AWS re:Invent

CMP201-L

Leadership session: Redefining Amazon EC2 for every workload

Dave Brown
Vice President
Amazon Elastic Compute Cloud (EC2)
Amazon Web Services

AWS re:Invent

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

aws

Broadest customer base for compute

STARTUP

zynga, lyft, Pinterest, SmugMug, Nextdoor, CODE, UrbanPro, ringDNA, zendesk, GAMESPARKS, dropcam, yelp, REDFIN, OOYALA, coinbase, wework, fuze, slack, Zillow, Spotify, 热云数据 More Than Data, asana, reddit, twilio, splunk>, EASY TAXI, SHAZAM, Zocdoc, mojo

ENTERPRISE

GE, Unilever, verizon, NETFLIX, HESS, DOW JONES, NASDAQ OMX, intuit, COMCAST, GoPro, NOVARTIS, Coca-Cola, CapitalOne, MATTEL, Nestlé, ThermoFisher SCIENTIFIC, BOSCH, SAMSUNG, Shell, NFL, Kellogg's, Pfizer

PUBLIC SECTOR

LOCKHEED MARTIN, CLARITAS GENOMICS, NASA, THE UNIVERSITY OF CHICAGO, American Heart Association. Learn and Live, YVR Vancouver International Airport Beyond, Every Day., Ministry of Justice, CDC Centers for Disease Control and Prevention, HARVARD MEDICAL SCHOOL, Cleveland Clinic, coursera, USDA, FDA, finra, Penn University of Pennsylvania, New York Public Library, LWV League of Women Voters, wefarm, ROYAL OPERA HOUSE, Landgate, 東京大学 The University of Tokyo, Ocean Conservancy, King County, KAPLAN

SI & ISV

CODESHIP, LOGGLY, GENESYS, freshdesk, 2ND WATCH, databricks, CHEF, wipro, Adobe, accenture, clearscale, dedalus mais que cloud computing, Booz | Allen | Hamilton strategy and technology consultants, Cognizant, Capgemini, CloudKinetics, leidos, slalom, ctp Cloud Technology Partners a Hewlett Packard Enterprise company, logicworks

Figure 1. Magic Quadrant for Cloud Infrastructure as a Service, Worldwide



Source: Gartner (July 2019)

AWS RECOGNIZED AS A CLOUD LEADER FOR THE 9TH CONSECUTIVE YEAR

Gartner, Magic Quadrant for Cloud Infrastructure as a Service, Worldwide, Raj Bala, Bob Gill, Dennis Smith, David Wright, July 2019. ID G00365830. Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose. The Gartner logo is a trademark and service mark of Gartner, Inc., and/or its affiliates, and is used herein with permission. All rights reserved.

“THERE IS NO COMPRESSION ALGORITHM FOR EXPERIENCE”

ANDY JASSY, CEO, AWS

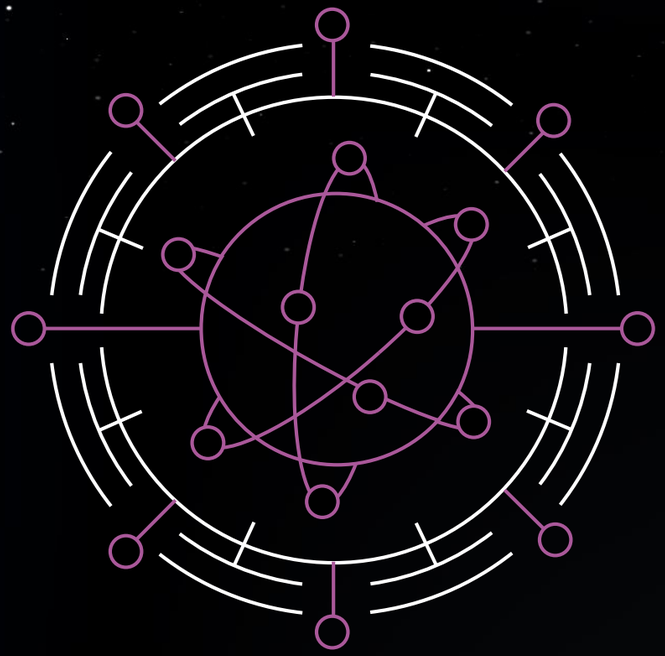
Best availability in the cloud

THE NEXT LARGEST CLOUD PROVIDER HAD



MORE DOWNTIME HOURS THAN AWS

AS MEASURED IN 2018



COMPUTE PLATFORM FOR EVERY WORKLOAD

BUILT TO SUPPORT VIRTUALLY EVERY WORKLOAD



World-class performance, security, and innovation



Broadest and deepest platform choice



Abstractions for managing cost and complexity



Enabling enterprise applications



Bringing the cloud closer to you

Broadest and deepest 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 and NVMe)

Accelerated computing
(GPUs and FPGA)

Networking
(up to 100 Gbps)

Bare Metal

Size
(Nano to 32xlarge)

OPTIONS

Amazon EBS

Amazon Elastic Inference

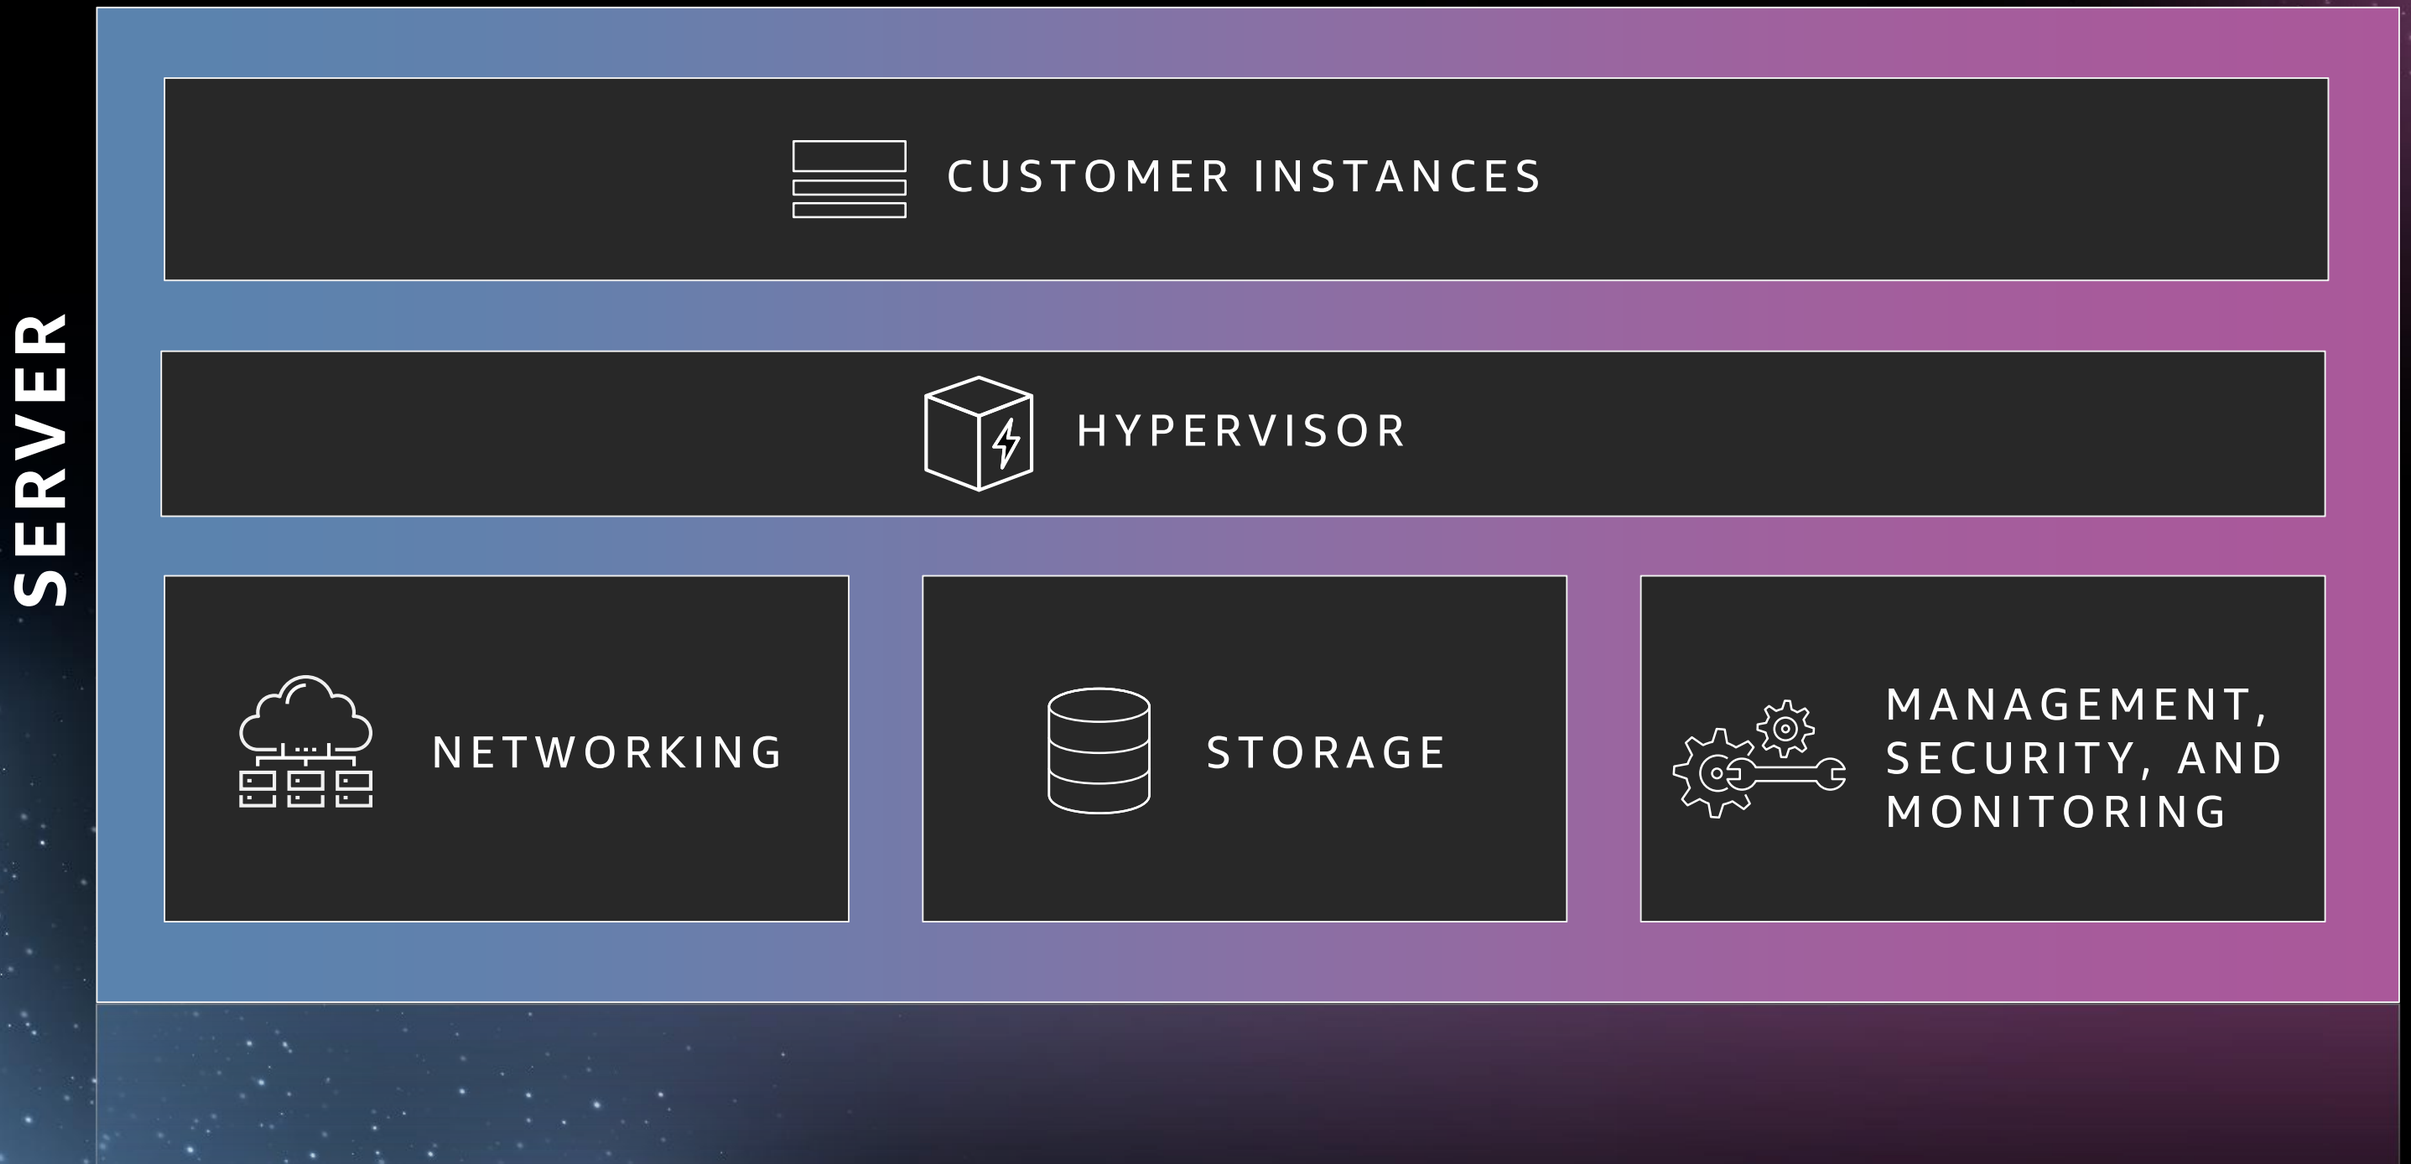
MORE THAN

270

INSTANCE TYPES

for virtually every workload and business need

Original EC2 host architecture



Original EC2 host architecture



INNOVATION

Unable to achieve the pace of innovation that our customers demanded



PERFORMANCE

Unable to achieve bare metal-like performance



PATCHING

Interruption to customer workloads when hypervisor needed to be patched



EFFICIENCY

30% of available compute capacity lost to the hypervisor

SERVER

CUSTOMER INSTANCES

HYPERVISOR



NETWORKING



STORAGE

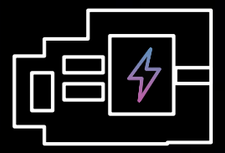


MANAGEMENT, SECURITY, AND MONITORING

Innovation enabled by AWS Nitro System

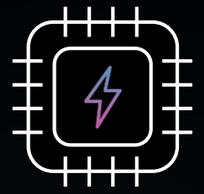
MODULAR BUILDING BLOCKS FOR RAPID DESIGN AND DELIVERY OF EC2 INSTANCES

NITRO CARD



Local NVMe storage, Amazon EBS, networking, monitoring, and security

NITRO SECURITY CHIP



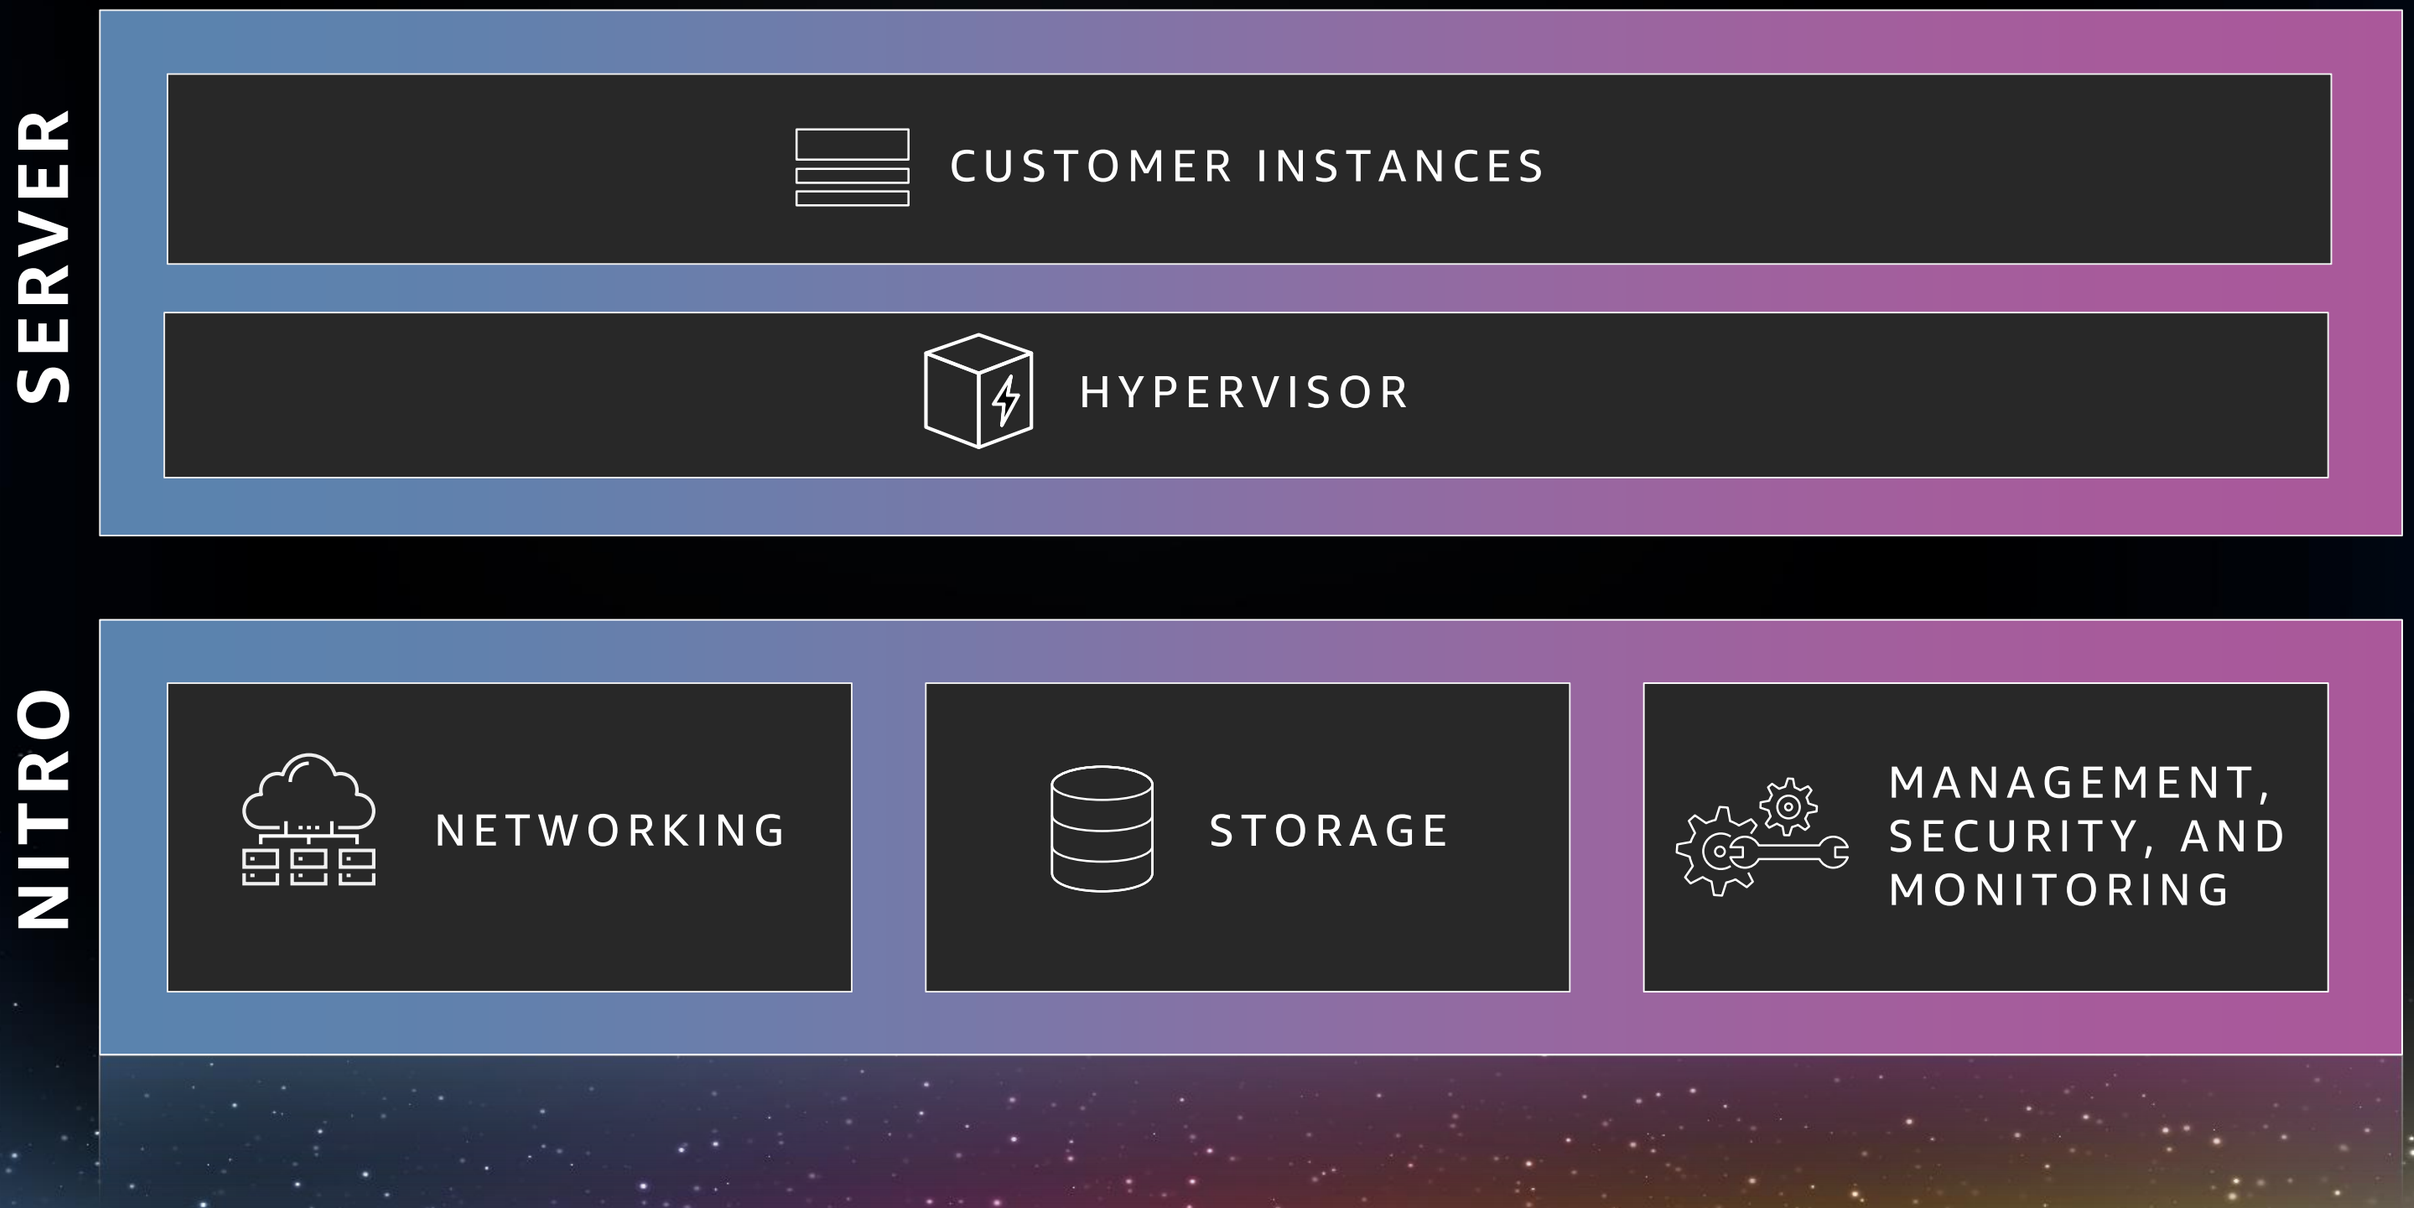
Integrated into motherboard
Protects hardware resources

NITRO HYPERVISOR



Lightweight hypervisor
Memory and CPU allocation
Bare Metal-like performance

The Nitro System architecture



AWS Nitro System enables . . .



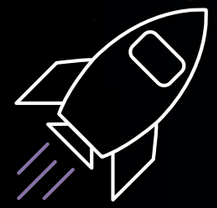
PERFORMANCE

Better performance across CPU, networking, and storage



SECURITY

Enhanced security that continuously monitors, protects, and verifies the instance hardware and firmware



INNOVATION

Building blocks can be assembled in many different ways, giving us the flexibility to design and rapidly deliver EC2 instances

AWS Nitro System enables . . .



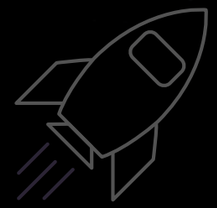
PERFORMANCE

Better performance across CPU, networking, and storage



SECURITY

Enhanced security that continuously monitors, protects, and verifies the instance hardware and firmware



INNOVATION

Building blocks can be assembled in many different ways, giving us the flexibility to design and rapidly deliver EC2 instances

Flying satellites

REAL-TIME SAMPLING OF SATELLITE DATA

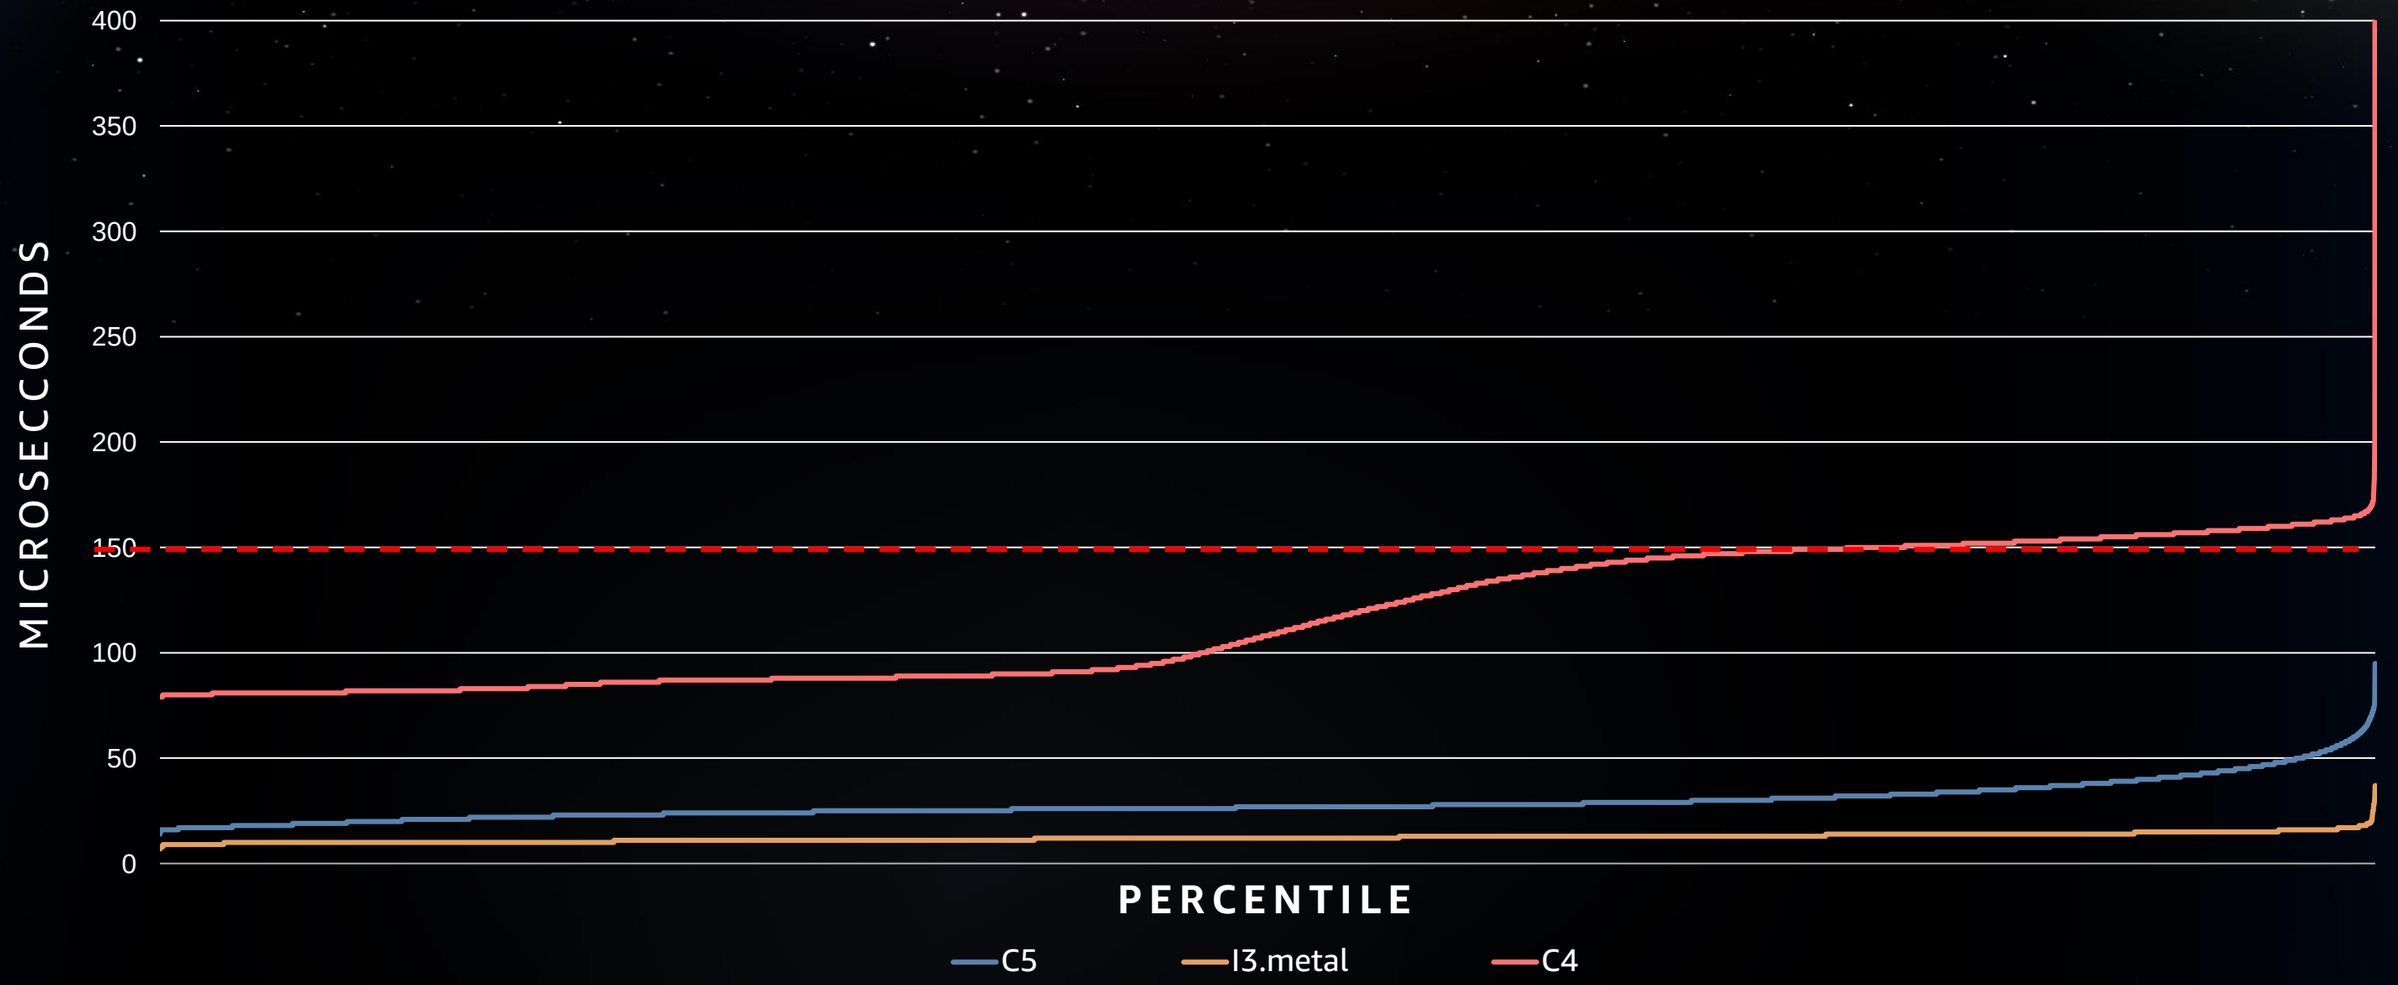
Customer migrated their satellite sampling workload to AWS

Satellite data needed to be sampled at regular intervals

Required no more than 150 µs latency on each sample

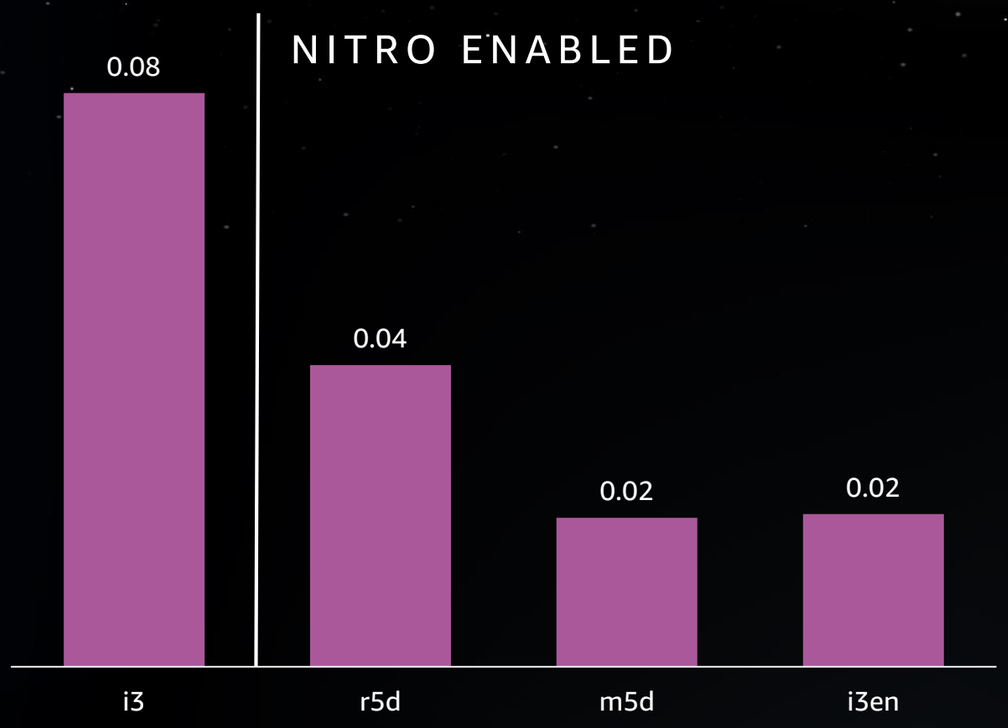


Nitro enables performance improvements

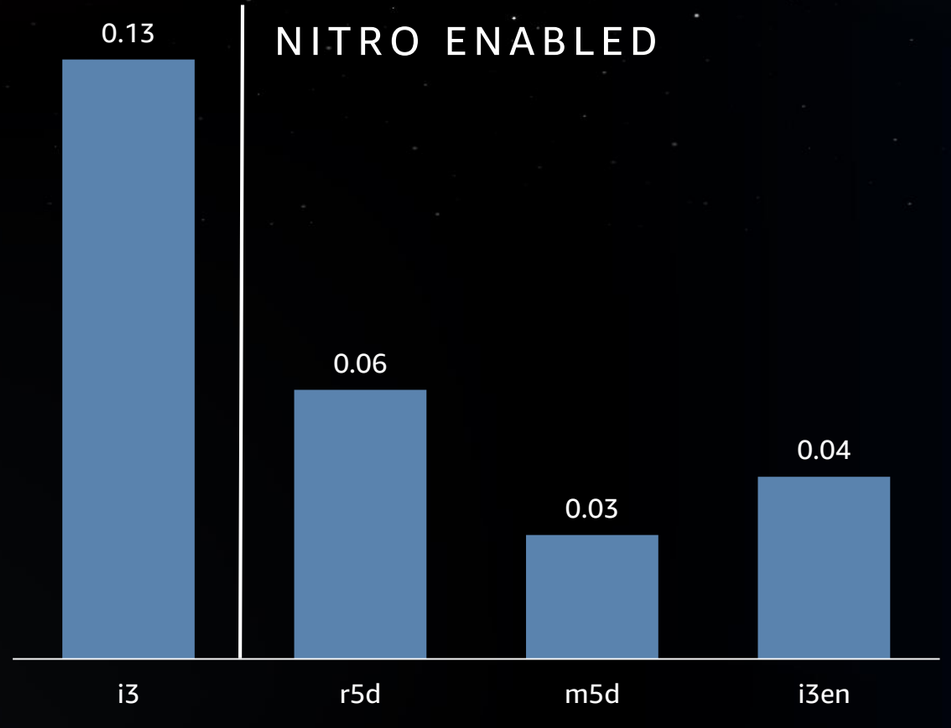


Nitro enables performance improvements

UP TO 4X IMPROVEMENT IN INSTANCE STORAGE LATENCIES



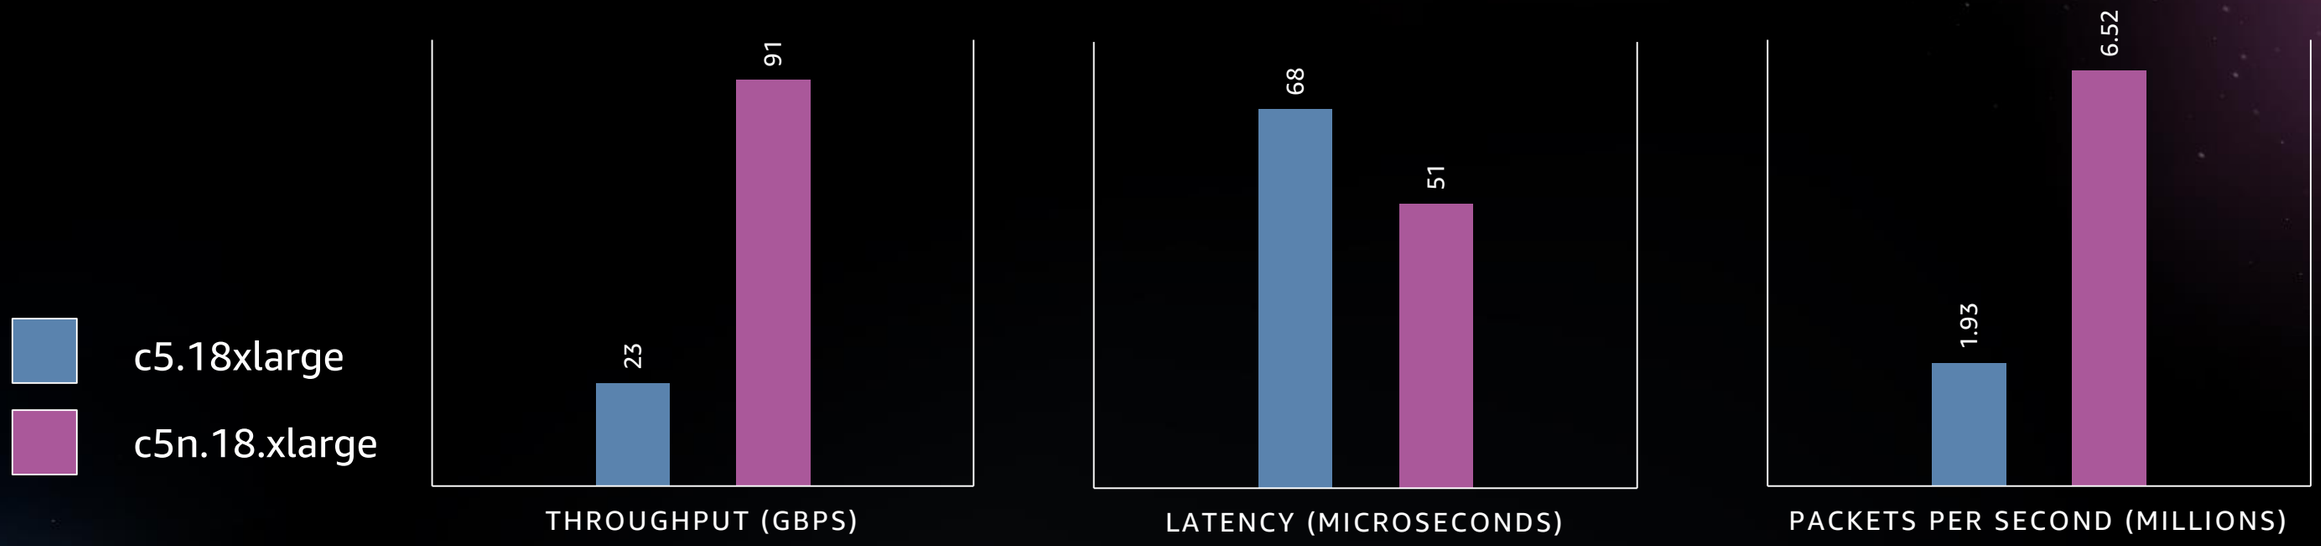
AVERAGE LATENCY (P50)



TAIL LATENCY (P99)

Nitro enables performance improvements

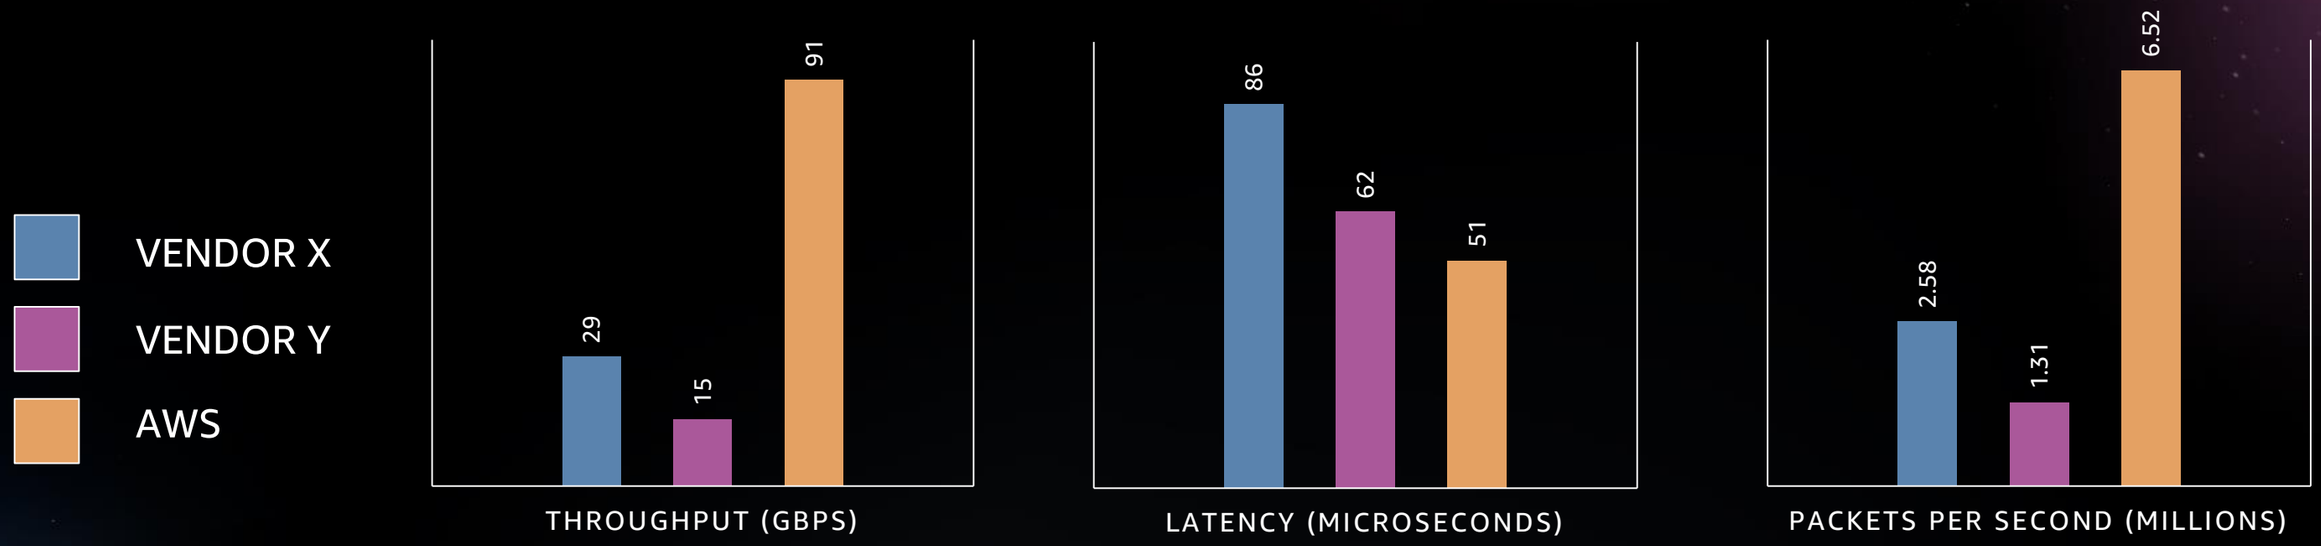
UP TO 4X IMPROVEMENT IN INSTANCE NETWORK THROUGHPUT



ENTERPRISE STRATEGY GROUP, PUBLISHED 2019

Nitro enables performance improvements

UP TO 4X IMPROVEMENT IN INSTANCE NETWORK THROUGHPUT



ENTERPRISE STRATEGY GROUP, PUBLISHED 2019

AWS Nitro System enables . . .



PERFORMANCE

Better performance across CPU, networking, and storage



SECURITY

Enhanced security that continuously monitors, protects, and verifies the instance hardware and firmware



INNOVATION

Building blocks can be assembled in many different ways, giving us the flexibility to design and rapidly deliver EC2 instances

Improved security with Nitro



Encryption of all communication channels

Passive communication design

Secure boot of components and system

Systems can be fully updated without downtime

No shell available on the Nitro System or server

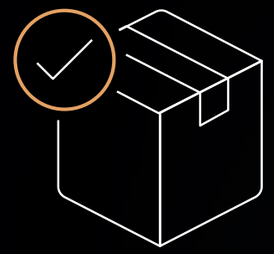
How to think about the challenge of protecting secrets on instances?

NEW!

COMING IN 2020

Nitro Enclaves

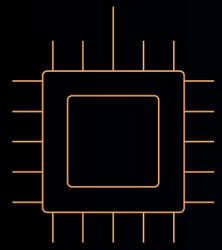
KEEP HIGHLY SENSITIVE DATA ISOLATED DURING APPLICATION PROCESSING



Create a completely isolated compute environment to process highly sensitive data



No storage, no ability to log in to the instance, and no network connectivity – has its own secure communication channel



Available on a range of instance types, including M5, C5, R5, and i3en

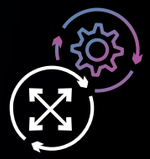
NEW!

COMING IN 2020

Nitro Enclaves

KEEP HIGHLY SENSITIVE DATA ISOLATED DURING APPLICATION PROCESSING





Secure isolation



Flexible resource allocation



Support for wide range of workloads

AWS Nitro System enables . . .



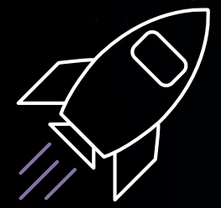
PERFORMANCE

Better performance across CPU, networking, and storage



SECURITY

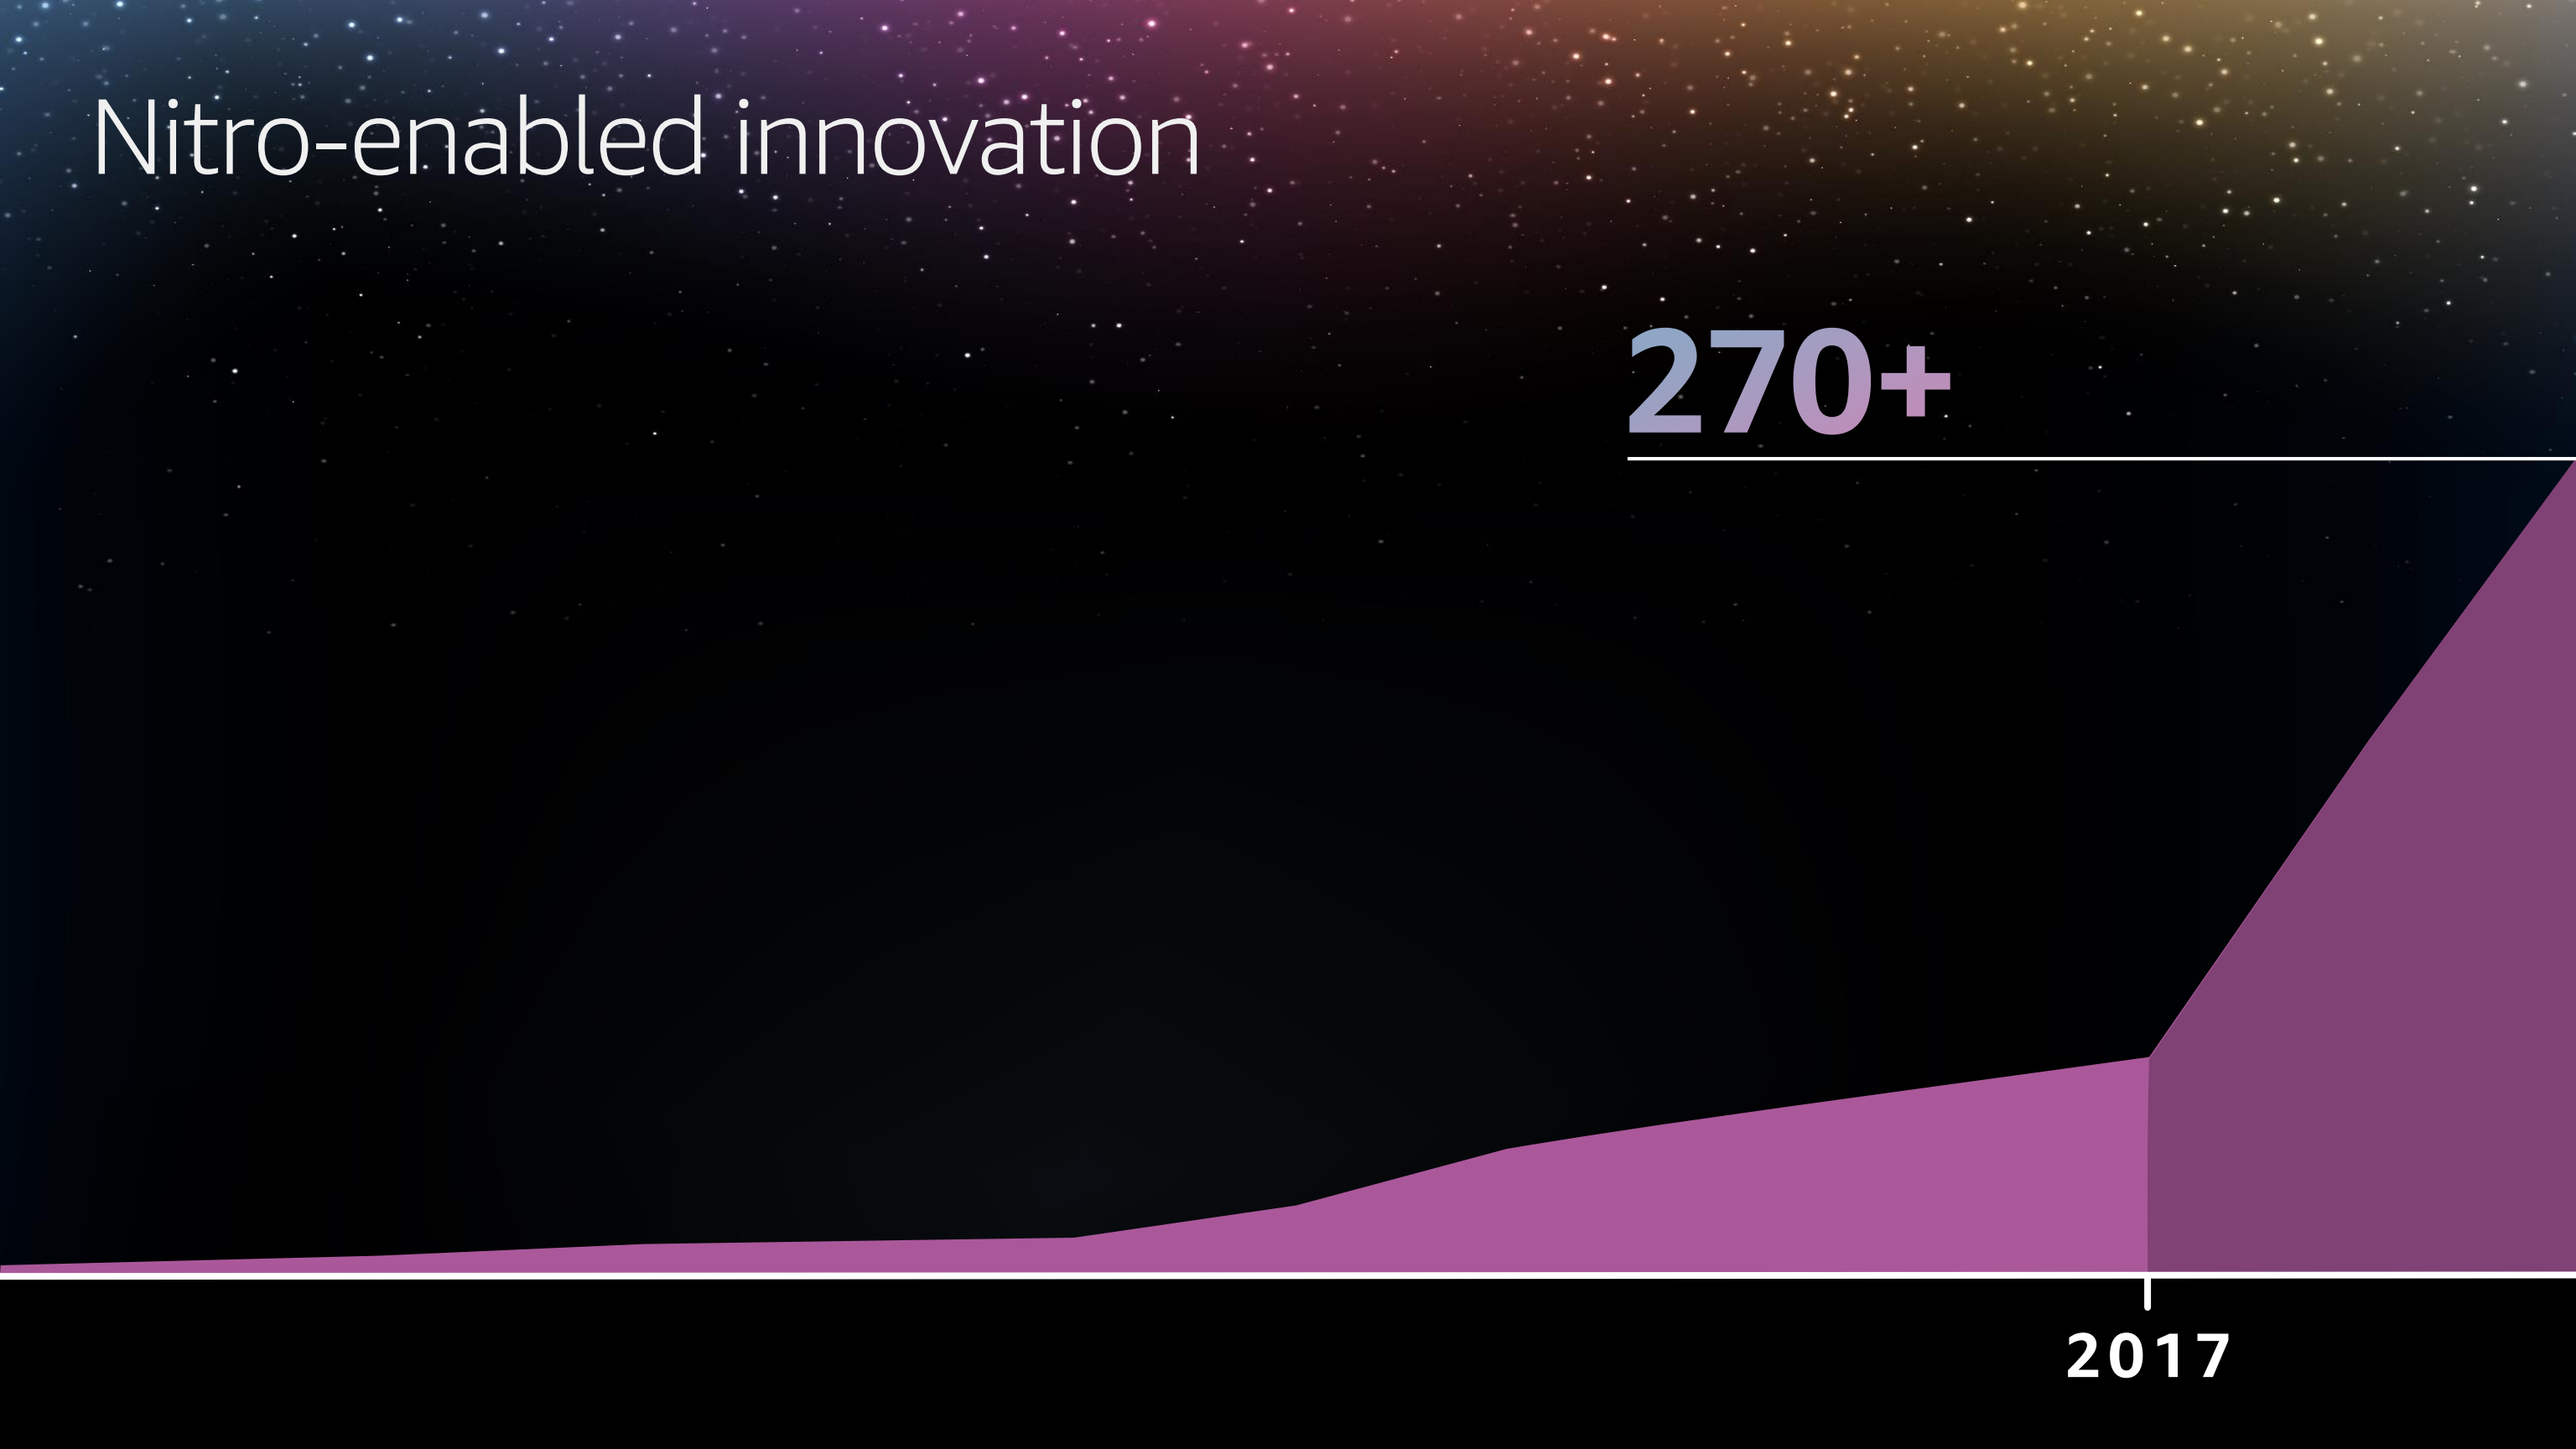
Enhanced security that continuously monitors, protects, and verifies the instance hardware and firmware



INNOVATION

Building blocks can be assembled in many different ways, giving us the flexibility to design and rapidly deliver EC2 instances

Nitro-enabled innovation



BUILT TO SUPPORT VIRTUALLY EVERY WORKLOAD



World-class performance, security, and innovation



Broadest and deepest platform choice



Abstractions for managing cost and complexity



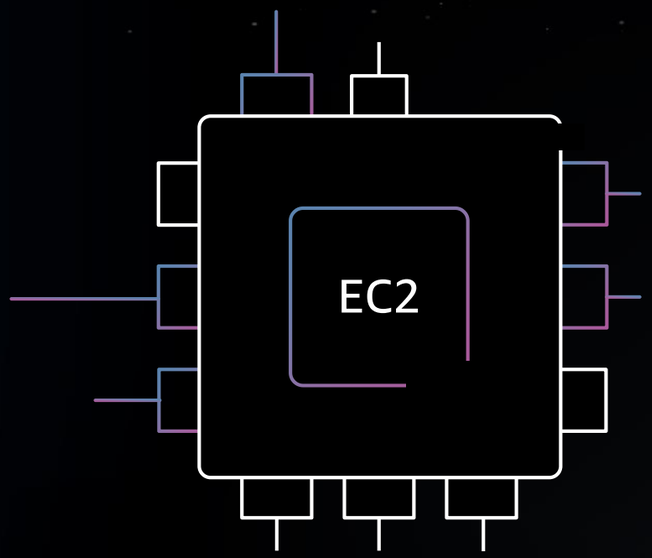
Enabling enterprise applications



Bringing the cloud closer to you

Broadest and deepest platform choice

MOST ADVANCED COMPUTE PLATFORM IN THE CLOUD



Linux and Windows

x86 and Arm architectures

General purpose and workload optimized

Bare metal, disk, networking capabilities

Packaged, custom, and community images (AMIs)

Multiple purchase options: On-Demand, RI, Spot, Savings Plans

Broadest choice of processors and architectures

RIGHT COMPUTE FOR THE WORKLOAD



INTEL ® Xeon Scalable (Skylake and Cascade Lake) processors



AMD EPYC processor



AWS Graviton Processor

Broadest choice of processors and architectures

RIGHT COMPUTE FOR THE WORKLOAD



INTEL ® Xeon Scalable (Skylake and Cascade Lake) processors

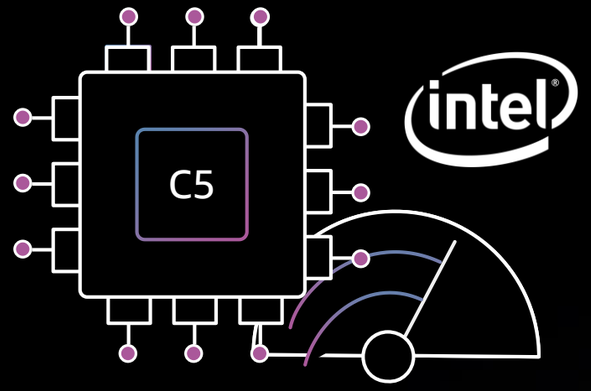


AMD EPYC processor



AWS Graviton Processor

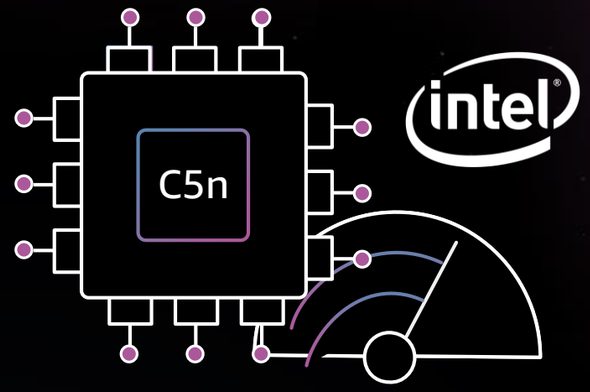
Maximum performance with optimized instances



HIGH-PERFORMANCE COMPUTE

High performance at a low price per vCPU ratio

2:1 memory to vCPU ratio
9 sizes with 2-96 vCPU



HIGH-PERFORMANCE NETWORKING

High-performance compute with 100 Gbps network bandwidth

Elastic Fabric Adapter (EFA) reduces network latencies to 15 µs

intel

POWERED BY INTEL® XEON® SCALABLE (SKYLAKE) PROCESSOR

Thunder from the cloud

A DATA REFINERY, BUILT TO UNDERSTAND OUR PLANET

Descartes Labs

Descartes Labs ran a 40,000 core High Performance Computing (HPC) cluster on AWS

Ranked 179th in TOP500 supercomputers

Only cost $5,000 to complete



Thunder from the cloud

A DATA REFINERY, BUILT TO UNDERSTAND OUR PLANET

"… WE DIDN'T ASK AMAZON TO GIVE OUR ENGINEERS ANY SPECIAL DISPENSATION, DISCOUNT,OR CUSTOM PLANNING OR SETUP. WE WANTED TO SEE IF WE COULD DO THIS ON OUR OWN …"



Descartes Labs

MIKE WARREN, FOUNDER & CTO, DESCARTES LAB

Descartes Labs ran a 40,000 core High Performance Computing (HPC) cluster on AWS

Ranked 179th in TOP500 supercomputers

Did not let AWS know beforehand

Expediting research

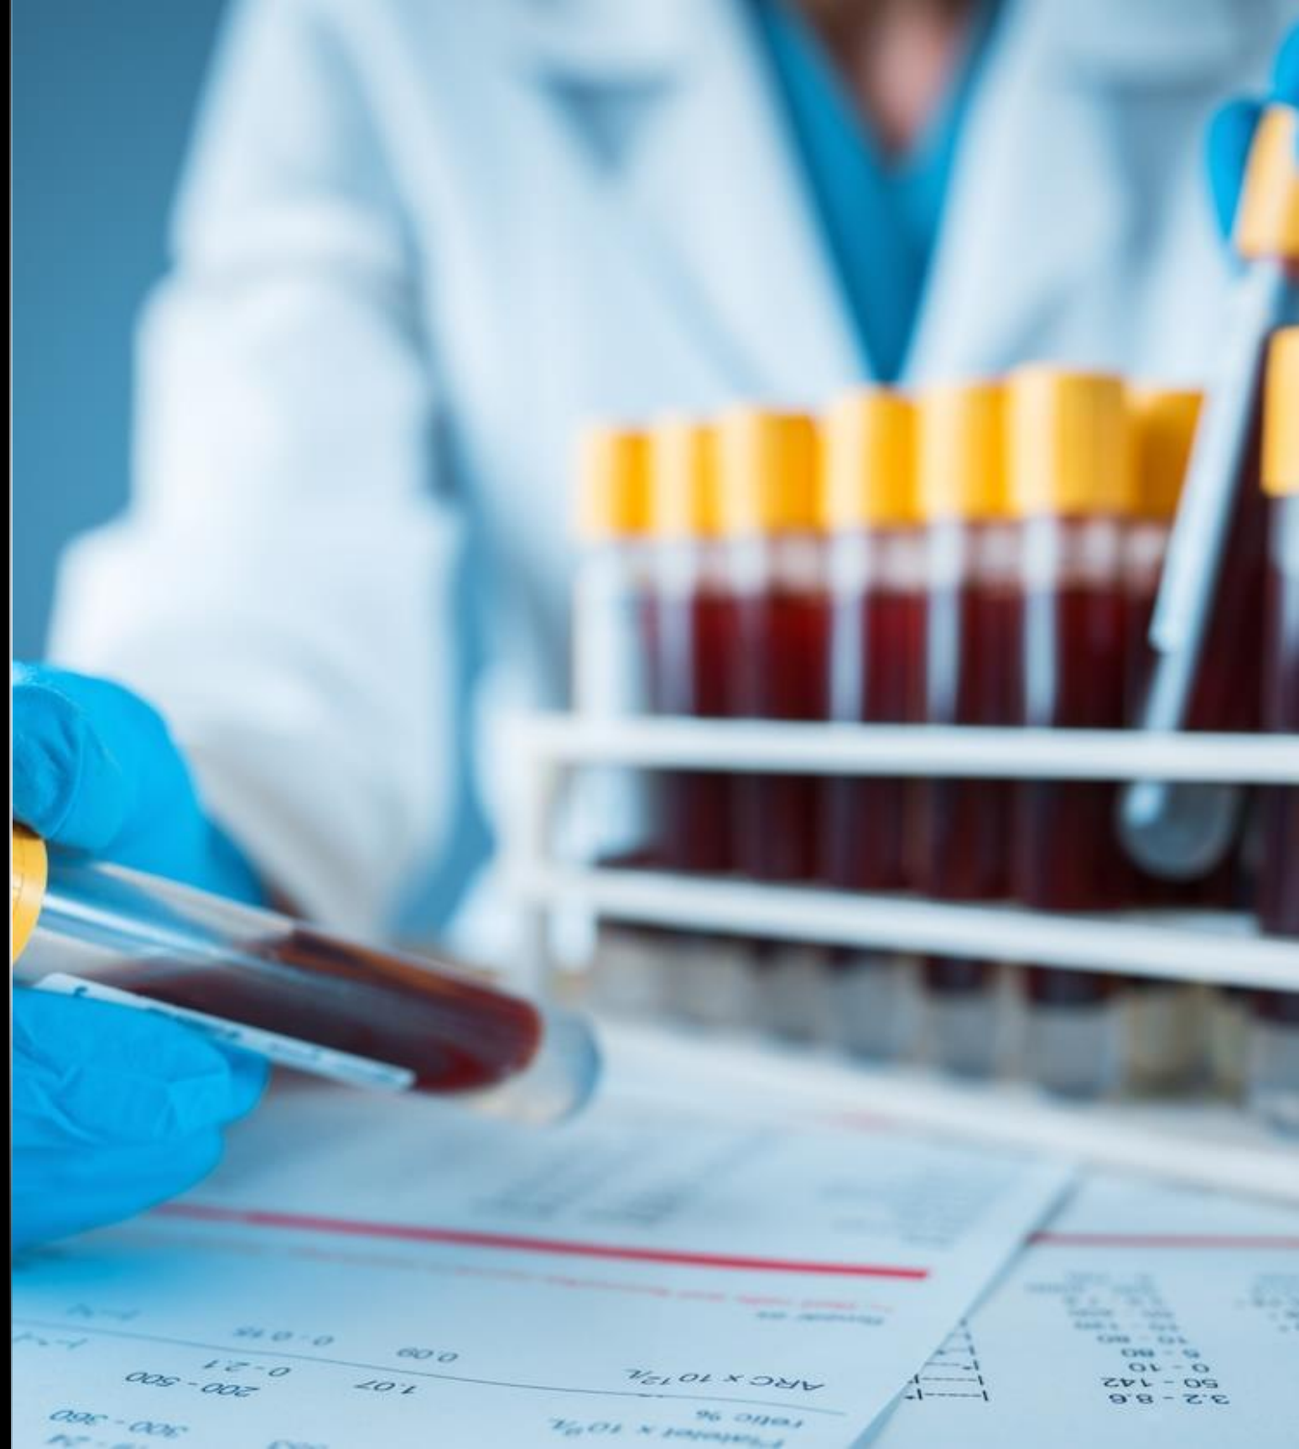
DEDICATED TO THE ELIMINATION OF CANCER AND RELATED DISEASES

FRED HUTCH
CURES START HERE™

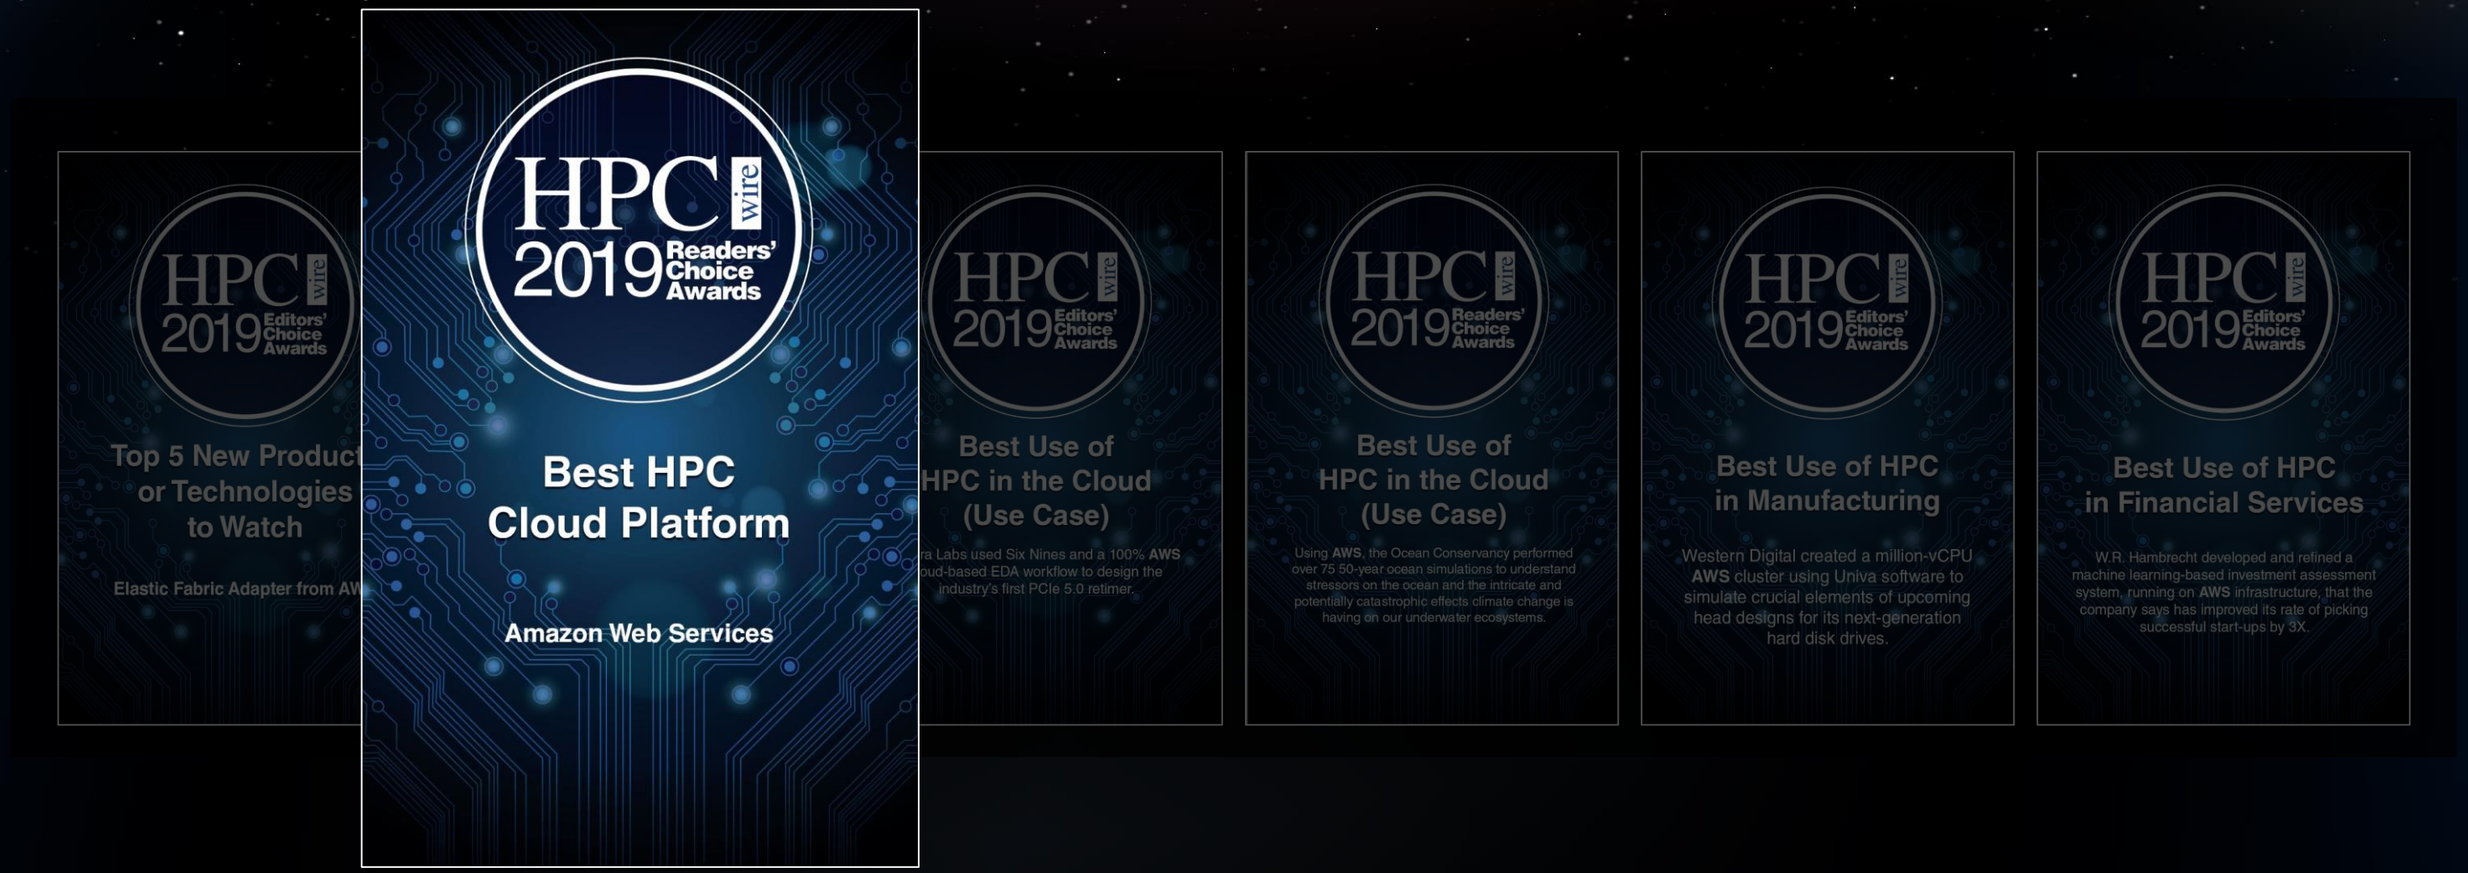
Fred Hutch decreased the time it took to analyze 10,000 biological samples from 7 years to 7 days!

“By accelerating our research processes on AWS, we can get closer to developing therapeutics to fight cancer.”

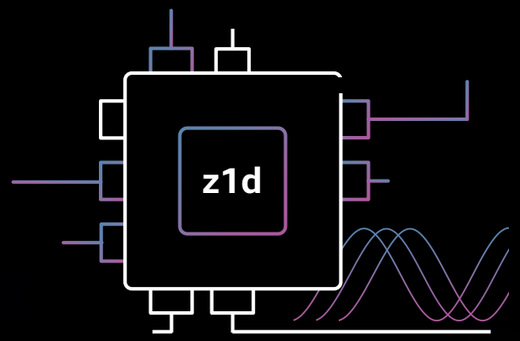
Sam Minot
PhD and Staff Scientist



HPC Wire: Best HPC Cloud Platform

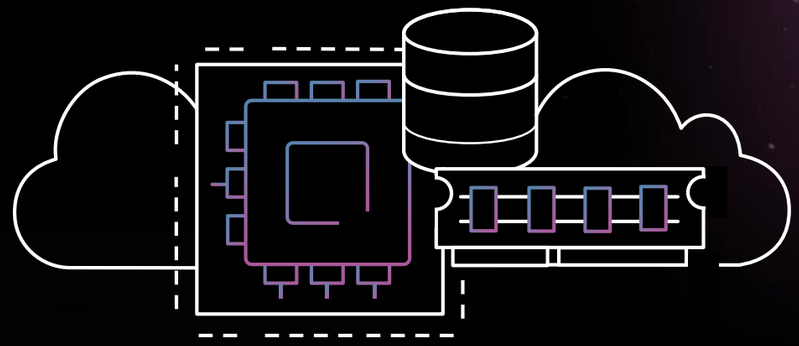


Maximum performance with optimized instances



HIGH-FREQUENCY

4.0 GHz sustained core frequency



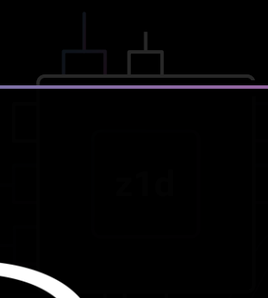
HIGH-MEMORY

6, 9, 12, 18 & 24 TiB memory

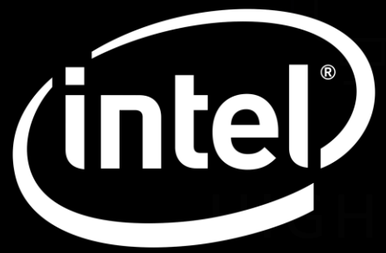
Fastest and largest SAP-certified in-memory instances

POWERED BY INTEL® XEON® SCALABLE (SKYLAKE) PROCESSOR

Maximum performance with optimized instances







14 YEARS OF PARTNERSHIP WITH
OVER 200 INSTANCE TYPES



HIGH-FREQUENCY

4.0 GHz sustained core frequency

HIGH-MEMORY

6, 9, 12, 18 & 24 TiB memory

Fastest and largest SAP-certified
in-memory instances



POWERED BY INTEL® XEON® SCALABLE (SKYLAKE) PROCESSOR

Broadest choice of processors and architectures

RIGHT COMPUTE FOR THE WORKLOAD



INTEL ® Xeon Scalable (Skylake and Cascade Lake) processors

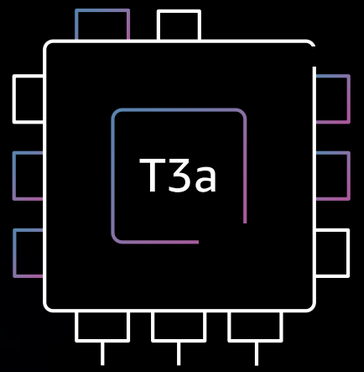


AMD EPYC processor



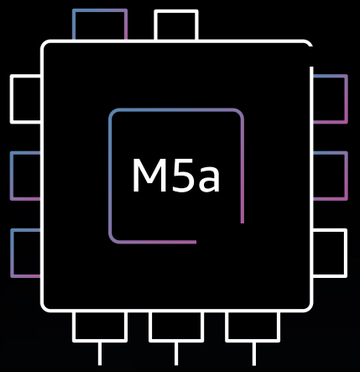
AWS Graviton Processor

Cost savings with AMD instances



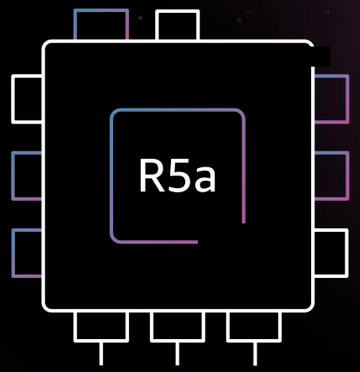
BURSTABLE

Baseline level of CPU performance with the ability to burst CPU usage at any time for as long as required



GENERAL PURPOSE

Balanced compute, memory, and networking resources for general purpose workloads



MEMORY OPTIMIZED

Fast performance for workloads that process large datasets in memory

AMD

POWERED BY AMD EPYC PROCESSOR

Cost savings with AMD instances



T3a

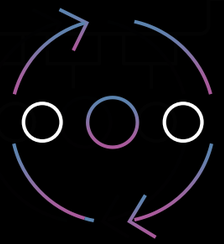


M5a



R5a

10% COST SAVINGS ON COMPUTE AND MEMORY



SEAMLESS APPLICATION MIGRATION

BURSTIBLE

Baseline level of CPU performance with the ability to burst CPU usage at any time for as long as required

GENERAL PURPOSE

Balanced compute, memory, and networking resources for general purpose workloads

MEMORY OPTIMIZED

Fast performance for workloads that process large data sets in memory

AMD

POWERED BY AMD EPYC PROCESSOR

AMD for microservices

TEN PERCENT REDUCTION ON COMPUTE COSTS

Expedia®

Expedia has migrated a number of microservices to AMD, where they were underutilizing compute

The 10% lower cost, application compatibility, and familiar instances were the reasons they decided to use AMD instances

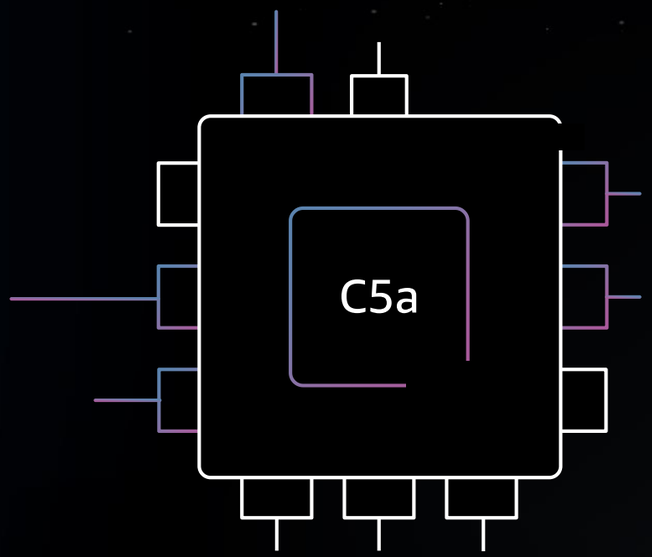


NEW!

COMING SOON

AMD Compute Optimized Instances

POWERED BY SECOND-GENERATION EPYC “ROME” PROCESSOR



Available in C5a and C5ad instances

8 instance sizes as well as bare metal

Up to 192 vCPUs, 384 GiB memory, 7.6 TiB of fast, local NVMe storage and 100 Gbps networking

Perfect for applications that require a lower vCPU-to-memory ratio

Broadest choice of processors and architectures

RIGHT COMPUTE FOR THE WORKLOAD



INTEL ® Xeon Scalable (Skylake and Cascade Lake) processors



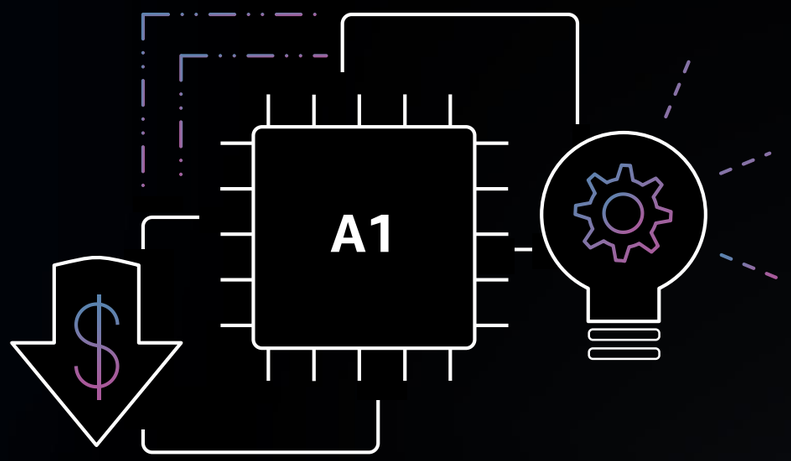
AMD EPYC processor



AWS Graviton Processor

A1 instance

POWERED BY AWS GRAVITON PROCESSOR



First Arm-based processor in major cloud provider

Built on 64-bit Arm Neoverse cores with AWS-designed 16 nm silicon

Up to 16 vCPUs, 10 Gbps enhanced networking, and 3.5 Gbps Amazon EBS bandwidth

Up to 40% cost savings with higher price performance

“THESE INSTANCES WILL ALSO APPEAL TO DEVELOPERS, ENTHUSIASTS, AND EDUCATORS ACROSS THE ARM DEVELOPER COMMUNITY”

AWS DOCUMENTATION

Igniting the Arm ecosystem

OPERATING SYSTEMS

Amazon Linux 2

Ubuntu 16.04, 18.04 and newer

RedHat Enterprise Linux 7.6 and 8.0

SUSE Linux Enterprise Server for Arm 15

Docker Desktop Community and Docker Enterprise Engine

CONTAINERS

Most Docker official images support arm64

Amazon ECS
AVAILABLE TODAY

Amazon EKS
PUBLIC PREVIEW

INTEGRATED SERVICES

AWS Marketplace

AWS Systems Manager

AWS CloudWatch

AWS CodeCommit

AWS Cloud9

AWS CodePipeline

Amazon Inspector

Amazon Batch

A1 for Big Data

EVOLVING INTO A MODERN BIG DATA COMPANY WITH A1

nielsen

Nielsen processes more than 12 billion batch jobs daily

By migrating to A1 instances, they experienced significant cost savings

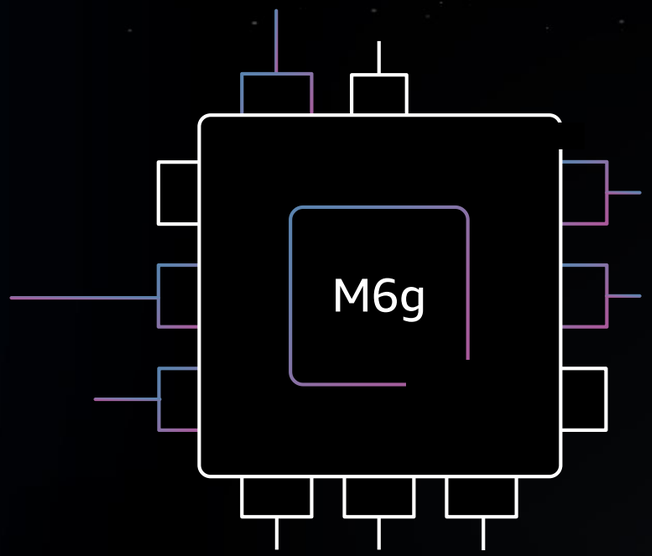


SO, COULD WE DO MORE?

NEW!

M6g, R6g, C6g instances

POWERED BY AWS GRAVITON2 PROCESSOR



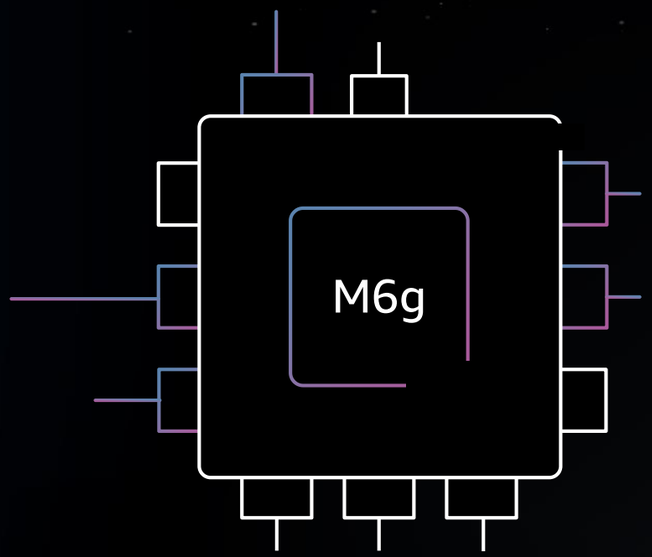
New generation of Arm-based instances powered by AWS Graviton2 processors offers 40% better price-performance than current x86-based instances

M6g	**C6g**	**R6g**
PREVIEW TODAY	COMING SOON	COMING SOON

NEW!

M6g, R6g, C6g Instances

POWERED BY AWS GRAVITON2 PROCESSOR



Customized 64-bit Neoverse cores with AWS-designed 7 nm silicon

Up to 64 vCPUs

25 Gbps enhanced networking

18 Gbps EBS bandwidth

40% price performance advantage over x86 5th generation instances

SO JUST HOW GOOD ARE THESE NEW GRAVITON2 EC2 INSTANCES?

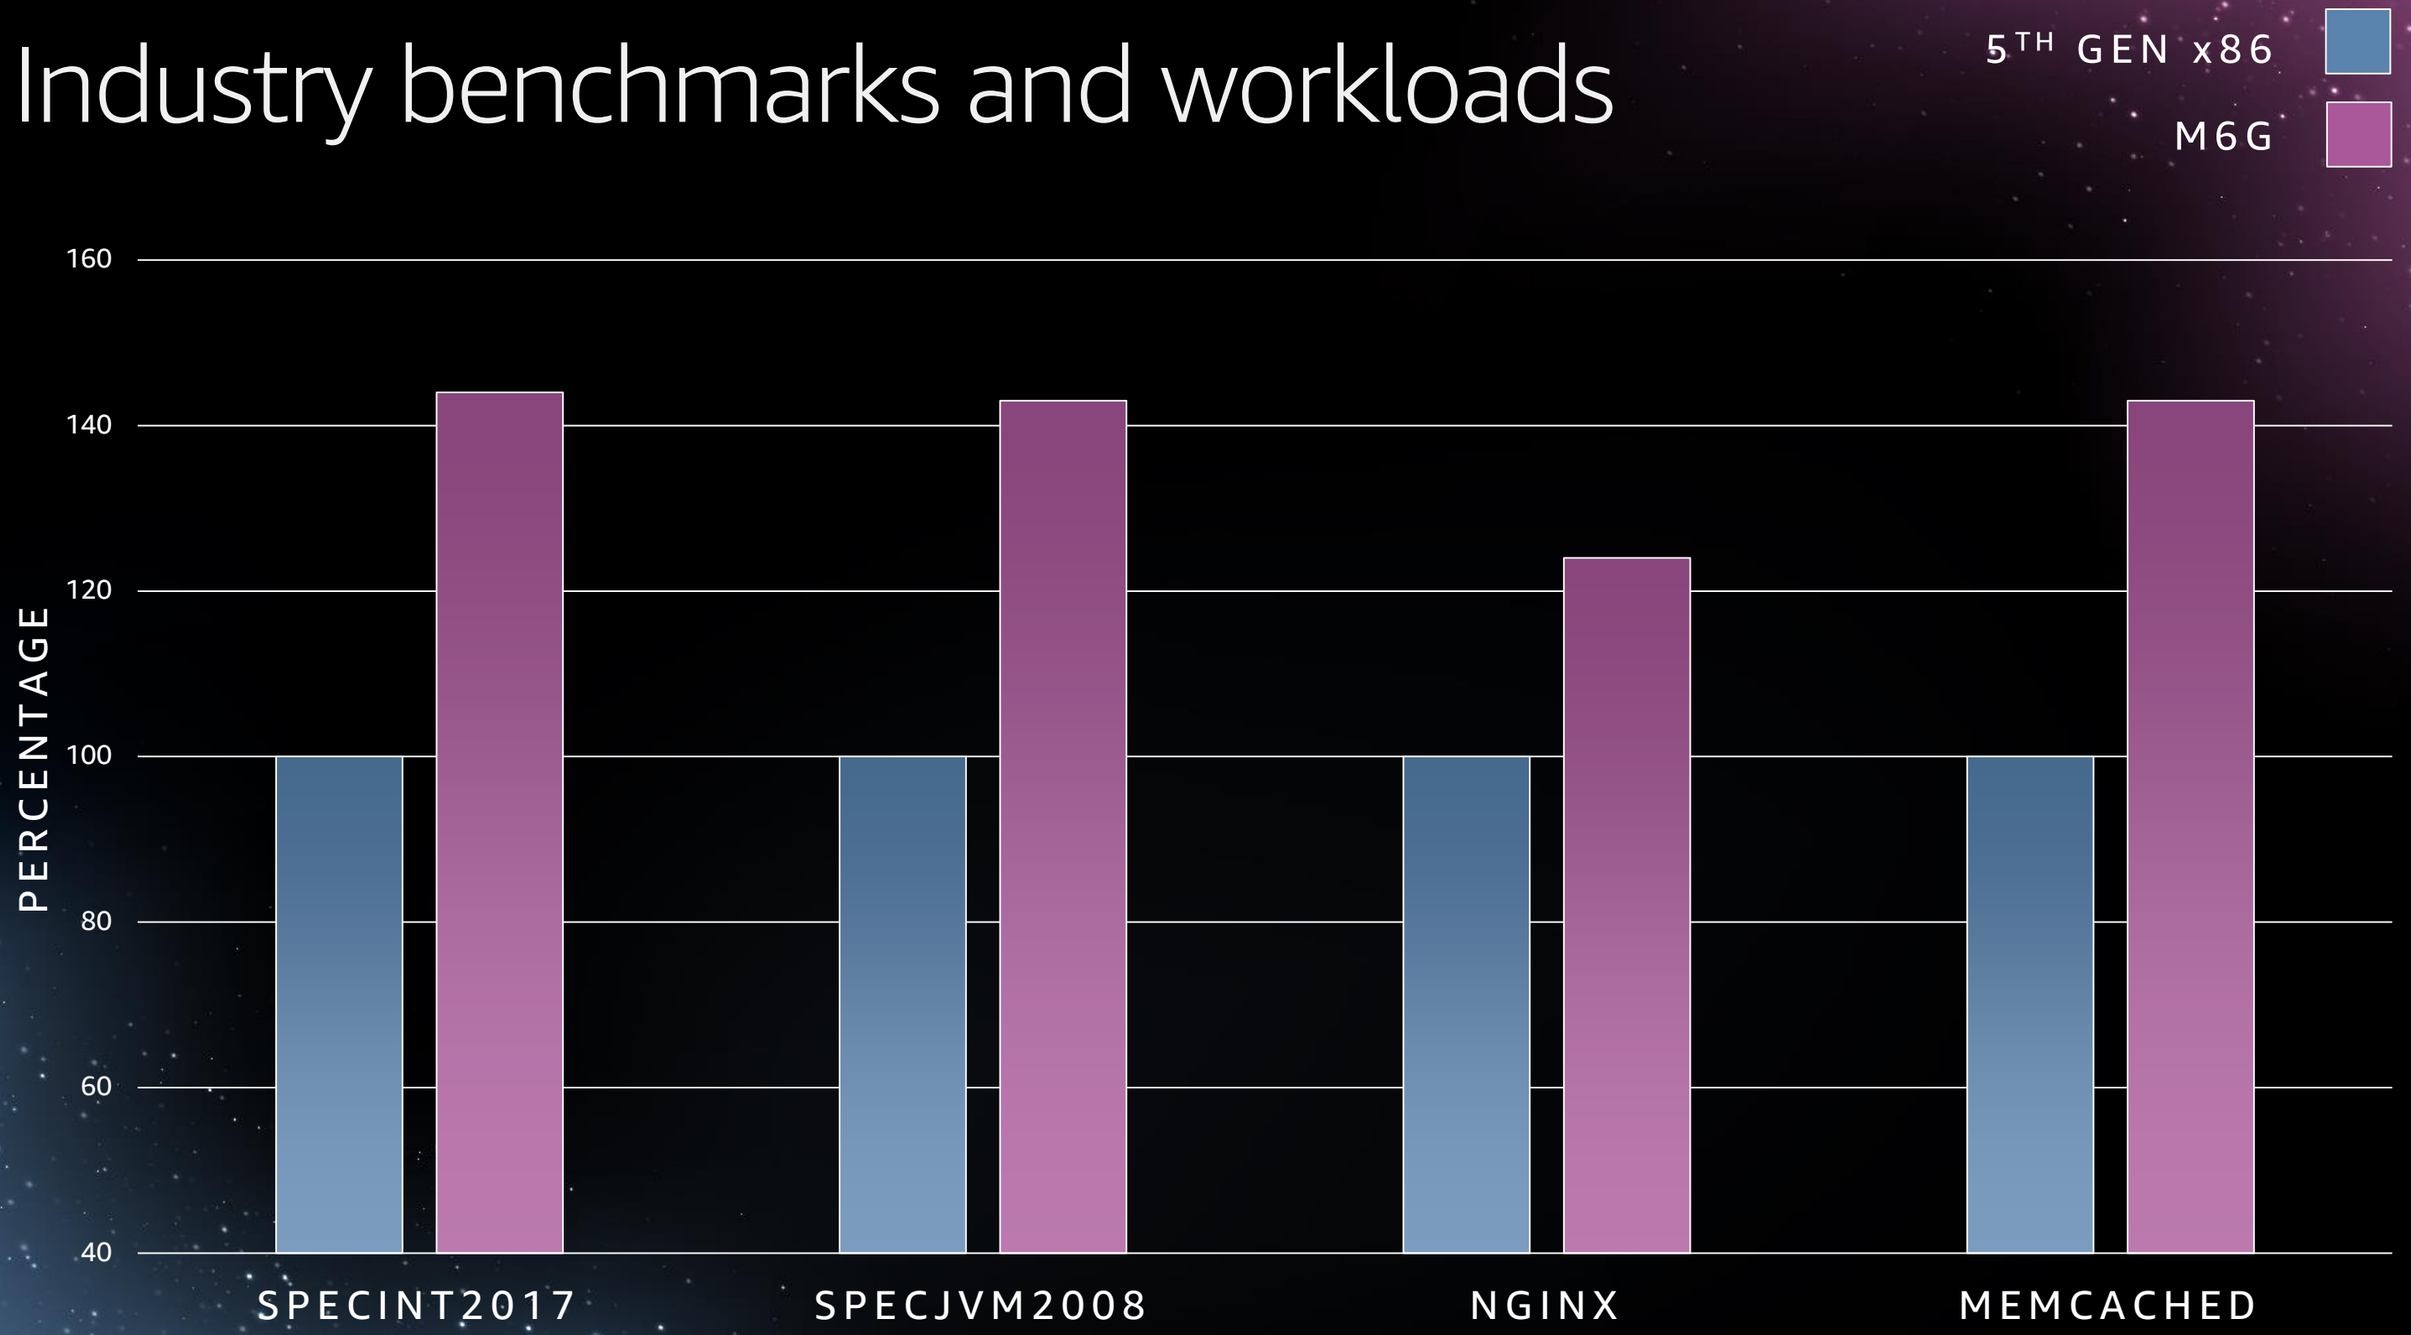
Comparing AWS Graviton2 to AWS Graviton

4x COMPUTE CORES

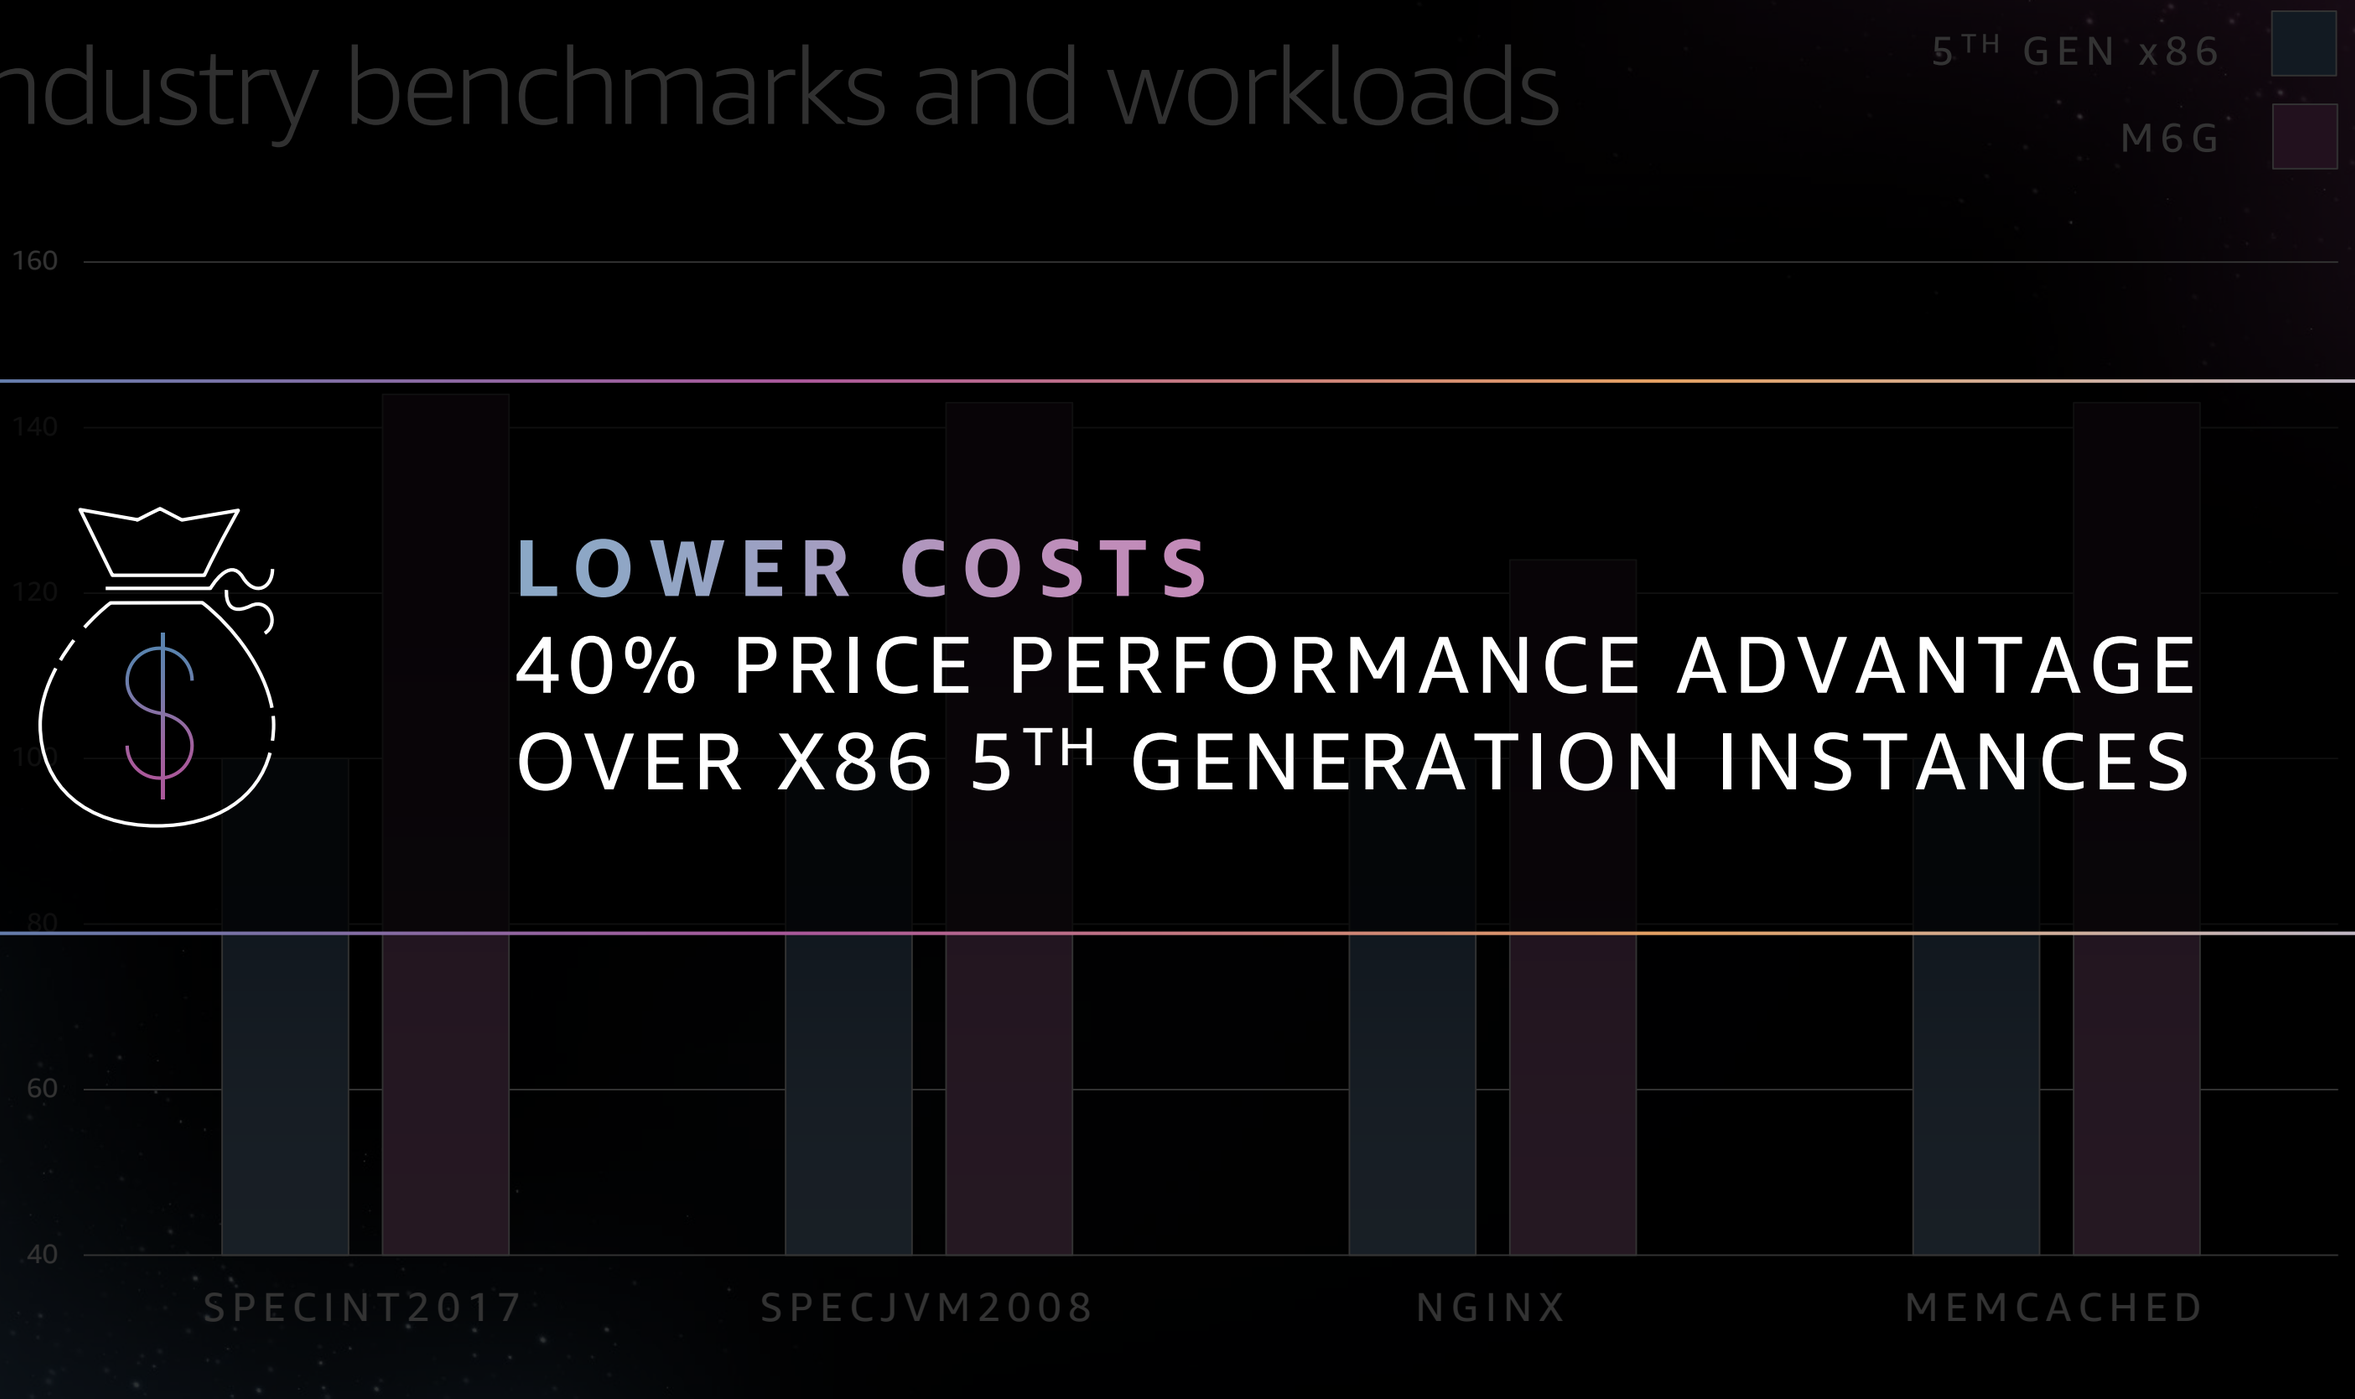
5x FASTER MEMORY

7x PERFORMANCE

Industry benchmarks and workloads



Industry benchmarks and workloads



LOWER COSTS

40% PRICE PERFORMANCE ADVANTAGE OVER X86 5TH GENERATION INSTANCES

Broadest choice of processors and architectures

RIGHT COMPUTE FOR THE WORKLOAD



INTEL ® Xeon Scalable (Skylake and Cascade Lake) processors



AMD EPYC processor



AWS Graviton Processor

Choice of accelerators for specialized workloads

G4

Cost-effective compute in the cloud for running GPU-enabled ML inference and graphics-intensive applications

P3dn

The fastest GPU instance available in the cloud for machine learning training

+

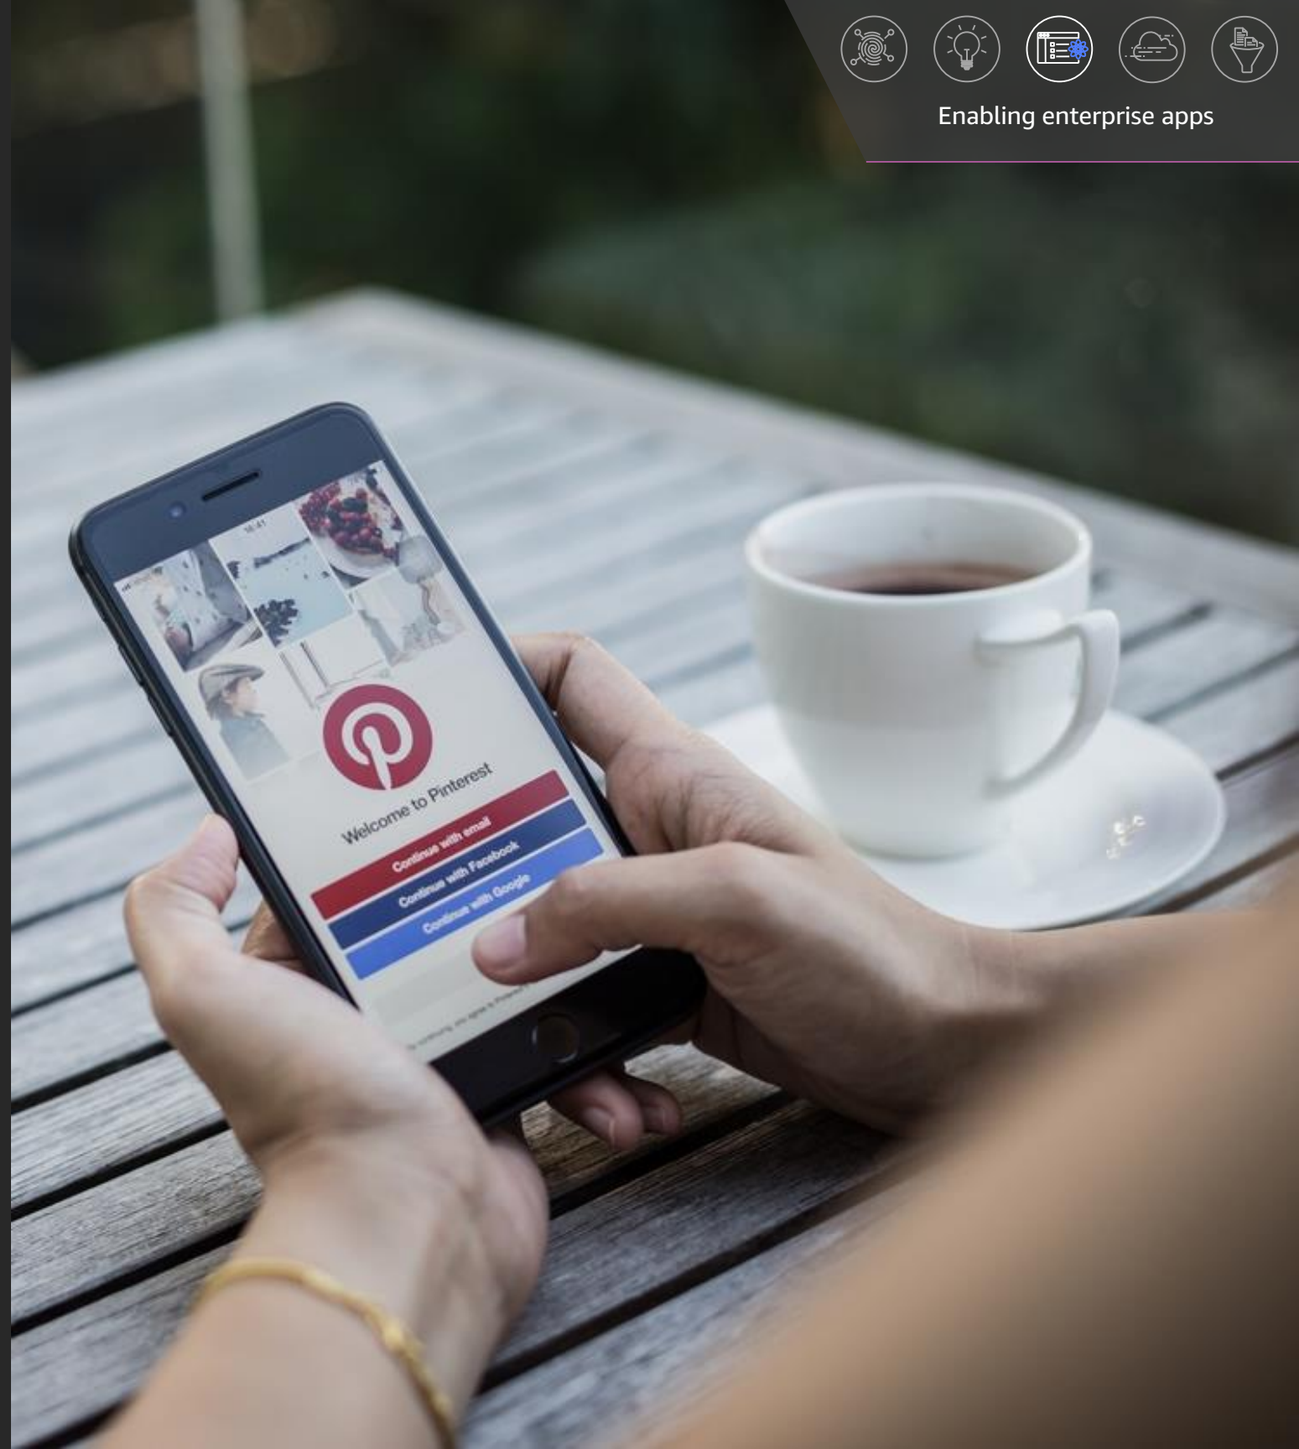
Elastic Inference

Attach low-cost GPU-powered acceleration

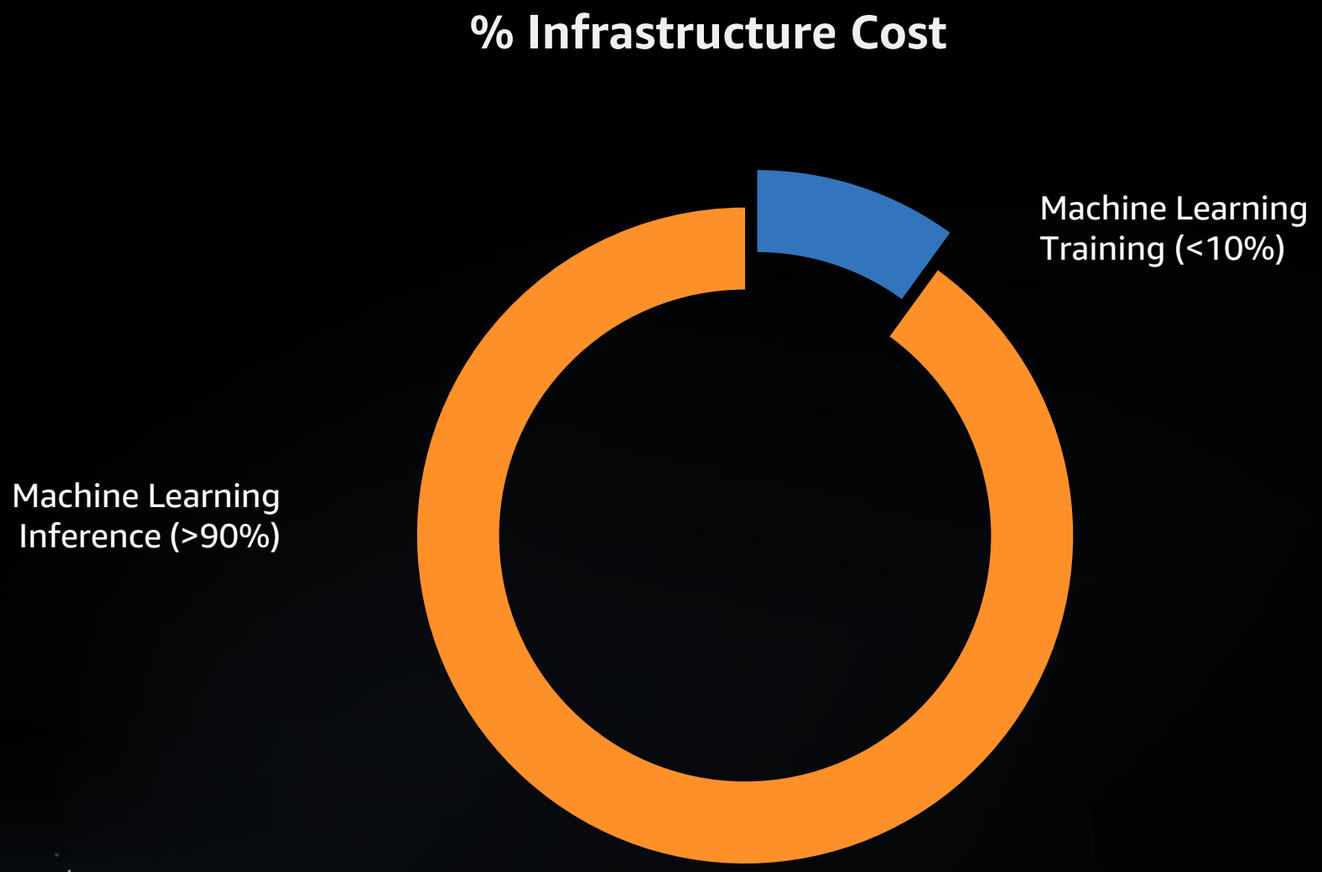


Using AI to enhance personalized experiences

Pinterest uses P3 instances to accelerate deep learning models, which enables fast and unique discovery experience for its >300 million active users



Inference account for majority of machine learning Infrastructure costs

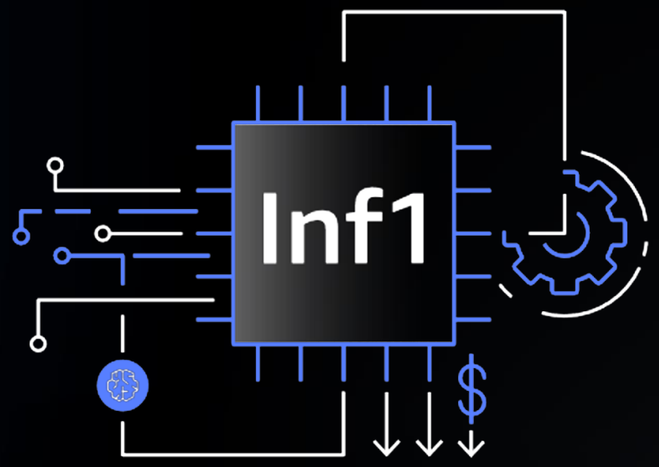


NEW!

AVAILABLE TODAY

Introducing Inf1 instances

THE FASTEST INFERENCE IN THE CLOUD



10x lower inference latency, 2x higher inference throughput than anything in the cloud

Up to 2,000 TOPS at sub-millisecond latency

2.5x lower cost-per-inference than any Amazon EC2 GPU instance

Integrated with popular ML frameworks like TensorFlow, PyTorch, and MXNET

BUILT TO SUPPORT VIRTUALLY EVERY WORKLOAD



World-class performance, security, and innovation



Broadest and deepest platform choice



Abstractions for managing cost and complexity



Enabling enterprise applications



Bringing the cloud closer to you

TOO MUCH CHOICE CAN ADD COMPLEXITY

Broadest and deepest platform choice

CATEGORIES

CAPABILITIES

OPTIONS

GPU compute

Graphics intensive

(up to 100 Gbps)

Bare Metal

Size
(Nano to 32xlarge)

business need

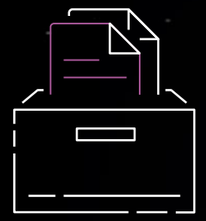
GREAT! BUT WHICH INSTANCE SHOULD I CHOOSE FOR MY WORKLOAD?

NEW!

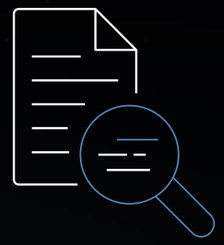
LAUNCHED IN NOVEMBER

Introducing EC2 Instance Discovery

NEW SEARCH AND DISCOVERY EXPERIENCE FOR LOCATING INSTANCE TYPES



Single source of truth for the latest instance types, attributes, regional and zonal offerings, and pricing



Compare the hardware attributes, pricing, and availability of each instance type to select the best option



Programmatic access to instance type information with two new APIs

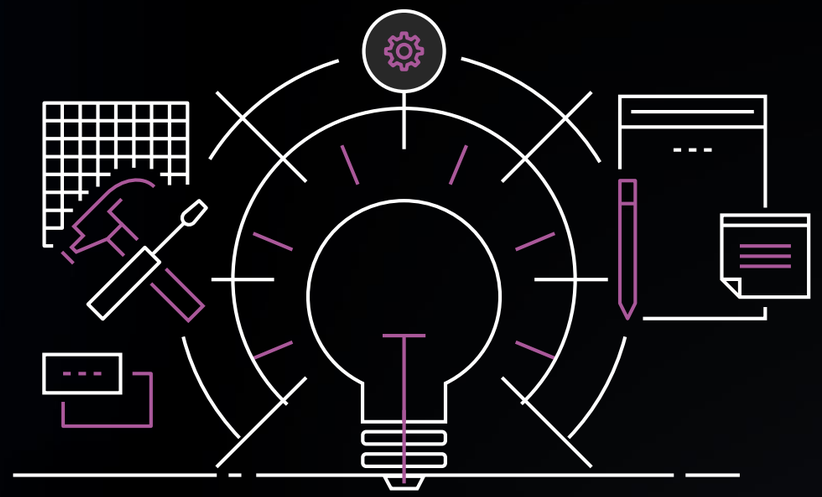
WHAT ABOUT OPTIMIZING INSTANCE SELECTION FOR RUNNING INSTANCES?

NEW!

AVAILABLE TODAY

Introducing AWS Compute Optimizer

DELIVERS INSTANCE TYPE AND AUTO SCALING GROUP RECOMMENDATIONS



Applies insights from millions of workloads to make recommendations

Saves time comparing and selecting optimal compute resources for your workload

Recommends 3 optimal instance options for EC2 and EC2 Auto Scaling groups from 140+ instances from M, C, R, T, and X families

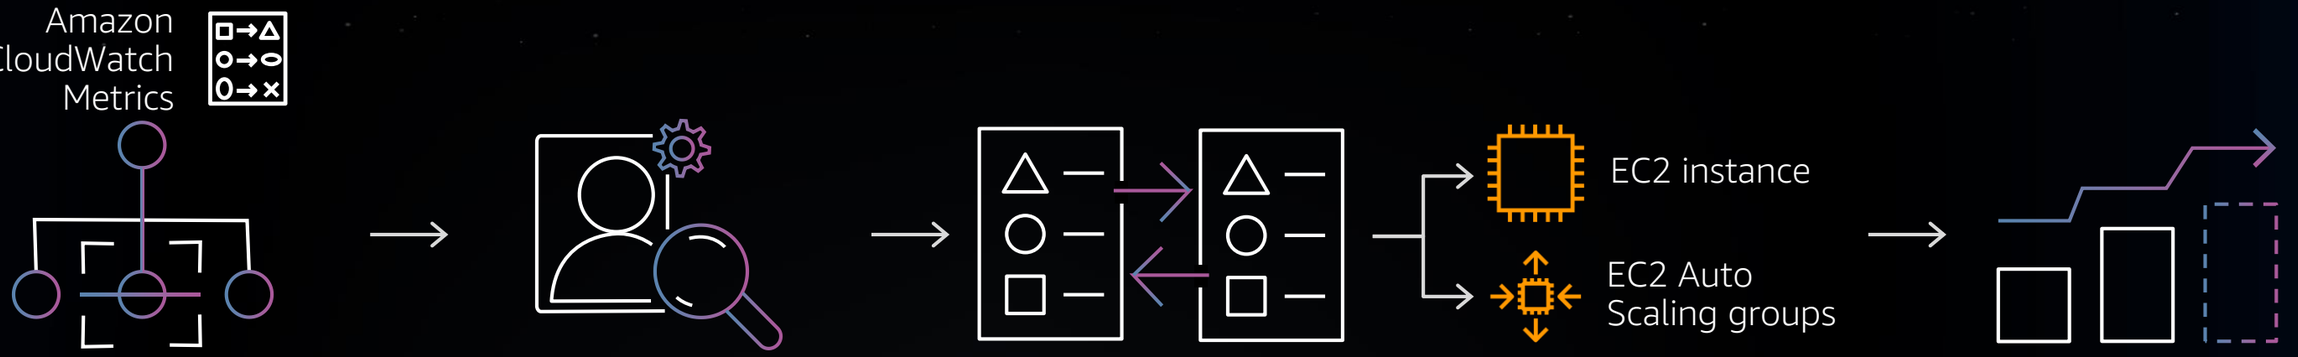
Available at no additional charge

NEW!

AVAILABLE TODAY

Introducing AWS Compute Optimizer

DELIVERS INSTANCE TYPE AND AUTO SCALING GROUP RECOMMENDATIONS



Scans your AWS infrastructure and uses machine learning to automatically identify optimal AWS resources for your workloads

Identifies workload characteristics and profile based on the data gathered

Matches the resource requirements of your workloads to optimal AWS compute resources

Helps you visualize what-if scenarios based on the recommended resources

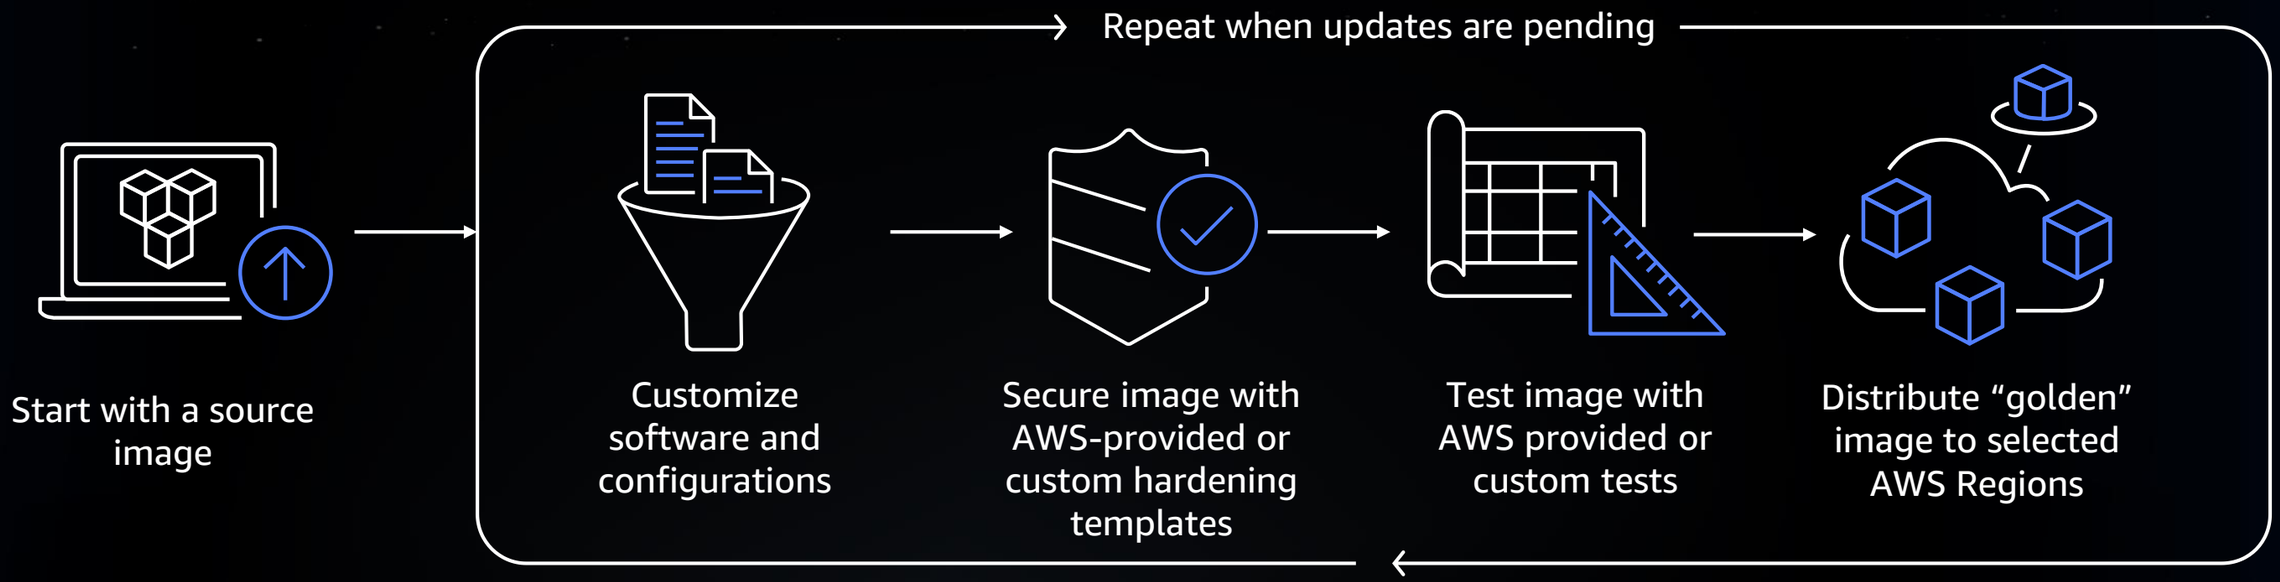
AMI … WHETHER IT'S "AH-MEE" or "AY-EM-I," WE NEED AN EASIER WAY TO BUILD IMAGES

NEW!

AVAILABLE TODAY

Introducing EC2 Image Builder

SIMPLIFIES THE CREATION, MAINTENANCE, VALIDATION, SHARING, AND DEPLOYMENT OF LINUX OR WINDOWS SERVER IMAGES

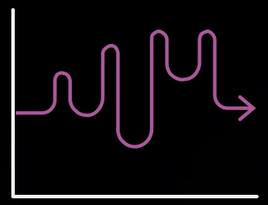


WE’RE ALWAYS LOOKING FOR WAYS TO REDUCE COSTS FOR CUSTOMERS

Ways we make it easy to reduce costs …

ON-DEMAND

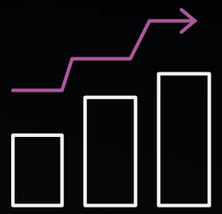
Pay for compute capacity by the second with no long-term commitments



Spiky workloads, to define needs

RESERVED INSTANCES

Make a 1- or 3-year commitment and receive a significant discount off On-Demand prices



Committed and steady-state usage

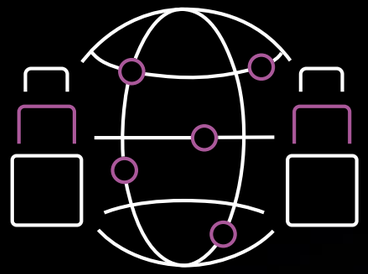
SPOT INSTANCES

Spare EC2 capacity at savings of up to 90% off On-Demand prices



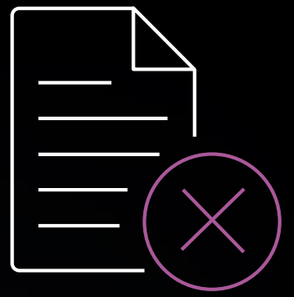
Fault-tolerant, flexible, stateless workloads

Still, needs to be simpler …



IMPROVED SCALE

Need an **easier way to manage commitments at scale** across different compute services and 200+ instance options across AWS 22+ Regions



IMPROVED FLEXIBILITY

Need a commitment option that **supports changes** to compute usage **without management overhead**

NEW!

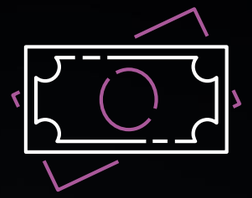
LAUNCHED IN NOVEMBER

Introducing AWS Savings Plans

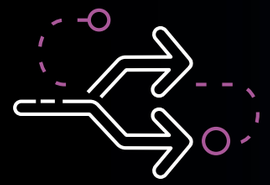
SAME GREAT SAVINGS AS RESERVED INSTANCES WITH MORE FLEXIBILITY



Easy to use with recommendations in AWS Cost Explorer

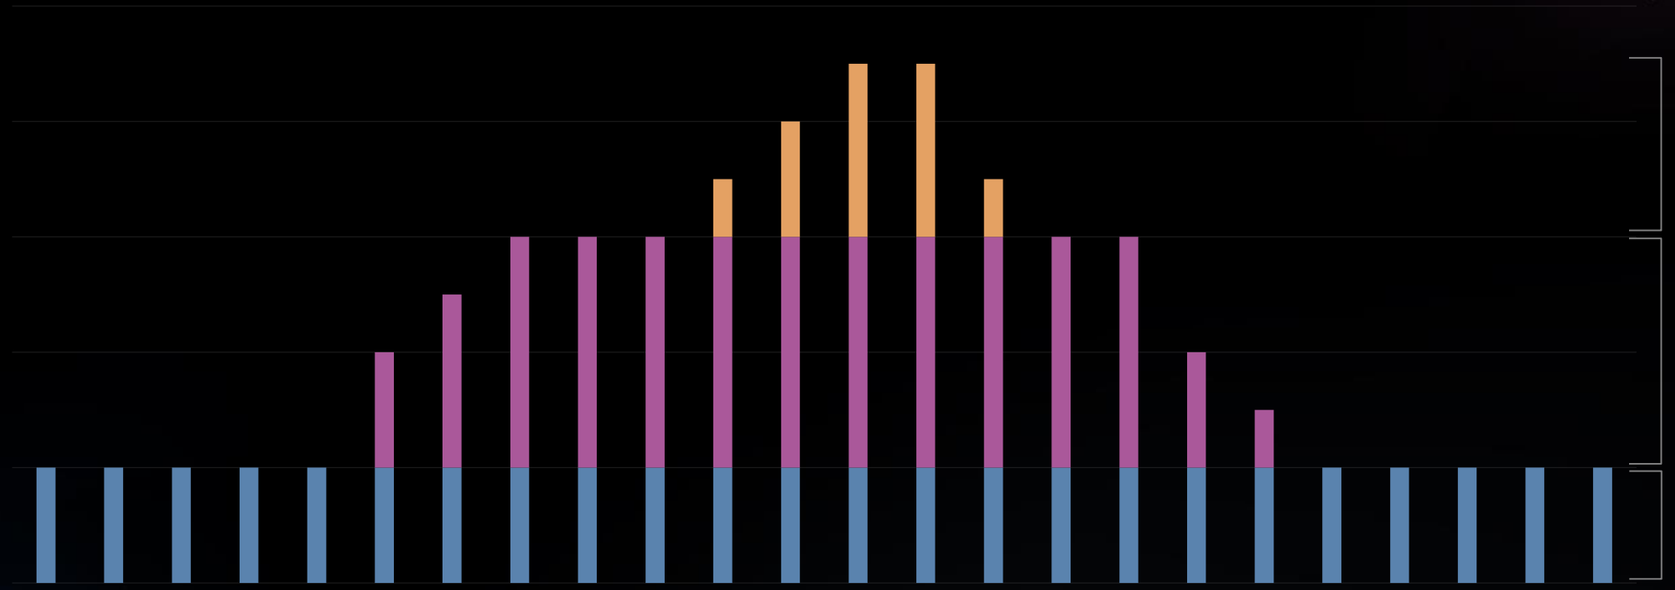


Significant savings of up to 72%



Flexible across instance family, size, OS, tenancy, or AWS Region; also applies to AWS Fargate usage

Simplifying capacity & cost optimization



SCALE USING **SPOT** FOR FLEXIBLE, FAULT-TOLERANT WORKLOADS

SCALE USING **ON-DEMAND** FOR NEW OR STATEFUL SPIKY WORKLOADS

USE **RIs** AND **SAVINGS PLANS** FOR KNOWN/STEADY-STATE WORKLOADS

AWS SERVICES MAKE THIS EASY AND EFFICIENT

- Amazon EC2 Auto Scaling
- EC2 Fleet
- Amazon Elastic Container Service (Amazon ECS)
- Amazon Elastic Kubernetes Service (Amazon EKS)
- AWS Thinkbox
- Amazon EMR
- AWS CloudFormation
- AWS Batch

BUILT TO SUPPORT VIRTUALLY EVERY WORKLOAD



World-class performance, security, and innovation



Broadest and deepest platform choice



Abstractions for managing cost and complexity



Enabling enterprise applications



Bringing the cloud closer to you

AWS is the best cloud for Windows

MOST CAPABILITIES

48

AWS services where Azure does not have comparable options

GREATER RELIABILITY

2x

Number of regions with multiple Availability Zones

BETTER PRICE-PERFORMANCE

2–3x

Performance advantage over Azure running SQL Server when comparing price of a workload

GREATER SECURITY

210

Security, compliance, and governance key features, **40** more than Azure

MIGRATION EXPERIENCE

11+

Years of migration experience

LOWER COSTS

442%

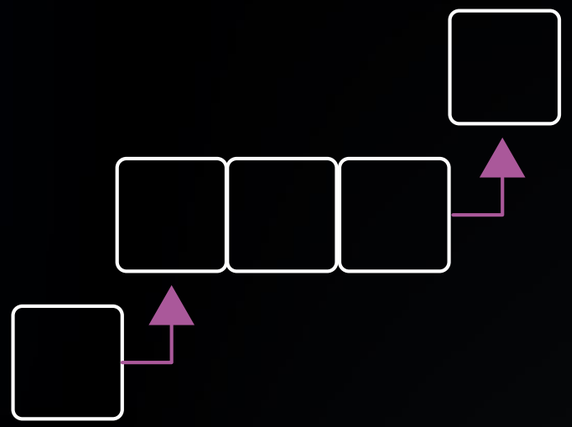
5-year return on investment

NEW!

AVAILABLE TODAY

Introducing Simplified Windows and SQL Server BYOL

NEW HOST MANAGEMENT CAPABILITIES TO SUPPORT EXISTING LICENSES



Dedicated Hosts now allow you to use your existing per-socket, per-core, or per-VM software licenses

Integrated with AWS License Manager, helping you manage your software licenses

Improved visibility into core utilization of Dedicated Host, helping you manage per-socket or per-core software licensing

BYOL on Dedicated Hosts made easier

Get elasticity of EC2 for Dedicated Hosts New



NEW!

AVAILABLE TODAY

AWS End-of-Support Migration Program for Windows

Helps customers upgrade legacy applications to run on newer, supported versions of Windows Server without any code changes

FUTURE-PROOF



Decouple from underlying OS



Low risk of failure on subsequent OS updates



Supports all OS version

REDUCED RISK



Improved security posture on supported, new OS

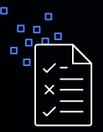


Isolate old runtimes

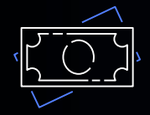


Compliance with industry regulations

COST-EFFECTIVE



No application refactoring or recoding cost



No extended support costs



Reduced operating costs

Modernizing using Windows on AWS



Enabling enterprise apps

3M Science. Applied to Life.™ · accenture · Adobe · arcadia HEALTHCARE SOLUTIONS · ARGO AI · ASU Arizona State University · Attraqt · Barracuda · Biblica

bmc · cardknox · Charter COMMUNICATIONS · cisco · Click Actual Intelligence. At Work. · CLOUDREACH · COMCAST BUSINESS · Dataminr · DAVINCI · Decisiv

dr.consulta o meu médico · DOMO · DOW JONES · Expedia · eMarketer · Earth Networks · gemalto a Thales company · GONG · HERALD LOGIC · DXC.technology

EICHER · HITACHI Inspire the Next · HSBC · infor · FICO · Information Builders · insights · IntelliGuard · Jobvite · KALTURA

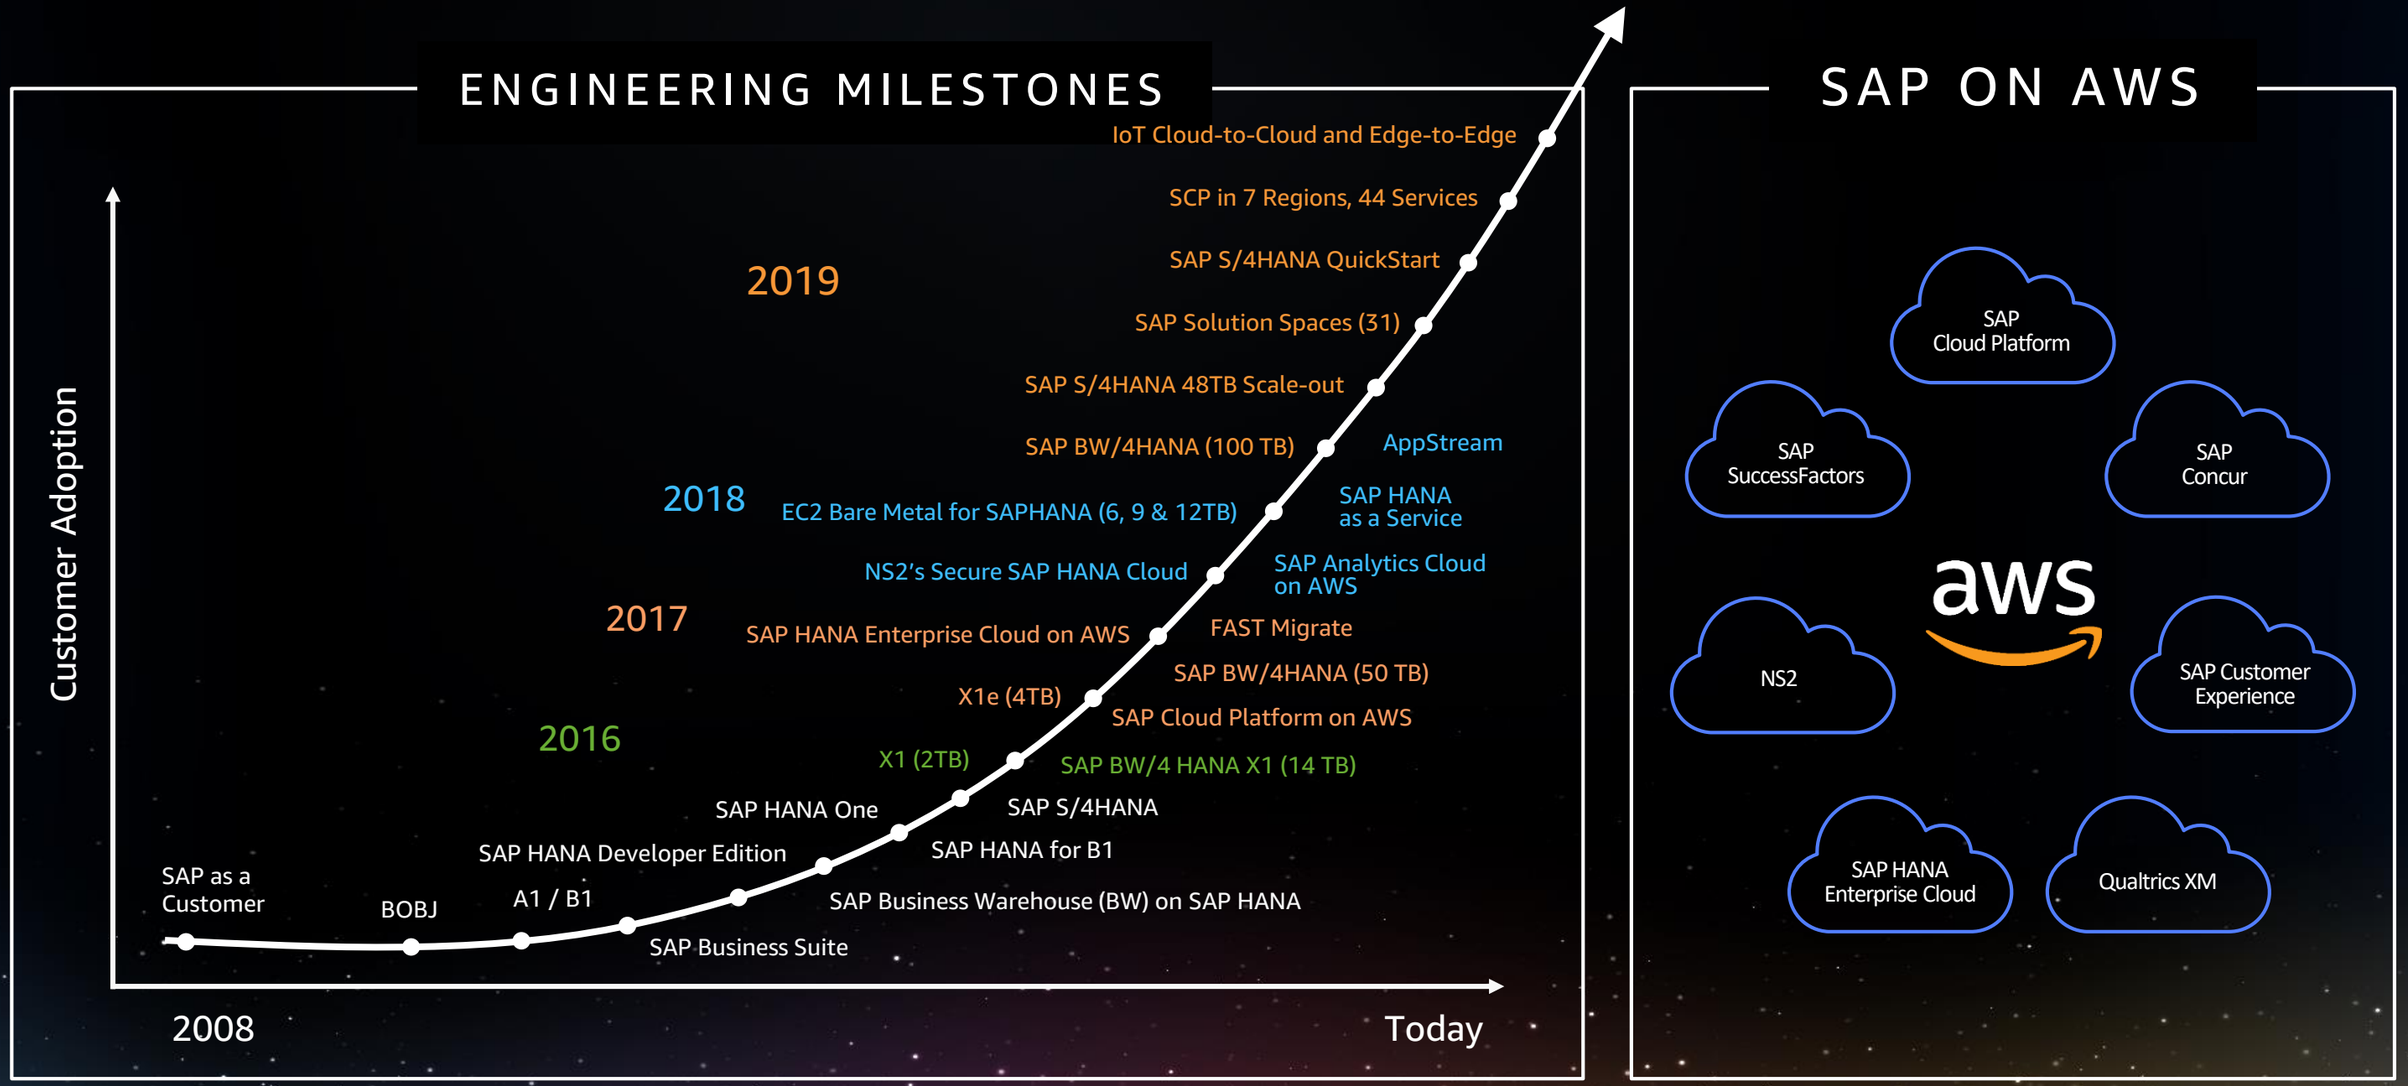
LADYDRIVER · KG Eduone · lucid INTERACTIVE · MATHNASIUM The Math Learning Center · McAfee · MZ Cloud Innovator MEGAZONE · mercado libre · MicroStrategy · MOTOROLA SOLUTIONS · MUFG

MYTEAM11 · NetApp · NEW INNOVATIONS · ohpen · OLA · origin · Pacific magazines We create magazines people love · paloalto NETWORKS · Pearson | eText · playwith INTERACTIVE

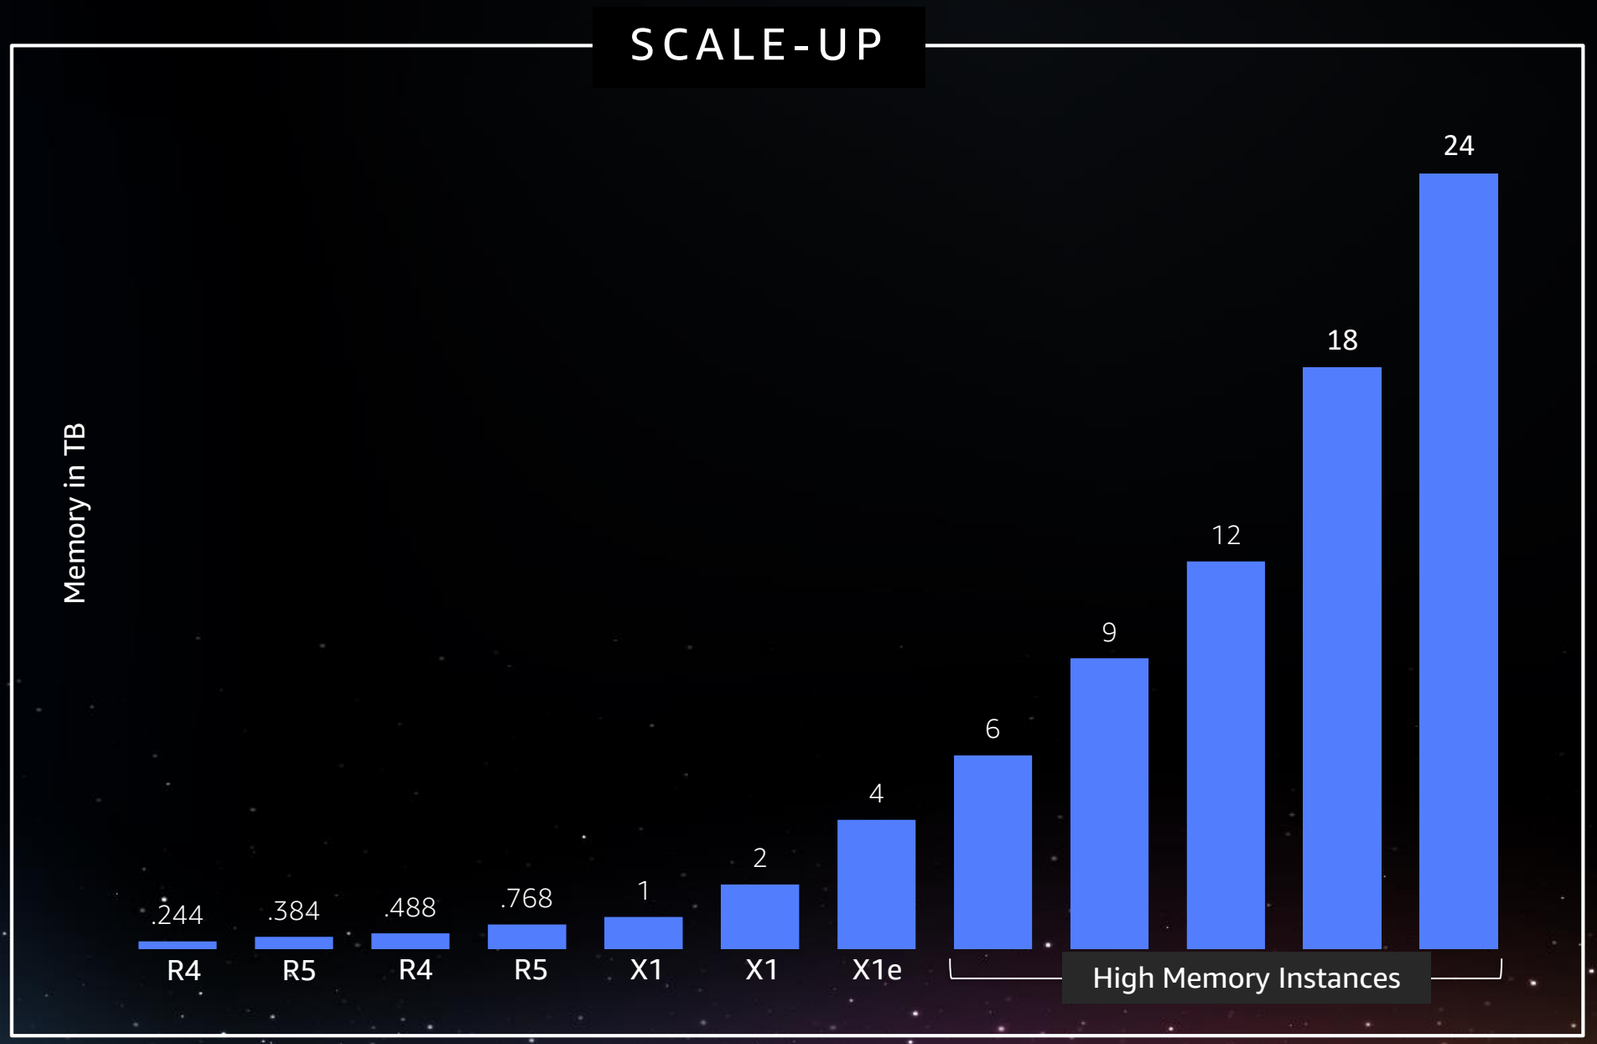
POTTERY BARN OUTLET · Quantcast · Red Hat · RYANAIR · REPAY Realtime Electronic Payments · salesforce · SallieMae Let's Make College Happen · SAMSUNG · SAP · shi

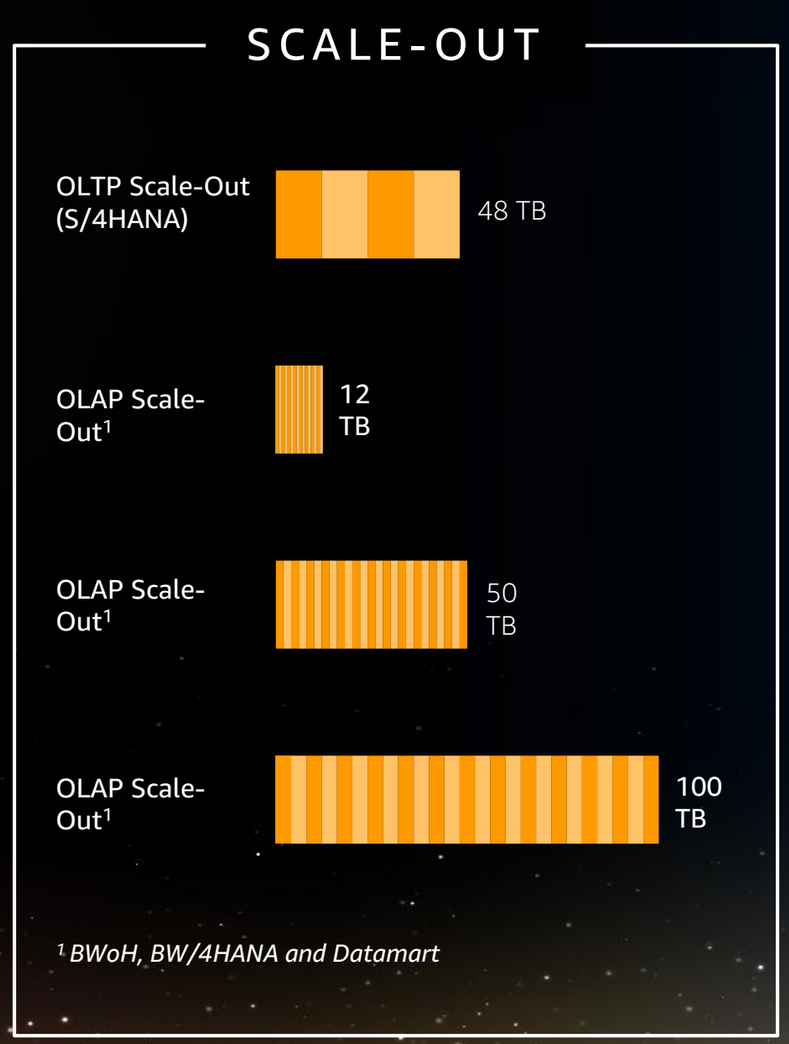
SIEMENS Ingenuity for life · SWIGGY · talend · TOYOTA RESEARCH INSTITUTE · Verisk 3E · VERITAS · xero · verizon · Zillow

SAP on AWS: Unmatched pace of innovation



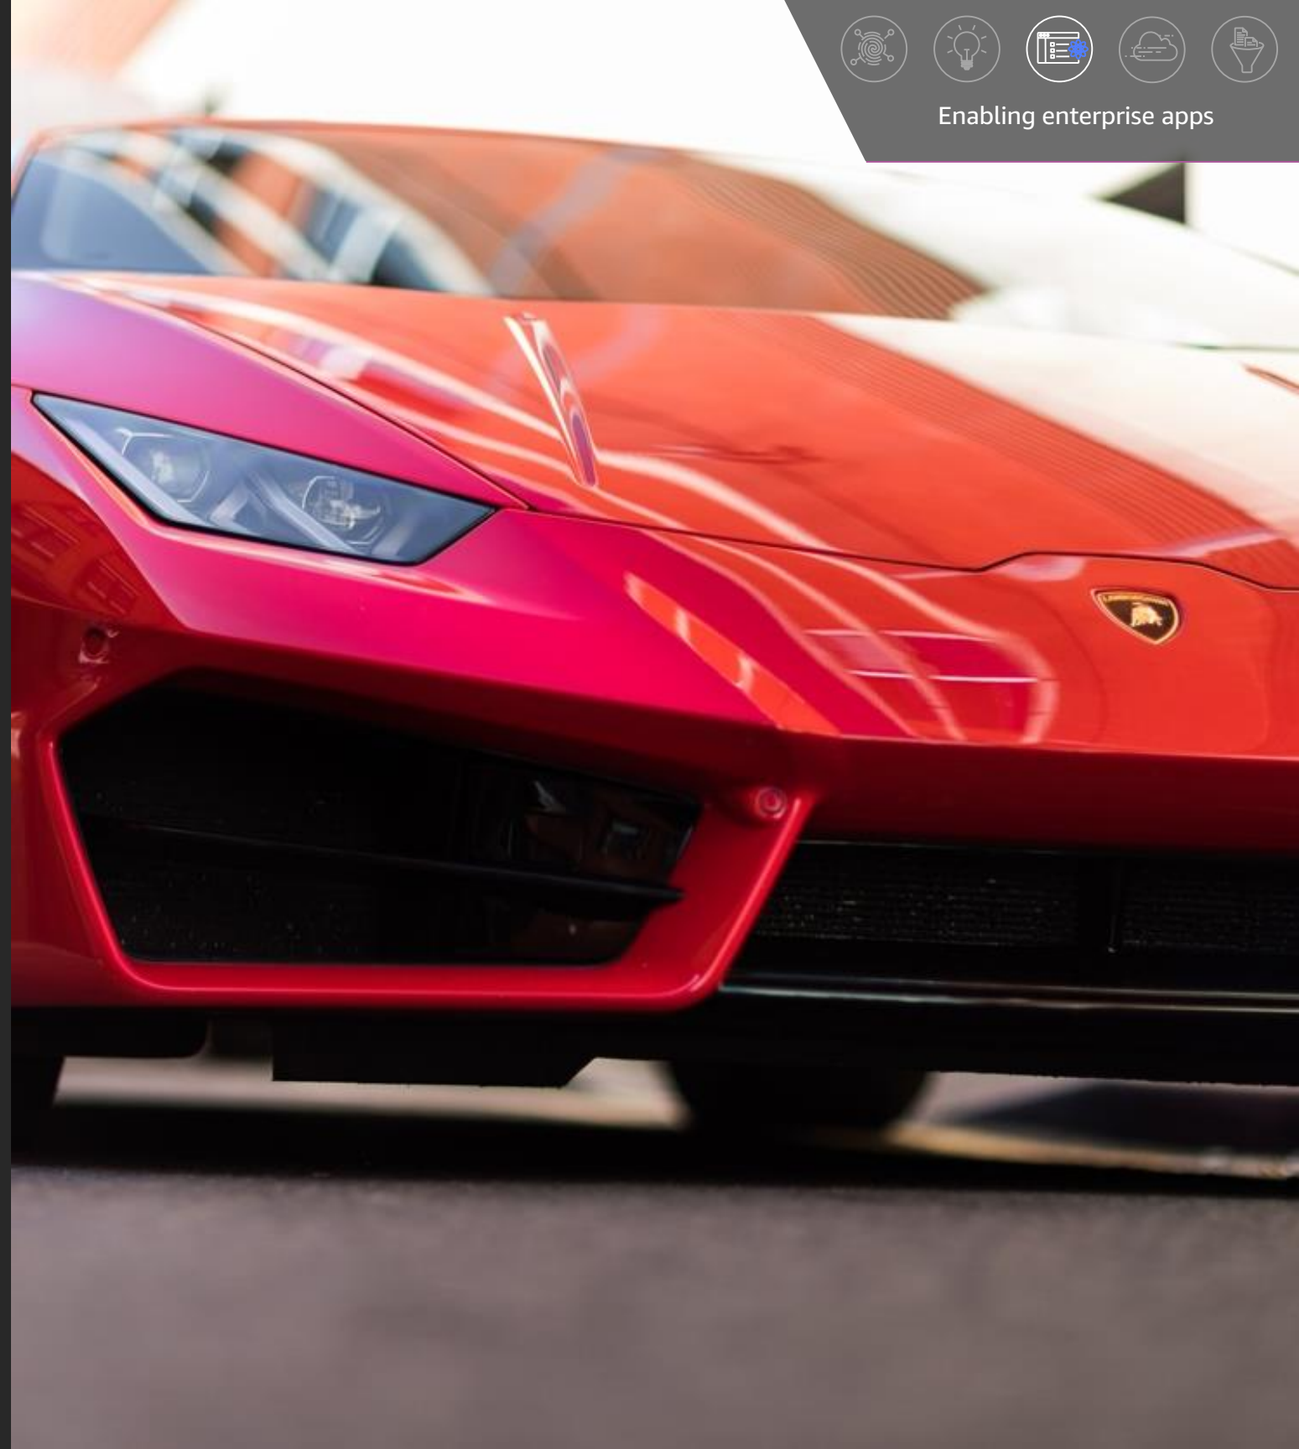
Amazon EC2 instances for SAP HANA





Reducing costs and increasing bandwidth

Lamborghini used AWS to cut its infrastructure costs by 50%, support a 250% burst of visitors, and reduce its time-to-market to close to 0



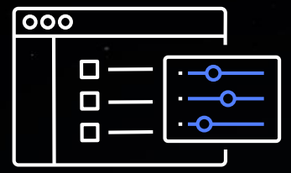
Enabling enterprise apps

VMware Cloud on AWS

VMware SDDC technologies you know and trust, delivered as a service on the world's most popular public cloud



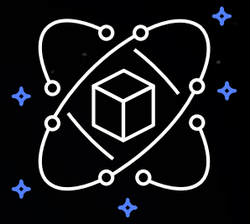
Rich VMware SDDC delivered as a cloud service on AWS



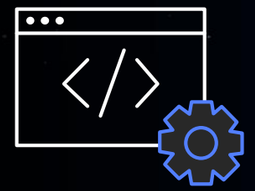
Consistency and familiarity of VMware technologies



Easy workload portability and hybrid capabilities



Direct access to the power of native AWS services



Existing and new apps with containers and VMs

VMware on AWS drives speed and innovation

S&P Global

S&P seamlessly migrated 150+ critical applications to AWS and VMware, so that they can now launch new environments in hours instead of weeks



BUILT TO SUPPORT VIRTUALLY EVERY WORKLOAD



World-class performance, security, and innovation



Broadest and deepest platform choice



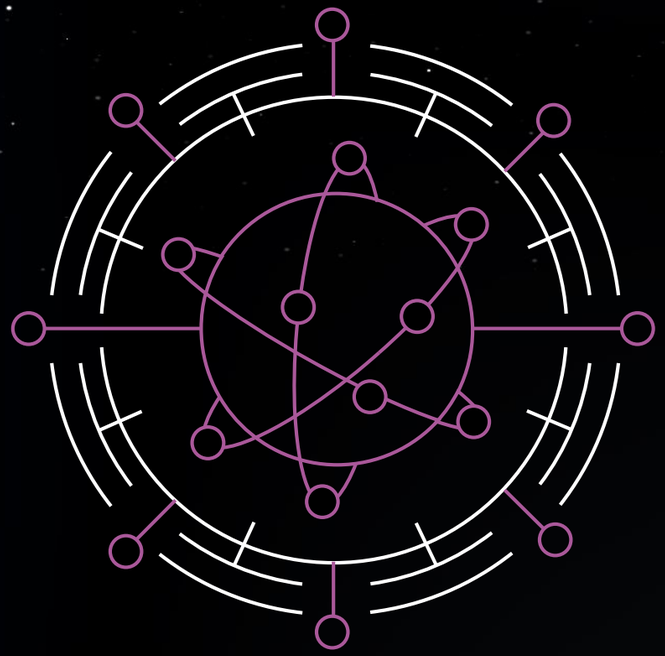
Abstractions for managing cost and complexity



Enabling enterprise applications



Bringing the cloud closer to you



COMPUTE PLATFORM FOR EVERY WORKLOAD

Some workloads need to remain on-premises



LATENCY

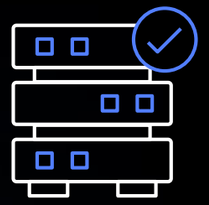
Applications that have very low latency requirements



LOCAL DATA

Applications that need to process large amounts of data locally

Customers want the **same** experience across on-premises and the cloud



Same reliable, secure, and high-performance infrastructure



Same operational consistency



Same services and APIs



Same tools for automation, deployments, and security controls



Same pace of innovation as in the cloud

NEW!

AVAILABLE TODAY

Introducing AWS Outposts

AWS INFRASTRUCTURE AND SERVICES IN YOUR ON-PREMISES LOCATION



Same AWS-designed infrastructure as in AWS data centers (built on AWS Nitro System)



Fully managed, monitored, and operated by AWS as if in AWS Regions



Single pane of management in the cloud providing the same APIs and tools as in AWS Regions

Application optimizations

REDUCING APPLICATION LATENCY AND COST WITH AWS OUTPOSTS

PHILIPS

Philips HealthSuite digital platform uses AWS Services within the region.

However, applications like our radiology PACS medical imaging offering analysis of high-resolution images are costly to send back and forth.

With AWS Outposts, they benefit from having a single set of management tools for applications running within the AWS Region and at healthcare sites.



AWS Outposts infrastructure

Industry standard 42U rack

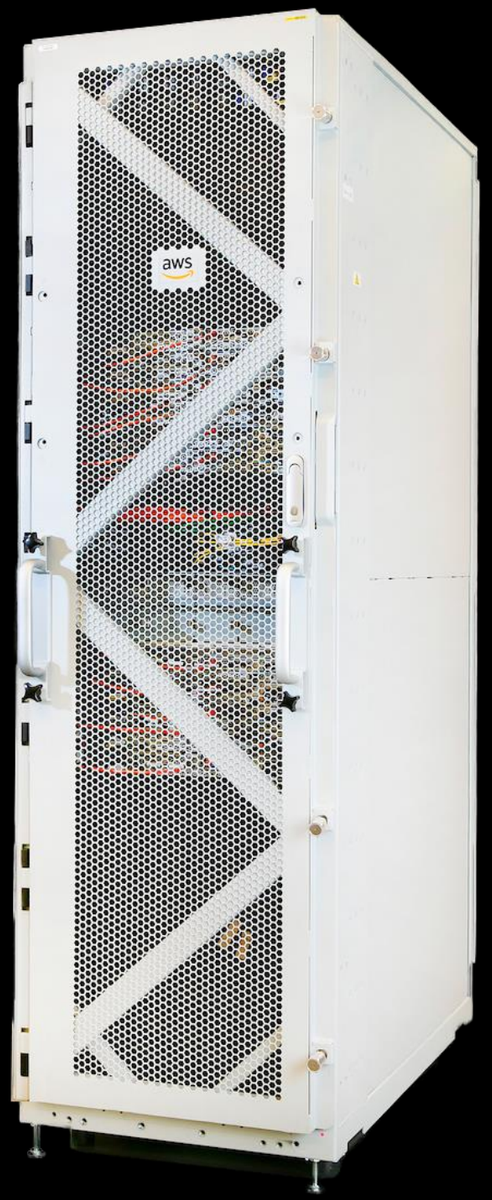
Fully assembled, ready to be rolled into final position

Installed by AWS, simply plug into power and network

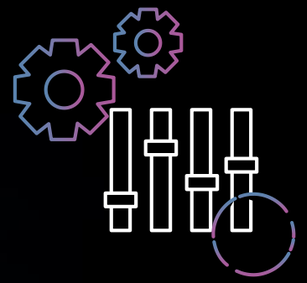
Centralized redundant power conversion unit and DC distribution system for higher reliability, energy efficiency, easier serviceability

Available in 5, 10, and 15 kVA power configurations

Available in 10, 40, and 100 Gbps networking

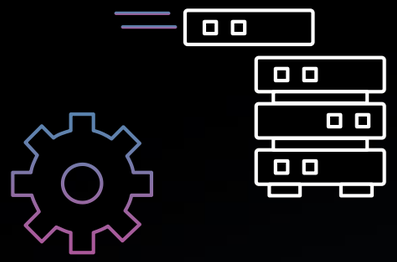


Getting started with AWS Outposts



1 ORDER

Select your compute and storage capacity



2 INSTALL

AWS delivers and installs the Outpost



3 LAUNCH

Use standard AWS APIs or AWS Management Console to launch and run AWS resources locally



4 OPERATE

Build and operate apps using native AWS services

APN Consulting Partners



Provide expertise and training to drive customer outcomes

Provide consulting services to modernize applications and implement solutions

APN Technology Partners

Validated technology solutions



Demonstrate integration with customer's on-premises workload
on AWS Outposts

Titans

Run latency-sensitive applications at the edge using AWS infrastructure and services

Extends AWS infrastructure services, APIs, and tools to where customers need it to support low-latency applications

Fully owned, managed, and supported by AWS

New type of Availability Zone that brings AWS compute, storage, networking, and select AWS services closer to where your end users are located

Los Angeles-based Titan connected to US West (Oregon)



The emergence of 5G networks

Low latency and quicker transfer of large volumes of data

Enables a variety of use cases including IoT, gaming, autonomous vehicles, and live media



Wavelength

Deliver ultra-low-latency applications to 5G devices using AWS

Extends AWS infrastructure, services, APIs, and tools to 5G networks

Delivers consistent AWS experience

Build once and deploy anywhere

AWS has partnered with a number of telcos around the world to launch the 5G edge zones

verizon

vodafone

KDDI

SK telecom



Bringing it all **together**

Recap of EC2 and Hybrid announcements

Pre-re:Invent announcements

EC2 A1 Bare Metal
EC2 M5dn, M5n, R5dn, R5n
EC2 T3a, M5ad, R5ad
EC2 High-Memory 18 and 24 TB (SAP-Certified)
EC2 C5a/C5ad instances
EC2 C5n instances
EC2 I3en
EC2 G4
Savings Plans
AWS Launch Wizard for SQL Server

Announced at re:Invent

EC2 M6g, C6g, R6g with AWS Graviton2
EC2 Inf1 with Inferentia
AWS Compute Optimizer
Nitro Enclaves
Image Builder
AWS Outposts
Titans
Wavelength

Thank you!

AWS re:Invent

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.

aws



Please complete the session survey in the mobile app.

© 2019, Amazon Web Services, Inc. or its affiliates. All rights reserved.