
AWS
re:Invent
NOV. 28 – DEC. 2, 2022 | LAS VEGAS, NV

PRT256

SPONSORED BY AMD

How AMD is taking the cloud to new heights, and what's next

Brad Smith

Corporate VP of Sales
AMD

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Agenda

Introduction

AMD technology leadership

What’s new from AMD

Cloud partners, solutions, and services

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

A Leader in High-Performance and Adaptive Computing



Pushing the Limits

27,000+ employees around the world push the limits of innovation to solve the world's most important challenges



Building the Best

Developing high-performance and adaptive computing technologies that move us forward



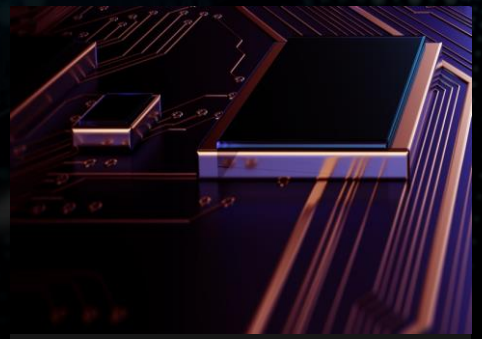
Transforming the World

Non-stop innovation for the world's creators, researchers, innovators, and explorers

AMD
together we advance_

The Next Five Years

Strategic Pillars



Compute Technology Leadership



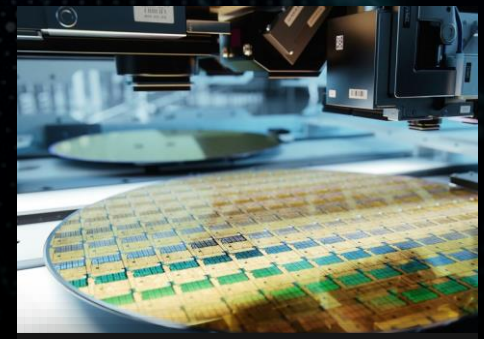
Expanding Data Center Leadership



Advancing Pervasive AI



Software Platforms & Developer Enablement



A Leading Provider of Custom Silicon and Solutions

AMD
together we advance_

AWS and AMD – Working Together

Sustainability

29% less power consumption and 29% lower greenhouse gas emissions[3] versus comparable instances

Cost Efficiency

Deliver up to **10–45% lower cost[2]** at same or better performance versus comparable instances

Performance

Faster application performance and time to solve using Amazon EC2 instances powered by world's highest performance server processor[1], AMD EPYC

Scale & Choice

Cloud scale

Worldwide availability

14 instance types

Larger vCPU sizes

x86 compatibility

1 AMD EPYC 7763; Results as of 01/28/2021 using SPECrate 2017_int_base – https://spec.org/cpu2017/results/res2020q2/cpu2017-20200525-22554.pdf
2 https://aws.amazon.com/ec2/amd/
3: See Endnote EN-1. Analysis based on AMD EPYC Bare Metal Server and Greenhouse Gas TCO Tool, v4.2, https://www.amd.com/en/processors/epyc-bare-metal-tco-tool

AMD
together we advance

AMD technology leadership:
Cloud product portfolio

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.



Our Path Forward

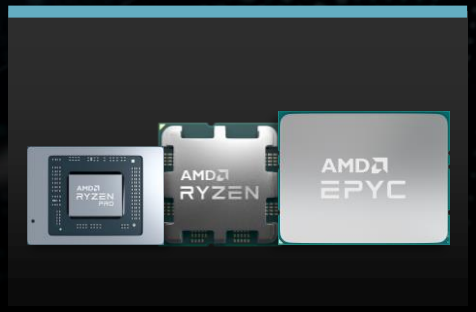
Superior Data Center Technology

Leadership **Performance**	World-Class **Efficiency**	Cutting-Edge **Security** Features

AMD EPYC AMD INSTINCT AMD PENSANDO AMD ALVEO AMD VERSAL

AMD
together we advance_

Superior Compute Technology



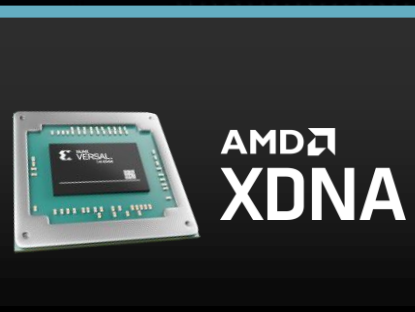
AMD EPYC

General purpose & high-performance computing



AMD INSTINCT

Technical computing and HPC



AMD VERSAL

Machine learning and AI



AMD ALVEO

Adaptive acceleration



AMD PENSANDO

Networking

AMD
together we advance_

AWS EC2 Instances Powered by AMD

Innovations in performance, cost efficiency, and x86 compatibility



Cost Efficiency

Up to 10–45% less cost[1]

At same or better performance versus comparable x86 instances

Performance

Up to 35%[2] higher vs. M5a

Up to 65% higher with HPC6a

Leading x86 price-performance[2]

Ease of migration

x86 compatibility

Application and instance size compatibility – makes it easier to migrate at reduced risk

[1]10% for Worldwide regions: https://aws.amazon.com/ec2/amd/; up to 45% For Asia Pacific (Mumbai) Region: https://aws.amazon.com/ec2/amd/in/

[2]https://aws.amazon.com/ec2/instance-types/m6a/

AMD
together we advance_



Right Size Your Workloads

AMD BASED INSTANCES

M6A|M5A
GENERAL PURPOSE

R6A|R5A
MEMORY OPTIMIZED

C6A|C5A
COMPUTE OPTIMIZED

T3A
BURST OPTIMIZED

G4AD
GPU OPTIMIZED

HPC6A
HPC OPTIMIZED

G5
GPU OPTIMIZED*

F1 and VT1
FPGA ACCELEATOR*

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

*Check with your AWS sales contact to confirm what cost savings may be available

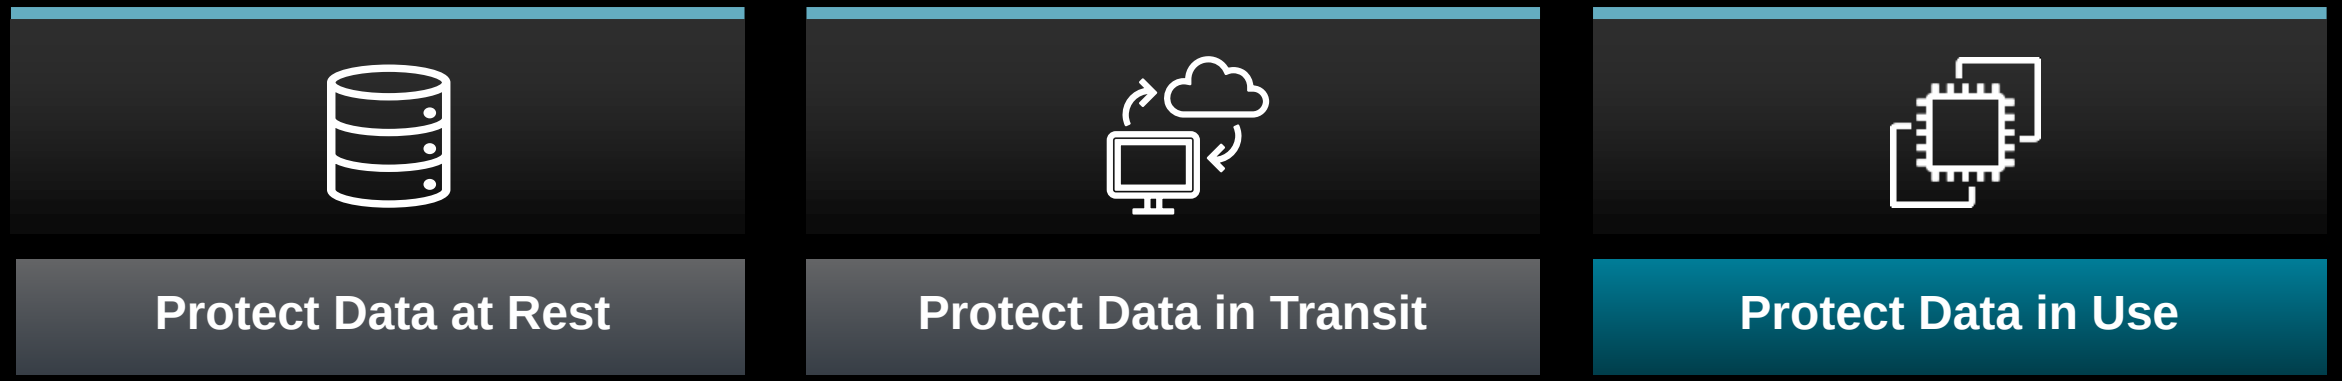
AMD
together we advance_

AMD technology leadership:
Confidential computing

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Protecting Your Applications & Data Holistically with AMD

Protect every surface of your environment from unintended access



Prevalent Security Model

AMD EPYC

SEV-SNP: Critical for cloud environments

Comprehensive Security Model

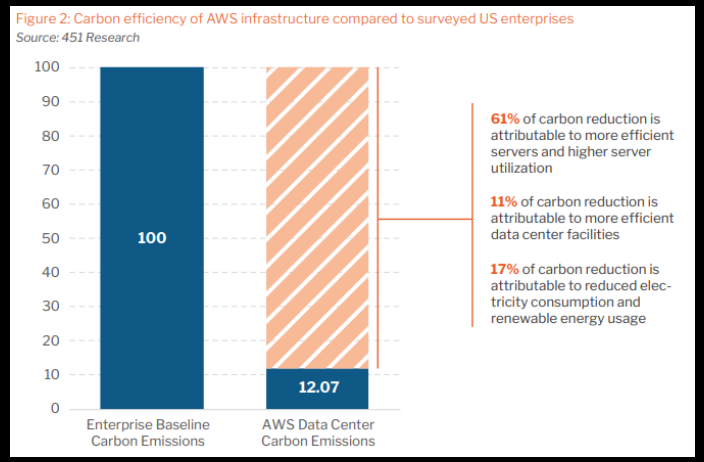
AMD
together we advance_

AMD technology leadership:
Cloud sustainability

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Cloud Sustainability: Working Together

Compared to the top 10% of most efficient enterprises surveyed, moving to AWS would deliver up to 80% reduction in carbon footprint[1]



AWS Infrastructure

More efficient servers

More efficient facilities

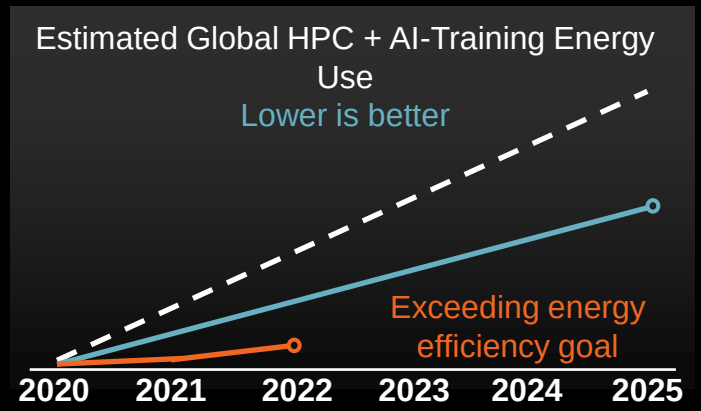
Less carbon emissions

AWS infrastructure is 3.6 times[1] more energy efficient than the median of US enterprise data centers (451 research)

Goal: Data center operations powered by 100% renewable energy by 2025

AMD Environmental Goals and Progress

On track[2] to achieve our 2025–2030 goals



30x25 Efficiency Goal

- Industry trend rate
- AMD 30x25 goal rate — Greater than 2.5x industry rate by 2025
- AMD actual energy efficiency rate

30x25 Goal: 30x increase in energy efficiency for AMD CPUs + GPUs powering servers for HPC and AI training (2020–2025)[2] — **6.8x** (as of mid-2022)[2]

50% reduction in GHG emissions from AMD operations (2020–2030) — **25%**

80% of AMD manufacturing suppliers source renewable energy by 2025 — **74%**

100% of AMD manufacturing suppliers have public GHG goals by 2025 — **74%**

[1]https://sustainability.aboutamazon.com/environment/the-cloud?energyType=true

[1]As of December 31, 2021 (unless otherwise noted). Source: AMD 2021-2022 Corporate Responsibility Report

EPYC CPUs have
superior performance/wattage and high core density

So your applications
use fewer servers
use less power

Resulting in
lower TCO
lower power consumption
lower GHG emissions

Compared to similar industry processor*

AMD CPU pricing Jan 2022. Intel Xeon Scalable CPU pricing from https://ark.intel.com, as of Jan 2022. All pricing is in USD. Analysis based on AMD EPYC Bare Metal Server and Greenhouse Gas TCO Tool, v4.2.

*Comparisons made using SPECrate 2017_int_base on 01/14/2022. SPEC, SPECrate and SPEC CPU are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information.

Includes hardware acquisition cost and 3-year opex

Processor X

17 servers



Similar industry processor
Link to SPEC result

AMD EPYC
estimated savings

32%
Lower TCO[1]

29%
Less power consumption

120 metric tons
Lower CO_2 emissions
($MTCO_2e$)

133,428 pounds of coal not burned
in USA

AMD

12 servers



AMD EPYC 64c 7763
SPEC score 861
Link to SPEC result

AMD technology leadership: High-performance computing

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

AMD EPYC AMD INSTINCT

Leading the Exascale Era

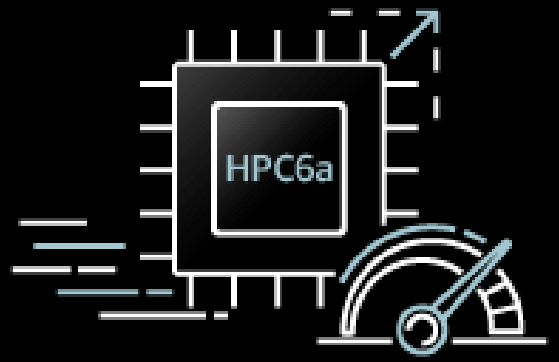


Powering World's #1 Supercomputer	Powering World's #1 Green Supercomputer	Powering World's #1 AI Supercomputer	95% Growth in TOP500 Systems Year over Year

TOP500, Green500, and HPL-AI lists, as of May 30,2022

High Performance Compute – Amazon EC2 HPC6a

HPC AWS instance powered by 3rd Gen AMD EPYC processor



High Core Density

Dual 48-core 3rd Gen AMD EPYC processors (**96 vCPUs**), up to 3.6 GHz boost frequency, and 384 GB of RAM[1]

Superior Performance

Up to **65% better price performance** over comparable Amazon EC2 x86 based, compute-intensive instances[2]

Available Now

AWS Regions:
US East (Ohio)
Europe (Stockholm)
AWS GovCloud (US-East)
AWS GovCloud (US-West)
And growing

aws PARTNER HPC Software Competency

AWS now offers a new HPC-dedicated instance family that delivers the **benefits of the cloud and on premises in a single interface[2]**

[1]Maximum boost for AMD EPYC processors is the maximum frequency achievable by any single core on the processor under normal operating conditions for server systems. EPYC-18
[2]https://aws.amazon.com/ec2/instance-types/hpc6/

AMD
together we advance_

What's new from AMD: AWS Region and instance expansion

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Global Availability of AMD-Powered Instances – And Growing

14 Unique EC2 Instance Types

97 different combinations of AMD-based instance types and sizes available globally in many AWS Regions

23 Geographic Regions

An AWS Region is a physical location in the world where AWS has multiple Availability Zones

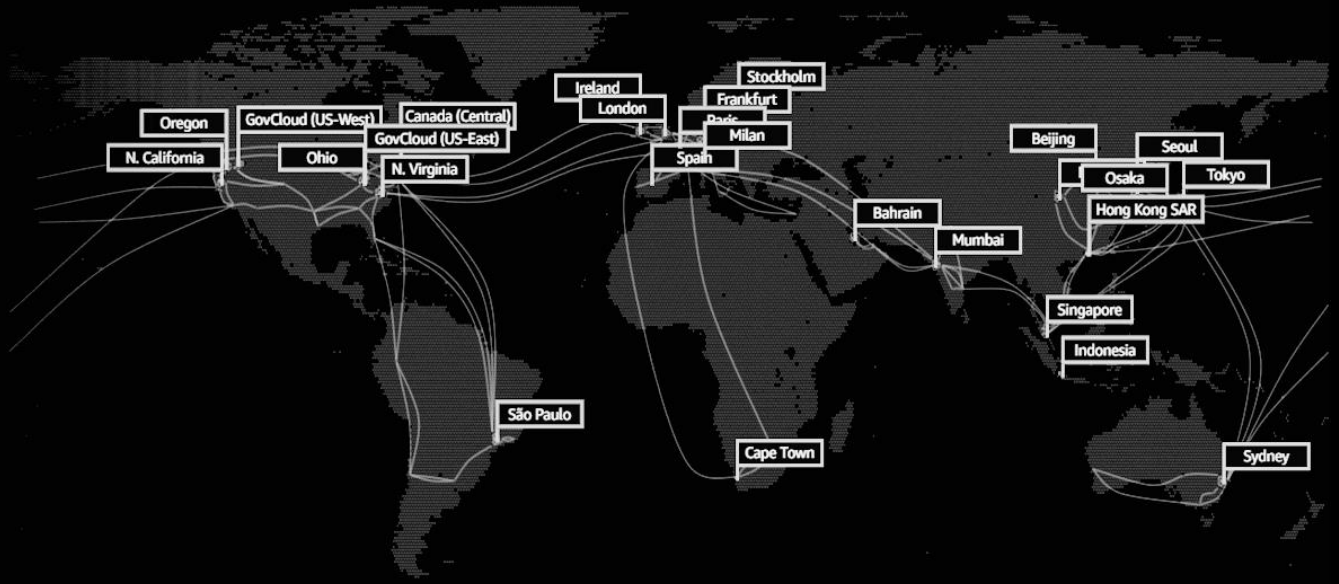
70 Availability Zones

Distinct locations that are engineered to be insulated from failures in other Availability Zones

We're Still Growing

Stay tuned for expanded availability

AMD EPYC AMD RADEON AMD VIRTEX AMD ALVEO



UP TO 45% COST SAVINGS IN ASIA PACIFIC (MUMBAI) REGION[1]
UP TO 10% COST SAVINGS IN OTHER REGIONS[2]

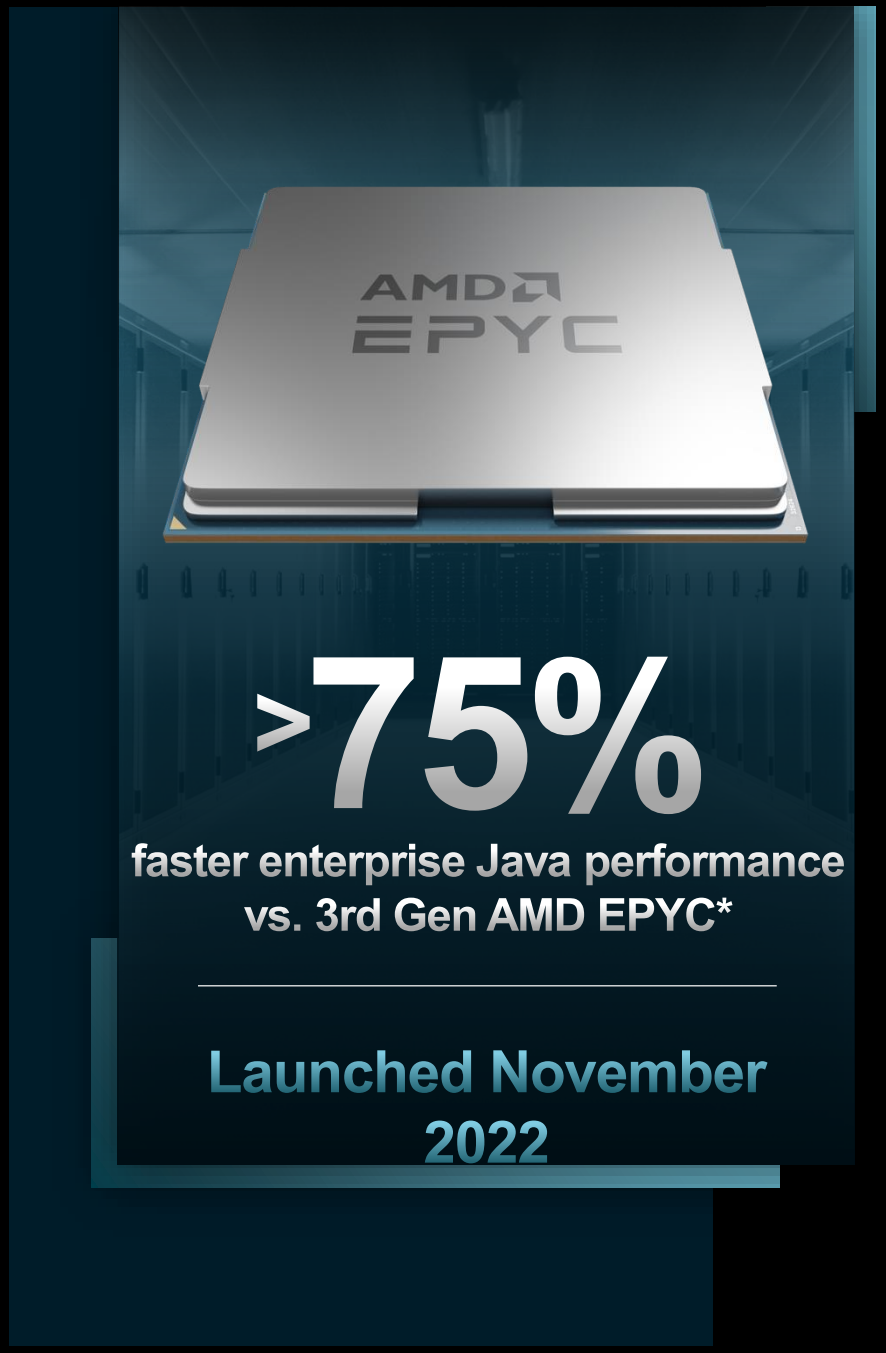
[1]https://aws.amazon.com/about-aws/whats-new/2019/11/amazon-ec2-amd-instances-are-now-available-in-asia-pacific-mumbai-aws-region/
[2]https://aws.amazon.com/ec2/amd/

For the most up to date availability zones visit – aws.amazon.com/ec2/amd/

AMD
together we advance_

What's new from AMD: Next-generation server CPU – Genoa

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.



AMD 4TH GEN EPYC

EXTENDING COMPUTE LEADERSHIP

>75%

faster enterprise Java performance vs. 3rd Gen AMD EPYC*

Launched November 2022

- Socket and Per-Core Performance
  Up to 96 Zen 4 cores in 5 nm
- Memory Bandwidth and Capacity
  12-channel DDR5
- Next-Generation I/O
  PCIe gen-5 | Memory expansion with CXL
- Advances in Confidential Computing
  Memory encryption | Direct and CXL attached

*96c Genoa vs. EPYC 7763. See Endnotes SP5-003, SP5-005, SP5-007. Preliminary data and projections subject to change

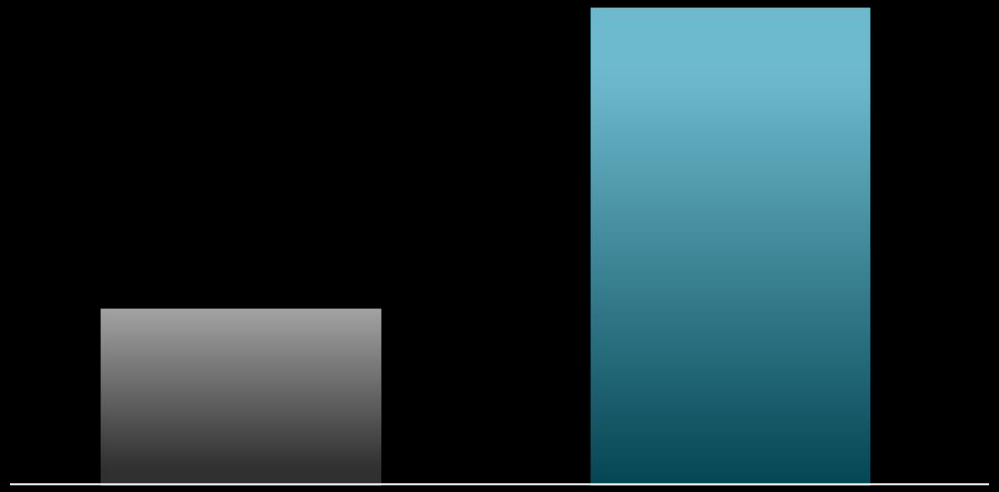
AMD
together we advance_

Faster business insights

Decision support workload derived from TPC Benchmark™ H with MySQL™

2.7x

Decision support queries



2x similar x86 top-of-stack processor

10 Containers of 8 cores each

2x 4th Gen AMD EPYC™ 9654

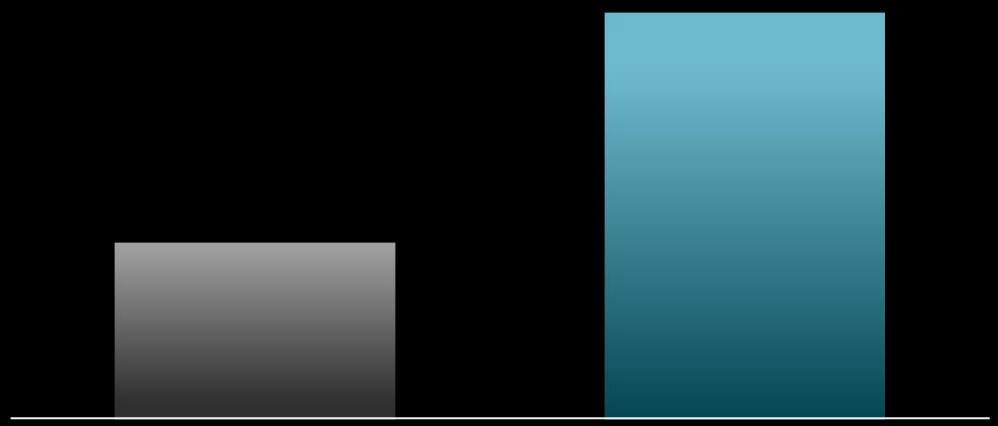
24 Containers of 8 cores each

More customer transactions

Transaction processing workload derived from TPC Benchmark™ C with MySQL™

2.3x

Online business transactions



2x similar x86 top-of-stack processor

10 Containers of 8 cores each

2x 4th Gen AMD EPYC™ 9654

24 Containers of 8 cores each

Endnotes: SP5-070, -071

AMD
together we advance_

Cloud partners, solutions, and services

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Solutions: Cloud Cost Optimization (FinOps)

Significant, easy & safe optimizations with a certified AWS Partner

AUREA CLOUDFIX

Customers can save 15-60% per AWS service*

- AWS specific cost optimizations
- High reliability, non-disruptive, low-to-no downtime fixes
- Easy implementation of fixes – deploy automatically or manually with one click


CLOUDFIX

Savings Assessment

PREPARED FOR:
Company ABC

Prepared for: John Smith

Date: 02/07/2022 Account type: Resource

Topline Metrics	Savings	Potential Savings
ROI: 467%	$518,874	9%
% Savings: 9%	Based on 84 AWS-recommended fixes.	$46,698

Readout

AWS issues 2,500 releases per year. Imagine your future savings if today CloudFix saved you: $518,874 Today's savings due to CloudFix

Relative to your overall AWS spend across all of your AWS instances, your savings opportunities are much greater. $3,000,000 Estimated overall AWS spend

CloudFix builds cost optimization in tandem with the future of the cloud and AWS. By applying straight extrapolation of the saving identified to date, CloudFix could, with minimal effort, reduce your annual AWS bill by: $270,000 Annual AWS savings

Fast installation, continuous scanning → automatic/manual deployment of easy fixes

- Fast, 5-minute installation
- Scans and analyzes your AWS environment to find quick cost optimizations
- Uses small, safe optimizations that help you find ways to improve cost-efficiency and performance
- Finds and implements AWS advisories announced each week

- **Start saving today**
  - Contact CloudFix to run a complimentary scan to help you determine the savings potential in your cloud environment
  - Visit AMD booth to connect with partners that can work with you to optimize your cloud spend today

*CloudFix, “AWS Cost Optimization You Can Trust” – https://cloudfix.aurea.com/aws/

AMD
together we advance_

Solutions: Automotive Product R&D Using HPC

Running HPC simulation on Rescale's ScaleX platform

rescale

Challenges automotive engineers and ITDMs face



Run more simultaneous simulations

Eliminate lengthy job queues

Find better solutions faster

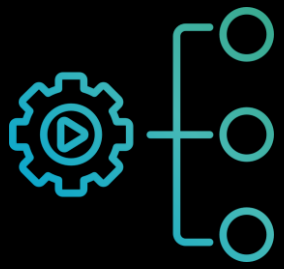
Run HPC simulation on Rescale's ScaleX platform on Amazon EC2 Hpc6a instances powered by 3rd Gen AMD EPYC processors



Simulate at unlimited scale



Innovate faster on the cloud



Optimize product design

AMD
together we advance_

Case Study: Accelerate Wind Designs More Efficient Rooftop Wind Turbines with Amazon EC2 Instances Powered by AMD

"I wish we had switched to Amazon EC2 instances powered by AMD EPYC processors a long time ago. We could have sped up our design process a lot and saved our company time."

Erika Boeing, CEO of Accelerate Wind

"Our 384-core Amazon EC2 Hpc6a instance powered by an AMD EPYC processor was 15 times faster than the 32-core on-premises system, with just 12 times more cores, taking a four-hour simulation on the workstation down to 12 minutes."

Adam Bennett, Aero-Mechanical Engineer at Accelerate Wind



CHALLENGES

Accelerate Wind wanted to realize its turbine designs more quickly, to make renewable wind energy more effective and affordable for commercial and industrial buildings

Its existing on-premises computational fluid dynamics processing provision took too long to iterate simulations and could not calculate 3D models fast enough for practical use

SOLUTION

With the assistance of AWS Partner TotalCAE, Accelerate Wind deployed 384-core Hpc6a instances powered by 3rd Gen AMD EPYC processors, providing 15 times faster CFD simulations than their 32-core on premises system from only 12 times the number of Hpc6a CPU cores.

RESULTS

Accelerate Wind was able to reduce 4-hour simulations to 12 minutes, allowing the company to iterate more quickly and optimize the price, power output, and reliability of its turbines and deliver on its promise of enabling cost-effective renewable wind energy for commercial and industrial buildings.

www.amd.com/en/case-studies

AMD
together we advance_

Case Study: TrueCar Modernizing Infrastructure Yields Cost Savings and Performance

"Instead of having to plan around a timeframe measured in months, we now use our AWS-powered Spacepods to spin up environments in as little as five minutes. This means we can innovate much more quickly."

Tommy McClung, CTO, TrueCar

AMD instances deliver 10–25% cost savings with x86-64 platform compatibility at equivalent or better performance



CHALLENGES

TrueCar wanted to operate more efficiently – both to save money and free up time for its engineers to innovate for the car buying experience

They needed to achieve cost savings while retaining full compatibility with TrueCar's existing workloads

SOLUTION

The company modernized its infrastructure with AWS and AMD-based Amazon EC2 instances

Instances from AMD are compatible with TrueCar's x86-64 architecture, are less expensive, and deliver the same or better performance than other options

RESULTS

10% lower cost while maintaining application performance

25% cost reduction from a combination of choosing the AMD instance family for its core infrastructure and right-sizing instances

Environment spin-up time improved from months to minutes

https://aws.amazon.com/solutions/case-studies/truecar/

AMD
together we advance_

Case Study: Blackboard Optimizes on Amazon EC2, Provides Video Conferencing to Millions of Teachers and Students

"By switching our preference to AMD-powered instances, we could handle 10% additional media processing using equivalent capacity."

Chris Cooksey, Sr Software Manager, Blackboard (now part of Anthology)

". . . we were actually seeing a lower CPU usage for an equivalent load on the AMD instance even though it had a smaller number of CPUs"

Gavin Llewellyn, Senior Software Engineer, Blackboard



CHALLENGES

Blackboard supports over 150 million users across education, business, and government agencies

When the COVID-19 pandemic challenged in-person learning worldwide, Blackboard saw concurrent usage of Blackboard Collaborate increase by 45 times within 4 weeks

SOLUTION

The company modernized its infrastructure with AWS and AMD-based Amazon EC2 instances

This only involved minor infrastructure reconfiguring using Terraform, testing the existing software on the new instance type, and comparing the performance

RESULTS

Supported a 4,800% increase in video conferencing usage

Saved ~28% on total media processing and batch recording costs

Improved media processing performance by 10%

https://aws.amazon.com/solutions/case-studies/blackboard-ec2-case-study/

AMD
together we advance_

Modernize & Save

Move workloads deployed on legacy instances to new AMD EPYC CPU-powered instances and leverage your opex savings

Optimize & Save

Move workloads deployed on current instances to new AMD EPYC CPU-powered instances and leverage your opex savings



Connect with Us

AMD booth – Ask the experts (booth 941)

Aiven, Altair, Aurea CloudFix, beSharp, CloudSaver, Cockroach Labs, Iternal, Red5 Pro, Rescale, TigerGraph, TotalCAE, YellowDog

Request a 1:1 meeting with AMD during AWS re:Invent or email: AWS@amd.com

Join us for our reception at Canaletto in The Venetian, tonight at 6:00 p.m.

Resources

Optimization programs with partners

Proof-of-concept support

Case studies and how-to guides

Visit amd.com/aws or aws.com/amd

AMD
together we advance_

Thank you!



Please complete the session survey in the **mobile app**

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Endnotes

SP5-070: MySQL® 8.0.17 DSS comparison based on AMD measured median scores on 2P 96-core EPYC 9654 compared to 2P 40-core Xeon Platinum 8380 running virtualized HammerDB TPROC-H SF1 (KVM Hypervisor Virtualization server environment with 4 streams, 4 virtual units, calculating throughput with 4 streams x 22 queries x 3600 divided by the slowest VU completion time in seconds) as of 11/10/2022. Configurations: 2x AMD EPYC 9654 (~126,980 TPROC-H tpm) vs. 2x Xeon Platinum 8380 (~47452 TPROC-H queries/hour) for ~2.68x the tpm performance.

SP5-071: MySQL® 8.0.17 OLTP comparison based on AMD measured median scores on 2P 96-core EPYC 9654 compared to 2P 40-core Xeon Platinum 8380 running virtualized HammerDB TPROC-C (KVM Hypervisor Virtualization server environment with 400 WH and 64 users) as of 11/10/2022. Configurations: 2x AMD EPYC 9654 (~126,980 TPROC-C tpm/~531,183 NOPM) vs. 2x Xeon Platinum 8380 (~47452 TPROC-C tpm/~224,126 NOPM) for ~2.37x the tpm/NOPM performance.

SP5-003: Estimated SPECrate®2017_int_base comparison based on internal AMD estimates. The 4th Gen EPYC score is based on engineering projections and the EPYC 7763 system score is a measured SPECrate score, using two different AMD reference systems as of 4/5/2022. Configurations: 2x 64-core AMD EPYC #### (280W TDP) versus 2x 64-core EPYC 7763 (280W TDP) for ~1.34x the performance/watt. OEM published scores and performance/CPU TDP Watt will vary based on system configuration and use of production silicon. SPEC®, SPEC CPU®, and SPECrate® are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information. SPEC - Standard Performance Evaluation Corporation

SP5-005: Server-side Java multiJVM workload demo comparison based on AMD measured testing as of 6/2/2022. Configurations: 2x 96-core AMD 4th Gen EPYC (pre-production silicon) on a reference system versus 2x 64-core EPYC 7763 on a reference system. Java version JDK18.To fit in demo time permitted, the full server-side Java run that typically takes about 20 minutes was reduced to a short-preset time where both systems were running at 99% CPU utilization, eliminating the “warm up” period data. OEM published scores will vary based on system configuration and use of production silicon.

SP5-007: 4th Gen EPYC CPUs have 12 channels of DDR5 memory vs. 3rd Gen Intel Xeon Scalable CPUs with 8 channels of DDR4 memory and announced standard “Sapphire Rapids” with 8 channels of DDR5 memory

EN-1: AMD CPU pricing Jan 2022. Intel Xeon Scalable CPU pricing from https://ark.intel.com, as of Jan 2022. All pricing is in USD. Analysis based on AMD EPYC Bare Metal Server and Greenhouse Gas TCO Tool, v4.2. Comparisons made using SPECrate 2017_int_base on 01/14/2022. SPEC, SPECrate and SPEC CPU are registered trademarks of the Standard Performance Evaluation Corporation. See www.spec.org for more information. Includes hardware acquisition cost and 3-year opex

AMD
together we advance_