
AWS
re:Invent

STG211

Extend Amazon S3 to on-premises environments with AWS Outposts

Paul Reed

Principal Product Manager

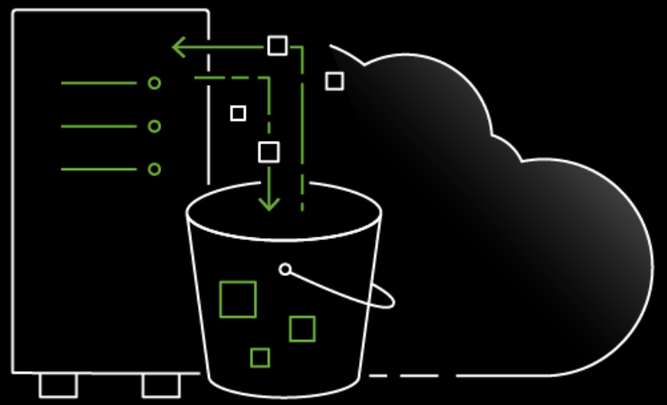
Amazon Web Services

AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Agenda

What it is → What it does → How to use it



Amazon S3 on Outposts

Introducing Amazon S3 on Outposts



What it is

What it does

How to use it

AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

AWS global infrastructure

AWS PROVIDES A MORE EXTENSIVE GLOBAL FOOTPRINT THAN ANY OTHER CLOUD PROVIDER

Region and number of availability zones (AZs)

GovCloud (US)
US-East (3), US-West (3)

US West
Oregon (4)
Northern California (3)

US East
N. Virginia (6), Ohio (3)

Canada
Central (3)

South America
São Paulo (3)

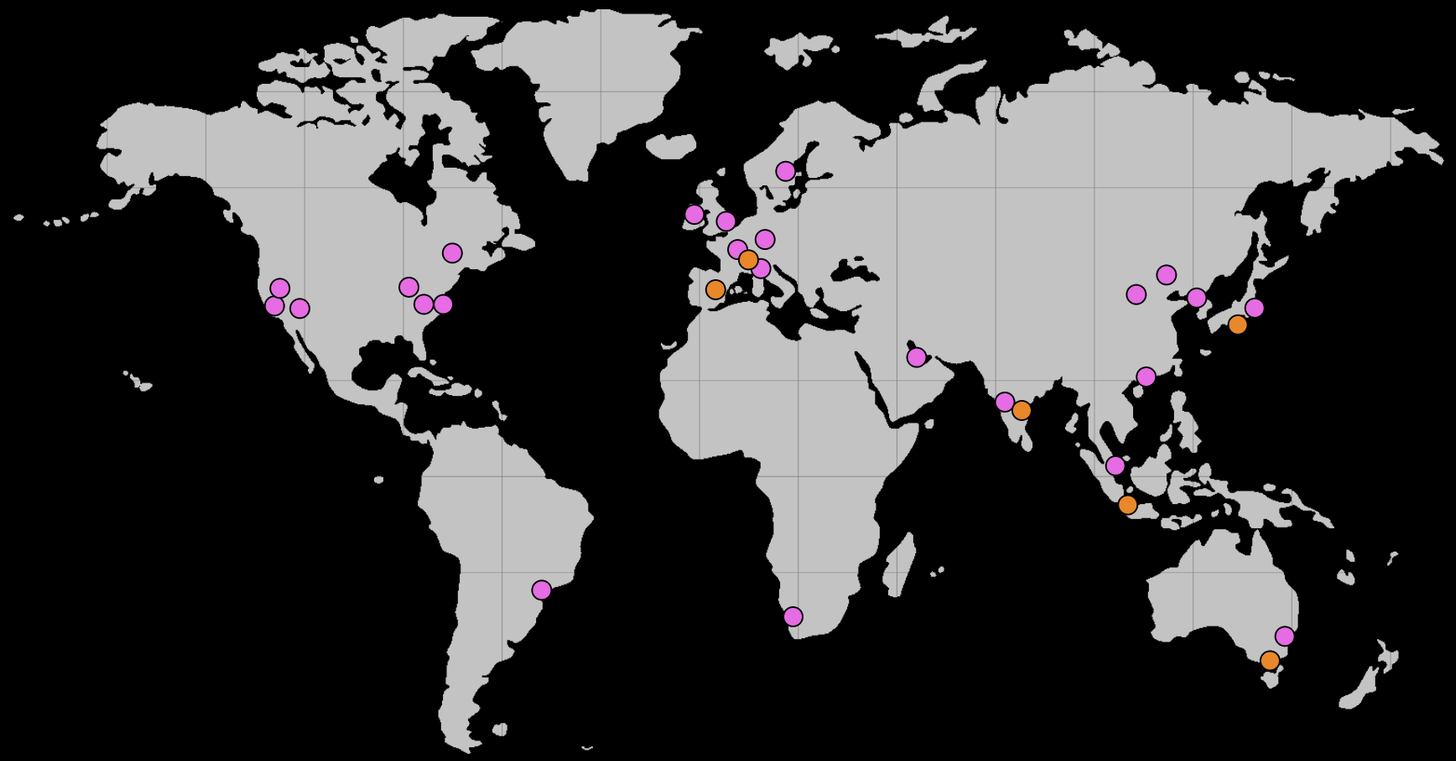
Africa
Cape Town (3)

Europe
Frankfurt (3), Paris (3),
Ireland (3), Stockholm (3),
London (3), Milan (3)

Middle East
Bahrain (3)

Asia Pacific
Singapore (3), Sydney (3),
Tokyo (4), Osaka-Local (1)*,
Seoul (4), Mumbai (3),
Hong Kong (3)

China
Beijing (2), Ningxia (3)



Announced regions

* Available to select AWS customers who request access; customers wishing to use the Asia Pacific (Osaka-Local) Region should speak with their sales representative

Application continuum: Ease of migration to cloud

Easier ← Move to public cloud → **Harder**

Web and enterprise apps



Most web and enterprise applications, such as email, collaboration, and intranet applications, are easily migrated to AWS Regions

Residency



Regulations and contracts dictate that data and infrastructure reside in specific countries

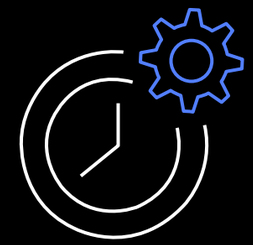
Local data processing



Large datasets that cannot be easily moved

Transcoding, filtering, caching, and alerting applied at the edge

Low latency



Equipment and processes sensitive to compute latency

Complex workloads that span a variety of host and storage systems

Customers want the same experience across their premises, the edge, and the cloud



Same reliable, secure, and high-performance infrastructure



Same operational consistency



Same services and APIs



Same tools for automation, deployment, and security



Same pace of innovation as in the cloud

Introducing AWS Outposts

Run AWS infrastructure and services on premises for a truly consistent hybrid experience

Standard 42U rack – ready to be rolled into final position

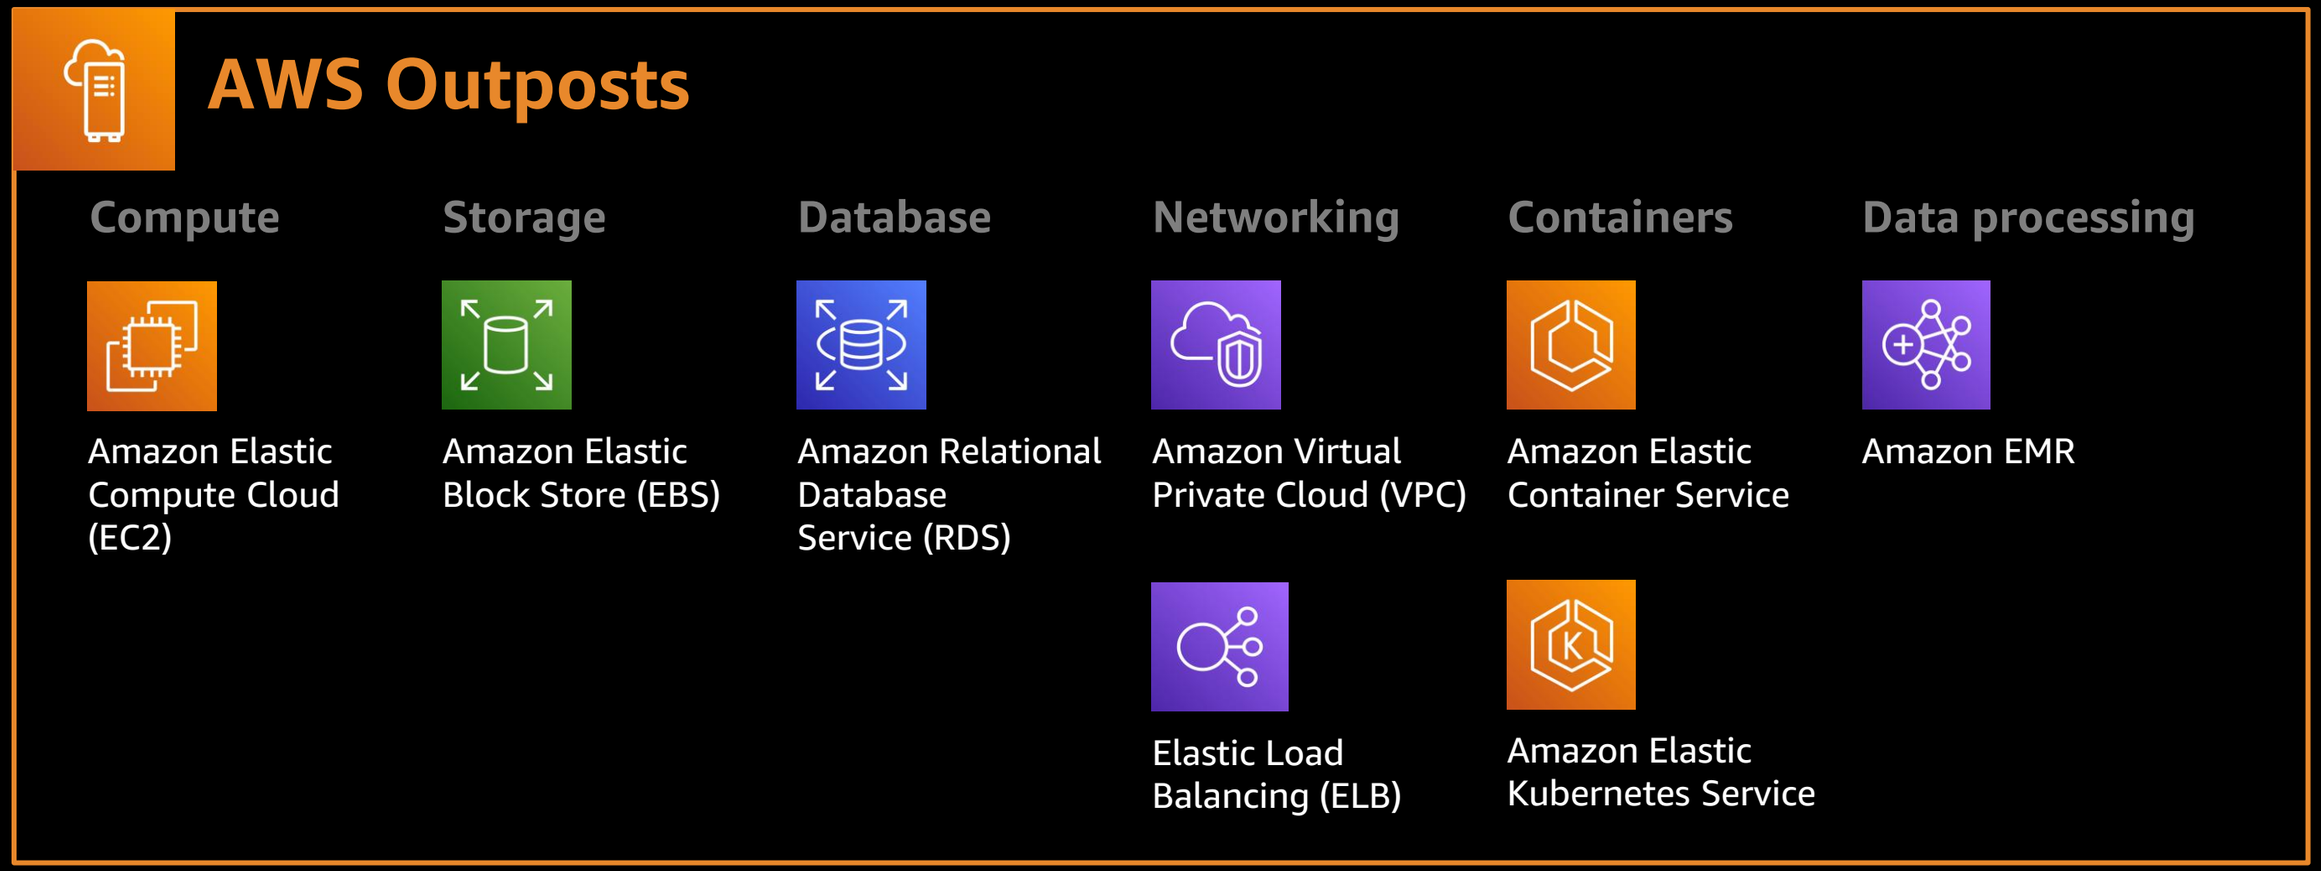
Redundant components including power units, top-of-rack switches, and hot spare hosts

Installed by AWS – simply plugged into power and network



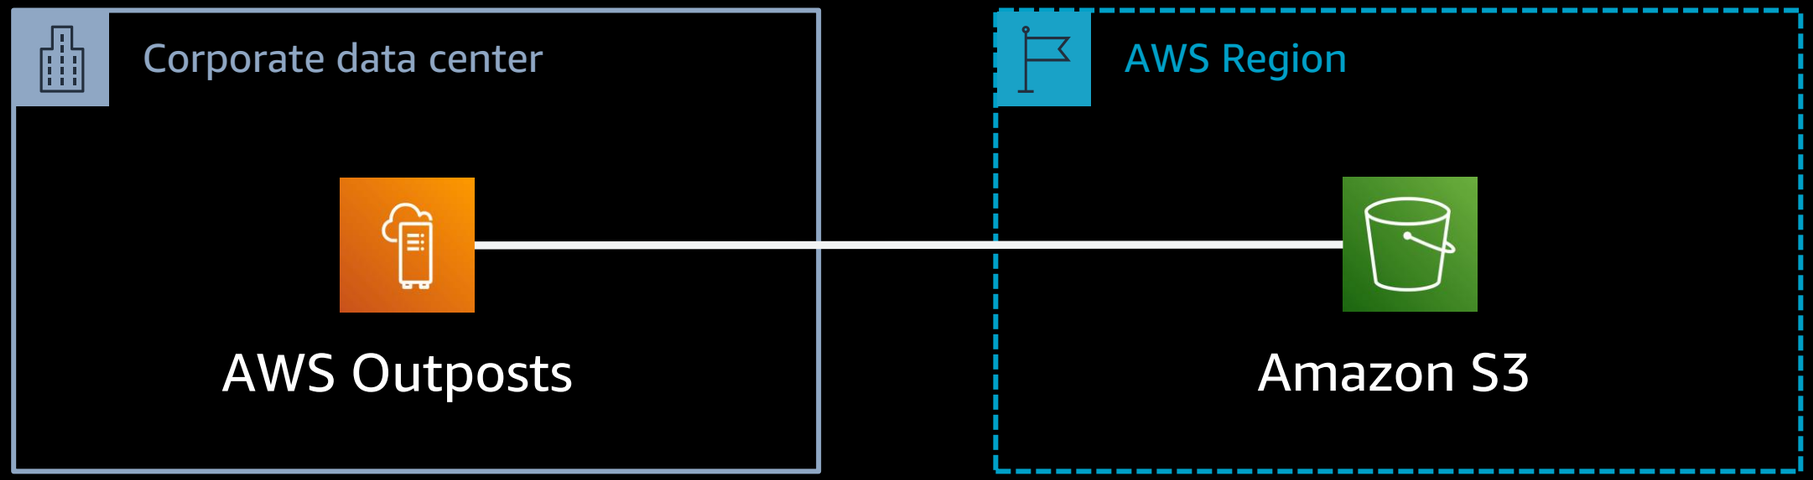
AWS Outposts lets customers run AWS services on their premises

Customers use AWS Outposts for a **broad spectrum of workloads**



AWS Outposts lets customers access AWS services in AWS Regions

Applications on AWS Outposts can **access Amazon S3 in AWS Regions over your network**

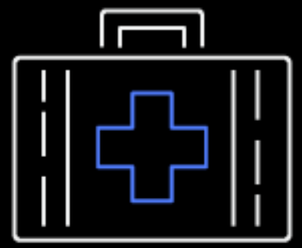


But . . . this does not solve all use cases

Some applications need to **store and access data locally** for processing and data residency purposes



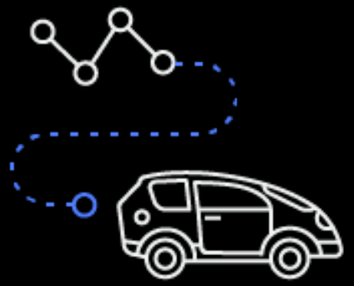
Manufacturing processes



Hospitals and healthcare facilities

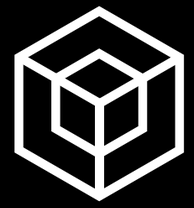


Medical research

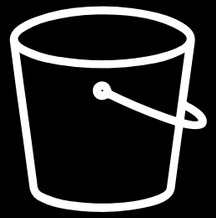


Autonomous vehicle development

Amazon S3 on Outposts brings Amazon S3 to your premises



Same application programming interfaces (APIs) and familiar features of Amazon S3 in the cloud

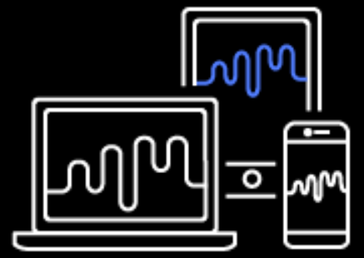


48 TB or 96 TB of storage per AWS Outpost



New Amazon S3 storage class that stores data durably across multiple devices and servers in your AWS Outpost

Amazon S3 on Outposts helps you . . .



Meet data residency requirements



Perform local data processing



Stage and validate applications for cloud migration

Meet data residency requirements with Amazon S3 on Outposts

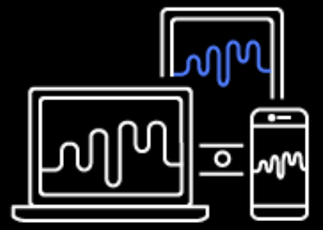


Meet data residency requirements

Store data in a location where there is not an AWS Region today

Build applications on AWS Outposts that use the Amazon S3 API, and meet local data residency requirements

Perform local data processing with Amazon S3 on Outposts



Perform local data processing

Applications such as medical imaging in hospitals, autonomous vehicle data capture, and manufacturing processes require local storage

Local storage minimizes data transfers and buffers applications from the network

Stage and validate applications with Amazon S3 on Outposts



Stage and validate applications for cloud migration

Build and test applications on premises that may eventually move to an AWS Region, and minimize the changes required

An intermediate step in your cloud migration journey is provided

Features of Amazon S3 on Outposts



What it is | What it does | How to use it

AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Same Amazon S3 APIs and familiar features



Amazon S3
on Outposts

- Integrated Amazon S3 API experience
- New storage class with encryption by default
- Object API endpoint on your AWS Outposts
- Amazon S3 security and access control features
- Object and bucket tagging
- Amazon S3 lifecycle expiration actions
- AWS CloudTrail logs and Amazon CloudWatch metrics

Integrated Amazon S3 API experience

Same programming model so applications function without modification

Requires latest AWS Software Development Kits (SDKs) and AWS CLI

New SDK for endpoint management specific to S3 on Outposts



AWS Tools and SDKs



AWS Command Line Interface (CLI)

New storage class for on-premises storage



Amazon S3 Outposts

Single storage class, Amazon S3 Outposts

Durably and redundantly store data across multiple devices and servers on your AWS Outposts

Always encrypts objects using server-side encryption with S3 managed keys (SSE-S3)

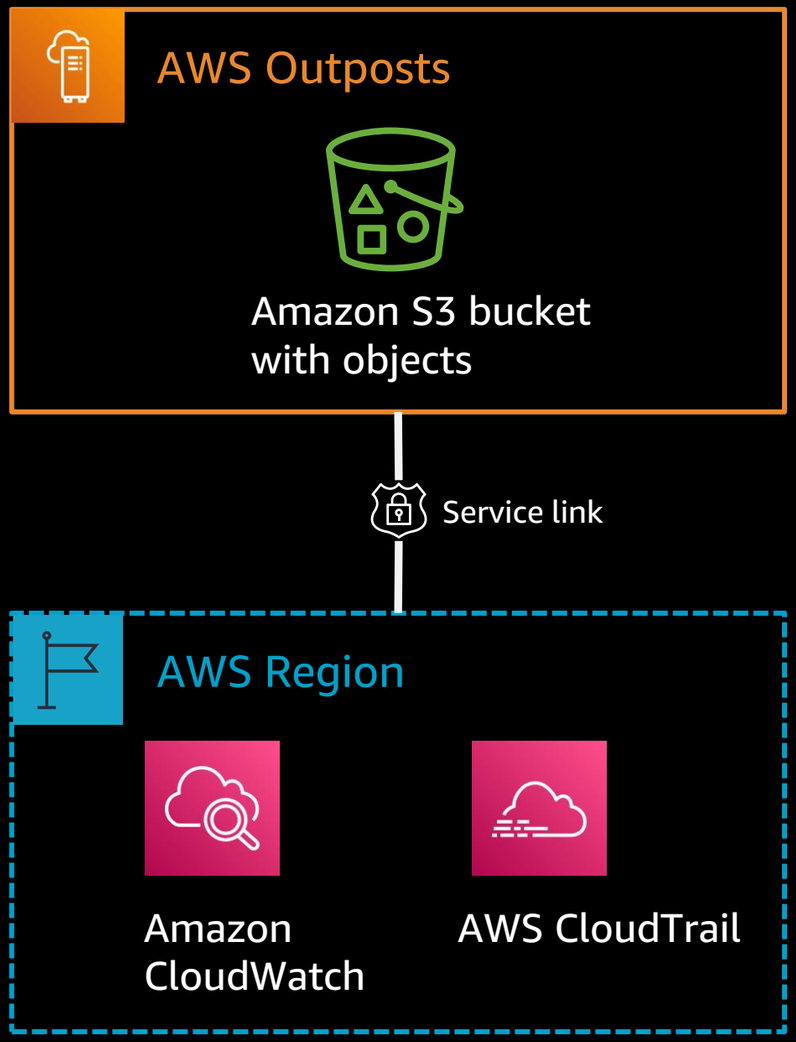
Data storage on-premises integrated with the AWS cloud

Data is stored on your AWS Outpost

- Object data
- Object user metadata
- Object tags

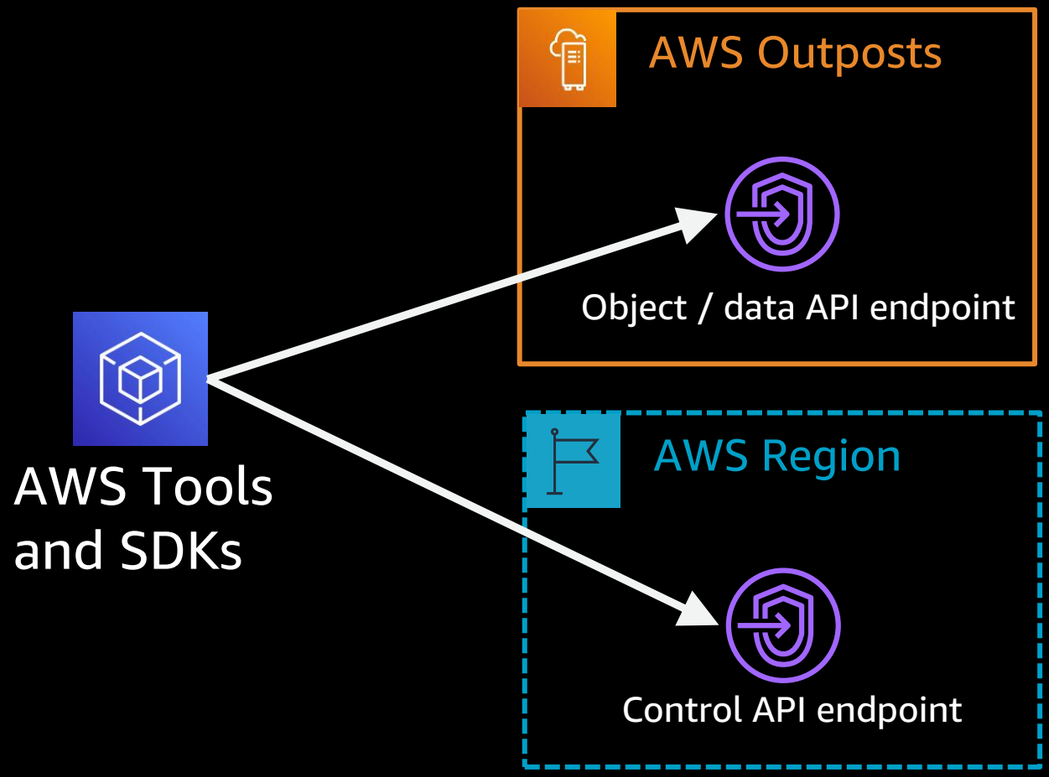
Buckets are created and managed from your AWS Outposts home region

Telemetry is published to the home region



Object API endpoint on your AWS Outposts

Separate endpoints for control and data API requests



AWS Outposts endpoint in your VPC for **access to objects** in buckets

Regional endpoint for **management and control** of buckets

SDK/CLI determines correct endpoint based on request and attributes

Familiar security and access control features



AWS Identity & Access Management (IAM) permissions and policies



Block Public Access (always enabled)



Amazon S3 Access Points (VPC restricted)



S3 Object Ownership (always enabled)

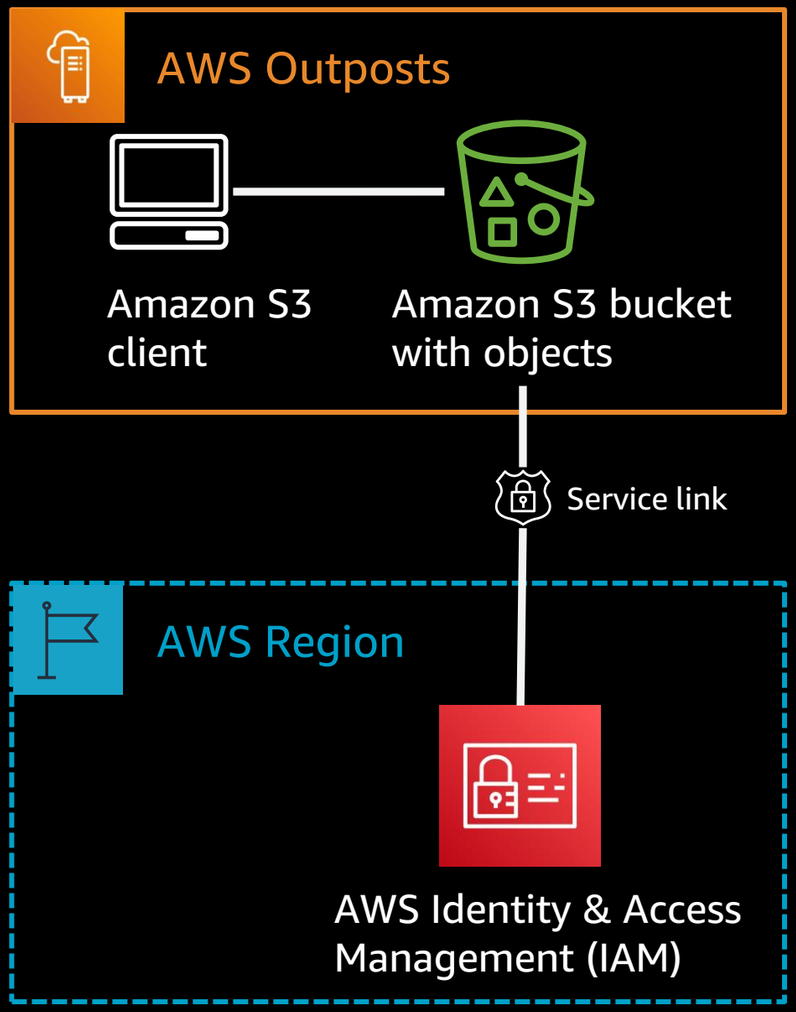
AWS Identity & Access Management: Controlling access to your data

API authentication uses AWS IAM in the AWS Region via the service link

- Both bucket and access point policies are enforced
- Policies use the s3-outposts:* namespace

Access denied if no connection to AWS IAM

- Data remains safely stored on the AWS Outposts



Using Amazon S3 on Outposts



What it is

What it does

How to use it

AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Storing data as objects on your AWS Outposts







Order and provision Amazon S3

Create and configure an Amazon S3 bucket

Store your data

Create and configure your Amazon S3 resources



Create a bucket on your AWS Outposts



Create an Amazon S3 Access Point restricted to a VPC on your AWS Outposts

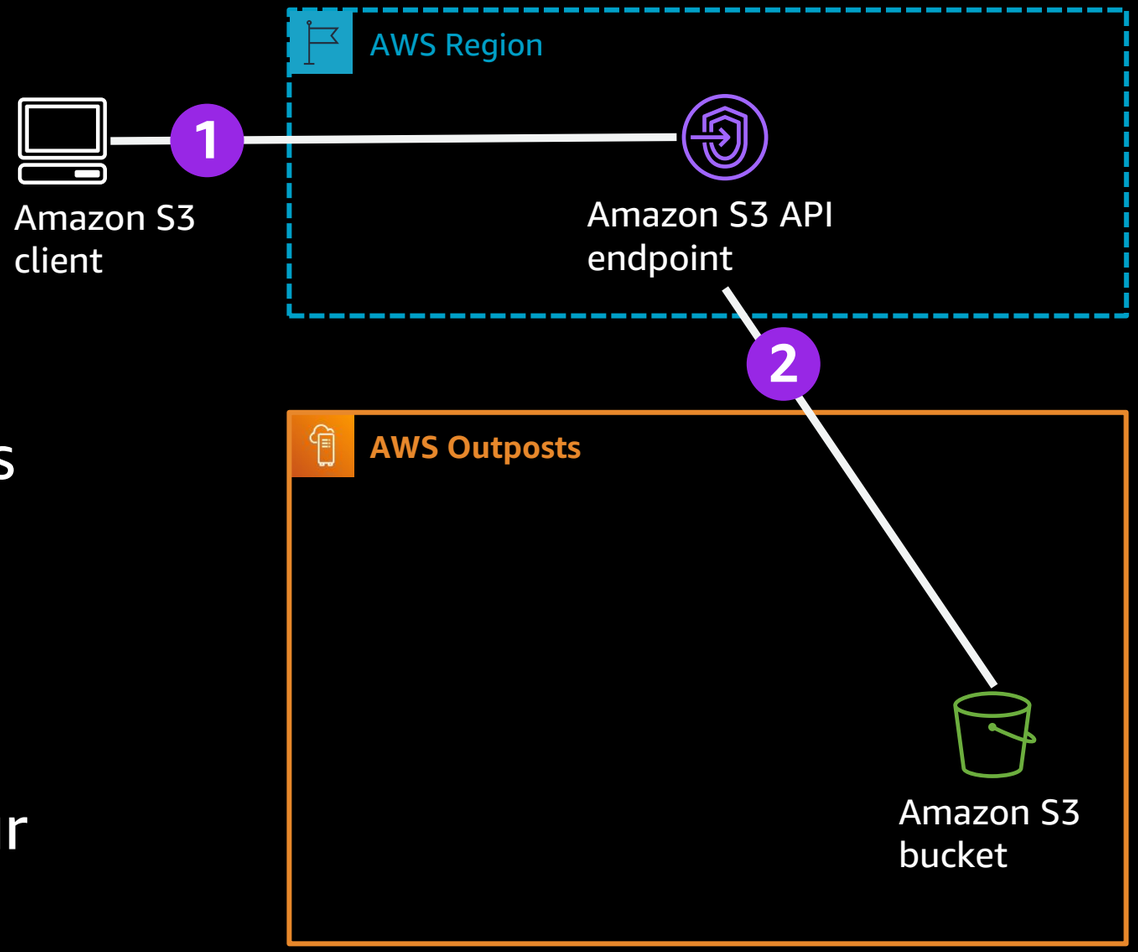


Create an endpoint in a subnet on your AWS Outposts

Create a bucket on your AWS Outposts

```
$ aws s3control create-bucket … --outposts-id {outpostsId}
```

Request includes ID of the AWS Outposts on which to store the bucket



1. CLI sees the AWS Outposts ID and routes request to regional endpoint

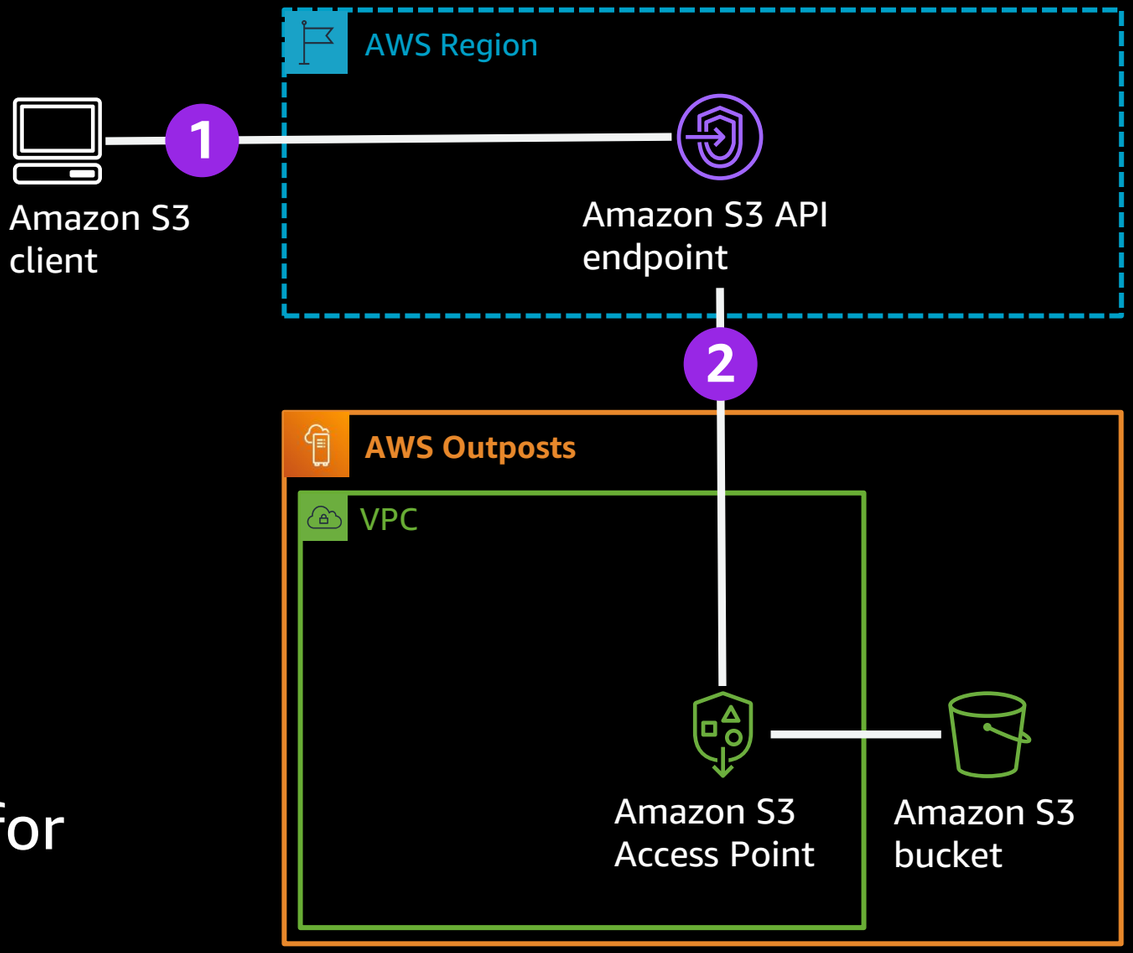
   Bucket ARN is returned

2. Bucket is asynchronously created on your AWS Outposts

Create an Amazon S3 Access Point

```
$ aws s3control create-access-point … --bucket {bucketArn} \
   --vpc-configuration …
```

Request includes a VPC on your AWS Outposts to restrict access



1. CLI routes request to the regional endpoint for Amazon S3 on Outposts

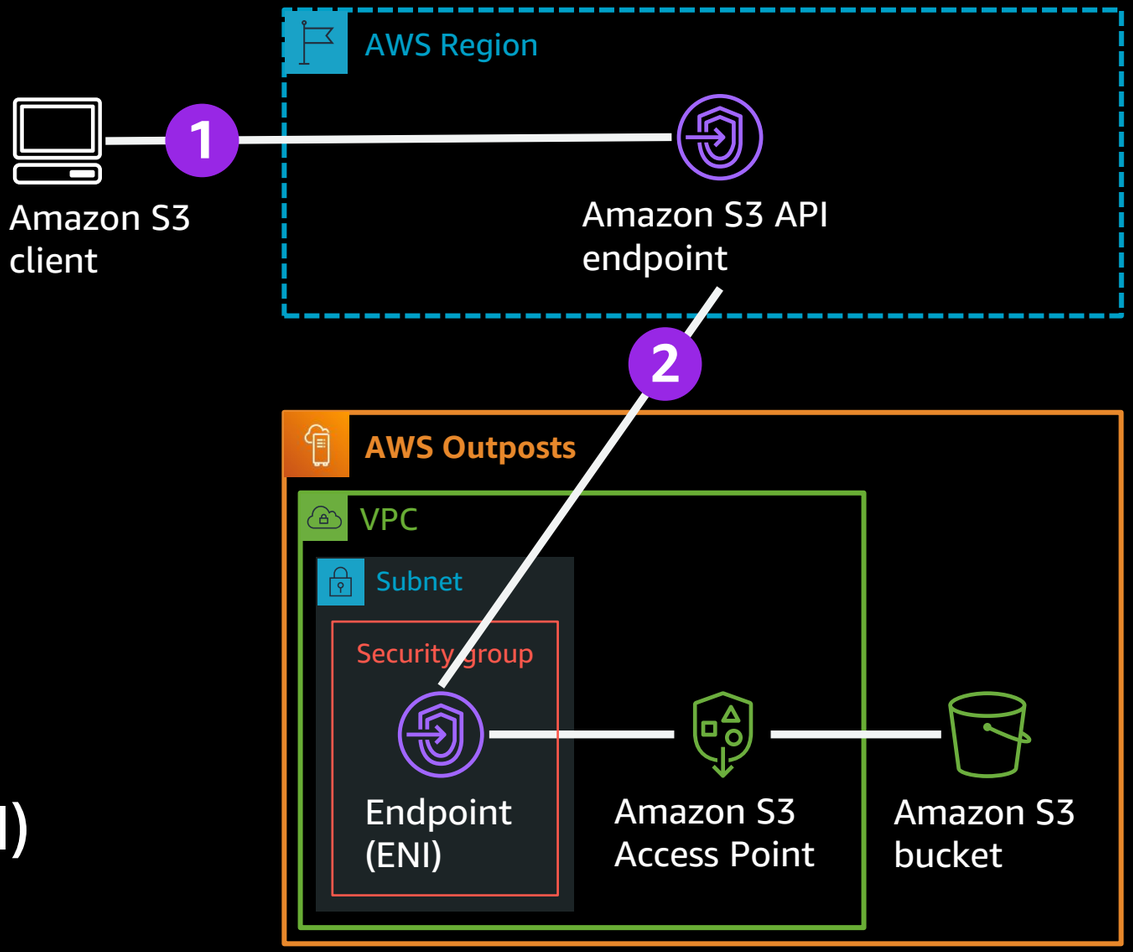
   Access point ARN is returned

2. Asynchronously creates an access point for the bucket

Create an endpoint on your AWS Outposts

```
$ aws s3outposts create-endpoint … --subnet-id … \
   --security-group-id …
```

Request includes subnet and security group for the endpoint on your AWS Outposts



1. CLI routes **s3outposts** request to the regional endpoint

   Endpoint ARN is returned

2. Asynchronously creates an endpoint (ENI) in your subnet

Store an object

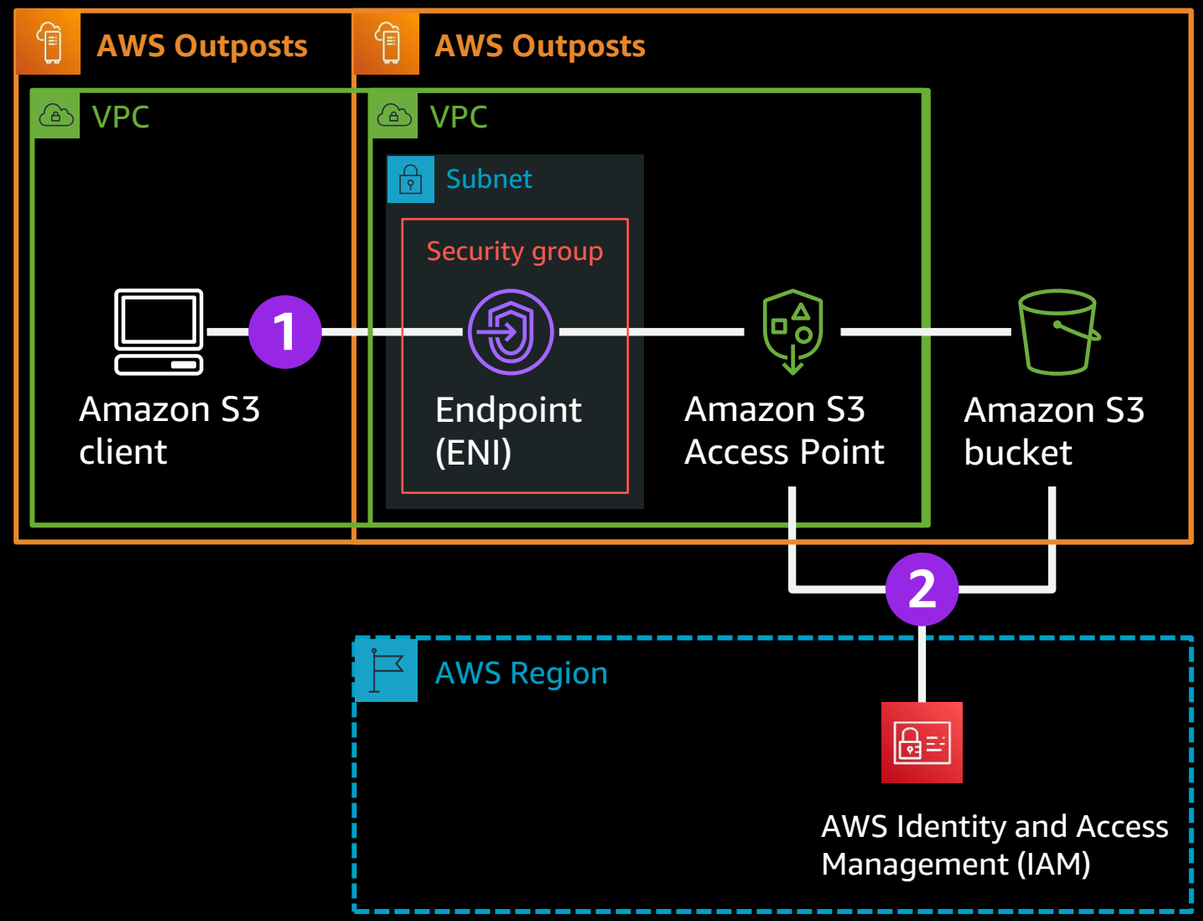
```
$ aws s3api put-object … --key … --bucket {accessPointArn}
```

Request uses the ARN of the access point in place of the bucket name

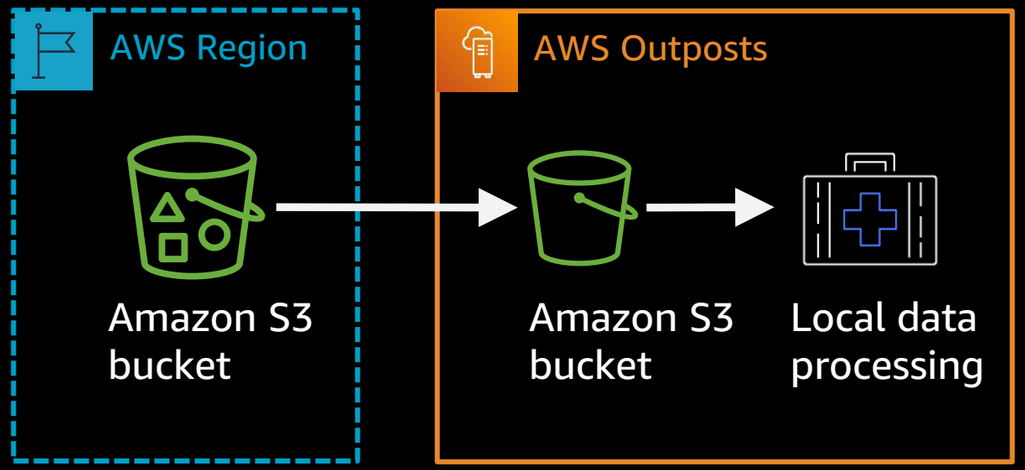
1. CLI uses access point ARN to resolve the endpoint on your AWS Outposts

2. Request is authenticated and authorized with VPC-restricted access point and bucket policies

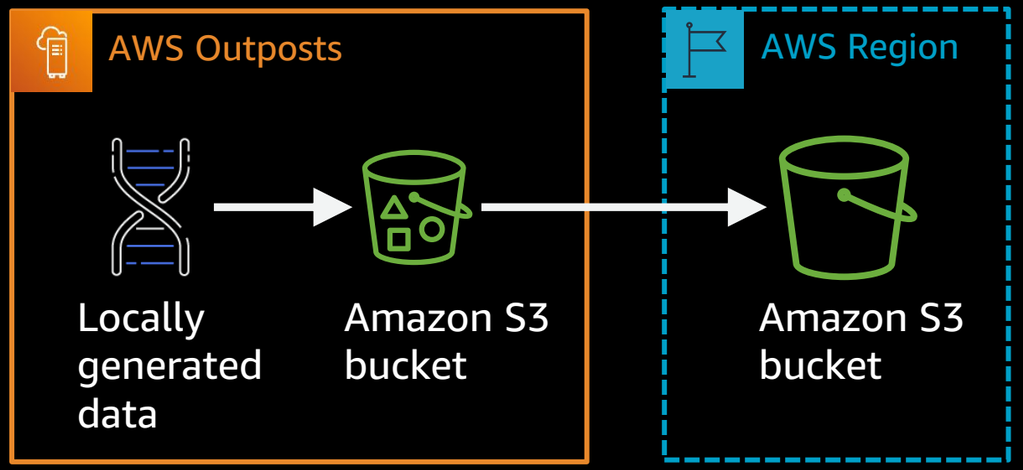
   Object is stored



Managing data movement between AWS Outposts and AWS Regions



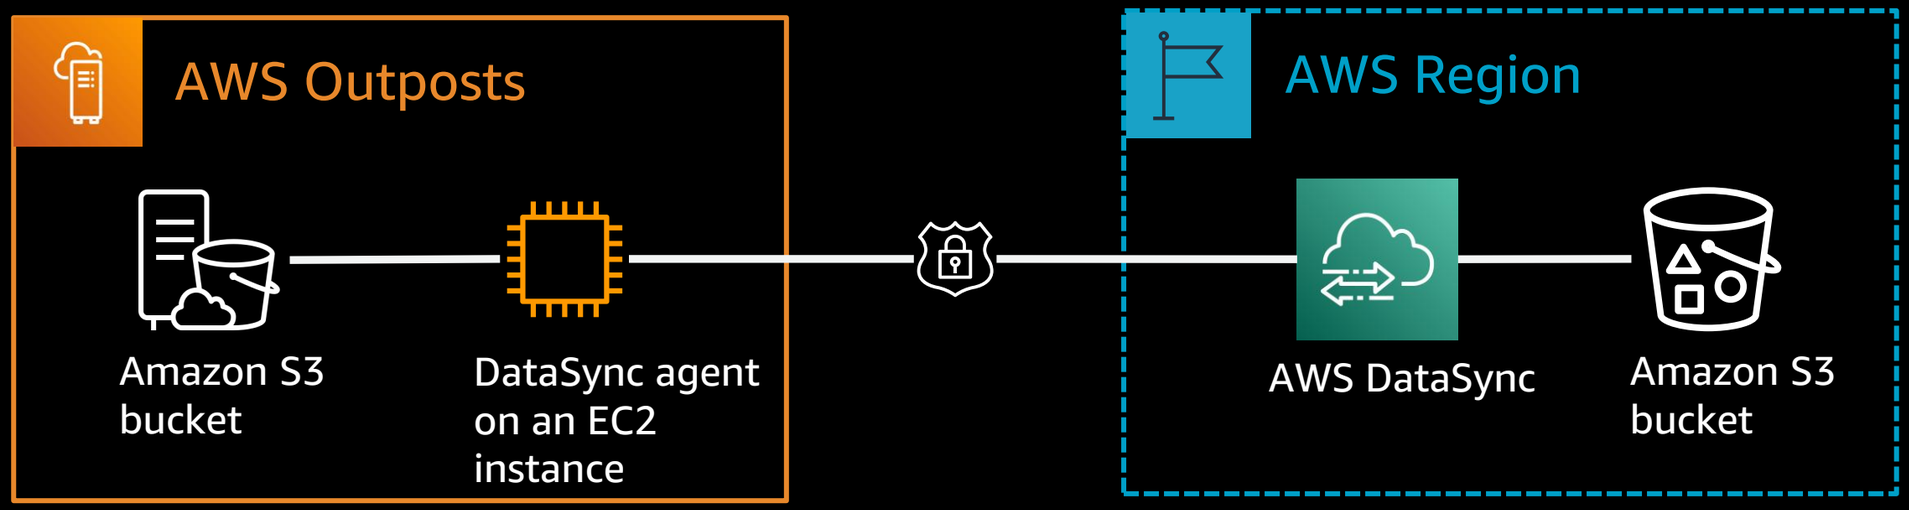
Bring data to your AWS Outposts for **local data processing**



Transfer data for further **analysis or archival in the cloud**

AWS DataSync automates and accelerates data movement

Securely and efficiently **copy objects between Amazon S3 buckets** on AWS Outposts and in AWS Regions



Summary



© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Amazon S3 on Outposts: Bringing Amazon S3 on premises



Emerging need for high-bandwidth data processing and data residency requirements



Customers want the same Amazon S3 experience on premises and the cloud to speed cloud migration



Amazon S3 on AWS Outposts delivers the Amazon S3 APIs and familiar features of the cloud



Simplifies IT

Grows IT efficiency

Amplifies productivity

Accelerates pace of innovation

Thank you!



AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.



Please complete the session survey

AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.