aws INNOVATE

MODERN APPLICATIONS EDITION

20 October, 2022

aws

Solving business challenges with application modernization

Akanksha Balani

AWS APJ Alliance head @ Intel
Global AI GTM Lead

aws

Agenda

- Tech Today

- Why Intel & AWS

- Workloads, Verticals and Solutions

- Learn | Engage | Innovate

Transformation with Cloud

Consumer	Health	Finance	Retail	Government	Energy	Transport	Industrial	Other
Smart Assistants	Enhanced Diagnostics	Algorithmic Trading	Support	Defense	Oil & Gas Exploration	Autonomous Cars	Factory Automation	Advertising
Chatbots	Drug Discovery	Fraud Detection Research	Experience	Data Insights	Smart Grid	Automated Trucking	Predictive Maintenance	Education
Search Personalization	Patient Care Research	Personal Finance	Marketing	Safety & Security	Operational Improvement	Aerospace Shipping	Precision Agriculture	Gaming
Augmented Reality	Sensory Aids	Risk Mitigation	Merchandising	Resident Engagement	Conservation	Search & Rescue	Field Automation	Professional & IT Services
Robots			Loyalty	Smarter Cities				Telco/Media
			Supply Chain					Sports
			Security					

What does Intel do with AWS?

17 years of engineering partnership

Digital transformation

Shared customer passion

High performance + low costs

World-class supply chain

Amazon EC2

Amazon S3

Amazon SageMaker

VMWare Cloud on AWS

AWS Deep Learning AMIs

SAP on AWS

AWS IoT Greengrass

AWS IoT Core

Amazon Echo Show

Amazon Echo Look

AWS DeepLens

AWS DeepRacer

Alexa Voice Service

AWS Outposts

connectivity

connectivity

aws | intel®

Cloud & data center

Things & devices

Key Workloads
Migrations/Modernization
AI/ML
HPC
Hybrid/Edge

> *Intel is a very deep partner of AWS and will be for a long time. That's not changing.*
>
> **Andy Jassy,** CEO, AWS

Greatest variety and availability to meet your global workload needs

aws | intel

275+ Intel instances

General purpose
T3 | M5 | M5n | M5zn | M5dn | **M6i**

Compute optimized
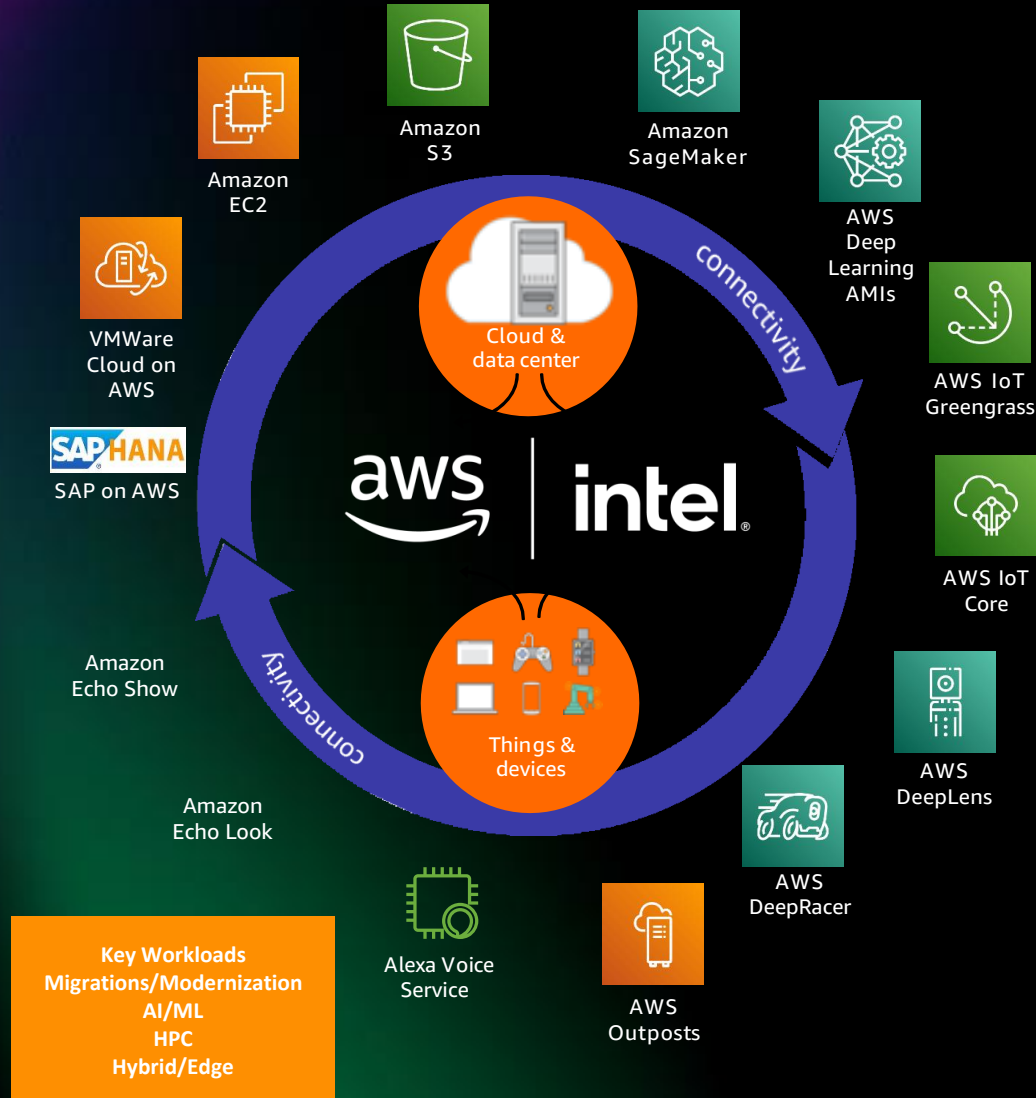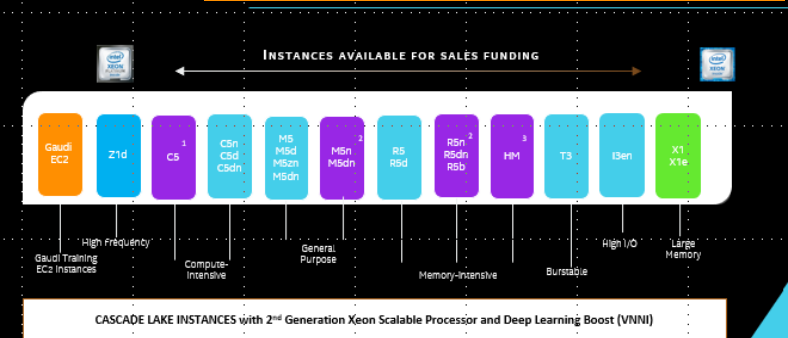C5 | C5n | C5d | C5dn | **C6i**

Memory optimized
R5 | R5n | R5b | X1e / X1 | High Memory | Z1d

Accelerated compute
Gaudi Instances | P3 | G4 | F1

Storage optimized
I3 | I3en | D3/D3en

Workloads

High Performance Computing (HPC)

Artificial Intelligence (AI) Machine Learning (ML) Big Data

SAP on AWS

VMware Cloud on AWS

Internet of Things (IoT)

Strategic Migration

INSTANCES AVAILABLE FOR SALES FUNDING

| Gaudi EC2 | Z1d | C5 | C5n C5d C5dn | M5 M5d M5zn M5dn | M5n M5dn | R5 R5d | R5n R5dn R5b | HM | T3 | I3en | X1 X1e |

Gaudi Training EC2 Instances

High Frequency

Compute-Intensive

General Purpose

Memory-Intensive

Burstable

High I/O

Large Memory

CASCADE LAKE INSTANCES with 2nd Generation Xeon Scalable Processor and Deep Learning Boost (VNNI)

Highlights of the past year

Career Launcher Rapidly Scales Learning Portal during Pandemic
Intel & AWS collaborate to help serve >160,000 students in India within two months on AWS.[1]

65x more parallel wildfire simulations
Intel & AWS work with RONIN to help increase fire fighting effectiveness in Australia. [4]

select
SOLUTIONS

+

AWS ParallelCluster

AWS as first CSP with verified Intel Select Solution [2]

Amazon EC2 M5zn instance – fastest Intel Xeon Scalable CPU in the Cloud [5]
Highest all-core turbo CPU performance with a frequency up to 4.5 GHz.

AWS announces DL1, M6i, C6i, DL1
AI instances with 40% better price/performance built on Habana Gaudi [3]
Intel's Habana & AWS co-engineered solution using up to 8 Gaudi accelerators

[1] https://www.intel.com/content/www/us/en/customer-spotlight/stories/career-launcher-customer-story.html
[2] https://docs.aws.amazon.com/parallelcluster/latest/ug/intel-select-solutions.html
[3] https://aws.amazon.com/ec2/instance-types/habana-gaudi/
[4] https://dpgresources.intel.com/asset-library/intel-aws-the-csiro-spark-intel-poc-summary/
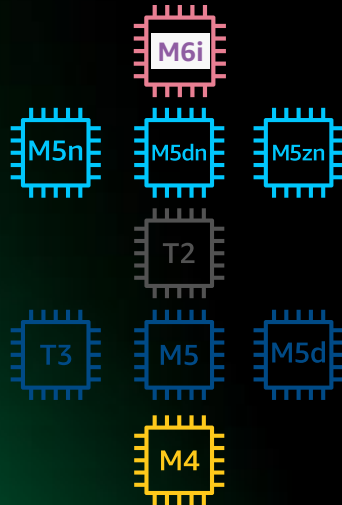[5] https://aws.amazon.com/blogs/aws/new-ec2-m5zn-instances-fastest-intel-xeon-scalable-cpu-in-the-cloud/

aws | intel.

Instance Types on Intel

275+ Intel instances

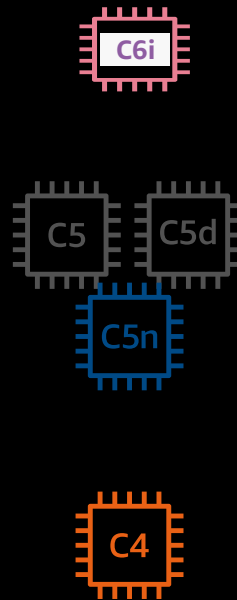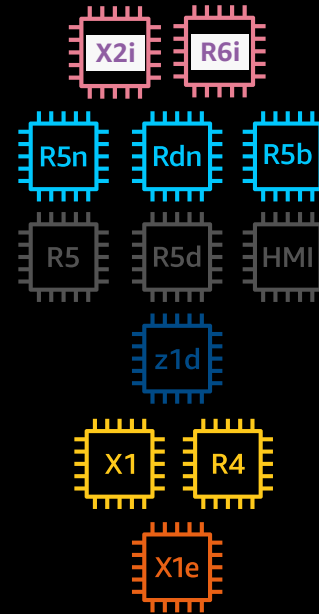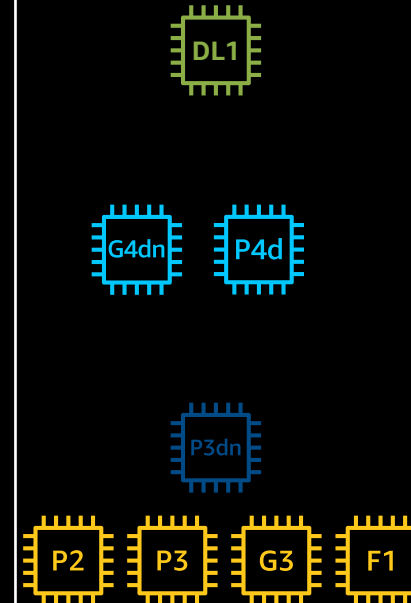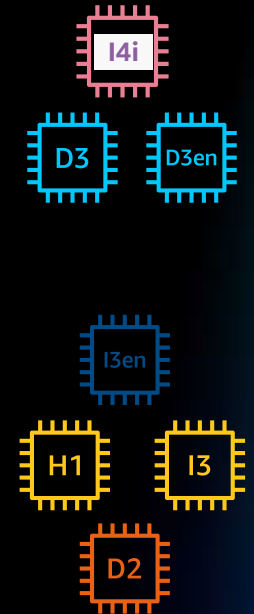|	General Purpose	Compute Optimized	Memory Optimized	Accelerated Compute	Storage Optimized
	General purpose instances provide a balance of compute, memory and networking resources, and can be used for a variety of diverse workloads.	Compute optimized instances are ideal for compute bound applications that benefit from high performance processors.	Memory optimized instances are designed to deliver fast performance for workloads that process large data sets in memory.	Accelerated computing instances use hardware accelerators, or co-processors, to perform functions more efficiently.	Storage optimized instances are designed for workloads that require high, sequential read and write access to very large data sets on local storage.
Habana Gaudi				DL1	
Ice Lake	M6i	C6i	X2i R6i		I4i
Cascade Lake	M5n M5dn M5zn		R5n Rdn R5b	G4dn P4d	D3 D3en
Skylake & Cascade Lake	T2	C5 C5d	R5 R5d HMI		
Skylake	T3 M5 M5d	C5n	z1d	P3dn	I3en
Broadwell	M4		X1 R4	P2 P3 G3 F1	H1 I3
Haswell		C4	X1e		D2

aws

AWS EC2 instance offerings – optimized by use case
Not exhaustive – focused on newer instances

Balanced workloads	Compute-intensive, HPC, data lakes, network appliances	High performance databases, in-memory databases	High IOPS at low cost	Accelerated WLs machine learning, 3D rendering
General Purpose	**Compute Optimized**	**Memory Optimized**	**Storage Optimized**	**Accelerated Computing**
M6i *(NEW)* Up to 128 vCPUs & 512 GB of memory - ICX	**C6i** *(NEW)* Up to 128 vCPUs & 256 GiB of memory – ICX	**X2i** *(NEW)* Memory-optimized & up to 4,096 GiB of memory- ICX&CLX	**I4i** *(NEW)* NVMe SSD Storage New size with up to 128 vCPUs and 1,024 GiB of memory-ICX	**P3dn** P3 with Local Host Attached NVMe SSDs and up to 100Gbps Network Bandwidth SKX – 24C
M5 Non-burstable CPU usage SKX/CLX – 24C	**C5** High performance $/performance optimized SKX – 18C CLX – 24C	**R6i** *(NEW)* Up to 128 vCPUs & 1,024 GiB of memory – ICX	**I3** NVMe SSD Storage and Bare Metal Instances BDX - 16C	**G4** 2 NVIDIA Tesla M60 GPUs per CPU CLX – 24C
M5d M5 with local host attached NVMe SSDs SKX/CLX – 24C	**C5d** C5 with local host attached NVMe SSDs SKX – 18C CLX – 24C	**R5, R5b** Up to 768GiB of Memory SKX - 24C	**I3en** NVMe SSD Storage and Bare Metal Instances SKX – 24C	**F1** 4 FPGAs per CPU BDX - 16C
M5zn M5 with local host attached NVMe SSDs SKX/CLX – 24C	**C5n** C5 with up to 100 Gbps network bandwidth SKX – 18C	**X1** One of the Lowest Price/GiB of RAM HSX - 16C (4 socket)	**H1** Compute and Memory Balanced, Up to 16TB HDD Storage BDX - 16C	**DL1** *(NEW)* Habana Gaudi AI/ML Up to 8 accelerators 40% better price/perf than current GPU-based instances
T3 Burstable CPU usage SKX/CLX – 24C		**X1e** X1 with Extended Memory Footprint HSX - 16C (4 socket)	**D3, D3en** Up to 366 TB HDD Storage, Lowest Price/Disk Throughput Perf CLX – 24C	
		Z1d Single threaded compute optimized with high memory SKX - 12C		
		Bare Metal 8 Socket Xeon with 6 TiB Memory up to 24 TiB; SKX/CLX - 28C		

https://aws.amazon.com/ec2/instance-types/

Intel offerings for customers in the cloud

	WORKLOAD BENCHMARKING	With decades of experience in workload benchmarking, Intel helps reduce the time to do baseline performance analysis on public cloud instances for comparisons across Xeon generations and competitive architectures.
	PERFORMANCE TOOLS	Intel supports a suite of tools to optimize workloads, raising the performance ceiling. Performance analysis and tuning can be conducted with Intel VTune with Flame Graphs and support for Java, Go, and Python. Cookbooks are in development for use of Vtune on cloud instances.
	RESOURCE OPTIMIZATION	Intel is investing to optimize cloud resources and costs. Examples include partnerships with Densify and Granulate to offer optimizers leveraging machine learning to increase utilization and performance at different scopes. Intel also offers Optimized Cloud Images to take the guess work out of configurations and deliver verifiable performance.
	OPEN SOURCE COLLABORATION	As a contributor to open source communities, Intel is enabling performance, security, and portability. Examples of Intel's engagements and contributions include Linux, Java, Kubernetes, Ceph, RocksDB, and Cassandra. Intel is also driving native support for AI acceleration through optimizations on common frameworks like TensorFlow.
	ARTIFICIAL INTELLIGENCE	Intel has invested in open-source AI frameworks like PyTorch and TensorFlow which now has native support for Xeon AI acceleration technologies (AVX, VNNI) on cloud instances. With hardware and software engineering expertise, Intel is here to provide consultation and software enablement of AI solutions in the cloud.

aws

Portfolio of innovation

Solutions
Solution Architects

Platforms · Finance · Healthcare · Energy · Industrial · Transport · Retail · Home · More…

Toolkits
App Developers

OpenVINO™ Toolkit	Deep Learning Developer Toolkit
OpenVINO Toolkit for inference deployment on CPU, processor graphics, FPGA & VPU using TF, Caffe* & MXNet*	*Optimized inference deployment for all Intel® Movidius™ VPUs using TensorFlow* & Caffe**

Libraries
Data Scientists

MACHINE LEARNING LIBRARIES

Python	R	Distributed
• Scikit-learn	• Cart	• MlLib (on Spark)
• Pandas	• RandomForest	• Mahout
• NumPy	• E1071	

DEEP LEARNING FRAMEWORKS

TensorFlow* · MXNet* · Caffe* · BigDL/Spark*   Caffe2* · PyTorch* · PaddlePaddle*

Foundation
Library Developers

ANALYTICS, MACHINE & DEEP LEARNING PRIMITIVES

Python	DAAL	MKL-DNN	clDNN
Intel distribution optimized for machine learning	*Intel® Data Analytics Acceleration Library (for machine learning)*	*Open-source deep neural network functions for CPU, processor graphics*	

Hardware
IT System Architects

FOUNDATION

intel ATOM inside · intel CORE 8th Gen · intel XEON PLATINUM inside

ACCELERATORS

intel STRATIX 10 inside · intel ARRIA 10 inside · intel MOVIDIUS inside · Movidius An Intel Company · intel GNA inside · habana An Intel Company

Inference

New 3rd Gen Intel Xeon Scalable Processor

- Higher workload performance
- Designed for reliability at scale
- New crypto acceleration
- Advanced security capabilities
- Total Memory Encryption (TME)

Up to
1.58x
Improvement in web microservices performance

Up to
40%
Performance improvement (Specrate2017_int_base)
on new Ice Lake SKU offerings vs. Cascade Lake

Up to
1.42X
More cores per processor
40-core Ice Lake vs. 28-core Cascade Lake

intel.
XEON

3rd Gen Intel® Xeon®
Scalable Processors

App Workload @ Intel

Enterprise Migration & Modernization

AI / ML

HPC

Edge: Hybrid / Outpost/C2E

Enterprise Migration & Modernization

AI / ML

HPC

Edge / Hybrid with Outpost

For every point of the cloud journey:

Continued support on typical workloads such as:
- SAP HANA
- VMWare Cloud
- MSFT
- Red Hat OpenShift
- Databricks

AI/ML

Amazon EC2 DL 1 Instances available for deep learning capabilities at 40% better price/performance

New 3rd Gen Xeon with DL Boost

AWS Deep Racer - for the developers

AWS Parallel Cluster

First CSP with verified Intel Select Solution for HPC workloads

Cloud to Edge and Hybrid Solutions

On premise fully managed rack configurations
Optimized for VMWare Cloud
Seamlessly extend AI/ML to on premise with Intel based gateways and technologies

Private Networks

Seamlessly extend virtual private cloud (VPC) on premise

aws

SAP-Certified Instances on AWS Powered by Intel® Xeon Processors

Amazon EC2 Memory Optimized instances powered by Intel Xeon processors, have highly scalable memory for memory-intensive enterprise applications that benefit from lower latency. There are more than 40 Intel-powered instances certified by SAP to run Business Warehouse on HANA (BW), Data Mart Solutions on HANA, Business Suite on HANA (SoH), Business Suite S/4HANA.

Workload	Instance Family	Instance Type	Use Cases	Notable Features *2 vCPUs / core*
SAP on AWS	Memory Optimized	R6i, R5, R5n, R5dn, R5b	• Mid-size in-memory and high performance databases, distributed web scale in-memory caches, real time big data analytics. • 10% price per GiB improvement and a ~20% increased CPU performance over R4	• 8 GiB/vCPU memory • Up to 100 Gbps networking (R5n)
	High Memory/ Performance Database	X1e, X1, X2i	• World's first, purpose-built instance for SAP in the cloud • One of the lower price per GiB of RAM within Amazon EC2 • Optimized for large-scale, enterprise-class and in-memory applications	• ~15 GiB / vCPU memory • Up to 1,952 GiB of DRAM memory • High freq Intel Xeon E7-8880 v3 (HSW)
	Enterprise Scale, In-Memory Database	High Memory	• First ever Amazon EC2 instances built on 8-socket hosts. • Large in-memory databases ; SAP HANA deployments • Set a new record of 480,600 SAPs[1]	• Up to 448 vCPUs • 6, 9 & 12 TB memory on Skylake • 18 & 24 TB memory on Cascade Lake

Recommended video - 6 Tips for Selecting the Best Cloud for SAP by Lemongrass, an AWS Premier Consulting Partner

[1] https://aws.amazon.com/blogs/awsforsap/new-sap-certifications-for-aws-instances-and-world-record-benchmark-results/

VMware Cloud on AWS
Unlock the full potential of your Hybrid Cloud with VMware and Intel

Intel and VMware enable optimized and unified hybrid cloud solutions on AWS. VMware Cloud is delivered on next-gen bare metal instances and powered by Intel Xeon processors. Optimized high I/O instances and featuring low-latency Non-Volatile Memory Express (NVMe) based SSDs, customers can integrate their clouds and extend business with VMC on AWS powered by Intel.

Instance type	Instance	Best Use Cases	Notable Features
Storage-Optimized	I3.metal	• Best for NoSQL databases, in-memory databases, scale-out transactional databases, data warehousing, Elasticsearch, and analytics workloads	• 72 vCPUs - Broadwell • Up to 15TB of NVMe SSD storage • 512 GiB memory
	I3en.metal	• Targeted for Disaster Recovery, large databases and HPC analytics. • Optimized for low latency, high random I/O performance, high sequential disk throughput • Lowest price per GB of SSD instance storage on Amazon EC2.	• 96 vCPUs - Intel 2nd Gen Xeon Scalable processor (custom CLX) • Up to 60 TB of NVMe storage • 768 GiB memory

Database, Data Analytics, Big Data
Unmatched breadth of Optimized Instances to Right Size your Instance

Intel data-specific optimizations result in higher cache capacity and faster increased memory bandwidth while instructions such as **Intel® Transactional Synchronization Extensions (Intel® TSX)** maximize performance while enabling consistent, low latencies. Instances featuring Intel 3rd Gen Xeon® Scalable® processors (ICX) benefit from **Intel® Mesh Architecture** for efficient and scalable data flow between all vCPUs.

Workload	Instance Family	Instance Type	Best Use Cases	Notable Features
Database, Data Analytics, Big Data	Memory Optimized	R6i, R5, R5n, R5b	• Fast performance for workloads that process large data sets in memory • Memory-intensive applications include high performance databases, enterprise applications, and real time big data analytics	• Up to 128 vCPUs • Up to 768 GiB memory, 100 Gbps nwkg • R5b delivers up to 60 Gbps and 260k IOPS of EBS
	Storage-Optimized	I3, I3en	• Designed for high, sequential read and write access to very large data sets on local storage • Deliver tens of 1000's of low-latency, random IOPS • NoSQL, in-memory and scale-out transactional databases, data warehousing, Elasticsearch, analytics, Disaster Recovery	• Up to 96 vCPUs (BDW, CLX) • Up to 768 GiB memory, 100 Gbps nwkg • Up to 60 TB of NVMe storage • I3en offers lowest price per GB of SSD
	Storage-Optimized	D3, D3en	• Optimized for high sequential I/O performance, disk throughput and low cost storage for very large data sets • Multi-node file storage systems such as Lustre, BeeGFS, GPFS, VxCFS, and GFS2. High Capacity data lakes	• Up to 48 vCPUs (CLX), 75 Gbps nwkg • Up to 336 TB of HDD storage • D3en offer lowest HDD storage costs
	General Purpose	M6i, M5, M5n, M5dn	• Balance of compute, memory and networking resources, and can be used for a variety of diverse workloads	• Up to 128 vCPUs • Up to 512 GiB memory, 100 Gbps nwkg
	Compute Optimized	C6i, C5, C5n, C5dn	• Optimized for compute-intensive workloads and deliver cost-effective high performance at a low price per compute ratio.	• Up to 128 vCPUs • Up to 192 GiB memory, 100 Gbps nwkg

Artificial Intelligence (AI) and Machine Learning (ML)

AWS and Intel make building custom machine learning and deep-learning models easier and more cost-effective than ever before. AWS C6i and M6i instances, powered by 3rd Gen Intel Xeon Scalable processors offer AVX-512 – and in some instances, DL Boost – to speed up inference and other data intensive workloads.

Workload	Instance Family	Instance Type	Best Use Cases	Notable Features
AI/ML	Compute Optimized	C6i, C5, C5d, C5n	• Optimized for compute-intensive workloads • Deliver cost-effective high performance at a low price per compute ratio	• Ice lake now available • Most instances on SKL with AVX-512 • Larger size instances (12, 24.xlarge, metal) on CLX w support for DL Boost 2 GiB / vCPU memory
	General Purpose compute	M6i, M5, M5d	• General purpose instance provides a balance of compute, memory, and network resources	• Ice Lake & Sky Lake with Intel AVX-512 • 4 GiB/vCPU memory
	General Purpose	M5zn	• AWS claims "Fastest Intel Xeon Scalable processors in the cloud.* " • M5zn instances are an ideal fit for applications that benefit from extremely high single-thread performance and high throughput, low latency networking, such as gaming, High Performance Computing, and simulation modeling for the automotive, aerospace, energy, and telecommunication industries.	• CLX with all-core turbo frequency of up to 4.5 GHz • AVX-512 and DL Boost • Available on bare metal • Up to 100 Gbps of network bandwidth on the largest instance size and bare metal variant

See https://aws.amazon.com/ec2/instance-types/
https://aws.amazon.com/blogs/aws/new-ec2-m5zn-instances-fastest-intel-xeon-scalable-cpu-in-the-cloud/

AWS Customer References

THORN and MEMSQL

Critical time saved

65%

Powered by Intel® Xeon® Scalable processors, Amazon EC2 C5 instances with Amazon S3 and Amazon Rekognition help law enforcement fight child trafficking.

18,119
Victims identified

5,791
Children identified

6,553
Traffickers identified

Thorn Finds More Human Trafficking Victims Faster

Empowering law enforcement to collaborate beyond jurisdictions

Need: Abusers use advanced technology to facilitate their exploitation of children, and law enforcement needed to turn the tables and find these children faster.

Solution: Spotlight software ingests >100,000 online escort ads/day, storing them in Amazon S3. Amazon Rekognition in combination with MemSQL, a scalable database for operational analytics, helps find photos that have been edited to defeat image-search engines. Spotlight's ML models run on Amazon EC2 C5 instances, powered by Intel® Xeon® Scalable processors.

Value: Spotlight cut the time it takes by 65% to help the U.S. and Canada identify 18,119 victims—with 5,791 children and 6,553 traffickers identified.

"Spotlight helps officers identify child sex-trafficking ads much faster than the old paper-and-pencil methods."

Brooke Istook
Director of Strategy and Operations, Thorn

Products and Solutions
Amazon Rekognition
Amazon S3
Intel® Xeon® Scalable processors

Industry
AI-ML/Image Analysis

Organization Size
1-50

Country
United States

Partners
AWS
Thorn
MemSQL

Learn more
Partner Story

aws

High Performance Computing (HPC)

Intel based Amazon EC2 instances power the most computationally demanding applications in a cost effective way at scale. Intel and AWS offer a comprehensive set of leading compute, networking, storage, and visualization technologies to give customers an ideal environment for HPC workloads. Coupled with an extensive partner ecosystem, customers are empowered to innovate more freely.

| Workload	Instance Family	Instance Family	Best Use Cases	Notable Features
HPC	Compute Networking Performance	C5n	• Ideal choice for HPC workloads, data lakes • Network appliances that can take advantage of improved network throughput and packet rate performance	• Better memory bandwidth than C5 • Up to 100 Gbps network bandwidth • 5.25 GB/core memory • AVX-512
	Compute Performance	C6i, C5	• Optimized for compute-intensive workloads • Deliver cost-effective high performance at a low price per compute ratio	• DL Boost (12.xlarge, 24.xlarge, bare metal) • AVX-512 • 4GB/core memory
	Fastest Compute	Z1d M5zn	• Z1d targets both memory- and compute-intensive apps • Z1d ideal for EDA, gaming, and certain relational database workloads with high per-core licensing costs • M5zn delivers high perf, high throughput plus low latency networking for gaming, simulation modeling, telco workloads	• High single thread performance with sustained all core frequency up to 4.5 GHz • Z1d = 16 GiB / vCPU memory • Z1d has up to 1.8 TB of instance storage • M5zn – up to 100 Gbps network bandwidth
	Balanced Networking	M5 (+M5n, M5dn)	• General purpose instance that provides a balance of compute, memory, and network resources • Good for many applications including web, application and gaming servers and small to mid-size databases.	• 8 GB/core memory • Up to 100 Gbps network bandwidth (M5n) • Up to 3.6 TB of instance storage

aws

C5n and EFA are enabling CAE/CFD codes to scale on AWS
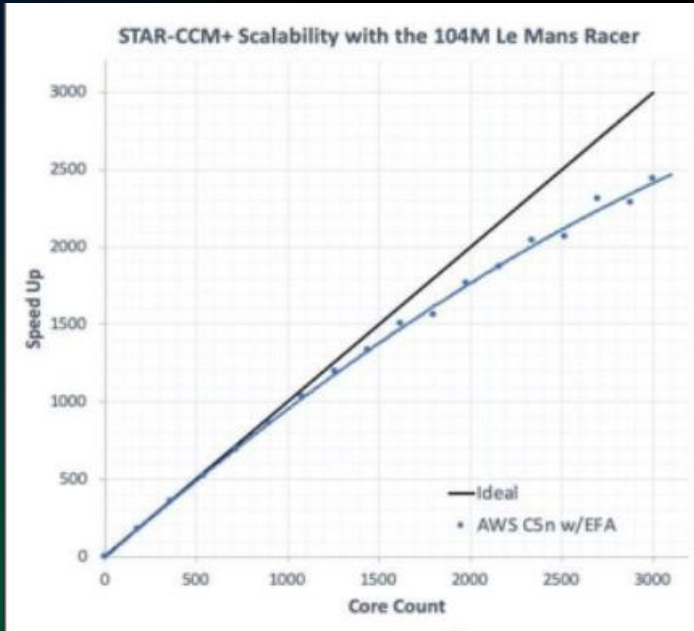
Simcenter STAR-CCM+
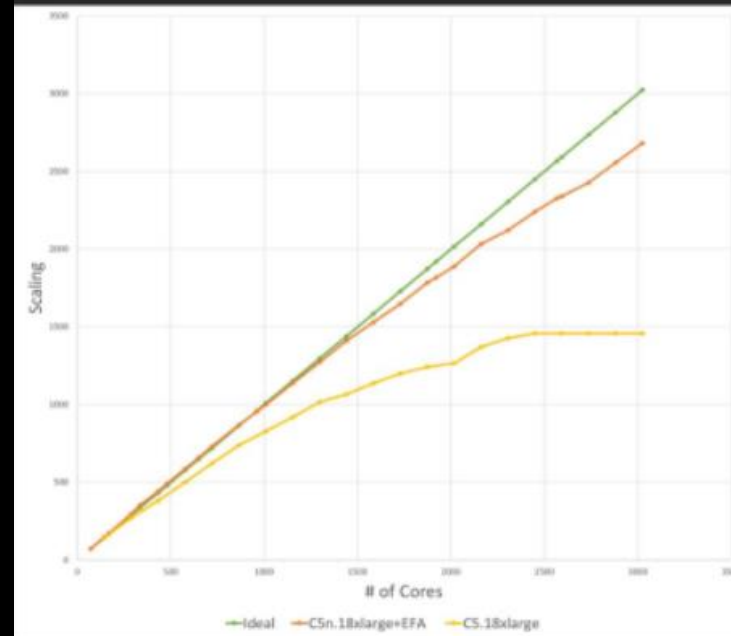
ANSYS Fluent
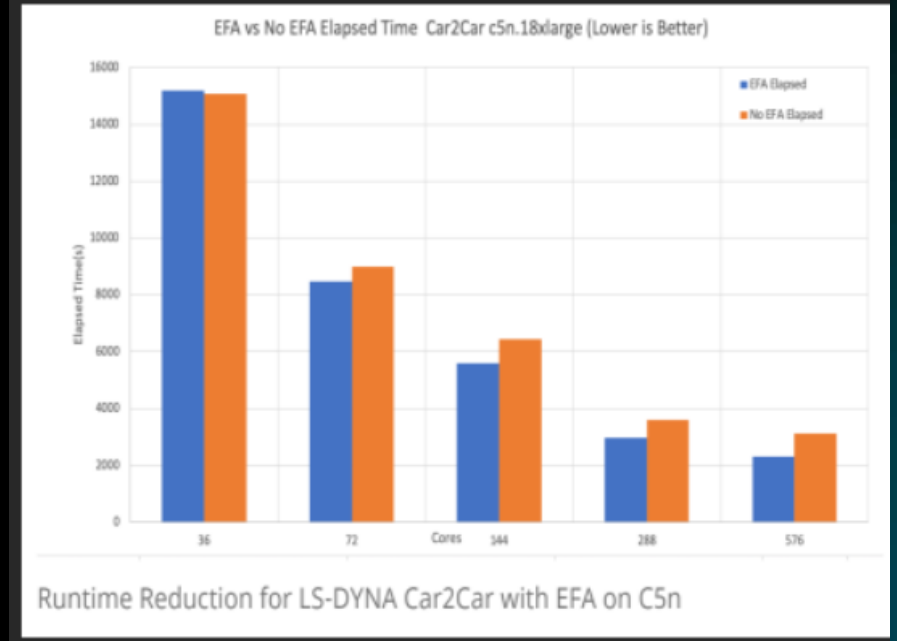
LS-DYNA

Fred Hutch Microbiome Researchers Use AWS to Perform 7 Years of Compute Time in 7 Days

Other benefits

- Processes data from more than 15,000 biological samples
- Increases resolution on microbiome samples to find links to improve health outcomes
- Enables collaboration with other scientific researchers

"Our goal is to accelerate our research processes on AWS so we can get closer to developing therapeutics to fight cancer."

Sam Minot, *PhD and Staff Scientist, Fred Hutch Microbiome Research Initiative*

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

AWS Outposts: proven and familiar AWS and Intel Technology for your hybrid environment

Same AWS-designed infrastructure as in AWS data centers built on Intel® Xeon® Scalable processors

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

If you're already running your applications on Intel® Xeon® servers on-prem you can expect consistent and robust application performance when you run them on Amazon Outposts.

Intel powered AWS Edge Offering

Customers can develop powerful and cost-effective IoT solutions with OpenVINO and Amazon solutions such as Amazon SageMaker and AWS IoT Greengrass on Intel powered gateways and edge devices.

Solution	Instance Family	Instance Type	Best Use Cases	Notable Features
Outposts	Rack	C6i, C5	• Optimized for compute-intensive workloads • Deliver cost-effective high performance at a low price per compute ratio	• DL Boost • 4GB/core memory
	Server	M6i, M5 (+M5n, M5dn)	• General purpose instance that provides a balance of compute, memory, and network resource.	• AVX-512 • 4 GiB/vCPU memory
Outposts	Developers can accelerate ML inference and image recognition with the Intel® Distribution of OpenVINO™ toolkit recently published in the AWS Marketplace. The Toolkit lets developers choose from various edge accelerators to speed up edge to cloud solutions.			
IoT (Edge Device)	▪ Intel IoT Market Ready Solutions: https://solutionsdirectory.intel.com/intel_iot_market_ready_solution ▪ AWS edge device catalogue: https://devices.amazonaws.com/search?hwpf=intel&page=1			

aws

Internet of Things (IOT)

Customers can develop powerful and cost-effective IoT solutions with OpenVINO and Amazon solutions such as Amazon SageMaker and AWS IoT Greengrass on Intel powered gateways and edge devices.

Workload	Instance Family	Instance Type	Best Use Cases	Notable Features
IoT (Cloud)	IoT (Cloud) Compute Performance	C6i, C5	• Optimized for compute-intensive workloads • Deliver cost-effective high performance at a low price per compute ratio	• DL Boost • 4GB/core memory
	General Purpose	M6i, M5 (+M5n, M5dn)	• General purpose instance that provides a balance of compute, memory, and network resource.	• AVX-512 • 4 GiB/vCPU memory
ML Inferencing w/ OpenVINO	Developers can accelerate ML inference and image recognition with the Intel® Distribution of OpenVINO™ toolkit recently published in the AWS Marketplace. The Toolkit lets developers choose from various edge accelerators to speed up edge to cloud solutions.			
IoT (Edge Device)	▪ Intel IoT Market Ready Solutions: https://solutionsdirectory.intel.com/intel_iot_market_ready_solution ▪ AWS edge device catalogue: https://devices.amazonaws.com/search?hwpf=intel&page=1			

Hospital Optimizes Efficiency with Datarwe on AWS and Intel

Gold Coast Health

Challenge

As Gold Coast University Hospital streamed and aggregated more data into its electronic medical records system, it wanted to bring in real-time streaming data and use bedside medical data to create a new application.

Solution

The hospital worked with Datarwe, an AWS Public Sector Partner, to deploy Clinical Data Nexus, an acute care medical research data platform that runs on AWS and uses Intel Xeon Processors.

Benefits

- Collects data on hundreds of patients daily
- Gains insight into patient outcome or risk prediction
- Gains better visibility into hospital supply and demand for ICU beds
- Drives new efficiencies

About Gold Coast University Hospital

Gold Coast University Hospital, based in Southport, Australia, is a 750-bed tertiary teaching hospital that serves the needs of most of the population of Australia's Gold Coast. The hospital offers facilities for cancer and cardiac services, neurosciences, trauma, and intensive care.
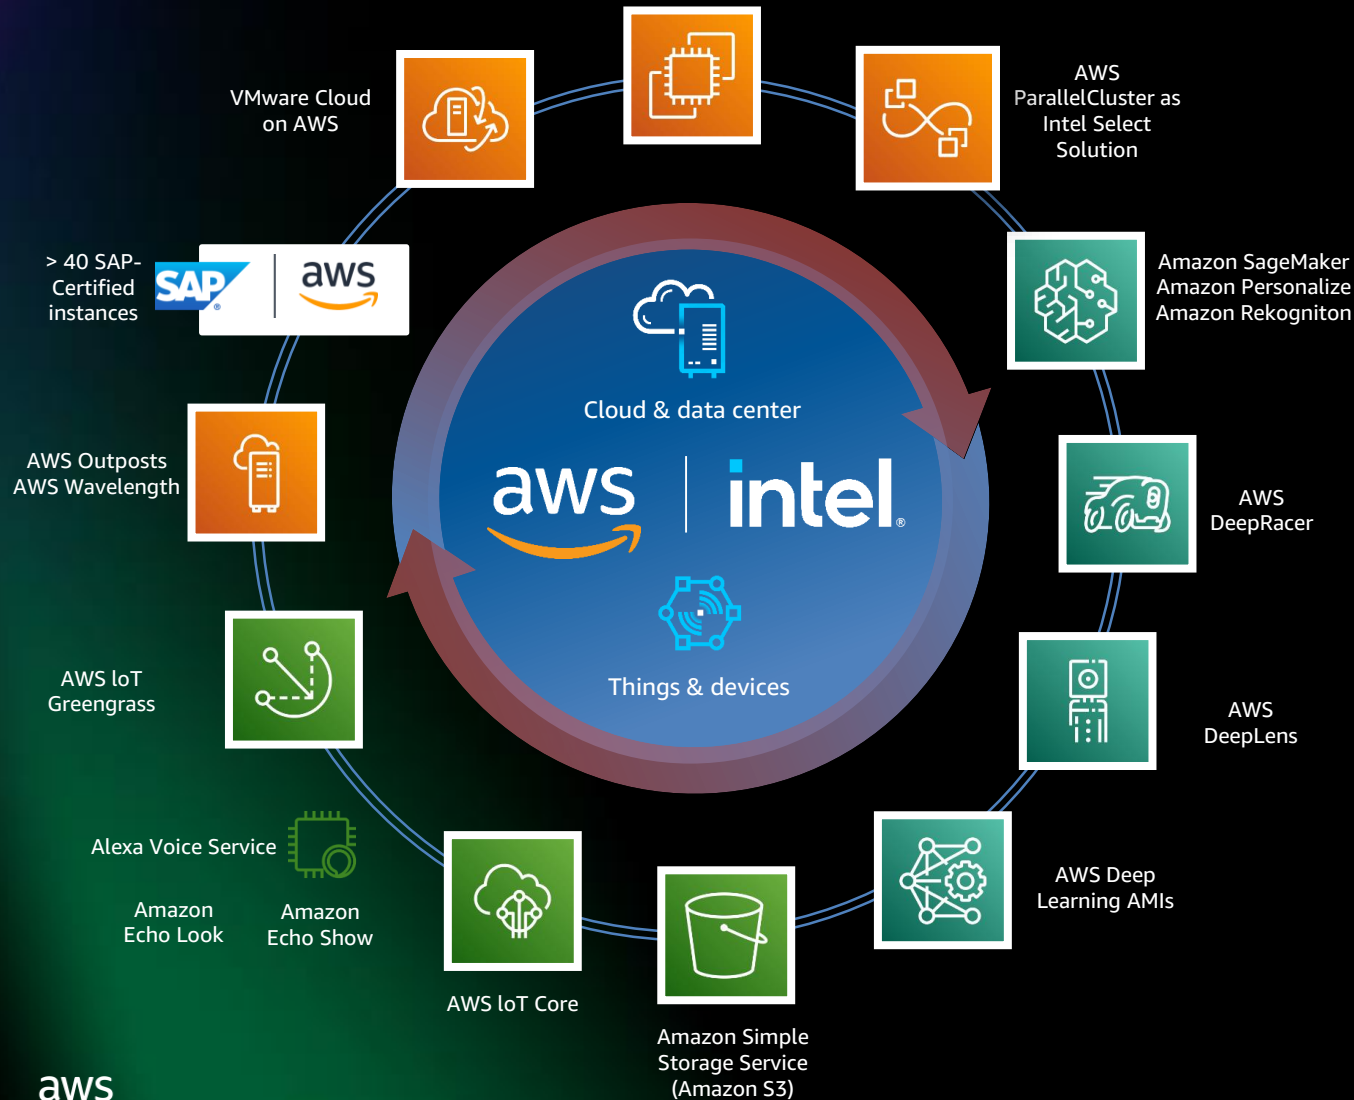
datarwe

About Datarwe

Datarwe, an AWS Public Sector Partner, is a data-driven technology company that provides an acute care medical research data platform as a service. The platform collects data from de-identified intensive care episodes and enriches the data to be research ready.

> "The world of medical IoT is in its infancy, and we have an enormous amount of equipment in hospitals that is not connected. The Datarwe solution, because of the power of AWS and Intel, will help us start unlocking the potential of medical IoT data.

> **- Dr. Brent Richards,** Medical Director of Innovation and Director Critical Care Research, Gold Coast University Hospital

AWS and Intel – Better together

Summary

Amazon EC2 > 250 instance types for database, SAP, VMware, AI, HPC and more

AWS ParallelCluster as Intel Select Solution

VMware Cloud on AWS

Amazon SageMaker Amazon Personalize Amazon Rekogniton

> 40 SAP-Certified instances

SAP | aws

Cloud & data center

aws | intel®

Things & devices

AWS DeepRacer

AWS Outposts AWS Wavelength

AWS DeepLens

AWS IoT Greengrass

AWS Deep Learning AMIs

Alexa Voice Service

Amazon Echo Look

Amazon Echo Show

AWS IoT Core

Amazon Simple Storage Service (Amazon S3)

- Close collaboration between Intel and AWS has resulted in best-in-class end-user experience and customer successes.

- Instance types with the best TCO on Intel to accelerate your customers' applications across a variety or workloads.

- Existing solutions for deployment with many successful outcomes delivering both high performance and cost savings.

Learn | Explore | Engage AI on Intel
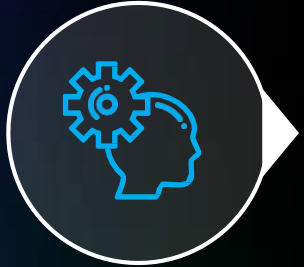
Learn ▶ More information at
https://aws.amazon.com/intel/ on AWS
& Intel

Explore ▶ New instances based on Intel Xeon
Scalable on AWS (M6i,C6i, R6i, M5, C5,
C5n, R5, T3)

Engage ▶ Contact your Intel/AWS representative
for access to Intel AI and POC
opportunities/case studies

CREATE WORLD CHANGING TECHNOLOGY THAT ENRICHES THE LIVES OF EVERY PERSON ON EARTH

aws

Thank you for attending AWS Innovate Modern Applications Edition

We hope you found it interesting! A kind reminder to **complete the survey.**
Let us know what you thought of today's event and how we can improve the event experience for you in the future.

aws-apj-marketing@amazon.com

twitter.com/AWSCloud

facebook.com/AmazonWebServices

youtube.com/user/AmazonWebServices

slideshare.net/AmazonWebServices

twitch.tv/aws

aws

Thank you!

aws