



AWS for virtually every application

Pragya Pandey (she/her)
Principal Product Marketing Manager
Amazon Web Services

Agenda


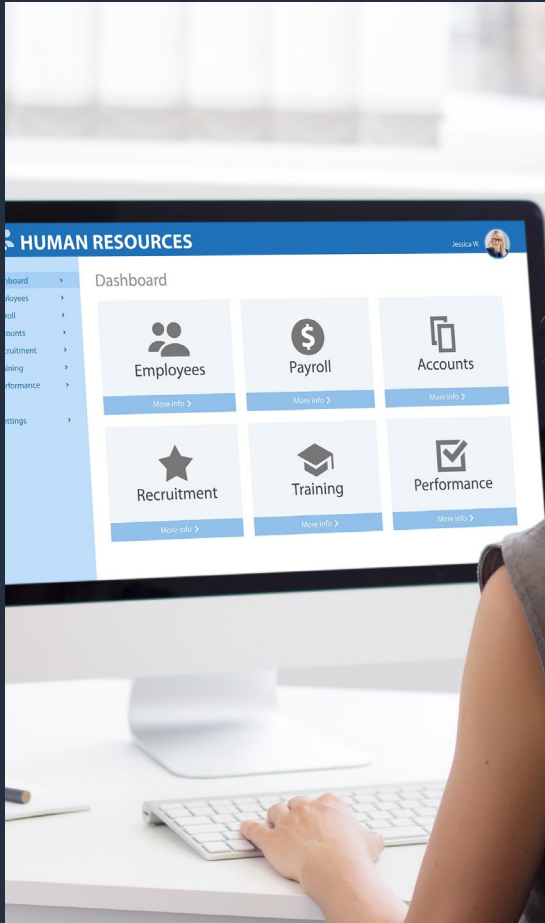
Diverse set of applications supporting businesses

How AWS supports a broad range of applications

Customers leveraging AWS for their applications


Next steps

Applications power businesses of all kinds



The way applications are built & run has evolved

BEFORE:



Monolithic design, on-premises deployment, limited capacity

NOW:



Microservices, cloud deployment, virtually infinite capacity

It's still **early days**



Our vision is to be the **best place**
to run your applications of **today** and
build your applications of **tomorrow**

Supporting virtually every application



Enterprise apps



**High performance
computing**



Modern applications



Low latency apps



IoT apps



Hybrid cloud

global footprint

cost

customer obsession

elasticity

sustainability

breadth of services

Why do organizations choose **AWS**?

agility

performance

reliability

productivity

innovation

security

community





Most secure and reliable cloud



Best performance and value from your infrastructure



Most capabilities for all your applications



AWS infrastructure and services wherever you need it



Most secure and reliable cloud



Best performance and value from your infrastructure



Most capabilities for all your applications



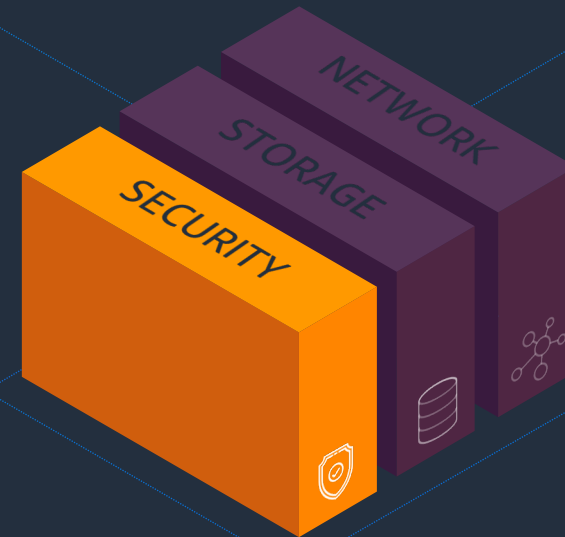
AWS infrastructure and services wherever you need it

Built-in security with the AWS Nitro System

OFFERING THE BEST SECURITY, PERFORMANCE, AND INNOVATION IN THE CLOUD



NITRO SYSTEM



NITRO



Security capabilities in our storage services

ACCESS, IDENTITY MANAGEMENT




Block public access to all objects at bucket or account level.
Use IAM policies to manage access to EFS in a cloud-native way, securely

ENCRYPTION



Encryption for data in transit, and data at rest across storage portfolio

AUDIT AND MONITOR



Use tools to like CloudWatch and Amazon S3 Lens to audit and monitor continuously

Security capabilities in our storage services

ACCESS, IDENTITY MANAGEMENT




Block public access to all objects at bucket or account level.
Use IAM policies to manage access to EFS in a cloud-native way, securely

ENCRYPTION



Encryption for data in transit, and data at rest across storage portfolio

AUDIT AND MONITOR



Use tools to like CloudWatch and Amazon S3 Lens to audit and monitor continuously

Building data resiliency for applications

ALWAYS PROTECTED

ROLLBACK




Quick rollback with file and object versioning, or with file and object replication

BACKUP



Backup using a fully managed service with AWS Backup. Backup data on Amazon EBS using EBS Snapshots

DISASTER RECOVERY



Recover on-premises applications to AWS, and cloud applications across AWS Regions with AWS Elastic Disaster Recovery

Zero Trust

A security model centered on the idea that access to data should not be solely made based on network location



NEW!

IN PREVIEW

AWS Verified Access

SECURE ACCESS TO CORPORATE APPLICATIONS WITHOUT A VPN



Improve security posture

Evaluates each user request in real-time using identity and device posture



Simplify operations

Onboard applications using a few clicks, manage all policies centrally



Increased mobility

Users access applications with a web-browser, without any additional agents

**Uniquely designed for
high network resilience**



The M10 logo consists of the letters "M10" in a bold, white, sans-serif font, centered within a solid blue rectangular background.

M10

"M10 Networks, Inc develops and deploys their M10 Ledger Platform, a service for developing and distributing central bank digital currencies and tokenized regulated liabilities, on AWS. The Ledger Platform uses AWS Nitro Enclaves to perform signature verification and cryptographic re-signing of batches of transactions. Using AWS Nitro Enclaves on AWS latest M6i instances, M10 is able to deliver a performant and cost effective solution for the digital currency market."

- Sascha Wise

M10 Founding Engineer



Most secure and reliable cloud



Best performance and value from your infrastructure



Most capabilities for all your applications





AWS infrastructure and services wherever you need it

Better performance with the Nitro System

Nitro performance for real-world workloads

AMAZON EC2 INSTANCES CAN DELIVER OVER 15% HIGHER THROUGHPUT PERFORMANCE





AWS Graviton

Tens of thousands of AWS customers use Graviton

Over 100 Graviton-based EC2 instances

Graviton3 has **25% better** compute performance than Graviton2

Up to **60% less** energy used by Graviton3-based EC2 instances

“ We have now found Graviton3 C7g instances to be 40% faster than the Graviton2 C6gn instances for those same simulations. ”

Pat Symonds

CTO at Formula 1 Management



DIRECTV

“ We found Graviton3 and the c7g series fit the sweet spot for price-to-performance, reliability, security, and availability for our golang-based microservice fleet. We have better performance with lower costs – 20% lower cost with p95 latency improvements up to 50% and cpu improvements up to 35%. ”

Jonathan Tronson




VP Software Engineering, DIRECTV





Supporting ML workloads

LARGEST SELECTION OF ML INSTANCES IN THE CLOUD

Training



Inference



Amazon EC2 Trn1 instances

HIGHEST PERFORMANCE AT LOWEST COST FOR TRAINING DEEP LEARNING MODELS

POWERED BY AWS TRAINIUM



16 Trainium accelerators with 512 GiB of HBM2 high speed memory, 1600 Gbps of networking throughput (Trn1n), and 8 TB of local NVMe storage

DEEP LEARNING PERFORMANCE



Lower time-to-train when compared to the latest comparable Amazon EC2 instances

DEPLOY IN AN ULTRACLUSTER



Deployable in Amazon EC2 UltraClusters, a supercomputer in the cloud with tens of thousands of Trainium accelerators with petabit scale networking bandwidth

Amazon EC2 Inf2 instances

NEW!

HIGH PERFORMANCE, ENERGY EFFICIENT, AND LOWEST COST INFERENCE

NOW AVAILABLE

POWERED BY AWS INFERENTIA



Optimized to deploy 100B+
parameter models at scale

12 Inferentia2, 384GB HBM2e

4x higher throughput and 10x
lower latency than Inf1

DEEP LEARNING PERFORMANCE



70% better price performance and
50% better performance/watt
than comparable Amazon EC2
instances

OPTIMIZED FOR DISTRIBUTED INFERENCE



192 GB/s of inter-accelerator
connectivity for efficient
distributed inference



Reduce total cost of ownership by

57%

With AWS serverless services



Building a culture of cost optimization

HOW DOES AWS FOCUS ON COST REDUCTION

MAKE IT CULTURAL




At Amazon, our “frugality” leadership principal ensures that we consider the cost implications of every design or solution

TOOLS AND METRICS



Use the tools and metrics available to you to measure and control usage within the organization

ACCOUNTABILITY



Establish a mechanism to ensure teams identify cost savings opportunities and then track these closely over time

Options and tools for cost optimization


LEVERAGE OPTIONS AND TOOLS AVAILABLE FOR COST REDUCTION

CHOICE




Leverage purchase options for compute and choose from comprehensive set of compute instances and storage classes

EFFICIENCY



Use compute capacity based on how your workloads scale up or down. Improve storage efficiency through compression and deduplication

OPTIMIZATION



Right-size workloads with machine learning-based recommendations. Use intelligent tiering for object and file storage based on usage pattern

Options and tools for cost optimization


LEVERAGE OPTIONS AND TOOLS AVAILABLE FOR COST REDUCTION

CHOICE




Leverage purchase options for compute and choose from comprehensive set of compute instances and storage classes

EFFICIENCY



Use compute capacity based on how your workloads scale up or down. Improve storage efficiency through compression and deduplication

OPTIMIZATION



Right-size workloads with machine learning-based recommendations. Use intelligent tiering for object and file storage based on usage pattern

Customers saving costs with AWS



\$1 million monthly savings with Spot Instances



30% cost savings with EC2 instances



\$63.5 millions saved in storage costs



Most secure and reliable cloud



Best performance and value from your infrastructure



Most capabilities for all your applications



AWS infrastructure and services wherever you need it

The **broadest** and **deepest**
services and capabilities

200+
fully featured services

200+
fully featured services

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, Apple)
Fast processors
(up to 4.5 GHz)
High memory footprint
(up to 24 TiB)
Instance storage
(HDD and NVMe)
Accelerated computing
(GPUs, ASICs, Video, FPGA)
Networking
(up to 1,600 Gbps)
Bare metal
Size
(Nano to 112xlarge)




OPTIONS

Amazon EBS
Amazon Elastic Inference



MORE THAN
600
INSTANCE TYPES
for virtually every
workload and
business need

Compute for virtually every application



Amazon EC2
Infrastructure
-as-a-service



Amazon ECS/EKS
Container-management
-as-a-service



AWS Fargate
Serverless
containers



AWS Lambda
Serverless
functions


Storage for virtually every application

Object



Amazon S3
and Amazon
S3 Glacier

Block



Amazon
Elastic
Block Store



Amazon
Elastic File
System



NetApp
ONTAP



Windows
File Server



Lustre



OpenZFS



Amazon
File Cache

Storage for virtually every application

Object



Amazon S3
and Amazon
S3 Glacier

Block



Amazon
Elastic
Block Store



Amazon
Elastic File
System



NetApp
ONTAP



Windows
File Server



Lustre



OpenZFS



Amazon
File Cache

Making it easier and faster for builders



AWS Amplify

Build full-stack web and mobile apps in hours. Easy to start, easy to scale

NEW!



Amazon CodeCatalyst (Preview)

Unified software development service to quickly build and deliver applications on AWS



Amazon CodeWhisperer

Build applications faster with the ML-powered coding companion



Most secure and reliable cloud



Best performance and value from your infrastructure



Most capabilities for all your applications



AWS infrastructure and services wherever you need it

Every application cannot move to the cloud



Low latency



Local data processing



Data residency



Migration
and modernization

Bringing cloud capabilities where customers need it

REGIONS



METRO AREAS
& TELCO
NETWORKS



ON-PREMISES



REMOTE & LIMITED
CONNECTIVITY
LOCATIONS



SMART
DEVICES





31 Geographical Regions

99 Availability Zones





● 32 Local Zones

● 21 Local Zones announced

Vancouver

Toronto

Querétaro

Bogotá

Lima

Rio de Janeiro

Santiago

Buenos Aires

Lisbon

Amsterdam

Brussels

Oslo

Berlin

Munich

Lagos

Nairobi

Johannesburg

Helsinki

Copenhagen

Hamburg

Warsaw

Prague

Vienna

Athens

Muscat

Bengaluru

Delhi

Chennai

Kolkata

Bangkok

Perth

Hanoi

Taipei

Manila

Brisbane

Auckland



Mindbody wanted to migrate a large portfolio of interdependent applications running in on-premises data centers to the cloud

Established a low-latency hybrid environment between their on-premises installations and AWS Local Zones

Migrated applications incrementally, drastically simplifying the migration process and enabling ongoing hybrid deployments





AWS Outposts family



Outposts rack

Bring the same AWS APIs, services, and features to virtually any data center or co-location space



Outposts servers

Run Outposts in locations with limited space or smaller capacity requirements



AWS for IoT

CHOICES FOR YOUR UNIQUE NEEDS



Foundational services

A broad and deep set of cloud services to help customers create custom IoT applications with the highest security, flexibility, and control



Purpose-built IoT services

Services that accelerate and simplify the time to build secure, reliable, and scalable IoT applications for specific industries and use cases



Vetted solutions

Ready-to-deploy code and configurations and customizable architectural guidance, built by AWS and AWS Partners

Delivering AWS wherever you need it





Most secure and reliable cloud



Best performance and value from your infrastructure



Most capabilities for all your applications



AWS infrastructure and services wherever you need it

Solving the world's hardest problems



High Performance Computing on AWS

Performance
at scale

AWS Nitro
System

Amazon EC2

Elastic Fabric
Adapter

Amazon FSx
for Lustre /
Amazon FSx
for OpenZFS

Job and
cluster
management

AWS Batch

AWS
ParallelCluster

NICE
DCV

Every type of app



Our vision is to be the best place to run your applications of today and build your applications of tomorrow

Next steps

Visit the AWS for Every Application Webpage

aws.amazon.com/aws-for-every-application/

Explore AWS migration solutions

aws.amazon.com/cloud-migration/



Thank you!