



Deep Dive: Hybrid Cloud Storage with AWS Storage Gateway


Paul Reed, Principal Product Manager
Peter Levett, Storage Specialist Solution Architect

June 26, 2018


Agenda

- Hybrid Cloud Storage
- The AWS Storage Gateway family
- File Gateway
 - Use case: Hybrid cloud data processing workflows with Amazon S3
 - Use case & Demo: Backing up Microsoft SQL Server to Amazon S3 with native SQL tools
- Volume Gateway
 - Use case: Flexible hybrid backup & recovery options with Amazon EBS Snapshots
- Tape Gateway
 - Use case: Migrating tape archives to Amazon S3 & Amazon Glacier

Complete Set of Data Building Blocks



Two Cloud Deployment Options



"All-In" Cloud

Fully deployed in
the cloud



Hybrid

Deployed on-premises
and in the cloud

So, how can you make storage hybrid?

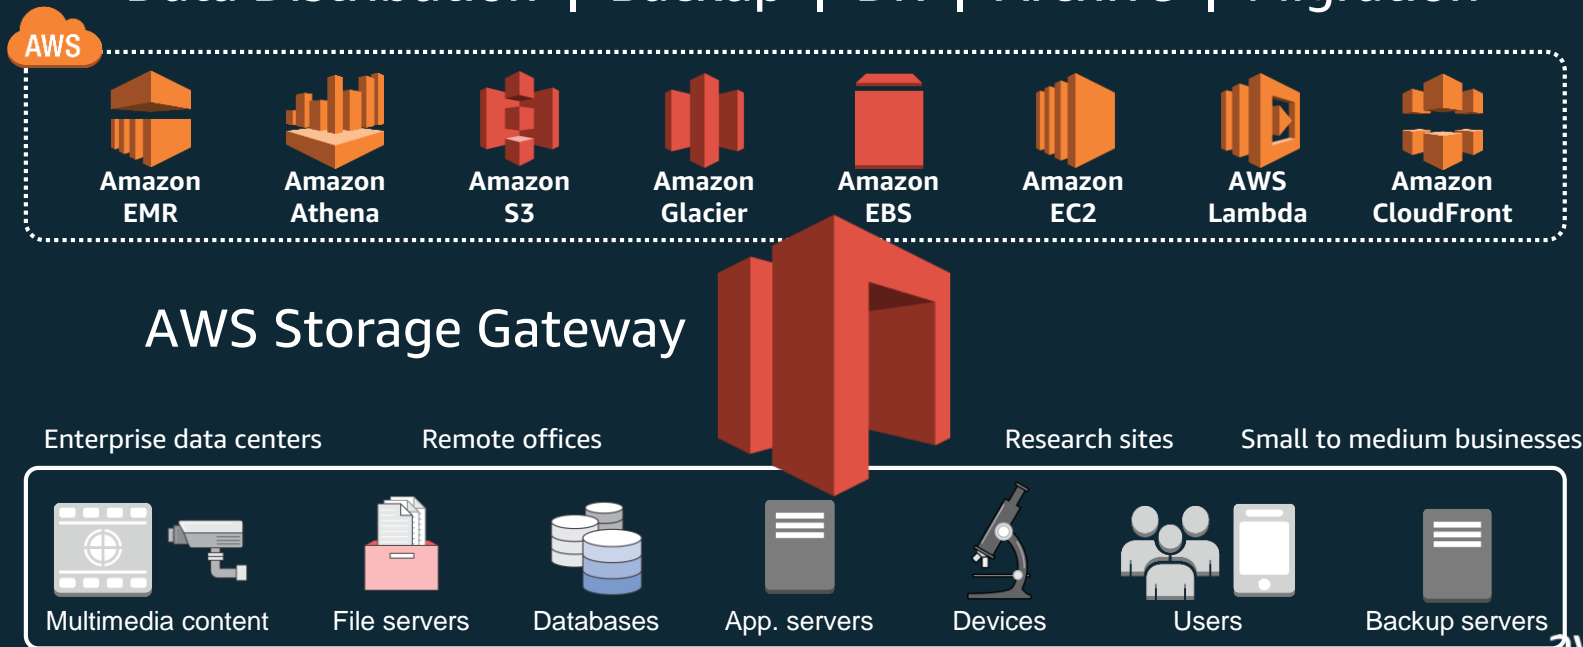
Gateways

So what are gateways?



AWS Storage Gateway enables a spectrum of hybrid use cases

Analytics | File Services | Production Tiering | Data Processing
Data Distribution | Backup | DR | Archive | Migration



AWS Storage Gateway Family



File Gateway

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



Volume Gateway


Cloud-backed block storage volumes presented on-premises with EBS snapshots and clones



Tape Gateway

Cloud storage for backups through a drop-in replacement for physical tape infrastructure

AWS Storage Gateway Family



Gateway provides applications

- Protocol conversion and device emulation
- Caching (read-through / write-back)
- Optimized data transfer

Native storage in AWS

- Objects in S3 (file)
- Snapshots in EBS (volume)
- Archival in Glacier (tape)

What's New for AWS Storage Gateway?

Last year

- 12 feature launches
- Available in all commercial Regions
- HIPAA-eligible
- **File Gateway** refresh cache for multi-site file sharing
- **Volume Gateway** cloning volumes for faster DR
- **Tape Gateway** support for Commvault and Arcserve, improved performance (pre-fetch and faster retrievals)


So far this year

- Available in AWS GovCloud (US)
- AWS CloudWatch Events for automation, support for Requester Pays and S3 One Zone - IA storage class
- Support for EMC NetWorker v9, NovaStor DataCenter
- KMS encryption for files, volumes, and tapes

File Gateway & Hybrid Use Cases for Data Processing and Backup

File Gateway

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



Reduce on-premises
storage infrastructure

Fully managed local
cache for low latency
access

Durability, scalability,
and reliability of
Amazon S3 storage

Last Week We Added SMB to File Gateway

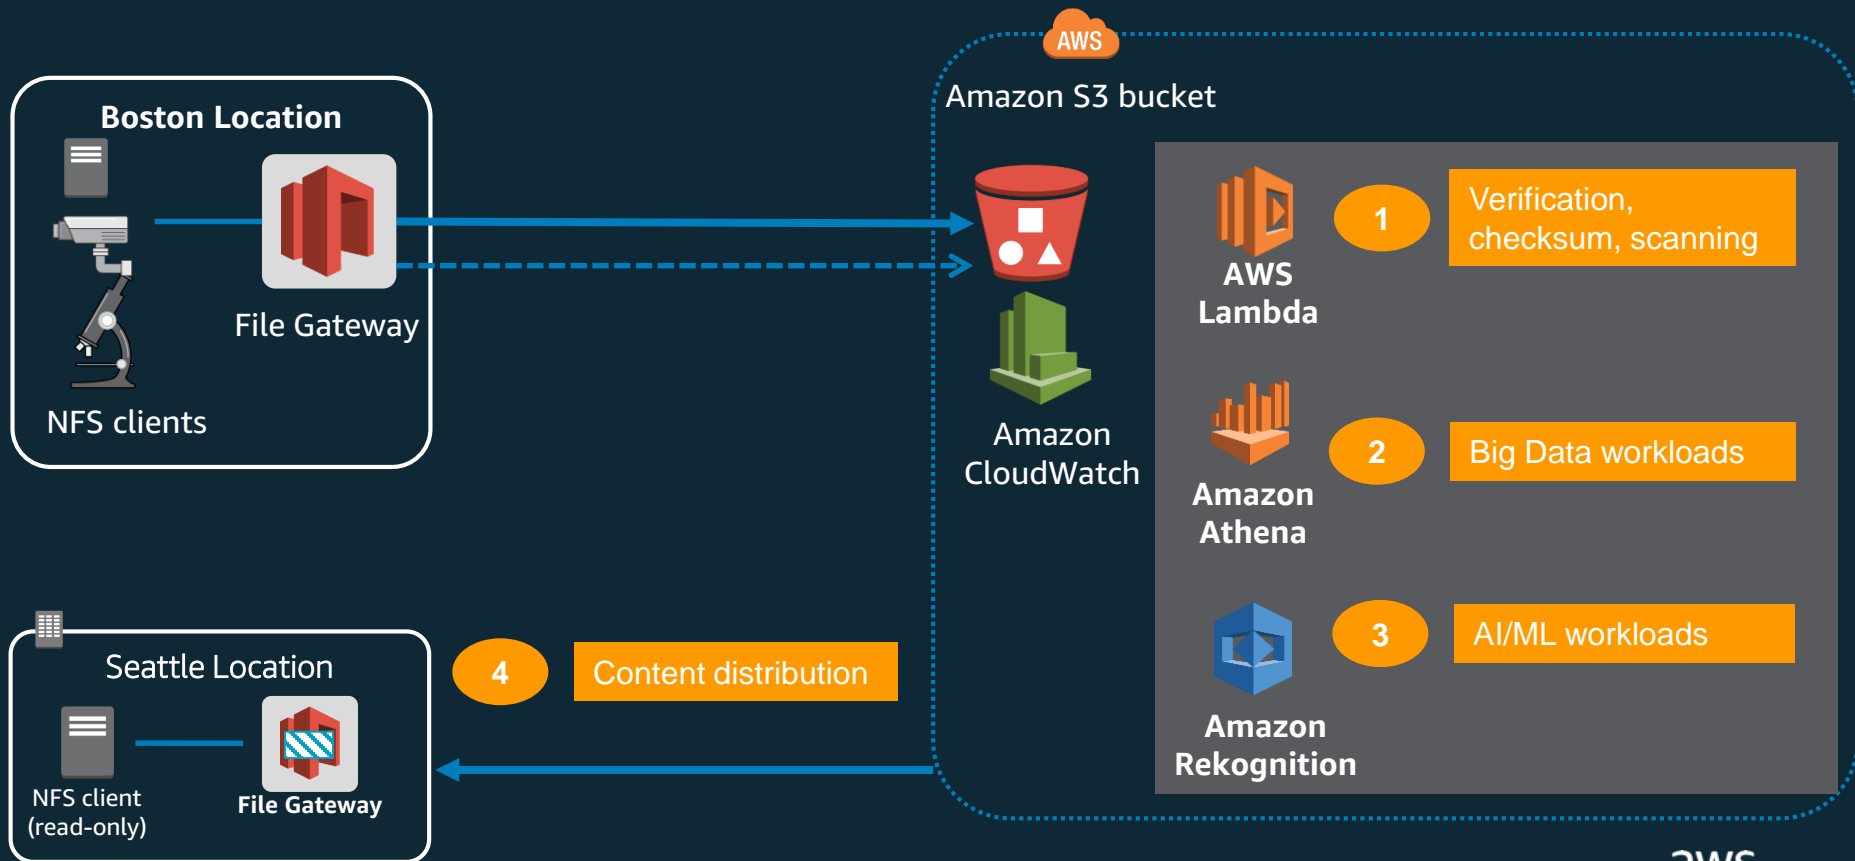
Just
Launched!



Store and access objects in Amazon S3 buckets
from file-based Windows applications


- Supports SMB v2 & v3
- Control access to SMB shares using Active Directory (AD)
- Control access to objects using POSIX ACLs (compatible with NTFS)

Hybrid File Use Cases: Backup, Analytics, Machine Learning



How it works

File Gateway – Mapping Files to Objects



File Gateway – Control Over Data Storage and Access



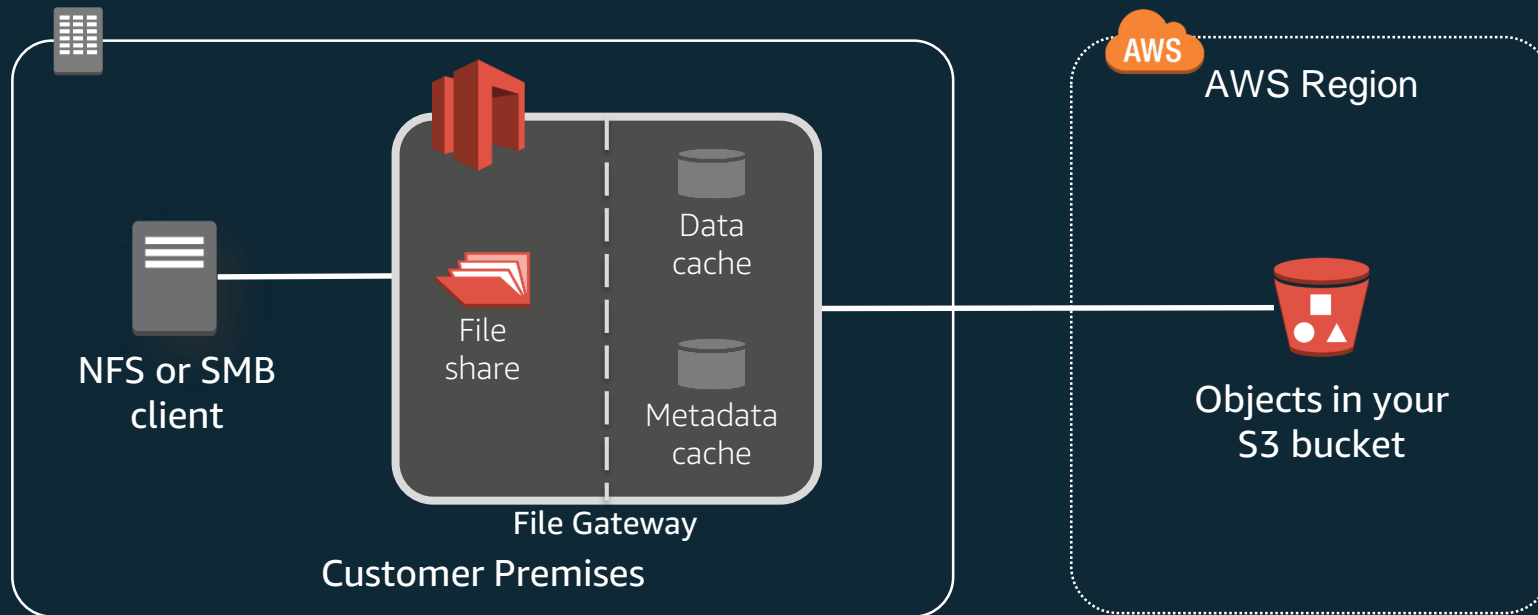
NFS or SMB file share options

- Restrict access by IP (NFS) or AD (SMB)
- Read-only/read-write
- Default ownership and permissions
- User squashing (NFS)

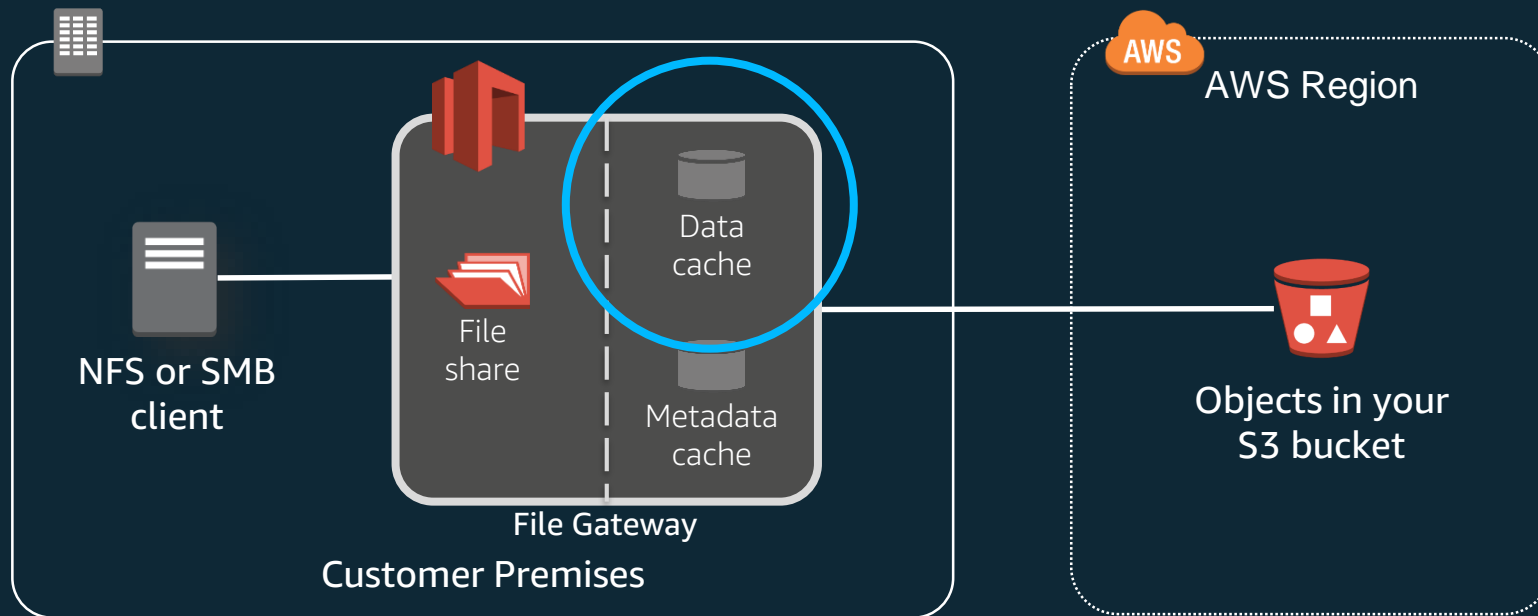
Amazon S3 options per bucket

- IAM role for access
- Storage class
- Object encryption with AWS KMS
- Guess MIME type

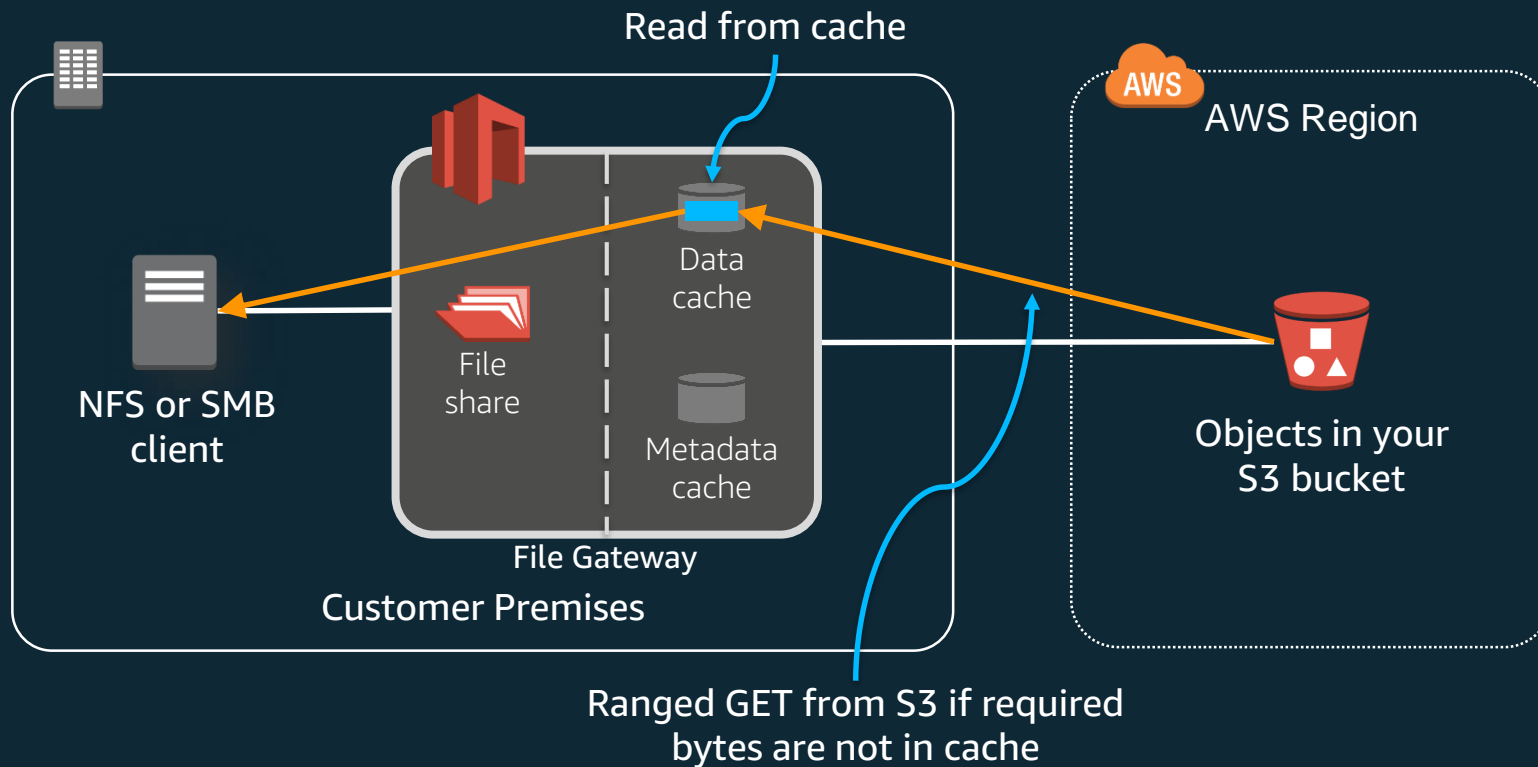
File Gateway – Local Caching For Low Latency Operations



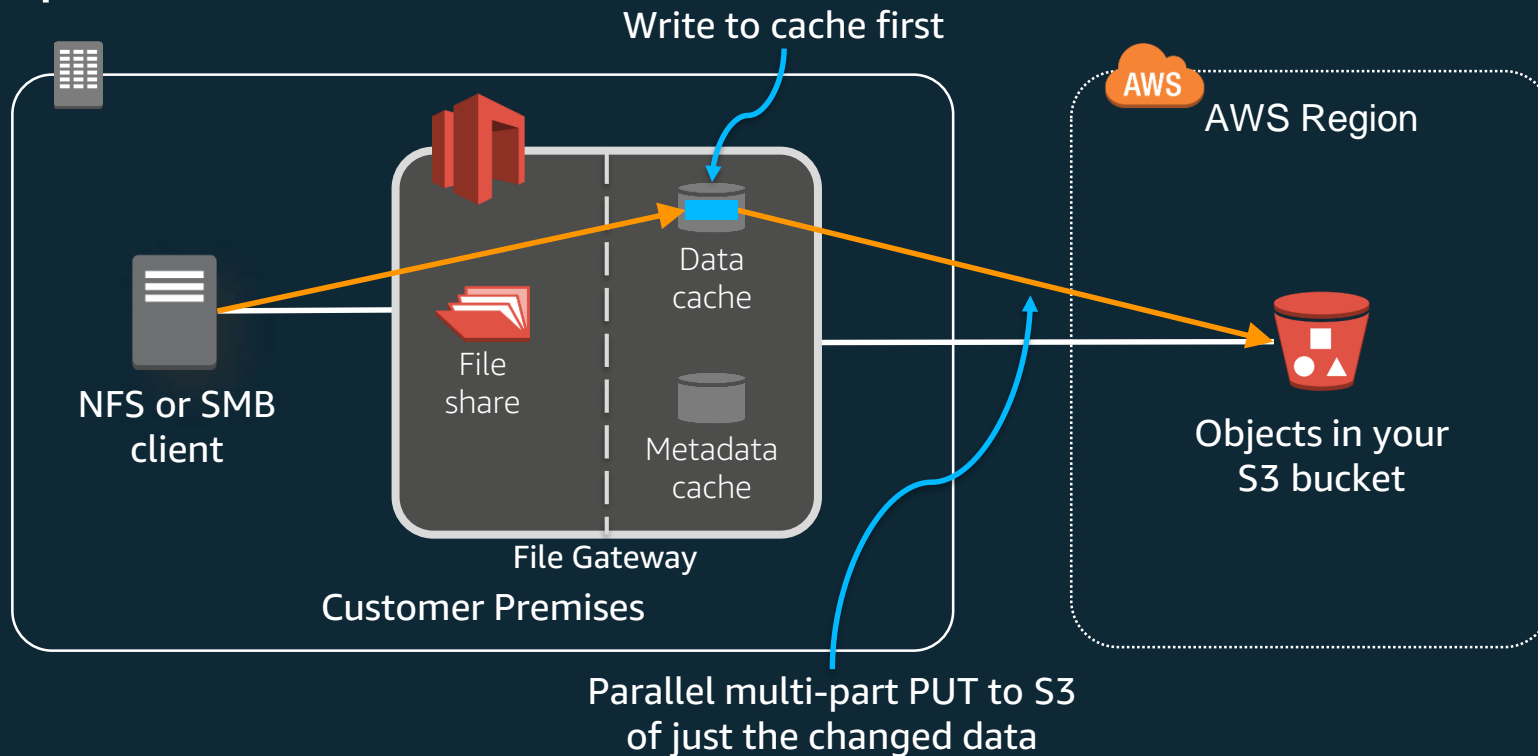
File Gateway – Data Cache For File Reads And Writes



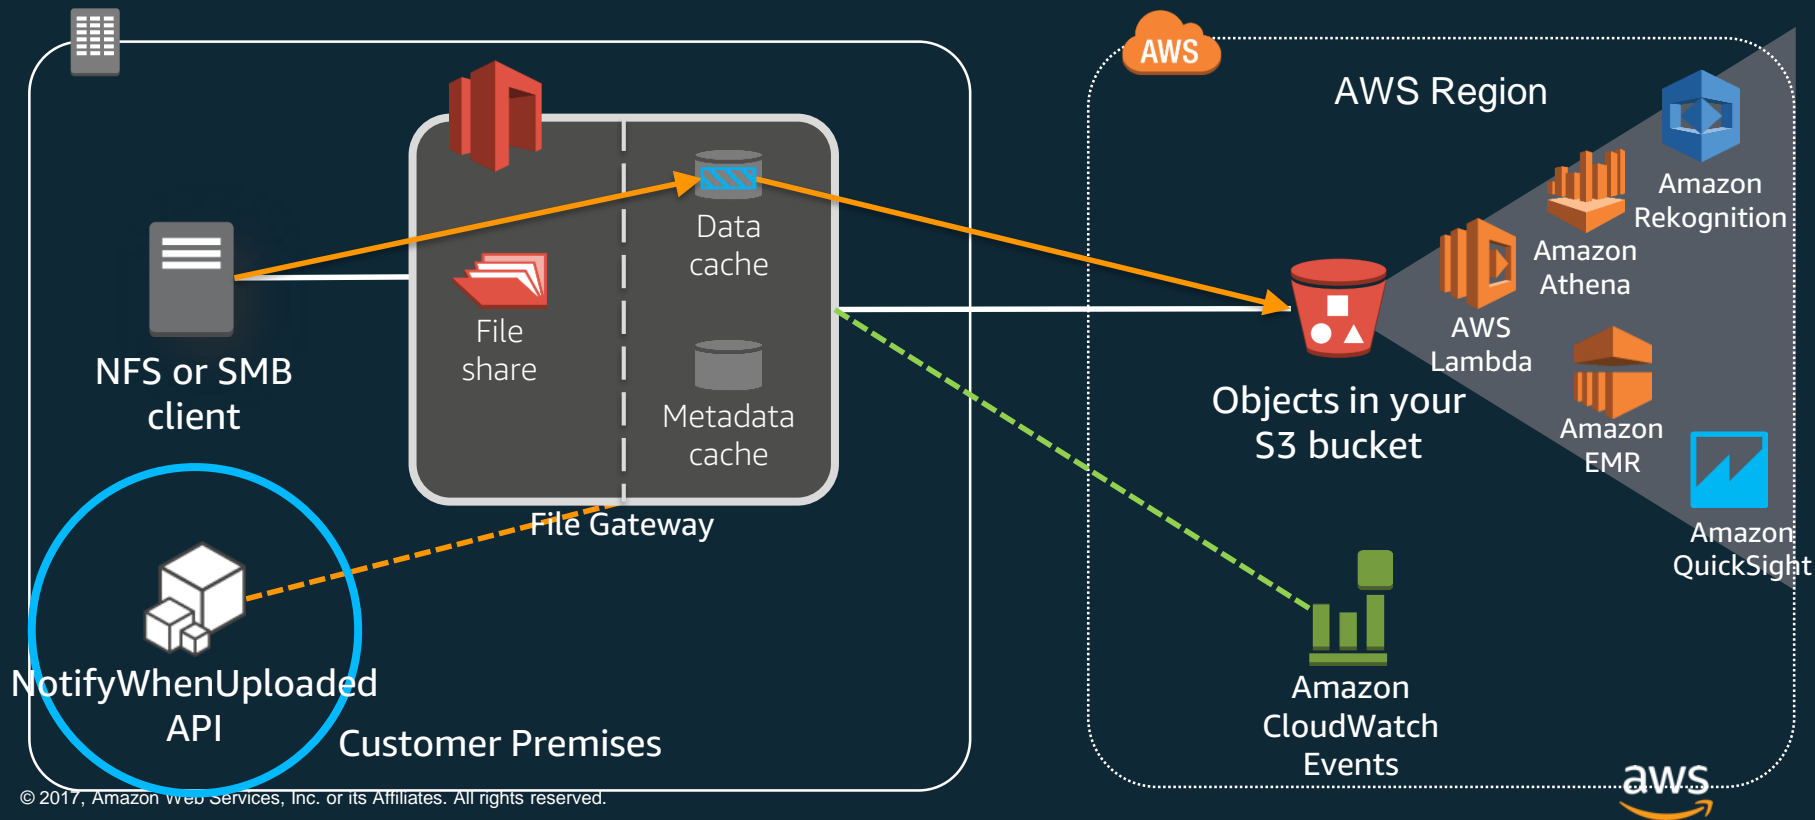
File Gateway – Data Read Through The Cache




File Gateway – Data Written To The Cache Then Uploaded



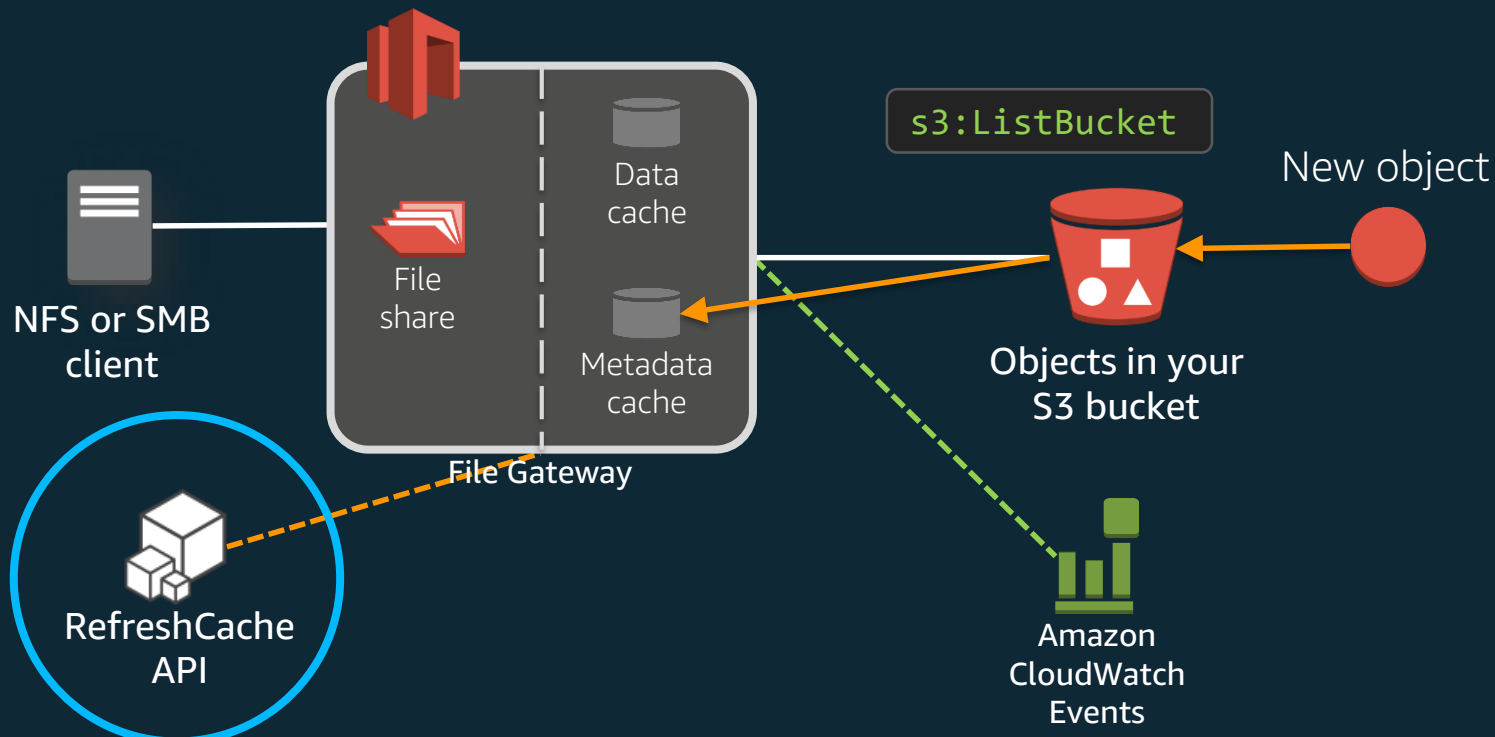
File Gateway – Upload Notifications for Automation



File Gateway – Metadata Cache



File Gateway – Updating the Metadata Cache



Use Case Demo: Backing Up Microsoft SQL Server with Native SQL Tools and File Gateway

Volume Gateway and Hybrid Recovery & Migration Use Cases

Volume Gateway


Cloud-backed block storage presented on-premises

- Tier snapshots or whole volumes to the cloud to reduce SAN/NAS mgt.
- Flexible recoveries in-cloud or on-premises with snapshots and clones
- Common uses: backup and restore, disaster recovery, data migrations




Volume Gateway – Stored Mode

Low-latency access to all your data with point-in-time backup to the cloud through EBS snapshots



Volume Gateway – Cached Mode

Reduce on-premises storage, caching frequently used data local to your application, with 100% of your data in the cloud



Restore, Recover, and Migrate Volumes with Amazon EBS Snapshots and Volume Clones

Volume Gateway – Data Protection with Snapshots and Clones



EBS snapshots

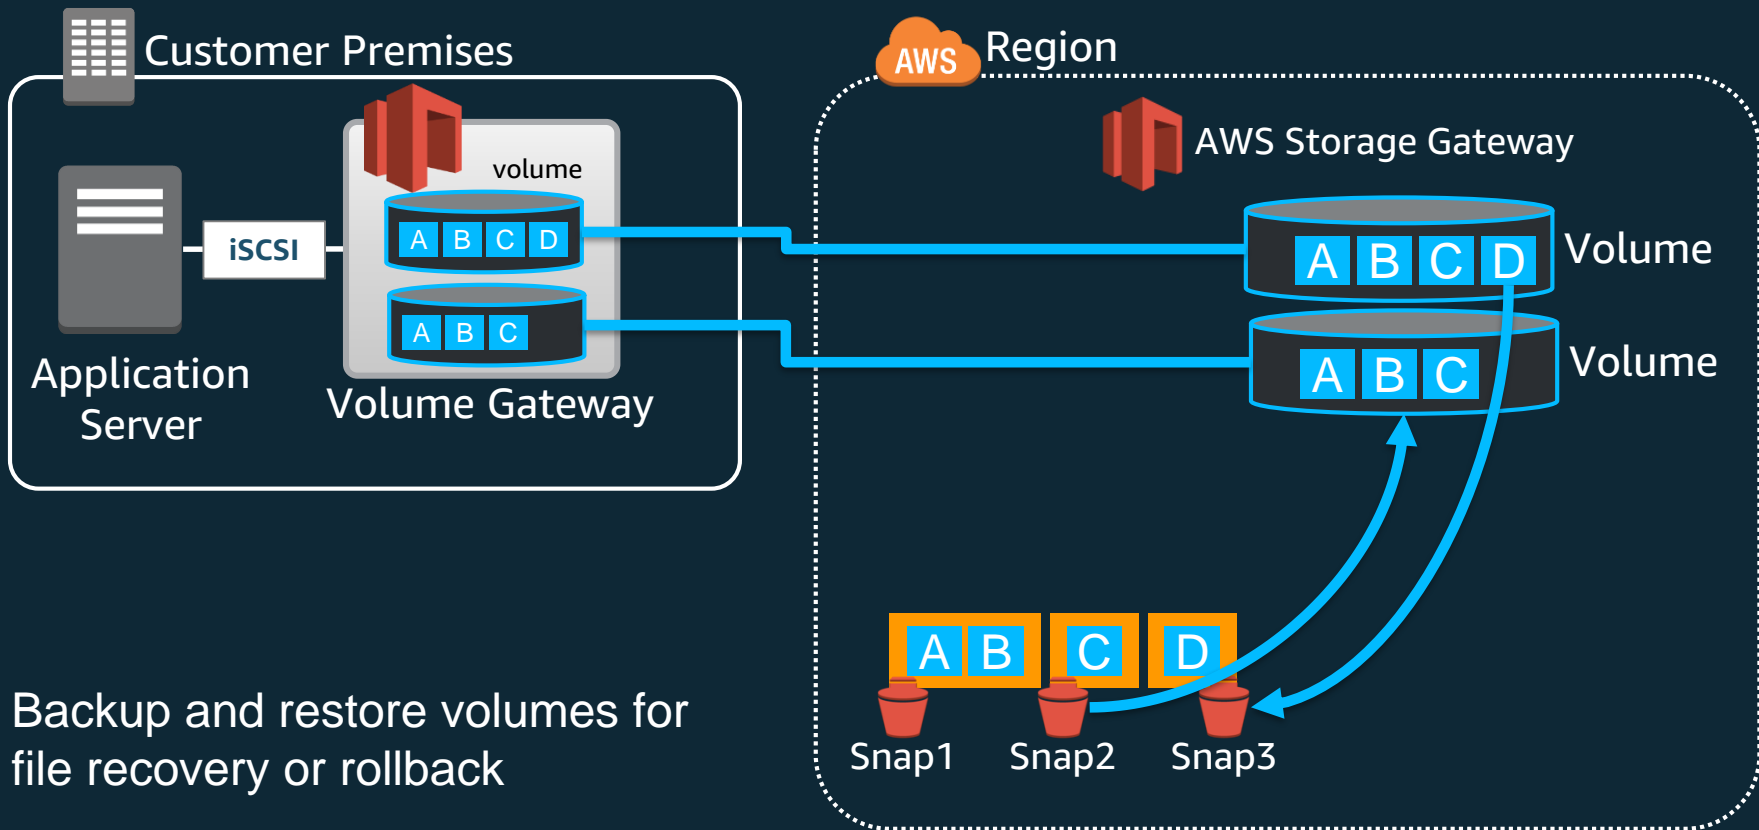
Point-in-time backups of a stored or caches volumes
Created on-demand or on a configurable schedule
Can use to create an EBS or Storage Gateway volume



Volume clones

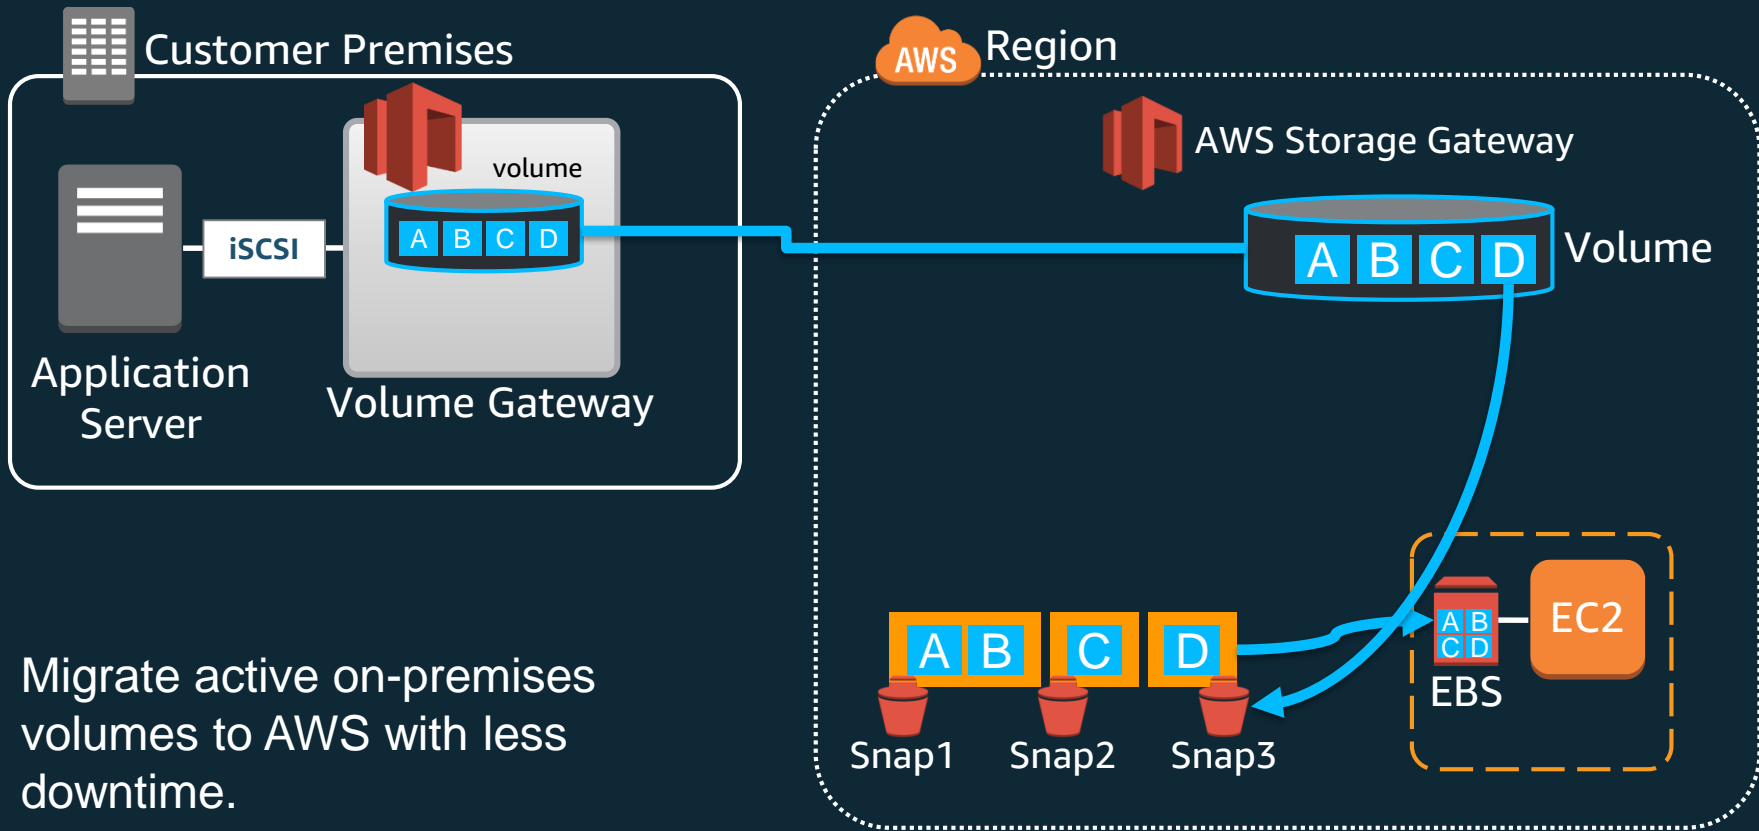
Instant real-time copy of a cached volume
Represents current state of volume stored in AWS
Can use to create Storage Gateway volume

Restore Data from Earlier Point-in-Time Snapshot



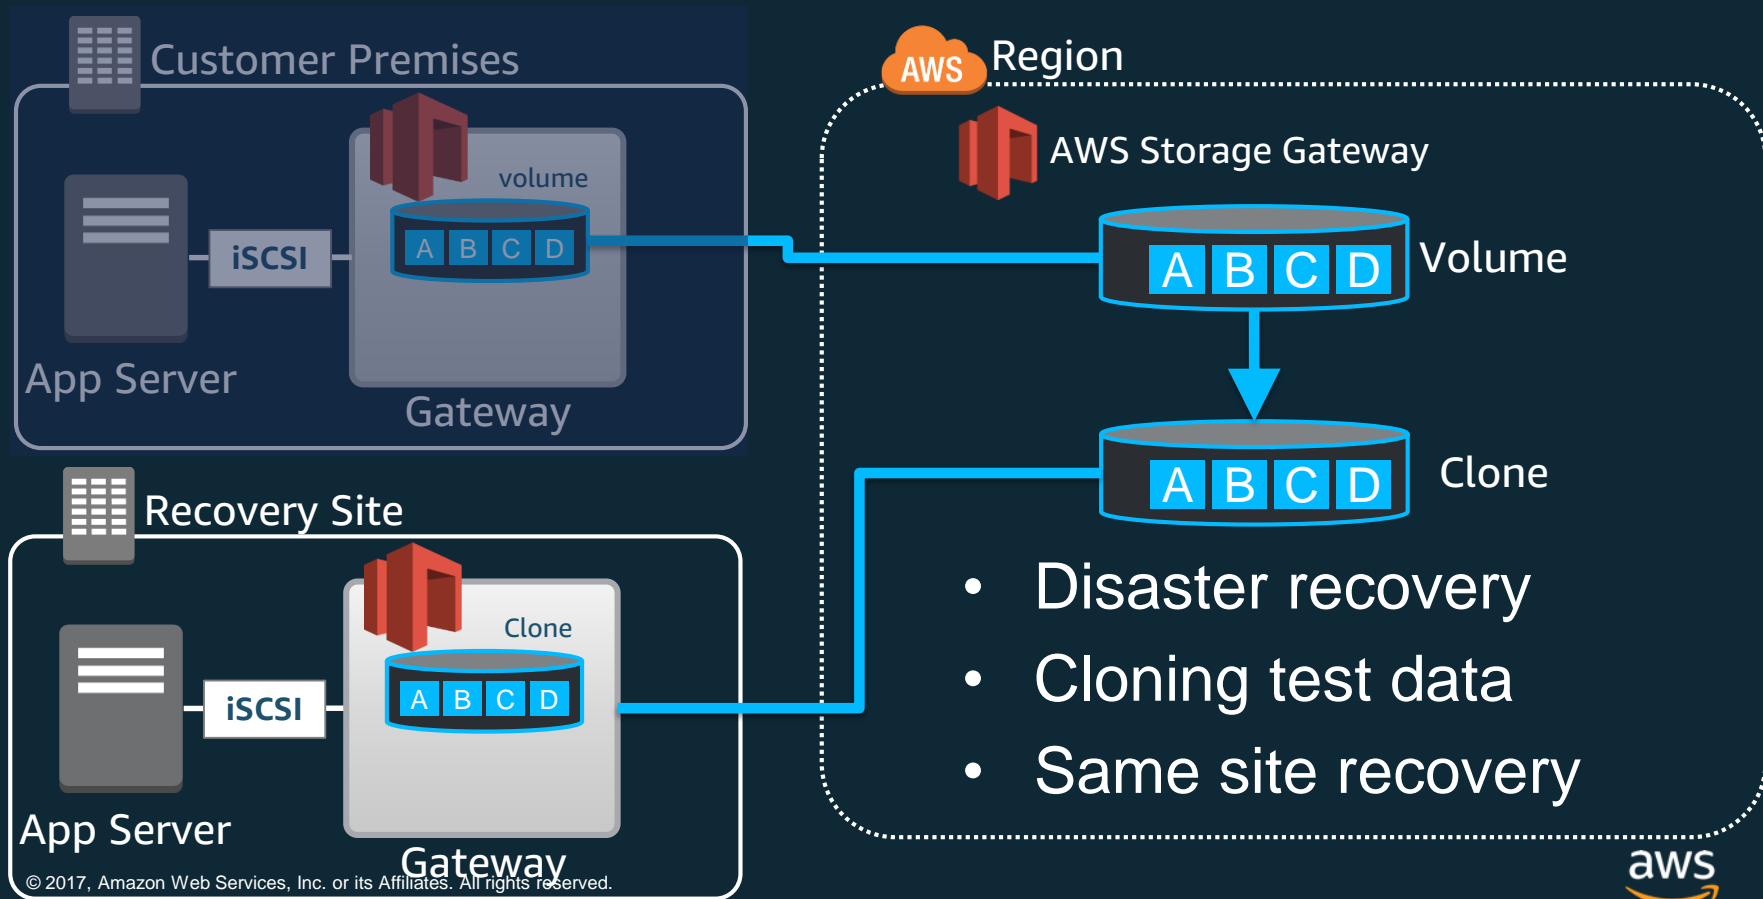
Backup and restore volumes for
file recovery or rollback

Migrate Data to EBS Using Latest Snapshot



Migrate active on-premises volumes to AWS with less downtime.

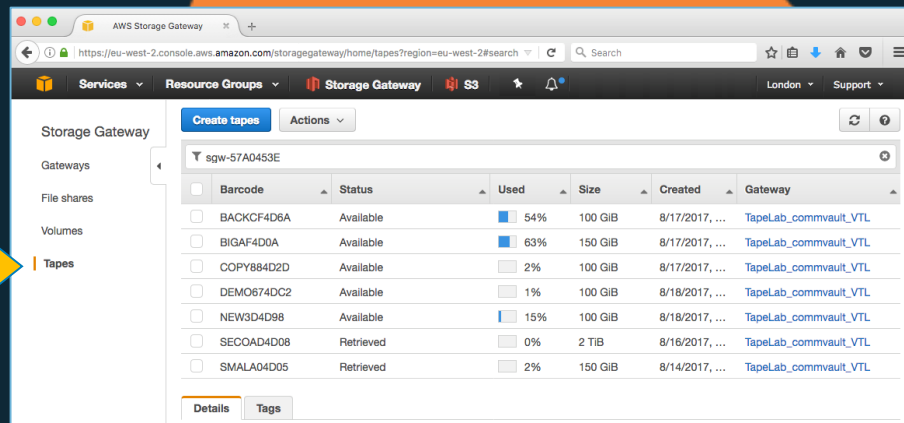
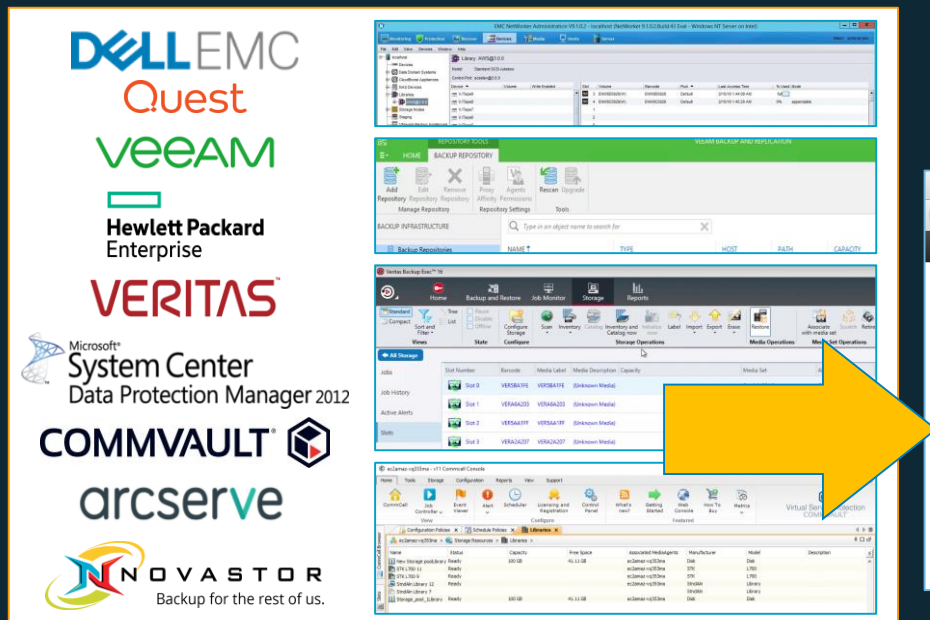
Disaster Recovery With Volume Cloning



Hybrid Cloud Virtual Tape Backups to Amazon S3 & Amazon Glacier with Tape Gateway


Tape Gateway – Drop-in Replacement for Tape Backups

- Emulates a tape library - virtual tapes on Amazon S3 and Amazon Glacier
- Works with common backup apps, to support existing backup workflows
- Low-cost: predictable costs and reduced management

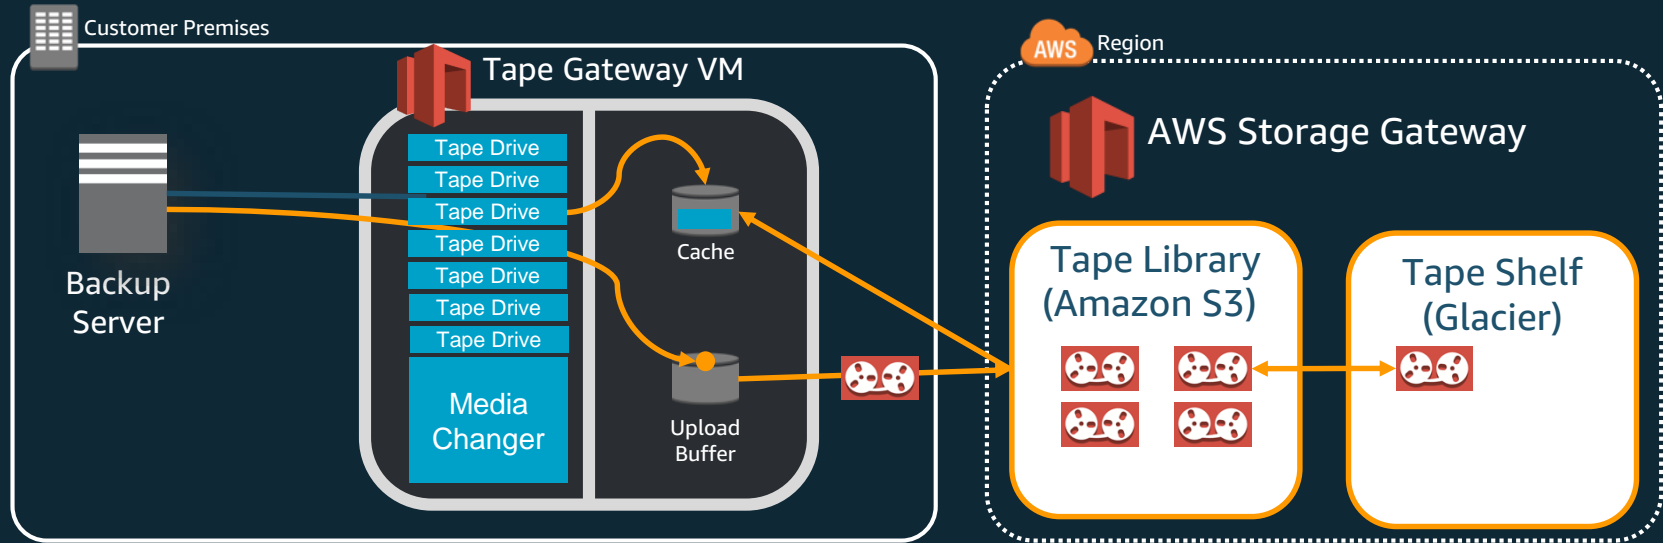


Tape Gateway – Cost Comparison, a Customer example

Every month, 100 TB written and 1 tape retrieval; 1-year retention



Tape Gateway – How the VTL works



- Emulates a physical tape device with a media changer and tape drives
- Scalable: Virtually unlimited tape storage in AWS
- Virtual tapes are written to S3 – data is in S3 when tape is in virtual library drive or slot
- 'Ejected' virtual tapes are marked read-only and moved to 'Tape Shelf' on Amazon Glacier
- Recovery: Retrieve tapes to library (3-5 hours) and read data to same or different gateway

Summary

Hybrid Storage Architectures with AWS

- When you need data and access on-premises with storage in-cloud
- AWS Storage Gateway - Seamlessly connect on-premises applications with AWS Cloud storage; No changes to existing applications
- Storage Gateway supports use cases throughout your cloud journey: Backups, tiered file & block storage, hybrid processing, migration

Thank you!

To learn more, visit aws.amazon.com/storagegateway