Architecting for Data Residency on AWS Outposts

John Mallory, AWS Outposts BD

David Filiatrault, AWS ProServe Outposts Lead

March 25, 2021

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

aws

What are Data Residency and Sovereignty?

Data Residency

The requirement that all customer content be processed and stored in an IT system that remains within a specific locality's borders

Data Sovereignty

The control and governance over who can and cannot have legal access to data, its location and usage

AWS Outposts may help you meet Data Residency & Sovereignty needs

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

aws

The most sensitive workloads run on AWS

"We determined that security in AWS is superior to our on-premises data center across several dimensions, including patching, encryption, auditing and logging, entitlements, and compliance."

—John Brady, CISO, FINRA (Financial Industry Regulatory Authority)

"AWS allowed us to scale our business to handle 6 million patients a month and elevate our security—all while maintaining HIPAA compliance-–as we migrated 100% to cloud in less than 12 months"

—Brian Lozada, CISO, Zocdoc.

"Amazon Web Services was the clear choice in terms of security and PCI DSS Level 1 compliance compared to an on-premises or co-location data center solution."

—Stefano Harak, online senior product manager for Vodafone Italy

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

aws

Infrastructure and services to elevate your security in the cloud

Inherit global security & compliance controls

Scale with superior visibility & control

Highest standards for privacy & data security

Automate & reduce risk with deeply integrated services

Largest ecosystem of security partners & solutions

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

aws

Highest standards for privacy and data security

Meet data residency requirements
Choose an AWS Region and AWS will not replicate it elsewhere unless you choose to do so

Encryption at scale
with keys managed by our AWS Key Management Service (KMS) or managing your own encryption keys with AWS CloudHSM using FIPS 140-2 Level 3 validated HSMs

Comply with local data privacy laws
by controlling who can access content, its lifecycle, and disposal

Access services and tools that enable you to **build compliant infrastructure** on top of AWS

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

But what if you can’t move your data to an AWS region?

Local Processing

No Local AWS Region

Low Latency

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

aws

Customers want the same experience across on-premises and the cloud

Same reliable, secure, and high-performance infrastructure

Same operational consistency

Same services and APIs

Same tools for automation, deployments, and security controls

Same pace of innovation as in the cloud

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

aws

AWS Outposts: Bringing AWS on-premises

Same AWS-designed infrastructure as in AWS data centers (built on AWS Nitro System)

Fully managed, monitored, and operated by AWS as if in AWS Regions

Single pane of management in the cloud providing the same APIs and tools as in AWS Regions

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

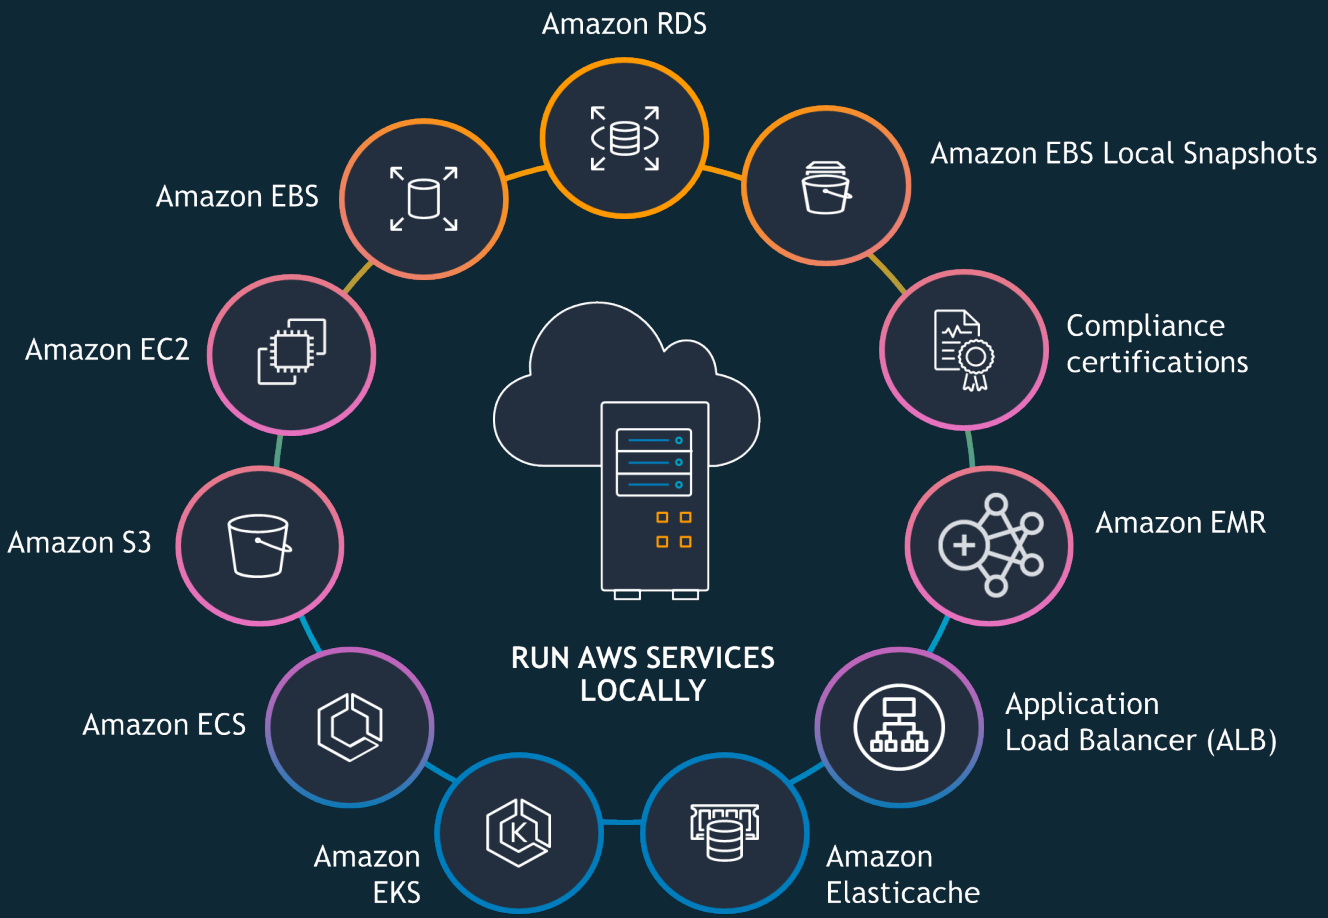

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

AWS Outposts rack

Industry standard 42U rack

Fully assembled, ready to be rolled into final position

Installed by AWS, simply plugged into power and network

Centralized redundant power conversion unit and DC distribution system for higher reliability, energy efficiency, easier serviceability

Redundant active components including top of rack switches

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

Transforming the user experience in banking

The challenge

- Transform the user experience for its banking customers while complying with data residency requirements

The solution

- New user experiences for mobile payments and banking
- Using Outposts to build and deploy modern containerised workloads while being able to store and process data on-premises
- Multiple AWS Outposts located in Abu Dhabi and Dubai to meet business continuity requirements
- Leveraging Control Tower to implement governance and guardrails from the AWS Region to AWS Outposts

Business outcome

- Accelerated digital transformation in the cloud and on-premises
- Increased developer and infrastructure team productivity by leveraging the same infrastructure, services, APIs, and tools across the cloud and on-premises

"First Abu Dhabi Bank (FAB) is the UAE's largest bank and one of the world's largest and safest financial institutions, offering corporate, investment, and personal banking services. It was important to us to deliver digital banking services, including e-wallet and mobile payments, to our customers, whilst keeping customer data secure and resident in the UAE. Multiple AWS Outposts in the UAE allow us to provide business continuity whilst leveraging the same API's and services running locally as in the AWS Region, to accelerate the pace of our innovation."

Yuri Misnik, Group Chief Technology Officer, FAB

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

Tipico uses AWS Outposts to scale iGaming expansion plan

Challenge

To expand its bookmaking business internationally, Tipico had to deploy infrastructure in colocation facilities and use different tools and services to run its application than it used in the cloud.

Solution

The company used AWS Outposts to rapidly and compliantly enter and expand in the US market while optimizing costs, simplifying its architecture, and gaining the agility for future rapid expansion.

Benefits

- Entered US market 5–10x faster
- Reduced bet slip process from 400–500 ms to 150 ms
- Onboards engineers in 24 hours instead of 1–2 weeks

> "AWS Outposts ideally fit our use case. We can use the AWS APIs, and we can use all we have in code; we just need to make some small adjustments."
>
> **– Thorsten Hanf**, head of enterprise operations, Tipico

Company: Tipico

Industry: Gaming

Country: Global

Website: www.tipico.com/us

About Tipico

Founded in 2004, Tipico is one of the leading gaming operators in Germany and a top gaming operator worldwide. Serving seven million customers, it offers safe, secure digital and mobile betting entertainment across 30 sports.

In Region shared responsibility model

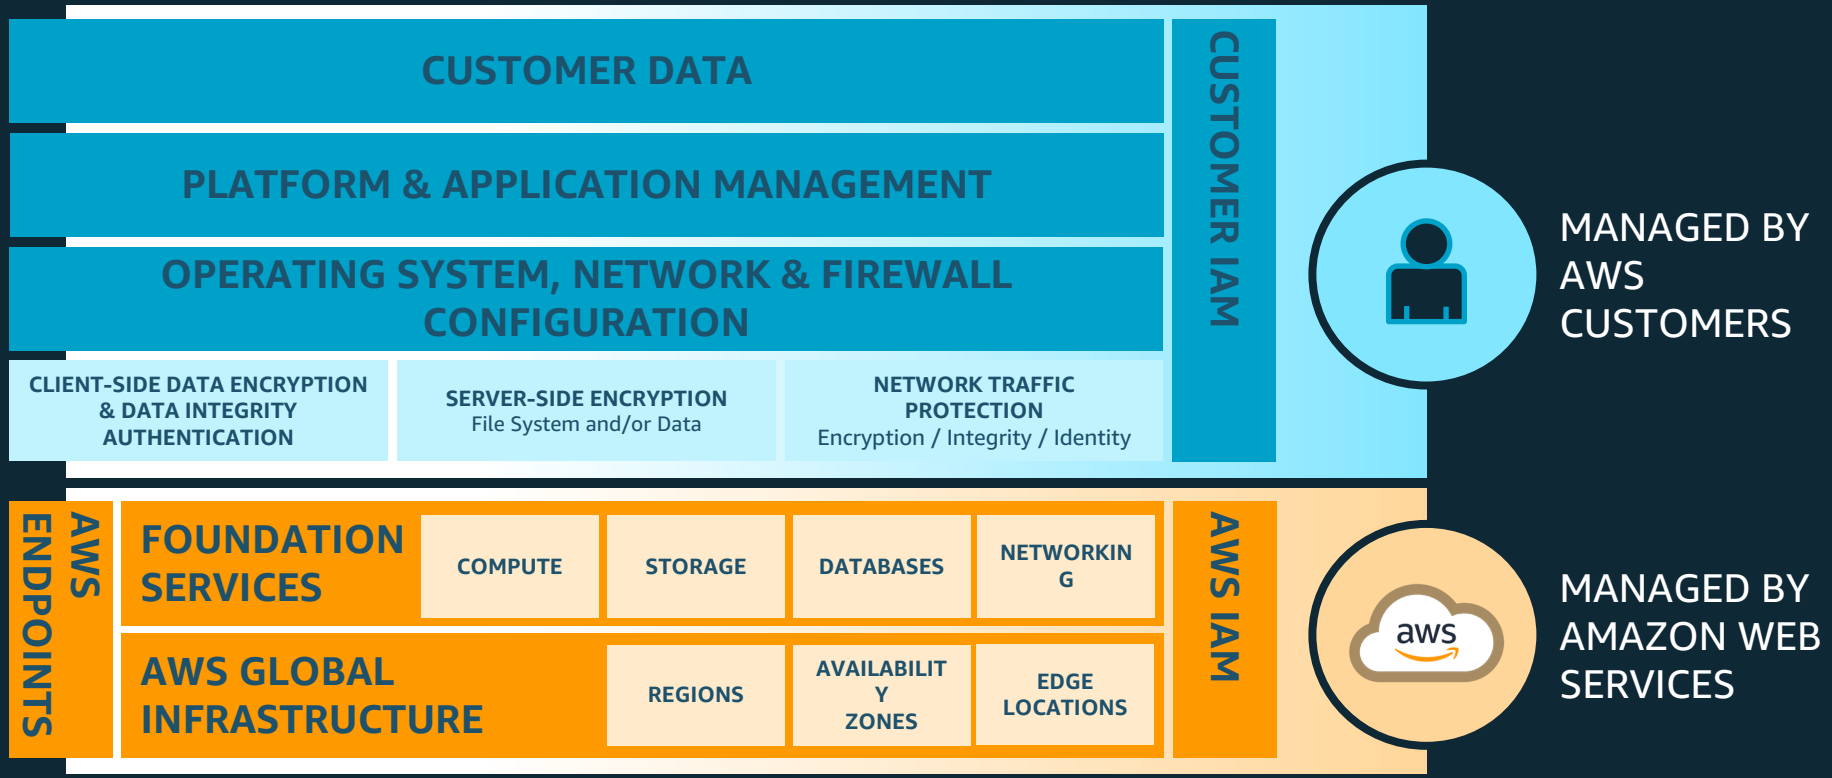

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

AWS Outposts shared responsibility model

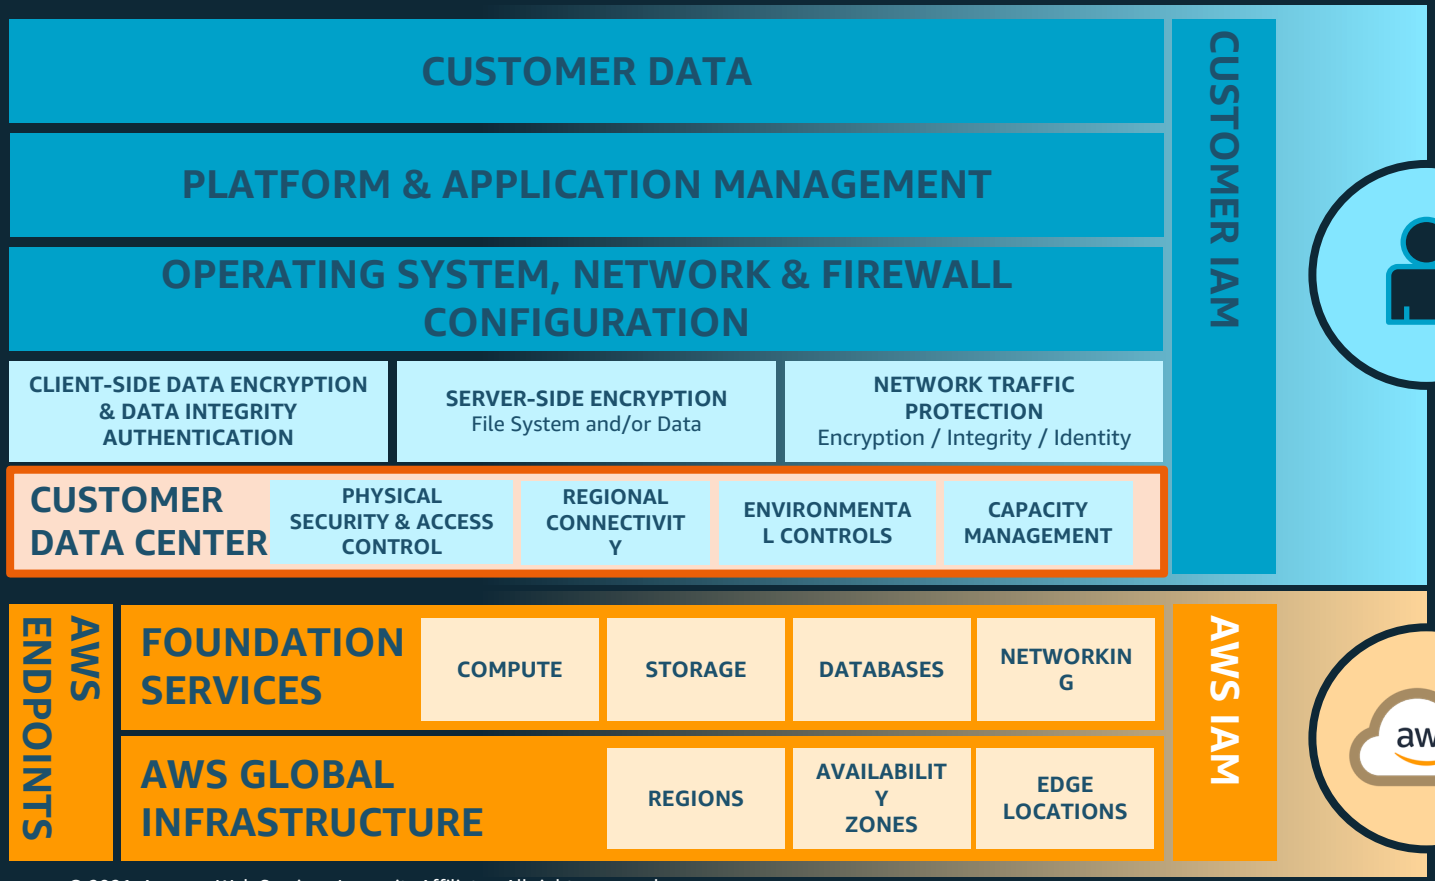

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

AWS Outposts logical security - Service Link specifics

Control plane stays within the region

VPN tunnels to anchor points within a single Availability Zone

TCP/UDP port 443 required

Each Outposts server makes a service link connection to the region

Service link is established outbound from the Outpost

Consists of Data plane and Management Plane

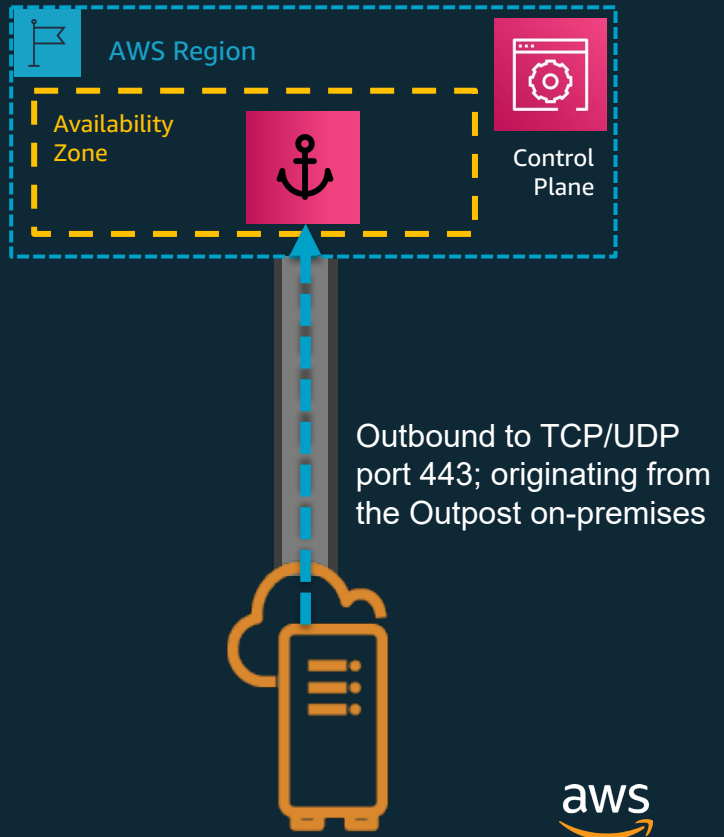

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

Management and control plane in the AWS region

Example data flows

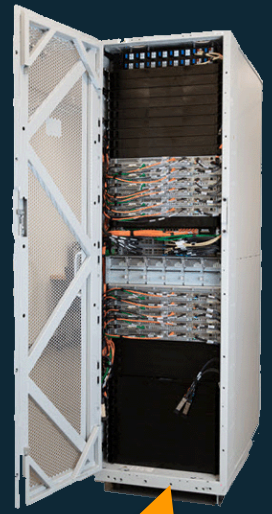

Telemetry including health, state, hardware metrics

RDS Snapshots and Backups

Cloudwatch

EBS snapshots
(when the OutPostArn is not specified)

APIs calls

Systems Management

DNS

Amazon Machine Images

EKS, ECS, EMR Control Plane

Region

Identity and Access Management

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

Seamlessly extend your regional VPC

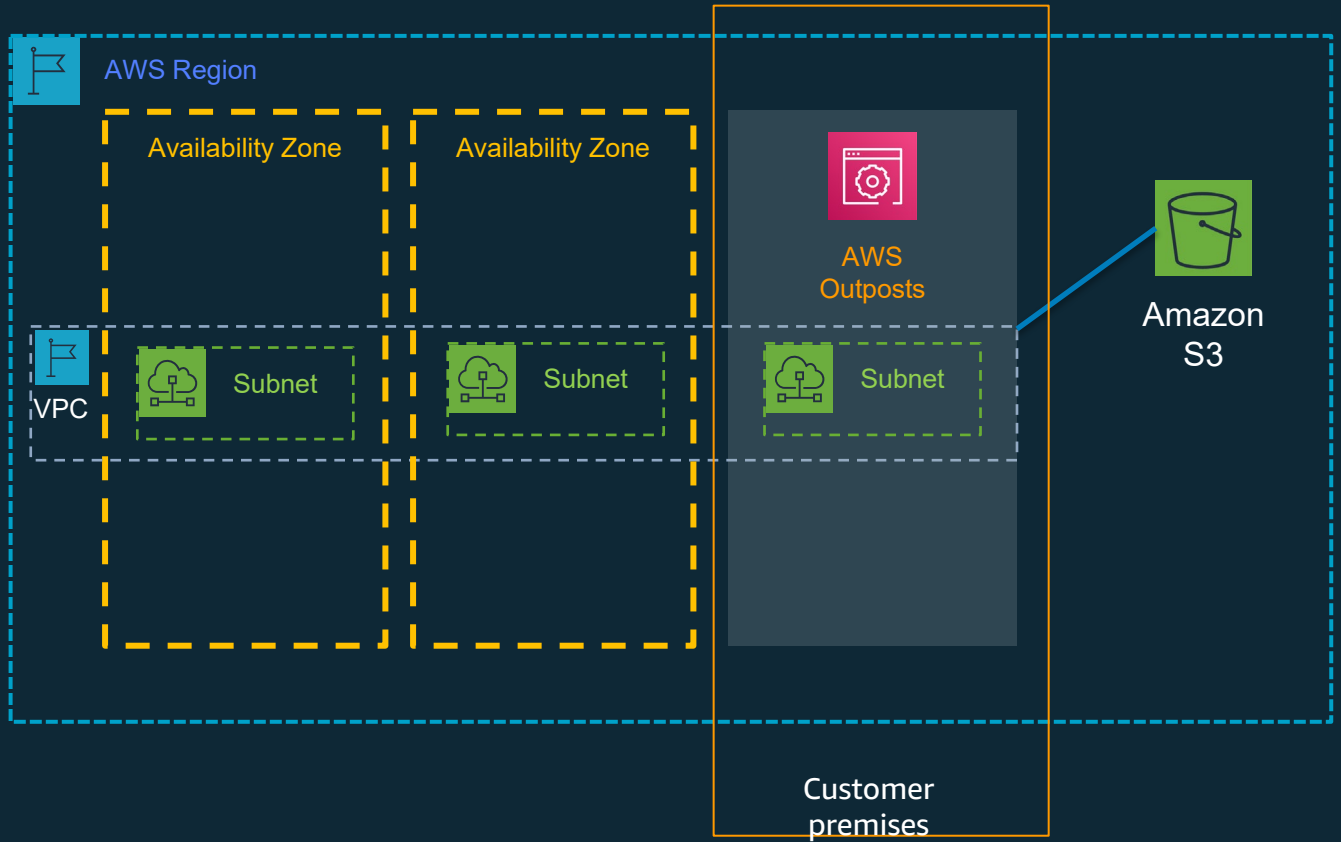

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

You can keep your VPCs contained within Outposts

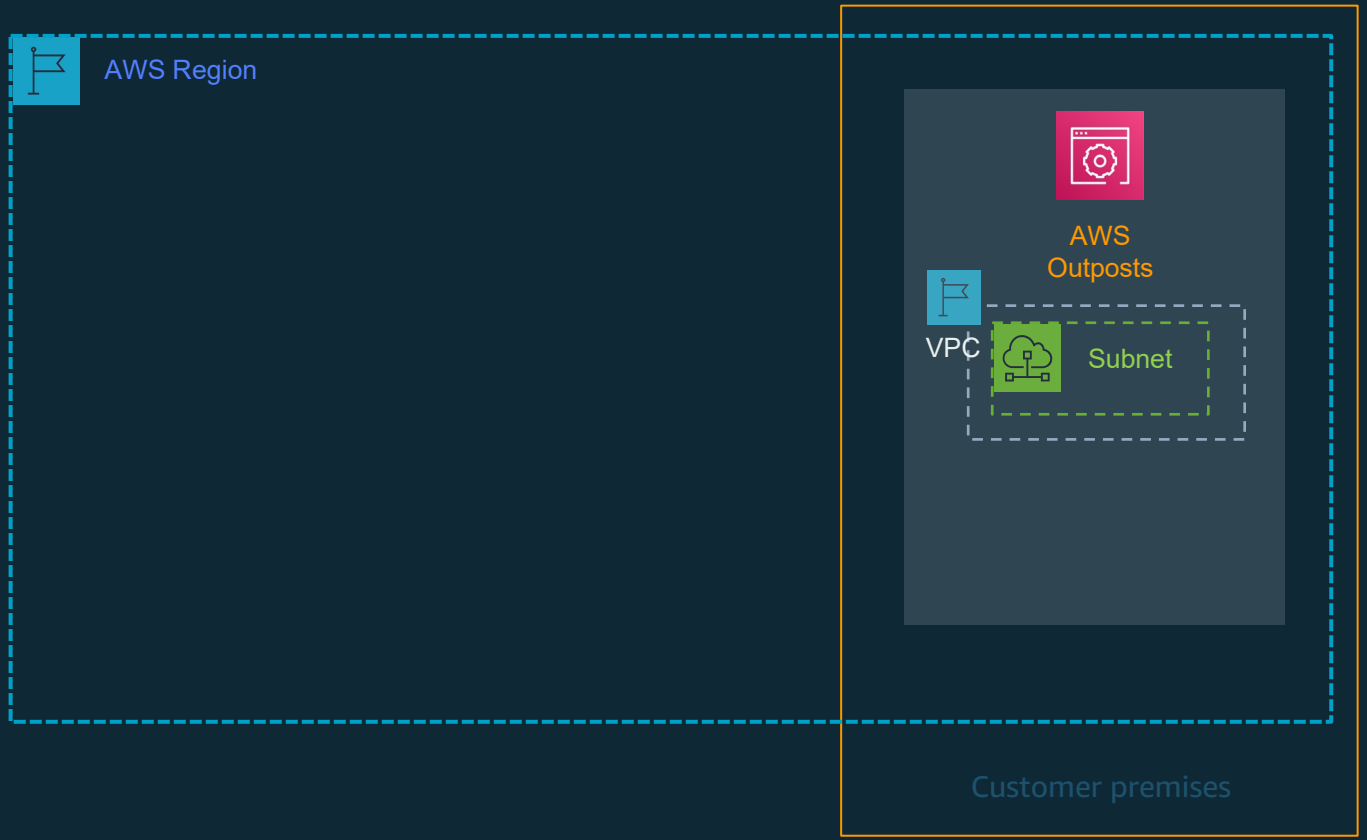

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

Customer data & processing can stay resident on Outposts

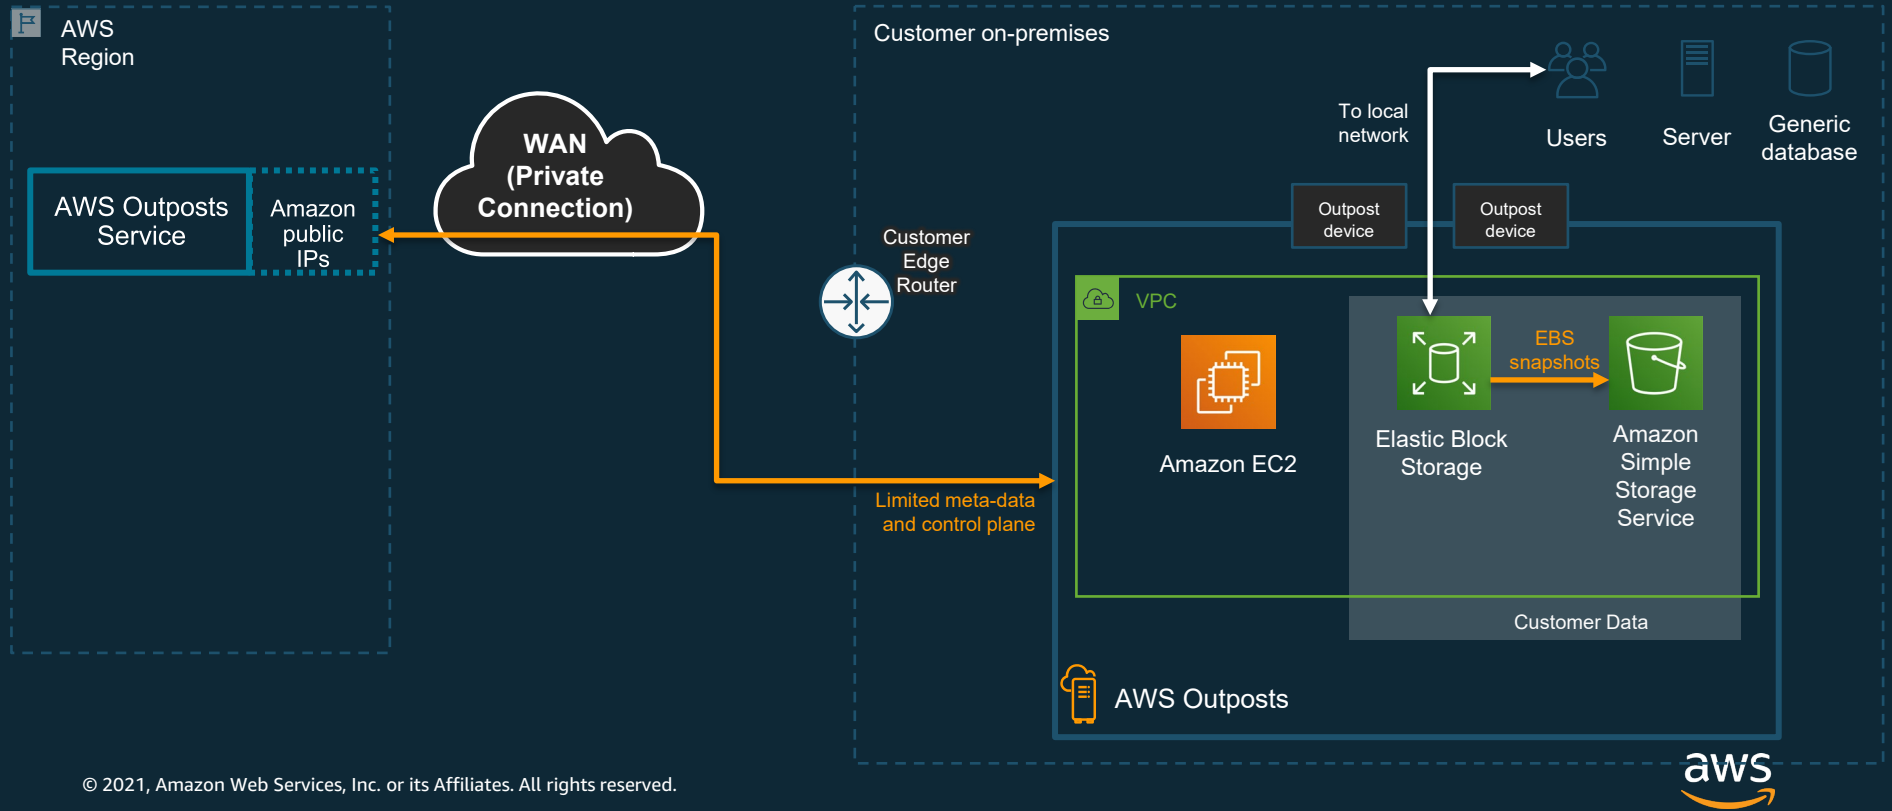

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

S3 on Outposts Architecture

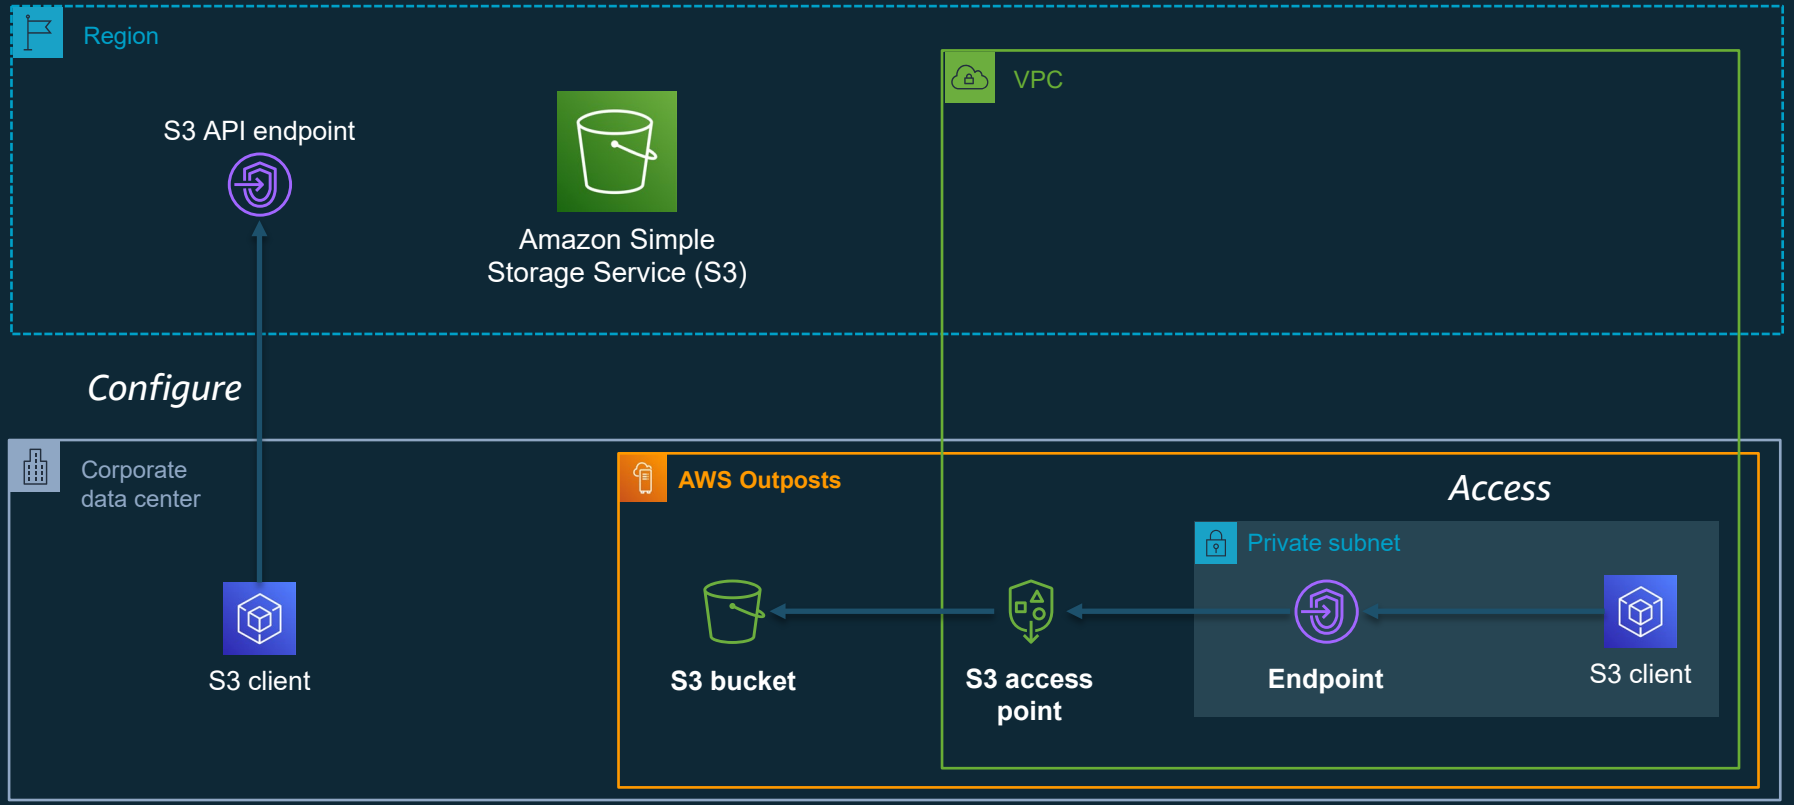

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

Where is my object data stored?

Data is always stored on the Outposts

- Object data
- Object system and user metadata
- Object tags

Buckets are created and managed in the Outposts home region

Telemetry is available in the home region

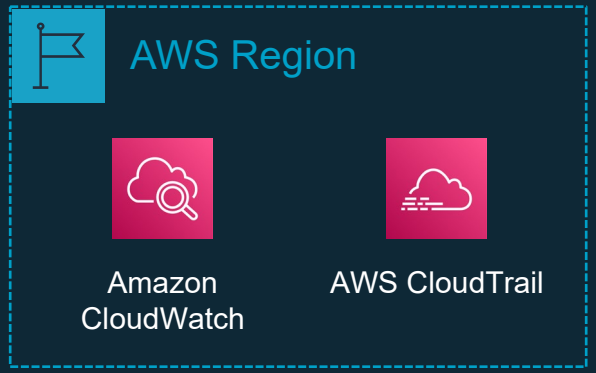

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

S3 security and access control features

IAM policies

Block Public Access
(always enabled)

S3 Access Points
(VPC restricted)

S3 Object Ownership
(always enabled)

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

aws

Create access point on Outpost

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

Request must include VPC configuration to restrict access

1. SDK uses routes request to the regional endpoint for S3 on Outposts
2. Service asynchronously creates access point for bucket

Response contains ARN of the access point

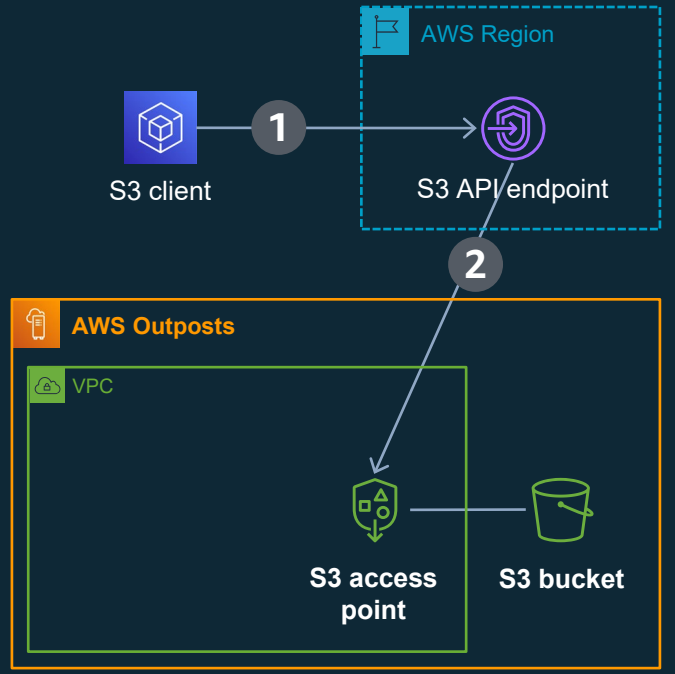

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

Create endpoint on Outpost

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

Request includes subnet and security group for endpoint on the Outpost

1. SDK routes *s3outposts* request to the regional endpoint
2. Service asynchronously creates endpoint (X-ENIs) in the specified subnet

Response contains ARN of the endpoint

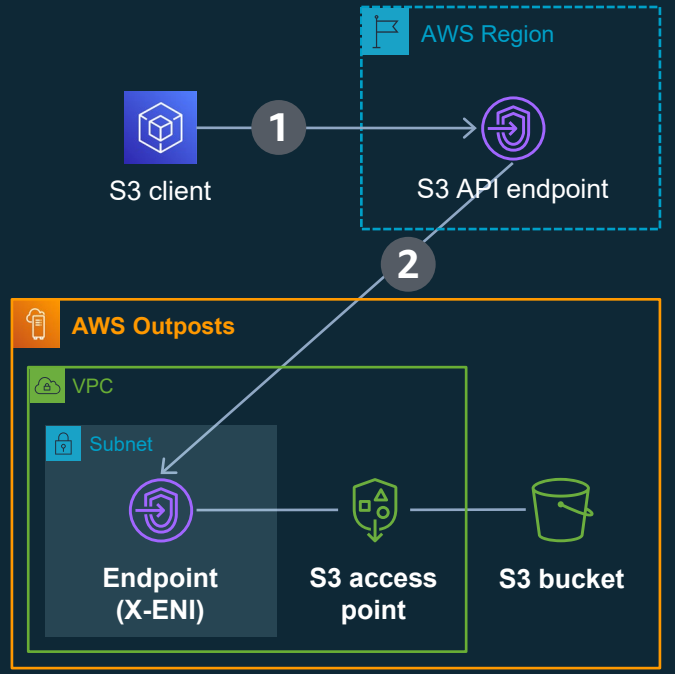

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

Store an object

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

Request uses the ARN of the access point as the bucket name

1. SDK uses access point ARN to route request the endpoint on the Outpost
2. Request is authenticated and authorized against access point and bucket policies

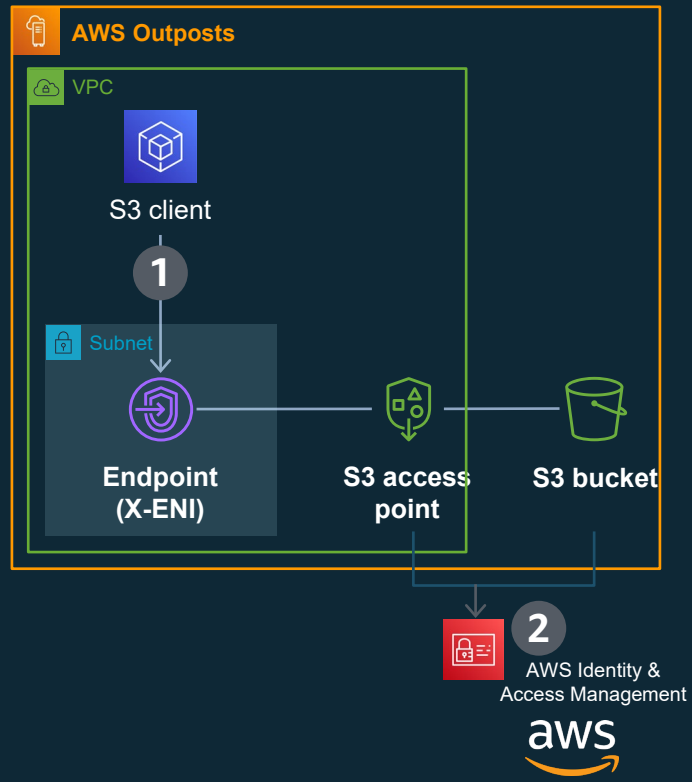

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

Amazon EBS local snapshots on AWS Outposts

Store snapshots of data, boot volume and EBS backed AMI's on- premises on Amazon S3 on Outposts

Requires Outposts to be provisioned with S3 on Outposts

Meet real-time local data processing and data residency needs with local backups

Create point in time incremental snapshots of block storage on Outposts using the EBS Snapshots API/CLIs

Copy regional snapshots/AMI to Outposts and hydrate EBS Volume or launch EC2 instances using local AMI on Outposts

All local snapshots on Outposts are encrypted by default

Control access to data and movement using AWS Identity and Access Management

Automation and Life cycle Policies through Amazon Data Life cycle Manager

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

aws

Encryption on AWS Outposts

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

Physical security

Nitro Security Keys

- Key material is wrapped to the external key stored in this removable device and is leveraged as part of a Trusted Platform Module (TPM)
- If the key is removed, any data on the server is cryptographically shredded per the NIST 800-88 standard
- Physical destruction enabled by turning a screw on the key
- Key has 1:1 relationship with server and can’t be replaced.
- Can’t remove server without removing the key

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

aws

Encryption at Rest – Envelope encryption primer

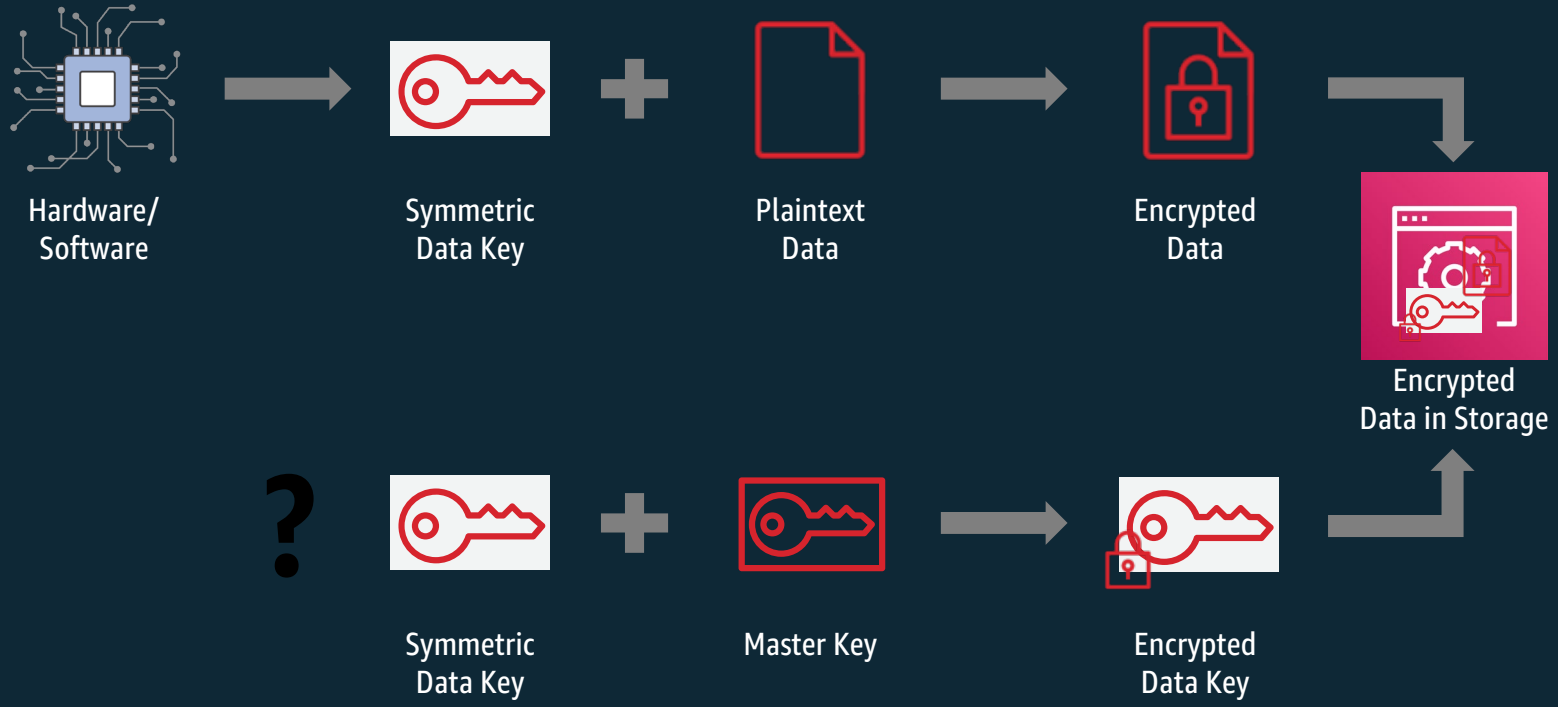

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

Amazon EBS Encryption

What is encrypted:

- EBS Encryption is enabled by default and can’t be disabled
- Data at rest inside the volume
- All data moving between the volume and the instance
- All snapshots created from the volume
- All volumes created from those snapshots

How is it encrypted:

- EBS encrypts your volume using an industry-standard AES-256 algorithm using a unique data key per volume.
- Your data key is stored on-disk with your encrypted data, but not before EBS encrypts it with your CMK. Your data key never appears on disk in plaintext.
- For more information, see Data keys in the *AWS Key Management Service Developer Guide*.

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

aws

Amazon S3 on AWS Outposts encryption

What is encrypted

- Data in transit and at Rest

How is it encrypted

- In Transit: Transport Layer Security (TLS)
- At Rest:
  - Industry-standard AES-256 algorithm
  - By default, server-side encryption with Amazon S3 managed encryption keys (SSE-S3).
    - S3 encrypts each object with a unique key. As an additional safeguard, it encrypts the key itself with a master key that it rotates regularly.
  - Optionally, customer can explicitly choose to store objects using server-side encryption with customer-provided encryption keys (SSE-C).
    - Amazon S3 does not store the encryption key you provide. Instead, it stores a randomly salted HMAC value of the encryption key to validate future requests.
    - The salted HMAC value cannot be used to derive the value of the encryption key or to decrypt the contents of the encrypted object. That means if you lose the encryption key, you lose the object.

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

Preventative Controls for Data Residency on Outposts leveraging Implementation of Service Control Policies and IAM Policies

Preventive controls can stop an action from ever succeeding.

The policies examples of how to prevent data from being copied from an Outpost to a Region

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

IAM Policy and SCP Overview

Identity-based policies

- Attach managed and inline policies to IAM identities (users, groups to which users belong, or roles).
- Identity-based policies grant or limit permissions to an identity.

Organizations SCPs

- Use an AWS Organizations service control policy (SCP) to define the maximum permissions for account members of an organization or organizational unit (OU).
- SCPs limit permissions that identity-based policies or resource-based policies grant to entities (users or roles) within the account, but do not grant permissions.

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

aws

EBS Snapshot: Deny snapshots from in-region storage

```
"Effect":"Deny",
"Action":”ec2:CreateSnapshot*”,
"Resource":"arn:aws:ec2:*::snapshot/*",
"Condition":{
    ”ArnLike":{
      "ec2:SourceOutpostArn”:"arn:aws:outposts:*:*:outpost/*”
    },
    "Null":{
      "ec2:OutpostArn":"true"
    }
}
```

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

aws

EBS Snapshot: Deny create of snapshots in-region: Test

Allowed:

- aws ec2 --profile test create-snapshot --volume-id vol-079f5e0a7a12c938e --outpost-arn arn:aws:outposts:us-west-2:{redacted}:outpost/op-0d4579457ff2dc345

Denied:

- aws ec2 --profile test create-snapshot --volume-id vol-079f5e0a7a12c938e

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

S3: Deny PutObject to region

```
{
   "Effect": "Deny"
    "Action": [
        "s3:PutObject"
    ],
    "Resource": ["arn:aws:s3:::*"]
}
```

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

aws

S3: Deny PutObject to region

Allowed:

- aws s3api put-object --bucket arn:aws:s3-outposts:${var.region}:{redacted}:outpost/op-01ac5d28a6a232904/accesspoint/test-data-residency --key testkey --body /etc/os-release

Denied:

- aws s3api put-object --bucket arn:aws:s3:${var.region}:{redacted}:accesspoint/test-data-residency --key testkey --body /etc/os-release
- aws s3api put-object --bucket test-data-residency --key testkey --body /etc/os-release

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

aws

Deny Launching Instances into in-region Subnets

```
"Effect":"Deny",
"Action": [
    "ec2:RunInstances" ,
    "ec2:CreateNetworkInterface"
],
"Resource": [
      "arn:aws:ec2:*:*:network-interface/*"
],
“Condition": {
    "ArnNotEquals": {
        "ec2:Subnet": "${data.aws_subnet.outpost-only.arn}"
      }
}
# Note: Terraform syntax for dynamic substitution of subnet id
```

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

aws

Deny copying of data to Region

```
"Effect": "Deny",
"Action": [
      "ec2:ModifyImageAttribute",
      "ec2:CopyImage",
      "ec2:CreateImage",
      "ec2:CreateInstanceExportTask",
      "ec2:ExportImage",
      "ec2:ImportImage",
      "ec2:ImportInstance",
      "ec2:ImportSnapshot",
      "ec2:ImportVolume",
      "rds:CreateDBSnapshot",
      "rds:CreateDBClusterSnapshot",
      "rds:ModifyDBSnapshotAttribute",
      "elasticache:CreateSnapshot",
      "elasticache:CopySnapshot",
      "datasync:Create*",
      "datasync:Update*" ],
"Resource": "*"
```

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

aws

Deny Direct Storage of data to Region

```
"Action": [
      "dynamodb:PutItem",
      "ec2:CreateTrafficMirrorSession",
      "es:Create*",
      "elasticfilesystem:C*",
      "storagegateway:Create*",
      "neptune-db:connect",
      "glue:CreateDevEndpoint",
      "datapipeline:CreatePipeline",
      "sagemaker:Create*",
      "redshift:CreateCluster*",
      "ses:Send*",
      "sqs:Send*",
      "mq:Create*",
      "ecr:Put*",
      "ecr:Create*",
      "ecr:Upload*"  ],
"Resource": "*",
"Effect": "Deny"
```

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

aws

Logging

1. By default, Outposts does not perform application level logging.
2. Cloudwatch Logs:
    1. If the CloudWatch Agent, or other usage of the logs:PutLogEvents API, is enabled, the log data will be sent to the region.
    2. In this case it is up to the customer to ensure that their applications logs do not contain any residency-impacted data
    3. If you can't accept that risk, then add the logs:PutLogEvents action in a Deny policy.
3. Services which can be configured to use Cloudwatch Logs:
    1. Systems Manager
        1. EC2 Instances with both the SSM Agent and CloudWatch Agent enabled
        2. Send-Command via the option cloud-watch-output-config
        3. Start-session
    2. ECS and Docker via the awslogs driver

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

aws

Additional controls: Network

1. Network API actions, for example ec2:CreateSubnet, should be subject to Segregation of Duties and Least Privilege to reduce risk of data transfer to in-region resources.
    1. Can’t create in-region resources, such as RDS instances, if there are no in-region subnets.
2. Delete in-region subnets if you don’t need them for specific reasons
3. If a workload does require in-region subnets:
    1. Use NACLs and security-groups to tightly control what components are allowed to communicate.
    2. If needed for VPC Endpoints, then have a IAM policy and/or SCP which denies ec2:Runstance, to those subnets.

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

Other Design Considerations

1. On-Premises Storage
    - Data can be stored outside of AWS Outposts
    - Outpost Service Ready Partner examples: Infinidat, NetApp, PureStorage, Zadara
2. On-premises Backups
    - When your data protection requirements mandate backups to be stored Outside of the Outpost, but your data residency requirements eliminate the option of in-region backups, then use application specific or OS Agent based backup solutions which enable targeting on-prem or colo based locations.
    - Outpost Service Ready Partner examples: Commvault, NetApp, Veritas
3. Disaster Recovery
    - Replication thru LGW:
        - Application level: database native replication, Redis replication
        - OS Level: CloudEndure Disaster Recovery

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

aws

AWS Outposts Data Residency model recap

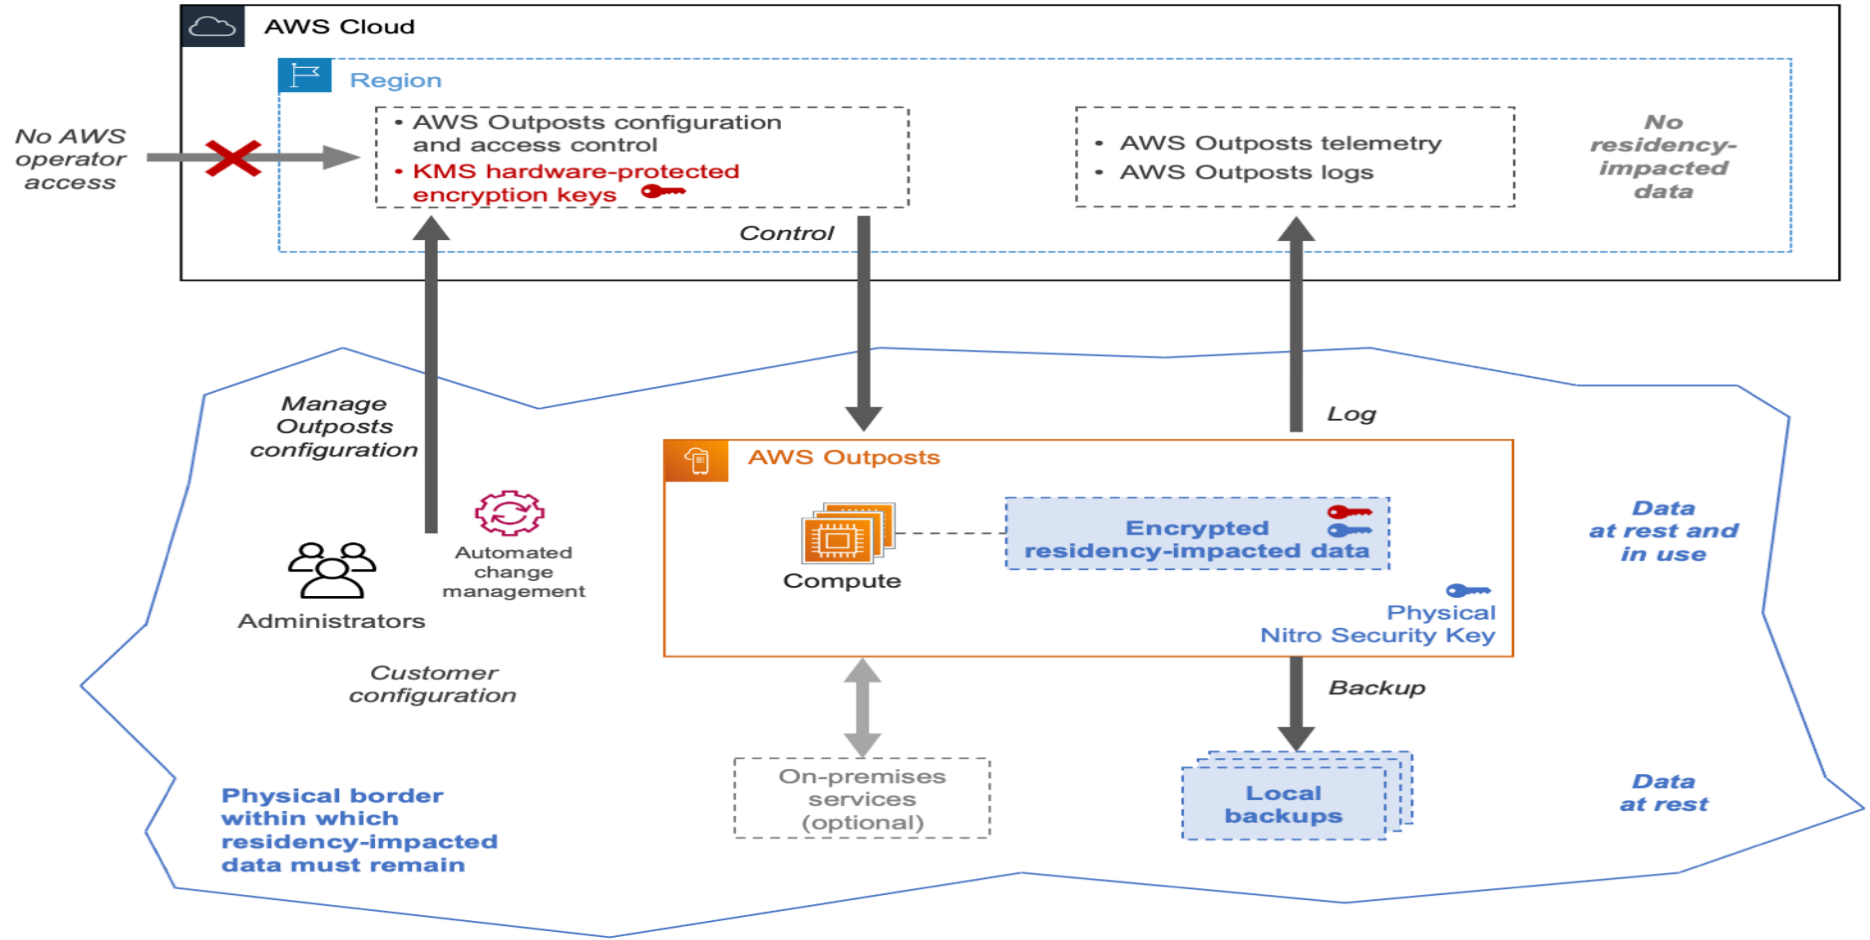

Conclusion

1. Outposts can be leveraged to meet Data Residency use cases
2. Use the policies discussed in this webinar as part of your Data Residency strategy.
3. If you would like to see additional webinar content for KMS Key Policies or Detective Controls relevant to Data Residency, please let us know.

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.

aws

Closing

How to Reach us:

John Mallory
johmallo@amazon.com

David Filiatrault
filiatra@amazon.com

Time for Q&A

© 2021, Amazon Web Services, Inc. or its Affiliates. All rights reserved.