
aws
SUMMIT
ANAHEIM | AUGUST 18, 2022

CMP301

AWS Graviton-based instances

Karsten Ploesser (he/him/his)
Principal Solutions Architect, ISV
Amazon Web Services

Marcin Bednarz (he/him/his)
Principal Solutions Architect, ISV
Amazon Web Services

Broadest and deepest compute platform choice

CATEGORIES

General purpose

Burstable

Compute intensive

Memory intensive

Storage (high I/O)

Dense storage

GPU compute

Graphics intensive

+

CAPABILITIES

Choice of processor
(AWS, Intel, AMD)

Fast processors
(up to 4.0 GHz)

High memory footprint
(up to 12 TiB)

Instance storage
(HDD, SSD, NVMe)

Accelerated computing
(GPUs and FPGA)

Networking
(up to 100 Gbps)

Bare Metal

Size
(Nano to 32xlarge)

+

OPTIONS

Amazon EBS

Amazon Elastic Inference

=

500+

INSTANCE TYPES

for virtually every workload and business need

AWS Graviton processors

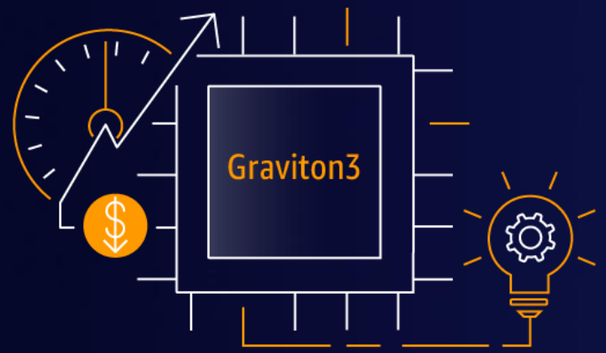

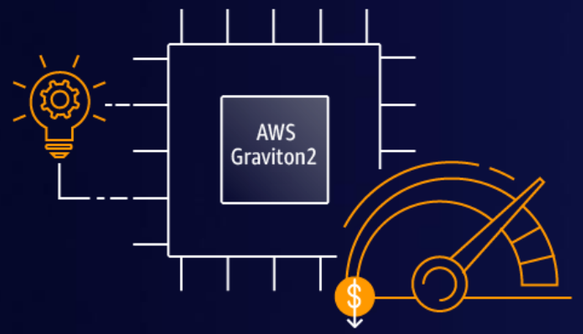

Custom AWS silicon with 64-bit Arm processor cores

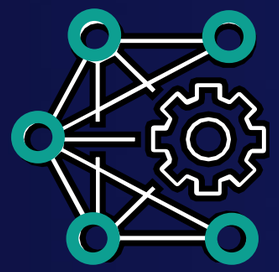

Targeted optimizations for cloud-native workloads

Rapidly innovate, build, and iterate on behalf of customers

Announcing AWS Graviton3 and Amazon EC2 C7g instances

SUPPORTING THE BEST PRICE PERFORMANCE FOR WORKLOADS IN AMAZON EC2

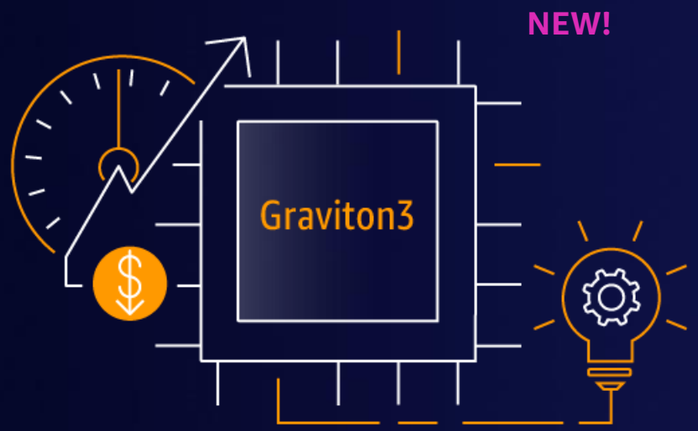

Up to 25% better performance compared to Graviton2

Up to 2x higher floating-point performance, up to 2x faster cryptographic workload performance, and up to 3x better machine learning performance compared to Graviton2

First generally available in the cloud to feature DDR5 memory

Up to 60% more energy efficient over comparable Amazon EC2 instances

C7g instances provide the best price performance for compute-intensive workloads in Amazon EC2

AWS Graviton: Broad workload applicability

Web and gaming servers

Open-source databases

High performance computing

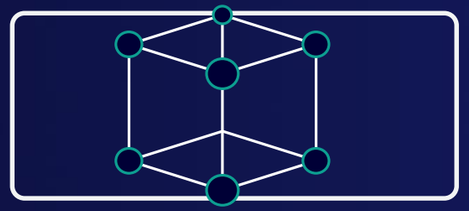

In-memory caches

Media encoding

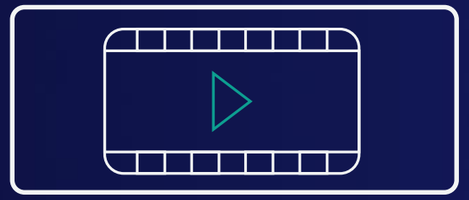

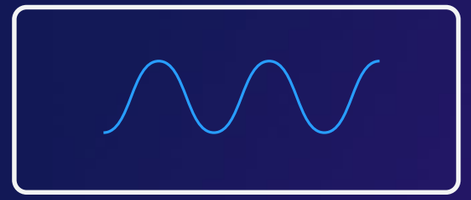

Electronic design automation

Analytics

Microservices

AWS Graviton2-based Amazon EC2 instances

UP TO 40% BETTER PRICE-PERFORMANCE OVER COMPARABLE X86-BASED INSTANCES

M6g, M6gd

General purpose workloads

T4g

Burstable general purpose workloads

C6g, C6gd, C6gn

Compute-intensive workloads

R6g, R6gd, X2gd

Memory-intensive workloads

Im4gn, Is4gen NEW!

Storage-intensive workloads

G5g NEW!

GPU-based graphics and machine learning workloads

AVAILABLE ACROSS 23 AWS REGIONS GLOBALLY*

* Not all instances available in all Regions

Components of AWS Graviton-based instances

Graviton processors

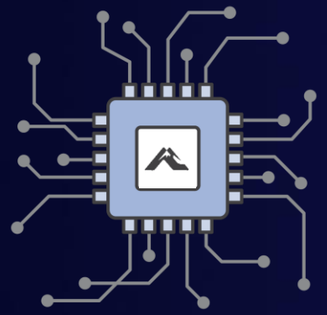

- Exceptional performance
- Reduced costs

AWS Nitro Security Chip

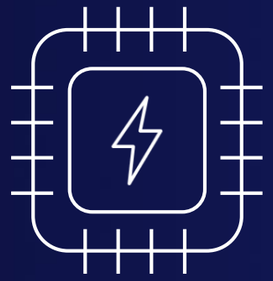

- Integrated into motherboard
- Protects hardware resources

AWS Nitro Card

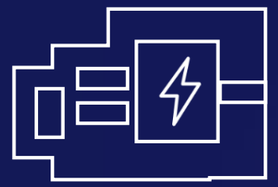

- Amazon EBS
- Elastic Network Adapter
- Monitoring and security

AWS Nitro Hypervisor

- Lightweight hypervisor
- Memory and CPU allocation
- Bare-metal-like performance

Exclusive purpose-built & modular building blocks

Honeycomb.io – Observations on Graviton2 instances

"A year out from our initial Graviton2 dogfood experiments, we can report that **92% of vCPUs in use by Honeycomb are Arm64**, spanning virtually all of our workloads and all of our environments"

honeycomb.io

"As of today, we've shifted 100% of our dogfood shepherd workload to run on C6g, using 35% fewer instances than we used with C5."

Liz Fong-Jones
Principal Developer Advocate,
Honeycomb.io

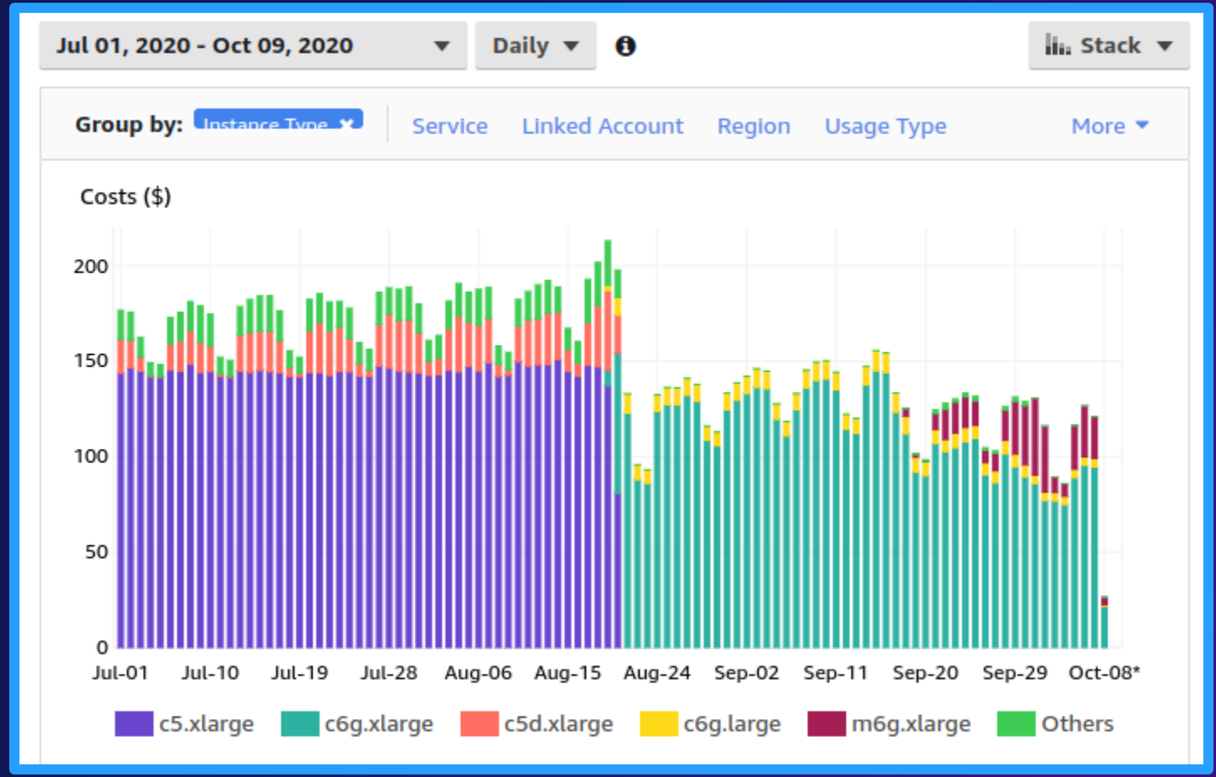

Source: https://twitter.com/lizthegrey/status/1314618333140971521?s=20

For more information, see Honeycomb's blogs on Amazon EC2 M6g instances (https://bit.ly/3SEDBIa) and one year of Graviton2 (https://bit.ly/3psG4YZ)

AWS Graviton ease of adoption

AS A RULE, THE MORE CURRENT YOUR SOFTWARE STACK THE BETTER

Difficulty	Workload	Actions
Virtually no effort	Amazon RDS, Aurora, ElastiCache, OpenSearch Service, MemoryDB & Neptune	Upgrade to latest and enjoy
Super easy	Amazon EMR	Typically, just works
Pretty easy	AWS Lambda	Typically, just works with Lambda managed runtimes or base images. Watch: JNI or Python-native modules
Quite easy	Linux – Interpreted and JIT'd languages (e.g., Java, PHP, Node.js)	Select Arm64 AMI and Install Bonus if containerized Watch: JNI or Python-native modules
More involved	Linux – Compiled languages (e.g., C/C++, Python, Go)	Select Arm64 AMI and compile Watch: port any intrinsics or assembly
Some work, high reward	Microsoft Windows – .NET	Migrate to Linux + .NET core on Arm64
Sorry, not yet	Microsoft Windows	Microsoft Windows Server not yet available for Arm64

AWS Graviton2: Industry benchmarks and workloads

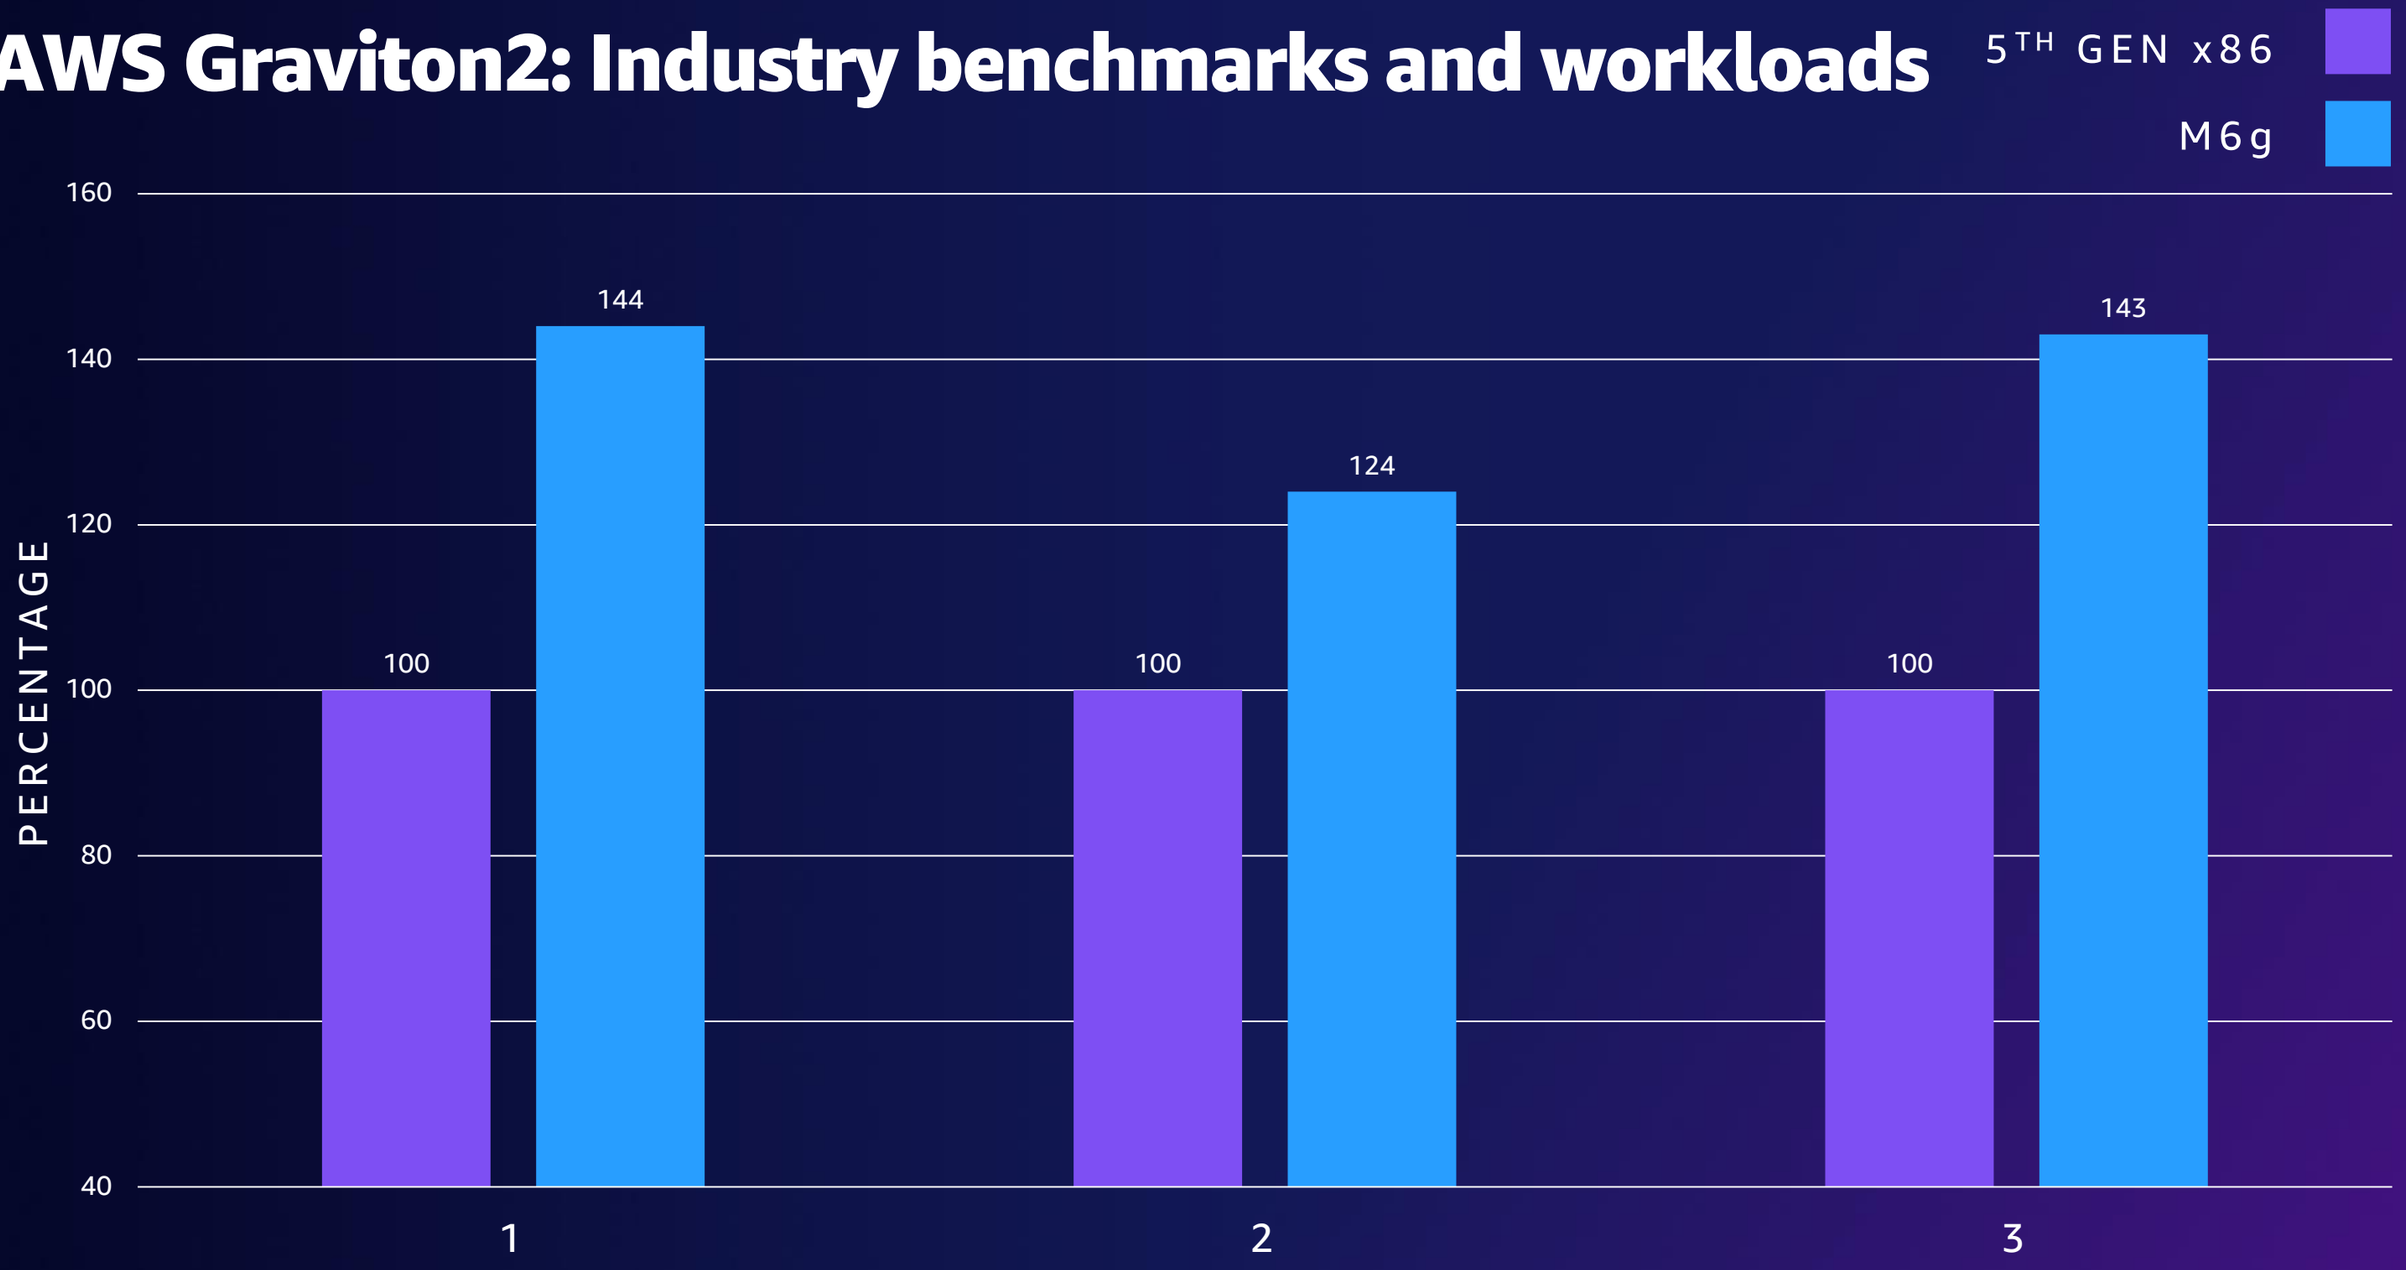

AWS Graviton Ready Program

CERTIFIED PARTNER SOLUTIONS FOR GRAVITON CUSTOMERS

AWS Graviton2 Processor Power Efficiency

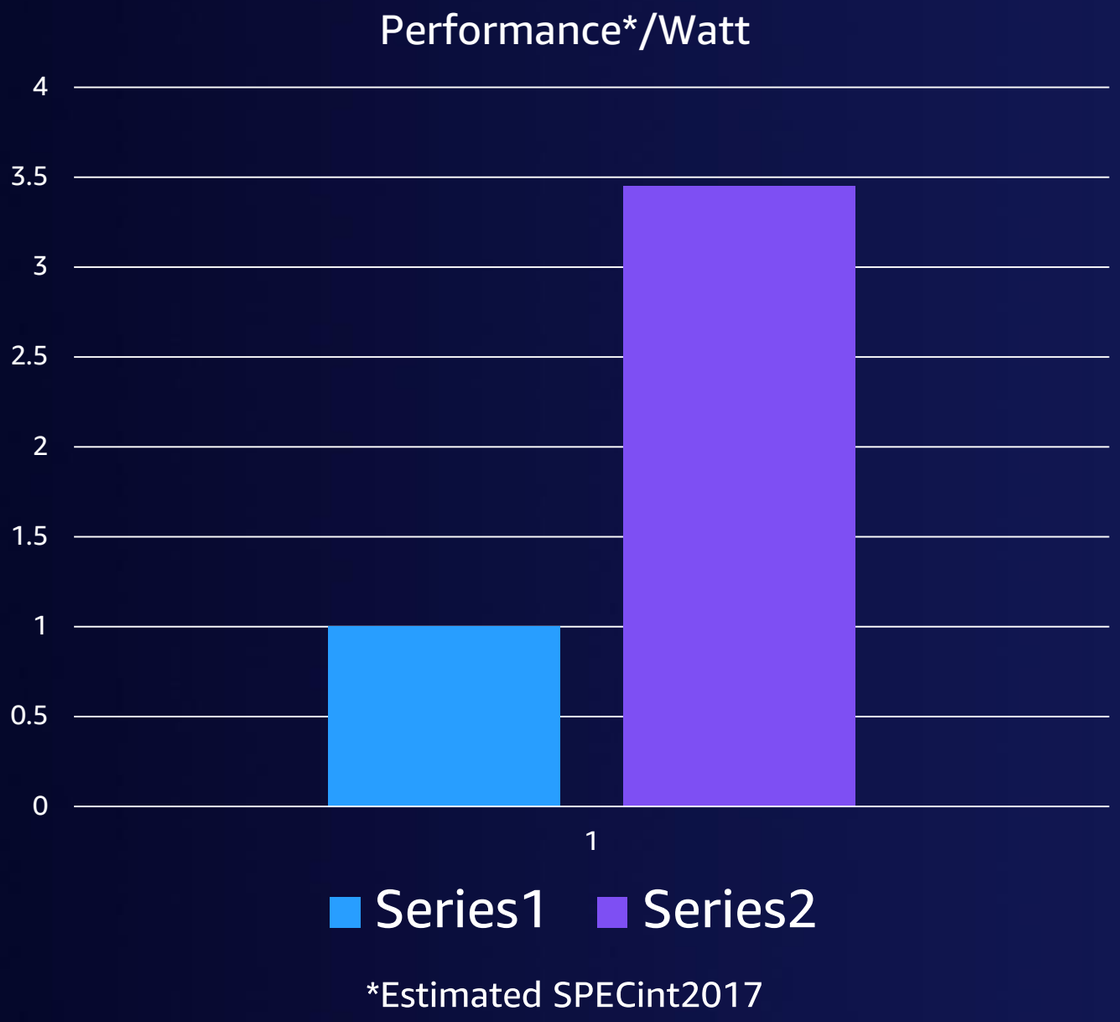

*Estimated SPECint2017

Lower power

→ Higher density

→ Lower costs

→ Lower carbon footprint

Summary

- Up to 40% better price-performance
- Broad spectrum of workloads
- Most applications easily run on multiple processor architectures
- Supported by popular Linux distributions
- Extensive ecosystem of independent software vendors (ISVs)
- AWS Graviton3-based C7g instances now generally available

Workshop URL

graviton2-workshop.workshop.aws

Learn in-demand AWS Cloud skills

AWS Skill Builder

Access **500+ free** digital courses and Learning Plans

Explore resources with a variety of skill levels and **16+** languages to meet your learning needs

Deepen your skills with digital learning on demand

Train now

AWS Certifications

Earn an industry-recognized credential

Receive Foundational, Associate, Professional, and Specialty certifications

Join the **AWS Certified community** and get exclusive benefits

Access **new** exam guides

Thank you!

Karsten Ploesser
linkedin.com/in/karstenploesser

Marcin Bednarz
linkedin.com/in/marcinbednarz