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

ADM201

Ad technology innovation with NBCUniversal and FreeWheel

Ryan McConville
EVP, Ad Platforms
NBCUniversal

Jeffrey Pinard
SVP/CIO, AdSales Technology
NBCUniversal

Xinyu Huang
VP of Data
FreeWheel

Stephanie Layser
Global Head, Publisher Ad Tech Solutions
AWS

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

NBCUniversal







Comcast

Local

Syndication

Films

Parks


One
Platform
NBCUniversal
YouTubeTV
Roku
COMCAST
CABLE ONE
sling
Samsung SMART TV
pluto tv
peacock
CONCERT
NBC NEWS
OXYGEN
sling
Spectrum
DIRECTV
tv
CNBC
NBCSports
VIZIO
AT&T tv NOW
fubo TV
NBC
FYF
verizon
firetv
ATLANTIC broadband
SONY
hulu
E!
chromecast
FANDANGO
COX
U-verse
usa
iOS
Rotten Tomatoes
GOLF
androidtv
WOW!
bravo
VUDU
RCN
android
Vox
XUMO
MSNBC
UNIVERSAL KIDS
XBOX
MEDIACOM
TELEMUNDO
dish
YouTube
ellation
BuzzFeed

For years, marketers have had to make a choice between...

The world of big tech

Social platforms

User-generated content

Online influencers

Big first-party data

Proprietary technology/automation

Outcome-based measurement

+

The world of big media

Premium video platforms

World-class, professionally produced content

Industry-leading talent in front of and behind the camera

Small panel-based data

Manual processes

Traditional demo measurement

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

With the rise of streaming and the digitization of TV, you don’t have to choose

= One Platform
NBCUniversal

NBCU’s One Platform is TV, modernized

The first global video platform to **combine** the power of professionally produced, culture-defining content **with** the benefits of big data, precision targeting, automated buying, and outcome-based measurement

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

One Platform helps marketers break with legacy across 3 key pillars

TO MAKE THE MOST OF THEIR PREMIUM VIDEO CAMPAIGNS

1 Audiences

Small data	Broad targets
Big data	Precision targets

2 Activation and automation

Linear and digital planned separately	Highly manual and siloed
Planning across platforms	Automated and interoperable

3 Measurement

Panel-based and platform-specific	Delivery only
Census-based and unified	Business outcomes

aws

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

One Platform is powered by three corresponding product suites

One Platform
NBCUniversal

NBCUnified

- NBCUnified Audiences
- NBCUnified Consumer Match
- NBCUnified Access

adsmart

- AdSmart Linear (DDL)
- AdSmart Streaming (Peacock AX)
- AdSmart Cross-Platform (XP)
- AdSmart Exchange (Prog Digital)
- AdSmart APIs (Self-Service Linear)

NBCU NBCUNIVERSAL MEASUREMENT PARTNER

- Certified Currencies
- Certified Outcomes
- Audience Insights Hub

aws

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

We have been building toward One Platform for years

LEADING THE MEDIA INDUSTRY WITH FIRST-TO-MARKET INNOVATION

2016

Launched data-driven linear

Launched premium video PMP

2017

Launched self-service linear (PTV)

2018

Launched multi-platform measurement

2020

Launched cross-platform optimization

2021

Launched privacy-safe data interoperability

2022

Launched enterprise data and identity platform (NBCUnified)

Launched Measurement Certification Program

aws

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

A look under the hood at the One Platform tech stack and where AWS comes in

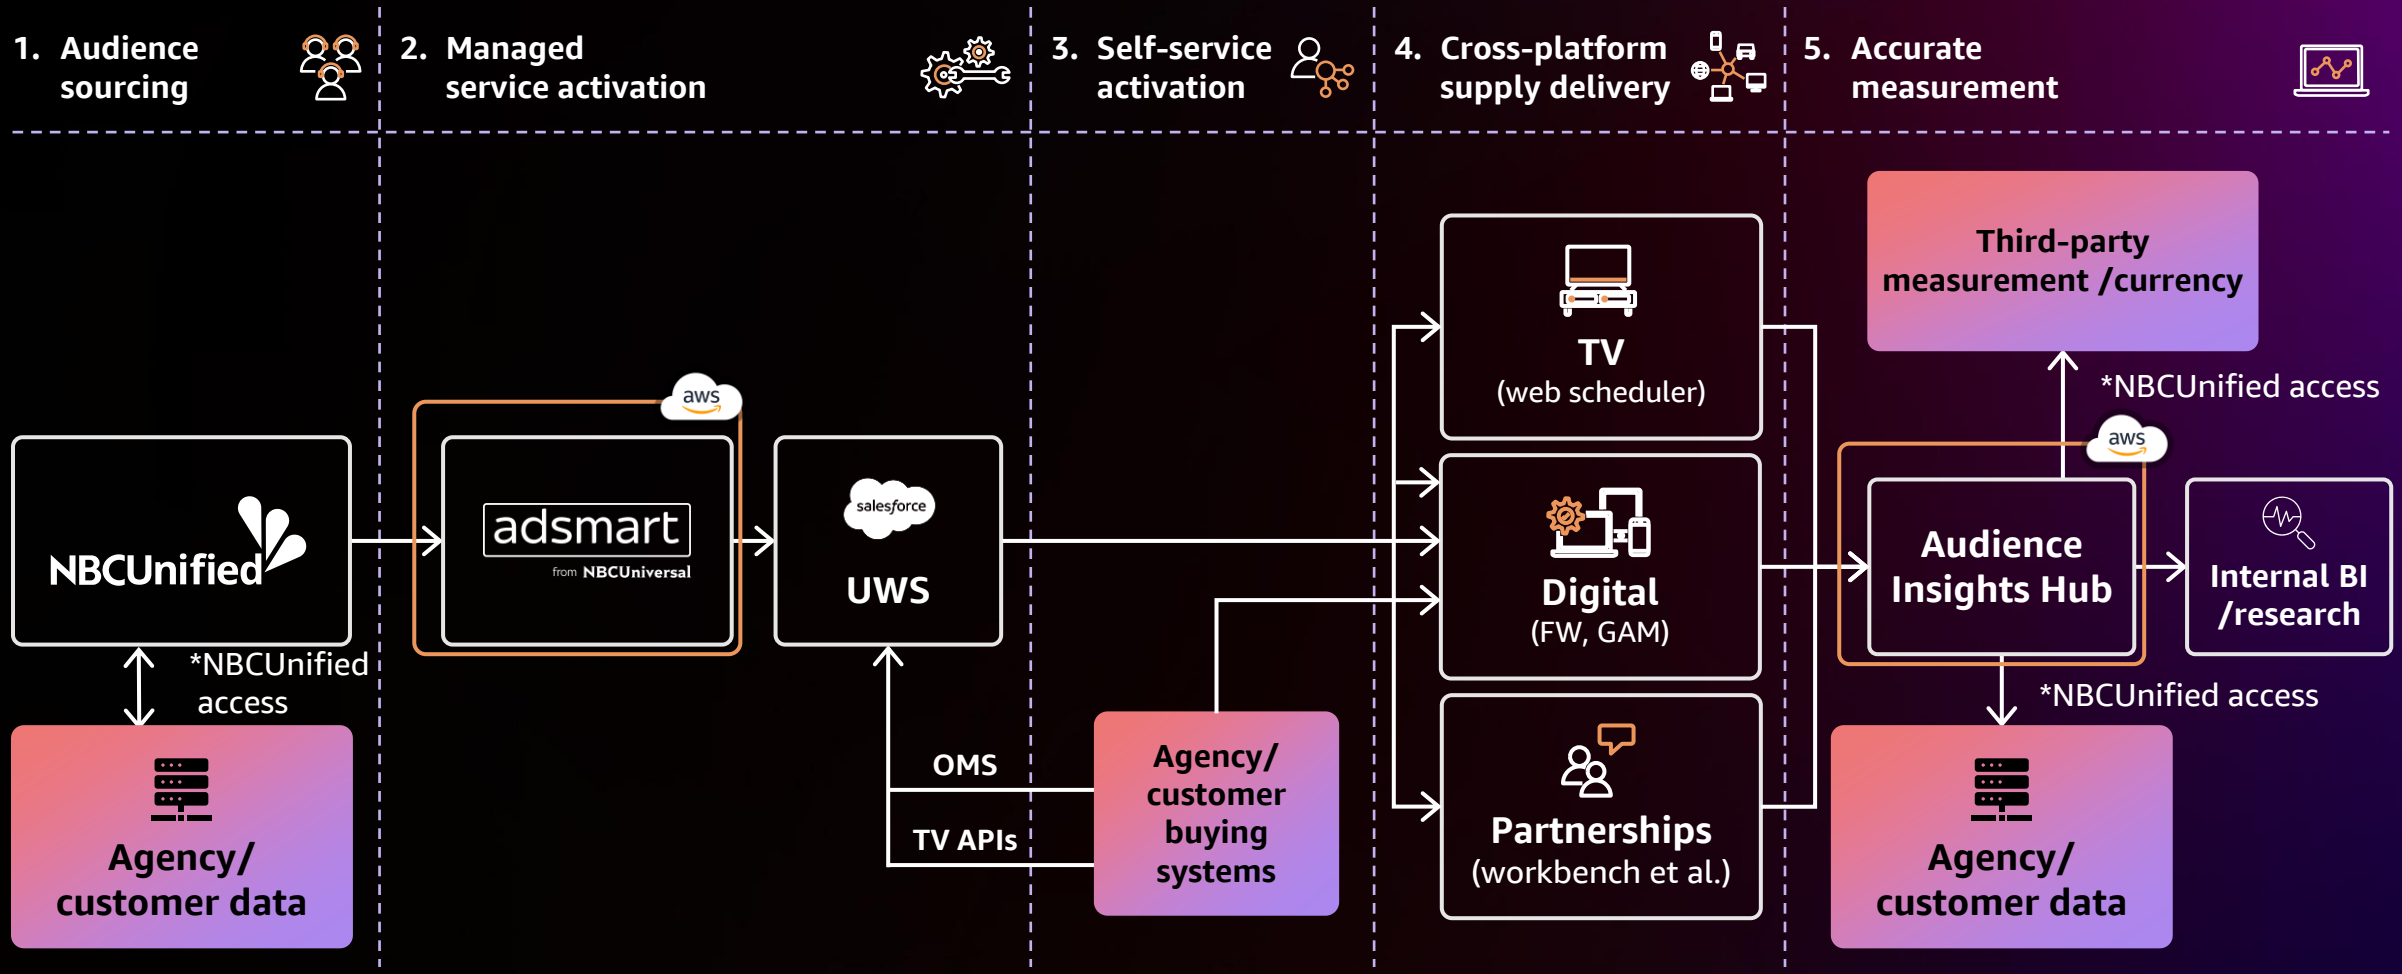

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Transforming TV ad planning, delivery, and monetization

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

NBCUniversal
Operations & Technology
TECHNICAL OPERATIONS

The challenge

Transform the way NBCU plans, delivers, and monetizes ads from the segregated monolithic digital and linear platforms of yesterday to a unified distributed planning, trafficking, and monetization cloud-based platform. This new platform will deliver better outcomes to NBCU and its clients by enabling a unified view of linear and digital inventory and delivery metrics.

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

Viewership data sizing

54TB

Comcast STB data

10TB

iSpot ad delivery data

20TB

FreeWheel ad logs

aws

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

Current state: Siloed digital and linear systems walled off from each other

RESULTING IN A TOTAL SEPARATION OF THE PLANNING AND TRAFFICKING OF EACH AD PLATFORM

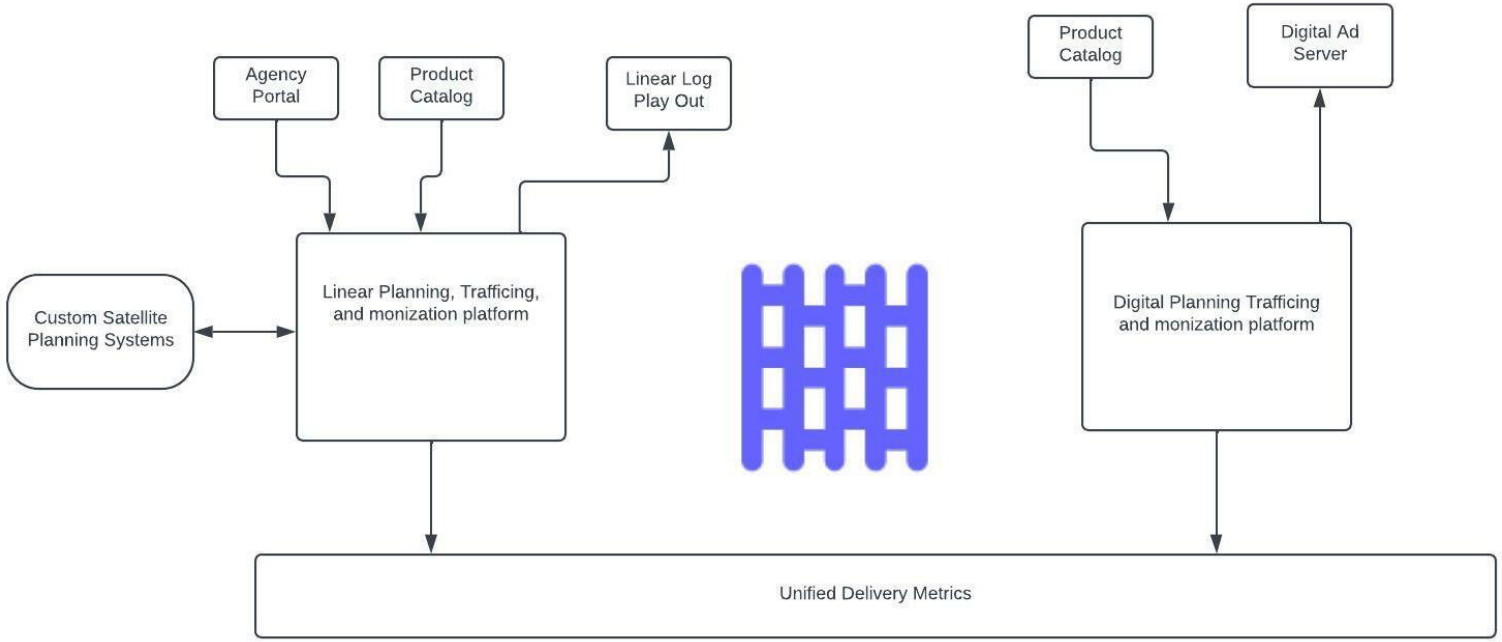

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Future state: Unified systems distributed by function

RESULTING IN A TOTAL SEPARATION OF THE PLANNING AND TRAFFICKING OF EACH AD PLATFORM

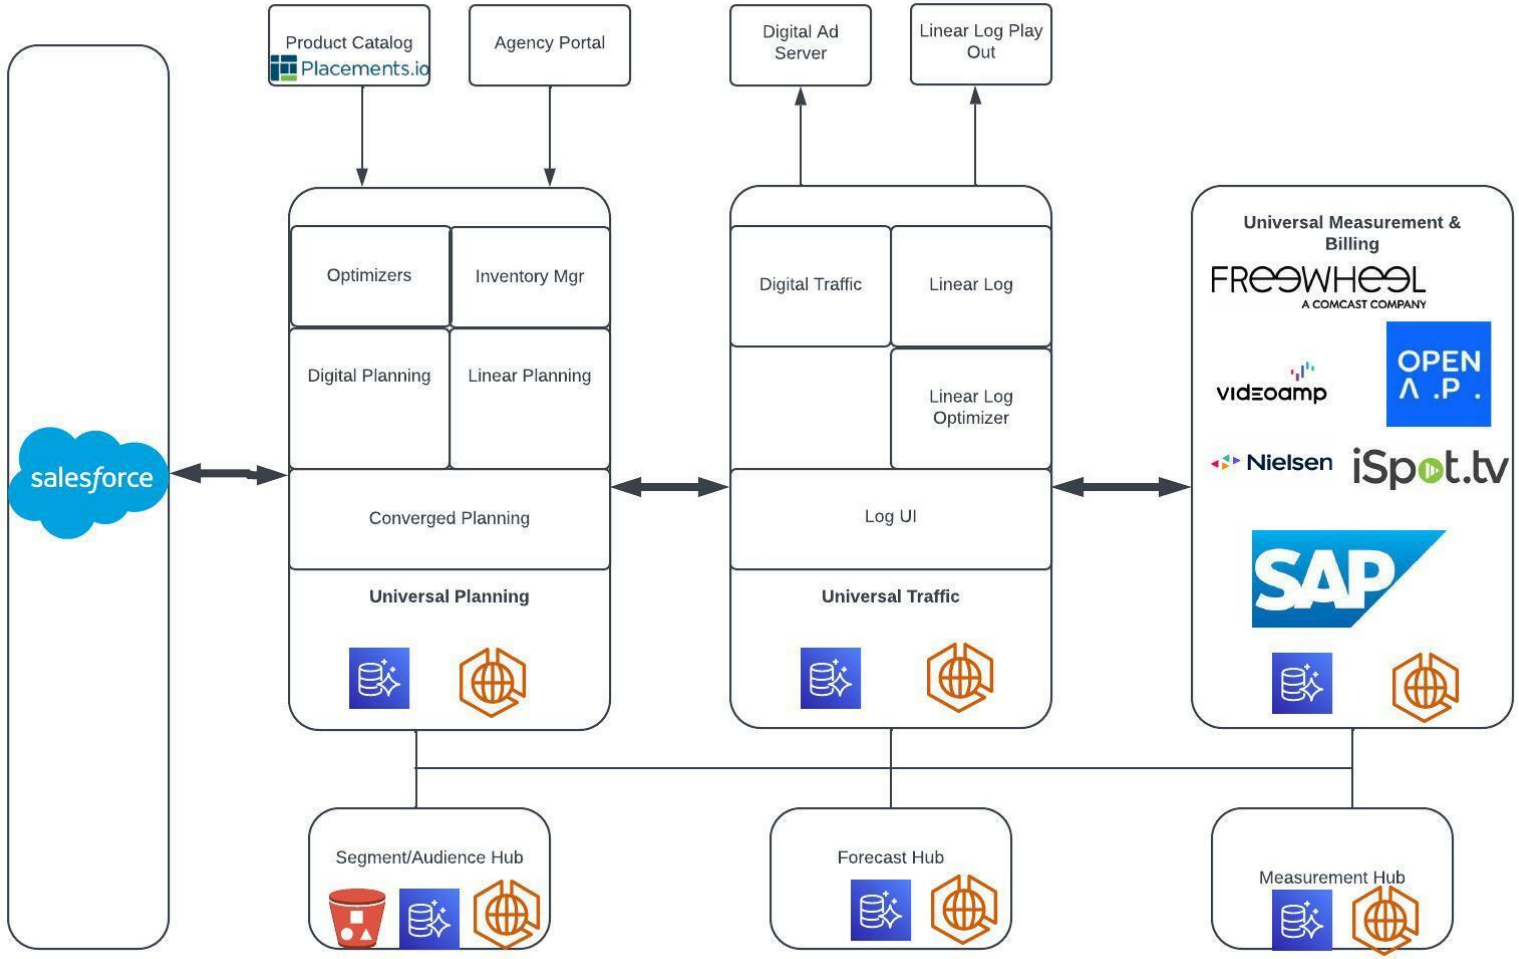

aws

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

Currency and forecasting pipelines

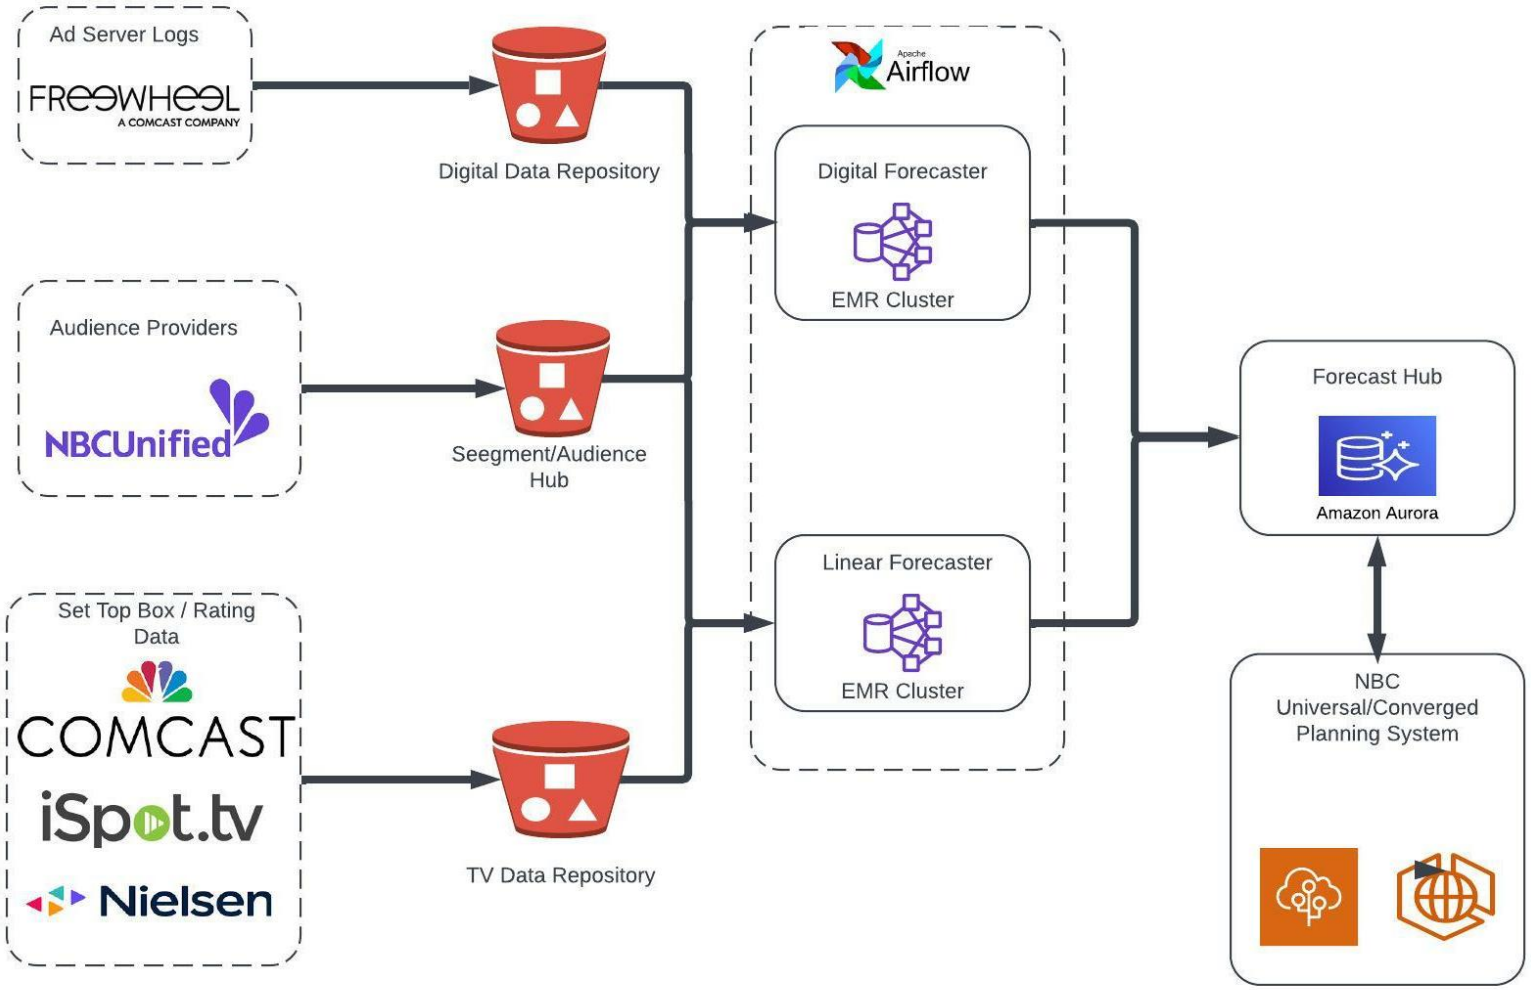

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

NBCU Ad Sales: Big data journey

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

Where we started

AoA Data Center: Hortonworks Big Data Cluster

Home of NBCU's Ad Sales data products:

- 2,432 cores
- 24TB RAM
- 2,688 TB storage

OnPrem ecosystem: Prior to June 2020

- Expensive to run and maintain
- Out of data center floor space
- Out of processing power
- Unable to react to exponentially growing business needs

Desired end state: June 2020–present

- Standardized S3 data ocean
- Modernized data access using purpose-built technology
- Metadata-driven infrastructure and code
- Design paradigms and IP protection

aws

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

Partnering with AWS

- ✓ Joint architectural reviews
- ✓ Individual well-architected solutions
- ✓ Migration acceleration program
- ✓ Weekly check-ins
- ✓ Event planning
- ✓ Quarterly cost optimization reviews

AWS Account Support & TAMS + **NBCUniversal Engineering**

One Team

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

AWS data lake journey

Moved our big data platform from on premises to cloud in 1 year

- 6 months of cost/performance optimization
- Readiness in time to handle Summer 2020 Olympic Games

Tailored machine configurations based on tasks

Now focusing on solving additional business challenges/use cases (clean rooms, data sharing, data democratization)

15.4TB

Interactive reporting data

4PB

Data in data lake

$35M+

Cumulative savings over 10 years

40%

Reduction in overall cost

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

Optimizations in AWS

Spot instances

- ✓ Spot instances with 99.65% initial success rate
- ✓ Machine equivalence matrix
- ✓ Failover to on-demand

Amazon EMR

- ✓ Ephemeral and persistent clusters
- ✓ Spot fleet instances

databricks

- ✓ Introduction of delta lake for S3 updates
- ✓ Additional cost savings for batch jobs

AWS Lambda

- ✓ Serverless components
- ✓ Efficient data streaming

aws

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

Optimizations in AWS global platform

Modernizing the code base

- ✓ Metadata-driven
- ✓ Basis of a unified development platform

Formation of DevOps team

- ✓ Introduction of Terraform
- ✓ Introduction of DataDog for monitoring

aws

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

Case study: 2020 Summer Olympics

AIRED IN SUMMER 2021

Overview

- Data volumes **increased 950% from Rio Games with no latency**
- AWS infrastructure **scaled without intervention**

Switch from Spot to On-Demand

- **No code releases** necessary
- **No failures** related to infrastructure
- **Seamless hardware** scaling

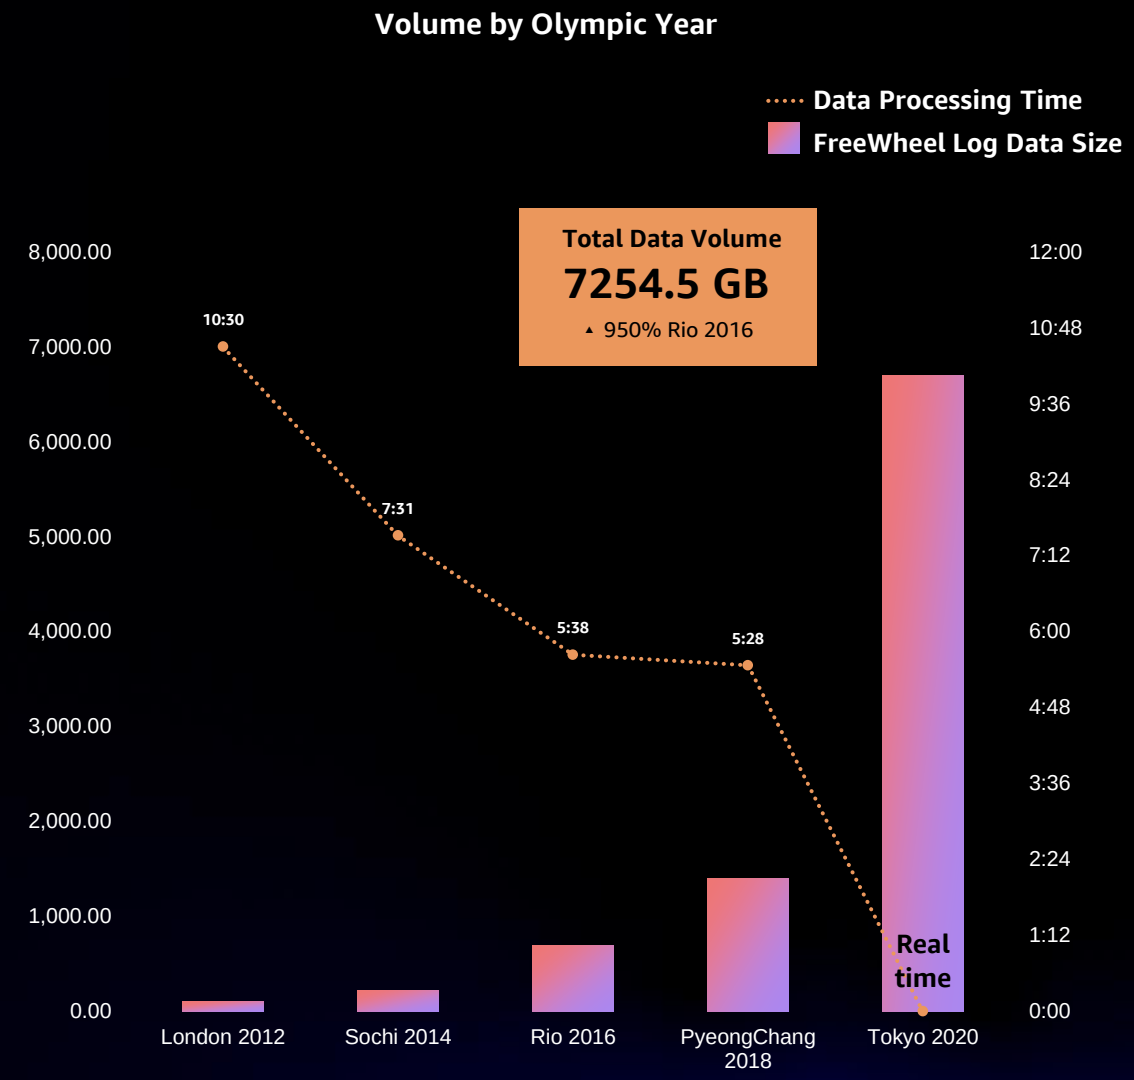

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Where we are now

100% NATIVE TO AWS

- ✓ Vendor agnostic
- ✓ Cost optimized
- ✓ Ever evolving

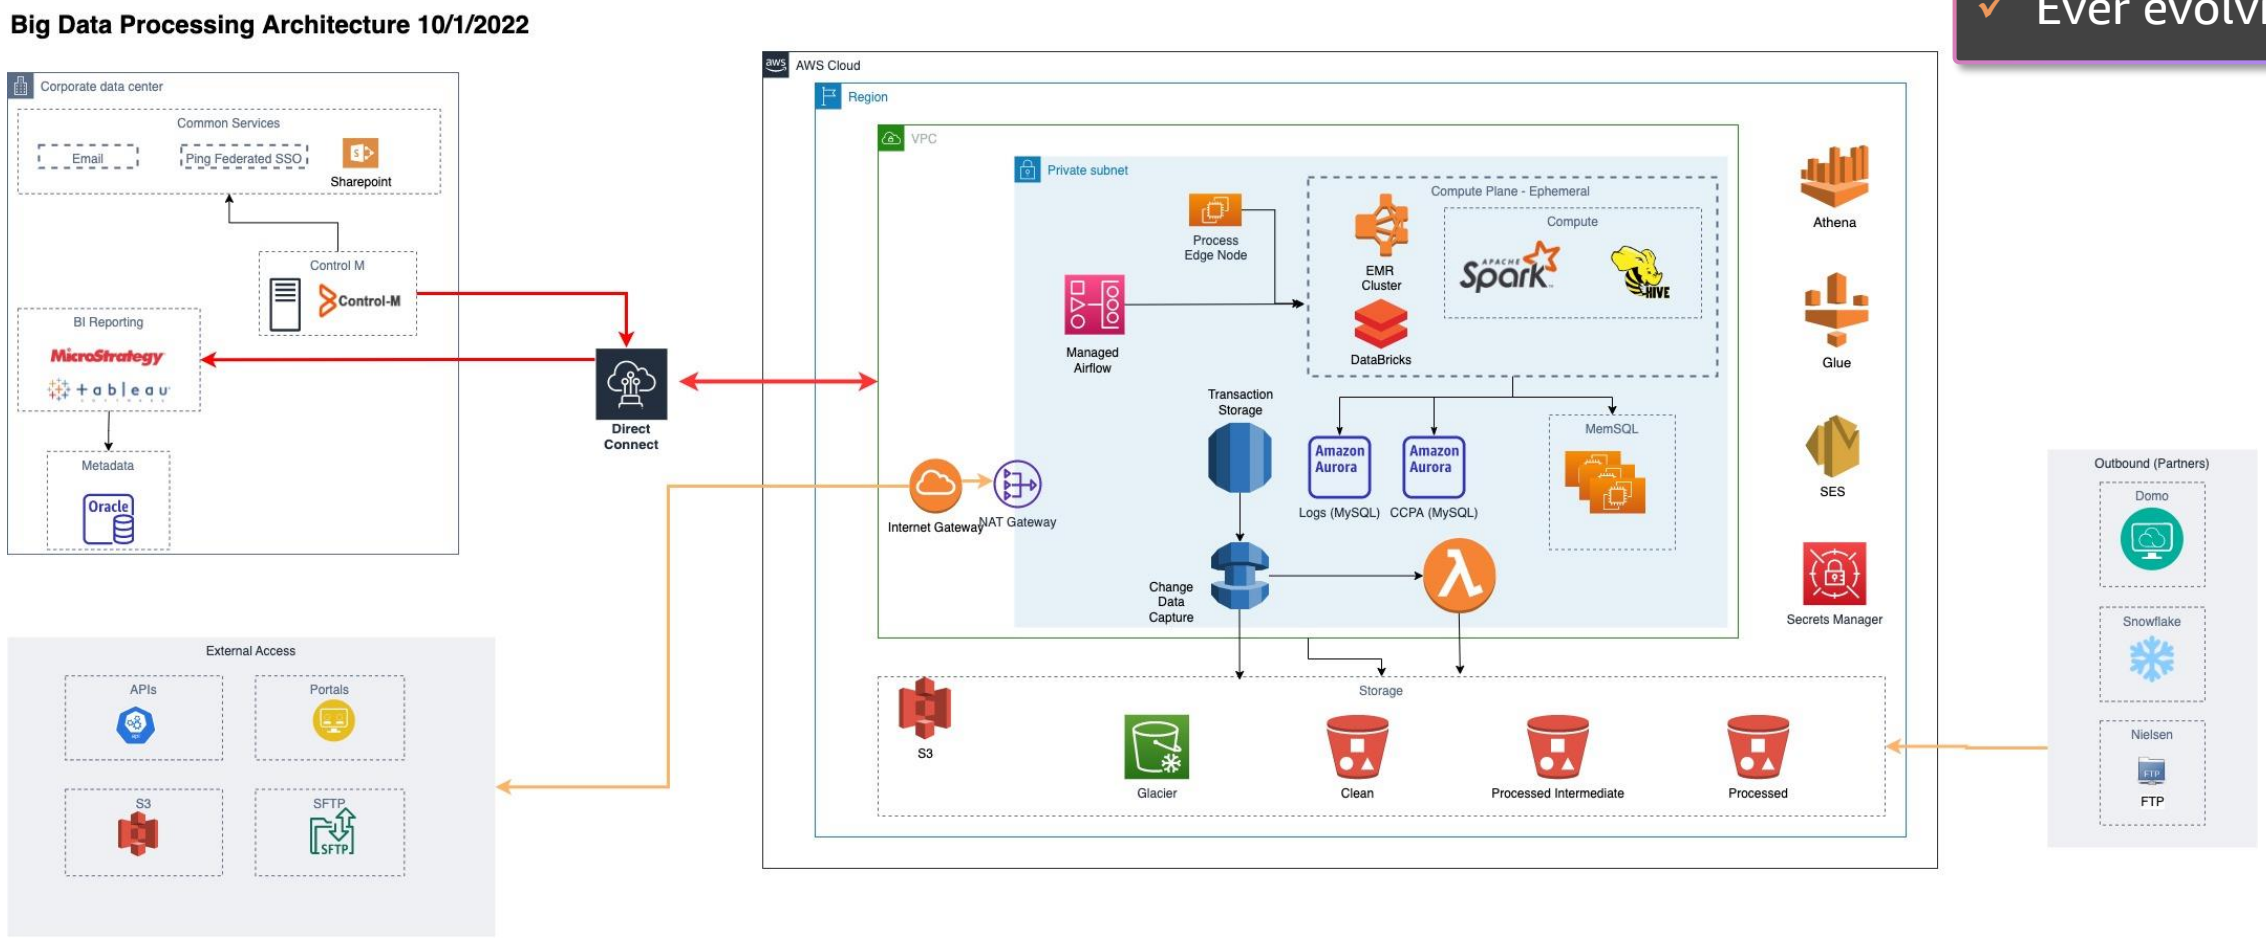

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Where we're headed: Data-as-a-Service

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

An architectural vision

COOPERATING INDEPENDENT SERVICE PLANES

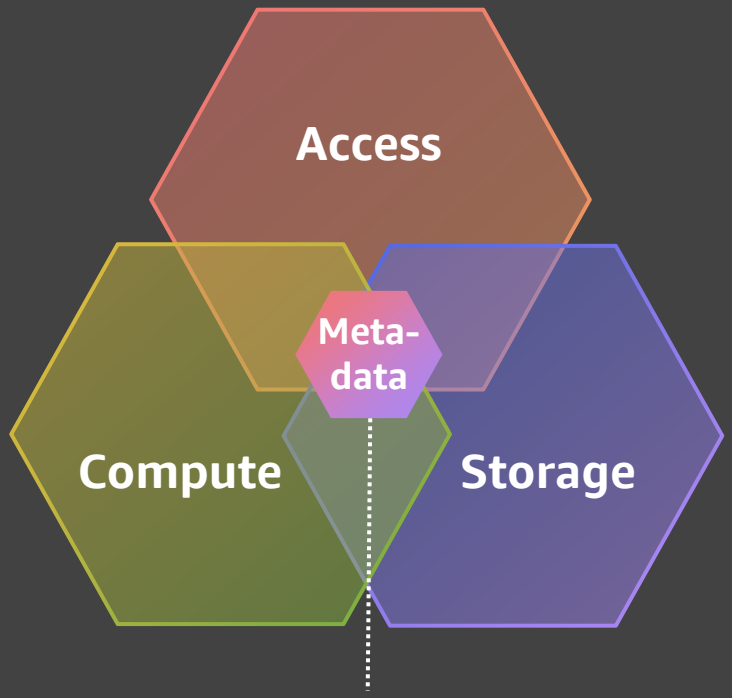

Fully configurable metadata driving and unifying all service planes

Access

- Optimized access to canonical data in the storage plane
- Customized to a variety data consumption patterns: JDBC, APIs, files
- Agnostic to front-end applications: BI tools, custom apps, etc.

Compute

- Cost-efficient ephemeral computational clusters
- Institutionalized business data transformations and mappings
- Parallelized development and continuous improvement

Storage

- Democratized canonical data (single source of the truth)
- Common data lake for all applications to store and exchange data
- Cost-efficient, secured and privacy-compliant data-at-rest
- Agnostic to data access patterns

aws

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

Key definition: Canonical data

DISTINGUISHING BETWEEN CANONICAL DATA AND PRIVATE DATA IS CRITICAL

A reliable published endpoint of data on which business analysis can be driven and technical dependencies may be built

- ✓ Cataloged
- ✓ Quality assured
- ✓ SLA enforced
- ✓ Change controlled
- ✓ Managed lifecycle

aws

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

Data-as-a-Service

AD SALES DATA MANAGEMENT APPROACH FOR HOLISTIC DATA LAKE STEWARDSHIP

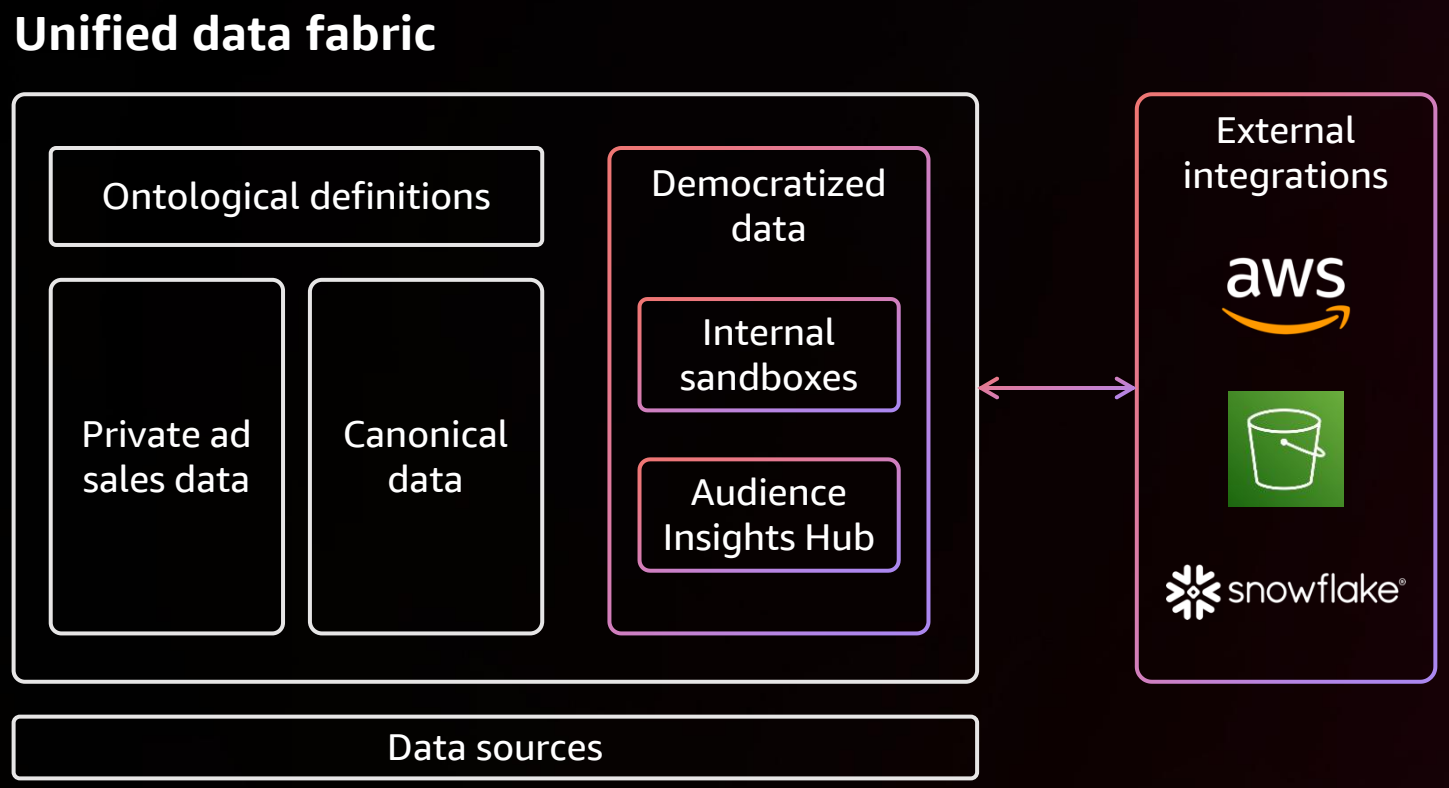

- ✓ **Single source of truth**
- ✓ **Lowest common denominator**
- ✓ **Exportable to a wide variety of technologies**
- ✓ **Democratizable**
  - Internally and externally
  - Clean room integrations with Snowflake, Safehaven, etc.

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

aws

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

FREEWHEEL

A COMCAST COMPANY

We are the platform that makes premium video advertising work, serving >90% top media companies

aws

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

FreeWheel infrastructure highlights

Every second

1.2M
concurrent requests

0.5M
ads served

Every day

20B
ad serving events logged

40TB
data processed

Every year

~10T
requests handled

20PB
data for reporting, analytics, ML

aws

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

Do more/faster for less

FREEWHEEL'S JOURNEY IN GRAVITON MIGRATION

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

FreeWheel Compute infrastructure on AWS

COMPUTE

Billions
yearly compute hours

10s of millions
yearly EC2 instance launches

EKS clusters scale up to thousands of EC2 instances to support daily operations and live events, handle 100x request spike within 2 seconds

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

Doing the same/more for less with Graviton

~20%
cost savings*

5–20%
similar or better performance*

*With already migrated workflows and systems

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

Test, measure, analyze, then decide

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

Instance type	Node number	Workload	CPU	Elapsed time(s)	Cluster cost ($/hour)
C5.9xlarge	8	TPC-H	288	1,805	12.3
C6g.8xlarge	9			1,644	9.8

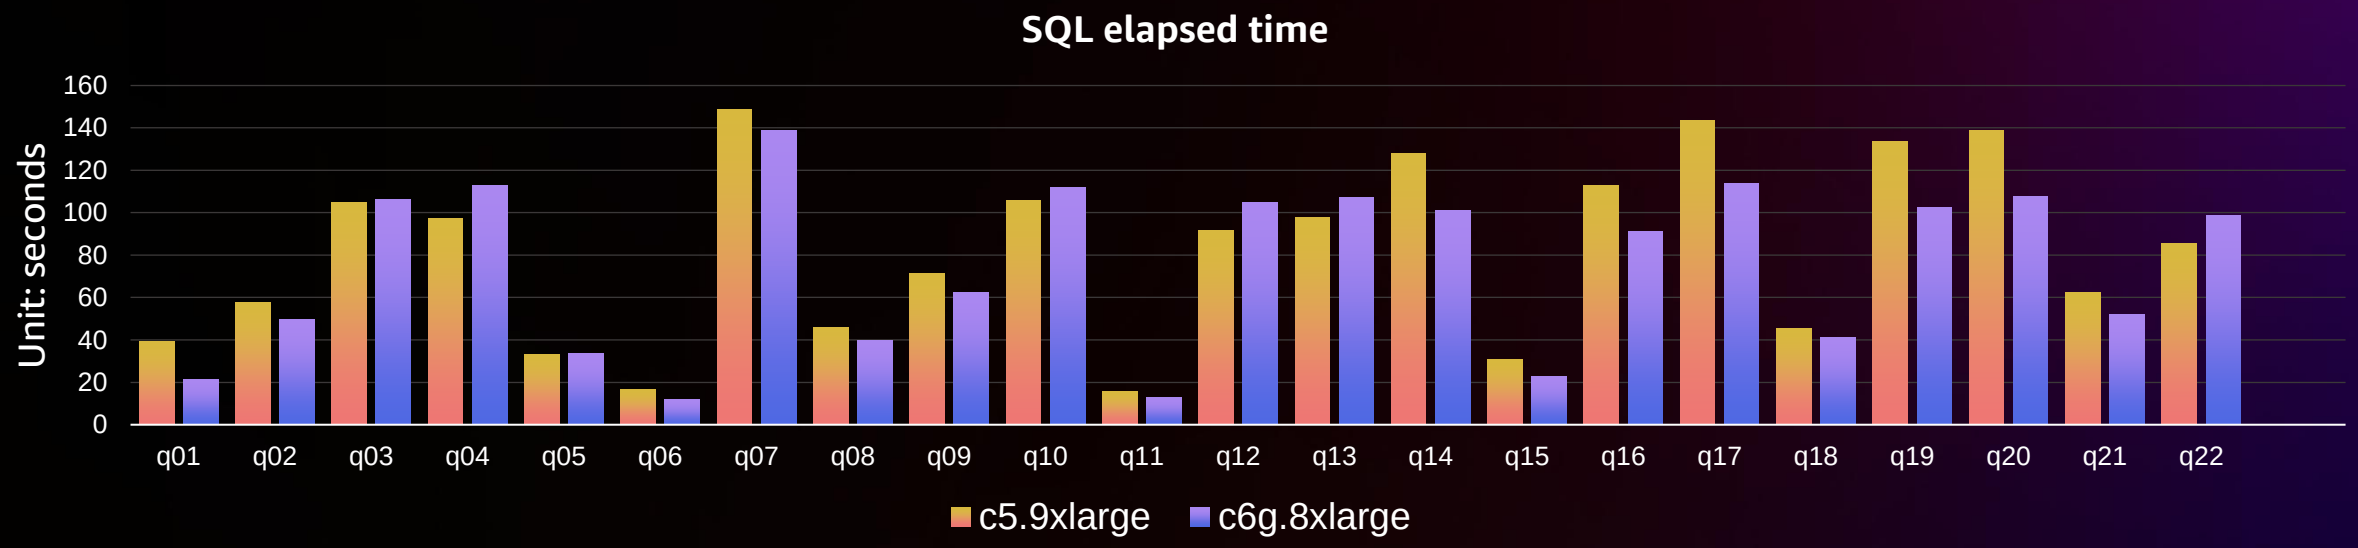

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Upgrade end-to-end

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

Instance type	Workload	Java	Nodes	CPU	Elapsed time(s)
C6g.8xlarge	TPC-H	JDK8	1	32	1,725
		Correto-11+Graal			1,821
		Correto-11			1,408

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

Avoid bifurcations

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

Check Graviton availabilities

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

Measure, test, analyze

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

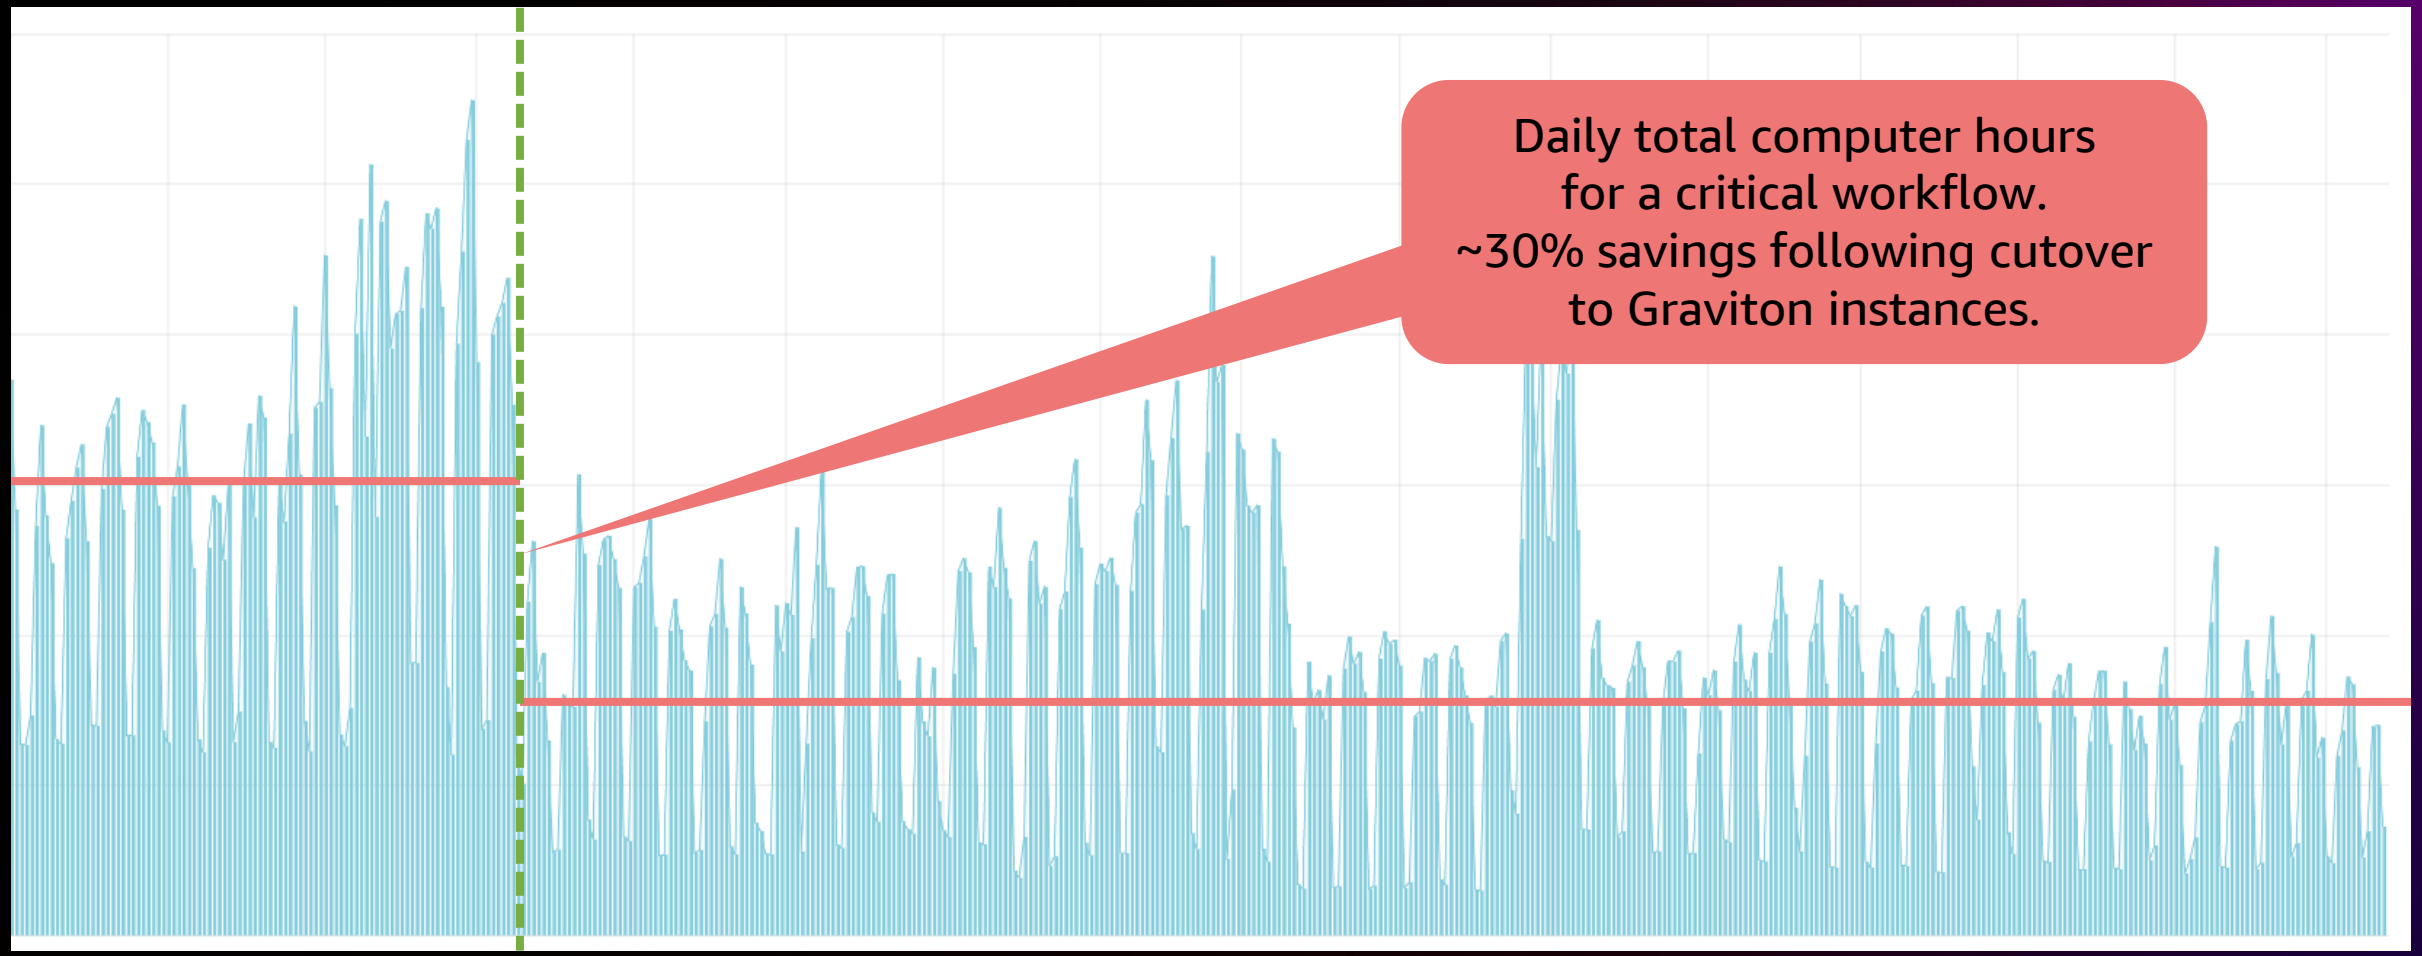

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Doing the same/more for less with Graviton

~20%
cost savings*

5–20%
similar or better performance*

*With already migrated workflows and systems

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

Flatten the storage cost curve

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

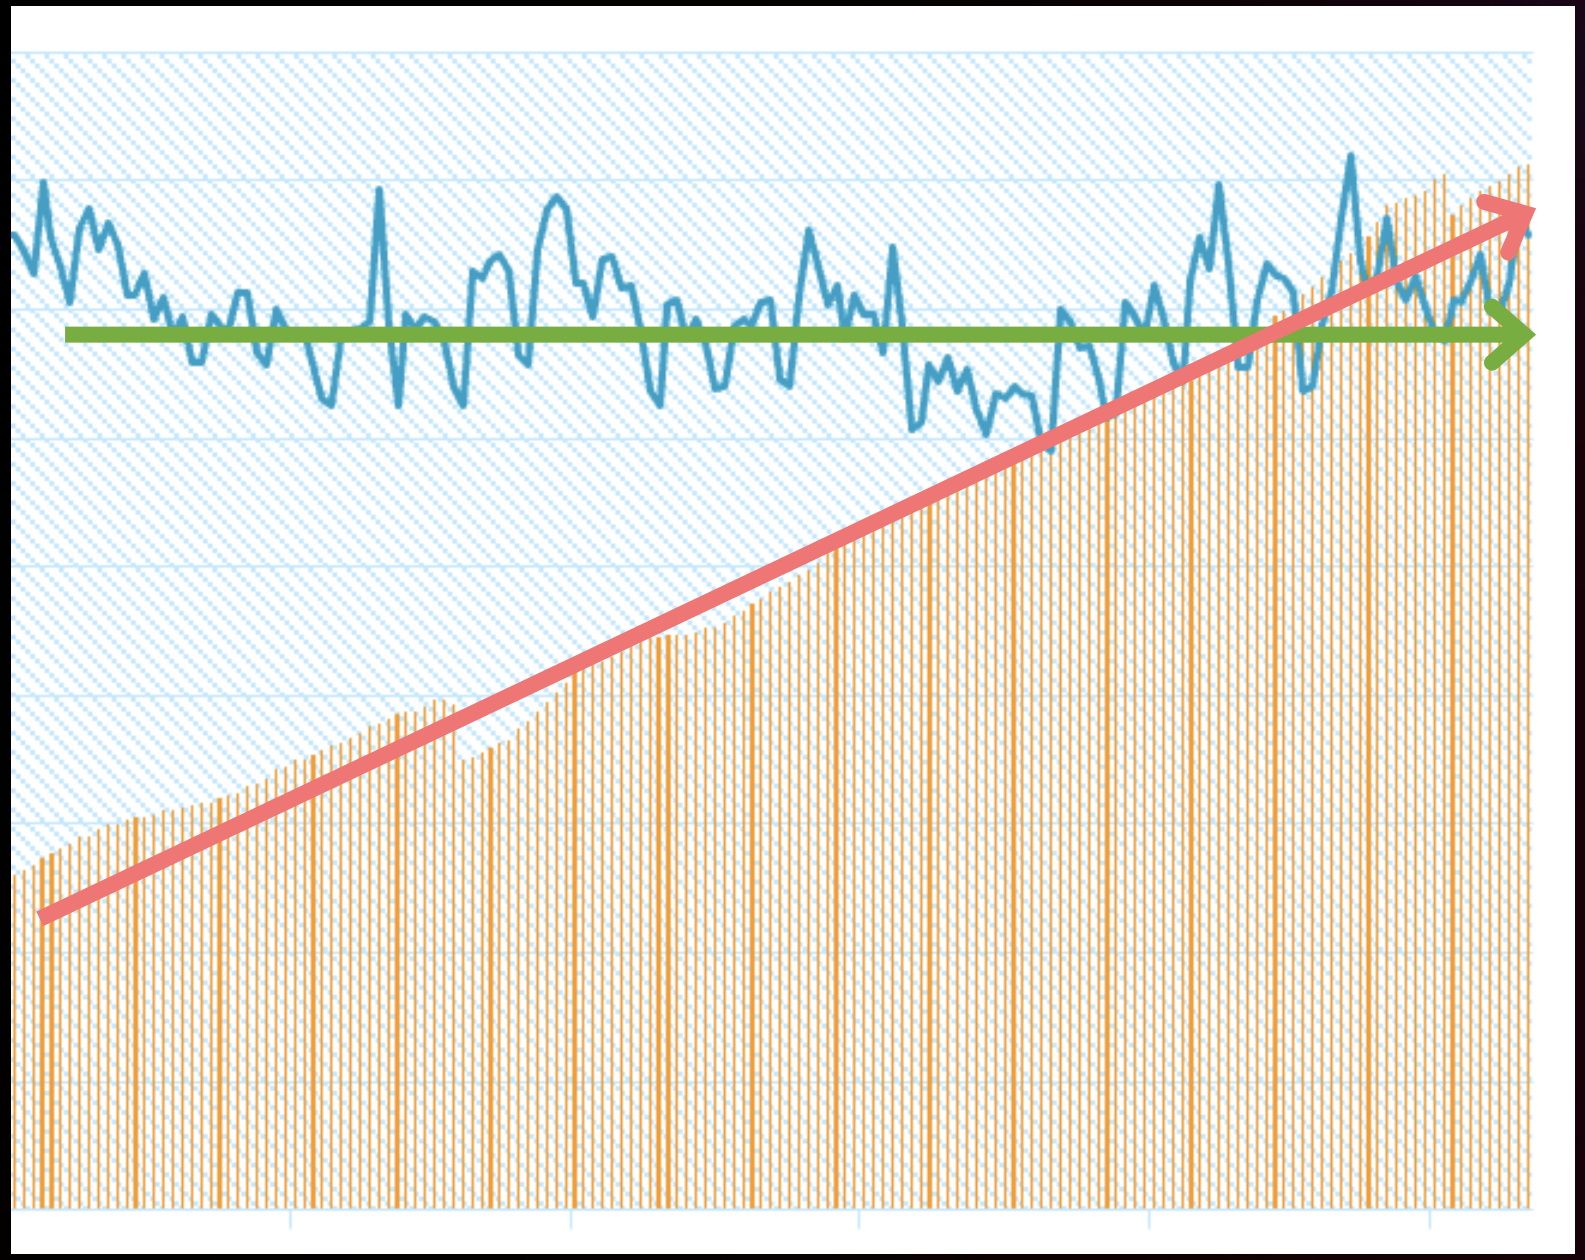

Data volume: +30%

S3 cost: ~+0%

aws

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

Can we do better?
With help from 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.