
AWS
re:Invent
NOV. 28 – DEC. 2, 2022 | LAS VEGAS, NV

CMP320

Accelerating semiconductor design, simulation, and verification on AWS

Mark Galbraith

VP Productivity Engineering
Arm Limited

Bill Cornwell

Vice President, CE – ASIC & Processor
Marvell Technology

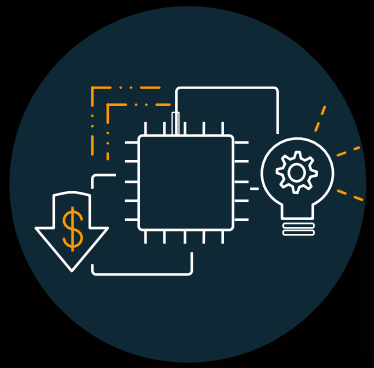
David Pellerin

Director of Industry Solutions, HiTech, Electronics & Semiconductor
Amazon Web Services

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

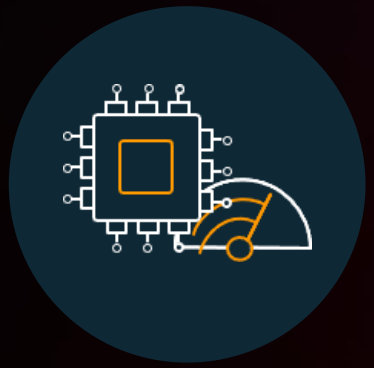
Industry Trends and Challenges



Supply Chain Constraints and Shortages



Increased Multi-Party ASIC/SoC Development



Balancing Power, Performance, Area, Cost (PPAC)



Engineering Staffing Shortages



Sustainability and ESH/ESG

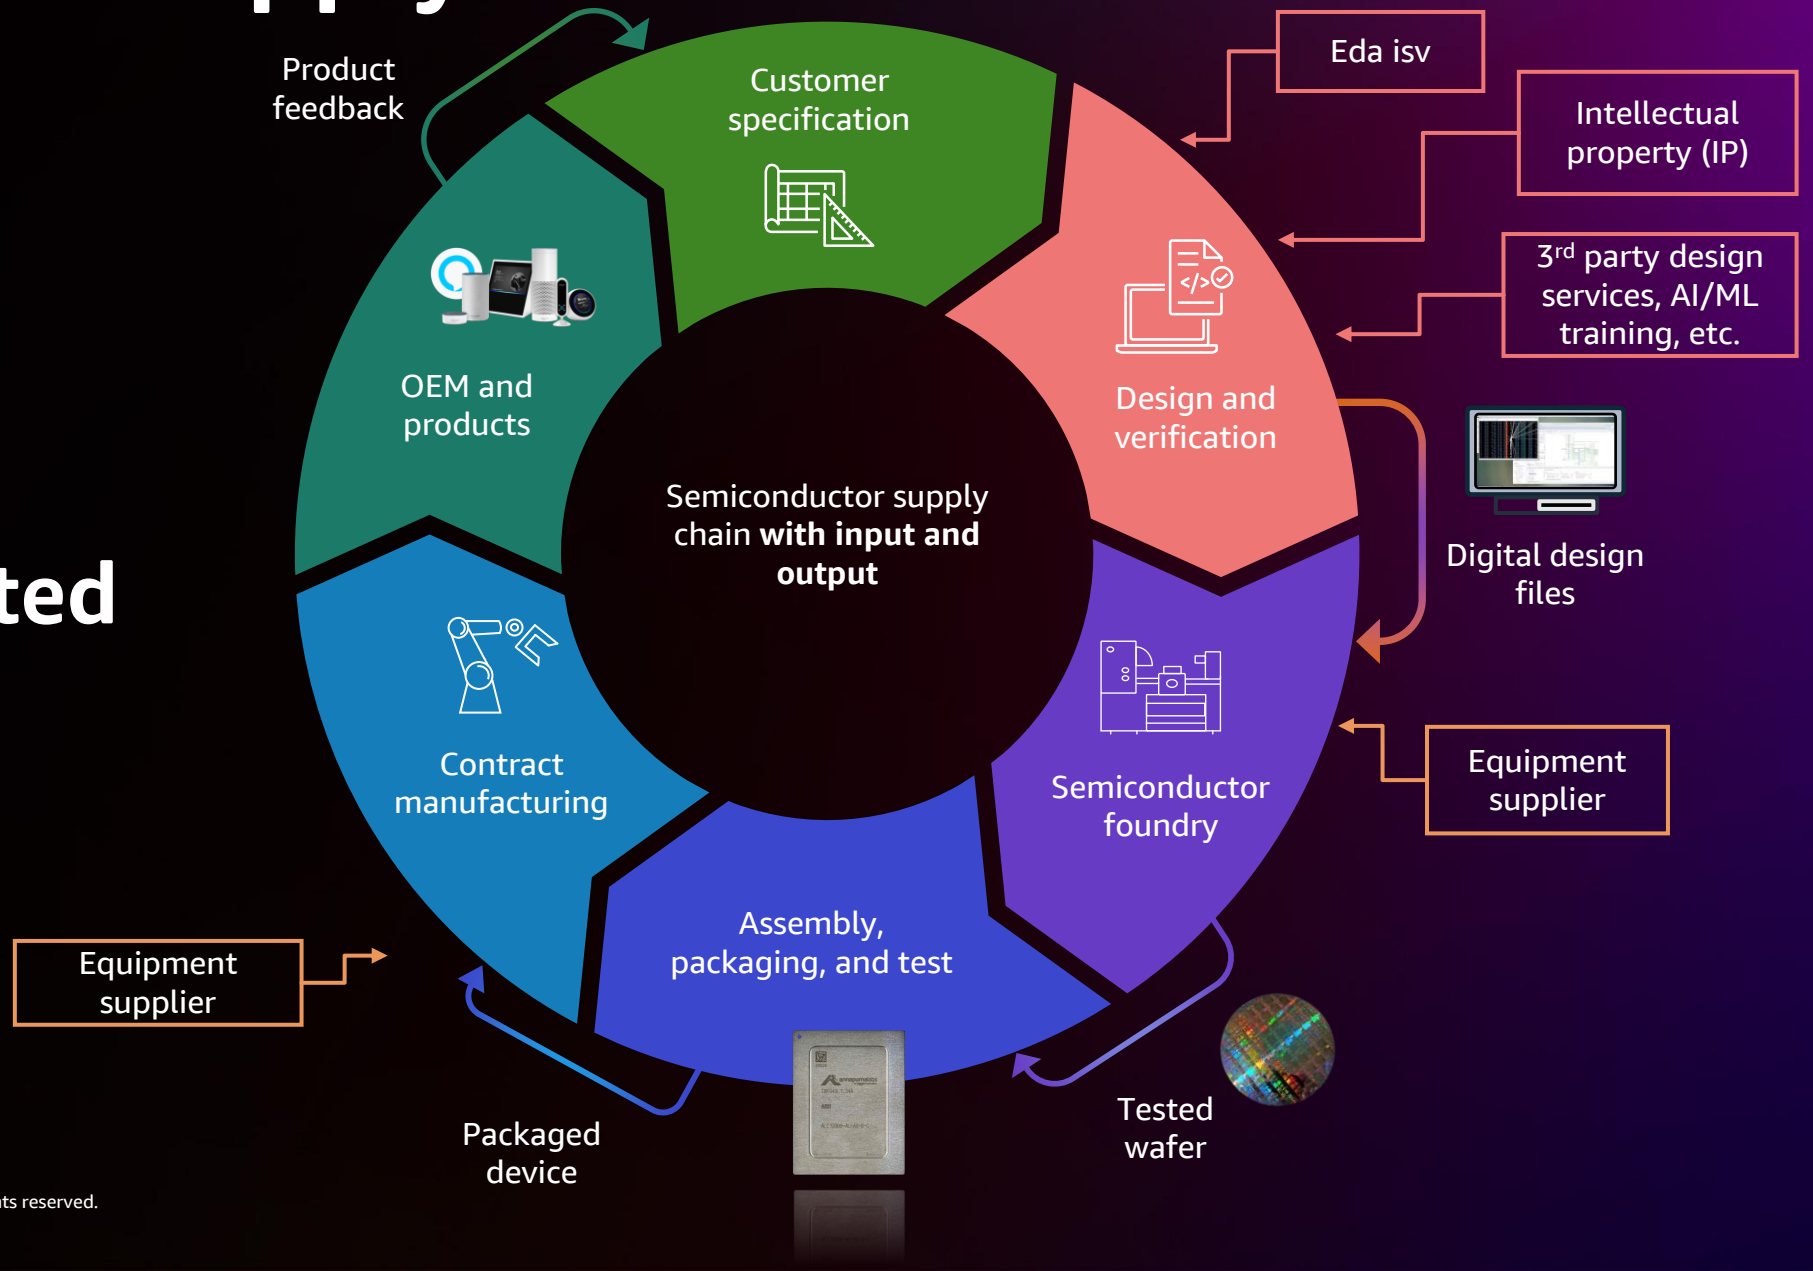
Electronic Design Automation, deployed on AWS Cloud, is helping to address these challenges across the entire semiconductor supply chain

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Semiconductor supply chain

Opportunities for **cloud-accelerated innovation**



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

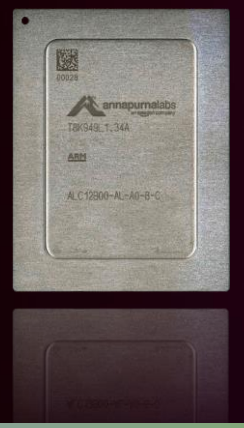
Amazon is a Fabless Semiconductor Company

We design our own silicon devices, and we source from a global supply chain

Amazon develops and uses semiconductor devices for

- AWS data center infrastructure
- Amazon fulfillment centers
- Consumer devices
- Robotics and AI
- Space and satellite infrastructure
- Autonomous vehicles
- And more





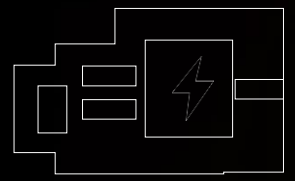


We value our semiconductor and EDA industry partnerships

aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

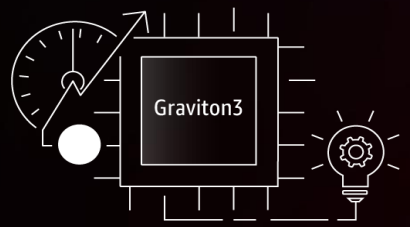
Silicon innovation at AWS

annapurnalabs
an amazon company



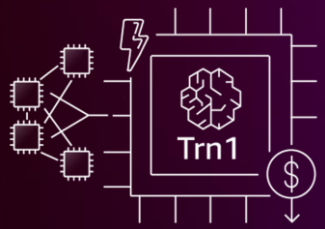
AWS Nitro System
AWS Nitro SSD

Hypervisor, network, storage, and security



AWS Graviton2
AWS Graviton3

Powerful and efficient, modern applications



AWS Inferentia
AWS Trainium

Machine learning, hardware, and software at scale

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Why EDA on AWS?

Innovate faster

Prototype, design, and verify complex systems-on-chip, using scalable cloud resources for Electronic Design Automation (EDA.)

Collaborate better

Work seamlessly and securely with third-party partners including IP providers, EDA software vendors, and manufacturing service providers (foundries, OSATs, contract, and original device manufacturers).

Reduce risk

Advanced silicon and system verification is hard, and getting harder. Mistakes can cost millions if not billions of dollars for a larger companies.

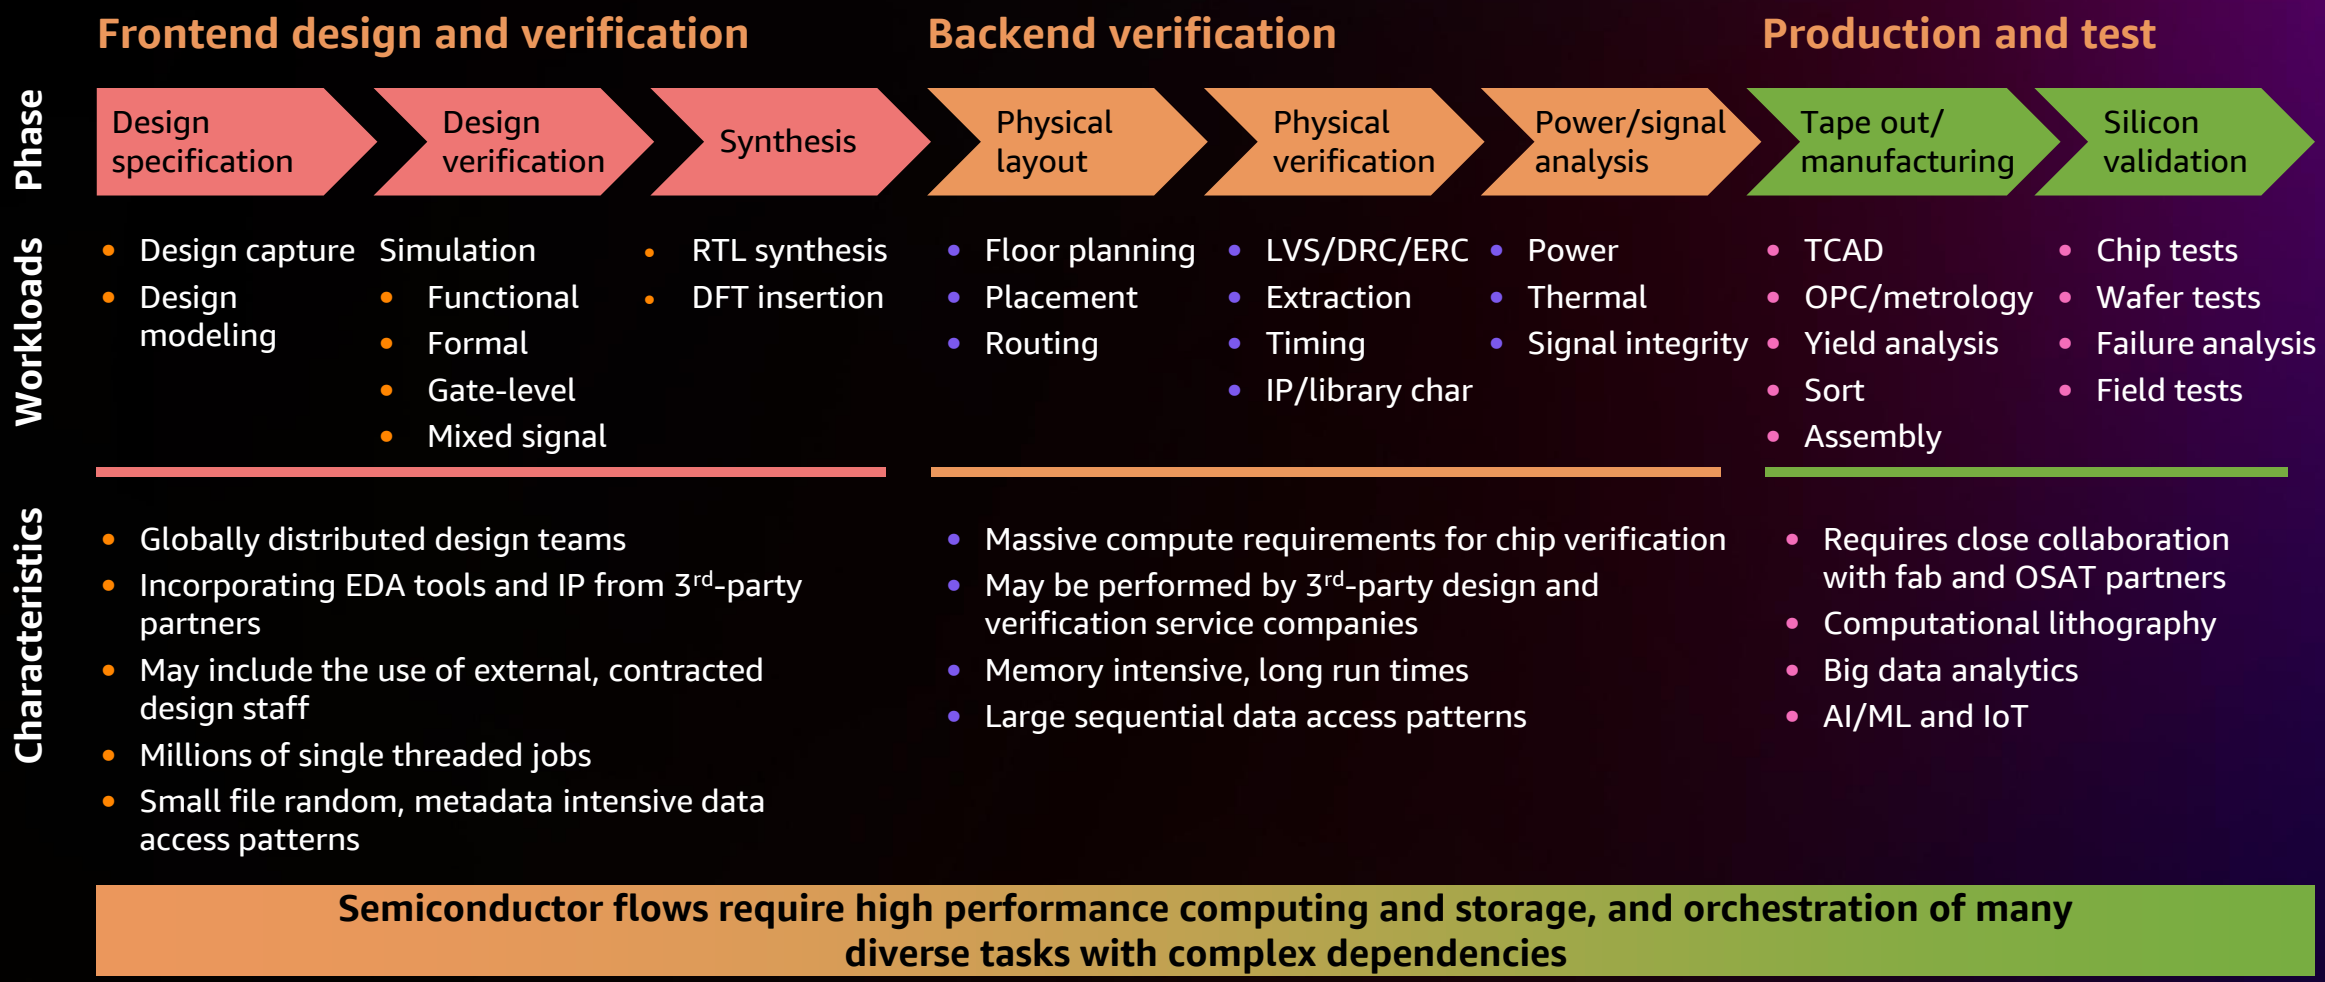
Reduce cost

Stop wasting CAPEX on IT, and stop wasting valuable engineering time.

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

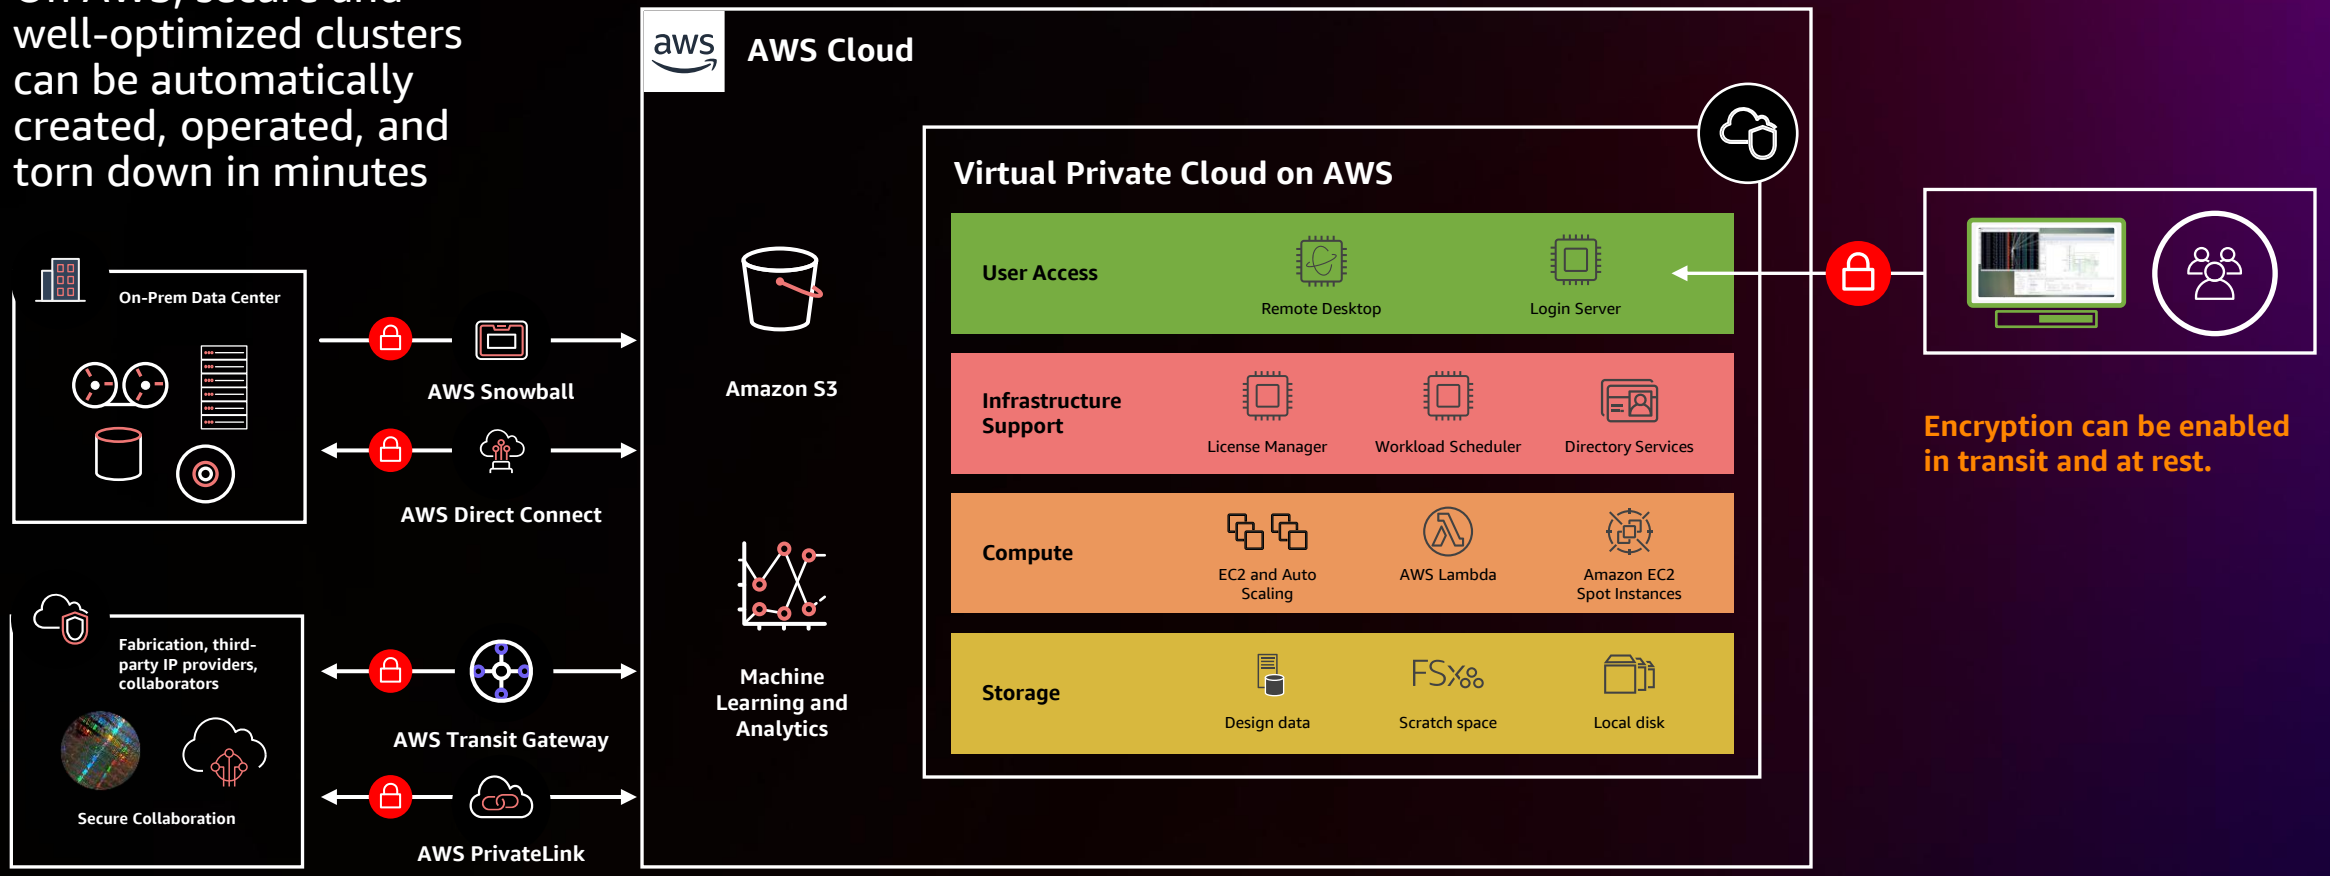
Semiconductor design flow



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Semiconductor design infrastructure on AWS

On AWS, secure and well-optimized clusters can be automatically created, operated, and torn down in minutes



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Broadest choice of processors and accelerators



INTEL Xeon
Scalable processors



AMD EPYC
processors



AWS Graviton
processors



Latest Tensor
Core GPUs



FPGAs for custom
hardware acceleration



Radeon Pro
graphics Processors



AWS
Inferentia/Trainium

aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Amazon EC2 Instances for EDA workloads

Latest Intel Ice Lake based instances

(R6i/X2idn/X2iedn)

- Up to 3.5 GHz all core turbo
- R6i instances offer up to 15% better compute price performance over equivalent Gen5 instances
- X2idn/X2iedn instances offer up to 50% better price performance than X1 instances
- Offer up to 1 TB/2 TB/4 TB of memory

High-frequency instances

(M5zn/z1d/X2iezn)

- Up to 4.5 GHz all core turbo, fastest in the cloud
- M5zn instances offer up to 45% higher compute performance over EC2 M5 instances
- X2iezn instances offer up to 55% better price performance than X1e instances
- Offer up to 192 GB/384 GB/1.5 TB of memory

AWS Graviton based instances

(R6g/R6gd/X2gd)

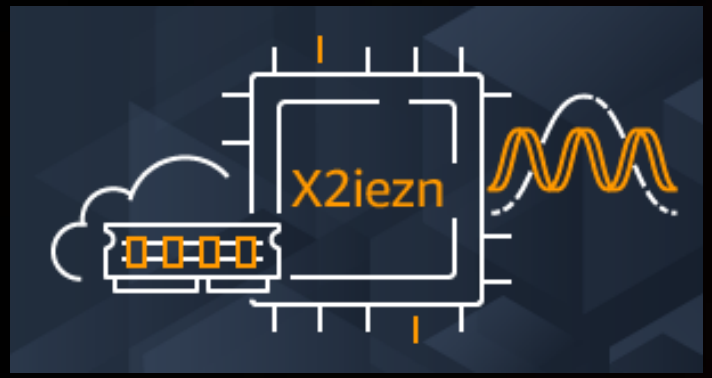
- Up to 40% better price/performance over R5 instances
- Offer up to 512 GB/1 TB of memory
- Provide the lowest cost per GiB of memory in Amazon EC2

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Amazon EC2 X2iezn instances

High frequency processing for electronic design automation



Up to 4.5 GHz all core turbo frequency, the fastest Intel Xeon Scalable processor in the cloud

2x more performance per vCPU compared to X1e to reduce license costs by up to 50%

X2iezn instances are a great fit for backend EDA workloads requiring both high performance and high memory

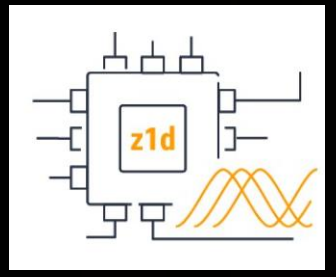
aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Amazon EC2 M5zn/z1d instances

High frequency processing for electronic design automation





Up to 4.5 GHz all core turbo frequency, the fastest Intel Xeon Scalable processor in the cloud

Reduce software license costs by right-sizing workloads on fewer vCPUs with 45% better price performance than M5 instances

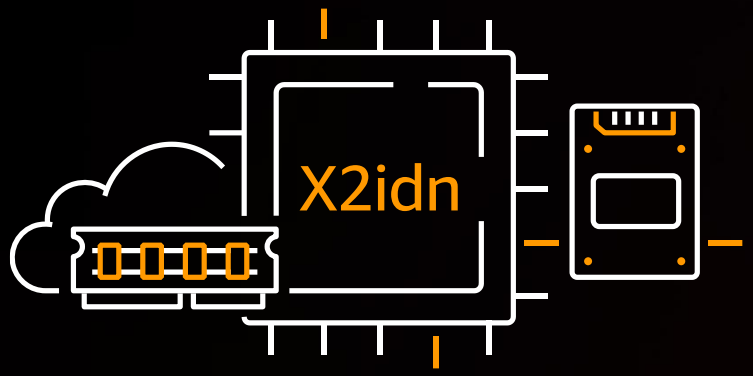
Great fit for frontend EDA workloads

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Amazon EC2 X2idn/X2iedn instances

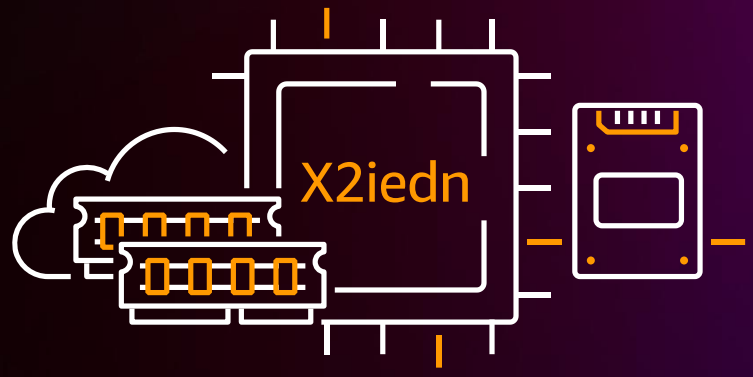
Next generation memory optimized instances



Up to 64 CPU cores, 2 TiB memory, 80 Gbps Amazon EBS, and 100 Gbps networking

Up to 50% better price performance than previous generation instances for memory intensive workloads

Ideal for in-memory databases, big data processing engines, and electronic design automation (EDA)



Up to 64 CPU cores, 4 TiB memory, 80 Gbps Amazon EBS, and 100 Gbps networking

Up to 50% better price performance than previous generation instances for memory intensive workloads

Ideal for large in-memory databases (SAP HANA), in-memory analytics, and high-performance databases

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

ARM-based instances: M6g, C6g, R6g, C6gn, X2gd

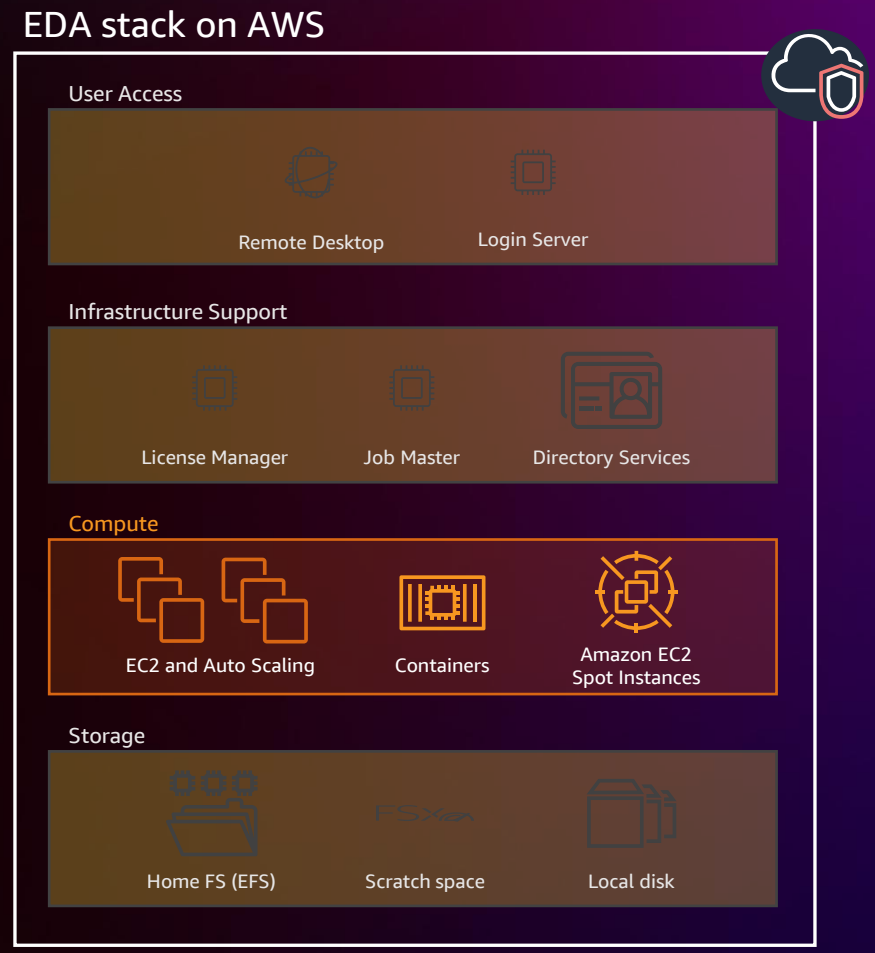
Powered by AWS Graviton processors

Improved price performance

- 64-bit Arm Neoverse N1 cores with AWS designed silicon
- Up to 64 physical cores
- Up to 512 GiB DDR4 memory
- Enhanced networking, up to 100 Gbps throughput
- Up to 40% price/performance advantage over x86 gen 5
- M6g: 4GB/core, C6g/C6gn: 2GB/core
- R6g: 8 GB/core, X2gd: 16 GB/core
- X2gd: 16 GB/core
- Tested and validated for major EDA applications

Featuring AWS Graviton Processor

annapurnalabs an amazon company

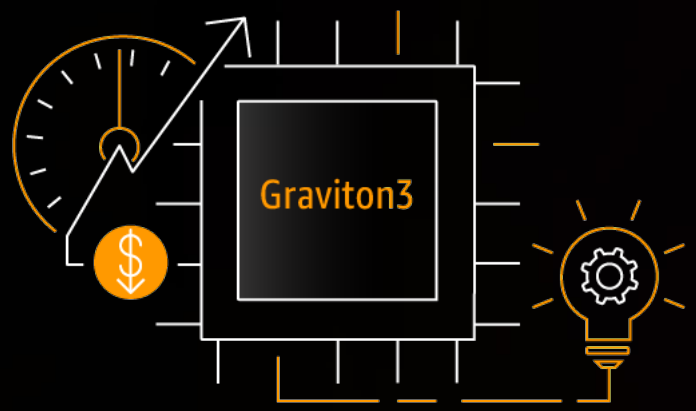


aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

AWS Graviton3 and Amazon EC2 C7g instances

Enabling the best price performance for workloads in Amazon EC2



Up to 25% better performance compared to Graviton2

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

First in the cloud to feature DDR5 memory

60% more energy efficient over comparable EC2 instances

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

https://aws.amazon.com/blogs/aws/new-amazon-ec2-c7g-instances-powered-by-aws-graviton3-processors/

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

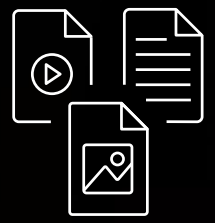
Solutions library for semiconductor EDA



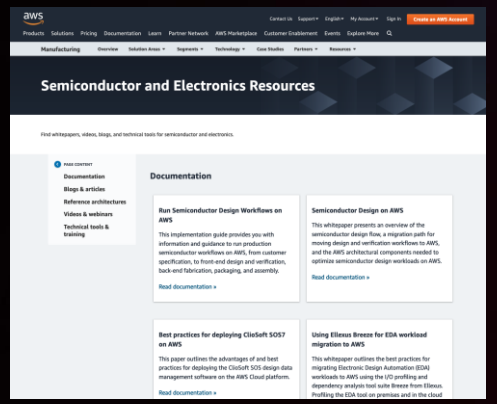
https://aws.amazon.com/solutions/semiconductor-electronics/electronic-design-automation/

aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

AWS resources for semiconductor design

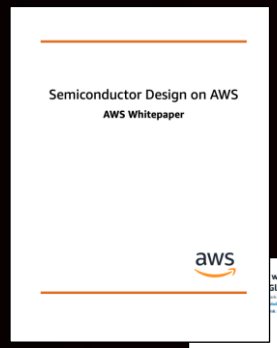


https://aws.com/semiconductor



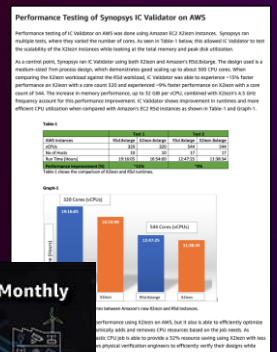
Under “Resources” link

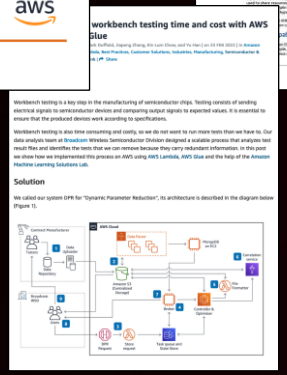
- White papers
- Blogs and articles
- Reference architectures
- Videos and webinars
- Technical tools and training

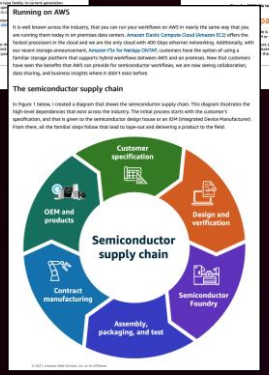














© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

CMP320

Accelerating semiconductor design, verification, and implementation

Mark Galbraith

VP of Productivity Engineering
Arm

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

70% of the world's population use Arm processor technology

Semiconductor IP & Solutions Business

arm

230bn+
Arm-based chips shipped to-date

29bn+
Arm-based chips shipped in 2021

Mobile Revolution



To Cloud Disruption



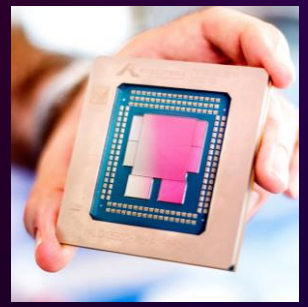
AWS (Annapurna Labs) is an Arm IP customer



AWS Nitro (Cortex-A72)



AWS Graviton2 (Neoverse N1)

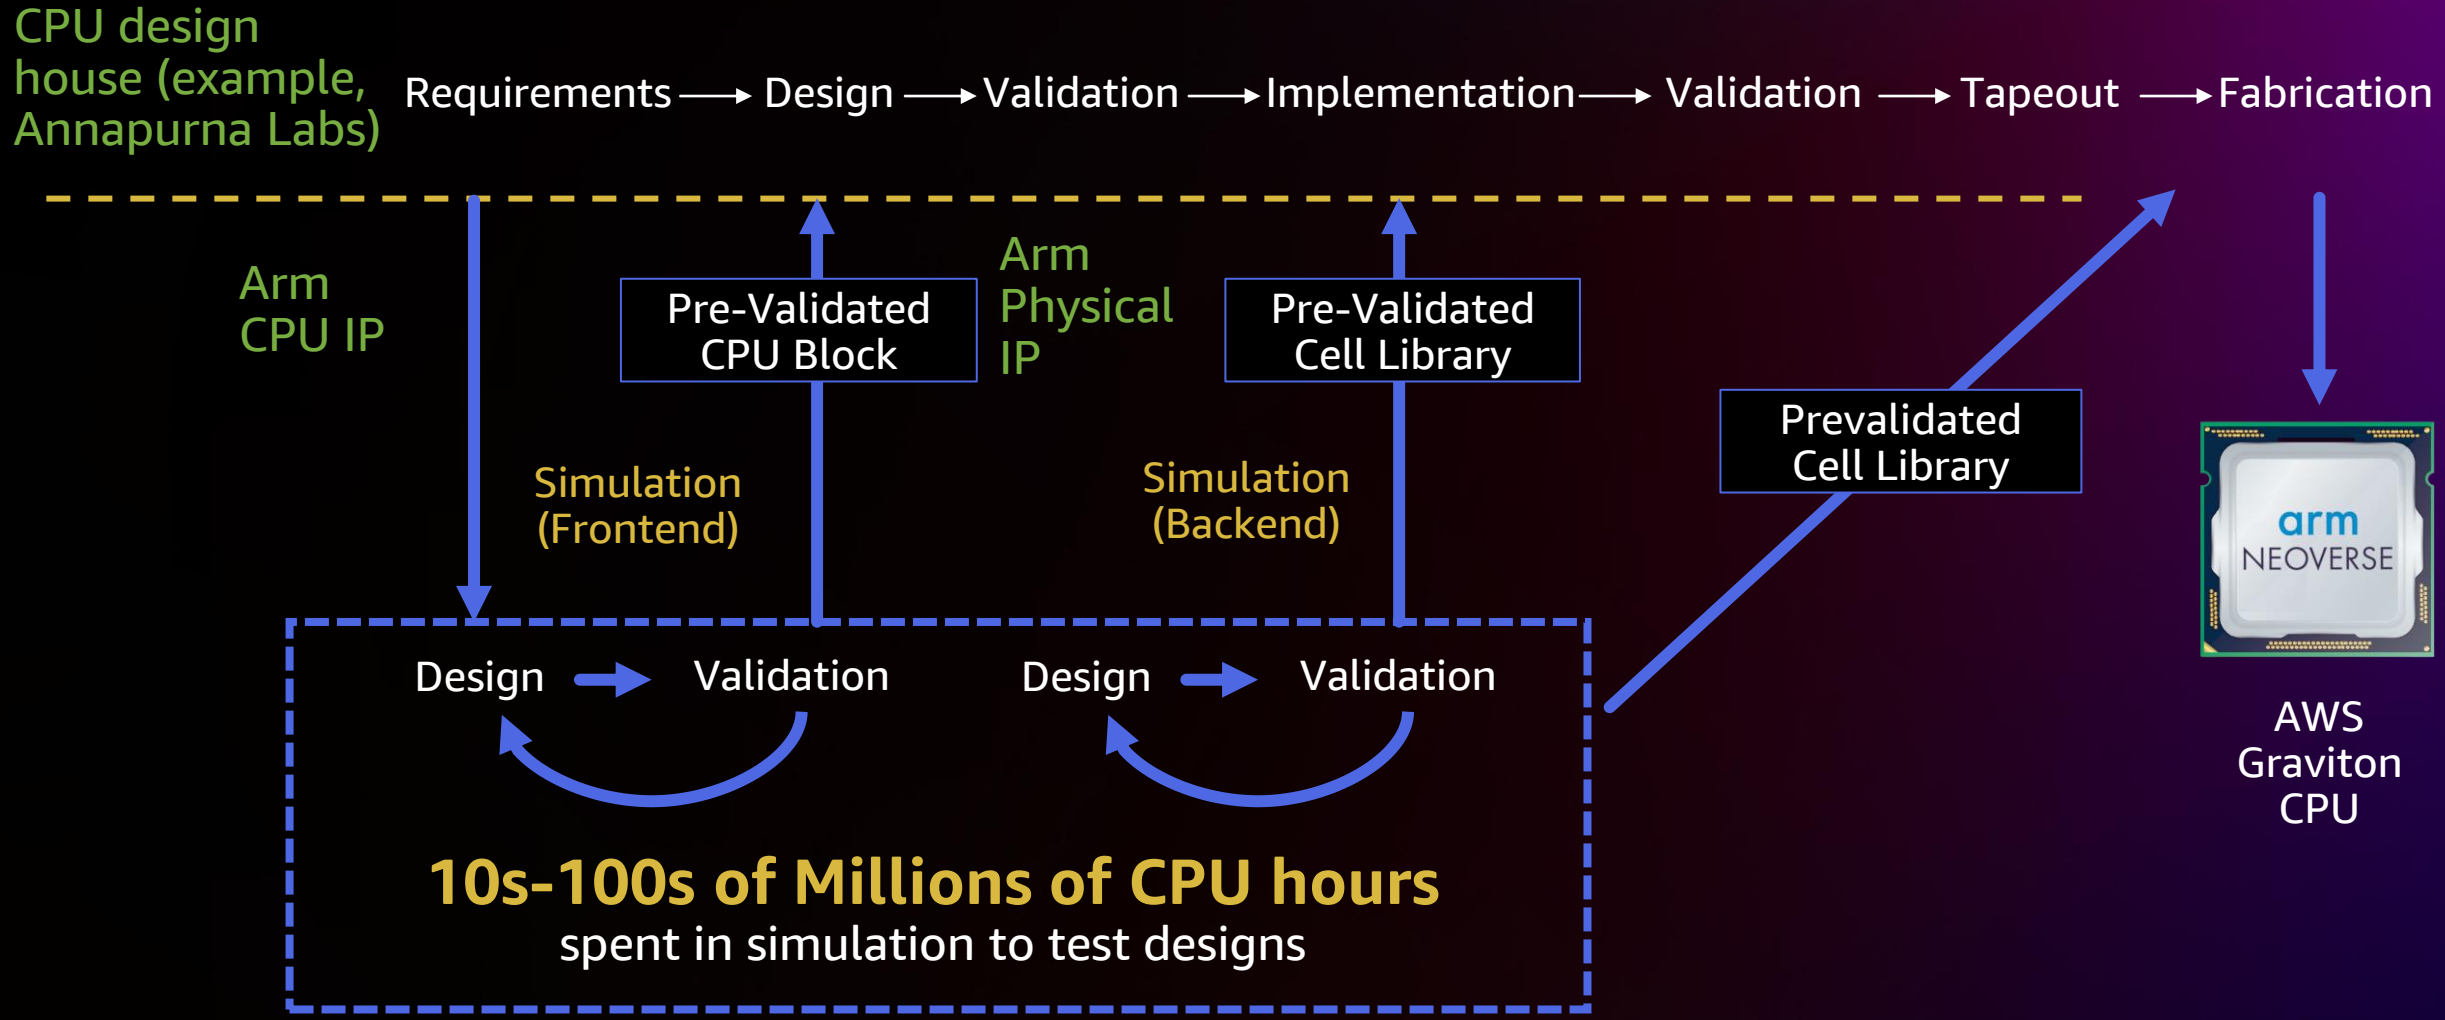


AWS Graviton3 (Neoverse V1)

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Chip Creation

Making a chip can cost $MM – design has to be right the first time



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Why Cloud for Arm IP Design and Verification?







Hardware efficiency

Scaling efficiency

Faster quality

Improve time-to-market

Maximize productivity

Reduce TCO

aws

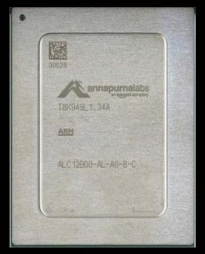
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

AWS Graviton's Expanding EC2 Portfolio

Arm-based instance types deliver improved performance, price, and sustainability

AWS Graviton2

- 40% Price/Performance[1]
- 3.45x Performance/Watt[1]
- 20% Lower cost[1]



EC2 Instance Types

General Purpose



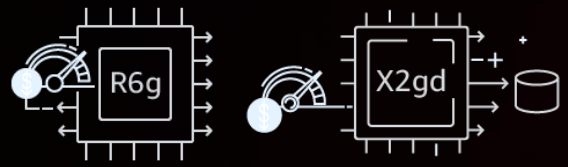
+ M6gd

Compute Optimized



+ C6gd, C6gn

Memory Optimized



+ R6gd

Storage Optimized

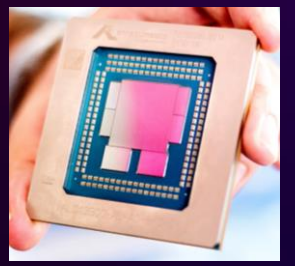


Accelerated Computing



AWS Graviton3

- 25% Performance[2]
- 3x ML Performance[2]
- 60% Less energy[3]

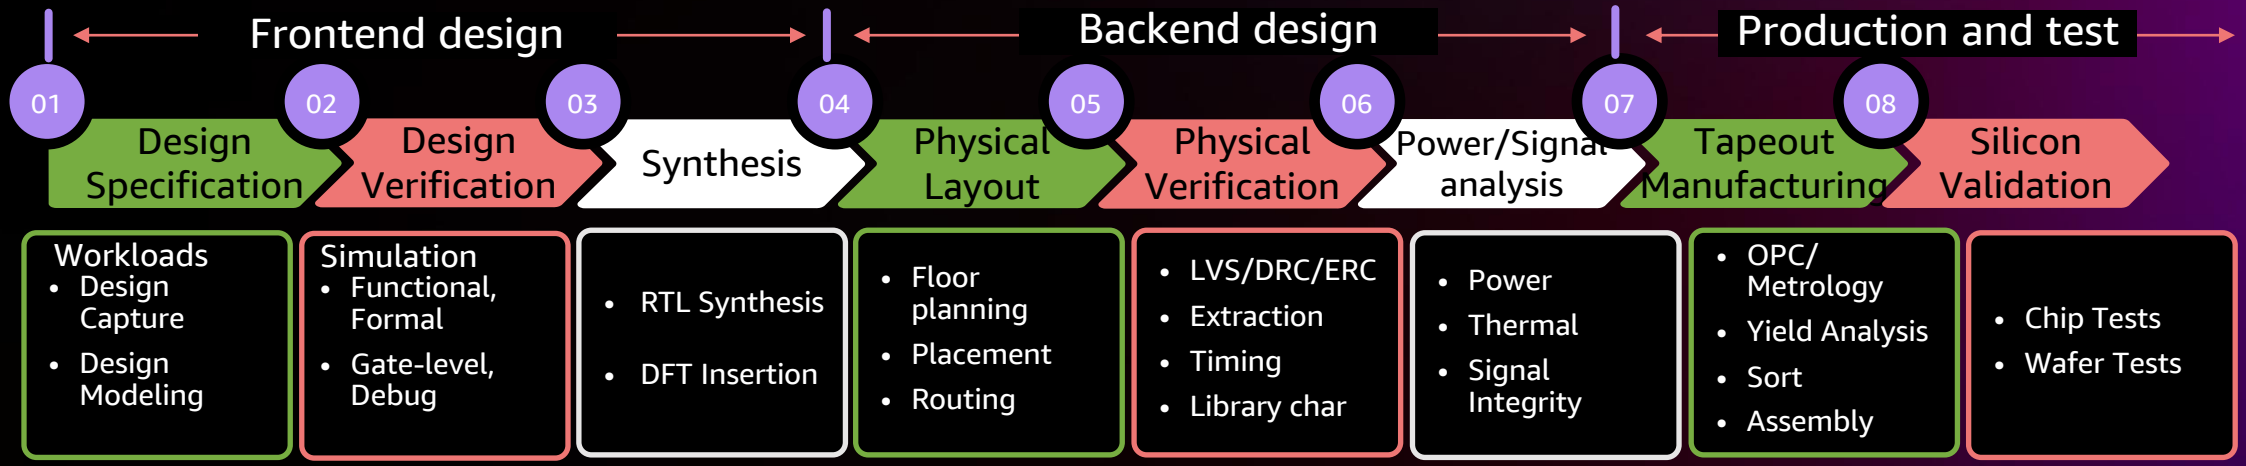


1. Compared to 5th Gen x86 instances
2. Compared to AWS Graviton2
3. Equal workload performance, compared to x86

Source: https://aws.amazon.com/ec2/graviton/

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

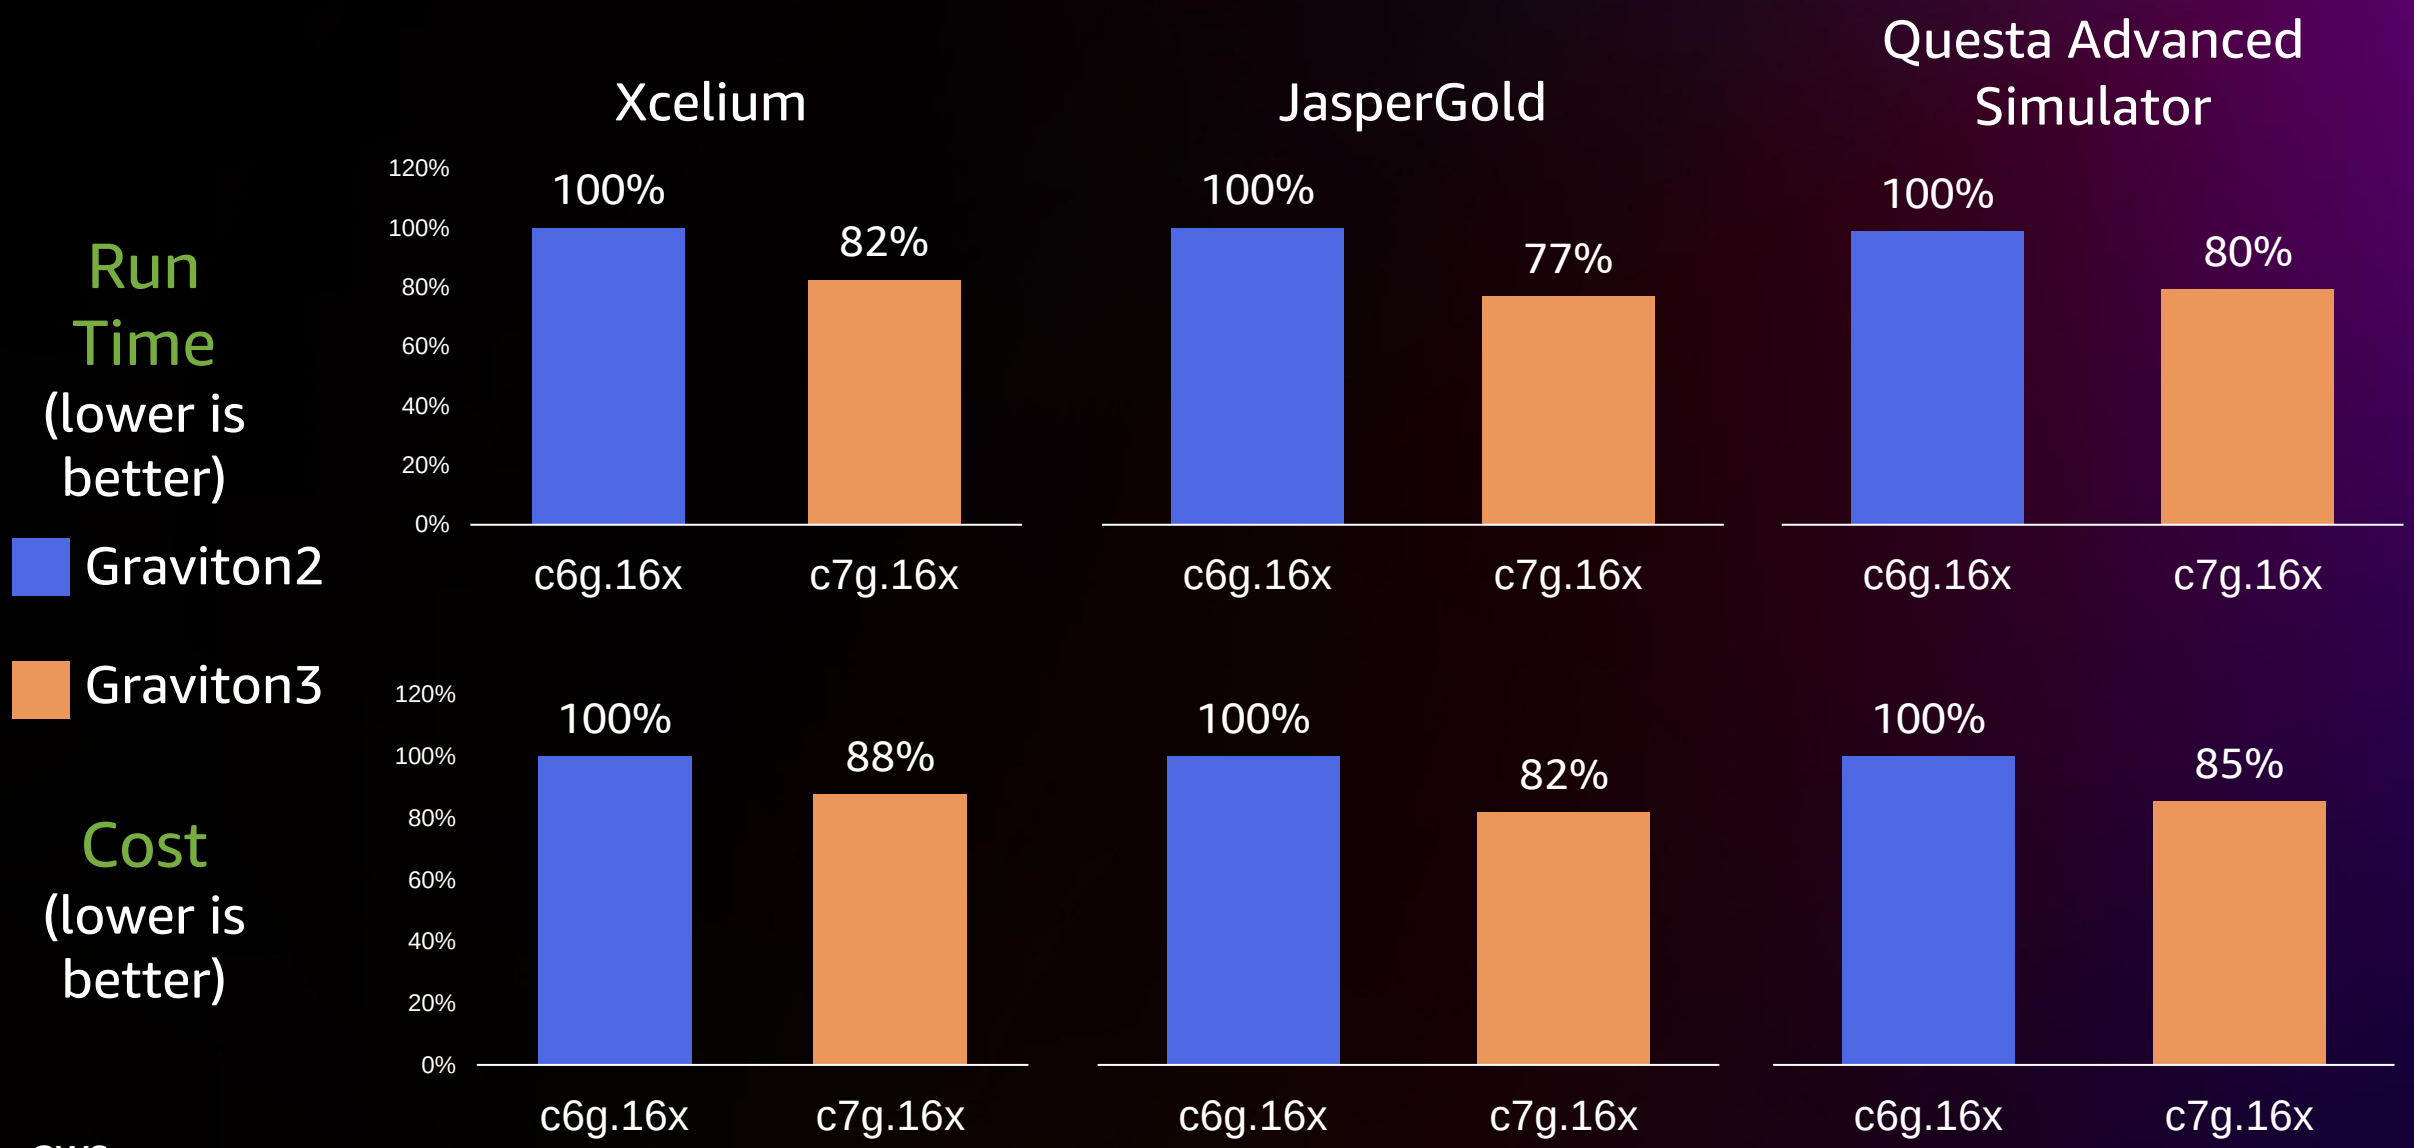
EDA Tools Availability for Arm



	Functional Simulation	Formal	Debug	Characterization/ Simulation
SYNOPSYS	VCS	VC Formal	Verdi	SiliconSmart → PrimeLib, CustomSim, Hspice
cadence	Xcelium	JasperGold	Indago	Liberate Trio, Spectre Simulation Platform, Spectre XPS
SIEMENS	Questa Prime	Questa PropCheck		Solido
Ansys				APL->RedHawk-SC

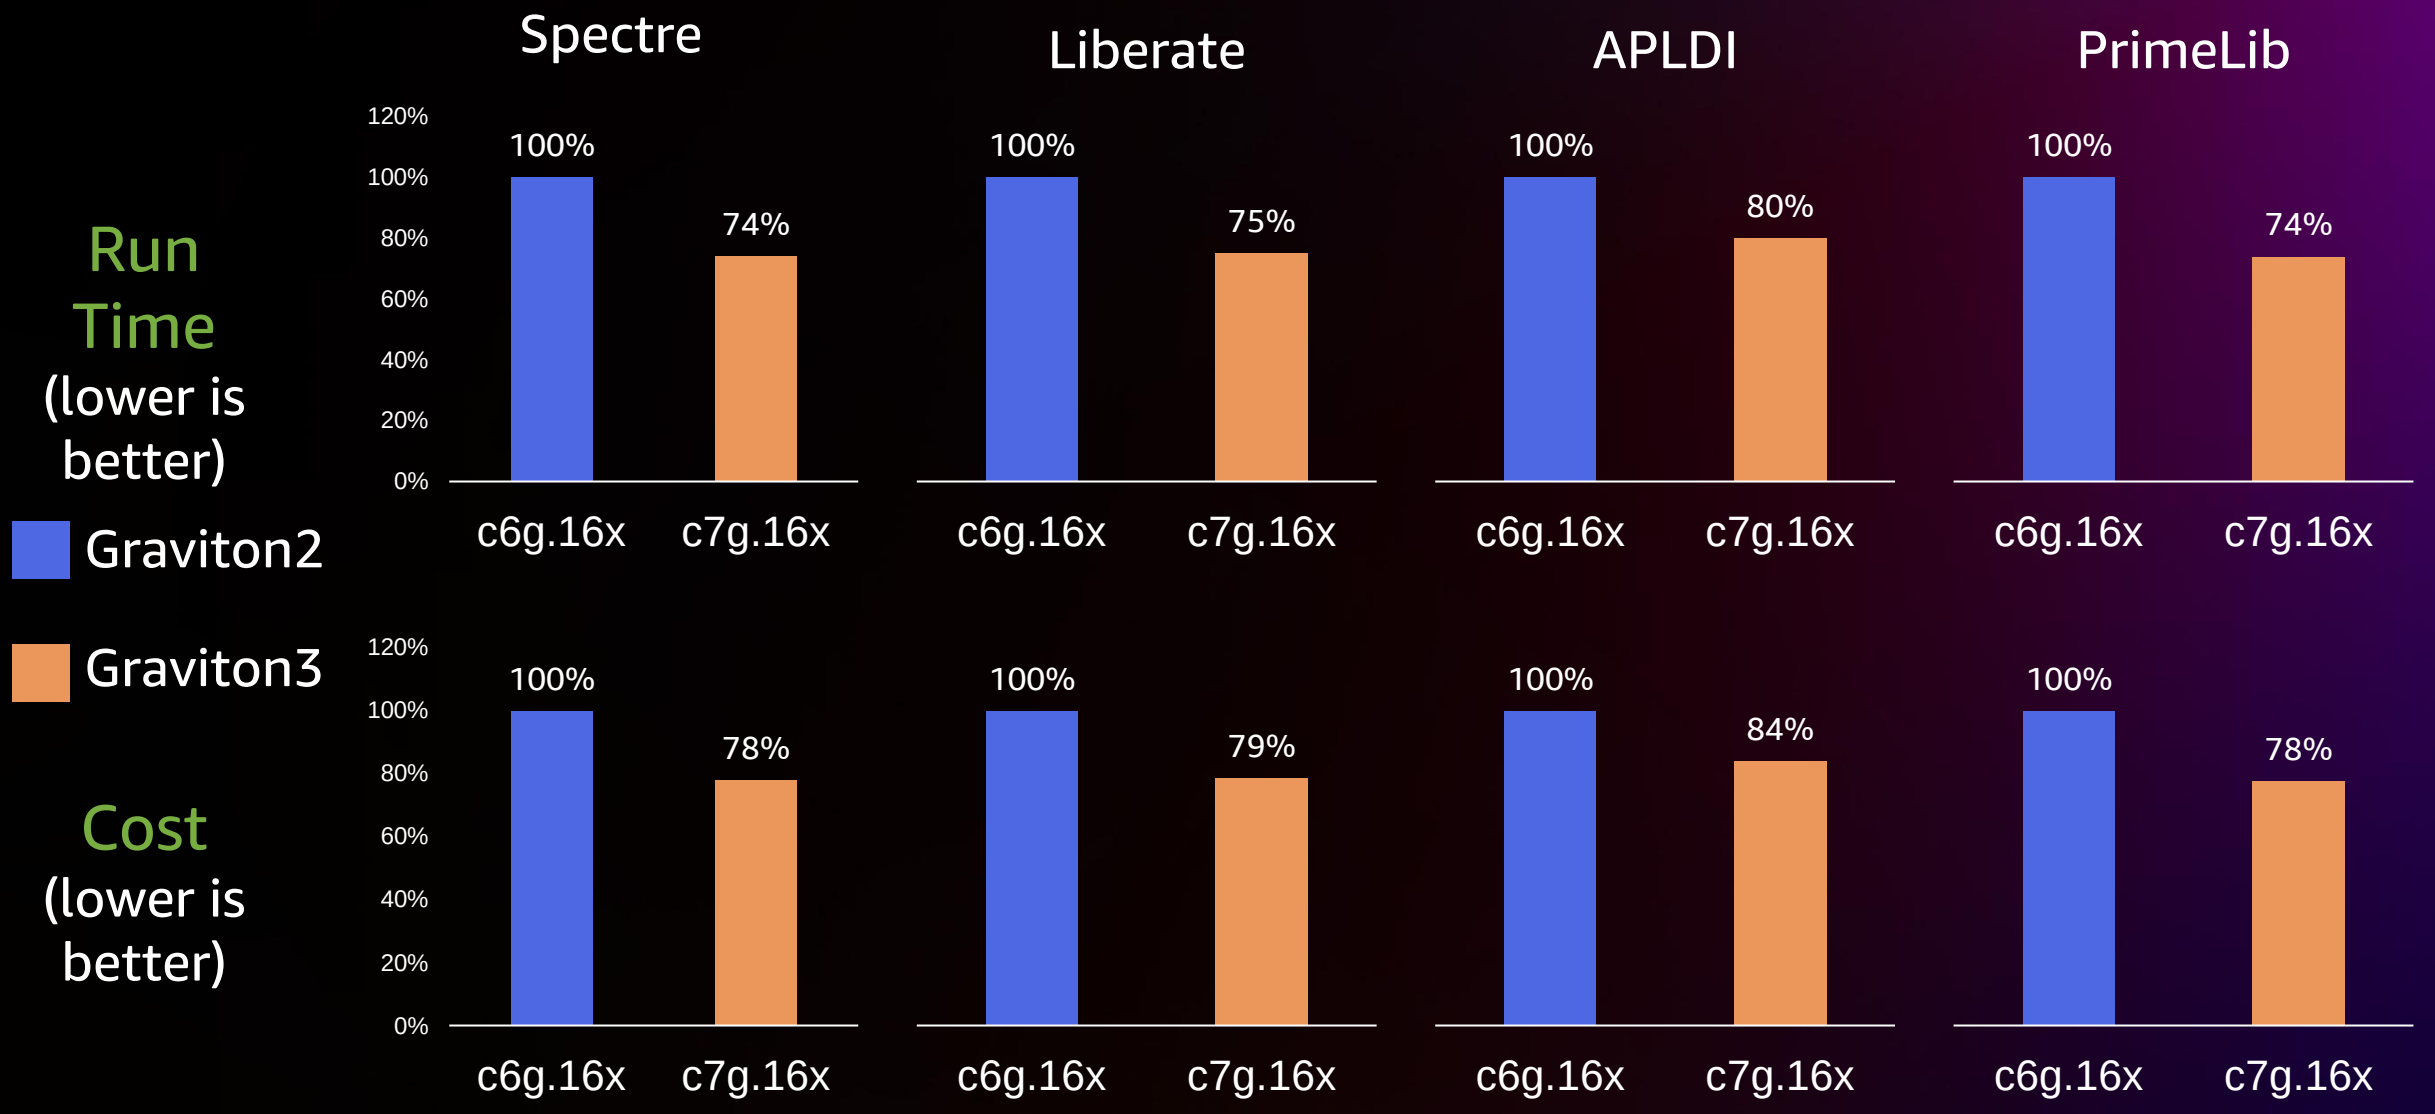
aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

AWS Graviton3 price/performance for FE EDA



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

AWS Graviton3 price/performance for BE EDA



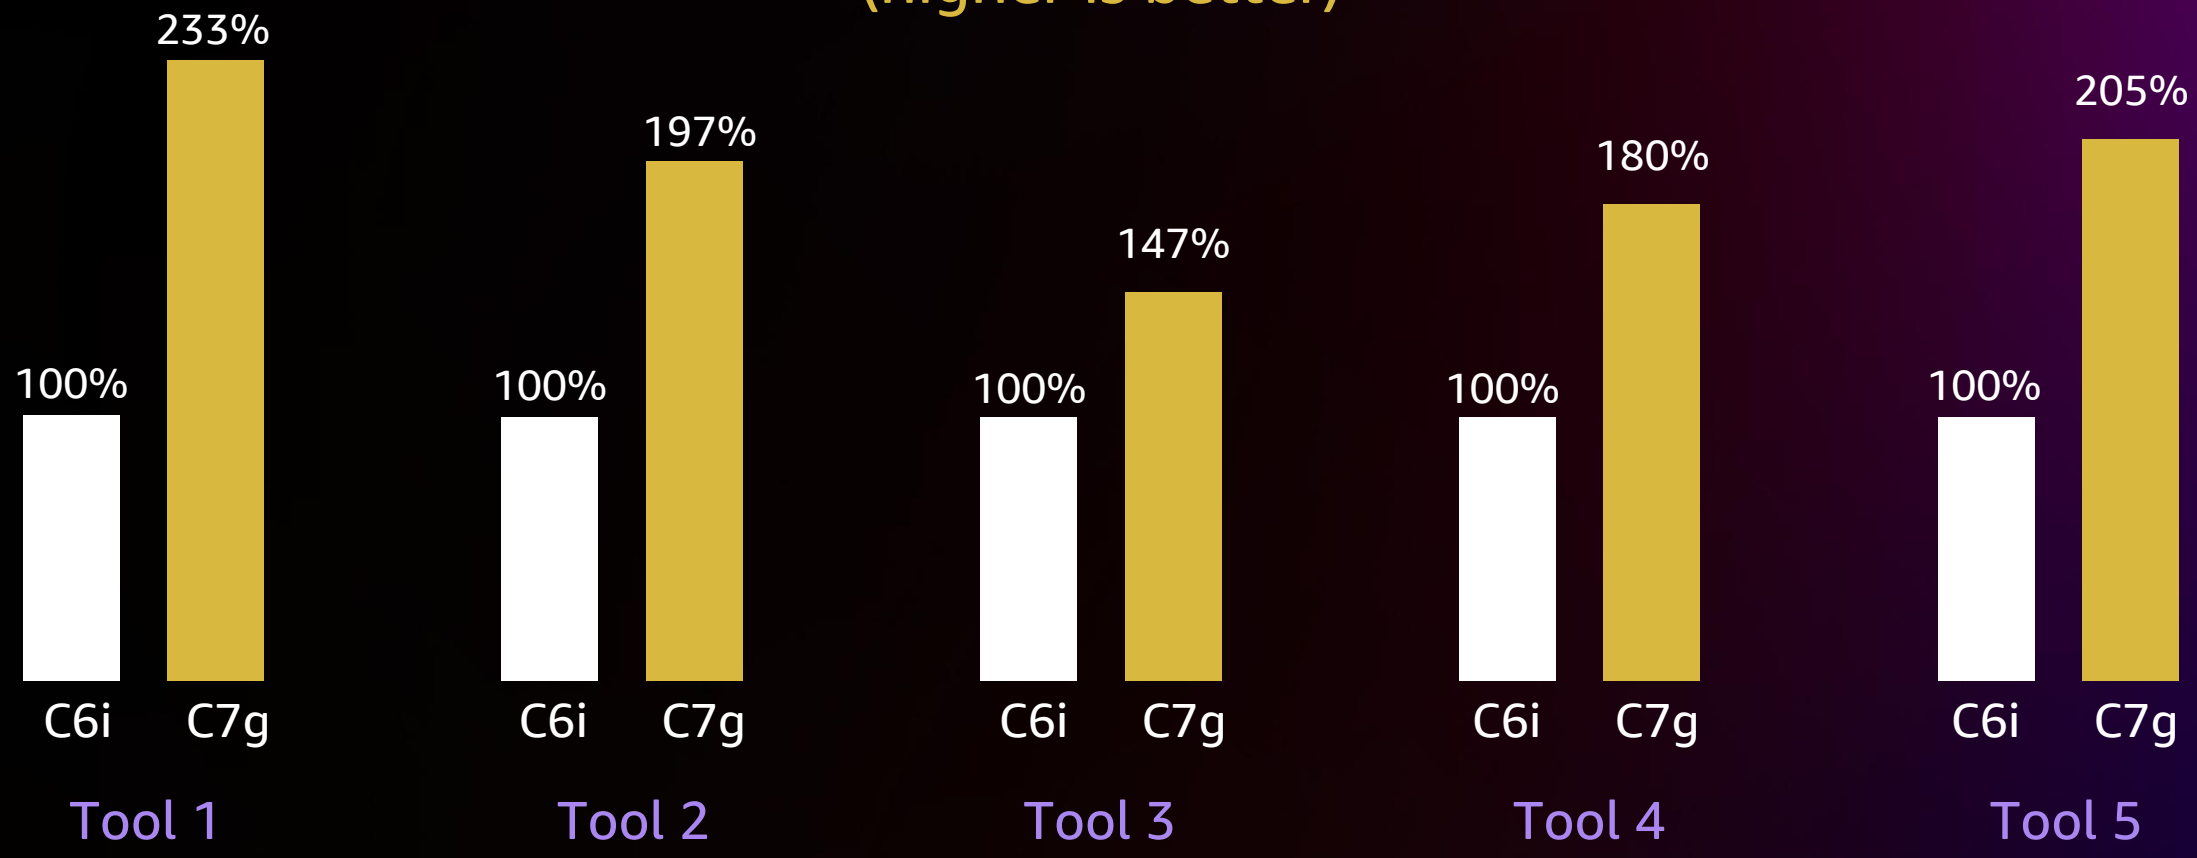
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

EDA Tools Performance

AWS Graviton3 (C7g) vs. latest x86 (C6i)

Ansys
cadence
SIEMENS
SYNOPSYS

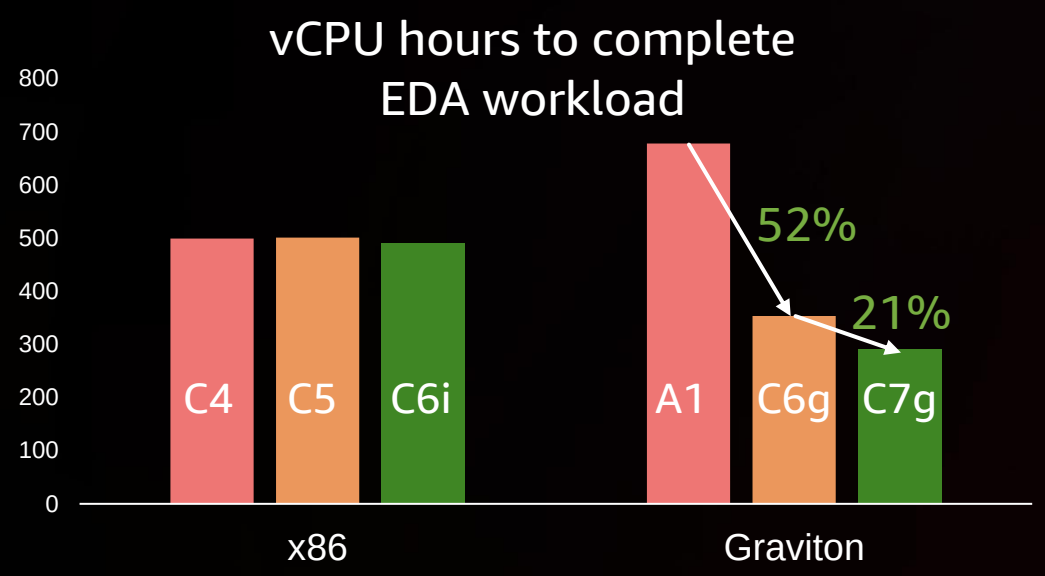
# of simulation jobs completed for the same cost
(higher is better)



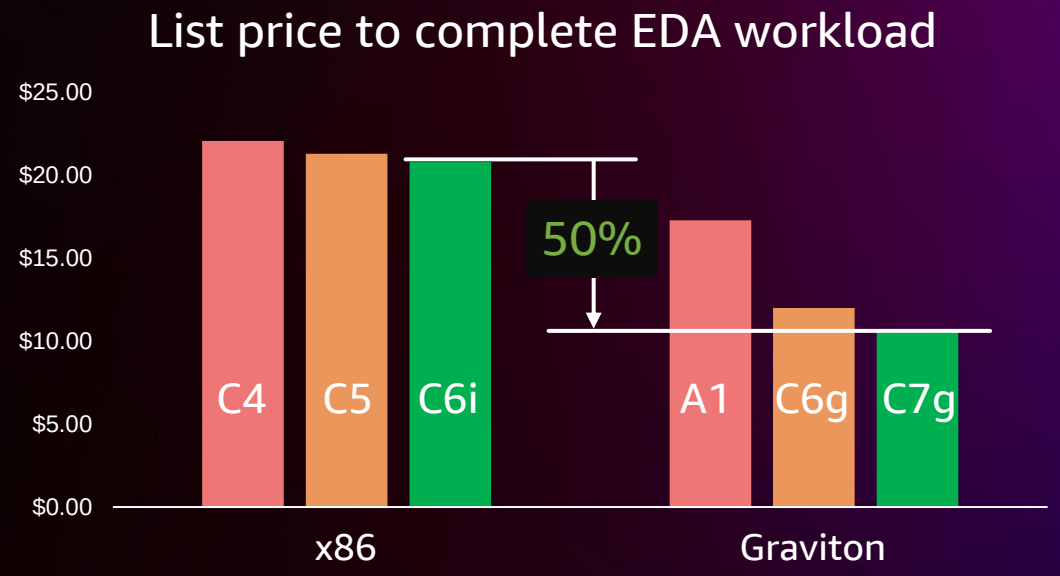
aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Price-performance gains with AWS Graviton



Graviton achieving significant performance improvements with each new generation



Graviton price/performance improvements continue to surpass traditional architectures

aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Physical IP design on cloud

Batch workloads on Arm Cloud

- Running production workloads on AWS since 2019
- 80% of PDG's usage on AWS is Graviton2 today
- Taking advantage of the performance and cost using Arm-based Graviton 2 family

Ansys
- APLDI

cadence
- Liberate Trio
- Spectre Simulation Platform
- Spectre XPS

SIEMENS
- Solido

SYNOPSYS
- PrimeLib
- SiliconSmart
- FineSim

Implementation workloads characteristics

- Project execution completely on cloud
- High run time and memory requirement
- Interactive session

Cloud foundation platform

- LSF on cloud
- On-prem-like latency for interactive sessions
- Compute instance available on cloud which provide 3-24 TB of memory per host.

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

On-premises environment



Job submission

Users

LDAP

Compliance, security

PDK EDA tools

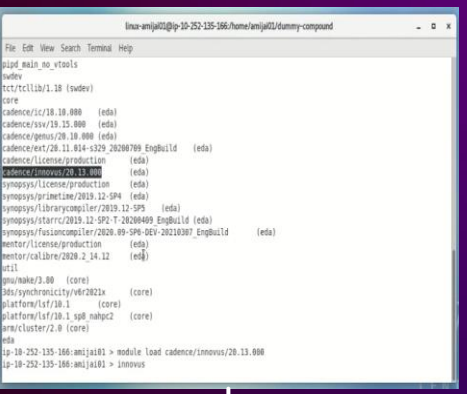
NetApp

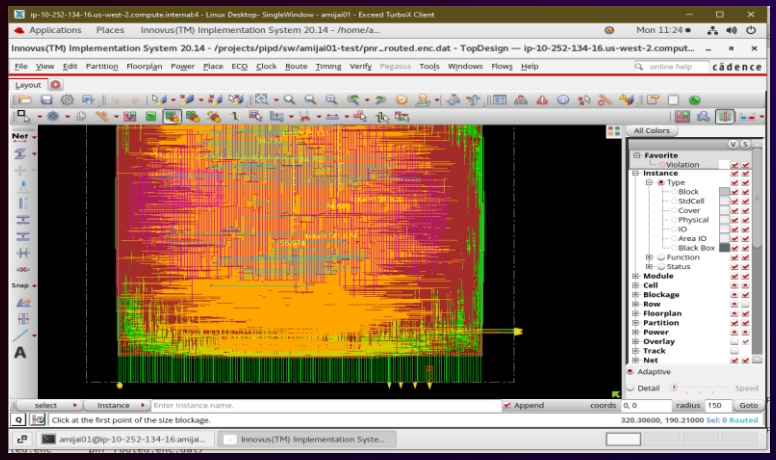
Master bsub

IBM LSF Cluster

EC2 Clients

License server

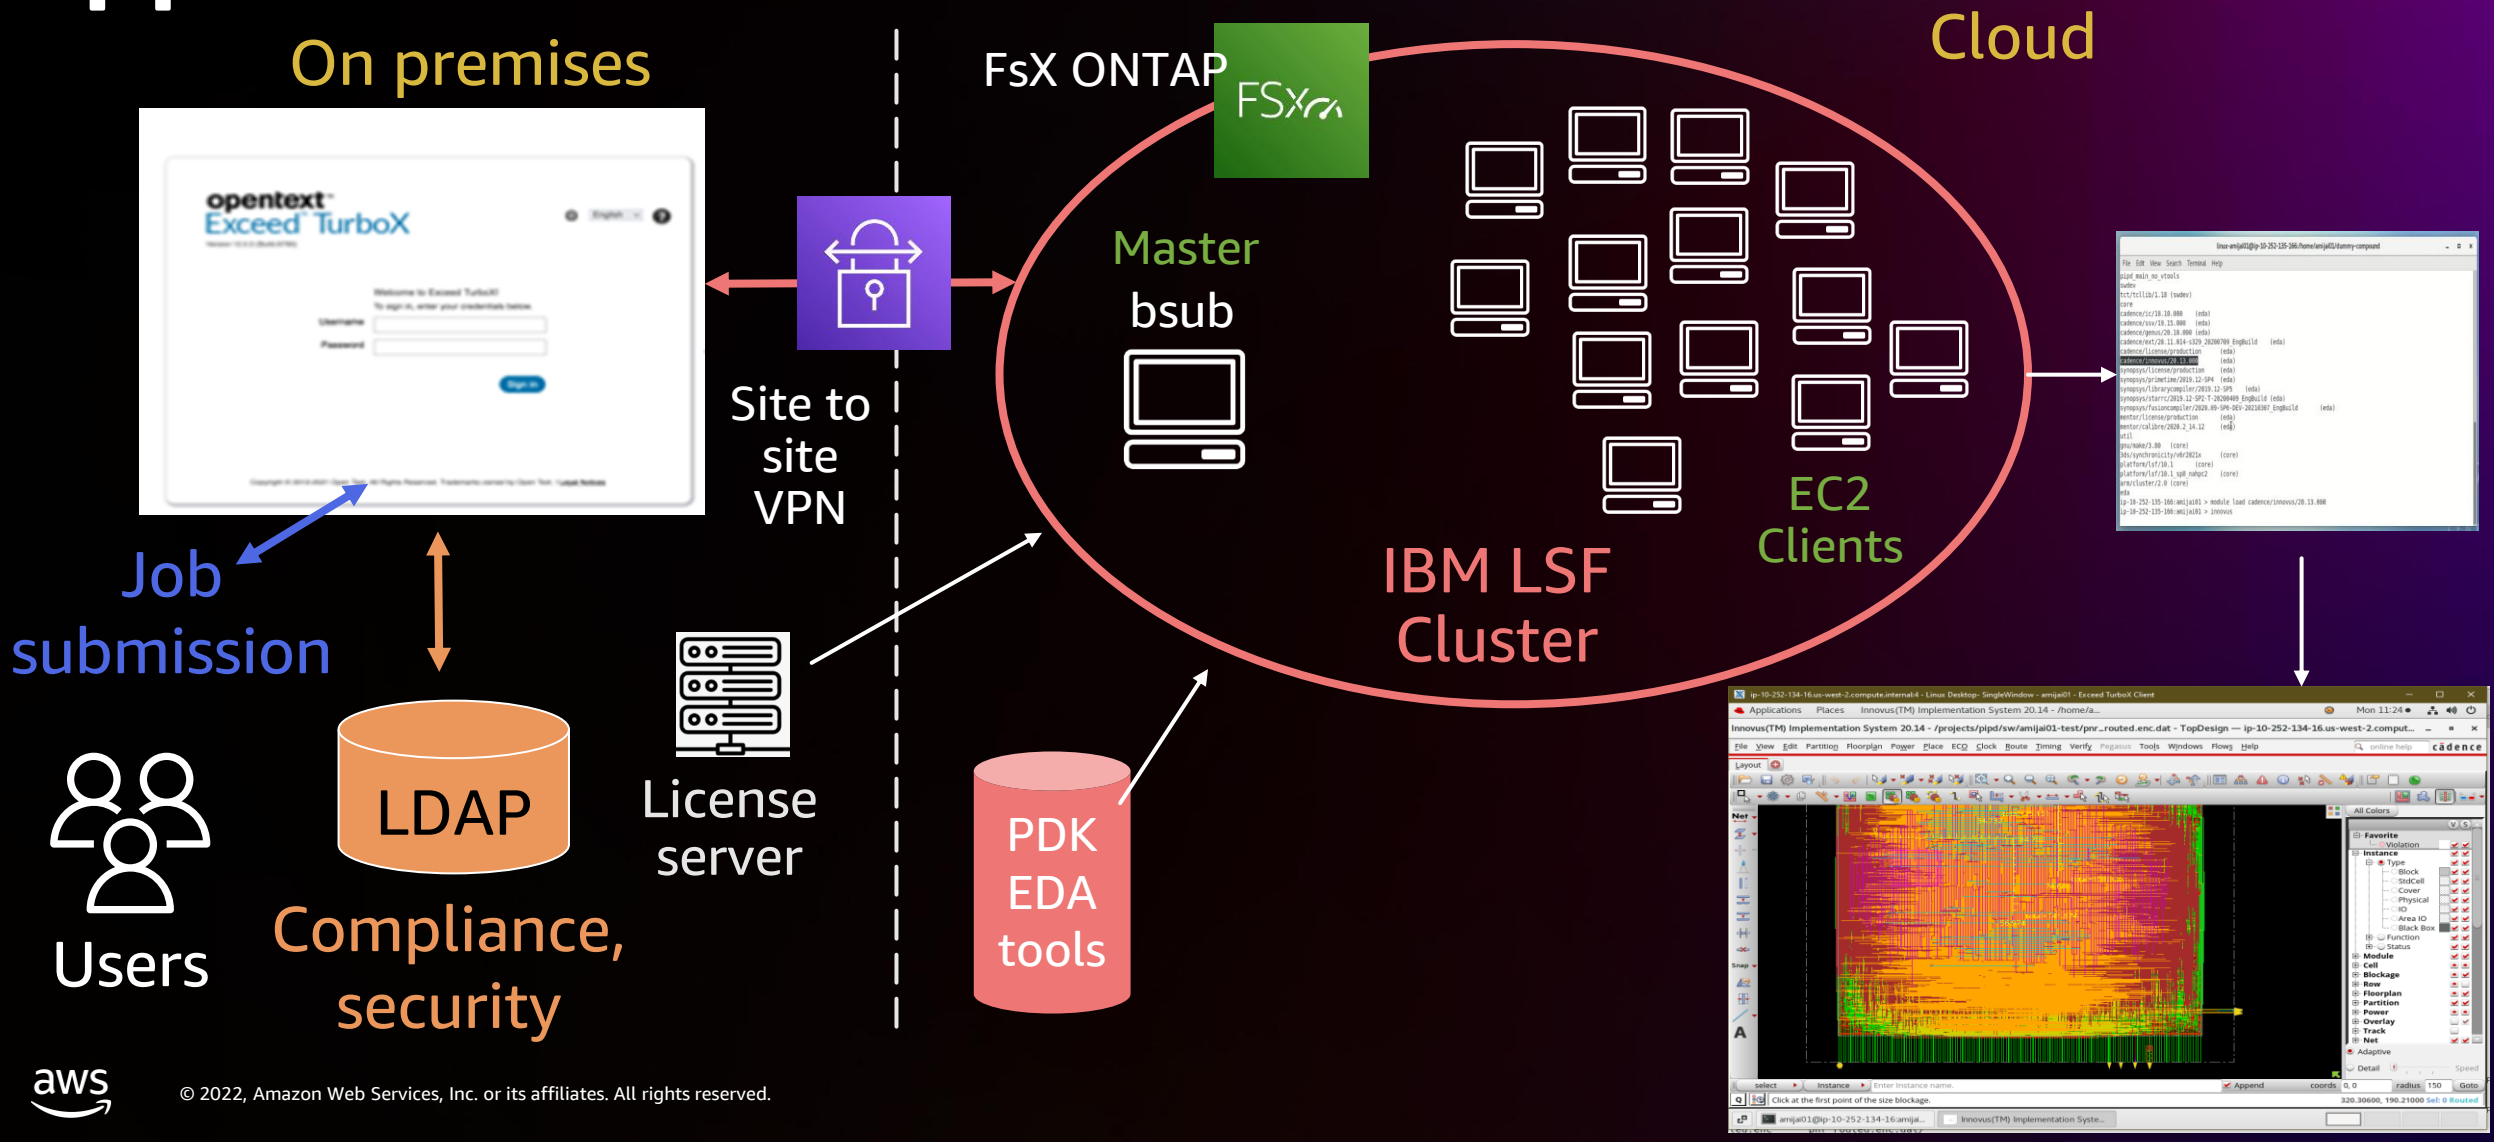




aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Cloud Foundation: Running on-premises applications in the cloud



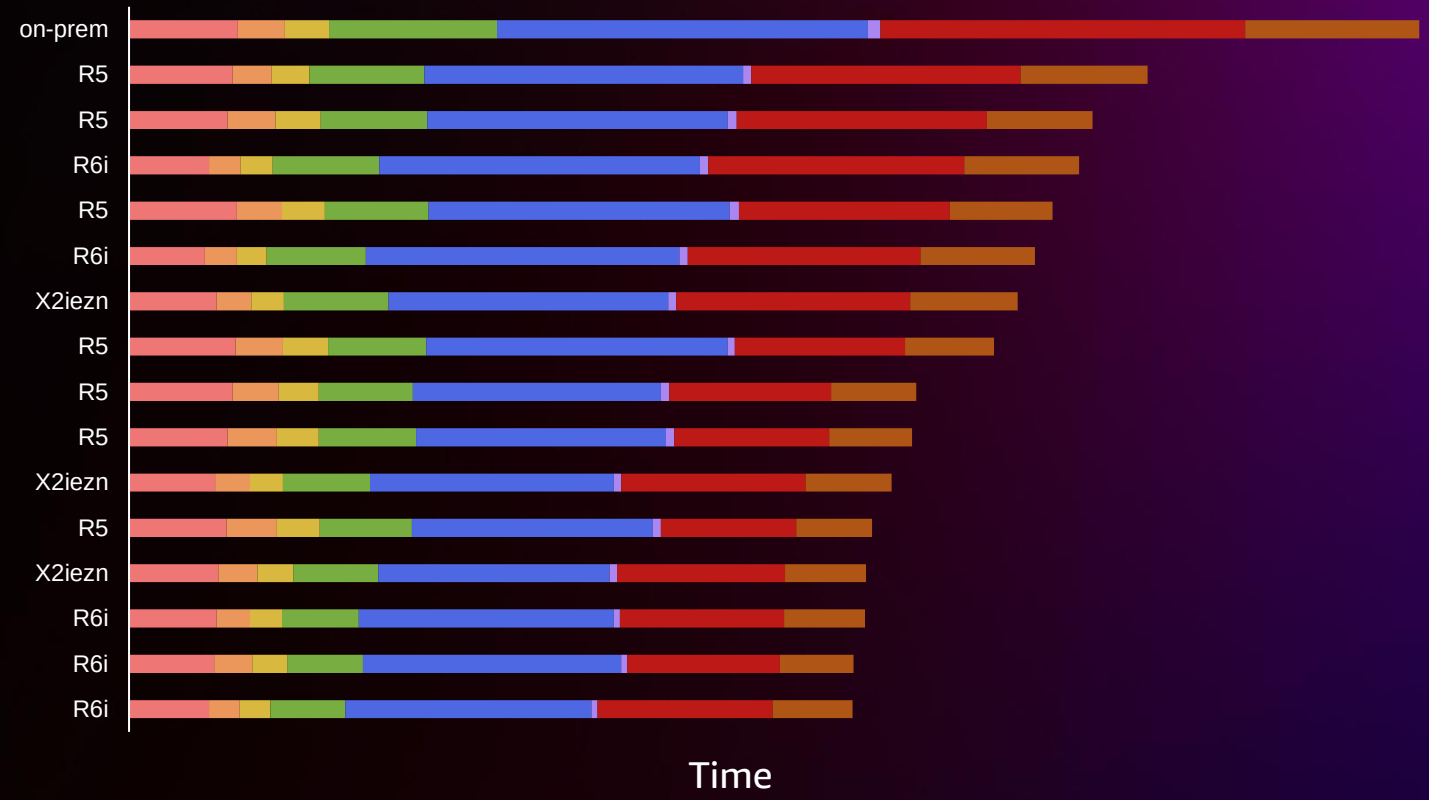
aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Choosing the best instance types

Cost of workload
Lower is better



Runtime for different phases of an implementation workload
Lower is better



Workload testing lets Arm identify the fastest instance and the instance offering the best performance-per-dollar

aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

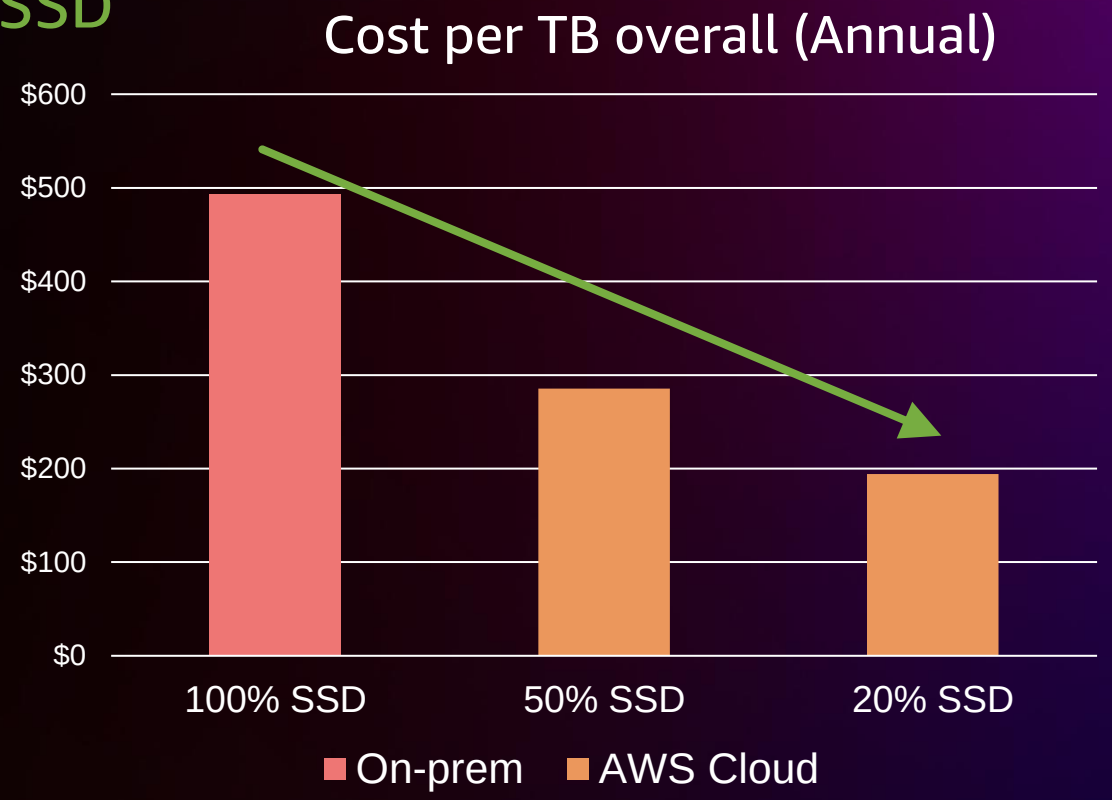
Optimizing Storage cost on cloud

Using right storage technology

- Various storage technologies available: Amazon EBS, Amazon EFS, Amazon FsX for Lustre, Amazon FsX ONTAP, Amazon S3
- FsX ONTAP found to be best suited for metadata heavy EDA workloads

Purpose	Technology
Local storage	EBS
Shared storage	FsX ONTAP + capacity pool tier
Back up	FsX ONTAP + capacity pool tier
Archival	S3

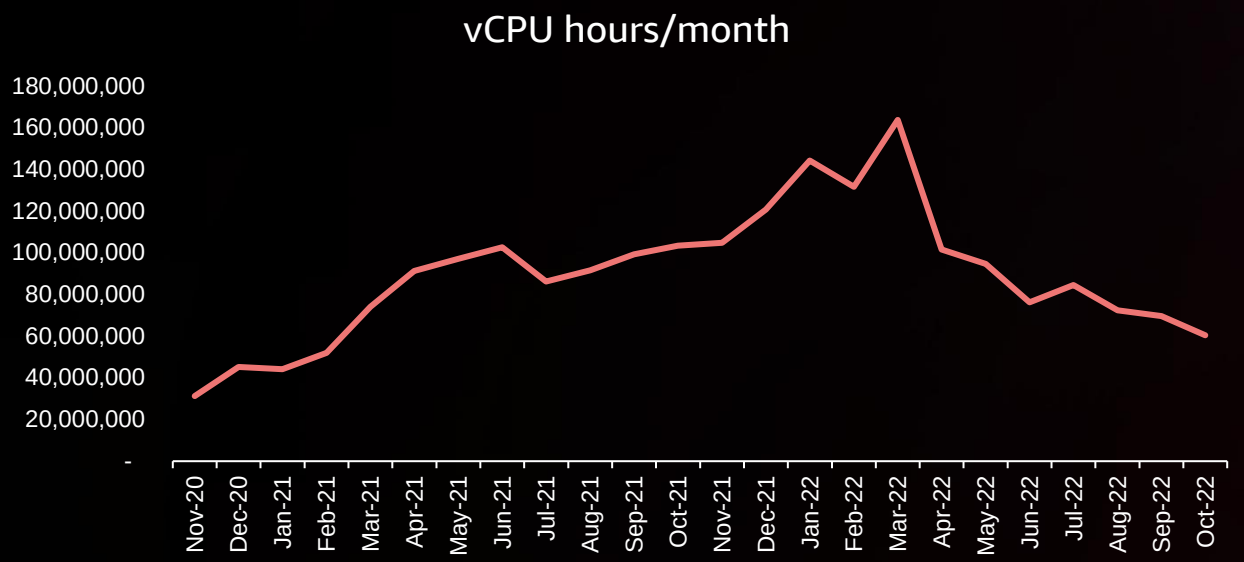
ONTAP cost reduction: 2.5X cost reduction using 20% SSD versus 100% SSD

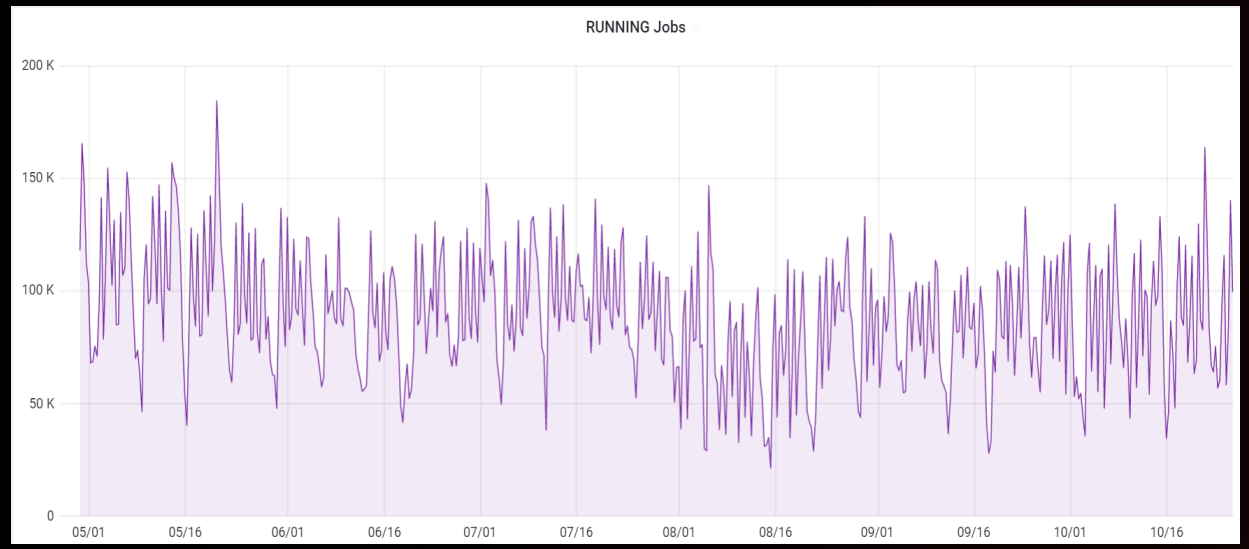


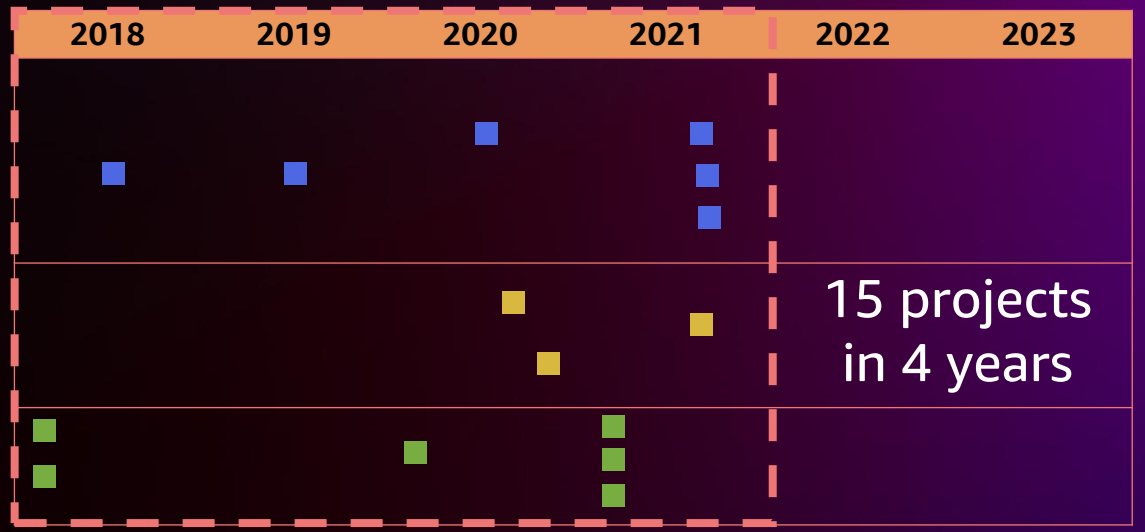
- Ireland region single AZ deployment
- Estimates based on 100TB storage 128Mbps throughput

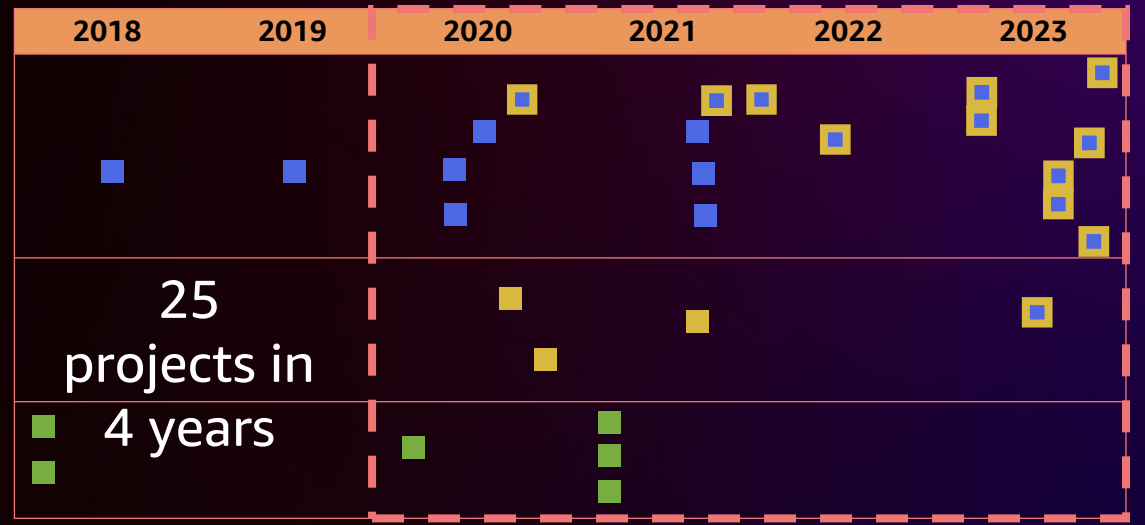
aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Impact of RTL verification in AWS





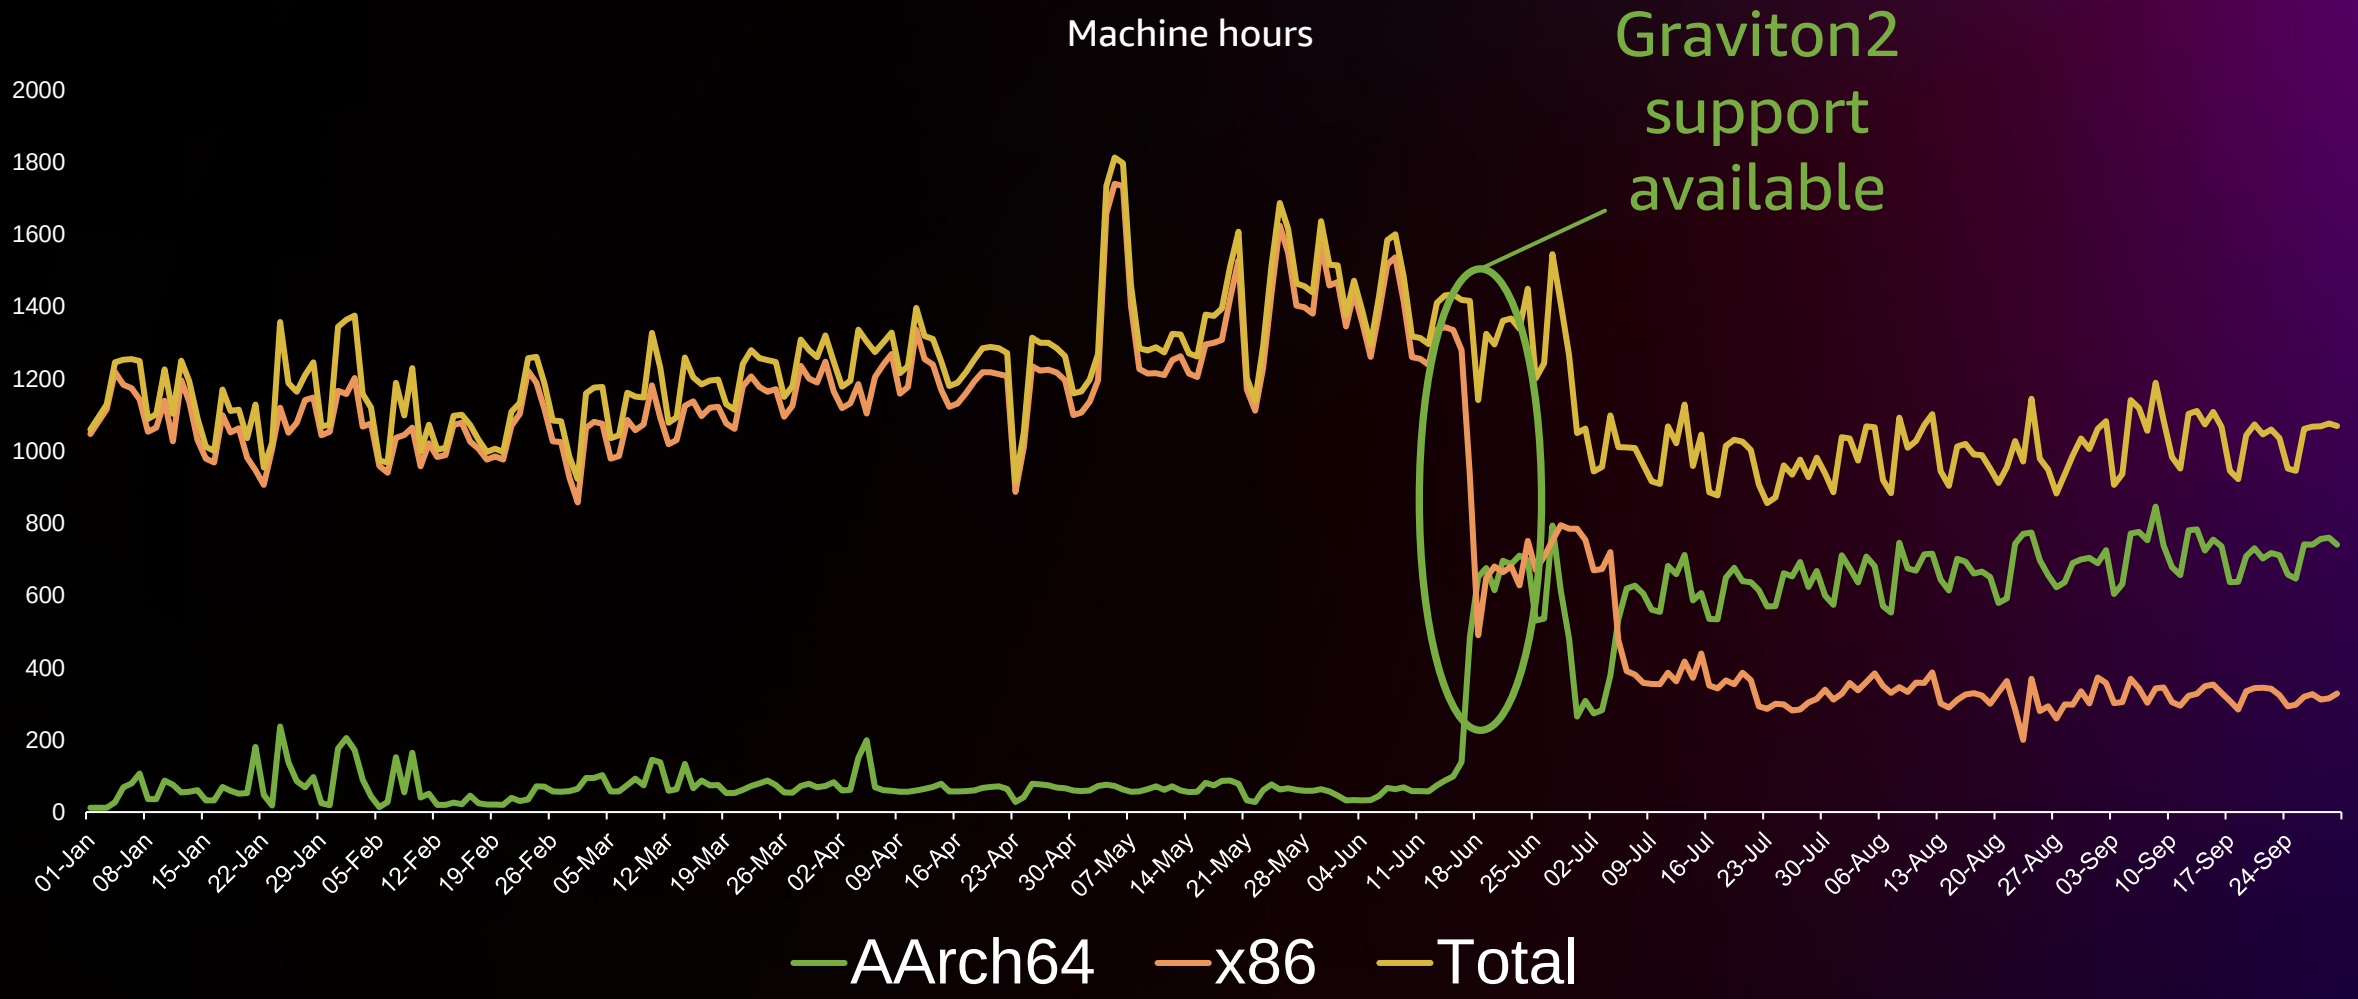




+66% added throughput with cloud

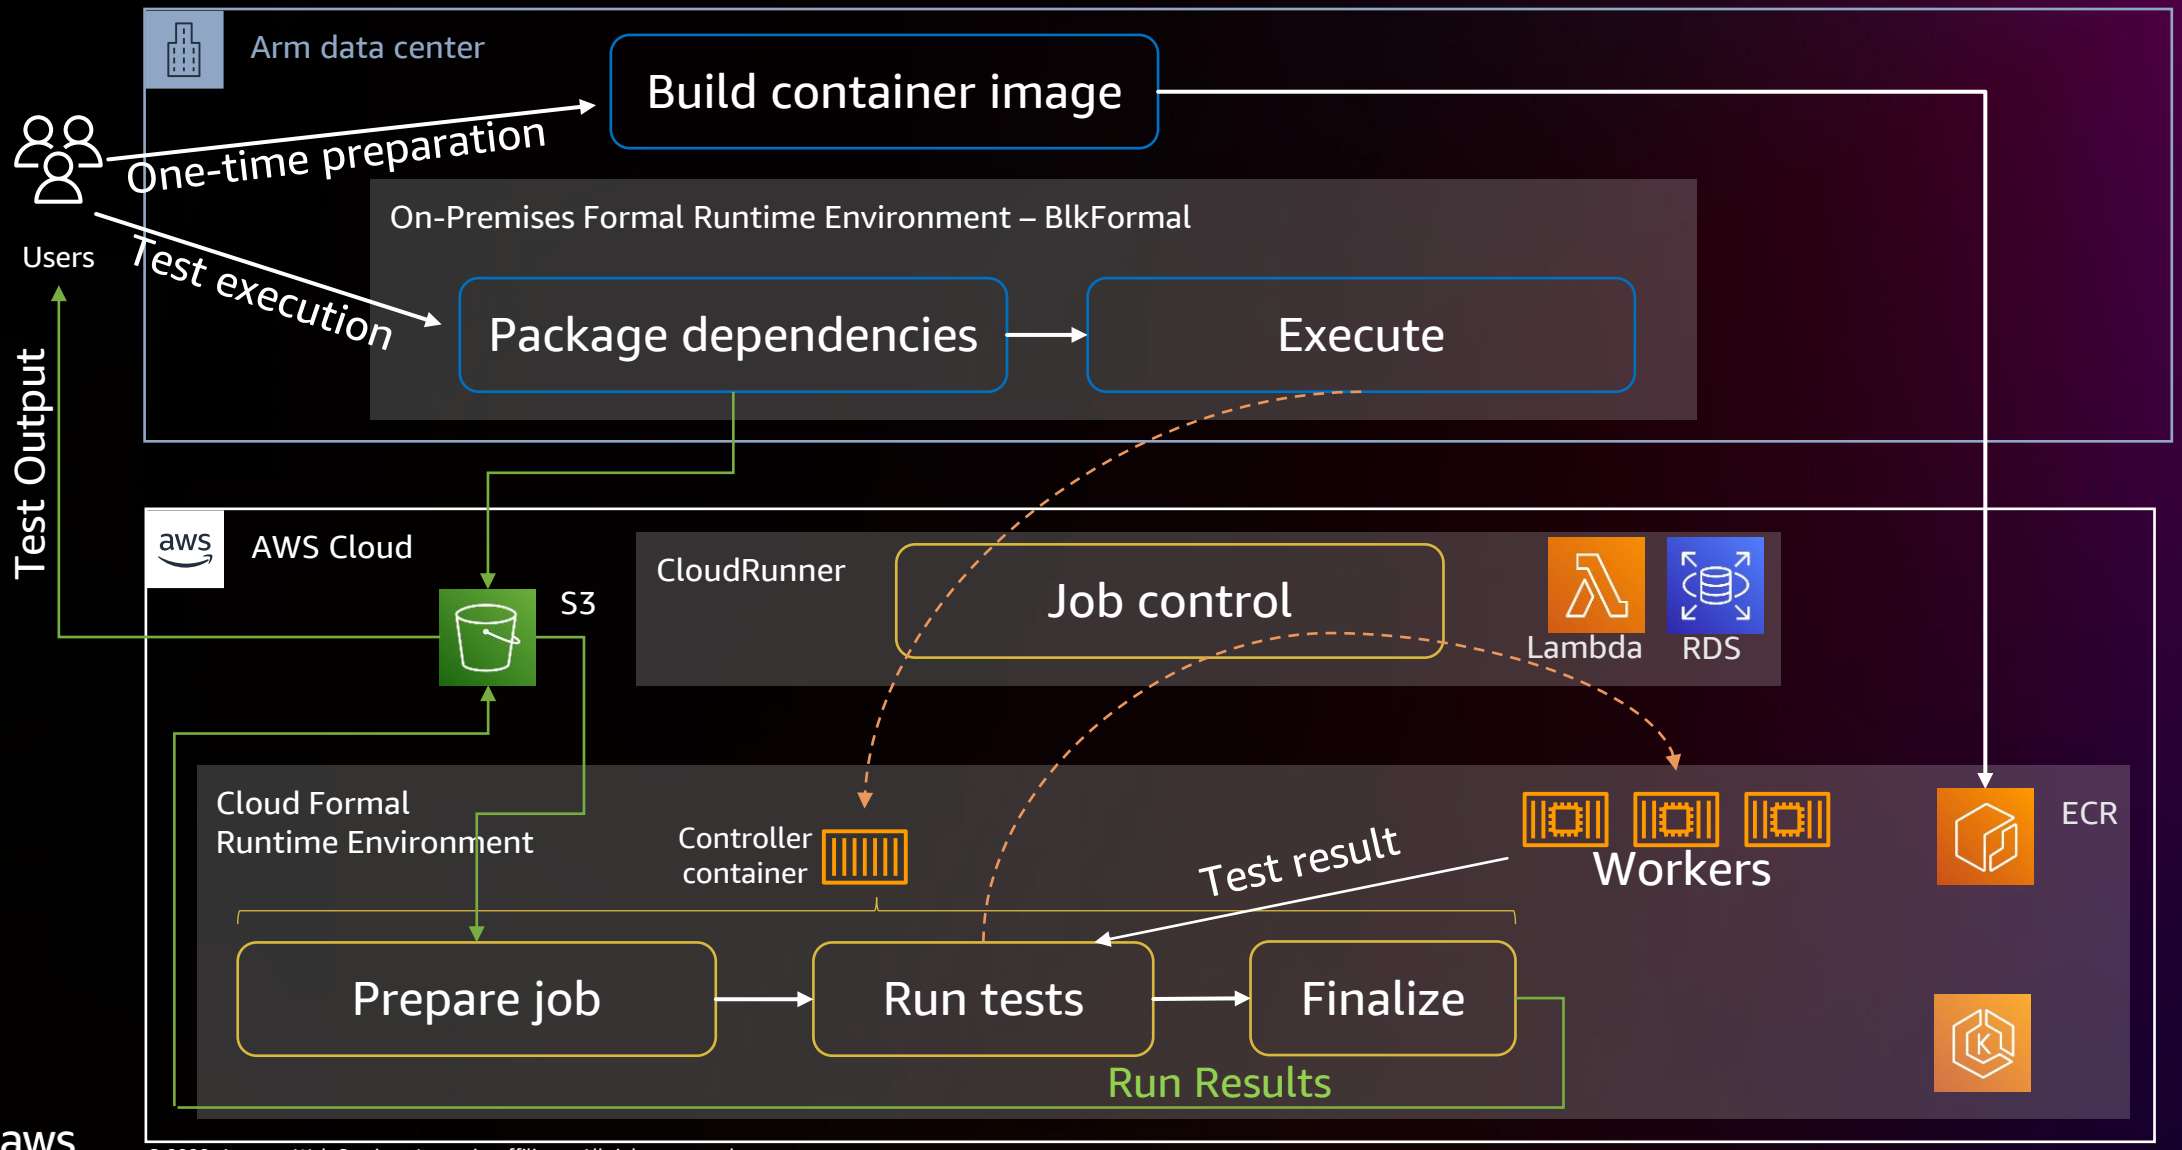
aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Case Study: Databricks



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Formal Flow IP Design in the Arm cloud



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Software developers experiences

On premises	AWS	Benefit
Could take 30 days to get a new VM provisioned	Cloud VM provisioned in seconds	Flexibility of cloud allows software teams to have the infrastructure they need, when they need it
Single common development environment	Multiple environments easy to deploy	Teams can establish test environments for multiple configurations of their software
Once provisioned, infrastructure is static	Instance configuration can be changed as needed	Occasional need for large memory can be satisfied when required, but then return to cheaper instances
Hardware procurement takes months to deliver	New technologies available on tap	Brand new processor features, like SVE-enabled Graviton3, can be integrated as soon as available

“So, all in all, I'd say we gained 5hrs/week/head with the switch to AWS. It felt like more than that, actually.”

Andrzej Warzniski
Flang compiler engineering

aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Arm’s cloud ecosystem is rapidly expanding


AWS Lambda


Amazon ElastiCache


Amazon Aurora


Amazon EMR


Amazon EKS


Amazon EC2


Amazon RDS


Amazon Neptune


Amazon MemoryDB for Redis


Amazon OpenSearch Service


AWS Fargate


Amazon ECR


Other cloud providers

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Summary

- Designing chips is resource intensive and expensive – but the AWS cloud helps us improve productivity and reduce cost
  - We leverage AWS Graviton wherever we can (better performance, lower cost than x86)
  - We are looking to move more workloads to Graviton3 (+10-20% more perf/$ over Graviton2)
  - With cloud, we can always find the optimal instance at the best price for the job
  - With cloud, we always have compute scalability which improves throughput
  - Our engineers are happier and more productive
- The EDA partners and tools ecosystem is ready for Arm
- Arm is now in every major cloud



aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

CMP320

Marvell EDA on Cloud

Bill Cornwell

Vice President, CE – ASIC & Processor
Marvell

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.



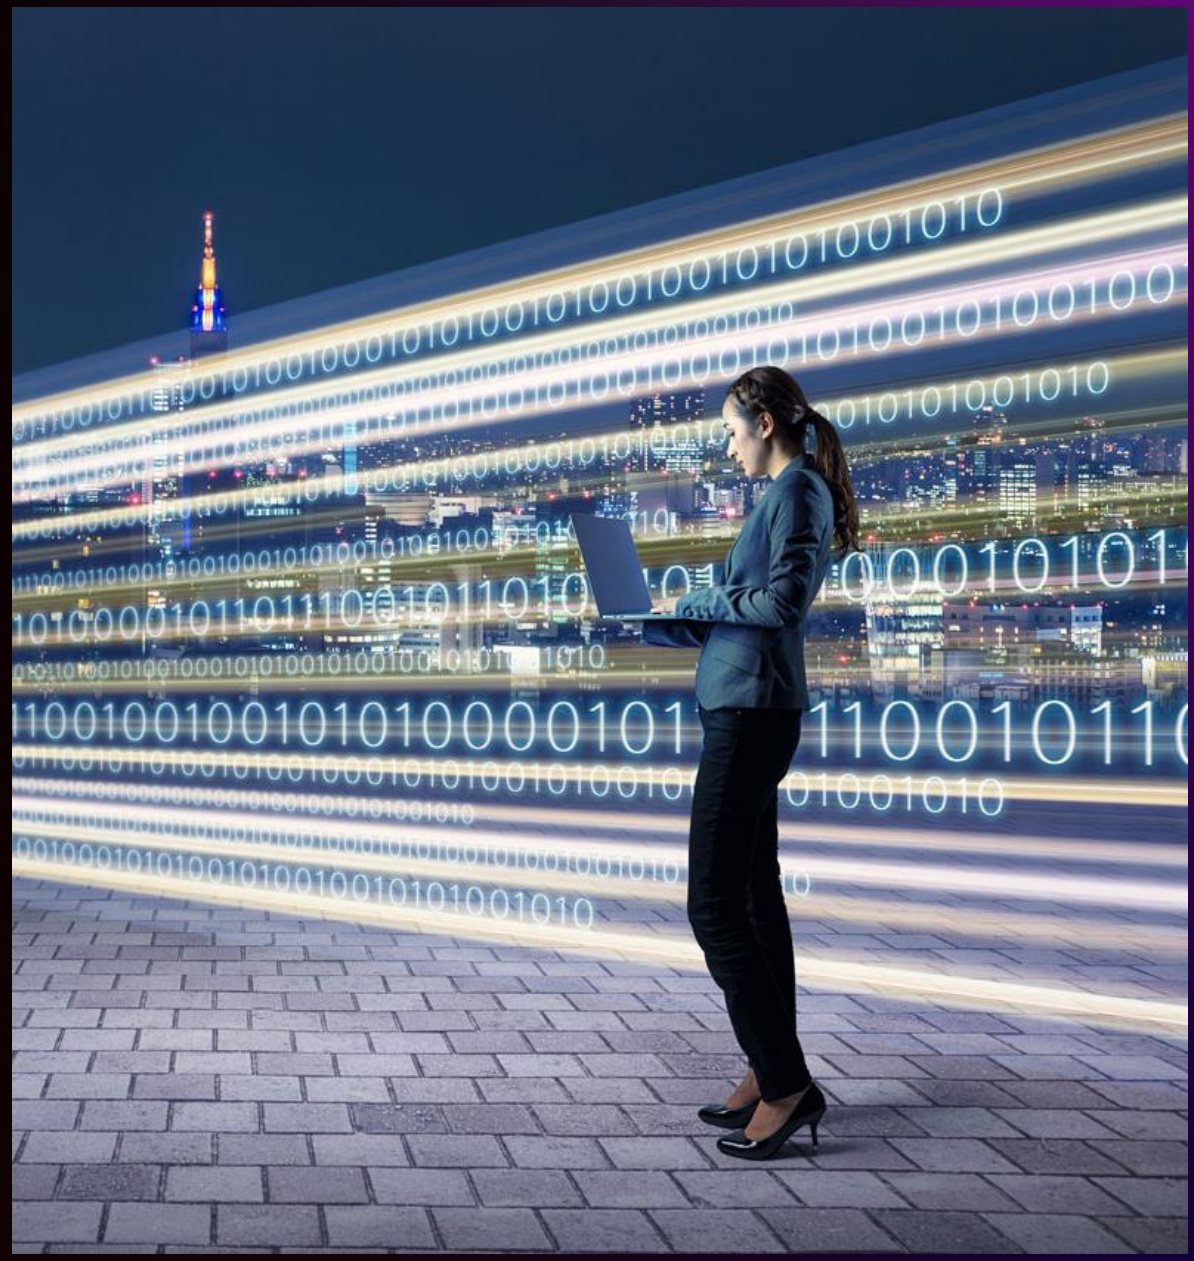
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Agenda

Design chaos

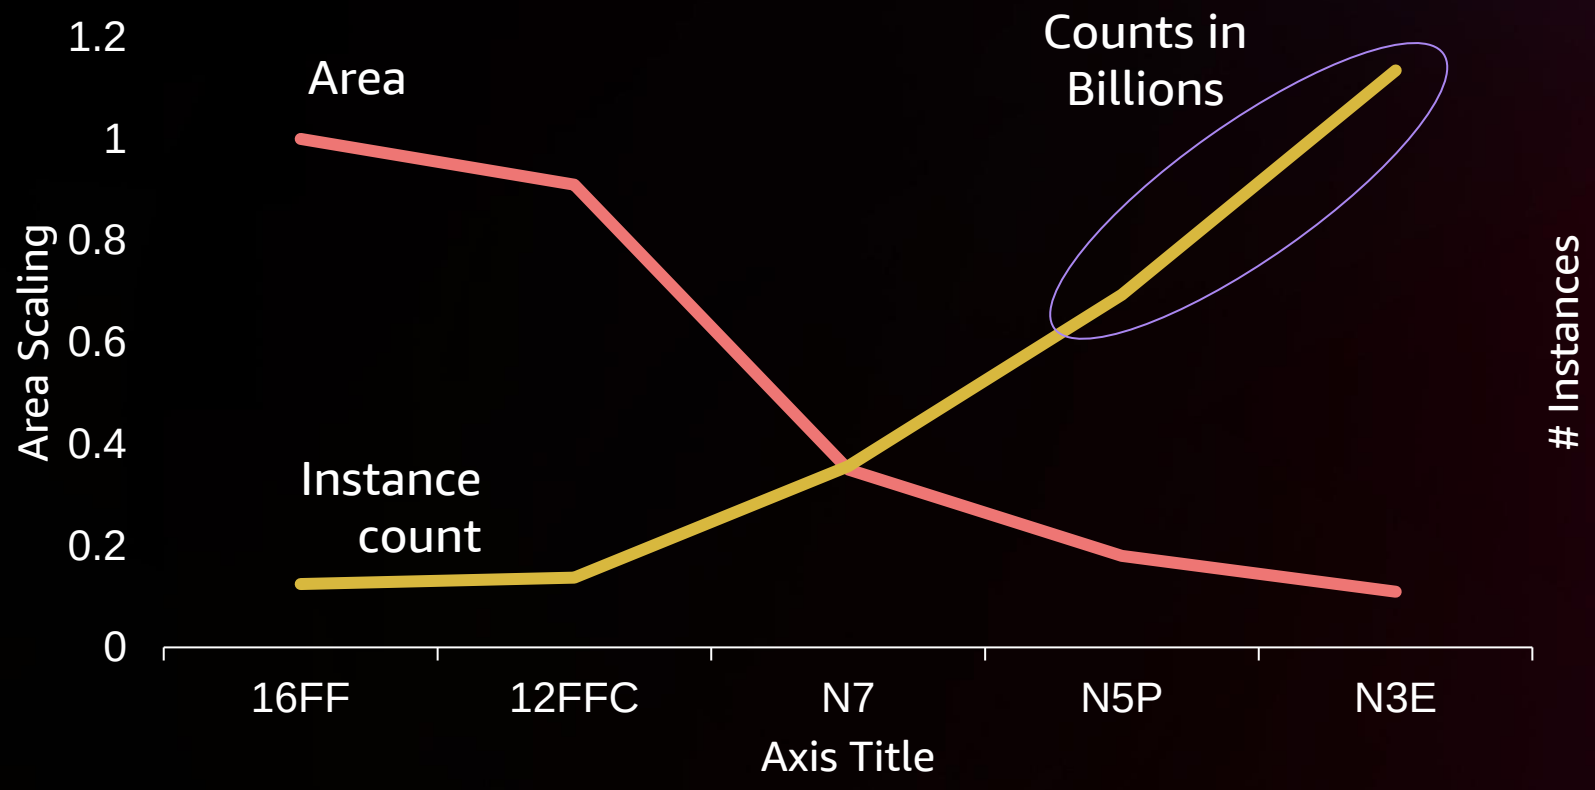
Scaling smart

Innovate faster



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Design Chaos

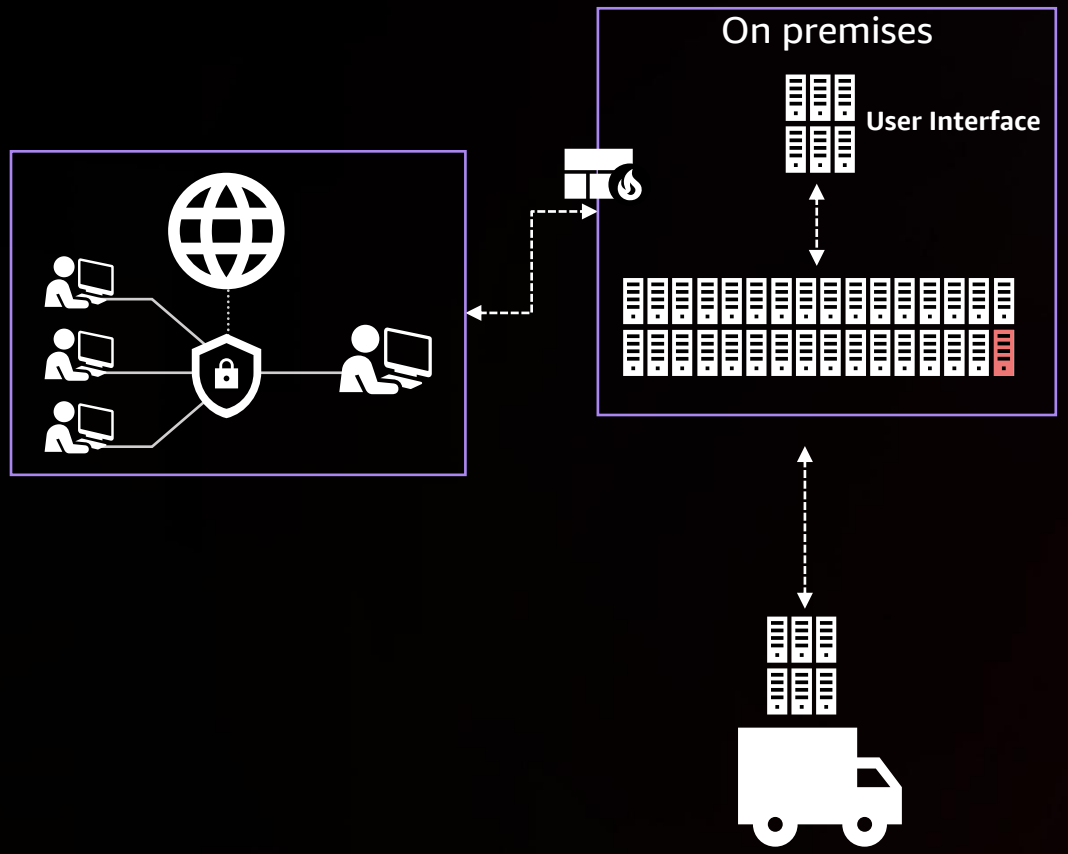


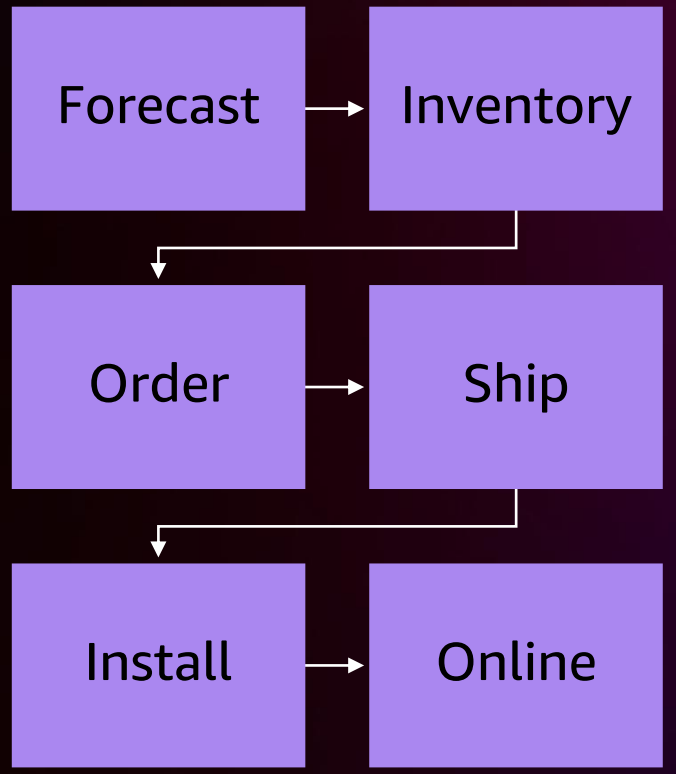
- Technology Scaling
- Increasing design complexity

- Increasing demand
  Machine Resources
  Engineering

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

On-premises experience

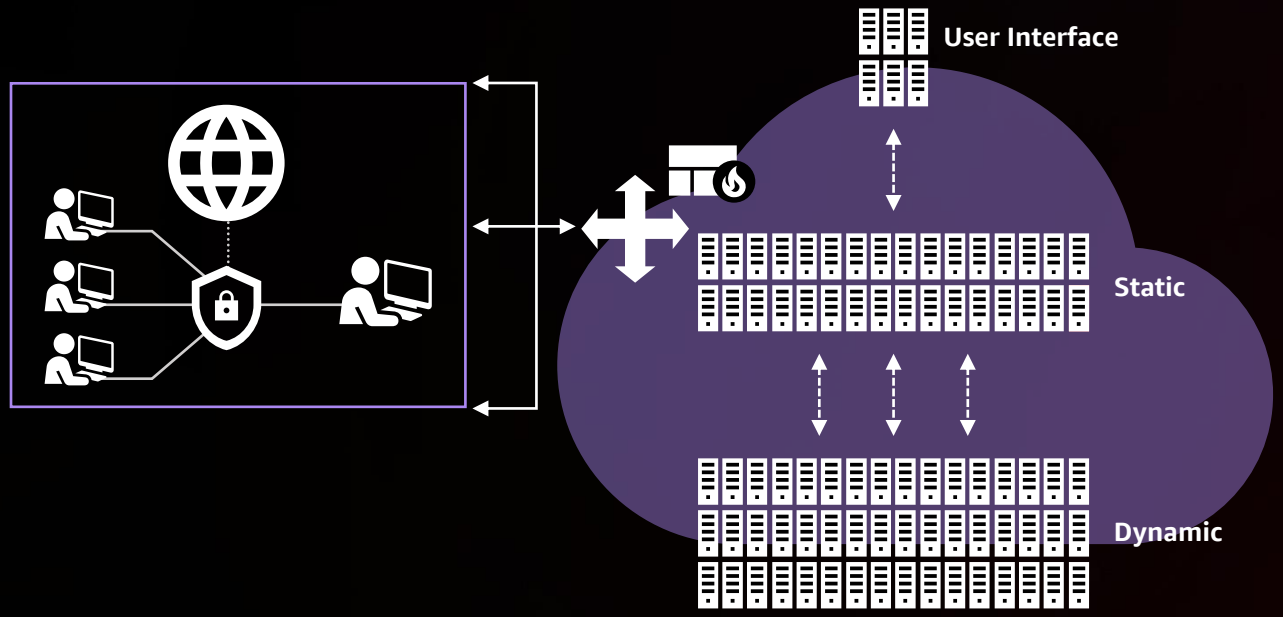


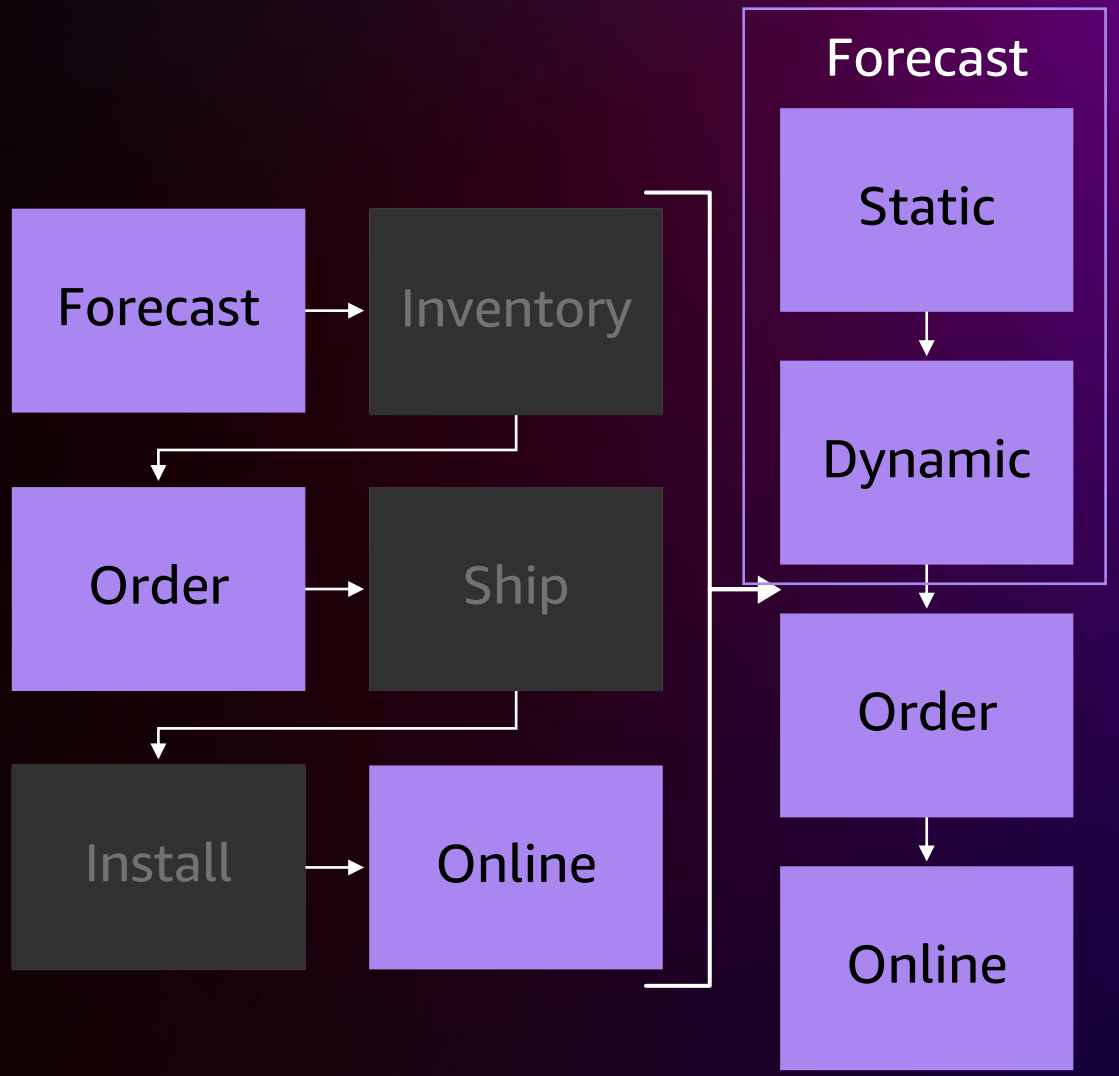


© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

EDA On Cloud

- Looks and feels like on-premises data center
- Instantaneous deployment of additional compute

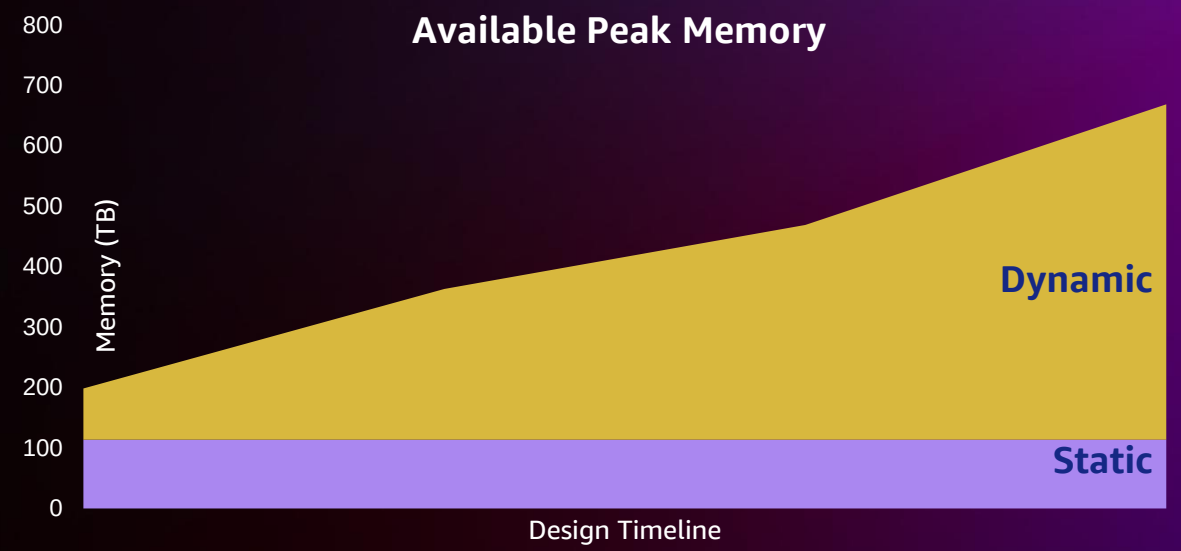




aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Forecasting

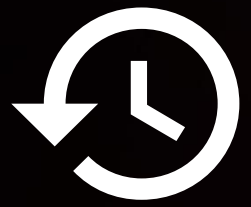
- Finding the balance
- Static Machines
  - Available all the time
  - Goal max utilization
- Dynamic Machines
  - Forecasted in Hours/Month
  - Available 20 minutes
- Storage
  - Desire high disk utilization without hitting 100%





aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Innovate faster



Beyond Machine TAT

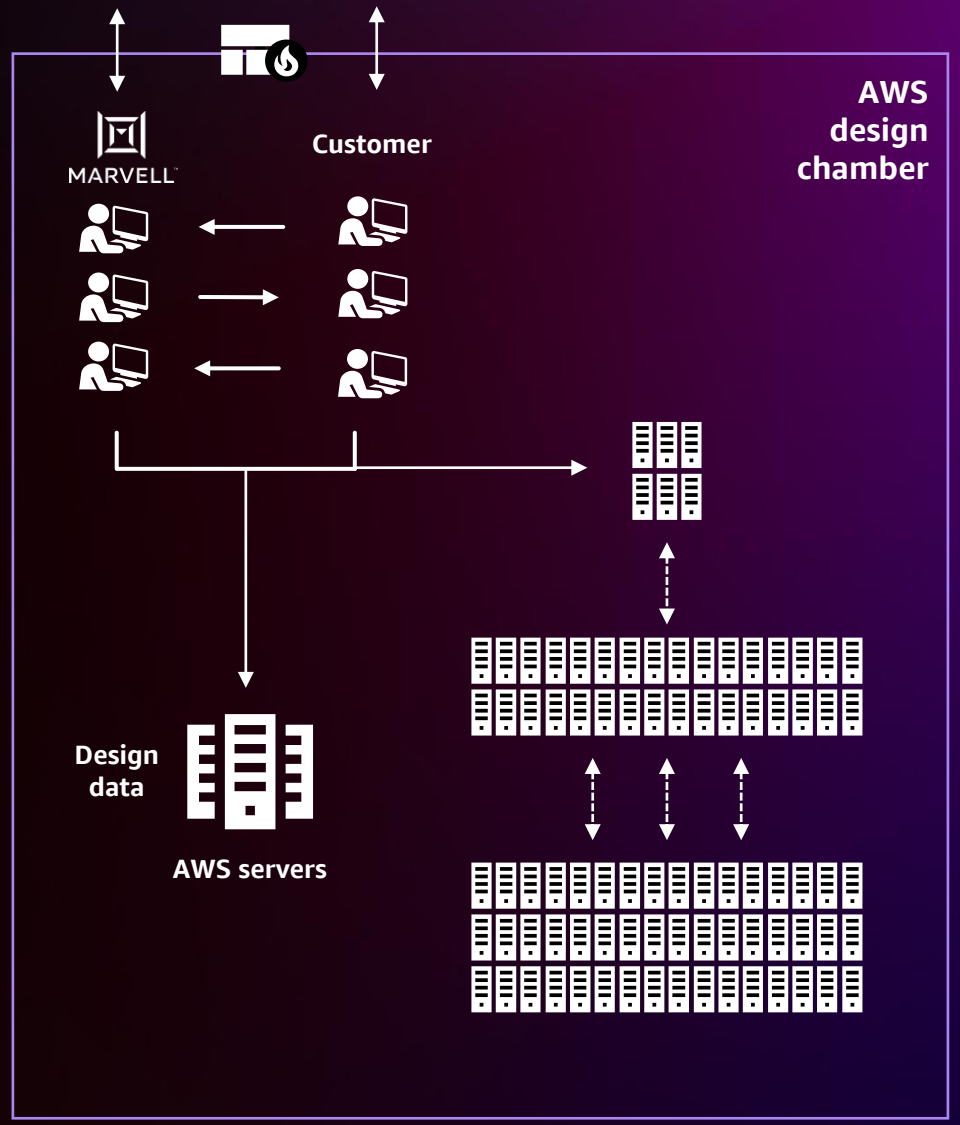


Shared cloud

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Shared cloud

- True partnership
- Shared Design Environment
- Direct Engagement
  - RTL/netlist handoffs
  - Report sharing
  - Updates/ECOs
- Joint team processing the design



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Cloud-optimized silicon

Every cloud application is unique

IaaS

PaaS

SaaS

AI/ML

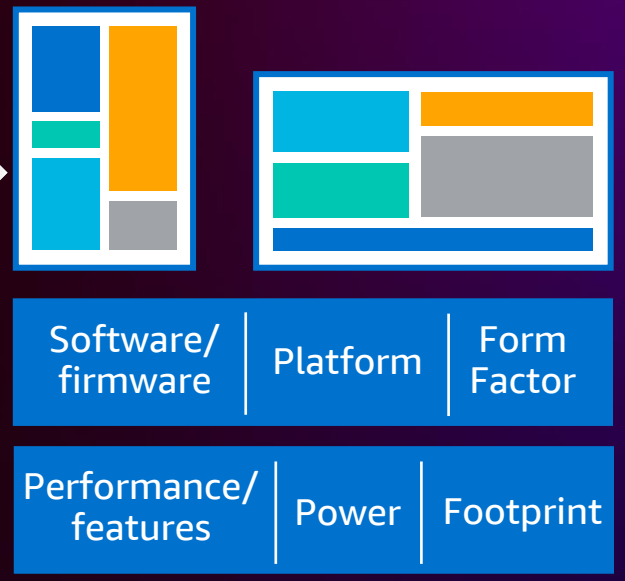
Search

Social media

Cloud-optimized silicon is necessary



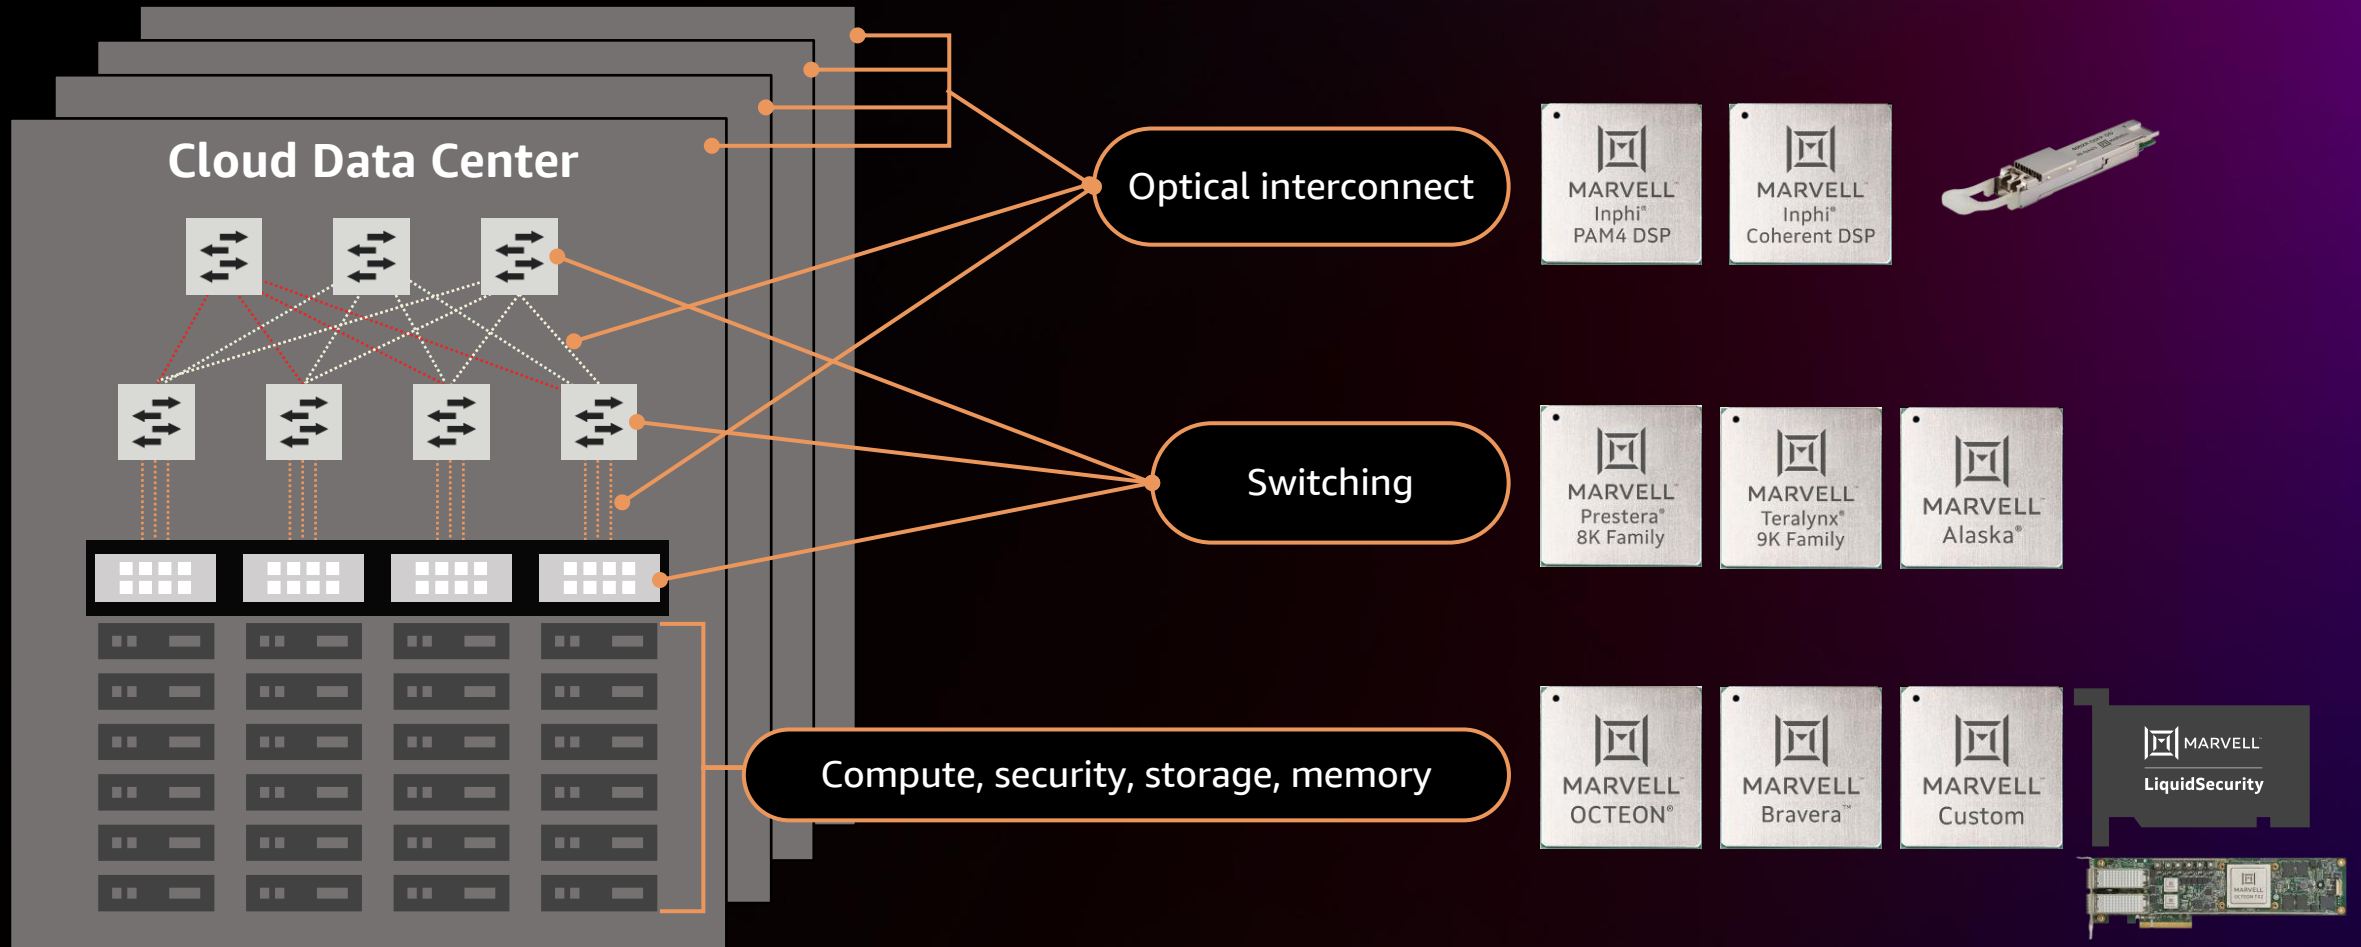
General-Purpose



aws

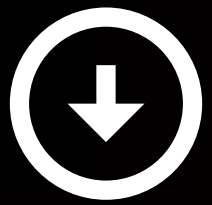
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Unique cloud-optimized silicon portfolio



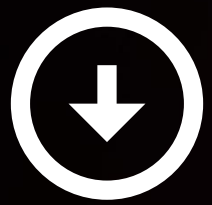
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Key Takeaways



Run times

New machines



Wait times

Dynamic resources



TAT

Time to market



Productivity

Innovate faster

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.



Summary

Successful partnership on your journey on designing silicon in the cloud

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Thank you!



Please complete the session survey in the **mobile app**

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.