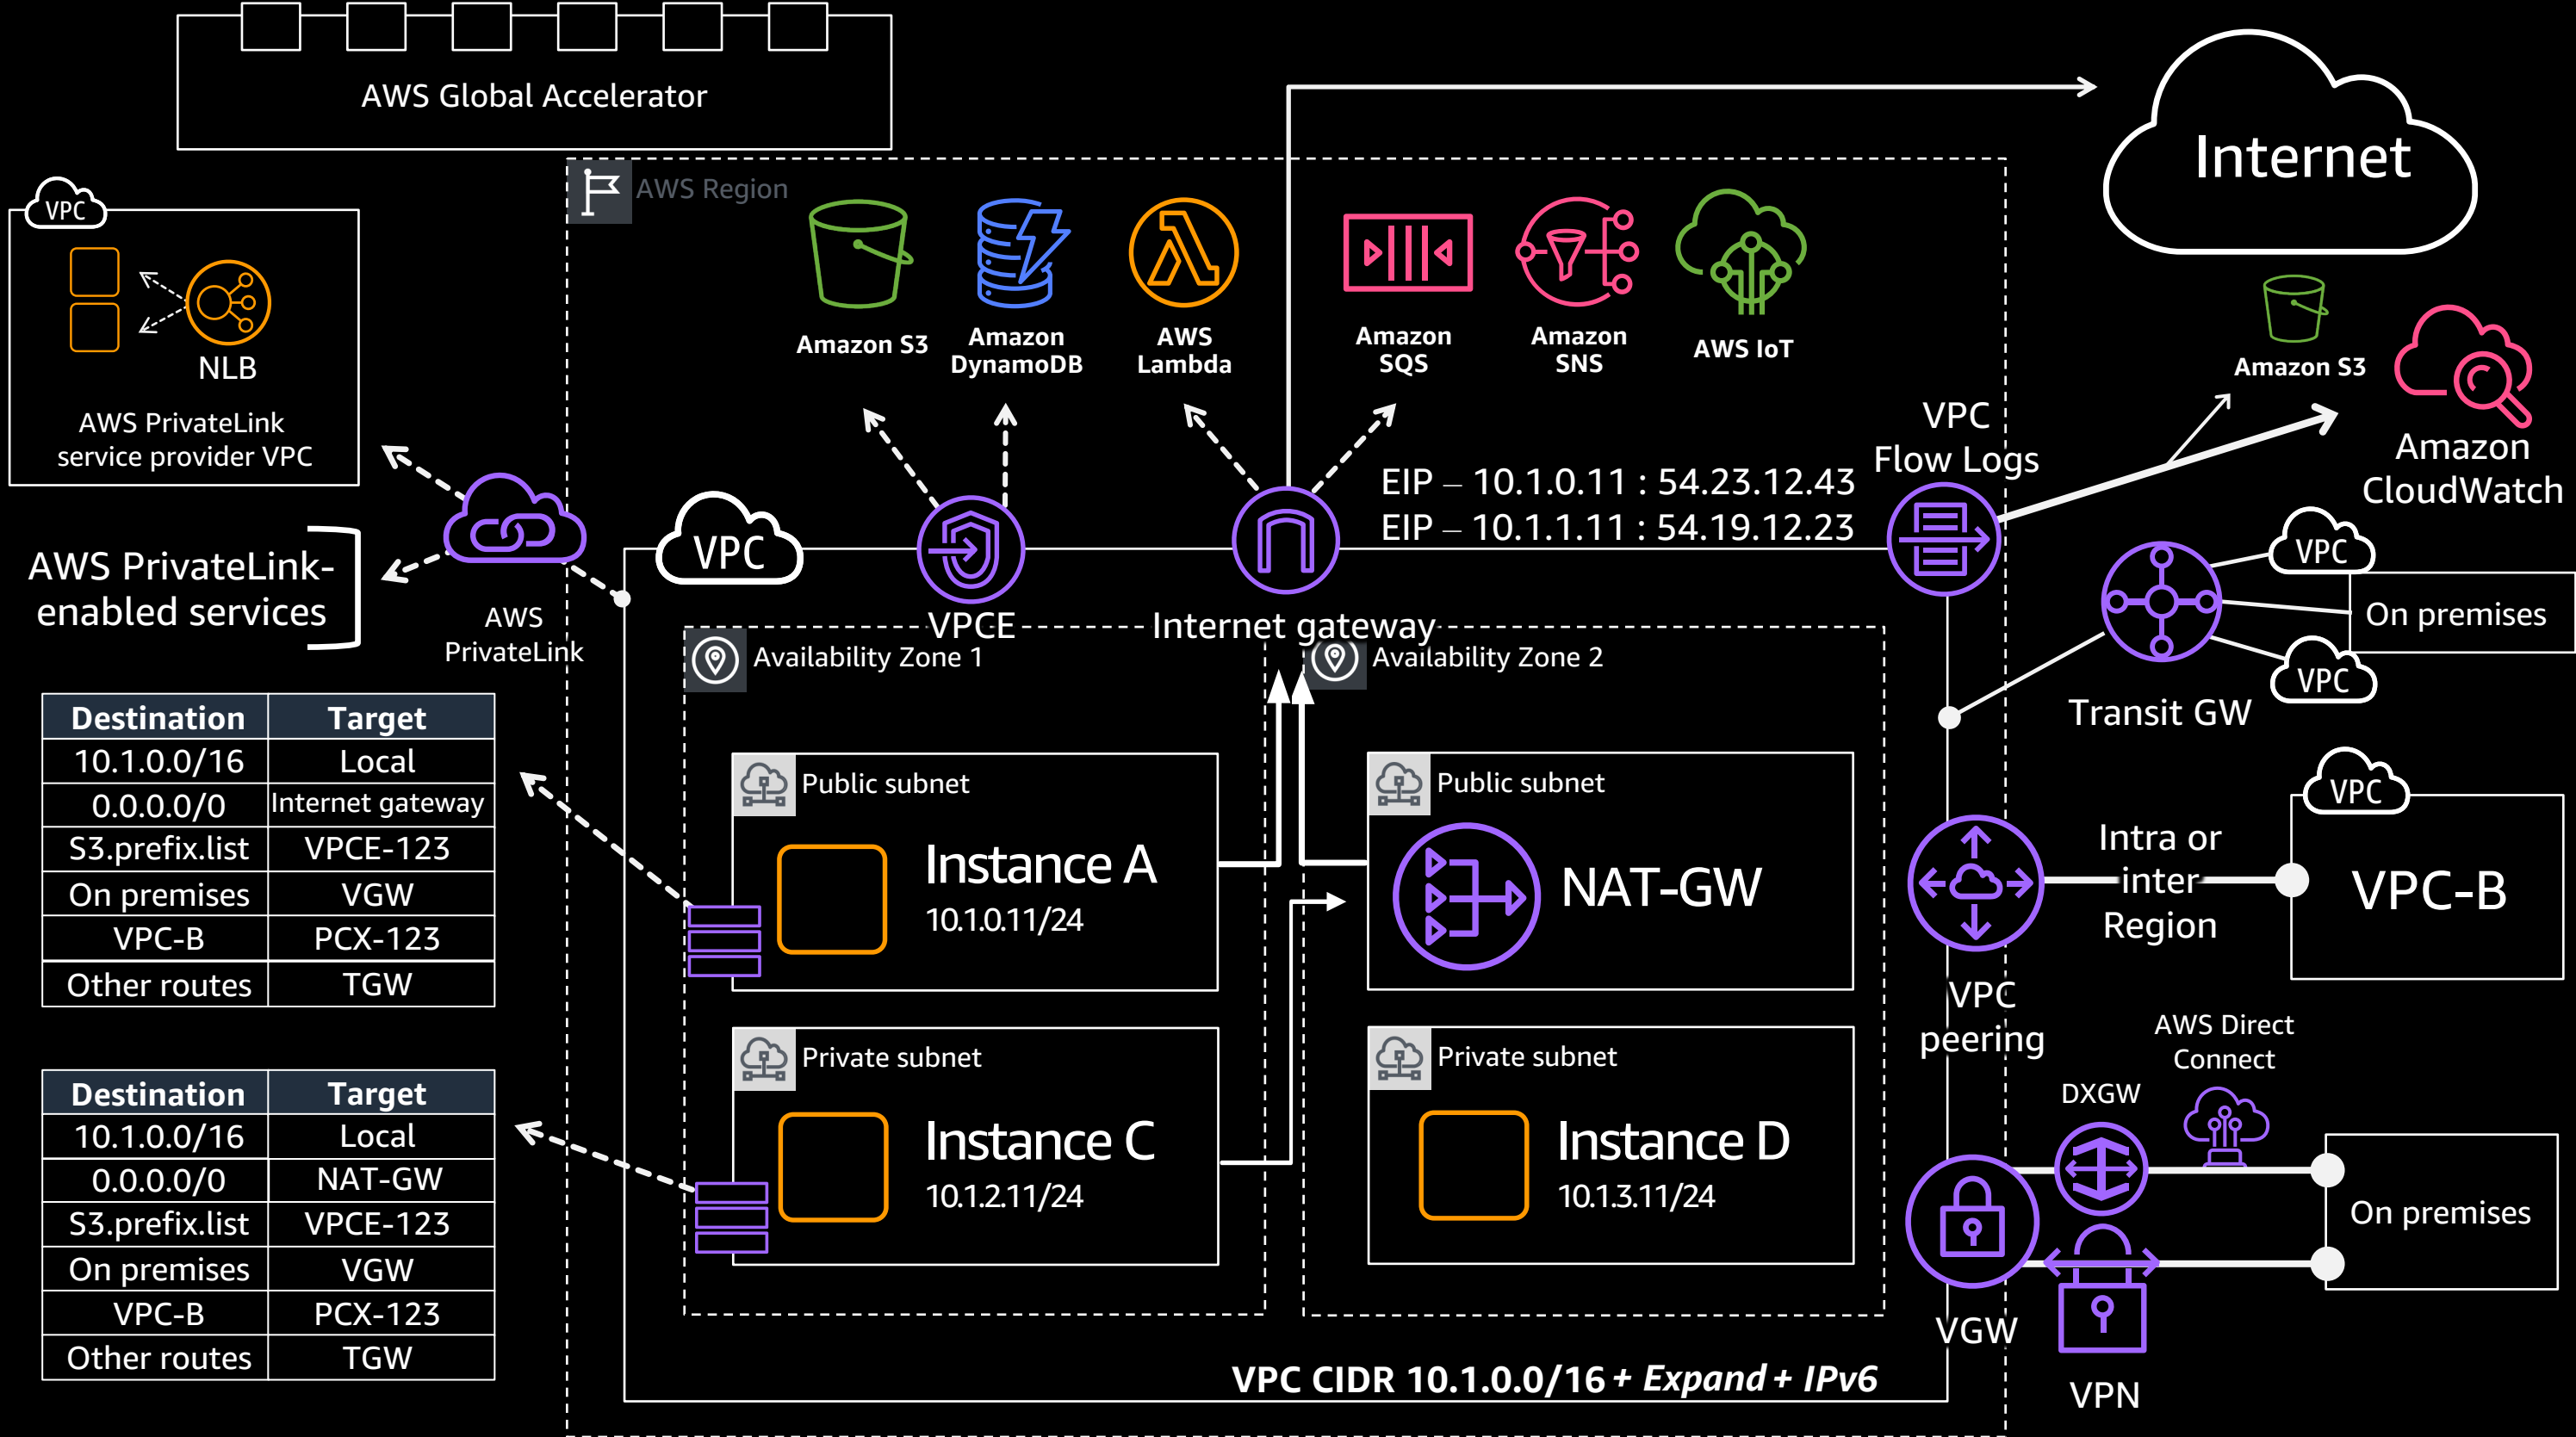


# Exploring the fundamentals of AWS networking

Francois van Rensburg  
Solutions Architect  
Amazon Web Services





That was the agenda  
for this session

# The fundamentals









## AWS Global Infrastructure

- 24 Regions with 76 Availability Zones
- 3 Regions coming soon:  
Spain, Jakarta and Osaka





- 205 Edge Locations 
- 11 Regional Edge Caches 







80+ Direct Connect  
Locations



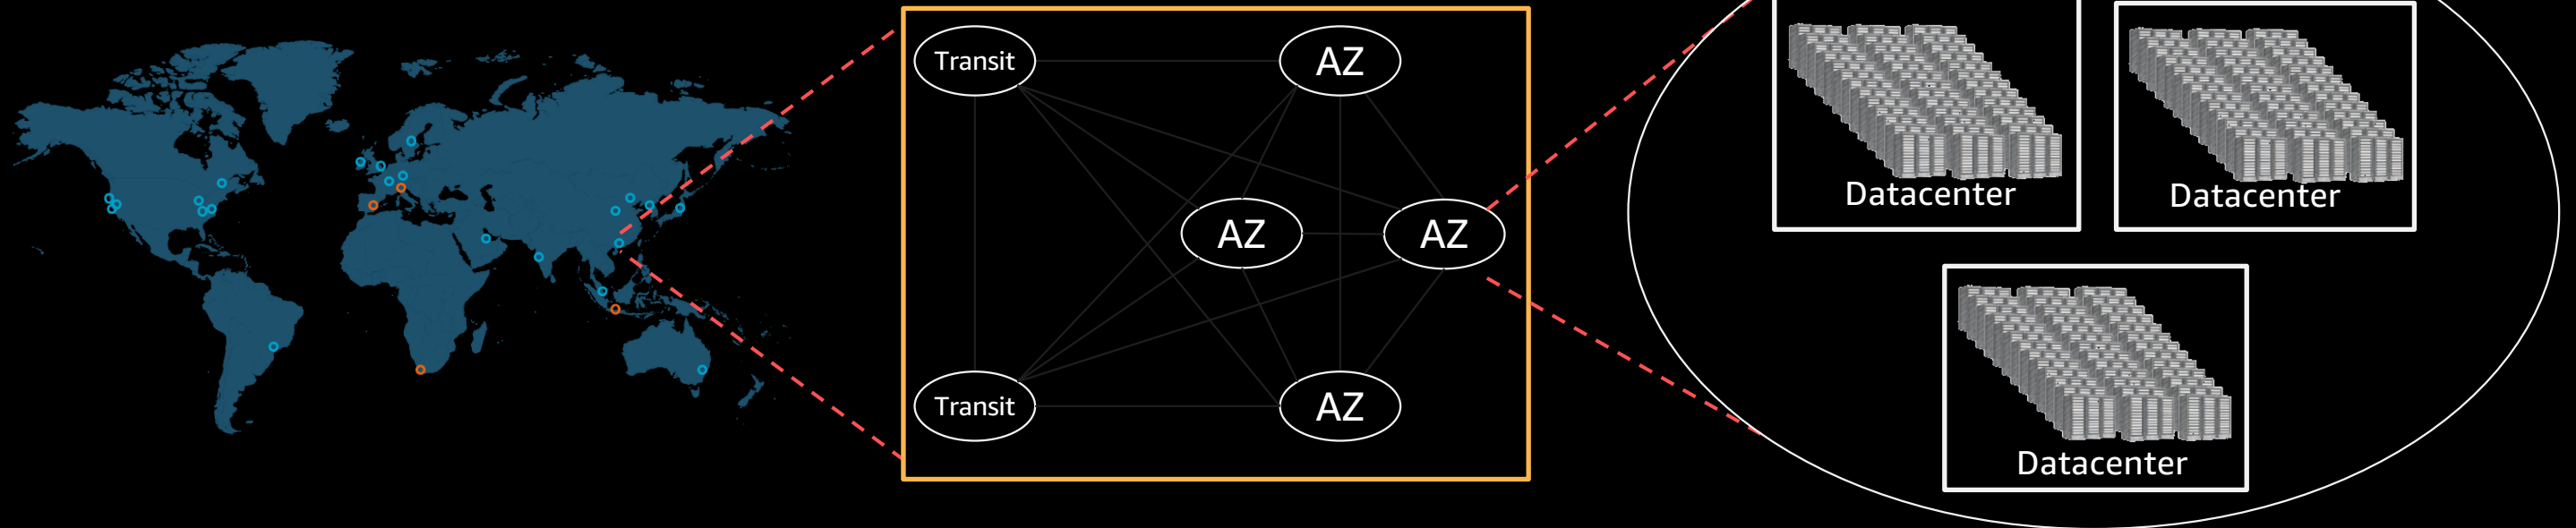


## AWS Global Network

- Redundant 100 GbE network
- Private network capacity between all AWS Region, except China

# AWS Region design

AWS Regions are comprised of multiple AZs for **high availability**, **high scalability**, and **high fault tolerance**. Applications and data are replicated in real time and consistent in the different AZs.



**A Region** is a physical location in the world where we have multiple **Availability Zones**.

**Availability Zones** consist of one or more discrete data centers, each with redundant power, networking, and connectivity, housed in separate facilities.

# VPC concepts and fundamentals



# What is a VPC?



# VPC concepts and fundamentals



IP  
addressing



Creating  
subnets



Routing in  
a VPC



DNS in-VPC  
with Amazon  
Route 53



Internet  
Access



# Choosing an IP address range

# Choosing an IP address range for your VPC



Avoid ranges that overlap with other networks to which you might connect

Size your VPC appropriately  
Largest /16 smallest /28

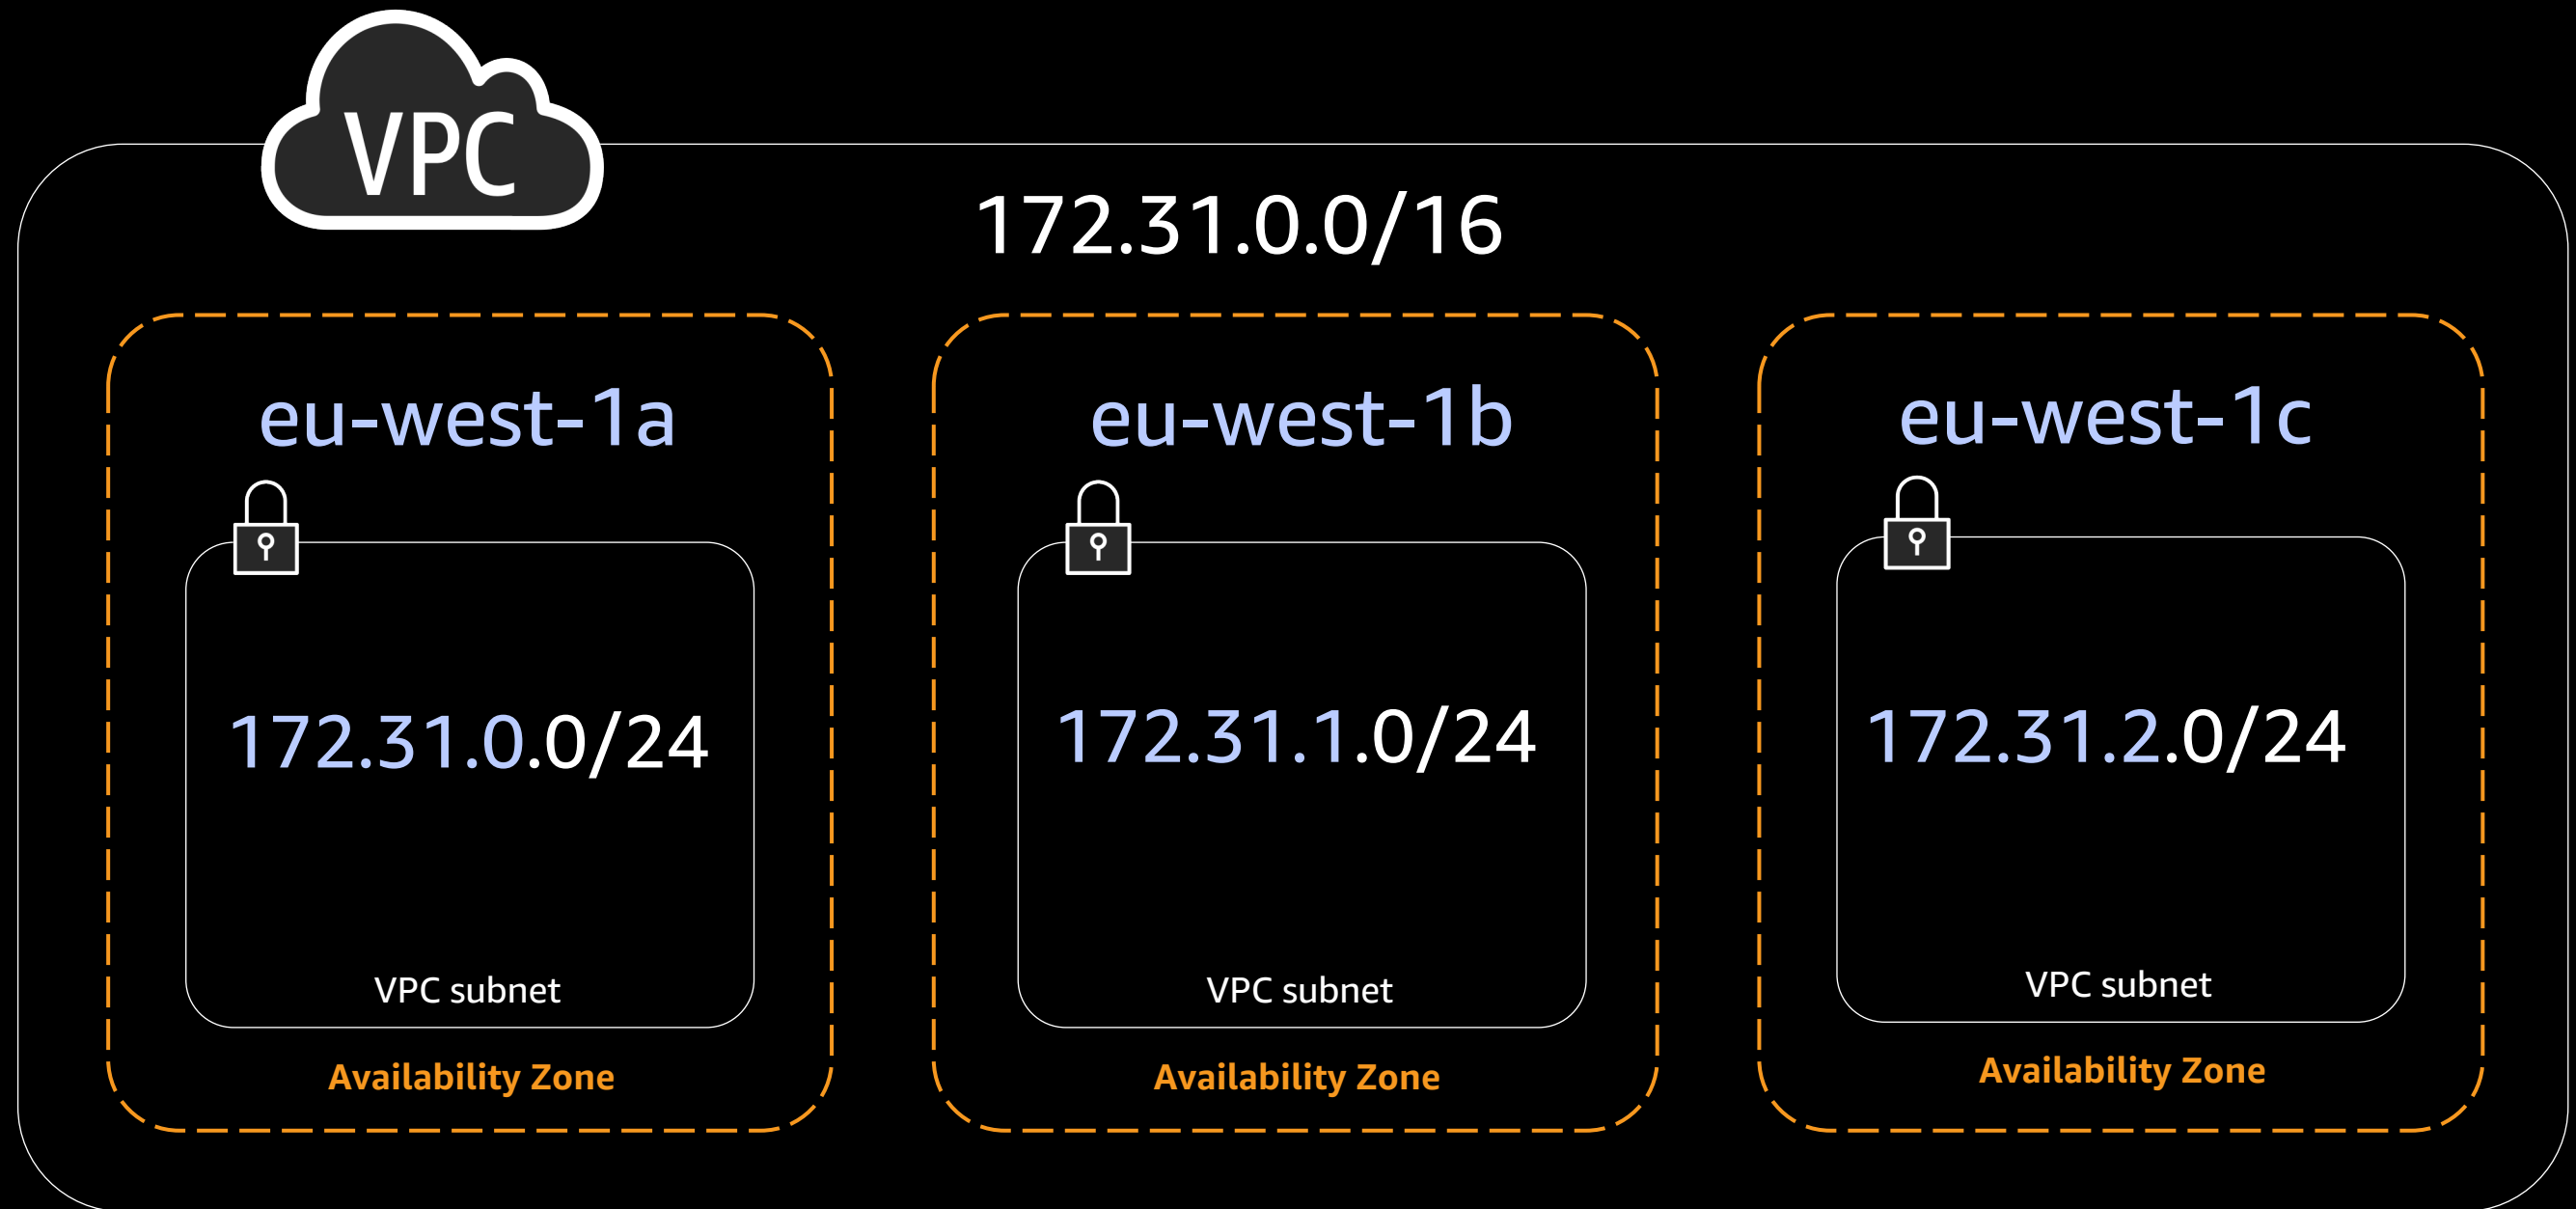
172.31.0.0/16

Recommended:  
RFC1918 range

Expandable


# Creating subnets in a VPC

# VPC subnets and Availability Zones

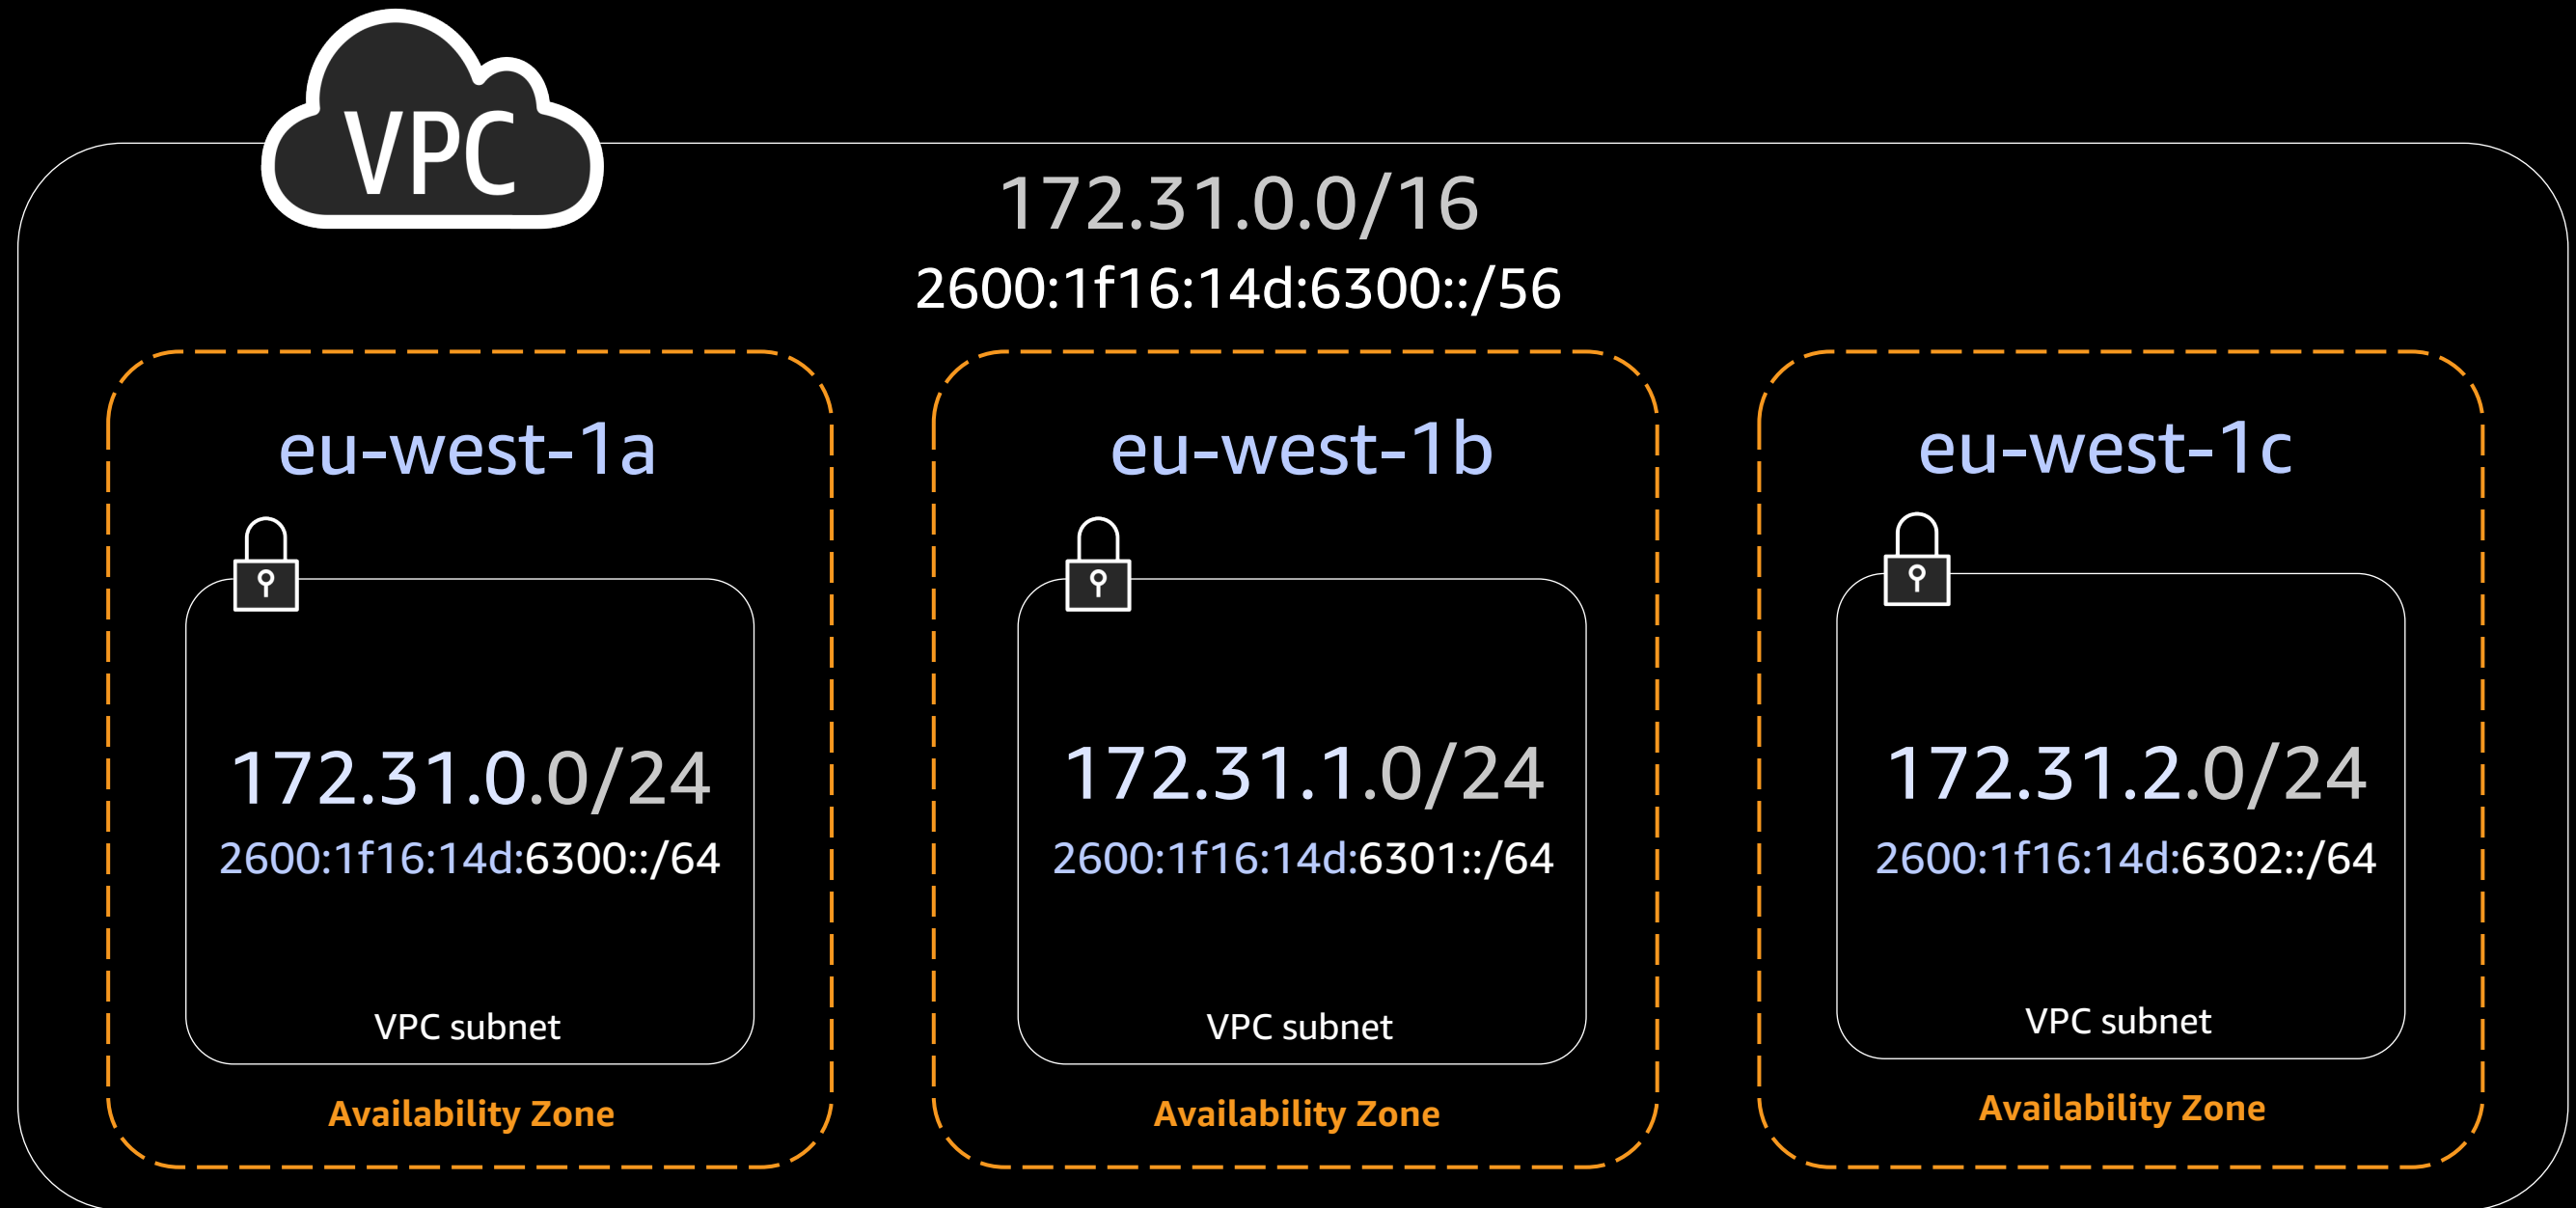


# IPv6 in your VPC

- Can have a dual-stack VPC by adding an IPv6 CIDR
- Fixed sizes for VPC and subnets
  - **/56 VPC** (4,722,366,482,869,645,213,696 addresses)
  - **/64 subnets** (18,446,744,073,709,551,616 addresses)

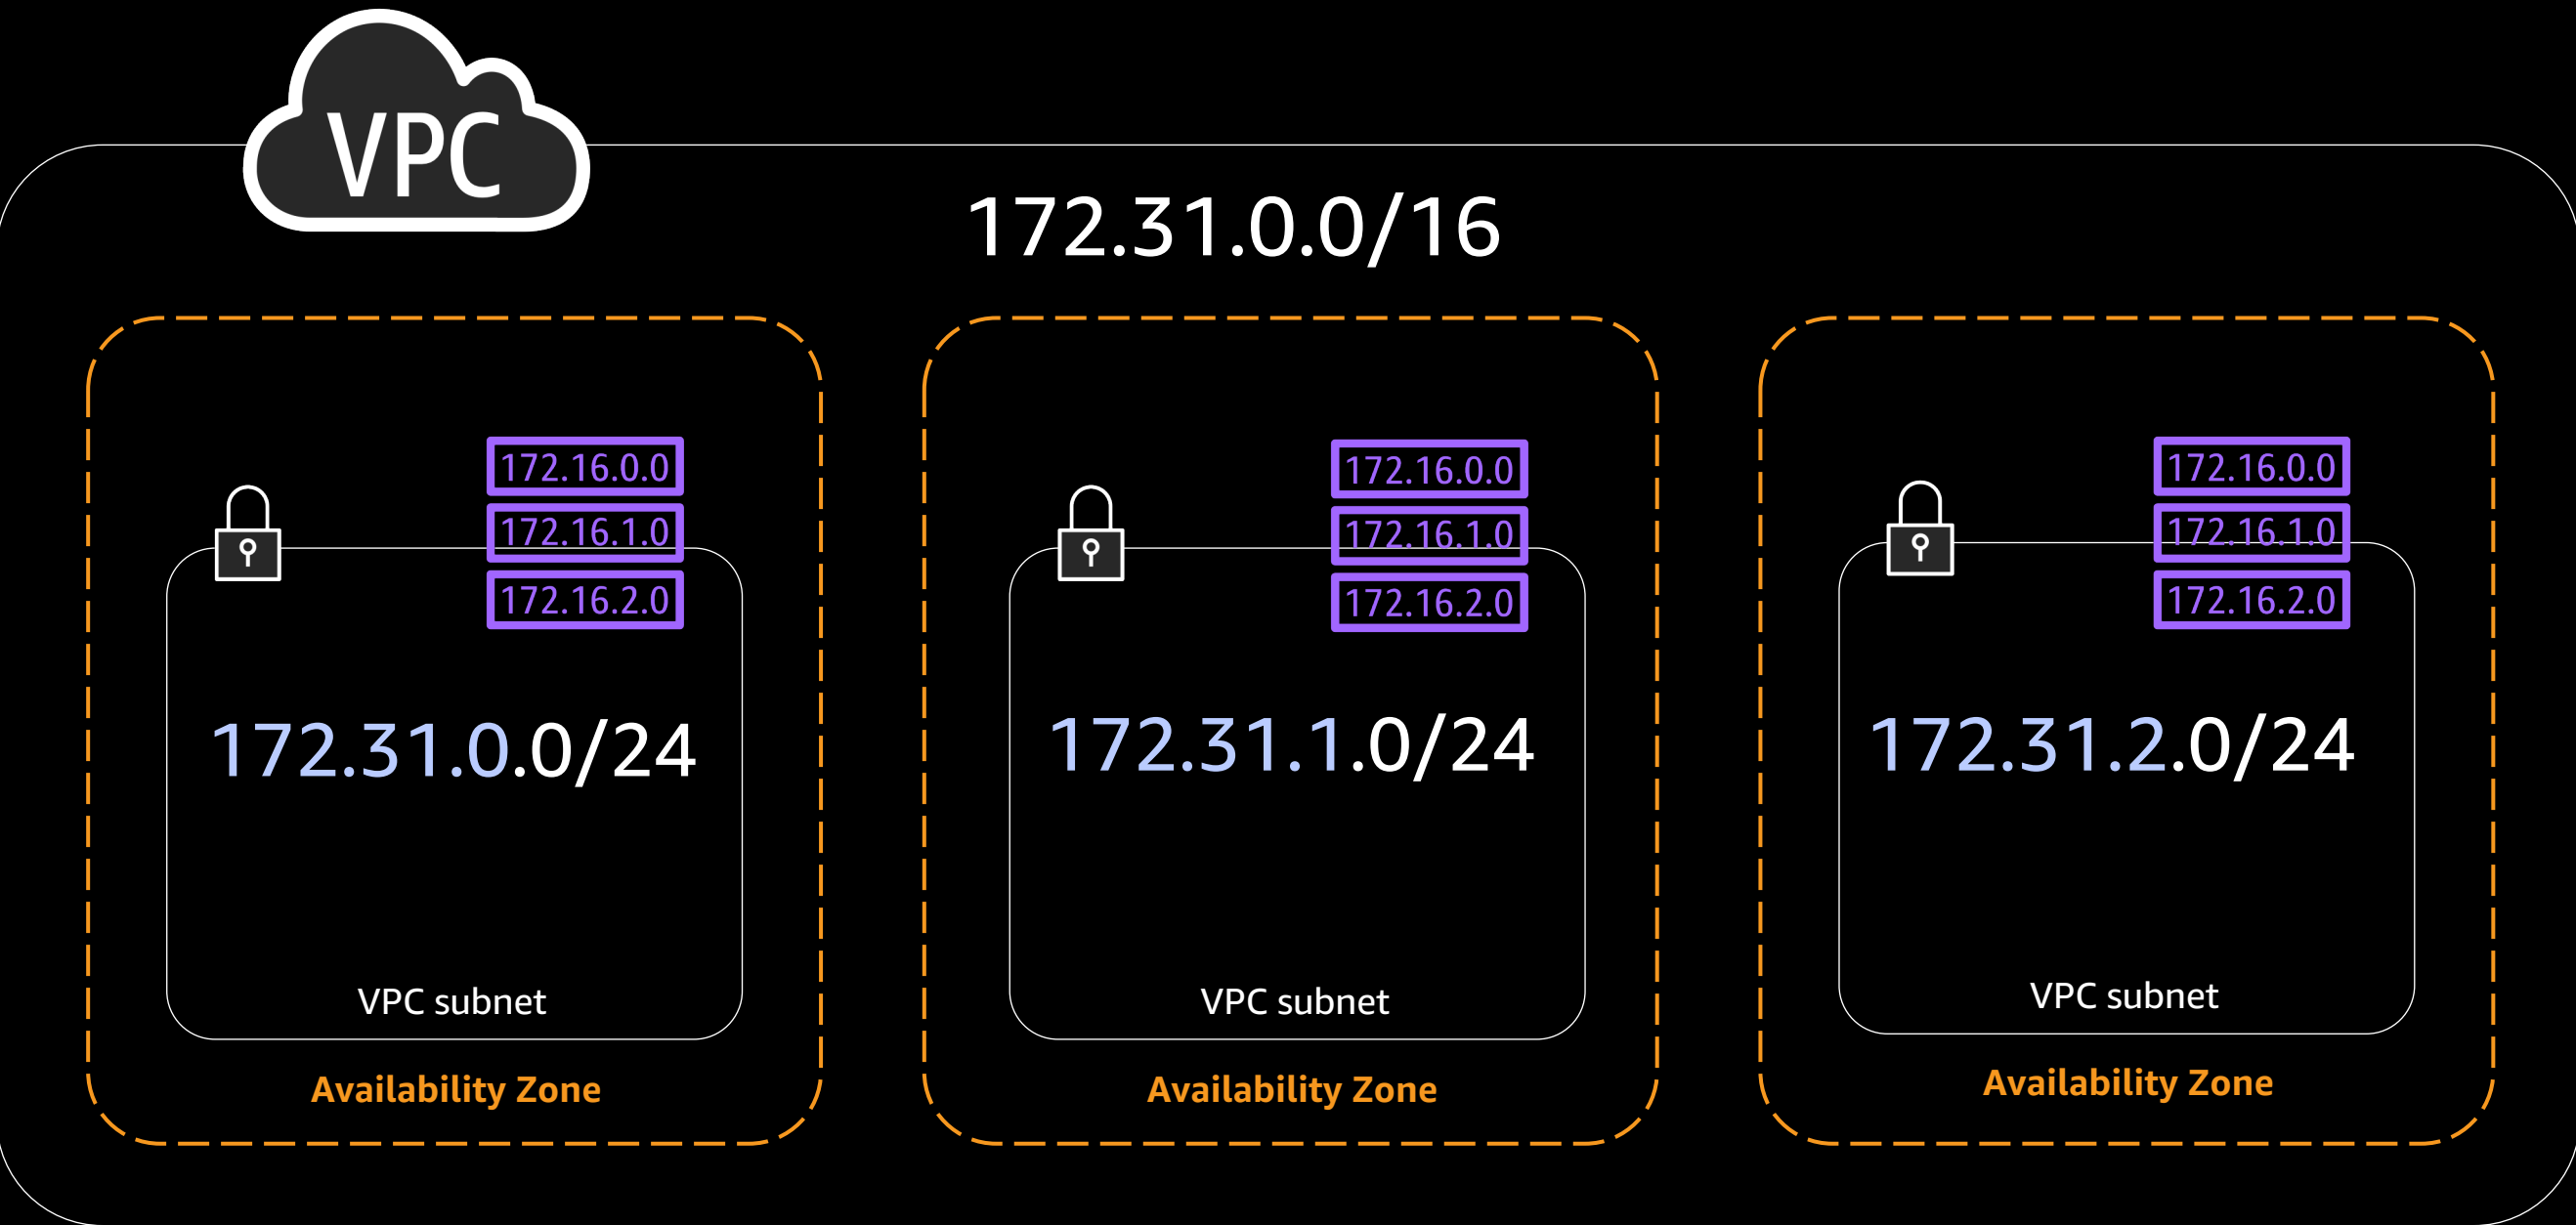


# VPC subnets and Availability Zones



# Routing in a VPC

# Route tables





Route Table: rtb-0ea57a71



- Summary
- Routes**
- Subnet Associations
- Route Propagation
- Tags

Edit routes

Traffic destined for my VPC stays in my VPC

Destination	Target	Status	Propagated
192.168.0.0/16	local	active	No
0.0.0.0/0	<a href="#">igw-062f547f</a>	active	No
10.0.0.0/16	<a href="#">pcx-4844e820</a>	active	No

# DNS in a VPC

# VPC DNS options

Name	VPC ID	State	IPv4 CIDR	IPv6 CIDR	DHCP options set
myVPC	vpc-0bcb5110cf0ce088b	available	172.31.0.0/16	2600:1f16:14d:6300::/56	dopt-c8cf28a1

**vpc-0bcb5110cf0ce088b | myVPC**

**Summary** | CIDR Blocks | Flow Logs | Tags

**VPC ID:** vpc-0bcb5110cf0ce088b | myVPC  
**State:** available  
**IPv4 CIDR:** 172.31.0.0/16  
**IPv6 CIDR:** 2600:1f16:14d:6300::/56  
**DHCP options set:** dopt-c8cf28a1  
**Route table:** rtb-0028d8ca88068723d

**Network ACL:** acl-0eb64...2bbc5a5  
**Tenancy:** Default

**DNS resolution:** yes  
**DNS hostnames:** yes

Have EC2 auto-assign DNS host names to instances

Use Amazon DNS server

# Amazon Route 53 Resolver for hybrid clouds

Step1  
Configure endpoints

Step2  
Configure inbound endpoint


Step3  
Configure outbound endpoint

Step4  
Create rule

Step5  
Review and create

## Configure endpoints


Endpoints provide the information that Resolver needs to route DNS queries from your VPCs to your network, from your network to your VPCs, or both.


 You are signed in to the following region: **us-west-2**  
To change your region use the region selector in the upper-right corner.


### Basic configuration

#### Direction of DNS queries [info](#)

You can configure endpoints for inbound DNS queries (to your VPC), outbound DNS queries (from your VPC), or both.

☒ Inbound and outbound  
Configure endpoints that allows DNS queries both to and from your VPC.  


☐ Inbound only  
Configure an endpoint that allows DNS queries to your VPC from an on-premises network or another VPC.  


☐ Outbound Only  
Configure an endpoint that allows DNS queries from your VPC to an on-premises network or another VPC.  


[Cancel](#) [Previous](#) [Next](#)

Conditional forwarding rules

Route 53 Resolver endpoints

# Internet Breakout


# The “5 Things” required for Internet traffic

1. Public IP Address
2. Internet Gateway Attached to a VPC
3. Route to an Internet Gateway
4. NACL Allow Rule
5. Security Group Allow Rule

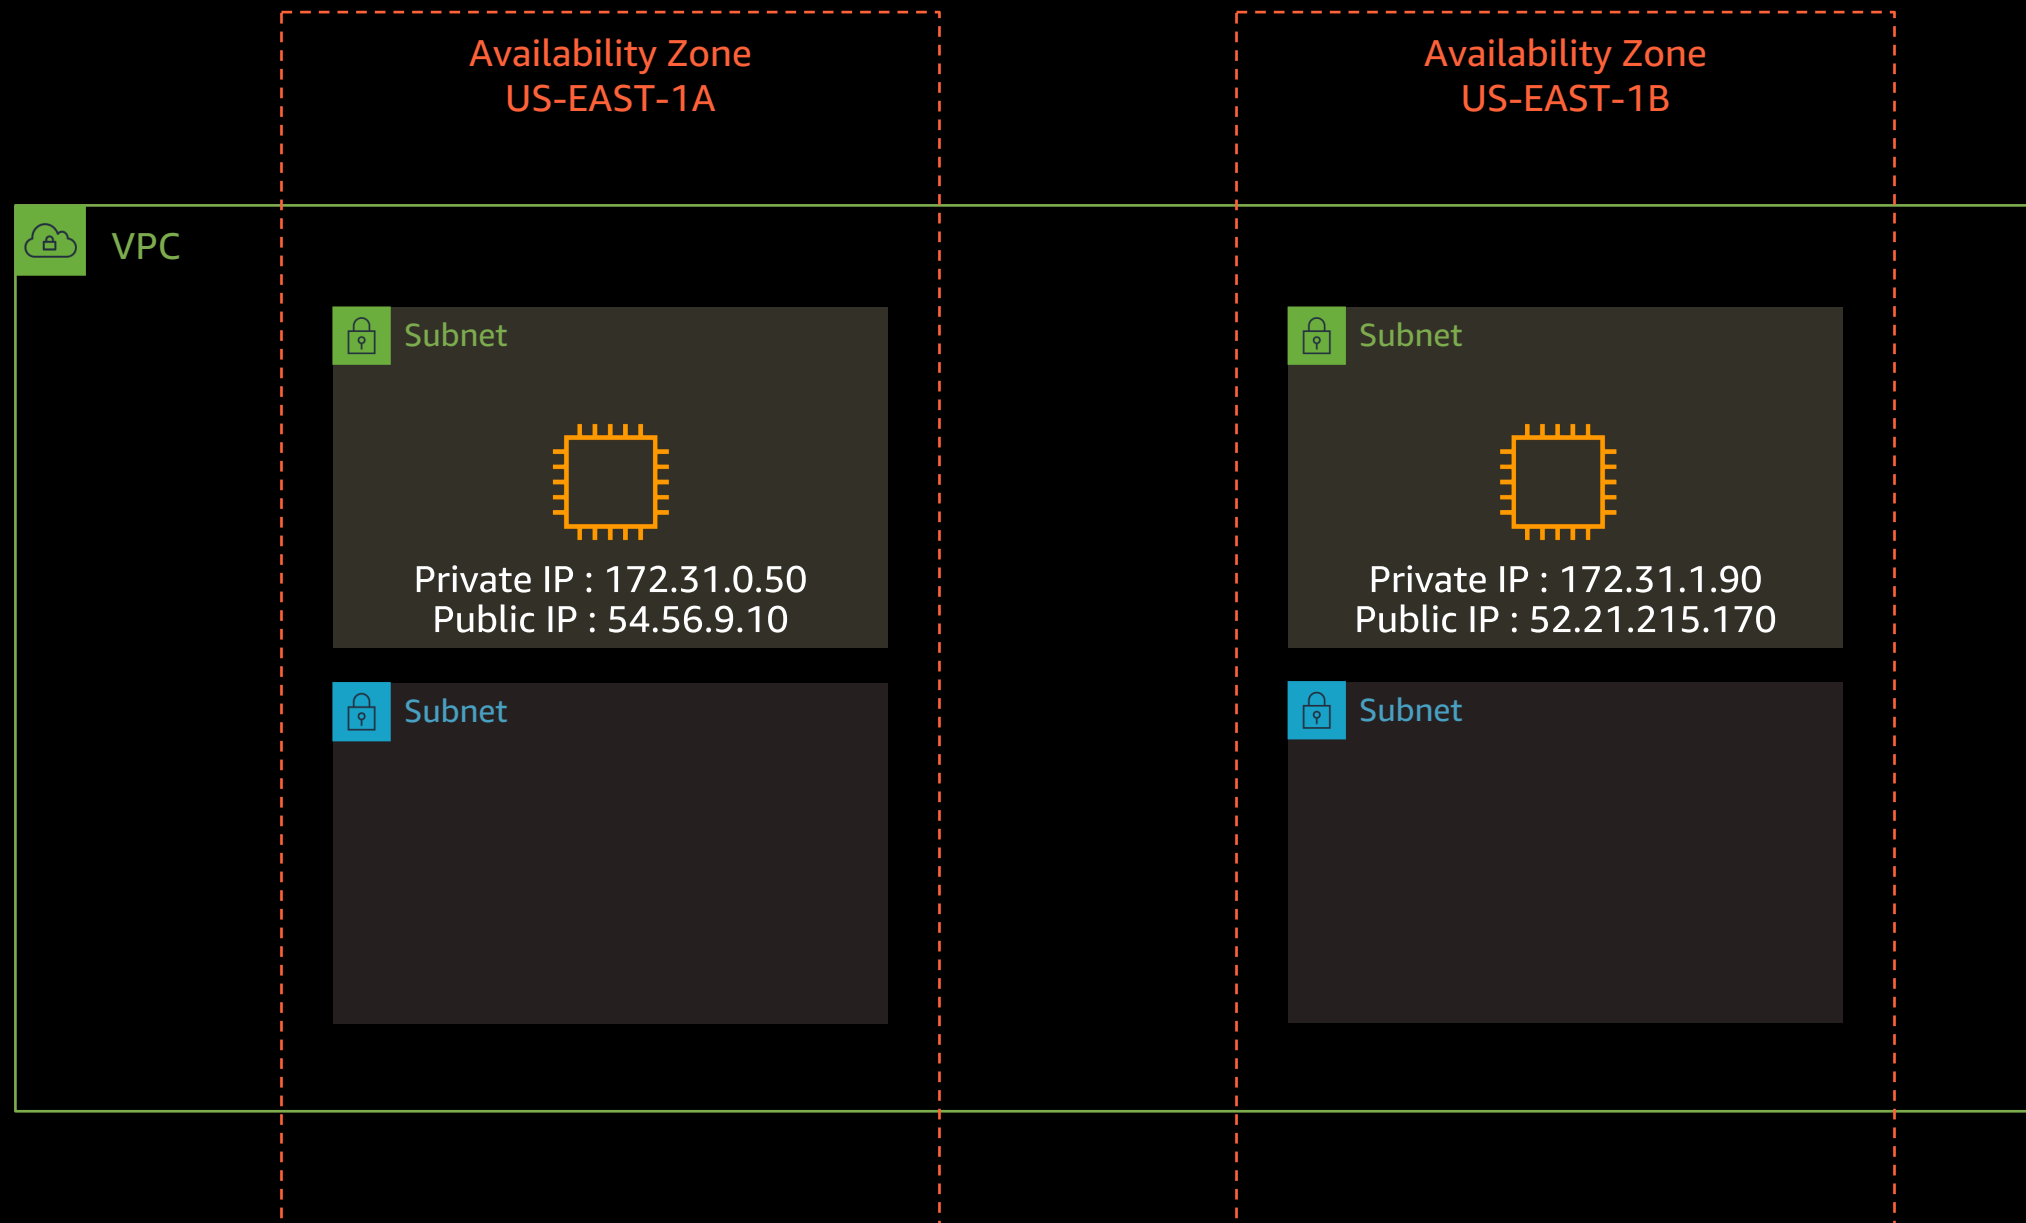


# Public IP addresses for your instances

- Auto-assign public IP addresses
- Elastic IP Addresses (EIP)
  - Amazon EIP Pool
  - Bring Your Own IP (BYOIP) Pool




# Public IP addresses





# Internet access



172.16.0.0

172.16.1.0

172.16.2.0

Create internet gateway

Actions

Filter by tags and attributes or search by keyword

<div></div>	Name	ID	State	VPC
<div></div>		igw-09ef761d872b...	attached	vpc-0bcb5110cf0c...

Destination	Target	Status	Propagated
172.31.0.0/16	local	Active	No
2600:1f16:14d:6300::/56	local	Active	No
0.0.0.0/0	igw-09ef761d872bd7540	Active	No
::/0	igw-09ef761d872bd7540	Active	No

*“To get to the IPv4 Internet (0.0.0.0/0) go via the Internet Gateway (IGW)”*

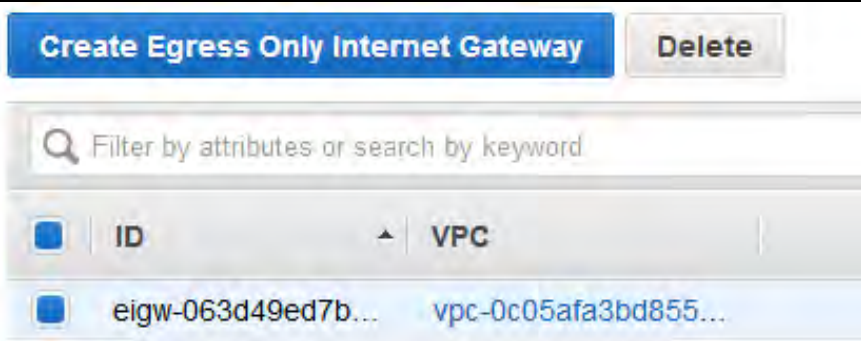
*“To get to the IPv6 Internet (::/0) go via the Internet Gateway (IGW)”*



# Internet access



172.16.0.0  
172.16.1.0  
172.16.2.0



Destination	Target	Status	Propagated
172.31.0.0/16	local	Active	No
2600:1f16:14d:6300::/56	local	Active	No
0.0.0.0/0	igw-09ef761d872bd7540	Active	No
::/0	eigw-063d49ed7bb0f8c36	Active	No

*“To get to the IPv6 Internet (::/0) go via the Egress Only Internet Gateway (EIGW)”*



# Different routes for different subnets

172.16.0.0

172.16.1.0

172.16.2.0

## Public subnet

Destination	Target	Status	Propagated
172.31.0.0/16	local	Active	No
2600:1f16:14d:6300::/56	local	Active	No
0.0.0.0/0	igw-09ef761d872bd7540	Active	No
::/0	igw-09ef761d872bd7540	Active	No

*“To get to the Internet go via the Internet Gateway (IGW)”*

## Private subnet

Destination	Target	Status	Propagated
172.31.0.0/16	local	Active	No
2600:1f16:14d:6300::/56	local	Active	No

*“To get to anything in the VPC – stay local. No route anywhere else.”*




# Public & private subnets

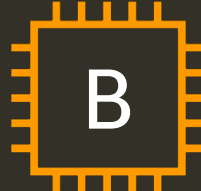
 Private subnet



A


Private IP : 172.31.128.75

 Public subnet



B

Private IP : 172.31.0.50  
Public IP : 54.56.9.10



Destination	Target	Status	Propagated
172.31.0.0/16	local	Active	No
2600:1f16:14d:6300::/56	local	Active	No

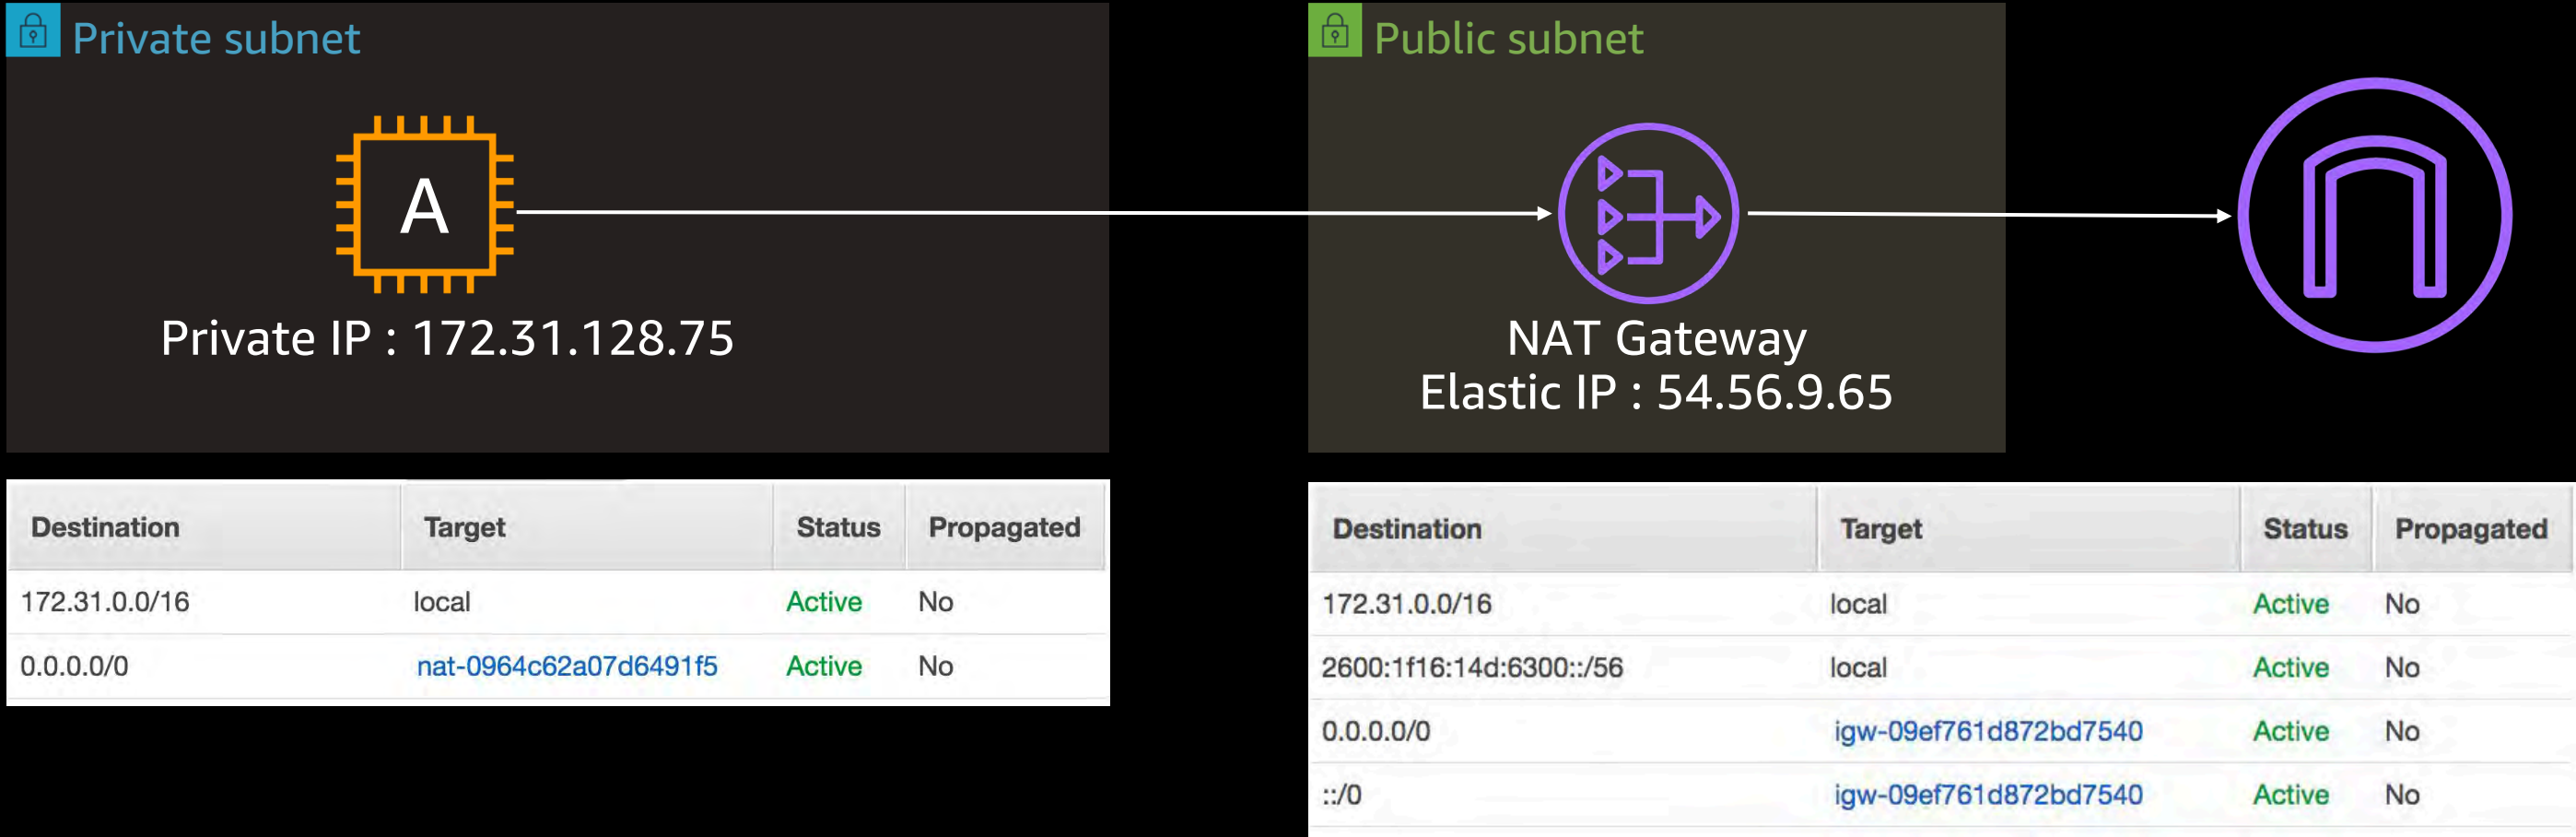
Destination	Target	Status	Propagated
172.31.0.0/16	local	Active	No
2600:1f16:14d:6300::/56	local	Active	No
0.0.0.0/0	igw-09ef761d872bd7540	Active	No
::/0	igw-09ef761d872bd7540	Active	No

*“Instance A has a path to and from Instance B.”*

*“Instance B has a path to and from the Internet.”*



# Network Address Translation (NAT) Gateway



*The Route Table for the Private Subnet says to send all IPv4 Internet Traffic to the NAT Gateway.*

*The NAT Gateway translates all traffic it receives such that it appears to come from itself.*

*The Route Table for the Public Subnet says to send all Internet Traffic to the Internet Gateway.*






# Network security

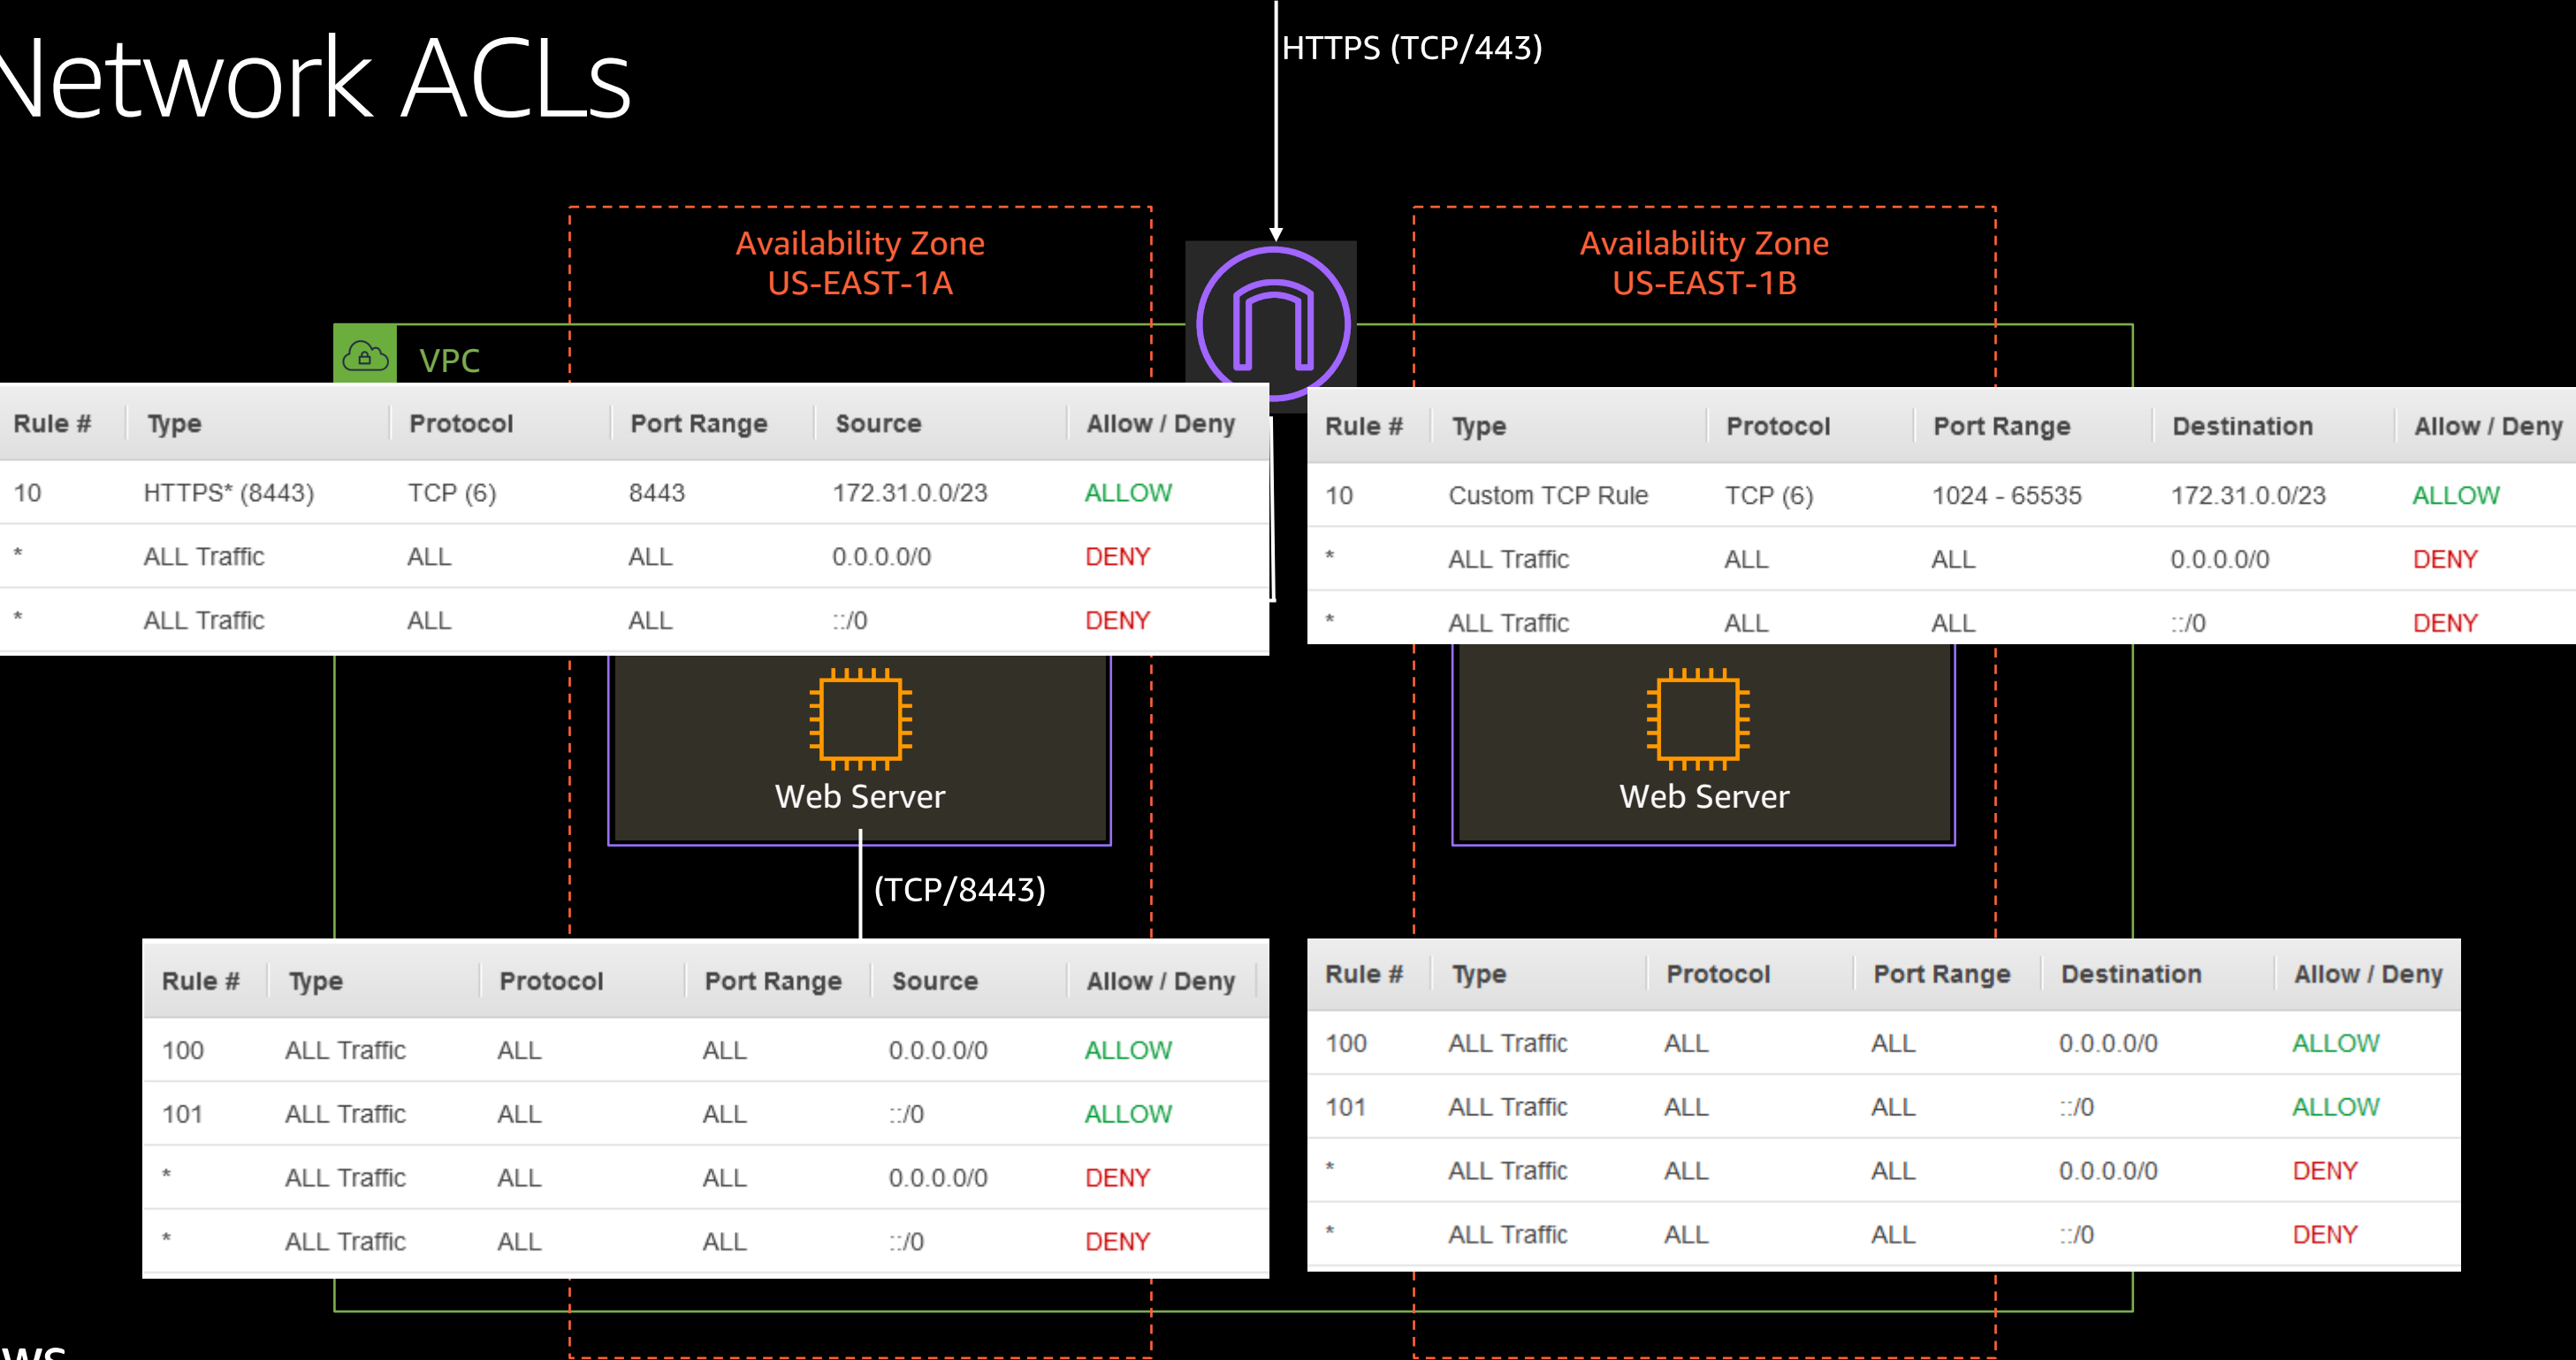


# Network security

- Network ACLs
- Security Groups
- VPC Flow Logs
- Amazon VPC Traffic Mirroring

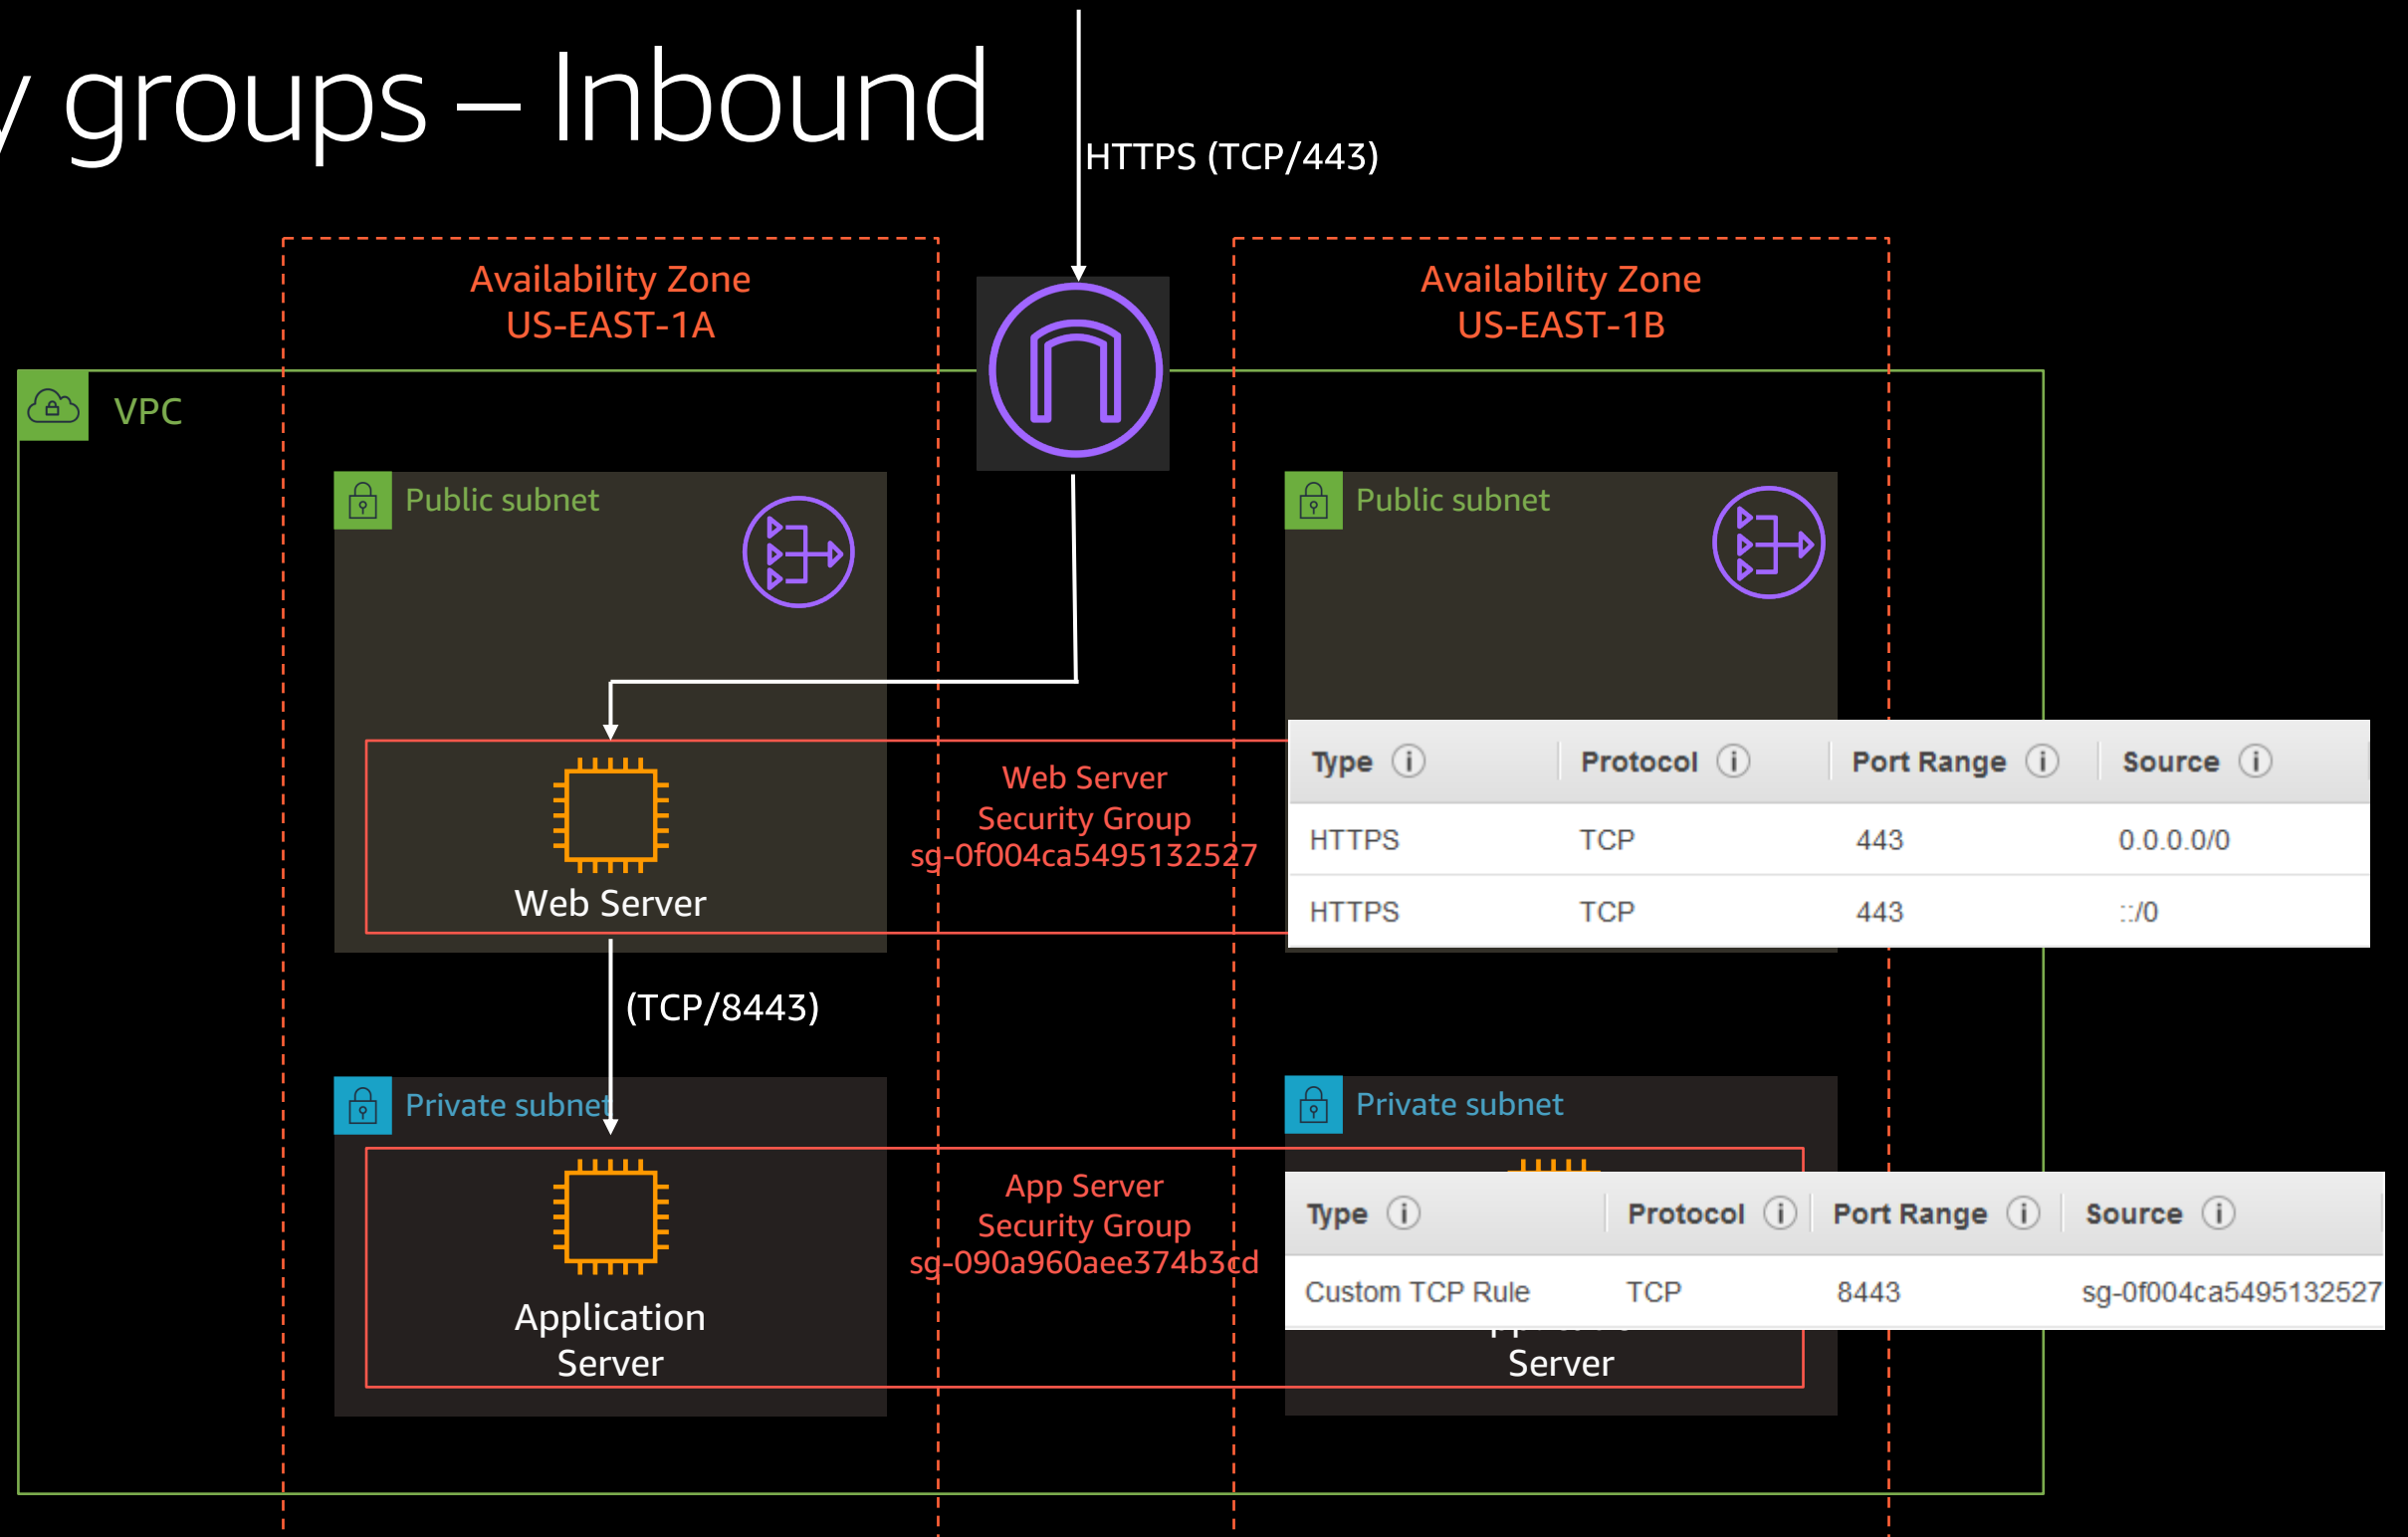


# Network ACLs

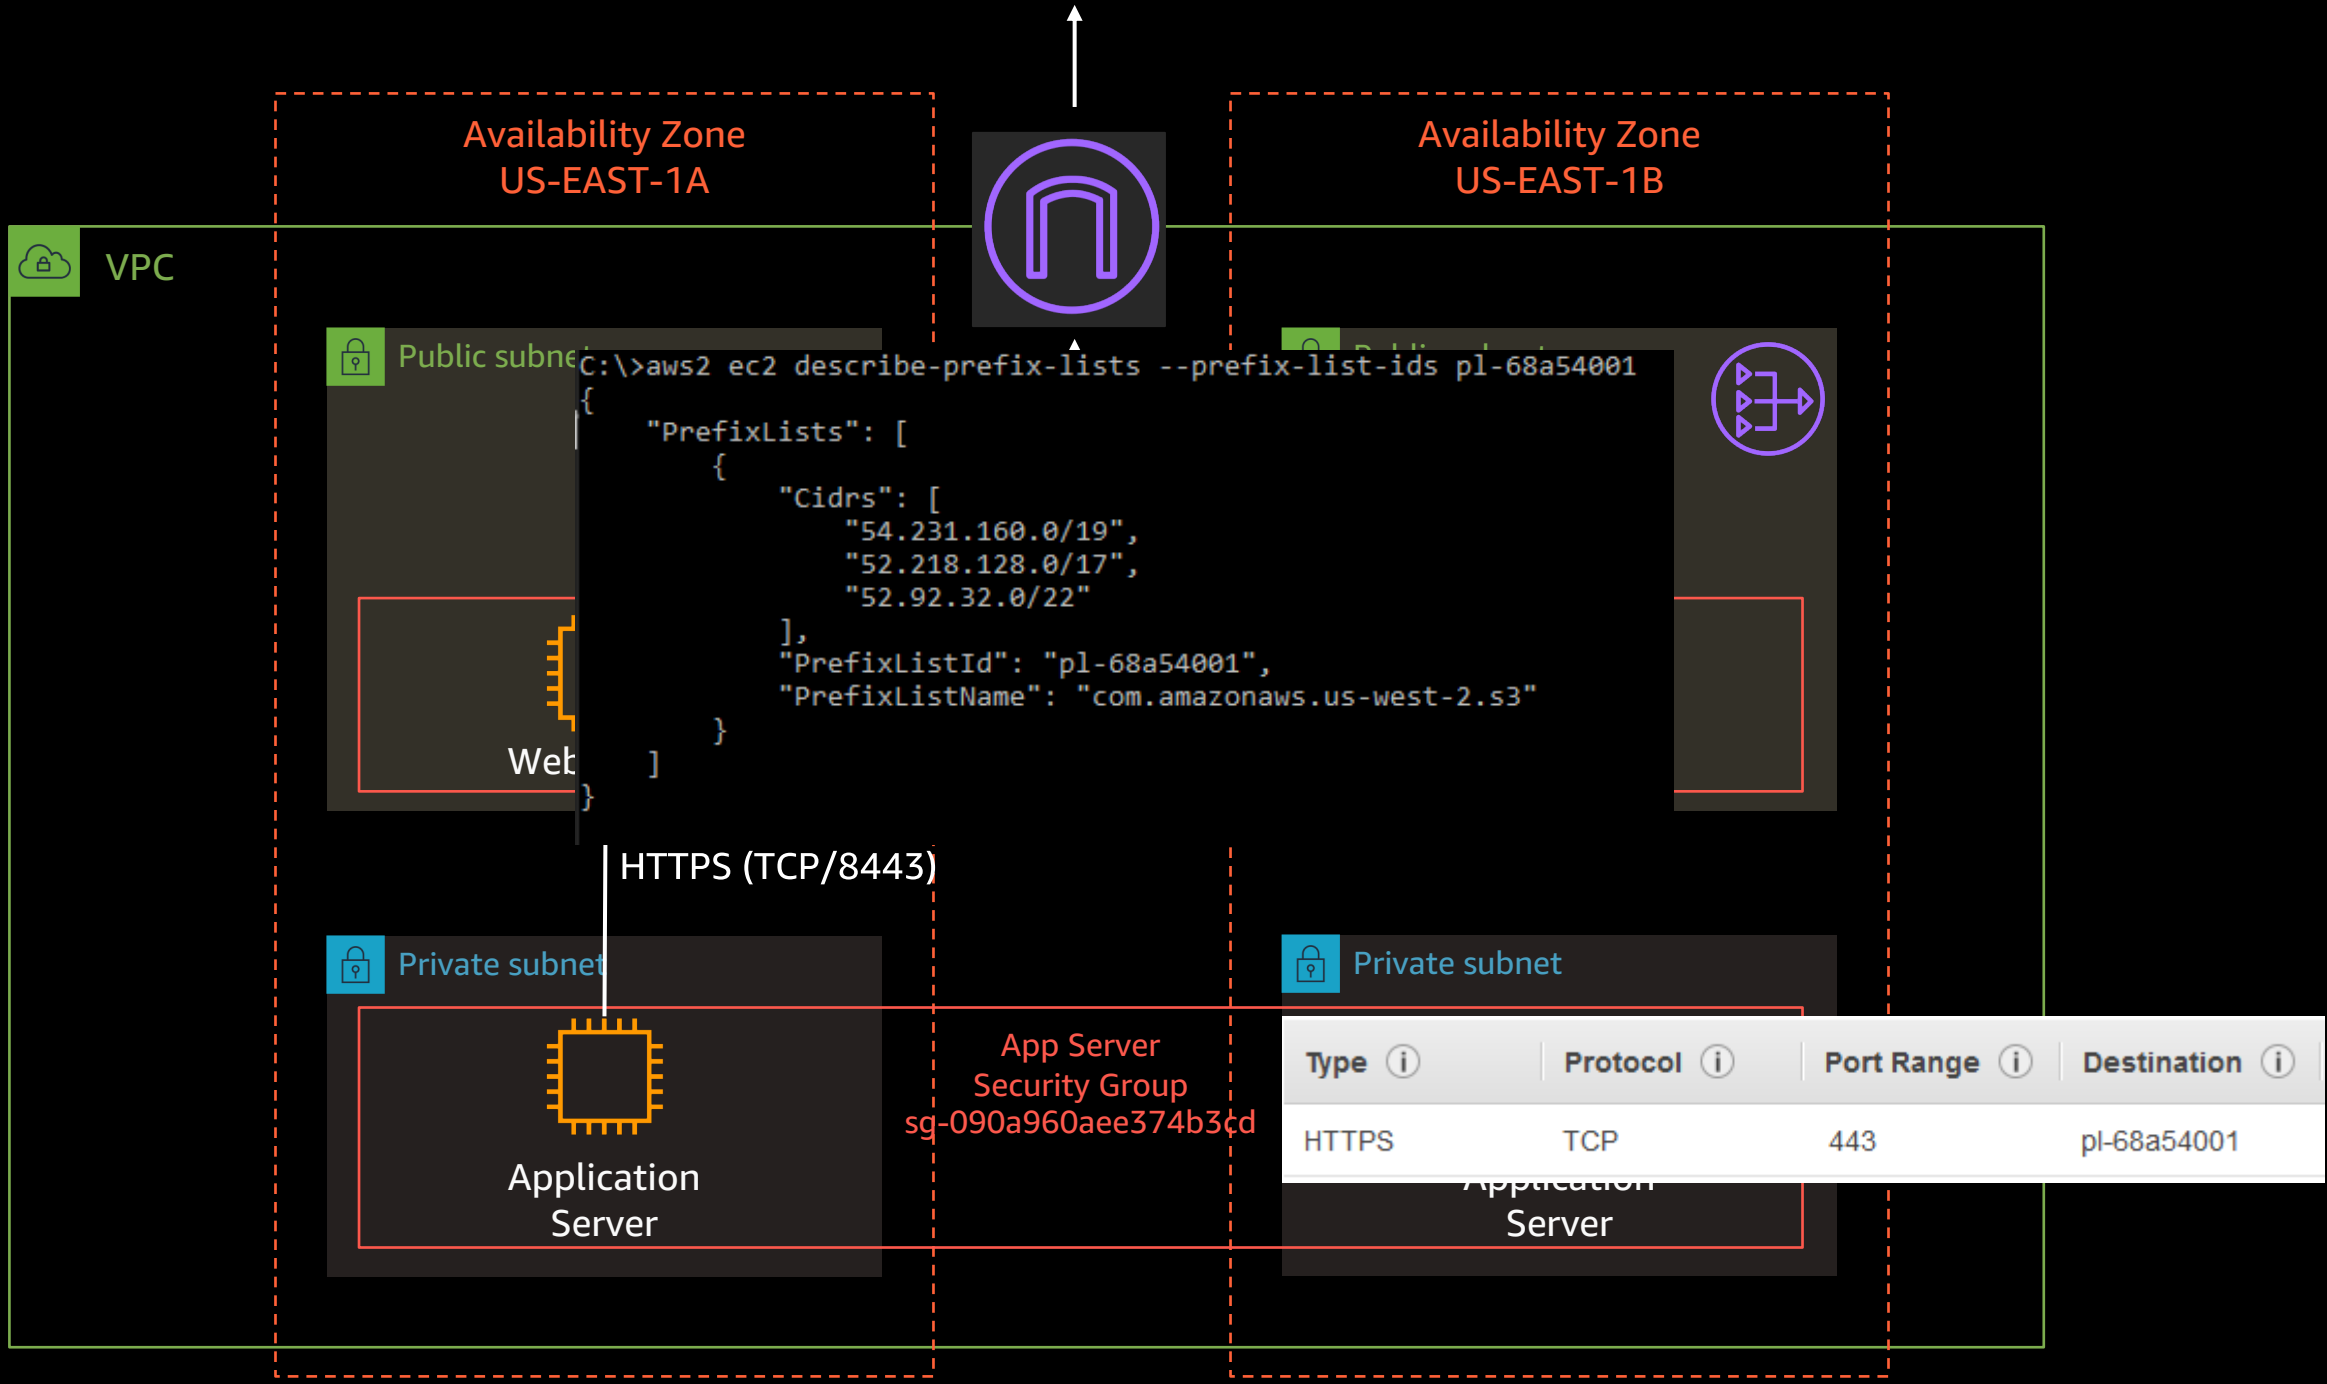




# Security groups – Inbound



# Security groups – Outbound



# VPC flow logs

- Amazon CloudWatch Logs or Amazon S3
- Does not impact throughput or latency
- Apply to VPC, Subnet, or ENI
- Accepted, Rejected, or All traffic

version	3
account-id	384767312345
interface-id	eni-0b62d5e000e412345
srcaddr	108.56.192.231
dstaddr	172.31.0.202
srcport	50565
dstport	80
protocol	6
packets	7
bytes	751
start	1573704396
end	1573704455
action	ACCEPT
log-status	OK
vpc-id	vpc-0af48868ceeb12345
subnet-id	subnet-02ab634d2e4c12345
instance-id	i-0a998a68301112345
tcp-flags	3
type	IPv4
pkt-srcaddr	108.56.192.231
pkt-dstaddr	172.31.0.202

# Amazon VPC traffic mirroring

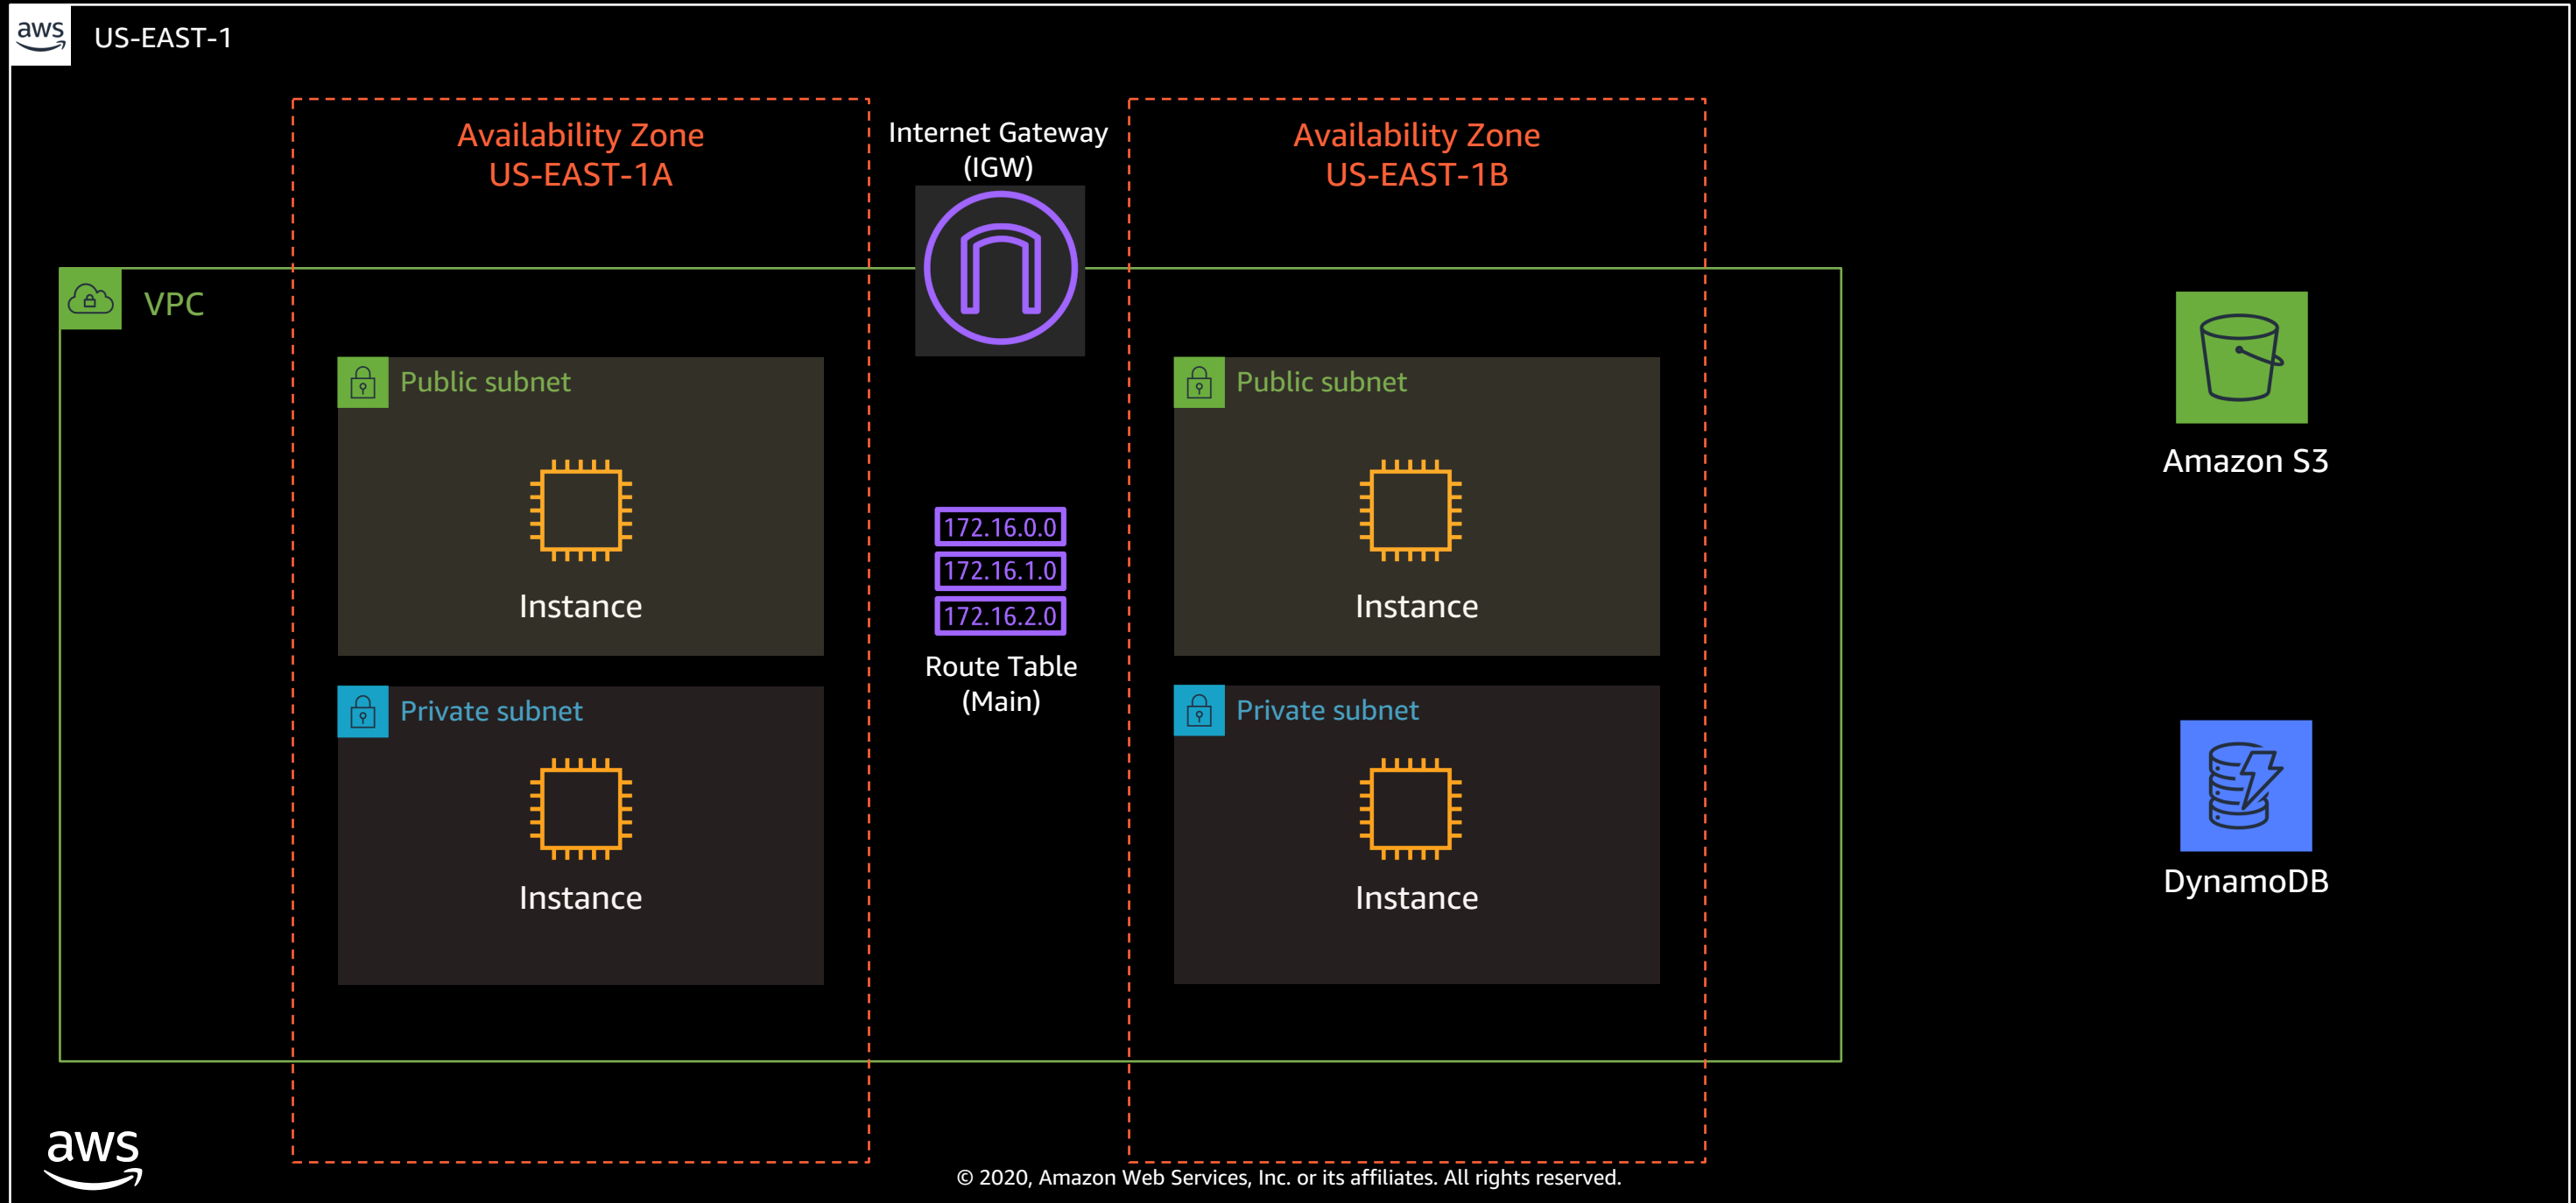
- Mirror to another ENI or Network Load Balancer with UDP listener
- Packet copy. Shares interface bandwidth.
- Traffic mirror filters to define “interesting traffic”
- Traffic mirror session is the combination of source, target, and filter



# VPC endpoints

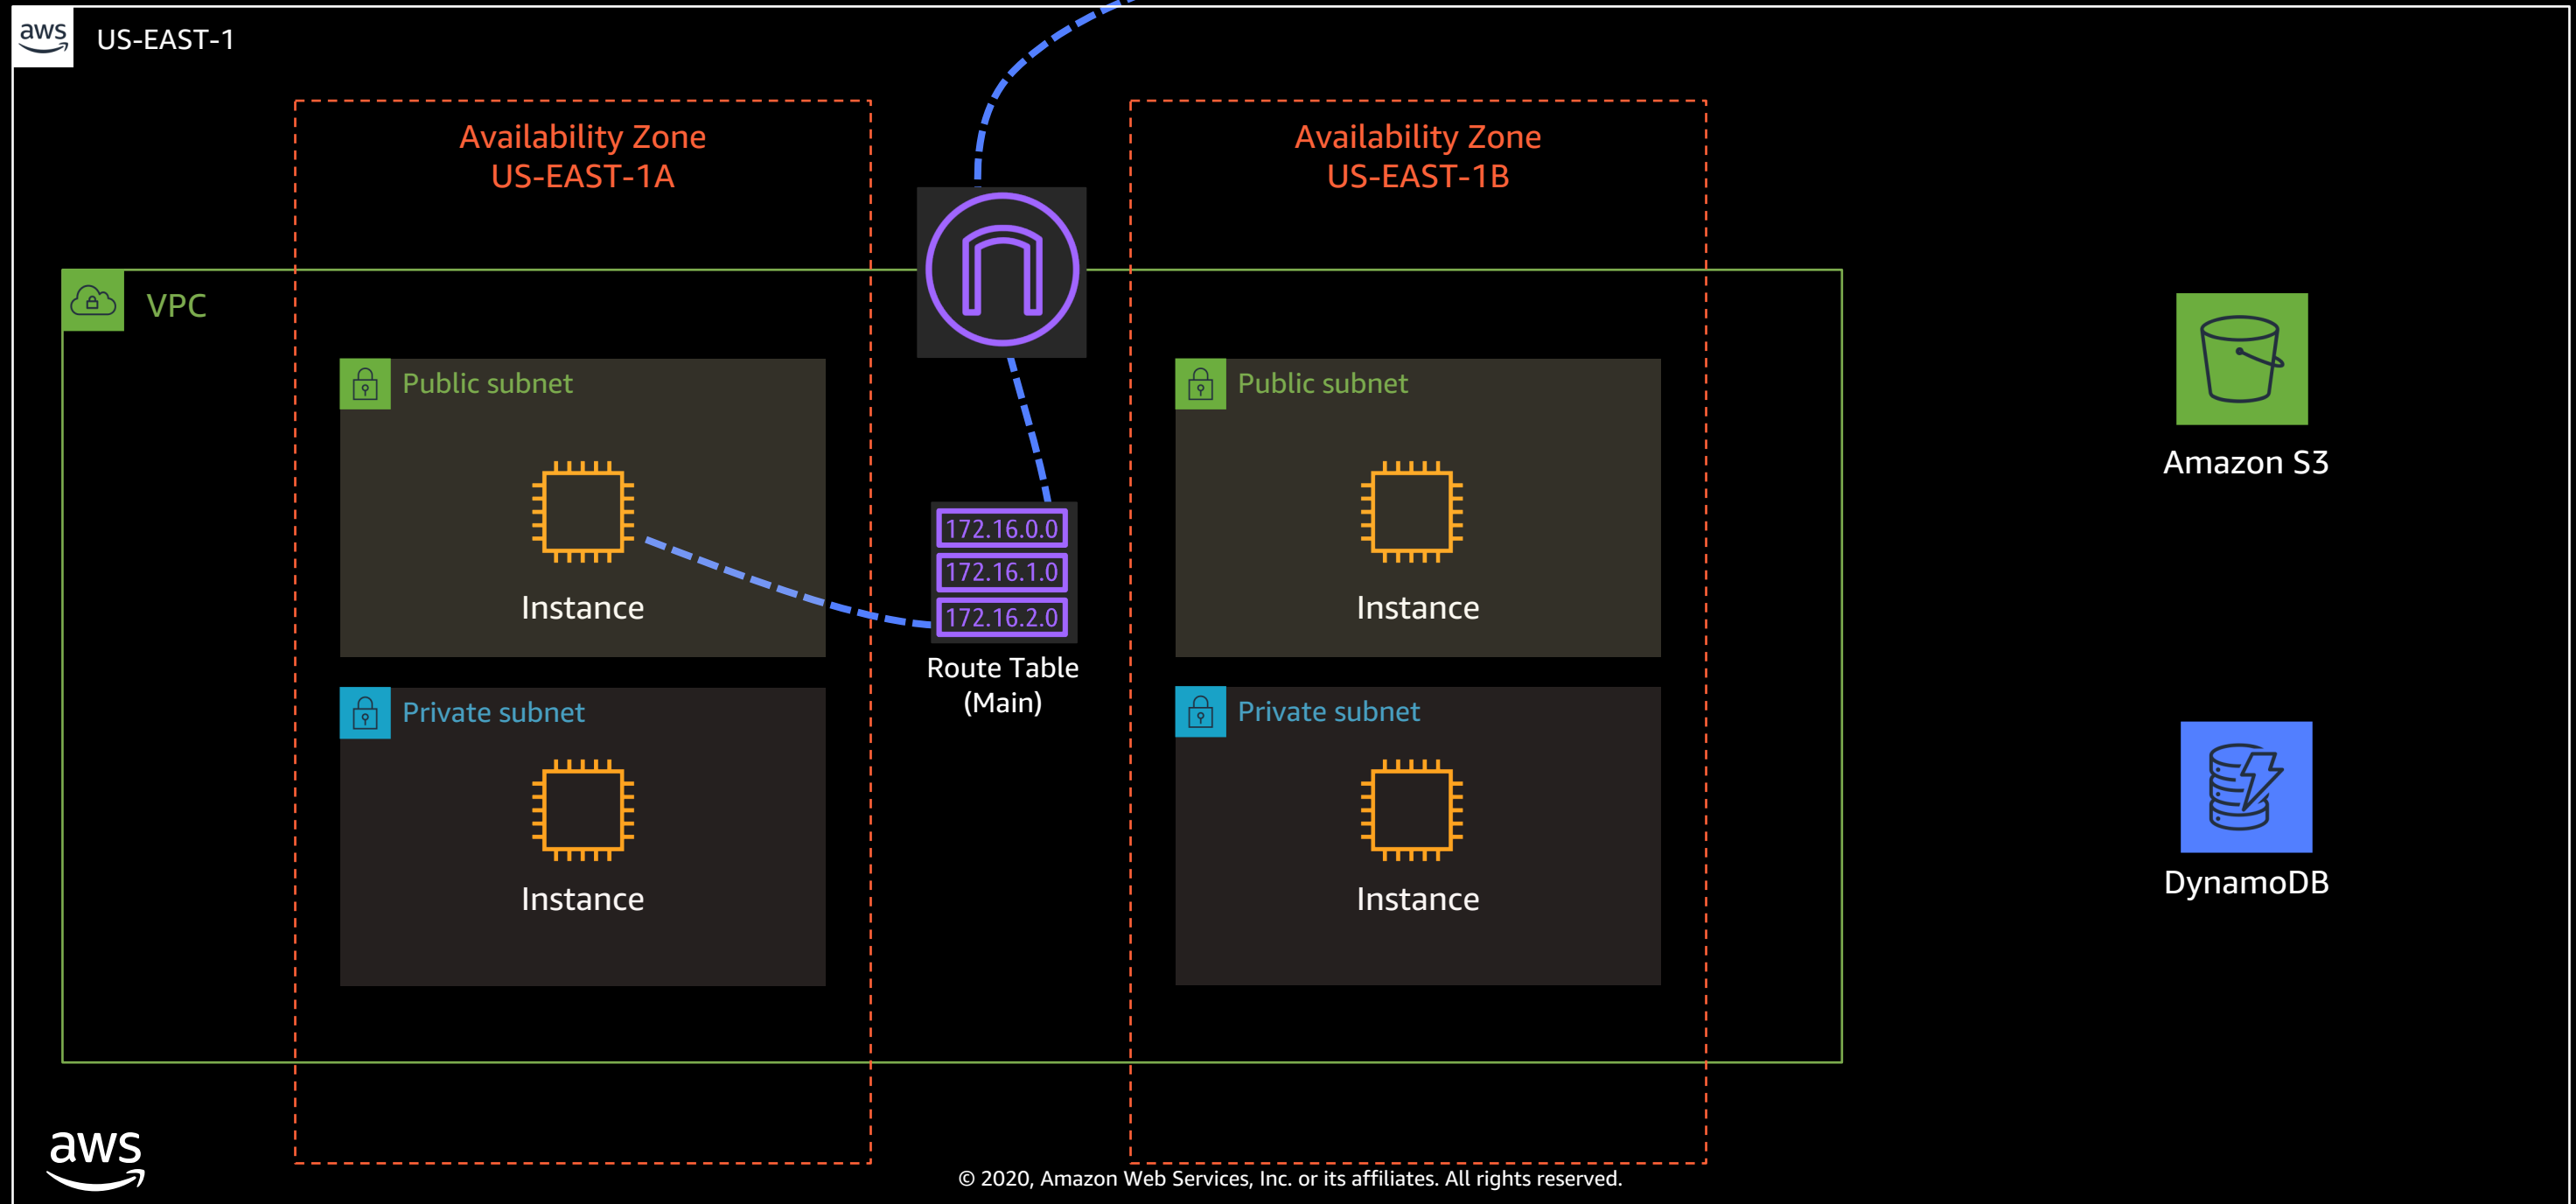


# Gateway VPC endpoints



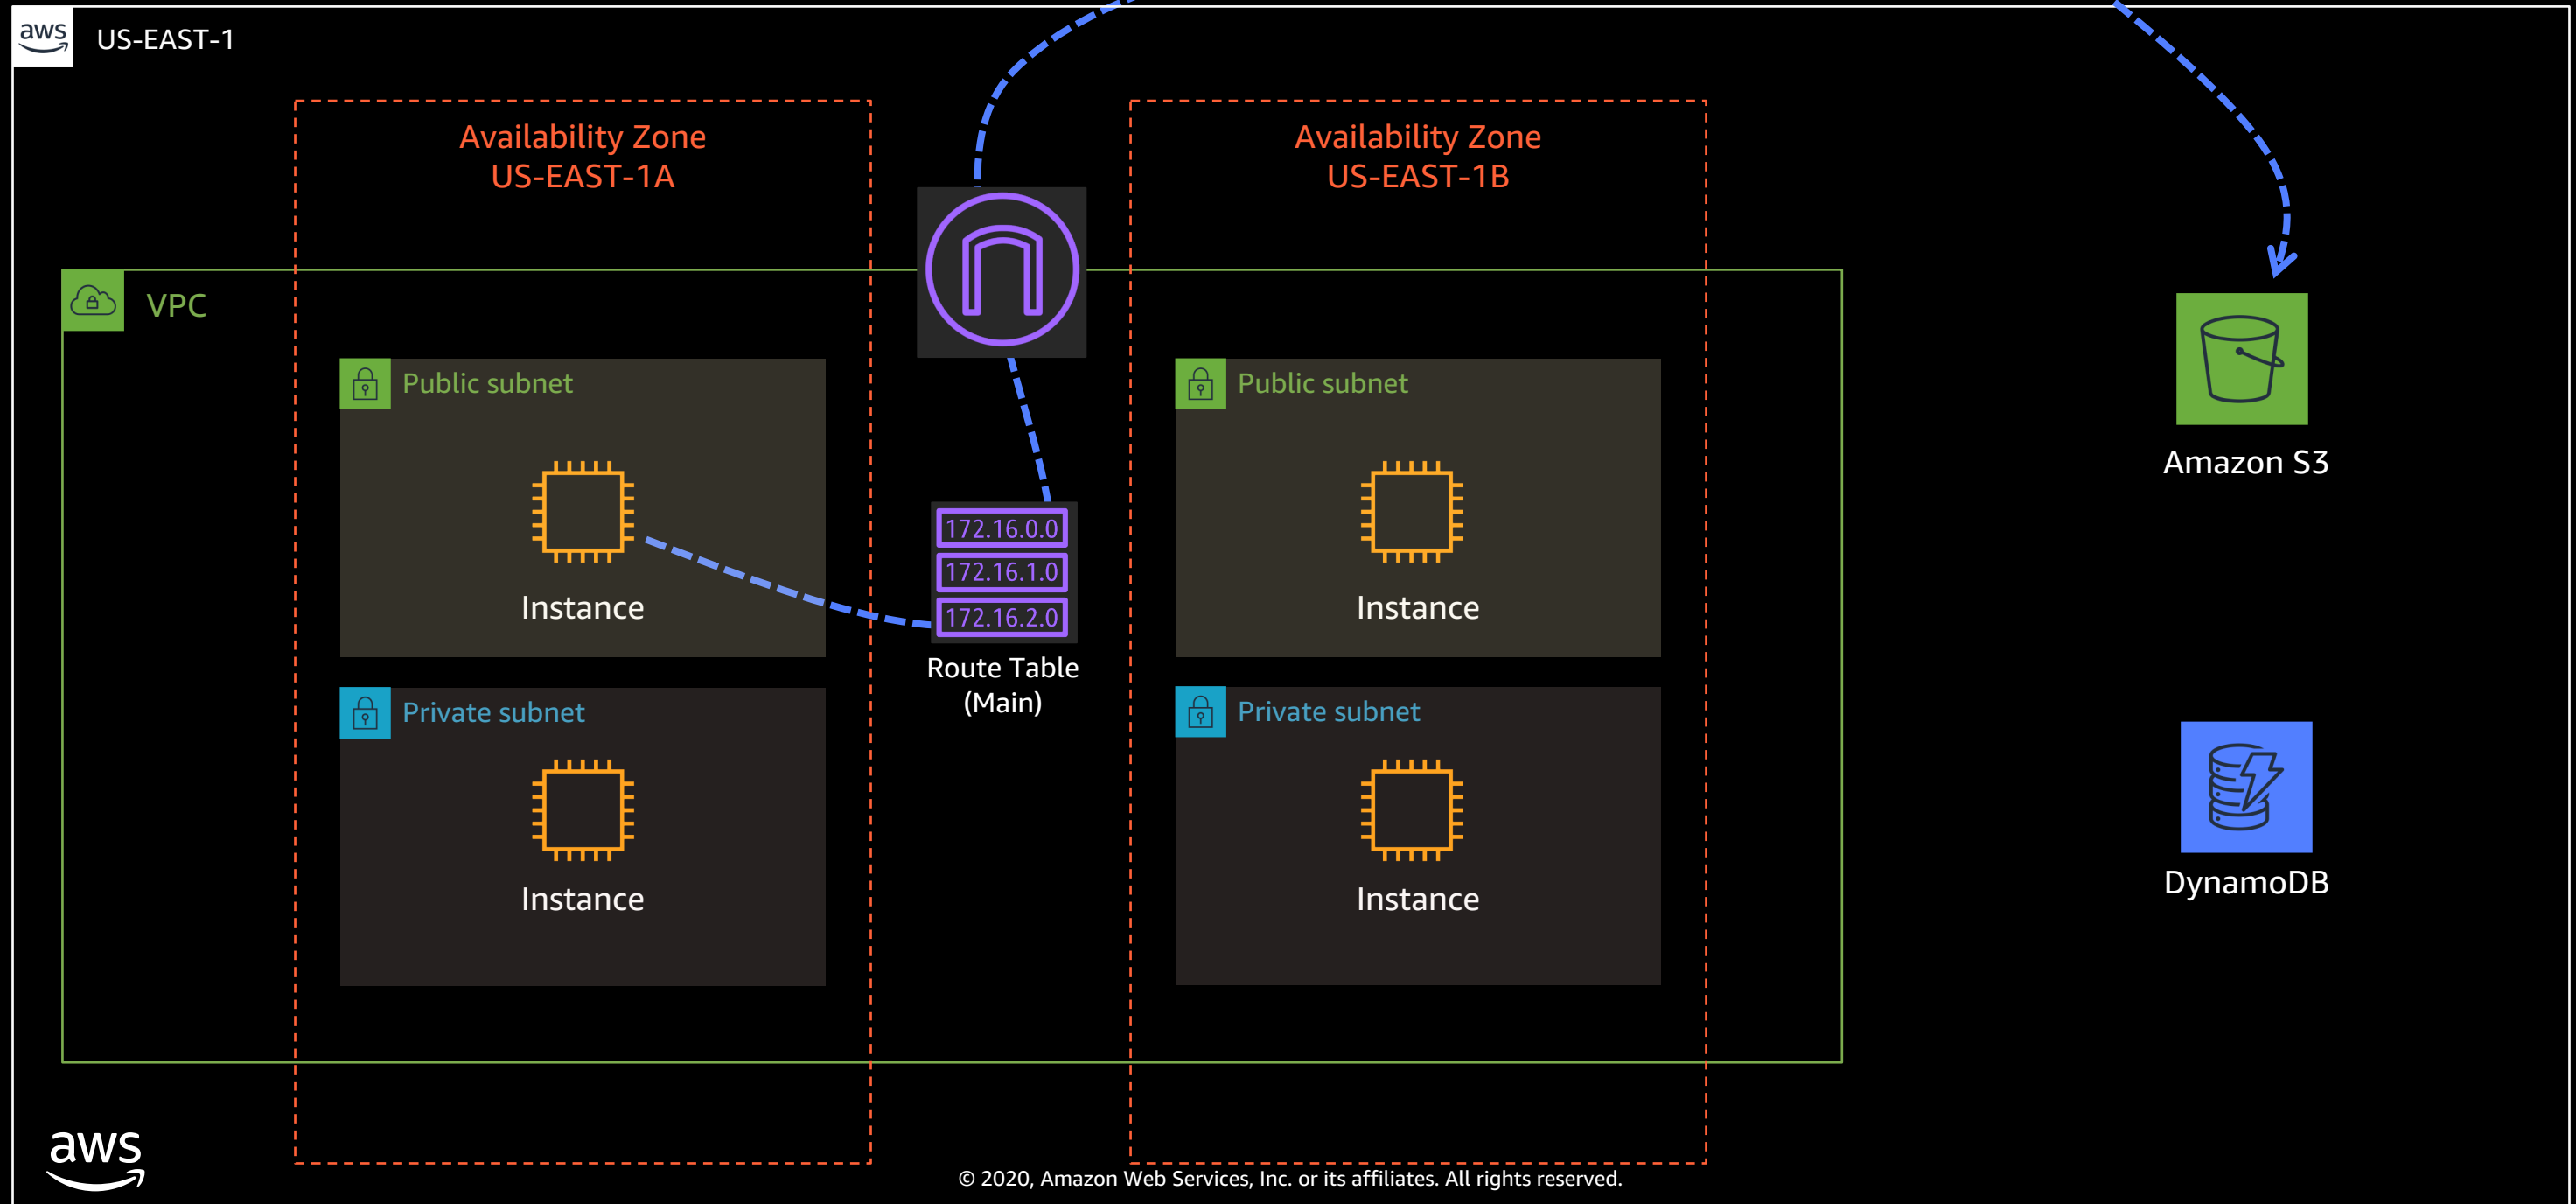
# Gateway VPC endpoints

s3.us-east-1.amazonaws.com  
52.216.229.141 ..... etc



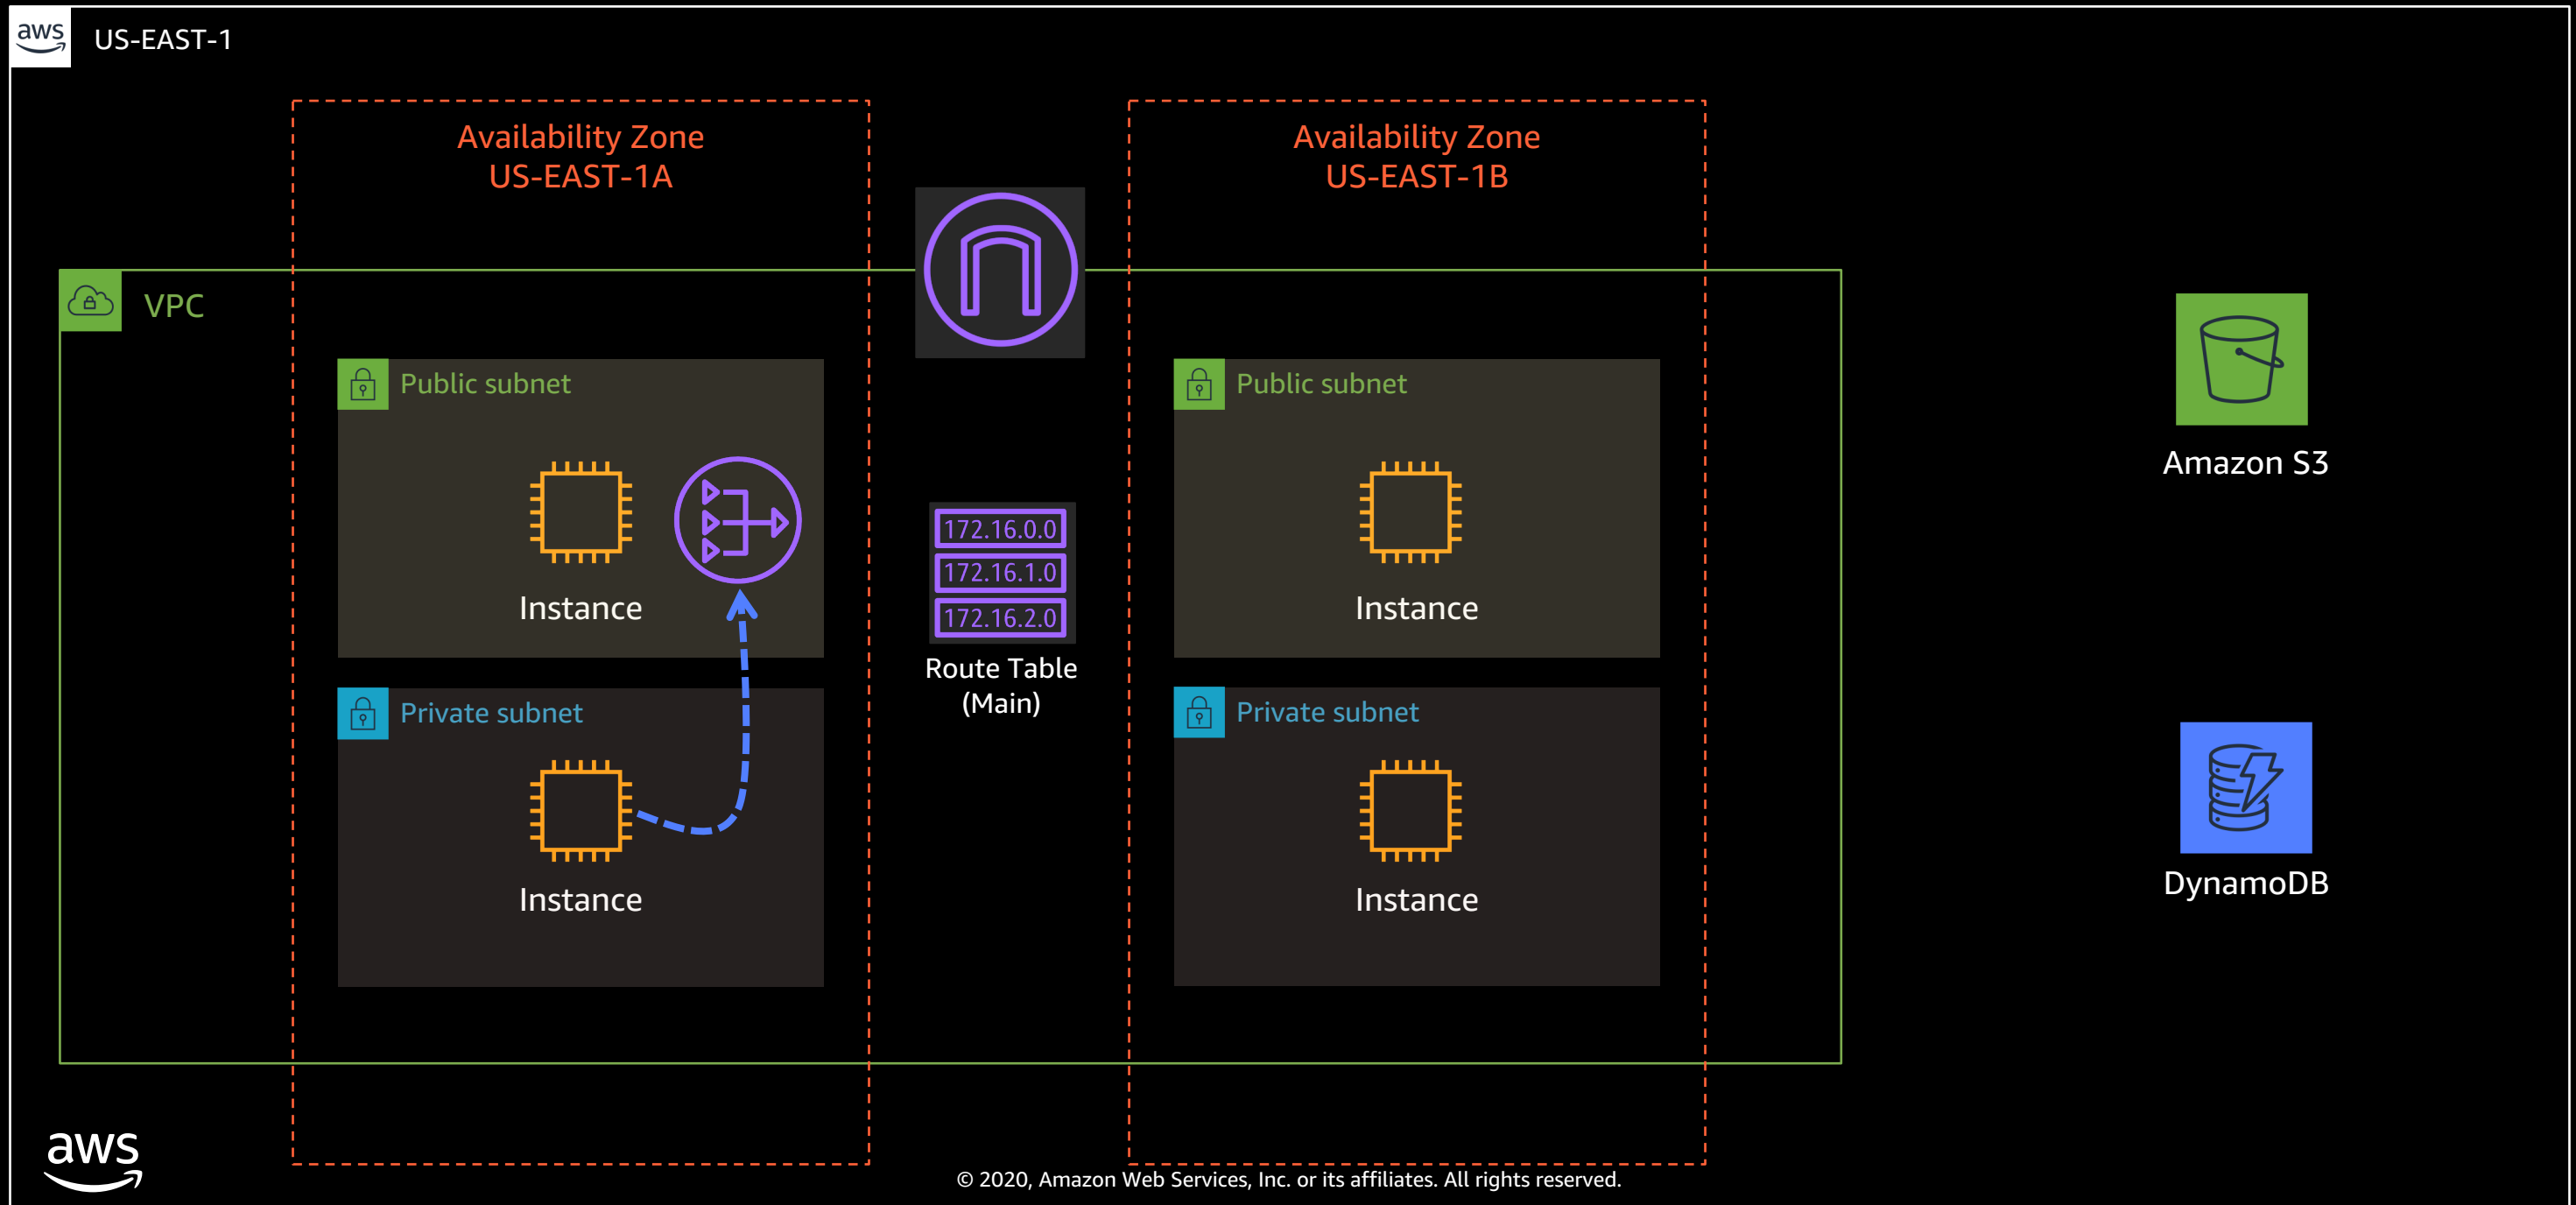
# Gateway VPC endpoints

s3.us-east-1.amazonaws.com  
52.216.229.141 ..... etc



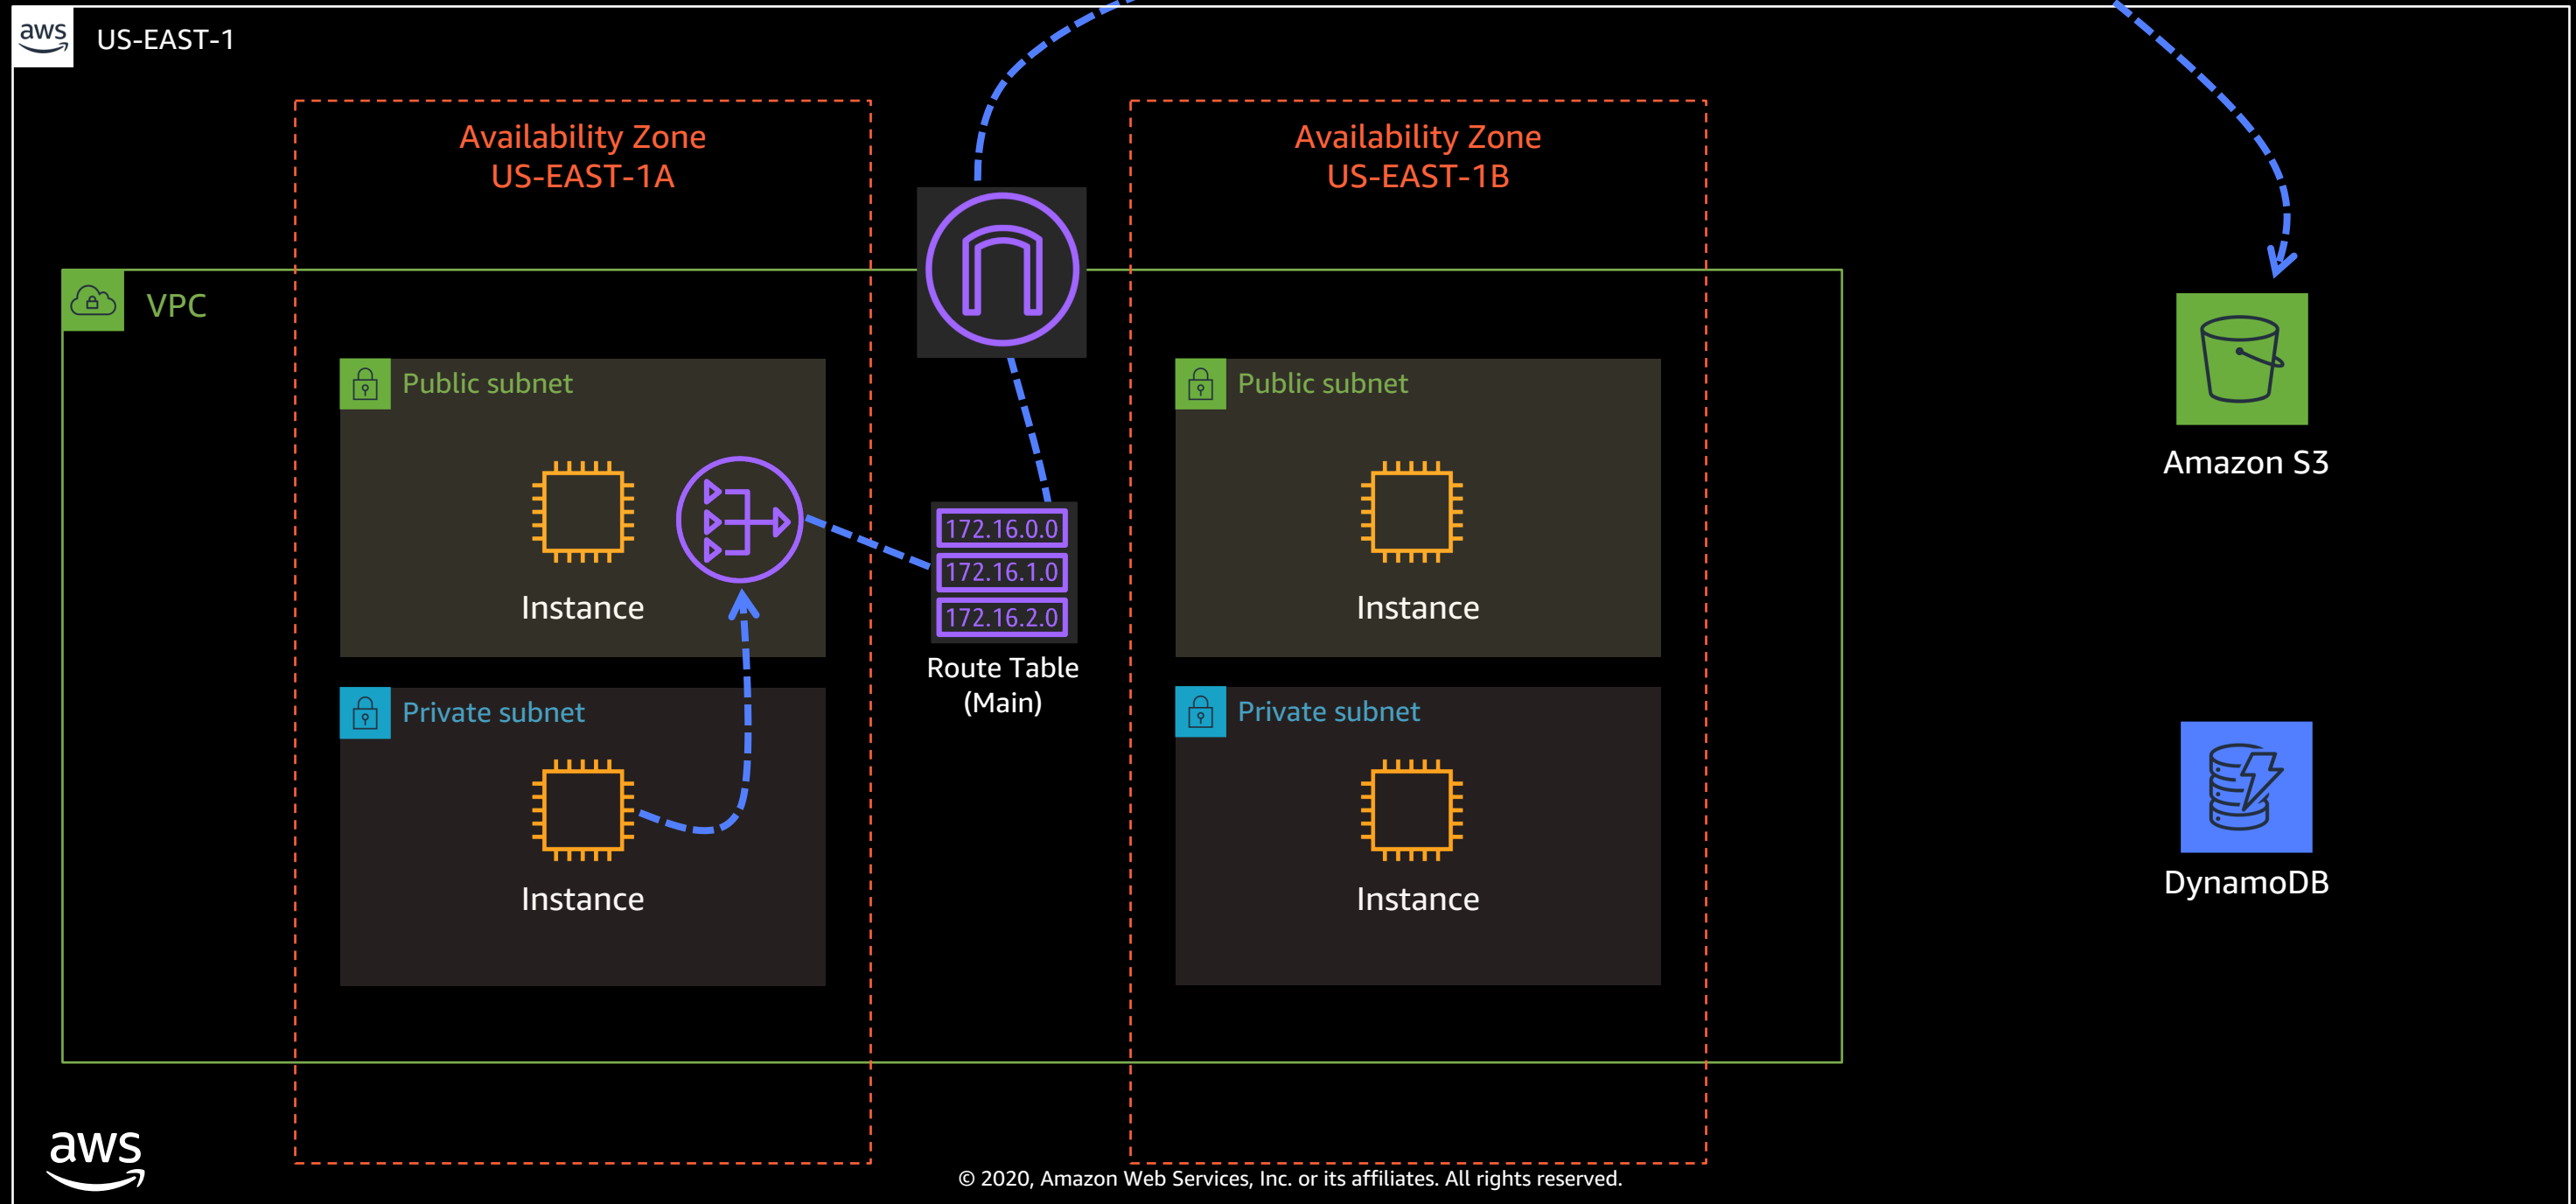


# Gateway VPC endpoints

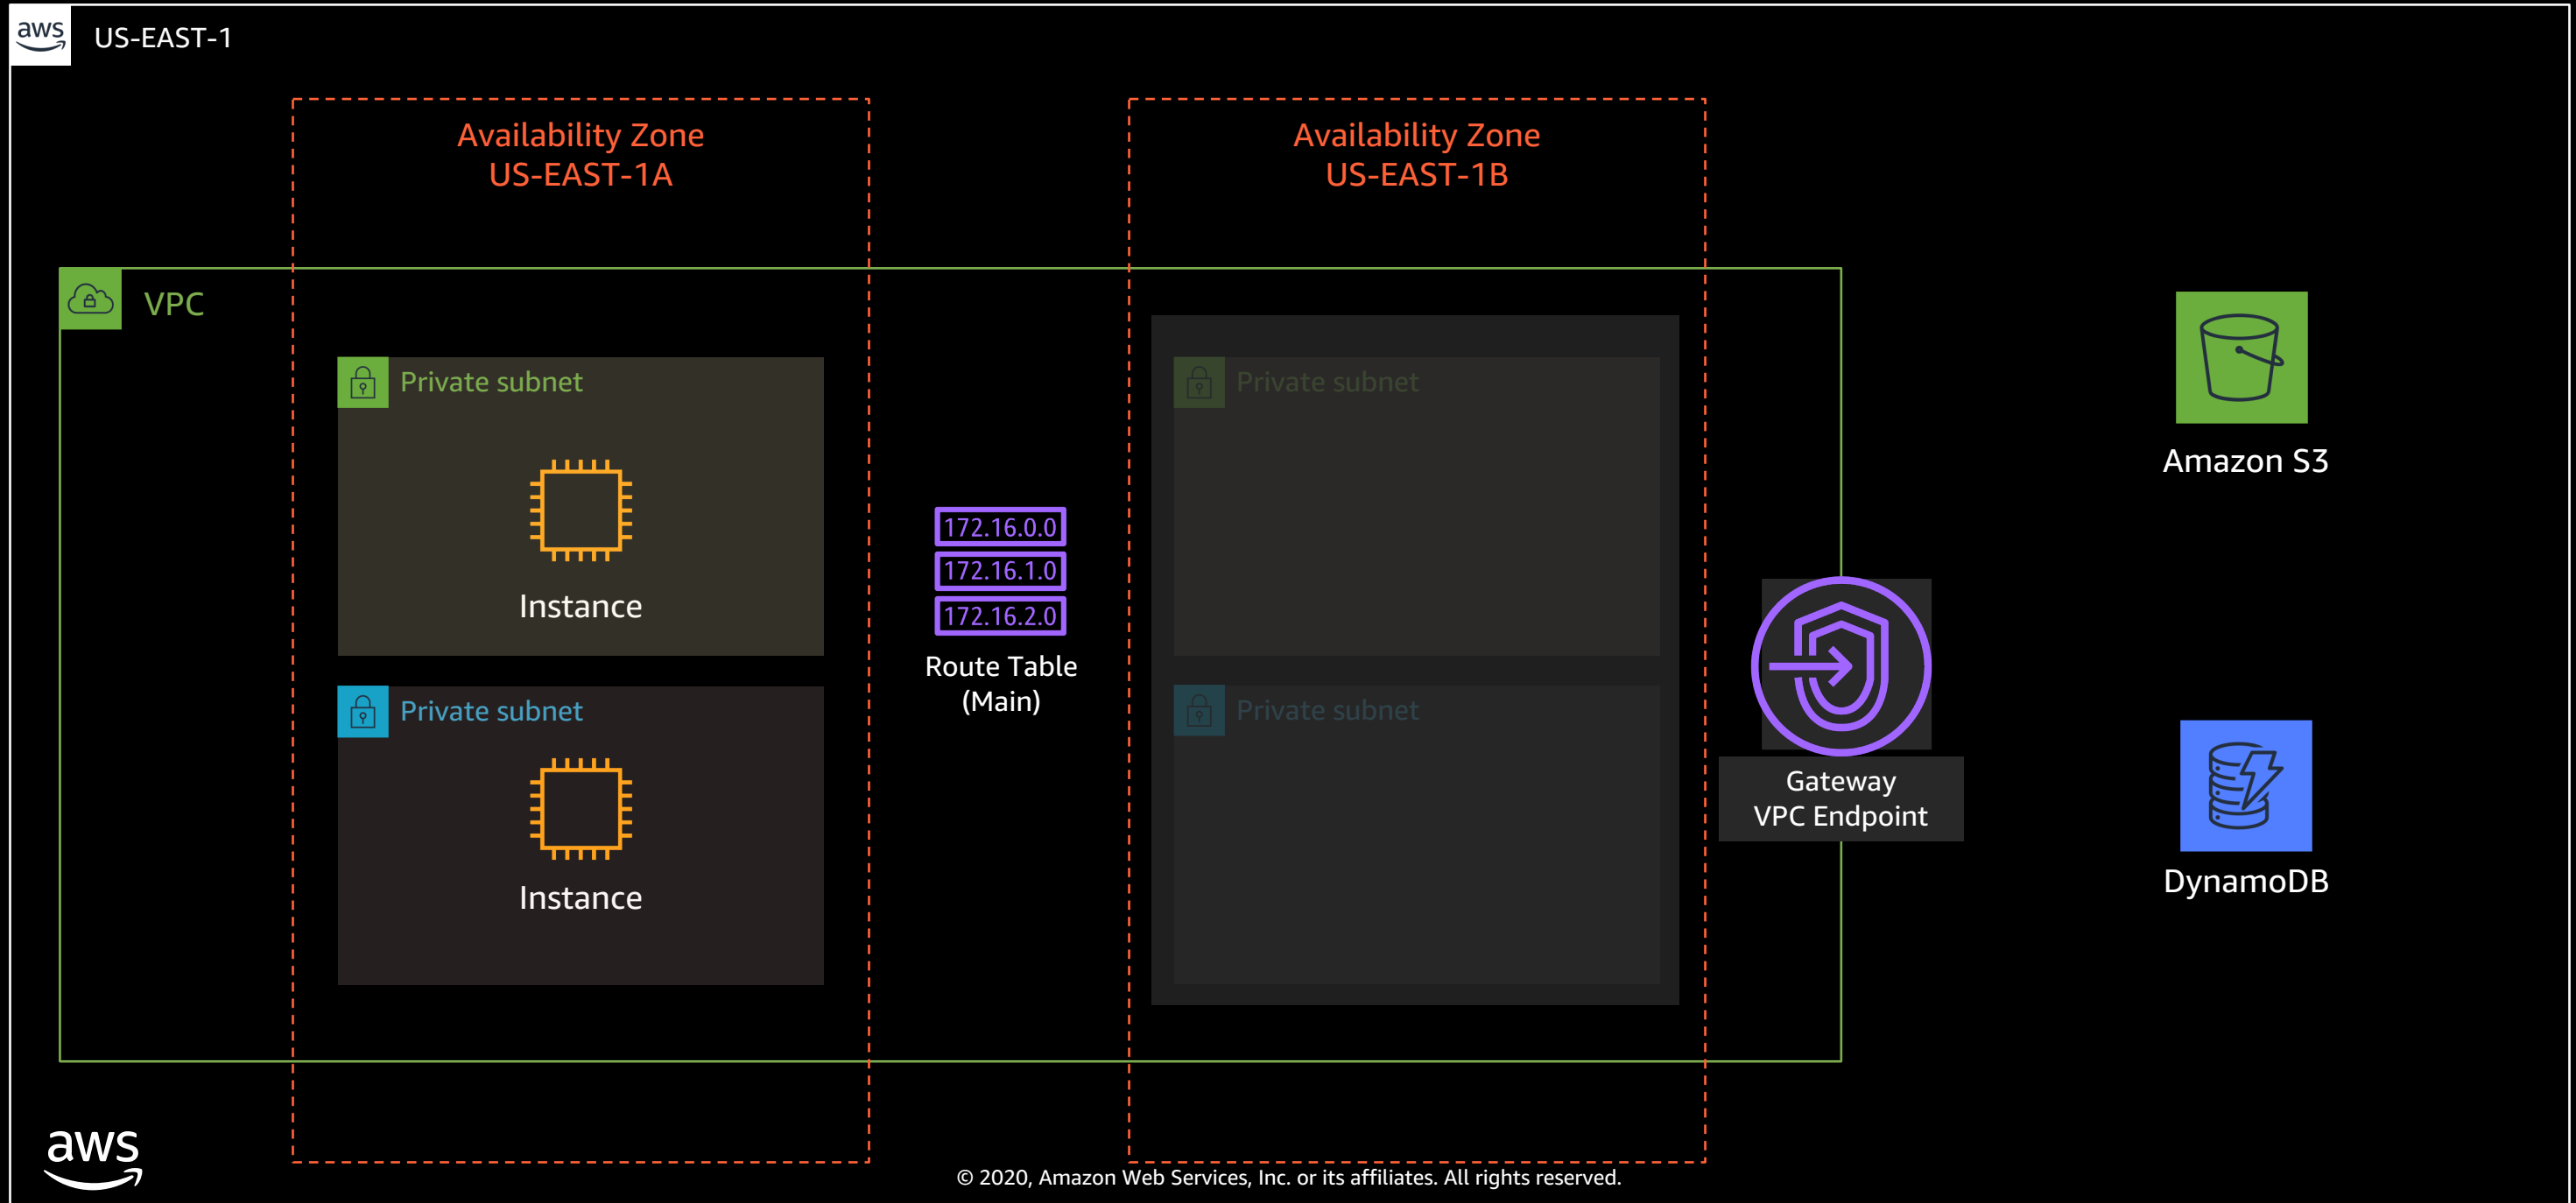


# Gateway VPC endpoints

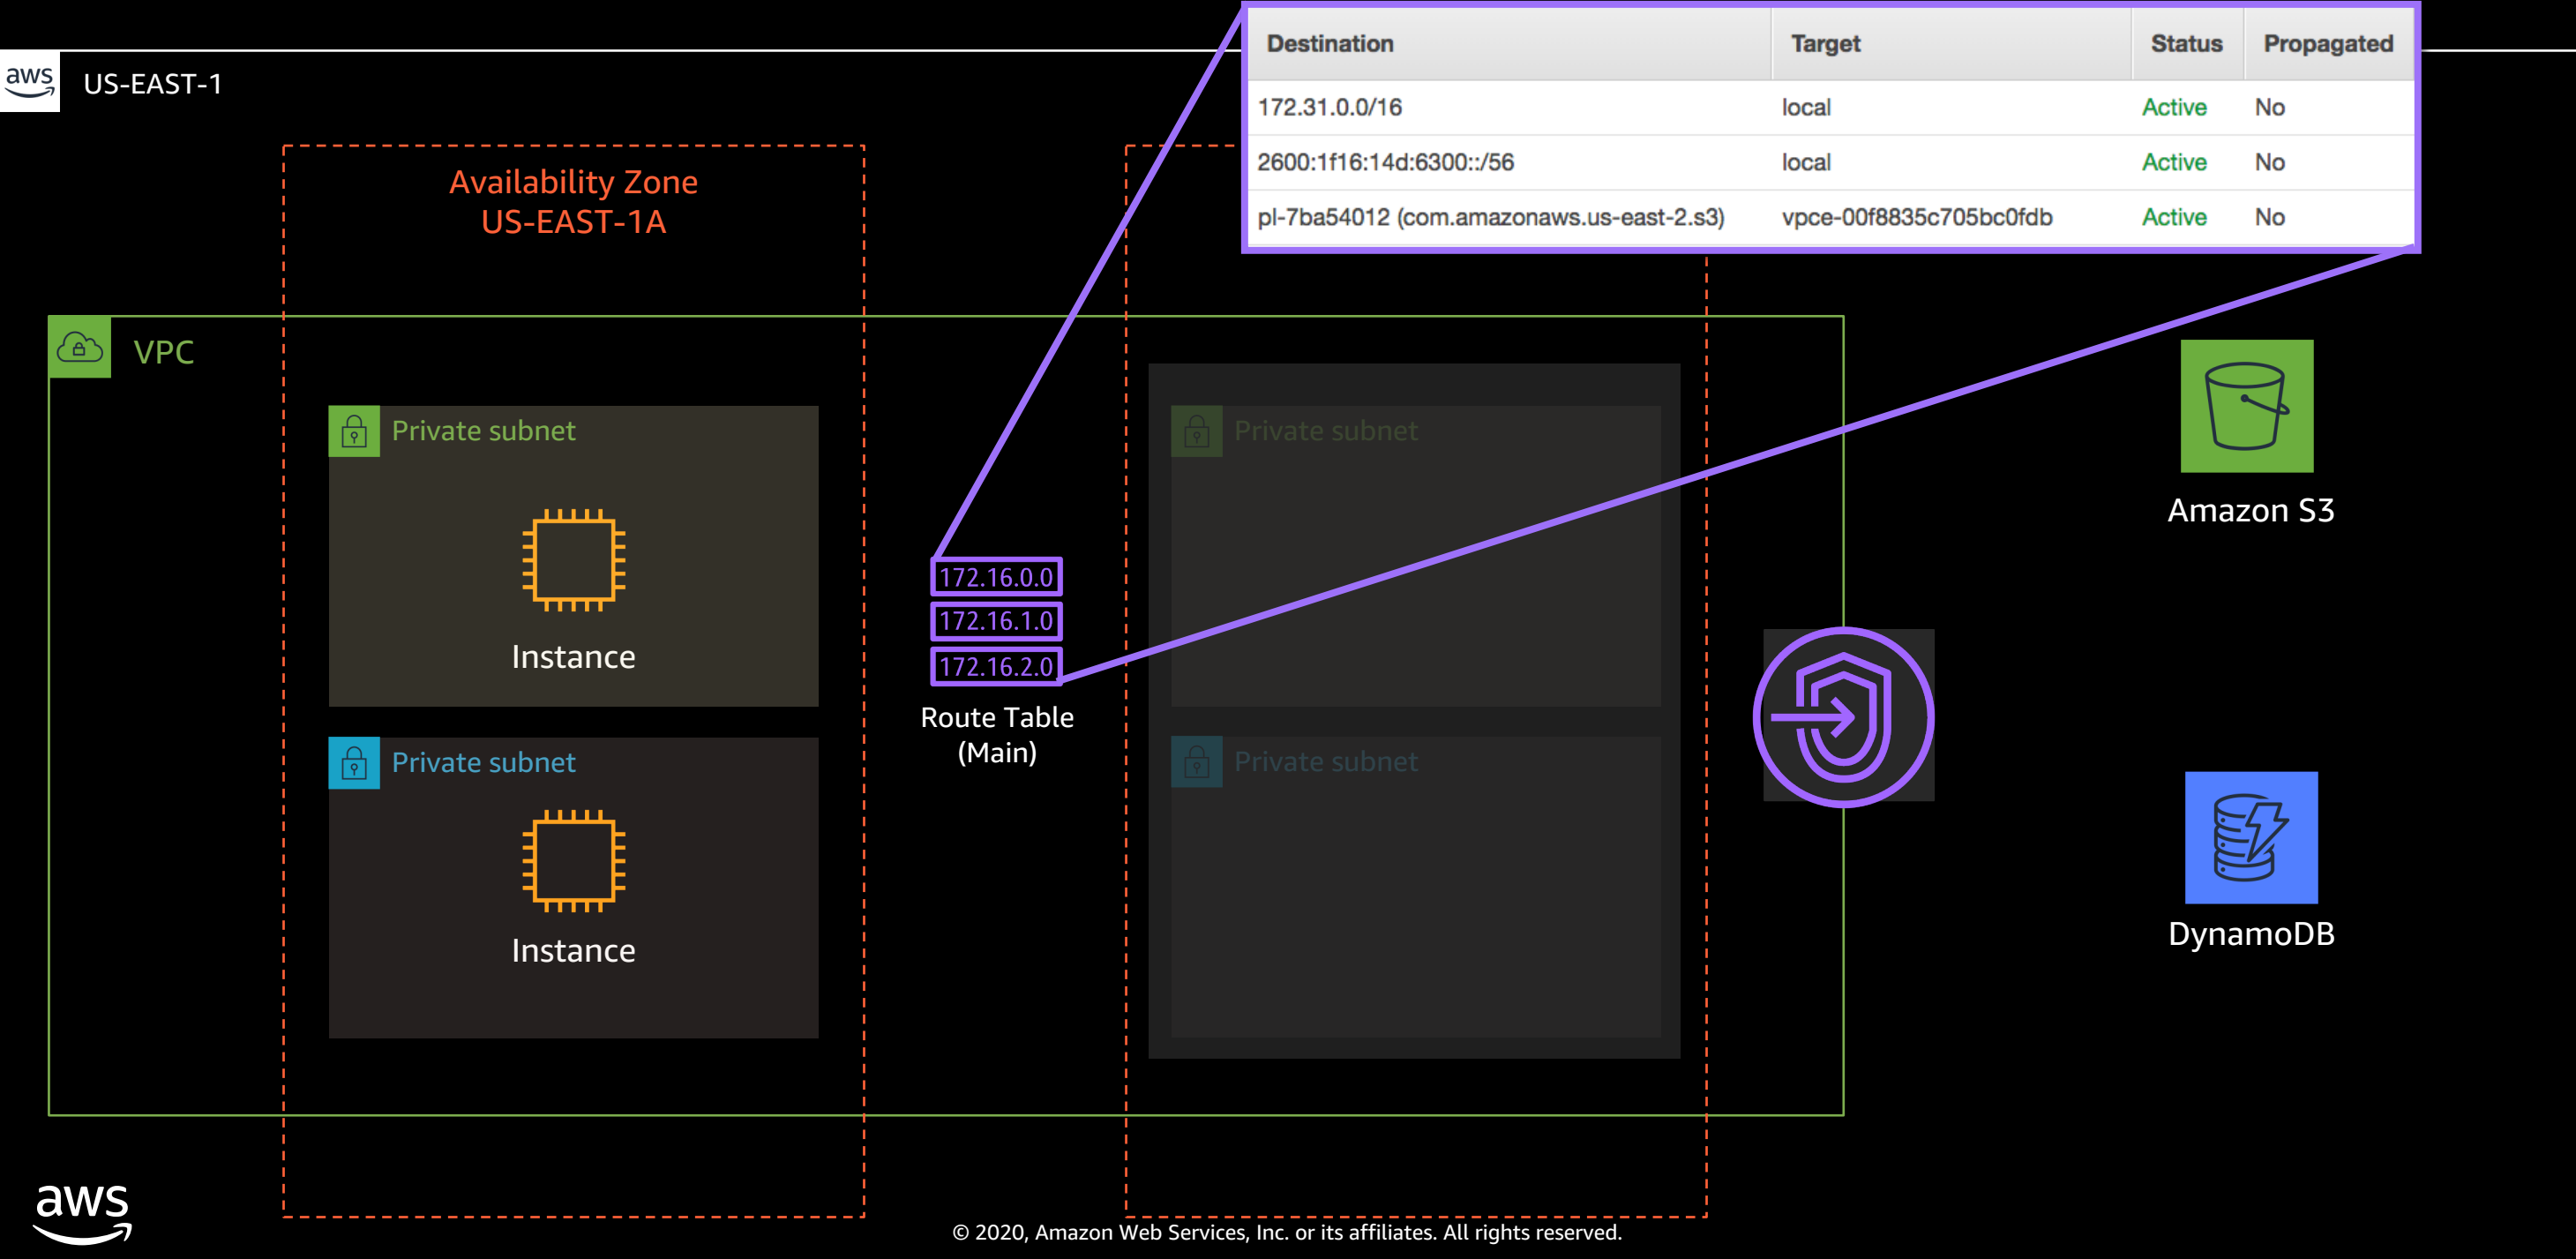
s3.us-east-1.amazonaws.com  
52.216.229.141 ... etc.



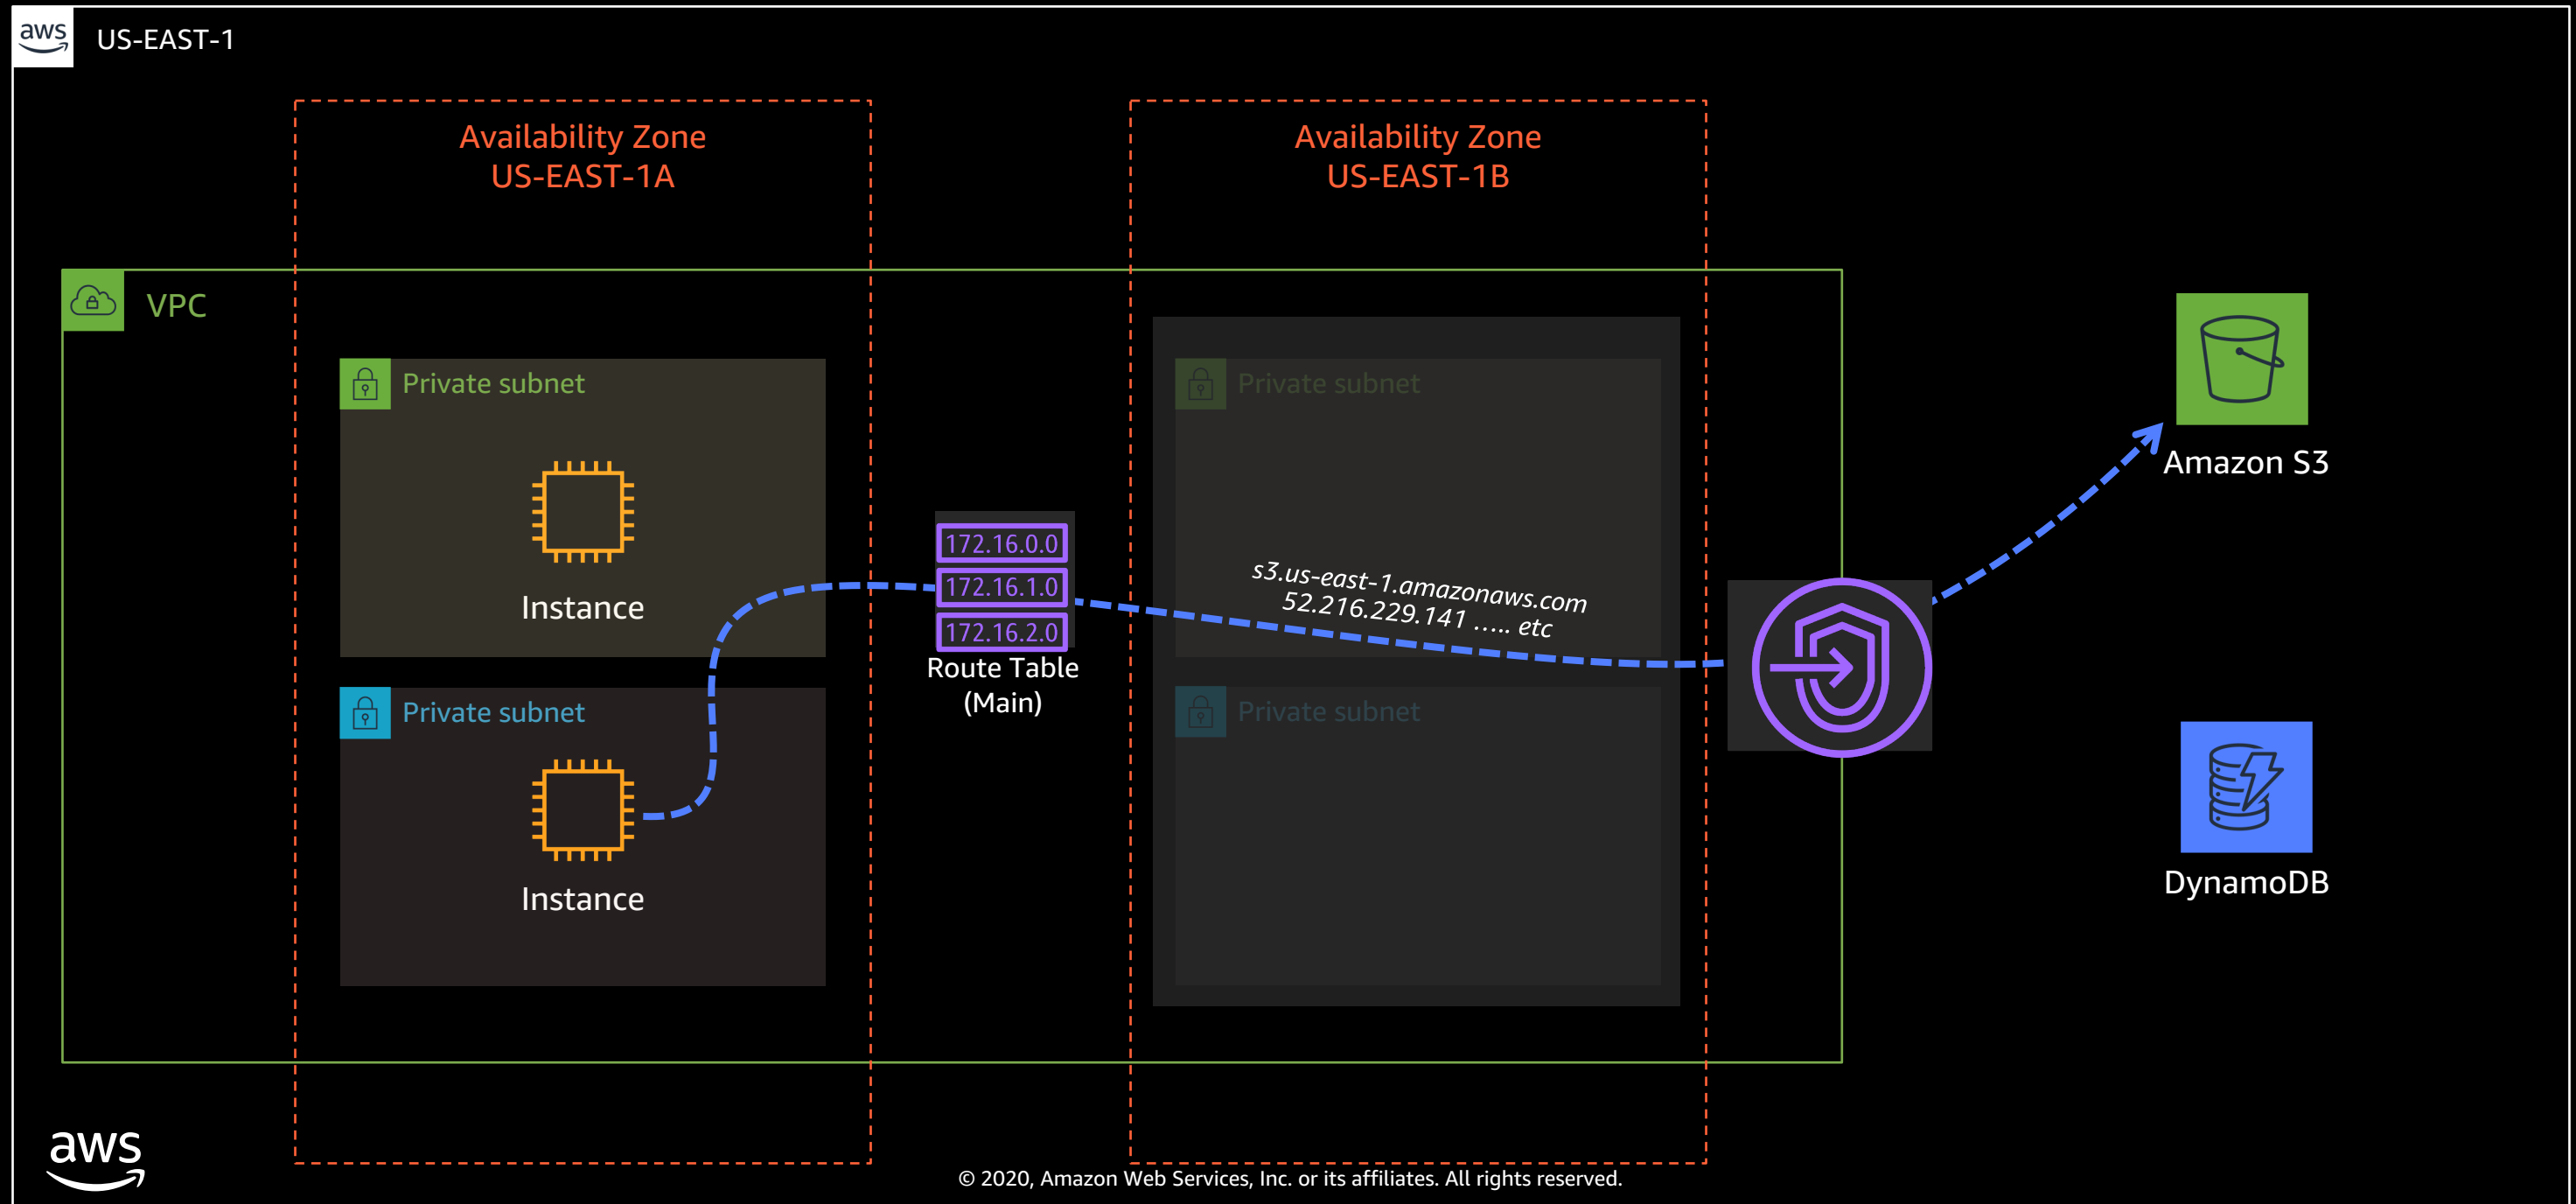
# Gateway VPC endpoints



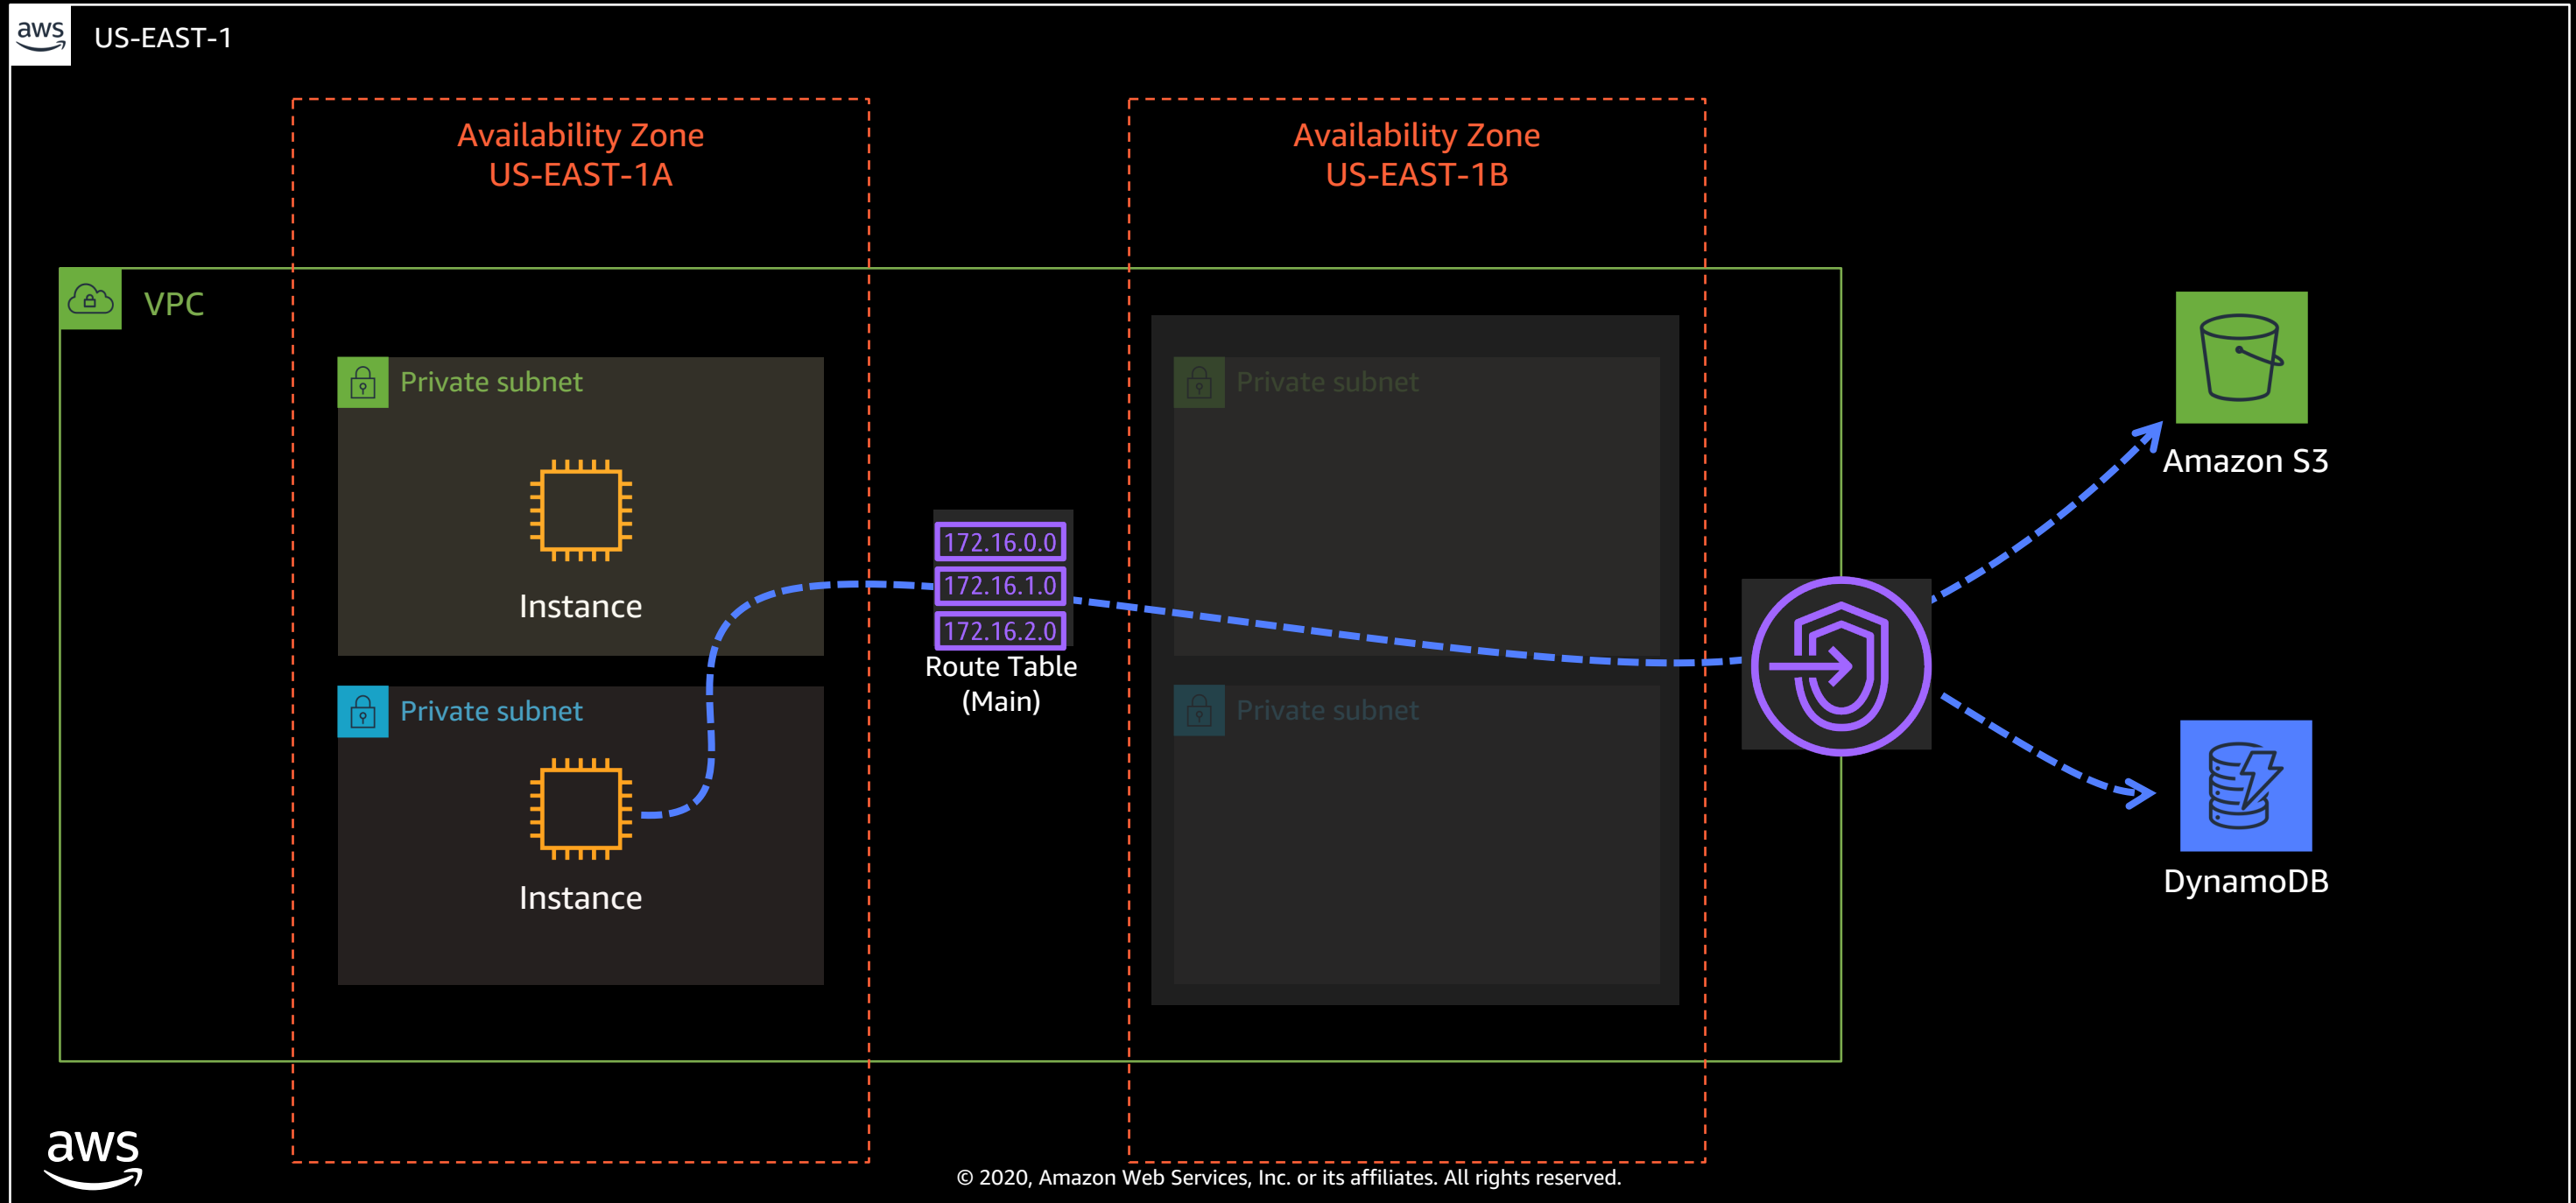
# Gateway VPC endpoints



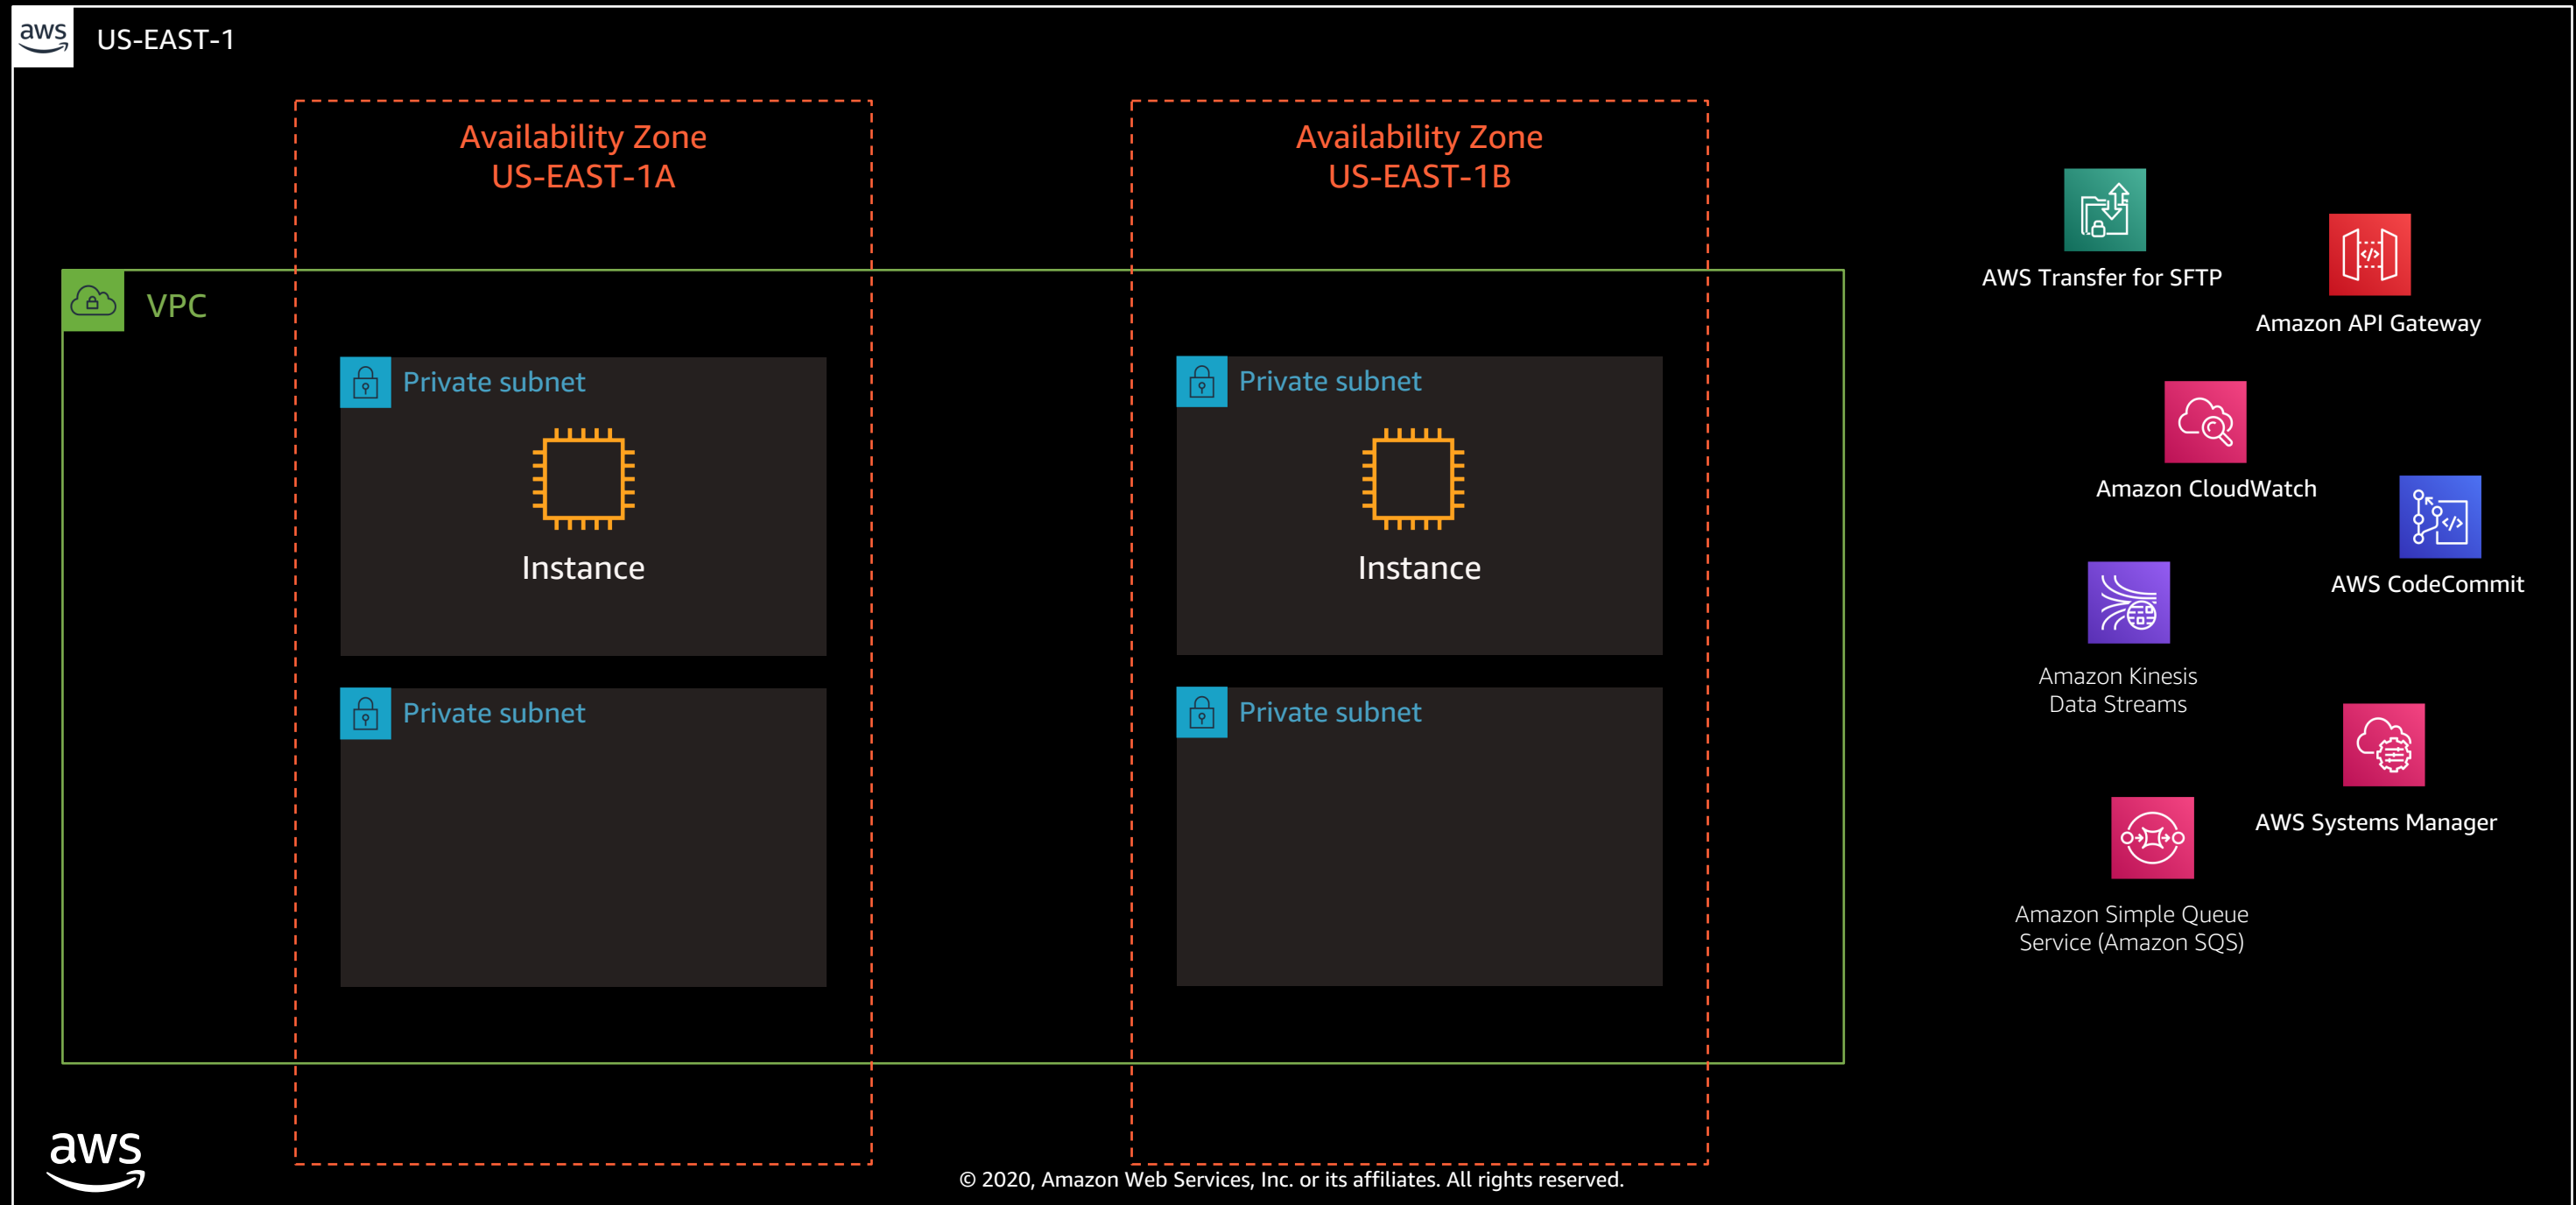
# Gateway VPC endpoints



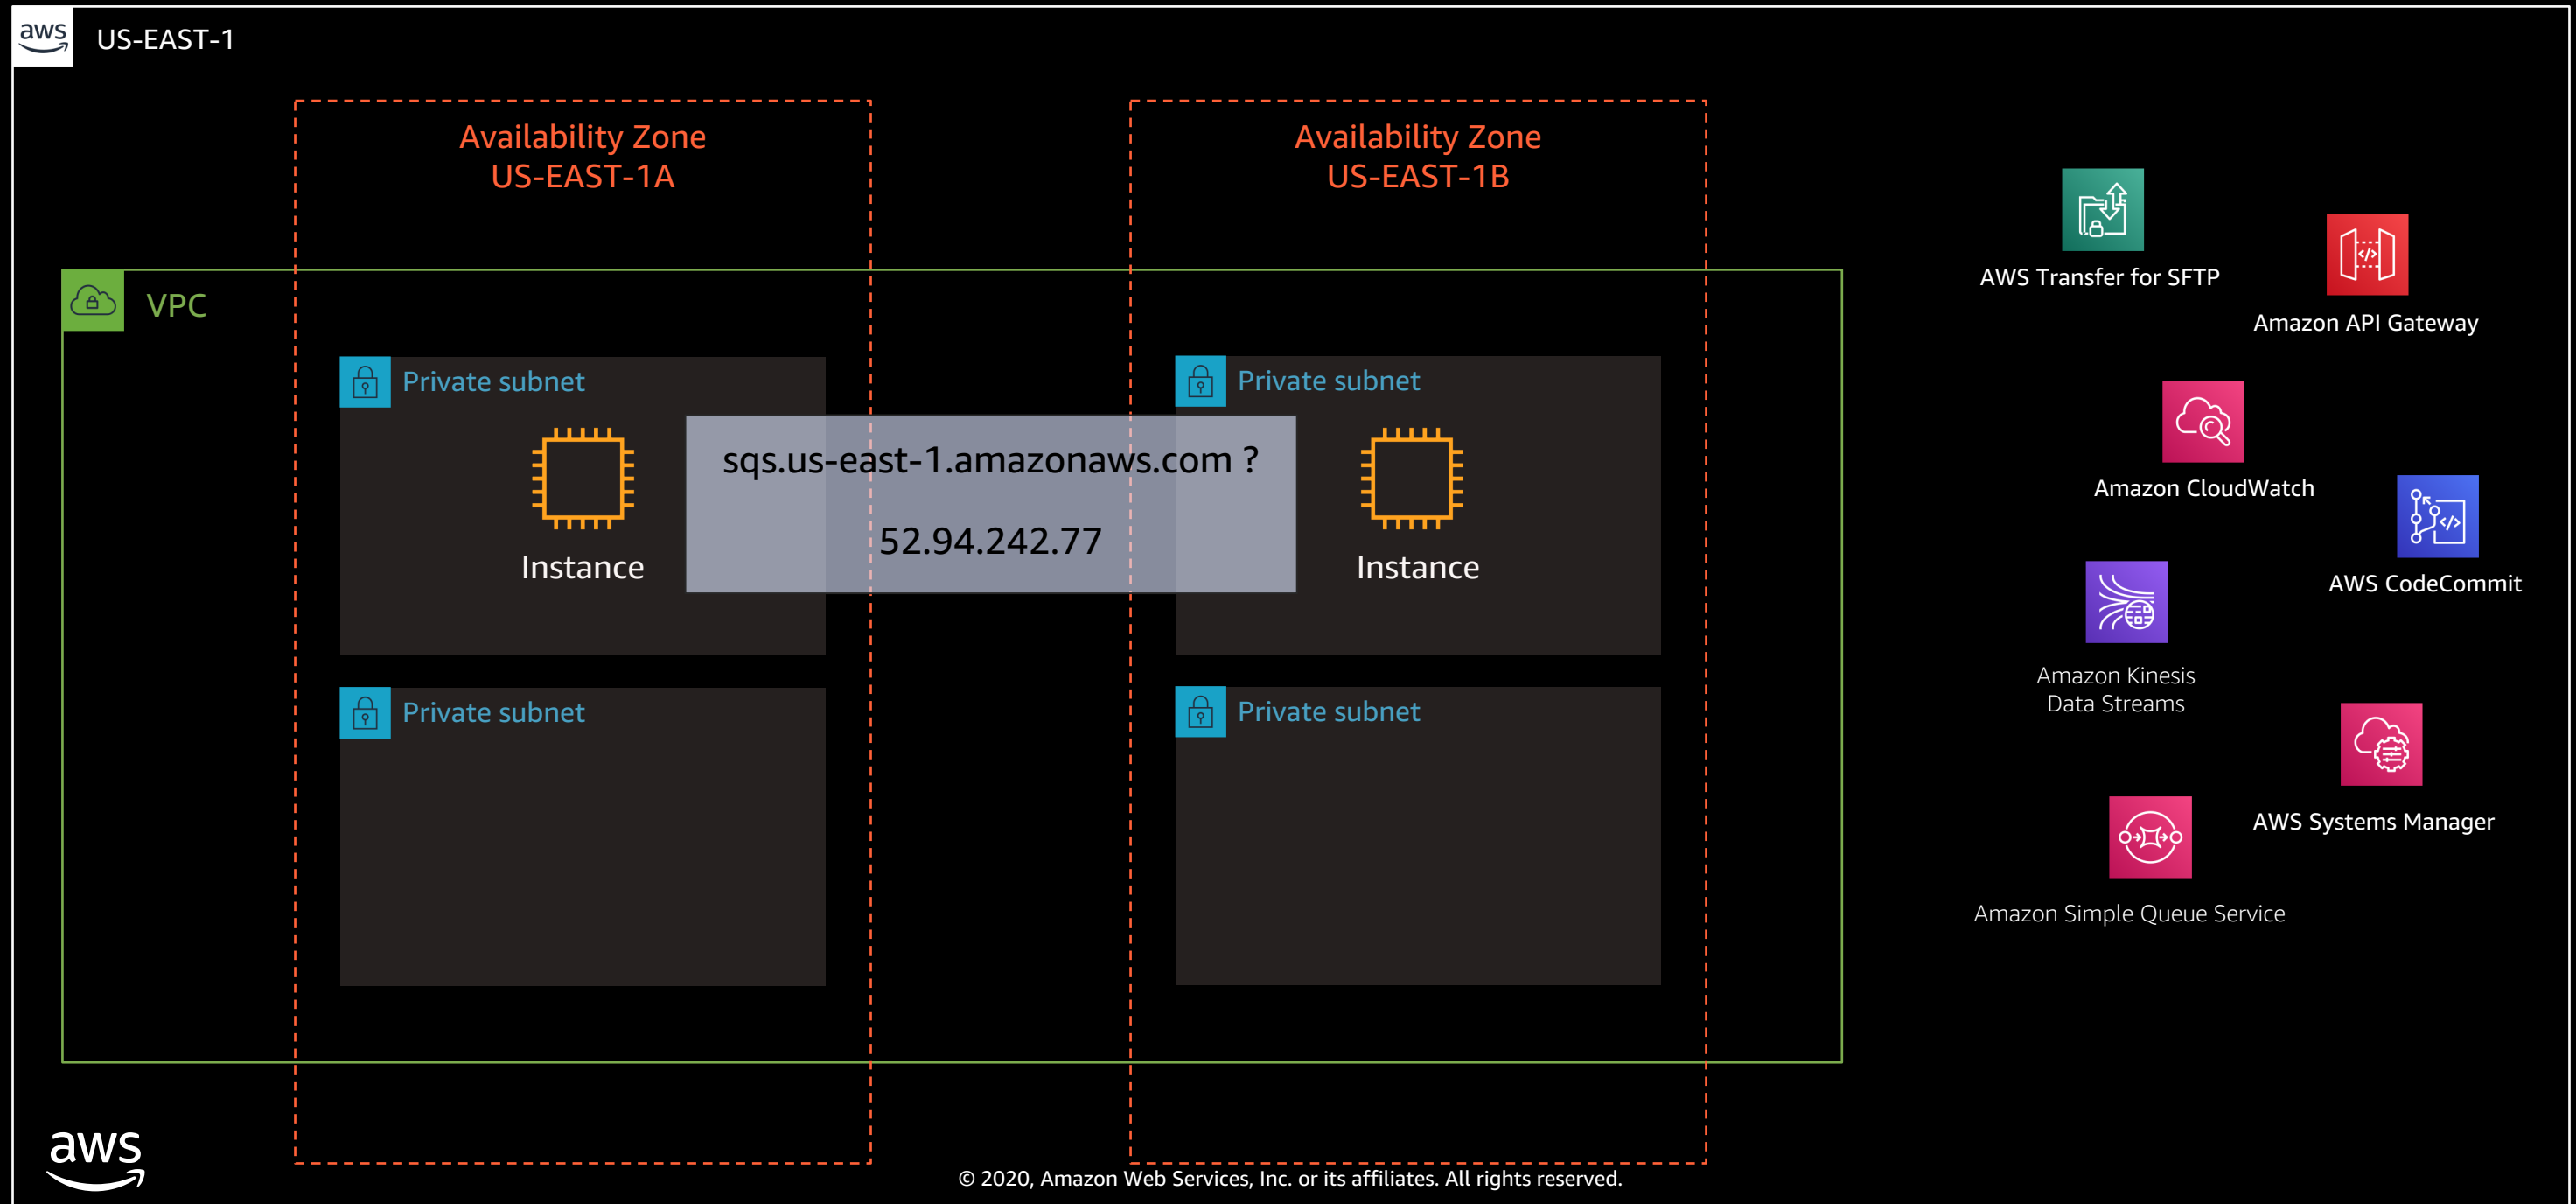
# Gateway VPC endpoints



# Interface VPC endpoints (AWS PrivateLink)

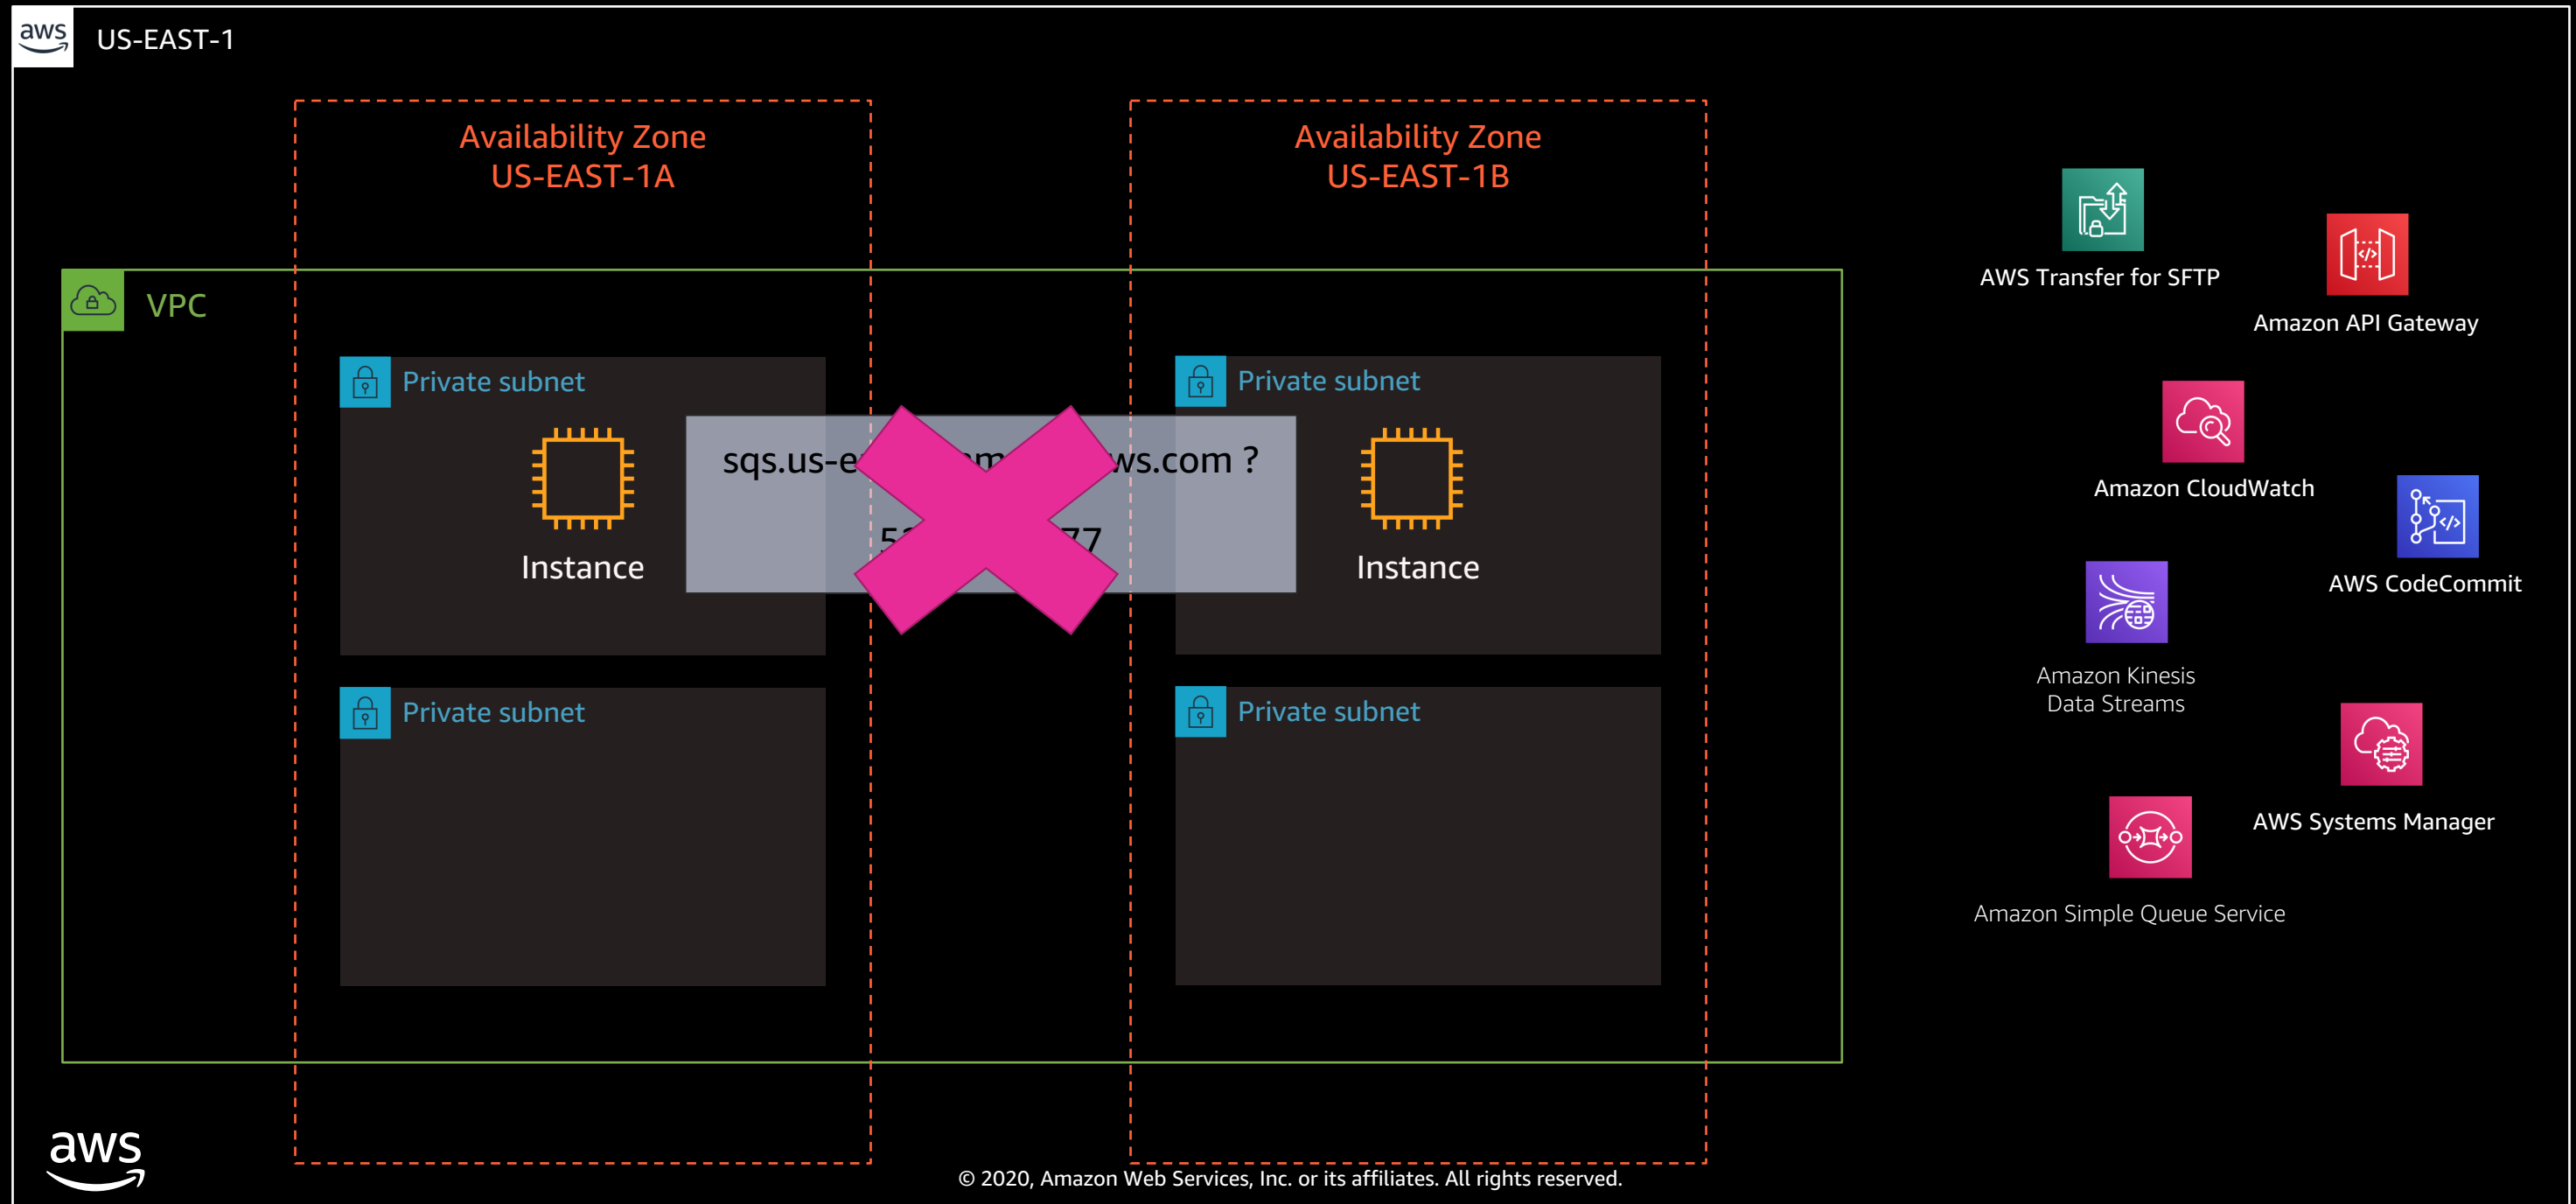


# Interface VPC endpoints (AWS PrivateLink)

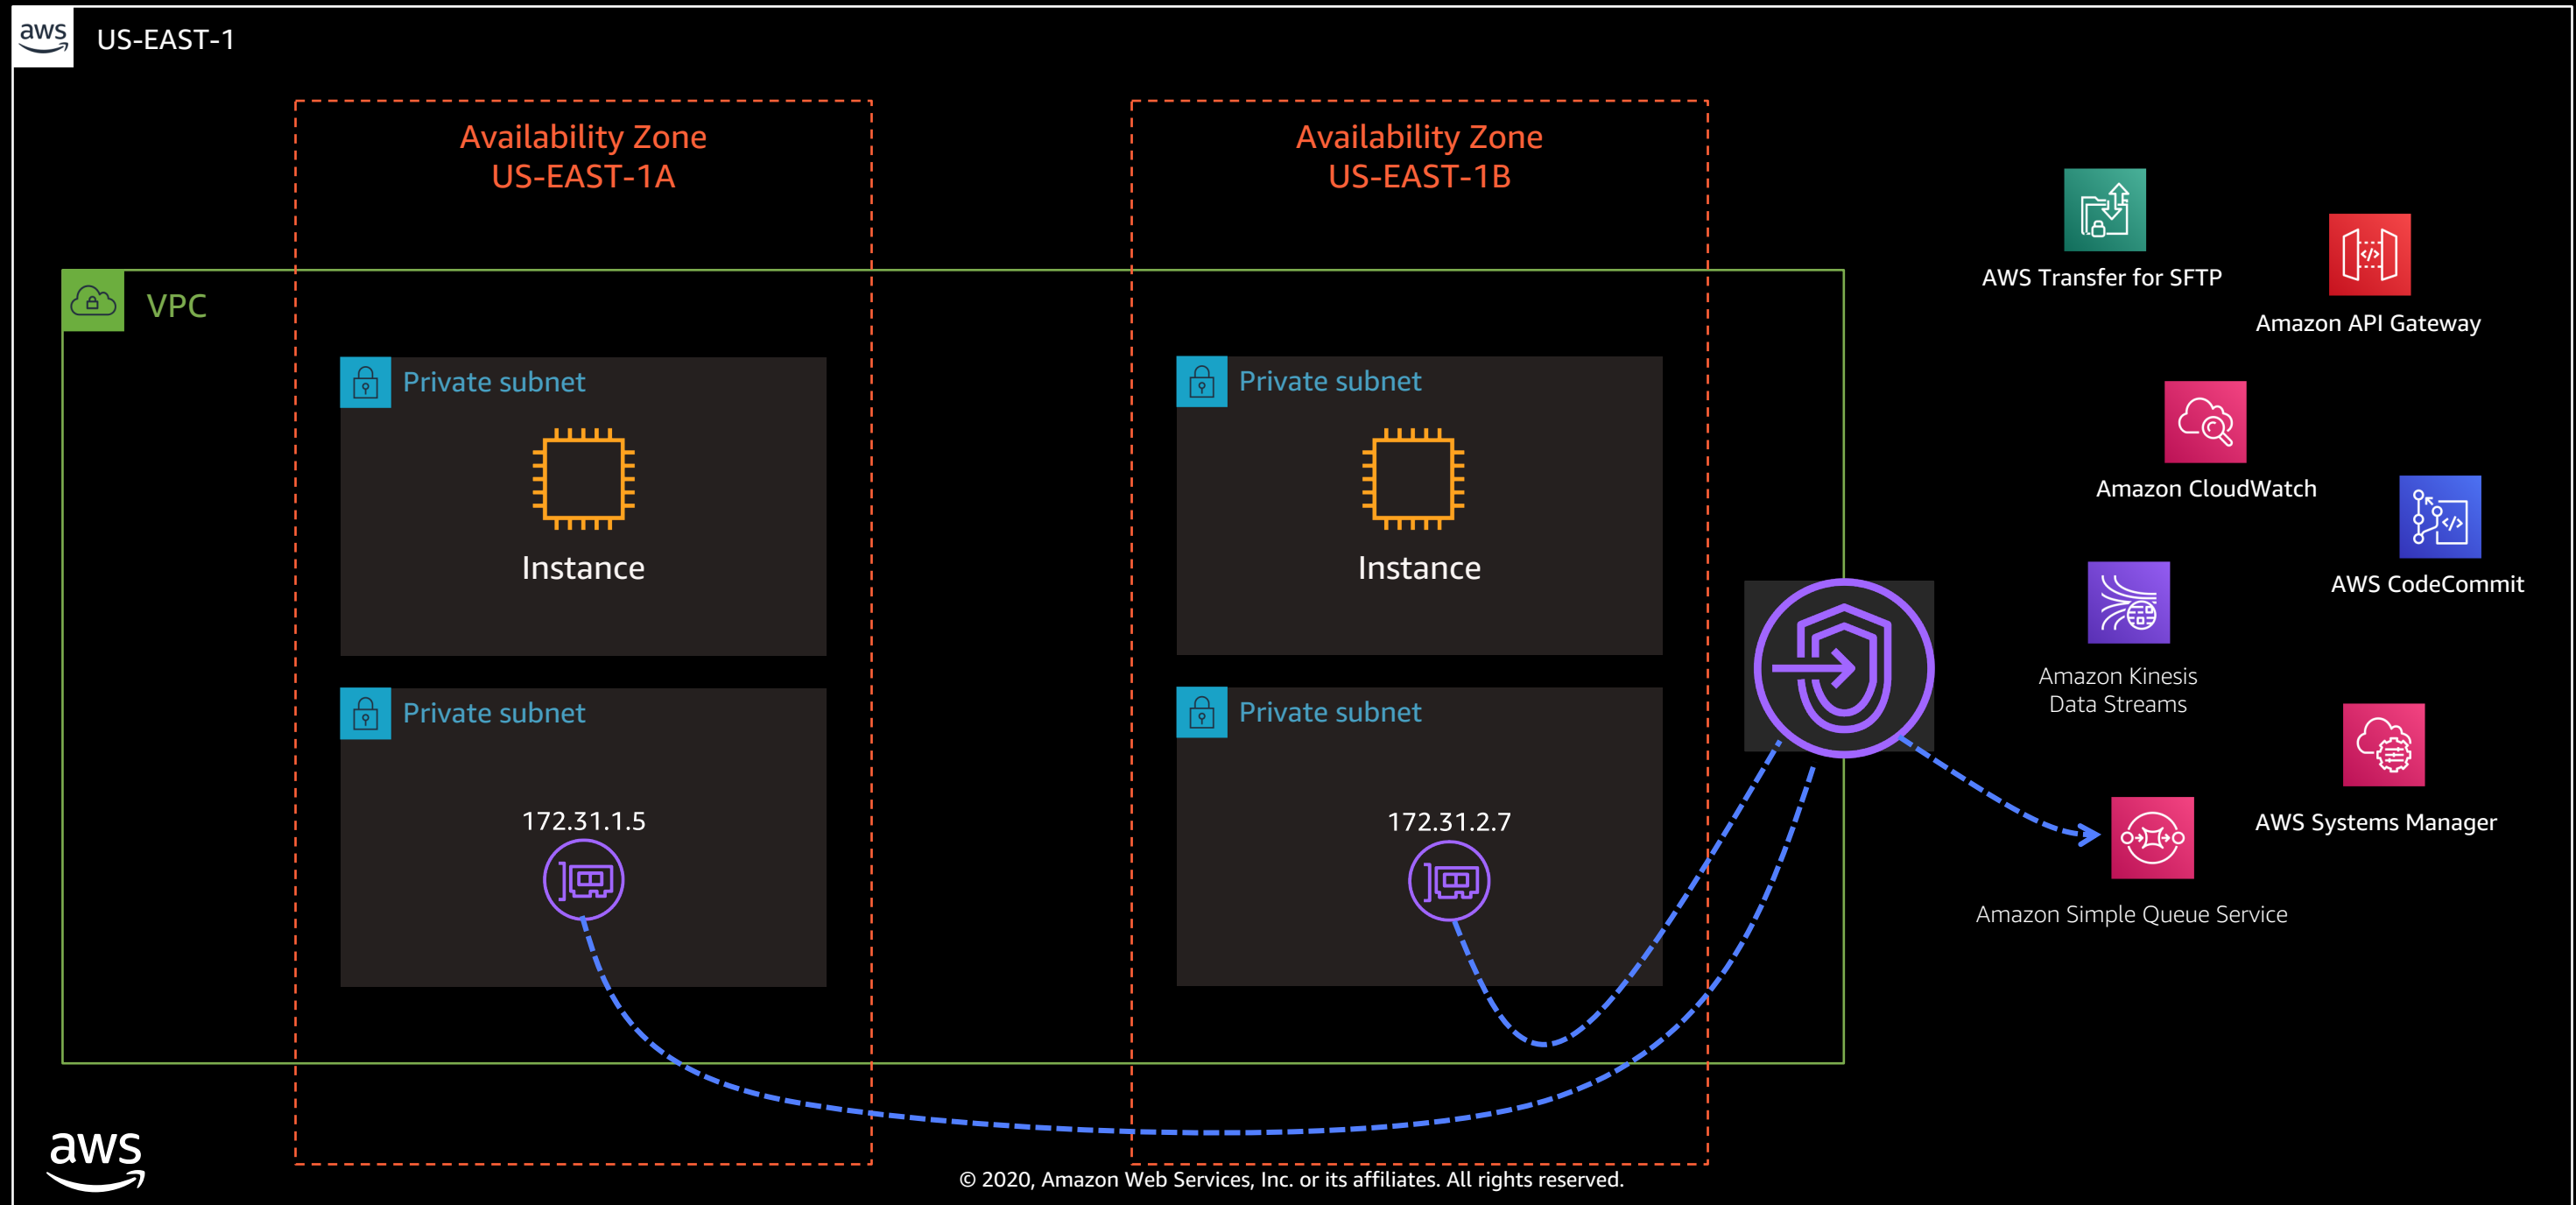




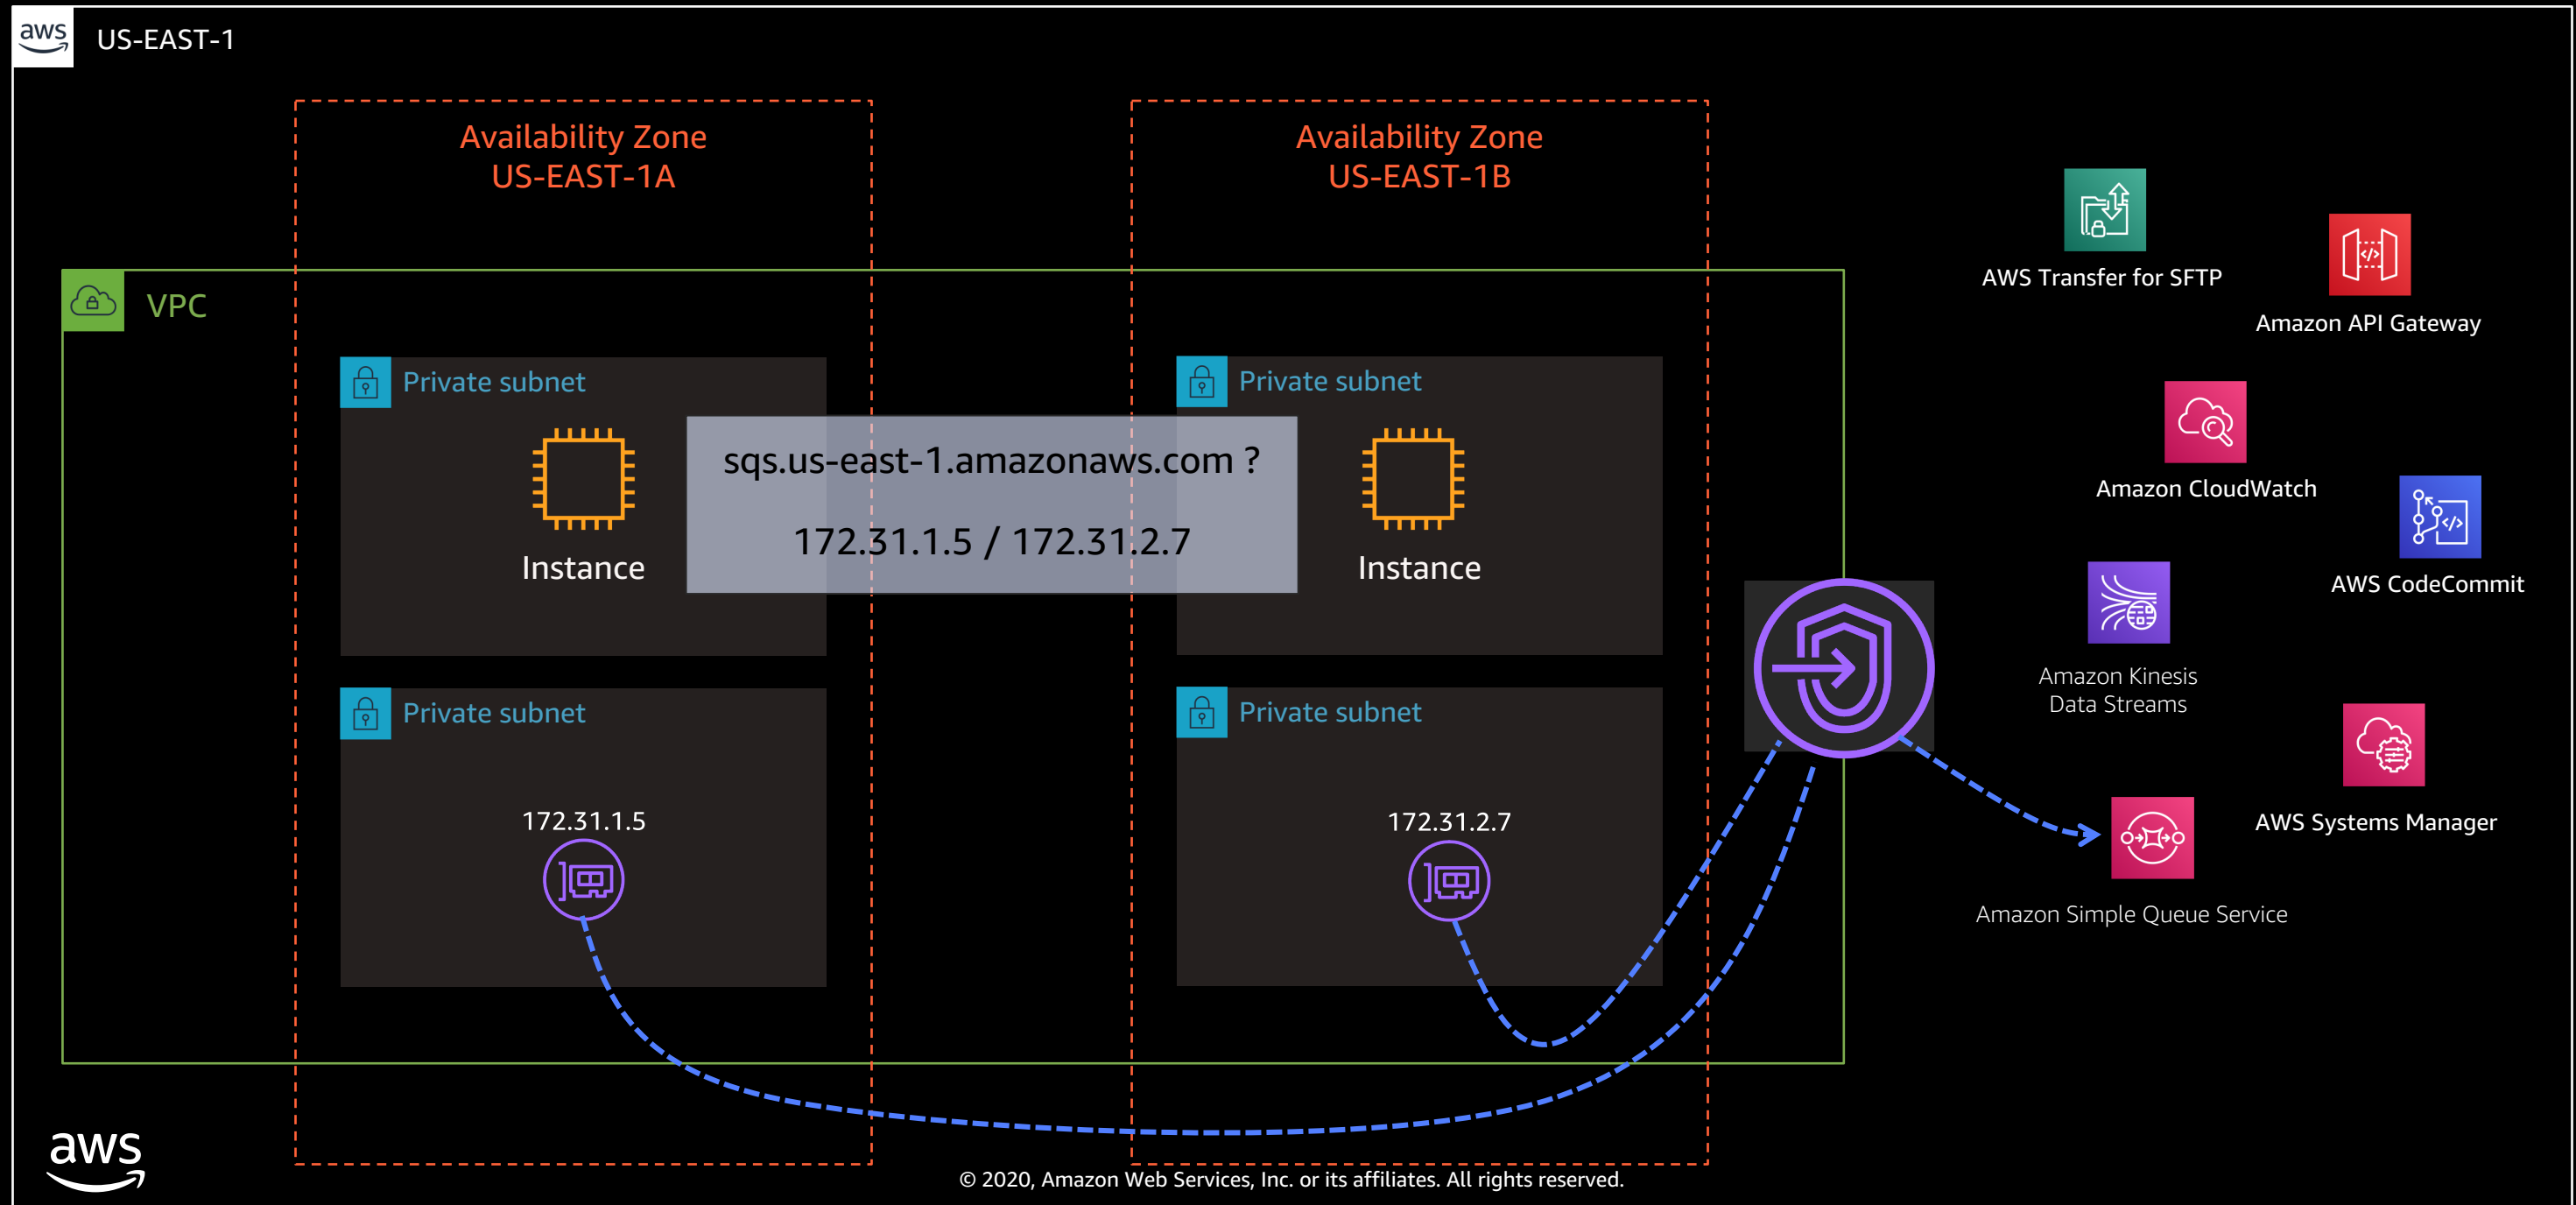
# Interface VPC endpoints (AWS PrivateLink)



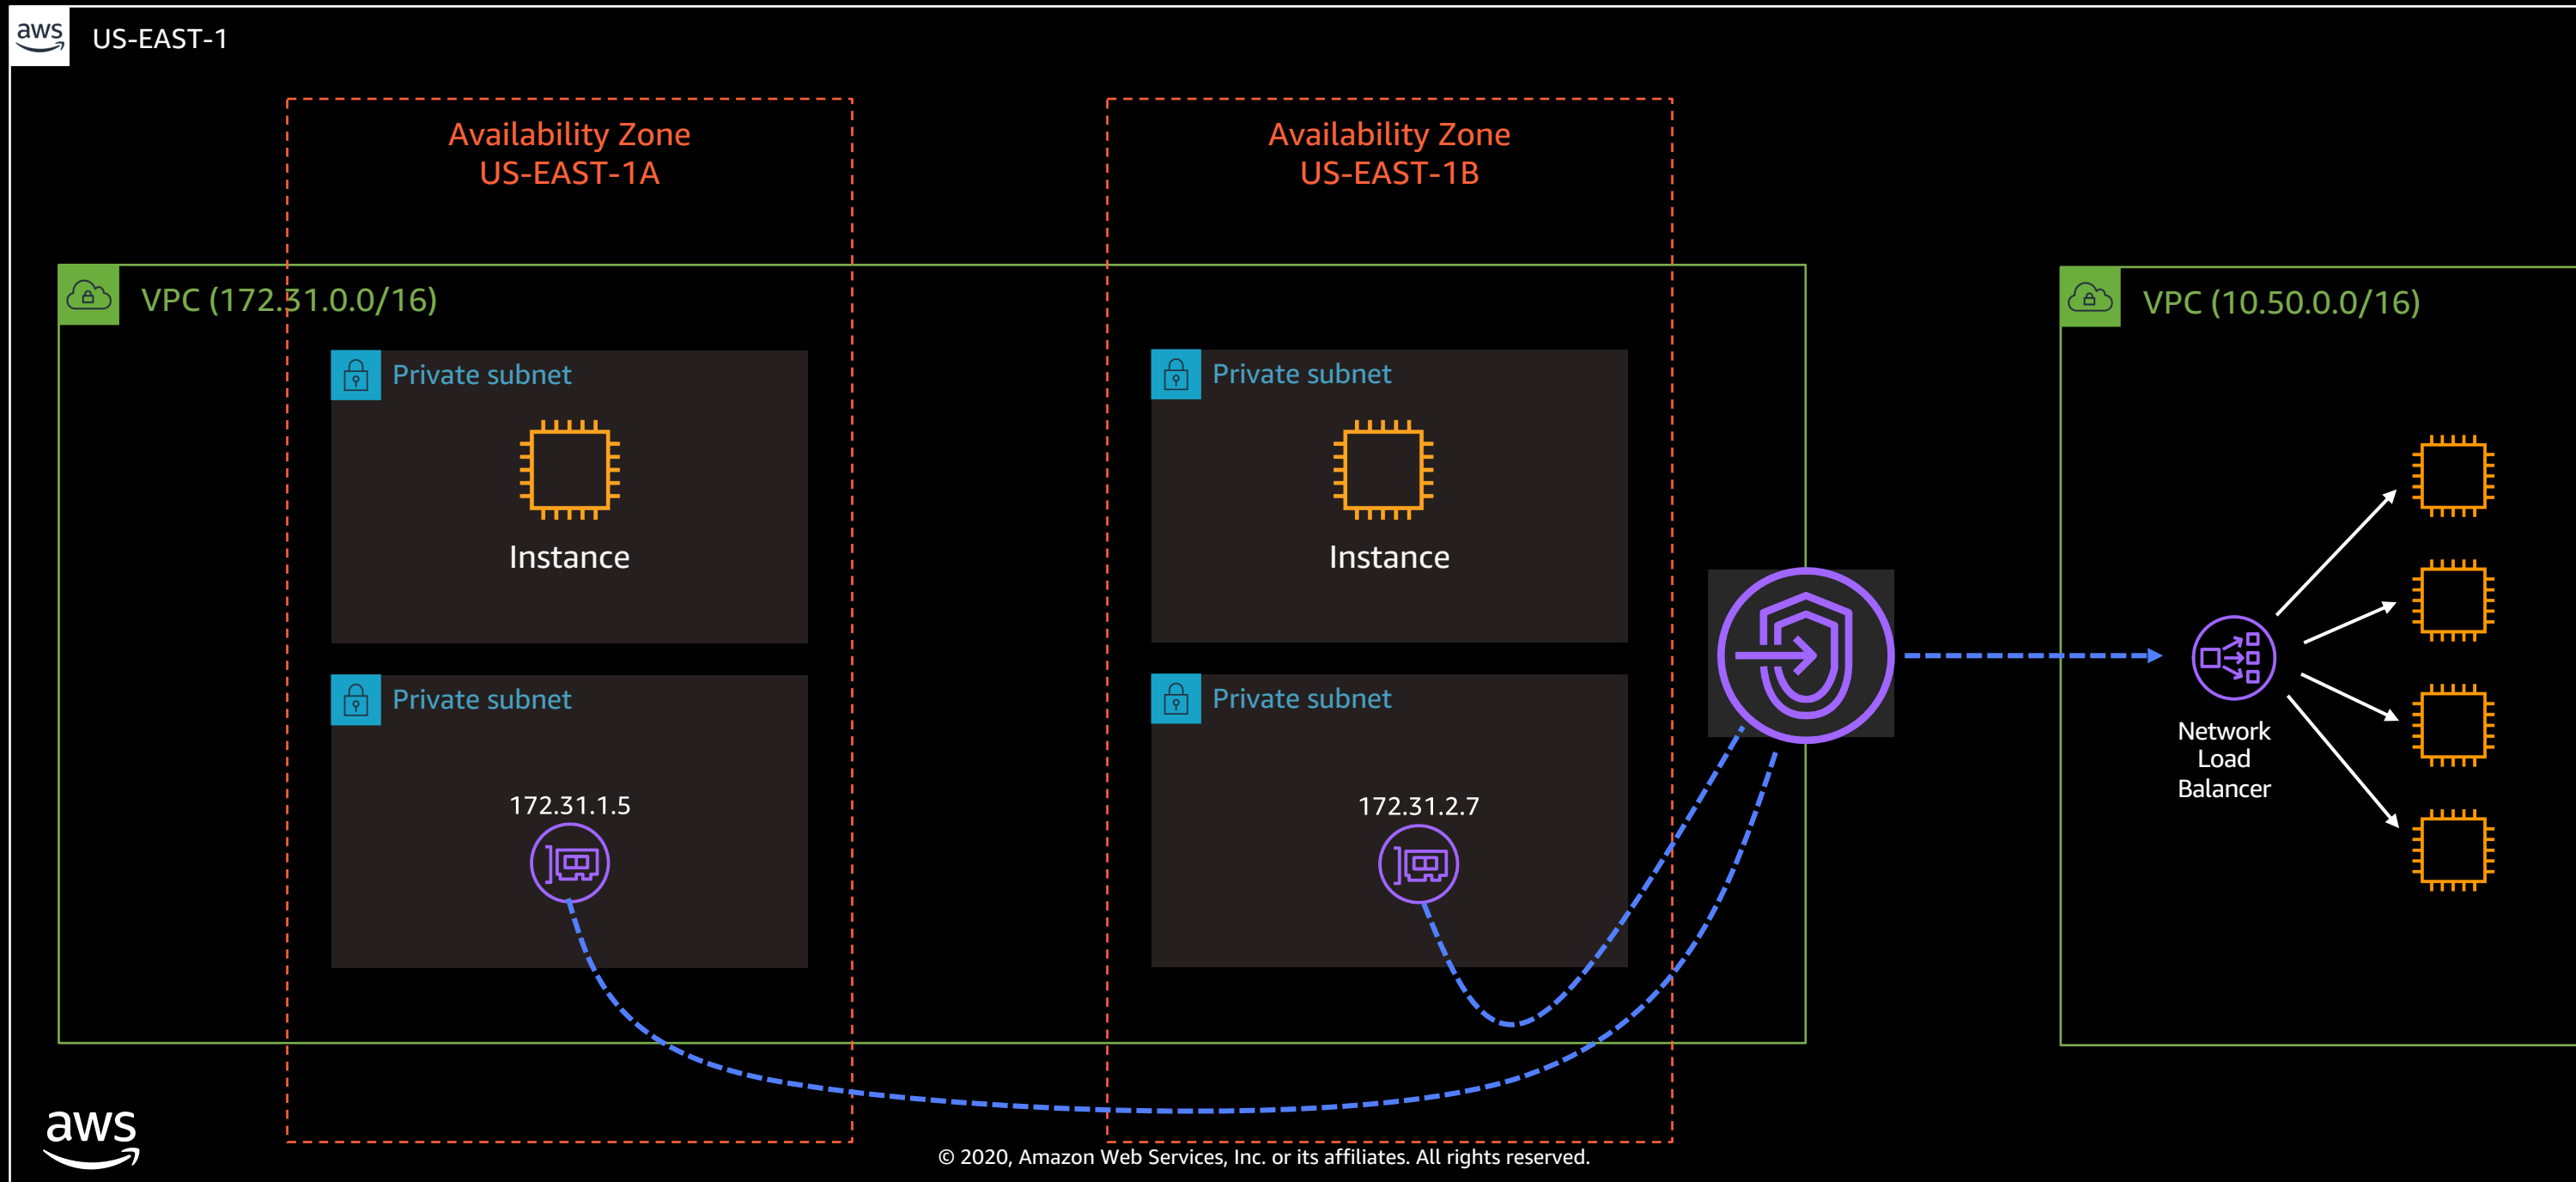
# Interface VPC endpoints (AWS PrivateLink)



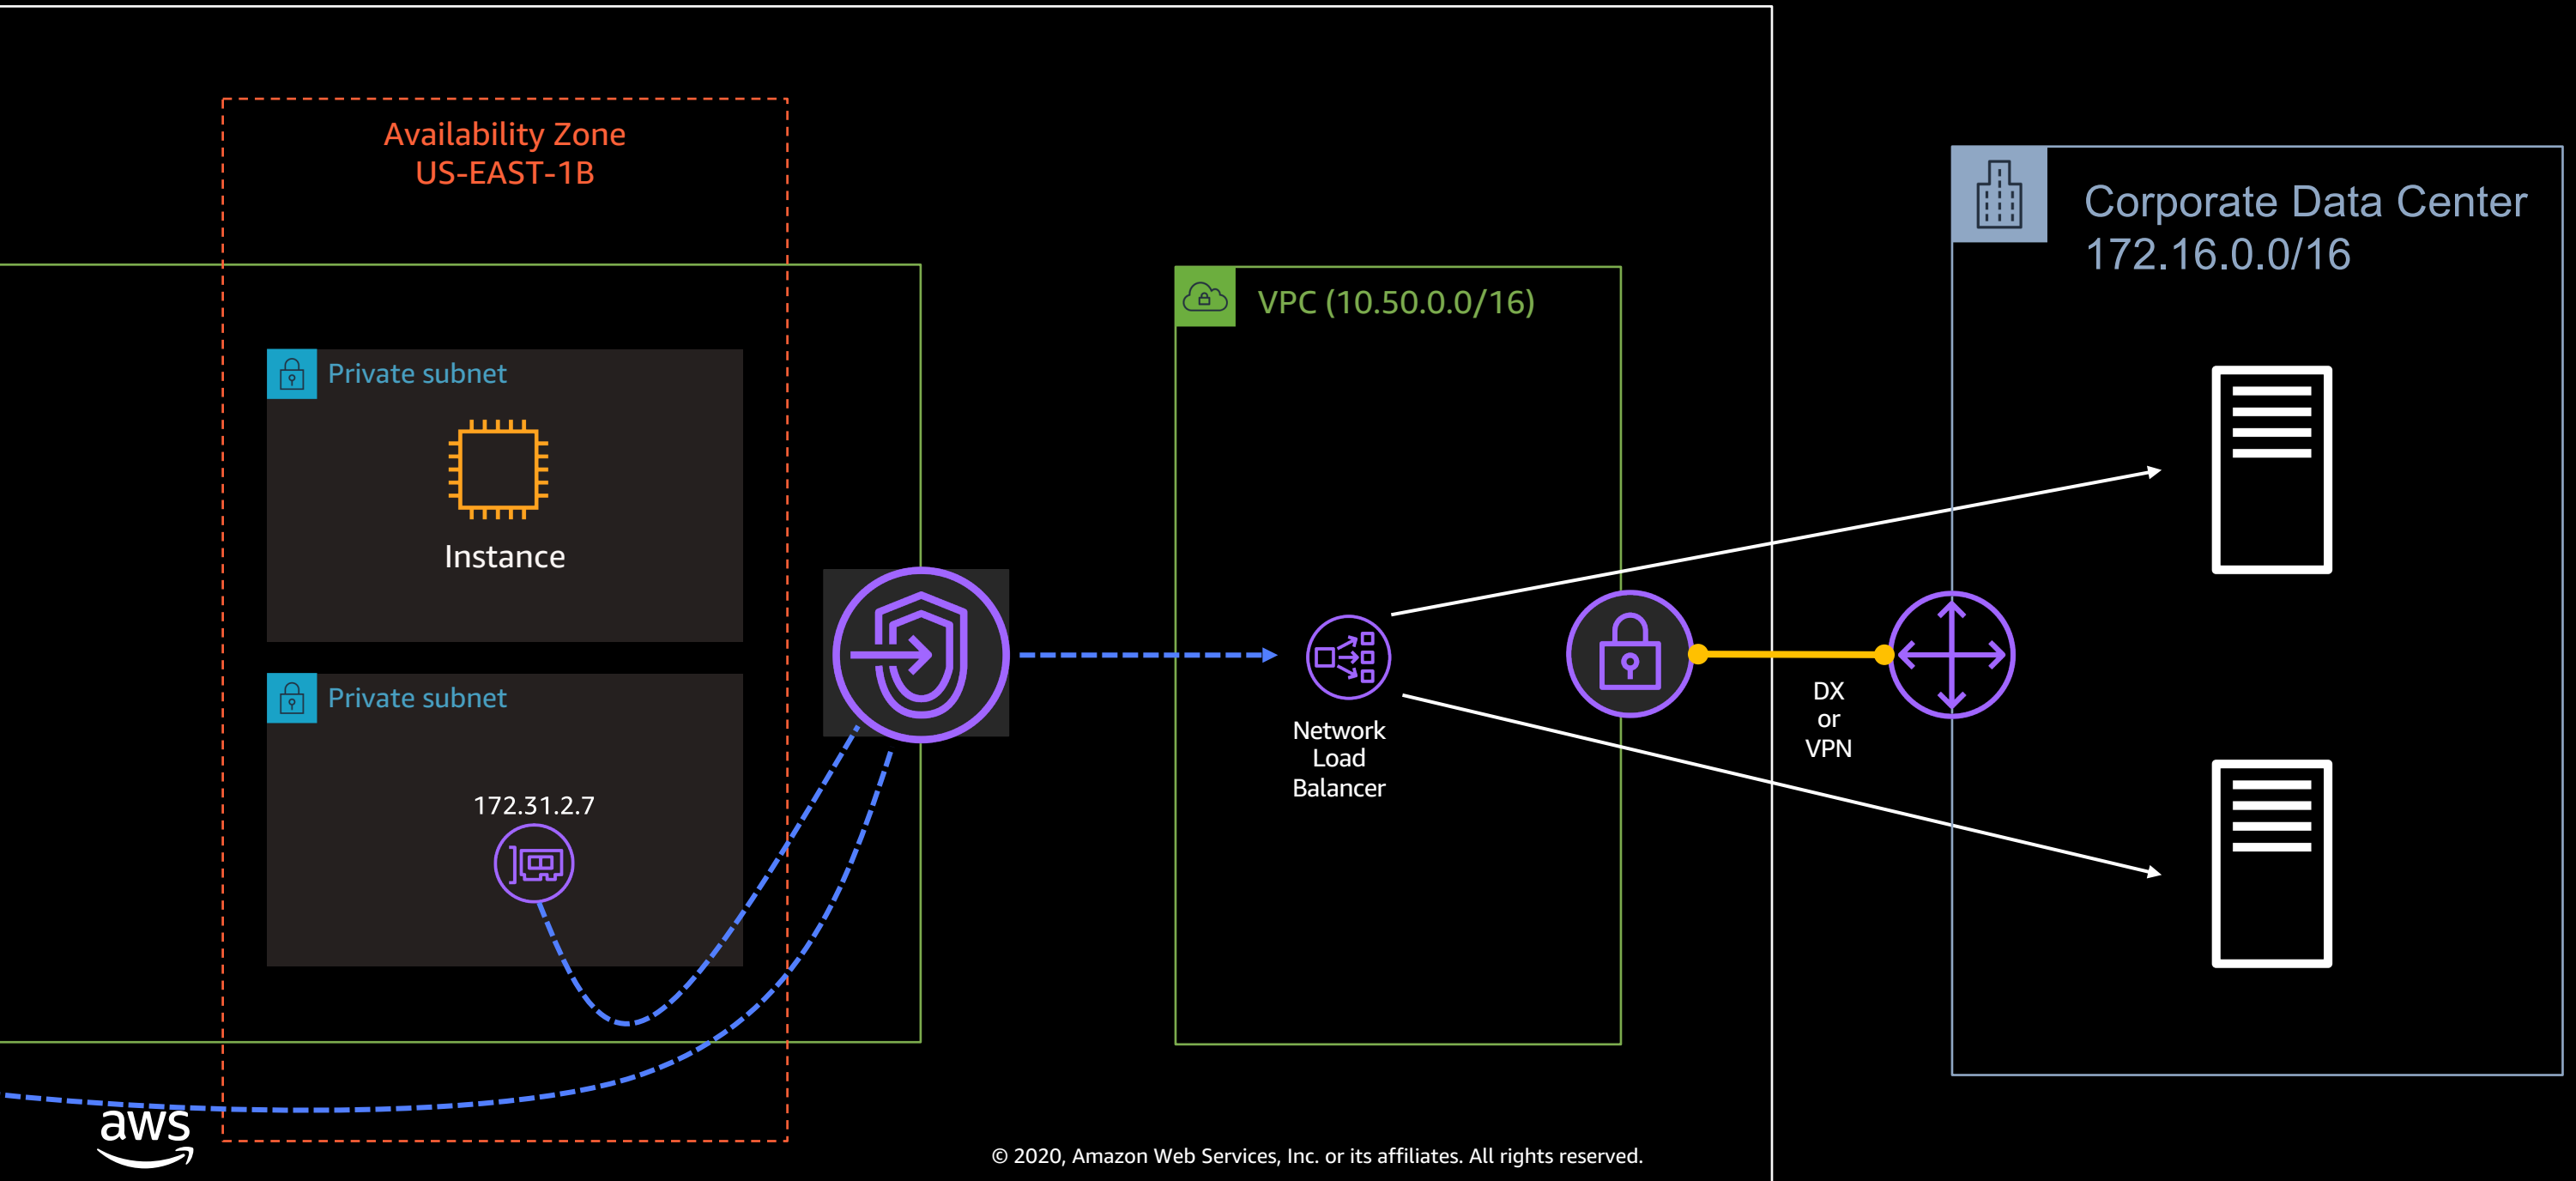
# Interface VPC endpoints (AWS PrivateLink)



# AWS PrivateLink – your own services



# AWS PrivateLink – Your own services – On-prem



# Endpoint policies



- A VPC endpoint policy is an AWS Identity and Access Management (IAM) resource policy that you attach to an endpoint
- An endpoint policy does not override or replace IAM user policies or service-specific policies (such as S3 bucket policies)

## Example for S3

- IAM policy at VPC endpoint: You may only access the “Data” bucket
- IAM policy at S3 bucket: Access to this bucket is only allowed from VPCE-X

# Shared VPCs



# Account and VPC segmentation

## Larger VPCs or accounts

- AWS Identity and Access Management (IAM)
- Strict security groups and routing
- Identifying resources with tags

### Policy and IAM

## Smaller VPCs or accounts

- Automation of infrastructure
- AWS Direct Connect and VPN standards
- Subnet and routing standards

### Infrastructure and networking

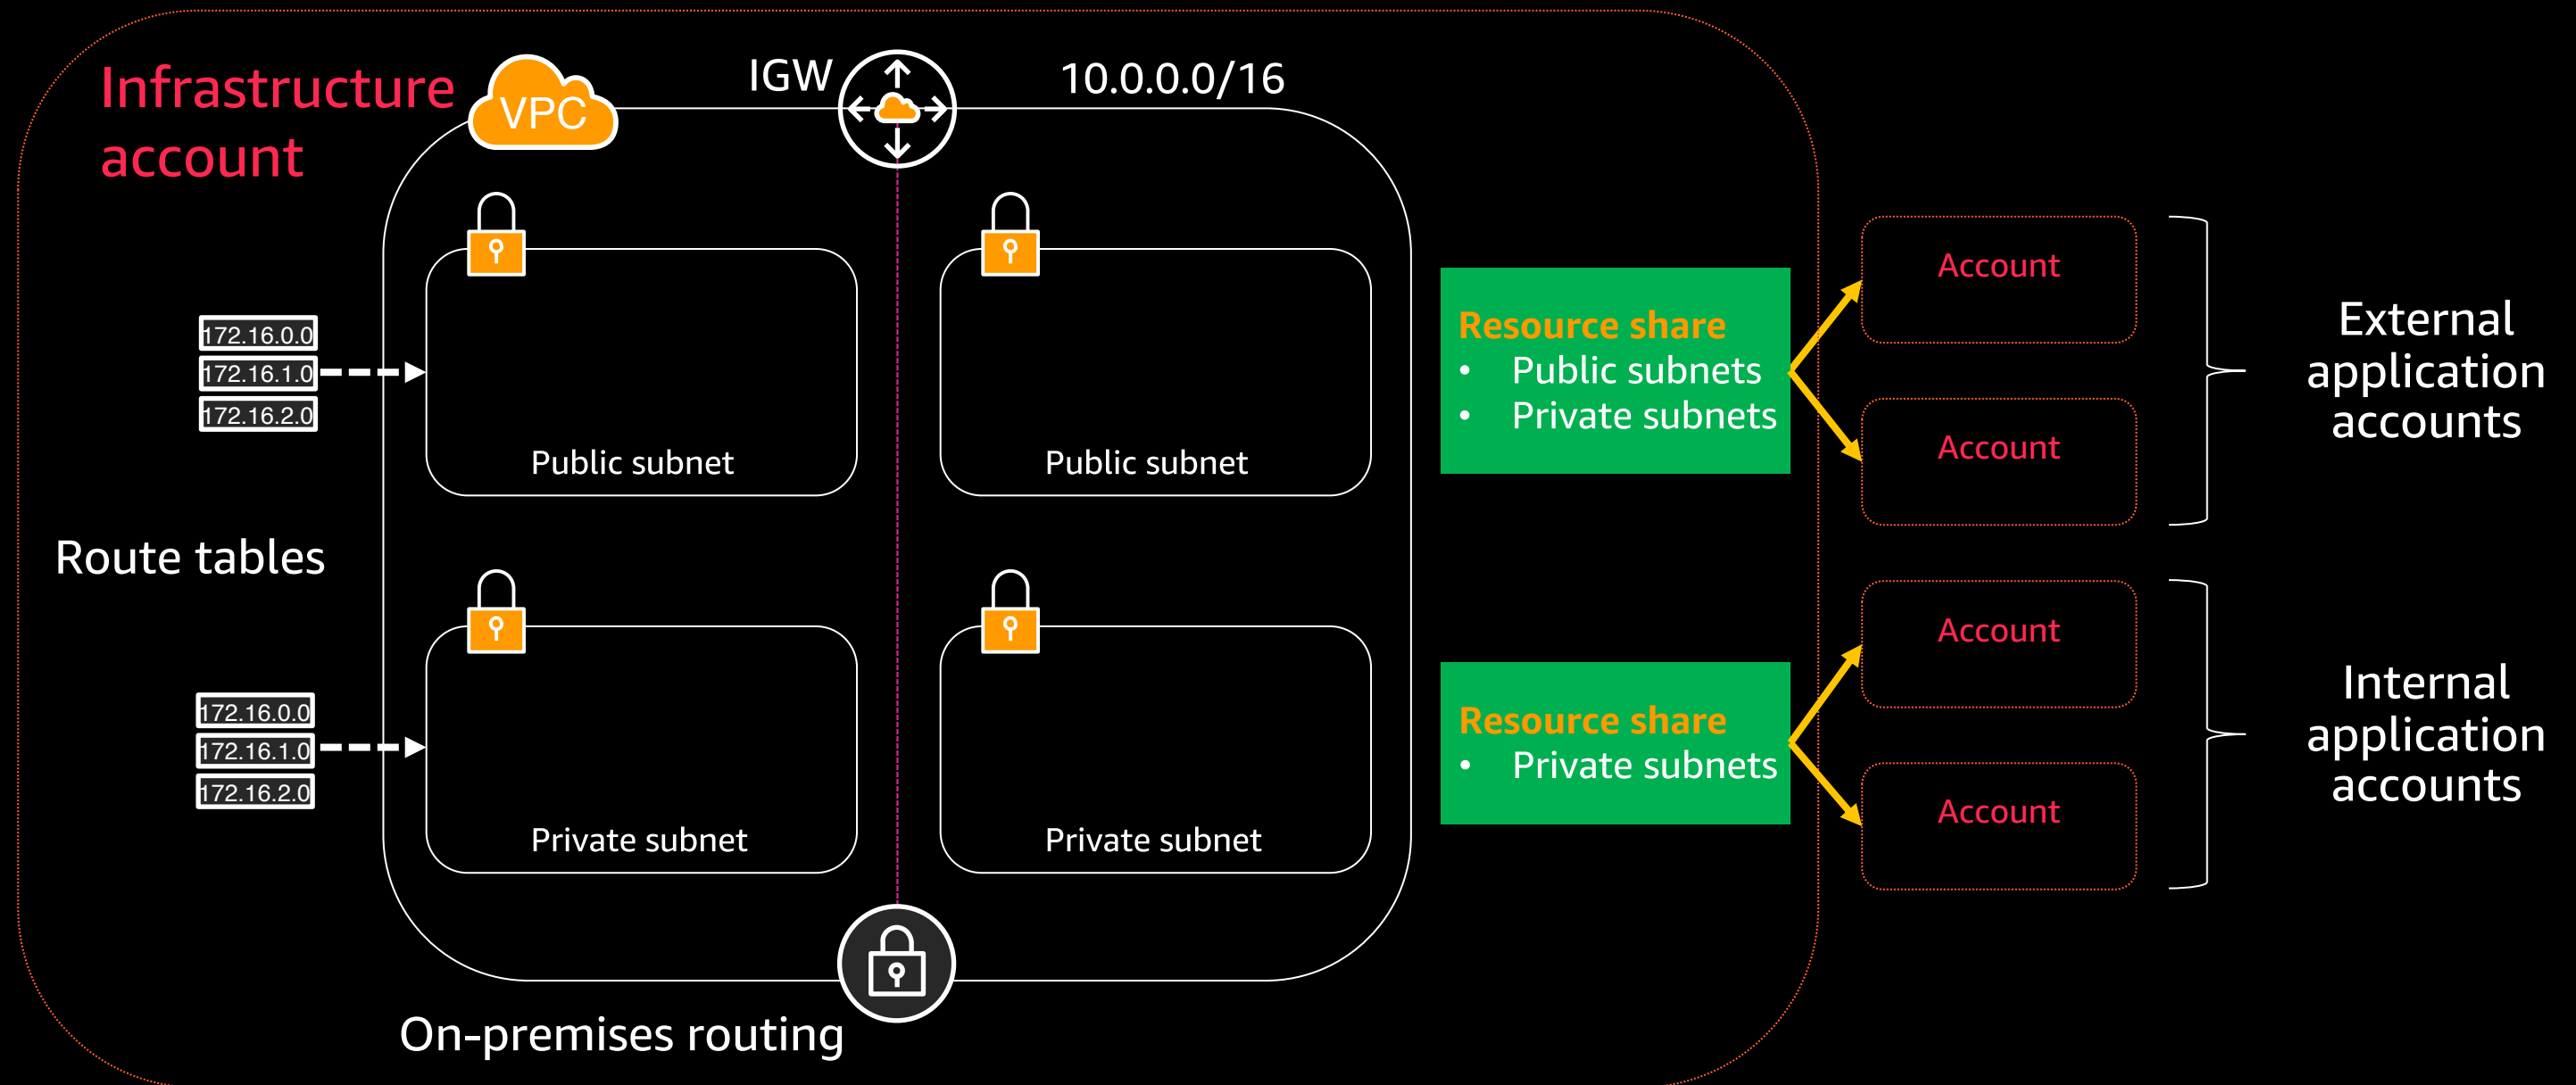


# Why not both?

Provide granular account control  
with centralized infrastructure

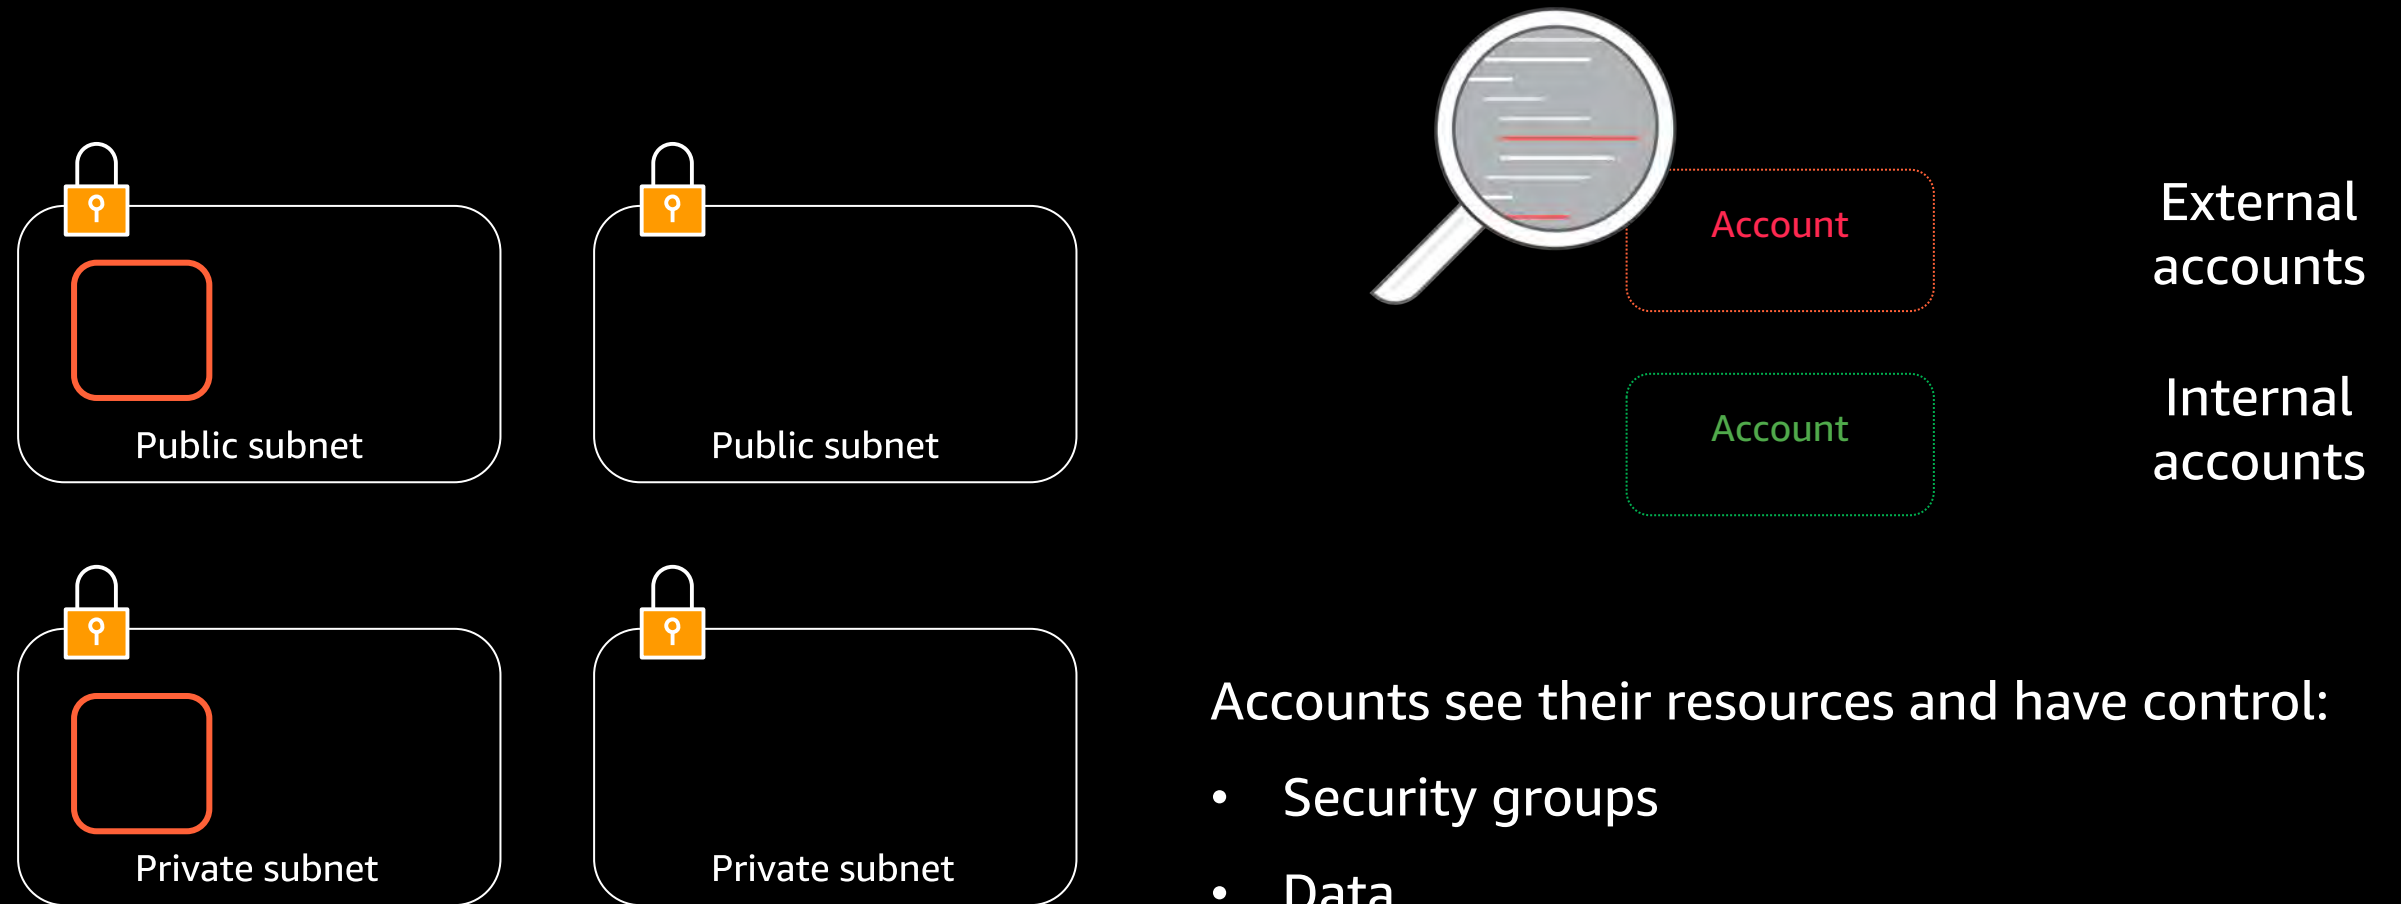
# VPC Sharing and Resource Access Manager

## Share subnets between accounts in an AWS Organization



# VPC Sharing and Resource Access Manager

Account owners only see subnets and their resources



Accounts see their resources and have control:

- Security groups
- Data
- Instance details
- Account configuration

# VPC Sharing and Resource Access Manager


Account owners only see subnets and their resources



Accounts see their resources and have control:


- Security groups
- Data
- Instance details
- Account configuration

# VPC Sharing benefits




## Separation of duties

- Infrastructure strictly controls routing, IP addresses, and VPC structure
- Developers own their resources, accounts, and security groups



## Less unused resources

- Higher density subnets, add up to 5 additional CIDRs
- More efficient use of VPN and AWS Direct Connect



## Decouple accounts and networks

- Account protection and billing without additional infrastructure
- Many accounts with fewer networks
- Avoid VPC peering charges

# Other account considerations

## One size does not need to fit all

- Example: production may use separate VPCs, development can use a shared VPC
- **AWS Transit Gateway can handle large amounts of VPCs if needed**

## VPC Sharing works within an AWS Organization

## VPC Sharing doesn't restrict resource utilization


- NAT gateways, VPN, subnet address space, and security groups have shared limits
- VPC Sharing doesn't change any VPC limits, only account limits
- Give highly scalable services like AWS Lambda dedicated IP space

# 15 min Q/A and Break

# Connectivity options for VPCs





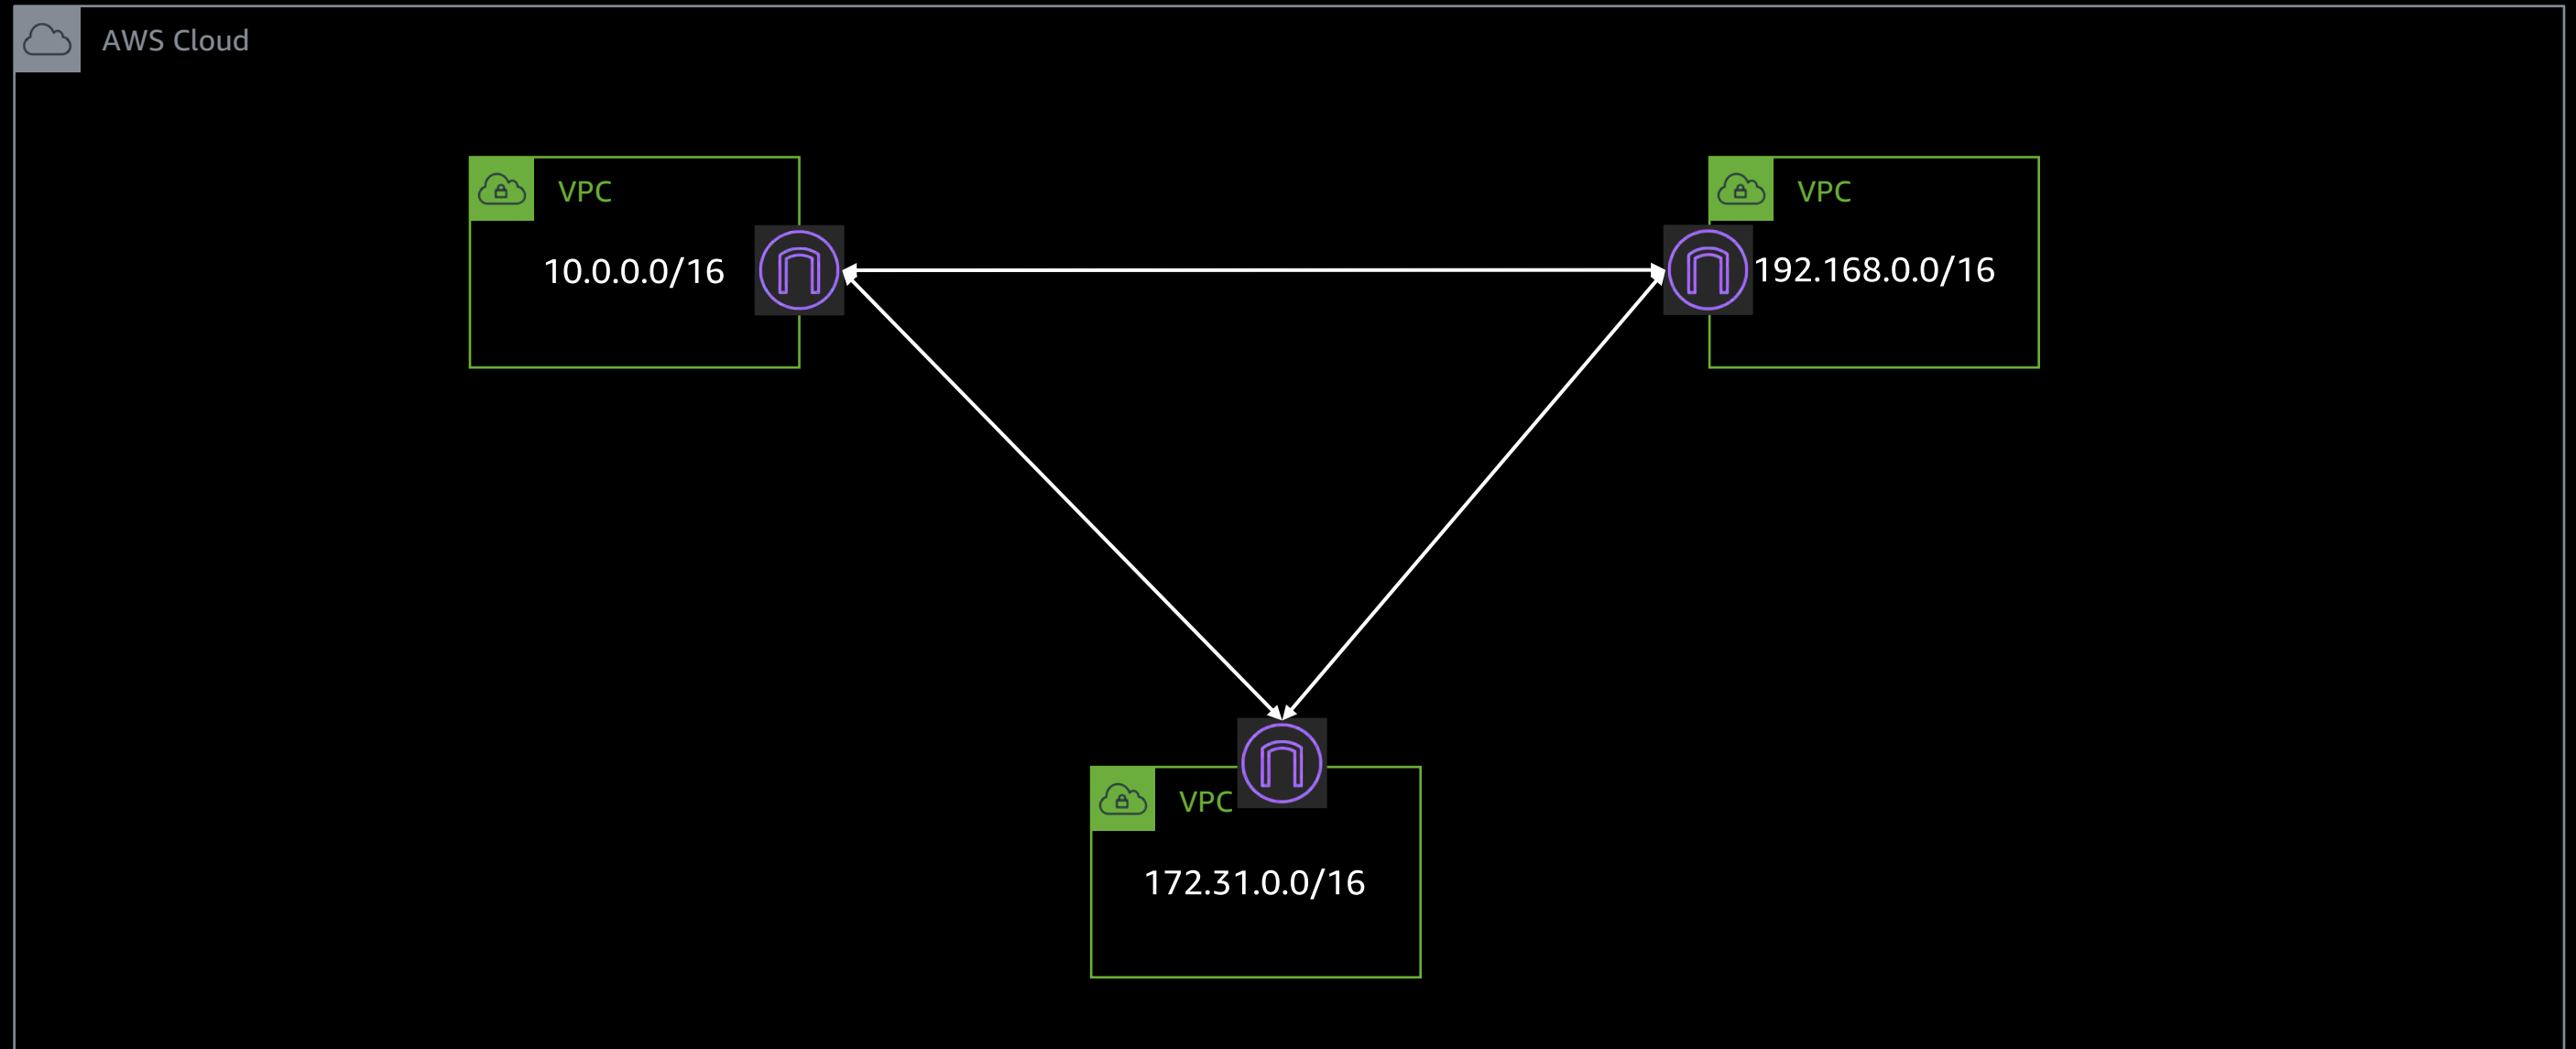


That was the agenda  
for this session

# Connecting to other VPCs



# Connecting between VPCs



# VPC peering – same region

AWS Cloud



VPC

10.0

### Create Peering Connection

Peering connection name tag

Select a local VPC to peer with

VPC (Requester)\*

CIDRs	CIDR	Status	Status Reason
	172.31.0.0/16	<span>●</span> associated	

Select another VPC to peer with

Account ☒ My account ☐ Another account

Region ☒ This region (us-east-1) ☐ Another Region

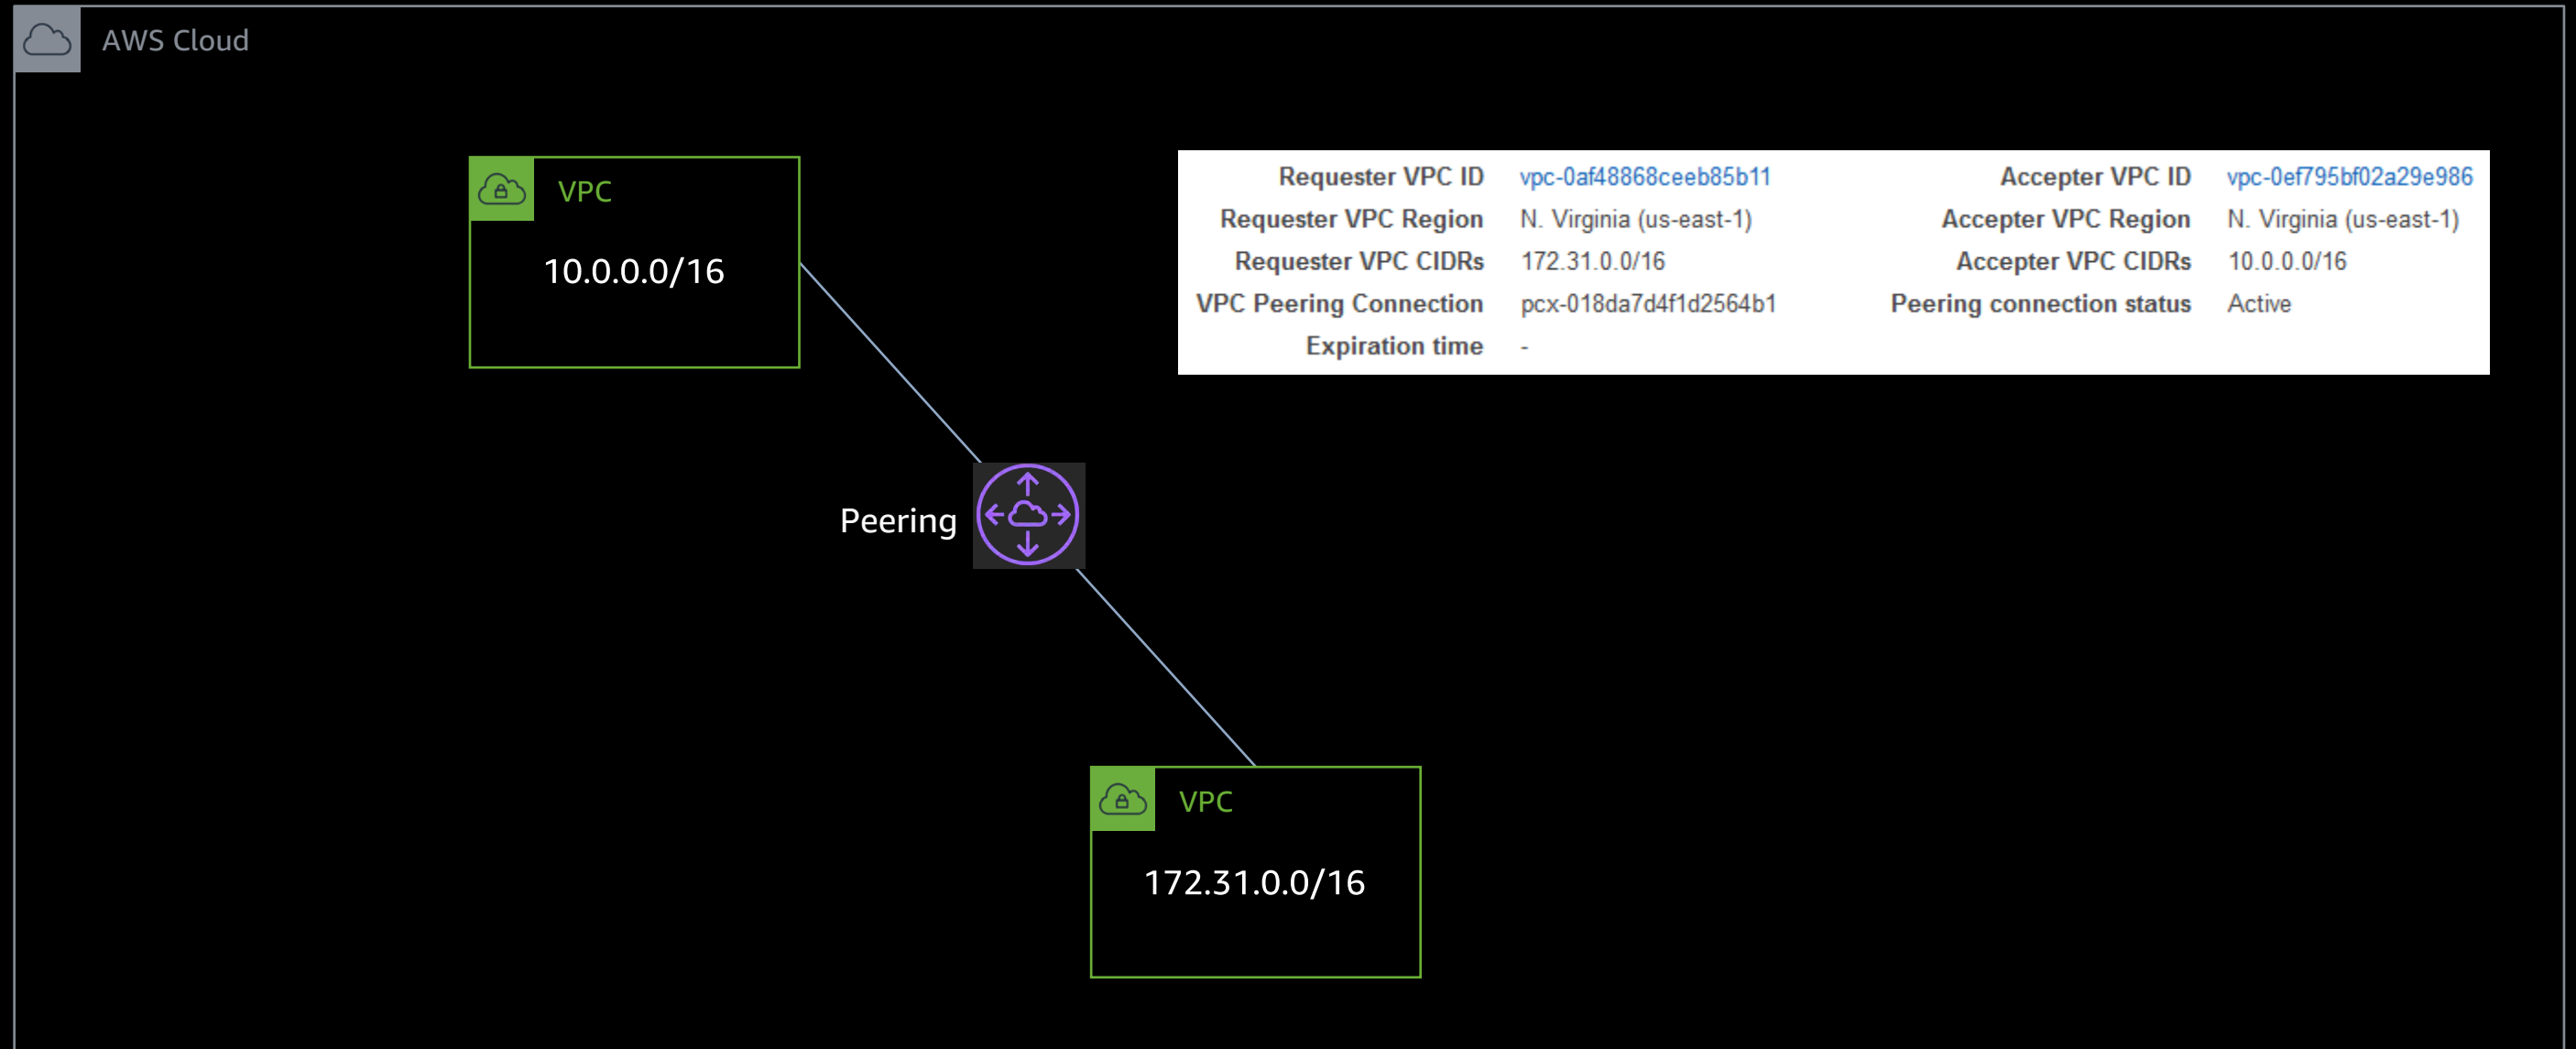
VPC (Acceptor)\*

CIDRs	CIDR	Status	Status Reason
	10.0.0.0/16	<span>●</span> associated	

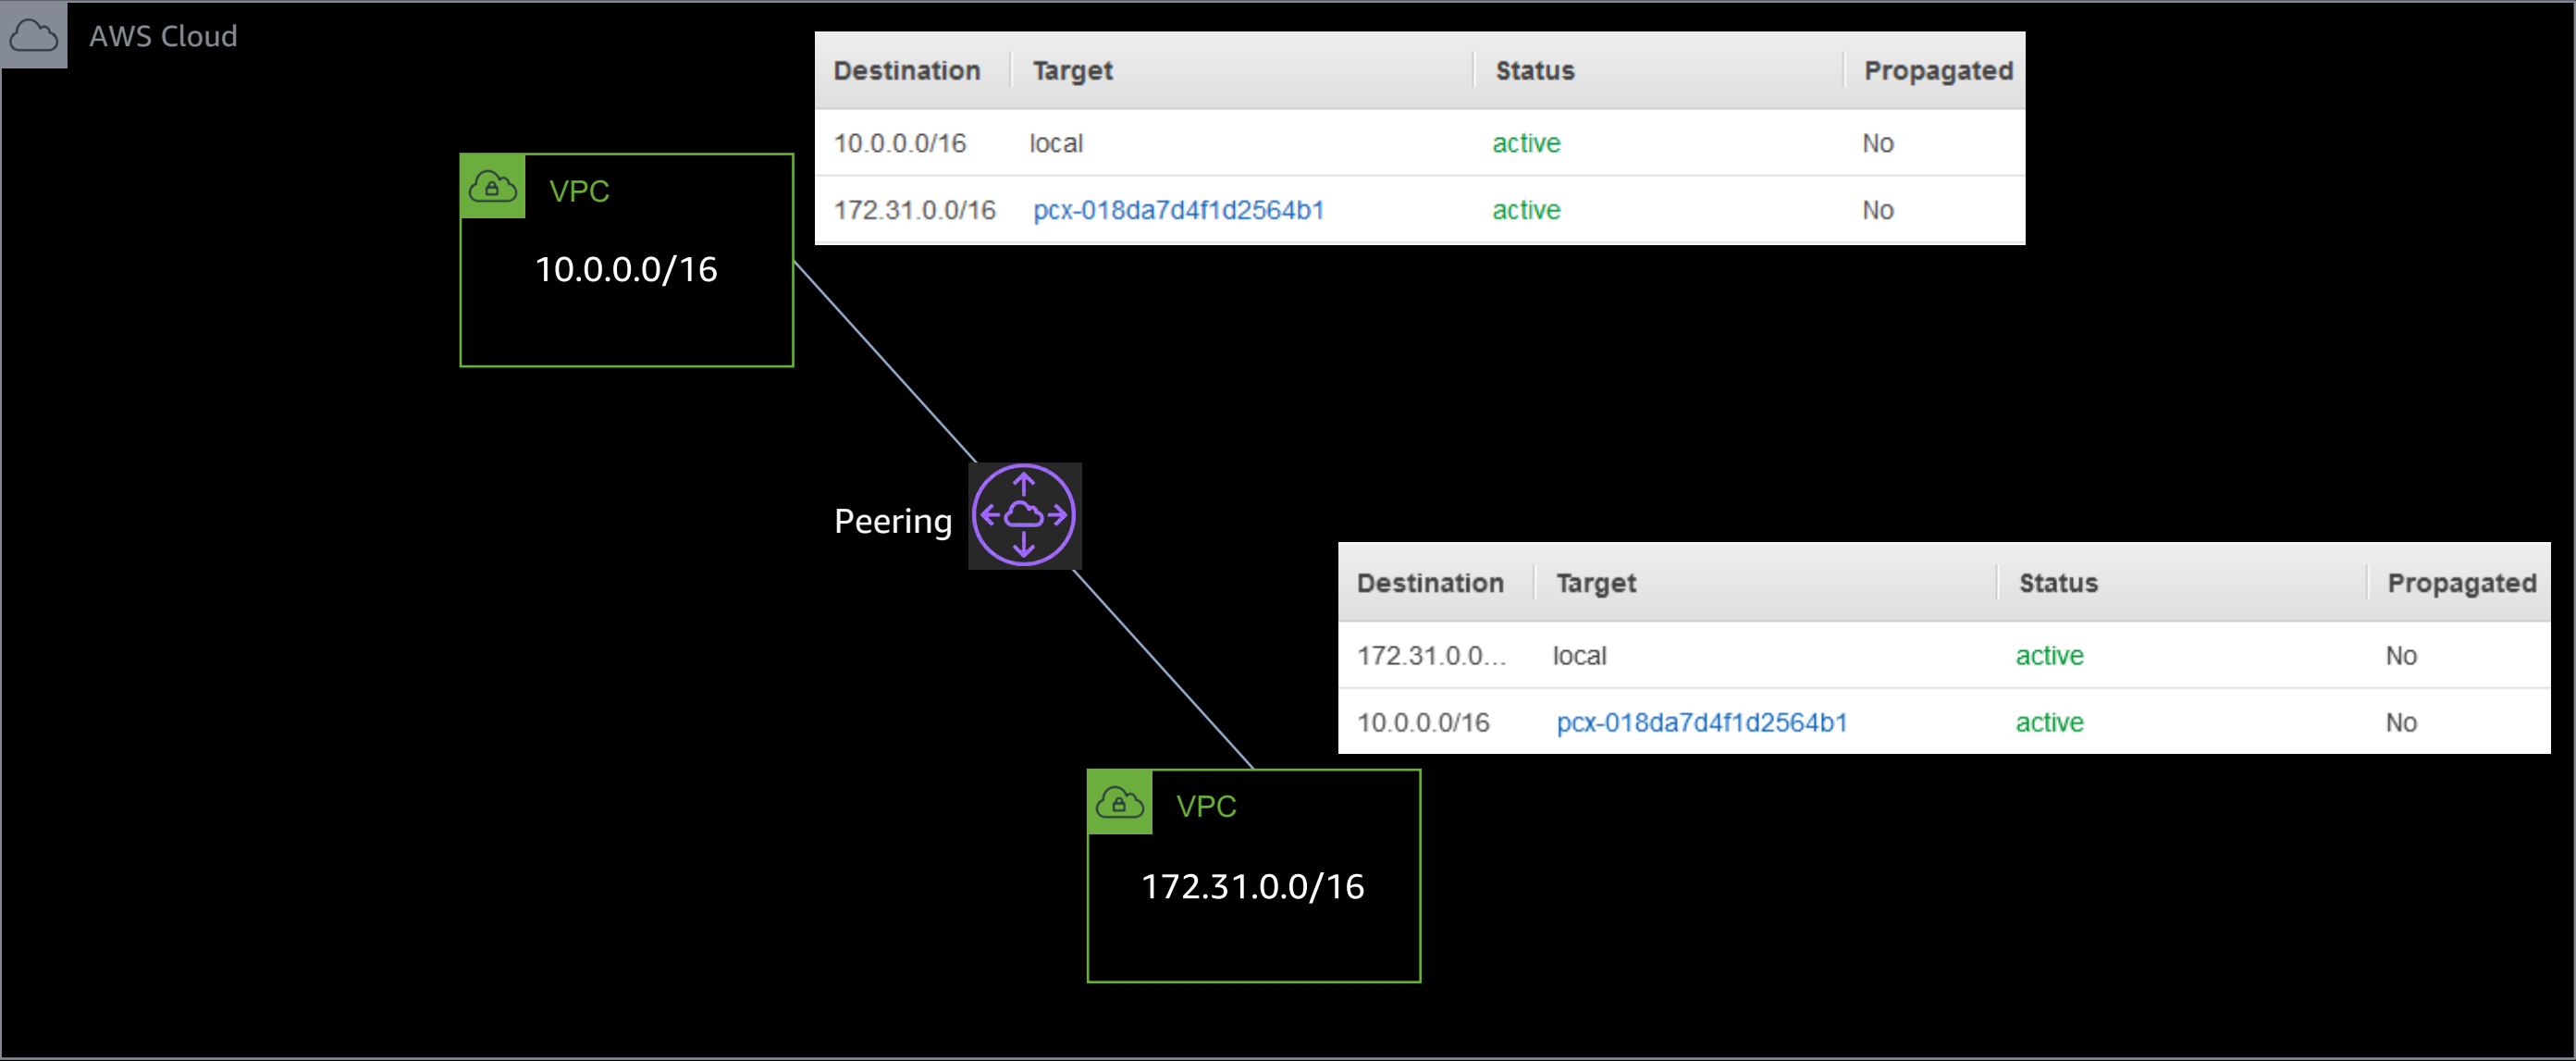
\* Required

[Cancel](#) [Create Peering Connection](#)

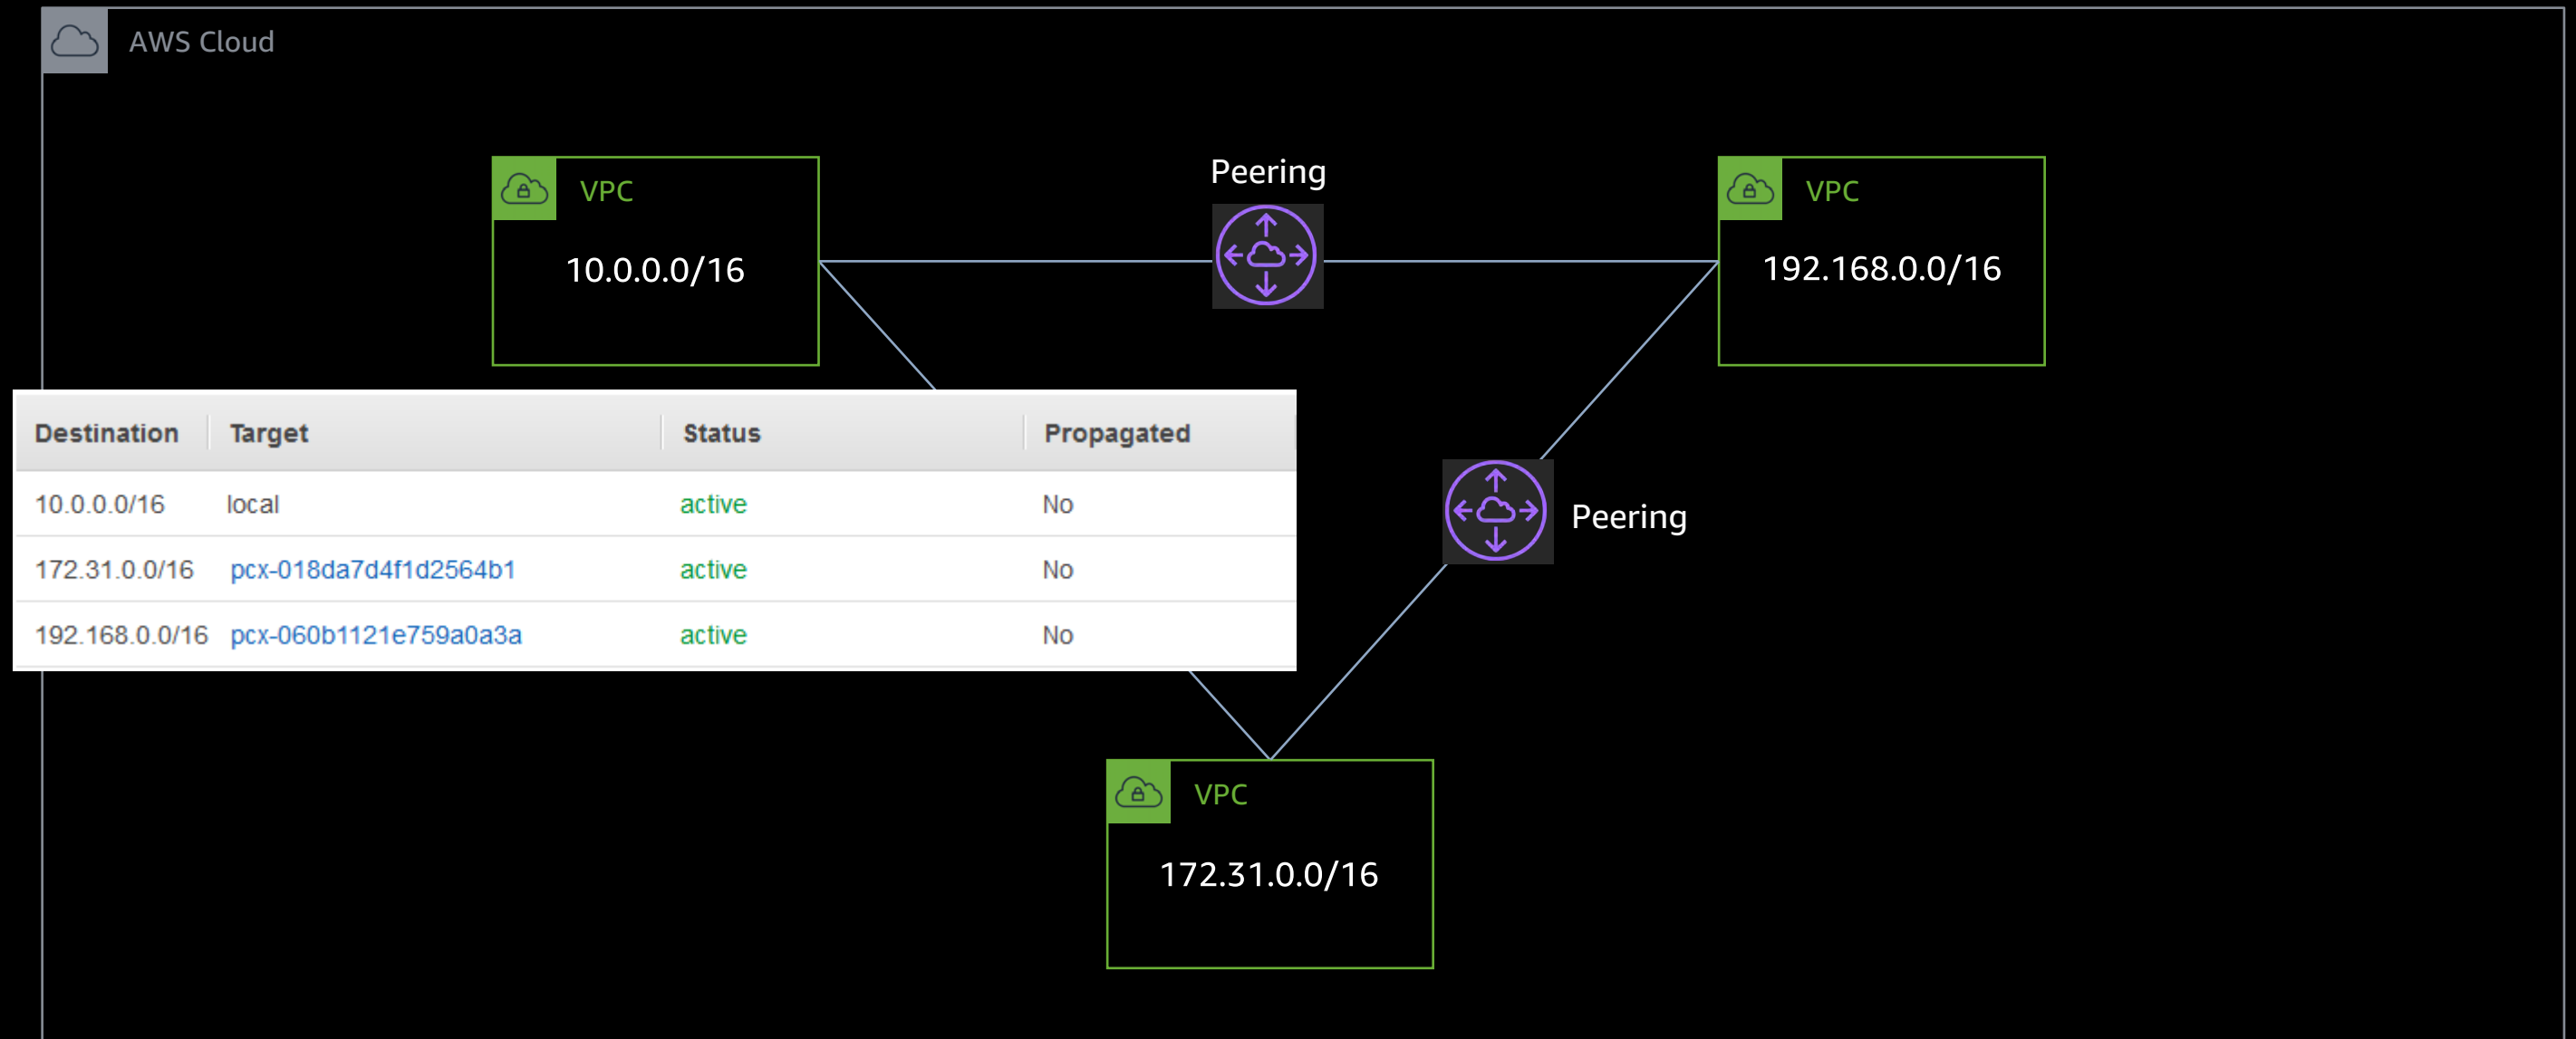
# VPC peering – same region



# VPC peering – same region

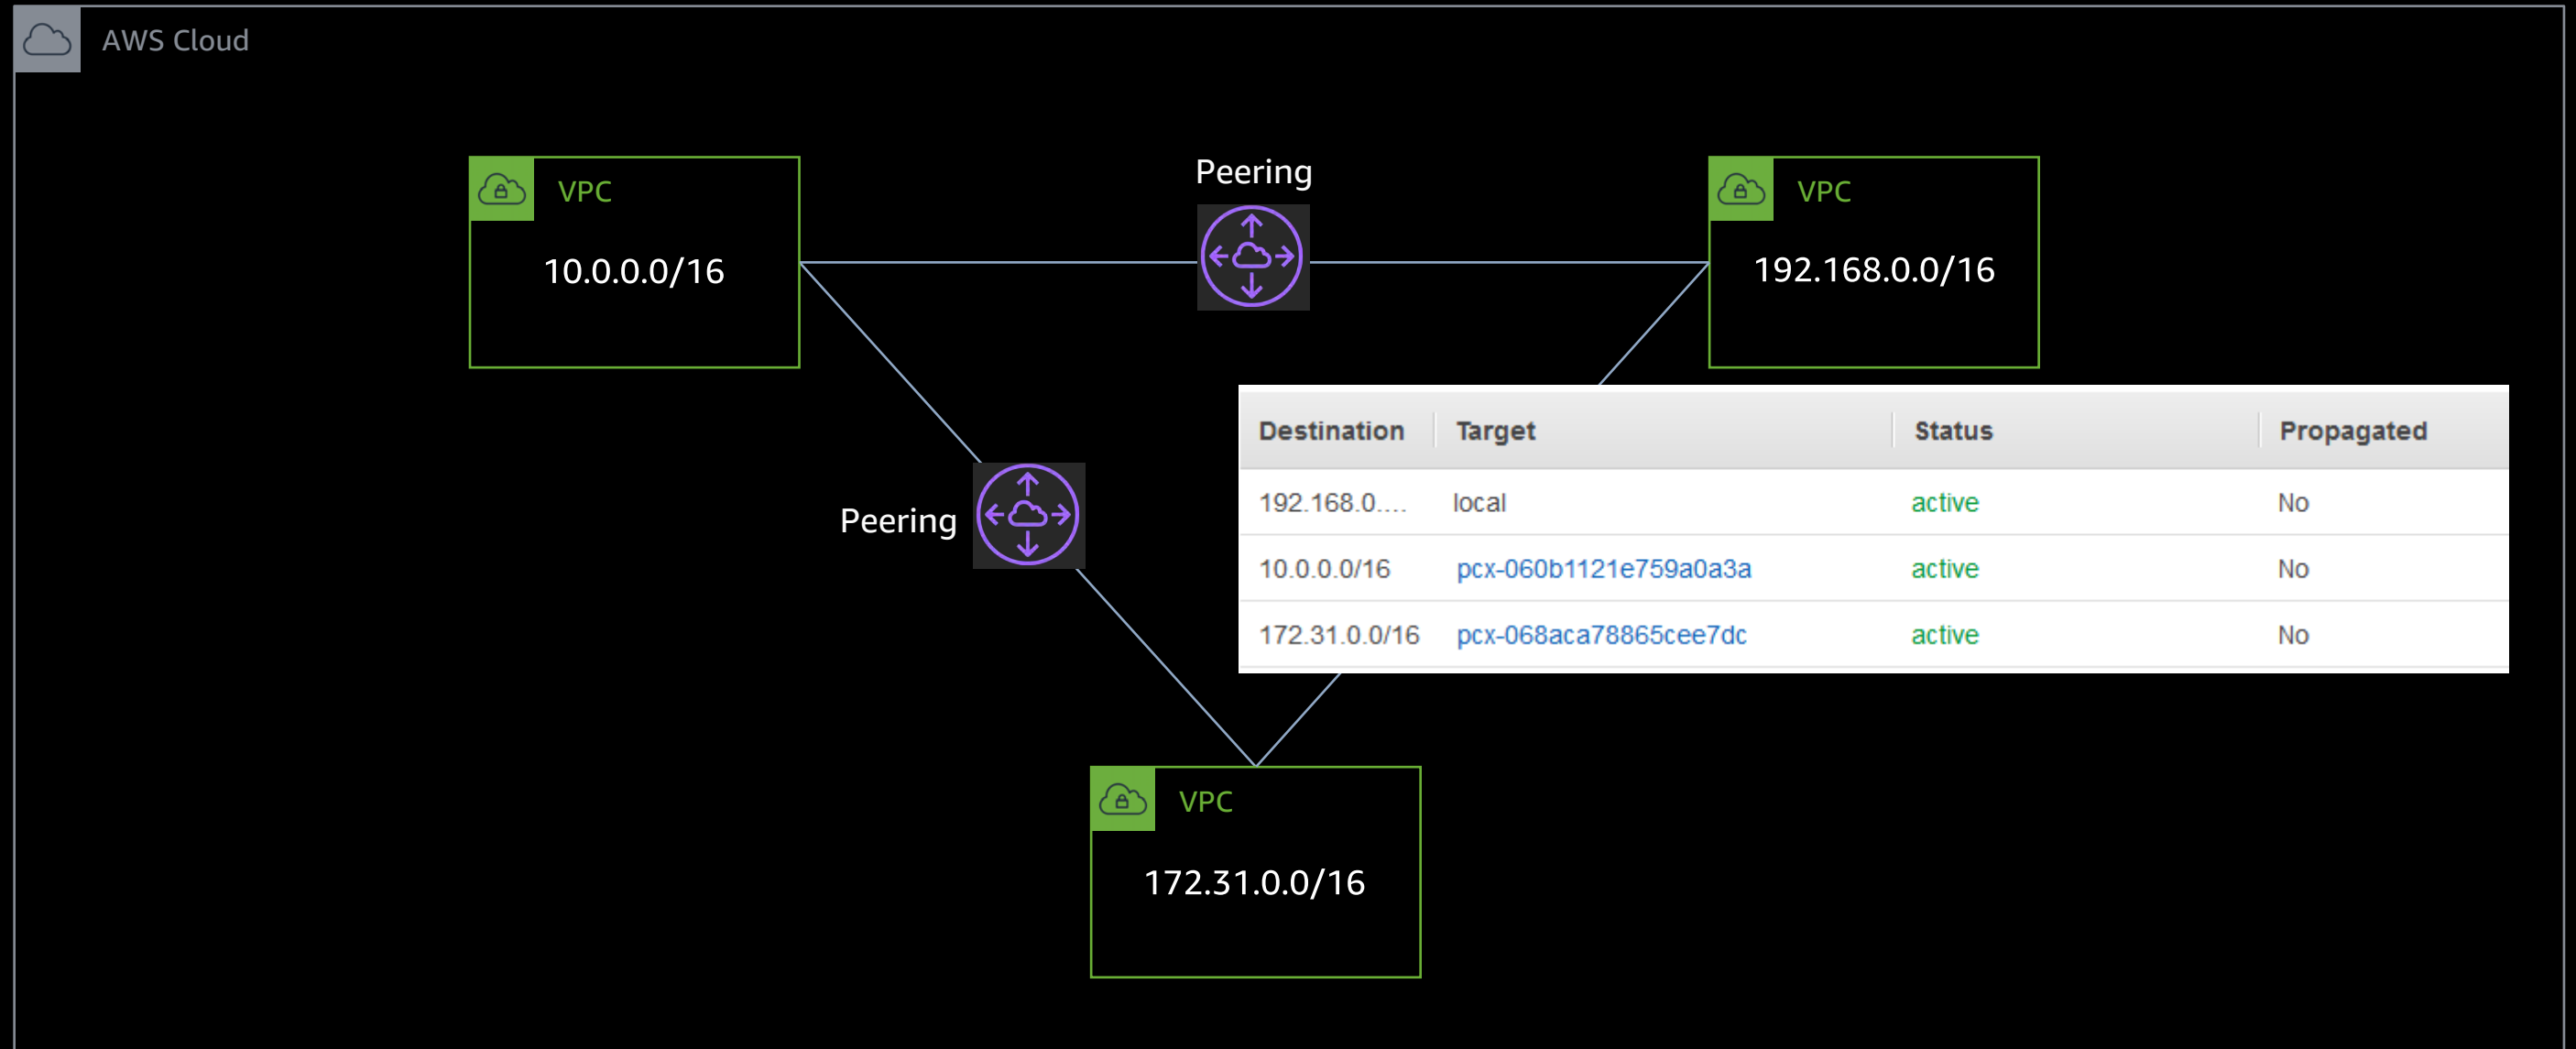


# VPC peering – same region

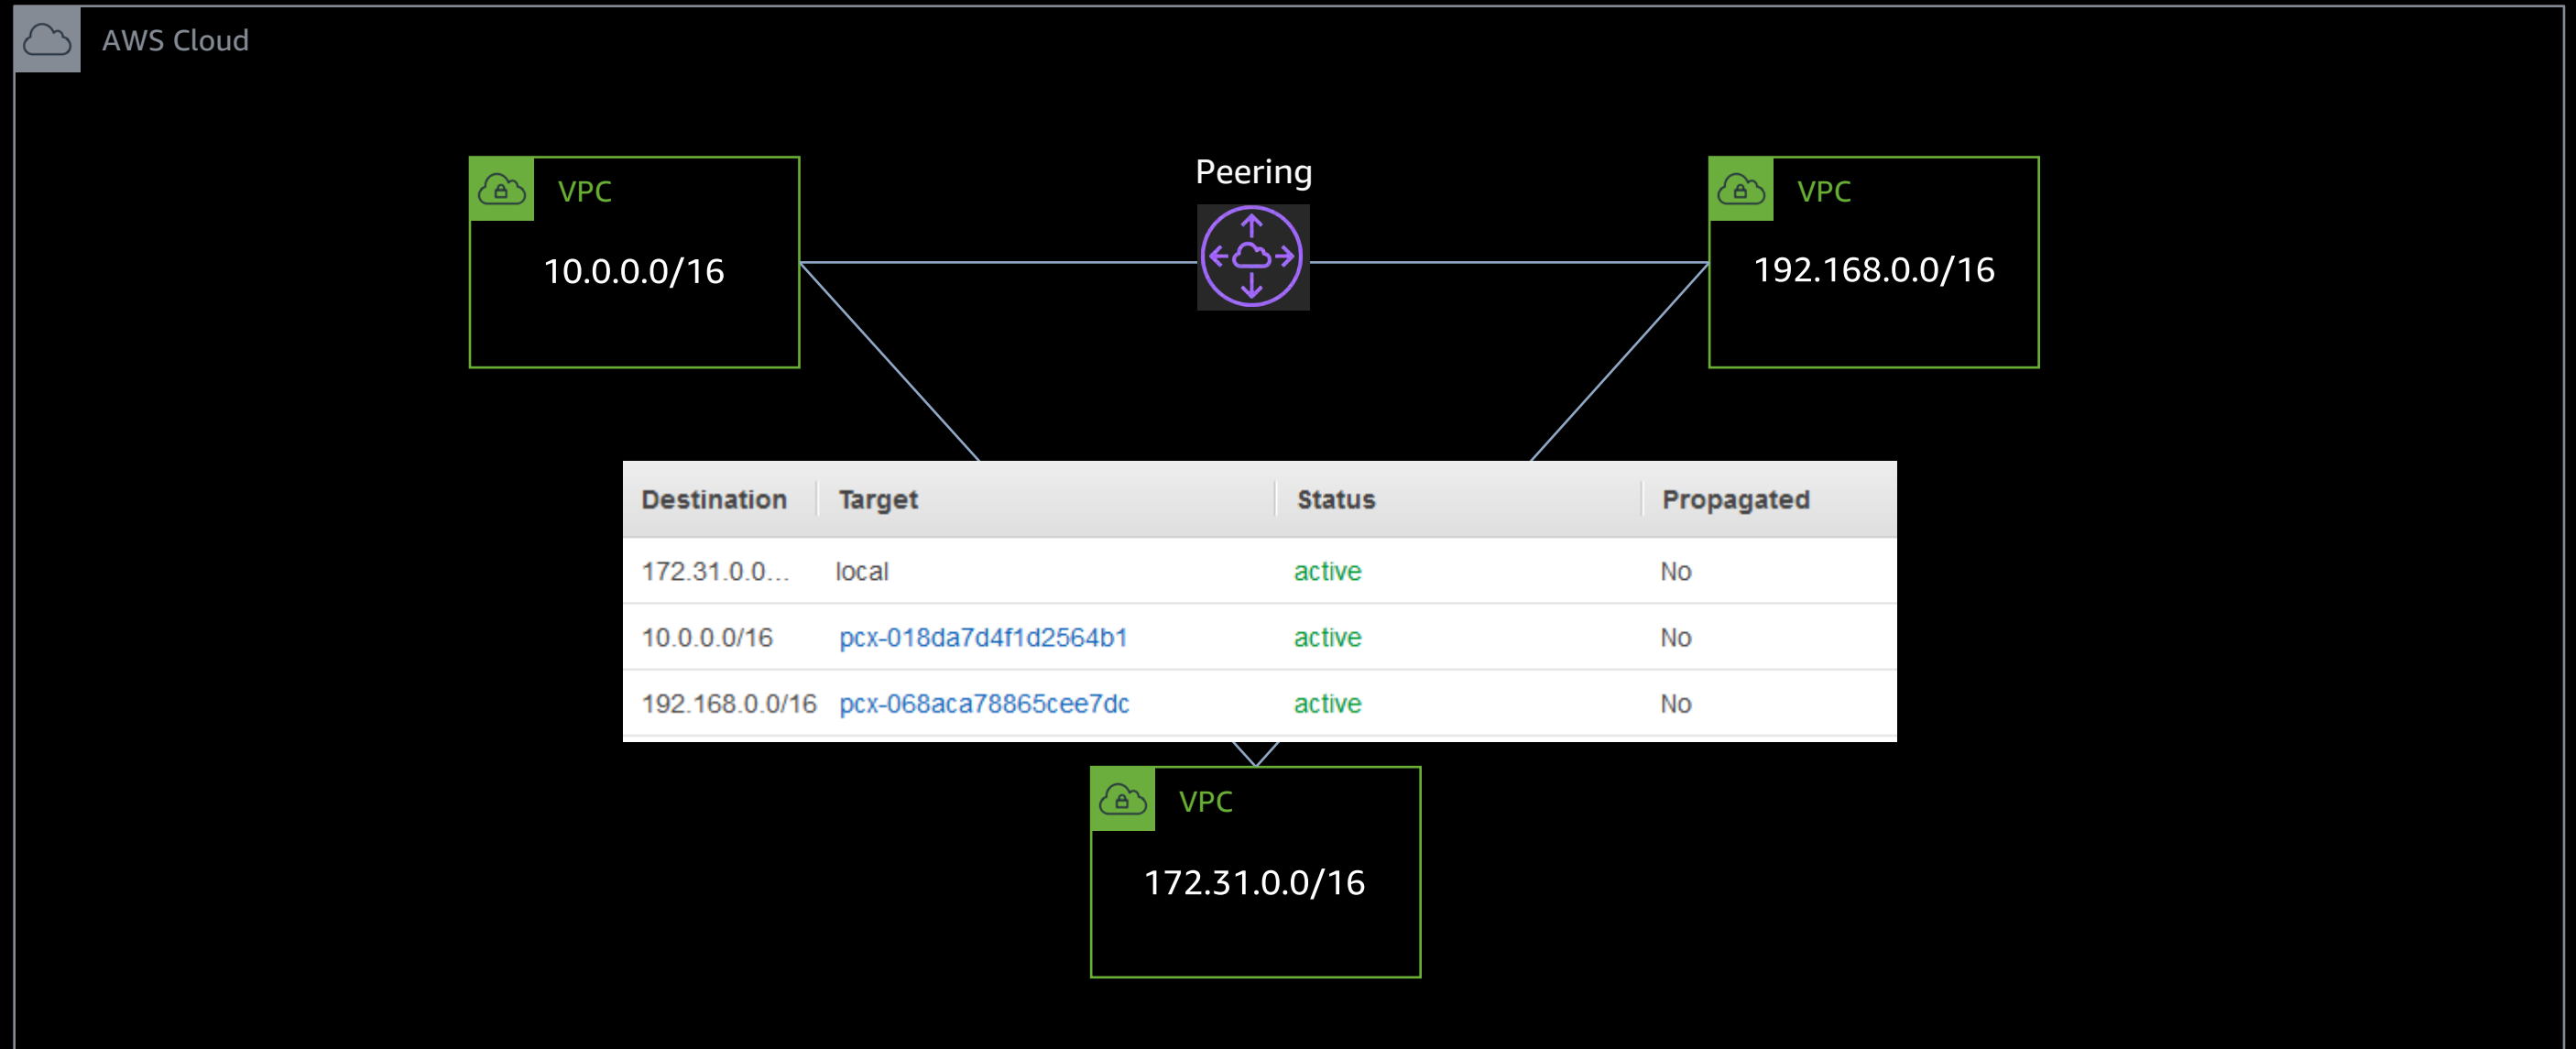




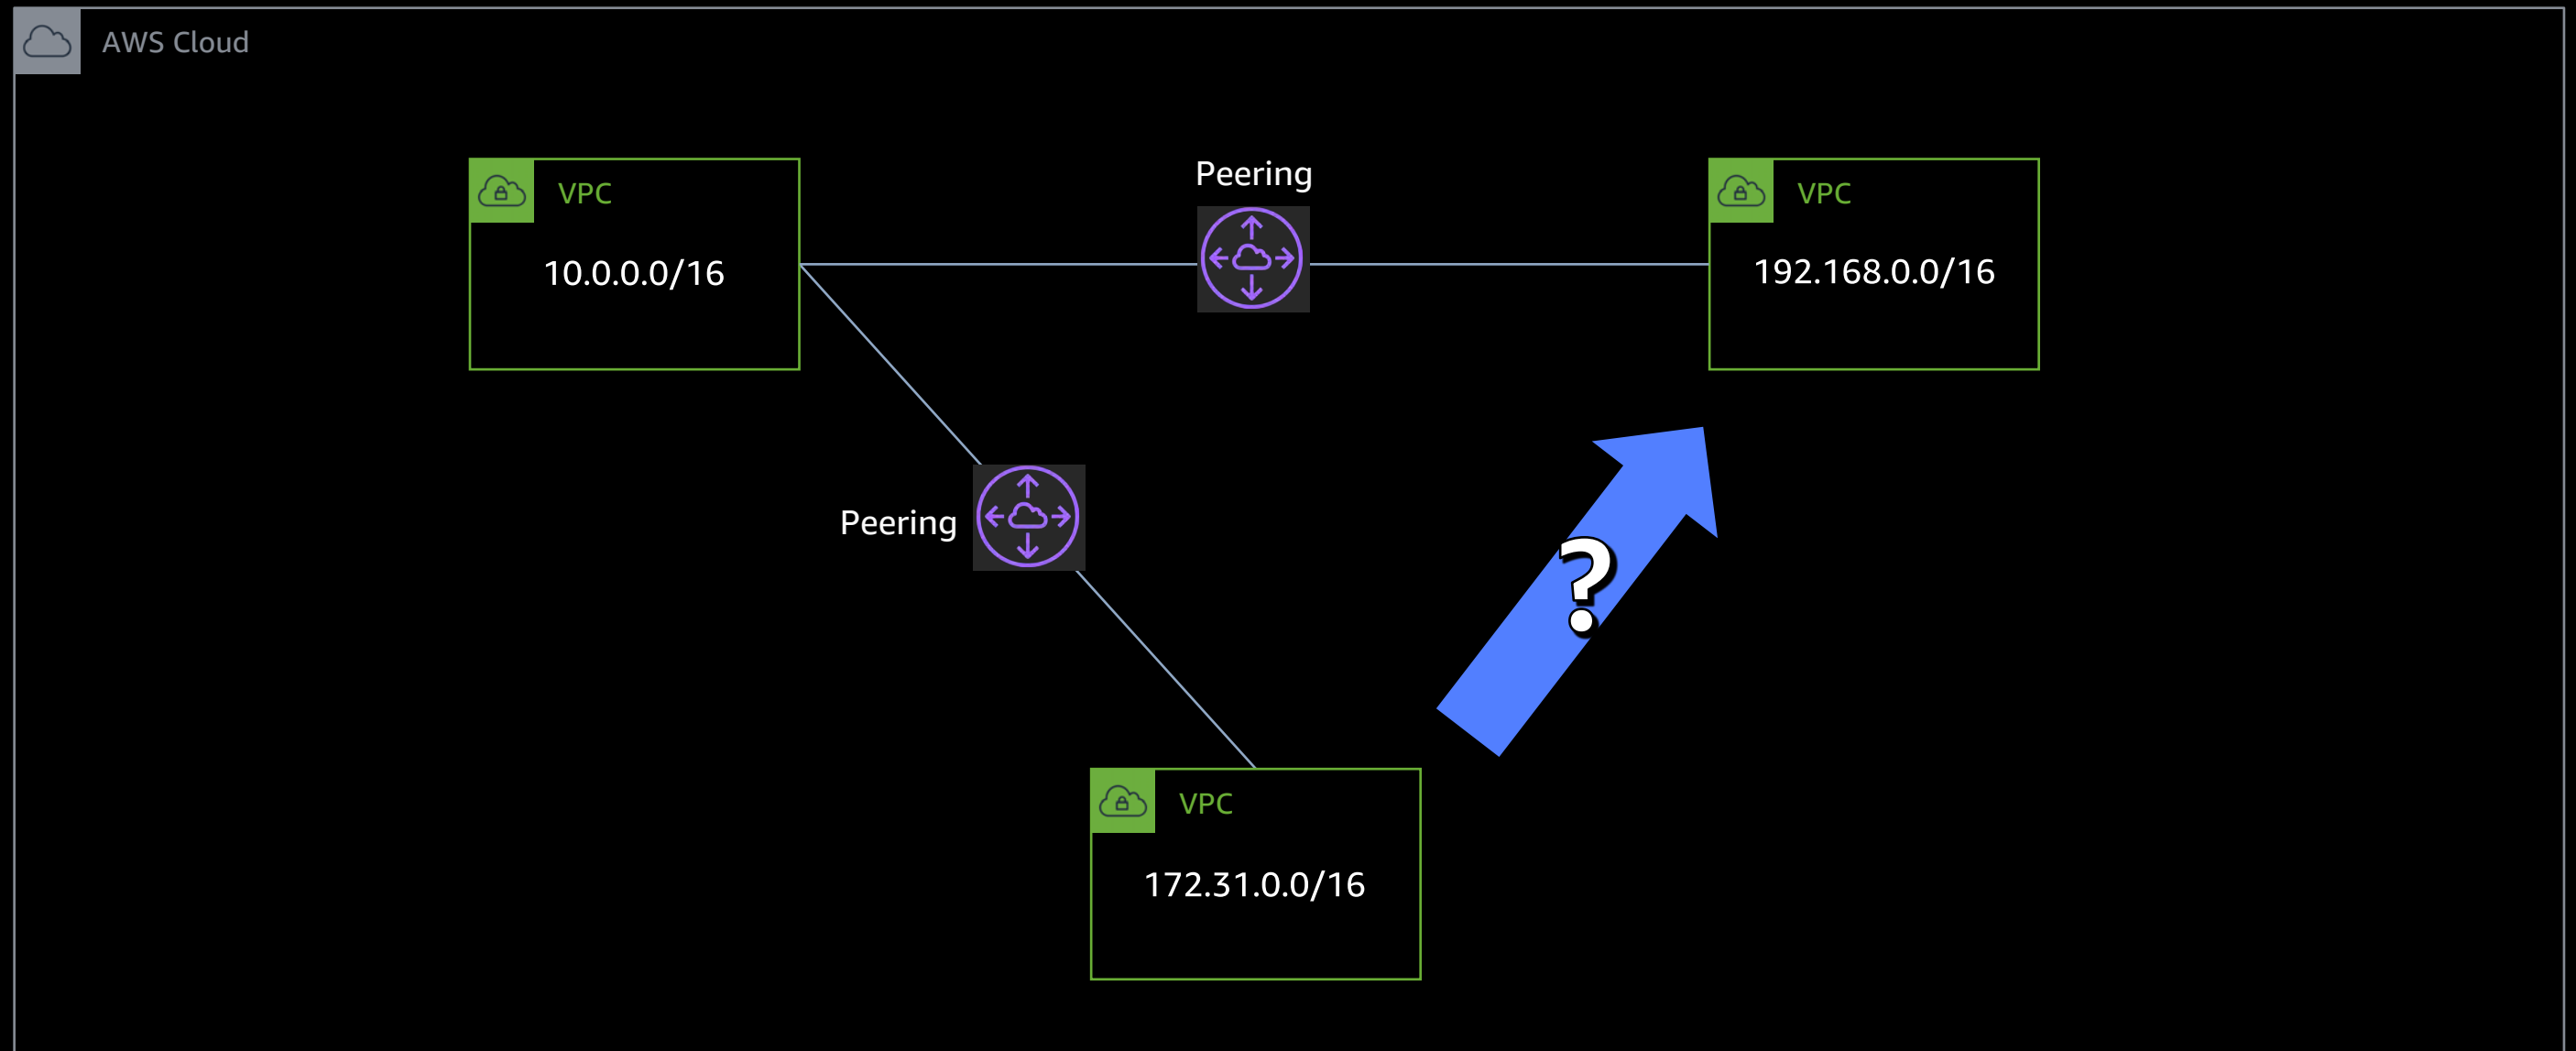
# VPC peering – same region



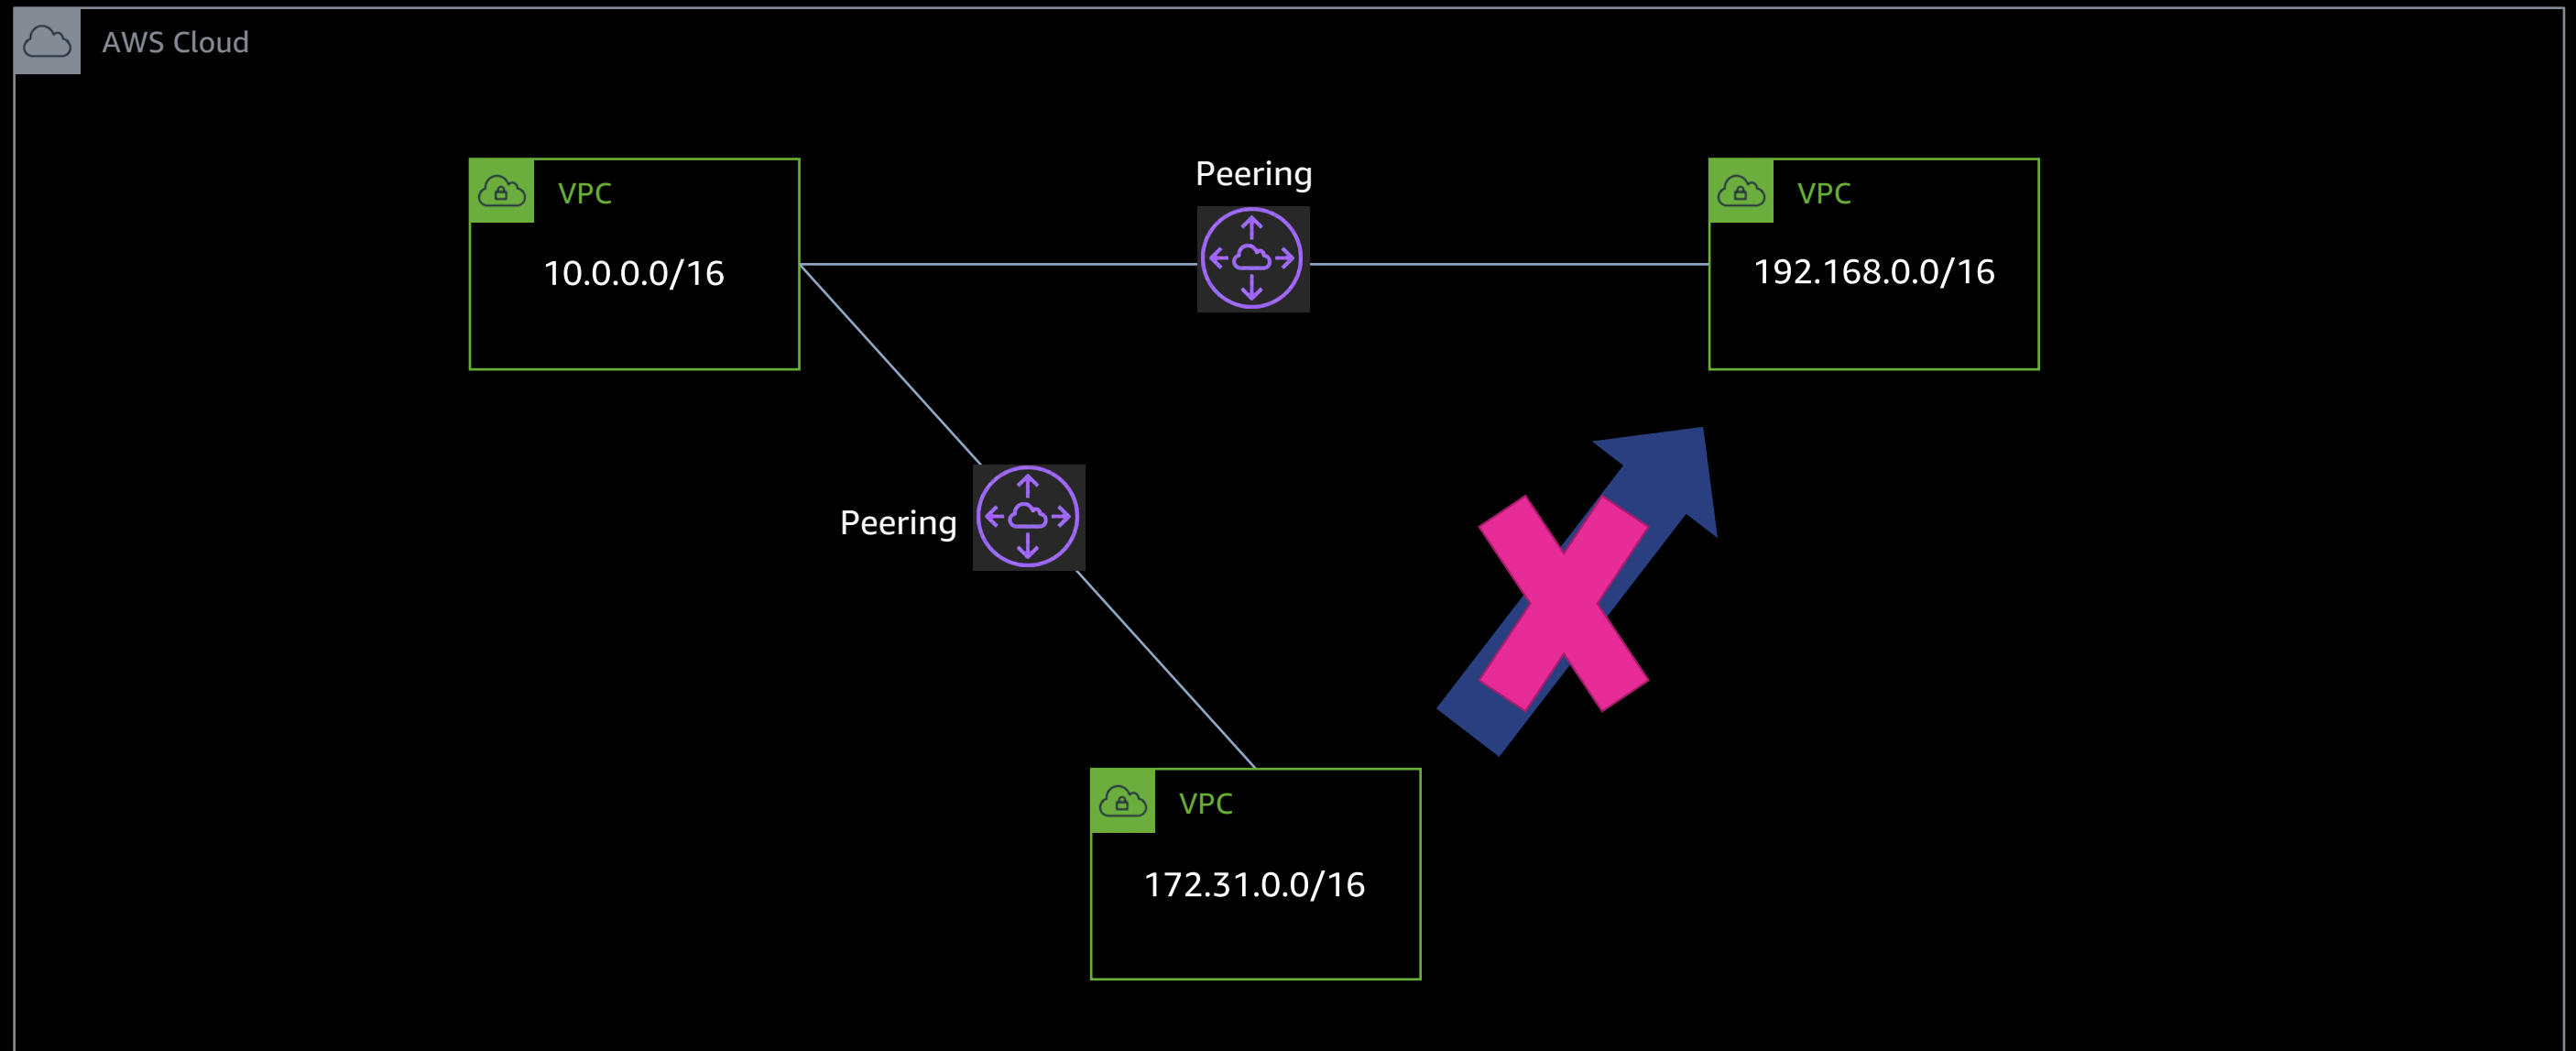
# VPC peering – same region



# VPC peering – same region



# VPC peering – same region



# VPC peering – different region

Create Peering Connection

Peering connection name tag

Select a local VPC to peer with

VPC (Requester)\*

CIDRs	CIDR	Status	Status Reason
	172.31.0.0/16	<span>●</span> associated	

Select another VPC to peer with

Account ☒ My account ☐ Another account

Region ☐ This region (us-east-1) ☒ Another Region

VPC (Acceptor)\*

\* Required

[Cancel](#) [Create Peering Connection](#)

# VPC peering – different account

Create Peering Connection

Peering connection name tag  ⓘ

Select a local VPC to peer with

VPC (Requester)\*  ↕ ↻

CIDRs	CIDR	Status	Status Reason
	172.31.0.0/16	● associated	

Select another VPC to peer with

Account ☐ My account  
☒ Another account

Account ID\*

Region ☐ This region (us-east-1)  
☒ Another Region

↕ ↻

VPC (Acceptor)\*

\* Required

Cancel Create Peering Connection

# VPC peering – things to know

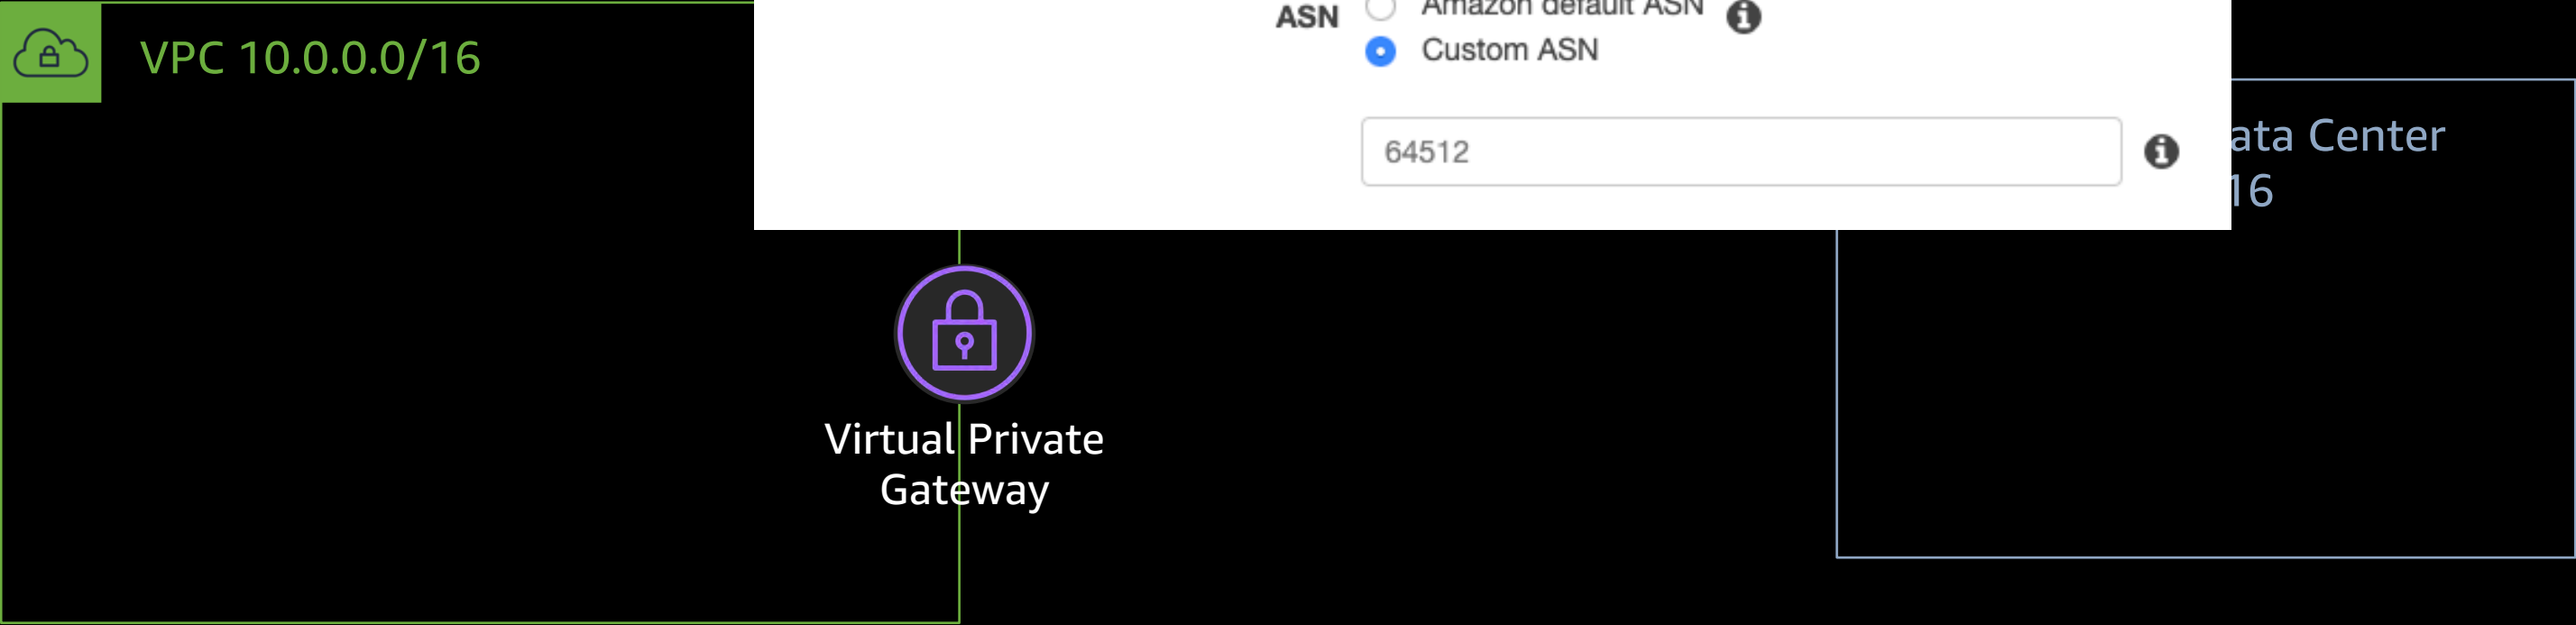
- **Can** reference security groups from the peer VPC in the same region
- **Can** enable DNS hostname resolution to return private IP addresses
- **Can** peer for both IPv4 & IPv6 addresses
- **Cannot** have overlapping IP addresses
- **Cannot** have multiple peers between the same pair of VPCs
- **Cannot** use jumbo frames across inter-region VPC peering

# Connectivity to on-premises networks





# AWS site-to-site VPN setup – VGW



### Create Virtual Private Gateway

A virtual private gateway is the router on the Amazon side of the VPN tunnel.

Name tag

myVGW

i

ASN

☐ Amazon default ASN

☒ Custom ASN

i

64512

i

# AWS site-to-site VPN – CGW

## Create Customer Gateway

Specify the Internet-routable IP address for your gateway's external interface; the address must be static and may be behind a device performing network address translation (NAT). For dynamic routing, also specify your gateway's Border Gateway Protocol (BGP) Autonomous System Number (ASN); this can be either a public or private ASN (such as those in the 64512-65534 range).

Name

myRouter



Routing



Dynamic



Static

BGP ASN\*

65000



IP Address

198.18.0.1

Certificate ARN

arn:aws:acm:us

IP Address not needed when  
Certificate is used

Corporate Data Center  
172.16.0.0/16


mer  
way

# AWS site-to-site VPN




1x VPN Connection = 2x VPN Tunnels

# AWS site-to-site VPN



1x VPN Connection = 2x VPN Tunnels

# AWS site-to-site VPN



1x VPN Connection = 2x VPN Tunnels

1x VPN Tunnel = 1.25Gbps

# AWS Direct Connect



# AWS Direct Connect

**Dedicated** network connection to AWS providing **consistent** performance and **reduced** bandwidth costs

# AWS Direct Connect – Physical connection

## Connection settings

### Name

A name to help you identify the connection.

myDirectConnect

Name must contain no more than 100 characters. Valid characters are a-z, 0-9, and - (hyphen)

### Location

The location in which your connection is located.

CoreSite NY1, New York, NY

### Sub location

Sub location within the location.


24th Floor

### Port speed

Desired bandwidth for the new connection.

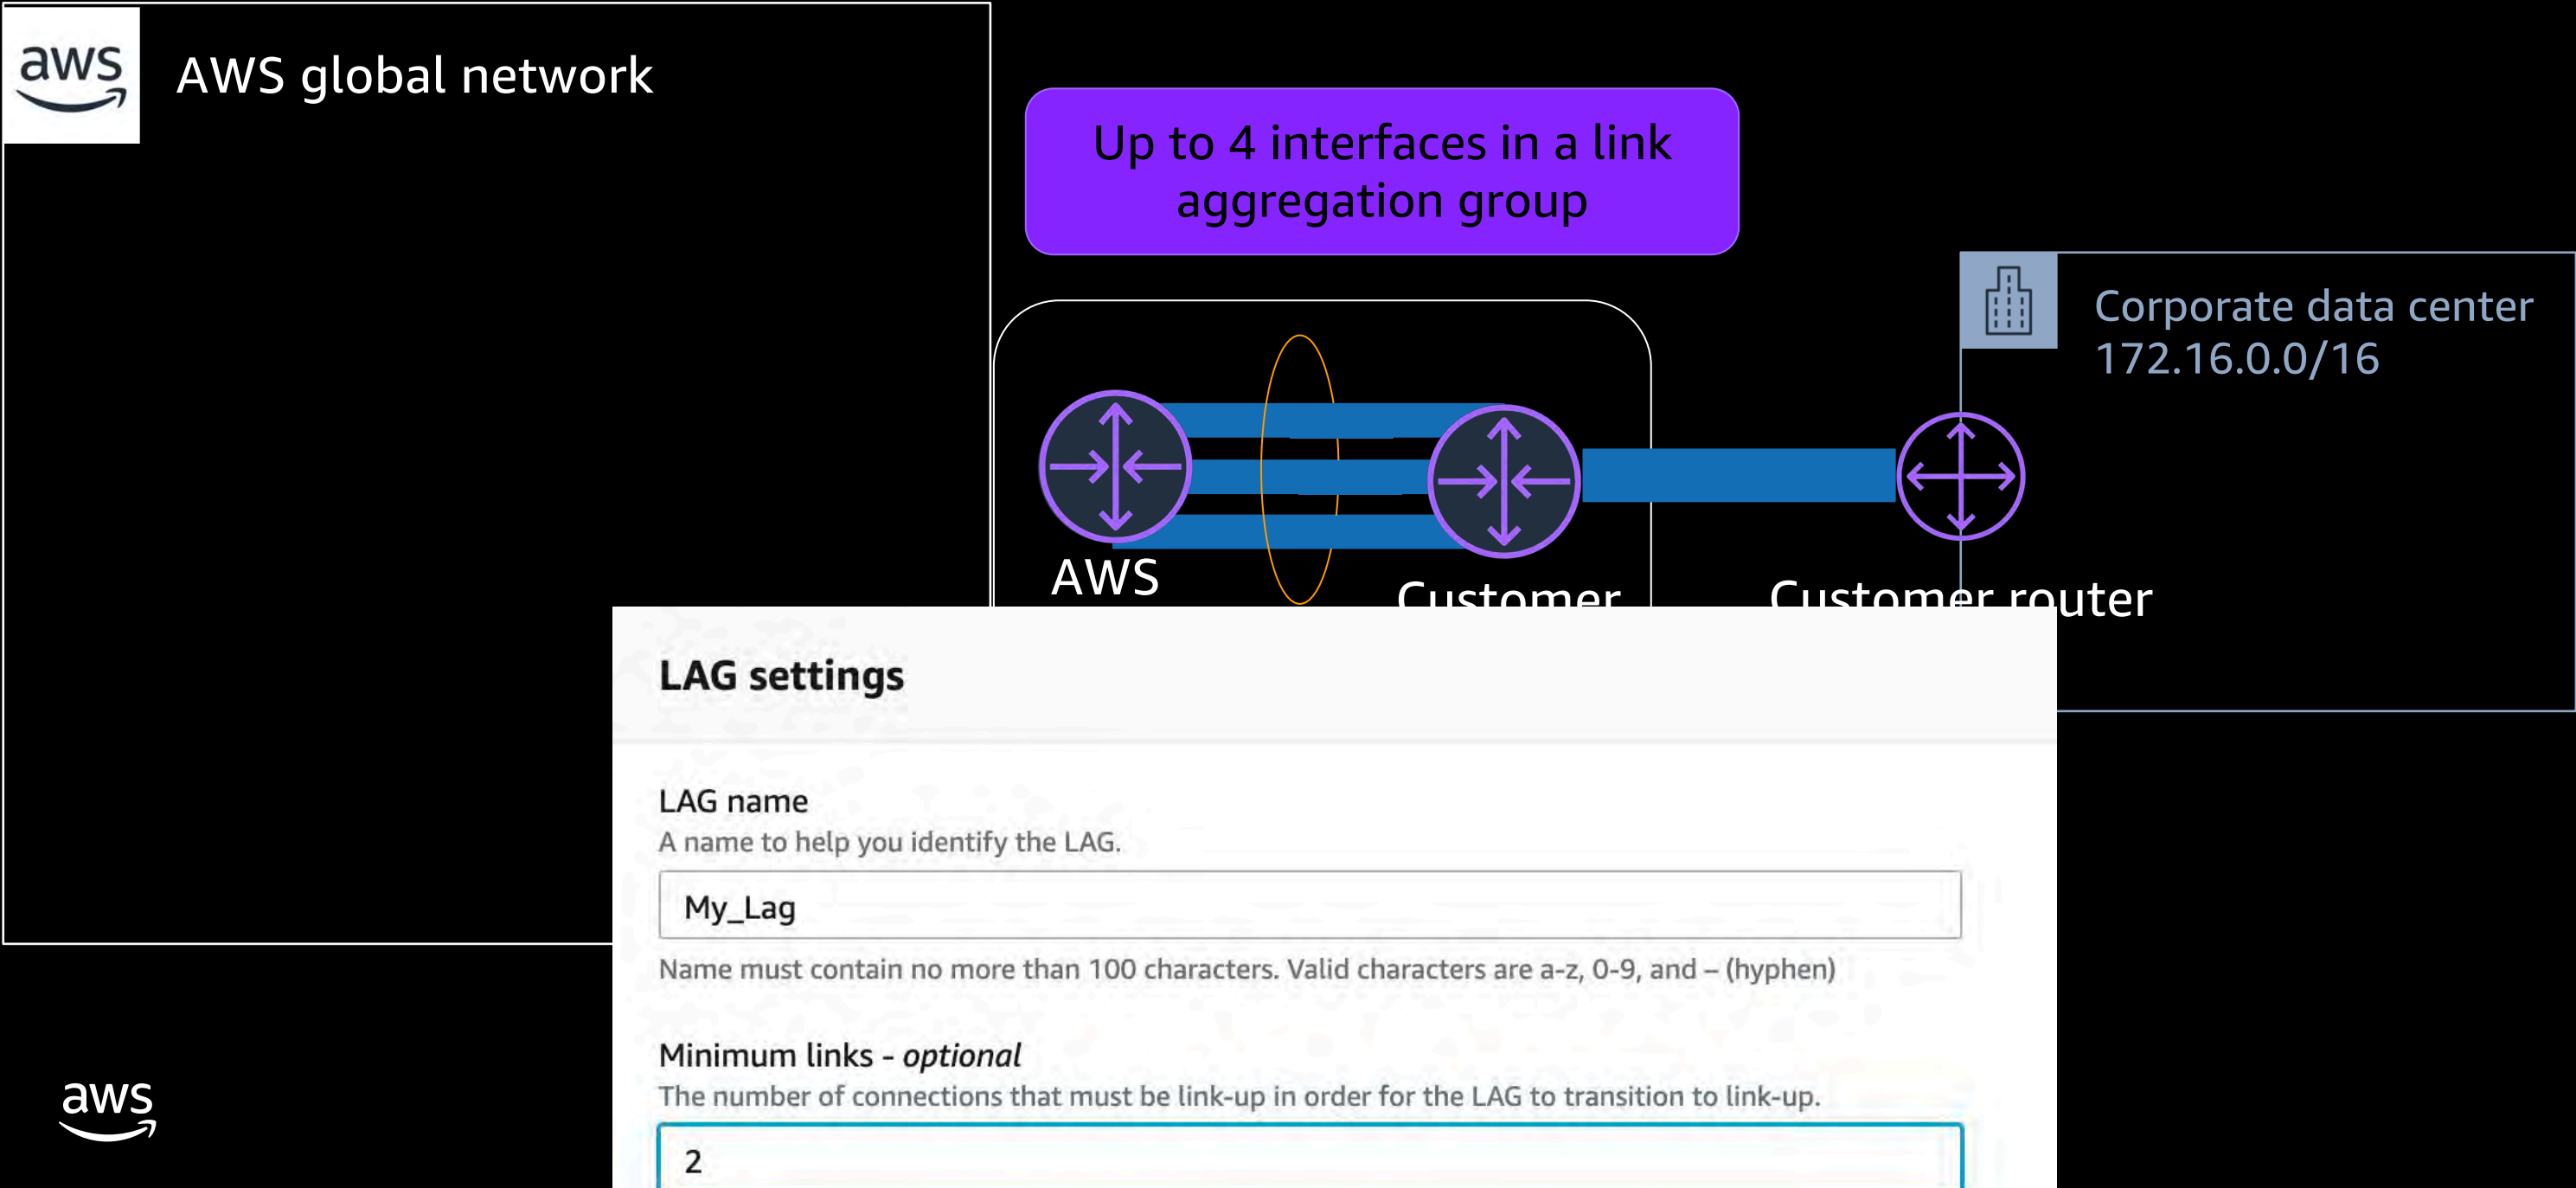
☒ 1Gbps

☐ 10Gbps





# AWS Direct Connect – Link aggregation



# AWS Direct Connect new features – Resiliency toolkit

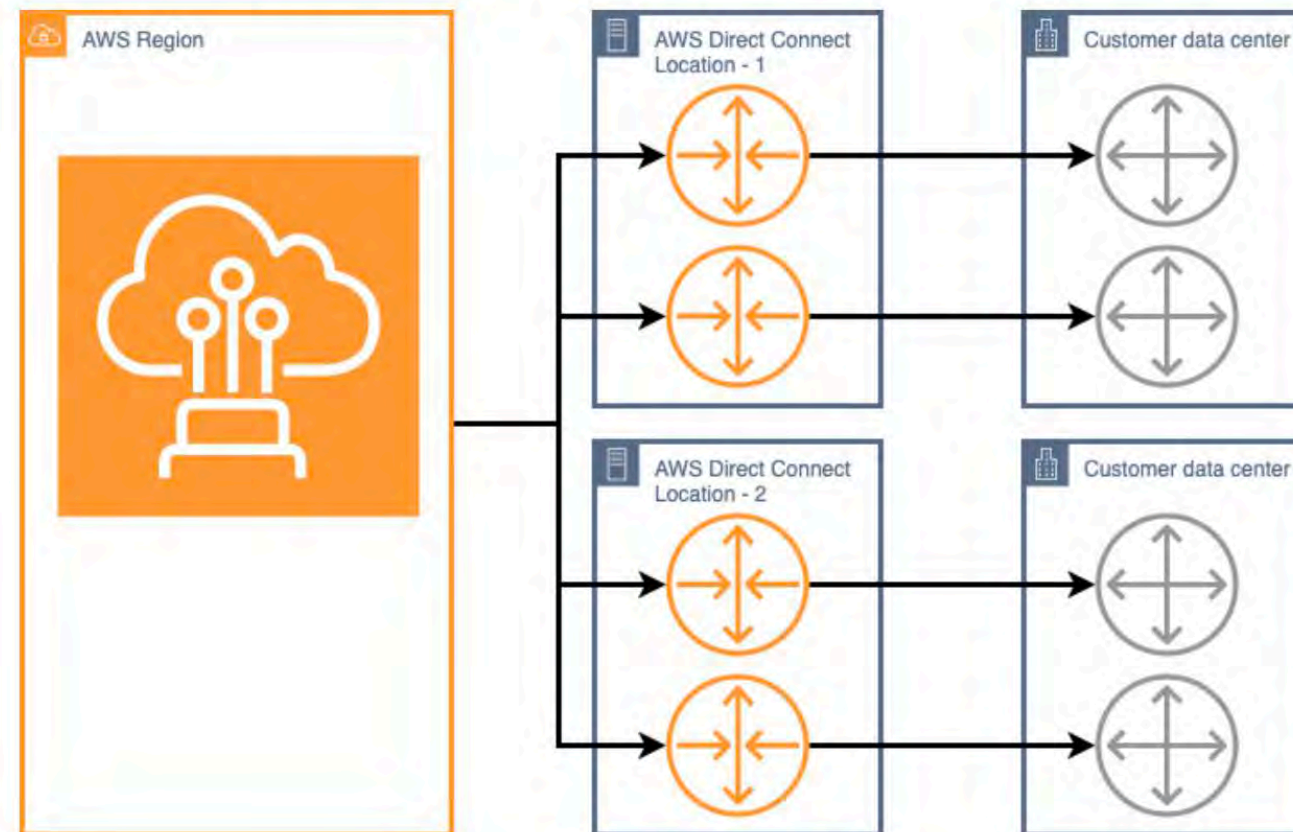
☒ **Maximum Resiliency**  
Maximum Resiliency for Critical Workloads

☐ **High Resiliency**  
High Resiliency for Critical Workloads

☐ **Development and Test**  
Non Critical Workloads or Development Workloads

## Maximum Resiliency

You can achieve maximum resiliency for critical workloads by using separate connections that terminate on separate devices in more than one location (as shown in the figure). This topology provides resiliency against device, connectivity, and complete location failures.

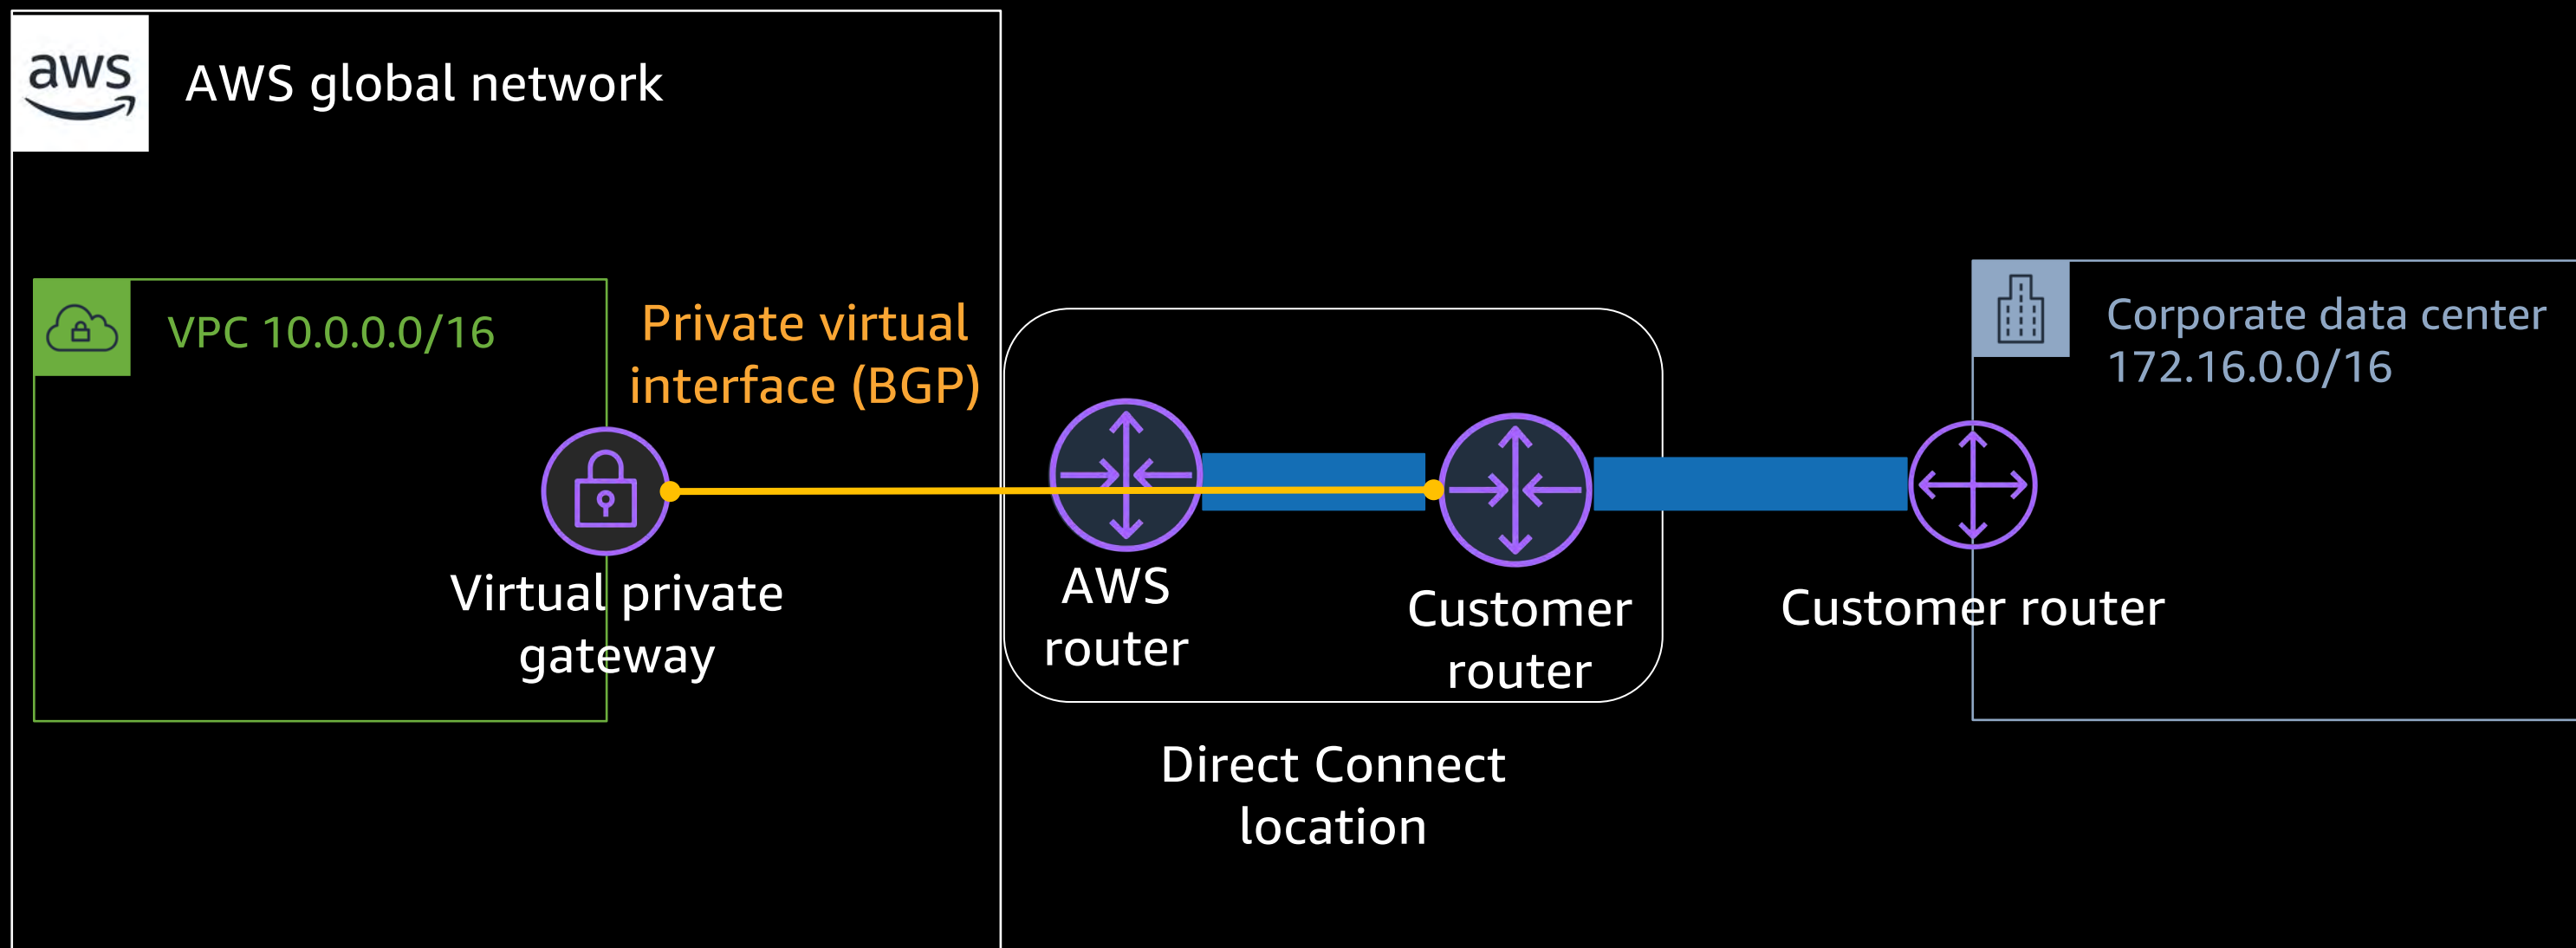


# AWS Direct Connect – Interface types

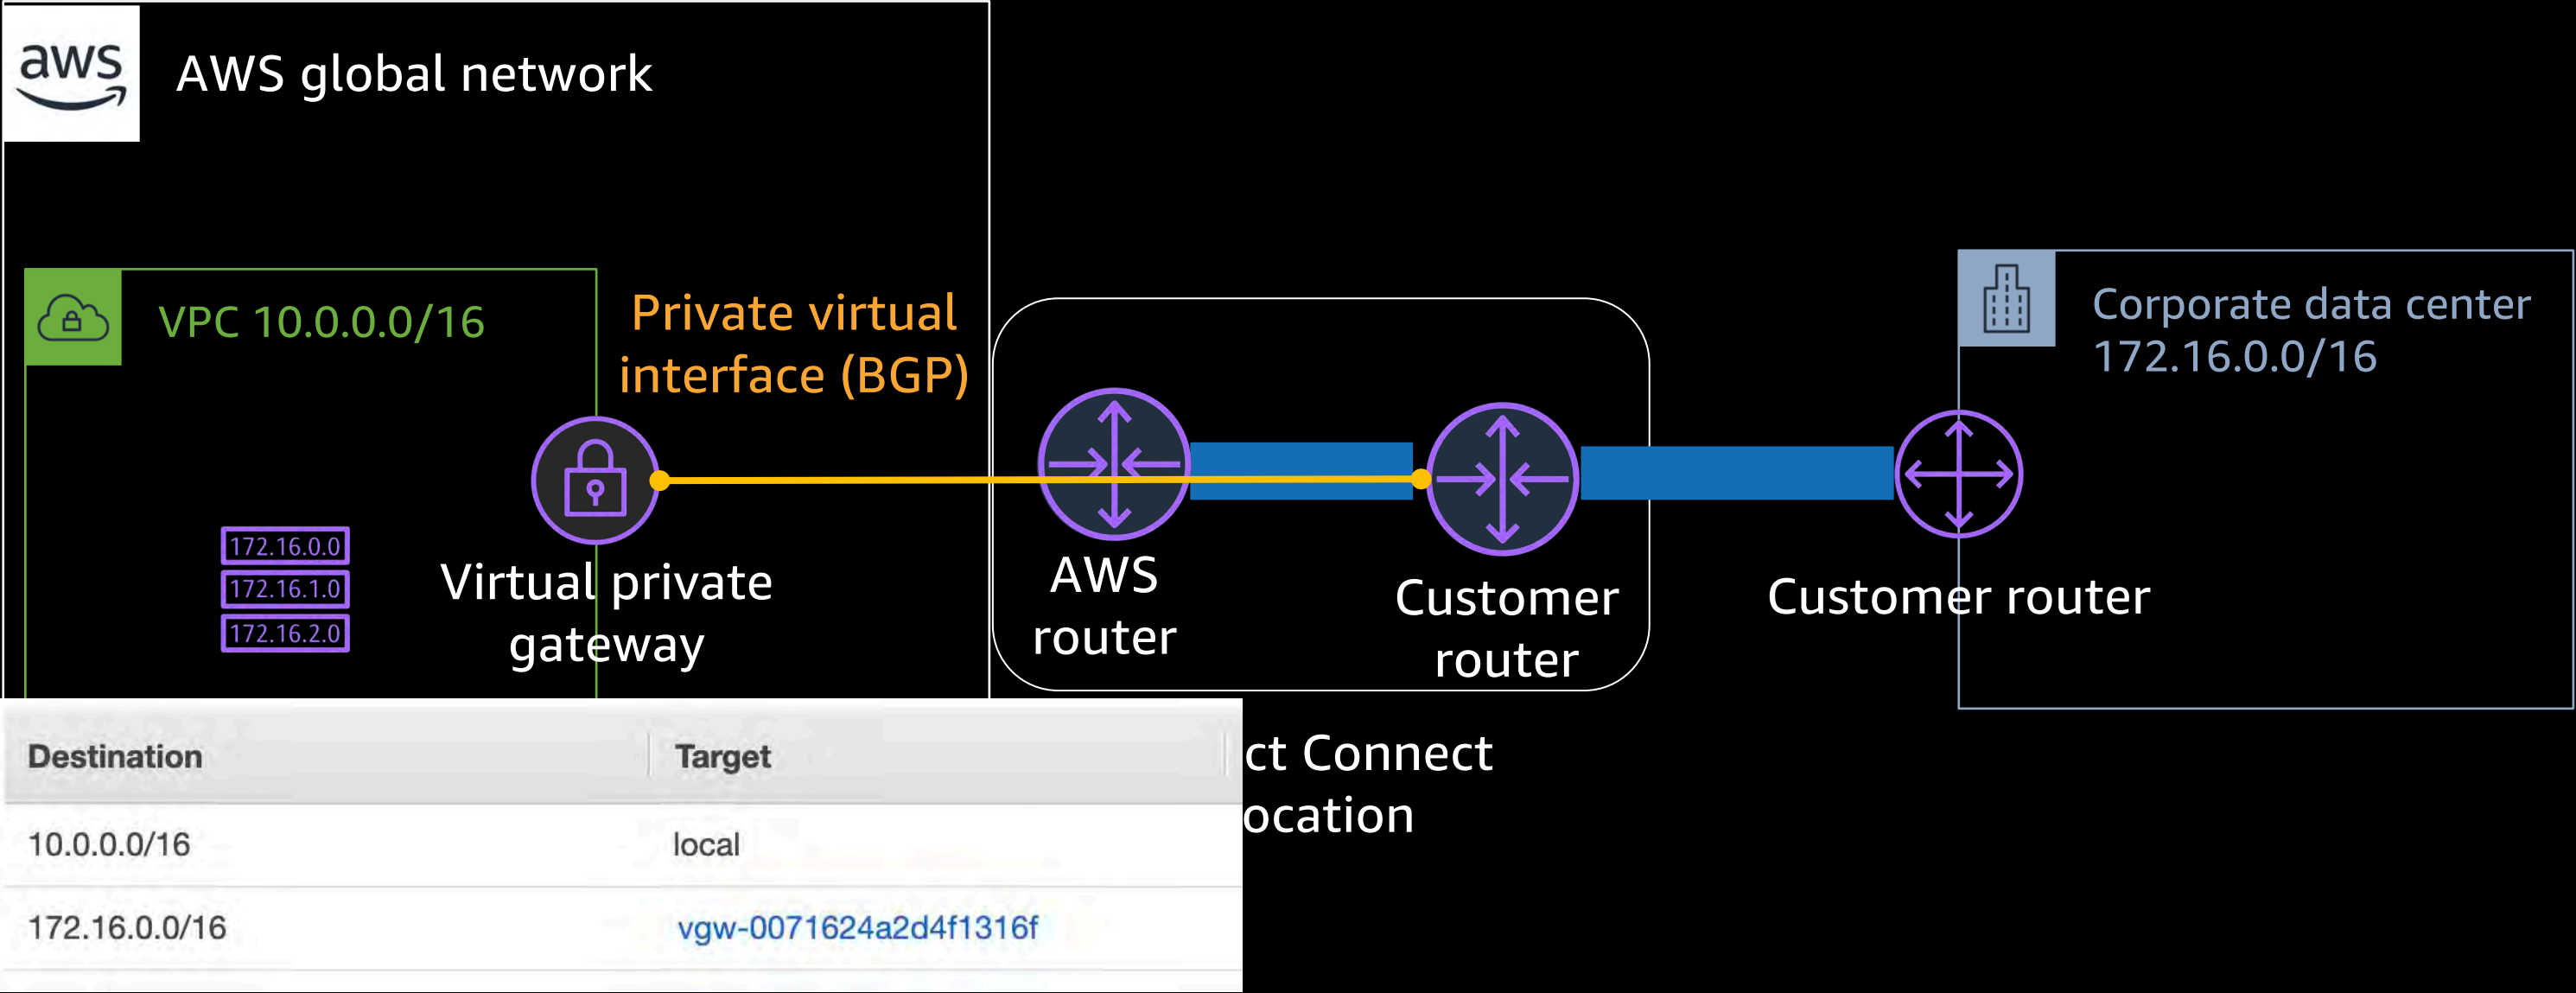
- **Private VIF** – Used to connect to Amazon VPCs using private IP addresses; directly or via Direct Connect gateway
- **Transit VIF** – Used to connect to transit gateways via Direct Connect gateway
- **Public VIF** – Used to access all AWS public services using public IP addresses

**All virtual interfaces are 802.1Q VLANs with BGP peering**

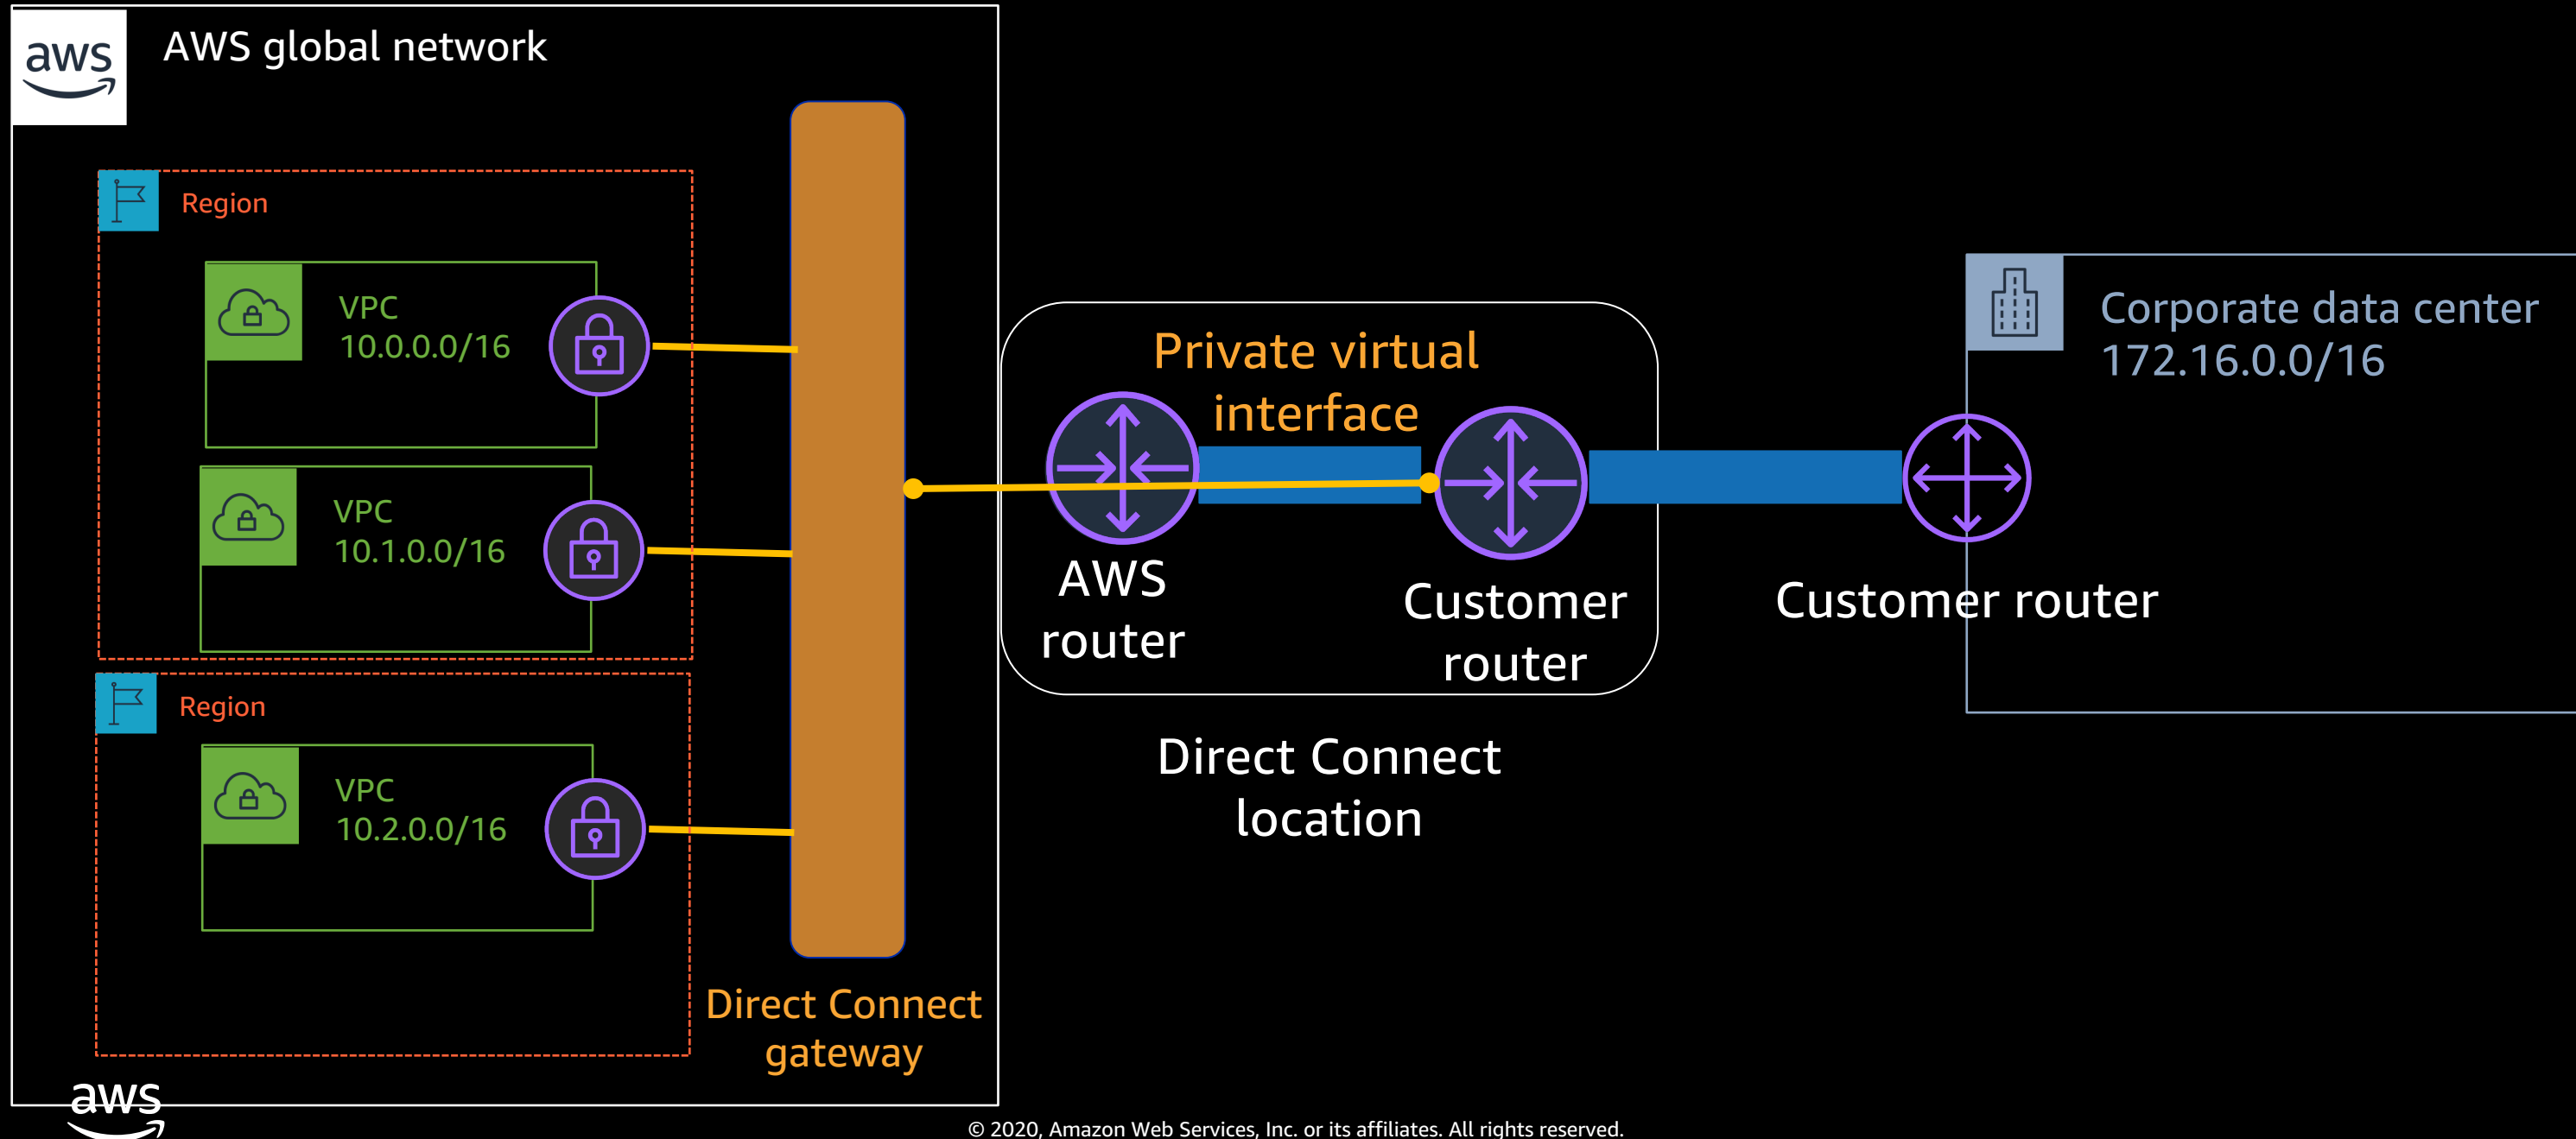
# AWS Direct Connect – Private VIF



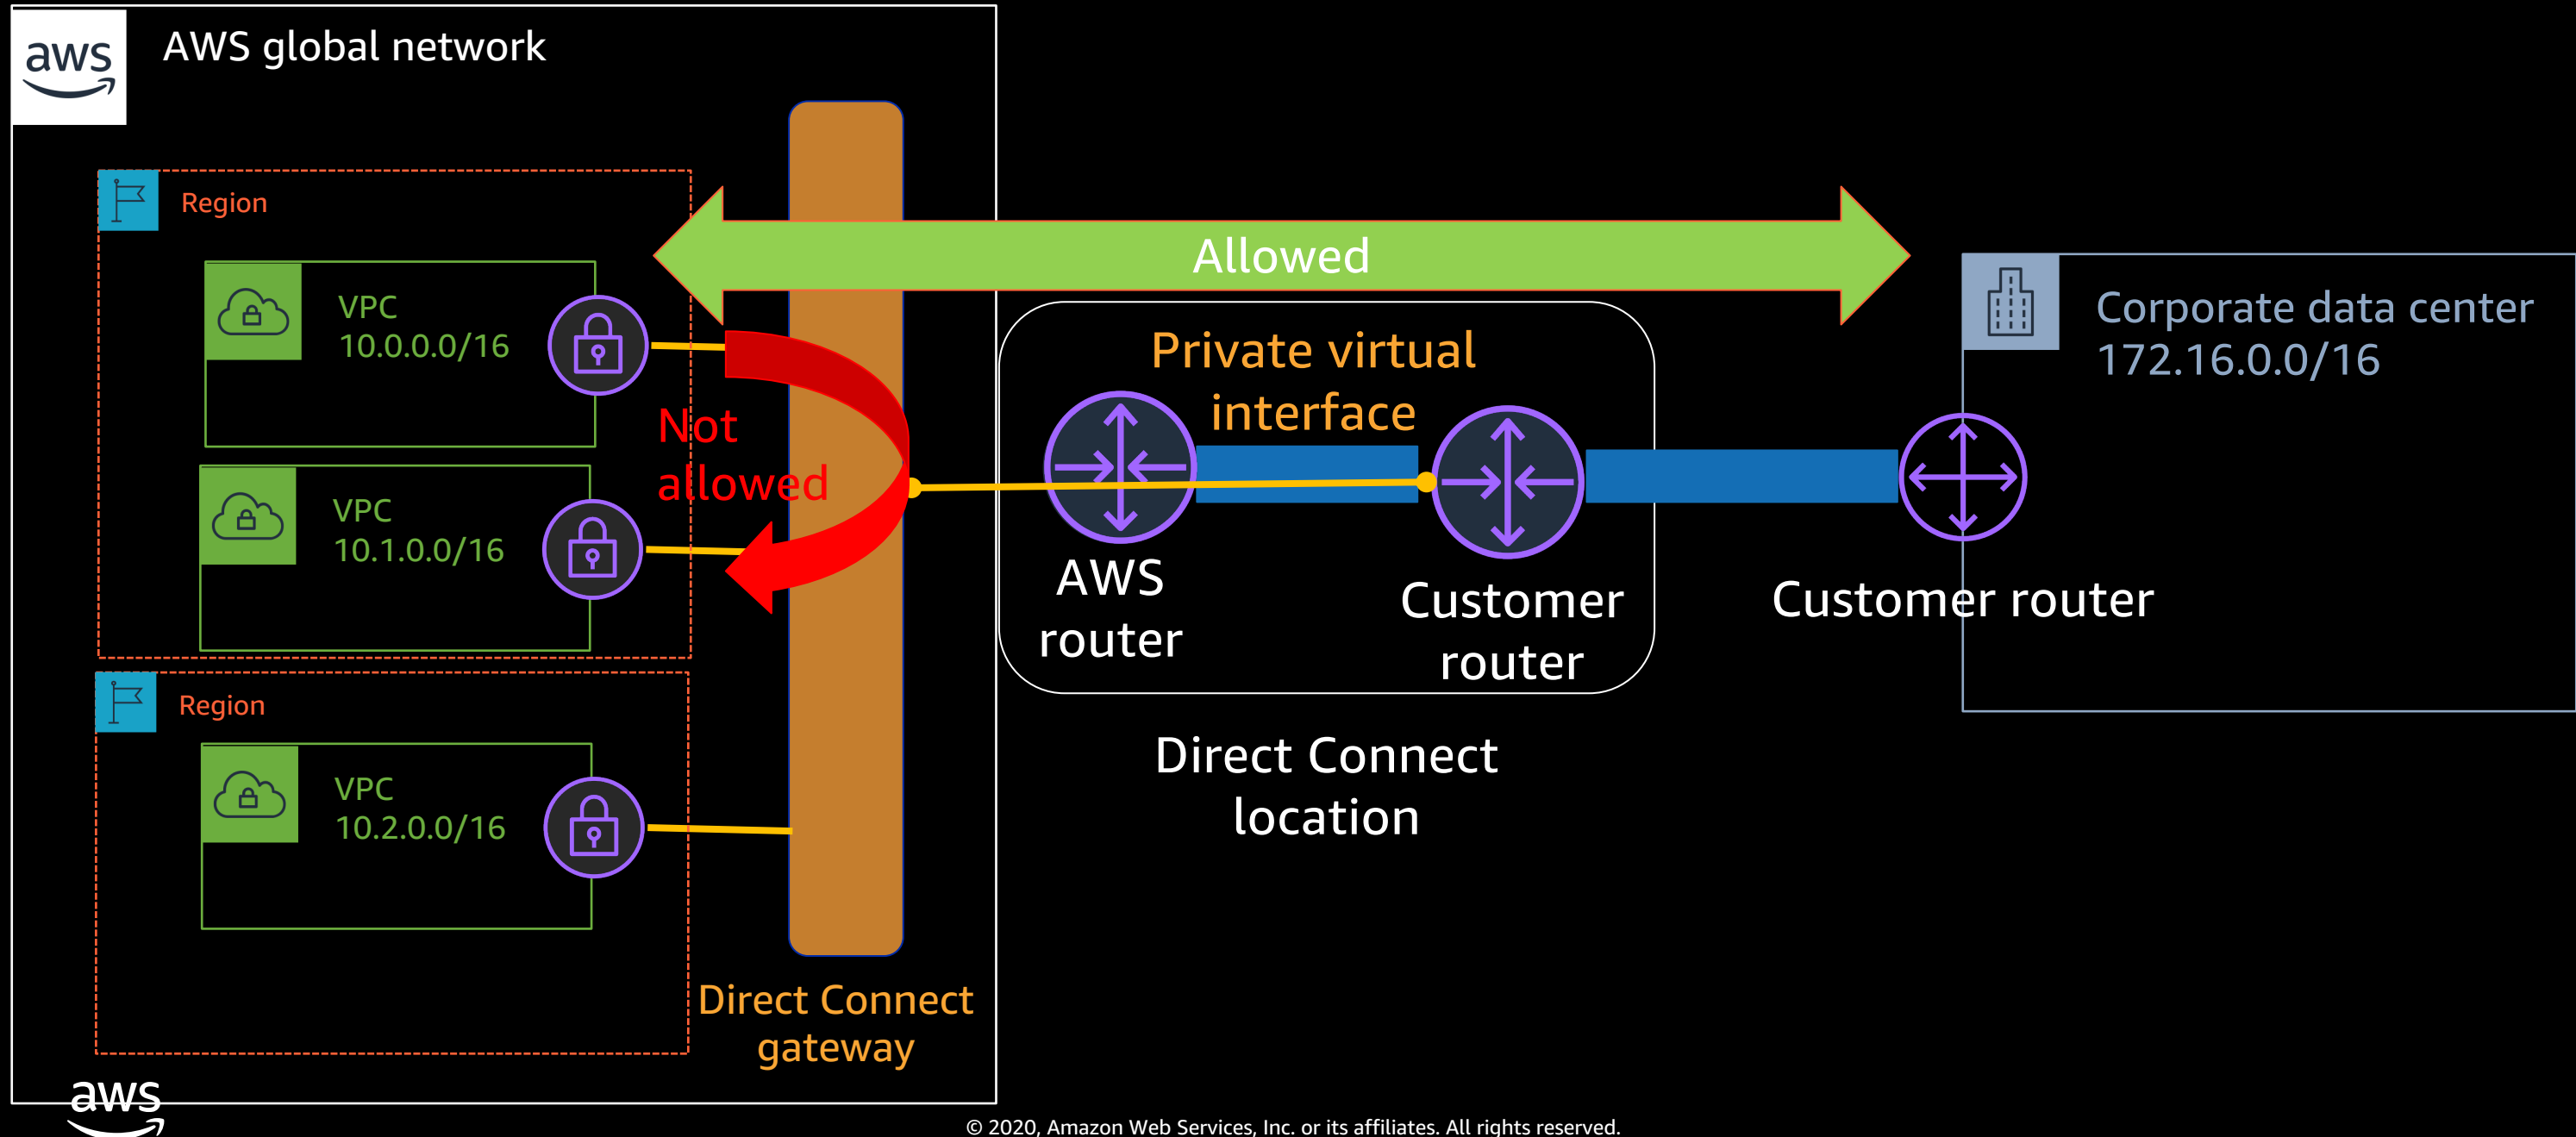
# AWS Direct Connect – Private VIF



# AWS Direct Connect gateway – Private VIF

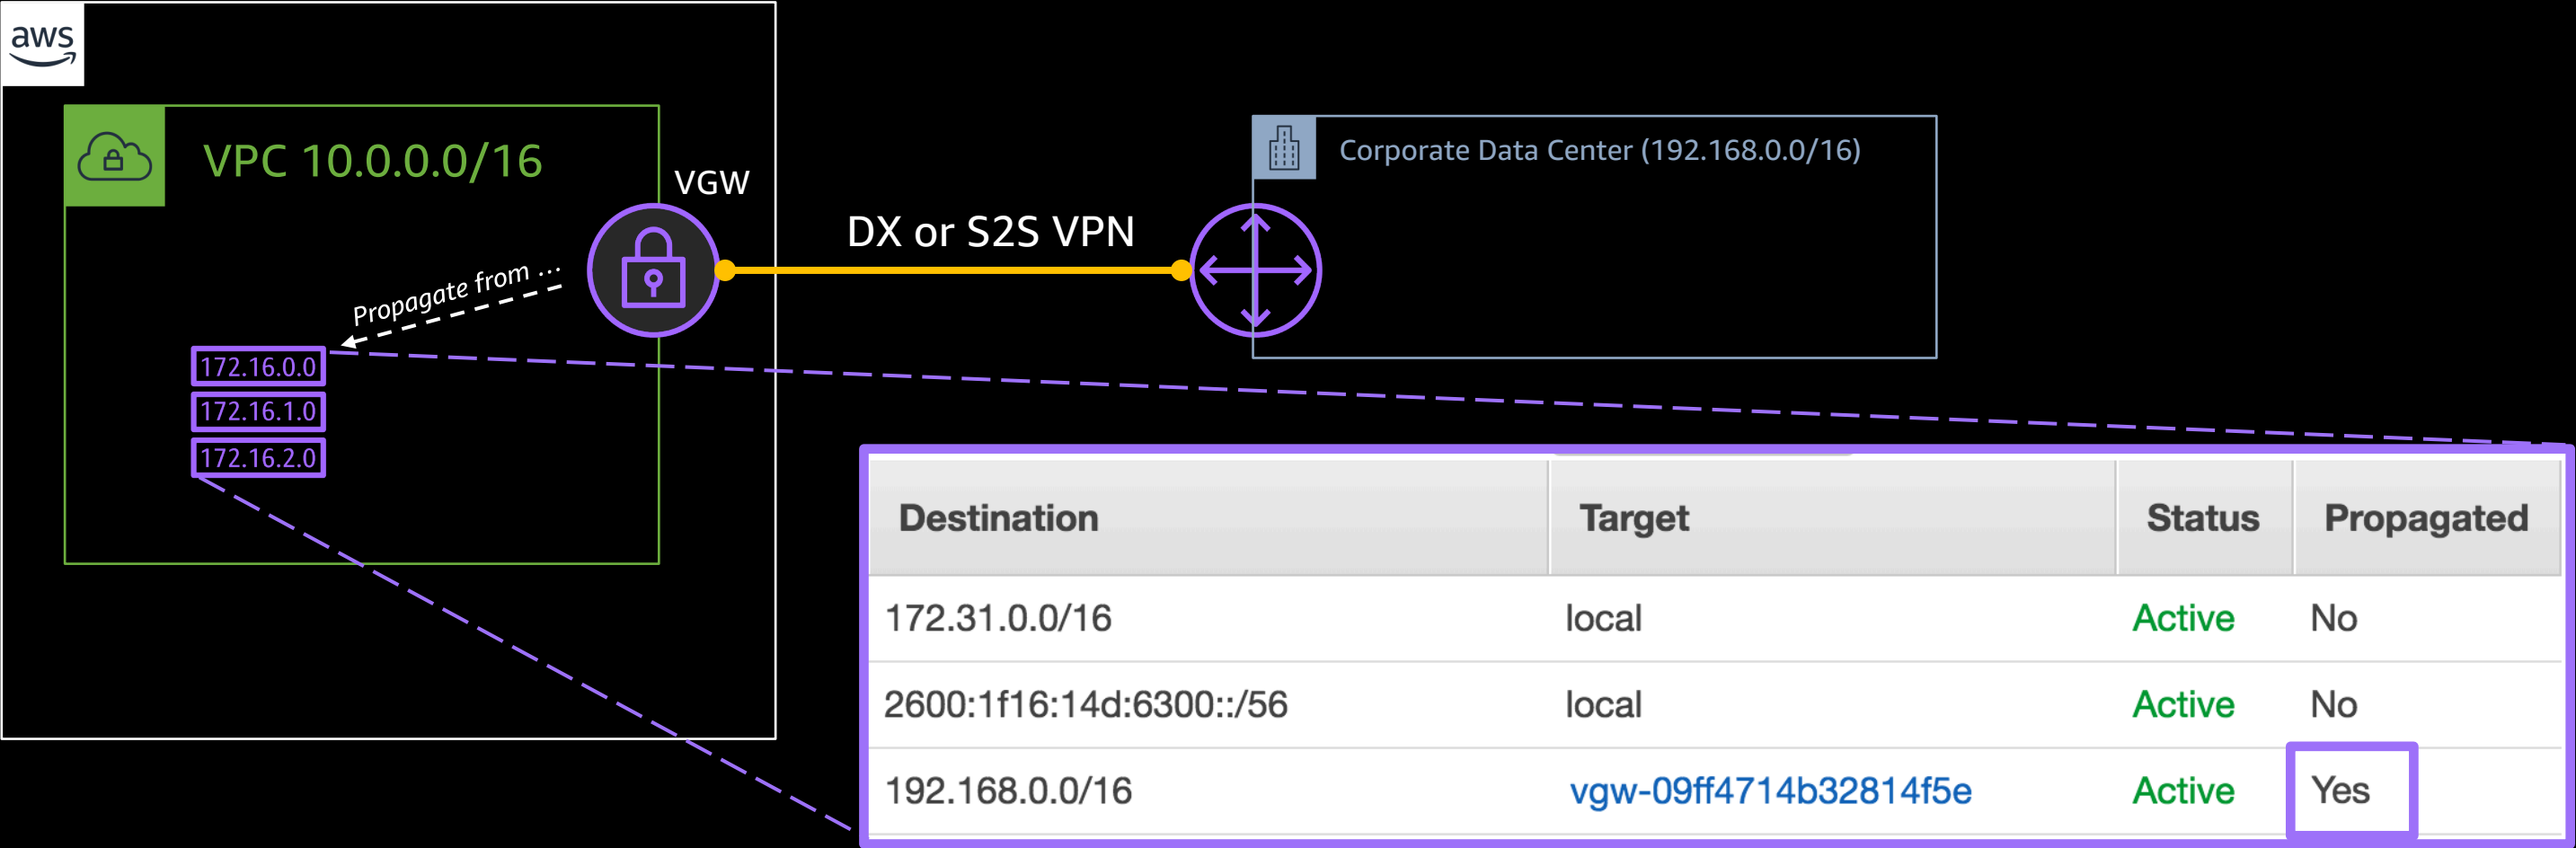


# AWS Direct Connect Gateway – Traffic flow



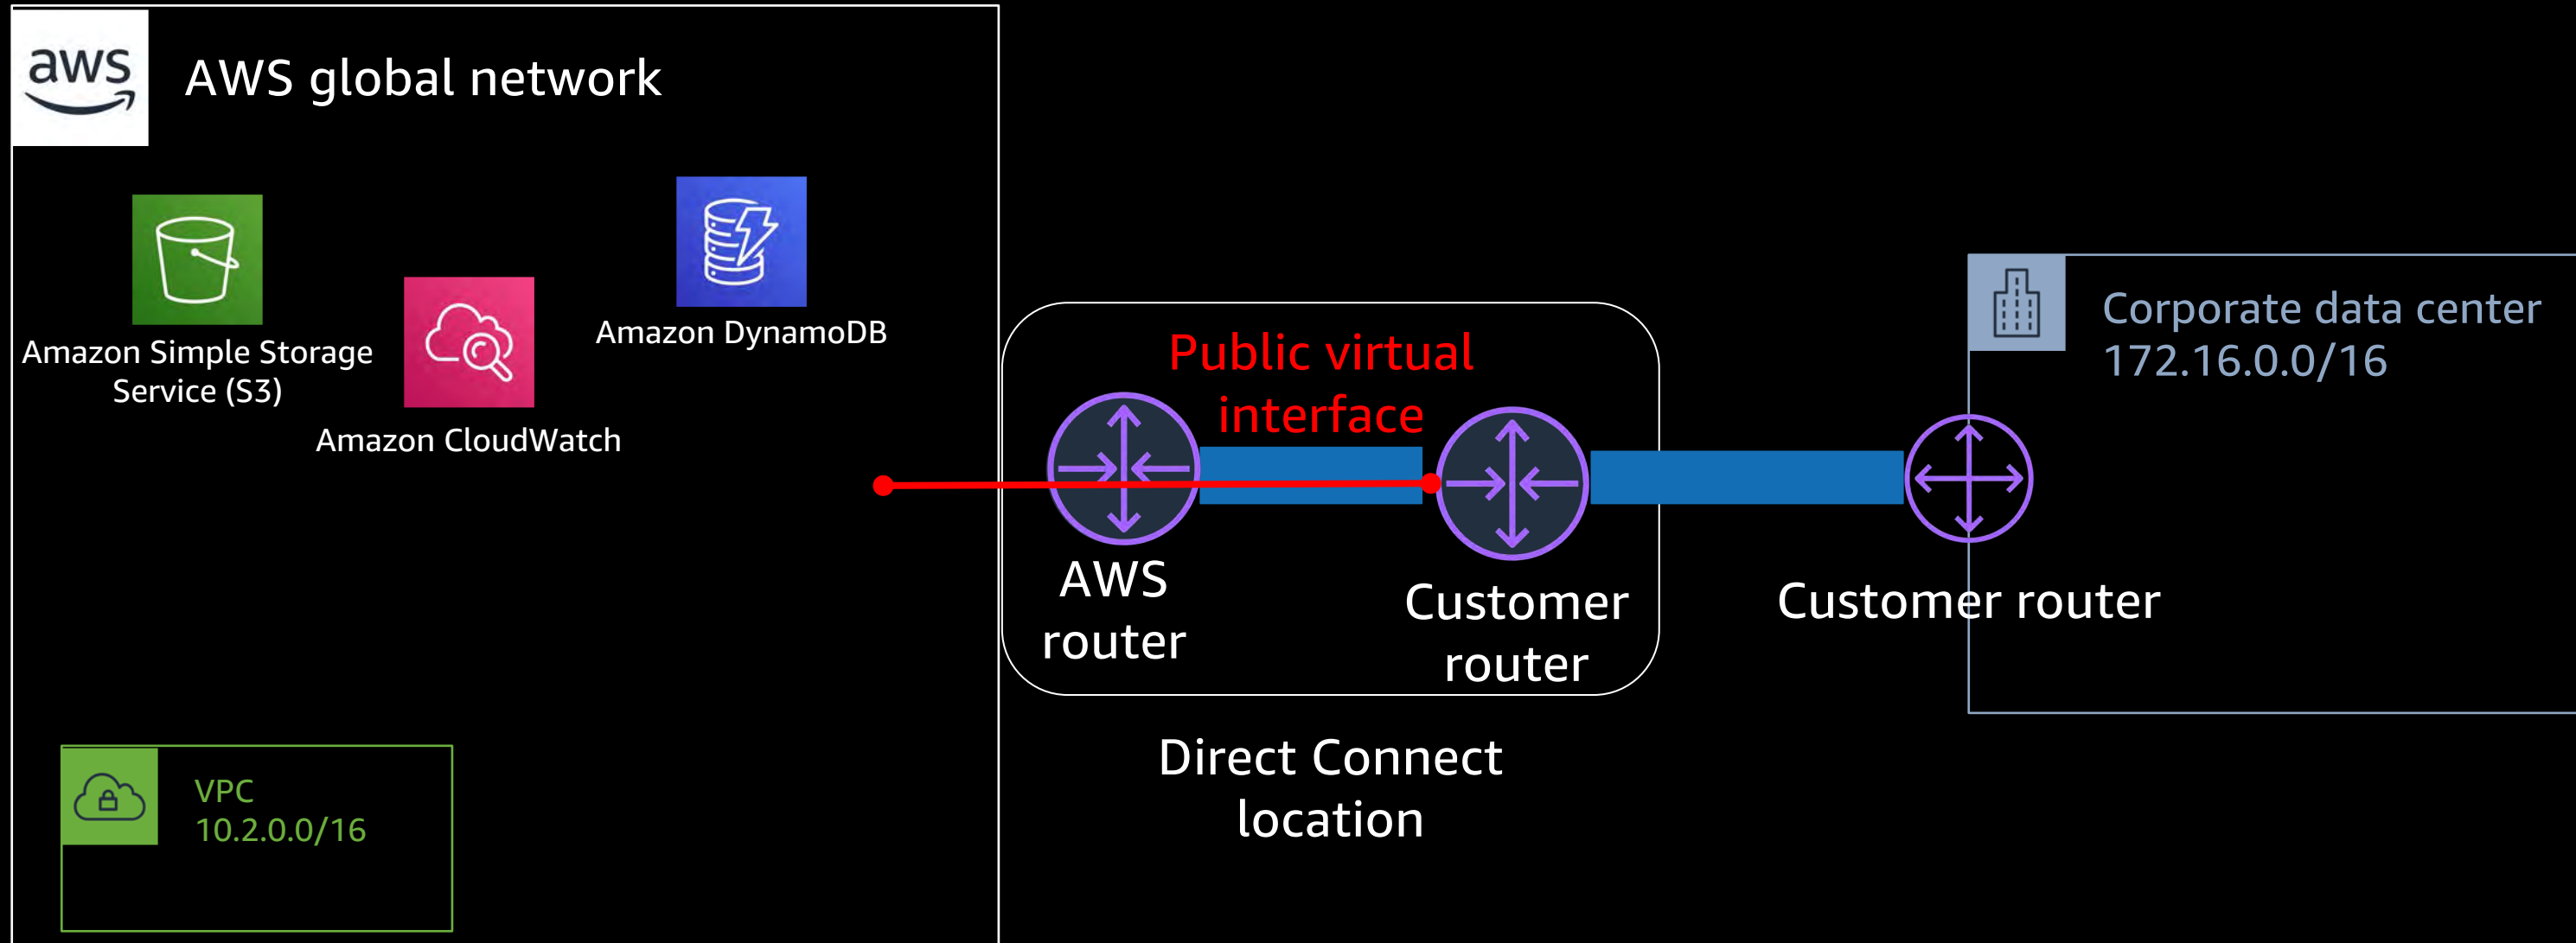
# Route propagation

- Enable propagation on the Route Table
- Automatically populates with anything the VGW learns via BGP

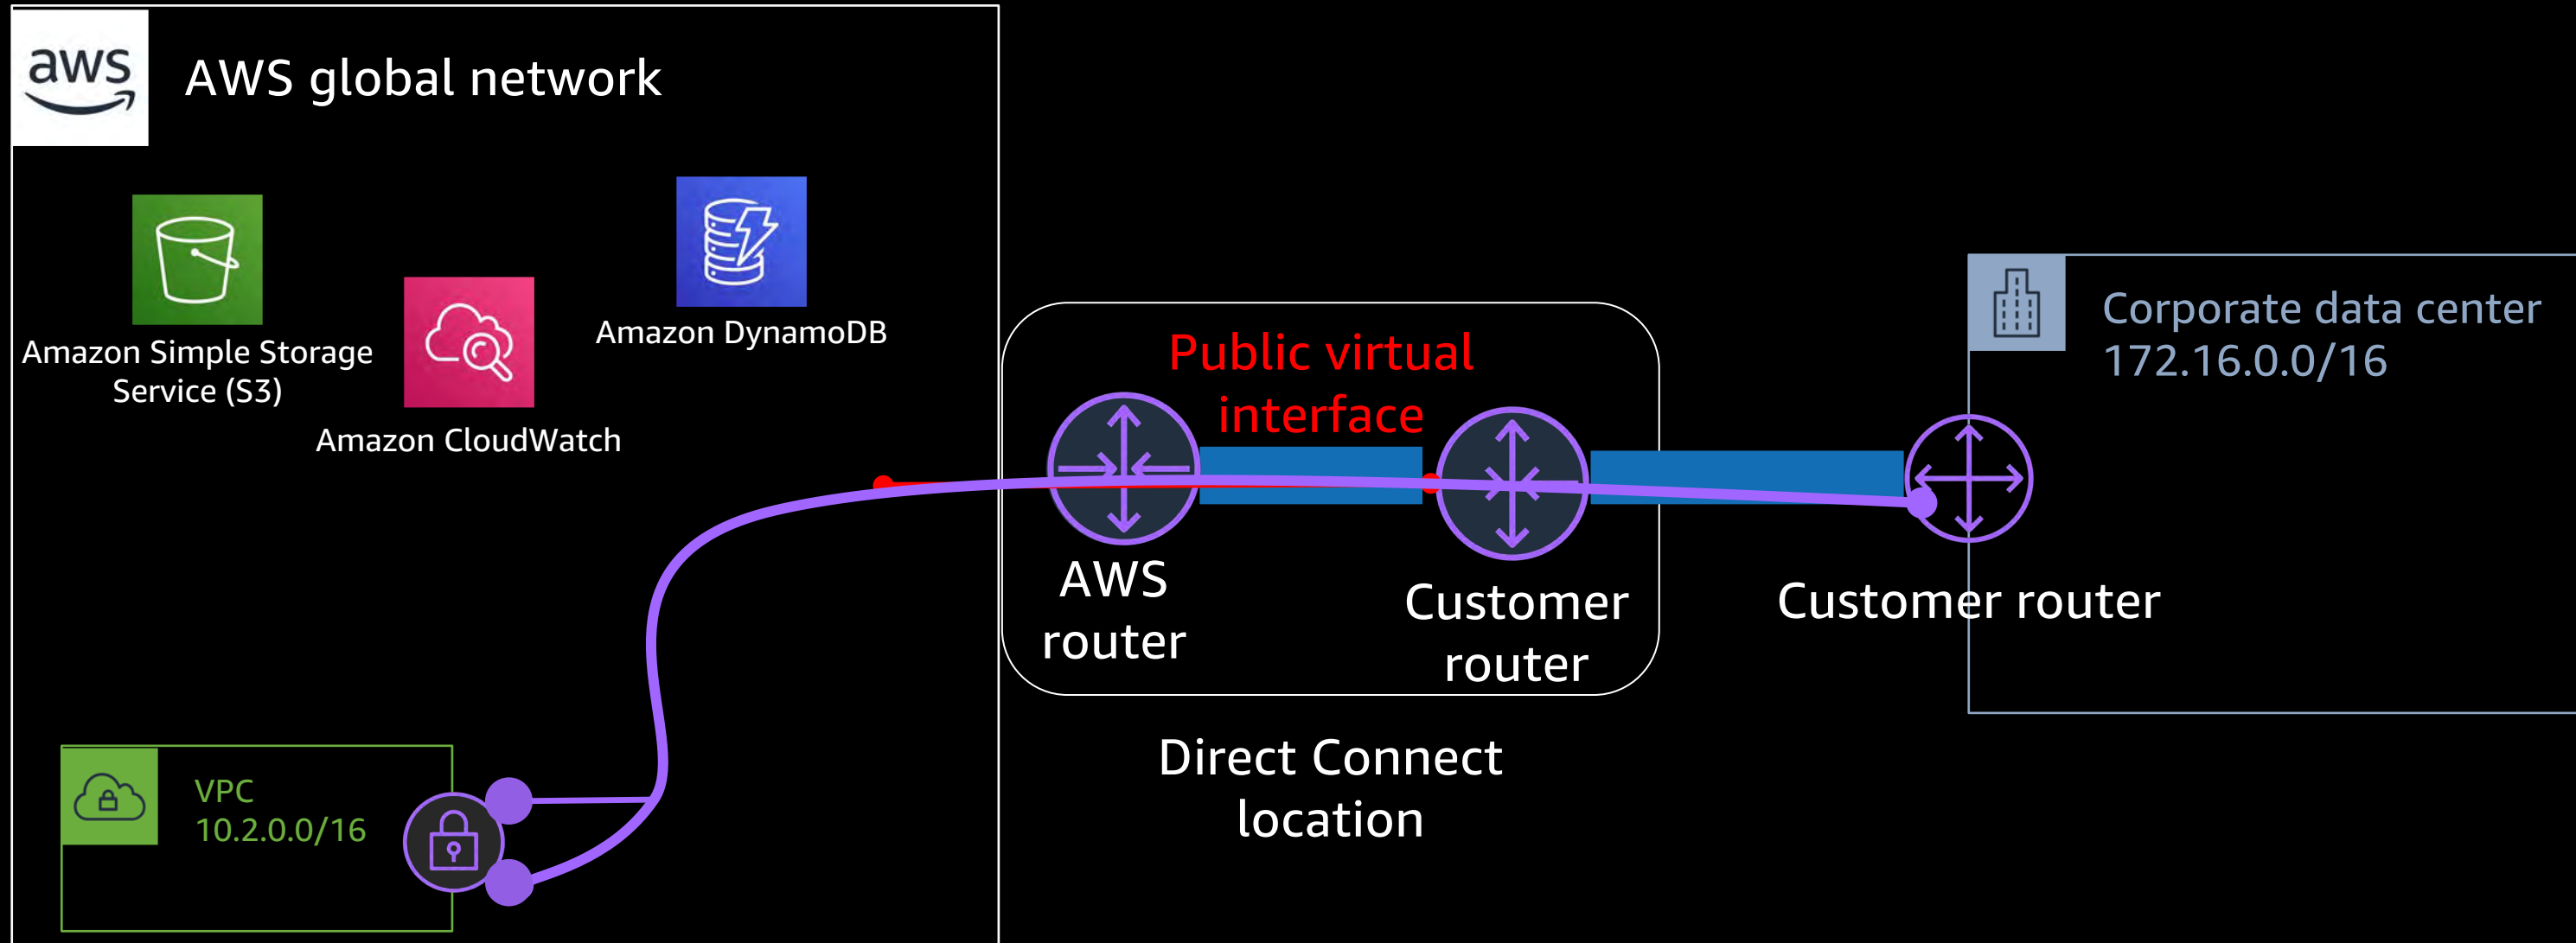




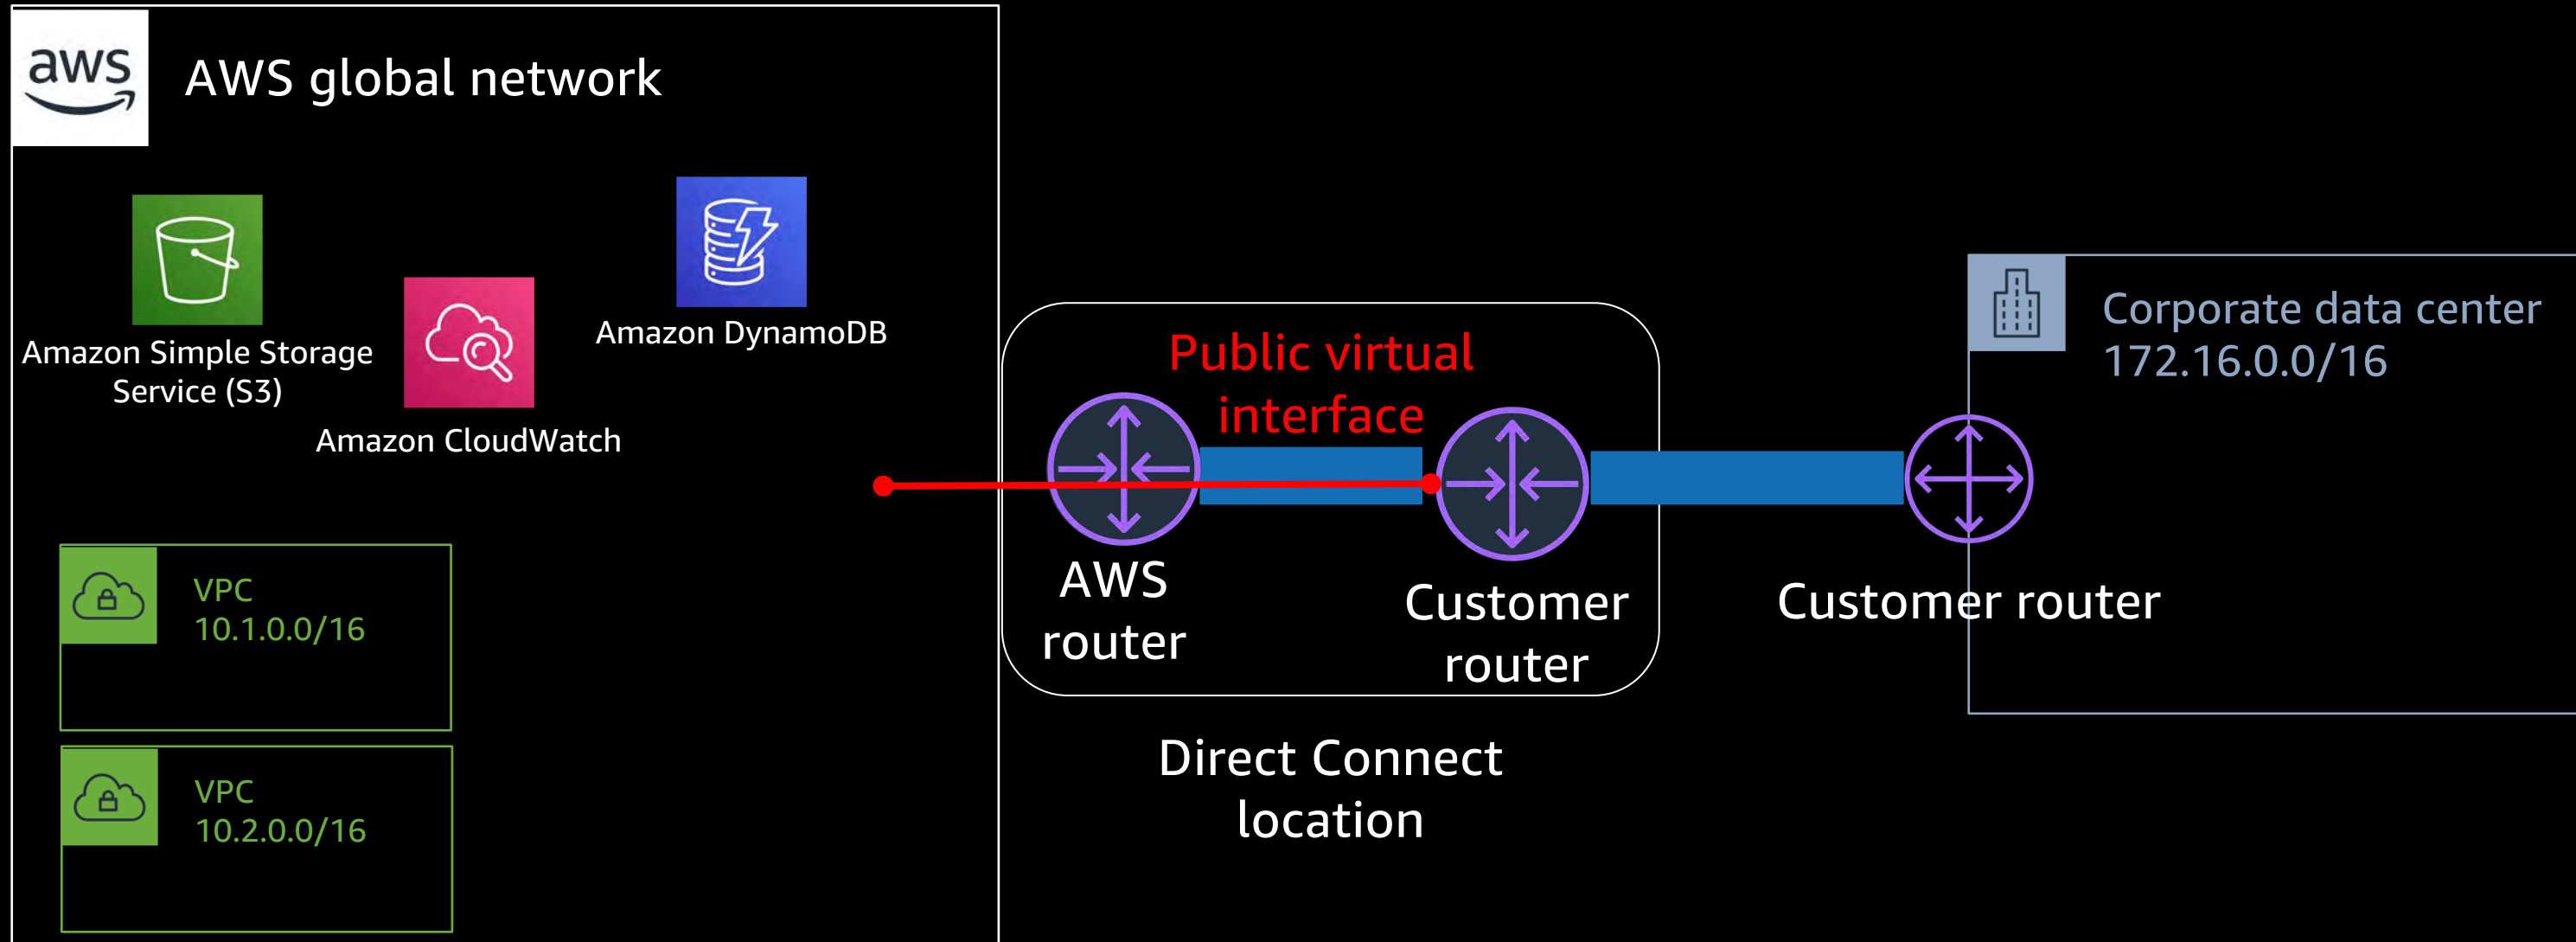
# AWS Direct Connect – Public VIF



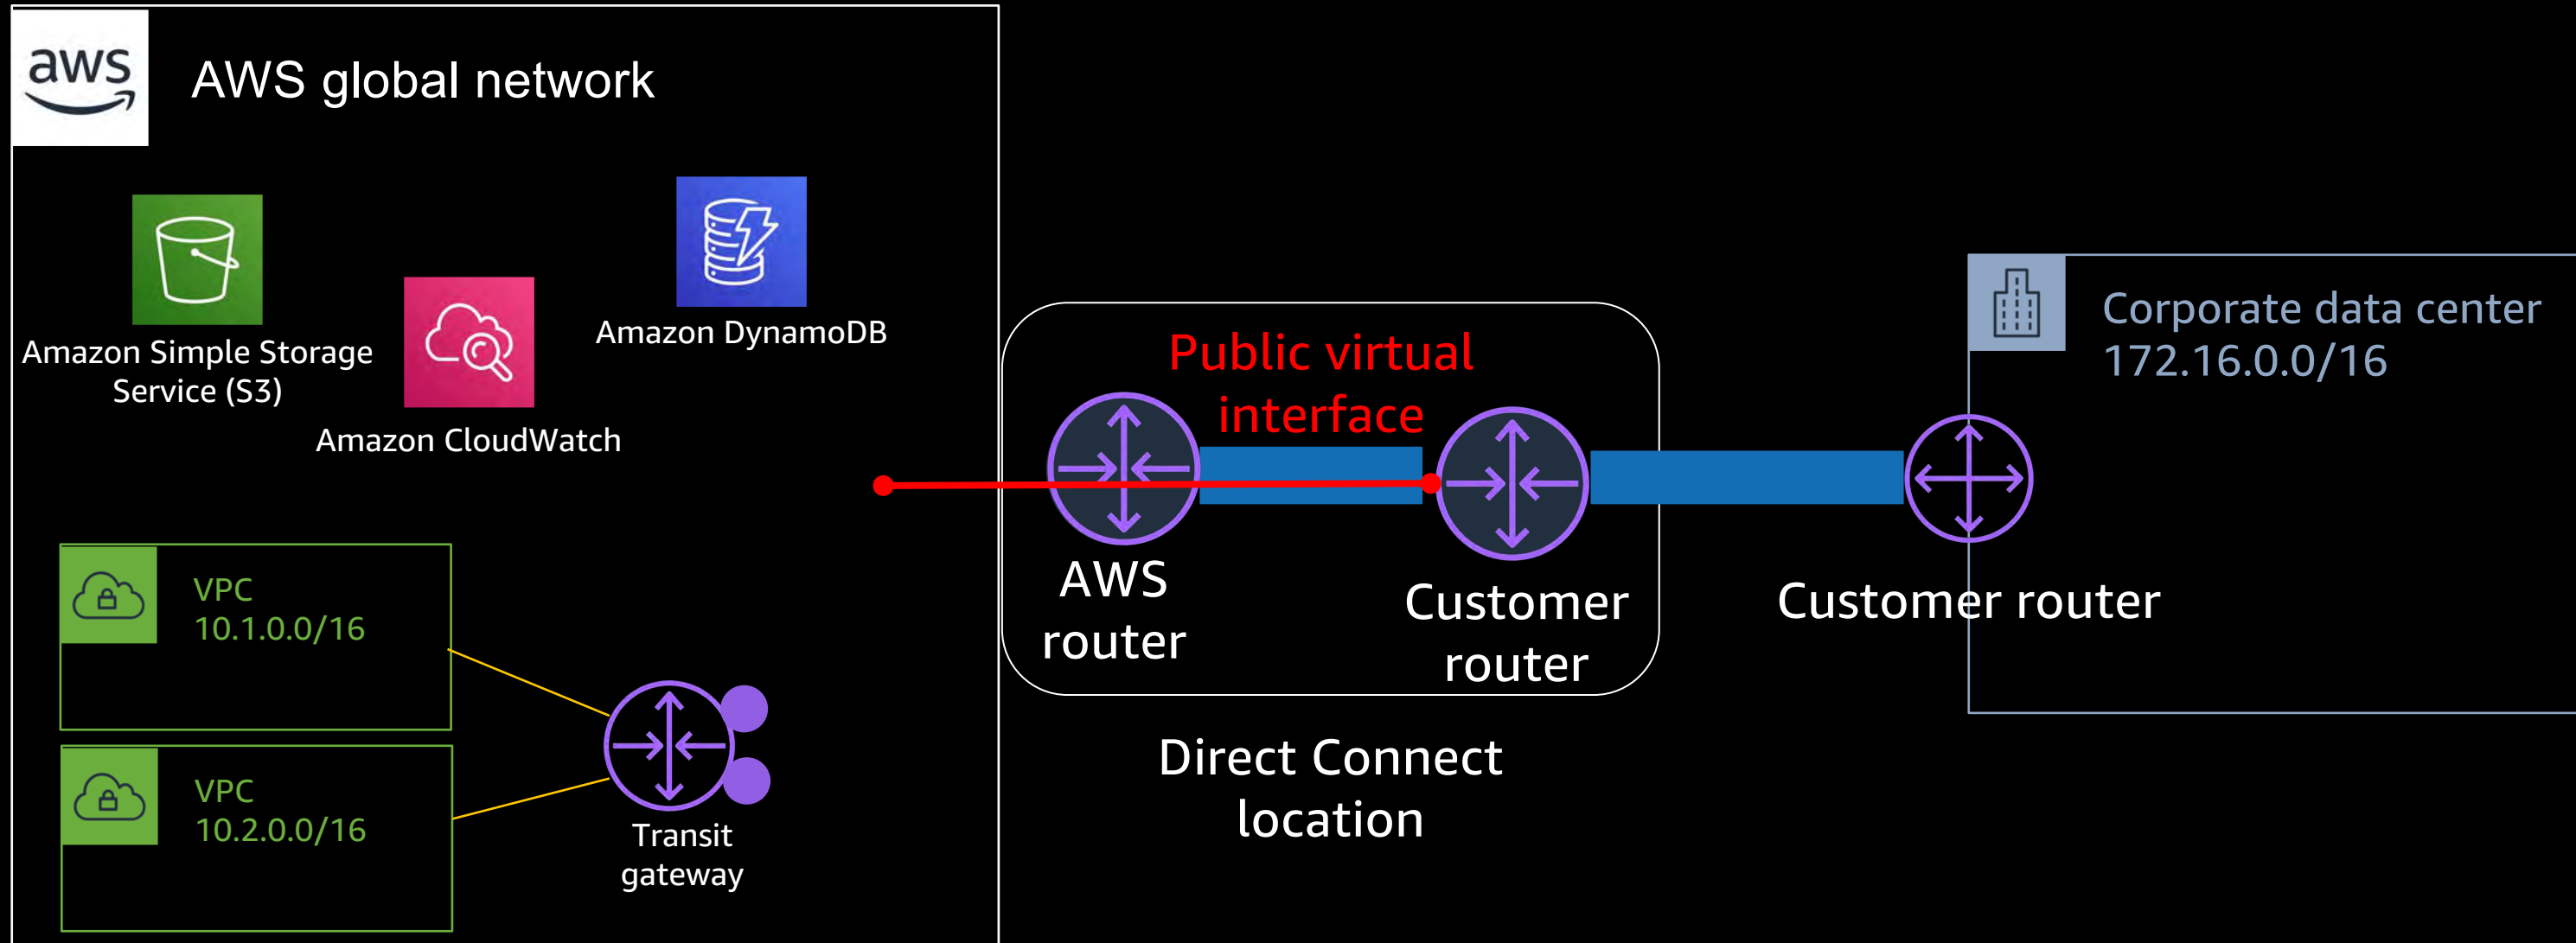
# AWS Direct Connect – Public VIF + AWS VPN



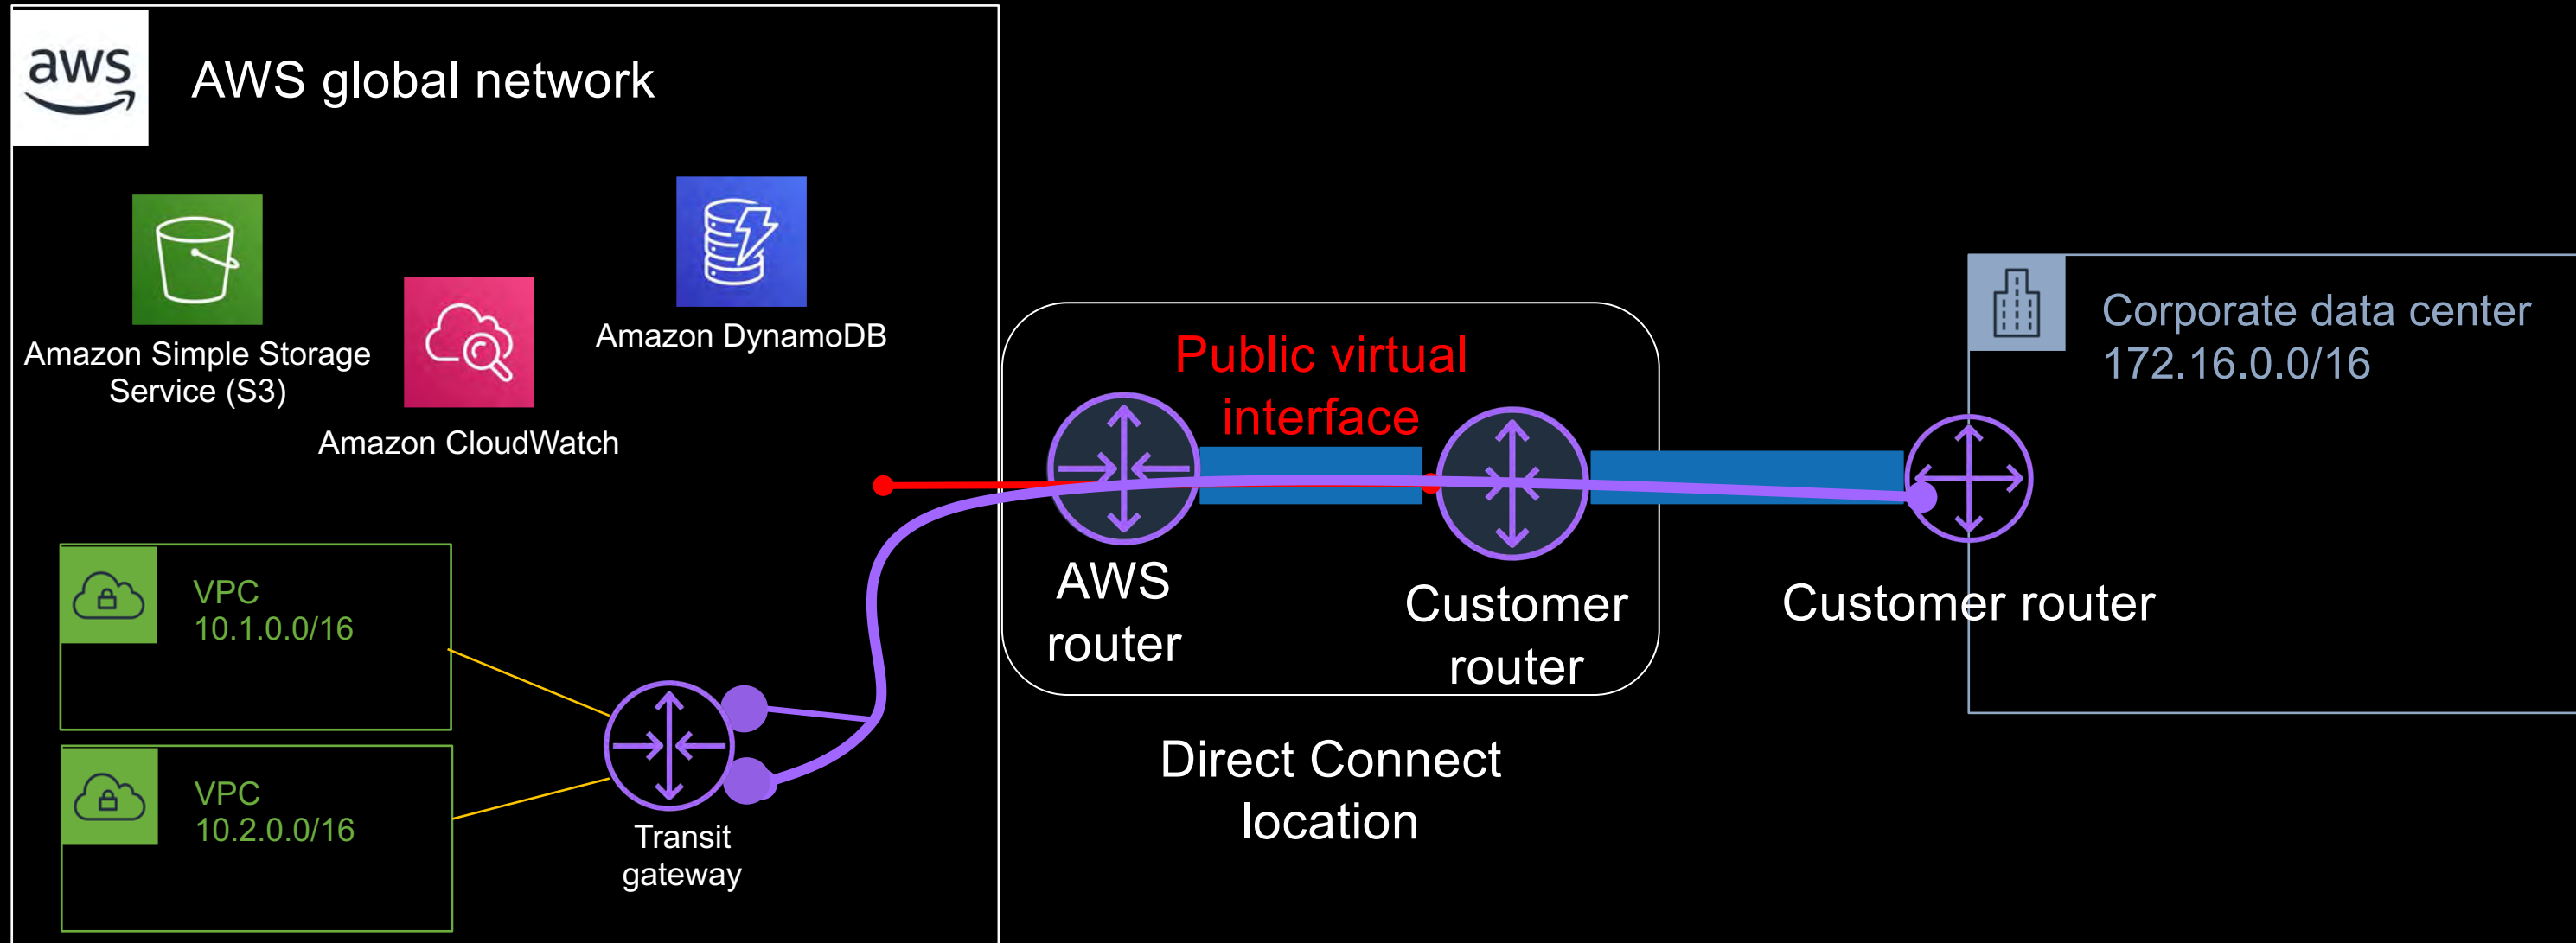
# AWS Direct Connect – Public VIF + AWS VPN



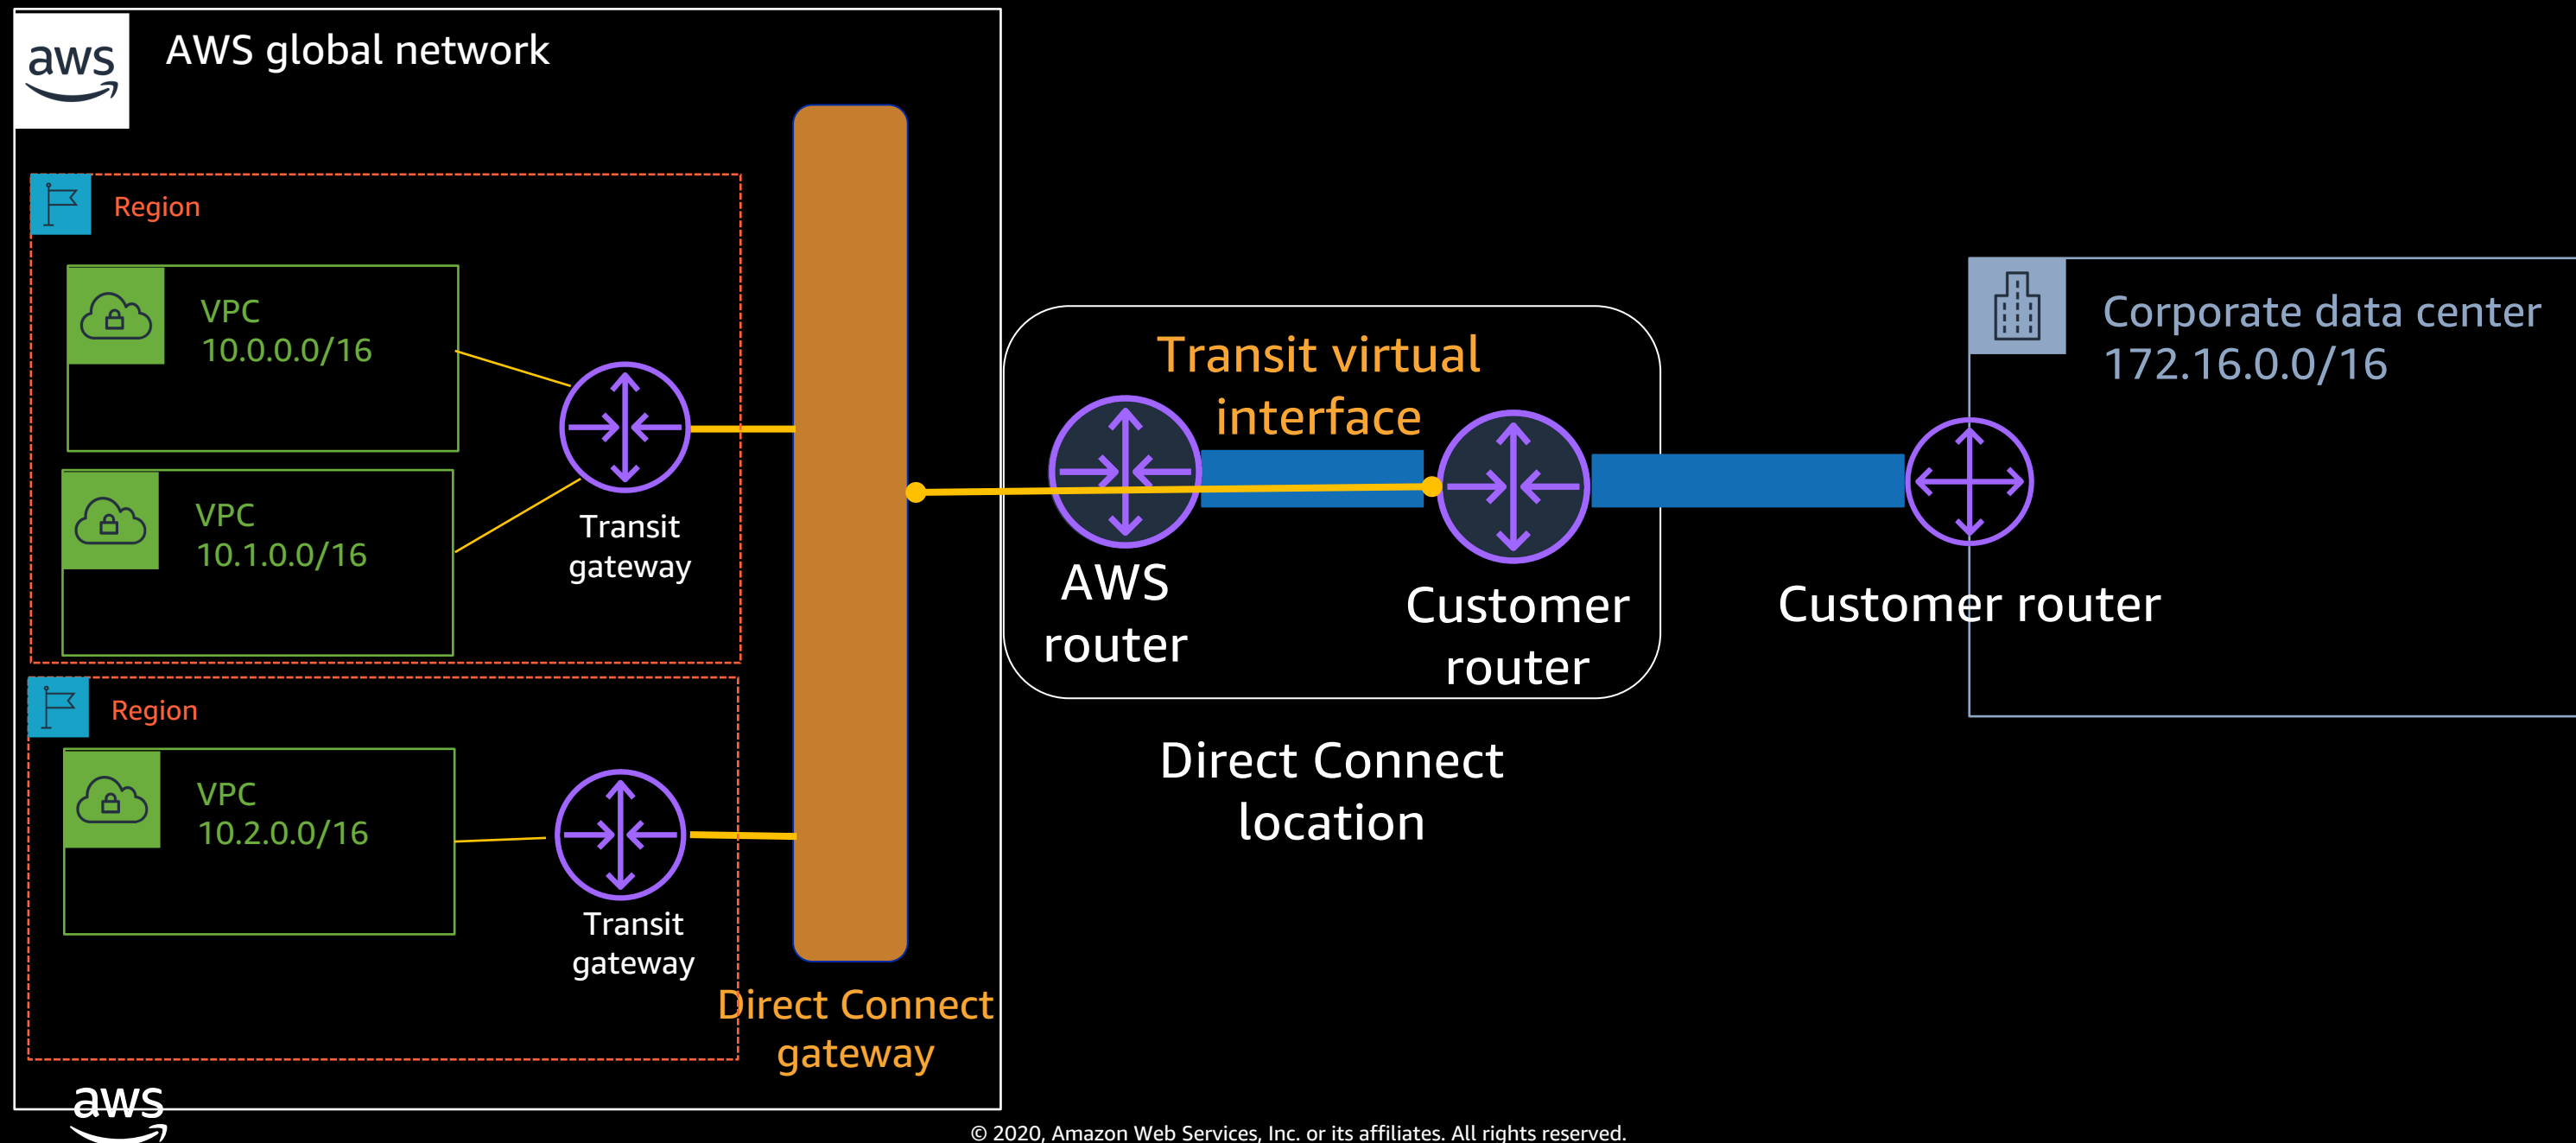
# AWS Direct Connect – Public VIF + AWS VPN



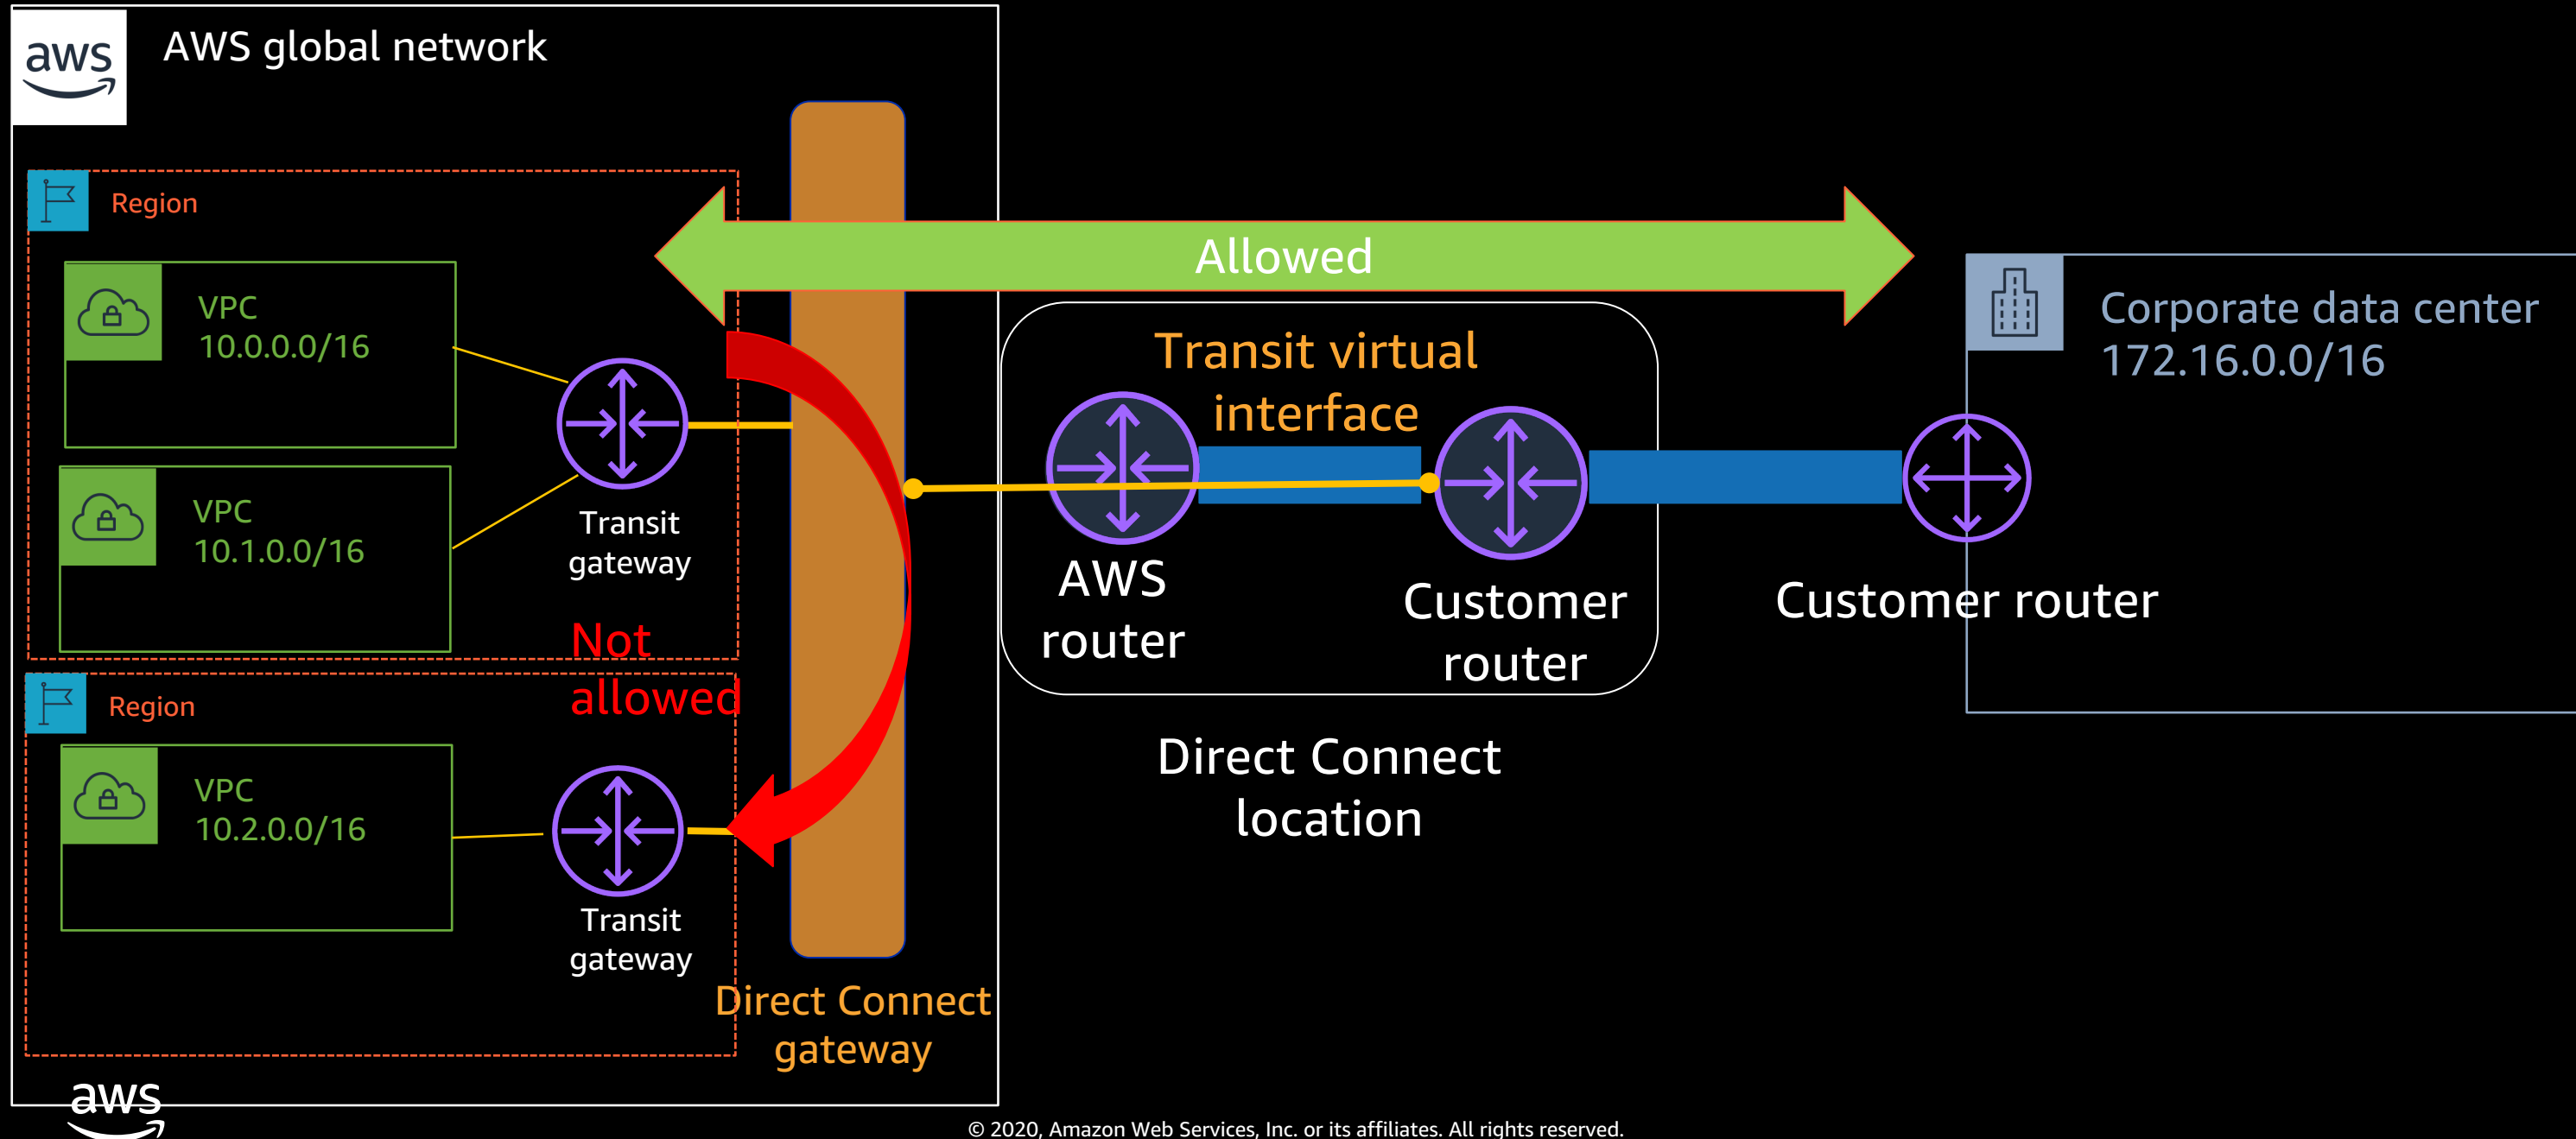
# AWS Direct Connect – Public VIF + AWS VPN



# AWS Direct Connect Gateway – Transit VIF




# AWS Direct Connect Gateway – Traffic flow




# AWS Direct Connect – Interface types


## Private VIF



## Transit VIF



## Public VIF




**Direct Connect gateway allows for connecting to resources in any AWS Region (except for China)**



# Direct Connect connection types


1 Gbps or 10 Gbps



50 VIFs  
**AND** 1 transit VIF

Dedicated  
connection

50 Mbps -> 10 Gbps



1 VIF OR  
1 transit VIF (1Gbps + only)

Hosted  
connection

50 Mbps -> 10 Gbps  
(can be oversubscribed)

1 VIF  
no transit VIF

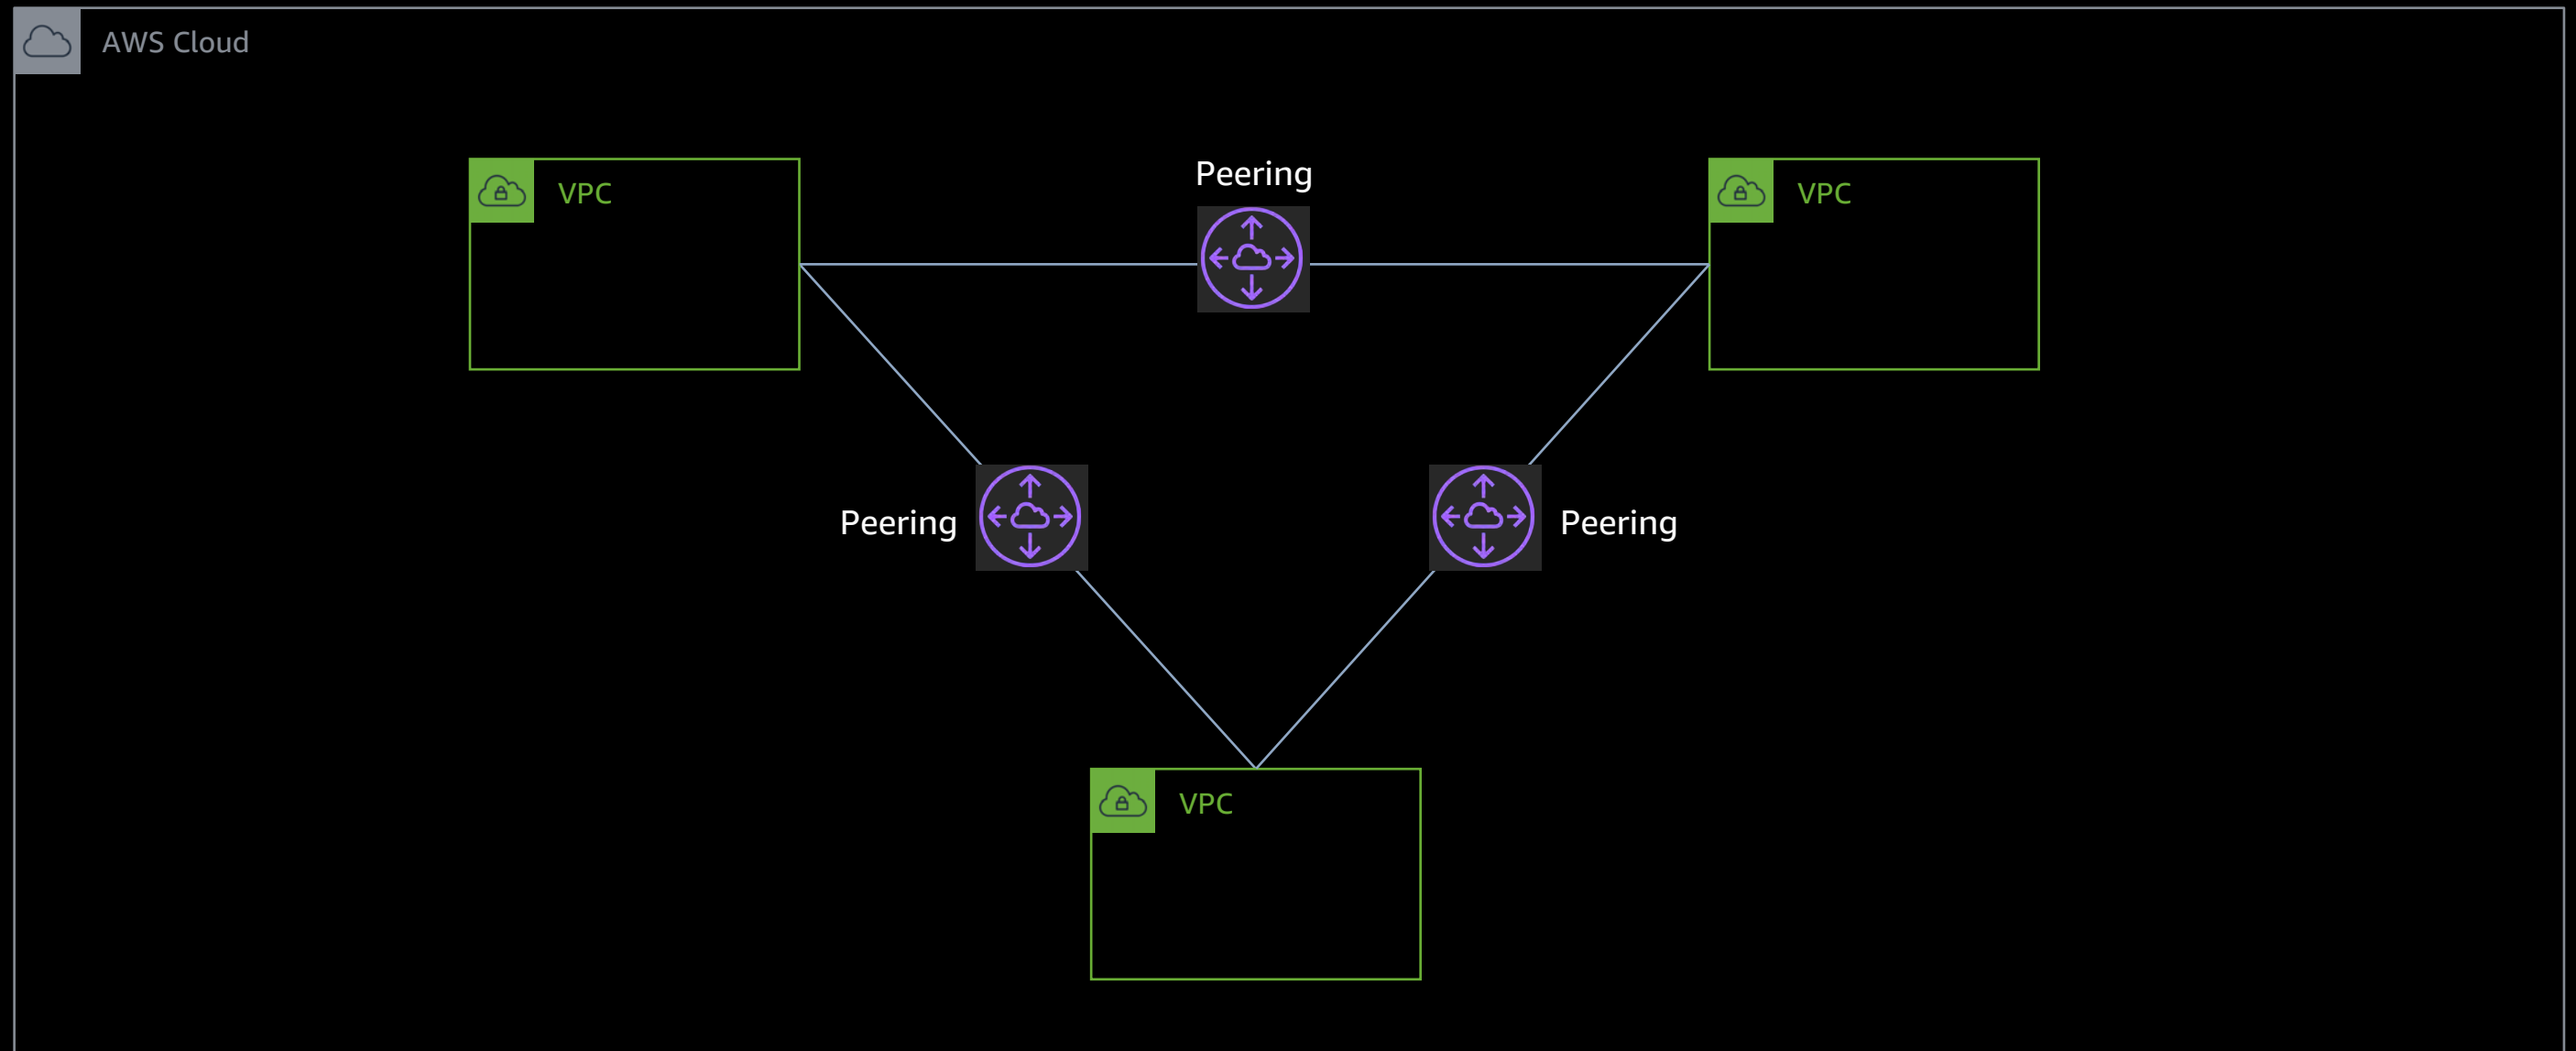
(Partner) Hosted  
virtual interface

**ONLY FROM PARTNERS**

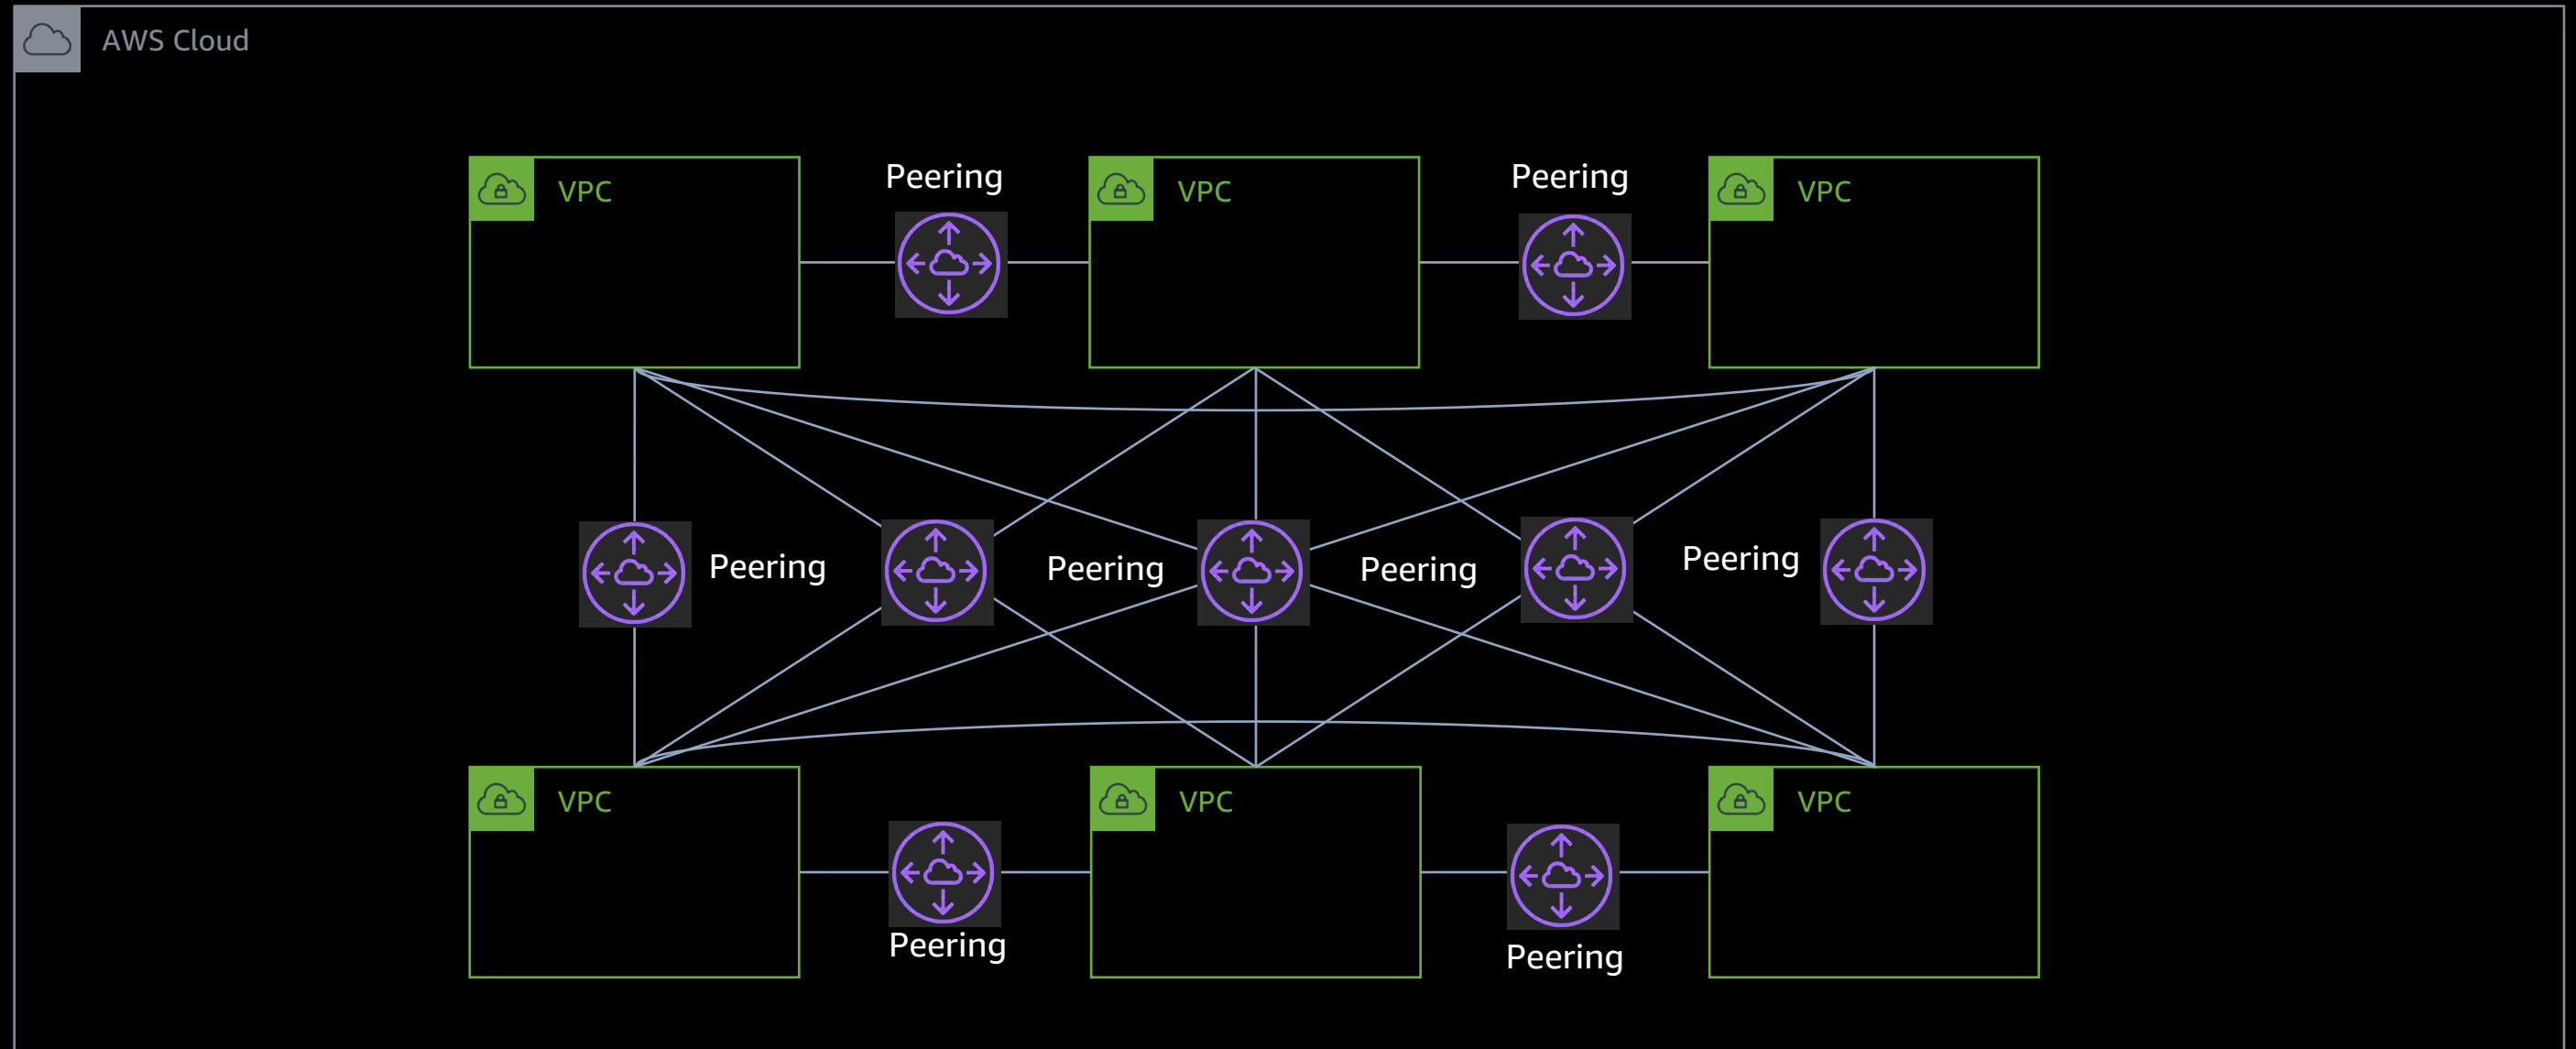
# AWS Transit Gateway



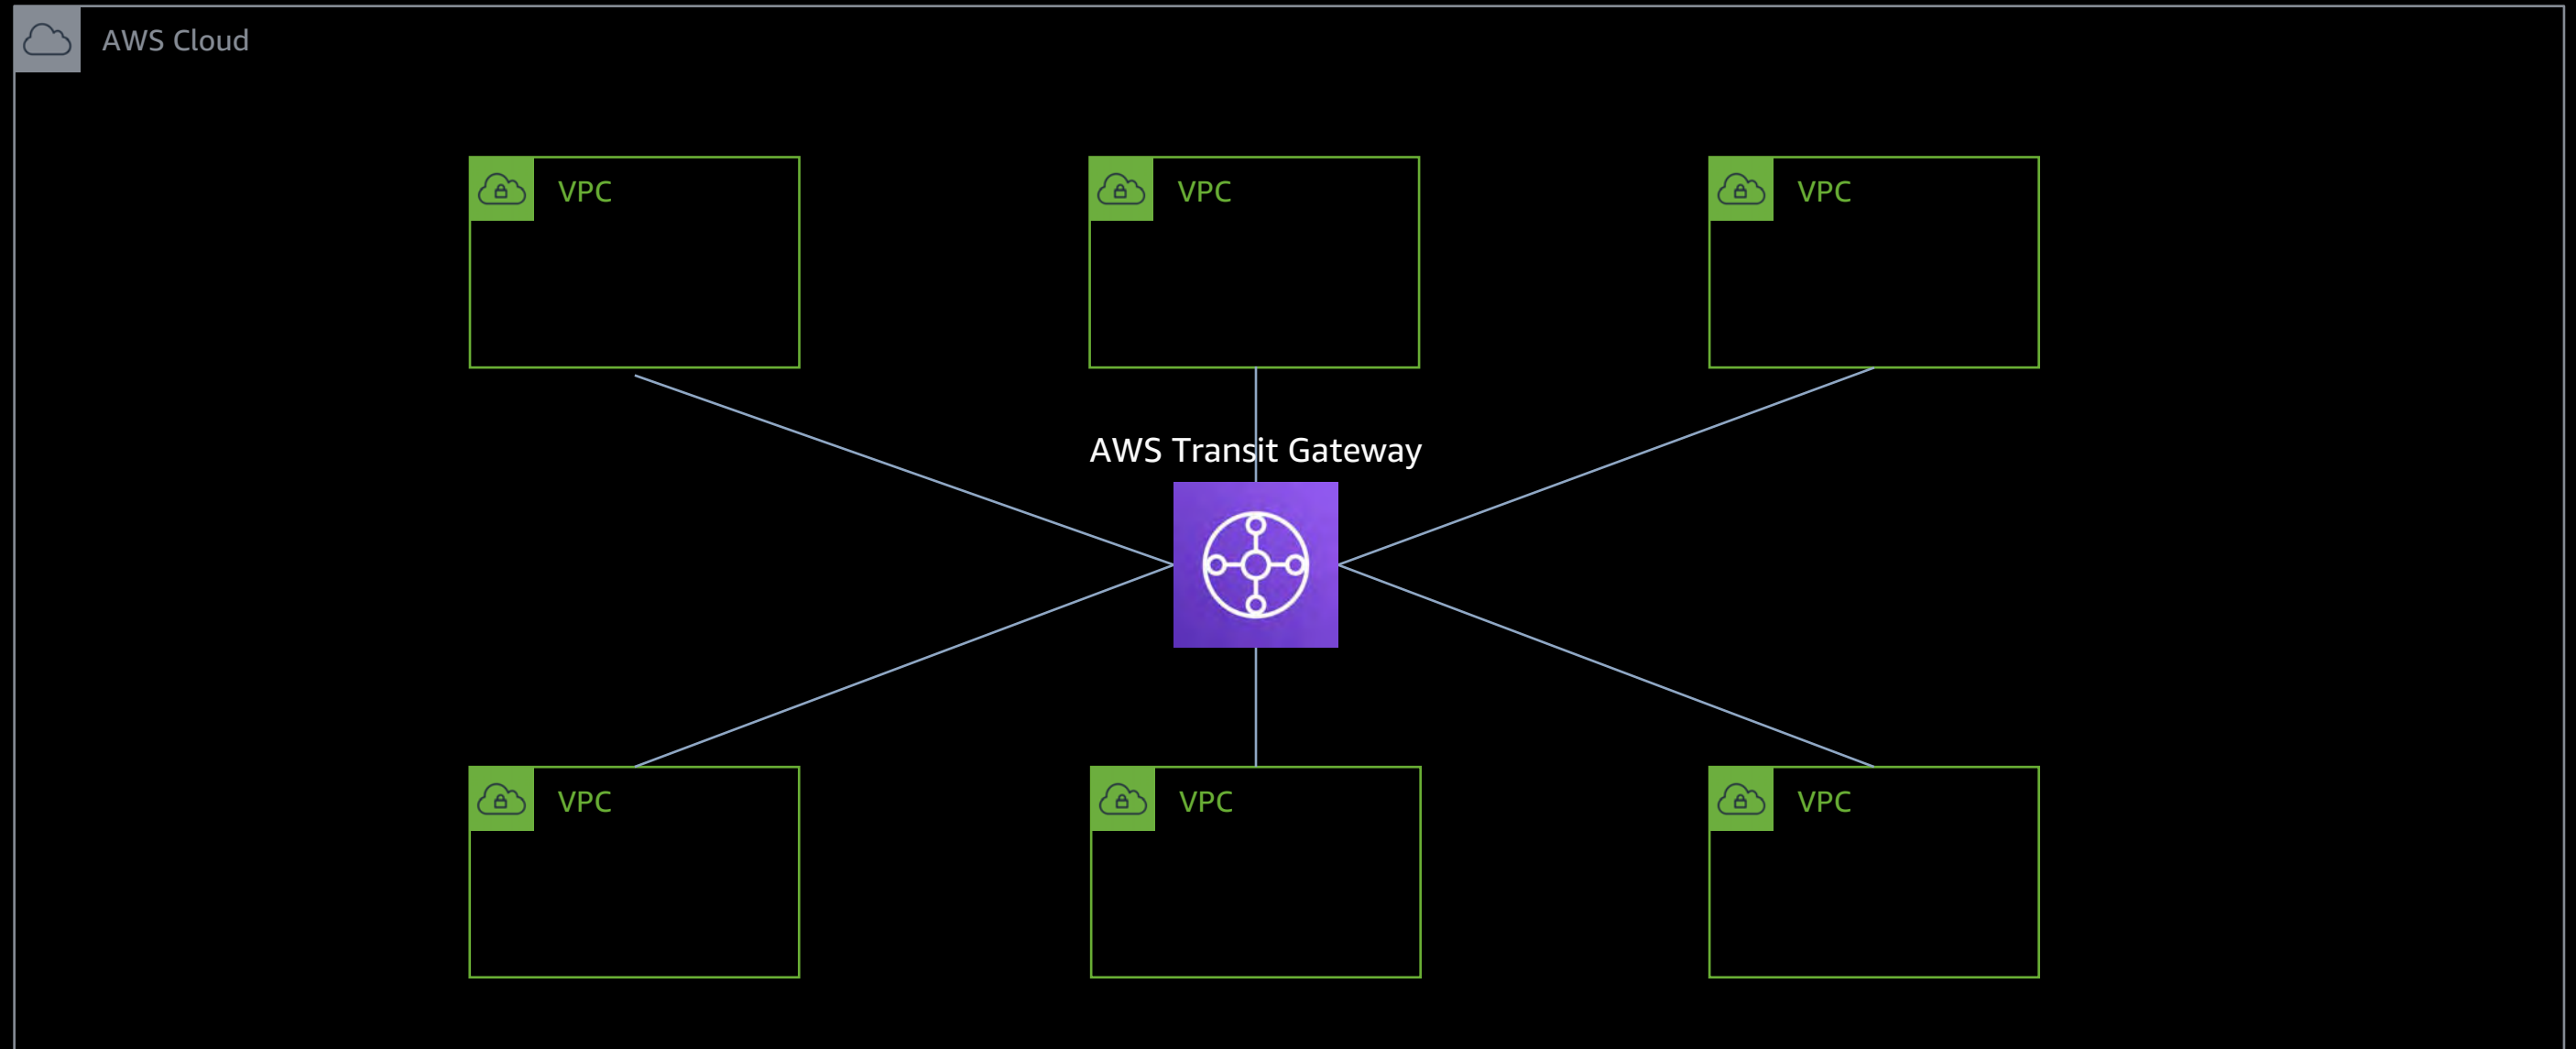
# Interconnecting VPCs at scale – VPC peering



# Interconnecting VPCs at scale – VPC peering



# Multiple VPCs access models – AWS Transit Gateway



## Attachment

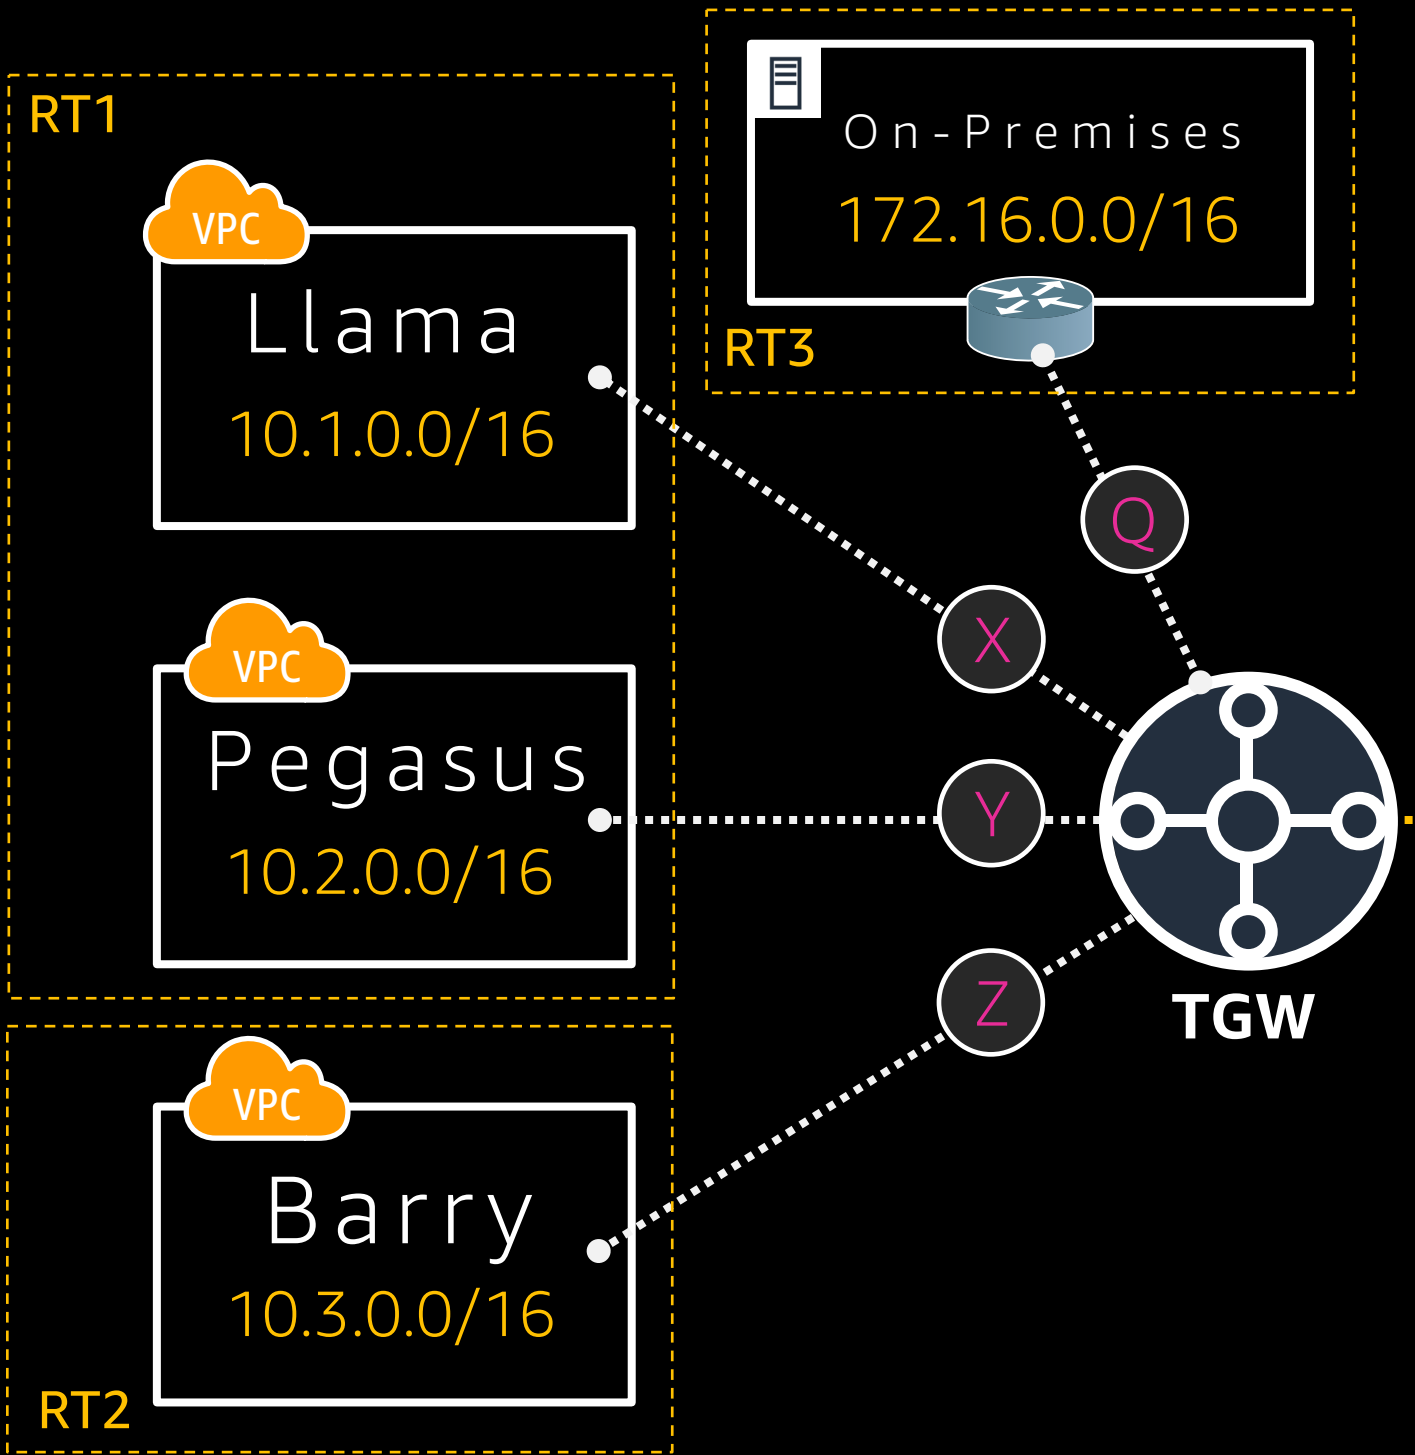
The connection from an Amazon VPC, VPN, and DX GW to a Transit Gateway

## Association

The route table used to route packets coming from an attachment

## Propagation

The route table where the attachment's routes are installed

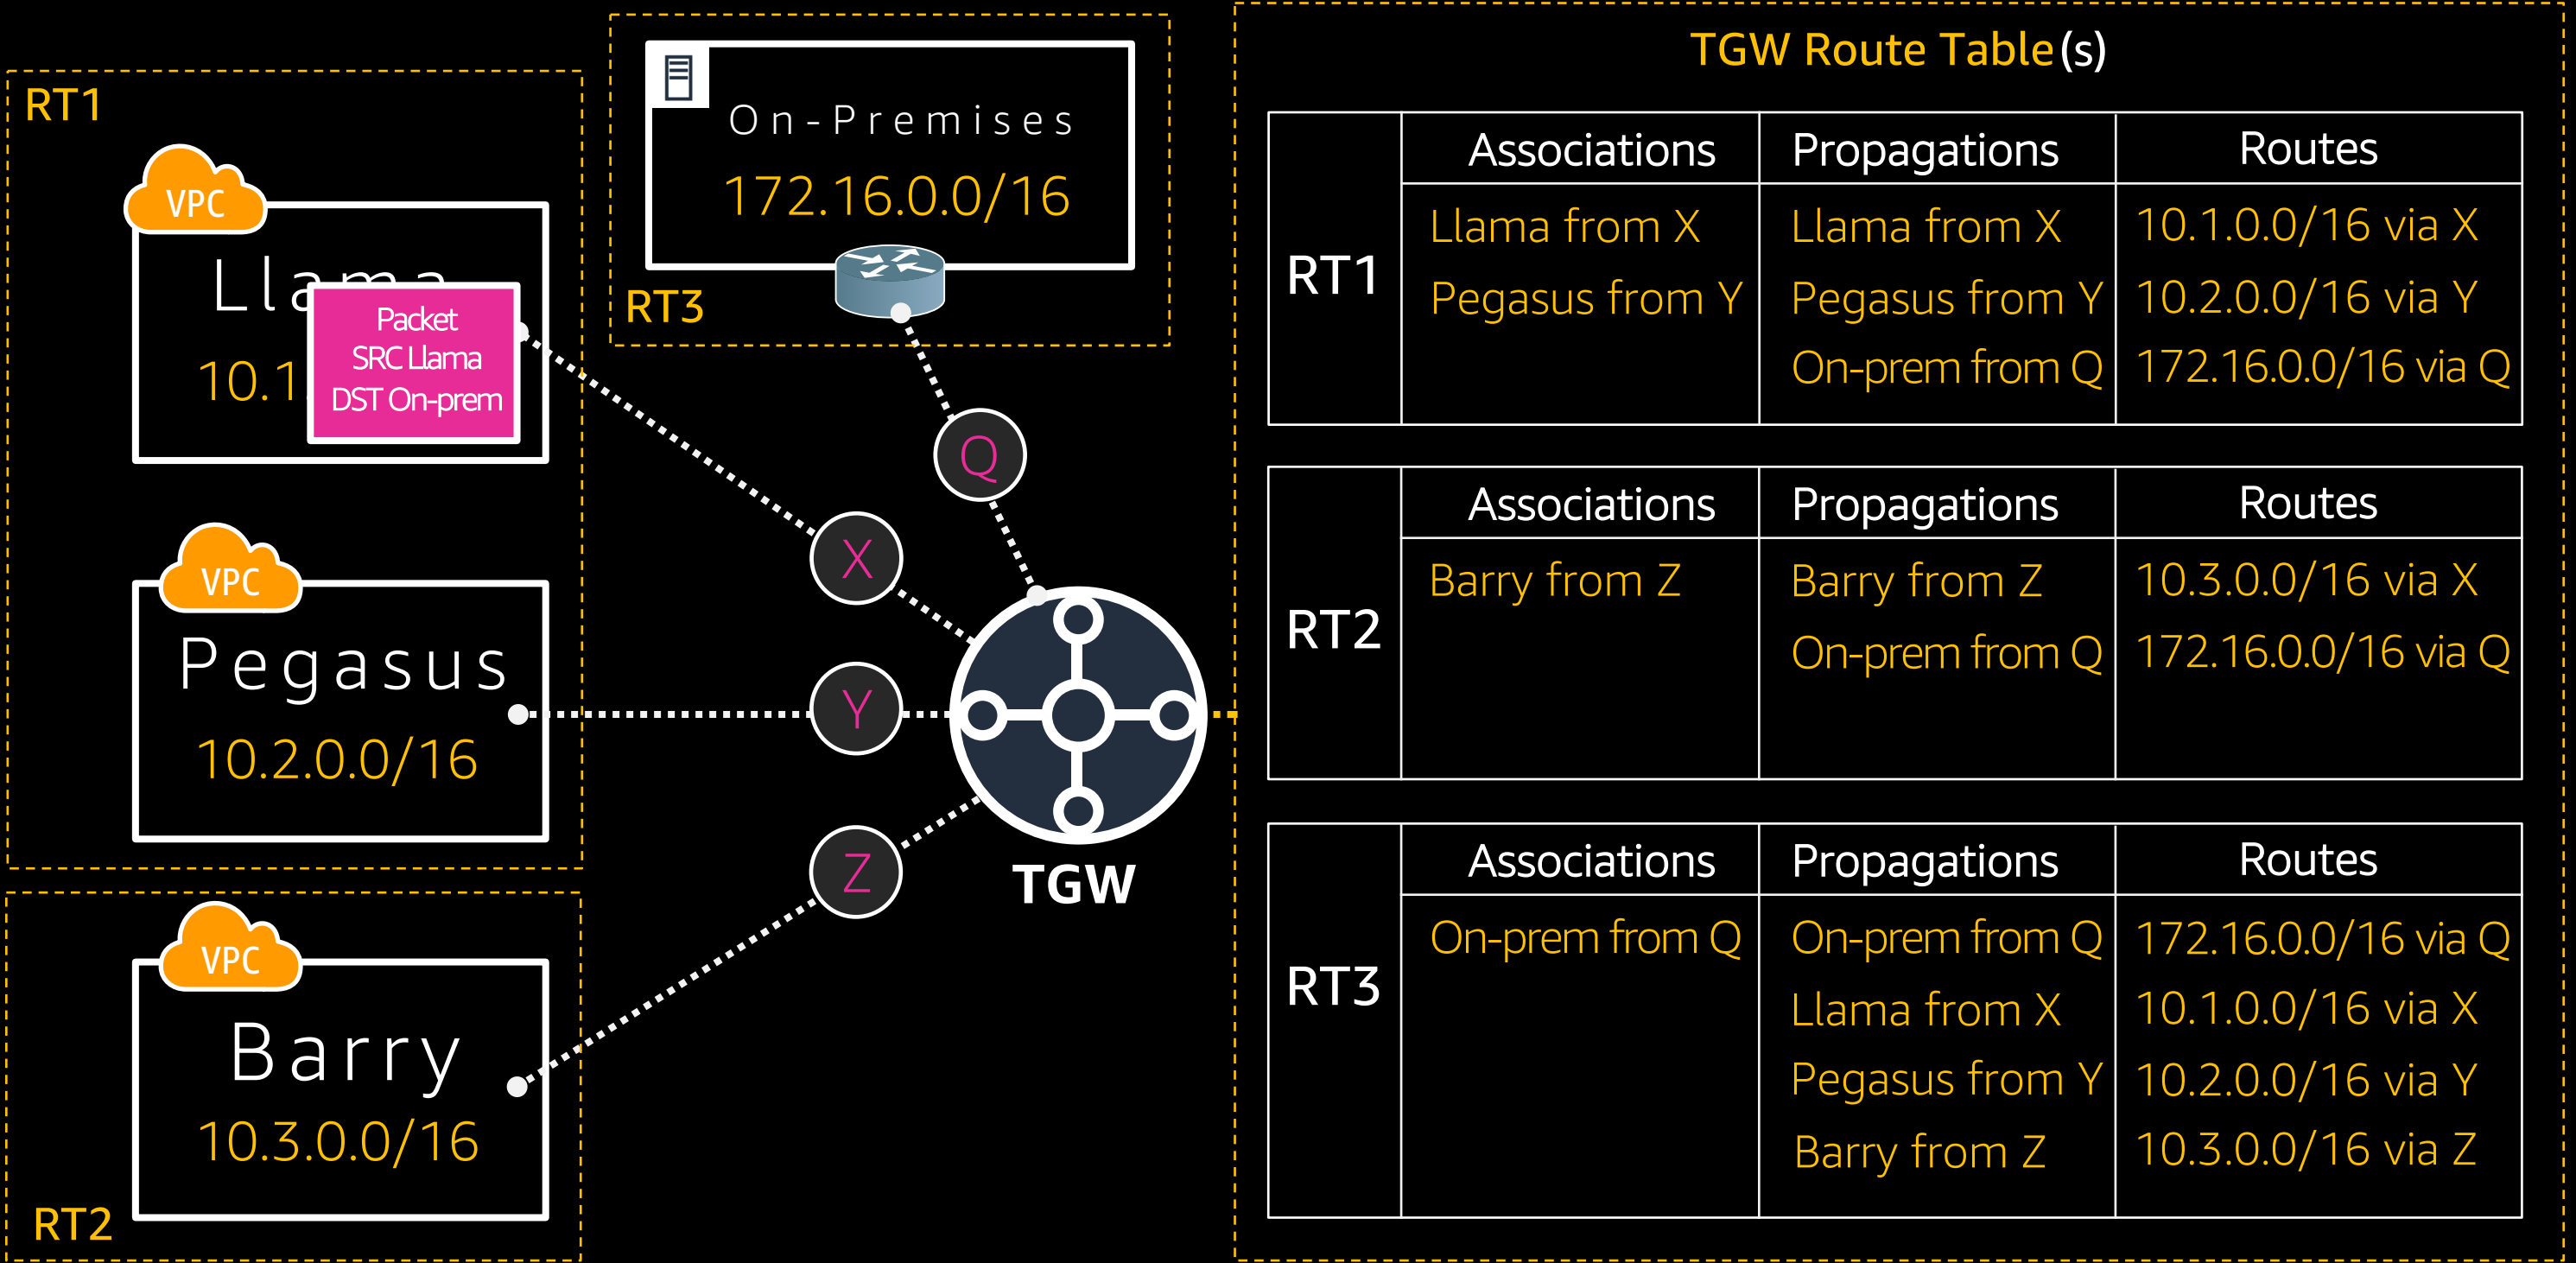


**TGW Route Table(s)**

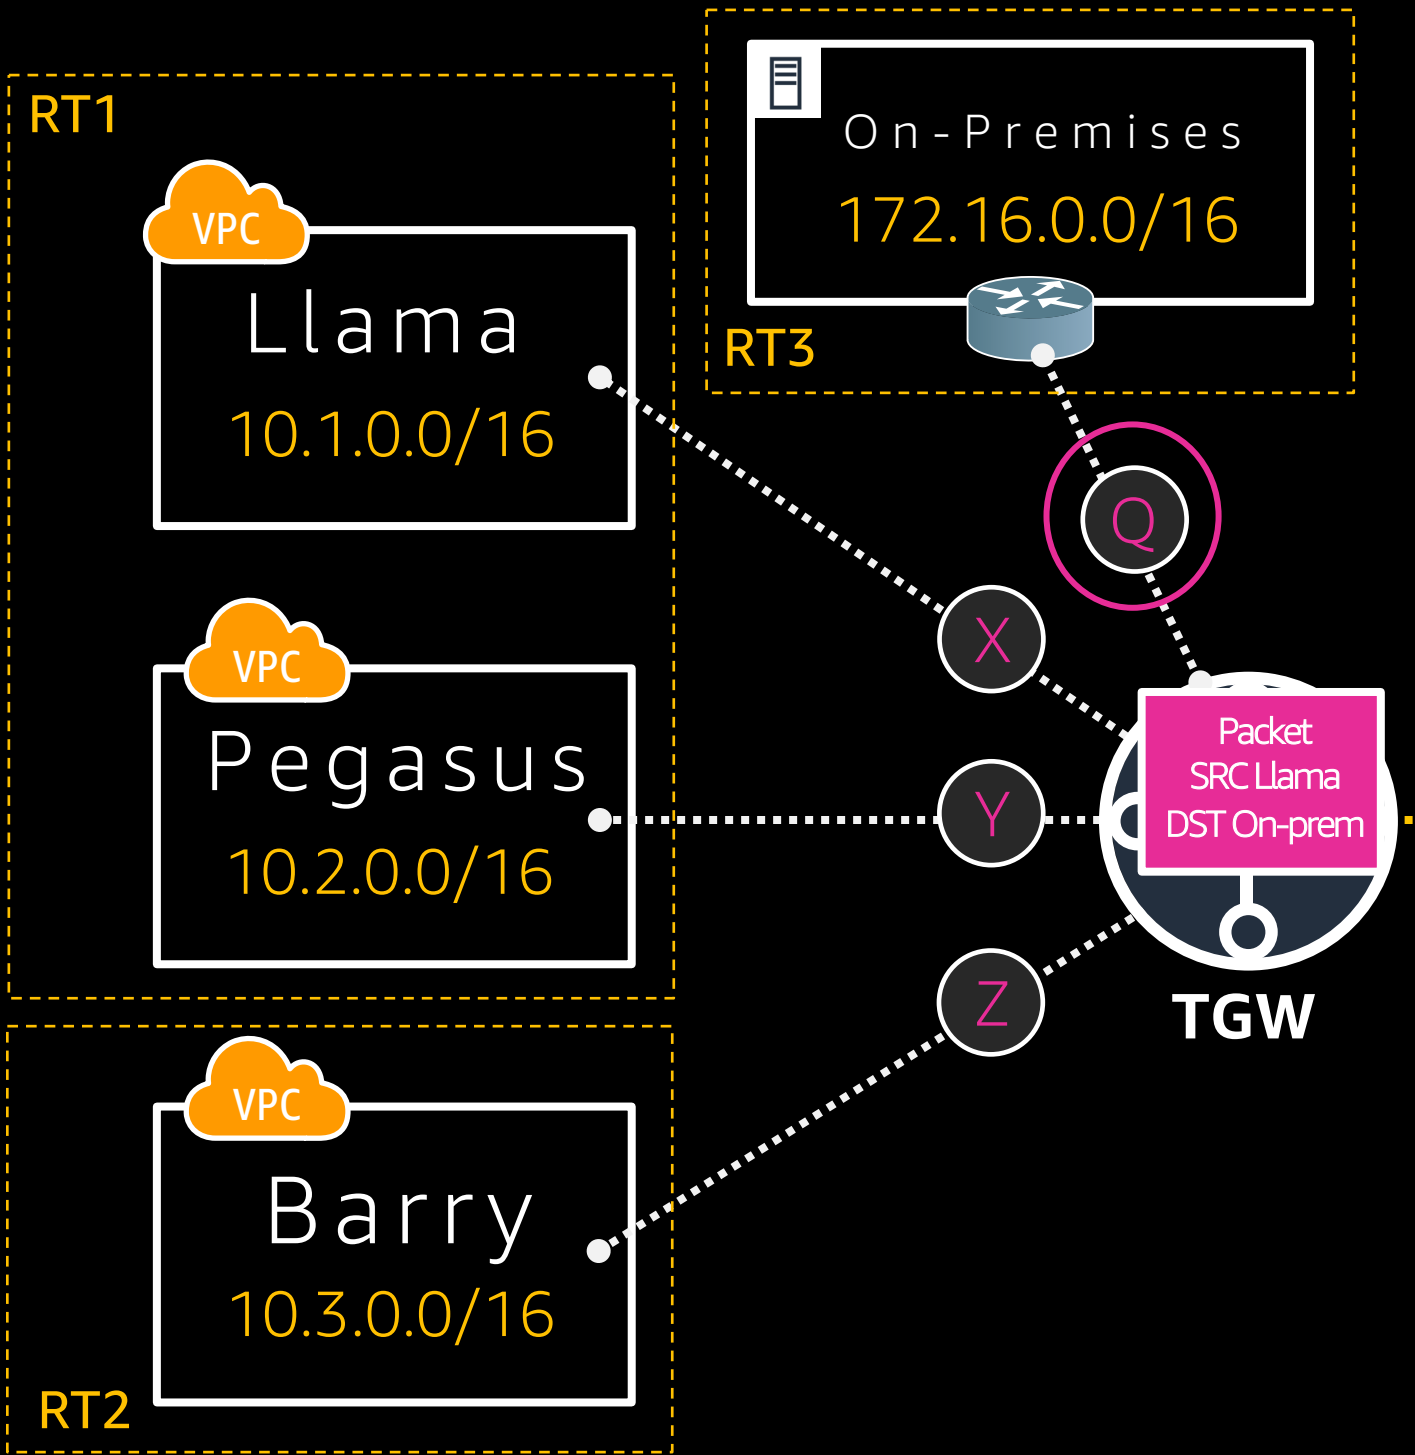
RT1	Associations	Propagations	Routes
	Llama from X Pegasus from Y	Llama from X Pegasus from Y On-prem from Q	10.1.0.0/16 via X 10.2.0.0/16 via Y 172.16.0.0/16 via Q

RT2	Associations	Propagations	Routes
	Barry from Z	Barry from Z On-prem from Q	10.3.0.0/16 via Z 172.16.0.0/16 via Q

RT3	Associations	Propagations	Routes
	On-prem from Q	On-prem from Q Llama from X Pegasus from Y Barry from Z	172.16.0.0/16 via Q 10.1.0.0/16 via X 10.2.0.0/16 via Y 10.3.0.0/16 via Z







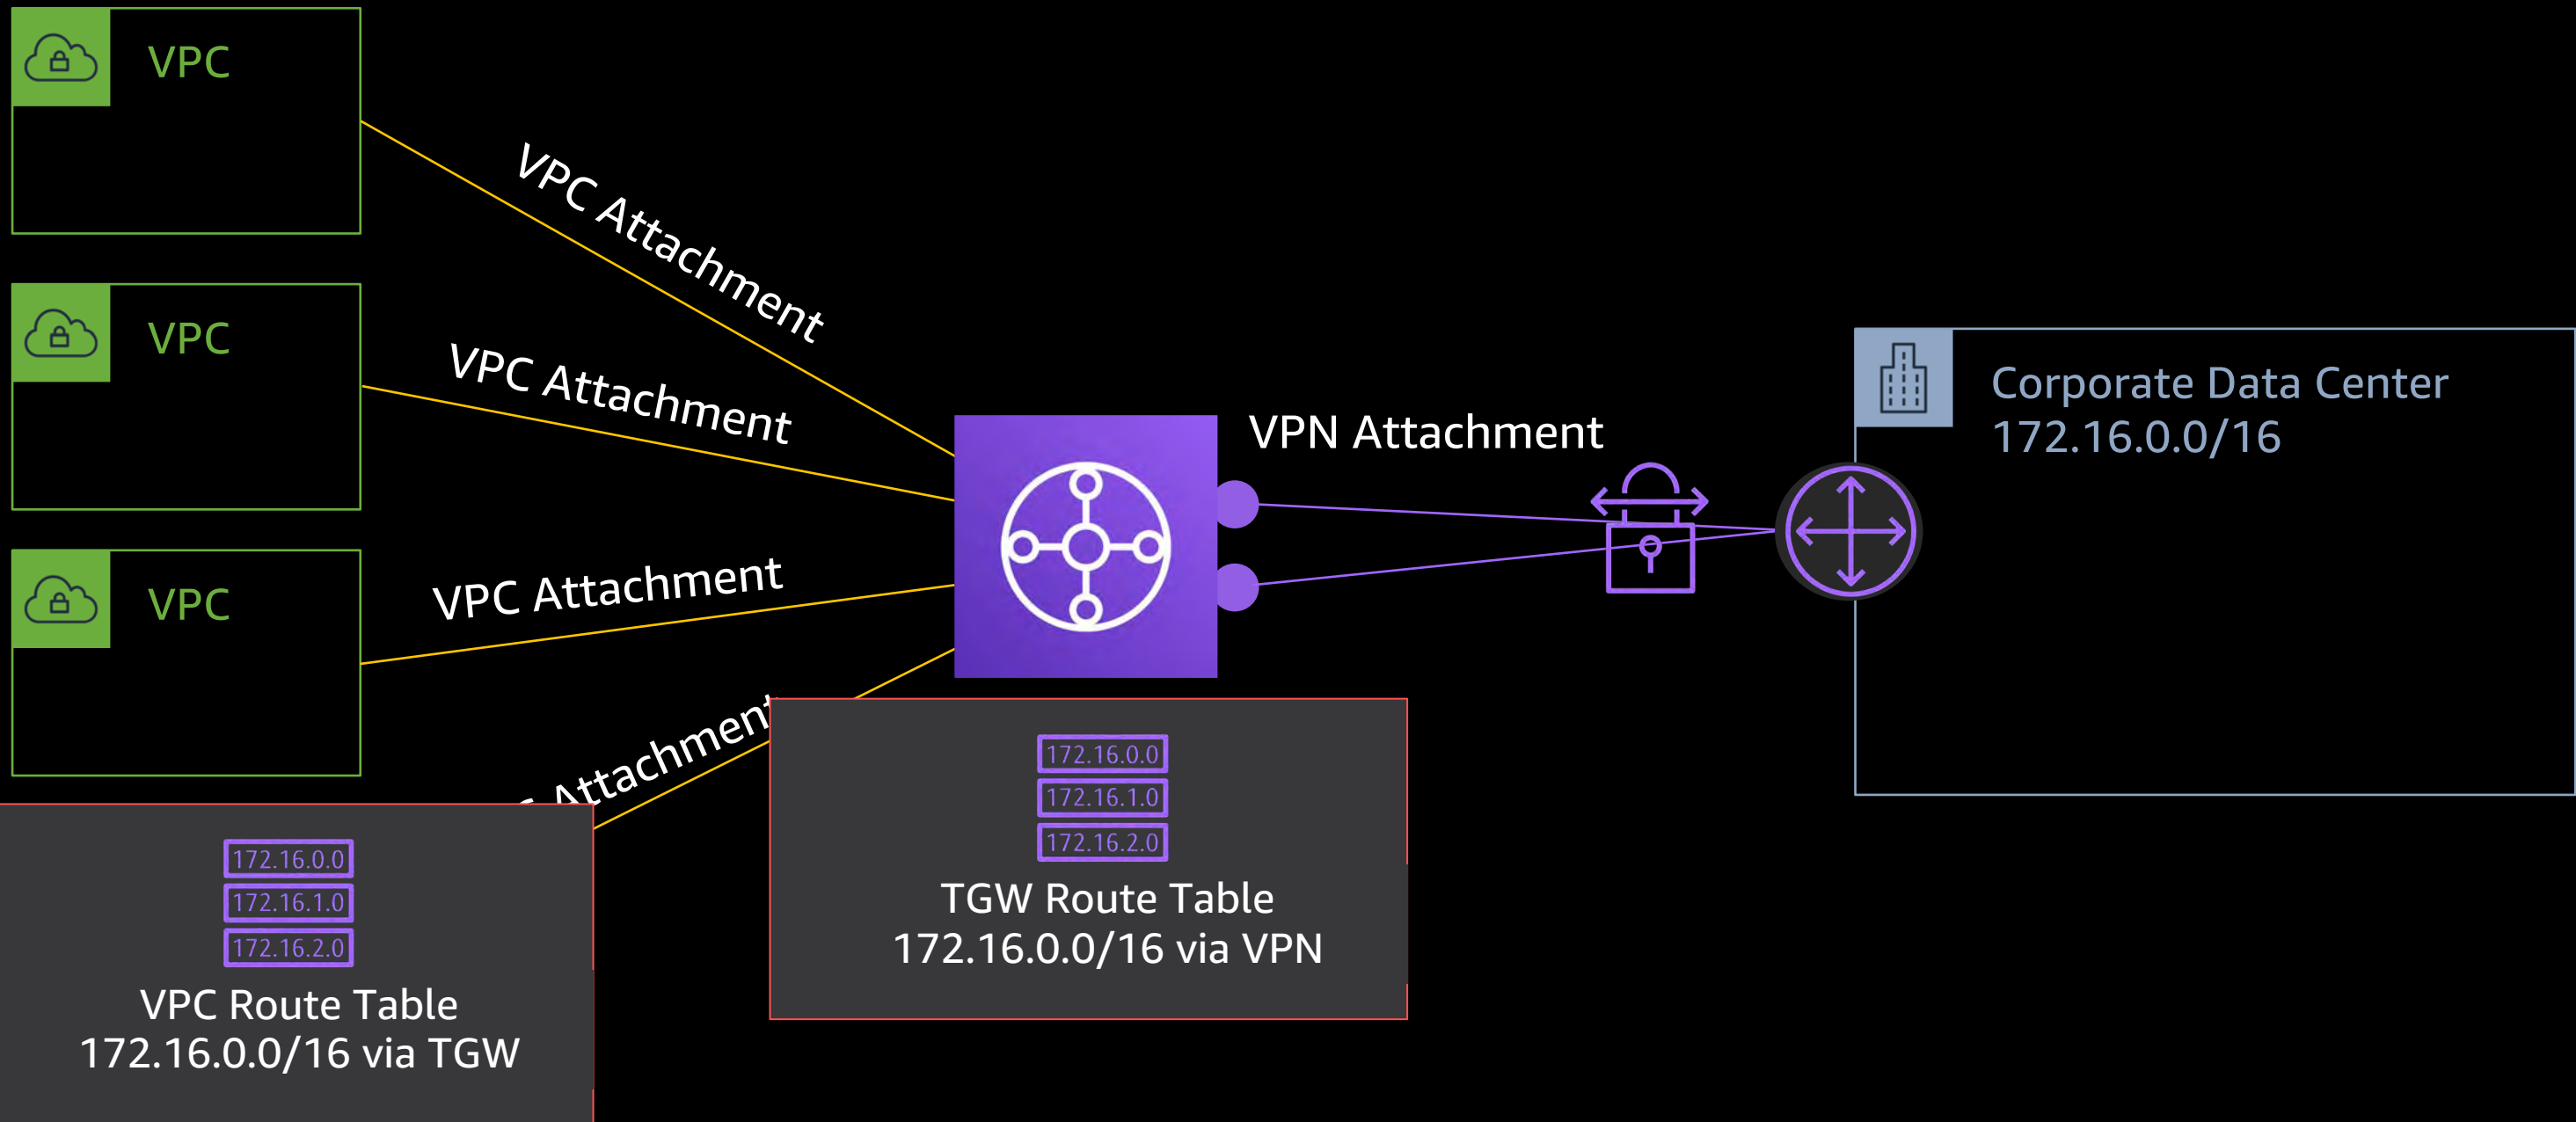
### TGW Route Table(s)

	Associations	Propagations	Routes
RT1	Llama from X	Llama from X	10.1.0.0/16 via X
	Pegasus from Y	Pegasus from Y	10.2.0.0/16 via Y
		On-prem from Q	172.16.0.0/16 via Q

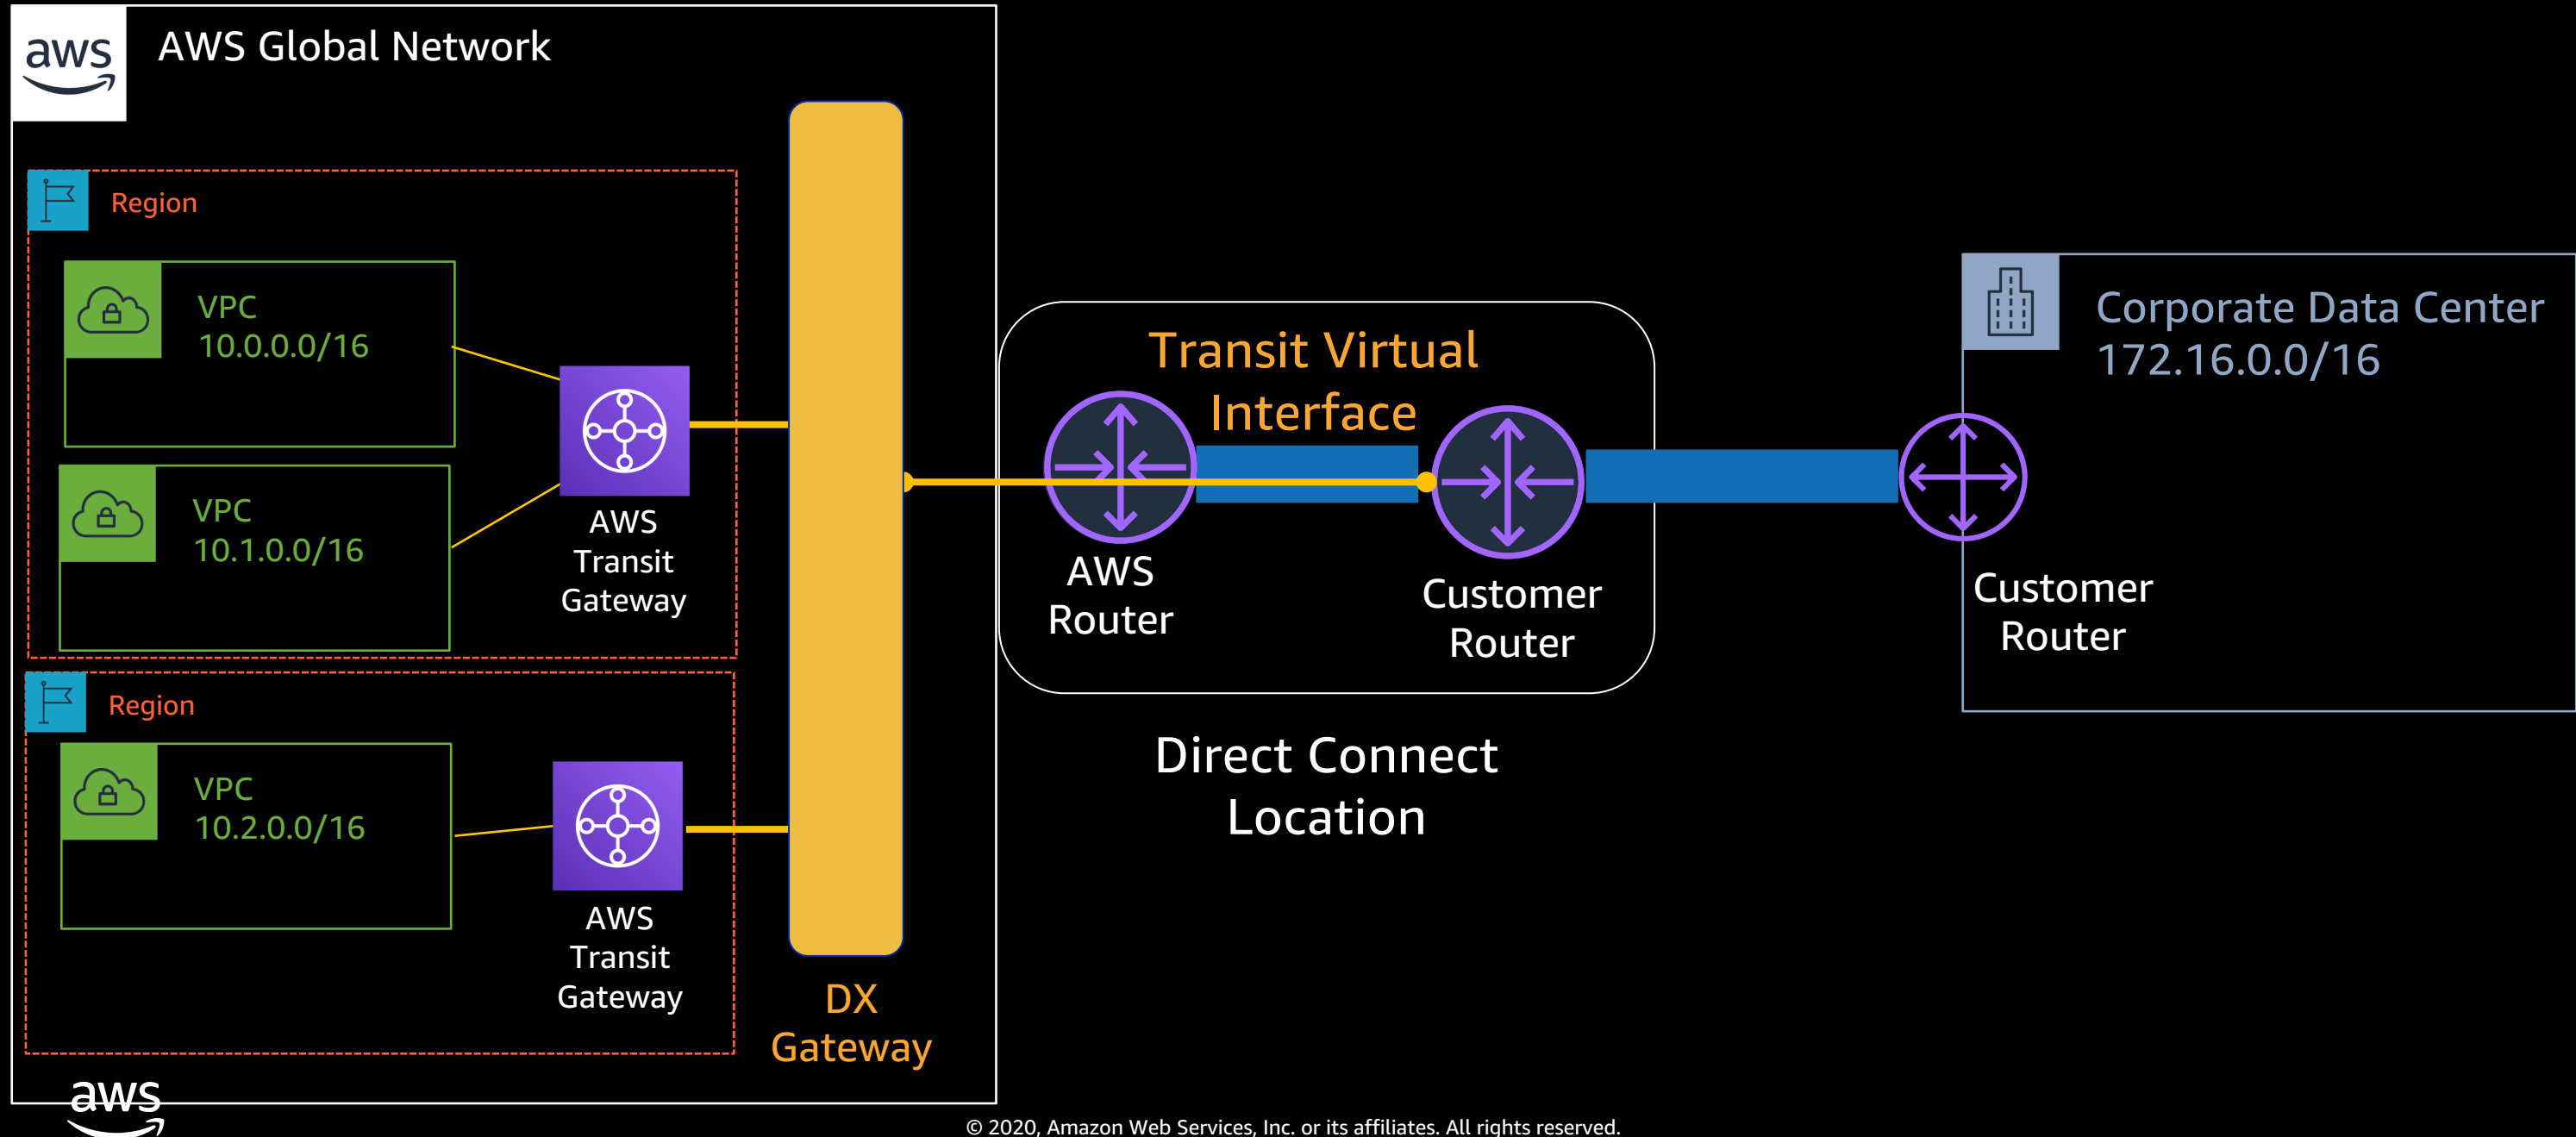
	Associations	Propagations	Routes
RT2	Barry from Z	Barry from Z	10.3.0.0/16 via X
		On-prem from Q	172.16.0.0/16 via Q

	Associations	Propagations	Routes
RT3	On-prem from Q	On-prem from Q	172.16.0.0/16 via Q
		Llama from X	10.1.0.0/16 via X
		Pegasus from Y	10.2.0.0/16 via Y
		Barry from Z	10.3.0.0/16 via Z

# AWS Transit Gateway with AWS site-to-site VPN



# AWS Transit Gateway with DX gateway



Transit Gateways per  
account/Transit Gateway  
attachments per Amazon VPC

5


Maximum burstable  
bandwidth per attachment

50 Gbps

Maximum bandwidth  
per VPN connection

\* 1.25 Gbps

\*With ECMP, you can distribute traffic over multiple tunnels,  
e.g., 8 tunnels = 10 Gbps



Routes per AWS  
Transit Gateway

10,000

**WOW!**

Number of AWS Transit Gateway  
attachments per Region per account

5,000

# Cross-Region connectivity?

AWS Transit Gateway allows for cross  
Region peering



...and there's more



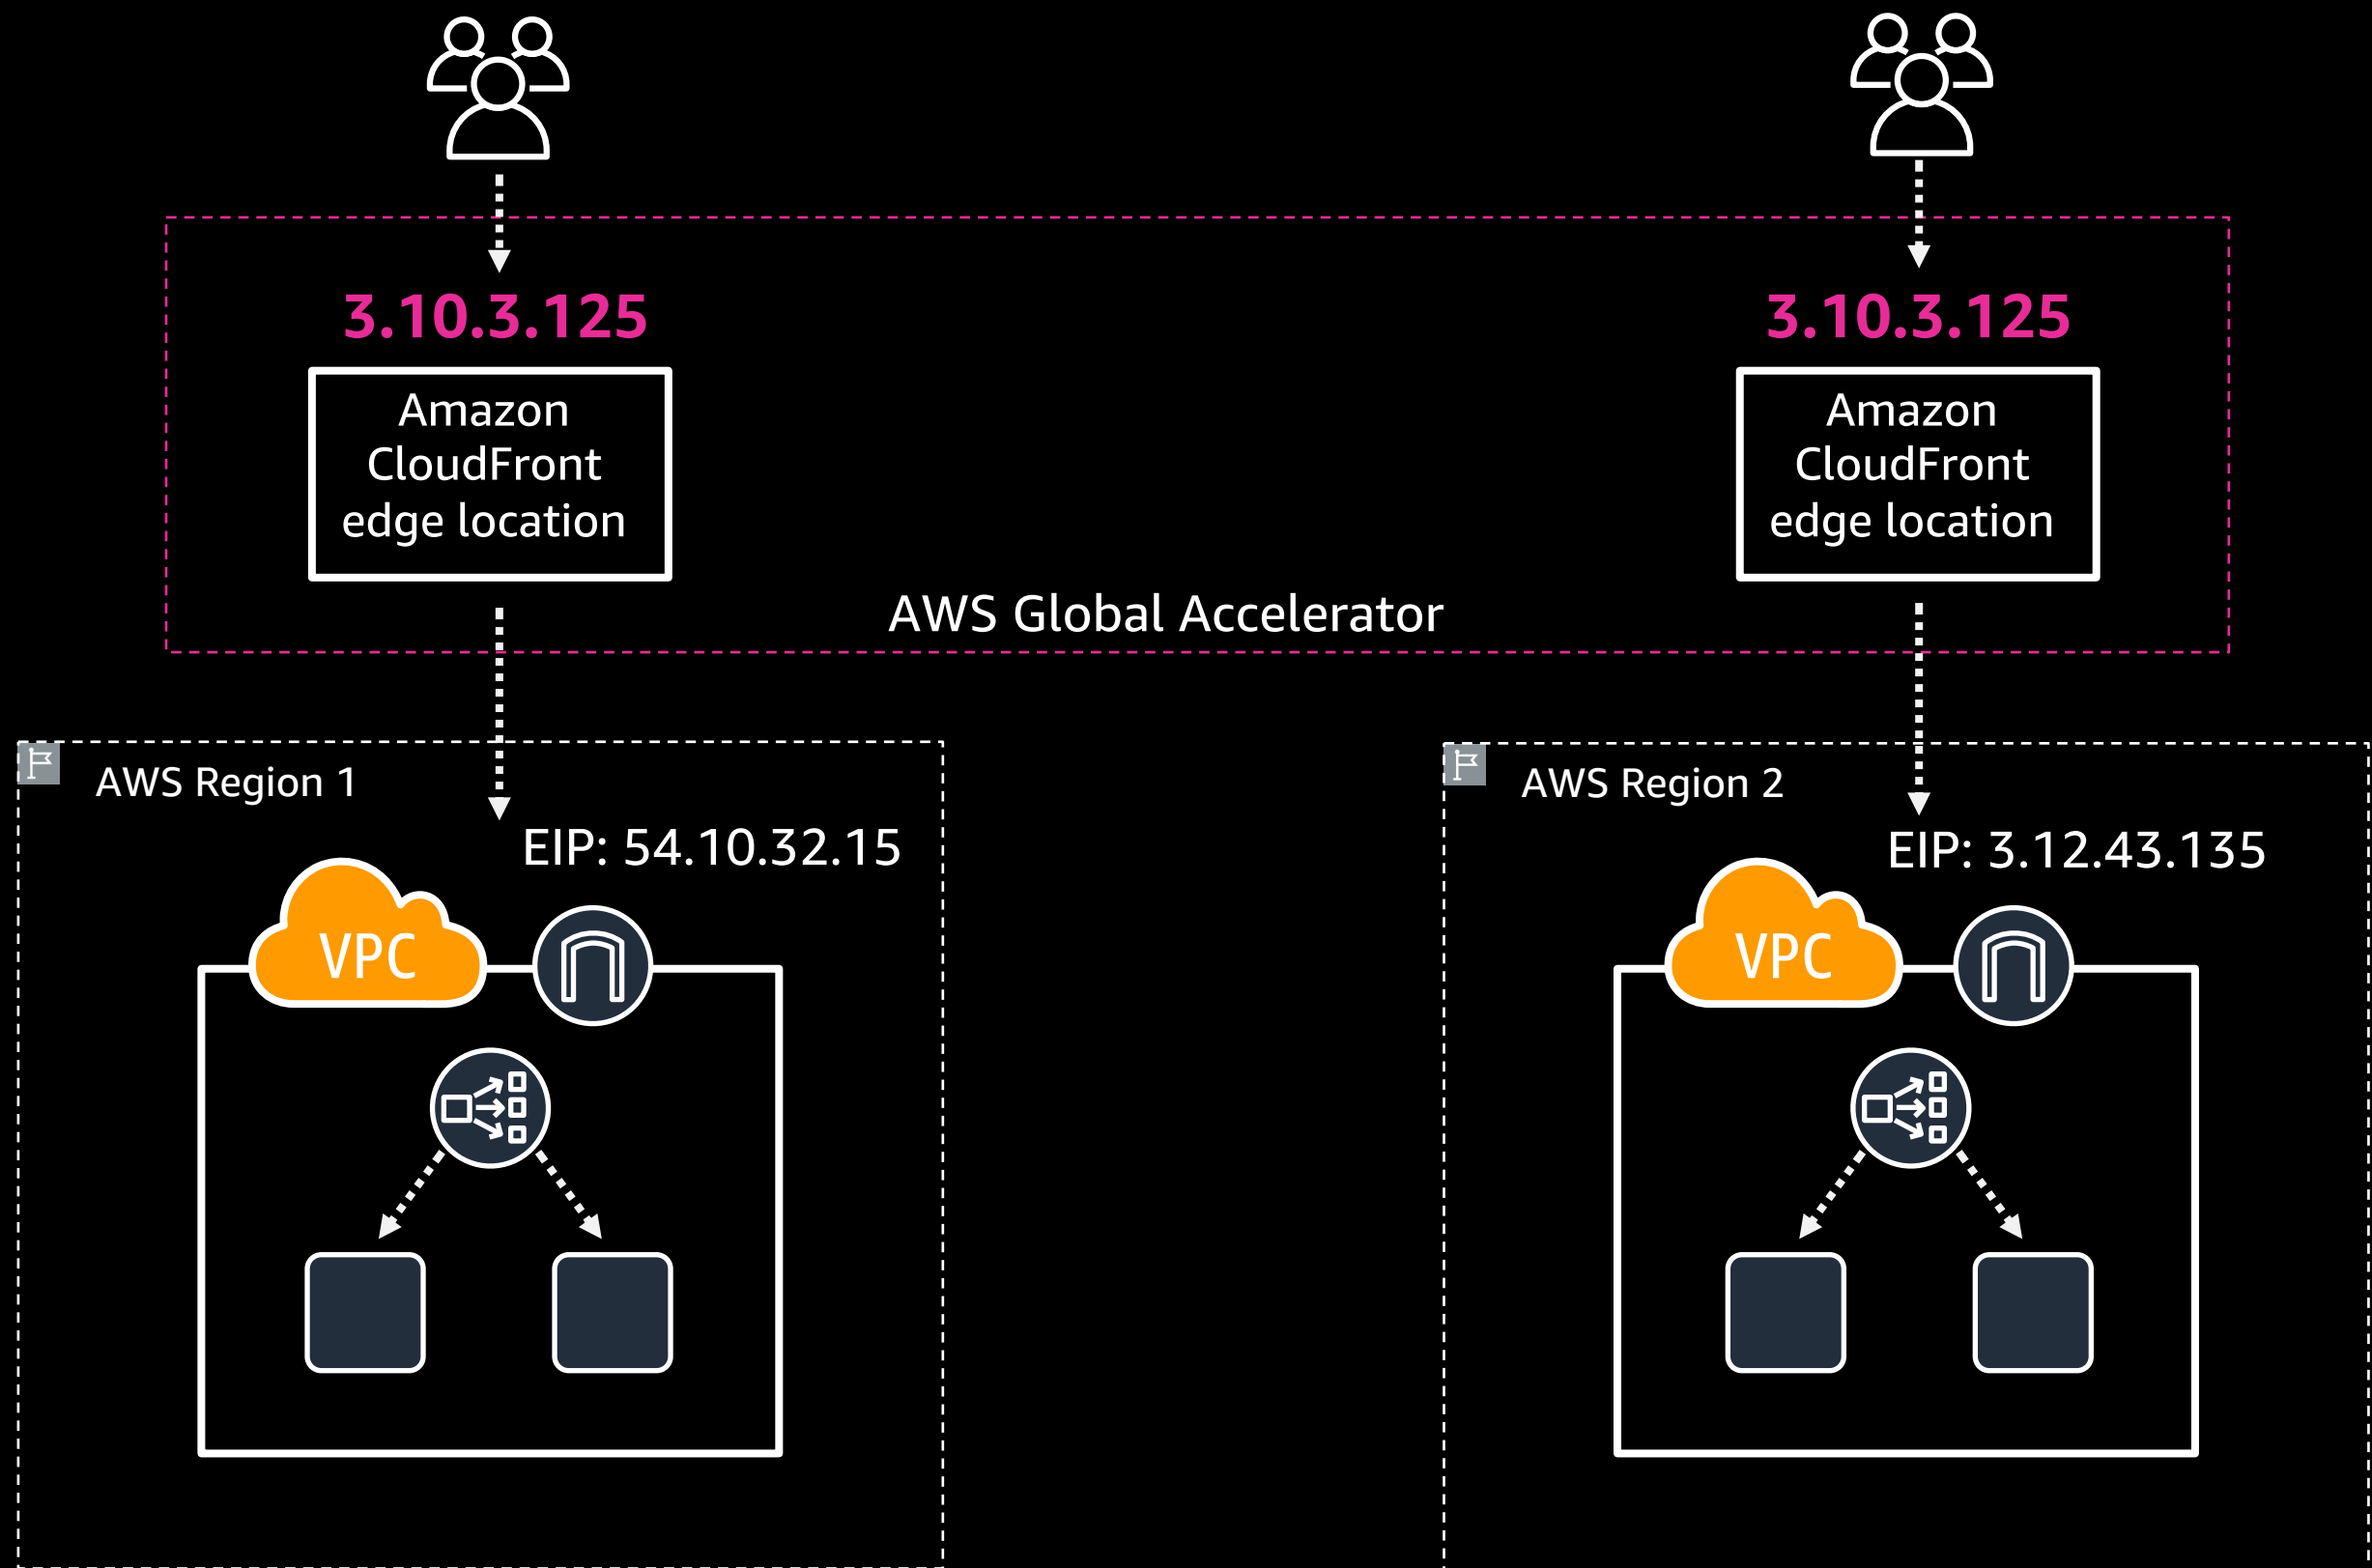


# AWS Global Accelerator

# Before



# After



## AWS global network

Traffic routed through AWS Global Accelerator traverses AWS global network (instead of the public internet)

## Client state

Applications can keep state, with connections routed to the same endpoint, after initial connection


## Static anycast IPs

AWS Global Accelerator uses static IP addresses as a fixed entry point to your applications, which are anycast from AWS edge locations

NEW! AVAILABLE TODAY


# AWS Accelerated Site-to-Site VPN

High availability and improved performance of Site-to-Site VPN




---

Securely connect  
multiple sites



---

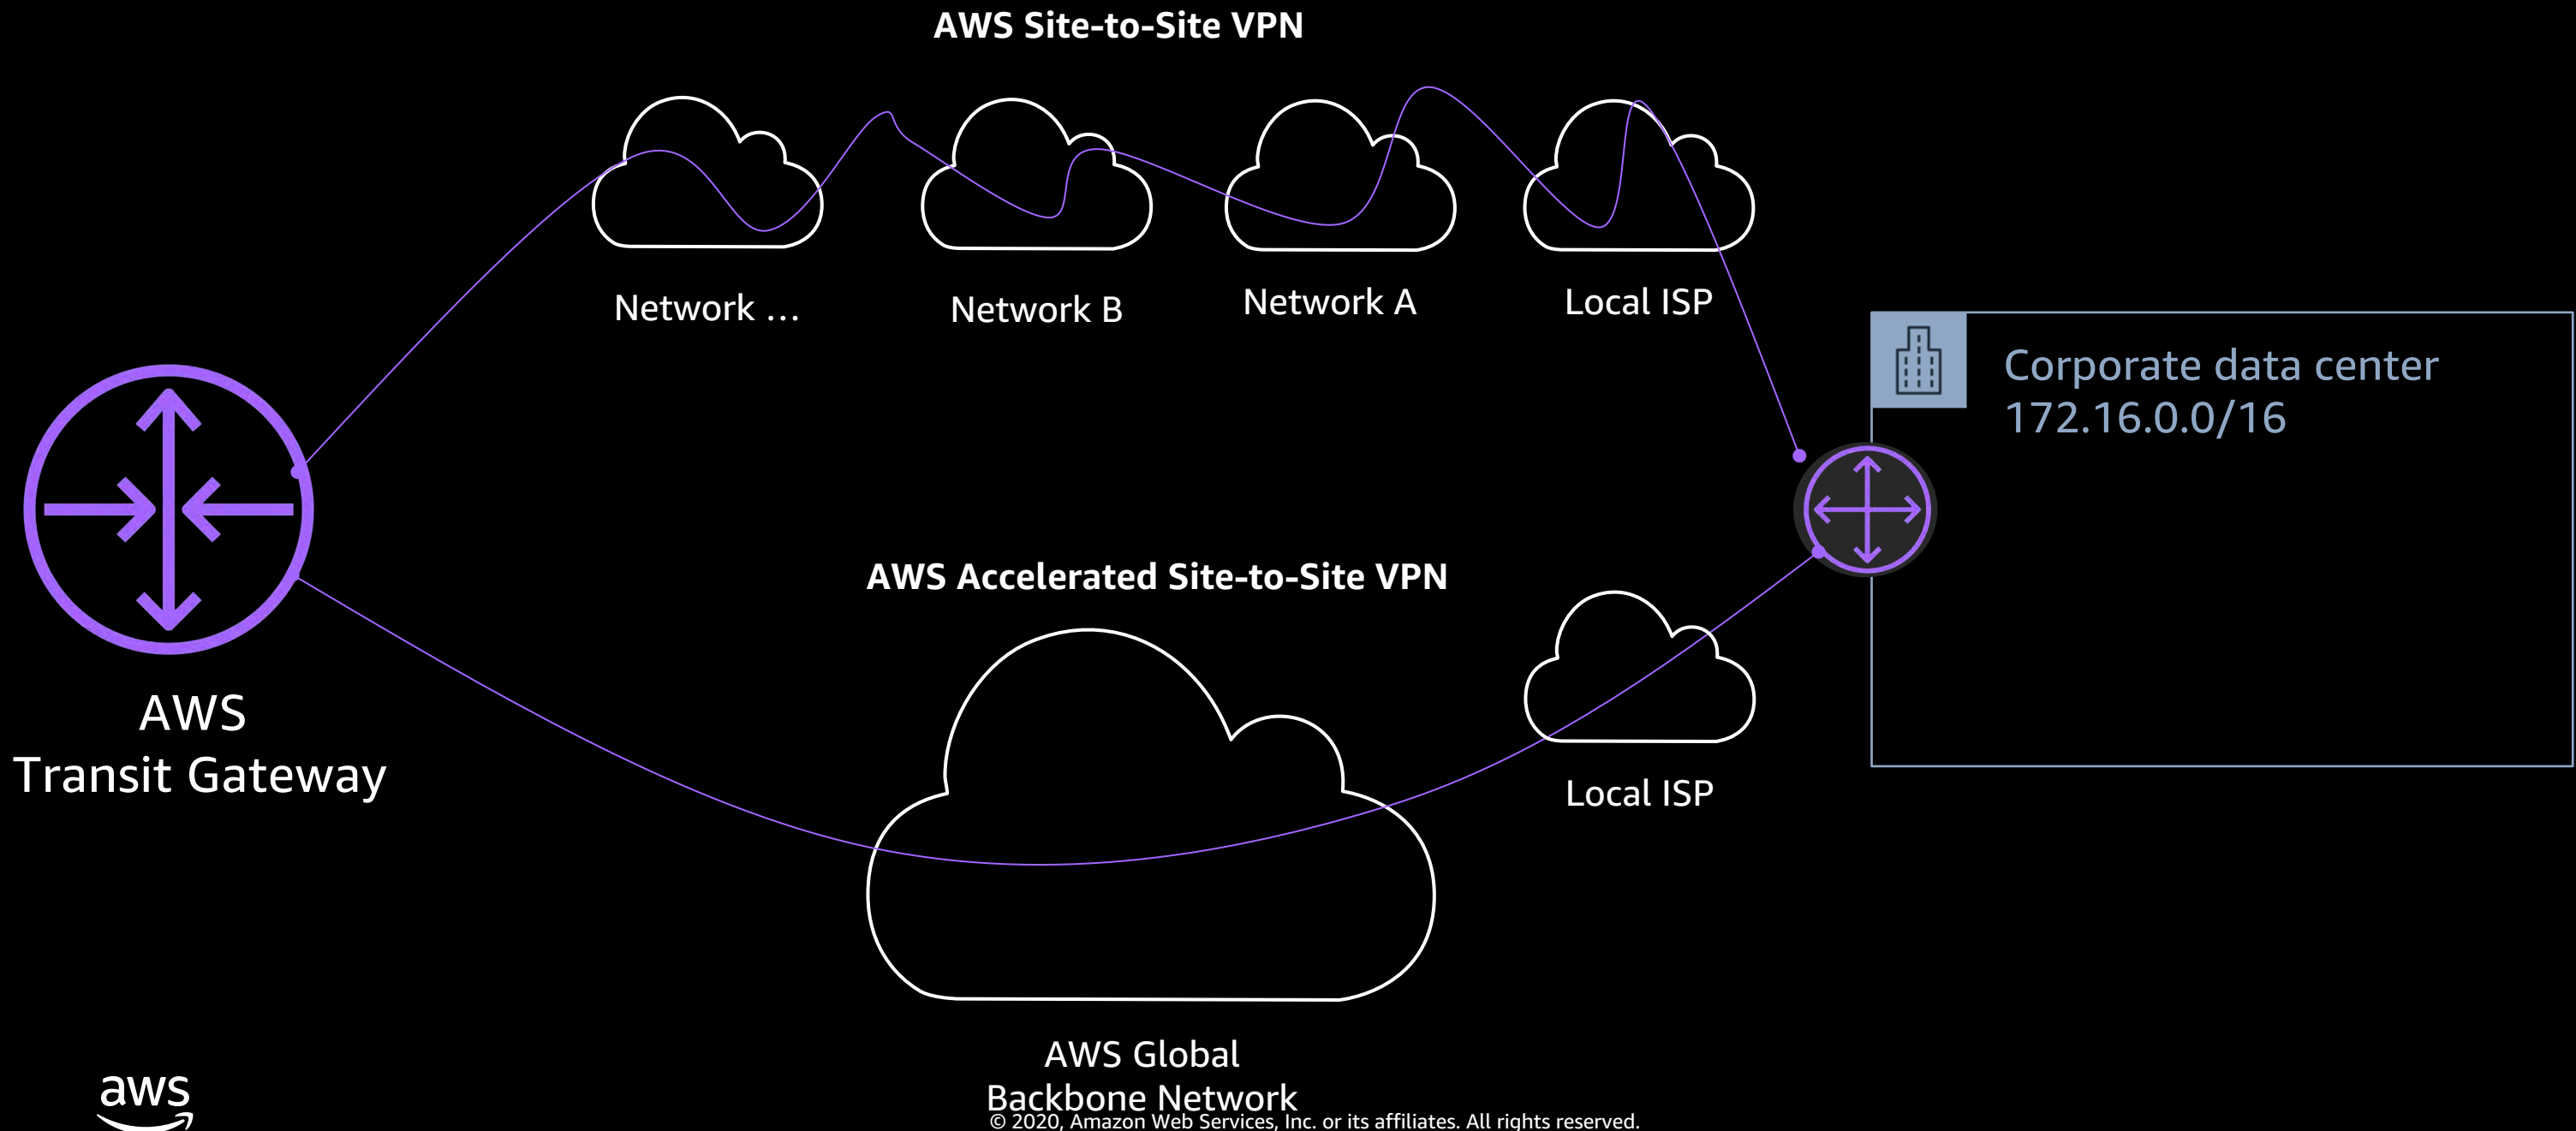
Improve performance of  
your VPN connections



---

Highly  
available

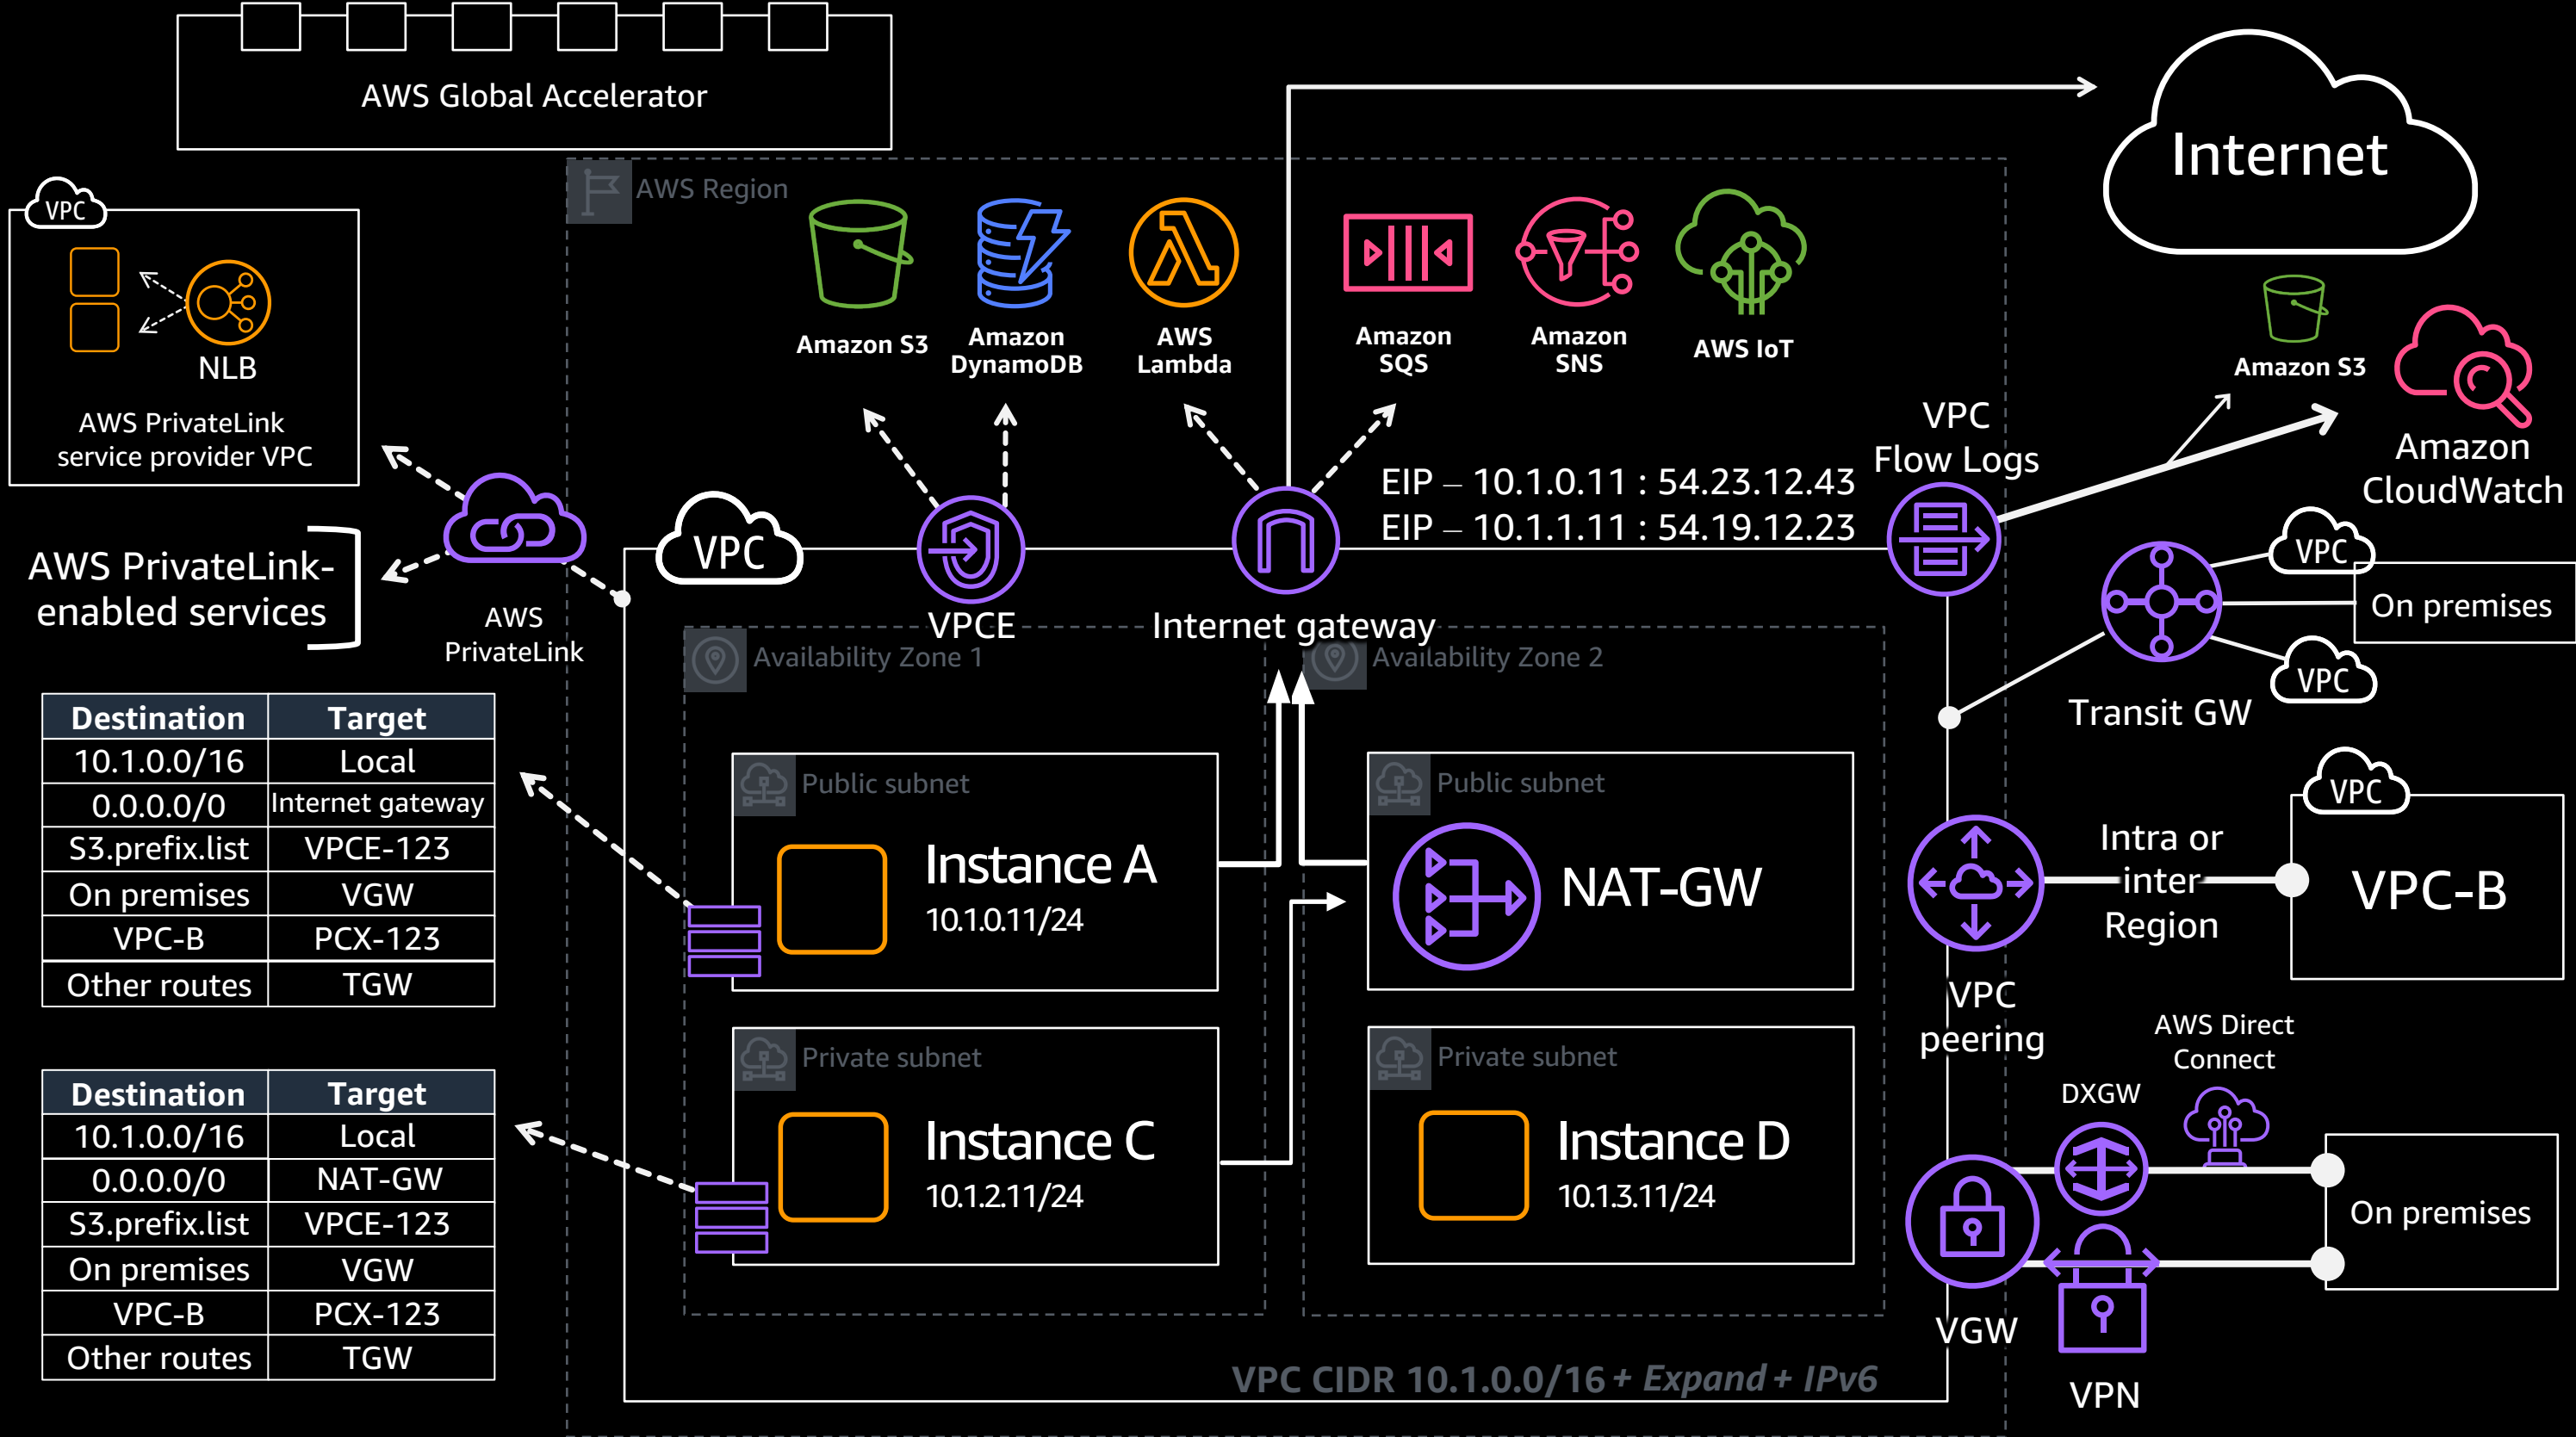
# AWS Accelerated Site-to-Site VPN





Let's bring it all back together...





# Thank you!

Francois van Rensburg  
rensburf@amazon.com

 @zaFvRensburg

 /francoisvr

Joseph Mutua  
mutuaj@amazon.com

