Advanced Networking on AWS

From One to Many - Evolving VPC Design

Nick Matthews, Partner Solutions Architect

January 2017

amazon web services

What to expect from this session:

- VPC Design
- Scalable and Available NAT
- Multiple VPCs
- VPC Endpoints
- Shared Service Hubs
- Transit VPC
- Direct Connect

Disclaimer:

Do Try This at Home!

Assuming you've heard of…

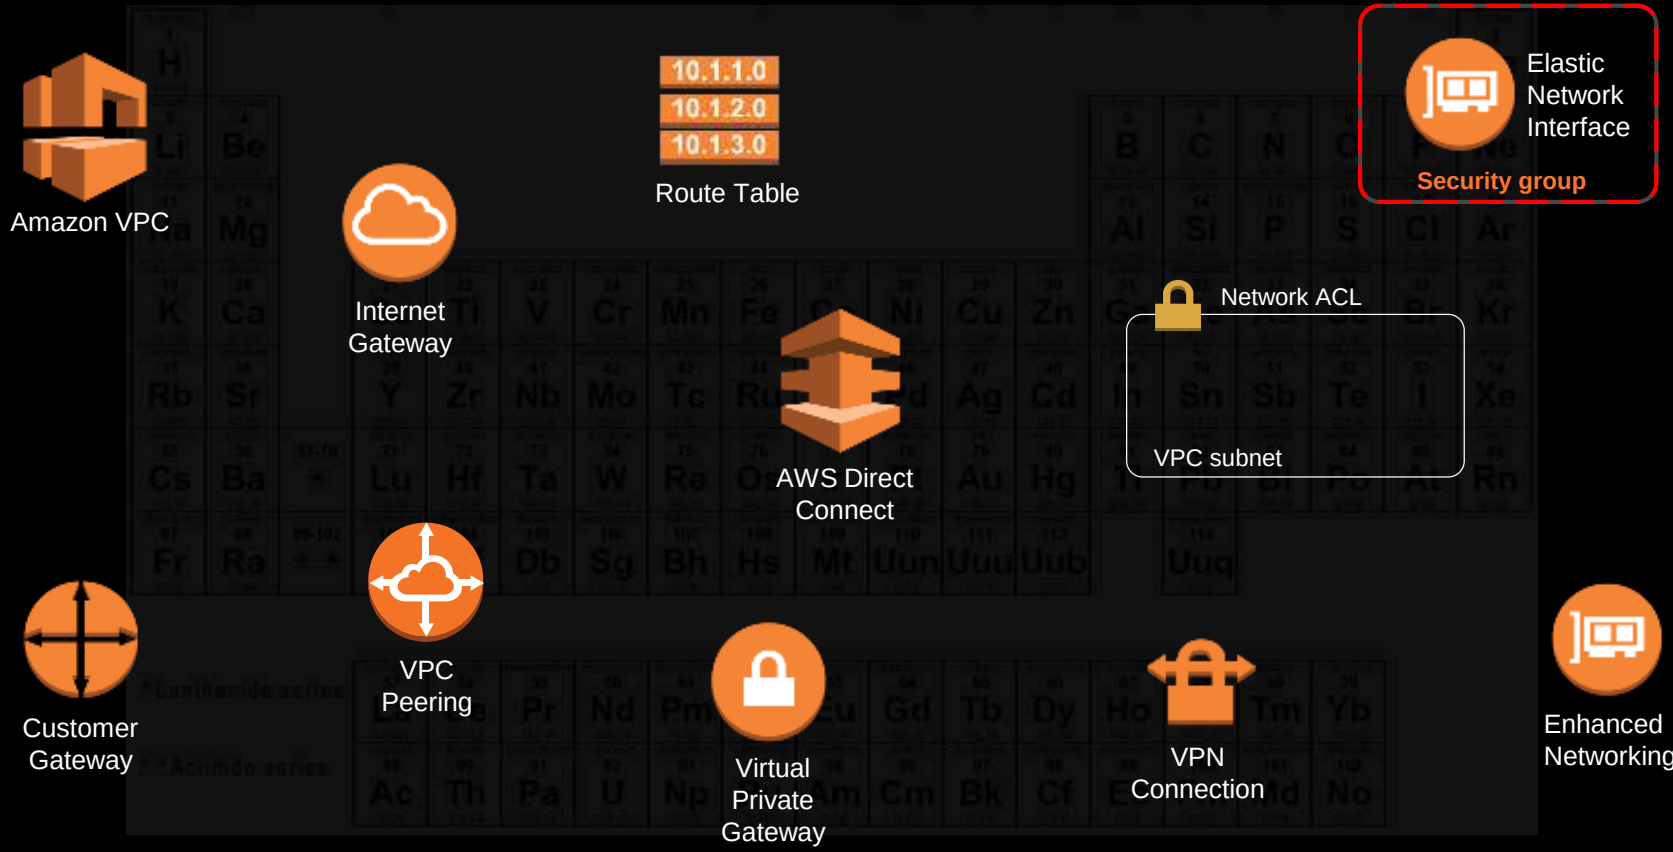

Related re:Invent Sessions

NET201 – Creating Your Virtual Data Center: VPC Fundamentals and Connectivity Options

NET305 – Extending Datacenters to the Cloud: Connectivity Options and Considerations for Hybrid Environments

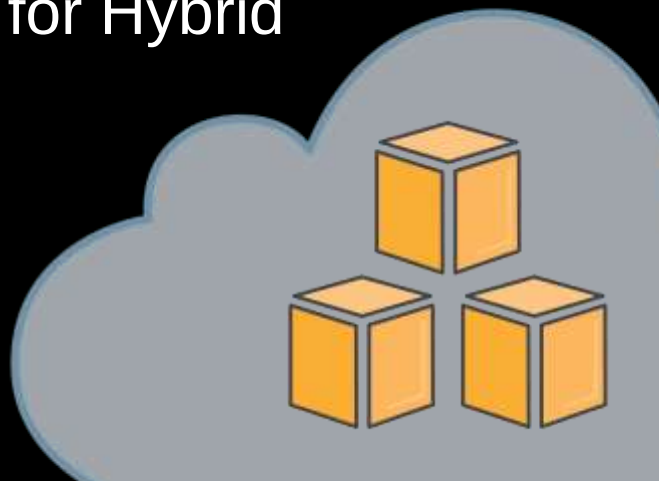

From one…

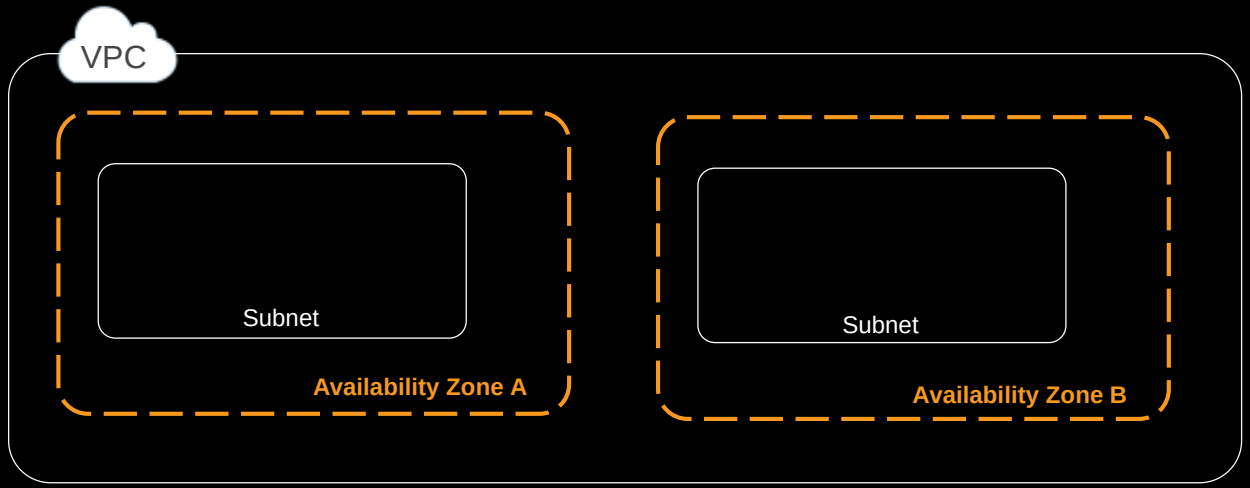

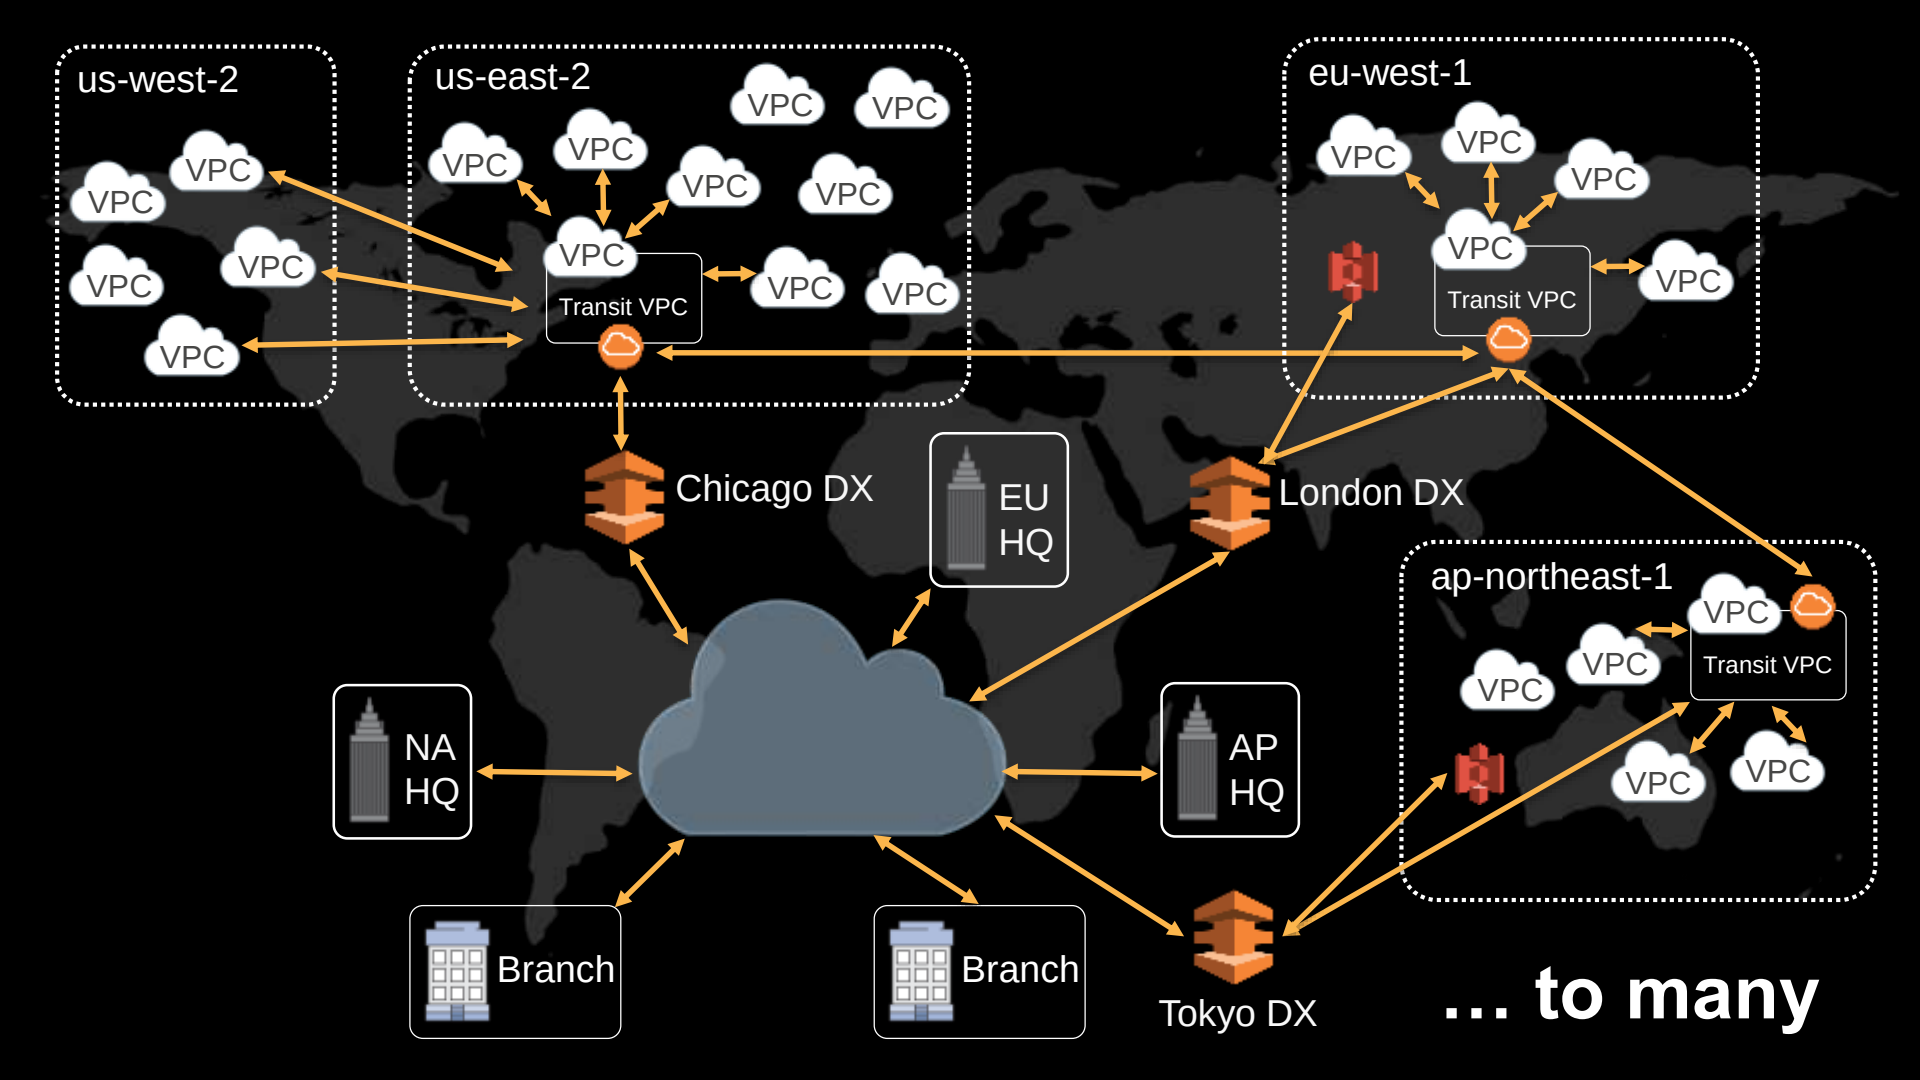
us-west-2
us-east-2
eu-west-1
VPC
Transit VPC
Chicago DX
London DX
EU HQ
NA HQ
AP HQ
ap-northeast-1
Branch
Branch
Tokyo DX
… to many

Choose a CIDR

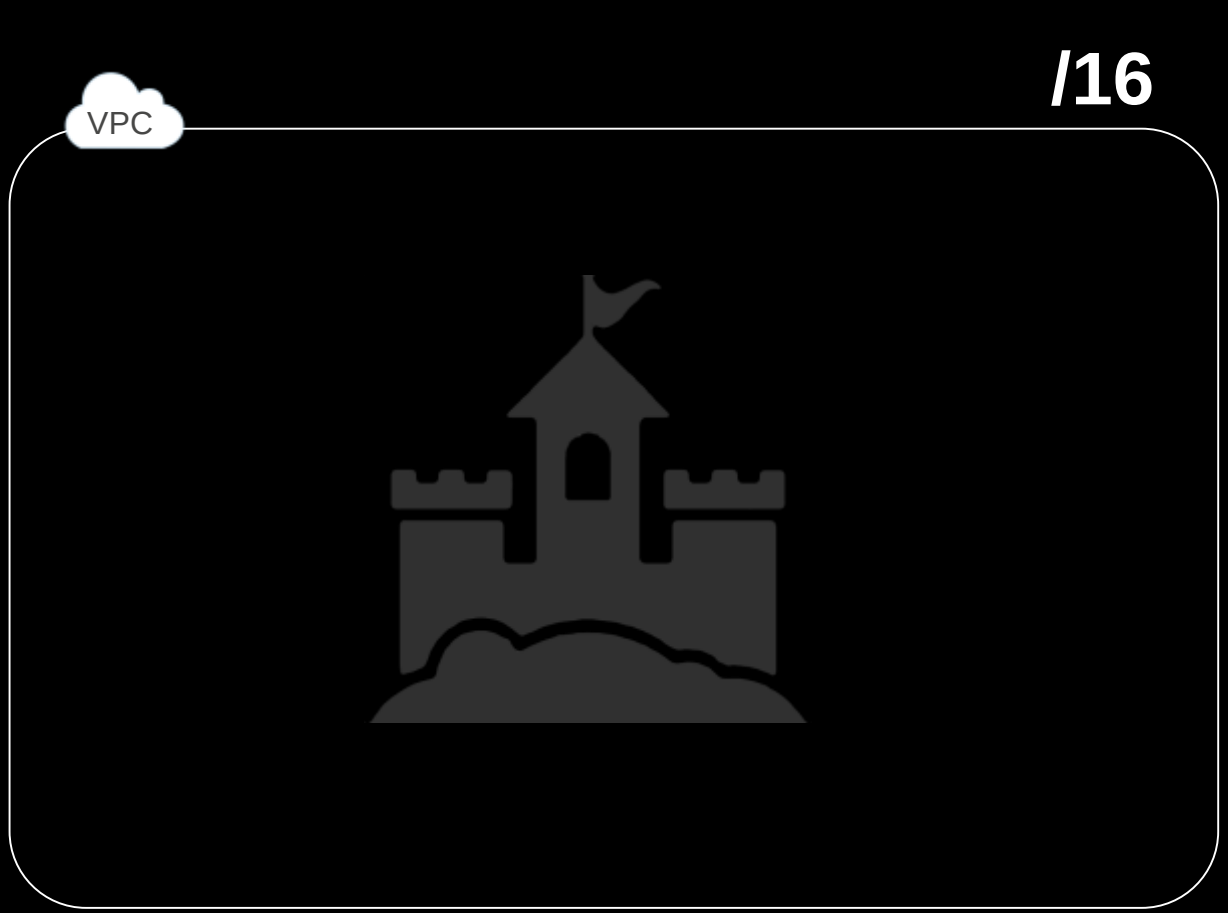

- CIDR fixed on VPC creation
- /16 down to /28
- Go Big

VPC IPv4 space design

- Plan for expansion to additional Availability Zones or regions
- Consider connectivity to corporate networks
- Don’t overlap IP space
- Save space for the future
- IPv4 space is required

Create subnets

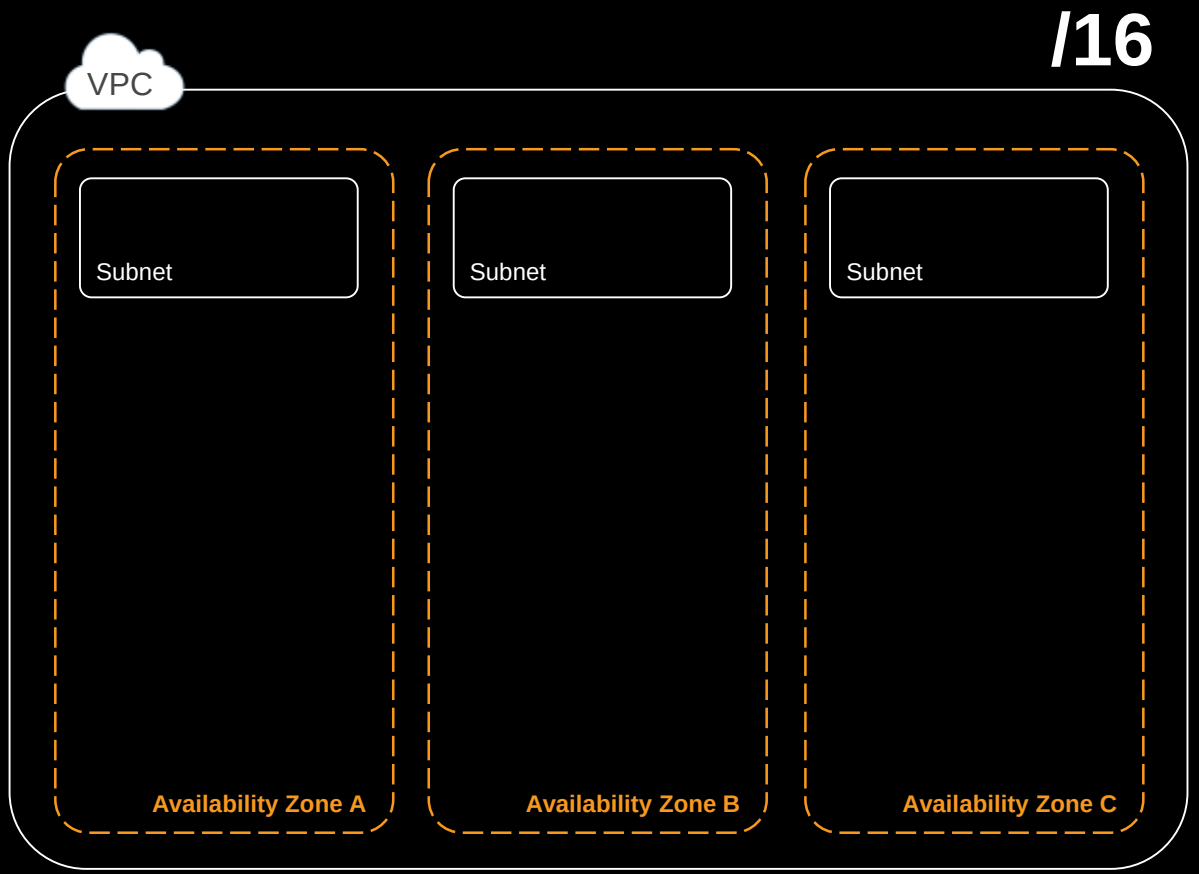

- Even distribution of IP space across AZs
- Use at least 2 AZs
- Subnets are AZ specific
- How big? How many?

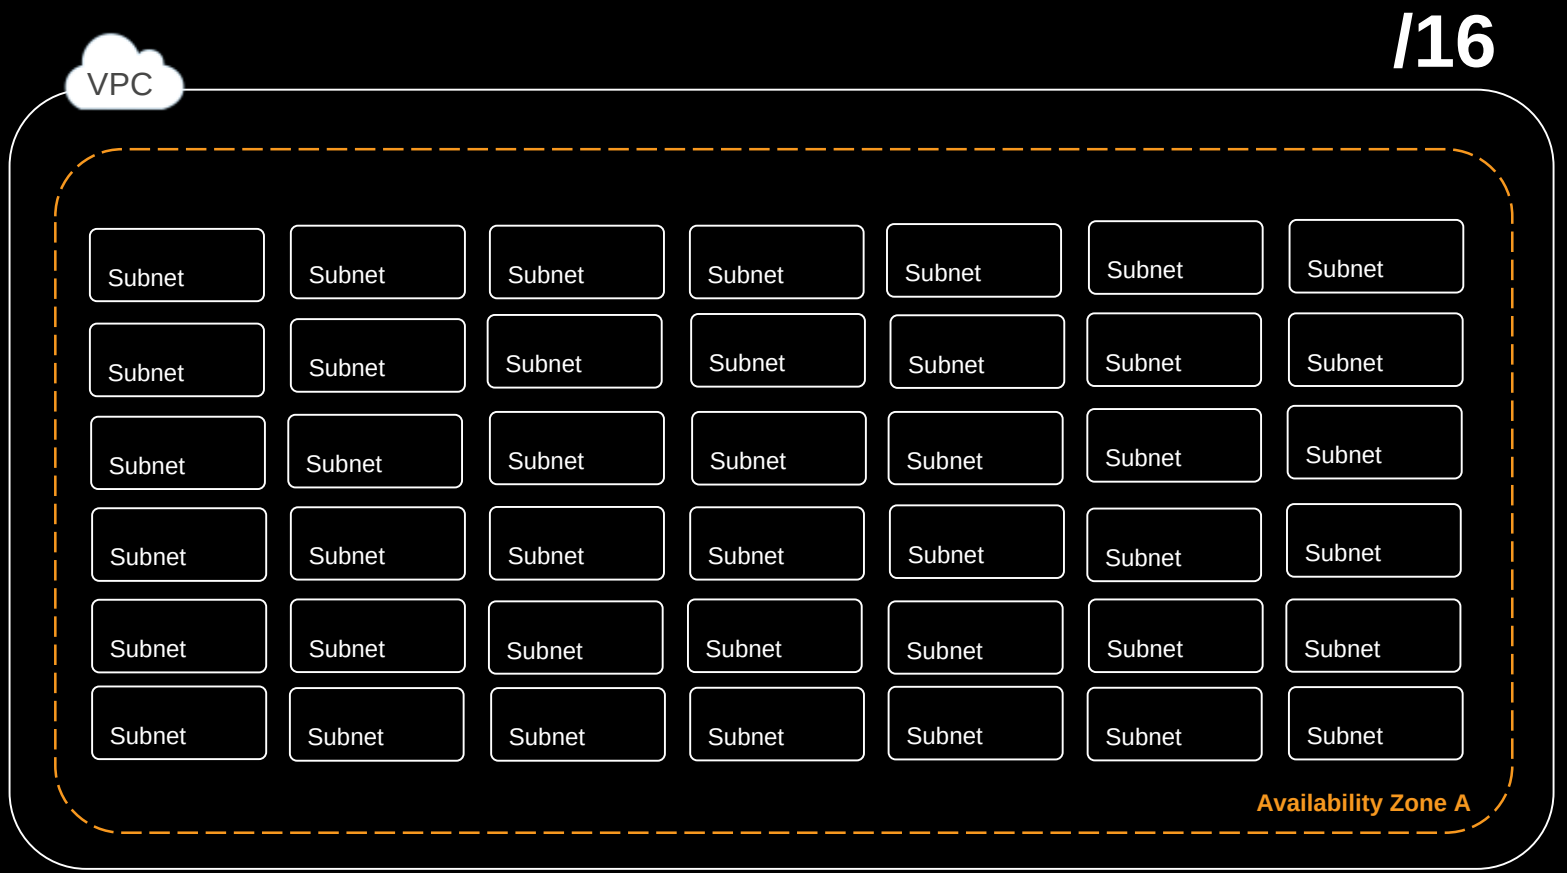
/16
VPC
Subnet
Availability Zone A

VPC subnet design

- Traditional switching limitations do not apply
- Consider large, mixed use subnets
- Use security groups to enforce isolation
- Use tags for grouping resources
- Use subnets as containers for routing policy

Related re:Invent Sessions

NET401 – Another Day, Another Billion Packets

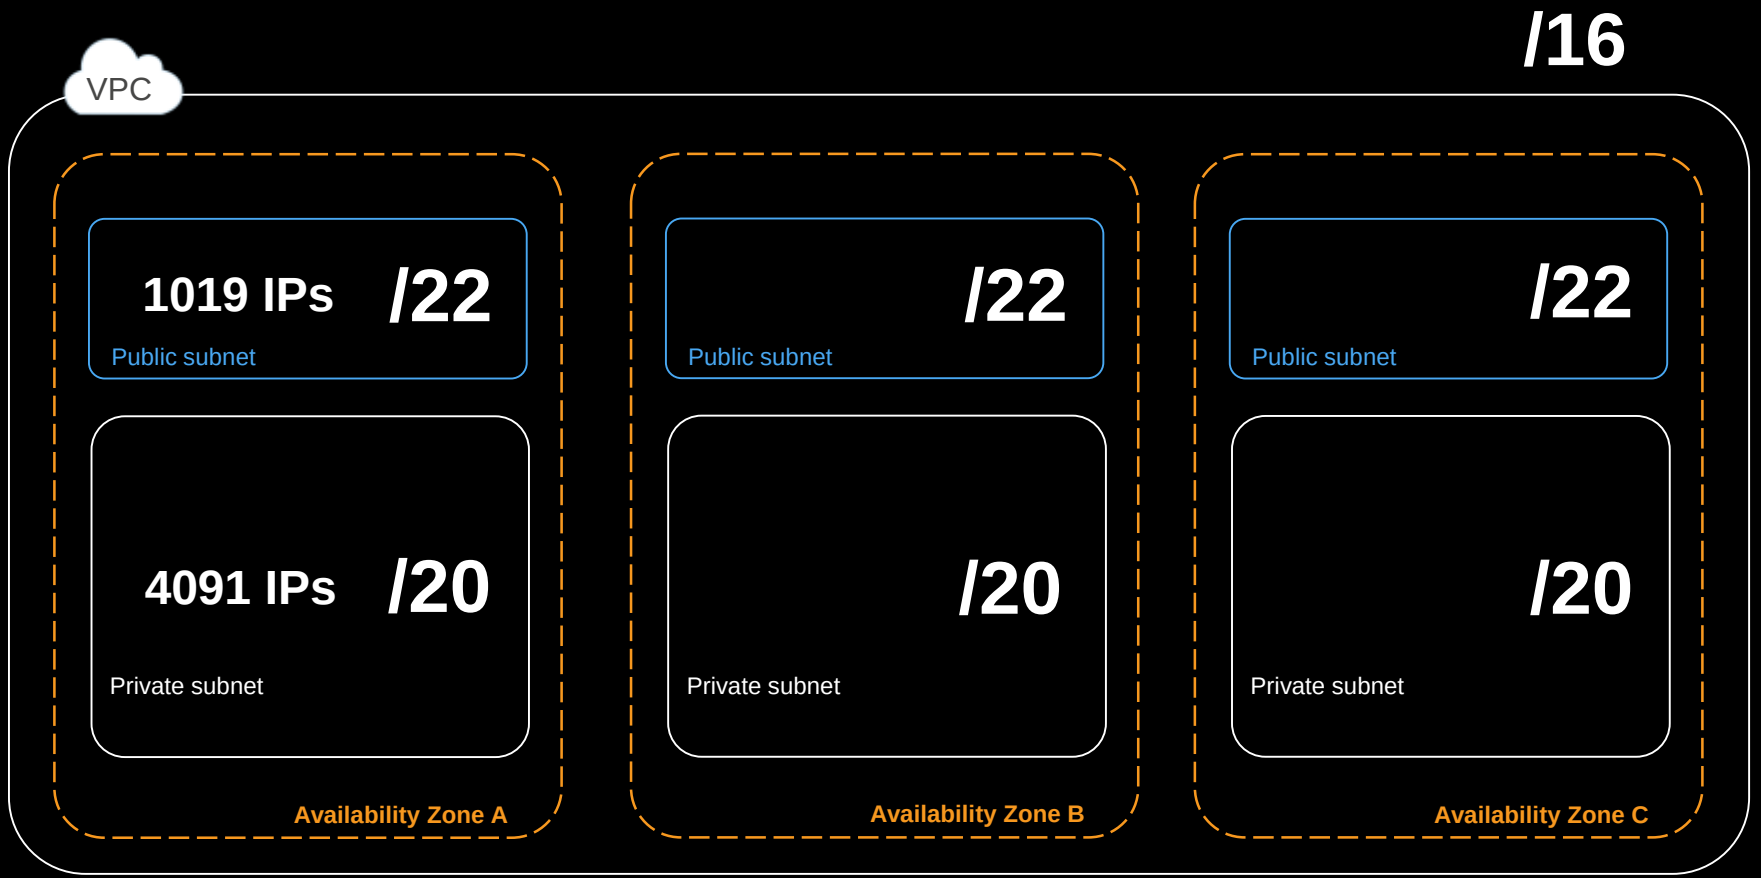
/16
VPC
1019 IPs
/22
Public subnet
4091 IPs
/20
Private subnet
Availability Zone A
/22
Public subnet
/20
Private subnet
Availability Zone B
/22
Public subnet
/20
Private subnet
Availability Zone C

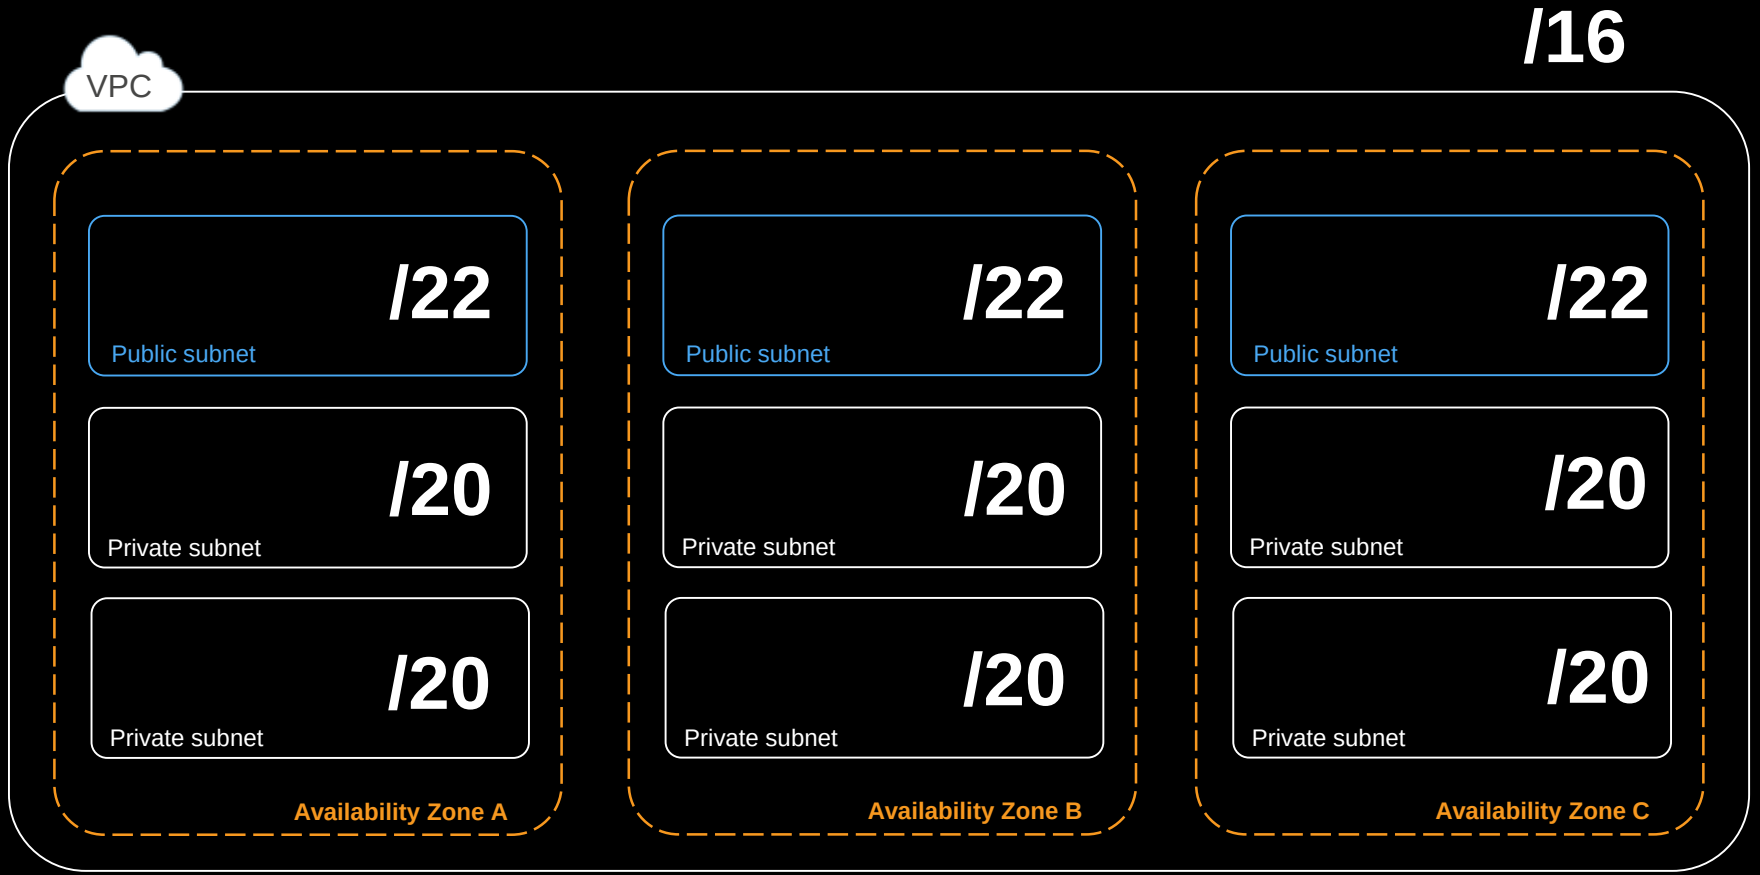
/16
VPC
/22
Public subnet
/20
Private subnet
/20
Private subnet
Availability Zone A
/22
Public subnet
/20
Private subnet
/20
Private subnet
Availability Zone B
/22
Public subnet
/20
Private subnet
/20
Private subnet
Availability Zone C

Routing Policy

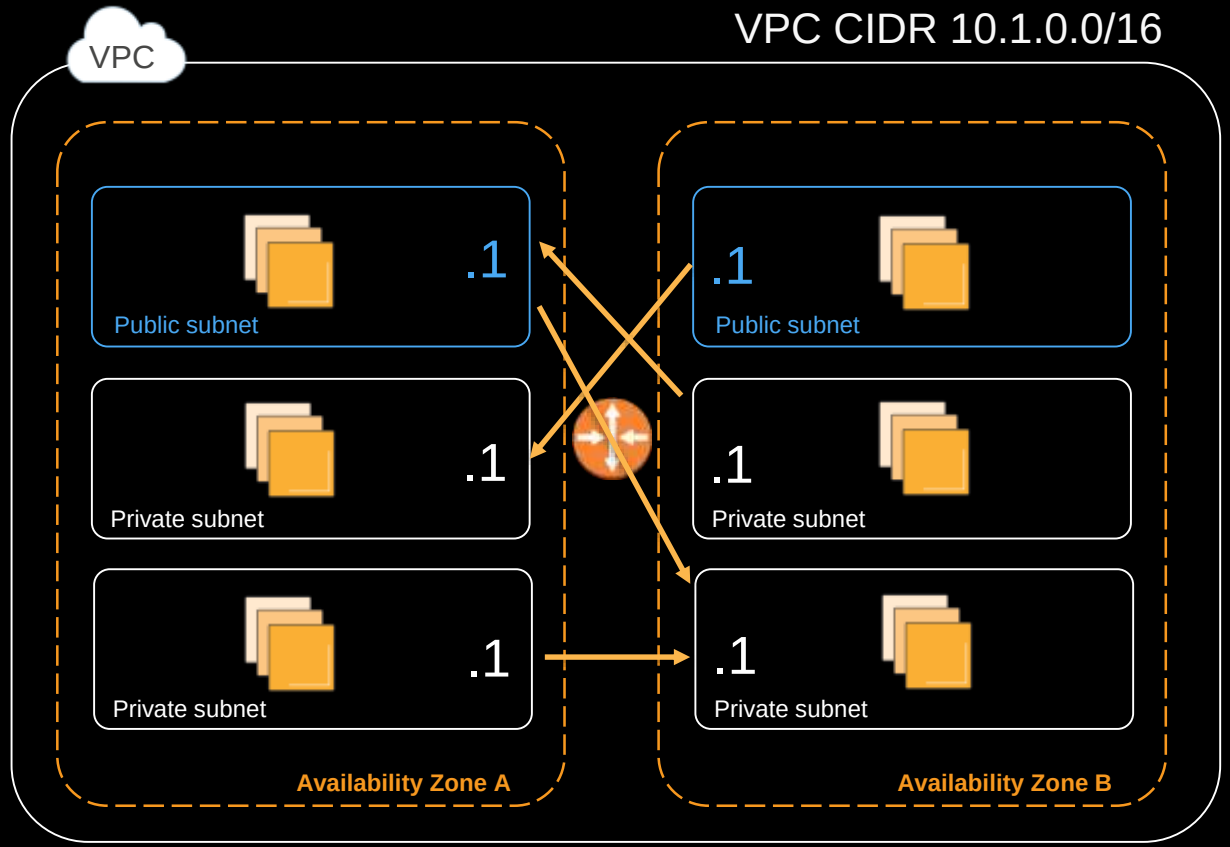

Main Route Table	
Destination	Target
10.1.0.0/16	Local

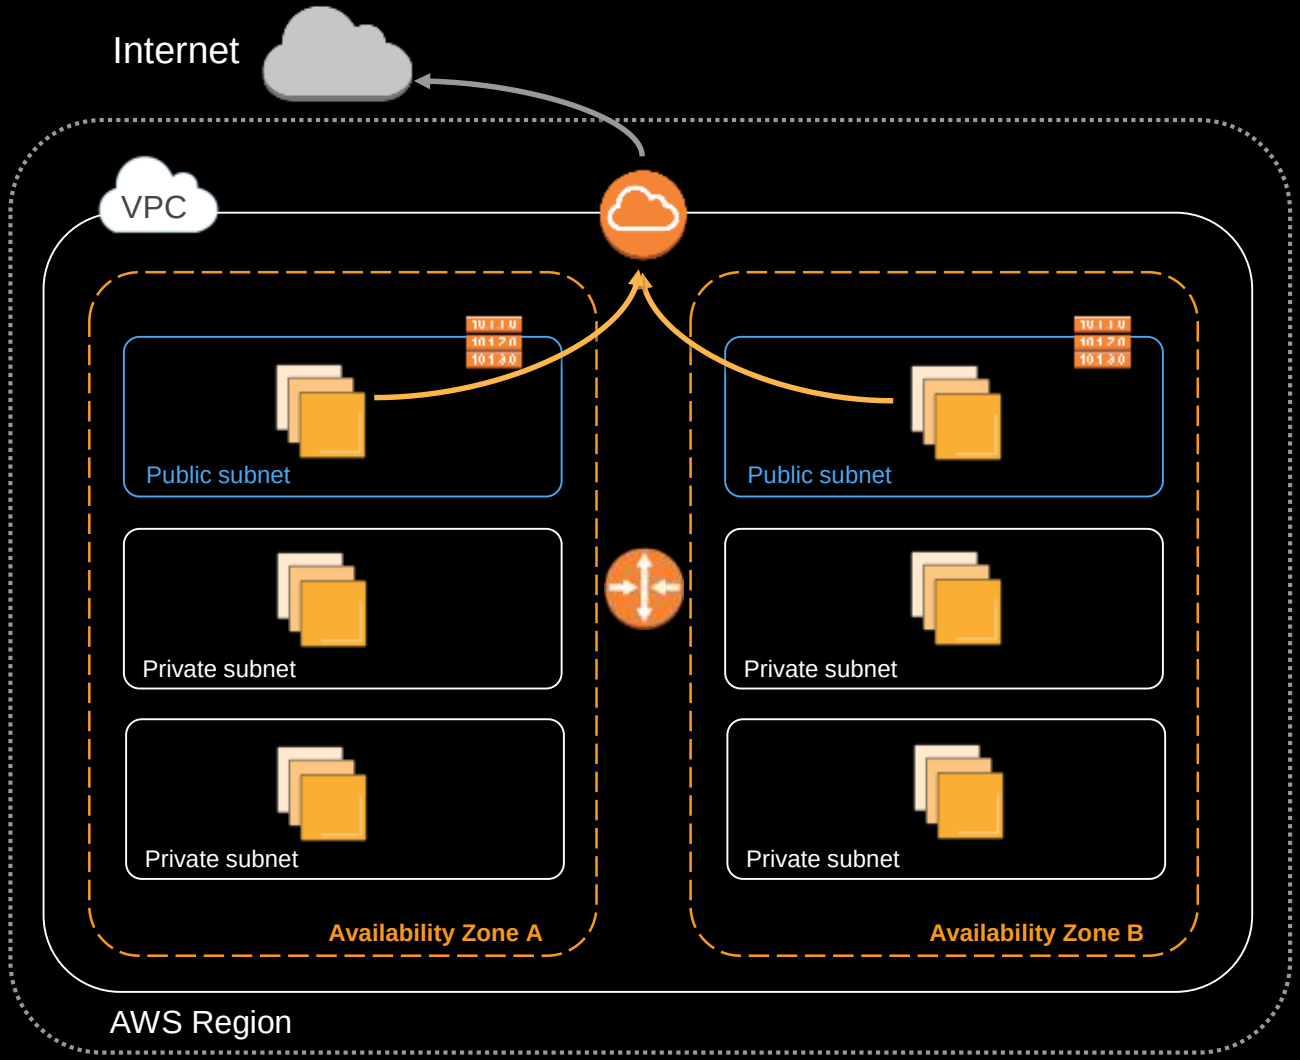

Routing Policy

Public Route Table	
Destination	Target
10.1.0.0/16	Local
0.0.0.0/0	IGW

Routing Policy

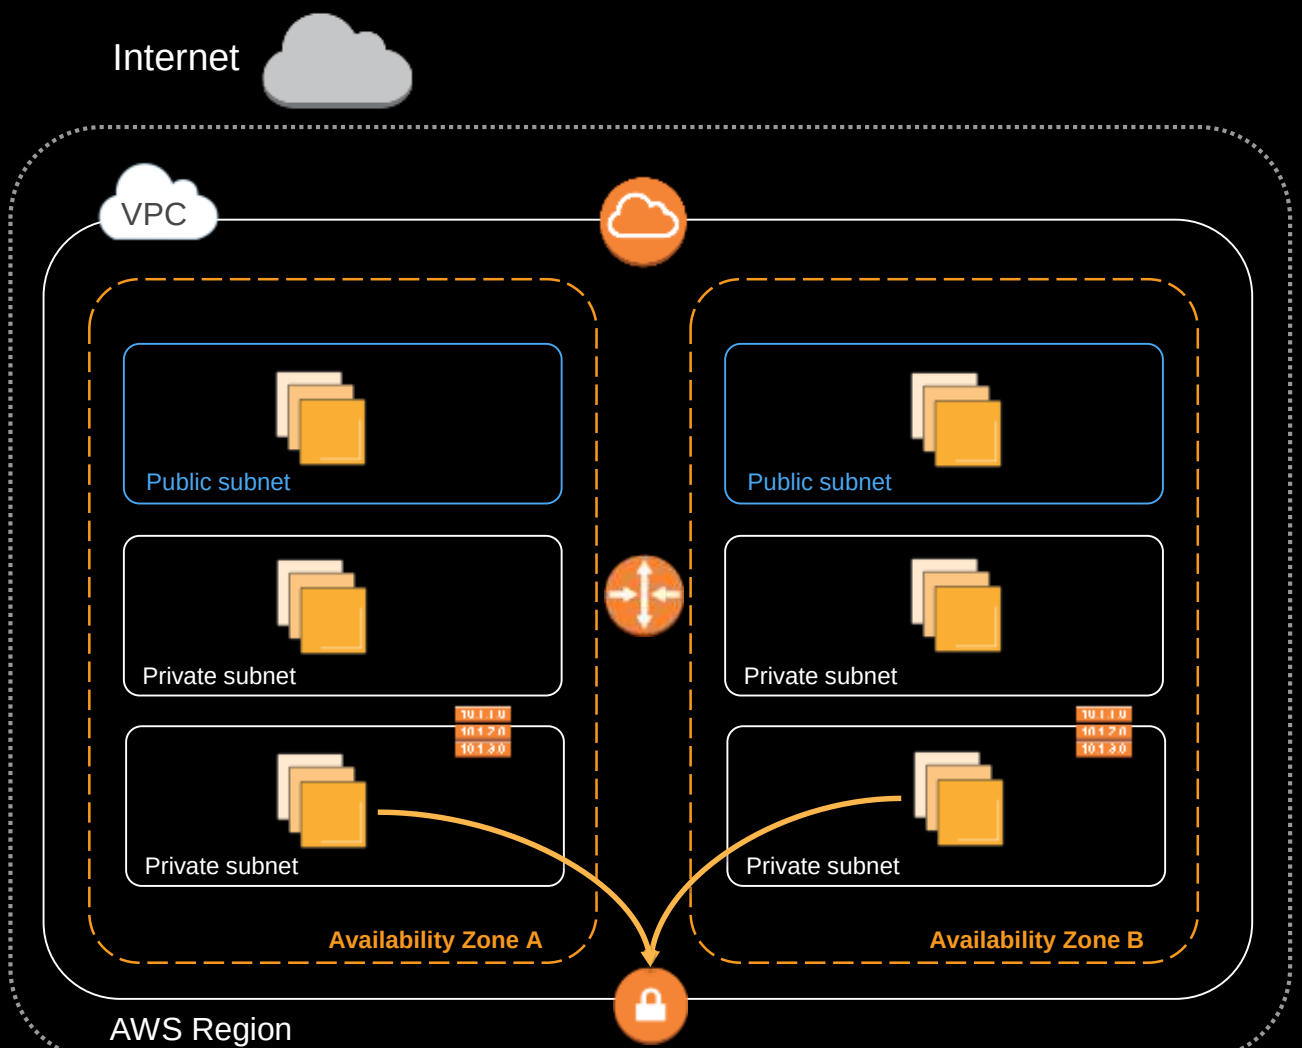

Private Route Table	
Destination	Target
10.1.0.0/16	Local
Corp CIDR	VGW

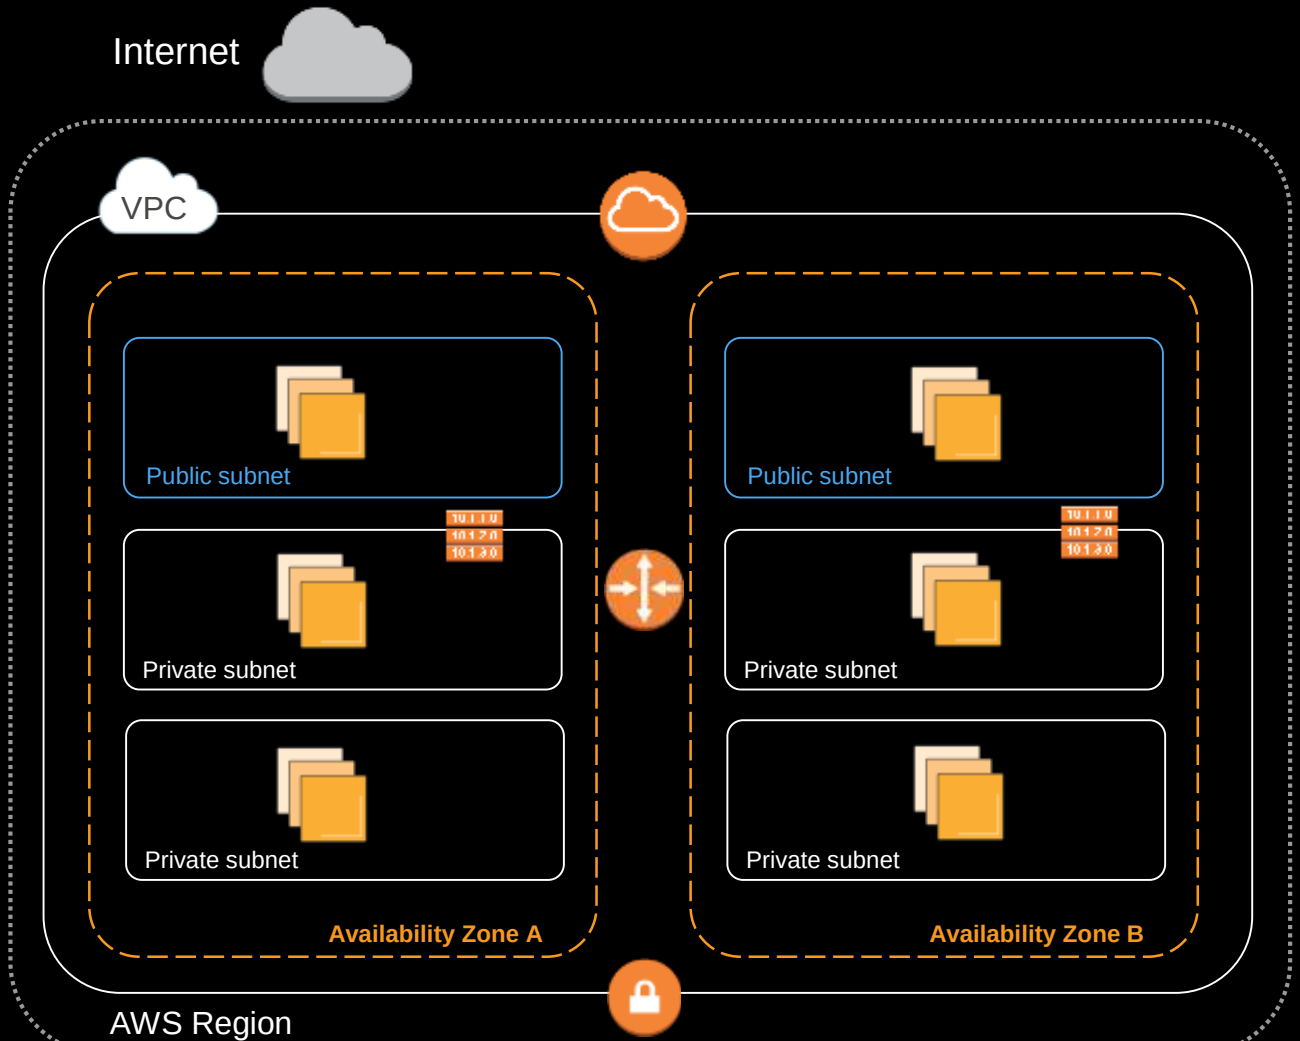

Routing Policy

Private Route Table	
Destination	Target
10.1.0.0/16	Local
0.0.0.0/0	???
Corp CIDR	VGW

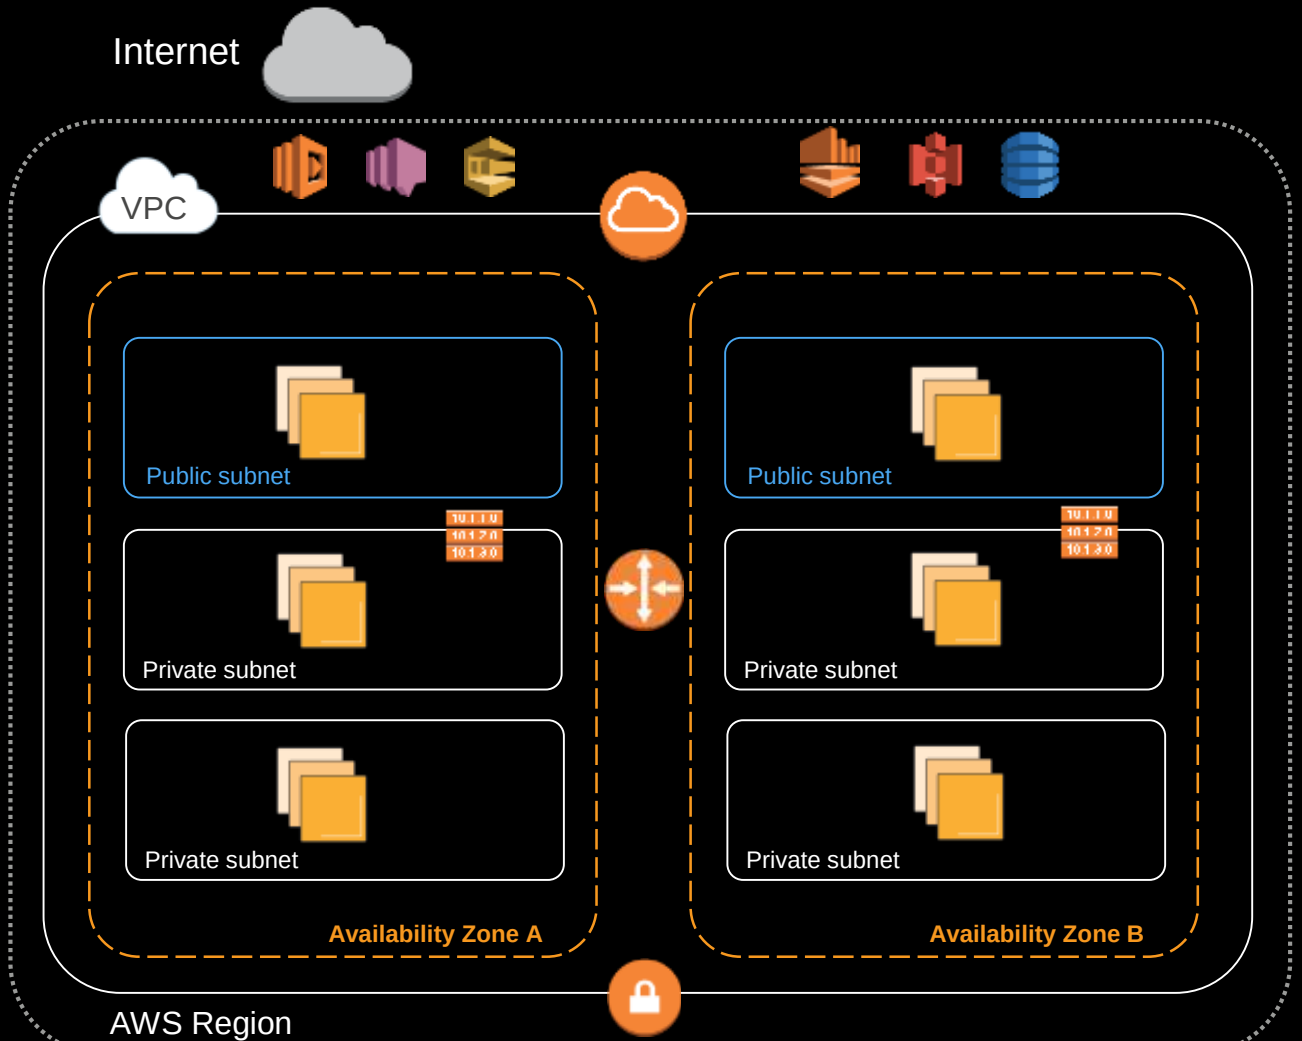

Routing Policy

Why go outside?

- AWS API endpoints
- Regional services
- Third-party services

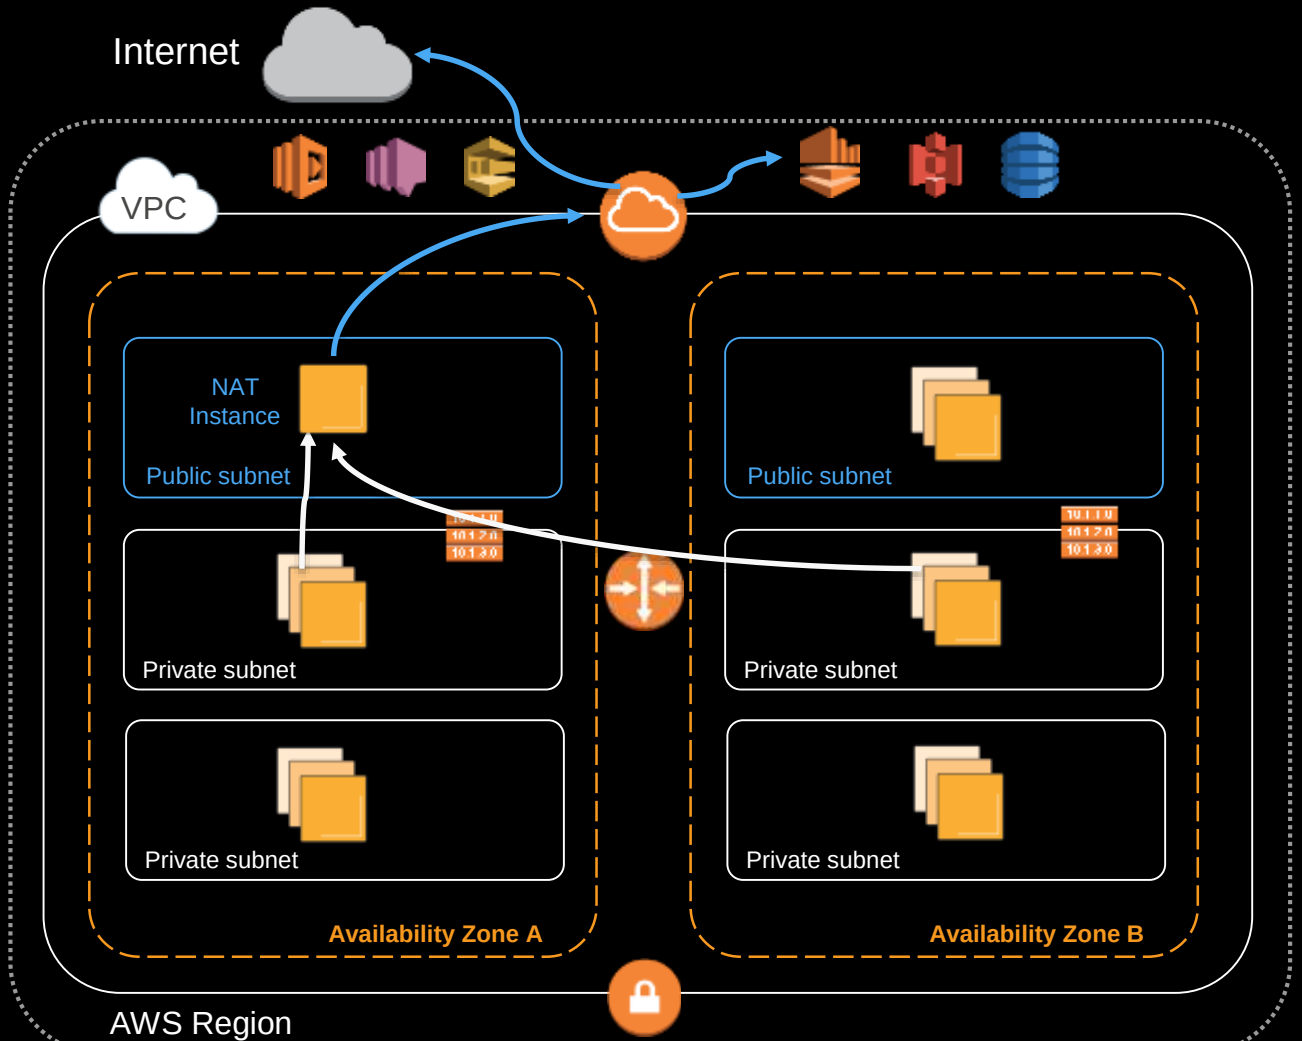

Routing Policy

Private Route Table	
Destination	Target
10.1.0.0/16	Local
0.0.0.0/0	NAT Instance
Corp CIDR	VGW

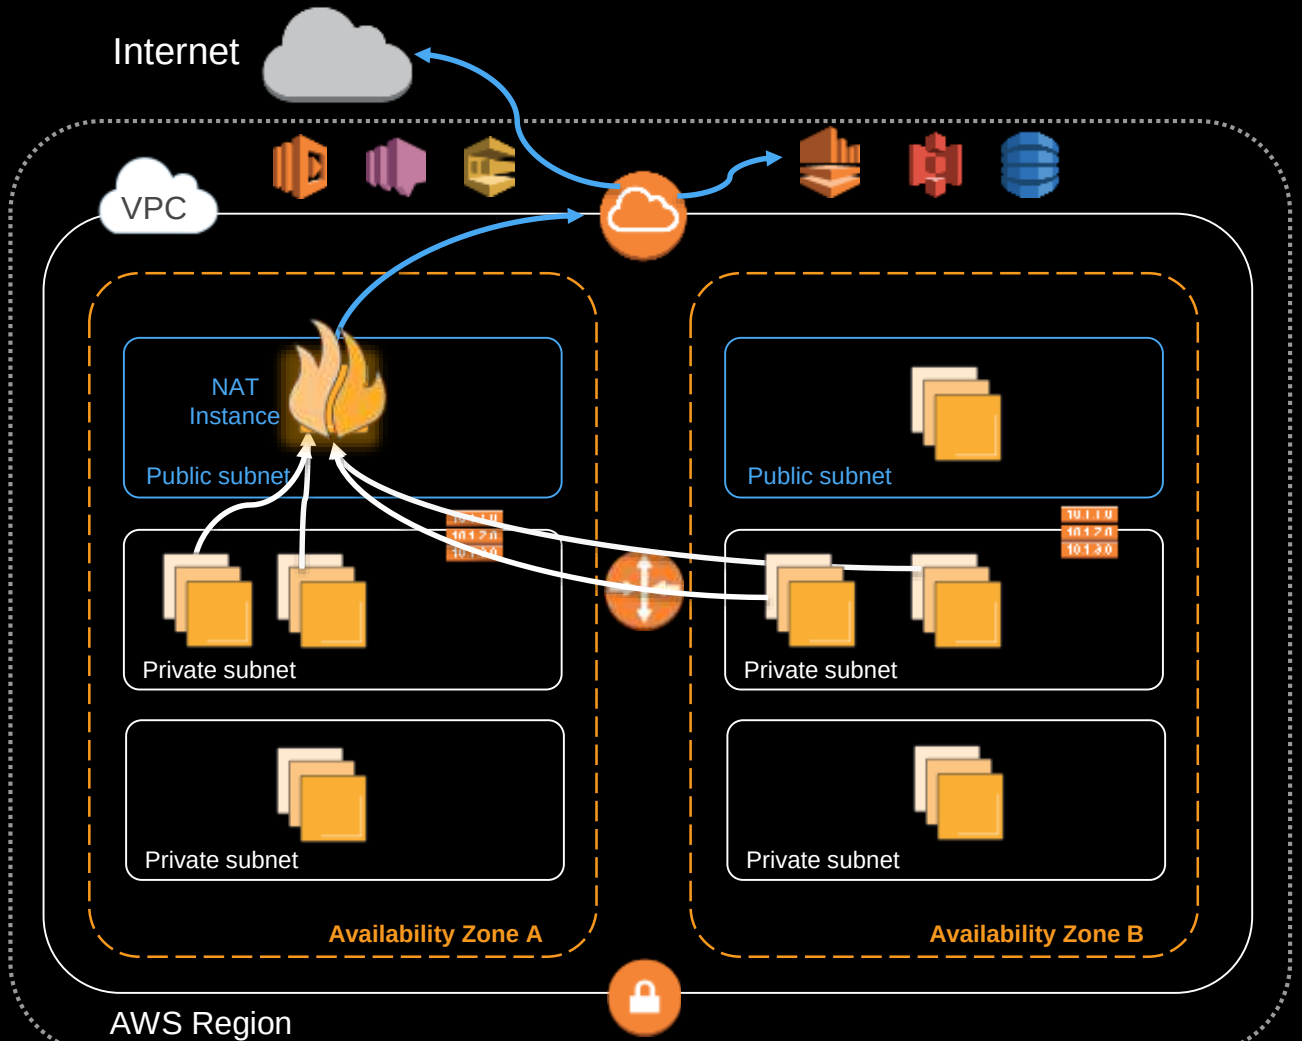

Routing Policy

Private Route Table	
Destination	Target
10.1.0.0/16	Local
0.0.0.0/0	*Black Hole*
Corp CIDR	VGW

Scalable and Available NAT

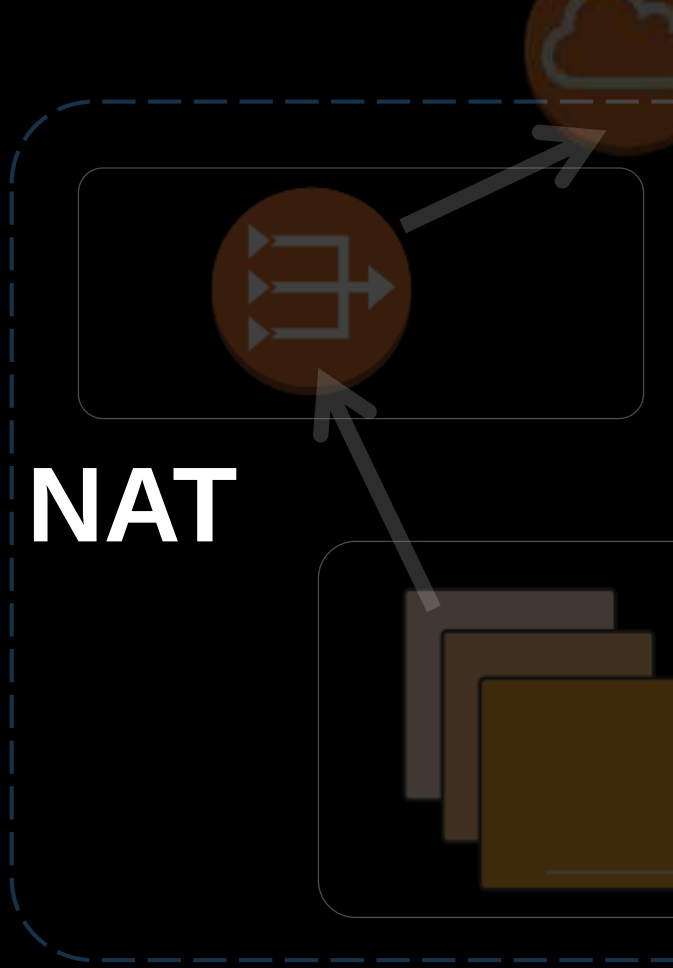

Evolving design requirements

- Public subnets for resources reachable from Internet
- Private subnets with egress only access to public network
- Scalable, highly available NAT
- One AWS account
- One VPC
- One region

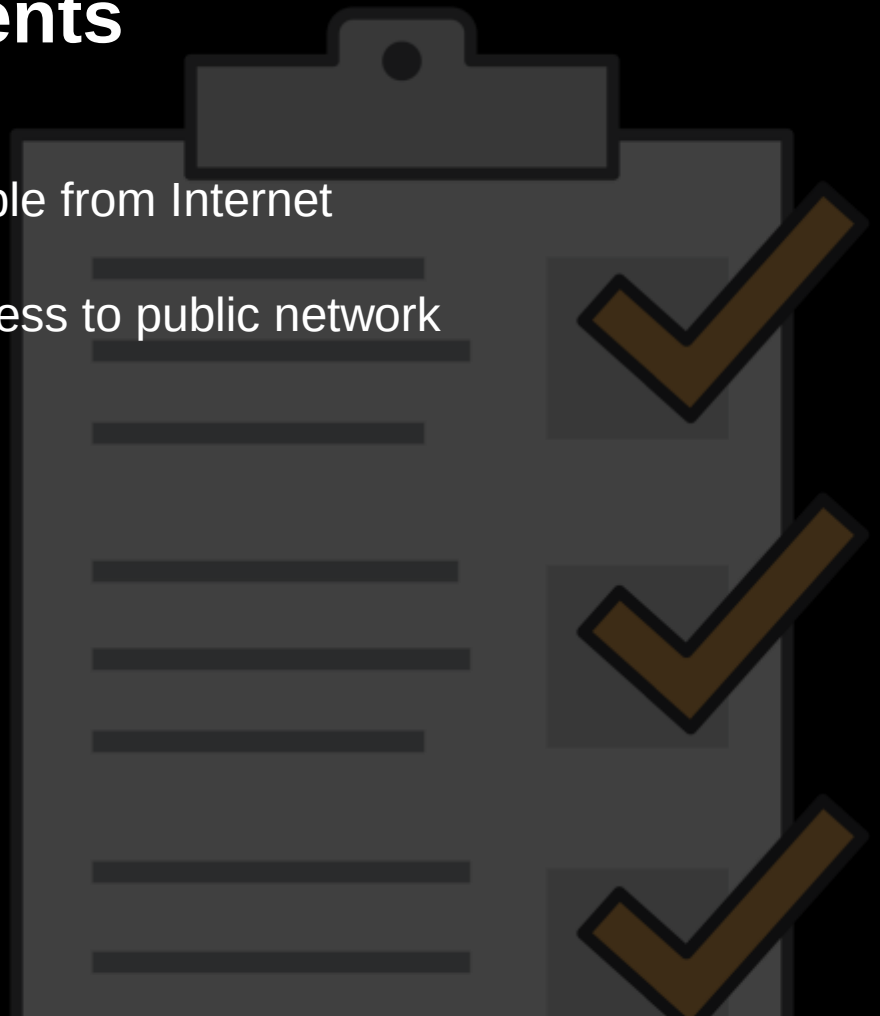

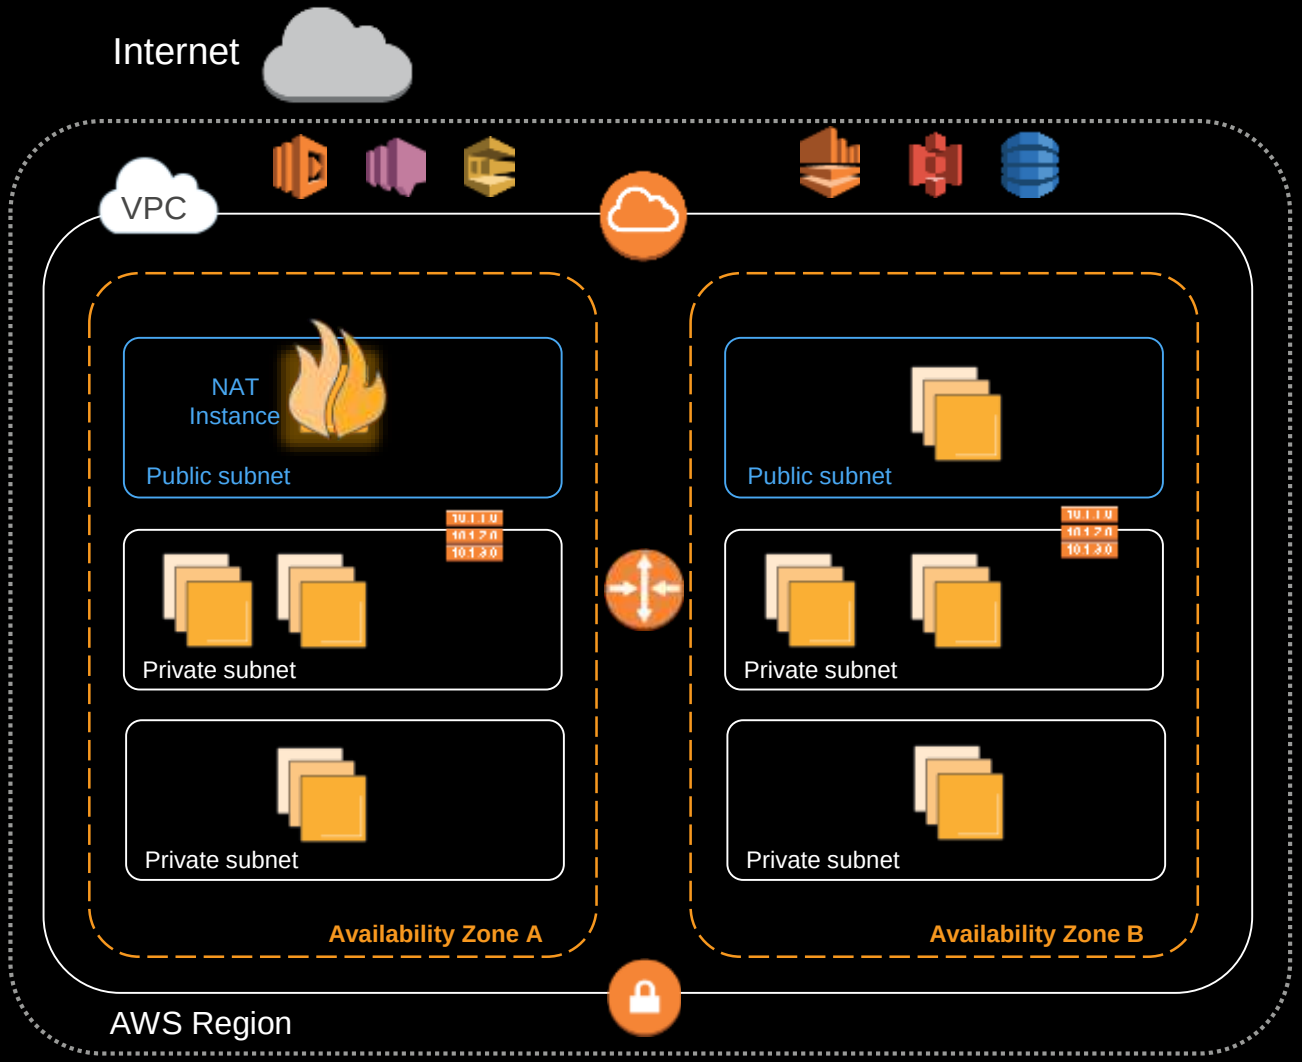

Private Route Table	
Destination	Target
10.1.0.0/16	Local
0.0.0.0/0	*Black Hole*
Corp CIDR	VGW

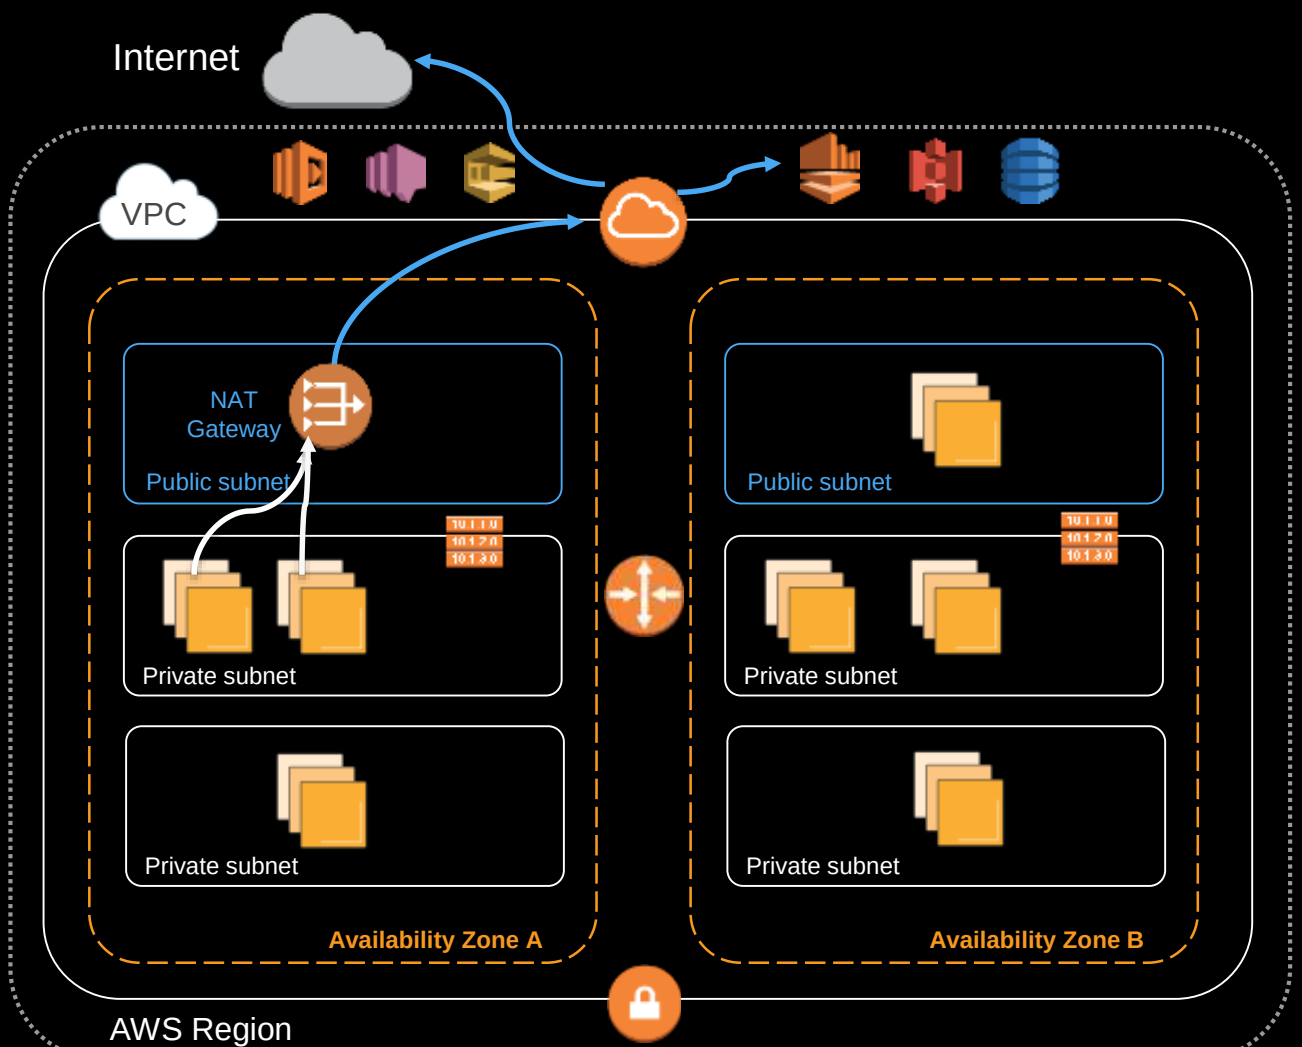

Deploy a NAT Gateway

Private Route Table	
Destination	Target
10.1.0.0/16	Local
0.0.0.0/0	NAT Gateway
Corp CIDR	VGW

Why a NAT Gateway?

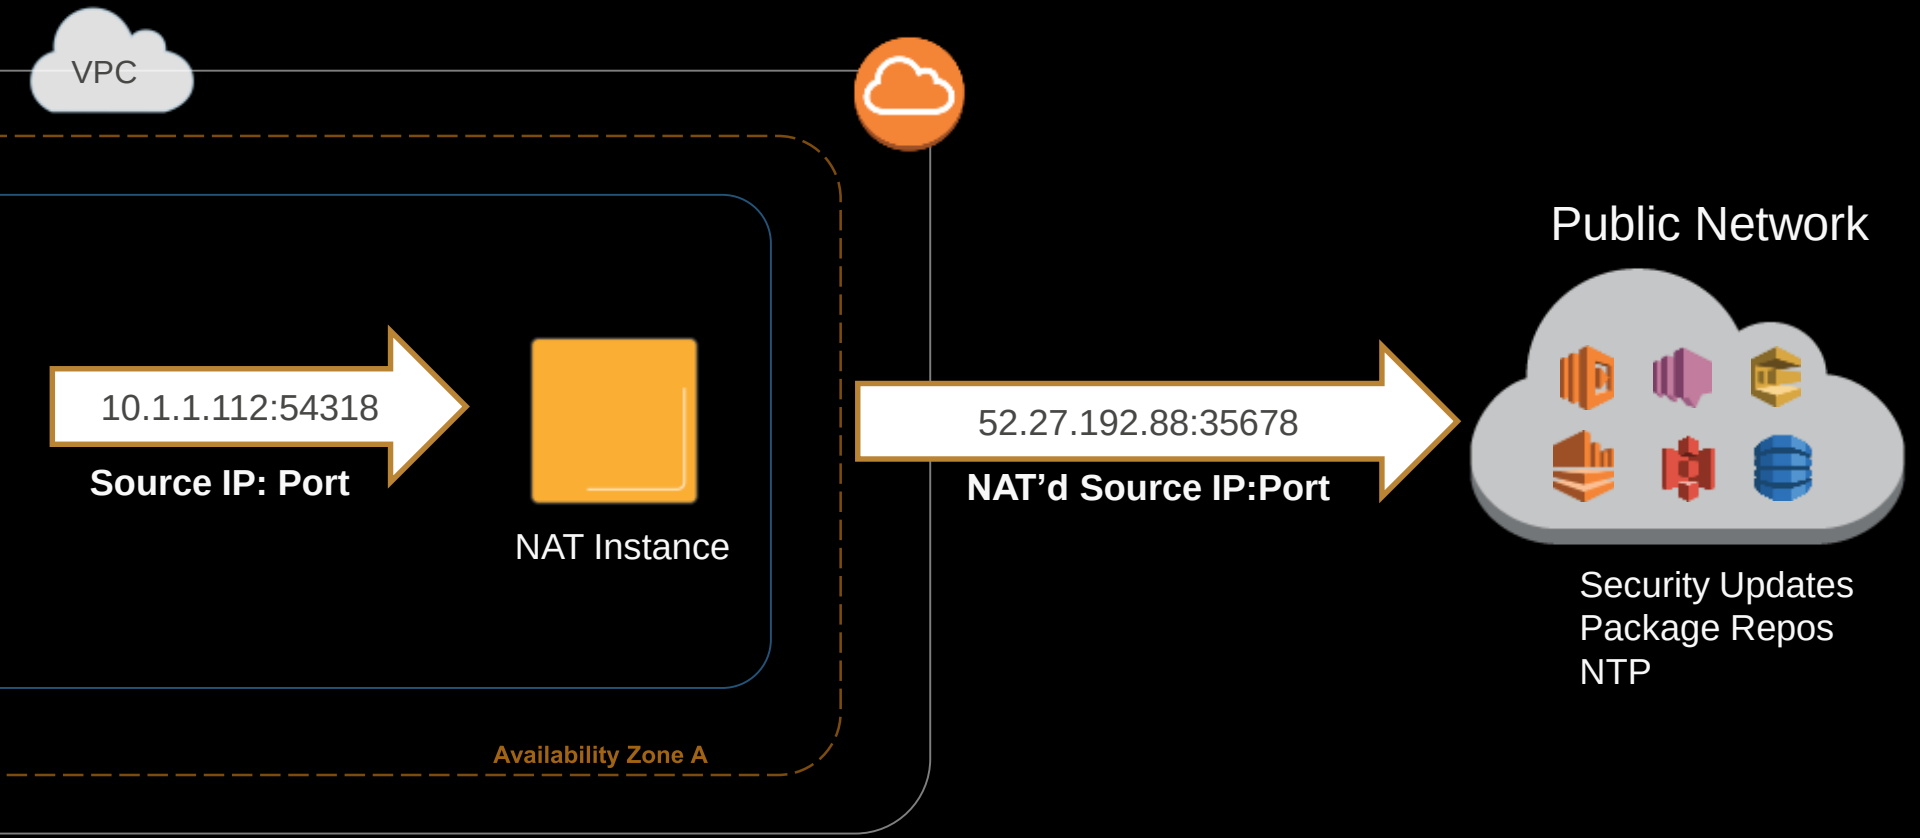

Why a NAT Gateway?

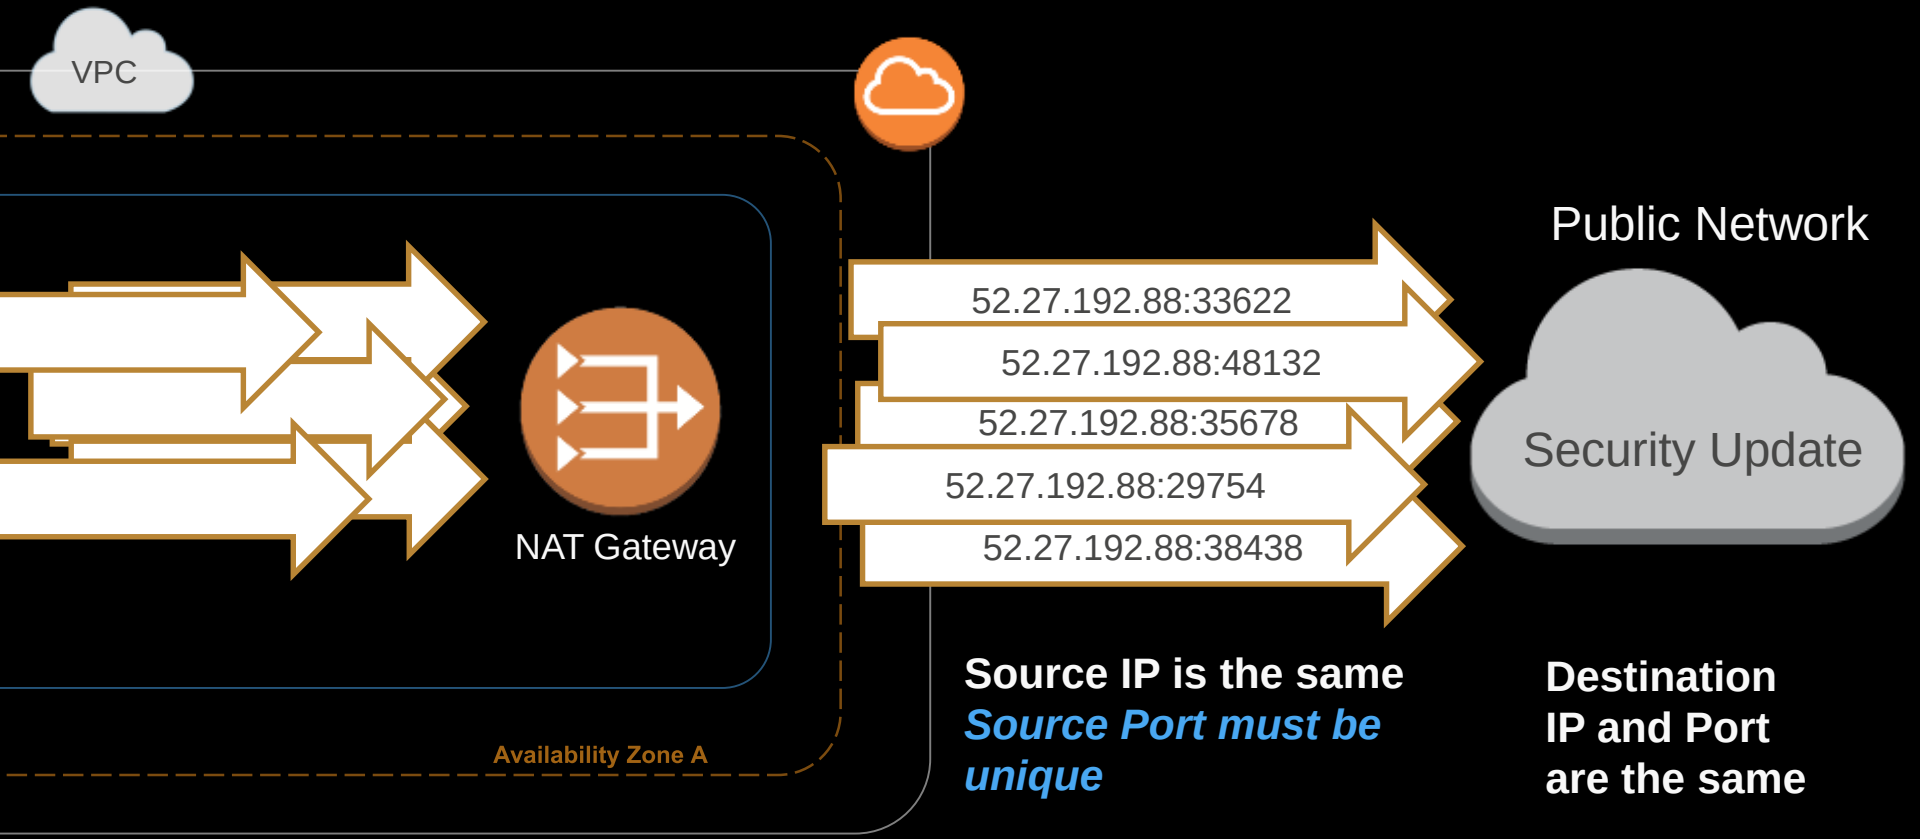

NAT Gateway: Securing Access

NAT Gateway ENI:

Network interface ID	Subnet ID	VPC ID	Security groups	Description	Status	IPv4 Public IP	Primary private IPv4 IP
eni-	subnet-	vpc-		Interface for NAT Gateway nat-056	in-use	13.54	10.2.1.89

1 Network ACLs still apply

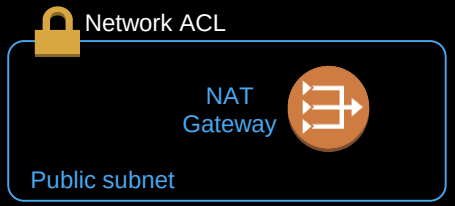

NAT Gateway: Securing Access

2

Use routing policy to control access to NAT Gateway

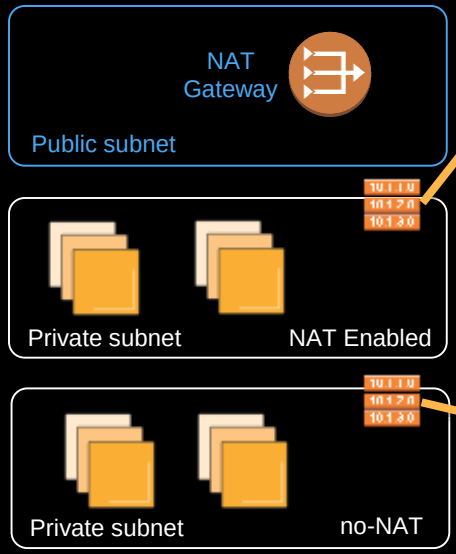

NAT Enabled Route Table	
Destination	Target
10.1.0.0/16	Local
0.0.0.0/0	NAT Gateway

no-NAT Private Route Table	
Destination	Target
10.1.0.0/16	Local

NAT Gateway: Securing Access

3

Use security groups to restrict outbound access for instances

Default VPC security group:

Outbound Rules			
Type	Protocol	Port Range	Destination
All traffic	All	0 - 65535	0.0.0.0/0

NAT Gateway: Securing Access

3

Use security groups to restrict outbound access for instances

Default VPC security group:

Outbound Rules			
Type	Protocol	Port Range	Destination
All traffic	All	0 - 65535	10.2.0.0/16

NAT Enabled VPC security group:

Outbound Rules			
Type	Protocol	Port Range	Destination
All traffic	All	0 - 65535	0.0.0.0/0

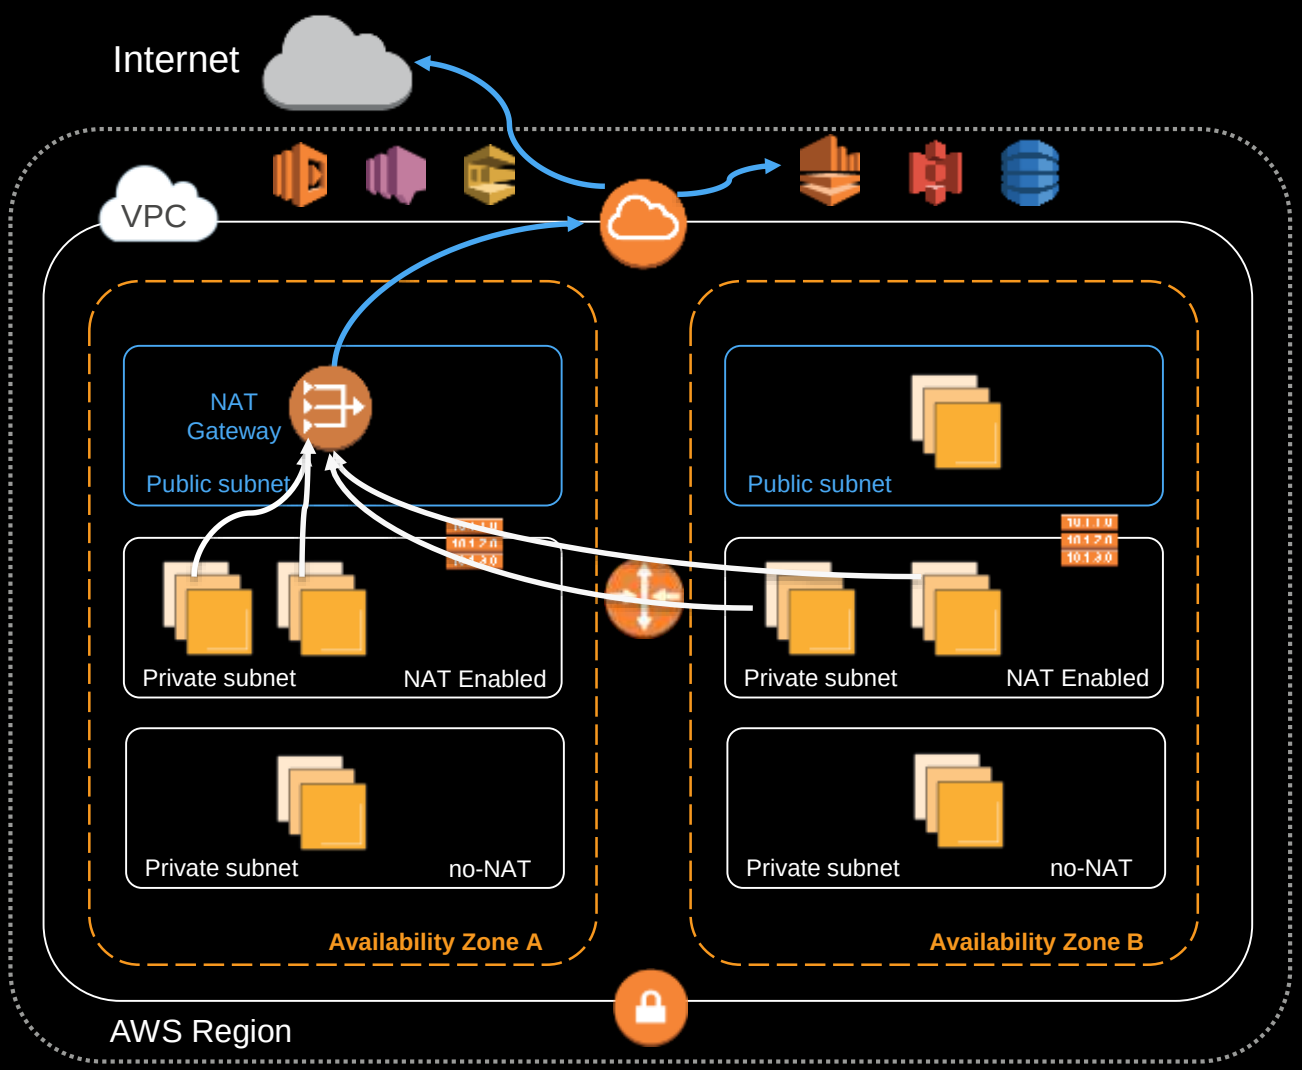

Deploy a NAT Gateway

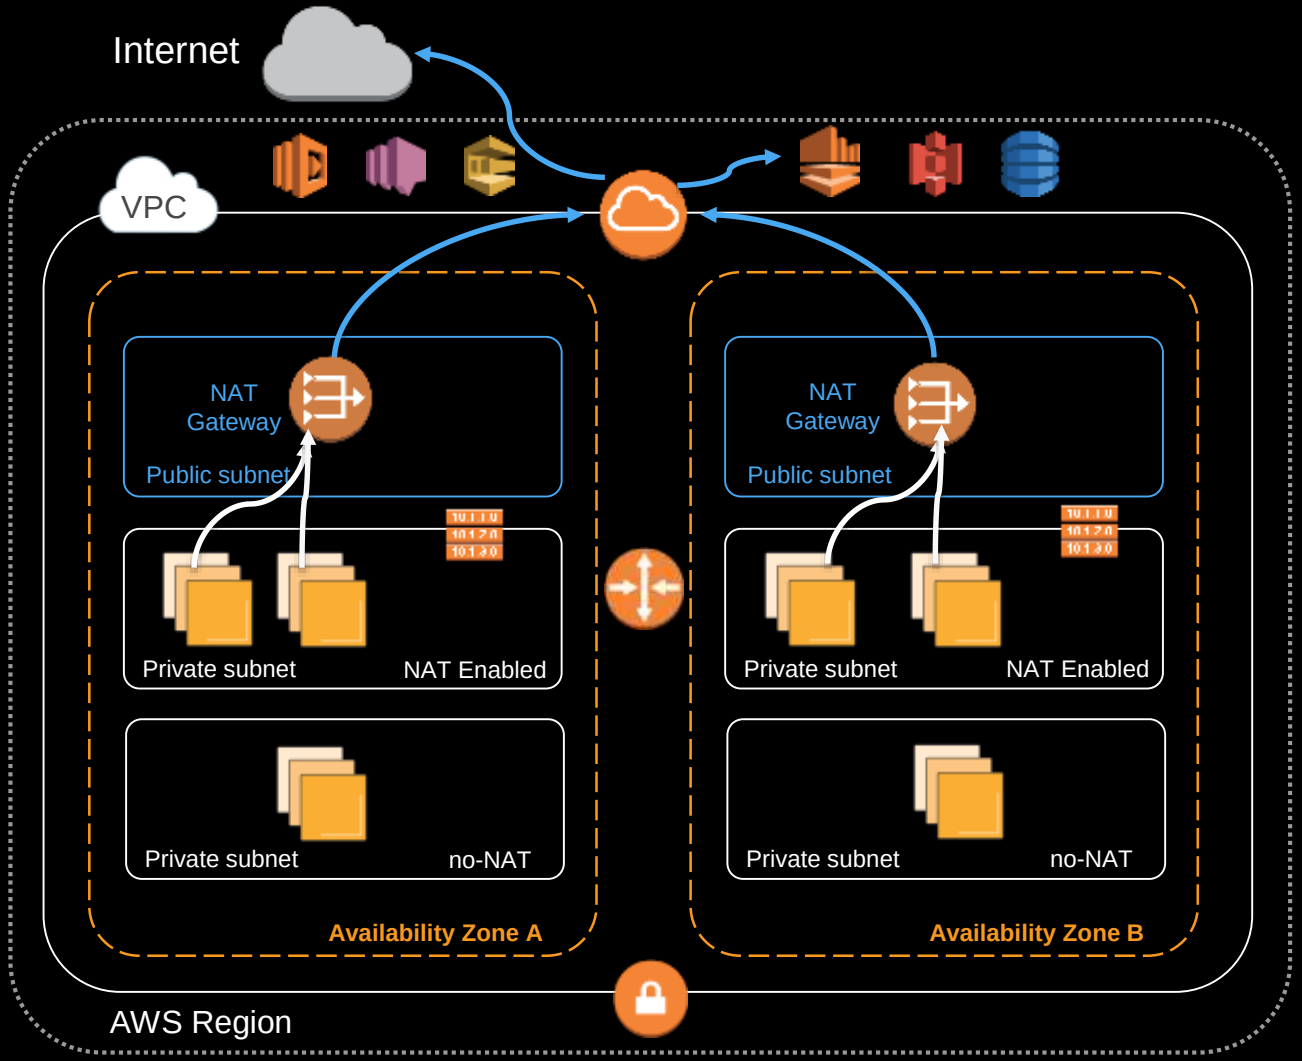

Deploy a NAT Gateway

Pro & Con: NAT Gateway

- Drop in replacement for NAT instance
- Fully managed
- Highly available and fault tolerant
- Scalable to 10 Gbps burst per gateway
- Supports VPC Flow Logs

- No higher level functions like IPS, UTM, URL Filtering, packet inspection, etc
- Cannot associate security group to gateway

Considering multiple VPCs

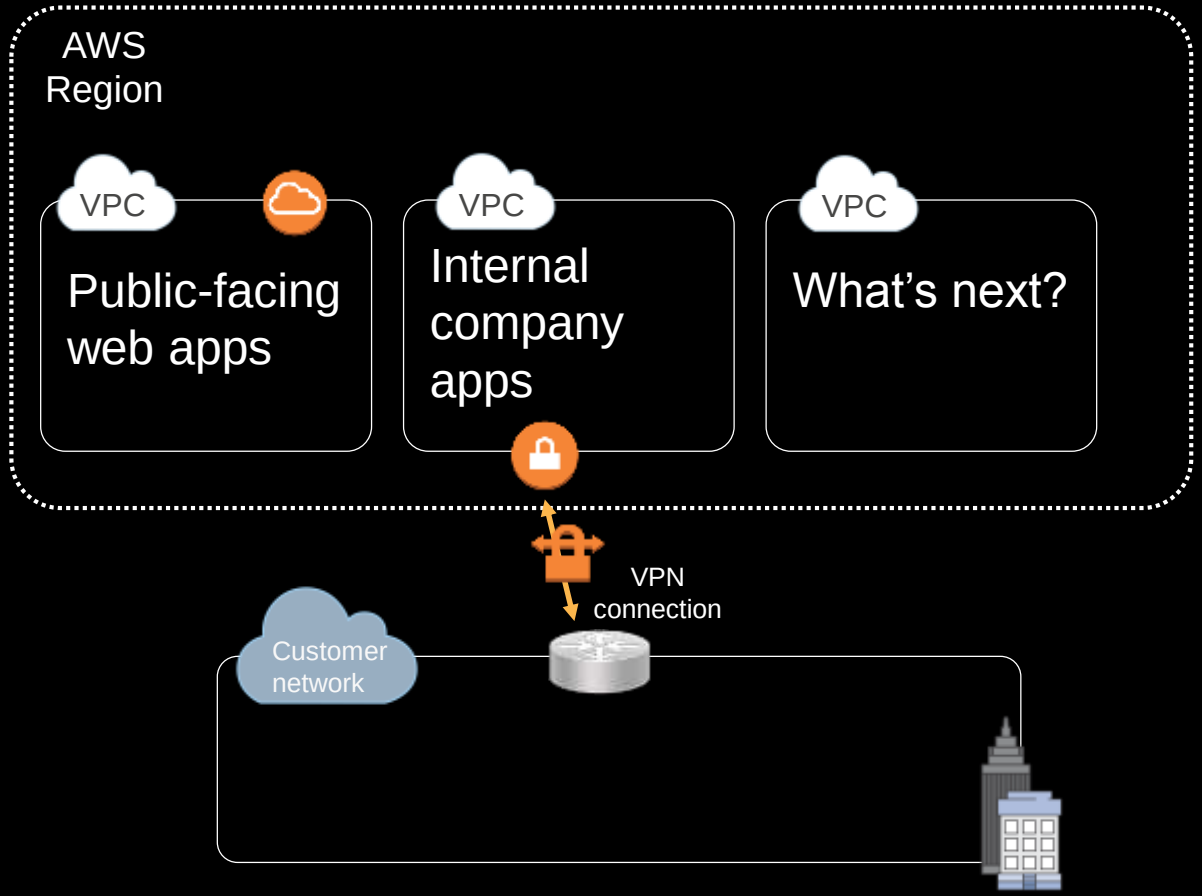


One VPC, Two VPC

VPC

Why not 1 big VPC?

Why not 1 AWS Account?

- Blast radius
- Account Limits
- API Limits

Considerations for one or many VPCs

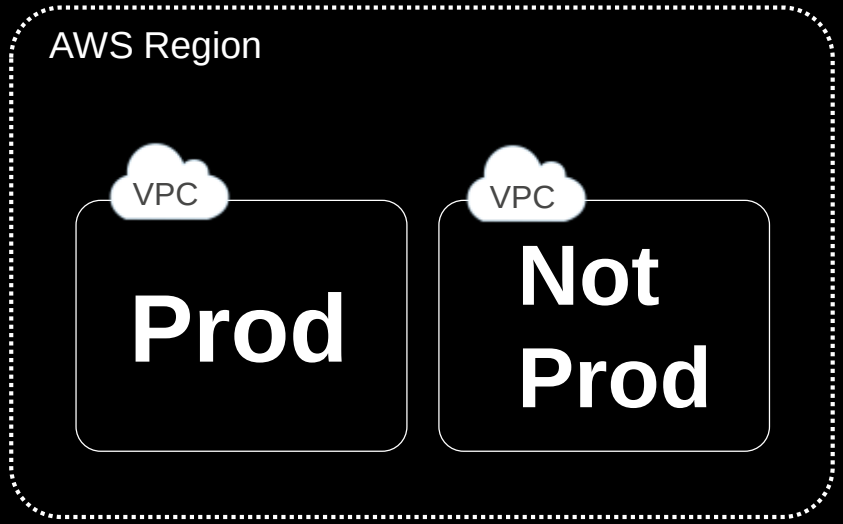

Considerations for one or many VPCs

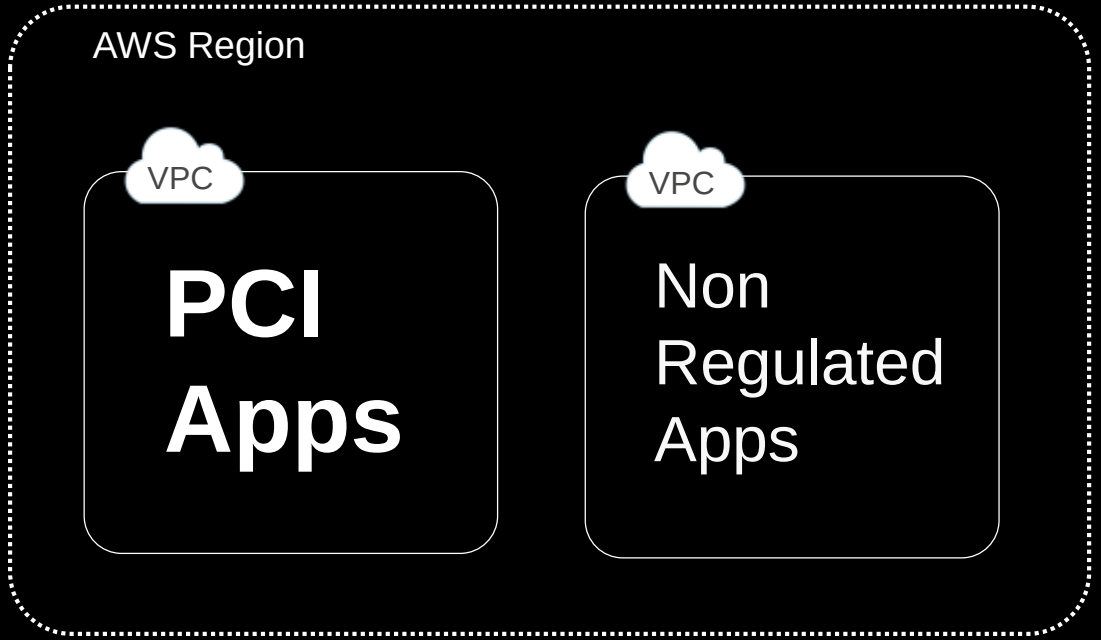

Considerations for one or many VPCs

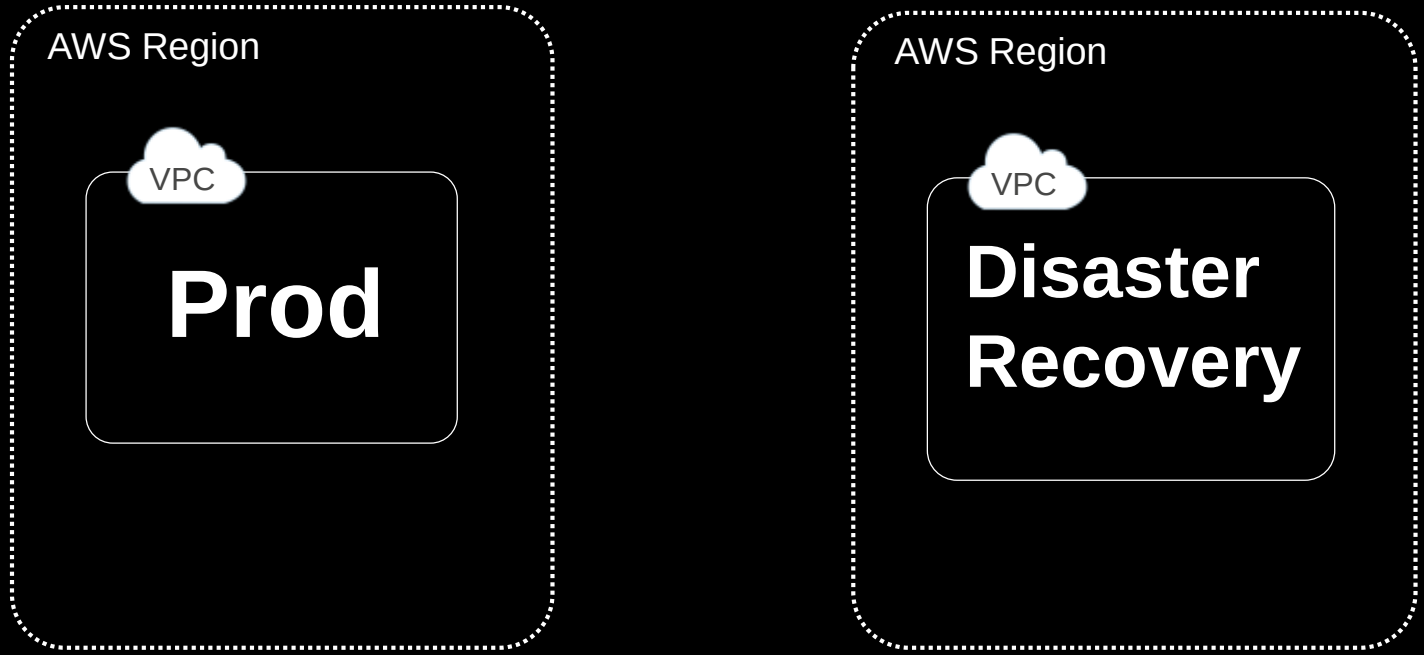

Considerations for one or many VPCs

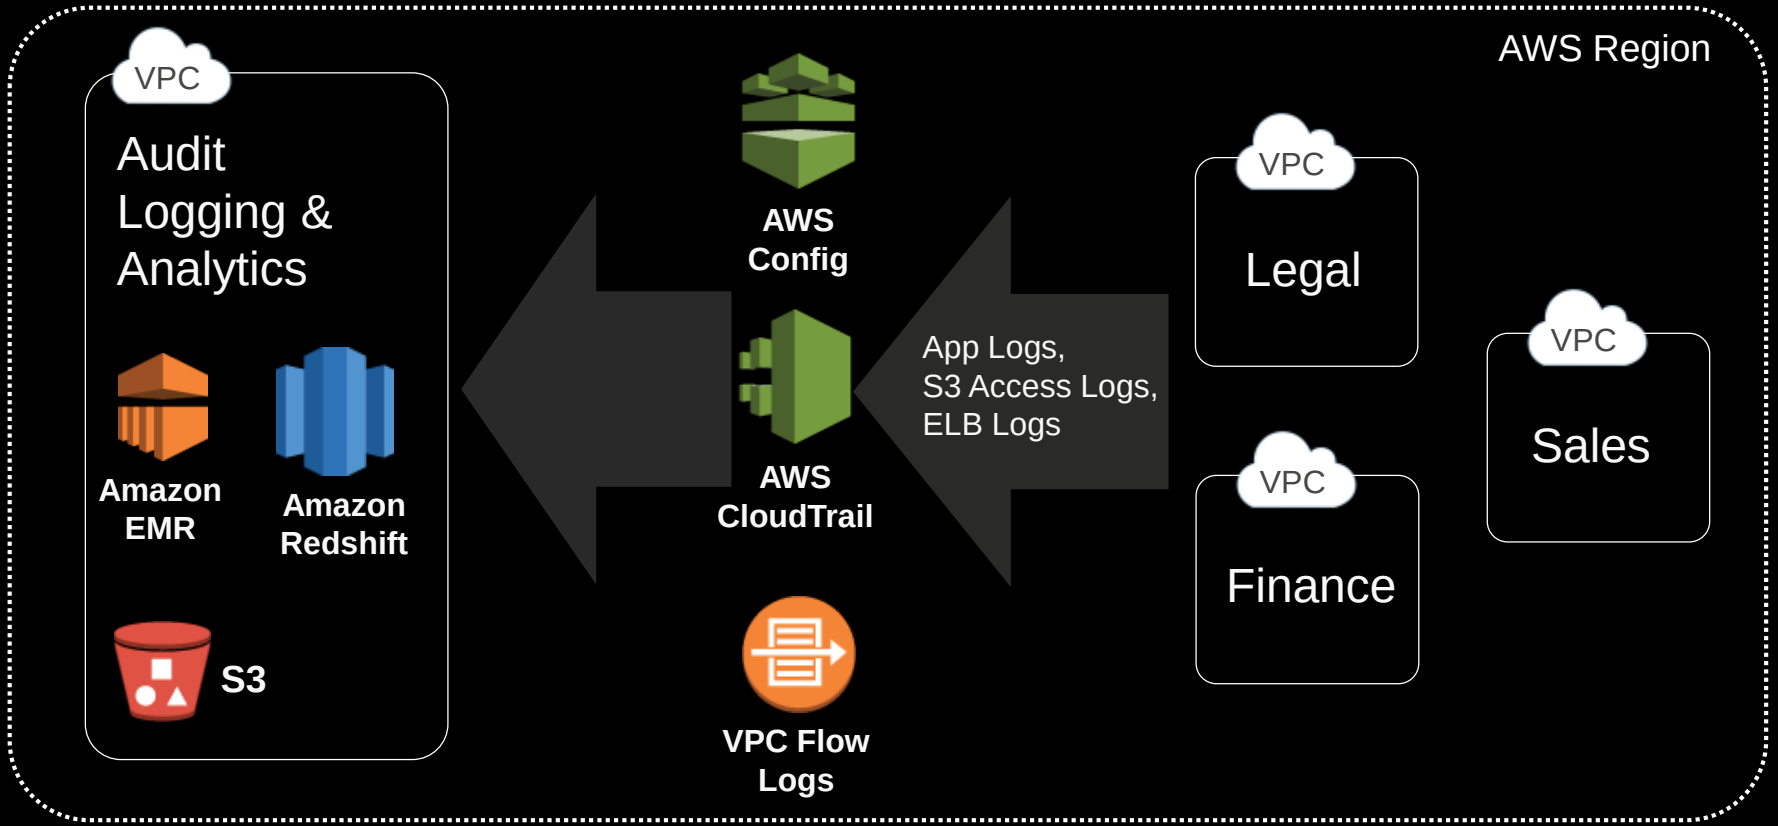

Internal application to VPC

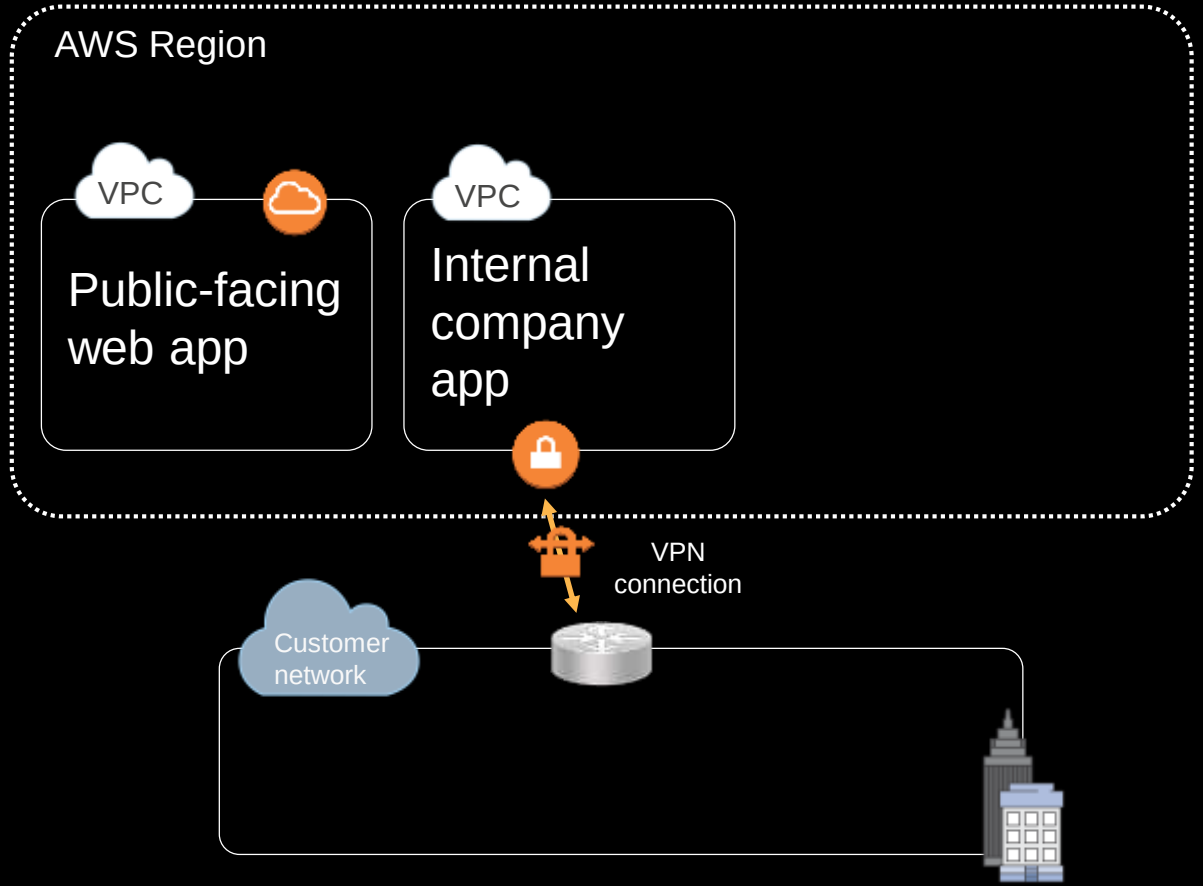

Internal application to VPC

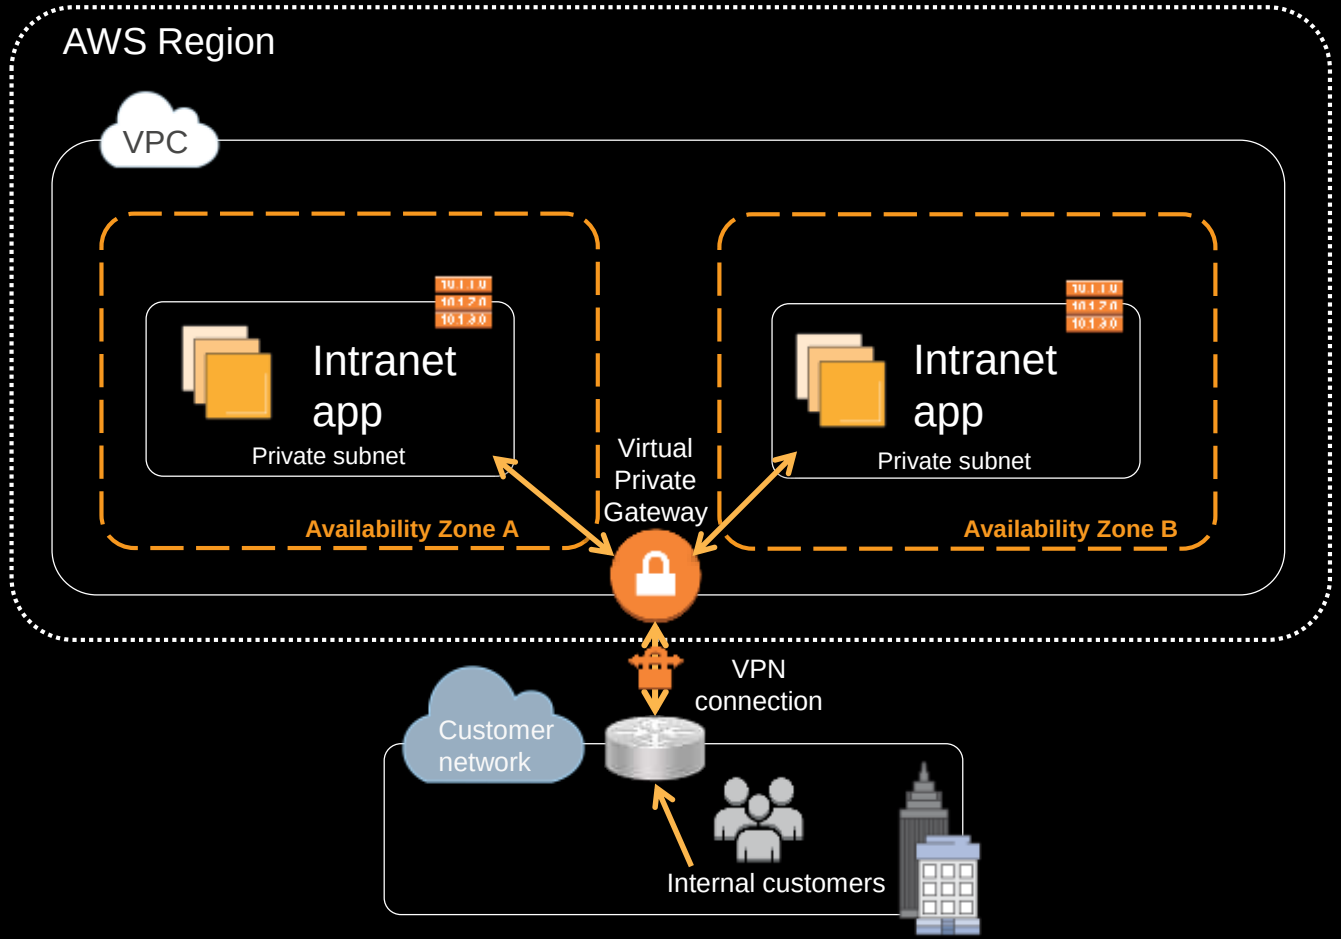

Private Route Table	
Destination	Target
10.1.0.0/16	Local
Corp CIDR	VGW

But apps will make heavy use of …

Amazon S3

…as a primary data store

VPC Egress Control

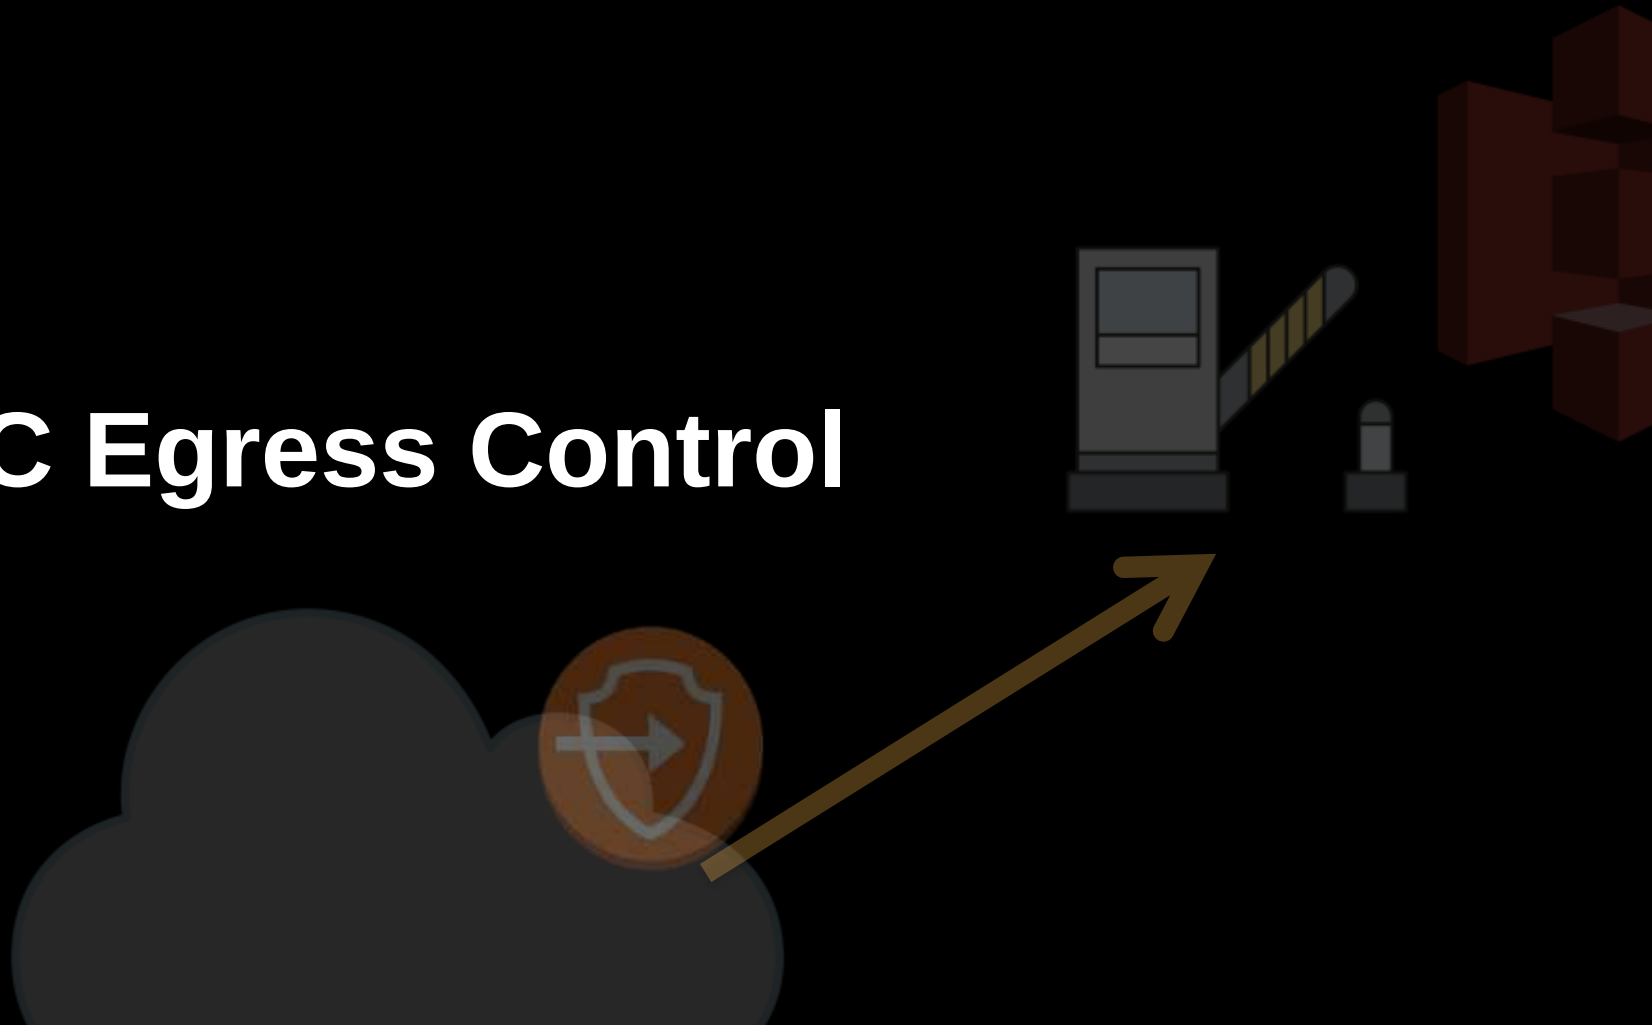

Evolving design requirements

- VPN connectivity to private-only VPC
- No egress in the VPC to public networks
- Private IP access to Amazon S3
- Content-specific access controls
- One AWS account
- One VPC
- One region

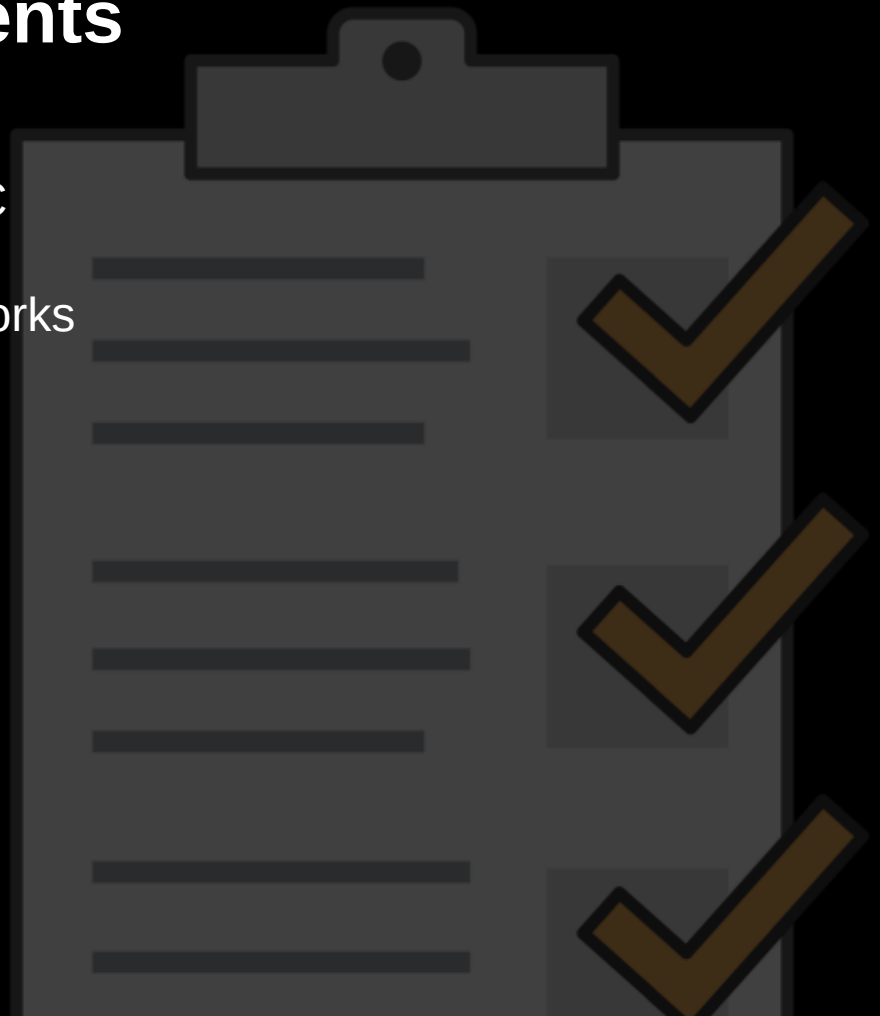

You really don't want to do this:

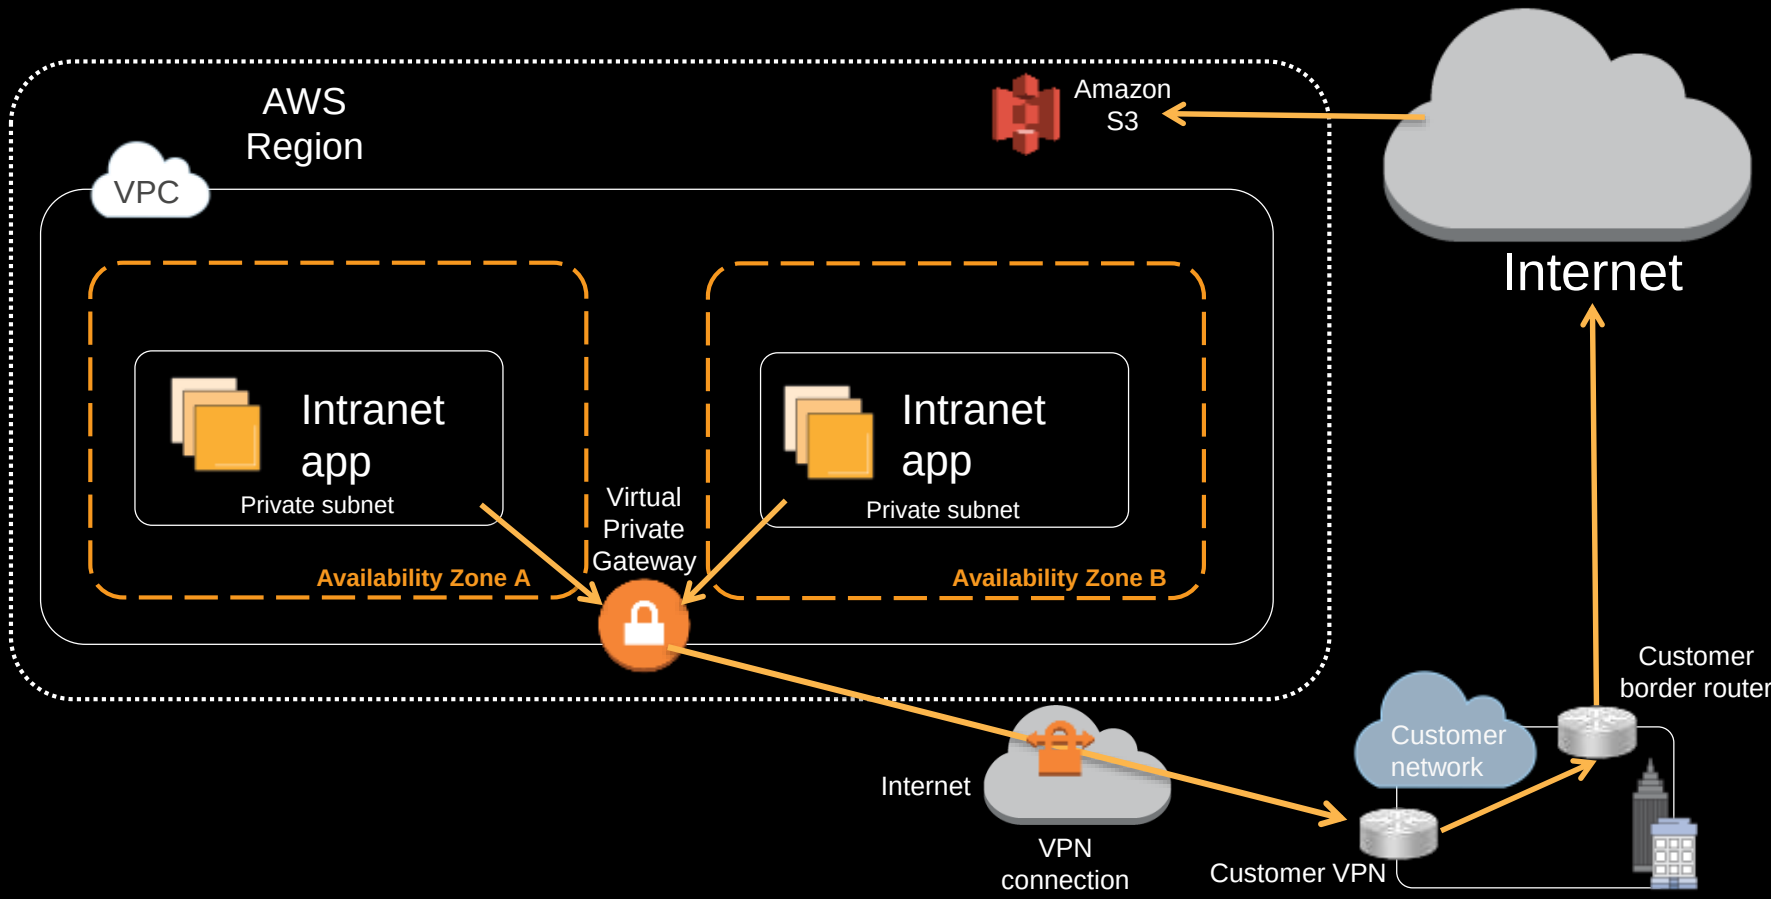

So do this instead:

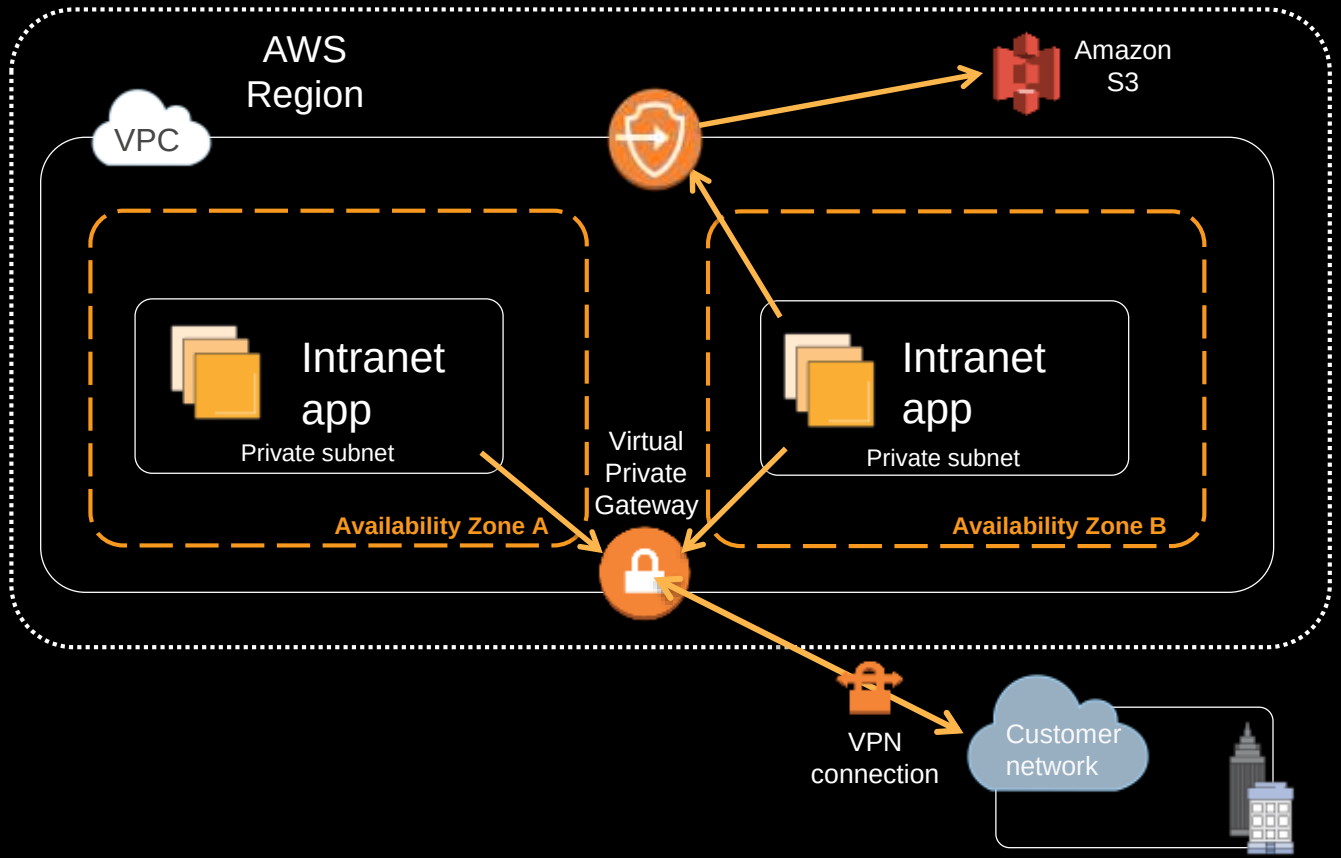

VPC Endpoints

- No IGW
- No NAT
- No public IPs
- Free
- Robust access control

Creating S3 VPC endpoint

```
aws ec2 create-vpc-endpoint
--vpc-id vpc-40f18d25
--service-name com.amazonaws.us-west-2.s3
--route-table-ids rtb-2ae6a24f
```

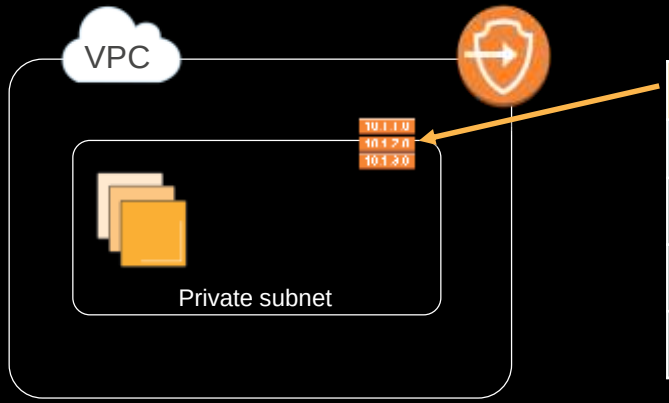

Route Table	
Destination	Target
10.1.0.0/16	Local
Corp CIDR	VGW
Prefix List for S3 us-west-2	VPCE

Prefix lists

```
aws ec2 describe-prefix-lists
PREFIXLISTS    pl-68a54001    com.amazonaws.us-west-2.s3
CIDRS     54.231.160.0/19
CIDRS     52.218.128.0/18
```

- Logical route destination target
- Dynamically translates to service IPs
- S3 IP ranges change over time
- S3 prefix lists abstract change

Prefix lists in Security Groups

… and use them in your outbound security group rules!

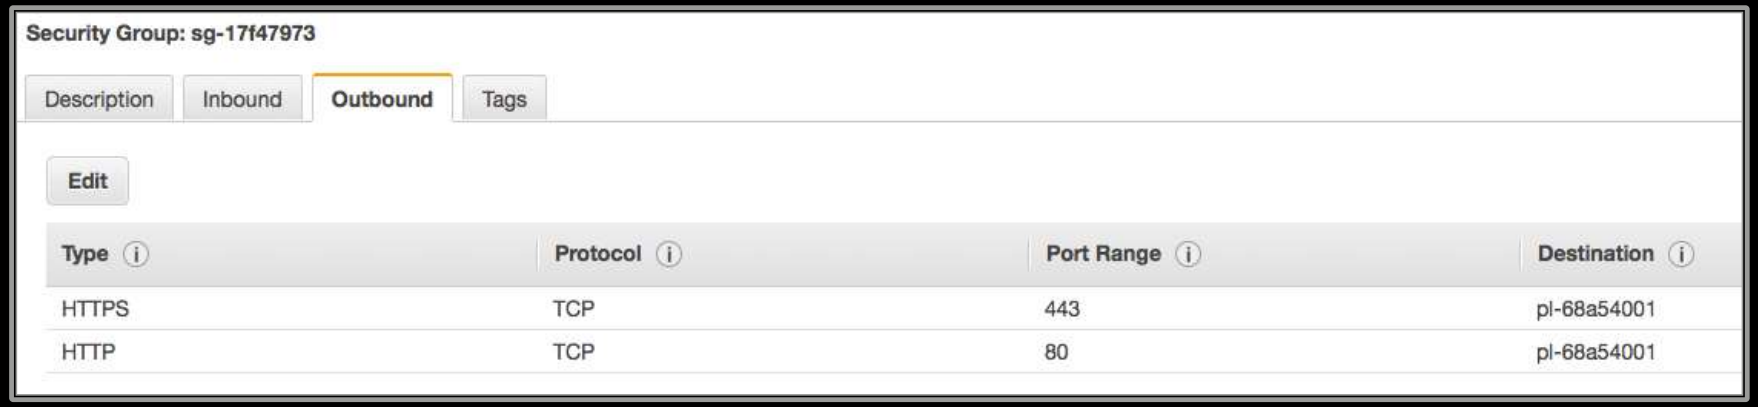

Which buckets can my endpoint access?

AWS Identity & Access Management (IAM) policy on VPCE:

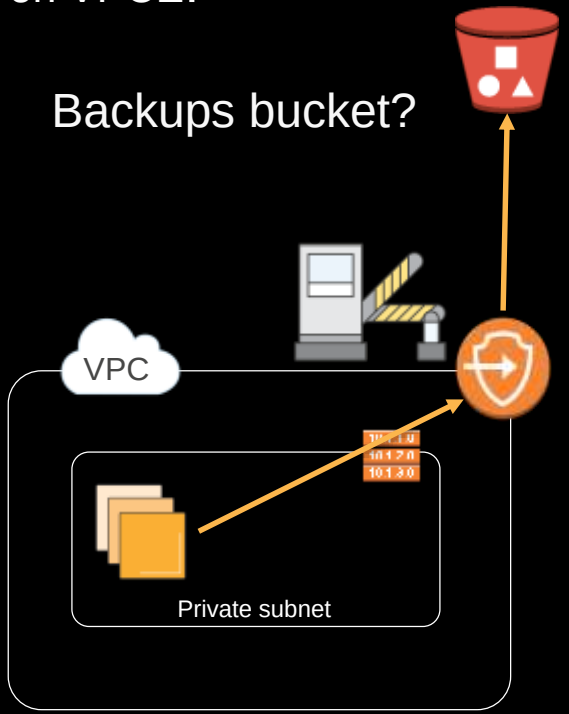

```
{
  "Statement": [
    {
      "Sid": "vpce-restrict-to-backup-bucket",
      "Principal": "*",
      "Action": [
        "s3:GetObject",
        "s3:PutObject"
      ],
      "Effect": "Allow",
      "Resource": ["arn:aws:s3:::backups-reinvent",
                   "arn:aws:s3:::backups-reinvent/*"]
    }
  ]
}
```

Which endpoints can access my bucket?

S3 bucket policy:

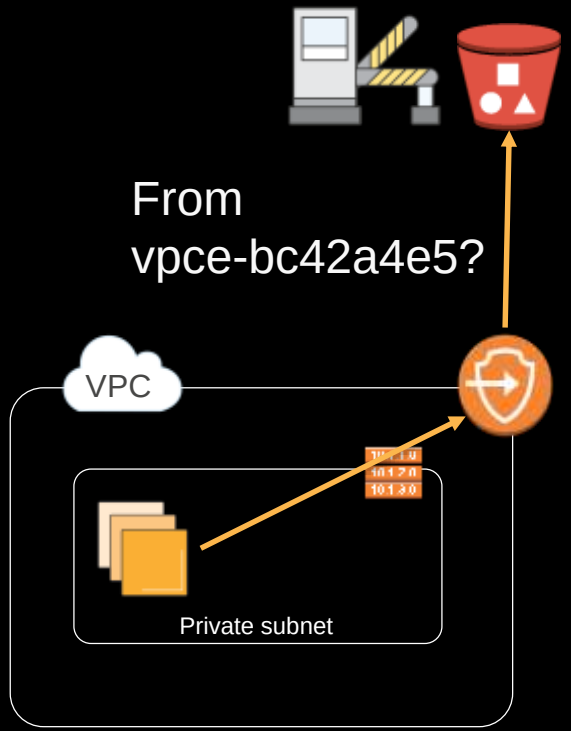


```
{
  "Statement": [
    {
      "Sid": "bucket-restrict-to-specific-vpce",
      "Principal": "*",
      "Action": "s3:*",
      "Effect": "Deny",
      "Resource": ["arn:aws:s3:::backups-reinvent",
                   "arn:aws:s3:::backups-reinvent/*"],
      "Condition": {
        "StringNotEquals": {
          "aws:sourceVpce": "vpce-bc42a4e5"
        }
      }
    }
  ]
}
```

Controlling VPC access to Amazon S3

Recap on security layers:

1. Route table association
2. VPCE policy
3. Bucket policy
4. Security groups with prefix list

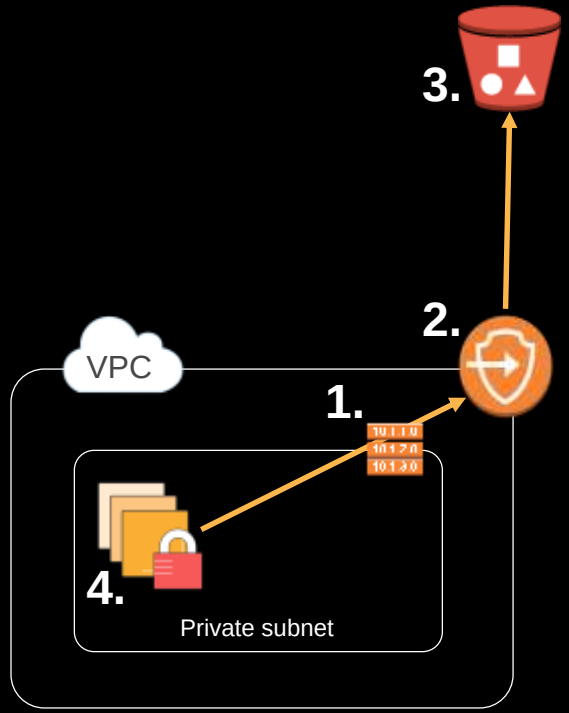

Endpoints in action

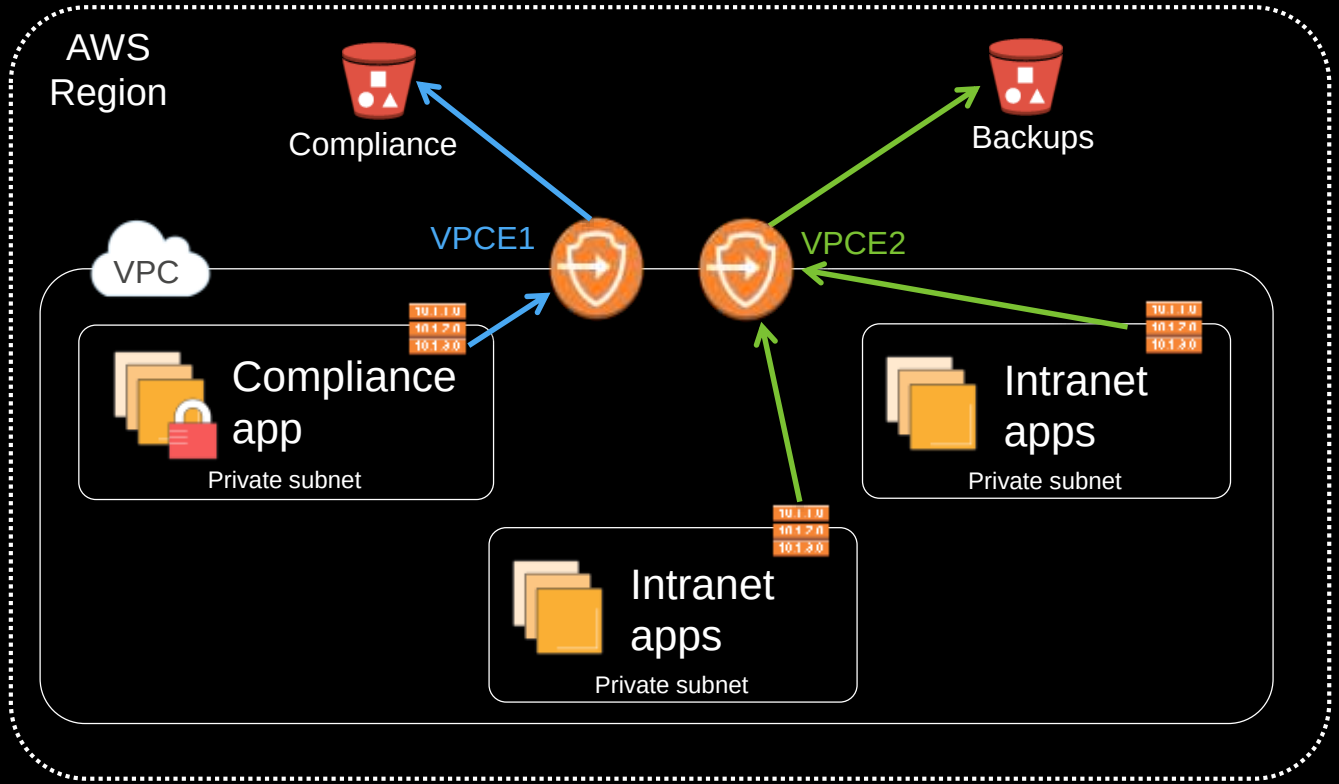

Pro & Con: VPC Endpoints

- Secure, highly scalable and highly available access to S3
- Fine grained control of access to content in S3 from VPC
- Control which VPCs/VPCEs can access which S3 buckets
- No public IPs required, source IPs kept private

- Bucket policy restricted to specific VPCs (or VPCEs) will disable S3 Console access
- Requires Amazon DNS enabled on VPC

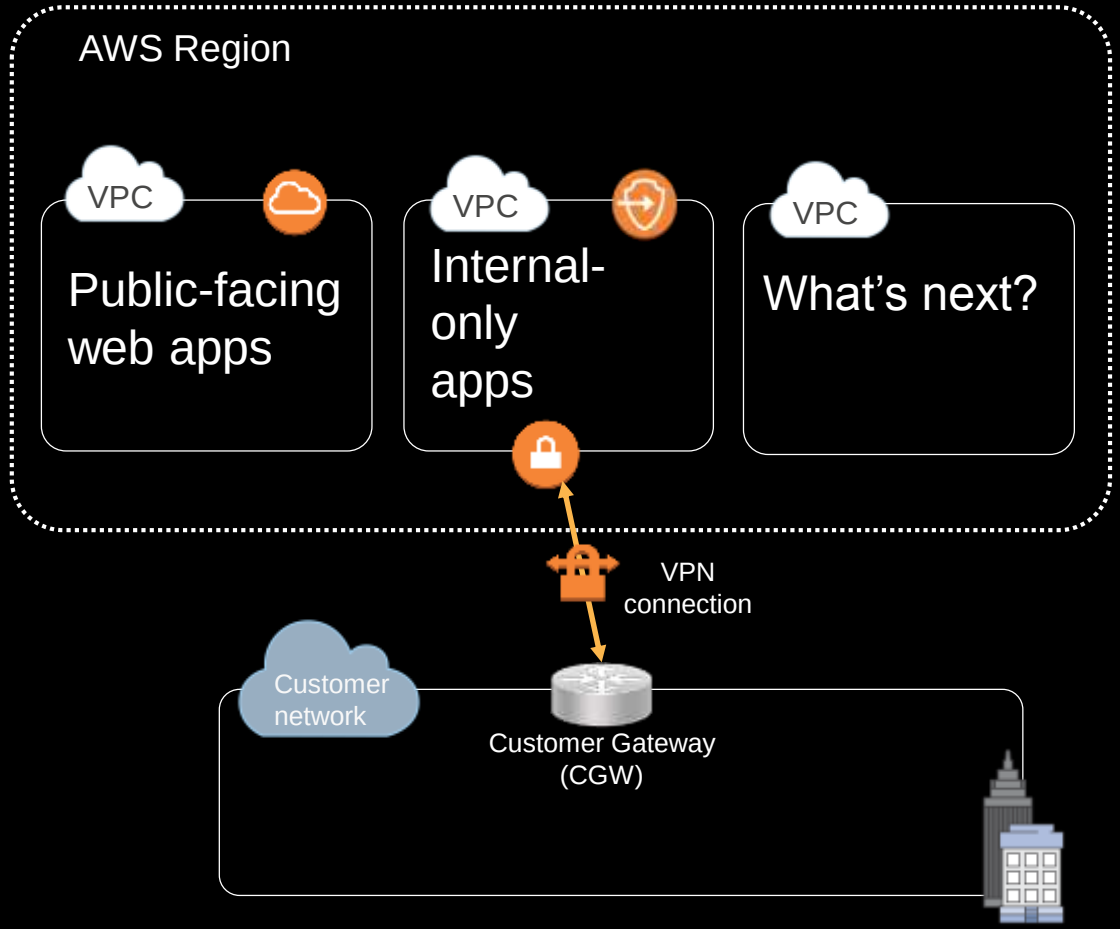
AWS Region
VPC
Public-facing web apps
VPC
Internal-only apps
VPC
What's next?
VPN connection
Customer network
Customer Gateway (CGW)

Shared Service Hubs

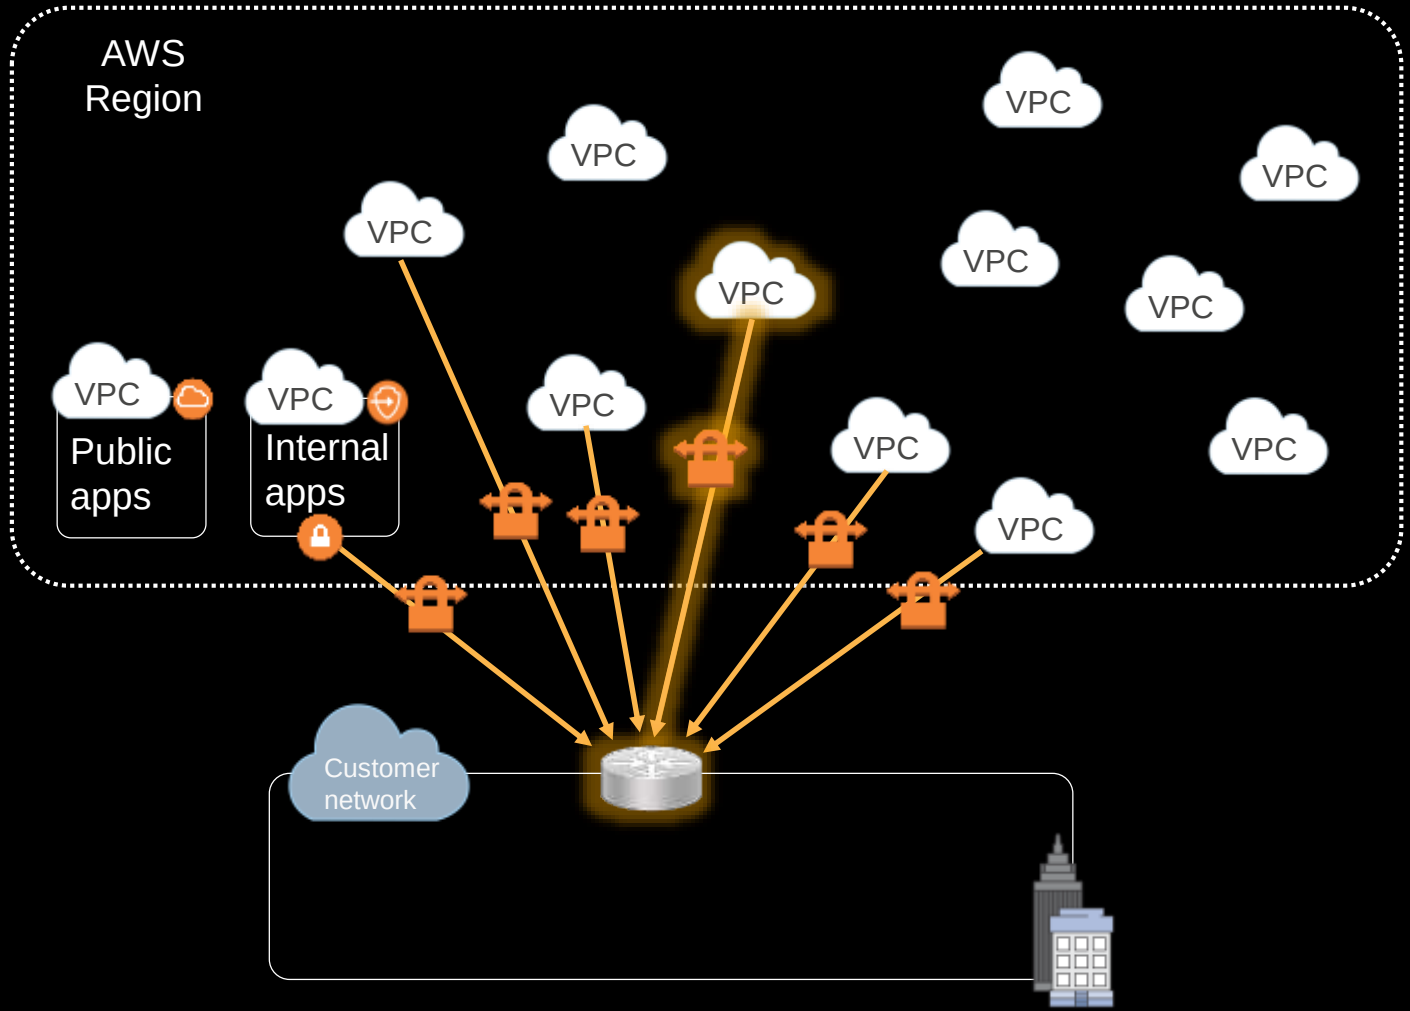
AWS
Region
VPC
VPC
VPC
VPC
VPC
VPC
VPC
VPC
VPC
VPC
VPC
VPC
VPC
Public
apps
VPC
Internal
apps
Customer
network

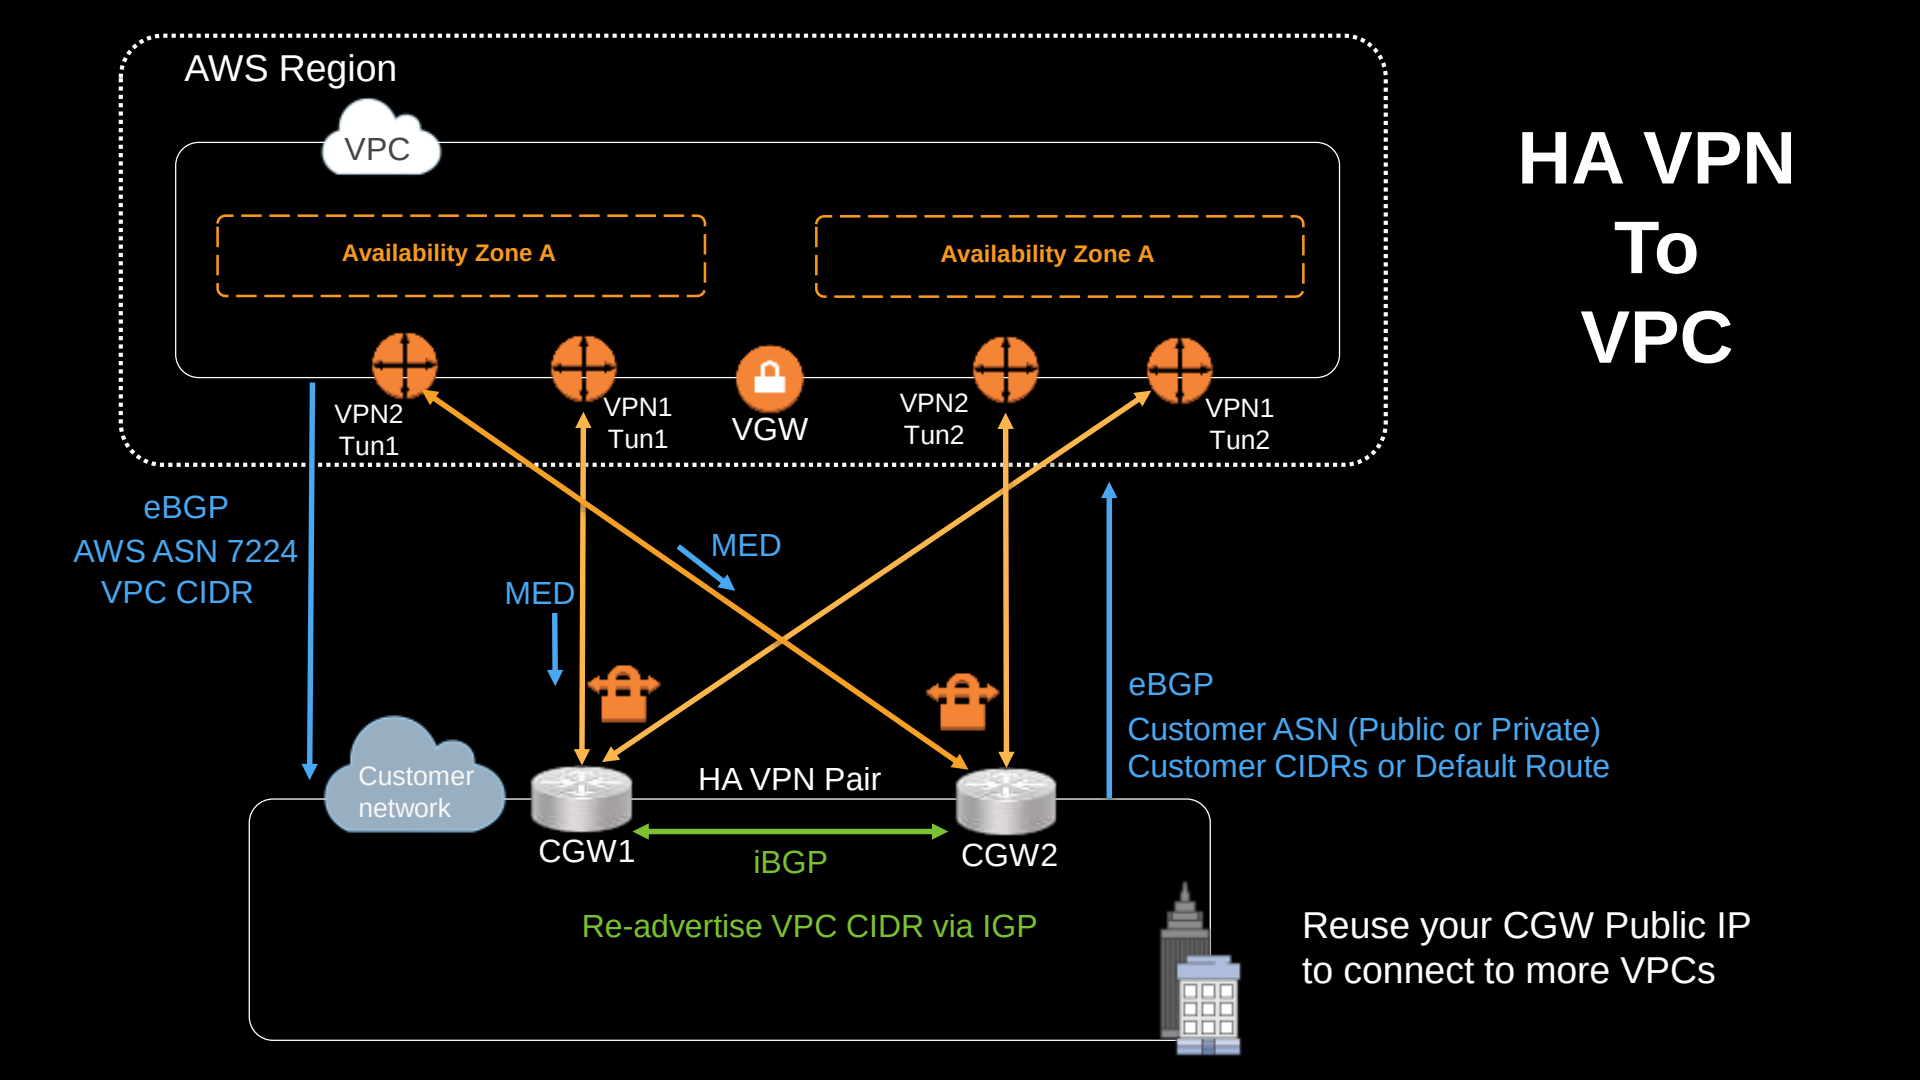
AWS Region
VPC
Availability Zone A
Availability Zone A
VPN2
Tun1
VPN1
Tun1
VGW
VPN2
Tun2
VPN1
Tun2
eBGP
AWS ASN 7224
VPC CIDR
MED
MED
eBGP
Customer ASN (Public or Private)
Customer CIDRs or Default Route
Customer
network
HA VPN Pair
CGW1
CGW2
iBGP
Re-advertise VPC CIDR via IGP
HA VPN
To
VPC
Reuse your CGW Public IP
to connect to more VPCs

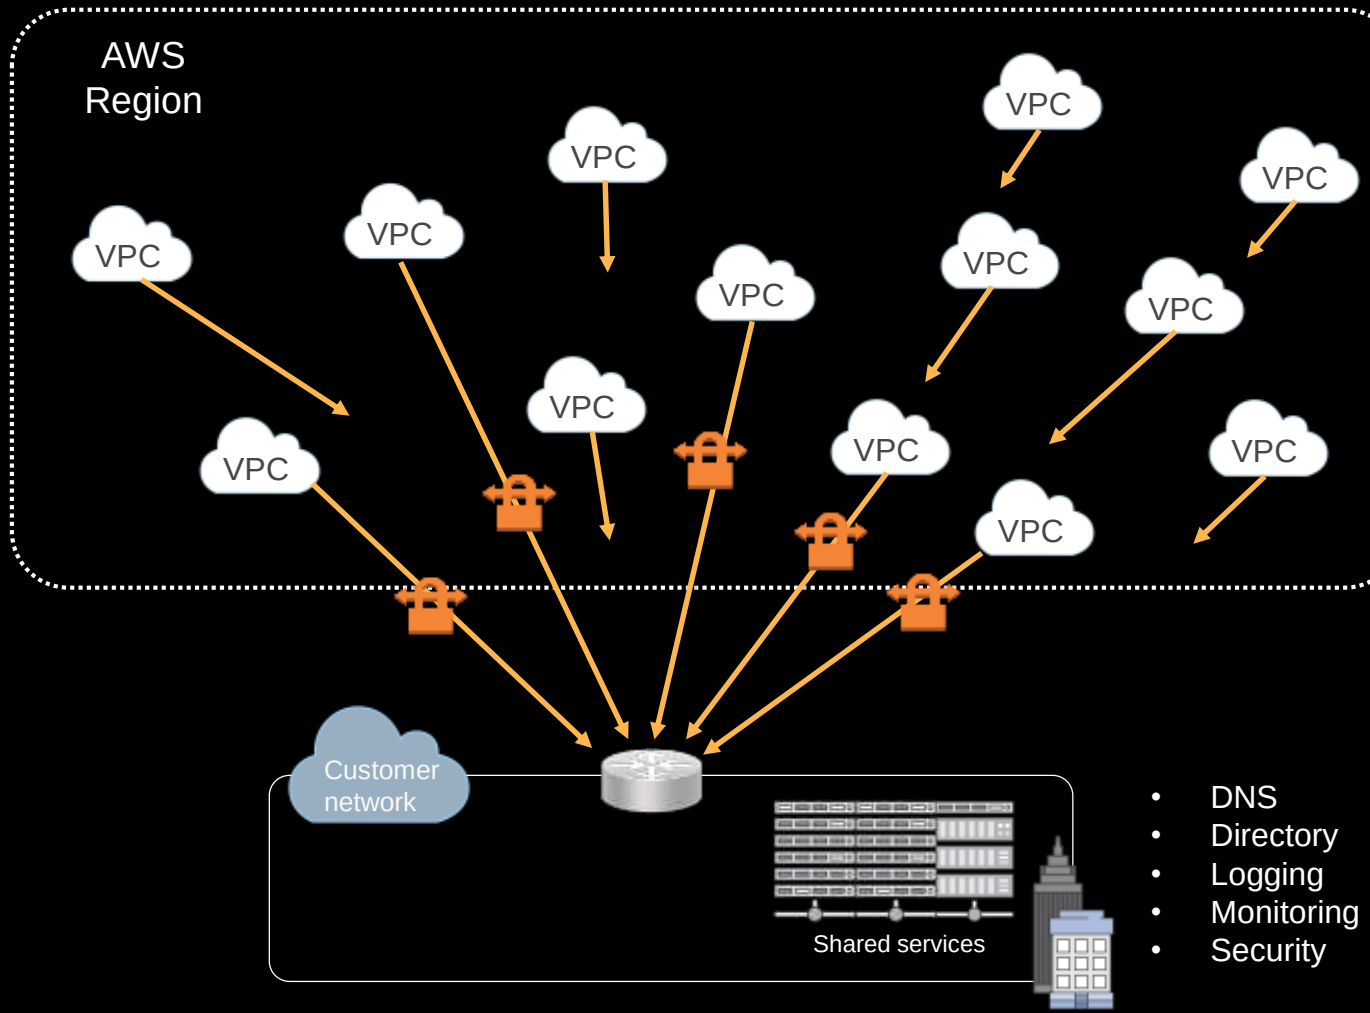
AWS
Region
VPC
VPC
VPC
VPC
VPC
VPC
VPC
VPC
VPC
VPC
VPC
VPC
VPC
Customer
network
Shared services
DNS
Directory
Logging
Monitoring
Security

Evolving design requirements

- Centralize network connectivity to and from cloud
- Centralize management, security, and common services
- Account owners in control of own VPC resources
- Many AWS accounts
- Many VPCs
- One region

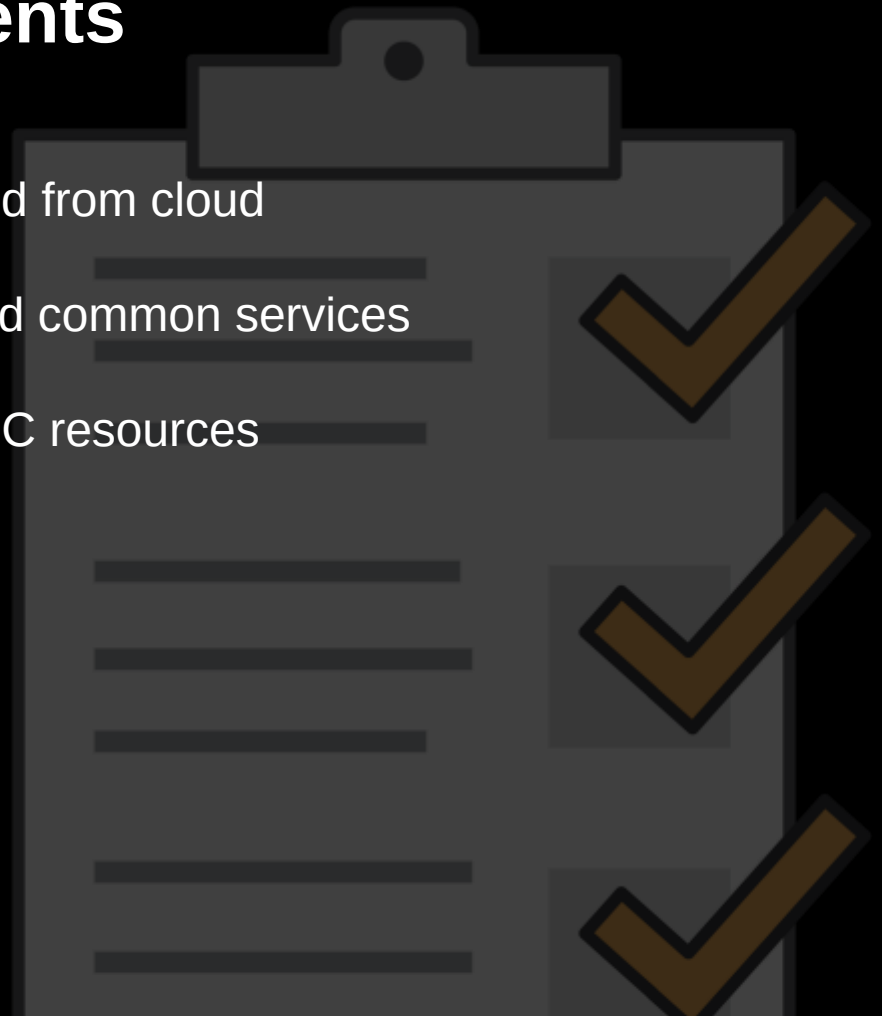

Hub and Spoke with Peering

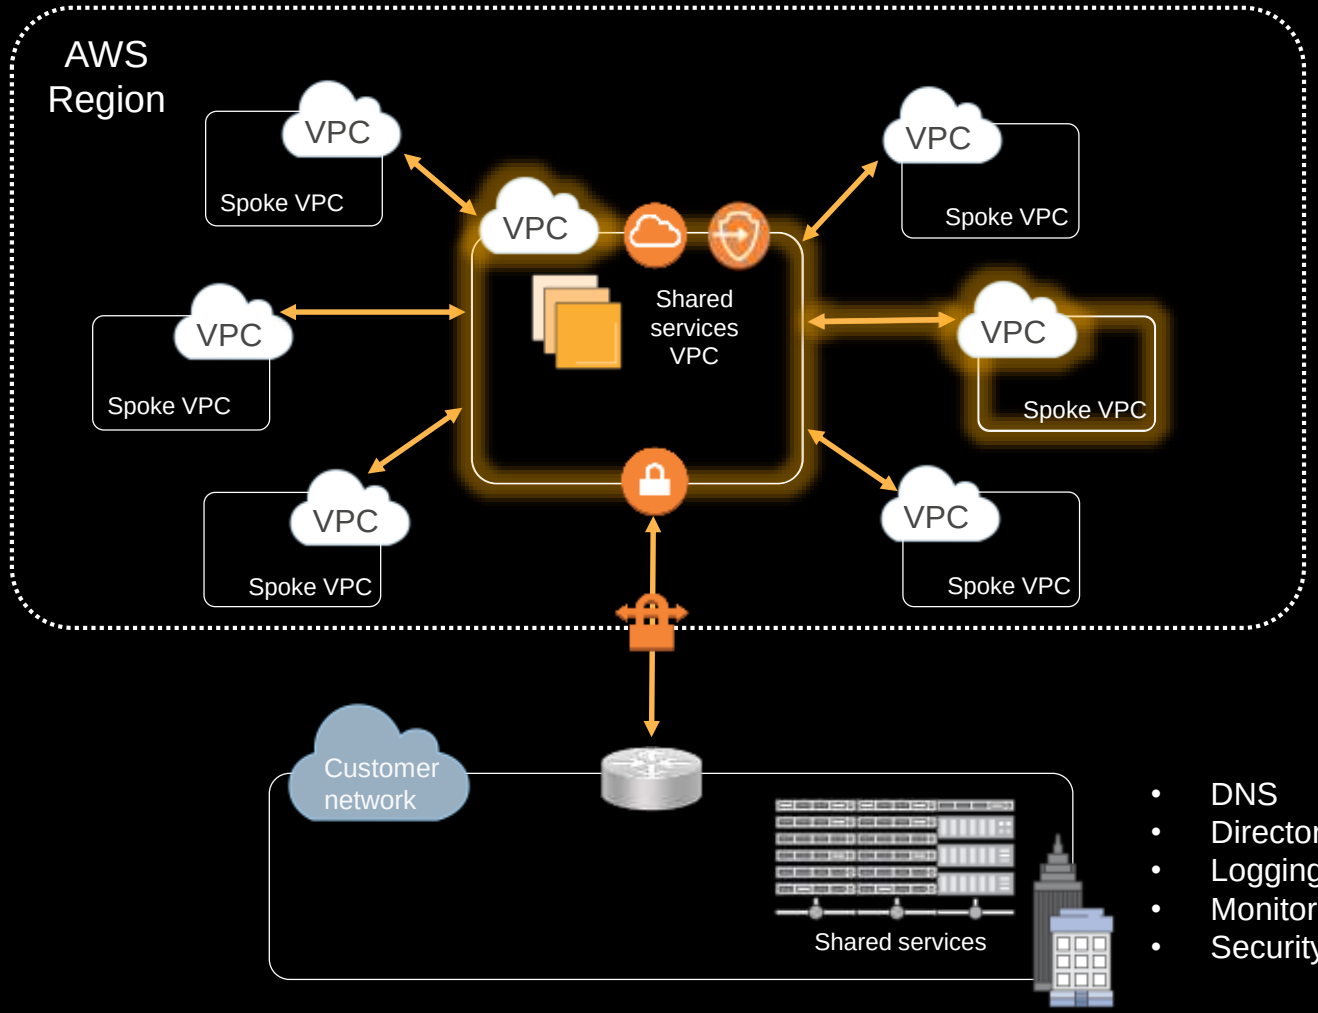

- DNS
- Directory
- Logging
- Monitoring
- Security

VPC peering

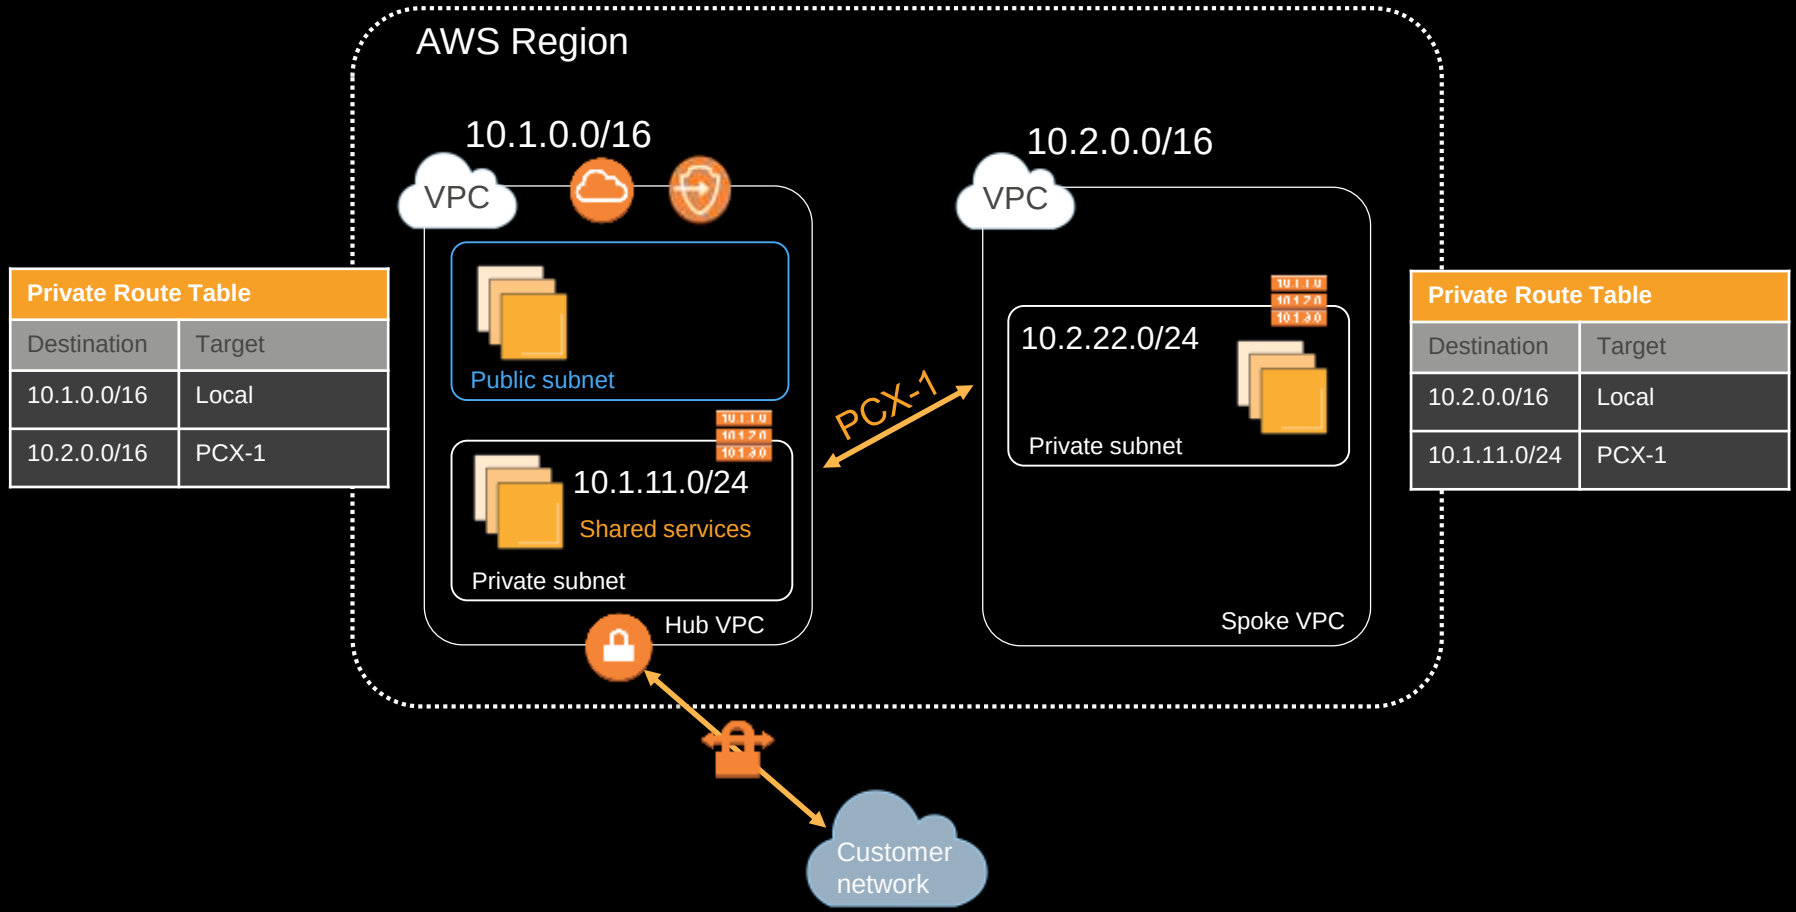

Private Route Table	
Destination	Target
10.1.0.0/16	Local
10.2.0.0/16	PCX-1

Private Route Table	
Destination	Target
10.2.0.0/16	Local
10.1.11.0/24	PCX-1

Edge-to-edge routing

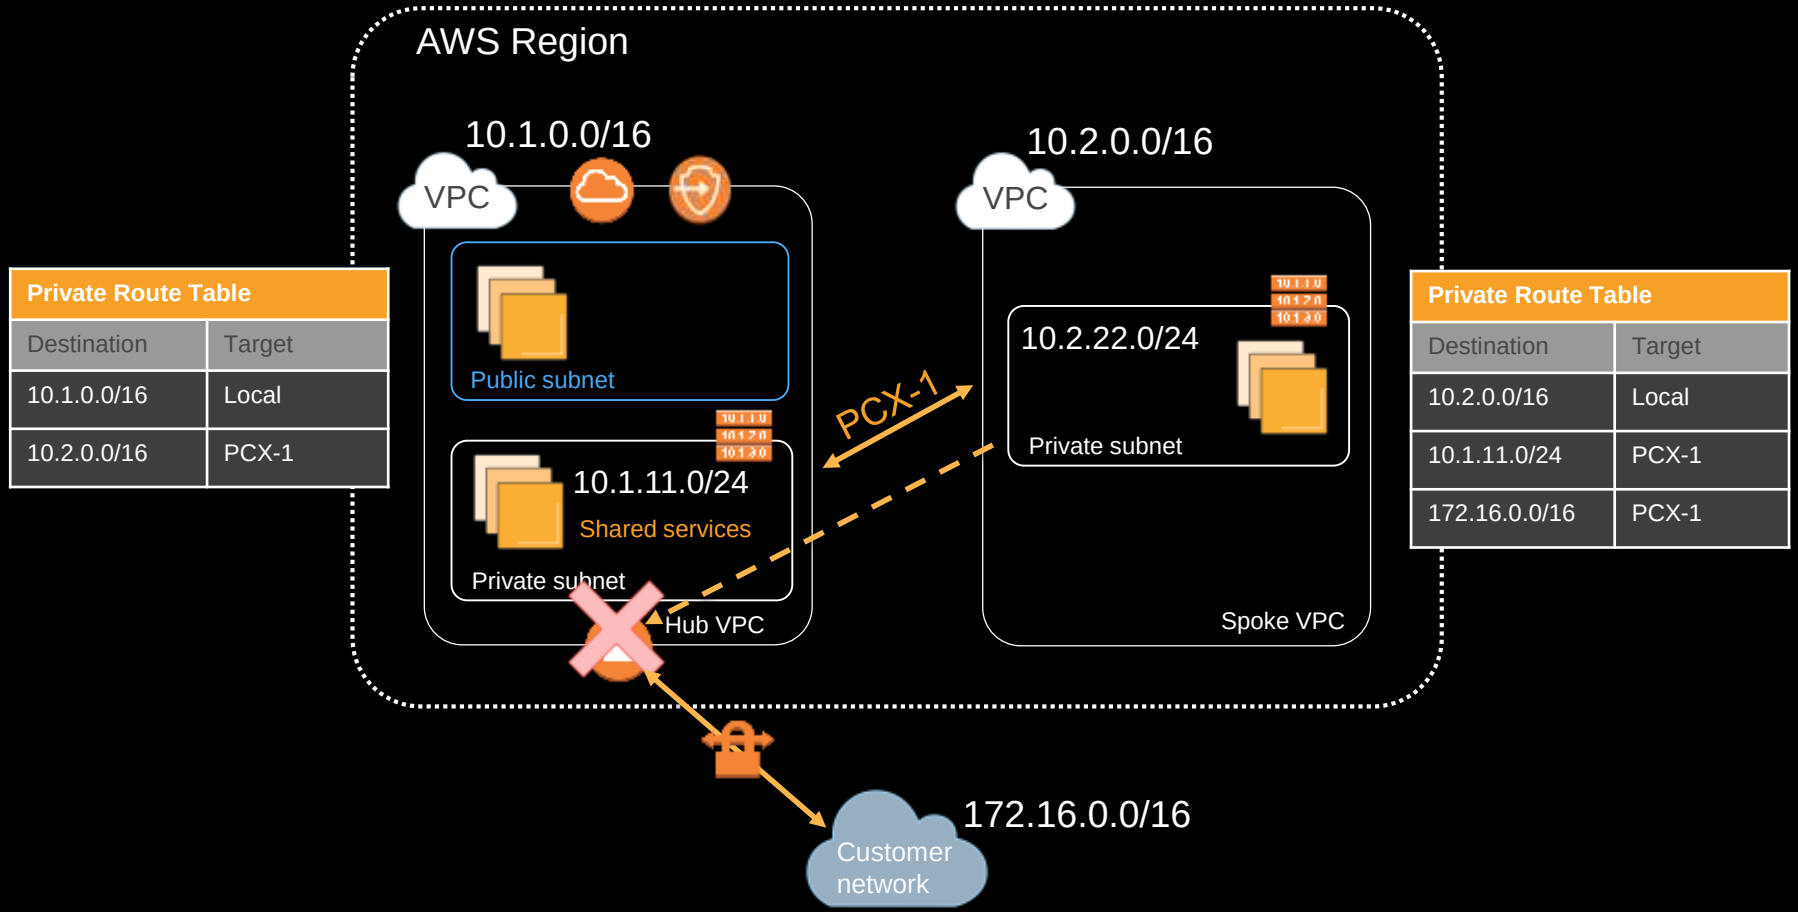

Private Route Table	
Destination	Target
10.1.0.0/16	Local
10.2.0.0/16	PCX-1

Private Route Table	
Destination	Target
10.2.0.0/16	Local
10.1.11.0/24	PCX-1
172.16.0.0/16	PCX-1

Edge-to-edge via proxy

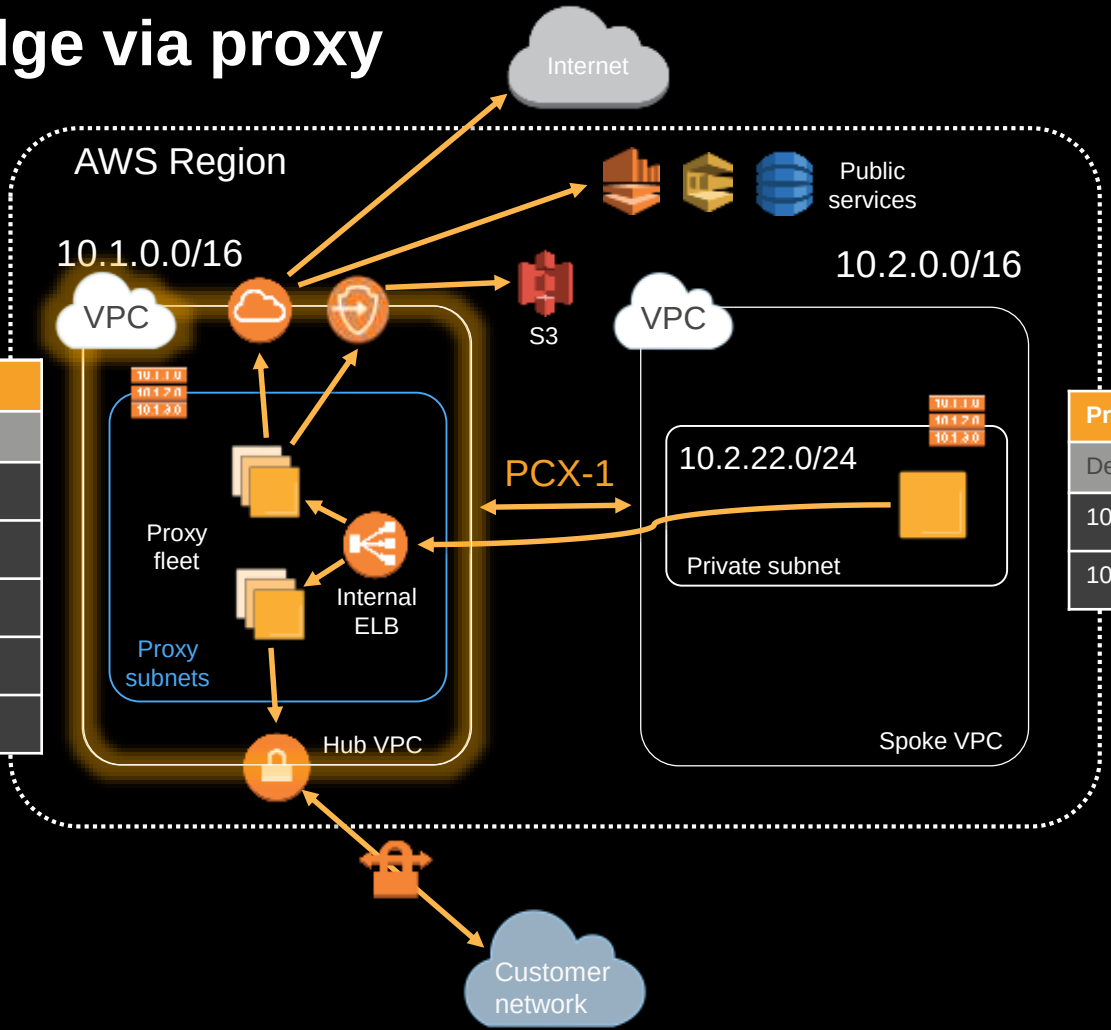

Proxy Route Table	
Destination	Target
10.1.0.0/16	Local
10.2.0.0/16	PCX-1
172.16.0.0/16	VGW
0.0.0.0/0	IGW
S3 Prefix List	VPCE

Private Route Table	
Destination	Target
10.2.0.0/16	Local
10.1.0.0/16	PCX-1

Pro & Con: Shared Services Hub and Spoke

- Minimizes on premises network change
- Reduces latency, cost of cloud applications accessing common services
- Provides spoke accounts control over own resources
- But controls and secures egress traffic from spokes
- Security Groups work across peers

- Cost and management of central proxy layer
- Not a transparent proxy
- Configuring end devices to use proxy
- Restricted to HTTP/S
- No transitive networking
- Peering data transfer cost

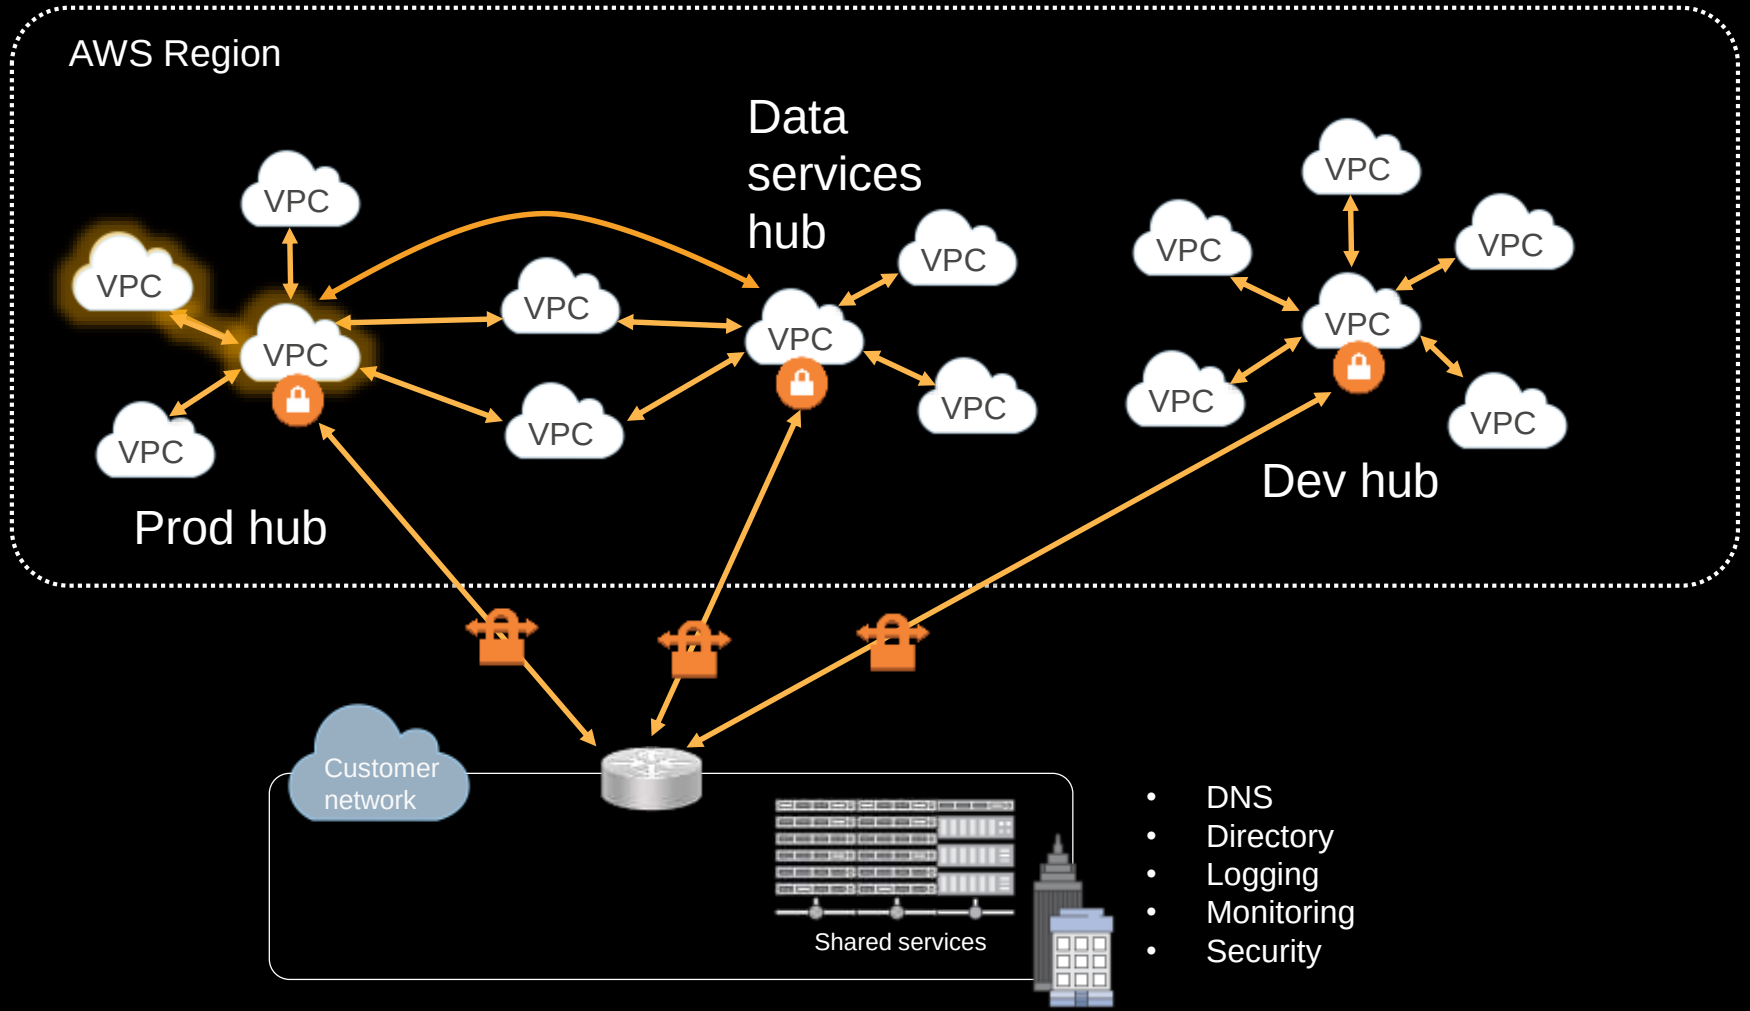
AWS Region
Data services hub
VPC
Prod hub
Dev hub
Customer network
Shared services
DNS
Directory
Logging
Monitoring
Security

Transit VPC

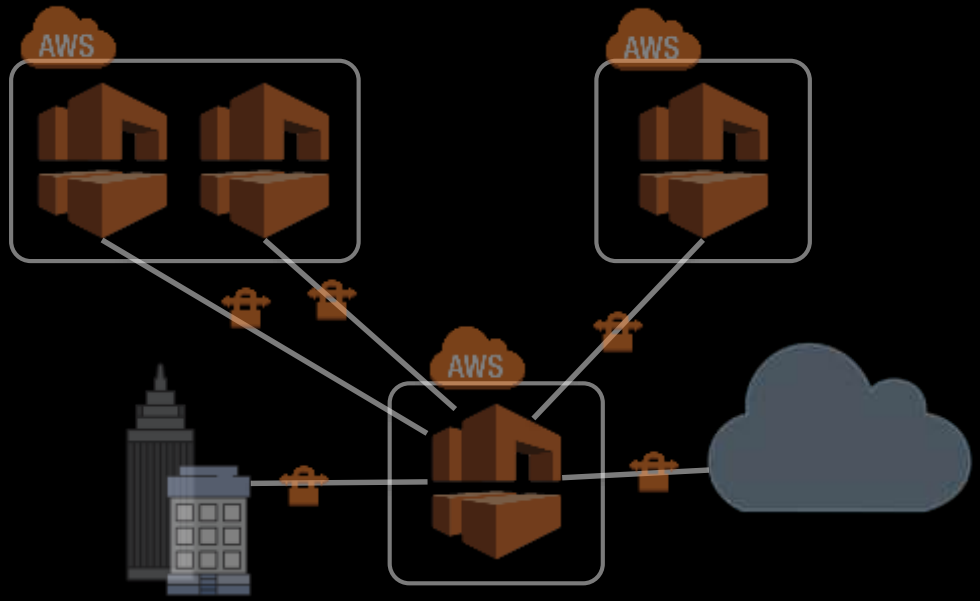

https://aws.amazon.com/answers/networking/transit-vpc/

Transit VPC

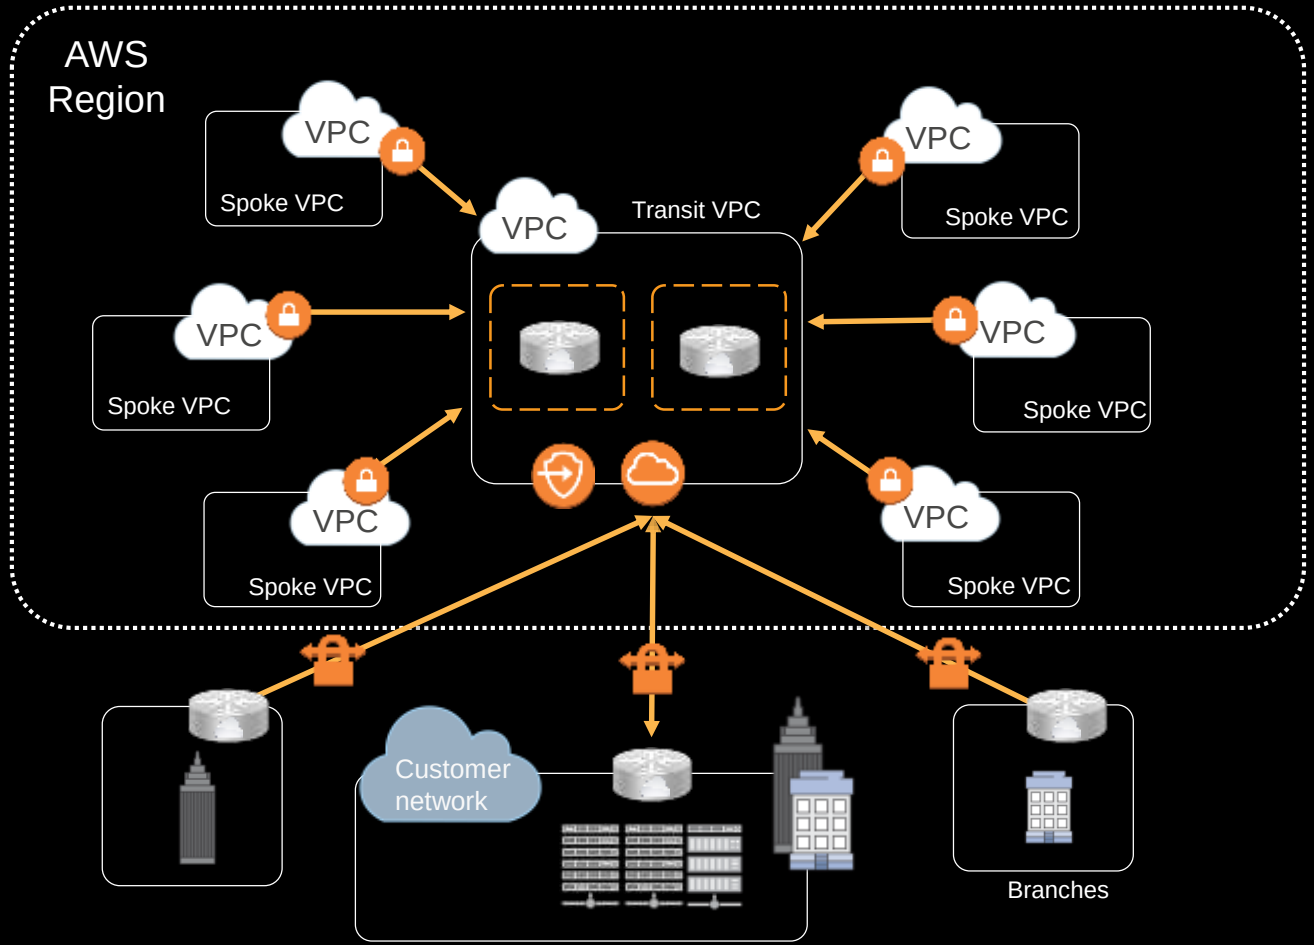
AWS Region
VPC
Spoke VPC
VPC
Transit VPC
VPC
Spoke VPC
VPC
Spoke VPC
VPC
Spoke VPC
VPC
Spoke VPC
VPC
Spoke VPC
Customer network
Branches

Pro & Con: Transit VPC

- End to End routing between VPCs in all regions and any other non-AWS network
- Central transit routers can perform higher level networking and security functions
- Spoke VGWs are HA by default
- Minimizes on premises networking changes
- Can minimize cost if replacing on premises or colo networking hardware

- Availability and management of transit router instances
- Licensing costs
- Cost of data transfer between transit, spokes and other networks

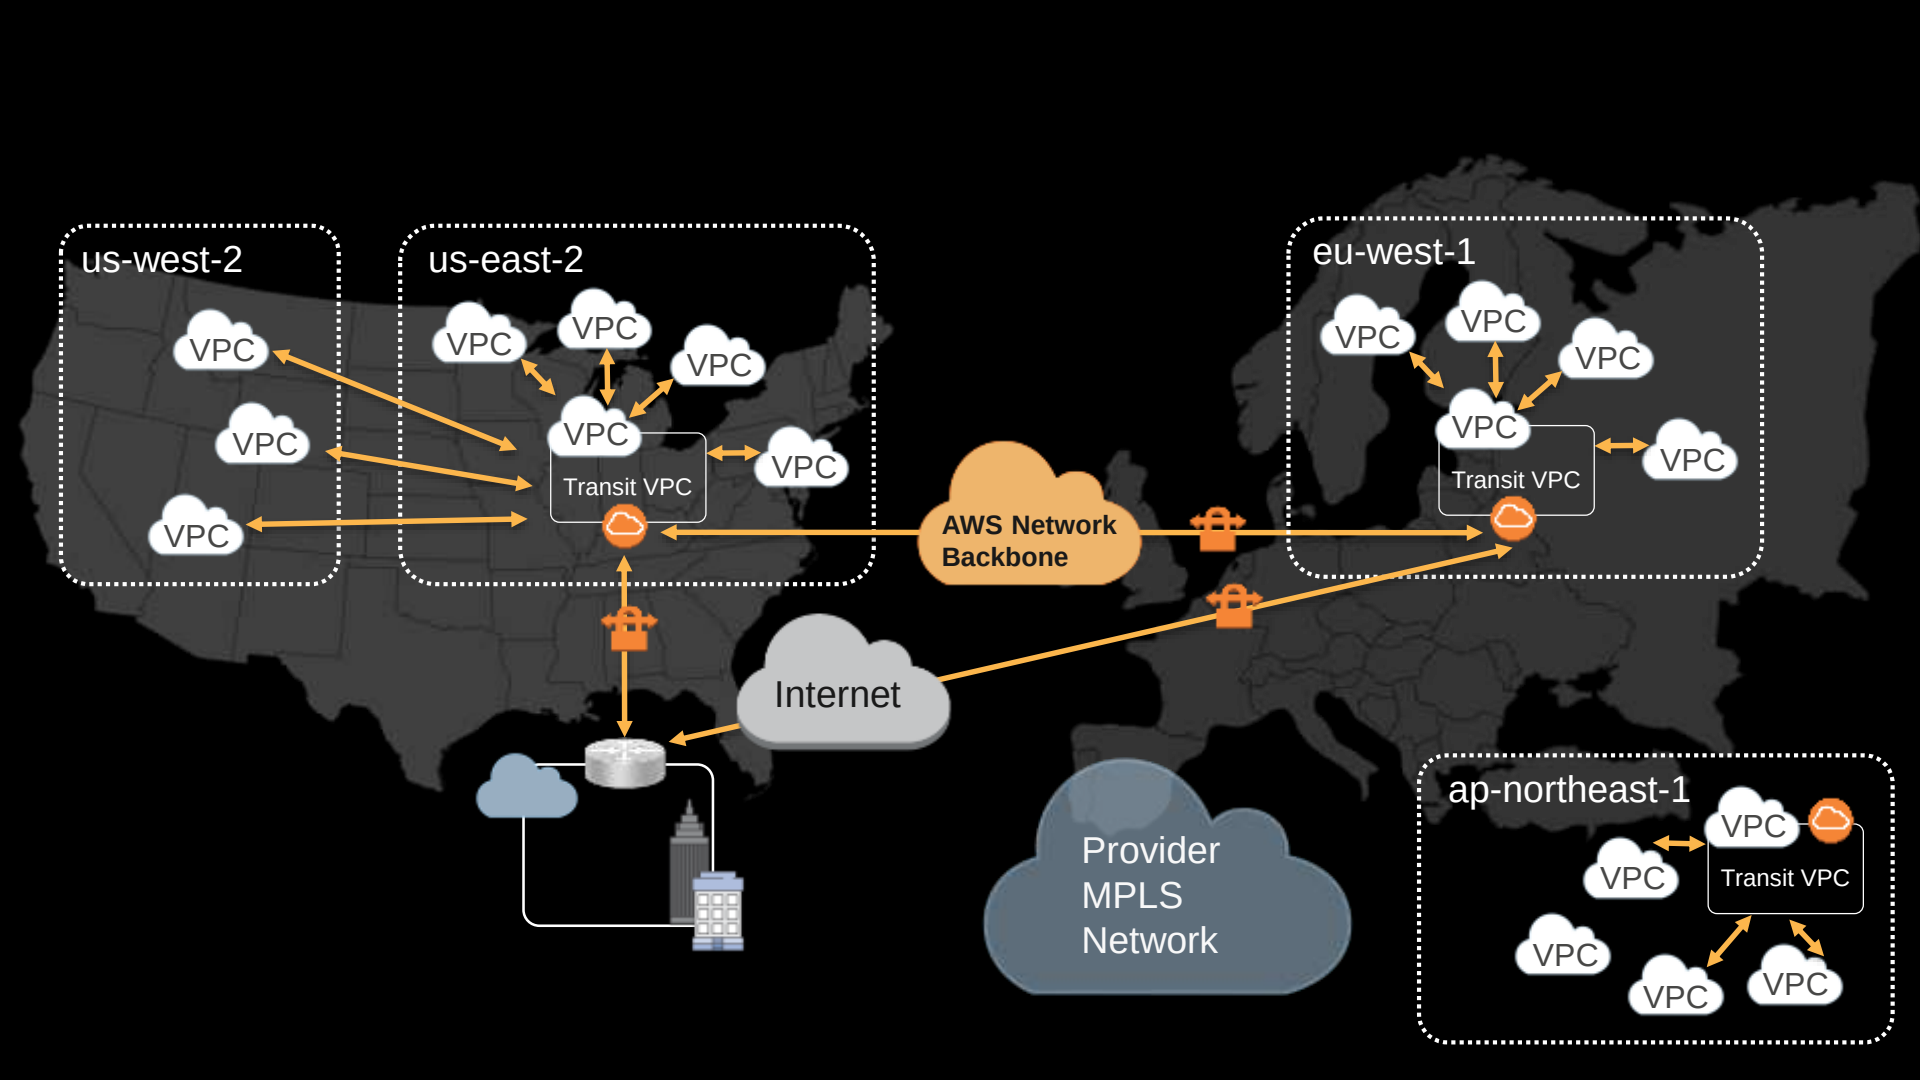
us-west-2
us-east-2
eu-west-1
ap-northeast-1
VPC
Transit VPC
AWS Network Backbone
Internet
Provider MPLS Network

Direct Connect

How do you provision AWS Direct Connect?

1. Build your AWS infrastructure
2. Create your Virtual Private Gateway (VGW) and attach to your Virtual Private Cloud (VPC)
3. Order an AWS Direct Connect from the console or through a Direct Connect Partner

at&t Level(3) verizon + More

4. Have your cross connect provisioned from the AWS router to your device or your partners device (or use a partners NNI)
5. Build connectivity if not already available through partner back to on-premises
6. Provision your Virtual interfaces (private or public) and start using your AWS Direct Connect.

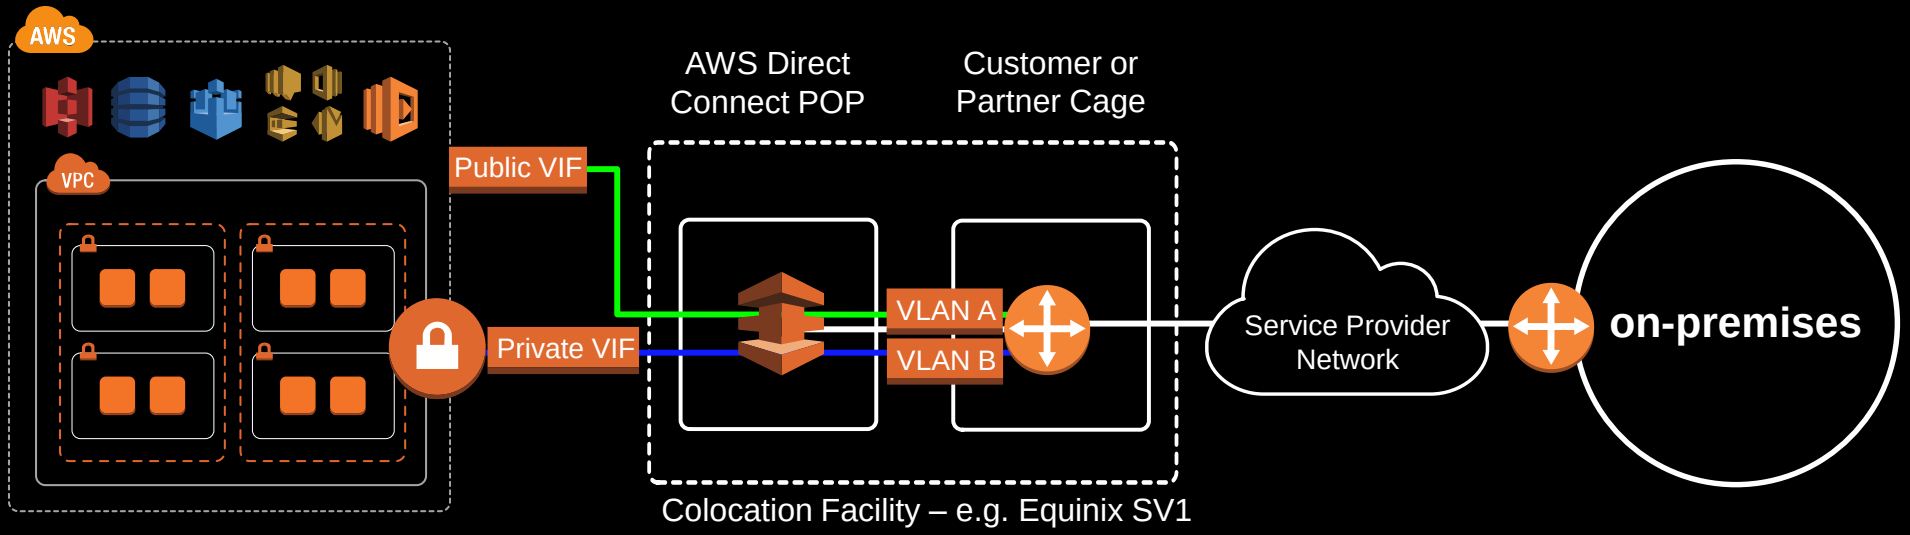

AWS Direct Connect (DX) in the United States

AWS Direct Connect (DX) in Canada

Partners:

- Cologix
- Console
- Global Telecom and Technology
- Hibernia Networks
- Level 3
- Netelligent
- Orange Business Services
- Sohonet
- Tata Communications
- Zayo Group

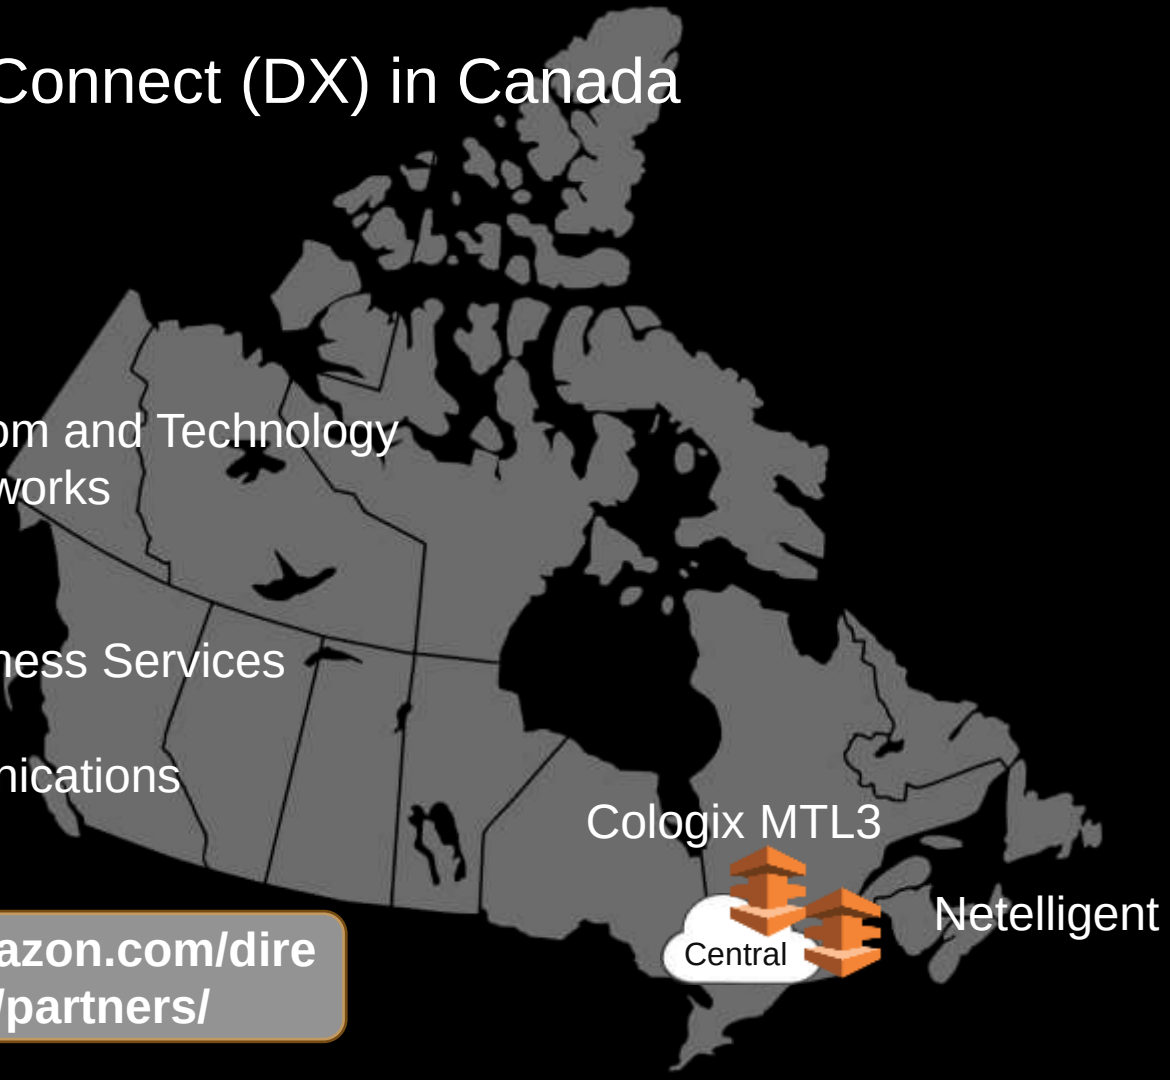

https://aws.amazon.com/directconnect/partners/

Direct Connect – physical connectivity

1) Customer presence in the same DX location

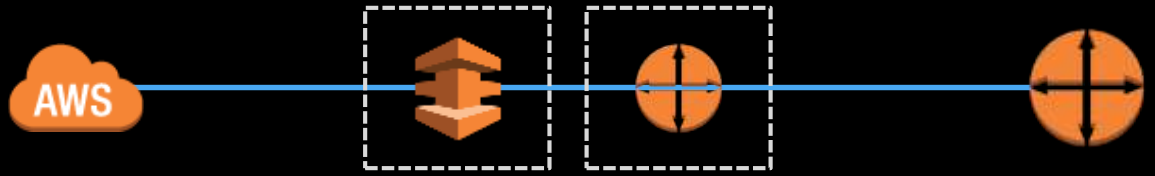

2) Circuit between customer data center and DX location

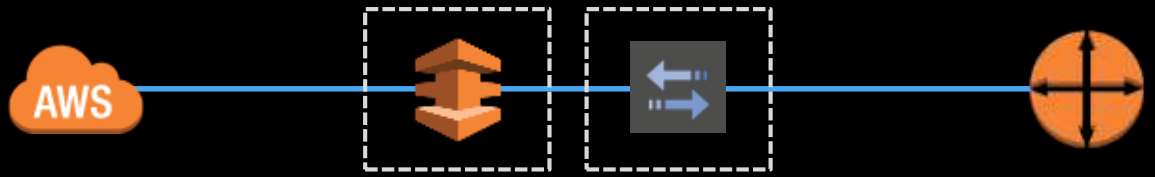

3) Service provider network extending to DX location

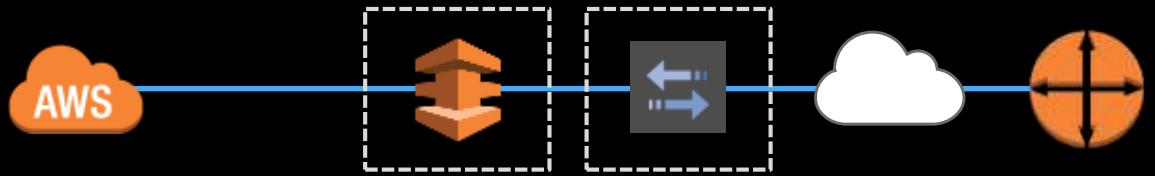

1) Customer router in colo

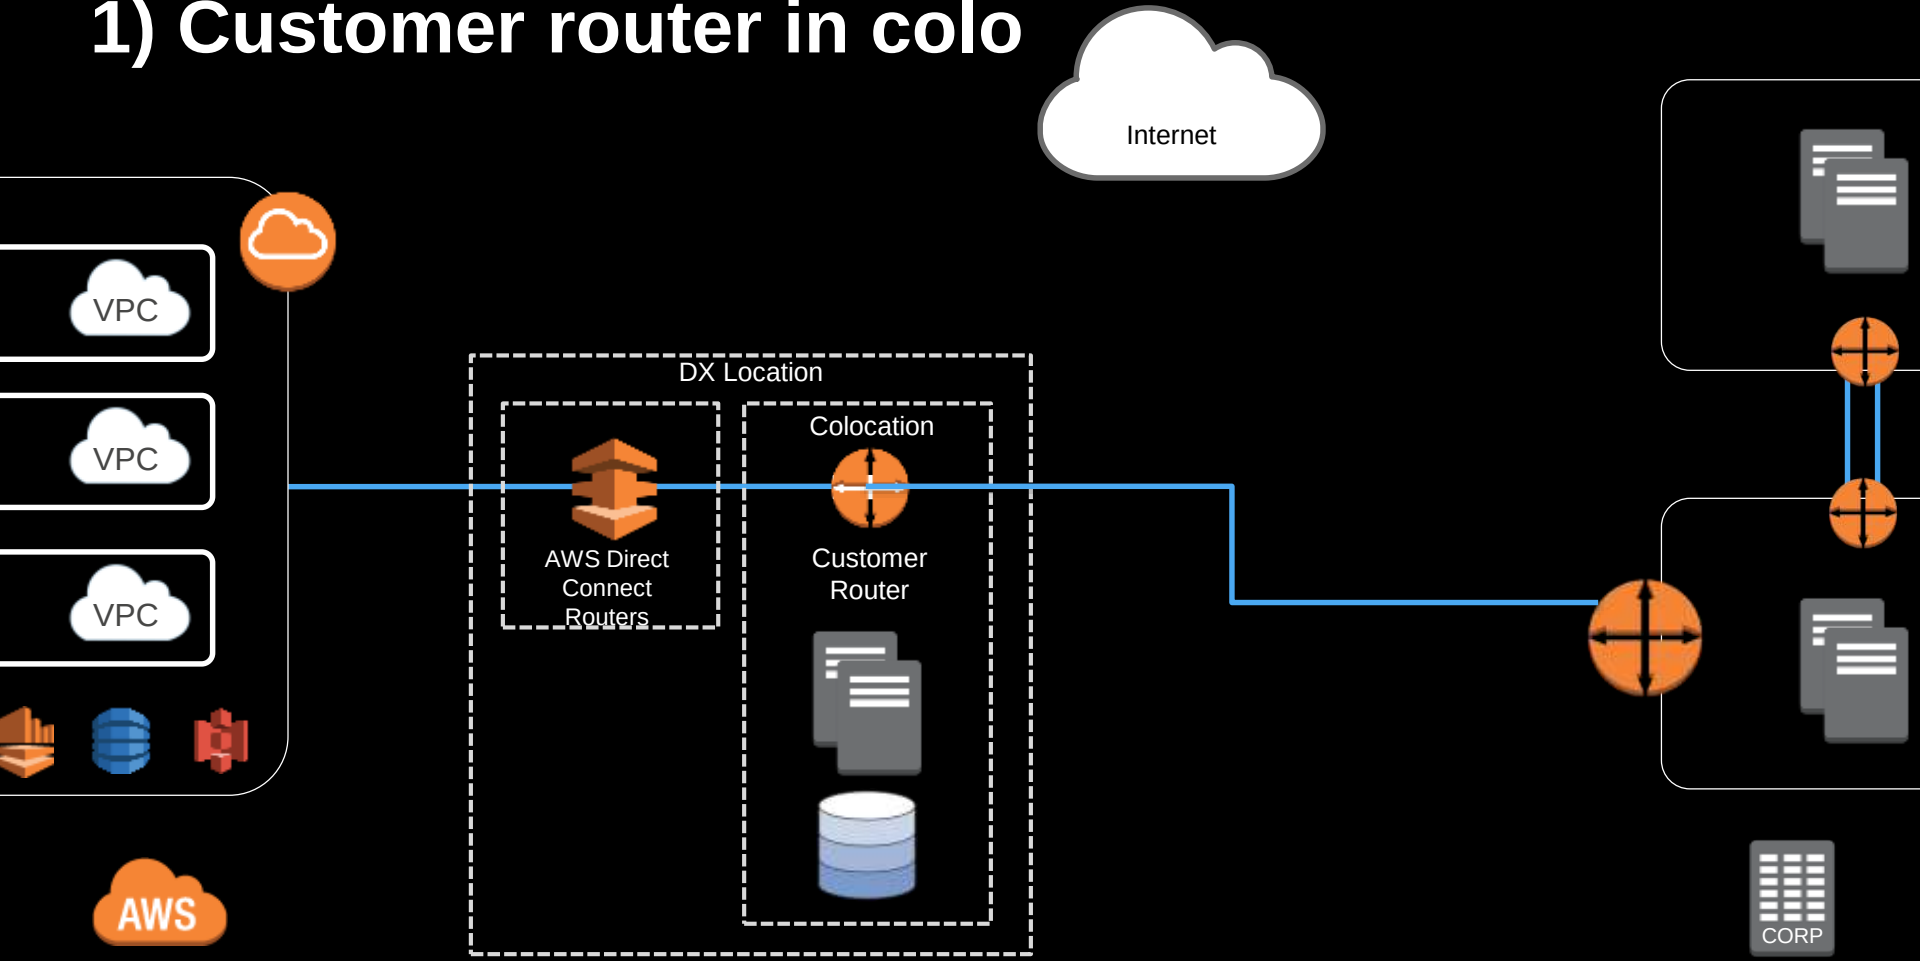

2) Partner-built circuit

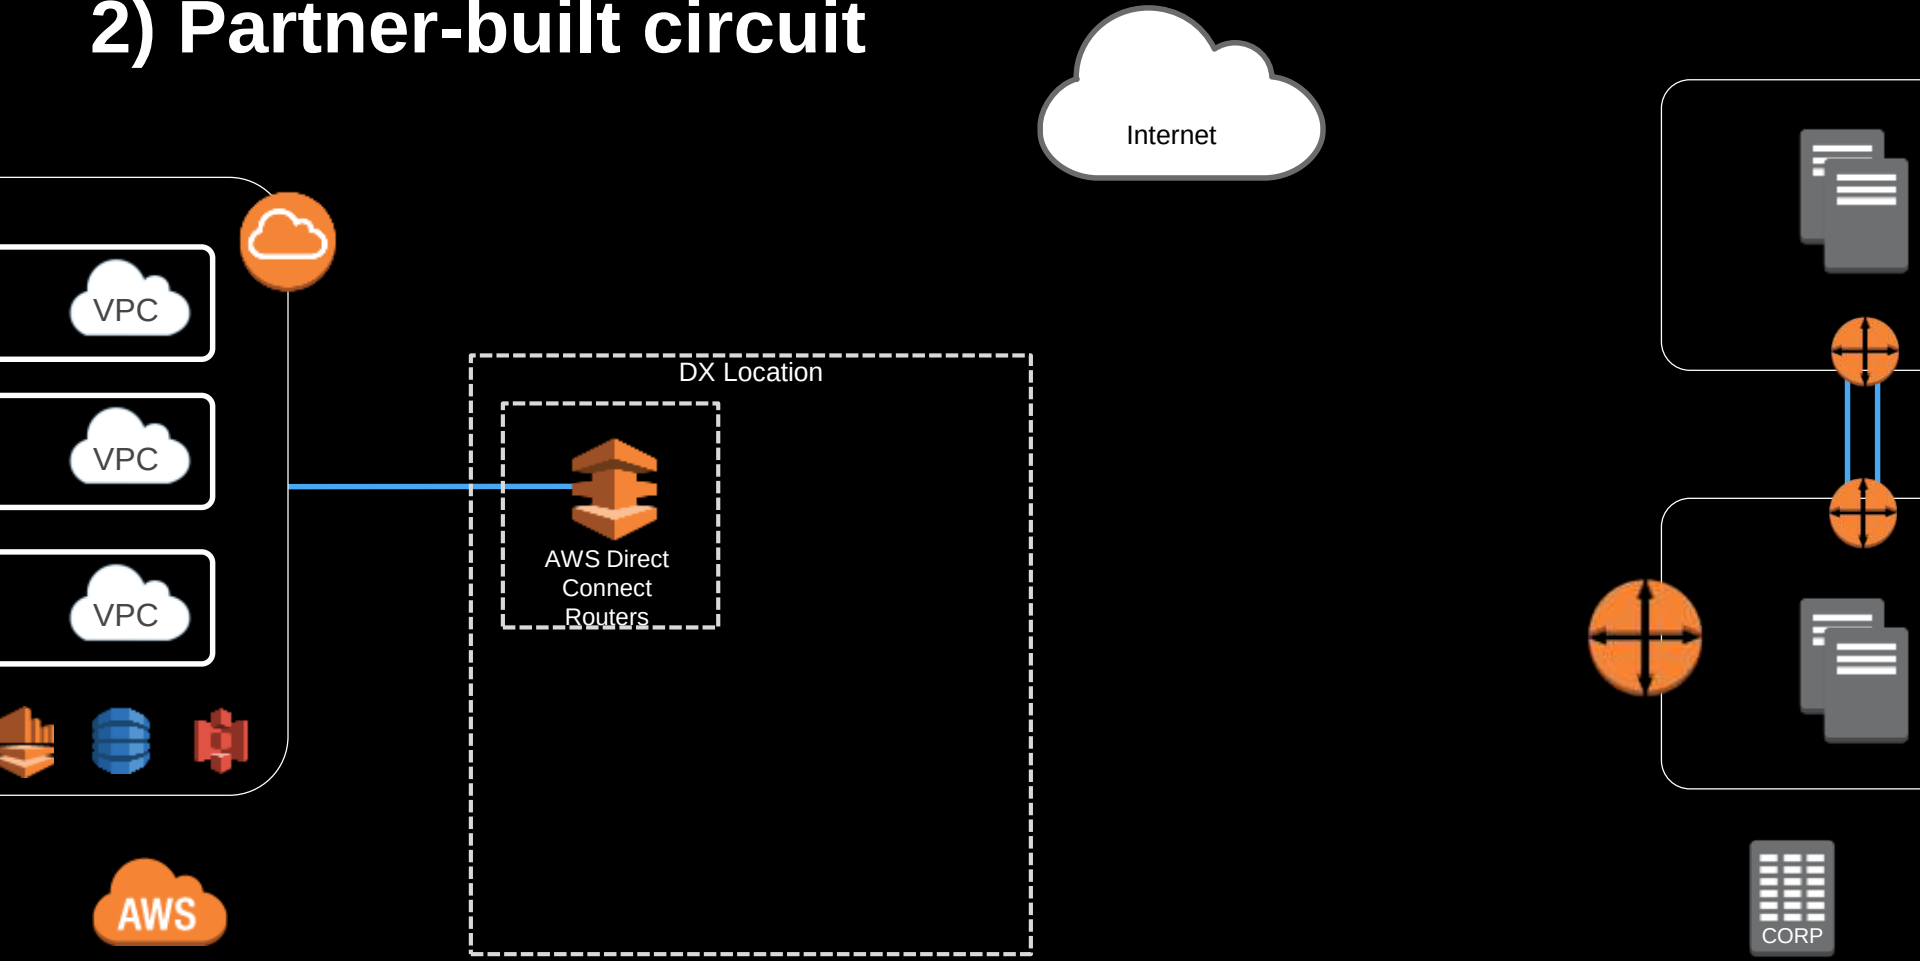

2) Partner-built circuit

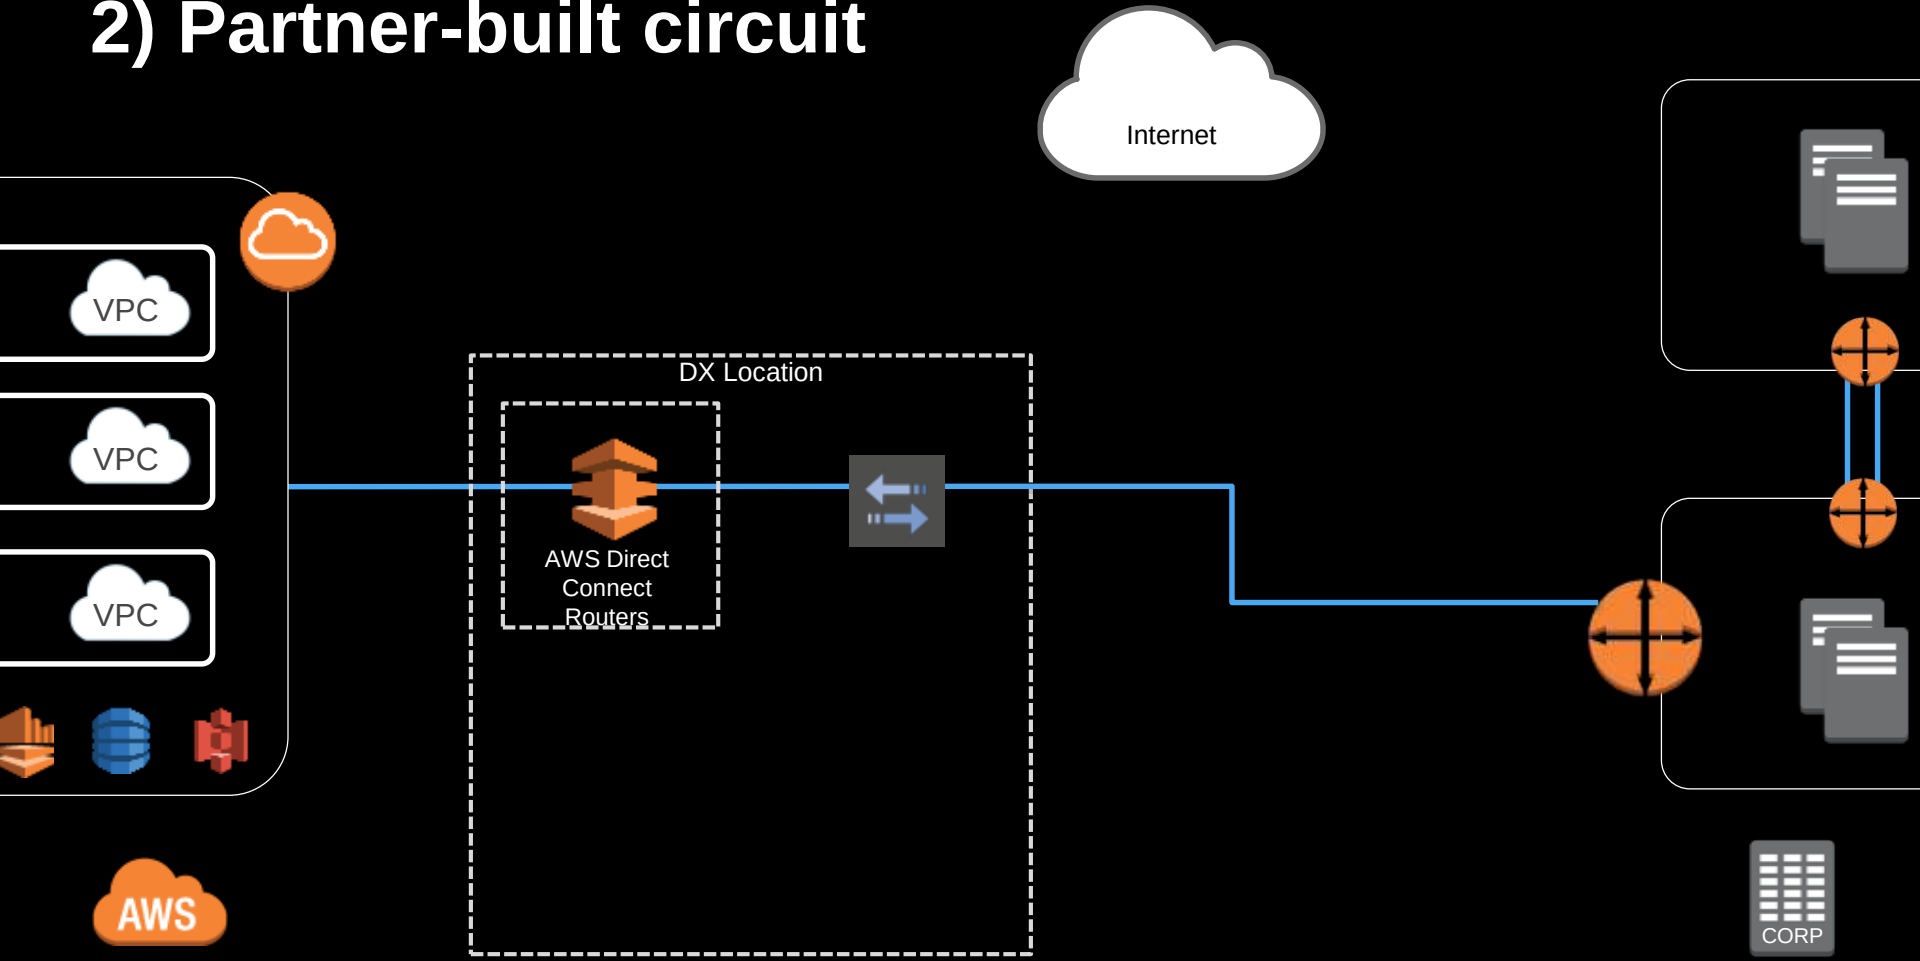

3) Service provider network

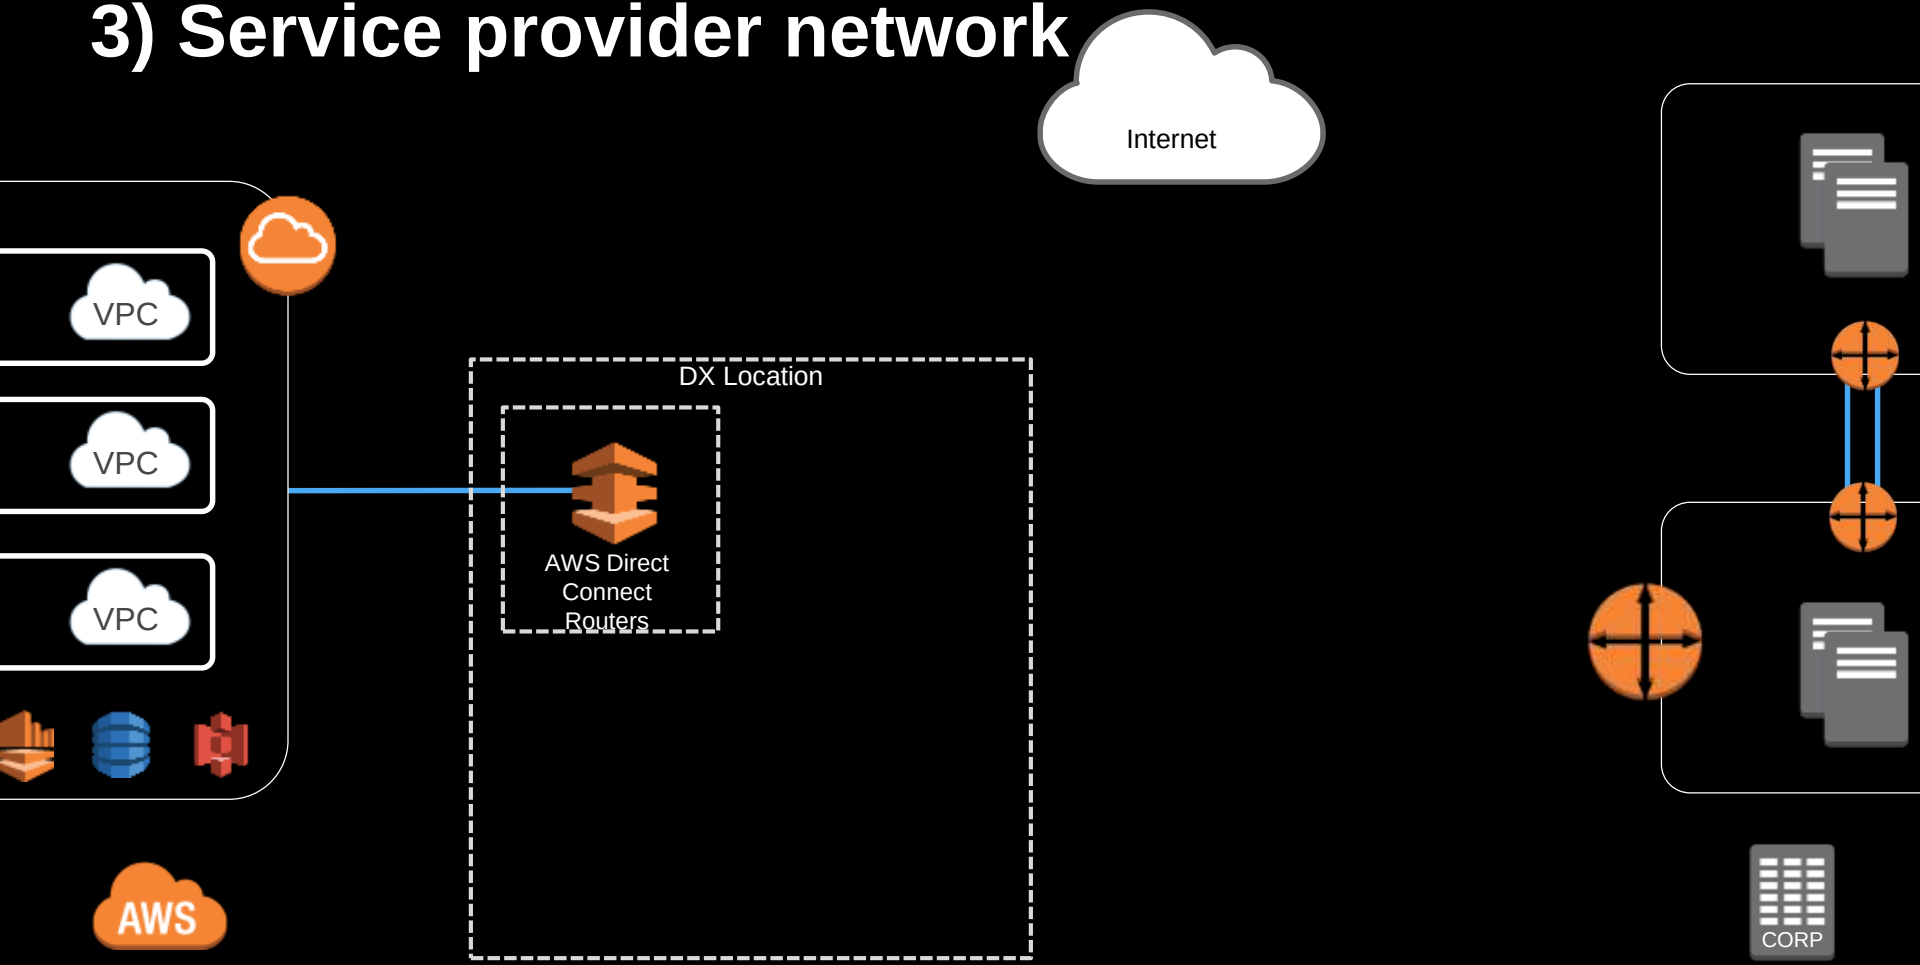

3) Service provider network

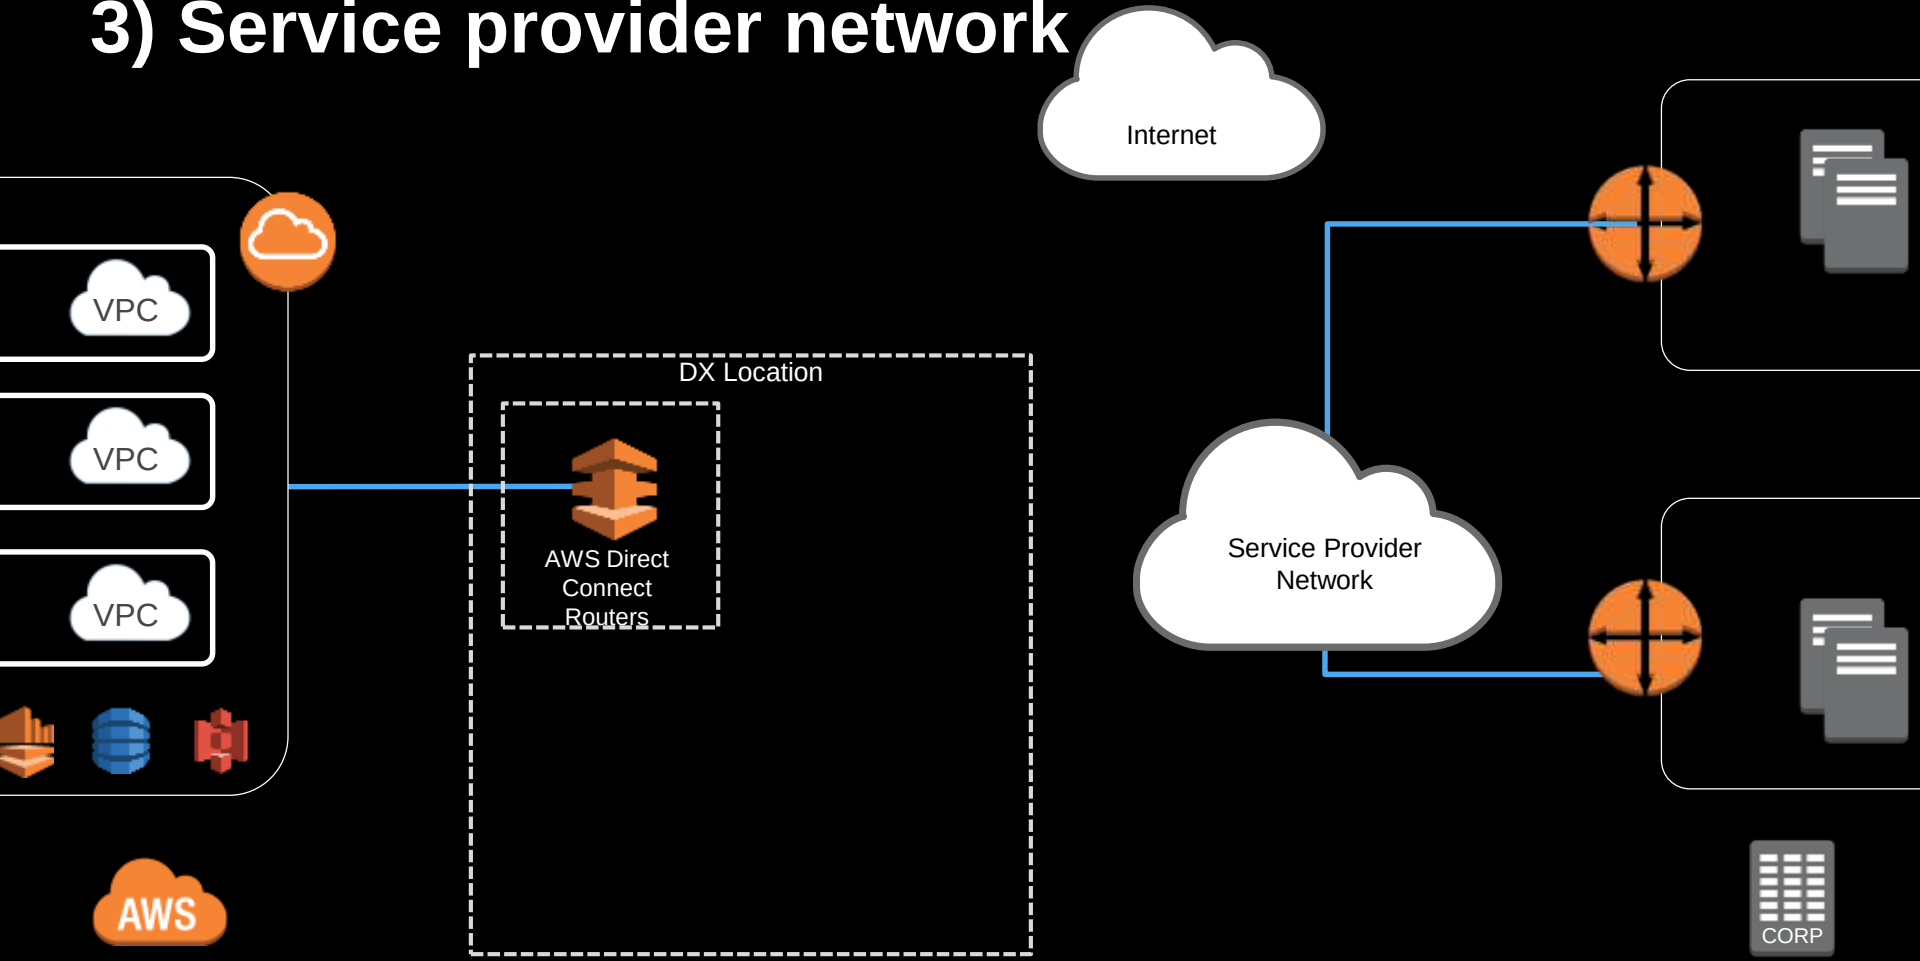

3) Service provider network

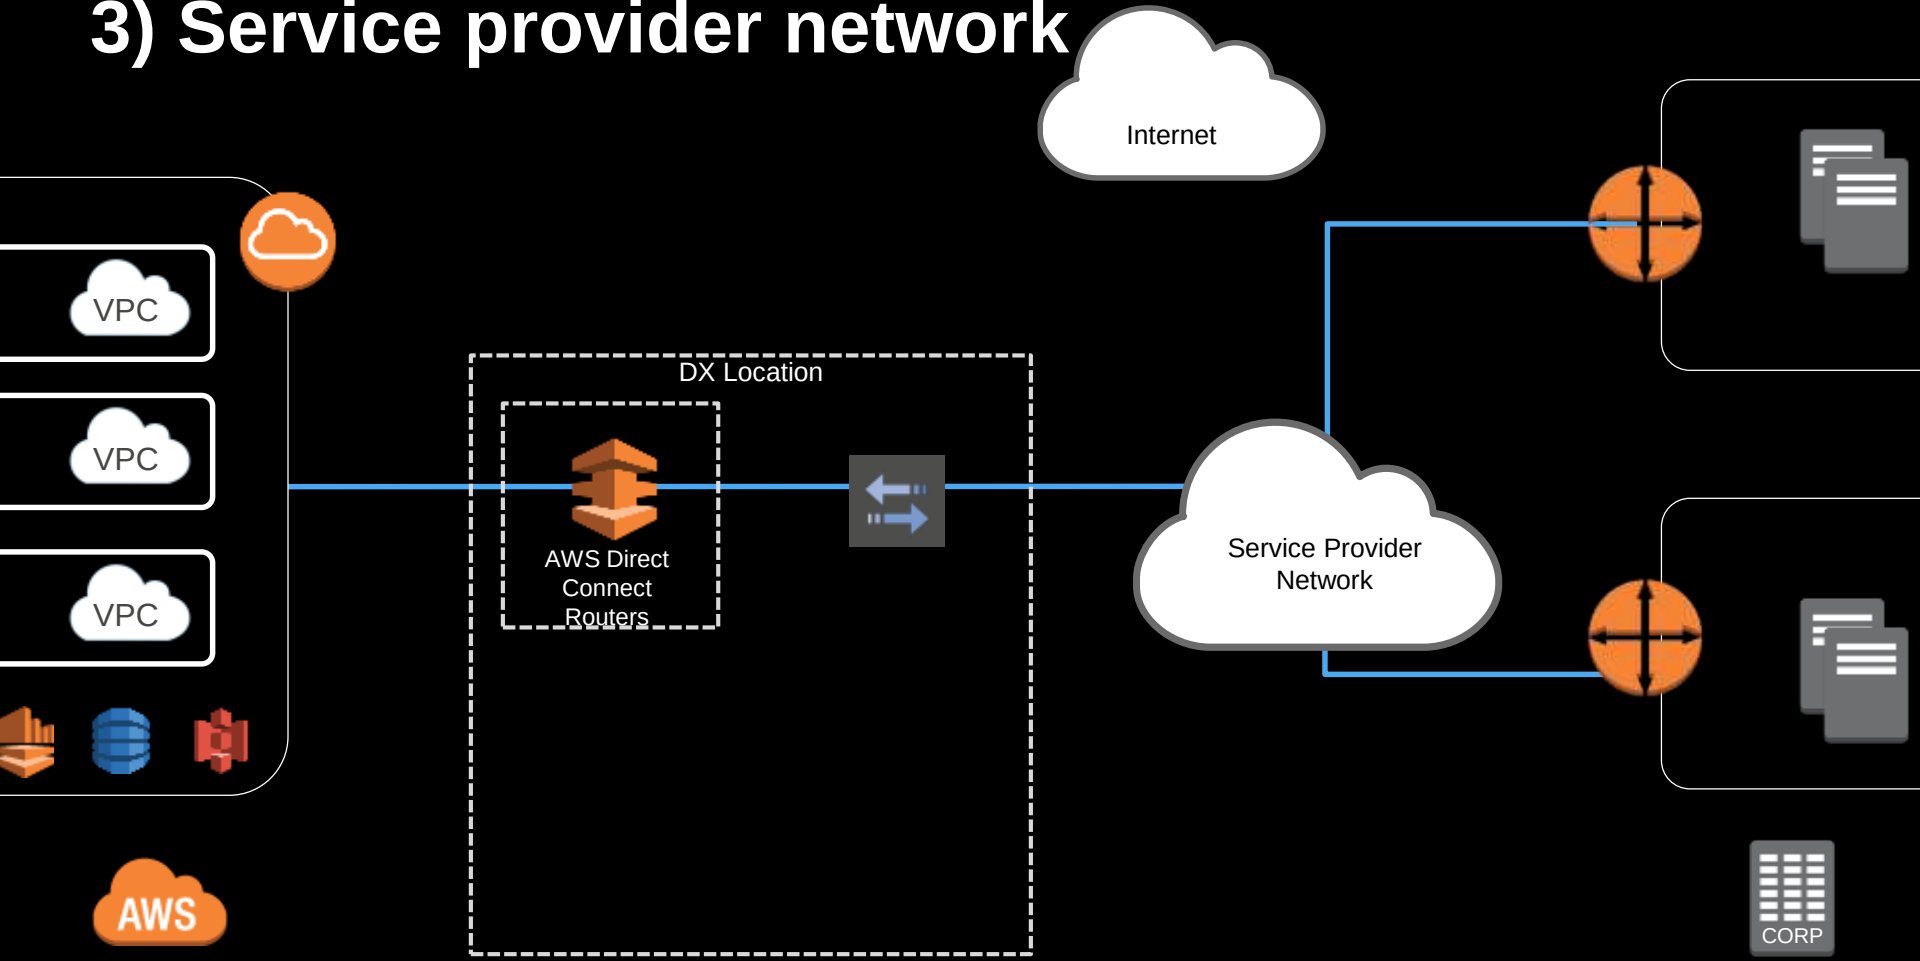

DX physical connectivity considerations

- [x] AWS account that owns the DX port?
- [x] Adding/removing virtual interfaces?
- [x] Routing ownership?
- [x] End-to-end costs?

Direct Connect cost considerations

- Port hour + data transfer
- Data in $0; data out differs by region
- Factor in circuit costs
- Calculate data center Internet costs (VPN)

Related re:Invent Sessions

GPSISV4 – Hybrid Architecture Design

NET305 – Extending Data Centers to the Cloud

NET402 – Deep Dive: AWS Direct Connect and VPNs

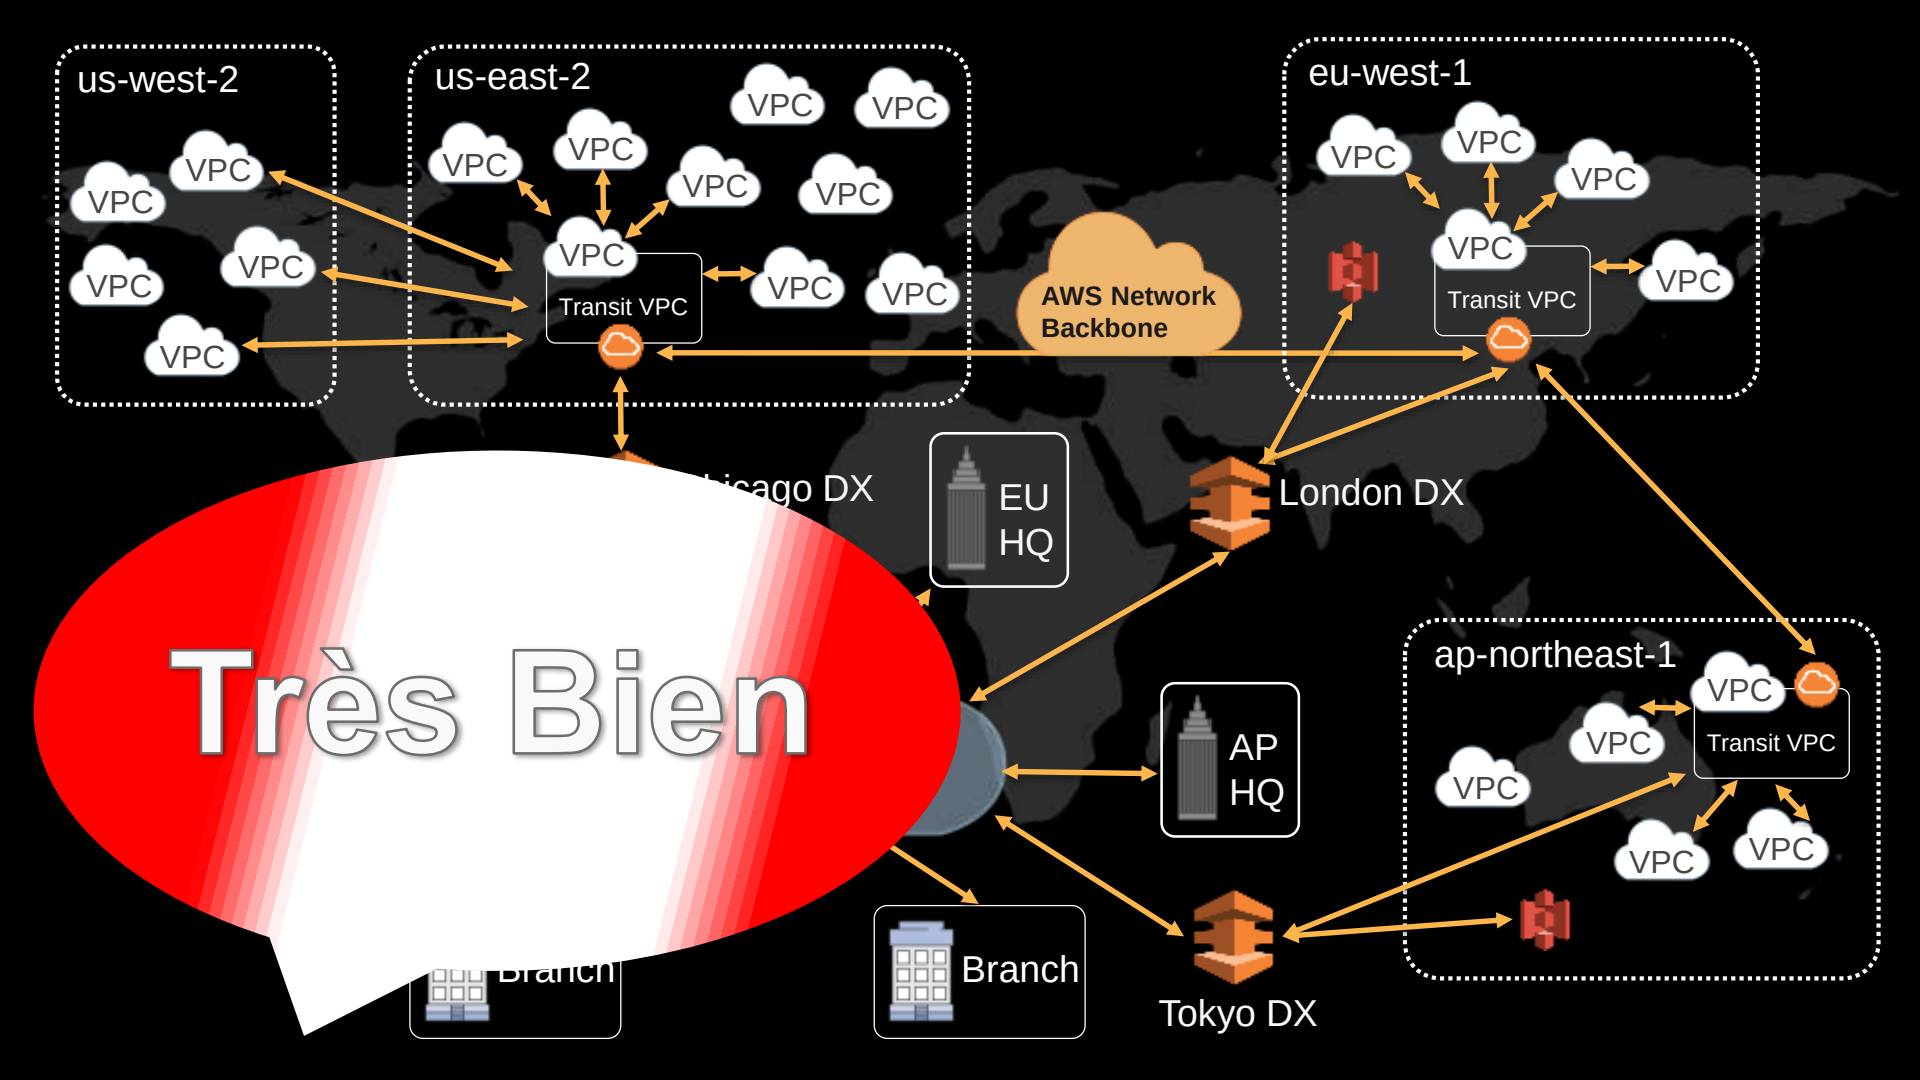
us-west-2
us-east-2
eu-west-1
VPC
Transit VPC
AWS Network Backbone
London DX
EU HQ
AP HQ
Tokyo DX
Branch
ap-northeast-1
Très Bien

Thank you!

amazon web services

Remember to complete your evaluations!