AWS re:Invent

STG208

Backup-and-restore and disaster-recovery solutions with AWS

Jody Kirk
Principal, Global Storage BD
Amazon Web Services

Creighton Swank
Cloud Architect
Caterpillar

Juan Mejia
VP, Data Center Management
BankUnited

Agenda

Why protect data

Modernizing backup using the cloud

Common customer backup models

Disaster recovery approaches

Q&A

Why do we backup data?

Minimize time to recover

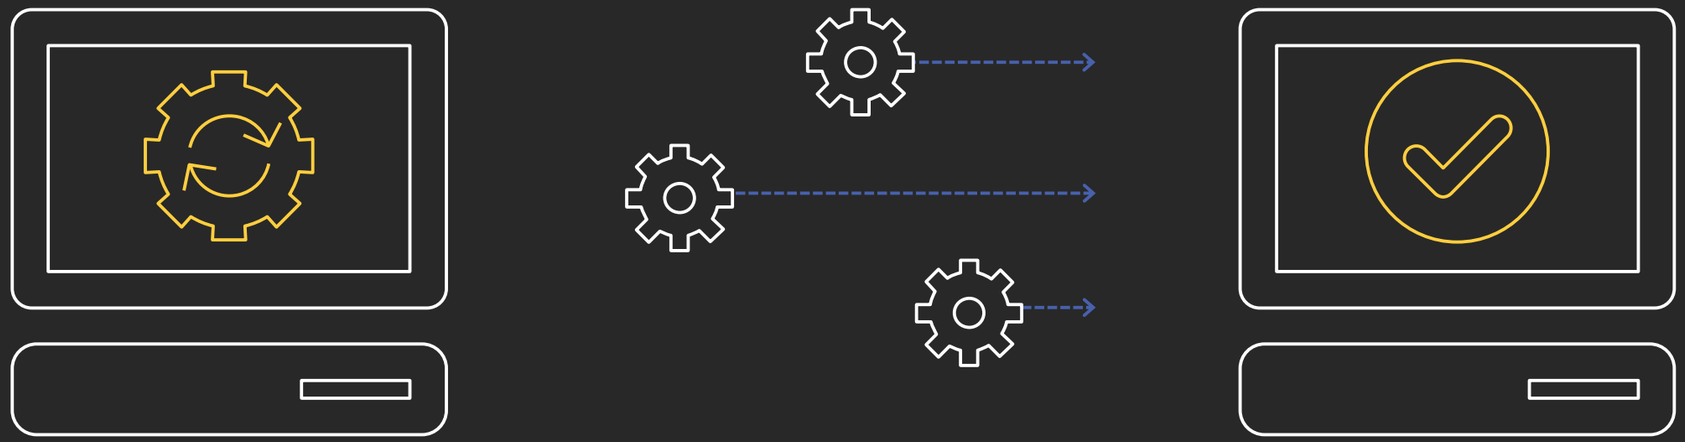

Why do we backup data?

Minimize data loss

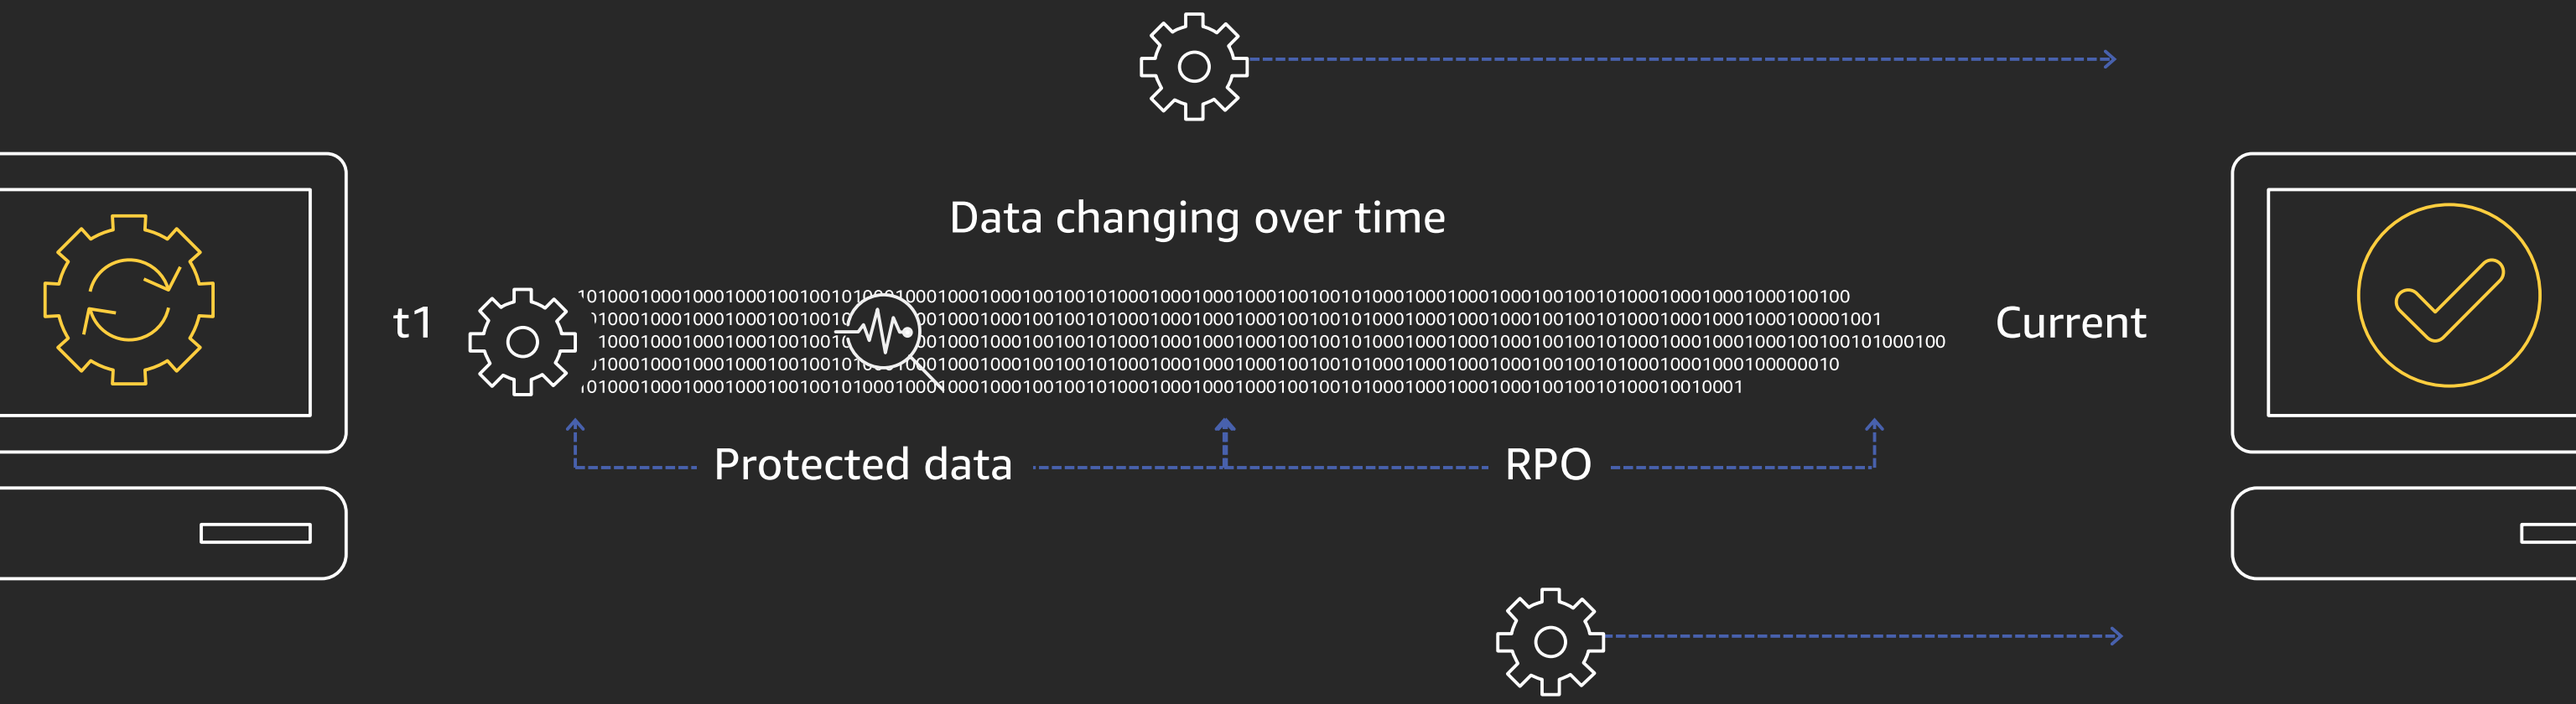

Why do we backup data?

Balance cost with liability

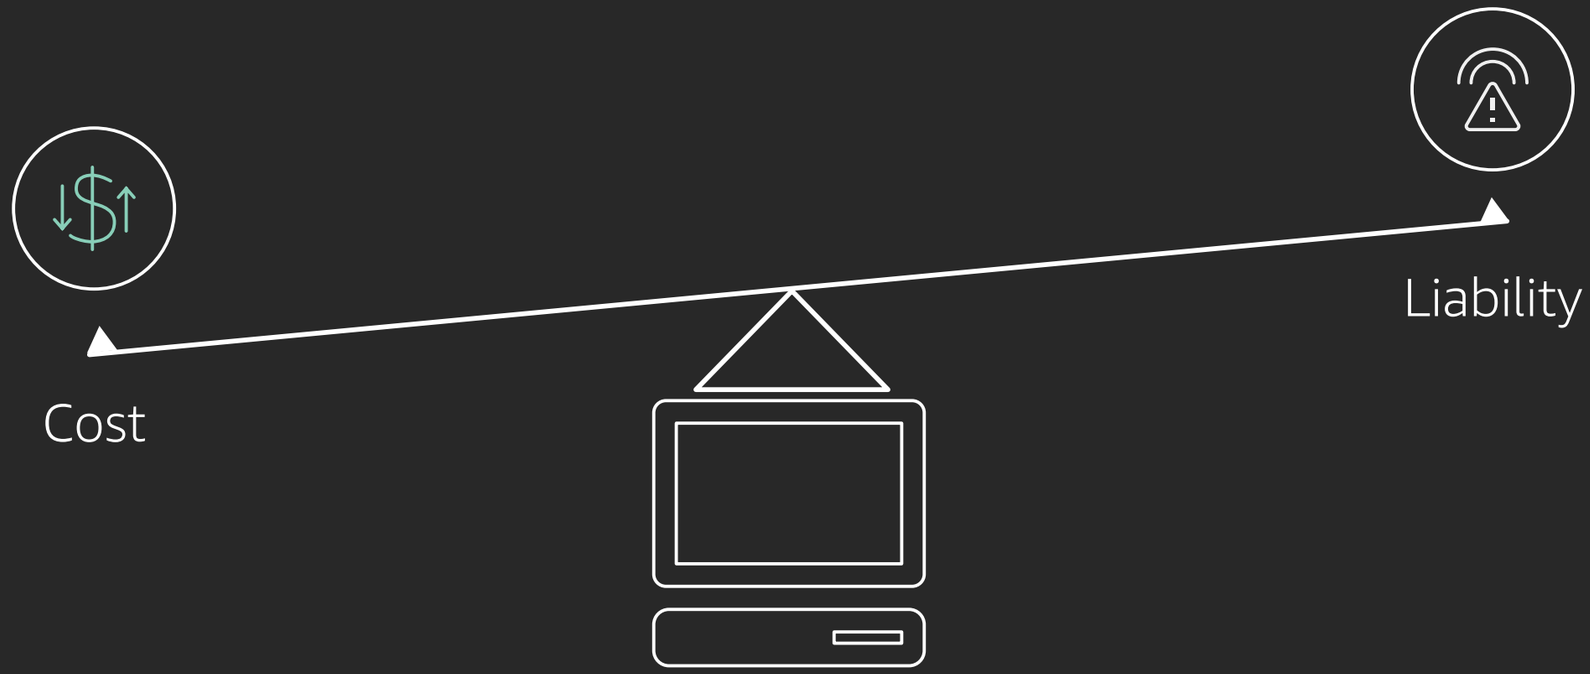

Initial questions to answer

How important are the applications to your business?

What is the associated recovery point and time for these applications?

How are you storing the data?

Where are you storing the data?

How are you restoring the application?

Traditional backup environment

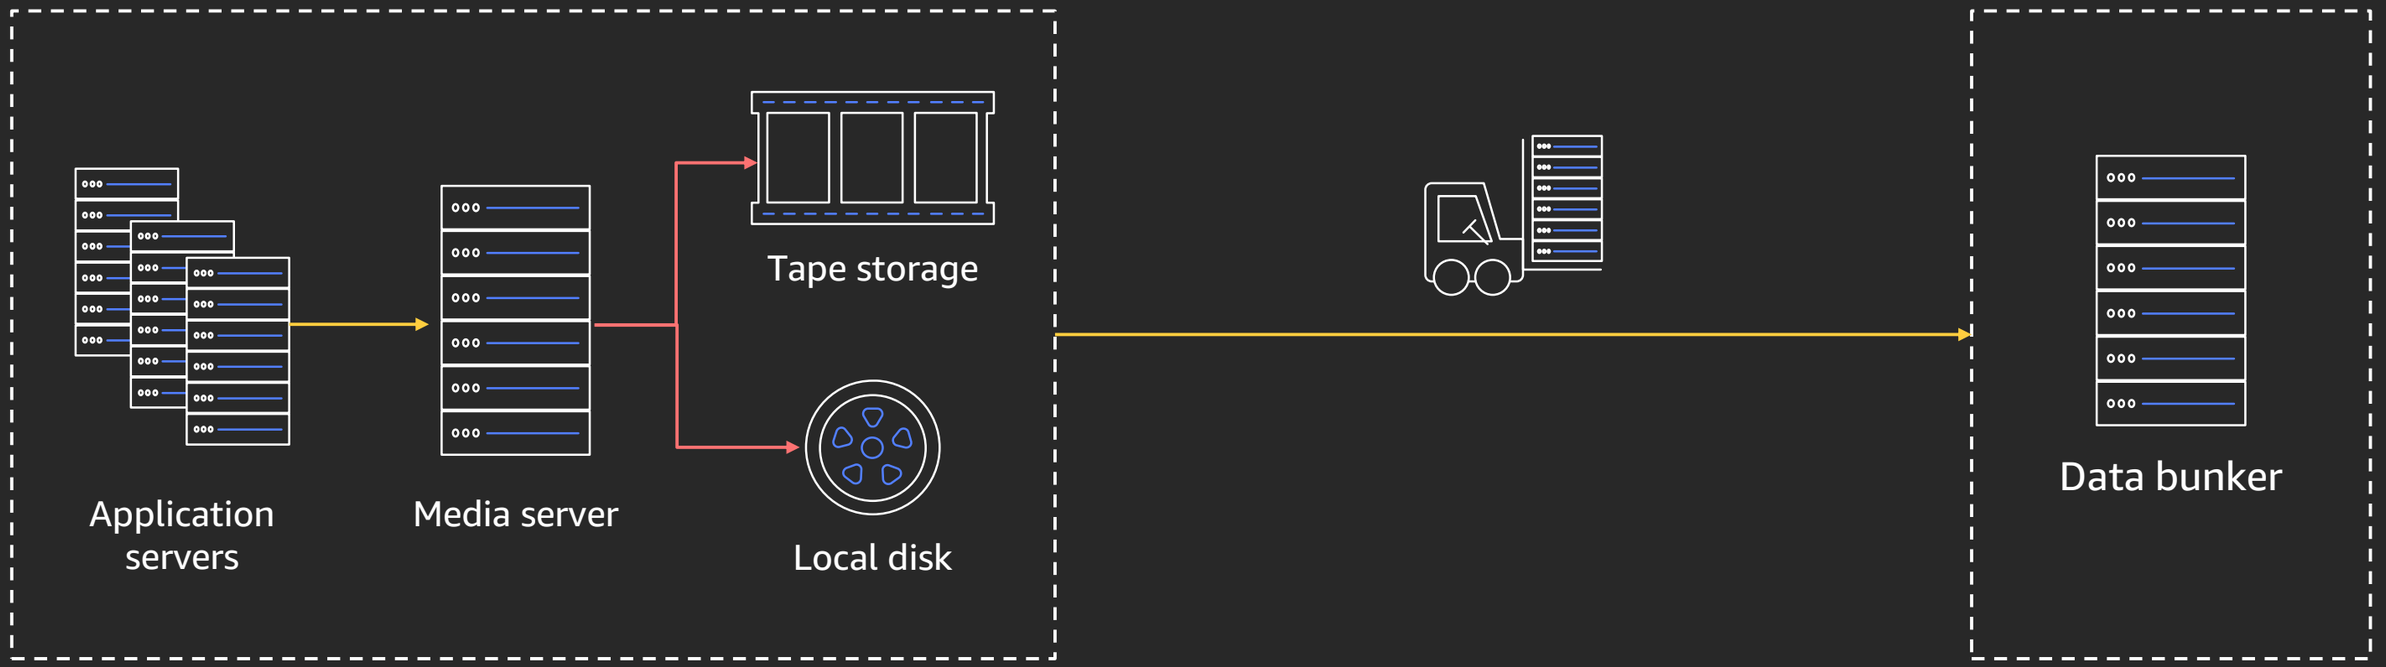

Modernizing backup architectures using cloud

API driven

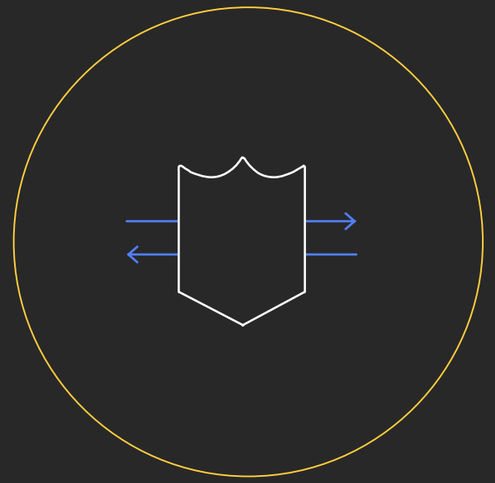

Simple, secure
and durable

Variable cost

Immediate cloud backup benefits

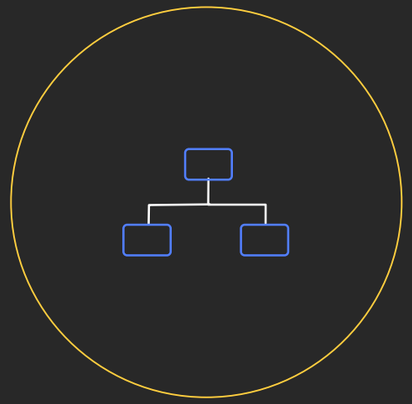

Leverage existing investments in infrastructure

…cloud as a backup target integrates with existing backup frameworks

Cost effective offsite storage alternatives

…with pay as you go pricing and no upfront capital investments

Elimination of physical tape backups and administration

…for a low-cost, highly scalable virtual alternative with nominal disruptions to existing systems

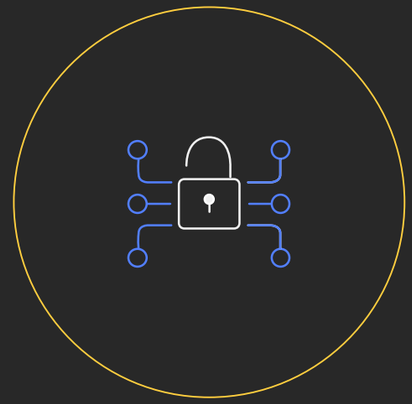

Unlocking insights from your data

…by applying analytics, artificial intelligence, and machine learning capabilities

Amazon S3 and Amazon S3 Glacier designed for 11 9s of durability (99.999999999%)

Region & number of Availability Zones

Announced Regions

Bahrain, Cape Town, Jakarta, and Milan

AWS Region

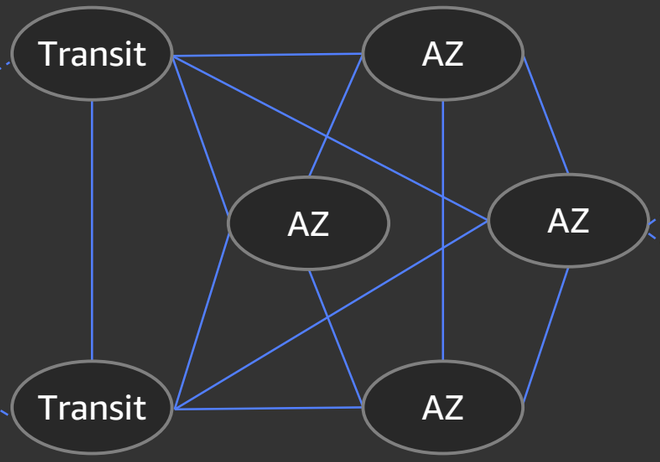

A Region is a physical location in the world where we have multiple **Availability Zones.**

AWS Availability Zone (AZ)

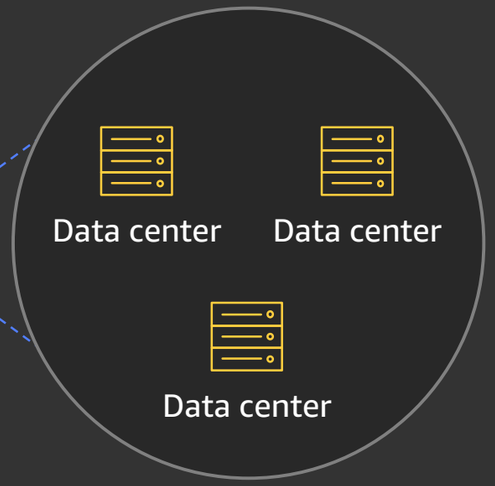

Availability Zones consist of one or more discrete data centers, each with redundant power, networking, and connectivity, housed in separate facilities.

Cross-region replication for compliance

Flexibility to replicate data:

- At the bucket, prefix, or object level
- From any Region to any Region
- To any storage class
- Across AWS accounts

Use cases:

- Compliance—store data 100s of miles apart
- Lower latency—distribute data to remote customers
- Security—create remote replicas managed in separate accounts

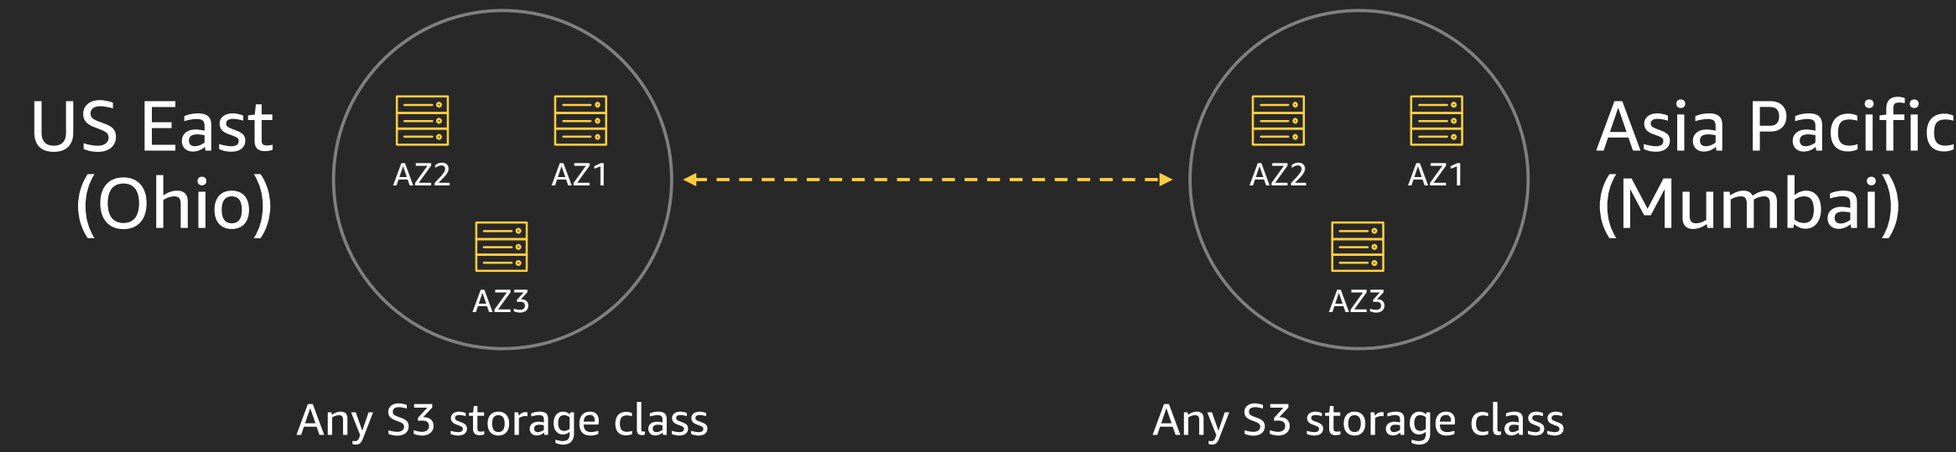

Backup patterns

File gateway

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

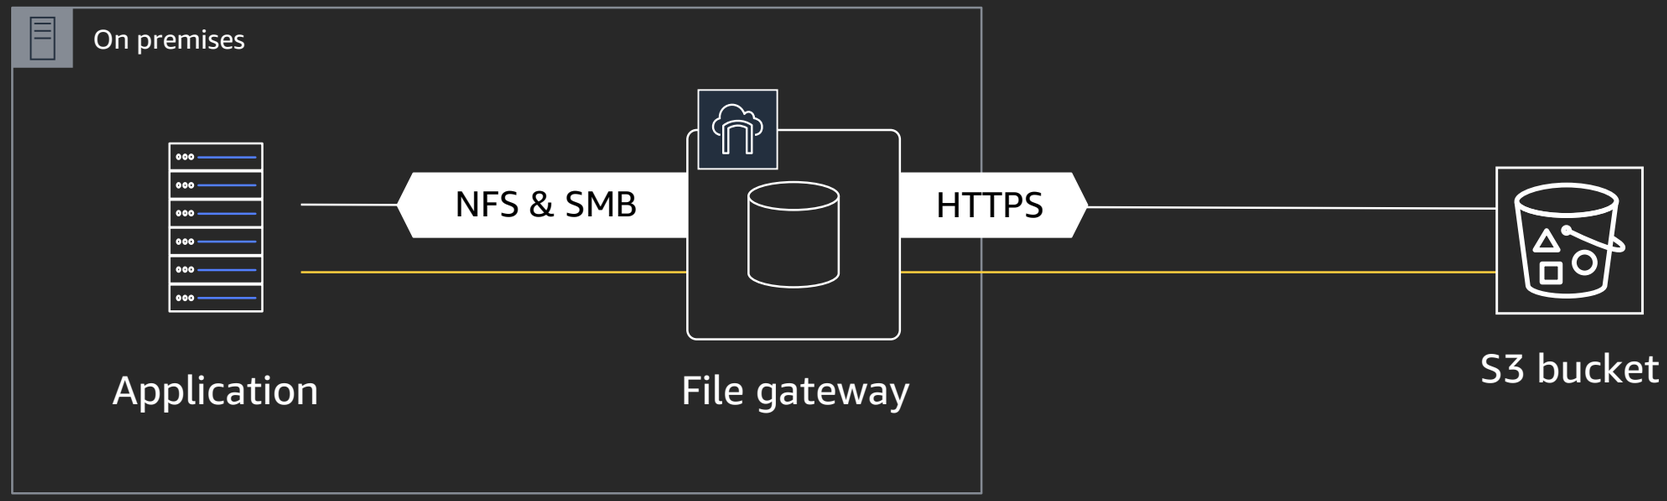

Use cases

- Back up on-premises data to the cloud
- Hybrid cloud workflows
- Low-latency on-premises access to cloud storage
- Content distribution and collaboration via the cloud

Celgene

Direct Supply

Tape gateway

Virtual tapes presented to on-premises backup applications

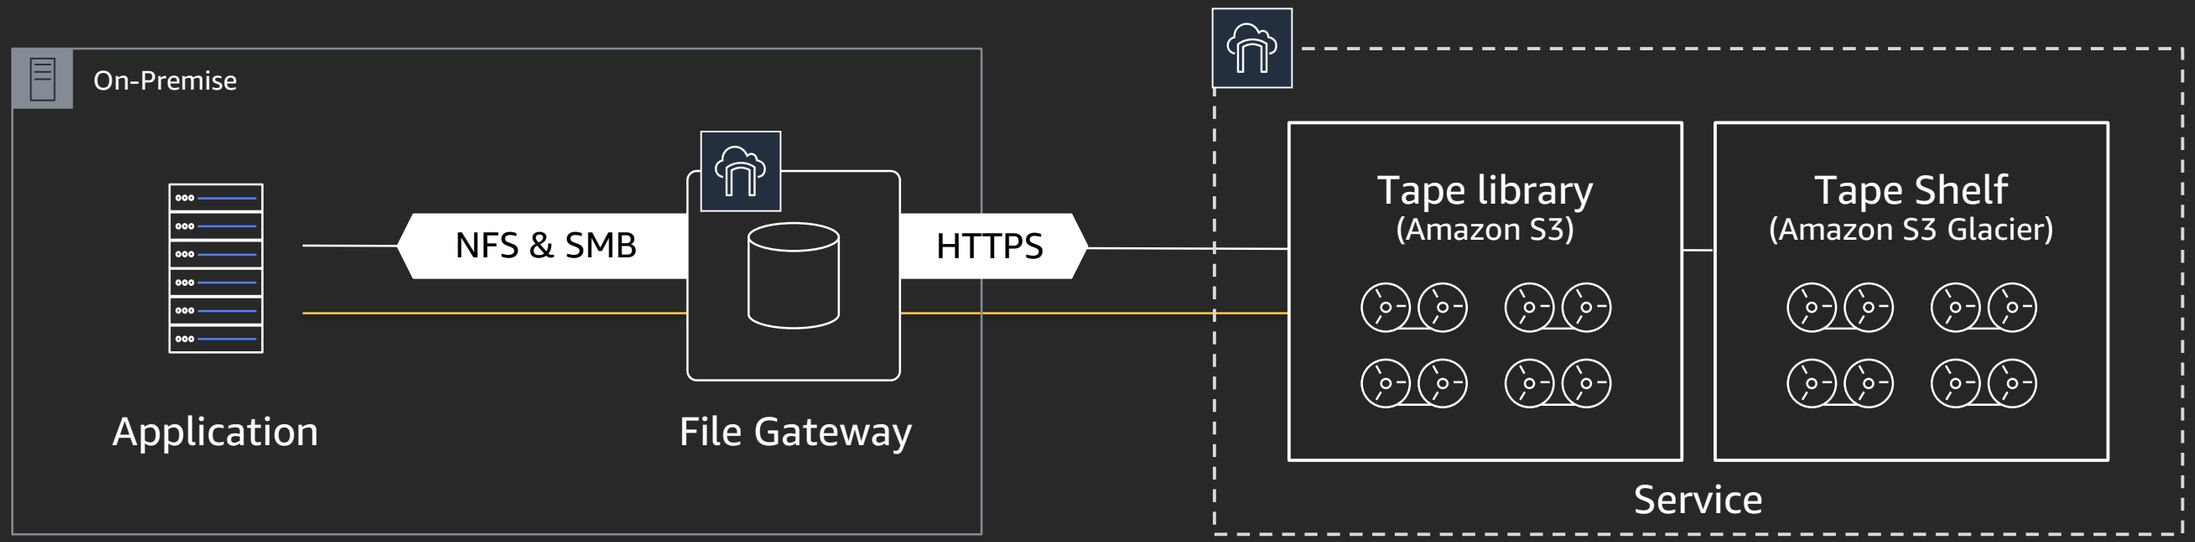

Use cases

- Tape based backup with existing backup apps
- Archive to Amazon S3 Glacier

ANALOG DEVICES
AHEAD OF WHAT'S POSSIBLE™

Southern OREGON UNIVERSITY

Volume gateway

Block storage on-premises backed by cloud storage

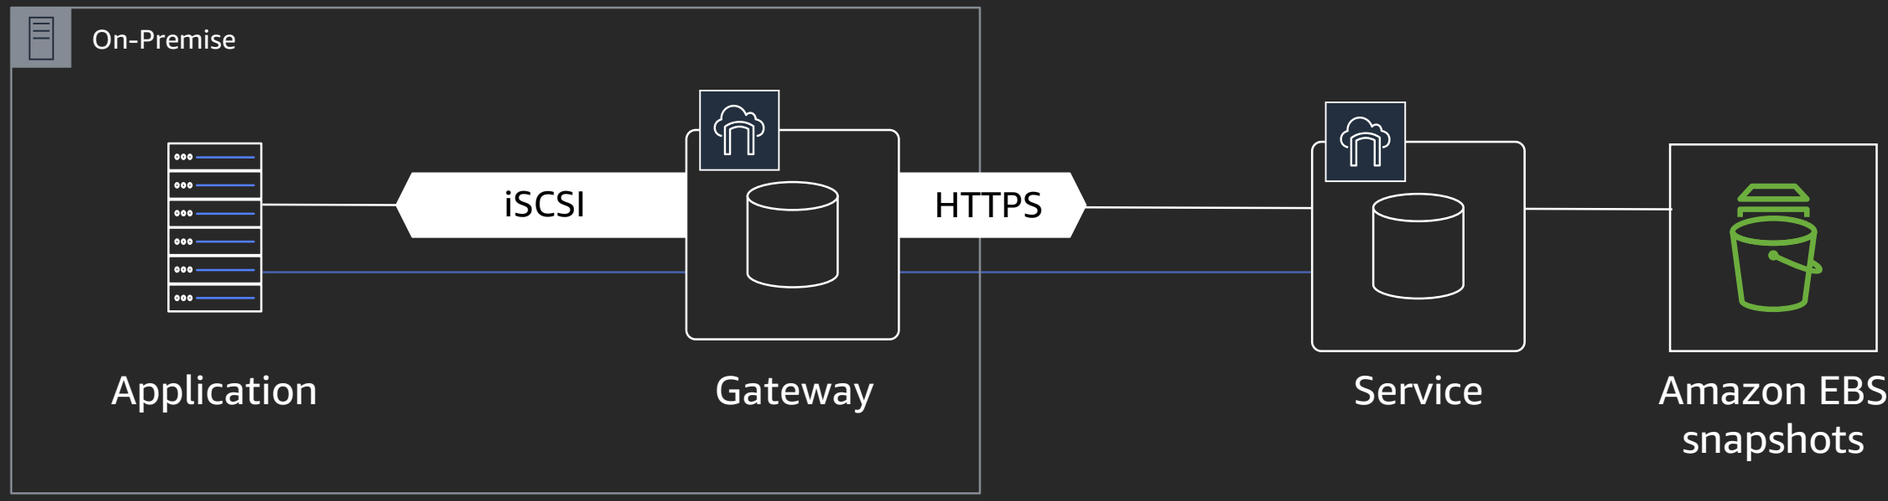

Use cases

- Backup on-premises data to the cloud
- Migration of volumes to the cloud
- DR to the cloud

Stemcell Technologies

iRobot

Amazon EBS Snapshots

Block storage data protection in AWS

Use cases

- Back up in the cloud
- Disaster recovery

Amazon EFS: database backups

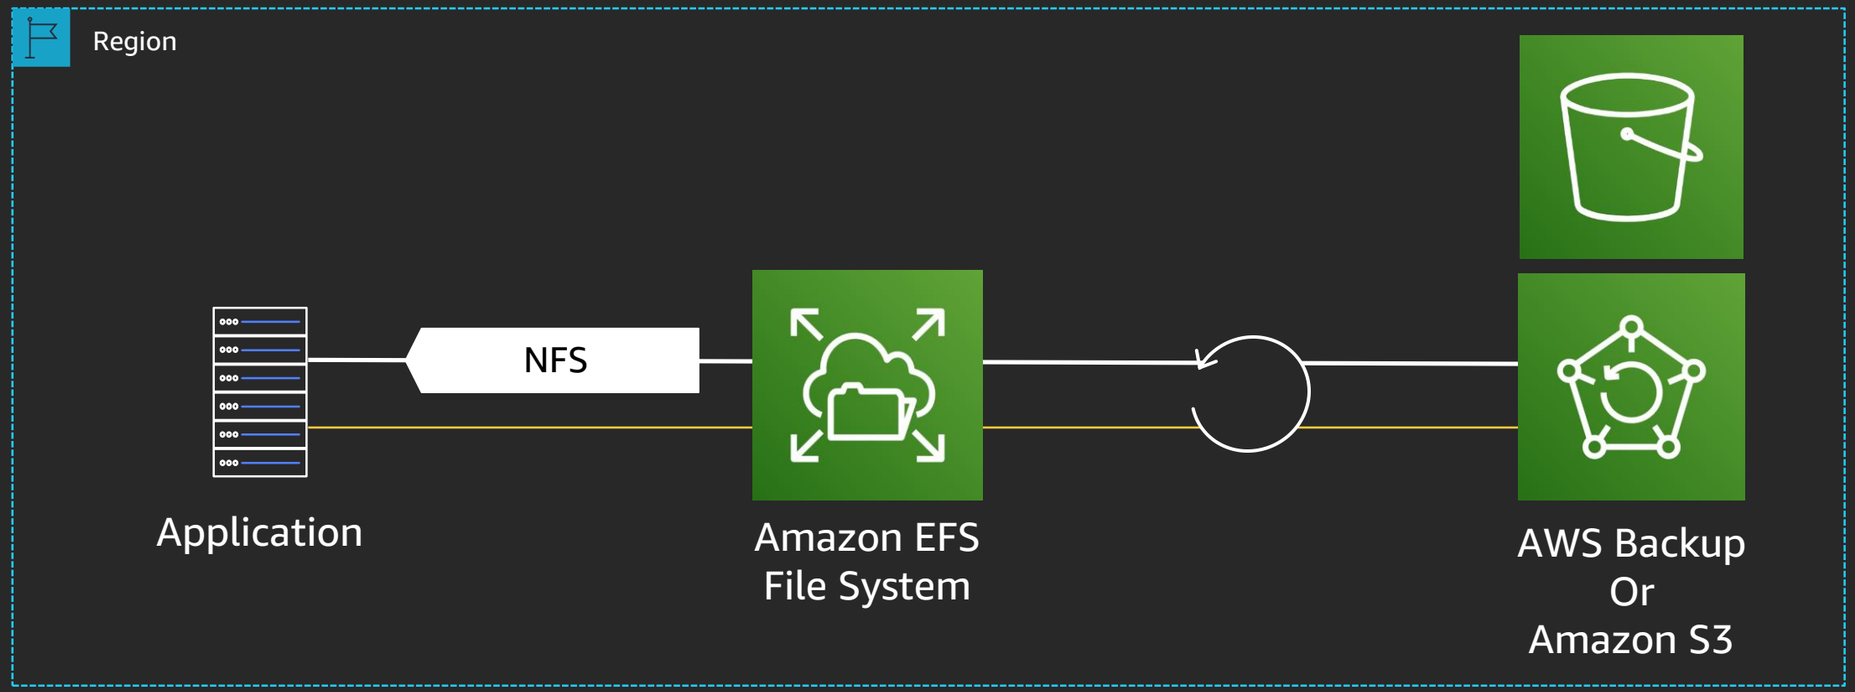

Use cases

- Backup in cloud databases for warm recovery
- Utilize database backup tools from Oracle, IBM, SAP
- Move to Amazon EFS cold storage backup using AWS Backup or copy to S3

Sysco™

Partner enabled solutions

Backup from on-premises to AWS

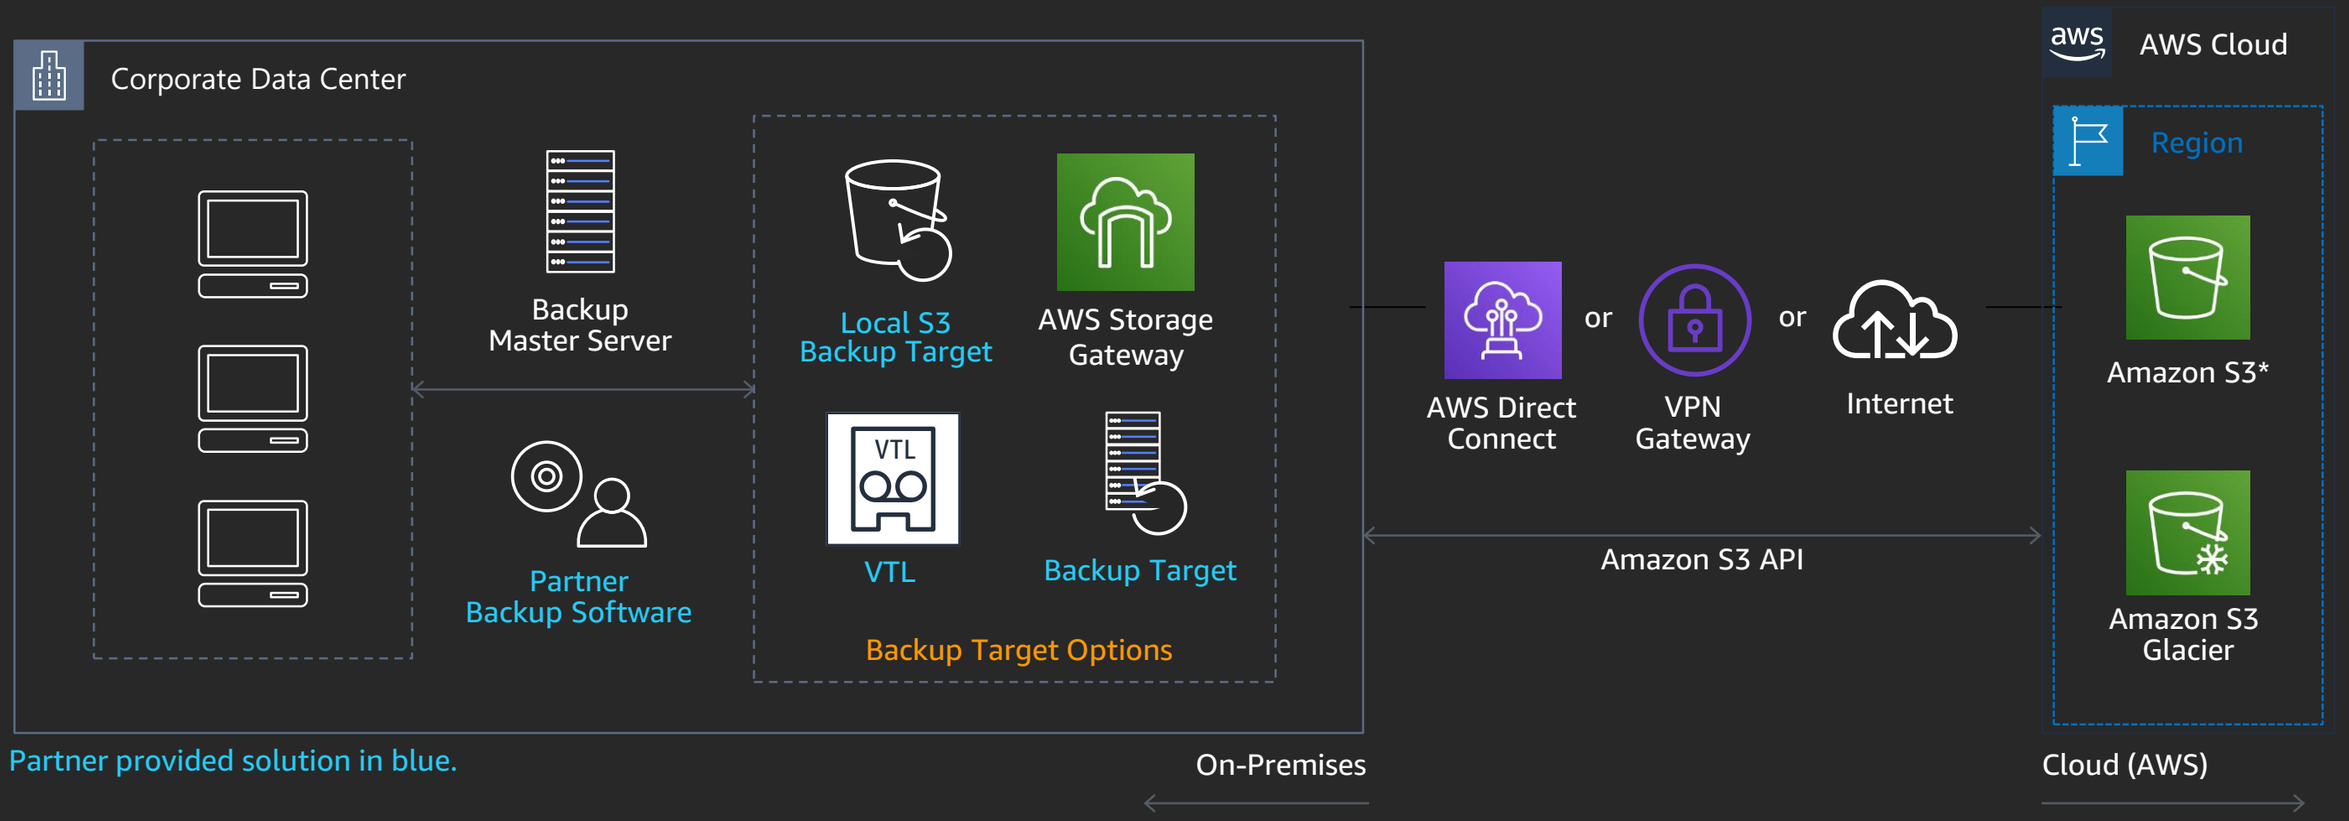

- Backup data flows from on-premises into AWS object storage services
- One of the most common use cases

Partners

CloudBerry Lab cloud storage tools · ctera · IBM · CLUMIO · PURESTORAGE
COMMVAULT · DELL EMC · NetApp · VEEAM
COHESITY · druva · rubrik · VERITAS

Backup on-premises remote and branch offices to AWS

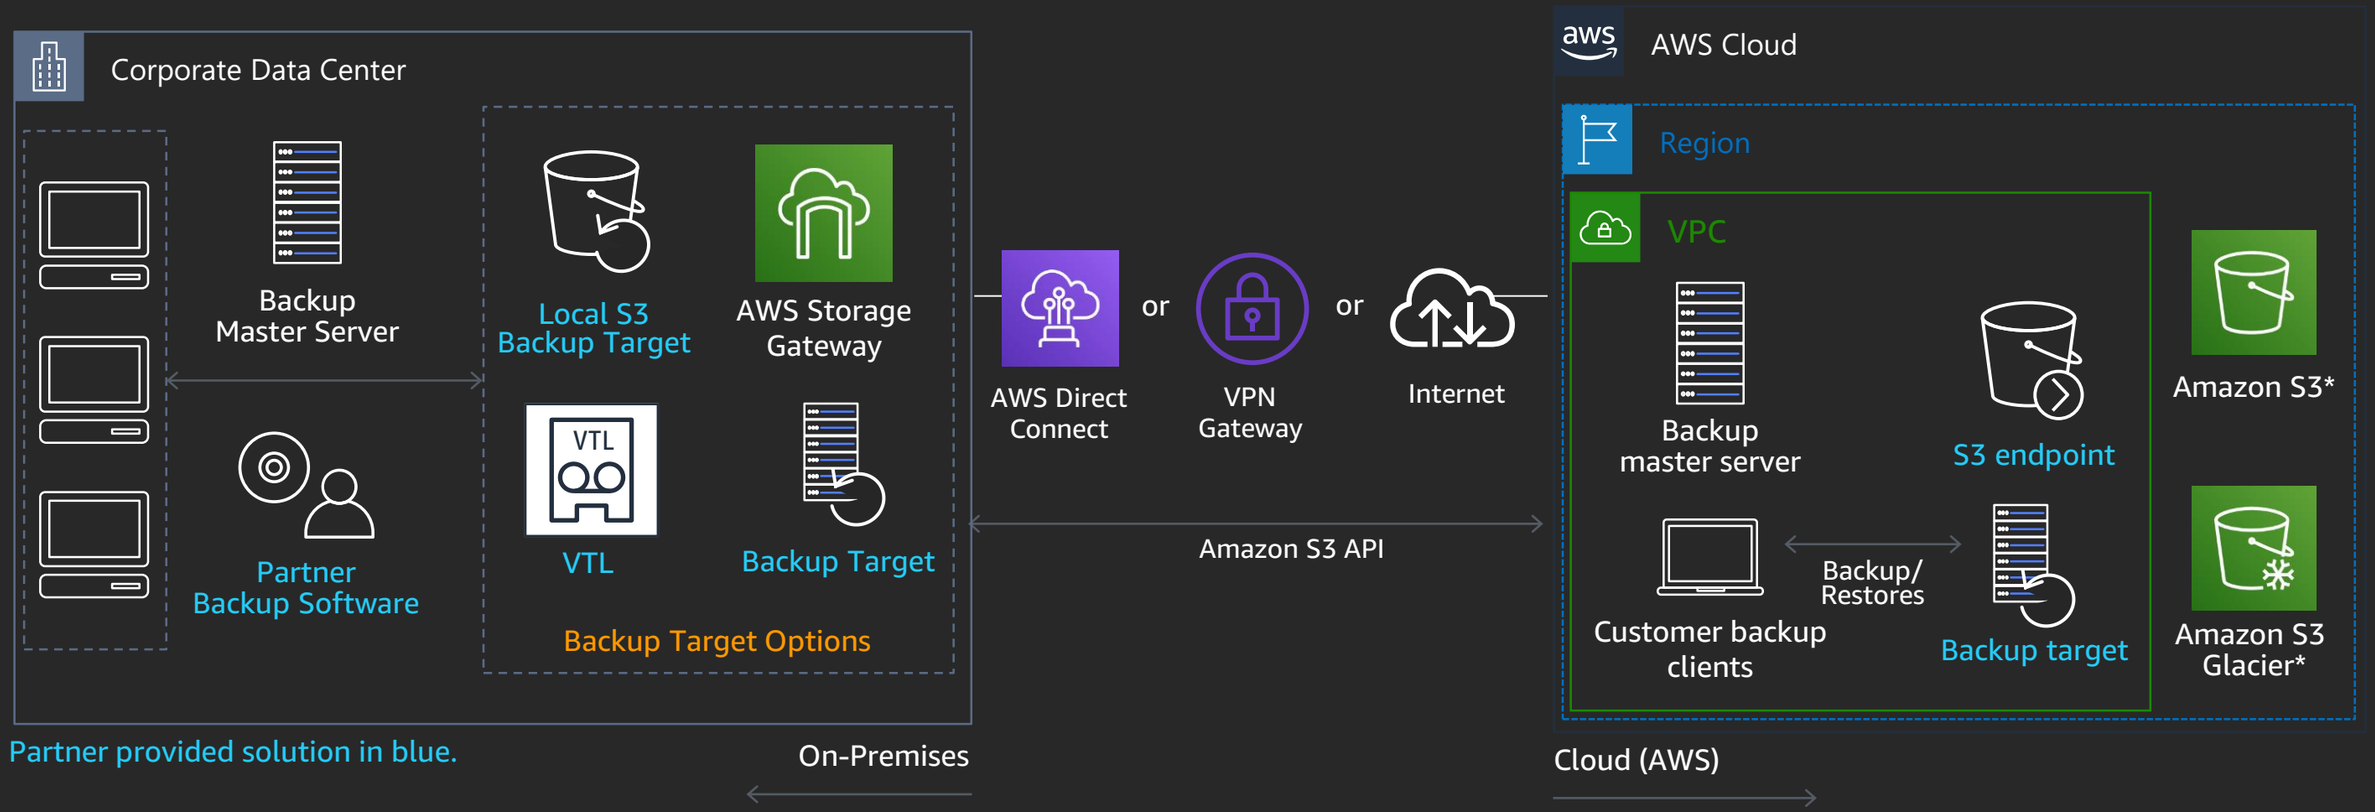

- Provides backup targets for on-premises while having central control in AWS
- Central backup server is protected from individual site issues
- Restores can be quickly redirected to other remote sites, or even within AWS

Partners

DELL EMC · PURESTORAGE · rubrik · COMMVAULT · druva · VEEAM · COHESITY · VERITAS · CLUMIO

SaaS backup model from on-premises to AWS

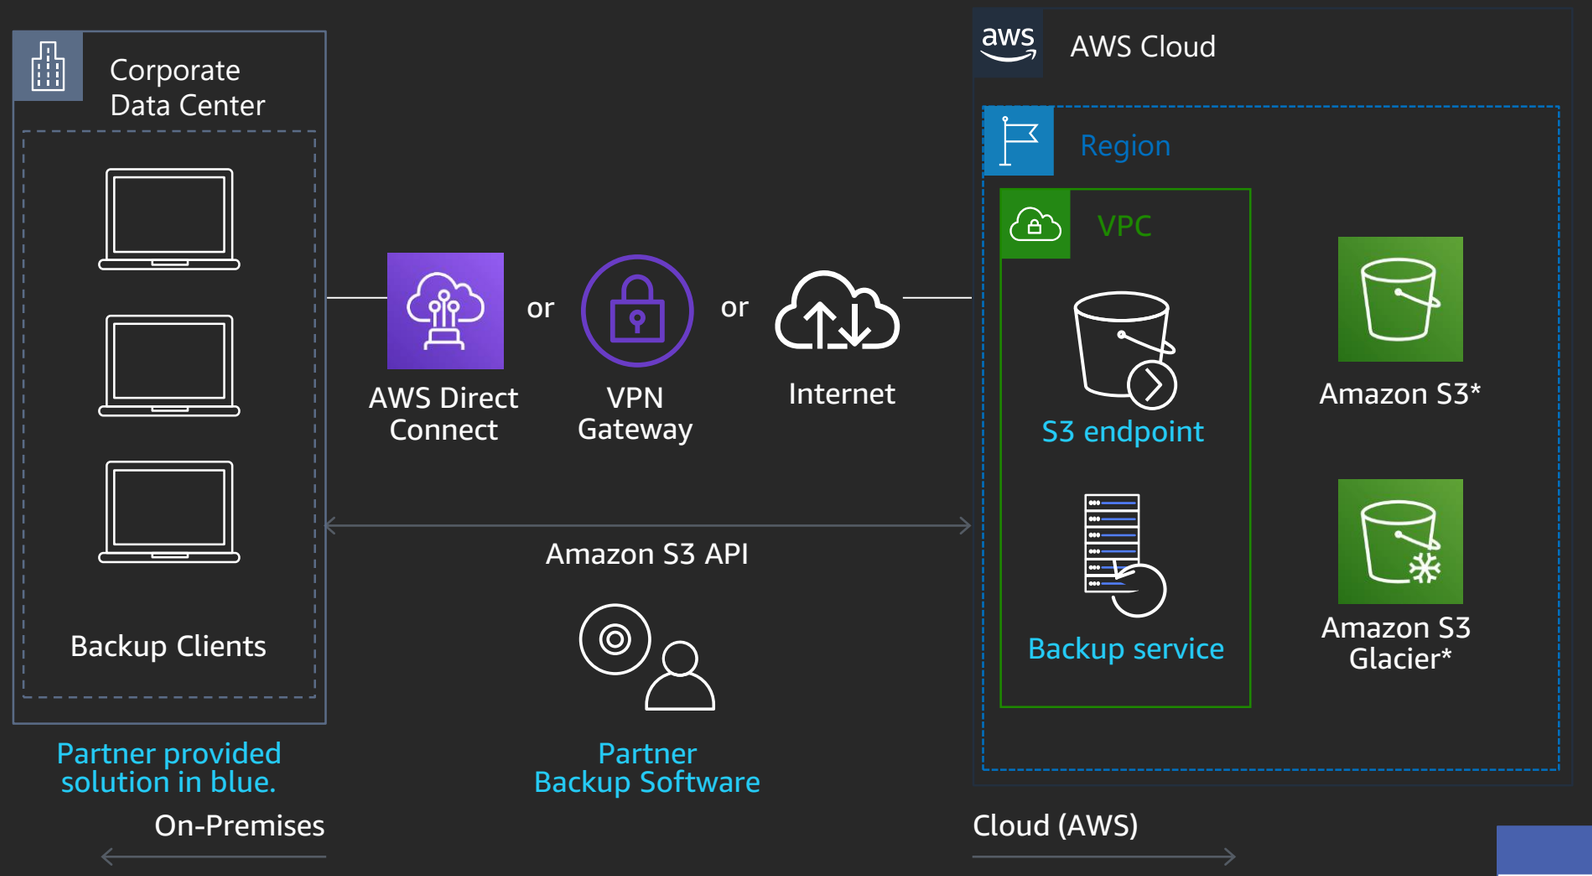

- Utility pricing
- No infrastructure requirement
- Scales on demand

Partners

COMMVAULT

druva

CLUMIO

Back up Amazon EC2 instances with on-premises hosted control

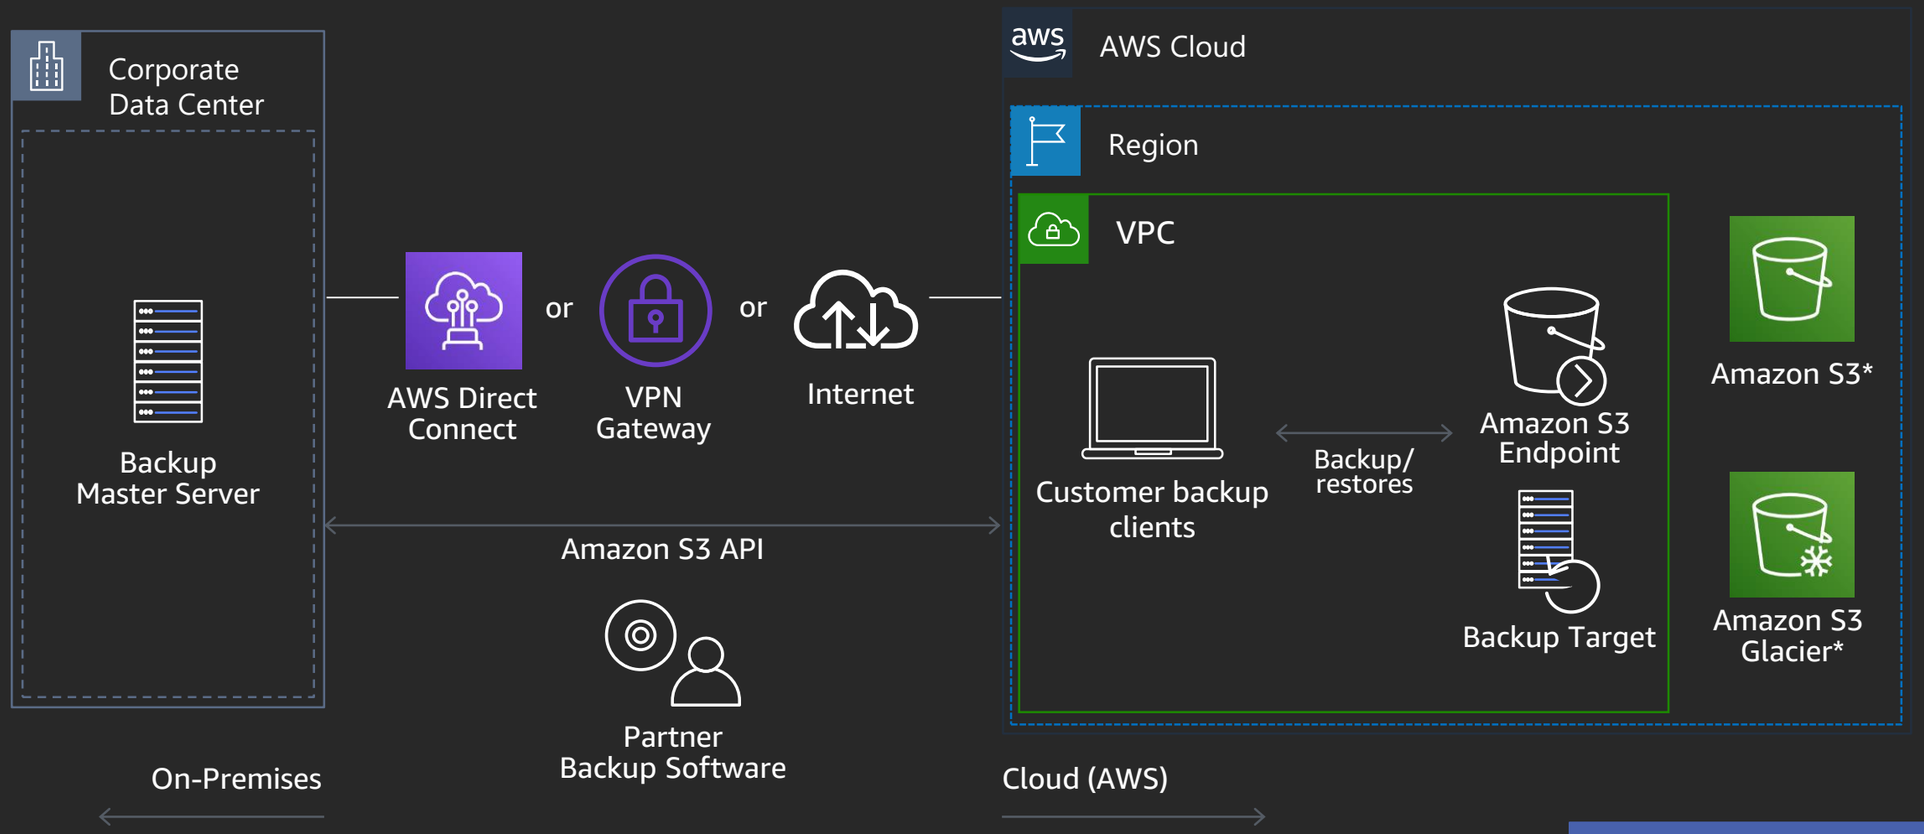

- Maintains local control of backup environment, including current operational practices
- Many AWS Regions can be backed up using a single central site
- Both AWS and on-premises environments can be controlled in a single pane of glass

*And associated Amazon S3/Amazon S3 Glacier storage classes

Partners

COMMVAULT

DELL EMC

VERITAS

Back up Amazon EC2 instances with AWS-hosted control

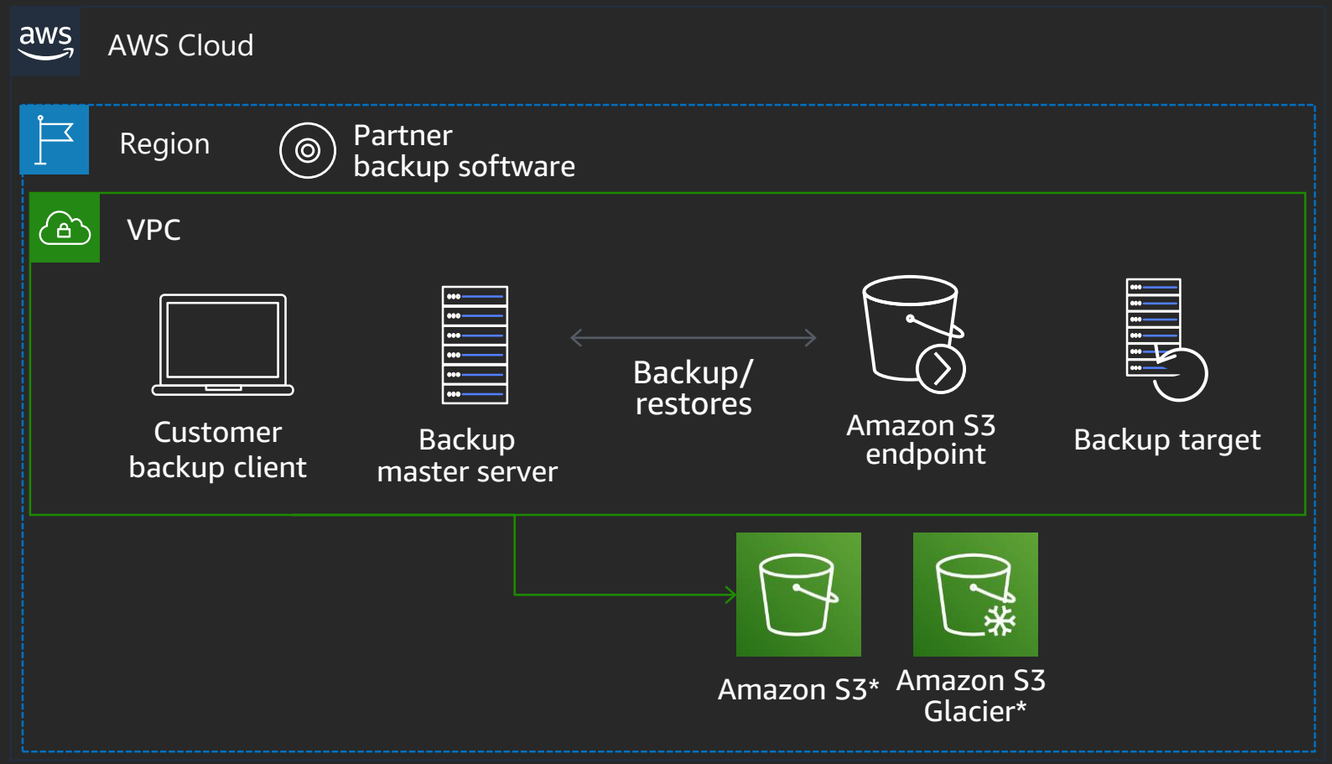

Partner provided solution in blue.

- Fully encapsulated within AWS, also protects Amazon EC2 and Amazon EBS
- Requires no external resources
- Easy to expand and provision

Partners

ctera · rubrik · PURESTORAGE
DELL EMC · NetApp · VERITAS
COHESITY · druva · CLUMIO

AWS Backup: centralize compliance, automate backup, work across services

AWS Backup

1. Simplified backup scheduling and lifecycle management across AWS services
2. Centrally manage backup activities, security, and reporting
3. Achieve consistency and meet compliance requirements

Amazon EBS

Amazon EFS

Amazon RDS

Amazon DynamoDB

AWS Storage Gateway

AWS Backup and APN backup partners

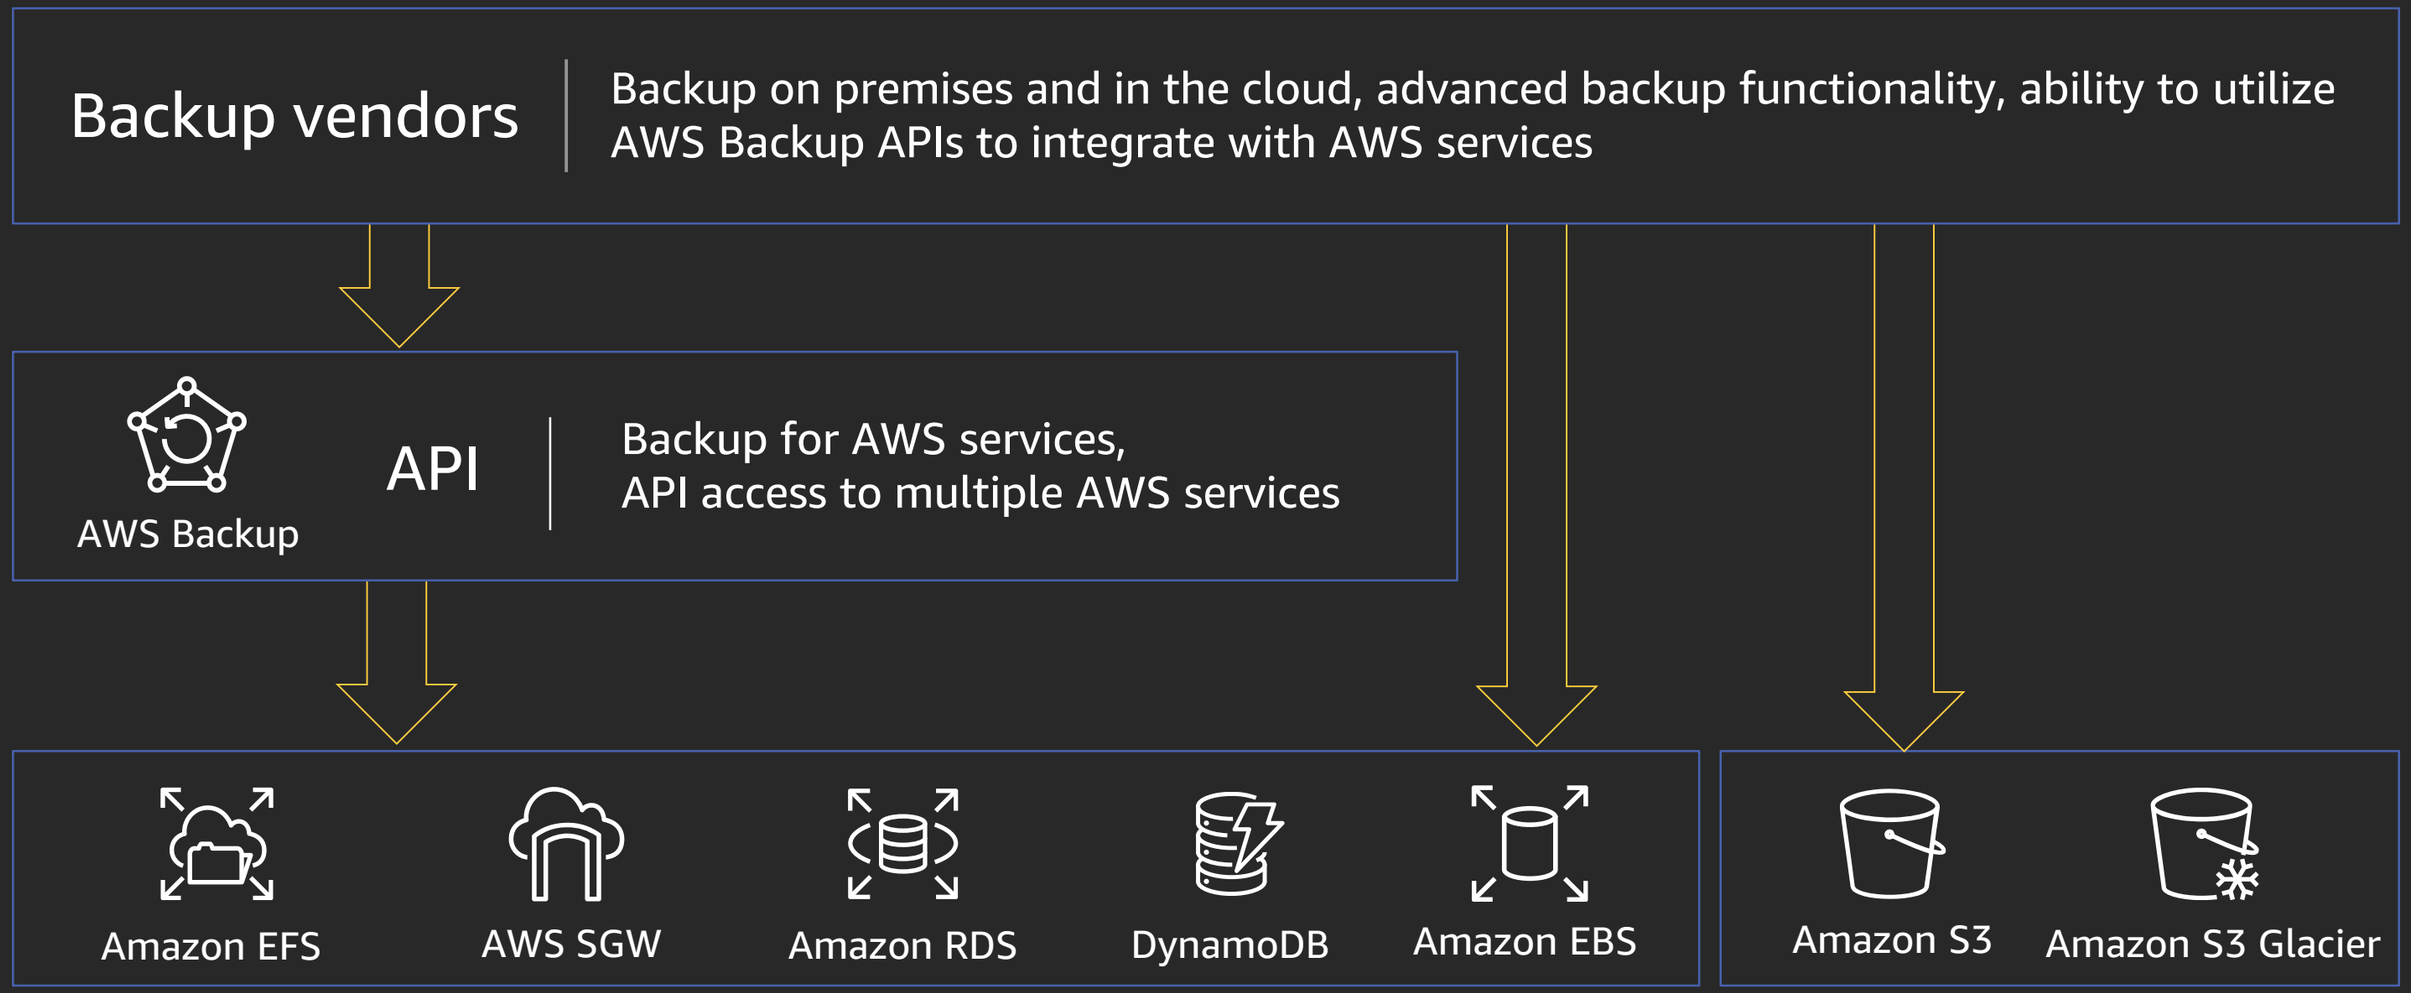

Disaster recovery

Every business requires data protection

Natural disasters

Infrastructure availability and access

Rogue actors

IDC estimates that up to half of all organizations would be unlikely to survive if hit by a disaster that rendered their datacenter unusable. In part, this is because, on average, fewer than half (48.4%) of applications are covered by a disaster recovery (DR) plan.

Source: IDC, 2019

Modernizing DR architectures using cloud

API driven

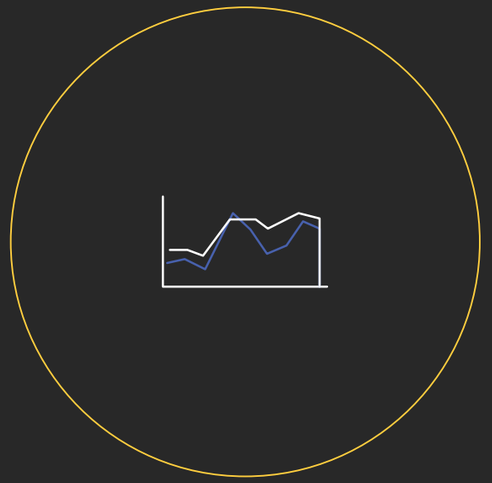

Low overhead

Variable cost

Cloud disaster recovery benefits

Traditional disaster recovery

- Massive upfront investment in hardware & high ongoing costs
- Data growth increases hardware and operational costs
- Hard to actually test without business disruption
- Management and infrastructure overhead for globally distributed businesses
- Vulnerable to cyber threats/hacking

DR in the cloud

- No upfront hardware investment needed
- Only pay for rightsized compute/storage when actually needed
- Lower IT management overhead
- (MUCH!) More automation
- Easy and repeatable testing
- Systems up in minutes (not hours/days!)

AWS offers four levels of backup and DR support across a spectrum of complexity and time

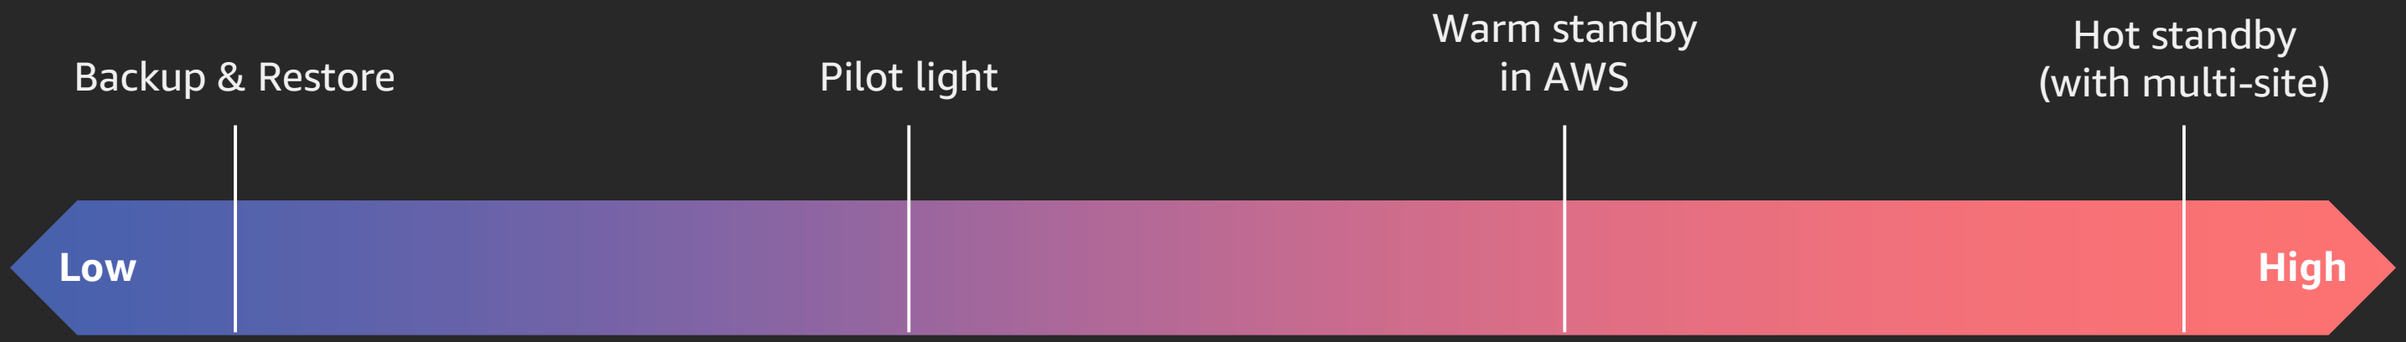

RPO/RTO: Hours

- Lower priority use cases
- Solutions: Amazon S3, AWS Storage Gateway
- Cost: $

RPO/RTO: 10s of Minutes

- Meeting lower RTO & RPO requirements
- Core services
- Scale AWS resources in response to a DR event
- Cost: $$

RPO/RTO: Minutes

- Solutions that require RTO & RPO in minutes
- Business critical services
- Cost: $$$

RPO/RTO: Real-time

- Auto-failover of your environment in AWS
- Cost: $$$$

Start with requirements

Identify applications to protect

Business impact analysis

Define RPO and RTO requirements

Compliance considerations

Availability concepts

High availability

Keep your applications running 24x7

Backup

Make sure your data is safe

Disaster recovery

Get your applications and data back after a major disaster

High availability—keeping your applications running 24x7

Between AZs

Amazon RDS HA

Amazon ElastiCache

Amazon S3

Amazon Aurora

Application level replication

EC2 instance recovery

Between Regions

Application level replication

Third party

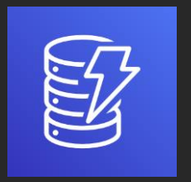

Amazon DynamoDB

Pilot light

Pilot light prep

www.example.com

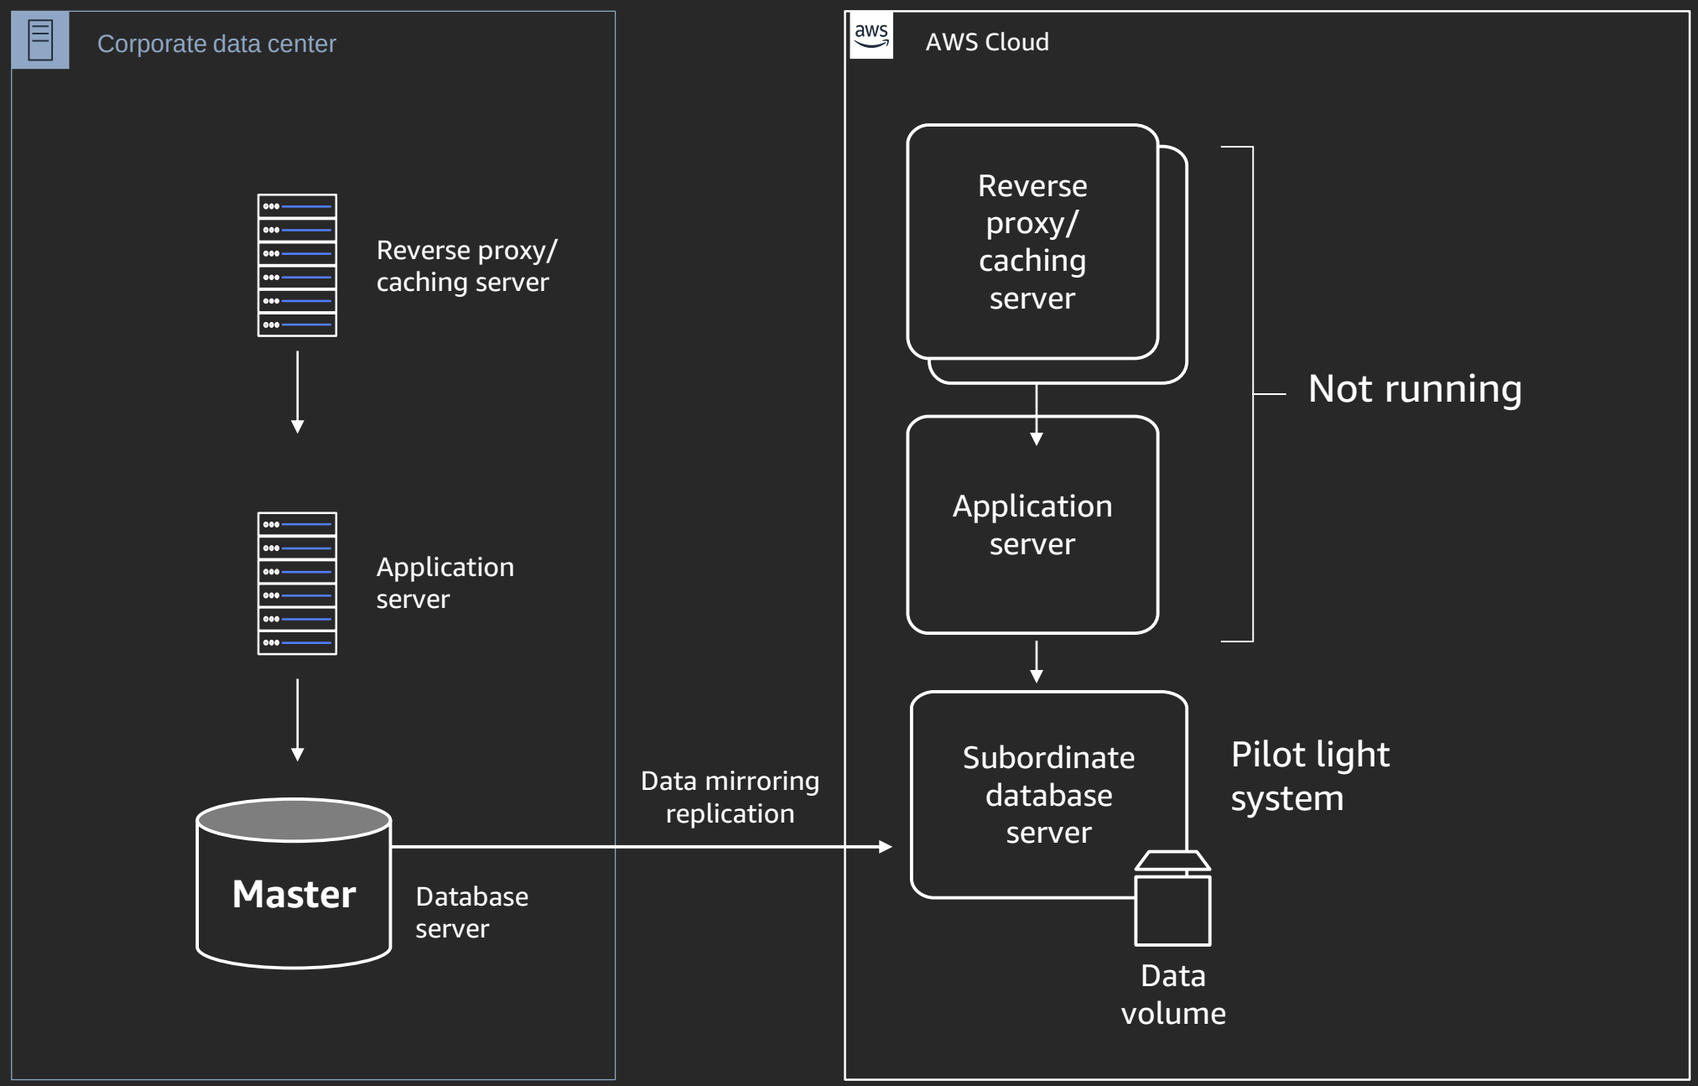

Pilot light recovery

www.example.com

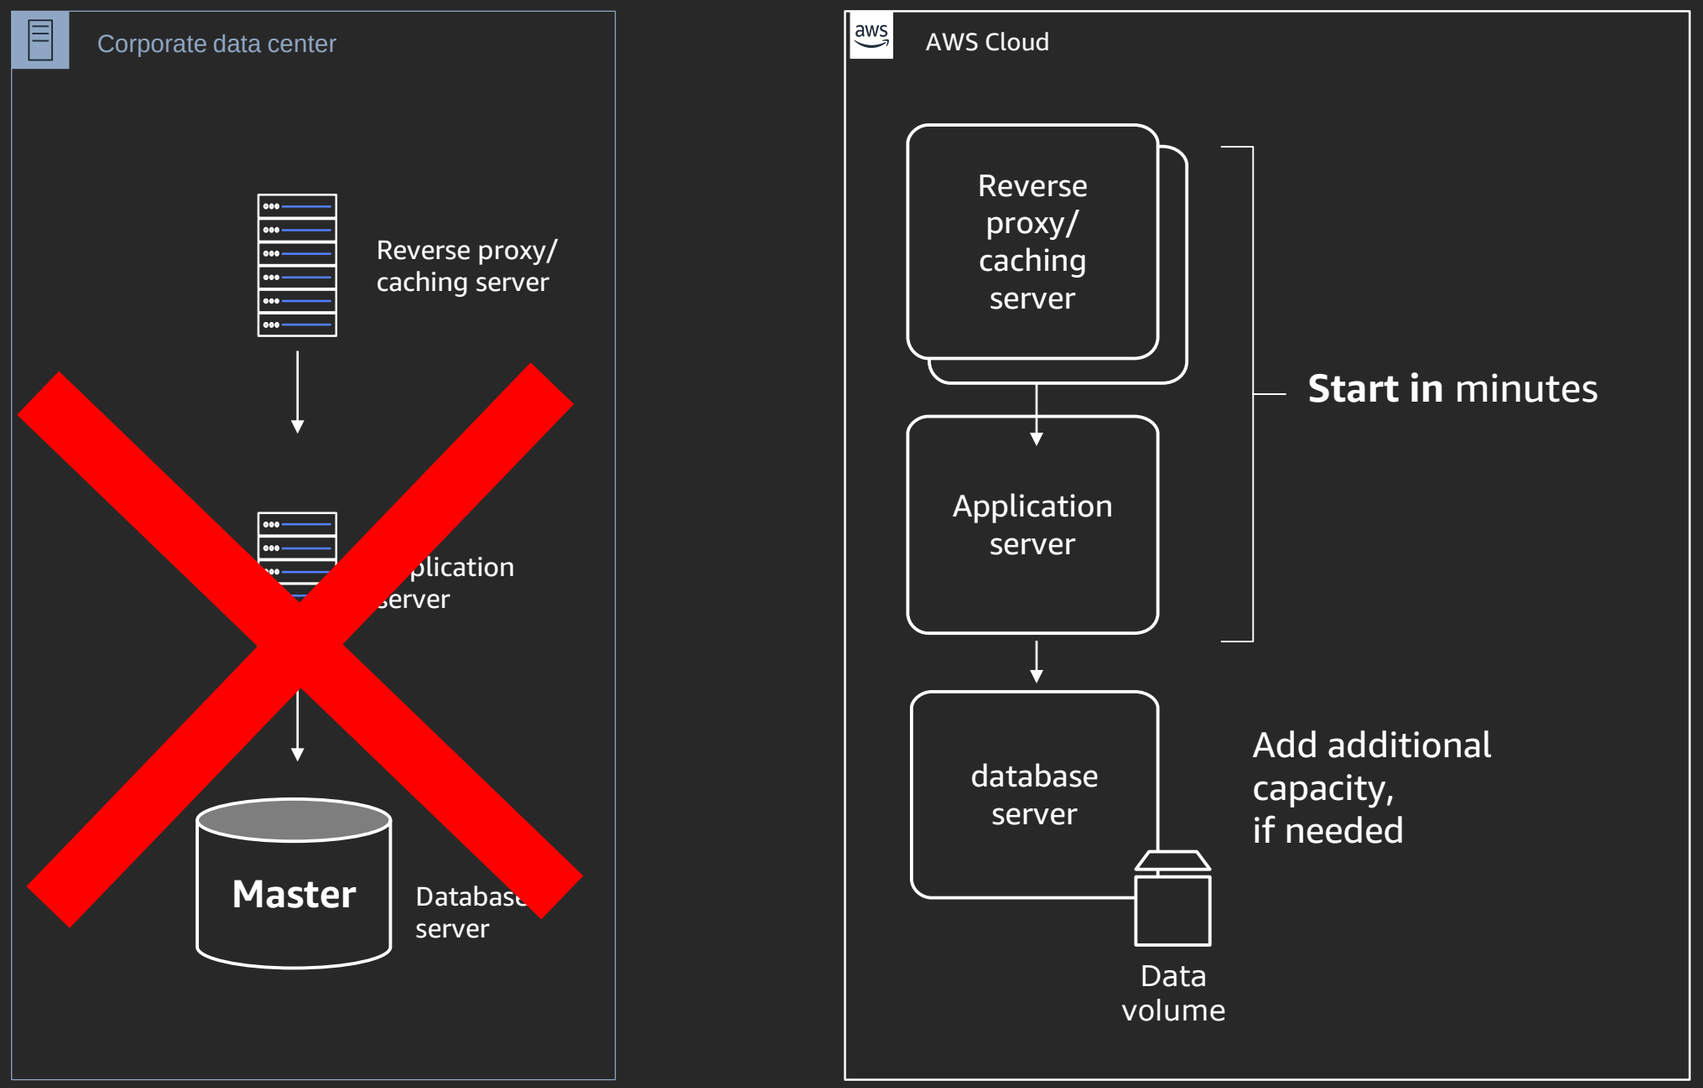

Warm standby

Warm standby prep

www.example.com

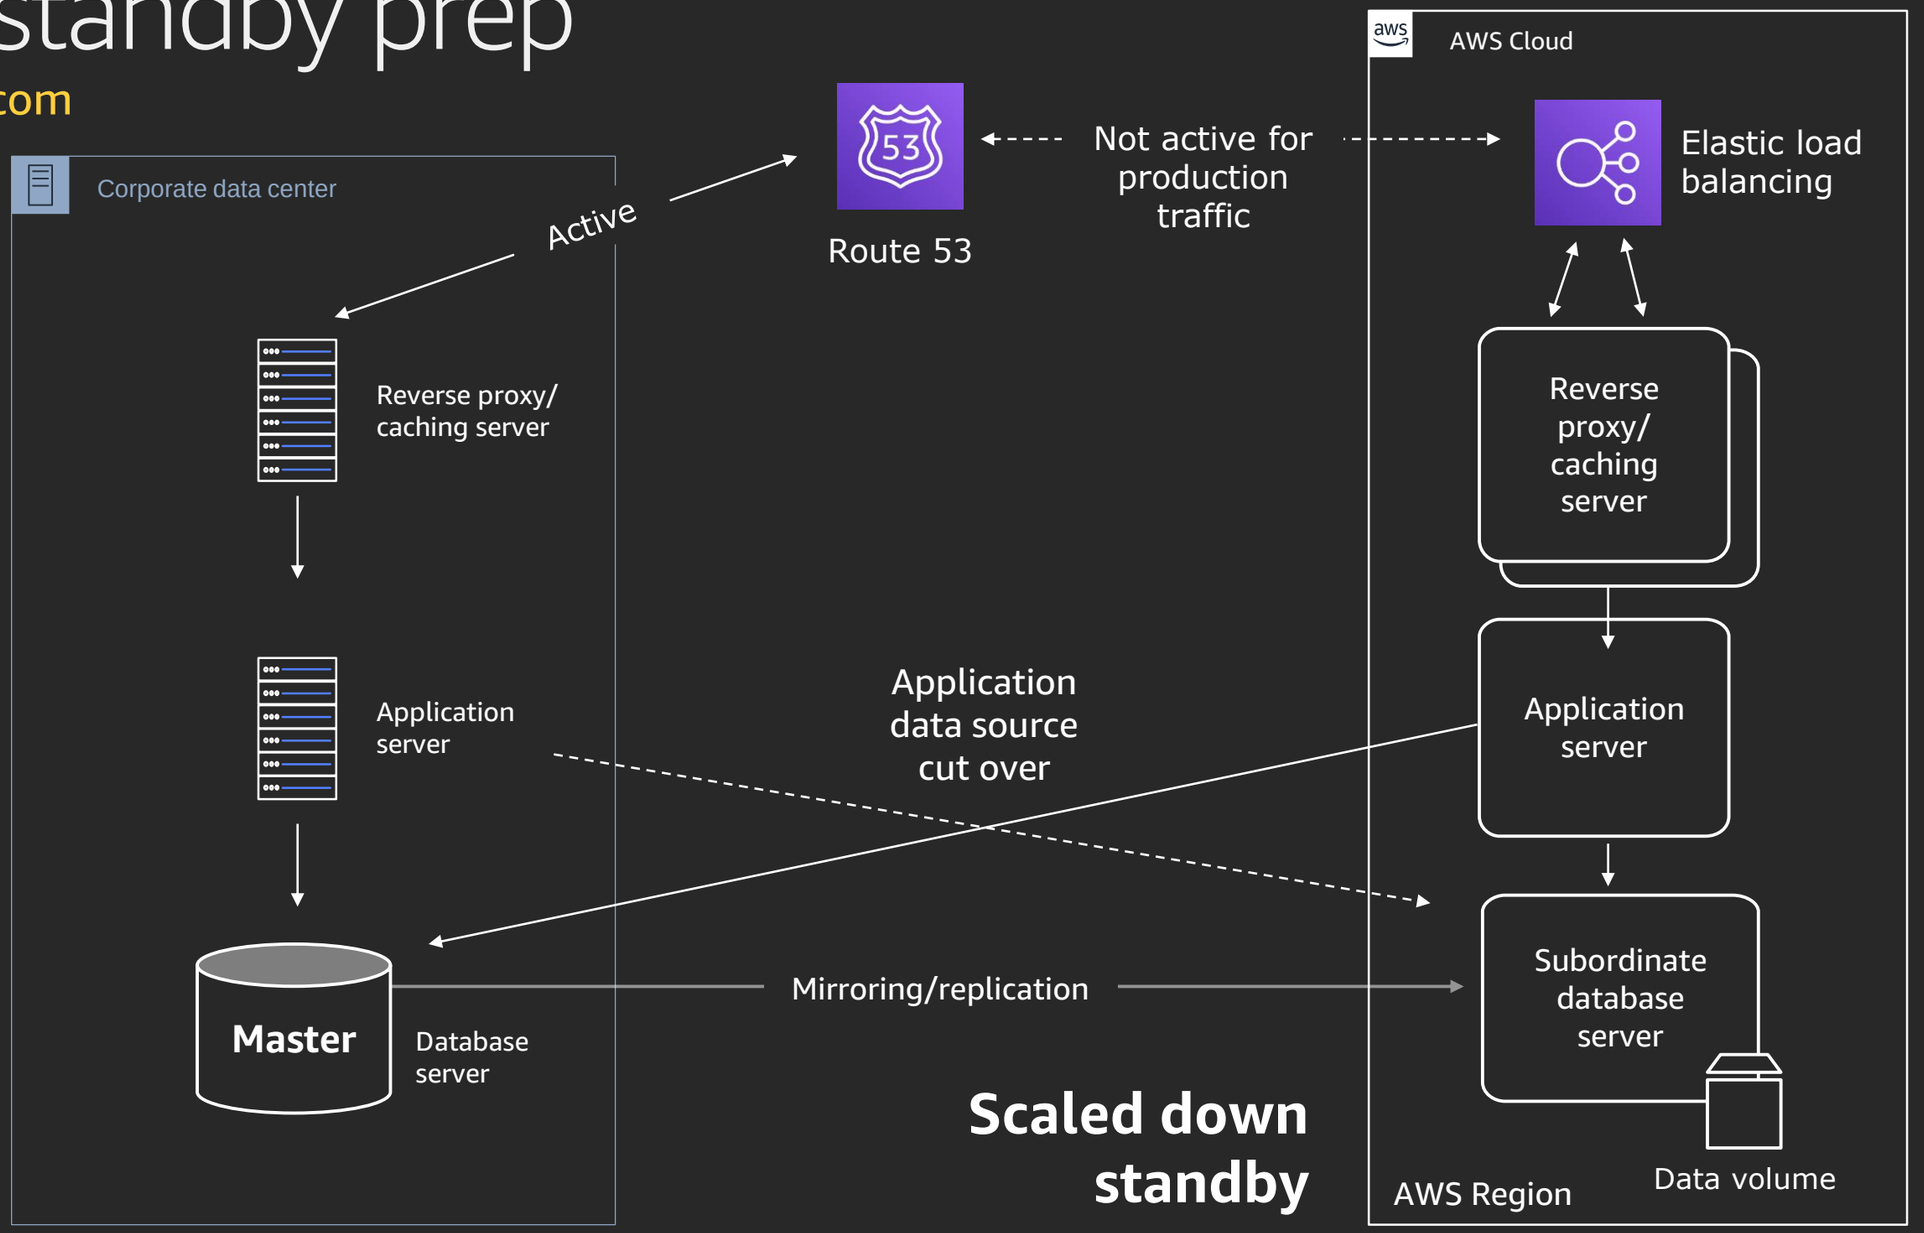

Warm standby recover

www.example.com

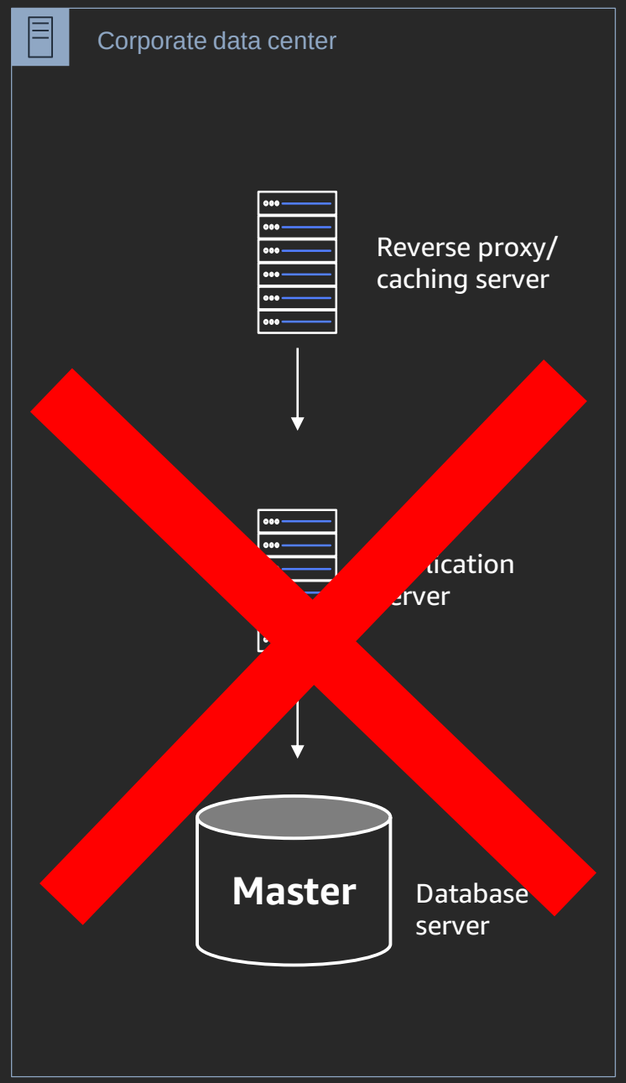

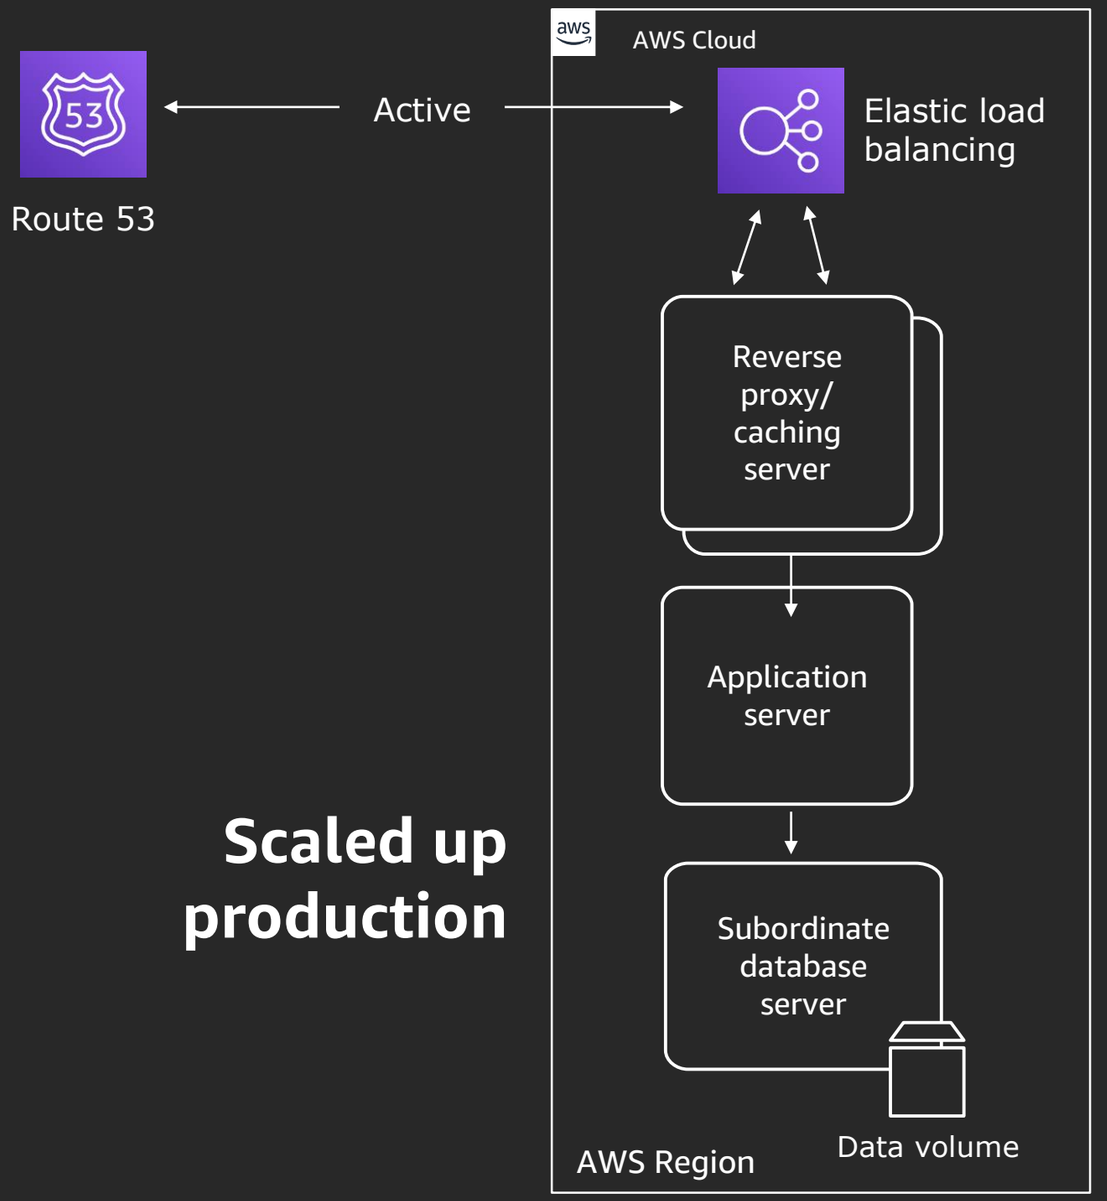

Scaled up production

Hot site

Hot site prep

www.example.com

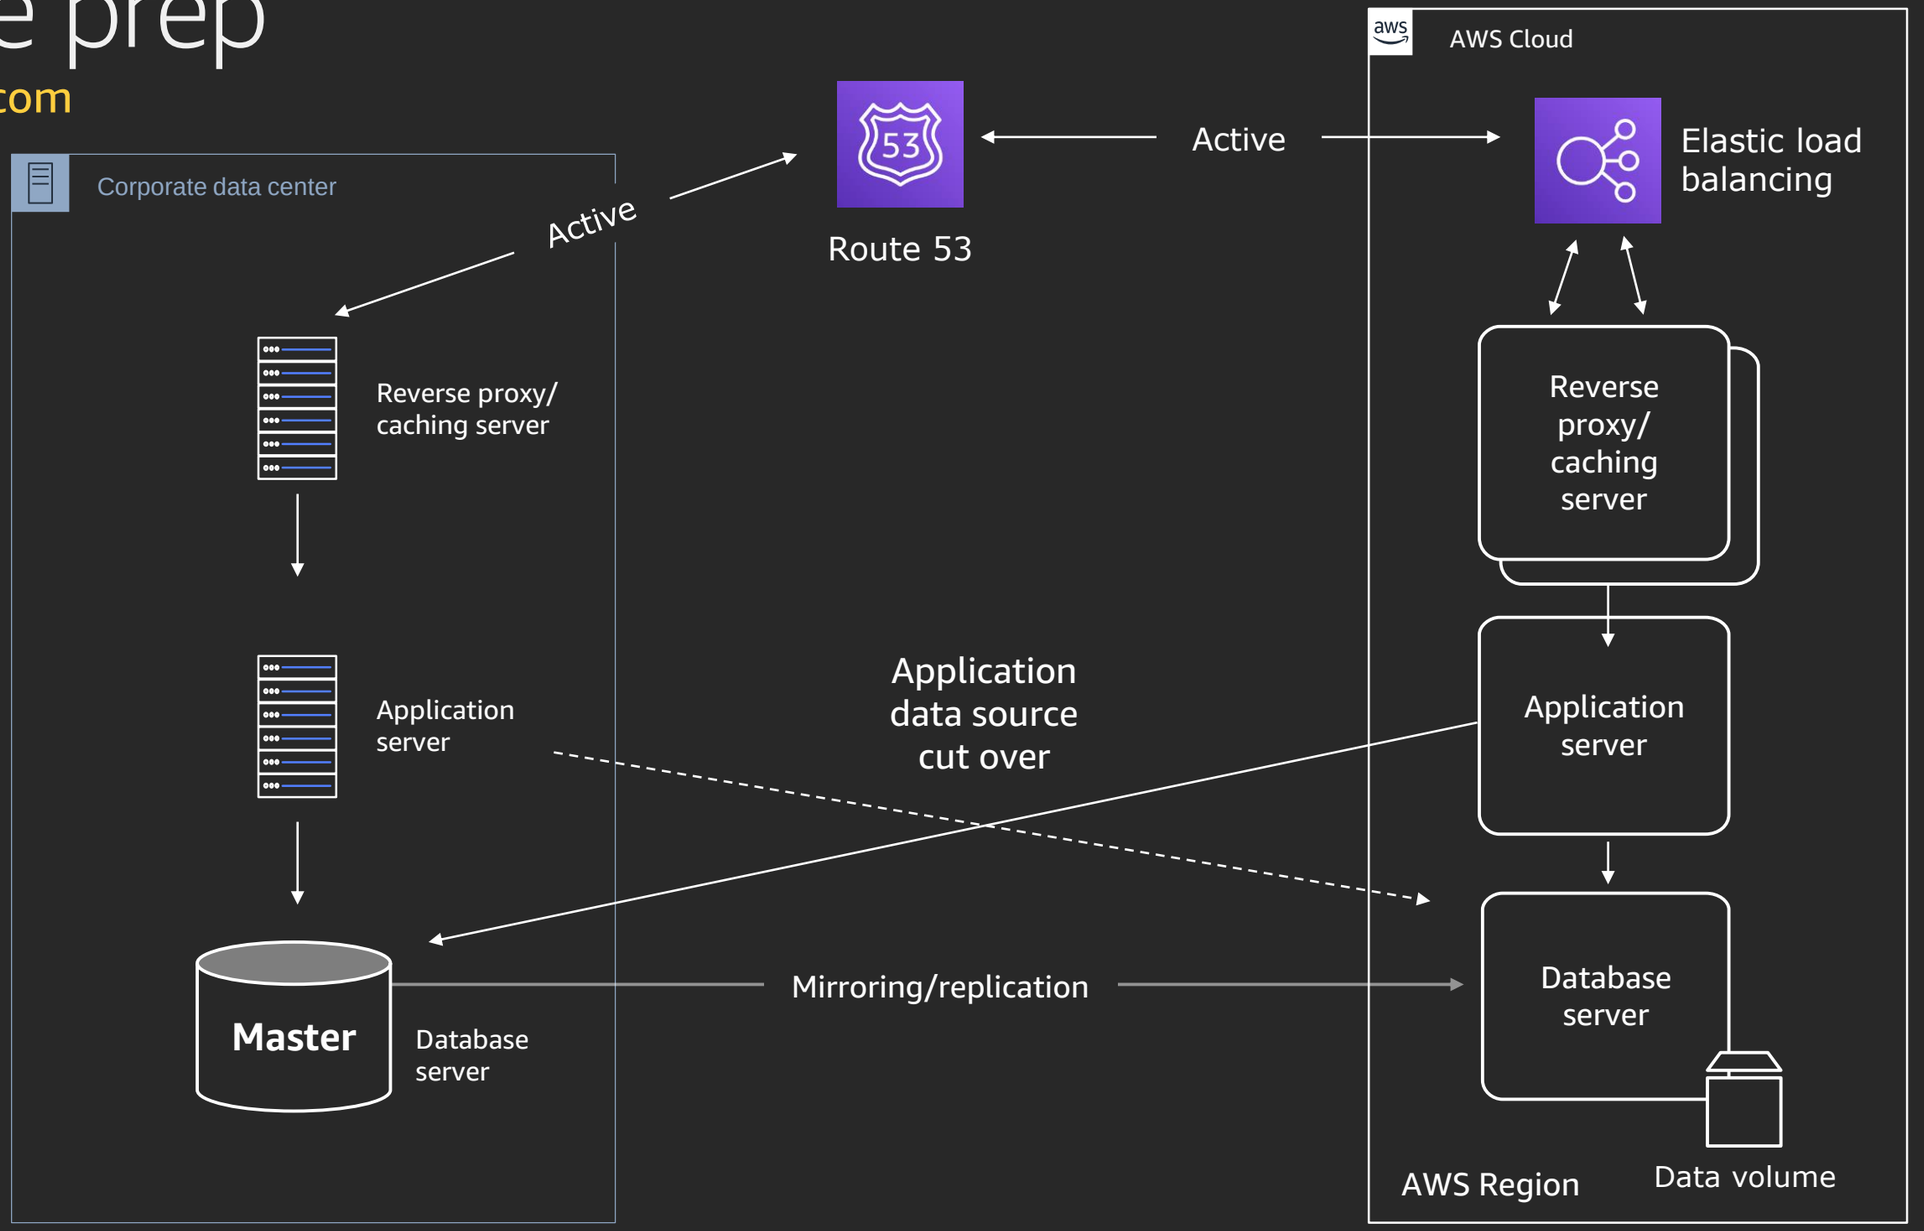

Hot site recovery

www.example.com

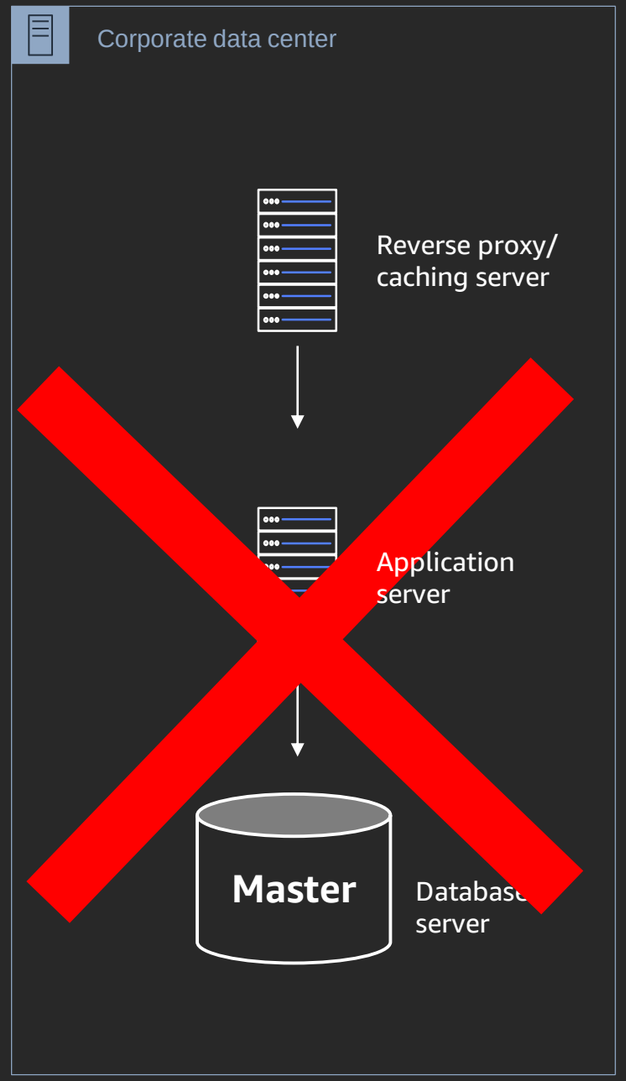

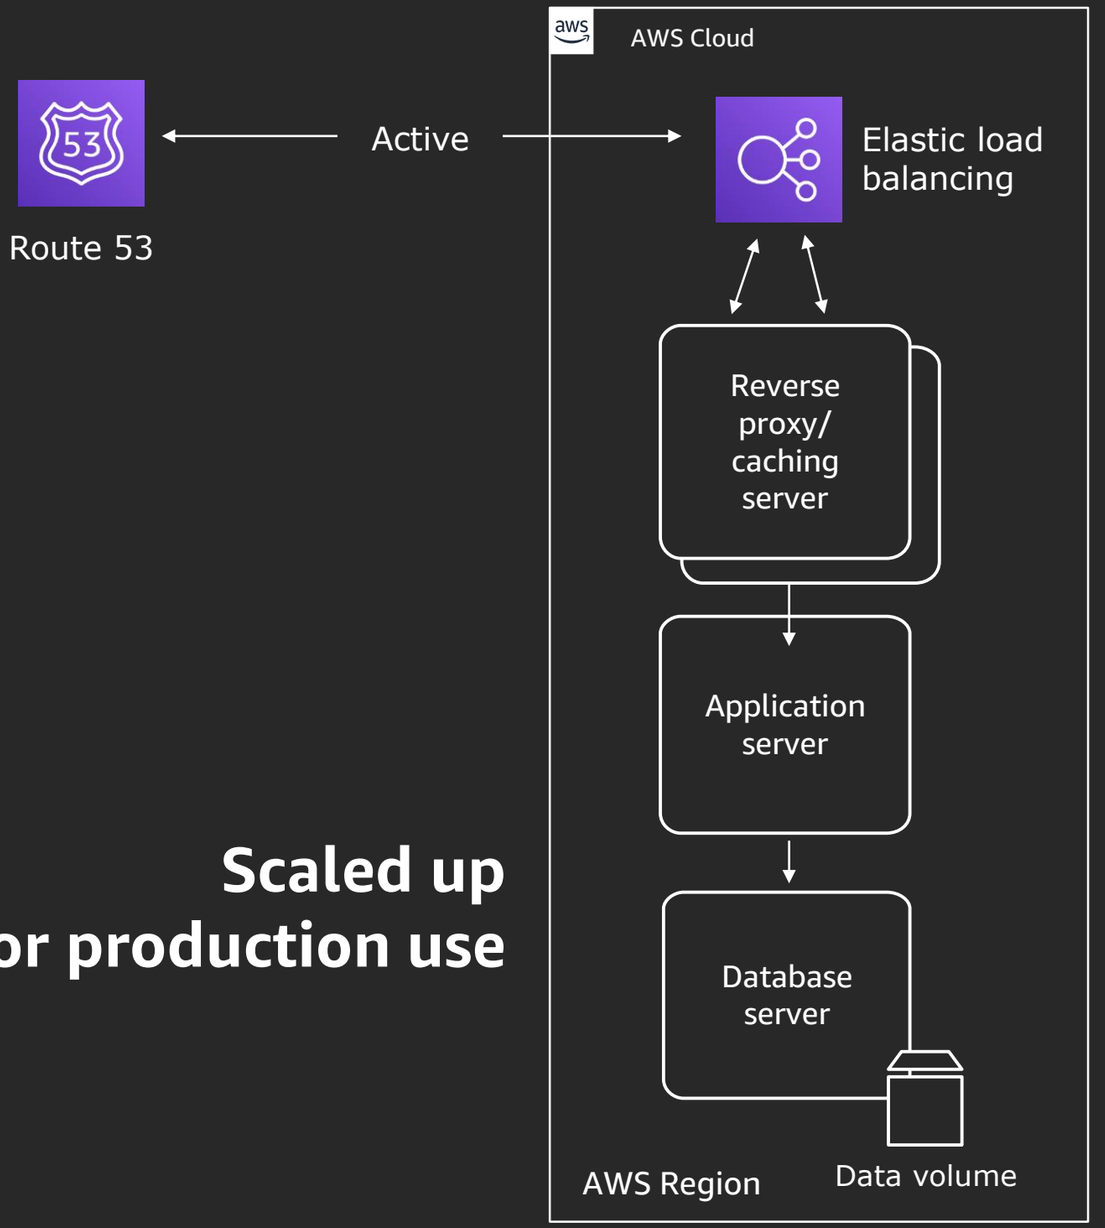

Scaled up for production use

CloudEndure: Better, faster, more affordable disaster recovery

Flexible	Reliable	Highly automated
 Replicate from any source	 Robust, predictable, non-disruptive continuous replication	 Minimal skill set required to operate
 Wide range of OS, application, and database support	 RPO: subsecond RTO: minutes	 Easy, non-disruptive DR tests
 Failback to cloud/on-prem	 Protection against ransomware, corruptions, and human errors	 Automated lightweight staging area reduces TCO

- Improve recovery objectives & reduce TCO
- Simple setup lets you start in minutes
- Same highly automated process for all workloads
- Minimizes complexity and reduces risk
- Easy failover and failback

CloudEndure: how does it work?

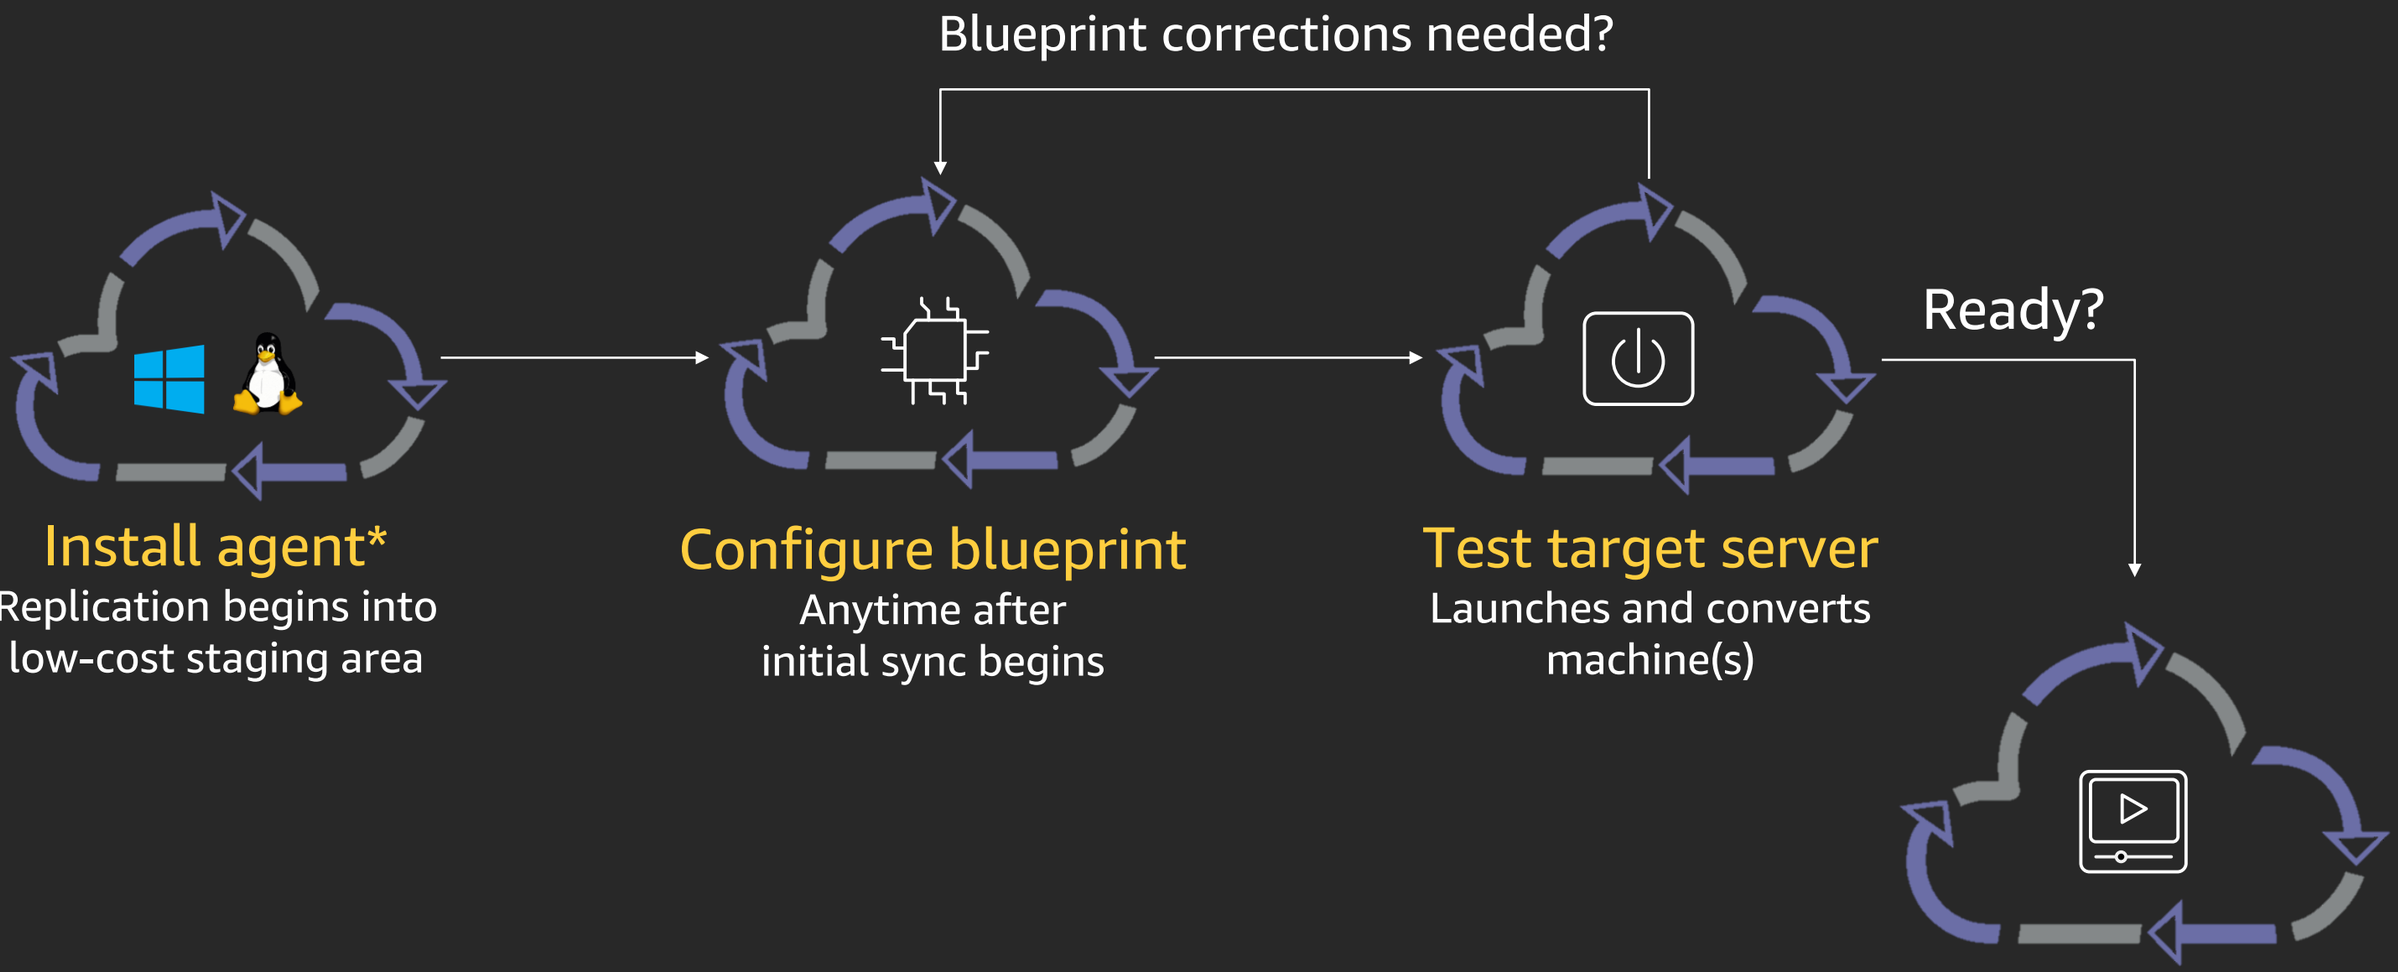

*No reboot, No performance impact, No application configuration

**May be modified anytime after the CloudEndure agent is installed

Lightweight staging

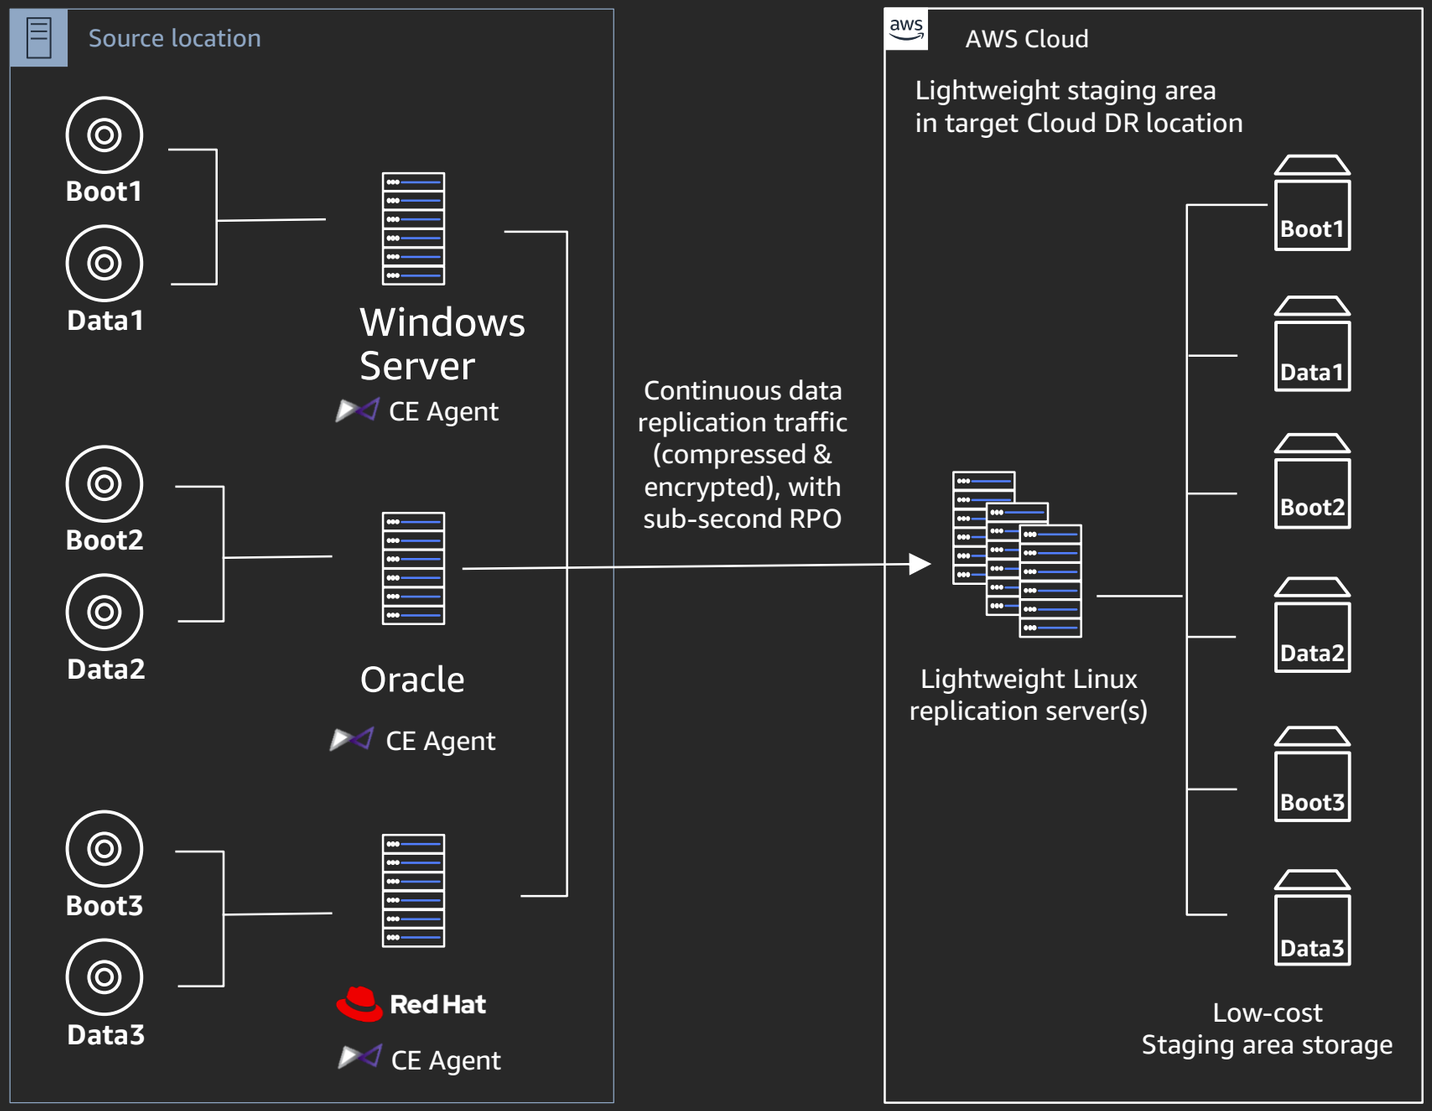

- Reduce DR site compute costs by 95%+
- Reduce DR site storage costs by 70%+
- Zero DR site duplicate OS license fees!
- Zero DR site software/DB license fees!
- Zero DR site networking equipment fees!
- Continuous replication with subsecond RPO

Disaster event or test

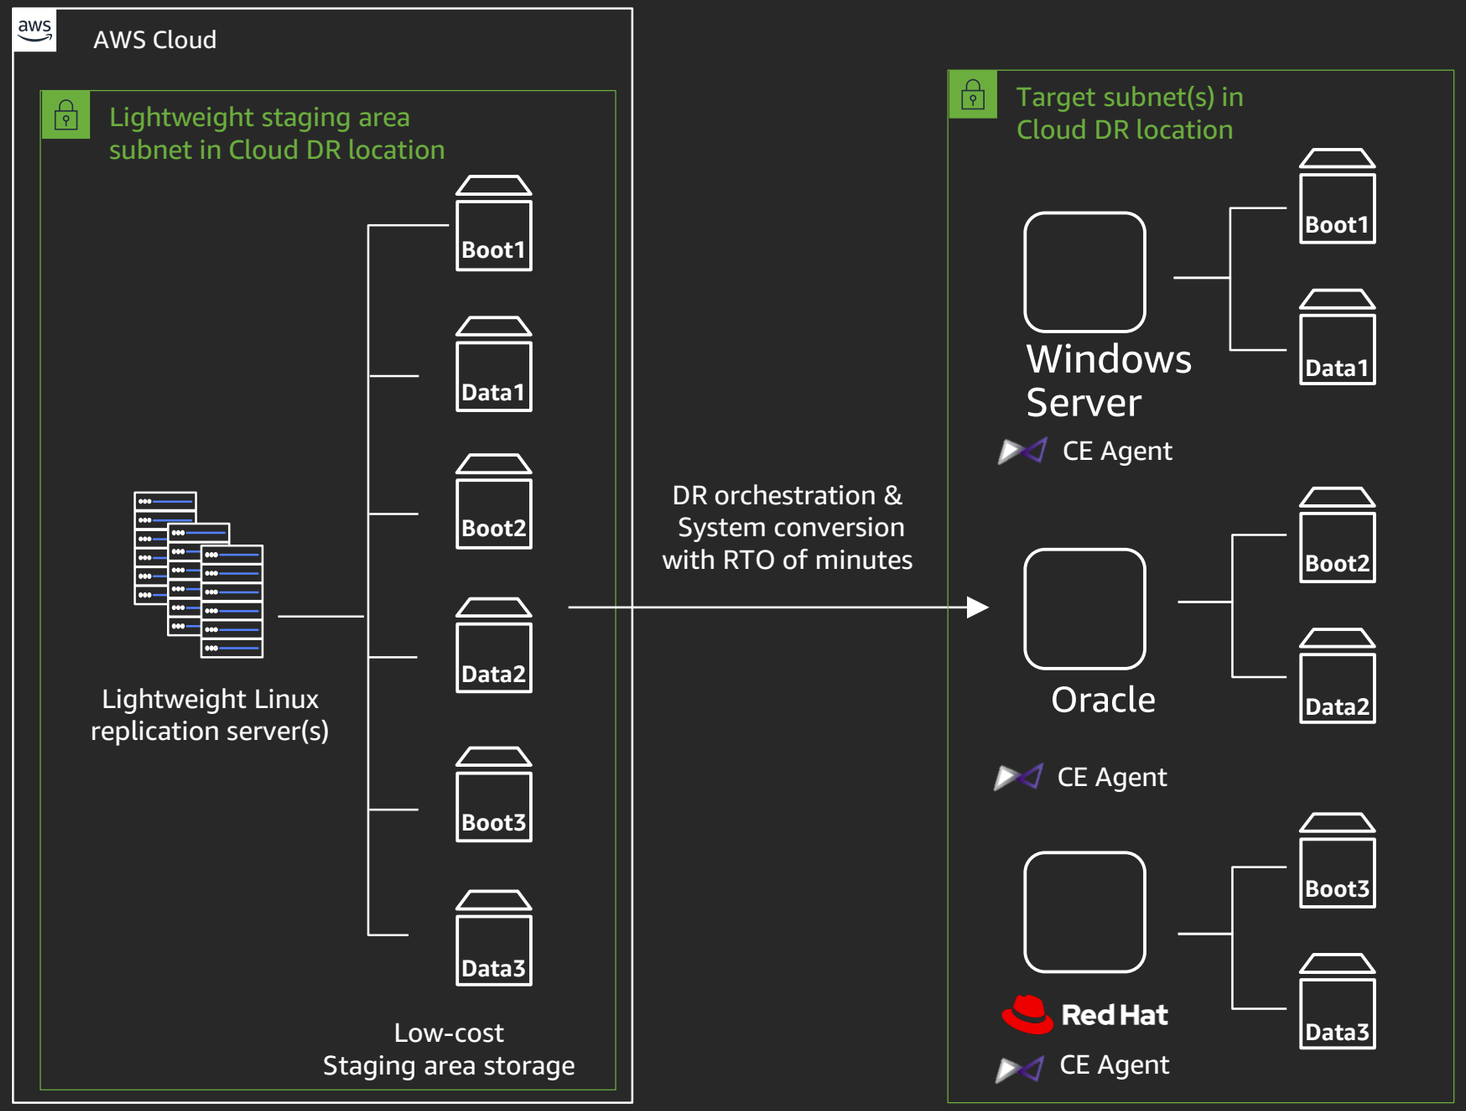

- Rapid machine recovery (RTO of minutes)
- Self-service DR dashboard
- Unlimited free non-disruptive DR tests
- Built-in fail-back to any infrastructure
- Enable one-click future migration
- Enable cross-region/cross-cloud DR

About Us

BankUnited Inc. (NYSE: BKU), headquartered in Miami Lakes, Florida, is a financial services company with $33.0 billion in total assets as of September 30, 2019.

About Me…

Juan Mejia
VP, Data Center Manager

- 24 Years of IT experience
- 18 Years at BankUnited
- Manage Storage, Hypervisor, and IT Business Continuity

Pre DR Modernization Architecture, Challenges, & Objectives

AWS re:Invent

aws

Prior DR Architecture

- Two (2) physical data centers
  - Primary – Miami, FL
  - Secondary – Long Island, NY
- Multiple forms/methods of replication
  - vSphere
  - Storage Array (ZFS, NFS, etc.)
  - Database Replication
- High touch failover
  - Documented manual procedures
  - High complexity
  - Requires heightened levels of coordination between the business and IT

Challenges

The primary challenges to the Bank’s data center topology are as follows:

Environmental

Environment prone to hurricanes and flooding

Geographical

Physical locations situated in direct flight paths and close to major highways

Facilities

Facilities requiring additional Investment (power, HVAC, etc.) to support ongoing DR operations

Objectives

When BankUnited decided to move the DR facilities to the cloud, the Bank was working to understand the following objectives:

Scalability

Ability to increase resources accordingly in order to meet customer/business demand

Elasticity

Enable resources to be made available when needed and for only as long as needed in order to meet demand

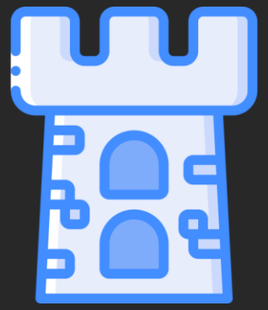

Resiliency

Increase the availability and survivability of the Bank's technology estate

Efficiency

Enable the Bank to realize its objectives without incurring significant idle costs

Enablement

Help enable the Bank's Digital Transformation

Pre DR Modernization Questions & Solutions

Problems to resolve

- Securely transport data to AWS
- Keep the data in sync once the initial seeding is complete
- Avoid paying for unused compute on top of the required storage
- Address multiple replication use cases
- Provision resources once a DR event has been enacted

"We must walk consciously only part way toward our goal, and then leap in the darkness to our success."

Henry David Thoreau

Modernization of DR architecture = CloudEndure

- **Securely transport data to AWS**
  Direct Connect and CloudEndure encryption
- **Keep the data in sync once the initial seeding is complete**
  CloudEndure's Continuous Data Replication engine provides real-time, asynchronous, block-level replication
- **Avoid paying for unused compute on top of the required storage**
  CloudEndure maintains ongoing replication into a low-cost "staging area" in the target infrastructure.

Modernization of DR architecture = CloudEndure

- **Address multiple replication use cases**
  CloudEndure replicates at the OS level enabling support of any type of source infrastructure

- **Provision resources once a DR event has been enacted**
  CloudEndure orchestration engine

- **Convert from on premises VMware hypervisor to Amazon EC2**
  CloudEndure machine conversion handles hypervisor and OS configuration, boot process changes, and guest agent installation

Post DR Modernization Architecture

Post DR Modernization Architecture

Benefits

- Leverage CloudEndure’s blueprint settings
- Single pane of glass for recovery
- Snapshots (48 hours of protection):
  - Every 10 mins an hour for 24 hours
  - Every hour for 24 hours

What changes in DR testing for BankUnited?

- Removed the need for constantly updating scripts to recover storage and VMs
- Removed the dependency on physical hardware

Next Steps

- Leverage CloudEndure for Region to Region replication as workloads now become production on AWS

“There is no substitute for victory.”

Douglas MacArthur

APN partner solutions for disaster recovery

Technology solutions vetted by the AWS Storage competency program

BCDR

COHESITY

druva

NetApp

rubrik

VERITAS

Zerto

Consulting

AHEAD

shi

CompuNet, Inc.

SIRIUS

REAN CLOUD

softchoice

HCL

SUNGARD AVAILABILITY SERVICES

rackspace

ONICA

PRESIDIO Future. Built.

PROACT

infiniti

Conclusion

Using the cloud for backup and disaster recovery provides compelling benefits

AWS provides services and tools for building reliable, fault-tolerant, and highly available systems in the cloud

AWS has partnered with industry leaders in backup and DR to help our mutual customers modernize their practices

Learn storage with AWS Training and Certification

Resources created by the experts at AWS to help you build cloud storage skills

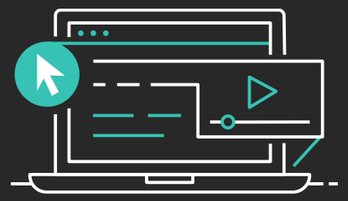

45+ free digital courses cover topics related to cloud storage, including:

- Amazon S3
- AWS Storage Gateway
- Amazon S3 Glacier
- Amazon Elastic File Storage (Amazon EFS)
- Amazon Elastic Block Storage (Amazon EBS)

Classroom offerings, like Architecting on AWS, feature AWS expert instructors and hands-on activities

Visit aws.amazon.com/training/path-storage/

aws training and certification

Thank you!

Jody Kirk

jodykirk@amazon.com

Please complete the session survey in the mobile app.