AWS Bootcamp

Vince Lo Faso, Solutions Architect

Scott Turvey, Solutions Architect

Agenda

8:30am – 9:00am: Check-in

9:00am – 12:00pm: AWS Bootcamp

- Bootcamp Introduction
- Introduction to Cloud Computing
- AWS Overview and Regions
- AWS Networking
  - VPC - Route 53
- AWS Compute
  - EC2 – ELB - ASG – IAM - Directory Services
- AWS Storage
  - EBS - S3 - Glacier
- AWS Databases
  - RDS – DynamoDB - Redshift
- Security Overview

12:00pm – 1:00pm: Break & Lunch

1:00pm – 3:00pm: Hands-On Lab - Building AWS Infrastructure

Bootcamp Introductions

Name:

Company:

Title/Work Area:

What was the 1st computer you used?

Introduction to Cloud Computing

IT and the Business

A man is flying in a hot air balloon and realizes he is lost. He reduces height and spots a man down below. He lowers the balloon further and shouts:

"Excuse me, can you tell me where I am?"

The man below says: "yes, you're in a hot air balloon hovering 30 feet above this field."

"You must work in the Information Technology," says the balloonist.

"I do," replies the man, "How did you know."

"Well" says the balloonist,, "everything you have told me is technically correct, but it's of no use to anyone."

The man below says, "you must work in business".

"I do" replies the balloonist, "but how did you know?"

"Well", says the man, ,"you don't know where you are, or where you're going, but you expect me to be able to help. You're in the same position you were before we met, but now it's my fault."

New IT Business Model

Cloud Computing

is *first and foremost* a

Business Model

Business Reasons for Adopting Cloud Computing

Not Good

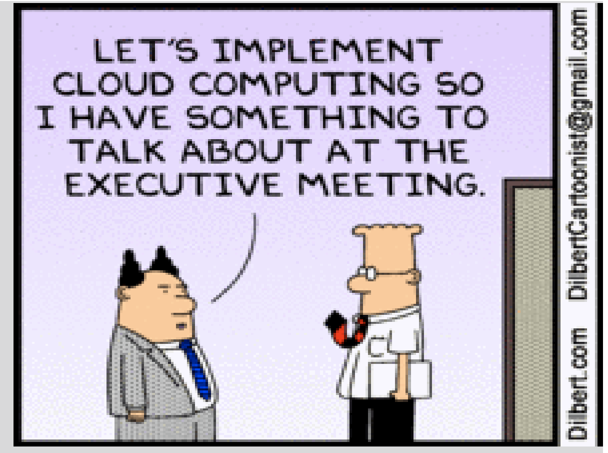
LET'S IMPLEMENT CLOUD COMPUTING SO I HAVE SOMETHING TO TALK ABOUT AT THE EXECUTIVE MEETING.
DilbertCartoonist@gmail.com
Dilbert.com

Good

- Move from capital expense to variable expense
- Increased agility
- Lower variable expense than they could achieve on their own
- Stop guessing capacity
- Remove undifferentiated heavy lifting
- Go global in minutes

Defining Cloud Computing

National Institute of Standards and Technology
U.S. Department of Commerce

NIST defined a well accepted, industry standard definition of Cloud Computing

url: http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf

Covers:

- 5 Key Characteristics of Cloud Computing
- 3 Service Model
- 4 Deployment Models

  plus
- 5 Cloud Actors
- A Cloud Reference Architecture
- Shared Security model

What is Cloud Computing?

NIST 5 Key Characteristics

#1 On-demand self service

“as easy as buying candy from a vending machine”

#5 Measured service

“pay only for what you consume”

#2 Broad network access

“access it anytime from anywhere”

#4 Rapid elasticity

“scale up and scale down in real-time”

#3 Resource pooling

“you’re not the only user”

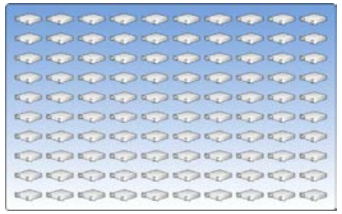

Why Amazon Web Services?

2011

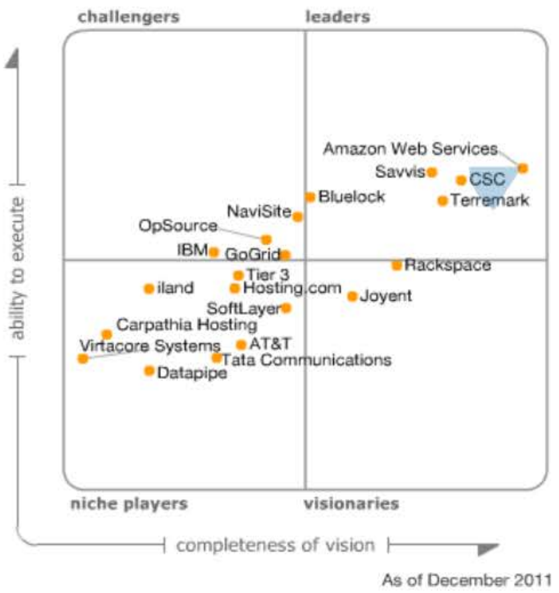

2012

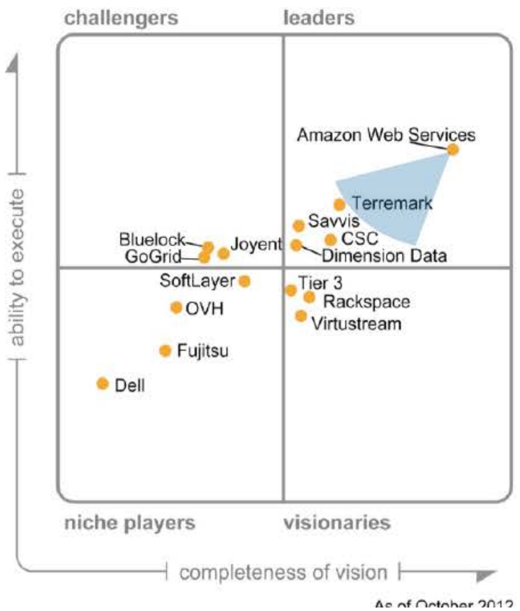

2013

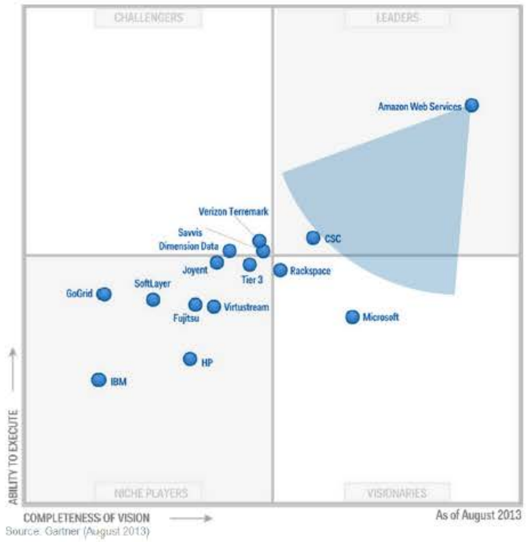

Source: Gartner (August 2013)

In 2013, IT research firm Gartner had to rescale its famed infrastructure-as-a- service "Magic Quadrant" to accommodate Amazon Web Services' enormous competitive lead.

"It is the overwhelming market share leader, with over *10 times more* cloud IaaS compute capacity in use *than the aggregate total of the other 14 providers* in this Magic Quadrant" Gartner Report May, 2015.

AWS Today - 2016

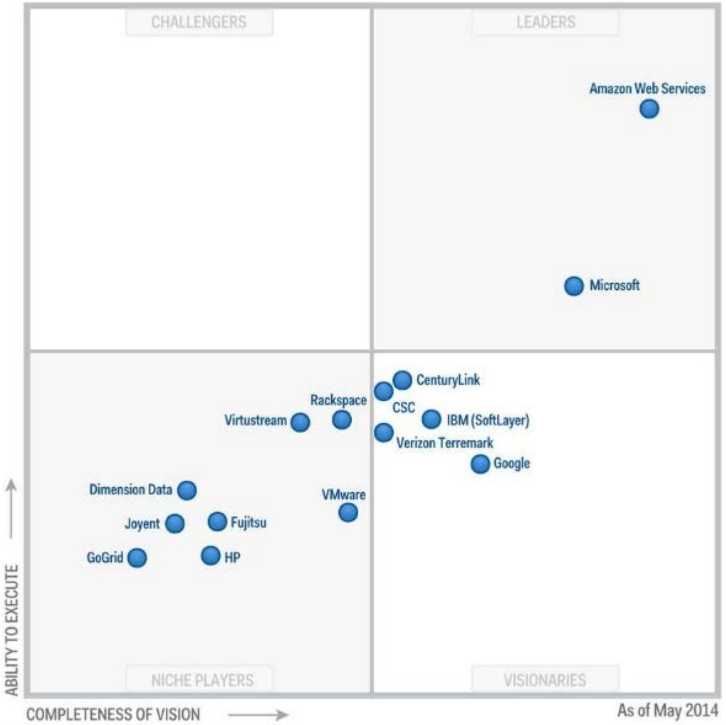

Source: Gartner (May 2014)

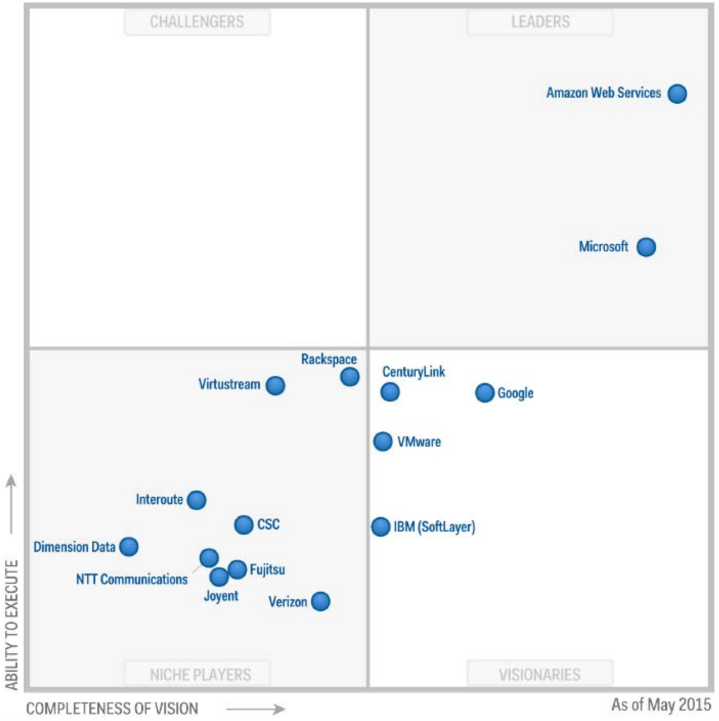

Source: Gartner (May 2015)

Magic Quadrant for Cloud Infrastructure as a Service

Public cloud market share

Amazon Web Services:	31%
Microsoft:	9%
IBM Cloud/SoftLayer:	7%
Google:	4%

per Synergy Research Group

AWS generating > $12 billion a year.

AWS Services and Regions Overview

AWS Services

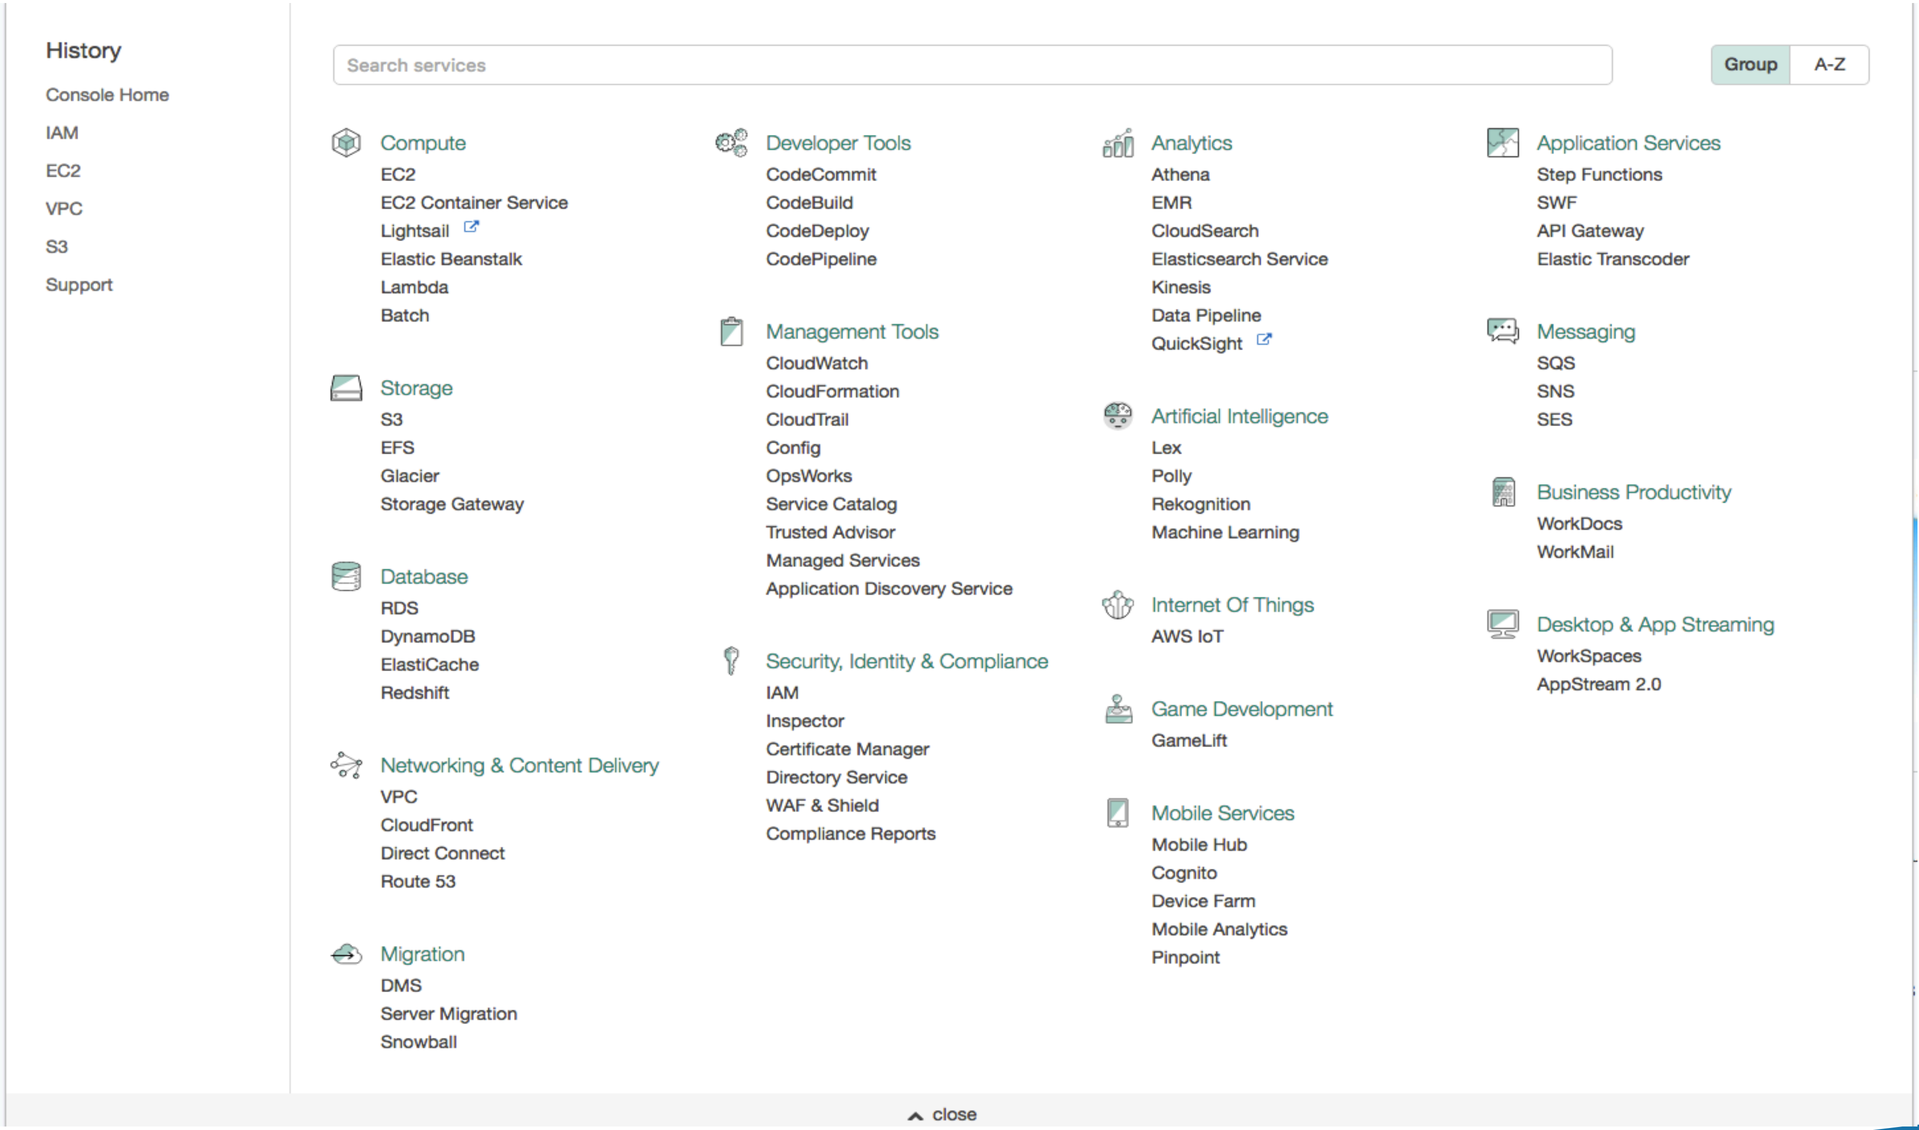

History

- Console Home
- IAM
- EC2
- VPC
- S3
- Support

Search services

Group A-Z

Compute
EC2
EC2 Container Service
Lightsail
Elastic Beanstalk
Lambda
Batch

Storage
S3
EFS
Glacier
Storage Gateway

Database
RDS
DynamoDB
ElastiCache
Redshift

Networking & Content Delivery
VPC
CloudFront
Direct Connect
Route 53

Migration
DMS
Server Migration
Snowball

Developer Tools
CodeCommit
CodeBuild
CodeDeploy
CodePipeline

Management Tools
CloudWatch
CloudFormation
CloudTrail
Config
OpsWorks
Service Catalog
Trusted Advisor
Managed Services
Application Discovery Service

Security, Identity & Compliance
IAM
Inspector
Certificate Manager
Directory Service
WAF & Shield
Compliance Reports

Analytics
Athena
EMR
CloudSearch
Elasticsearch Service
Kinesis
Data Pipeline
QuickSight

Artificial Intelligence
Lex
Polly
Rekognition
Machine Learning

Internet Of Things
AWS IoT

Game Development
GameLift

Mobile Services
Mobile Hub
Cognito
Device Farm
Mobile Analytics
Pinpoint

Application Services
Step Functions
SWF
API Gateway
Elastic Transcoder

Messaging
SQS
SNS
SES

Business Productivity
WorkDocs
WorkMail

Desktop & App Streaming
WorkSpaces
AppStream 2.0

close

Core Services

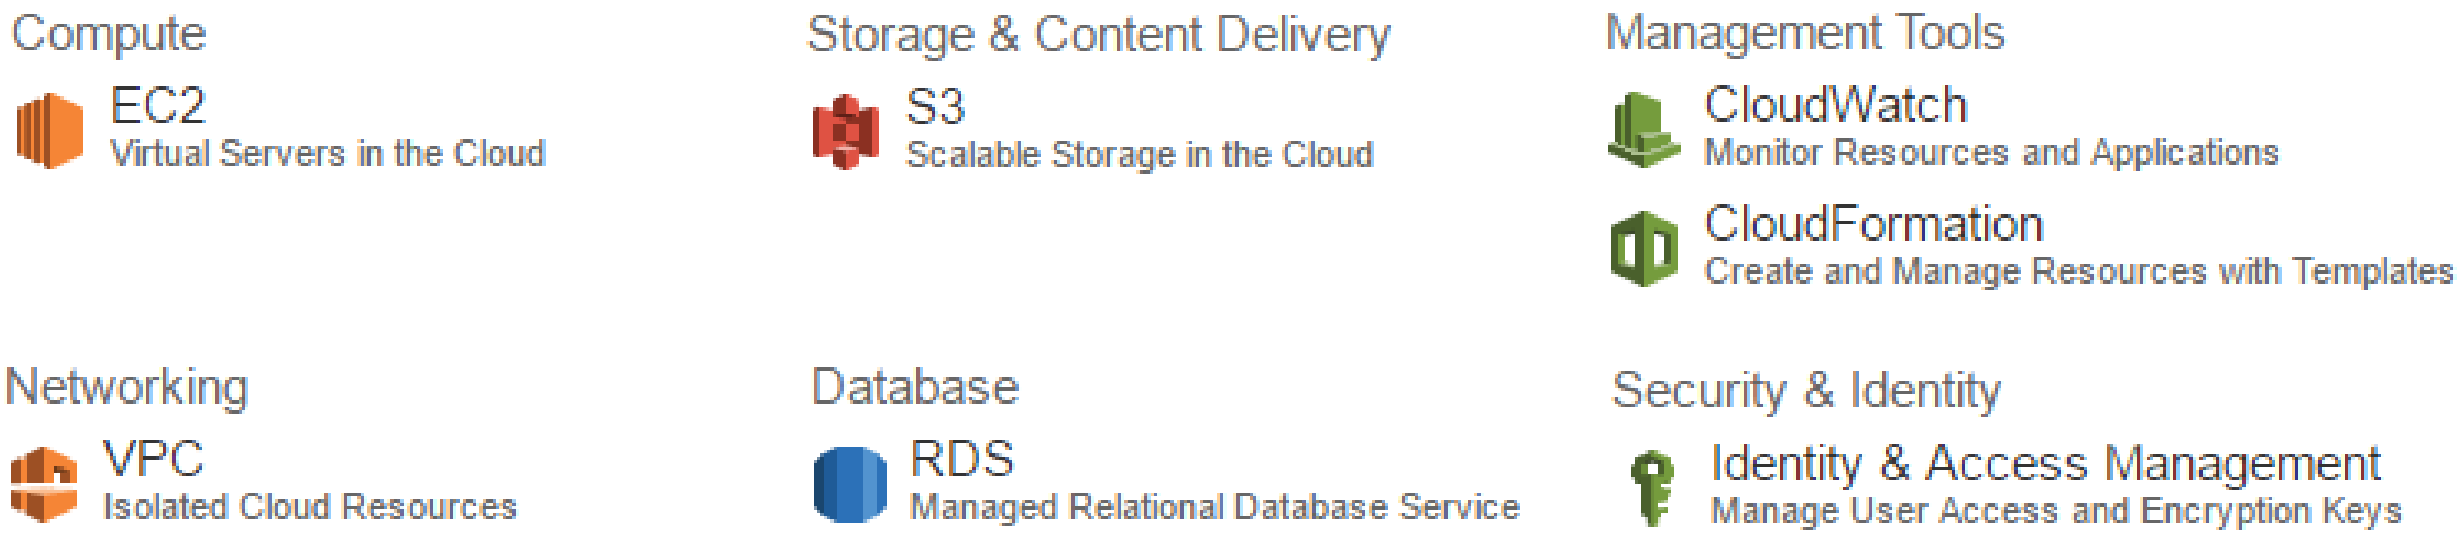

Compute

EC2
Virtual Servers in the Cloud

Networking

VPC
Isolated Cloud Resources

Storage & Content Delivery

S3
Scalable Storage in the Cloud

Database

RDS
Managed Relational Database Service

Management Tools

CloudWatch
Monitor Resources and Applications

CloudFormation
Create and Manage Resources with Templates

Security & Identity

Identity & Access Management
Manage User Access and Encryption Keys

AWS Regions – Global Infrastructure

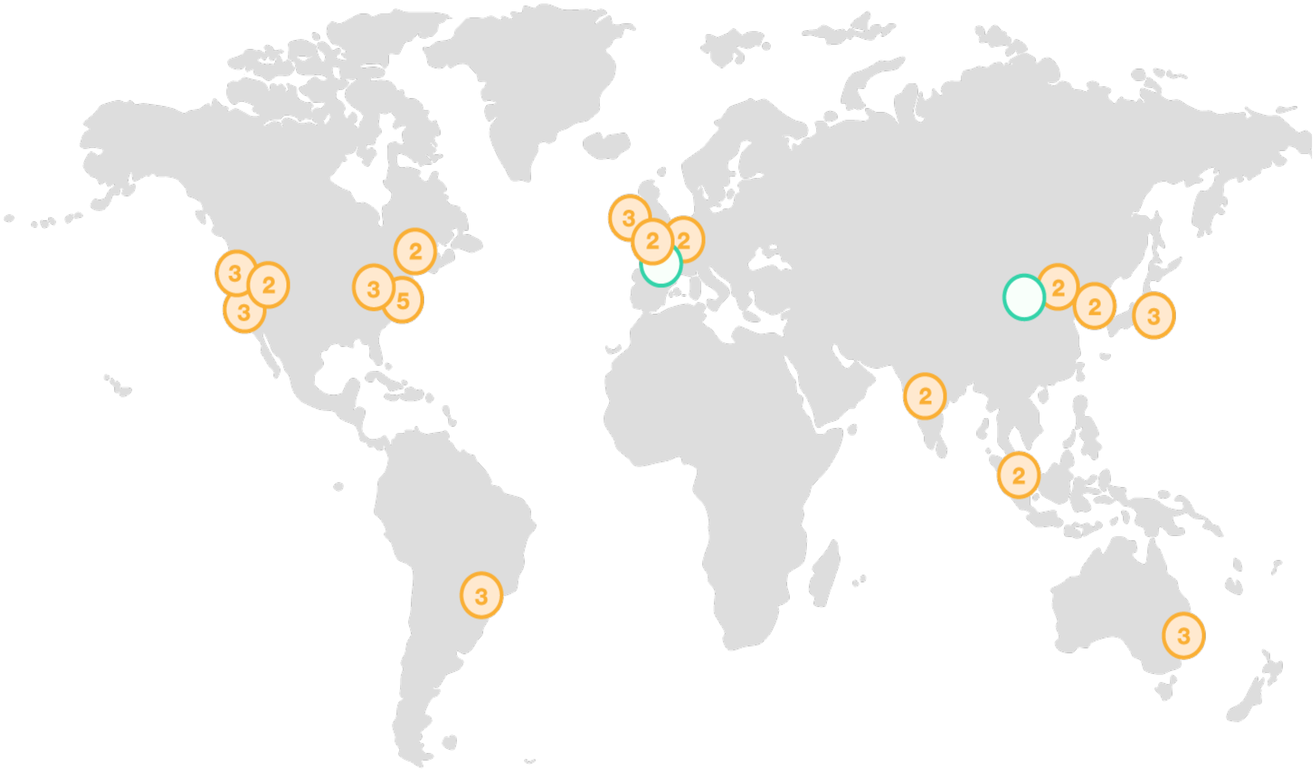

#

Region & Number of Availability Zones

AWS GovCloud (2)

US West
Oregon (3), Northern California (3)

US East
Northern Virginia (5), Ohio (3)

Canada
Central (2)

South America
São Paulo (3)

Europe
Ireland (3), Frankfurt (2), London (2)

Asia Pacific
Singapore (2), Sydney (3), Tokyo (3), Seoul (2), Mumbai (2)

China
Beijing (2)

○

New Region (coming soon)

Paris

Ningxia

Regions and Availability Zones

- **Global Resources**
  - » IAM Users
  - » Route 53 Records
- **Regional Resources**
  - » S3 Buckets
  - » VPCs
  - » ELB
  - » EIPs
- **AZ Resources**
  - » EBS Volumes
  - » EC2 Instances
  - » RDS Instances
  - » Subnets
  - » ENIs

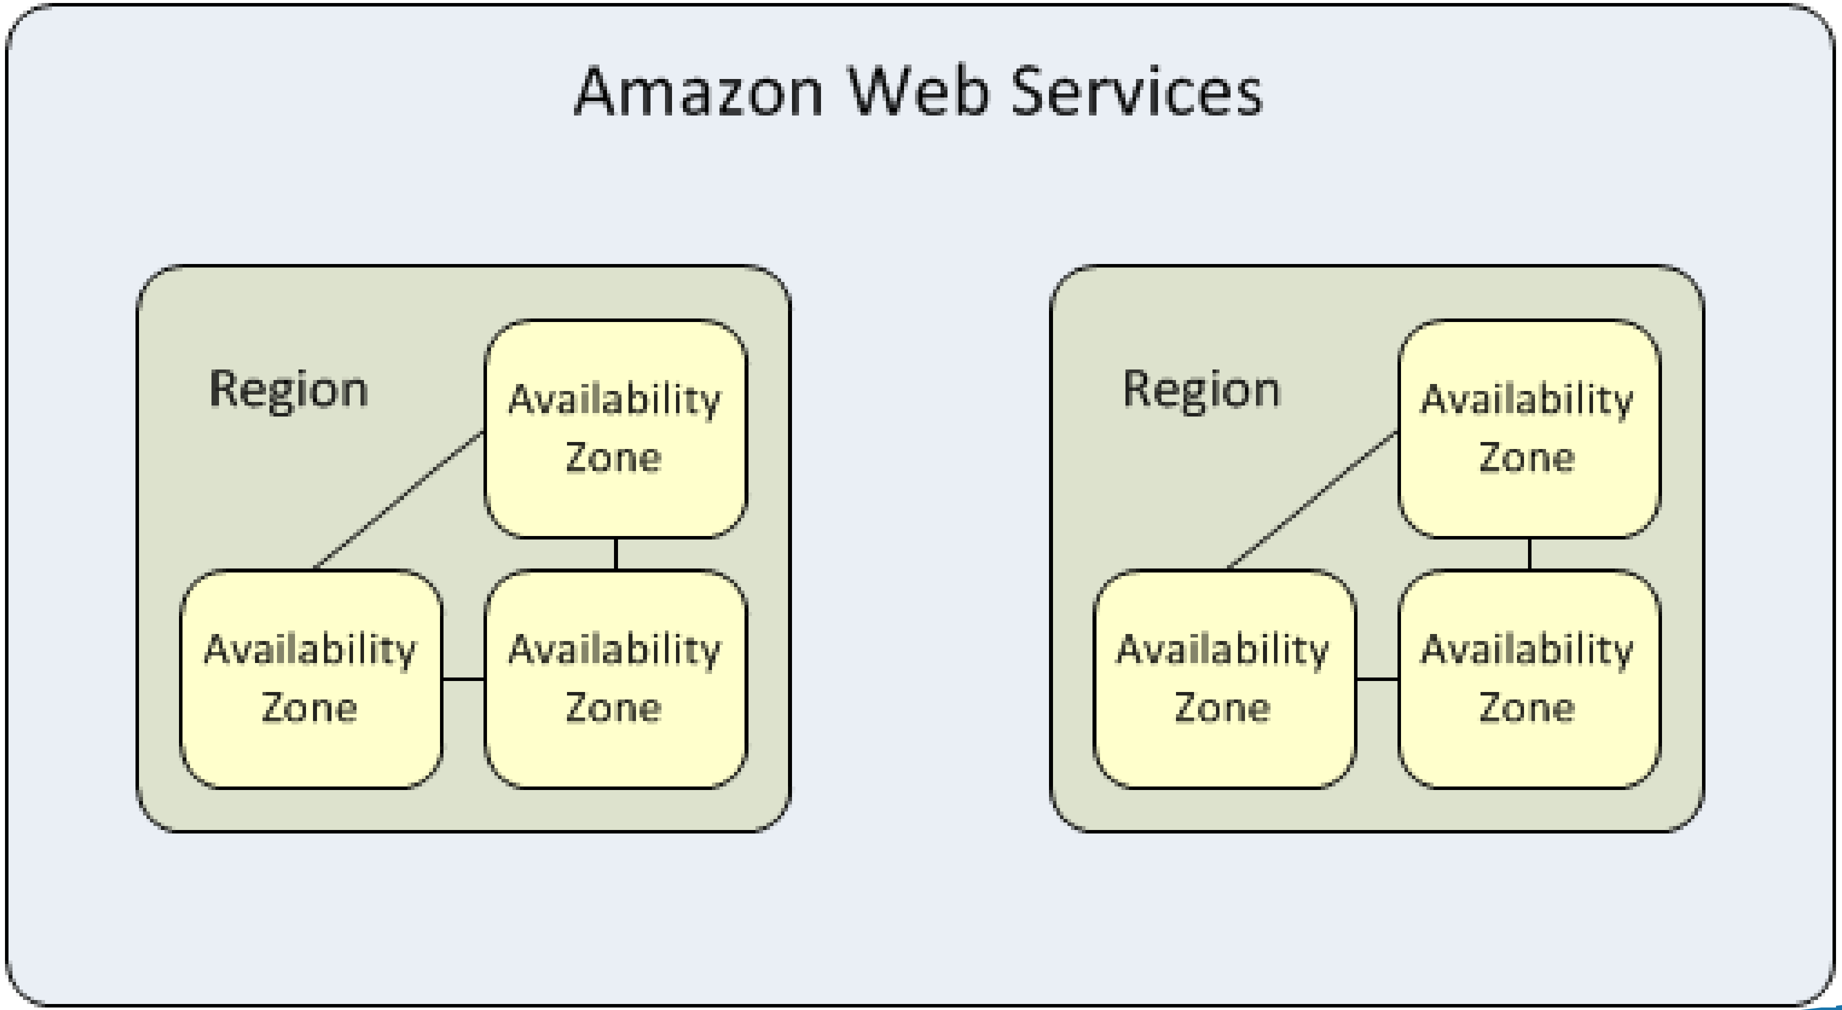

AWS Networking Overview

VPC - Virtual Private Cloud

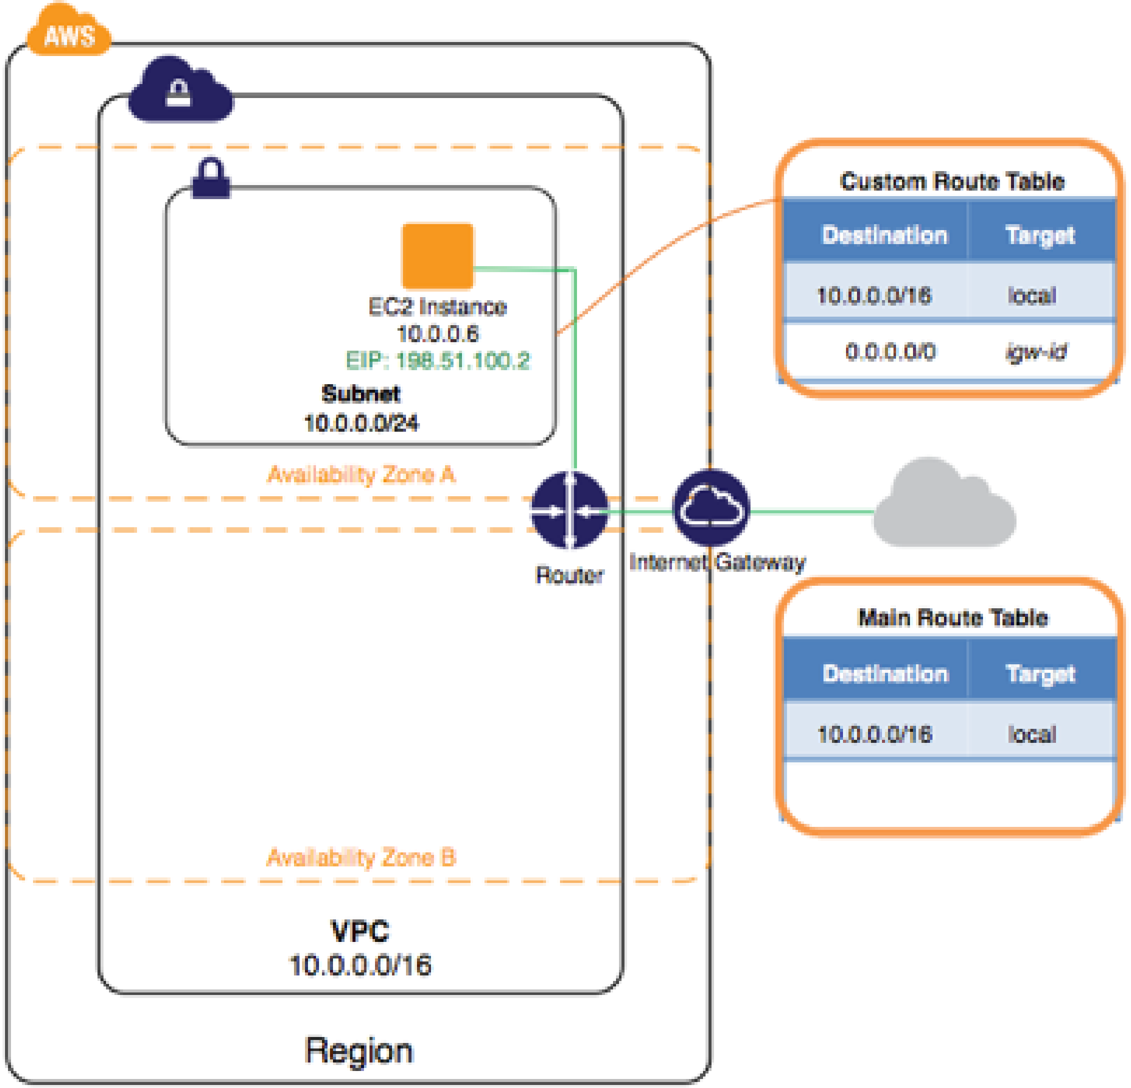

Custom Route Table

Destination	Target
10.0.0.0/16	local
0.0.0.0/0	*igw-id*

Main Route Table

Destination	Target
10.0.0.0/16	local

VPC - Virtual Private Cloud

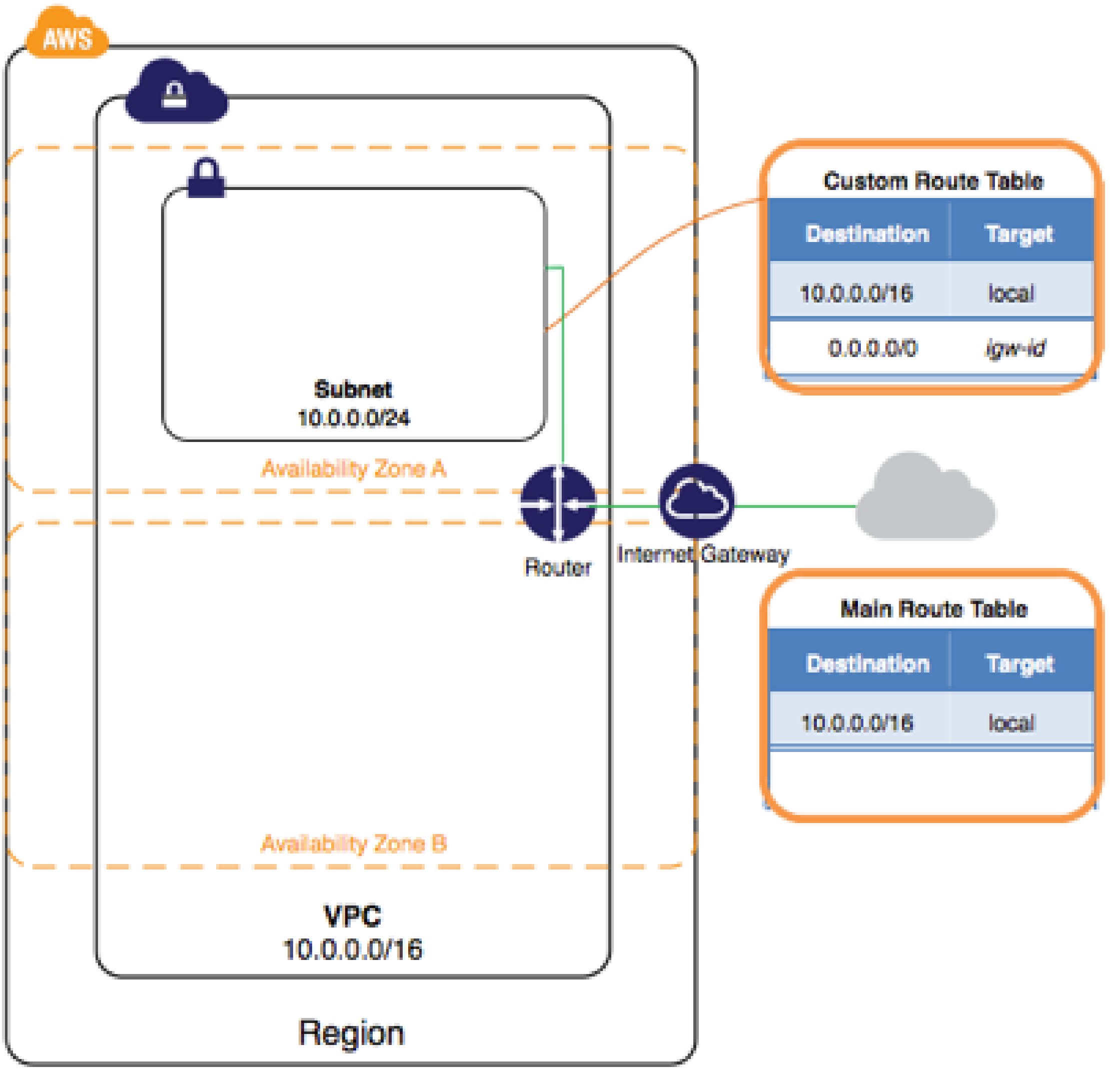

Custom Route Table

Destination	Target
10.0.0.0/16	local
0.0.0.0/0	*igw-id*

Main Route Table

Destination	Target
10.0.0.0/16	local

Virtual Private Cloud – NACL – Traffic Flow

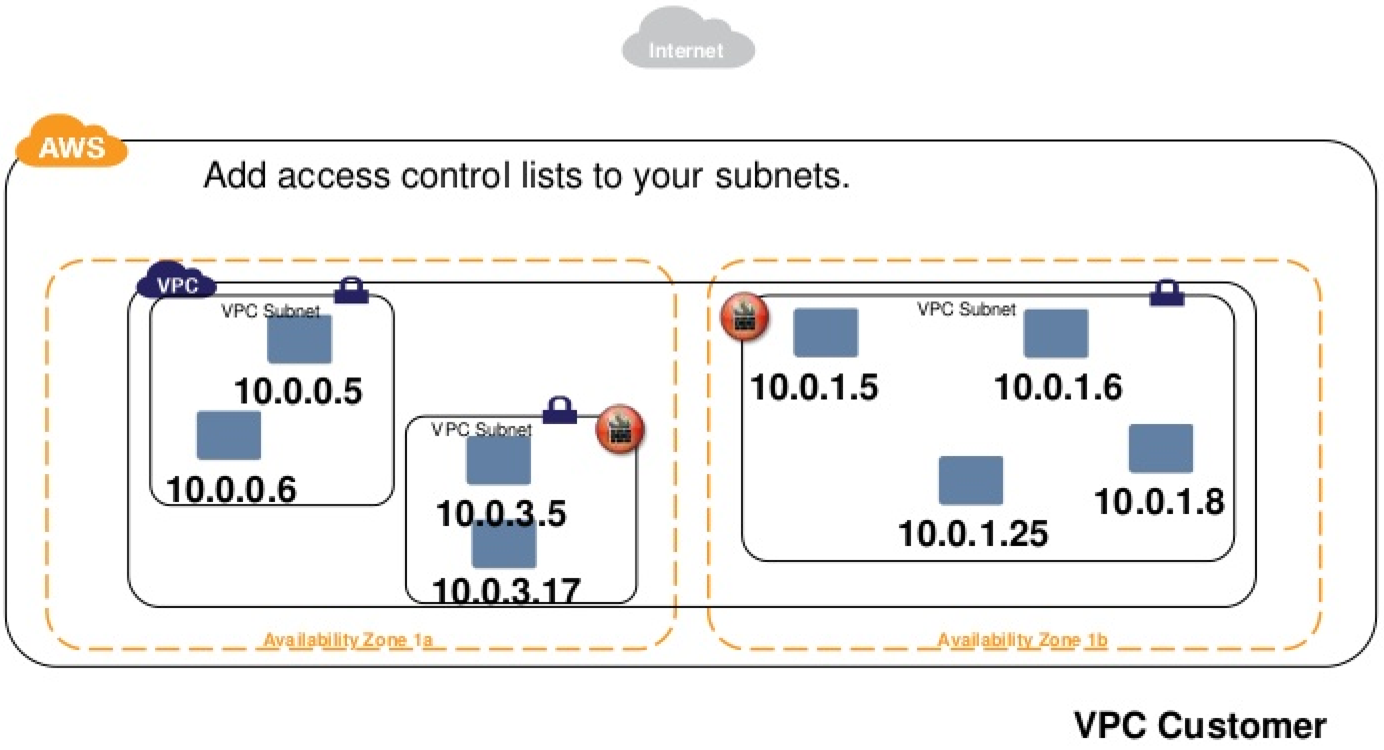

Security Groups

- Security Groups are similar to a firewall rule
- They can be associated to resources independent of a subnet or CIDR range
- Security Groups are limited only to the VPC in which you create them

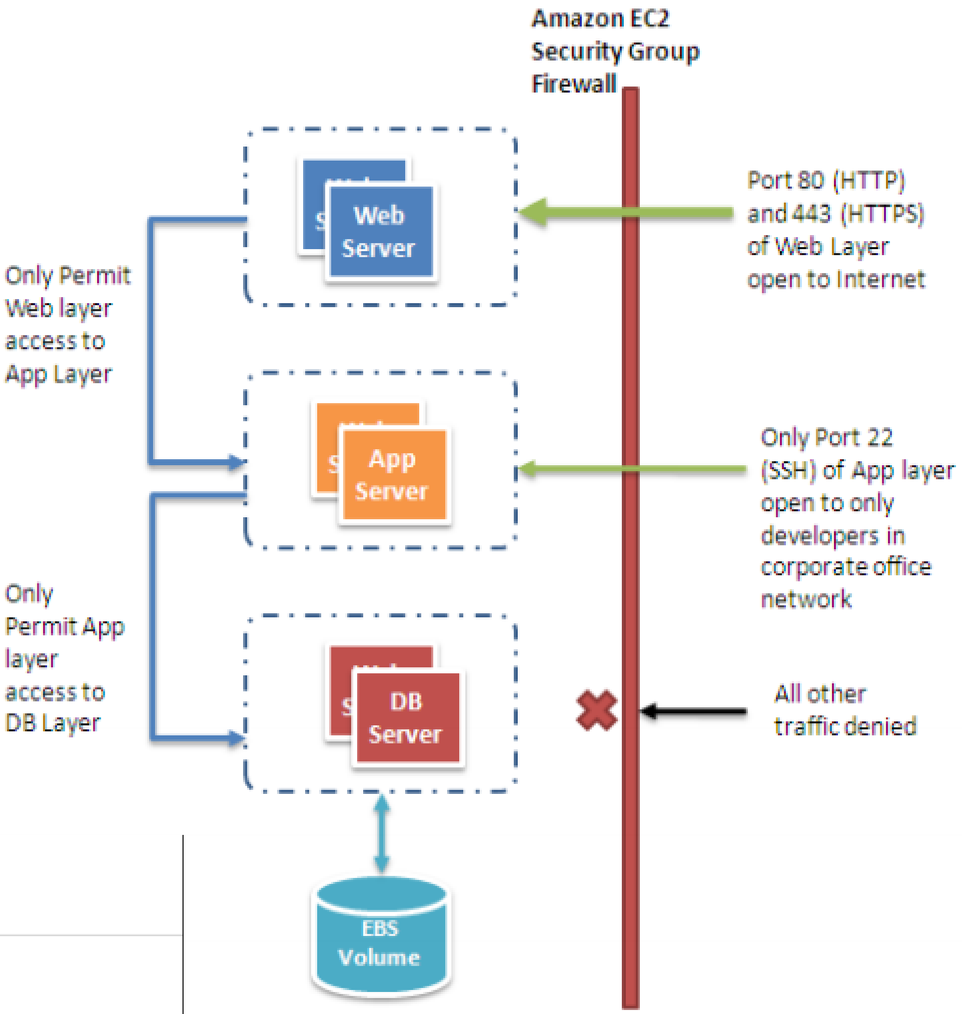

Example of Security Group configuration

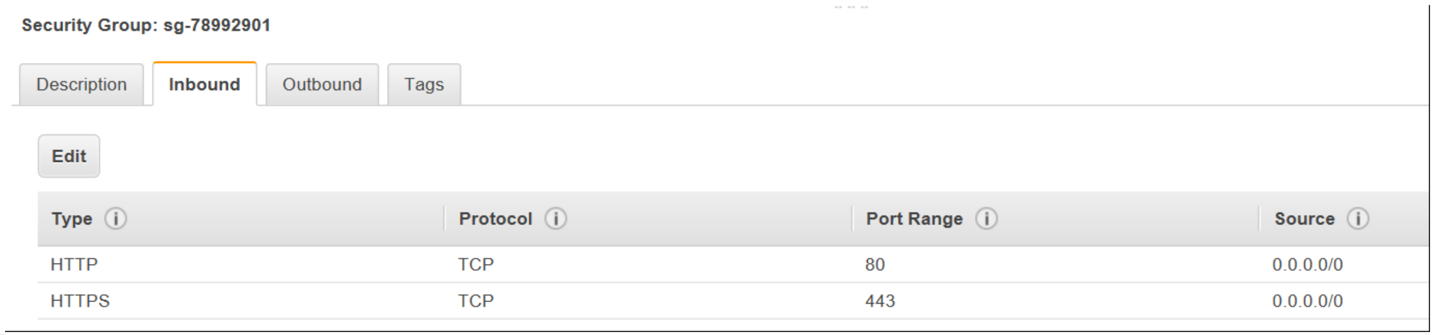

Security Group: sg-78992901

Description | **Inbound** | Outbound | Tags

Edit

Type	Protocol	Port Range	Source
HTTP	TCP	80	0.0.0.0/0
HTTPS	TCP	443	0.0.0.0/0

Security Groups

Deny by default

- IP Whitelisting
  - Specify a CIDR block that is allowed to access resources in your AWS environment.
  - This can be as large or small as you desire, giving it extreme flexibility.
  - Specifying a 32 bit block will whitelist a single IP (50.99.20.230/32)
- Allow port and protocol
  - You can allow TCP, UDP, ICMP or a combination of all three

Security Groups

- SG trust relationships
  - SGs can establish trust relationships
    - These trust relationships link resources and security policies
  - Not required to specify an IP address
  - Trust relationships are only valid within a VPC
- Ingress and Egress
  - VPC security groups have both ingress and egress
  - Security groups are stateful

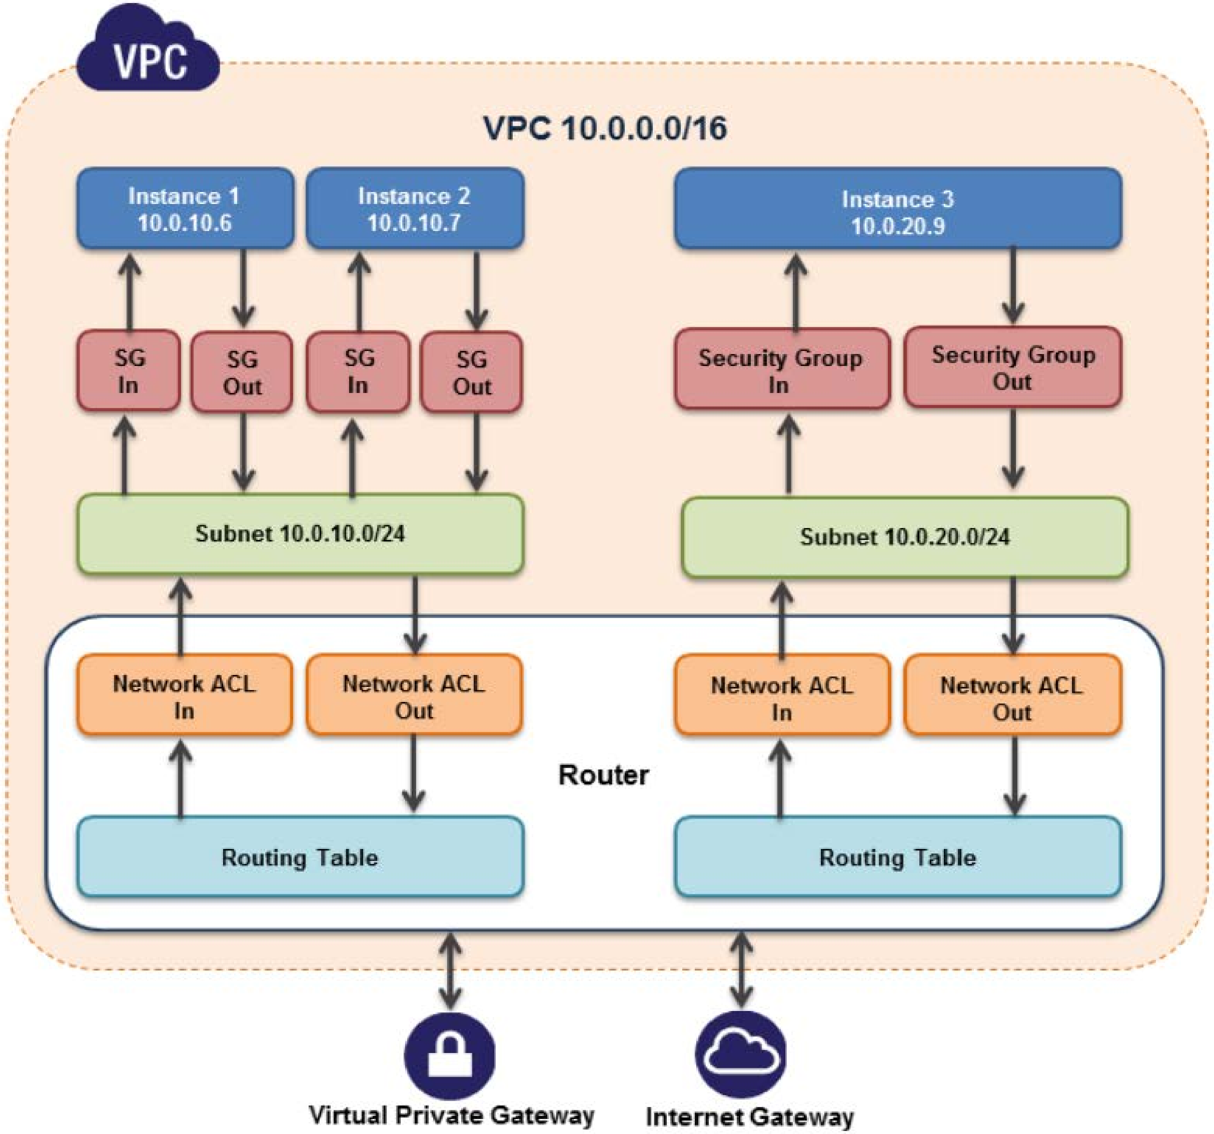

Security Group/Firewall Rules

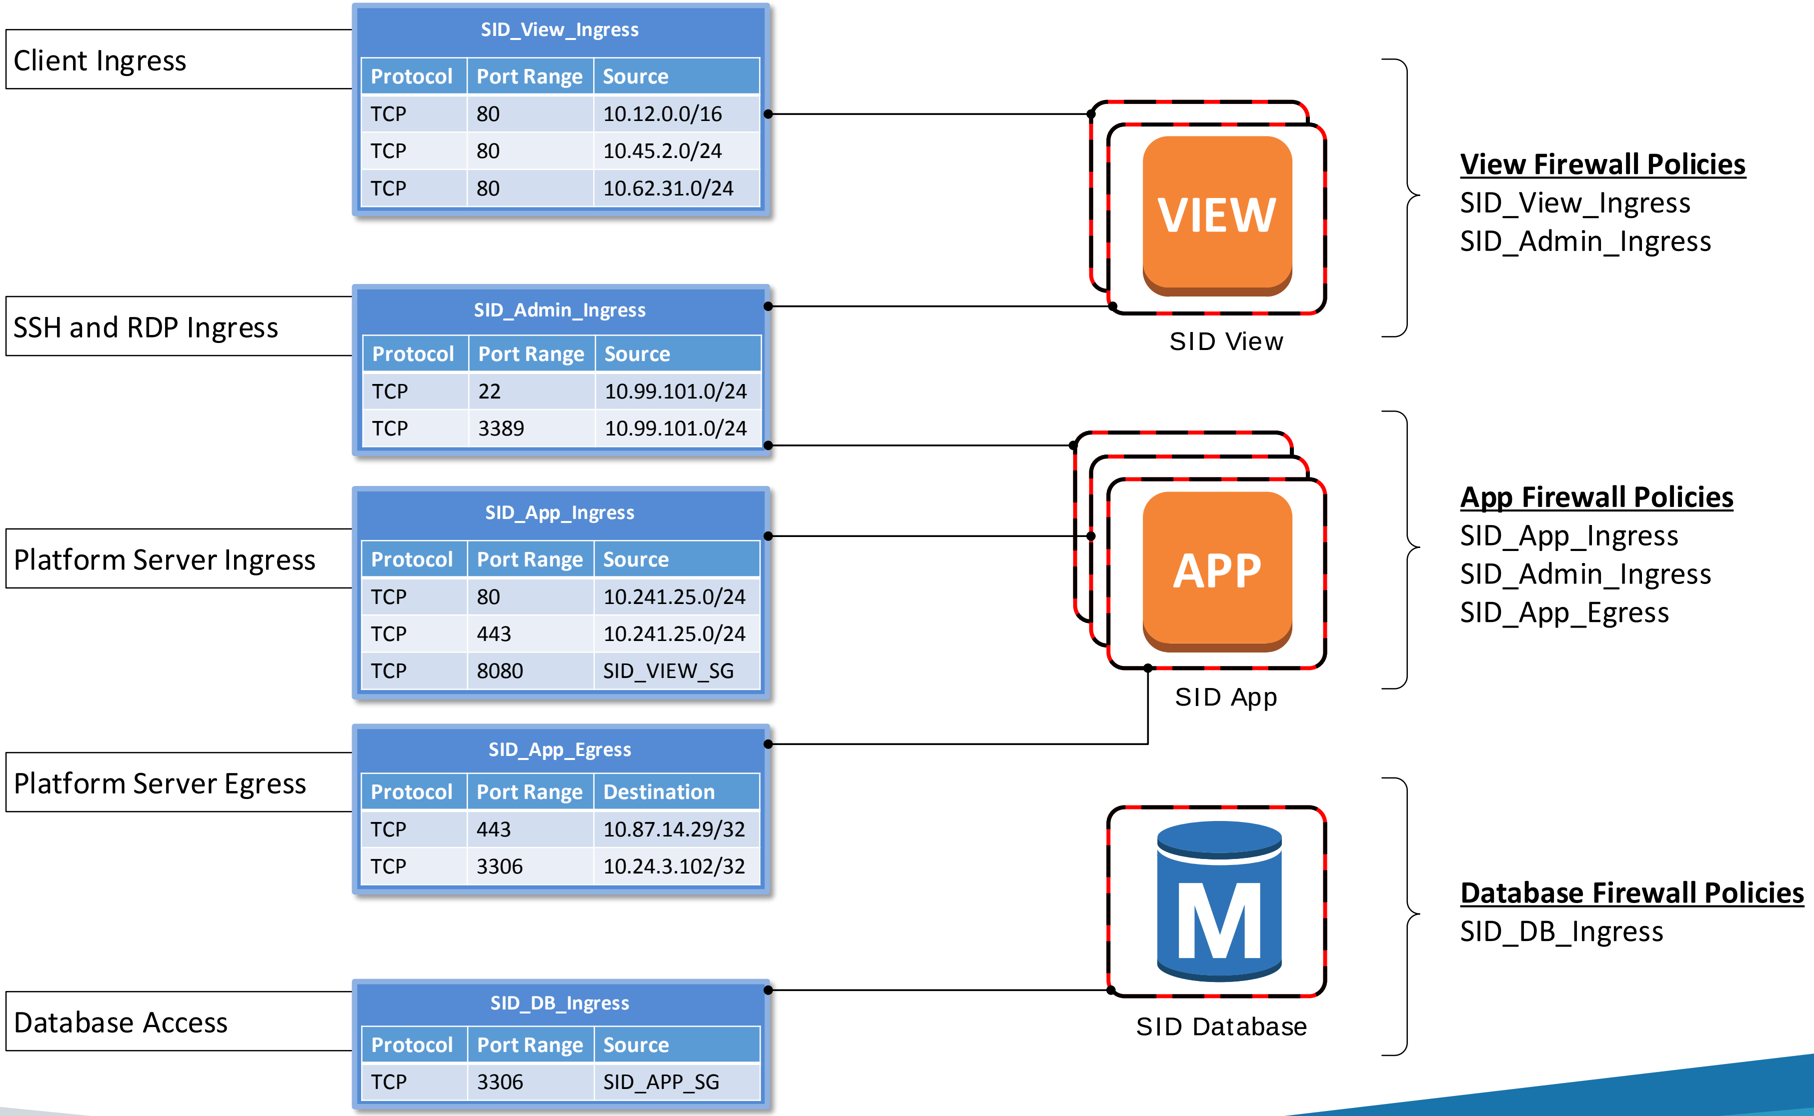

Client Ingress

SID_View_Ingress

Protocol	Port Range	Source
TCP	80	10.12.0.0/16
TCP	80	10.45.2.0/24
TCP	80	10.62.31.0/24

SSH and RDP Ingress

SID_Admin_Ingress

Protocol	Port Range	Source
TCP	22	10.99.101.0/24
TCP	3389	10.99.101.0/24

Platform Server Ingress

SID_App_Ingress

Protocol	Port Range	Source
TCP	80	10.241.25.0/24
TCP	443	10.241.25.0/24
TCP	8080	SID_VIEW_SG

Platform Server Egress

SID_App_Egress

Protocol	Port Range	Destination
TCP	443	10.87.14.29/32
TCP	3306	10.24.3.102/32

Database Access

SID_DB_Ingress

Protocol	Port Range	Source
TCP	3306	SID_APP_SG

View Firewall Policies

SID_View_Ingress
SID_Admin_Ingress

App Firewall Policies

SID_App_Ingress
SID_Admin_Ingress
SID_App_Egress

Database Firewall Policies

SID_DB_Ingress

Virtual Private Cloud – On-Premises Connection

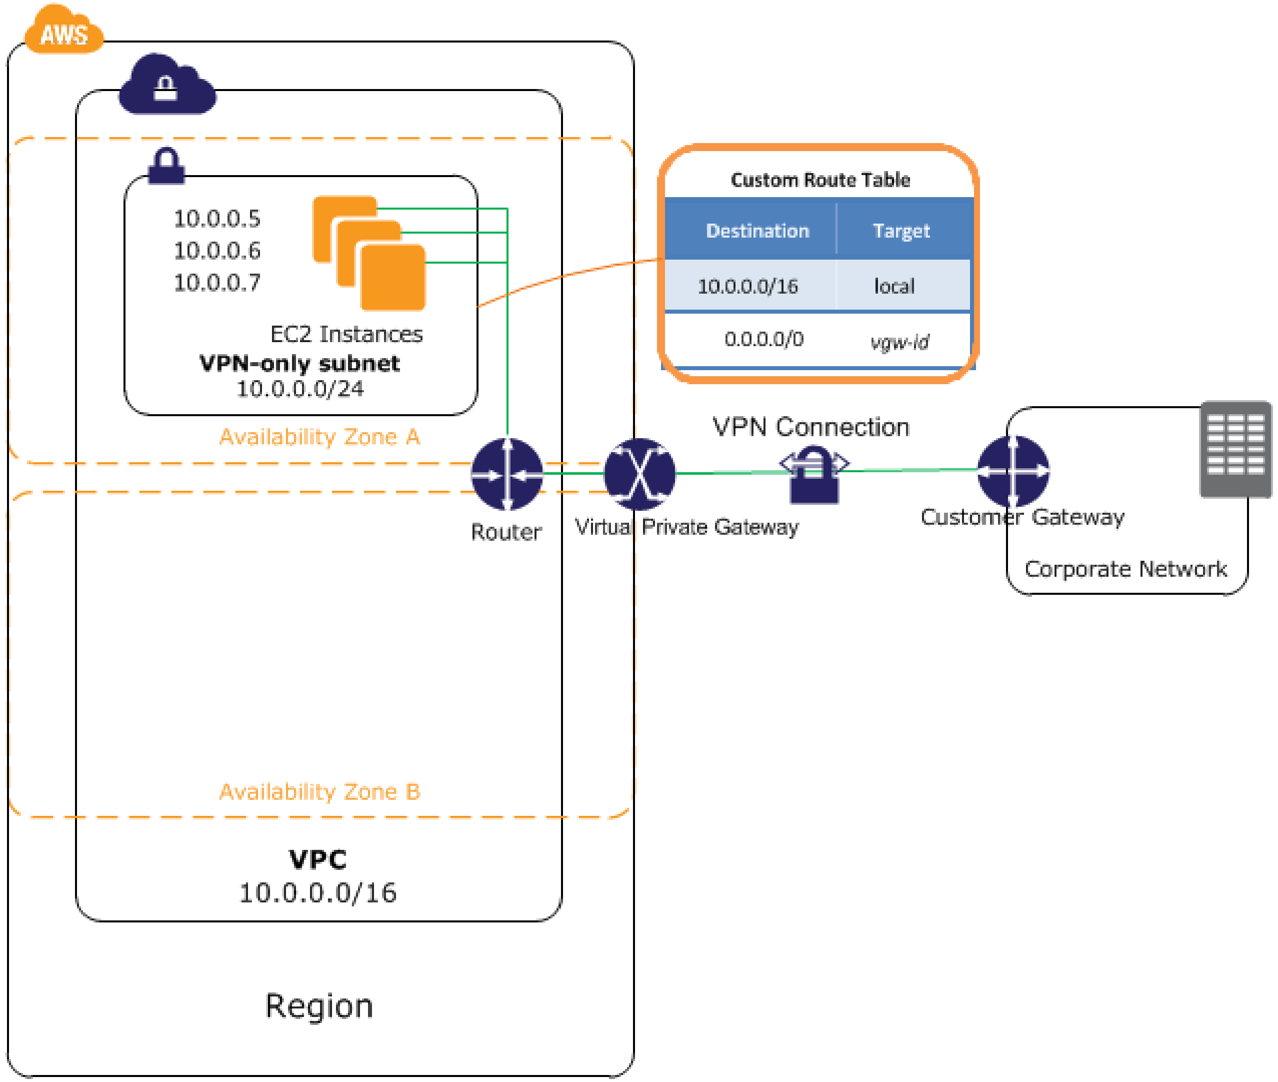

Custom Route Table

Destination	Target
10.0.0.0/16	local
0.0.0.0/0	*vgw-id*

Virtual Private Cloud – Overview

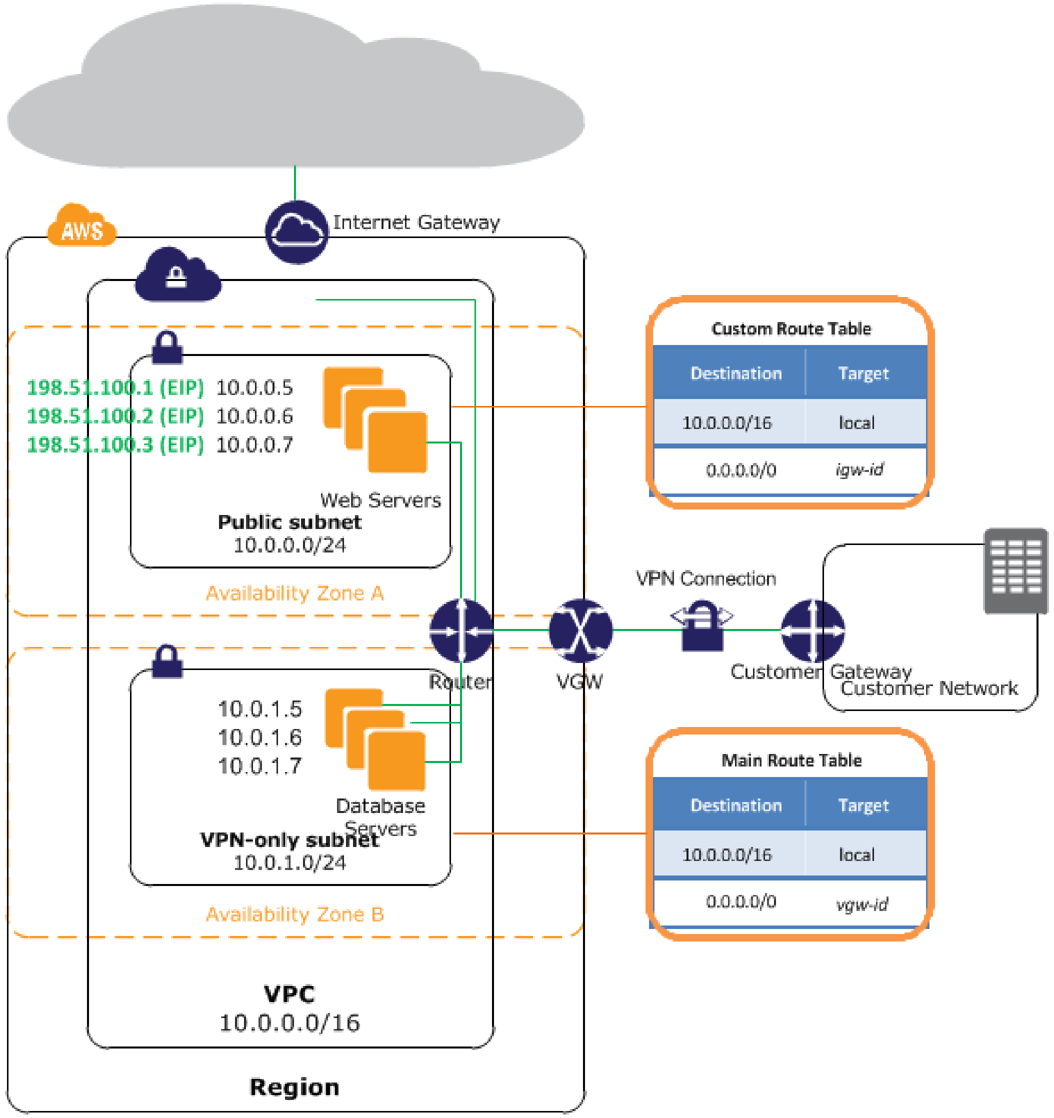

Custom Route Table

Destination	Target
10.0.0.0/16	local
0.0.0.0/0	*igw-id*

Main Route Table

Destination	Target
10.0.0.0/16	local
0.0.0.0/0	*vgw-id*

Virtual Private Cloud – Network and Subnets

- **Network Topology**
  - Private address space
    - Any range is valid, but we suggest a non-routable CIDR
    - Public CIDR ranges are only reachable via a Virtual Private Gateway
    - CIDR ranges can be as large as a /16 to as small as a /28
- **Subnets**
  - Public subnets have a 0.0.0.0/0 route to the Internet Gateway (IGW)
    - Instances that require a public IP need to reside in a public subnet
  - Private subnets do not have an outbound route through the IGW
    - NAT instances are commonly used as an outbound gateway for private instances
  - Subnets cannot span AZ's, but subnets can share routing tables, which provides similar functionality.

Virtual Private Cloud – Route Tables

- **Route Tables**
  - Can be applied to multiple subnets
  - Typical routing entries
    - 10.0.0.0/16 = Local
    - 0.0.0.0/16 = Internet Gateway (Public Subnet)
    - -or-
    - 0.0.0.0/16 = eni-12345678 (Private Subnet)

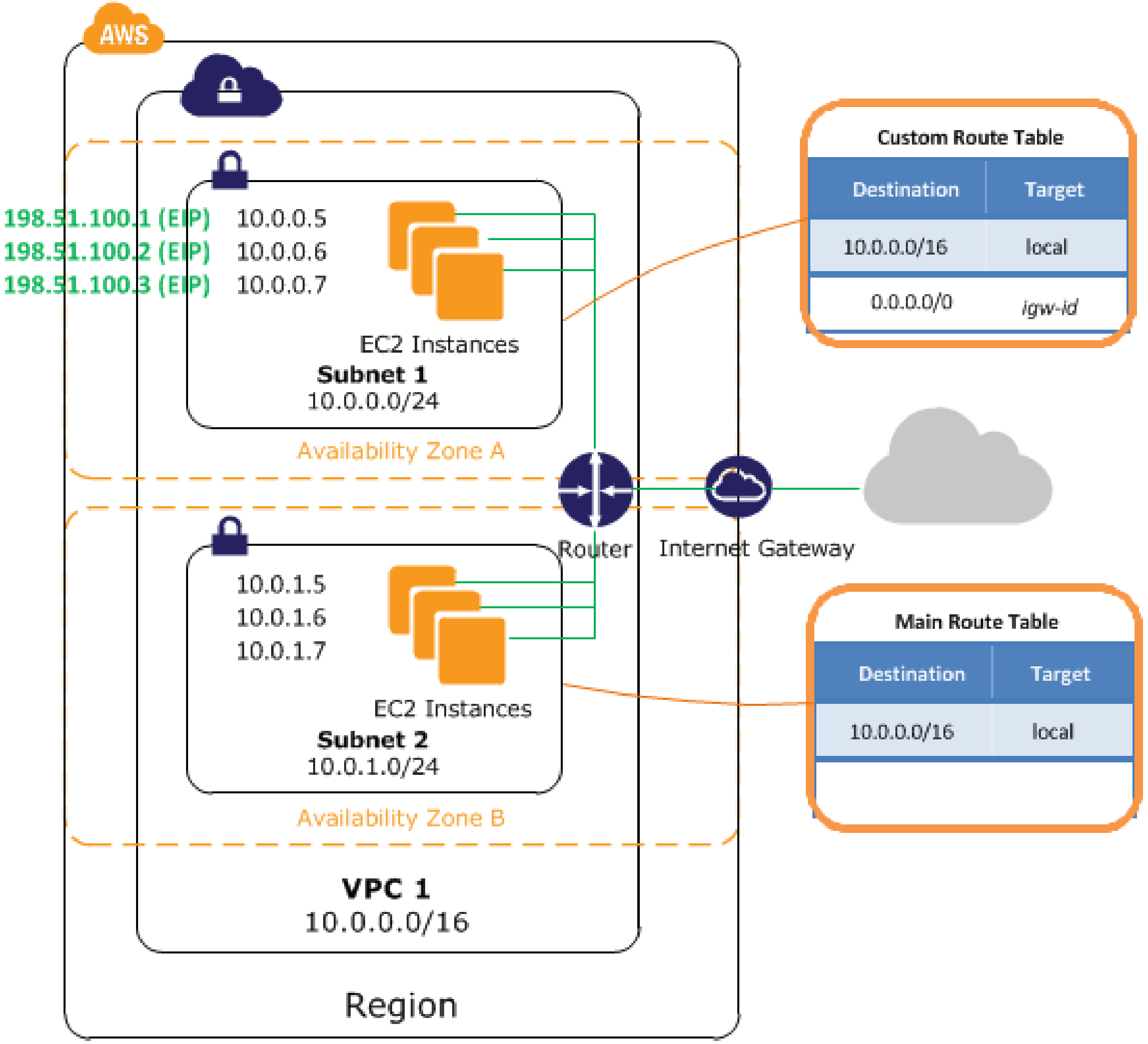

Custom Route Table

Destination	Target
10.0.0.0/16	local
0.0.0.0/0	*igw-id*

Main Route Table

Destination	Target
10.0.0.0/16	local

Virtual Private Cloud – Peering

- **Peering**
  - VPC -> VPC peering
  - Unique CIDR
  - VPN solutions
    - OpenVPN
    - OpenSwan

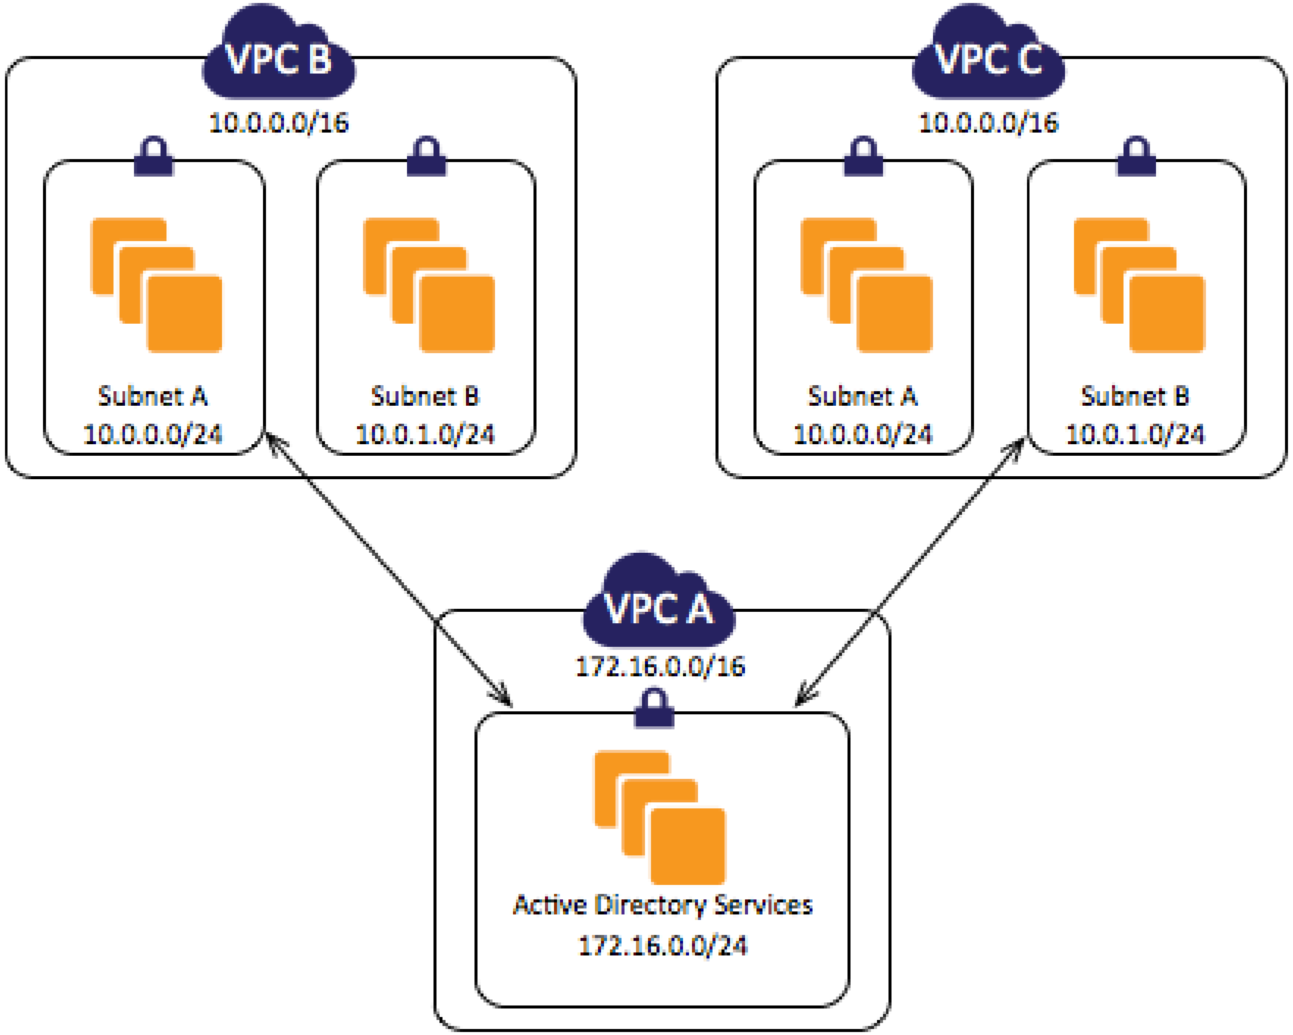

Route53 - Basic Feature Set

- Zone Creation
- Zone Import
  - Import your zone file from a previous provider
  - Delegate this zone to the AWS name servers
- Record Types
  - A
  - CNAME
  - TXT,MX,DKIM
  - Alias
  - S3 buckets and ELBs can be an alias target, allows zone apex magic

Route53 - Advanced Feature Set

- Weighted Resource Record Sets
- Health Checks
- Global Load Balancer
  - Using weighted record sets, you can create a pool of endpoints from which to balance traffic
  - Enabling a health-check on this pool allows for a DNS based load balancer which can be applied to any resource (AWS or non-AWS)
- Latency Resource Record Sets
- Geolocation Resource Record Sets

Route53 – Global Failover

- Global Failover Pattern
- Uses R53 Health Checks

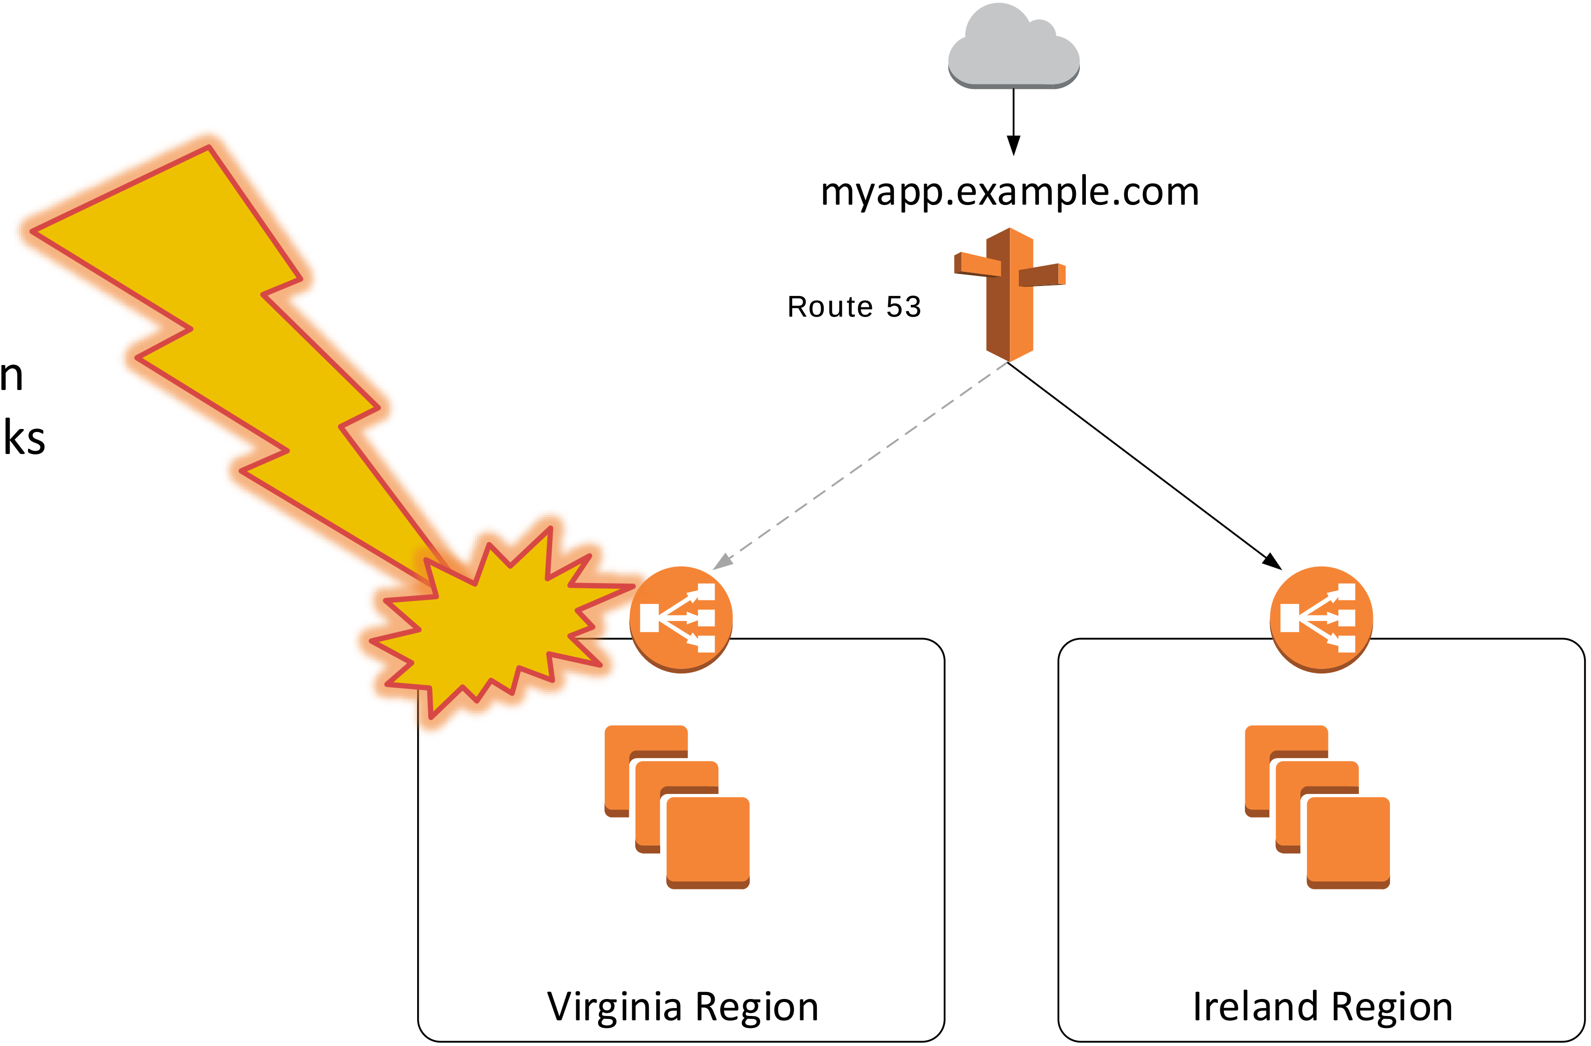

AWS Compute Overview

EC2 – Elastic Cloud Compute

- AMI
  - Instances are based on an Amazon Machine Image
  - You can create new AMIs from a running instance
  - AMIs are stored in S3 for 11 9's of durability
  - AMIs are unique to each region

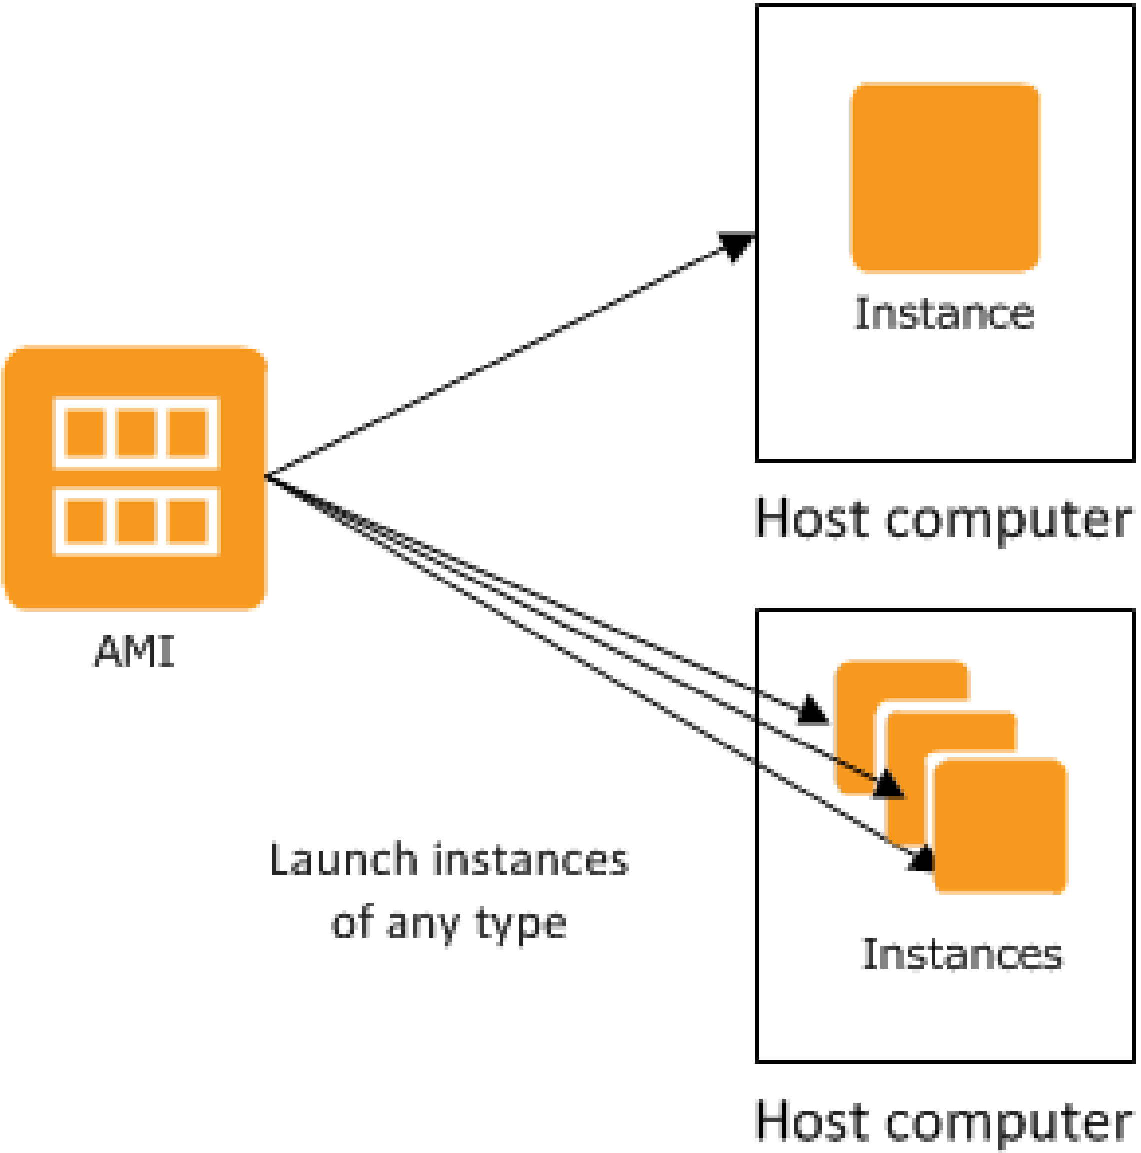

EC2 - Instance Types

Choosing the correct instance type for the required workload

- T2 for light weight general purpose – but with burstable performance
- M4 for general purpose
- R3, X1 for memory and database heavy application
- C3, C4 for compute heavy applications
- G2, P2 for GPU intensive applications
- I2, D2 for storage heavy applications (random)
- HS1 for storage heavy applications (sequential)

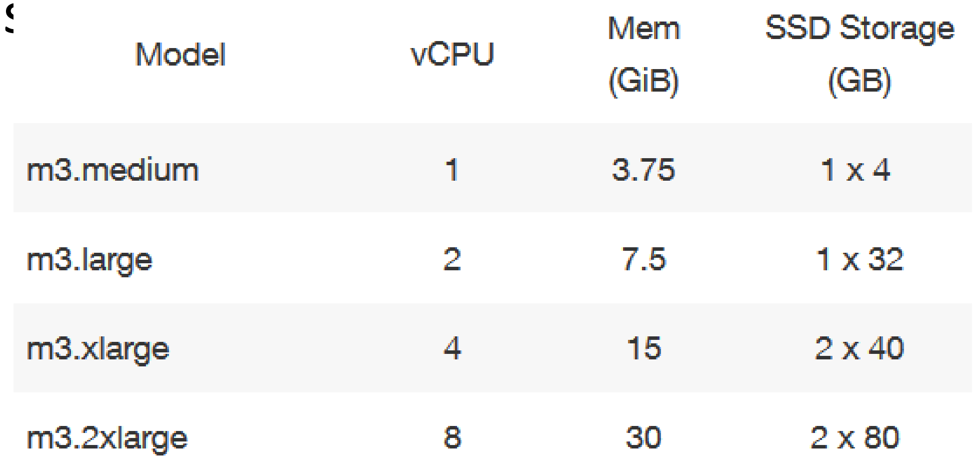

Model	vCPU	Mem (GiB)	SSD Storage (GB)
m3.medium	1	3.75	1 x 4
m3.large	2	7.5	1 x 32
m3.xlarge	4	15	2 x 40
m3.2xlarge	8	30	2 x 80

Example of M3 family of instance type

EC2 - Running Instances

Running instances

- Instances are launched into an existing VPC subnet
- CloudWatch monitoring is enabled by default
  - CPU Utilization, Network I/O are the primary data points of interest
  - Memory and Disk require an additional script that will post a to a custom CloudWatch metric
- Status checks
  - OS check
  - Network reachability check

EC2 - Monitoring

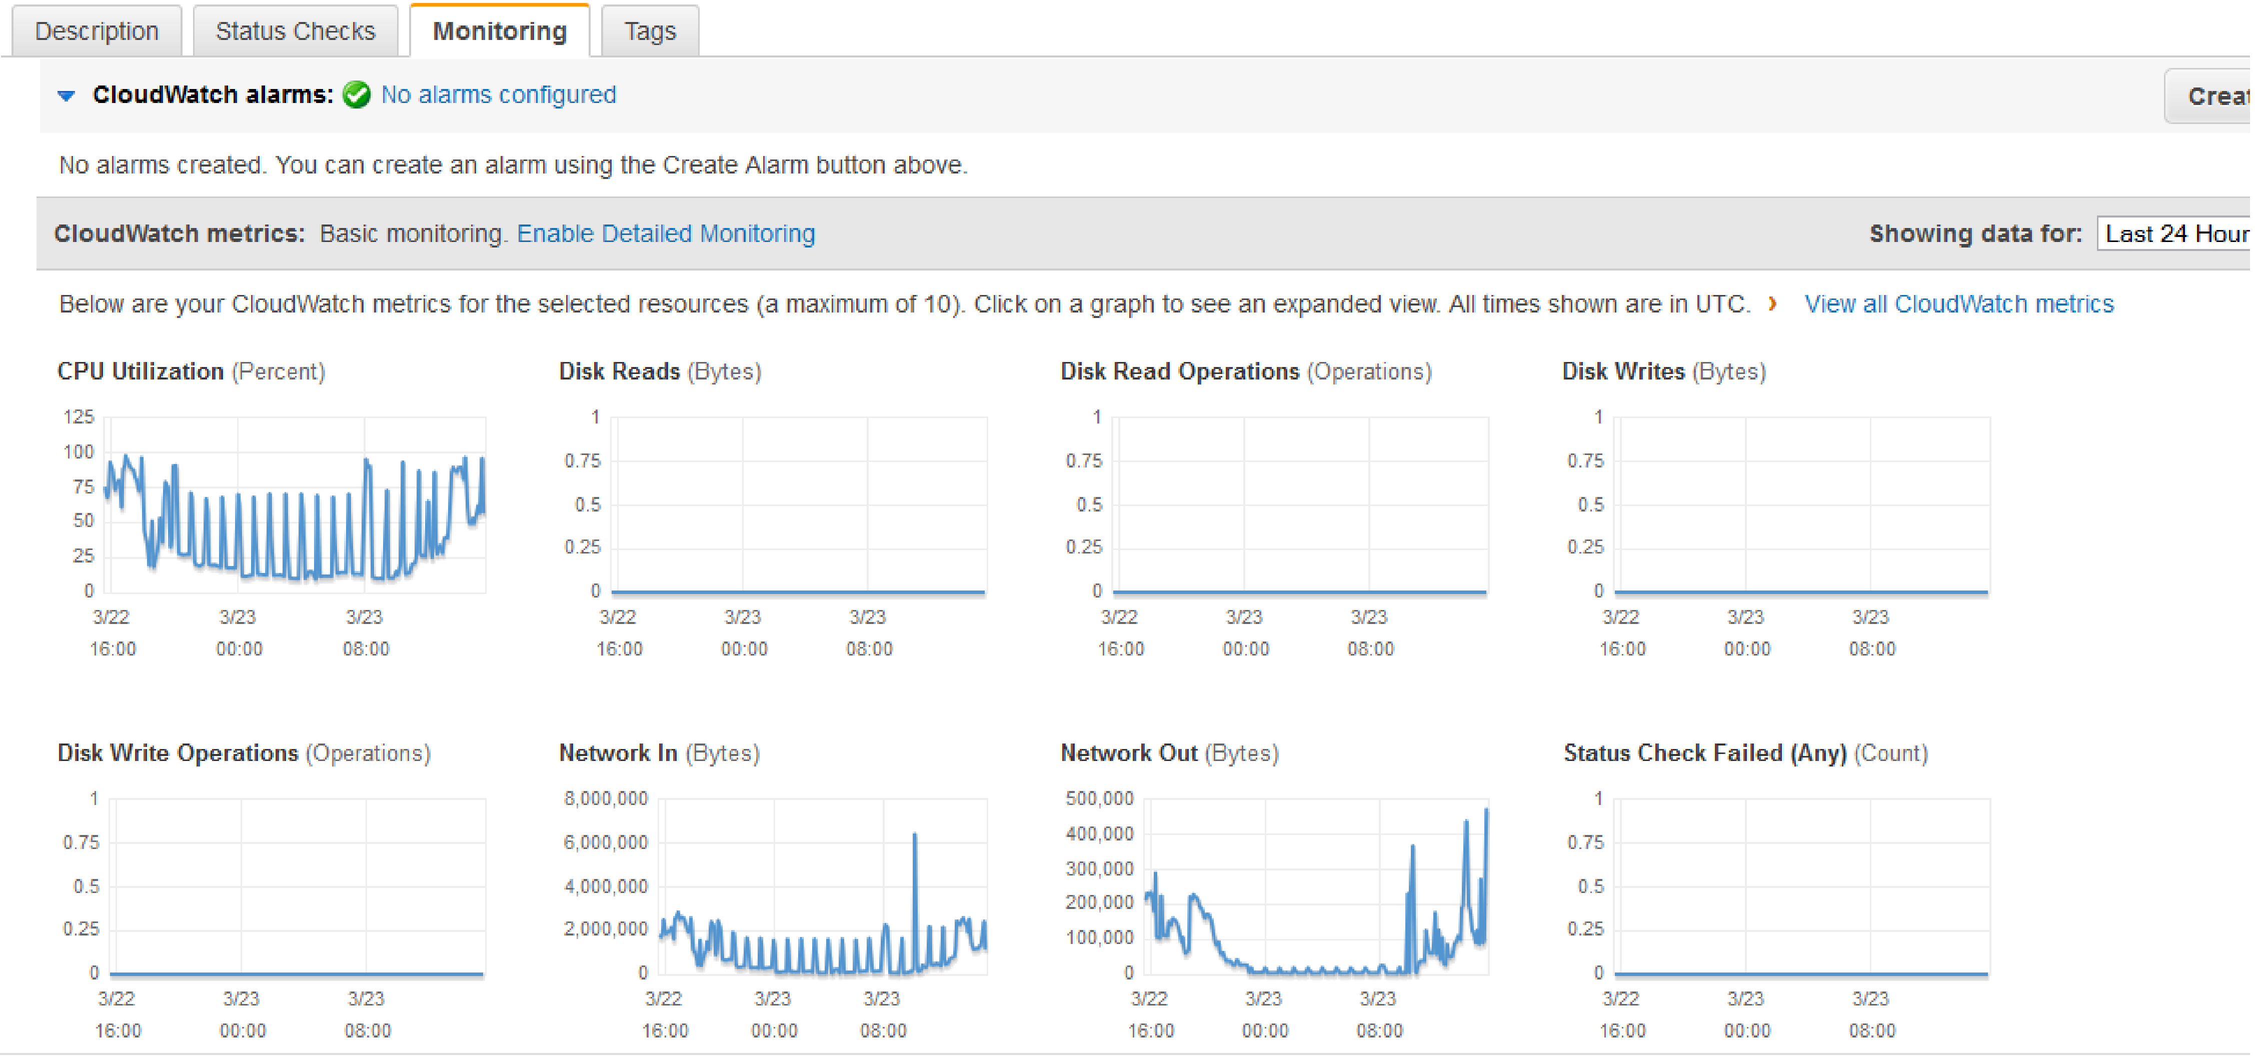

CloudWatch alarms: ✓ No alarms configured

No alarms created. You can create an alarm using the Create Alarm button above.

CloudWatch metrics: Basic monitoring. Enable Detailed Monitoring

Showing data for: Last 24 Hour

Below are your CloudWatch metrics for the selected resources (a maximum of 10). Click on a graph to see an expanded view. All times shown are in UTC. › View all CloudWatch metrics

EC2 - Bootstrapping

- User Data
  - Provides a hook to inject scripting into any standard instance you decide to launch
    - These include the Amazon Linux, Windows and Ubuntu AMIs
    - User Data can only be modified while the instance is stopped
  - Suggested patterns
    - Install security updates
      - yum update -y
    - Install middleware
      - yum install -y httpd
      - chkconfig httpd on
    - Download and execute a remote script
      - Assign an IAM Profile to the EC2 instance
      - Aws s3 cp s3://mybucket/myscript.sh /tmp/myscript.sh
      - ./tmp/myscript.sh

EC2 - Pricing

1 > 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

2 > Reserved Instance (RI)

- 1 or 3 year commitment
  - Pay for EC2 hourly at reduced rates (from On Demand rates)
- Payment Options
  - No Upfront payment: no CapEx, lower hourly rate than On Demand
  - Partial Upfront payment: some CapEx, lower hourly rate than No Upfront
  - All Upfront payment: larger Capex, lowest hourly rate possible

EC2 - Pricing

3 > Spot

- Useful for “worker pool” scenarios
  - o Transcode, map reduce task nodes
- Can be lost as soon as someone is willing to pay more for that instance

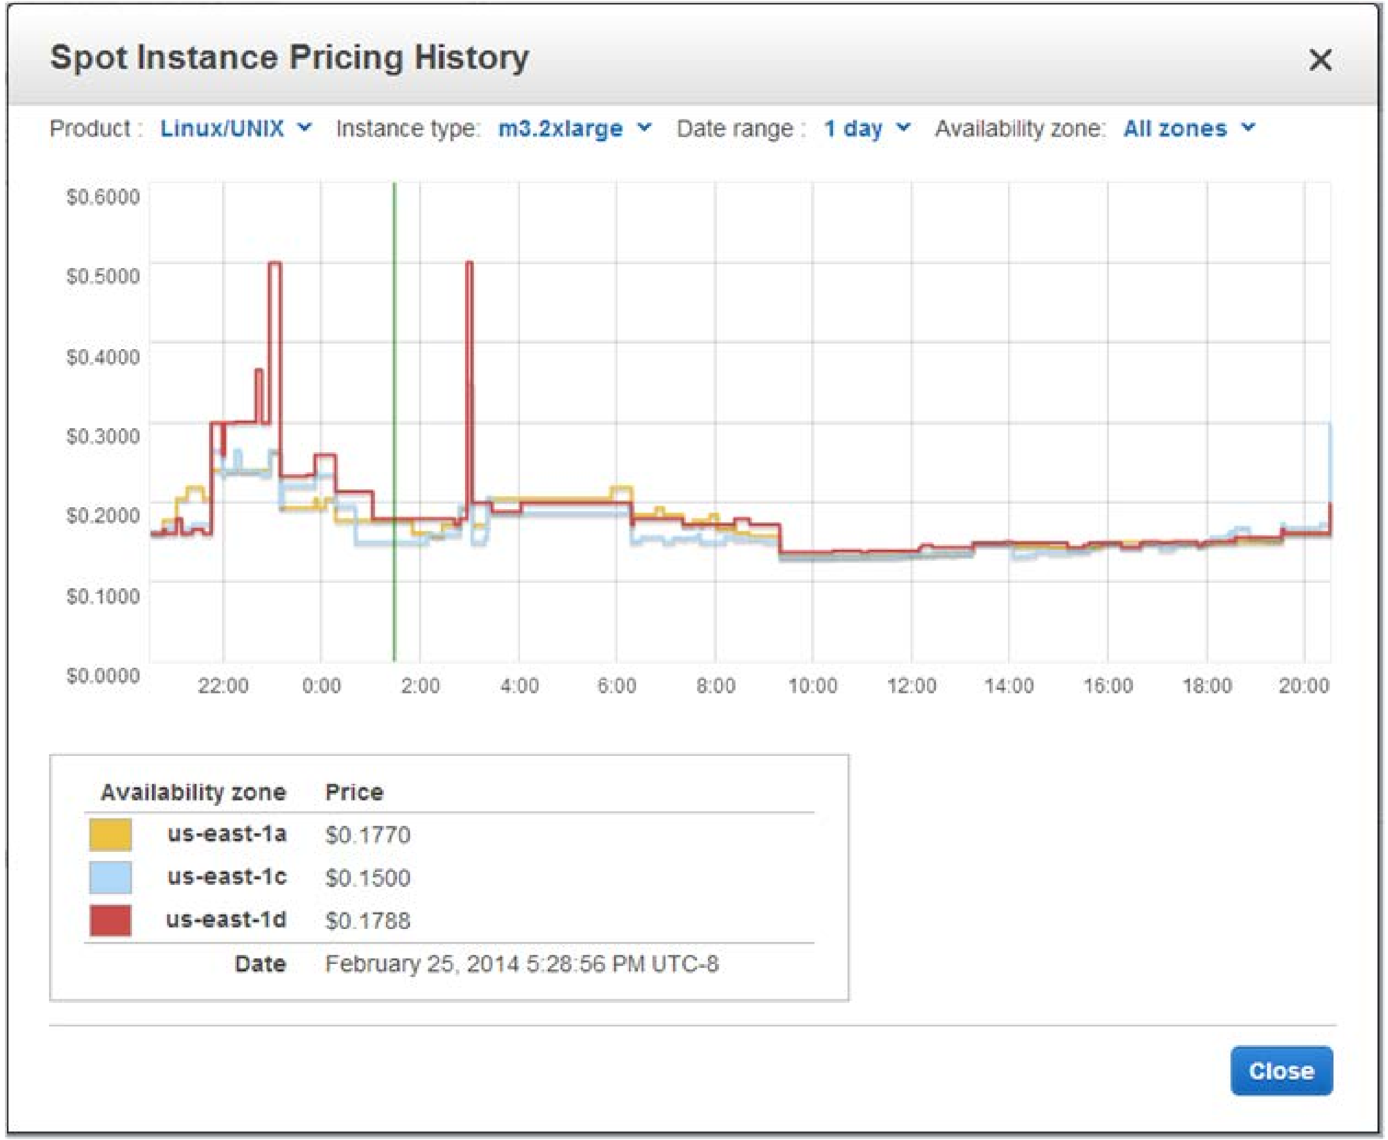

Spot Instance Pricing History

Product : Linux/UNIX
Instance type: m3.2xlarge
Date range : 1 day
Availability zone: All zones

	Availability zone	Price
🟨	us-east-1a	$0.1770
🟦	us-east-1c	$0.1500
🟥	us-east-1d	$0.1788
	Date	February 25, 2014 5:28:56 PM UTC-8

AWS Elastic Load Balancing

ELB - Elastic Load Balancer

Elastic Pool of Virtual Load Balancers

- **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

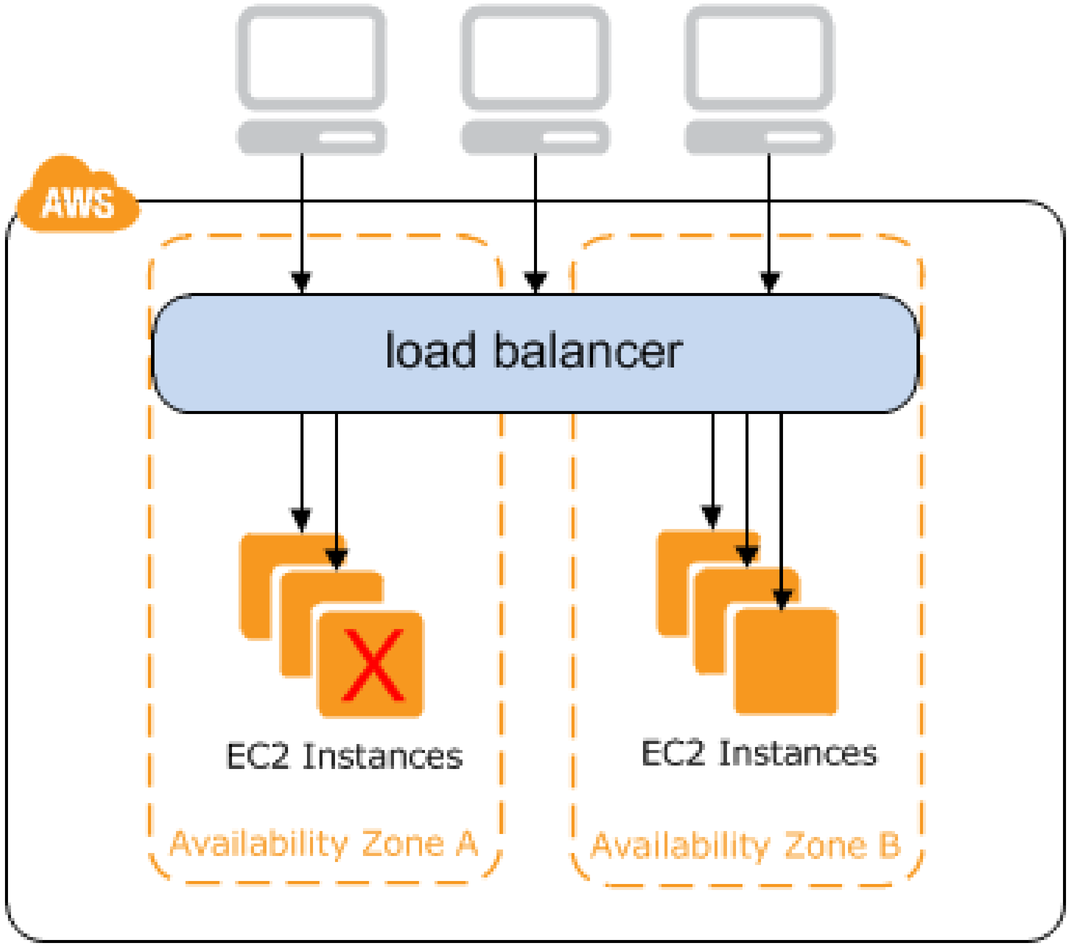

ELB – Spans Multiple Availability Zones

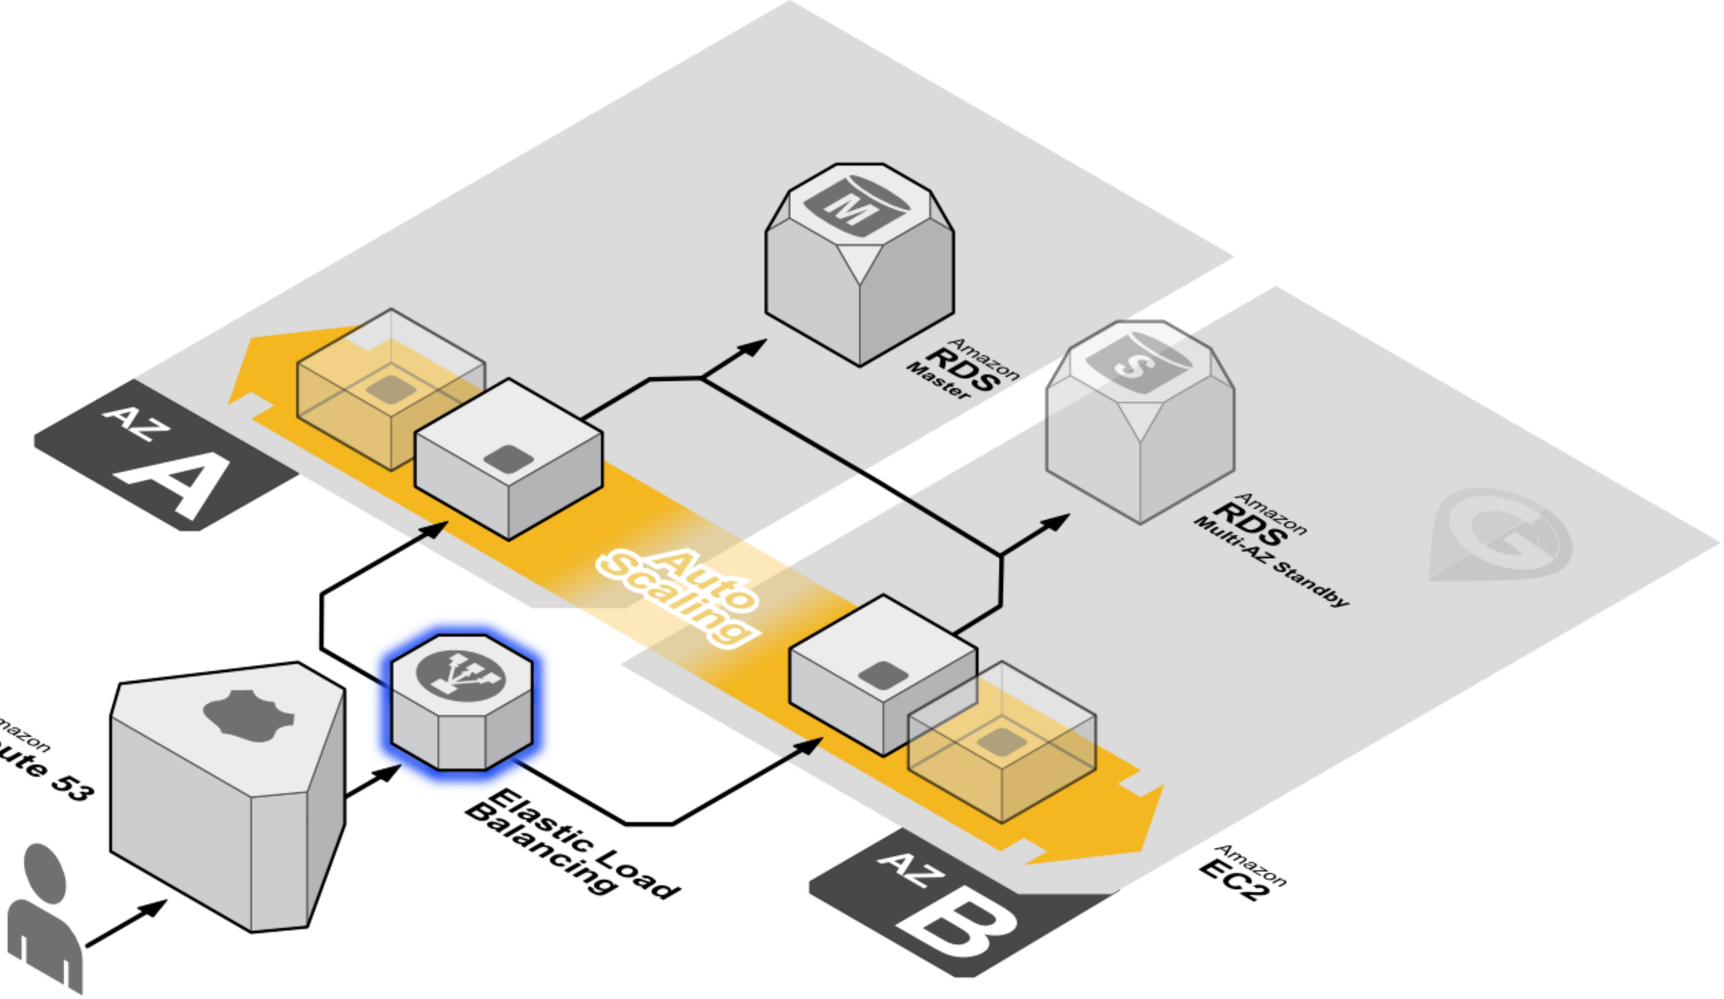

Auto Scaling - Overview

- **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

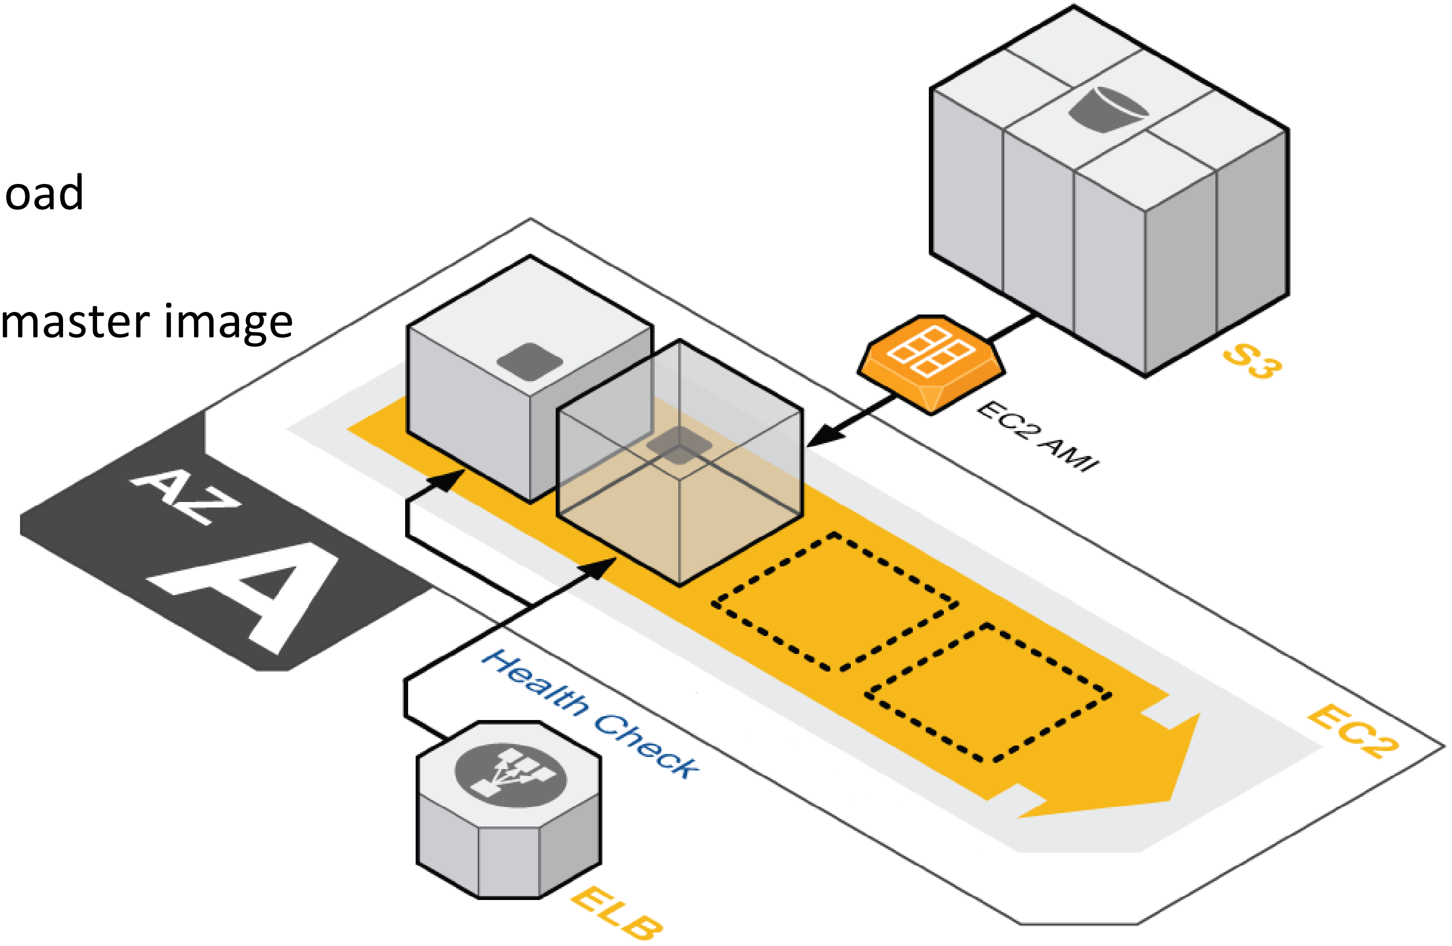

Auto Scaling - Components

- **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

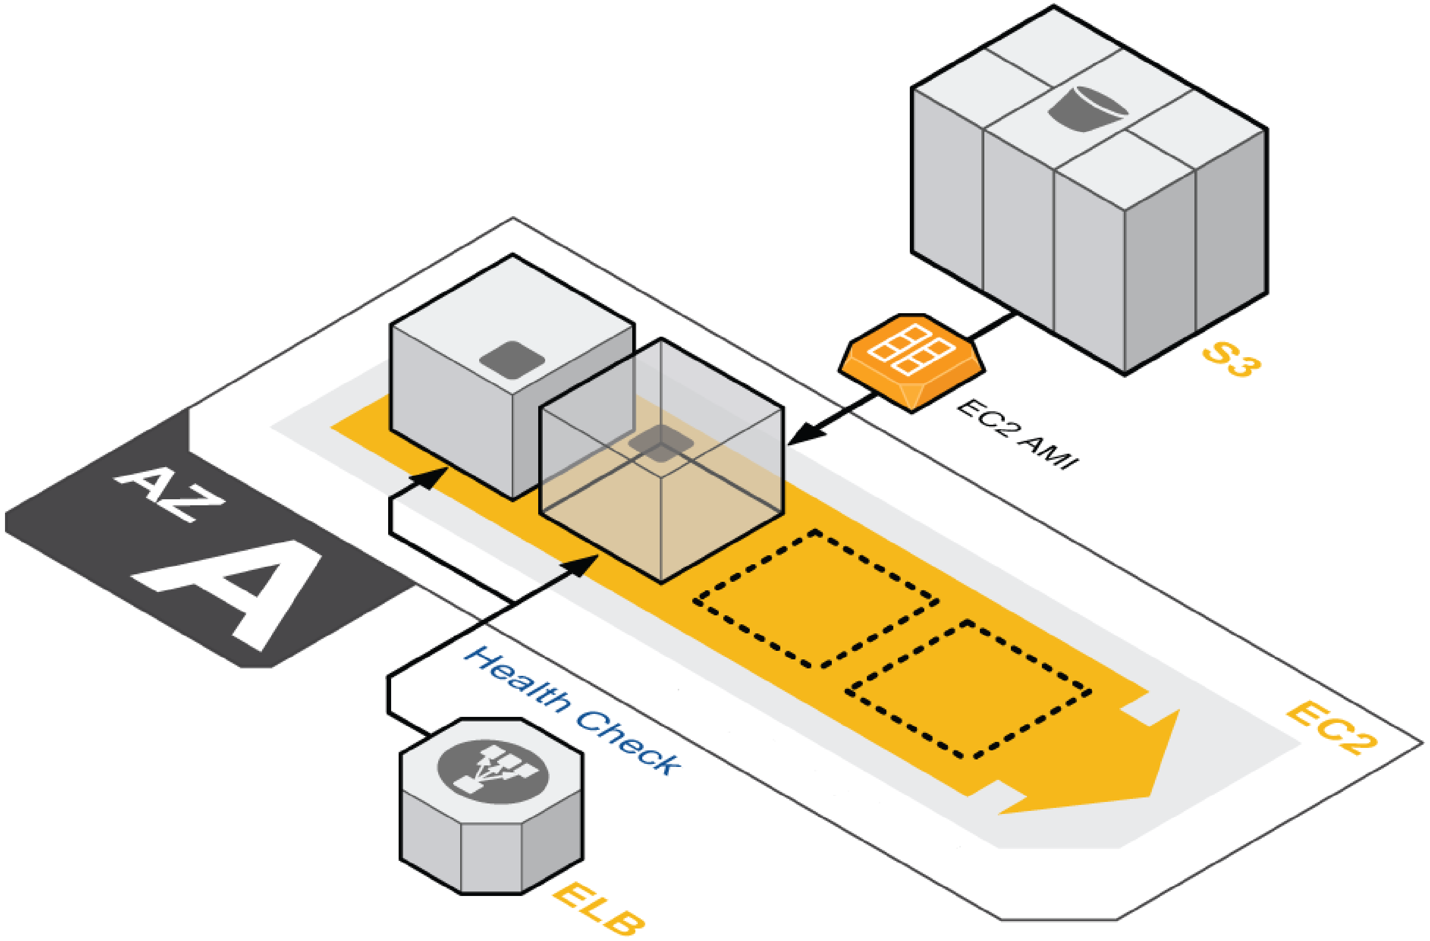

Auto Scaling - Multi-AZ

- **Multi-AZ Auto Scaling**
    - Highly Available
    - Production Standard
    - Spans Datacenters

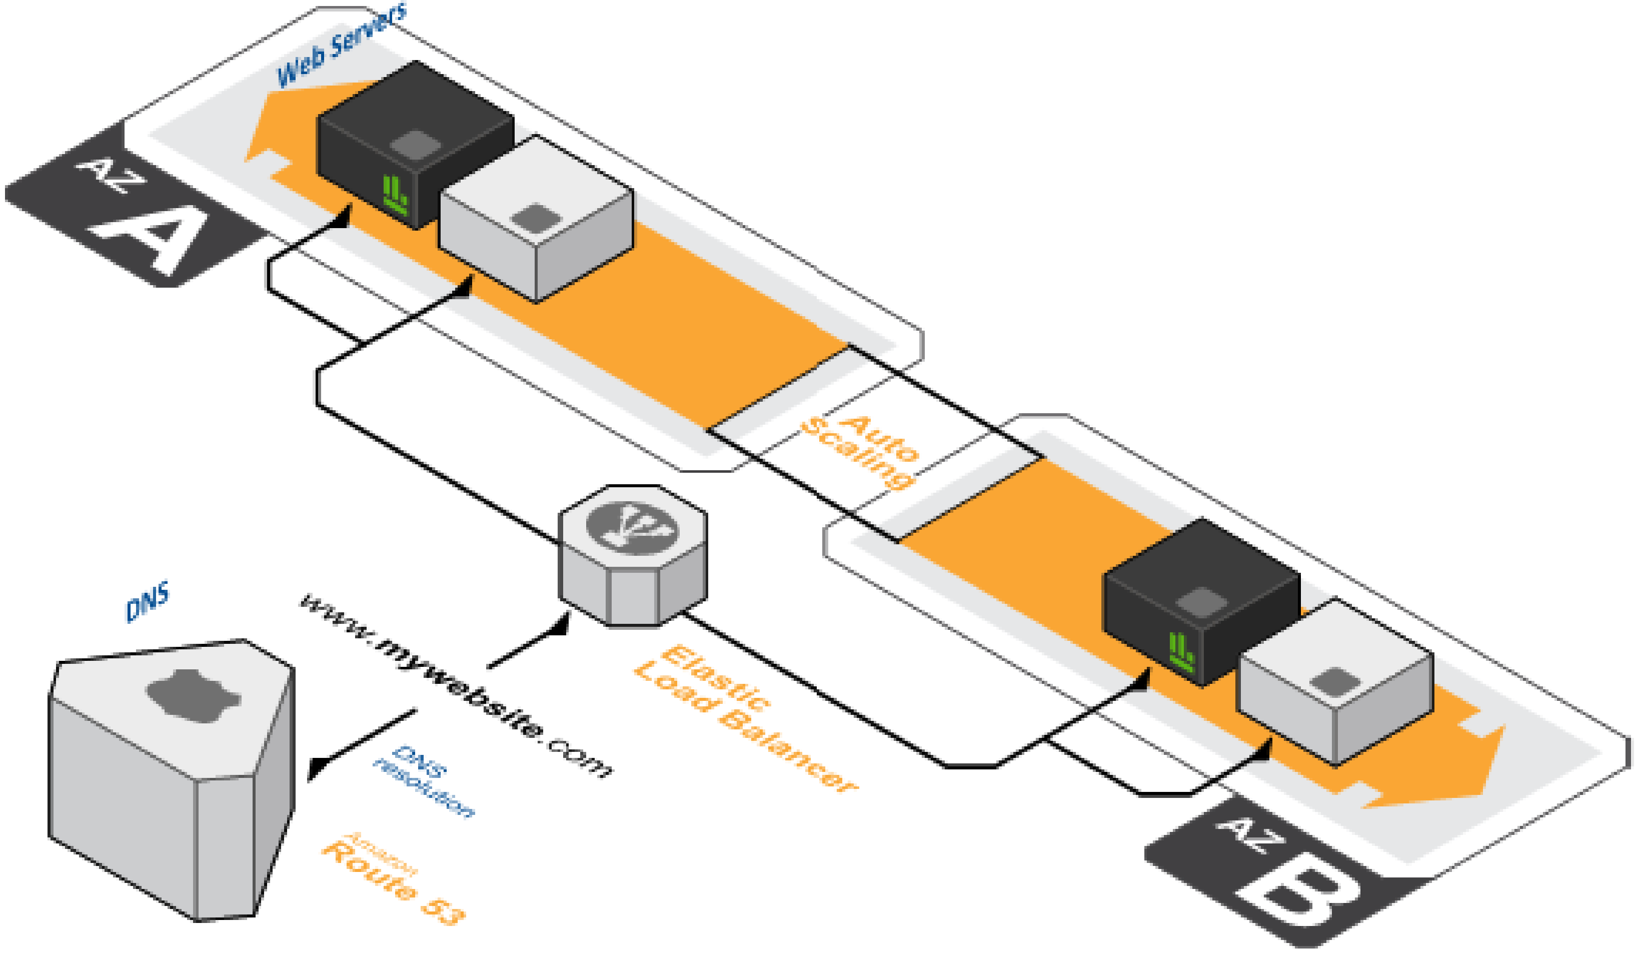

Auto Scaling - CloudWatch

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

AWS Directory Service

Directory Service - Overview

Three types of directory services:

- Microsoft AD
  - A managed Microsoft Active Directory service running on Windows Server 2012 R12
  - Highly availability (multi-AZ), patched, and monitored
  - Can support up to 50,000 users
  - Fully functional MS AD

- Simple AD
  - Powered by Samba 4 Active Directory
  - Users and Groups can be created directly in the AWS console
  - Windows servers can auto-join this domain as they would in an AD environment
  - Can support 500 users

- AD Connector
  - Connect your on-prem AD to your AWS account
  - Associate AD users/groups with IAM users/groups
  - Windows servers can auto-join this domain as they would in an AD environment
  - Manage the AWS console using your AD credentials

Directory Service – AD Connector

- Active Directory Connector instances are launched into your VPC
- AD Connectors communicate with on-prem AD servers
- AD credentials are no longer necessary when joining instances to a domain (Auto-Join)

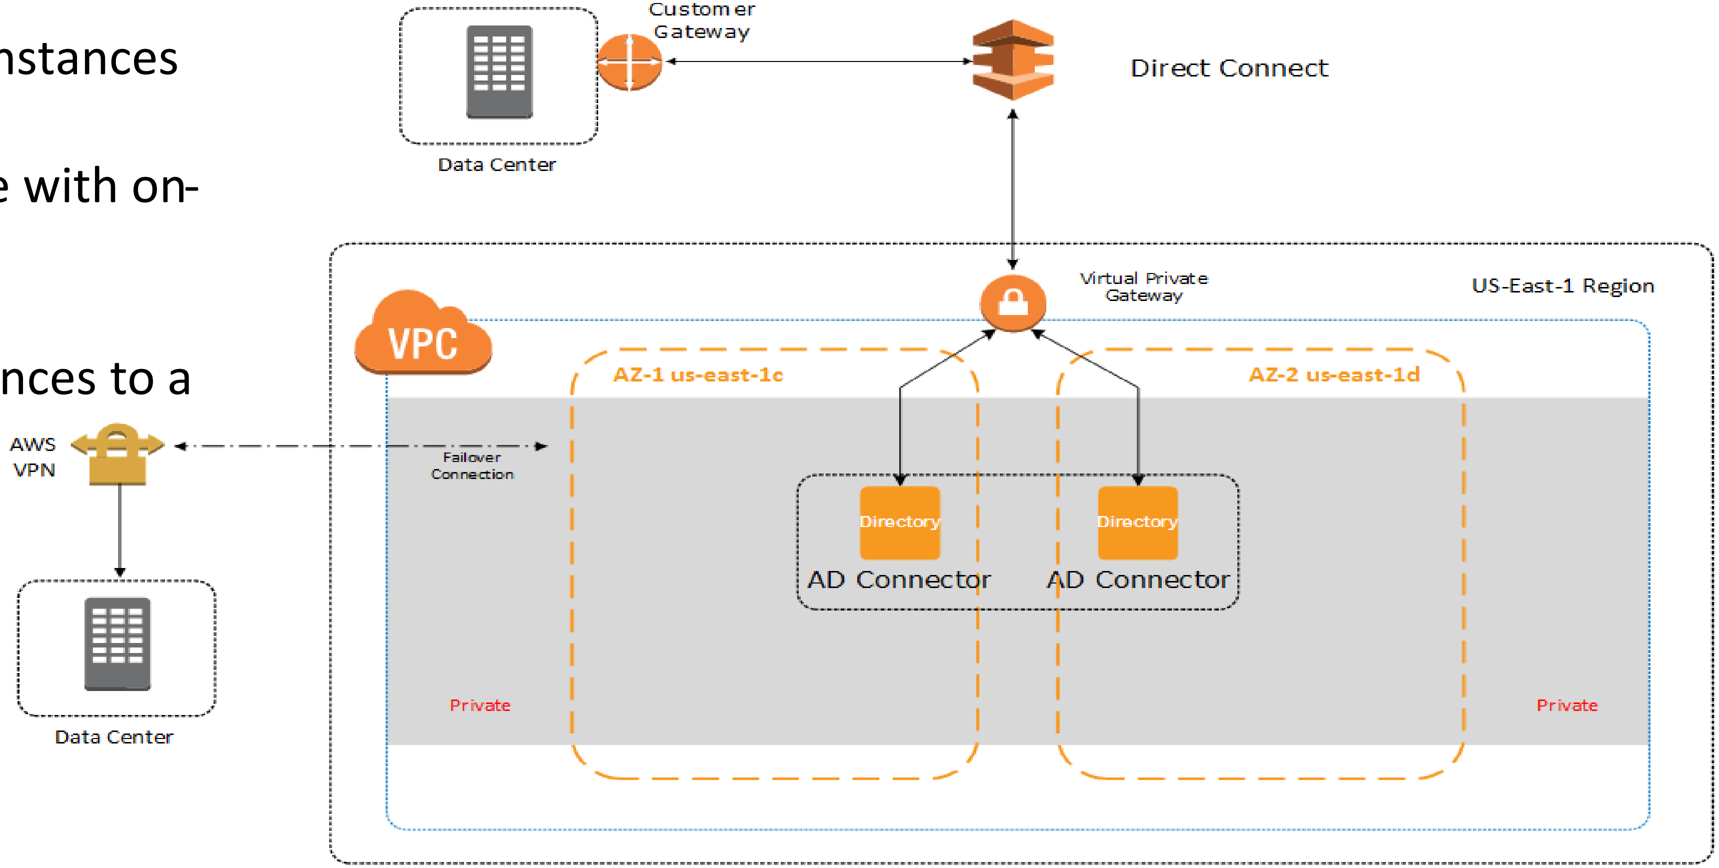

AD Connector - Single Sign On Flow

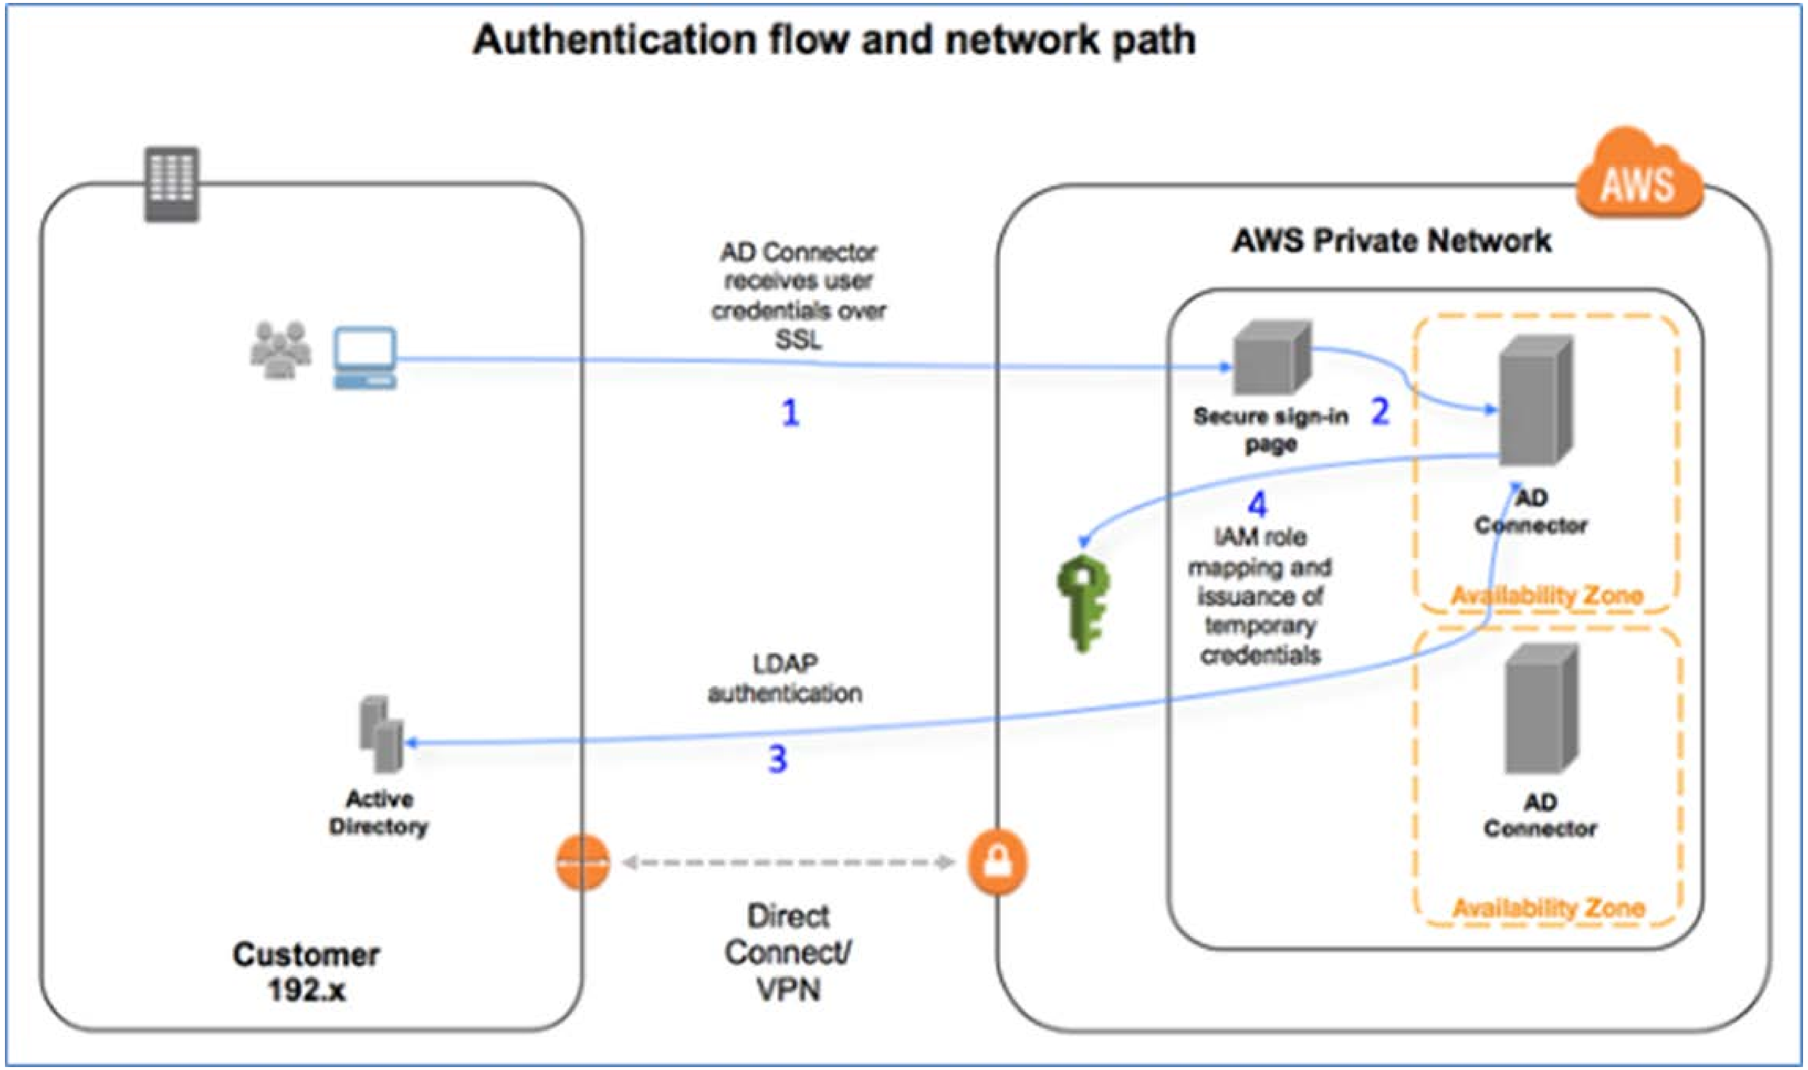

Identity & Access Management

IAM Users

- Identity and Access Management
- Create Users and Groups
- Establish Trust Relationships
- Govern Access via Policy Documents

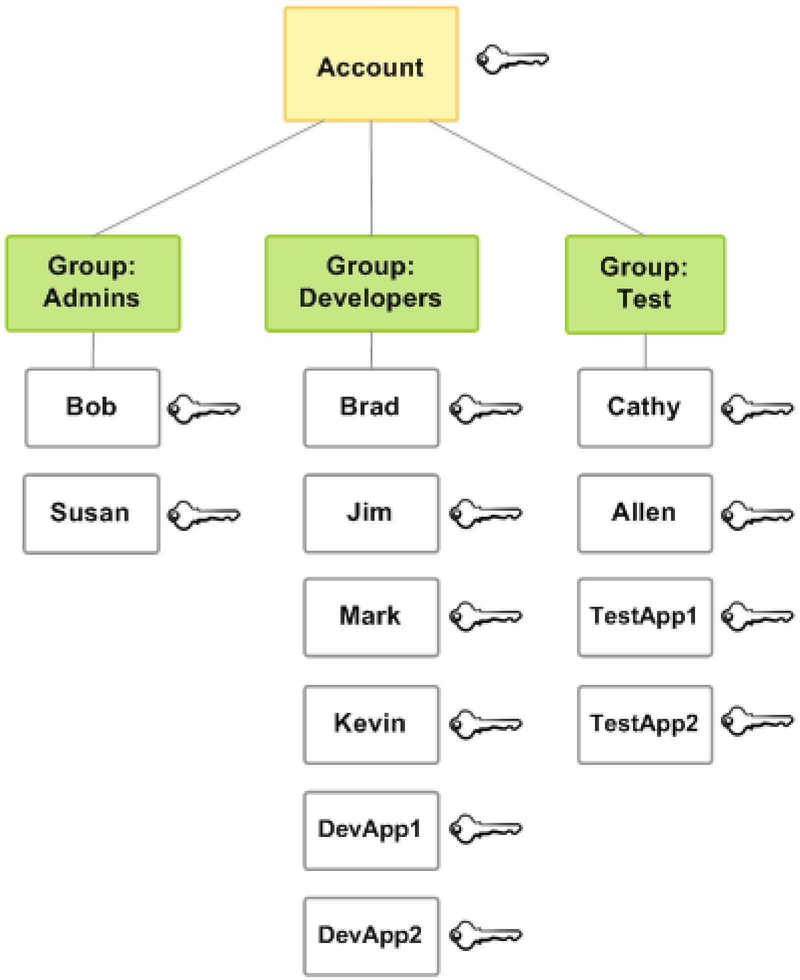

IAM - Groups, Roles & Instance Profiles

- Deny by default
  - Explicit allow required to grant access
  - Explicit deny always trumps an explicit allow
- Users/Groups
  - Policies can be applied at the group or user level
- Roles
  - Policies can be applied to roles
- Instance Profile
  - Assumes role
  - Credentials are stored in instance metadata
  - Only Access Key ID and Temporary Token

IAM - Instance Profiles

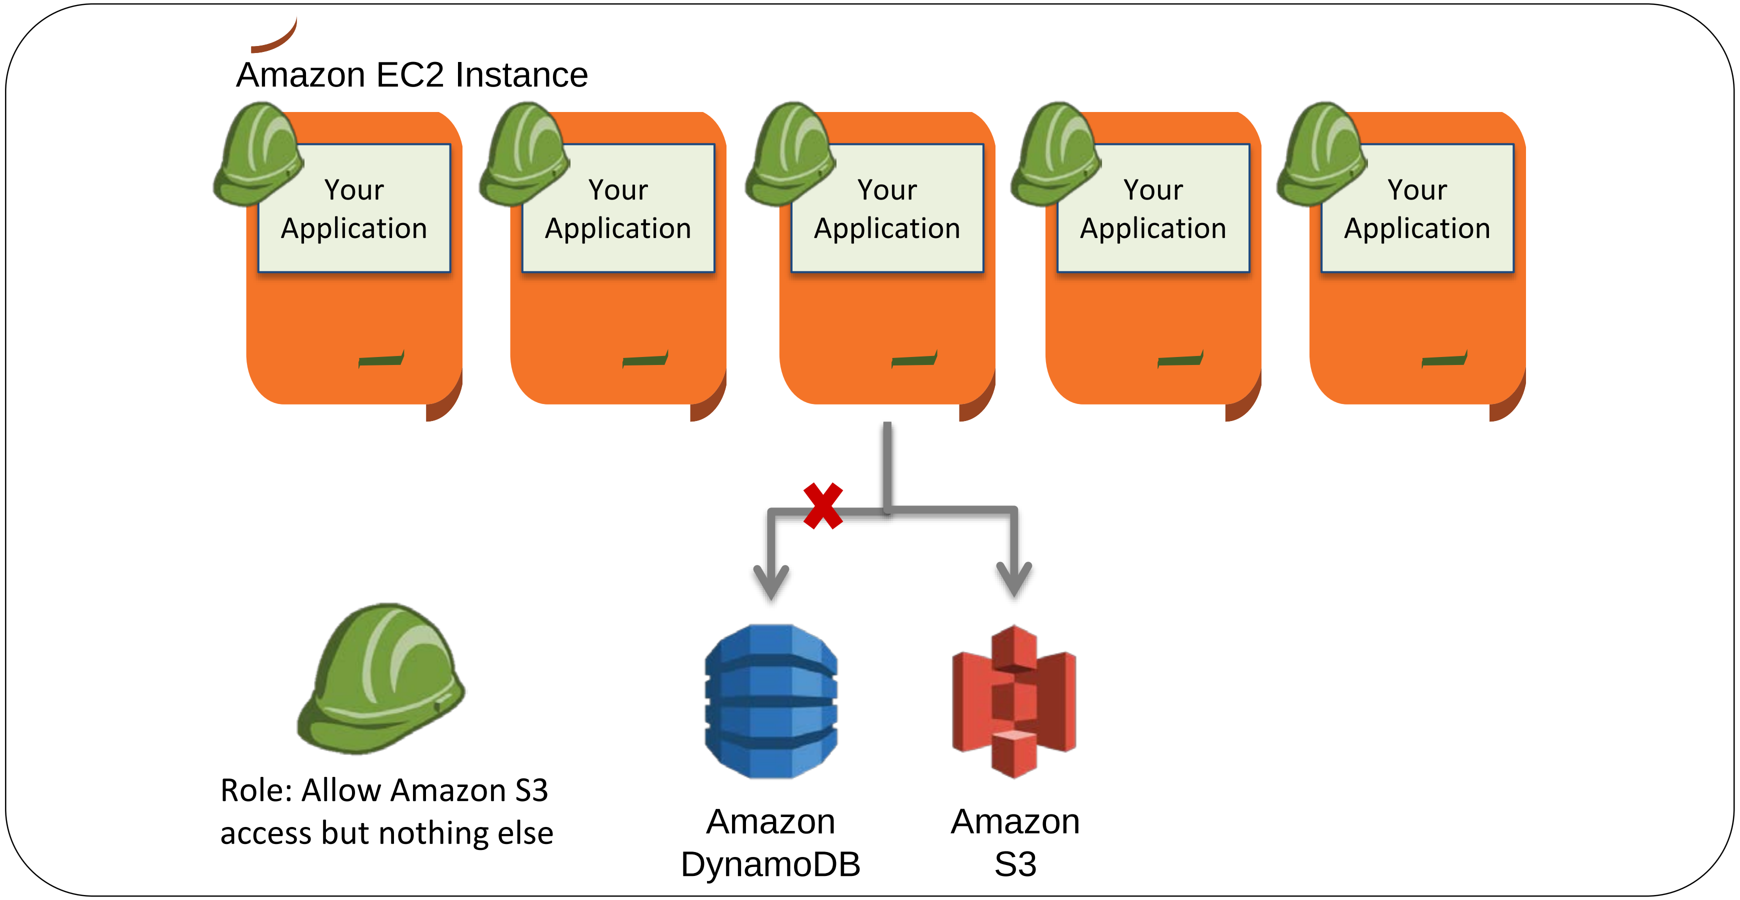

IAM – AWS Master Account

AWS Account

- Master/Root Account Permissions
  - Always treat the master account credentials as if they could launch an ICBM
- Allow by default
- MFA

AWS Storage Overview

Storage Services

- **EBS** – Elastic Block Store - (not actually a ”service”)
- **S3** – Simple Storage Service - (object storage)
  - **Standard**
  - **Standard I/A** – Infrequent Access
  - **Glacier** – Archival/Long-term
- **AWS Storage Gateway**
  - **Gateway-cached volumes** – store primary data in AWS and cache most recently used data locally
  - **Gateway-stored volumes** – store entire dataset onsite and asynchronously replicate data back to S3
  - **Gateway-virtual tape library** – store your virtual tapes in either S3 or Glacier
- **EFS** – Elastic File System

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

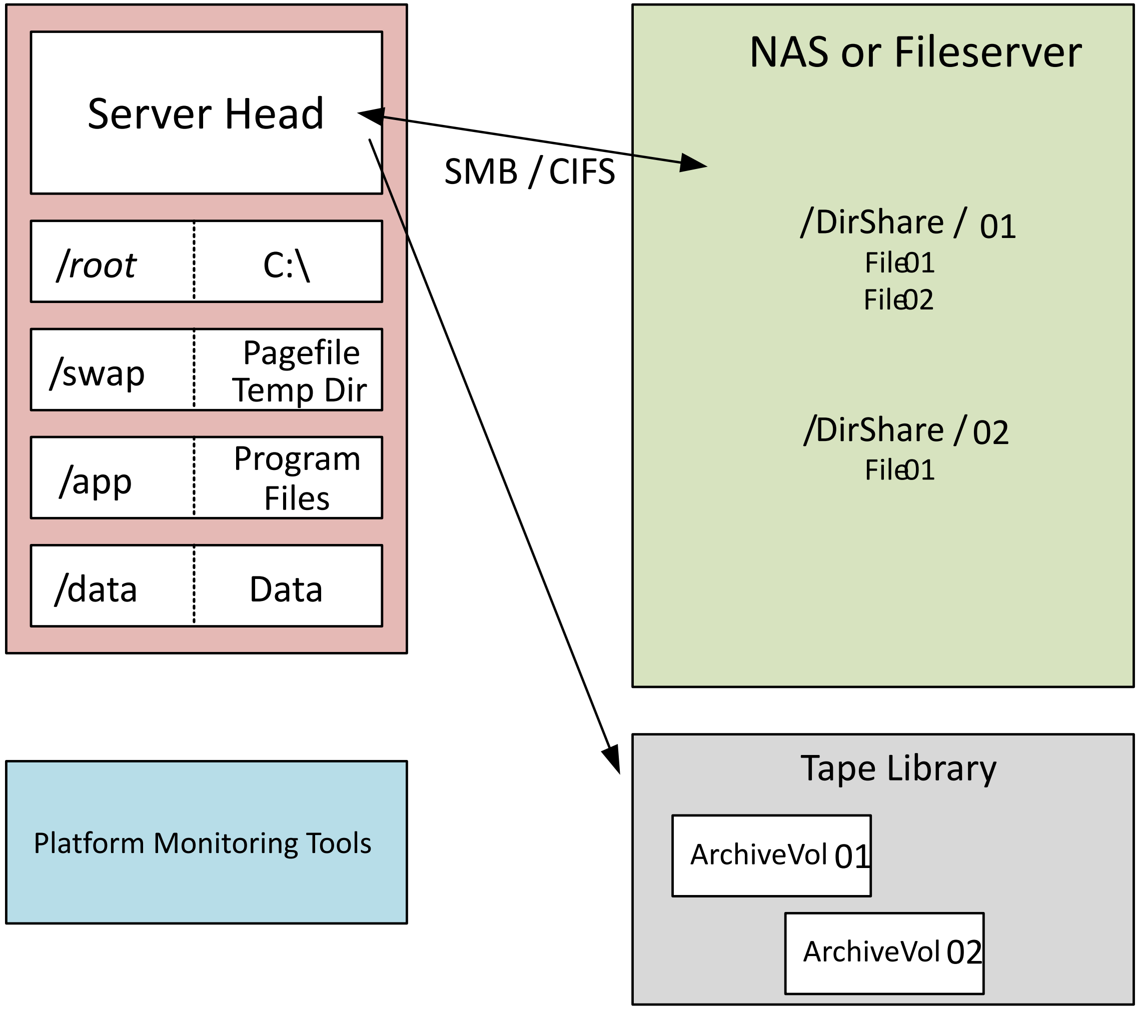

AWS Instance Volumes and Data Storage

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

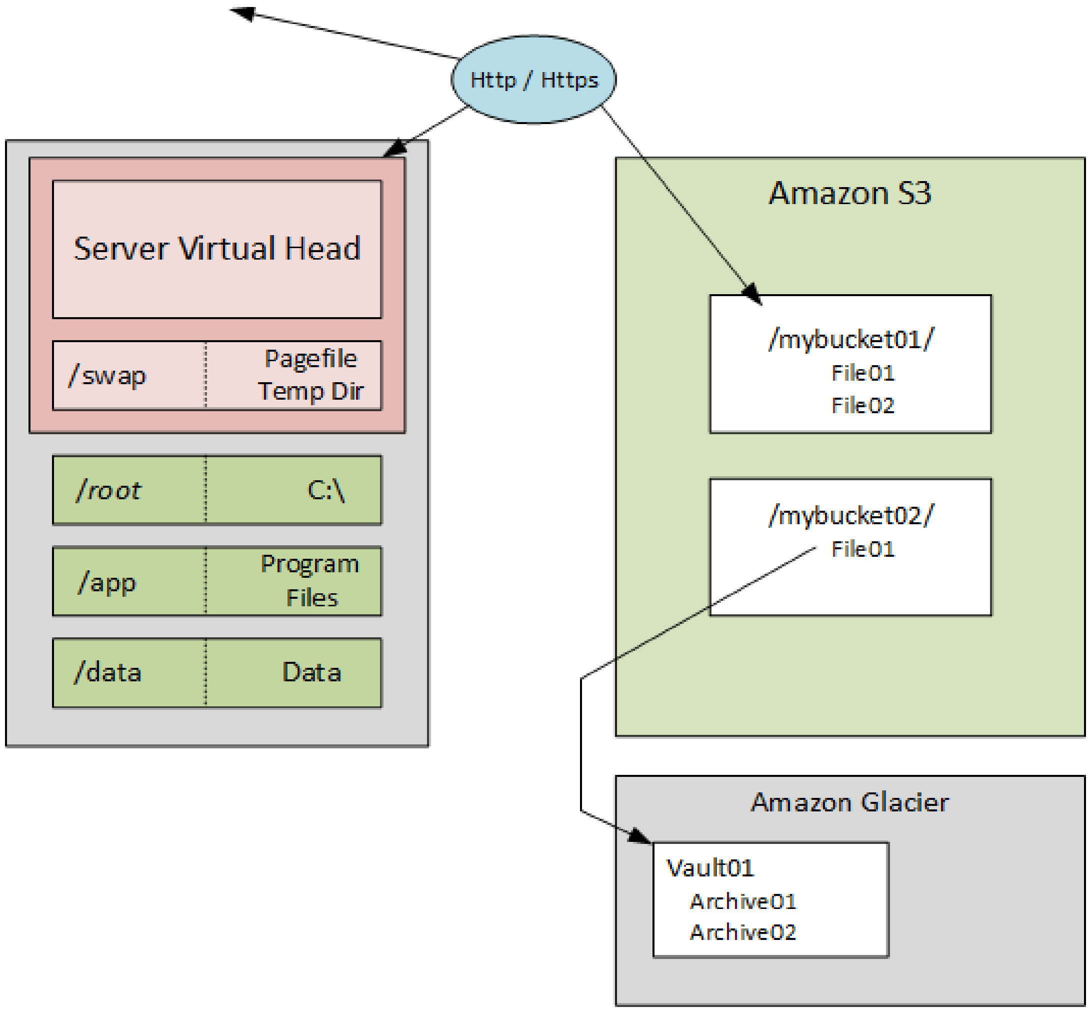

EBS - Elastic Block Store

Block storage ideal for creating versatile OS volumes

- Define type, size and optionally I/O capacities [within service limits]
- Magnetic, SSD and Provisioned IOPS
- Mount to a single instance, similar to local drive
- Simplified Encryption options

Persistent and durable

- Redundant copies stored in single AZ
- Not permanently bound to a server instance and will survive server crash or shutdown

Snapshot capabilities for point-in-time backups

- Resizing and duplicating volumes
- Moving across AZs; Exporting across Regions

Performance metrics available through CloudWatch

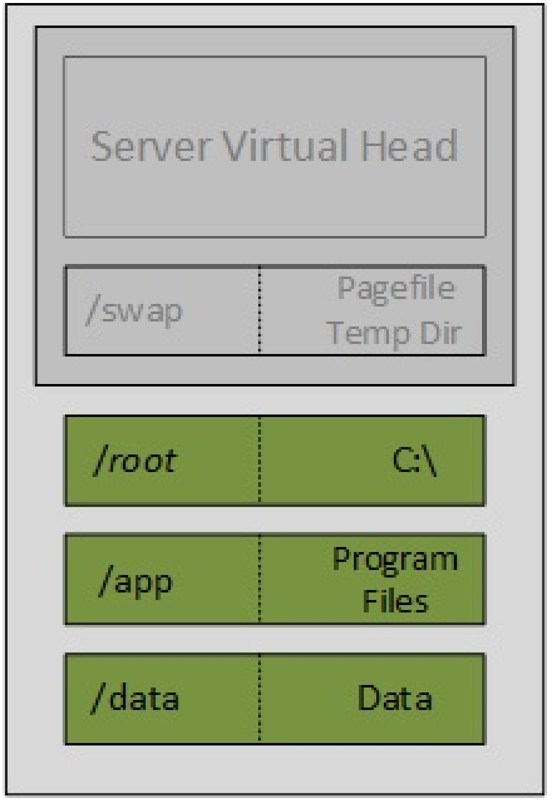

Elastic Block Store (EBS) – Best Practices

Recommended for applications

- Making frequent data changes
- Requiring consistent I/O performance
- Needing to persist data beyond server instance stop/start cycles
- Requiring fine-grain control of raw, unformatted data blocks

Define appropriate configuration options

- EBS Optimized instances can handle higher I/O bandwidth
- Underlying technology (Magnetic, General Purpose (SSD), Provisioned IOPS (SSD)

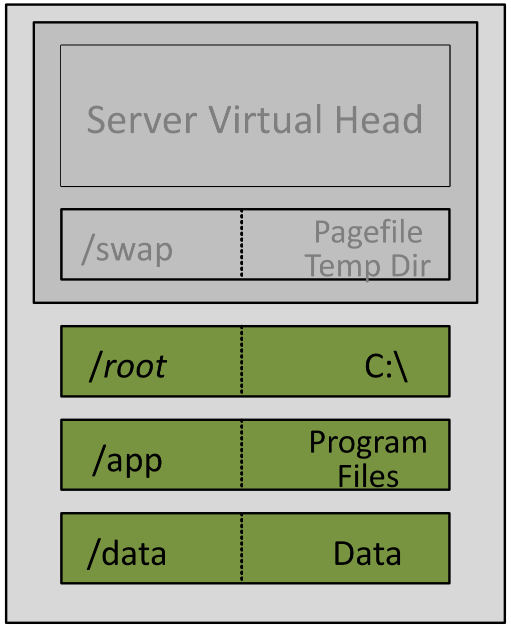

Ephemeral Drives (EC2 Instance Store) Overview

Block device attached to the host machine

- Available to server instance
- May be mounted and used for temporary storage
- No additional usage charges for disk space or I/O

Not redundant: no built-in RAID or snapshot function

Data loss will result if any of the following occur:

- Host server or instance crash
- Instance termination
- Disk failure

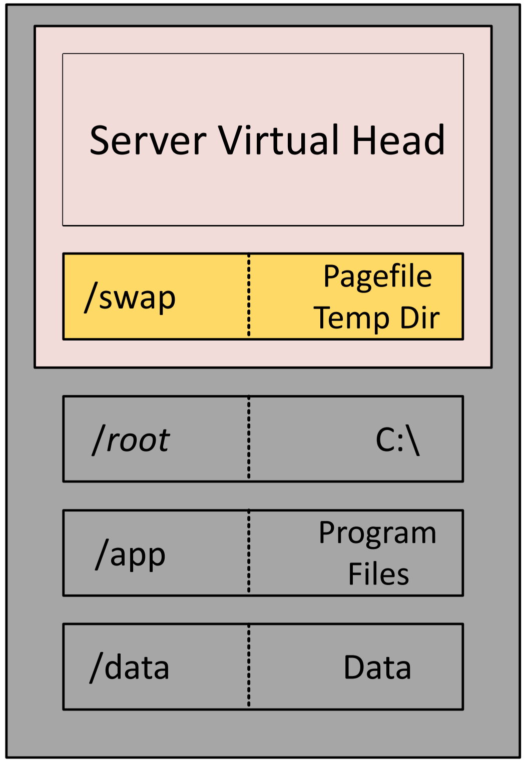

S3 - Simple Storage Service

Object storage container with virtually unlimited capacity

- Store files (objects) in containers (buckets)
- Redundant copies for high durability and reliability
- Available on the internet via REST requests directly or through SDK
- Multiple strategies to secure contents
  - Set permissions, access policies and optionally require MFA
  - Encryption: Server (simplified) or Client-side
  - Audit logging (optional) will record all access requests via APIs
- Built-in tools for managing versioning, object lifecycle and creating static websites
- Provides 99.999999999% durability (11 ‘9s’)
- Provides 99.99% availability

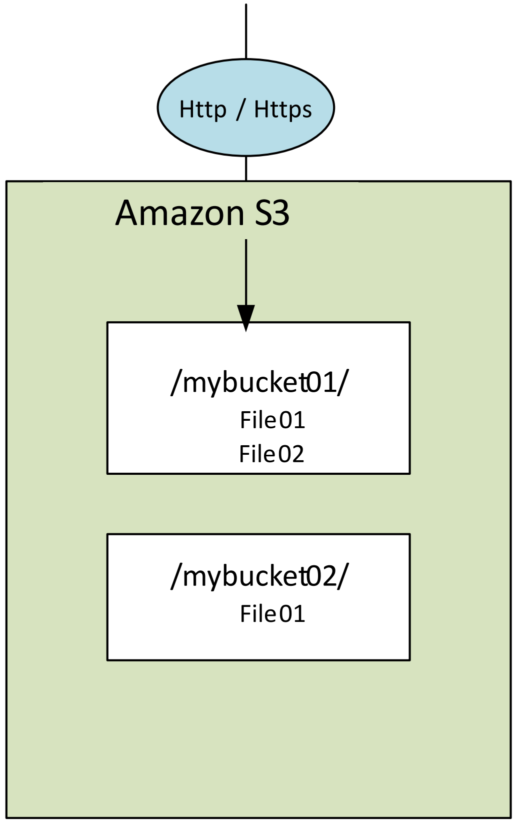

Amazon Glacier - Overview

Storage service optimized for reliable and low cost storage of archive data

- Data objects are securely archived, however not immediately accessible
- Create vaults (containers) to hold archives (any file based object)
- Upload archives programmatically
- Submit requests to retrieve archives. Available in about 4 hours
- Cost is approximately $.01/GB/Month plus modest API and retrieval charges [if applicable]

EFS – Elastic File System

Fully managed file server storage

- Uses NFS (v4.1) protocol
- Linux server only, Windows support planned for future release
- Can be mounted by 1,000s of EC2s
- Can be accessed from on-prem Data Center if using Direct Connect
- Highly available, redundant across multiple AZs

EFS – Comparing EFS and EBS

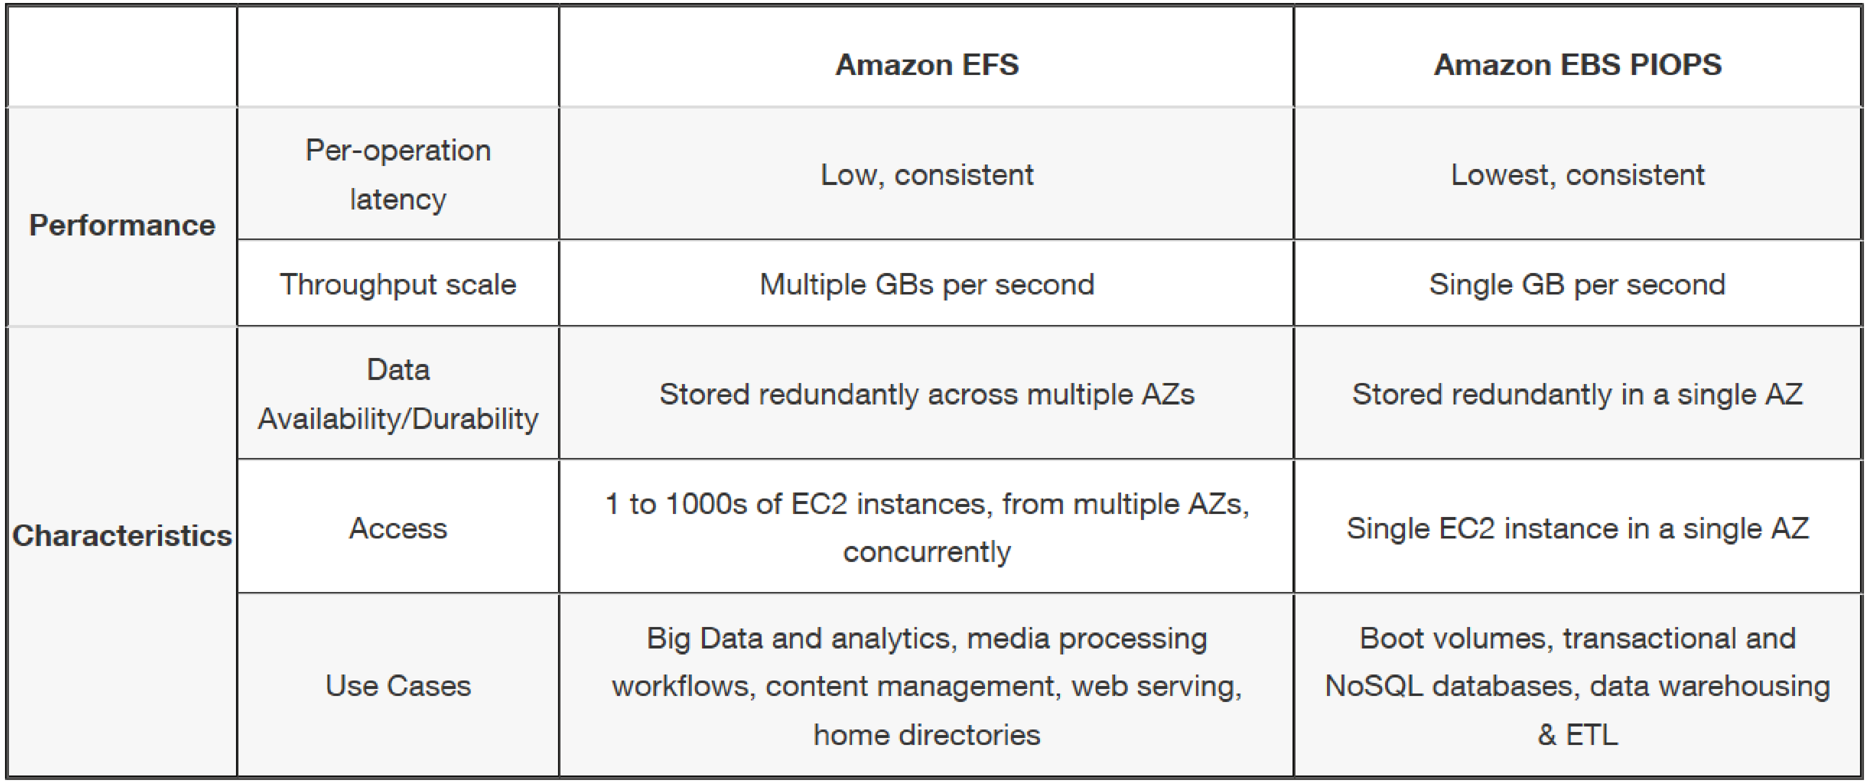

		Amazon EFS	Amazon EBS PIOPS
Performance	Per-operation latency	Low, consistent	Lowest, consistent
	Throughput scale	Multiple GBs per second	Single GB per second
Characteristics	Data Availability/Durability	Stored redundantly across multiple AZs	Stored redundantly in a single AZ
	Access	1 to 1000s of EC2 instances, from multiple AZs, concurrently	Single EC2 instance in a single AZ
	Use Cases	Big Data and analytics, media processing workflows, content management, web serving, home directories	Boot volumes, transactional and NoSQL databases, data warehousing & ETL

AWS Database Overwiew

AWS Structured Data Services

Deploying structured data systems (for example SQL, NoSQL and Data Warehouse applications) in a traditional environment may be complex, costly, and time consuming.

Amazon provides a set of structured data services with the following advantages:

- Simple to deploy, operate and scale
- Many common administrative and operational tasks are automated
- Pay-as-you-go pricing
- Support for a wide variety of standard and emerging application models

RDS - Relational Data Services

Fully managed relational database service offering popular platforms with the following key advantages:

- Amazon manages resource redundancy, software patching, backups, failure detection and recovery
- Ability to configure specific resources to cost-effectively scale your application
- Pay-as-you-go model offering included license or license portability [see fine print to ensure license compliance]
- Streamlined management options to easily configure highly available topologies, create database snapshots and deploy test instances

Relational Data Services

Key Concepts

- Database Instance
- Database Storage
- DB Instance Class

- 5 Platforms

1. Oracle
2. MS SQL
3. MySQL
4. PostgreSQL
5. MariaDB

AWS Aurora

AWS's version of MySQL database that is tailored for cloud environment

Key features:

- Architected for 99.99% availability
- Enterprise performance (5x) at 1/10 the cost
- Compatible with MySQL 5.6
- Automatically grows storage as needed, up to 64 TB
- Easy migration from MySQL
- Up to 15 Aurora Replicas in a region
- Cross-region replication
- Encryption in-transit and at rest
- Continuous backup to S3 (11 9's data durability)
- Fully managed

DynamoDB

Fully managed NoSQL database service offering the following key advantages:

- Seamless and virtually unlimited scalability conveniently managed automatically by Amazon
- Ability to define specific resource allocation limits to ensure predictable performance while containing costs
- Easy administration and well-supported development model
- Integration with other core Amazon data services (for example Redshift and EMR)

Redshift

Fully managed Enterprise-class data warehouse service offering the following advantages:

- High performance, massively parallel columnar storage architecture providing streamlined scalability
- Mainstream SQL query syntax allowing for rapid platform adoption
- Flexible node type and RI options allowing for workload alignment and cost efficiency

AWS Security Overview

Security and Compliance – Shared Security Model

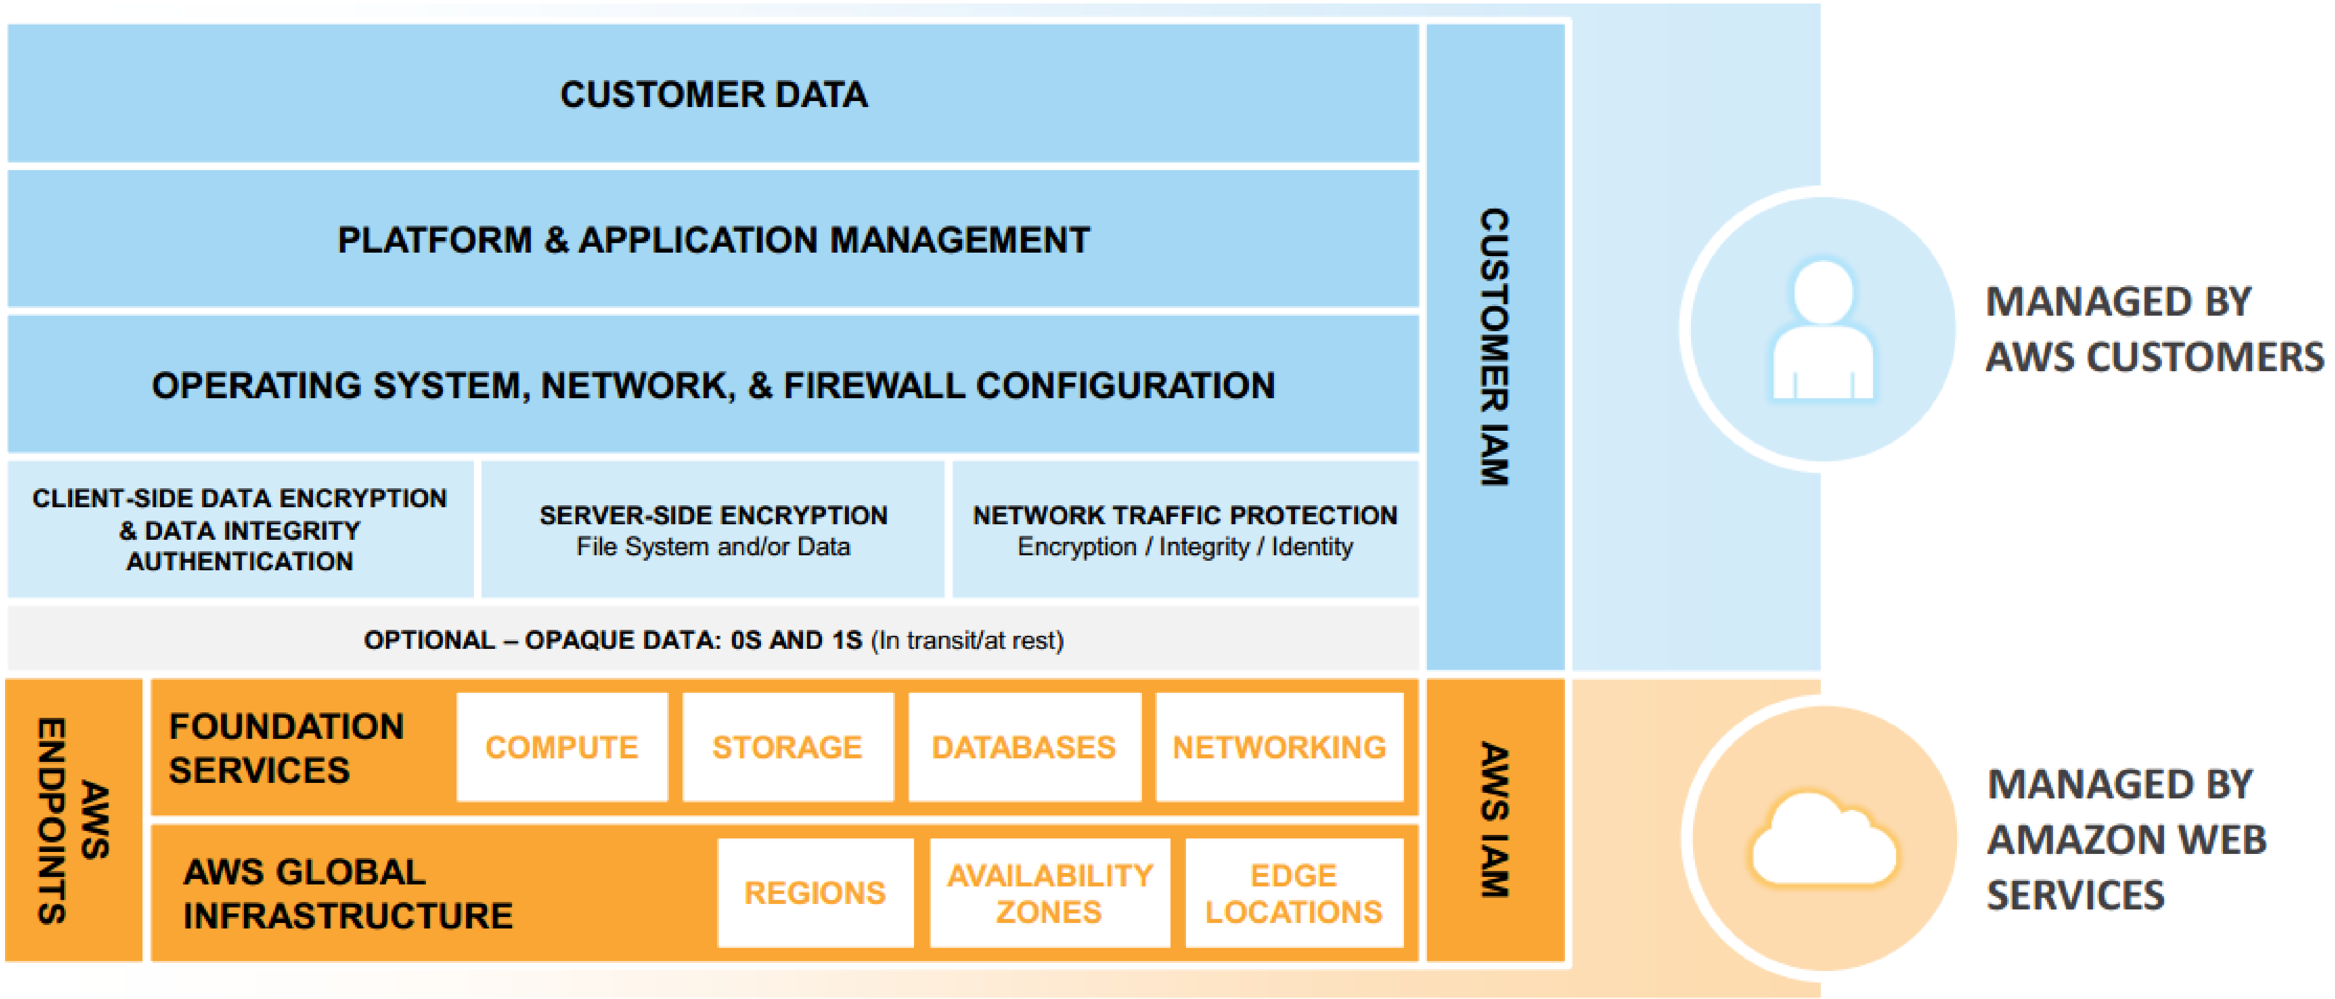

CUSTOMER DATA

PLATFORM & APPLICATION MANAGEMENT

OPERATING SYSTEM, NETWORK, & FIREWALL CONFIGURATION

CLIENT-SIDE DATA ENCRYPTION & DATA INTEGRITY AUTHENTICATION	SERVER-SIDE ENCRYPTION File System and/or Data	NETWORK TRAFFIC PROTECTION Encryption / Integrity / Identity

OPTIONAL – OPAQUE DATA: 0S AND 1S (In transit/at rest)

CUSTOMER IAM

AWS ENDPOINTS

FOUNDATION SERVICES	COMPUTE	STORAGE	DATABASES	NETWORKING
AWS GLOBAL INFRASTRUCTURE	REGIONS	AVAILABILITY ZONES	EDGE LOCATIONS	

AWS IAM

MANAGED BY AWS CUSTOMERS

MANAGED BY AMAZON WEB SERVICES

Security - Security Group/Firewall Rules

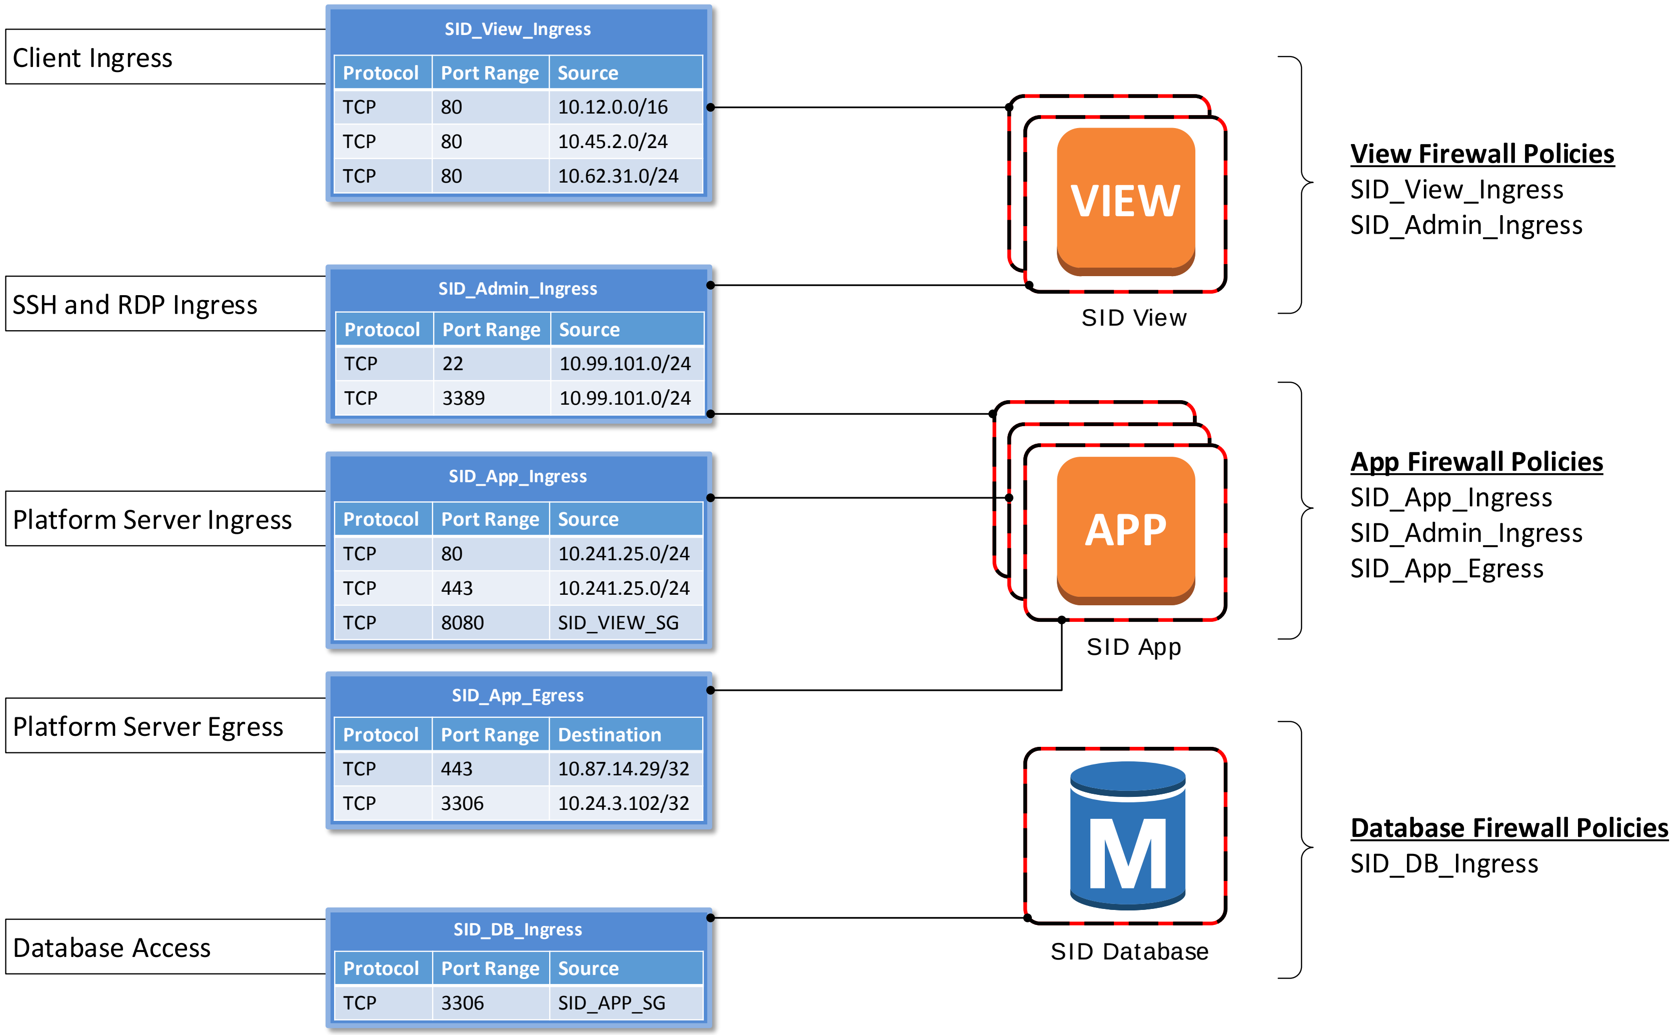

Client Ingress

SID_View_Ingress

Protocol	Port Range	Source
TCP	80	10.12.0.0/16
TCP	80	10.45.2.0/24
TCP	80	10.62.31.0/24

SSH and RDP Ingress

SID_Admin_Ingress

Protocol	Port Range	Source
TCP	22	10.99.101.0/24
TCP	3389	10.99.101.0/24

Platform Server Ingress

SID_App_Ingress

Protocol	Port Range	Source
TCP	80	10.241.25.0/24
TCP	443	10.241.25.0/24
TCP	8080	SID_VIEW_SG

Platform Server Egress

SID_App_Egress

Protocol	Port Range	Destination
TCP	443	10.87.14.29/32
TCP	3306	10.24.3.102/32

Database Access

SID_DB_Ingress

Protocol	Port Range	Source
TCP	3306	SID_APP_SG

View Firewall Policies
SID_View_Ingress
SID_Admin_Ingress

App Firewall Policies
SID_App_Ingress
SID_Admin_Ingress
SID_App_Egress

Database Firewall Policies
SID_DB_Ingress

AWS Certifications

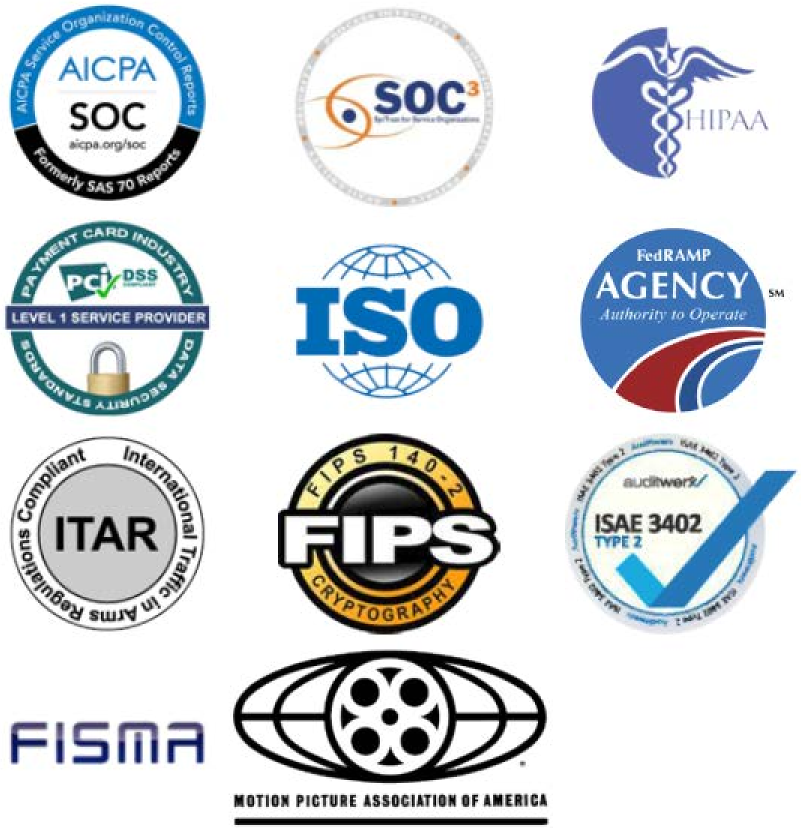

Contact Us

Vince Lo Faso
Solutions Architect
vlofaso@2ndwatch.com

Scott Turvey
Solutions Architect
sturvey@2ndwatch.com

Cameron Hatten
Regional Territory Manager
chatten@2ndwatch.com

General Information
1-888-317-7920
info@2ndwatch.com
www.2ndwatch.com

Locations

SEATTLE
NEW YORK
VIRGINIA
ATLANTA
PHILADELPHIA
HOUSTON
LIBERTY LAKE
LOS ANGELES
CHICAGO

Thank You | Questions?