aws storage

Migrating File Data to AWS: Demo & Technical Guidance

John Kennedy – Principal Product Manager, AWS
Ed Laura – Solutions Architect, AWS

© 2020, Amazon Web Services, Inc. or its Affiliates.

Today's Program

- Overview of the cloud adoption journey
- Introduction to your cloud storage options
- AWS services for migration and hybrid cloud storage
- Deep Dive – Demo
  - On-premises File Server migration using AWS DataSync and AWS Storage Gateway

© 2020, Amazon Web Services, Inc. or its Affiliates.

aws storage

You have on-premises data and applications…

…and want to use storage and services in the cloud

aws storage

© 2020, Amazon Web Services, Inc. or its Affiliates.

Cloud migration storage challenges

On-premises

I want to run my existing applications without change (databases, files, backups)…

I need local access to data in the cloud…

Compliance requires that I integrate with enterprise security and management tools…

I need reliable, easy connectivity to the cloud…

In-cloud

…but leverage the infinite scale of the cloud

…while accessing the power of cloud compute and analytics

…but I want to manage & monitor from a single pane of glass

…even when my data is in many locations

aws storage

© 2020, Amazon Web Services, Inc. or its Affiliates.

Stages of the Cloud Adoption Journey

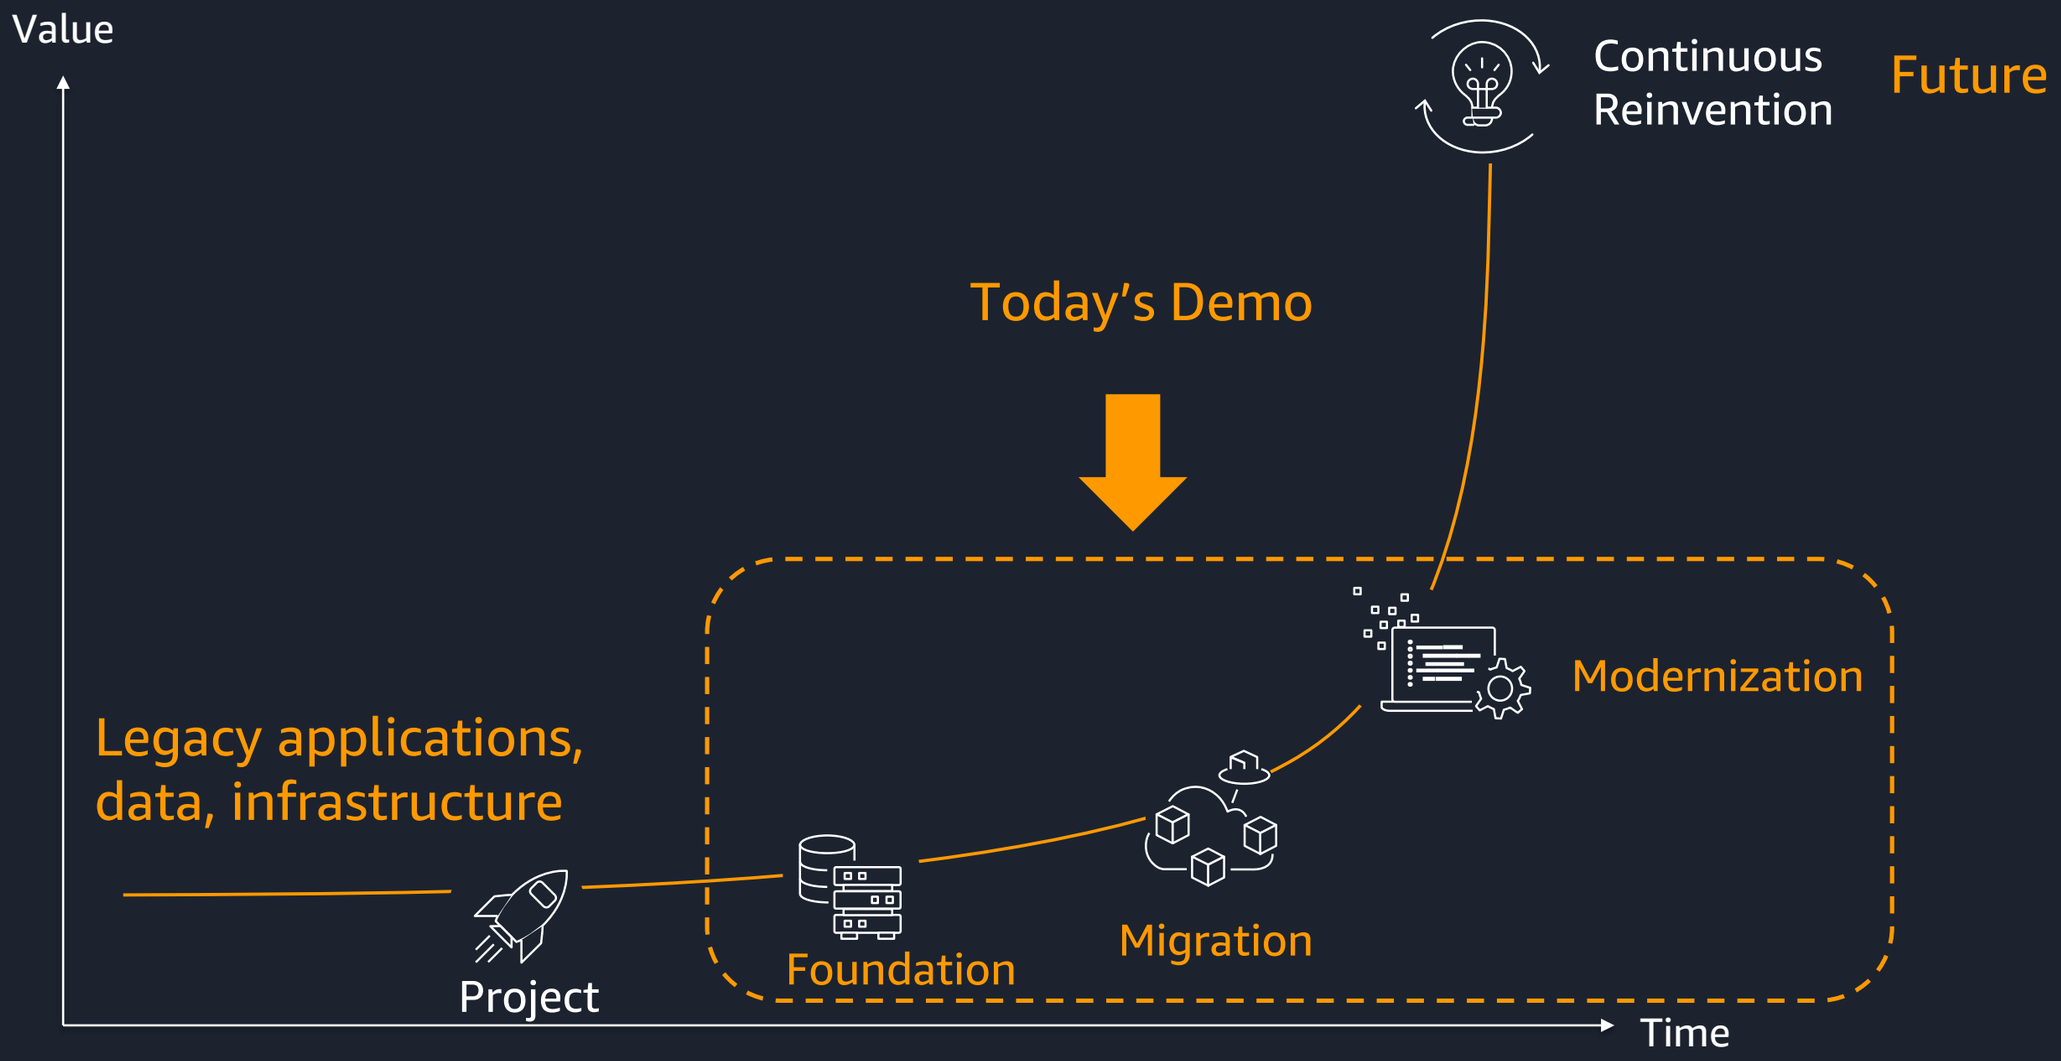

aws storage

© 2020, Amazon Web Services, Inc. or its Affiliates.

Use Cases – Migrating File Data to AWS

Storage Gateway

- Move backups to the cloud
- Shift on-premises storage to cloud-backed file shares
- Low latency access for on-premises applications to cloud data

DataSync

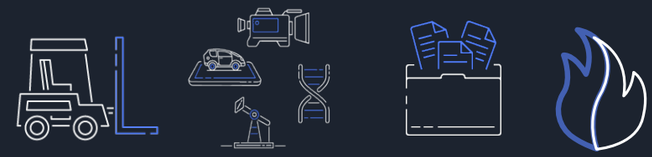

- Migration of active application data
- Transfers for timely in-cloud processing
- Archive to free up on-premises storage capacity
- Replication for data protection and recovery

© 2020, Amazon Web Services, Inc. or its Affiliates.

aws storage

Storage Gateway family

File Gateway

Store and access objects in Amazon S3 from file-based applications with local caching

File-based applications work without change

Volume Gateway

Block storage on-premises backed by cloud storage with local caching, Amazon EBS snapshots, and clones, integrated with AWS Backup

SAN-like storage with cloud recovery

Tape Gateway

Drop-in replacement for physical tape infrastructure backed by cloud storage with local caching

Easily switch tape backups to AWS

aws storage

© 2020, Amazon Web Services, Inc. or its Affiliates.

Hybrid Cloud - AWS Storage Gateway

On-premises access to virtually unlimited cloud storage

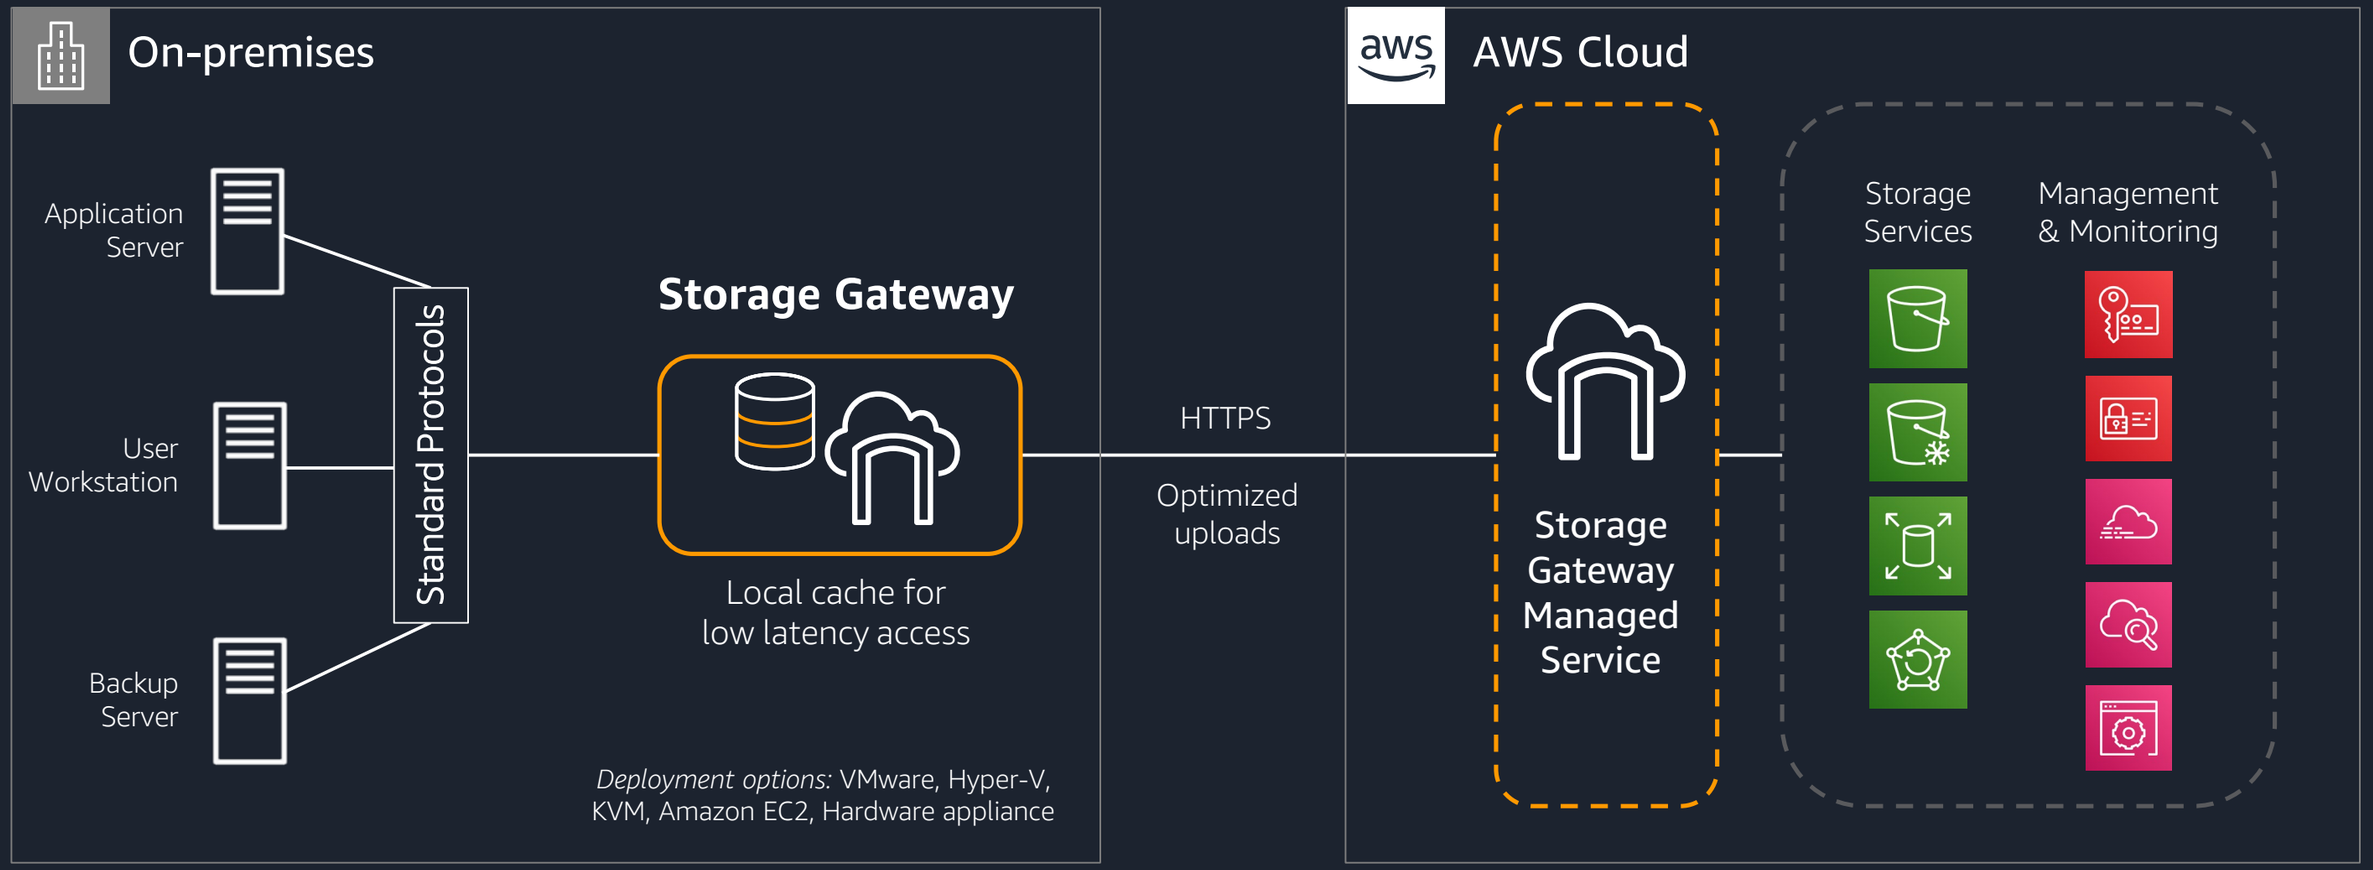

aws storage

© 2020, Amazon Web Services, Inc. or its Affiliates.

AWS DataSync

Online transfer service that simplifies, automates, and accelerates moving data between on-premises storage and AWS

Fast data transfer

Easy to use

Secure and reliable

Cloud integrated

Cost-effective

The speed and reliability of network acceleration software with the cost-effectiveness of open source tools

aws storage

© 2020, Amazon Web Services, Inc. or its Affiliates.

AWS DataSync: How it works

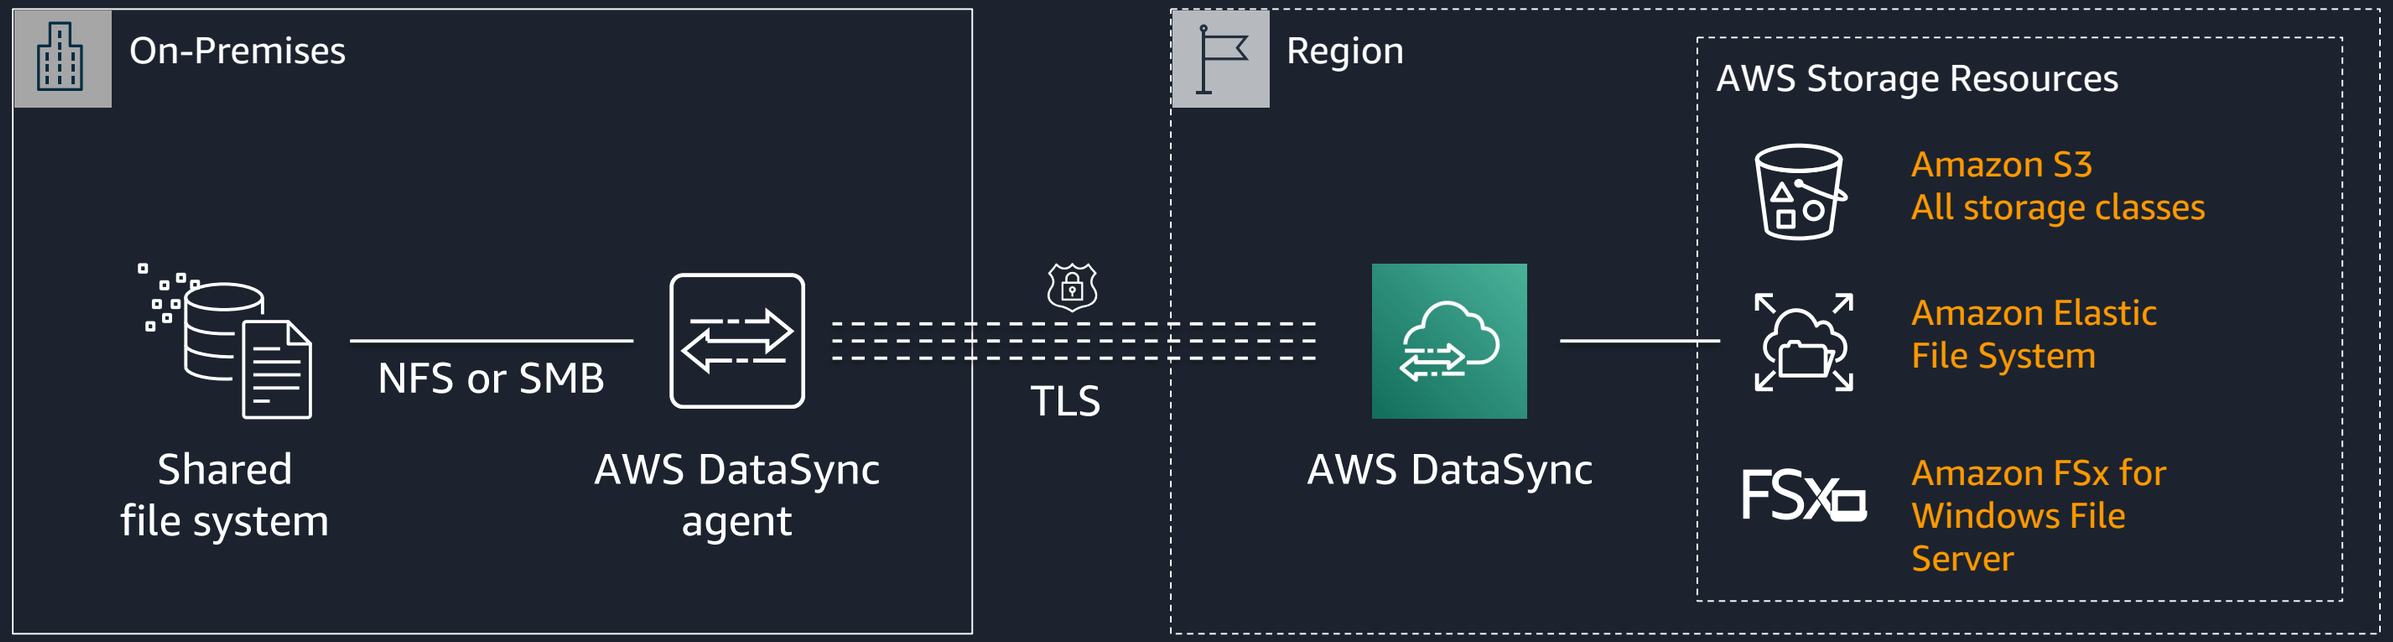

Deploy agent on VMware or EC2 for efficient access to local NFS or SMB server → Secure highly parallel transfers using optimized network protocol → Fully managed service scales to send or receive data from agent → Optimized reads and writes to Amazon S3, Amazon EFS or Amazon FSx for Windows File Server

aws storage

© 2020, Amazon Web Services, Inc. or its Affiliates.

AWS DataSync - In-Cloud Storage Destinations

Amazon S3

- Object storage service
- Offers industry-leading scalability, data availability, security, and performance
- 99.999999999% (11 9's) of durability
- Use Cases include websites, mobile applications, backup and restore, archive, enterprise applications, IoT devices, and big data analytics

Amazon Elastic File System (EFS)

- Simple, scalable, fully managed elastic NFS file system
- Available for use with AWS Cloud Services
- Built to scale to petabytes of data
- Use Cases include home directories and business-critical applications requiring an NFS filesystem

Amazon FSx for Windows File Server

- Fully managed native Microsoft Windows File System
- Available for use with AWS Cloud Services
- Full support for SMB protocol, Windows NTFS, and Microsoft AD Integration
- Use Cases include home directories, media workflows, and business applications

aws storage

© 2020, Amazon Web Services, Inc. or its Affiliates.

COX AUTOMOTIVE™

"The cool thing about having file gateway, is that anybody that expects to write via NFS today, can continue to do that tomorrow, and next month and next quarter…So you can actually start gaining in the intermediate phase before you're in the end state…and we think that's really valuable."

Dan Boisvert
Lead Software Engineer

Problem

- Stacks of disk arrays on-premises were expensive and required a lot of space
- Complex architecture and cache hierarchy
- Many readers via NFS

Solution

- AWS DataSync to transfer bulk data and active datasets to cloud
- File Gateway for local access to cloud data
- Active/active multi-region and versioning with lifecycles

Outcome

- Saved ~85% on storage, per location
- Storage engineers focused on high-value activities

aws storage

© 2020, Amazon Web Services, Inc. or its Affiliates.

DEMO – On-Premises NFS File Server Migration

Goal

Migrate on-premises NFS server data to Amazon S3, validate access to Amazon S3 from File Gateway, and shut down the on-premises NFS server

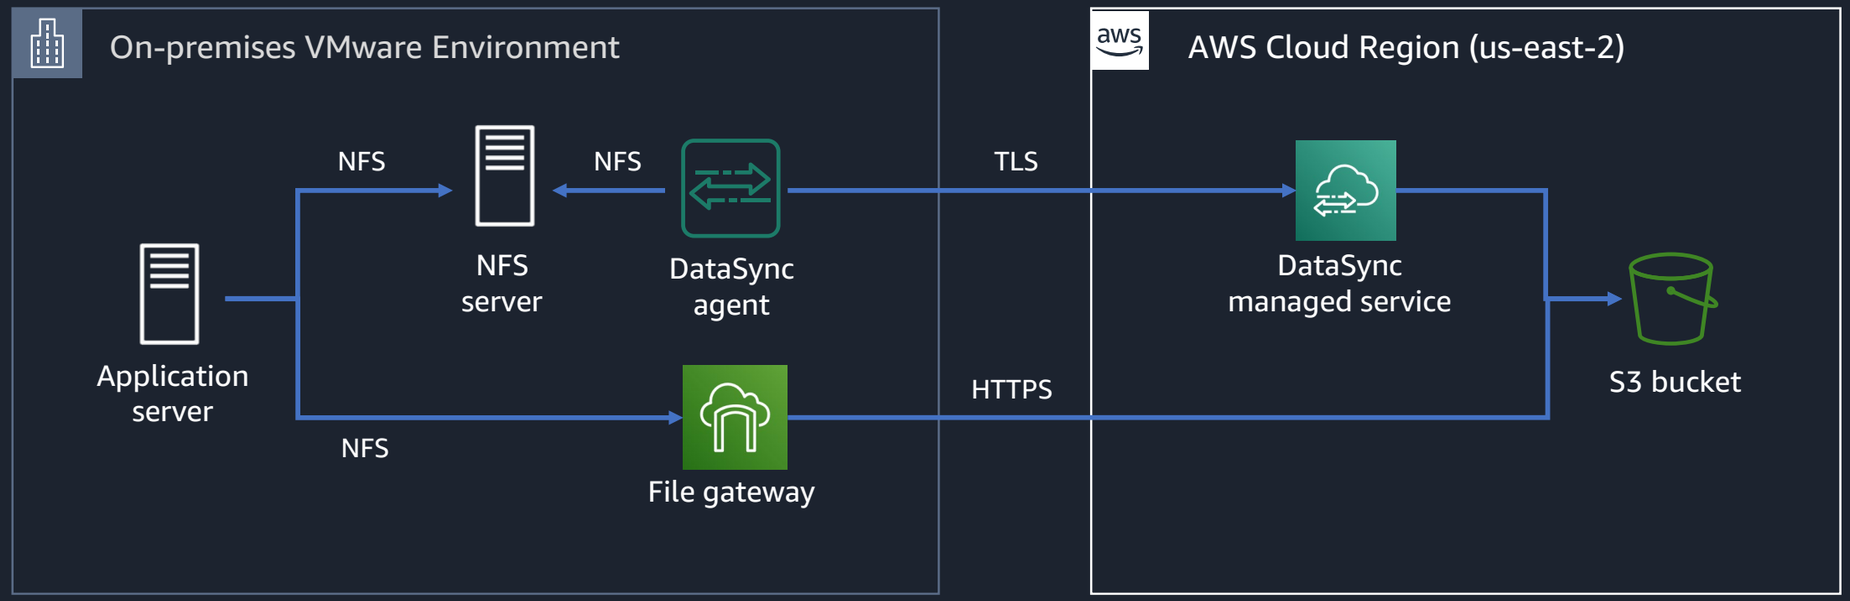

© 2020, Amazon Web Services, Inc. or its Affiliates.

aws storage

Initial Sync Task with DataSync

Step 1: Copy file data to Amazon S3 using AWS DataSync

Create a DataSync task via the DataSync agent, create source and destination locations, trigger the task to copy files to S3

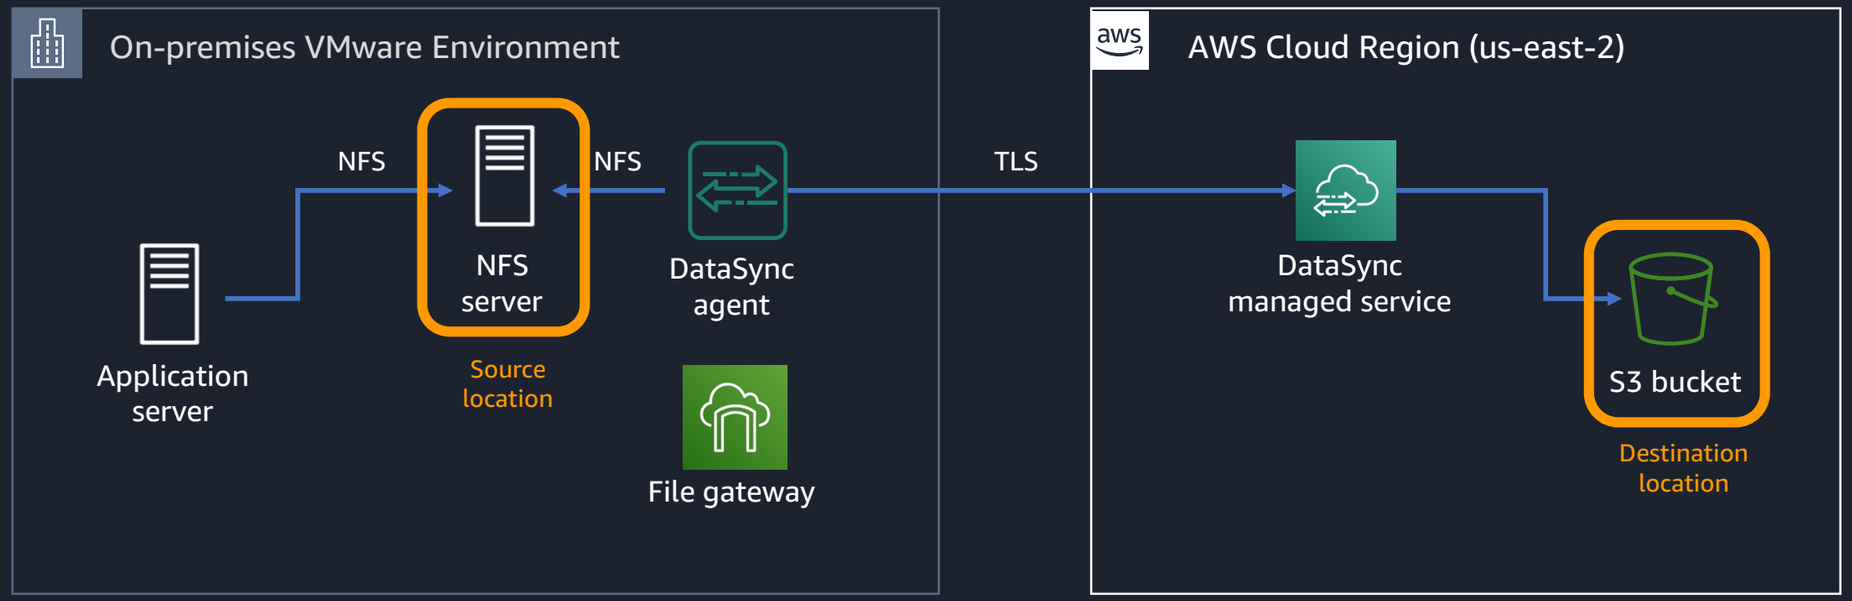

© 2020, Amazon Web Services, Inc. or its Affiliates.

aws storage

Export the new NFS Fileshare with Storage Gateway

Step 2: Access the data in the S3 bucket on-premises using File Gateway

Create an NFS share, mount the share on the application server

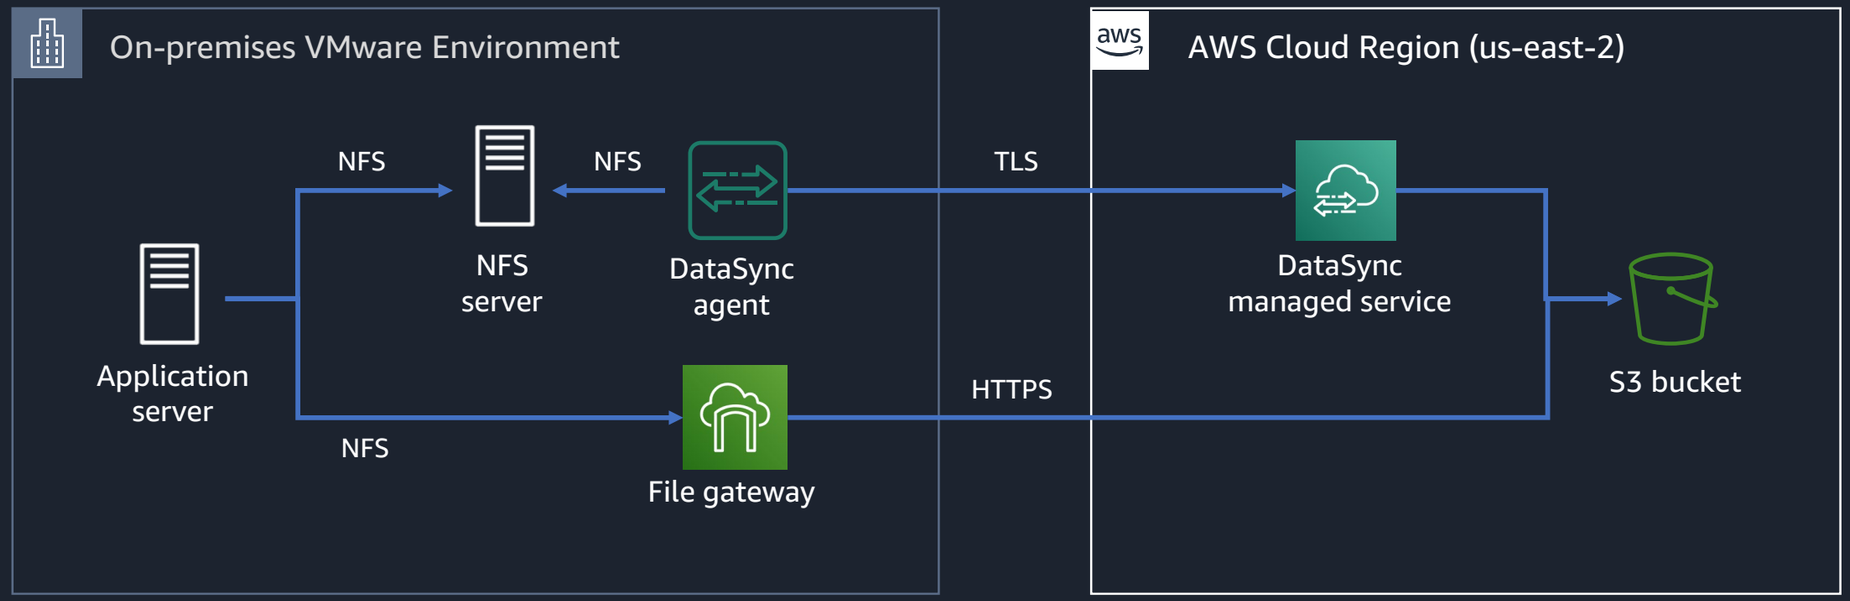

© 2020, Amazon Web Services, Inc. or its Affiliates.

aws storage

Delta Sync

Step 3: Perform an incremental transfer to get remaining data

Rerun the DataSync task to copy remaining files from the NFS server to the S3 bucket

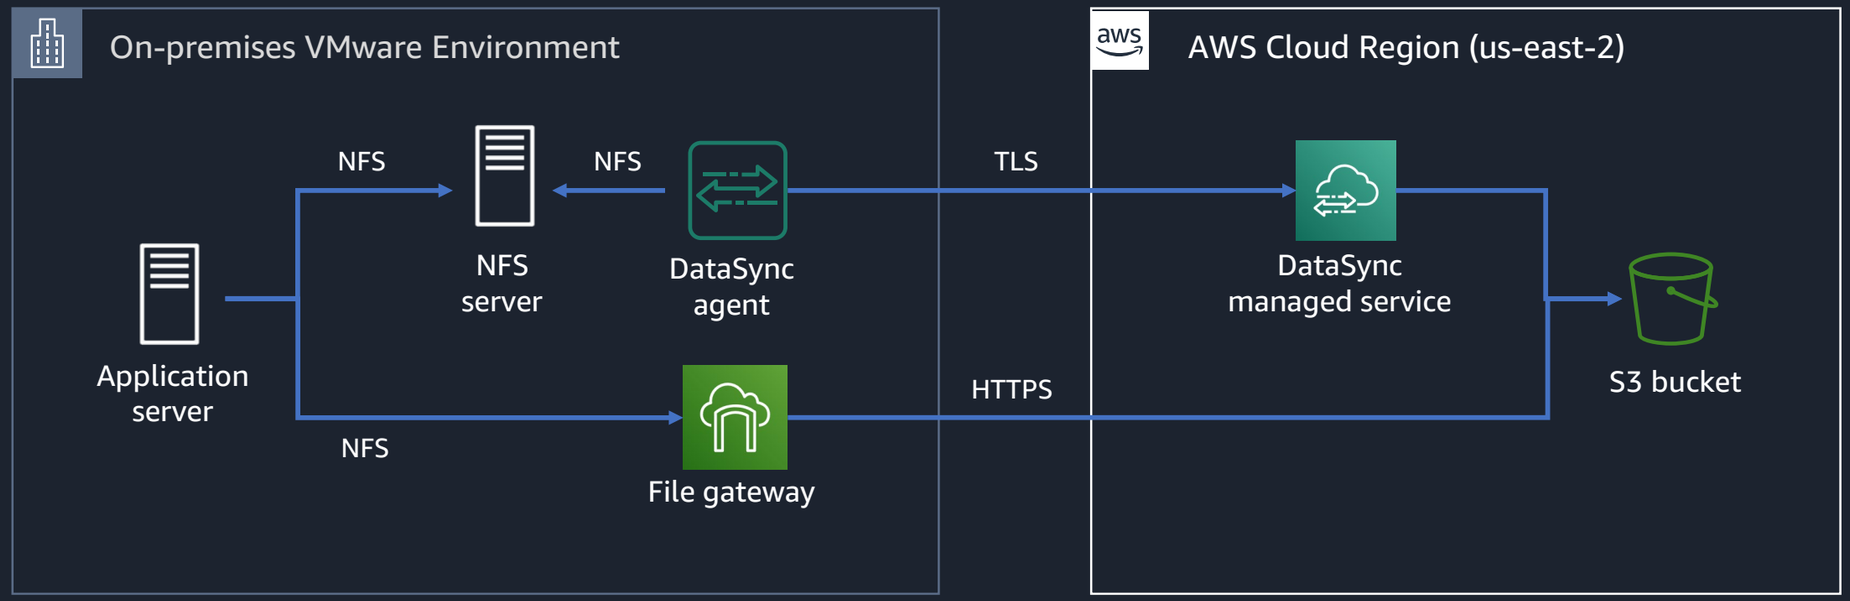

© 2020, Amazon Web Services, Inc. or its Affiliates.

aws storage

Final Cutover

Step 4: Cutover to File Gateway and shut down the NFS server

Unmount the NFS server on the application server, all traffic now goes to File Gateway

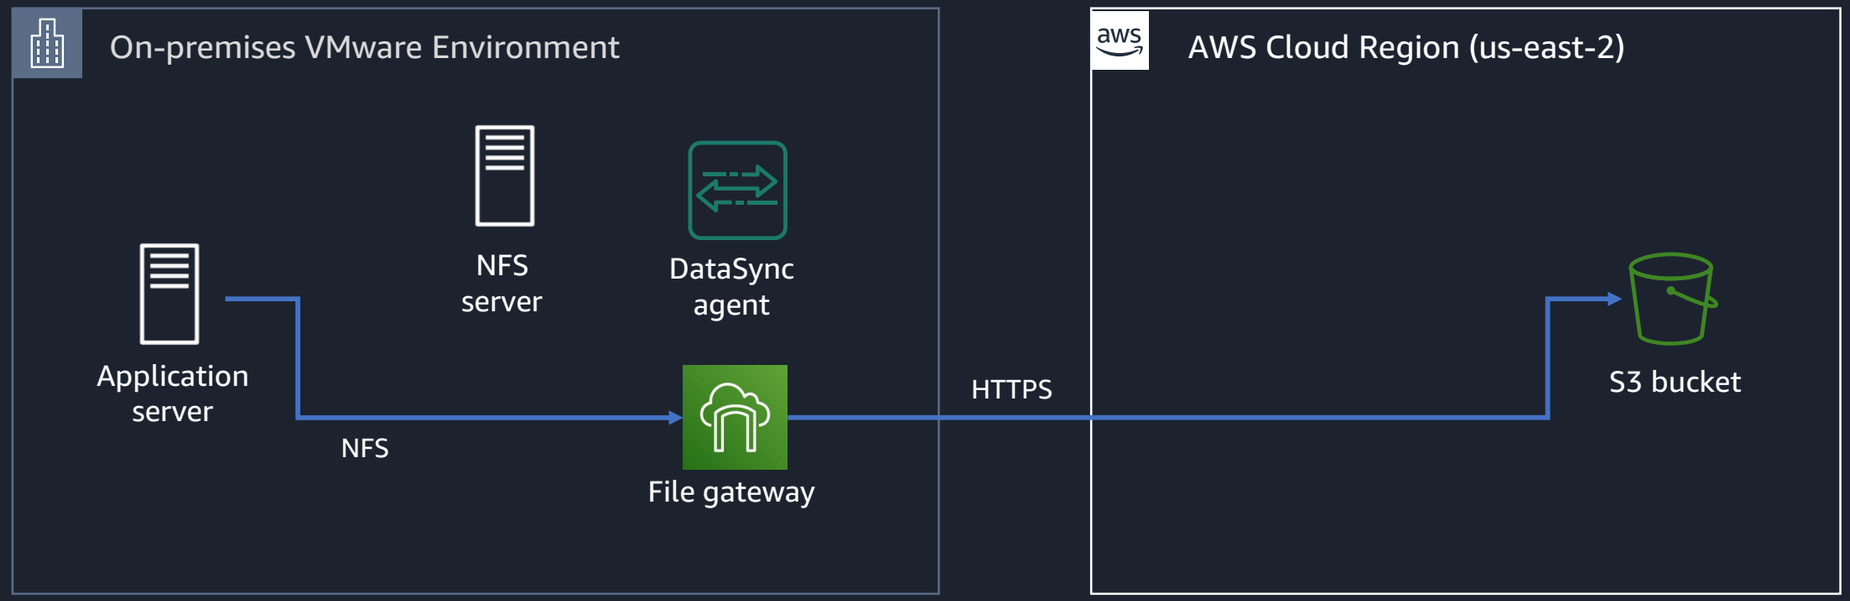

© 2020, Amazon Web Services, Inc. or its Affiliates.

aws storage

Review – On-Premises NFS File Server Migration

Migration Complete

Migrated on-premises NFS server data to Amazon S3 with AWS DataSync, validated NFS access to Amazon S3 from AWS File Gateway, and shut down the on-premises NFS server

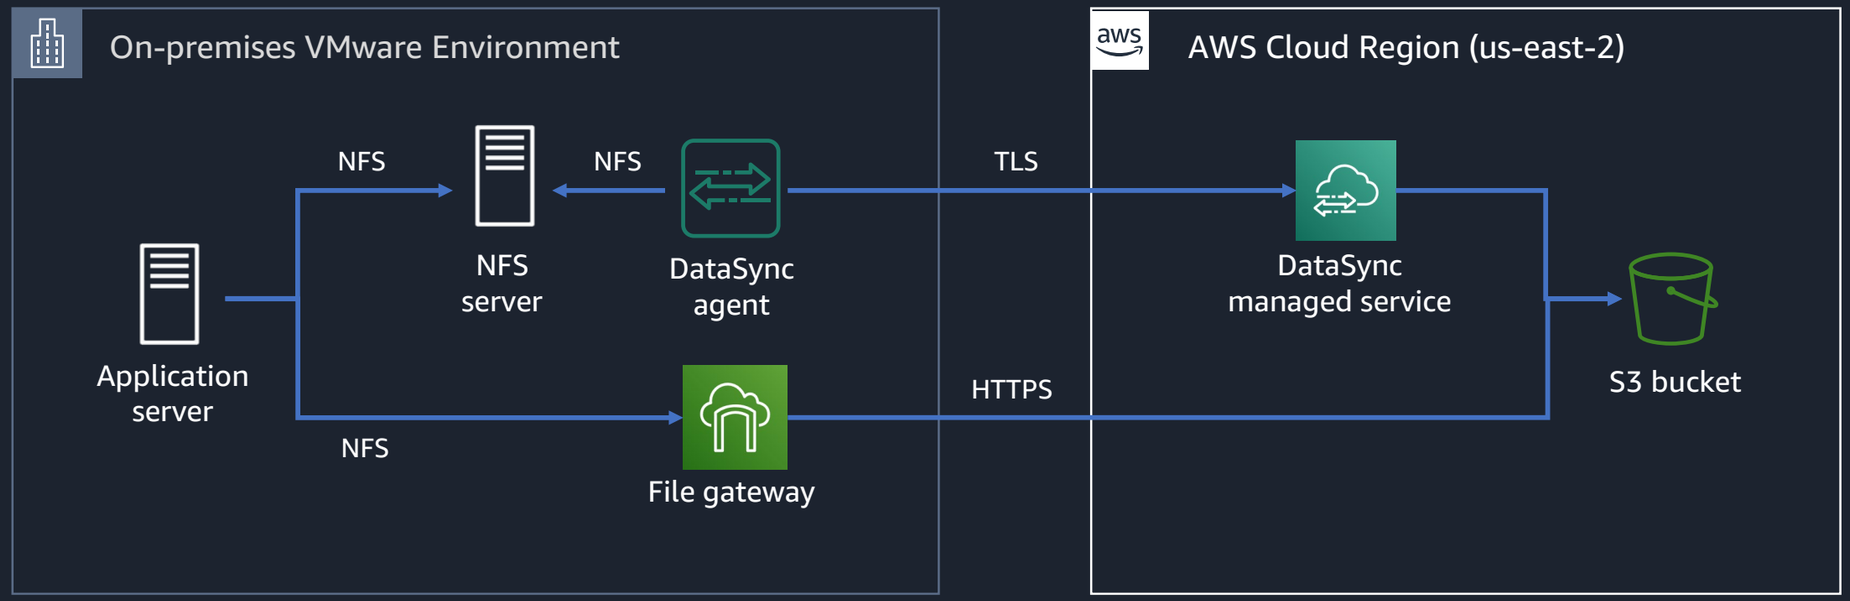

© 2020, Amazon Web Services, Inc. or its Affiliates.

aws storage

Learn More

- AWS DataSync - aws.amazon.com/datasync
- AWS Storage Gateway - aws.amazon.com/storagegateway
- Amazon S3 - aws.amazon.com/s3/
- Amazon EFS - aws.amazon.com/efs/
- Amazon FSx for Windows File Server - aws.amazon.com/fsx/windows/
- Migration workshop - github.com/aws-samples/aws-datasync-migration-workshop

© 2020, Amazon Web Services, Inc. or its Affiliates.

aws storage

Q&A

Thank you!

© 2020, Amazon Web Services, Inc. or its Affiliates.

aws storage