



INNOVATE

DATA EDITION

19 August 2021

Solving business challenges with your data

Akanksha Balani

AWS APJ Alliance Head, Intel
Global AI and HPC GTM Lead



Agenda

- Data today
- Why Intel & AWS
- Successful data analytics with Intel & AWS
- Learn | Engage | Innovate on data

Data transformation across industries



Consumer

Smart Assistants
Chatbots
Search
Personalization
Augmented Reality
Robots



Health

Enhanced Diagnostics
Drug Discovery
Patient Care
Research
Sensory Aids



Finance

Algorithmic Trading
Fraud Detection
Research
Personal Finance
Risk Mitigation



Retail

Support
Experience
Marketing
Merchandising
Loyalty
Supply Chain
Security



Government

Defense
Data Insights
Safety & Security
Resident Engagement
Smarter Cities



Energy

Oil & Gas Exploration
Smart Grid
Operational Improvement
Conservation



Transport

Autonomous Cars
Automated Trucking
Aerospace
Shipping
Search & Rescue



Industrial

Factory Automation
Predictive Maintenance
Precision Agriculture
Field Automation



Other

Advertising
Education
Gaming
Professional & IT Services
Telco/Media
Sports

over
93% of the world's
data runs on Intel

81M+

Intel® Xeon®
processors deployed
in the past 3 years

20+

years of delivering
incredible performance
gen on gen

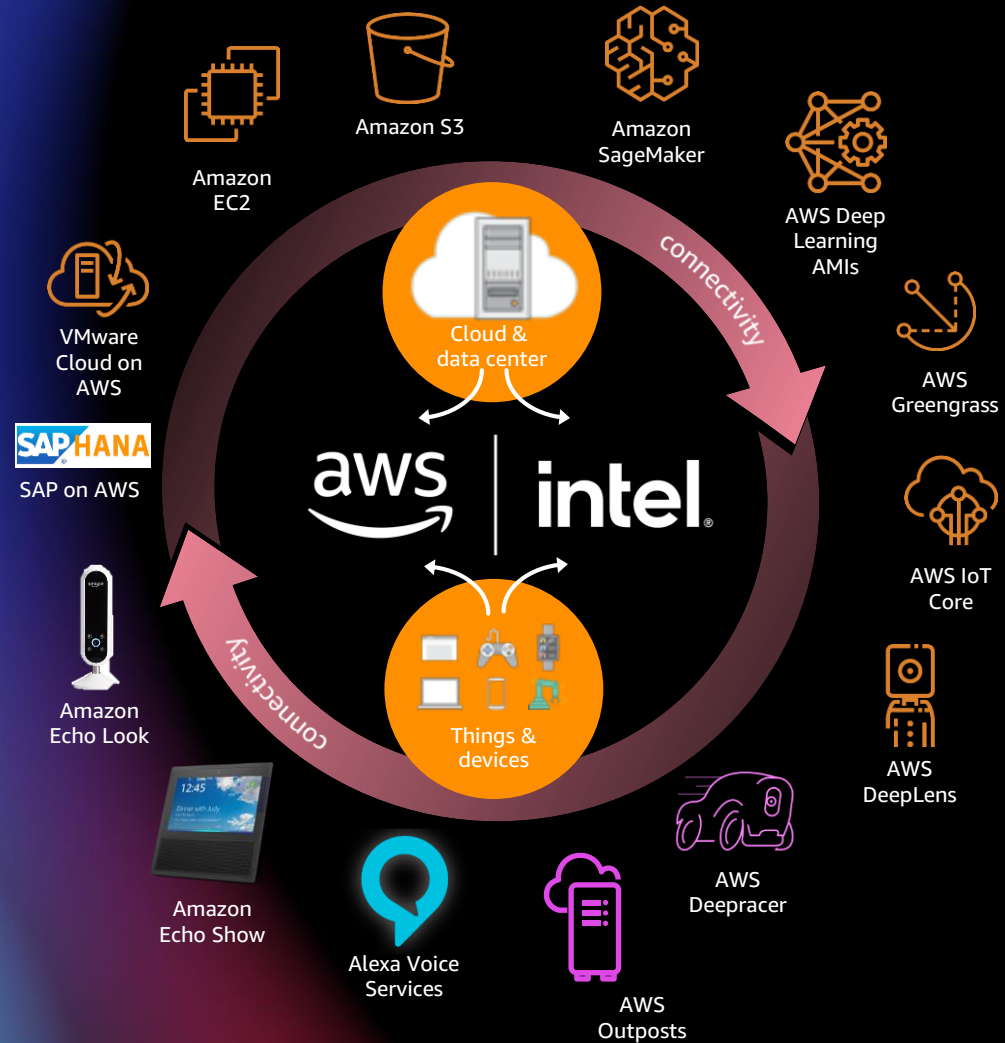
#1

Nearly all early
commercial vRAN
deployments are
running on Intel

50%+

of core network
workloads virtualized
in 2020, the majority
running on Intel
Xeon processors

What does Intel do with Amazon?



Greatest variety and availability to meet your global workload needs

aws | intel

250+ Intel instances

INSTANCES AVAILABLE FOR SALES FUNDING

Workloads	Instances	Category
High Performance Computing (HPC)	Gaudi EC2	High Frequency
Artificial Intelligence (AI) Machine Learning (ML) Big Data	Z1d, C5	Compute-Intensive
SAP on AWS	C5n, C5d, C5dn, M5, M5d, M5zn, M5dn	General Purpose
VMware Cloud on AWS	M5n, M5dn, R5, R5d, R5n, R5dn, R5b, R5bn	Memory-Intensive
Internet of Things (IoT)	HM1	Burstable
Strategic Migration	T3, T3n, I3en, X1, X1e	High I/O, Large Memory

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

Key Workloads

- Strategic migrations
- AI/ML
- HPC
- SAP
- 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



Highlights of the past year

Career Launcher Rapidly Scales Learning Portal during Pandemic

Intel & AWS partner to help serve >160,000 students in India within two months on AWS.¹

65x more parallel wildfire simulations

Intel & AWS partner with RONIN to help increase fire fighting effectiveness in Australia.⁴



AWS ParallelCluster

AWS as first CSP with verified Intel Select Solution²

Amazon EC2 M5zn instance – fastest Intel Xeon Scalable CPU in the Cloud⁵
Highest all-core turbo CPU performance with a frequency up to 4.5 GHz.

AWS announces AI instances with 40% better price/perf built on Habana Gaudi³

Intel's Habana & AWS co-engineer 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/>



Data analytics portfolio



Solutions
Solution Architects



Platforms



Finance



Healthcare



Energy



Industrial



Transport



Retail



Home



More...



Toolkits
App Developers

OpenVINO™ Toolkit

OpenVINO Toolkit for inference deployment on CPU, processor graphics, FPGA & VPU using TF, Caffe* & MXNet*

Deep Learning Developer Toolkit

Optimized inference deployment for all Intel® Movidius™ VPUs using TensorFlow* & Caffe*



Libraries
Data Scientists

MACHINE LEARNING LIBRARIES

Python

- [Scikit-learn](#)
- [Pandas](#)
- [NumPy](#)


R

- [Cart](#)
- [RandomForest](#)
- [E1071](#)

Distributed

- [MLLib \(on Spark\)](#)
- [Mahout](#)

DEEP LEARNING FRAMEWORKS



TensorFlow*



MXNet*




Caffe*



BigDL/Spark*



Caffe2*



PyTorch*



PaddlePaddle*



Foundation
Library Developers

ANALYTICS, MACHINE & DEEP LEARNING PRIMITIVES

Python

Intel distribution optimized for machine learning

DAAL

Intel® Data Analytics Acceleration Library (for machine learning)

MKL-DNN


Open-source deep neural network functions for CPU, processor graphics

cIDNN

FOUNDATION



ACCELERATORS



← Inference →



Hardware
IT System Architects



† Formerly the Intel® Computer Vision SDK

*Other names and brands may be claimed as the property of others.

All products, computer systems, dates, and figures are preliminary based on current expectations, and are subject to change without notice.

AWS data acceleration with Intel

**AWS C5 Deep Learning AMI
Optimized for Intel CPU
(Training)**

7.4x

**faster than training on the stock
Tensorflow 1.6 binaries**

**AWS C5 Deep Learning AMI
Optimized for Intel CPU
(Inference)**

12x

**faster than default configuration for NMT
Inference with MxNet**

**AWS SageMaker Machine Learning
Optimized for Intel CPU
(Training & Inference)**

10x

**Machine Learning Algorithms
optimized for IA CPU**



Amazon SageMaker



Amazon Rekognition




Amazon Forecast



Amazon Personalize

**AWS MARKETPLACE: Deep learning AMIs | Multiple containers with
OneDNN-optimized frameworks**

Customer programs with Intel



How to engage

Participate in an Intel/AWS workshops at the AWS Programs

Explore 16 Intel-optimized software libraries and frameworks on AWS Marketplace

Learn from other customers like FootAsylum, Thorn, Zignal Labs, Krispy Kreme, GE, Thomson Reuters, ASIC: <https://aws.amazon.com/solutions/case-studies/>

Connect with a partner – DataRobot, Data Bricks, IntellHQ, Intellify, Peak.AI, C3.AI, H2O.AI, Slalom

Success Stories



Analyzed 3 TB of Plant Breeding and Acclimatization Genomics Data in Hours to study the diversity of wheat.



Worked with educators to address the needs of Delhi's schools during Covid-19. Within two months a solution was scaled to serve over 160,000 students.



Australia's CSIRO research agency was able to increase by 60x the number of parallel wildfire simulations with 98% utilization of large Amazon EC2 C5-based instances.



Researchers democratized data and drove ML models on genomic sequencing for endangered species.



Delivered more than 100 dashboards using anonymized government and public COVID-19 data points to prevent disease transmission.



AI solution with the Intel® Distribution of OpenVINO™ toolkit
Helps identify, detect, and respond in real time to hazards along Singapore's coastline.

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

AI/ML Image Analysis

“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

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.

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](#)



See link for details. Results may vary.

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

AWS Customer References



and



8,400%

return on ad spend

The AI system runs on Amazon EC2 C5 instances, powered by Intel® Xeon® Scalable processors.

AI/ML, Big Data
Online Marketing

"Intel gives us the ability to scale, so we can give our AI solutions the computational power they need. This is critical for our business."

Mylo Portas,
Head of Retail, Peak AI

Boost Sales with AI-based Customer Personalization

Expands to serve >160,000 students within two months on AWS.

Products and Solutions

[Peak.ai](#)
[Amazon EMR – Big Data Platform using Apache Spark](#)
[Intel® Xeon® Scalable processors](#)

Peak, an AWS Advanced Consulting Partner, implements AWS-powered solutions that enabled Footasylum to deliver advanced personalization in marketing communications via AI-powered customer segmentation.

Need: Deliver a personalized, connected retail experience to increase return on advertising and revenue via email marketing.

Solution: Customer purchasing data is stored in Amazon S3 and processed using Apache Spark running on Amazon EMR. The AI system, powered by Intel® Xeon® Scalable processors, runs on Amazon EC2 C5 instances to train its ML models and run predictions based on the models.

Value: Predictions and prescriptions target groups of customers in a way to achieve an 8,400% return on ad spend (ROAS), and a 28% revenue boost via email marketing.

Industry

AI-ML/Online
Marketing

Organization Size

1-50

Country

United States

Partners

[AWS](#)
[Footasylum](#)

Learn more

[Partner Story](#)



See link for details. Results may vary.

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Data in healthcare



Introducing Habana Gaudi-based Instances

COMING IN 2021

ML Training Powered by New Habana Gaudi Processors from Intel



New EC2 instances built specifically for ML training and powered by up to 8 new Habana Gaudi processors from Intel

Will deliver up to 40% lower cost-to-train deep learning models over GPU-based instances

Will allow customers to iterate and train models more frequently

Benefit from full stack of Amazon EC2 services—DL AMIs, DLC for containerized applications, ultimately SageMaker


Developers can implement Gaudi-based instances via AWS ECS and EKS for containerized applications

Will support common frameworks like TensorFlow and PyTorch

Wide range of ML workloads for applications including, NLP, image classification, object detection, recommendation systems

For efficient scaling across multiple Gaudi-based EC2 Instances, support for AWS Elastic Fabric Adapter

Introducing

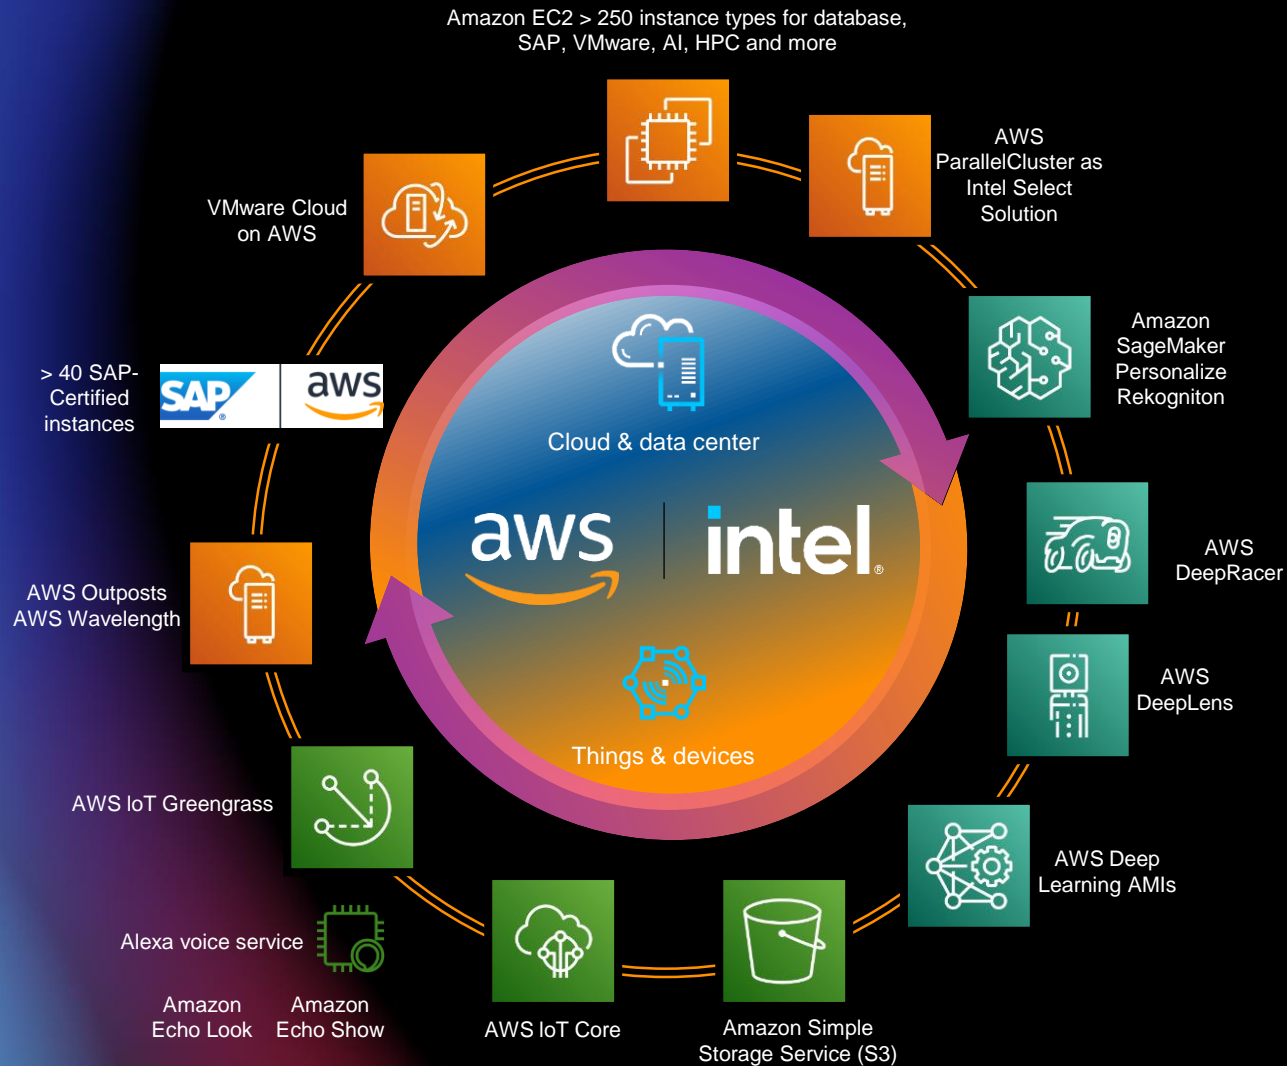


Introducing the GAUDI® AI Processor – powering better price-performance in upcoming Amazon EC2 Instances

Eitan Medina
CBO, Habana Labs



Summary



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

Thank you!

Akanksha Balani

AWS APJ Alliance head @ Intel

Global AI and HPC GTM Lead



Thank you for attending AWS Innovate – Data 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