

Hybrid Cloud Storage with AWS Storage Gateway & Amazon S3

Paul Reed, Principal Product Manager

August 29, 2018

Why are we here?

You have data ...

and you want to store it in the cloud and access it on-premises



We'll be talking about data that is:


- Connected to on-premises applications
- Part of continuous hybrid workflows
- Backup for critical systems
- In bulk storage (e.g., NAS filers)

Why move data storage to the cloud?

Retirement



Management

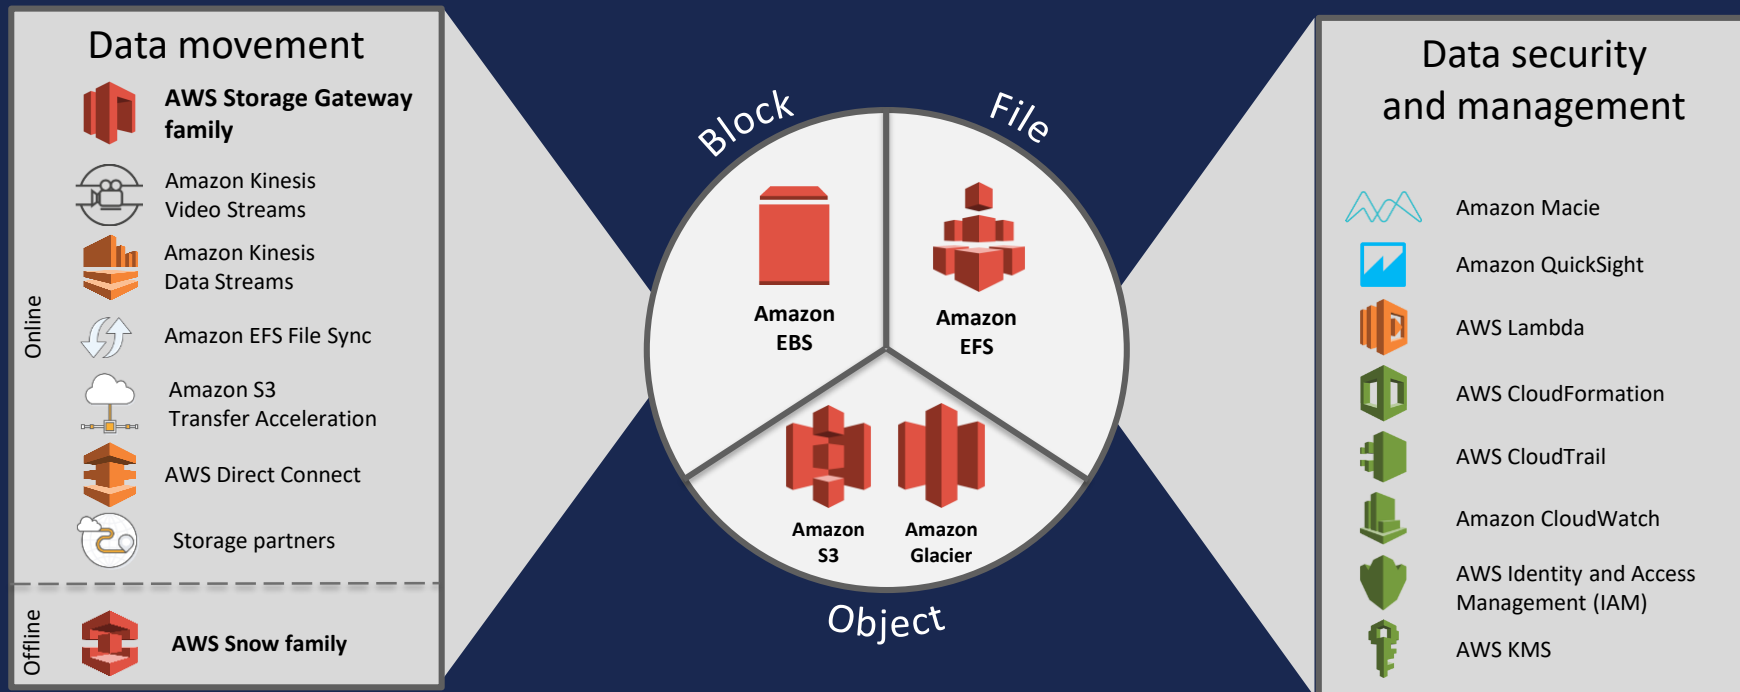


Cost

Date		INVOICE #Nth	
To: You, CIO & AP		Ship To	
Your expensive colo space or data center		Big Iron Storage Vendor	
Instructions		Tel (508) 512-4086 176 South RR St, Sunnyvale, CA	
Pay now, pay next year, and serv again in 3 years.			
Quantity	Description	Unit Price	Total
	4 big storage arrays, 2 for primary DC, 1 for DR, 1 for archive for prod, DR, and 1 for all the rest of the stuff you store	CMG, security? But after 75% discount...	\$80 a bit
	Installation services you could probably do yourself if you had the time & money, or you couldn't because it's black magic, duct tape and shoe strings.	More than you want to pay	The same, because you can't discount people
	Fancy new software feature packages	A bit, but it's worth it. Fast enter sum to do	\$50 a bit
	Integration with that other software thing you'll never use	A couple grand you won't notice	We'll discount it some
	Support you're not really happy with	10-15% of top line	Worth the cost
	On-site spare hard drives for faster replacement, because moving parts is, big drives take a long time to re-build and RAID isn't good enough		We'll call it a grand
	On-site spare controller, because stuff happens to chips and heat sinks and code		A couple grand
	On-site spare power supplies, because stuff happens to batteries and fans		We'll call it a grand
	Subtotal	\$	Yup, still a bit
	Shipping & Handling		Yes, it comes in a box
	Total Due		OUCH!

Due upon receipt
Thank you for your business!

Complete set of data building blocks



So, how can you make storage hybrid?

Gateways



AWS Storage Gateway enables a spectrum of hybrid use cases

Analytics | File services | Machine learning | Data processing
Data distribution | Backup | DR | Archive | Migration



Storage Gateway




Enterprise data centers

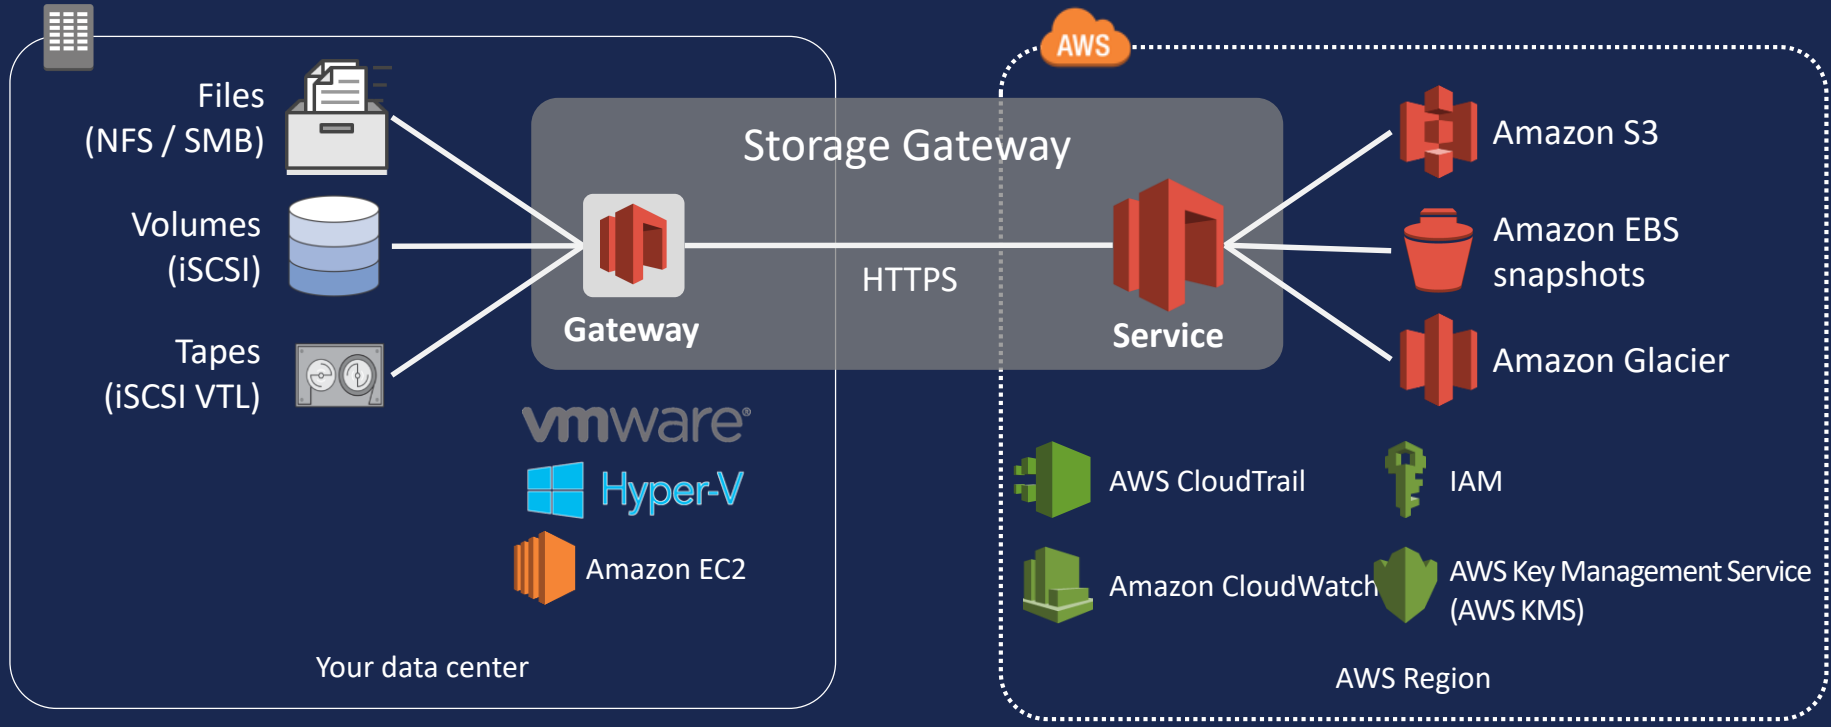
Remote offices

Research sites

Small to medium businesses



AWS Storage Gateway



AWS Storage Gateway family

Hybrid storage service enabling applications to seamlessly use AWS storage



File Gateway

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



Volume Gateway

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



Tape Gateway

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



AWS Storage Gateway key features



Standard storage protocols

AWS storage accessible without applications
needing to be modified



Fully managed cache

Frequently used data cached locally
for low-latency access



Optimized data transfer

Secure upload of changed data and downloads
requested data



Durable storage

On-premises application data natively stored in
Amazon S3, Amazon Glacier, and Amazon EBS
snapshots

File Gateway for hybrid cloud file workloads

File Gateway

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

Now supports SMB!



Per file share options

Restrict access by client IP (NFS) or Active Directory (SMB) users/groups

POSIX permissions for object-level access*

Read-only/read-write

* Compatible subset of NTFS

Per S3 bucket options

IAM role

Object storage class


Encryption with AWS KMS

Guess MIME type, requester pays, bucket owner ACL, etc.

File Gateway: Fully managed cache for low latency access

Some of your data
cached on gateway

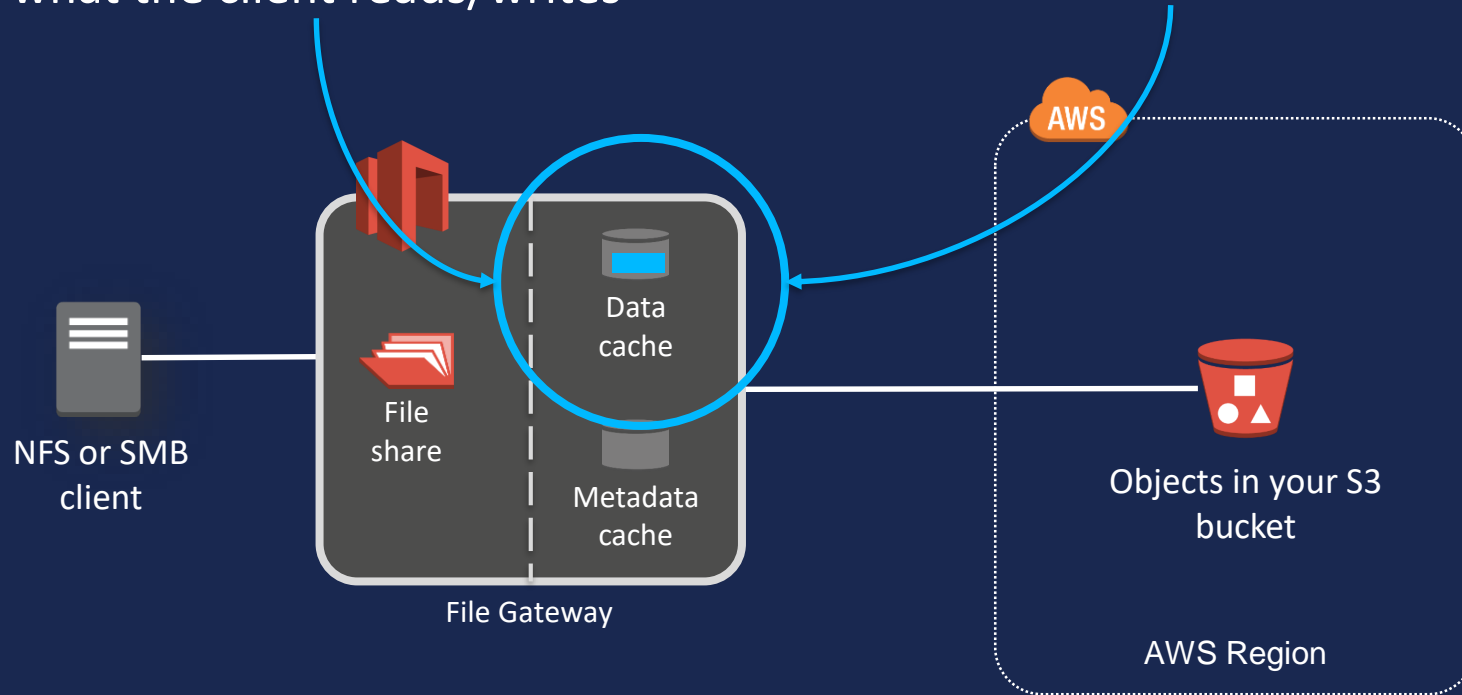
All of your data durably
stored in Amazon S3



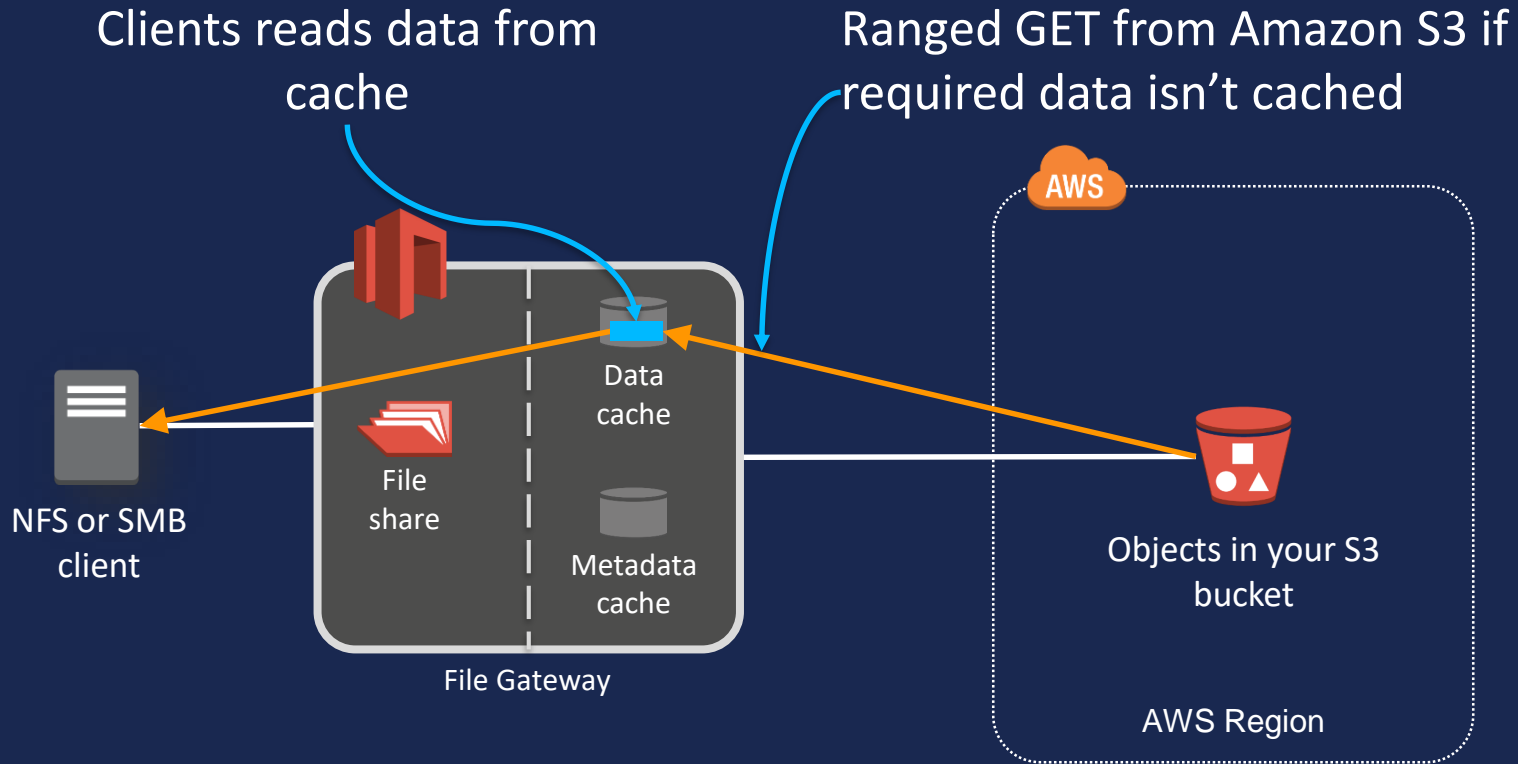
File Gateway: Data cache for file reads and writes

Parts of files cached based on what the client reads/writes

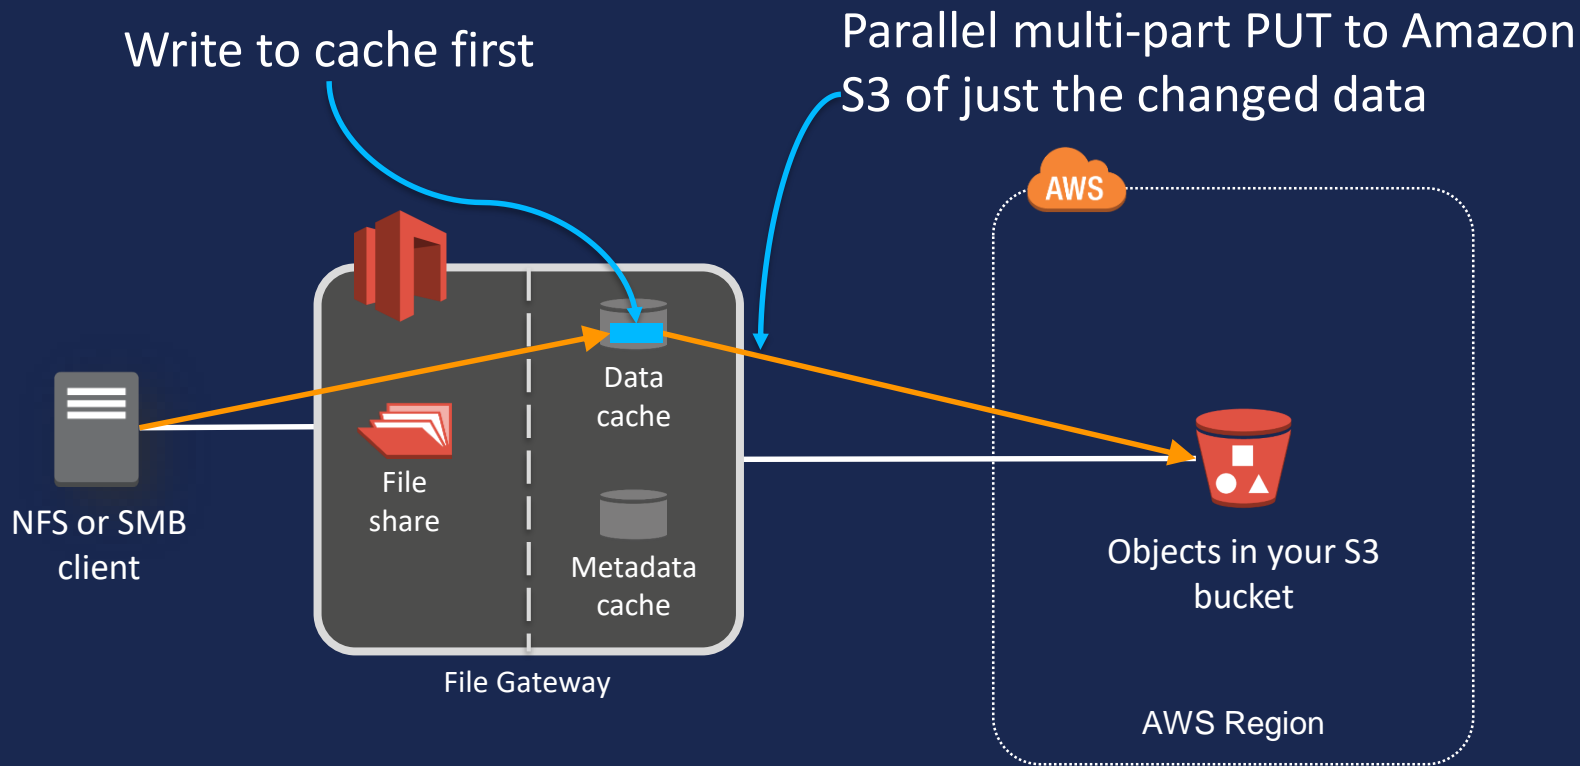
Data cached until space needed for more recently accessed files



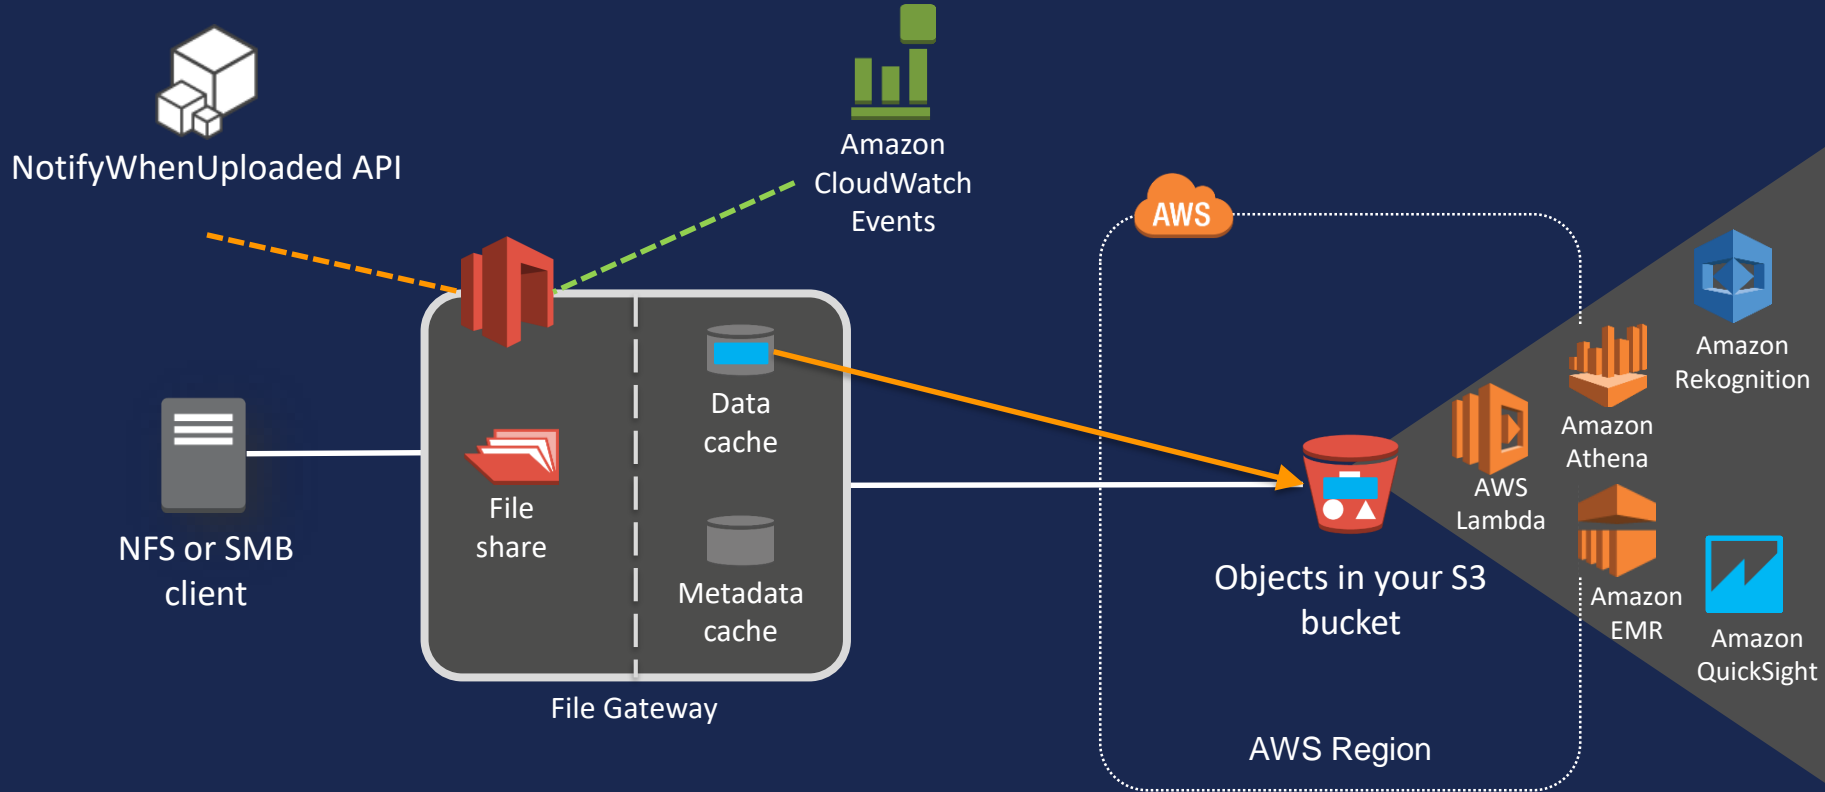
File Gateway: Data is read from the cache



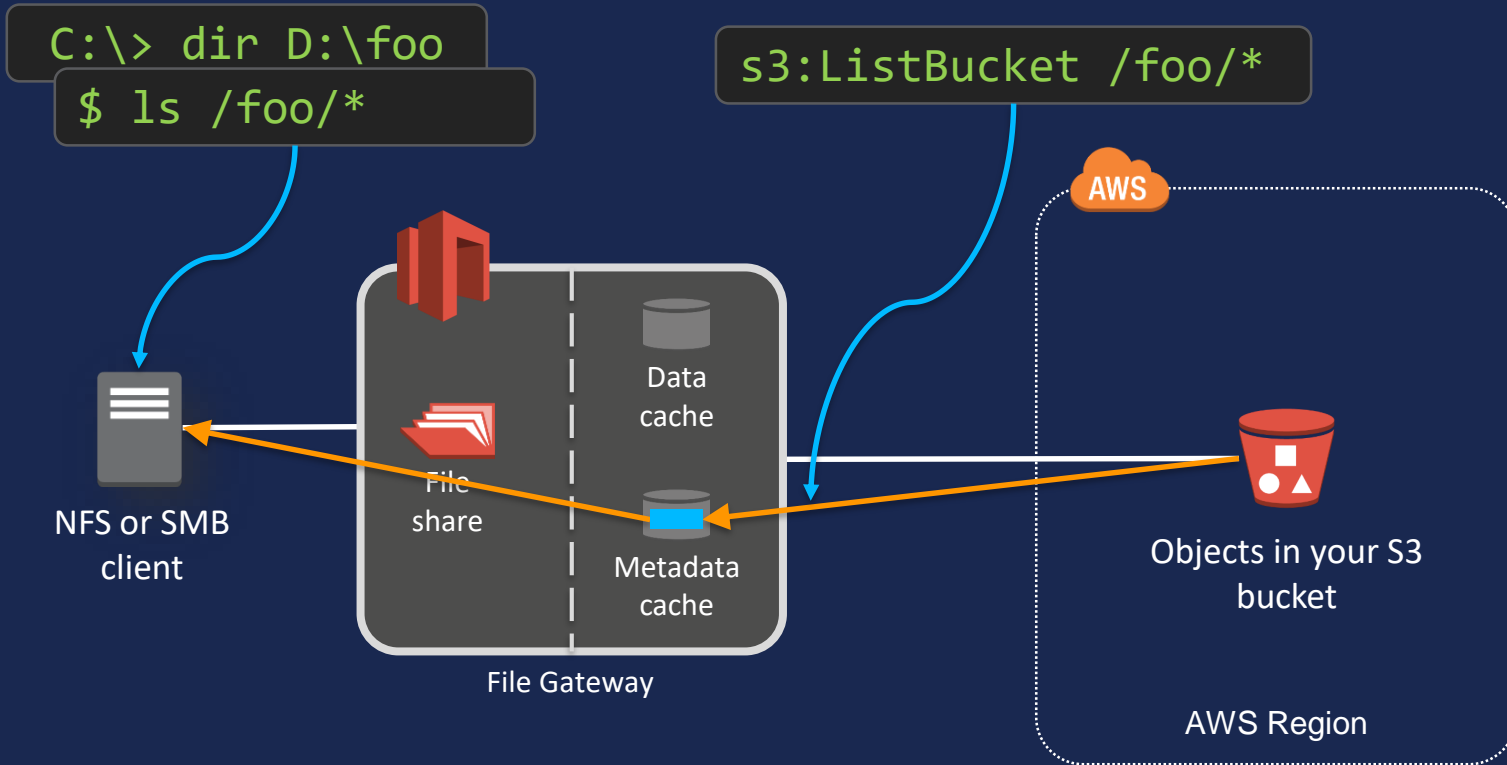
File Gateway: Data is written to cache then uploaded



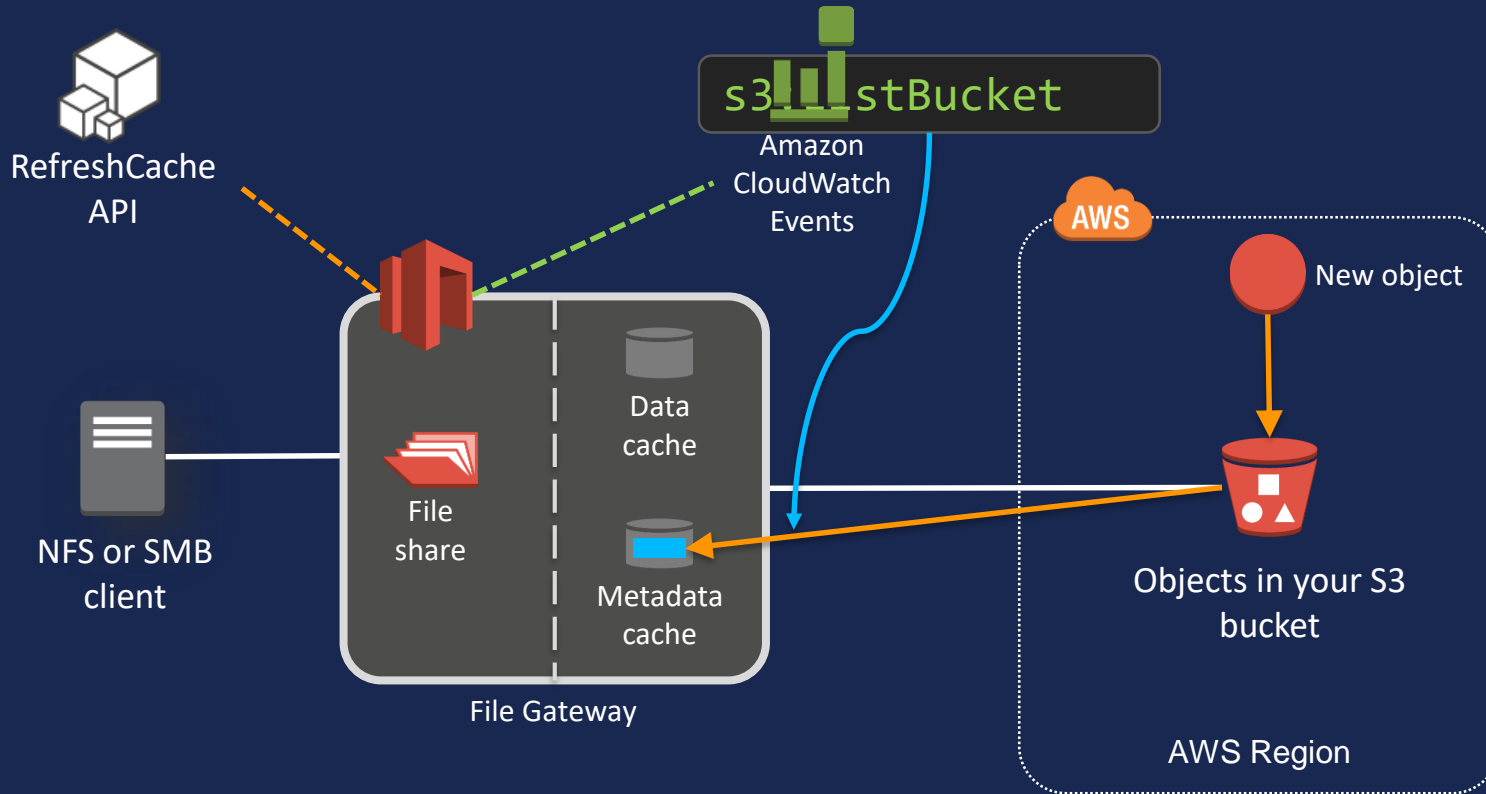
File Gateway: How do I know my data is uploaded?



File Gateway: Metadata cache for folder operations

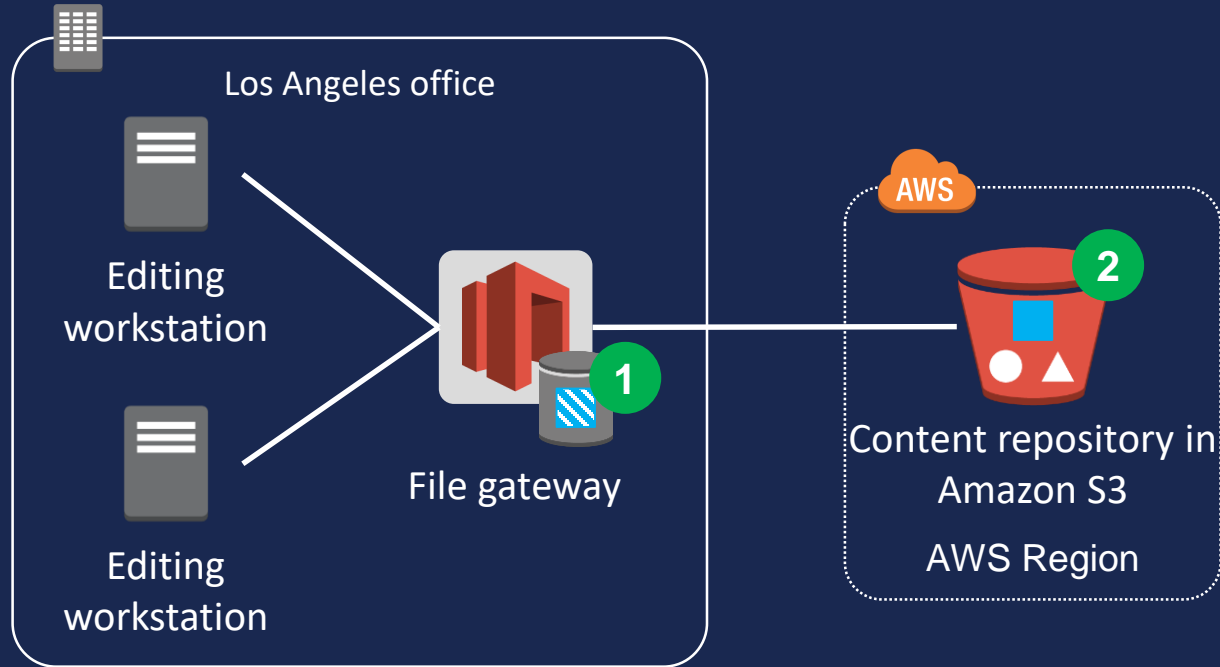


File Gateway: Updating the metadata cache



Hybrid File Storage Use Cases

Hybrid file use case: Low-latency content editing



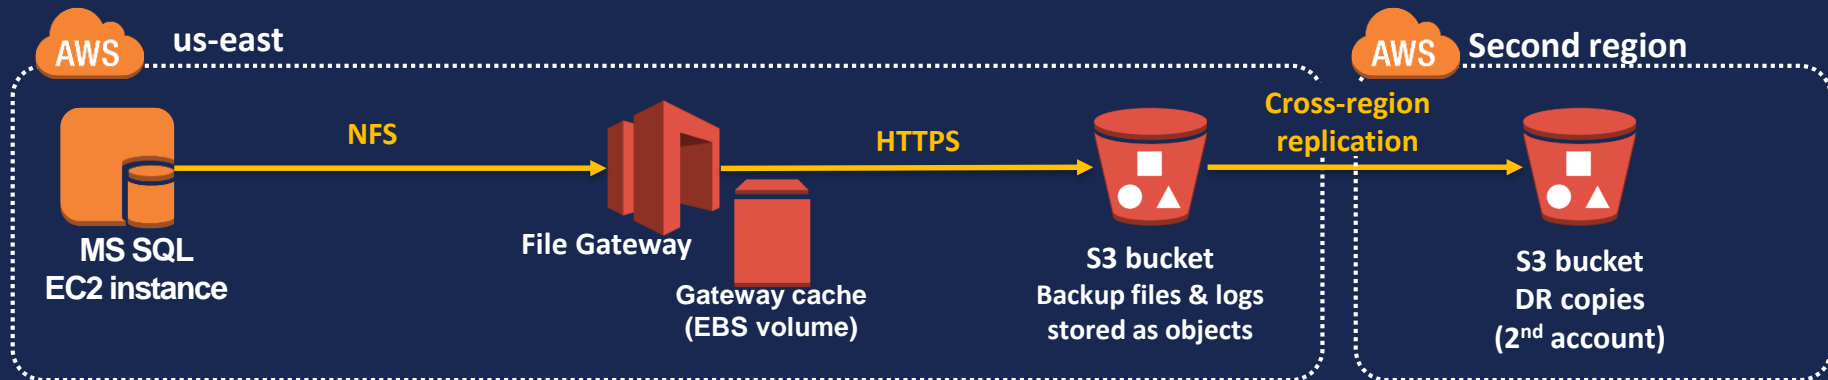
- 1 Low latency access for data actively being edited
- 2 Changes are all durably stored in master repository in Amazon S3

Hybrid file use case: Content distribution

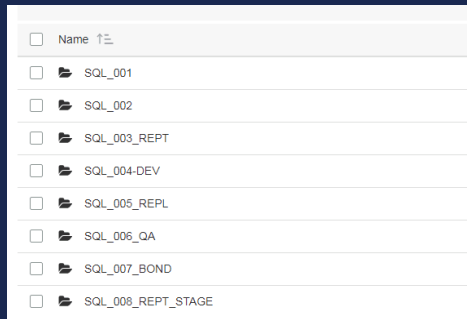


- 1 Application in Seattle writes files, which are uploaded to Amazon S3 by gateway
- 2 After refresh cache, files are visible to applications in Boston
Local cache improves access performance

Hybrid file use case: SQL Server backups



```
mount -o nolock  
[gateway-IP] : / [bucket-  
name] / [sql-instance]  
[WindowsDriveLetter] :
```

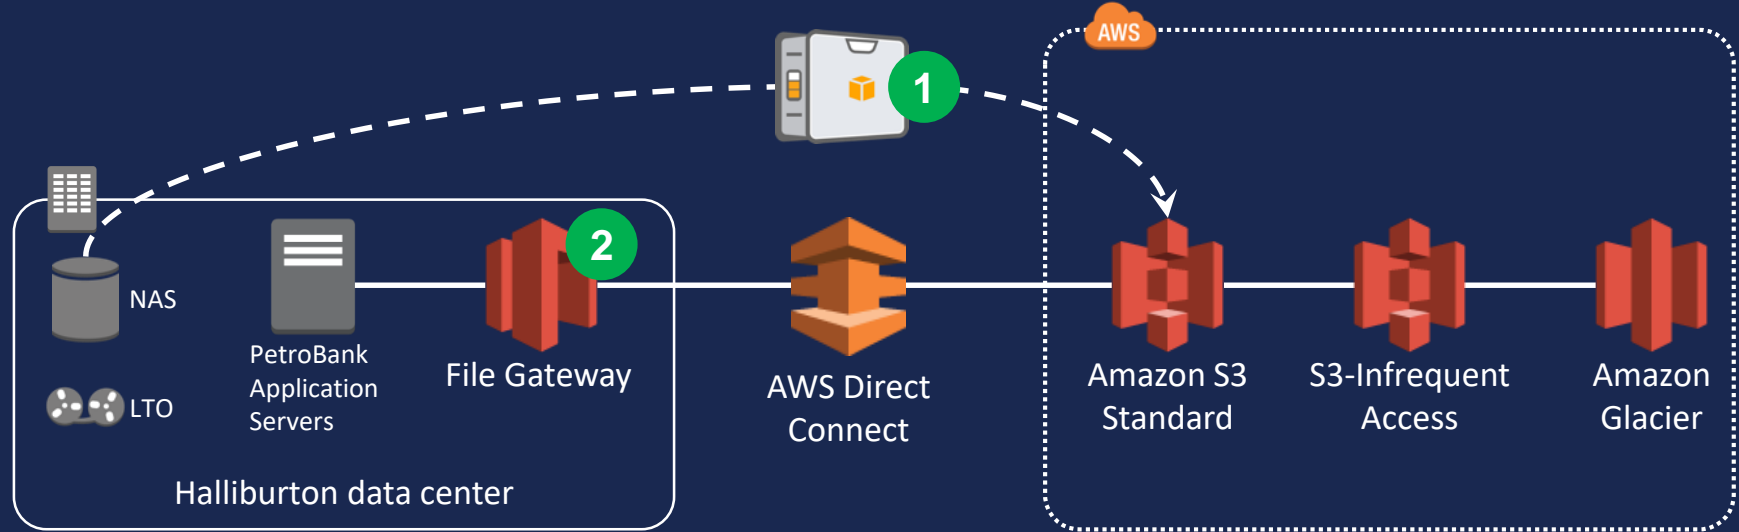


Hybrid file use case: Active archive migration

Cost-effective storage in AWS with local data access

HALLIBURTON

Landmark




- 1 Use Snowball to ship data from on-premises offline archives
- 2 Online access to all data through gateway
Minimal on-premises storage reduces cost
Time-to-data reduced by days or weeks

Hybrid volumes for recovery & migration with Volume Gateway and Amazon EBS snapshots

Volume Gateway

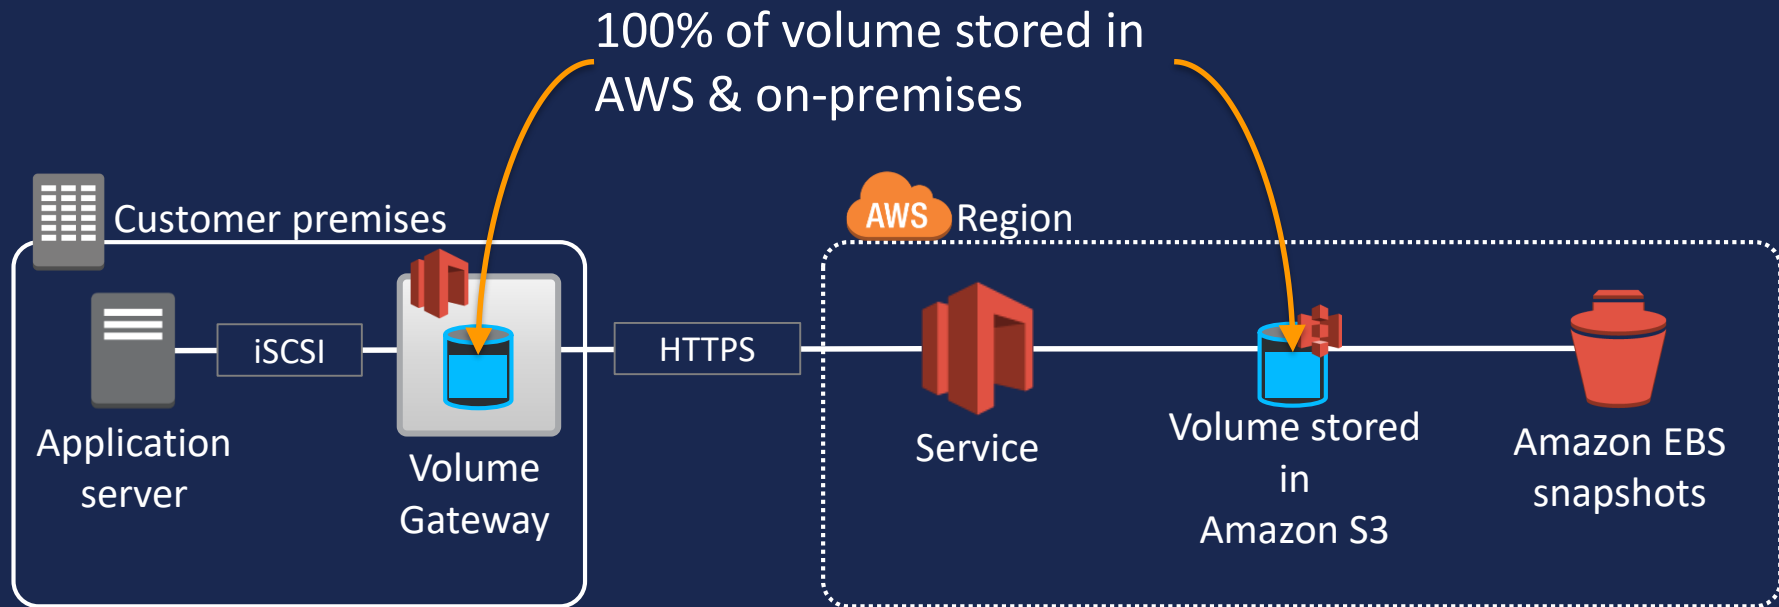
Cloud-backed block storage presented on-premises

- 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 (DR), data migrations



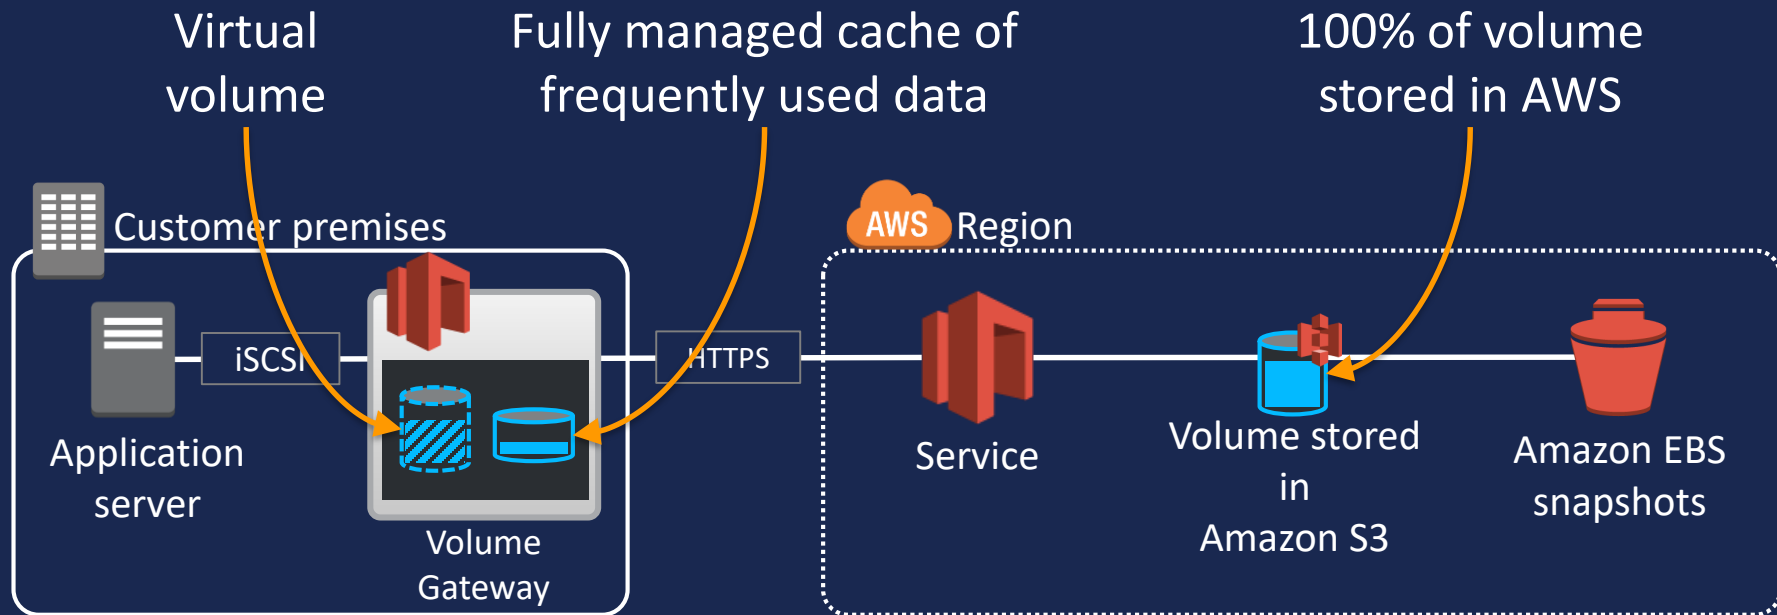
Volume Gateway: Stored mode

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



Volume Gateway: Snapshots and clones for data protection



EBS snapshots

Point-in-time backups of a stored or caches volumes

Created on-demand or on a configurable schedule

Restore either as an Amazon EBS or a Storage Gateway volume



Volume clones

Instant real-time copy of a cached volume

Represents current state of volume stored in AWS

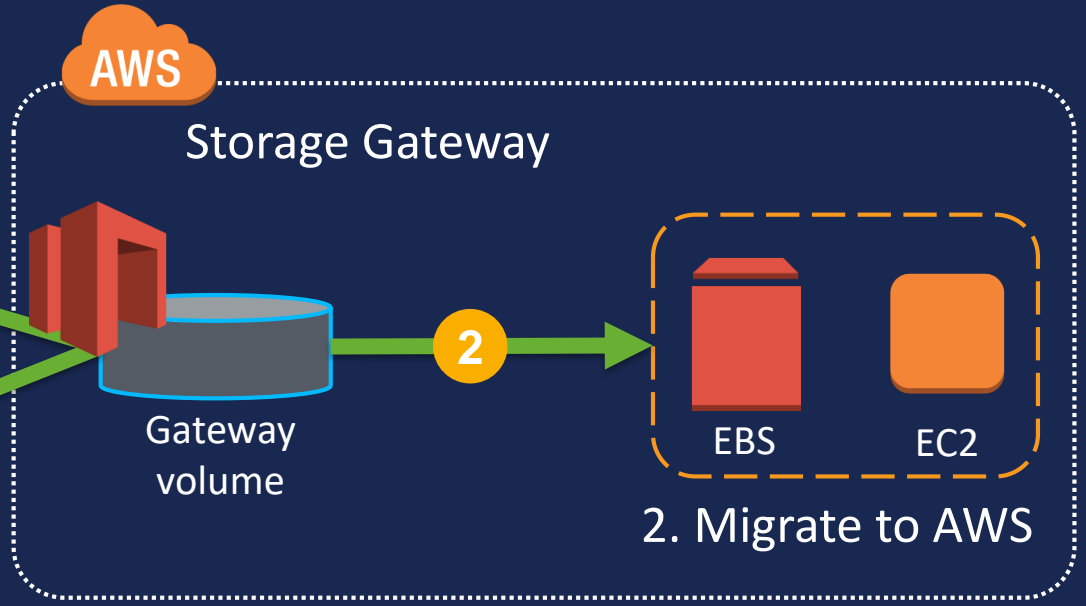
Restore as a Storage Gateway volume

Flexible data backup, recovery & migration use cases

1. Restore to your data center



1



Storage Gateway

AWS

Gateway volume

EBS

EC2

2. Migrate to AWS

3. Recover to a Disaster Recovery (DR) site



3

Hybrid use case: Protect & recover volumes in-cloud



Canada's largest biotech firm

“It made no sense to keep buying big disk siloes, especially as we opened up new global offices, and now we can recover in the cloud from a snapshot if we ever had to.”

Adam Leggett
IT manager

- Data sovereignty required local hot files & tape archives in 10 global offices
- **Volume Gateway** eliminated 50-hour backup windows and tape archive systems
- Cut on-premises; storage capex 40%; reduced RTO from 48 hours to 10 minutes
- Meets cloud strategy while retaining local ownership and data sovereignty
- Enabled data center exit in next 6–12 months



Migrate tape backup workflows to AWS 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

DELL EMC

Quest

VEEAM



Hewlett Packard
Enterprise

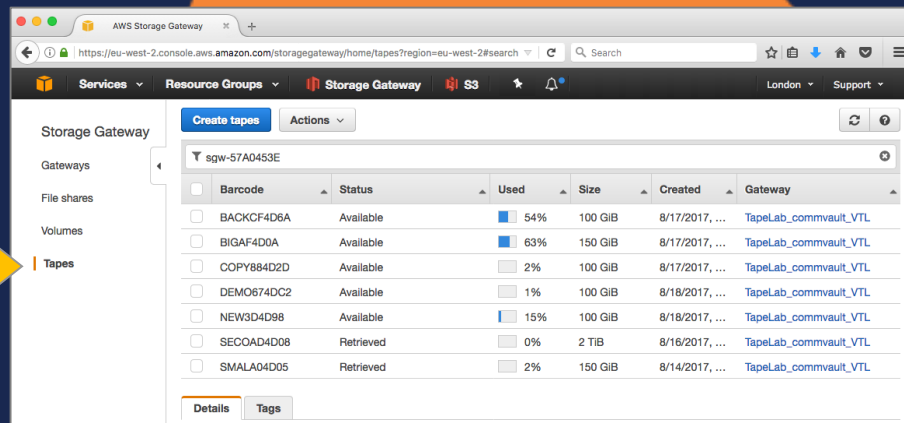
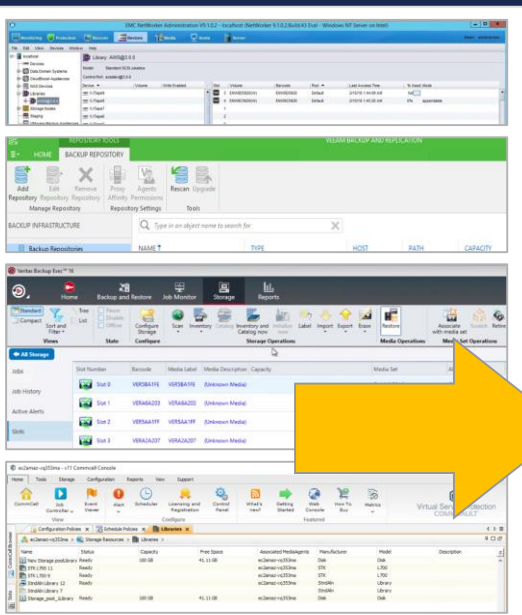
VERITAS

Microsoft
System Center
Data Protection Manager 2012

COMMVAULT

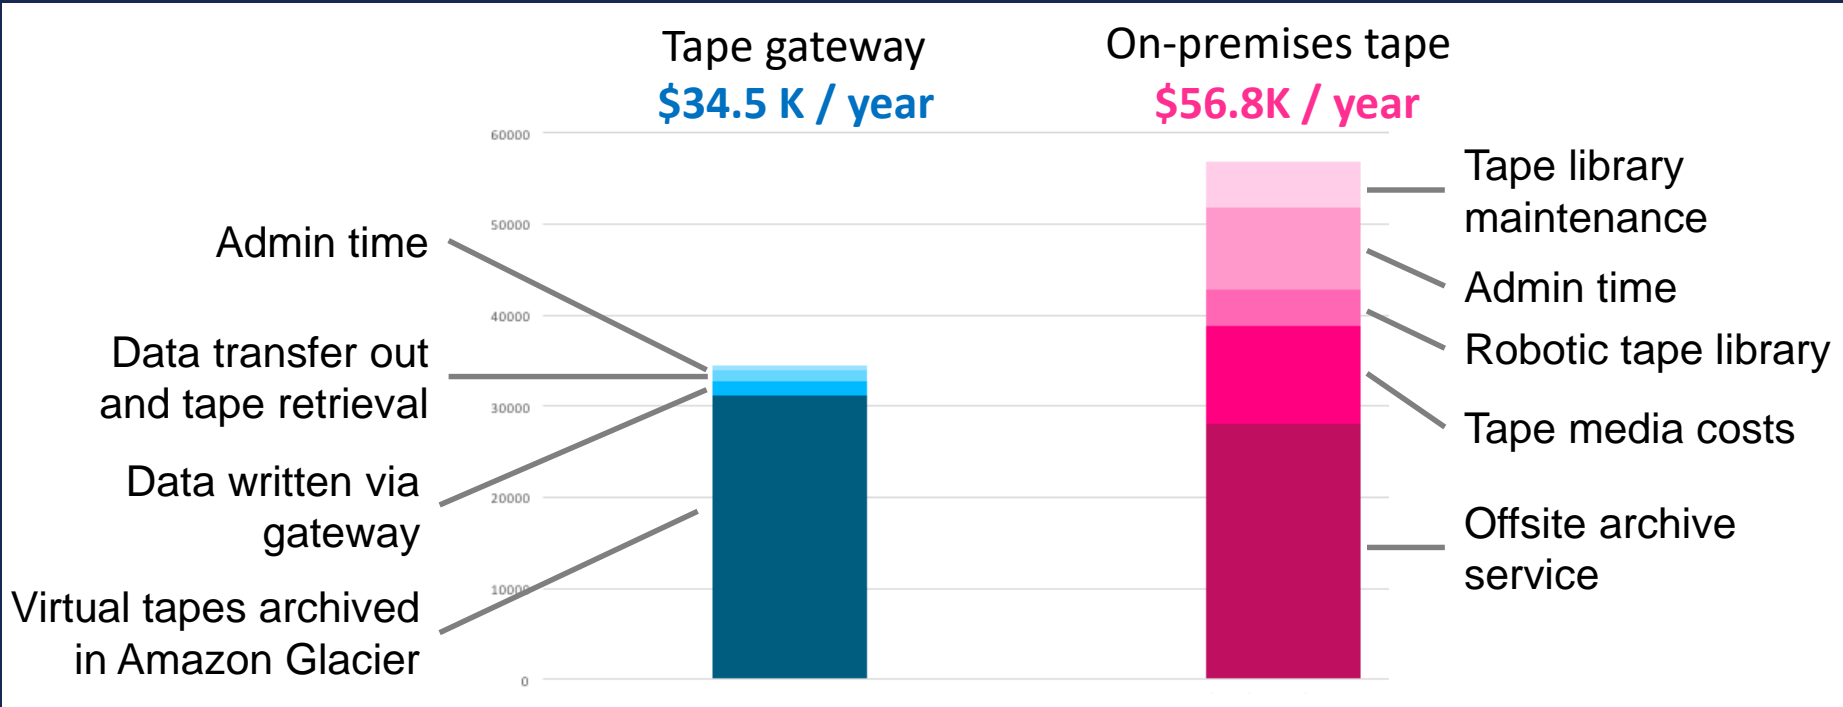
arcserve

NOVASTOR
Backup for the rest of us.

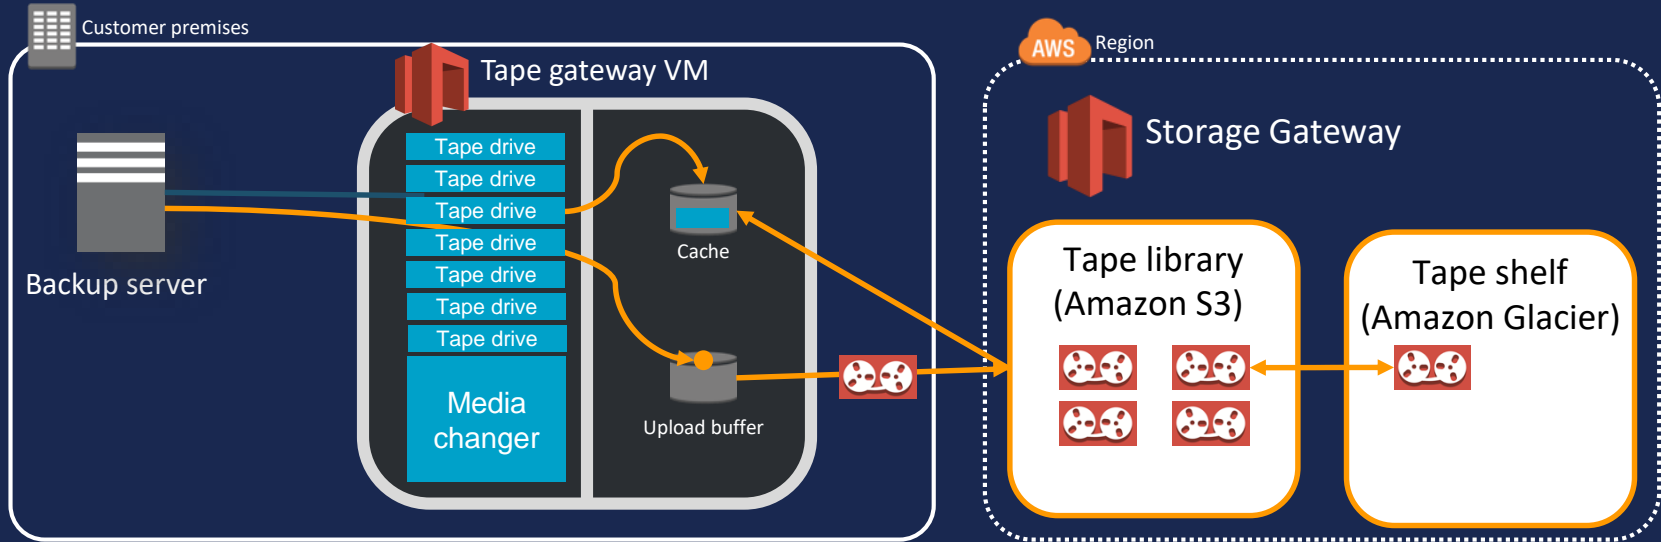


Tape gateway: Cost comparison, a customer example

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 written to Amazon 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

Backup, archive, and disaster recovery

Cost effective storage in AWS with local or cloud restore



“Using AWS Storage Gateway, we switched from physical to virtual tape backup simply by dropping the gateway’s virtual appliance into our existing **Veeam** workflow.

Setting it all up took three hours, at most.

We can now provision virtual tapes on AWS with the click of a button.”



Summary: AWS Storage Gateway enables you to ...

Reduce on-premises storage infrastructure with AWS

Move file, block, and tape storage to the cloud

Enable a range of backup and hybrid workloads



Questions?

Thank you.