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

ANT338

Scaling data processing with Amazon EMR at the speed of market volatility

Sambath Velusamy
Director, Technology
FINRA

Meenakshi S Ponn Shankaran
Senior Big Data Architect, Professional Services
AWS

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

Agenda

Who is FINRA, and what is FINRA’s mission?

What is the Consolidated Audit Trail (CAT)?

What are the challenges in implementing the CAT to process billions of events daily?

How we optimized the CAT application?

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

FINRA – Who We Are



Who do we work with:

620K+ Brokers

3400 Broker-Dealer Firms

32 Exchanges

How we process:

300K+ Amazon EC2 instances per day,

9,000+ Amazon EMR clusters per day

FINRA CAT Volumes:

600B+ Peak Market Events

130K+ Files per Day

500+ PB of storage

FINRA process

2 trillion + data validations per day

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

Consolidated audit trail

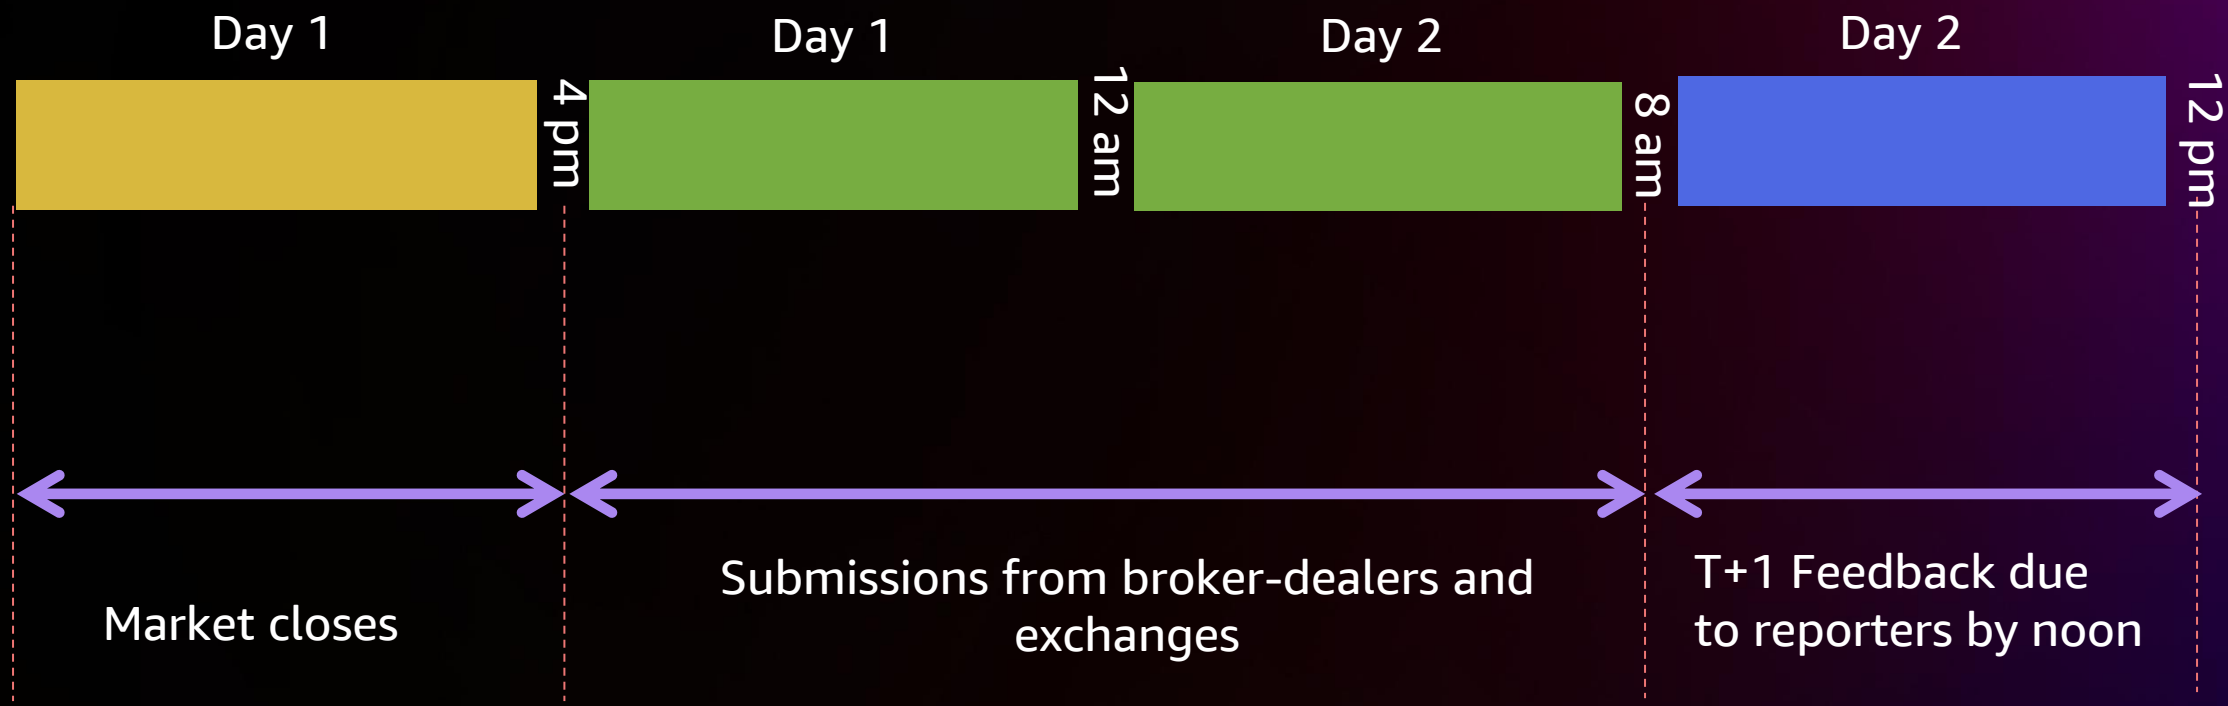
SEC created Rule 613 to implement a comprehensive consolidated audit trail that would allow regulators to efficiently and accurately track all activity throughout the U.S. markets in National Market System (NMS) securities.



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

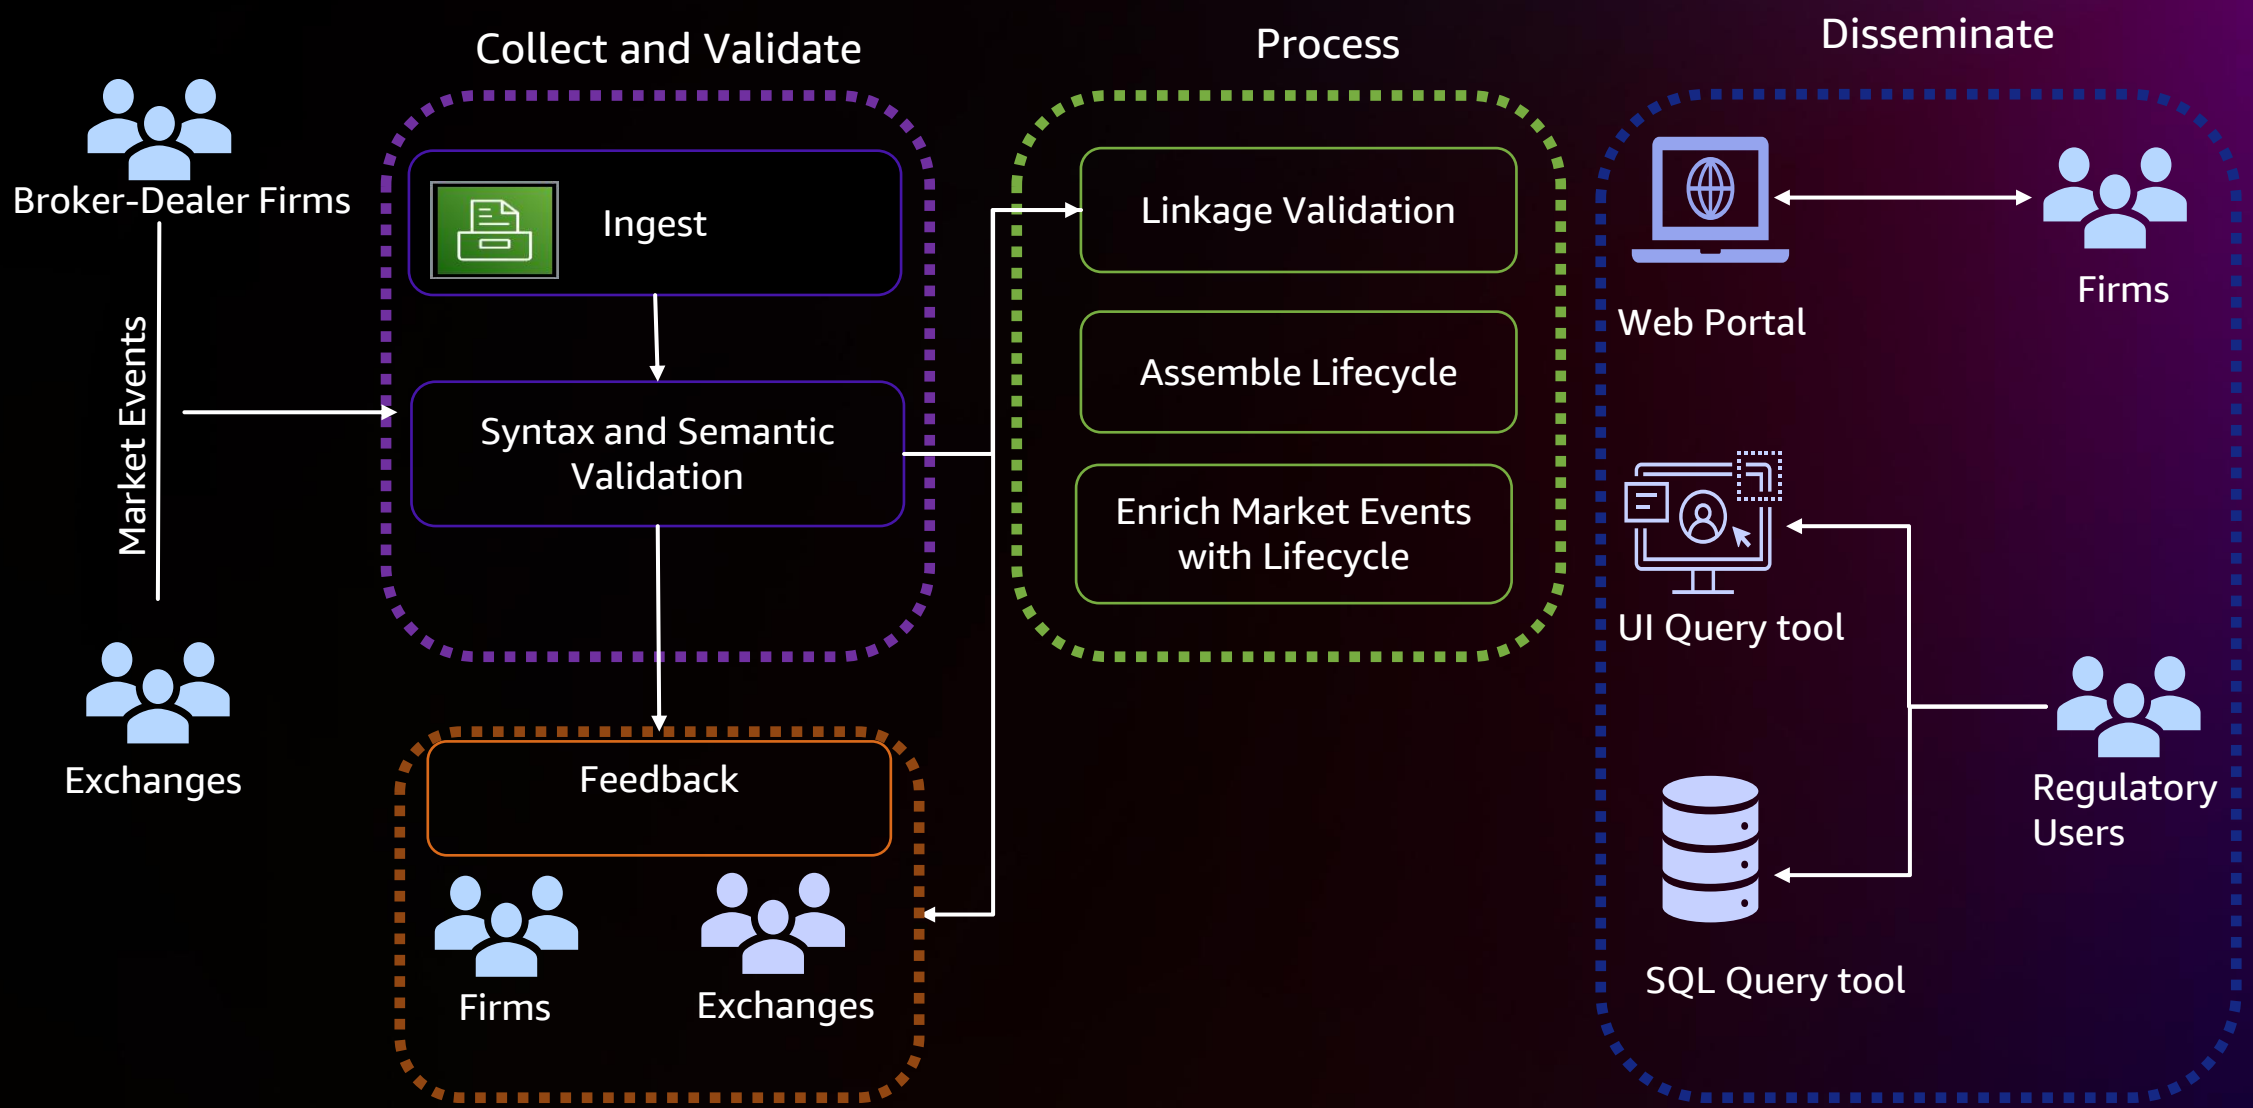
CAT SLA

All the trade events for the Day T has to be processed and feedbacks for linkage errors need to be reported back to broker-dealers and exchanges by T+1 noon



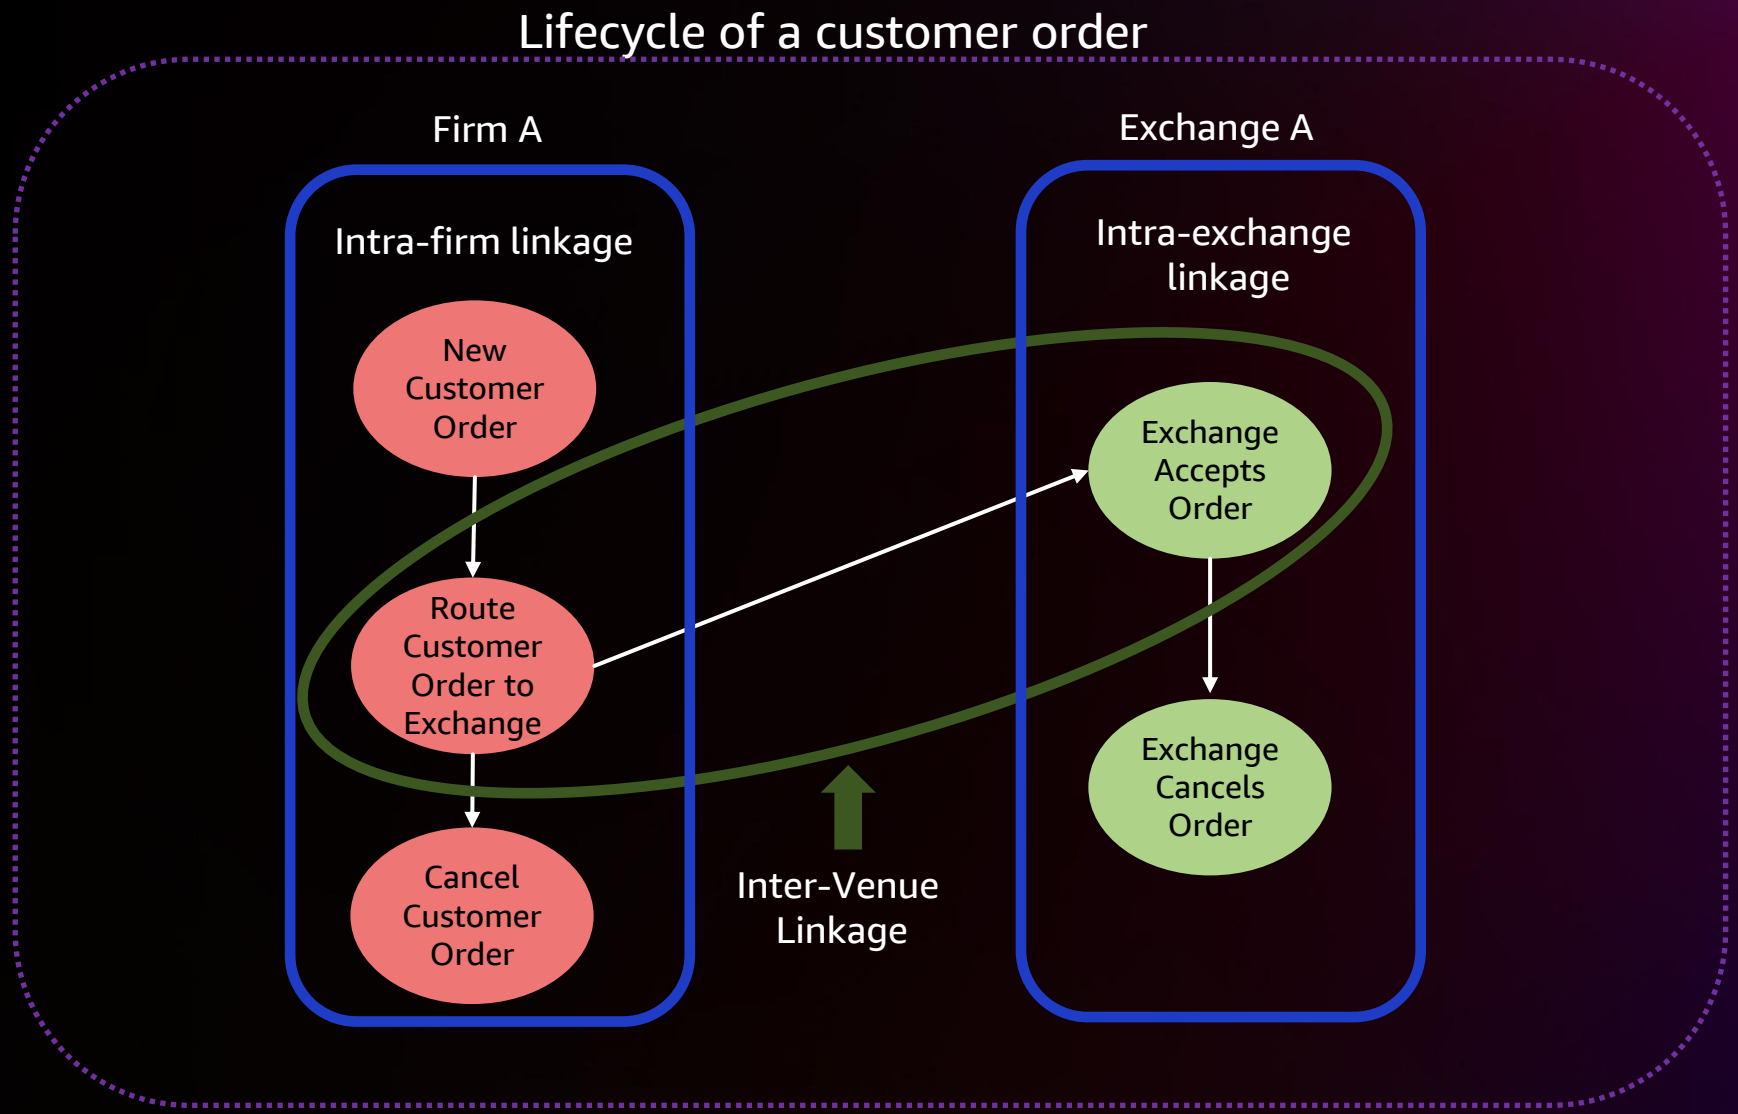
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

CAT process flow



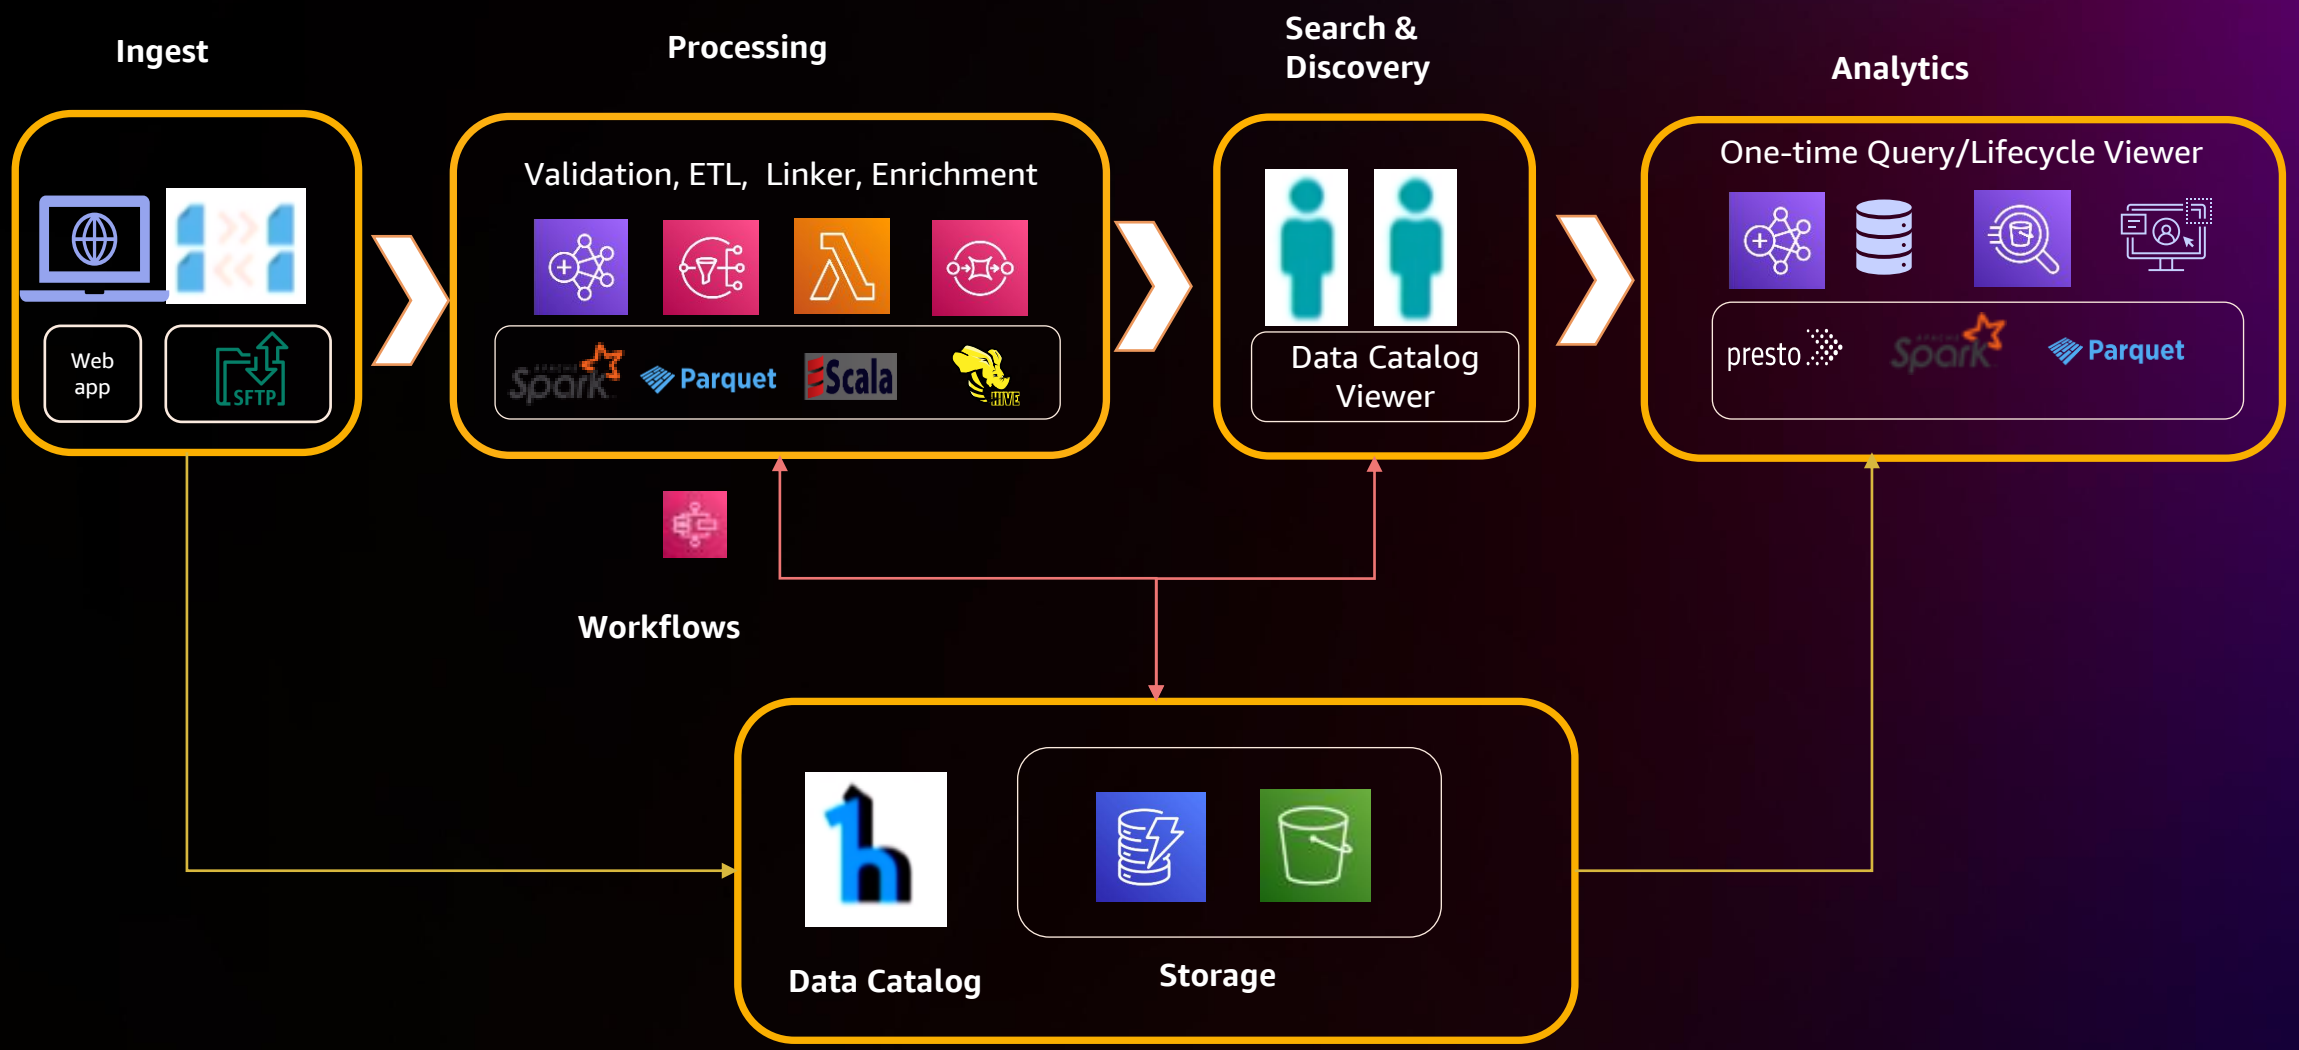
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

CAT lifecycle – A simple example



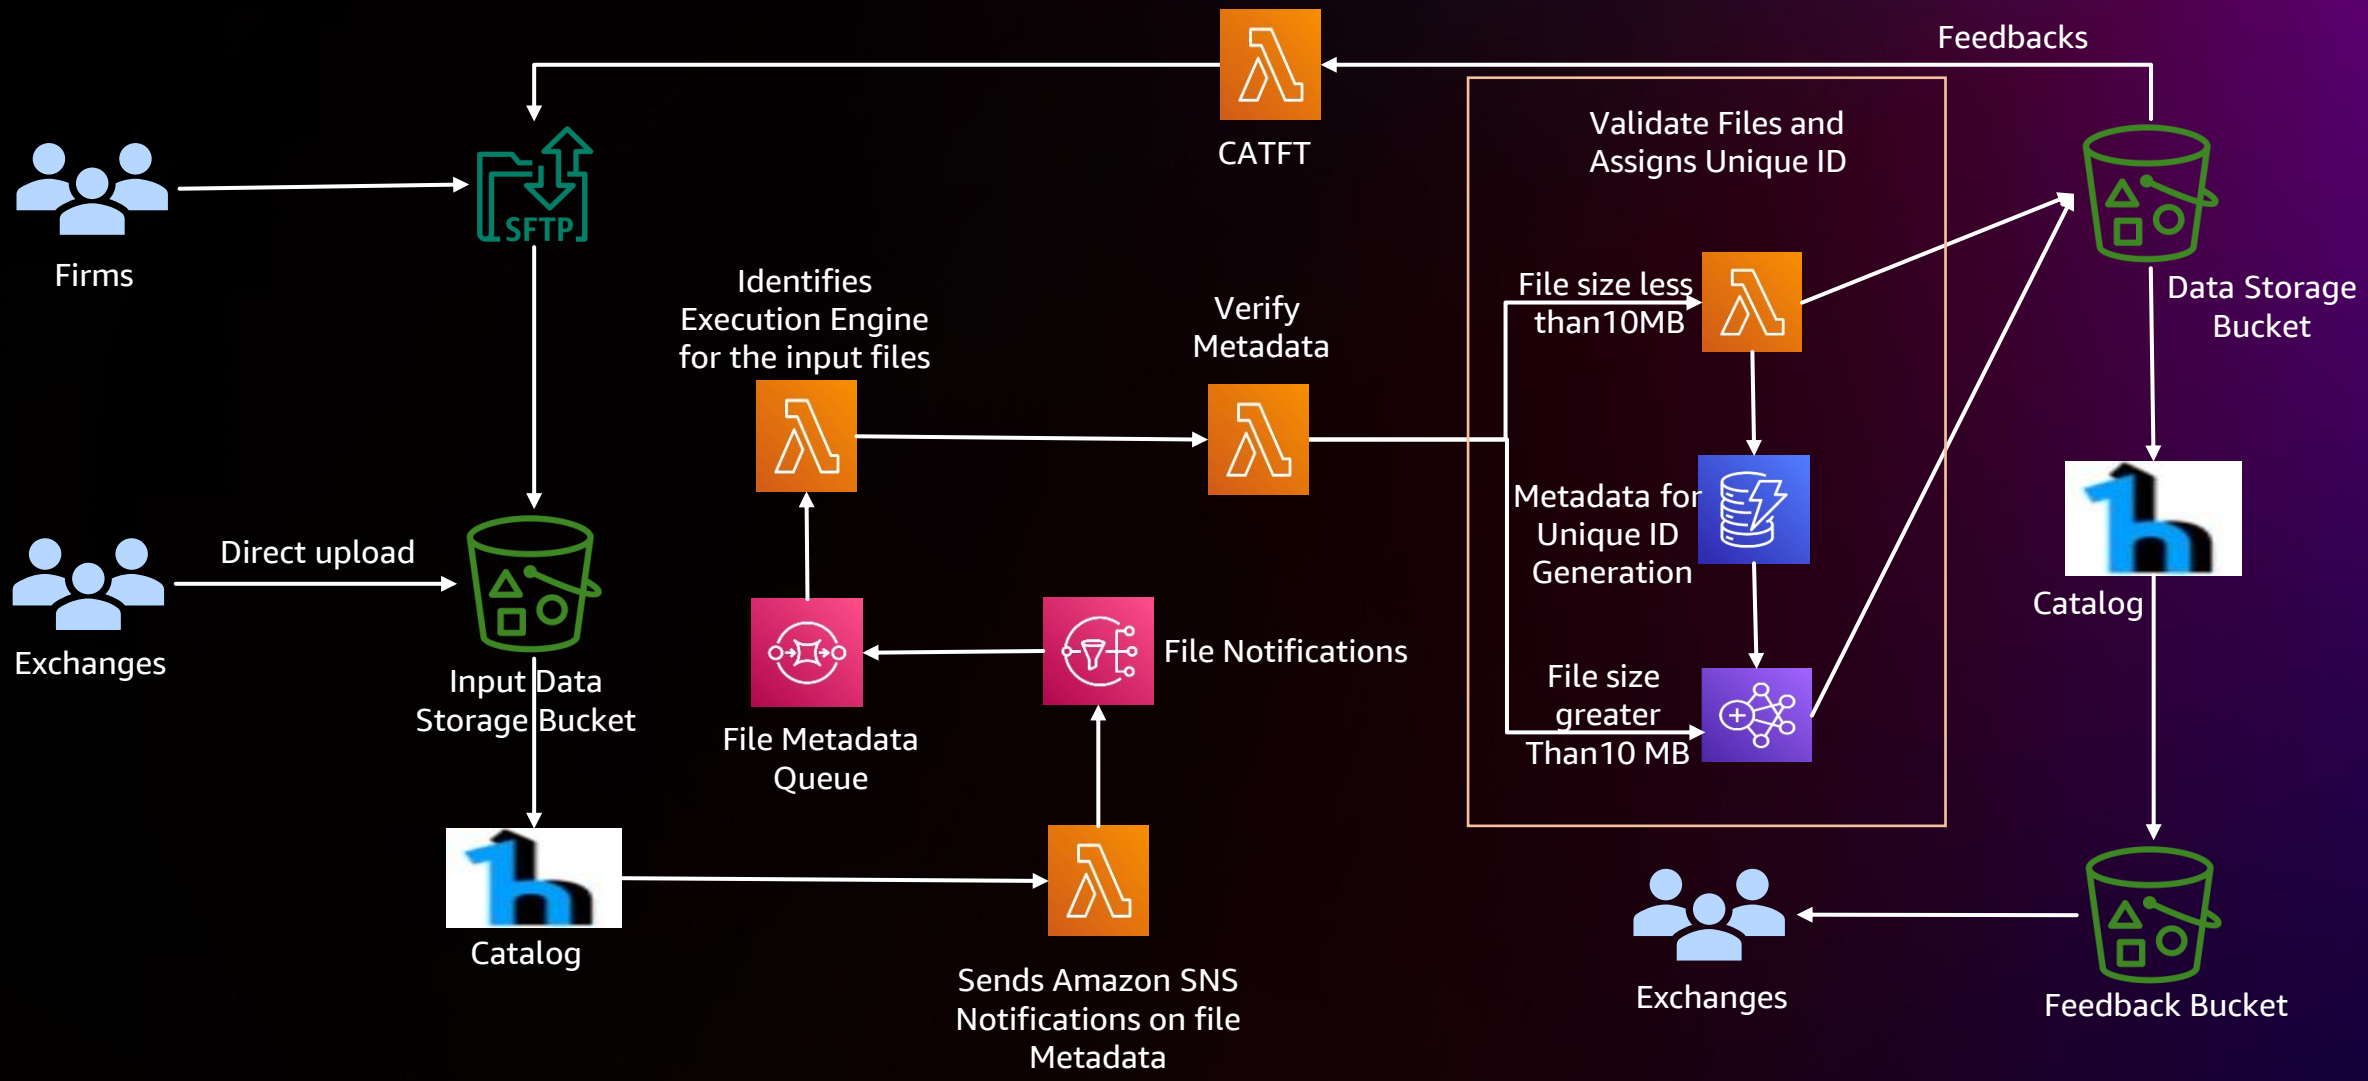
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

CAT technology stack



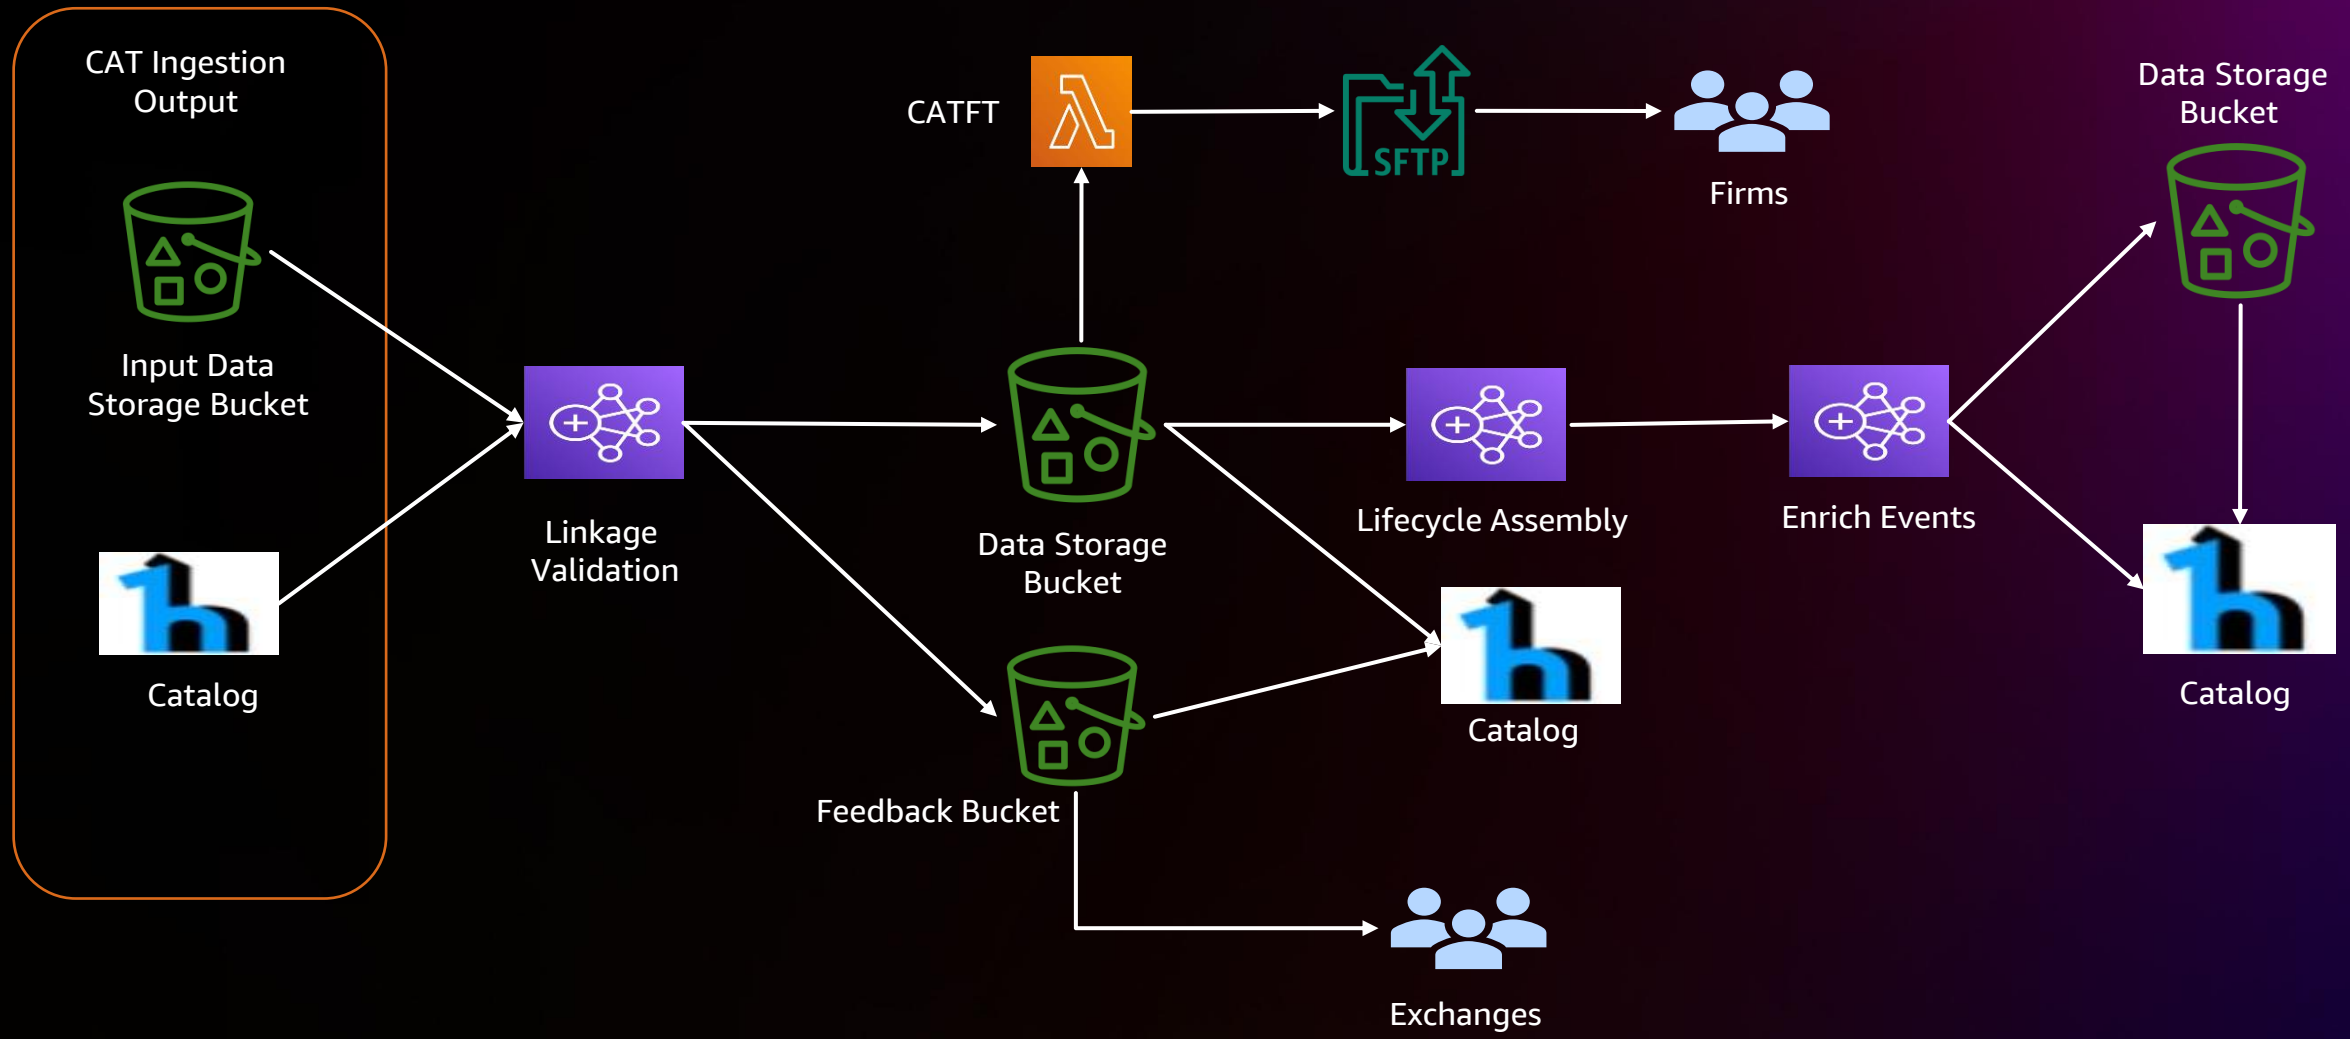
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

CAT ingestion architecture



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

CAT processing architecture



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

CAT challenges in meeting SLA

Performance



CAT has an aggressive SLA of 4 hours to link 1 Trillion market events and report the errors to firms and exchanges

Resiliency



Resiliency of CAT processing is important as there is no time to waste in order to consistently achieve the 4-hour SLA

Scalability



CAT application must be scalable to process fluctuating market volumes within the defined SLA. For example, 130K files where the size of any file can vary from 1 KB to few GB's depending on market activity

Cost



Compute and storage cost increases when the market volume goes up. We must continuously optimize the code and infrastructure to keep the cost same or less for increasing market volumes.

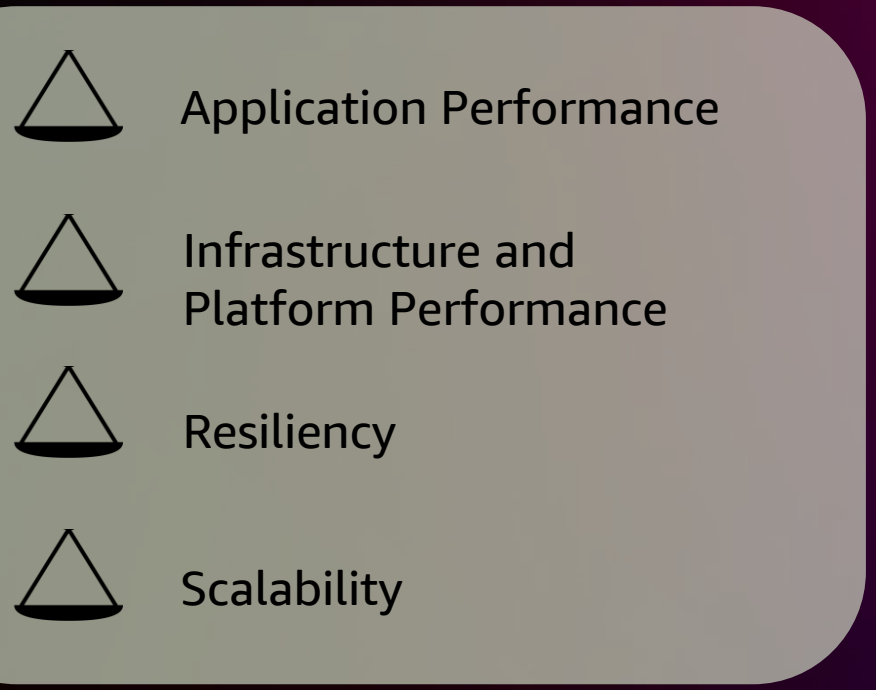
aws

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

Achieving SLA is a balancing act



Balancing critical factors



- Adding more capacity is not going to solve performance or scalability problem
- Application architecture and application code needs to be optimized to effectively utilizes the scaled infrastructure

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

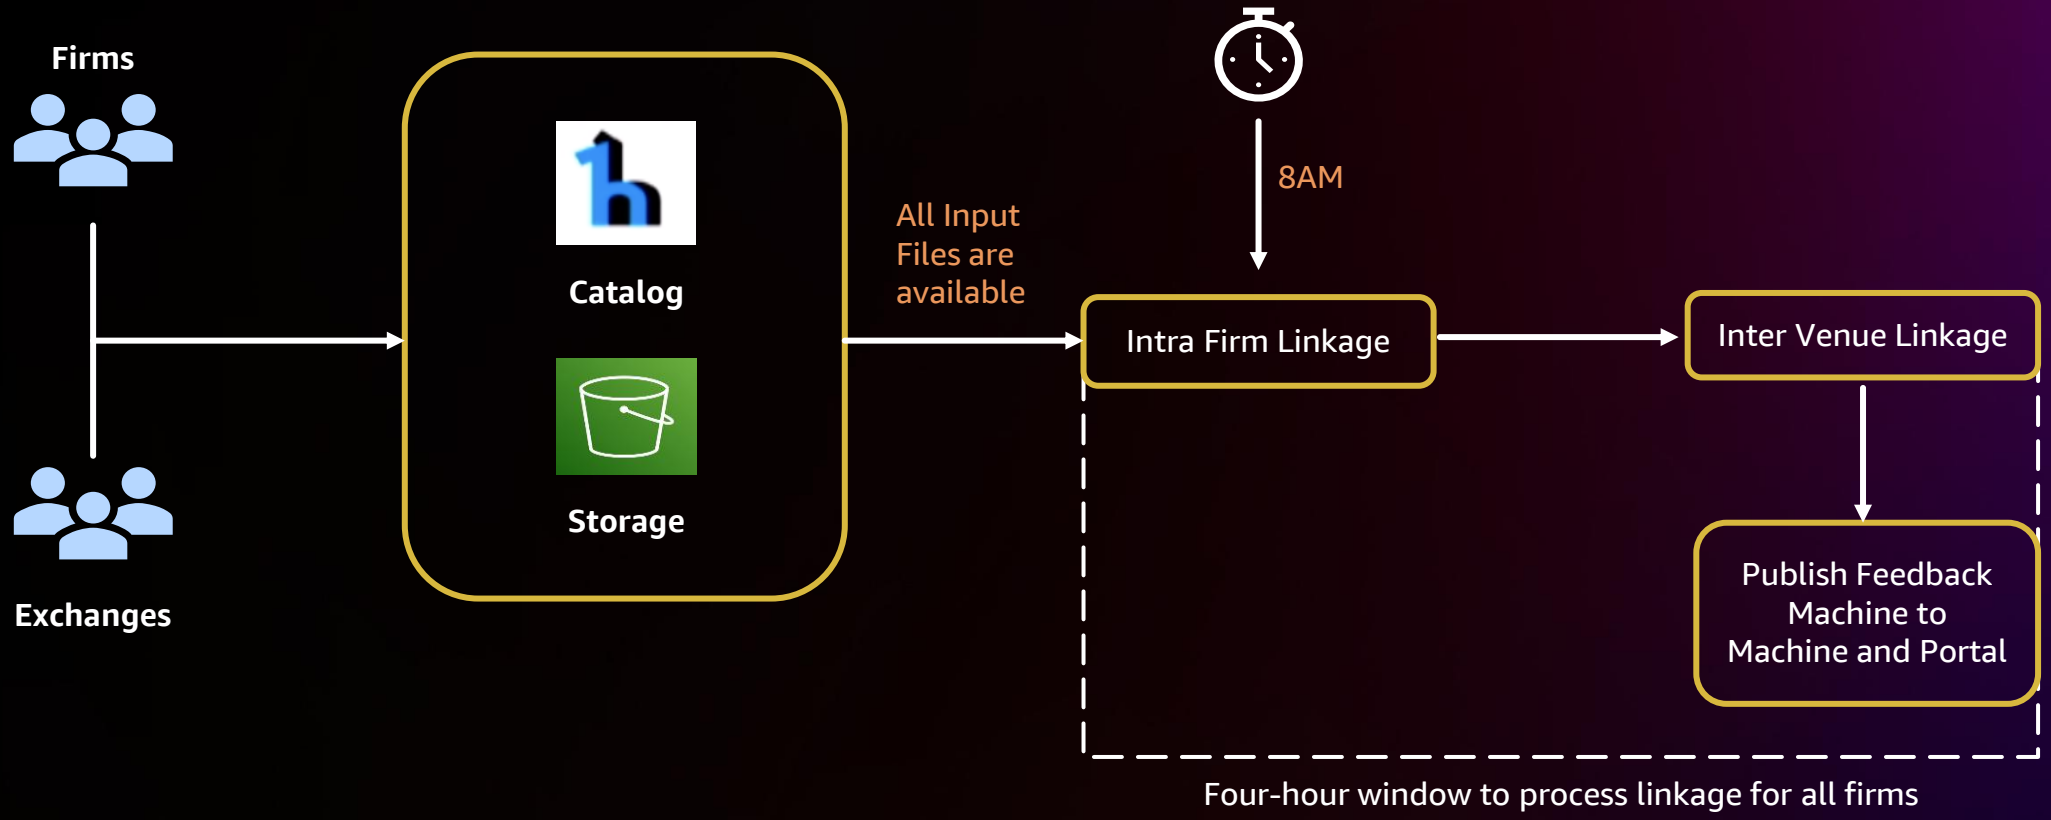
Application optimizations – Lessons learned

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

Non-progressive architecture

Performance

- Linkage for all firms starts at 8 AM after all input files are available.
- Four-hour window to complete processing irrespective of market volume.

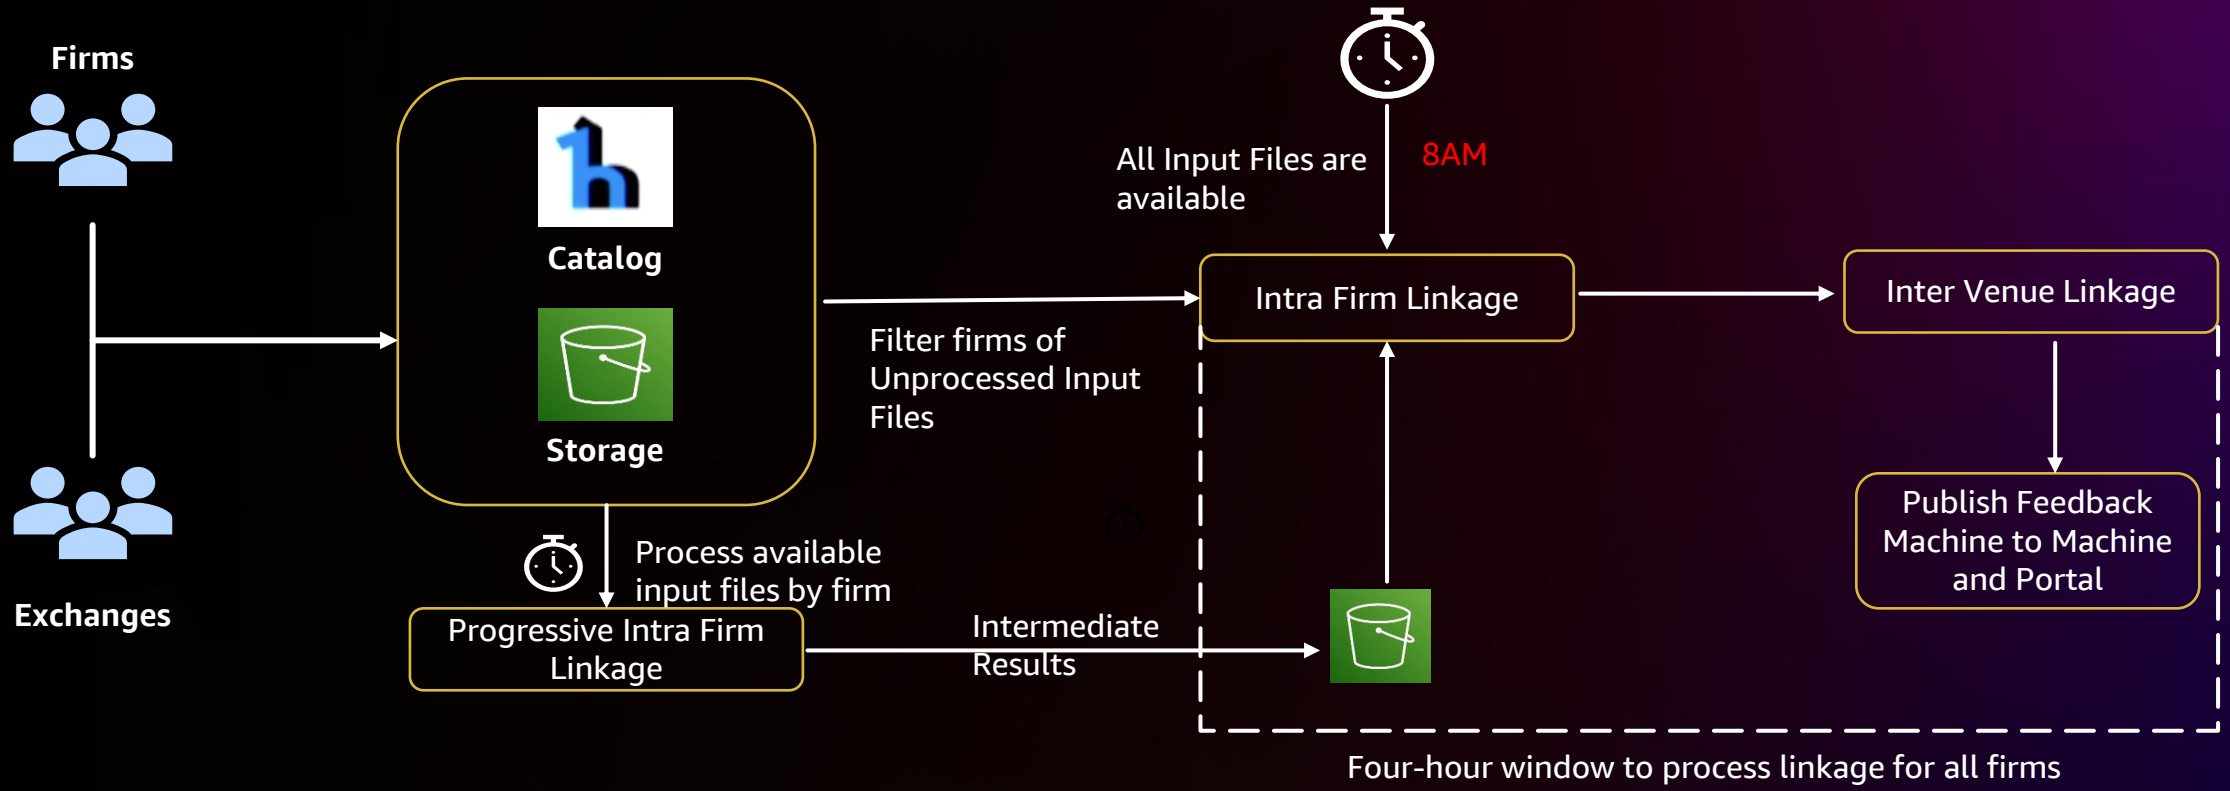


© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Progressive architecture

Performance

- Progressive – Linkage for firms starts with what we have and consolidate at 8 AM, rather than waiting till 8 AM for all the inputs
- At 8 AM only firms of unprocessed files from the last execution are linked again, and as a result volumes reduced by 60–80%



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

Performance impact of Apache Spark version



Performance

WHERE WERE WE?

- Workloads were running on clusters of size ranging from few nodes to 1800 nodes with Amazon EMR 5.27-5.29/Apache Spark 2.4
- C5.18xlarge and C5.24xlarge were the most used instance types
- Performance impacted by job resiliency as a result of open-source bugs and large clusters

aws

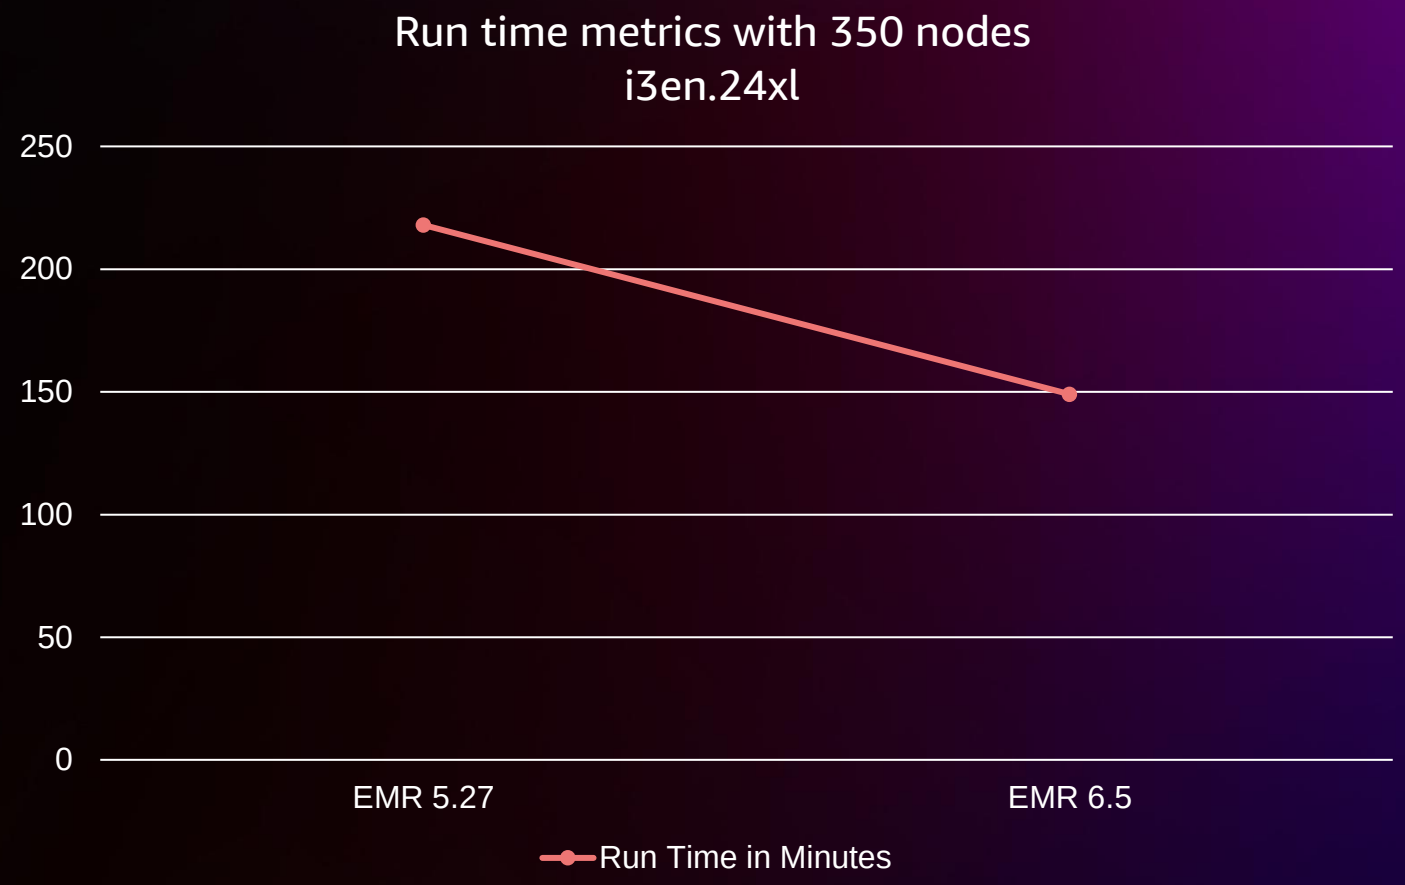
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Amazon EMR 6.5/Apache Spark 3.1.2 – Upgrade impact

Performance

WHERE WE LANDED?

- Upgraded to EMR 6.5, Spark 3.1.2, Graviton2 instances
- 30% performance improvement observed with improved resiliency and security compliance (IMDSv2)
- New capabilities like AWS Graviton support, multiple parallel Amazon EMR steps and better shuffle performance



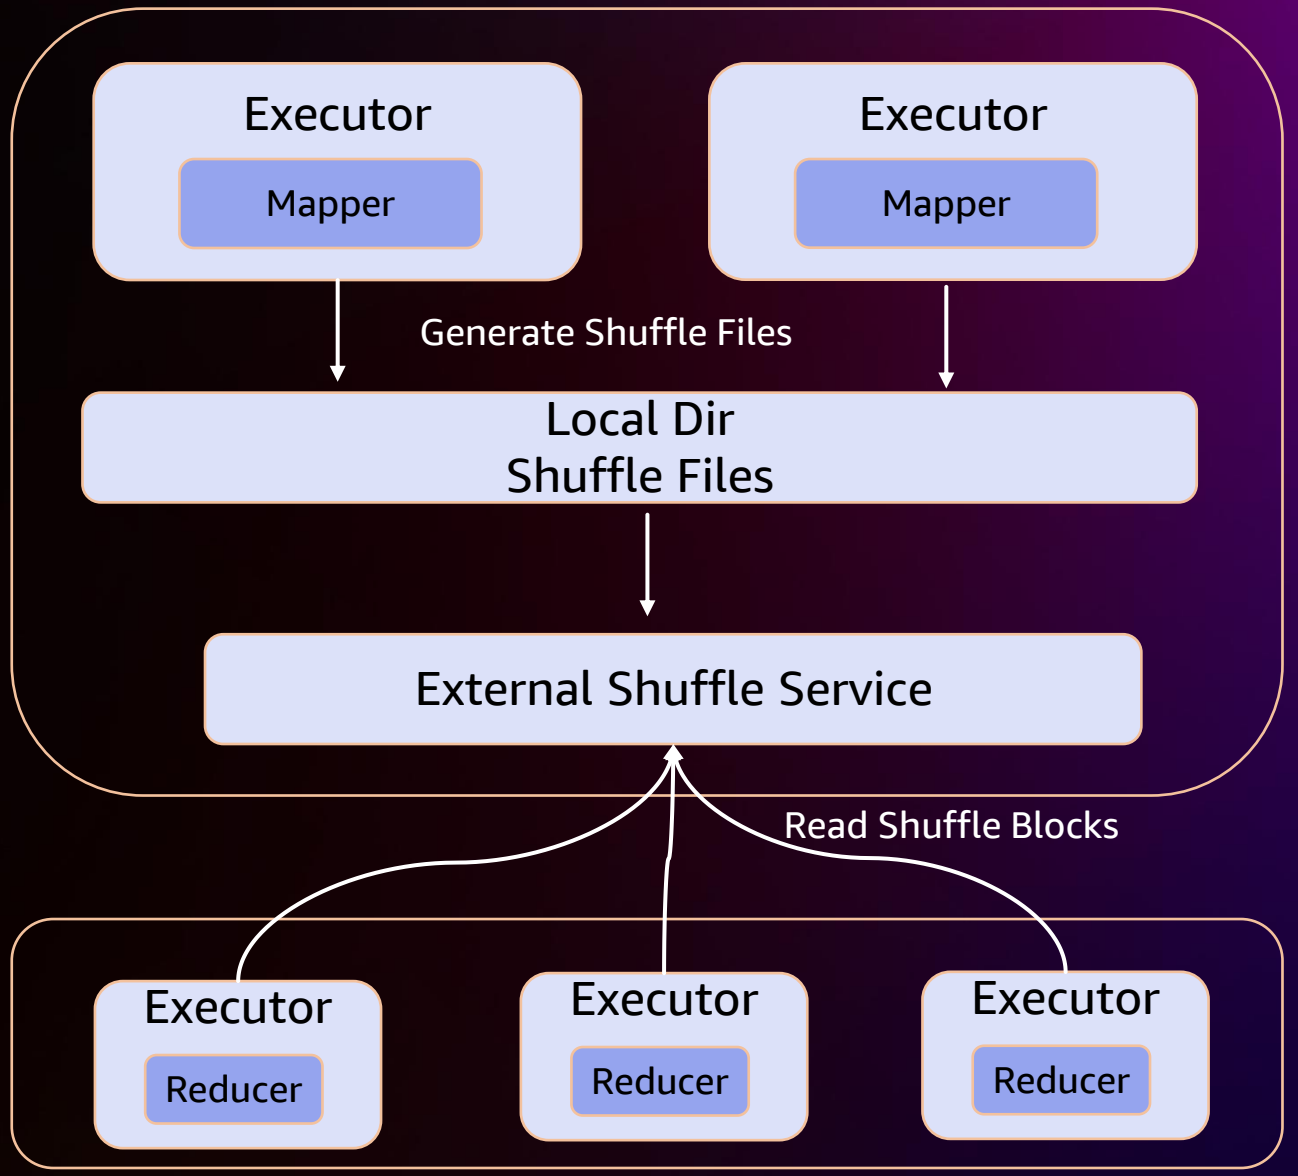
aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

CAT application challenge

Performance

APACHE SPARK SHUFFLE LIMITATION

- Linkage Validation requires grouping by multiple keys
- Any group by operation will trigger a data shuffle greater than 100 TB
- Shuffle is very heavy IO as it uses local disk for reading and writing shuffle files
- With IO intensive shuffles Amazon EBS GP2 IOPS was saturated with large clusters



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

NVME for Apache Spark Shuffle



Performance

Huge shuffle require higher iops to improve resiliency

Non-Volatile Memory Express (NVMe) SSD-backed instance offers

- Very high random I/O performance
- High sequential read throughput
- Provides high IOPS

Benchmarked the application with i3en, r5d and r6gd Amazon EC2 instances backed by NVMe SSD disk

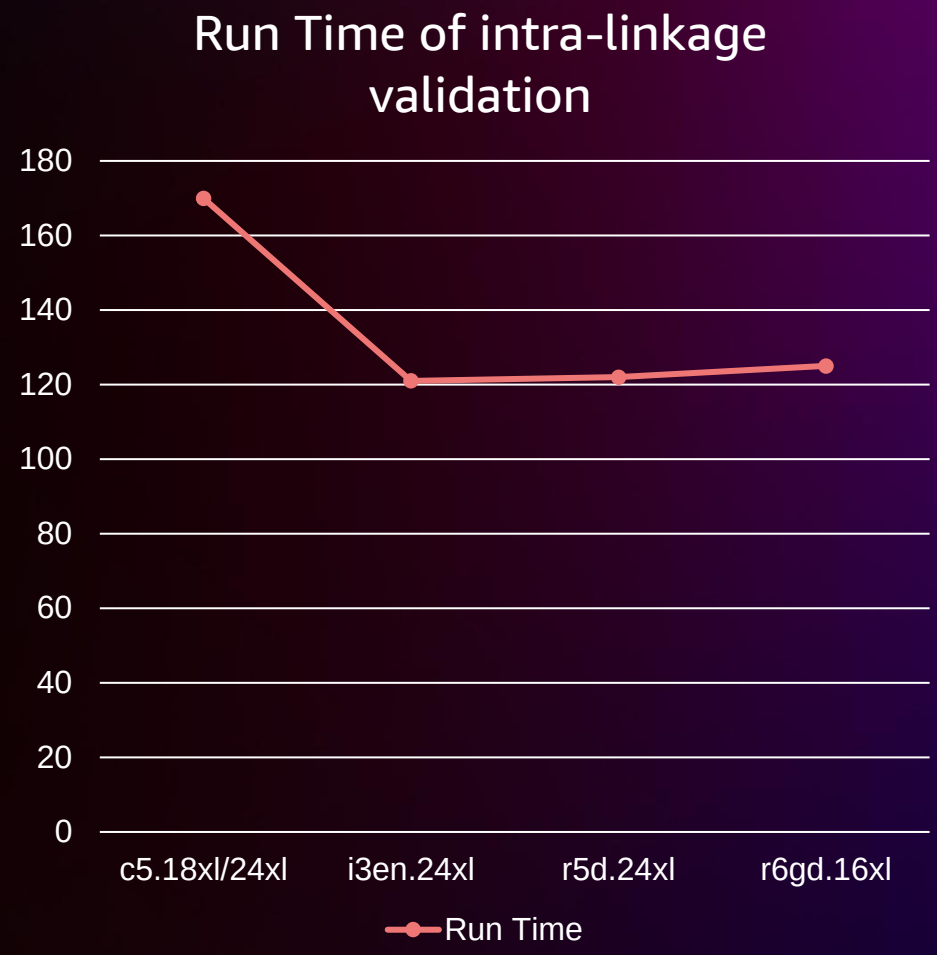
aws

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

AWS Graviton instance with NVMe for Apache Spark Shuffle

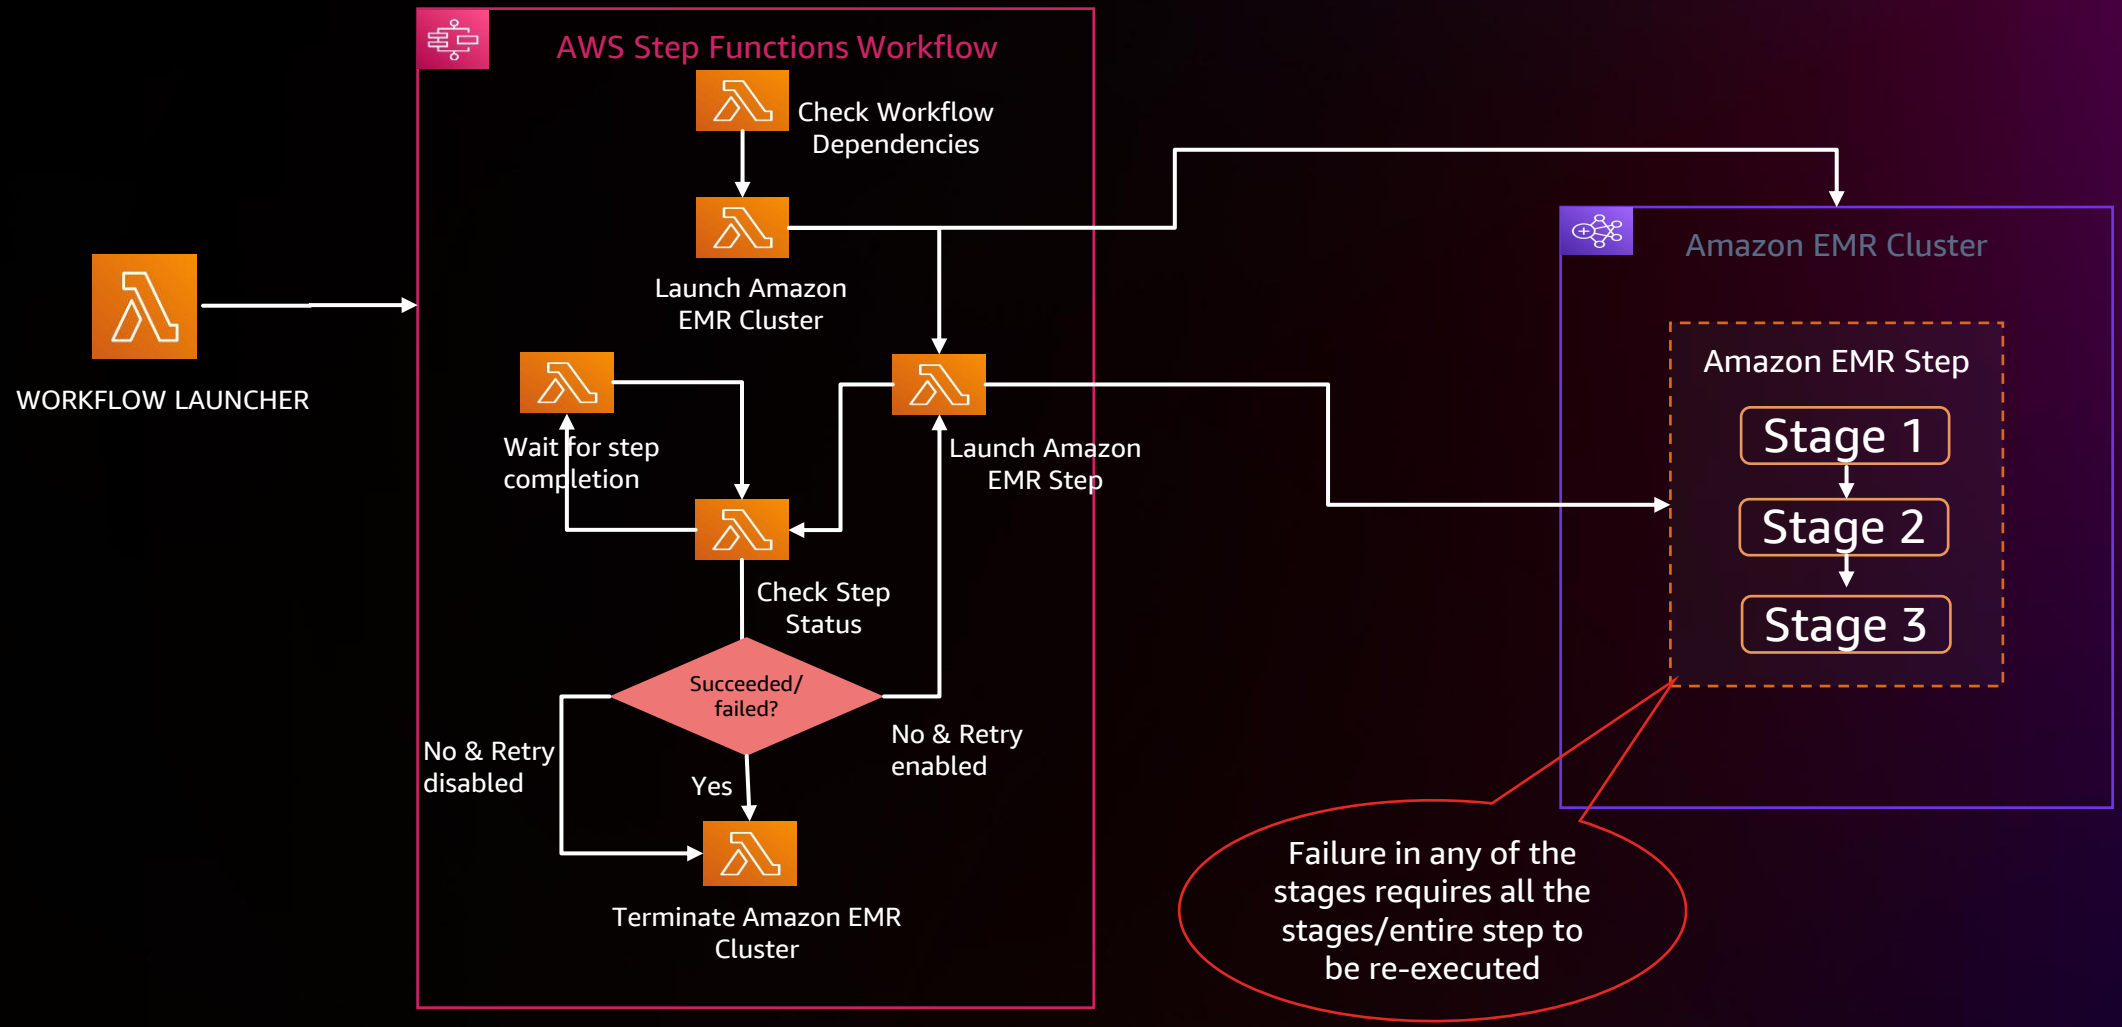
Performance

- Amazon EMR upgrade enabled us to use AWS Graviton based instances
- AWS Graviton processors are designed by AWS to deliver the best price performance for big data workloads running in Amazon EC2
- AWS Graviton instance R6GD.16xl with NVMe storage was chosen for benchmark
- 30% improvement in performance was observed with AWS Graviton



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

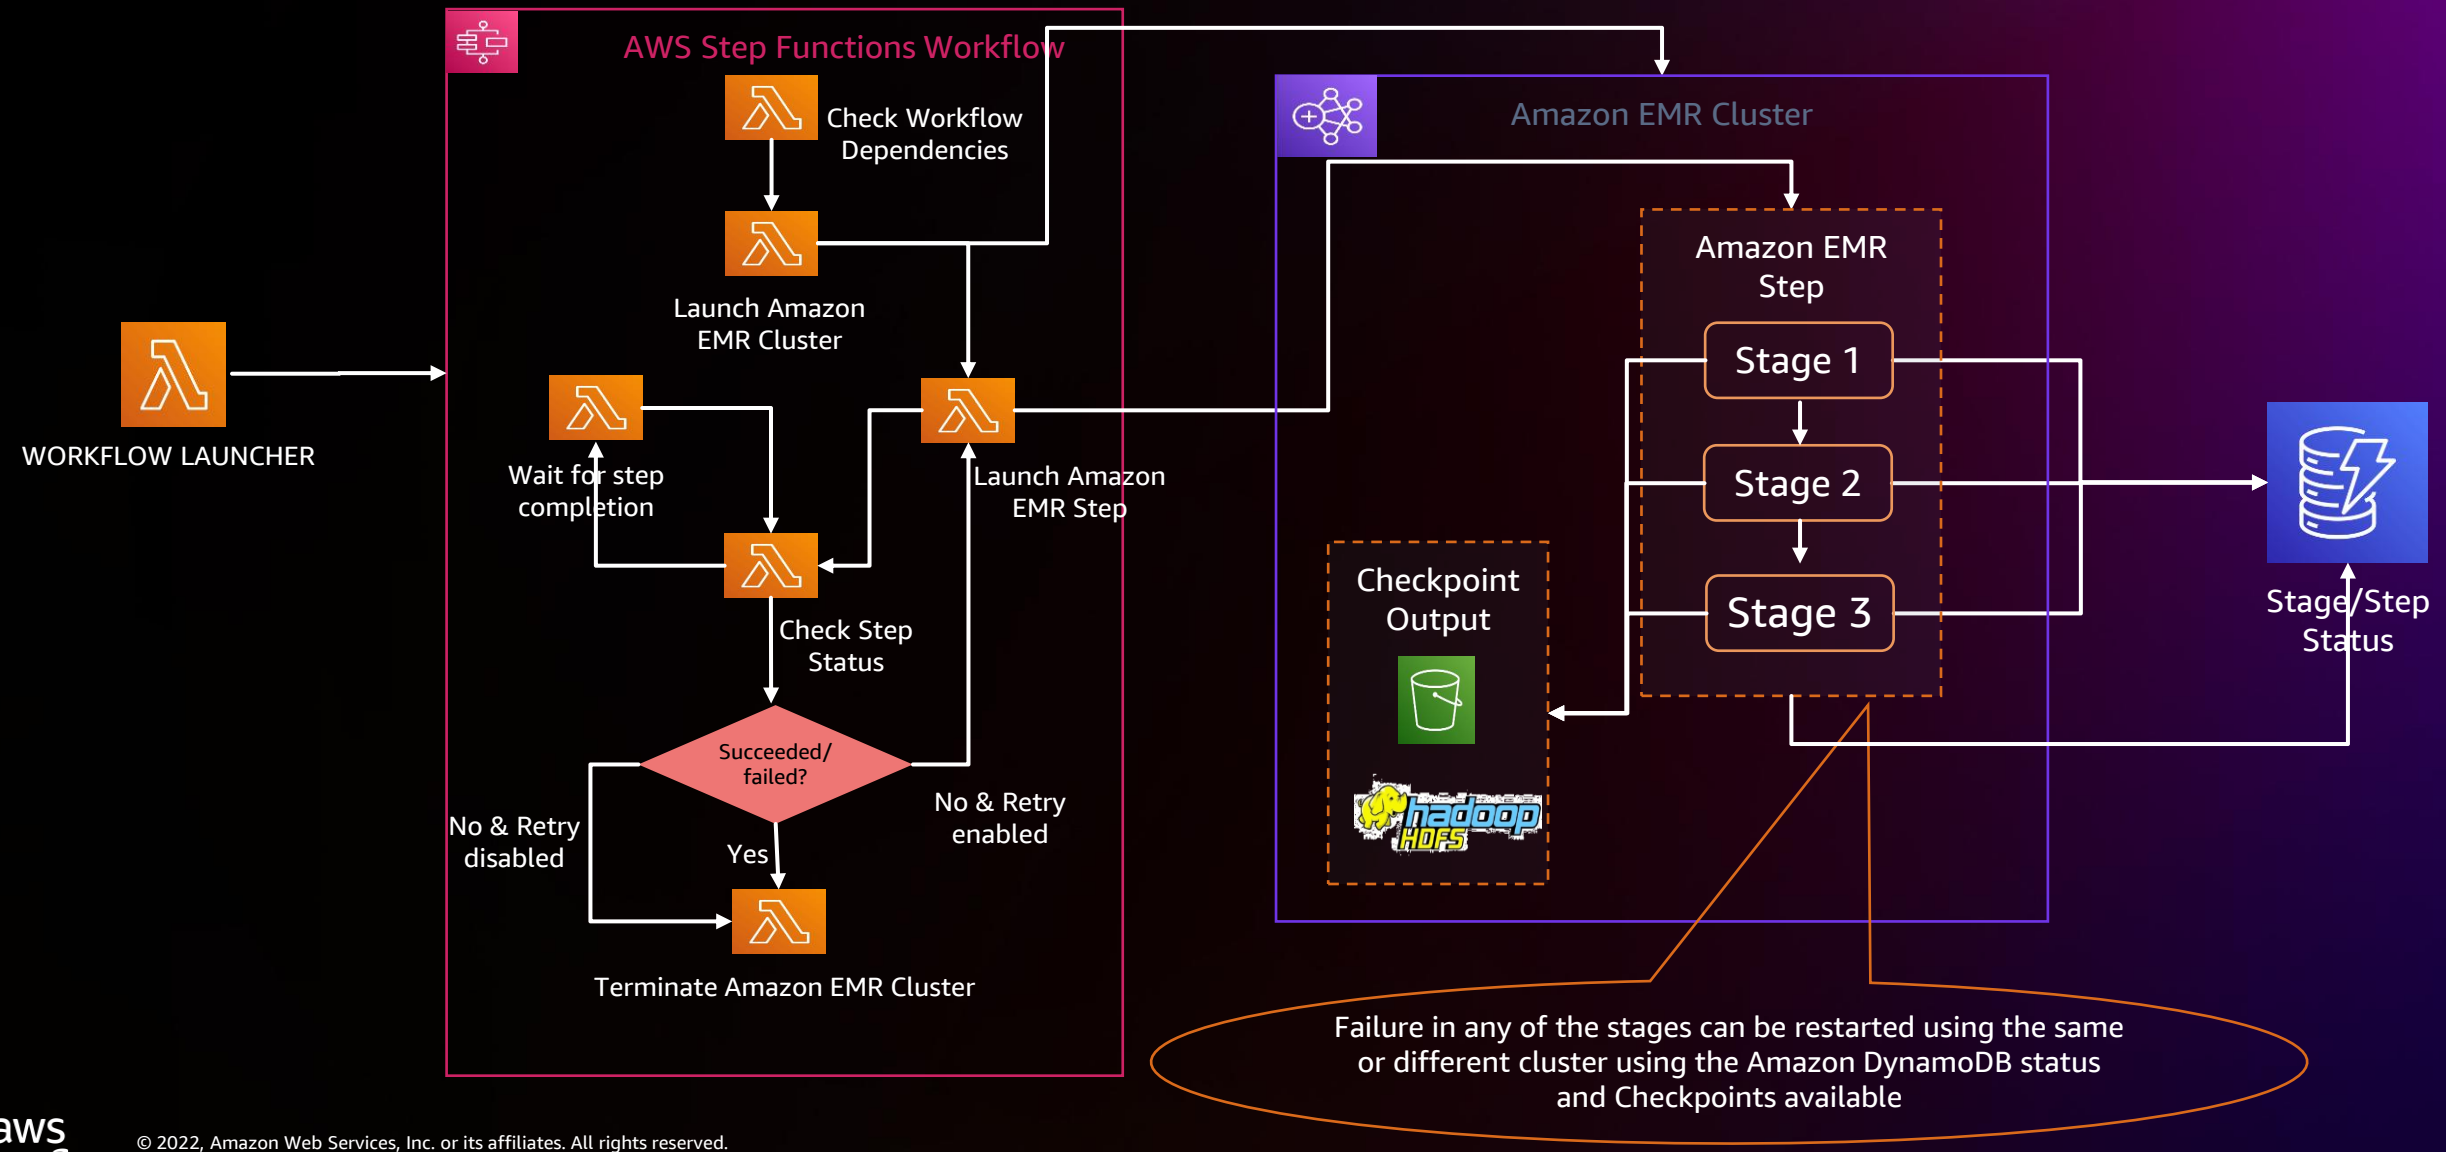
Resiliency – No restarts within Amazon EMR Steps



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

Resiliency by checkpointing and restart

Resiliency



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

Resiliency by checkpointing and restart



Checkpoints

- Checkpointing at regular intervals helps to quickly recover from failures
- We checkpoint the output after every logical step
- Each step is not more than 30 mins

Restarts

- Restart from the stage where the job failed by persisting the stage/workflow status on Amazon DynamoDB
- By checkpointing to HDFS we can restart on the same cluster
- By checkpointing to Amazon S3 we can restart on a different cluster
- Enable auto restart for known hardware/Apache Spark failures

aws

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

AWS service and platform optimizations – Lessons learned

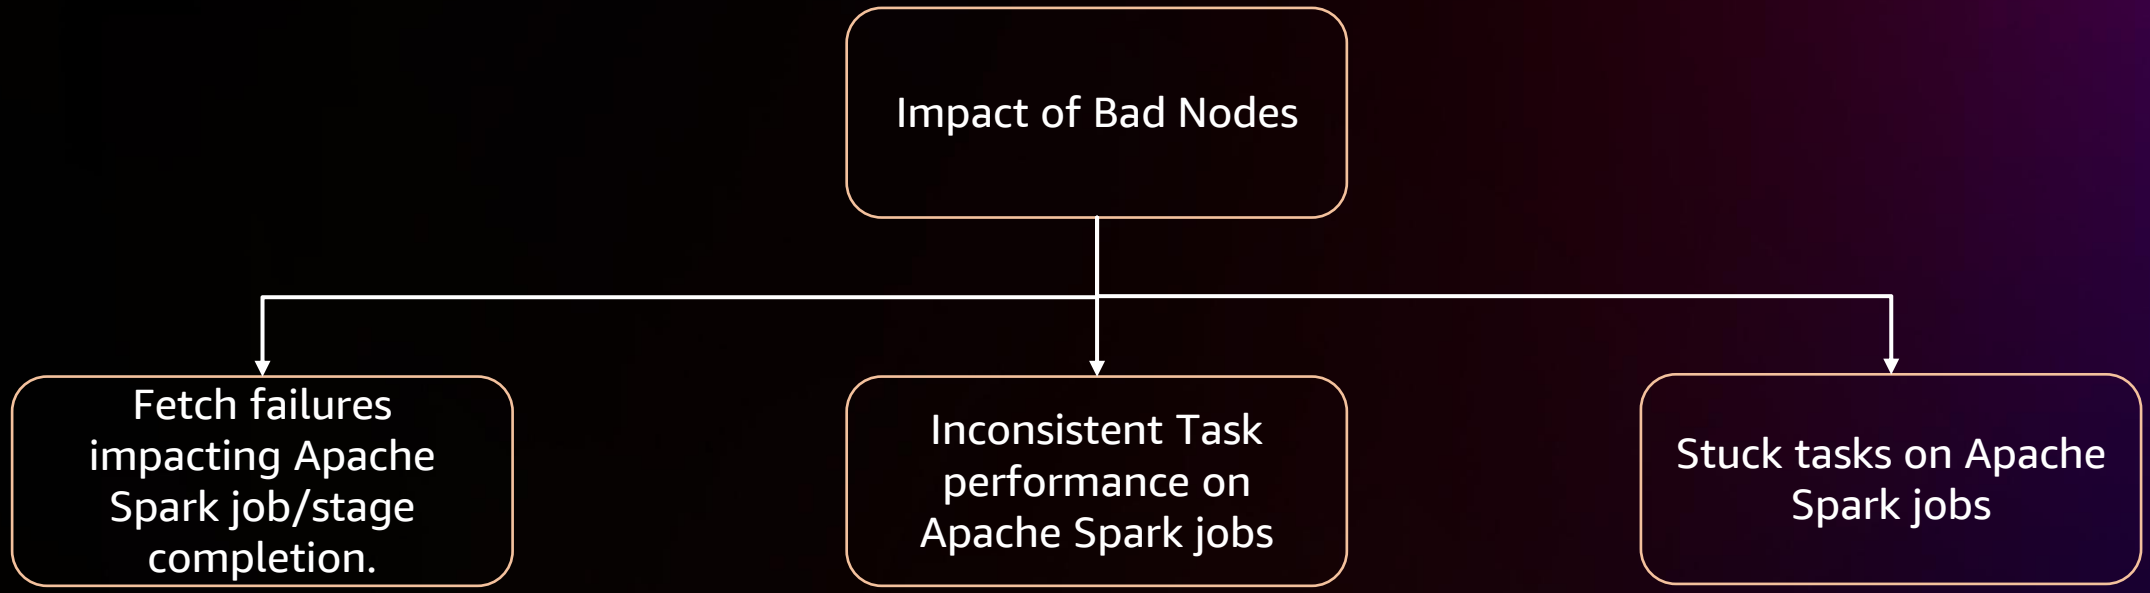
Meenakshi S. Ponn Shankaran
Senior Big Data Architect, Professional Services
AWS

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

SLA Impact of bad nodes in the Cluster

Resiliency

- ~300k compute nodes daily – At some point we will land on every server in an Availability Zone
- Common to land on hardware with known/unknown issues [Memory Errors/CPU]
- Even with Amazon EMR and Apache Spark has built in fault tolerance, processing was impacted by 15–20 mins due to task retries, slow running or straggler task



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

Workarounds to remediate bad or slow nodes

Resiliency

Graviton Migration

Workload migration to AWS Graviton reduced the hardware errors significantly.

Reduced error rates resulted in improved stability

Upgrades

Amazon EMR upgrade to reduce the impact of open-source bugs

Amazon Linux version upgrade to reduce kernel bugs

Cluster Configuration Optimization

Enable and Optimize Apache Spark Speculation feature

Apache Spark Configuration optimization [exclude hosts]

Enhanced Monitoring/ Remediation plan

Custom built automation using AWS SSM agent to remove instances impacting performance

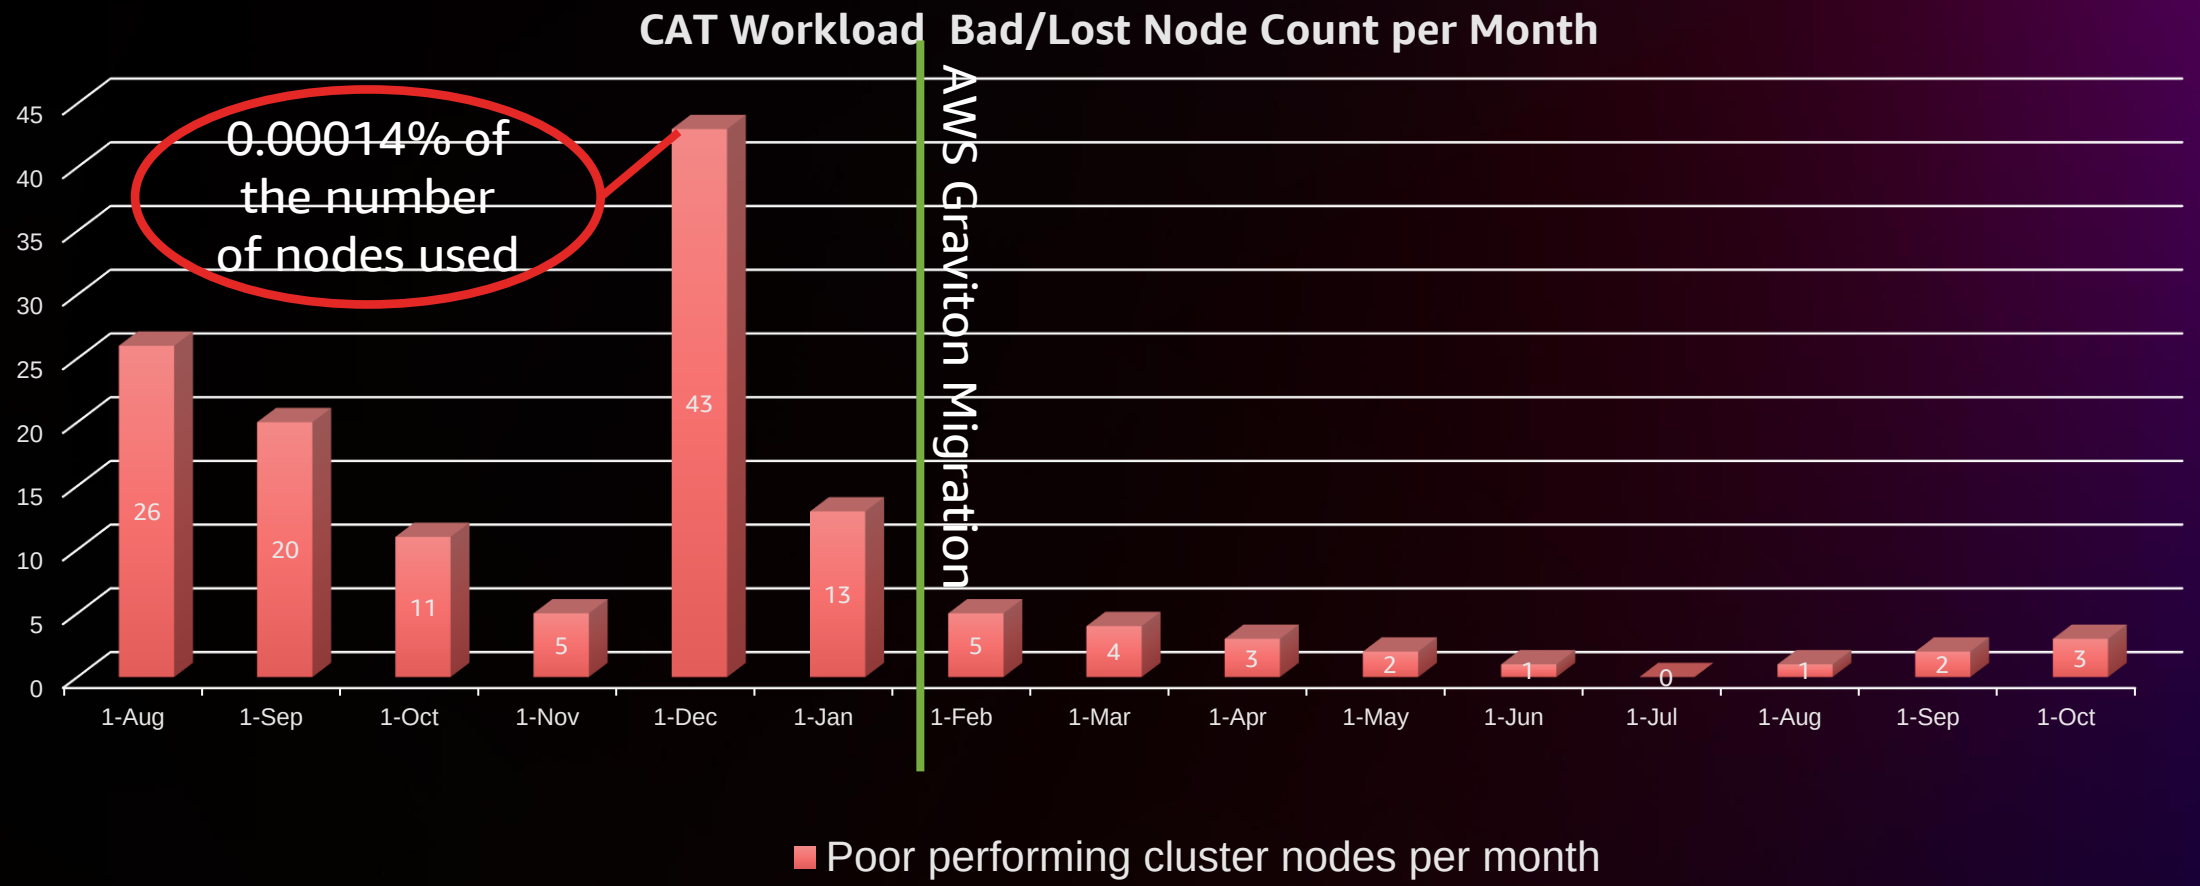
Enhanced service level monitoring

aws

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

Bad/slow nodes trends after workarounds

Resiliency



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

Handling burst of reads/writes in an Amazon S3 partition

Scalability



- Amazon S3 – 3,500 put/5,500 get requests per second per prefix
- FINRA's Amazon S3 prefixes are partitioned by trade date and will look like this

  s3:<bucket>/cat/processing_date=mmddyyy/brokerevents/brokerid=abcd/
- Folders inside the trade date prefix will need high throughput due to traffic burst generated by multiple workloads running on Amazon EMR
- Amazon S3 traffic heat generated at processing date level due to multiple parallel jobs reading and writing to the same date level partition
- Jobs impacted with 503 (s3 slow down) errors which causes jobs to fail

aws

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

Handling burst of reads/writes in an Amazon S3



Scalability

Amazon S3 Pre-Partitioning

Collaborate with Amazon S3 team for better performance

Pre-partition reduced 503 errors and improved the resiliency of the job

Additive Increase Multiplicative Decrease (AIMD) Retry Policy

Worked with EMR team to introduce new retry policy – Additive Increase Multiplicative Decrease (AIMD)

Enabling AIMD retry policy reduced error by 79% and eventually reduced job failures

Manage Workloads

Reduce Amazon S3 read traffic by eliminating Amazon S3 reads from unneeded prefixes

Use Apache Parquet with optimized file size (between 128 MB to 1 GB)

Design application to avoid concurrent burst of reads/writes in a single Amazon S3 partition

aws

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

Targeted/suspendable ODCR to guarantee capacity



Scalability

CAT Needs Capacity Guarantee

- CAT needs to complete processing the markets events in the four-hour window to provide feedback to the firms and exchanges
- To meet the processing deadlines, guaranteed capacity is very important.
- No insufficient capacity exceptions (ICE) when provisioning

On-Demand Capacity Reservations

- CAT mission critical application relies on OnDemand Capacity reservations to reduce the impact of ICE
- Targeted on demand reservations are used to target the reservations against specific workflows by specifying the reservations on the Amazon EMR cluster definition
- Reservations are suspended over the weekends to reduce cost and resumed when the processing starts

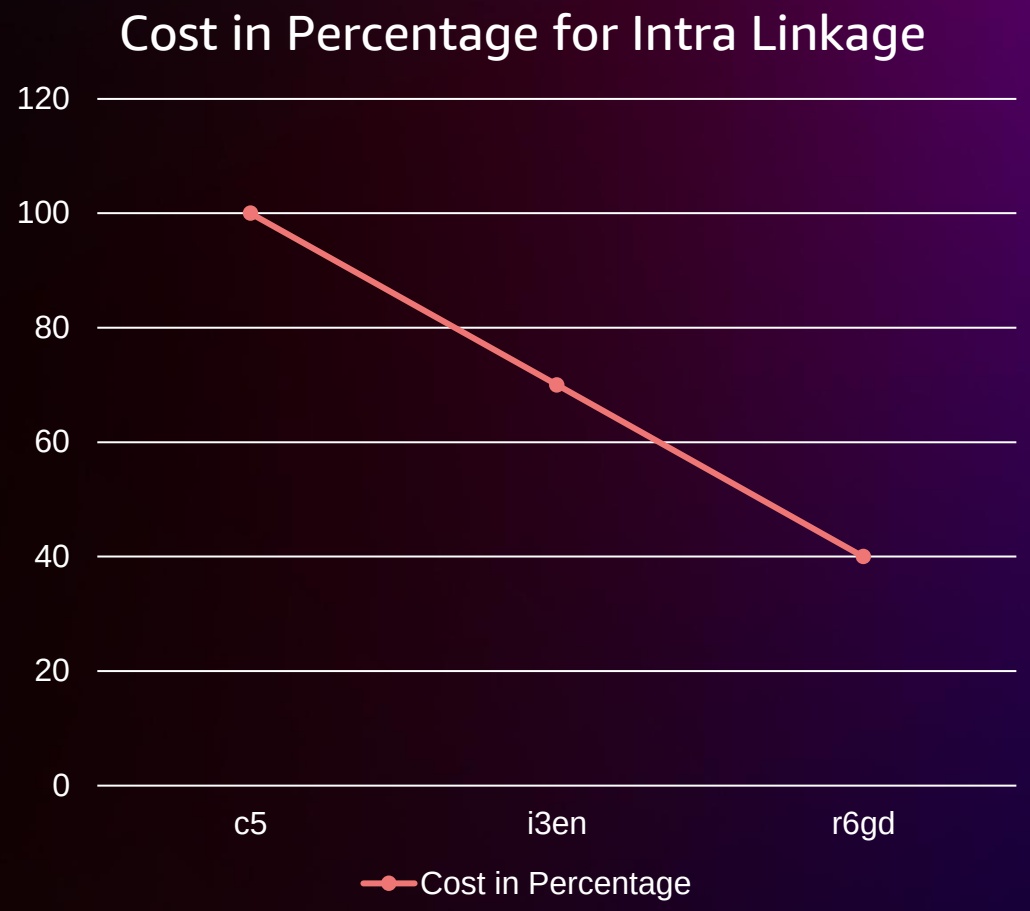
aws

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

AWS Graviton adaption for cost optimization

Cost

- With the use of the Graviton instance type, the compute instance hours have reduced significantly
- 60% better price performance was observed on Graviton instance R6GD.16xl with NVMe storage



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

Instance type diversification for cost/capacity optimization

Cost

Spot/On-Demand Capacity impact

- CAT Amazon EMR cluster fleet definitions had 1-5 Amazon EC2 instances types
- Spot capacity was not available at times forcing us to switch to on-demand
- Sometimes pool depth of the requested instance type is not sufficient, impacting cluster provisioning using on-demand
- Cost increased for workloads that can withstand spot interruptions and are not mission critical

Instance Type Diversification & Allocation Strategy

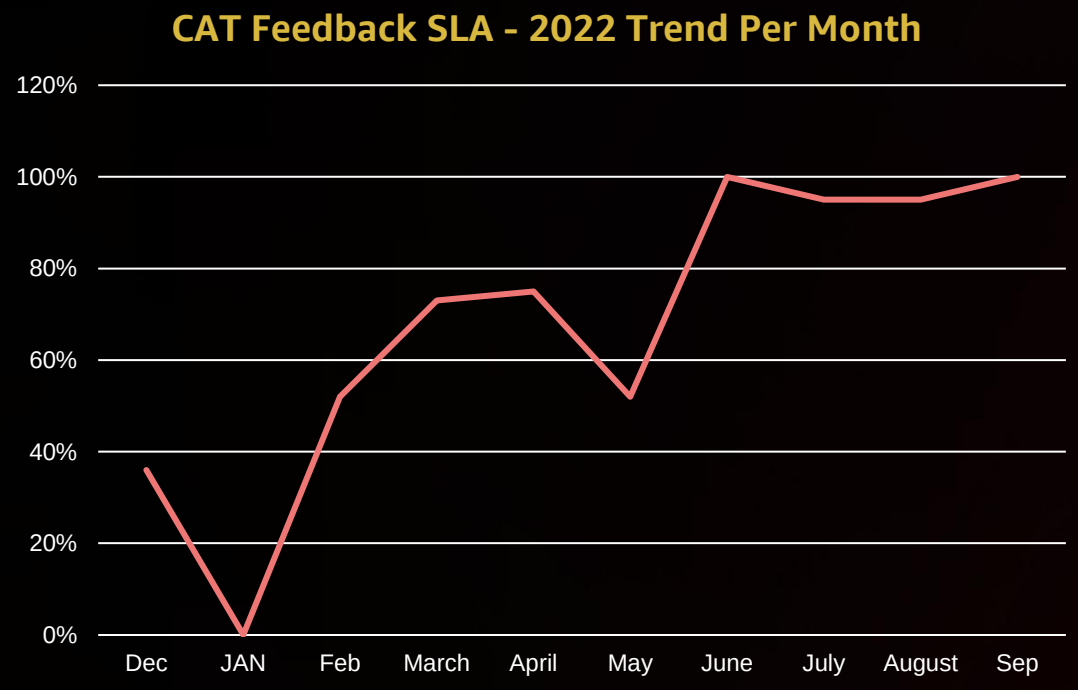
- Amazon EMR currently supports up to 30 instance types per instance fleet. CAT cluster definitions are taking advantage of this feature to diversify instance types where we defined 5 to 15 instance types in the fleet
- Enabled allocation strategy – lowest-price for On-Demand Instances; capacity-optimized for Spot Instance on the Amazon EMR Cluster specifications
- Cluster creation failures due to Insufficient capacity exceptions reduced and spot usage increased

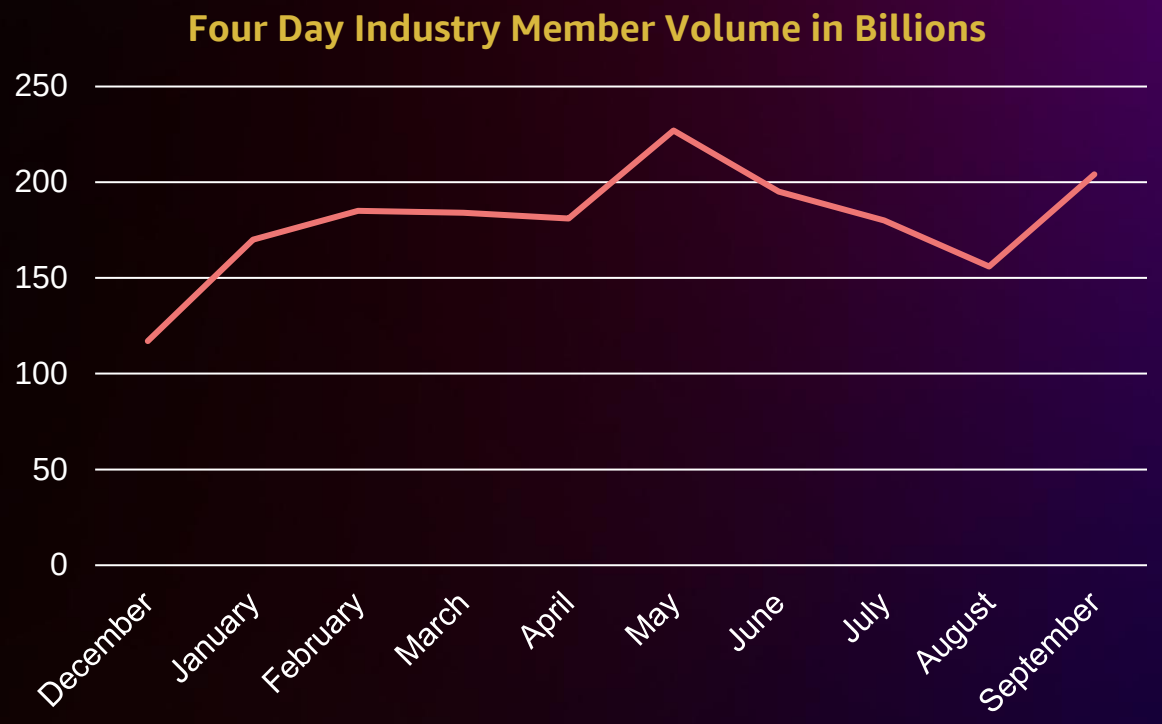
aws

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

Success metrics

- Improved success rate in achieving SLA
- SLA success rate increased from ~40% to 100% as market events volume increase by ~80–90%

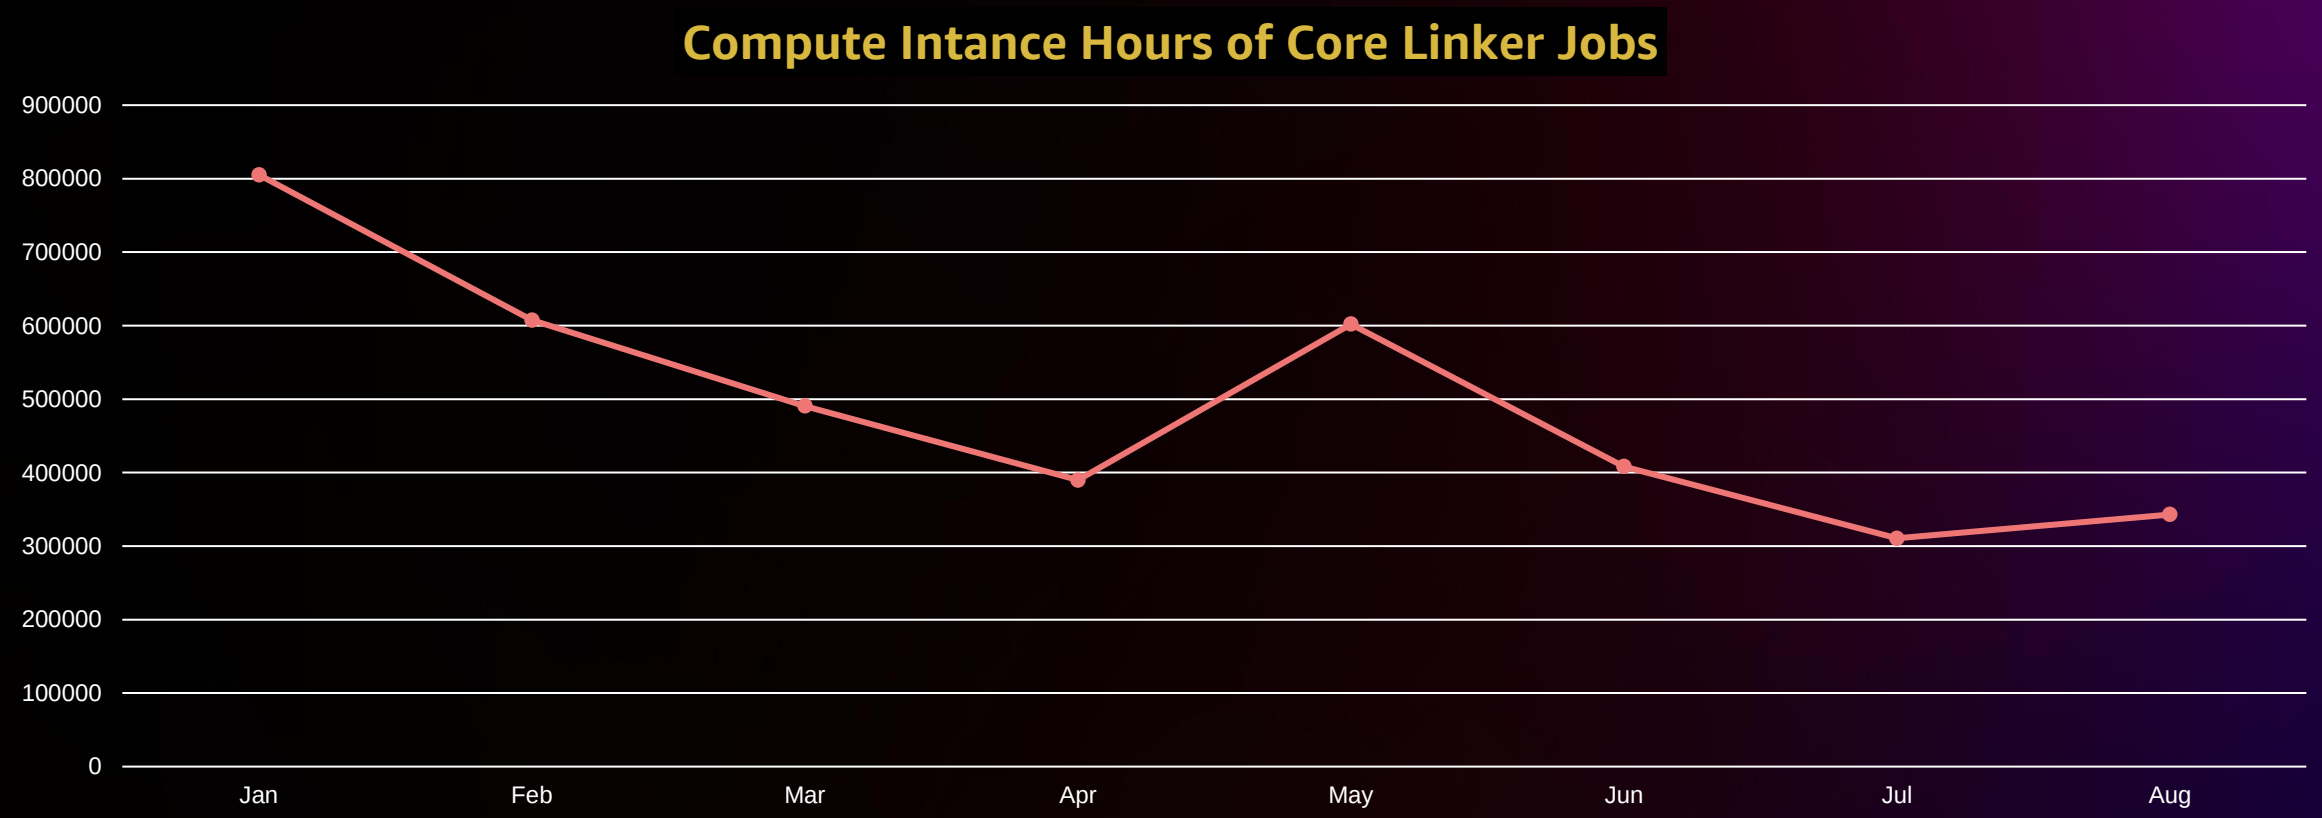




aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Success metrics

Compute instance hours for the core linker jobs have reduced by more than 50% after all the optimization and after migrating to Graviton instances



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

Future state architecture



Evaluate Amazon EMR on Amazon EKS and Amazon EMR Serverless



Adopt AWS Graviton 3 to improve price & performance



Adopt latest Amazon EMR releases for improved cost & performance



Evaluate Apache Iceberg



Adopt Amazon EBS GP3 for improved cost & performance



Application modernization – Application improvements to process by firm or group of firms in parallel

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.