

Cloud Training

Quản trị triển khai điện toán đám mây với AWS

Instructor: Hieu Vu



Agenda

- Module 1 Cloud Computing Implementation
- Module 2 Maintenance - Backup and Restore
- Module 3 Security and Implementation
- Module 4 Identify and Management
- Module 5 DNS, Caching and Performance Optimization
- Module 6 Troubleshooting in AWS

Module 1: Cloud Computing Implementation



Agenda

- AWS Introduction
- AWS Compute Overview
- AWS Networking Overview
- AWS Elastic Load Balancing

What is AWS?

AWS provides a highly reliable, scalable, low-cost infrastructure platform in the cloud that powers millions of businesses in over 190 countries around the world.

■ Benefits

- Low Cost
- Elasticity & Agility
- Open & Flexible
- Secure
- Global Reach



The advertisement features the AWS logo in the top left corner. The background is a dark blue globe with a network of white lines representing global connectivity. The text reads: "The Most Extensive, Reliable and Secure Global Cloud Infrastructure Available". Below this text is a button that says "SEE HOW WE DO IT >>". At the bottom, there is a small paragraph of text: "The Amazon Web Services (AWS) Global Infrastructure delivers a cloud infrastructure companies can depend on—no matter their size, changing needs, or challenges. The AWS Global Infrastructure is designed and built to deliver the most flexible, reliable, scalable, and secure cloud computing environment with the highest quality global".

Figure 1. Magic Quadrant for Cloud Infrastructure as a Service, Worldwide



Source: Gartner (July 2019)

AWS Recognized as a Cloud Leader for the 9th Consecutive Year

Gartner, Magic Quadrant for Cloud Infrastructure as a Service, Worldwide, Raj Bala, Bob Gill, Dennis Smith, David Wright, July 2019. ID G00365830. Gartner does not endorse any product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose. The Gartner logo is a trademark and service mark of Gartner, Inc., and/or its affiliates, and is used herein with permission. All rights reserved.



The AWS Platform

- Account Support
- Support
- Managed Services
- Professional Services
- Partner Ecosystem
- Training & Certification
- Solution Architects
- Account Management
- Security & Pricing Reports
- Technical Acct. Management

Marketplace Business Applications	Mgmt. Tools Monitoring	Analytics Query Large Data Sets	Dev Tools	Artificial Intelligence	IoT	Mobile	Enterprise Applications	Game Development
DevOps Tools	Auditing	Elasticsearch						
Business Intelligence	Service Catalog	Business Analytics						
Security	Server Management	Hadoop/Spark			Rules Engine	Build, Test, Monitor Apps	Document Sharing	
Networking	Configuration Tracking	Real-time Data Streaming	Private Git Repositories	Voice & Text Chatbots	Local Compute and Sync	Push Notifications	Email & Calendaring	
Database & Storage	Optimization	Orchestration Workflows	Continuous Delivery	Machine Learning	Device Shadows	Build, Deploy, Manage APIs	Hosted Desktops	
SaaS Subscriptions	Resource Templates	Managed Search	Build, Test, and Debug	Text-to-Speech	Device Gateway	Device Testing	Application Streaming	3D Game Engine
Operating Systems	Automation	Managed ETL	Deployment	Image Analysis	Registry	Identity	Backup	Multi-player Backends
Migration	Application Discovery	Application Migration	Data Migration	Database Migration	Server Migration			
Hybrid	Data Integration	Integrated Networking	Identity Federation	Resource Management	VMware on AWS	Devices & Edge Systems		
Application Services	Transcoding	Step Functions	Messaging					
Security	Identity & Access	Key Storage & Management	Active Directory	DDoS Protection	Application Analysis	Certificate Management	Web App. Firewall	
Database	Aurora	MySQL	PostgreSQL	Oracle	SQL Server	MariaDB	Data Warehousing	NoSQL
Storage	Object Storage	Archive	Exabyte-scale Data Transport	Block Storage	Managed File Storage			
Compute	Virtual Machines	Simple Servers	Web Applications	Auto Scaling	Batch	Containers	Event-driven Computing	
Networking	Isolated Resources	Dedicated Connections	Global CDN	Load Balancing	Scalable DNS			
Infrastructure	Regions	Availability Zones	Points of Presence					



22 Regions

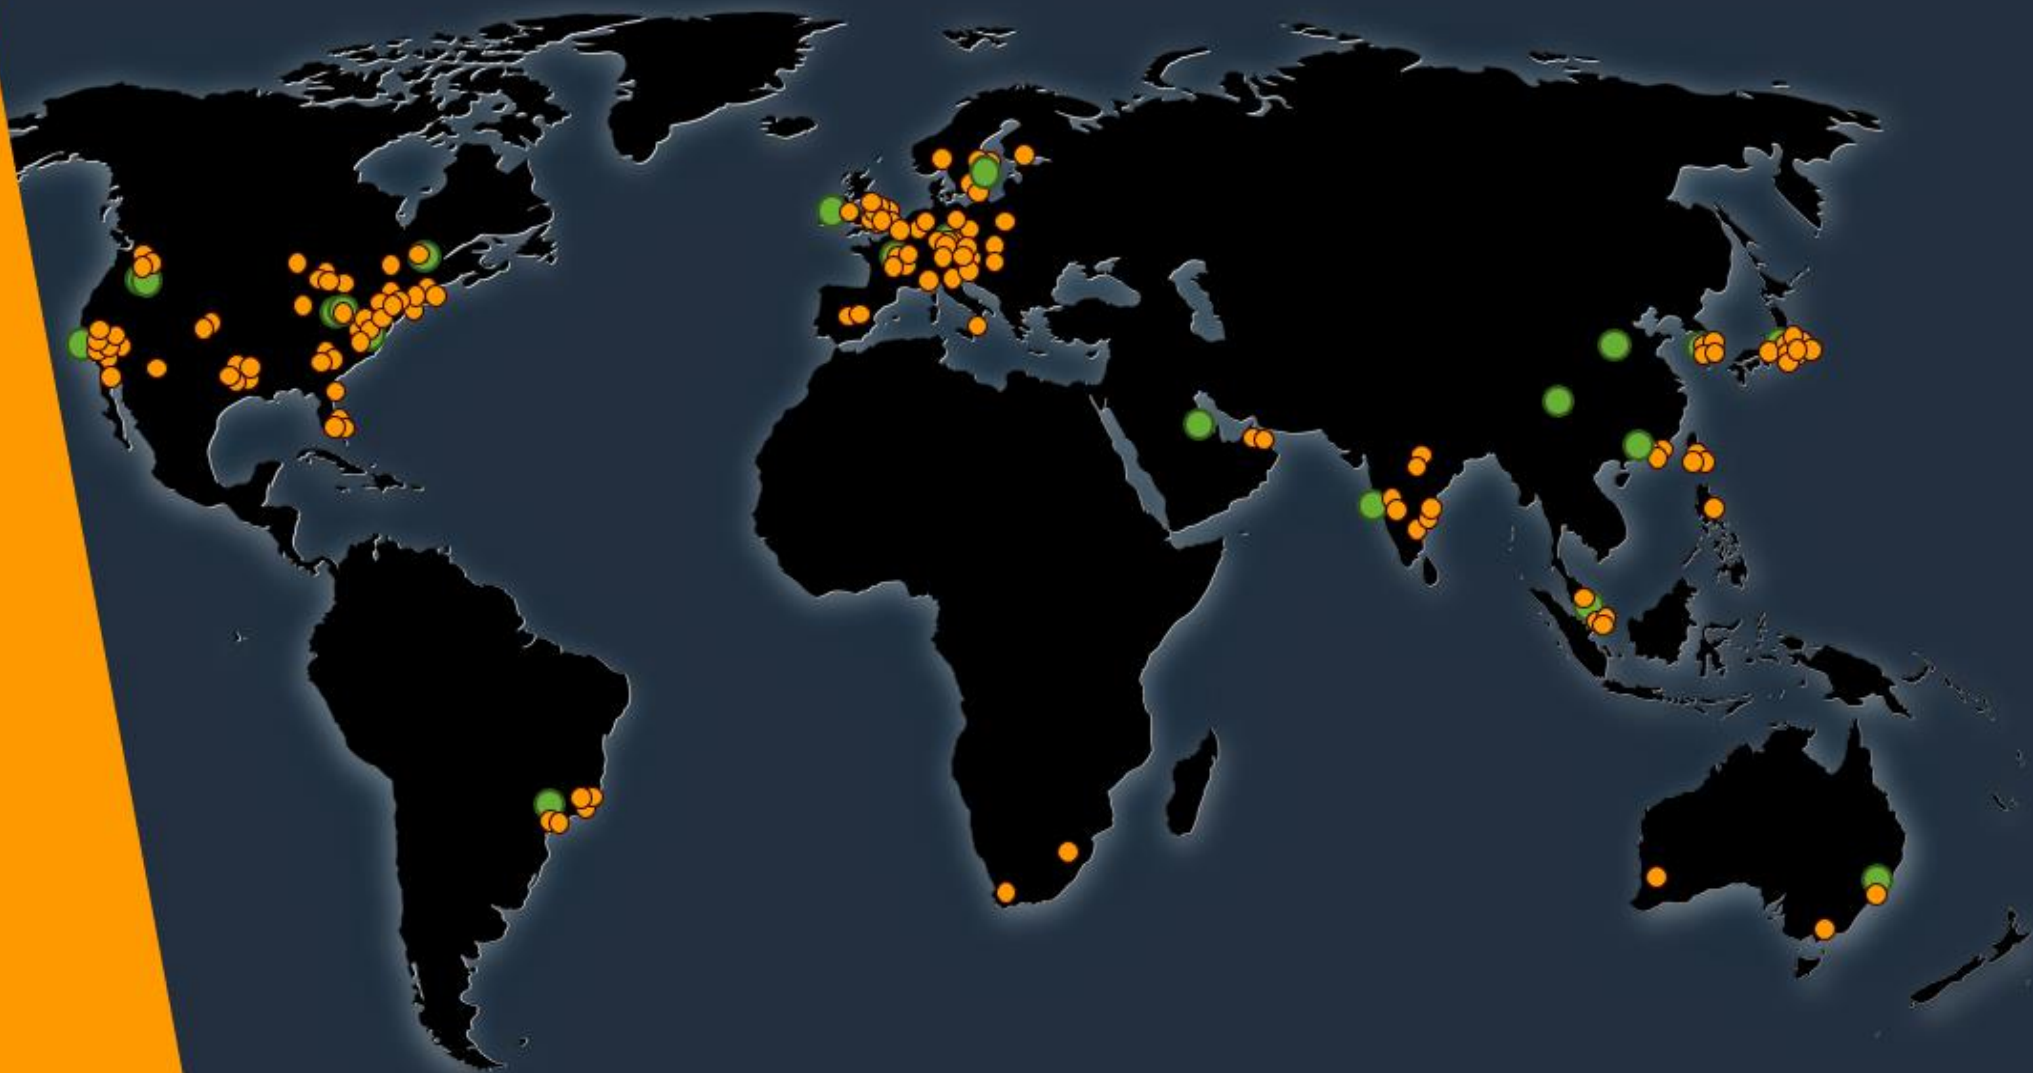


VIỆN ĐIỆN TỬ - VIỄN THÔNG
School of Electronics and Telecommunications



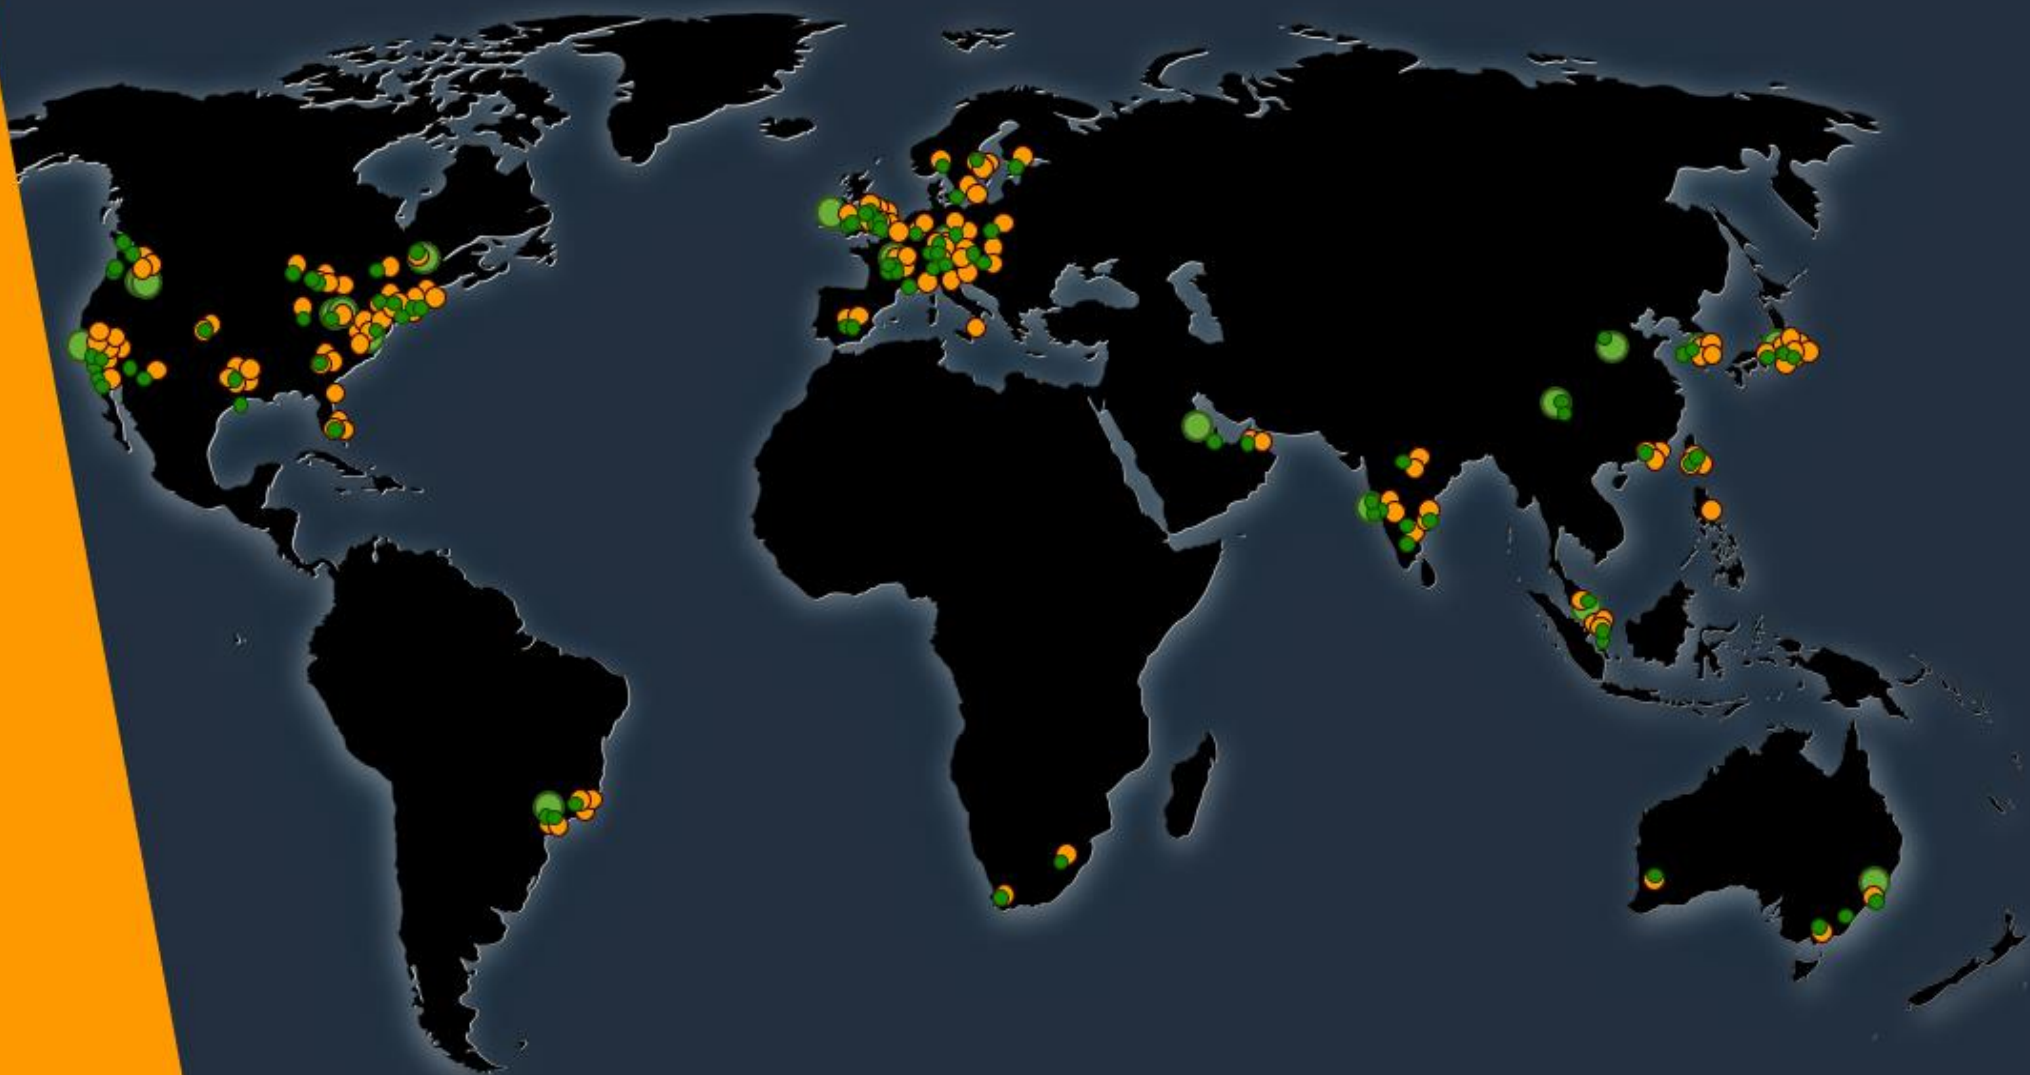
210

Amazon
CloudFront
Points of
Presence



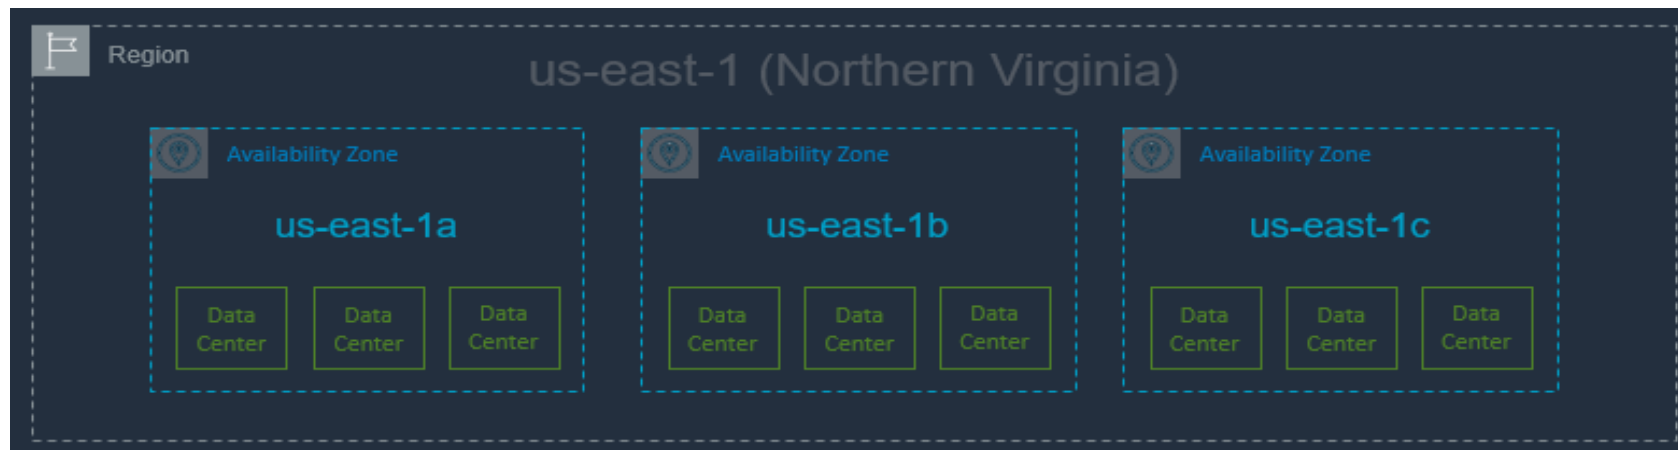
100

AWS Direct
Connect
locations

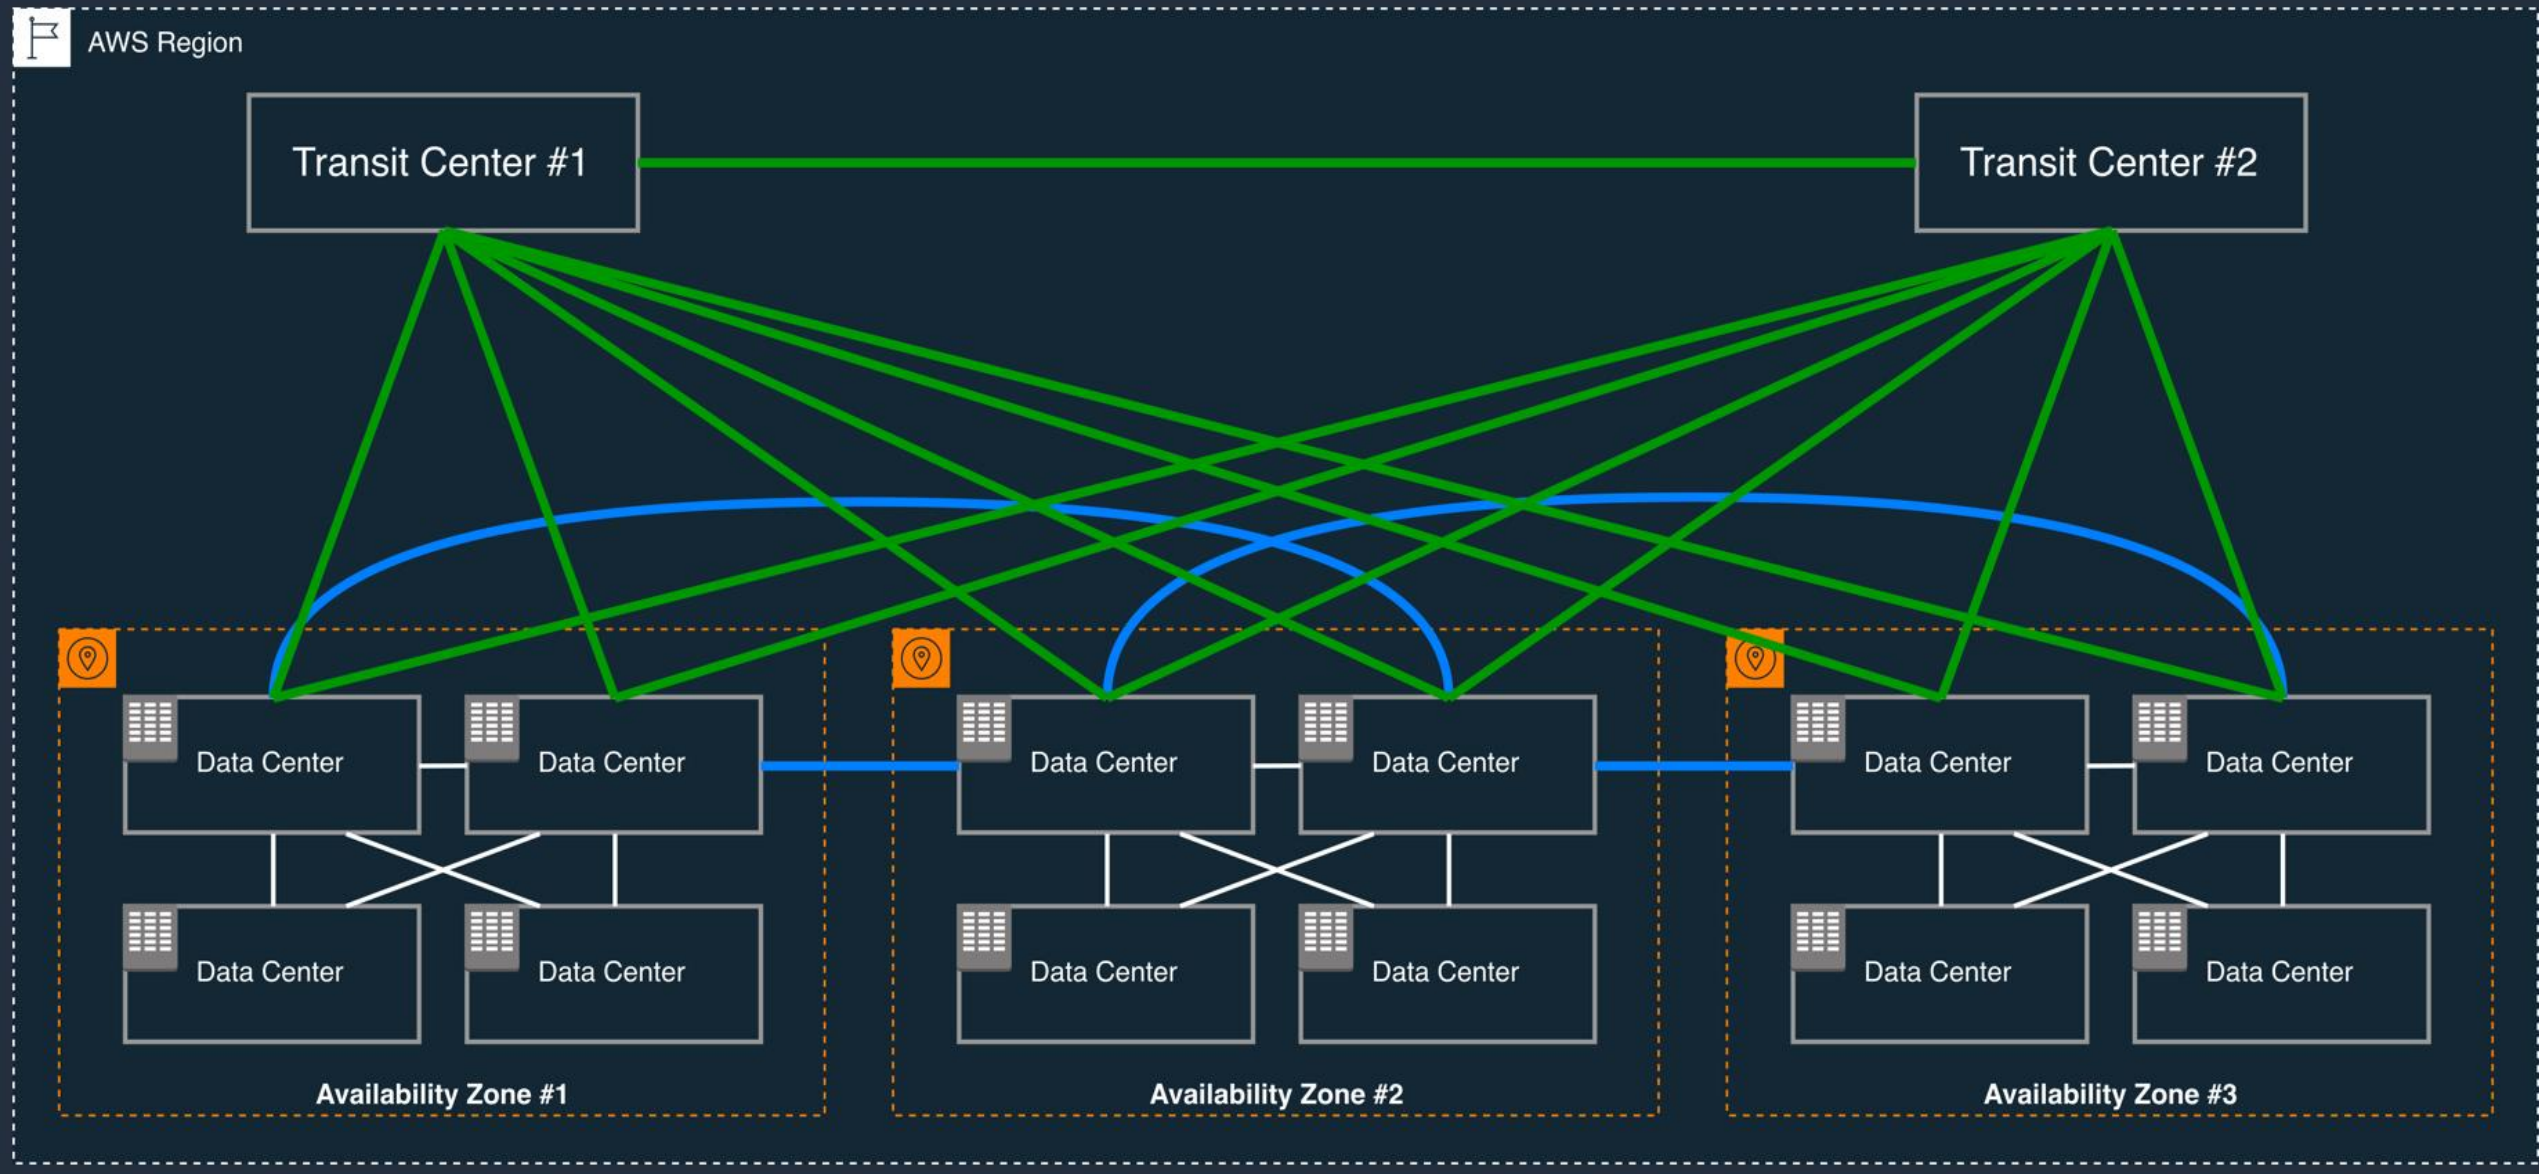


Availability Zones

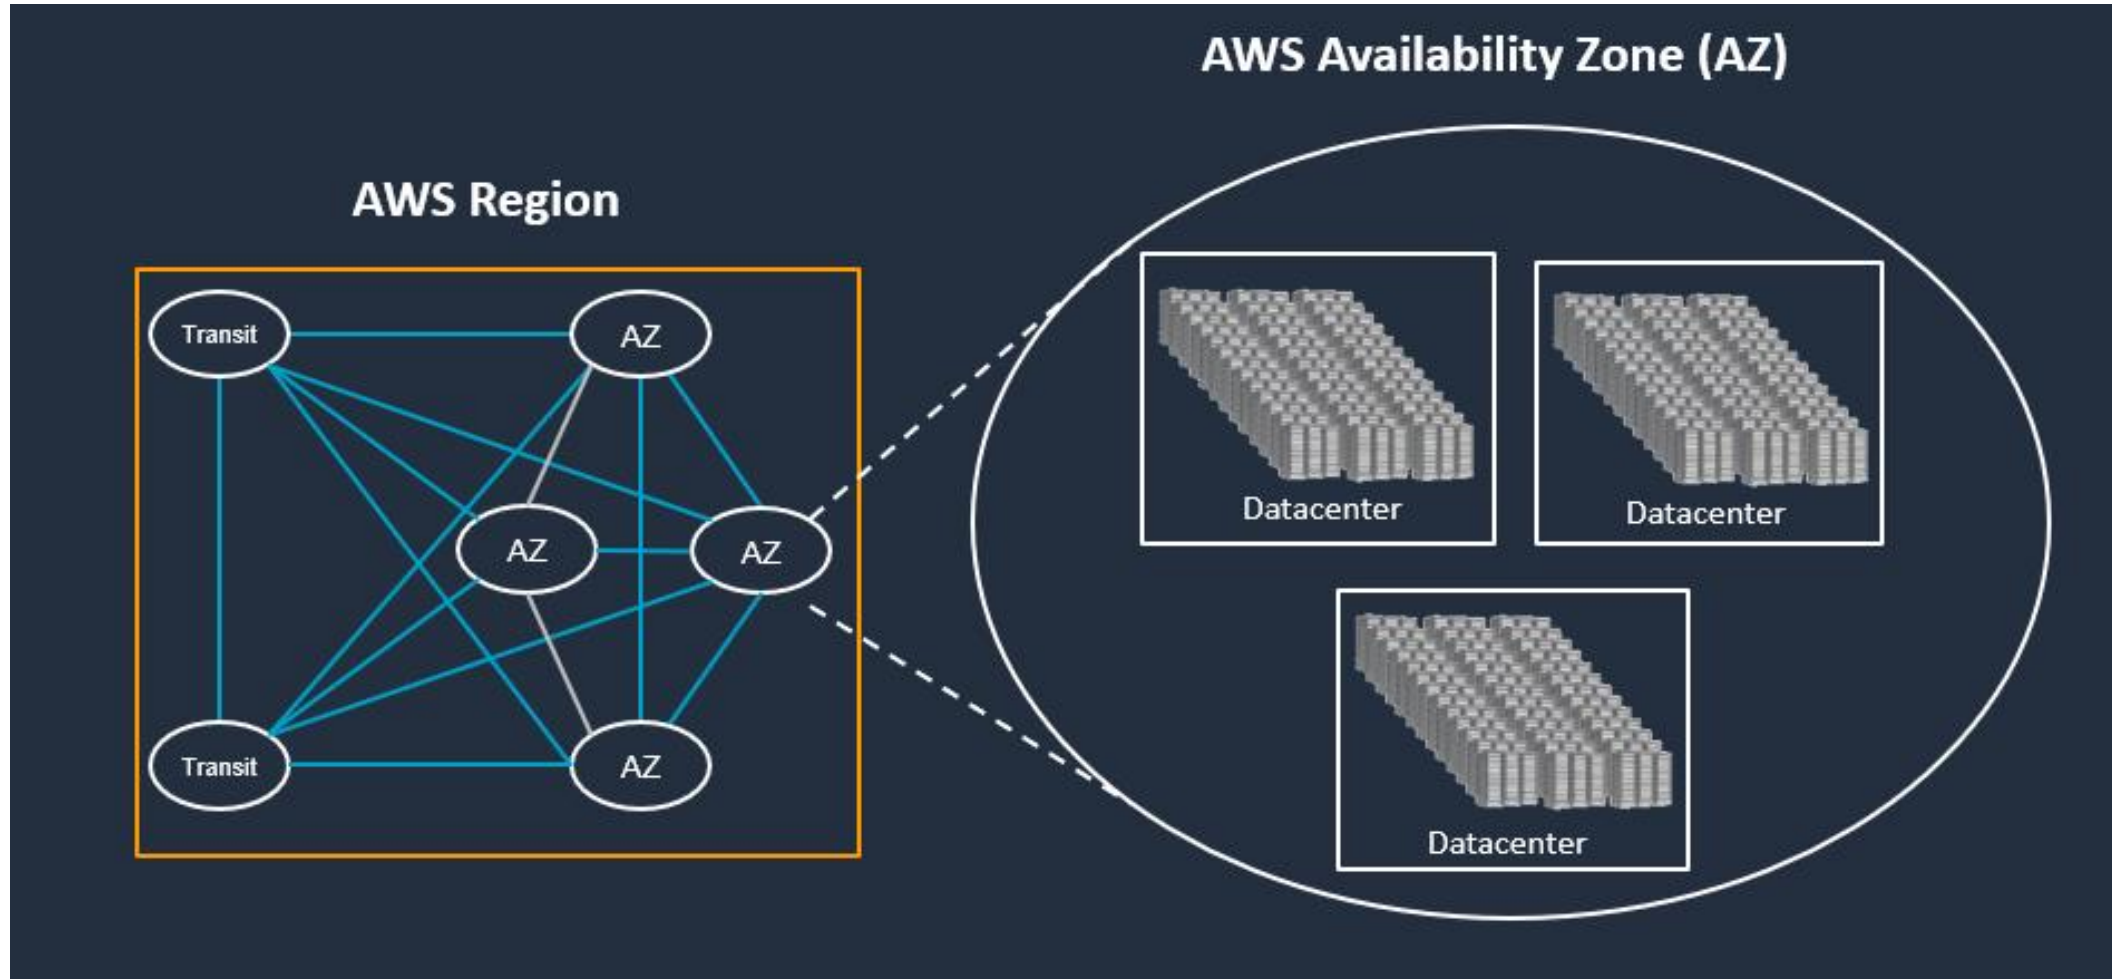
- A region is comprised of multiple Availability Zones (typically 3)
- Fully independent partitions on isolated fault lines, flood plains, and power grids
- Each AZ: redundant power and redundant dedicated network
- Each AZ: typically multiple data centers
- Between AZs: high throughput, low latency (<10ms) network
- Between AZs: physical separation < 100km (60mi)



Availability Zones



Availability Zones



AWS Cloud Services



Services ▲

Search for services, features, marketplace products, and docs

[Alt+S]



cloud_user @ 2122-3795-6973 ▼

N. Virginia ▼

Support

★ Favorites



Resource Groups & Tag Editor

Recently visited

Console Home

All services

Compute

EC2

Lightsail ↗

Lambda

Batch

Elastic Beanstalk

Serverless Application Rep...

AWS Outposts

EC2 Image Builder

AWS App Runner

Containers

Elastic Container Registry

Elastic Container Service

Elastic Kubernetes Service

Red Hat OpenShift Service ...

Storage

S3

EFS

FSx

Customer Enablement

AWS IQ ↗

Support

Managed Services

Activate for Startups

Robotics

AWS RoboMaker

Blockchain

Amazon Managed Blockchain

Satellite

Ground Station

Quantum Technologies

Amazon Braket

Management & Governance

AWS Organizations

CloudWatch

Machine Learning

Amazon SageMaker

Amazon Augmented AI

Amazon CodeGuru

Amazon DevOps Guru

Amazon Comprehend

Amazon Forecast

Amazon Fraud Detector

Amazon Kendra

Amazon Lex

Amazon Personalize

Amazon Polly

Amazon Rekognition

Amazon Textract

Amazon Transcribe

Amazon Translate

AWS DeepComposer

AWS DeepLens

AWS DeepRacer

AWS Panorama

Amazon Monitron

AWS Cost Management

AWS Cost Explorer

AWS Budgets

AWS Marketplace Subscript...

AWS Application Cost Profiler

Front-end Web & Mobile

AWS Amplify

Mobile Hub

AWS AppSync

Device Farm

Amazon Location Service

AR & VR

Amazon Sumerian

Application Integration

Step Functions

Amazon AppFlow

Amazon EventBridge

Amazon MQ



VIỆN ĐIỆN TỬ - VIỆN THÔNG
School of Electronics and Telecommunications



BKACAD


AWS Core Infrastructure and Services

Traditional Infrastructure

Amazon Web Services



Security



Networking



Servers



Storage & Database



AWS Foundation Services

Compute

- Amazon EC2
- Amazon EC2 Container Registry
- Amazon EC2 Container Service
- Amazon Lightsail
- Amazon VPC
- AWS Batch
- AWS Elastic Beanstalk
- AWS Lambda
- Elastic Load Balancing

Network

- Amazon CloudFront
- Amazon Route 53
- Amazon VPC
- AWS Direct Connect
- Elastic Load Balancing

Storage

- Amazon EFS
- Amazon Glacier
- Amazon S3
- AWS Snowball
- AWS Storage Gateway






































Security & Identity

- Amazon Inspector
- AWS Artifact
- AWS Certificate Manager
- AWS CloudHSM
- AWS Directory Service
- IAM
- AWS KMS
- AWS Organizations
- AWS Shield
- AWS WAF

Applications

- Amazon WorkDocs
- Amazon WorkMail
- Amazon AppStream
- Amazon WorkSpaces

AWS Platform Services

Databases	Analytics	Application Services	Management Tools	Developer Tools	Mobile Services	Internet of Things
 Amazon DynamoDB	 Amazon Athena	 Amazon API Gateway	 Amazon CloudWatch	 AWS CodeBuild	 Amazon API Gateway	 AWS IoT
 Amazon ElastiCache	 Amazon CloudSearch	 Amazon AppStream 2.0	 AWS CloudFormation	 AWS CodeCommit	 Amazon Cognito	 AWS Greengrass
 Amazon RDS	 Amazon EMR	 Amazon Elastic Transcoder	 AWS CloudTrail	 AWS CodeDeploy	 Amazon Mobile Analytics	
 Amazon Redshift	 Amazon ES	 Amazon SWF	 AWS Config	 AWS CodePipeline	 Amazon Pinpoint	
	 Amazon Kinesis	 AWS Step Functions	 AWS Managed Services	 AWS X-Ray	 AWS Device Farm	
	 Amazon QuickSight		 AWS OpsWorks		 AWS Mobile Hub	
	 Amazon Redshift		 AWS Service Catalog			
			 AWS Trusted Advisor			


AWS Compute Overview



VIỆN ĐIỆN TỬ - VIỄN THÔNG
School of Electronics and Telecommunications



Choices for Compute



Amazon EC2

Virtual server instances
in the cloud



Amazon ECS, EKS, and Fargate


Container management service
for running
Docker on a managed
cluster of EC2



AWS Lambda

Serverless compute
for stateless code execution in
response to triggers

Amazon EC2



Amazon EC2

Linux | Windows

Arm and x86 architectures


General purpose and workload optimized

Bare metal, disk, networking capabilities

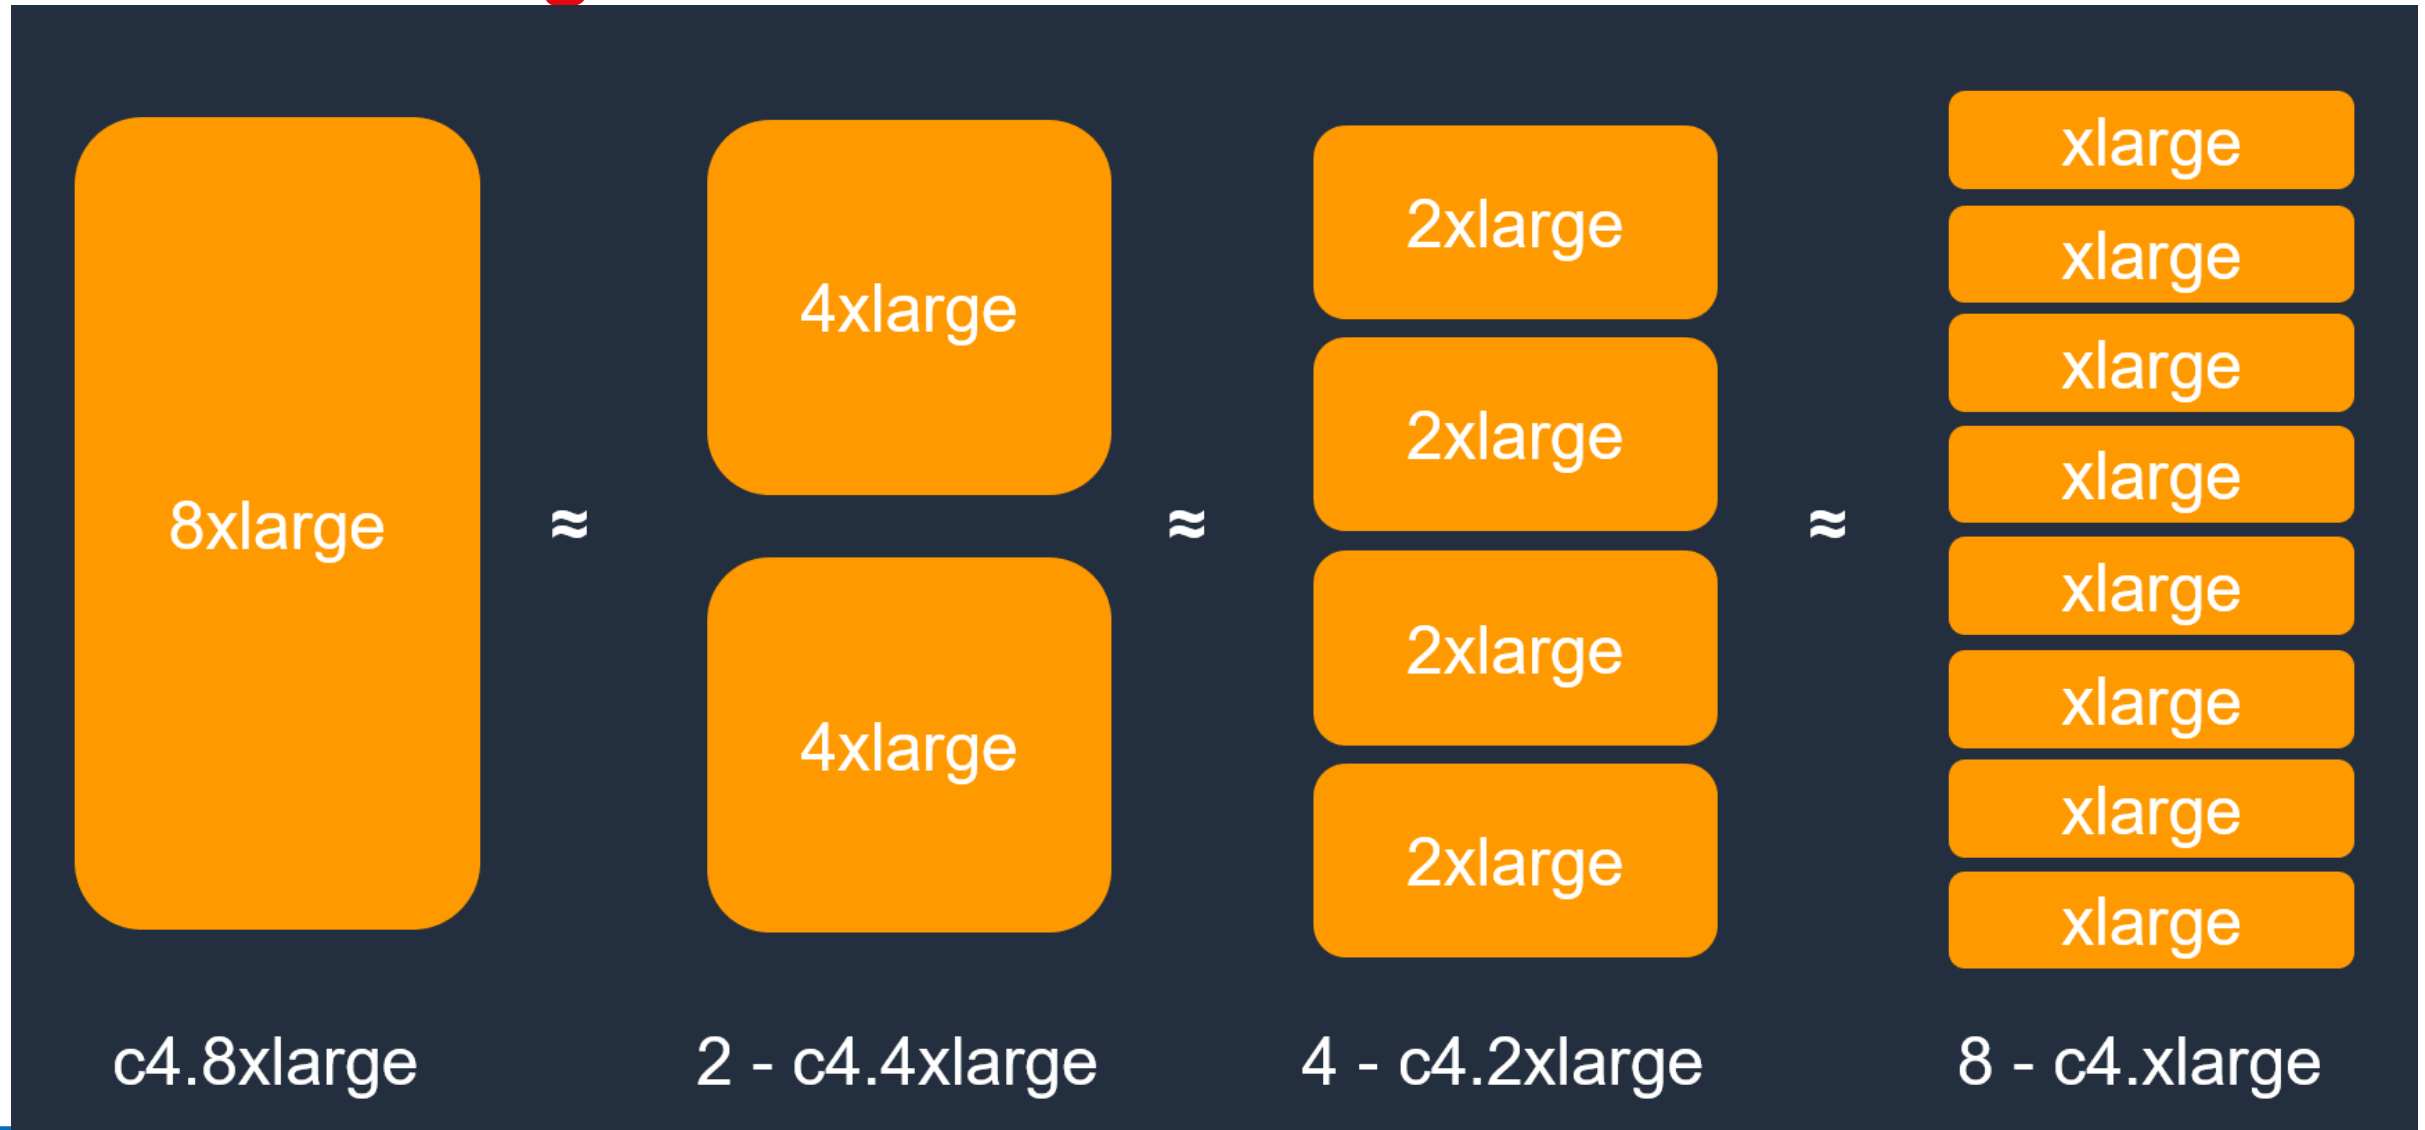
Packaged | Custom | Community AMIs

Multiple purchase options: On-demand, RI, Spot

EC2 Terminology



Instance Sizing



EC2 Naming Explained

Instance generation

c5n.xlarge

Instance
family

Attribute

Instance size

Instance Types



EC2 Elastic GPUs

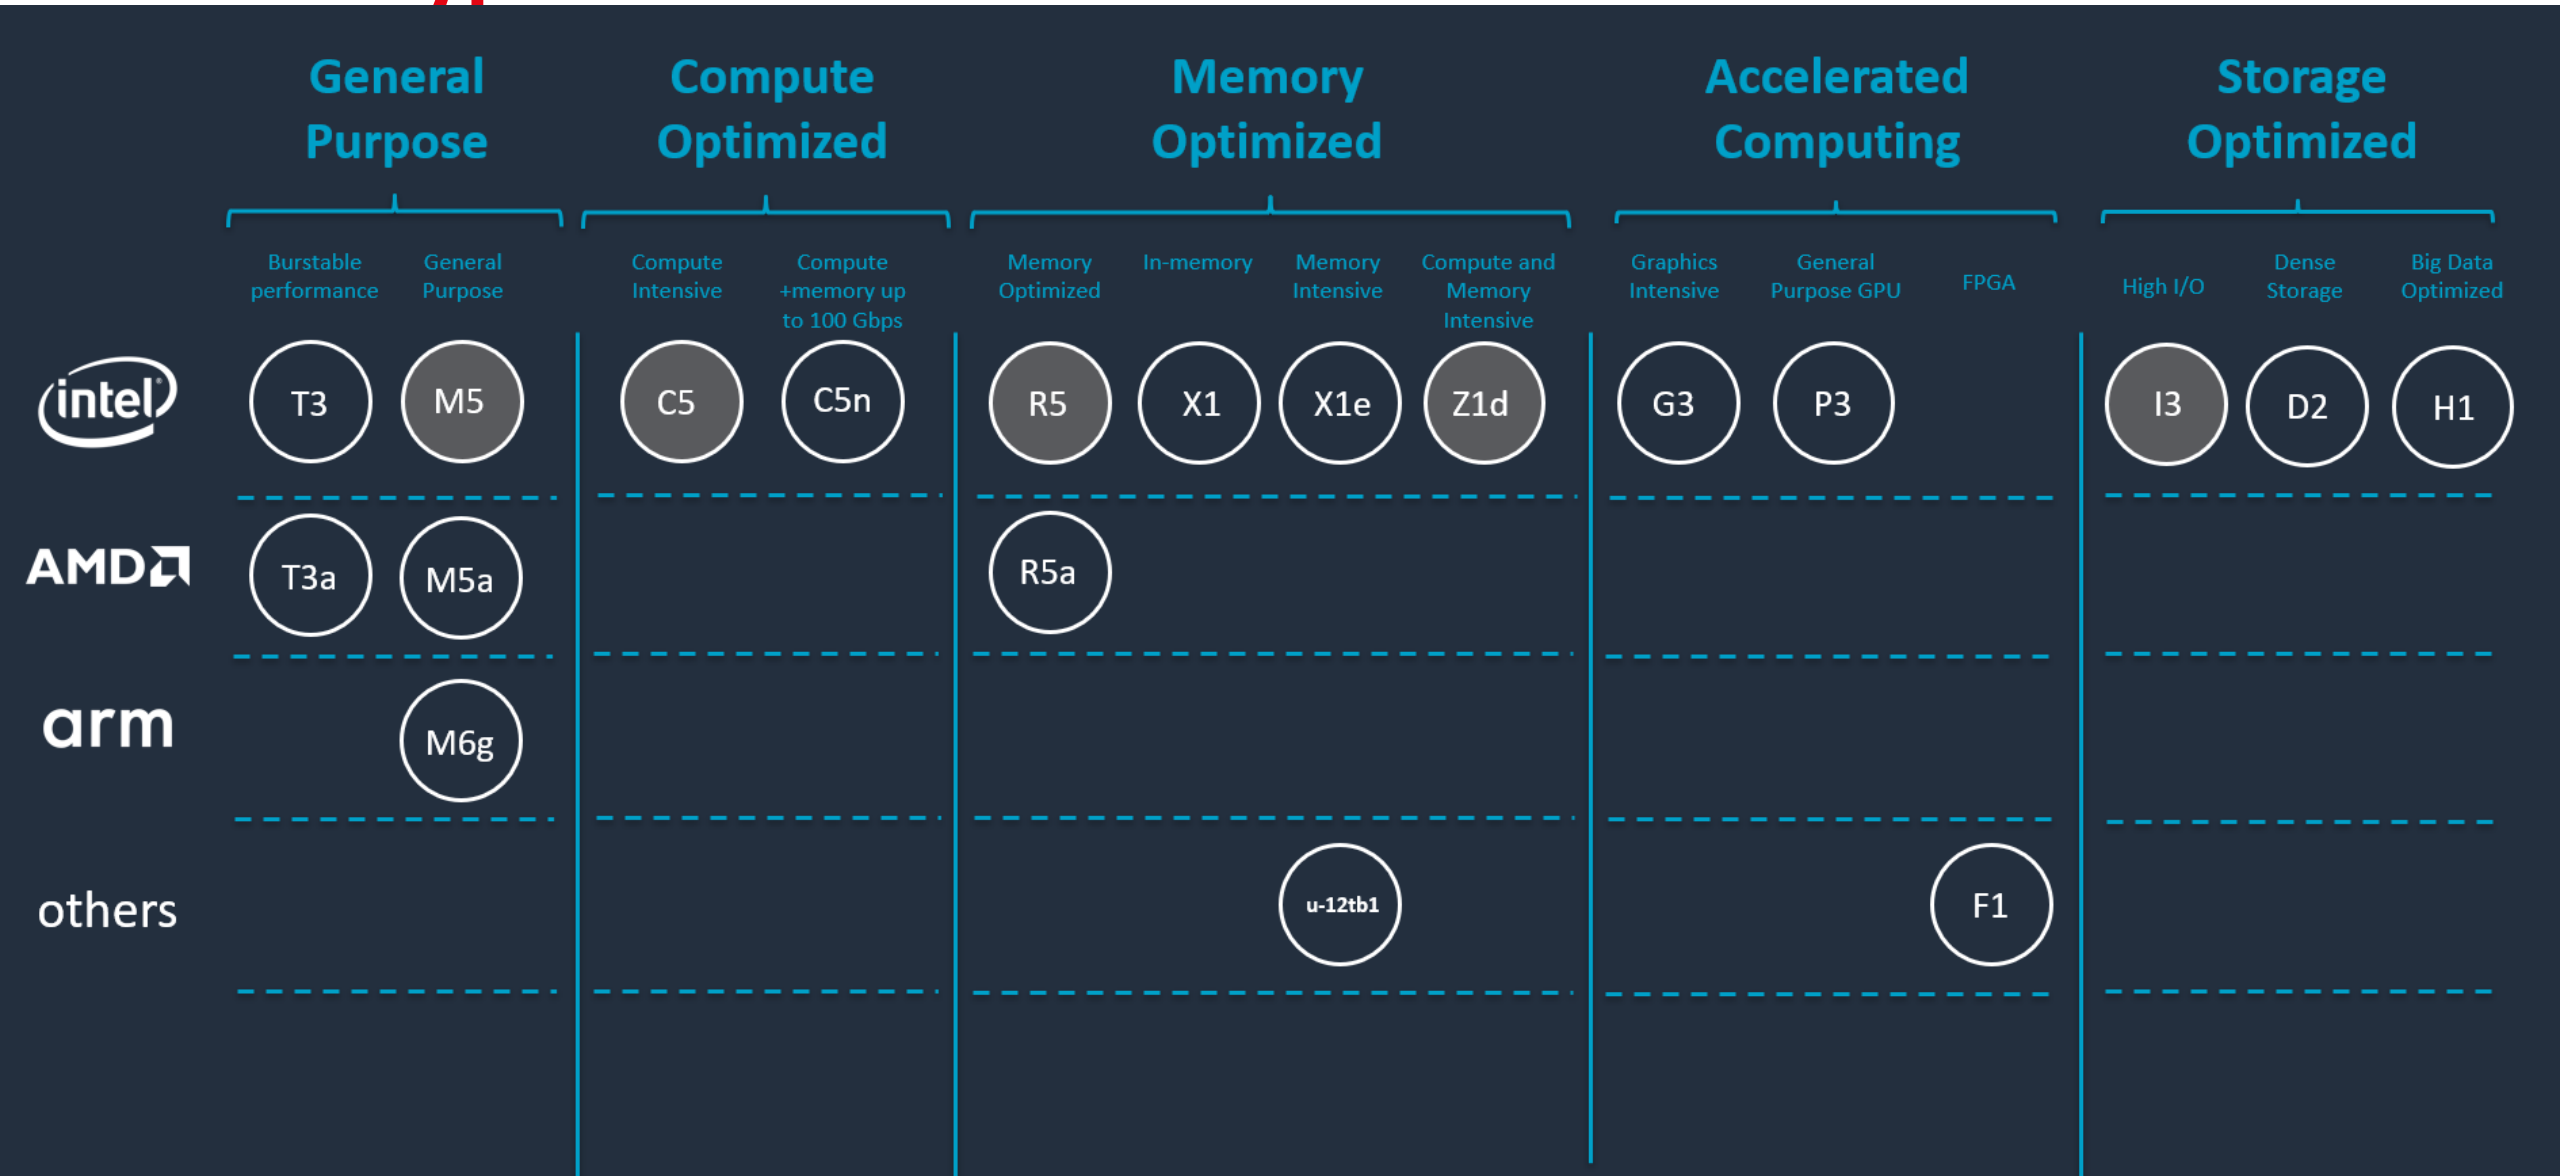
- Graphics acceleration for EC2 instances



EC2 Fleet


- Simplified provisioning
- Massive scale
- Flexible capacity allocation

Instance Types



EC2 Operating Systems Supported

- Windows 2003R2/2008/2008R2/2012/2012R2/2016/2019
- Amazon Linux
- Debian
- Suse
- CentOS
- Red Hat Enterprise Linux
- Ubuntu




for more OSes see: <https://aws.amazon.com/marketplace/b/2649367011>

EC2 AMI

■ AMI

- Instances are based on an Amazon Machine Image
- You can create new AMIs from a running instance
- AMIs are stored in S3
- AMIs are unique to each region



EC2 AMI

AWS Console

1. Choose AMI 2. Choose Instance Type 3. Configure Instance 4. Add Storage 5. Add Tags 6. Configure Security Group 7. Review

Step 1: Choose an Amazon Machine Image (AMI)

An AMI is a template that contains the software configuration (operating system, application server, and applications) required to launch your instance. You can select an AMI provided by AWS, our user community, or the AWS Marketplace; or you can select one of your own AMIs.

Quick Start 1 to 35 of 35 AMIs

- Amazon Linux 2 AMI (HVM), SSD Volume Type - ami-04681a1dbd79675a5** 64-bit
Amazon Linux 2 comes with five years support. It provides Linux kernel 4.14 tuned for optimal performance on Amazon EC2, systemd 219, GCC 7.3, Glibc 2.26, Binutils 2.29.1, and the latest software packages through extras.
Root device type: ebs Virtualization type: hvm ENA Enabled: Yes
- Amazon Linux AMI 2018.03.0 (HVM), SSD Volume Type - ami-0ff8a91007f77f867** 64-bit
The Amazon Linux AMI is an EBS-backed, AWS-supported image. The default image includes AWS command line tools, Python, Ruby, Perl, and Java. The repositories include Docker, PHP, MySQL, PostgreSQL, and other packages.
Root device type: ebs Virtualization type: hvm ENA Enabled: Yes
- Red Hat Enterprise Linux 7.5 (HVM), SSD Volume Type - ami-6871a115** 64-bit
Red Hat Enterprise Linux version 7.5 (HVM), EBS General Purpose (SSD) Volume Type
Root device type: ebs Virtualization type: hvm ENA Enabled: Yes

AWS Marketplace

aws marketplace

View Categories Migration Mapping Assistant Your Saved List Sell in AWS Marketplace Amazon Web Services Home Help

Operating Systems (336 results) showing 1 - 10

- CentOS 7 (x86_64) - with Updates HVM** 64-bit
★★★★★ (58) | Version 1805_01 | Sold by Centos.org
This is the Official CentOS 7 x86_64 HVM image that has been built with a minimal profile, suitable for use in HVM instance types only. The image contains just enough packages...
Linux/Unix, CentOS 7 - 64-bit Amazon Machine Image (AMI)
- CentOS 6 (x86_64) - with Updates HVM** 64-bit
★★★★★ (33) | Version 1805_01 | Sold by Centos.org
This is the Official CentOS 6 x86_64 HVM image that has been built with a minimal profile. The image contains just enough packages to run within AWS, bring up an SSH Server...
Linux/Unix, CentOS 6 - 64-bit Amazon Machine Image (AMI)
- Debian GNU/Linux 8 (Jessie)** 64-bit
★★★★★ (86) | Version 8.7 | Sold by Debian
Debian is a computer operating system composed of software packages released as free and open source software primarily under the GNU General Public License along with other...
Linux/Unix, Debian 8.6+1 - 64-bit Amazon Machine Image (AMI)
- CentOS 6.5 (x86_64) - Release Media** 64-bit
★★★★★ (55) | Version 6.5 - 2013-12-01 | Sold by Centos.org
This is the Official CentOS 6.5 x86_64 image that has been built with a minimal profile. The image contains just enough packages to run within AWS, bring up an SSH Server...

Use the AMI ID to launch through the API or AWS Command Line Interface (AWS CLI)


```
aws ec2 run-instances --image-id ami-04681a1dbd79675a5 --instance-type c4.8xlarge --count 10 --key-name MyKey
```



EC2 Security Groups

■ Security Group Rules

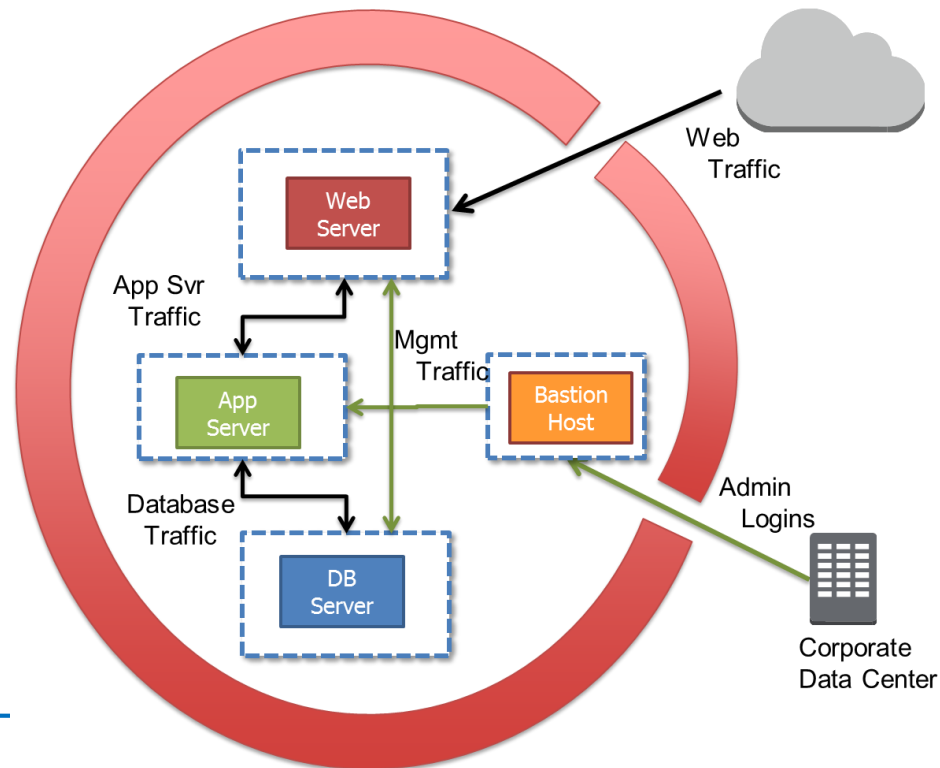
- Name
- Description
- Protocol
- Port range
- IP address, IP range, Security Group name



Tiered EC2 Security Groups

■ Hierarchical Security Group Rules

- Dynamically created rules
- Based on Security Group membership
- Create tiered network architectures



“Web” Security Group:

TCP 80 0.0.0.0/0
TCP 22 “Mgmt”

“App” Security Group:


TCP 8080 “Web”
TCP 22 “Mgmt”

“DB” Security Group:

TCP 3306 “App”
TCP 22 “Mgmt”

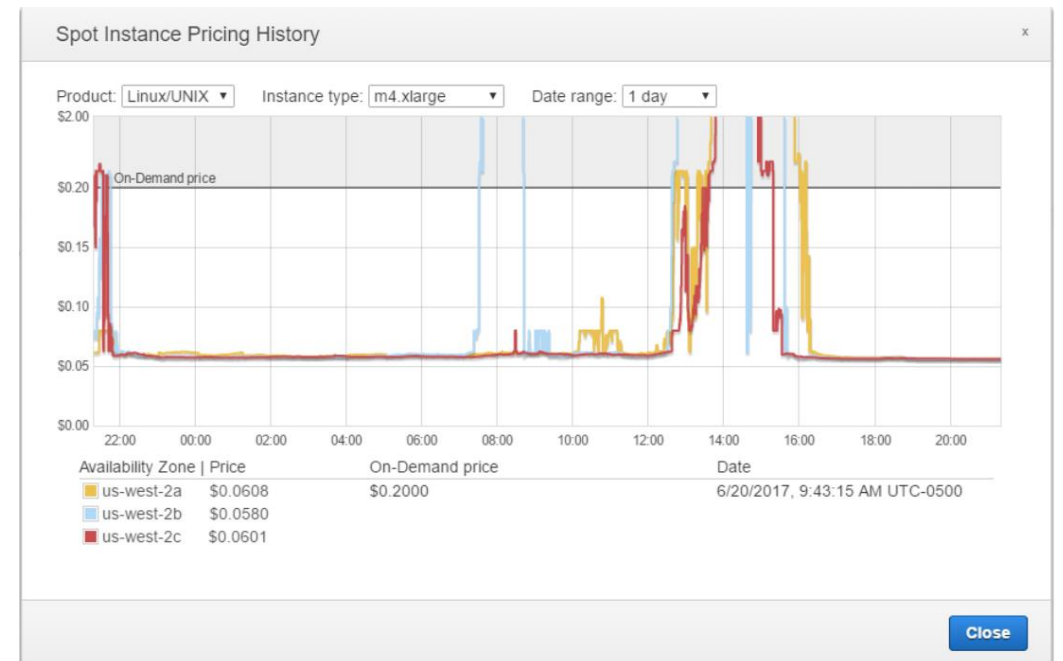
“Mgmt” Security Group:

TCP 22 163.128.25.32/32



EC2 Pricing

- On Demand Instance
 - This is the most common and flexible pricing option
 - Pay only for what you use
 - Stopped instances will not accrue hourly compute costs
 - Pay by the instance hour
- Reserved Instance (RI)
 - 1 or 3 year commitment
- Spot
 - Useful for “worker pool” scenarios
 - Transcode, map reduce task nodes



EC2 Purchasing Options

On-Demand

Pay for compute capacity by **the second** with no long-term commitments


For Spiky workloads or to define needs



Reserved Instances

Make a 1 or 3-year commitment and receive a **significant discount** off On-Demand prices

For committed utilization



Spot Instances

Spare EC2 capacity at **savings of up to 90%** off On-Demand prices

For time-insensitive or transient workloads
Need to be Fault-tolerant, stateless



Savings Plans

Commit to a \$/h spend and **share discount** across compute options and regions


For committed utilization



To optimize EC2, combine all three purchase options!



EC2 - capacity and cost optimization



Scale using
Spot,
On-Demand,
or both

Use **Reserved Instances**
for known/steady-state
workloads

AWS services make this easy and efficient



Amazon EC2
Auto Scaling



EC2 Fleet



Amazon Elastic
Container Service



Amazon Elastic
Container Service
for Kubernetes



AWS
Thinkbox



Amazon
EMR



AWS
CloudFormation

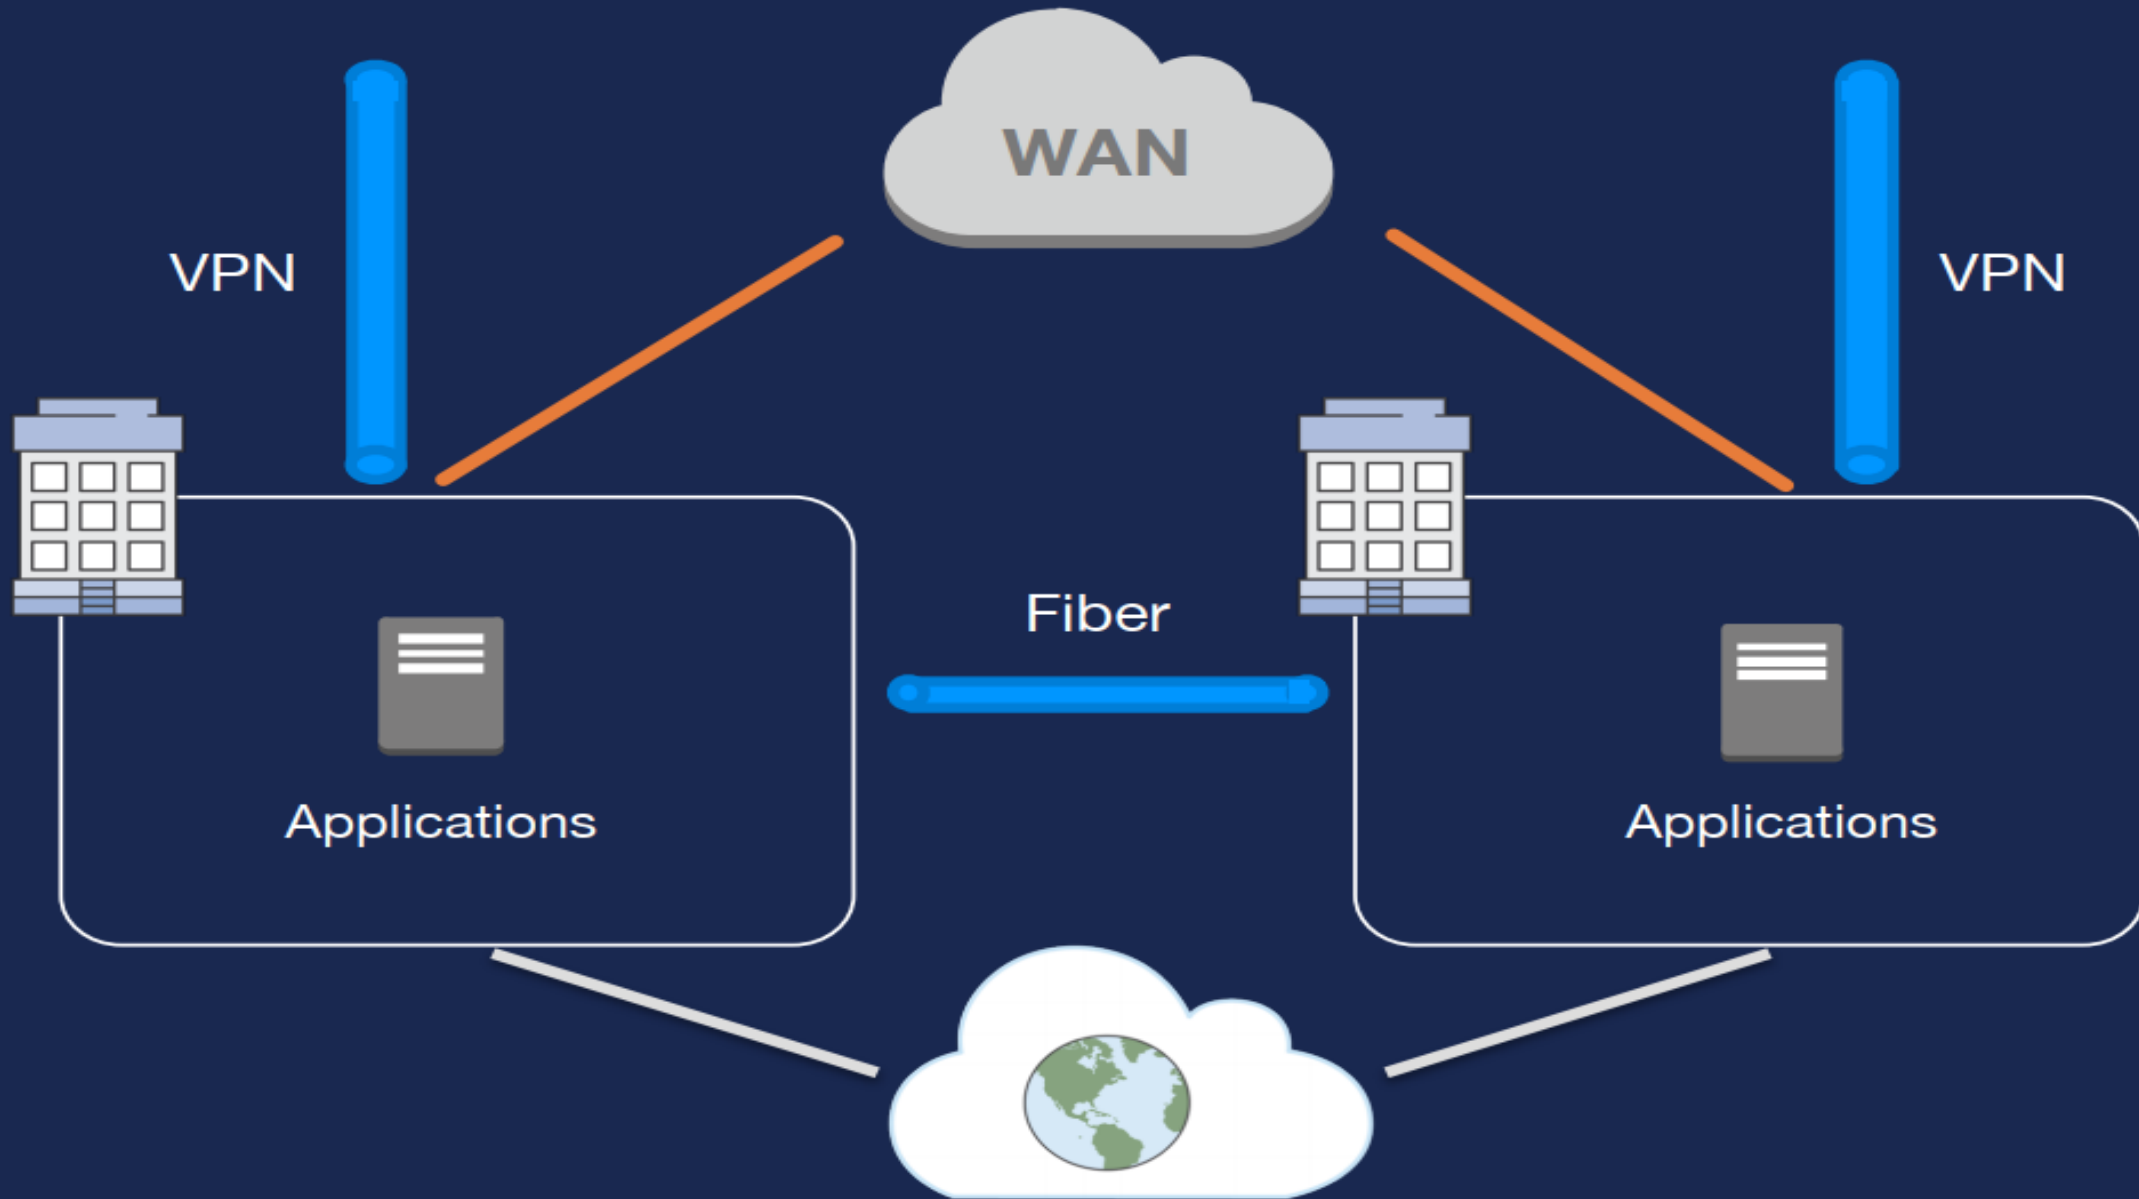


AWS Batch

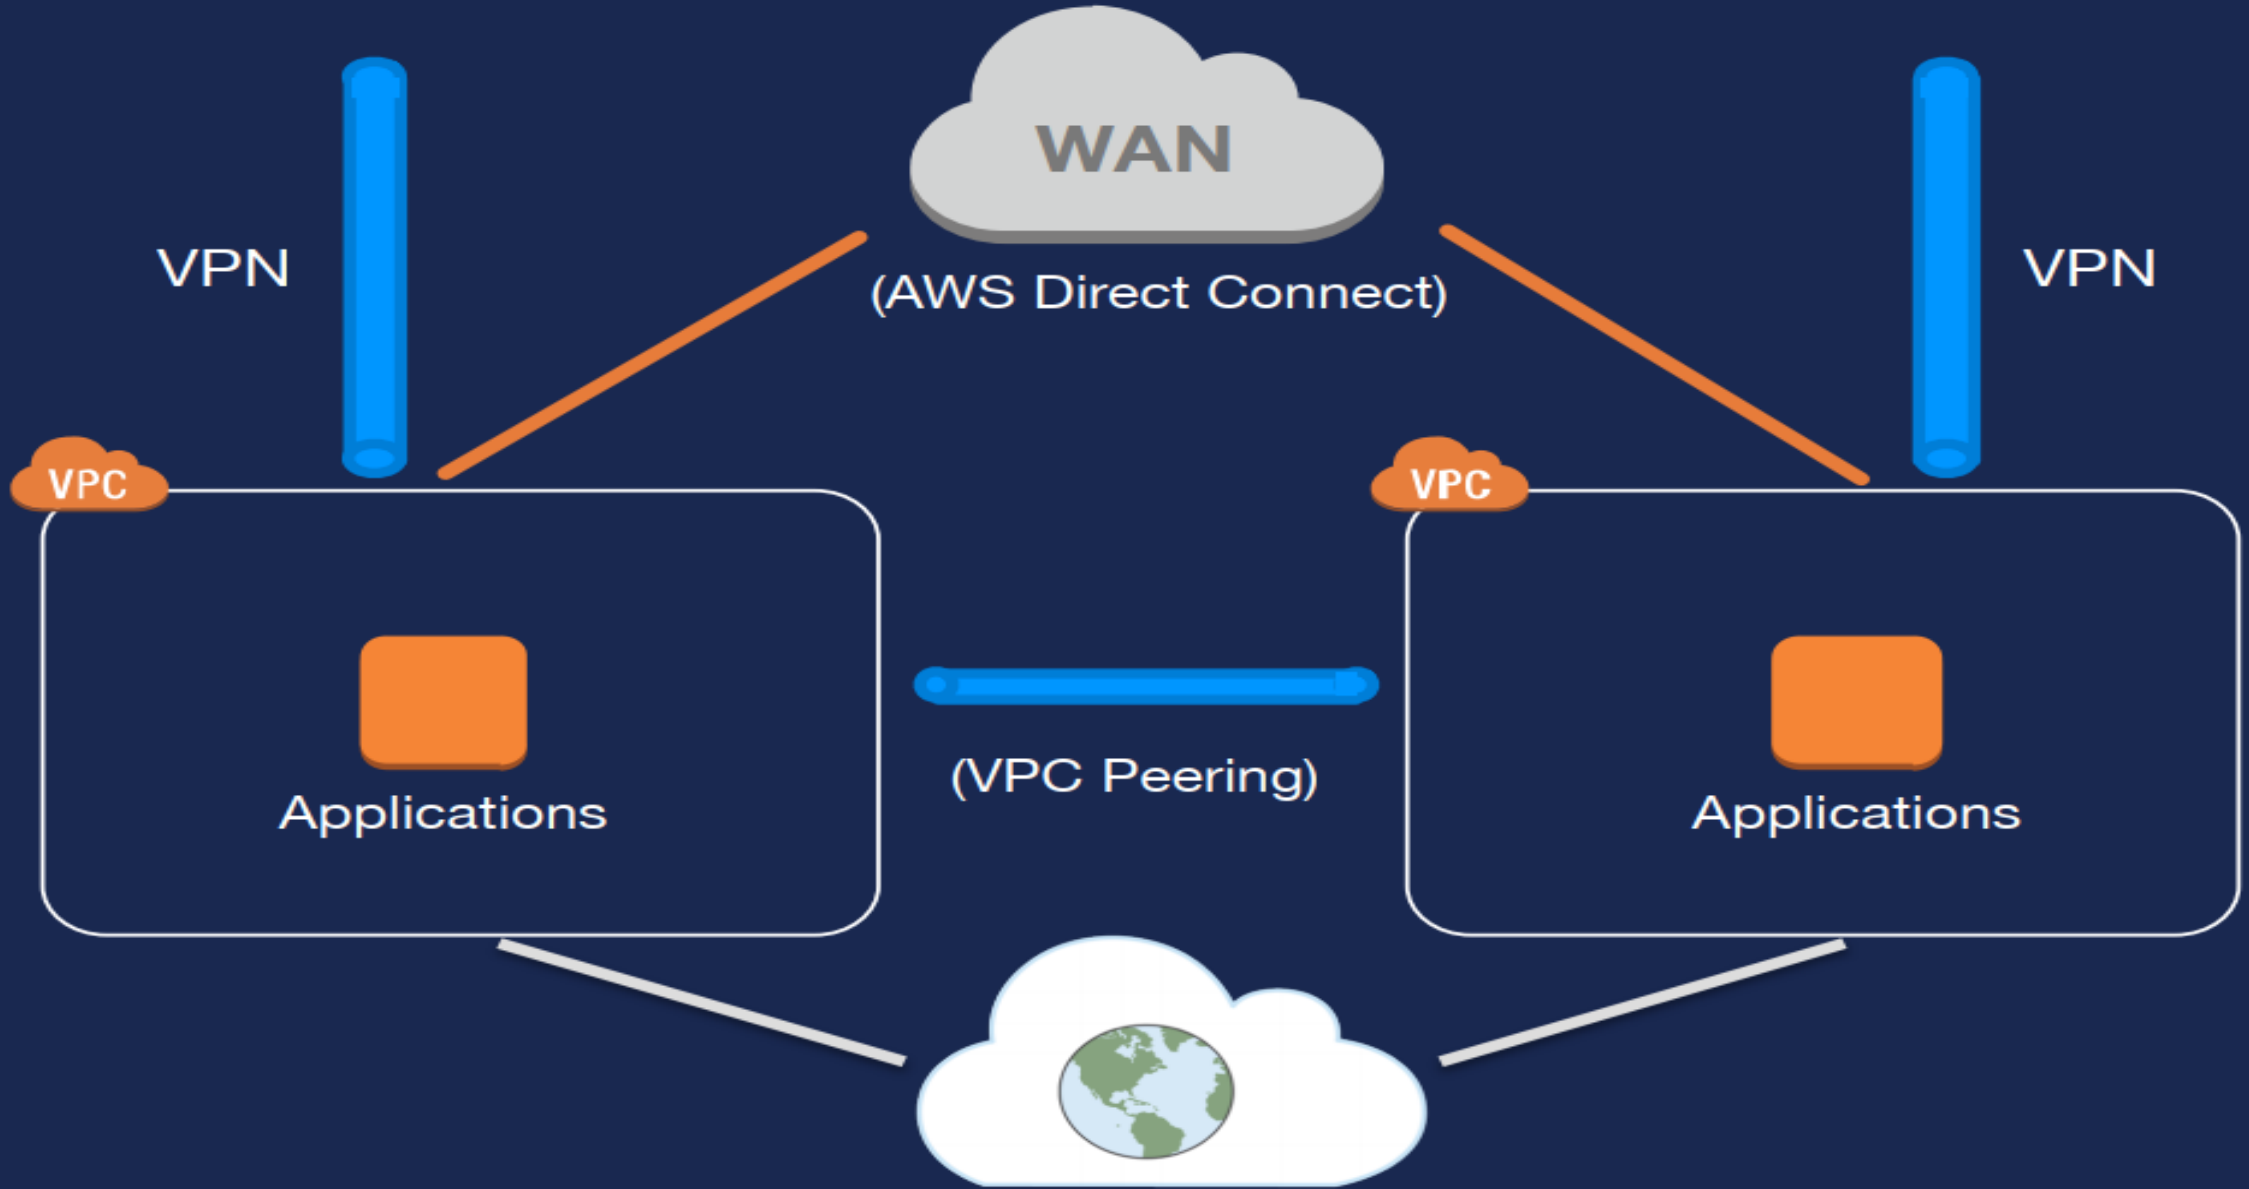
AWS Networking Overview



Traditional Network

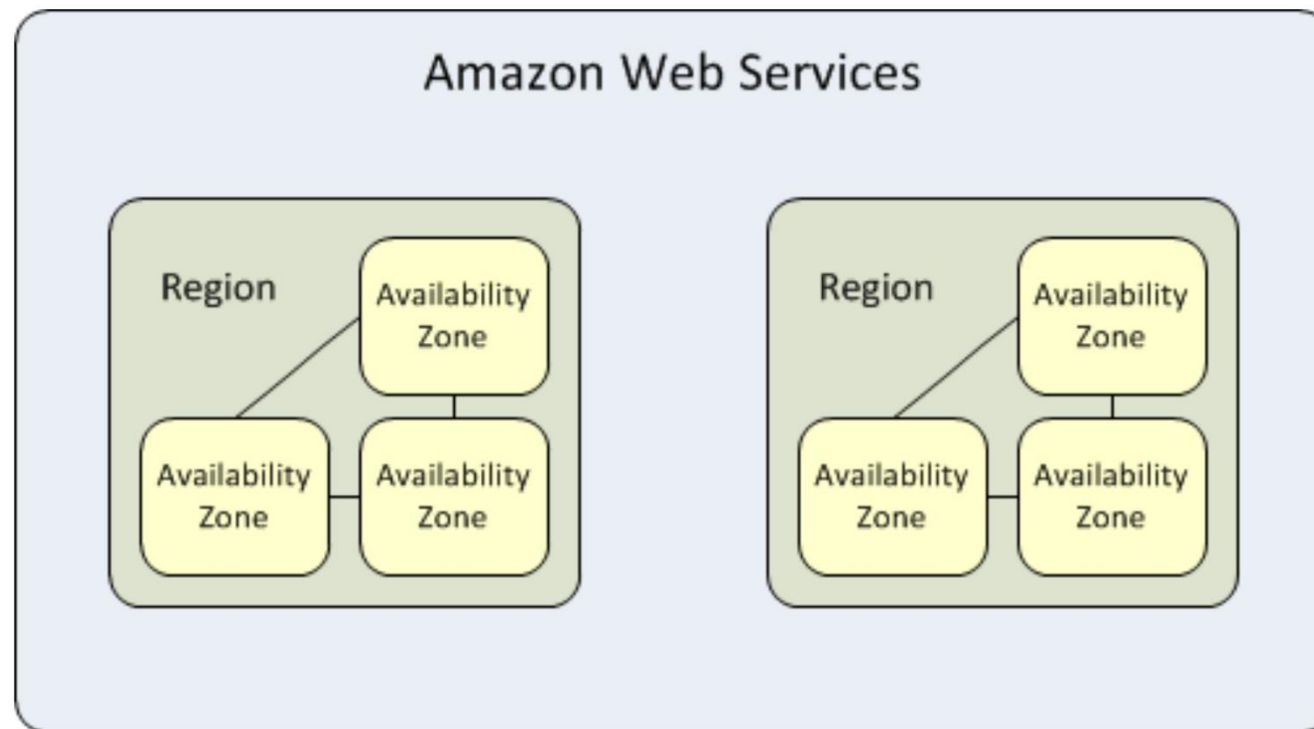


AWS Network



Region & Availability Zones

- Global Resources
 - IAM Users
 - Route 53 Records
- Regional Resources
 - S3 Buckets
 - VPCs
 - ELB (Elastic Load Balancing)
 - EIPs (Elastic IP Addresses)
- AZ Resources
 - EBS Volumes
 - EC2 Instances
 - RDS Instances
 - Subnets
 - ...

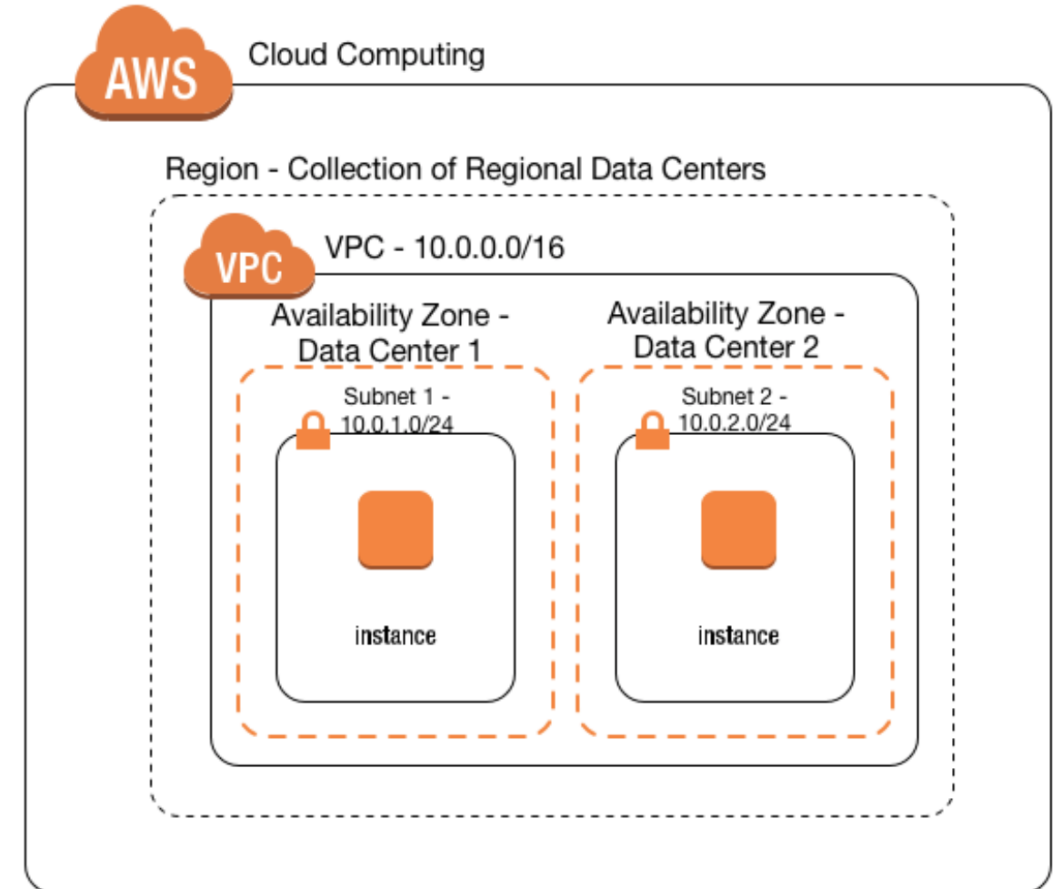


AWS Networking Components

“Your Virtual Datacenter in the Cloud”

■ Essential Components:

- Subnets
- Route Tables
- Network ACLs
- Security Groups
- Internet Gateways
- NAT Gateways
- Virtual Private Gateways



Amazon VPC – Virtual Private Cloud

- Provision a logically isolated section of the AWS Cloud where you can launch AWS resources in a virtual network that you define.

Bring your own network



IP Addresses



Subnets



Network Topology



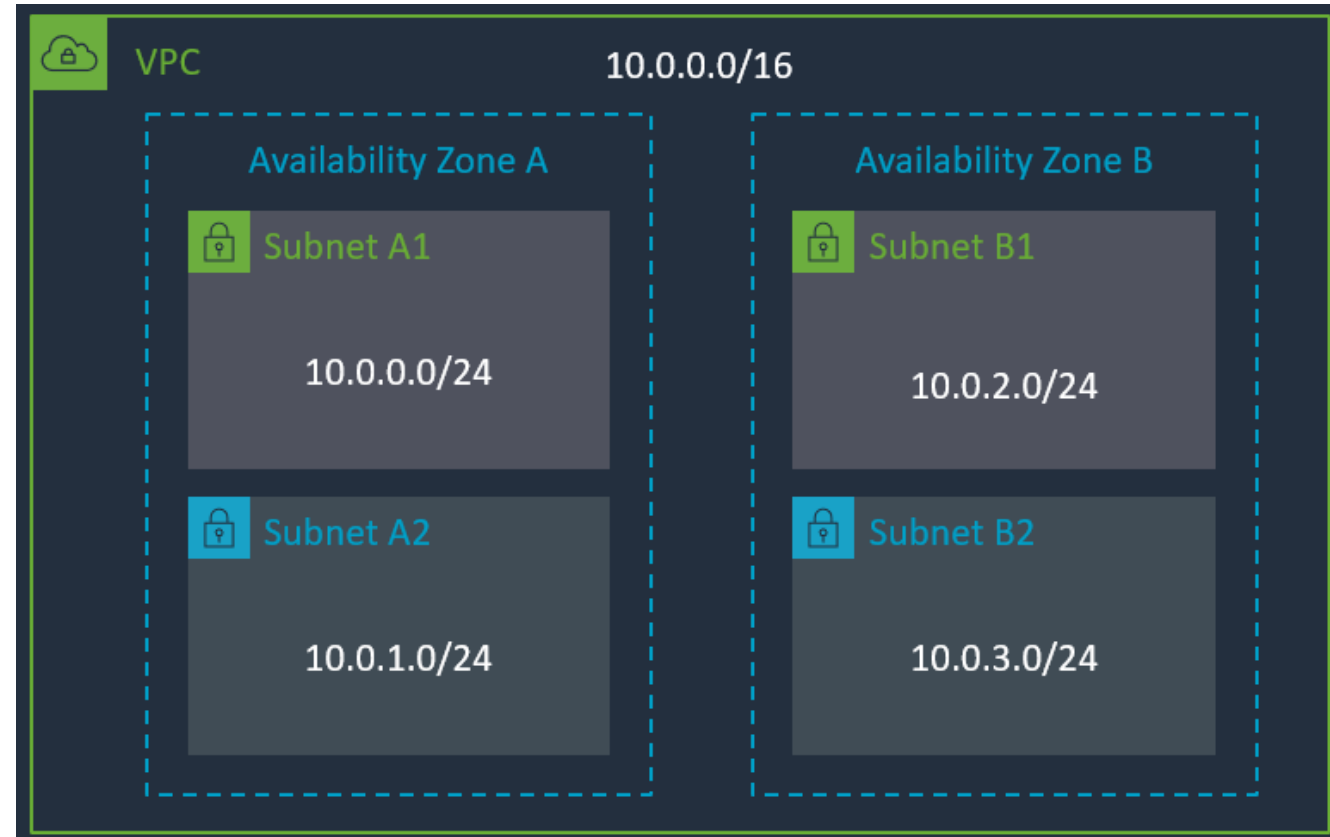
Routing Rules



Security Rules

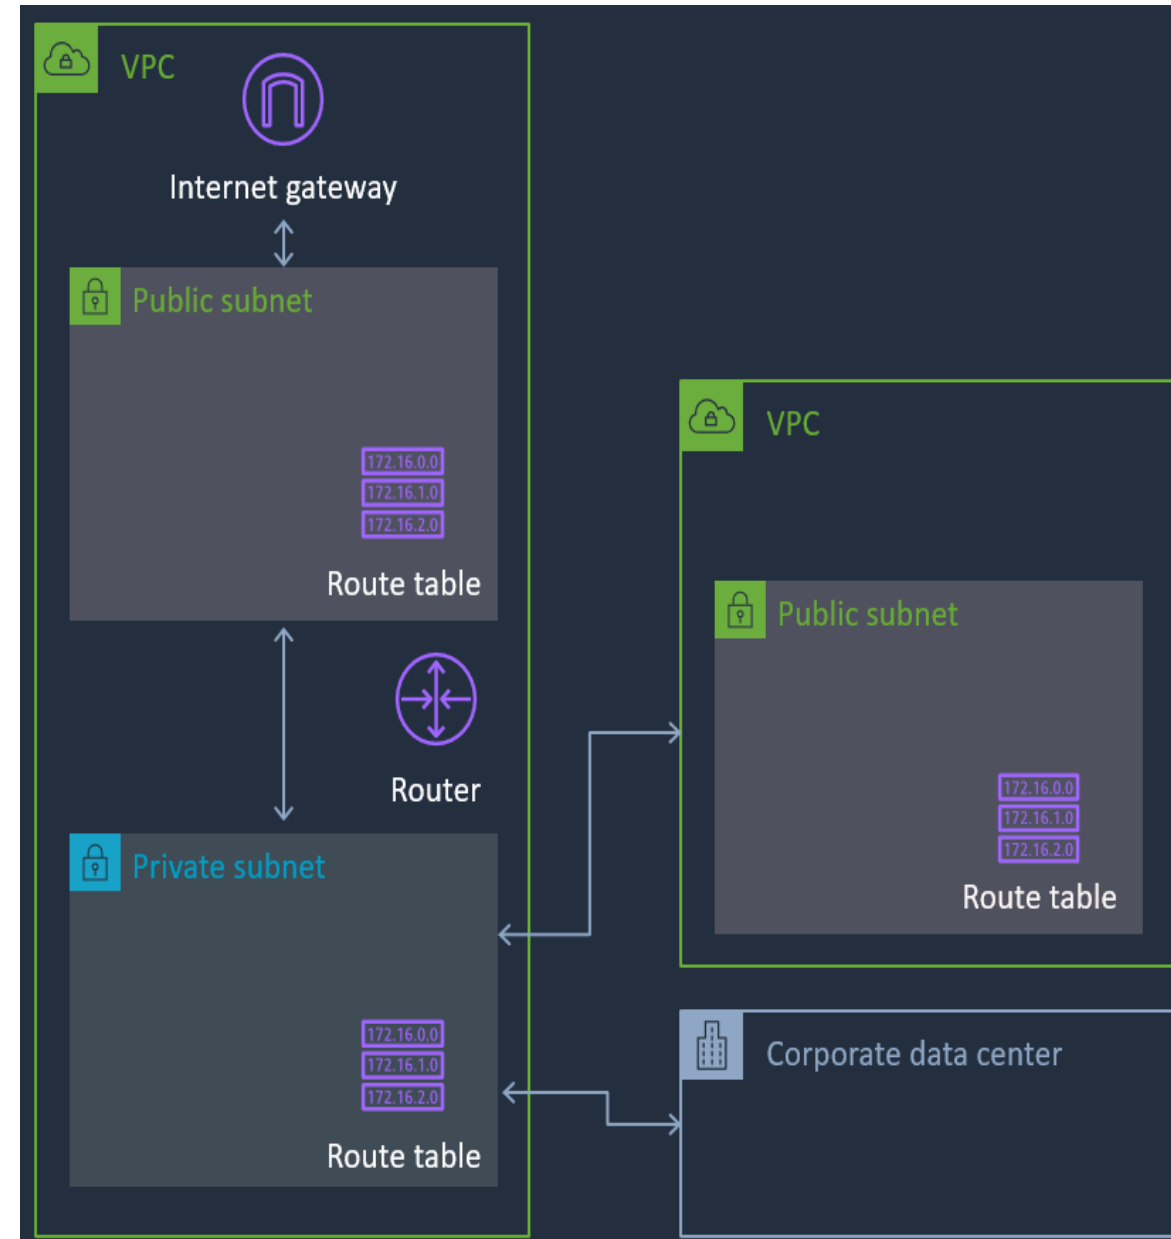
Subnet

- You can add one or more subnets in each Availability Zone
- AZs provides fault isolations
- Subnets are allocated as a subset of the VPC CIDR range



Subnet and Route Table

- Each subnet can have a unique Route Table
- Route Tables direct traffic out of the VPC, towards:
 - Internet Gateway
 - Virtual Private Gateway
 - VPC Endpoints
 - Direct Connect
 - VPC Peering
 - AWS Transit Gateway
- Subnets are named “Public Subnets” when connected to an Internet Gateway



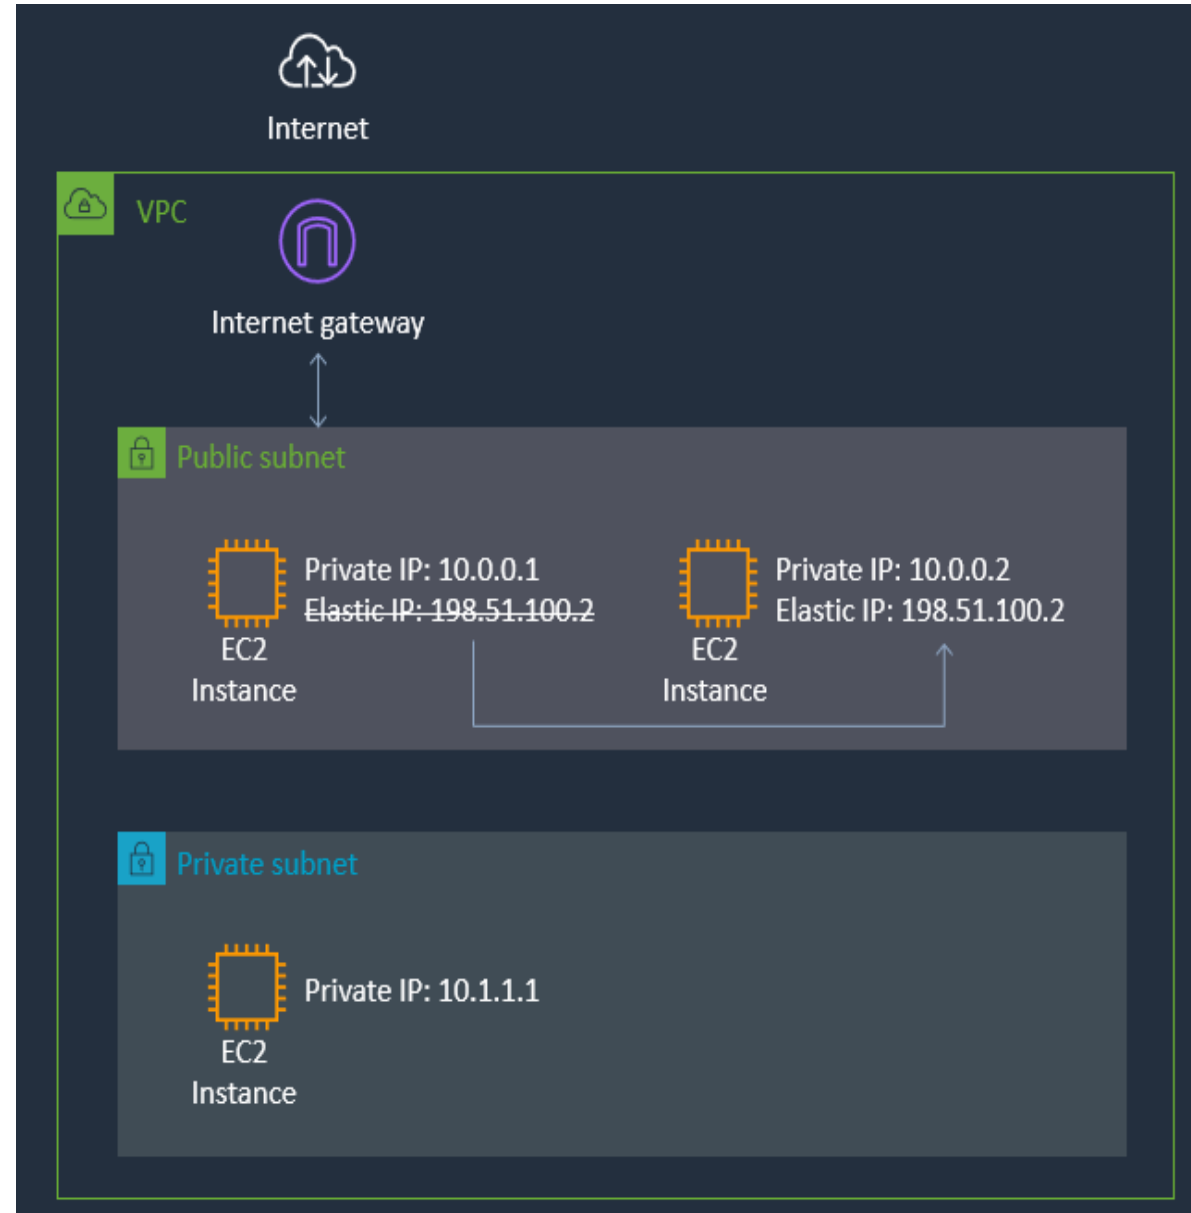
Internet Gateway

- Horizontally scaled, redundant, highly available VPC component
- Connect your VPC Subnets to the Internet
- Must be referenced on the Route Table
- Performs NAT between Public and Private IP Addresses



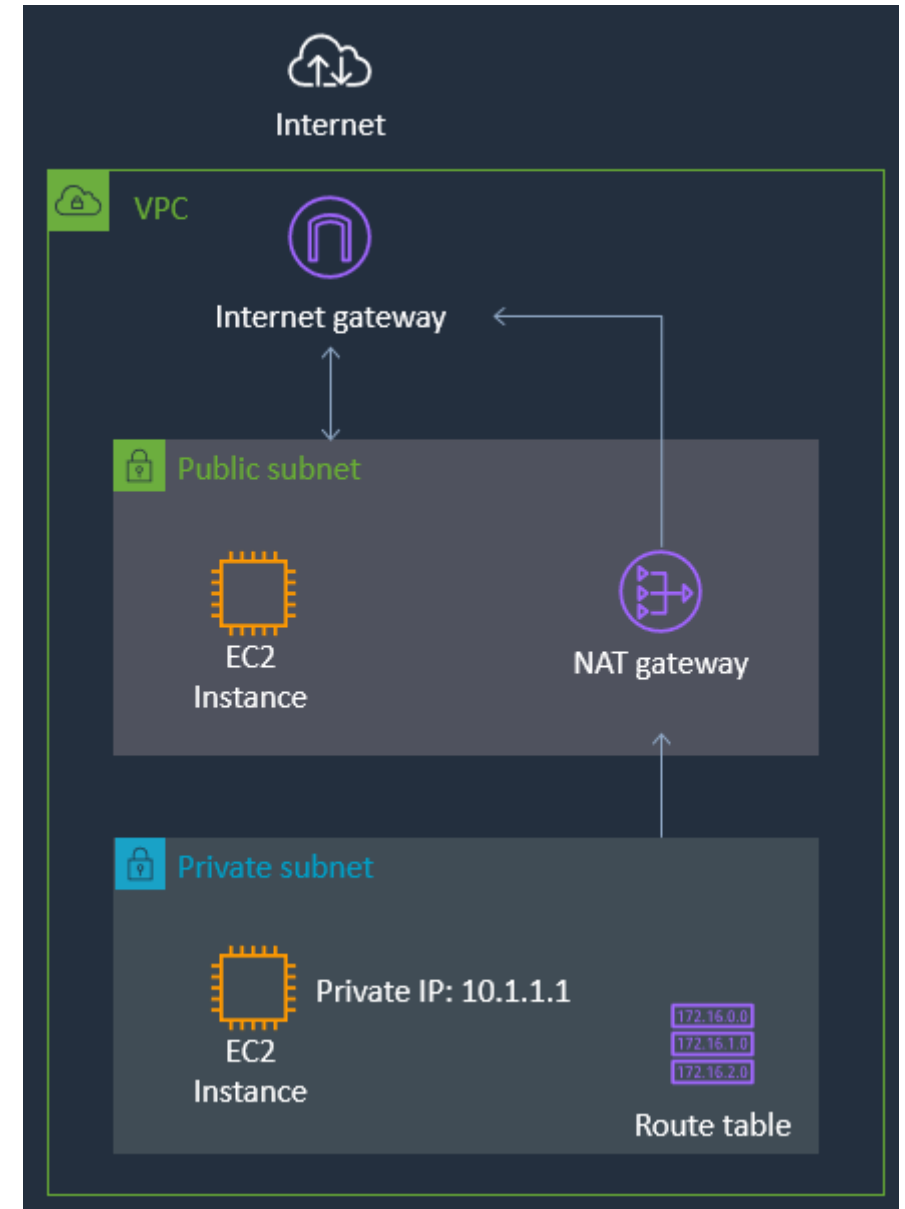
Elastic IP Address

- Static, Public IPv4 address, associated with your AWS account
- Can be associated with an instance or network interface
- Can be remapped to another instance in your account
- Useful for redundancy when Load Balancers are not an option



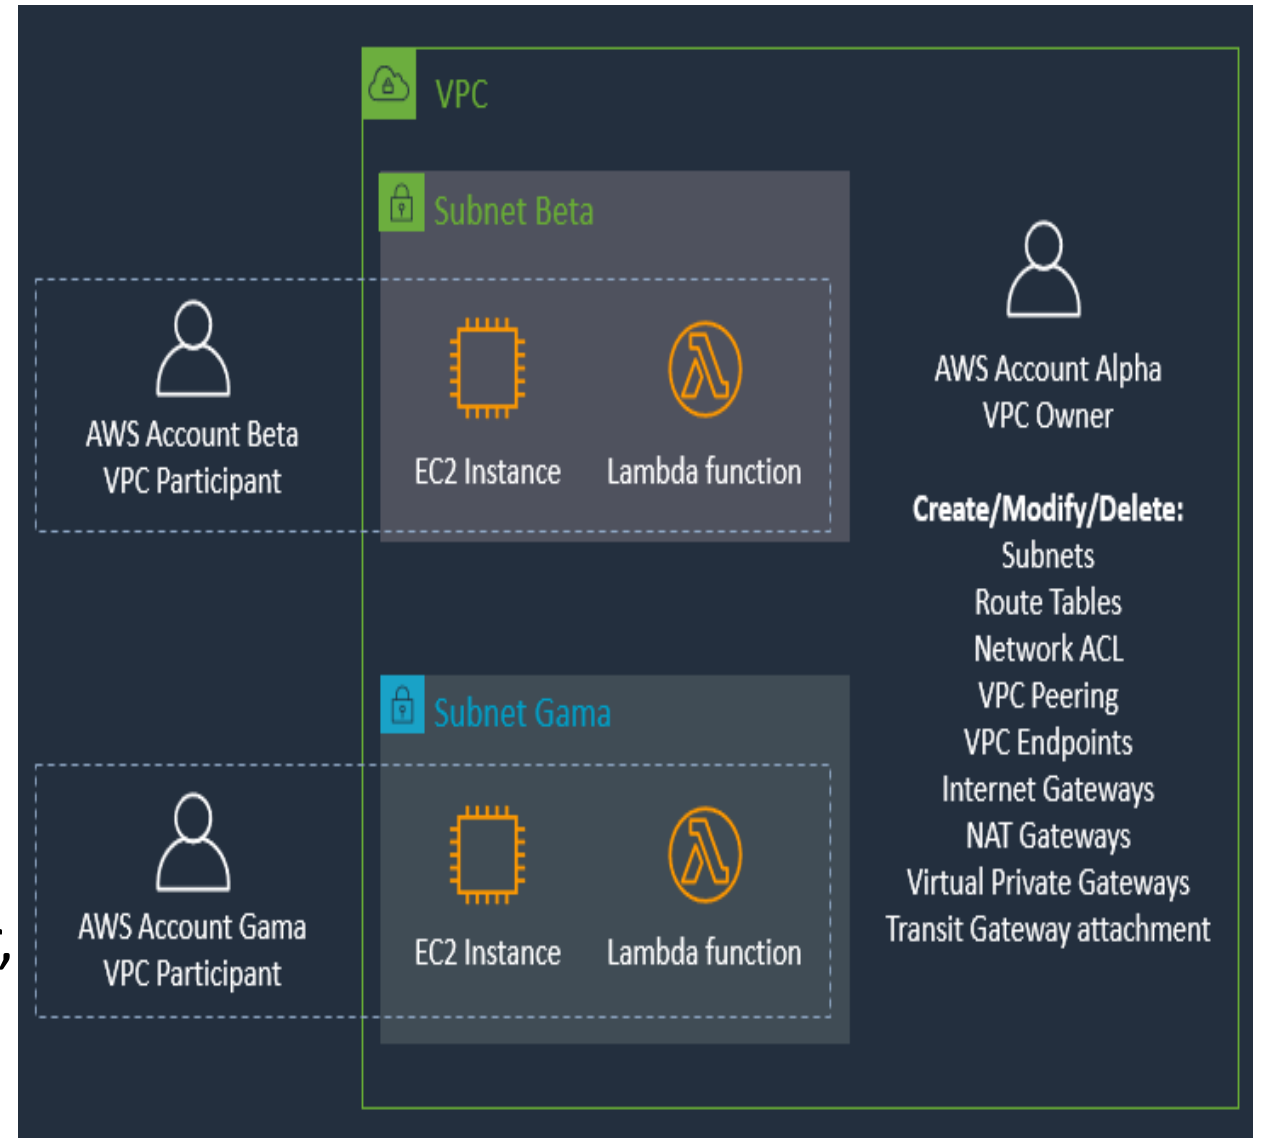
NAT Gateway

- Enable outbound connection to the internet
- No incoming connection - useful for OS/packages updates, public web services access
- Fully managed by AWS
- Highly available
- Up to 10Gbps bandwidth
- Supports TCP, UDP, and ICMP protocols
- Network ACLs apply to NAT gateway's traffic



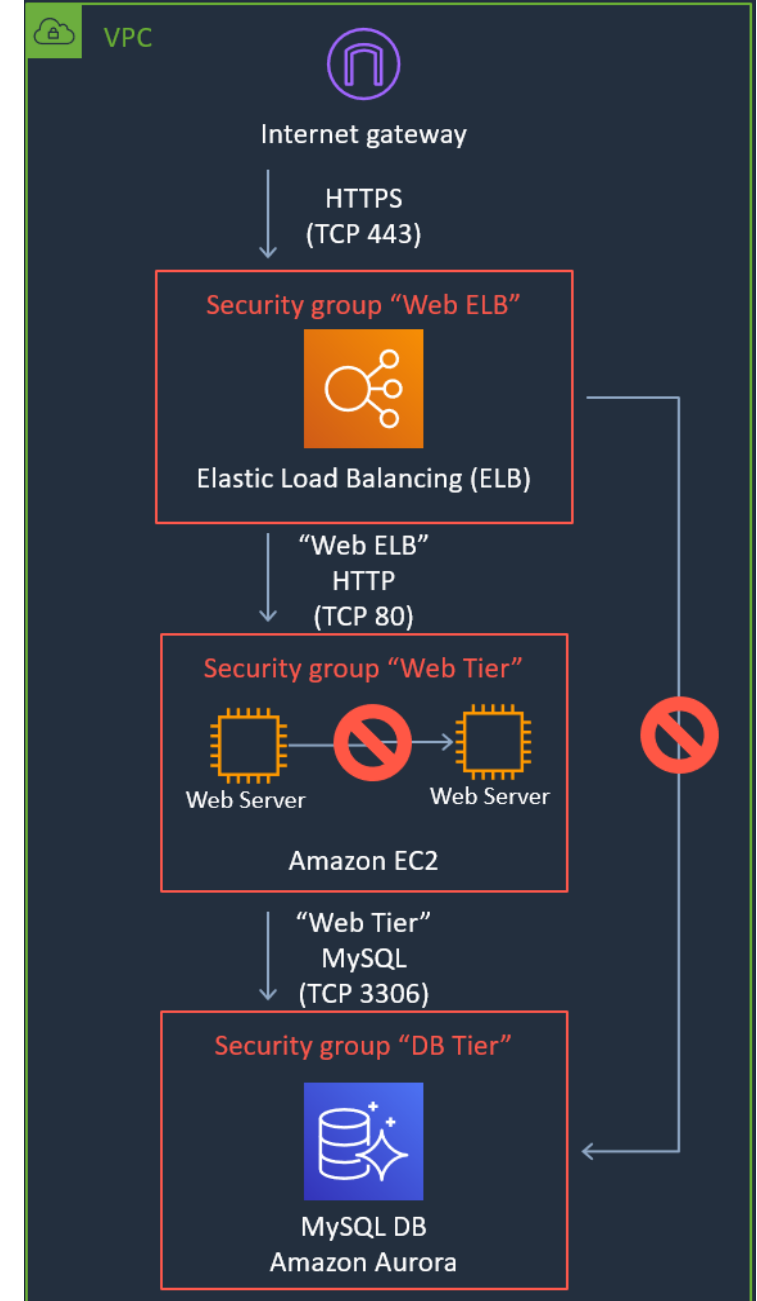
Shared VPC

- VPC Owner can create and edit VPC Components
- VPC Participants can launch resources in their assigned Subnets
- Each participant pays for their own resources and data transfer costs
- Based on AWS Resource Access Manager, under AWS Organizations



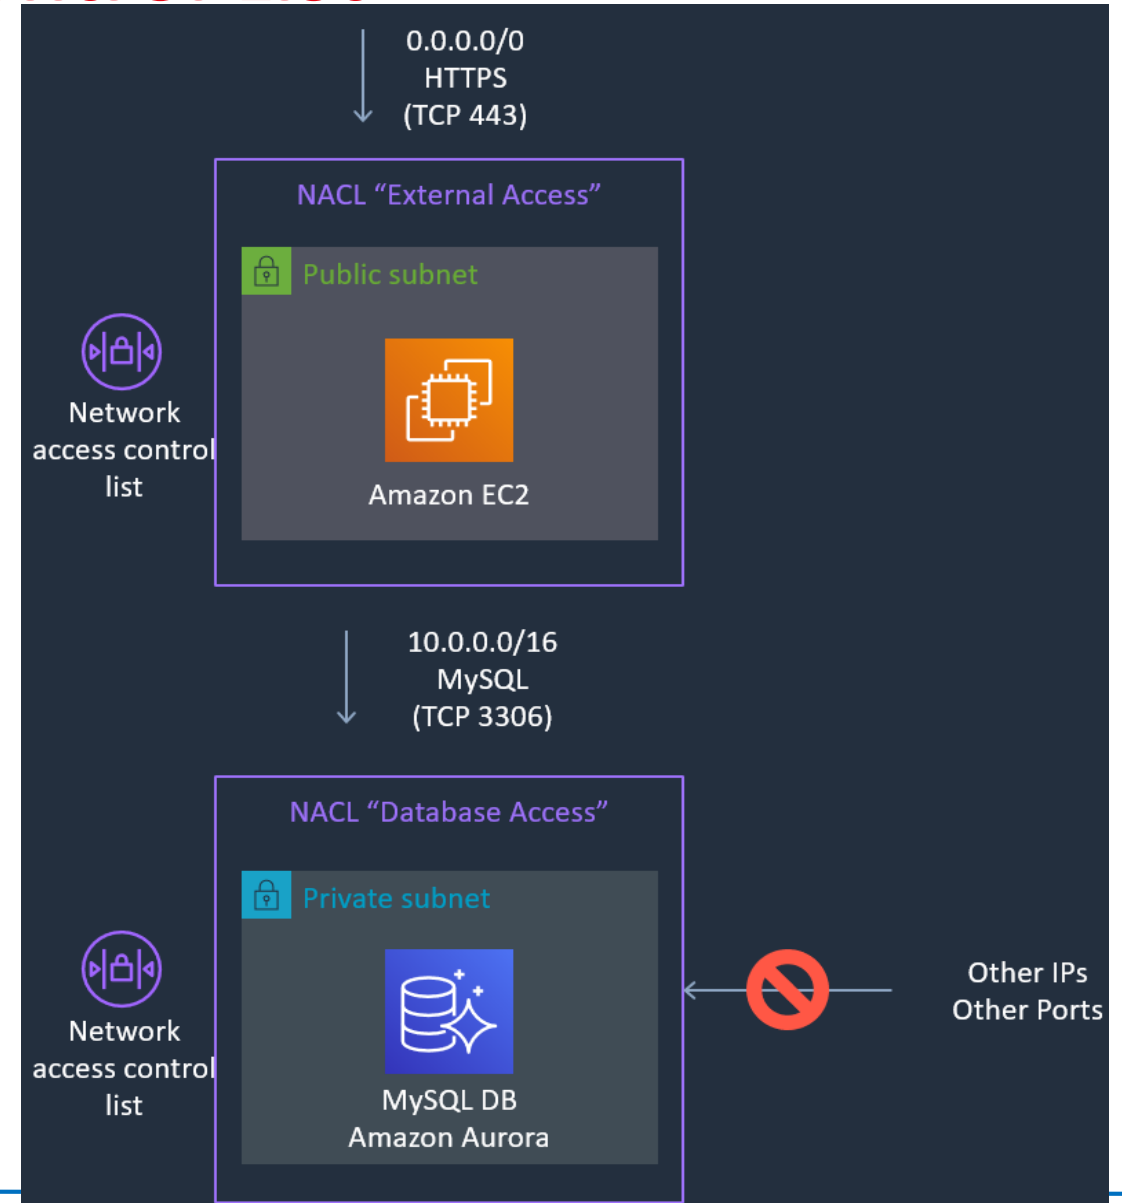
VPC Security - Security Groups

- Virtual stateful firewall
- Inbound and Outbound customer defined rules
- Instance/Interface level inspection
 - Micro segmentation
 - Mandatory, all instances have an associated Security Group
- Can be cross referenced
 - Works across VPC Peering
- Only supports allow rules
 - Implicit deny all at the end



VPC Security - Network Access Control List

- Inbound and Outbound
- Subnet level inspection
- Optional level of security
- By default, allow all traffic
- Stateless
- IP and TCP/UDP port based
- Supports allow and deny rules
- Deny all at the end



VPC Connectivity Options

- VPC Endpoints
- VPC Peering
- AWS Transit Gateway

Connect Data Center to AWS

- AWS Virtual Private Network
- AWS Direct Connect
- AWS Direct Connect Gateway
- AWS DX Gateway + AWS Transit Gateway
- AWS Client VPN

AWS Elastic Load Balancing



ELB - Elastic Load Balancer

■ Public Side


- Consists of an endpoint which is the equivalent to a traditional VIP
- Does not use a static IPv4, but rather an Alias/CNAME
- The endpoint will not always resolve to the same IP

■ Private Side


- Minimum of one virtual ELB node per AZ

■ Certificate Termination

- Only one SSL certificate per ELB
- Multi-Domain certificates are valid




ELB – Spans Multiple Availability Zones



Auto Scaling - Overview


■ Auto Scaling Key Features

- Adds or removes servers based on load
- Self-healing pool of resources
- Every instance is based on a “gold” master image



Auto Scaling - Components

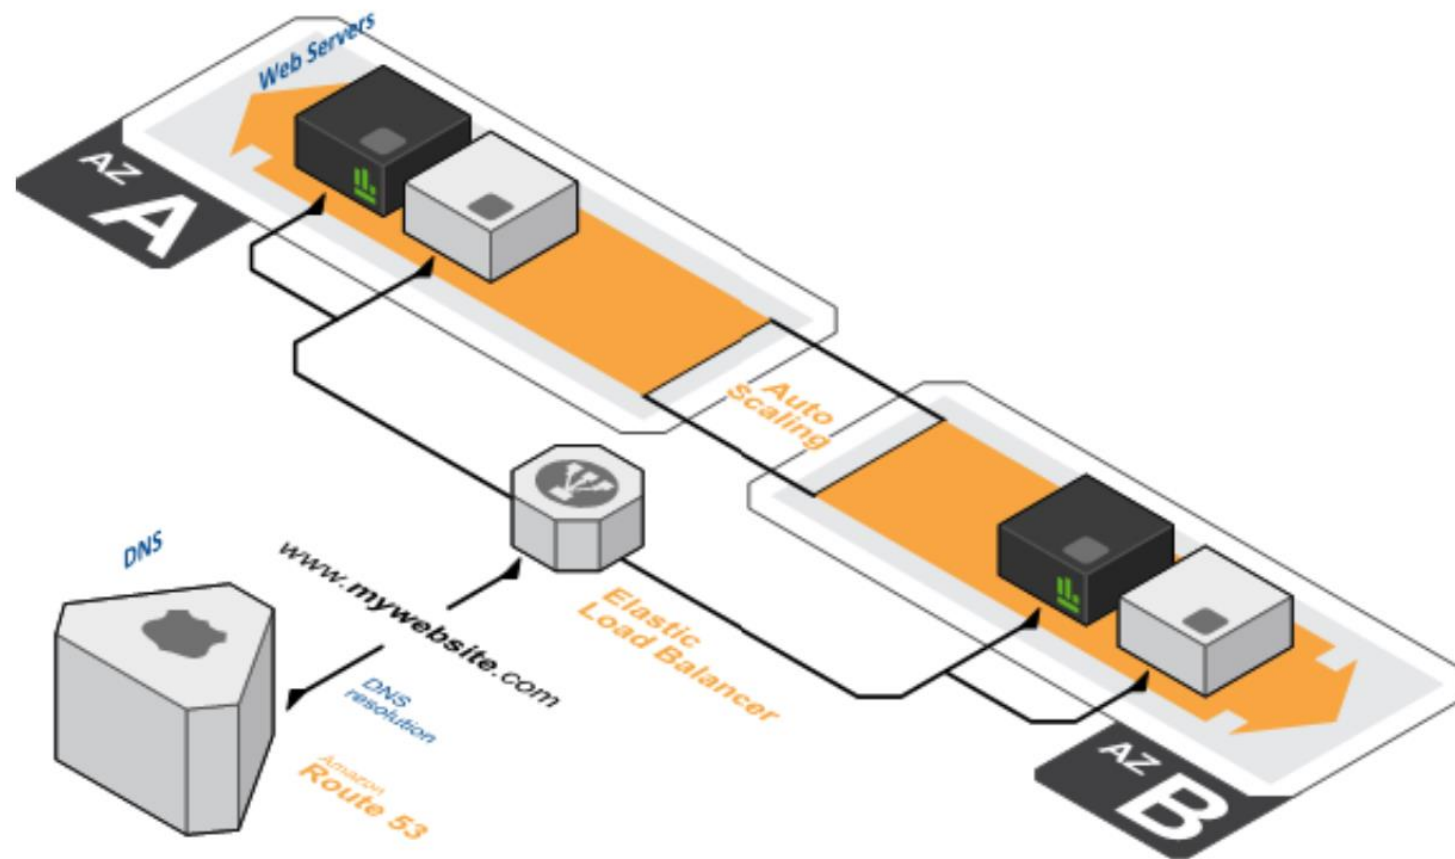
- Auto scaling group
 - Instance location
 - Subnet
 - Load Balancer
 - Number of instances
 - Min
 - Max
 - Desired
- Launch config
 - Instance details
 - ◇ Size
 - ◇ PEM key
 - ◇ IAM Profile
 - ◇ Security Group(s)
 - ◇ User data



Auto Scaling - Multi-AZ

■ Multi-AZ Auto Scaling

- Highly Available
- Production Standard
- Spans Datacenters



Auto Scaling – Cloud Watch

- CloudWatch is the final piece of the auto scaling puzzle. You can create alarms based on instance metrics which trigger auto scaling actions.
- Scaling policies
 - Scale up alarm
 - Execute policy when: CPU is greater than 60%
 - Take the action: Add 2 instances
 - And then wait: 10 minutes
 - Scale down alarm
 - Execute policy when: CPU is less than 20%
 - Take the action: Remove 2 instances
 - And then wait: 10 minutes

Q & A



VIỆN ĐIỆN TỬ - VIỄN THÔNG
School of Electronics and Telecommunications



Module 2: Maintenance - Backup and Restore



Agenda

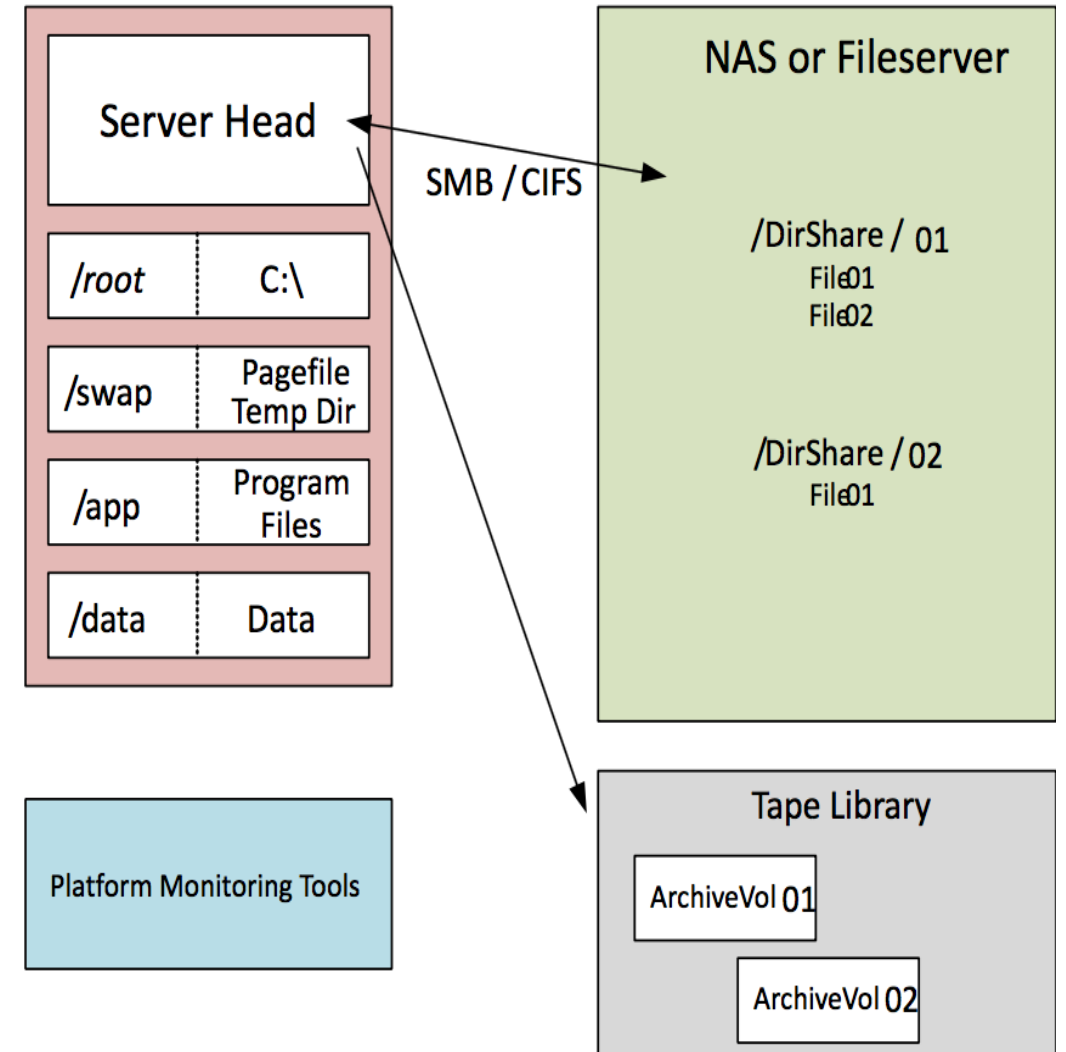
- AWS Storage Services
- AWS Database Services
- AWS Elastic Load Balancing

AWS Storage Services



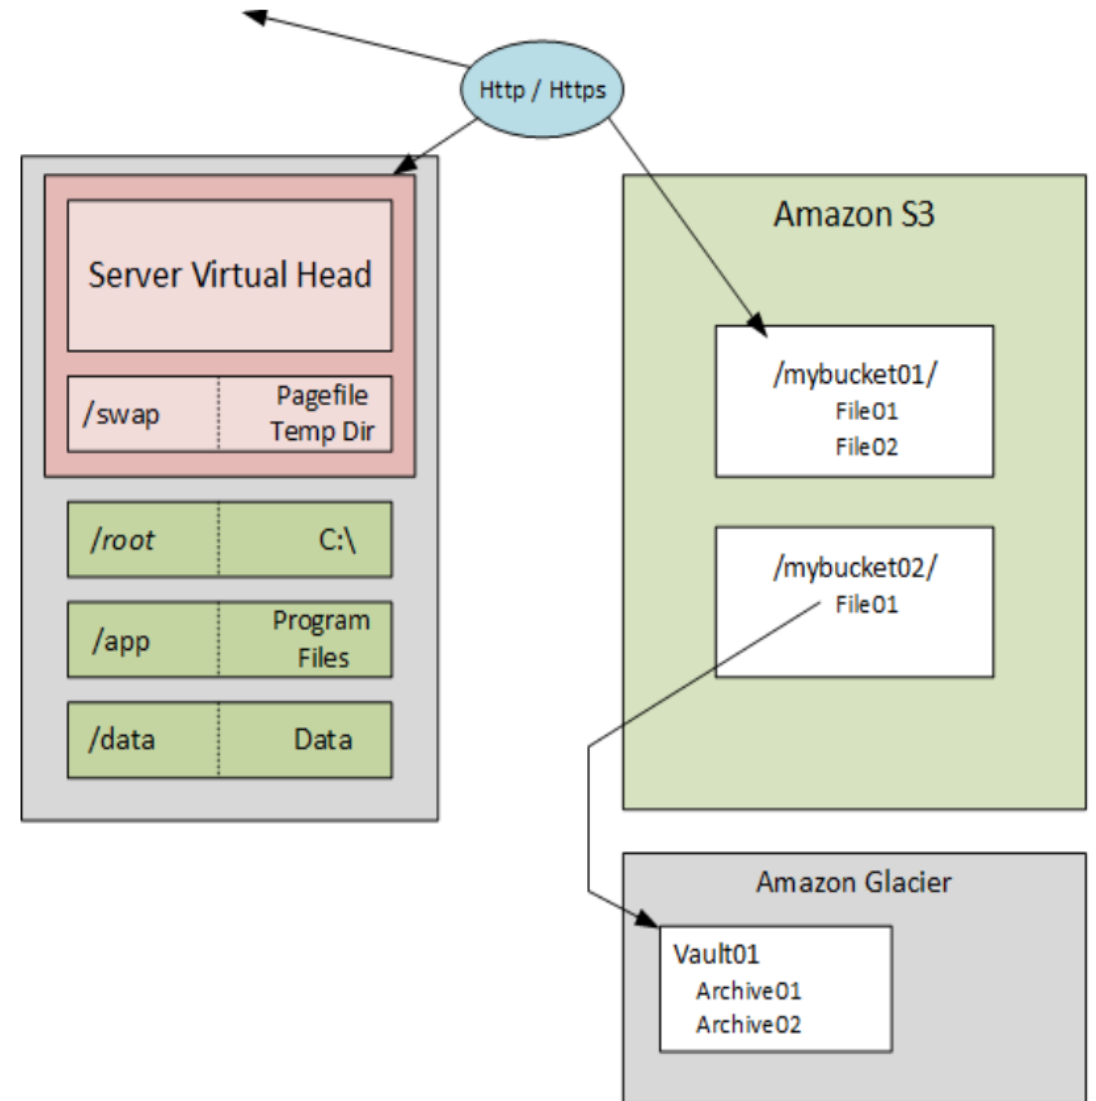
Traditional Platform - Storage Architecture

- In the old days...
 - Hardware acquisition and datacenter space required advanced planning
 - Disk space and I/O allocation juggling for the entire application lifecycle
 - Volume and file redundancy not built-in
 - Capital commitment and refresh budget considerations



AWS Instance Volumes and Data Storage

- The new [improved] way of doing things...
 - Elastic pay-as-you-go model
 - Redundancy and snapshot utilities built-in
 - New APIs and tools simplify application development, administration and data lifecycle management



AWS Storage Services & Content Delivery



S3



Glacier



Elastic
Block
Storage



Storage
GateWay



CloudFront

Amazon Simple Storage Service (S3)



S3



Glacier



Elastic
Block
Storage



Storage
GateWay

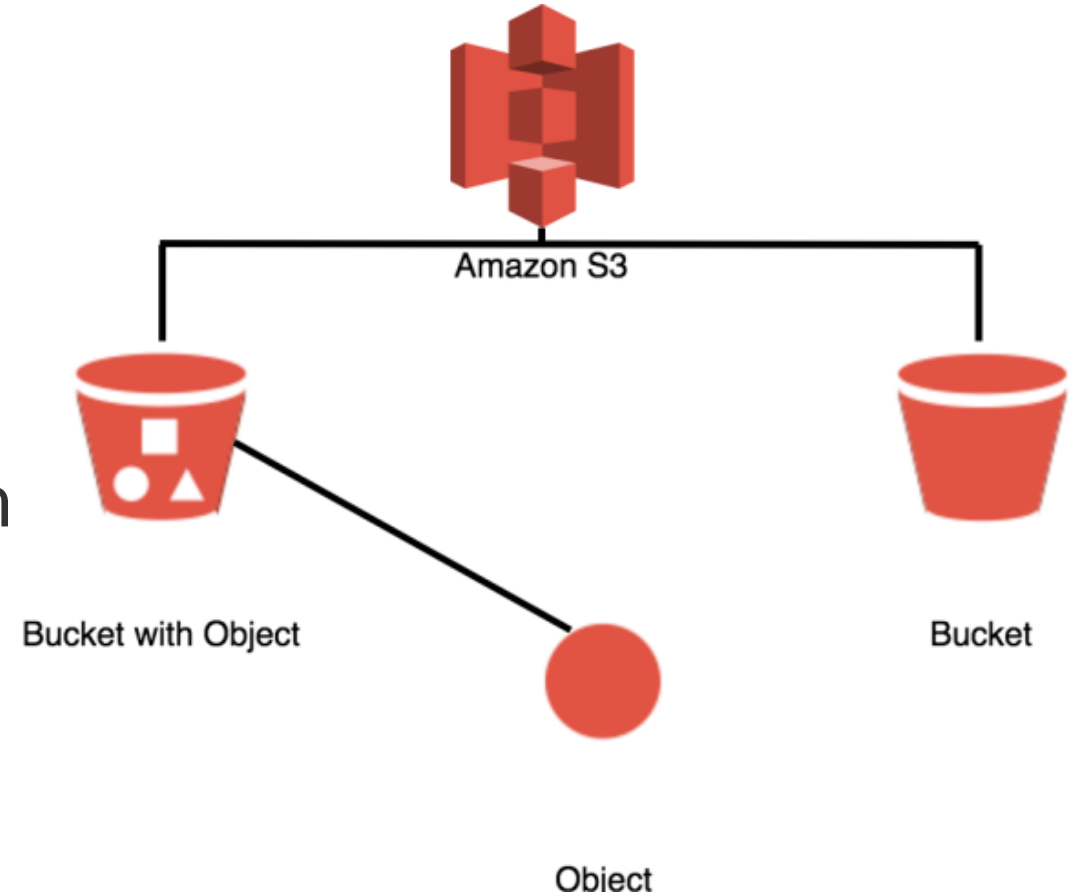


CloudFront

- Storage for the Internet
- Natively online, HTTP access
- Store and retrieve any amount of data, any time, from anywhere on the web
- High Scalable, reliable, fast and durable

Amazon S3 Concept

- Amazon S3 stores data as objects within buckets
- An object is comprised of a file and optionally any metadata that describes that file
- You can have up to 100 buckets in each account
- You can control access to the bucket and its objects



Amazon Glacier



S3



Glacier



Elastic
Block
Storage



Storage
GateWay



CloudFront

- Extremely low-cost storage
- Secure, durable storage for **data archiving and backup**
- Optimized for data that is **infrequently accessed**

Amazon Elastic Block Store (EBS)



S3



Glacier



Elastic
Block
Storage





Storage
GateWay



CloudFront

- **Persistent block level storage** volumes offering consistent and low-latency performance
- Automatically replicated within its Availability Zone
- Snapshots stored durably in Amazon S3

Amazon EBS vs Amazon S3



	Amazon EBS	Amazon S3
Paradigm	Block storage with file system	Object store
Performance	Very fast	Fast
Redundancy	Across multiple servers in an Availability Zone	Across multiple facilities in a Region
Security	EBS Encryption – Data volumes and Snapshots	Encryption
Access from the Internet?	No (1)	Yes (2)
Typical use case	It is a disk drive	Online storage

(1) Accessible from the Internet if mounted to server and set up as FTP, etc.

(2) Only with proper credentials, unless ACLs are world-readable

Amazon Storage Gateway



S3



Glacier



Elastic
Block
Storage



Storage
GateWay



CloudFront

- Connect an On-premises software appliance with cloud-based storage
- Securely upload data to the AWS cloud for cost effective backup and rapid disaster recovery
- Mirror your on-premises data to Amazon EC2 instances

Amazon Cloud Front



S3




Glacier



Elastic
Block
Storage



Storage
GateWay



CloudFront

- Easy and cost effective way to **distribute content** to end users
- **Low latency, high data transfer speeds**
- Deliver your entire website, including dynamic, static, streaming, and interactive content using a global network of edge locations

AWS Database Services




AWS Database Services




Managed Relational Database Service

Amazon RDS




Petabyte-scale Data Warehouse

Amazon Redshift




In-Memory Key Value Store

Amazon ElastiCache




Fully Managed Key-Value and Document Database

Amazon DynamoDB



Fully Managed Graph Database

Amazon Neptune



Cloud-Native Relational Database

Amazon Aurora

[Preview]



Fully Managed Time Series Database

Amazon Timestream


[Preview]



Fully Managed Ledger Database

Amazon QLDB

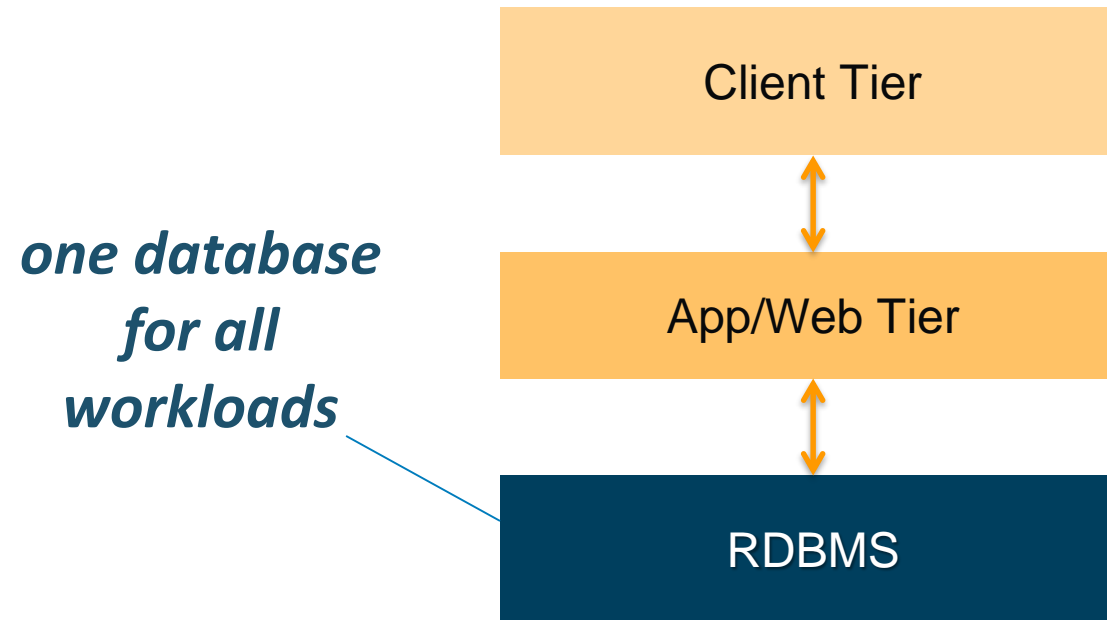
[Now GA]



MongoDB Compatible Document Database

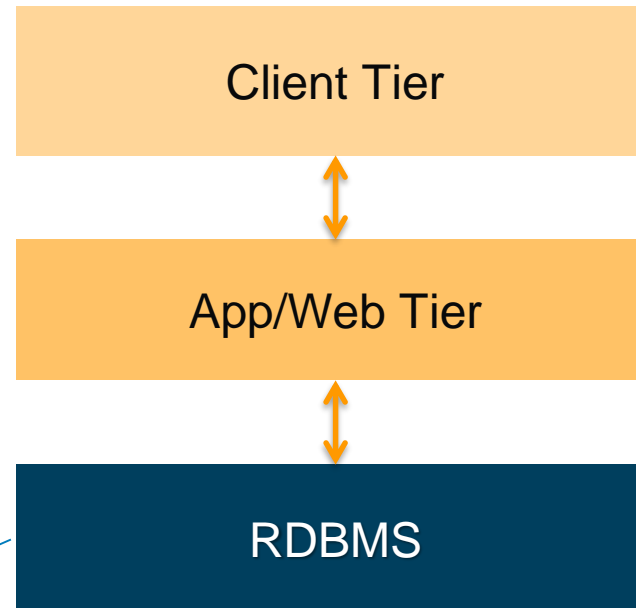
Amazon DocumentDB

Traditional Database Architecture



Traditional Database Architecture

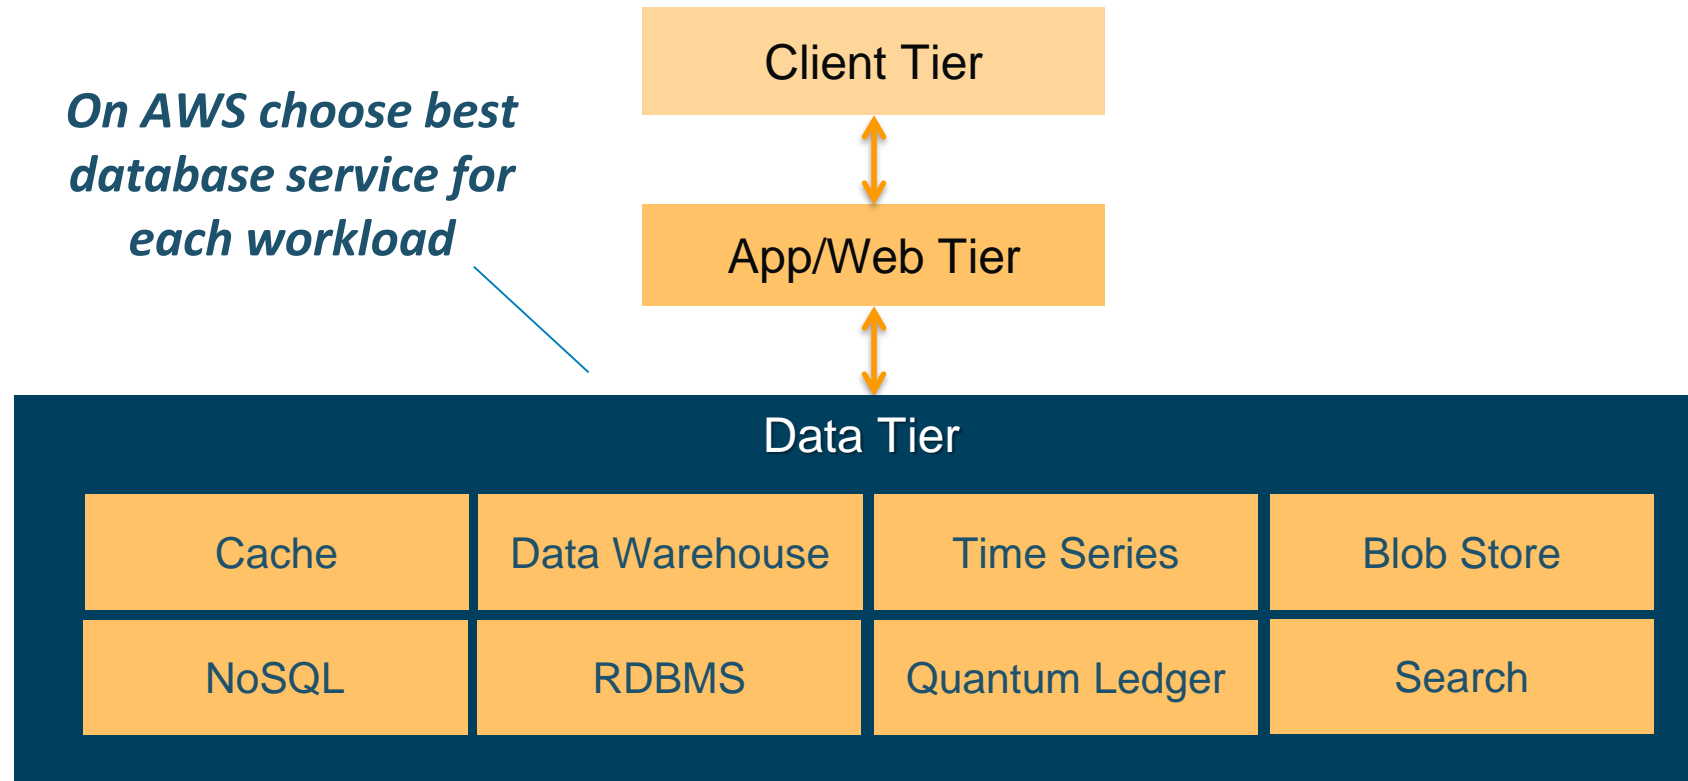
Key-value access
Complex queries
OLAP transactions
Analytics



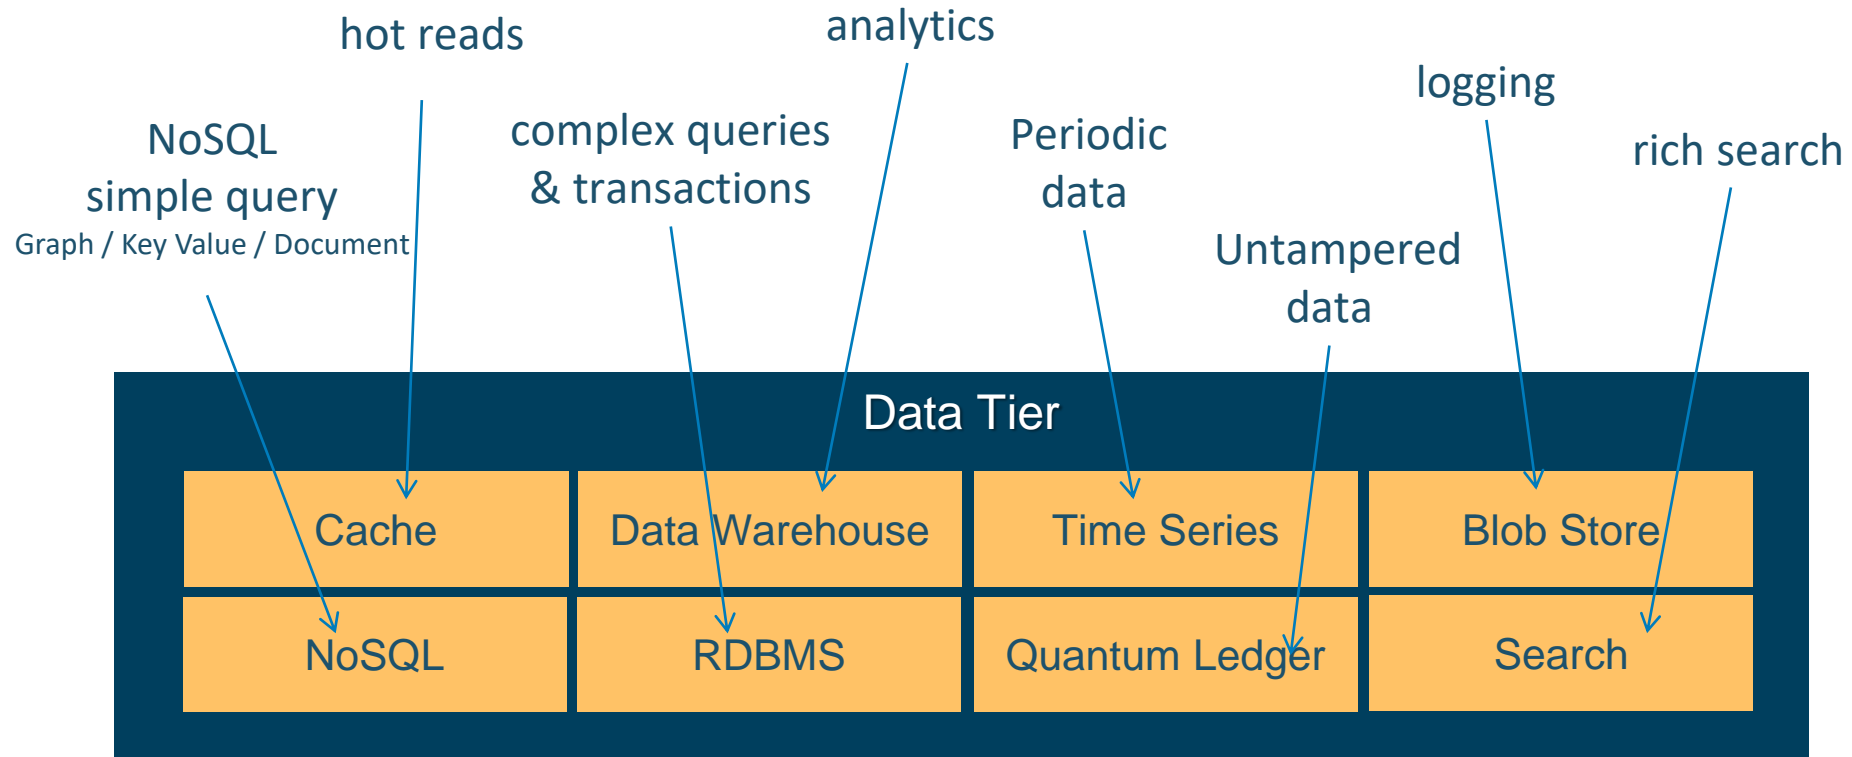
All forced into the relational database



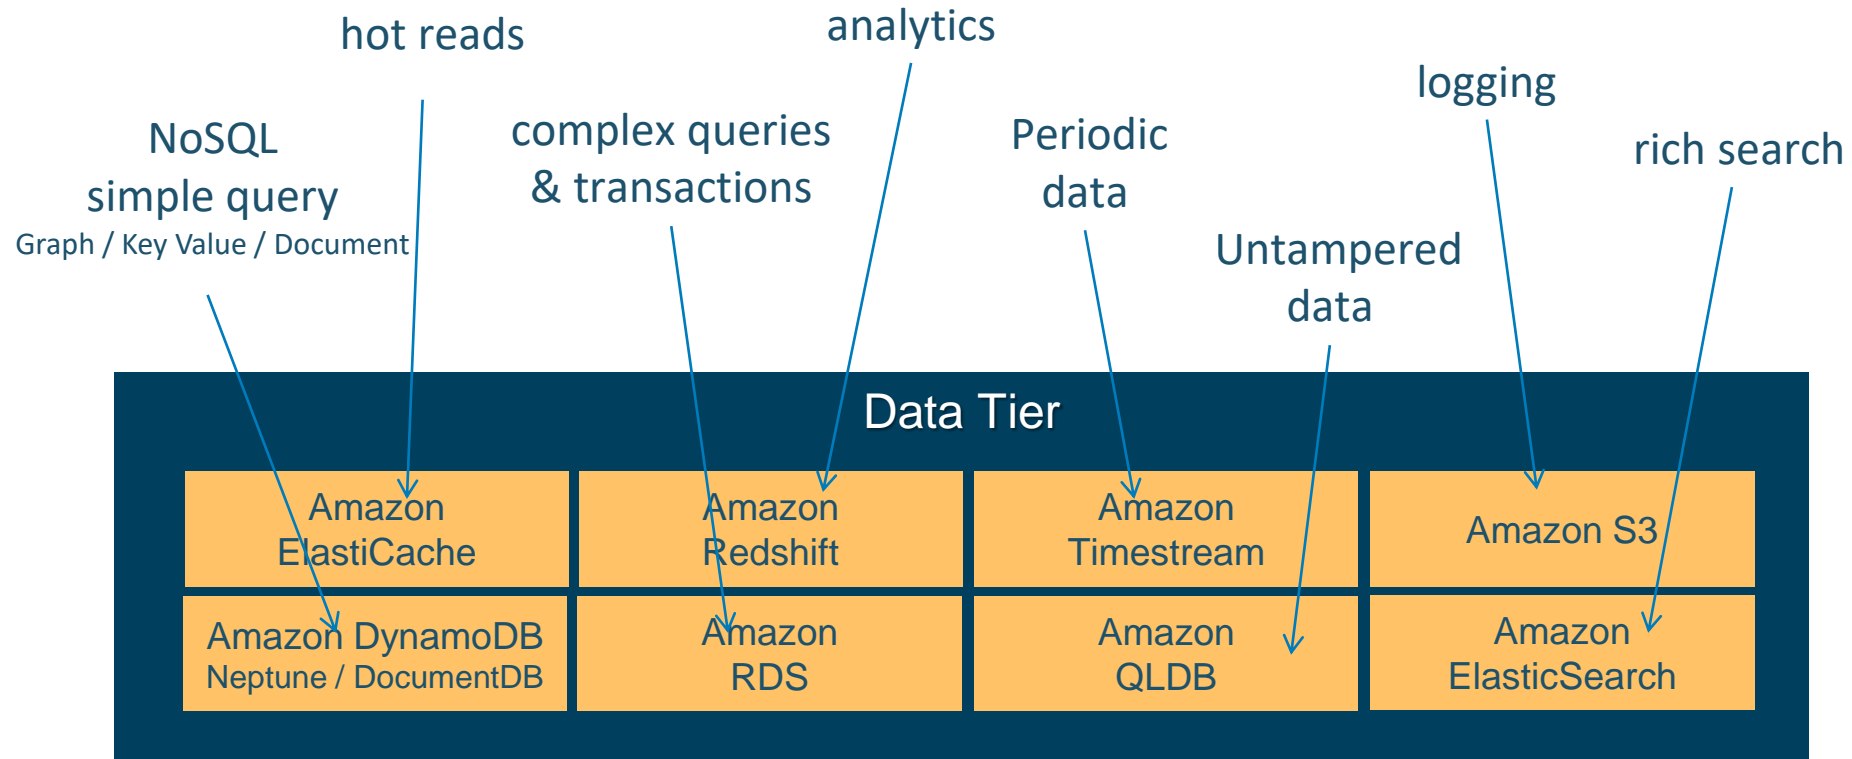
AWS Data Tier Architecture




Workload Driven Data Store Selection



AWS Database Services for the Data Tier





- Easy to setup, operation and scale a relational database
- **Automatically patches** the database software and backup your database
- Ability to **scale the compute resources** or storage capacity associated with your relational database instance via a single API call

Amazon RDS

Managed relational database service with a choice of popular database engines

Amazon
Aurora



Microsoft SQL Server

ORACLE



Easy to administer

Easily deploy and maintain hardware, OS and DB software; built-in monitoring



Performant & scalable

Scale compute and storage with a few clicks; minimal downtime for your application



Available & durable

Automatic Multi-AZ data replication; automated backup, snapshots, and failover



Secure and compliant

Data encryption at rest and in transit; industry compliance and assurance programs



If you host your databases on-premises...

- App optimization
- Scaling
- High availability
- Database backups
- DB s/w patches
- DB s/w installs
- OS patches
- OS installation
- Server maintenance
- Rack & stack
- Power, HVAC, net



you



If you host your databases in Amazon EC2...

- App optimization
- Scaling
- High availability
- Database backups
- DB s/w patches
- DB s/w installs
- OS patches
- OS installation
- Server maintenance
- Rack & stack
- Power, HVAC, net



- OS installation
- Server maintenance
- Rack & stack
- Power, HVAC, net



If you choose Amazon RDS...

App optimization

Scaling

High availability

Database backups

DB s/w patches

DB s/w installs

OS patches

OS installation

Server maintenance

Rack & stack

Power, HVAC, net



Scaling

High availability

Database backups

DB s/w patches

DB s/w installs

OS patches

OS installation

Server maintenance

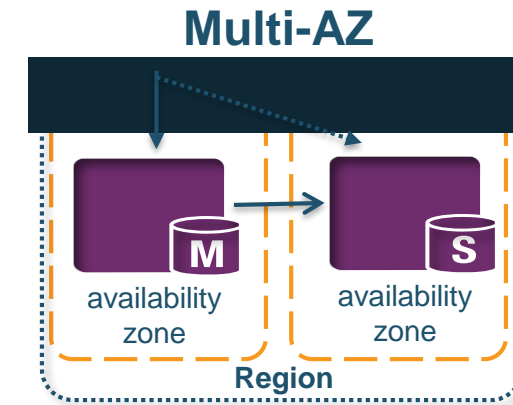
Rack & stack

Power, HVAC, net




Key Amazon RDS Features


Amazon RDS Configuration	Improve Availability	Increase Throughput	Reduce Latency
Push-Button Scaling		✓	
Multi AZ	✓		
Read Replicas		✓	
Provisioned IOPS		✓	✓




Push-Button Scaling



Read Replicas



Provisioned IOPS





Amazon
DynamoDB

- NoSQL database
- Seamless scalability
- Zero admin
- Single-digit millisecond latency
- Multi-Master
- Multi-Region
- Store any amount of data with **no limits**
- Fast, predictable performance using **SSDs**
- Easily provision and change the **request capacity** needed for each table


Amazon DynamoDB



Fully managed



Consistently fast at any scale



Highly available and durable



Secure

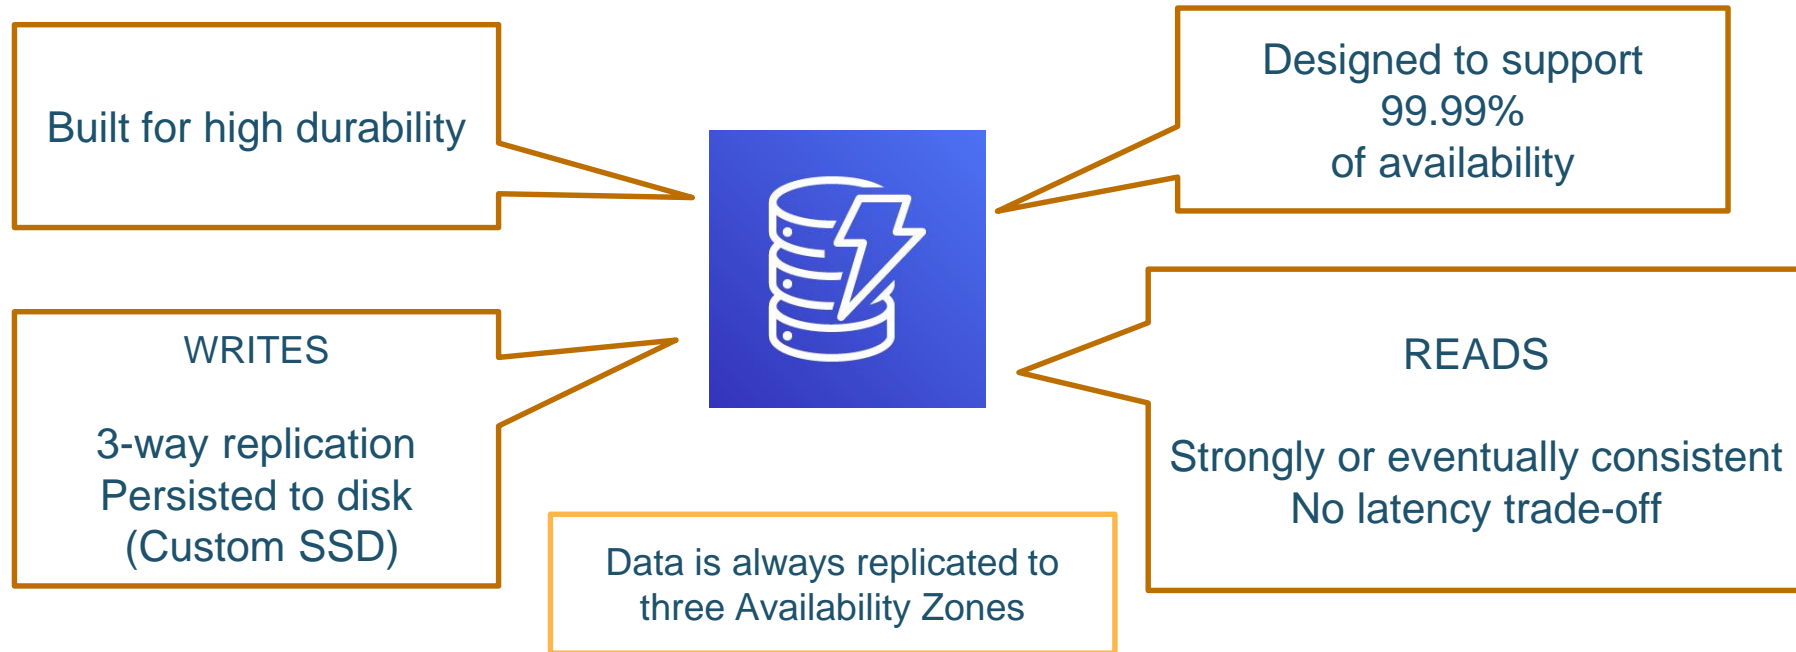


Integrates with AWS Lambda, Amazon Redshift, and more

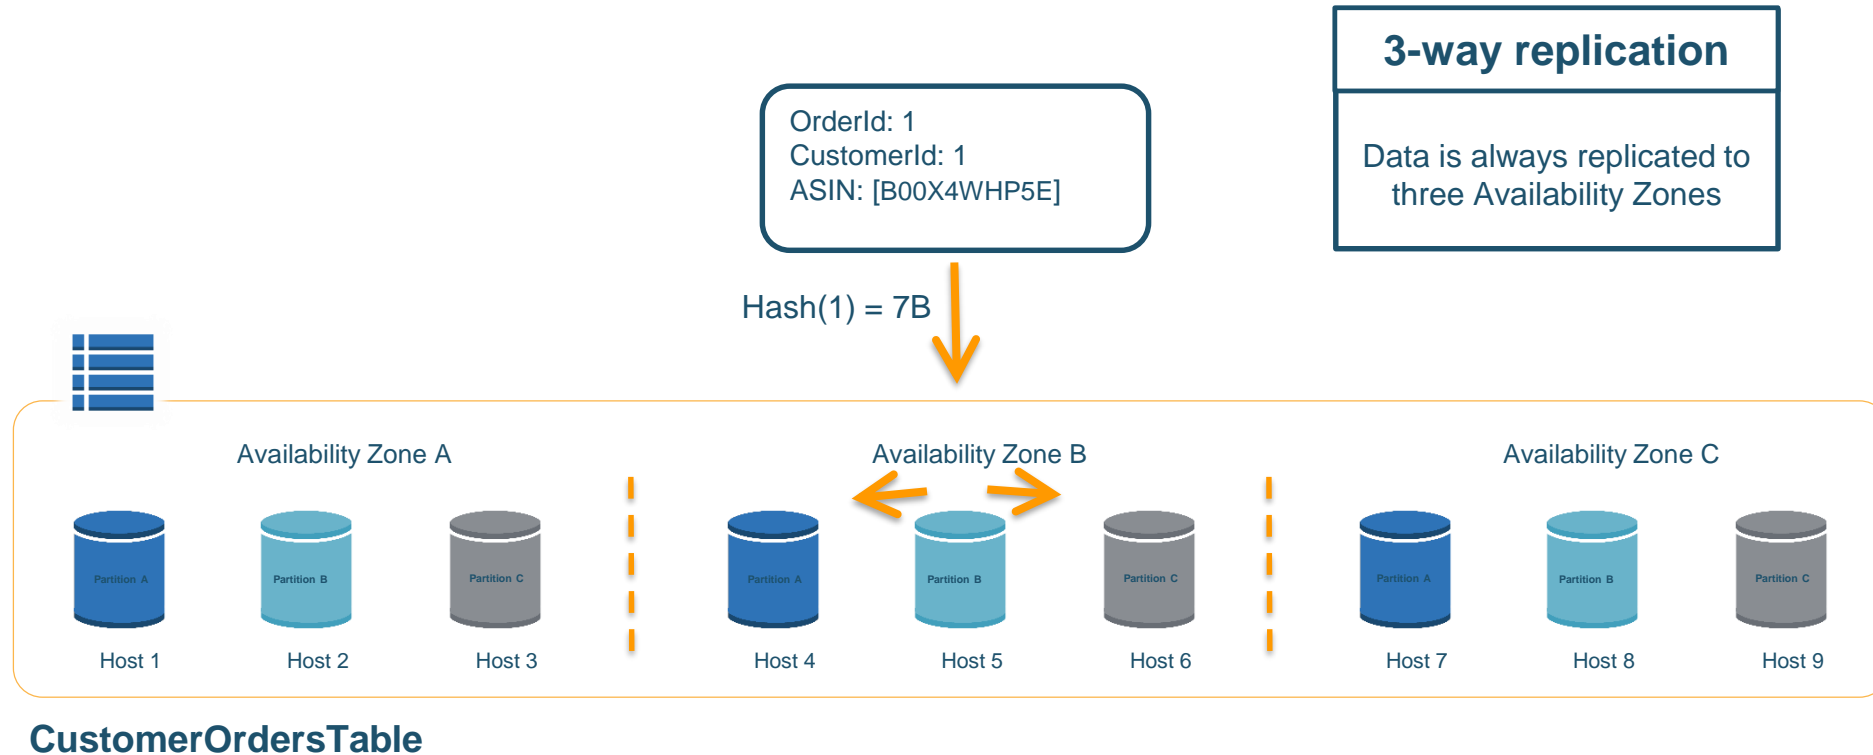


Cost-effective

Highly available and durable




Highly available and durable



Backup and Restore

The only cloud database to provide on-demand and continuous backups



On-demand backups for long-term data archival and compliance



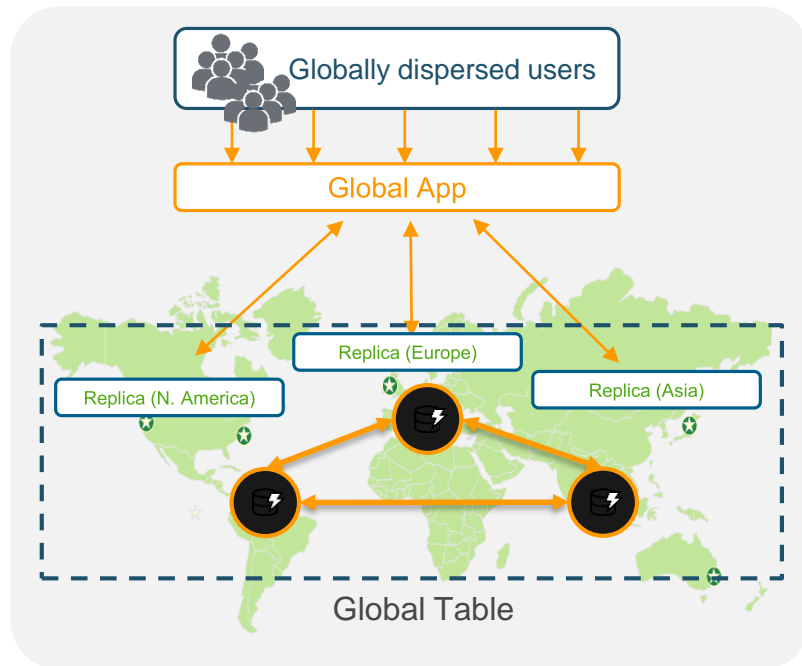
Point in time restore for short term retention and data corruption protection (35 days)



Point in time recovery with restore times in a few hours depending on table size

Global Table

The first fully-managed, multi-master, multi-region database



Build high performance, globally distributed applications

Low latency reads & writes to locally available tables

Disaster proof with multi-region redundancy

Easy to setup and no application re-writes required

Fully-managed, Redis or Memcached compatible, low-latency, in-memory data store



Amazon Elastic Cache




Extreme Performance

In-memory data store and cache for sub-millisecond response times



Fully Managed

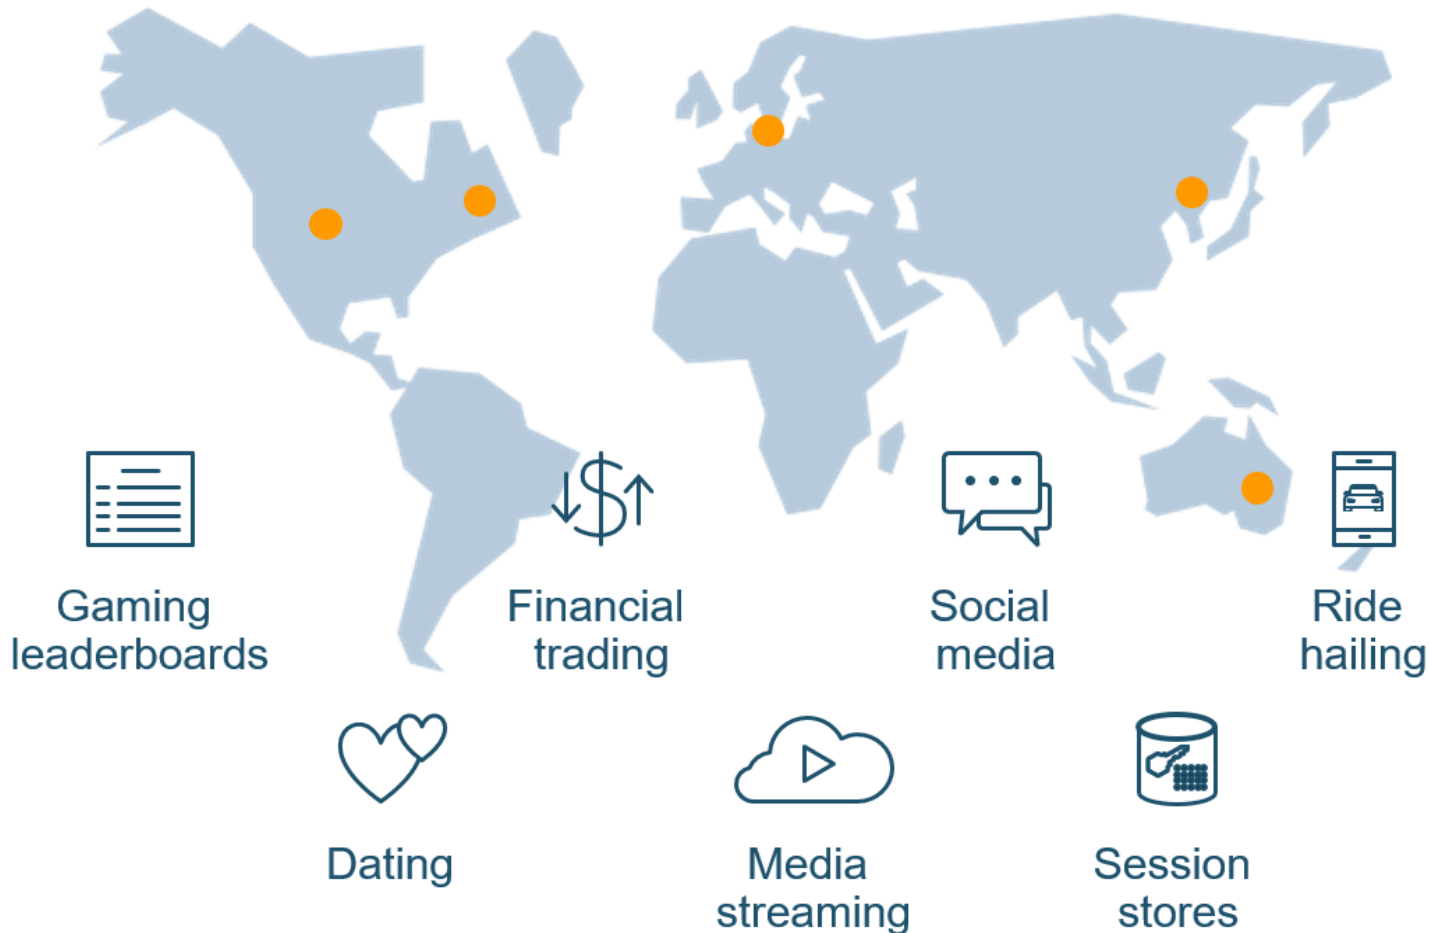
AWS manages all hardware and software setup, configuration, monitoring



Easily Scalable

Read scaling with replicas. Write and memory scaling with sharding. Non disruptive scaling

Internet-scale apps need low latency and high concurrency



Users	1M+
Data volume	TB-PB-EB
Locality	Global
Performance	Milliseconds to microseconds
Request Rate	Millions
Access	Mobile, IoT, Devices
Scale	Up-Out-In
Economics	Pay as you go
Developer access	Instant API access

Amazon ElasticCache

- In-memory cache in the cloud
- Improve latency and throughput for read-heavy workloads
- Supports open-source caching engines
 - Memcached
 - Redis
- Fully managed
- Multi-AZ



Examples

- Caching of MySQL database query results
- Caching of post-processing results
- Caching of user session and frequently accessed data



ElasticCache Redis

#1 Key-Value Store*

Fast in-memory data store in the cloud. Use as a database, cache, message broker, queue

Fully Managed & Hardened

AWS manages hardware, software, setup, configuration, monitoring, failure recovery, and backups

Secure & Compliant

VPC for cluster isolation, encryption at rest/transit, HIPAA compliance

Highly Available & Reliable

Read replicas, multiple primaries, multi-AZ with automatic failover

Easily Scalable

Cluster with up to 6.1 TiB of in-memory data

Read scaling with replicas

Write and memory scaling with sharding

Scale out or in

ElasticCache Memcached



Fully Managed Memcached

Fast in-memory data store in the cloud. Use as a cache to reduce latency and improve throughput



Secure & Hardened

VPC for cluster isolation



Easily Scalable

Sharding to scale in-memory cache with up to 20 nodes and 8.14 TiB per cluster



Amazon
RedShift

*for as low as
\$934/TB per year*

- Petabyte scale
- Massively parallel
- Columnar Store
- Relational data warehouse
- Fully managed = no admin
- Amazon Redshift manages all the work needed
- Simple way to scale a cluster to improve performance
- Continuously monitors the health of the cluster

Amazon Redshift Highlight

- Redshift is a managed data warehouse intended for analytics workloads
- Patching, backup/restore, and resize are fully managed by the service
- It uses a distributed, massively parallel architecture that scales horizontally to meet throughput requirements
- Redshift uses a c-store architecture, but still supports ANSI SQL including Transactions and Foreign Keys
- You can implement any type of data model on Redshift, but some types of data models scale better than others
- Redshift is extremely cost effective, and can offer similar performance for 1/10th the cost of Oracle, Teradata, or Netezza (as low as \$1000/TB)



Amazon
Aurora

- Serverless database, high performance
- Storage Auto-scaling, auto failover
- Fully manage by AWS
- Low latency read replicas
- Parallel Query

Amazon Aurora

MySQL and PostgreSQL compatible relational database built for the cloud

Performance and availability of commercial-grade databases at 1/10th the cost




Performance & scalability

5x throughput of standard MySQL and 3x of standard PostgreSQL; scale-out up to 15 read replicas




Availability & durability

Fault-tolerant, self-healing storage; six copies of data across three AZs; continuous backup to S3



Highly secure

Network isolation, encryption at rest/transit

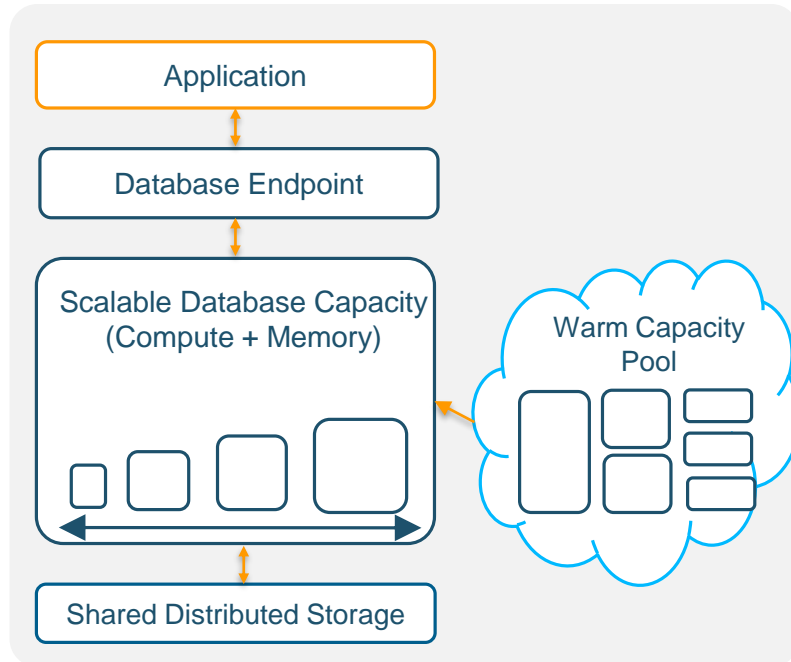


Fully managed

Managed by RDS: no hardware provisioning, software patching, setup, configuration, or backups

Aurora Serverless

On-demand, auto-scaling database for applications with variable workloads



Starts up on demand, shuts down when not in use

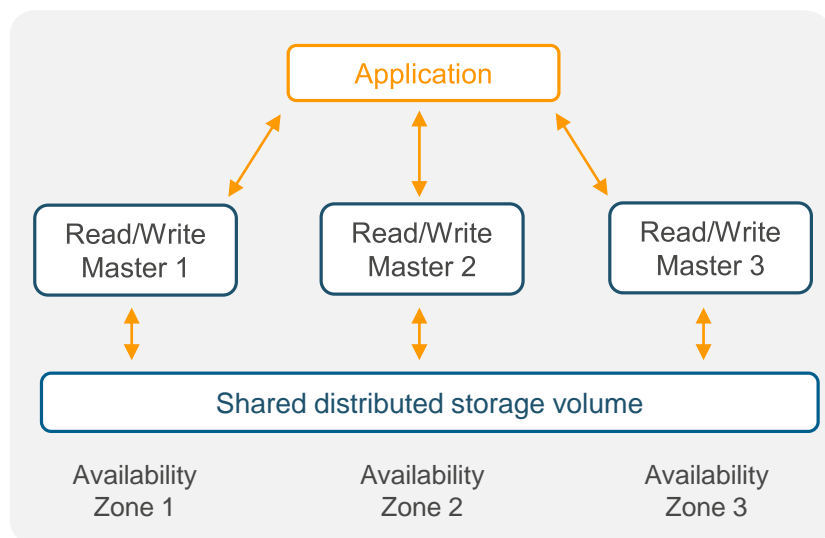
Automatically scales with no instances to manage

Pay per second for the database capacity you use

Aurora Multi Master

First relational database service with scale-out reads and writes across multiple data centers

Scale out both reads **and** writes



Zero application downtime from ANY instance failure

Zero application downtime from ANY AZ failure

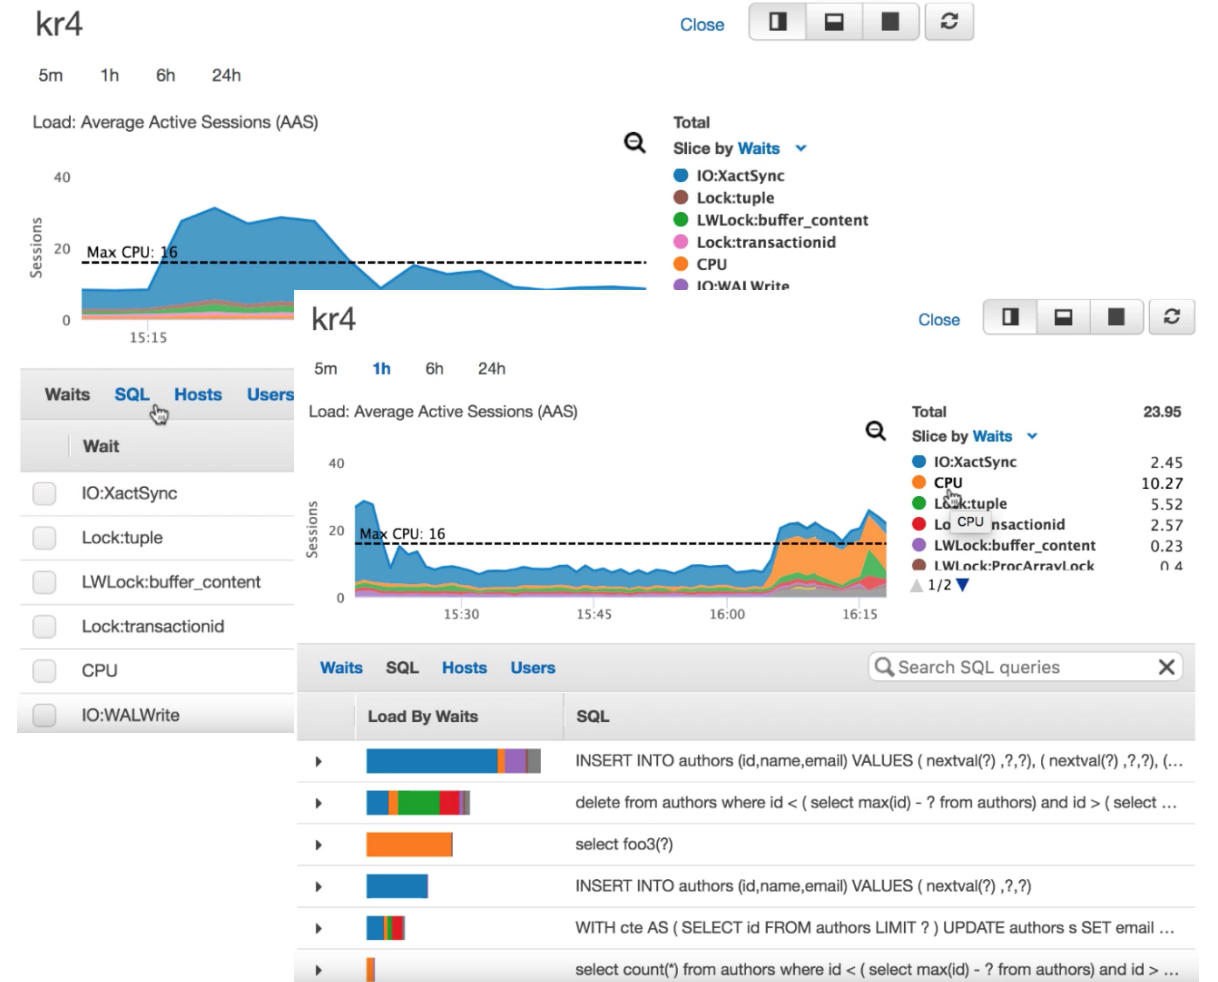
Faster write performance and higher scale

Sign up for single-region multi-master preview today;
multi-region multi-master **coming**

Performance Insights for Aurora

Analyze and troubleshoot your database performance

- Supports PostgreSQL and MySQL
- Expands on existing Amazon RDS monitoring features to analyze issues and performance
- Easy bottleneck identification – keep track of performance metrics such as high CPU consumption, lock waits, I/O latency, and SQL statements



Aurora Global Database

Faster disaster recovery and enhanced data locality

Promote read-replica to a master for faster recovery **in the event of disaster**

Bring data close to your customer's applications in **different regions**

Promote to a master for **easy migration**



Module 3: Security and Implementation



VIỆN ĐIỆN TỬ - VIỄN THÔNG
School of Electronics and Telecommunications



Agenda

- AWS Security Overview
- Infrastructure Security
- AWS Monitoring & Logging

AWS Security Overview



Broad Accreditations & Certifications

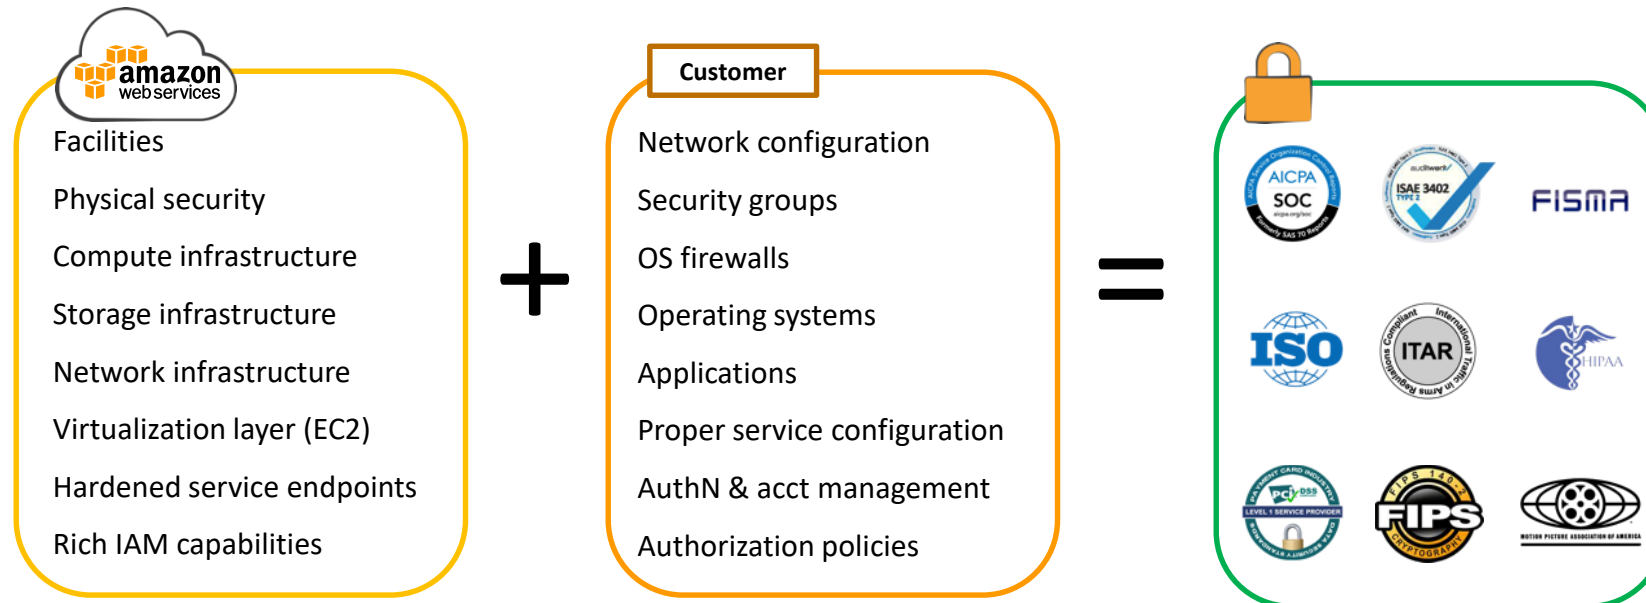


Glacier Vault Lock
& SEC Rule 17a-4(f)

See <https://aws.amazon.com/compliance/programs/> for full list

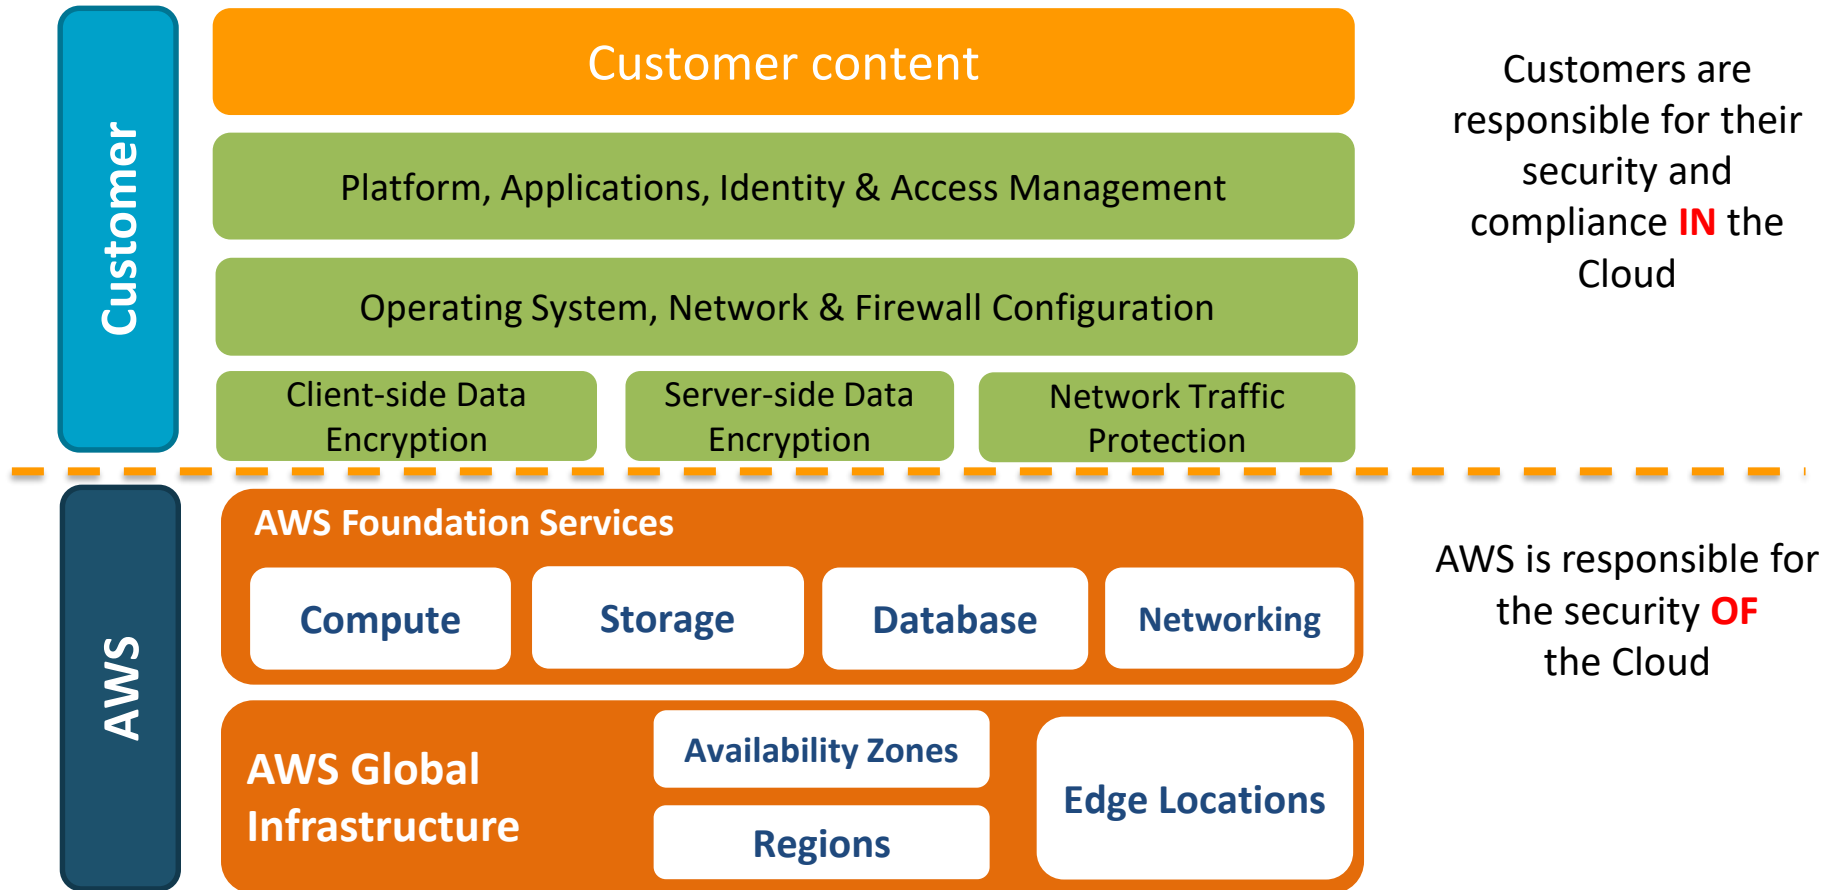


AWS Shared Responsibility Model



- Scope of responsibility depends on the type of service offered by AWS: **Infrastructure, Container, Abstracted Services**
- Understanding who is responsible for what is critical to ensuring your AWS data and systems are secure!

Shared Responsibility Model



Meet your own security objectives



AWS Responsibilities

Physical Security of Data Center

- Amazon has been building large-scale data centers for many years.
- **Important attributes:**
 - Non-descript facilities
 - Robust perimeter controls
 - Strictly controlled physical access
 - Two or more levels of two-factor authentication
- **Controlled, need-based access.**
- **All access is logged and reviewed.**
- **Separation of Duties**
 - Employees with physical access don't have logical privileges.



AWS Responsibilities

EC2 Security

- **Host (hypervisor) operating system**
 - Individual SSH keyed logins via bastion host for AWS admins
 - All accesses logged and audited
- **Guest (EC2 Instance) operating system**
 - Customer controlled (customer owns root/admin)
 - AWS admins cannot log in
 - Customer-generated keypairs
- **Stateful firewall**
 - Mandatory inbound firewall, default deny mode
 - Customer controls configuration via Security Groups



Network Security

- IP Spoofing prohibited at host OS level.
- Packet sniffing (promiscuous mode) is ineffective (protected at hypervisor level).
- Unauthorized Port Scanning a violation of TOS and is detected/blocked.
- Inbound ports blocked by default.

AWS Responsibilities

Configuration Management

- Most updates are done in such a manner that they will not impact the customer.
- Changes are authorized, logged, tested, approved, and documented.
- AWS will communicate with customers, either via email, the AWS Service Health Dashboard (<http://status.aws.amazon.com/>), or the AWS Personal Health Dashboard (<https://phd.aws.amazon.com/>) when there is a potential for service being affected.

Built for “Continuous Availability”

- **Scalable, fault tolerant services.**
- **All availability zones (AZs) are always on.**
 - There is no “Disaster Recovery Datacenter”
 - All managed to the same standards
- **Robust Internet connectivity**
 - Each AZ has redundant, Tier 1 ISP Service Providers
 - Resilient network infrastructure

AWS Responsibilities

Disk Management

- Proprietary disk management prevents customers from accessing each other's data.
- Disks wiped prior to use.
- Disks can also be encrypted by the customer for additional security.

Storage Device Decommissioning

- All storage devices go through process using techniques from:
 - DoD 5220.22-M (“National Industrial Security Program Operating Manual”).
 - NIST 800-88 (“Guidelines for Media Sanitization”).
- Ultimately devices are:
 - Degaussed.
 - Physically destroyed.

Identity and Access Management

“...the management of individual **principals**, their **authentication**, **authorization**, and **privileges** ...with the goal of increasing security and productivity while decreasing cost, downtime and repetitive tasks.”



AAA with AWS

Authenticate

IAM Username/Password
Access Key
(+ MFA)
Federation

Authorize

IAM Policies

Audit

CloudTrail

Considerations for Layers of Principals


Applications

- Identities: Application Users, Application Administrators



Operating Systems

- Identities: Developers, and/or Systems Engineers



Amazon Web Services

- Identities: Developers, Solutions Architects, Testers, Software/Platform
- Interaction of AWS Identities:
 - Provisioning/deprovisioning EC2 instances and EBS storage.
 - Configuring Elastic Load Balancers.
 - Accessing S3 Objects or data in DynamoDB.
 - Accessing data in DynamoDB.
 - Interacting with SQS queues.
 - Sending SNS notifications.



AWS Principals

Account Owner ID (Root Account)

- Access to all subscribed services.
- Access to billing.
- Access to console and APIs.
- Access to Customer Support.



IAM Users, Groups and Roles

- Access to specific services.
- Access to console and/or APIs.
- Access to Customer Support (Business and Enterprise).



Temporary Security Credentials

- Access to specific services.
- Access to console and/or APIs.

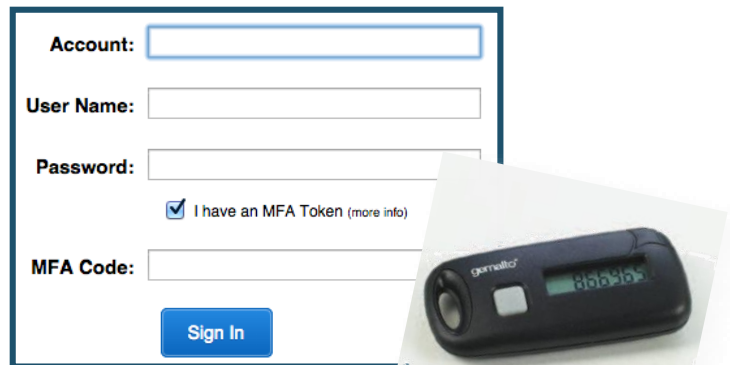


AWS Identity Authentication

Authentication: How do we know you are who you say you are?

AWS Management Console

Login with **Username/Password** with optional **MFA** (recommended)



Account:

User Name:

Password:

I have an MFA Token (more info)

MFA Code:

Sign In

For time-limited access: a **Signed URL** can provide temporary access to the Console

API access

Access API using **Access Key + Secret Key**, with optional MFA

ACCESS KEY ID

Ex: AKIAIOSFODNN7EXAMPLE

SECRET KEY

Ex: UtnFEMI/K7MDENG/bPxRfiCYEXAMPL



For time-limited access: Call the AWS Security Token Service (STS) to get a temporary AccessKey + SecretKey + session token

AWS Authorization and Privileges

Authorization: What are you allowed to do?

Account Owner (Root)

- Privileged for all actions.

Note: Always associate the account owner ID with an MFA device and store it in a secured place!

IAM Policies

- Privileges defined at User and Resource Level

You are accessing the security credentials page for your AWS account. The account credentials provide unlimited access to your AWS resources.

To help secure your account, follow an [AWS best practice](#) by creating and using AWS Identity and Access Management (IAM) users with limited permissions.

[Continue to Security Credentials](#) [Get Started with IAM Users](#)

Don't show me this message again

Permissions

This view shows all policies that apply to this User. This includes policies that are assigned to groups that this User belongs to.

User Policies

There are no policies attached to this user.

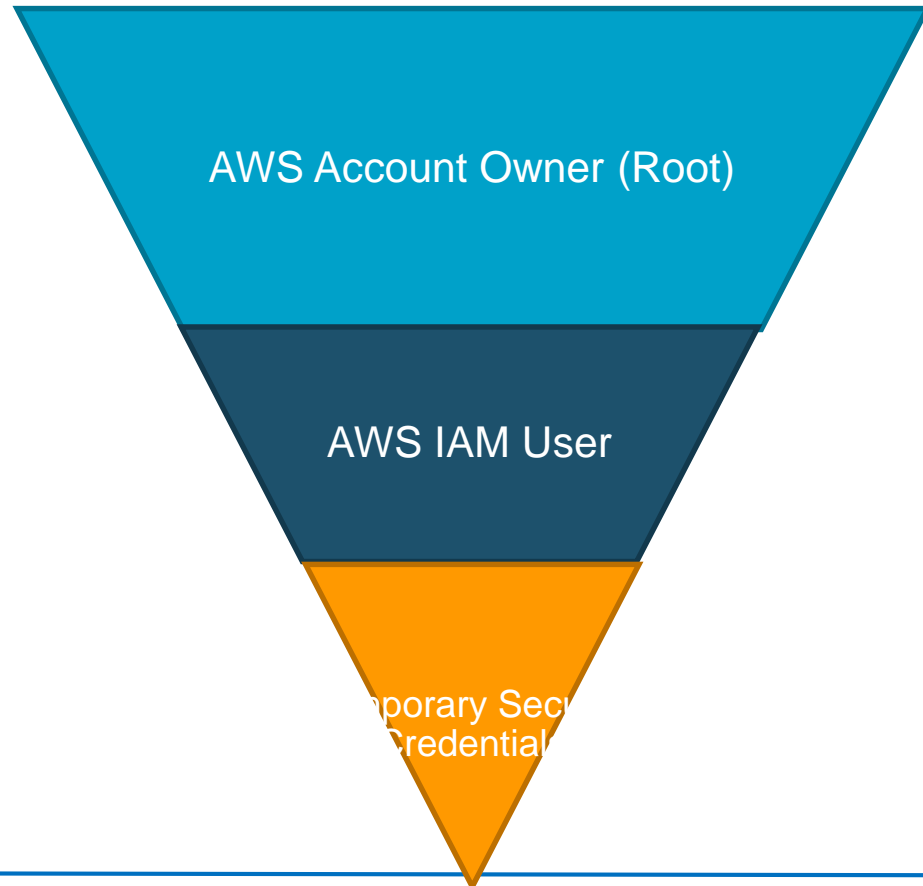
[Attach User Policy](#)

Group Policies

Policy Name	Group Name
AdministratorAccess-Administrators-201408161823 Show	Administrators
AdministratorAccess-Demo-201410281057 Show	Demo

AWS IAM Hierarchy of Privileges

Enforce principle of least privilege with Identity and Access Management (IAM) users, groups, and policies and temporary credentials.



Permissions	Example
Unrestricted access to all enabled services and resources.	Action: * Effect: Allow Resource: * (implicit)
Access restricted by Group and User policies	Action: ['s3:*', 'sts:Get*'] Effect: Allow Resource: *
Access restricted by generating identity and further by policies used to generate token	Action: ['s3:Get*'] Effect: Allow Resource: 'arn:aws:s3:::mybucket/*'

AWS Identity and Access Management (IAM)

Securely control access to AWS services and resources for your users.

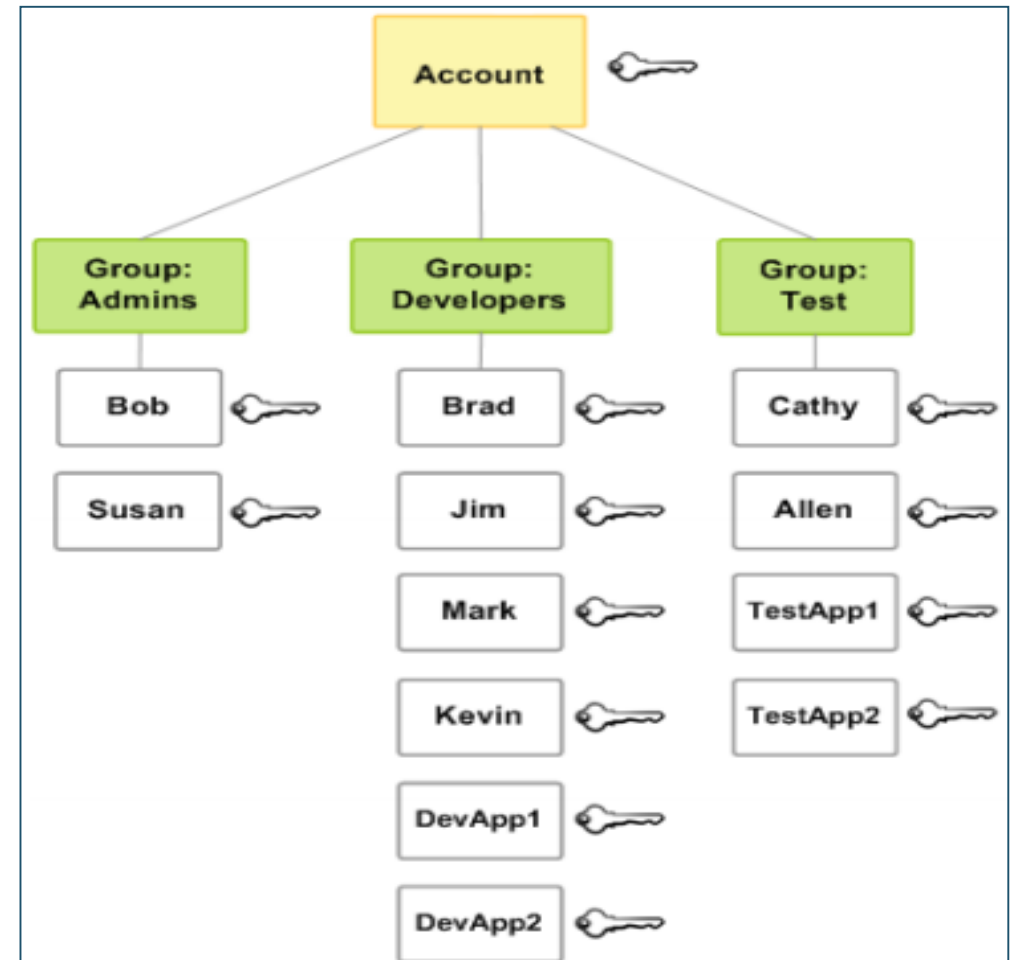
Username/
User

Manage groups of
users

Centralized
Access Control

Optional Configurations:

- Password for console access.
- Policies for controlling access AWS APIs.
- Two methods to sign API calls:
 - X.509 certificate
 - Access/Secret Keys
- Multi-factor Authentication (MFA)




Identity and Access Management

Common approaches for Applications and Operating Systems

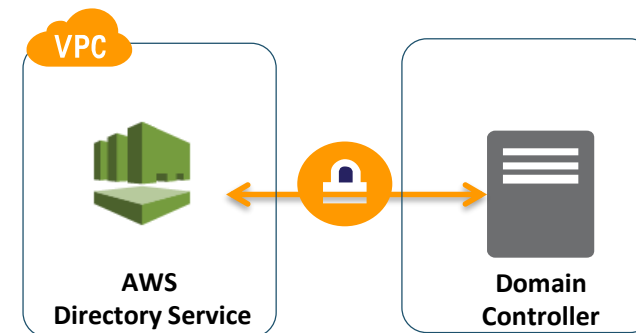
Local User Databases

- Local Password (passwd) files
- Local Windows admin accounts
- User Databases



LDAP Directories

- On-premise accessed over VPN.
- Replicated to AWS (read-only or read/write)
- Federated (one-way trusts, ADFS).
- Managed Samba-based directories via AWS Directory Services.



AWS Directory Service

Managed service for Active Directory

Use your existing Corporate Credentials for

- AWS-based applications
- AWS Management Console



Microsoft AD

Based on Microsoft Active Directory in Windows Server 2012 R2. Supports adding trust relationships with on-premises domains. Extend your schema using MS AD



Simple AD

A Microsoft Active-Directory compatible directory powered by Samba 4.



AD Connector

Connect to your on-premises Active Directory. Integrates with existing RADIUS MFA solutions.

AWS Encryption

Protecting data in-transit and at-rest.



Encryption In-Transit

HTTPS

SSL/TLS

VPN / IPSEC

SSH

Encryption At-Rest

Object


Database

Filesystem

Disk

*Details about encryption can be found in the AWS Whitepaper,
[“Securing Data at Rest with Encryption”](#).*

AWS Certificate Manager

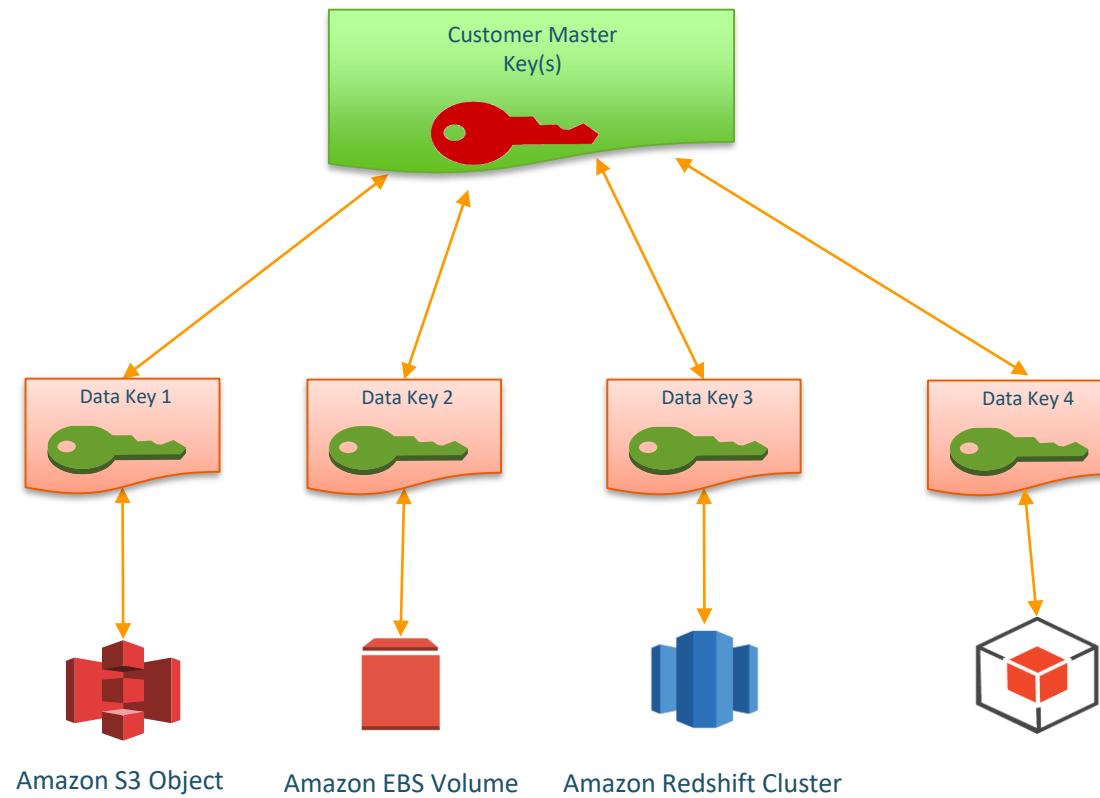


AWS Certificate Manager (ACM) makes it easy to provision, manage, deploy, and renew SSL/TLS certificates on the AWS platform.

AWS Key Management Service



Managed service to securely create, control, rotate, and use encryption keys.

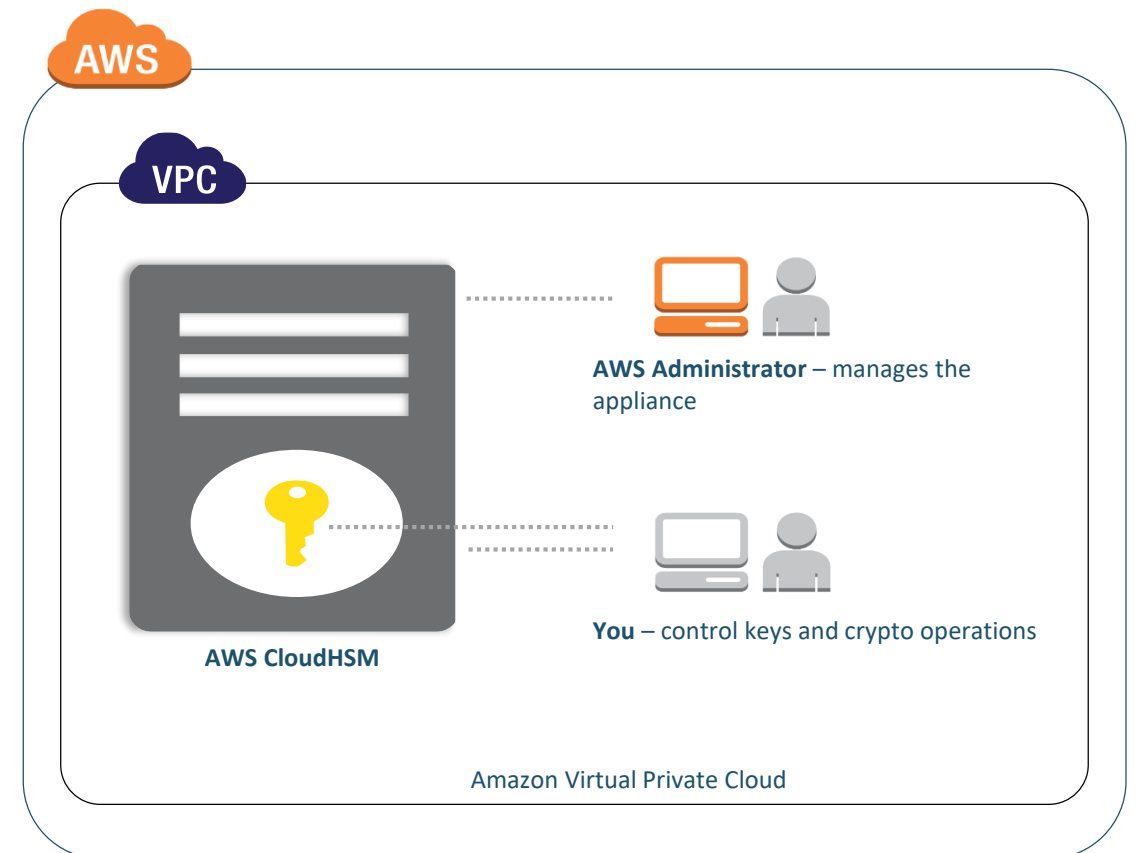


AWS CloudHSM

Help meet compliance requirements for data security by using a dedicated Hardware Security Module appliance with AWS.


- Dedicated, single-tenant hardware device
- Can be deployed as HA and load balanced

- Customer use cases:
 - Oracle TDE
 - MS SQL Server TDE
 - Setup SSL connections
 - Digital Rights Management (DRM)
 - Document Signing




Configuration Management - Amazon Inspector

- Vulnerability Assessment Service
 - Built from the ground up to support DevSecOps
 - Automatable via APIs
 - Integrates with CI/CD tools
 - On-Demand Pricing model
 - Static & Dynamic Rules Packages
 - Generates Findings



AWS WAF



Web Traffic Filtering with Custom Rules

Create custom rules that can block, allow or monitor requests based on IP address, HTTP headers, or a combination of both.



Malicious Request Blocking

AWS WAF can recognize and block common web application security risks like SQL injection (SQLi) and cross-site scripting (XSS).



Active monitoring & tuning

Monitor and configure the requests that are being blocked and allowed by the Web ACL rules.

AWS CloudTrail

Web service that records AWS API calls for your account and delivers logs.

Who?	When?	What?	Where to?	Where from?
Bill	3:27pm	Launch Instance	us-west-2	72.21.198.64
Alice	8:19am	Added Bob to admin group	us-east-1	127.0.0.1
Steve	2:22pm	Deleted DynamoDB table	eu-west-1	205.251.233.176












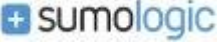

















```
{
  "Records": [
    {
      "eventVersion": "1.0",
      "userIdentity": {
        "type": "IAMUser",
        "principalId": "EX_PRINCIPAL_ID",
        "arn": "arn:aws:iam::123456789012:user/Alice",
        "accountId": "123456789012",
        "accessKeyId": "EXAMPLE_KEY_ID",
        "userName": "Alice",
        "sessionContext": {
          "attributes": {
            "mfaAuthenticated": "false",
            "creationDate": "2014-03-25T18:45:11Z"
          }
        }
      },
      "eventTime": "2014-03-25T21:08:14Z",
      "eventSource": "iam.amazonaws.com",
      "eventName": "AddUserToGroup",
      "awsRegion": "us-east-1",
      "sourceIPAddress": "127.0.0.1",
      "userAgent": "AWSConsole",
      "requestParameters": {
        "userName": "Bob",
        "groupName": "admin"
      },
      "responseElements": null
    },
    ...additional entries
  ]
}
```

Amazon Macie

Leverage Amazon Macie to help prevent data loss in AWS.

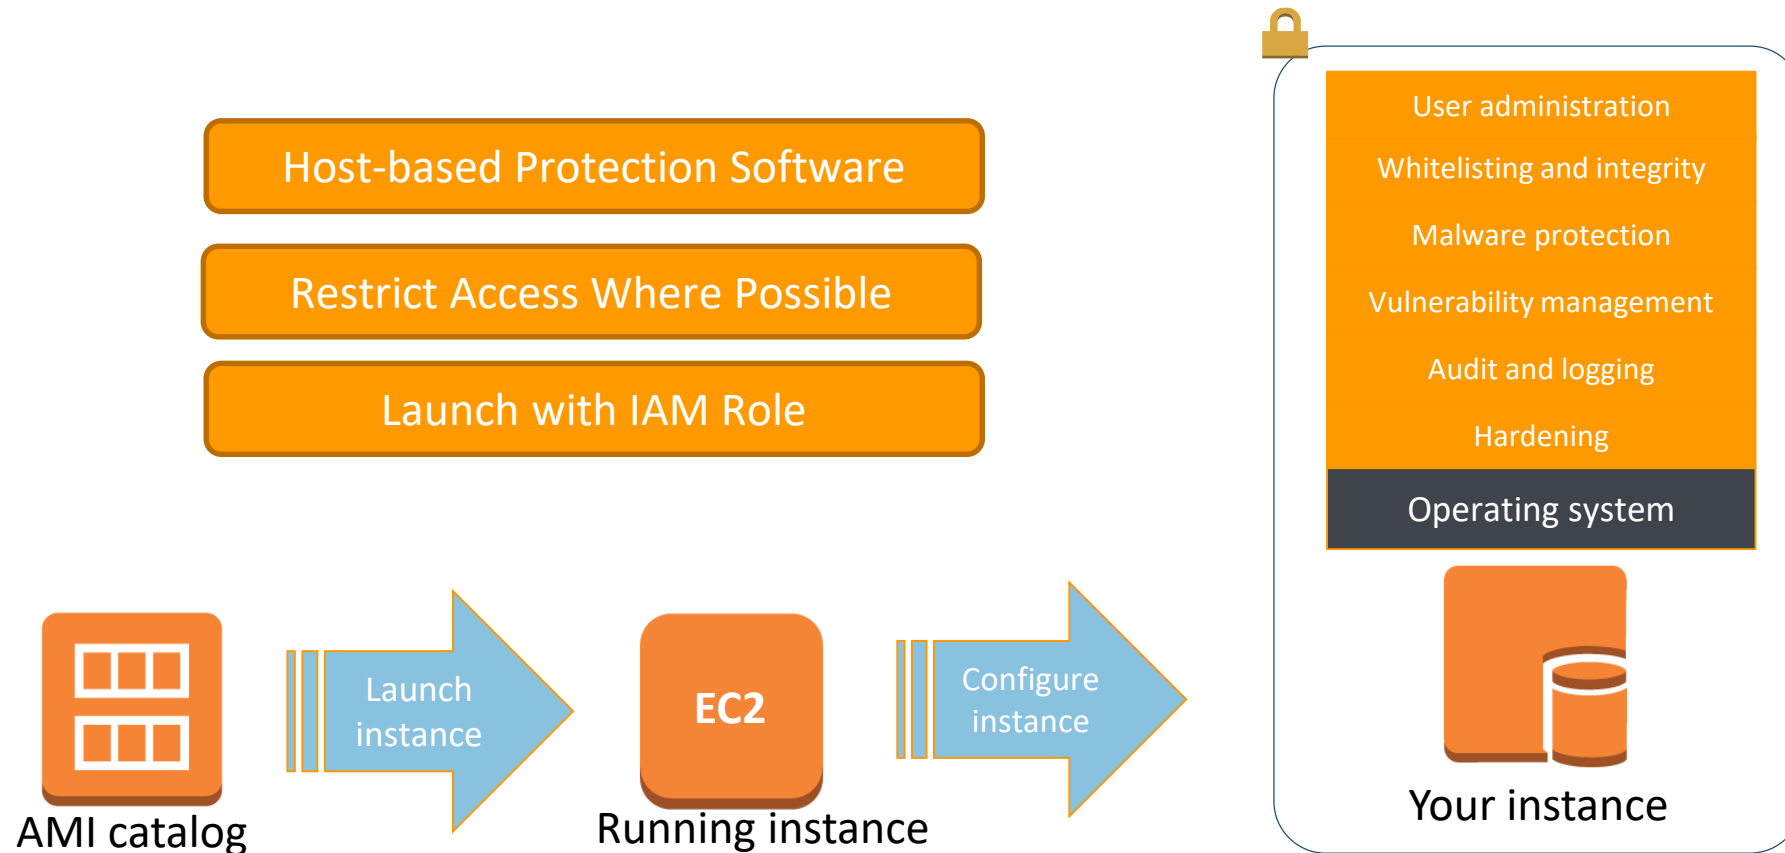
The screenshot displays the Amazon Macie console interface. On the left is a navigation sidebar with options: ALERTS, DASHBOARD, REPORTS, RESEARCH, SETTINGS, and INTEGRATIONS. The main dashboard area features four key metrics: Critical Assets (0.00% of all) with a value of 2 (Risk level 8 to 10), Total Events (681k count of events), Total User Sessions (147k count of user sessions), and Total users (44) represented by four icons with counts 4, 3, 23, and 14. Below these is a 'Minimum Risk: 6' slider and 'Total Matching Themes: unique risky Themes'. A section titled 'Amazon S3 content for selected time range - minRisk: (6)' contains a donut chart and a legend. The legend includes categories like 'All Data', 'Range: 0 - 6 months ago', 'Range: beyond 6 months ago', 'Amazon Access Key Headers', 'Confidential Markings', 'Large number of IPv4 addresses', 'Proprietary Markings', 'aws_access_key', 'aws_credentials_context', 'aws_secret_key', 'email/all', 'json/aws_cloudtrail_logs', and 'json/other'. On the right, the 'ALERTS' section shows a list of alerts. The top alert is 'S3 Bucket uses IAM policy to grant read rights to Everyone' (100 alerts), categorized as CUSTOM_ALERT and DLP, with 11 minutes ago, 0 comments, and 0 views. The second alert is identical. The third alert is 'Access Denied In Secure Account' (50 alerts), also categorized as CUSTOM_ALERT and DLP, with 30 minutes ago, 0 comments, and 0 views.

AWS Marketplace Security Partners

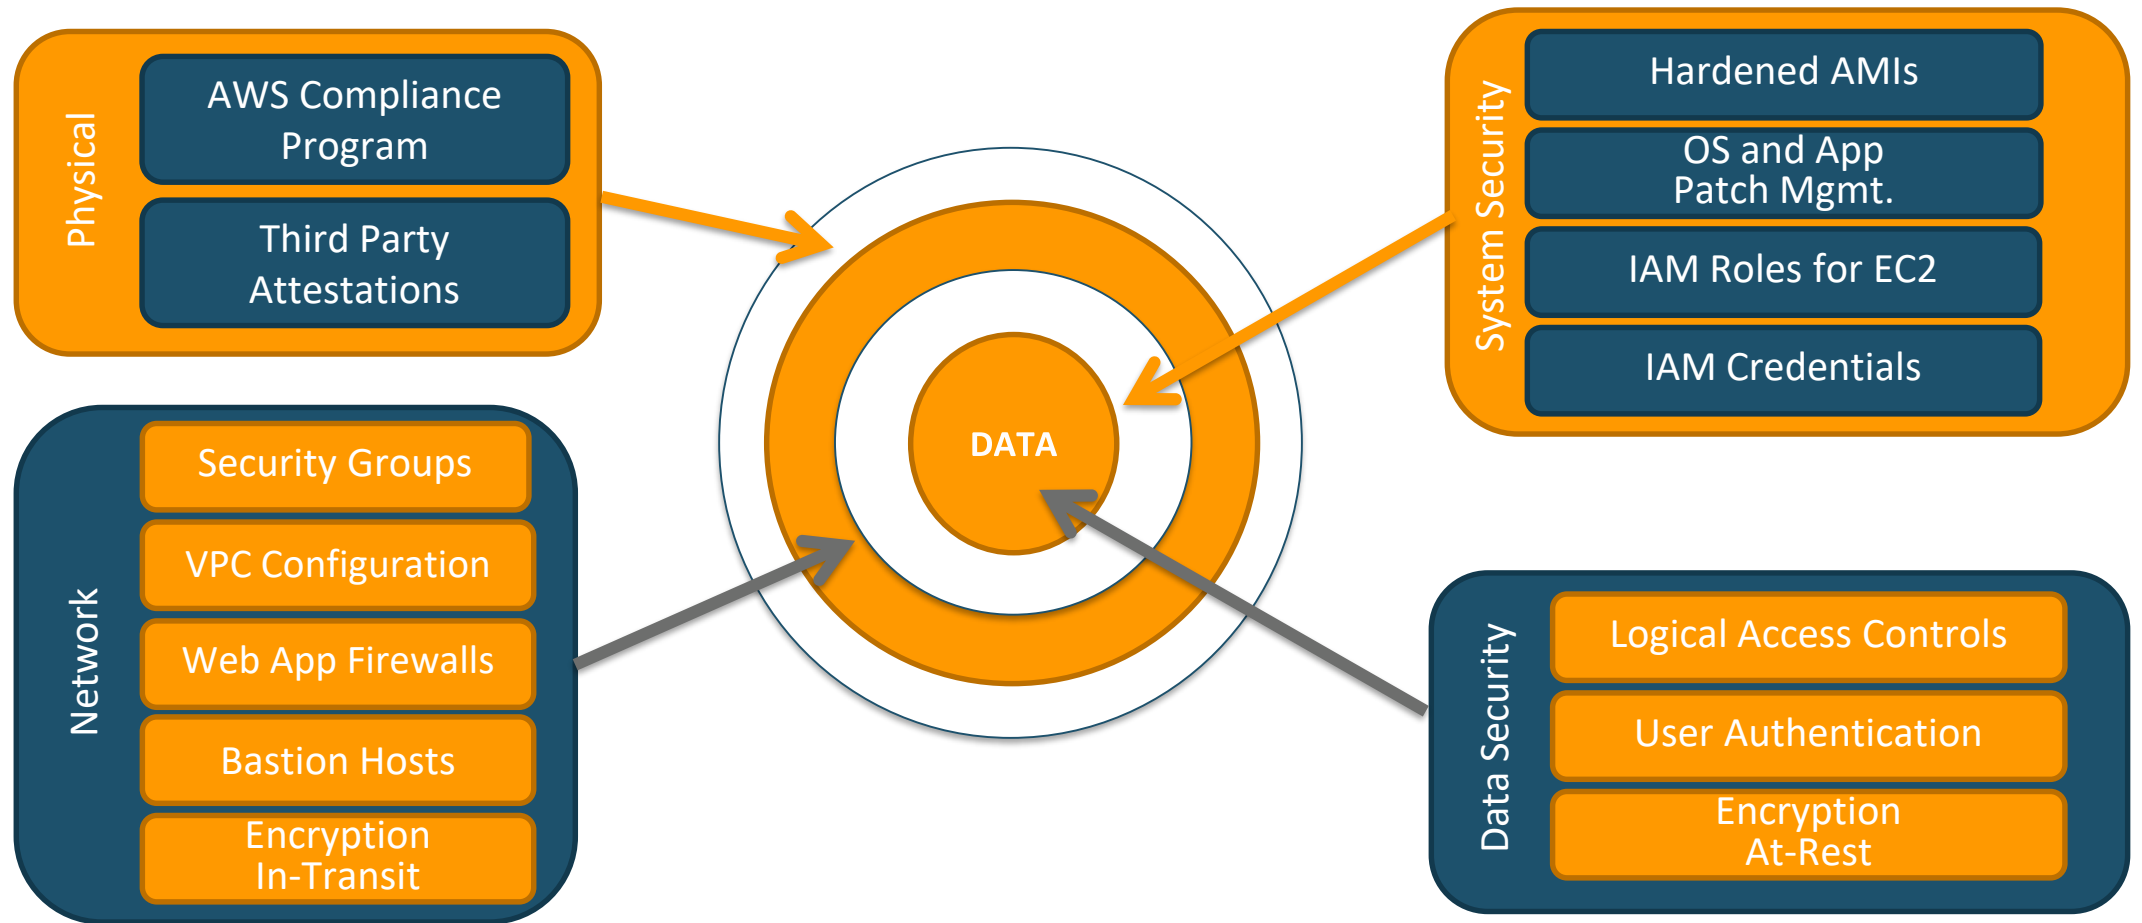
Infrastructure Security	Logging & Monitoring	Identity & Access Control	Configuration & Vulnerability Analysis	Data Protection
         	    	    	    	     

Enforce consistent security on your hosts

Configure and harden EC2 instances based on security and compliance needs.



Defense-in-Depth



AWS Security Center

Comprehensive security portal to provide a variety of security notifications, information and documentation.

<http://aws.amazon.com/security>



Security Whitepapers

- Overview of Security Process
- AWS Risk and Compliance
- AWS Security Best Practices

Security Bulletin

Security Resources

Vulnerability Reporting

Penetration Testing

Requests

Report Suspicious Emails



AWS Security Center

Security Resources

<http://aws.amazon.com/security/security-resources/>

Developer Information, Articles and Tutorials,
Security Products, and Whitepapers

Menu | amazon web services | Products | More | English | My Account | Sign In to the Console

Security Resources

Contact AWS Sales

For more information about security features that AWS provides or how to stay safe in the cloud, click on a link below.

Developer Documents

See how to configure important security settings within EC2, S3, and other AWS services:

- Signing AWS API Requests
- Using Encryption in S3
- Configuring EC2 Security Groups
- Server Access Logging in S3
- List of Secure Endpoints
- AWS Security Credentials
- Turning on CloudTrail Logging

Articles & Tutorials

Read tips and tricks from AWS experts on using certain tools and features to architect and configure AWS services securely:

- Amazon S3 Bucket Public Access Considerations

AWS Security Blog

<http://blogs.aws.amazon.com/security/>

Subscribe to the blog – it's a great way to stay up-to-date on
AWS security and compliance.

amazon web services | Security Blog

Stay up to date on security and compliance in AWS

The IAM Console Now Helps Prevent You From Accidentally Deleting In-Use Resources

January 13, 2016 | Kai Zhao | Announcements | Access Management | IAM | Permissions | Policies | Resource deletion

Deleting unused resources can help to improve the security of your AWS account and make your account easier to manage. However, if you have ever been unsure of whether an AWS Identity and Access Management (IAM) user or role was being used actively, you probably erred on the side of caution and kept it.

Starting today, the IAM console shows [service last accessed data](#) as part of the process of deleting an IAM user or role. Now you have additional data that shows you when a resource was last active so that you can make a more informed decision about whether or not to delete it.

Read More →

January 13, 2016 | Permalink | Comments (0) | Share | | | |

Adhere to IAM Best Practices in 2016

January 6, 2016 | Craig Liebendorfer | Announcements | Best Practices | credentials | EC2 | IAM | least privilege | MFA | password | permissions | policy conditions | roles | users

As another new year begins, we encourage you to review our recommended AWS Identity and Access Management (IAM) best practices. Following these best practices can help you maintain the security of your AWS resources. You can learn more by watching the [IAM Best Practices to Live By presentation](#) that Anders Samuelsson gave at AWS re:Invent 2015, or you can click the following links that will take you to IAM documentation, blog posts, and videos.

Read More →

Follow us on Twitter | RSS

Latest Blog Entries

- The IAM Console Now Helps Prevent You From Accidentally Deleting In-Use Resources
- Adhere to IAM Best Practices in 2016
- The Most Popular AWS Security Blog Posts in 2015
- AWS ISO 27001 Certification Increases Total In-Scope Services to 33
- Another Way to Remove Unnecessary Permissions Your IAM Policies by Using Service Last Accessed Data
- How to Automatically Update Your Security Groups for Amazon CloudFront and AWS WAF by Using AWS Lambda
- AWS Certification Update – ISO 9001 Has 10 New Services in Scope
- How to Set Up SSO to the AWS Management Console for Multiple Accounts by Using AD FS and SAML 2.0



AWS Compliance

List of compliance, assurance programs and resources:

<http://aws.amazon.com/compliance/>.



Glacier Vault Lock & SEC Rule 17a-4(f)



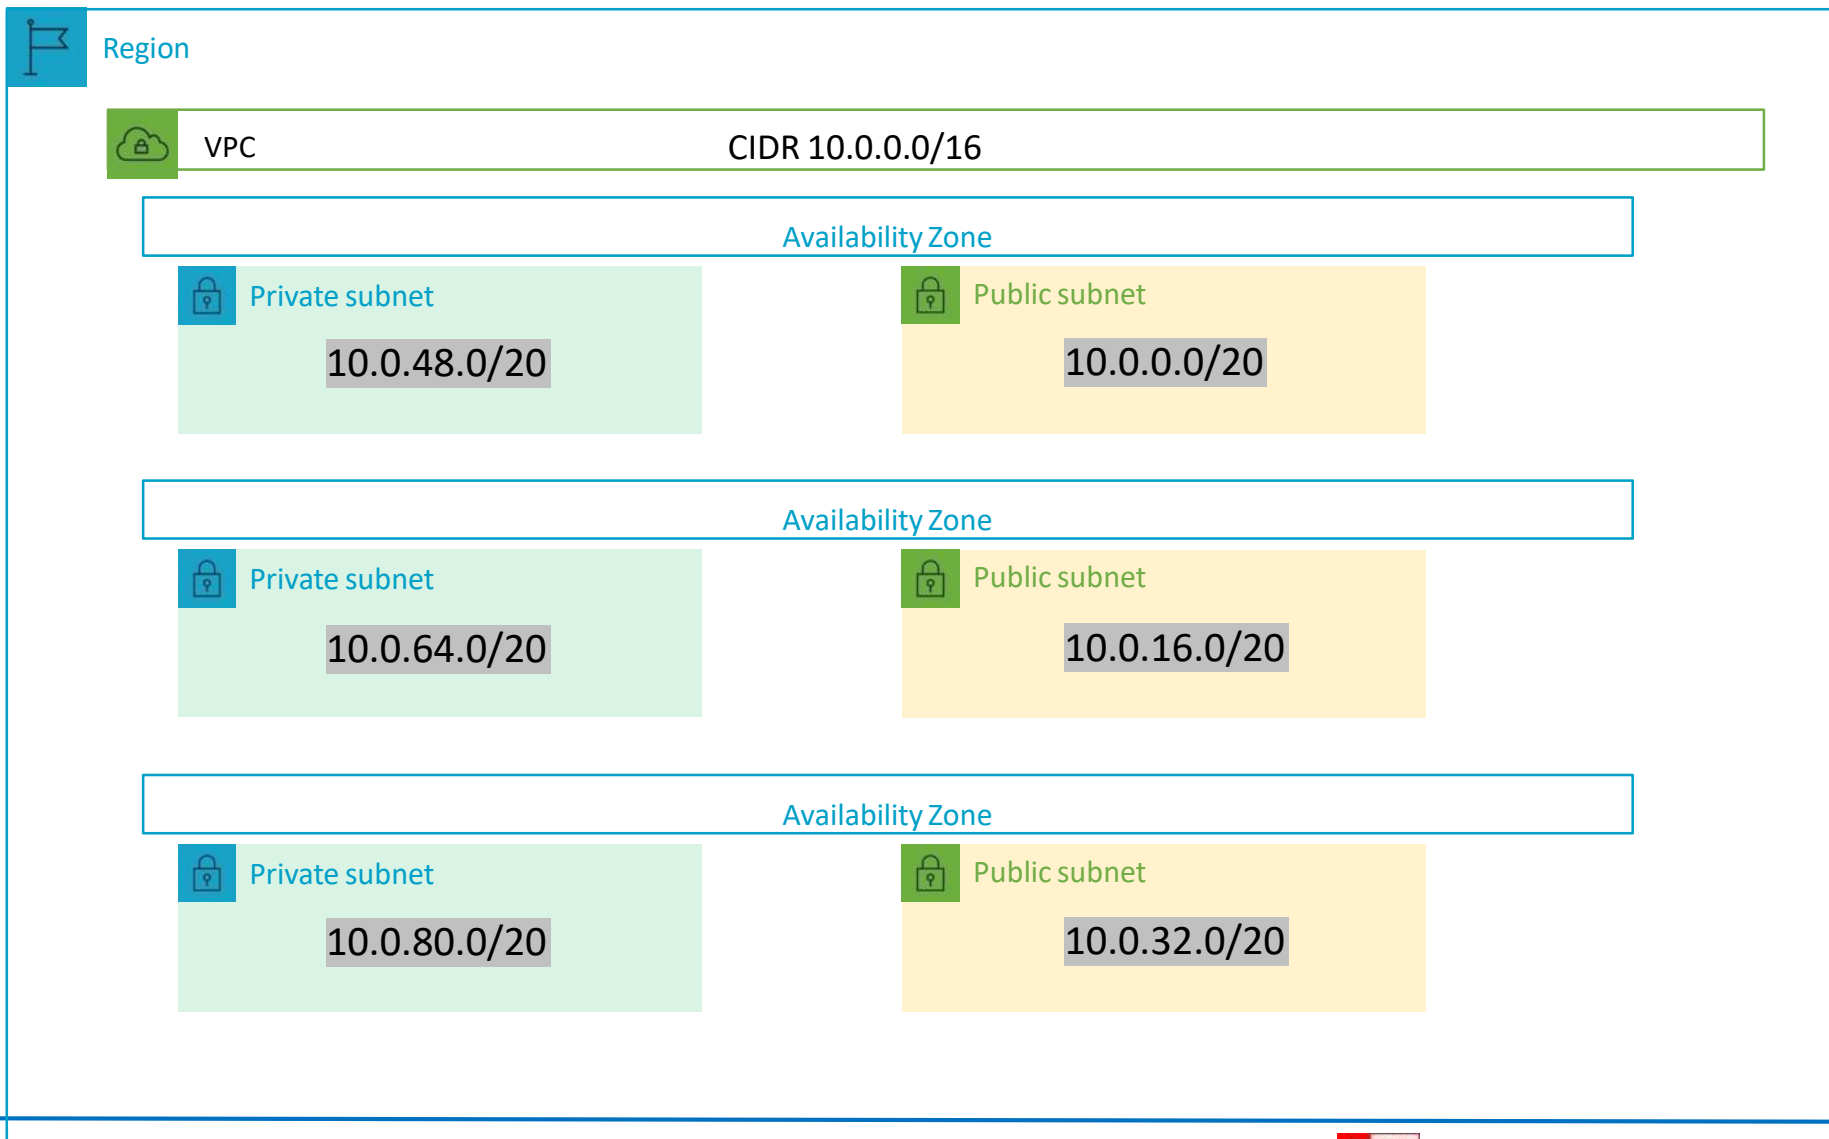
27018



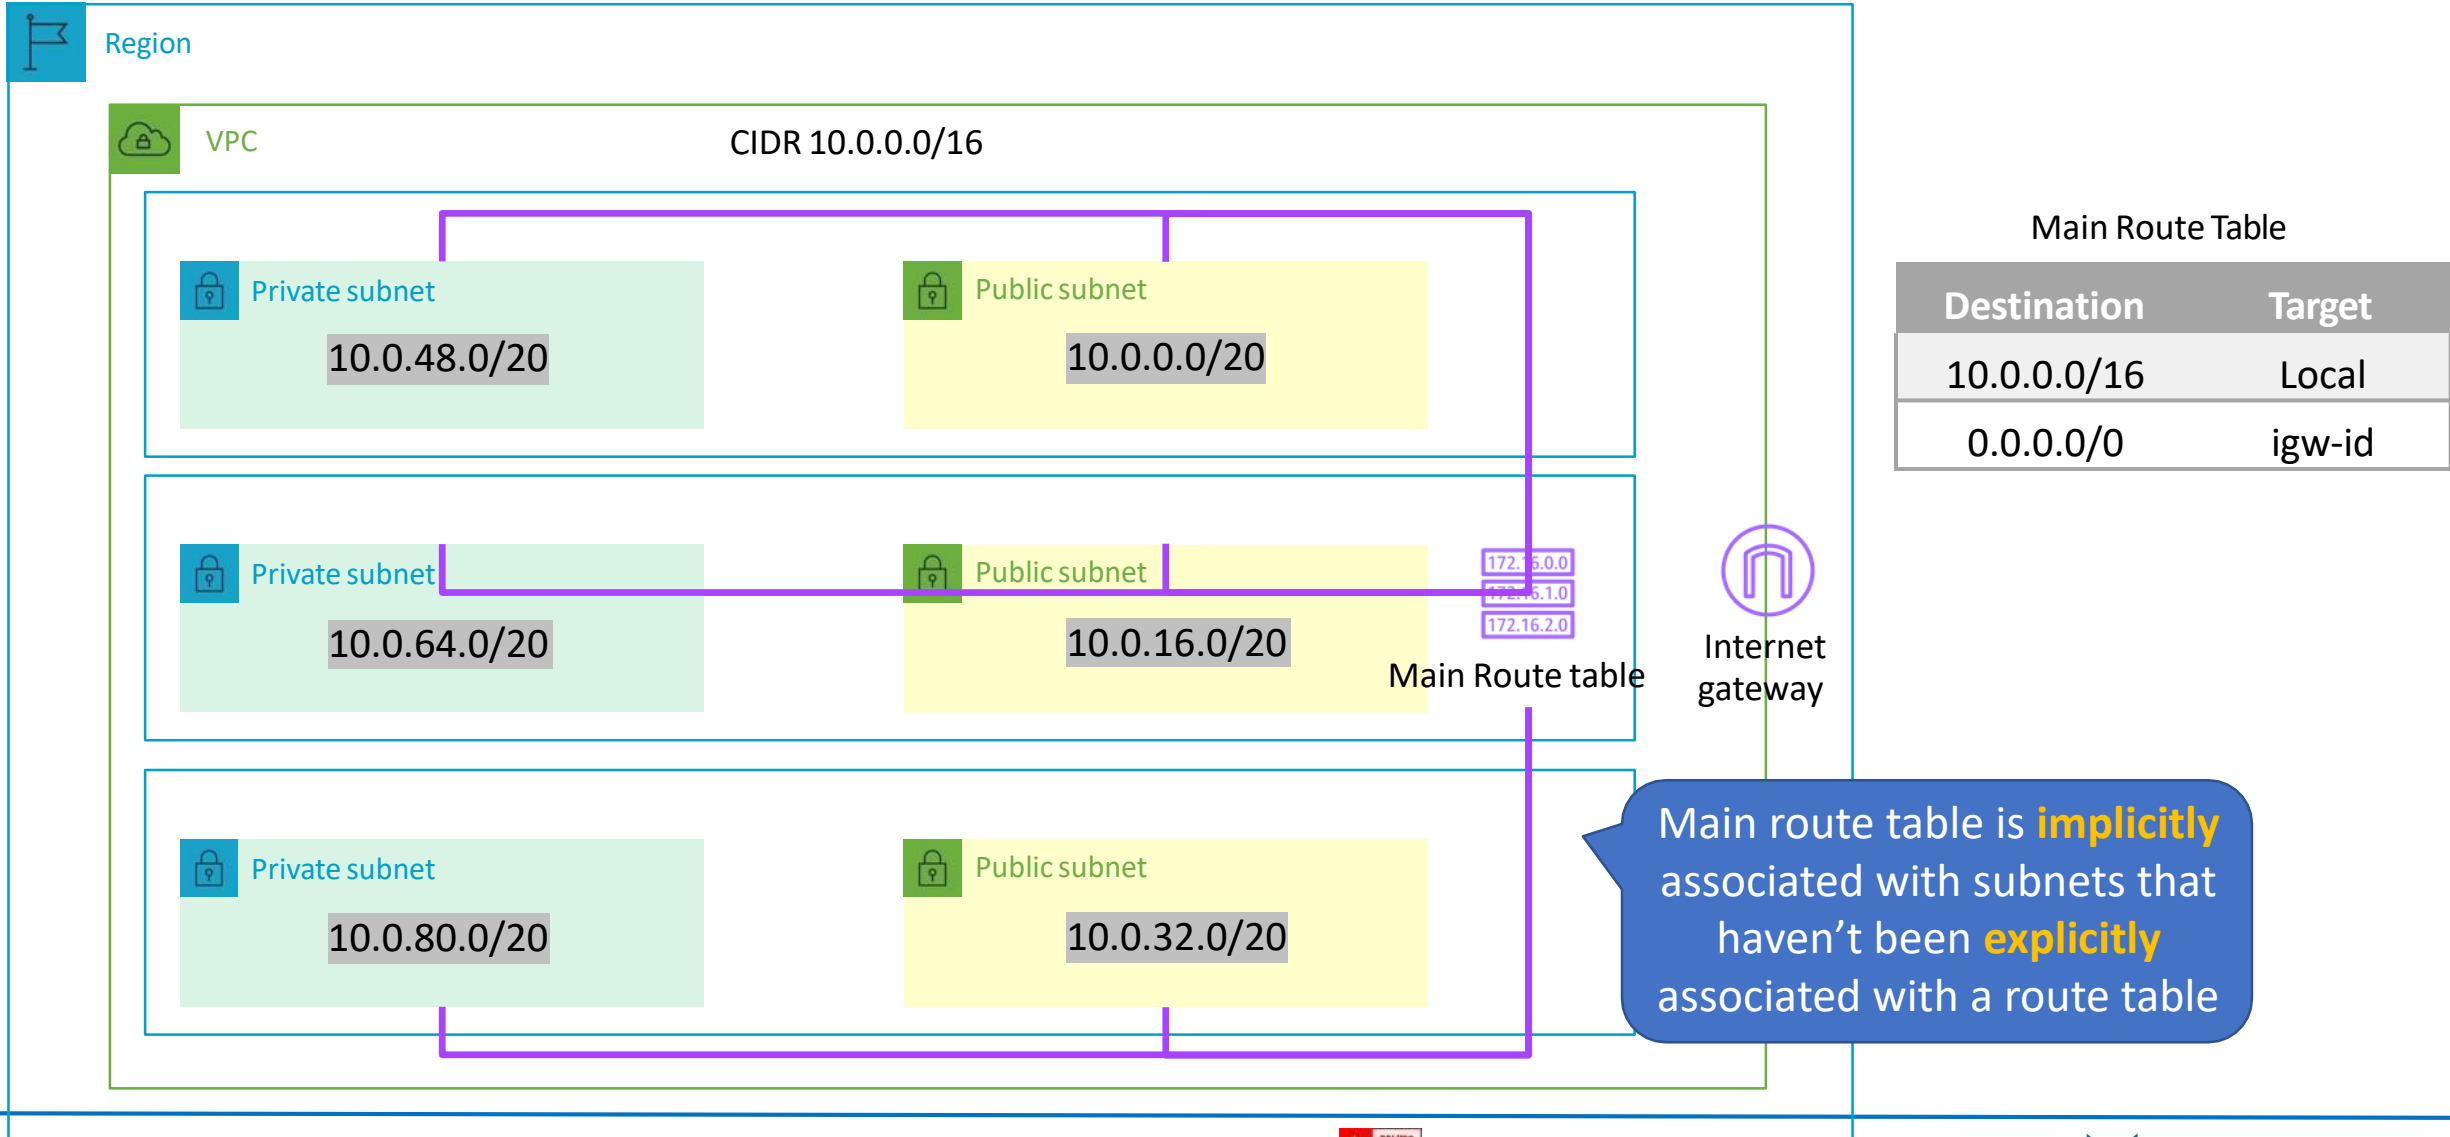
Infrastructure Security



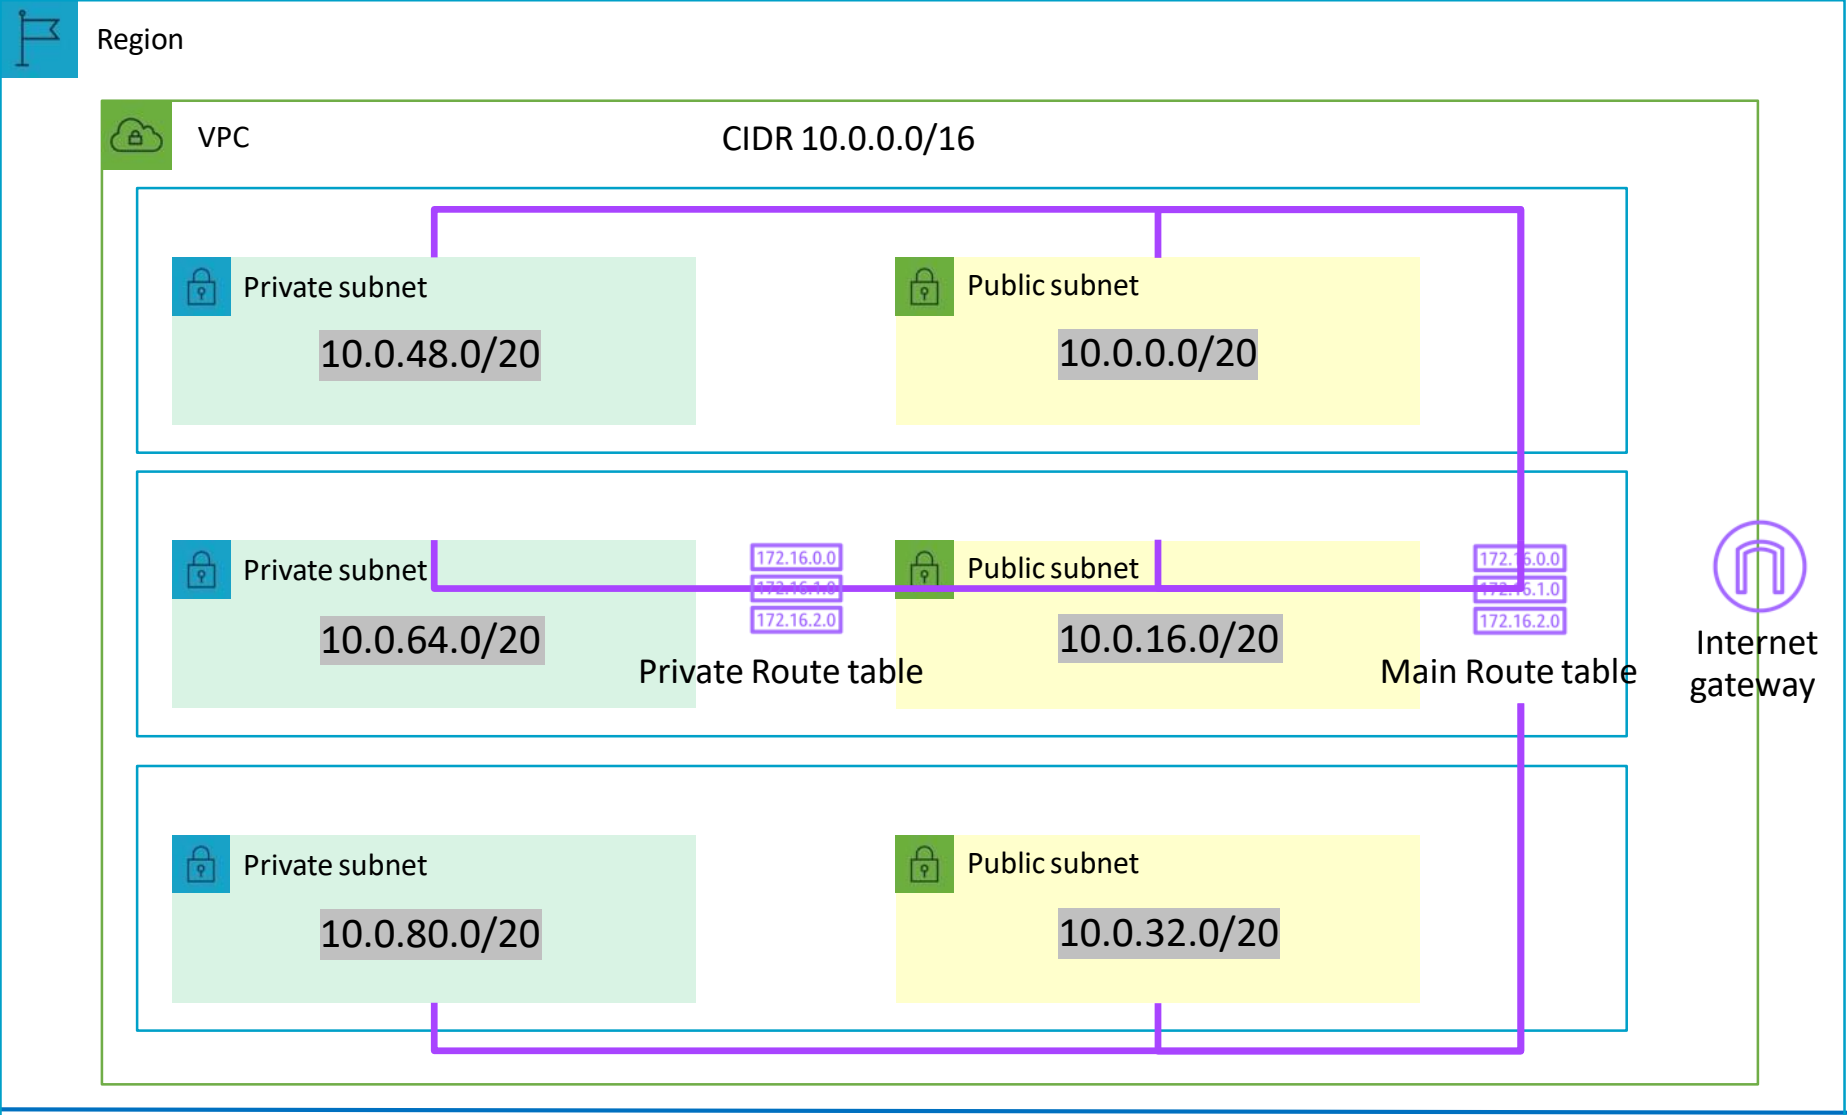
Create a Custom VPC



VPC Routing Deep Dive



VPC Routing Deep Dive



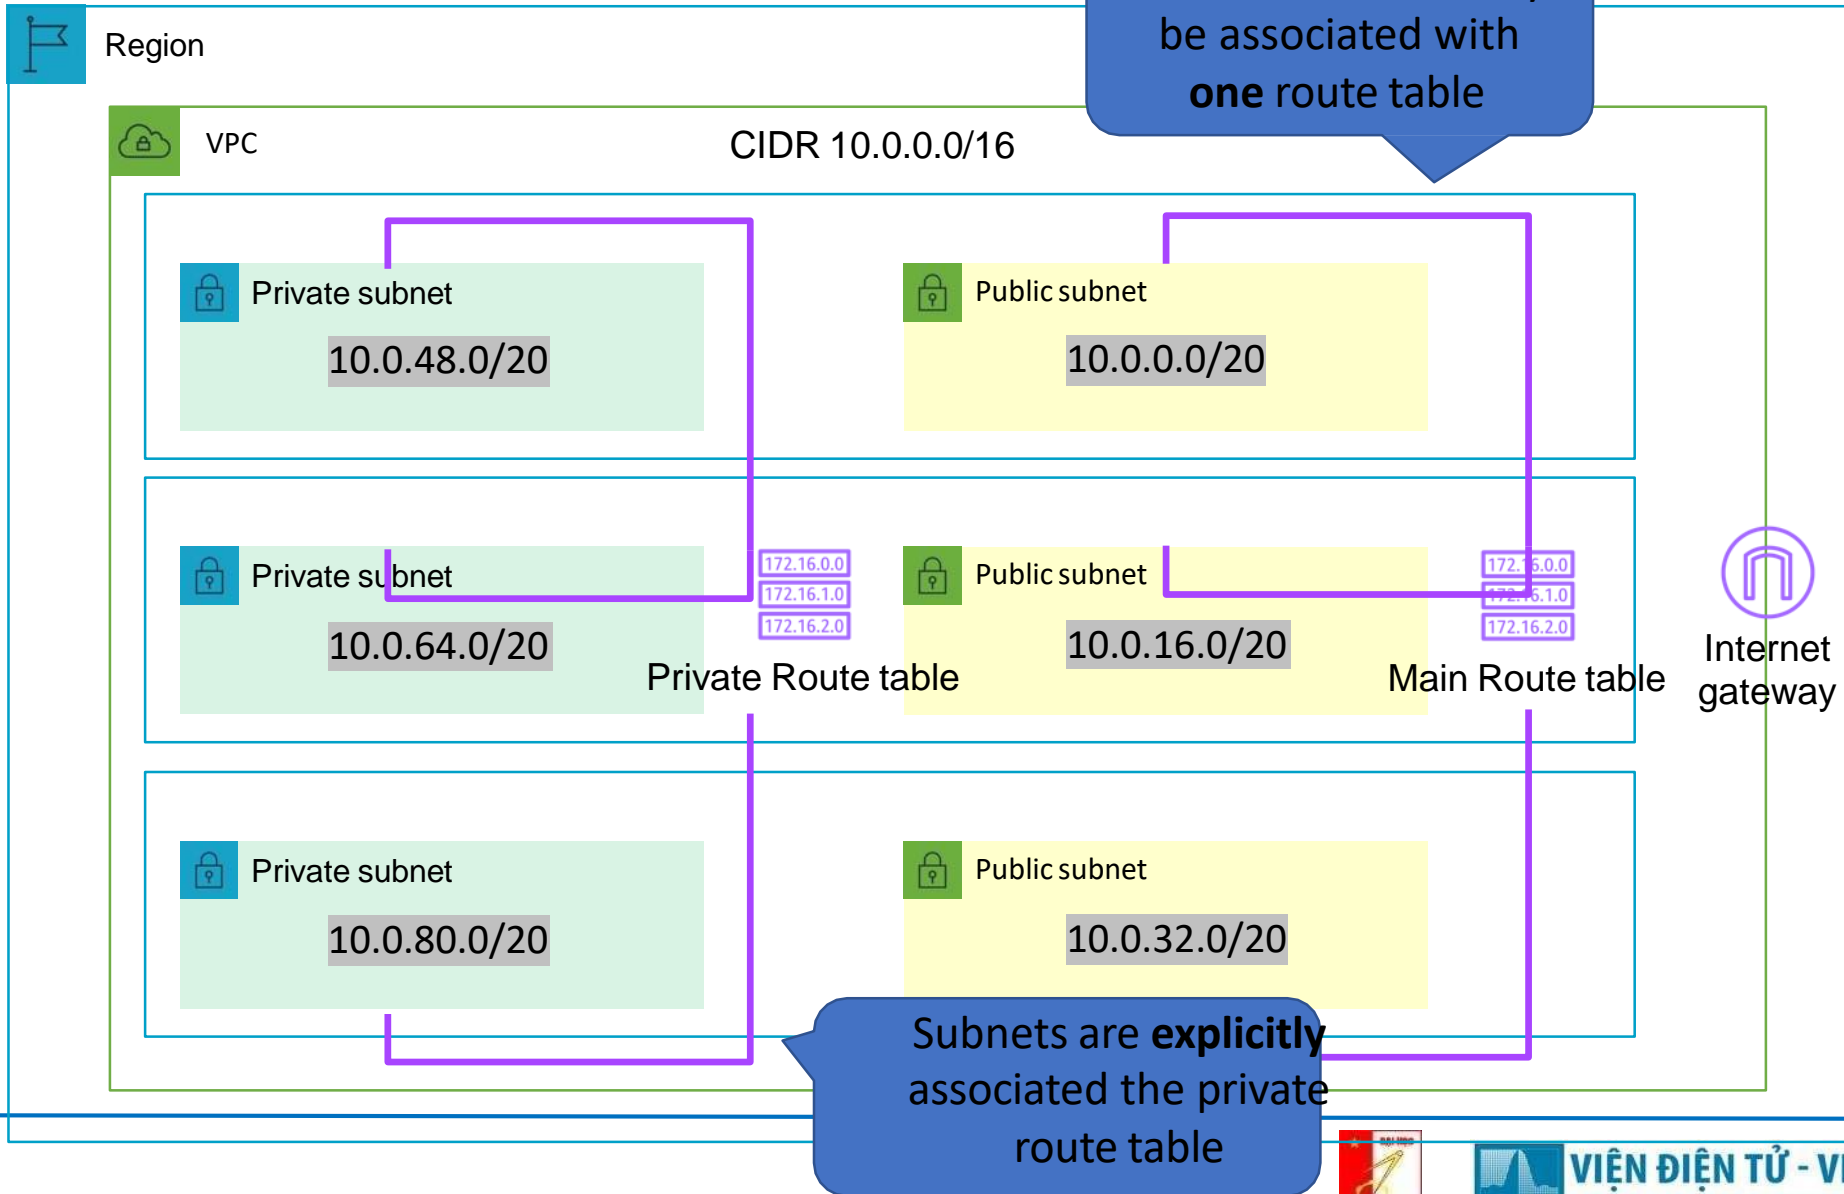
Main Route Table

Destination	Target
10.0.0.0/16	Local
0.0.0.0/0	igw-id

Private Route Table

Destination	Target
10.0.0.0/16	Local

VPC Routing Deep Dive



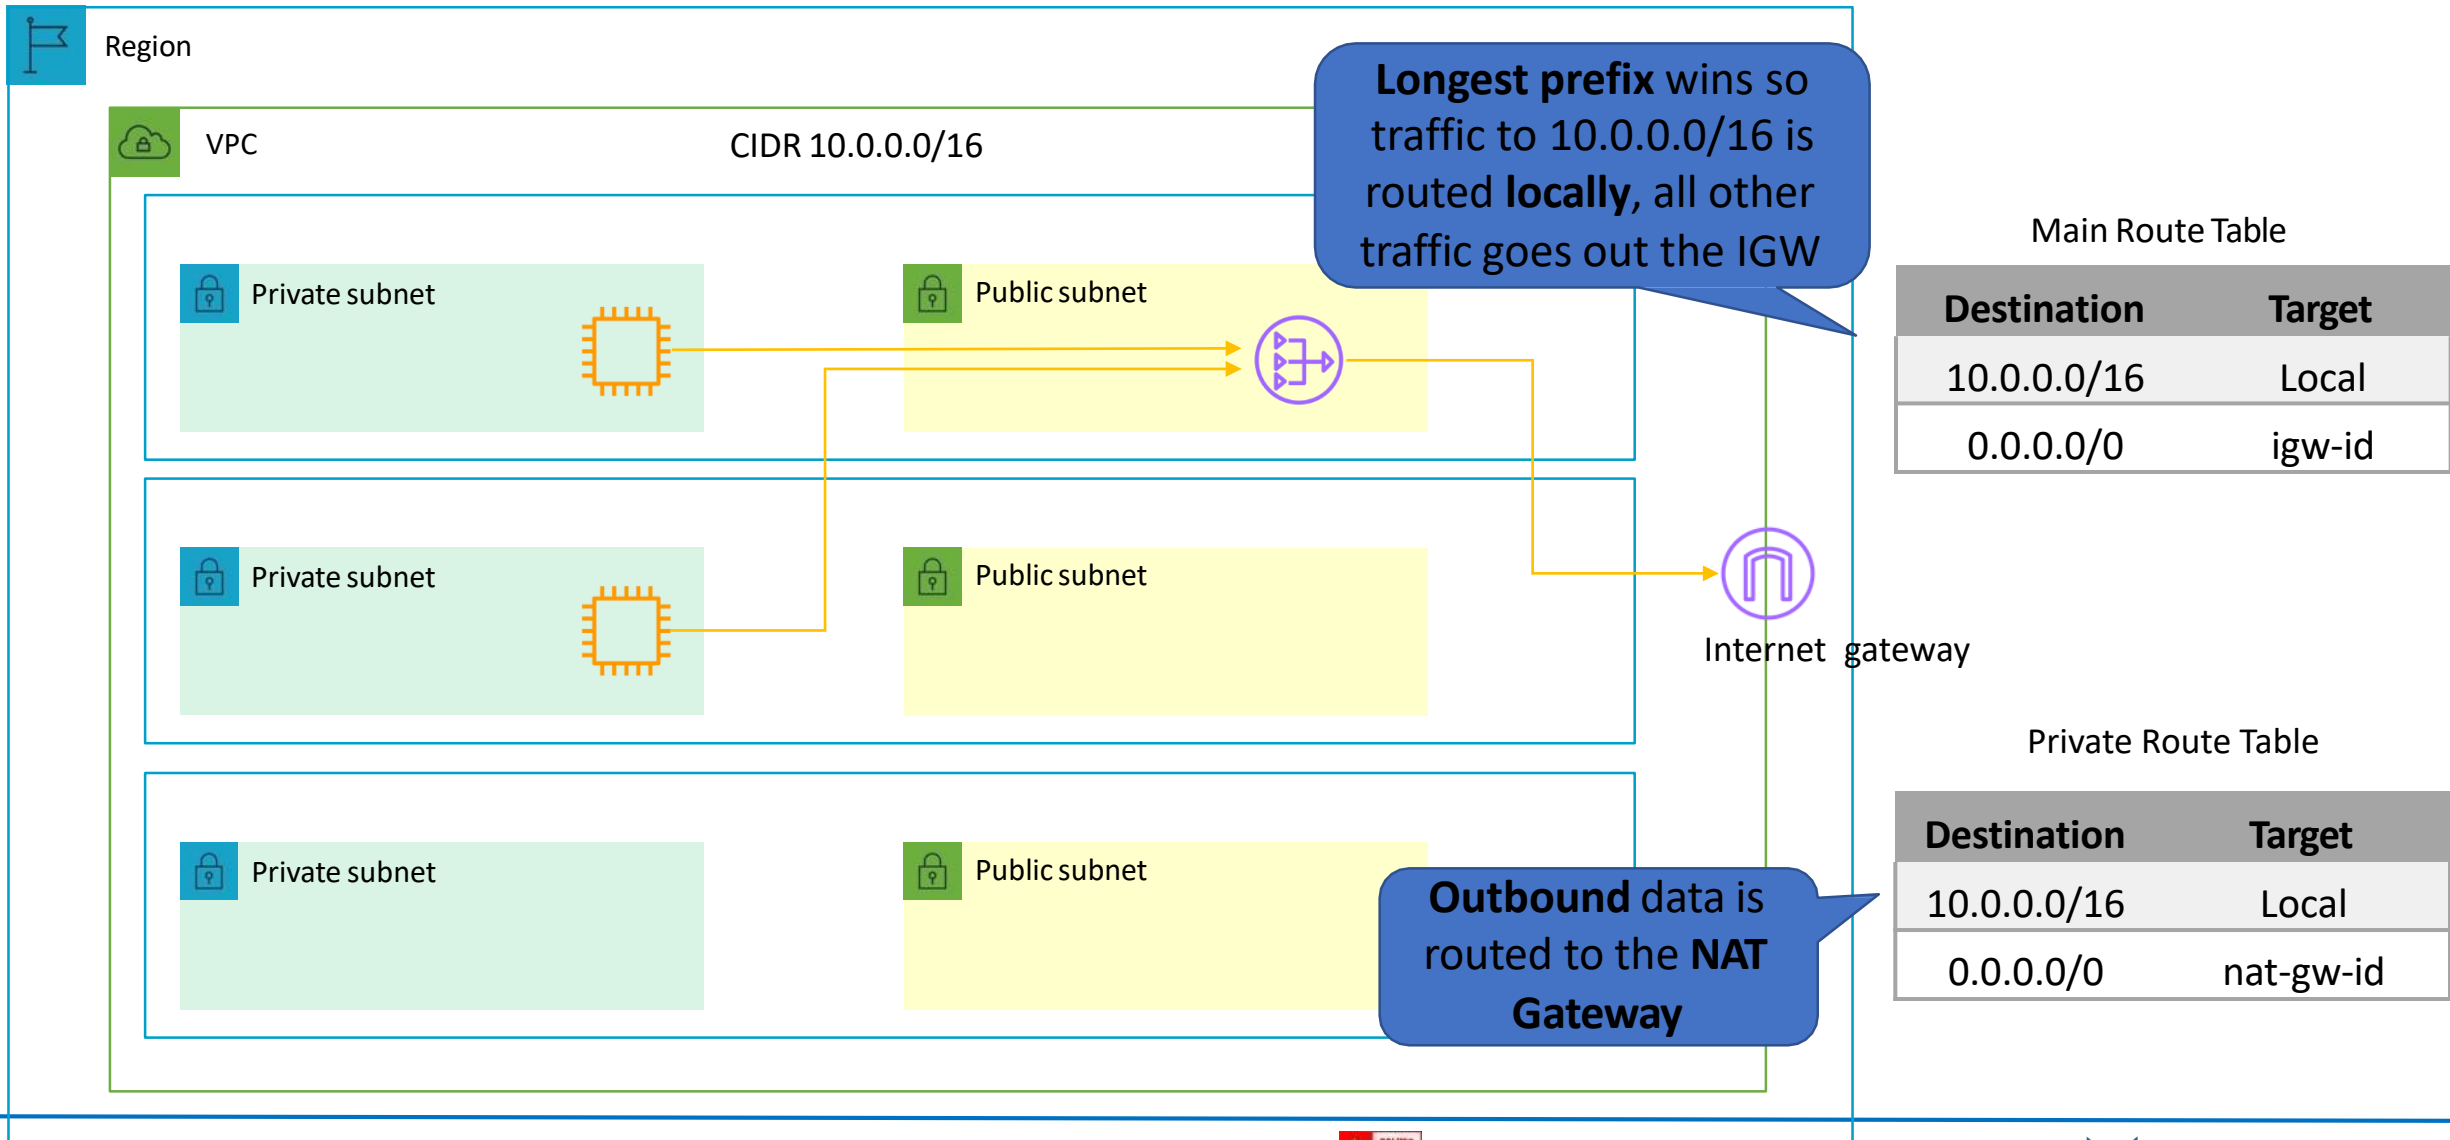
Main Route Table

Destination	Target
10.0.0.0/16	Local
0.0.0.0/0	igw-id

Private Route Table

Destination	Target
10.0.0.0/16	Local


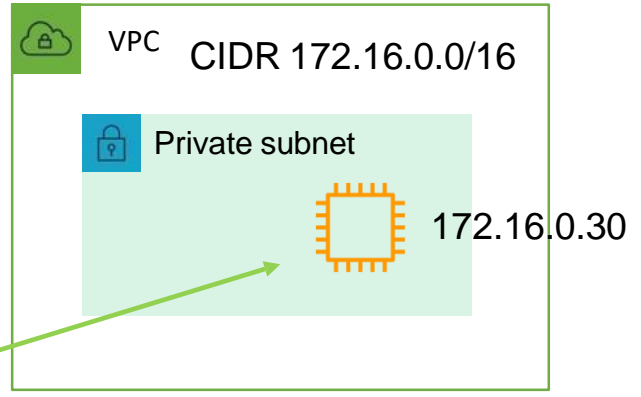

VPC Routing Deep Dive



VPC Routing Deep Dive

Longest prefix wins so all 172.16.0.0 traffic goes via **peer 1** except traffic to 172.16.0.15 which goes via **peer 2**

Destination	Target
10.0.0.0/16	Local
172.16.0.0/16	vpc-peer-1
172.16.0.15/32	vpc-peer-2



172.16.0.0.30 ->

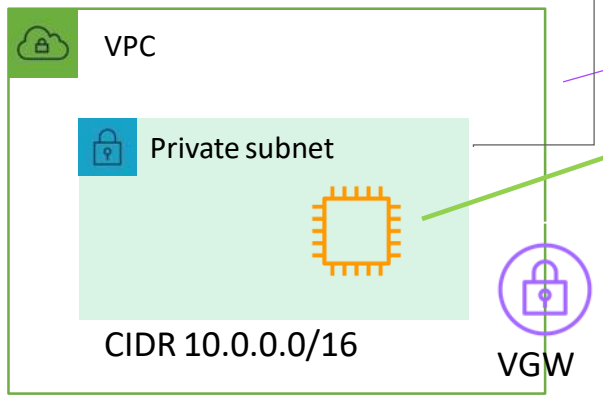
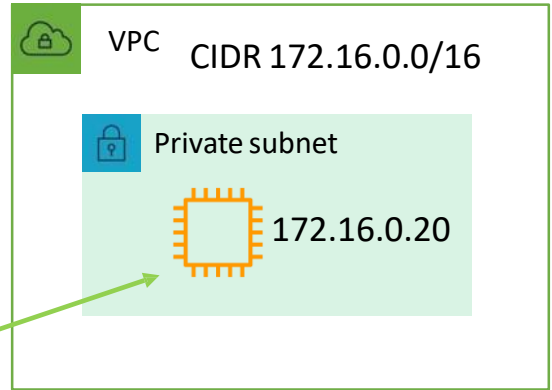
172.16.0.0.15 ->



VPC Routing Deep Dive


Static routes are preferred over propagated routes

Destination	Target
10.0.0.0/16	Local
172.16.0.0/16	vpc-peer-1
172.16.0.0/16	vgw-conn-1



Traffic to 172.16.0.20 gets routed to EC2 instance

Routes learned and propagated by BGP to route table



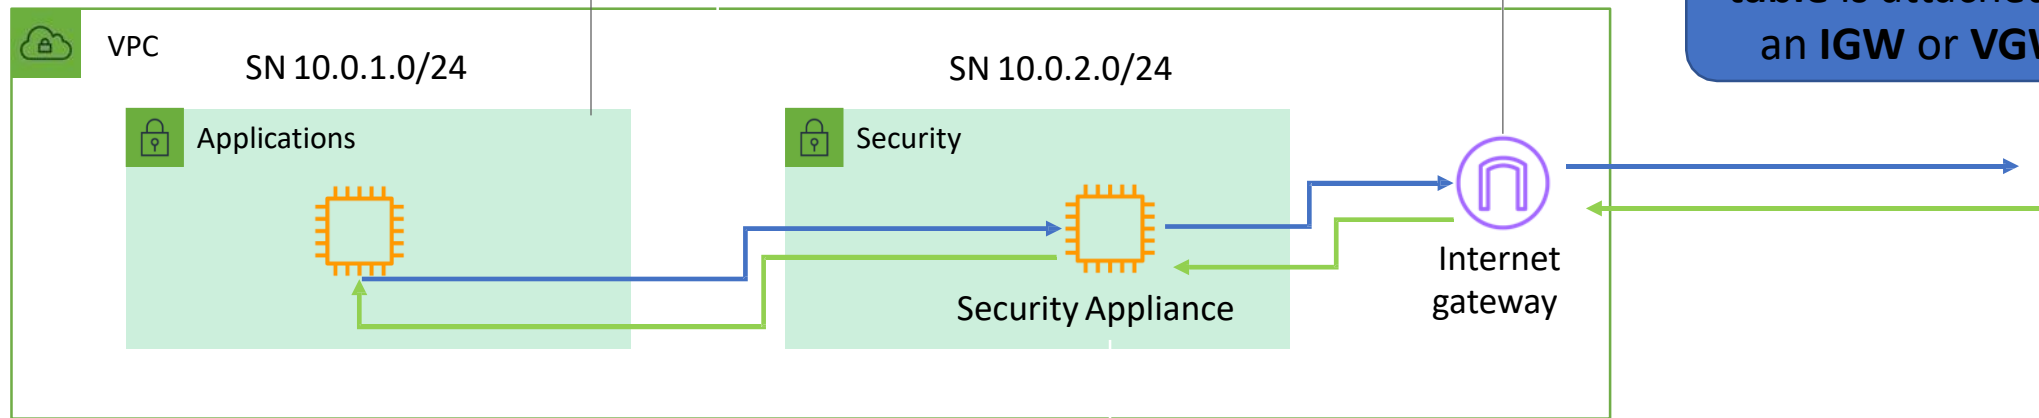
Gateway Route Tables

0.0.0.0/0 points to the ENI ID of the security appliance

Destination	Target
10.0.0.0/16	Local
0.0.0.0/0	eni-id-sec

Destination	Target
10.0.0.0/16	Local
10.0.1.0/24	eni-id-sec

A Gateway route table is attached to an IGW or VGW



All **outbound** traffic forwarded to **IGW**

Destination	Target
10.0.0.0/16	Local
0.0.0.0/0	igw-id

IPv4 and IPv6 Routing

Destination	Target
10.0.0.0/16	Local
2001:db8:1234:1a00::/56	Local
172.31.0.0/16	pcx-11223344556677889
0.0.0.0/0	igw-12345678901234567
::/0	eigw-aabbccdde1122334

IPv6 traffic within the VPC is routed **locally**

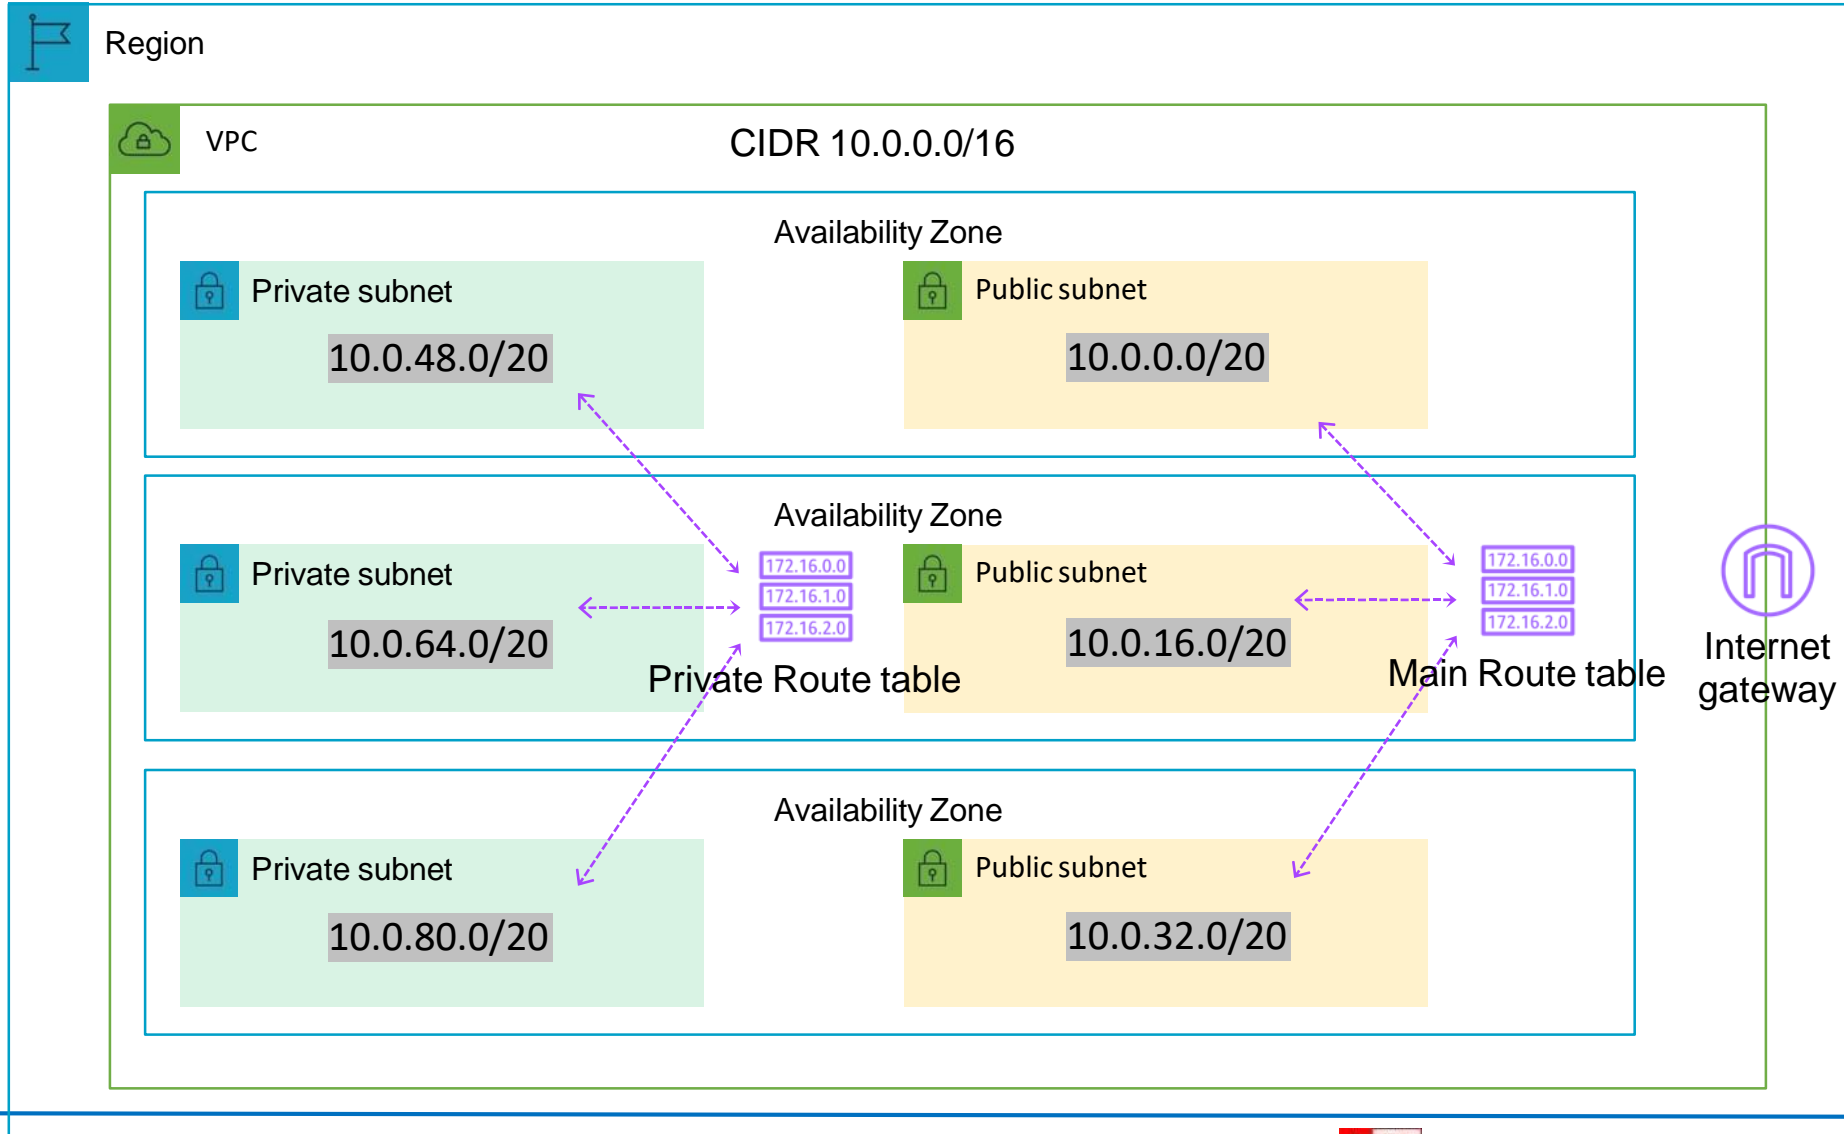
IPv4 traffic within the VPC is routed **locally**

Traffic that doesn't match a more specific route goes via the **IGW**

IPv4 traffic for 172.31.0.0/16 network goes via a **peering connection**

IPv6 traffic that doesn't match a more specific route goes via the **EIGW**

Create a Custom VPC - Configure Routing



Main Route Table

Destination	Target
10.0.0.0/16	Local
0.0.0.0/0	igw-id

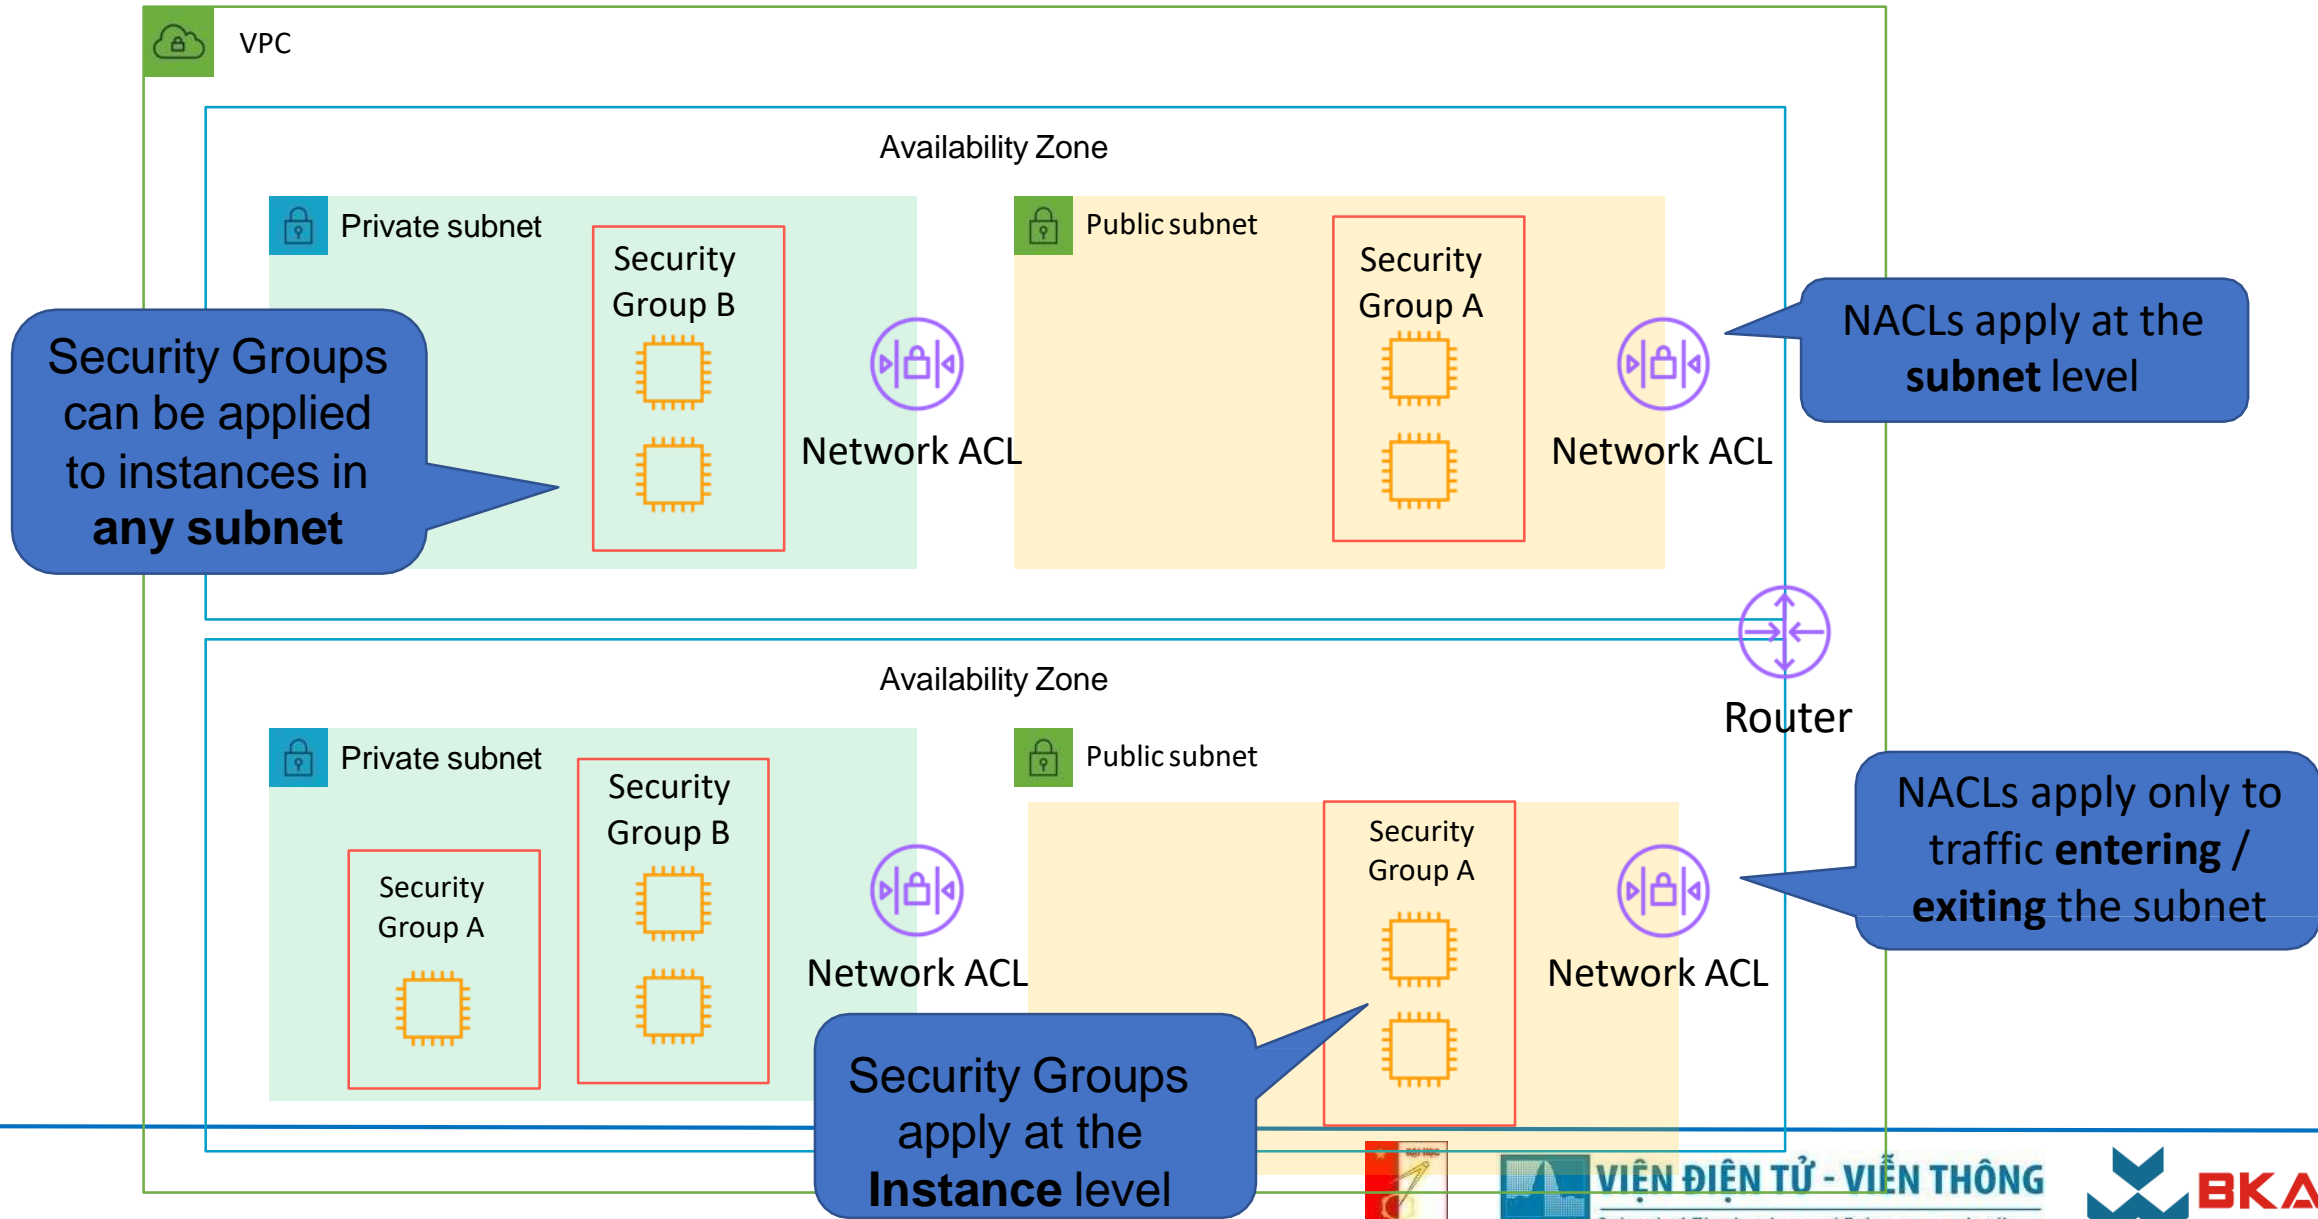
Private Route Table

Destination	Target
10.0.0.0/16	Local

Security Groups and Network ACLs

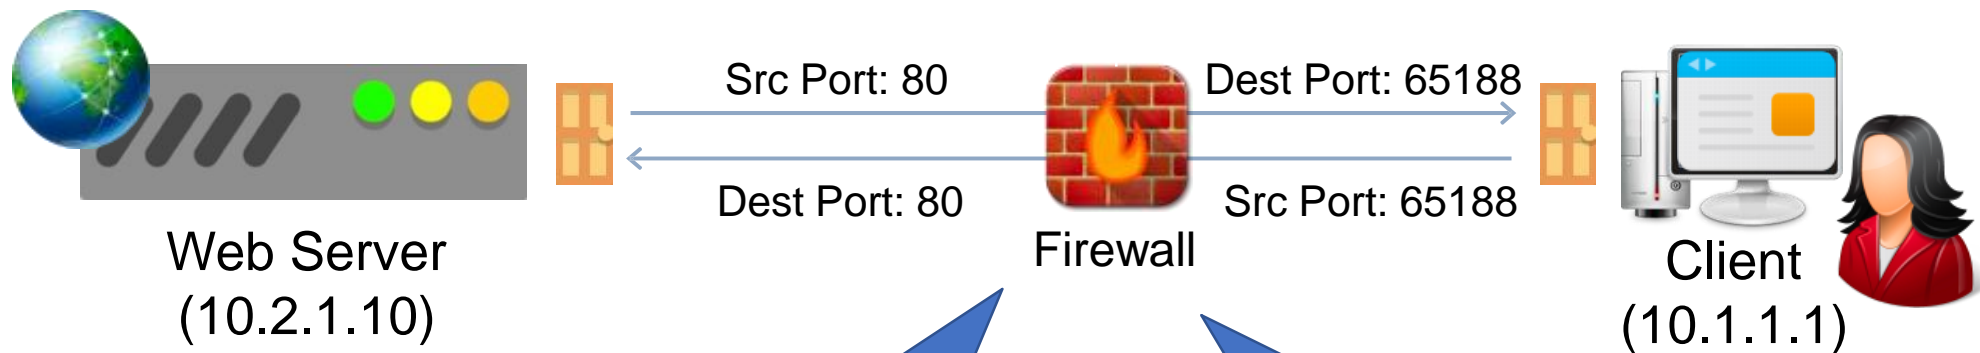


Security Groups and Network ACLs



Stateful vs Stateless Firewalls

PROTOCOL	SOURCE IP	DESTINATION IP	SOURCE PORT	DESTINATION PORT
HTTP	10.1.1.1	10.2.1.10	65188	80
HTTP	10.2.1.10	10.1.1.1	80	65188



A stateful firewall allows the return traffic automatically

A stateless firewall checks for an allow rule for both connections

Security Group Rules

Security groups support **allow** rules only


Inbound rules

Separate rules are defined for outbound traffic

Type	Protocol	Port range	Source
SSH	TCP	22	0.0.0.0/0
RDP	TCP	3389	0.0.0.0/0
RDP	TCP	3389	::/0
HTTPS	TCP	443	0.0.0.0/0
HTTPS	TCP	443	::/0
All ICMP - IPv4	ICMP	All	0.0.0.0/0

A source can be an **IP address** or **security group ID**


Security Groups Best Practice

 Public subnet(s)

Security group – **PublicALB**

Inbound: Protocol/Port HTTP/80 Source: 0.0.0.0/0

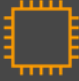
Outbound: Protocol/Port HTTPS:80 Destination: **PublicEC2**



Internet-facing ALB

Security group – **PublicEC2**

Inbound: Protocol/Port HTTP/80 Source: **PublicALB**


Outbound: Protocol/Port HTTPS/8080 Destination: **PrivateALB**


Web Front-End

 Private subnet(s)

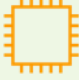
Security group – **PrivateALB**

Inbound: Protocol/Port HTTP/8080 Source: **PublicEC2** Outbound: Protocol/Port HTTPS/8080 Destination: **PrivateEC2**


Internal ALB

Security group – **PrivateEC2**

Inbound: Protocol/Port HTTP/8080 Source: **PrivateALB**


Application Layer

Network ACLs

Inbound Rules

Rule #	Type	Protocol	Port Range	Source	Allow / Deny
100	ALL Traffic	ALL	ALL	0.0.0.0/0	ALLOW
101	ALL Traffic	ALL	ALL	::/0	ALLOW
*	ALL Traffic	ALL	ALL	0.0.0.0/0	DENY
*	ALL Traffic	ALL	ALL	::/0	DENY

Outbound Rules

Rule #	Type	Protocol	Port Range	Destination	Allow / Deny
100	ALL Traffic	ALL	ALL	0.0.0.0/0	ALLOW
101	ALL Traffic	ALL	ALL	::/0	ALLOW
*	ALL Traffic	ALL	ALL	0.0.0.0/0	DENY
*	ALL Traffic	ALL	ALL	::/0	DENY

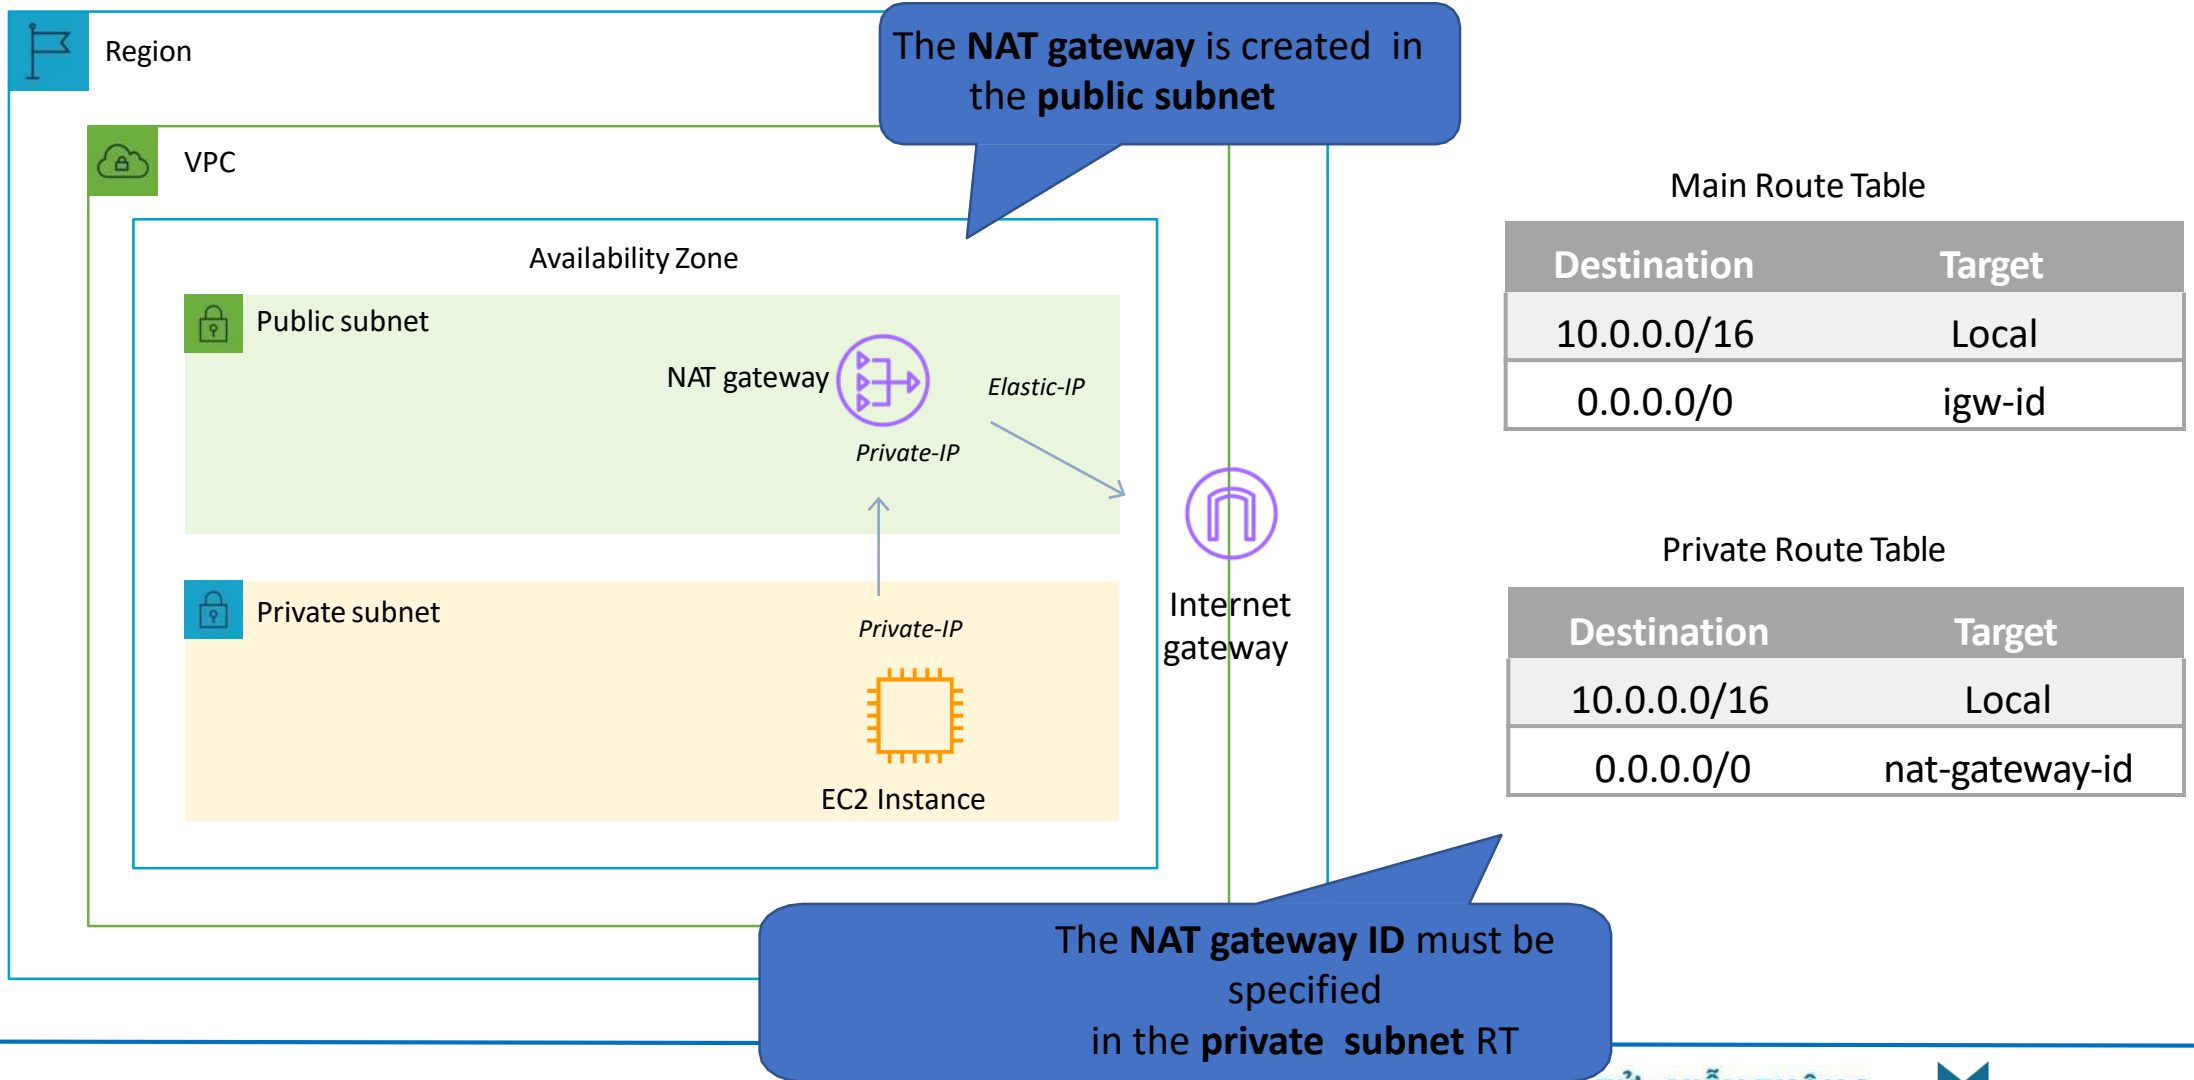
NACLs have an explicit deny

Rules are processed in order

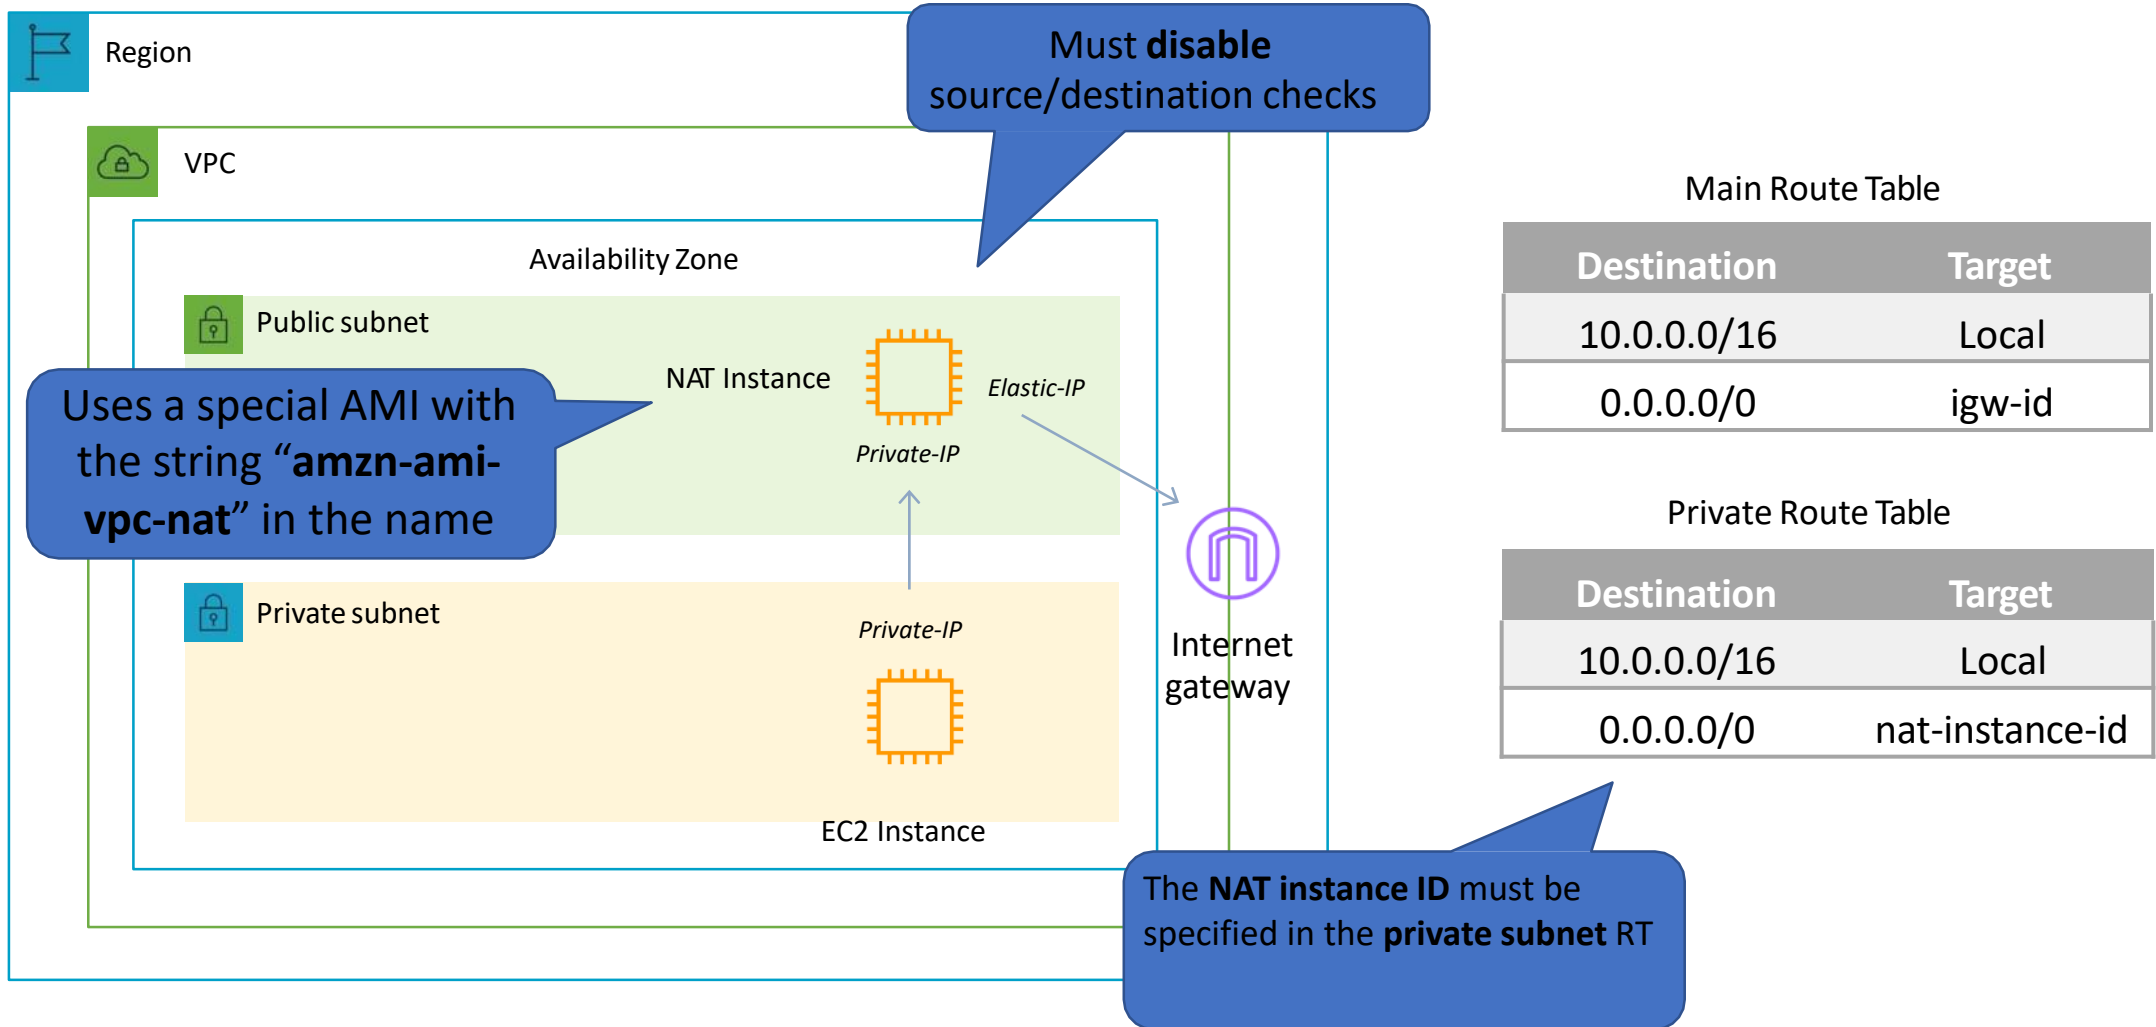
NAT Gateways and NAT Instances



NAT Gateways



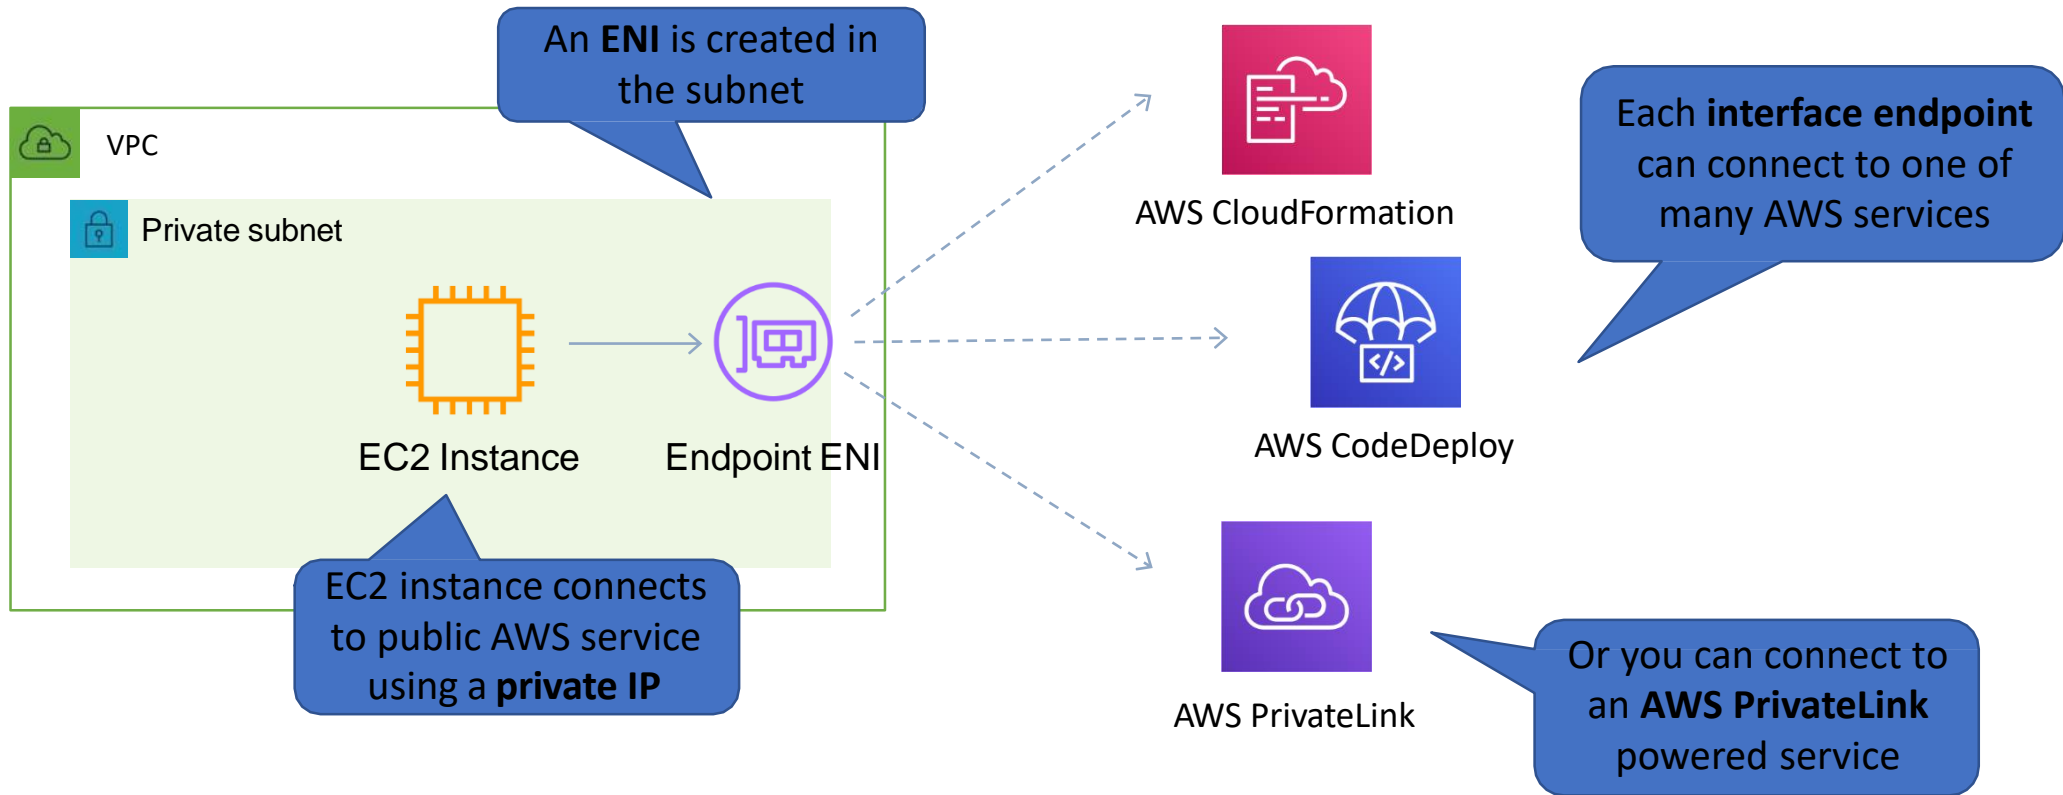
NAT Instances



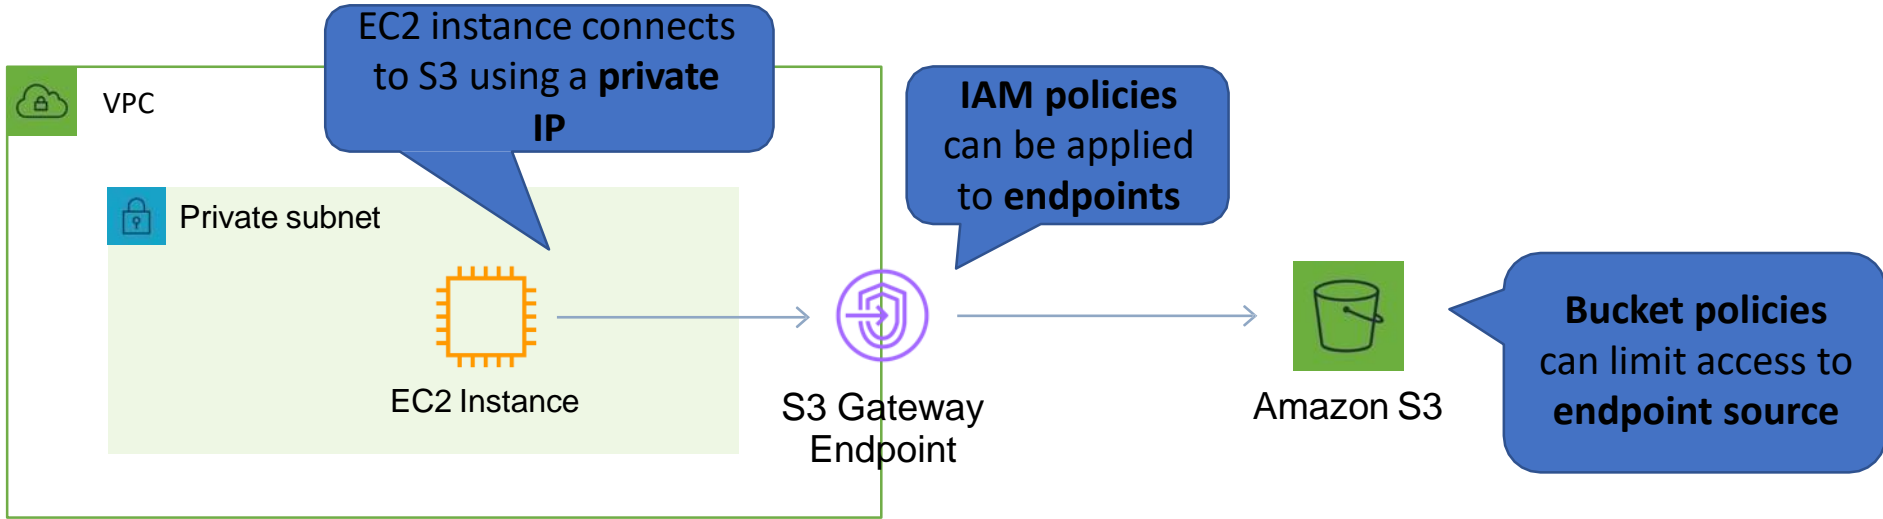
NAT Instance vs NAT Gateway

NAT Instance	NAT Gateway
Managed by you (e.g. software updates)	Managed by AWS
Scale up (instance type) manually and use enhanced networking	Elastic scalability up to 45 Gbps
No high availability – scripted/auto-scaled HA possible using multiple NATs in multiple subnets	Provides automatic high availability within an AZ and can be placed in multiple AZs
Need to assign Security Group	No Security Groups
Can use as a bastion host	Cannot access through SSH
Use an Elastic IP address or a public IP address with a NAT instance	Choose the Elastic IP address to associate with a NAT gateway at creation
Can implement port forwarding through manual customisation	Does not support port forwarding

VPC Interface Endpoints



VPC Gateway Endpoints



Route Table

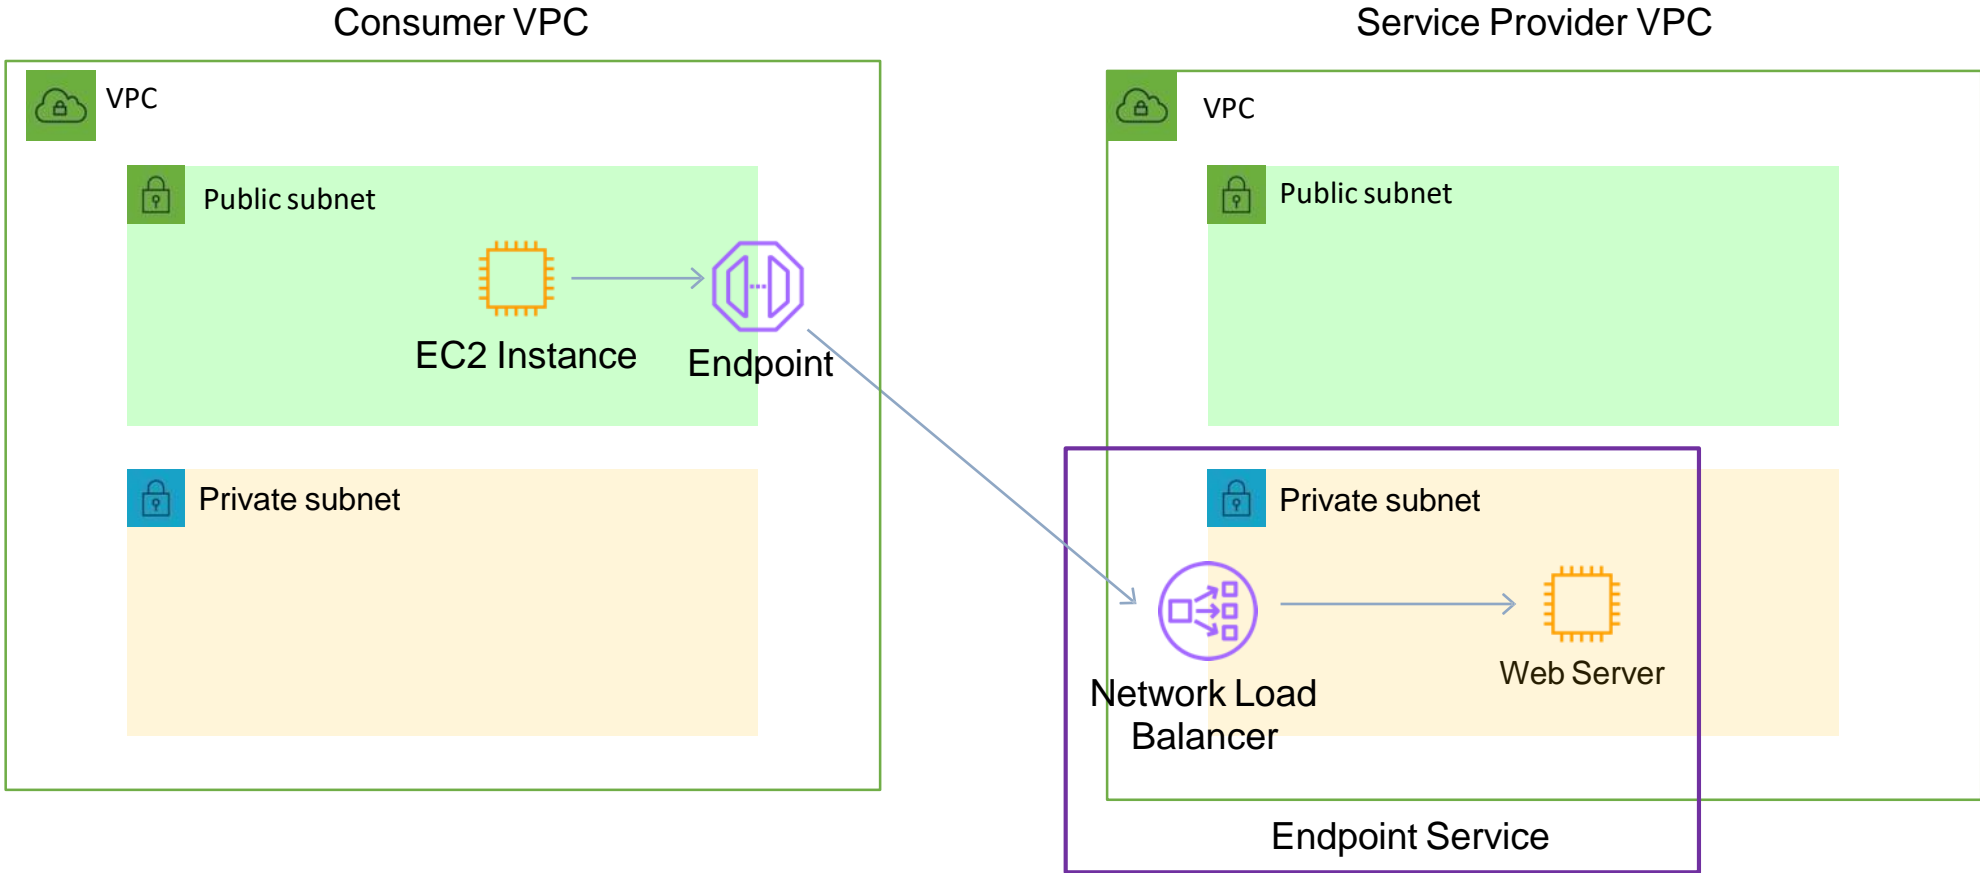
Destination	Target
<i>pl-6ca54005 (com.amazonaws.ap-southeast-2.s3, 54.231.248.0/22, 54.231.252.0/24, 52.95.128.0/21)</i>	<i>vpce-ID</i>

A route table entry is required with the prefix list for S3 and the gateway ID

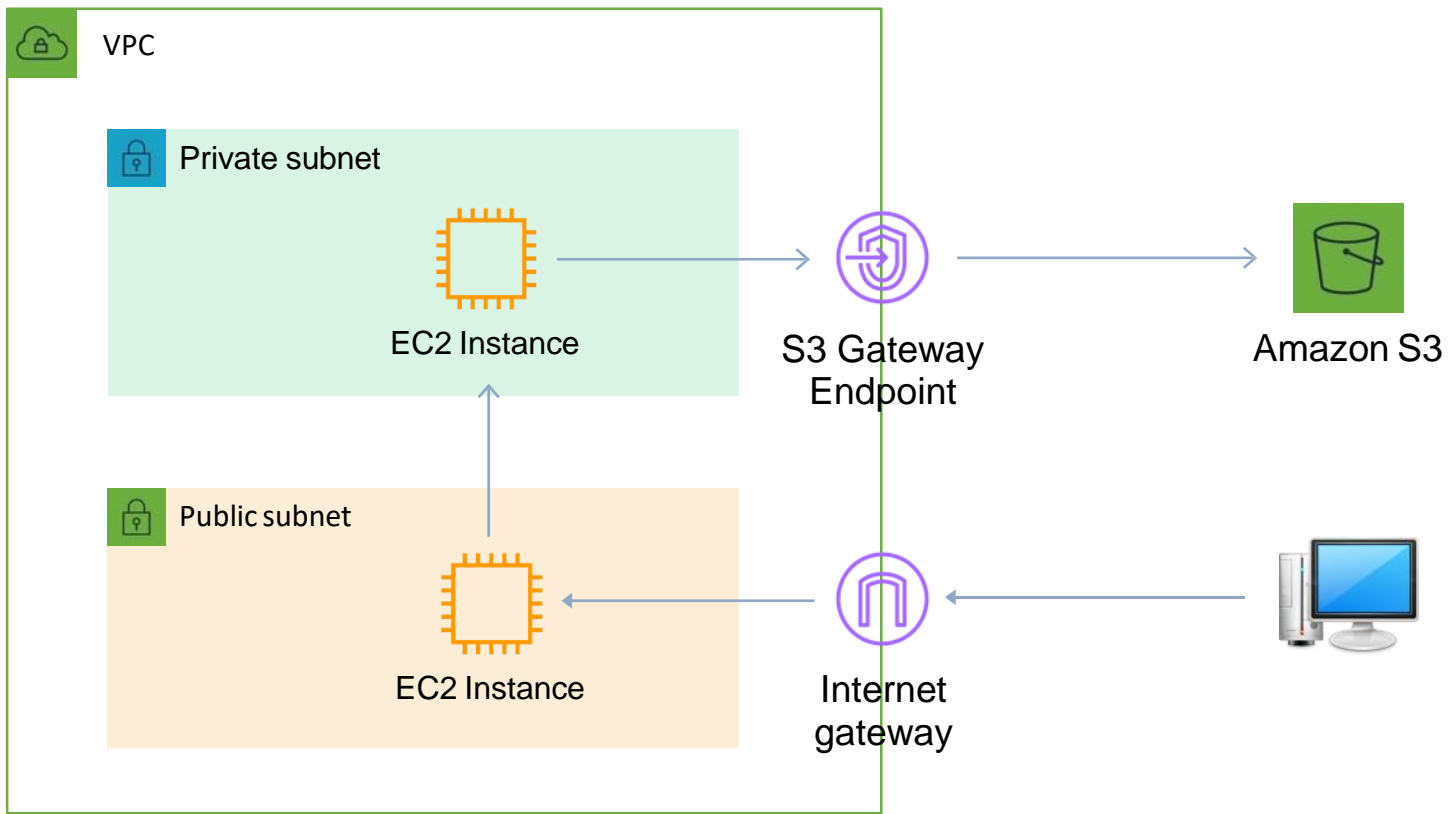
VPC Endpoints

	Interface Endpoint	Gateway Endpoint
What	Elastic Network Interface with a Private IP	A gateway that is a target for a specific route
How	Uses DNS entries to redirect traffic	Uses prefix lists in the route table to redirect traffic
Which services	API Gateway, CloudFormation, CloudWatch etc.	Amazon S3, DynamoDB
Security	Security Groups	VPC Endpoint Policies

Service Provider Model



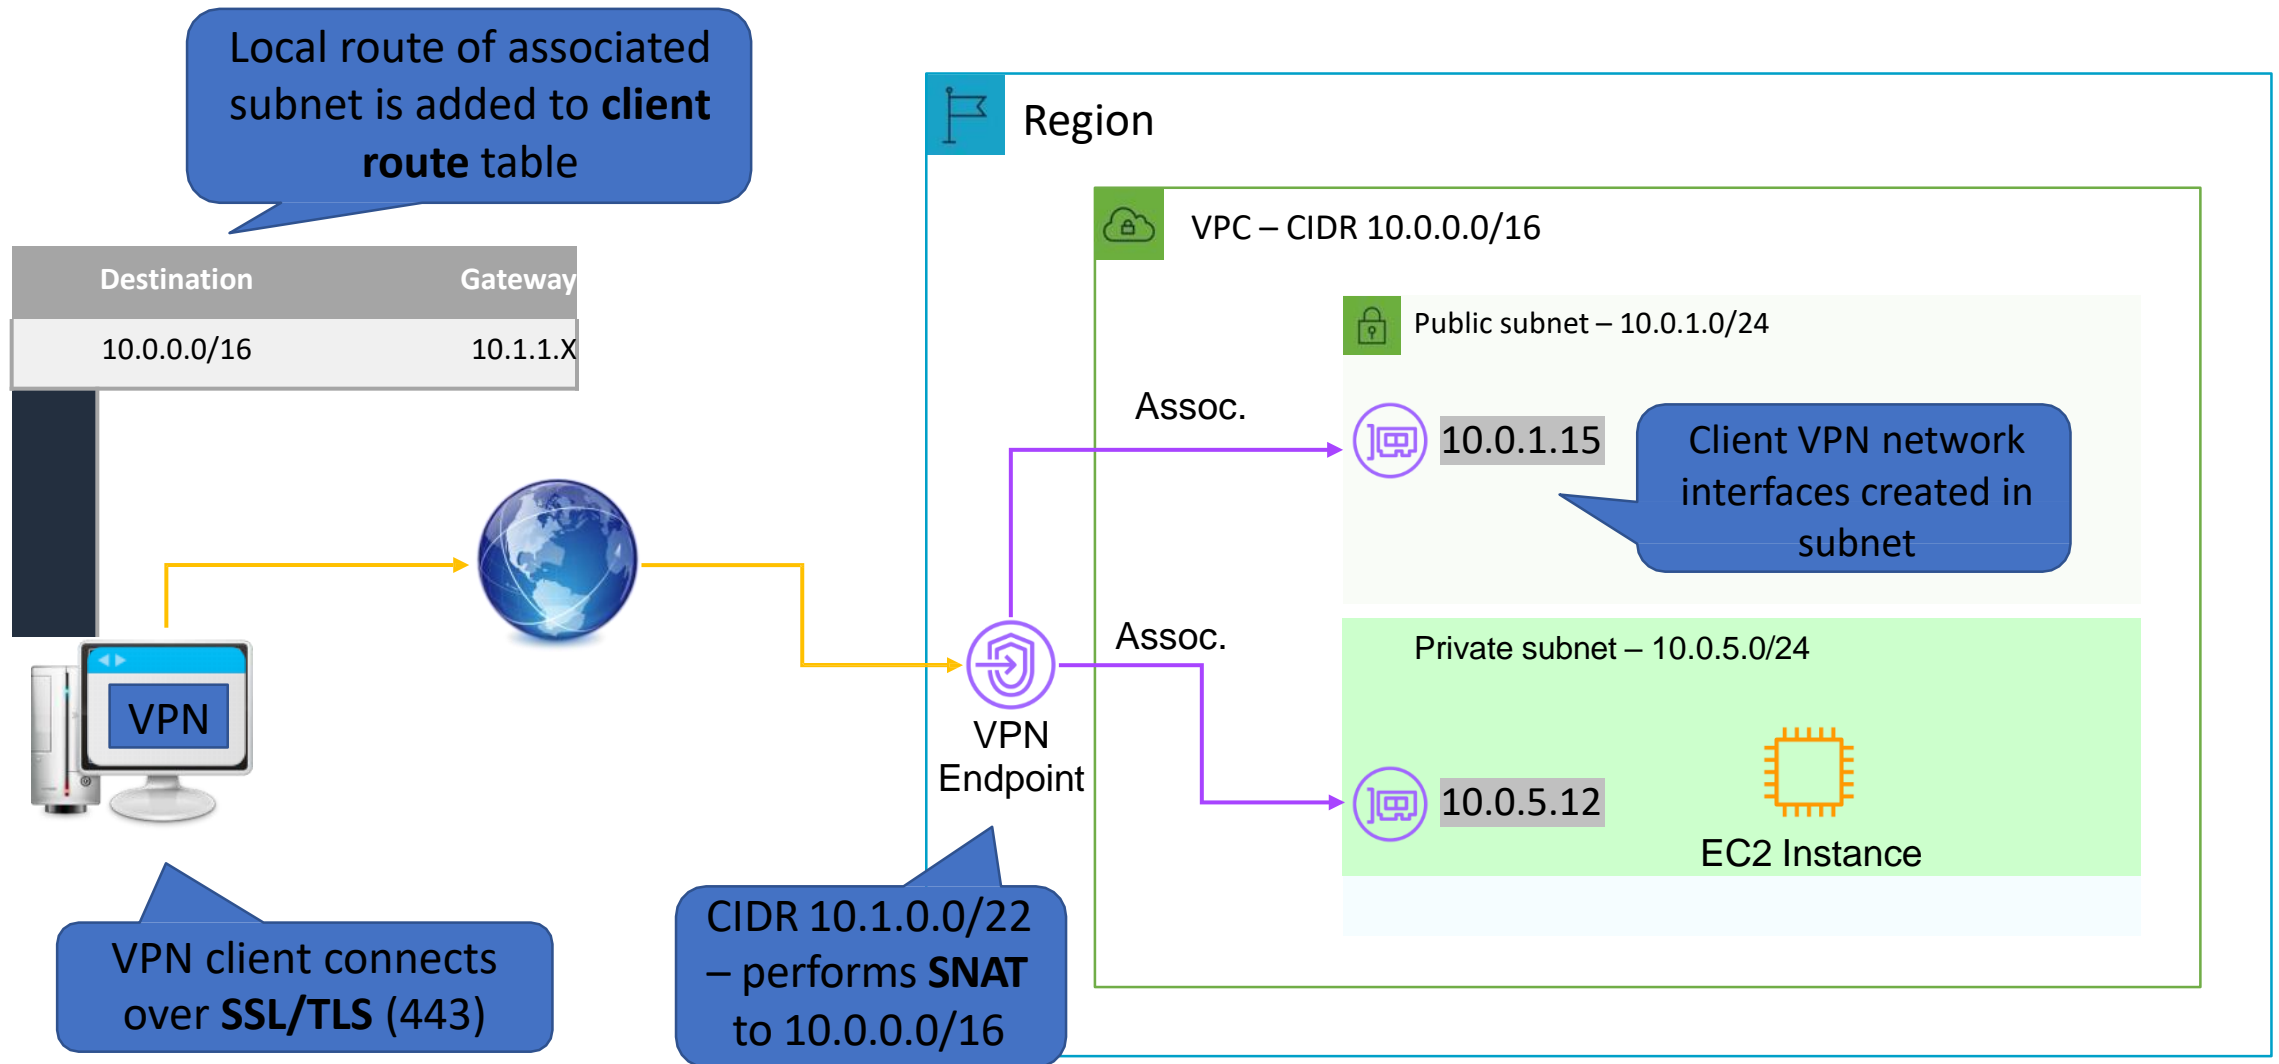
VPC Gateway Endpoints



Private Subnet Route Table

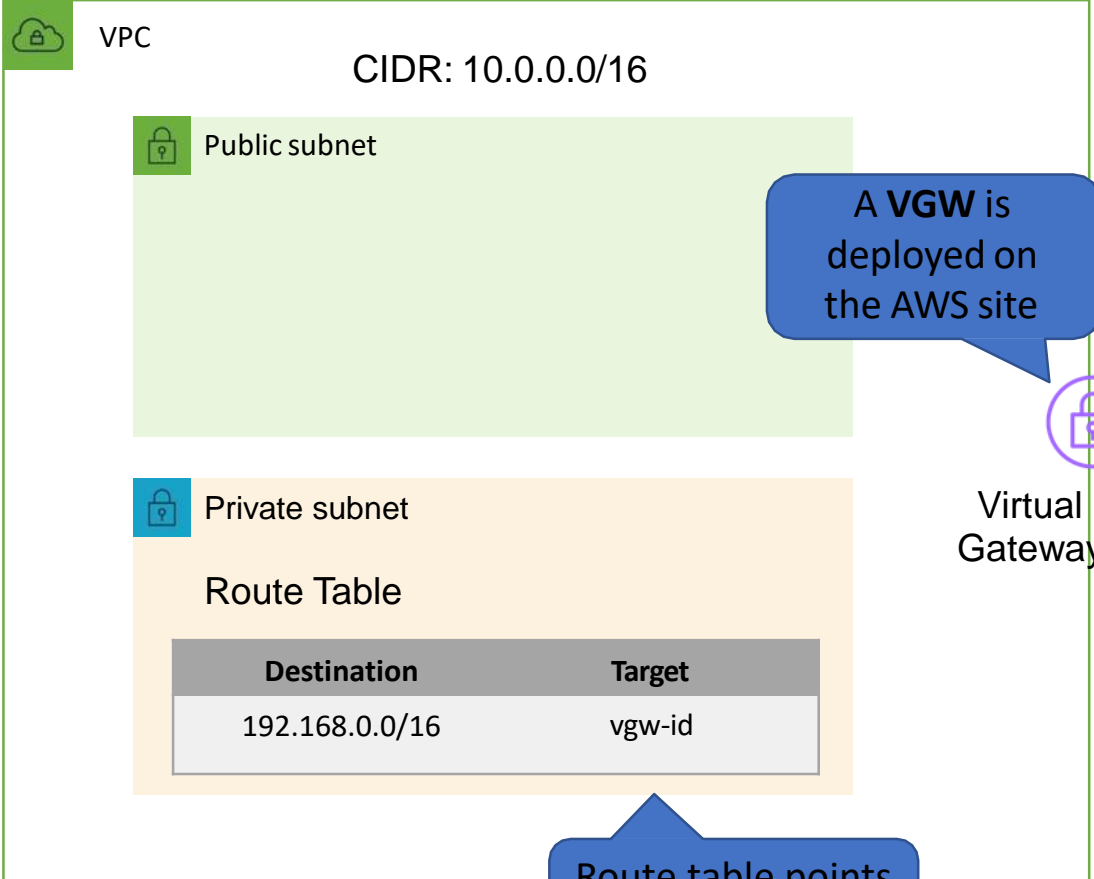
Destination	Target
pl-6ca54005 (com.amazonaws.ap-southeast-2.s3, 54.231.248.0/22, 54.231.252.0/24, 52.95.128.0/21)	vpce-ID

AWS Client VPN

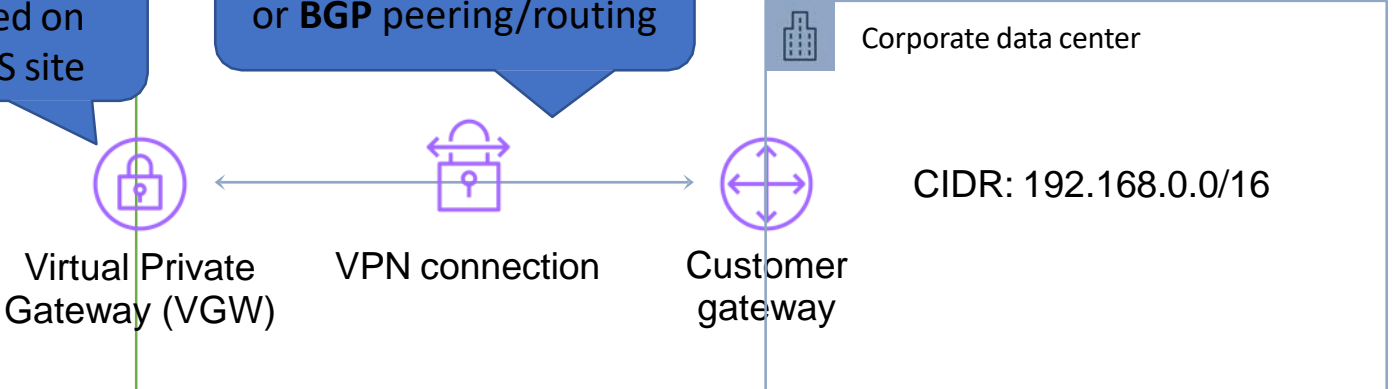


AWS Site-to-Site VPN

AWS VPN is a managed IPsec VPN



Supports **static routes** or **BGP peering/routing**



Route table points to the **VGW**

A **customer gateway** is deployed on the **customer side**

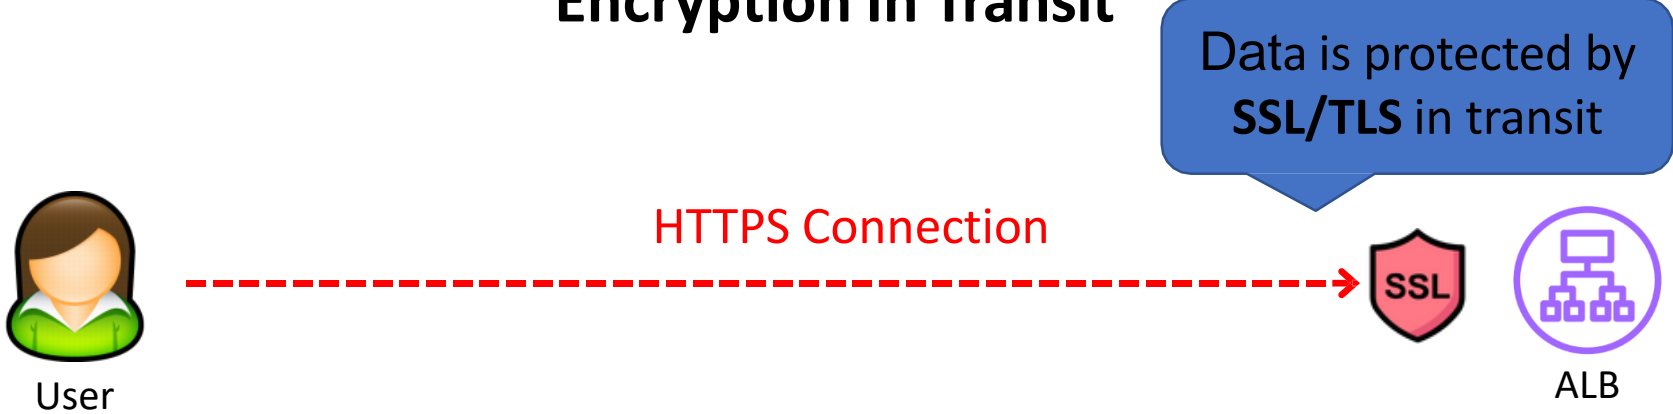


AWS Certificate Manager (ACM)

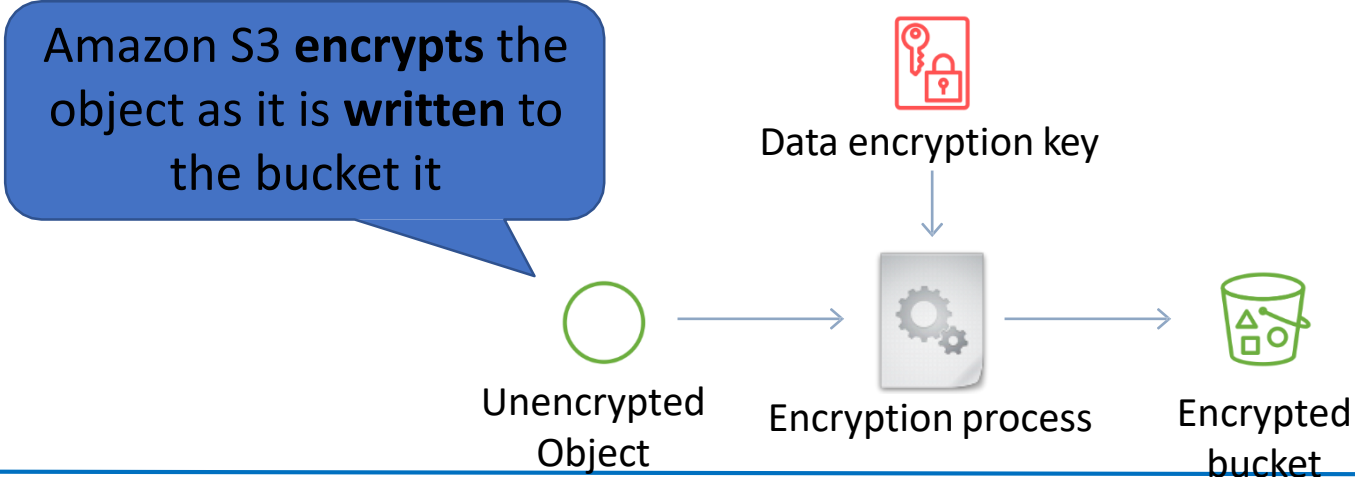


Encryption In Transit vs At Rest

Encryption In Transit




Encryption At Rest



Asymmetric Encryption

- Asymmetric encryption is also known as public key cryptography
- Messages encrypted with the public key can only be decrypted with the private key
- Messages encrypted with the private key can be decrypted with the public key
- Examples include SSL/TLS and SSH

Public key Private key



AWS Certificate Manager (ACM)

- Create, store and renew SSL/TLS X.509 certificates
- Single domains, multiple domain names and wildcards
- Integrates with several AWS services including:
 - **Elastic Load Balancing**
 - **Amazon CloudFront**
 - **AWS Elastic Beanstalk**
 - **AWS Nitro Enclaves**
 - **AWS CloudFormation**

AWS Certificate Manager (ACM)

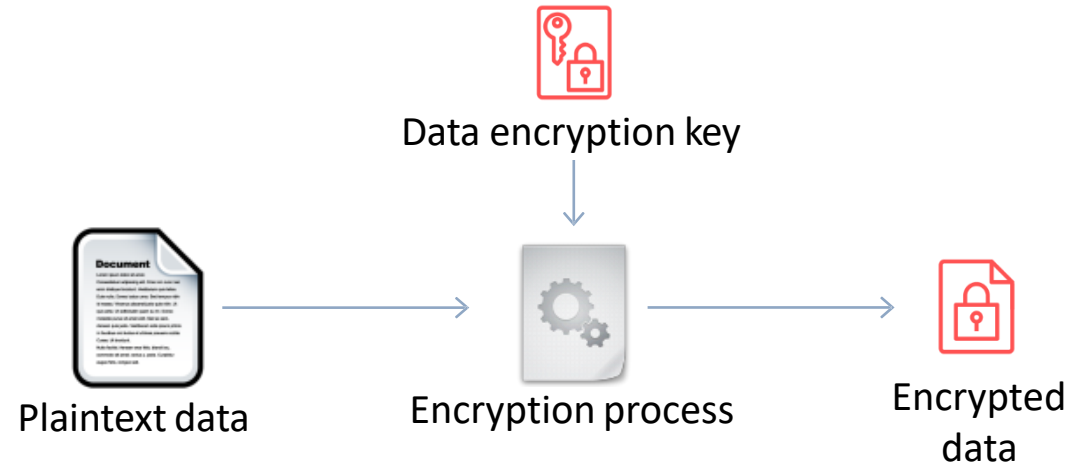
- **Public certificates** are signed by the AWS public Certificate Authority
- You can also create a Private CA with ACM
- Can then issue private certificates
- You can also import certificates from third-party issuers

AWS Key Management Service (KMS)

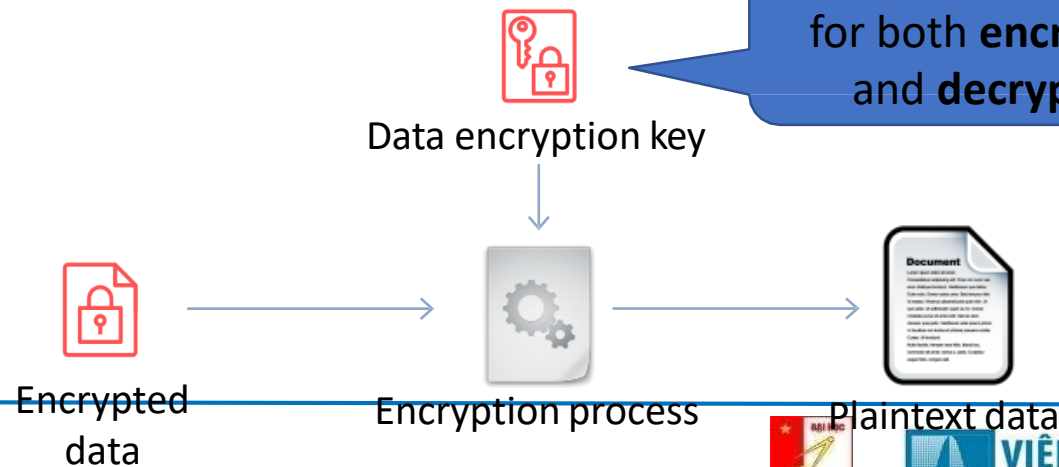


Symmetric Encryption

Encryption



Decryption



AWS Key Management Service (KMS)

- Create and managed **symmetric** and **asymmetric** encryption keys
- The **customer master keys** (CMKs) are protected by hardware security modules (HSMs)

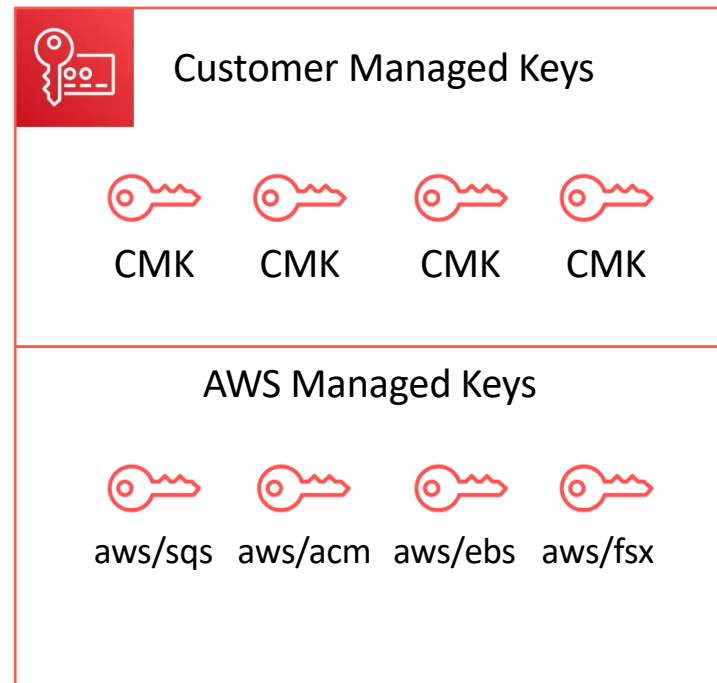
(HSMs)



Developer



AWS KMS



Developer creates **customer managed** customer master keys (CMKs) in AWS KMS

Customer Master Keys (CMKs)

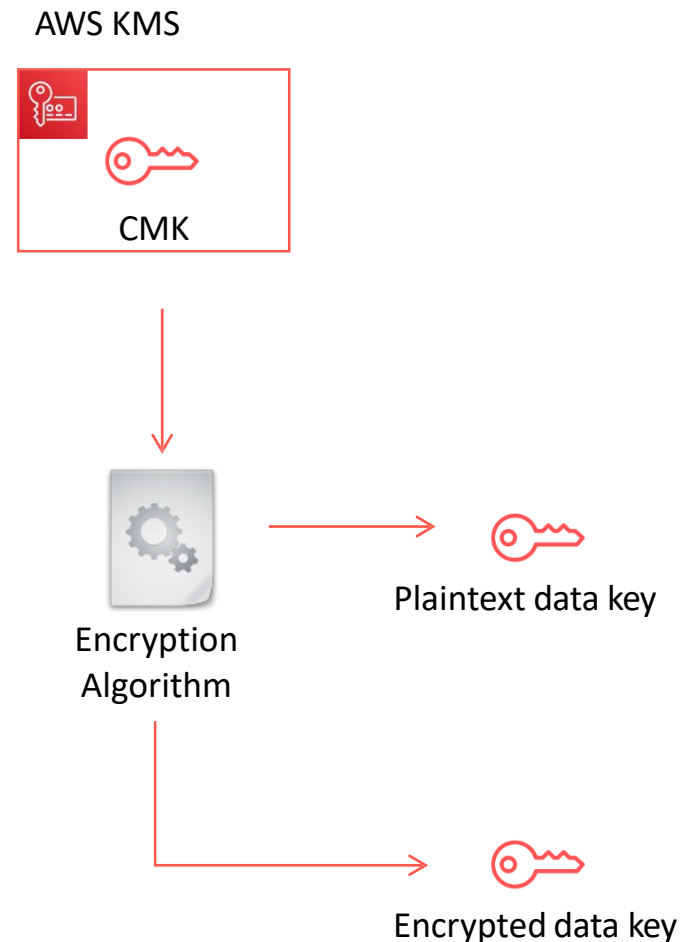
- Customer master keys are the primary resources in AWS KMS
- The CMK also contains the key material used to encrypt and decrypt data
- CMKs are created in AWS KMS. Symmetric CMKs and the private keys of asymmetric CMKs never leave AWS KMS unencrypted
- By default, AWS KMS creates the key material for a CMK
- Can also import your own key material
- A CMK can encrypt data up to 4KB in size
- A CMK can generate, encrypt and decrypt Data Encryption Keys (DEKs)

AWS Managed CMKs

- Created, managed, and used on your behalf by an AWS service that is integrated with AWS KMS
- You cannot manage these CMKs, rotate them, or change their key policies
- You also cannot use AWS managed CMKs in cryptographic operations directly; the service that creates them uses them on your behalf

Data Encryption Keys

- Data keys are encryption keys that you can use to encrypt data, including large amounts of data and other data encryption keys
- You can use AWS KMS customer master keys (CMKs) to generate, encrypt, and decrypt data keys
- AWS KMS does not store, manage, or track your data keys, or perform cryptographic operations with data keys
- You must use and manage data keys outside of AWS KMS

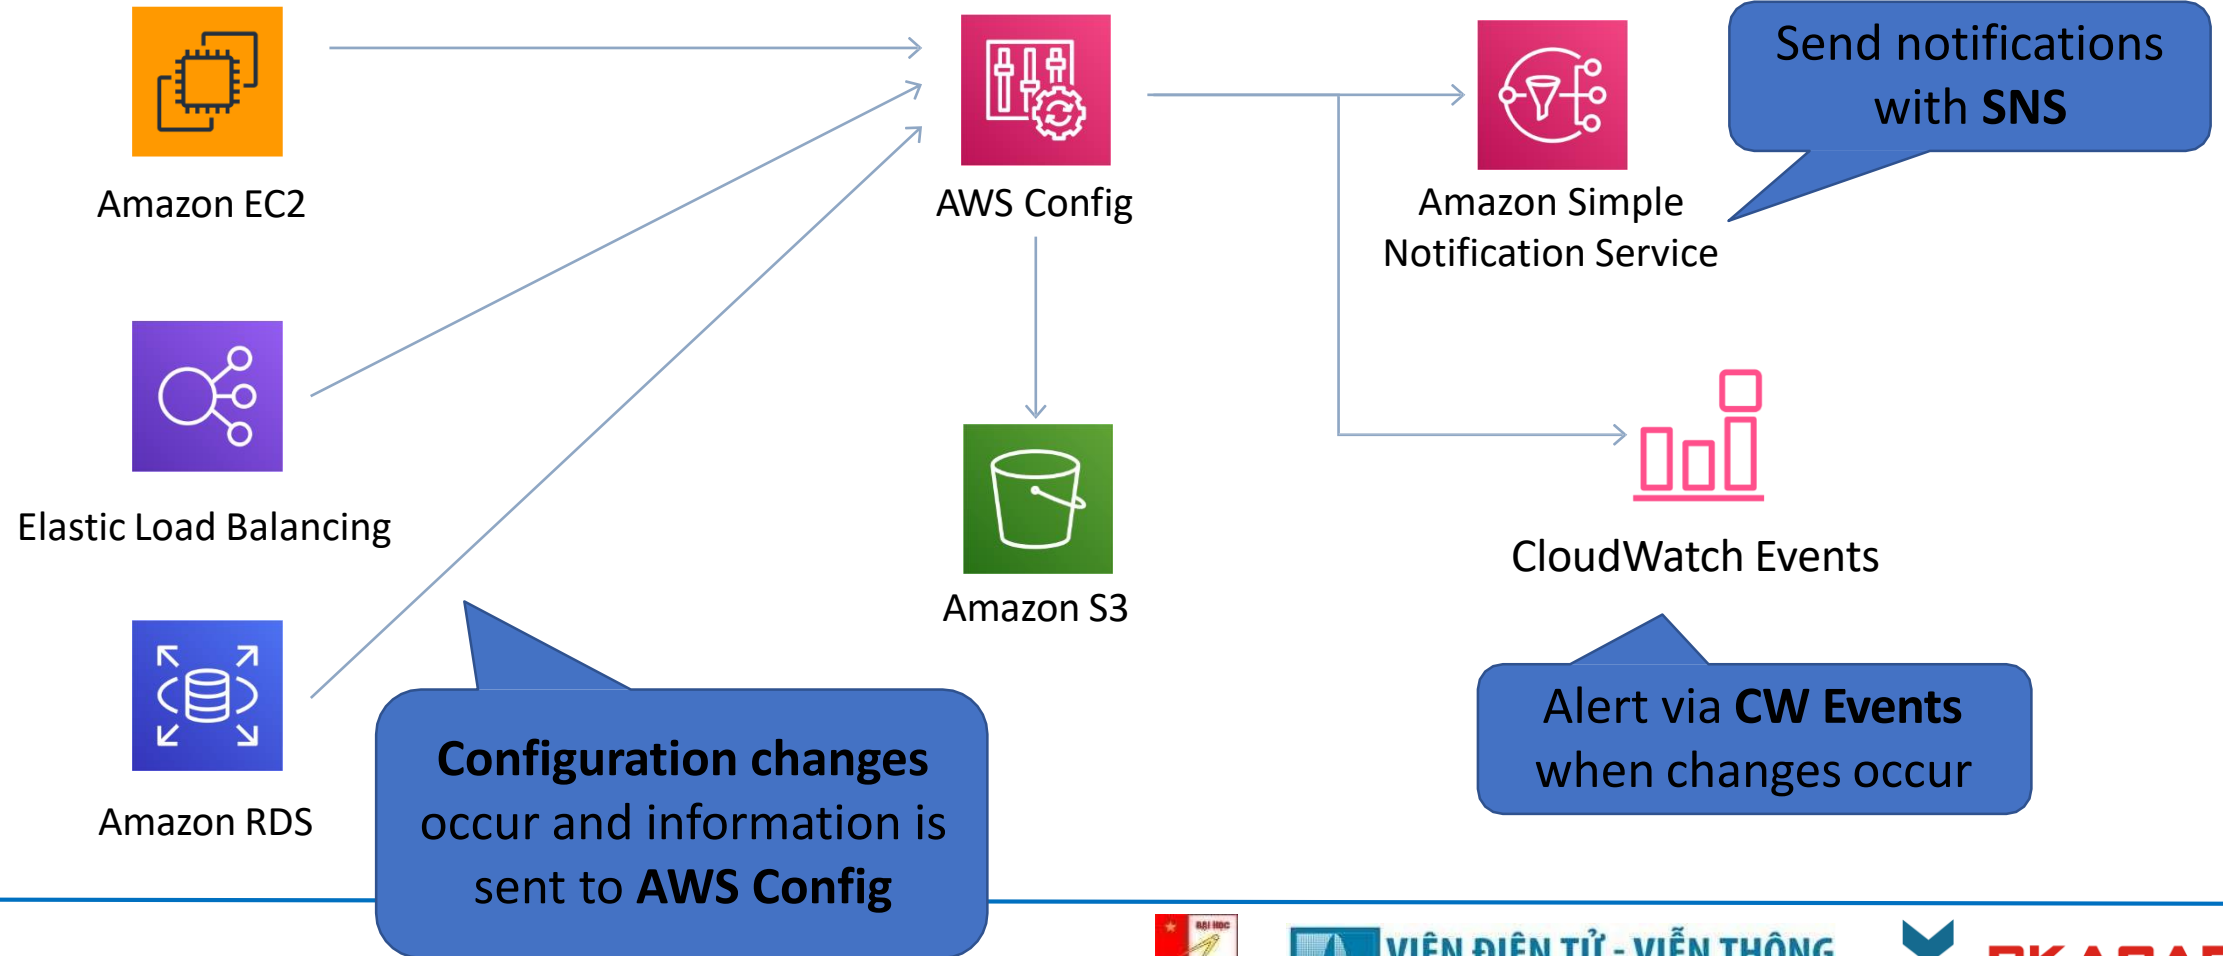


Customer Master Keys (CMKs)

Type of CMK	Can view	Can manage	Used only for my AWS account	Automatic rotation
Customer managed CMK	Yes	Yes	Yes	Optional. Every 365 days
AWS managed CMK	Yes	No	Yes	Required. Every 1095 days
AWS owned CMK	No	No	No	Varies

AWS Config

Example Services:



AWS Config

Example Rule	Description
s3-bucket-server-side-encryption-enabled	Checks that your Amazon S3 bucket either has S3 default encryption enabled or that the S3 bucket policy explicitly denies put-object requests without server side encryption
restricted-ssh	Checks whether security groups that are in use disallow unrestricted incoming SSH traffic
rds-instance-public-access-check	Checks whether the Amazon Relational Database Service (RDS) instances are not publicly accessible
cloudtrail-enabled	Checks whether AWS CloudTrail is enabled in your AWS account

Amazon Inspector

- Runs assessments that check for security exposures and vulnerabilities in EC2 instances
- Can be configured to run on a schedule
- Agent must be installed on EC2 for host assessments
- Network assessments do not require an agent

Amazon Inspector

Network Assessments

- Assessments: Network configuration analysis to check for ports reachable from outside the VPC
- If the Inspector Agent is installed on your EC2 instances, the assessment also finds processes reachable on port
- Price based on the number of instance assessments

Amazon Inspector

Host Assessments

- Assessments: Vulnerable software (CVE), host hardening (CIS benchmarks), and security best practices
- Requires an agent (auto-install with SSM Run Command)
- Price based on the number of instance assessments

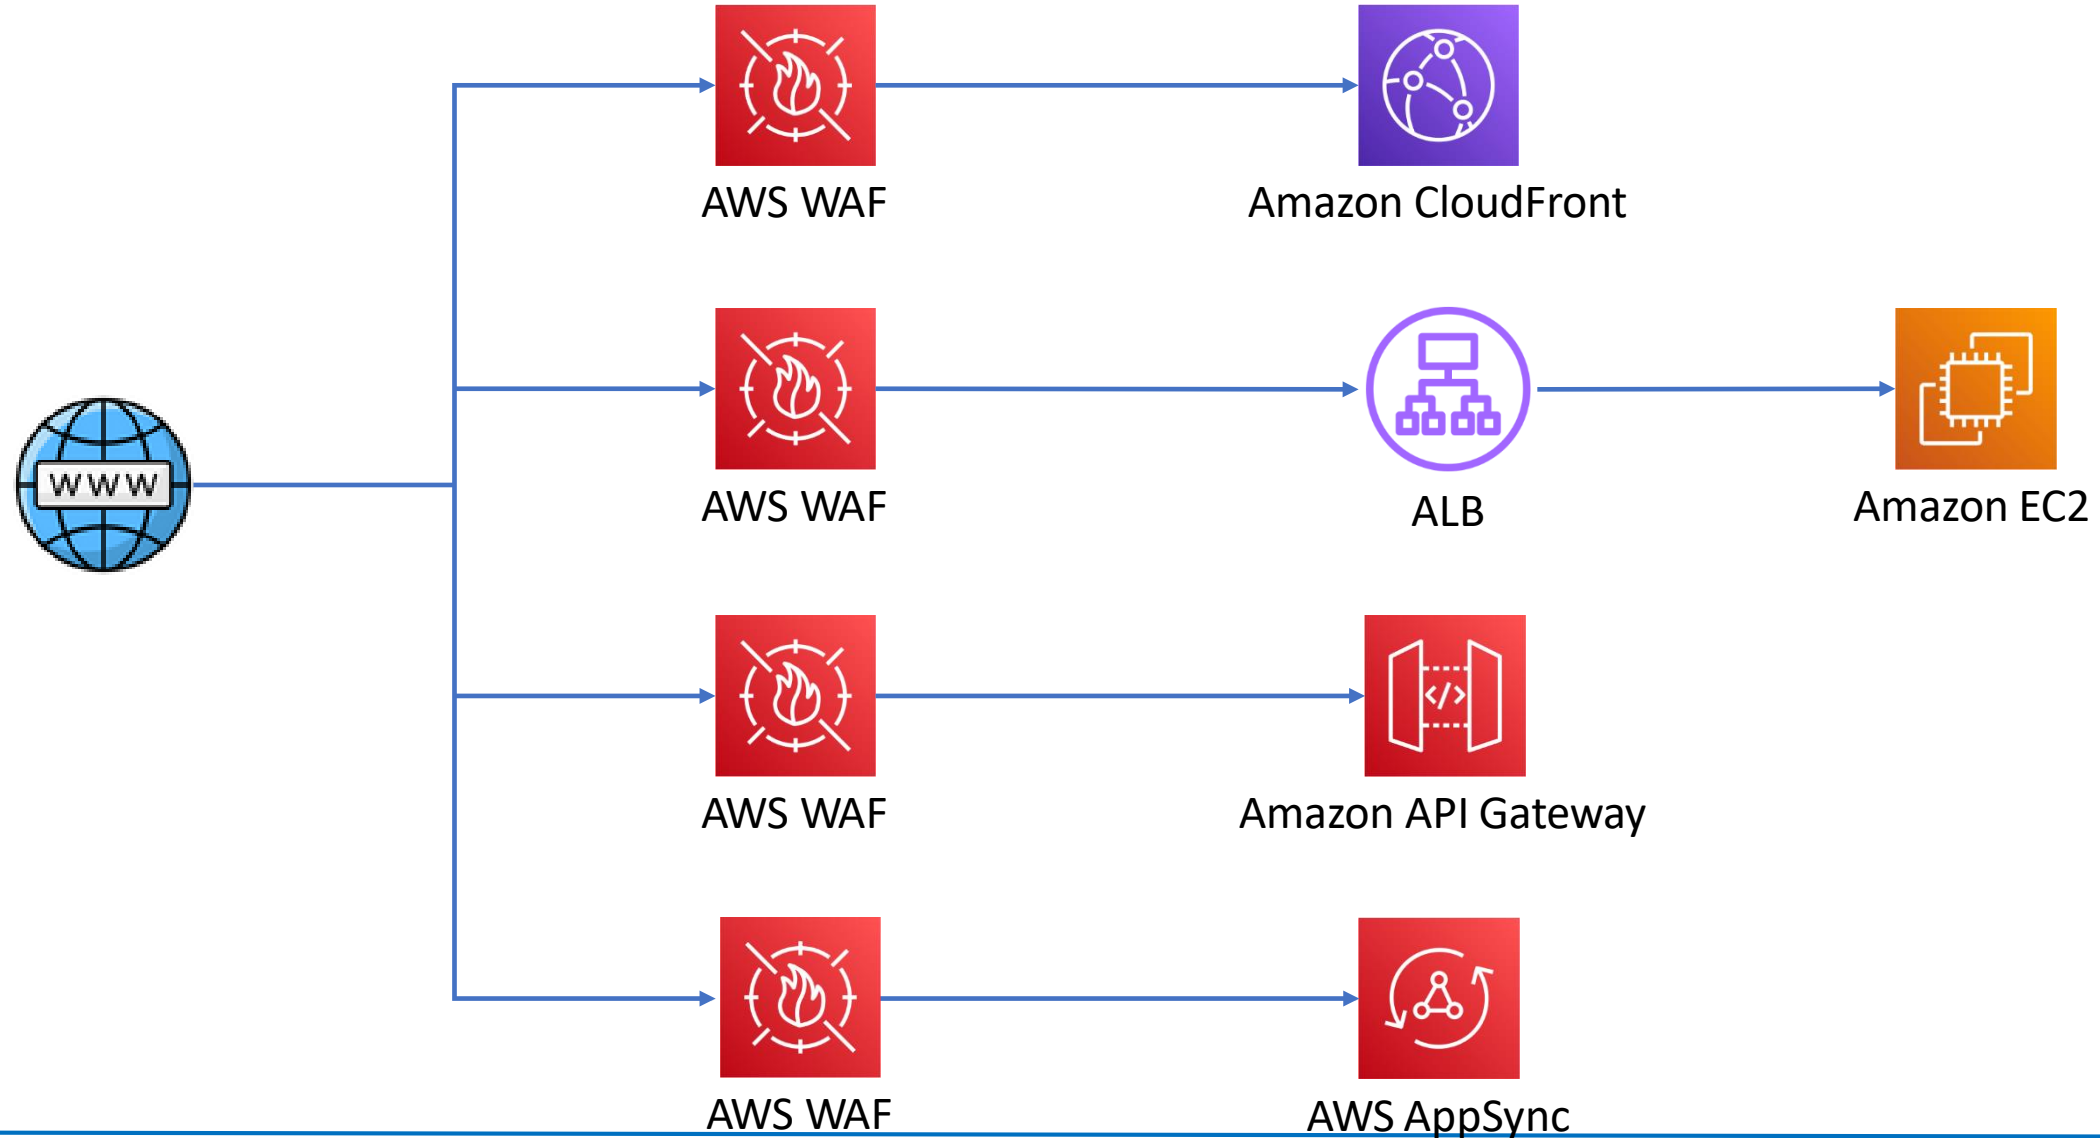
AWS Web Application Firewall (WAF)



AWS WAF

- AWS WAF is a web application firewall
- WAF lets you create rules to filter web traffic based on conditions that include IP addresses, HTTP headers and body, or custom URIs
- WAF makes it easy to create rules that block common web exploits like SQL injection and cross site scripting

AWS WAF



AWS WAF

- **Web ACLs** – You use a web access control list (ACL) to protect a set of AWS resources
- **Rules** – Each rule contains a statement that defines the inspection criteria, and an action to take if a web request meets the criteria
- **Rule groups** – You can use rules individually or in reusable rule groups

Rule type



IP set

Use IP sets to identify a specific list of IP addresses.



Rule builder

Use a custom rule to inspect for patterns including query strings, headers, countries, and rate limit violations.



Rule group

Use a rule group to combine rules into a single logical set.

AWS WAF

- **IP Sets** - An IP set provides a collection of IP addresses and IP address ranges that you want to use together in a rule statement
- **Regex pattern set** - A regex pattern set provides a collection of regular expressions that you want to use together in a rule statement

AWS WAF

A **rule action** tells AWS WAF what to do with a web request when it *matches* the criteria defined in the rule:

- **Count** – AWS WAF counts the request but doesn't determine whether to allow it or block it. With this action, AWS WAF continues processing the remaining rules in the web ACL
- **Allow** – AWS WAF allows the request to be forwarded to the AWS resource for processing and response
- **Block** – AWS WAF blocks the request and the AWS resource responds with an HTTP 403 (Forbidden) status code

AWS WAF

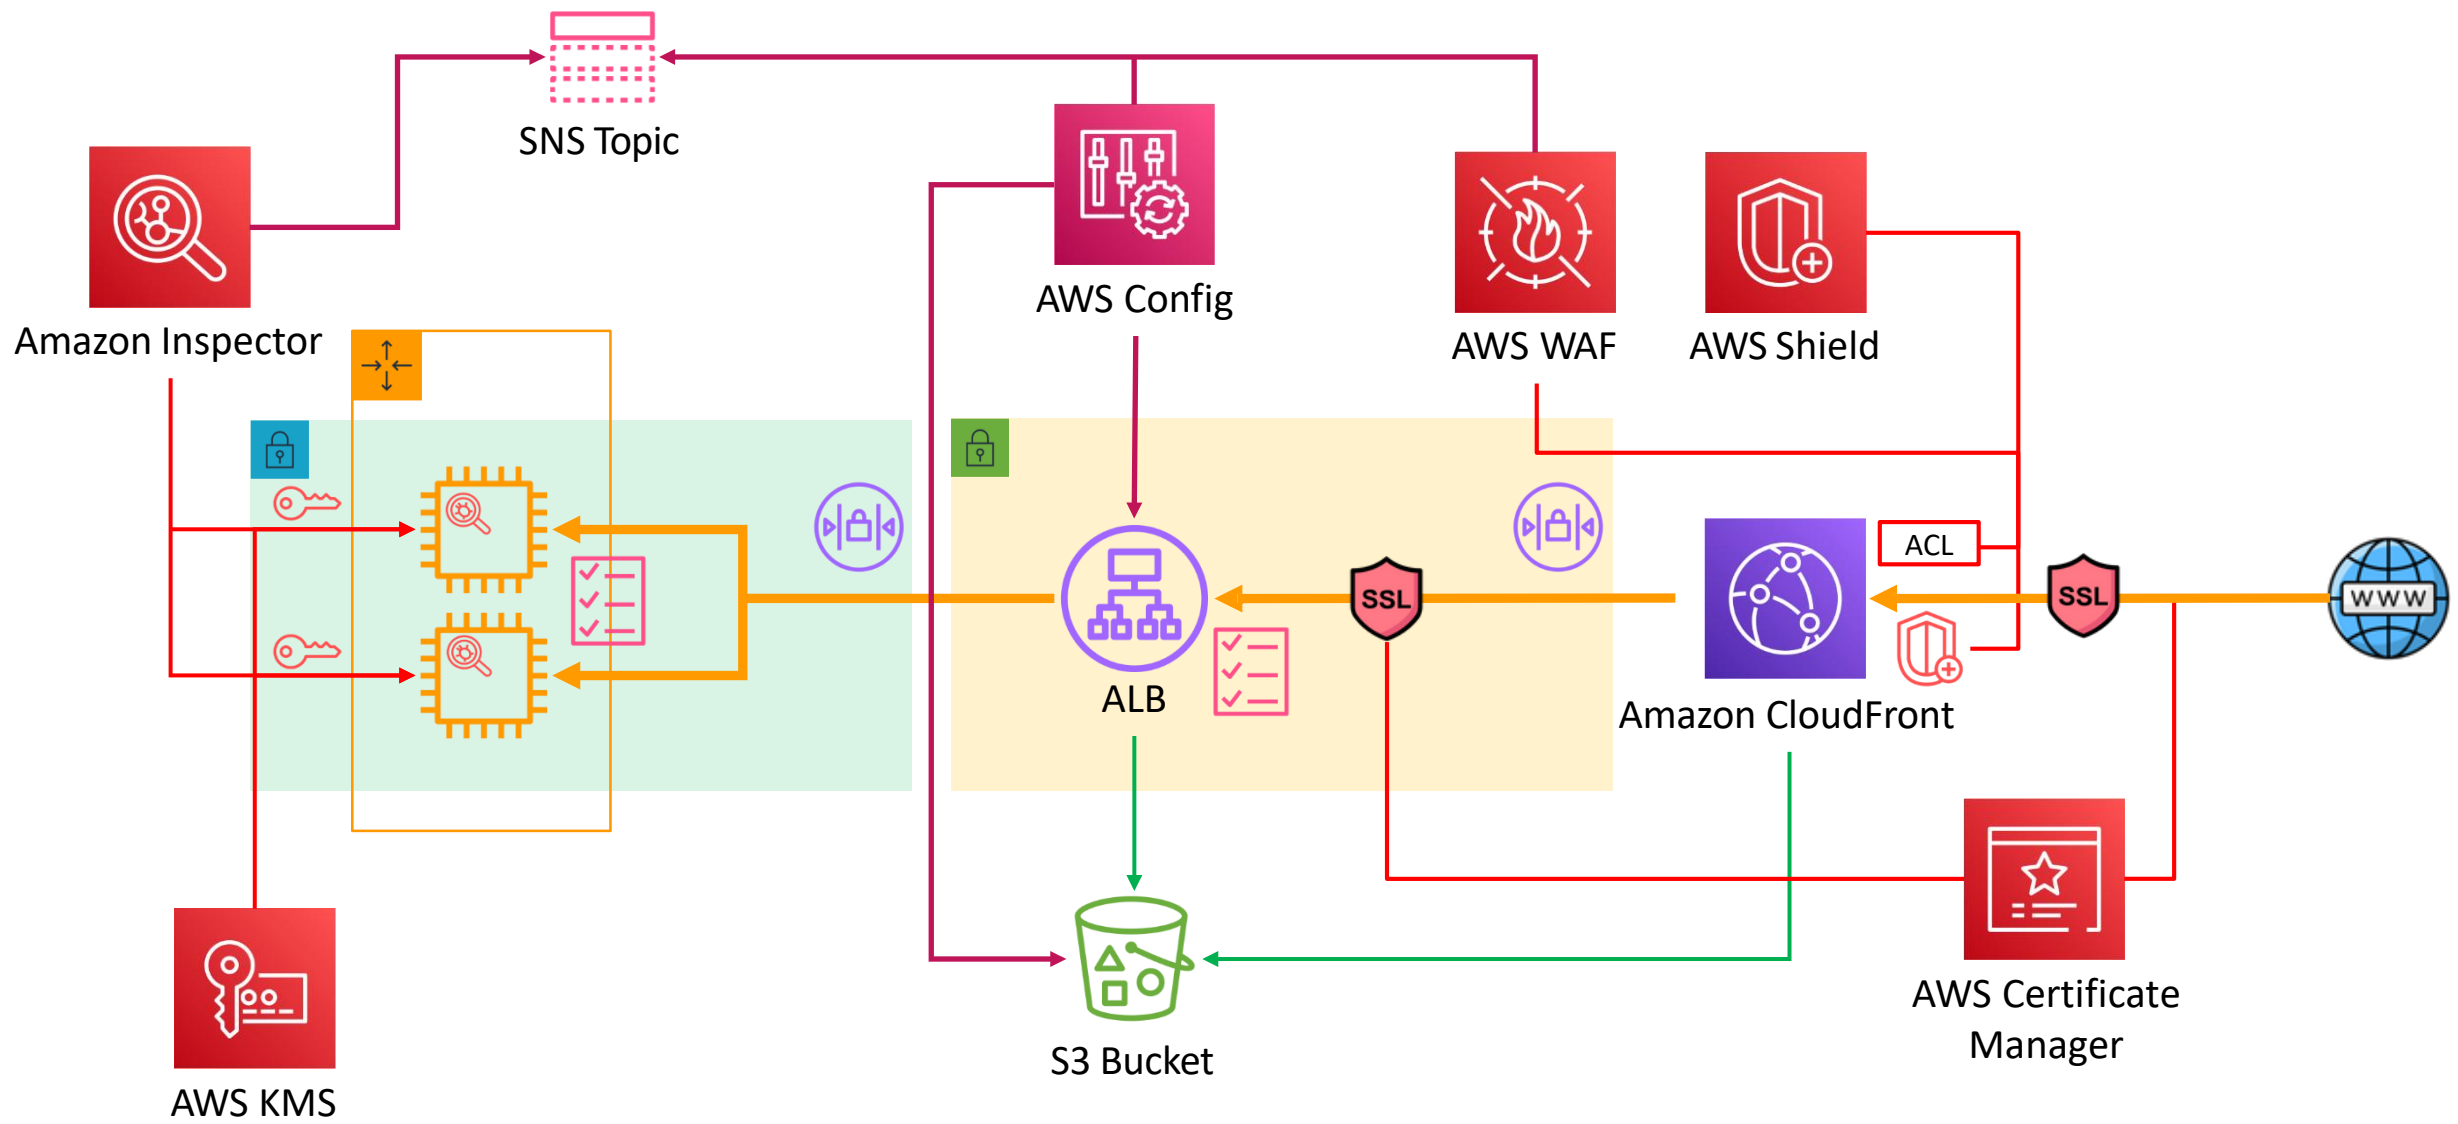
Match statements compare the web request or its origin against conditions that you provide

Match Statement	Description
Geographic match	Inspects the request's country of origin
IP set match	Inspects the request against a set of IP addresses and address ranges
Regex pattern set	Compares regex patterns against a specified request component
Size constraint	Checks size constraints against a specified request component
SQLi attack	Inspects for malicious SQL code in a specified request component
String match	Compares a string to a specified request component
XSS scripting attack	Inspects for cross-site scripting attacks in a specified request component

AWS Shield

- AWS Shield is a managed Distributed Denial of Service (DDoS) protection service
- Safeguards web application running on AWS with always-on detection and automatic inline mitigations
- Helps to minimize application downtime and latency
- Two tiers –
 - **Standard** – no cost
 - **Advanced** - \$3k USD per month and 1 year commitment
- Integrated with Amazon CloudFront (standard included by default)

Build a Secure Multi-Tier Architecture



AWS GuardDuty

- Intelligent threat detection service
- Detects account compromise, instance compromise, malicious reconnaissance, and bucket compromise
- Continuous monitoring for events across:
 - **AWS CloudTrail Management Events**
 - **AWS CloudTrail S3 Data Events**
 - **Amazon VPC Flow Logs**
 - **DNS Logs**

Architecture Patterns - Security



Architecture Patterns – Security

Requirement

Need to enable custom domain name and encryption in transit for an application running behind an Application Load Balancer

Company records customer information in CSV files in an Amazon S3 bucket and must not store PII data

For compliance reasons all S3 buckets must have encryption enabled and any non-compliant buckets must be auto remediated

Solution

Use AWS Route 53 to create an Alias record to the ALB's DNS name and attach an SSL/TLS certificate issued by Amazon Certificate Manager (ACM)

Use Amazon Macie to scan the S3 bucket for any PII data

Use AWS Config to check the encryption status of the buckets and use auto remediation to enable encryption as required

Architecture Patterns – Security

Requirement

EC2 instances must be checked against CIS benchmarks every 7 days

Website running on EC2 instances behind and ALB must be protected against well known web exploits

Need to block access to an application running on an ALB from connections originating in a specific list of countries

Solution

Install the Amazon Inspector agent and configure a host assessment every 7 days

Create a Web ACL in AWS WAF to protect against web exploits and attach to the ALB

Create a Web ACL in AWS WAF with a geographic match and block traffic that matches the list of countries

AWS Monitoring & Logging

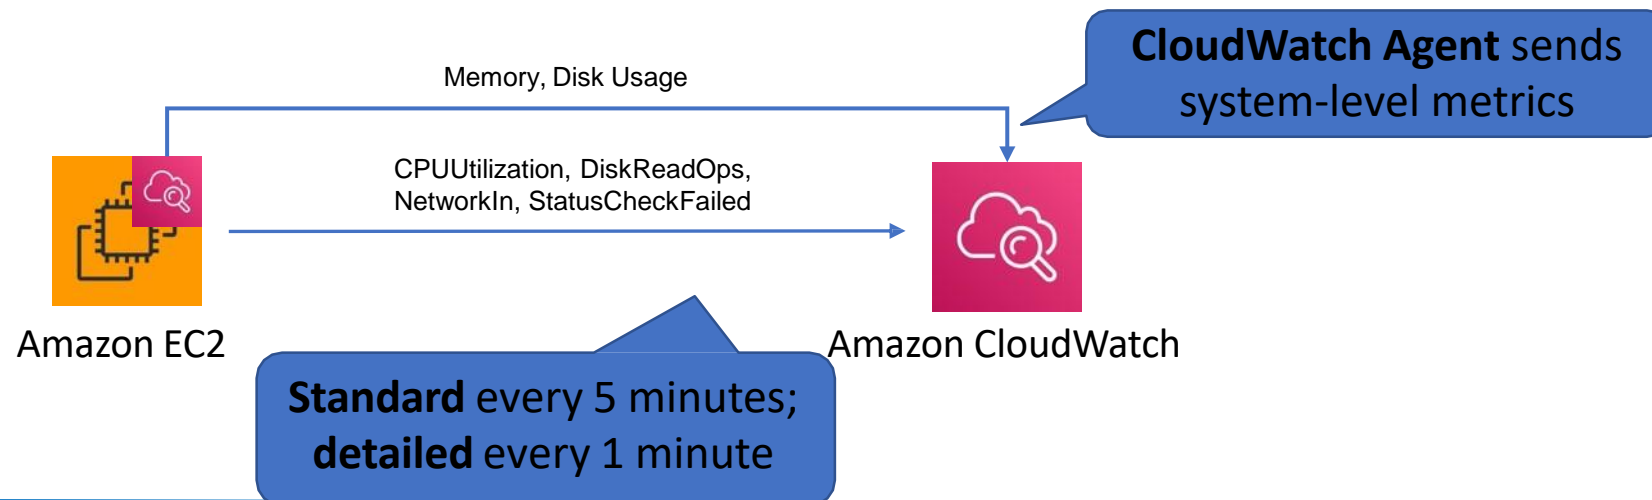


Amazon CloudWatch

- **CloudWatch Metrics** – services send time-ordered data points to CloudWatch
- **CloudWatch Alarms** – monitor metrics and initiate actions
- **CloudWatch Logs** – centralized collection of system and application logs
- **CloudWatch Events** – stream of system events describing changes to AWS resources and can trigger actions

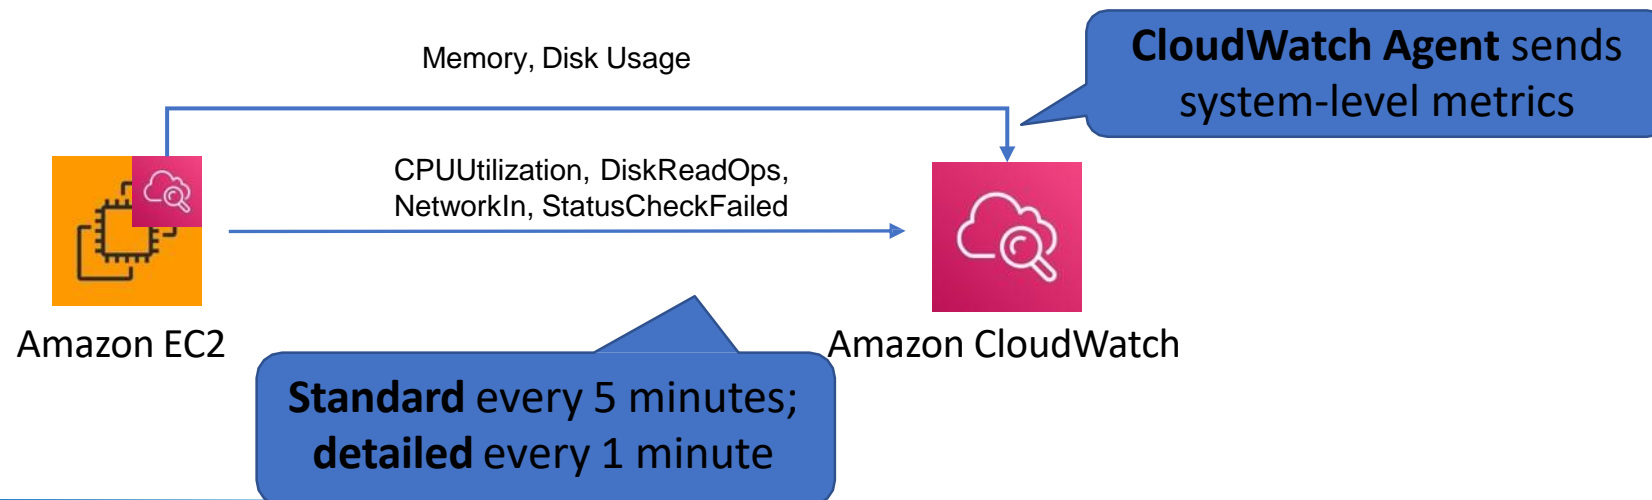
Amazon CloudWatch Metrics

- Metrics are sent to CloudWatch for many AWS services
- EC2 metrics are sent every **5 minutes** by default (free)
- Detailed EC2 monitoring sends every **1 minute** (chargeable)
- Unified CloudWatch Agent sends system-level metrics for EC2 and on-premises servers
- System-level metrics include memory and disk usage



Amazon CloudWatch Metrics


- Can publish custom metrics using CLI or API
- Custom metrics are one of the following resolutions:
 - **Standard resolution** – data having a one-minute granularity
 - **High resolution** – data at a granularity of one second
- AWS metrics are standard resolution by default



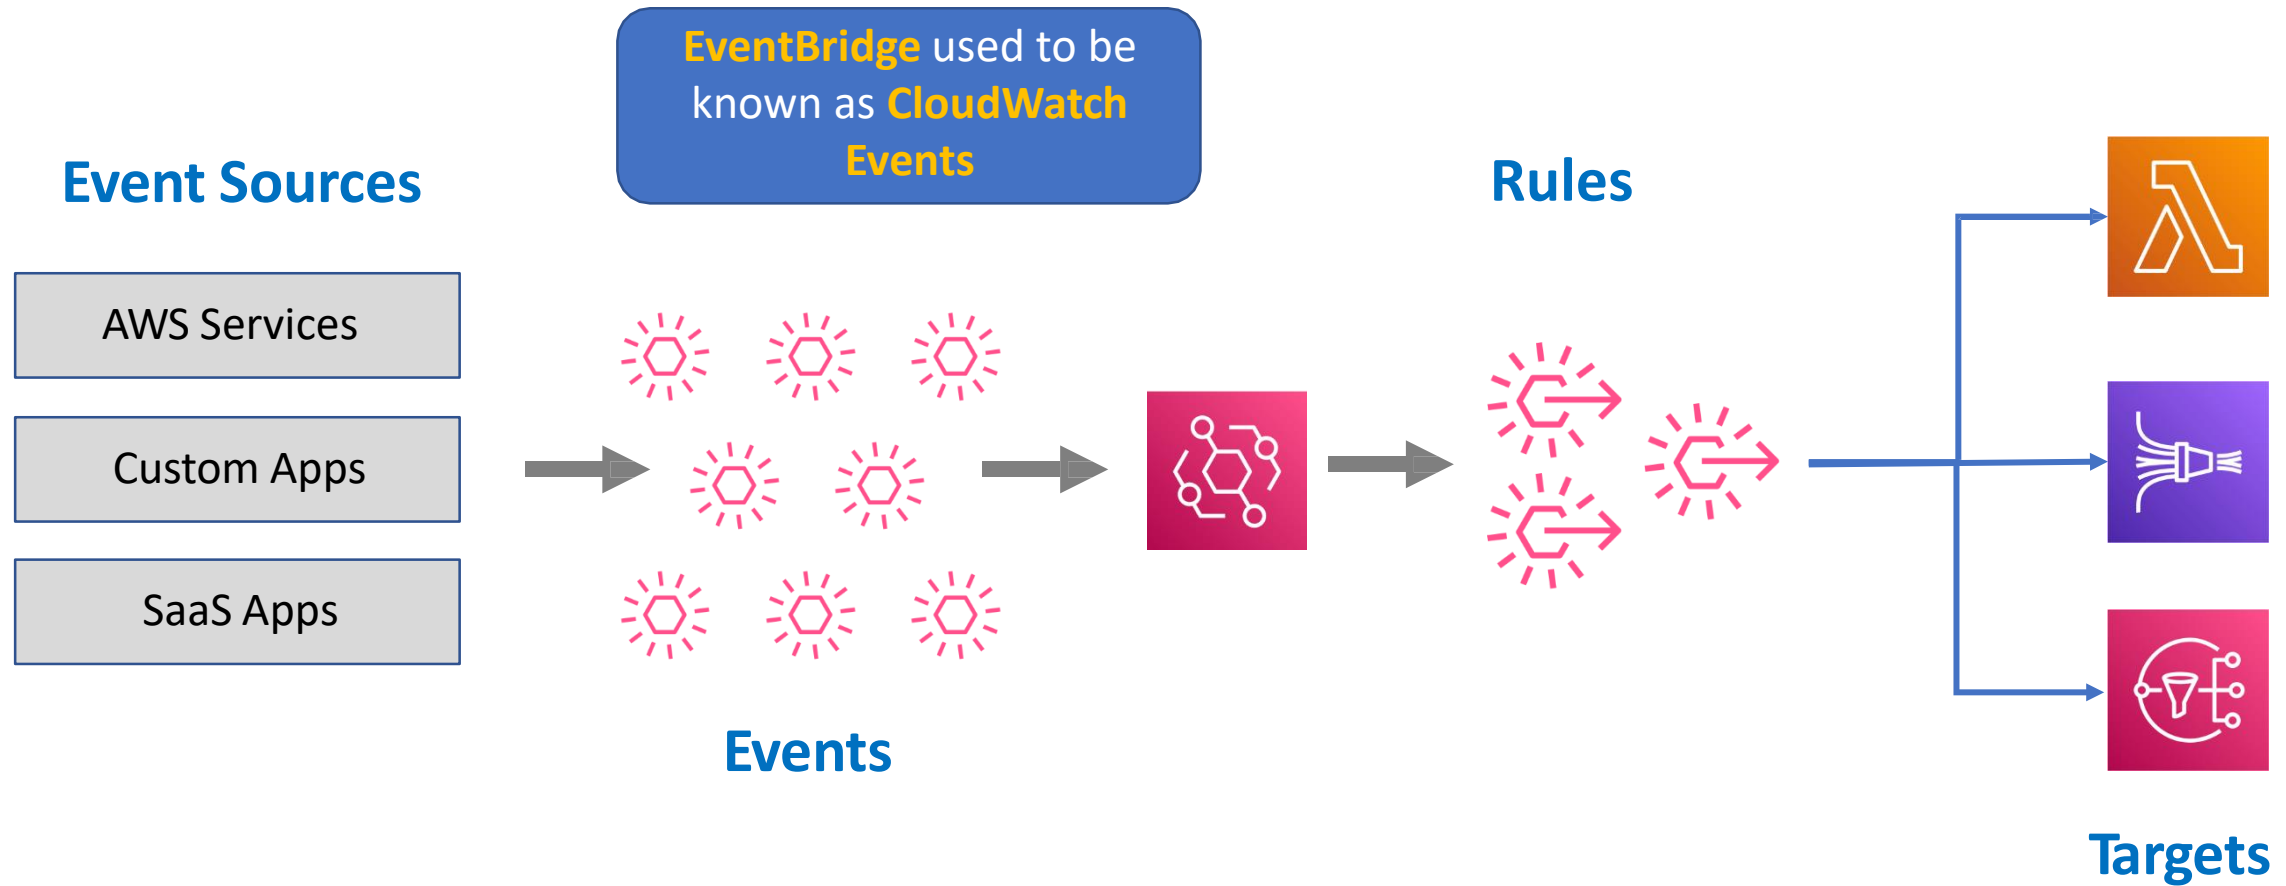
Amazon CloudWatch Alarms

Two types of alarms

- **Metric alarm** – performs one or more actions based on a single metric
- **Composite alarm** – uses a rule expression and takes into account multiple alarms
- **Metric alarm states:**
 - **OK** – Metric is within a threshold
 - **ALARM** – Metric is outside a threshold
 - **INSUFFICIENT_DATA** – not enough data

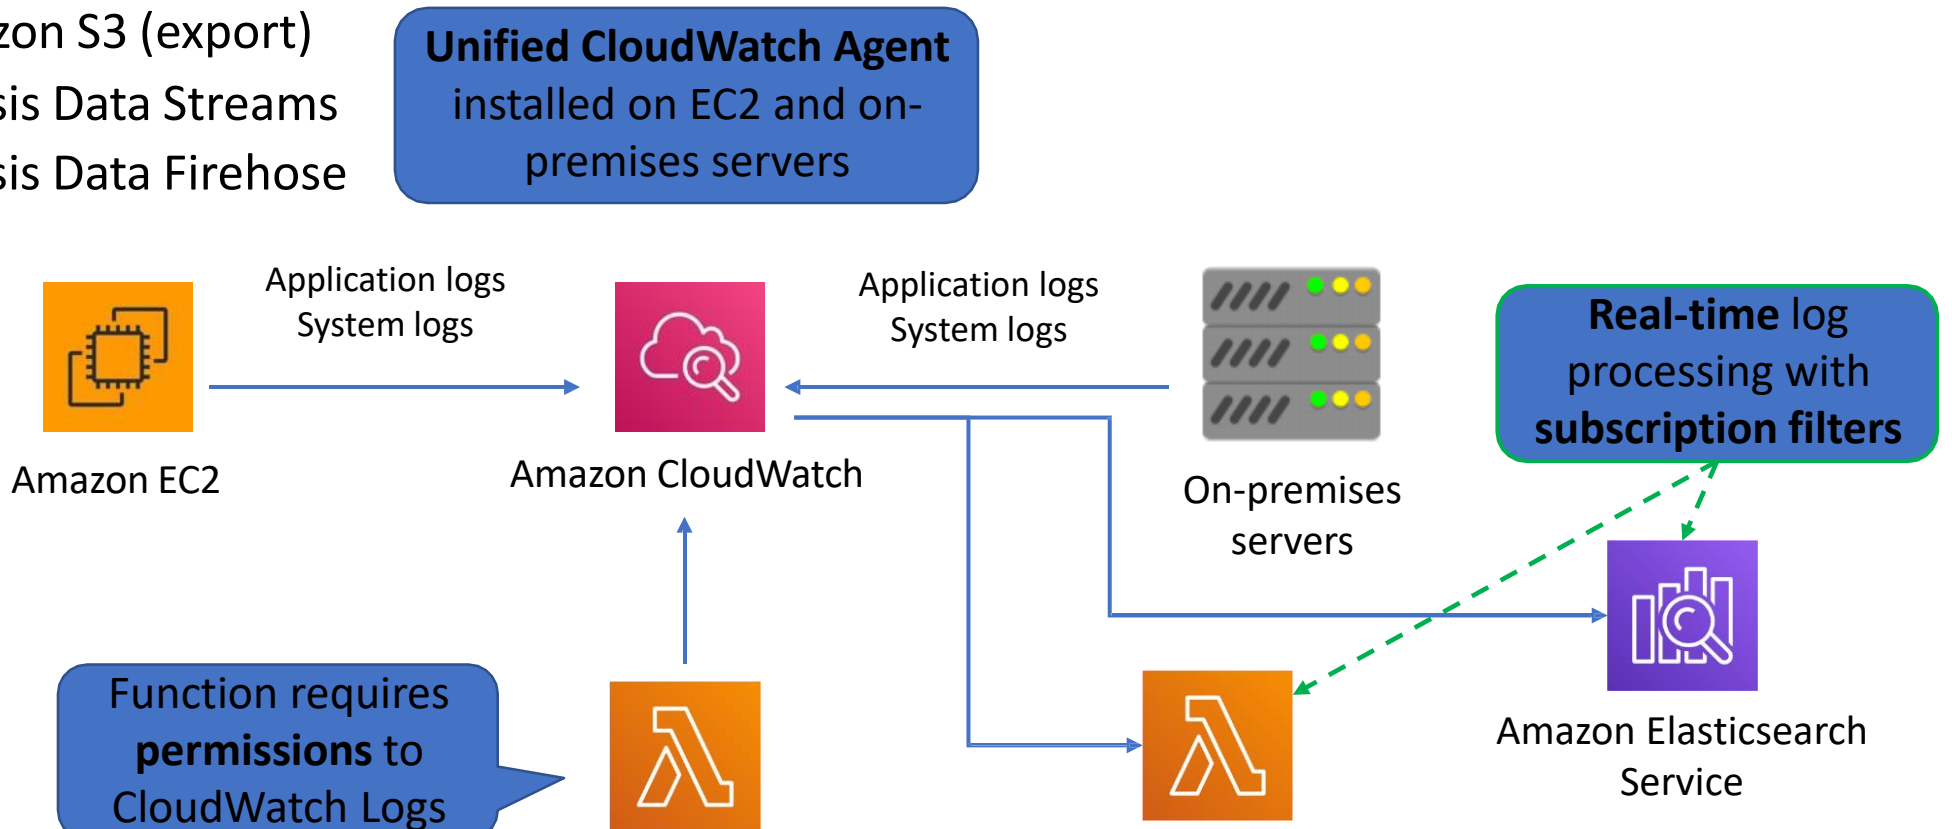


Amazon CloudWatch Events / EventBridge



Amazon CloudWatch Logs

- Gather application and system logs in CloudWatch
- Defined expiration policies and KMS encryption
- Send to:
 - Amazon S3 (export)
 - Kinesis Data Streams
 - Kinesis Data Firehose



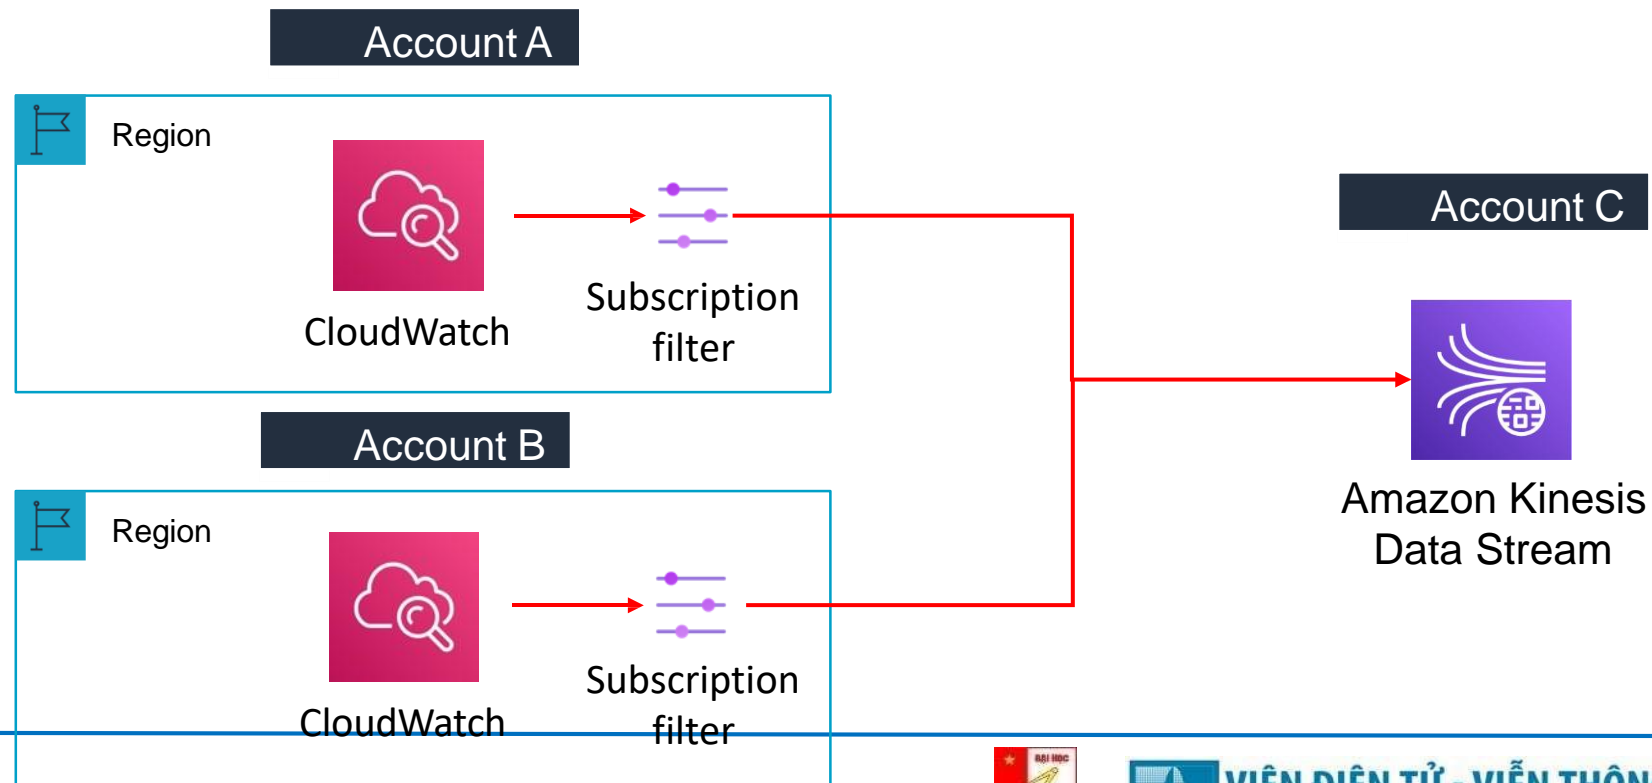
AWS Lambda




AWS Lambda

Cross-Account Log Data Sharing

- Share CloudWatch Logs across accounts
- Kinesis Data Streams is the only supported destination
- **Log data sender** – sends log data to the recipient
- **Log data recipient** – sends data to a Kinesis Data stream



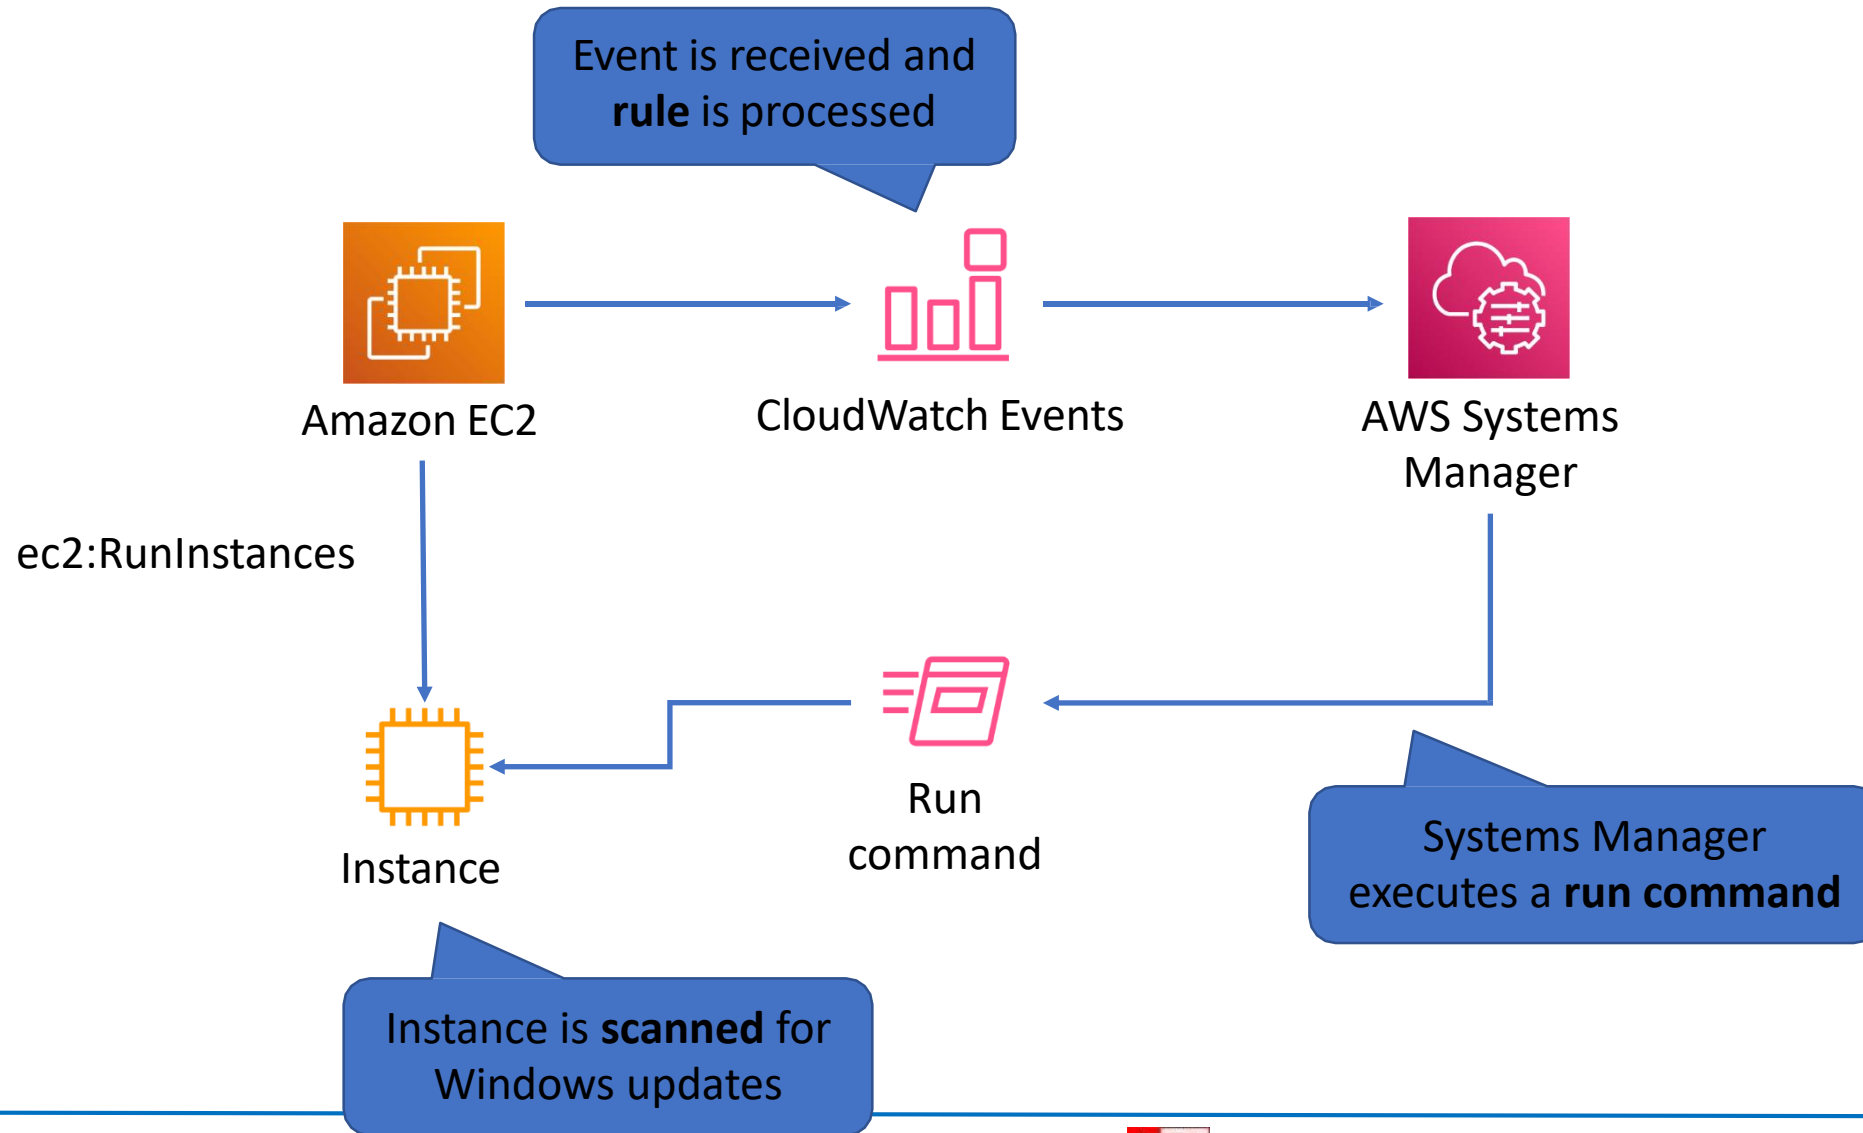
Export CloudWatch Logs to S3



Logs are **exported** to the S3 bucket

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": "s3:GetBucketAcl",
      "Effect": "Allow",
      "Resource": "arn:aws:s3:::mybucket",
      "Principal": { "Service": "logs.us-east-1.amazonaws.com" }
    },
    {
      "Action": "s3:PutObject",
      "Effect": "Allow",
      "Resource": "arn:aws:s3:::mybucket/*",
      "Condition": { "StringEquals": { "s3:x-amz-acl": "bucket-owner-full-control" } },
      "Principal": { "Service": "logs.us-east-1.amazonaws.com" }
    }
  ]
}
```

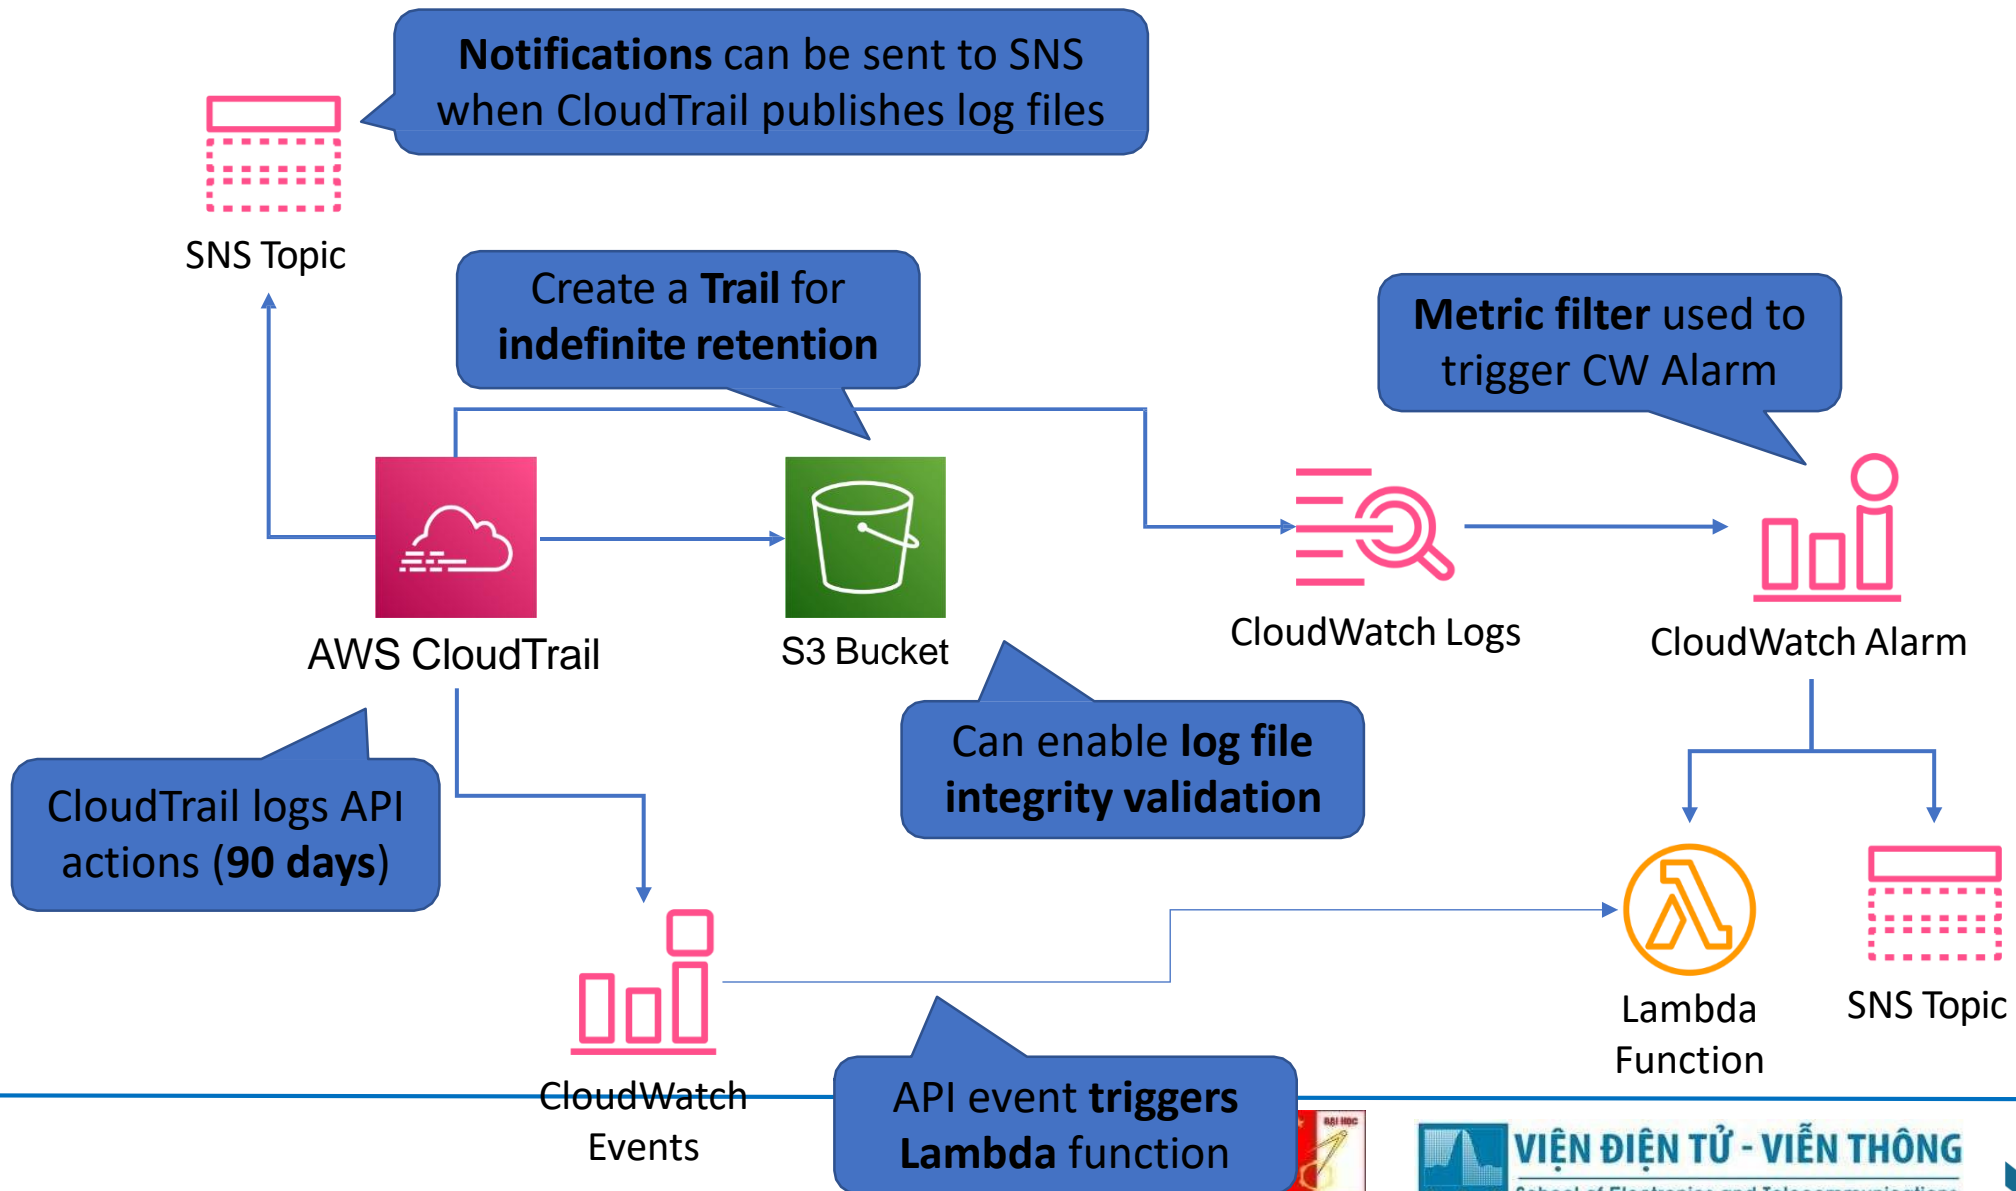
Trigger SSM on Instance Launch



AWS CloudTrail Use Cases



AWS CloudTrail



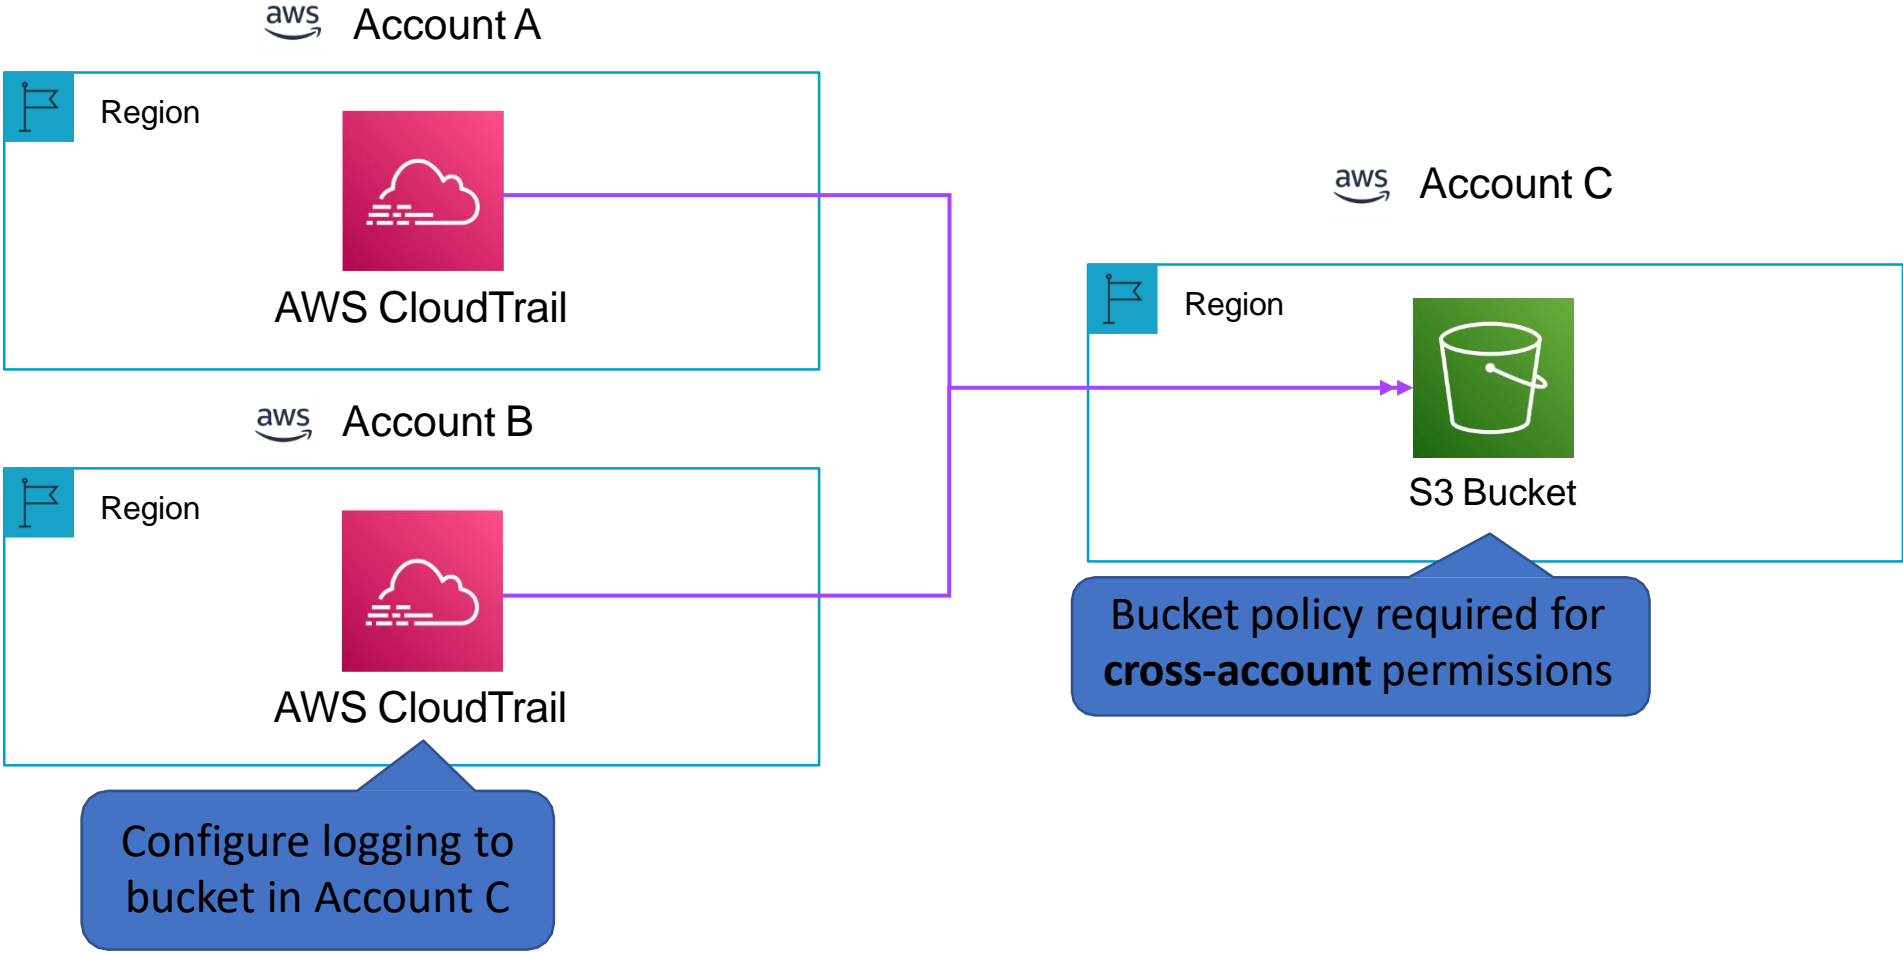
AWS CloudTrail

- CloudTrail logs **API activity** for auditing
- By default, management events are logged and retained for 90 days
- A **CloudTrail Trail** logs any events to S3 for indefinite retention
- Trail can be within Region or all Regions
- CloudWatch Events can be triggered based on API calls in CloudTrail
- Events can be streamed to CloudWatch Logs

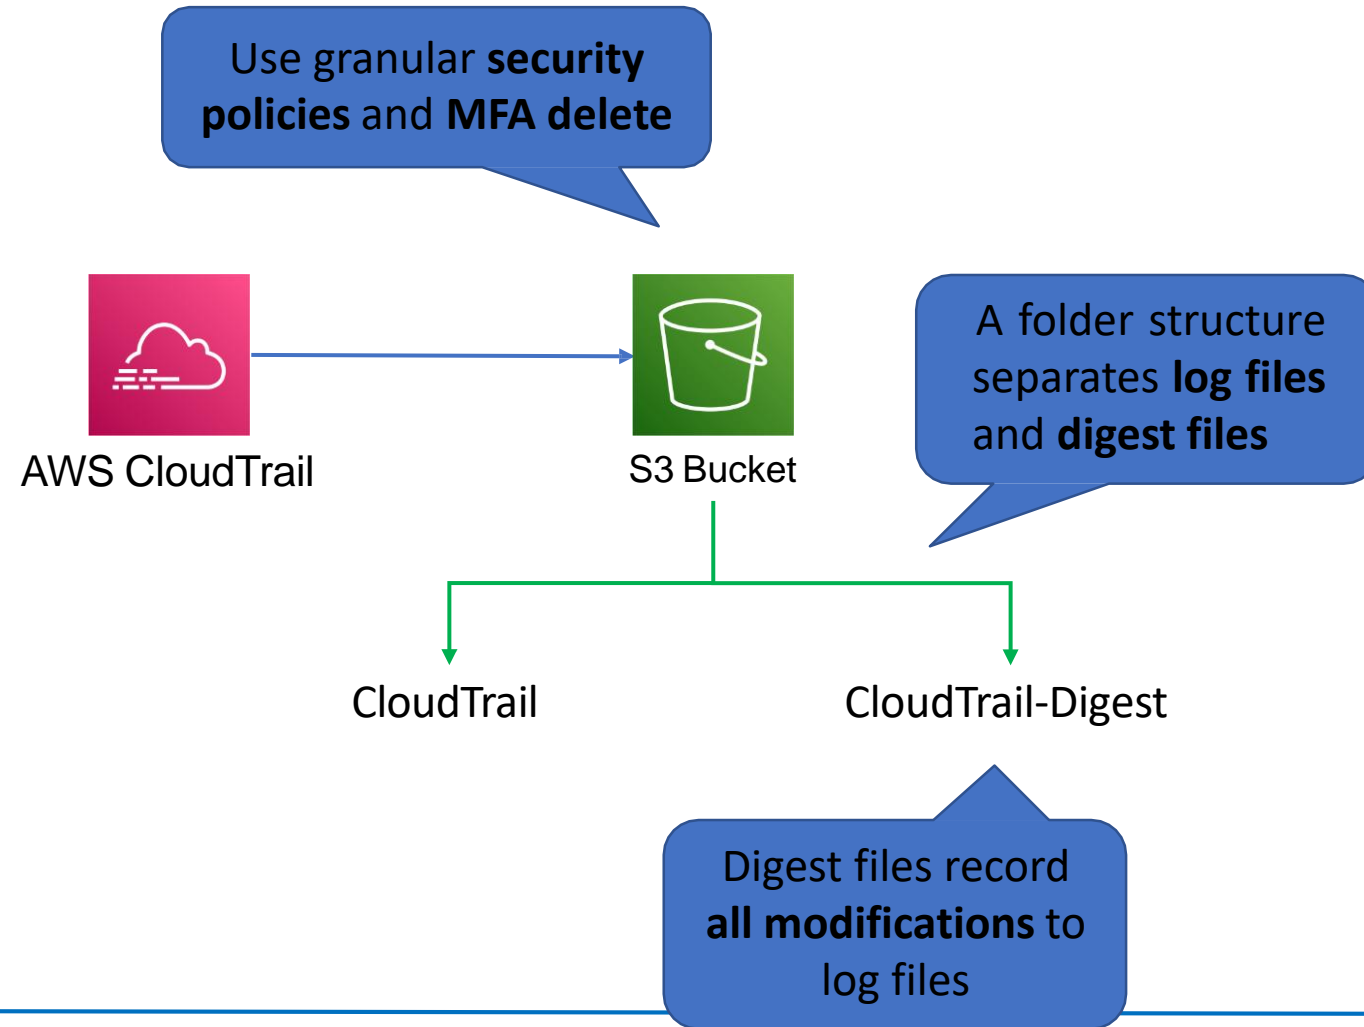
CloudTrail – Management and Data Events

- **Management events** provide information about management operations that are performed on resources in your AWS account
- **Data events** provide information about the resource operations performed on or in a resource

CloudTrail – Multi Account and Region



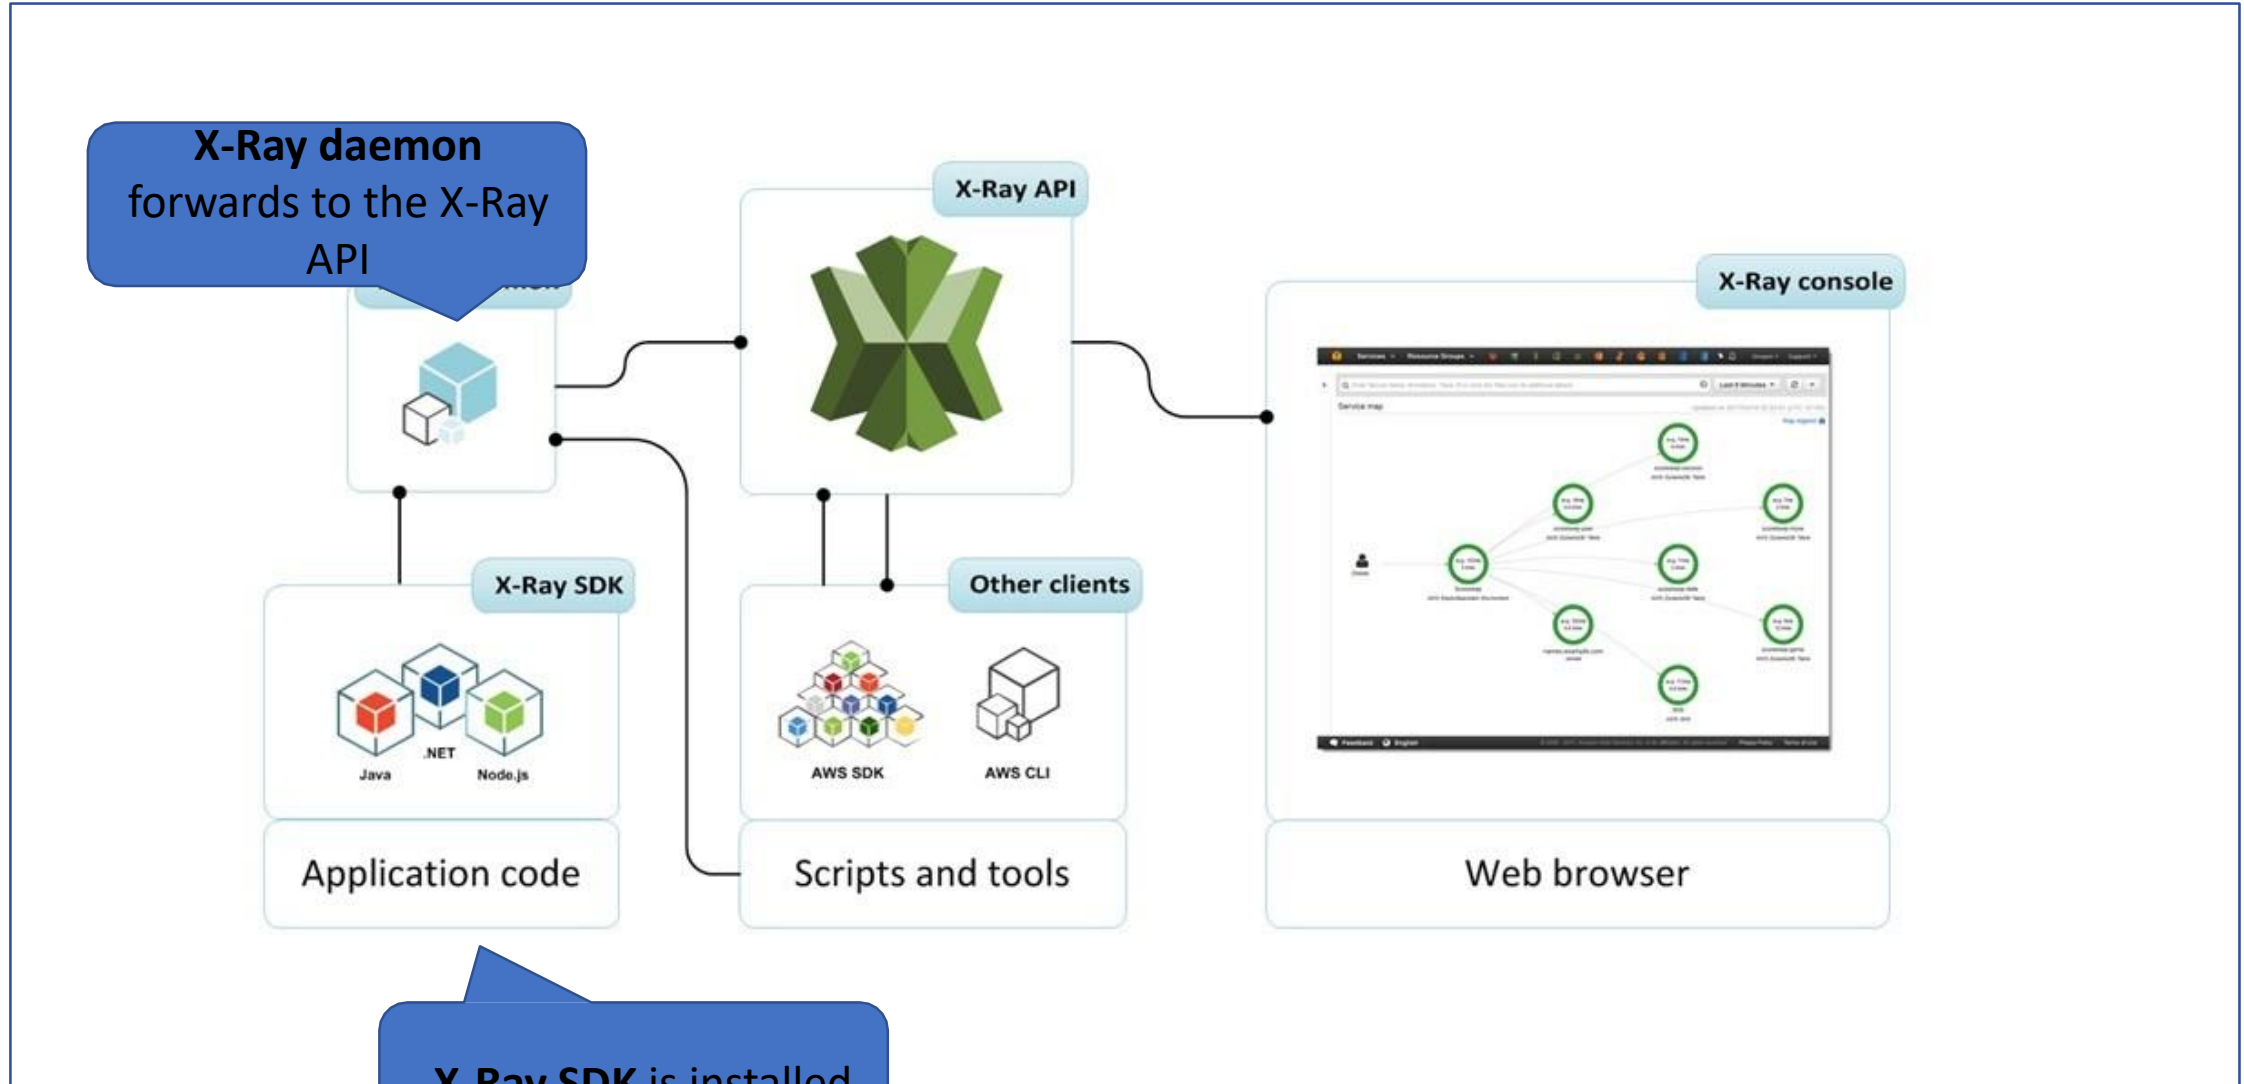
Enable CloudTrail Log File Validation



AWS X-Ray

- Analyze and debug production, distributed applications, such as those built using a microservices architecture
- AWS X-Ray supports applications running on:
 - Amazon EC2
 - Amazon ECS
 - AWS Lambda
 - AWS Elastic Beanstalk
- X-Ray SDK must be integrated with application along with X-Ray agent

AWS X-Ray



X-Ray SDK is installed in app and forwards to the X-Ray daemon

AWS X-Ray SDK

- The **X-Ray SDK** is installed in your application and forwards to the X-Ray daemon which forwards to the X-Ray API.
- You can then visualize what is happening in the X-Ray console.
- The X-Ray SDK provides:
 - **Interceptors** to add your code to trace incoming HTTP requests.
 - **Client handlers** to instrument AWS SDK client that your application uses to call other AWS services.
 - **An HTTP client** to use to instrument calls to other internal and external HTTP web services.

Amazon VPC Flow Logs

Log network traffic for Amazon VPC, subnet or single interfaces

- Stores logs in AWS CloudWatch Logs
- Can be enabled on
 - Amazon VPC, a subnet, or a network interface
 - Amazon VPC & Subnet enables logging for all interfaces in the VPC/subnet
 - Each network interface has a unique log stream
- Flow logs do not capture real-time log streams for your network interfaces
- Filter desired result based on need
 - All, Reject, Accept
 - Troubleshooting or security related with alerting needs?
 - Think before enabling All on VPC, will you use it?

VPC Flow Logs

- Agentless
- Enable per ENI, per subnet, or per VPC
- Logged to AWS CloudWatch Logs
- Create CloudWatch metrics from log data
- Alarm on those metrics

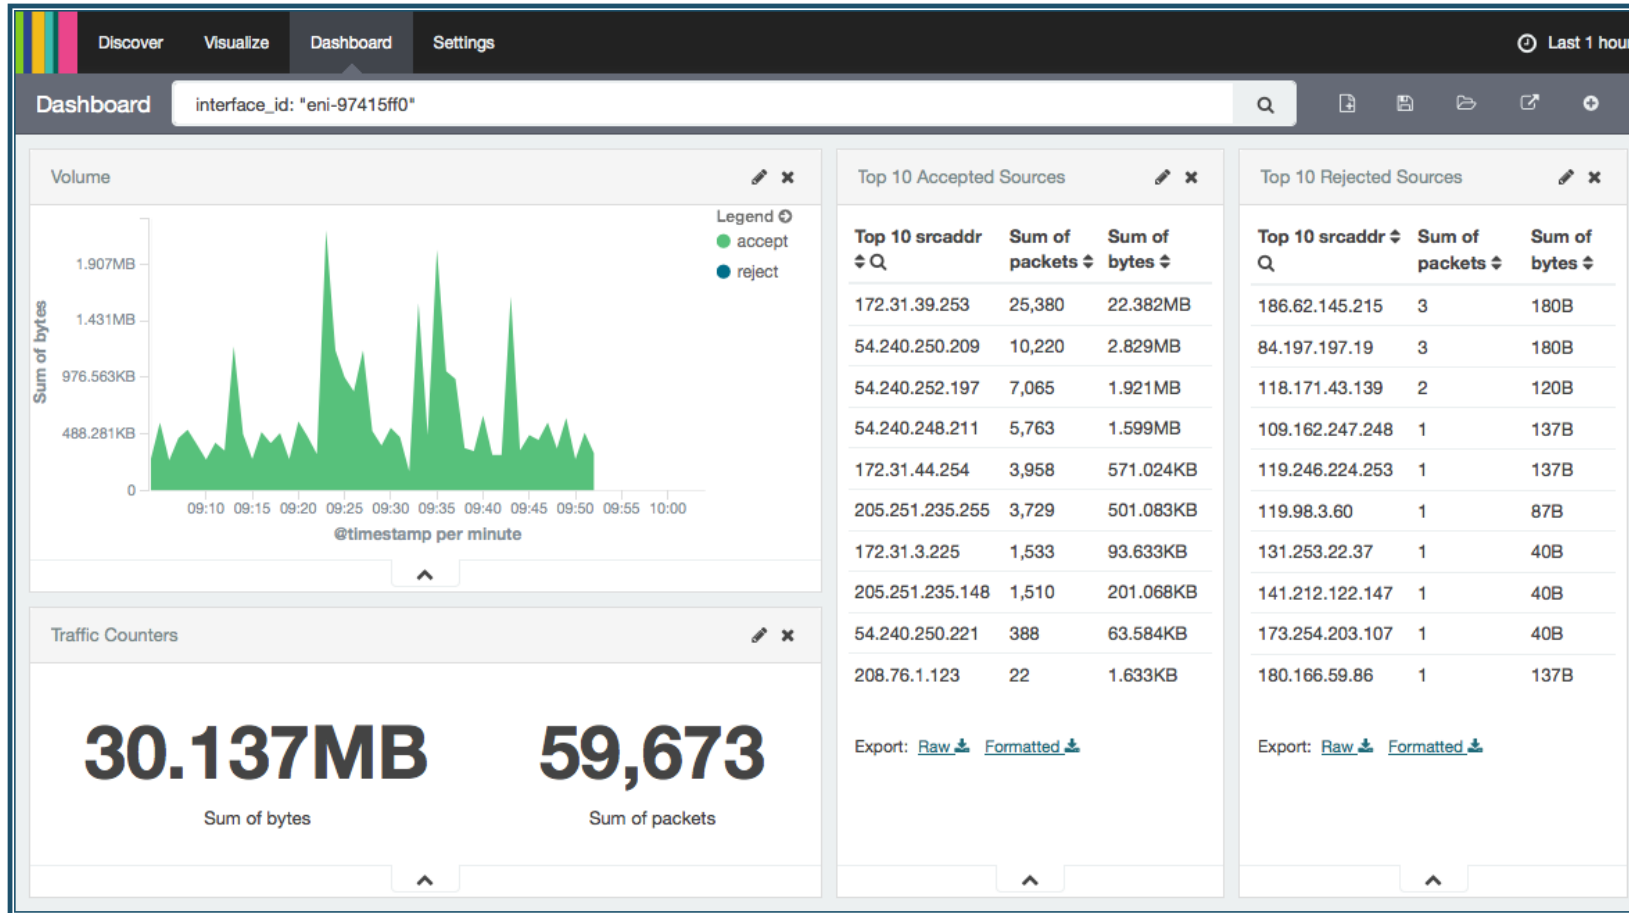
Interface Source IP Source port Protocol Packets

AWS account

Event Data	Interface	Source IP	Source port	Destination IP	Destination port	Protocol	Bytes	Packets	Start/end time	Accept or reject	
▶ 2 41747	eni-b30b9cd5	119.147.115.32	10.1.1.179	6000	22	6	1	40	1442975475 - 1442975535	REJECT OK	
▼ 2 41747	eni-b30b9cd5	169.54.233.117	10.1.1.179	21188	80	6	1	40	1442975535 - 1442975595	REJECT OK	
▼ 2 41747	eni-b30b9cd5	212.7.209.6	10.1.1.179	3389	3389	6	1	40	1442975596 - 1442975655	REJECT OK	
▼ 2 41747	eni-b30b9cd5	189.134.227.225	10.1.1.179	3966	4	23	6	2	180	1442975656 - 1442975716	REJECT OK
▼ 2 41747	eni-b30b9cd5	77.85.113.238	10.1.1.179	0	0	1	1	100	1442975656 - 1442975716	REJECT OK	
▼ 2 41747	eni-b30b9cd5	10.1.1.179	198.60.73.8	512	123	17	1	76	1442975776 - 1442975836	ACCEPT OK	

Destination IP Destination port Bytes Start/end time

VPC Flow Logs



- Amazon Elasticsearch Service
- Amazon CloudWatch Logs subscriptions

Managing, Monitoring & Processing Logs

CloudWatch Logs

- Near real-time, aggregate, monitor, store, and search

Amazon Elasticsearch Service Integration (or ELK stack)

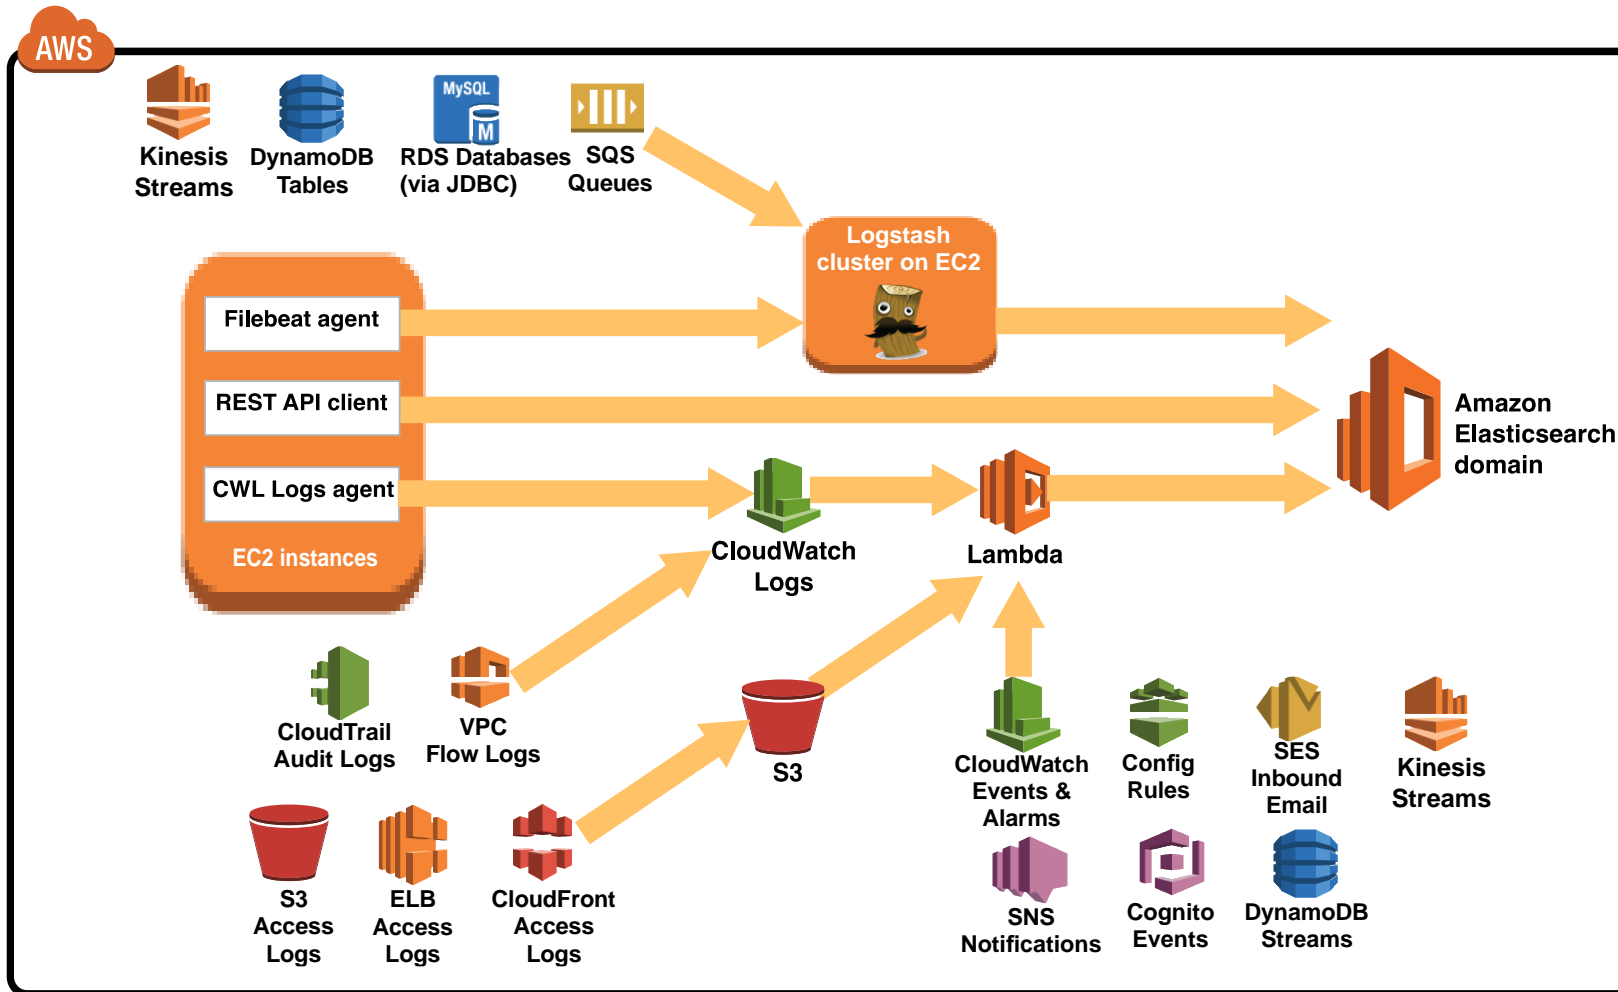
- Analytics and Kibana interface

AWS Lambda & Amazon Kinesis Integration

- Custom processing with your code

Export to S3

- SDK & CLI batch export of logs for analytics



Arrow direction indicates general direction of data flow

AWS Technology Partner solutions integrated with CloudTrail



Architecture Patterns – Monitoring, Logging and Auditing

Requirement

Need to stream logs from Amazon EC2 instances in an Auto Scaling Group

Need to collect metrics from EC2 instances with a 1 second granularity

The application logs from on-premises servers must be processed by AWS Lambda in real time

Solution

Install the unified CloudWatch Agent and collect log files in Amazon CloudWatch

Create a custom metric with high resolution

Install the unified CloudWatch Agent on the servers and use a subscription filter in CloudWatch to connect to a Lambda function

Architecture Patterns – Monitoring, Logging and Auditing

Requirement

CloudWatch Logs entries must be transformed with Lambda and then loaded into Amazon S3

CloudWatch Logs entries must be analyzed and stored centrally in a security account

Access auditing must be enabled and records must be stored for a minimum of 5 years. Any attempts to modify the log files must be identified

Solution

Configure a Kinesis Firehose destination, transform with Lambda and then load into an S3 bucket

Use cross-account sharing and configure a Kinesis Data Stream in the security account to collect the log files then use Lambda to analyze and store

Create a trail in CloudTrail that stores the data in an S3 bucket and enable log file integrity validation

Architecture Patterns – Monitoring, Logging and Auditing

Requirement

API activity must be captured from multiple accounts and stored in a central security account

Need to trace and debug application with distributed components

Solution

Use CloudTrail in each account to record API activity and use cross-account access to a security account to store the log files in a central S3 bucket

Use AWS X-Ray to trace and debug the application

Cost Monitoring

Frameworks



Tools



Best Practices



AWS Bill Development Process



- Consumption Data is aggregated across all linked accounts, based on CloudWatch entries
- RI discounts, Spot Discounts, EDP Discounts, and non-use charges are applied based on the aggregated set of purchases across the linked accounts

Billing Report & Enable Billing Alert



Services ▾

Resource Groups ▾



Global ▾

Support ▾

- Dashboard
- Bills
- Cost Explorer
- Budgets
- Reports
- Cost Allocation Tags
- Payment Methods
- Payment History
- Consolidated Billing
- Preferences
- Credits
- Tax Settings

Preferences ?

Billing Preferences

Receive PDF Invoice By Email

Turn on this feature to receive a PDF version of your invoice by email. Invoices are generally available within the first three days of the month.

Cost Management Preferences

Receive Free Tier Usage Alerts

Turn on this feature to receive email alerts when your AWS service usage is approaching, or has exceeded, the AWS Free Tier usage limits. If you wish to receive these alerts at an email address that is not the primary email address associated with this account, please specify the email address below.

Email Address:

Receive Billing Alerts

Turn on this feature to monitor your AWS usage charges and recurring fees automatically, making it easier to track and manage your spending on AWS. You can set up billing alerts to receive email notifications when your charges reach a specified threshold. Once enabled, this preference cannot be disabled. [Manage Billing Alerts](#) or [try the new budgets feature!](#)

Receive Billing Reports

Turn on this feature to receive ongoing reports of your AWS charges once or more daily. AWS delivers these reports to the Amazon S3 bucket that you specify where indicated below. For consolidated billing customers, AWS generates reports only for paying accounts. Linked accounts cannot sign up for billing reports.

Save to S3 Bucket:

Verify

Note: You must apply appropriate permissions to your S3 bucket [sample policy](#)

Save preferences



SCHOOL OF ELECTRONICS AND TELECOMMUNICATIONS



Budgets Management

Create budget



Budget details

Name*

Include costs related to

Tag


Project

Period Monthly

Start date*

End date*

Monthly amount*



Billing Alarms

Notifications (optional)

You can create a billing alarm to receive e-mail alerts when your current or forecasted AWS charges exceed a threshold you choose.

Budget notifications are enabled and processed as Amazon CloudWatch Alarms. New and existing CloudWatch customers will receive 10 free CloudWatch Alarms each month; each additional alarm will be charged at \$0.10 per month. To view more information on pricing, please visit [Amazon CloudWatch Pricing](#). ✕

Notify me when costs exceed % of budgeted costs

Notify me when costs exceed % of budgeted costs ✕

+ Add new alert

Email contacts*


Billing Alert Response

Respond to billing alerts in CloudWatch.


When an alarm is triggered:

1. Email the project team, and the budget approver (AWS console)
2. Open a Service Management Ticket in your ITSM system
3. Open a Ticket in your Finance System
4. Resize the project's IT resources
5. Cull/Reduce the projects IT resources
6. Shut down the projects IT resources

Tools to Manage Billing Data



Using Trusted Advisor



Trusted Advisor

Cost Optimization



5 2 0

\$1,223.92

Potential monthly savings

Performance



8 0 0

Security



6 1 4

Fault Tolerance



13 0 2

Cost Monitoring Summary

- Setup Sensible Billing Alarms for your Organization
- Proactively Monitor Your Account Billing Usage
- Leverage AWS Partner tools
- Leverage Trusted Advisor reports to:
 - Identify Idle Resources and Turn Off Unused Instances
 - Identify Under-utilized Resources and Resize them
 - Identify Baseline Consumption needs to support RI commitments
- Review new Discount and Technology Options on a Monthly basis



AWS Architecting Best Practice



VIỆN ĐIỆN TỬ - VIỄN THÔNG
School of Electronics and Telecommunications



Architecting Approaches for AWS

Lift-and-shift

- **Deploy existing apps in AWS with minimal re-design**
- Good strategy if starting out on AWS, or if application can't be re-architected due to cost or resource constraints
- Primarily use core services such as EC2, EBS, VPC

Cloud-optimized

- **Evolve architecture for existing app to leverage AWS services**
- Gain cost and performance benefits from using AWS services such as Auto Scaling Groups, RDS, SQS, and so on

Cloud-native architecture

- **Architect app to be cloud-native from the outset**
- Leverage the full AWS portfolio
- Truly gain all the benefits of AWS (security, scalability, cost, durability, low operational burden, etc)

Cloud Architecture Best Practices

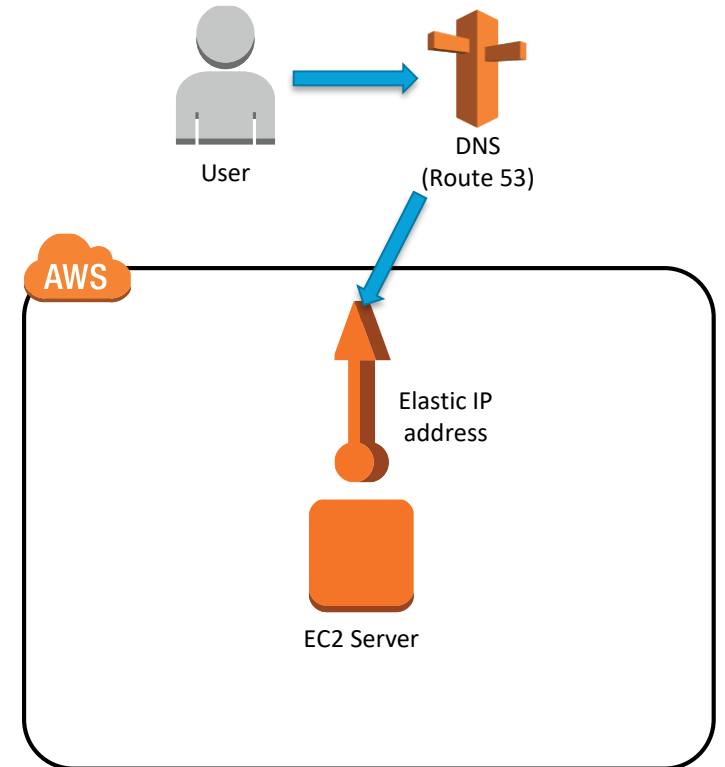
1. Design for failure and nothing fails
2. Build security in every layer
3. Leverage different storage options
4. Implement elasticity
5. Think parallel
6. Loose coupling sets you free
7. Don't fear constraints

Design for Failure and Nothing Fails

Design for Failure: A Single User

Single Points of Failure:

- A single Elastic IP
 - Gives a server a static Public IP address
- A single Amazon Elastic Compute Cloud (EC2) instance
 - Full stack on single host
 - Web application
 - Database
 - Management, etc...




Design for Failure and Nothing Fails

Design for Failure: Difficulties Scaling to Many Users

We could potentially get to a few hundred to a few thousand users depending on application complexity and traffic, but...

There may be difficulty scaling to many more users due to:

- **All eggs in one basket**
- **No failover or redundancy**

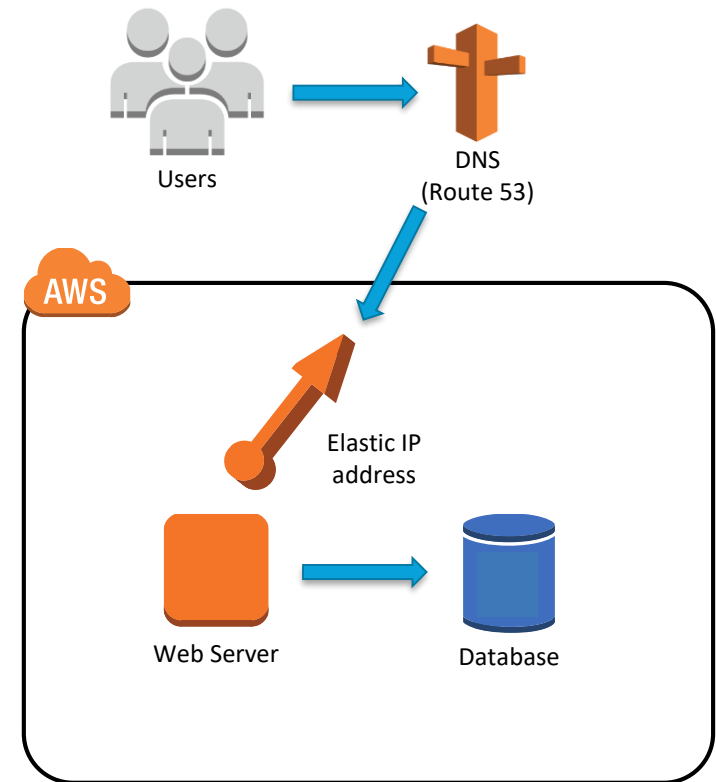


Design for Failure and Nothing Fails

Design for Failure: Solving “All Eggs in One Basket”

Separate single EC2 Server into web and database tiers:

- Web Server on EC2
- Database on EC2 or RDS
 - Amazon Relational Database Service (RDS) can take care of management overhead such as patching, backups, and failure detection

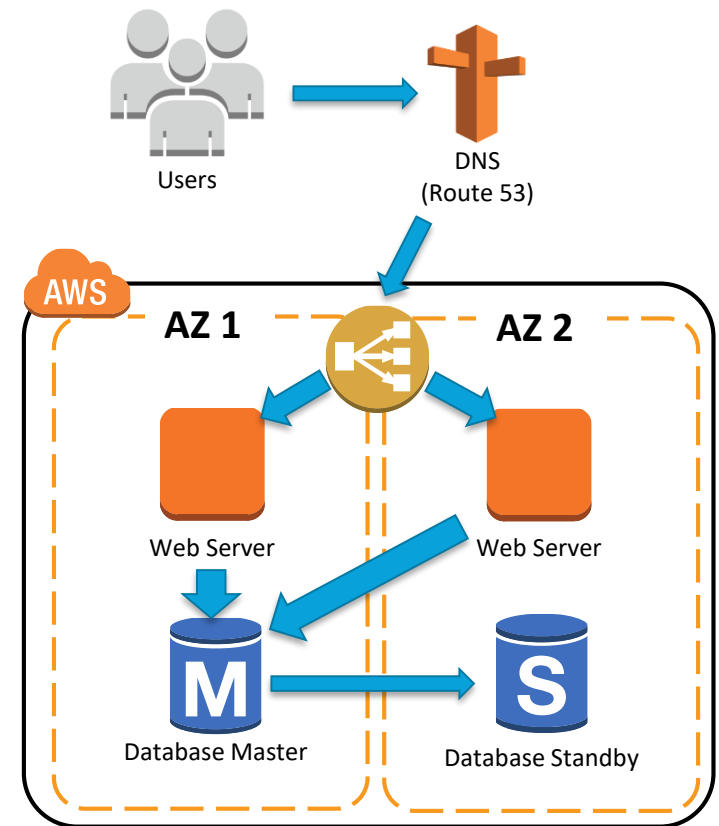


Design for Failure and Nothing Fails

Design for Failure: Solving No Failover/Redundancy

Leverage **multiple Availability Zones** for redundancy and high availability.

- Use an **Elastic Load Balancer (ELB)** across AZs for availability and failover
- If using RDS, use the Multi-AZ feature for managed **replication** and a **standby** instance
 - If not, use failover and replication features native to your database engine

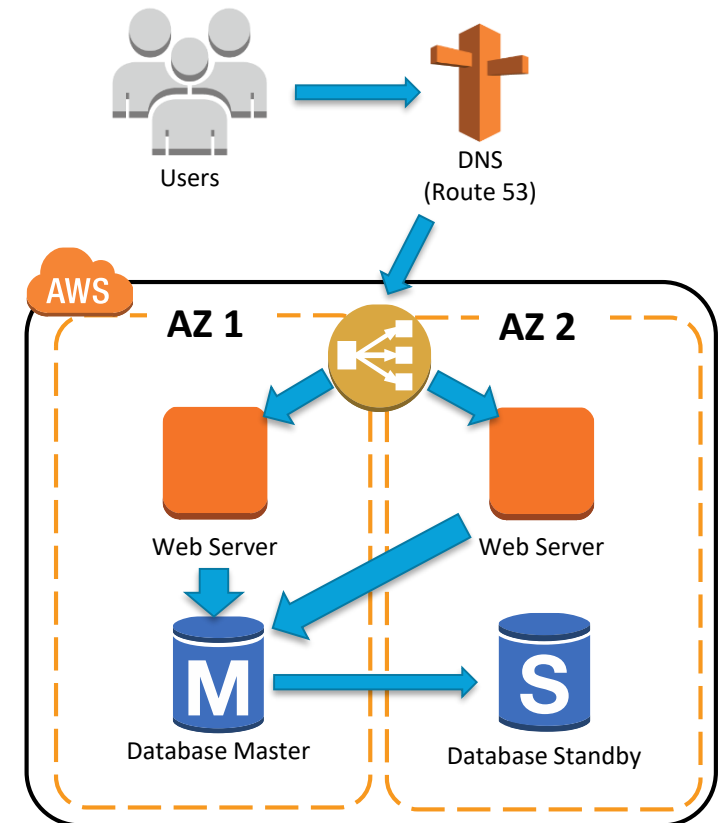


Design for Failure and Nothing Fails

Design for Failure: Best Practices

Best Practices:

- Eliminate single points of failure
- Use multiple Availability Zones
- Use Elastic Load Balancing
- Do real-time monitoring with CloudWatch
- Create a database standby across Availability Zones



Design for Failure and Nothing Fails

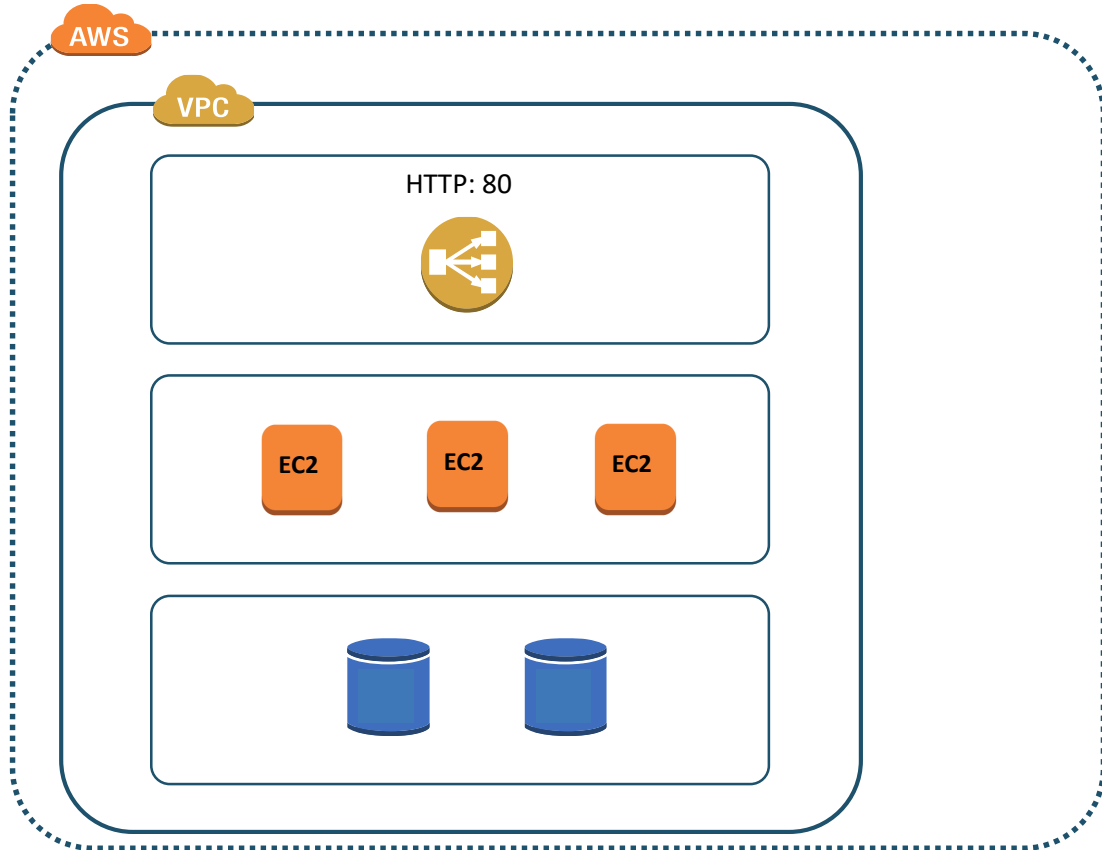

Avoid single points of failure

Assume **everything fails** and design backwards

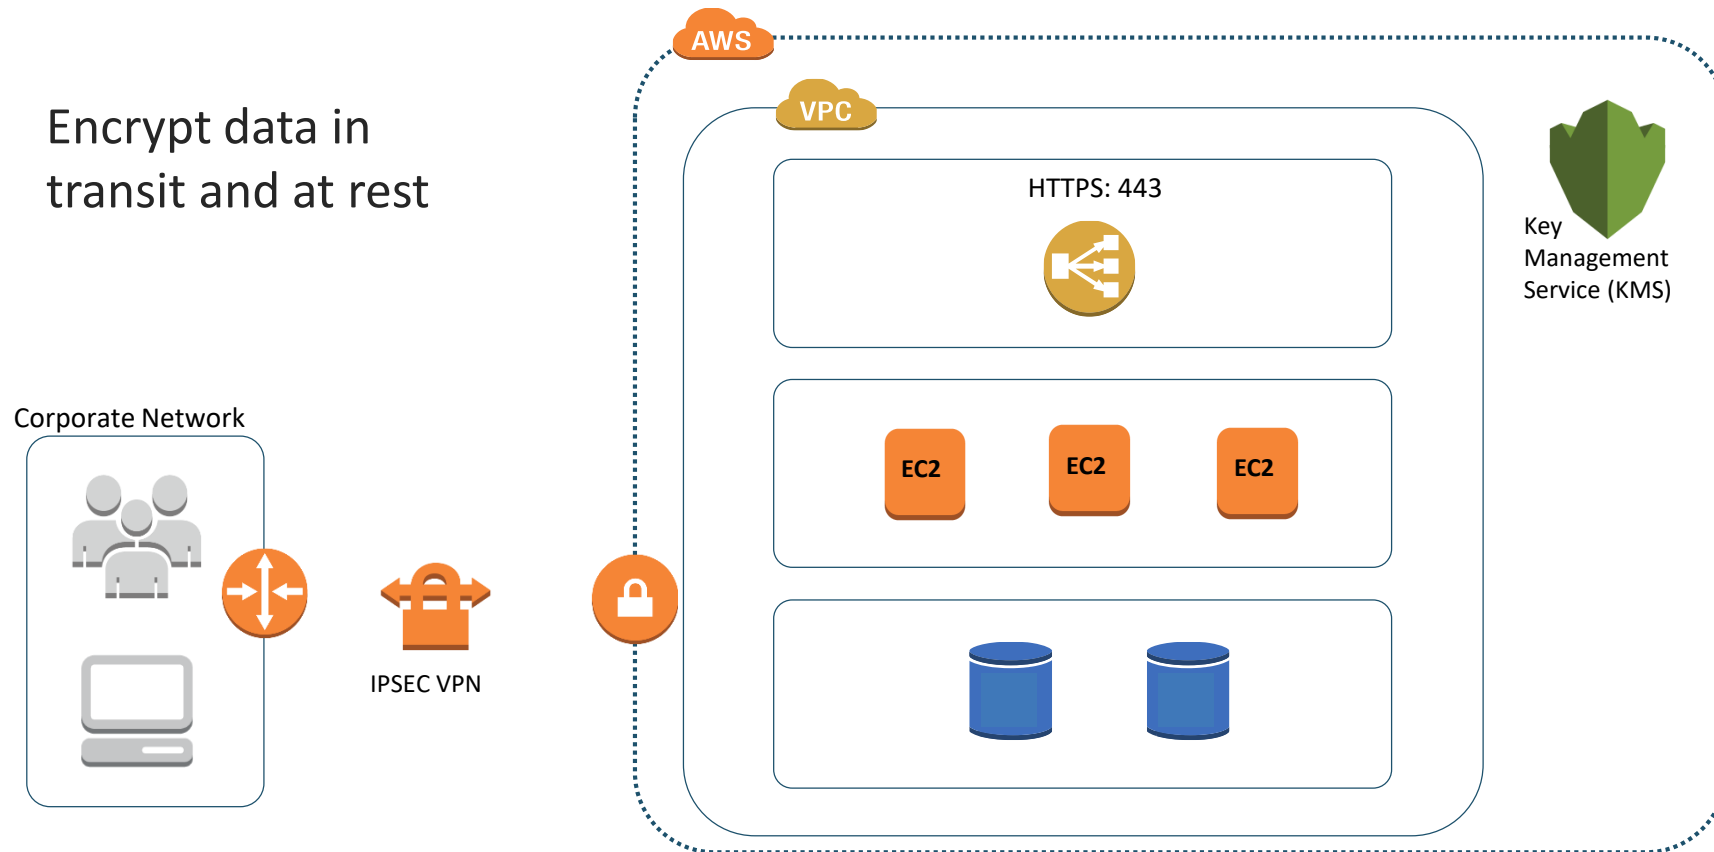
- When, not if, an individual component fails, the application does not fail
 - Think of your servers as **cattle, not pets**
- Leverage Route 53 DNS **Pilot-light** or **Warm-standby** strategies to implement Disaster Recovery
- **Auto Scaling groups** can be used to detect failures and self-heal, thus protecting against AZ level outages

Build Security in Every Layer

Corporate Network

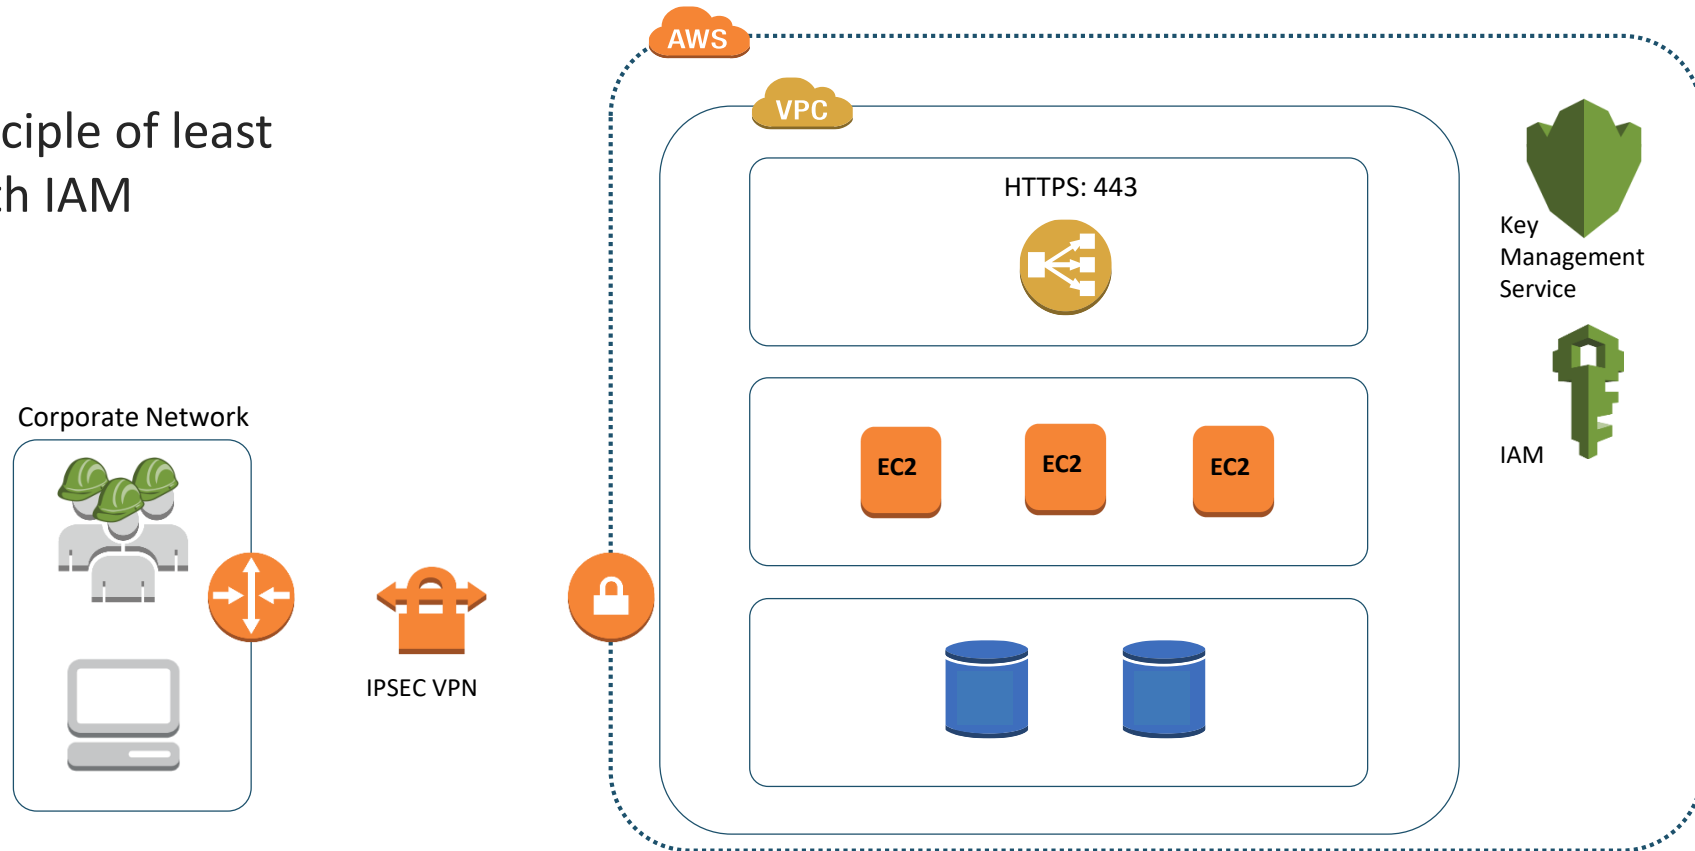


Build Security in Every Layer




Build Security in Every Layer

Enforce principle of least privilege with IAM



Build Security in Every Layer

Create firewall rules with Security Groups and NACLs



Build Security in Every Layer

More Tools for your Security Toolbox:

- Amazon Inspector
- Amazon Certificate Manager
- AWS Shield
- AWS Web Application Firewall (WAF)
- Amazon Macie
- Amazon GuardDuty
- AWS Config

Leverage Many Storage Options

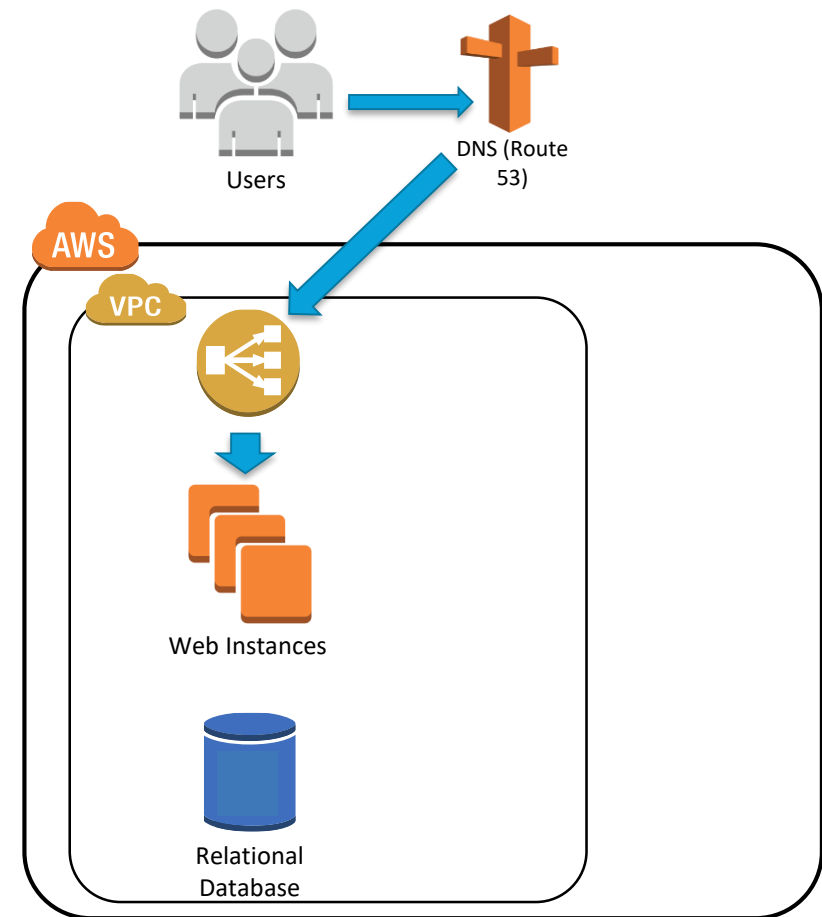
One size does NOT fit all

- **Amazon Elastic Block Storage (EBS)** – persistent block storage
- **Amazon EC2 Instance Storage** – ephemeral block storage
- **Amazon RDS** – managed relational database
- **Amazon CloudFront** – content distribution network
- **Amazon S3** – object/blob store, good for large objects
- **Amazon DynamoDB** – non-relational data (key-value)
- **Amazon ElastiCache** – managed Redis or Memcached

Leverage Many Storage Options

Current State:

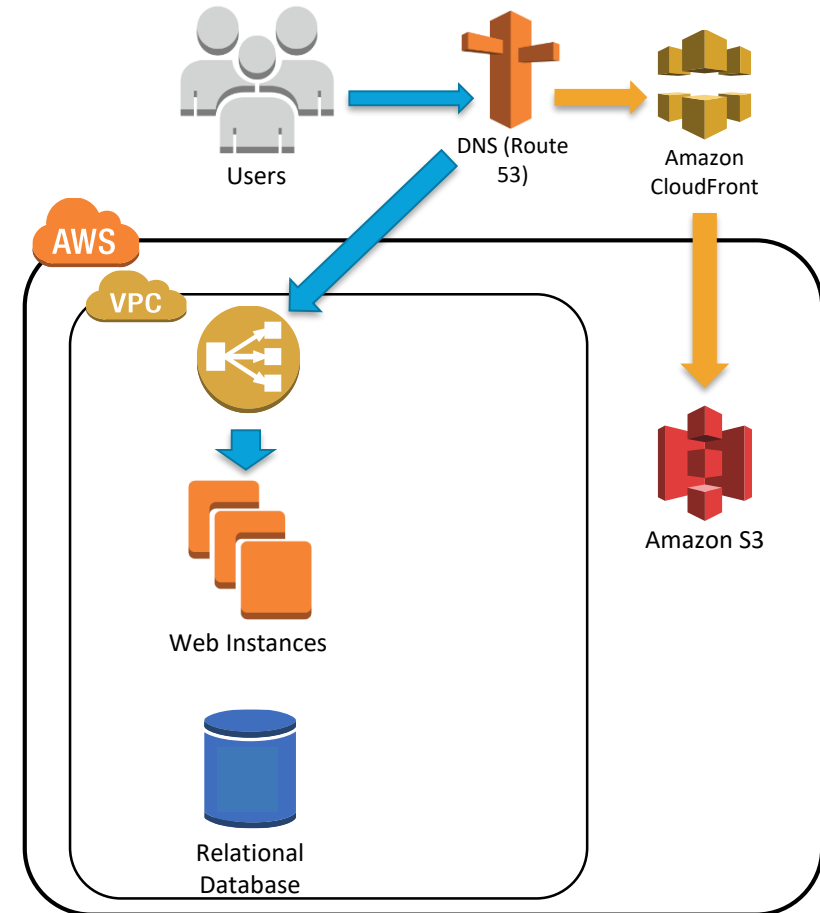
- All load handled by one stack
 - Elastic Load Balancer (ELB)
 - EC2 Web App cluster
 - Relational Database
- No caching layer(s)
- All persistent data in database or Web instances' Elastic Block Storage (EBS) volumes



Leverage Many Storage Options

Offload and cache requests for static assets:

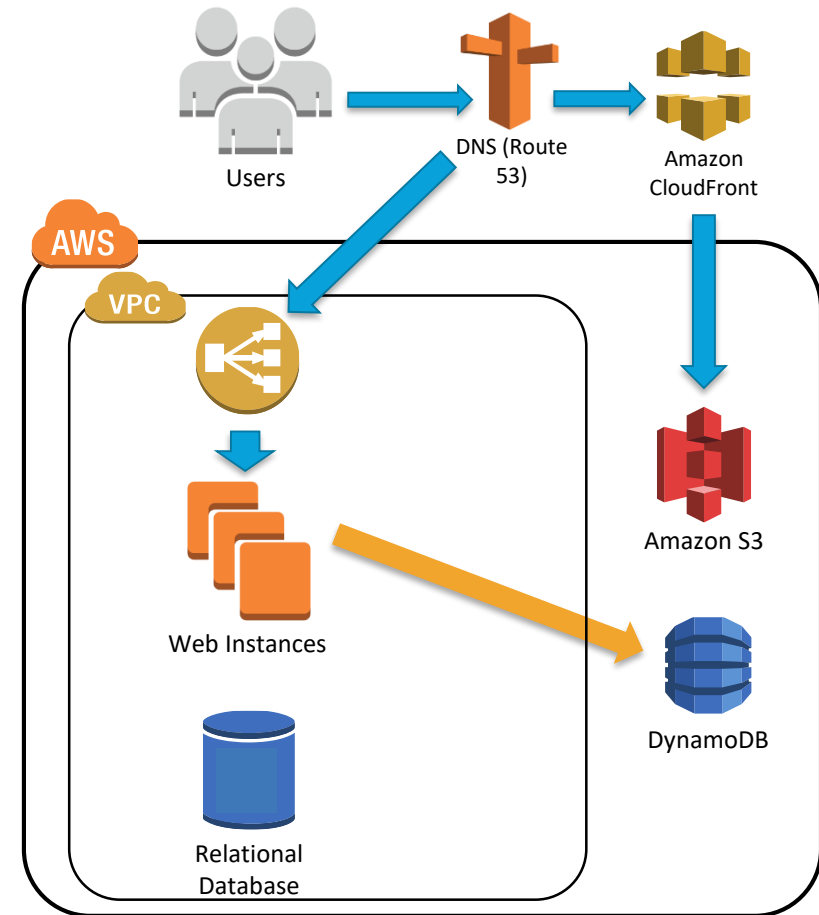
- Store large/static objects in Simple Storage Service (S3)
- Use a Content Delivery Network (CDN) like CloudFront to cache responses using points of presence all around the world



Leverage Many Storage Options

Save user session data in a database to avoid interrupting the user experience if a web host becomes unresponsive:

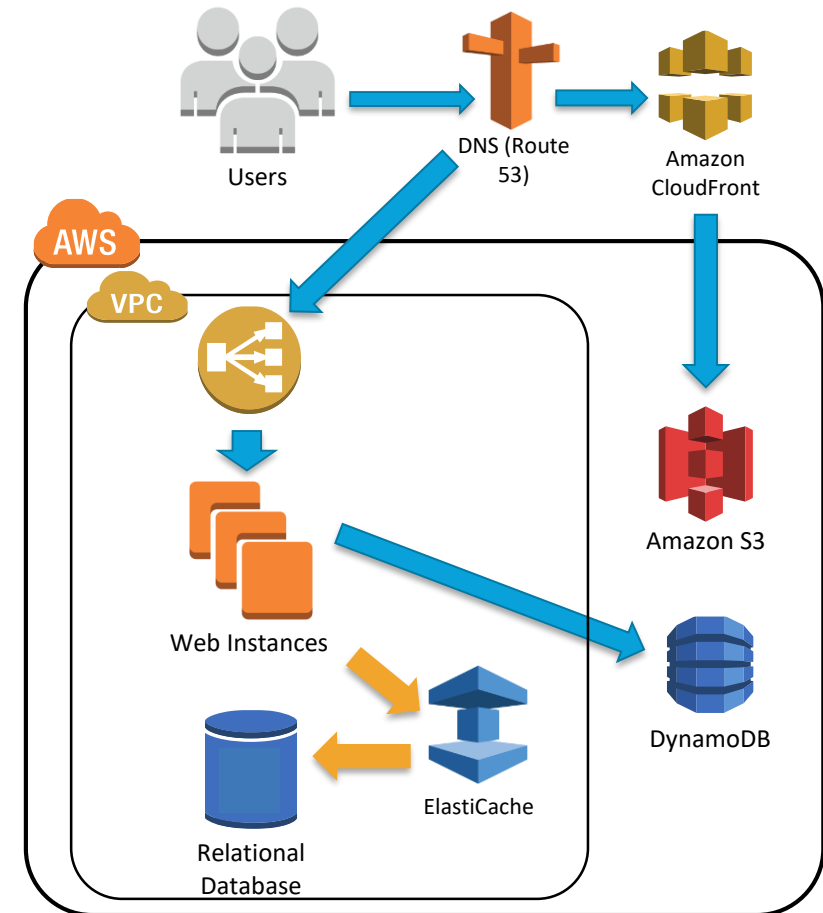
- Store session/state data in **DynamoDB**, a managed NoSQL key-value store



Leverage Many Storage Options

Cache frequent queries to shift the load off of your database:

- Put **ElastiCache** as a caching layer between the web hosts and the database



Implement Elasticity

How:

- Write **Auto Scaling policies** with your specific application access patterns in mind
- Prepare your application to **be flexible**: don't assume the health, availability, or fixed location of components
- Architect **resiliency to reboot** and relaunch
 - When an instance launches, it should ask *"Who am I and what is my role?"*
- Leverage highly **scalable, managed services** such as S3 and DynamoDB

Think Parallel

Scale Horizontally, Not Vertically

- Decouple compute from state/session data
- Use ELBs to distribute load
- Break up big data into pieces for distributed processing
 - AWS Elastic Map Reduce (EMR) – managed Hadoop

Think Parallel

Faster doesn't need to mean more expensive

- With EC2 On Demand, the following will cost the same:
 - 12 hours of work using 4 vCPUs
 - 1 hour of work using 48 vCPUs
- Right Size your infrastructure to your workload to get the best balance between cost and performance


Think Parallel

Parallelize using native managed services

- Get the best performance out of S3 with parallelized reads/writes
 - Multi-part uploads (API) and byte-range GETs (HTTP)
- Take on high concurrency with Lambda
 - Initial soft limit: 1000 concurrent requests per region

Loose Coupling Sets You Free: Queueing

Use Amazon Simple Queue Service (SQS) to pass messages between loosely coupled components




Loose Coupling Sets You Free: Don't Reinvent the Wheel

Nearly everything in AWS is an API call.

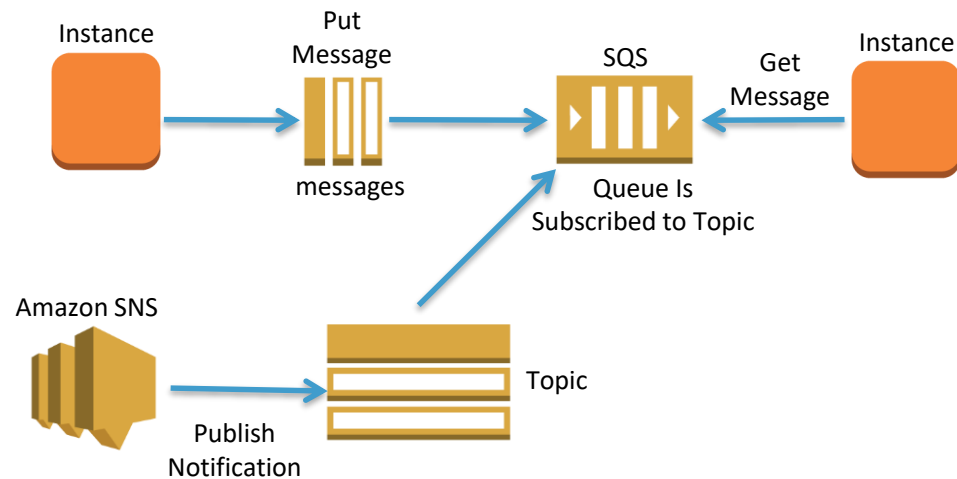
Leverage AWS Native Services for...

- Queuing
- Transcoding
- Search
- Databases
- Email
- Monitoring
- Metrics
- Logging
- Compute




Loose Coupling Sets You Free

Using **SNS** and **SQS** to asynchronously scale:



Using **Lambda** triggers to decouple actions:



Don't Fear Constraints

Rethink traditional architectural constraints

Need more RAM?

- Don't: vertically scale
- Do: distribute load across machines or a shared cache

Need better IOPS for database?

- Don't: rework schema/indexes or vertically scale
- Do: create read replicas, implement sharding, add a caching layer

Hardware failed or config got corrupted?

- Don't: waste production time diagnosing the problem
- Do: "Rip and replace" – stop/terminate old instance and relaunch

Need a Cost Effective Disaster Recovery (DR) strategy?

- Don't: double your infrastructure costs when you don't need to
- Do: implement Pilot Light or Warm Standby DR stacks

Q & A



VIỆN ĐIỆN TỬ - VIỄN THÔNG
School of Electronics and Telecommunications



Module 4: Identify & Management



VIỆN ĐIỆN TỬ - VIỄN THÔNG
School of Electronics and Telecommunications




Agenda

- AWS Account and Organizations
- Identity Management and Permissions
- Deployment and Management

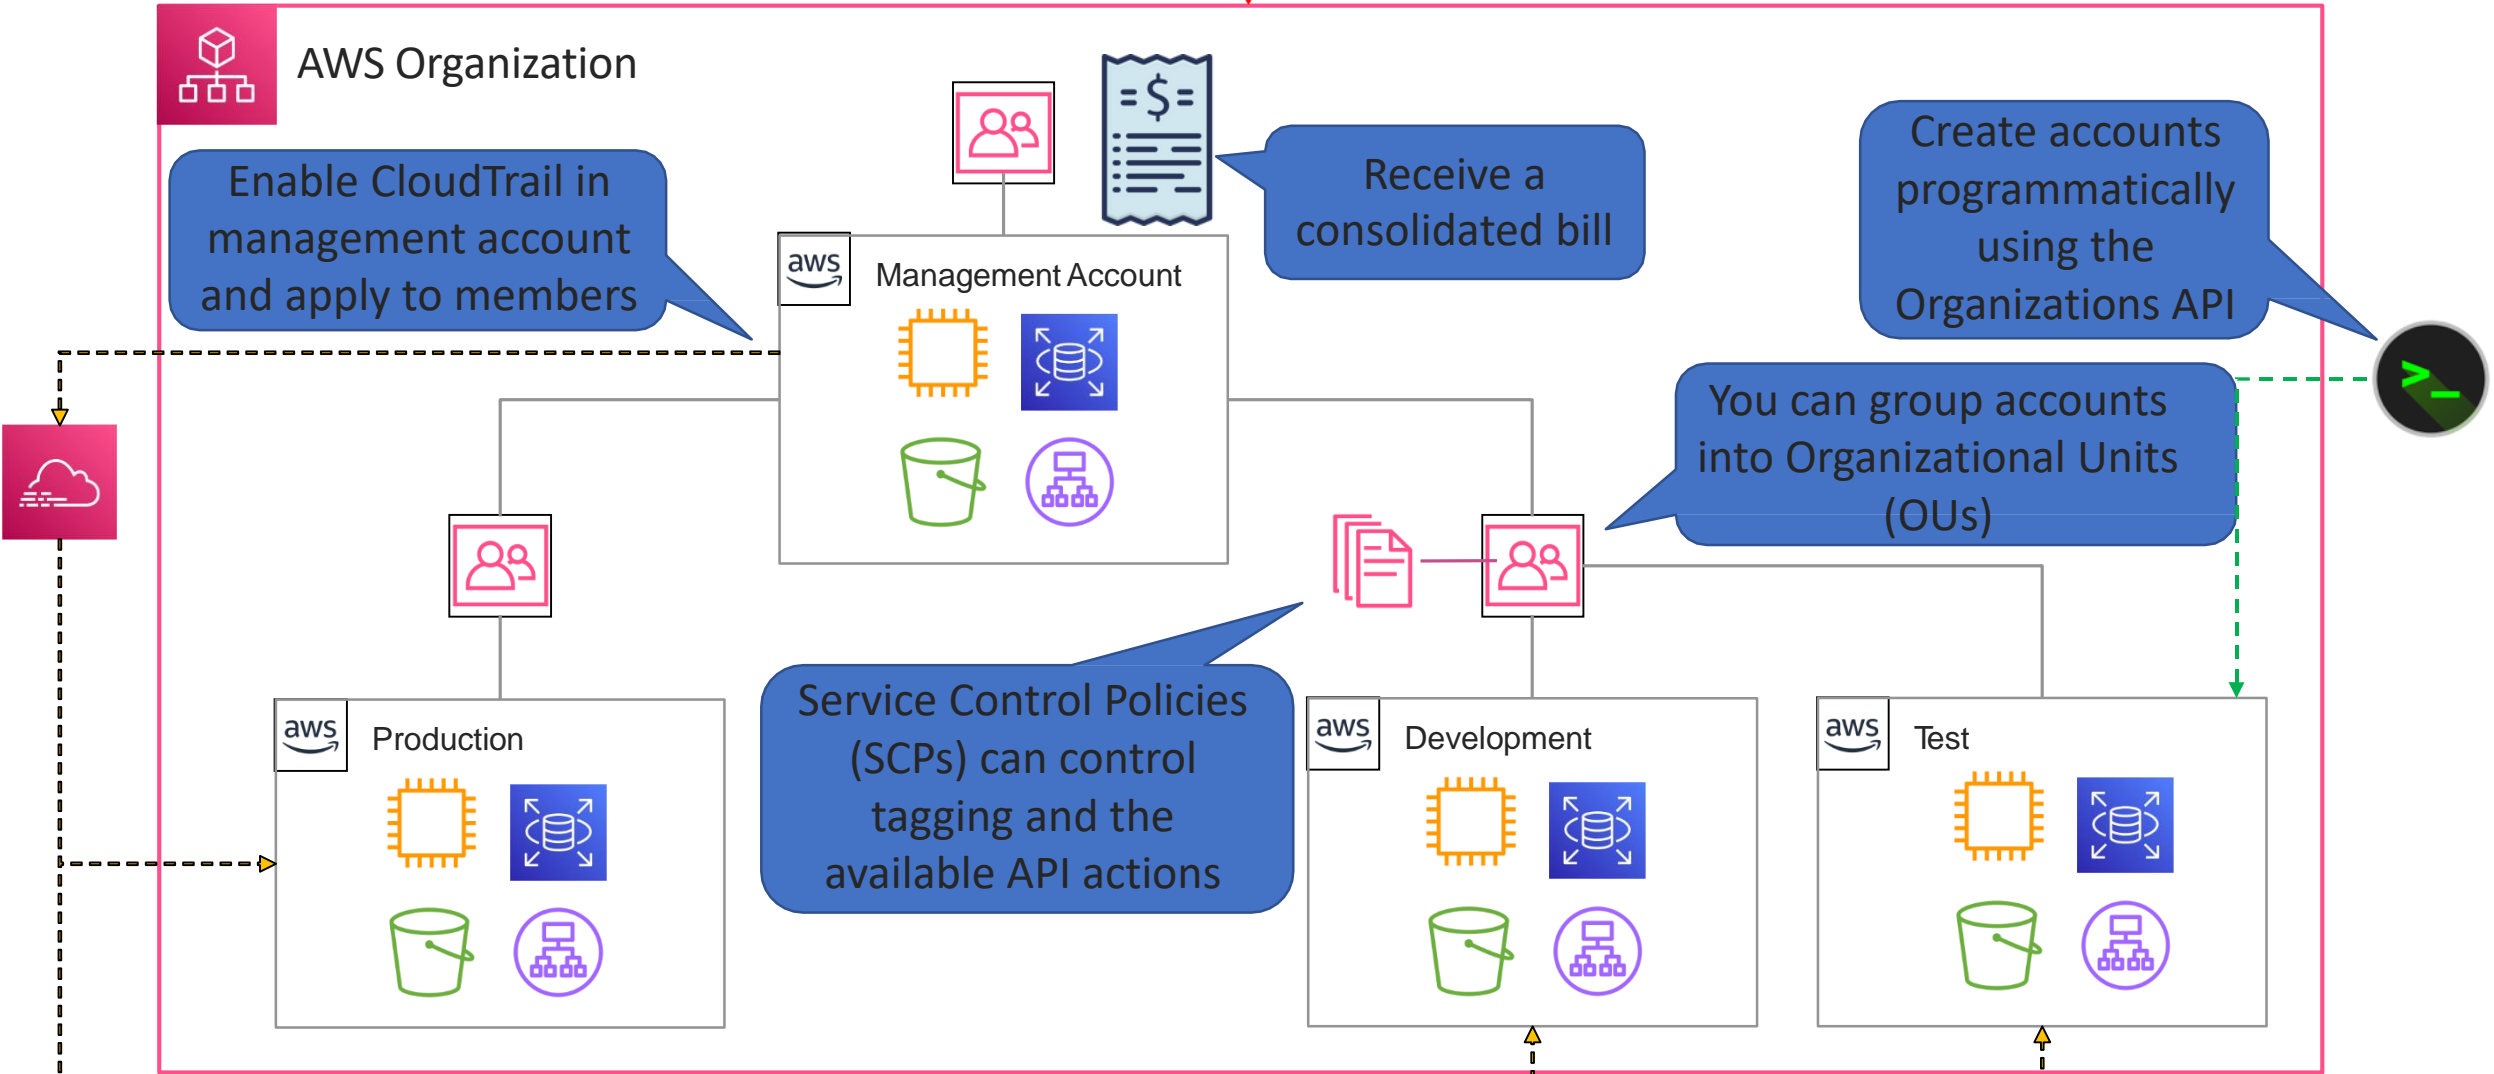
AWS Account and Organizations



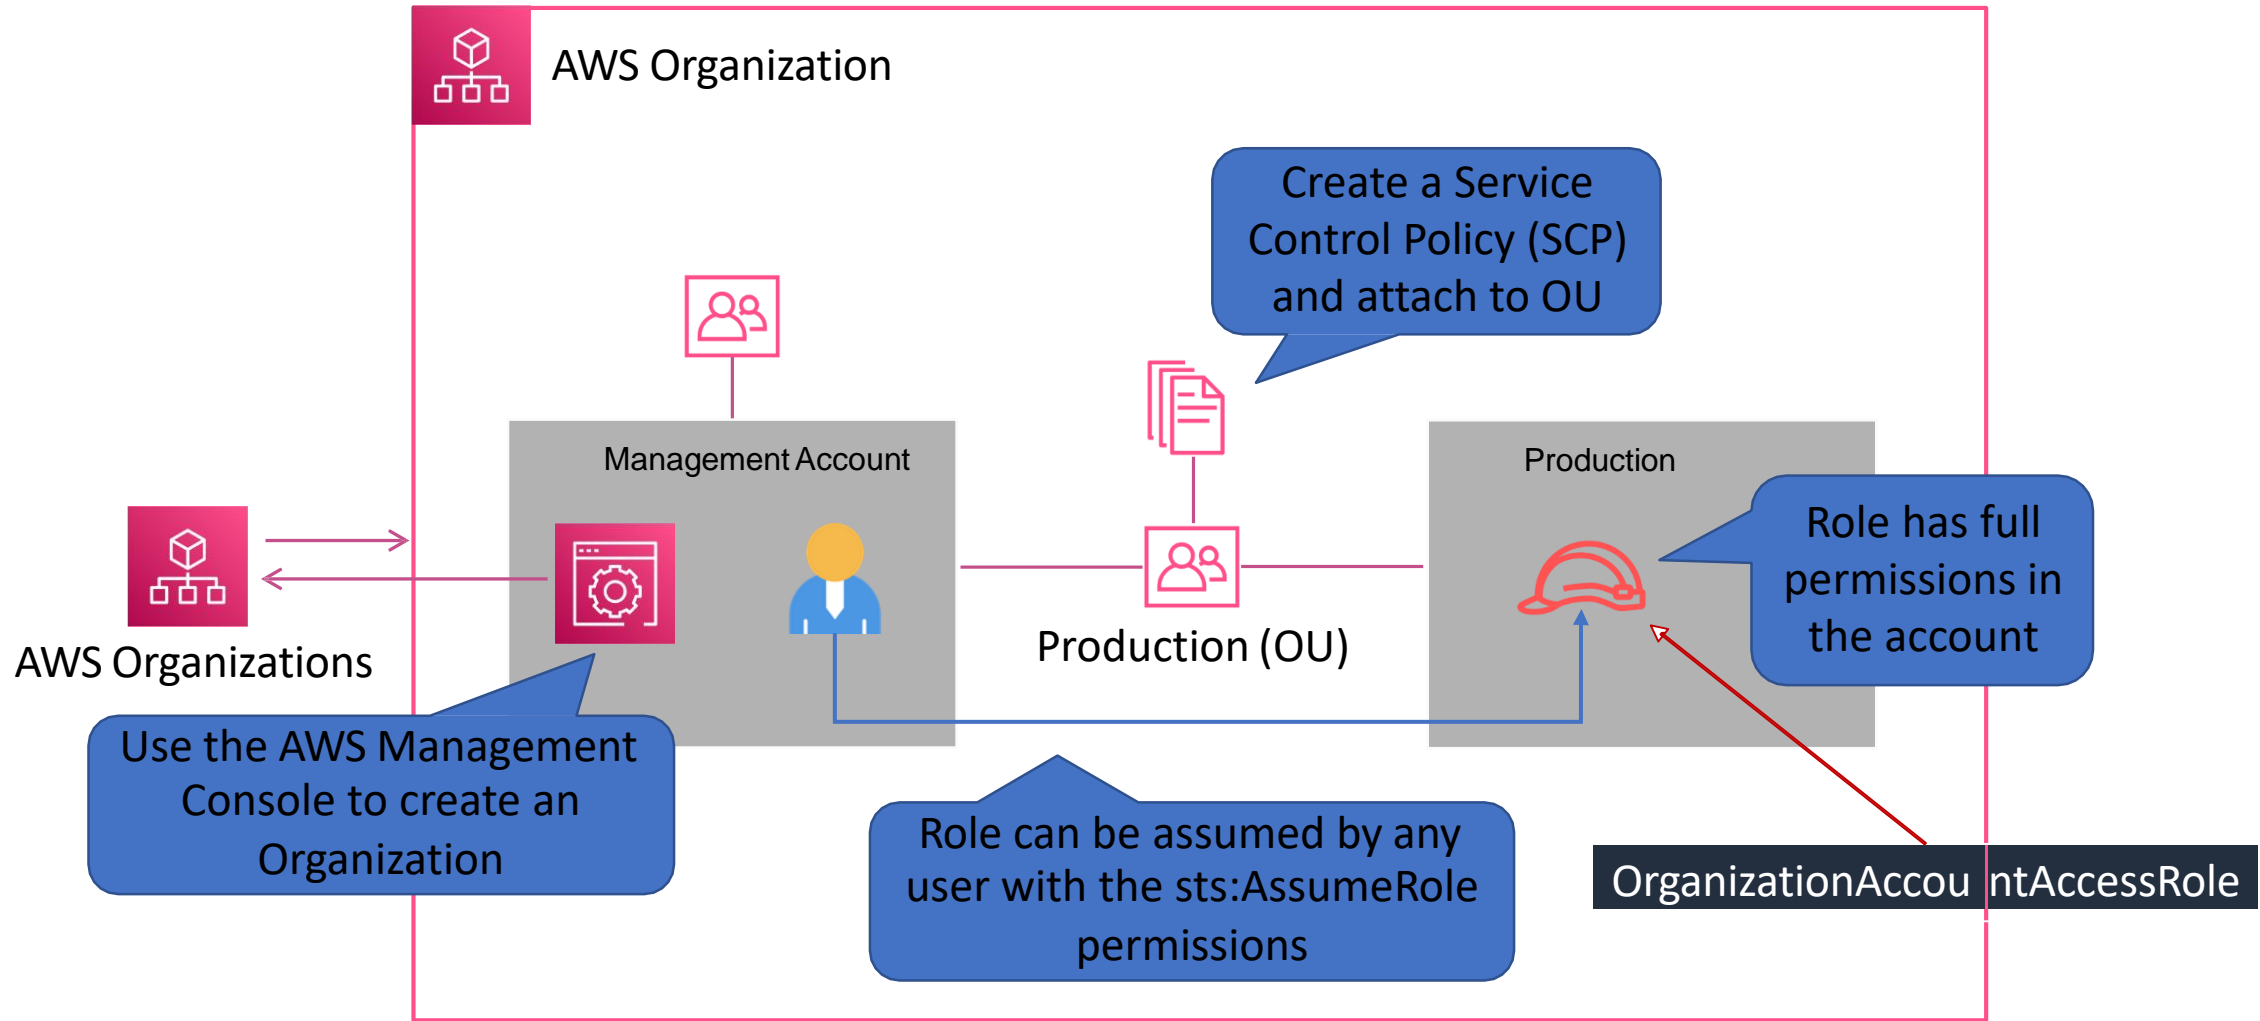
AWS Organizations



Enable AWS SSO using on-prem directory

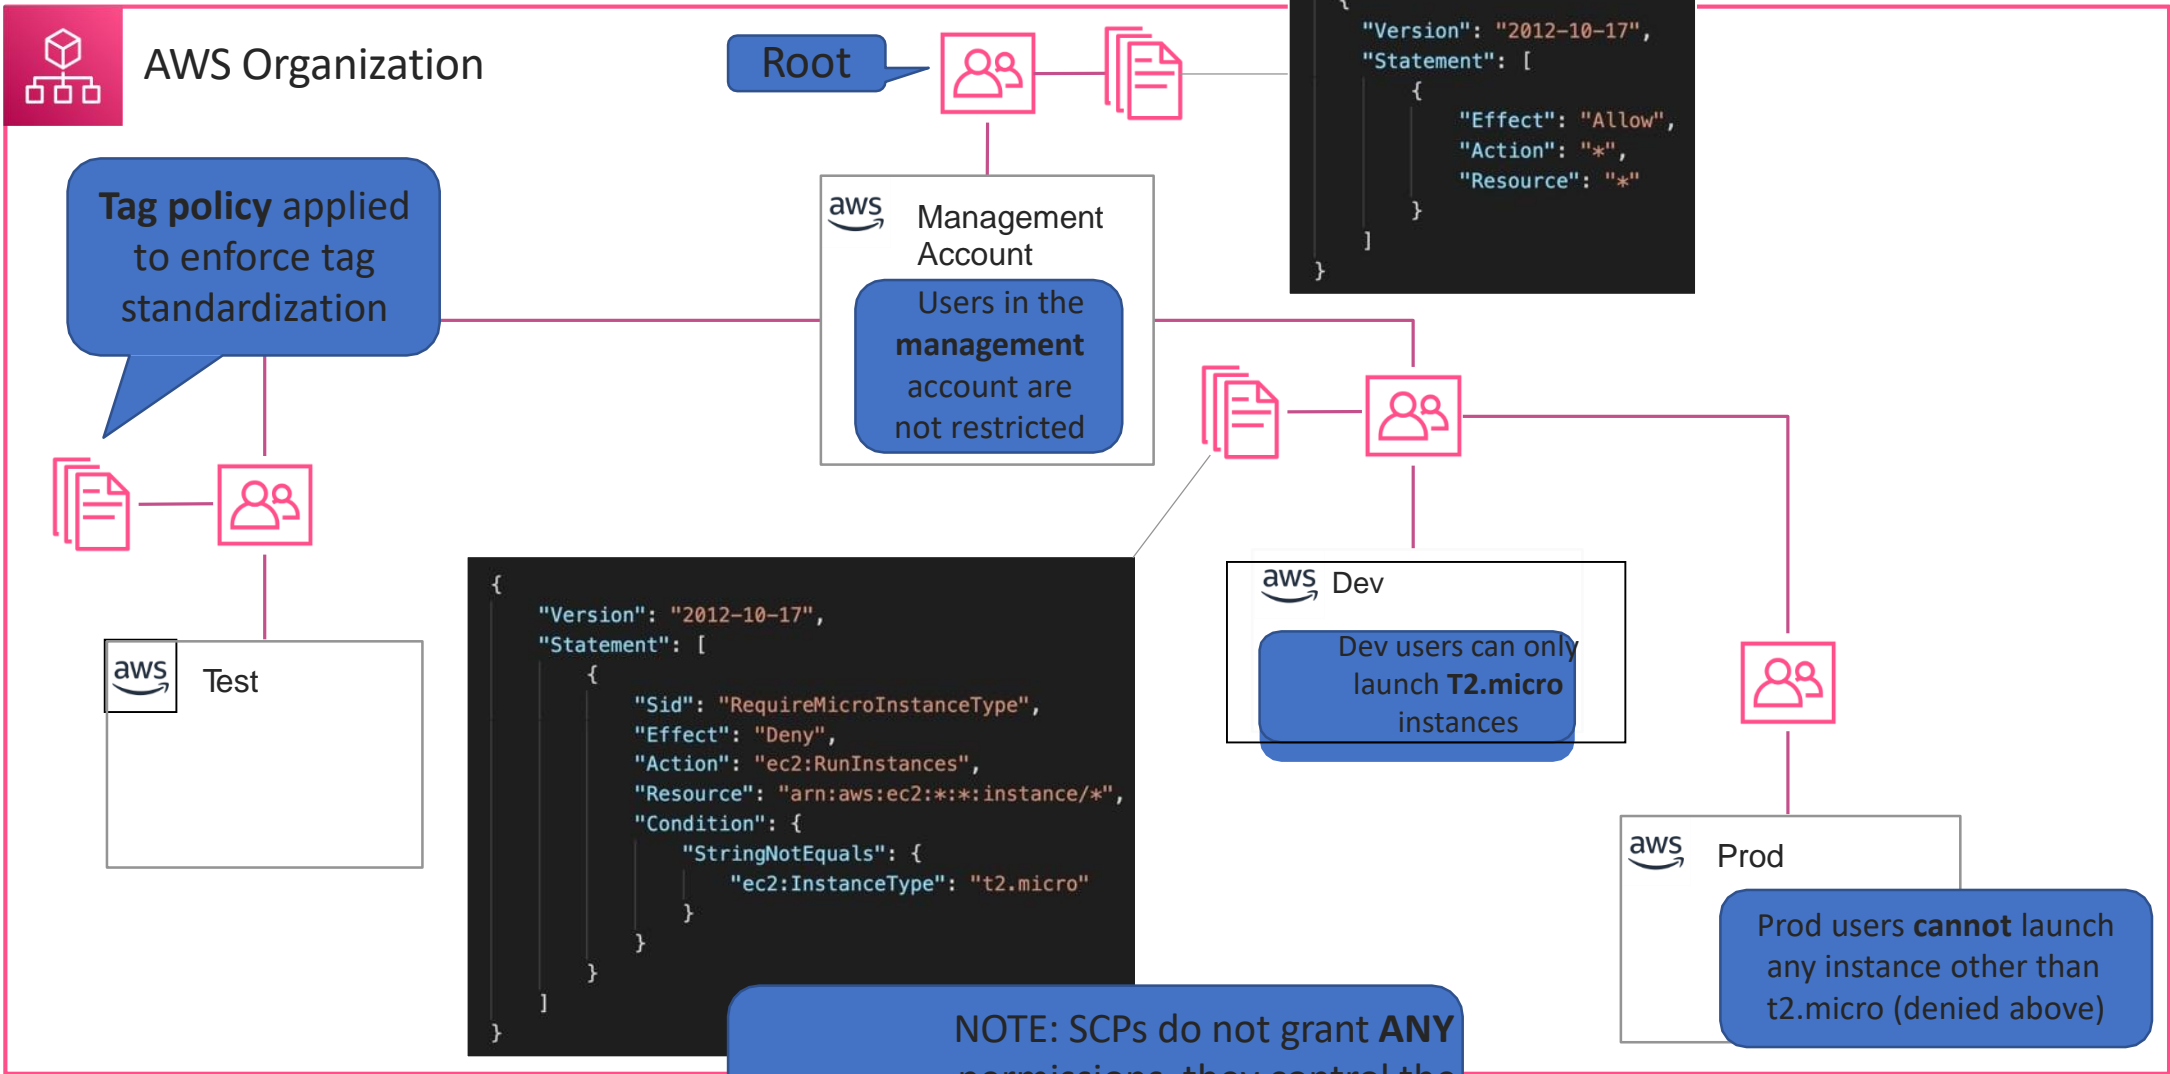


Account Configuration



Service Control Policies

SCPs control the maximum available permissions

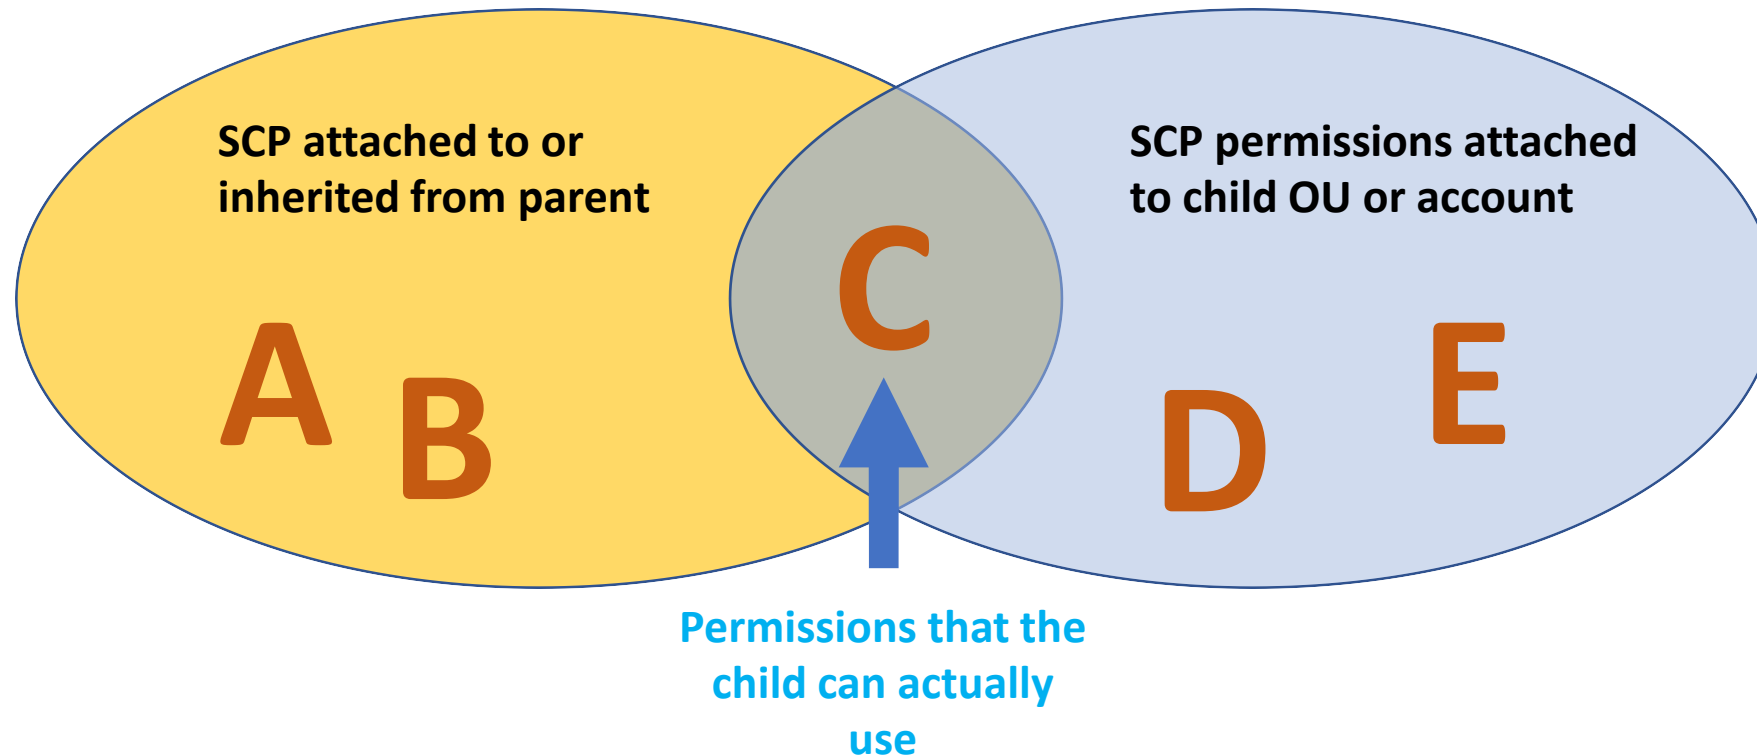


NOTE: SCPs do not grant ANY permissions, they control the AVAILABLE permissions

SCP Strategies and Inheritance



SCP Strategies and Inheritance



SCP Strategies and Inheritance

Deny List Strategy

- The `FullAWSAccess` SCP is attached to every OU and account
- Explicitly allows all permissions to flow down from the root
- Can explicitly override with a deny in an SCP
- This is the default setup

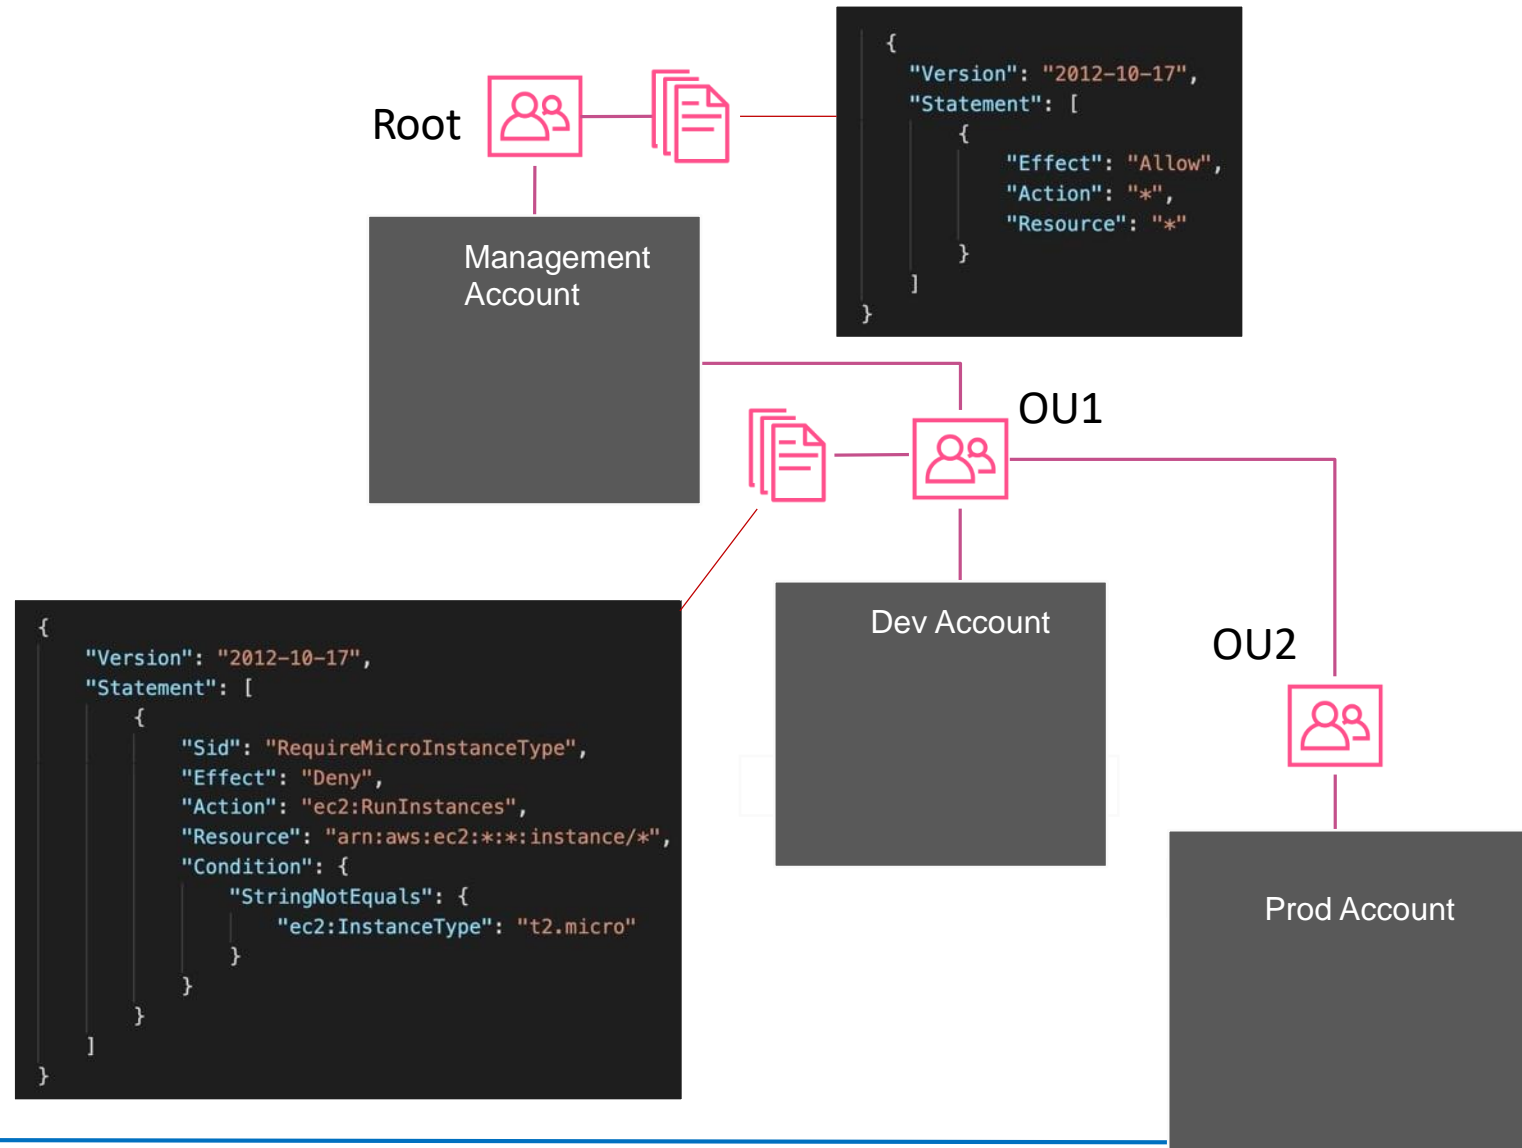
Note: An **explicit deny** overrides any kind of **allow**

Allow List Strategy

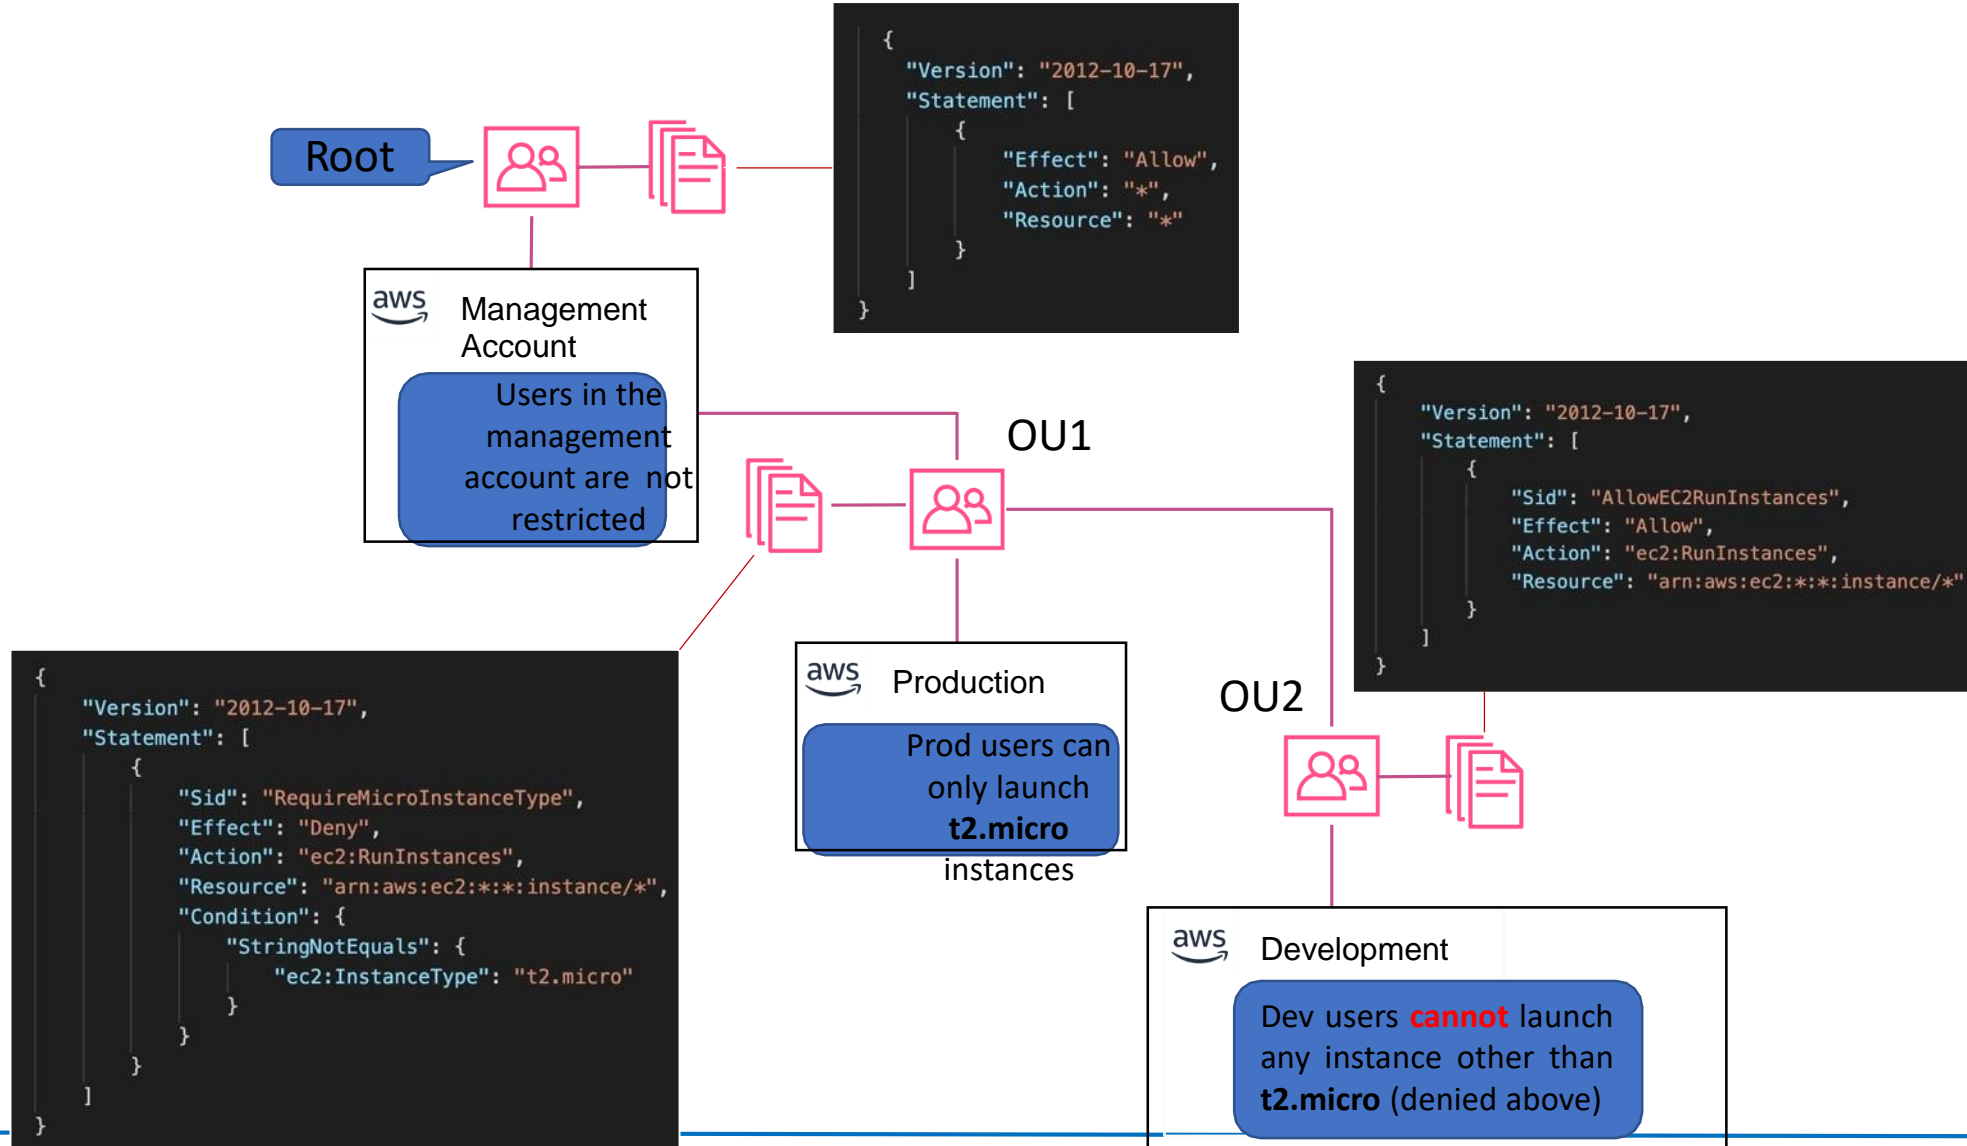
- The `FullAWSAccess` SCP is **removed** from every OU and account
- To allow a permission, SCPs with allow statements must be added to the account and every OU above it including root
- Every SCP in the hierarchy must explicitly allow the APIs you want to use

Note: An **explicit allow** overrides an **implicit deny**

Service Control Policies



Service Control Policies



Identity Management and Permissions



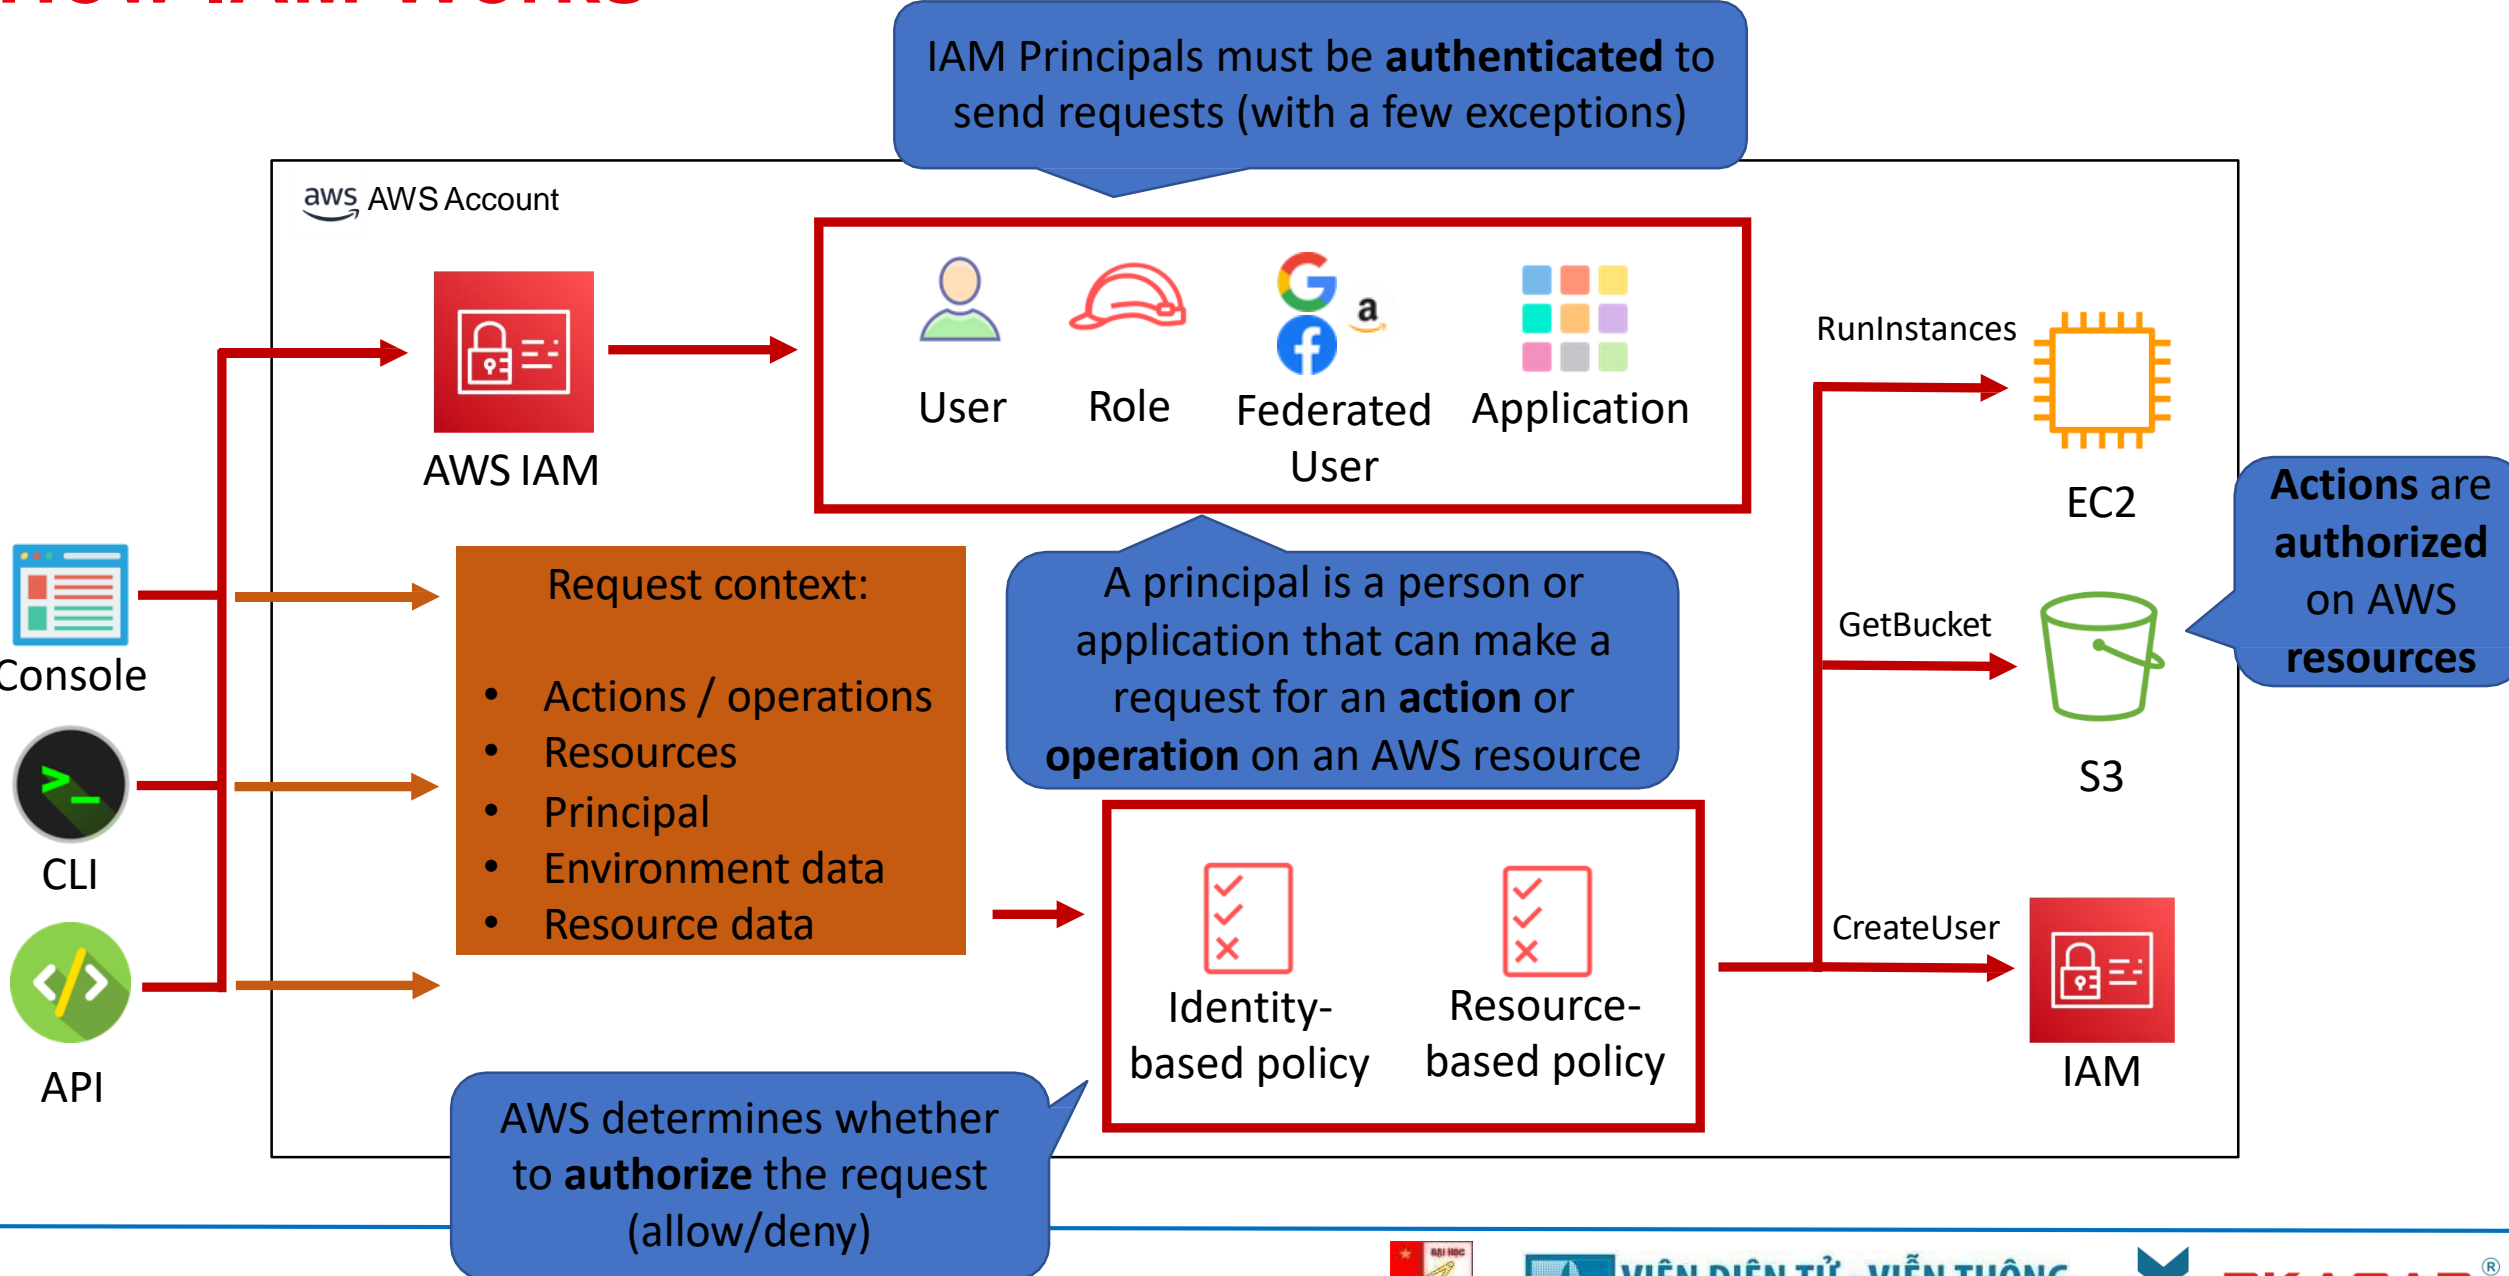
How IAM Works



VIỆN ĐIỆN TỬ - VIỄN THÔNG
School of Electronics and Telecommunications



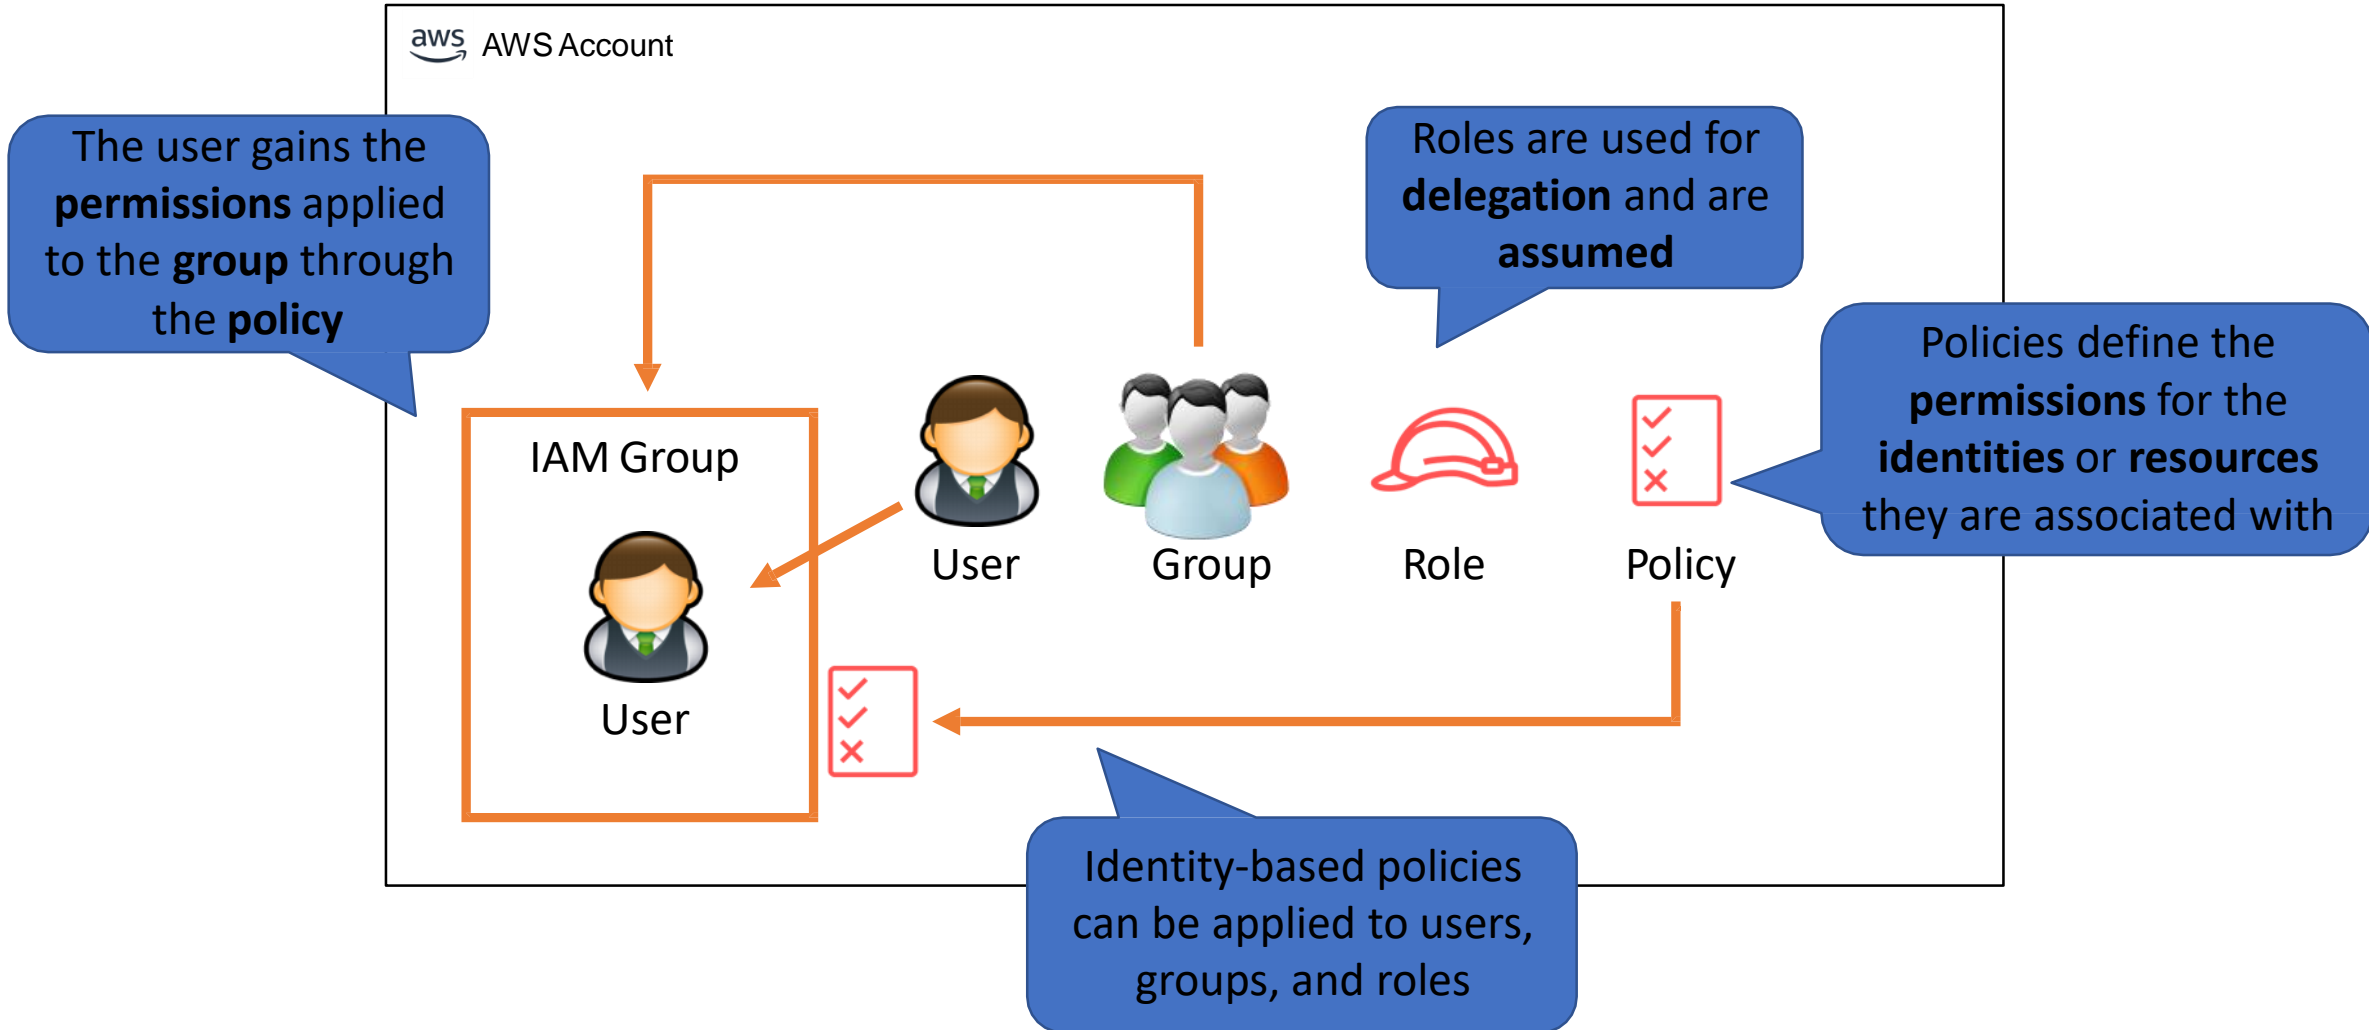
How IAM Works



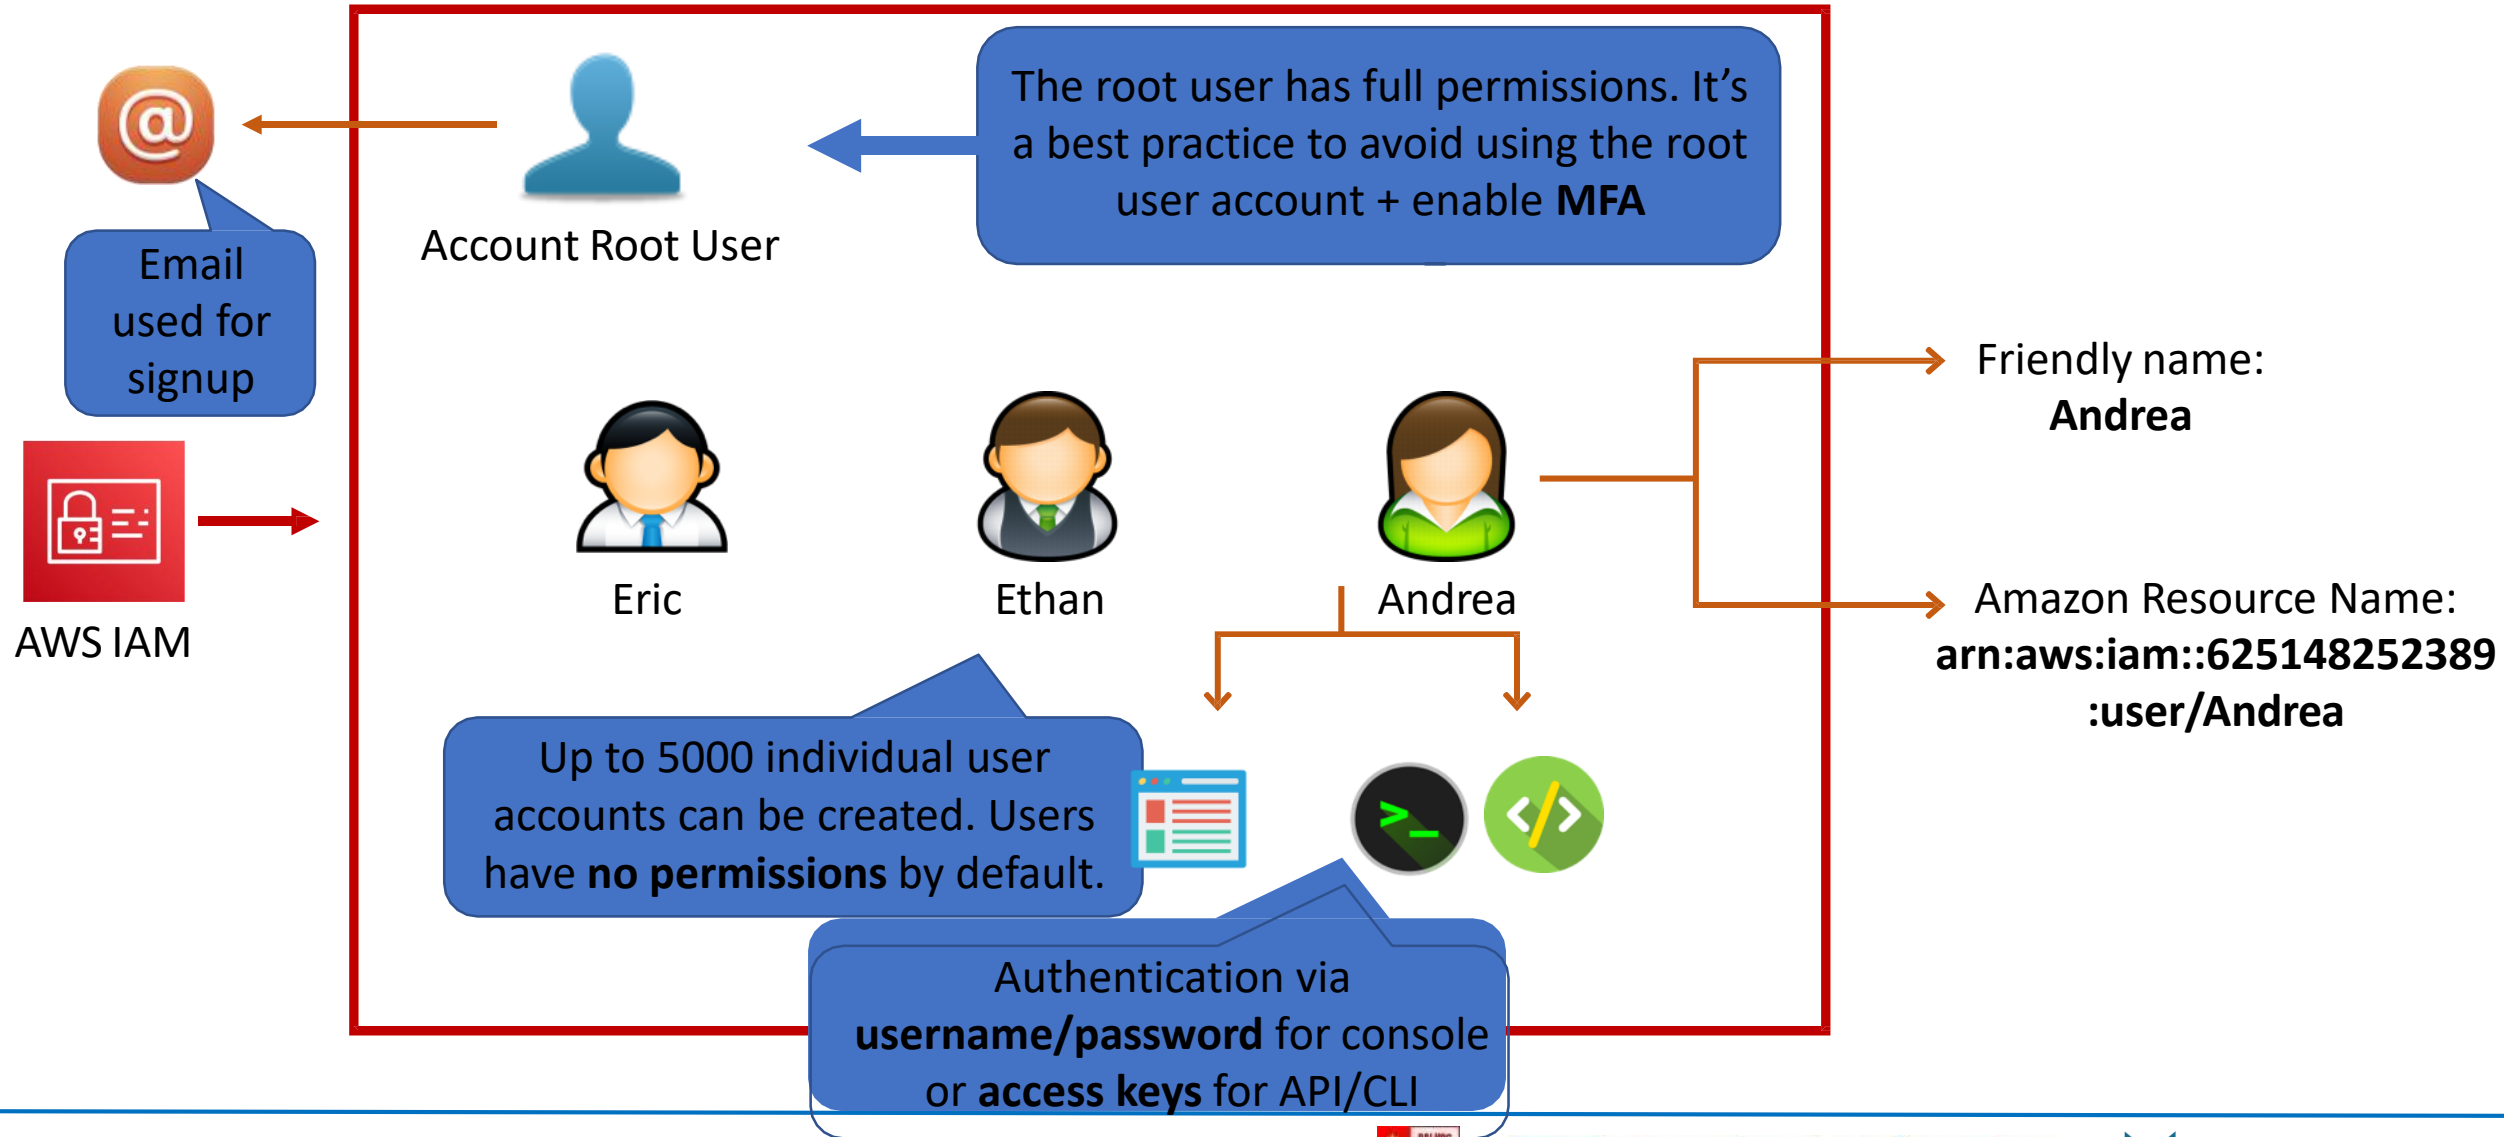
Overview of Users, Groups, Roles and Policies




Users, Groups, Roles and Policies




IAM Users



IAM Groups



Admin Group




Eric Sunil

Development Group



Ethan Lee


Operations Group




Andrea

The user gains the **permissions** applied to the **group** through the **policy**

Groups are collections of users. Users can be members of up to 10 groups

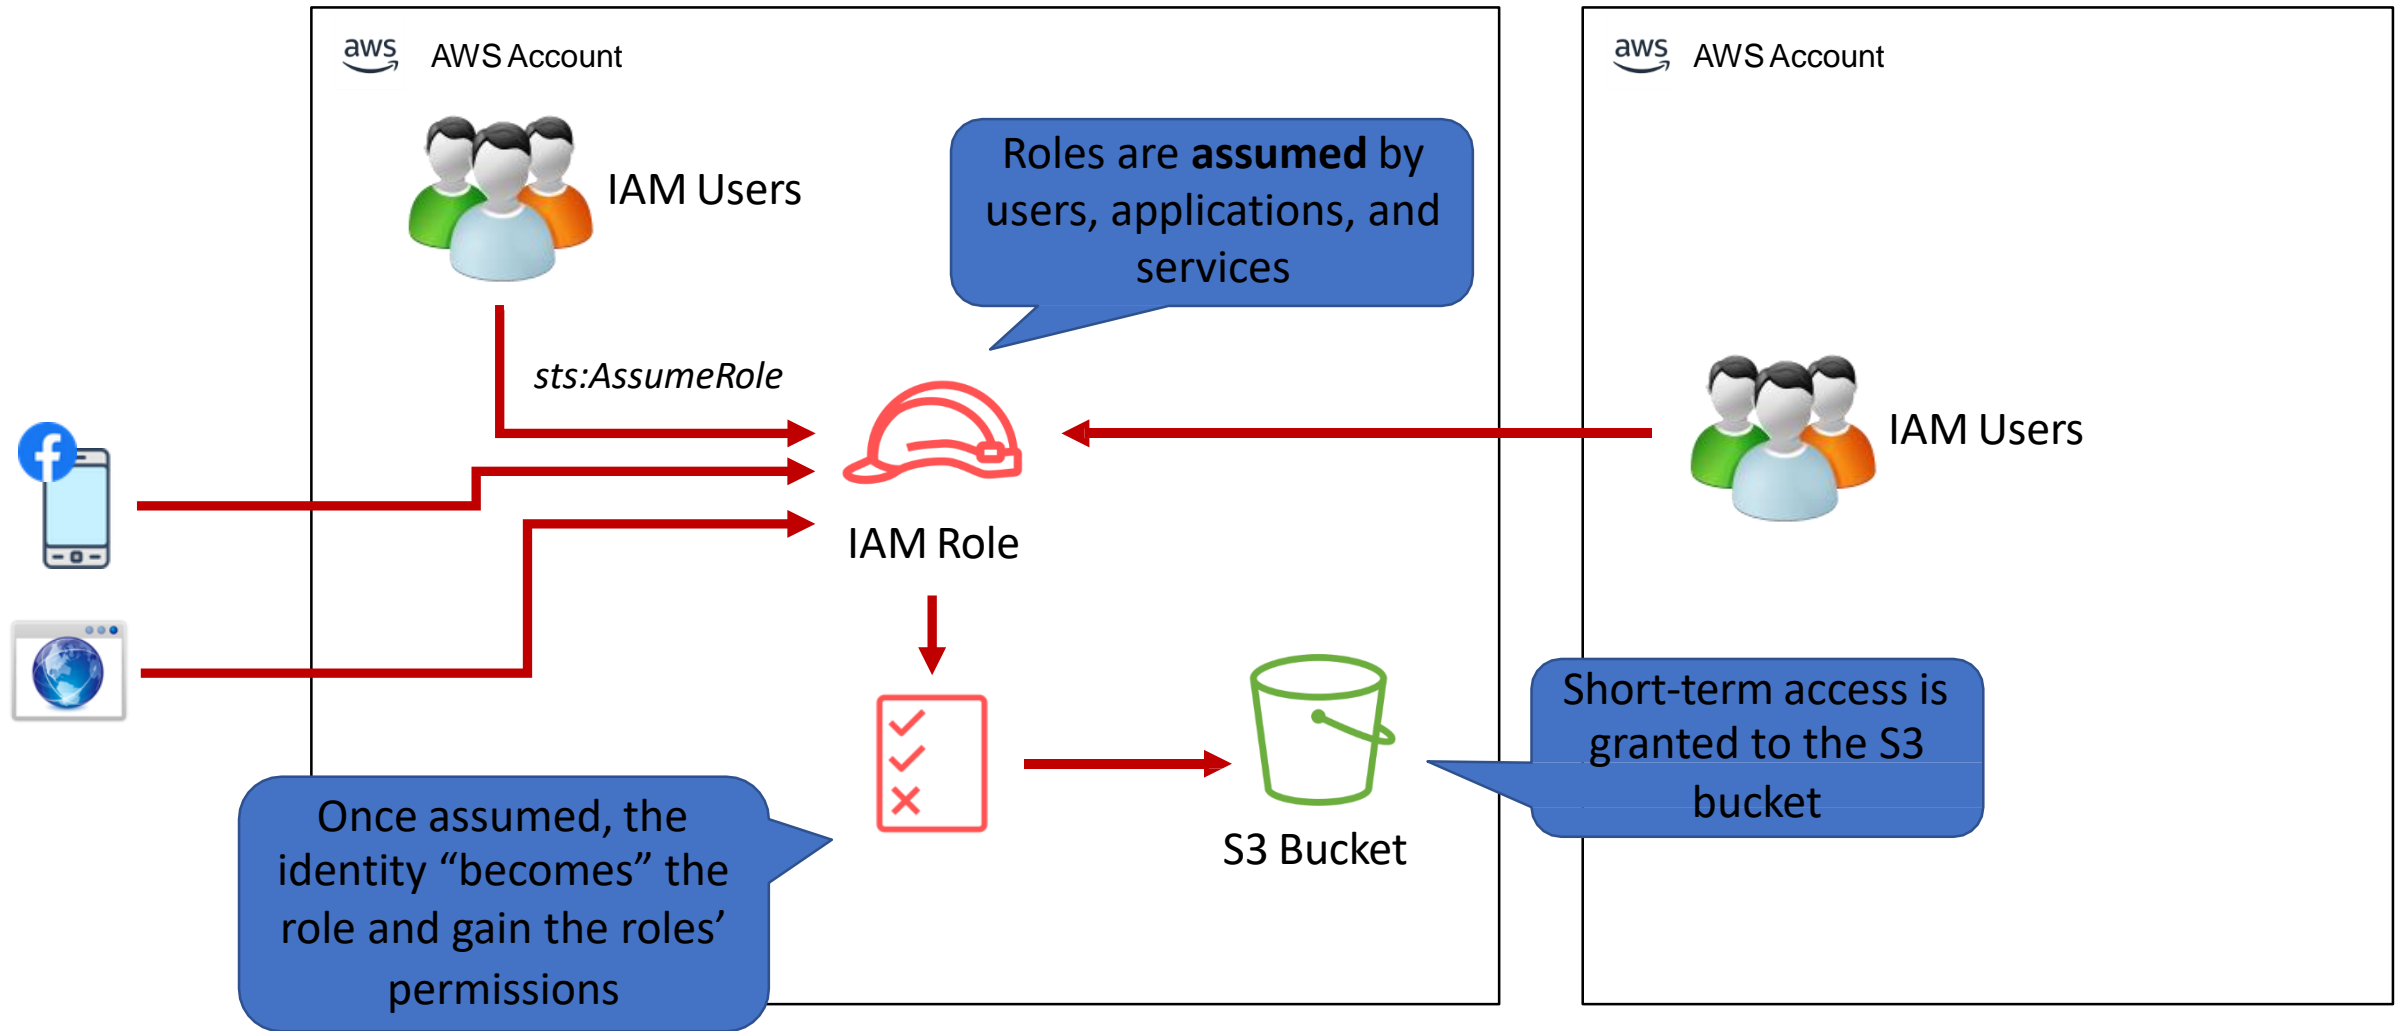


The main reason to use groups is to apply **permissions** to users using **policies**



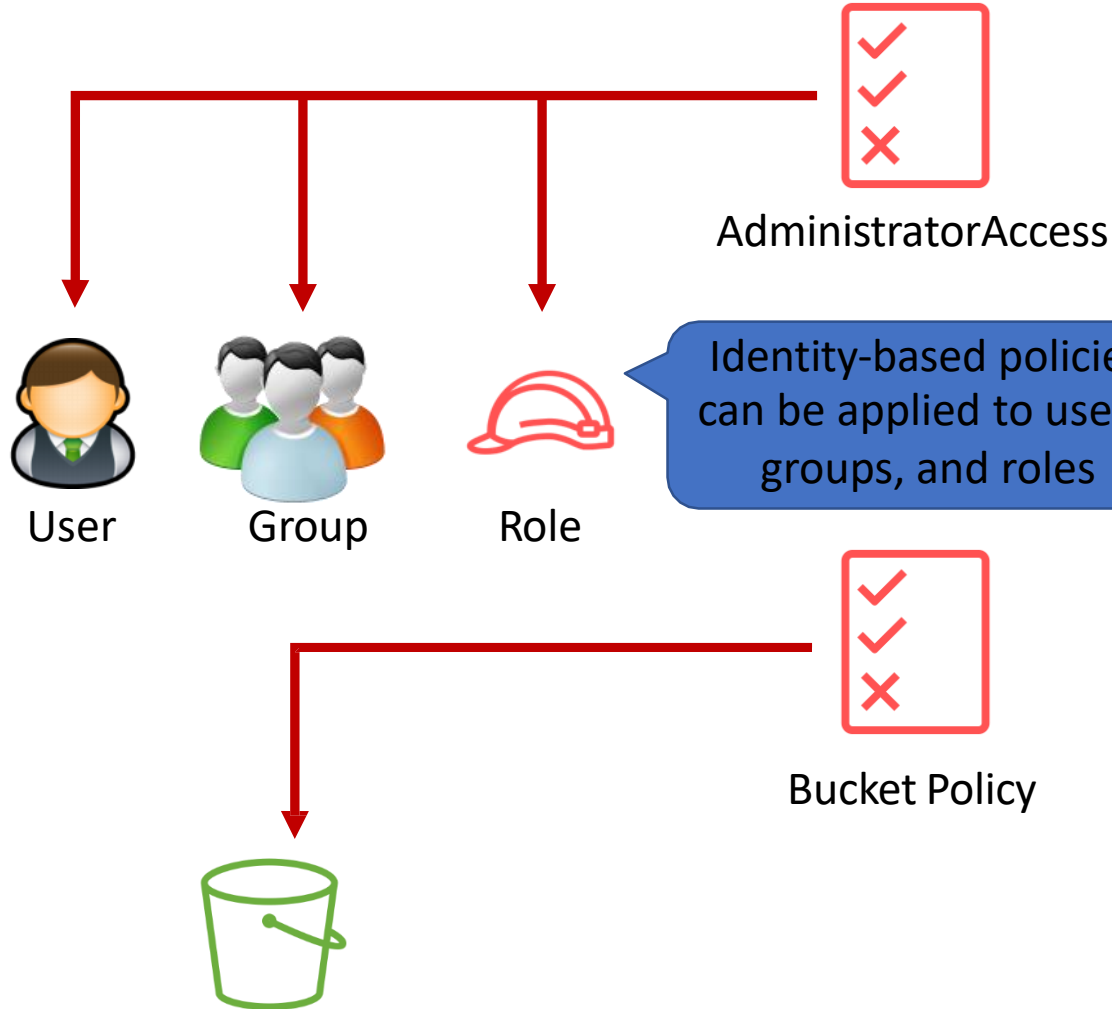
IAM Roles

An IAM role is an IAM identity that has specific permissions



IAM Policies

Policies are documents that define permissions and are written in JSON



Identity-based policies can be applied to users, groups, and roles

```
{  
  "Version": "2012-10-17",  
  "Statement": [  
    {  
      "Effect": "Allow",  
      "Action": "*",  
      "Resource": "*"  
    }  
  ]  
}
```

All permissions are implicitly denied by default

```
{  
  "Version": "2012-10-17",  
  "Id": "Policy1561964929358",  
  "Statement": [  
    {  
      "Sid": "Stmt1561964454052",  
      "Effect": "Allow",  
      "Principal": {  
        "AWS": "arn:aws:iam::515148227241:user/Paul"  
      },  
      "Action": "s3:*",  
      "Resource": "arn:aws:s3:::dctcompany",  
      "Condition": {  
        "StringLike": {  
          "s3:prefix": "Confidential/*"  
        }  
      }  
    }  
  ]  
}
```

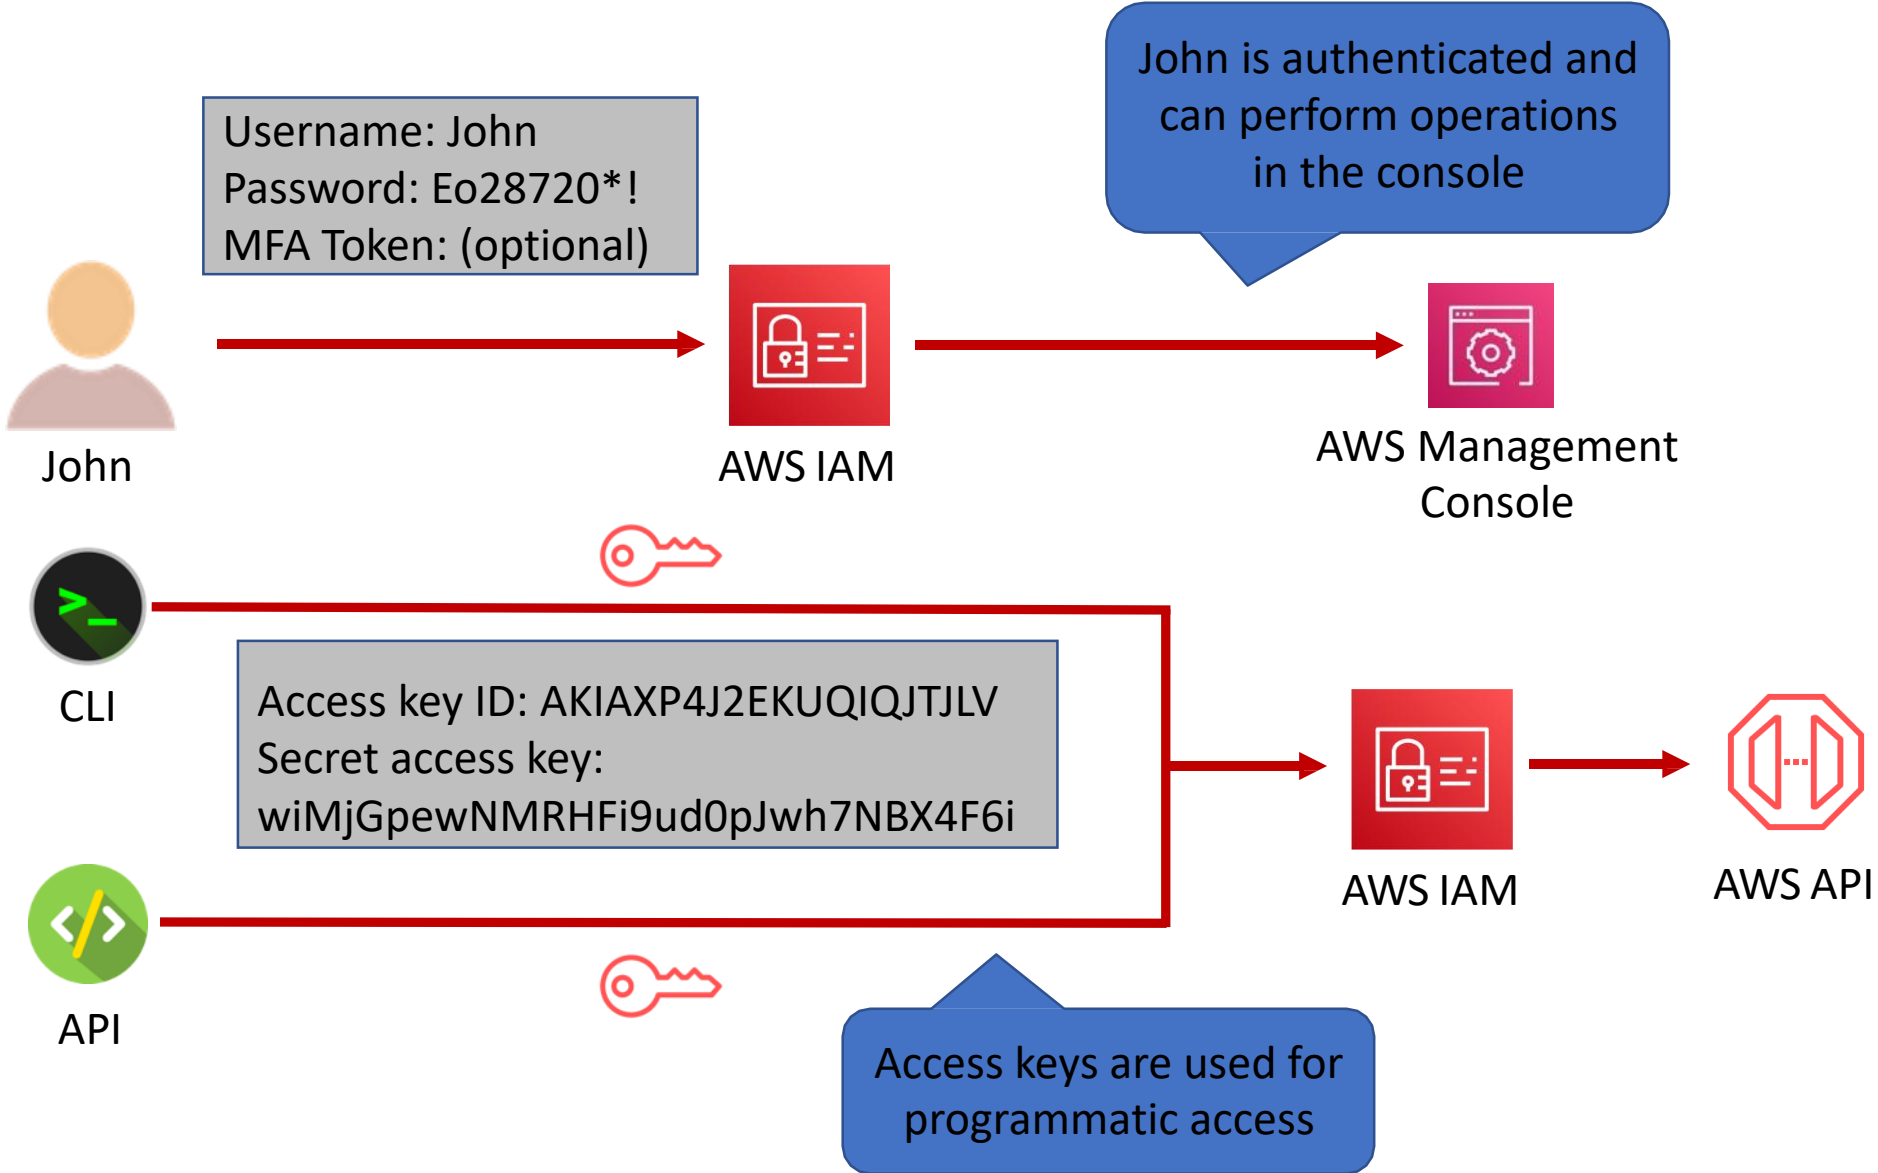
Resource-based policies apply to resources such as S3 buckets or DynamoDB tables

S3 Bucket

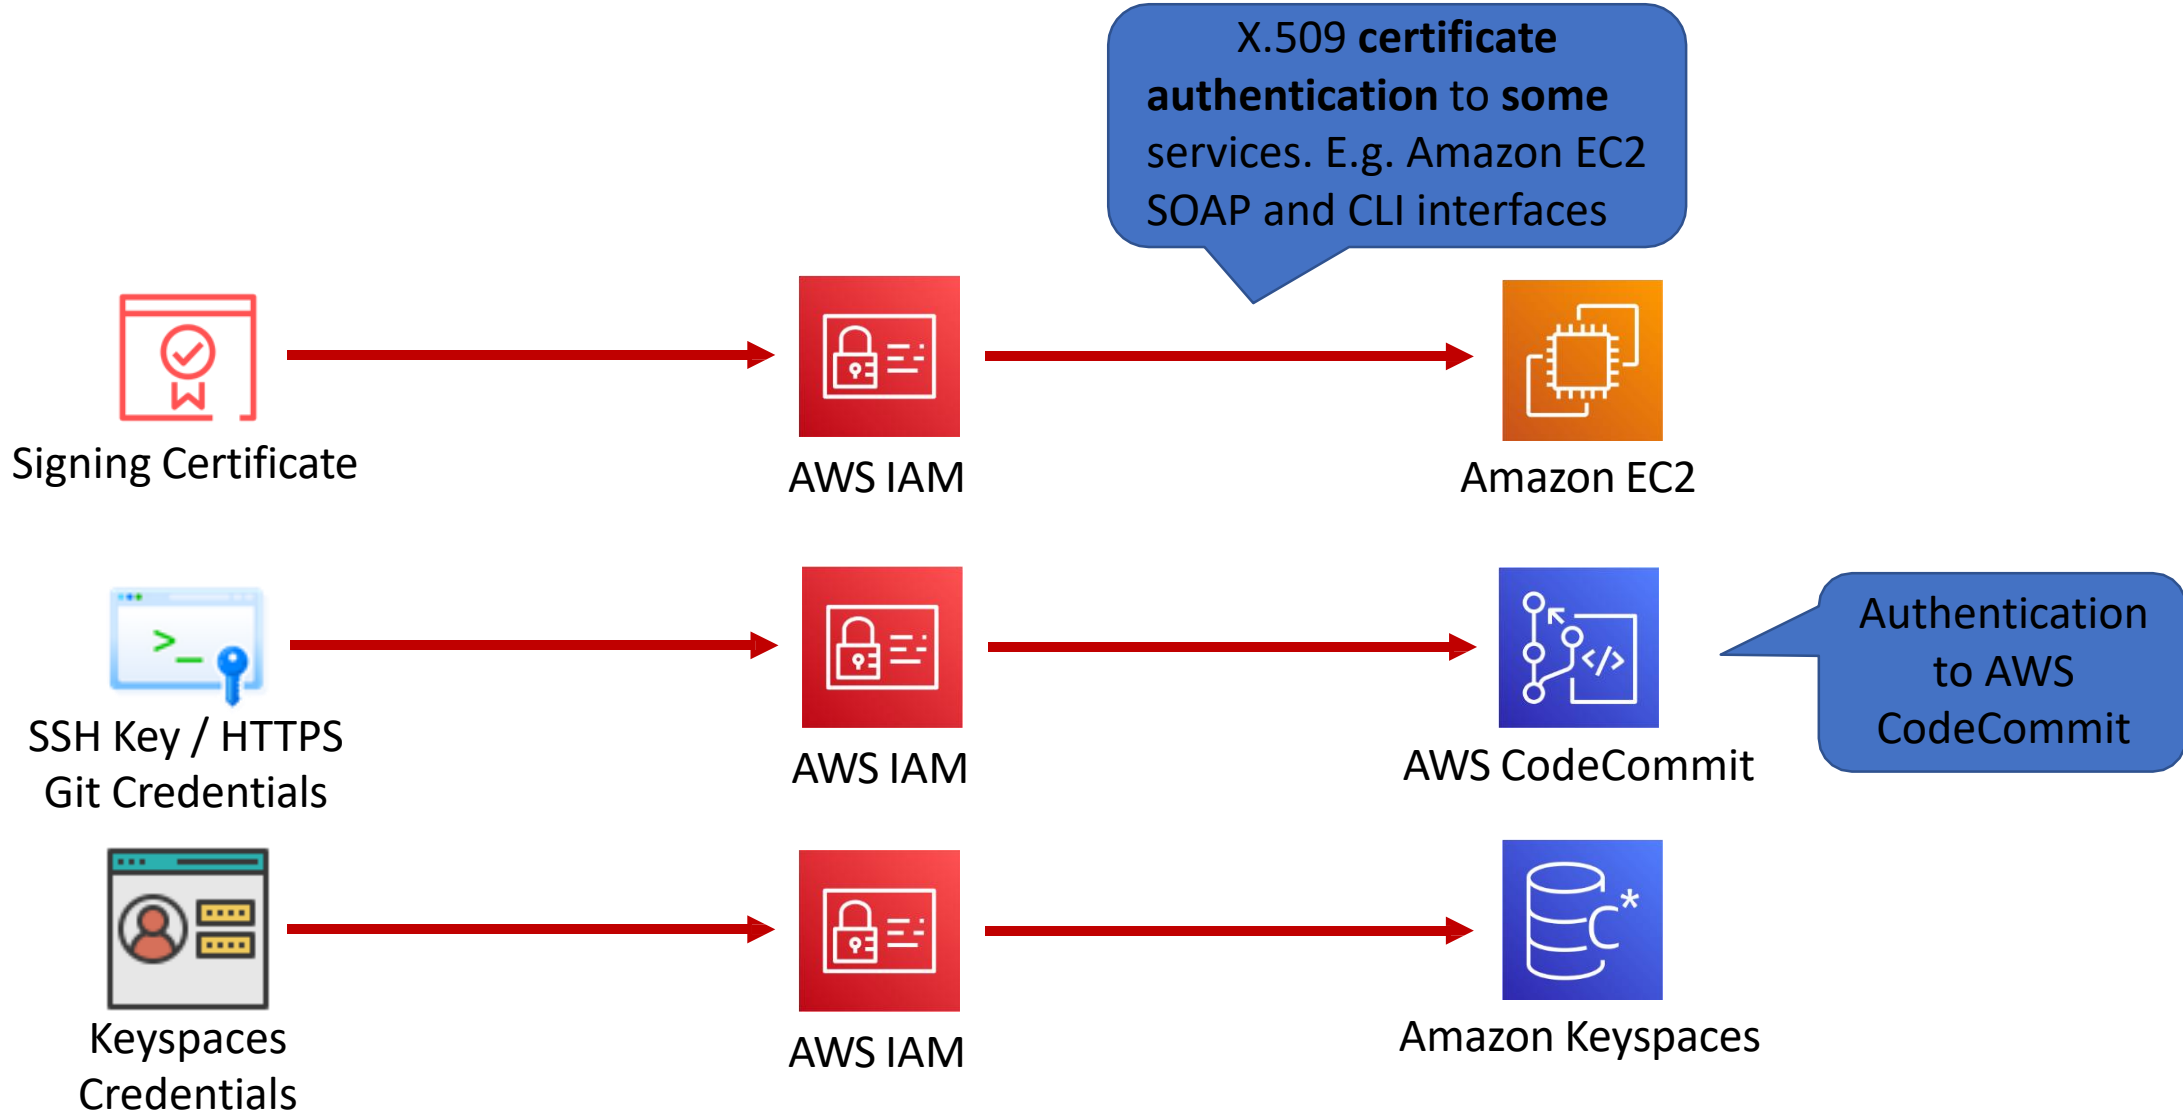
IAM Authentication Methods



IAM Authentication Methods



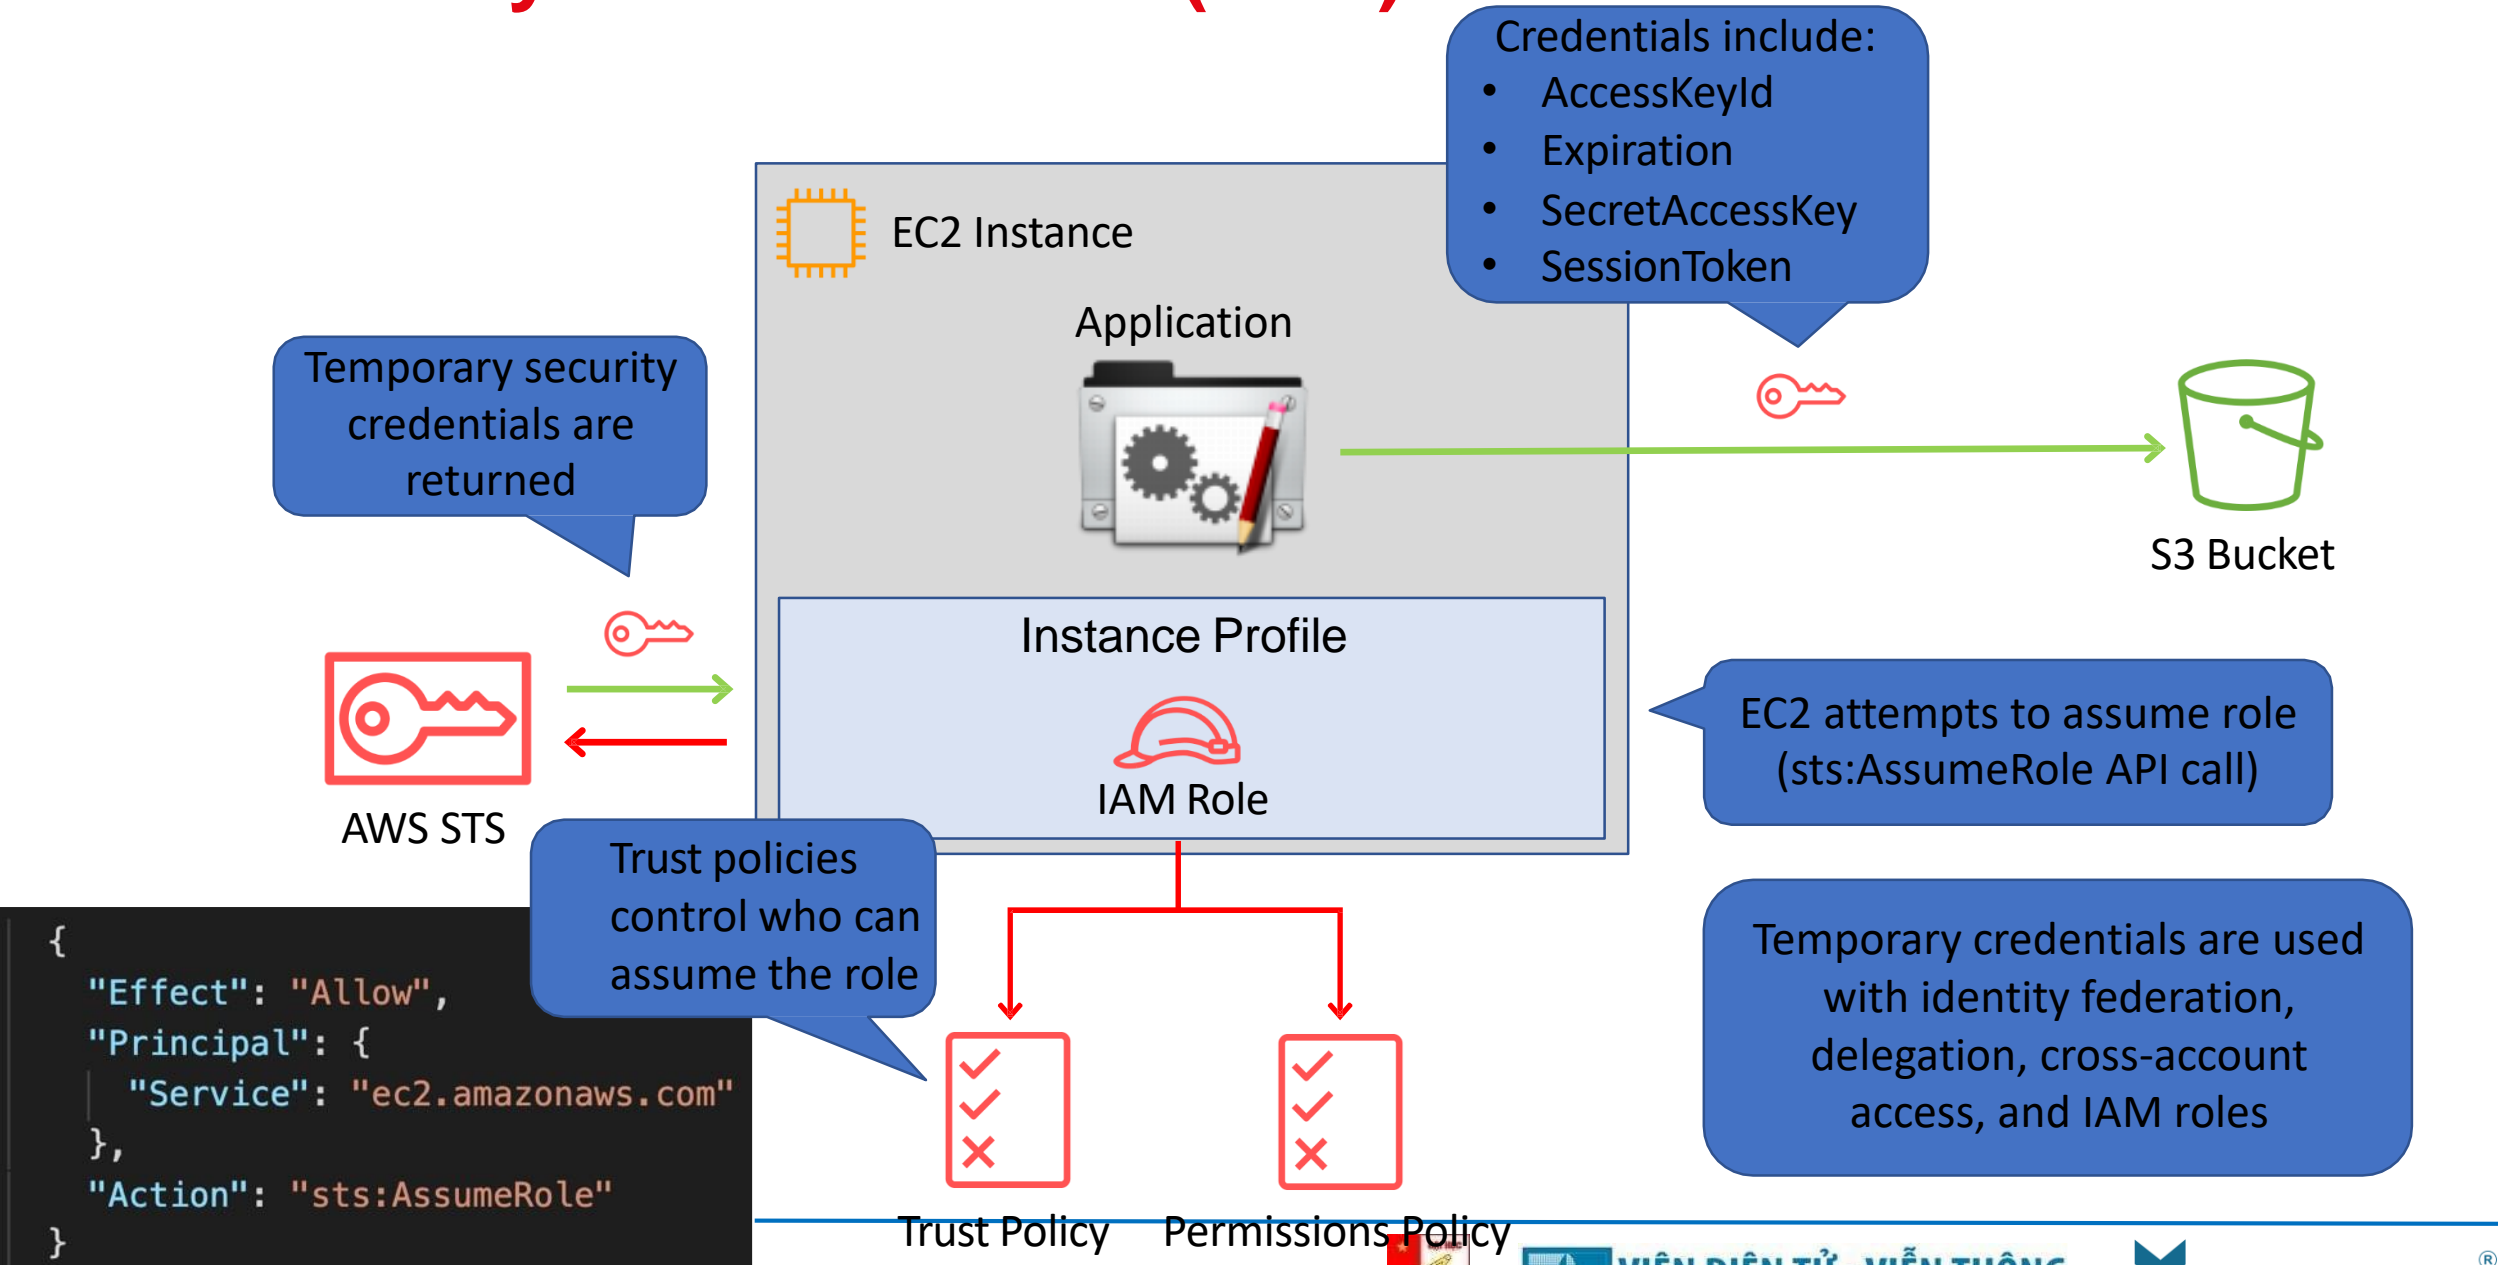
IAM Authentication Methods



AWS Security Token Service (STS)



AWS Security Token Service (STS)



```
{  
  "Effect": "Allow",  
  "Principal": {  
    "Service": "ec2.amazonaws.com"  
  },  
  "Action": "sts:AssumeRole"  
}
```



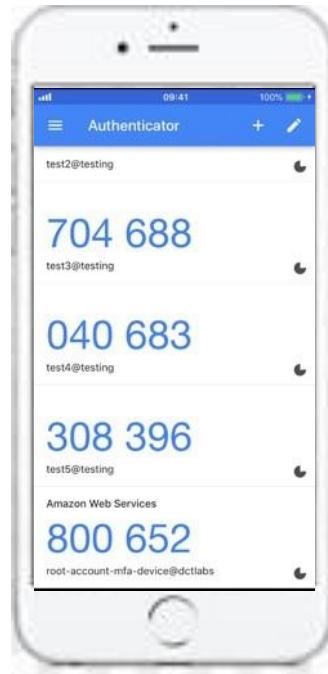
Multi-Factor Authentication

Something you know:

EJpX!*21p9%

Password

Something you have:



Something you are:



Multi-Factor Authentication

Something you **know**:



IAM User

EJPx!*21p9%

Password

Something you **have**:




Physical MFA



e.g. Google Authenticator on
your smart phone

Virtual MFA

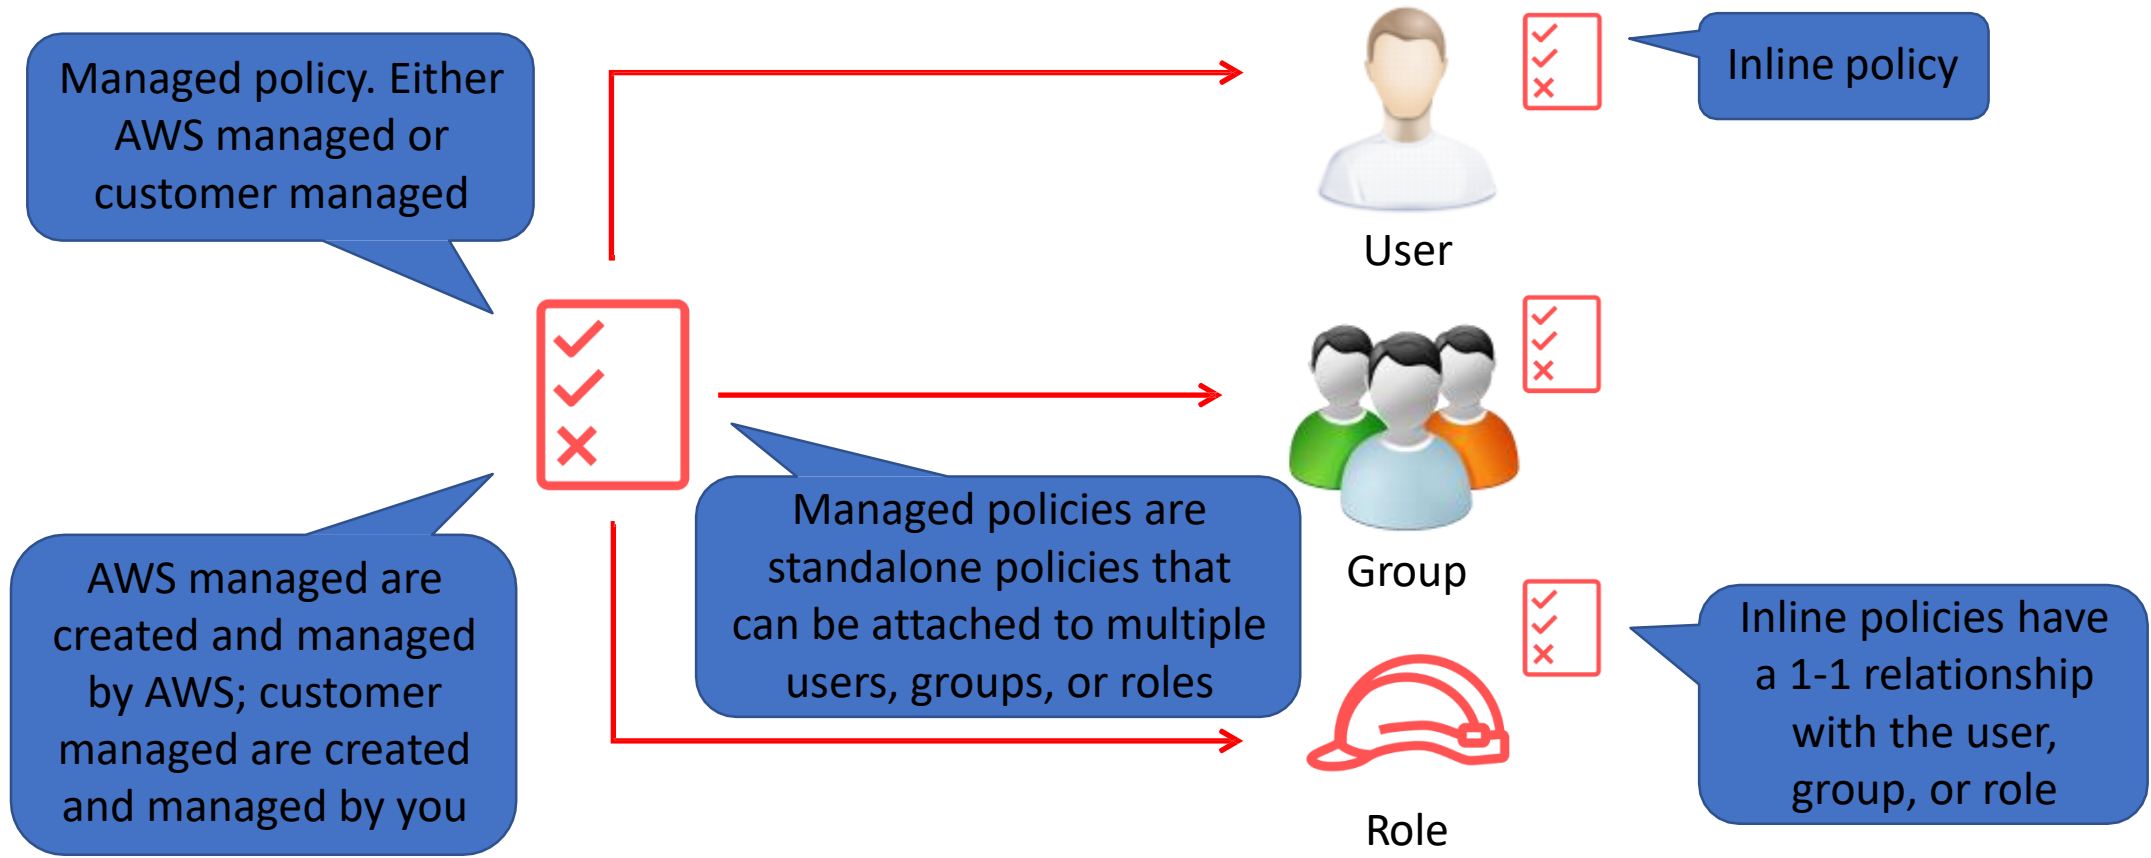


Identity-Based Policies and Resource-Based Policies



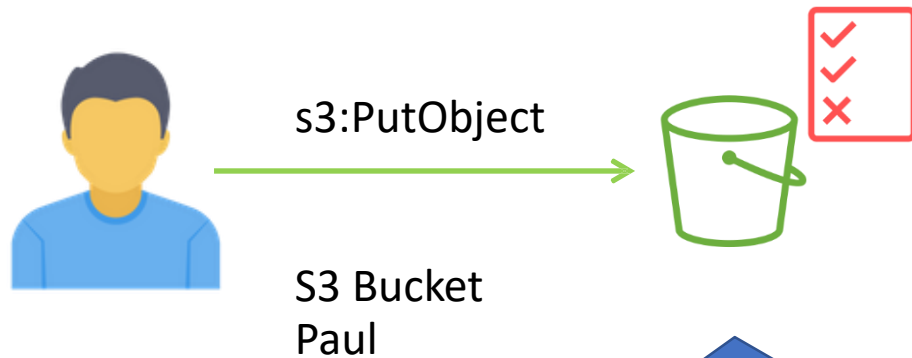
Identity-Based IAM Policies

Identity-based policies are JSON permissions policy documents that control what actions an identity can perform, on which resources, and under what conditions



Resource-Based Policies

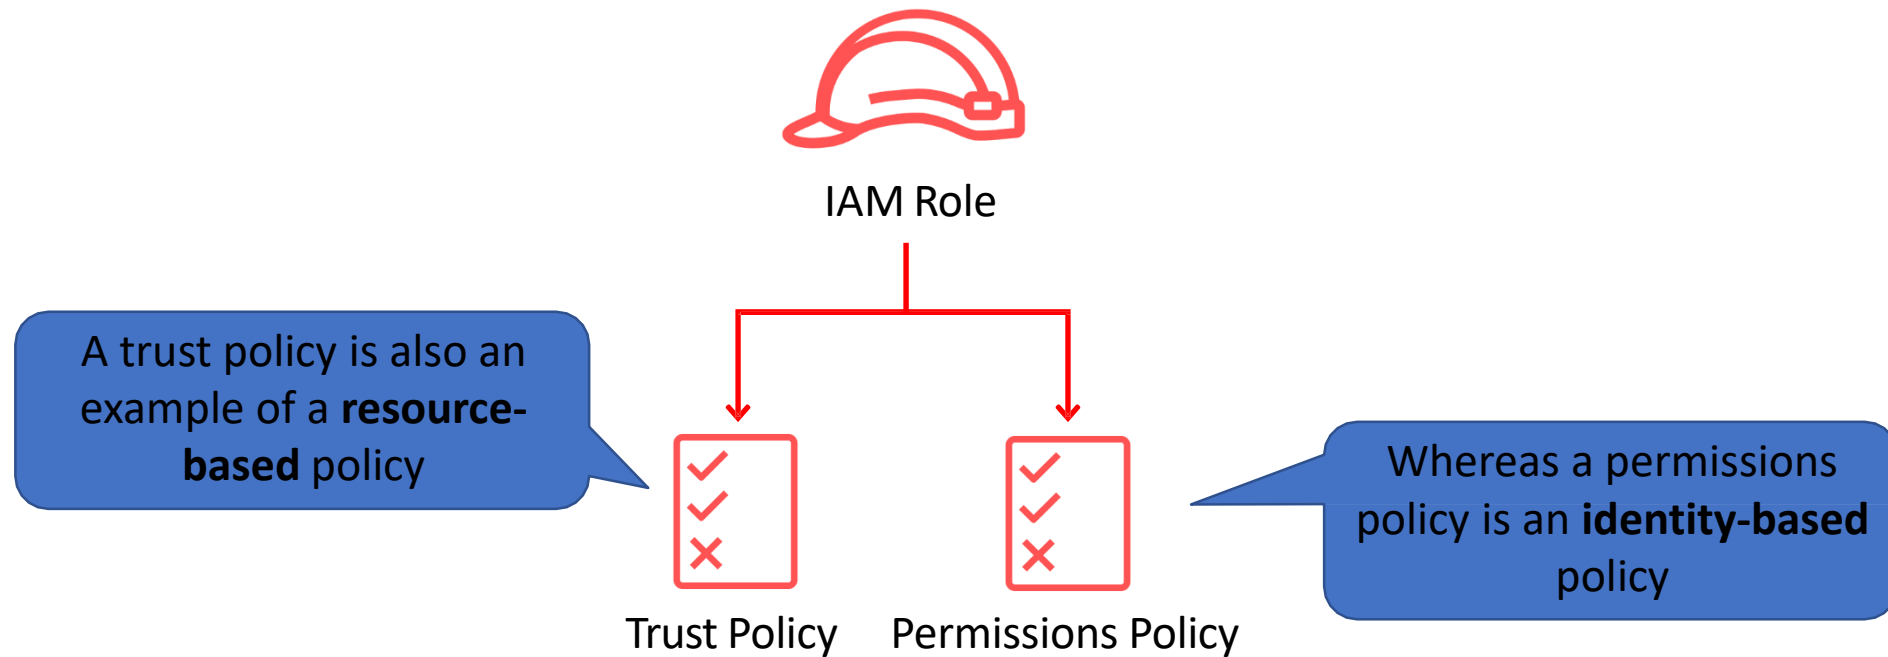
Resource-based policies are JSON policy documents that you attach to a resource such as an Amazon S3 bucket



Resource-based policies grant the specified **principal** (Paul) **permission** to perform specific **actions** on the **resource**

```
{
  "Version": "2012-10-17",
  "Id": "Policy1561964929358",
  "Statement": [
    {
      "Sid": "Stmt1561964454052",
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::515148227241:user/Paul"
      },
      "Action": "s3:*",
      "Resource": "arn:aws:s3:::dctcompany"
    }
  ]
}
```

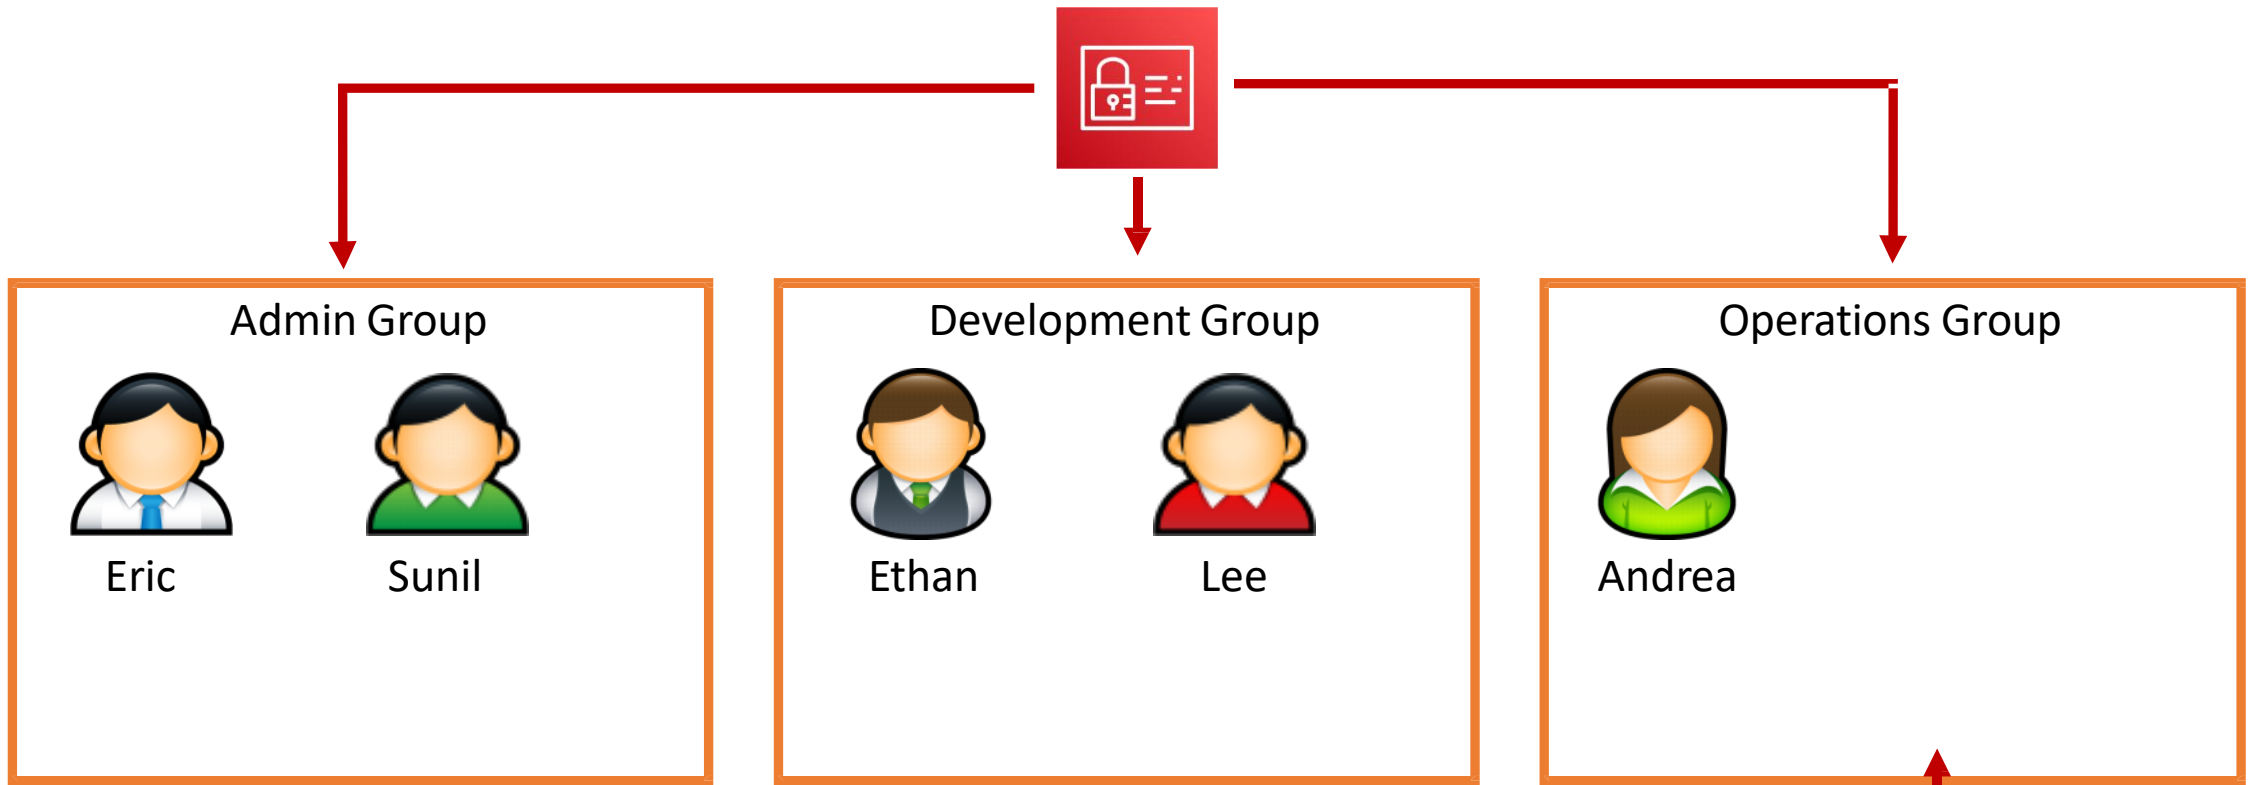
Resource-Based Policies



Access Control Methods - RBAC & ABAC




Role-Based Access Control (RBAC)



Users are assigned permissions through policies attached to groups



Groups are organized by job function



Best practice is to grant the minimum permissions required to perform the job

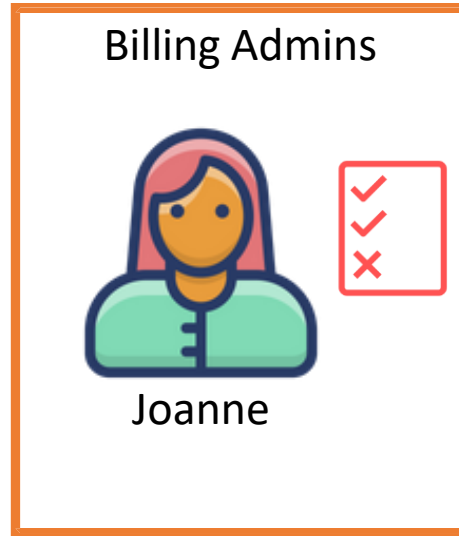


Role-Based Access Control (RBAC)

Job function policies:

- Administrator
- Billing
- Database administrator
- Data scientist
- Developer power user
- Network administrator
- Security auditor
- Support user
- System administrator
- View-only user

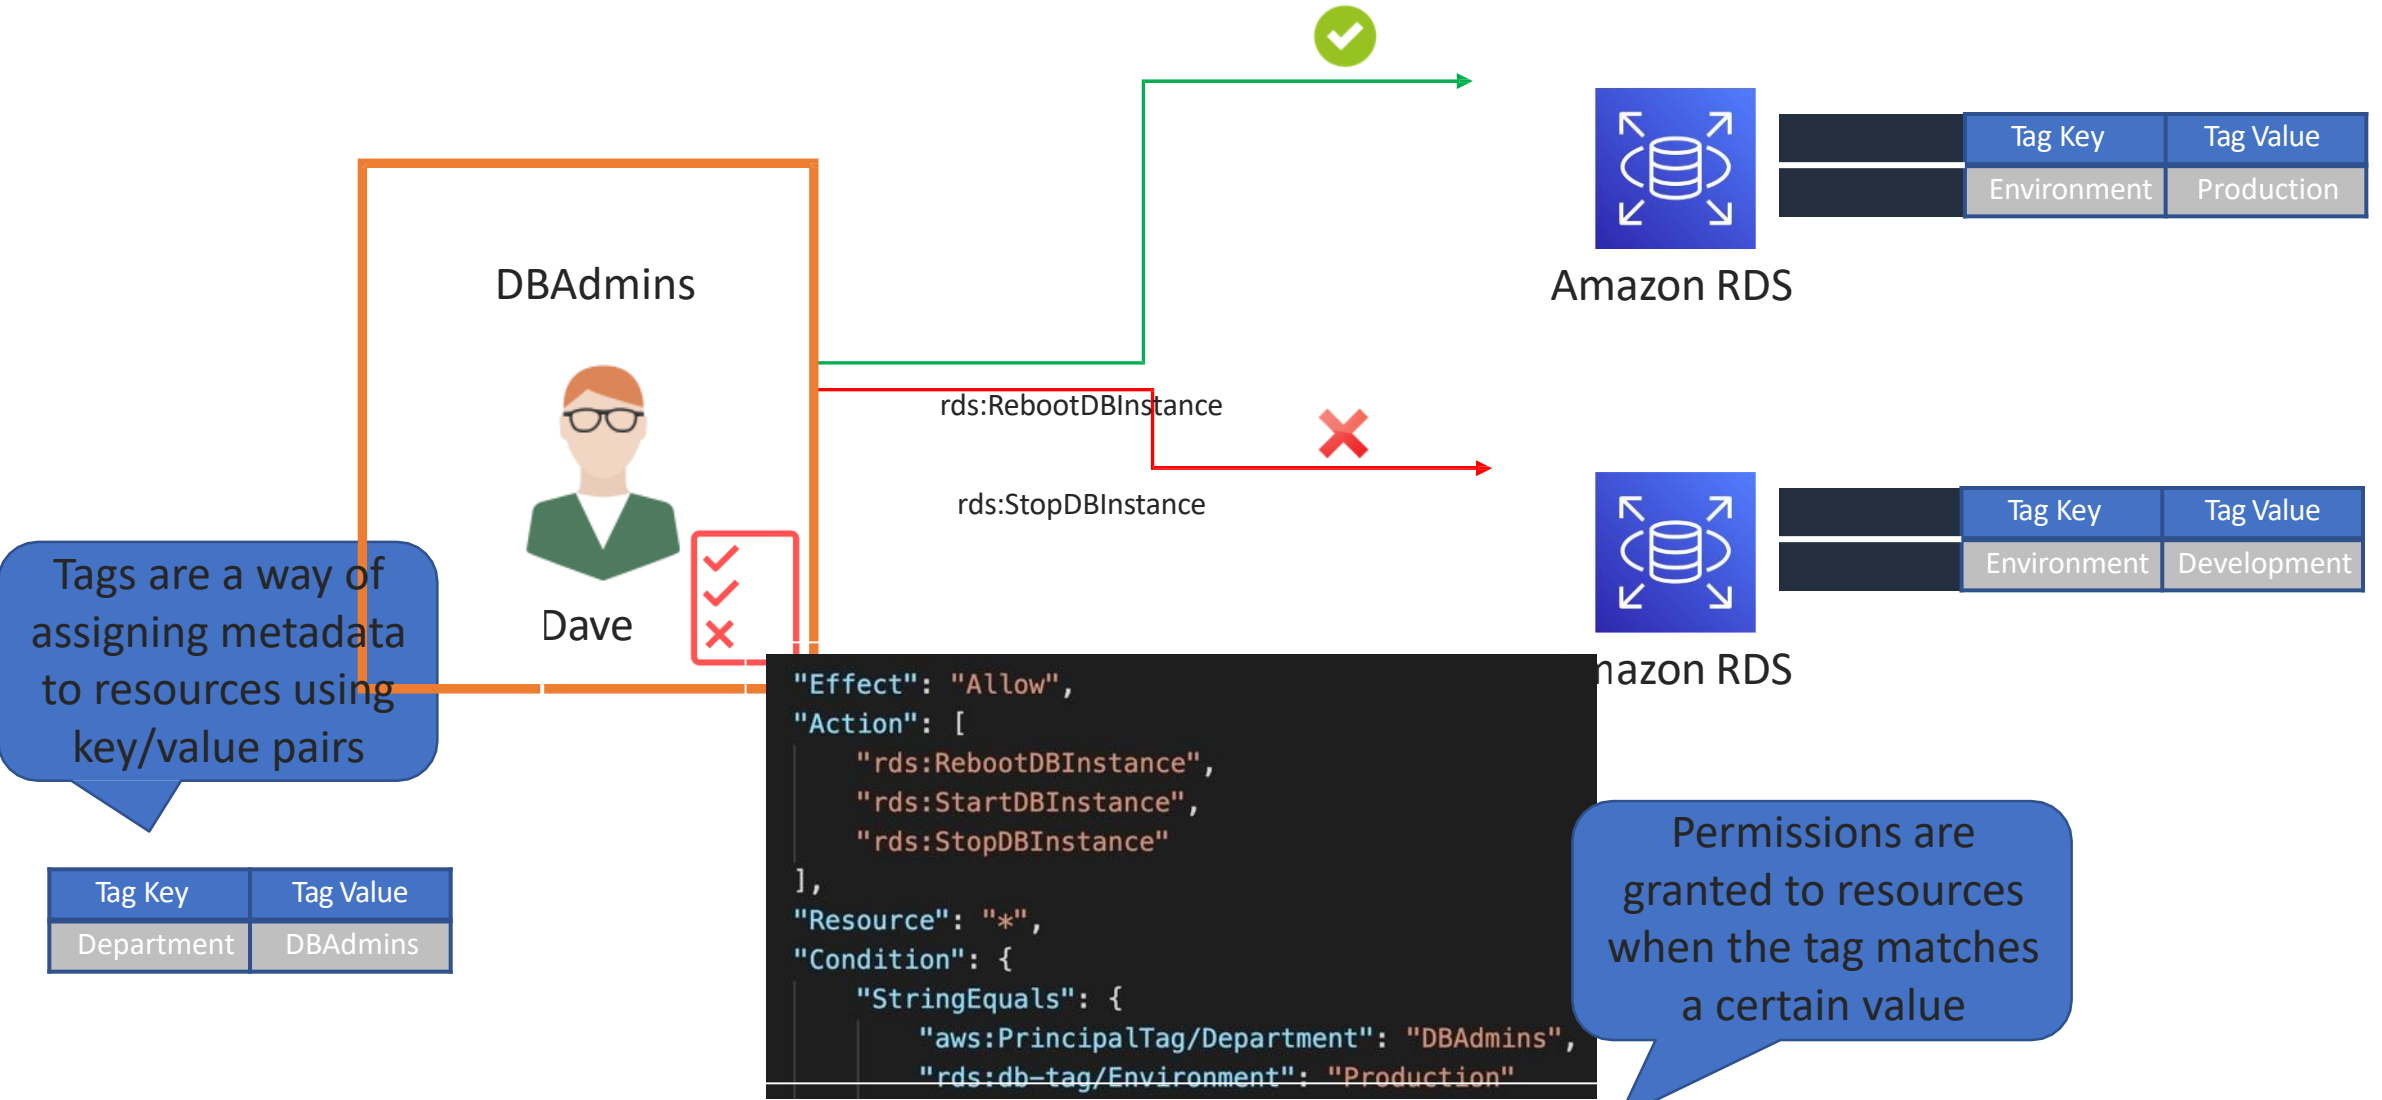
The Billing managed policy is attached to the group



AWS managed policies for job functions are designed to closely align to common job functions in the IT industry

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "aws-portal:*Billing",
        "aws-portal:*Usage",
        "aws-portal:*PaymentMethods",
        "budgets:ViewBudget",
        "budgets:ModifyBudget",
        "ce:UpdatePreferences",
        "ce:CreateReport",
        "ce:UpdateReport",
        "ce>DeleteReport",
        "ce:CreateNotificationSubscription",
        "ce:UpdateNotificationSubscription",
        "ce>DeleteNotificationSubscription",
        "cur:DescribeReportDefinitions",
        "cur:PutReportDefinition",
        "cur:ModifyReportDefinition",
        "cur>DeleteReportDefinition",
        "purchase-orders:*PurchaseOrders"
      ],
      "Resource": "*"
    }
  ]
}
```

Attribute-Based Access Control (ABAC)



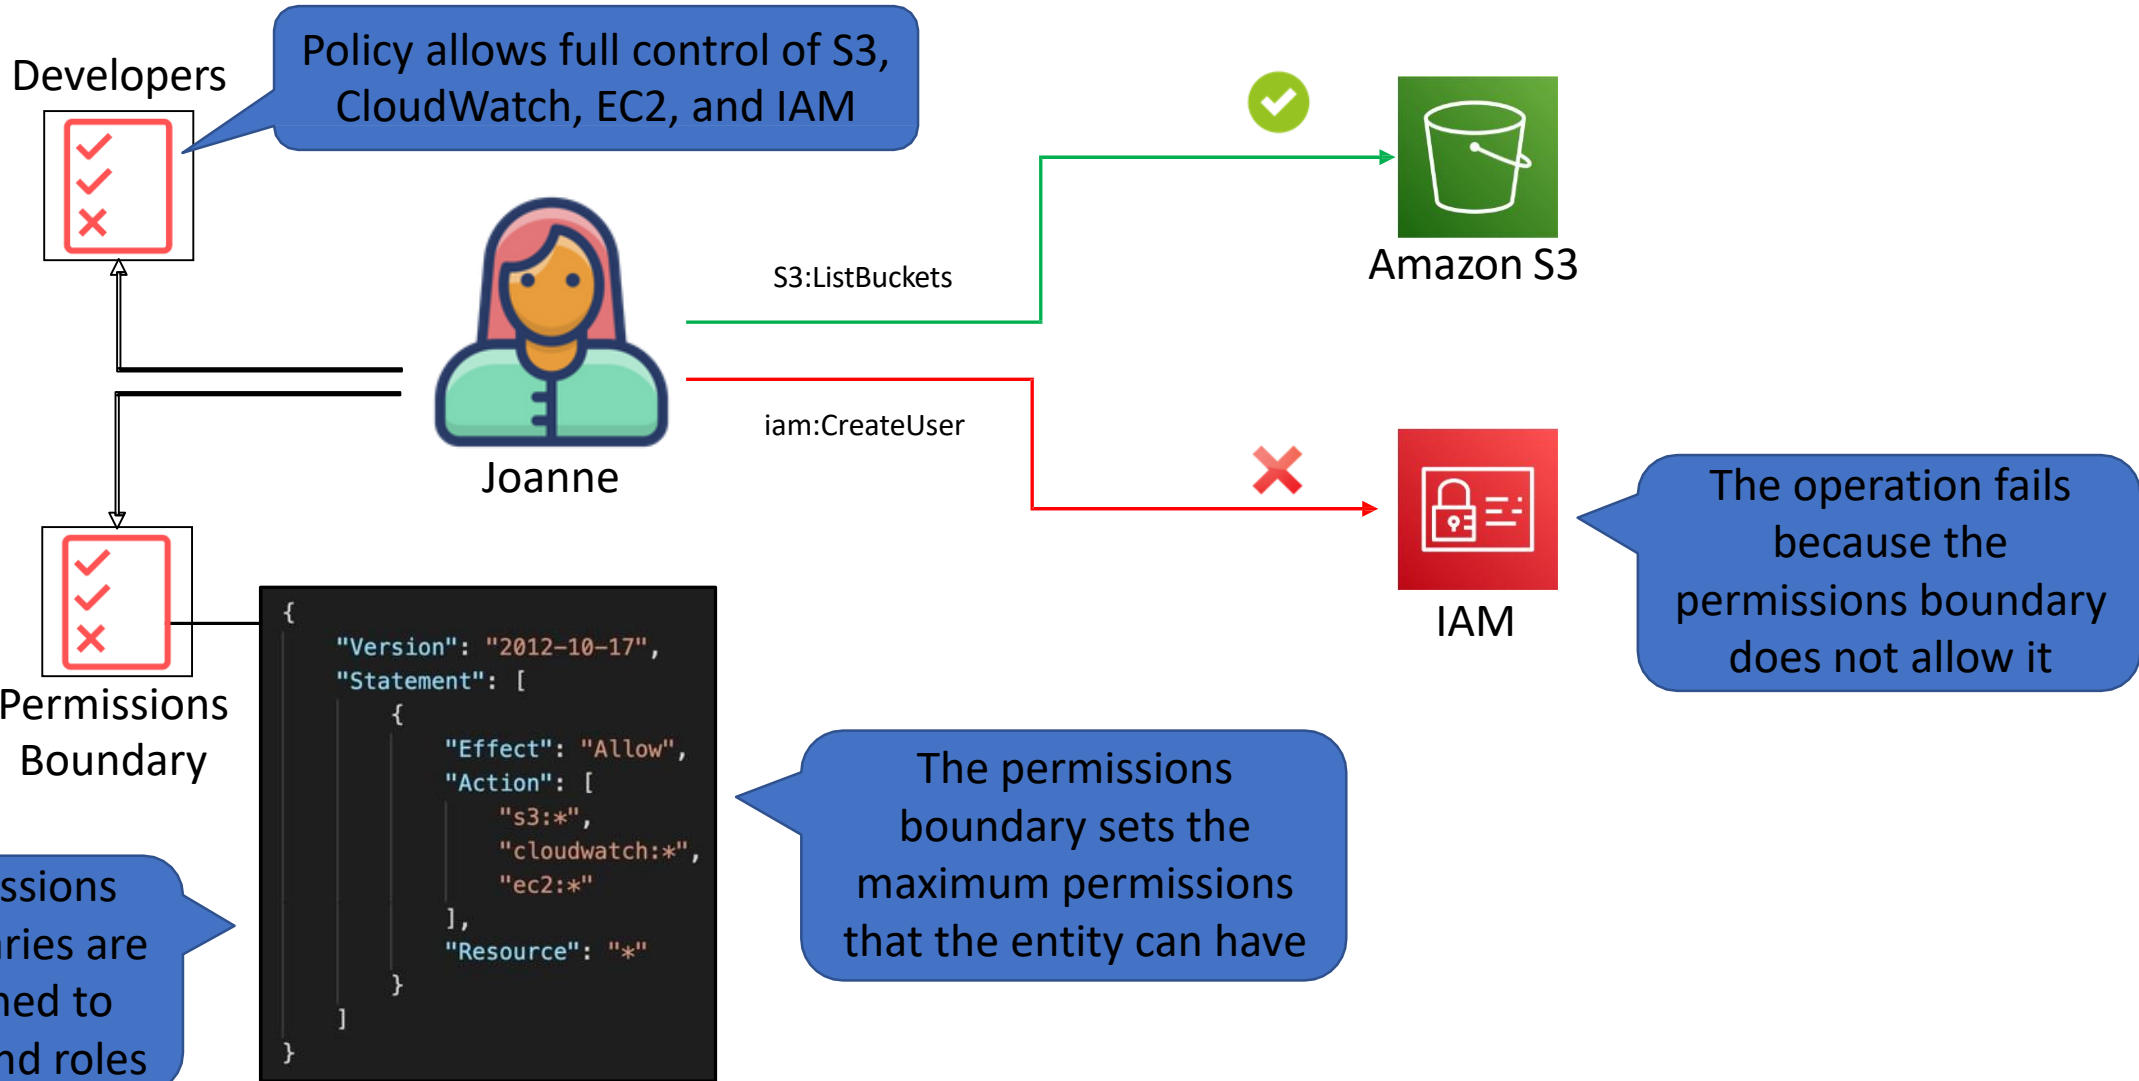
Permissions Boundaries



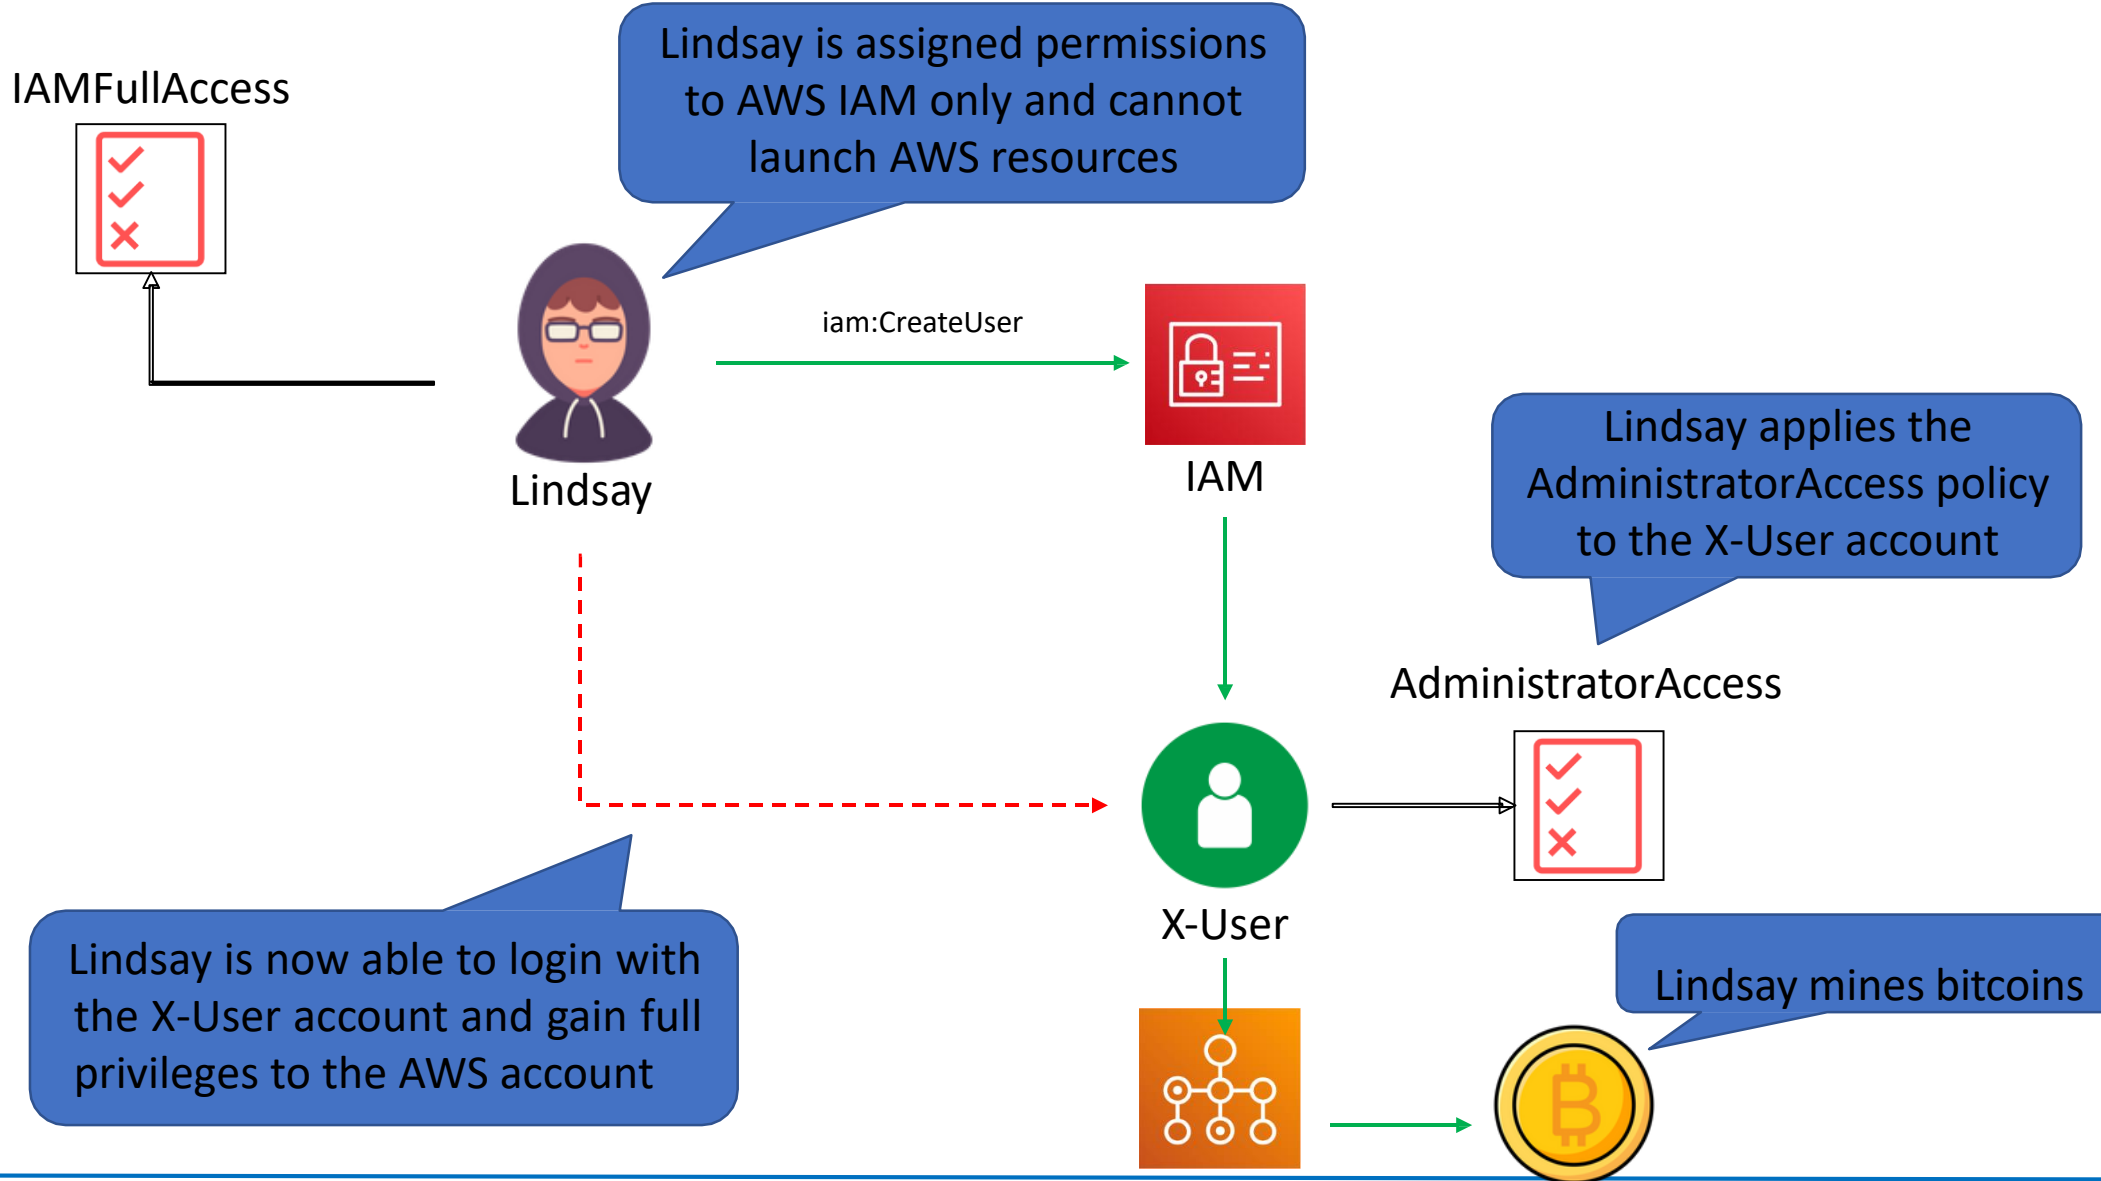
VIỆN ĐIỆN TỬ - VIỄN THÔNG
School of Electronics and Telecommunications



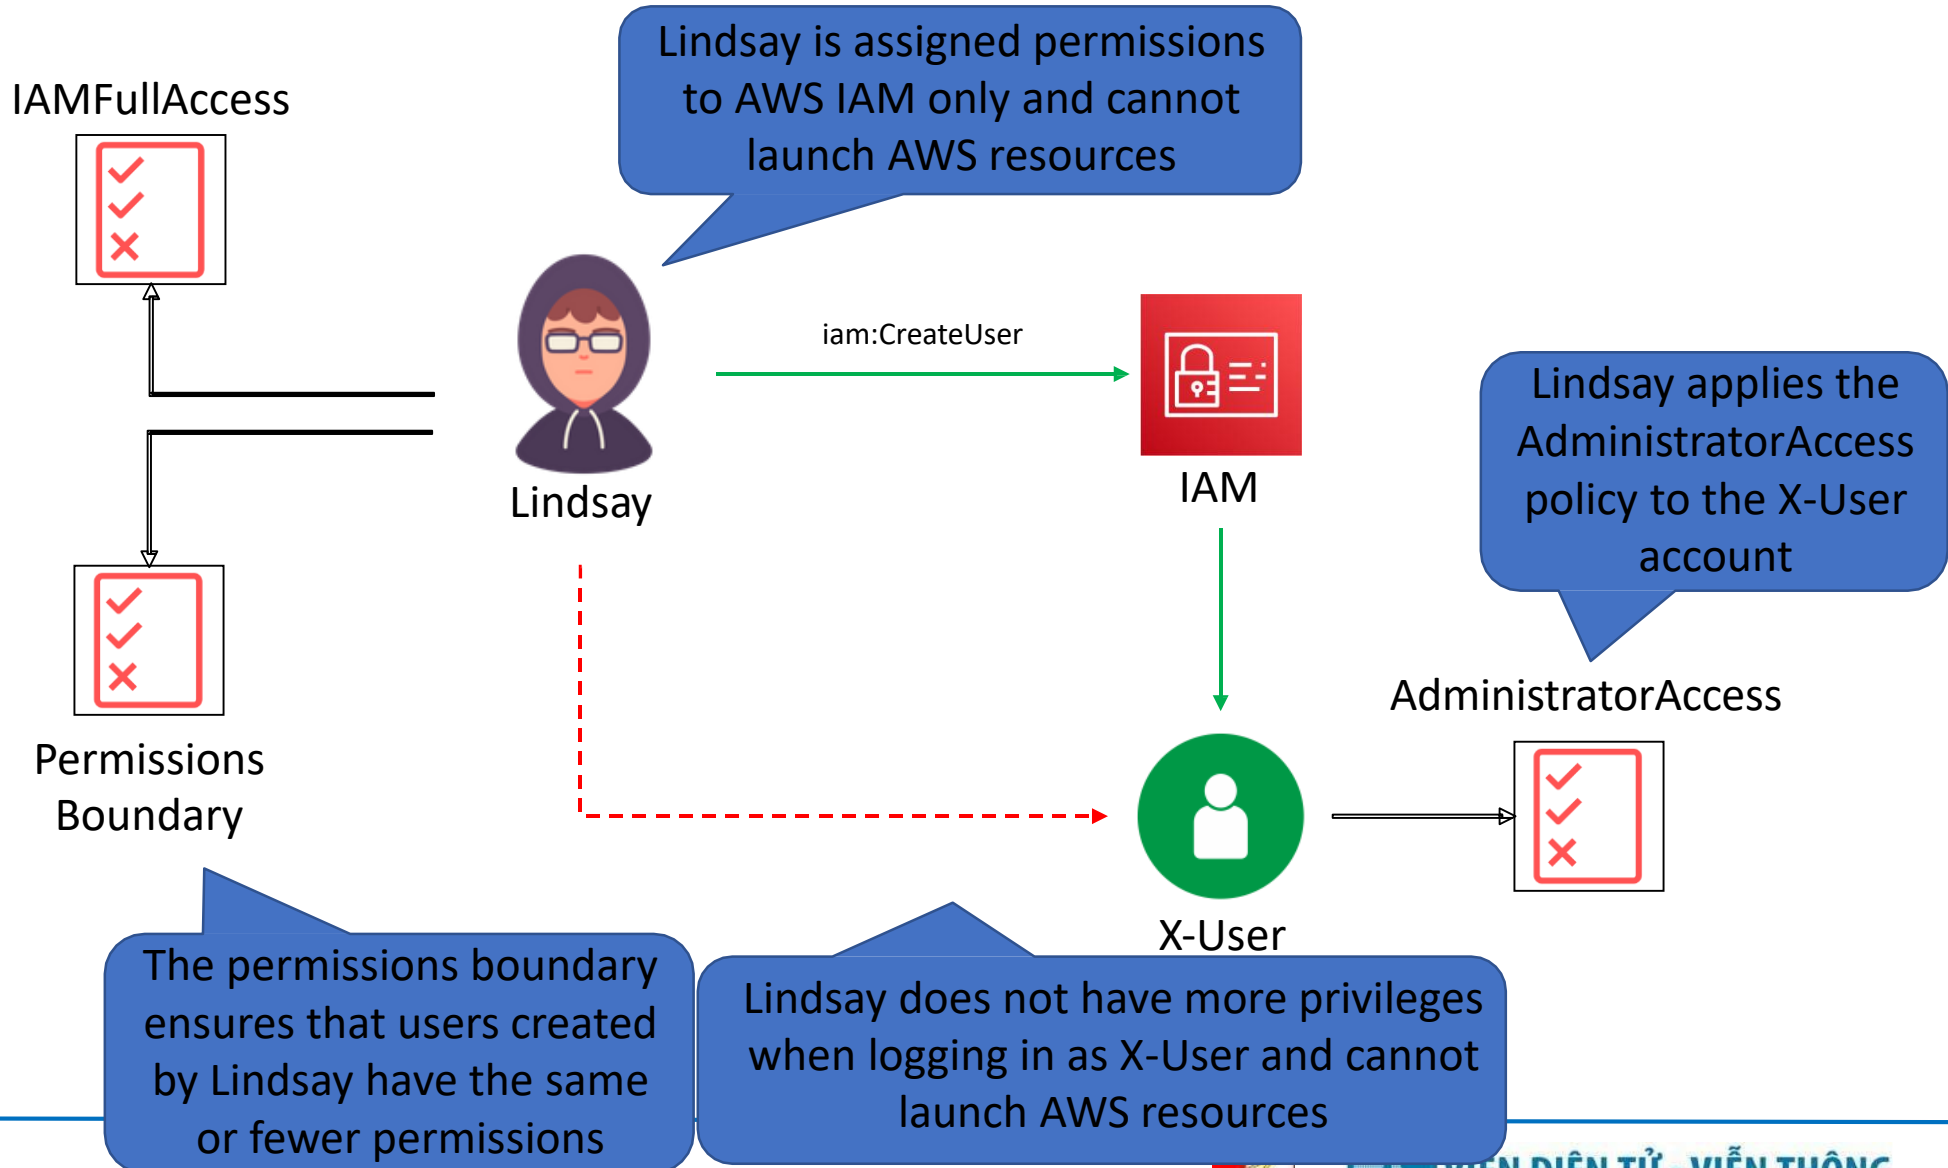
Permissions Boundaries



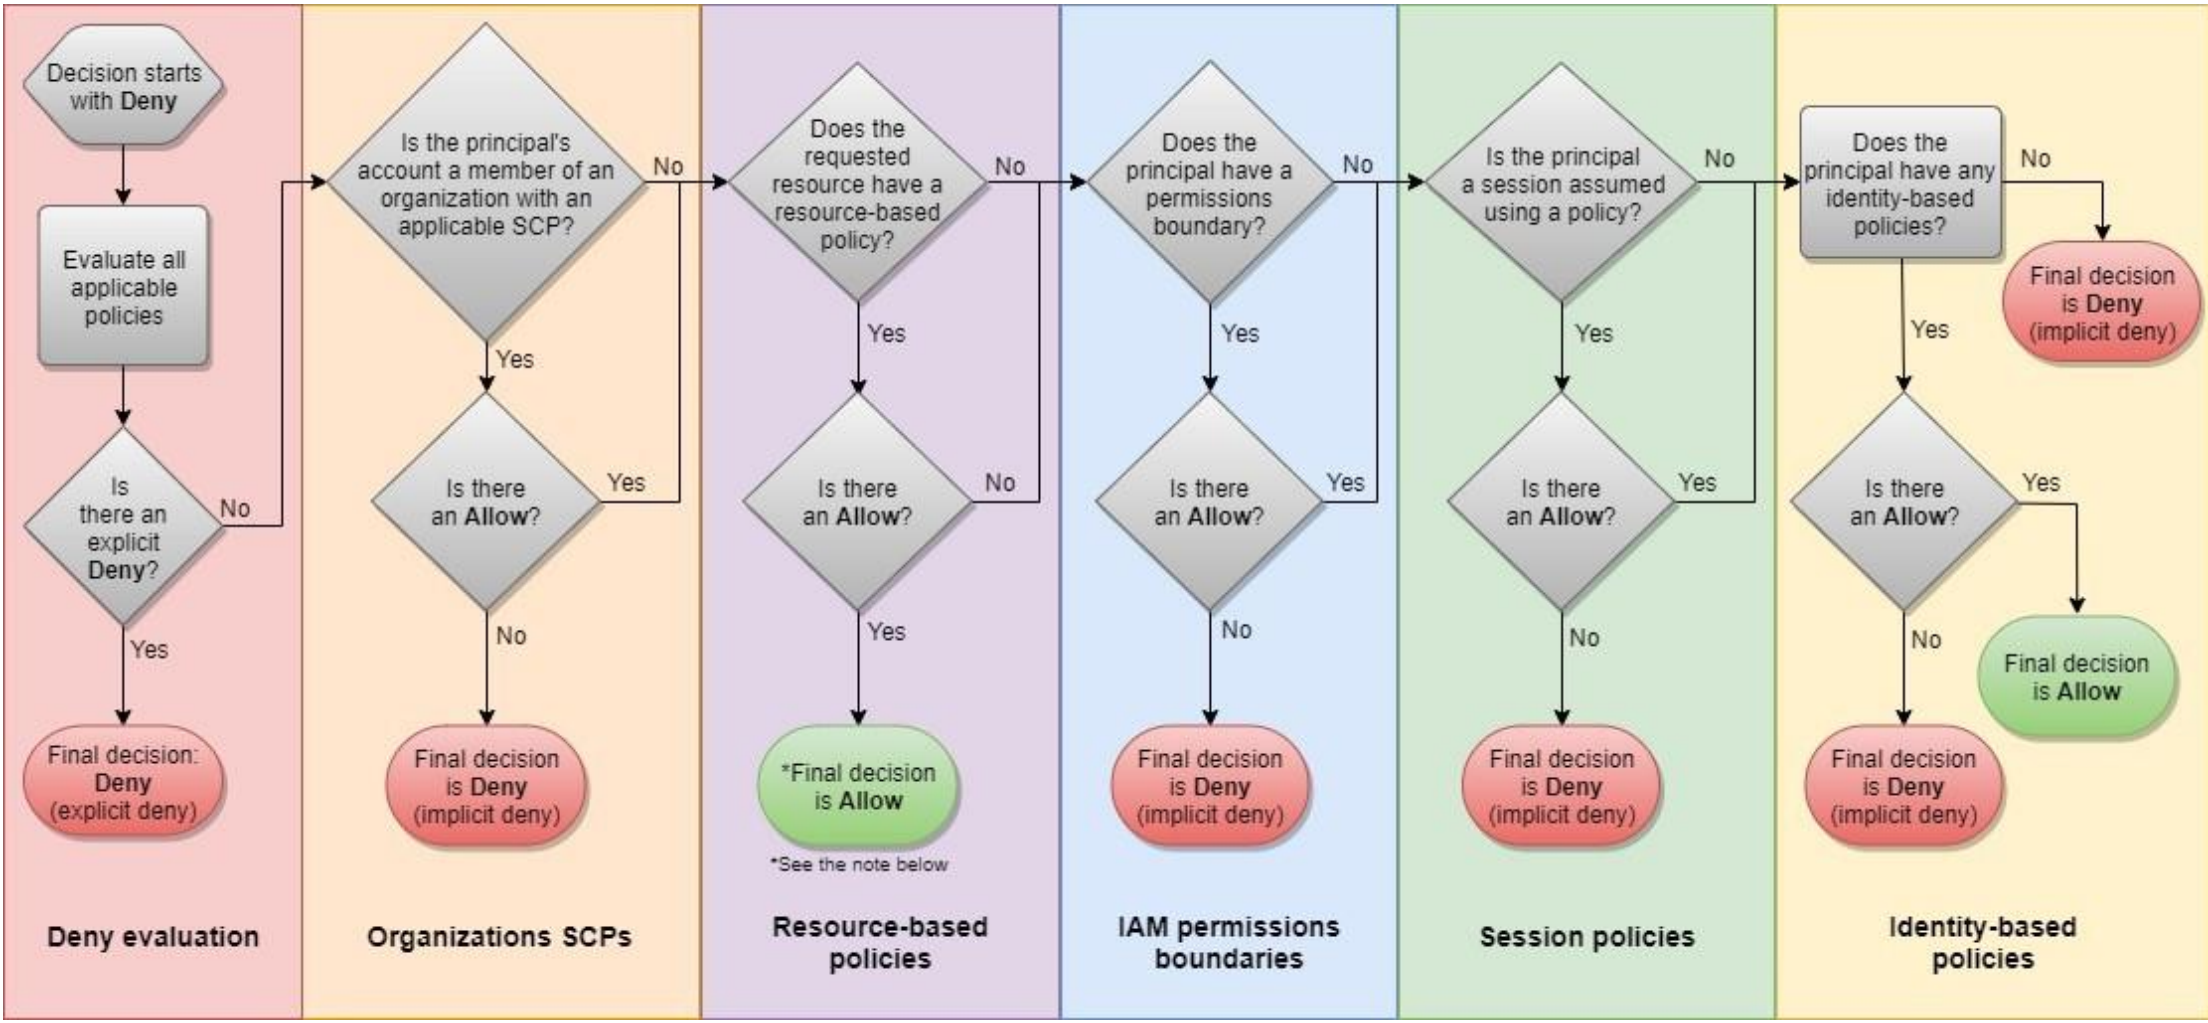
Privilege Escalation



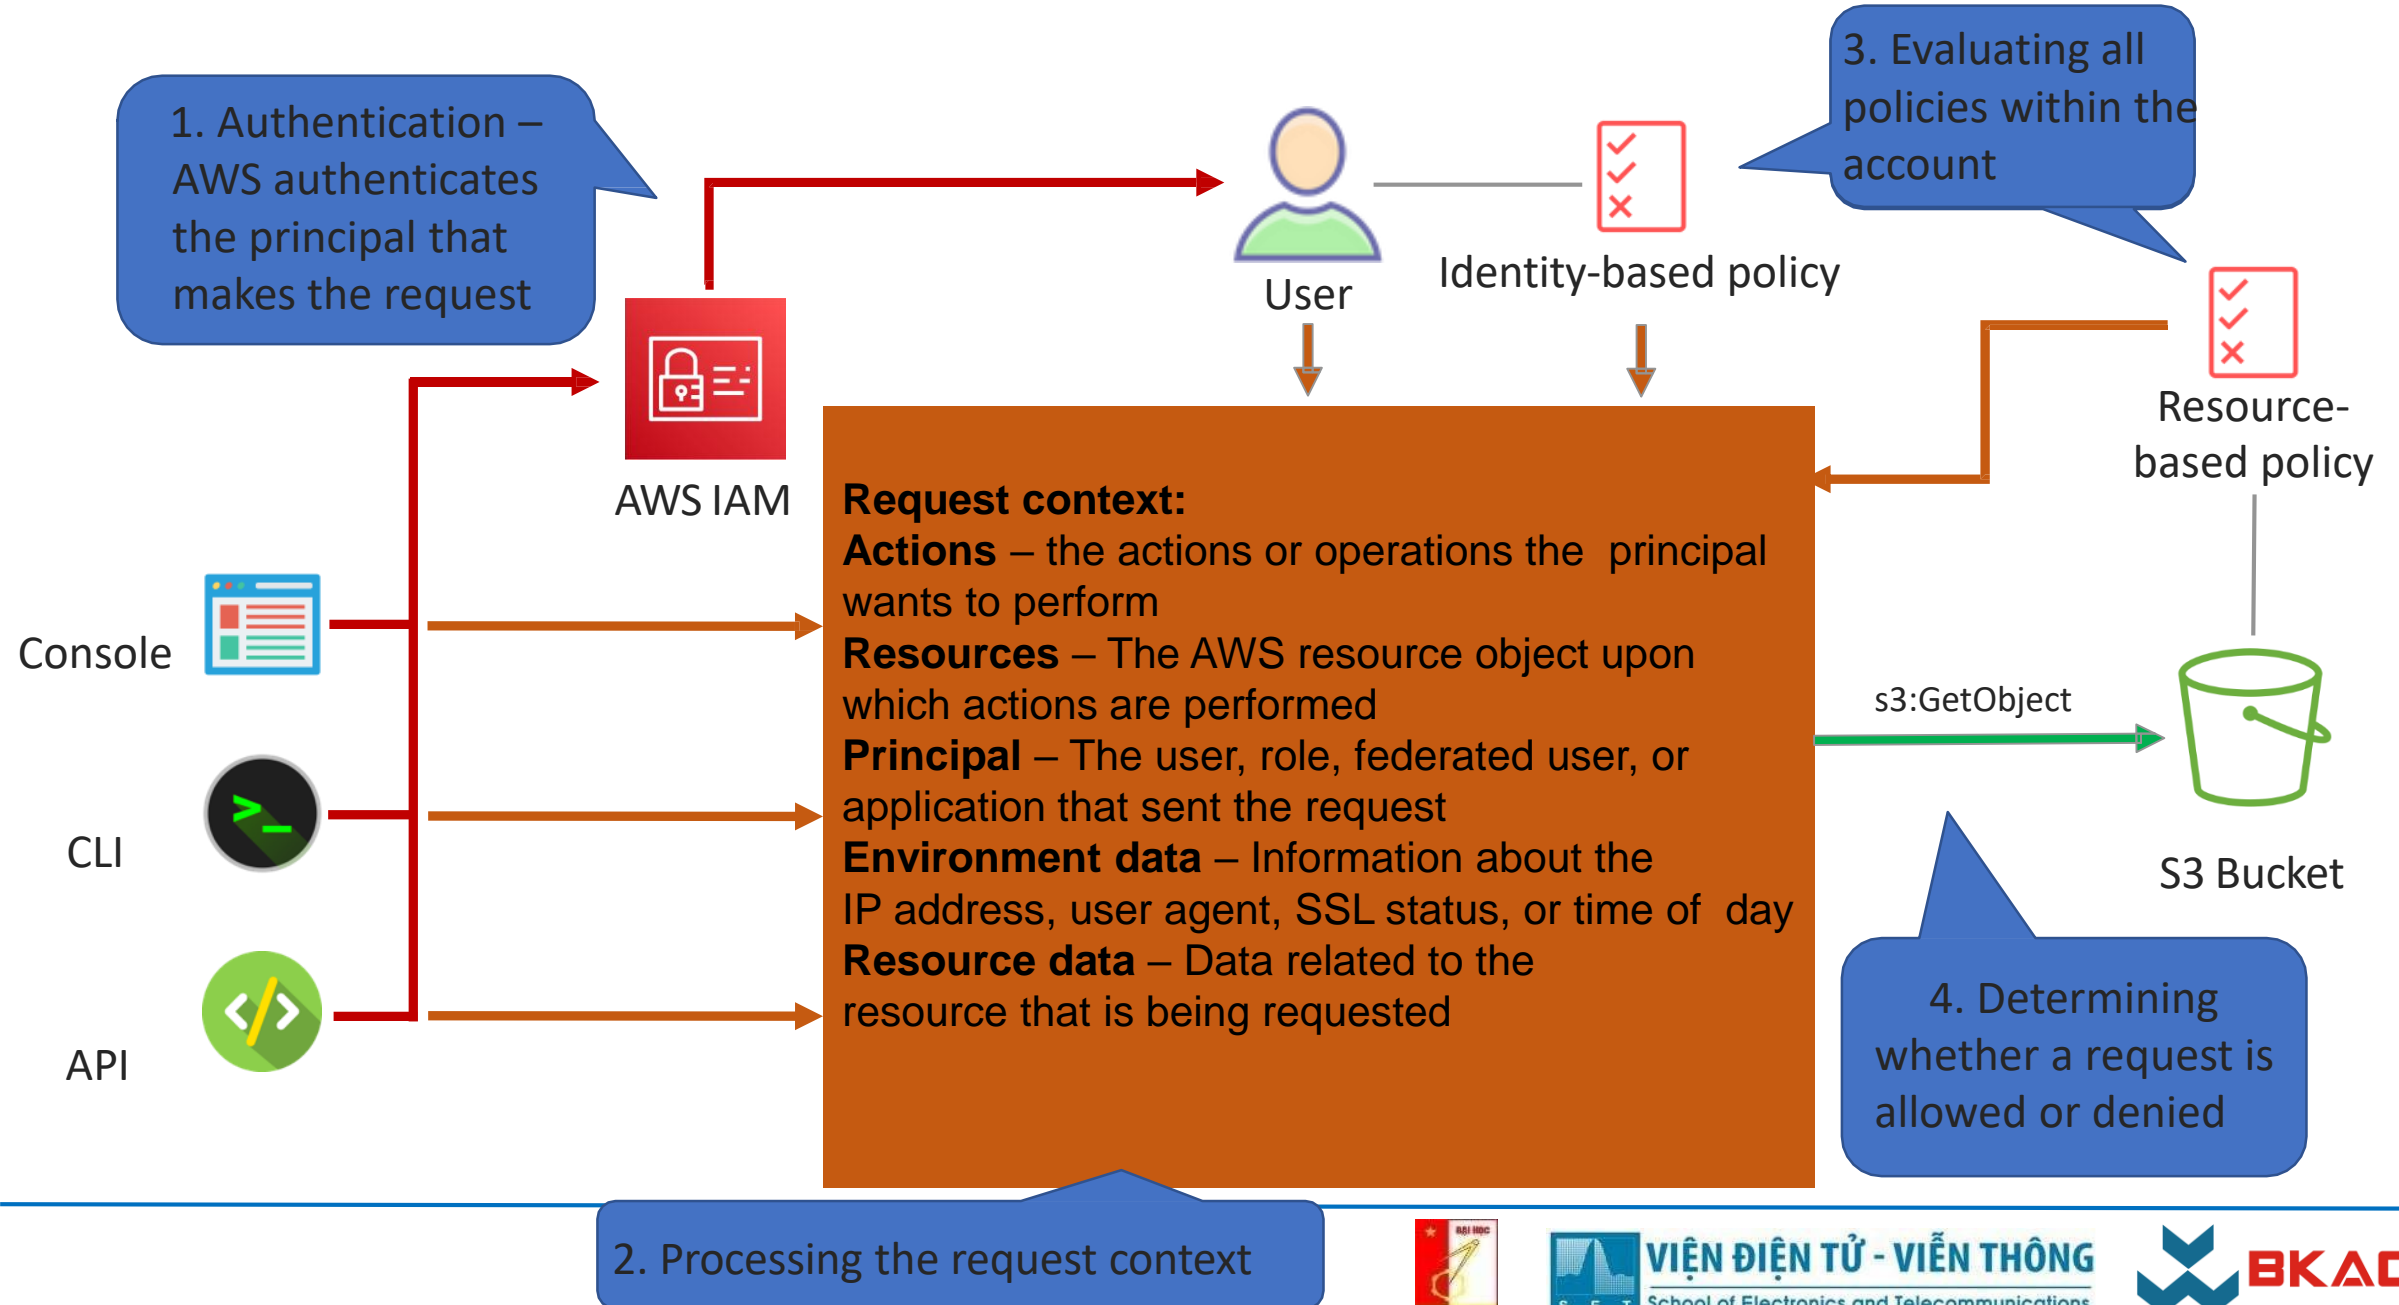
Preventing Privilege Escalation



Evaluation Logic



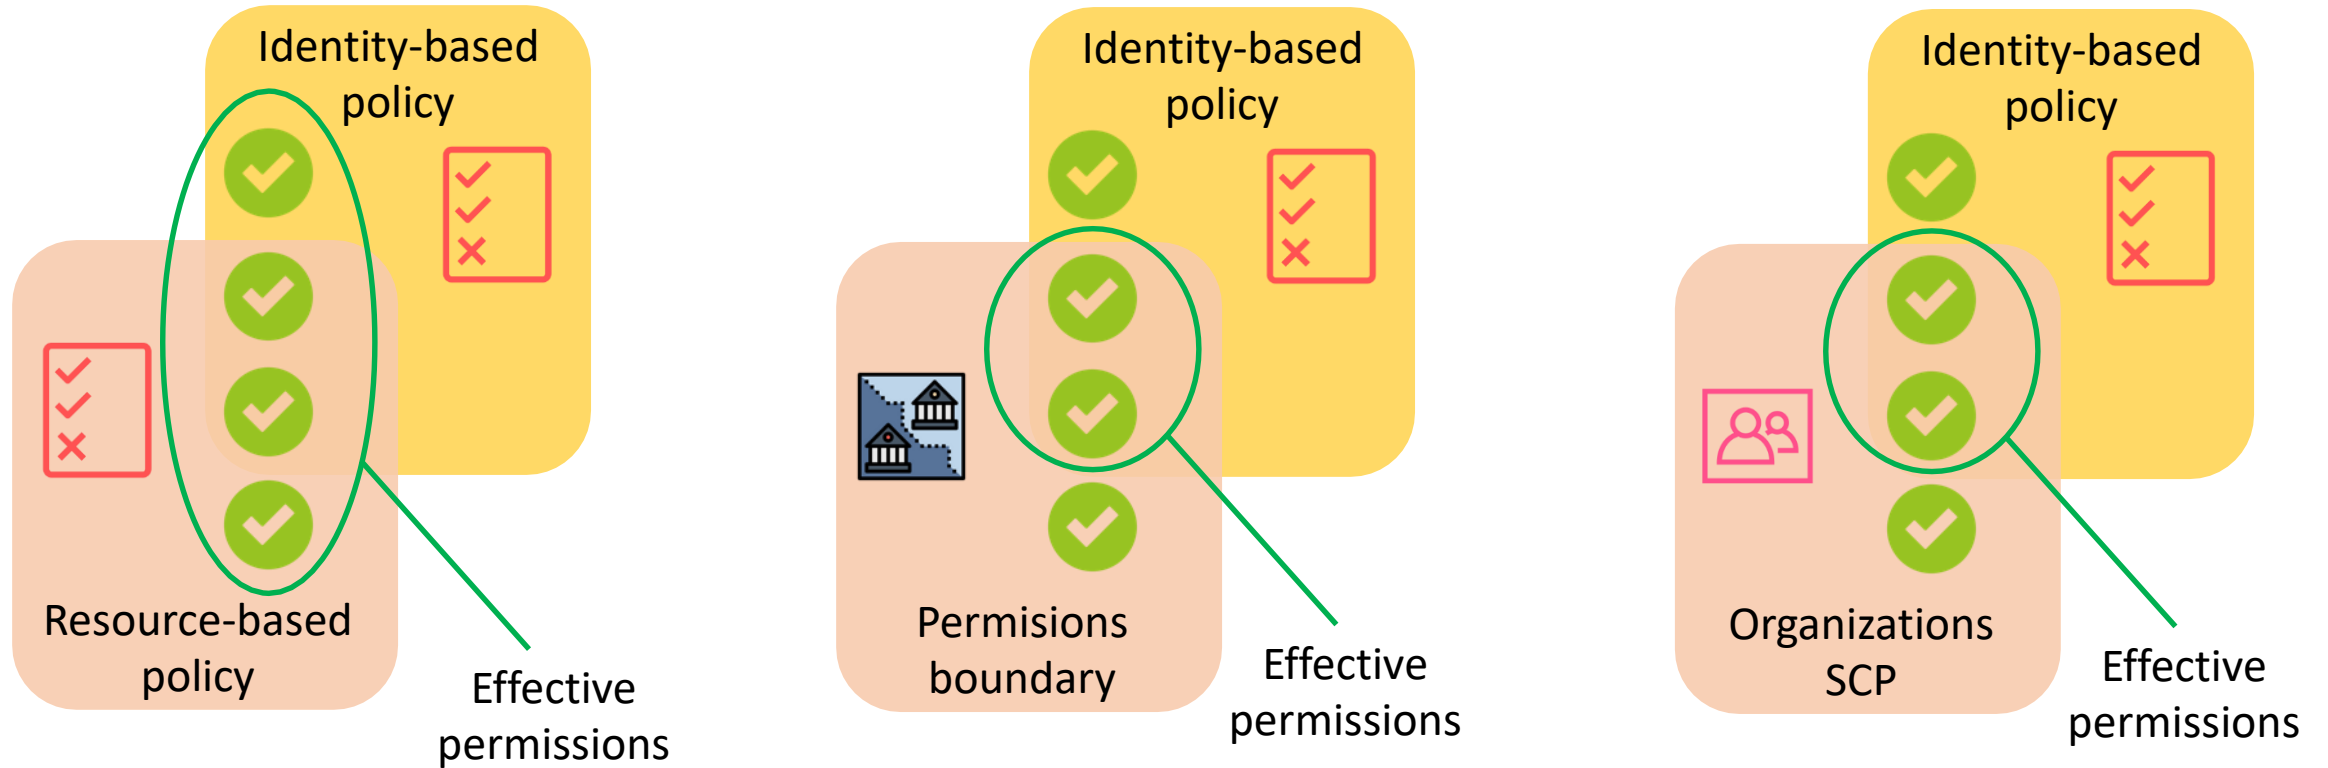
Steps for Authorizing Requests to AWS



Types of Policy

- **Identity-based policies** – attached to users, groups, or roles
- **Resource-based policies** – attached to a resource; define permissions for a principal accessing the resource
- **IAM permissions boundaries** – set the maximum permissions an identity-based policy can grant an IAM entity
- **AWS Organizations service control policies (SCP)** – specify the maximum permissions for an organization or OU
- **Session policies** – used with AssumeRole* API actions

Evaluating Policies within an AWS Account



Determination Rules

1. By default, all requests are implicitly denied (though the root user has full access)
2. An explicit allow in an identity-based or resource-based policy overrides this default
3. If a permissions boundary, Organizations SCP, or session policy is present, it might override the allow with an implicit deny
4. An explicit deny in any policy overrides any allows

IAM Policy Structure



IAM Policy Structure

An IAM policy is a JSON document that consists of one or more statements

The **Action** element is the specific API action for which you are granting or denying permission

```
{  
  "Statement": [{  
    "Effect": "effect",  
    "Action": "action",  
    "Resource": "arn",  
    "Condition": {  
      "condition": {  
        "key": "value"  
      }  
    }  
  }  
]  
}
```

The **Effect** element can be Allow or Deny

The **Resource** element specifies the resource that's affected by the action

The **Condition** element is optional and can be used to control when your policy is in effect

IAM Policy Example 1

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "*",
      "Resource": "*"
    }
  ]
}
```

The AdministratorAccess policy uses wildcards (*) to allow all actions on all resources

IAM Policy Example 2

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": ["ec2:TerminateInstances"],
      "Resource": ["*"]
    },
    {
      "Effect": "Deny",
      "Action": ["ec2:TerminateInstances"],
      "Condition": {
        "NotIpAddress": {
          "aws:SourceIp": [
            "192.0.2.0/24",
            "203.0.113.0/24"
          ]
        }
      },
      "Resource": ["*"]
    }
  ]
}
```

The specific API action is defined

The effect is to deny the API action if the IP address is not in the specified range

IAM Policy Example 3

```
{
  "Version": "2012-10-17",
  "Id": "ExamplePolicy01",
  "Statement": [
    {
      "Sid": "ExampleStatement01",
      "Effect": "Allow",
      "Principal": {
        "AWS": "*"
      },
      "Action": [
        "elasticfilesystem:ClientRootAccess",
        "elasticfilesystem:ClientMount",
        "elasticfilesystem:ClientWrite"
      ],
      "Condition": {
        "Bool": {
          "aws:SecureTransport": "true"
        }
      }
    }
  ]
}
```

You can tell this is a resource-based policy as it has a principal element defined

The policy grants read and write access to an EFS file systems to all IAM principals ("AWS ": "*")

Additionally, the policy condition element requires that SSL/TLS encryption is used

IAM Policy Example 4

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": ["s3:ListBucket"],
      "Effect": "Allow",
      "Resource": ["arn:aws:s3:::mybucket"],
      "Condition": {"StringLike": {"s3:prefix": ["${aws:username}/*"]}}
    },
    {
      "Action": [
        "s3:GetObject",
        "s3:PutObject"
      ],
      "Effect": "Allow",
      "Resource": ["arn:aws:s3:::mybucket/${aws:username}/*"]
    }
  ]
}
```

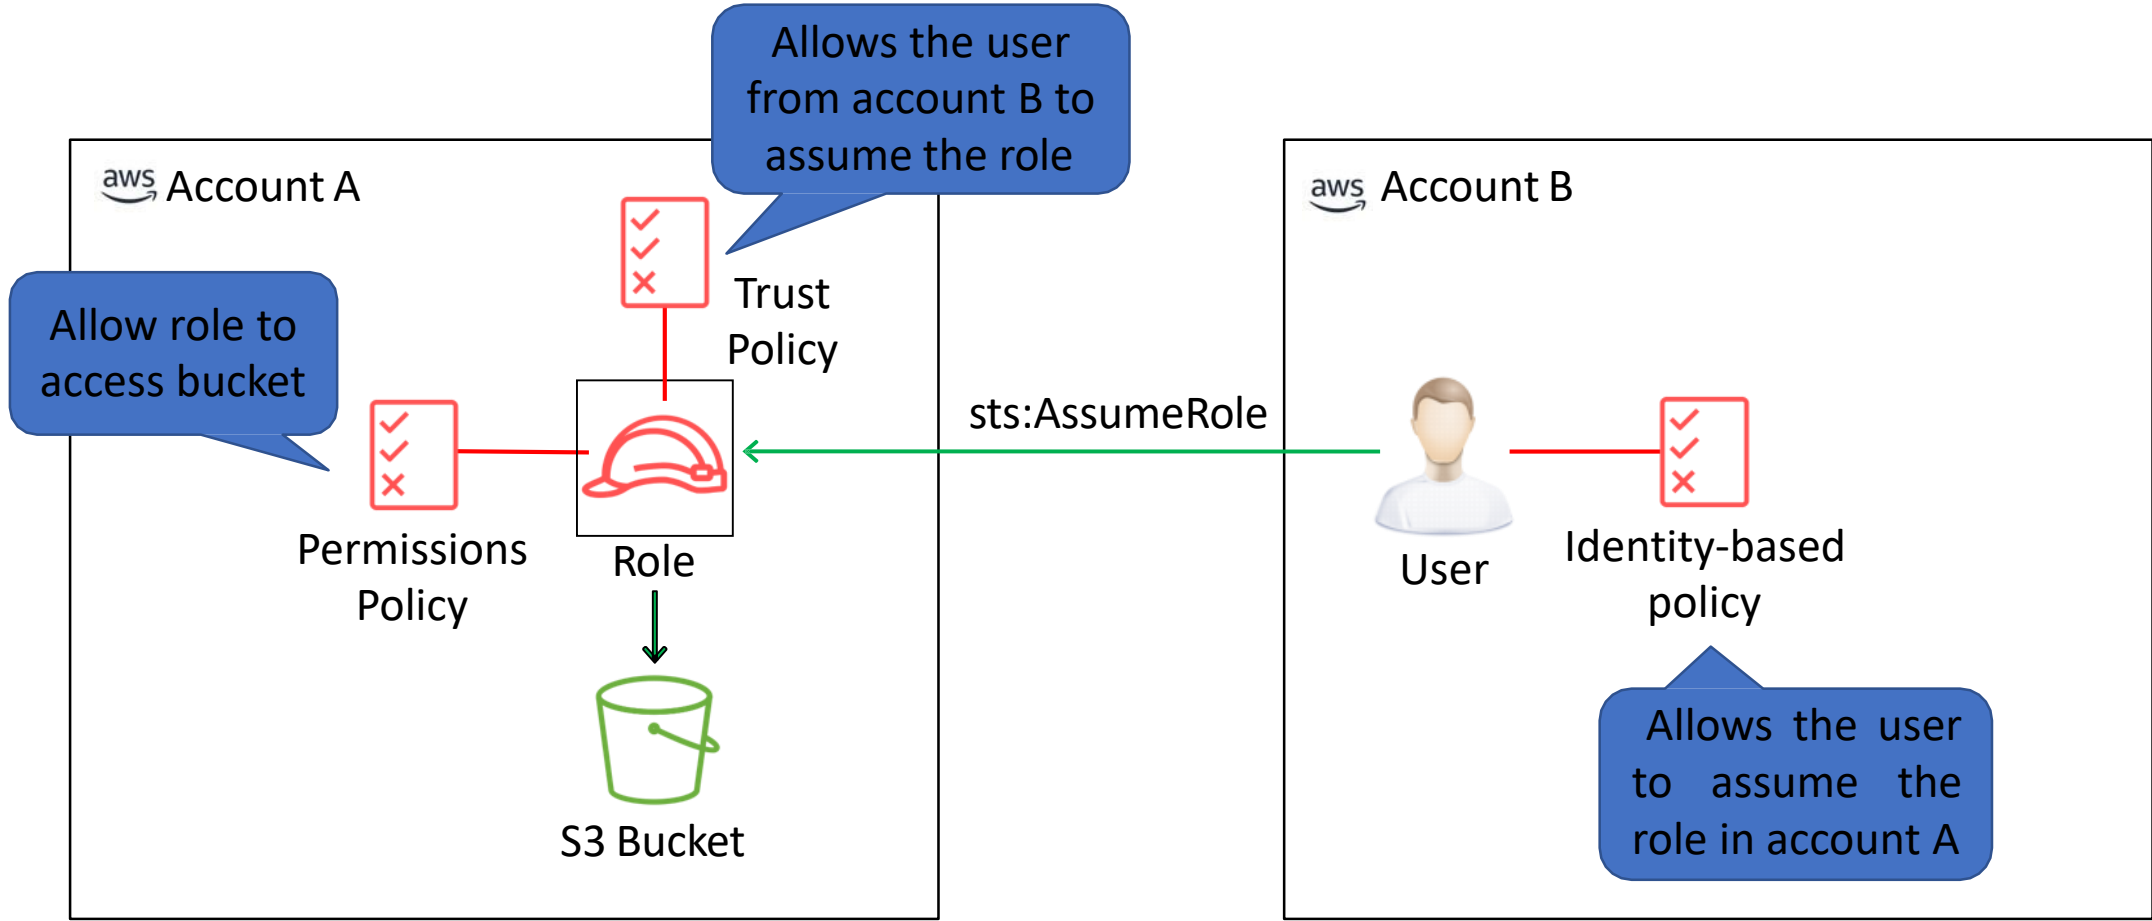
A variable is used for the s3:prefix that is replaced with the user's friendly name

The actions are allowed only within the user's folder within the bucket

Use Cases for IAM Roles



Use Case: Cross Account Access

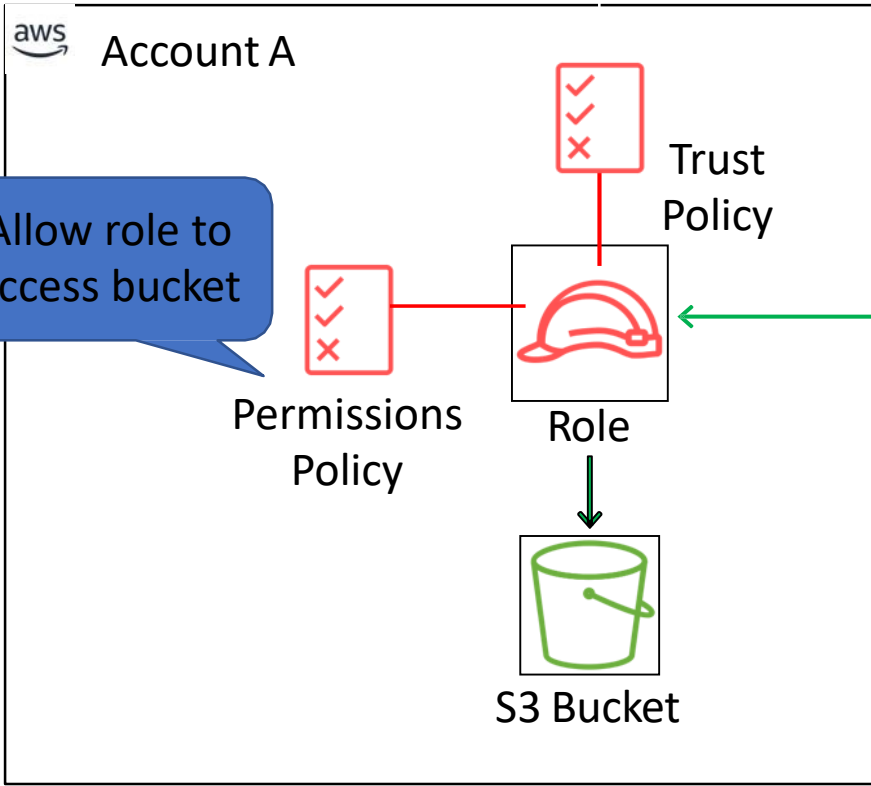


Use Case: Cross Account Access (3rd Party)

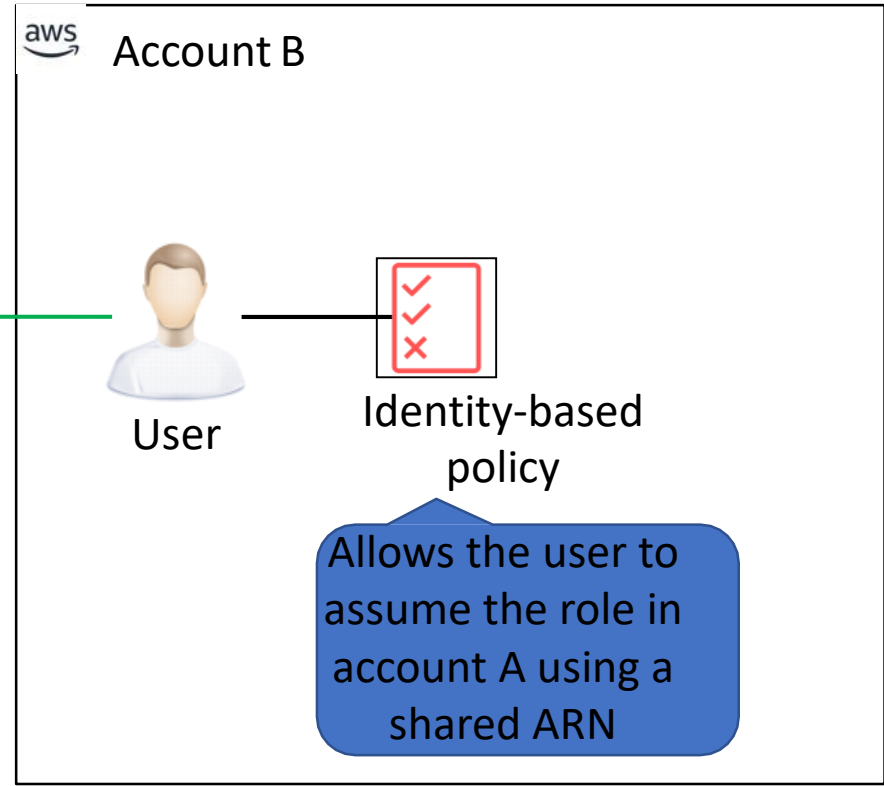
The trust policy condition requires the external ID

```
"Statement": {  
  "Effect": "Allow",  
  "Action": "sts:AssumeRole",  
  "Principal": {"AWS": "3rd party AWS Account ID"},  
  "Condition": {"StringEquals": {"sts:ExternalId": "12345"}}  
}
```

Allow role to access bucket

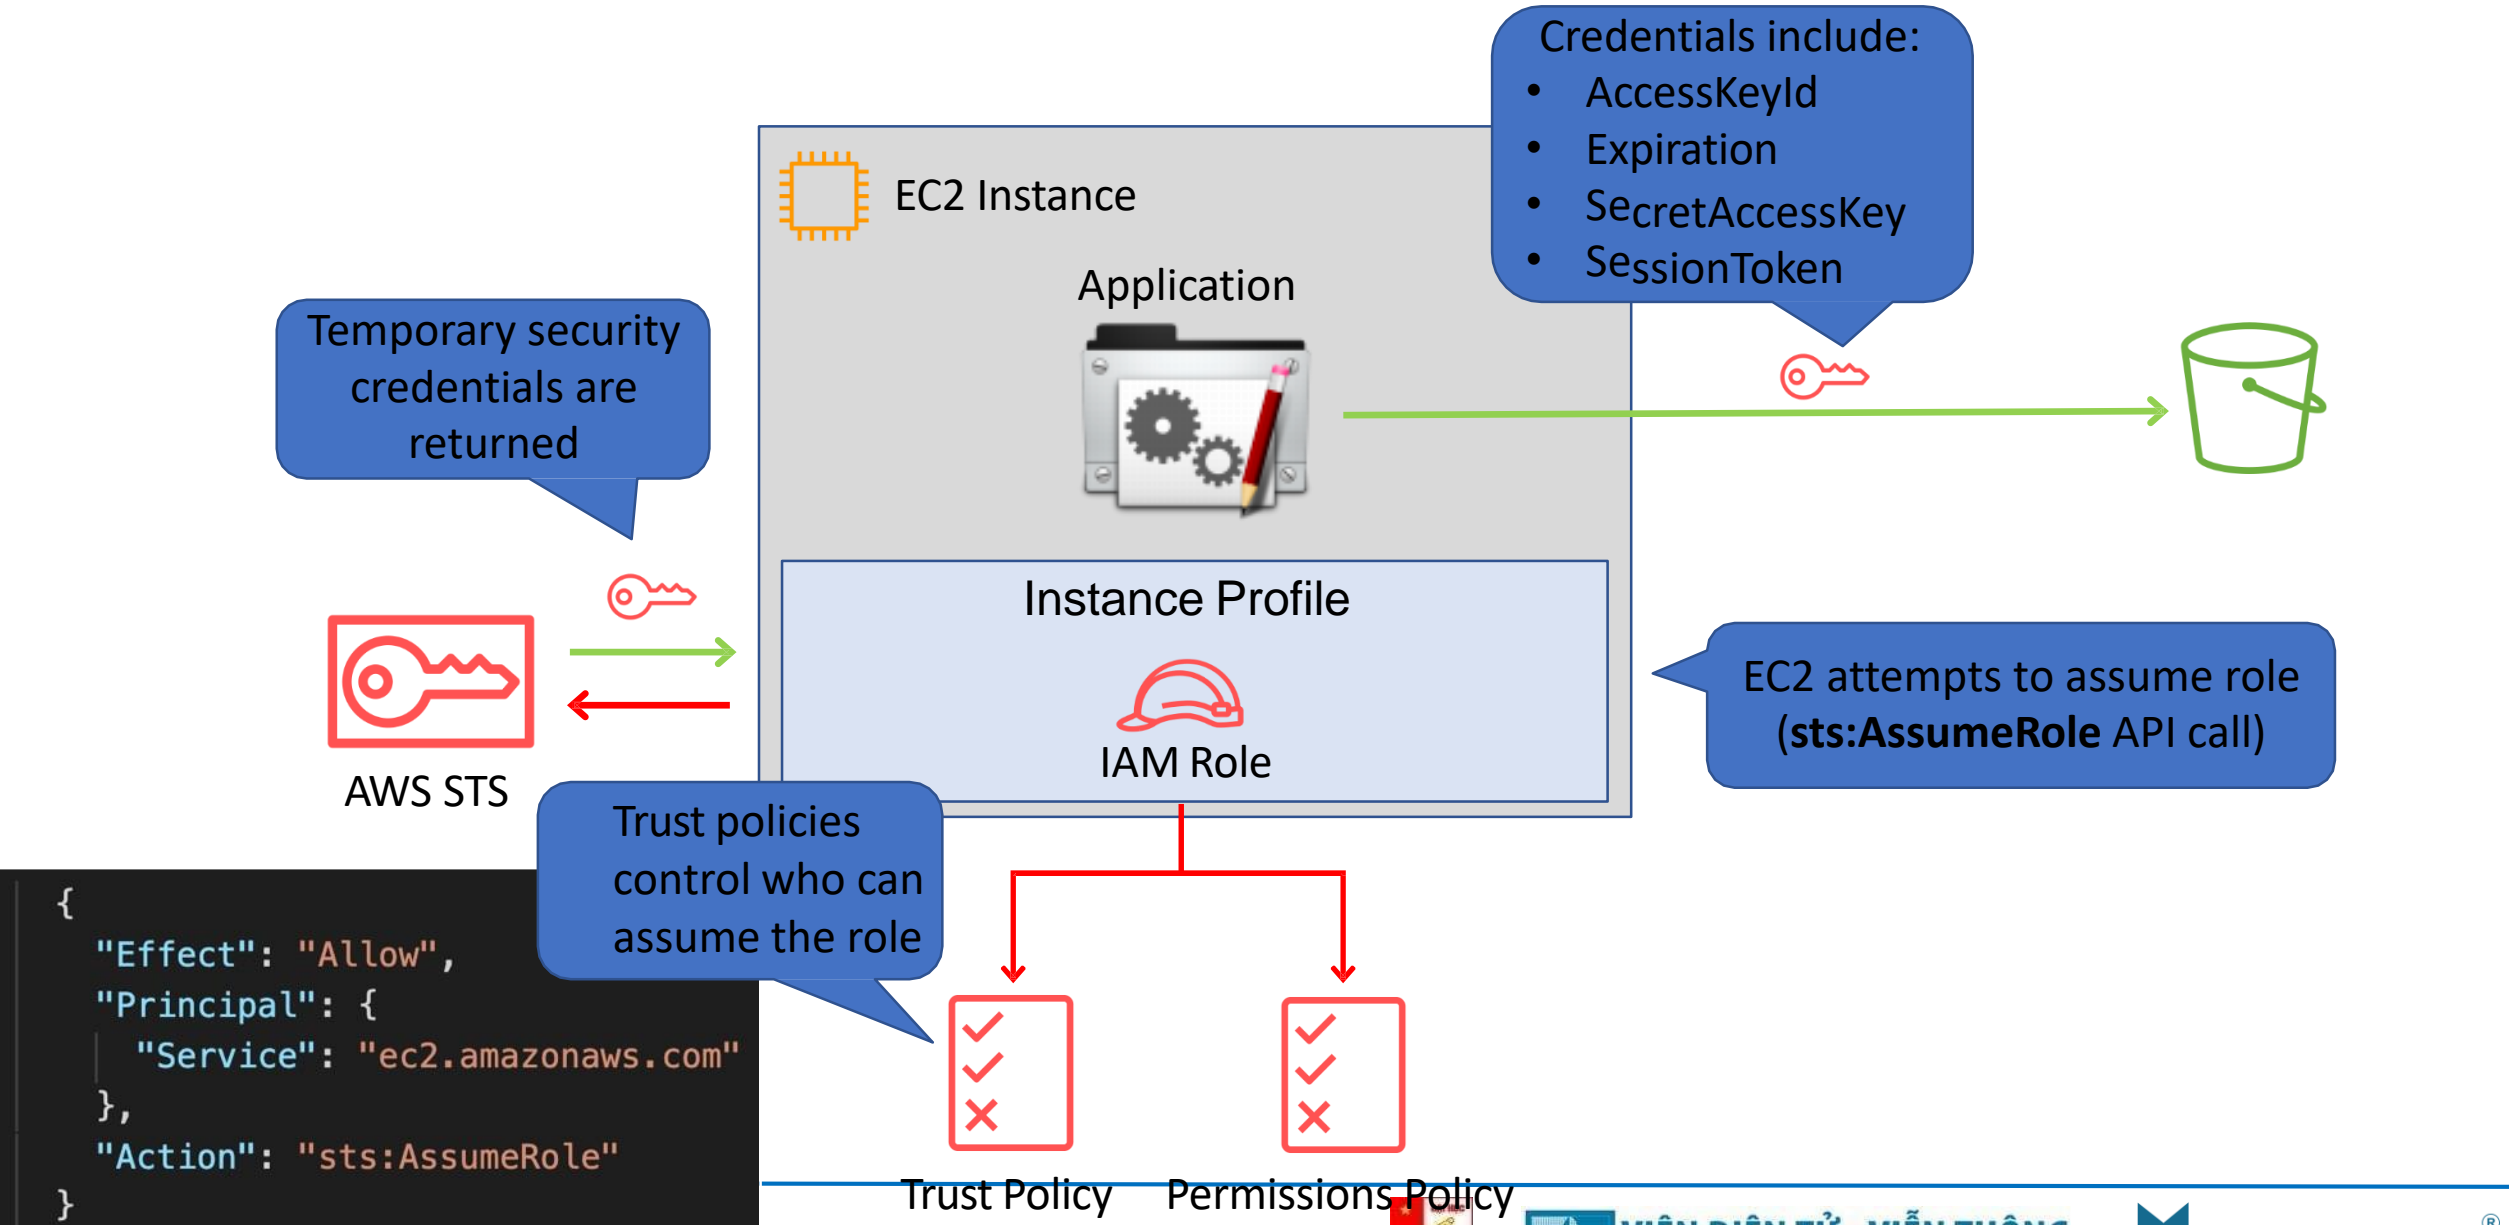


sts:AssumeRole with external ID



Allows the user to assume the role in account A using a shared ARN

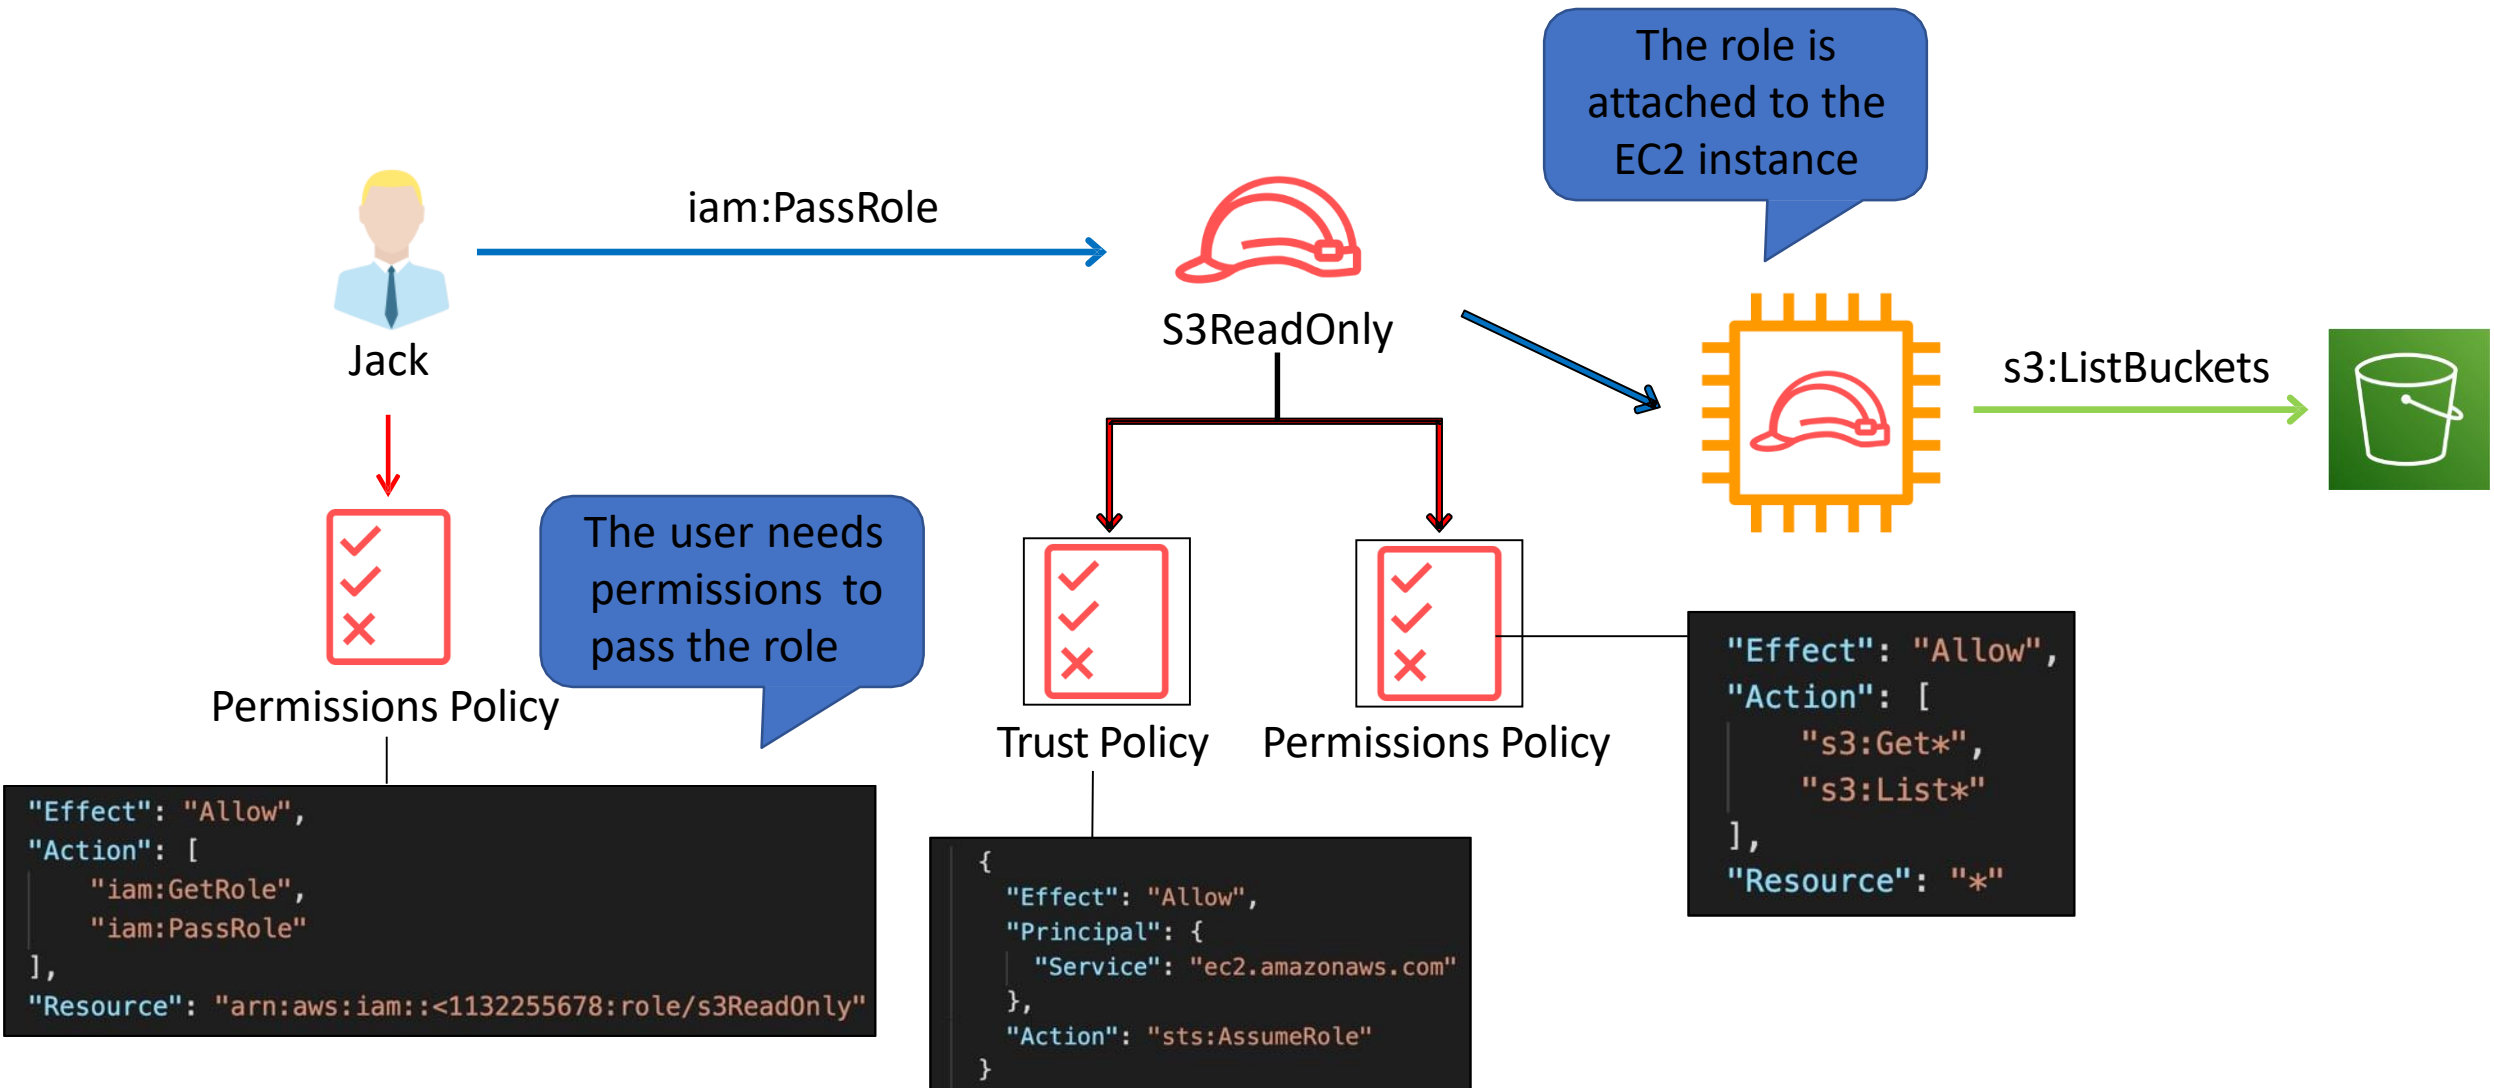
Use Case: Delegation to AWS Services



Amazon EC2 Instance Profile



Attach Role to EC2 Instance



AWS IAM Best Practices



AWS IAM Best Practices

- Lock away your AWS account root user access keys
- Create individual IAM users
- Use groups to assign permissions to IAM users
- Grant least privilege
- Get started using permissions with AWS managed policies
- Use customer managed policies instead of inline policies
- Use access levels to review IAM permissions
- Configure a strong password policy for your users
- Enable MFA

AWS IAM Best Practices

- Use roles for applications that run on Amazon EC2 instances
- Use roles to delegate permissions
- Do not share access keys
- Rotate credentials regularly
- Remove unnecessary credentials
- Use policy conditions for extra security
- Monitor activity in your AWS account

Q & A



VIỆN ĐIỆN TỬ - VIỄN THÔNG
School of Electronics and Telecommunications



Module 5 DNS, Caching, and Performance Optimization



Agenda

- Amazon Route 53
- Amazon CloudFront

Amazon Route 53 Hosted Zones



Amazon Route 53 Hosted Zones

Name	Type	Value	TTL
example.com	A	8.1.2.1	60
dev.example.com	A	8.1.2.2	60



Amazon Route 53

This is an example of a **public hosted zone**

A **hosted zone** represents a set of records belonging to a domain

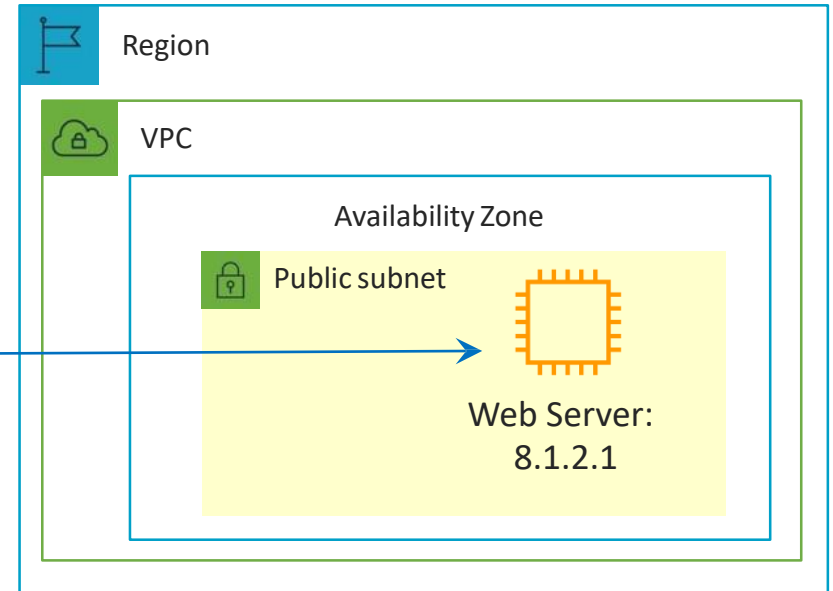
What's the address for example.com?

example.com

Address is 8.1.2.1



HTTP GET to 8.1.2.1



Amazon Route 53 Hosted Zones

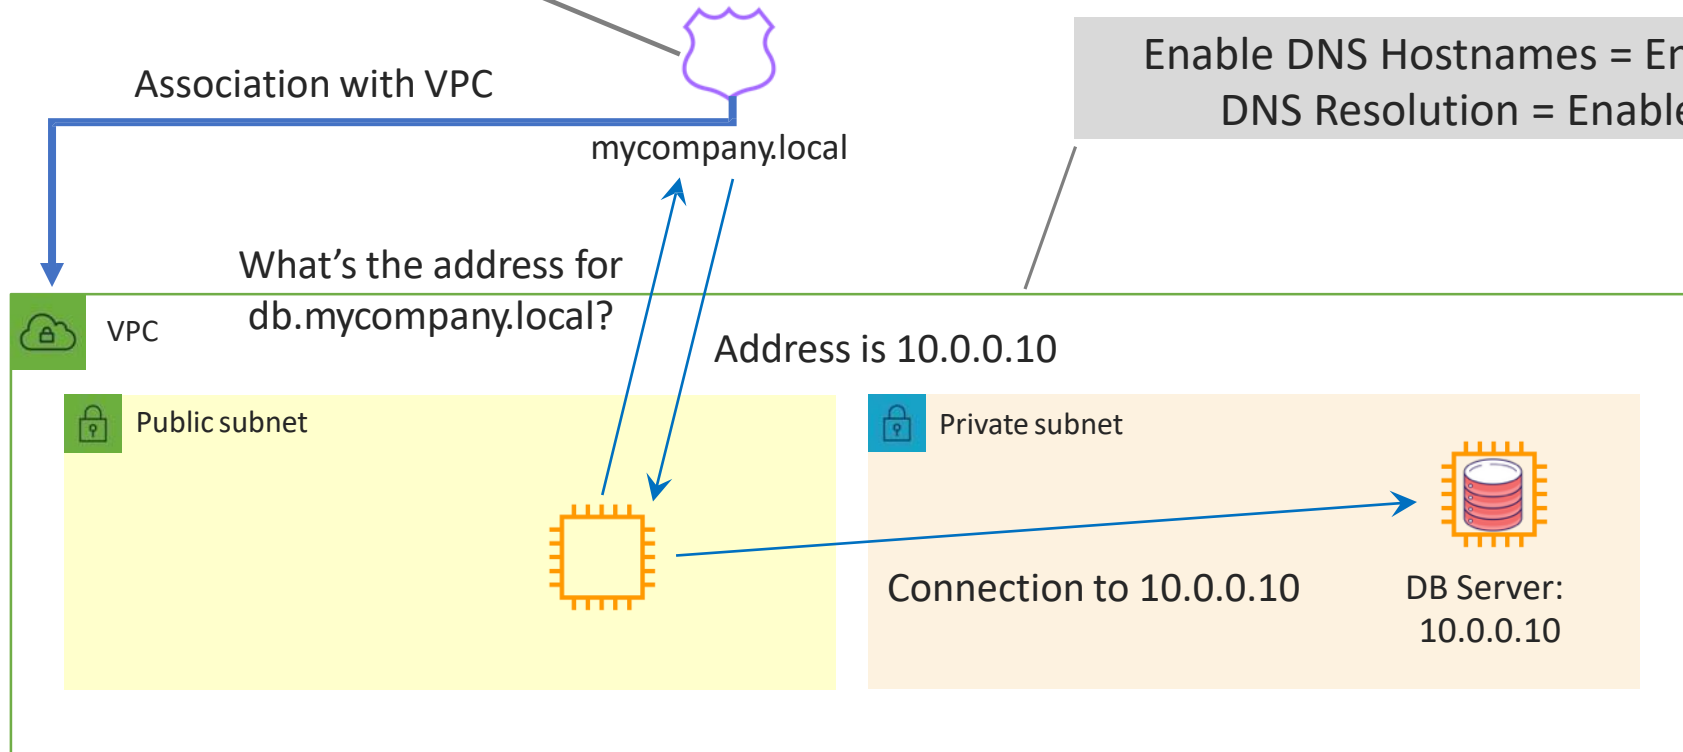
Name	Type	Value	TTL
db.mycompany.local	A	10.0.0.10	60
app.mycompany.local	A	10.0.0.11	60



Amazon Route 53

This is an example of a **private hosted zone**

Enable DNS Hostnames = Enabled
DNS Resolution = Enabled



Migration to/from Route 53

- You can migrate from **another DNS provider** and can import records
- You can migrate a hosted zone to **another AWS account**
- You can migrate from Route 53 to **another registrar**
- You can also associate a Route 53 hosted zone with a **VPC in another account**
 - Authorize association with VPC in the second account.
 - Create an association in the second account

Route 53 Routing Policies



VIỆN ĐIỆN TỬ - VIỆN THÔNG
School of Electronics and Telecommunications



Amazon Route 53 Routing Policies


Routing Policy	What it does
Simple	Simple DNS response providing the IP address associated with a name
Failover	If primary is down (based on health checks), routes to secondary destination
Geolocation	Uses geographic location you're in (e.g. Europe) to route you to the closest region
Geoproximity	Routes you to the closest region within a geographic area
Latency	Directs you based on the lowest latency route to resources
Multivalue answer	Returns several IP addresses and functions as a basic load balancer
Weighted	Uses the relative weights assigned to resources to determine which to route to

Amazon Route 53 – Simple Routing Policy

Name	Type	Value	TTL
simple.dctlabs.com	A	1.1.1.1	60
		2.2.2.2	
simple2.dctlabs.com	A	3.3.3.3	60



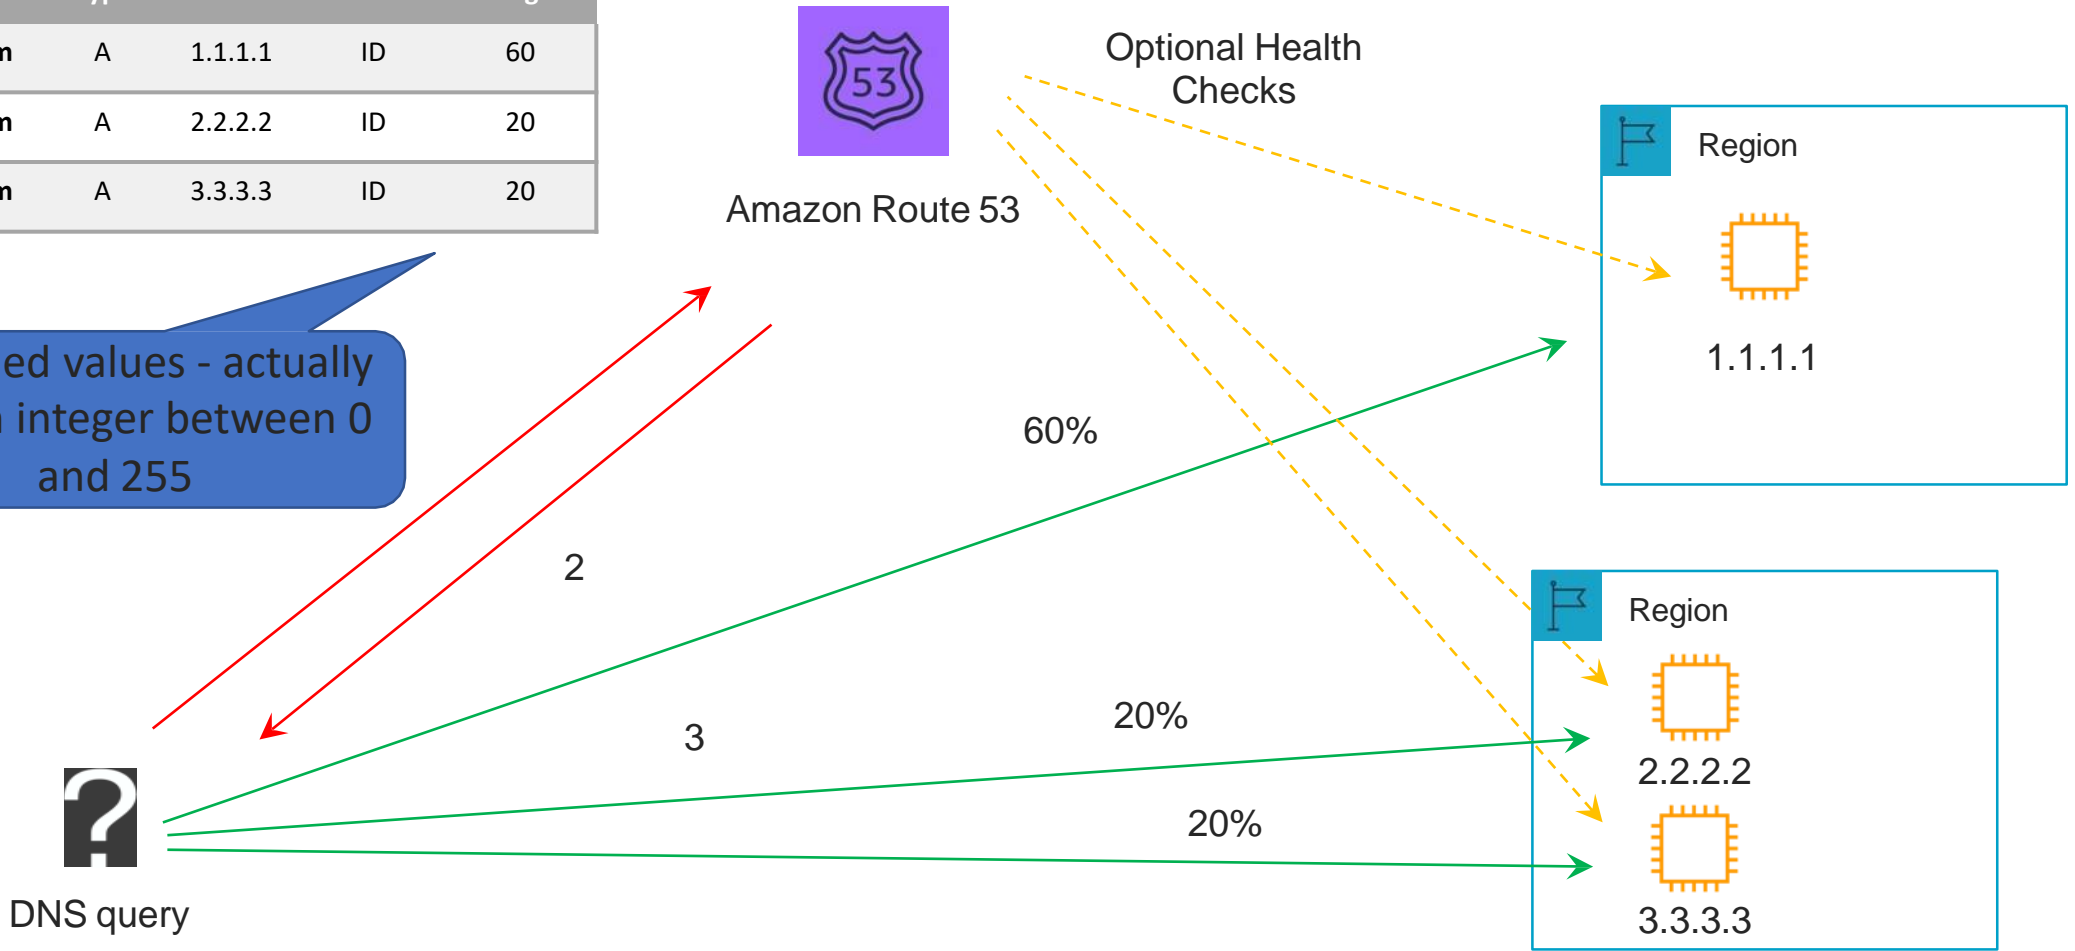
Amazon Route 53



Amazon Route 53 – Weighted Routing Policy

Name	Type	Value	Health	Weight
weighted.dctlabs.com	A	1.1.1.1	ID	60
weighted.dctlabs.com	A	2.2.2.2	ID	20
weighted.dctlabs.com	A	3.3.3.3	ID	20

Simplified values - actually uses an integer between 0 and 255



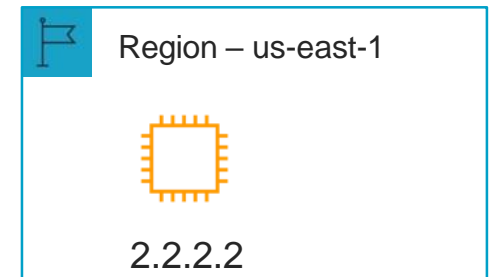
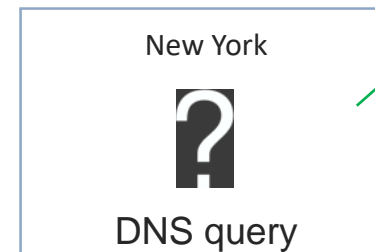

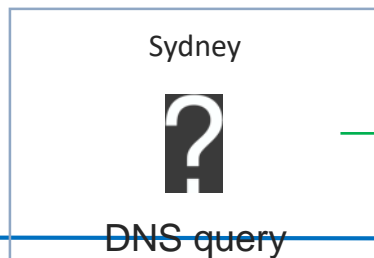
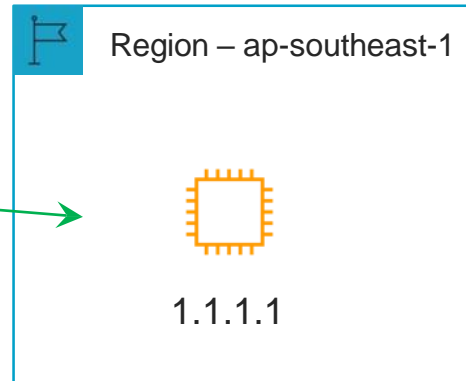

Amazon Route 53 – Latency Routing Policy

Name	Type	Value	Health	Region
latency.dctlabs.com	A	1.1.1.1	ID	ap-southeast-1
latency.dctlabs.com	A	2.2.2.2	ID	us-east-1
latency.dctlabs.com	A	alb-id	ID	ap-southeast-2



Amazon Route 53

Optional Health Checks



Amazon Route 53 – Failover Routing Policy

Name	Type	Value	Health	Record Type
failover.dctlabs.com	A	1.1.1.1	ID	Primary
failover.dctlabs.com	A	alb-id		Secondary




Amazon Route 53

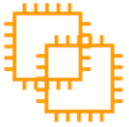

Health check is required on Primary


DNS query

Region – us-east-1


1.1.1.1

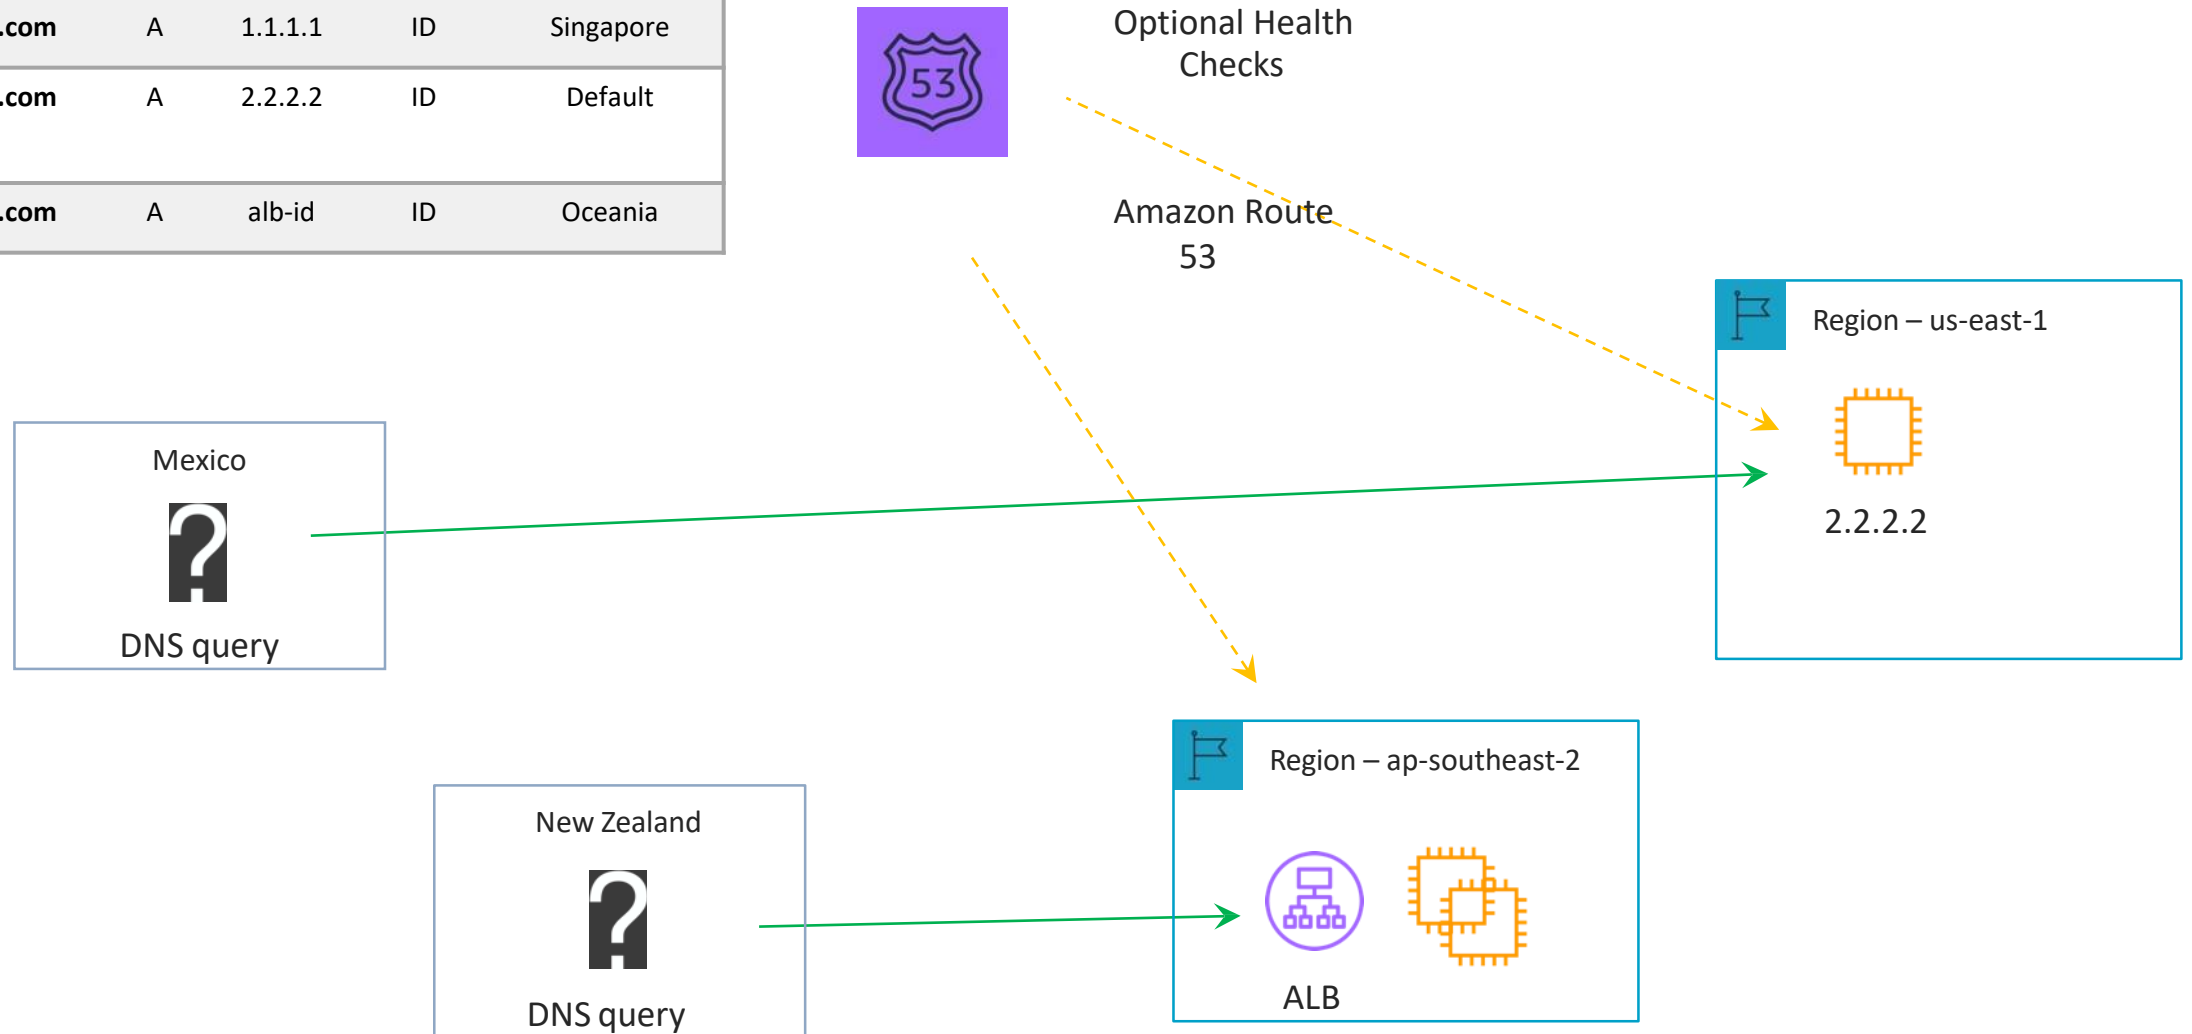
Region – ap-southeast-2


ALB

ap-southeast-2 is the secondary Region

Amazon Route 53 – Geolocation Routing Policy

Name	Type	Value	Health	Geolocation
geolocation.dctlabs.com	A	1.1.1.1	ID	Singapore
geolocation.dctlabs.com	A	2.2.2.2	ID	Default
geolocation.dctlabs.com	A	alb-id	ID	Oceania




Amazon Route 53 – Multivalue Routing Policy

Name	Type	Value	Health	Multi Value
multivalue.dctlabs.com	A	1.1.1.1	ID	Yes
multivalue.dctlabs.com	A	2.2.2.2	ID	Yes
multivalue.dctlabs.com	A	3.3.3.3	ID	Yes



Amazon Route 53

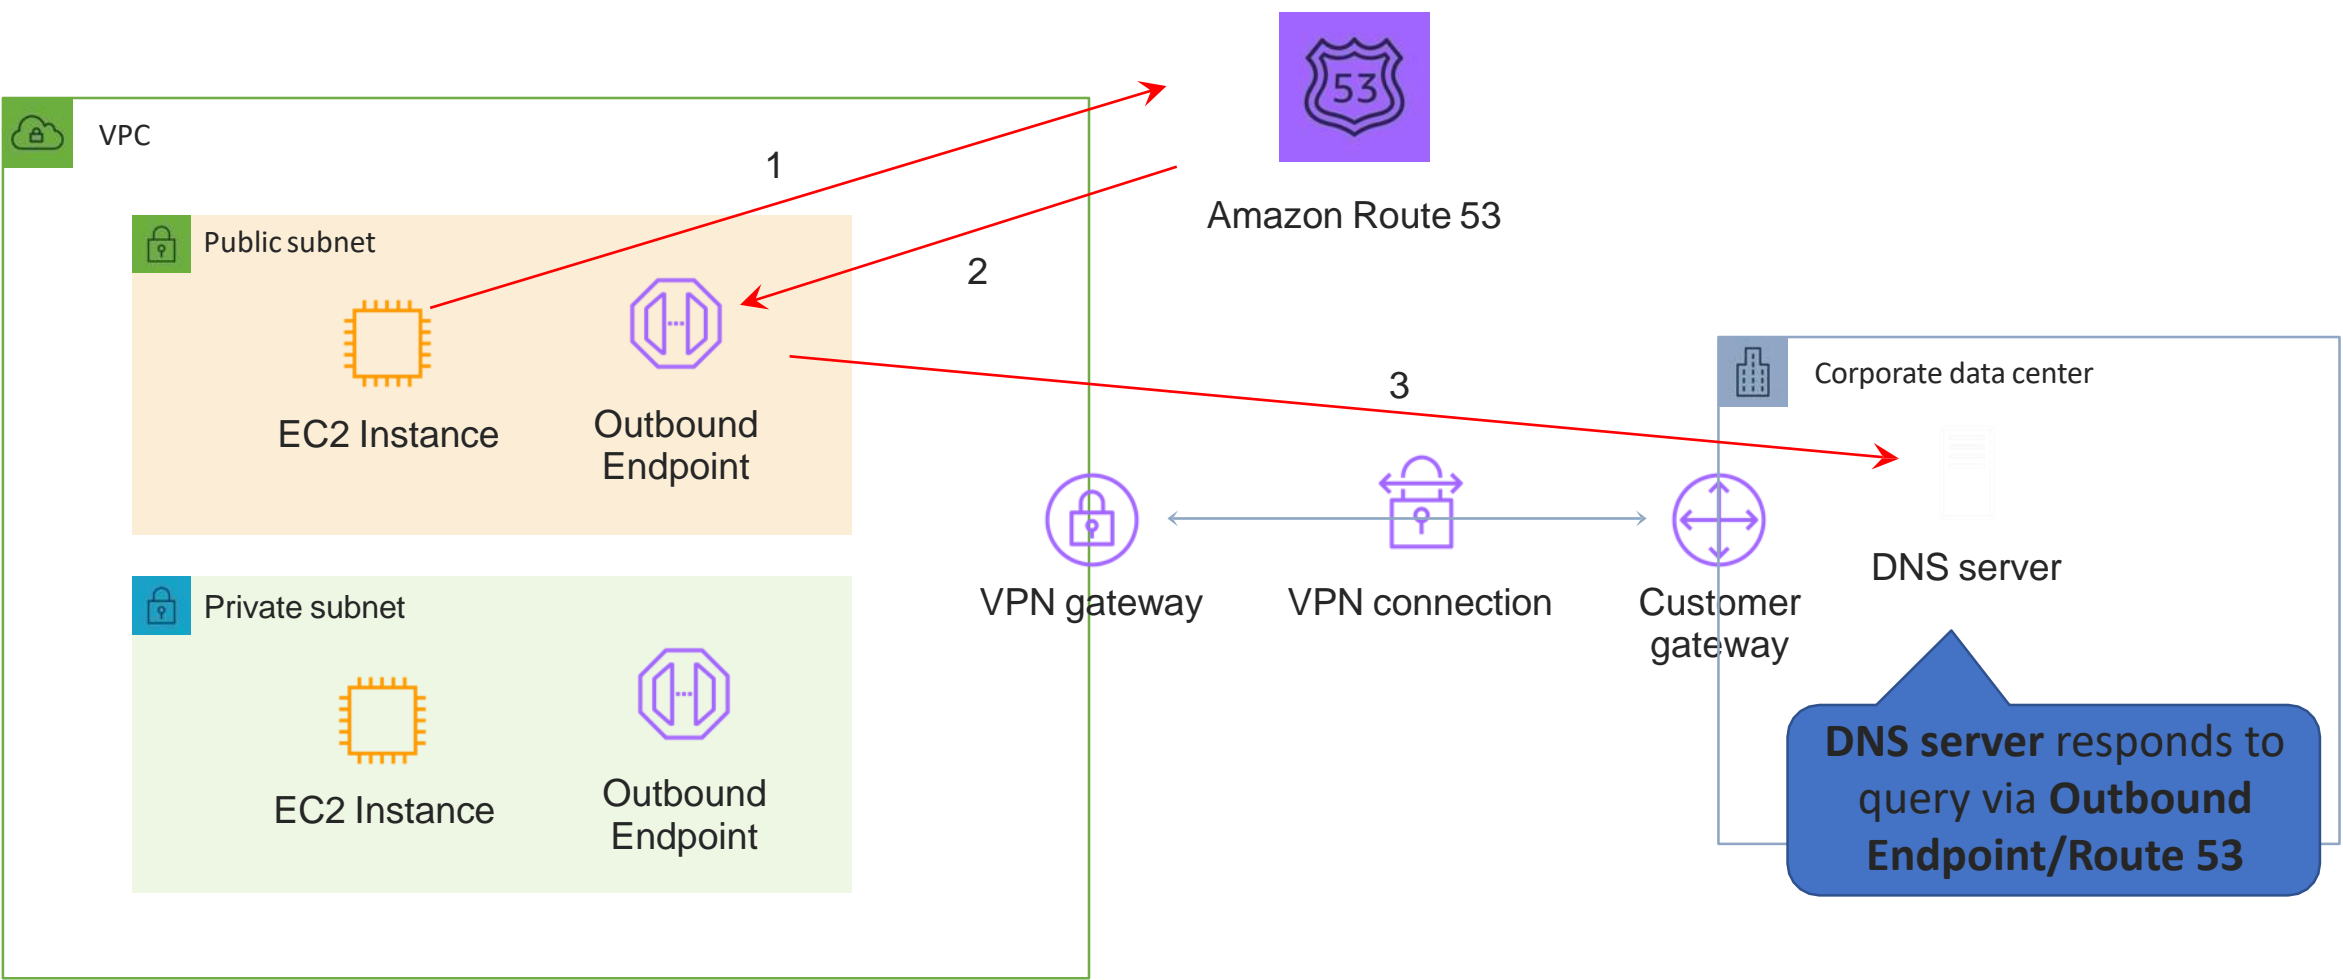
Health Checks:
returns healthy
records only



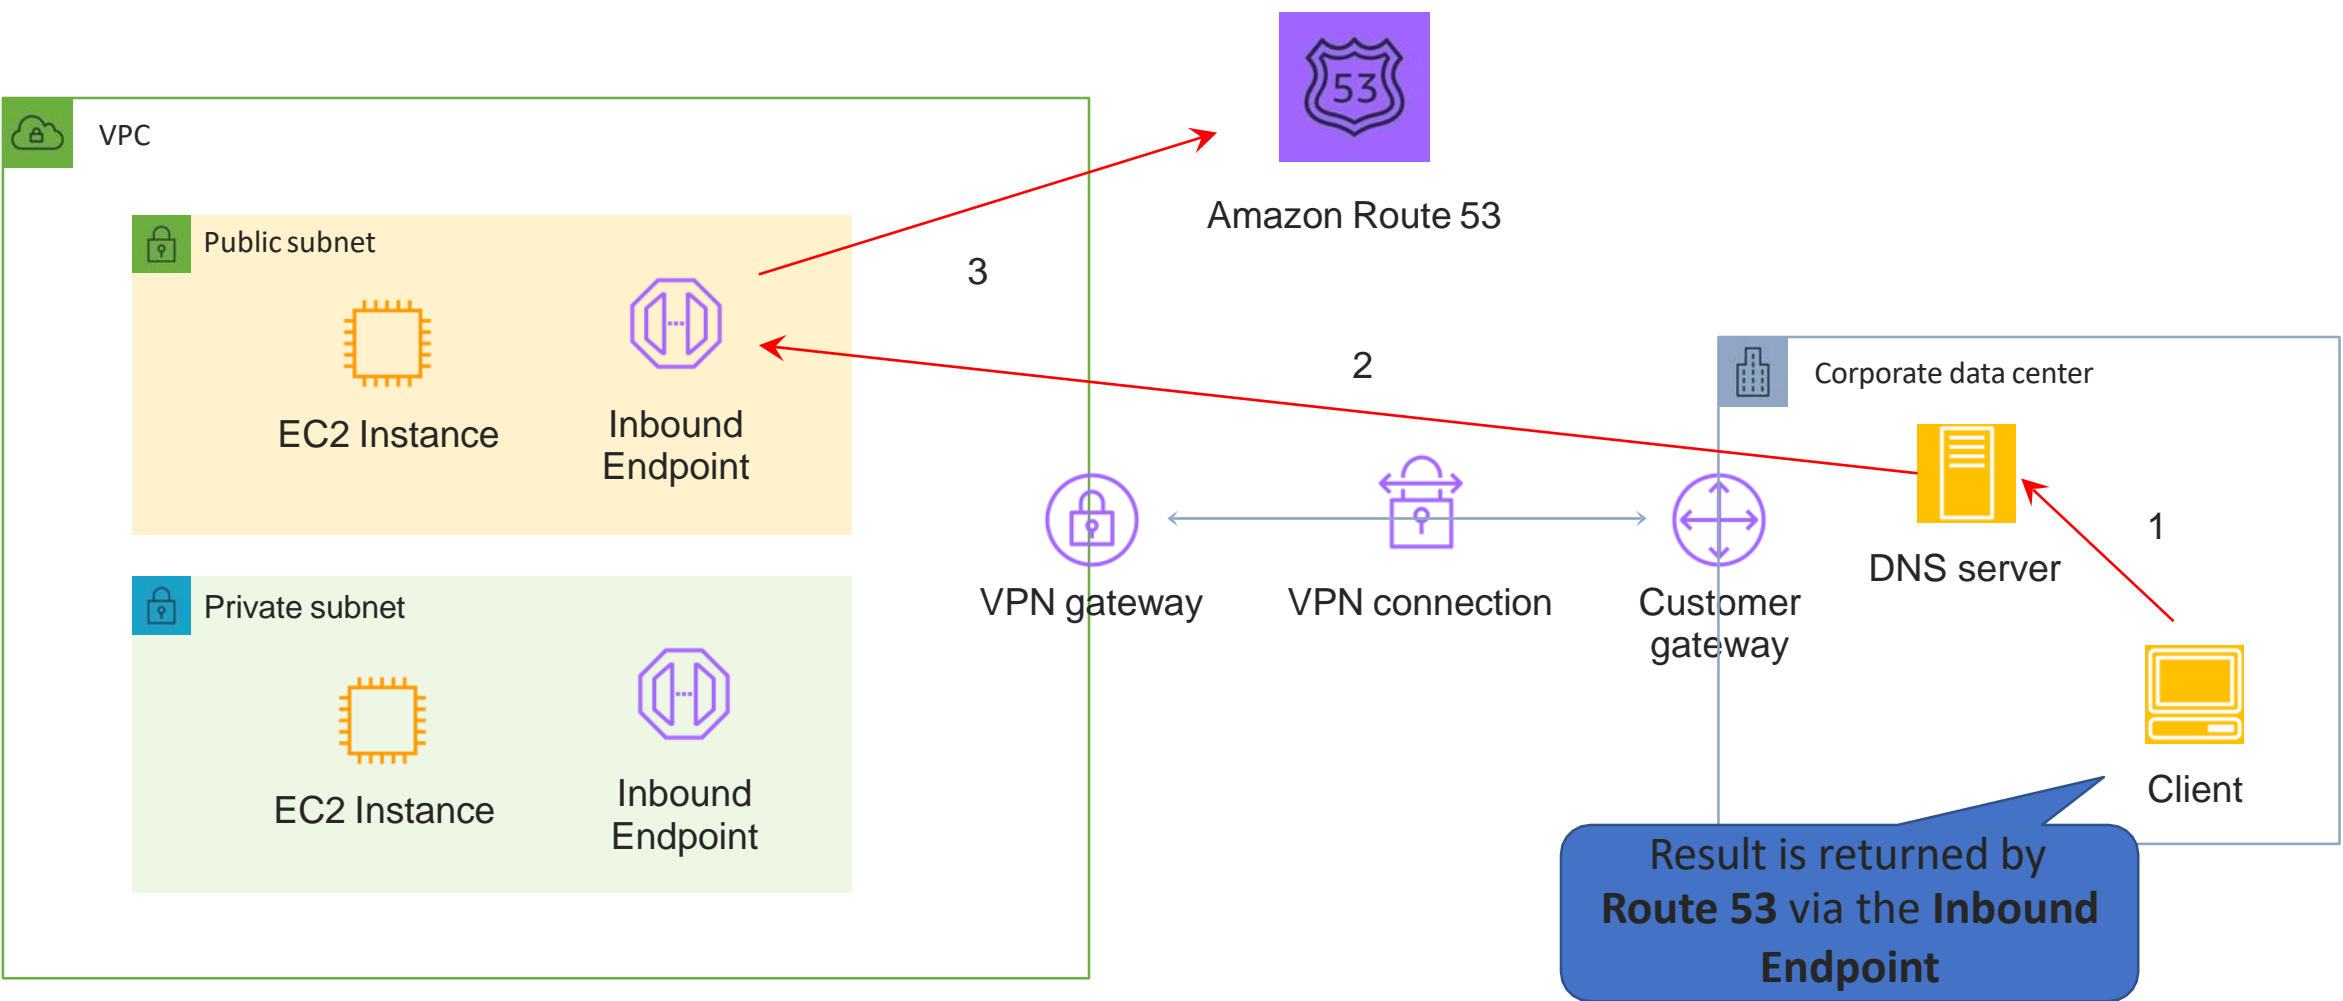
Amazon Route 53 Resolver



Route 53 Resolver – Outbound Endpoints



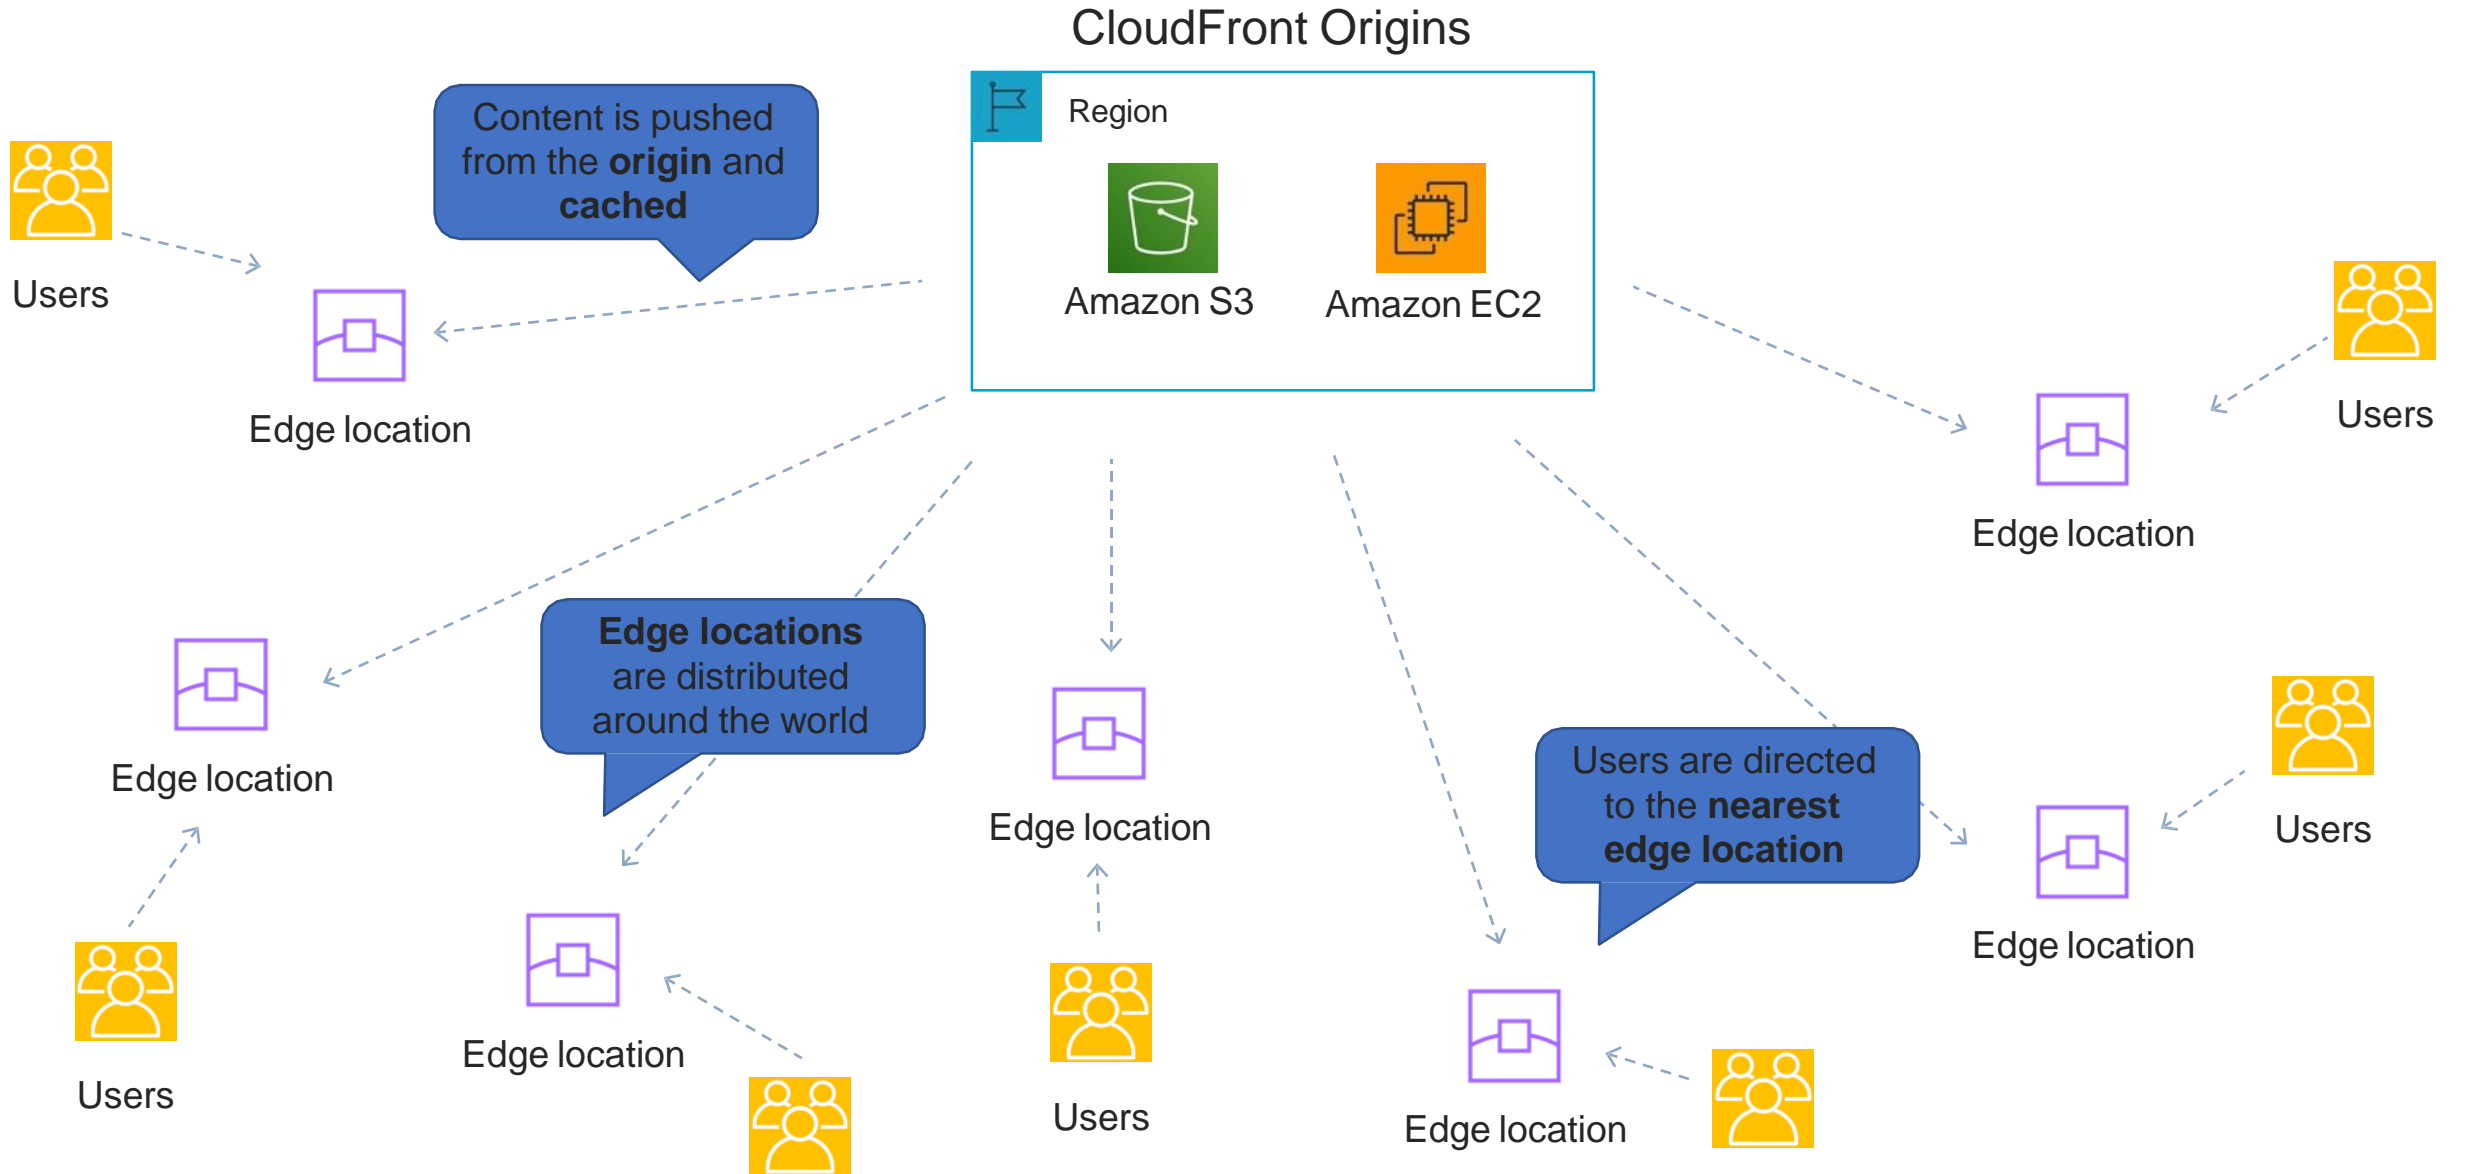
Route 53 Resolver – Inbound Endpoints



Amazon CloudFront Origins and Distributions



Amazon CloudFront

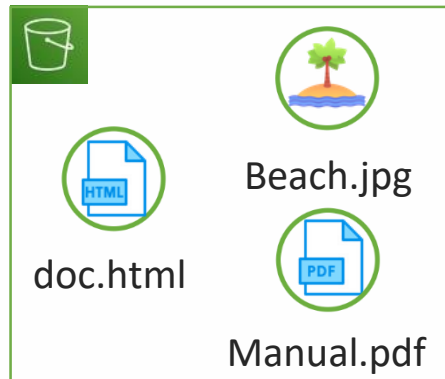


CloudFront Origins and Distributions

CloudFront Distribution

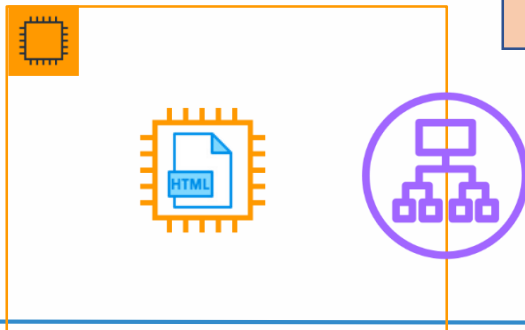
Name: **d1schtd9zdwrn1.cloudfront.net**

S3 Origin



S3 static websites
can also be origins

Custom Origin



Behaviors

Path Pattern
Viewer Protocol
Policy Cache
Policy
Origin Request
Policy

RTMP distributions were discontinued
so only web distributions are currently
available

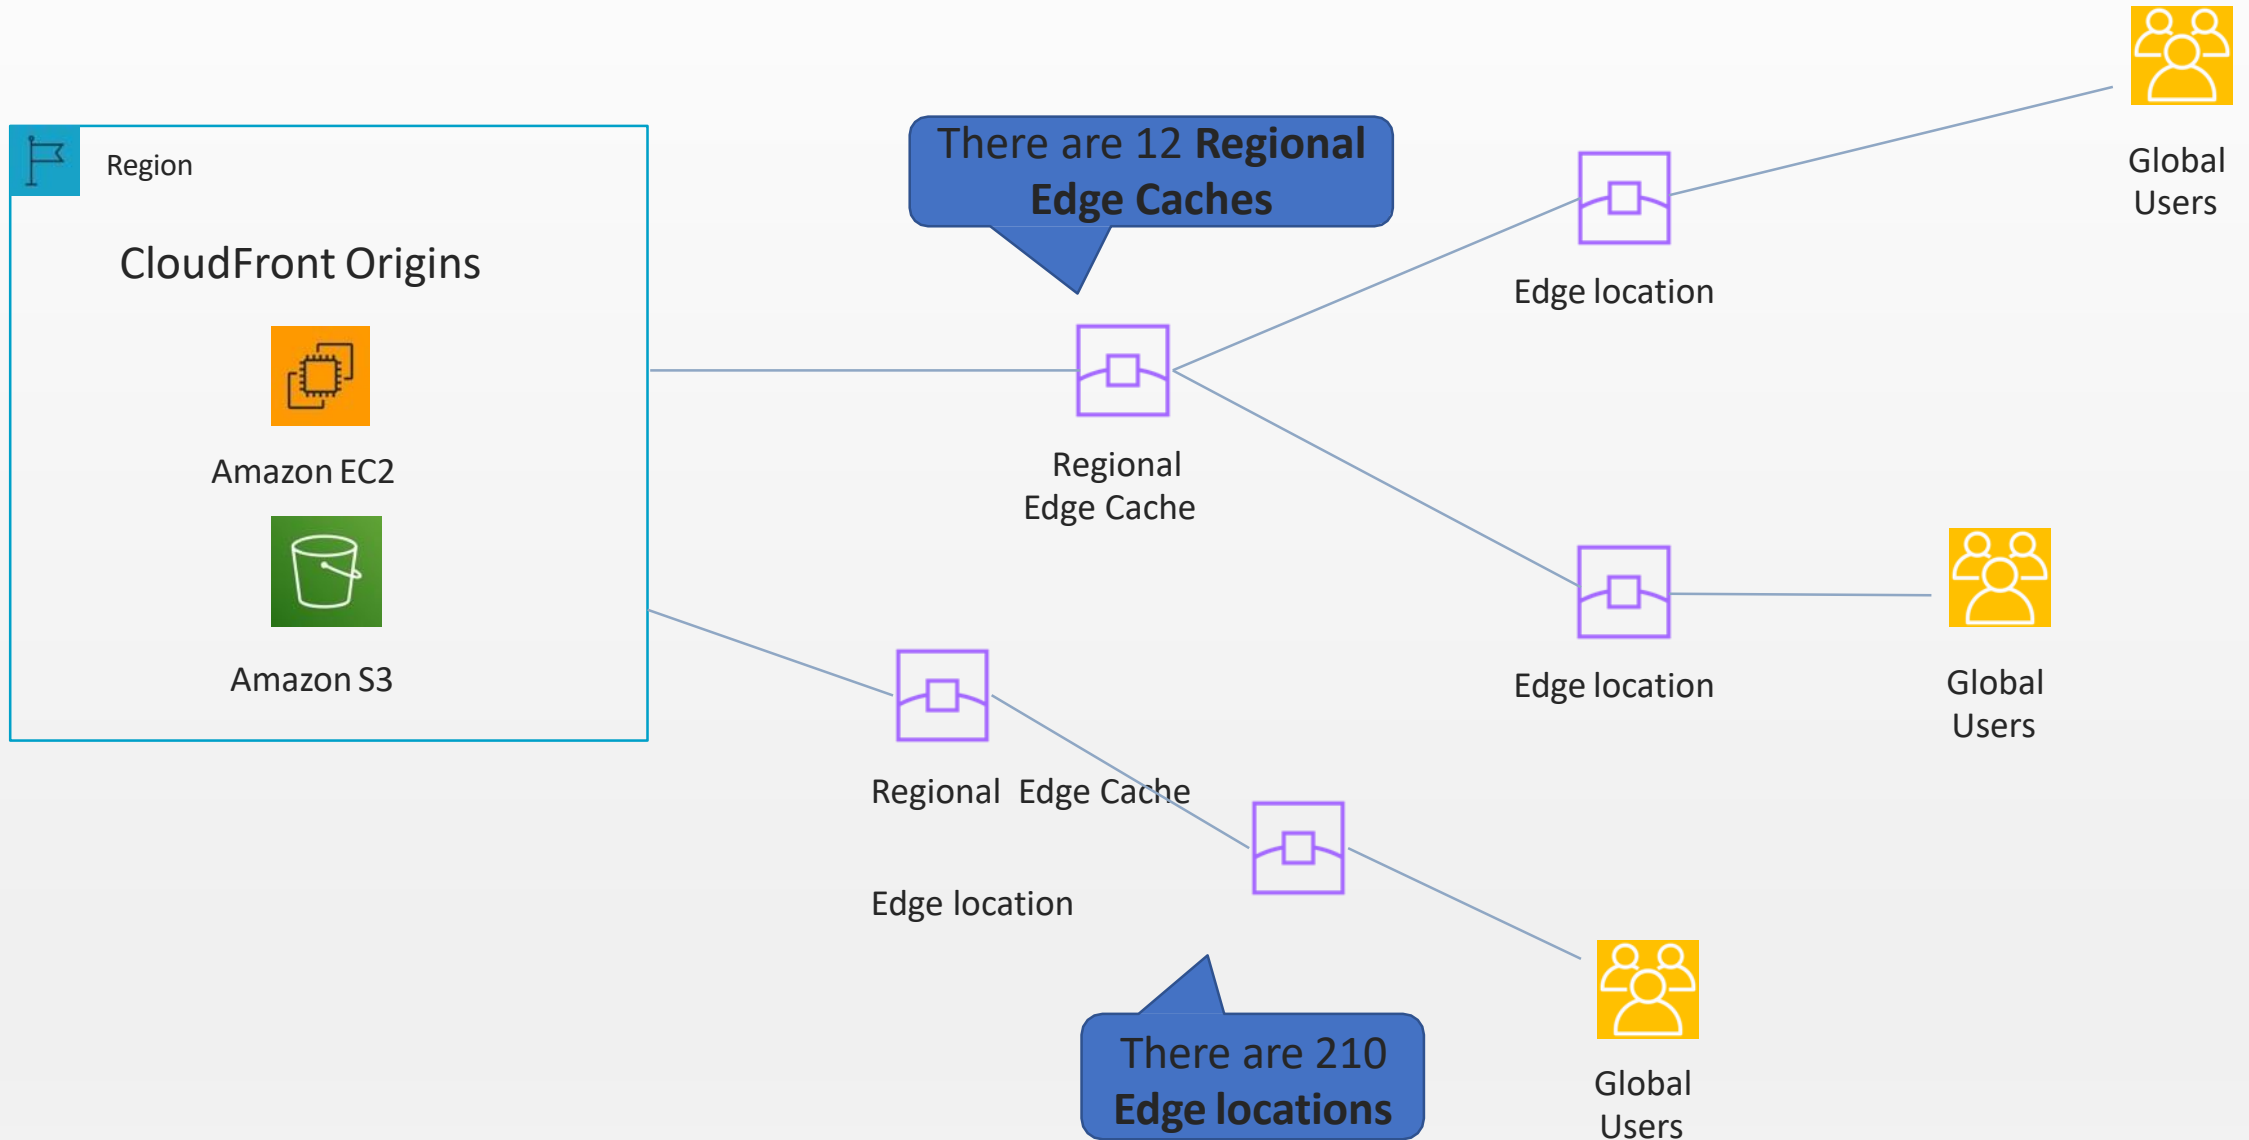
CloudFront Web Distribution:

- Speed up distribution of static and dynamic content, for example, .html, .css, .php, and graphics files.
- Distribute media files using HTTP or HTTPS.
- Add, update, or delete objects, and submit data from web forms.
- Use live streaming to stream an event in real time.

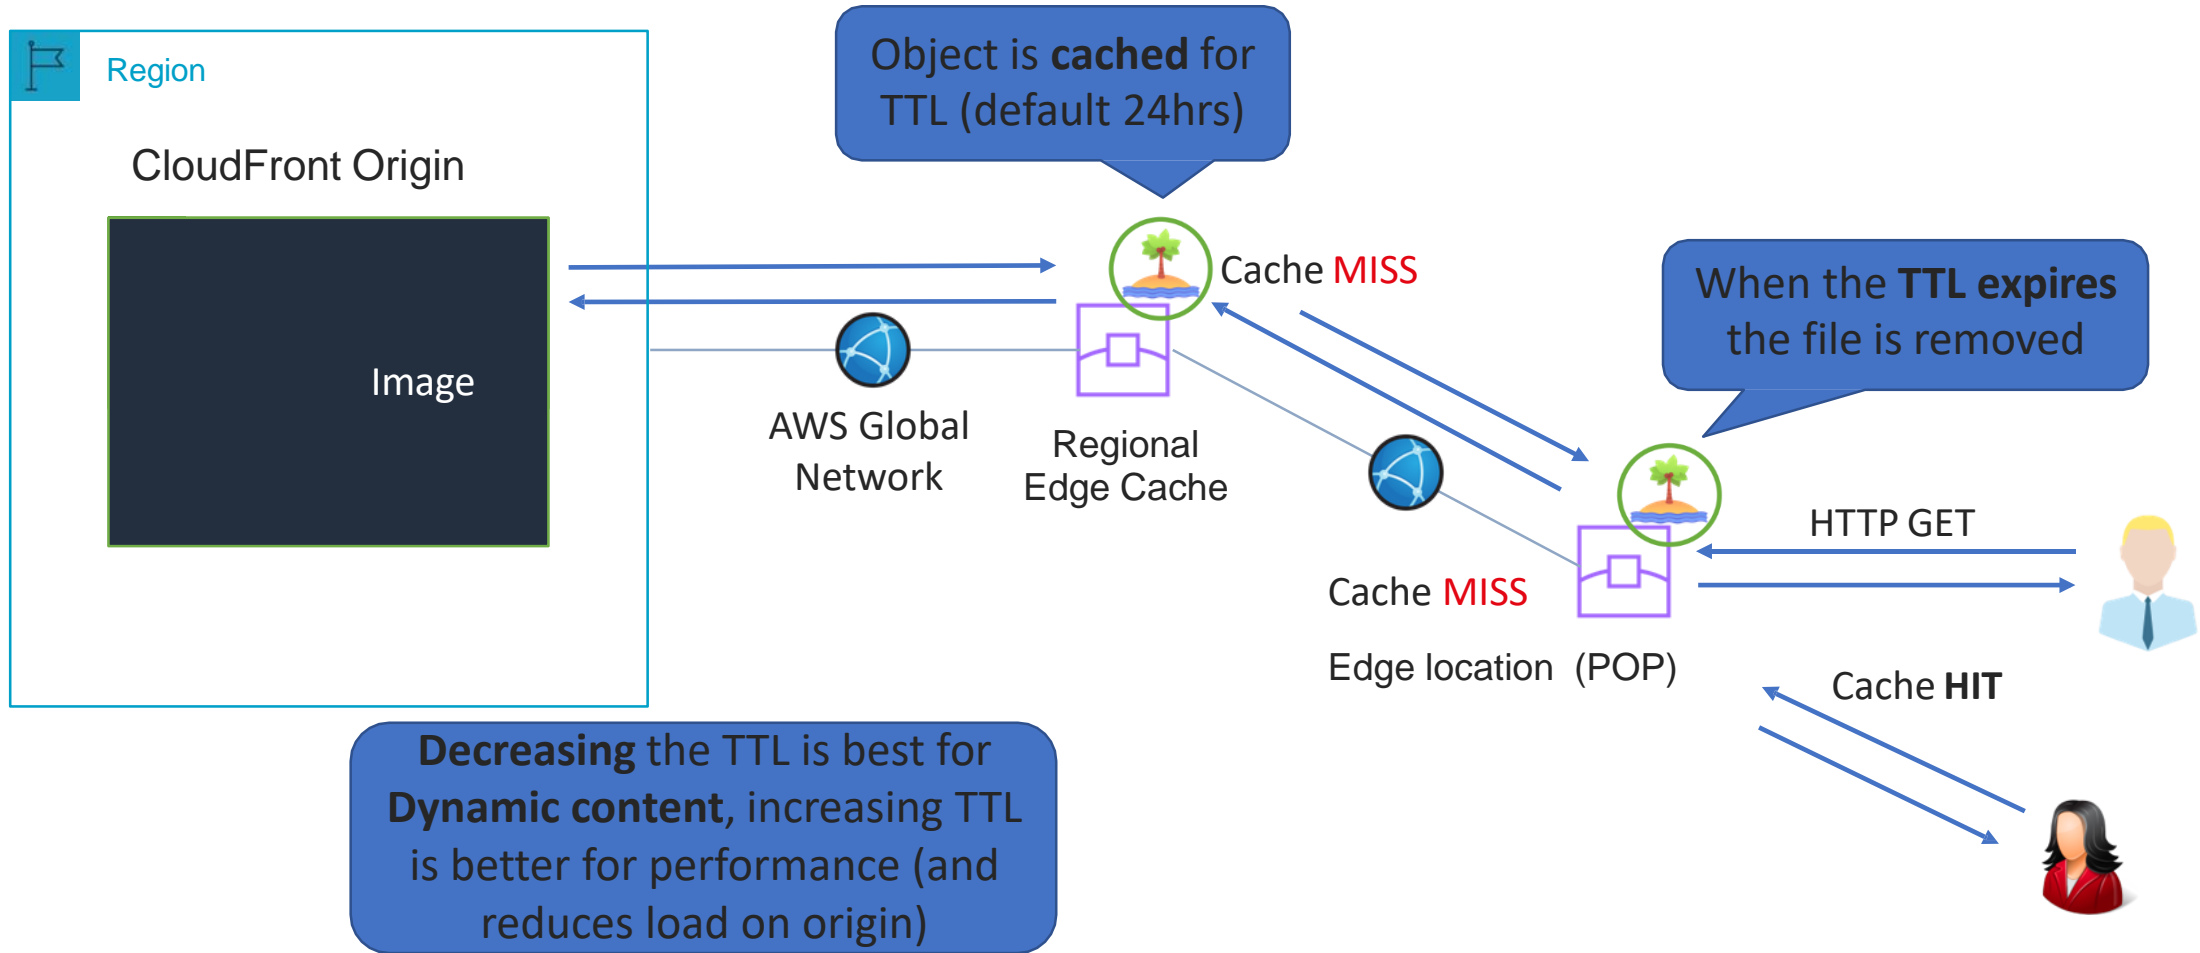
Amazon CloudFront Caching and Behaviors



Amazon CloudFront Caching



Amazon CloudFront Caching

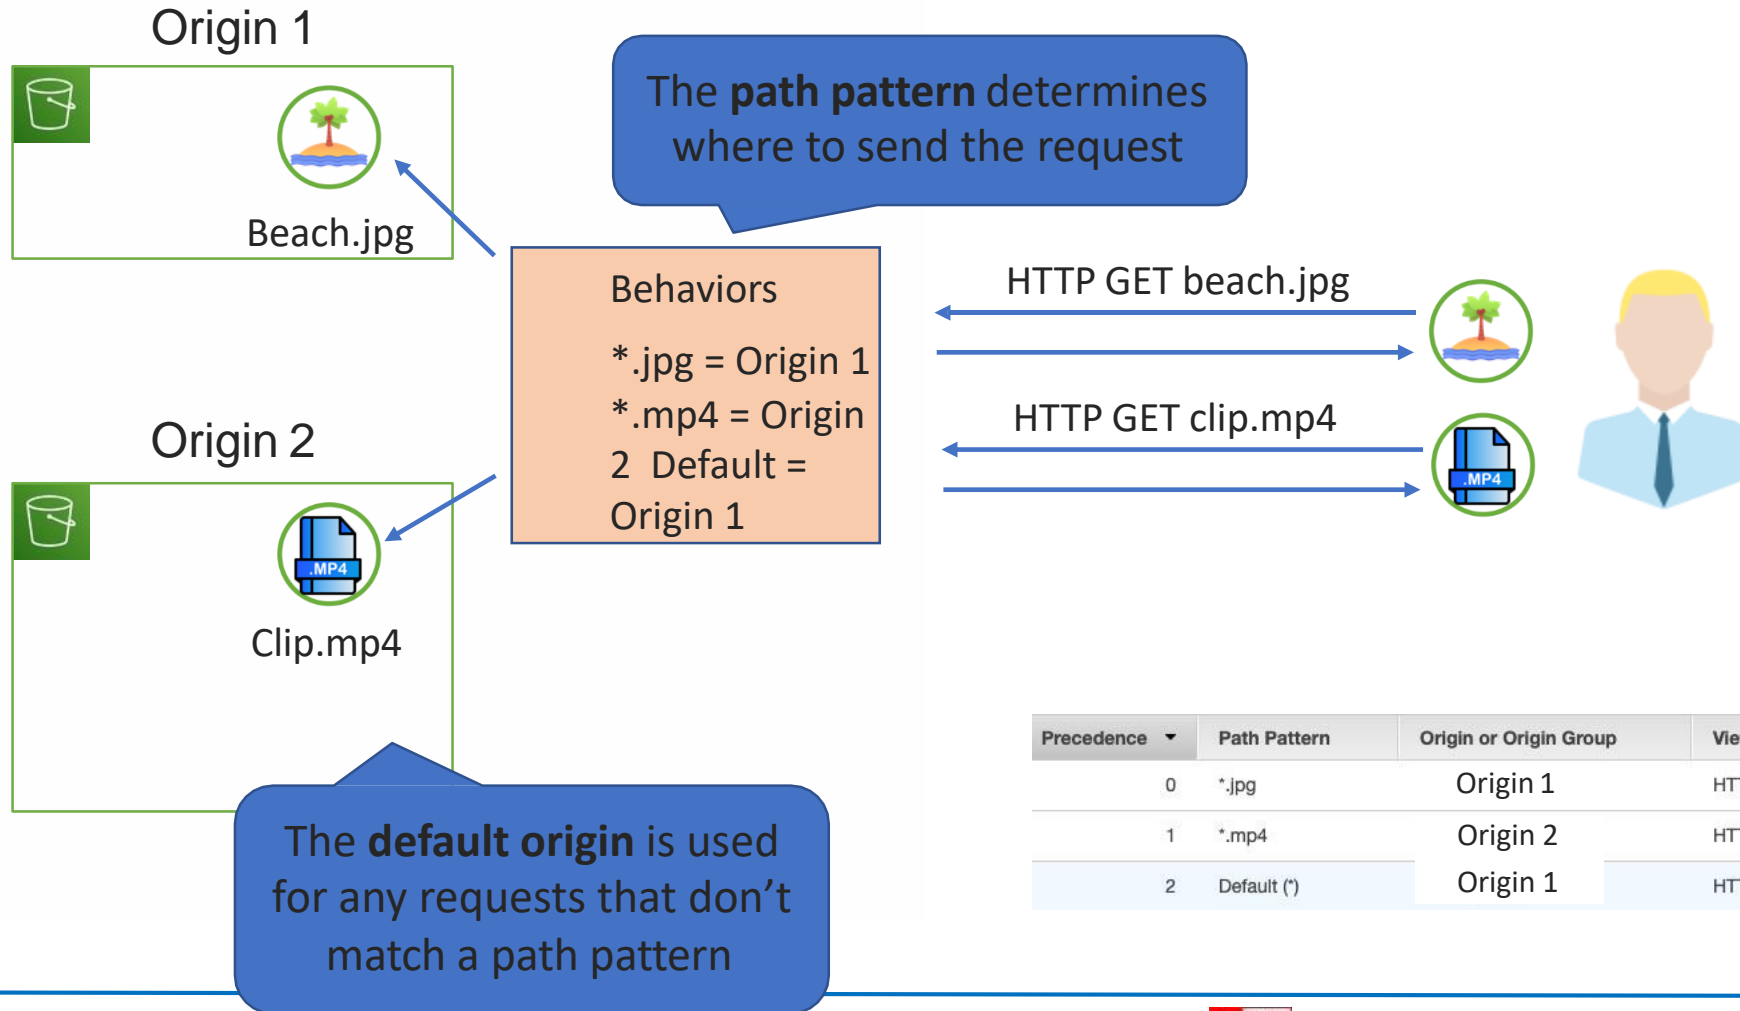


Amazon CloudFront Caching

- You can define a maximum **Time To Live (TTL)** and a default TTL
- TTL is defined at the **behavior** level
- This can be used to define different TTLs for different file types (e.g. png vs jpg)
- After expiration, CloudFront checks the origin for any new requests (check the file is the latest version)
- Headers can be used to control the cache:
 - `Cache-Control max-age= (seconds)` - specify how long before CloudFront gets the object again from the origin server
 - `Expires` - specify an expiration date and time

CloudFront Path Patterns

CloudFront Distribution



Precedence	Path Pattern	Origin or Origin Group	Viewer Protocol Policy	Cache Policy Name
0	*.jpg	Origin 1	HTTP and HTTPS	Managed-CachingOptimized
1	*.mp4	Origin 2	HTTP and HTTPS	Managed-CachingOptimized
2	Default (*)	Origin 1	HTTP and HTTPS	Managed-CachingOptimized

Caching Based on Request Headers

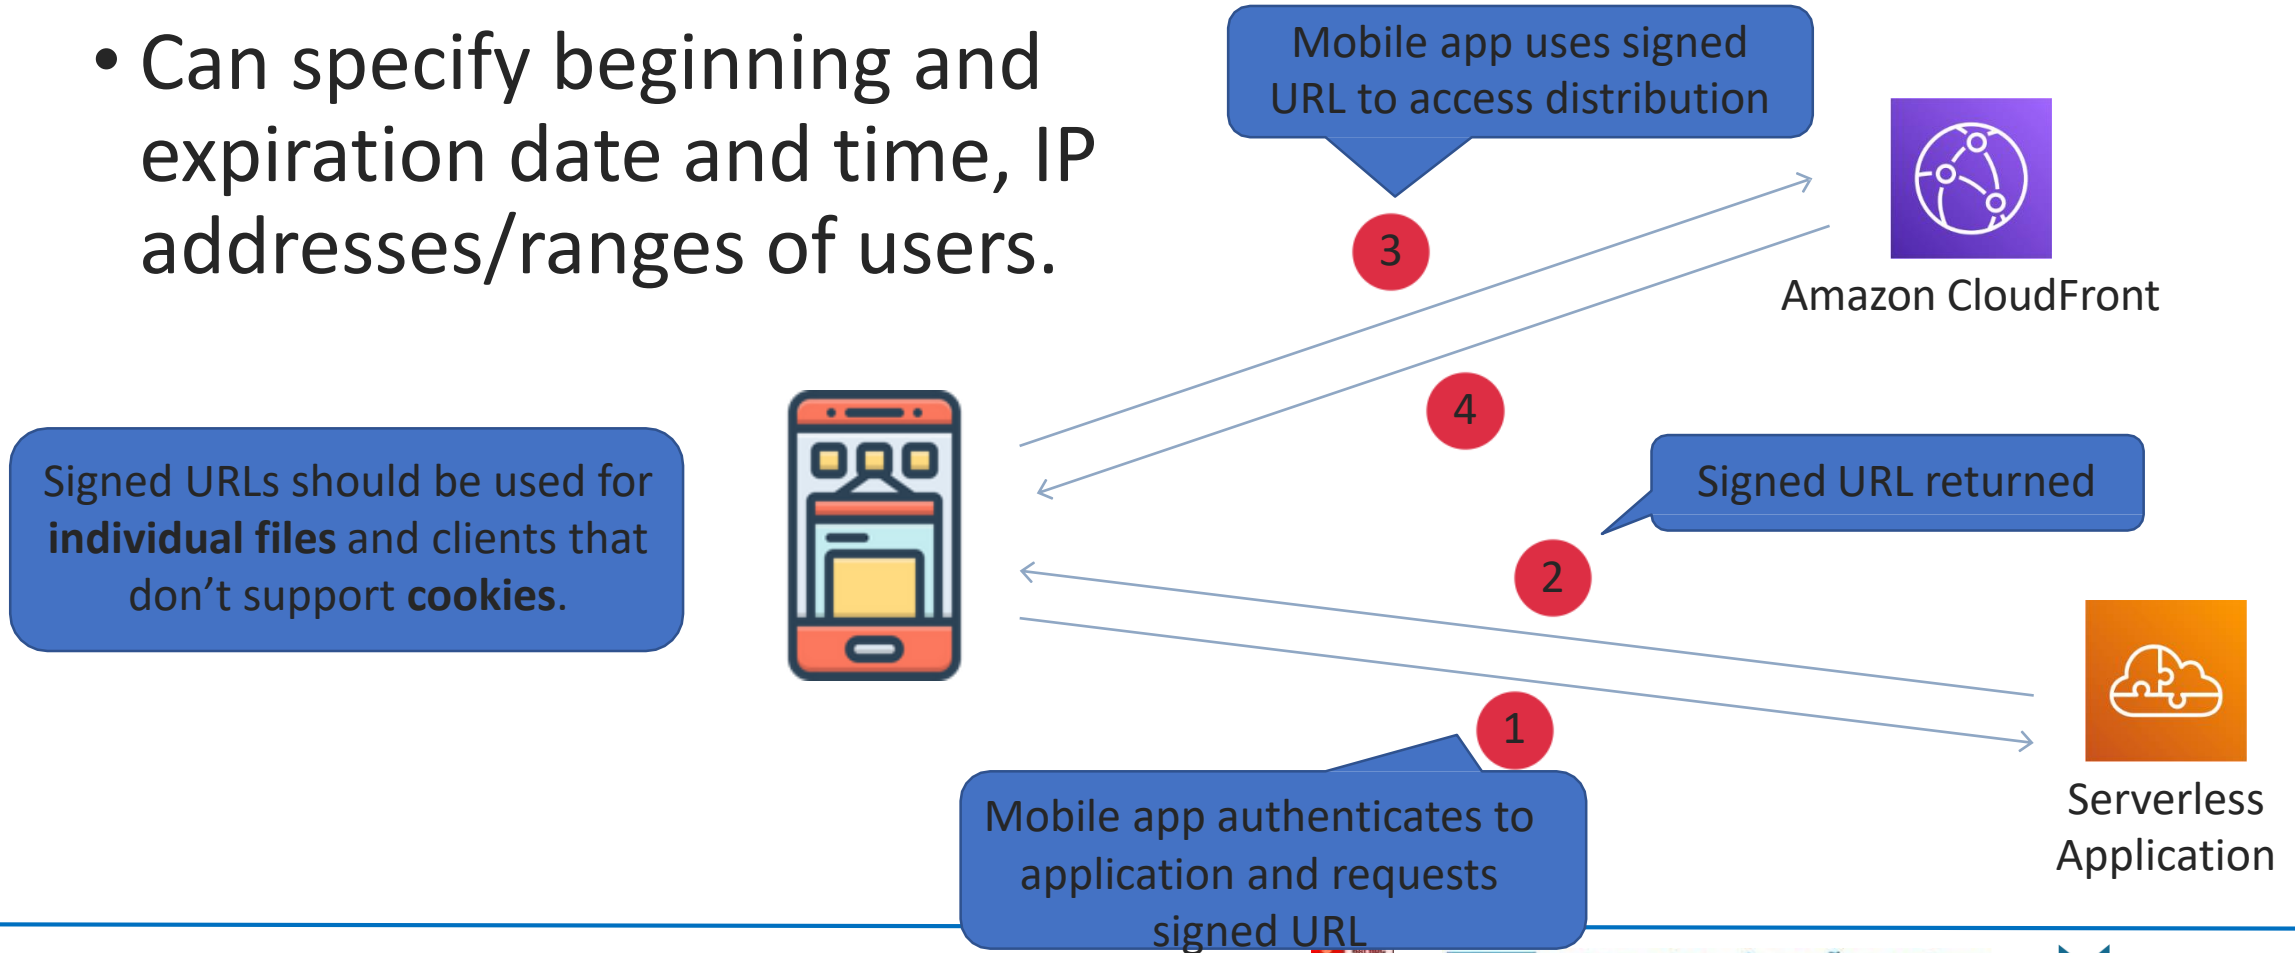
- You can configure CloudFront to forward **headers** in the **viewer request** to the origin
- CloudFront can then cache multiple versions of an object based on the values in one or more request headers
- Controlled in a behavior to do one of the following:
 - Forward all headers to your origin (objects are not cached)
 - Forward a whitelist of headers that you specify
 - Forward only the default headers (doesn't cache objects based on values in request headers)

CloudFront Signed URLs and OAI



CloudFront Signed URLs

- Signed URLs provide more control over access to content.
- Can specify beginning and expiration date and time, IP addresses/ranges of users.



CloudFront Signed Cookies

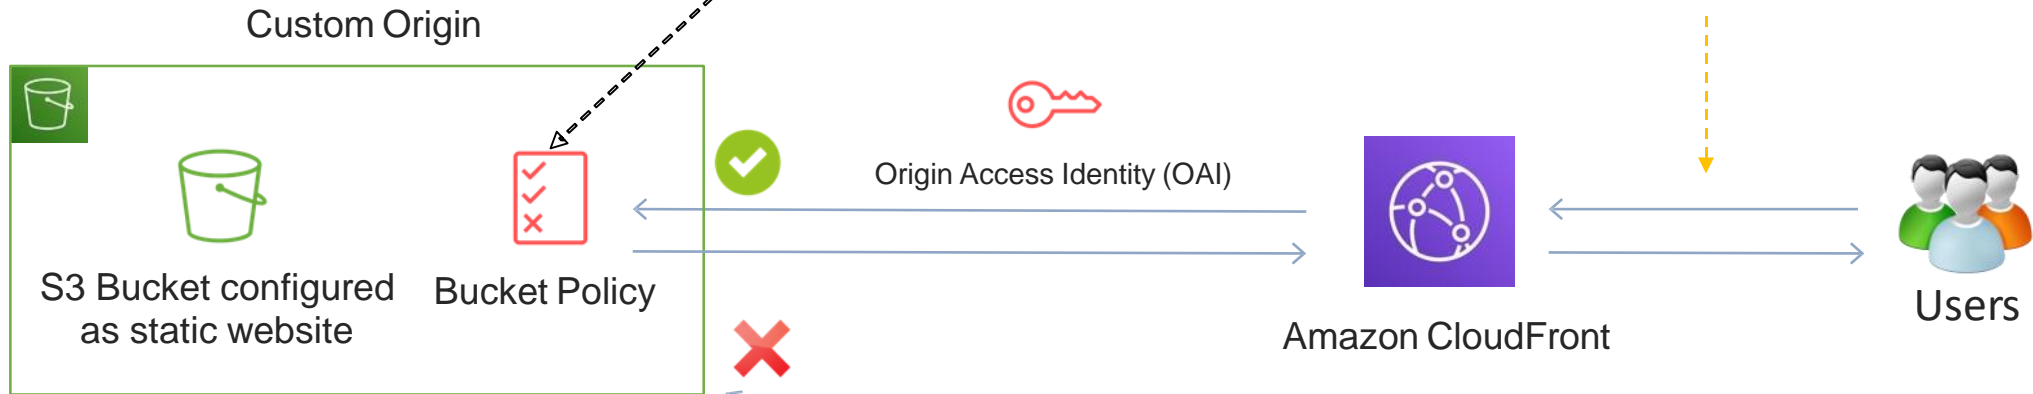
- Similar to Signed URLs
- Use signed cookies when you don't want to change URLs
- Can also be used when you want to provide access to **multiple restricted files** (Signed URLs are for individual files)

CloudFront Origin Access Identity (OAI)

```
{
  "Version": "2008-10-17",
  "Id": "PolicyForCloudFrontPrivateContent",
  "Statement": [
    {
      "Sid": "1",
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::cloudfront:user/CloudFront Origin Access Identity E11A2JL2H306JJ"
      },
      "Action": "s3:GetObject",
      "Resource": "arn:aws:s3:::mybucket/*"
    }
  ]
}
```

Policy restricts access to the OAI

HTTP GET
<https://d1schtd9zdwrn1.cloudfront.net>



GET <https://mybucket.s3.amazonaws.com/beach.jpg>



VIỆN ĐIỆN TỬ - VIỄN THÔNG
School of Electronics and Telecommunications

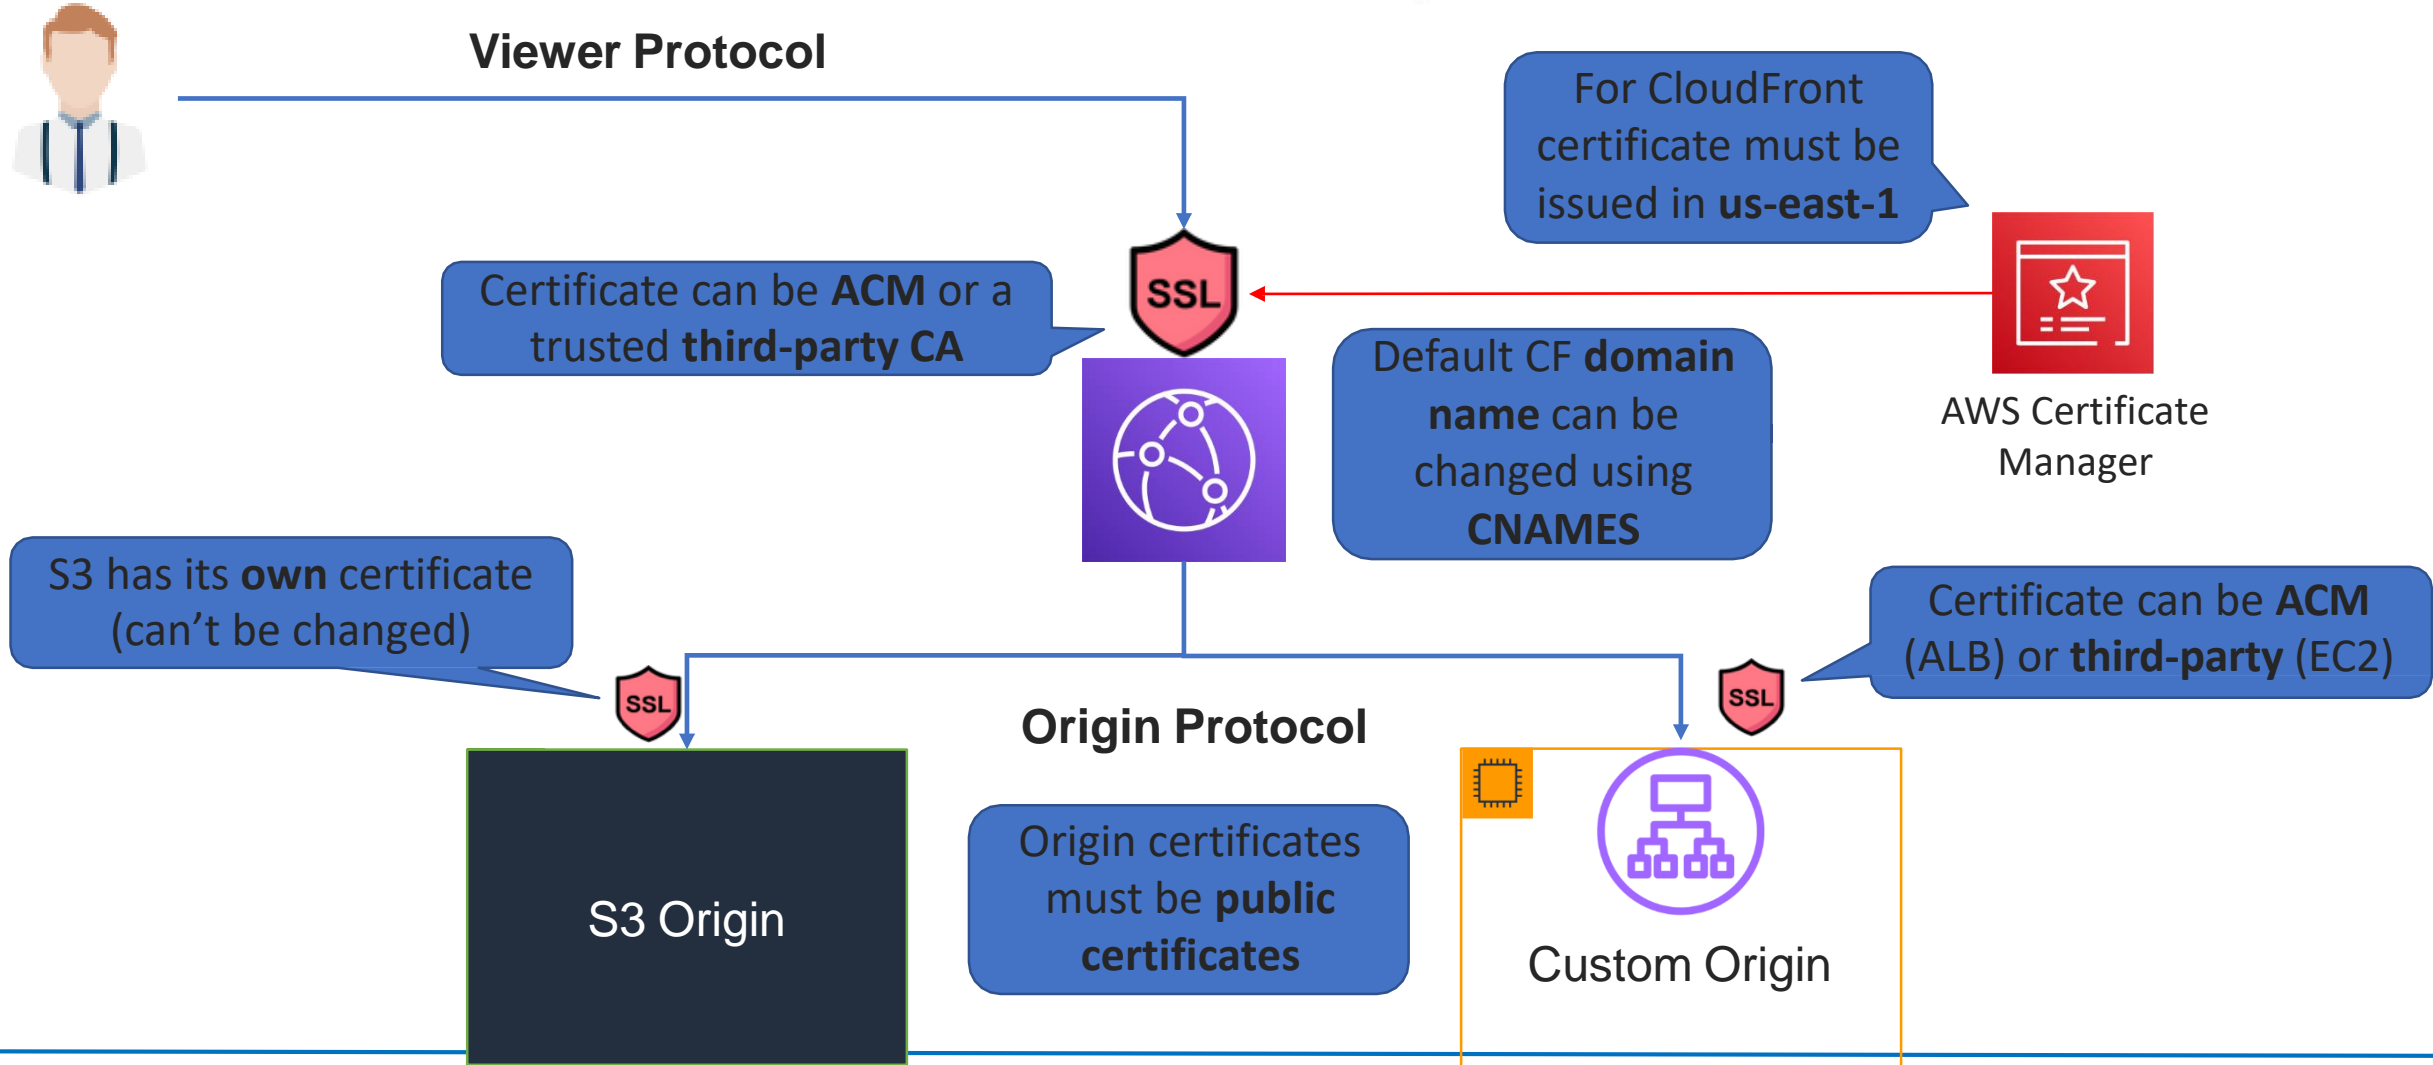


CloudFront SSL/TLS and SNI

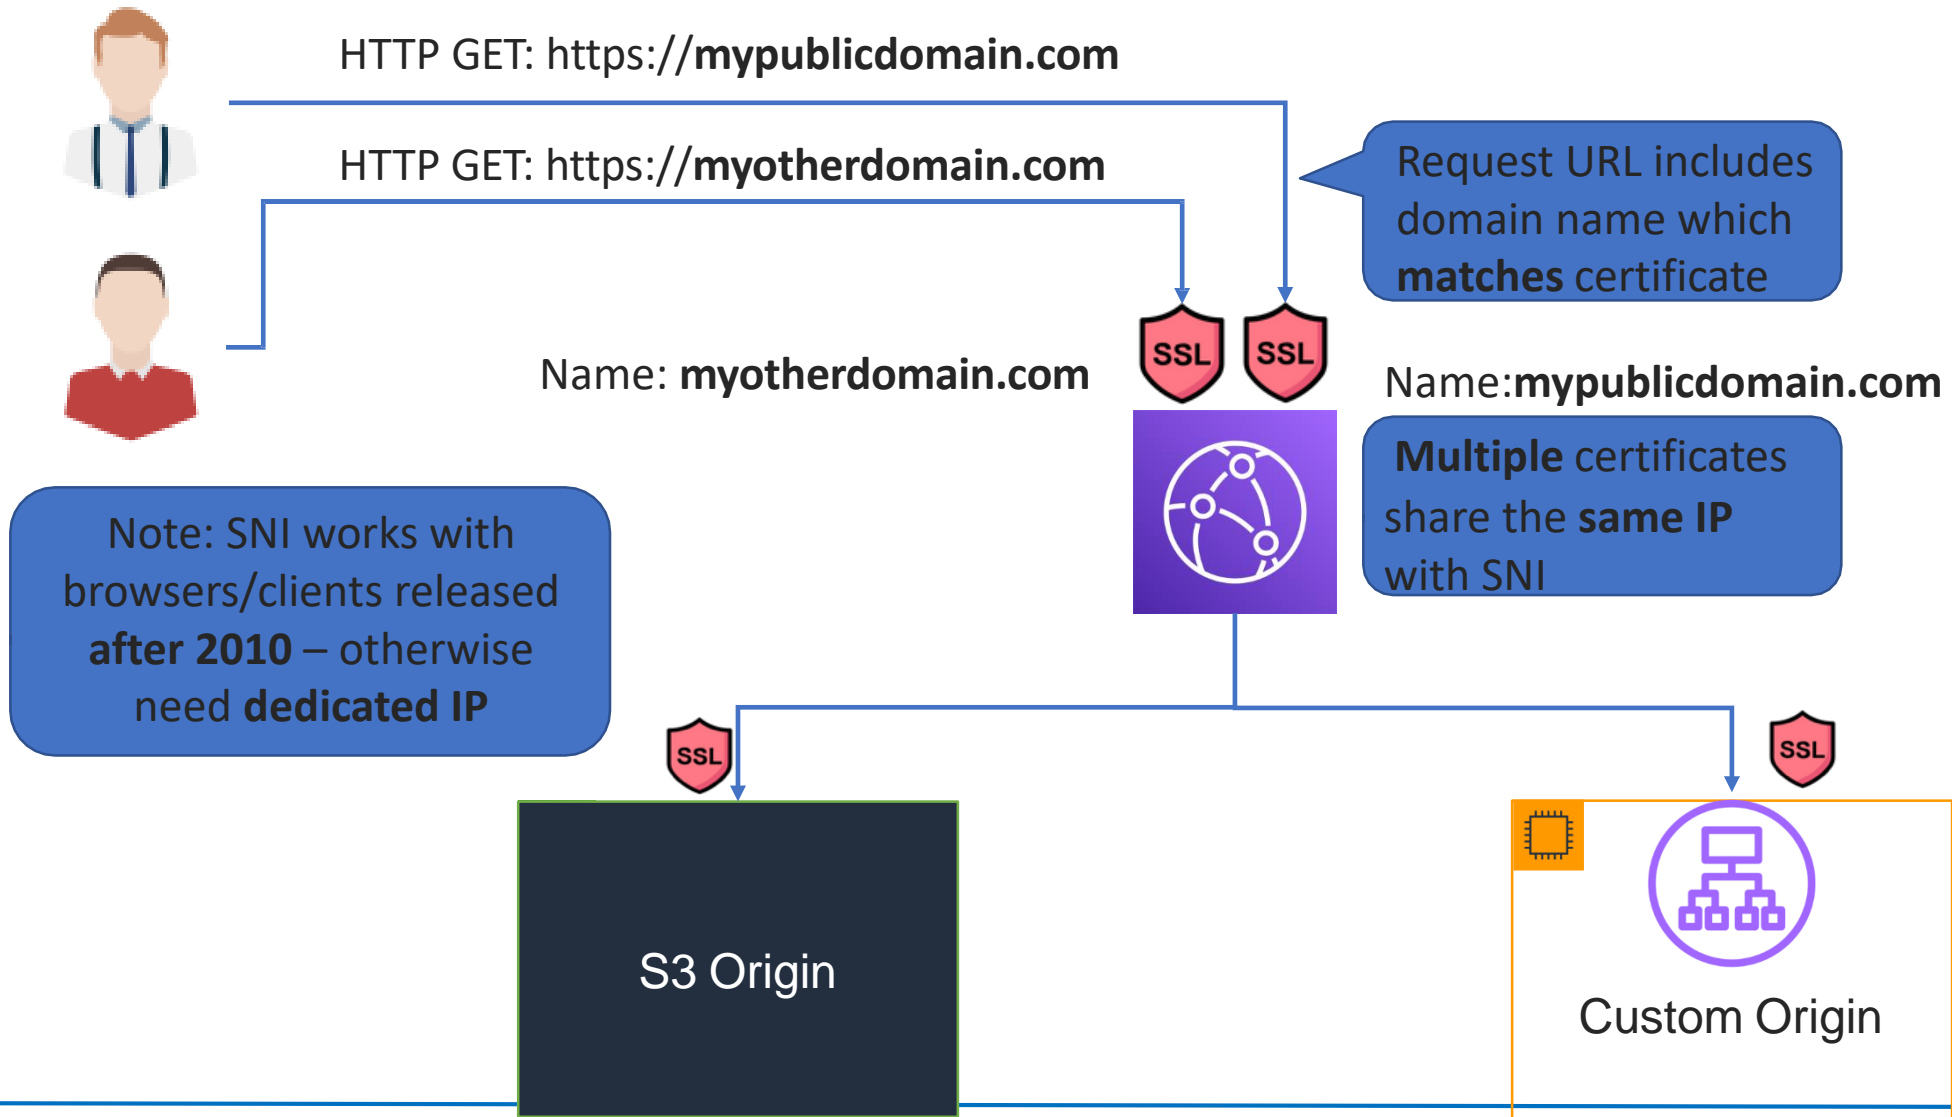


CloudFront SSL/TLS

- Viewer Protocol Policy
- HTTP and HTTPS
 - Redirect HTTP to HTTPS
 - HTTPS Only



CloudFront Server Name Indication (SNI)



Q & A



VIỆN ĐIỆN TỬ - VIỄN THÔNG
School of Electronics and Telecommunications



Module 6 Troubleshooting in AWS



Troubleshooting in AWS

- Application Domains
- Infrastructure Domains
- AWS Service Domains

AWS Service Domains

Check status AWS Services

<https://status.aws.amazon.com/>

Open Support Case

<https://docs.aws.amazon.com/awssupport/latest/user/case-management.html#creating-a-support-case>

Infrastructure Domains

<https://docs.aws.amazon.com/awssupport/latest/user/troubleshooting.html>



Troubleshooting EC2

InstanceLimitExceeded error

Cause

- You get the InstanceLimitExceeded error when you try to launch a new instance or restart a stopped instance
- You have reached the limit on the number of instances that you can launch in a Region

Solution

Request an instance limit increase on a per-region basis

Troubleshooting EC2

Instance terminates immediately

Your instance goes from the pending state to the terminated state

Cause

- An EBS snapshot is corrupted
- You've exceeded your EBS volume limits
- The root EBS volume is encrypted and you do not have permissions to access the KMS key for decryption

Solution: Get the termination reason then take one of corresponding actions

Open the **Amazon EC2** console

In the navigation pane, choose **Instances**, and select the instance.

On the first tab, find the reason next to **State transition reason**

Troubleshooting EC2

Connecting to your instance have problems

Causes

- Users, IP address
- Network (ACL, SG...)
- Keypairs

Solution:

Verify information about User, IP

Check ACL, SG to verify permission and networks

Check keypairs with at least permission (ReadOnly)

Troubleshooting EC2

System status checks

Instance status checks

<input type="checkbox"/>	Name	Instance ID	Instance state	Instance type	Status check	Alarm status	Avail
<input checked="" type="checkbox"/>	-	i-0c0186a12aab3741d	Running	t2.large	1/2 checks ...	No alarms +	eu-w
<input type="checkbox"/>	-	i-0138edcaf722db475	Running	m4.large	2/2 checks ...	No alarms +	eu-w
<input type="checkbox"/>	-	i-02c65b735153975ec	Running	t3.medium	2/2 checks ...	No alarms +	eu-w

Instance: i-0c0186a12aab3741d

Details | Security | Networking | Storage | **Status checks** | Monitoring | Tags

Status checks [Info](#)

Status checks detect problems that may impair i-0c0186a12aab3741d from running your applications.

System status checks

✔ System reachability check passed

Instance status checks

✘ Instance reachability check failed

Check failure at

2020/12/16 17:30 GMT+2 (about 1 month)

Need assistance?

If your instance is unreachable for more than 20 minutes, the **Open support case** button becomes available so that you can contact the Support Center.

[Open support case](#)

Visit the [Support Center](#) or post a question to the [Discussion Forums](#)



Troubleshooting EC2

System status checks

Cause

- Loss of network connectivity
- Loss of system power
- Software issues on the physical host
- Hardware issues on the physical host that impact network reachability

Solution:

Stop and Start EC2 instances

Terminate EC2 instances

Troubleshooting EC2

Instance status checks (0/2, 1/2)

Causes

- Failed system status checks
- Incorrect networking or startup configuration
- Exhausted memory
- Corrupted file system
- Incompatible kernel

Solution:

Stop and Start EC2 instances

Terminate EC2 instances

Troubleshooting ELB

A registered target is not in service

Cause

- Network not allow traffic
- The ping path does not exist
- The connection times out

Troubleshooting ELB

Clients cannot connect to an internet-facing load balancer

Cause

- Your internet-facing load balancer is attached to a private subnet
- A security group or network ACL does not allow traffic

Troubleshooting ELB

The load balancer generates an HTTP error

- **HTTP 400: Bad request** The client sent a malformed request that does not meet the HTTP specification.
- **HTTP 401: Unauthorized** Problems with Authenticate
- **HTTP 403: Forbidden** AWS WAF web access control list (web ACL) to monitor requests to your Application Load Balancer and it blocked a request.

<https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-troubleshooting.html>

Q & A



VIỆN ĐIỆN TỬ - VIỄN THÔNG
School of Electronics and Telecommunications



Thank You!



VIỆN ĐIỆN TỬ - VIỄN THÔNG
School of Electronics and Telecommunications

