

AWS CERTIFIED

CLOUD

PRACTITIONER



Technology



AMBER ISRAELSEN

Developer | Technical Trainer

Course Outline

Course
Introduction

Cloud Concepts

Security and
Compliance

Cloud
Technology and
Services

Billing, Pricing
and Support

Preparing for
the Exam



aws Services [Alt+S] Ohio

AWS Management Console

AWS services

▼ Recently visited services

- EC2
- Support
- VPC
- AWS License Manager
- Cloud9
- S3
- Security Hub
- Artifact
- Config
- Secrets Manager
- RDS
- Key Management Service
- IAM
- Lambda
- Alexa for Business

► All services

Build a solution

Get started with simple wizards and automated workflows.

- Launch a virtual machine**
With EC2
2-3 minutes
- Build a web app**
With Elastic Beanstalk
6 minutes
- Build using virtual servers**
With Lightsail
1-2 minutes
- Register a domain**
With Route 53
3 minutes
- Connect an IoT device**
With AWS IoT
5 minutes
- Start migrating to AWS**
With AWS MGN
1-2 minutes

Stay connected to your AWS resources on-the-go

AWS Console Mobile App now supports four additional regions. Download the AWS Console Mobile App to your iOS or Android mobile device. [Learn more](#)

Explore AWS

- Build Apps Faster with GraphQL**
AWS AppSync uses GraphQL APIs to query data from multiple data sources in a single request. [Get started](#)
- Calling All Java and Python Developers**
Join the AWS BugBust challenge to bust one million bugs. [Learn more](#)
- AWS Cloud Training**
Comprehensive training that accelerates and broadens cloud adoption. [Learn more](#)
- AWS Certification**
Propel your career forward with AWS Certification. [Learn more](#)

Feedback English (US) © 2022, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences



The screenshot shows a web-based IDE interface. At the top, there's a browser address bar with the URL `n/cloud9/ide/43a7df44b13e4d65960e19f2973fff79` and an Incognito window. The IDE has a menu bar with 'Go', 'Run', 'Tools', 'Window', and 'Support'. Below the menu, there are buttons for 'Preview', 'Run', 'Share', and a settings gear. The main editor area shows a file named `sharingfiles.py` with the following Python code:

```
1 import boto3
2
3 url = boto3.client('s3').generate_presigned_url(
4     ClientMethod='get_object',
5     Params={'Bucket': 'BUCKET_NAME', 'Key': 'OBJECT_KEY'},
6     ExpiresIn=3600)
7
8 print('The URL is ' + url)
9
```

At the bottom right of the editor, it says '9:1 Python Spaces: 4'. Below the editor is a terminal window showing the output of a pip install command:

```
Requirement already satisfied: python-dateutil<3.0.0,>=2.1 in /usr/local/lib/python3.7/site-packages (from boto3<2.8.2,>=1.23.26->boto3) (2.8.2)
Requirement already satisfied: urllib3<1.27,>=1.25.4 in /usr/local/lib/python3.7/site-packages (from botocore<1.24.0,>=1.23.26->boto3) (1.26.7)
Requirement already satisfied: six>=1.5 in /usr/local/lib/python3.7/site-packages (from python-dateutil<3.0.0,>=2.1->botocore<1.24.0,>=1.23.26->boto3) (1.16.0)
Installing collected packages: botocore, s3transfer, boto3
Successfully installed boto3-1.20.26 botocore-1.23.26 s3transfer-0.5.0
TTTAdmin:~/environment $
```

SOFTWARE DEVELOPER KIT (SDK)



```
Command Prompt - aws help
aws
^^^

Description
*****

The AWS Command Line Interface is a unified tool to manage your AWS
services.

Synopsis
*****

aws [options] <command> <subcommand> [parameters]

Use *aws command help* for information on a specific command. Use *aws
help topics* to view a list of available help topics. The synopsis for
each command shows its parameters and their usage. Optional parameters
are shown in square brackets.

Options
*****

"--debug" (boolean)

Turn on debug logging.
-- More --
```

COMMAND LINE INTERFACE (CLI)



aws Services [Alt+S] Ohio

AWS CloudShell Actions

us-east-2

```
Preparing your terminal...  
[cloudshell-user@ip-10-0-125-111 ~]$ Try these commands to get started:  
aws help or aws <command> help or aws <command> --cli-auto-prompt  
[cloudshell-user@ip-10-0-125-111 ~]$
```

AWS CLOUDSHELL

Feedback English (US) © 2022, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences



DEMO

Installing the AWS CLI on Windows



DEMO

Installing the AWS CLI on Linux



DEMO

Installing the AWS CLI on Mac



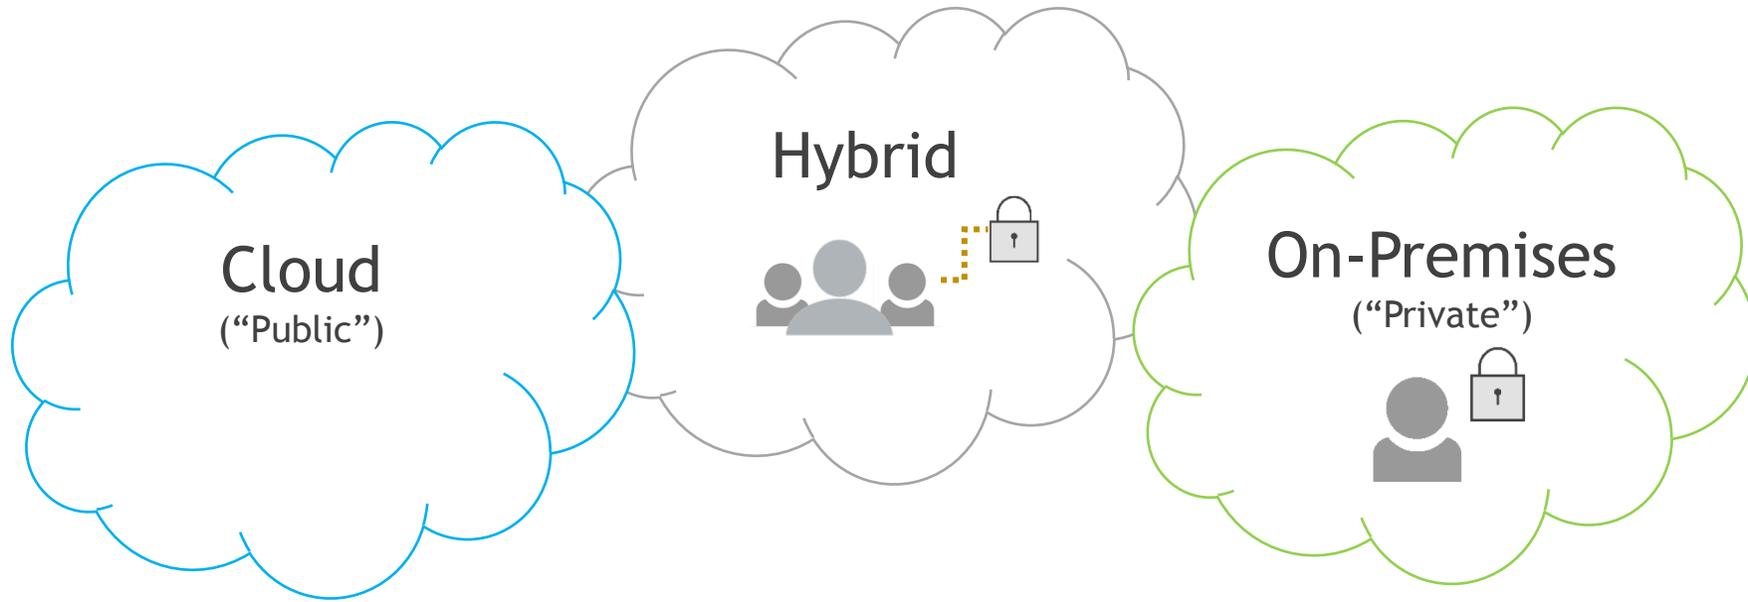
DEMO

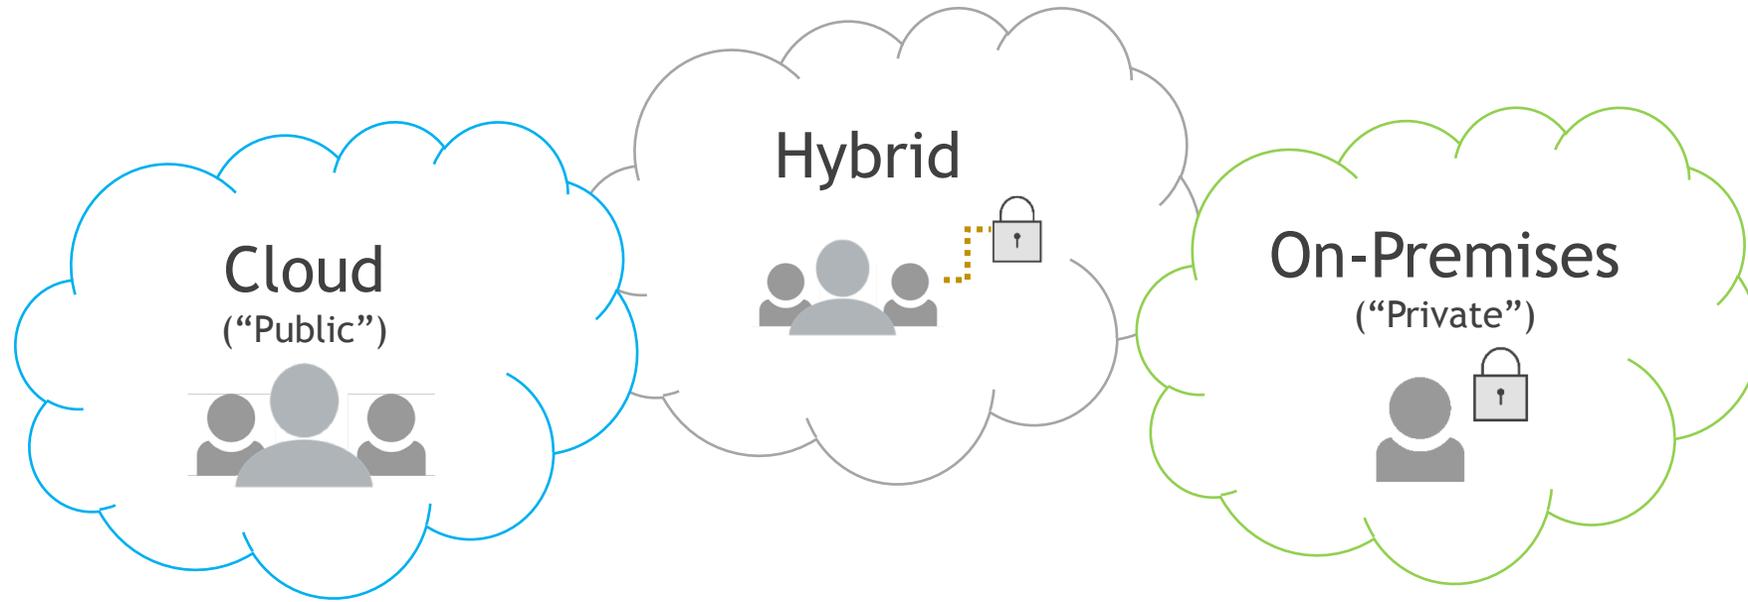
Configuring the AWS CLI on Windows, Mac, Linux



DEMO

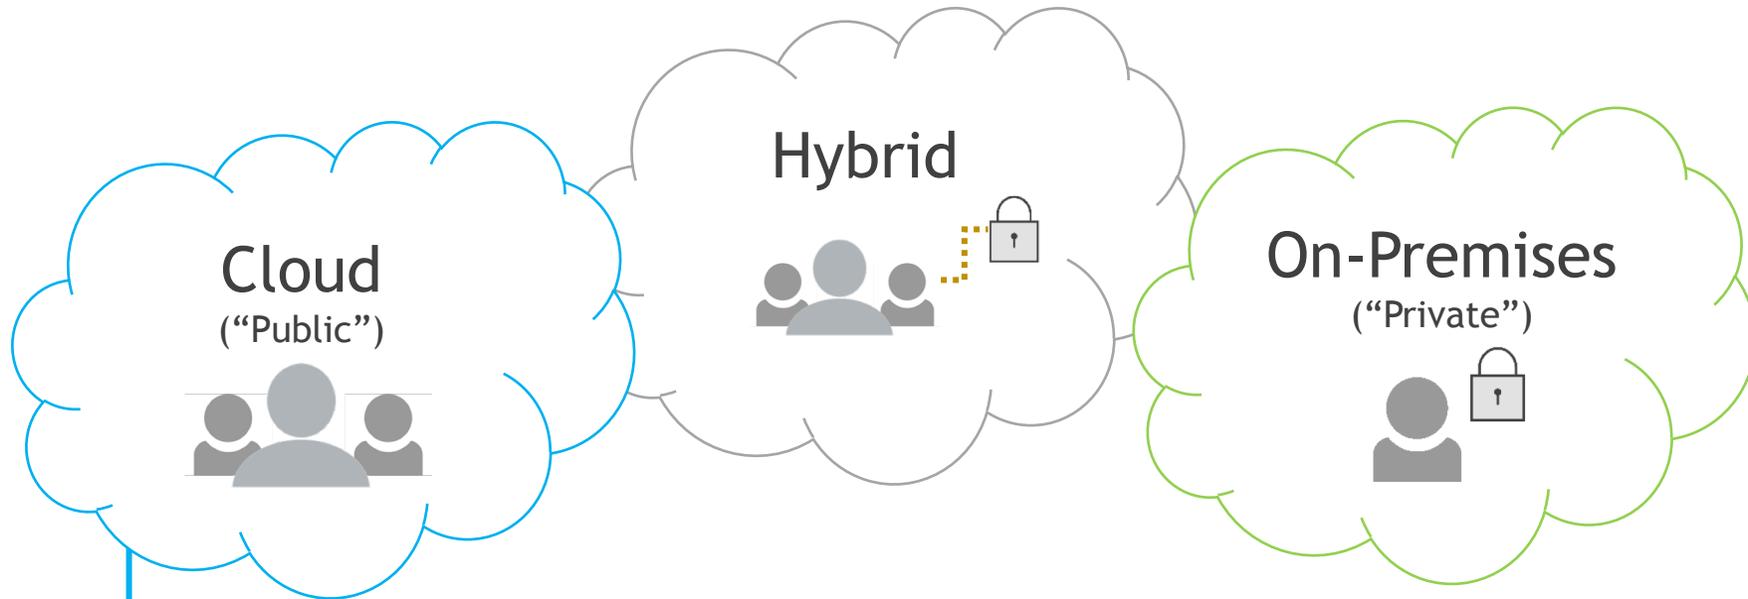
Working with AWS CloudShell





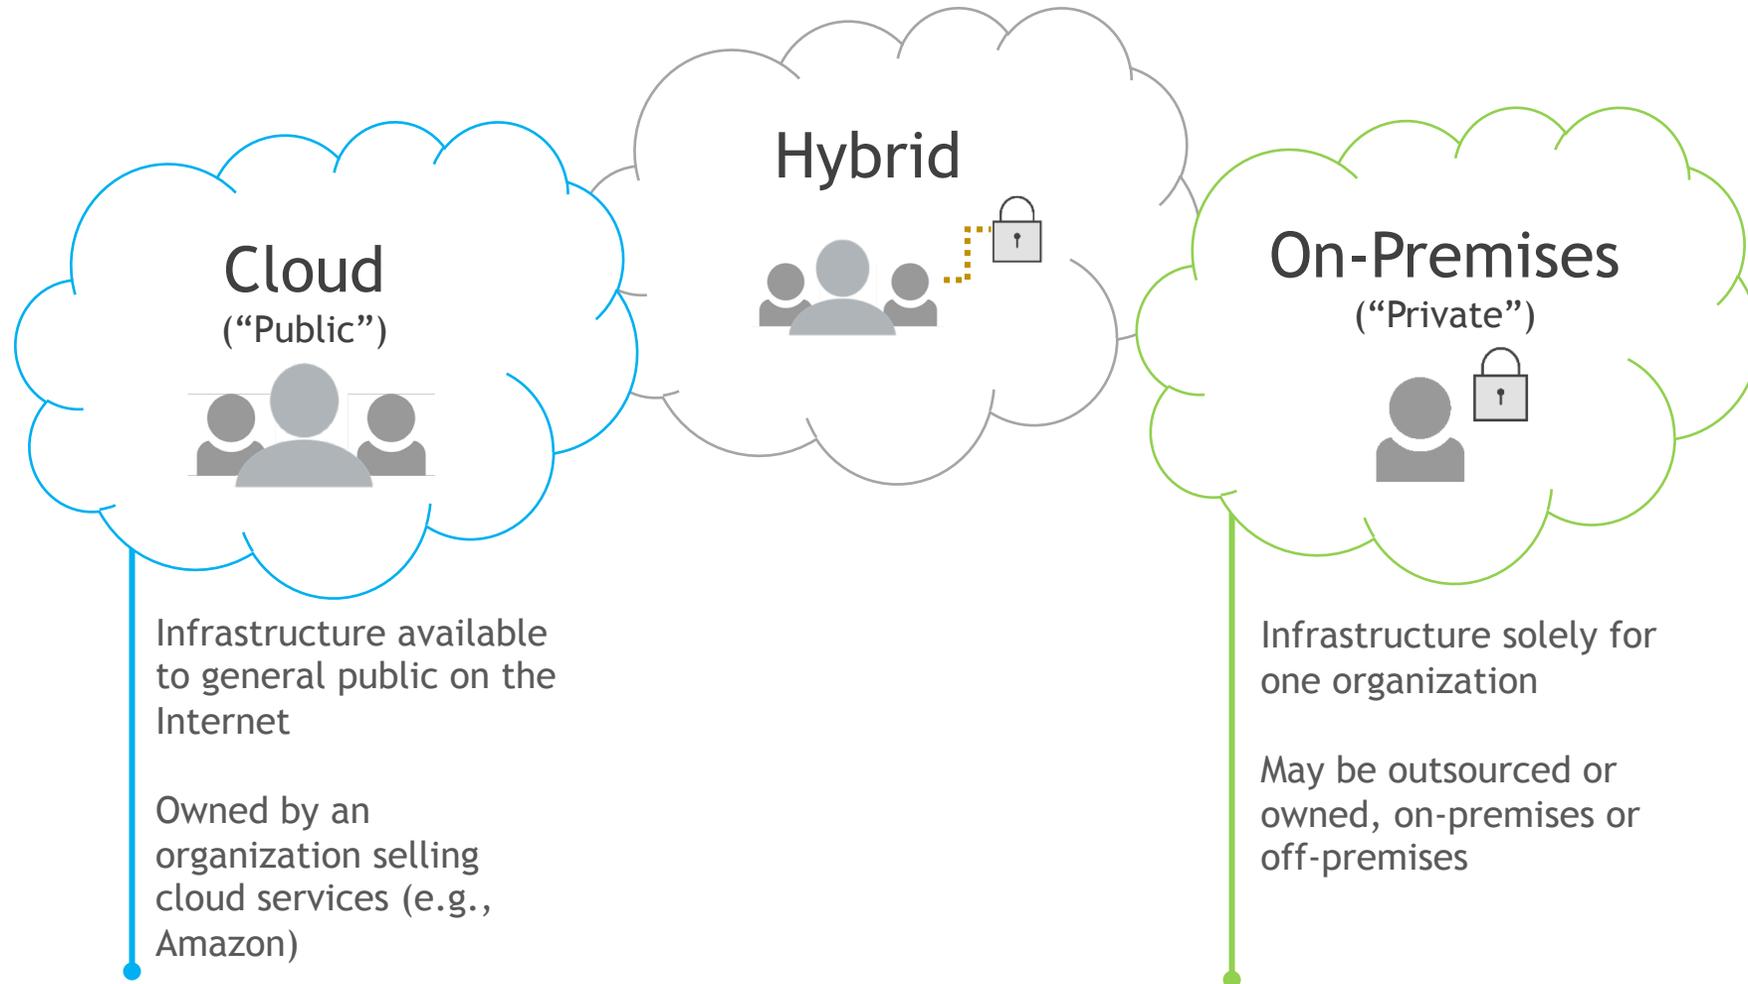
- Cost savings
- Scalability
- Flexibility

- Privacy
- Total Control
- Dedicated Resources



Infrastructure available to general public on the Internet

Owned by an organization selling cloud services (e.g., Amazon)



Cloud ("Public")



Infrastructure available to general public on the Internet

Owned by an organization selling cloud services (e.g., Amazon)

Hybrid

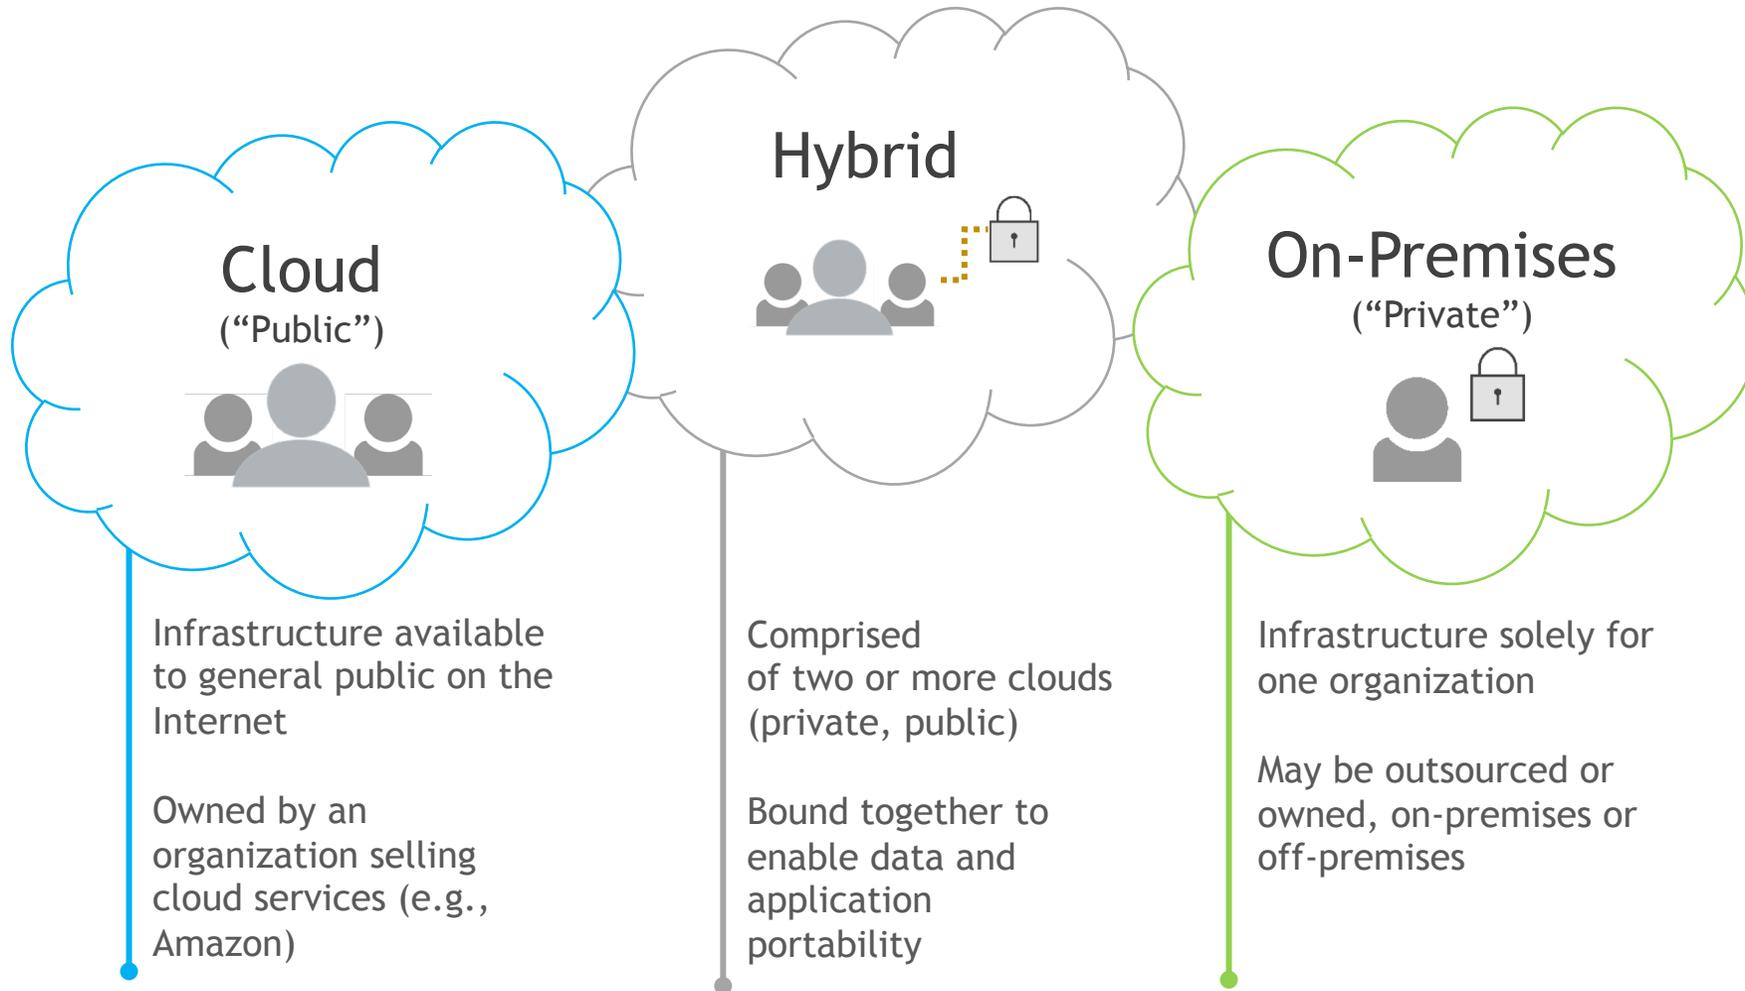


On-Premises ("Private")



Infrastructure solely for one organization

May be outsourced or owned, on-premises or off-premises



Cloud ("Public")



Infrastructure available to general public on the Internet

Owned by an organization selling cloud services (e.g., Amazon)

Hybrid



Comprised of two or more clouds (private, public)

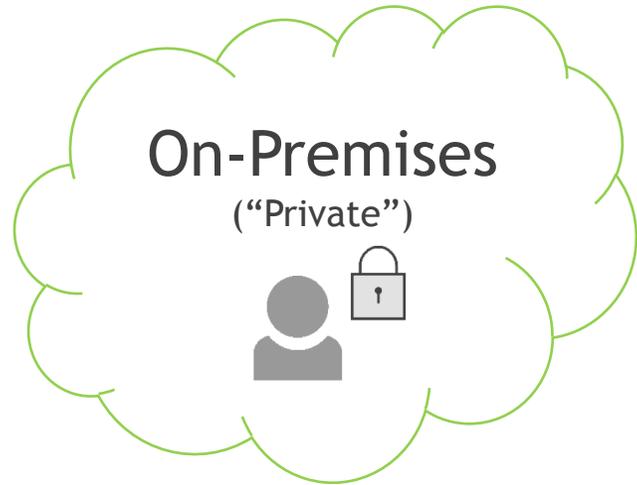
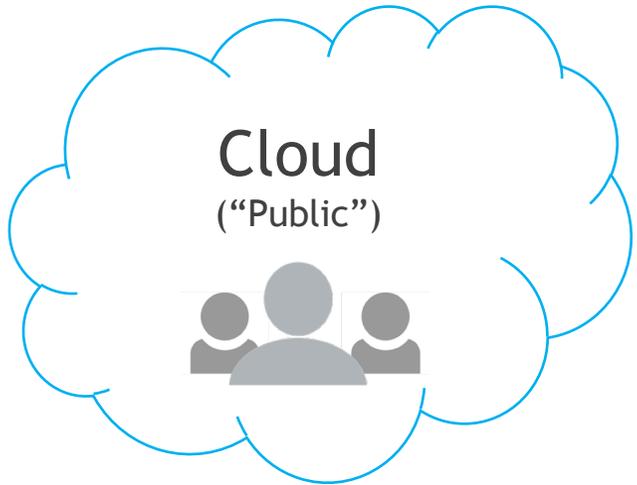
Bound together to enable data and application portability

On-Premises ("Private")



Infrastructure solely for one organization

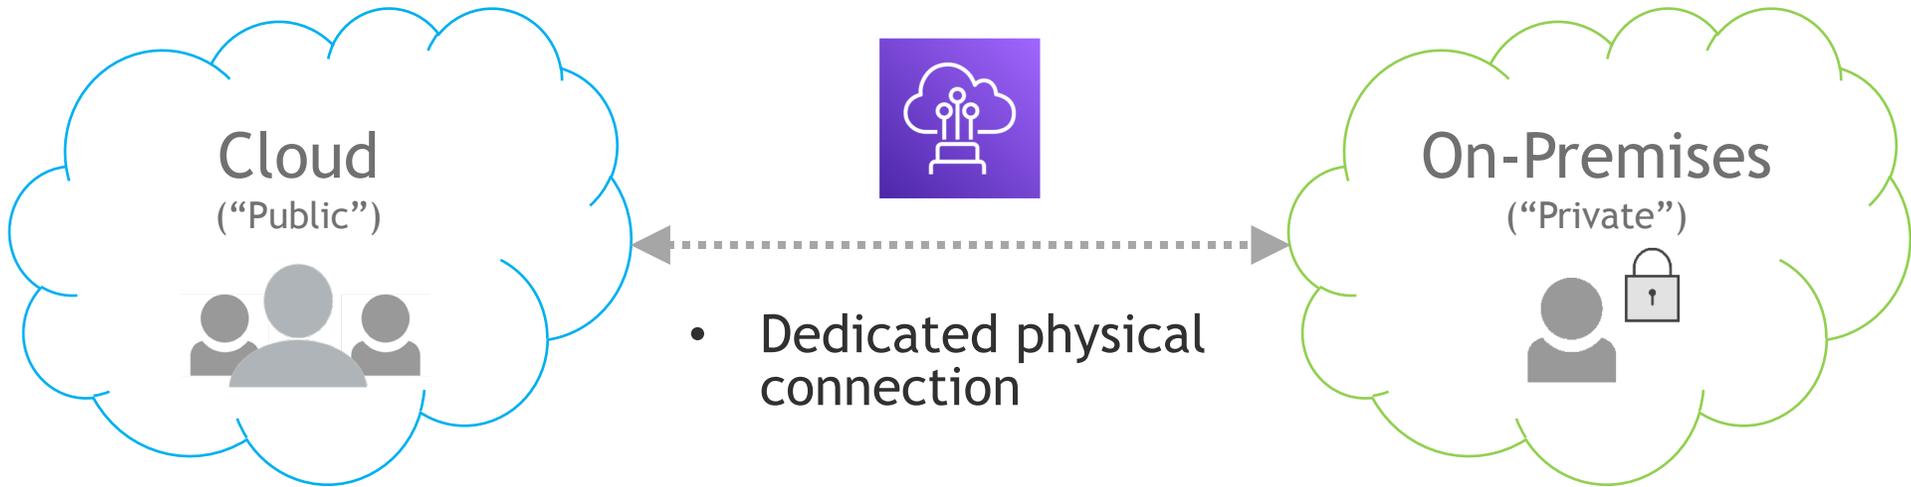
May be outsourced or owned, on-premises or off-premises



Site-to-Site VPN

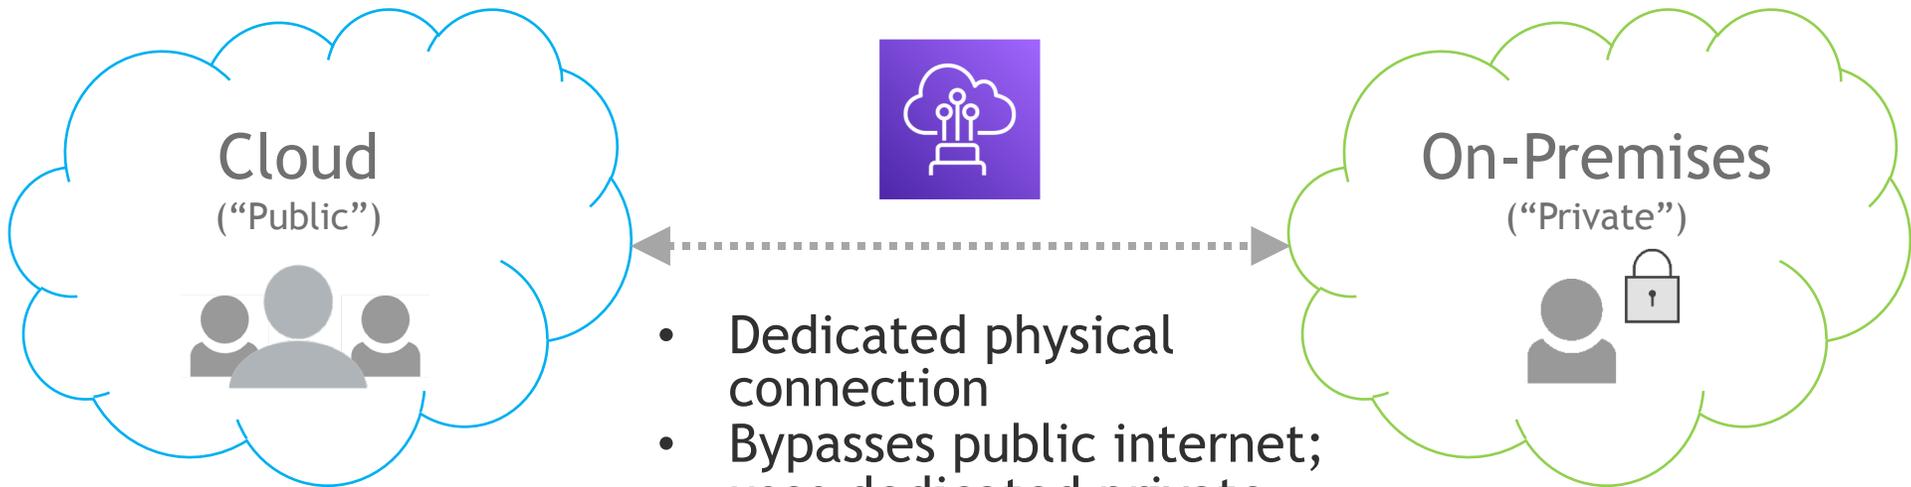


AWS Direct Connect





AWS Direct Connect



- Dedicated physical connection
- Bypasses public internet; uses dedicated private network (not encrypted by default)
- Faster and more expensive
- Takes longer to set up

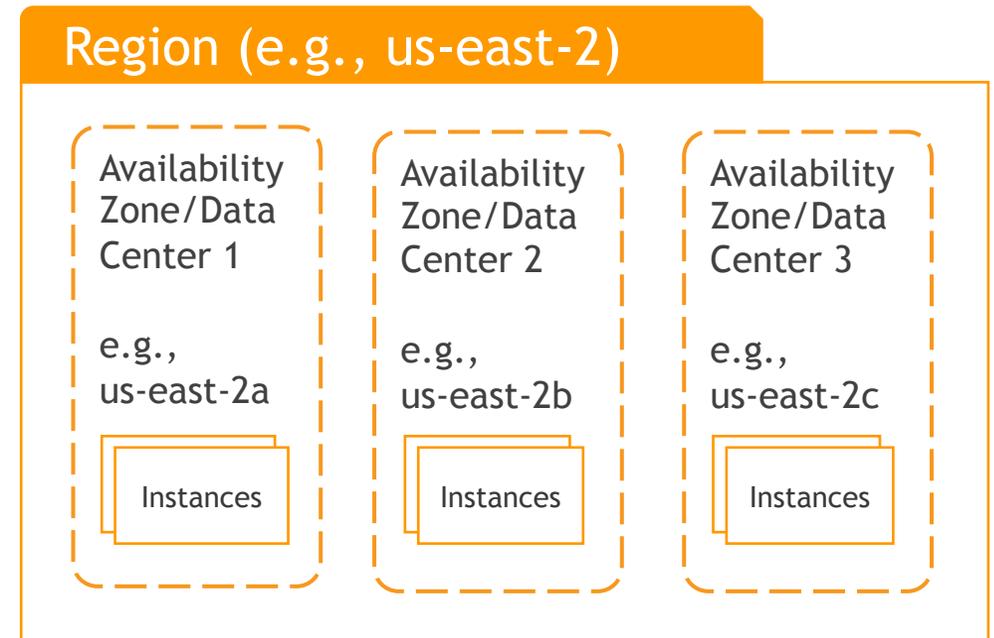
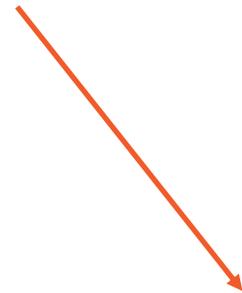
Availability Zone

A single **data center** or group of data centers within a region, each with redundant power, networking and connectivity, housed in separate facilities

Network ⓘ	vpc-0e90de444882ac5f3 (default) ⌵
Subnet ⓘ	No preference (default subnet in any Availability Zone) ⌵
Auto-assign Public IP ⓘ	No preference (default subnet in any Availability Zone)
Hostname type ⓘ	subnet-0c317952da0b1a69c Default in us-east-2a
DNS Hostname ⓘ	subnet-0c866e5580054dbd0 Default in us-east-2b
	subnet-07a3c236d6e2f20be Default in us-east-2c

Regions and Availability Zones

	Ohio ▲
US East (N. Virginia)	us-east-1
US East (Ohio)	us-east-2
US West (N. California)	us-west-1
US West (Oregon)	us-west-2
Africa (Cape Town)	af-south-1
Asia Pacific (Hong Kong)	ap-east-1
Asia Pacific (Jakarta)	ap-southeast-3
Asia Pacific (Mumbai)	ap-south-1
Asia Pacific (Osaka)	ap-northeast-3
Asia Pacific (Seoul)	ap-northeast-2
Asia Pacific (Singapore)	ap-southeast-1
Asia Pacific (Sydney)	ap-southeast-2
Asia Pacific (Tokyo)	ap-northeast-1
Canada (Central)	ca-central-1
Europe (Frankfurt)	eu-central-1
Europe (Ireland)	eu-west-1
Europe (London)	eu-west-2
Europe (Milan)	eu-south-1
Europe (Paris)	eu-west-3
Europe (Stockholm)	eu-north-1
Middle East (Bahrain)	me-south-1
South America (São Paulo)	sa-east-1



Edge Location

A site that Amazon CloudFront uses to store cached copies of your content closer to your customers for faster delivery

AWS Technology Concepts: Important Points to Remember

WAYS TO WORK WITH AWS

- AWS Management Console (browser)
- Programmatic access
 - From your local machine - both require access key IDs and secrets
 - Software Developer Kit (SDK)
 - Command Line Interface (CLI)
 - AWS CloudShell
 - Browser-based; doesn't require access key ID

DEPLOYMENT MODELS

- Public cloud
- On-premises
- Hybrid
 - Connectivity: Site-to-Site VPN (over public internet) or Direct Connect (dedicated physical connection)

Regions, Availability Zones and Edge Locations

	Ohio ▲
US East (N. Virginia)	us-east-1
US East (Ohio)	us-east-2
US West (N. California)	us-west-1
US West (Oregon)	us-west-2
Africa (Cape Town)	af-south-1
Asia Pacific (Hong Kong)	ap-east-1
Asia Pacific (Jakarta)	ap-southeast-3
Asia Pacific (Mumbai)	ap-south-1
Asia Pacific (Osaka)	ap-northeast-3
Asia Pacific (Seoul)	ap-northeast-2
Asia Pacific (Singapore)	ap-southeast-1
Asia Pacific (Sydney)	ap-southeast-2
Asia Pacific (Tokyo)	ap-northeast-1
Canada (Central)	ca-central-1
Europe (Frankfurt)	eu-central-1
Europe (Ireland)	eu-west-1
Europe (London)	eu-west-2
Europe (Milan)	eu-south-1
Europe (Paris)	eu-west-3
Europe (Stockholm)	eu-north-1
Middle East (Bahrain)	me-south-1
South America (São Paulo)	sa-east-1

Region (e.g., us-east-2)

Availability
Zone/Data
Center 1

e.g.,
us-east-2a

Instances

Availability
Zone/Data
Center 2

e.g.,
us-east-2b

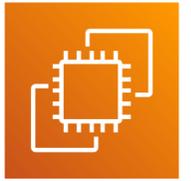
Instances

Availability
Zone/Data
Center 3

e.g.,
us-east-2c

Instances

- Edge locations are used by CloudFront to get content closer to customers
- High availability = 2 or more AZs
- Withstand environmental disaster = deploy to 2 or more regions



Elastic Compute Cloud (EC2)

A virtual server that you “rent” in AWS
Called an “Instance”

Easy to scale

Reliable

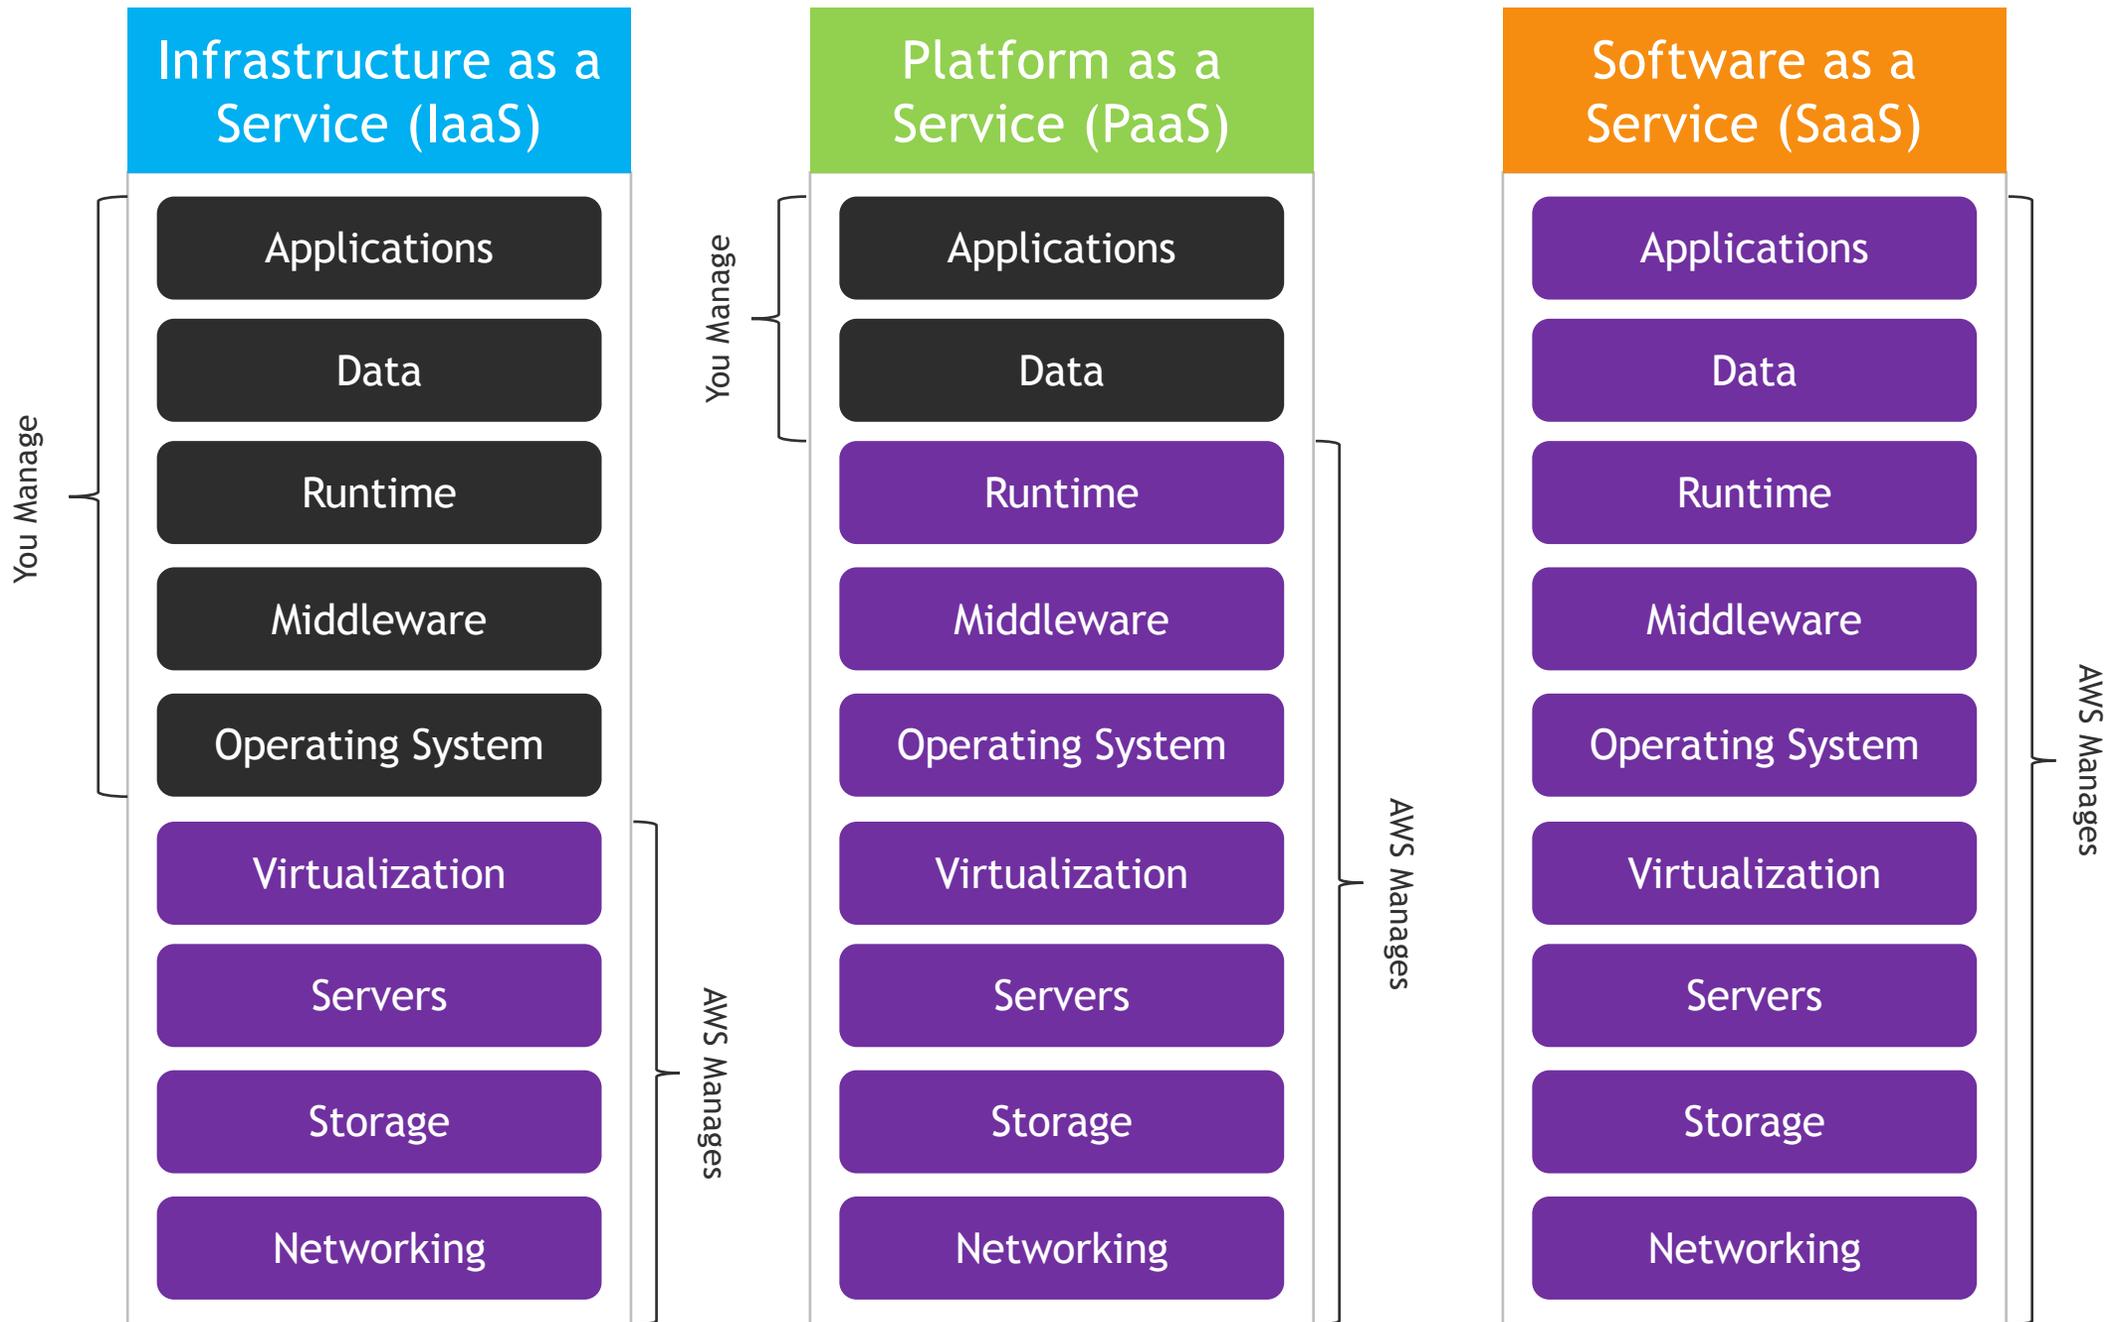
Use with other
services

Only pay for what
you use

Shared Responsibility Model for EC2

AWS Responsibility

Customer (Your) Responsibility



Shared Responsibility Model for EC2

AWS Responsibility

Physical hardware and network

Hypervisors (allow you to create virtual machines)

Elastic Block Store (EBS) storage system

Customer (Your) Responsibility

Software running on the instance, including operating system and applications (maintenance, patching, etc.)

Security groups and network access control lists

Network architecture within the Virtual Private Cloud (VPC)

EC2 Launch Types/Purchasing Options

TYPE	DESCRIPTION
On-Demand Instances	<ul style="list-style-type: none">• Default option• Pay by the second while instance is running (minimum 60 seconds)
Reserved Instances	<ul style="list-style-type: none">• Reserve instances for 1-3 years (e.g., database server)• Savings of up to 70%
Spot Instances	<ul style="list-style-type: none">• Bid on unused EC2 capacity• Savings up to 90%, but can lose instance at any time
Dedicated Hosts	<ul style="list-style-type: none">• Book an entire physical server and bring your own licenses• Can be purchased on-demand or reserved
Dedicated Instances	<ul style="list-style-type: none">• Ensures no other AWS customer shares your hardware• Pay by the hour
Capacity Reservations	<ul style="list-style-type: none">• Reserve capacity for instances in a specific availability zone for any duration



DEMO

Compute: Creating an EC2 Instance

Secure Shell (SSH)

Protocol used to securely log into other computers and run commands from the command line

Uses TCP protocol on Port 22



Options for Connecting to EC2 Instances

YOUR OPERATING SYSTEM	EC2 INSTANCE OPERATING SYSTEM	OPTIONS TO CONNECT
Linux	Linux	SSH client
Mac	Linux	SSH client
Windows 10 or later	Linux	PowerShell or cmd
Windows 8 or earlier	Linux	PuTTY
Windows (any version)	Windows	Remote Desktop Protocol (RDP)
Any (uses browser)	Linux	EC2 Instance Connect <i>Usually easiest option</i>



DEMO

SSH Connection from Linux



DEMO

SSH Connection from Mac



DEMO

SSH Connection from Windows 10 or Higher



DEMO

SSH Connection from Windows 8 or Earlier (PuTTY)



DEMO

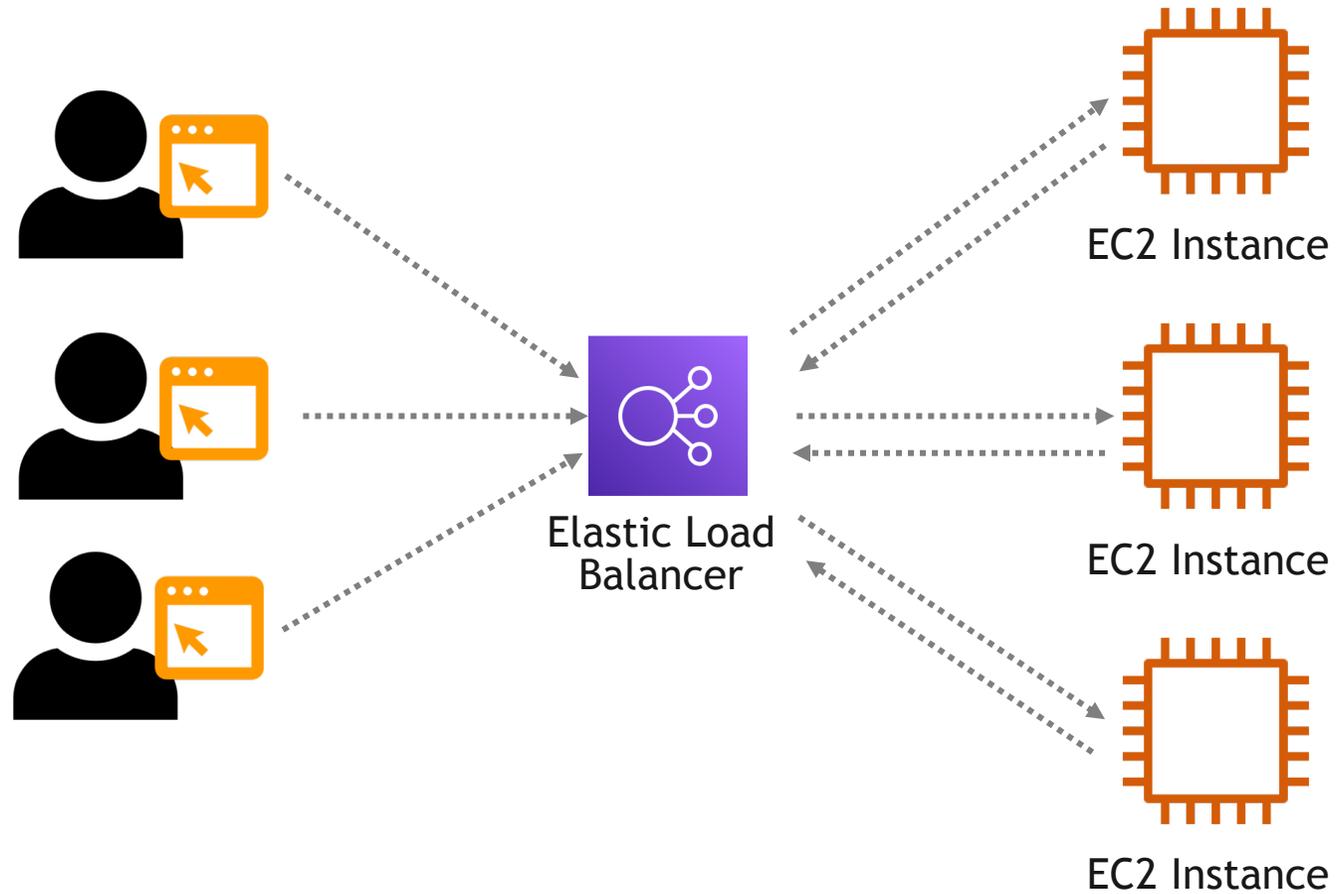
EC2 Instance Connect with Linux



DEMO

Remote Desktop with Windows

Elastic Load Balancing



Benefits of Using Elastic Load Balancing

Managed service

- AWS guarantees uptime
- AWS handles maintenance, upgrades, etc.

Distributes the load across multiple instances, and across availability zones

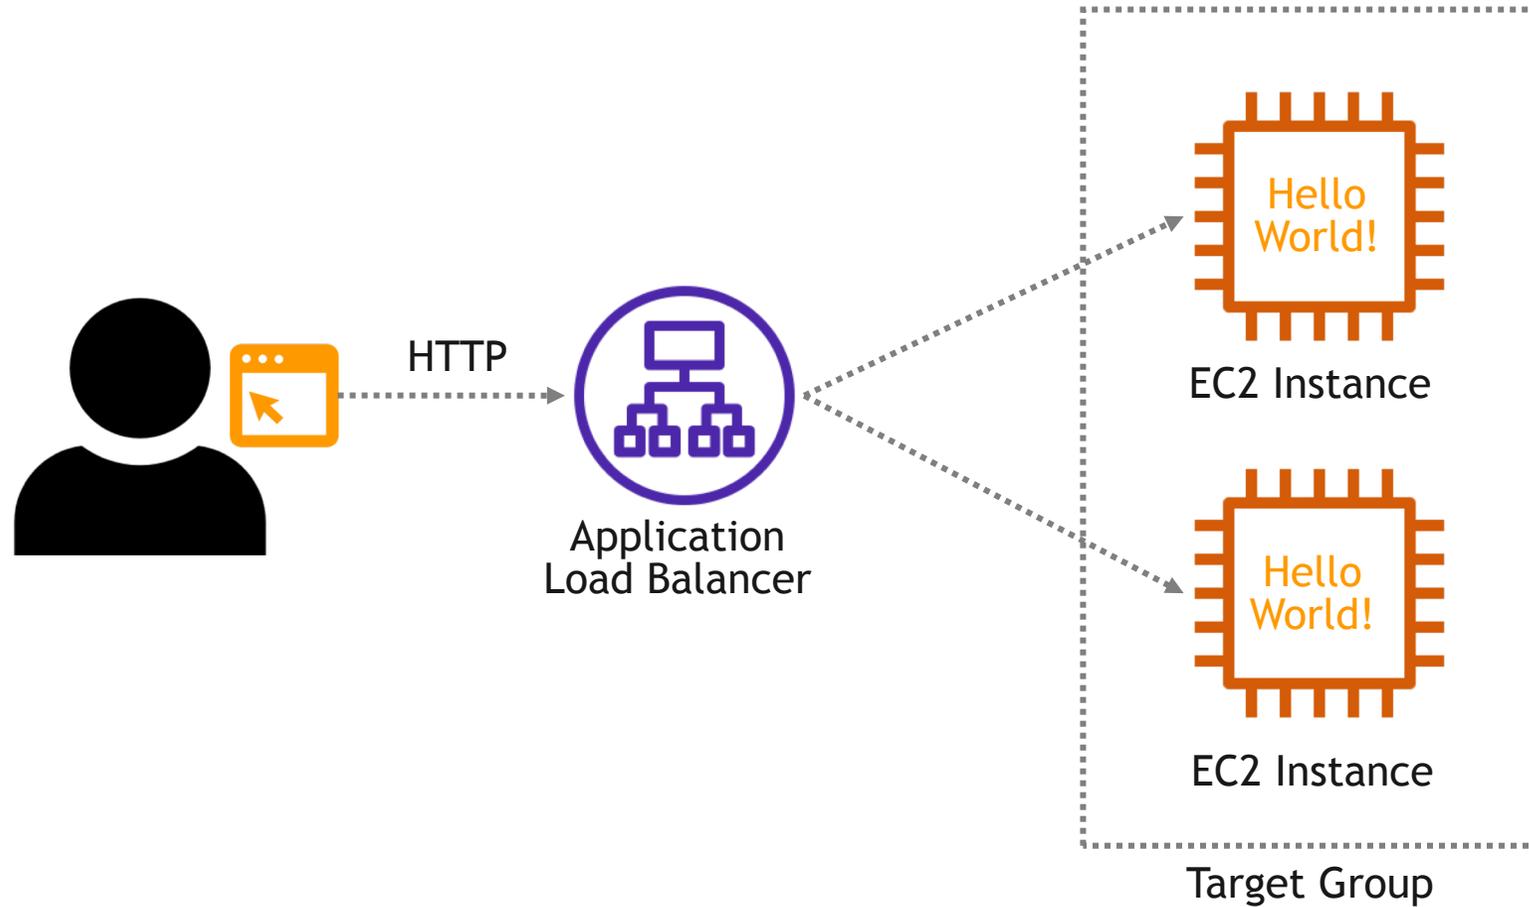
Can handle instance failures by doing regular health checks

Only exposes a single endpoint (DNS) to your application

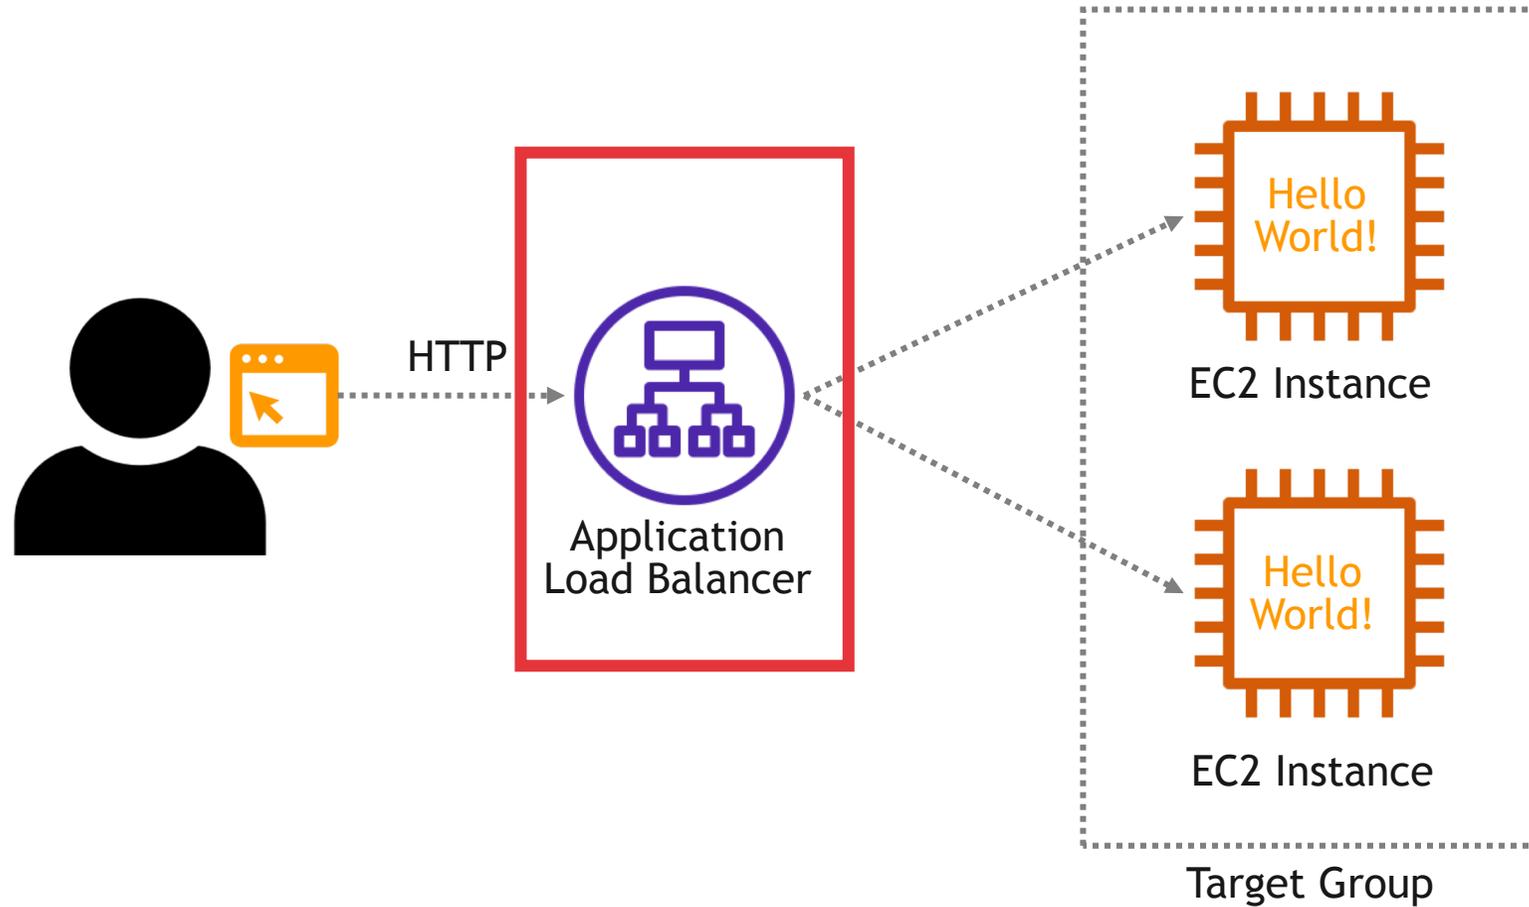
Types of Load Balancers

TYPE	DESCRIPTION
Application Load Balancer (ALB)	<ul style="list-style-type: none">• Handles HTTP/HTTPS traffic• Layer 7
Network Load Balancer	<ul style="list-style-type: none">• Handles TCP, UDP, TLS• Ultra-high performance, ultra-low latencies• Layer 4
Gateway Load Balancer	<ul style="list-style-type: none">• Used to deploy and manage third-party virtual appliances such as firewalls, intrusion detection and prevention systems
Classic Load Balancer	<ul style="list-style-type: none">• Previous generation load balancer for use with the EC2-Classic network

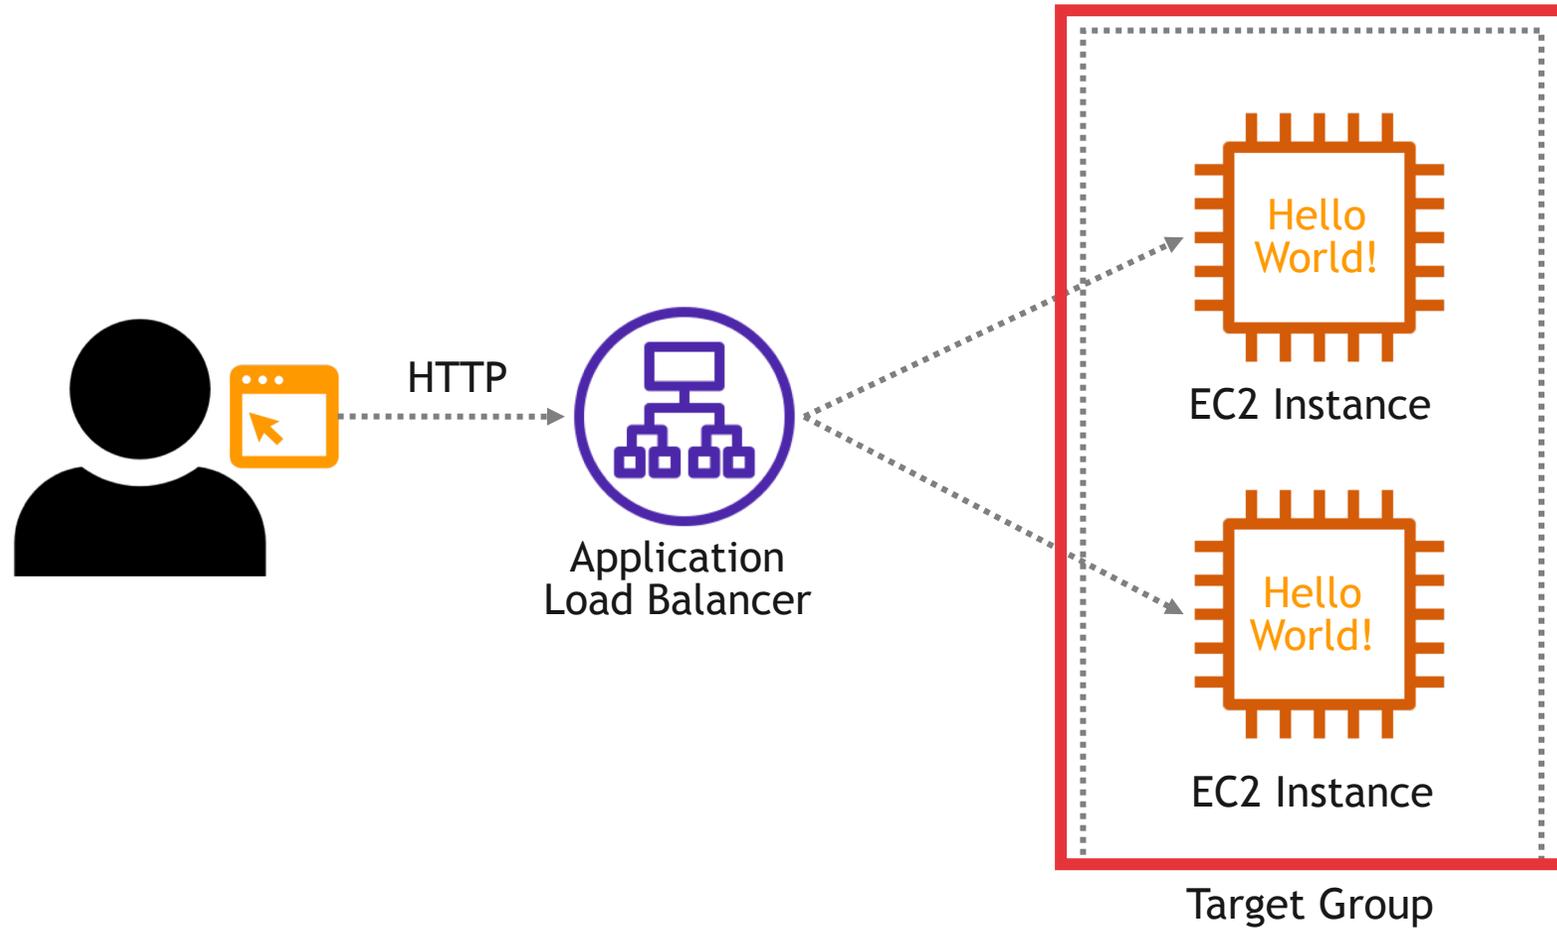
Demo: Application Load Balancer We're Building



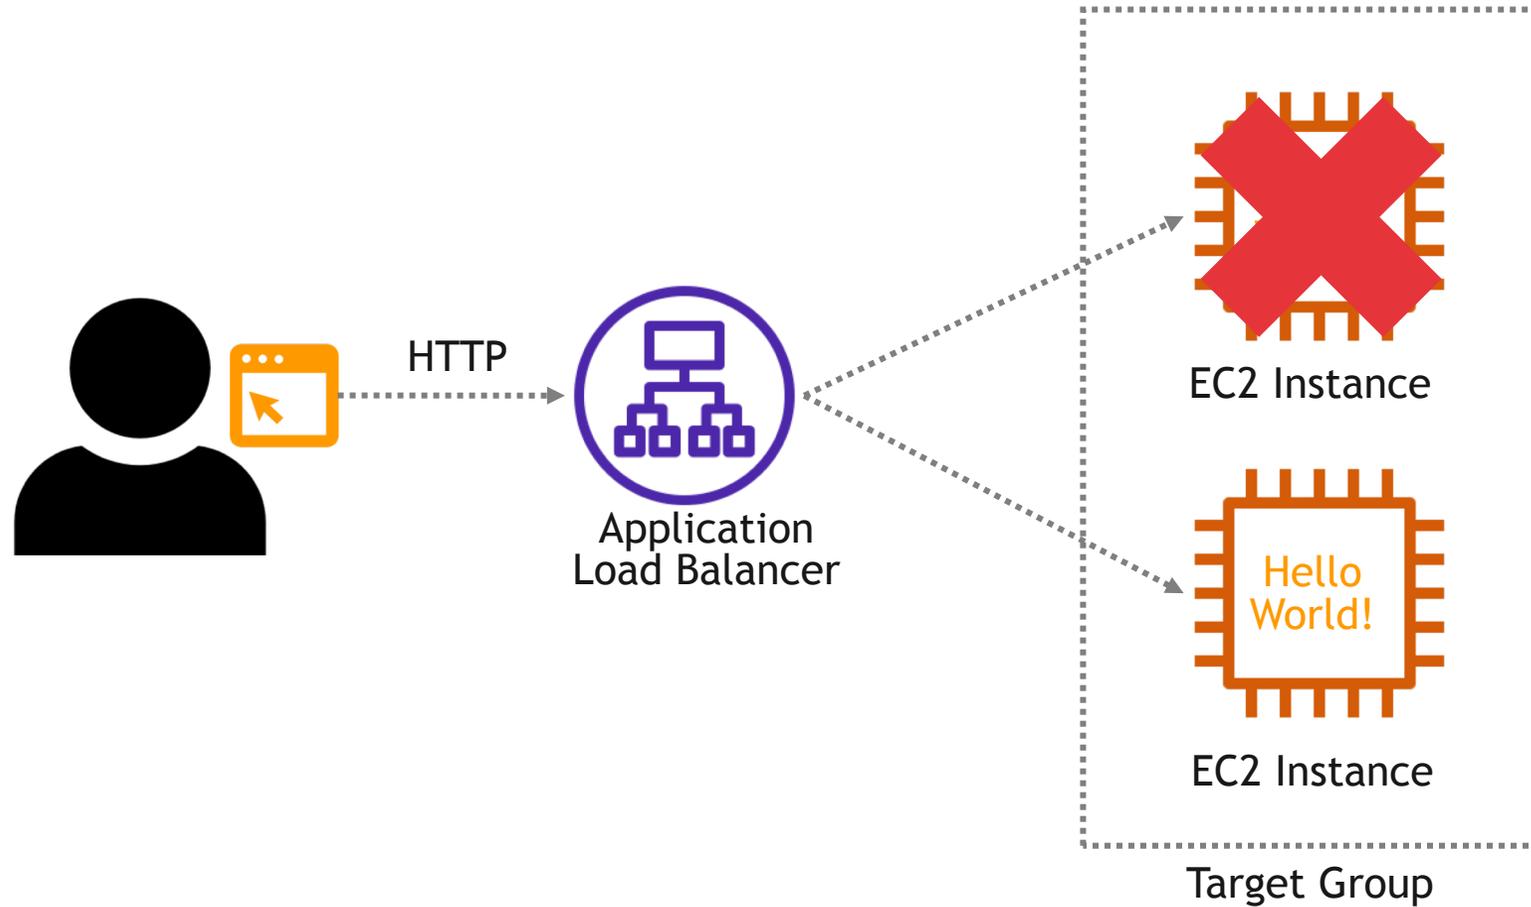
Demo: Application Load Balancer We're Building



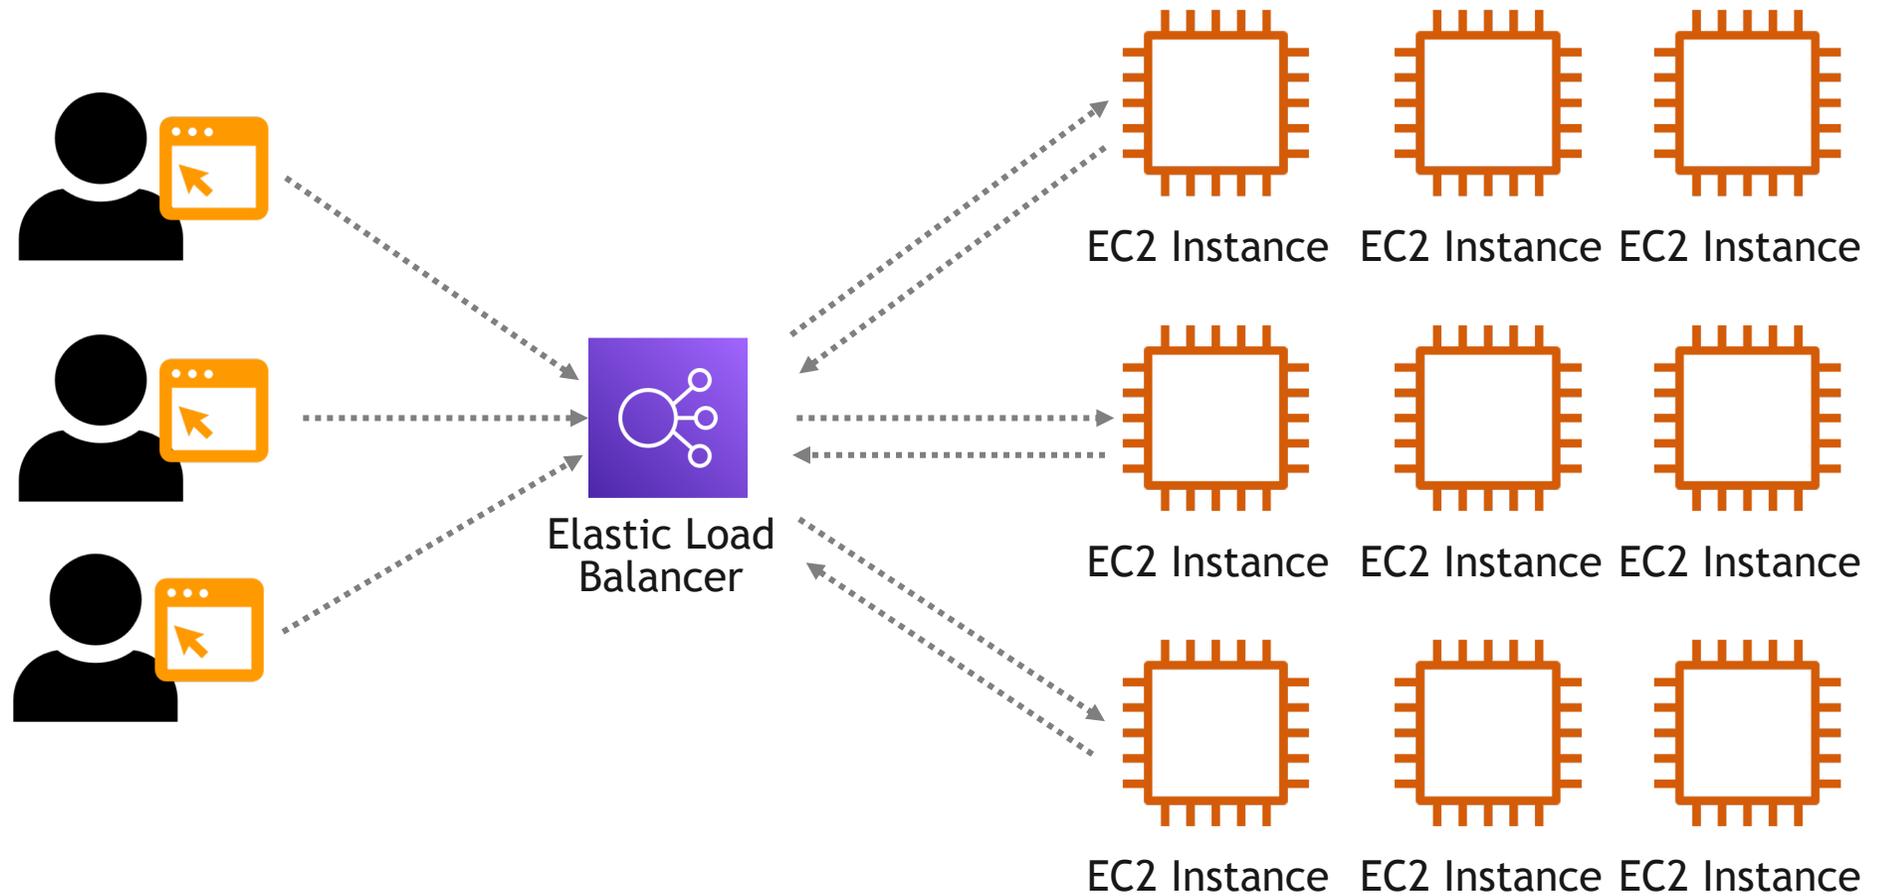
Demo: Application Load Balancer We're Building



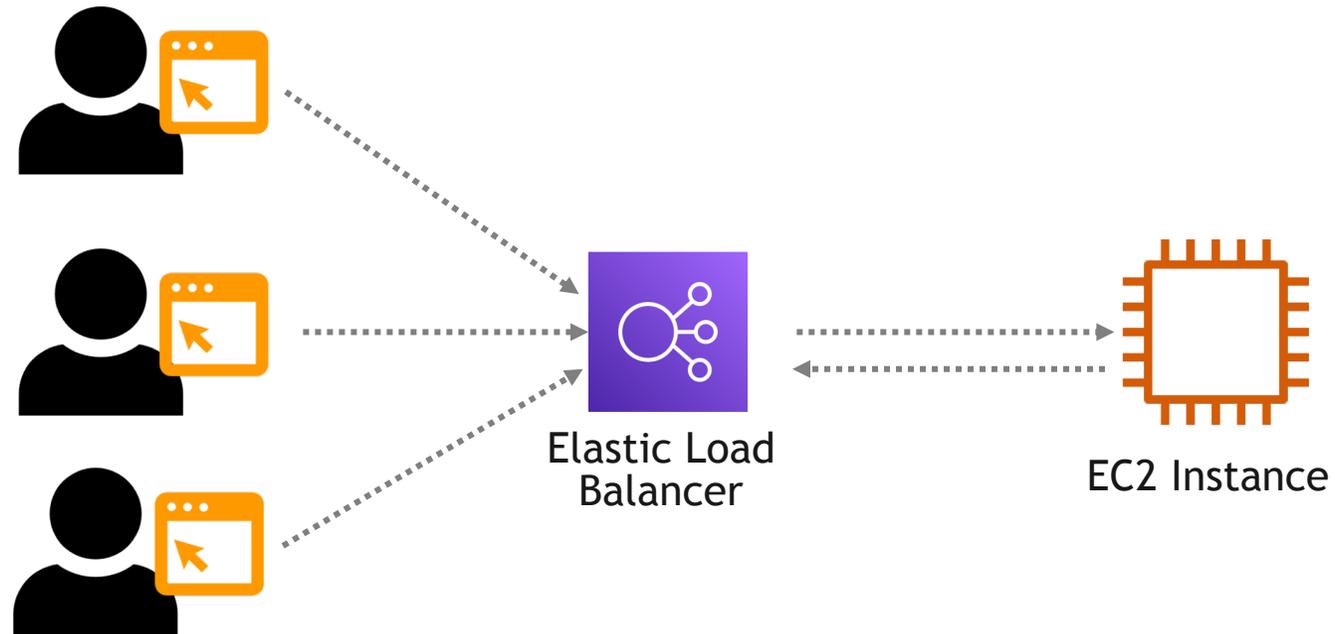
Demo: Application Load Balancer We're Building



Improving Our Load Balancing



Improving Our Load Balancing



Auto Scaling Groups (ASG)

Automatically scale out/in instances based on the load

- Specify instance numbers: desired capacity, minimum and maximum
- Registers new instances with the load balancer

Replace unhealthy instances automatically

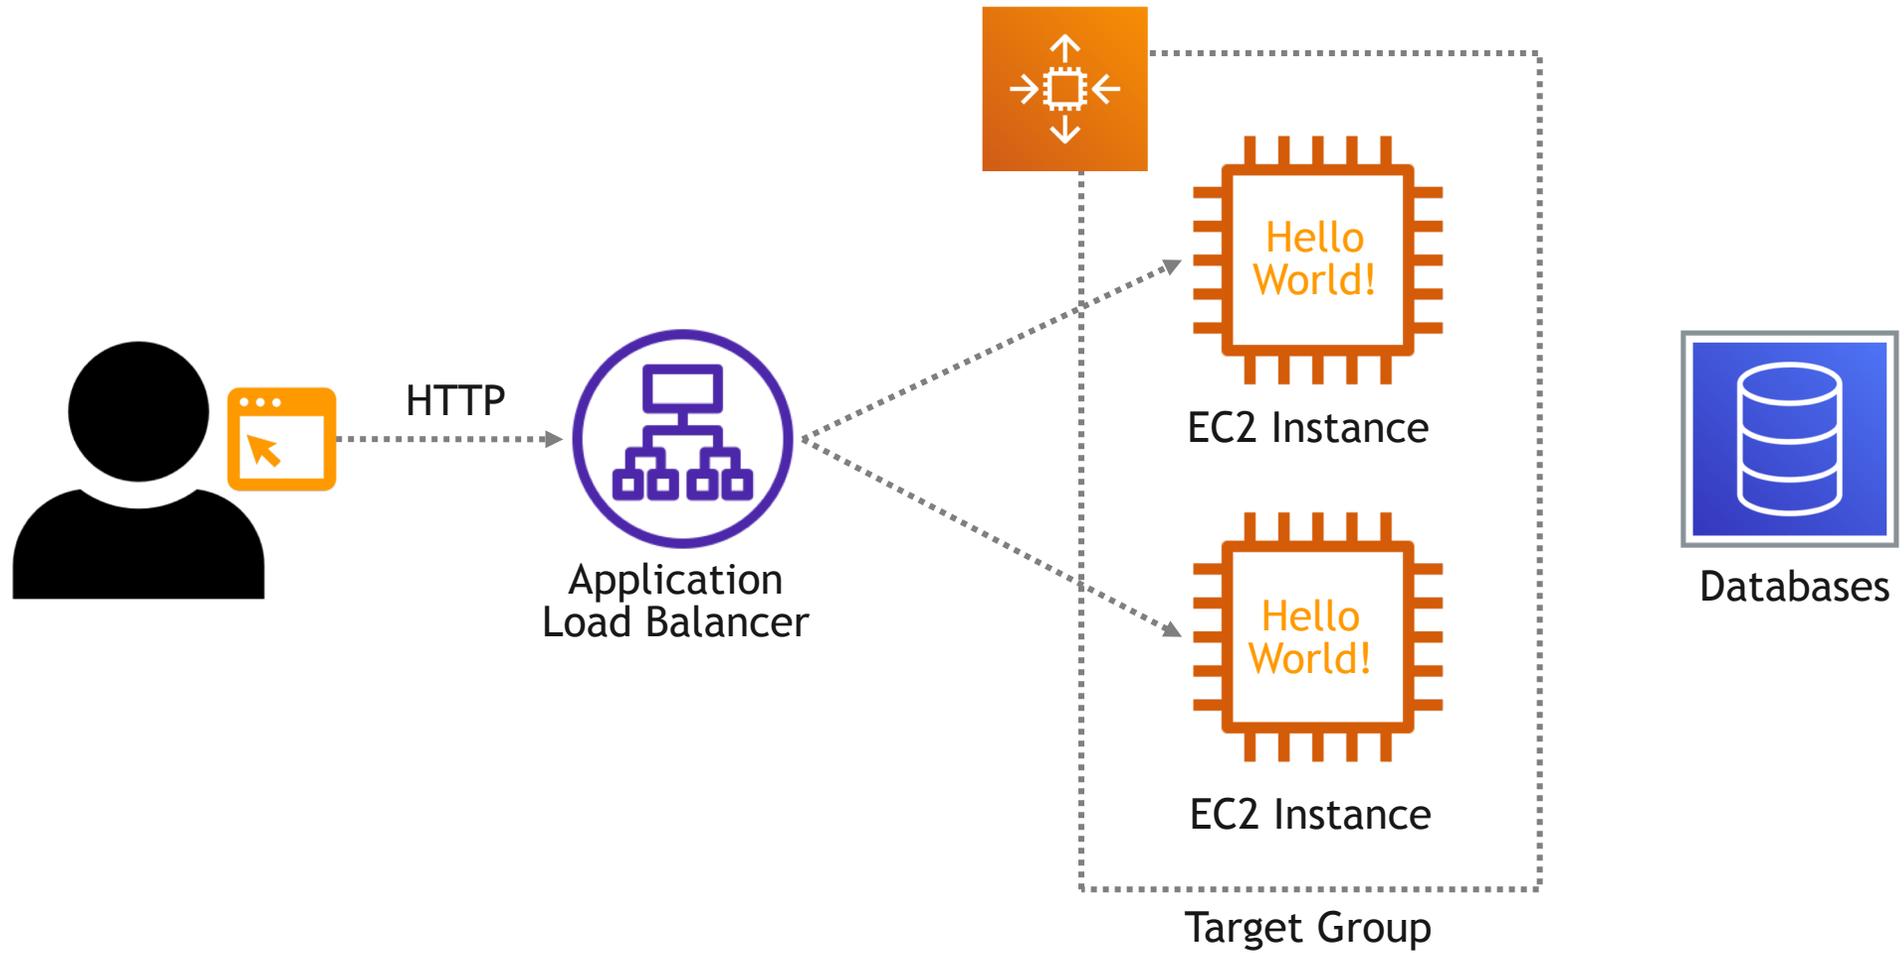
Run at optimal capacity, meaning cost savings



DEMO

Creating an Auto Scaling Group

All the Things



All the Things

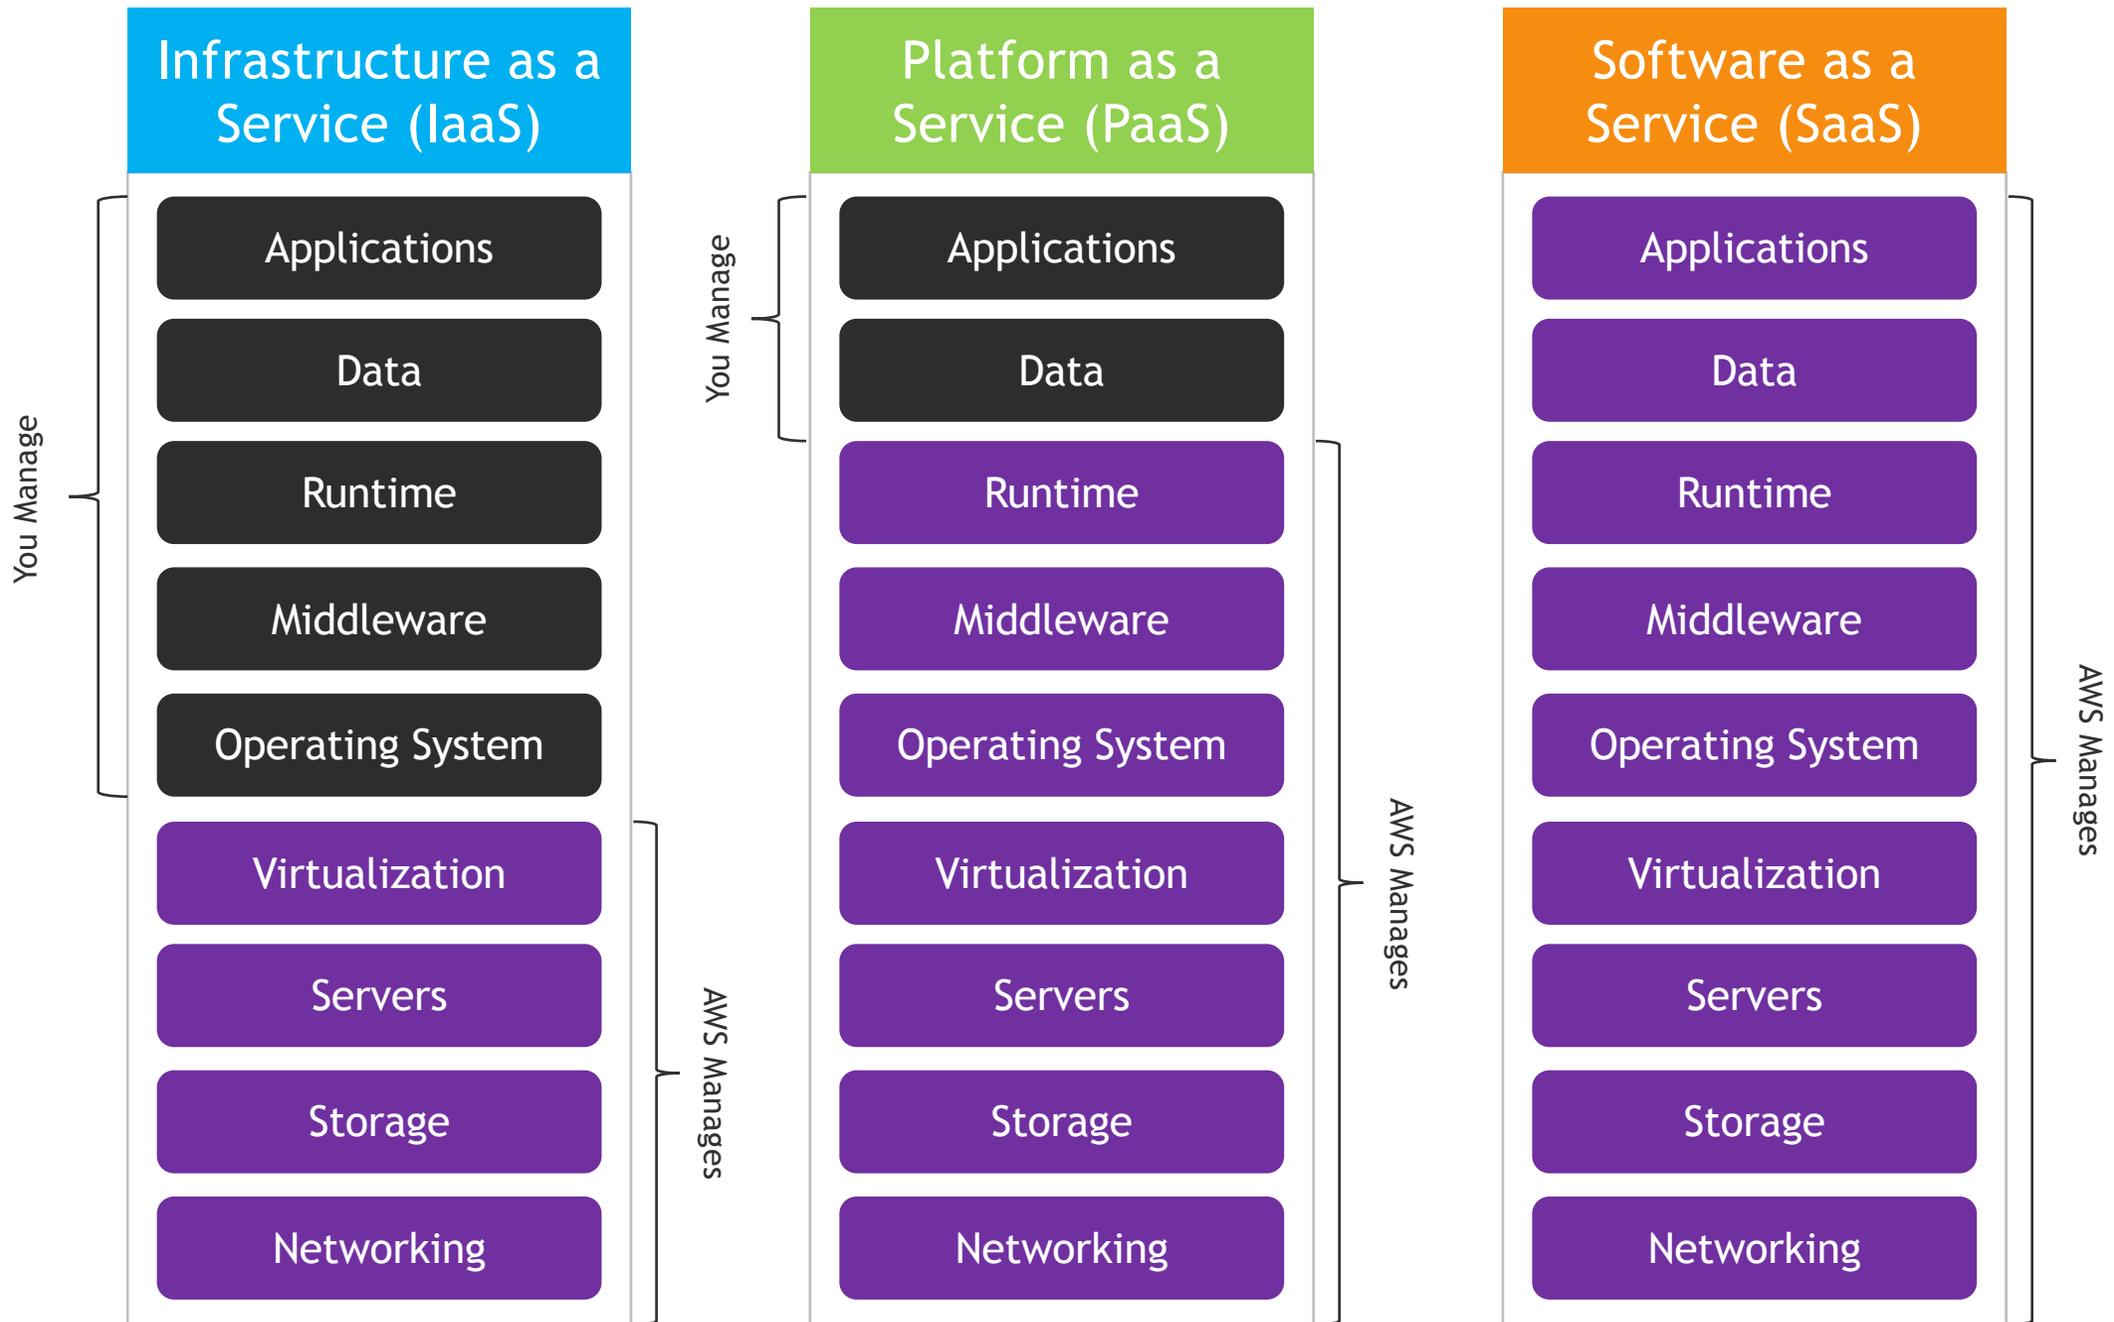
“I just want to
write code.”







AWS
Elastic
Beanstalk





AWS Elastic Beanstalk

Platform as a Service (PaaS)

- Infrastructure is handled for you
 - EC2 instances, load balancers, auto-scaling groups, database, etc.
- Uses CloudFormation templates to deploy the infrastructure

Health monitoring available in the Console

Support for many platforms

- Node, PHP, Java, Python, Go, .NET, Ruby, Docker, etc.



DEMO

Creating an Elastic Beanstalk Application

APPLICATION

ENVIRONMENTS

- Infrastructure (EC2 instances, load balancers, auto-scaling groups, etc.)
- Runs a single application version at a time
- An application can have many environments (like dev and test)

APPLICATION VERSIONS

- Application code
- Stored in S3
- An application can have many application versions (supporting rollback, etc.)

SAVED CONFIGURATIONS

- Defines how an environment and its resources behave
- Can be used to launch new environments
- An application can have many saved configurations





AWS
Batch

Managed service for running batch jobs

- Deep learning, big data, media encoding, image processing, email daily reports, etc.

You create the batch jobs (scripts, business logic)

- Defined as Docker images

AWS handles everything else (instances, scheduler, queues, retries, etc.)

- Processing at scale, optimized for the workload



Amazon Lightsail

AWS “lite”

- Meant for people with little to no cloud experience

Good for simple applications and websites

Pricing

- Predictable
- Less expensive



Amazon WorkSpaces

Desktops in the cloud

- Employees and contractors can easily access applications and desktops from anywhere

Windows and Linux operating systems, with common app bundles (e.g., LibreOffice, Microsoft Office, browsers, antivirus, etc.)

Replaces on-premises Virtual Desktop Infrastructure (VDI)



AWS
Lambda



Upload video

Create small thumbnail

Create large thumbnail

Create HD video



AWS Lambda

Serverless computing

AWS Lambda

~~Serverless computing~~

Runs on a server you don't have to buy or manage

AWS Lambda

Code runs on a server you don't have to buy or manage

AWS Lambda

Code runs in response to some event, on a server you don't have to buy or manage

AWS Lambda

Can also be thought of as “Scripts” or “Functions”



Upload video

Create small thumbnail

Create large thumbnail

Create HD video



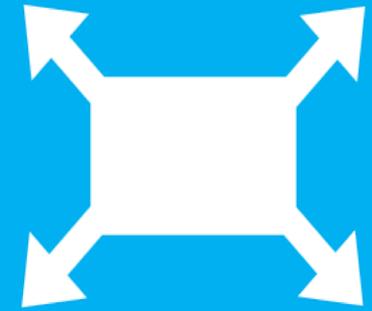


1

Single purpose



Right-sized



Easy to scale
(up AND down)

Shared Responsibility Model for Lambda

AWS Responsibility

Underlying infrastructure

Operating systems

Application platform

Customer (Your) Responsibility

Security of your code

Storage and access of data

Identity and access
management

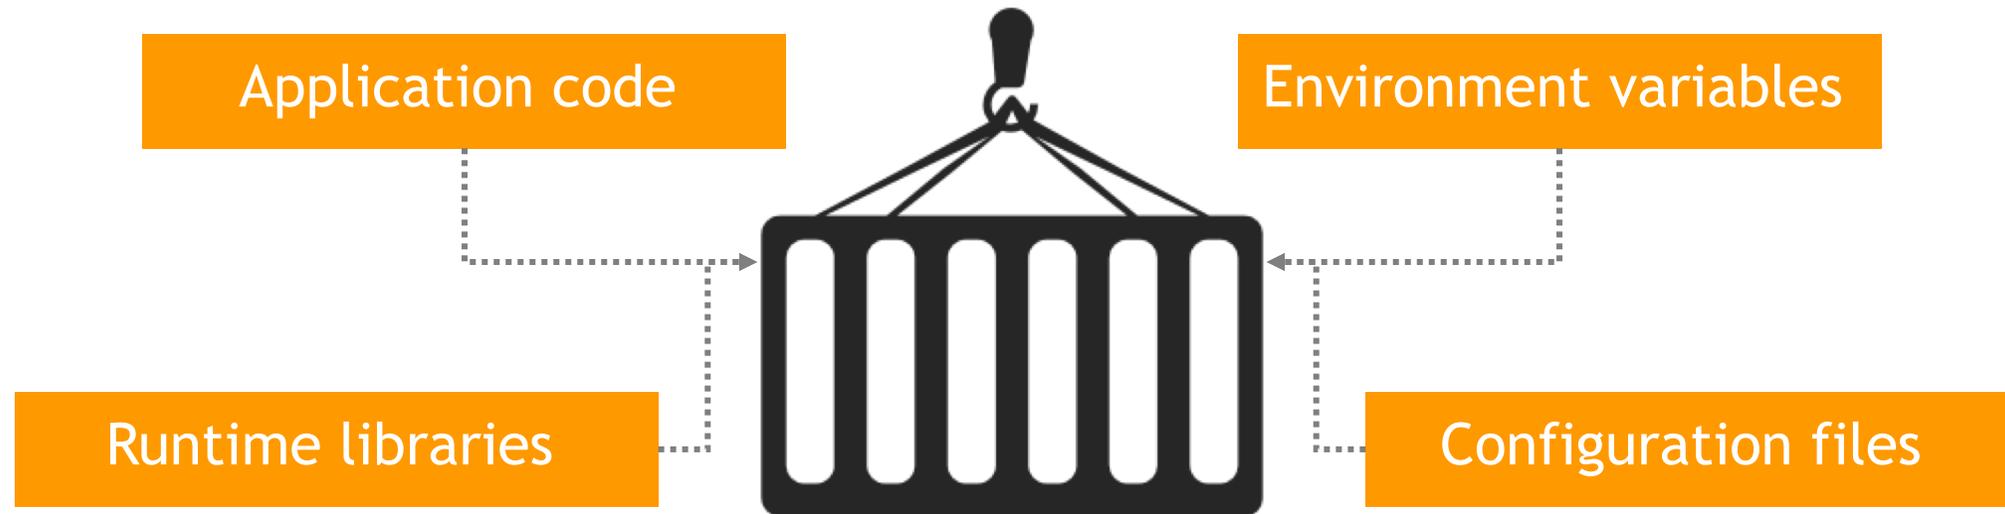




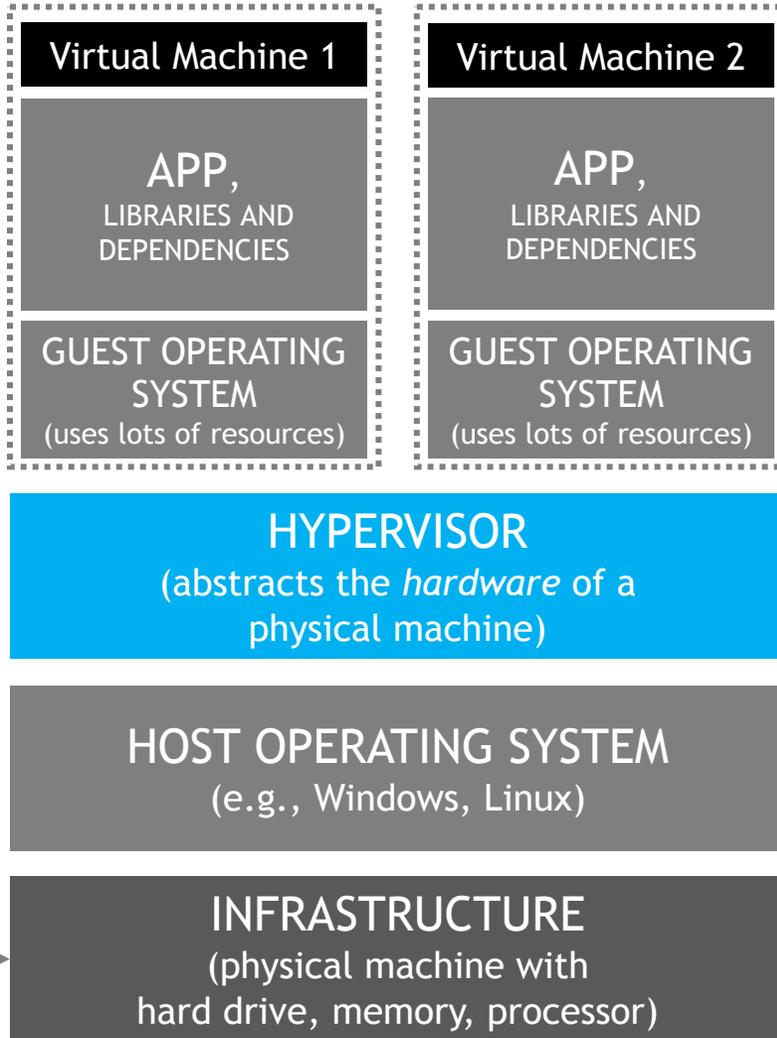
CONTAINERS OVERVIEW



Containers

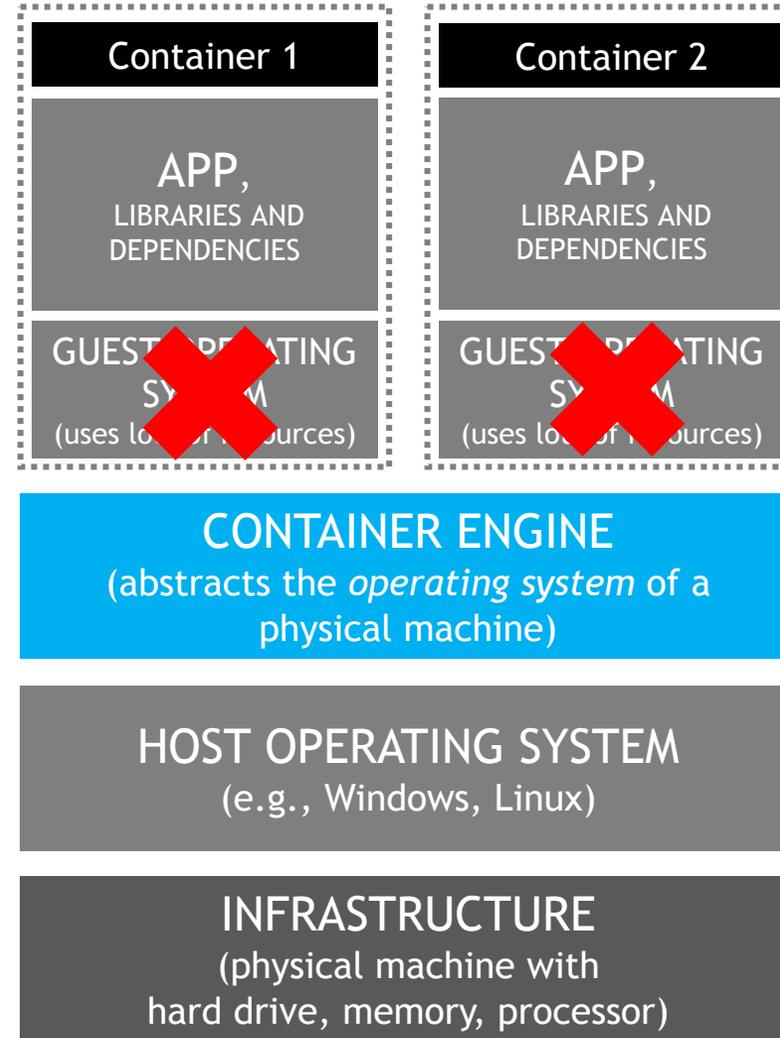


Each VM is a slice of the underlying hardware

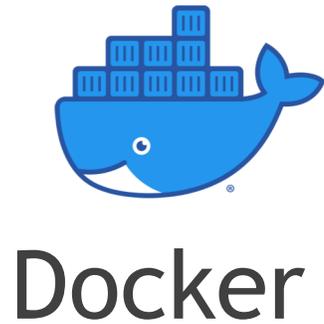


Virtual Machines

Each container is a slice of the underlying operating system



Containers



Underlying container technology used in AWS

Docker **images** are read-only templates/blueprints used to create Docker containers

- Created only once
- Stored in Docker repositories

Docker **containers** are a running instance of an image

- Create multiple containers from the same image



Amazon Elastic Container Service (ECS)

Used to run, stop and manage Docker containers in AWS

Integrated with many AWS services, including IAM, load balancers, auto scaling, VPC, etc.

Created by AWS



Amazon Elastic Kubernetes Service (EKS)

“Kubernetes as a service” on AWS

- Kubernetes is an open-source container orchestration platform
- Automates deploying, scaling, management and scheduling of containers
- Originally built by Google
- Has large community and support

AWS installs, operates, and maintains the Kubernetes cluster/nodes

Two Options for Running ECS/EKS

EC2 Instances

You provision and maintain the EC2 instances that the containers run on

Fargate



Serverless solution

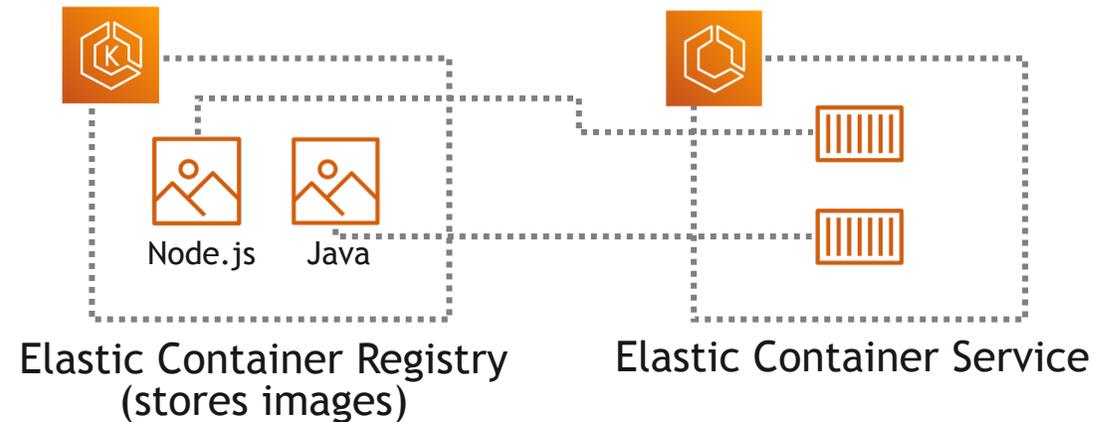
AWS handles the underlying infrastructure



Amazon Elastic Container Registry (ECR)

Private repository on AWS for Docker images

ECS, EKS pull images from here to build containers



Compute: Important Points to Remember

EC2

- Primary compute service used to create virtual servers
- Infrastructure as a service; YOU are responsible for operating systems, patching, access
- Different pricing options (on-demand, reserved, spot, etc.)
- SSH is a secure way to log into and administer remote computers
- AWS Marketplace allows partners to market and sell their software to AWS customers

LOAD BALANCING AND AUTO-SCALING

- Load balancing distributes load across multiple instances, across multiple availability zones (i.e., high availability)
- Auto-scaling automatically scales out/in instances based on the load

ELASTIC BEANSTALK

- Platform as a Service that allows you to “just write code”; all underlying work “magically” happens for you

Compute: Important Points to Remember

AWS LIGHTSAIL

- “AWS lite,” used to create simple websites without having to understand all of the underlying infrastructure
- Interface quite different from “regular” AWS

AMAZON WORKSPACES

- Desktops in the cloud
- Easily create workstations for employees and contractors (without having to set up/ship physical laptops or desktops)

AWS BATCH

- Managed service that allows you to schedule jobs for processing (e.g., media encoding, deep learning, big data analytics)

Compute: Important Points to Remember

AWS LAMBDA

- Serverless functions or scripts
- You're responsible for security of your code, data, and IAM; AWS does everything else

CONTAINERS

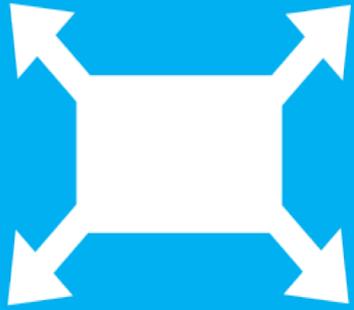
- Light-weight compute that lets you package up everything needed to run an application (no dedicated operating system required); built on Docker technology
- Services: Elastic Container Service (ECS) or Elastic Kubernetes Service (EKS)
- Fargate is a serverless solution for running containers (creates EC2 instances for you)
- Container images stored in Elastic Container Registry (ECR)

Simple Storage Service (S3)

Inexpensive object (think “file”) storage in AWS

“Unlimited” storage capacity

Organized by “buckets”



Easy to Scale



Reliable (99.99%)



Inexpensive



DEMO

Creating an S3 Bucket



DEMO

Hosting a Static Website on S3

S3 Storage Classes

<p style="text-align: center;"> FREQUENT (HOT) ← → ARCHIVE (COLD) FREQUENCY OF ACCESS </p>					
S3 Standard	S3 Standard-IA	S3 Intelligent-Tiering	S3 One Zone-IA	S3 Glacier	S3 Glacier Deep Archive
Frequently accessed data	Long-lived, infrequently accessed data	Long-lived data with changing or unknown access patterns	Long-lived, infrequently accessed, non-critical data	Long-term data archiving with retrieval times ranging from minutes to hours	Archiving rarely accessed data with a default retrieval time of 12 hours
≥ 3 AZs	≥ 3 AZs	≥ 3 AZs	1 AZ	≥ 3 AZs	≥ 3 AZs
No minimum storage duration	30-day minimum storage duration	30-day minimum storage duration	30-day minimum storage duration	90-day minimum storage duration	180-day minimum storage duration
				Must first restore objects before accessing them	Must first restore objects before accessing them

Cross-Region Replication

Store data closer
to your
customers

Store data closer
to your compute

Compliance
reasons

S3 Transfer Acceleration

Speeds up content transfers by as much as 50-500% for larger objects

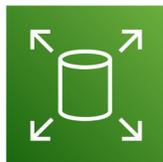
Routes traffic through CloudFront edge locations

Enabled at the bucket level



ELASTIC BLOCK STORE (EBS)





Elastic Block Store (EBS)

Storage for an EC2 instance

- Think of this as a virtual hard drive
 - Attached over the network (not physically)
- Root volume created for an EC2 instance; can add additional volumes as needed
- A volume can only be attached to one instance at a time
- Must be in the same availability zone as the instance

Default persistence settings

- Persists when an instance is shut down
- Deletes when an instance is terminated



Elastic Block Store (EBS)

Can create a snapshot of a volume and then restore it

- Stored in S3
- Useful for backup/recovery or copying to another availability region



DEMO

Creating and Attaching an EBS Volume



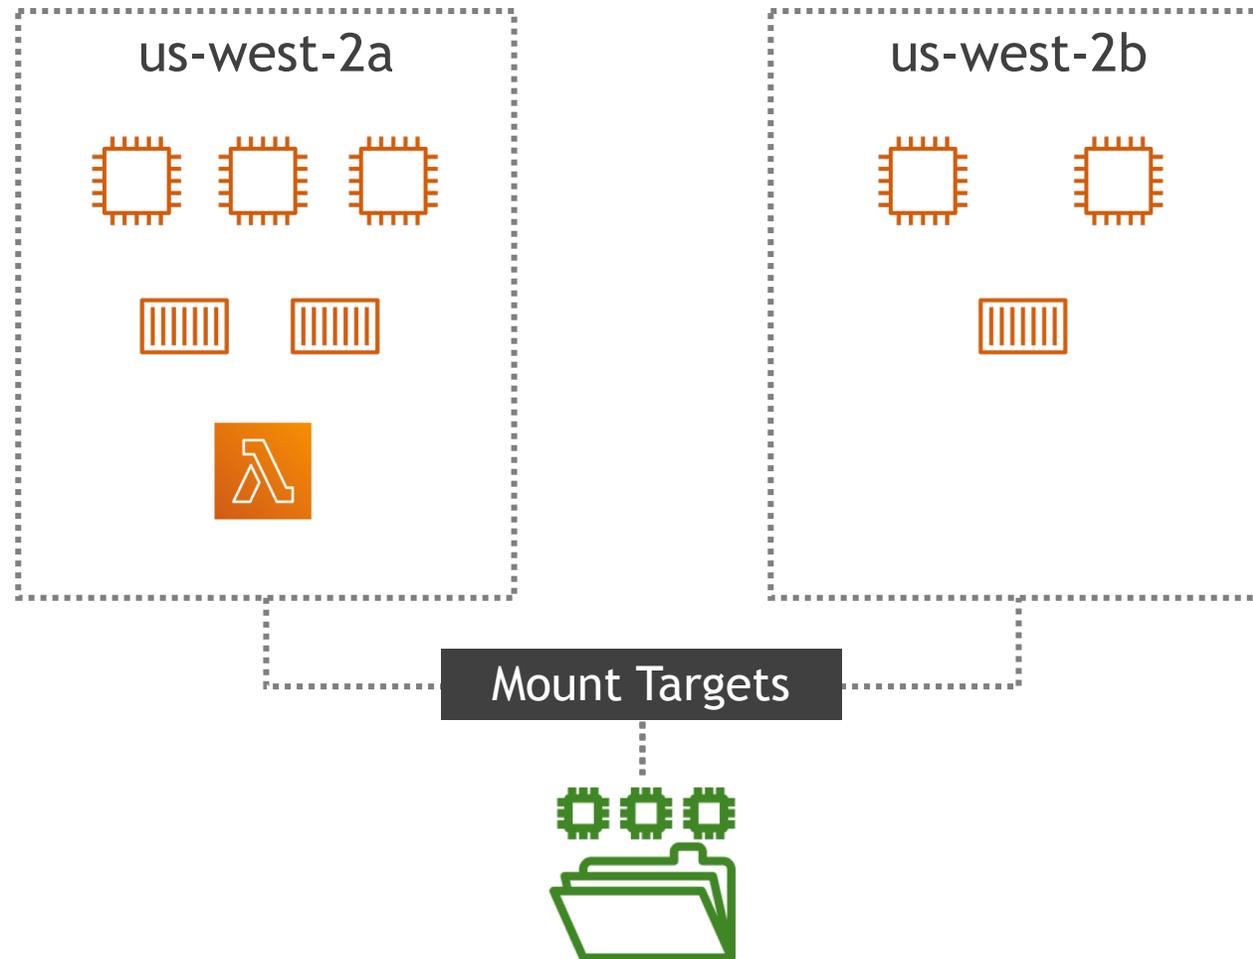
Elastic File System (EFS)

Storage for an EC2 instances,
containers, Lambda functions,
on-premises servers

- Can be used by **MULTIPLE** things at a time
- Can be used across **MULTIPLE** availability zones

Only pay for the storage you use (no
up-front provisioning)

Elastic File System (EFS)





Step 4: Add Storage

Your instance will be launched with the following storage device settings. You can attach additional EBS volumes and instance store volumes to your instance, or edit the settings of the root volume. You can also attach additional EBS volumes after launching an instance, but not instance store volumes. [Learn more](#) about storage options in Amazon EC2.

Volume Type	Device	Snapshot	Size (GiB)	Volume Type	IOPS	Throughput (MB/s)
Root	/dev/sda1	snap-047fe1cb9f3c81d95	30	General Purpose SSD (gp2)	100 / 3000	N/A

Add New Volume

Free tier eligible customers can get up to 30 GB of EBS General Purpose (SSD) or Magnetic storage. [Learn more](#) about free usage tier eligibility and usage restrictions.

Shared file systems

You currently don't have any file systems on this instance. Select "Add file system" button below to add a file system.

Add file system



Fully managed third-party file systems

- Benefits of the cloud (replication, high availability, etc.), but also features specific to the third-party (Active Directory integration, NTFS, etc.)

Amazon FSx for Windows File Server

- Built on Windows Server

Amazon FSx for Lustre

- High-performance file system

Storage Gateway

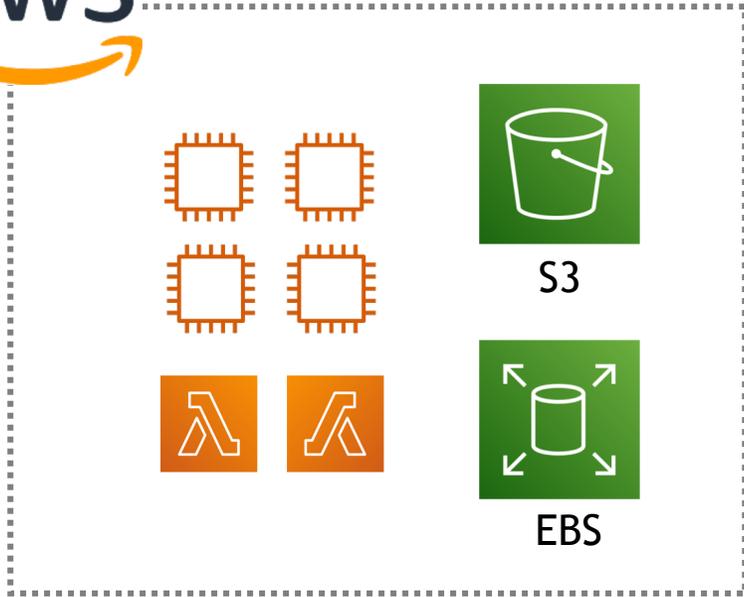


Storage Gateway



Storage Gateway

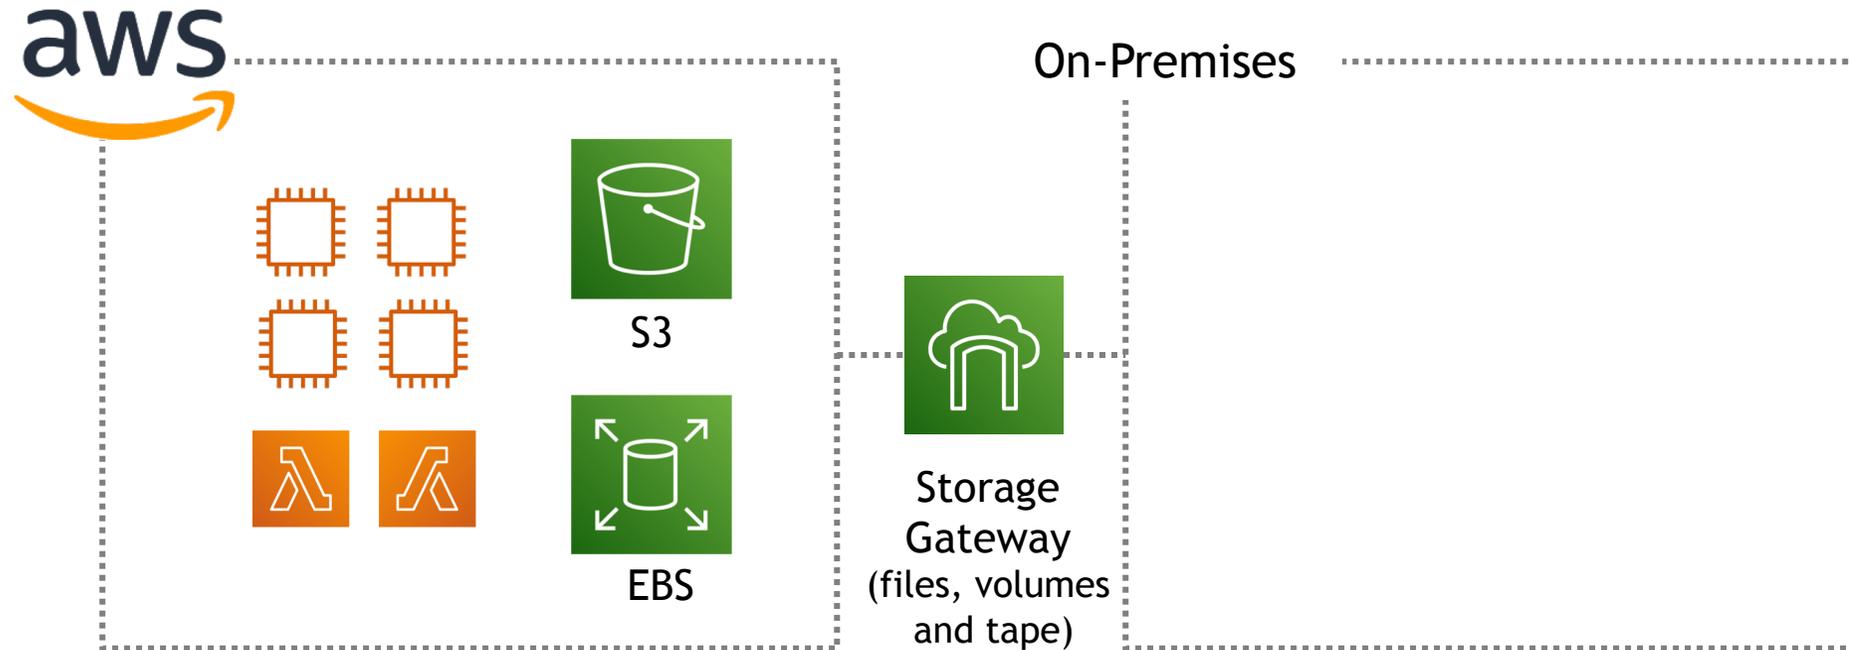
aws



On-Premises



Storage Gateway



A photograph of a server room with rows of server racks on both sides of a central aisle. The racks are dark with some blue lights visible. The ceiling has recessed lighting. A semi-transparent white banner is overlaid across the middle of the image.

Data. Lots and lots of data.

Use Cases

Migrating an on-premises data center

Backing up or restoring massive amounts of data

Edge computing and storage

- Limited internet

AWS Snow Family

AWS SNOWCONE



AWS SNOWBALL



AWS SNOWMOBILE



Backup Features for Various AWS Services

	Backup Features
Elastic Block Storage (EBS)	<ul style="list-style-type: none">• Snapshots of EBS volumes
Relational Database Service (RDS)	<ul style="list-style-type: none">• Automated backup• Snapshots• Read replicas
DynamoDB	<ul style="list-style-type: none">• On-demand table backups
Storage Gateway	<ul style="list-style-type: none">• Back up on-premises files, applications, databases, and volumes to AWS



AWS Backup

Manage backups in a single place

Configure schedules and retention periods

Storage: Important Points to Remember

S3

- Inexpensive, unlimited, reliable object storage (think: files, photos, logs, etc.)
- Storage classes let you optimize cost by storing data according to frequency of access
 - e.g., S3 Glacier can be used to store data you rarely access, but it takes some time to retrieve
- Versioning and object lock help prevent accidental deletion of objects

ELASTIC BLOCK STORE (EBS)

- Virtual hard drives for your EC2 instances
- Can only be attached to one instance at a time; default is to delete them when an instance is terminated
- Can create snapshots and then restore them (for backup/recovery purposes)

ELASTIC FILE SYSTEM (EFS)

- Storage that can be used by multiple services (EC2, Lambda, on-premises servers, etc.)
- Only pay for what you use (no up-front provisioning)

Storage: Important Points to Remember

AMAZON FSx

- Third-party file systems
- FSx for Windows File Server - built on Windows Server
- FSx for Lustre - high-performance file system

STORAGE GATEWAY

- Used in a hybrid architecture, to store data in AWS from on-premises

SNOW FAMILY OF PRODUCTS

- Physical storage devices used to transfer massive amounts of data (Snowcone, Snowball, Snowmobile)



Virtual Private Cloud (VPC)

Your own private cloud/network within the cloud

Isolates your resources from everyone else's

We'll focus on the
basics of networking



VPC



VPC

VPC

VPC

VPC

VPC

VPC



Internet Gateway

VPC

VPC



VPC

VPC

VPC



Internet Gateway

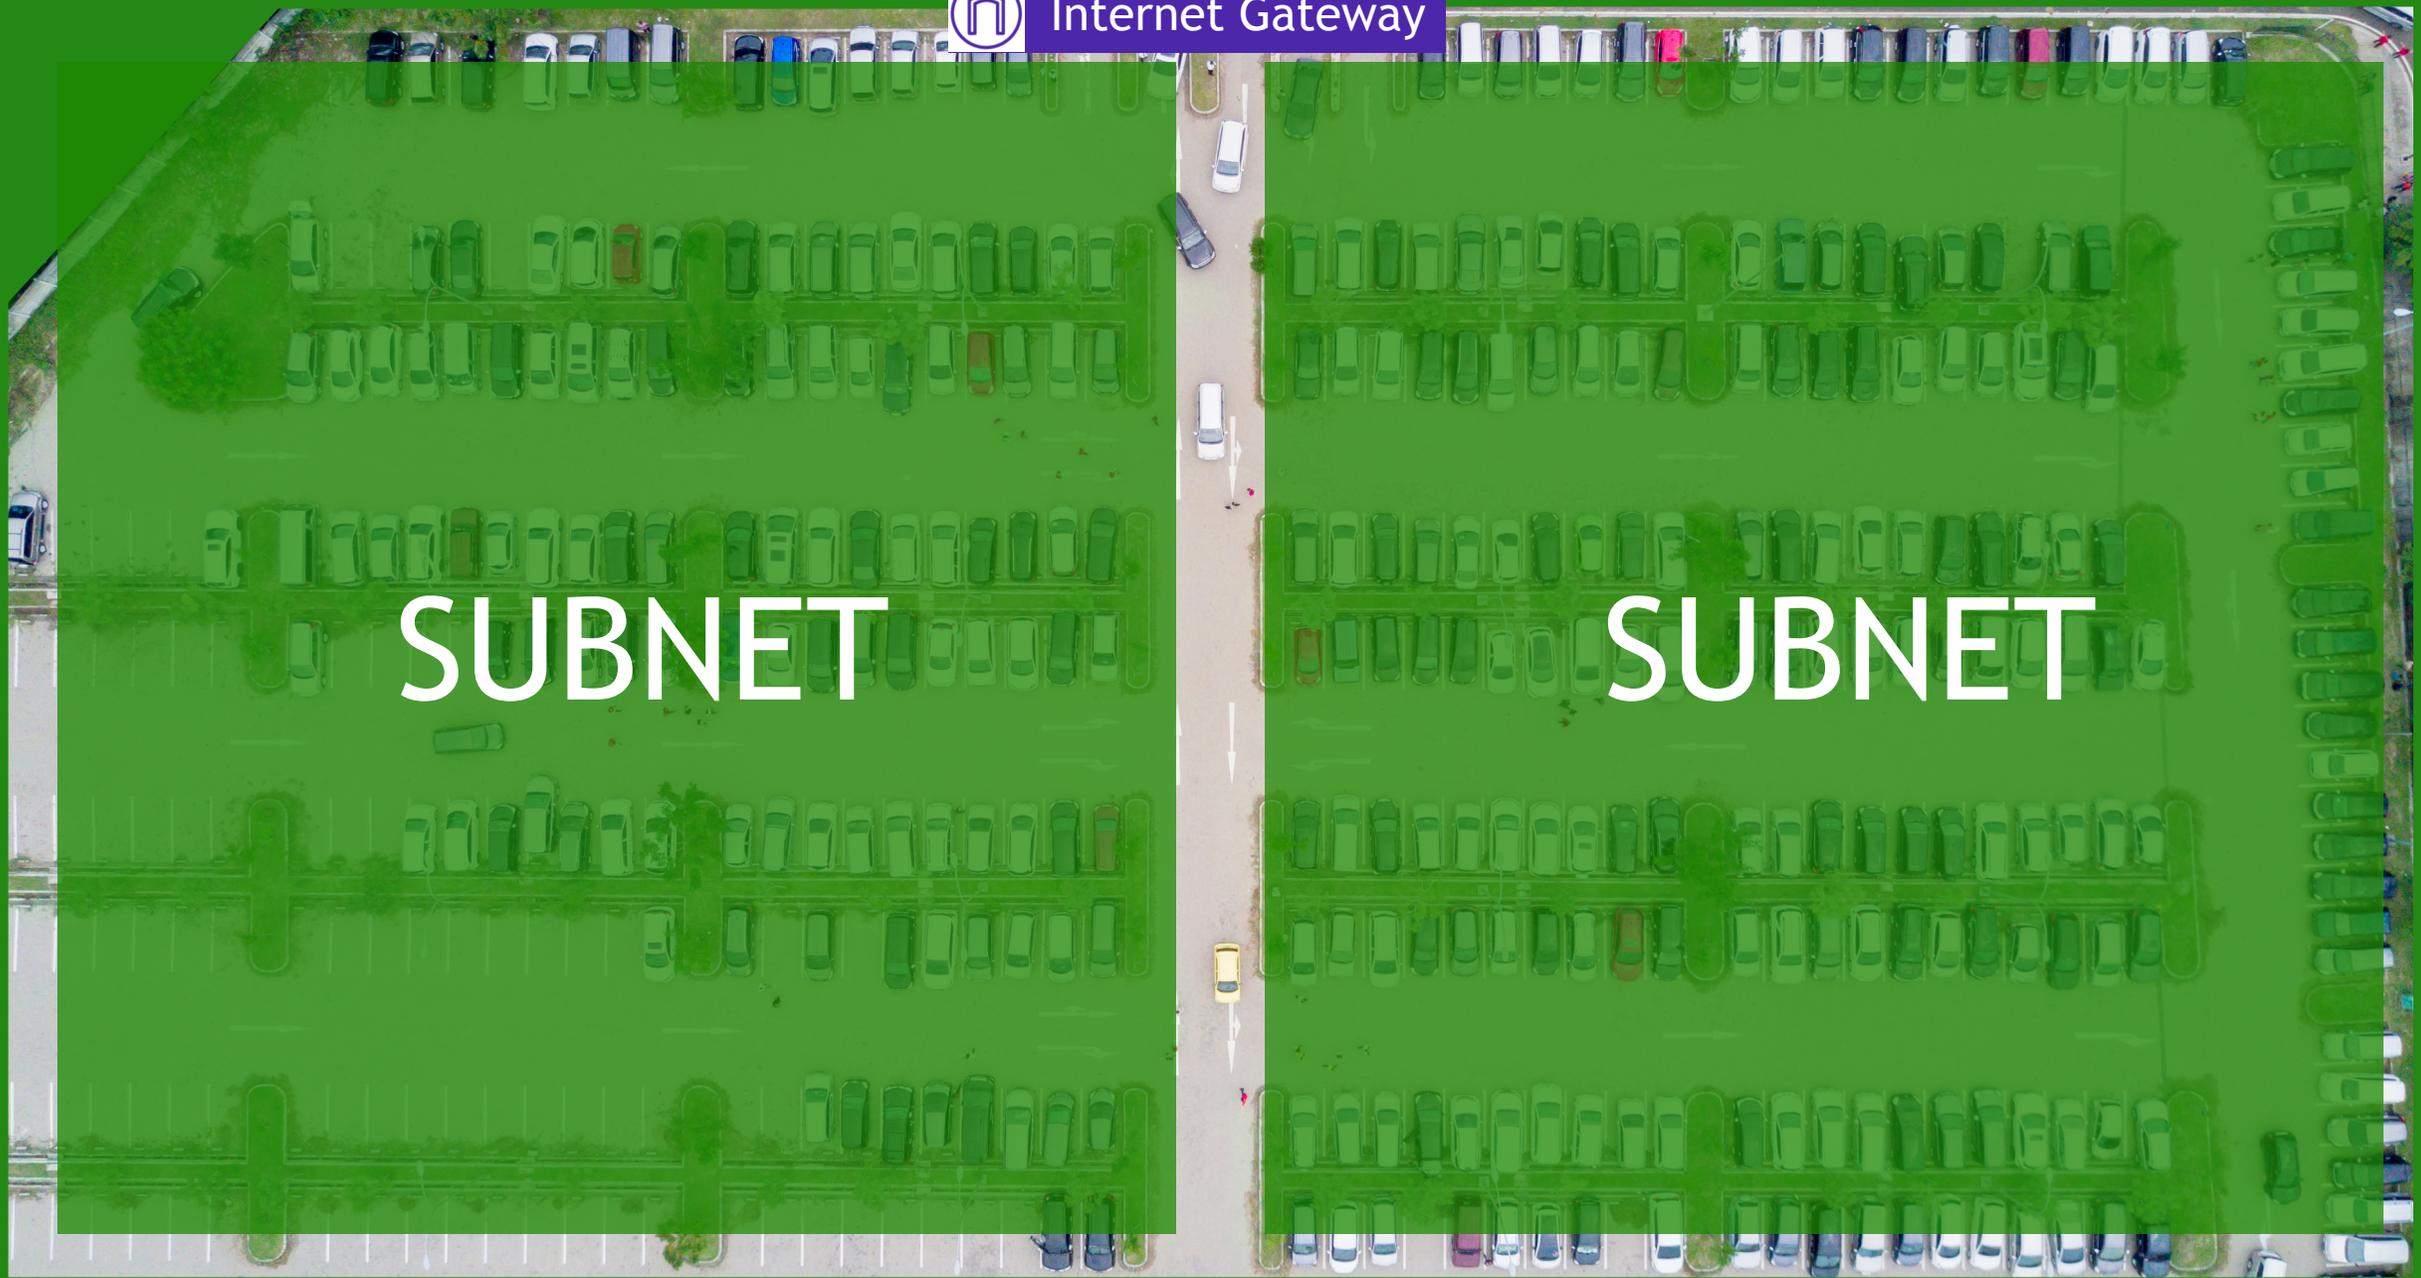
VPC

SUBNET

SUBNET

VPC

VPC



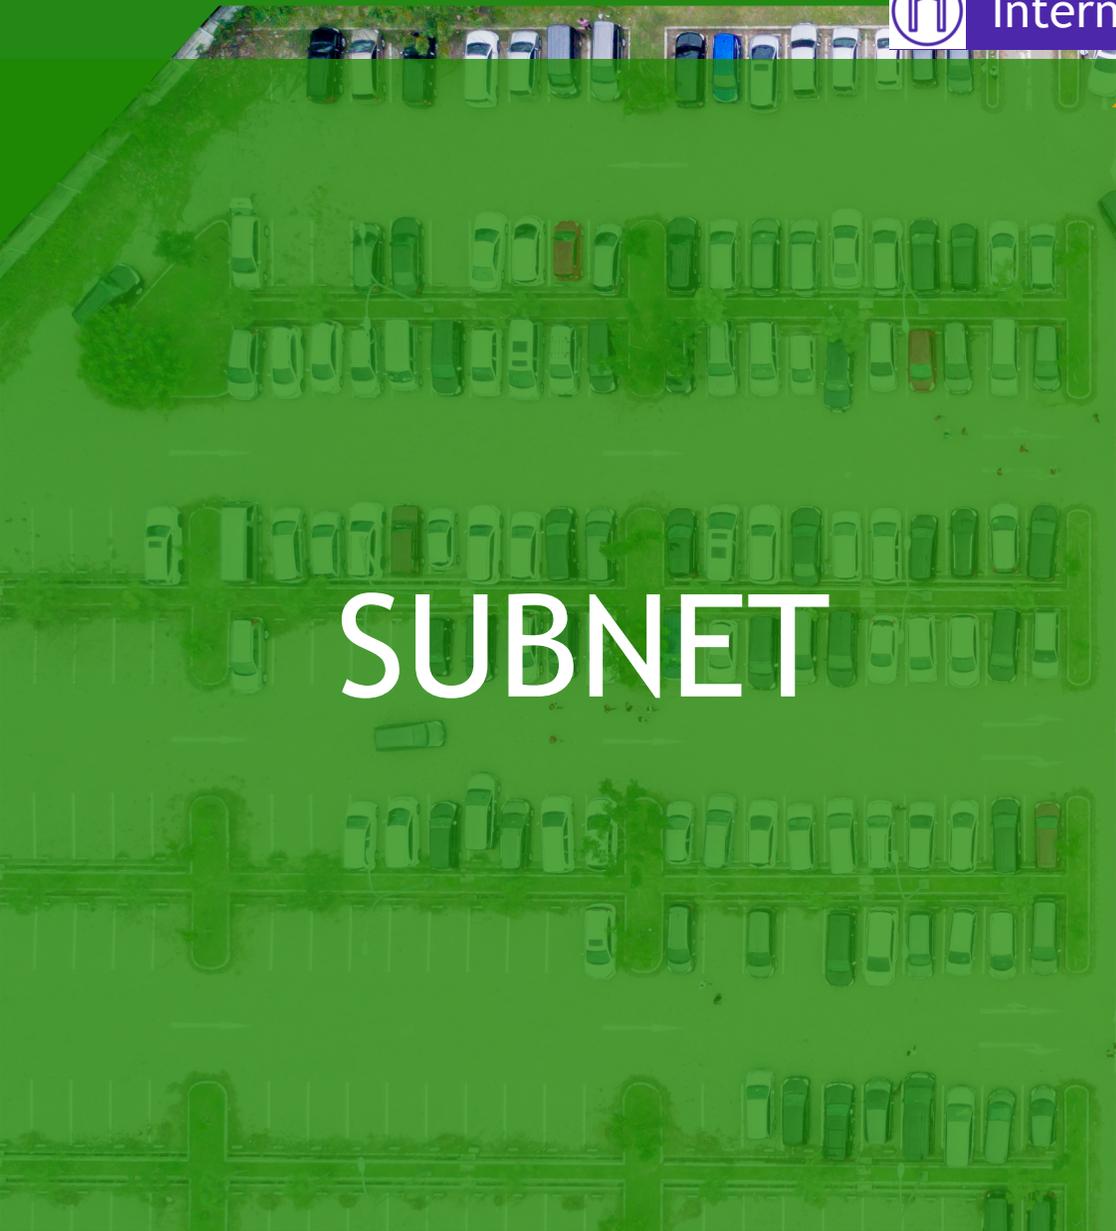
VPC

VPC



Internet Gateway

VPC



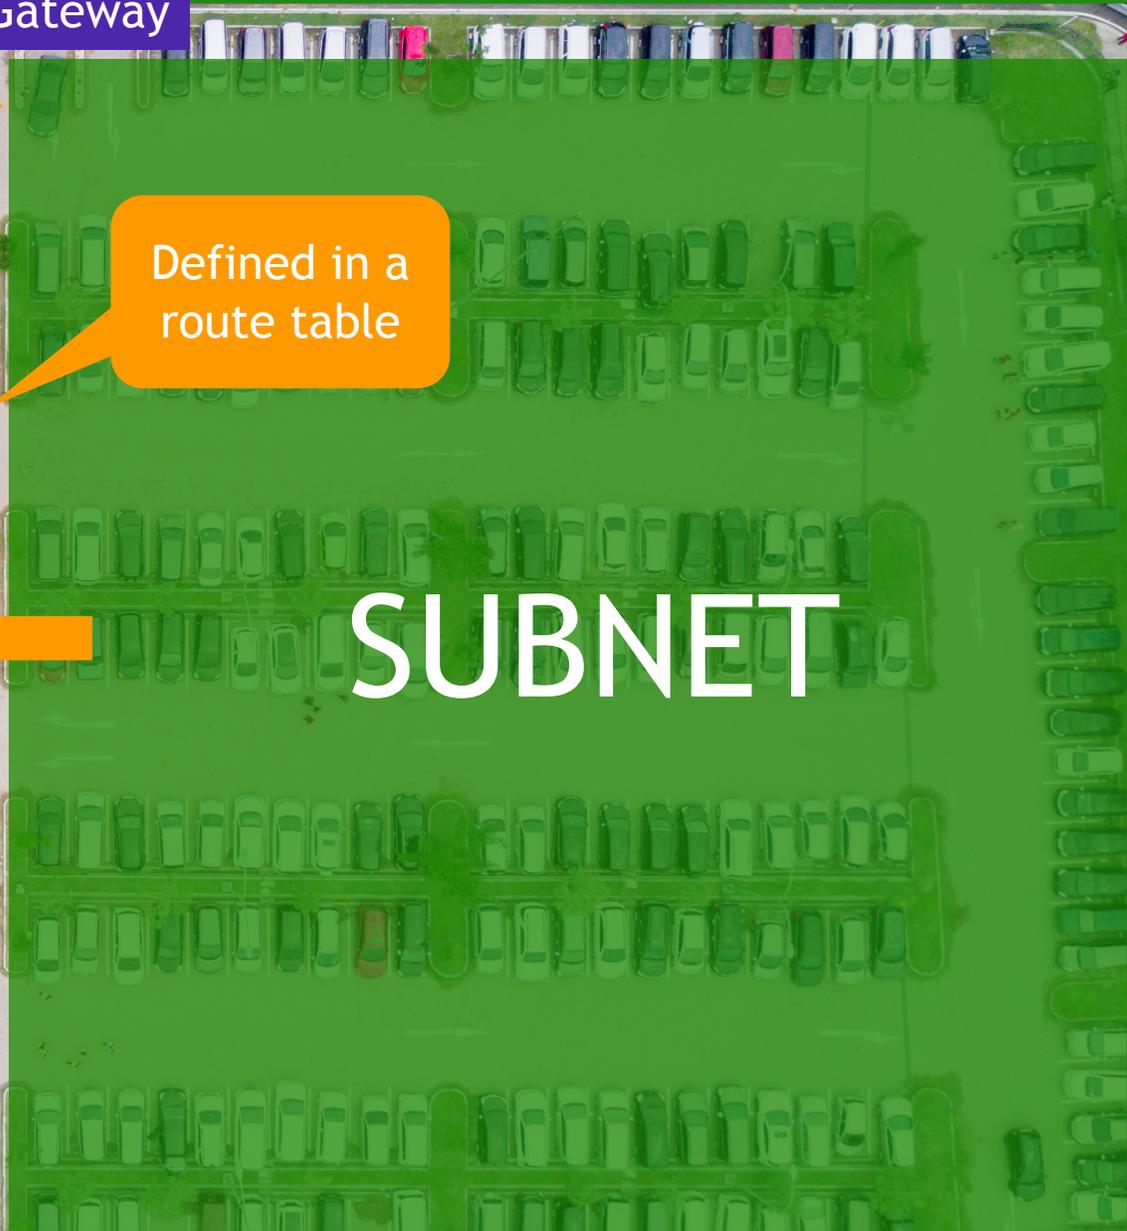
SUBNET

VPC



Route

Defined in a route table



SUBNET

VPC

VPC

VPC



Internet Gateway

VPC

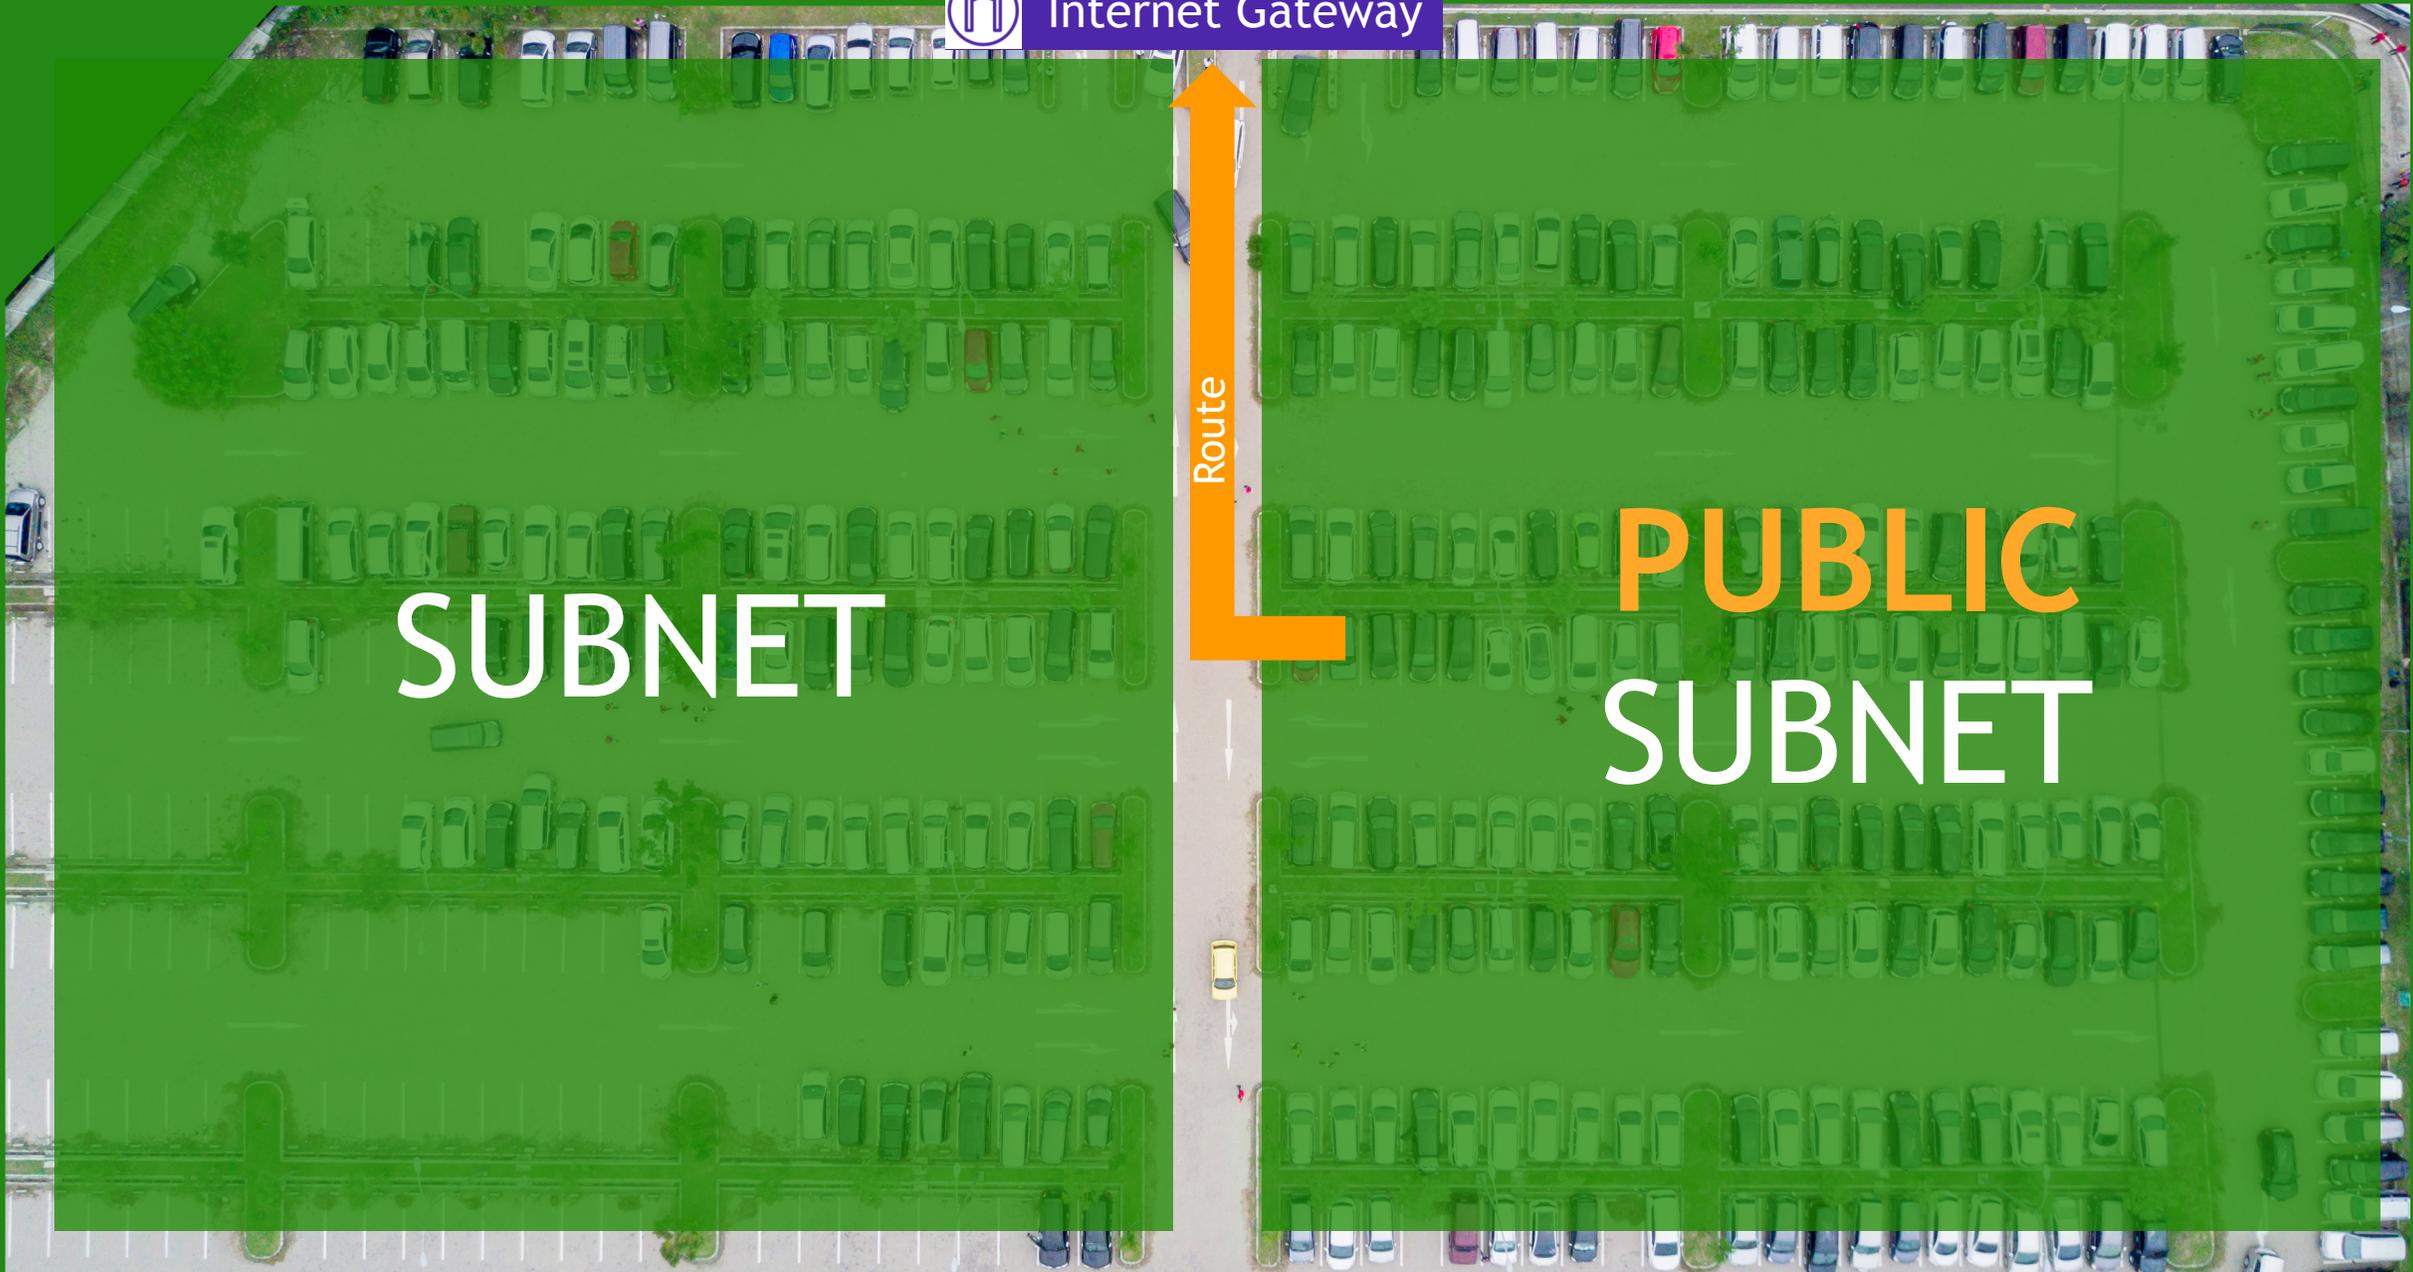
SUBNET

Route

PUBLIC
SUBNET

VPC

VPC



VPC

VPC



Internet Gateway

VPC

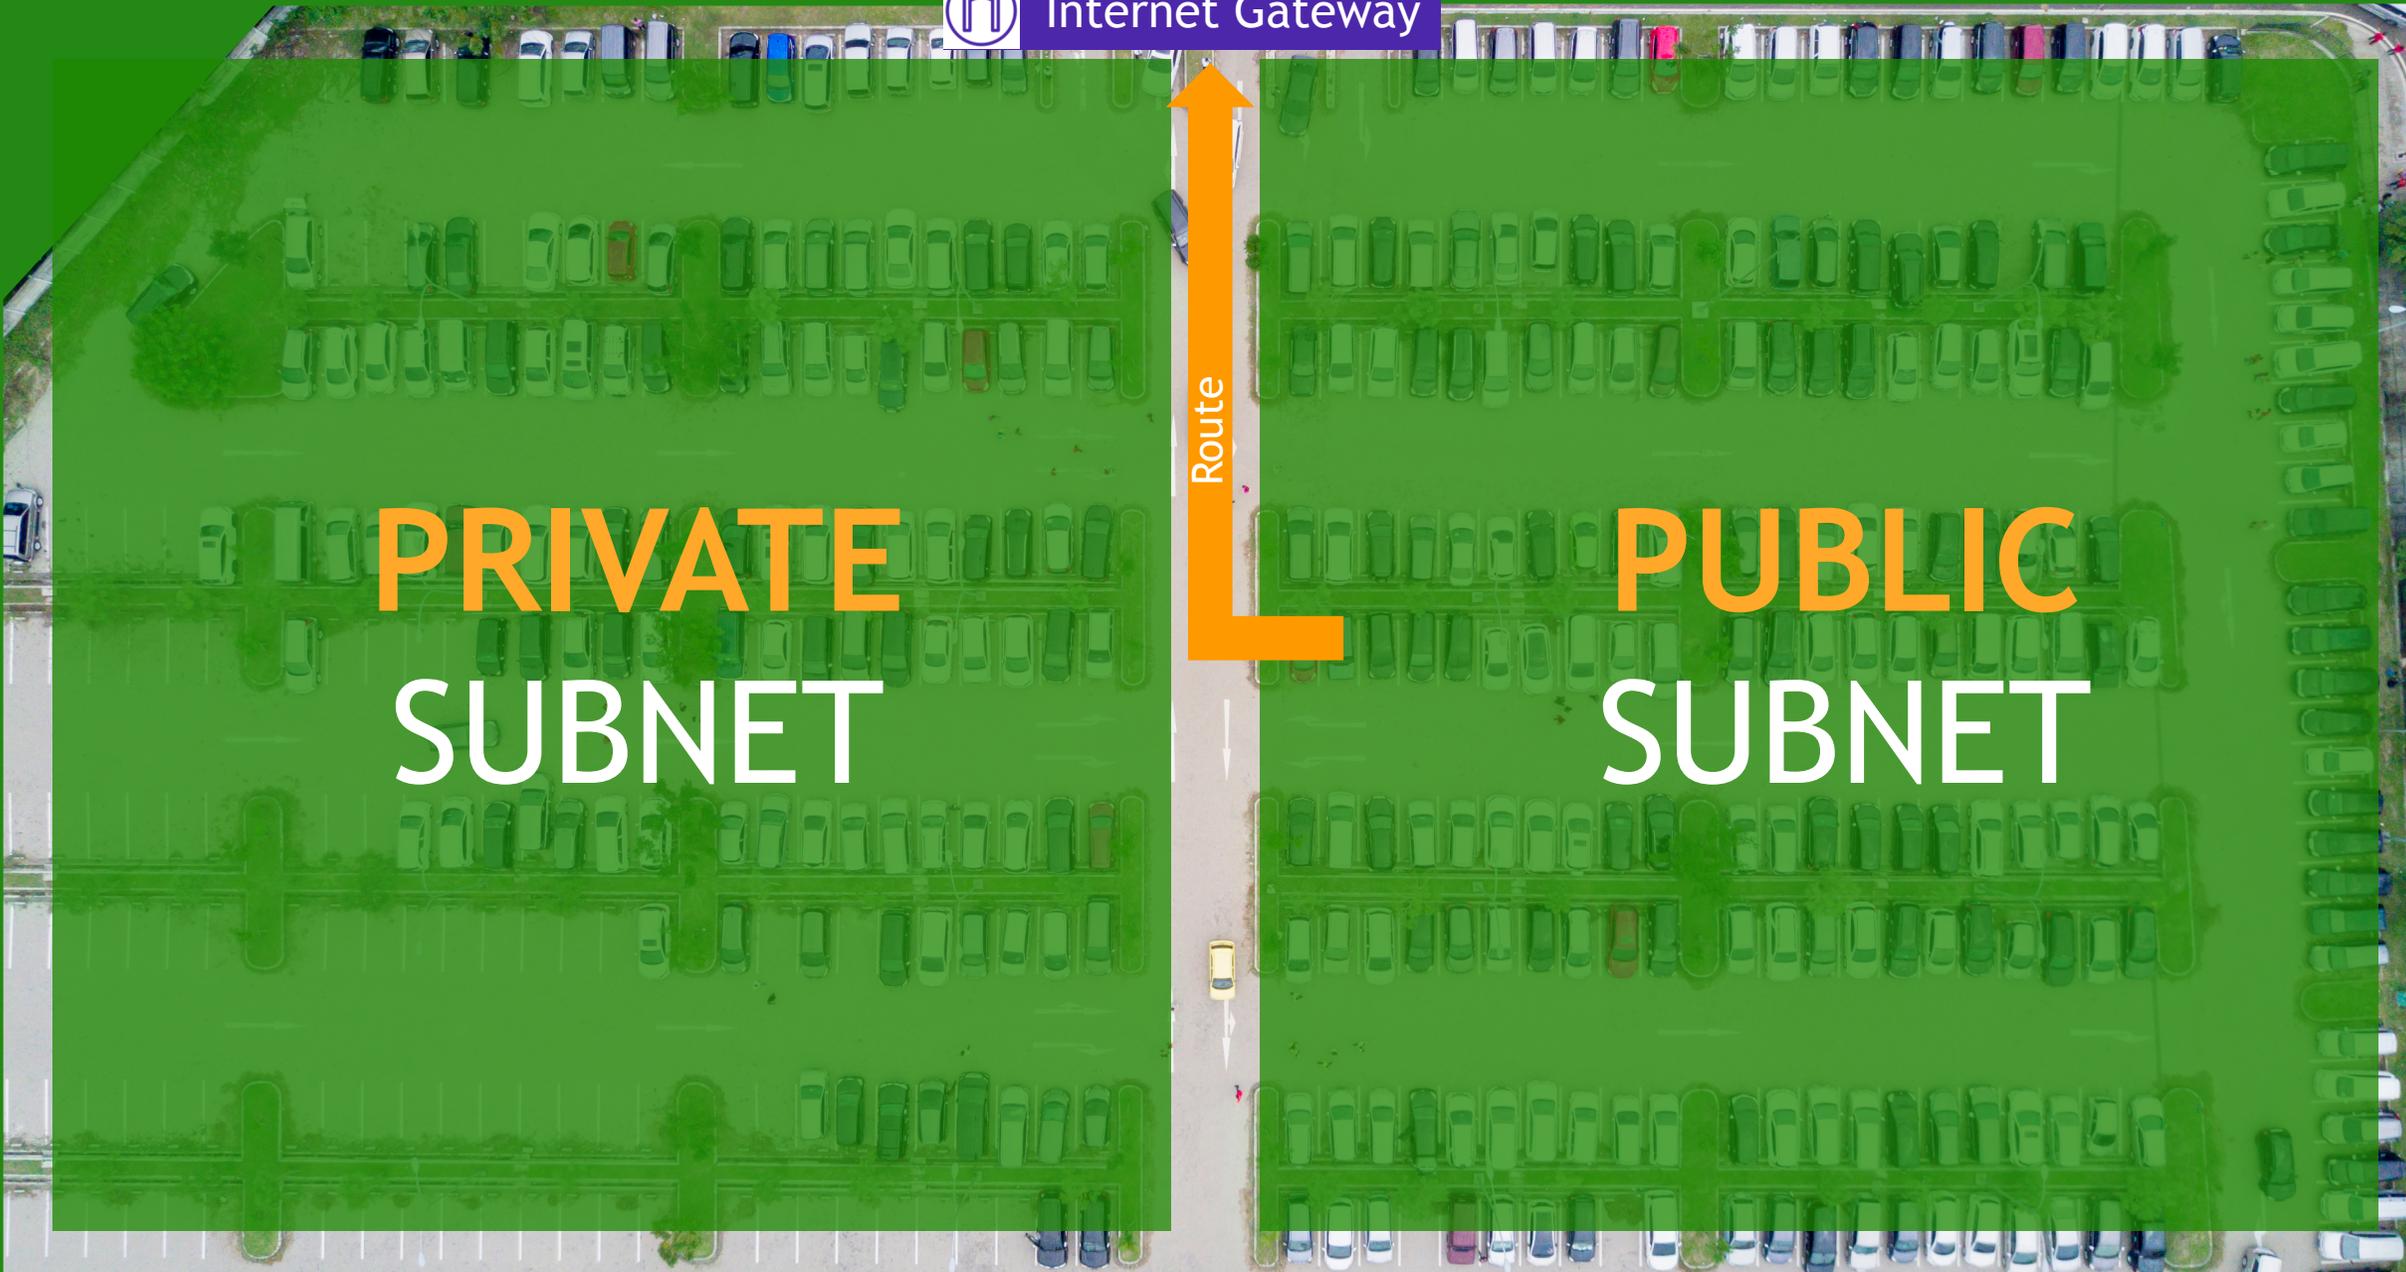
PRIVATE
SUBNET

Route

PUBLIC
SUBNET

VPC

VPC





Region

Availability Zone 1

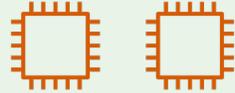
Availability Zone 2



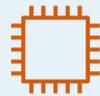
VPC (regional, spans AZs)



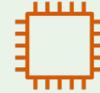
Public subnet
(tied to AZ, accessible from internet)



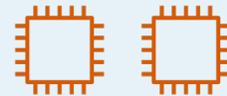
Private subnet
(tied to AZ, NOT accessible from internet)



Public subnet
(tied to AZ, accessible from internet)



Private subnet
(tied to AZ, NOT accessible from internet)





Classless Inter-Domain Routing (CIDR)

Notation for describing blocks of IP addresses



Region

Availability Zone 1

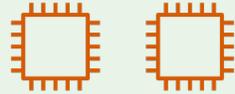
Availability Zone 2



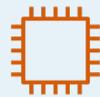
VPC (regional, spans AZs)



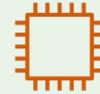
Public subnet
(tied to AZ, accessible from internet)



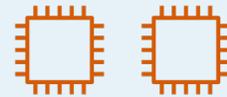
Private subnet
(tied to AZ, NOT accessible from internet)



Public subnet
(tied to AZ, accessible from internet)



Private subnet
(tied to AZ, NOT accessible from internet)



VPC CIDR RANGE: 10.0.0.0/16



Region

Availability Zone 1

Availability Zone 2



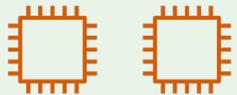
VPC (regional, spans AZs)



Internet Gateway



Public subnet
(tied to AZ, accessible from internet)



NAT Gateway



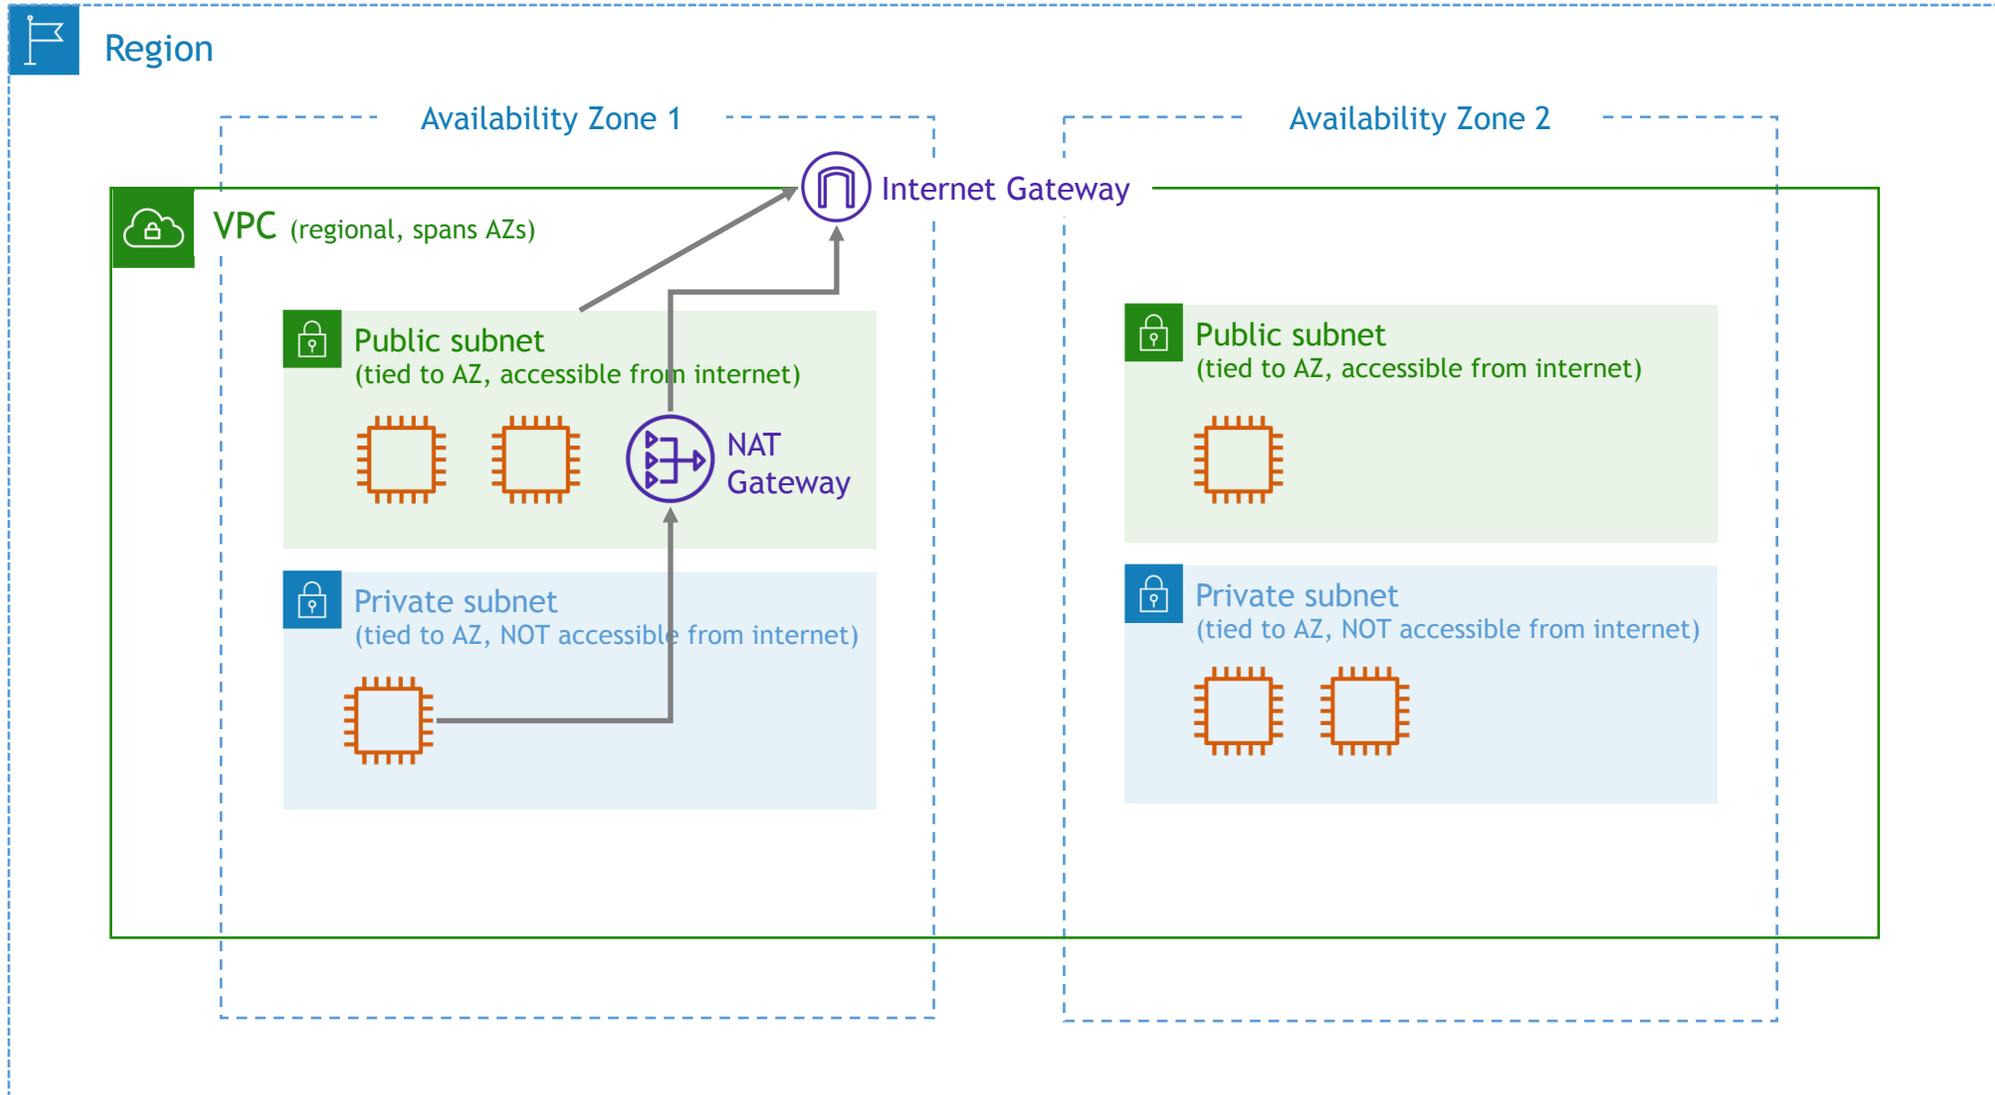
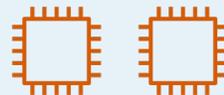
Private subnet
(tied to AZ, NOT accessible from internet)



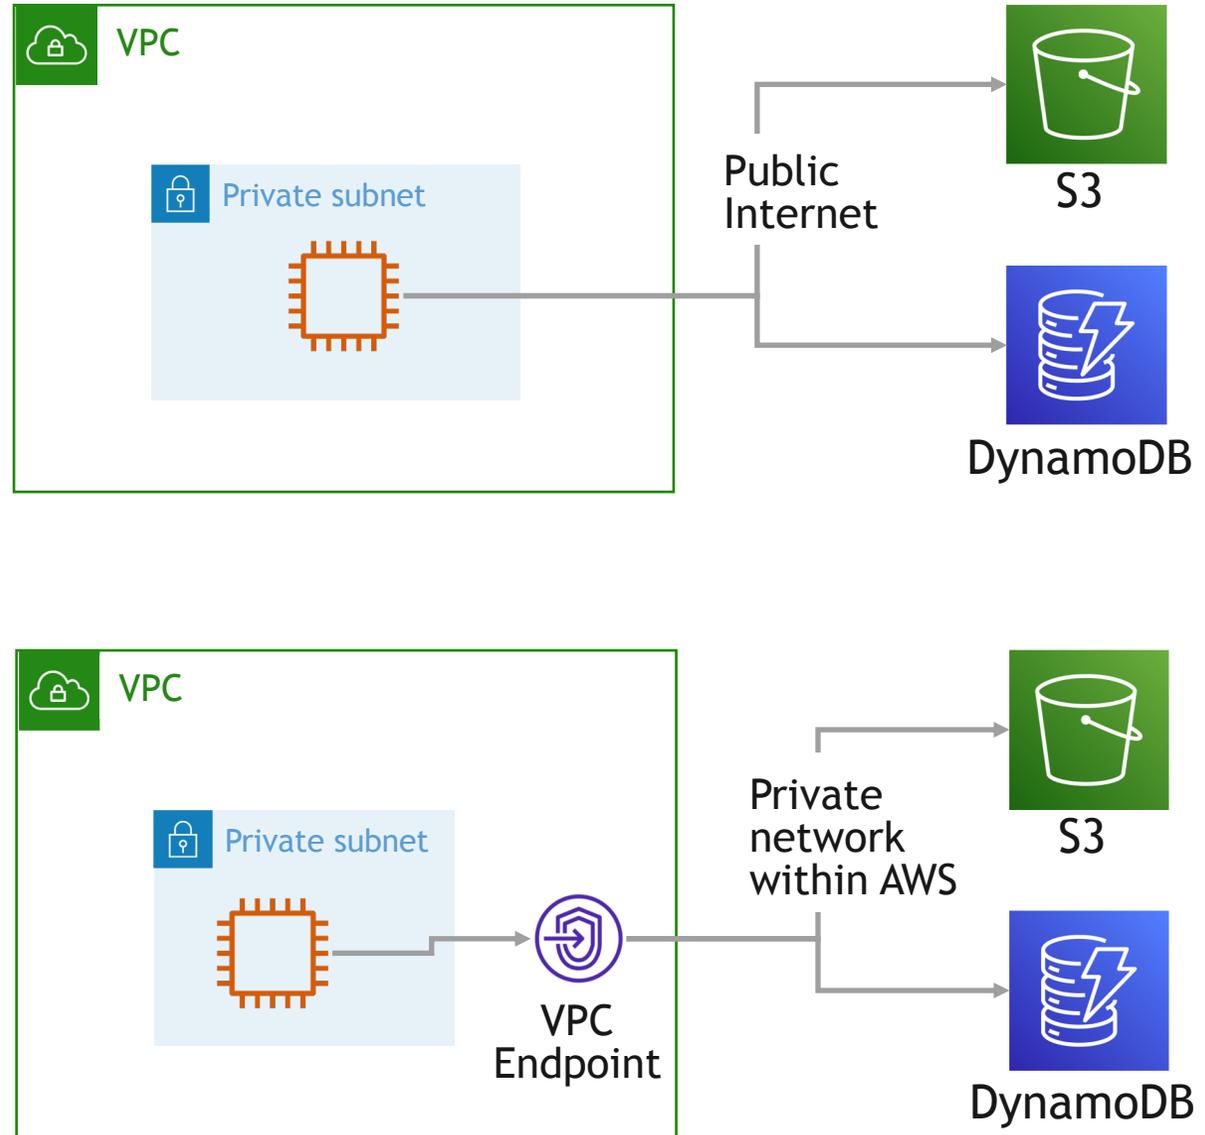
Public subnet
(tied to AZ, accessible from internet)



Private subnet
(tied to AZ, NOT accessible from internet)



VPC Endpoints



Two Types of VPC Endpoints

Gateway

Used to connect to:

S3

DynamoDB

Interface

Used to connect to all other
AWS services

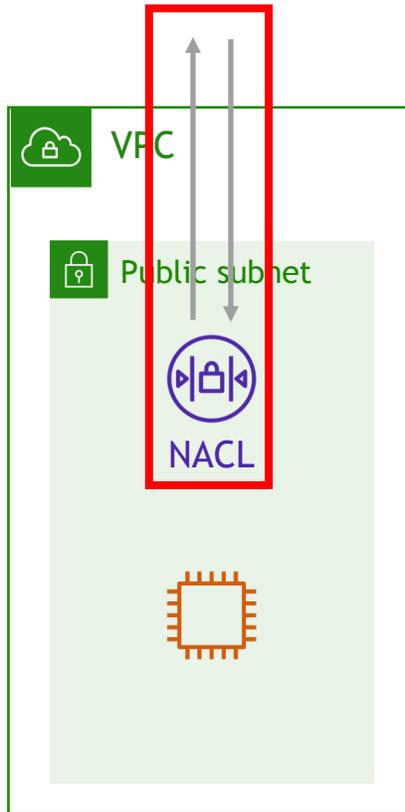
Network ACLs and Security Groups

NETWORK ACLs

Firewall that controls traffic in/out of a *subnet*

Rules for *Allow* and *Deny*

Rules include IP addresses (only)

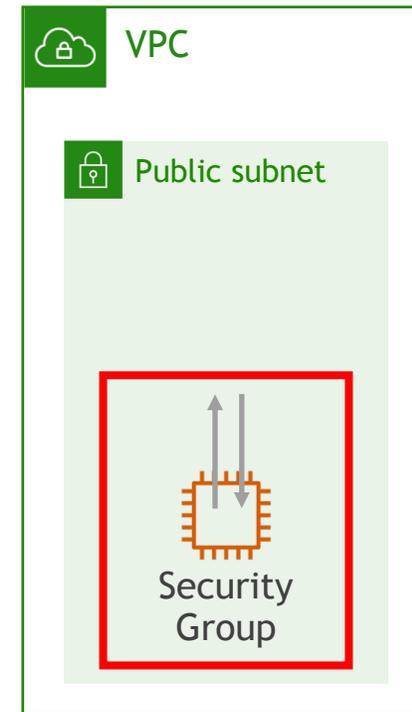


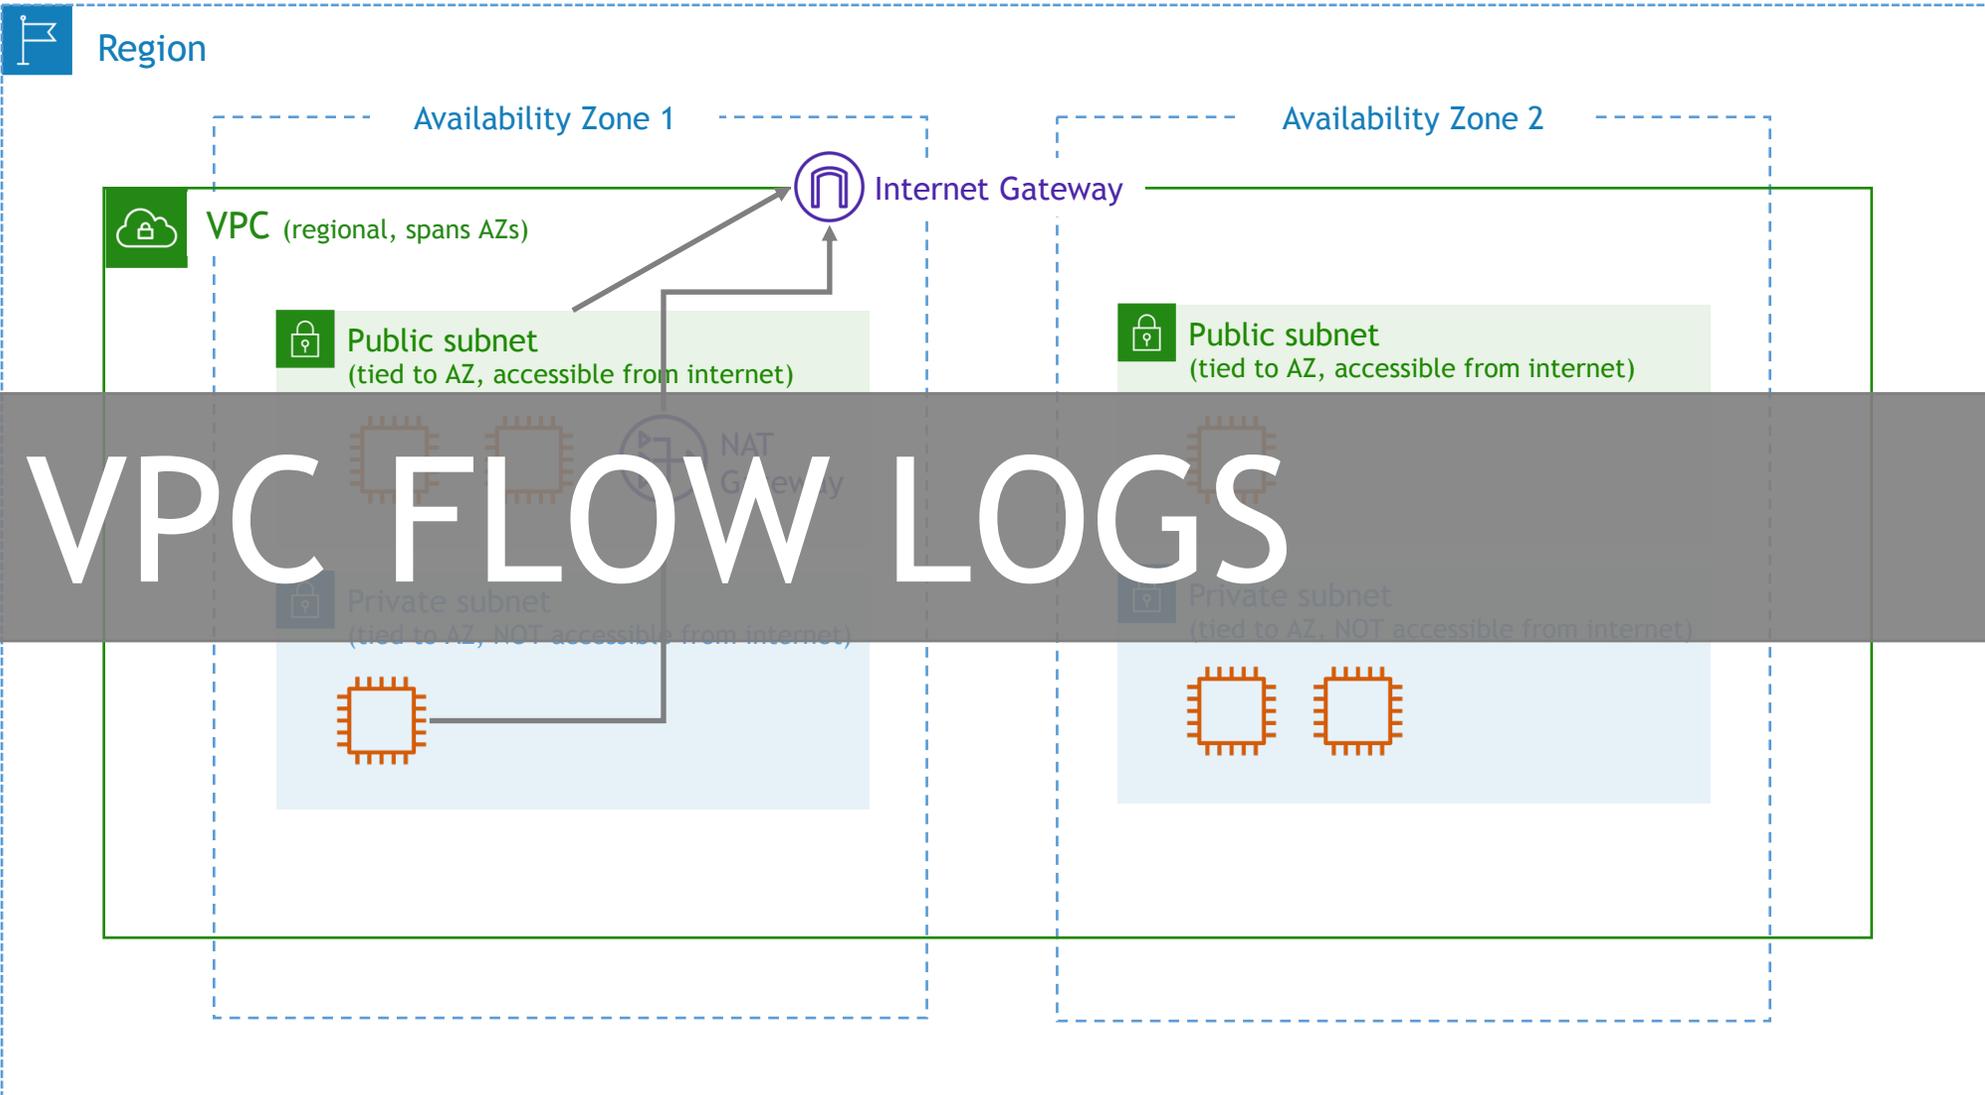
SECURITY GROUPS

Firewall that controls traffic in/out of an *EC2 instance*

Rules for *Allow* (only)

Rules include IP addresses *AND* other security groups





VPC FLOW LOGS



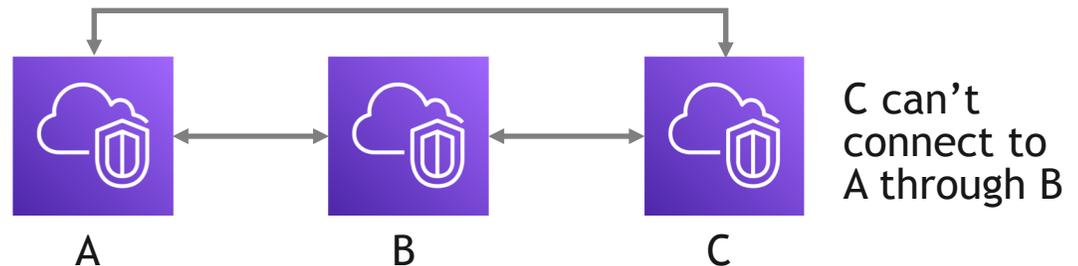
VPC Peering

Peer VPC A with VPC B

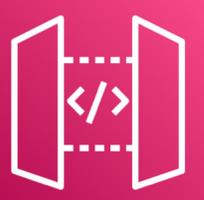
- Behave as a single network

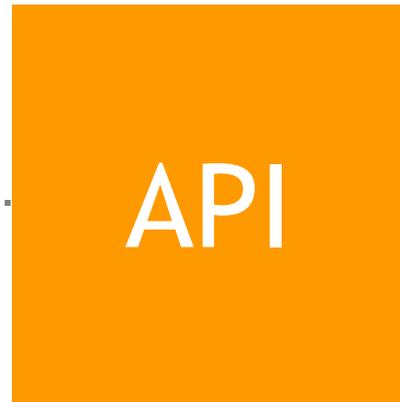
Caveats

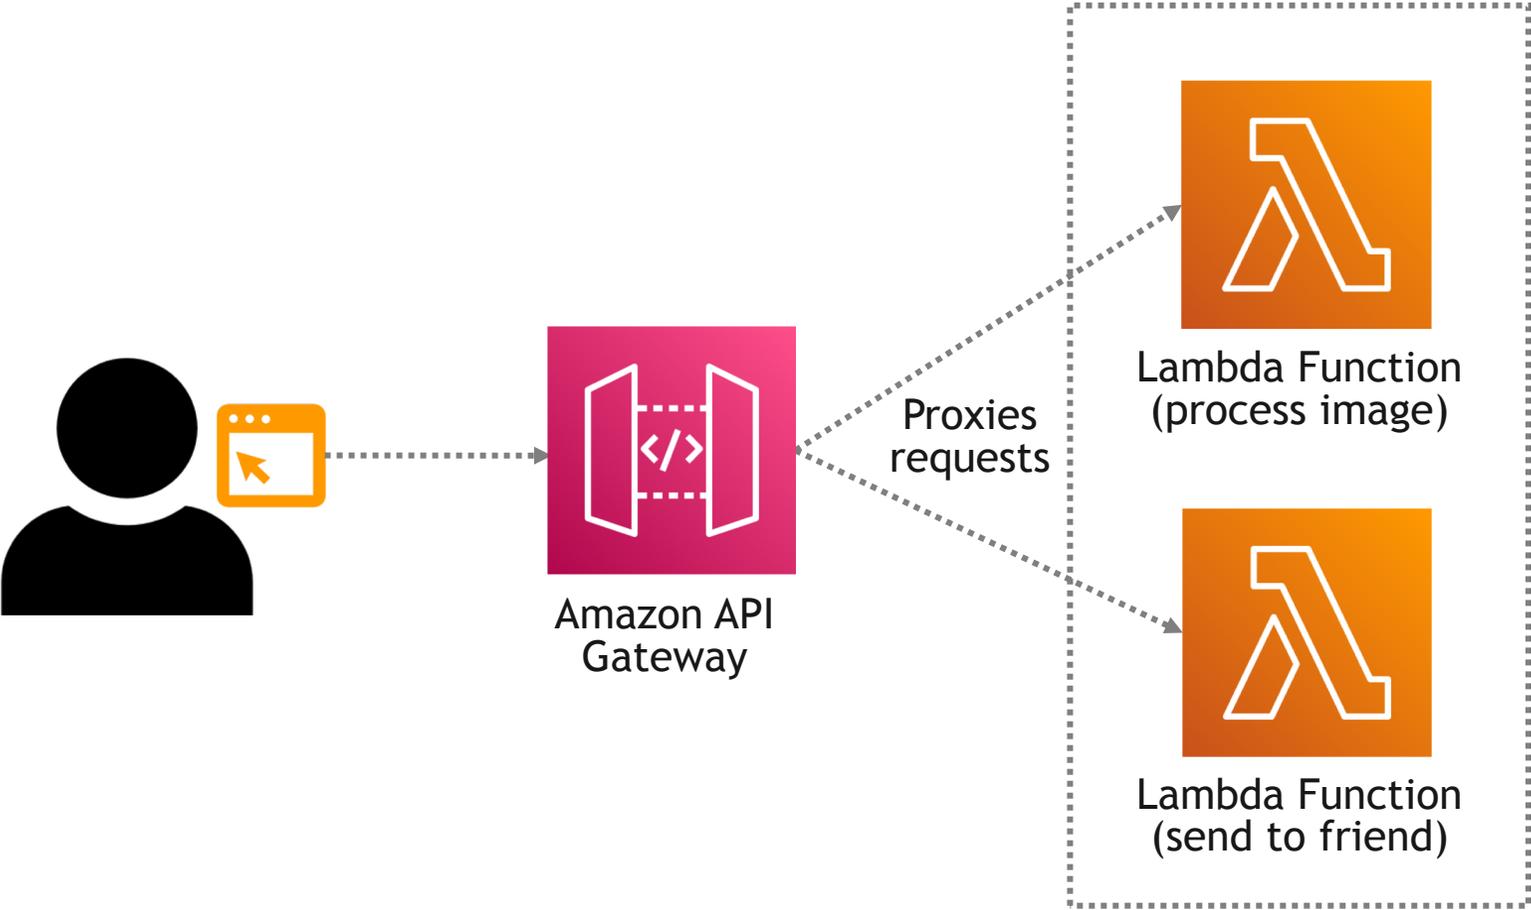
- VPCs can't have overlapping CIDR ranges
- Not transitive (each VPC must directly link to the other)

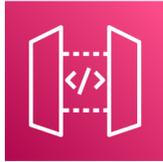




 **AMAZON API GATEWAY**







Amazon API Gateway

Supports REST and WebSocket APIs

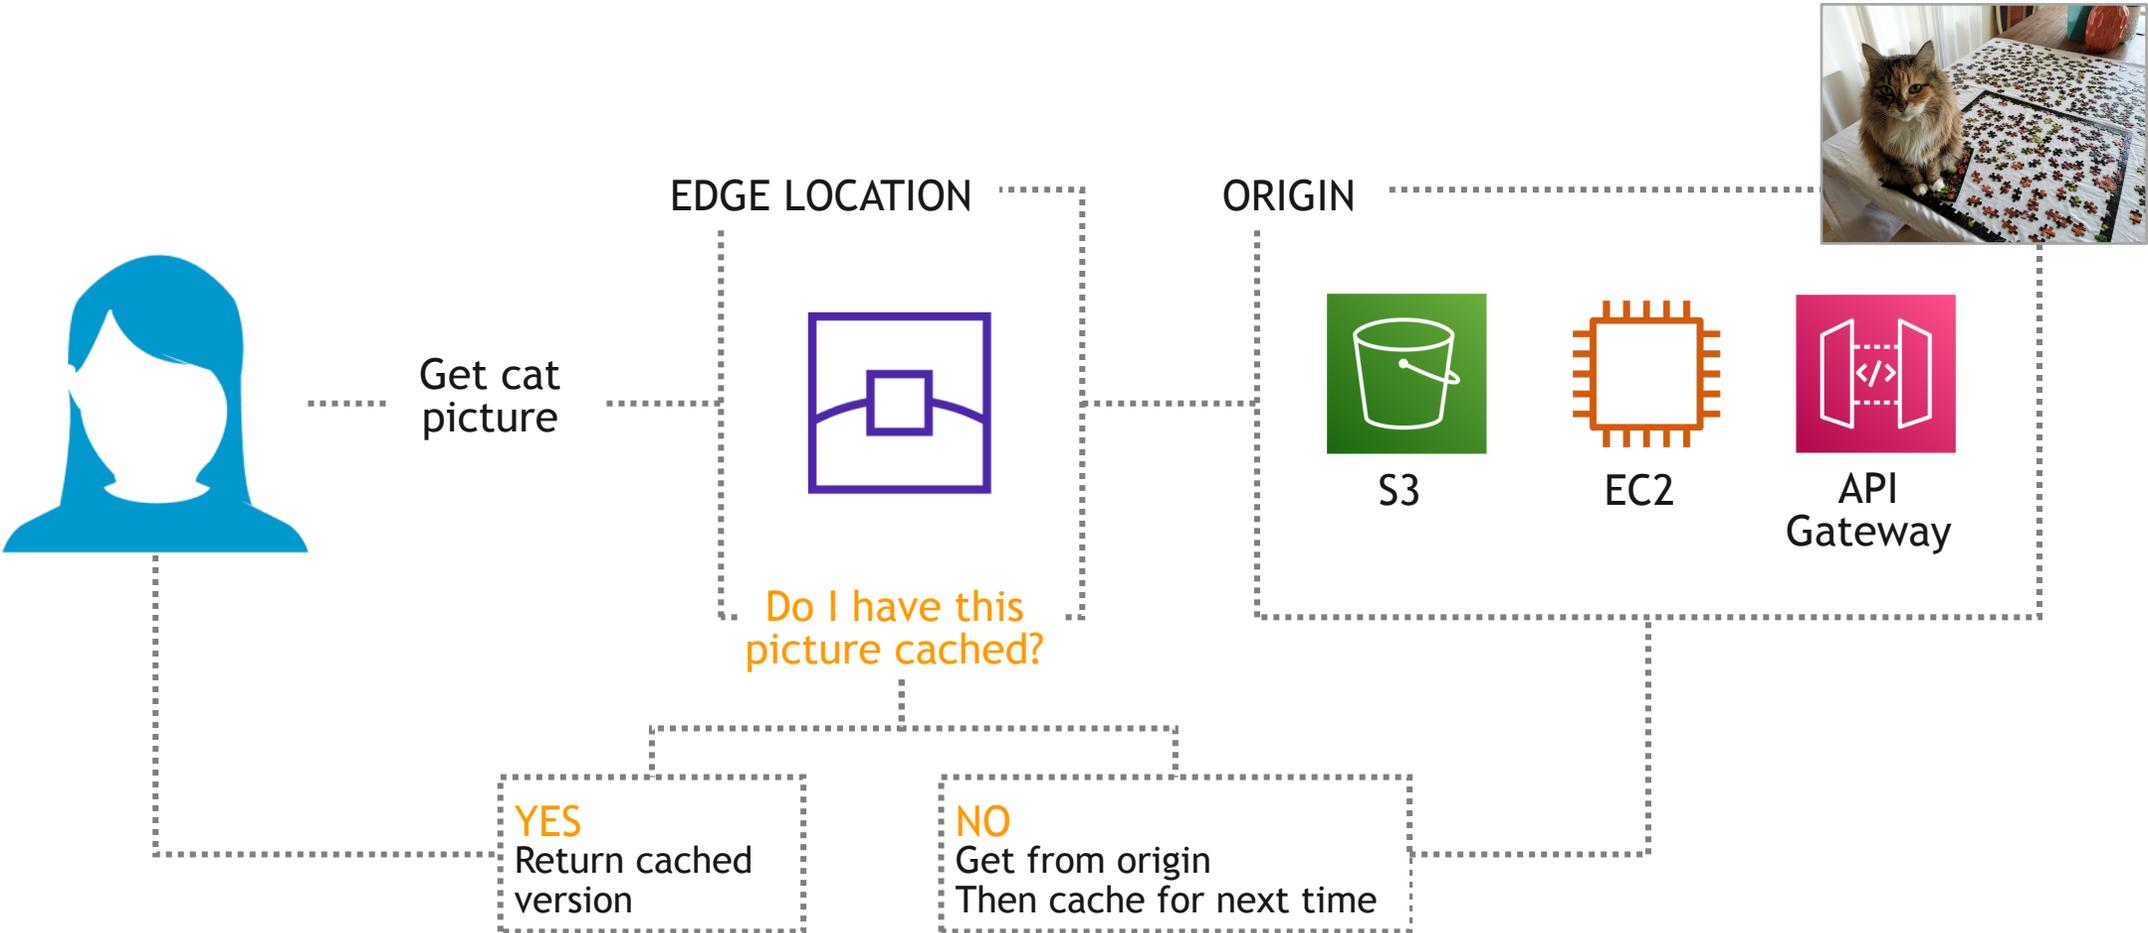
Supports monitoring, authentication, security, throttling and more

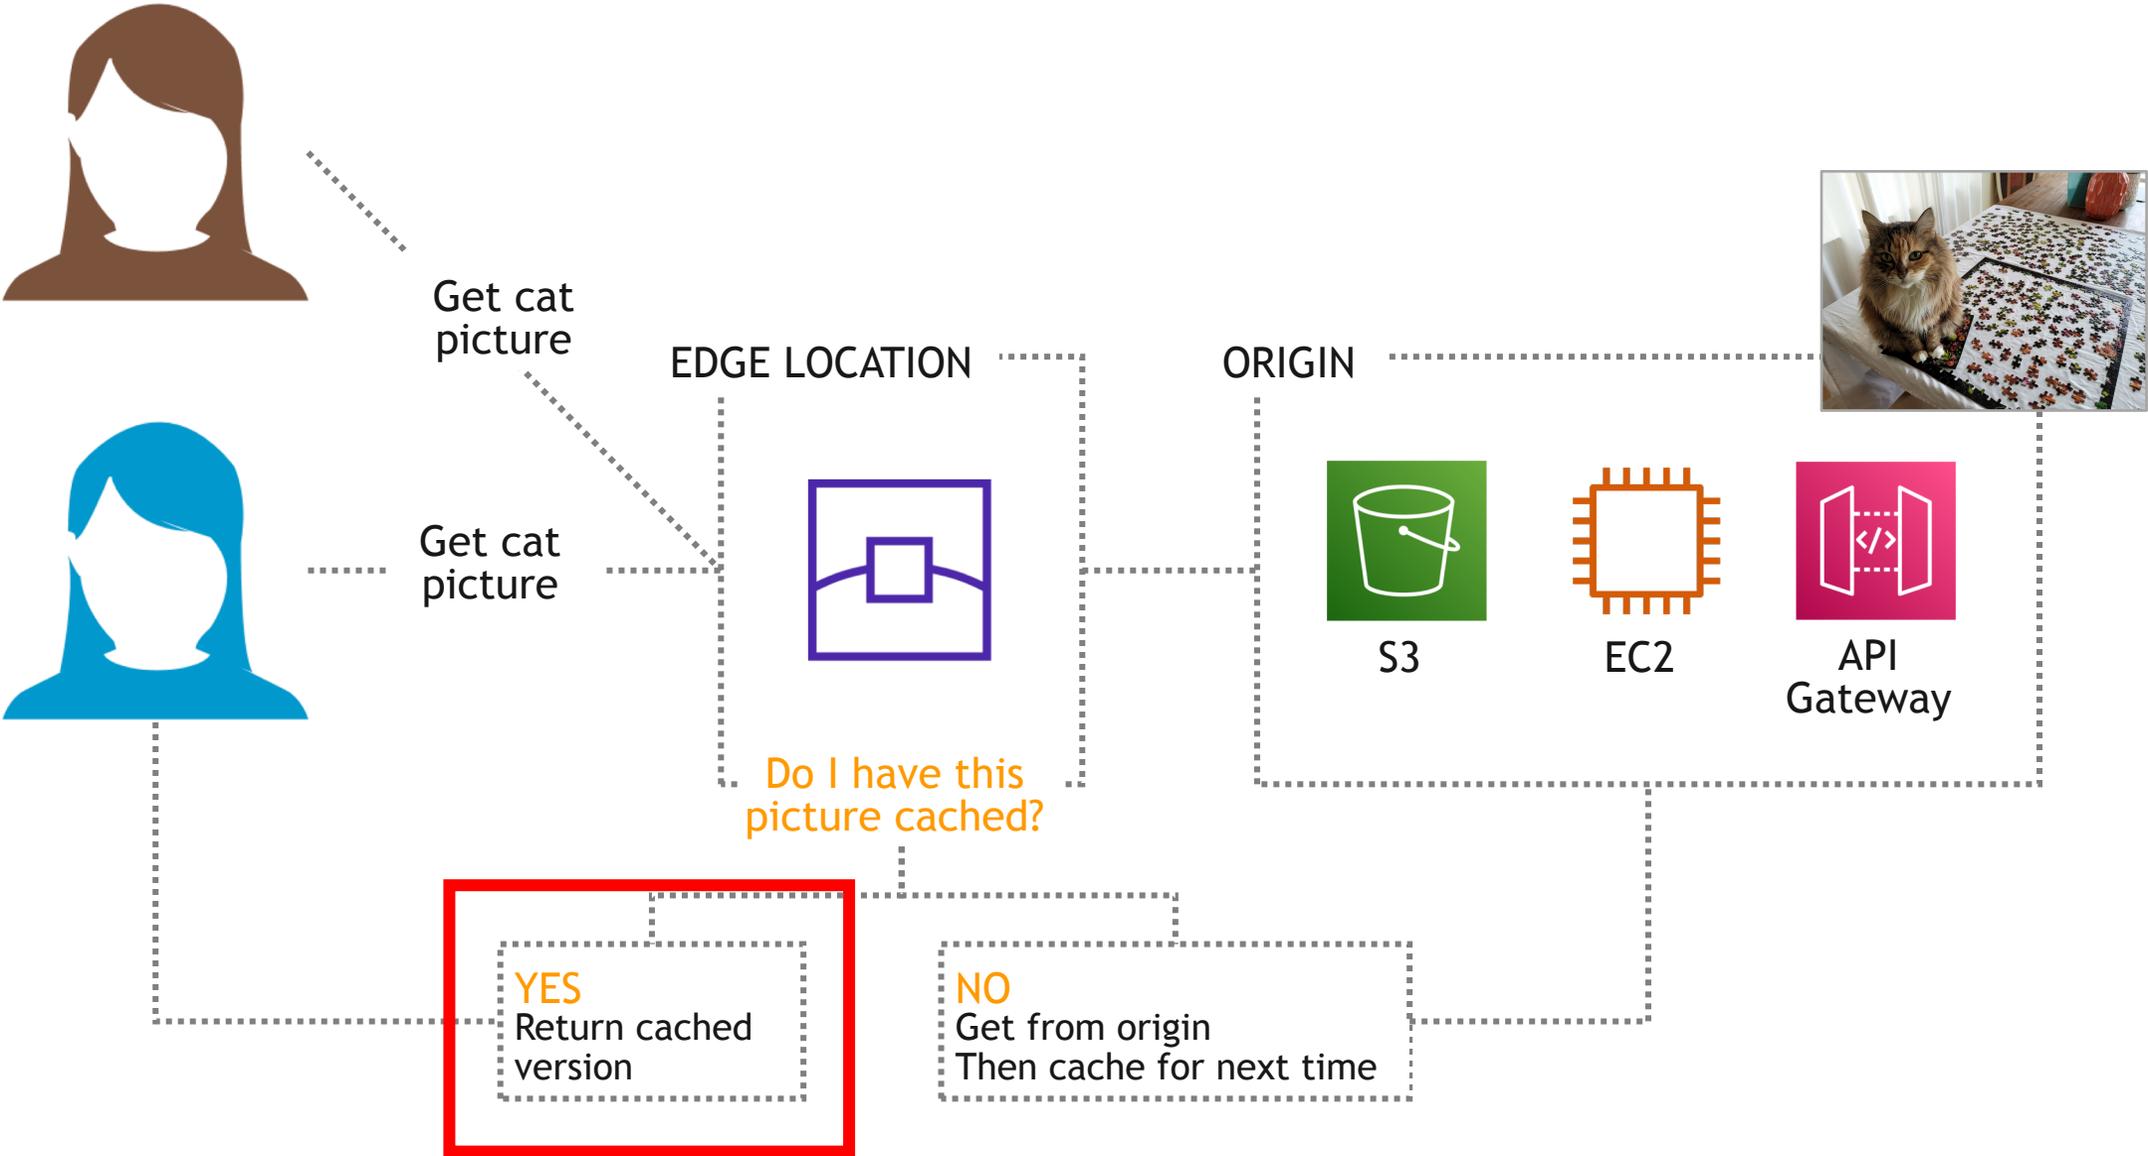


Amazon
CloudFront

Content Delivery Network (CDN)

- Geographically distributed
- Delivers content to users faster by caching

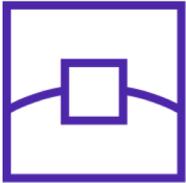






Get cat picture

EDGE LOCATION



ORIGIN



S3



puzzle_cat.jpg
index.html

CloudFront vs. AWS Global Accelerator

CLouDFRONT

Use AWS's global network and edge locations to improve speeds

Content delivery network (e.g., media files)

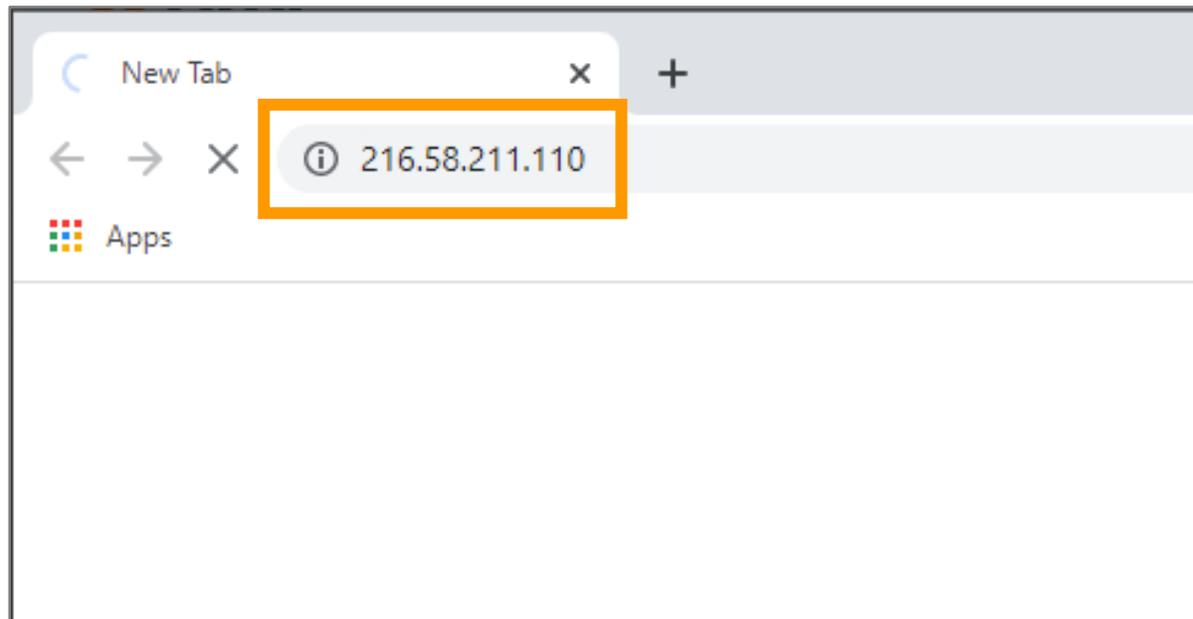
Serves cached content from edge locations

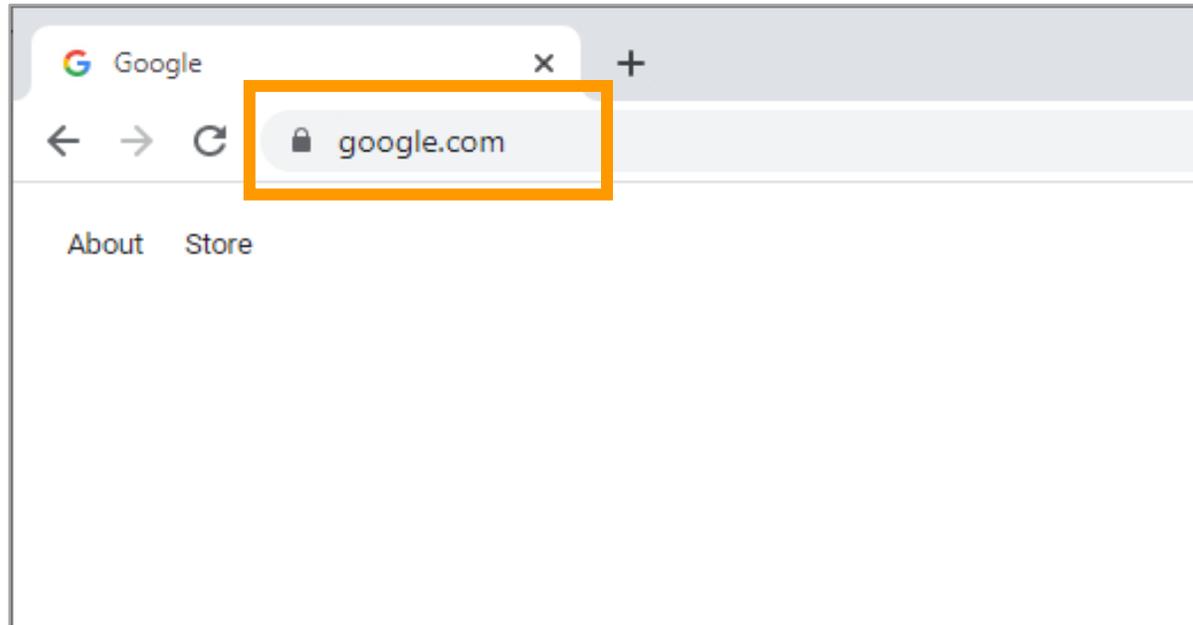
AWS GLOBAL ACCELERATOR

Use AWS's global network and edge locations to improve speeds

Works for a wide variety of applications (HTTP, TCP, UDP)

No caching; improvement comes from network routing of traffic (moving off the public internet)





Domain Name Service (DNS) Service

Maps IP addresses (216.58.211.110) to domain names (google.com)

Similar to a phone book



Amazon Route 53

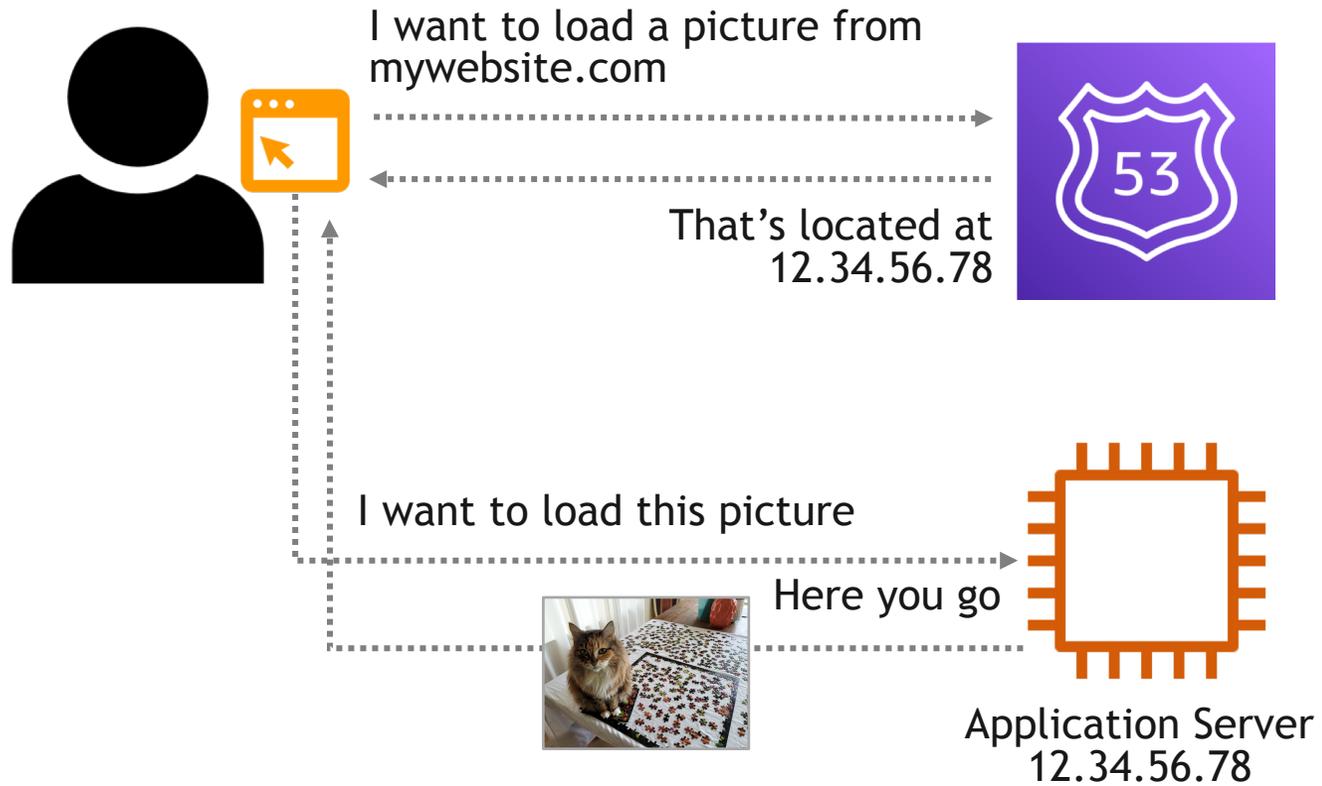
Managed DNS service

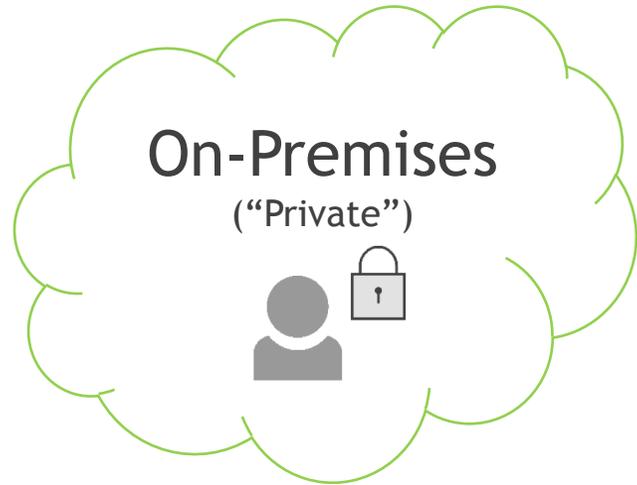
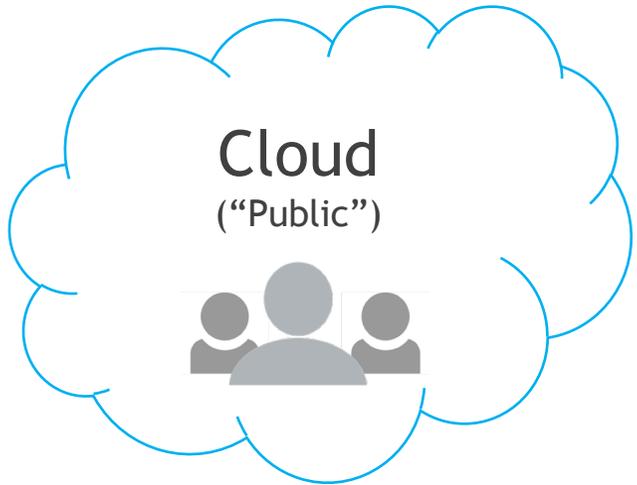
- DNS server requests are addressed to TCP/UDP port 53

Primary functions

- Domain name registration (mywebsite.com)
- DNS routing
- Health checking

DNS Routing

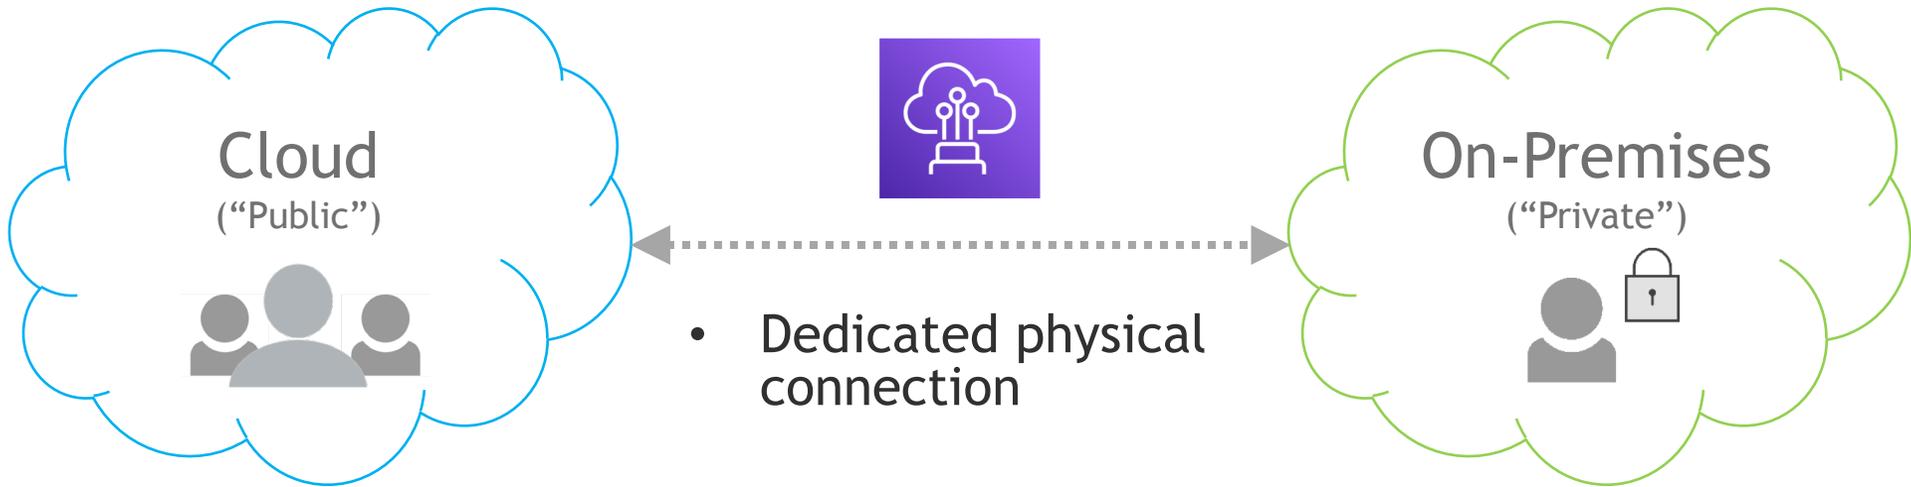




Site-to-Site VPN

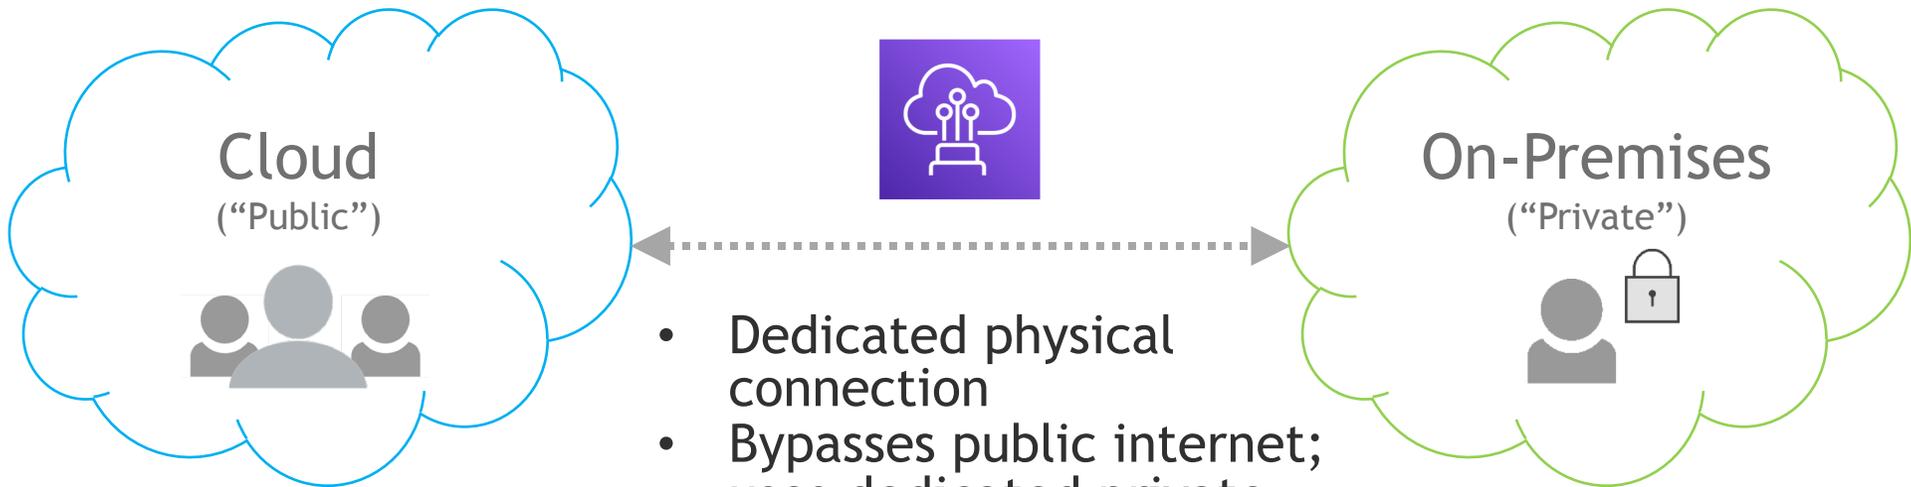


AWS Direct Connect





AWS Direct Connect



- Dedicated physical connection
- Bypasses public internet; uses dedicated private network (not encrypted by default)
- Faster and more expensive
- Takes longer to set up

Networking and Content Delivery: Important Points to Remember

VIRTUAL PRIVATE CLOUD (VPC)

- Your own private cloud/network within the cloud that isolates your resources
- Split into subnets
 - Public: has a route to an Internet Gateway
 - Private: does NOT have a route to an Internet Gateway; if you need to access the internet, use a NAT Gateway
- VPCs are tied to a region and span multiple availability zones
- Network access control lists (ACLs) control traffic at the subnet level
- Security groups control traffic at the EC2 instance level
- VPC endpoints allow you to access other AWS services through a private network (vs. going across the public internet)

Networking and Content Delivery: Important Points to Remember

CLLOUDFRONT

- Content delivery network (CDN) that's geographically distributed to delivery content faster by caching
- Commonly used to deliver media files (videos, images)

GLOBAL ACCELERATOR

- Uses global network of edge locations to improve speeds for a variety of applications
- Does not use caching; improvement comes from network routing of traffic (moving off the public internet)

ROUTE 53

- Managed domain name service (DNS)
- Used to map IP addresses (216.58.211.110) to domain names (google.com)
- Can register domain names (mywebsite.com)

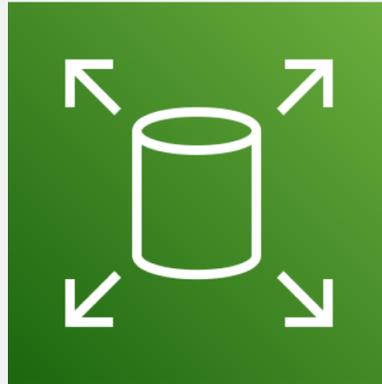


DATABASES IN AWS

So Many Ways to Store Data!



S3



EBS



EFS

Retrieve all orders for
this customer for the
last two years

Purpose-Built Databases in AWS

DATABASE TYPE	USE CASE	AWS SERVICE
Relational	<ul style="list-style-type: none">Traditional apps, CRM, eCommerce	<ul style="list-style-type: none">Amazon RDS, Aurora, Redshift
Non-Relational or NoSQL (key-value)	<ul style="list-style-type: none">High-traffic web apps, eCommerce, Gaming	<ul style="list-style-type: none">Amazon DynamoDB
In-memory	<ul style="list-style-type: none">Caching, session management	<ul style="list-style-type: none">Amazon ElastiCache
Graph	<ul style="list-style-type: none">Fraud detection, social networking, recommendation engines	<ul style="list-style-type: none">Amazon Neptune

Managed Database Service

A database where Amazon does all the underlying administrative work (backups, patching, recovery, etc.)

vs.

You installing the database directly on a server and managing it

Purpose-Built Databases in AWS

DATABASE TYPE	USE CASE	AWS SERVICE
Relational	<ul style="list-style-type: none">Traditional apps, CRM, eCommerce	<ul style="list-style-type: none">Amazon RDS, Aurora, Redshift
Non-Relational or NoSQL (key-value)	<ul style="list-style-type: none">High-traffic web apps, eCommerce, Gaming	<ul style="list-style-type: none">Amazon DynamoDB
In-memory	<ul style="list-style-type: none">Caching, session management	<ul style="list-style-type: none">Amazon ElastiCache
Graph	<ul style="list-style-type: none">Fraud detection, social networking, recommendation engines	<ul style="list-style-type: none">Amazon Neptune

Relational Databases

CUSTOMERS

CustomerID	FirstName	LastName	EmailAddress
1	Dwight	Schrute	dwight.schrute@email.com
2	Pam	Beasley	pam.beasley@email.com
3	Jim	Halpert	jim.halpert@email.com
4	Michael	Scott	michael.scott@email.com

ORDERS

OrderID	CustomerID	OrderDate	ShippingAddress
845	1	1/6/22	123 Main Street
846	3	12/30/21	555 Avenue F
847	4	12/28/21	628 Goldenrod
848	2	1/1/22	987 State Street

Structured Query Language (SQL)

Used to query a relational database

```
SELECT FirstName from CUSTOMERS
```

Purpose-Built Databases in AWS

DATABASE TYPE	USE CASE	AWS SERVICE
Relational	<ul style="list-style-type: none">Traditional apps, CRM, eCommerce	<ul style="list-style-type: none">Amazon RDS, Aurora, Redshift
Non-Relational or NoSQL (key-value)	<ul style="list-style-type: none">High-traffic web apps, eCommerce, Gaming	<ul style="list-style-type: none">Amazon DynamoDB
In-memory	<ul style="list-style-type: none">Caching, session management	<ul style="list-style-type: none">Amazon ElastiCache
Graph	<ul style="list-style-type: none">Fraud detection, social networking, recommendation engines	<ul style="list-style-type: none">Amazon Neptune

Engine type [Info](#)

Amazon Aurora



MySQL



MariaDB



PostgreSQL



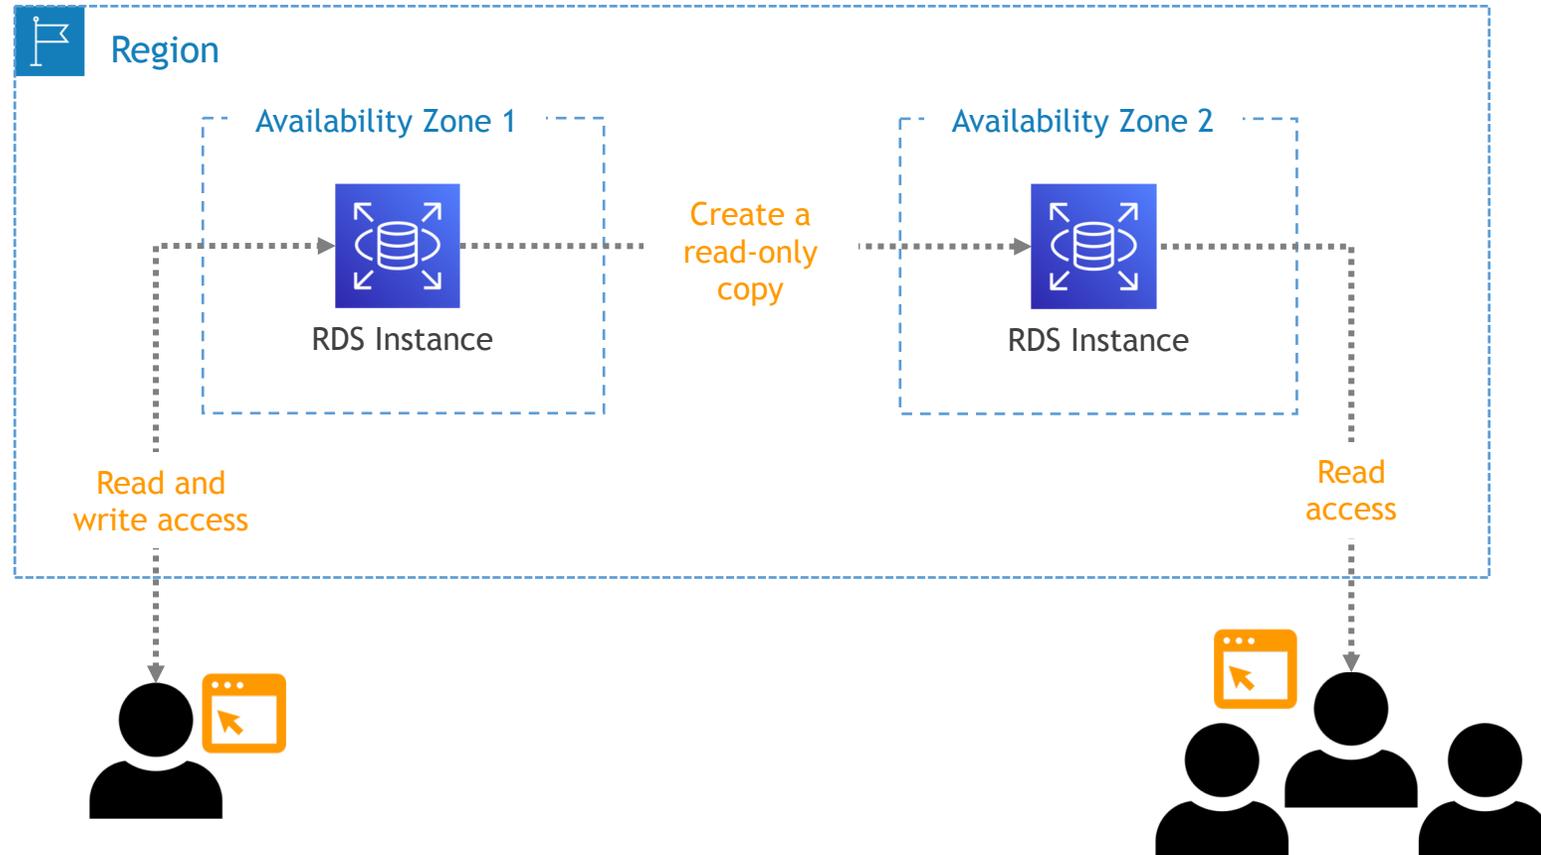
Oracle

ORACLE®

Microsoft SQL Server

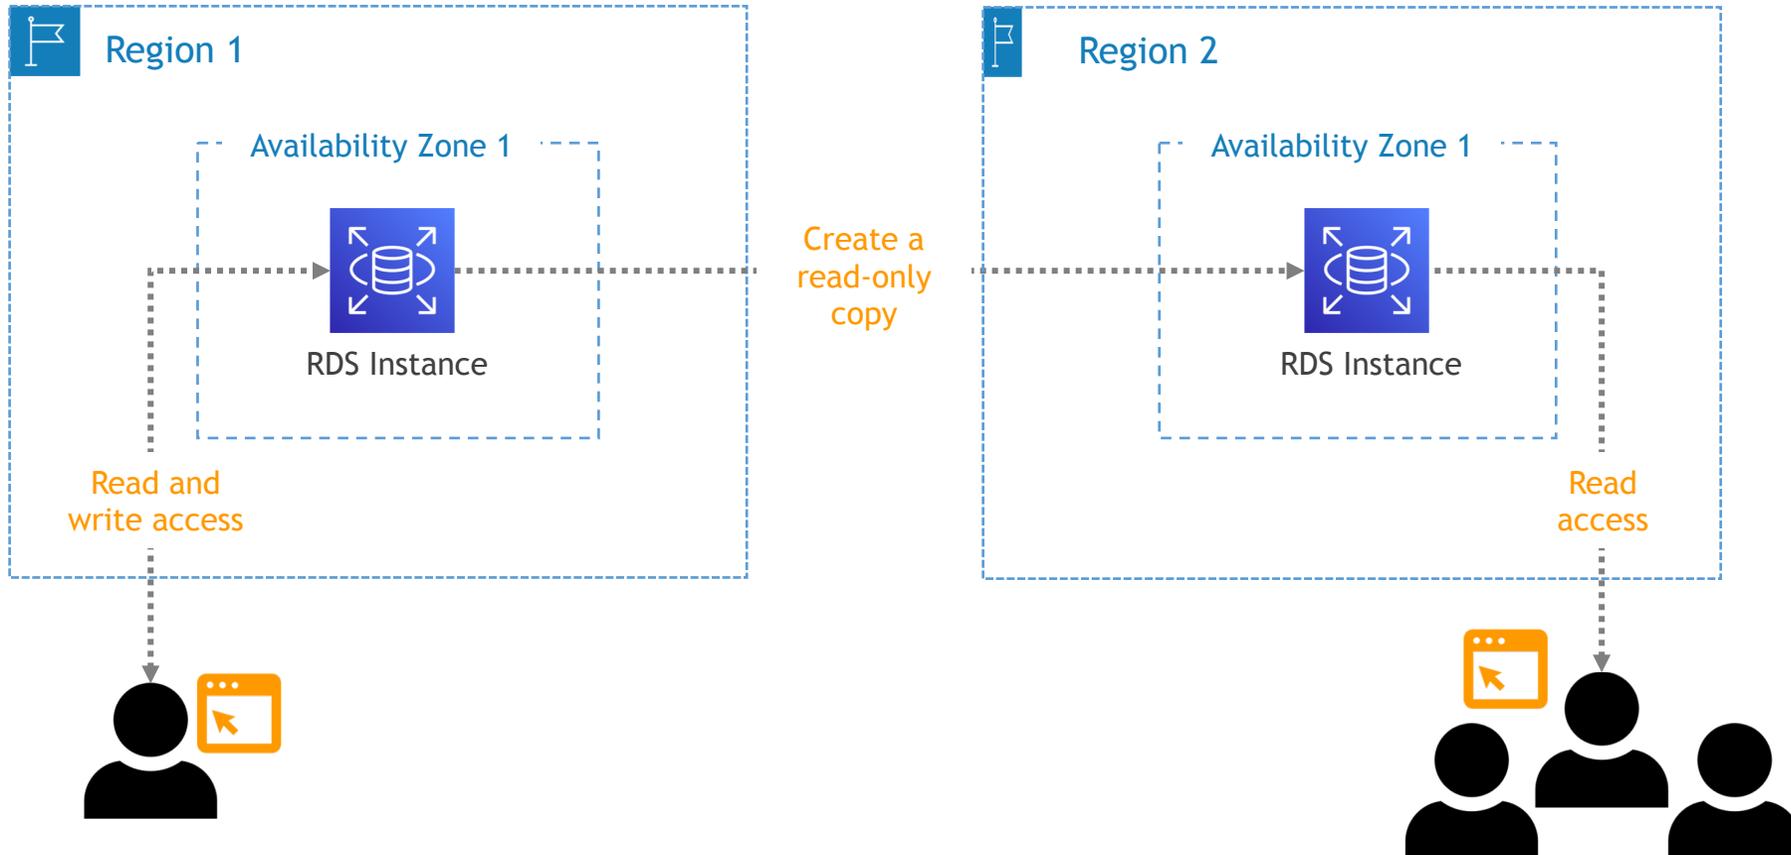


Read Replicas



But wait,
there's more!

CROSS-REGION Read Replicas



Shared Responsibility Model for RDS

AWS Responsibility

Underlying infrastructure
(instance and database engine)

Maintenance and patching

Customer (Your) Responsibility

Network access (through
security groups)

Setting up and managing
database users and permissions



DEMO

Amazon RDS:

Creating and Connecting to a *MySQL* Database



Data warehousing service

- Can handle massive amount of data (exabytes), structured and unstructured

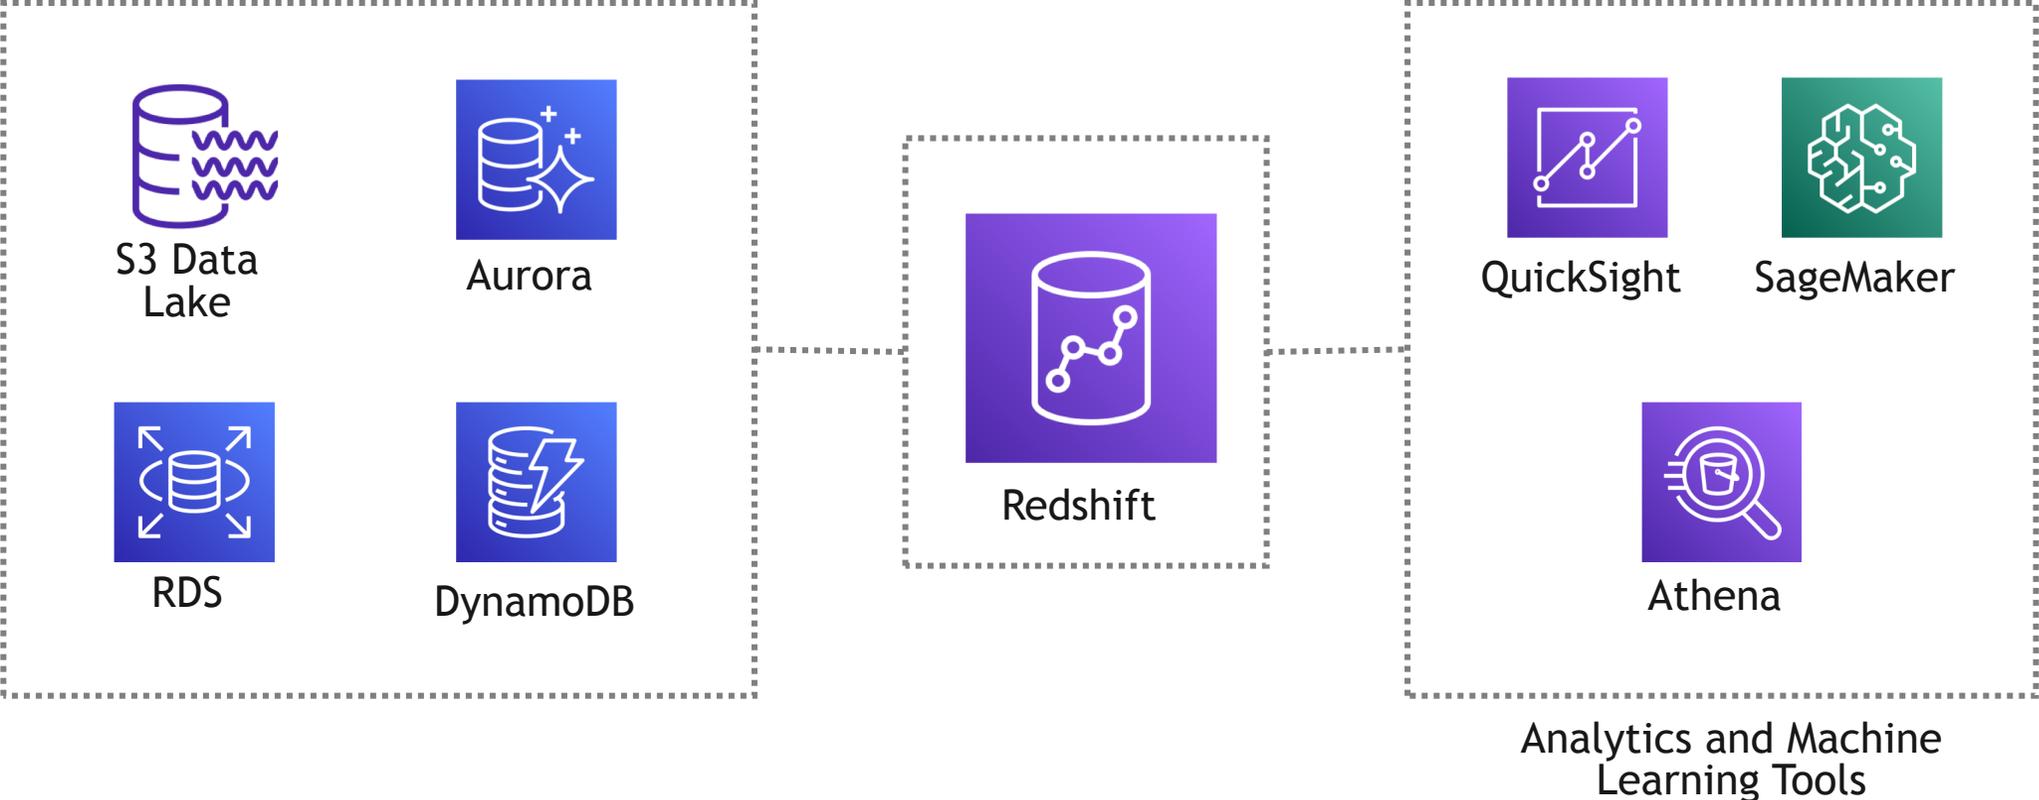
OLAP-style (Online Analytical Processing) columnar database based on PostgreSQL

- Can use SQL queries

Use cases

- Real-time analytics
- Log analysis
- Combining multiple data sources

Amazon Redshift





Data warehousing service

- Can handle massive amount of data (exabytes), structured and unstructured

OLAP-style (Online Analytical Processing) columnar database based on PostgreSQL

- Can use SQL queries

Use cases

- Real-time analytics
- Log analysis
- Combining multiple data sources

Purpose-Built Databases in AWS

DATABASE TYPE	USE CASE	AWS SERVICE
Relational	<ul style="list-style-type: none">Traditional apps, CRM, eCommerce	<ul style="list-style-type: none">Amazon RDS, Aurora, Redshift
Non-Relational or NoSQL (key-value)	<ul style="list-style-type: none">High-traffic web apps, eCommerce, Gaming	<ul style="list-style-type: none">Amazon DynamoDB
In-memory	<ul style="list-style-type: none">Caching, session management	<ul style="list-style-type: none">Amazon ElastiCache
Graph	<ul style="list-style-type: none">Fraud detection, social networking, recommendation engines	<ul style="list-style-type: none">Amazon Neptune

Non-Relational/NoSQL Databases

JSON

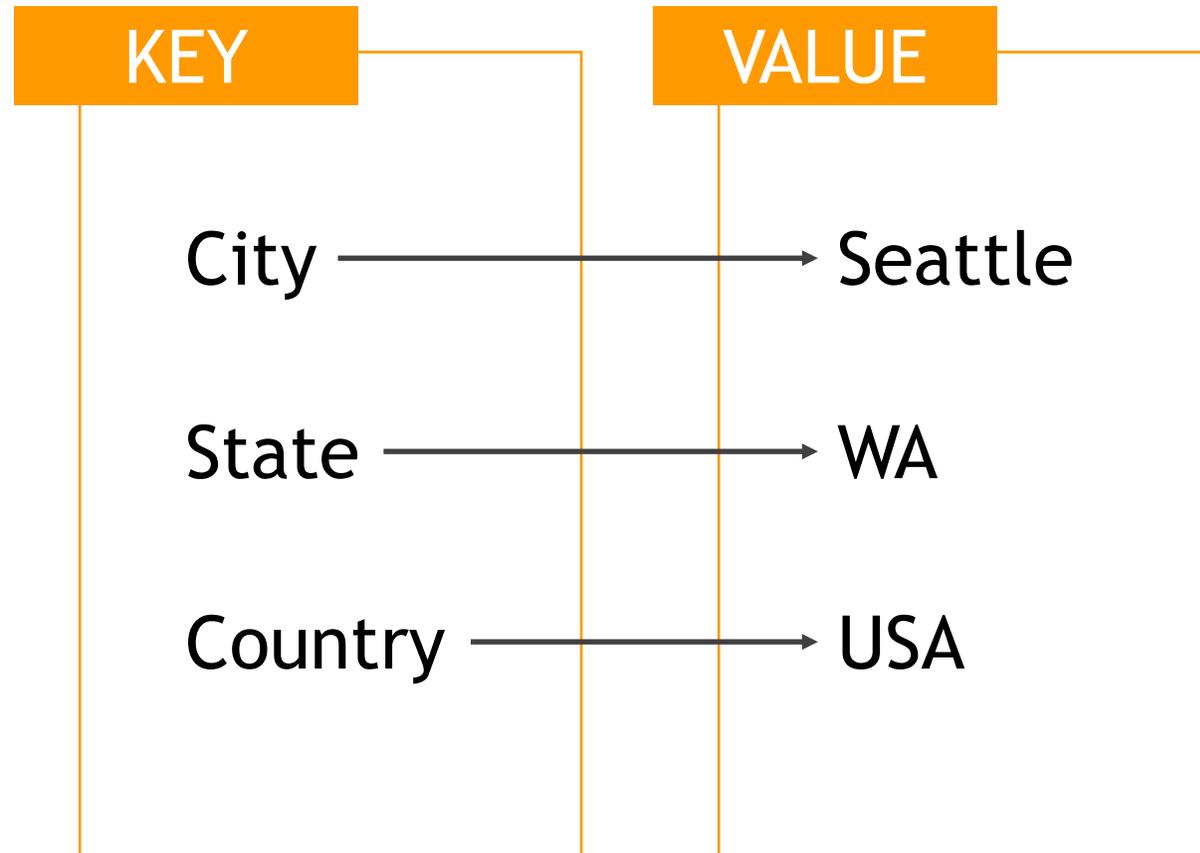
```
{  
  "first name": "Jeremy",  
  "last name": "Dunn",  
  "age": 27,  
  "hobbies": [  
    "Hiking",  
    "Surfing",  
    "Kayaking"  
  ],  
  "address": {  
    "number": "2401",  
    "street": "Monterey Avenue",  
    "city": "Monterey",  
    "state": "CA"  
  }  
}
```

Supports arrays

Flexible schema

And nesting...

Key-Value Database





DynamoDB Benefits and Use Cases

High performance, very fast for read and write operations

Massively scalable and serverless

Highly available across three Availability Zones

Shared Responsibility Model for DynamoDB

AWS Responsibility

Underlying infrastructure
(serverless)

Maintenance and patching

Customer (Your) Responsibility

Managing data

Setting up and managing
database users and permissions



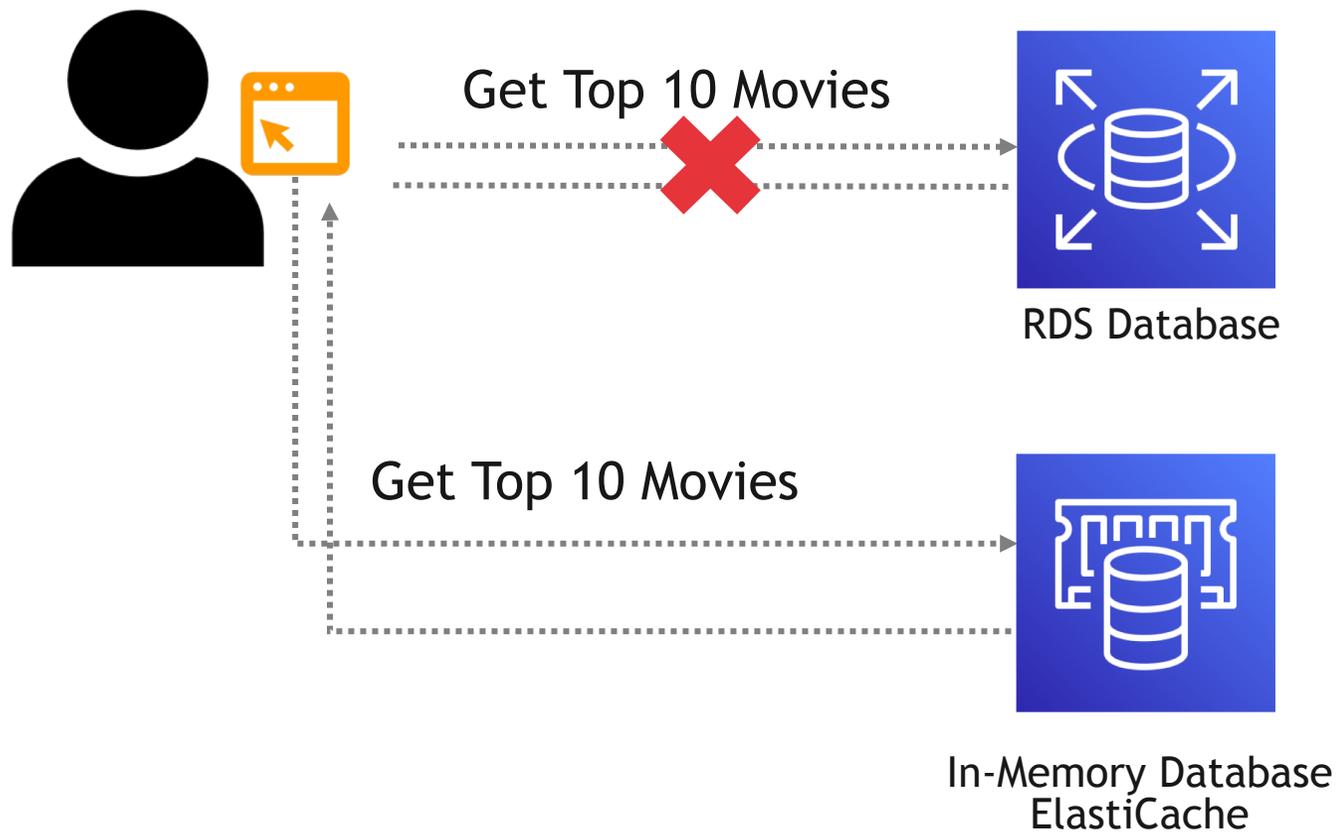
DEMO

Creating a DynamoDB Table

Purpose-Built Databases in AWS

DATABASE TYPE	USE CASE	AWS SERVICE
Relational	<ul style="list-style-type: none">Traditional apps, CRM, eCommerce	<ul style="list-style-type: none">Amazon RDS, Aurora, Redshift
Non-Relational or NoSQL (key-value)	<ul style="list-style-type: none">High-traffic web apps, eCommerce, Gaming	<ul style="list-style-type: none">Amazon DynamoDB
In-memory	<ul style="list-style-type: none">Caching, session management	<ul style="list-style-type: none">Amazon ElastiCache
Graph	<ul style="list-style-type: none">Fraud detection, social networking, recommendation engines	<ul style="list-style-type: none">Amazon Neptune

In-Memory Database





ElastiCache Benefits and Use Cases

Managed caching service

- Removes the complexity of deploying and managing a distributed cache environment
- Compatible with Redis or Memcached engines
- Integration with other AWS services

Use cases

- Read-intensive web applications
- Real-time apps that require fast data access

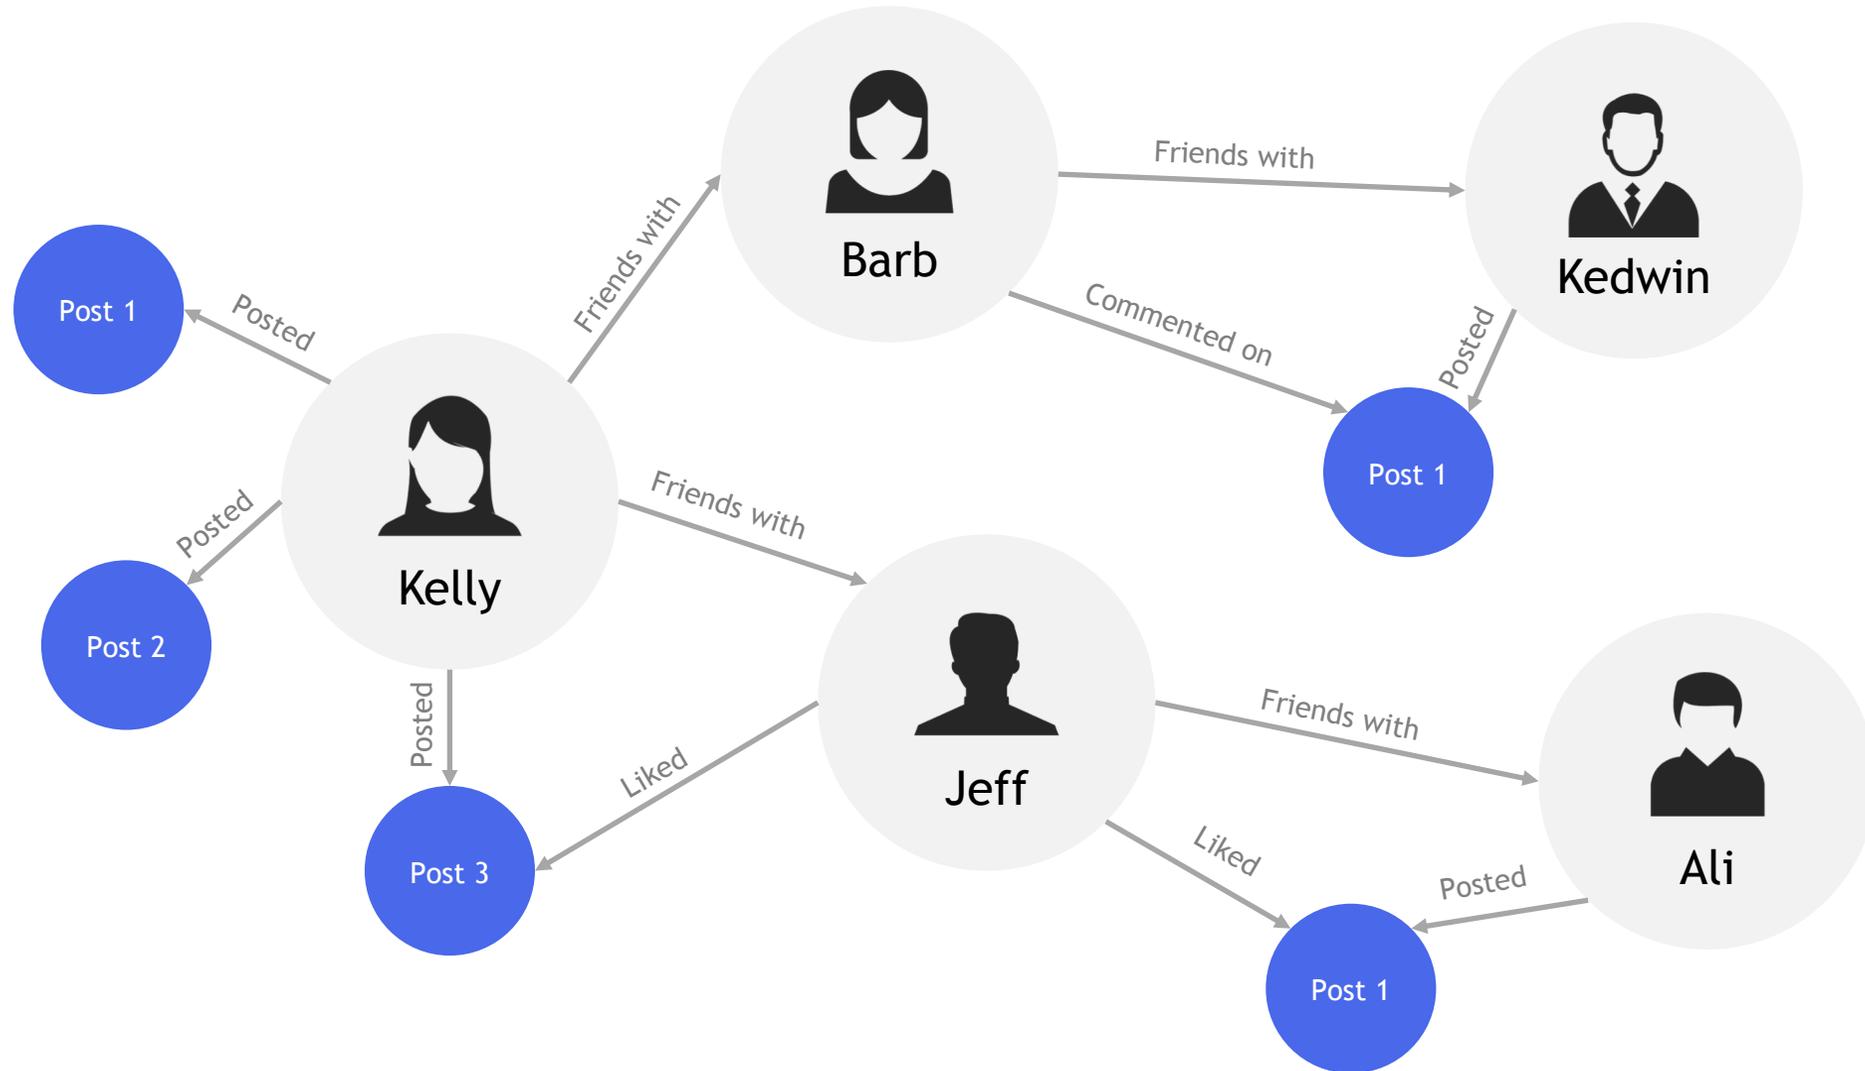
Purpose-Built Databases in AWS

DATABASE TYPE	USE CASE	AWS SERVICE
Relational	<ul style="list-style-type: none">Traditional apps, CRM, eCommerce	<ul style="list-style-type: none">Amazon RDS, Aurora, Redshift
Non-Relational or NoSQL (key-value)	<ul style="list-style-type: none">High-traffic web apps, eCommerce, Gaming	<ul style="list-style-type: none">Amazon DynamoDB
In-memory	<ul style="list-style-type: none">Caching, session management	<ul style="list-style-type: none">Amazon ElastiCache
Graph	<ul style="list-style-type: none">Fraud detection, social networking, recommendation engines	<ul style="list-style-type: none">Amazon Neptune

Graph Database



Neptune



Databases: Important Points to Remember

RELATIONAL DATABASE SERVICE (RDS)

- Relational database service with six engines available (Aurora, MySQL, PostgreSQL, etc.)
- Managed database service
 - AWS does all the underlying administrative work
 - You are responsible for network access and managing database users and permissions

REDSHIFT

- Used for data warehousing and analytics
 - A central repository of data that pulls from many other sources (e.g., S3, databases, etc.)

DYNAMODB

- Non-relational or NoSQL database
 - Data stored as key/value pairs
 - Massively scalable and highly performant

Databases: Important Points to Remember

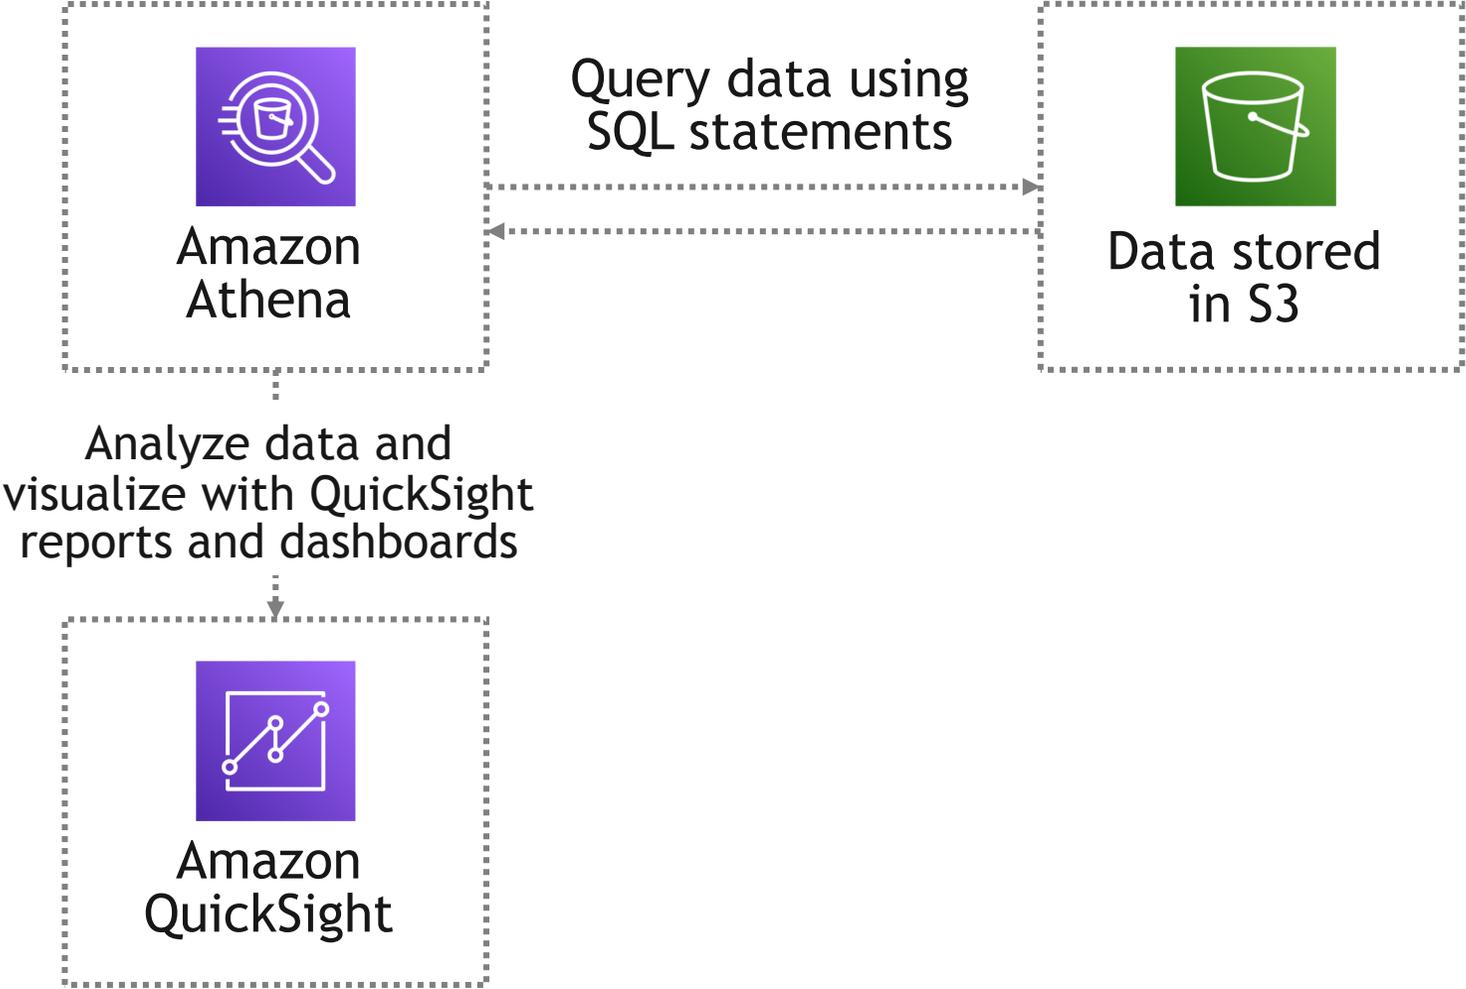
ELASTICACHE

- In-memory database used for caching and session management
- Used for read-intensive web applications, such as gaming or media streaming

NEPTUNE

- Graph database
- Used for social media, fraud detection and recommendation engines

Amazon Athena





Amazon Athena

Serverless

- The underlying instances are handled by AWS
- Fast querying due to parallel execution

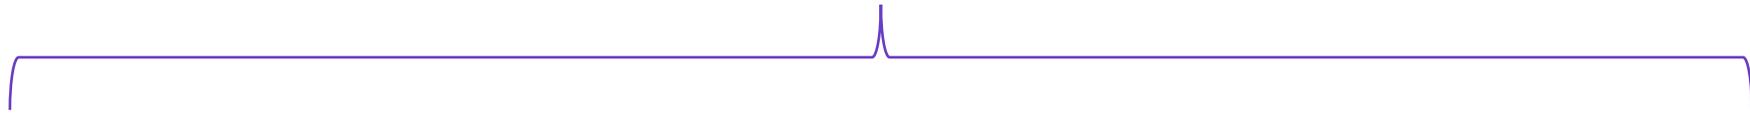
Use cases

- Querying AWS logs (CloudWatch, VPC Flow Logs, CloudTrail, etc.)
- Business intelligence/analytics
- Querying S3 using SQL statements

Big Data Stream (Real-Time) Processing



Amazon Kinesis



Kinesis
Video Streams

Kinesis
Data Streams

Kinesis
Data Firehose

Kinesis
Data Analytics

The Kinesis Family of Services

SERVICE	HIGHLIGHTS
 Video Streams	<ul style="list-style-type: none">• Ingests, processes/streams and stores streaming video and audio data
 Data Streams	<ul style="list-style-type: none">• Ingests, processes and stores streaming data, breaking it into “shards”• Data has to be processed (e.g., Lambda, Data Analytics) before storing (which is optional)
 Data Firehose	<ul style="list-style-type: none">• Ingests, processes and stores streaming data, without “shards”• Can stream directly to storage (processing is optional)
 Data Analytics	<ul style="list-style-type: none">• Analyzes streaming data• Enables SQL querying and custom Java applications



Serverless

- The underlying instances are handled by AWS

Integrates with a variety of AWS services (Athena, Redshift, RDS, S3, etc.)

Use cases

- Business intelligence/analytics
- Data visualization and dashboards



EMR = Elastic MapReduce

Hosted Hadoop framework

- Hadoop is an open-source framework to store and process huge amounts of data in scalable clusters (i.e., collection of computers)
- Also supports Apache Spark, Hive, Presto and other workloads

Use cases

- Big data processing
- Machine learning

Analytics: Important Points to Remember

ATHENA

- Used to query data in S3 using SQL queries

KINESIS FAMILY OF PRODUCTS

- Used for stream or real-time processing of data (such as stock prices or video streaming)

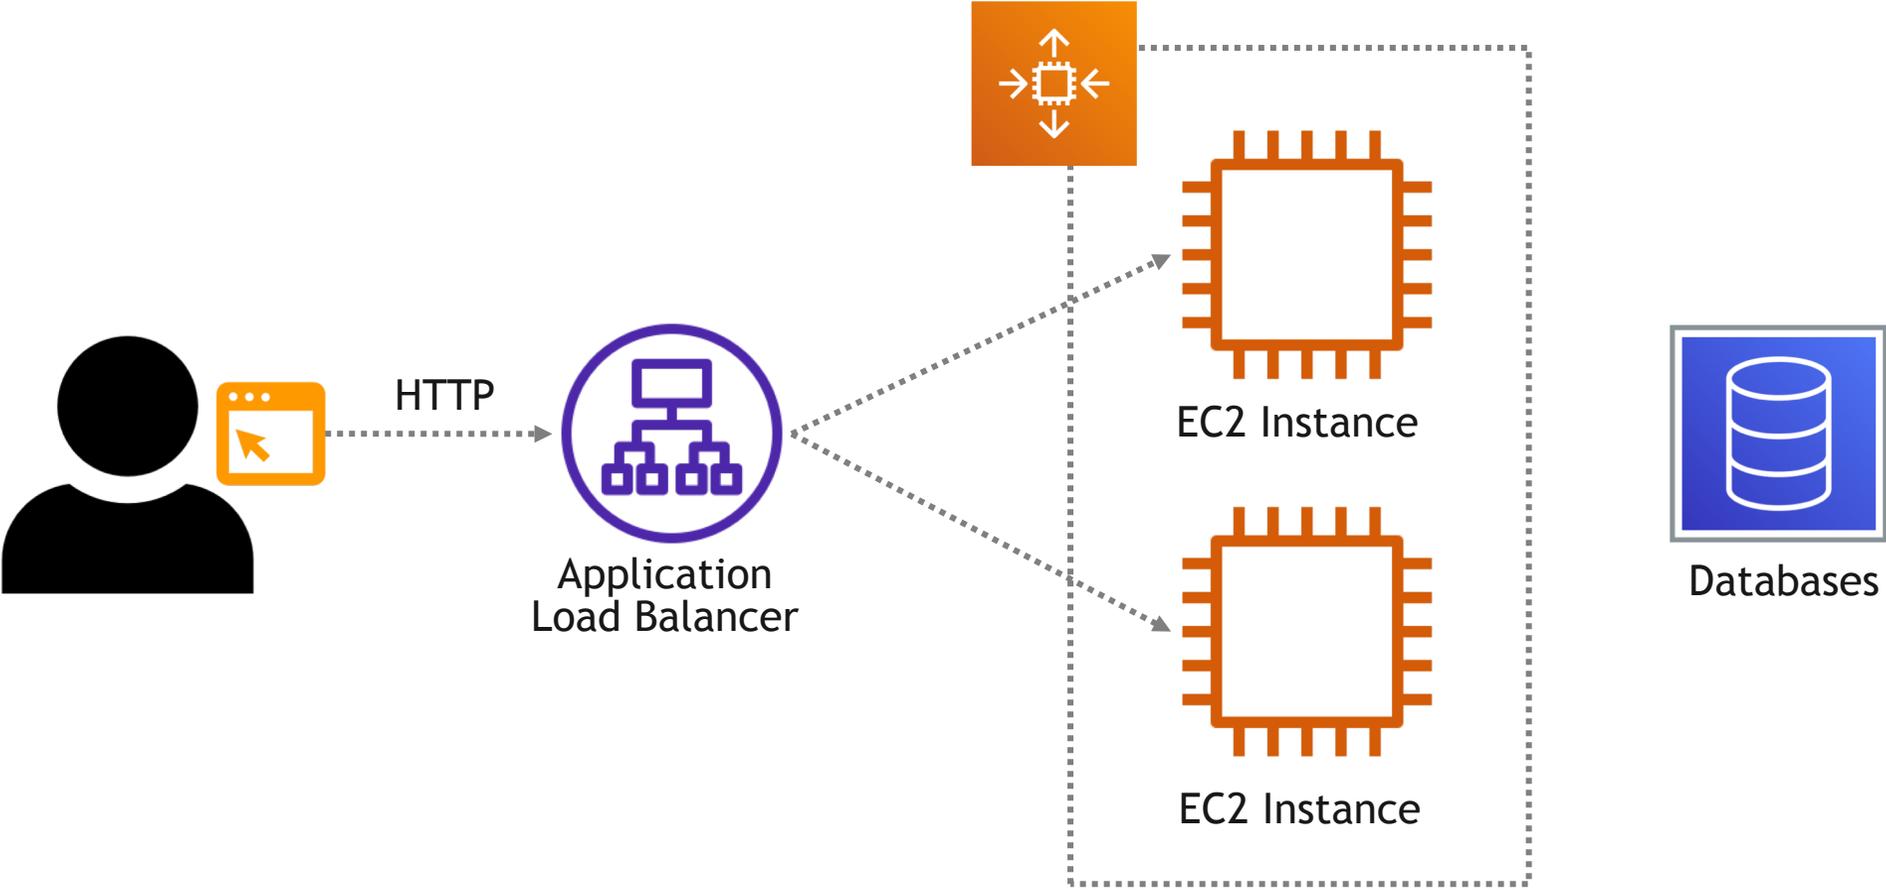
QUICKSIGHT

- Used for business intelligence/analytics to visualize your data using reports and dashboards

EMR (ELASTIC MAPREDUCE)

- A hosted Hadoop framework that also supports Spark, Hive and Presto

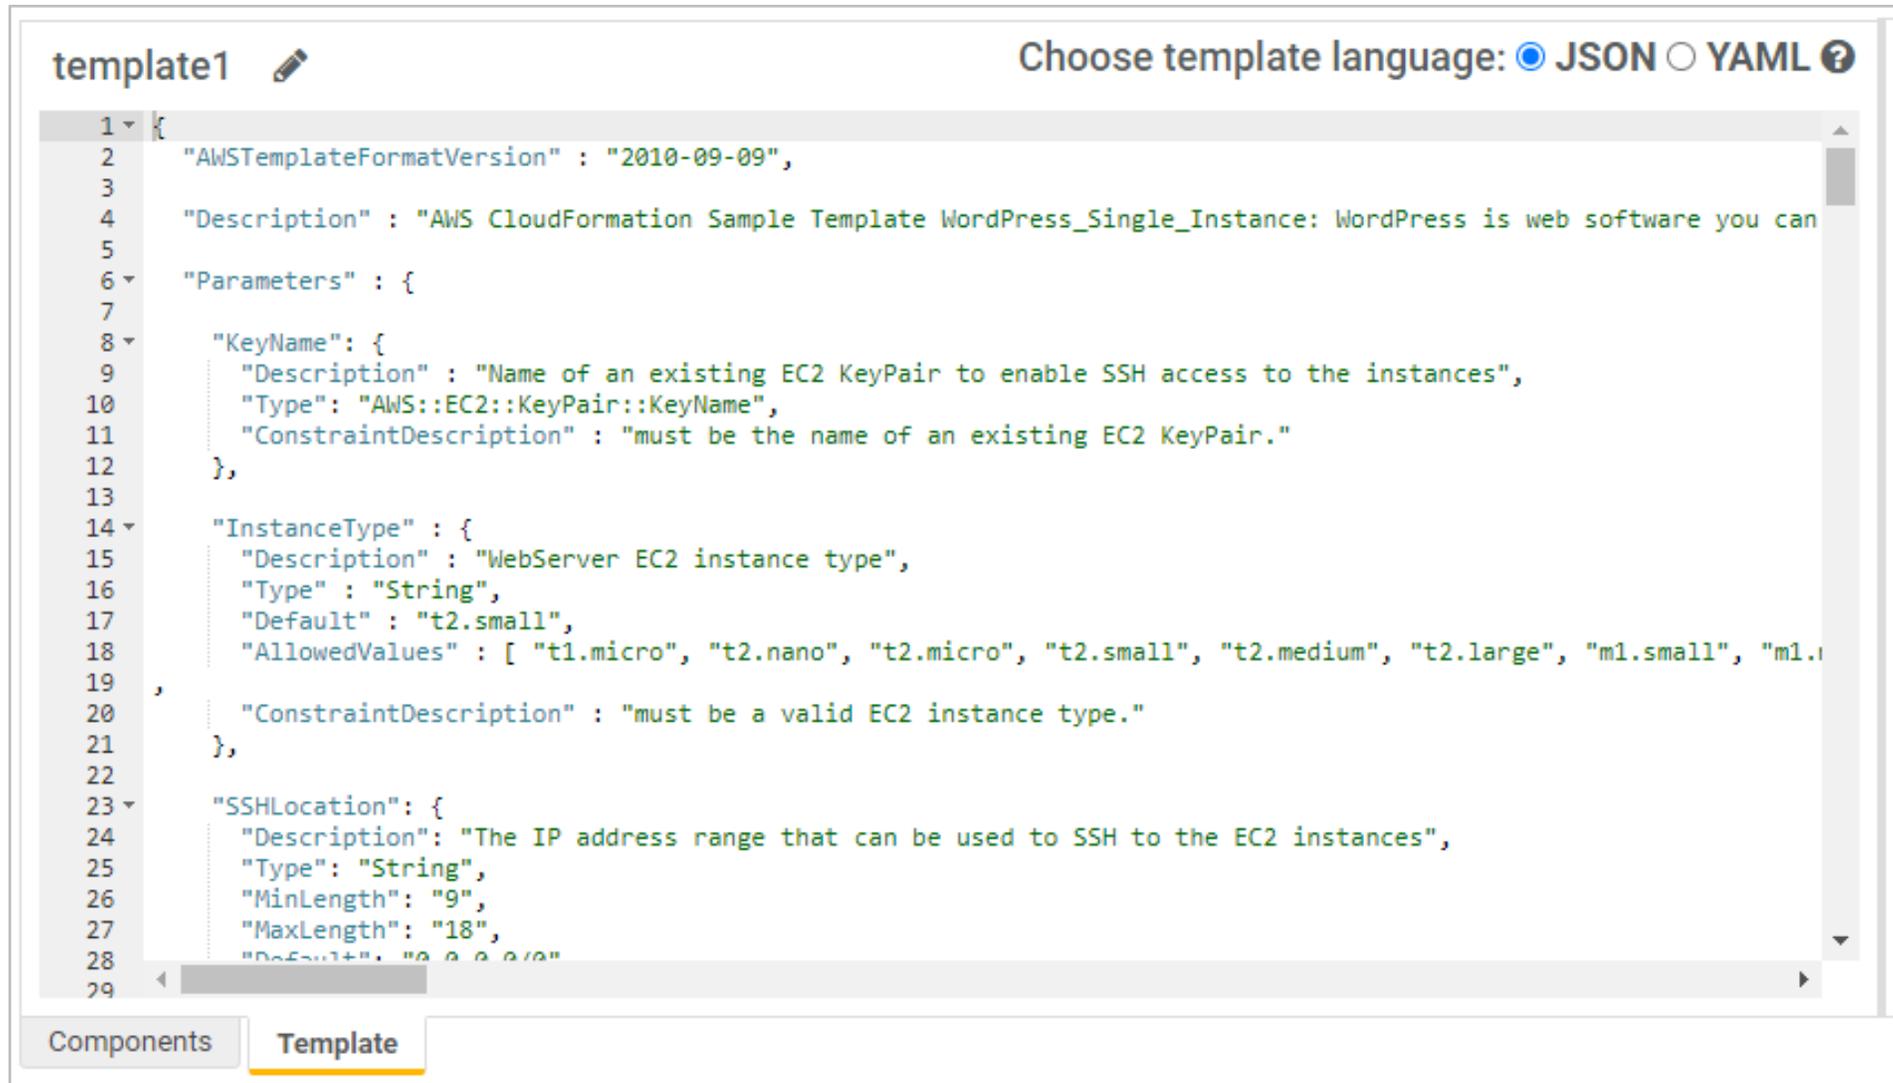
CloudFormation: Infrastructure as Code



I WANT:

- 2 EC2 instances with this configuration
- An auto-scaling group and load balancer
- 1 RDS database
- 1 VPC with 1 public subnet and 1 private subnet

CloudFormation Template



```
template1  Choose template language:  JSON  YAML 
1 {
2   "AWSTemplateFormatVersion" : "2010-09-09",
3
4   "Description" : "AWS CloudFormation Sample Template WordPress_Single_Instance: WordPress is web software you can
5
6   "Parameters" : {
7
8     "KeyName": {
9       "Description" : "Name of an existing EC2 KeyPair to enable SSH access to the instances",
10      "Type": "AWS::EC2::KeyPair::KeyName",
11      "ConstraintDescription" : "must be the name of an existing EC2 KeyPair."
12    },
13
14   "InstanceType" : {
15     "Description" : "WebServer EC2 instance type",
16     "Type" : "String",
17     "Default" : "t2.small",
18     "AllowedValues" : [ "t1.micro", "t2.nano", "t2.micro", "t2.small", "t2.medium", "t2.large", "m1.small", "m1.l
19
20     "ConstraintDescription" : "must be a valid EC2 instance type."
21   },
22
23   "SSHLocation": {
24     "Description": "The IP address range that can be used to SSH to the EC2 instances",
25     "Type": "String",
26     "MinLength": "9",
27     "MaxLength": "18",
28     "Default": "0.0.0.0/0"
29
```

Components **Template**

Benefits of CloudFormation

Huge time savings over manual work

Repeatable

Changes are also made as code

Can delete all resources at once

Estimate costs based on a template

Can be used cross-account and cross-region



DEMO

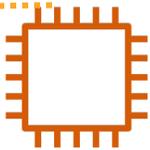
Deploying a CloudFormation Template



SERVICE CATALOG

SERVICE CATALOG

Web Development Portfolio



EC2



Application Load Balancer



Database

Data Analysis Portfolio



S3



Athena



QuickSight

Portfolios made up of products



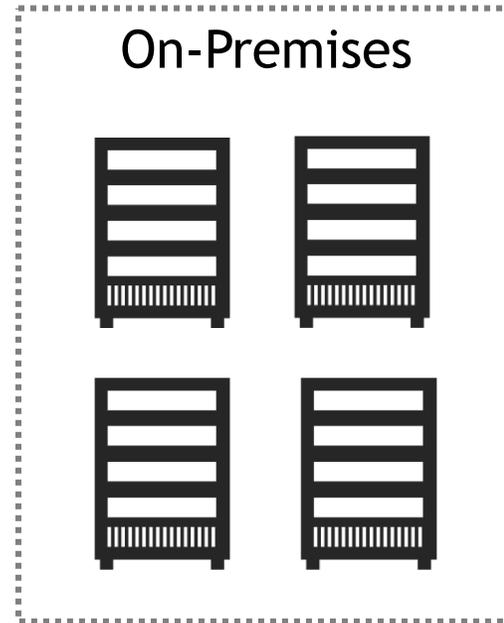
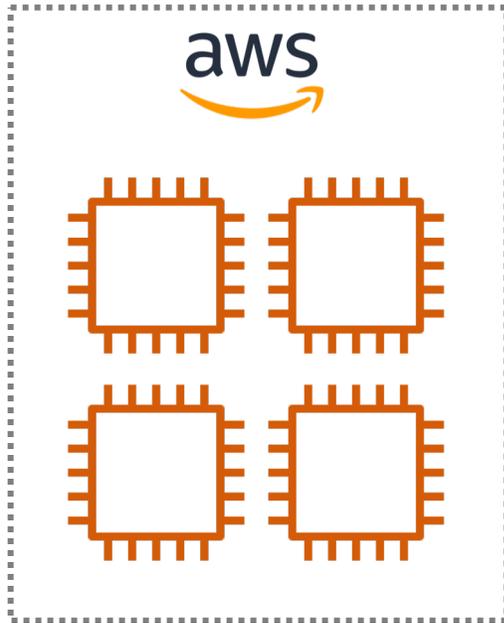
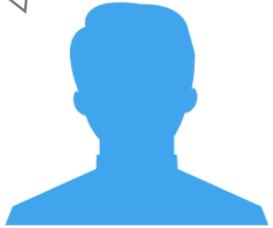


AWS Service Catalog

IT-approved resources that teams can provision

Engineering teams can provision things from the catalog

I need to patch all my servers





AWS Systems Manager (SSM)

Manage your fleet of servers at scale

- EC2 instances and on-premises servers
- Supports Linux and Windows

Common use cases

- Patching and maintenance
- “Run Command” to install applications, capture log files, perform configuration changes
- Parameter Store to manage global configuration settings

Deploying/Managing Infrastructure: Important Points to Remember

CLOUDFORMATION

- “Infrastructure as Code,” defined in a JSON or YAML template
- Repeatable way to deploy, update and delete resources
- Can be used cross-account and cross-region

SERVICE CATALOG

- An approved “catalog” of products that developers can use
- Organized into portfolios (e.g., web development, data analysis)

SYSTEMS MANAGER (SSM)

- Manage a fleet of servers at scale
- Commonly used for patching and maintenance



Deploy my
awesome
app



NOW WHAT?

- What does website traffic look like?
- How is performance?
- How much bandwidth is my app using?
- Are compute resources optimized?



Amazon CloudWatch

Provides visibility into AWS resources and applications

- Near-real-time metrics such as CPU, memory, disk, GPU utilization
- Dashboards offer at-a-glance views
- Alarms can trigger notifications for metrics
 - Example: Billing alarm to notify you when you've hit a certain amount



DEMO

Creating CloudWatch Alarms



Records user activity and API usage on your AWS account

- Examples: Create an IAM role, terminate an EC2 instance, log into the Console

Primarily used for auditing

Automatically enabled

Saves a history of API calls made through:

- AWS Management Console
- CLI
- SDKs
- Higher-level AWS services (like CloudFormation)

Logs can be saved to CloudWatch or S3

- Needed for storage longer than 90 days



CloudWatch vs. CloudTrail



CLOUDWATCH

Performance monitoring of applications and resources

Metrics such as CPU, memory, disk, GPU utilization

CLOUDTRAIL

Captures user activity and API calls on an AWS account

Activity such as creating new resources, terminating resources, logins, etc.



DEMO

Working with CloudTrail

AWS Service Health Dashboard

The screenshot shows the AWS Service Health Dashboard in a web browser. The page title is "AWS Service Health Dashboard" and the URL is "status.aws.amazon.com". The AWS logo is prominently displayed at the top left, followed by the text "SERVICE HEALTH DASHBOARD". Below this, there is a navigation link for "Amazon Web Services" and a sub-link for "Service Health Dashboard". A message encourages users to get a personalized view of AWS service health, with a button labeled "Open the Personal Health Dashboard".

Current Status - Jan 13, 2022 PST

Amazon Web Services publishes our most up-to-the-minute information on service availability in the table below. Check back here any time to get current status information, or subscribe to an RSS feed to be notified of interruptions to each individual service. If you are experiencing a real-time, operational issue with one of our services that is not described below, please inform us by clicking on the "Contact Us" link to submit a service issue report. All dates and times are Pacific Time (PST/PDT).

North America	South America	Europe	Africa	Asia Pacific	Middle East	Contact Us
Recent Events Details RSS						
✔ No recent events.						
Remaining Services Details RSS						
✔ Alexa for Business (N. Virginia) Service is operating normally RSS						
✔ Amazon API Gateway (Montreal) Service is operating normally RSS						
✔ Amazon API Gateway (N. California) Service is operating normally RSS						
✔ Amazon API Gateway (N. Virginia) Service is operating normally RSS						
✔ Amazon API Gateway (Ohio) Service is operating normally RSS						
✔ Amazon API Gateway (Oregon) Service is operating normally RSS						
✔ Amazon AppFlow (Montreal) Service is operating normally RSS						
✔ Amazon AppFlow (N. California) Service is operating normally RSS						
✔ Amazon AppFlow (N. Virginia) Service is operating normally RSS						

status.aws.amazon.com

Personal Health Dashboard

The screenshot displays the AWS Personal Health Dashboard in a browser window. The address bar shows the URL `phd.aws.amazon.com/phd/home#/dashboard/open-issues`. The dashboard header includes the AWS logo, a search bar, and user information for Amber Israelsen. The main content area is titled "Dashboard" and shows a refresh status of "Last refreshed less than 1 min ago".

Overview

Metric	Value	Time Period
Open issues	0	Past 7 days
Scheduled changes	0	Upcoming and past 7 days
Other notifications	0	Past 7 days

Set up alerts

Use Amazon CloudWatch Events to create rules and receive notifications for events that might affect your AWS infrastructure.

[Create rule](#)

Navigation tabs at the bottom: **Open issues** | Scheduled changes | Other notifications

Footer: Feedback | English (US) | © 2008 - 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved. | Privacy Policy | Terms of Use | Cookie preferences

phd.aws.amazon.com

Monitoring: Important Points to Remember

CLOUDWATCH

- Used for performance monitoring of applications and resources
 - e.g., CPU, memory, disk, GPU utilization

CLOUDTRAIL

- Captures user activity and API calls
 - e.g., who created an EC2 instance, who terminated an EC2 instance, who logged into the Console

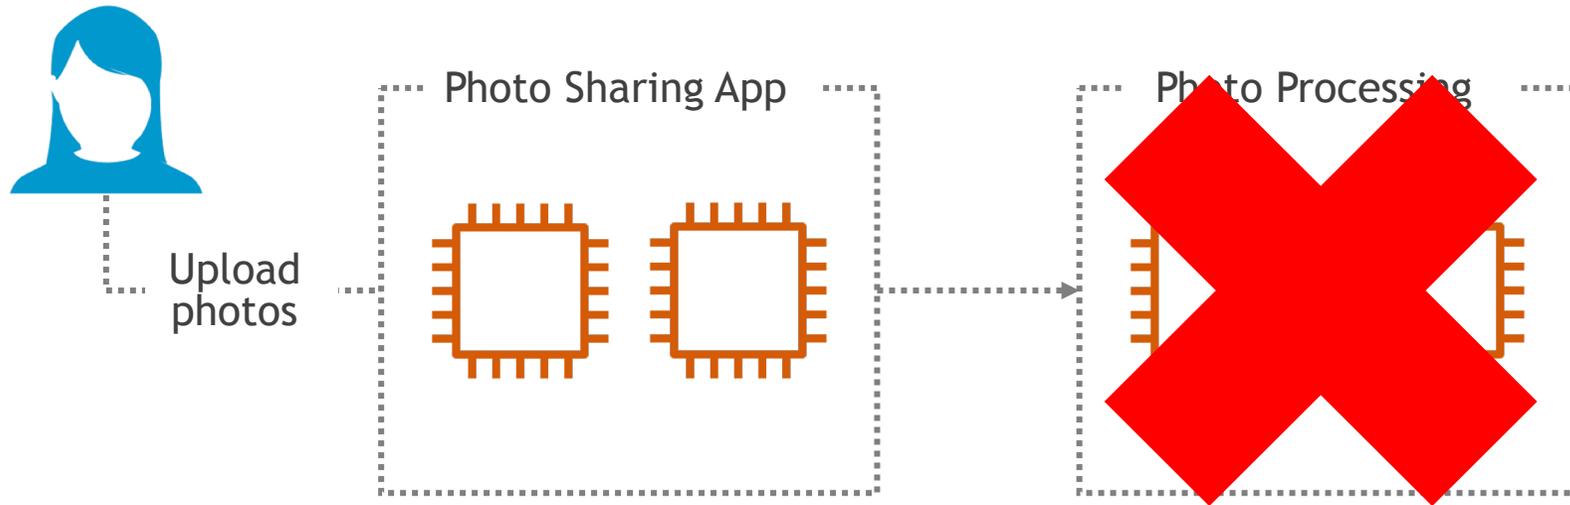
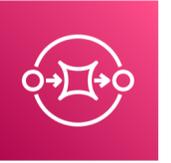
AWS SERVICE HEALTH DASHBOARD

- Status of all services around the world
- Public site available for anyone to check (status.aws.amazon.com)

PERSONAL HEALTH DASHBOARD

- View alerts about services that underlie YOUR applications
- Must be logged into your account (pdh.aws.amazon.com)

Simple Queue Service (SQS)



Simple Queue Service (SQS)



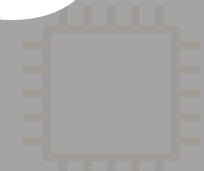
Tightly
coupled



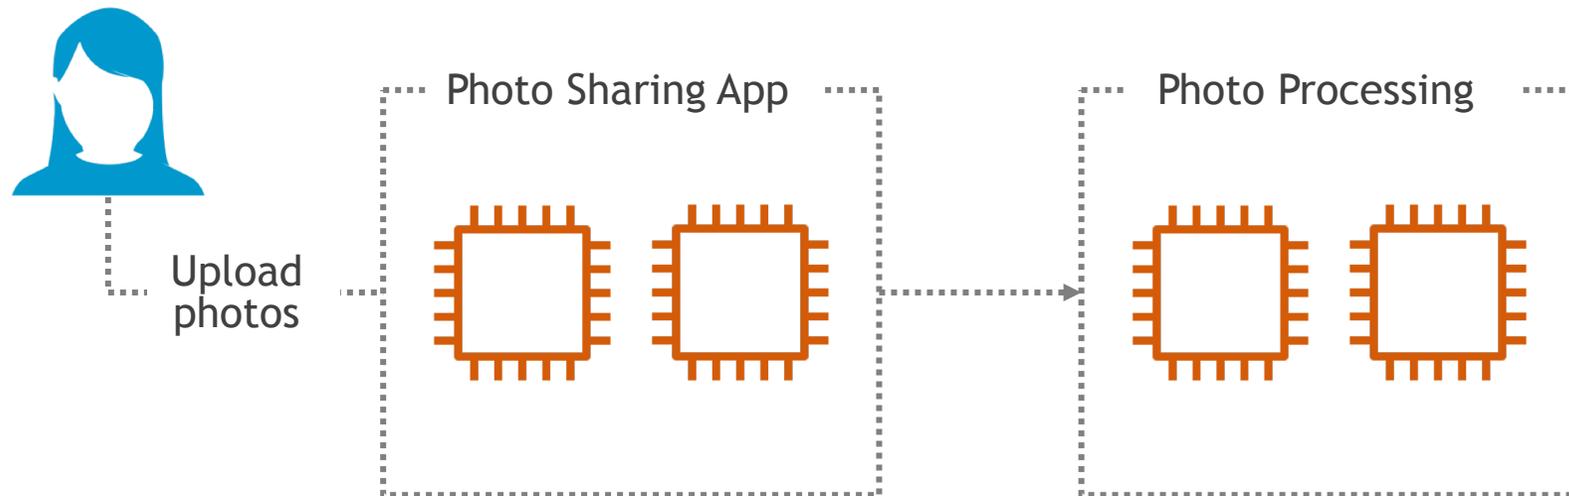
Upload
photos

Photo Storage App

Photo Processing

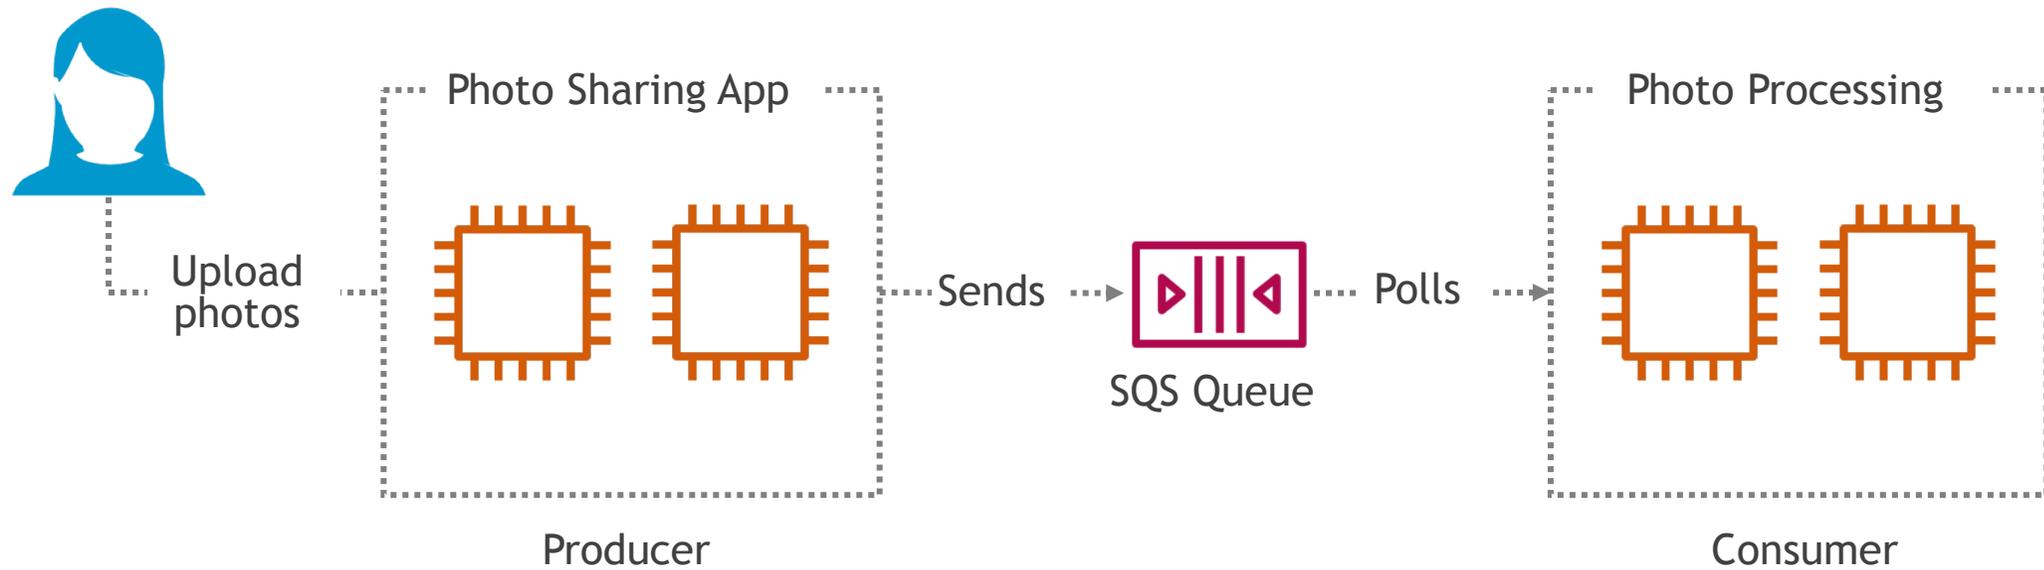


Simple Queue Service (SQS)



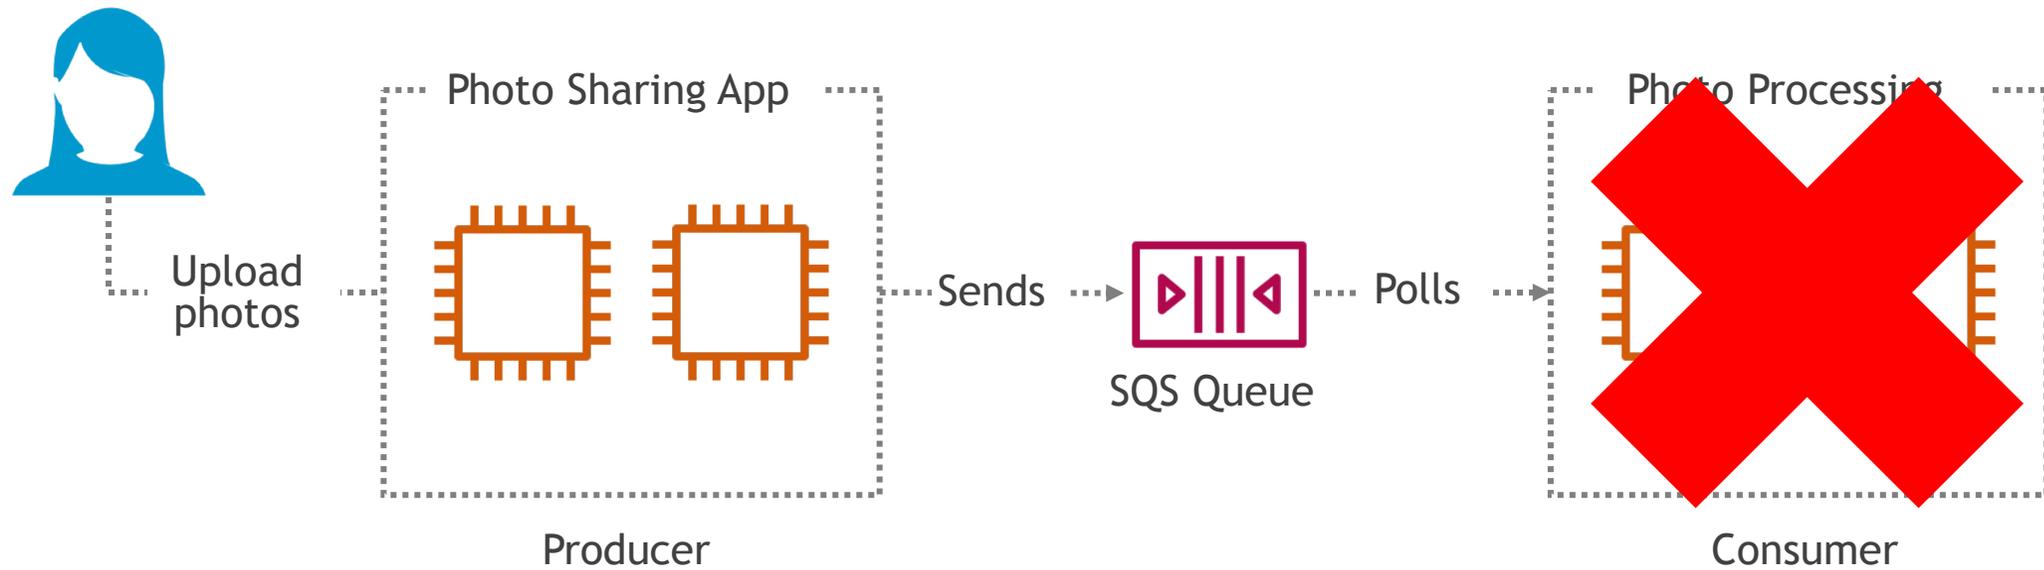
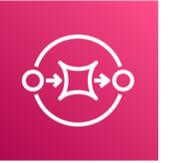
Simple Queue Service (SQS)

Decoupled



Simple Queue Service (SQS)

Decoupled





Amazon Simple Queue Service (SQS)

Supports multiple producers and multiple consumers

Messages deleted after they are read by the consumer

- Must happen within 14 days (the max storage time)

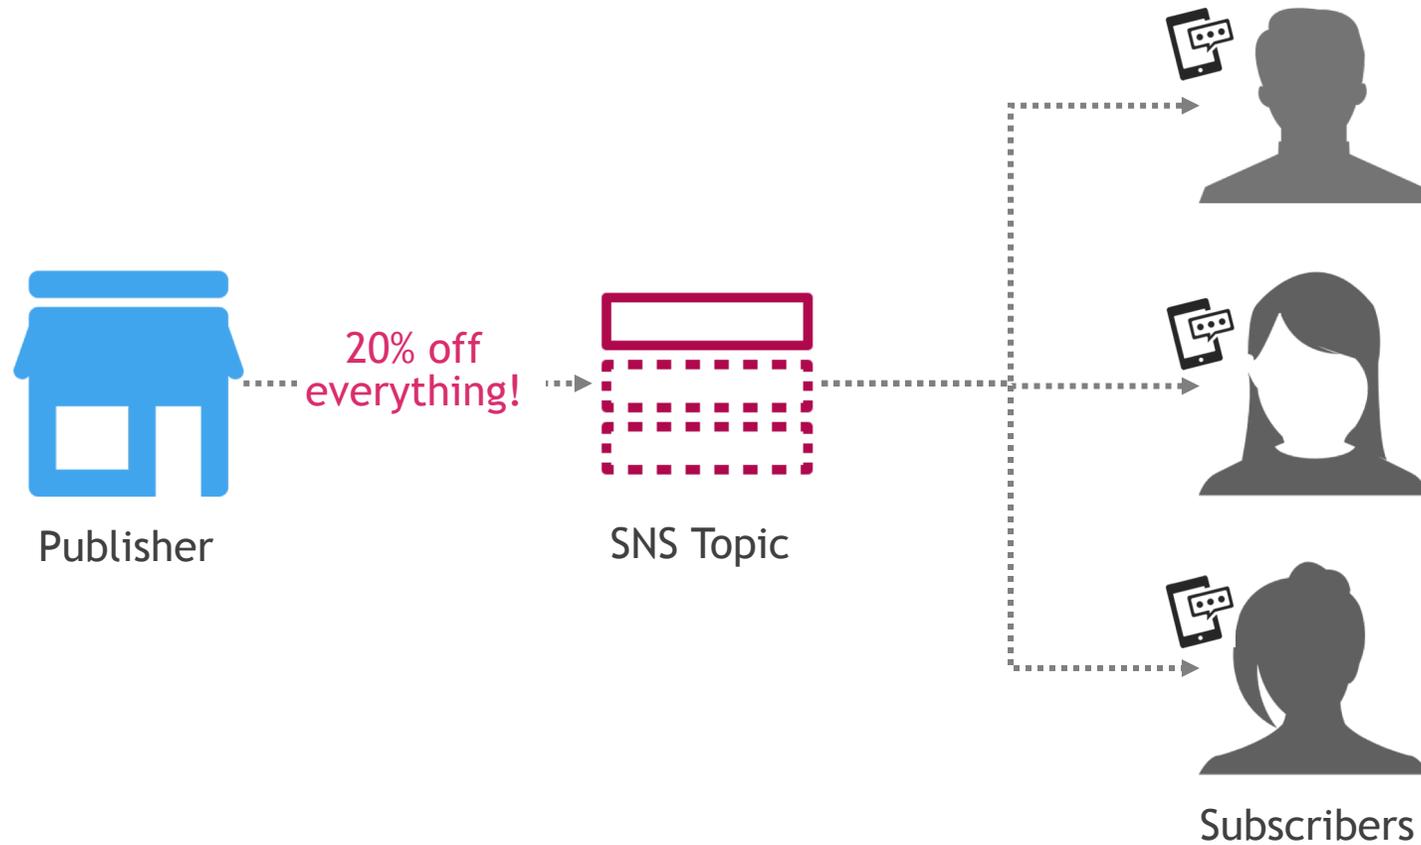


DEMO

Sending and Receiving Messages with SQS

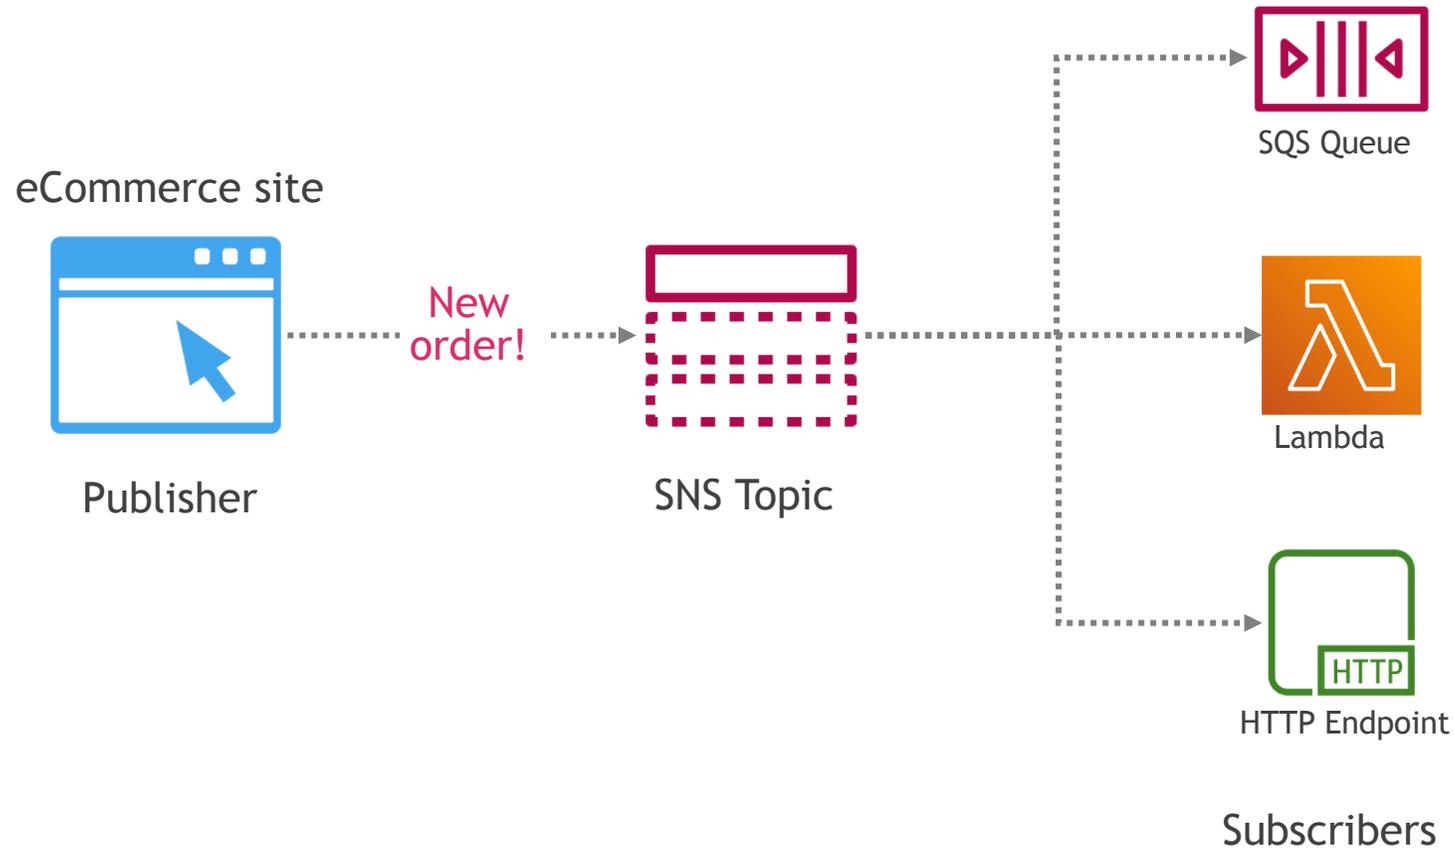
Simple Notification Service (SNS)

Pub/Sub



Simple Notification Service (SNS)

Pub/Sub





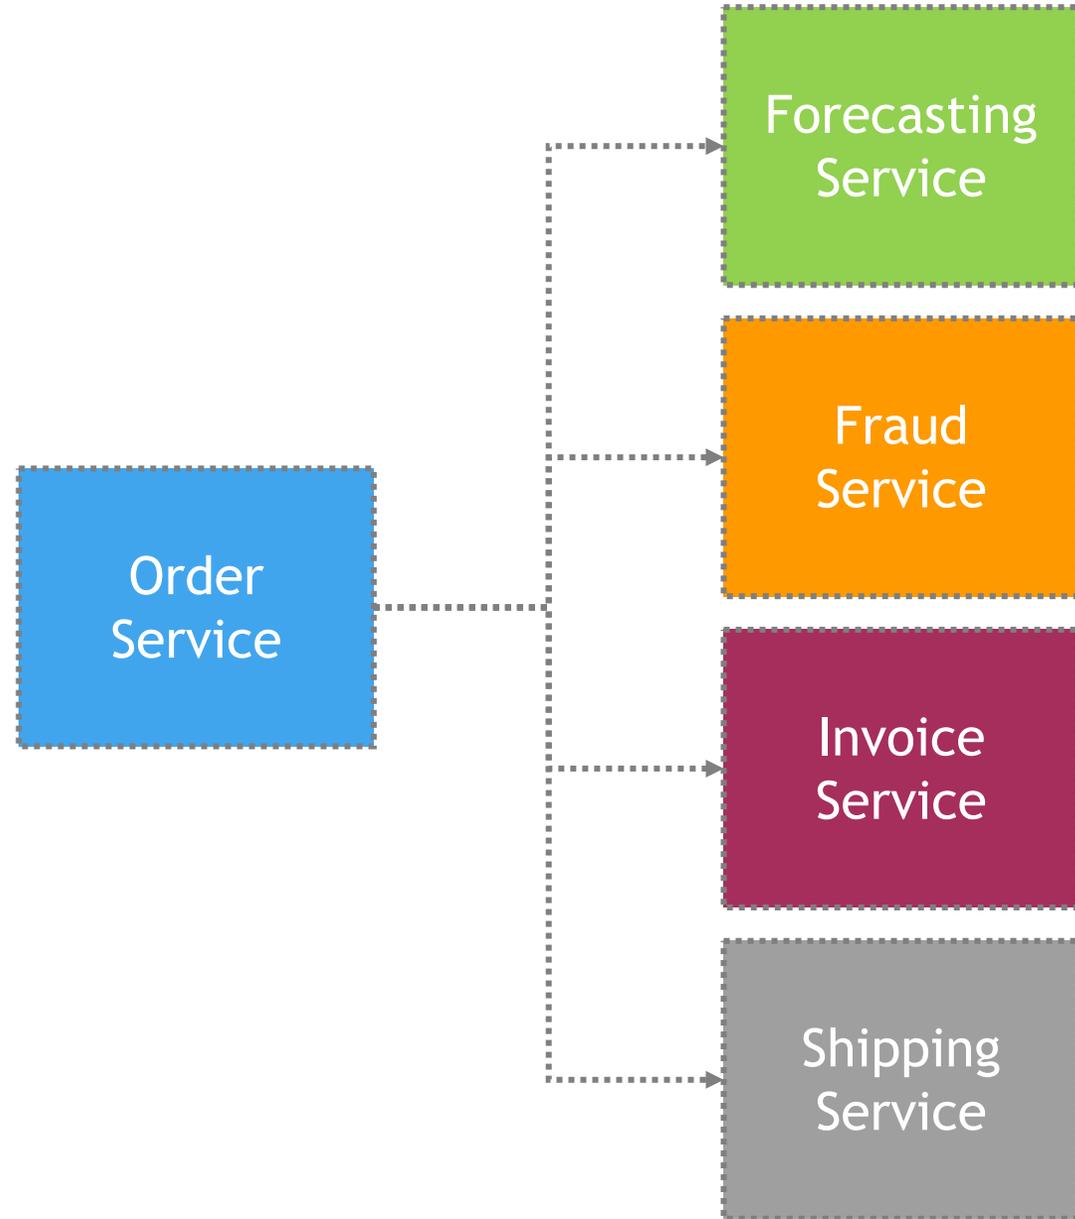
DEMO

Sending an Email Using Amazon SNS



Amazon EventBridge (formerly CloudWatch Events)

Build decoupled,
event-driven architectures





Amazon EventBridge (formerly CloudWatch Events)

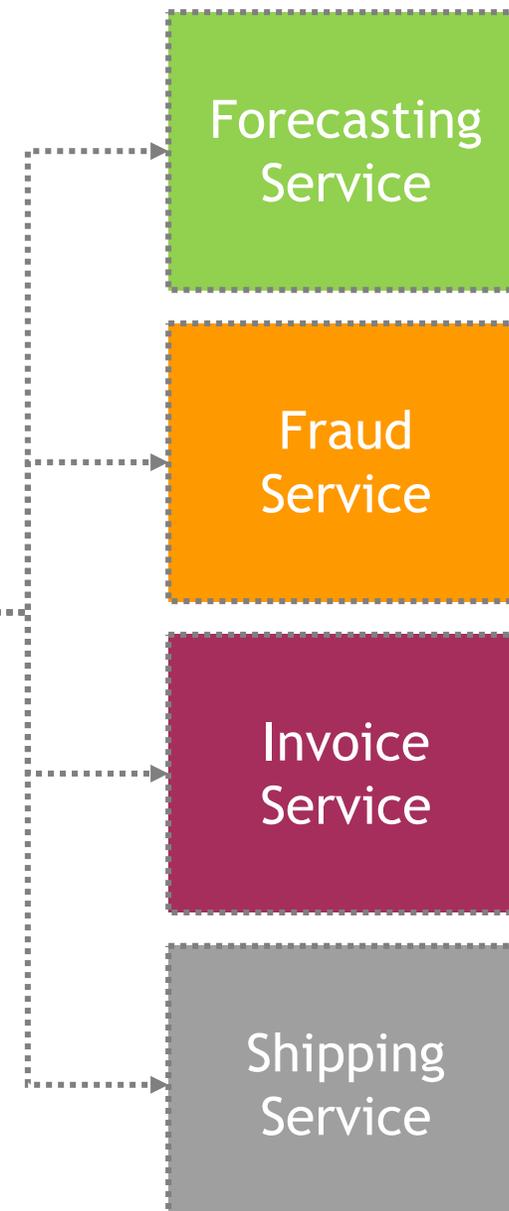
Pub/Sub



Publisher



Event Bus



Subscribers



Amazon

EventBridge
(formerly
CloudWatch
Events)

Pub/Sub

Order
Service

Publisher



Event Bus

Forecasting
Service

Fraud
Service

Invoice
Service

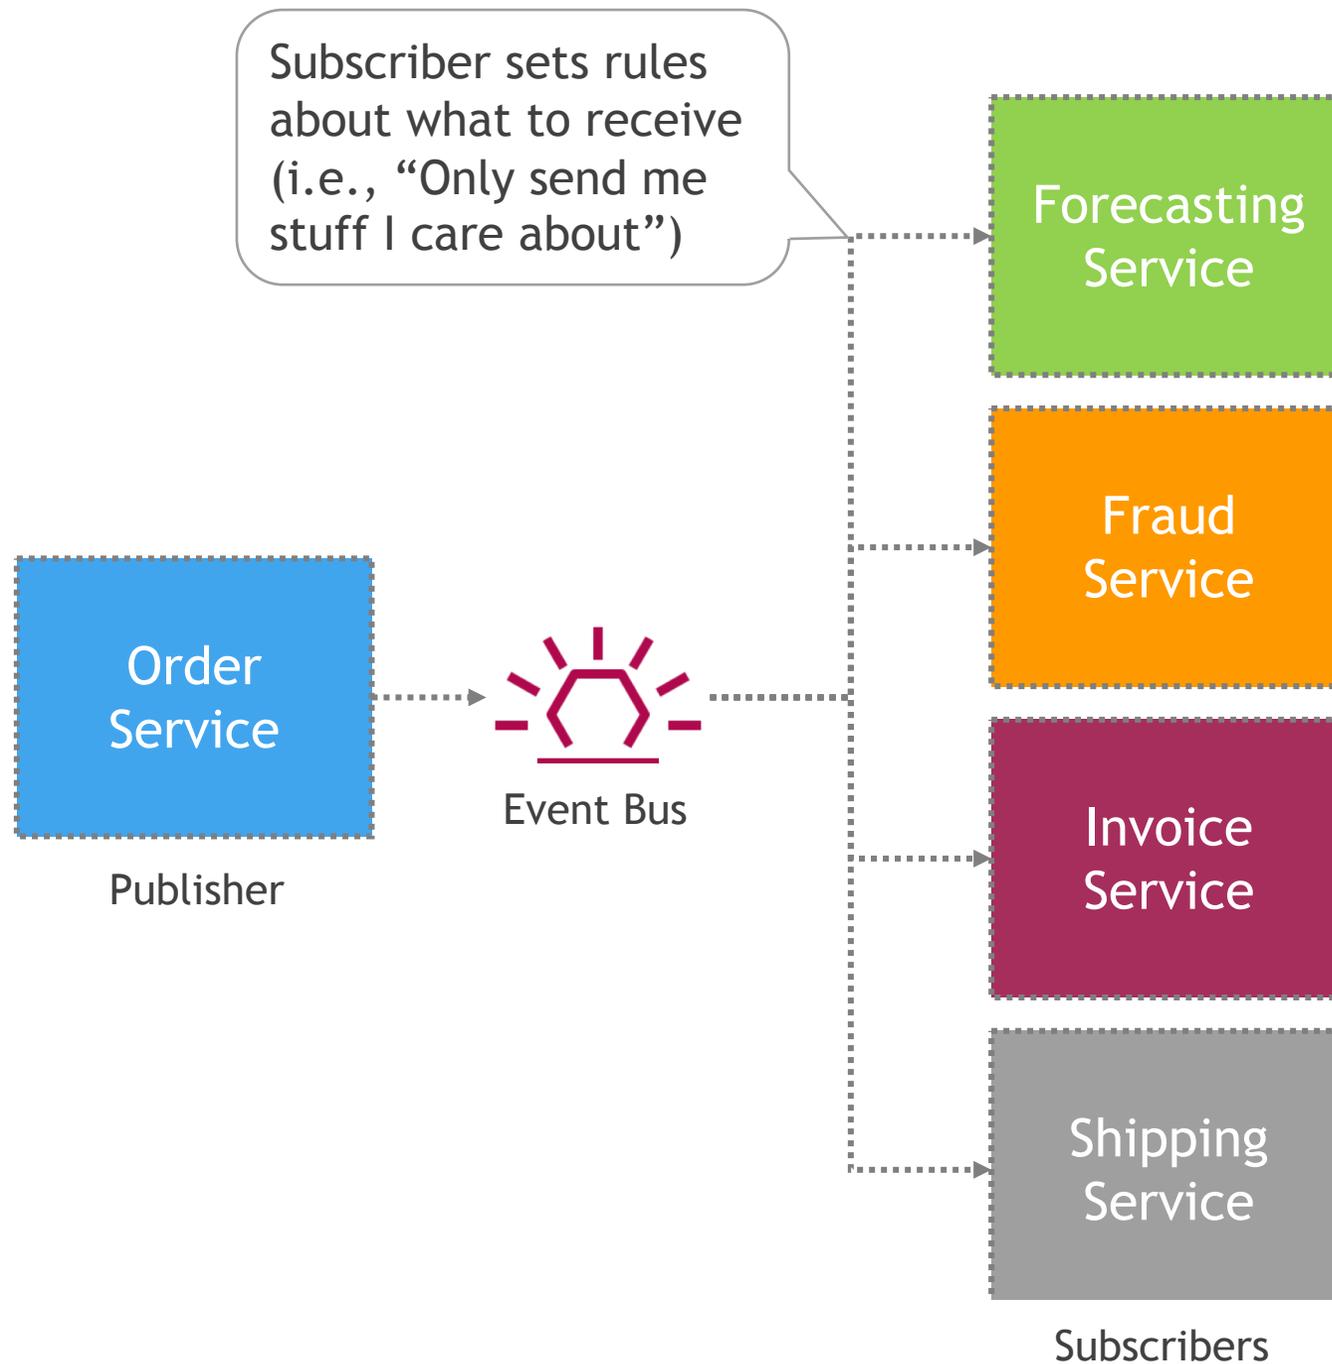
Shipping
Service

Subscribers

Isn't this the same thing as SNS?
Or SQS?



Amazon EventBridge (formerly CloudWatch Events)





Amazon EventBridge (formerly CloudWatch Events)

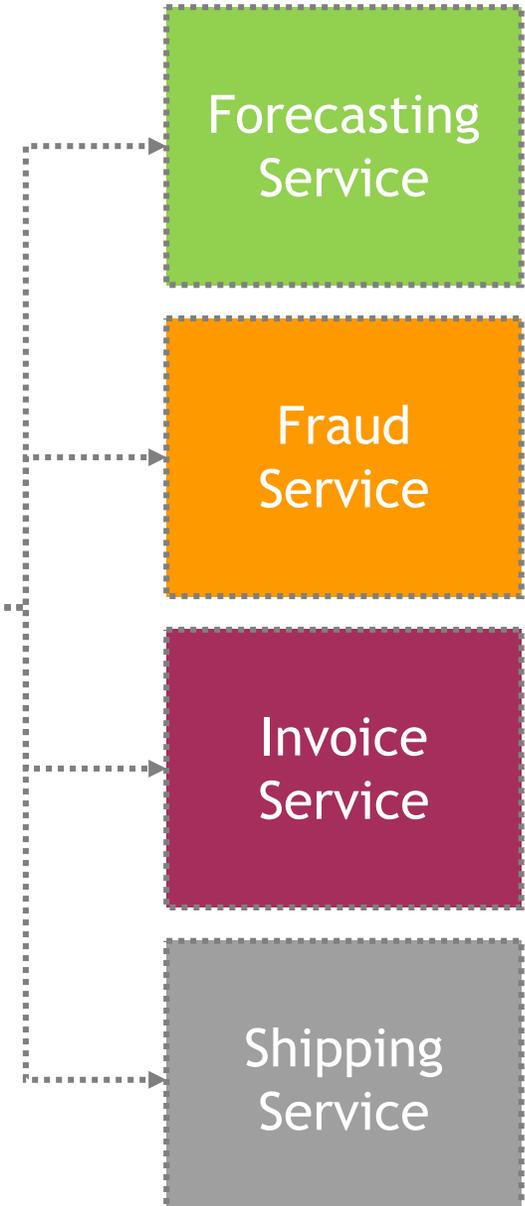
Schema Registry to
define up-front what
the event will look like



Publisher



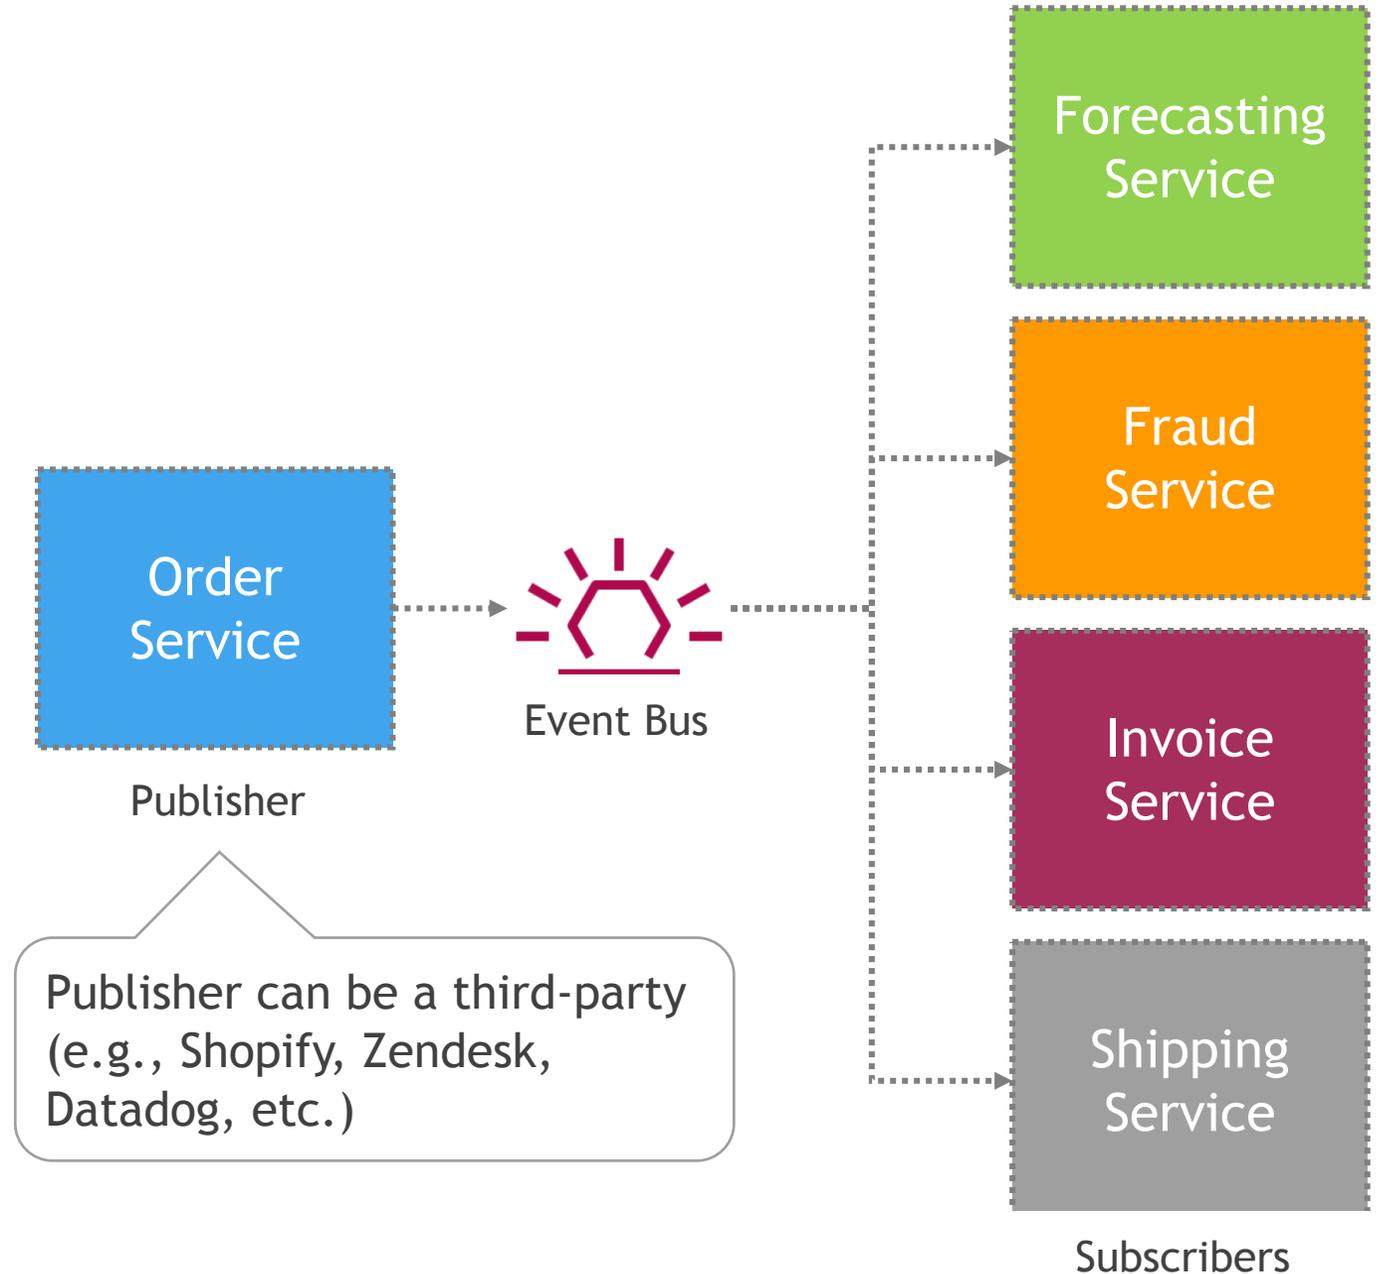
Event Bus



Subscribers



Amazon EventBridge (formerly CloudWatch Events)





Amazon EventBridge (formerly CloudWatch Events)

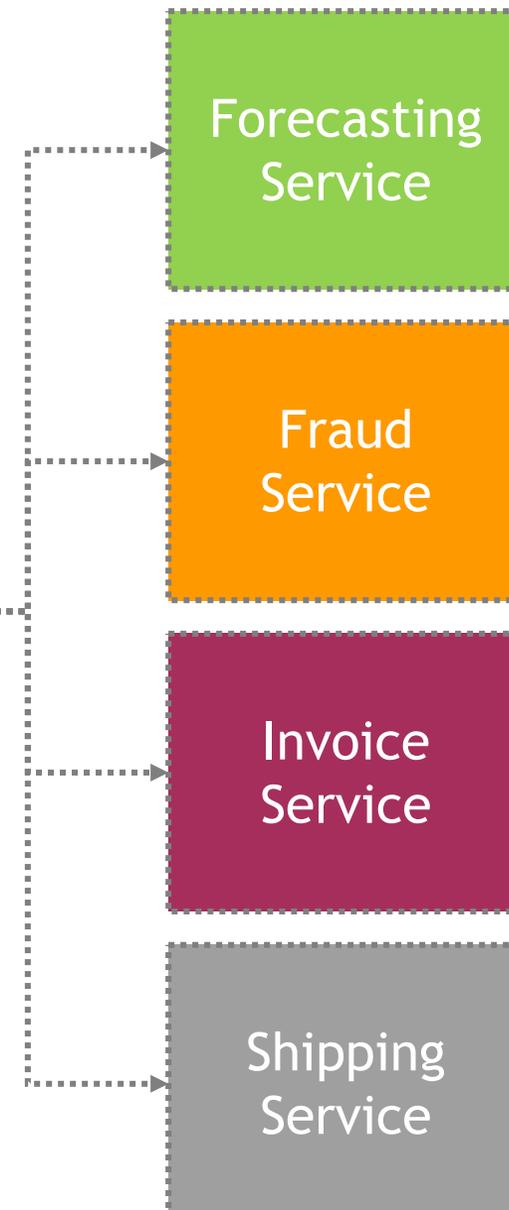
Schedule events
(e.g., every hour, call
a Lambda function to
write to a log)



Publisher



Event Bus



Subscribers

Application Integration: Important Points to Remember

SIMPLE QUEUE SERVICE (SQS)

- Key way to decouple components
- Producer sends to the queue; consumer polls from the queue

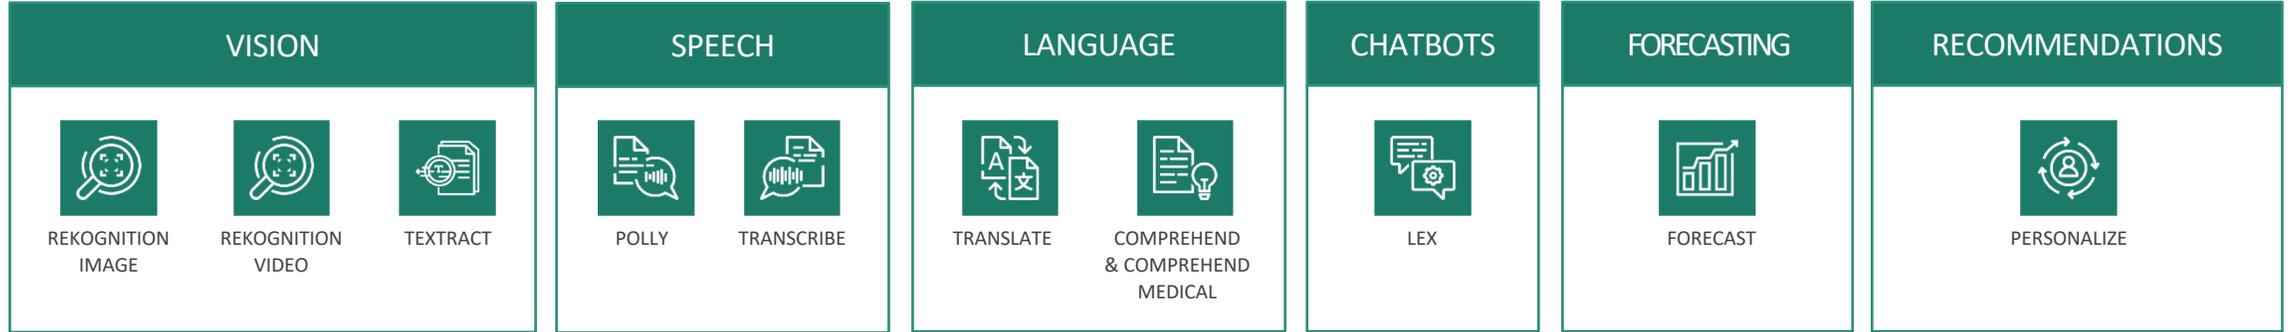
SIMPLE NOTIFICATION SERVICE (SNS)

- Send notifications by email, text, HTTP, Lambda
- “Pub-sub” model where a publisher publishes to a SNS topic, and subscribers subscribe to receive notifications

EVENTBRIDGE

- Some features were formerly called CloudWatch Events
- Used to built event-driven architectures (also a “pub-sub” model)
 - Subscribers set rules about what to receive
 - Schemas defined up front
 - Publishers can be third parties (e.g., Shopify, Zendesk)
 - Events can be scheduled

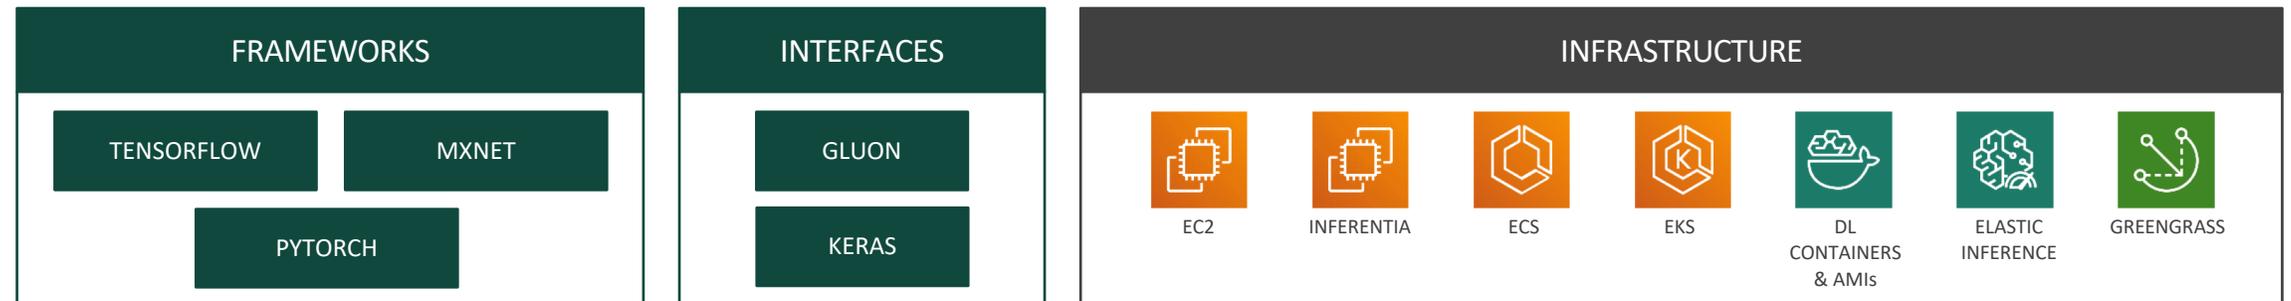
AI SERVICES



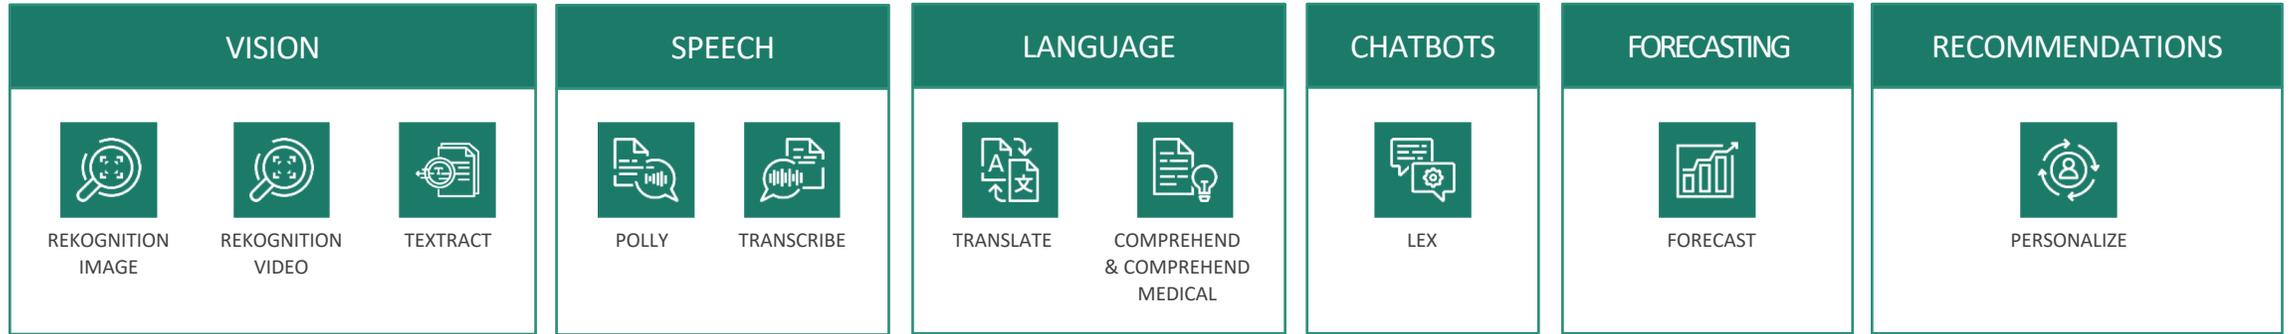
ML SERVICES



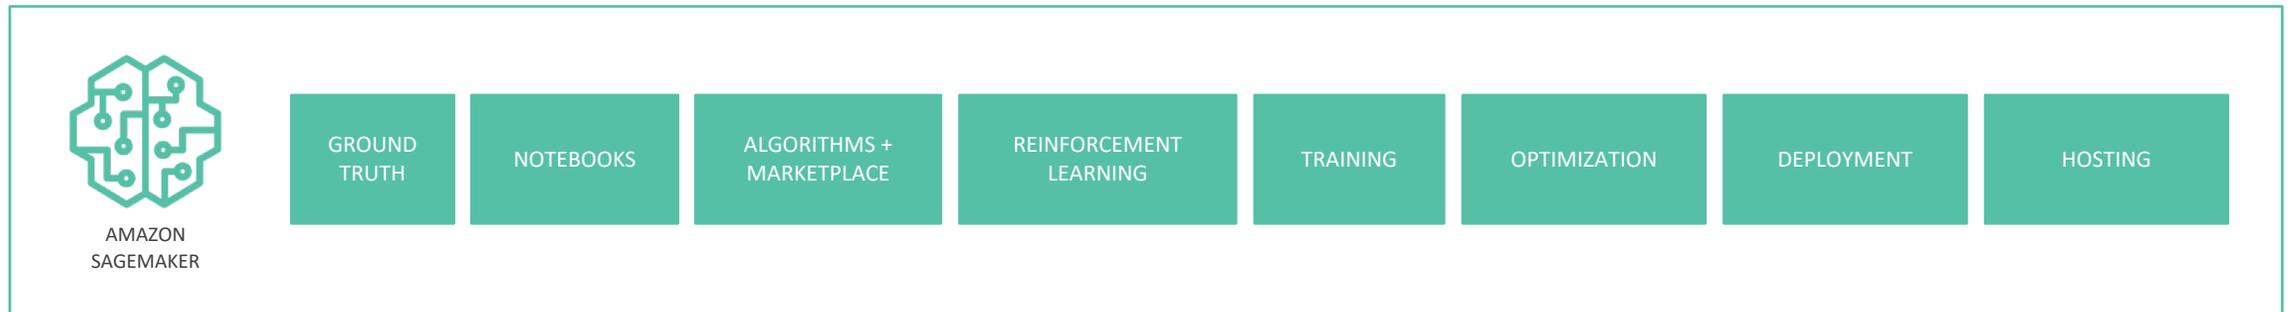
ML FRAMEWORKS + INFRASTRUCTURE



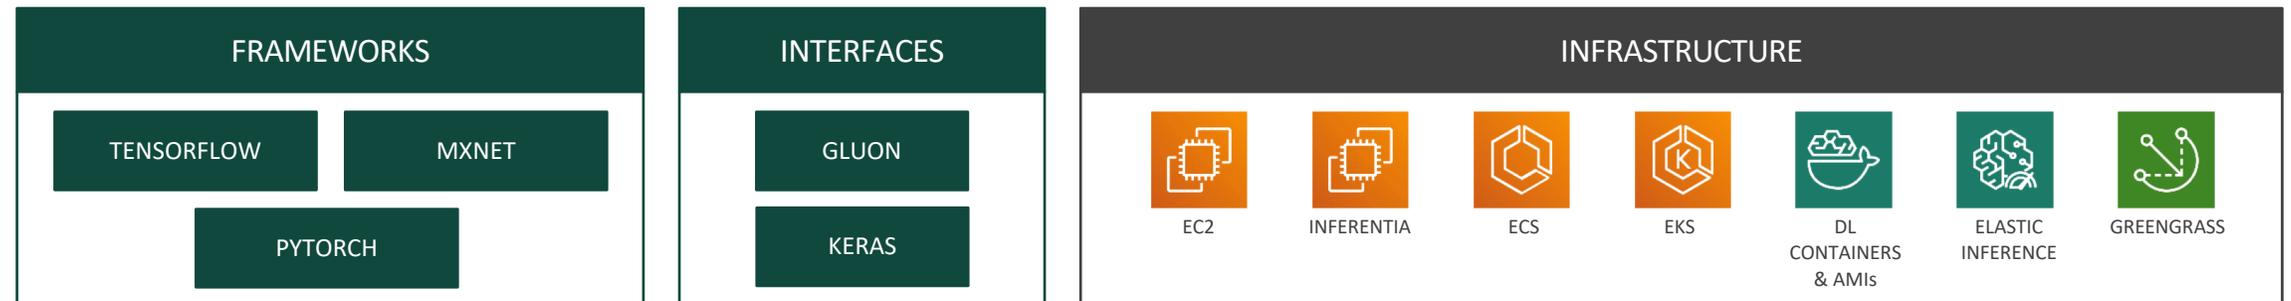
AI SERVICES



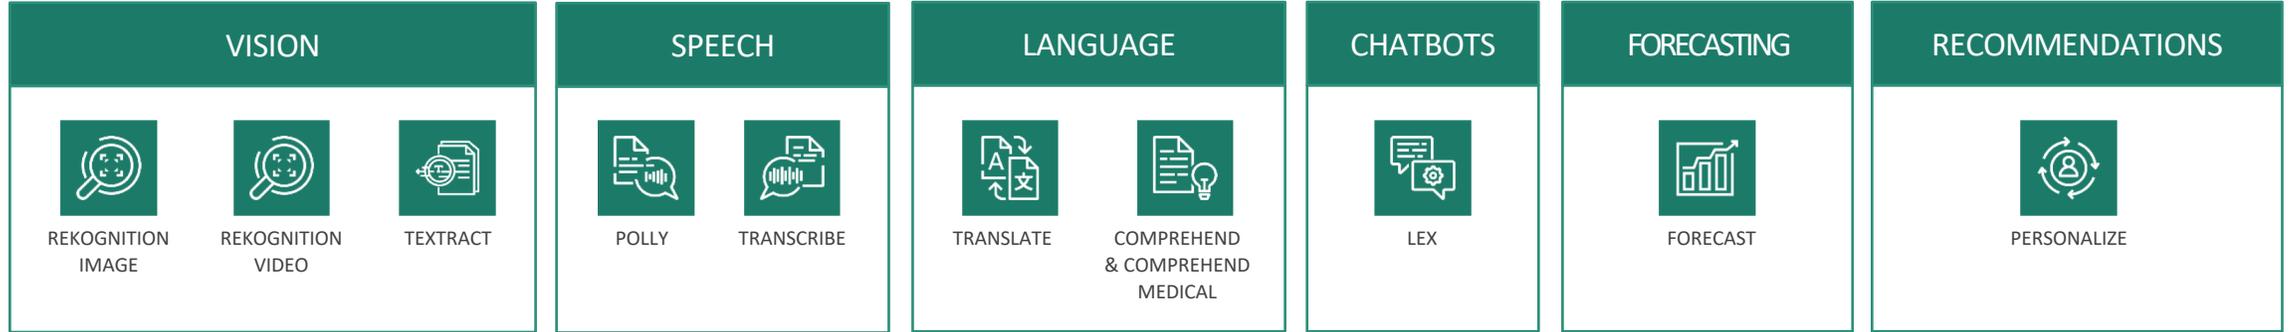
ML SERVICES



ML FRAMEWORKS + INFRASTRUCTURE



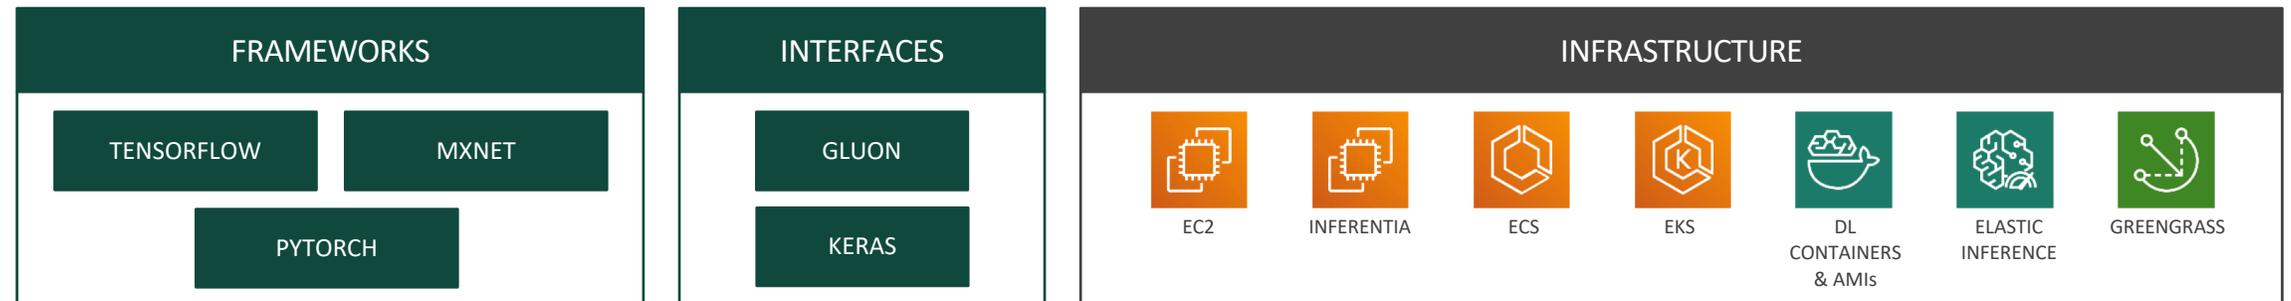
AI SERVICES



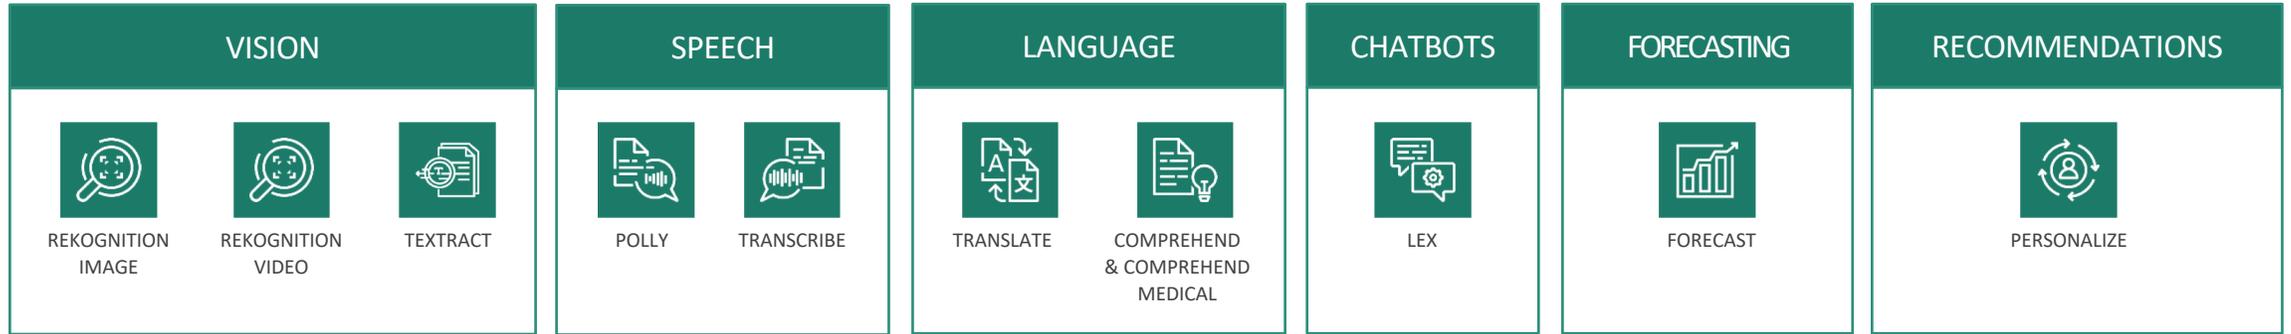
ML SERVICES



ML FRAMEWORKS + INFRASTRUCTURE



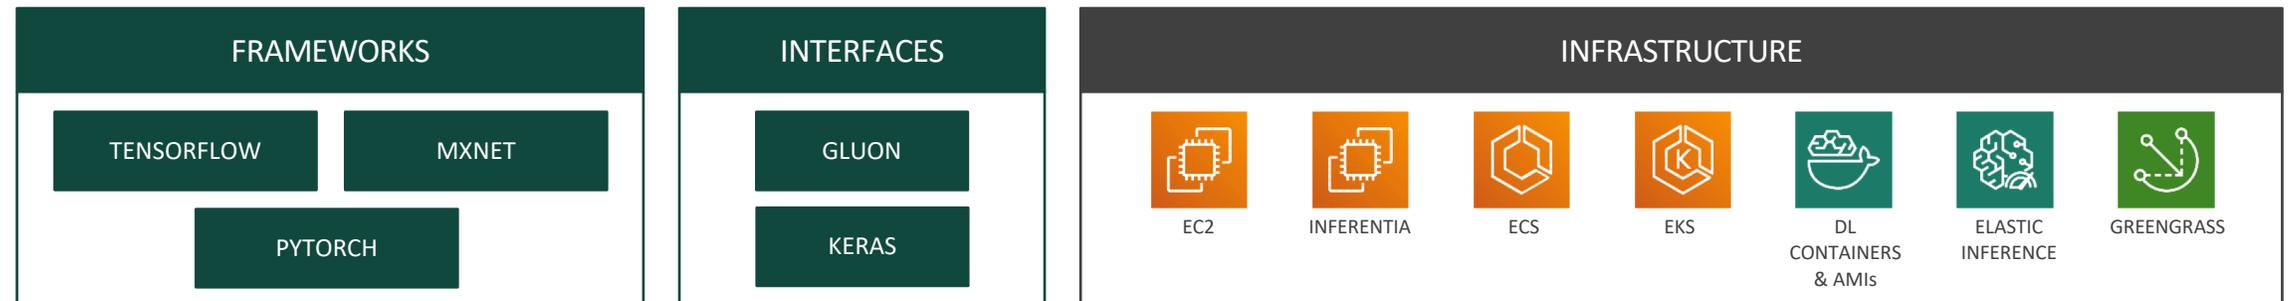
AI SERVICES



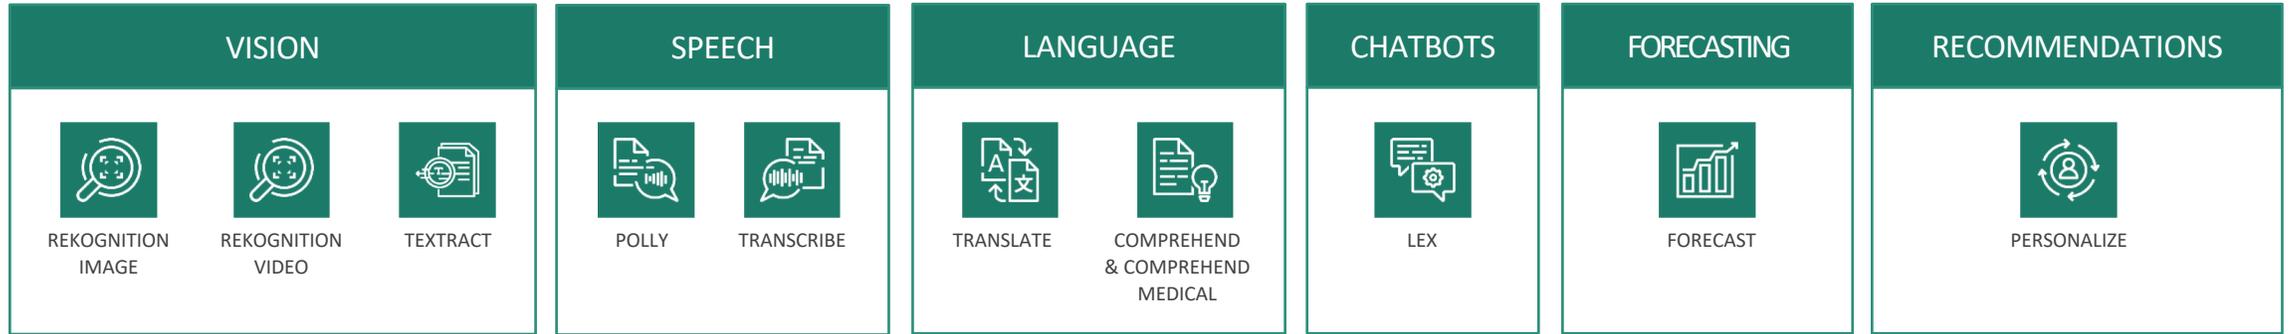
ML SERVICES



ML FRAMEWORKS + INFRASTRUCTURE



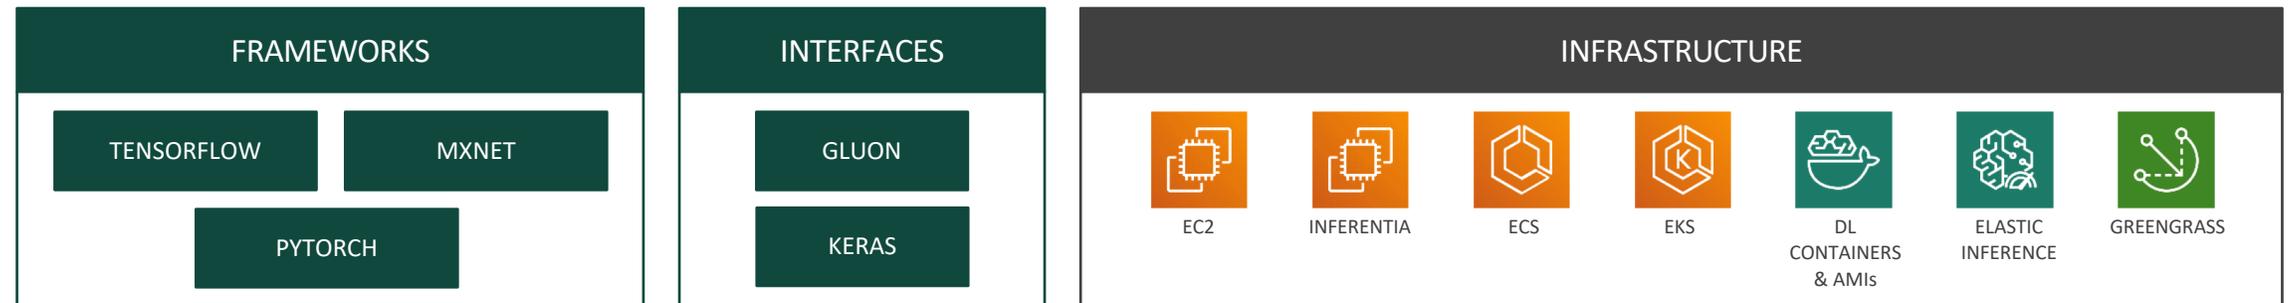
AI SERVICES



ML SERVICES



ML FRAMEWORKS + INFRASTRUCTURE





Rekognition Image

- Facial recognition
- Image search
- Sentiment analysis
- Image moderation
- Face-based verification

Rekognition Video

- Video search
- Video moderation

Textract

- Document import and processing
- Text extraction for natural language processing (NLP)



Polly

- Text-to-speech
- E-books and training for the visually impaired
- Announcement systems

Transcribe

- Speech-to-text
- Closed captioning and subtitles



Translate

- Natural-sounding language translation
- Localization of applications
- Multilingual sentiment analysis

Comprehend and Comprehend Medical

- Natural language process (NLP) of text
- Make sense of unstructured data such as emails, social media, reviews, and support tickets



Lex

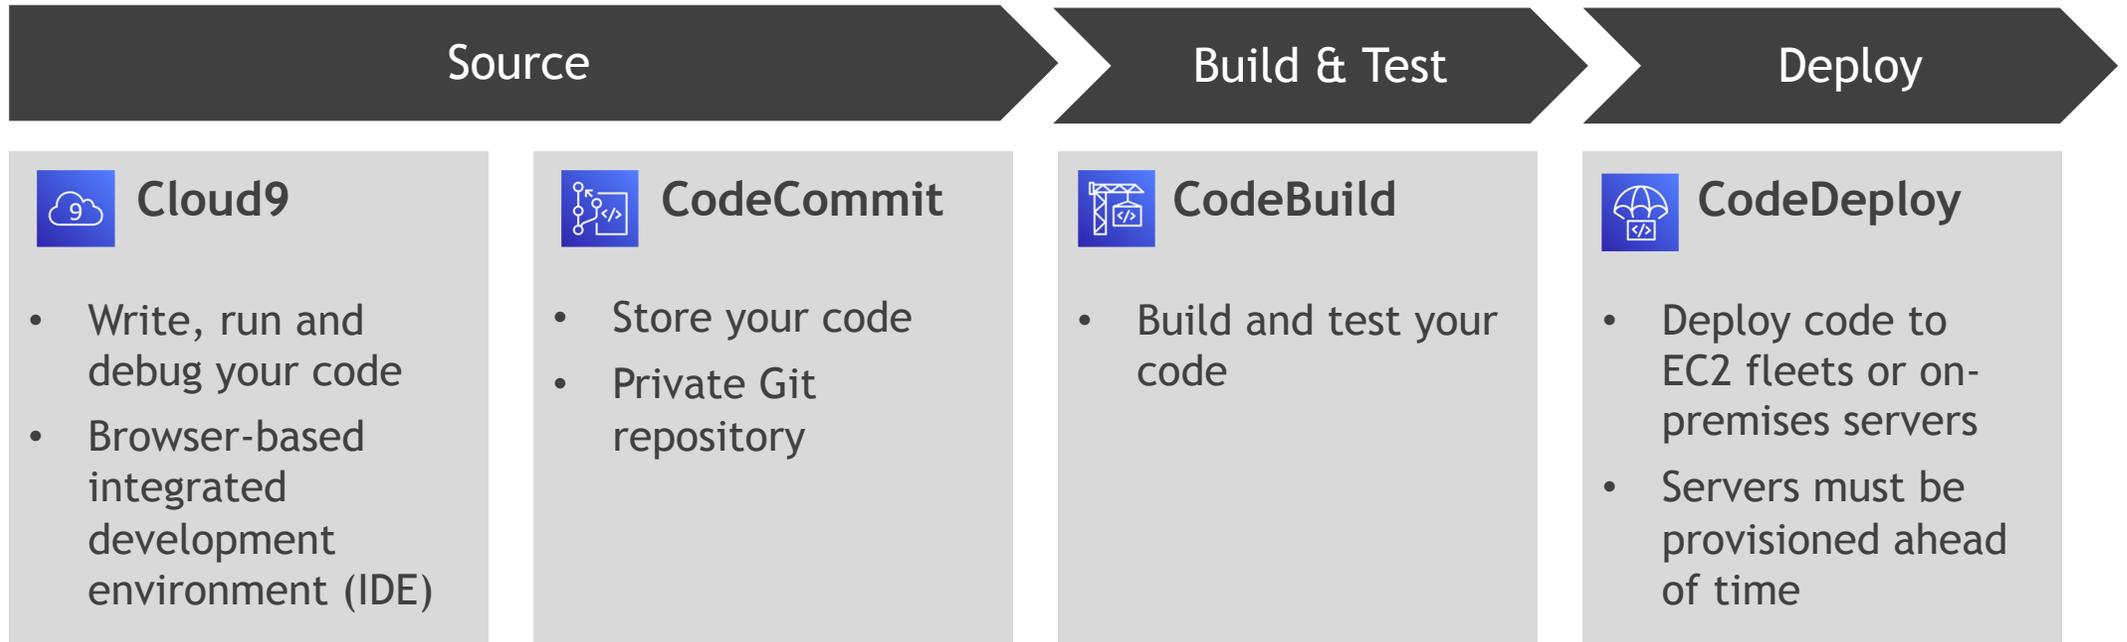
- Conversational interfaces, using the same technology as Amazon Alexa
- Call center bots
- Informational bots
- Productivity bots

Building and Deploying Software



CodePipeline

Automates and orchestrates these pieces:



The “rockSTAR” of the developer tools!



CodeStar

A unified interface to set up and work with all this:



CodePipeline

Automates and orchestrates these pieces:





DEMO

Working with CodeStar and Cloud9





AWS Trusted Advisor

Inspects your AWS environment and provides real-time recommendations based on best practices

- Cost optimization
- Performance
- Security
- Fault tolerance
- Service limits/quotas

SUPPORT PLAN	DETAILS
BASIC Free to all customers	<ul style="list-style-type: none"> • AWS Trusted Advisor service quota checks, limited security checks • Use of AWS Personal Health Dashboard • Access to whitepapers, documentation and support communities
DEVELOPER \$ For experimenting and testing	<ul style="list-style-type: none"> • AWS Trusted Advisor service quota and security checks • Business hours email address to Cloud Support Associates • Unlimited cases, one primary contact <ul style="list-style-type: none"> • General guidance < 24 hours; System impaired < 12 hours
BUSINESS \$\$ Minimum recommendation minimum for production	<ul style="list-style-type: none"> • All Trusted Advisor checks • 24x7 phone, email, and chat access to Cloud Support Engineers • Unlimited cases, unlimited contacts <ul style="list-style-type: none"> • General guidance < 24 hours; System impaired < 12 hours • Production system impaired < 4 hours; Production system down < 1 hour
ENTERPRISE \$\$\$ Recommended for mission-critical workloads	<ul style="list-style-type: none"> • All Trusted Advisor checks • 24x7 phone, email, and chat access to Cloud Support Engineers • Access to a Technical Account Manager (TAM) • Concierge Support Team for billing and account questions • Unlimited cases, unlimited contacts <ul style="list-style-type: none"> • General guidance < 24 hours; System impaired < 12 hours • Production system impaired < 4 hours; Production system down < 1 hour • Business/mission-critical system down < 15 minutes

Additional Resources

AWS Professional Services

- AWS's consulting service
- aws.amazon.com/professional-services

AWS Knowledge Center

- FAQs
- aws.amazon.com/premiumsupport/knowledge-center

Whitepapers

- aws.amazon.com/whitepapers

Blog

- aws.amazon.com/blogs/aws