



INNOVATE

DATA EDITION

19 August 2021

Migration myth busting on moving your data to AWS

Naim Mucaj

Solutions Architect

Amazon Web Services



Common questions & comments

Do I need a dedicated network connection?

It's too much data, I don't have enough bandwidth

Which transfer method do we choose?

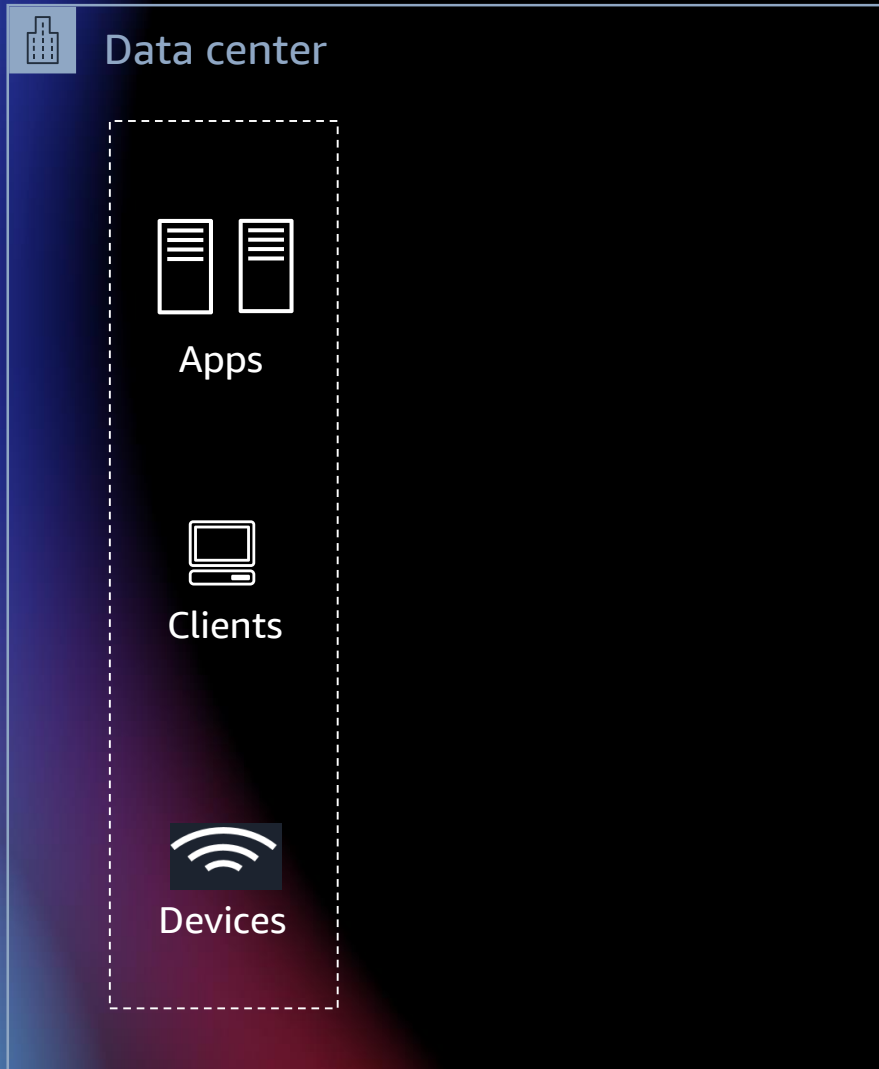
We don't have connectivity at those sites

It will take too long to transfer


Does AWS have performance or cost optimized storage?

Will I have to use DIY tools and manage the complexity?

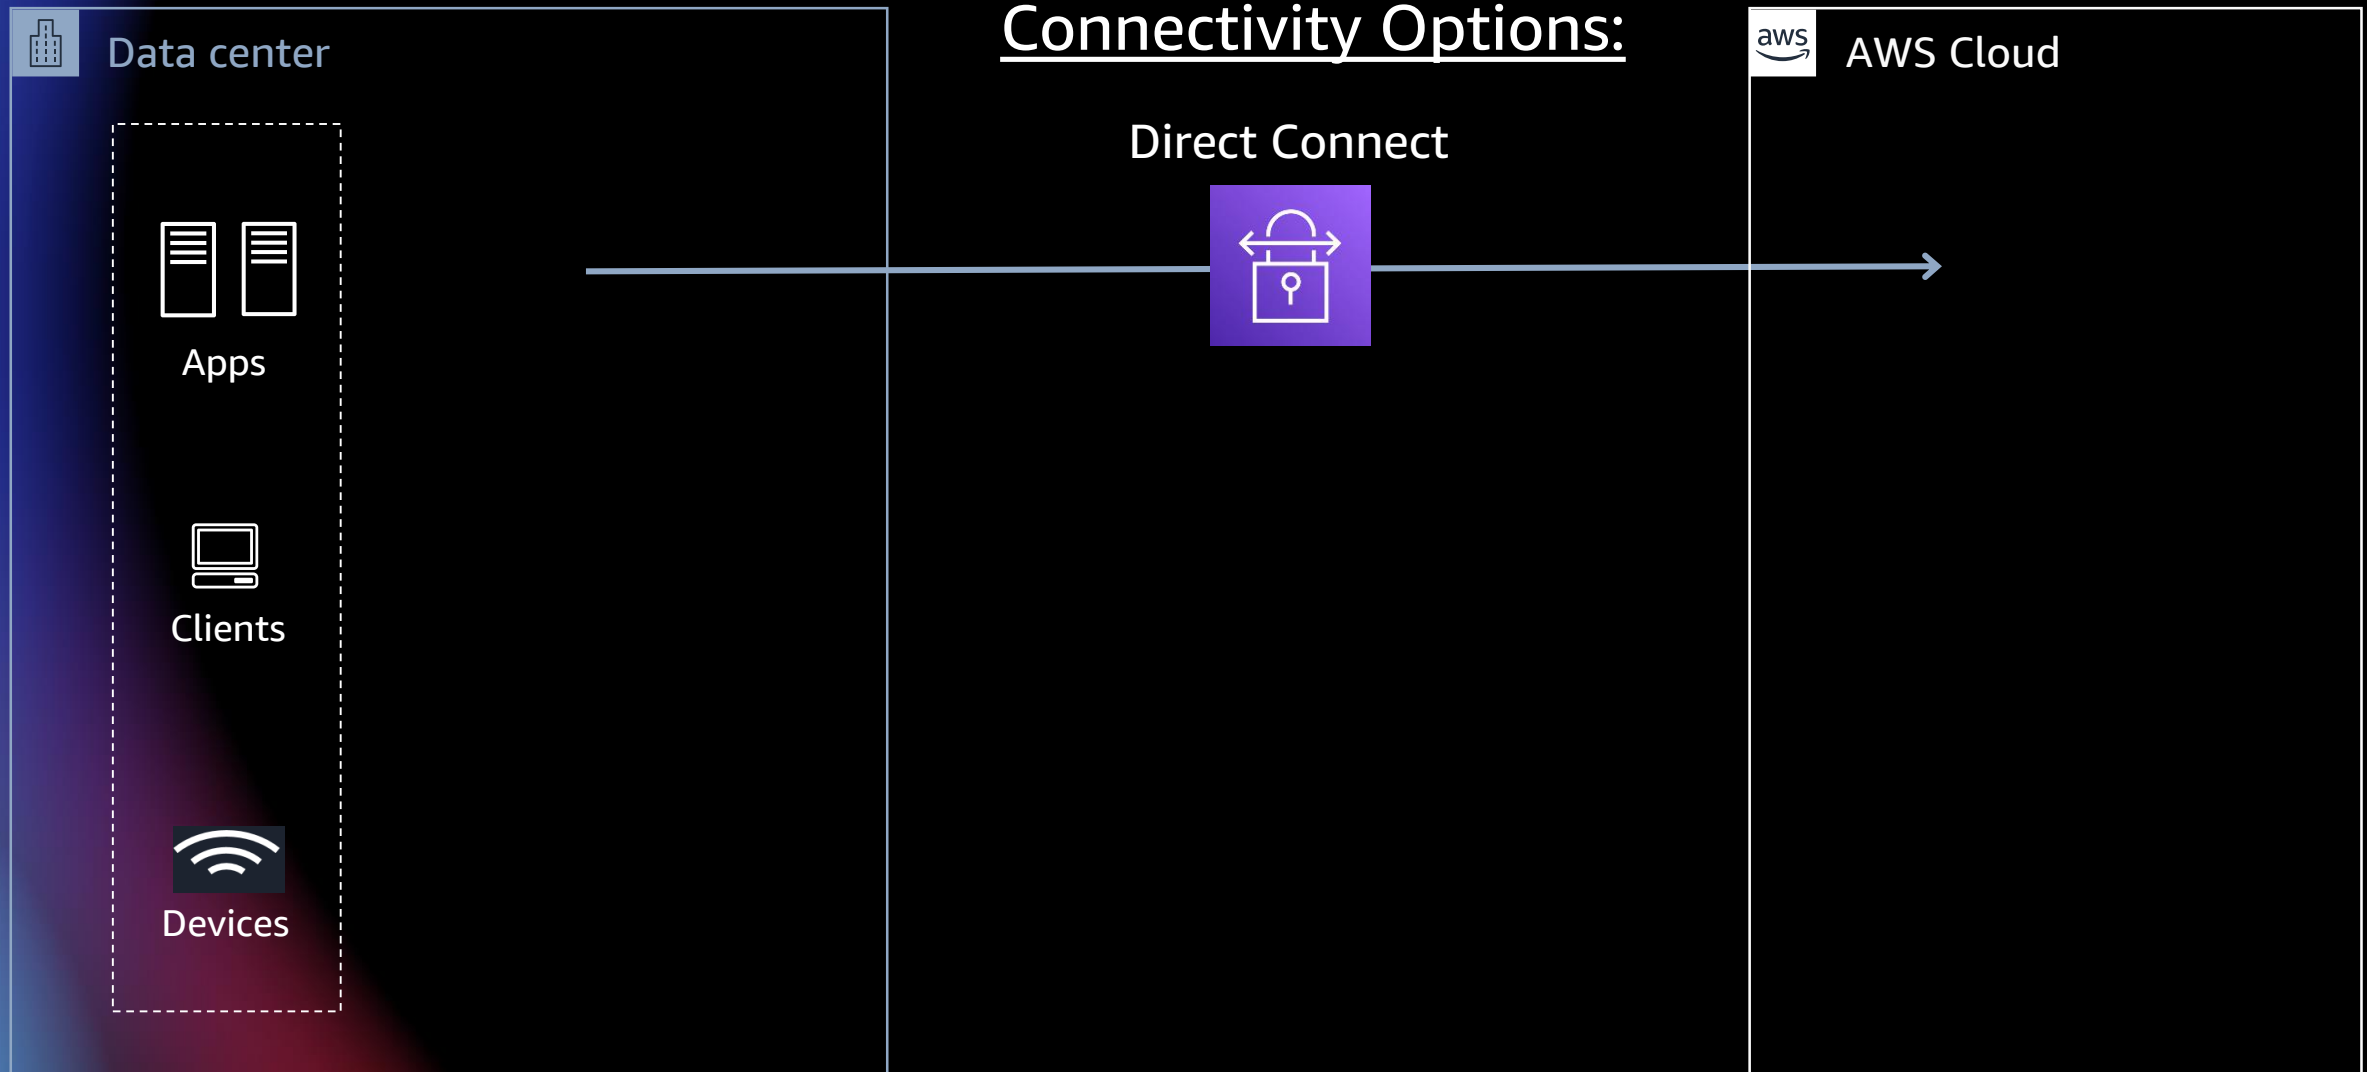
Do you need a dedicated network connection?



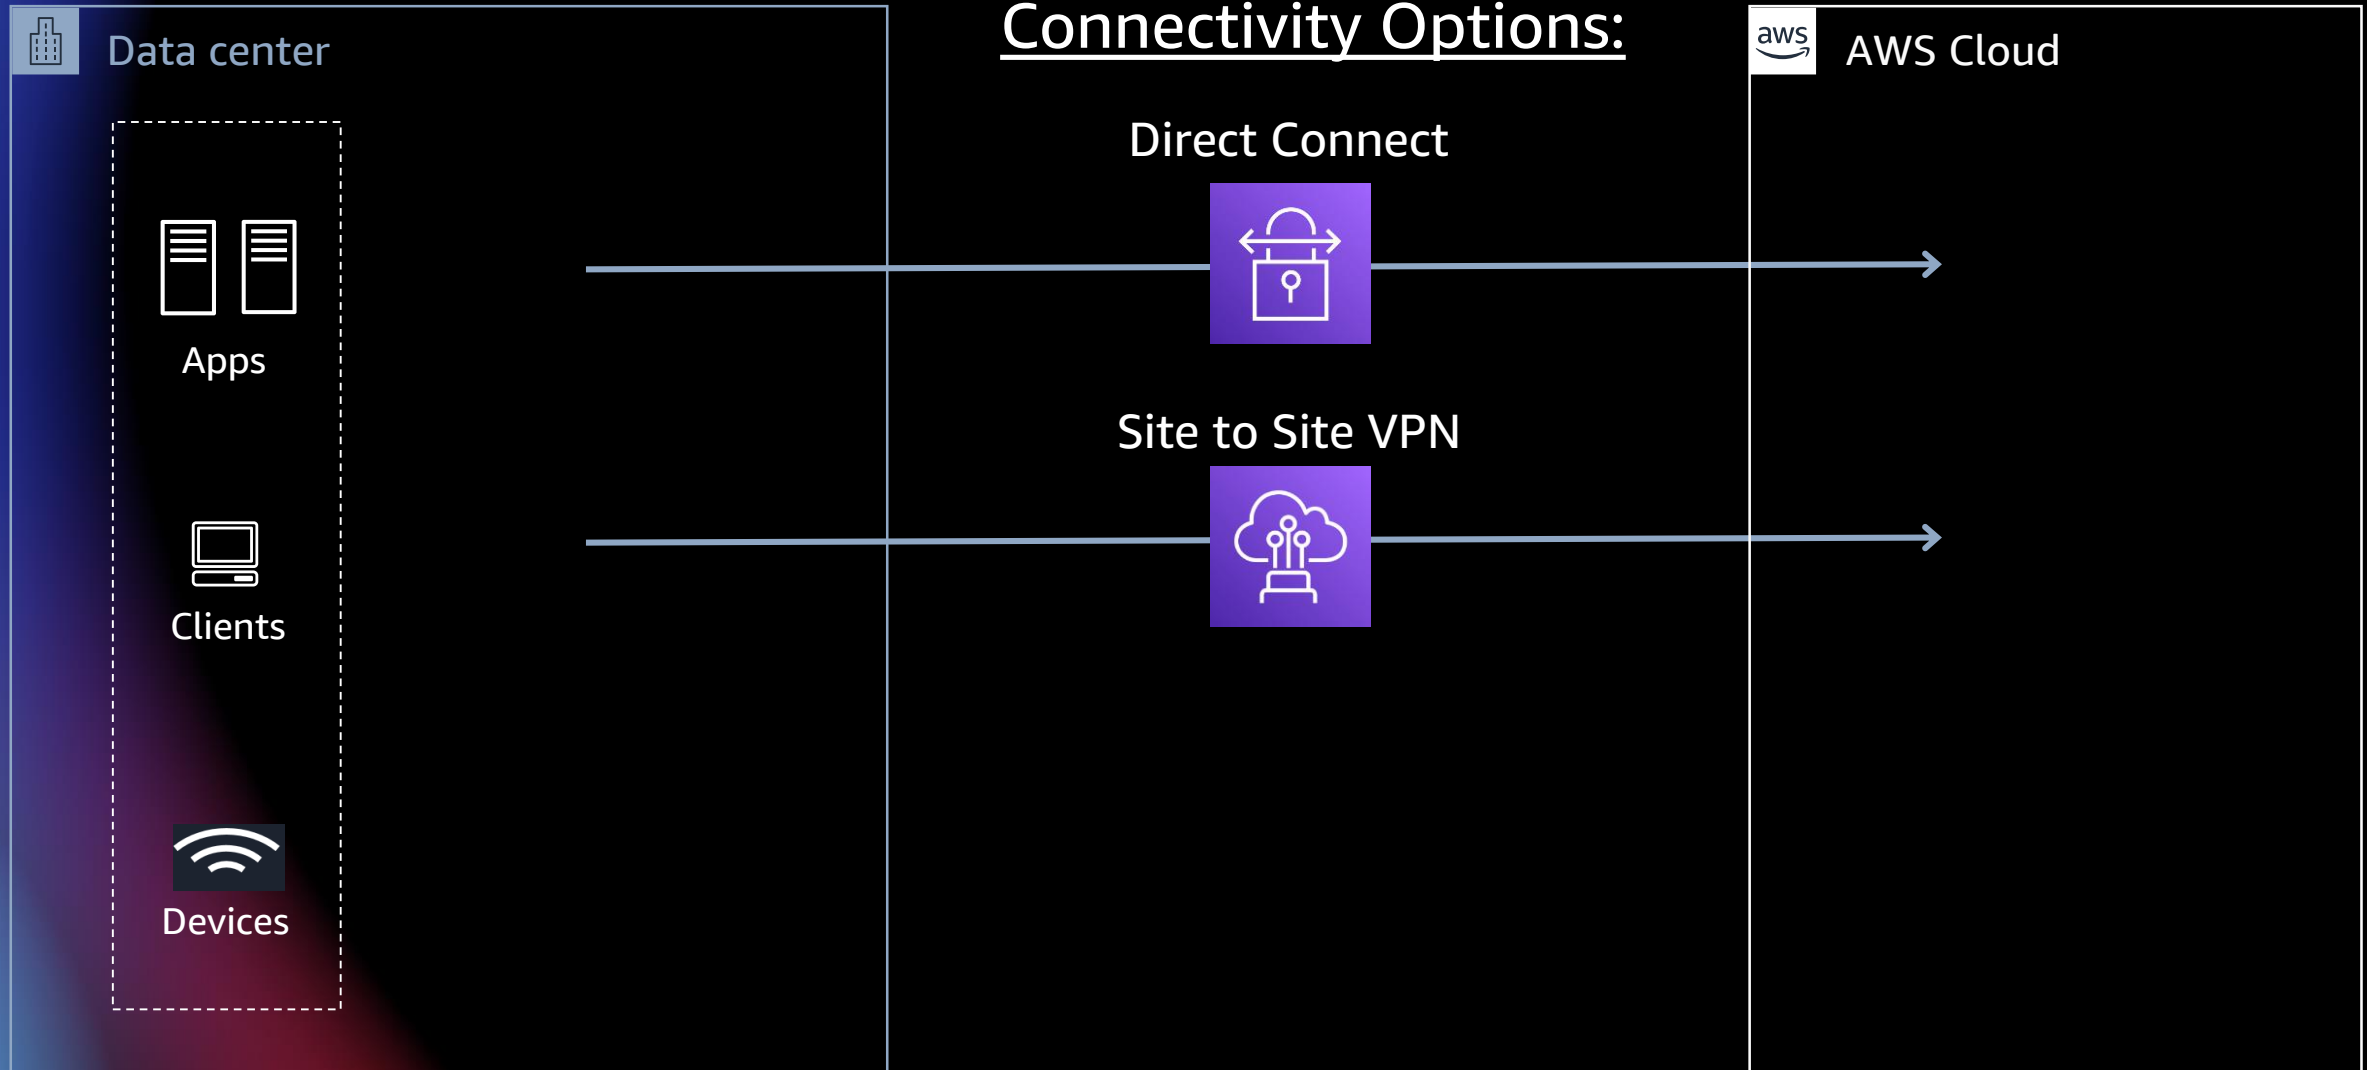
Connectivity Options:



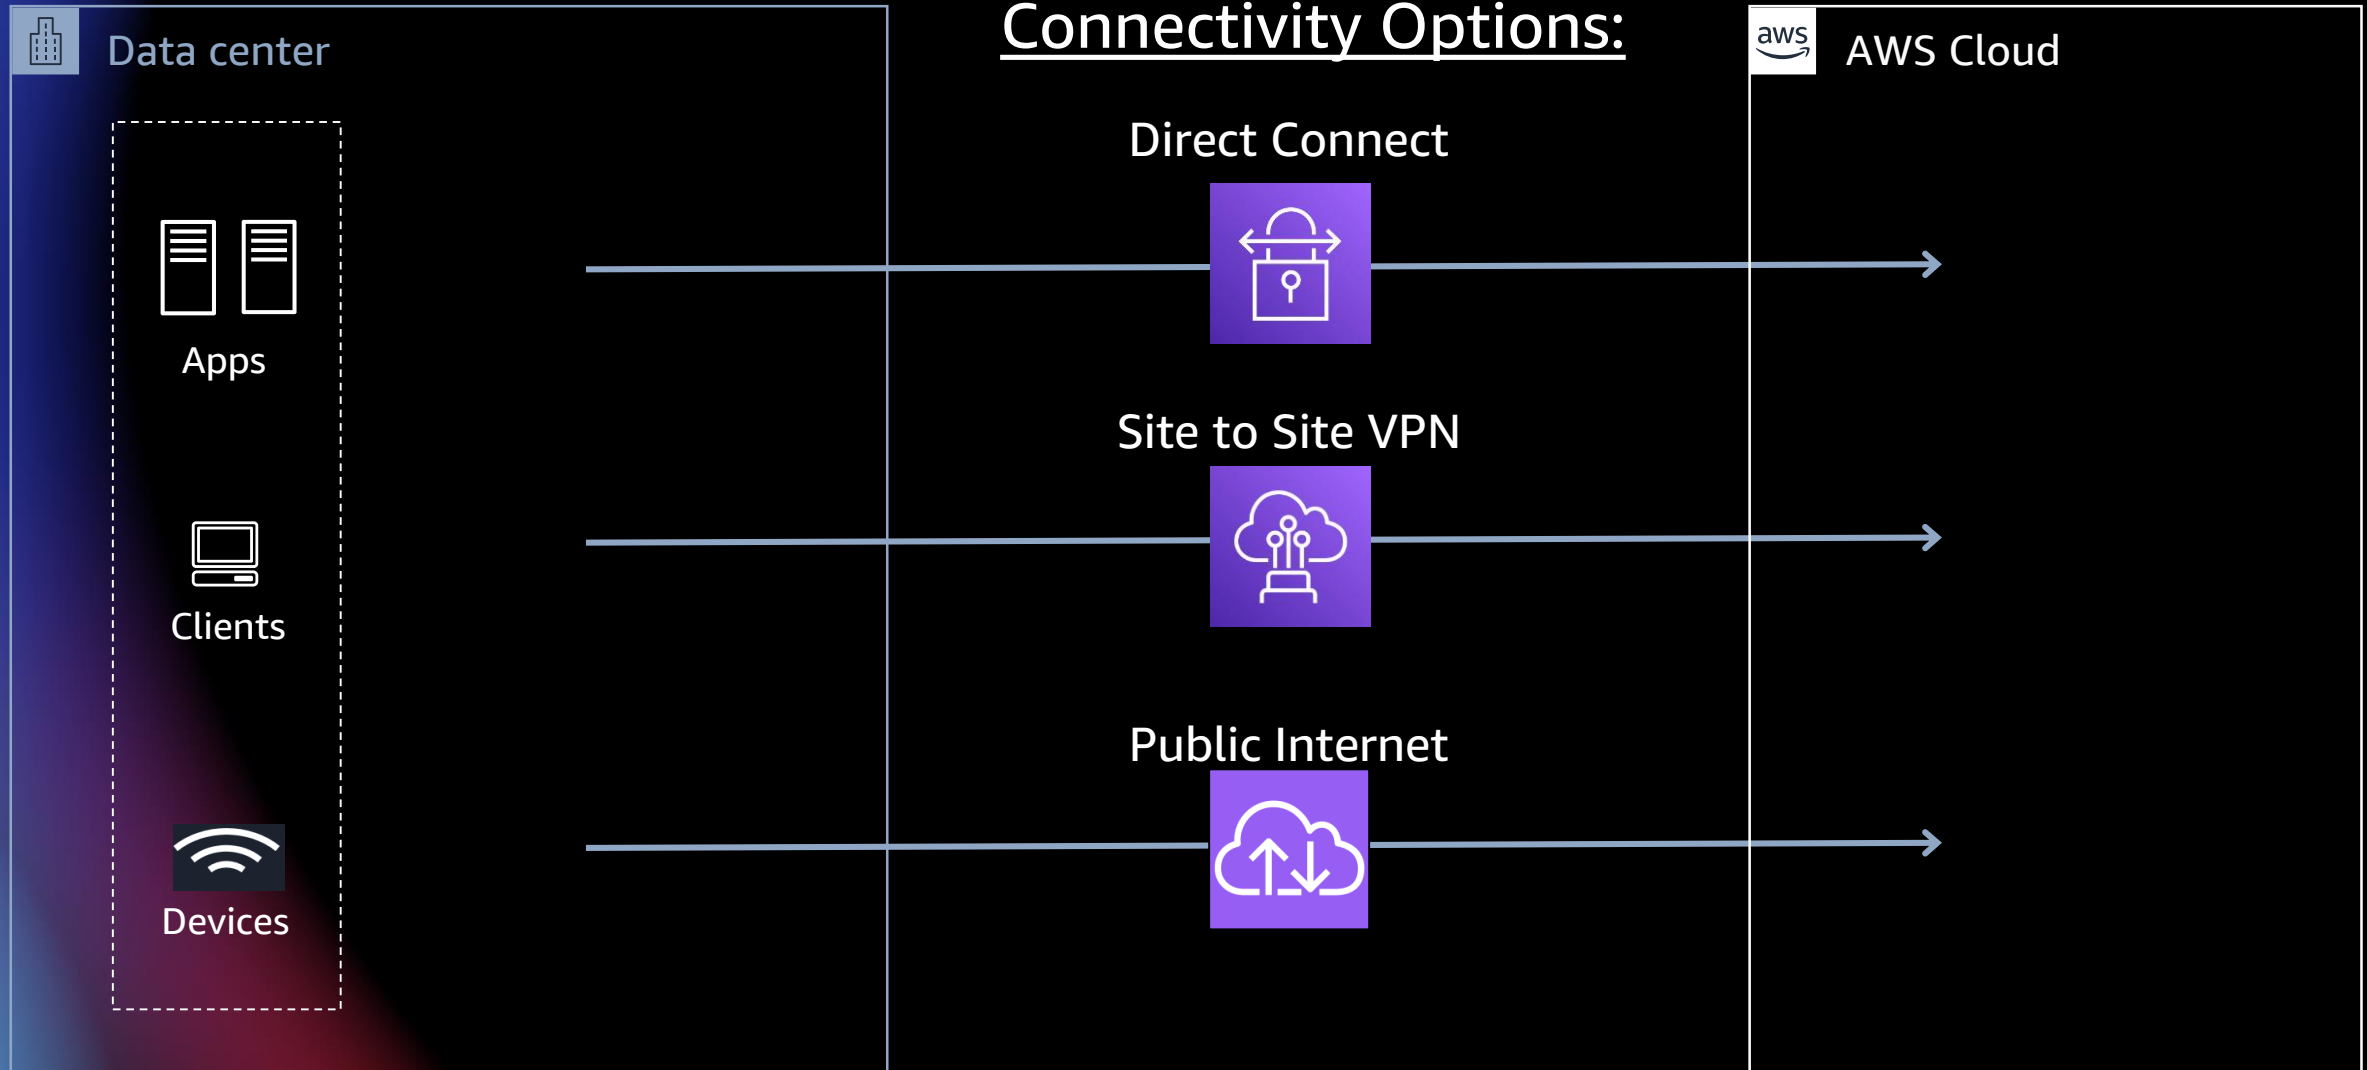
Do you need a dedicated network connection?



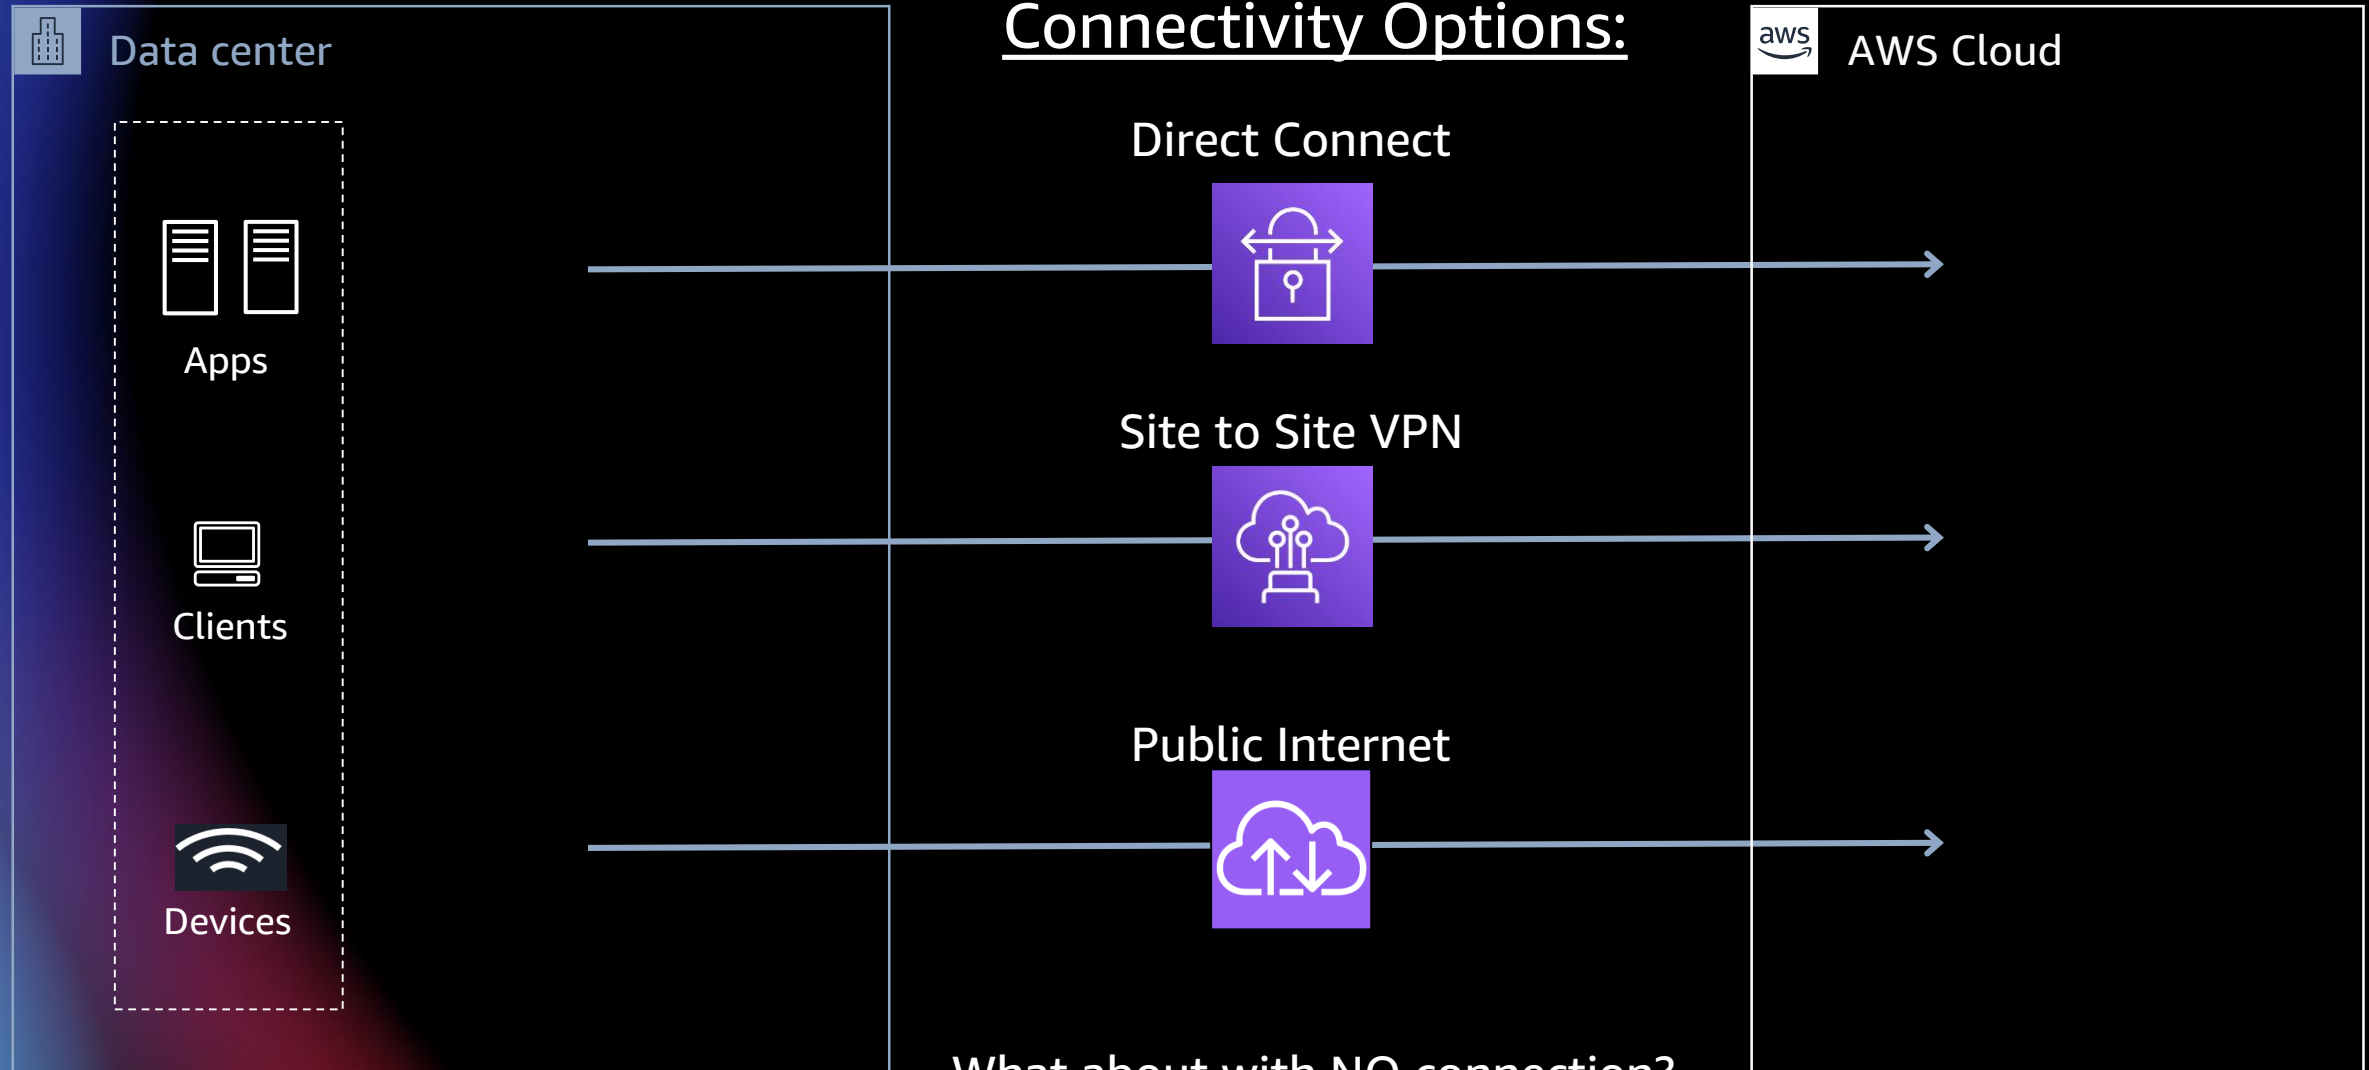
Do you need a dedicated network connection?



Do you need a dedicated network connection?




Do you need a dedicated network connection?



What about with NO connection?

Data Movement: AWS Snow Family



AWS Snowcone



AWS Snowball Edge



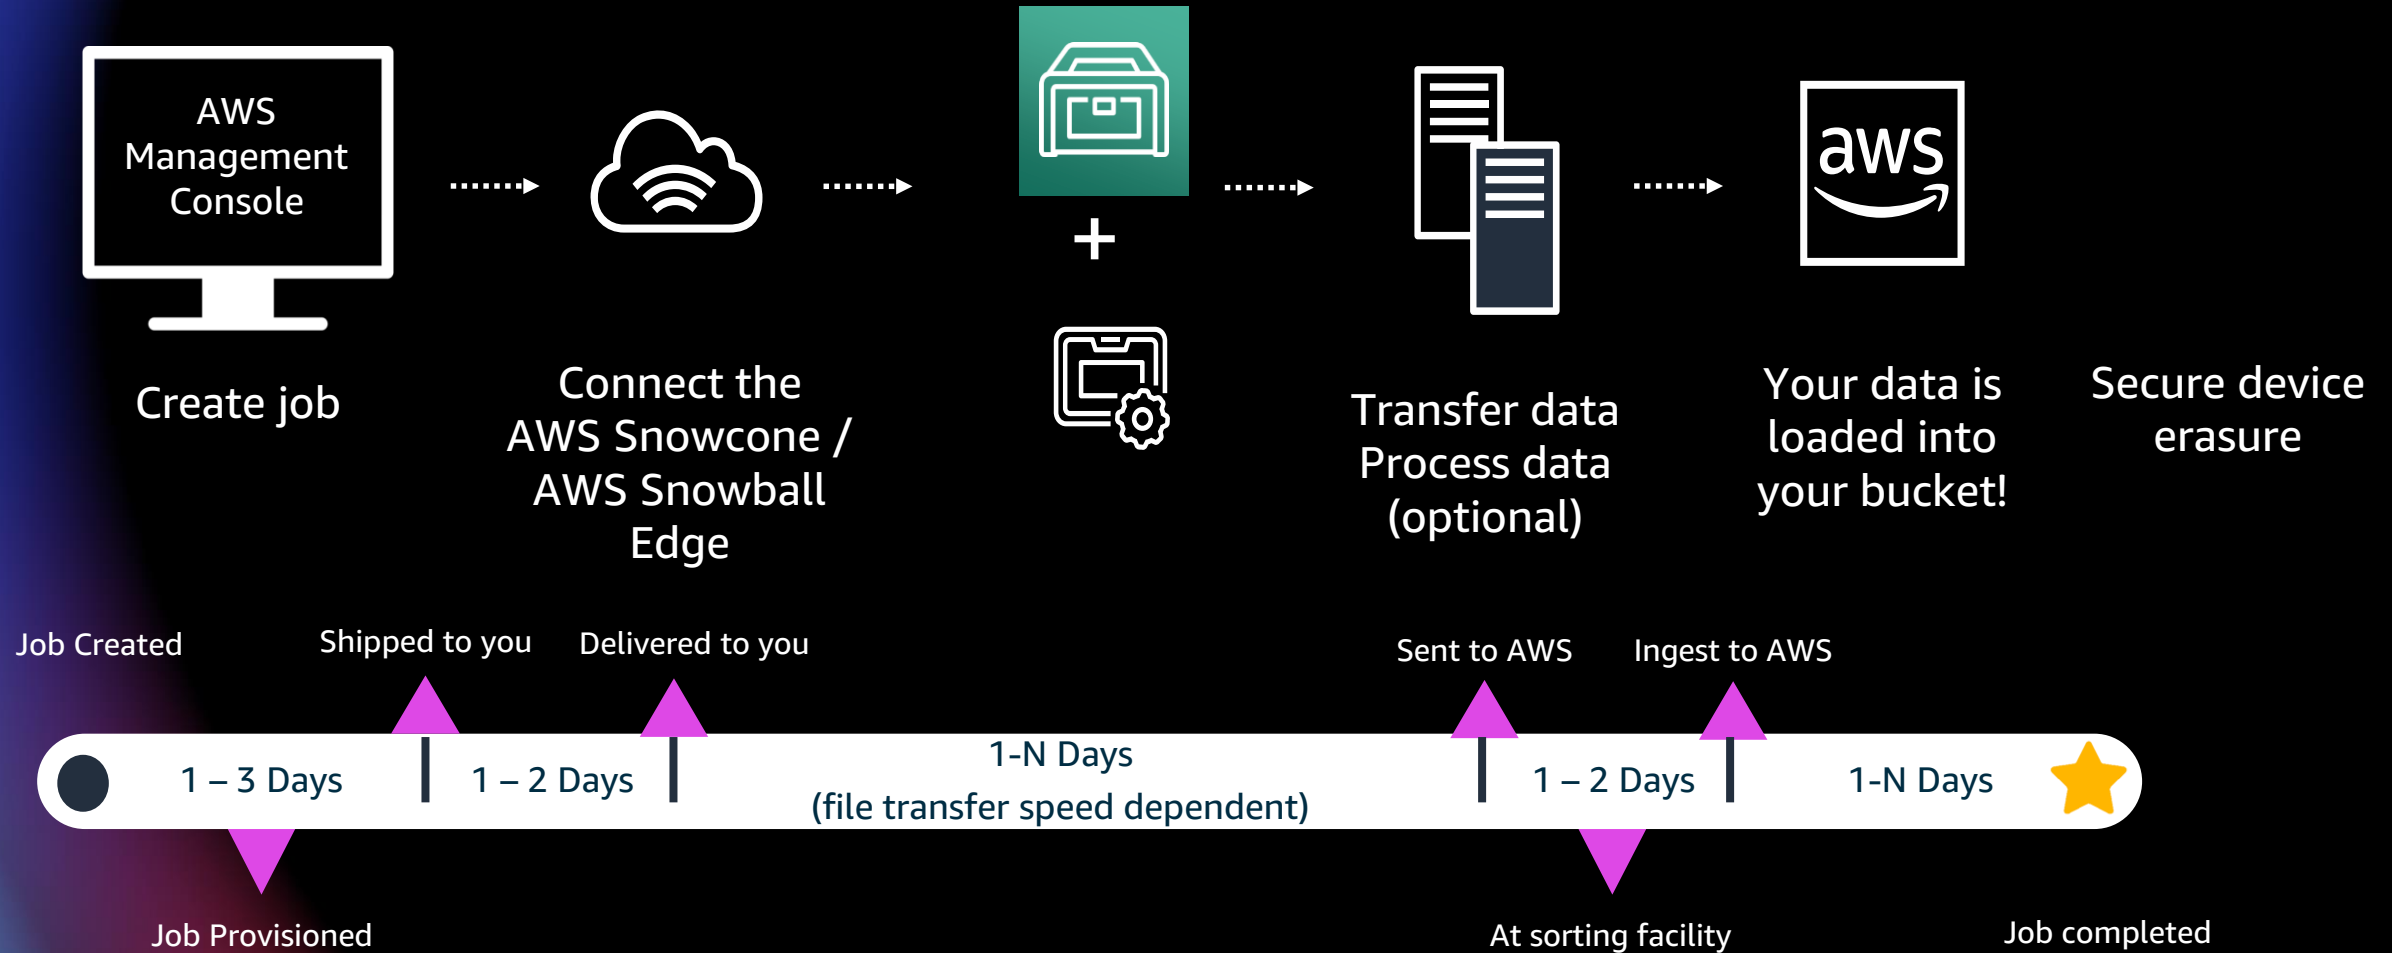
AWS Snowmobile

AWS Snow Family for data collection & data movement



	AWS Snowcone	AWS Snowball Edge Storage Optimized	AWS Snowmobile
Migration size	Up to 24 TB, online and offline	Up to petabytes, offline	Up to exabytes offline
Form factor	Rugged 8.5 G impact cases that are rain and dust resistant, E Ink label for shipping automation		45-foot container, scheduled delivery
Security	256-bit encryption, tamper detection		Encryption, security staff, GPS tracking, video surveillance, alarms
Storage capacity	8 TB usable	80 TB usable	<100 PB
DataSync agent	Pre-installed	-	-
Compute	2 vCPU, 4 GB RAM usable	40 vCPU, 80 GB RAM, 1 TB SSD usable	-
Onboard computing options	AWS IoT Greengrass functions Amazon Elastic Compute Cloud (Amazon EC2) AMIs		
Wireless	Wi-Fi	-	-
Portable or Mobile use	Battery based operation	-	-
Clustering	-	Up to 15 nodes	-

How data gets imported to AWS with Snow Family



Challenge

LotteON provides 39 million members with a unified storefront to 7 Lotte-owned stores.

Lotte needed to transform and transfer 140 million files consisting of product images and media assets.

Solution

Lotte used AWS Snowball Edge for secure, and fast offline data transfer.

Lotte did not have to provision additional network capacity.

AWS Snowball Edge enabled Lotte to meet its timeline to support a consistent shopping experience of 39 million Lotte members.

Benefits

The data was leveraged by 70+ AWS services, including Amazon S3, Amazon EKS, Amazon DynamoDB, AWS Lambda, Amazon Elasticsearch, Amazon Redshift, Amazon EMR, and Amazon Athena.

LOTTE ON

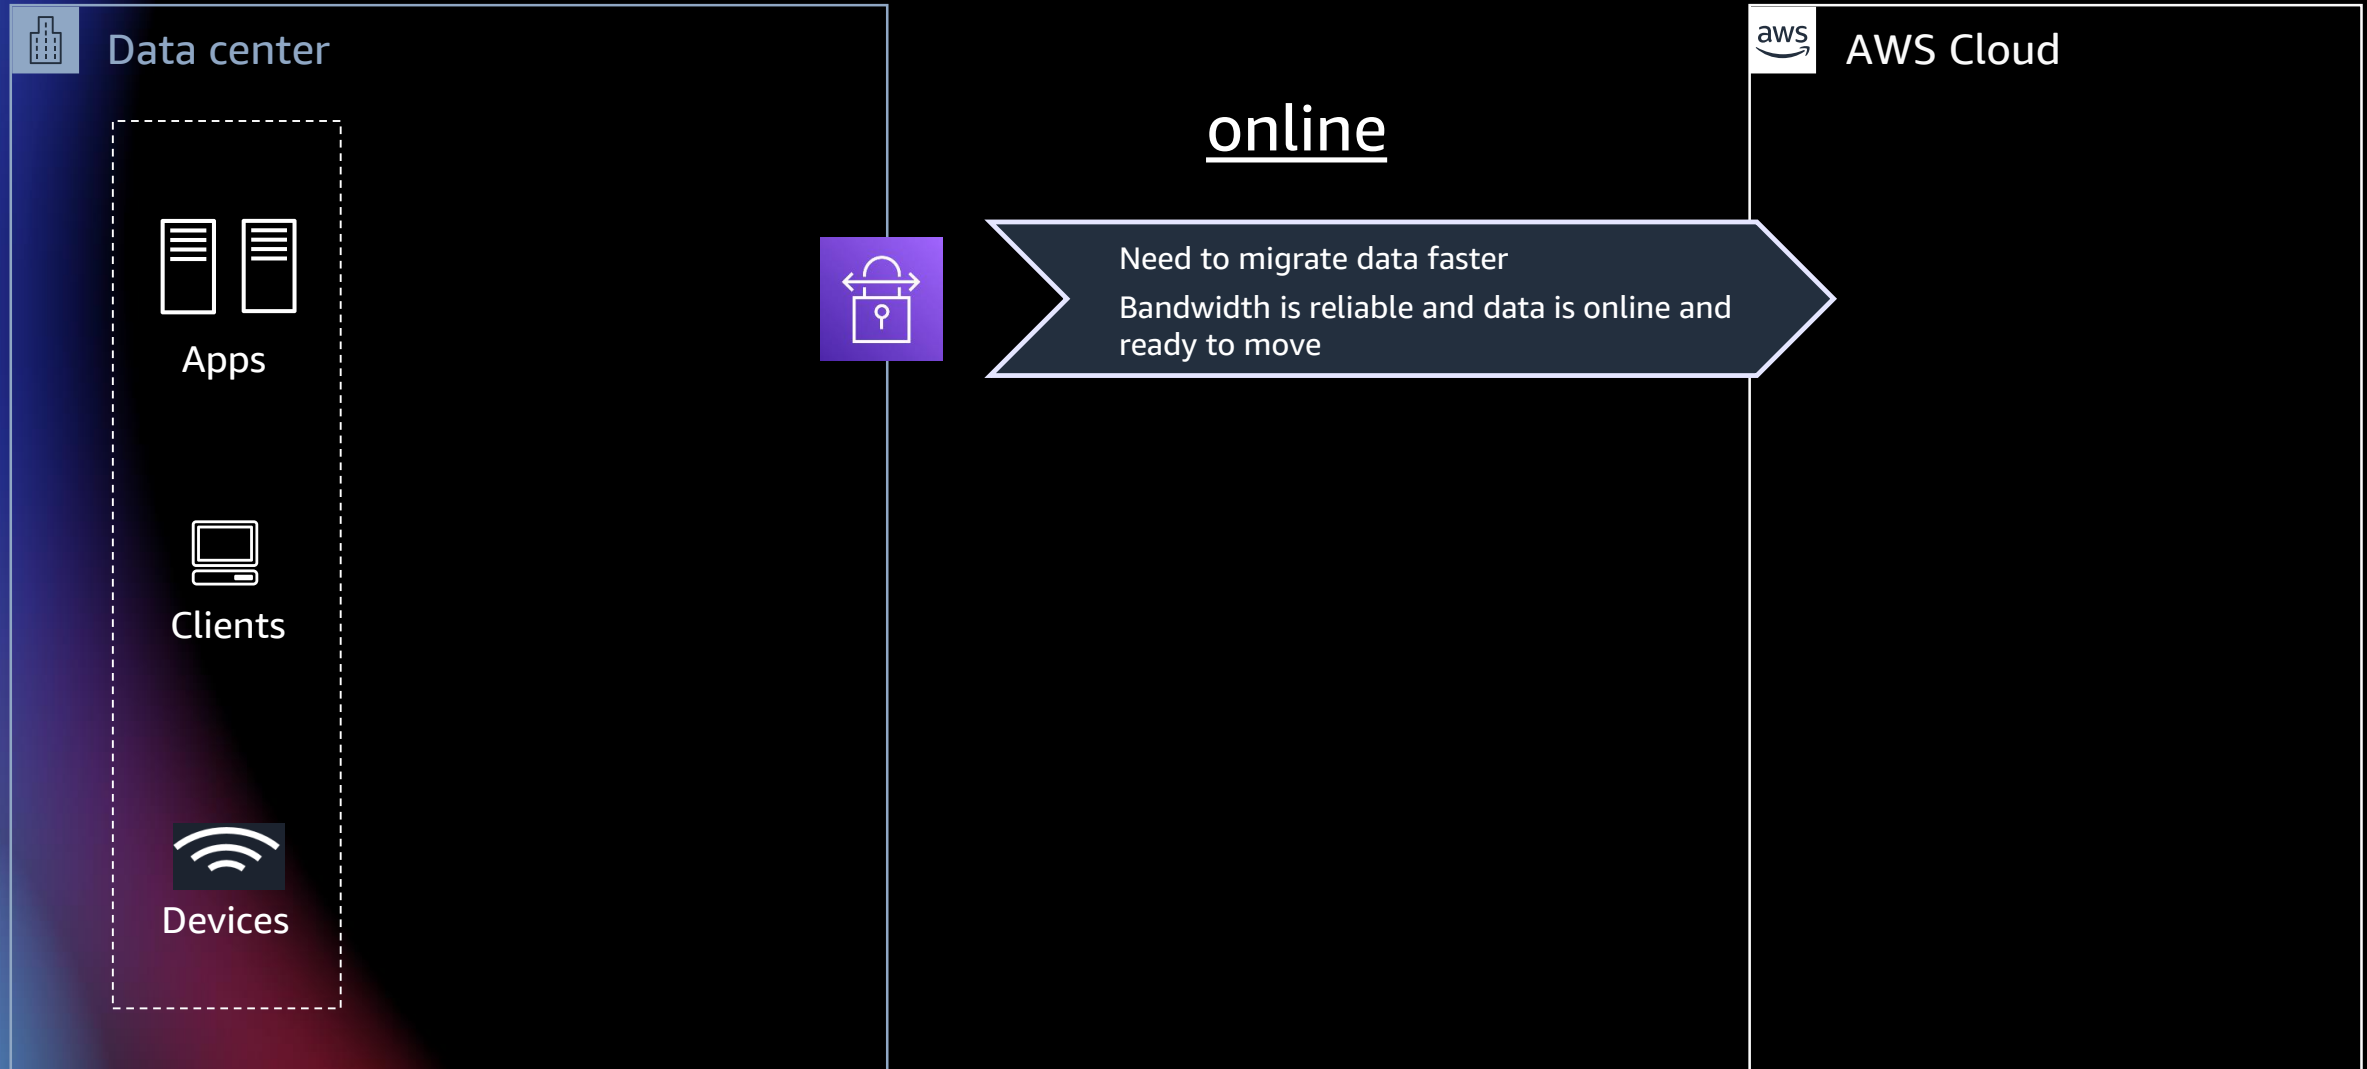
“With AWS Snowball, LotteON was quickly and reliably able to transfer 140 million product images and assets by combining 7 individual web stores within 2 weeks. Snowball helps meet the tight project launching schedule without worrying about network”

Hyeongkil Na,

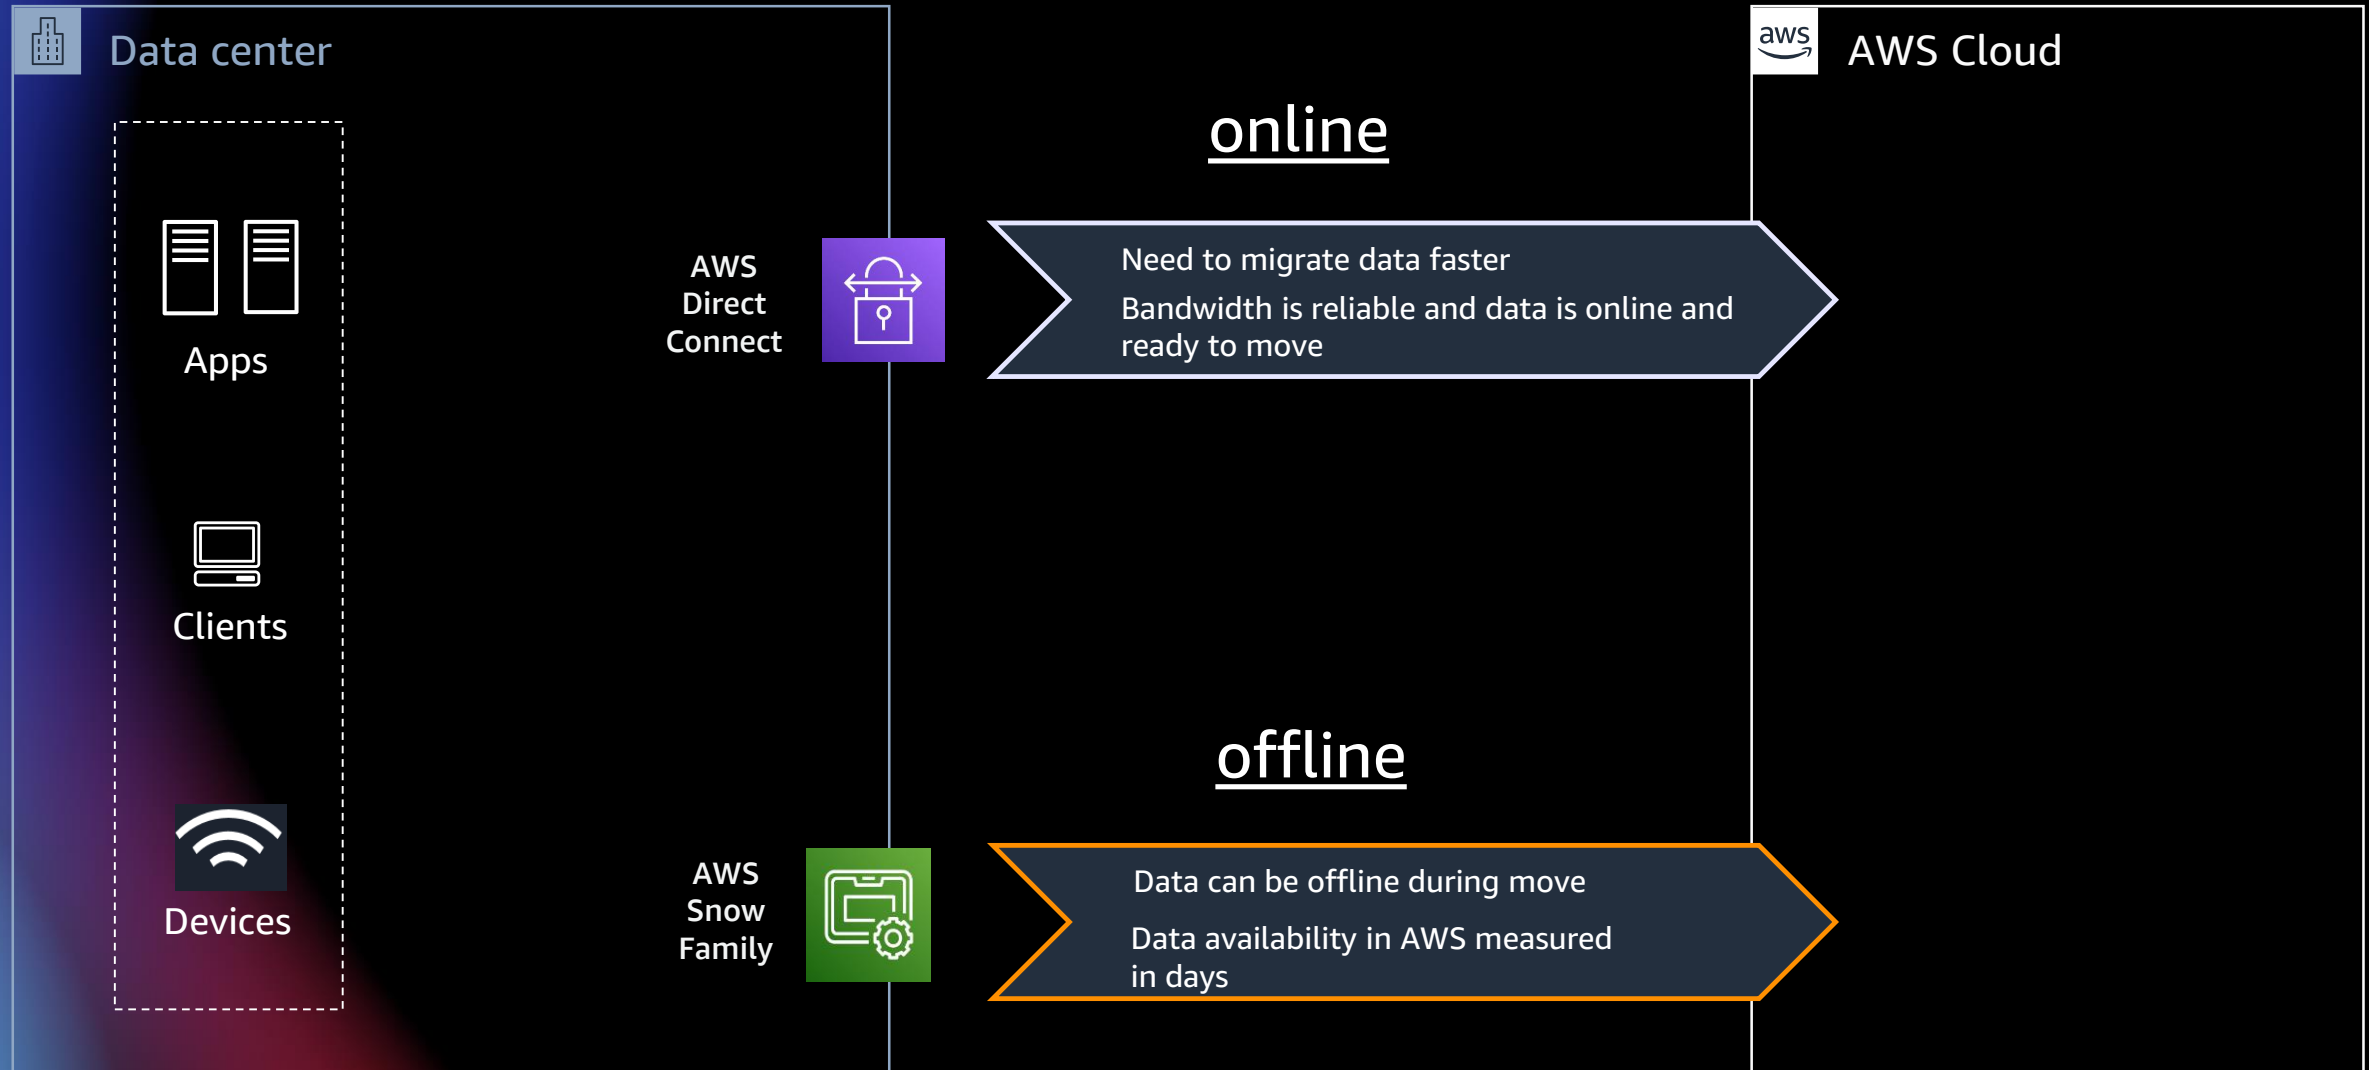
Platform Operation Team leader – Lotte e-commerce



Do I choose online or offline Migration?



Do I choose online or offline Migration?



Time to transfer and migrate data

Time to transfer and migrate data

User scenario: Ahmad's environment consists of a total of 250 TB of data

Time to transfer and migrate data

User scenario: Ahmad's environment consists of a total of 250 TB of data

Data:

100 TB is considered active data (File Shares, Virtual Machines)

150 TB is considered archive data (Historical Backup Data)

Connectivity:

Ahmad's connectivity is currently 1Gbps

Time to transfer and migrate data

User scenario: Ahmad's environment consists of a total of 250 TB of data

Data:

100 TB is considered active data (File Shares, Virtual Machines)

150 TB is considered archive data (Historical Backup Data)

Connectivity:

Ahmad's connectivity is currently 1Gbps

When do you need your data?

	Time to transfer		
	@100Mbps	@1Gbps	@10Gbps
10 TB	12 days	30 hours	3 hours
100 TB	124 days	12 days	30 hours
1 PB	3 years	124 days	12 days

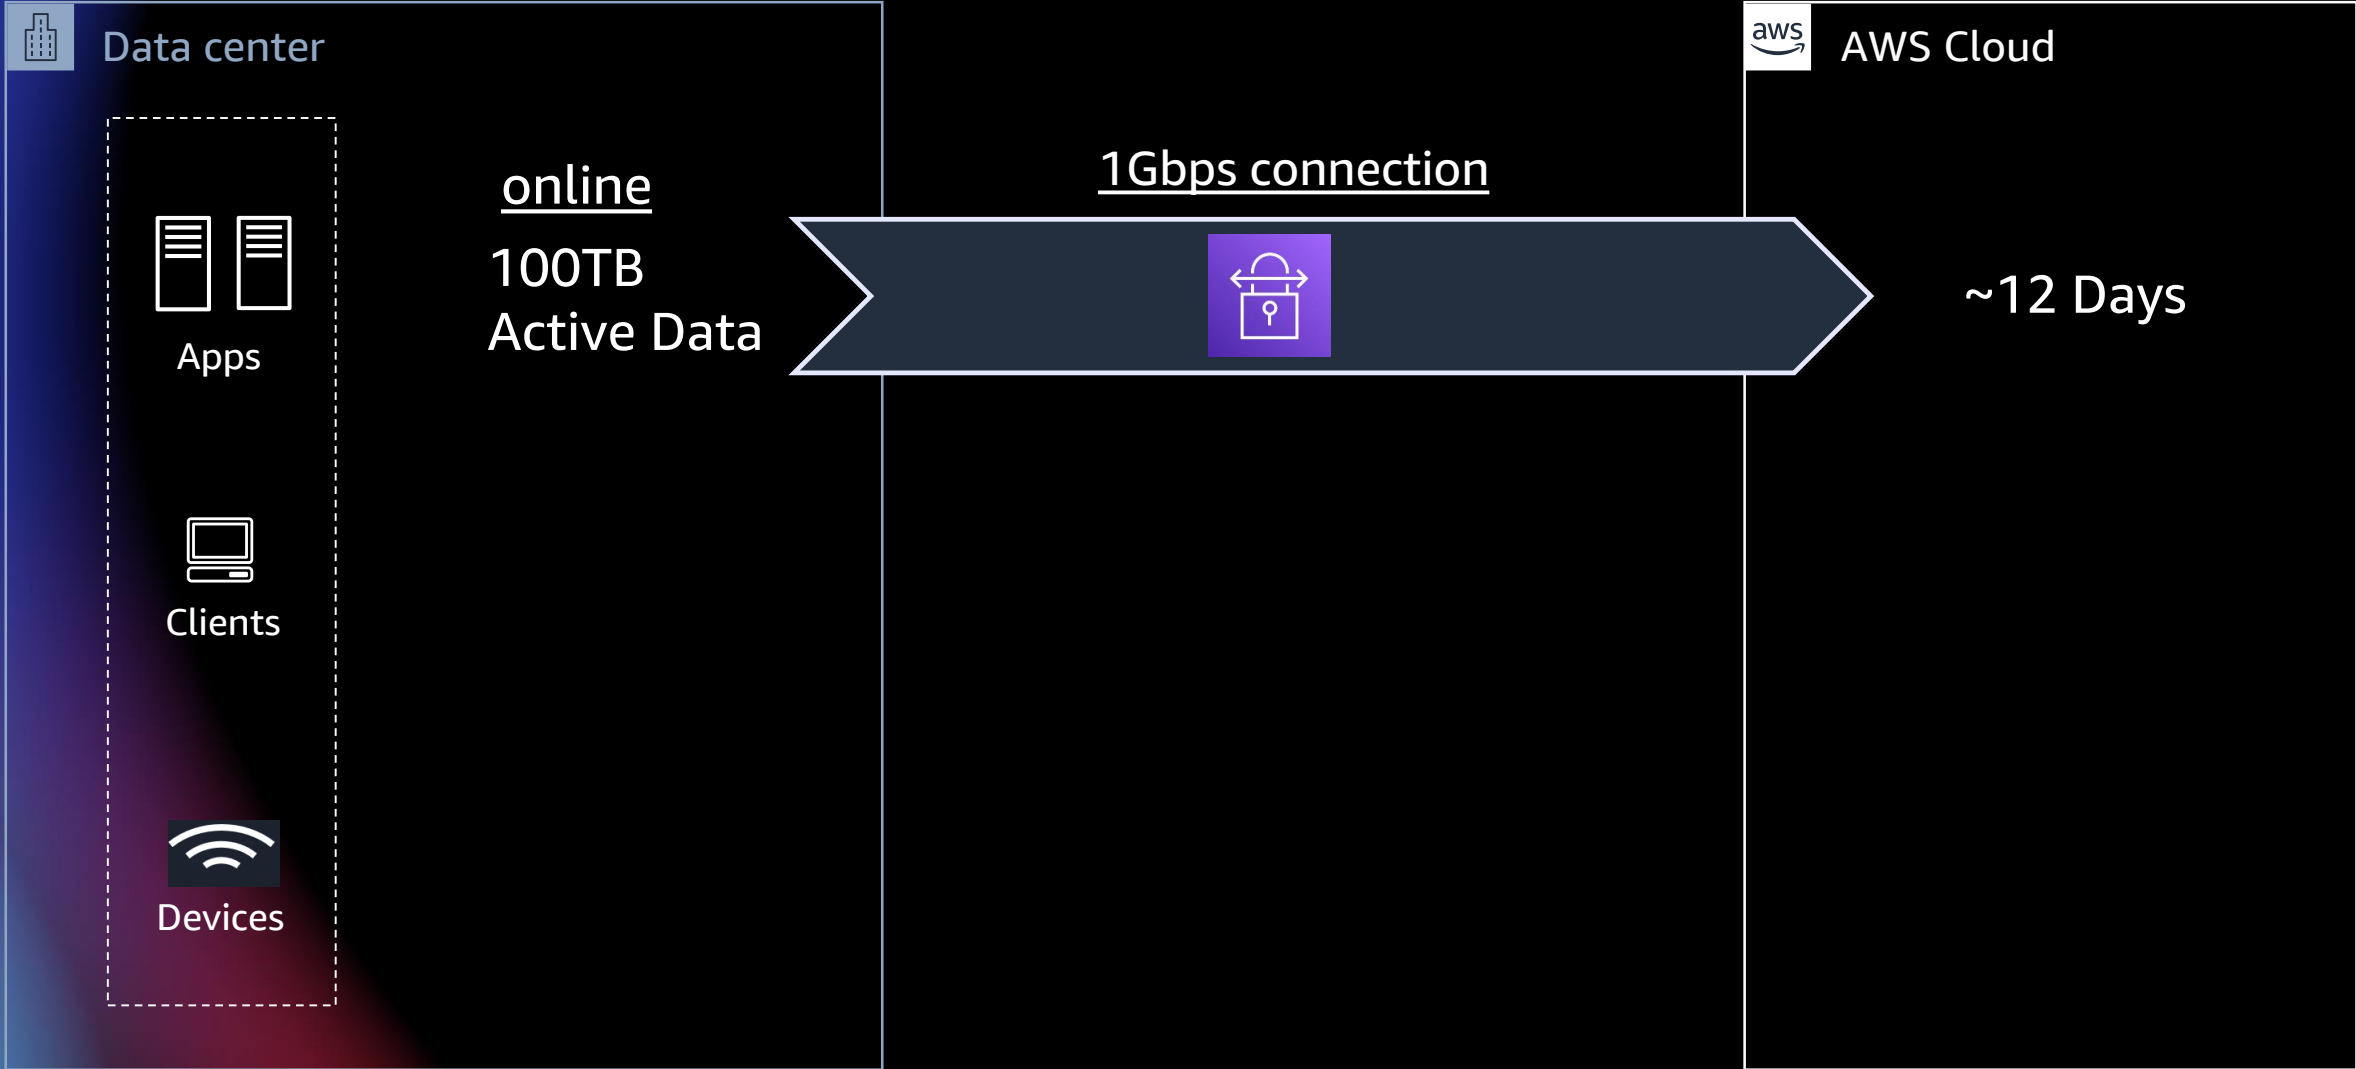
(assumes ~25% network overhead)

When do you need your data?

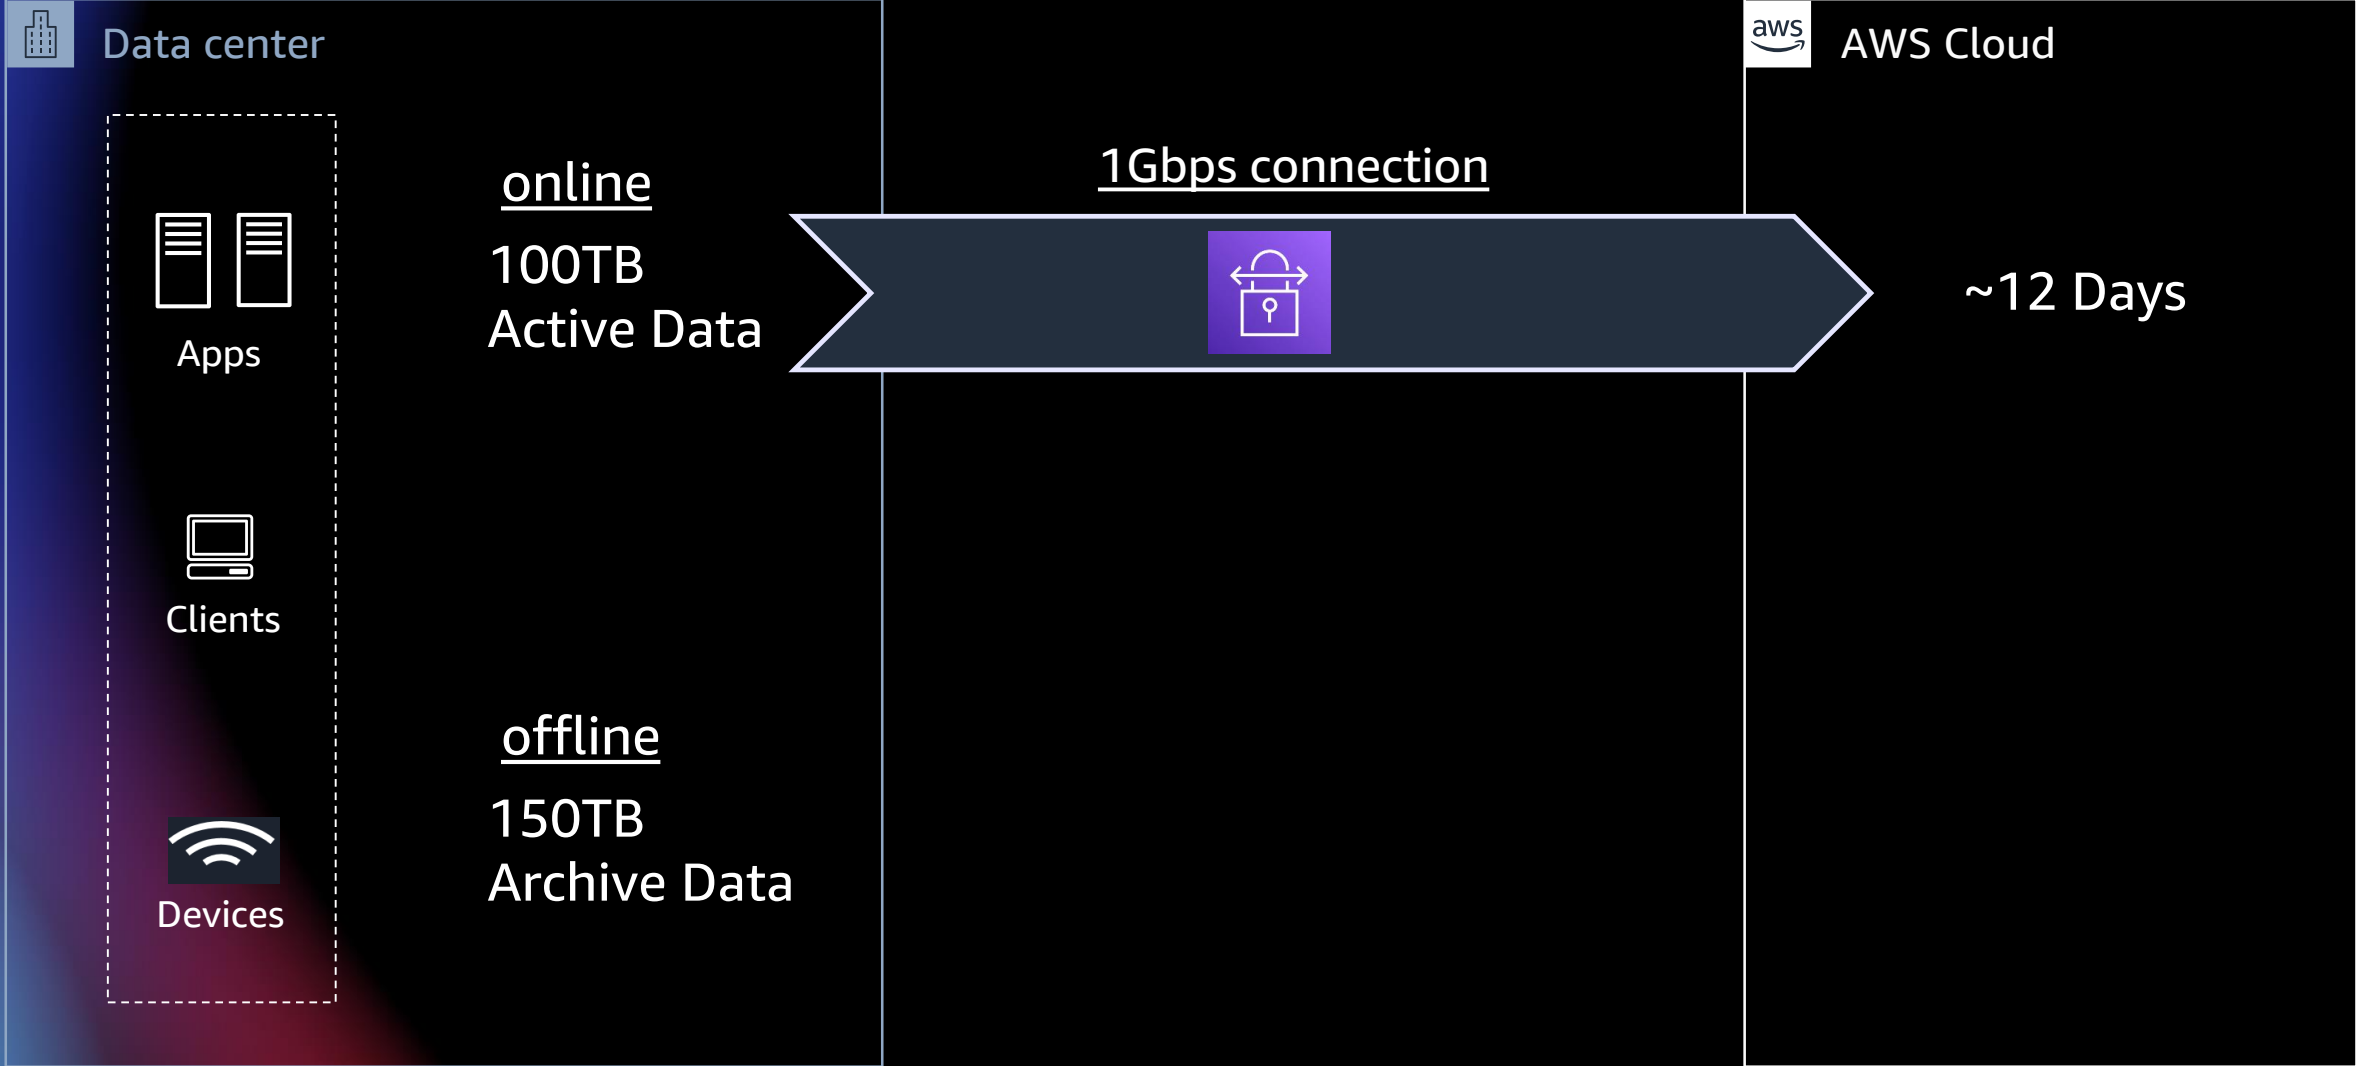
	Time to transfer		
	@100Mbps	@1Gbps	@10Gbps
10 TB	12 days	30 hours	3 hours
100 TB	124 days	12 days	30 hours
1 PB	3 years	124 days	12 days

(assumes ~25% network overhead)

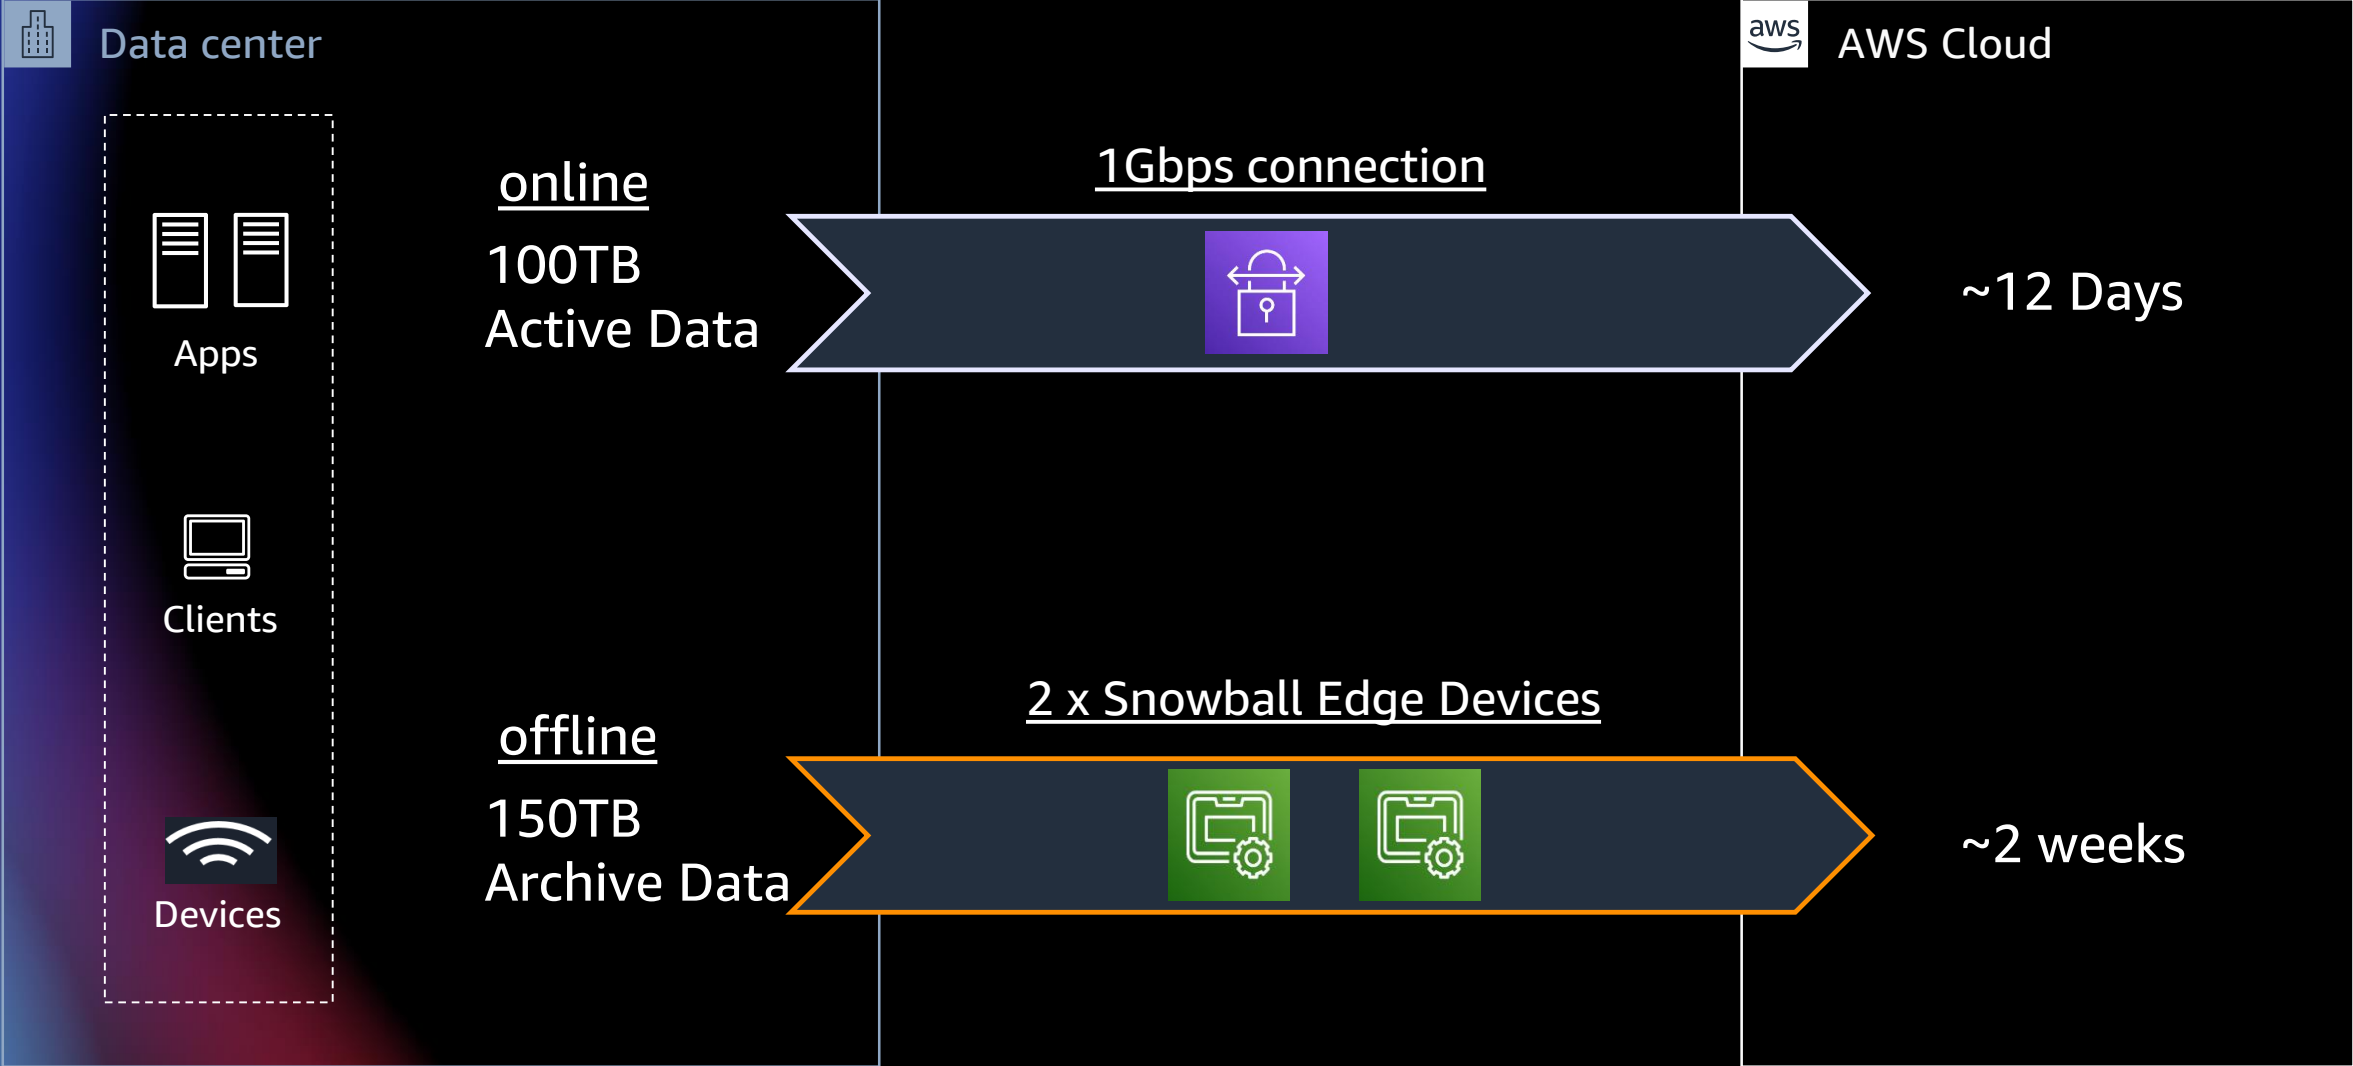
Time to transfer and migrate data



Time to transfer and migrate data



Time to transfer and migrate data



The most complete cloud storage portfolio

Object



Amazon Simple Storage Service (Amazon S3) And S3 Glacier

Block



Amazon Elastic Block Store (Amazon EBS)

File



Amazon Elastic File System (Amazon EFS)



Amazon FSx for Windows File Server



Amazon FSx For Lustre


Broadest and deepest data transfer and edge processing



AWS Storage Gateway



AWS DataSync



AWS Transfer Family



AWS Snowball



AWS Snowcone



AWS Snowmobile


The most complete cloud storage portfolio

Object



Amazon S3
And S3 Glacier

Block



Amazon EBS

File



Amazon EFS



Amazon FSx for
Windows File Server



Amazon FSx
For Lustre

Broadest and deepest data transfer and edge processing



AWS Storage
Gateway



AWS
DataSync



AWS Transfer
Family



AWS
Snowball



AWS
Snowcone



AWS
Snowmobile

The most complete cloud storage portfolio

Object



Amazon S3
And S3 Glacier

Block



Amazon EBS

File



Amazon EFS




Amazon FSx for
Windows File Server



Amazon FSx
For Lustre

Broadest and deepest data transfer and edge processing



AWS Storage
Gateway



AWS
DataSync



AWS Transfer
Family



AWS
Snowball



AWS
Snowcone



AWS
Snowmobile


The most complete cloud storage portfolio

Object



Amazon S3
And S3 Glacier

Block



Amazon EBS

File



Amazon EFS



Amazon FSx for
Windows File Server



Amazon FSx
For Lustre

Broadest and deepest data transfer and edge processing



AWS Storage
Gateway



AWS
DataSync



AWS Transfer
Family



AWS
Snowball



AWS
Snowcone



AWS
Snowmobile


The most complete cloud storage portfolio

Object



Amazon S3
And S3 Glacier

Block



Amazon EBS

File



Amazon EFS




Amazon FSx for
Windows File Server



Amazon FSx
For Lustre

Broadest and deepest data transfer and edge processing



AWS Storage
Gateway



AWS
DataSync



AWS Transfer
Family



AWS
Snowball



AWS
Snowcone



AWS
Snowmobile

The most complete cloud storage portfolio

Object



Amazon S3
And S3 Glacier

Block



Amazon EBS

File



Amazon EFS




Amazon FSx for
Windows File Server



Amazon FSx
For Lustre

Broadest and deepest data transfer and edge processing



AWS Storage
Gateway



AWS
DataSync



AWS Transfer
Family



AWS
Snowball



AWS
Snowcone



AWS
Snowmobile

Streamline data and application migrations

AWS provides a range of data migration services matched to your migration needs




AWS Transfer Family

Regular file transfers directly into and out of Amazon S3 or Amazon EFS using SFTP, FTPS, or FTP




AWS DataSync

Move data between on-premises storage systems and AWS, as well as between AWS storage services



CloudEndure Migration

Simplifies, expedites, and automates large-scale migrations from physical, virtual, and cloud-based infrastructure to AWS



AWS Storage Gateway

Hybrid cloud storage for on-premises access to virtually unlimited cloud storage



AWS DataSync

Quickly, easily, and securely transfer your data to
and from AWS

What is AWS DataSync?

Online data transfer service

that simplifies, automates, and accelerates copying file and object data to and from AWS storage



Fast data transfer



Easy to use



Secure and reliable



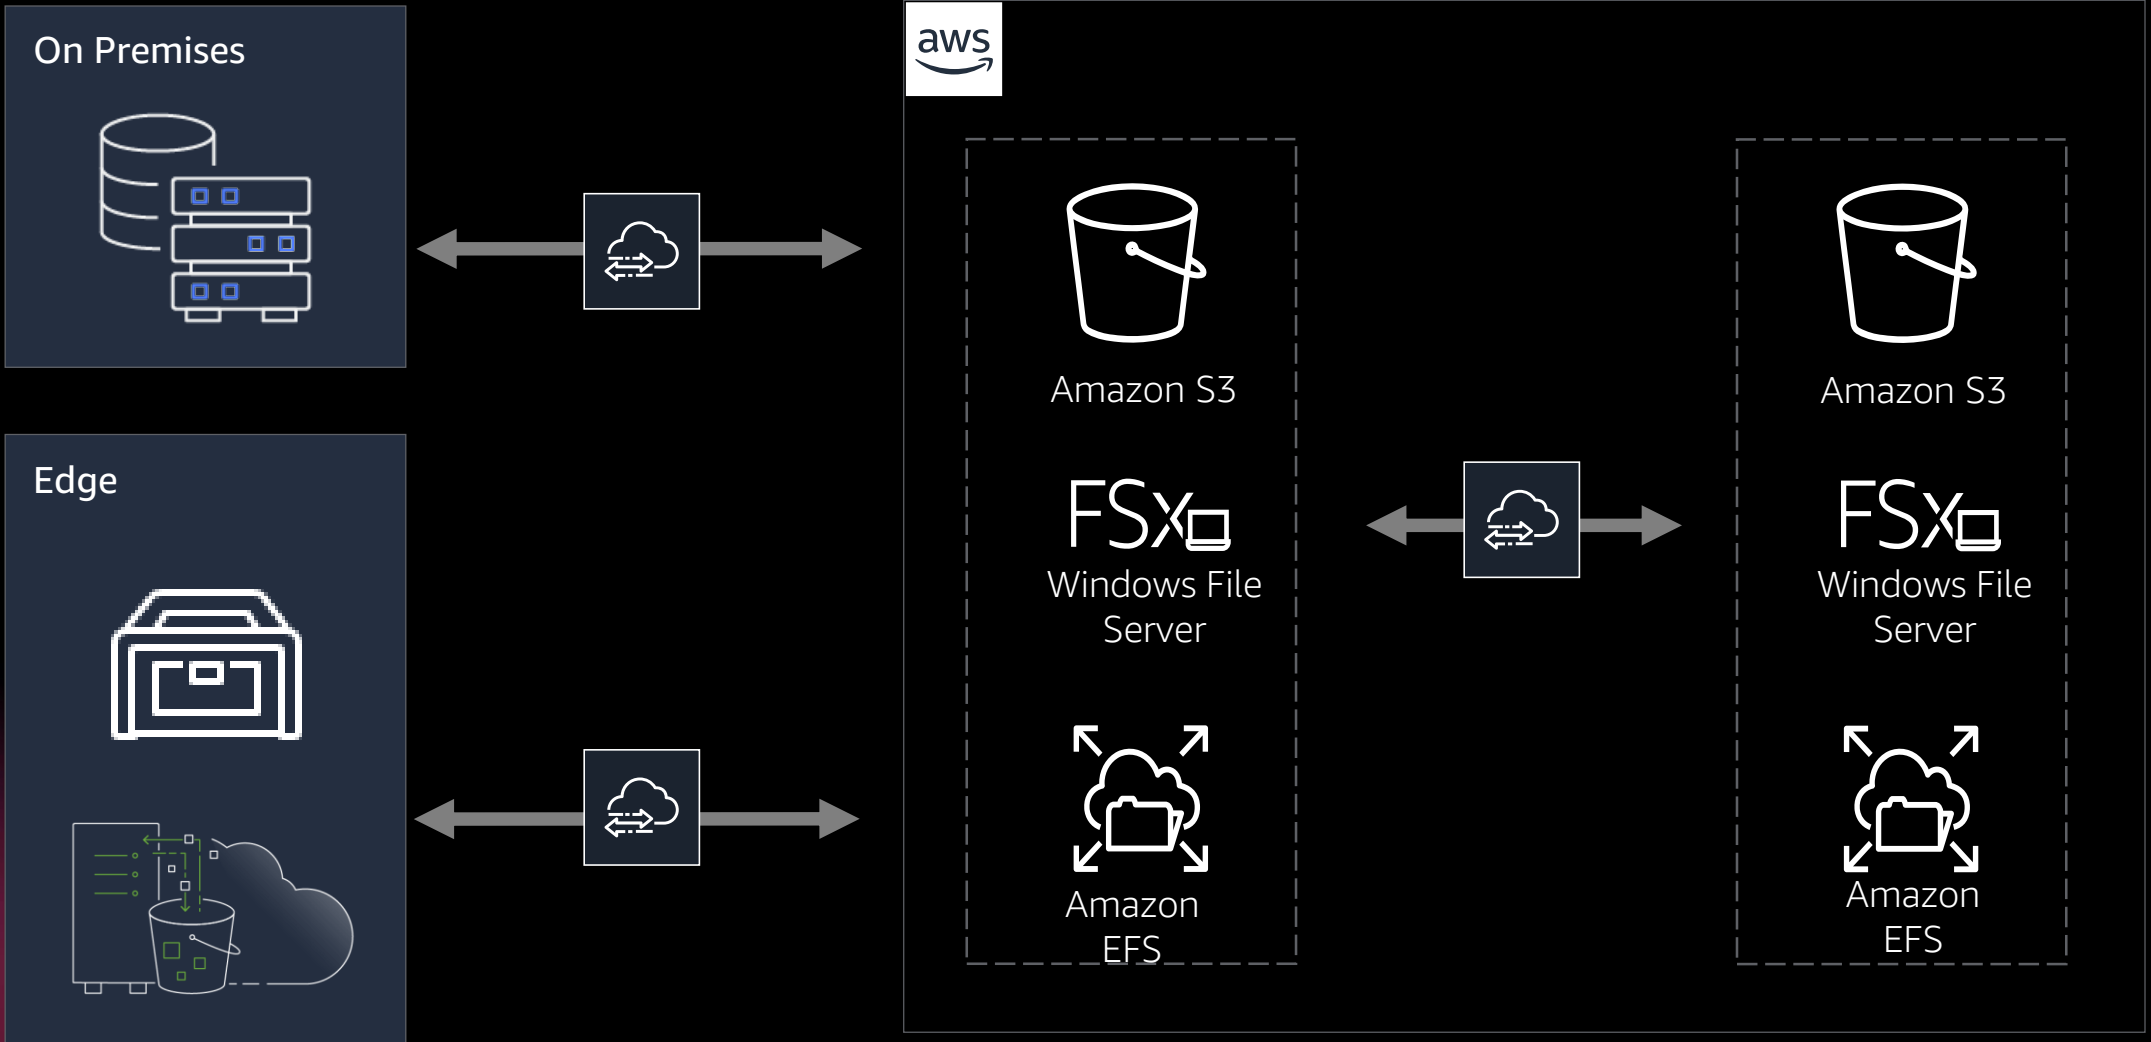
Fully managed



Cost-effective

What can you do with DataSync?

Quickly, easily, and securely transfer your data to and from AWS storage services





“Our petabyte scale data migration journey from on-premises to AWS was accomplished swiftly with minimal effort and was completely self-managed with AWS DataSync. This solution is a game changer!”

Satish Kumar
Infrastructure Engineer

Problem

Wanted to retire multi-petabyte on-premises Data Domain storage system

Data retention policies required data to be retained for many years

Solution

Used Amazon S3 for low cost, pay-as-you-go model as well as versioning support

Used DataSync to move data to S3

Outcome

Successfully transferred dataset to S3 with full byte-for-byte verification

Decommissioned on-premises Data Domain

“Migrating hundreds of TB of data to Amazon S3 with AWS DataSync”



CloudEndure Migration

Highly automated lift-and-shift application migration solution

CloudEndure Migration: a lift-and-shift solution

Simplifies, expedites, and reduces the cost of migrating applications to AWS



Phase One: Lift-and-shift

Quickly take advantage of AWS:

- Cost savings
- Productivity
- Resilience
- Agility



Phase Two: Modernize

After applications are on AWS,
it's easier to:

- Replatform
- Refactor

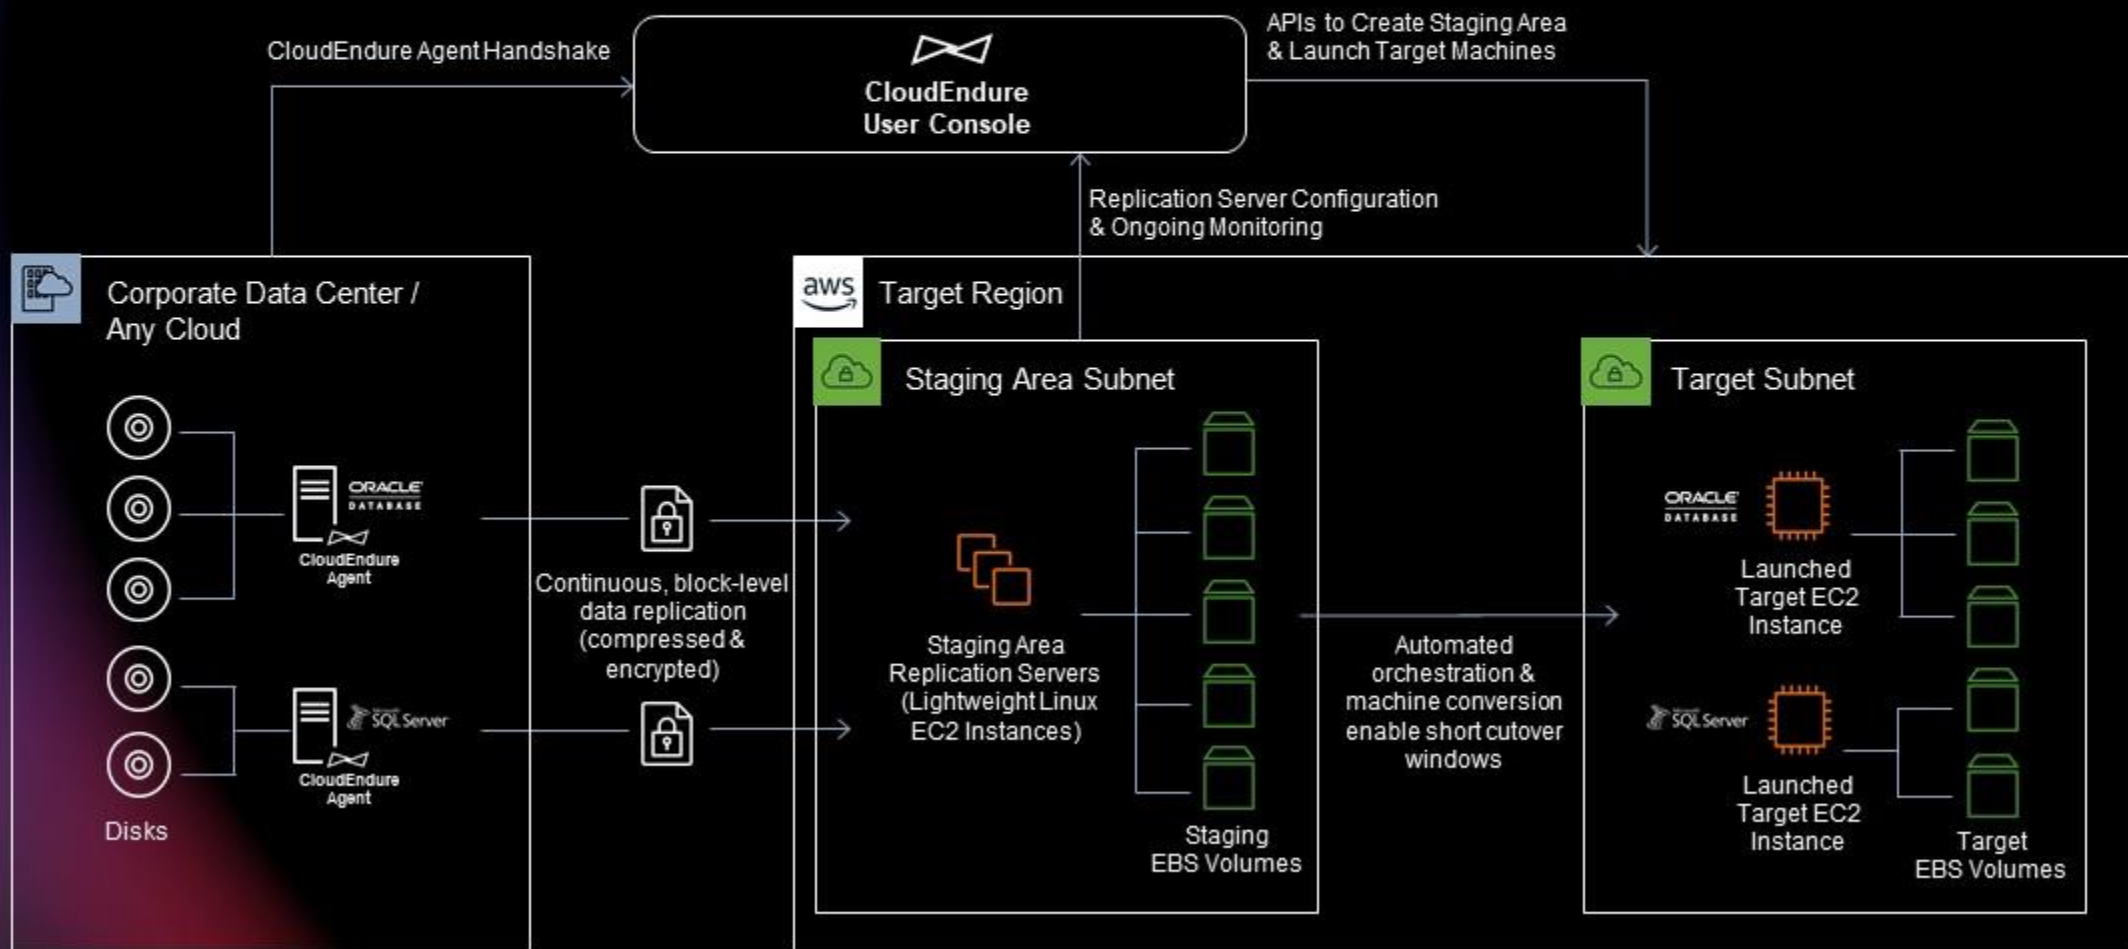
Why lift-and-shift (rehost)?

Lift-and-shift is fast, cost-effective, and minimally disruptive migration.

- Migrate applications to the cloud with **minimal or no changes**.
- Applications are effectively “lifted” from the existing environment and **“shifted” as is to the AWS Cloud**.
- Lift-and-shift **maintains the architecture constructs** of your source servers so after migration there are no significant changes to business processes, monitoring, and management interfaces.
- Lift-and-shift can be **automated with tools such as CloudEndure Migration**.
- Use AWS services after migration to **modernize applications**.

How CloudEndure Migration works

- Continuously replicates any application or database from any source into AWS
- Business outcome: self-service, rapid, reliable migration with minimal business disruption



- Implemented a two-phase migration strategy: “Lift and Shift” and then “Replatform.”
- Used CloudEndure Migration to lift and shift more than 1,500 servers to AWS.
- Averaged 120 servers/month.
- Workloads included Microsoft SharePoint sites, SQL Server, Java web applications, custom software, and legacy applications.



inmarsat

“If we had started with replatforming, it could have taken years. Our lift-and-shift strategy, using CloudEndure Migration, allowed us to move 65% of our IT infrastructure in just 6 months.”

–Tim Brown, Cloud and Infrastructure Operations Director, Inmarsat



AWS Transfer Family


Simple and seamless frequent B2B file transfers to Amazon S3 and Amazon Elastic File Service (EFS) using SFTP, FTPS, and FTP

AWS Transfer Family

Fully managed service for file transfers over SFTP, FTPS, and FTP for Amazon S3 and Amazon EFS




Seamless migration
of existing workflows



Fully managed, high
availability (HA), and
globally available




Secure and Compliant



Native integration
with AWS services



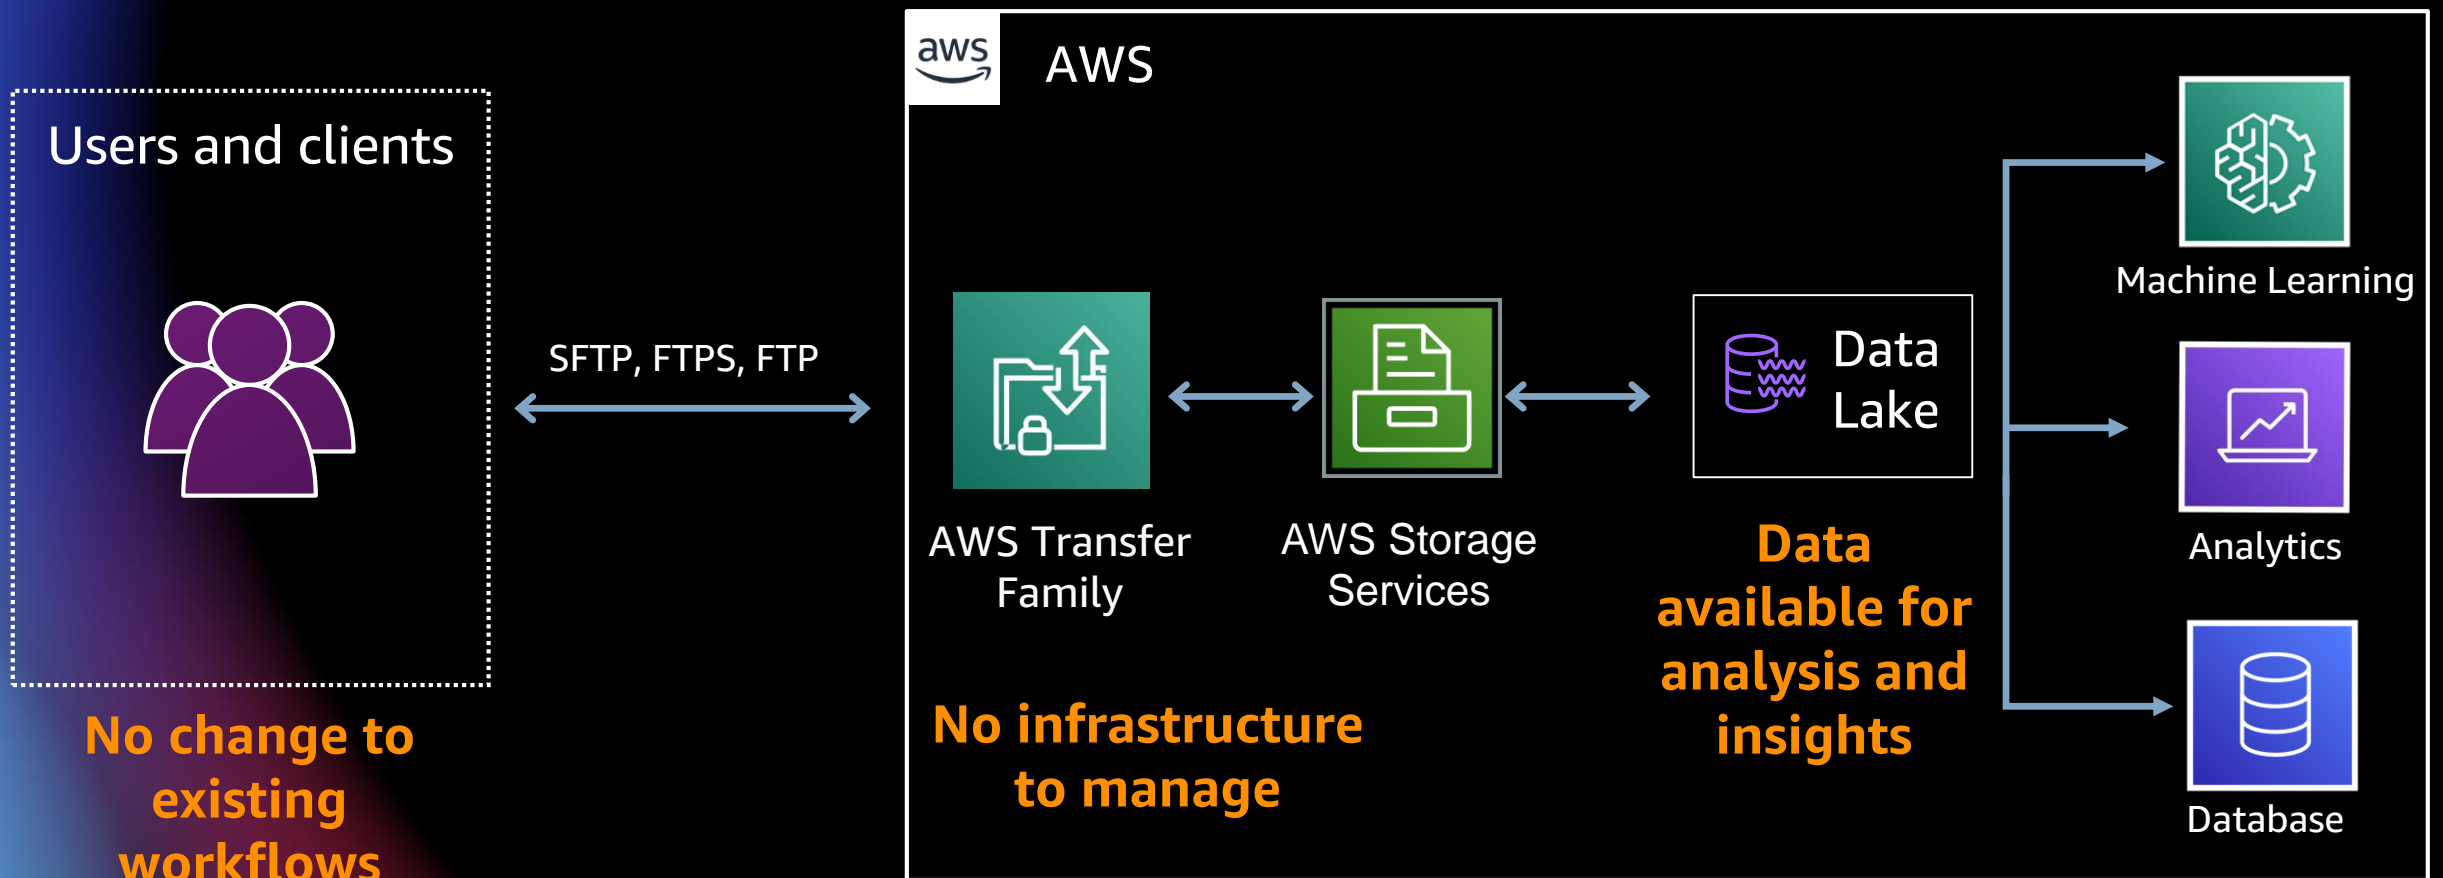
Cost Effective



Simple
to use

How AWS Transfer Family works

A high level overview of the AWS Transfer Family architecture





AWS Storage Gateway

On-premises access to virtually unlimited cloud storage

Hybrid cloud storage with Storage Gateway



On-premises

In-cloud



Easy on-ramp to the cloud

Fast time to deployment

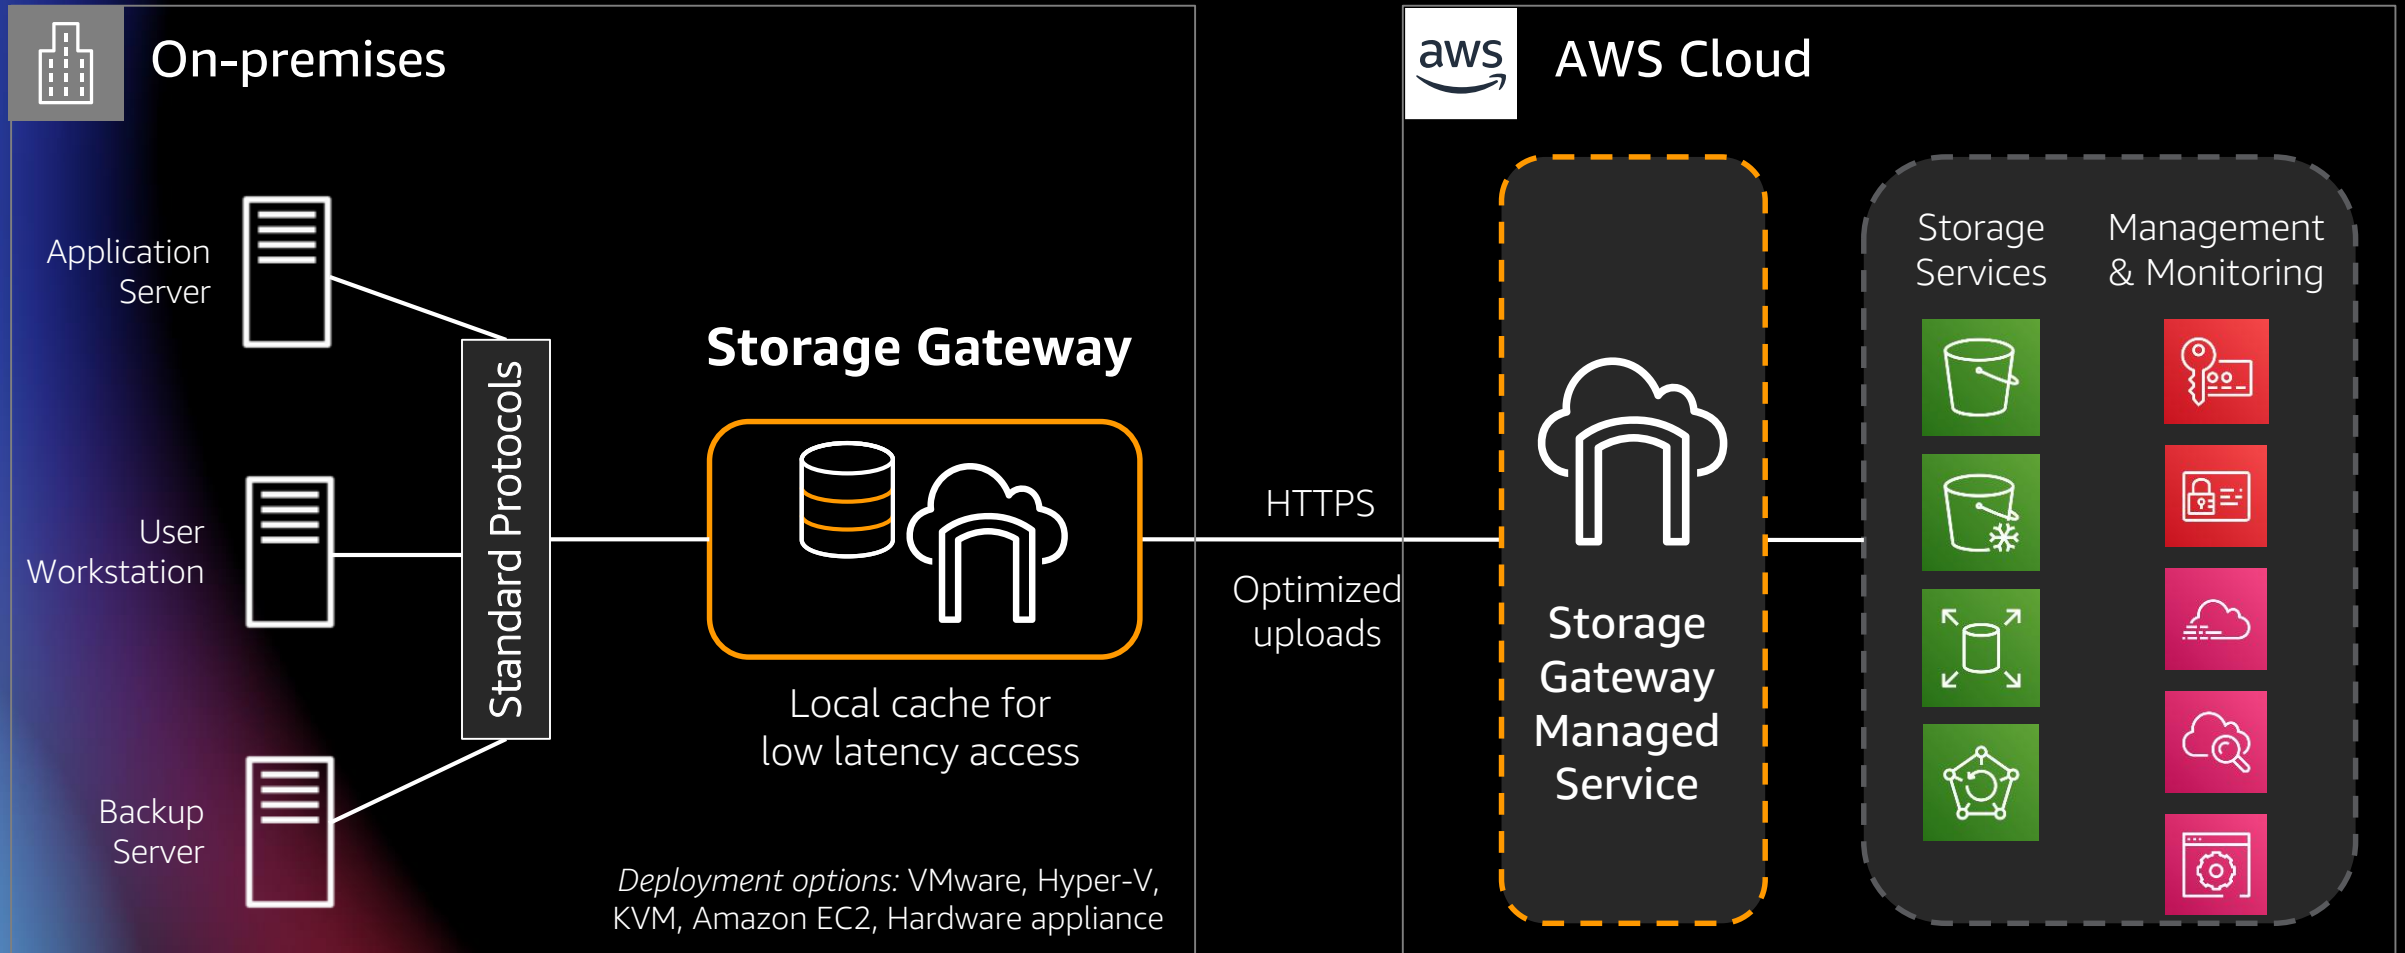
Agility & scale of the cloud

No changes to applications


Centralized management & control

AWS Storage Gateway

On-premises access to virtually unlimited cloud storage



Storage Gateway family



S3 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 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

Storage Gateway family



S3 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 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



TransferWise

"Using AWS, we can spin up compute and storage resources where and when we need to, much faster than we ever could before. AWS will help us continue to drive our global expansion."

Thomas Hewer, Lead Technologist,
TransferWise



Problem

Self-managed on-premises data centers could not provide the scale or availability for application environments needed to compliment the company's rapid global growth

Solution

- Migrated all its on-premises backups and databases to AWS using AWS Storage Gateway
- Replaced tape backups with Tape Gateway and wrote database backups to Volume Gateway, which created Amazon EBS Snapshots that were managed by AWS Backup
- Moved hundreds of PostgreSQL and MySQL databases into a combination of Amazon RDS and Amazon EC2 with Amazon EBS

Outcome

- Kept pace with its business growth by scaling web and mobile applications with AWS
- Reduced time and overhead of managing data centers
- Increased confidence in backups with automated backups to Amazon EFS and Amazon RDS

Recap

- Connectivity options to AWS
- Offline Transfers with Snow Family
- Various Migration and Transfer Methods, both Online and Offline
- Complete Set of Cloud Storage – Cost and Performance Optimized
- Set of Hybrid and Migration Tools to help you with your Data Migrations

Additional resources

Get hands-on experience with the AWS online data migration workshop

<https://github.com/aws-samples/aws-online-data-migration-workshop>

AWS Snow Family : <https://aws.amazon.com/snow>

AWS DataSync : <https://aws.amazon.com/datasync>

AWS Storage Gateway : <https://aws.amazon.com/storagegateway>

AWS Transfer Family : <https://aws.amazon.com/aws-transfer-family>

AWS Migration Service (CloudEndure) : <https://aws.amazon.com/application-migration-service/>

AWS Storage : <https://aws.amazon.com/products/storage/>



Visit the AWS Data Resource Hub

Dive deeper with these resources, get inspired and learn how you can use data to make better decisions and innovate faster.

- Building a winning data strategy
- The new leadership mindset for data & analytics
- Harness data to reinvent your organization
- Put your data to work with a modern analytics approach
- Breaking free from on-premises database constraints
- Cloud storage adoption: From cost optimization to agility & innovation
- A strategic playbook for data, analytics, and machine learning
- ... and more!



<https://tinyurl.com/aws-data-resource>

Visit resource hub



AWS Training and Certification

Empower your teams with comprehensive training

By building skills with AWS Training and Certification, businesses and individuals can see the bigger picture understanding the reasoning behind every data point. As training progresses and teams become data-fluent, previously hidden insights come into view.

Build data skills to
unlock any insight

Leverage free digital training

Learn how to harness the world's most valuable resource: data. Access digital and virtual instructor-led courses on data analytics and databases built by the experts at AWS and start your learning journey to become data-driven.

[Take a digital course »](#)



Get certified

Earn industry-recognized credibility and set tangible goals for success with industry-recognized certifications, like *AWS Certified Data Analytics – Specialty*.

[Learn more »](#)



Ramp-up your skills

Deep dive into new topics and focus on knowledge gaps at your own pace with the *AWS Ramp-Up Guide: Database* and *AWS Ramp-Up Guide: Data Analytics*. With a wide range of whitepapers, blog posts, videos, webinars and peer resources available for data professionals to leverage for independent learning.

[Download ramp-up guides »](#)

Thank you for attending AWS Innovate – Data Edition

We hope you found it interesting! A kind reminder to **complete the survey**.
Let us know what you thought of today's event and how we can improve the event experience for you in the future.



aws-apj-marketing@amazon.com



twitter.com/AWSCloud



facebook.com/AmazonWebServices



youtube.com/user/AmazonWebServices



slideshare.net/AmazonWebServices



twitch.tv/aws

Thank you!