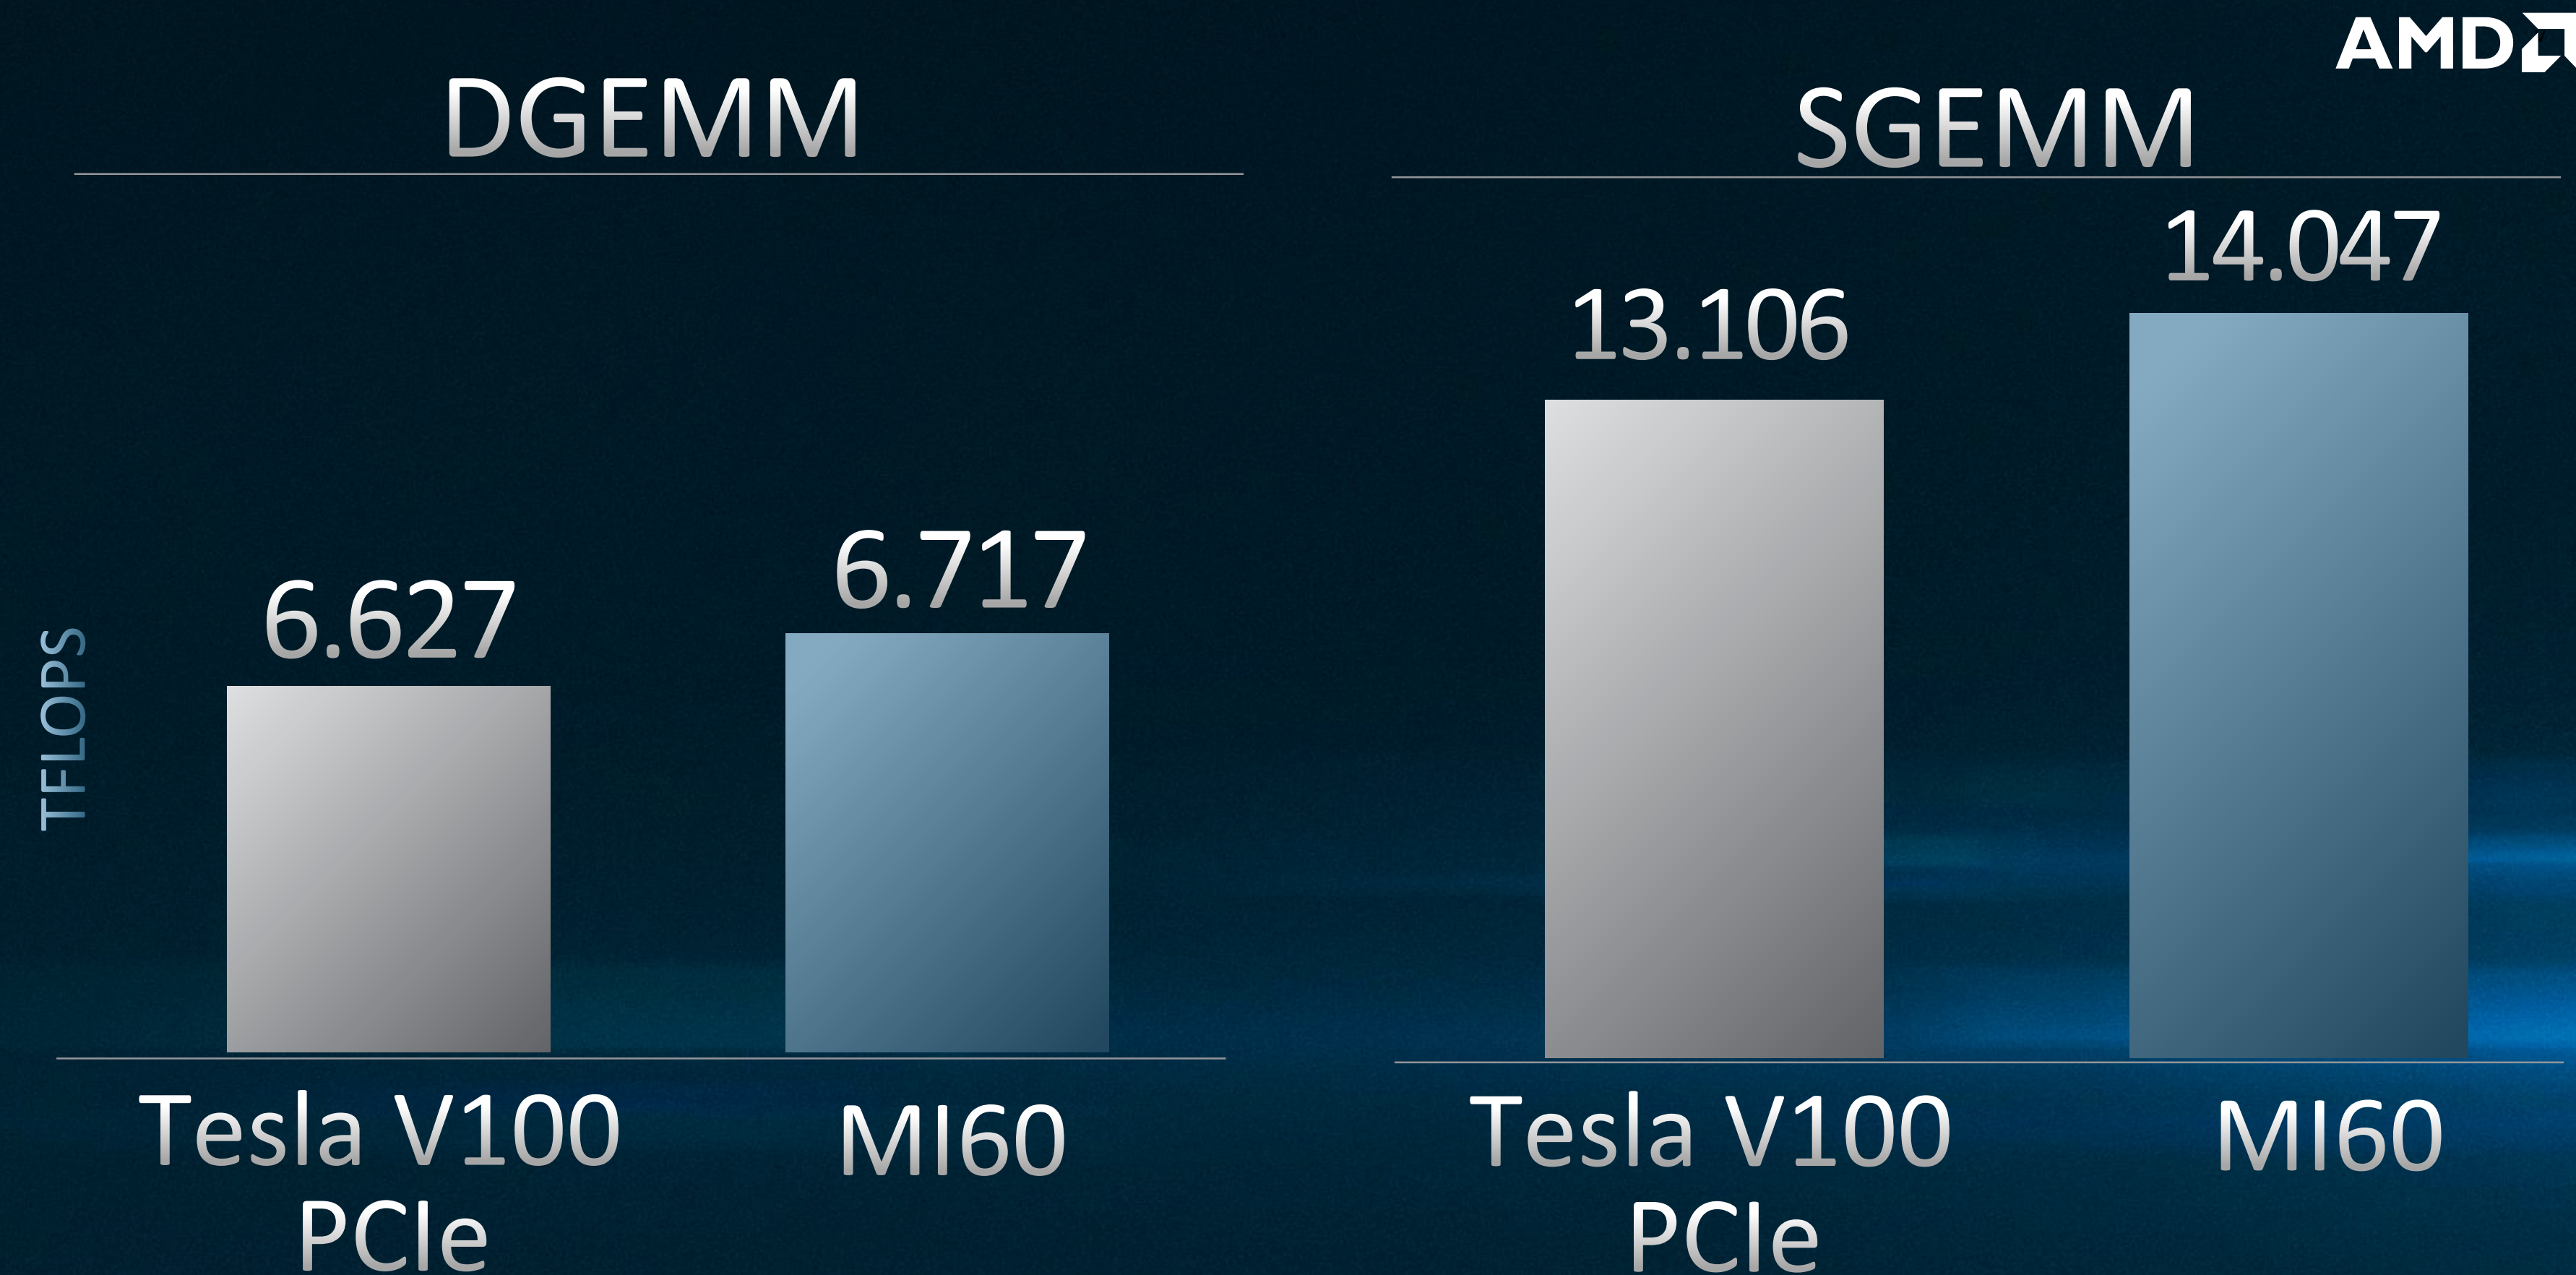
Excellent Performance Scalability

RESNET-50 TRAINING  
MI60 MULTI-GPU SCALING  
(PCIe, CPU PARAMETER-SERVER, 1/2/4/8 GPU)



# AMD RADEON INSTINCT™ MI60

Competitive Performance



# AMD RADEON INSTINCT™ MI60

Competitive Performance



RESNET-50 Training

IMAGES PER SECOND

Comparable Performance

vs. Tesla V100 PCIe

# SIMULATED WORLDS FOR MACHINE LEARNING

Real-time Data  
Capture is Inefficient

Simulation Can Produce  
Massive Datasets

Safe for Corner  
Cases and Variations

The logo for highwai, featuring the word "highwai" in a light blue, lowercase, sans-serif font. A stylized white swoosh underline starts under the 'h' and curves under the 'ai'.

**Peter McGuinness**  
CEO

STEP 1  
SIMULATION

Generate Data for Neural  
Network Training



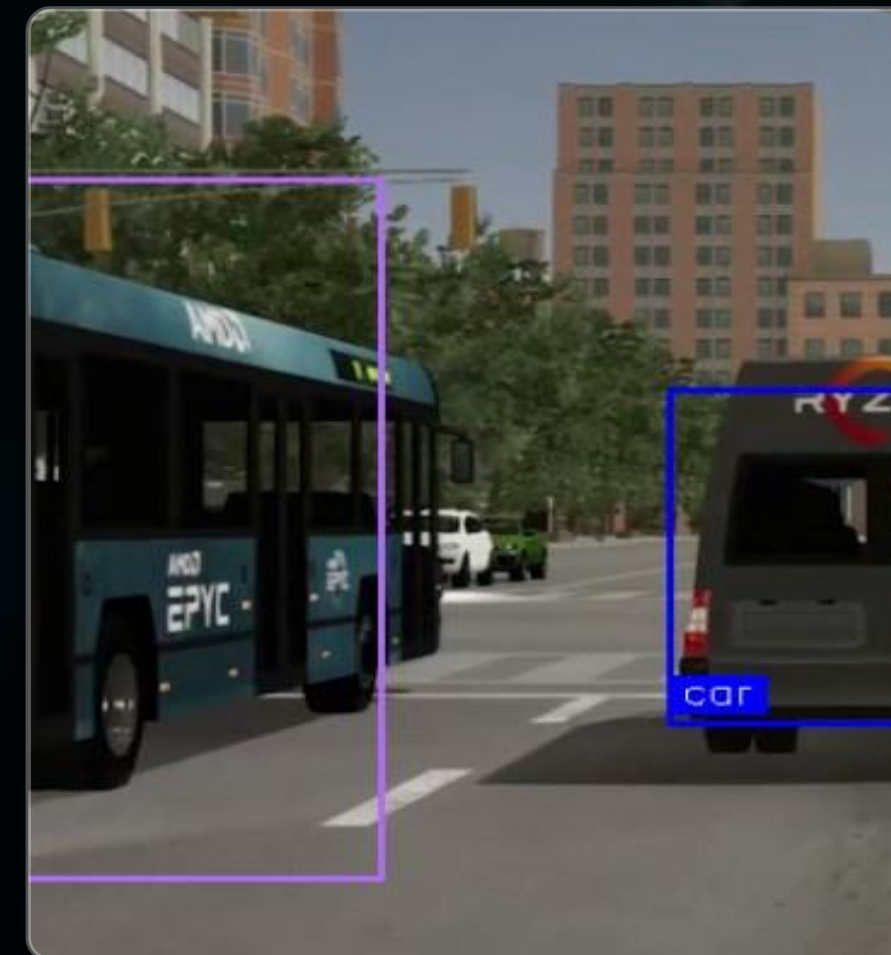
STEP 2  
TRAINING

Train Neural Networks  
with Simulated Data



STEP 3  
VALIDATION

Verify Training with  
Simulated and Real Data



STEP 4  
DEPLOYMENT

Apply to  
Real-world Problems



AMD RADEON INSTINCT™ MI60

# MACHINE LEARNING DEMONSTRATION

# AMD RADEON INSTINCT™ MI60

## SHIPPING THIS QUARTER

WORLD'S FIRST  
7nm GPU

FIRST PCIe® 4.0 CAPABLE GPU +  
INFINITY FABRIC LINKS

EFFICIENT 1 TB/S  
MEMORY BANDWIDTH

WORLD'S ONLY  
HARDWARE-VIRTUALIZED GPU

7.4 TFLOPS FP64  
FOR HPC

14.7 TFLOPS FP32  
FOR TRAINING

118 TOPS INT4  
FOR INFERENCE

# CHOICE AND INNOVATION FOR THE FUTURE DATACENTER GPU ROADMAP

## “MI-NEXT”

### AMD Radeon Instinct™ MI25

14nm GPU  
 PCIe® 3.0  
 16 GB of HBM2 Memory

### AMD Radeon Instinct™ MI60

First 7nm GPU  
 First PCIe® 4.0 Capable GPU  
 32 GB of HBM2 Memory  
 Infinity Fabric Links

Higher Performance  
 Increased Connectivity  
 Software Compatibility

## Predictable Product Cadence with Generational Performance Gains

The background of the slide is a deep space scene. At the bottom, the curved horizon of the Earth is visible, with a bright sun or star rising directly behind it, creating a lens flare and illuminating the scene. The sky is a dark, starry void. The AMD logo is rendered in a metallic, blue-tinted 3D style, while the words 'NEXT HORIZON' are in a clean, white, sans-serif font.

# AMD

## NEXT HORIZON

# END NOTES

**RIV-6:** As of Oct 22, 2018. The results calculated for Radeon Instinct MI60 designed with Vega 7nm FinFET process technology resulted in 118 TOPS INT4, 59 TOPS INT8, 29.5 TFLOPS half precision (FP16), 14.8 TFLOPS single precision (FP32) and 7.4 TFLOPS double precision (FP64) peak theoretical floating-point performance. The results calculated for Radeon Instinct MI50 designed with Vega 7nm FinFET process technology resulted in 107 TOPS INT4, 53.6 TOPS INT8, 26.8 TFLOPS peak half precision (FP16), 13.4 TFLOPS peak single precision (FP32) and 6.7 TFLOPS peak double precision (FP64) floating-point performance. The results calculated for Radeon Instinct MI25 GPU based on the “Vega10” architecture resulted in 24.6 TFLOPS peak half precision (FP16), 12.3 TFLOPS peak single precision (FP32) and 768 GFLOPS peak double precision (FP64) floating-point performance. AMD TFLOPS calculations conducted with the following equation for Radeon Instinct MI25, MI50, and MI60 GPUs: FLOPS calculations are performed by taking the engine clock from the highest DPM state and multiplying it by xx CUs per GPU. Then, multiplying that number by xx stream processors, which exist in each CU. Then, that number is multiplied by 2 FLOPS per clock for FP32, 4 FLOPS per clock for FP16, 8 FLOPS per clock for INT8, and 16 FLOPS per clock for INT8 to determine TFLOPS or TOPS. The FP64 TFLOPS rate for MI50 and MI60 is calculated using 1/2 rate. The FP64 TFLOPS rate for MI25 is calculated using 1/16th rate.

**RIV-8:** Testing Conducted by AMD performance labs on October 31, 2018, on a system comprising of Dual Intel Xeon Gold 6132, 256GB DDR4 system memory, Ubuntu 16.04.5 LTS, AMD Radeon Instinct MI25 graphics, AMD Radeon Instinct MI60 graphics, ROCm 19.224 driver, TensorFlow 1.11. Benchmark application: Resnet50 FP16 batch size 256. AMD Radeon Instinct MI25 = 179 images/s. AMD Radeon Instinct MI60 = max 498.99 images/s. Performance differential:  $498.99/179 = \text{up to } 2.8\text{x}$  more performance than Radeon Instinct MI25. Server manufacturers may vary configurations, yielding different results. Performance may vary based on use of latest drivers and optimizations.

**RIV-11:** Testing Conducted by AMD performance labs on October 31, 2018, on a system comprising of Dual Intel Xeon Gold 6132, 256GB DDR4 system memory, Ubuntu 16.04.5 LTS, AMD Radeon Instinct MI60 graphics running at 1600e/500m, ROCm 19.224 driver, TensorFlow 1.11. Benchmark application: Resnet50 FP32 batch size 256. 1x AMD Radeon Instinct MI60 = 278.63 images/s, 2x Radeon Instinct MI60 = 553.98 images/s. Performance differential:  $553.98/278.63 = 1.99\text{x}$  times more performance than 1x Radeon Instinct MI60. 4x Radeon Instinct MI60 = 1109.24 images/s. Performance differential:  $1109.24/278.63 = 3.98\text{x}$  times more performance than 1x Radeon Instinct MI60. 8x Radeon Instinct MI60 = 2128.33 images/s. Performance differential:  $2128.33/278.63 = 7.64\text{x}$  times more performance than 1x Radeon Instinct MI60. Server manufacturers may vary configurations, yielding different results. Performance may vary based on use of latest drivers and optimizations.

# END NOTES

**RIV-12:** Testing Conducted by AMD performance labs on October 31, 2018, on a system comprising of Dual Socket Intel Xeon Gold 6130, 256GB DDR4 system memory, Ubuntu 16.04.5 LTS, AMD Radeon Instinct MI25 graphics, AMD Radeon Instinct MI60 graphics, ROCm 19.211 driver. Benchmark application: rocBLAS DEGEMM N=M=K=5760. AMD Radeon Instinct MI25 GFLOPS = 763. AMD Radeon Instinct MI60 TFLOPS = 6.717. Performance differential:  $6717/763 = 8.8x$  more performance than Radeon Instinct MI25. Server manufacturers may vary configurations, yielding different results. Performance may vary based on use of latest drivers and optimizations. Calculated on Oct 22, 2018, the Radeon Instinct MI60 GPU resulted in 7.4 TFLOPS peak theoretical double precision floating-point (FP64) performance. AMD TFLOPS calculations conducted with the following equation: FLOPS calculations are performed by taking the engine clock from the highest DPM state and multiplying it by xx CUs per GPU. Then, multiplying that number by xx stream processors, which exist in each CU. Then, that number is multiplied by 1/2 FLOPS per clock for FP64. TFLOP calculations for MI60 can be found at <https://www.amd.com/en/products/professional-graphics/instinct-mi60>.

**RIV-13:** Testing Conducted by AMD performance labs on October 31, 2018, on a system comprising of Dual Socket Intel Xeon Gold 6130, 256GB DDR4 system memory, Ubuntu 16.04.5 LTS, NVIDIA Tesla V100 PCIe with CUDA 10.0.130 and CUDNN 7.3, AMD Radeon Instinct MI60 graphics, ROCm 19.224 driver, TensorFlow 1.11. Benchmark application: Resnet50 FP32 batch size 256. NVIDIA Tesla V100 PCIe = 357 images/s. AMD Radeon Instinct MI60 = 334 images/s. Server manufacturers may vary configurations, yielding different results. Performance may vary based on use of latest drivers and optimizations.

**RIV-14:** Testing Conducted by AMD performance labs on October 31, 2018, on a system comprising of Dual Socket Intel Xeon Gold 6130, 256GB DDR4 system memory, Ubuntu 16.04.5 LTS, NVIDIA Tesla V100 PCIe with CUDA 10.0.130 and CUDNN 7.3, AMD Radeon Instinct MI60 graphics with ROCm 19.224 driver. Benchmark application: cuBLAS/rocBLAS SGEMM N=M=K=5760. AMD NVIDIA Tesla V100 PCIe TFLOPS = 13.106. AMD Radeon Instinct MI60 TFLOPS = 14.047. Server manufacturers may vary configurations, yielding different results. Performance may vary based on use of latest drivers and optimizations.

# DISCLAIMER & ATTRIBUTION

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.

The information contained herein is for informational purposes only, and is subject to change without notice. Timelines, roadmaps, and/or product release dates shown in these slides are plans only and subject to change. [Insert codename, i.e. “Vega”, “Radeon Vega”, “Fiji”, and “Polaris”] are codenames for AMD architectures, and are not product names. GD-122.

## ATTRIBUTION

© 2018 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, Radeon, Radeon Instinct and combinations thereof are trademarks of Advanced Micro Devices, Inc. in the United States and/or other jurisdictions. Microsoft is a registered trademark of Microsoft Corporation in the US and other jurisdictions. Other names are for informational purposes only and may be trademarks of their respective owners.

# CAUTIONARY STATEMENT

This presentation contains forward-looking statements concerning Advanced Micro Devices, Inc. (AMD) including, but not limited to the expected datacenter total addressable market; and AMD's expected opportunities in the datacenter market, 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 29, 2018.

The logo features the word "AMD" in a large, metallic blue, 3D font. To the right of "AMD" is a square icon composed of four blue, 3D rectangular blocks arranged in a larger square. Below "AMD" is the word "NEXT HORIZON" in a white, 3D font with a blue-to-white gradient. The entire logo is set against a background of a dark space with a bright sun rising over the horizon of a planet, creating a lens flare effect.

# AMD

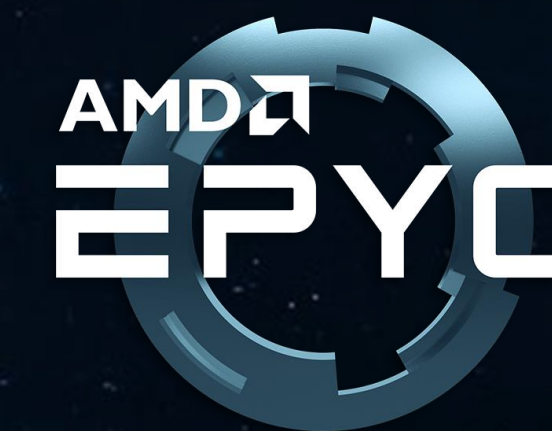
## NEXT HORIZON

**FORREST NORROD**

Senior Vice President and General Manager  
Datacenter and Embedded Solutions Group

THE DATACENTER REPRESENTS AN  
INCREDIBLE OPPORTUNITY

\$17B  
TAM



\*AMD ESTIMATE FOR 2021

# DRAMATIC CHANGE IN THE DATACENTER

## THE MODERN DATACENTER REQUIRED NEW THINKING



### A DECADE AGO

- Top 10 Customers: <10%
- Scale Up
- Physical
- On-Premise
- Proprietary Software

### TODAY

- Top 7 Customers: ~40%
- Scale Out
- Virtual → Container
- Hybrid – Cloud
- Open Source

# AMD EPYC™ LEADERSHIP FOR THE MODERN DATACENTER



Core Density

Memory Bandwidth

Leadership I/O

Advanced Security Features



## Cloud Service Providers



IaaS/PaaS



Media



Social



SaaS



## Enterprise IT



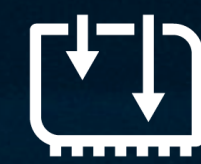
Virtualization



SDS/HCI



Hadoop



NoSQL



## High Performance Compute



Design & Simulation



Research & Academia



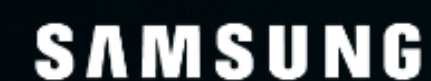
Machine Learning



Supercomputing

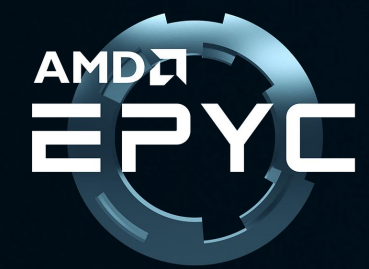
# AMD EPYC™ LEADERSHIP

# AMD EPYC™ YEAR ONE MOMENTUM





# AMD EPYC™ DESIGNED FOR CLOUD SERVICES



Accelerating and Securing the Journey to the Cloud



ROBUST AMD EPYC™-BASED  
CLOUD SERVICE OFFERINGS



# DIFFERENTIATED CLOUD SERVICES



LV Series for Demanding Storage and  
I/O Intensive Applications

Hb Series for High  
Performance Computing

Highest Throughput of Any  
Cloud Provider



Microsoft Azure



# GENERAL PURPOSE CLOUD SERVICES



Exceptional Virtual Machine  
Costs & Easy to Adopt

New “M” & “T” Instances for  
General Purpose Workloads

New “R” Instances for Memory  
Optimized Performance



# ENTERPRISE APPS AND BARE METAL



Up to 66% Less  
Cost/Core Hour

Virtual Machines & Bare Metal Shapes  
for Maximum Price/Performance

Full Support for Oracle  
Enterprise Applications

# ORACLE®

## CLOUD



# ENTERPRISE HOSTING PARTNERS

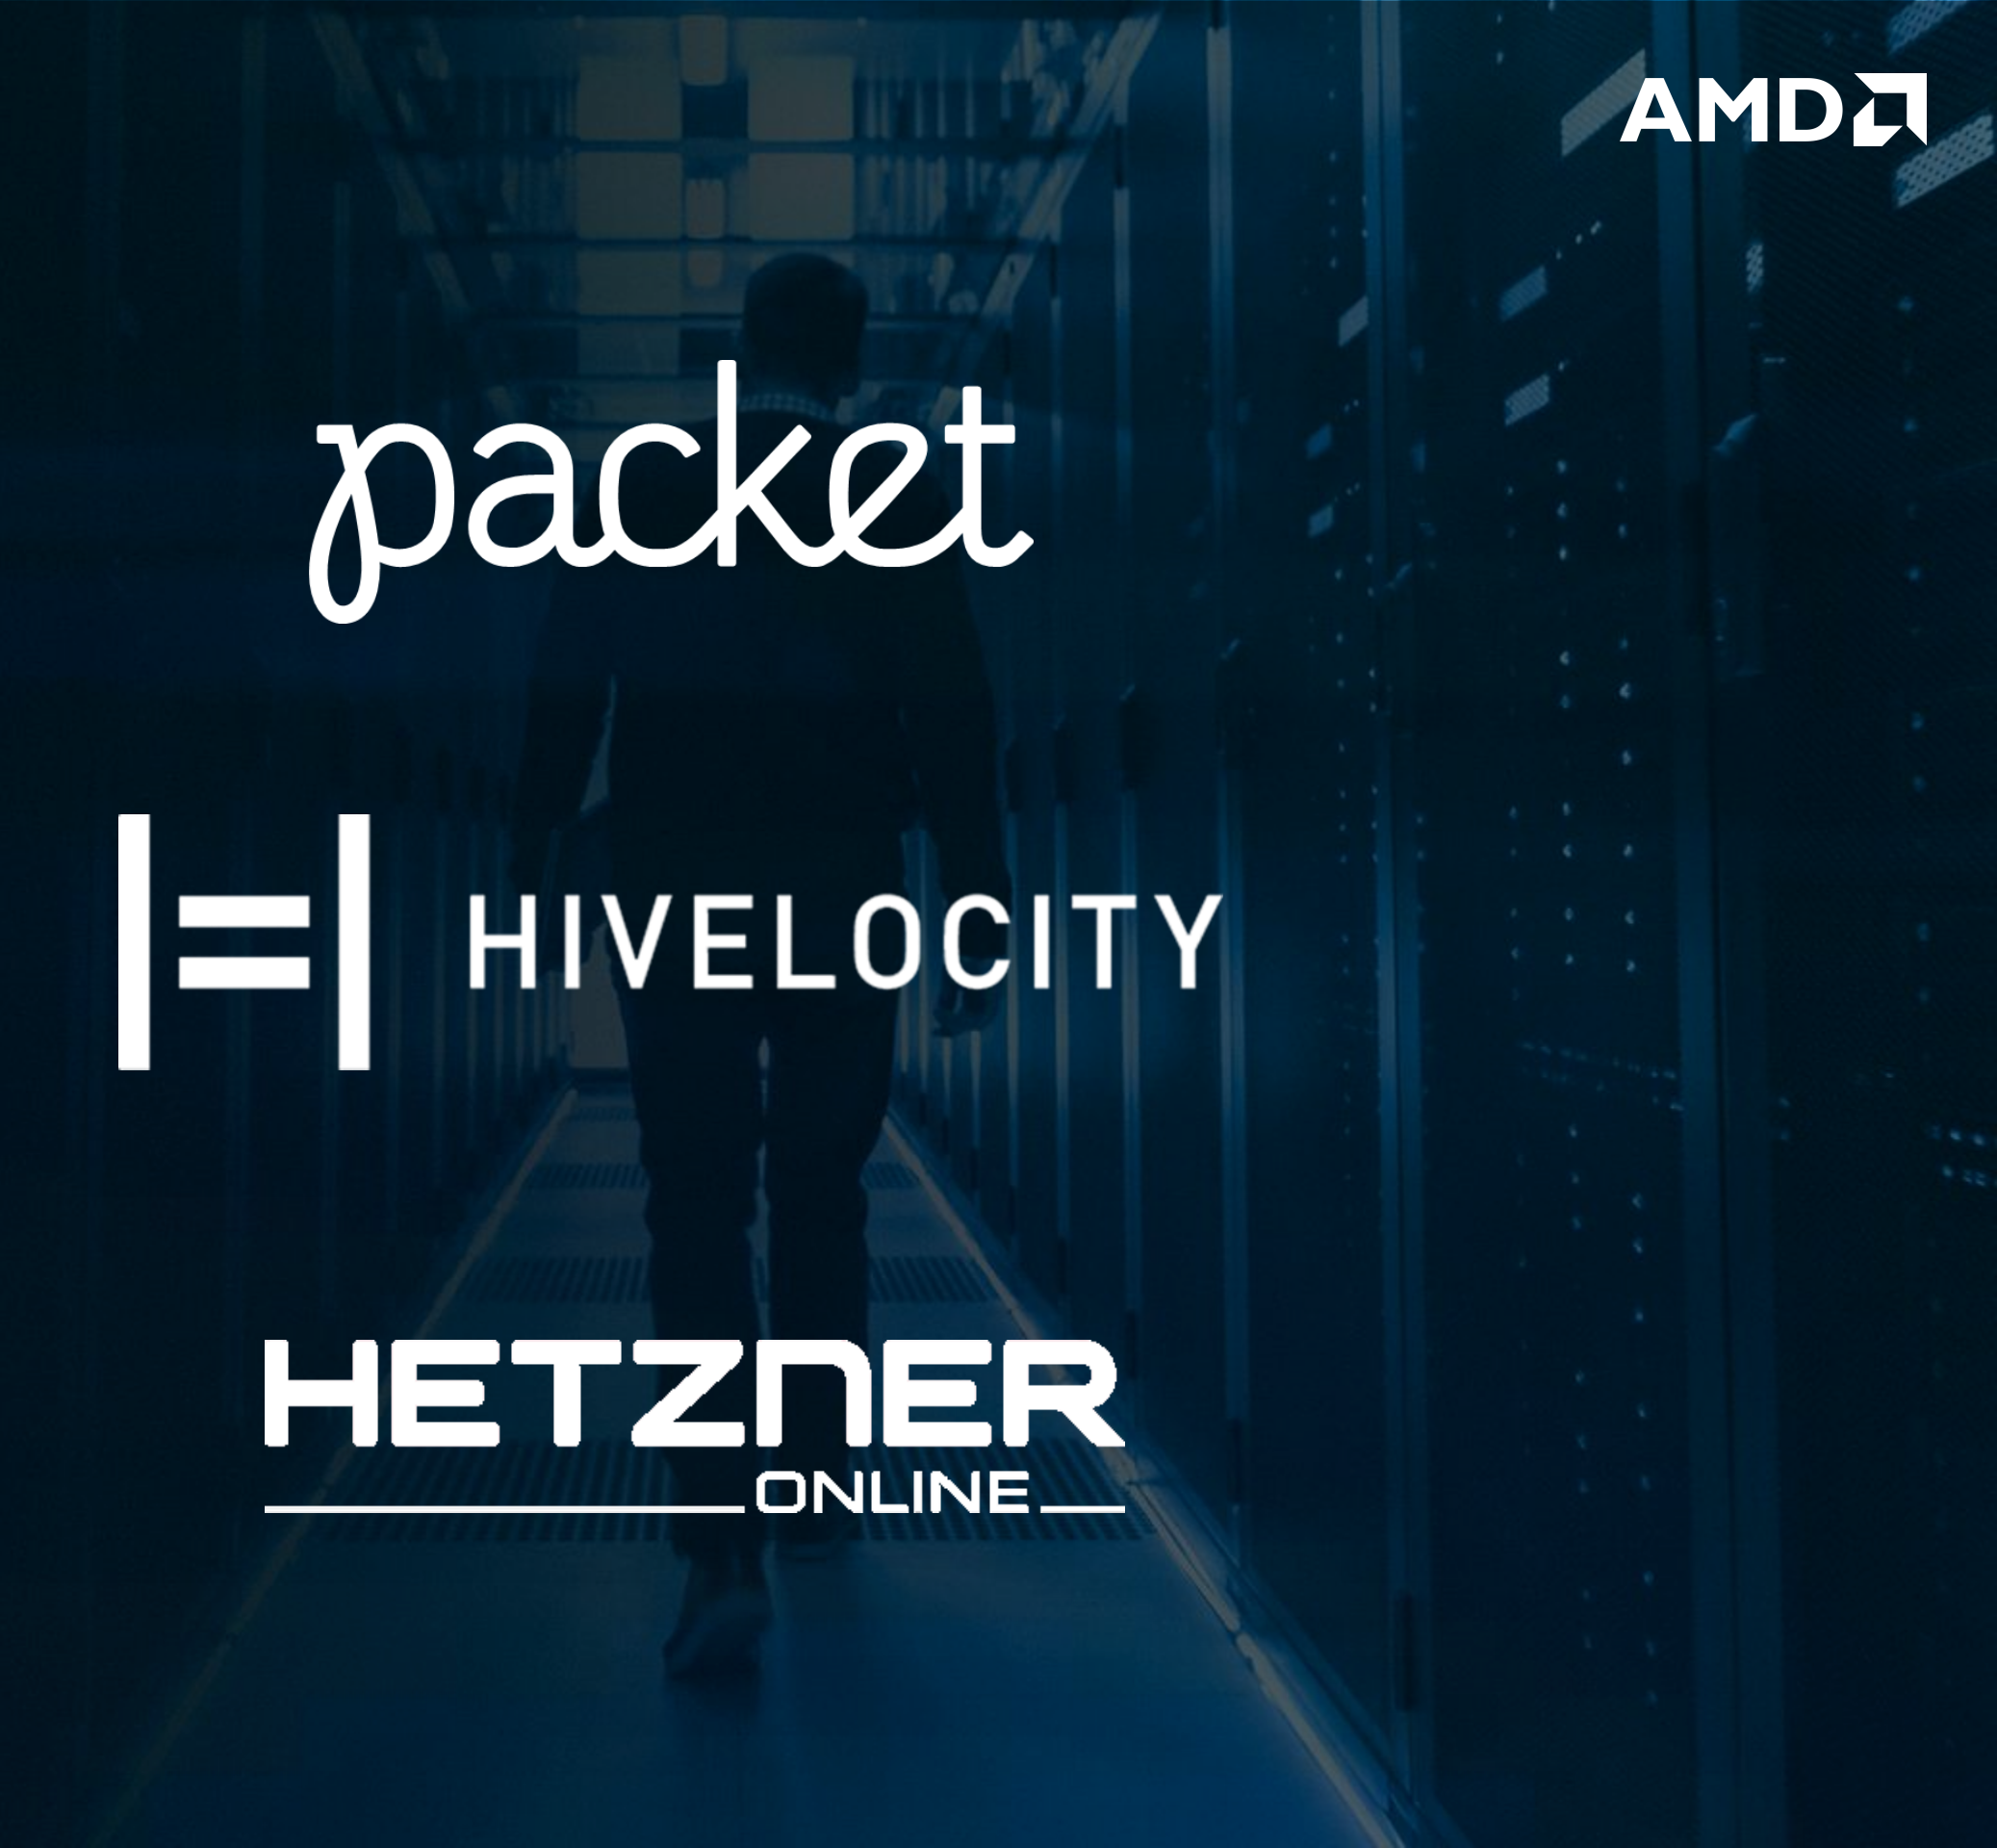


Wide and Robust Range of Enterprise Hosting and Managed Services based on AMD EPYC

packet

HIVELOCITY

HETZNER  
ONLINE



# AI, BIG DATA, CLOUD “ABC” DATACENTERS



Unrivaed Scale and Efficiency for AI,  
Big Data, and Cloud Computing Services

Single Socket Efficiencies Delivering  
Dual Socket Performance

Baidu 百度

# INFRASTRUCTURE AS A SERVICE



Exceptional Performance and 30%  
Lower Cost Per Virtual Machine

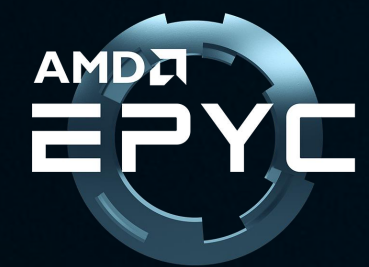
Balanced Computing, Memory  
and Network Resources

Extremely Cost Effective for a Wide  
Range of Enterprise Applications

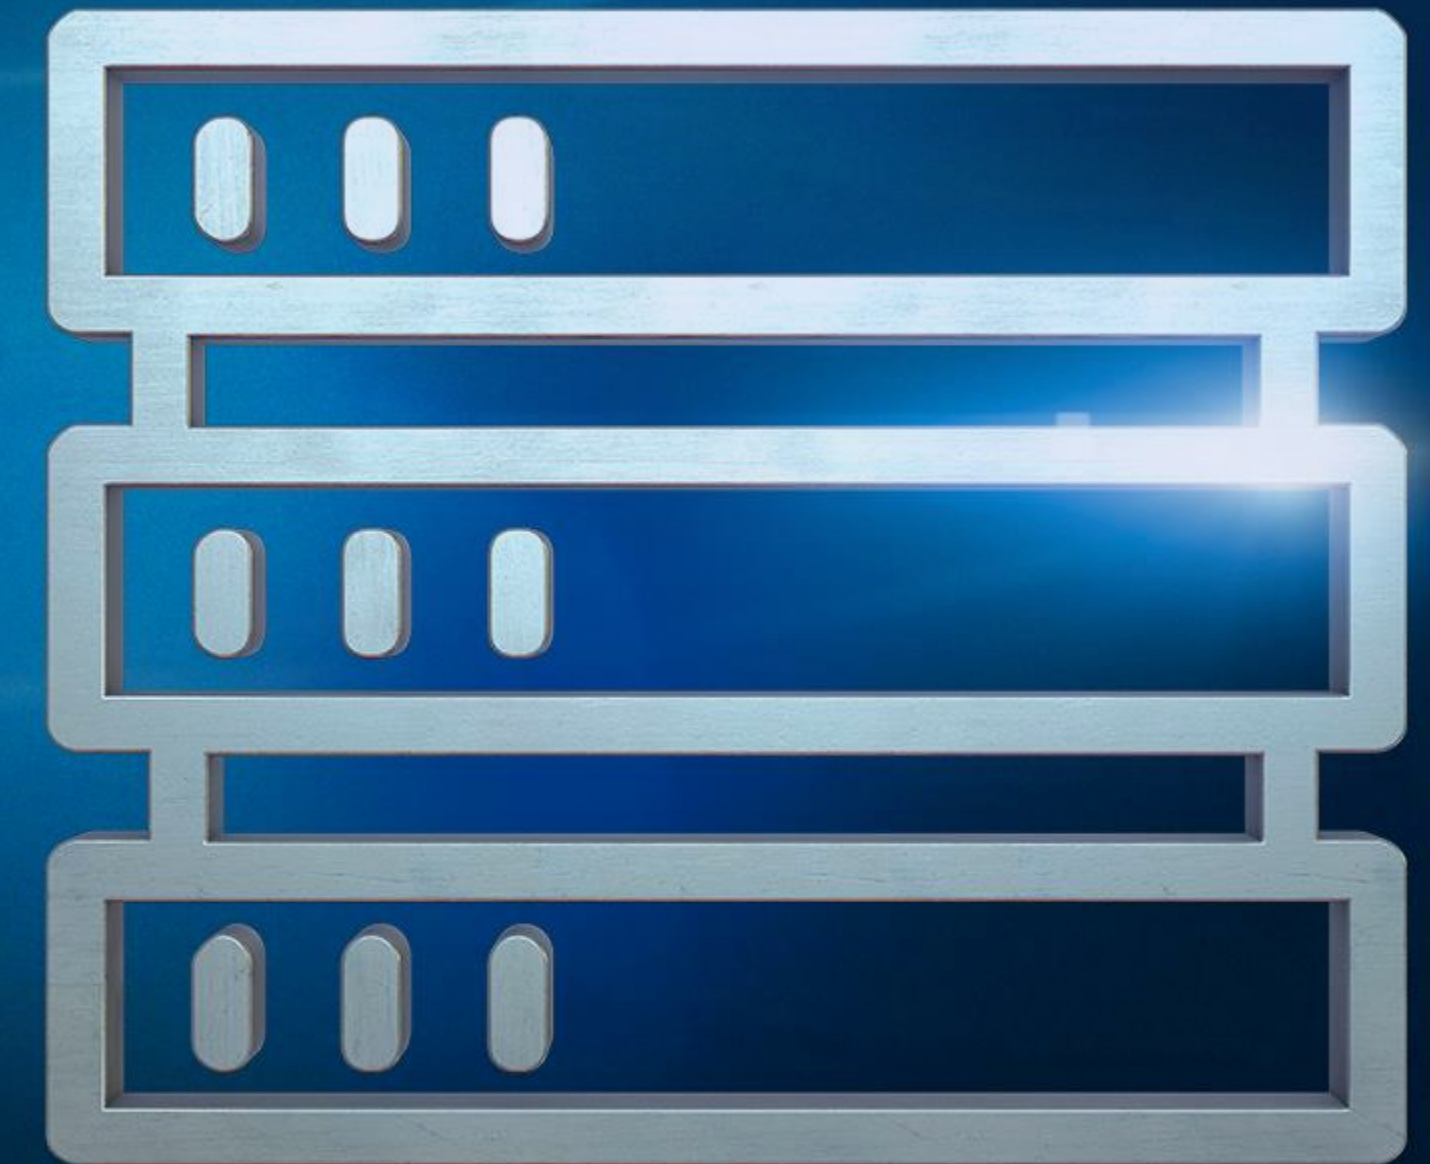


## Tencent Cloud

# AMD EPYC™ DESIGNED FOR VIRTUALIZATION



Reclaiming Budget to Revitalize and Modernize the Datacenter



# AMD EPYC™ VIRTUALIZATION LEADERSHIP AND INNOVATION

## SUPERIOR TCO

Best-in-Class  
Price Performance

## HIGHER DENSITY

Up to 2.7X Greater Virtual  
Machine Density

REVITALIZING THE HEART OF THE  
**VIRTUALIZED DATACENTER**

# REVITALIZING THE HEART OF THE VIRTUALIZED DATACENTER

## TYPICAL 4-YEAR OLD DEPLOYMENT

96 Dual Socket 2U Servers

~1900 Cores/Virtual Machines

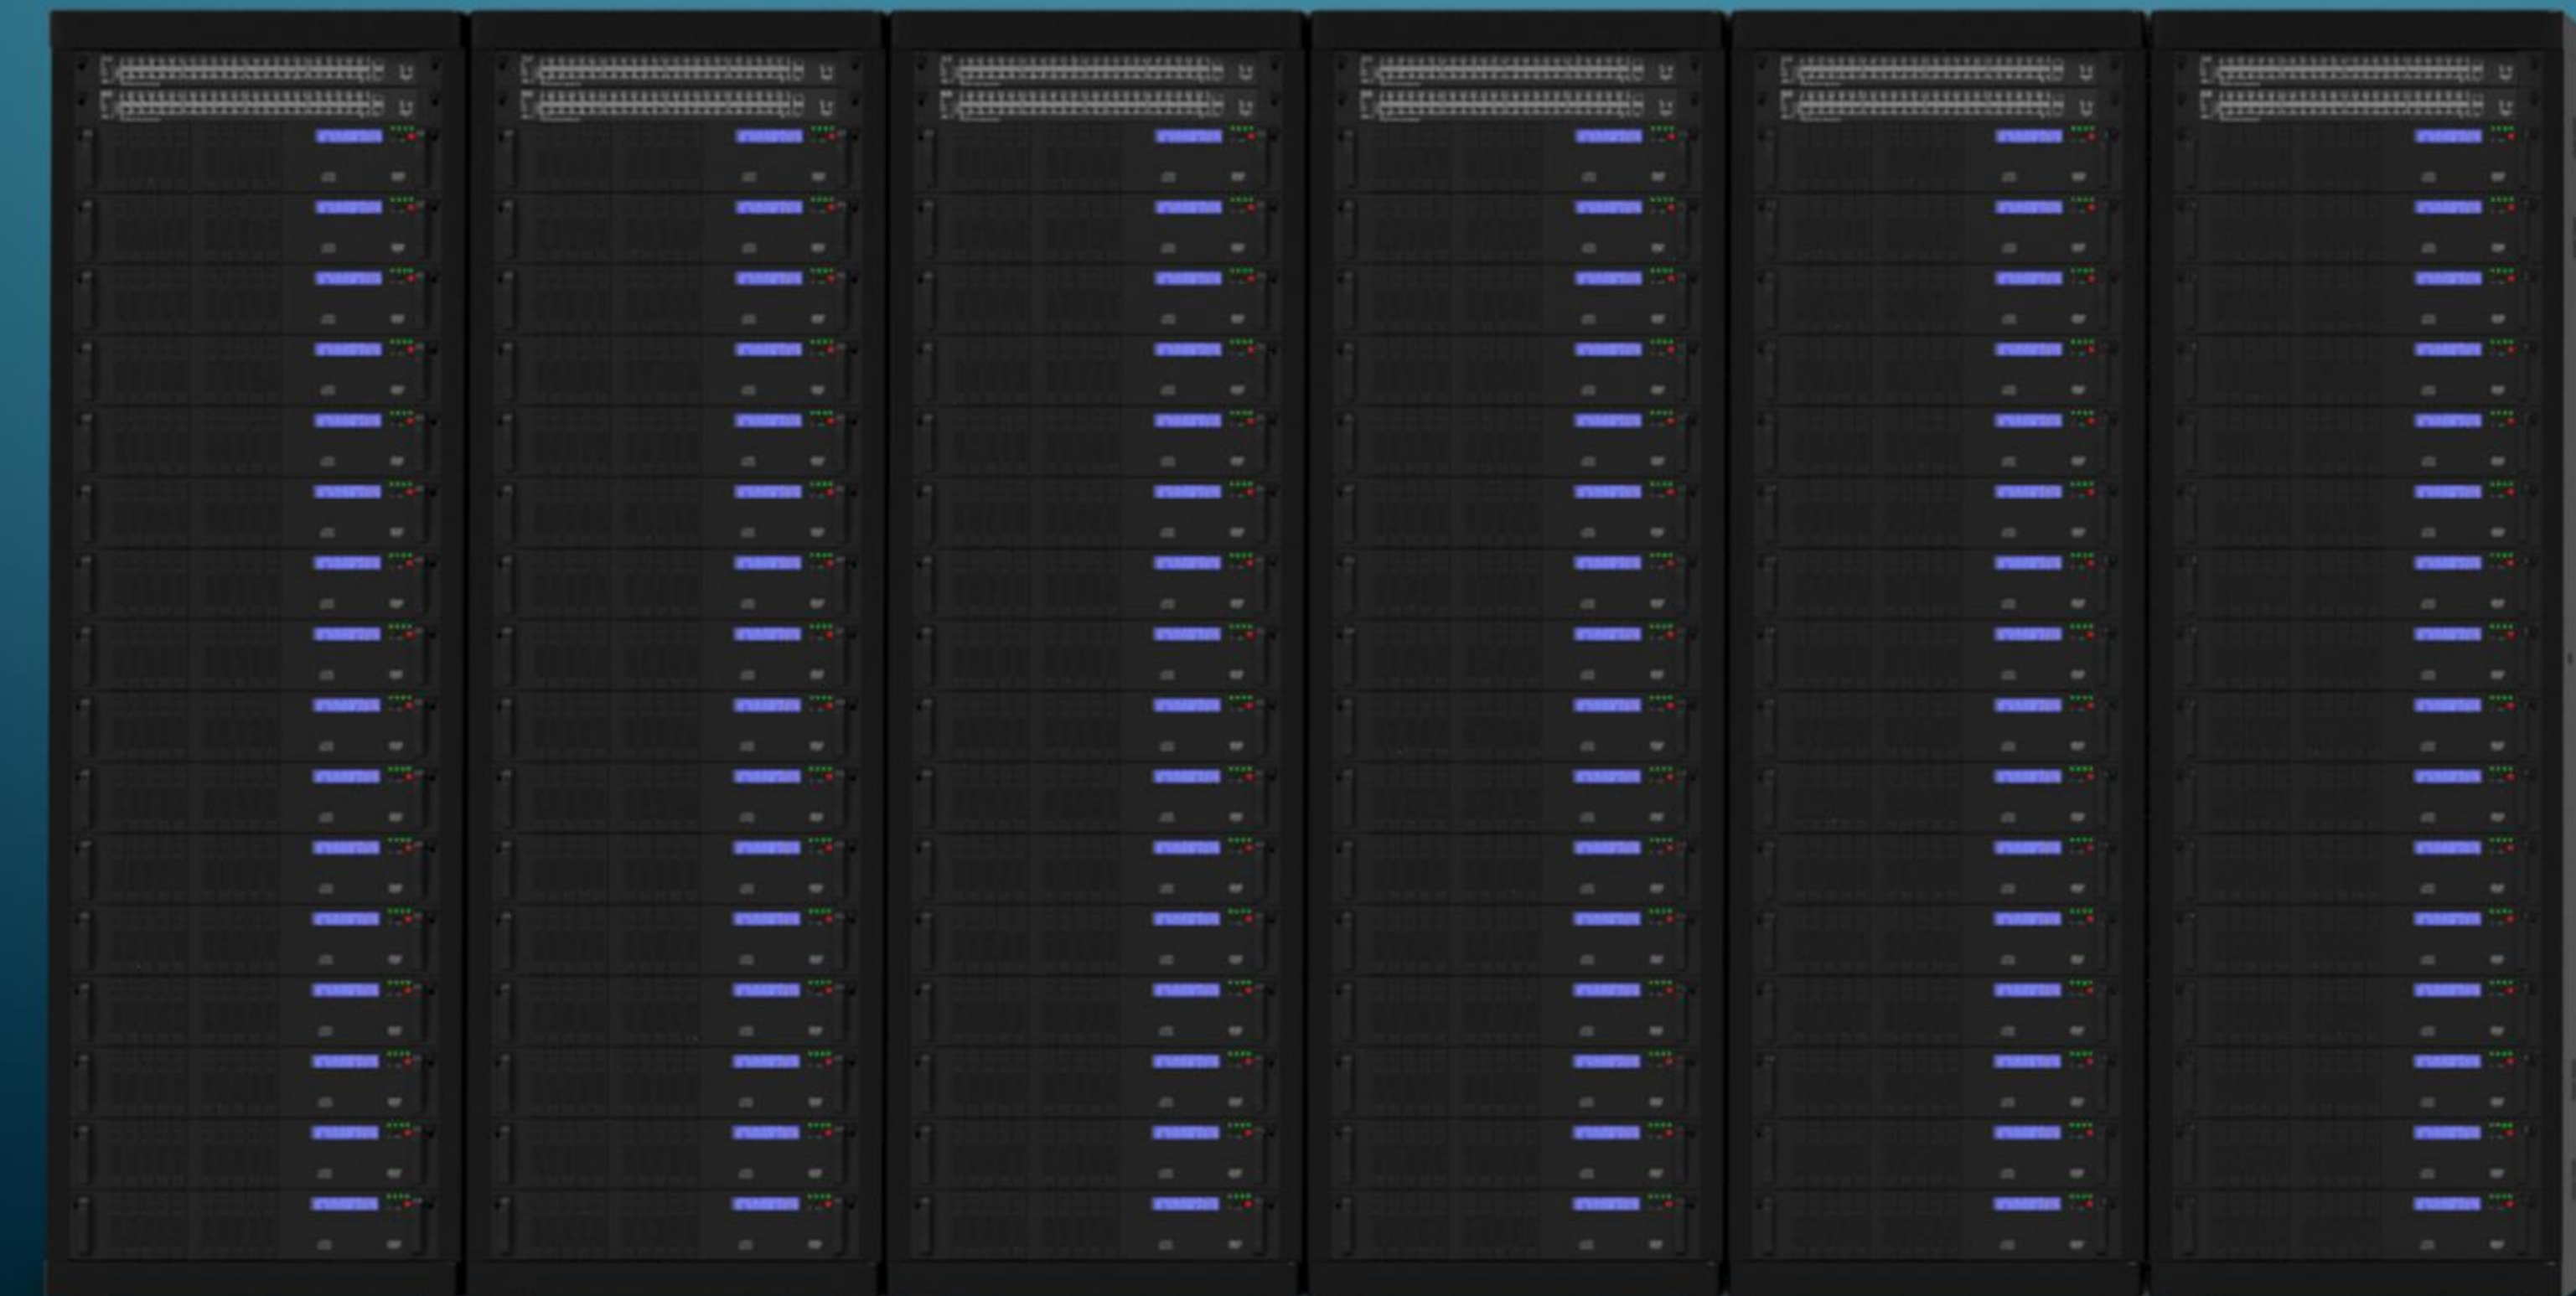
Intel Xeon E5-2660v3

~70%

ADMINISTRATIVE  
COSTS

~30%

SERVER  
PERFORMANCE



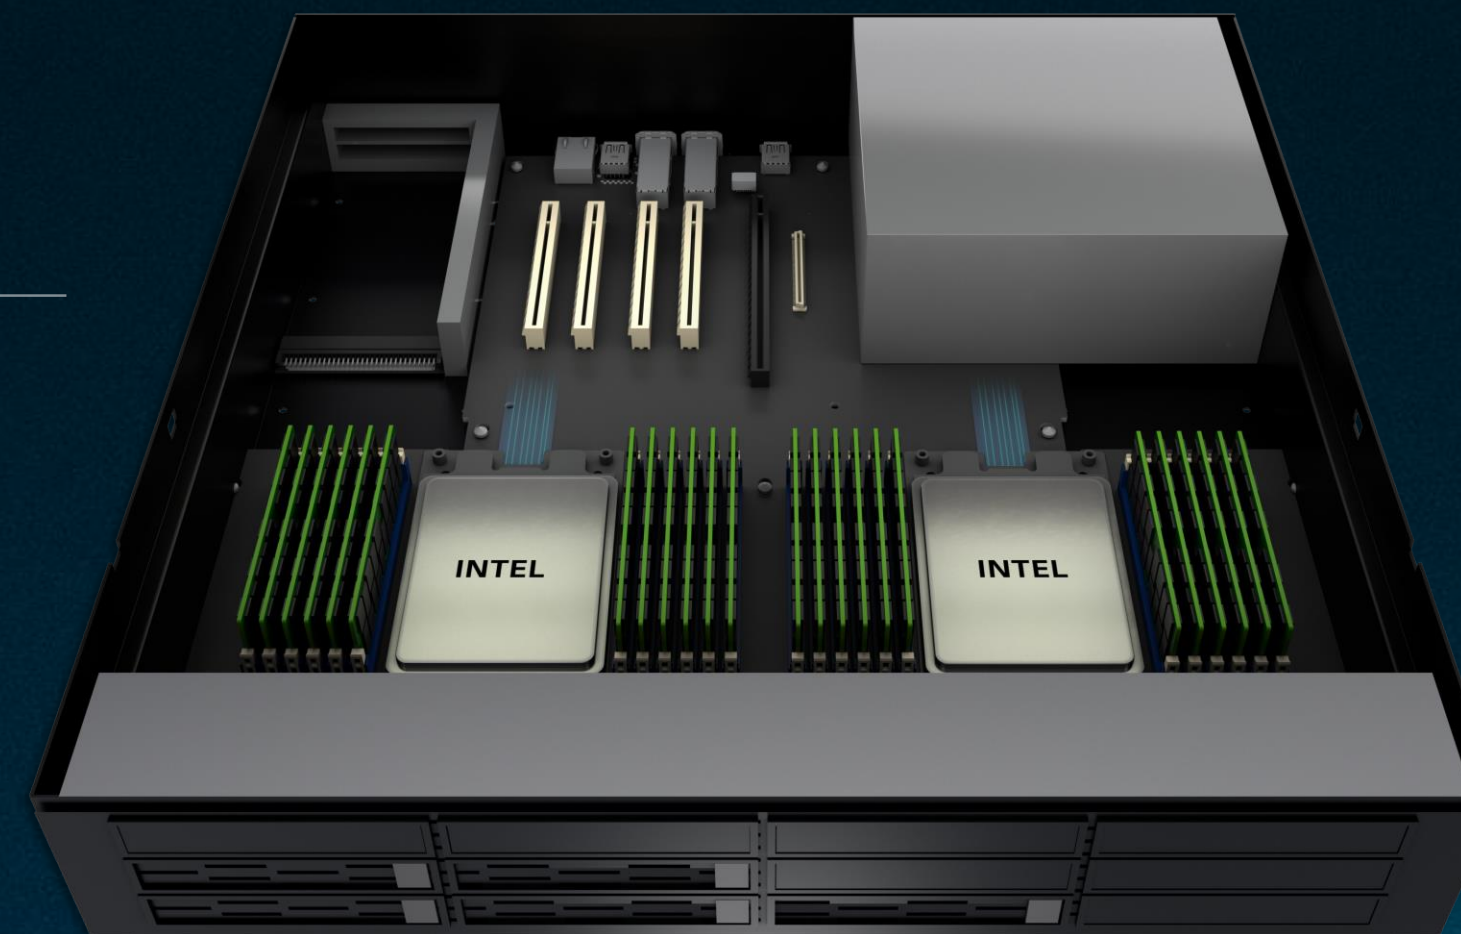
# RETHINKING THE VIRTUAL SERVER

## TRADITIONAL DUAL SOCKET

Dual Intel Gold 5118  
Processors

24 Cores

1.5TB

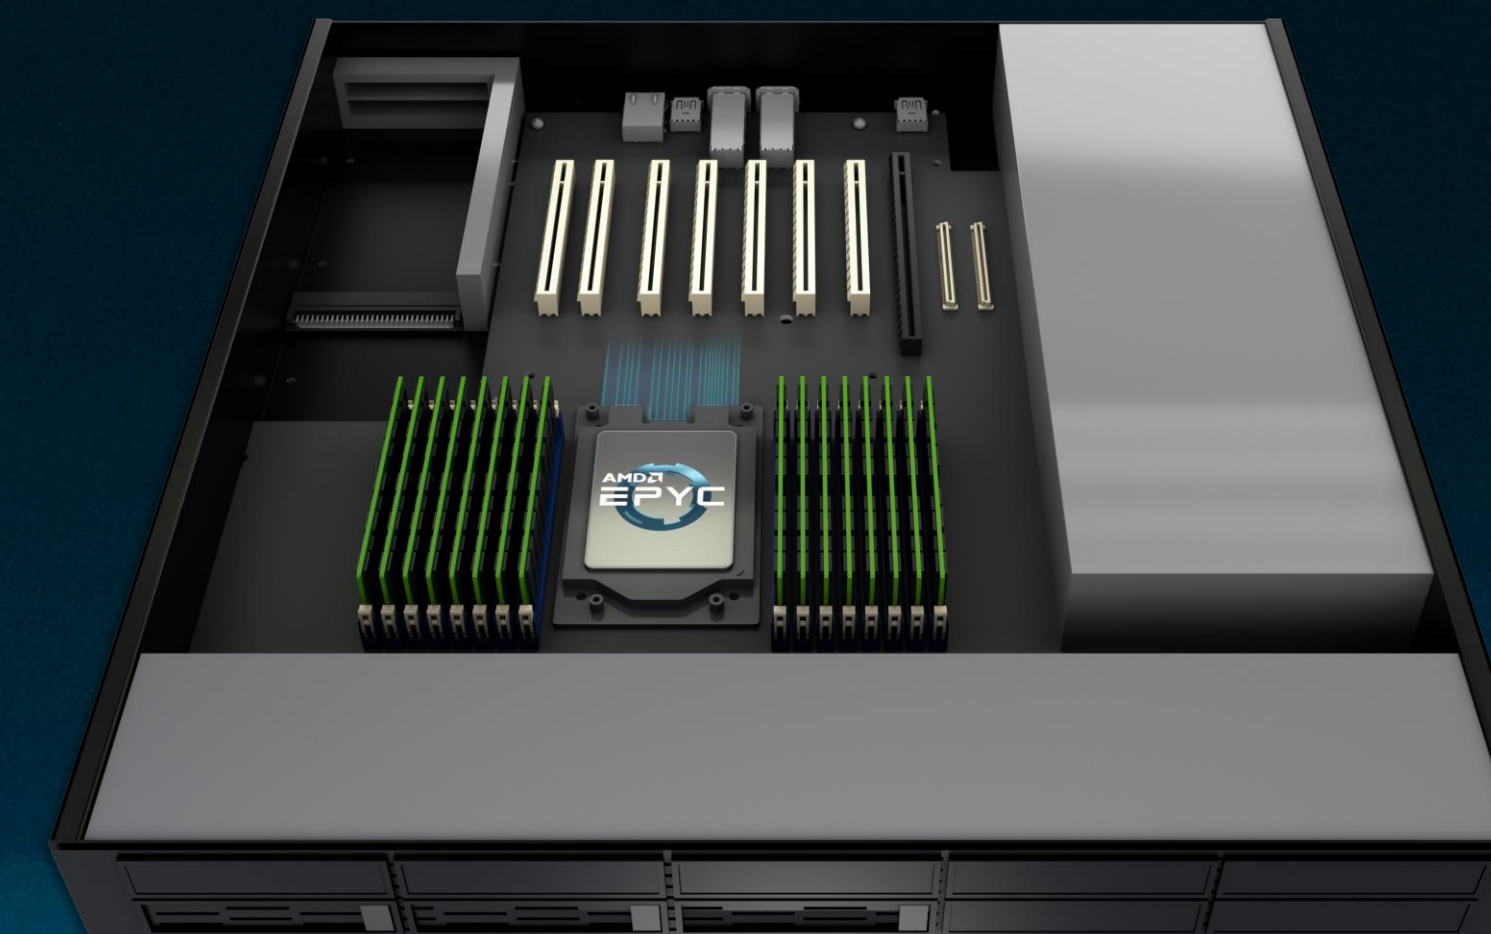


## NO COMPROMISE SINGLE SOCKET EPYC™

Single Socket AMD EPYC 7551P  
Processor

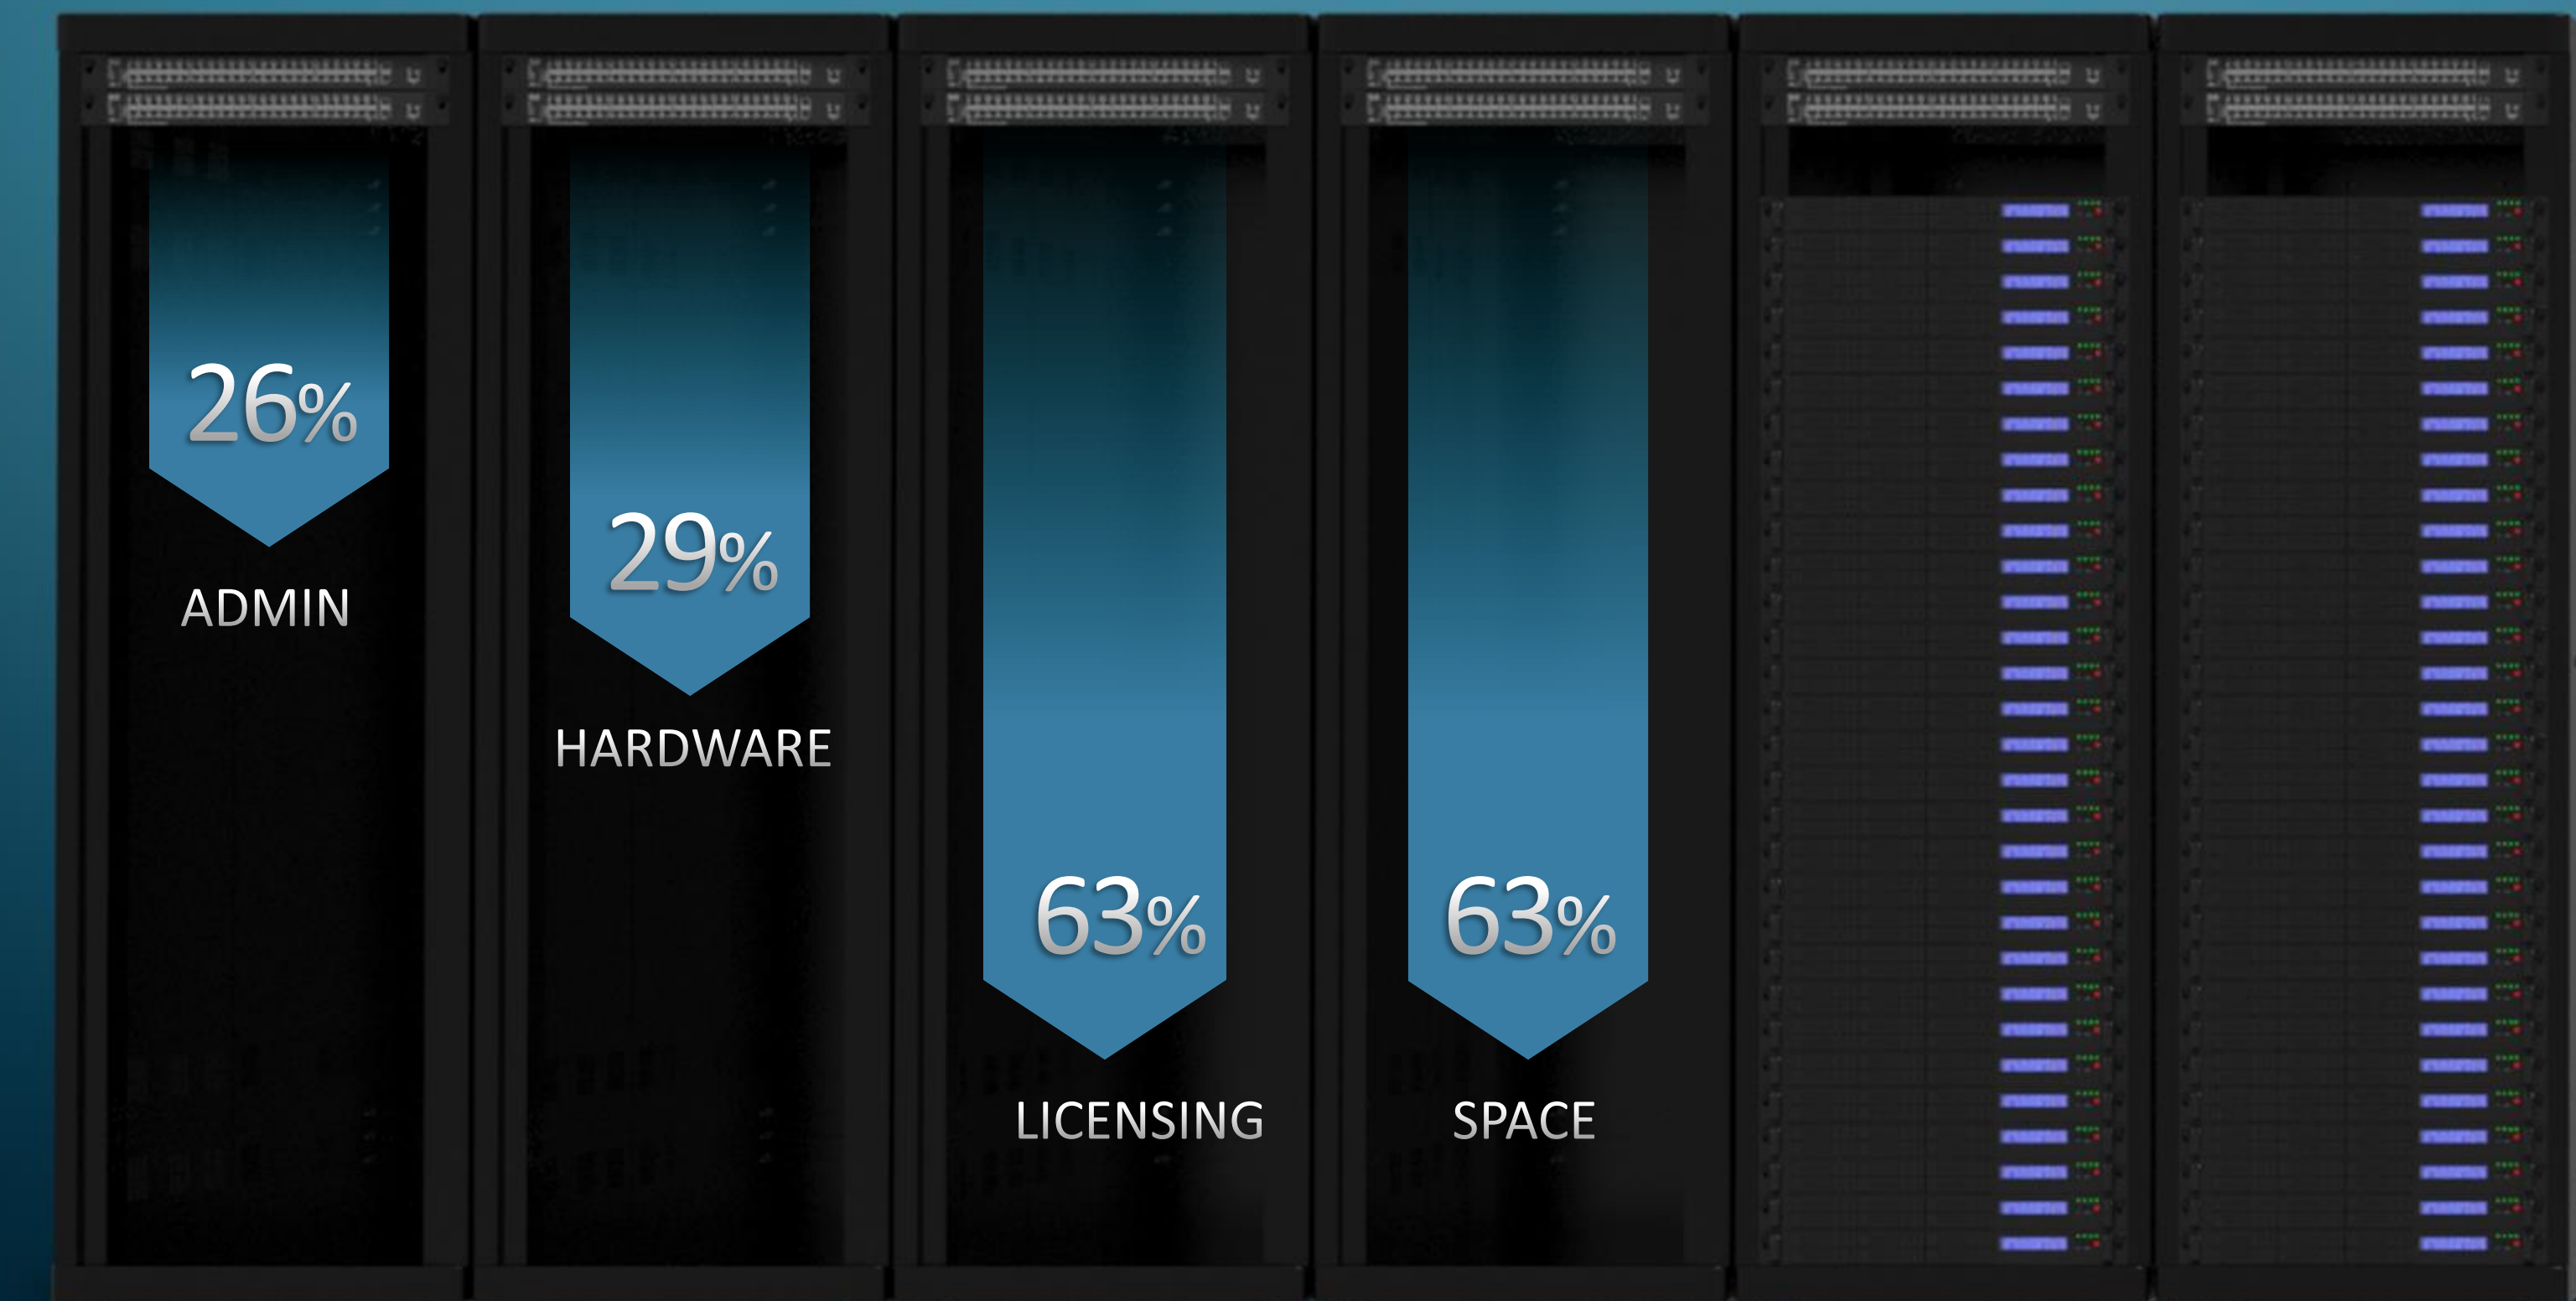
32 Cores

2TB



# REVITALIZING THE HEART OF THE VIRTUALIZED DATACENTER

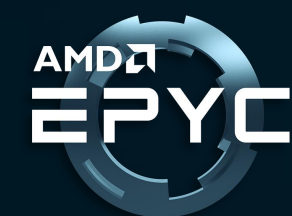
EPYC SINGLE SOCKET = DISRUPTIVE TCO



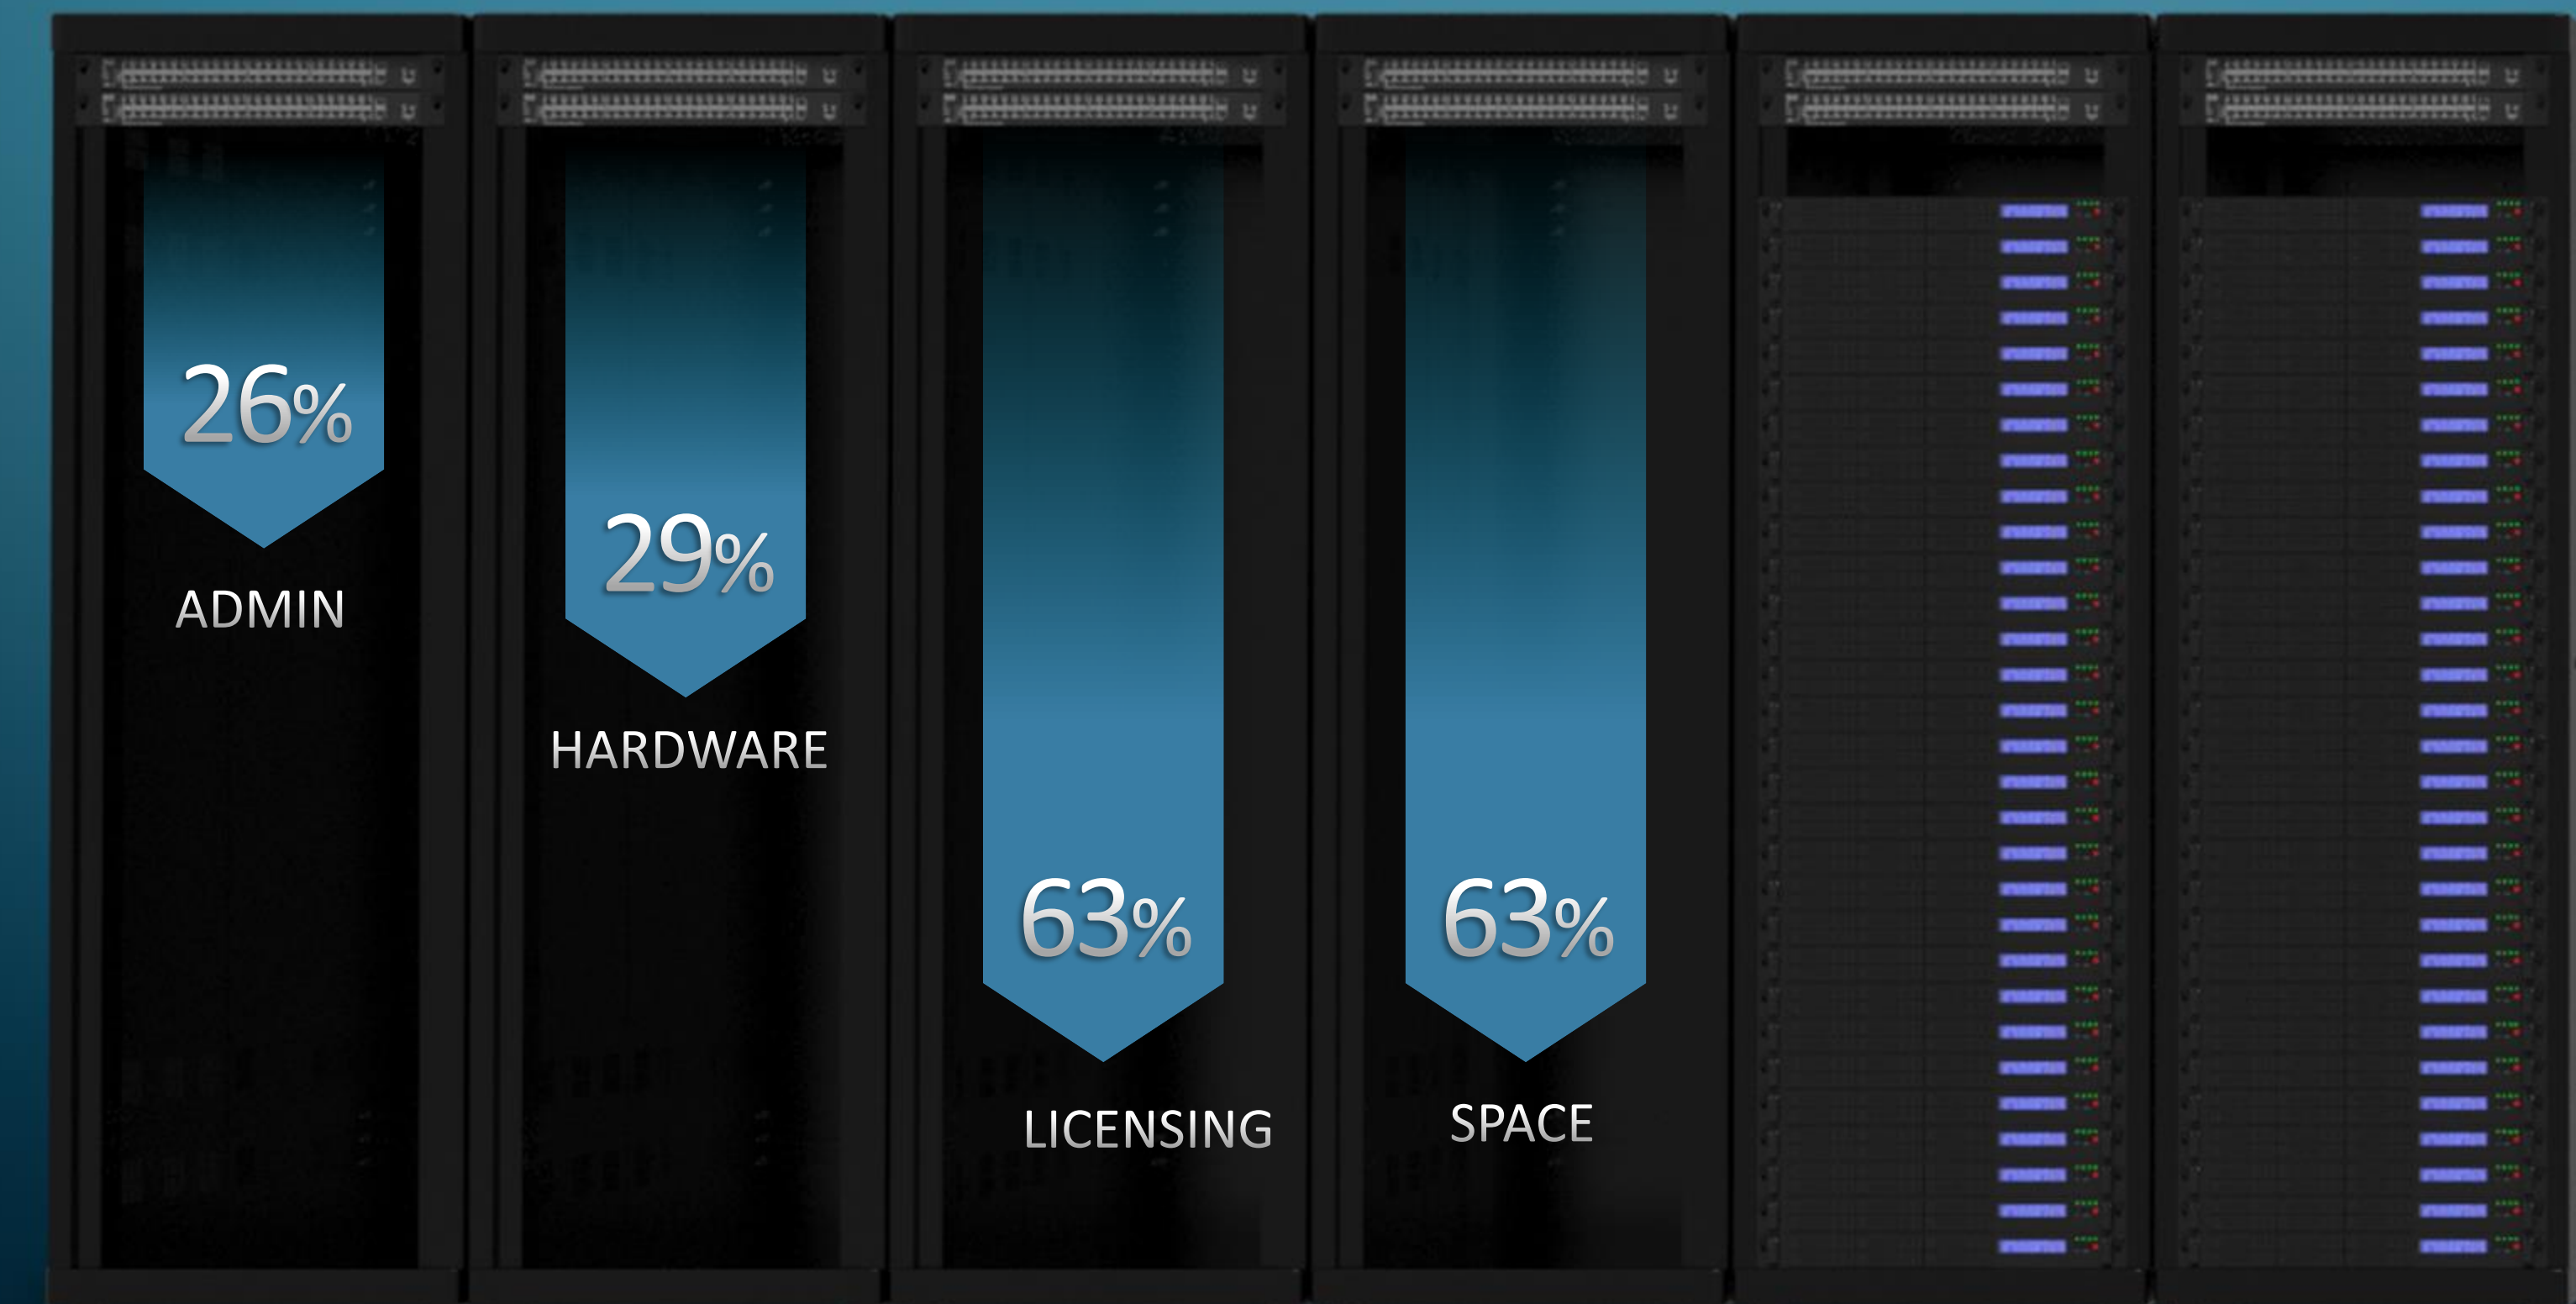
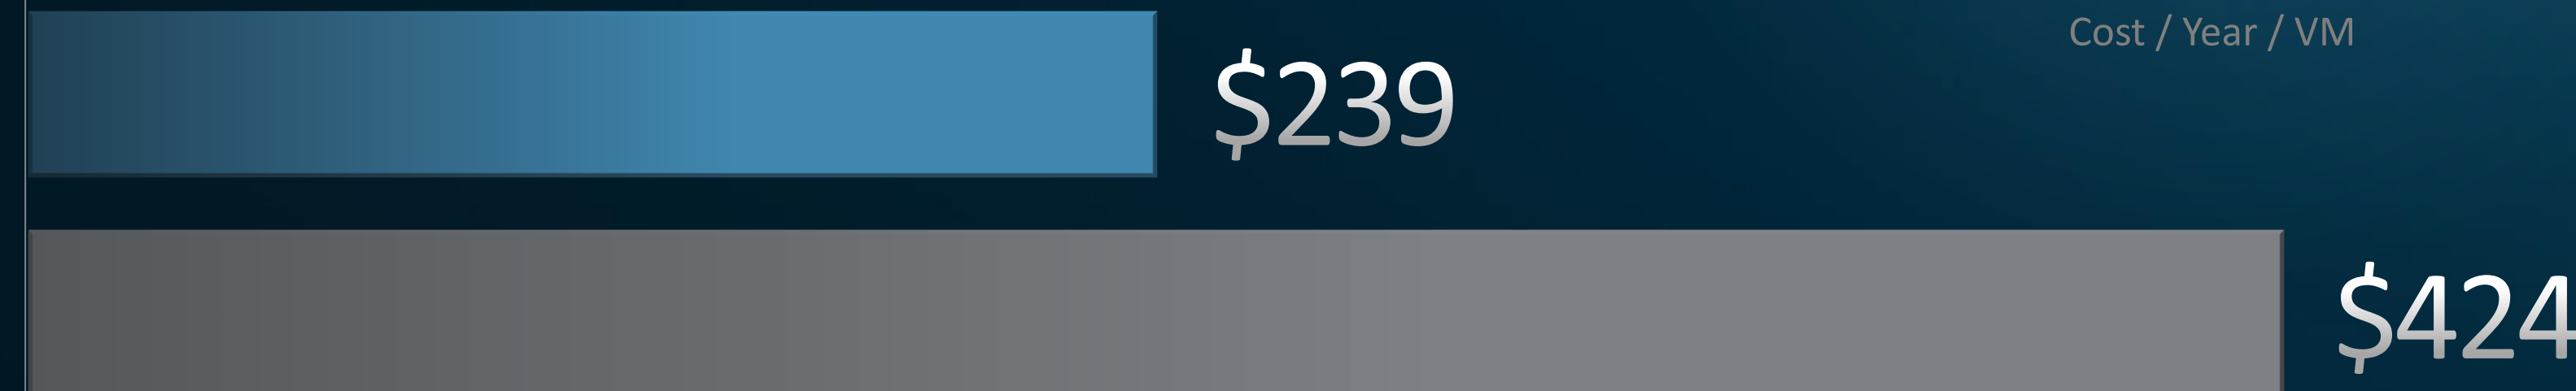
# REVITALIZING THE HEART OF THE VIRTUALIZED DATACENTER

REDUCE TCO UP TO **45%**

Cost / Year / VM

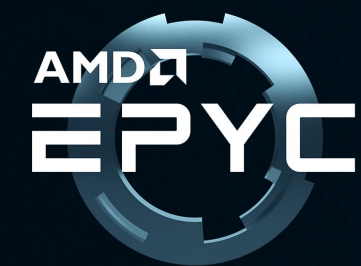


Intel



AMD EPYC™ DESIGNED FOR

# HIGH PERFORMANCE COMPUTING



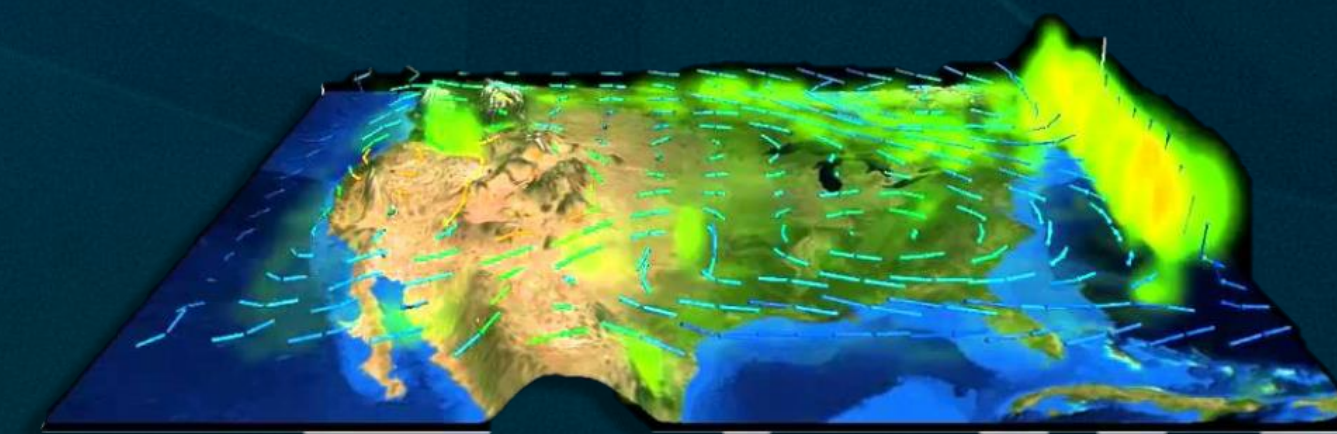
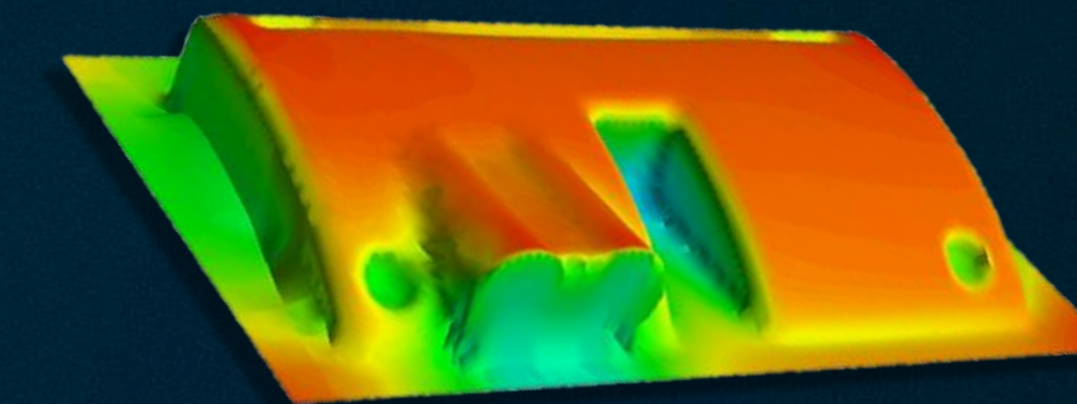
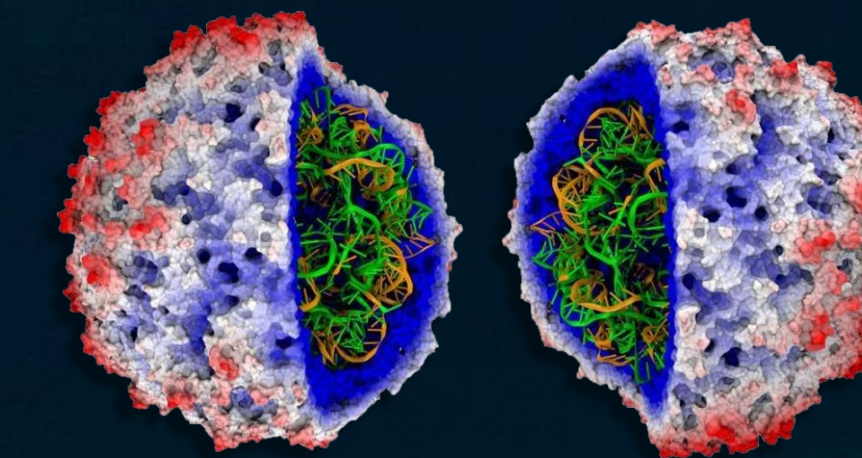
Superior Performance for Today's Commercial and Scientific Workloads



#1

FLOATING POINT  
PERFORMANCE

SPECrate2017\_fp\_peak (2P)

44% Faster WRF  
Weather Modeling25% Faster ANSYS  
Fluent41% Faster NAMD  
Molecular Dynamics

# Peter Ungaro

President and CEO

—  
CRAY®

PARTNERSHIP

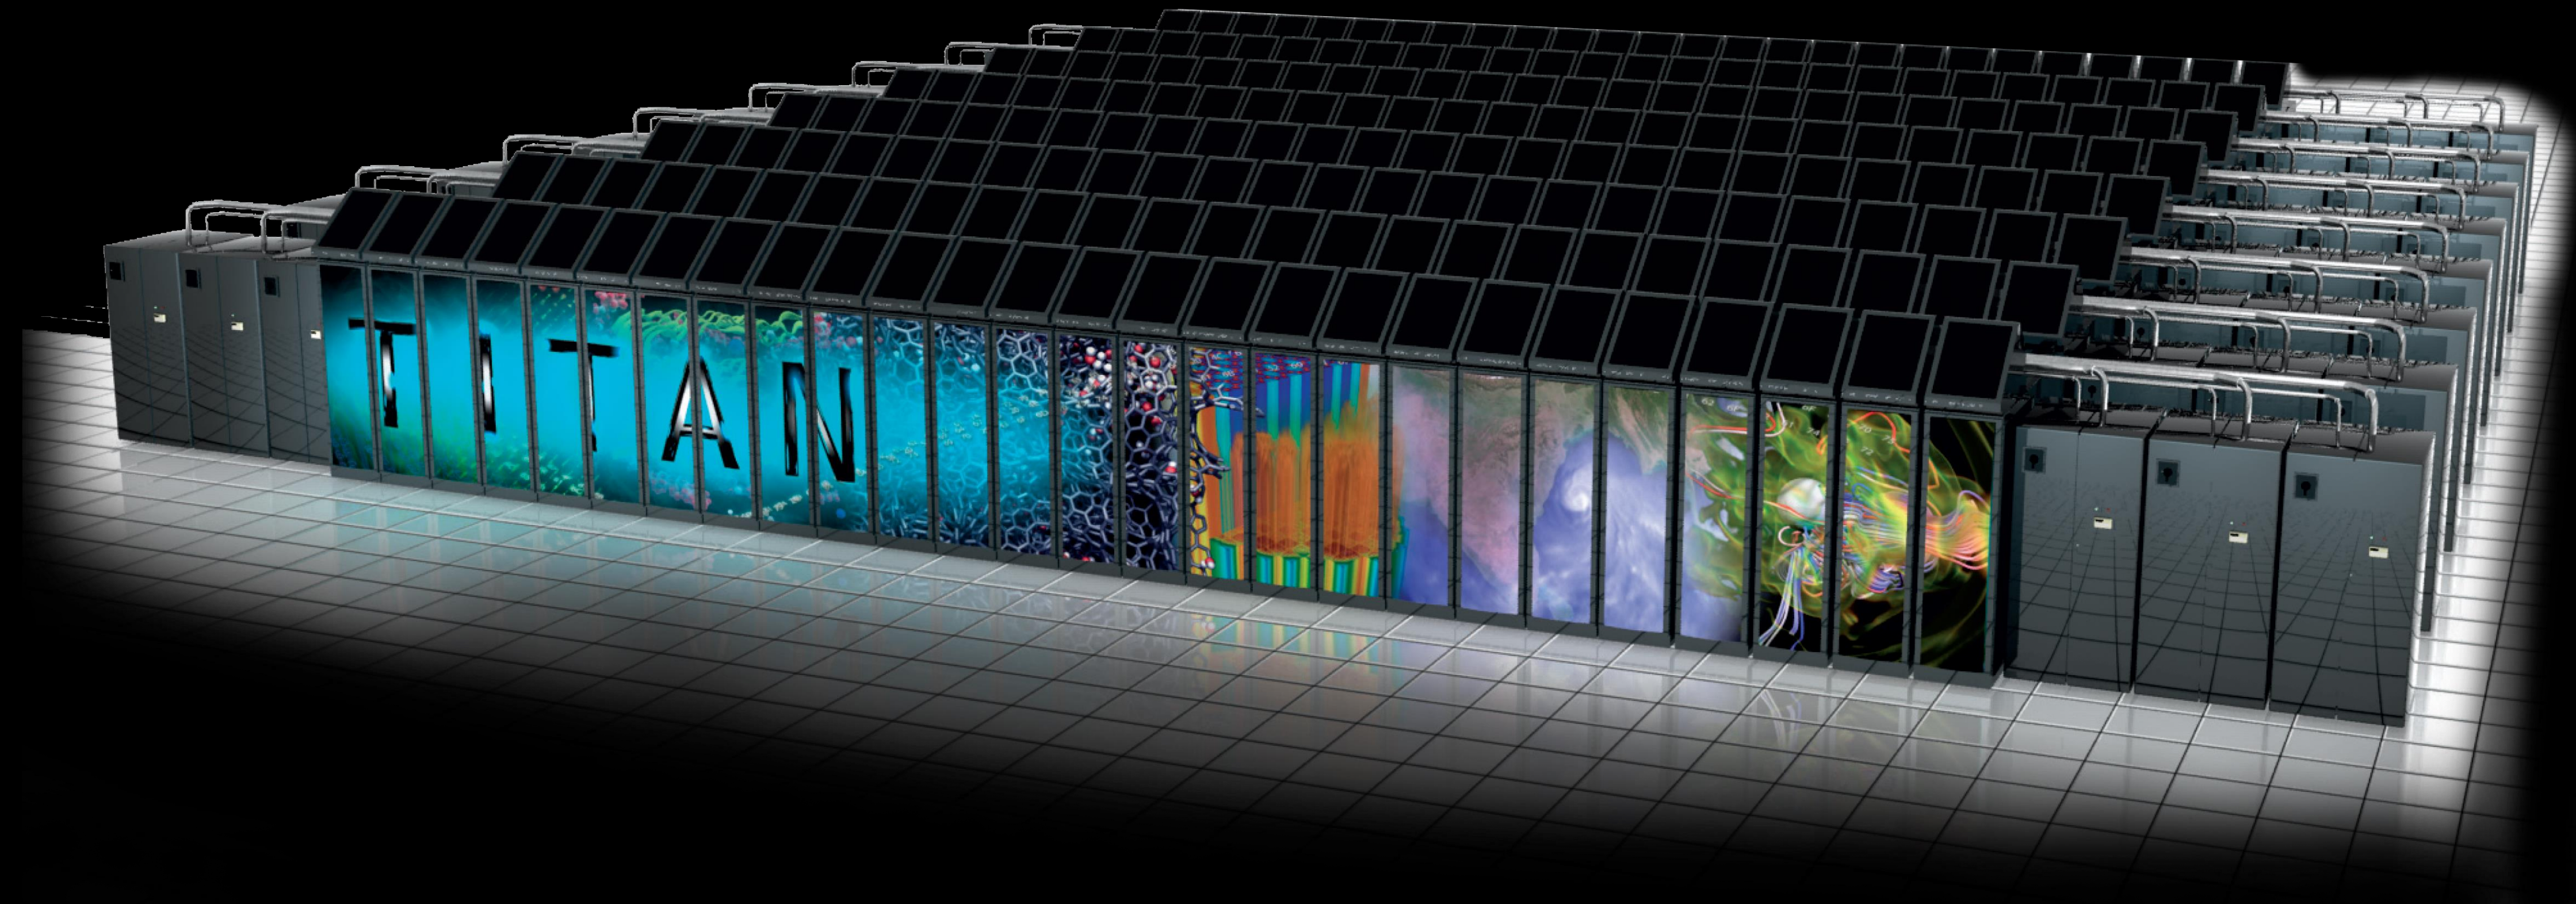
---

CRAY®

AMD

CRAY

AMD



CRAY®



Haas  
F1® Team





CRAY



LAWRENCE BERKELEY  
NATIONAL LABORATORY

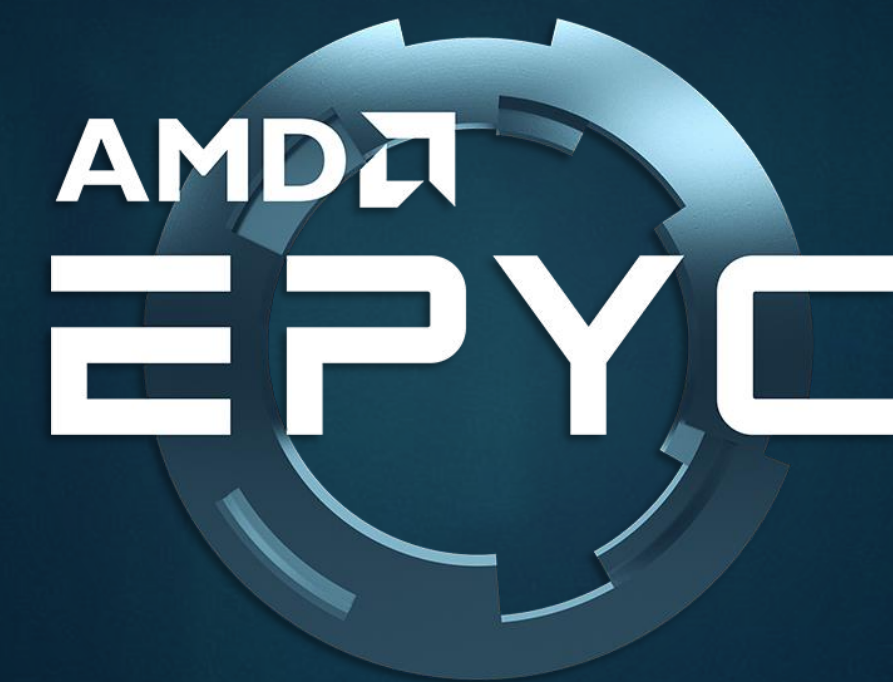


NATIONAL ENERGY RESEARCH  
SCIENTIFIC COMPUTING CENTER

# THE ROAD TO EXASCALE

CRAY<sup>®</sup>

AMD



WE'RE JUST  
GETTING STARTED

# ENDNOTES

NAP-42 – AMD EPYC™ 7601 processor supports up to 8 channels of DDR4-2667, versus the Xeon Platinum 8180 processor at 6 channels of DDR4-2667.

NAP-43 – AMD EPYC 7601 processor includes up to 32 CPU cores versus the Xeon Platinum 8180 processor with 28 CPU cores.

NAP-44 – A single AMD EPYC™ 7601 processor offers up to 2TB/processor (x 2 = 4TB), versus a single Xeon Platinum 8180 processor at 768Gb/processor (x 2 = 1.54TB). NAP-56 – AMD EPYC™ processor supports up to 128 PCIe® Gen 3 I/O lanes (in both 1 and 2-socket configuration), versus the Intel® Xeon® SP Series processor supporting a maximum of 48 lanes PCIe® Gen 3 per CPU, plus 20 lanes in the chipset (max of 68 lanes on 1 socket and 116 lanes on 2 socket).

Microsoft Azure: “Highest Throughput of any cloud” Based on “33 percent more connectivity than available two-socket solutions to address an unprecedented number of NVMe drives directly”. <https://www.amd.com/en/press-releases/microsoft-azure-becomes-2017dec05>

“Oracle : Up to 66% less cost per core hour” Source: <https://blogs.oracle.com/cloud-infrastructure/announcing-the-launch-of-amd-epyc-instances> Cost savings not verified independently by AMD

Baidu: Single Socket Efficiencies Delivering Dual Socket Performance Source: <https://www.amd.com/en/press-releases/baidu-deploys-amd-2017dec13>

Tencent: 30% lower cost per VM Source: <https://segmentnext.com/2018/06/25/amd-epyc-sa1-server/> Cost savings not verified independently by AMD

# ENDNOTES

“Superior TCO, Best in class price/performance & Higher Density” Best in class price performance based on SPECvirt\_sc2013 results and configuration details for the HPE ProLiant DL385 Gen10 and Lenovo ThinkSystem SR650. The HPE ProLiant DL385 scored 2,958 with 168 vm’s at a price of \$30.32 per vm versus the Lenovo ThinkSystem SR650 scoring 3,376 with 189VM’s at a price of \$39.69 per VM. Configuration and details for the HPE ProLiant DL385 Gen10: 2 x 32-core AMD EPYC 7601 processors, 16 x 64 GB PC4-2666 MHz DDR4 LRDIMMs, 6 x 1.6 TB 12G SAS SSDs, 1 x HPE SA P408i-a and 1 x HPE SA P408i-p adapters, 4 x 10GbE 2-port 562 SFP+ adapters. Price of \$89,699. Specvirt\_sc2013 score of 2,958@168 VMs. Price Performance of \$30.32. Configuration and details for the Lenovo ThinkSystem SR650: 2 x 28-core Xeon Platinum 8180 processors, 24 x 32 GB PC4-2666 MHz RDIMMs, 8 x 3.84 TB 12G SAS SSDs, 2 x ThinkSystem RAID 930-8i-2GB 12Gb Adapters, 2 x Intel X710-DA2 10GbE 2-port adapters. Specvirt\_sc2013 score of 3,376@189 VMs. Price Performance of \$30.32 \$39.69. NAP 105

“70% increased maintenance and 30% decreased performance” Source: Opinion July of 2016 “Why Upgrade Your Server Infrastructure Now” Data not independently verified by AMD

# ENDNOTES

Revitalizing the Heart of the Virtualized Datacenter - 26% Less Administration Costs; 29% Less Hardware Costs; 63% Less Licensing Costs; 63% Less Space; Reduce TCO up to 45%; \$239 per vm per year; \$424 per vm per year: Compares 3-year total cost of ownership for 1920 virtual machines of 1 virtual machine per core with 8GB of DRAM per core with similar performance. (80) PowerEdge R740 Configuration Chassis with up to 8 x 3.5" SAS/SATA Hard Drives for 2CPU Configuration, with (2) Intel® Xeon® Gold 5118 2.3G, 12C/24T, 10.4GT/s, 16.5M Cache, Turbo, HT (105W) DDR4-2400, Performance Optimized Memory Capacity of (12) 16GB RDIMM, 2666MT/s, Dual Rank DIMMs, Hard Drives including (1) 120GB SSD SATA Boot 6Gbps 512n 2.5in Hot-plug Drive & (6) 480GB SSD SAS Mix Use 12Gbps 512n 2.5in Hot-plug Drives, No Operating System 3 Years Basic Hardware Warranty Repair: 5x10 HW-Only, 5x10 NBD Onsite Deployment Services, No Installation. Configuration priced on 8/30/2018 at [www.dell.com](http://www.dell.com) for \$14,271 per server; \$1,141,680 total Hardware Acquisition Cost. Facilities and administration costs include \$91,450 for power and cooling, \$3,240 for space, and \$231,549 for IT administration costs for a total of \$326,239 over 3 years - calculated based on AMD EPYC™ TCO calculator at <https://uatfast.valuestoryapp.com/AMD/sales>. SW Licensing costs include (80) vSphere with Ops Manager Ent Plus, 2 CPU License, 5Yrs ProSupport + Subscription for \$12.163 each for a total SW licensing costs of \$973,040. Pricing obtained at [www.dell.com](http://www.dell.com) on 8/30/18. **Total 3 yr. Cost of Ownership of \$2,440,959. Total Cost per year per VM of \$424.** (60) PowerEdge R6415 Configuration Chassis with up to 8 x 3.5" SAS/SATA Hard Drives for 2CPU Configuration, with (1) AMD EPYC™ 7551P 2.00GHz/2.55GHz, 32C/64T, 64M Cache (180W) DDR4-2666, Performance Optimized Memory Capacity of (8) 32GB RDIMM, 2666MT/s, Dual Rank DIMMs, Hard Drives including (1) 120GB SSD SATA Boot 6Gbps 512n 2.5in Hot-plug Drive & (6) 400GB SSD SAS Mix Use 12Gbps 512n 2.5in Hot-plug Drives, No Operating System 3 Years Basic Hardware Warranty Repair: 5x10 HW-Only, 5x10 NBD Onsite Deployment Services, No Installation. Configuration priced on 8/30/18 at [www.dell.com](http://www.dell.com) for \$13,585 per server; \$815,100 total Hardware Acquisition Cost. Facilities and administration costs include \$23,760 for power and cooling, \$2,160 for space, and \$171,518 for IT administration costs for a total of \$197,438 over 3 years - calculated based on AMD EPYC™ TCO calculator at <https://uatfast.valuestoryapp.com/AMD/sales>. SW Licensing costs include (60) vSphere with Ops Manager Ent Plus, 1 CPU License, 5Yrs ProSupport + Subscription for \$6,077 each for a total SW licensing costs of \$364,620. Pricing obtained at [www.dell.com](http://www.dell.com) on 8/30/18. Total 3yr Cost of Ownership of \$1,377,158. Total Cost per vm per year of 239.

Superior TCO: of up to 45% at 320 VM's and 44% at 1920 VM's

2.8 greater density: 80 2RU Intel servers = 160 RU. 60 1RU AMD EPYC servers is 60RU. 63% Less Space and 2.7X greater VM density

Hardware costs: \$1,141,680 vs \$815,100 = 29% less

Administration costs: \$231,549 vs \$171,518 = 26% less

Licensing costs: \$973,040 vs \$364,620 = 63% less NAP-104

# ENDNOTES

#1 SPECrate2017\_fp\_peak: AMD has #1 spec\_fp\_peak scores for both 1P and 2P. 1P at <https://www.spec.org/cpu2017/results/res2018q2/cpu2017-20180426-05035.html>. 2P at <https://www.spec.org/cpu2017/results/res2018q2/cpu2017-20180319-04087.html>. SPECrate®2017\_fp\_peak as of Oct 30, 2018. <https://www.spec.org/cpu2017/results/>. NAP106

44% WRF: Internal AMD testing using the WRF v.3.9.1.1 benchmark, 12km model size, compiled with gcc 7.3.0, using OpenMPI v.3.1.1 and NetCDF v.1.1.3 comparing a Supermicro AS -1123US-01-AM036 configured with 2x EPYC 7601 CPUs, 16x16GB DDR4 2666MHz DIMMs and Ubuntu 18.04 vs a Supermicro SYS-1029U-TRTP configured with 2x Intel Xeon Platinum 8180 CPUs, 24x16GB DDR4 2666MHz DIMMs and Ubuntu 18.04

25% ANSYS: Based on Dell internal testing using the ANSYS Fluent benchmark test in November 2017, comparing a Dell PowerEdge R7425 configured with an AMD 7601, 32-core processor to a traditional 2-socket server with an Intel 6148, 20 core processor. Actual performance will vary based on configuration, usage and manufacturing variability. Not verified by AMD.

41% NAMD: Results from AnandTech testing - <https://www.anandtech.com/show/11544/intel-skylake-ep-vs-amd-epyc-7000-cpu-battle-of-the-decade/21>. Results not independently verified by AMD

## DISCLAIMER

The information contained herein is for informational purposes only, and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for particular purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD's products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale. Timelines, roadmaps, and/or product release dates shown in these slides are plans only and subject to change. “Vega” , “Naples”, “Rome”, “Zen”, “Zen2” are codenames for AMD architectures and are not product names. GD-122

## ATTRIBUTION

©2018 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, Radeon, EPYC and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

# CAUTIONARY STATEMENT

This presentation contains forward-looking statements concerning Advanced Micro Devices, Inc. (AMD) including, but not limited to the features, functionality, availability, timing, deployment, and expected benefits of AMD future products and technologies, including “Rome,” “Naples,” “Milan” and AMD Radeon Instinct™ MI60; the benefits and expectations of 7nm process technology; and AMD’s technology 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 Quarterly Report on Form 10-Q for the quarter ended September 29, 2018.

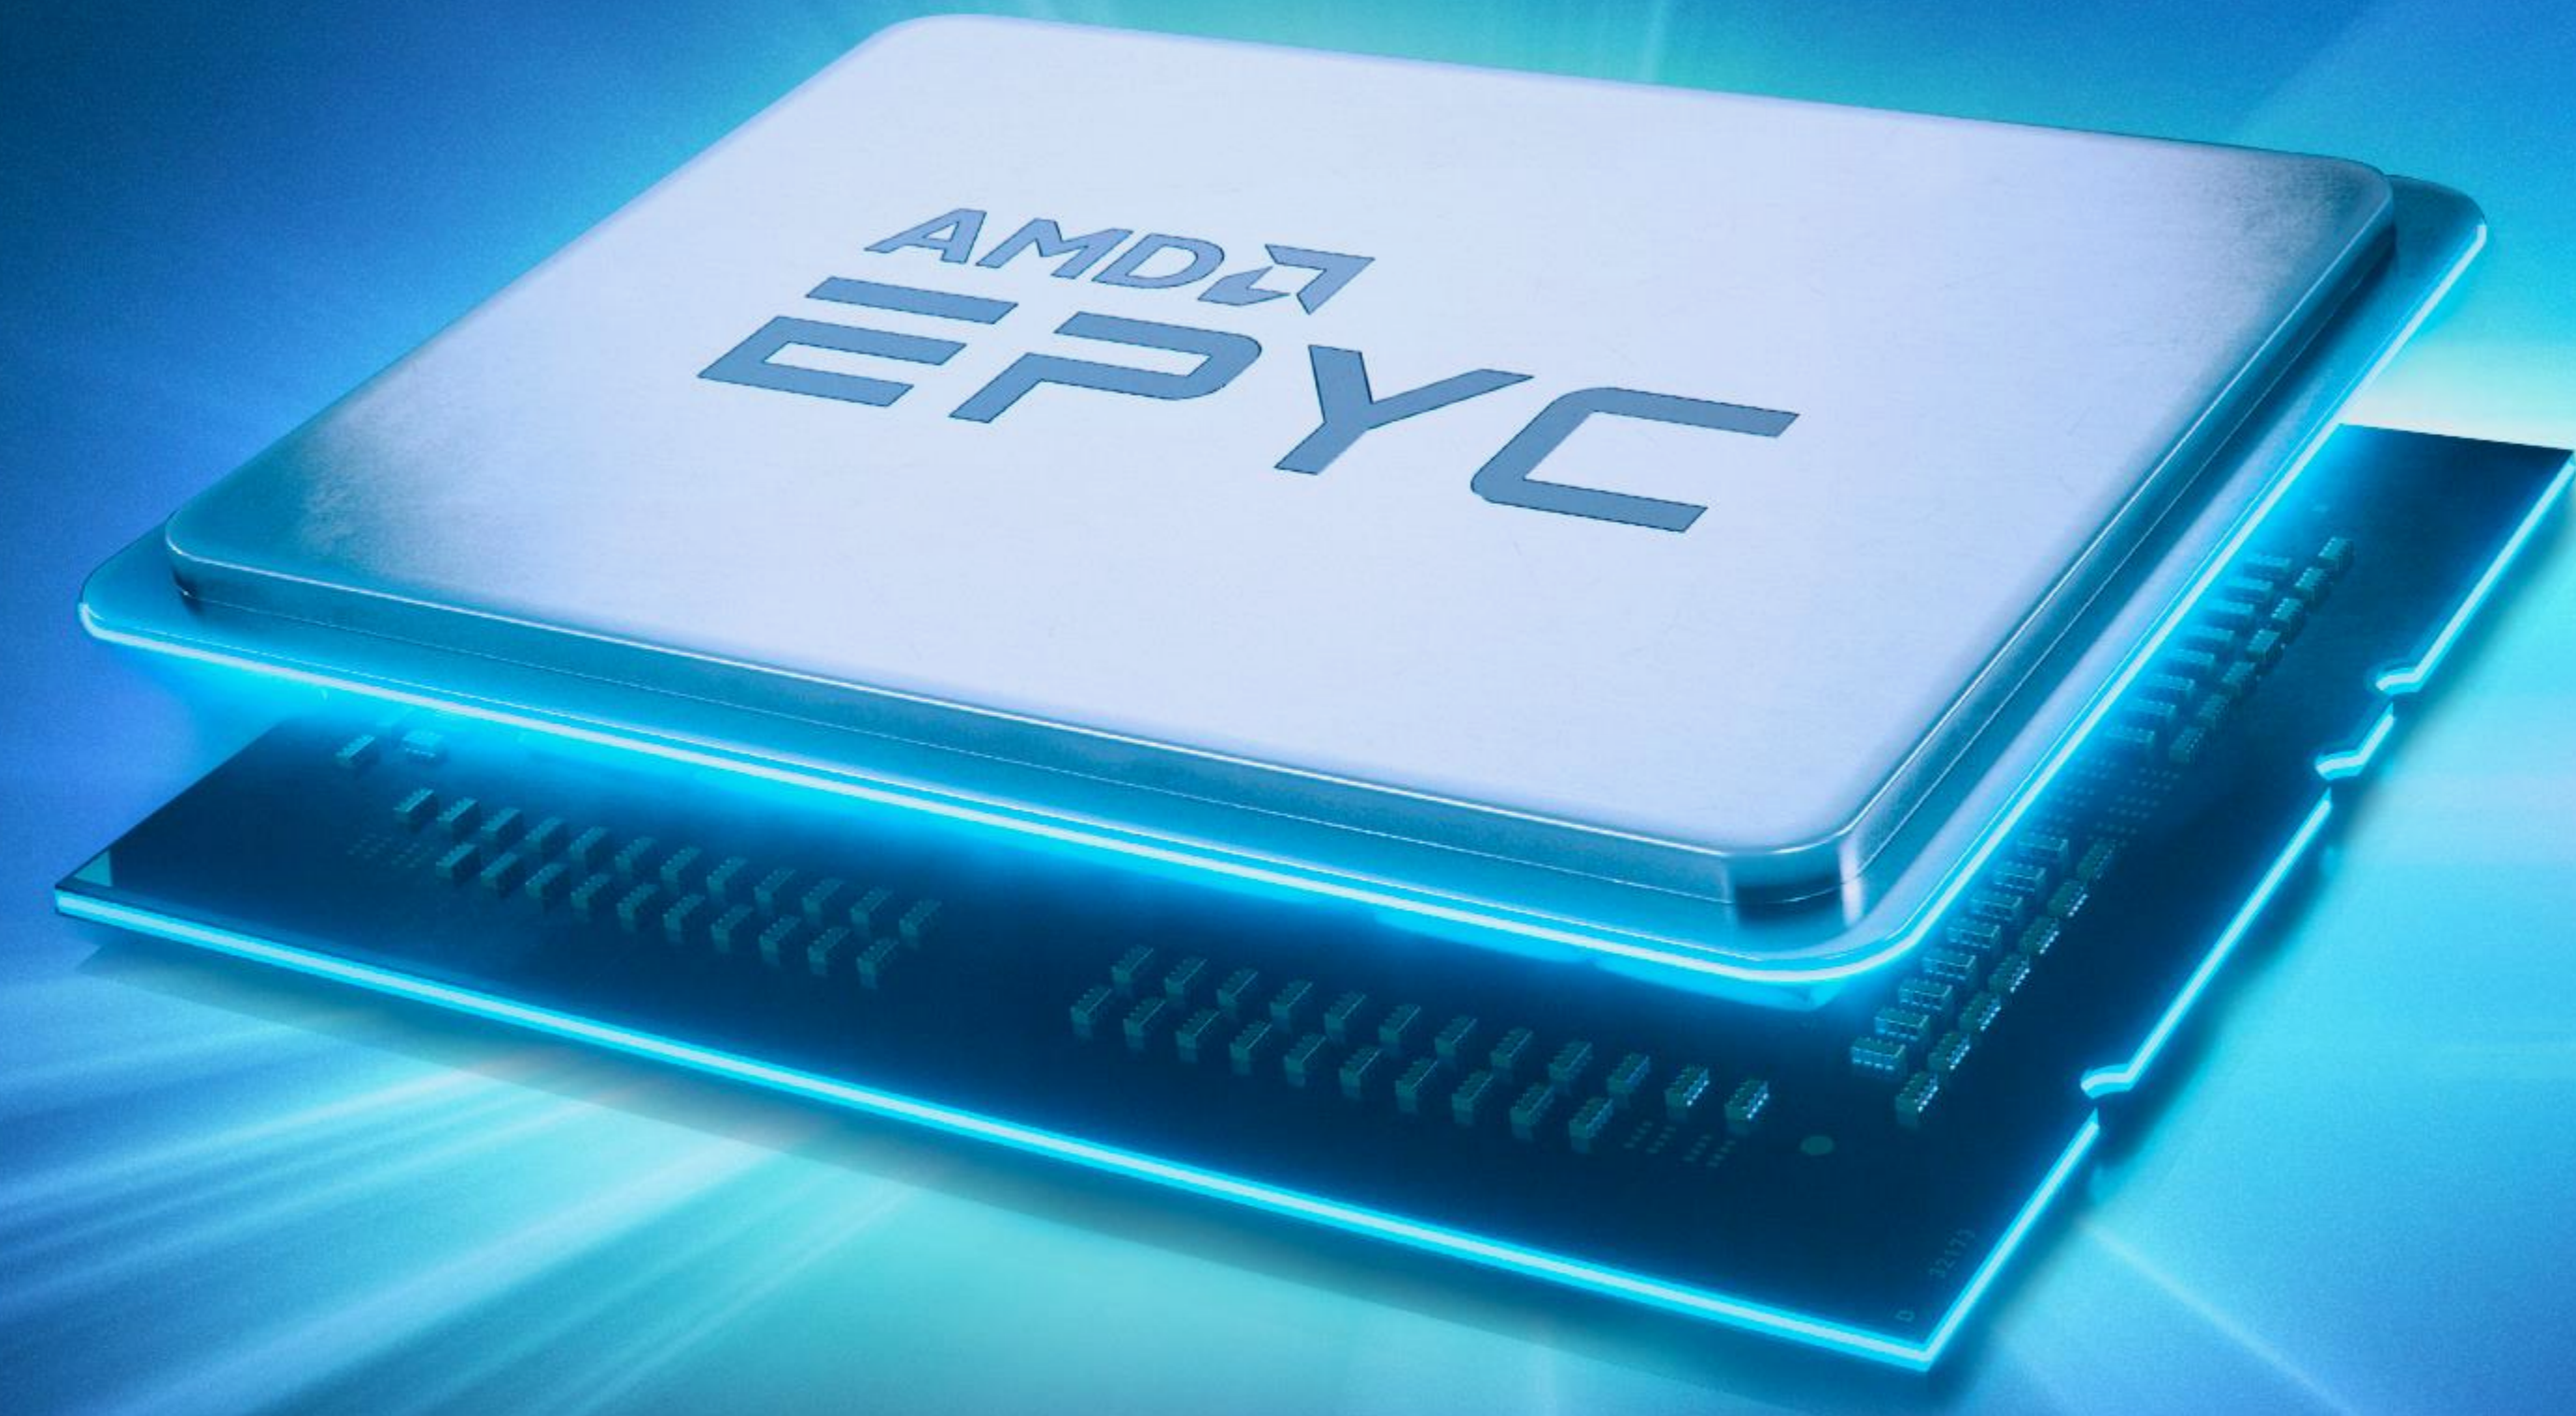
# AMD

## NEXT HORIZON

DR. LISA SU  
President and CEO

AMD EPYC™ “ROME” PREVIEW

# THE WORLD'S FIRST 7nm DATACENTER CPU



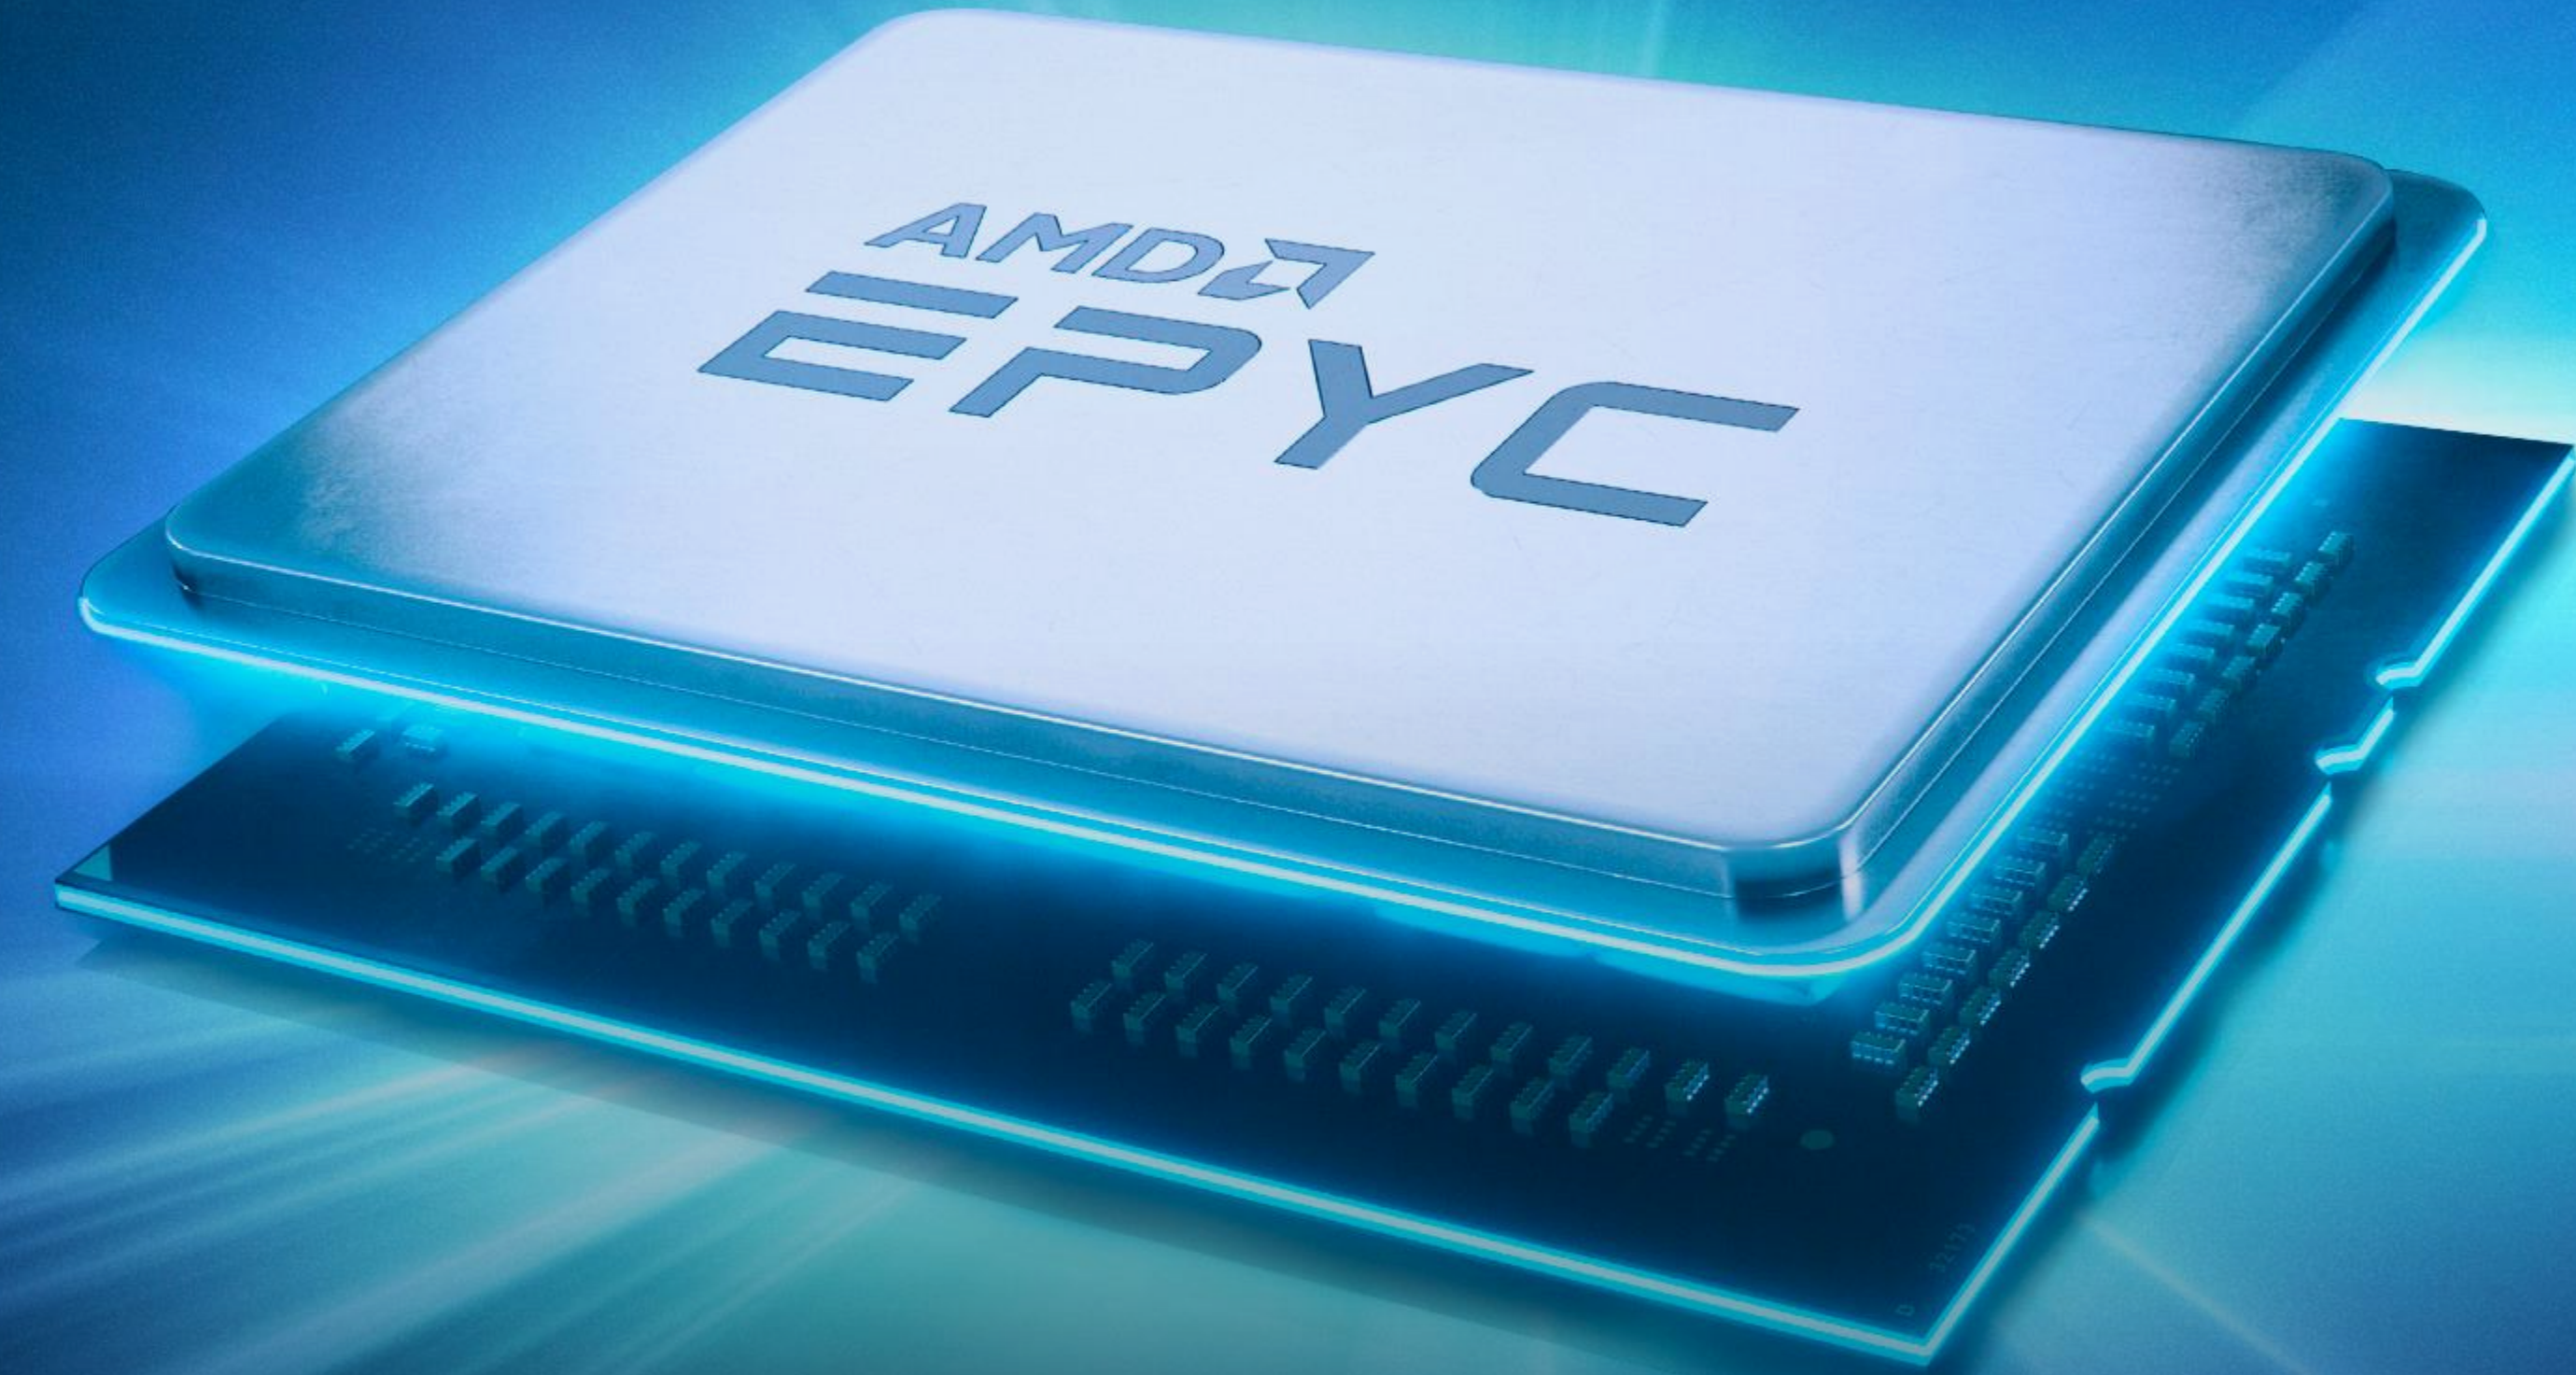
AMD EPYC™ “ROME” PREVIEW

# THE WORLD'S FIRST 7nm DATACENTER CPU

Up to 64  
“Zen 2” Cores

Increased  
Instructions-Per-Cycle

Leadership Compute, I/O  
and Memory Bandwidth



AMD EPYC™ “ROME” PREVIEW

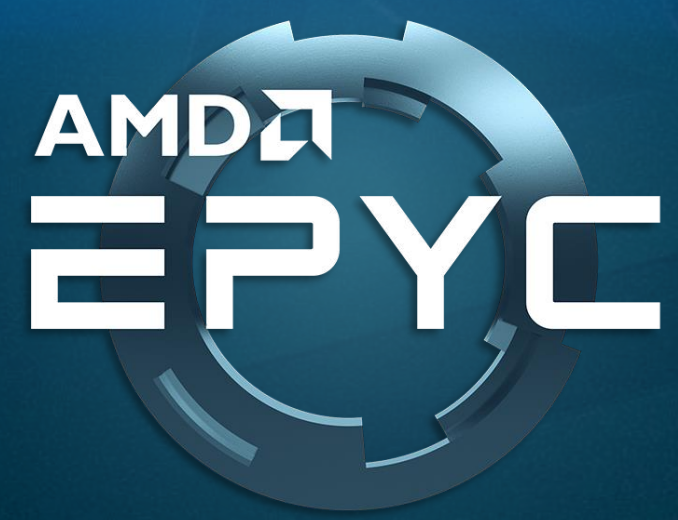
# LEADING THE ECOSYSTEM FORWARD

---

First PCIe® 4.0 Capable  
x86 Server CPU

Doubles the  
Bandwidth/Channel

Dramatically Improves  
Accelerator Performance



# THE ROAD TO “ROME” IS THROUGH “NAPLES”

Socket Compatible  
with Existing  
“Naples” Platforms

Forward Compatible  
with Next-Generation  
“Milan” Platforms

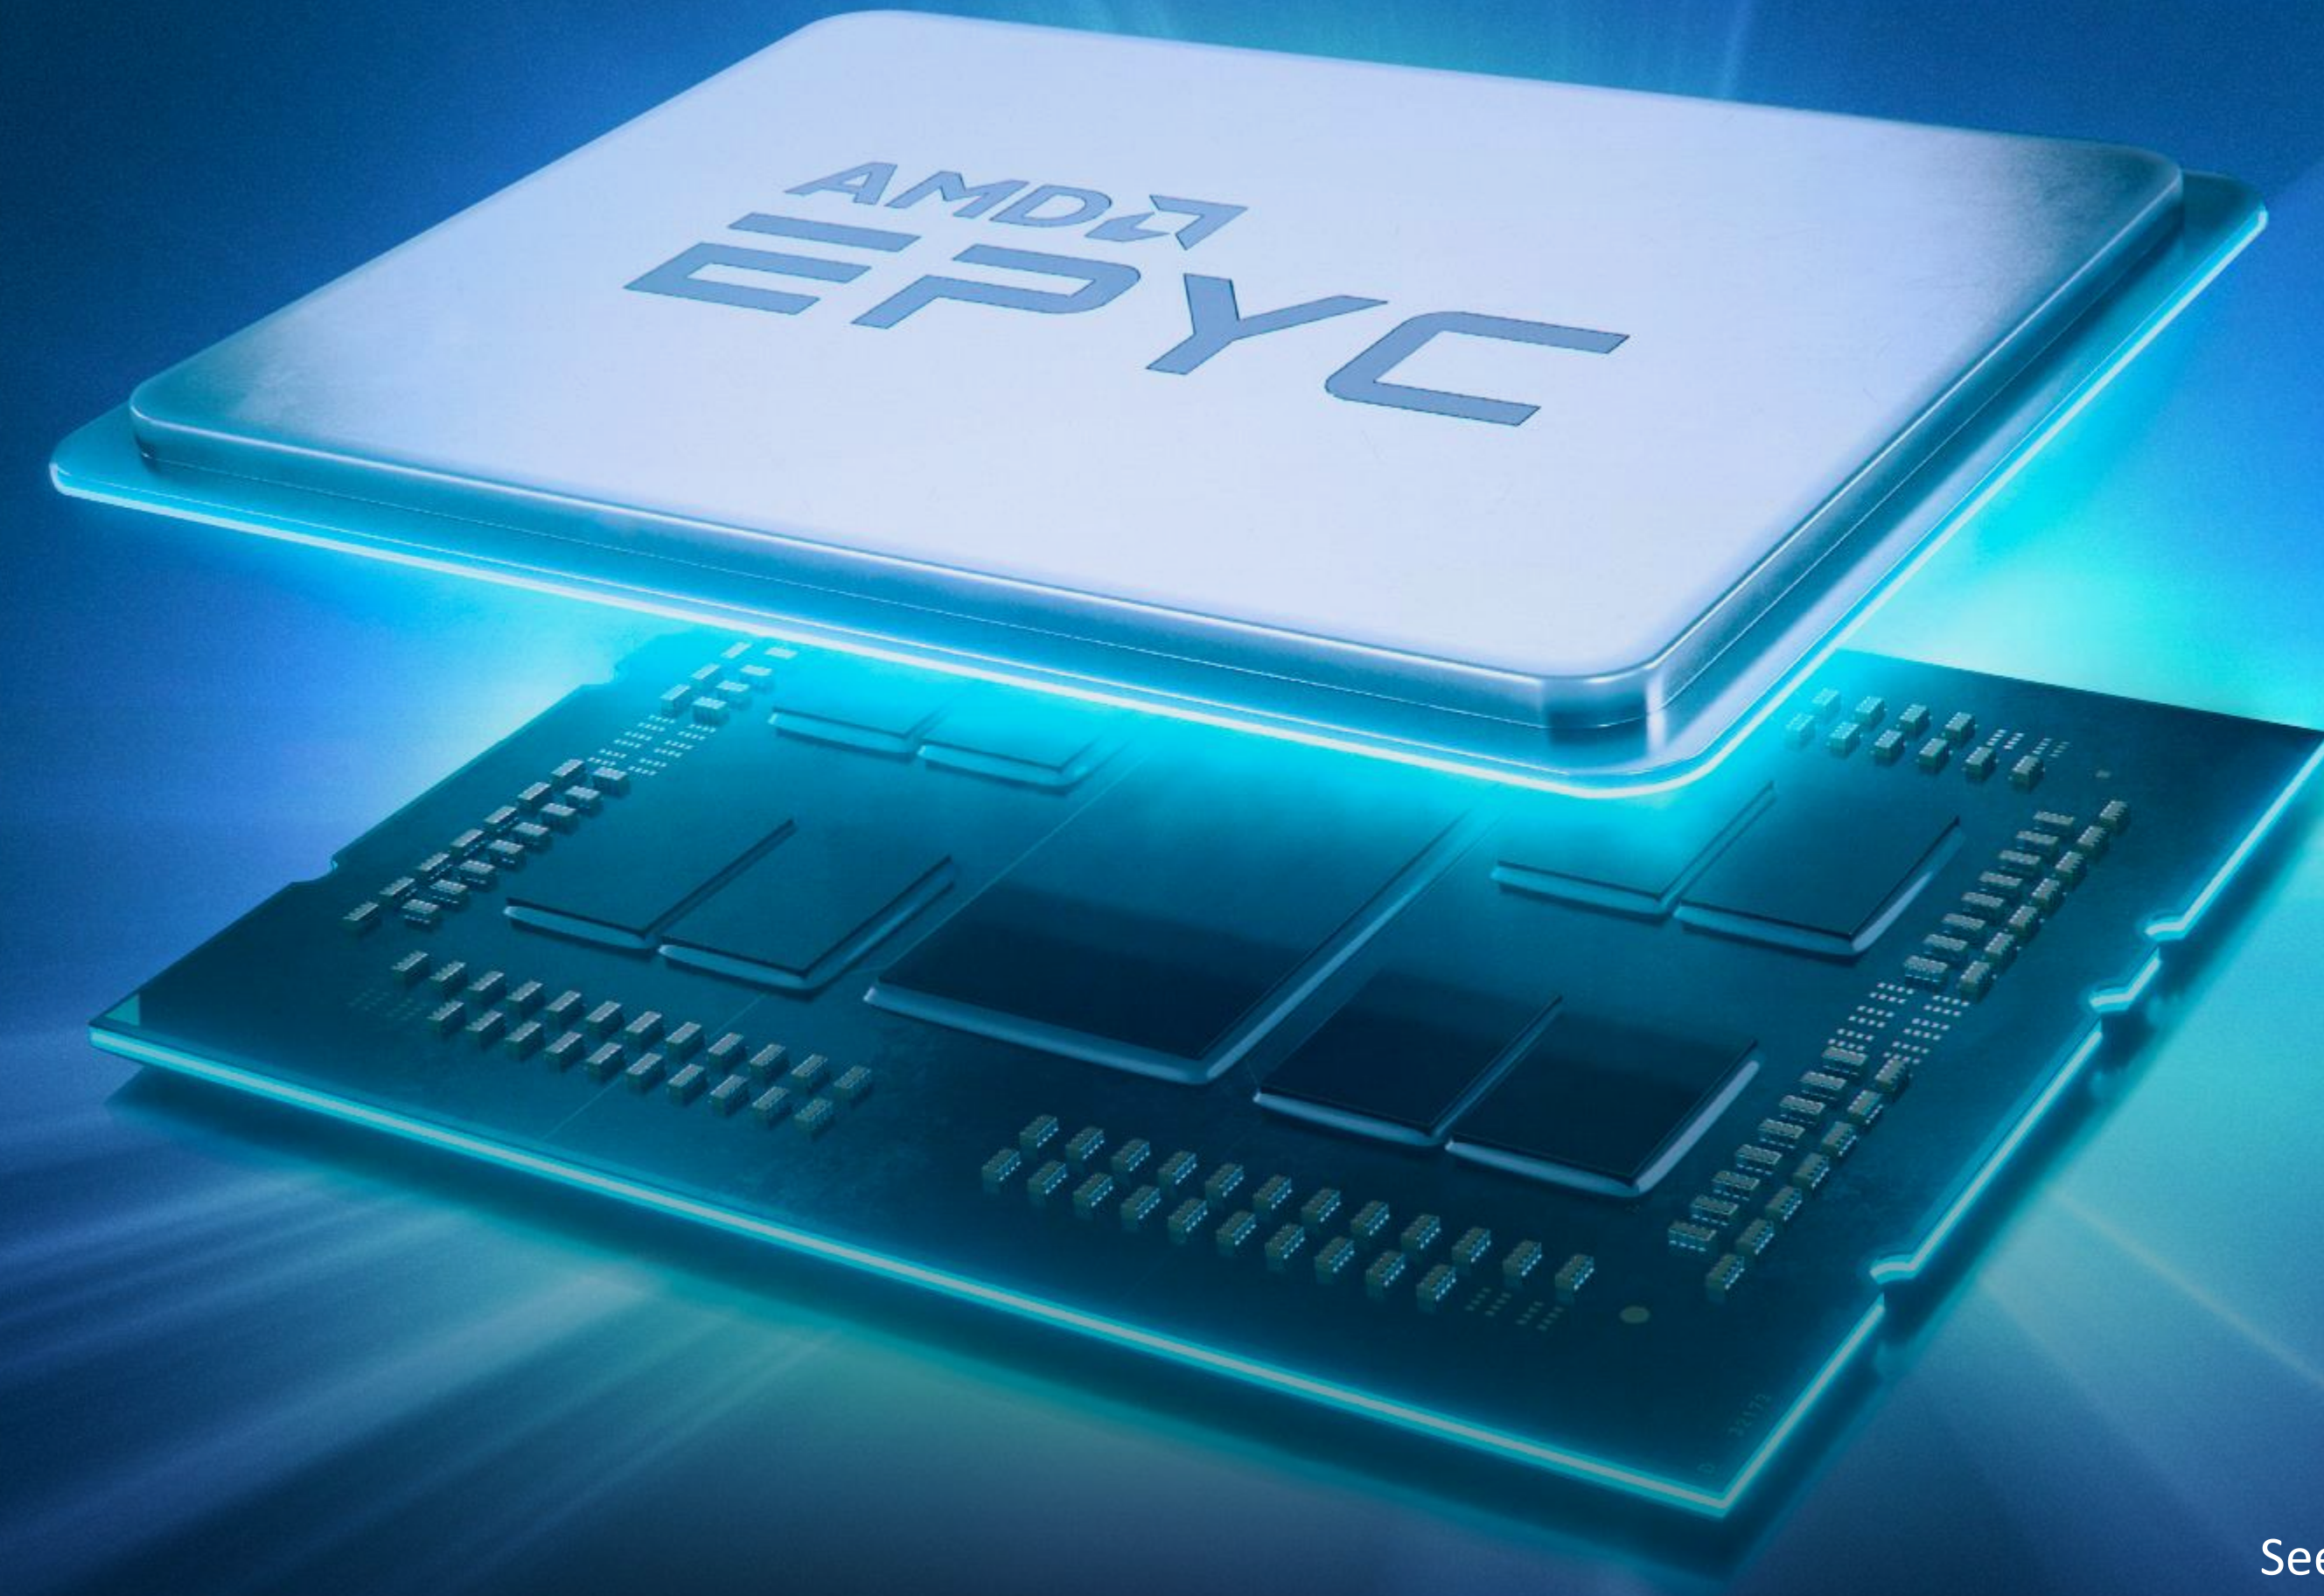
AMD EPYC™ “ROME” PREVIEW  
THE WORLD’S FIRST  
7nm DATACENTER CPU

2x

Performance Per Socket  
*vs. Previous Generation*

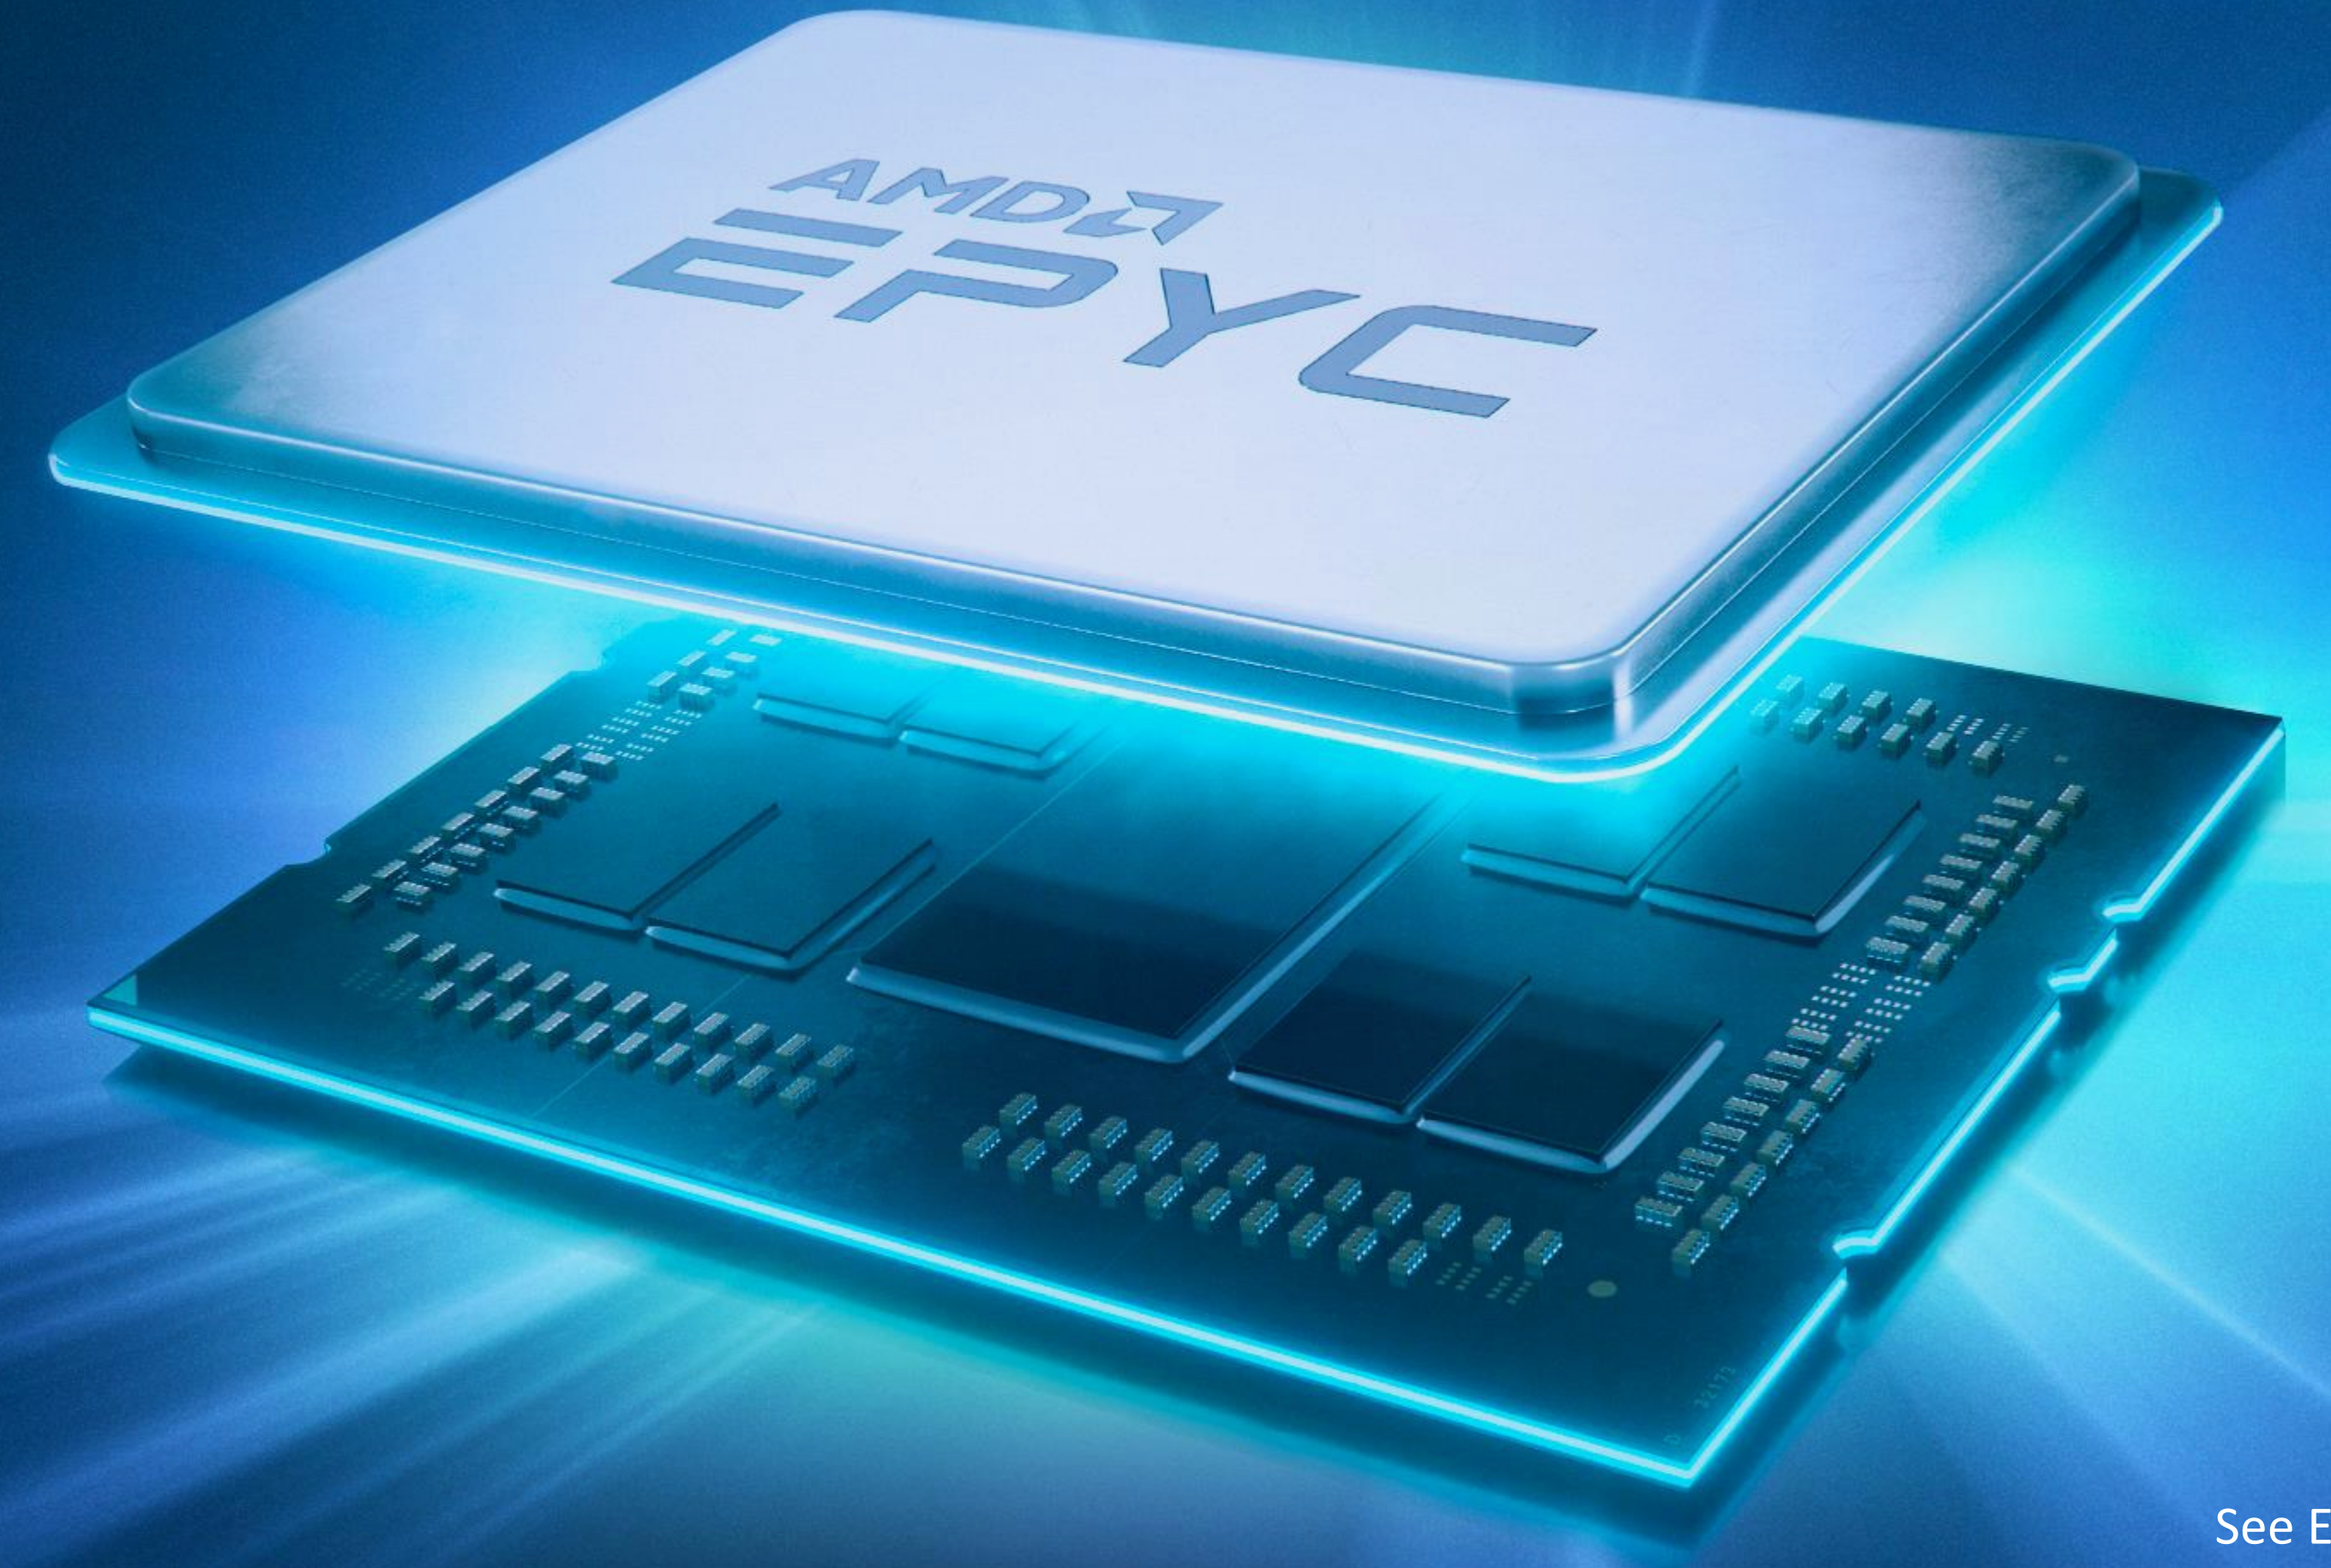
4x

Floating Point Per Socket  
*vs. Previous Generation*



AMD EPYC™ “ROME” PREVIEW

# THE WORLD'S FIRST 7nm DATACENTER CPU



WORLD'S FIRST 7nm DATACENTER CPU AND GPU DEMO

---

# AMD EPYC™ “ROME” WITH RADEON INSTINCT™ MI60

“Live Demo”

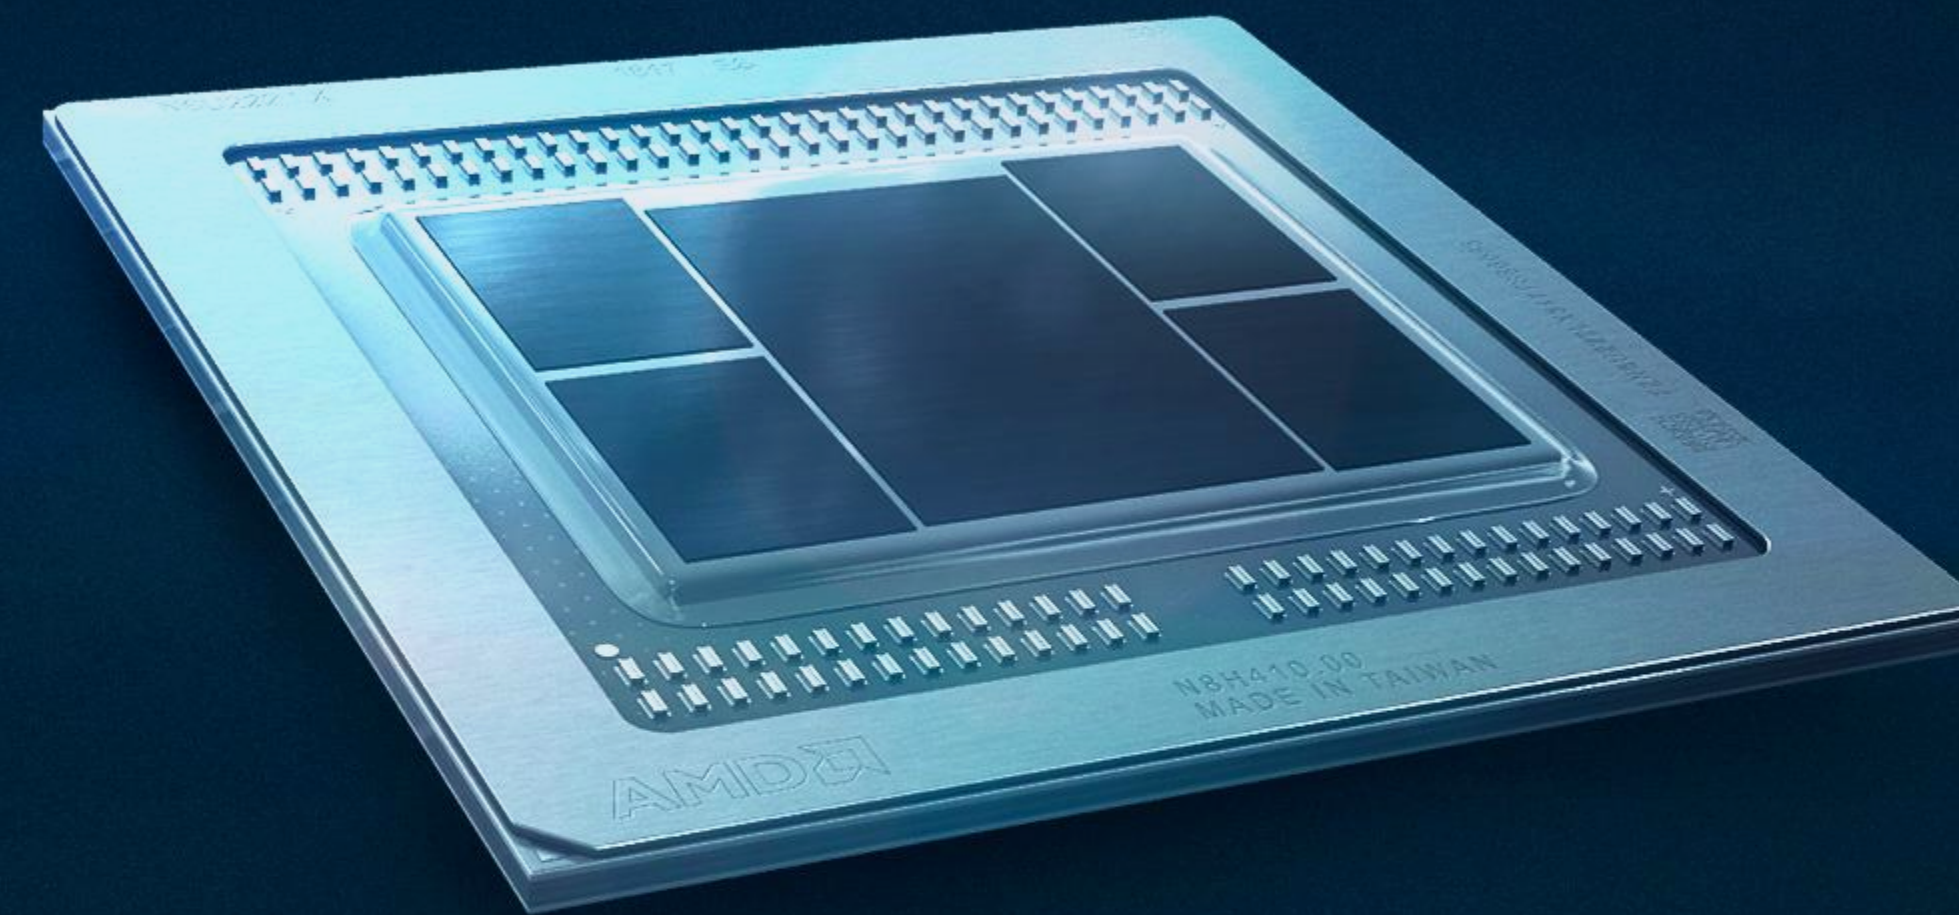
AMD EPYC™ "Naples" and  
AMD Radeon Instinct™ MI25

# WORLD'S FIRST 7NM DATACENTER CPU AND GPU DEMO

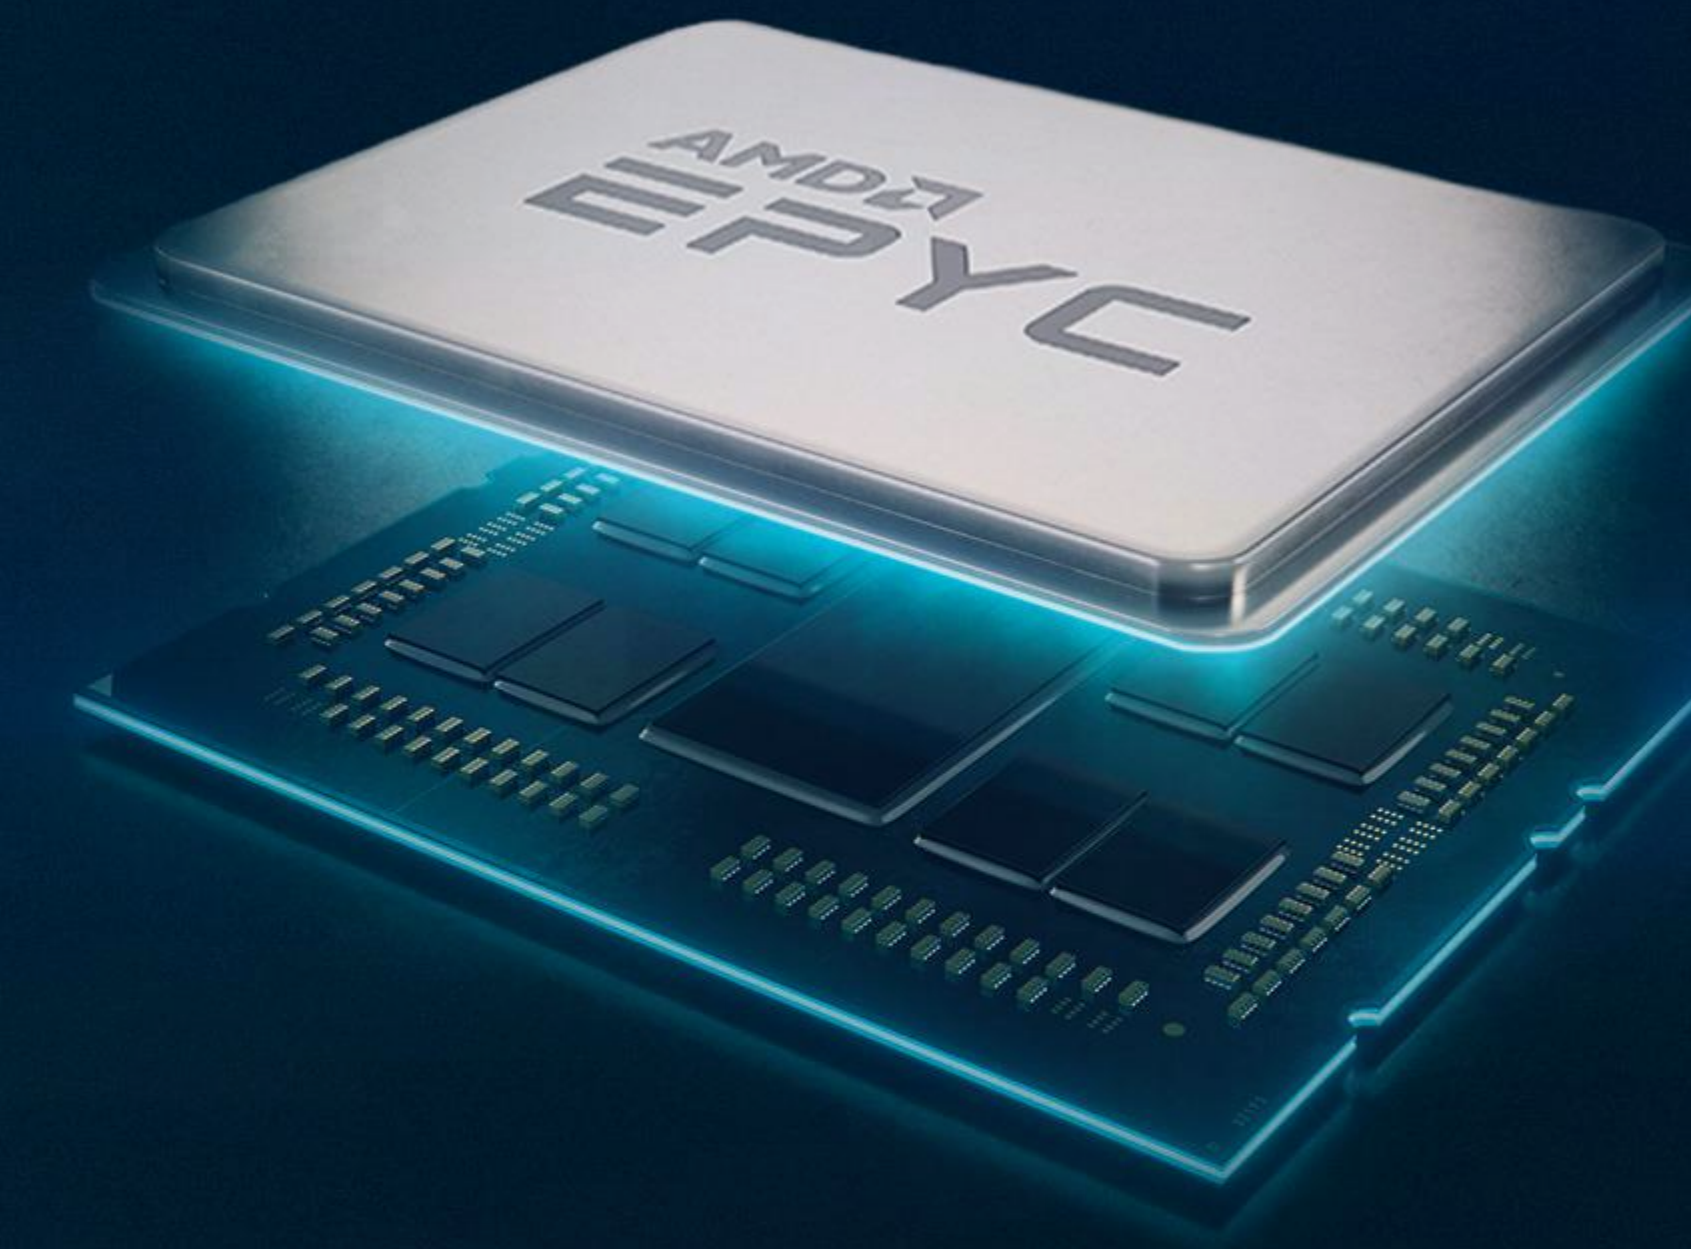
“Live Demo”

AMD EPYC™ "ROME" and  
AMD Radeon Instinct™ MI60

# ONLY AMD DESIGNS DATACENTER CPU AND GPU ARCHITECTURES TOGETHER



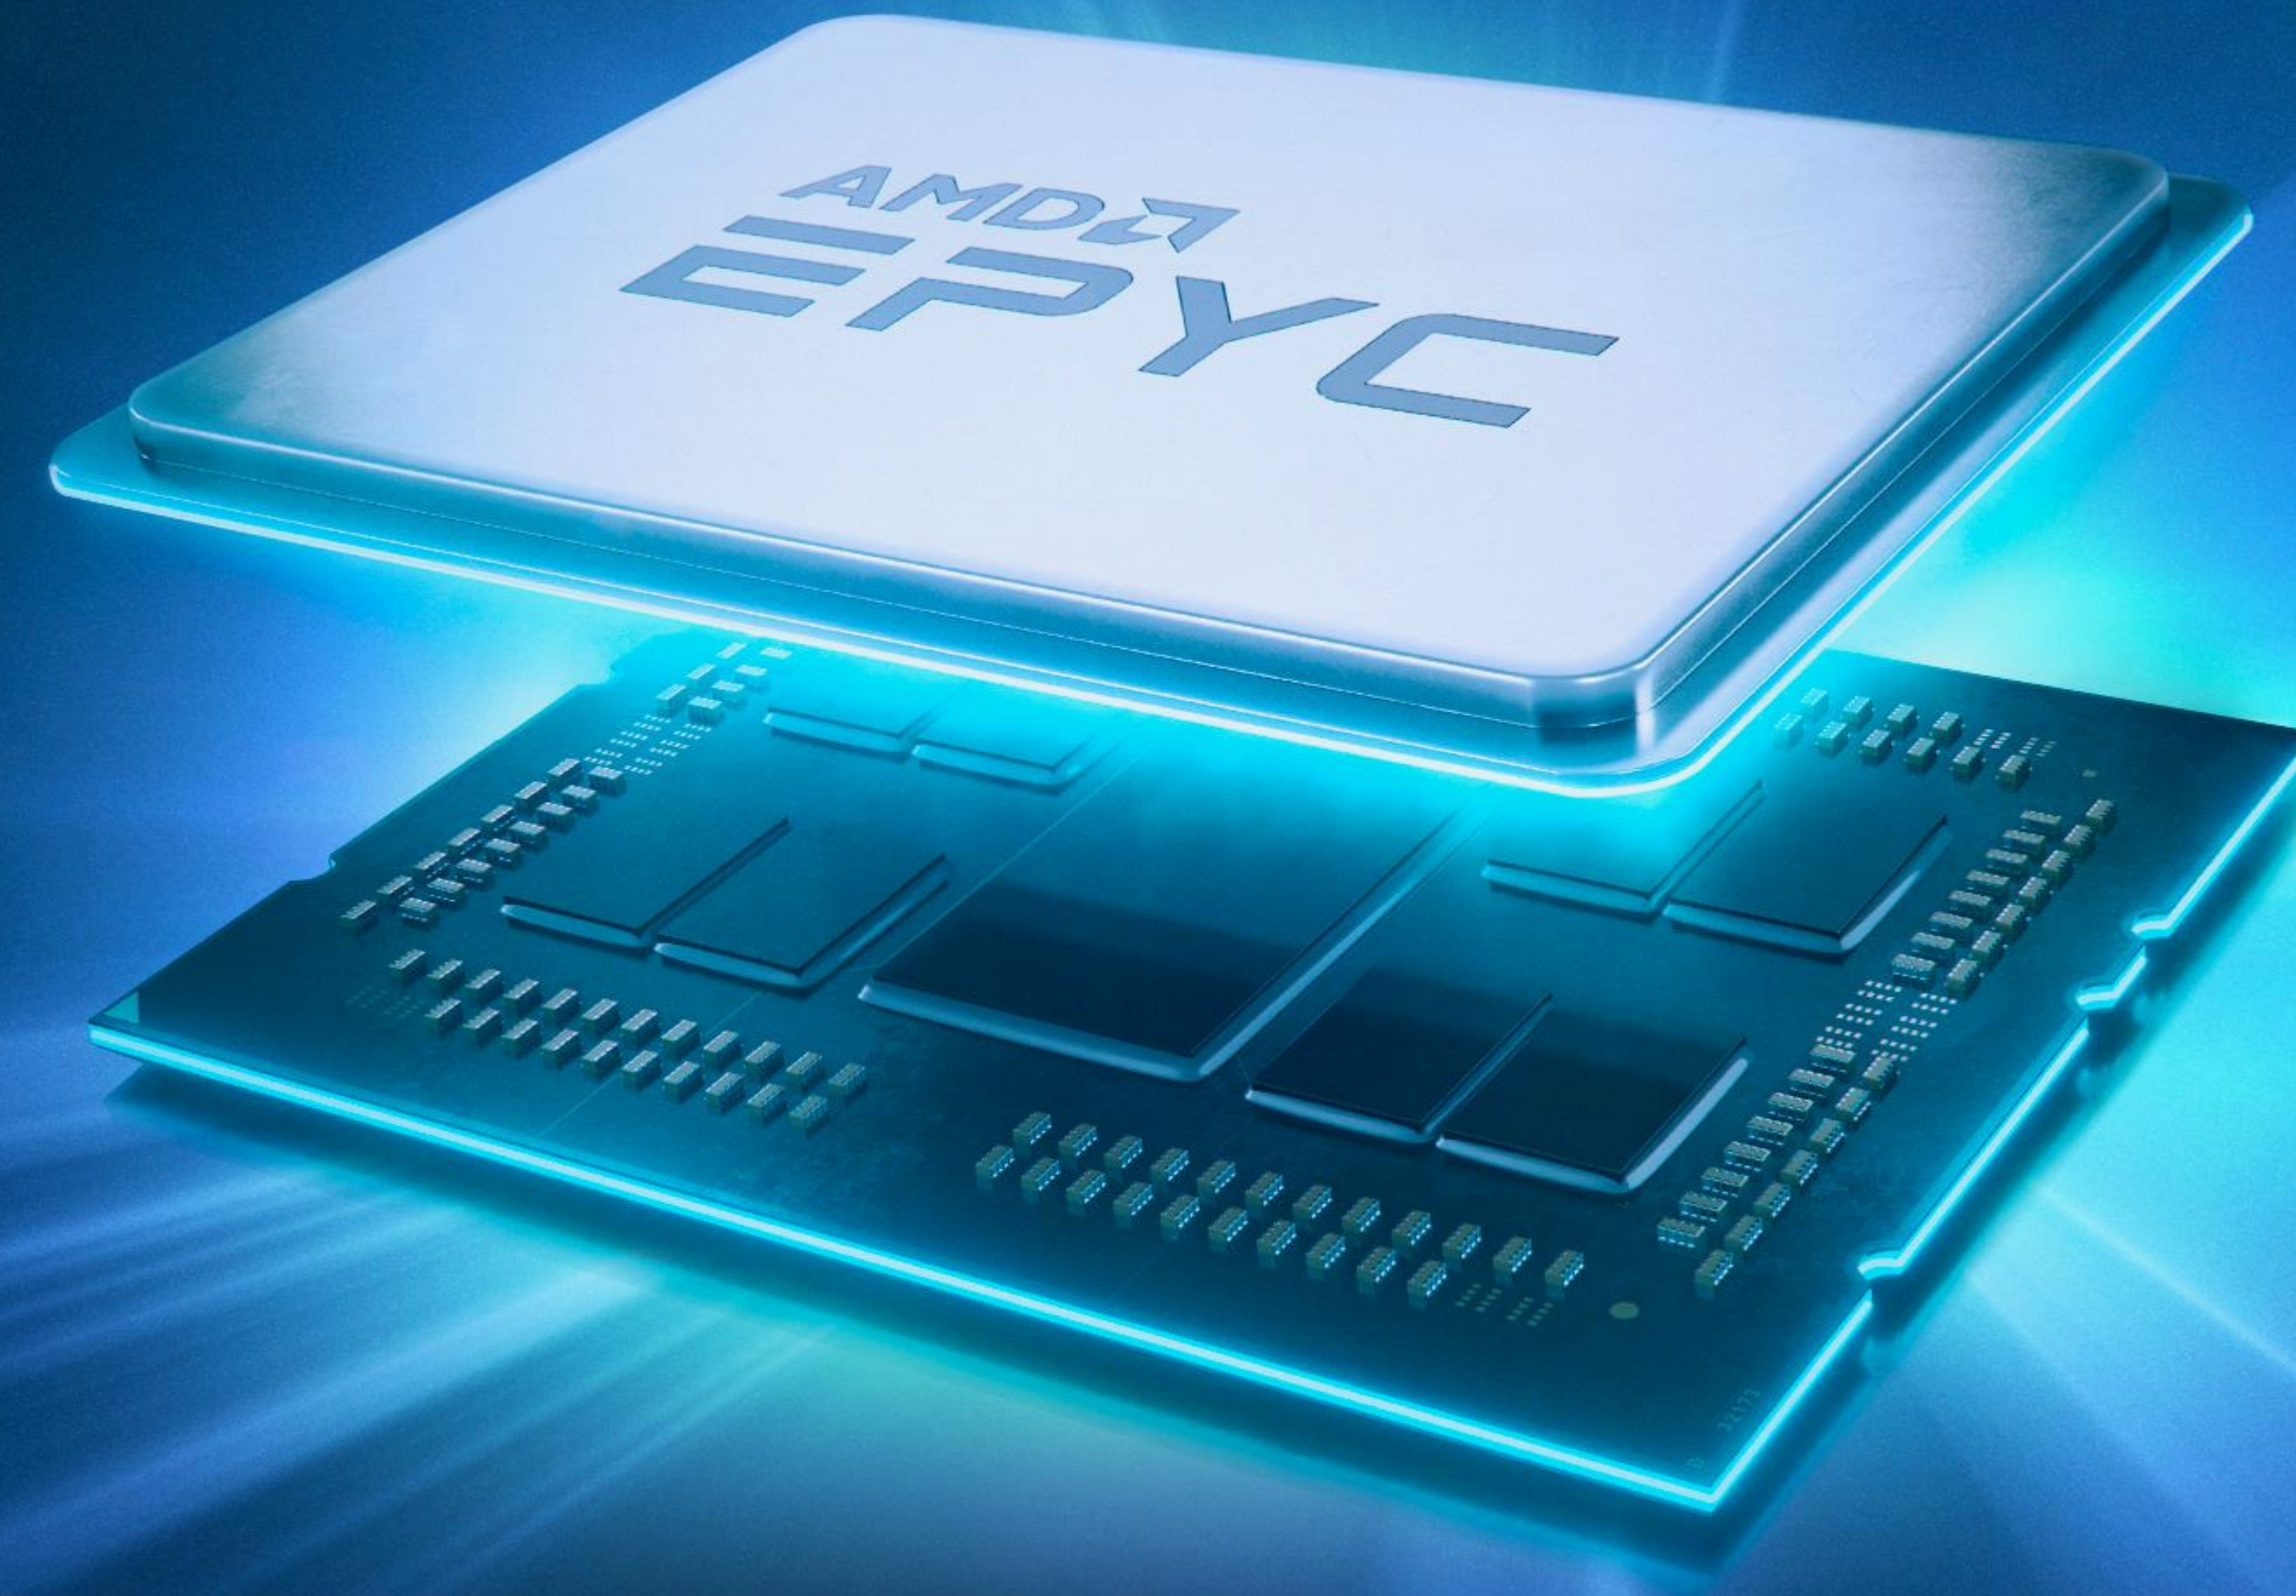
World's First 7nm GPU  
AMD RADEON INSTINCT MI60

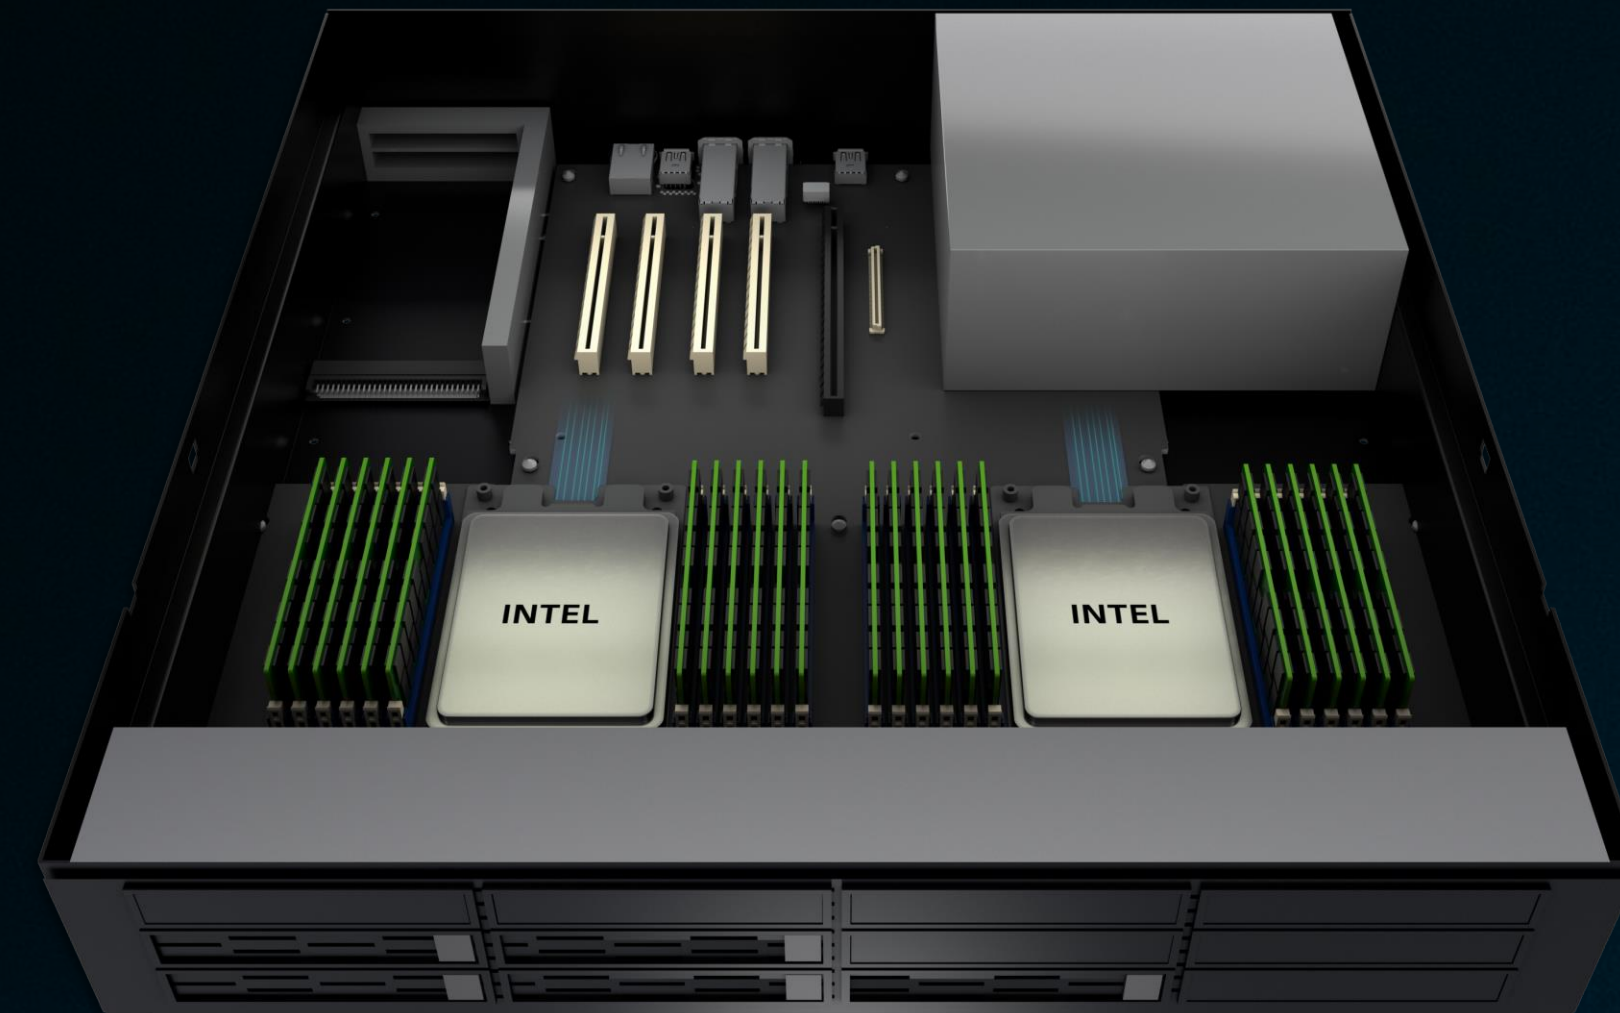


World's First 7nm Datacenter CPU  
"Rome"

WORLD'S FIRST COMPETITIVE DEMO WITH "ROME"

# AMD EPYC™ "ROME" VS. BEST- IN-CLASS "SKYLAKE"





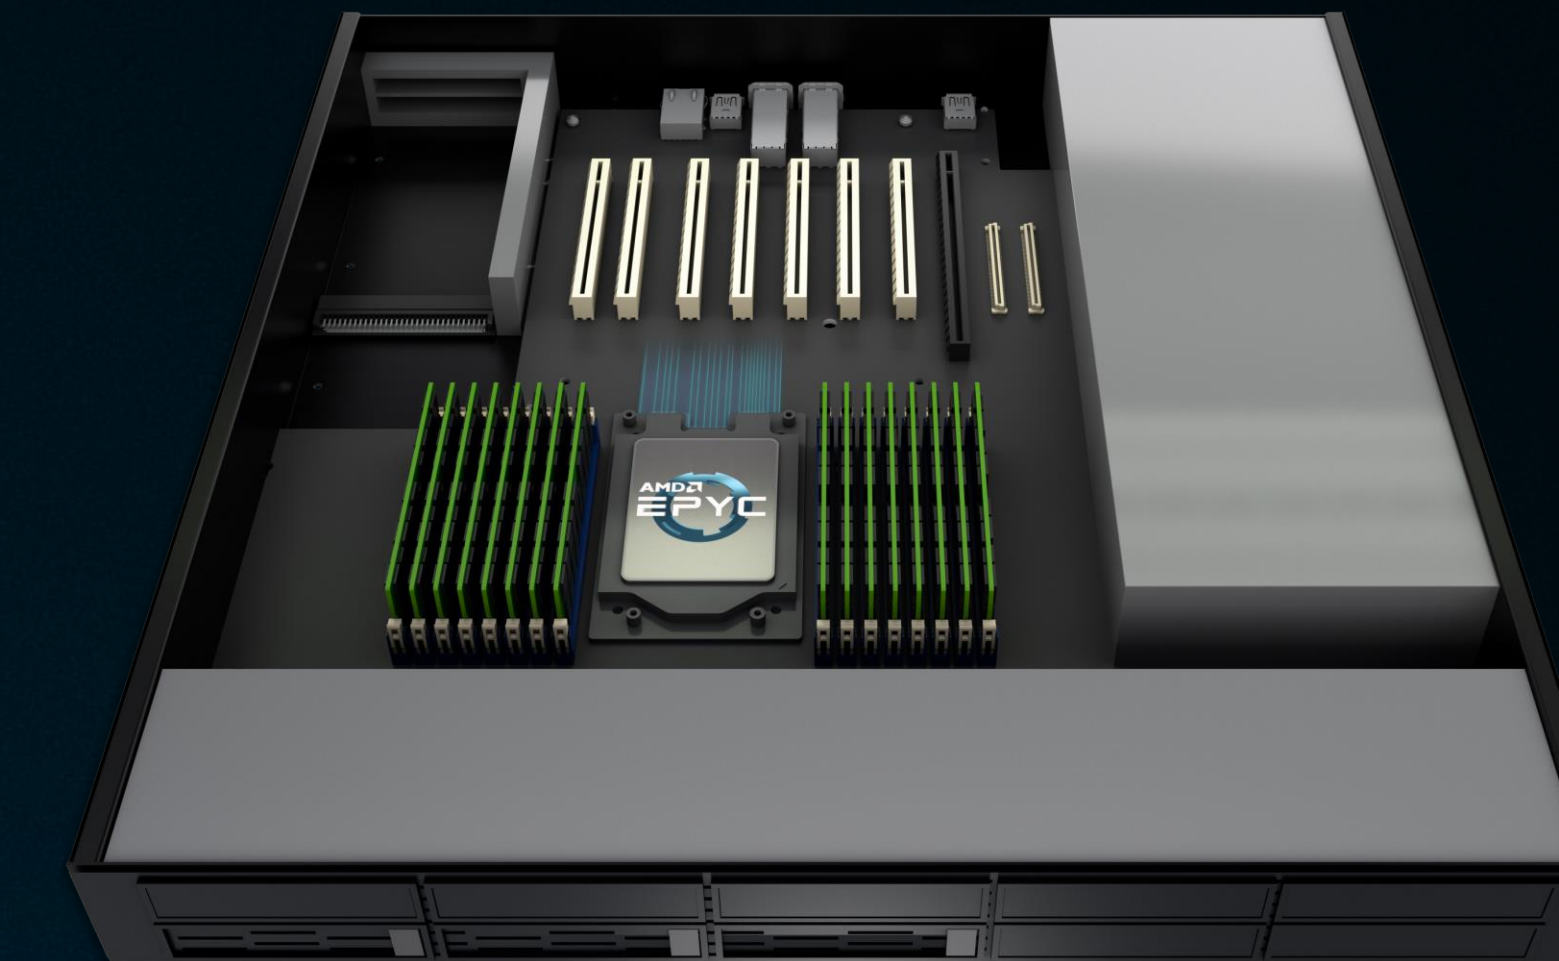
DUAL INTEL 8180M  
PLATINUM PROCESSORS

Up To 3 TB of DRAM

96 Lanes of PCIe® 3.0

# AMD EPYC™ “ROME”

HEAD-TO-HEAD PREVIEW



SINGLE SOCKET  
AMD EPYC “ROME”

Up To 4 TB of DRAM

128 Lanes of PCIe® 4.0

“LIVE DEMO”

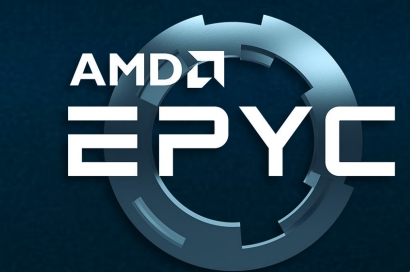
Intel Xeon Scalable 8180M  
TWO SOCKET

# AMD EPYC™ “ROME”

## HEAD-TO-HEAD PREVIEW

“C-RAY” INDUSTRY STANDARD BENCHMARK

“LIVE DEMO”



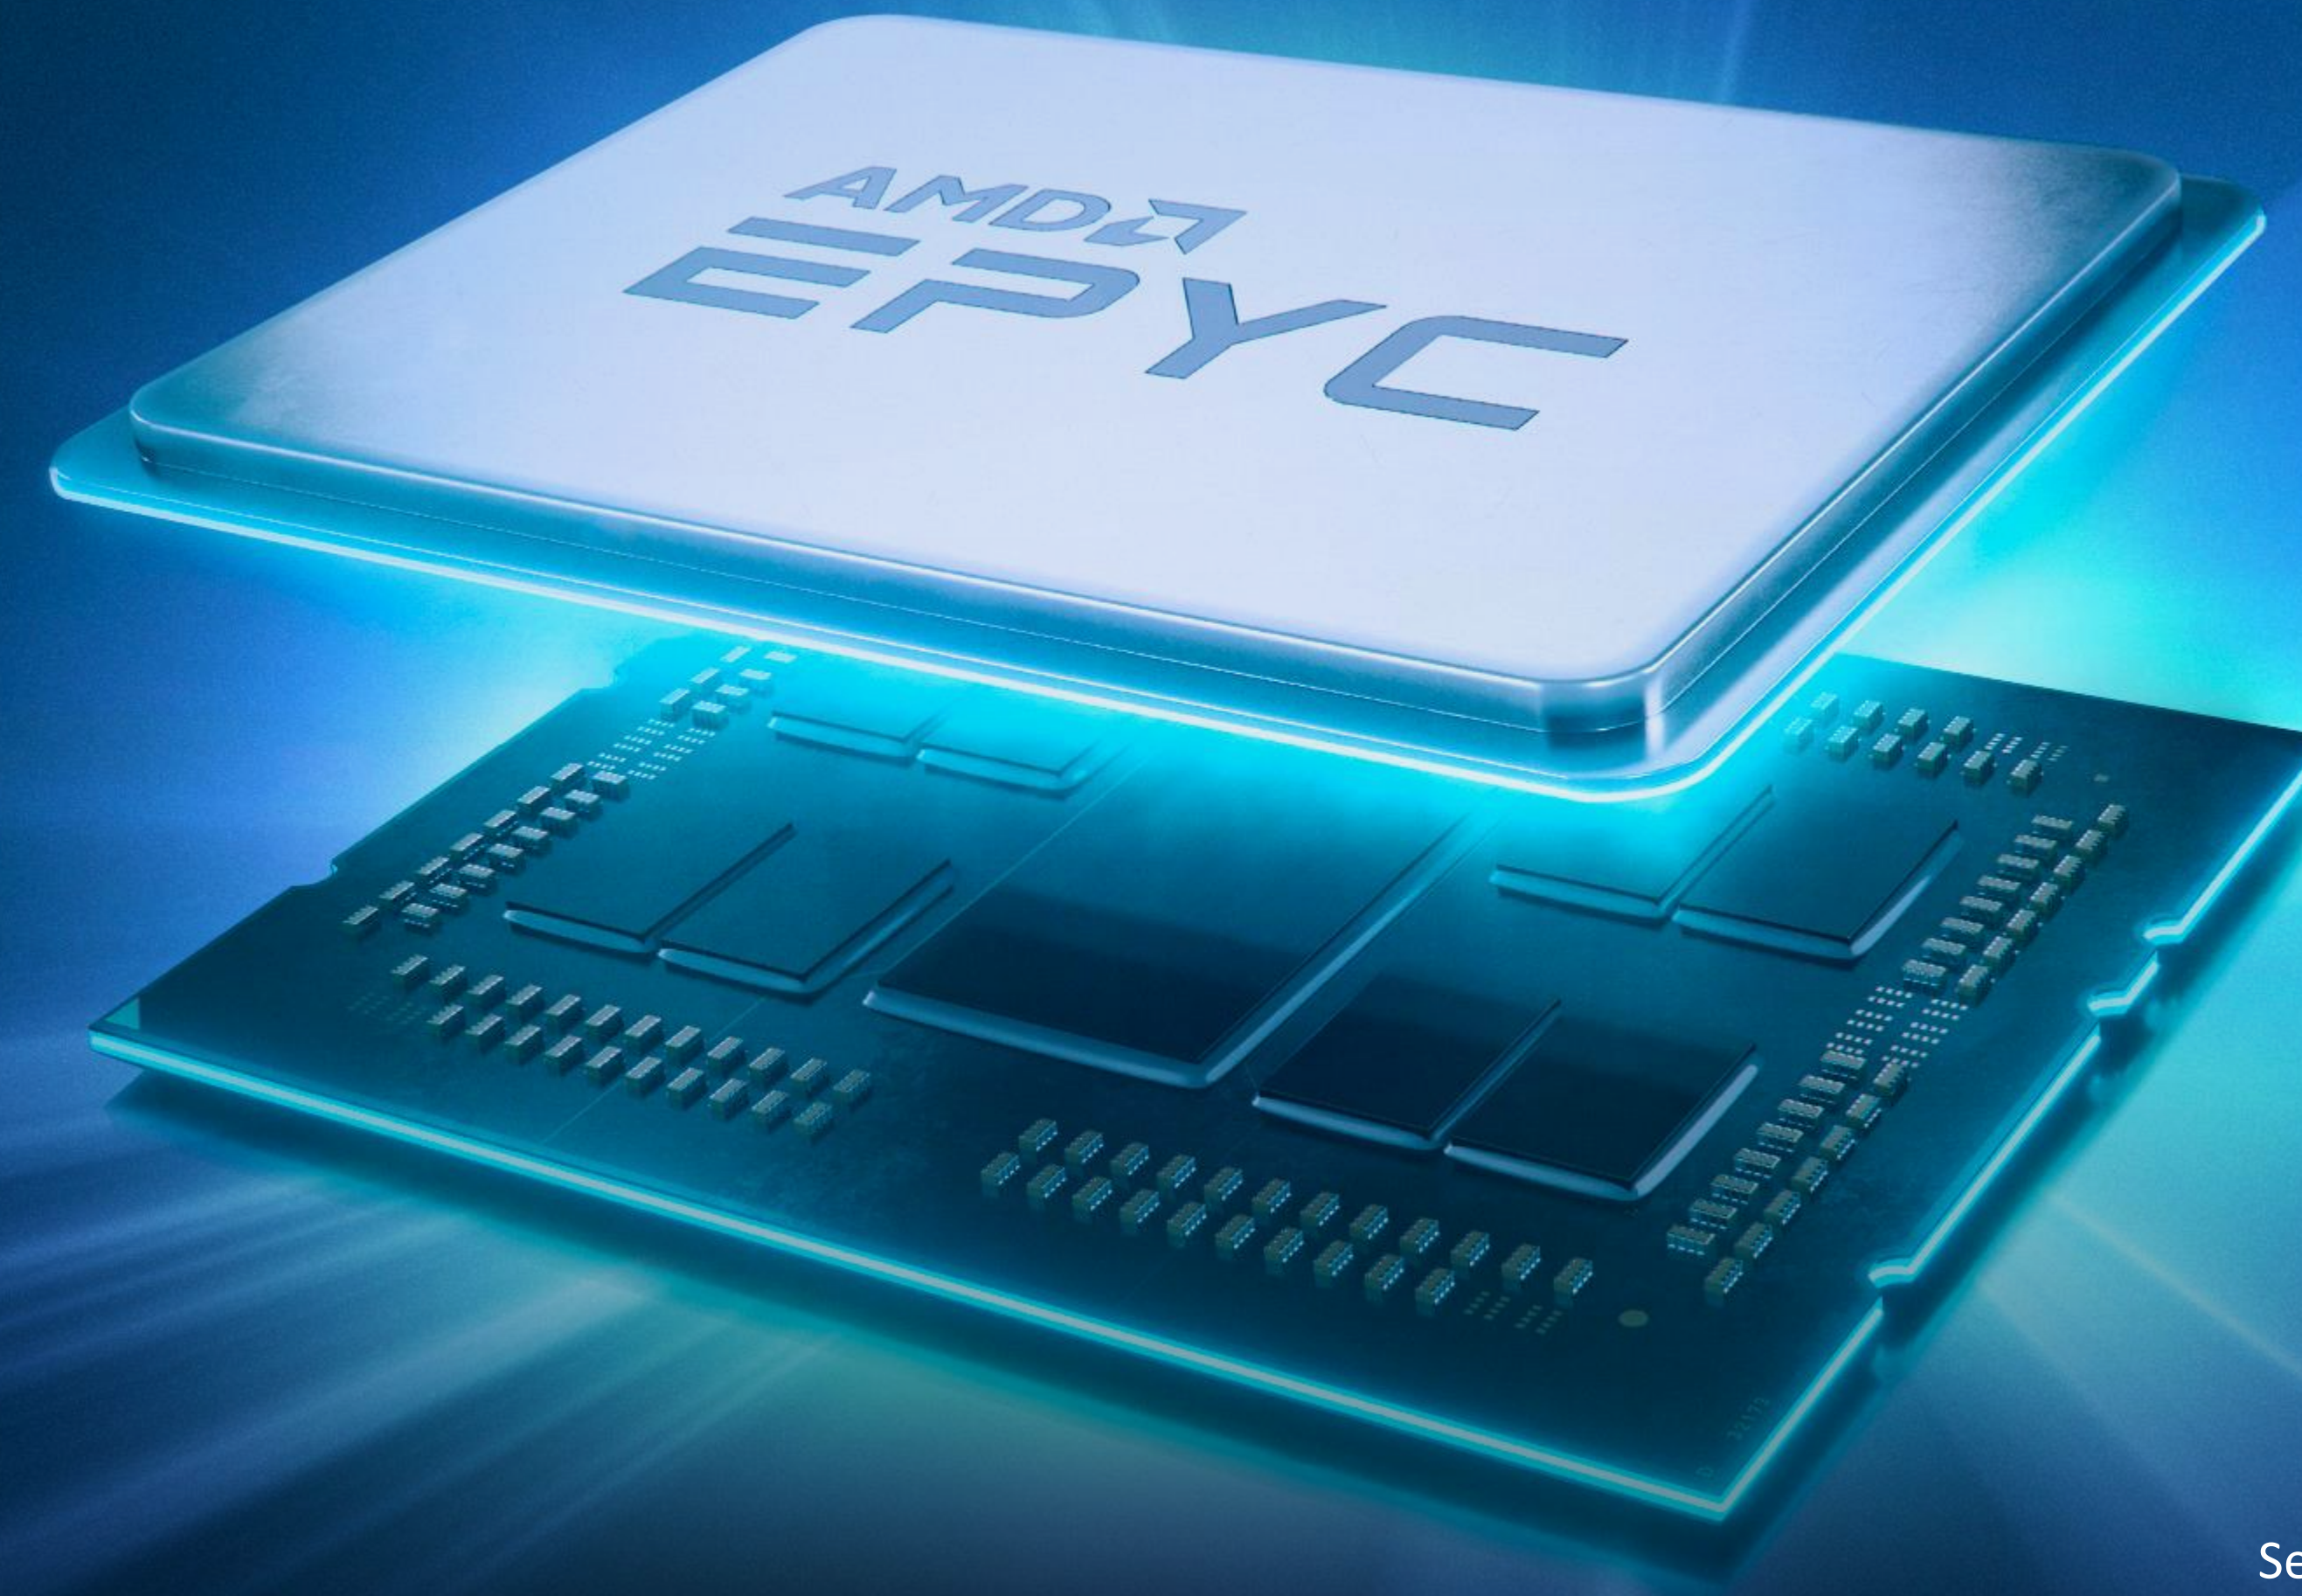
“Rome” Preview  
ONE SOCKET

AMD EPYC™ “ROME” PREVIEW

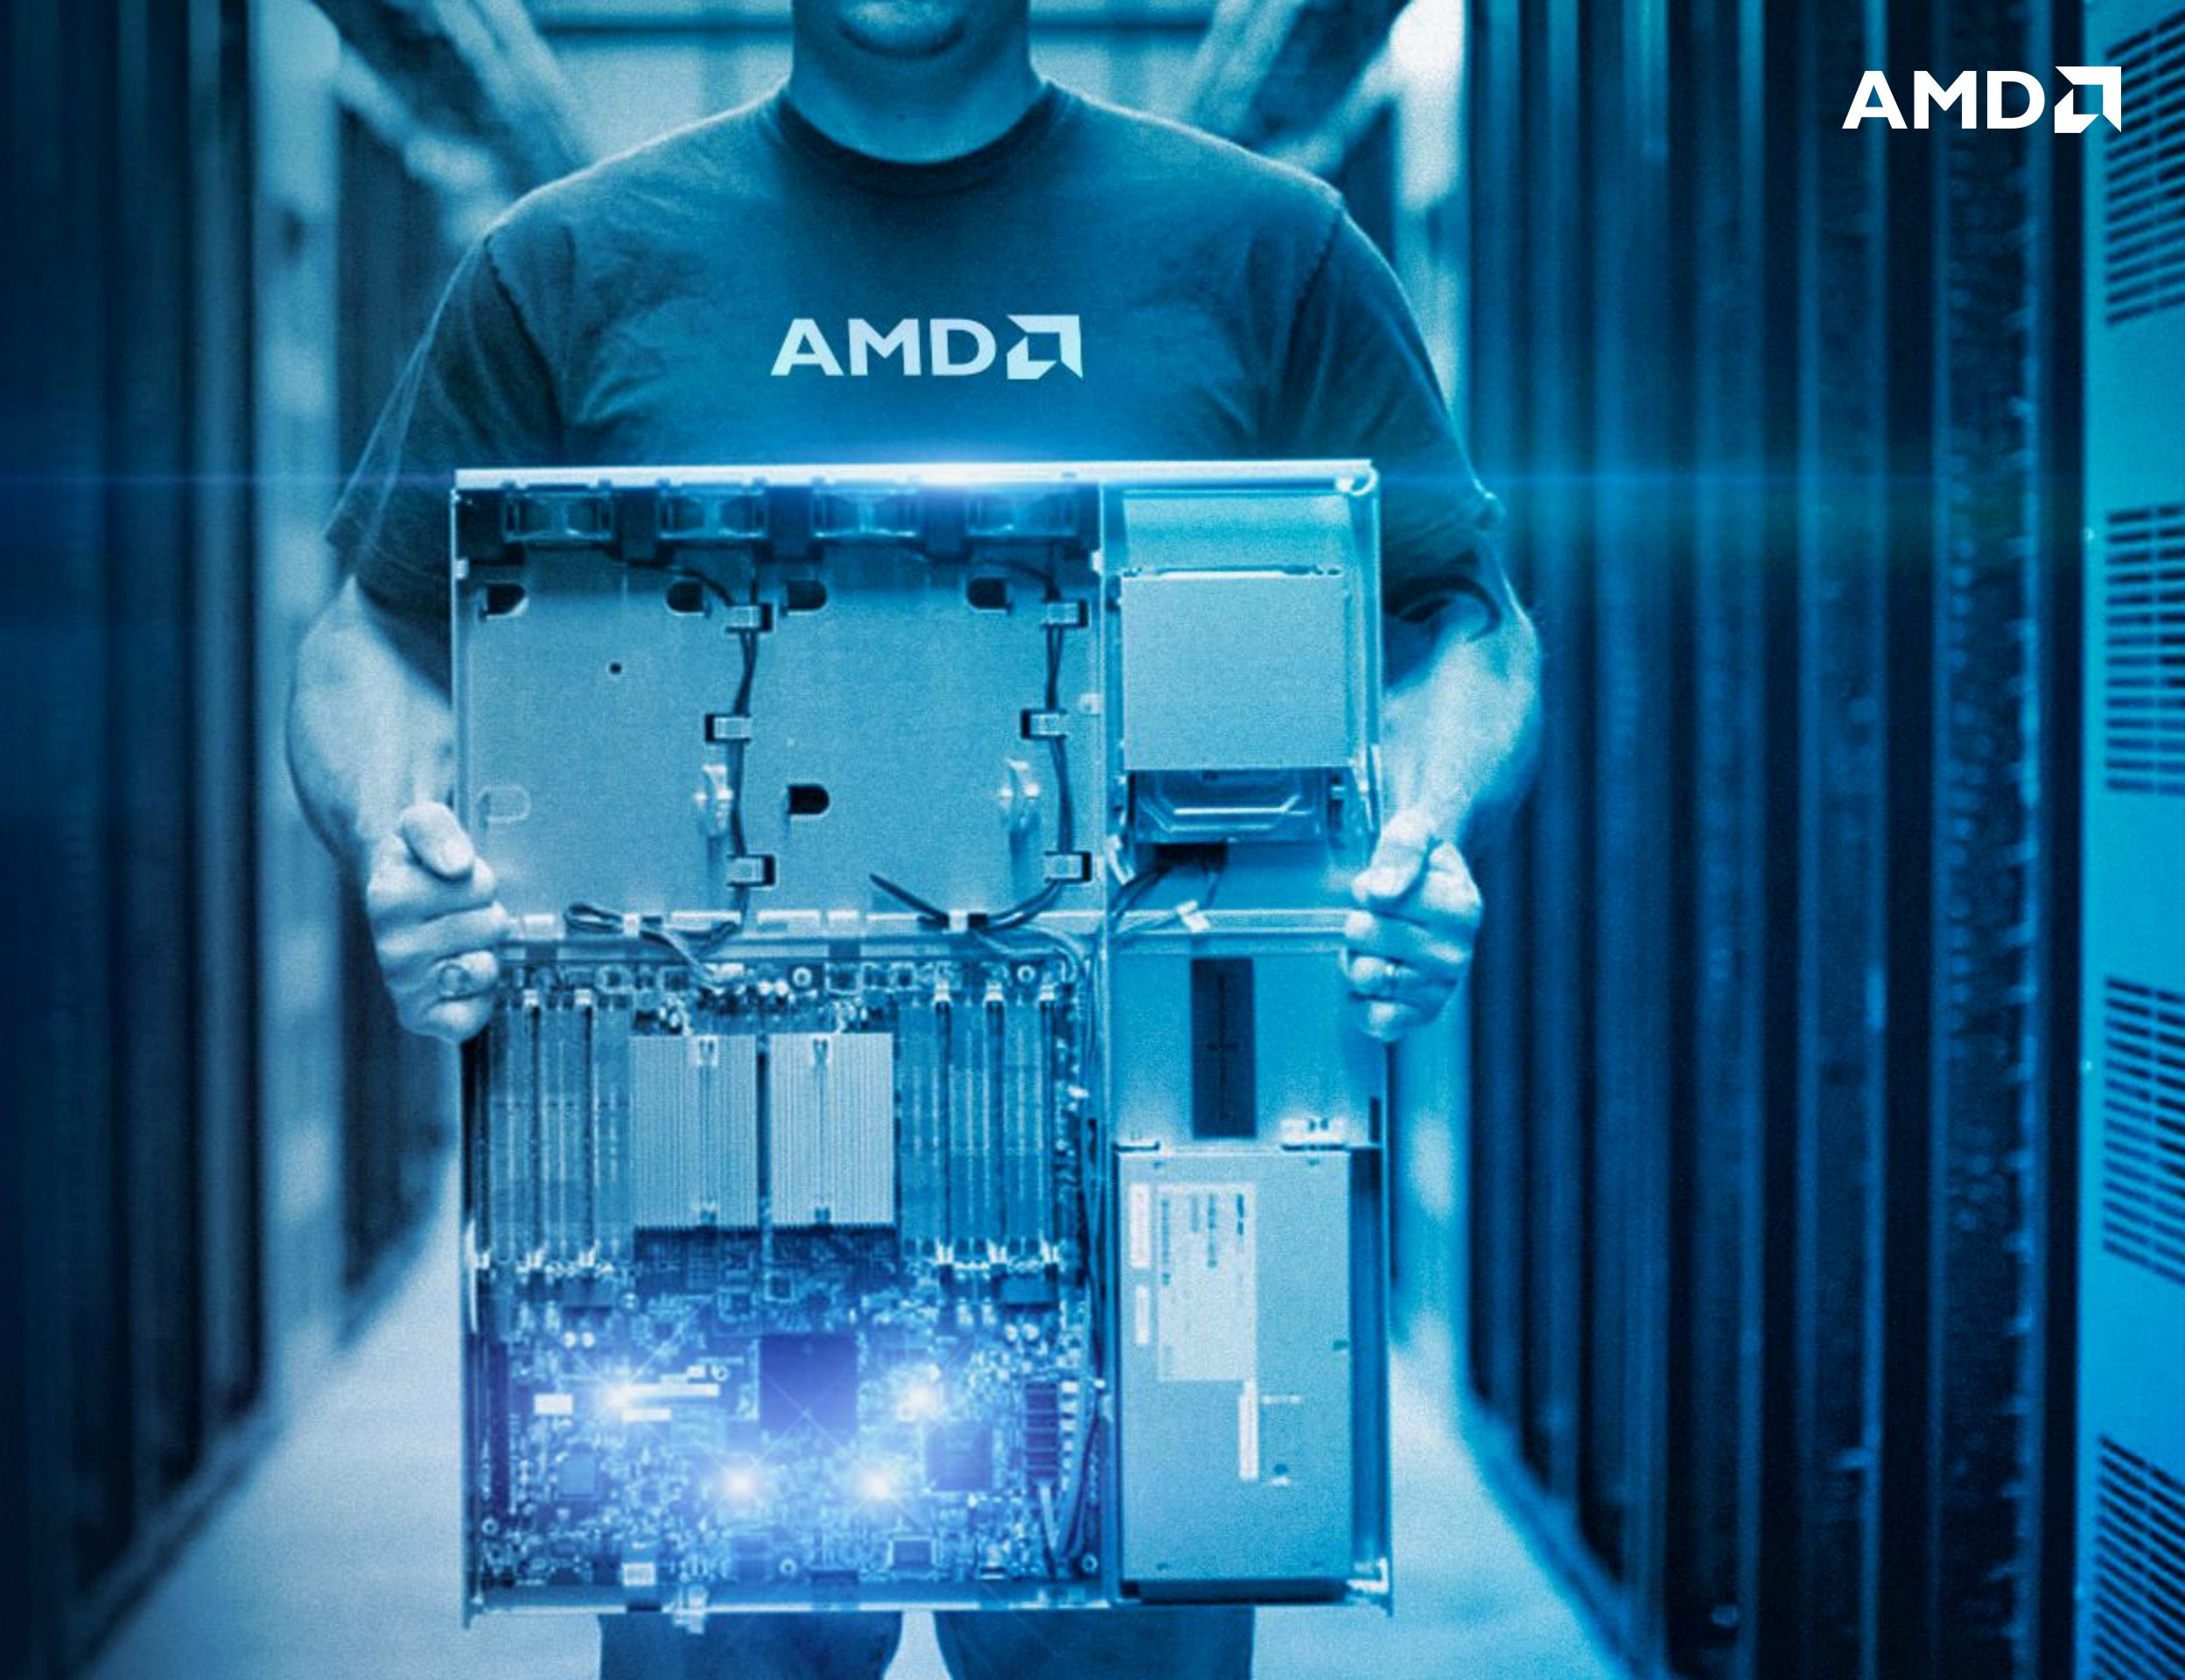
# WORLD'S FIRST 7nm DATACENTER CPU

---

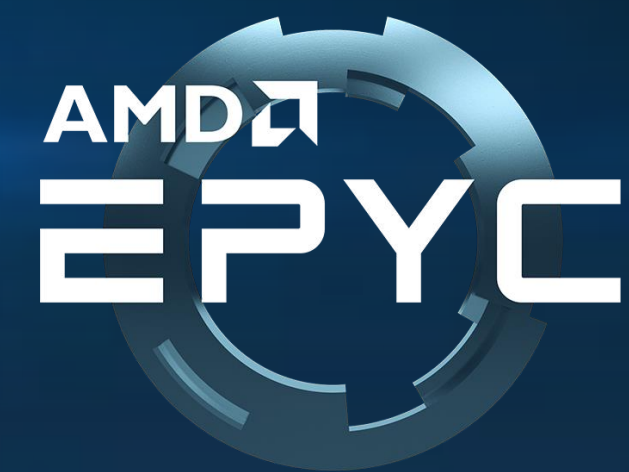
AVAILABLE 2019



# A NEW SET OF HIGH PERFORMANCE COMPUTE REQUIREMENTS DEMANDS NEW INDUSTRY LEADERS



# AMD DATACENTER ROADMAP



AMD  
RADEON  
INSTINCT

"NAPLES"

"Zen"  
Shipping

"ROME"

"Zen 2"  
Available 2019

"MILAN"

"Zen 3"  
On Track

MI25

14nm GPU  
Shipping

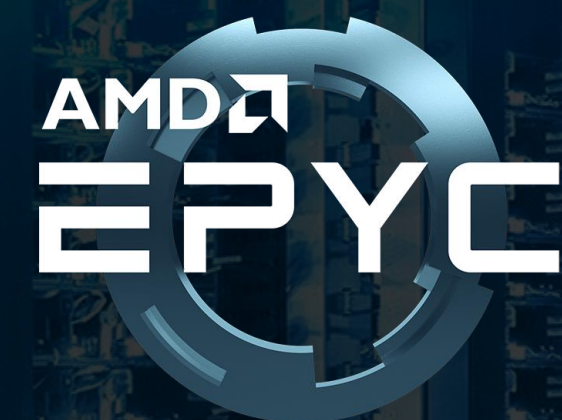
MI60

First 7nm GPU  
Available Q4

"MI-NEXT"

On Track

AGGRESSIVE LONG-TERM COMMITMENT



AMD  
RADEON  
INSTINCT



Open Compute

TODAY, AMD DELIVERS TOTAL  
DATACENTER COMMITMENT

# ENDNOTES

Slide 3, 4, 5, 7, 8

**Claim:** “Rome” is the world’s first 7nm datacenter CPU.

**Substantiation:** Based on AMD’s internal competitive assessment of all currently shipping or publicly disclosed X86 Server CPU products as of 9/6/2018

.

# ENDNOTES

## Slide 4

**Claim:** Increased instructions-per-cycle.

**Estimated increase in instructions per cycle (IPC) is based on AMD internal testing for “Zen 2” across microbenchmarks, compared to prior “Zen 1” generation CPU using combined floating point and integer benchmarks.**

**Claim:** Leadership compute, I/O and memory bandwidth.

**Substantiation:**

NAP-42 – AMD EPYC™ 7601 processor supports up to 8 channels of DDR4-2667, versus the Xeon Platinum 8180 processor at 6 channels of DDR4-2667.

NAP-43 – AMD EPYC 7601 processor includes up to 32 CPU cores versus the Xeon Platinum 8180 processor with 28 CPU cores.

NAP-44 – A single AMD EPYC™ 7601 processor offers up to 2TB/processor (x 2 = 4TB), versus a single Xeon Platinum 8180 processor at 768Gb/processor (x 2 = 1.54TB). NAP-56 – AMD EPYC™ processor supports up to 128 PCIe® Gen 3 I/O lanes (in both 1 and 2-socket configuration), versus the Intel® Xeon® SP Series processor supporting a maximum of 48 lanes PCIe® Gen 3 per CPU, plus 20 lanes in the chipset (max of 68 lanes on 1 socket and 116 lanes on 2 socket).

Based on “Zen2” design parameters versus “Zen1” and currently shipping products – core count increase from 32 to up to 64 per socket. Based on “Zen2” design parameters versus “Zen1” and currently shipping products – core count increase from 32 to up to 64 per socket. Memory bandwidth with “Zen 2” design parameters including increased memory speed across eight memory channels, I/O leadership extending to PCIeGen4.

**Claim:** Dramatically Improves Accelerator Performance

**Substantiation:** Estimates based on AMD design and AMD internal testing versus currently shipping AMD EPYC processors.

# ENDNOTES

## Slide 5

**Claim:** “Rome” doubles the bandwidth/channel.

**Substantiation:** Per Silicon Labs, provider of the PCIe Gen 4 solutions. (website here: [https://www.silabs.com/community/blog.entry.html/2015/12/04/what\\_is\\_pcie\\_gen4a-oobp](https://www.silabs.com/community/blog.entry.html/2015/12/04/what_is_pcie_gen4a-oobp)). PCIe Gen4 is a new standardized data transfer bus that will double the data transfer rate per lane of the prior Gen3 revision from 8.0 GT/s (gigatransfers/second) to 16.0 GT/s. This means that a single PCIe Gen4 interconnection will allow data rate transfers of up to 2GB/s (gigabytes/second), and a full 16 slot PCIe Gen4 interconnection for graphics cards and high-end solid state drives will allow data transfer rates of up to 32GB/s.

## Slide 7

**Claim:** “Rome” delivers 2x performance-per-socket v. previous generation.

**Substantiation:** Product not yet available, details to be shared at launch.

Claim: “Rome deliver 4x floating point per socket vs. previous generation

Substantiation: AMD Internal Estimates.

# DISCLAIMER & ATTRIBUTION

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.

Vega, Polaris and Zen are codenames for AMD architecture and are not product names.

## **ATTRIBUTION**

© 2018 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, EPYC, Ryzen, Radeon, and combinations thereof are trademarks of Advanced Micro Devices, Inc. in the United States and/or other jurisdictions.

Microsoft and DirectX are registered trademarks of Microsoft Corporation in the US and other jurisdictions. Xbox is a registered trademark of Microsoft Corporation.

“PlayStation”, the “PS” Family logo, and “PSP” are registered trademarks of Sony Interactive Entertainment Inc.

Other names are for informational purposes only and may be trademarks of their respective owners.