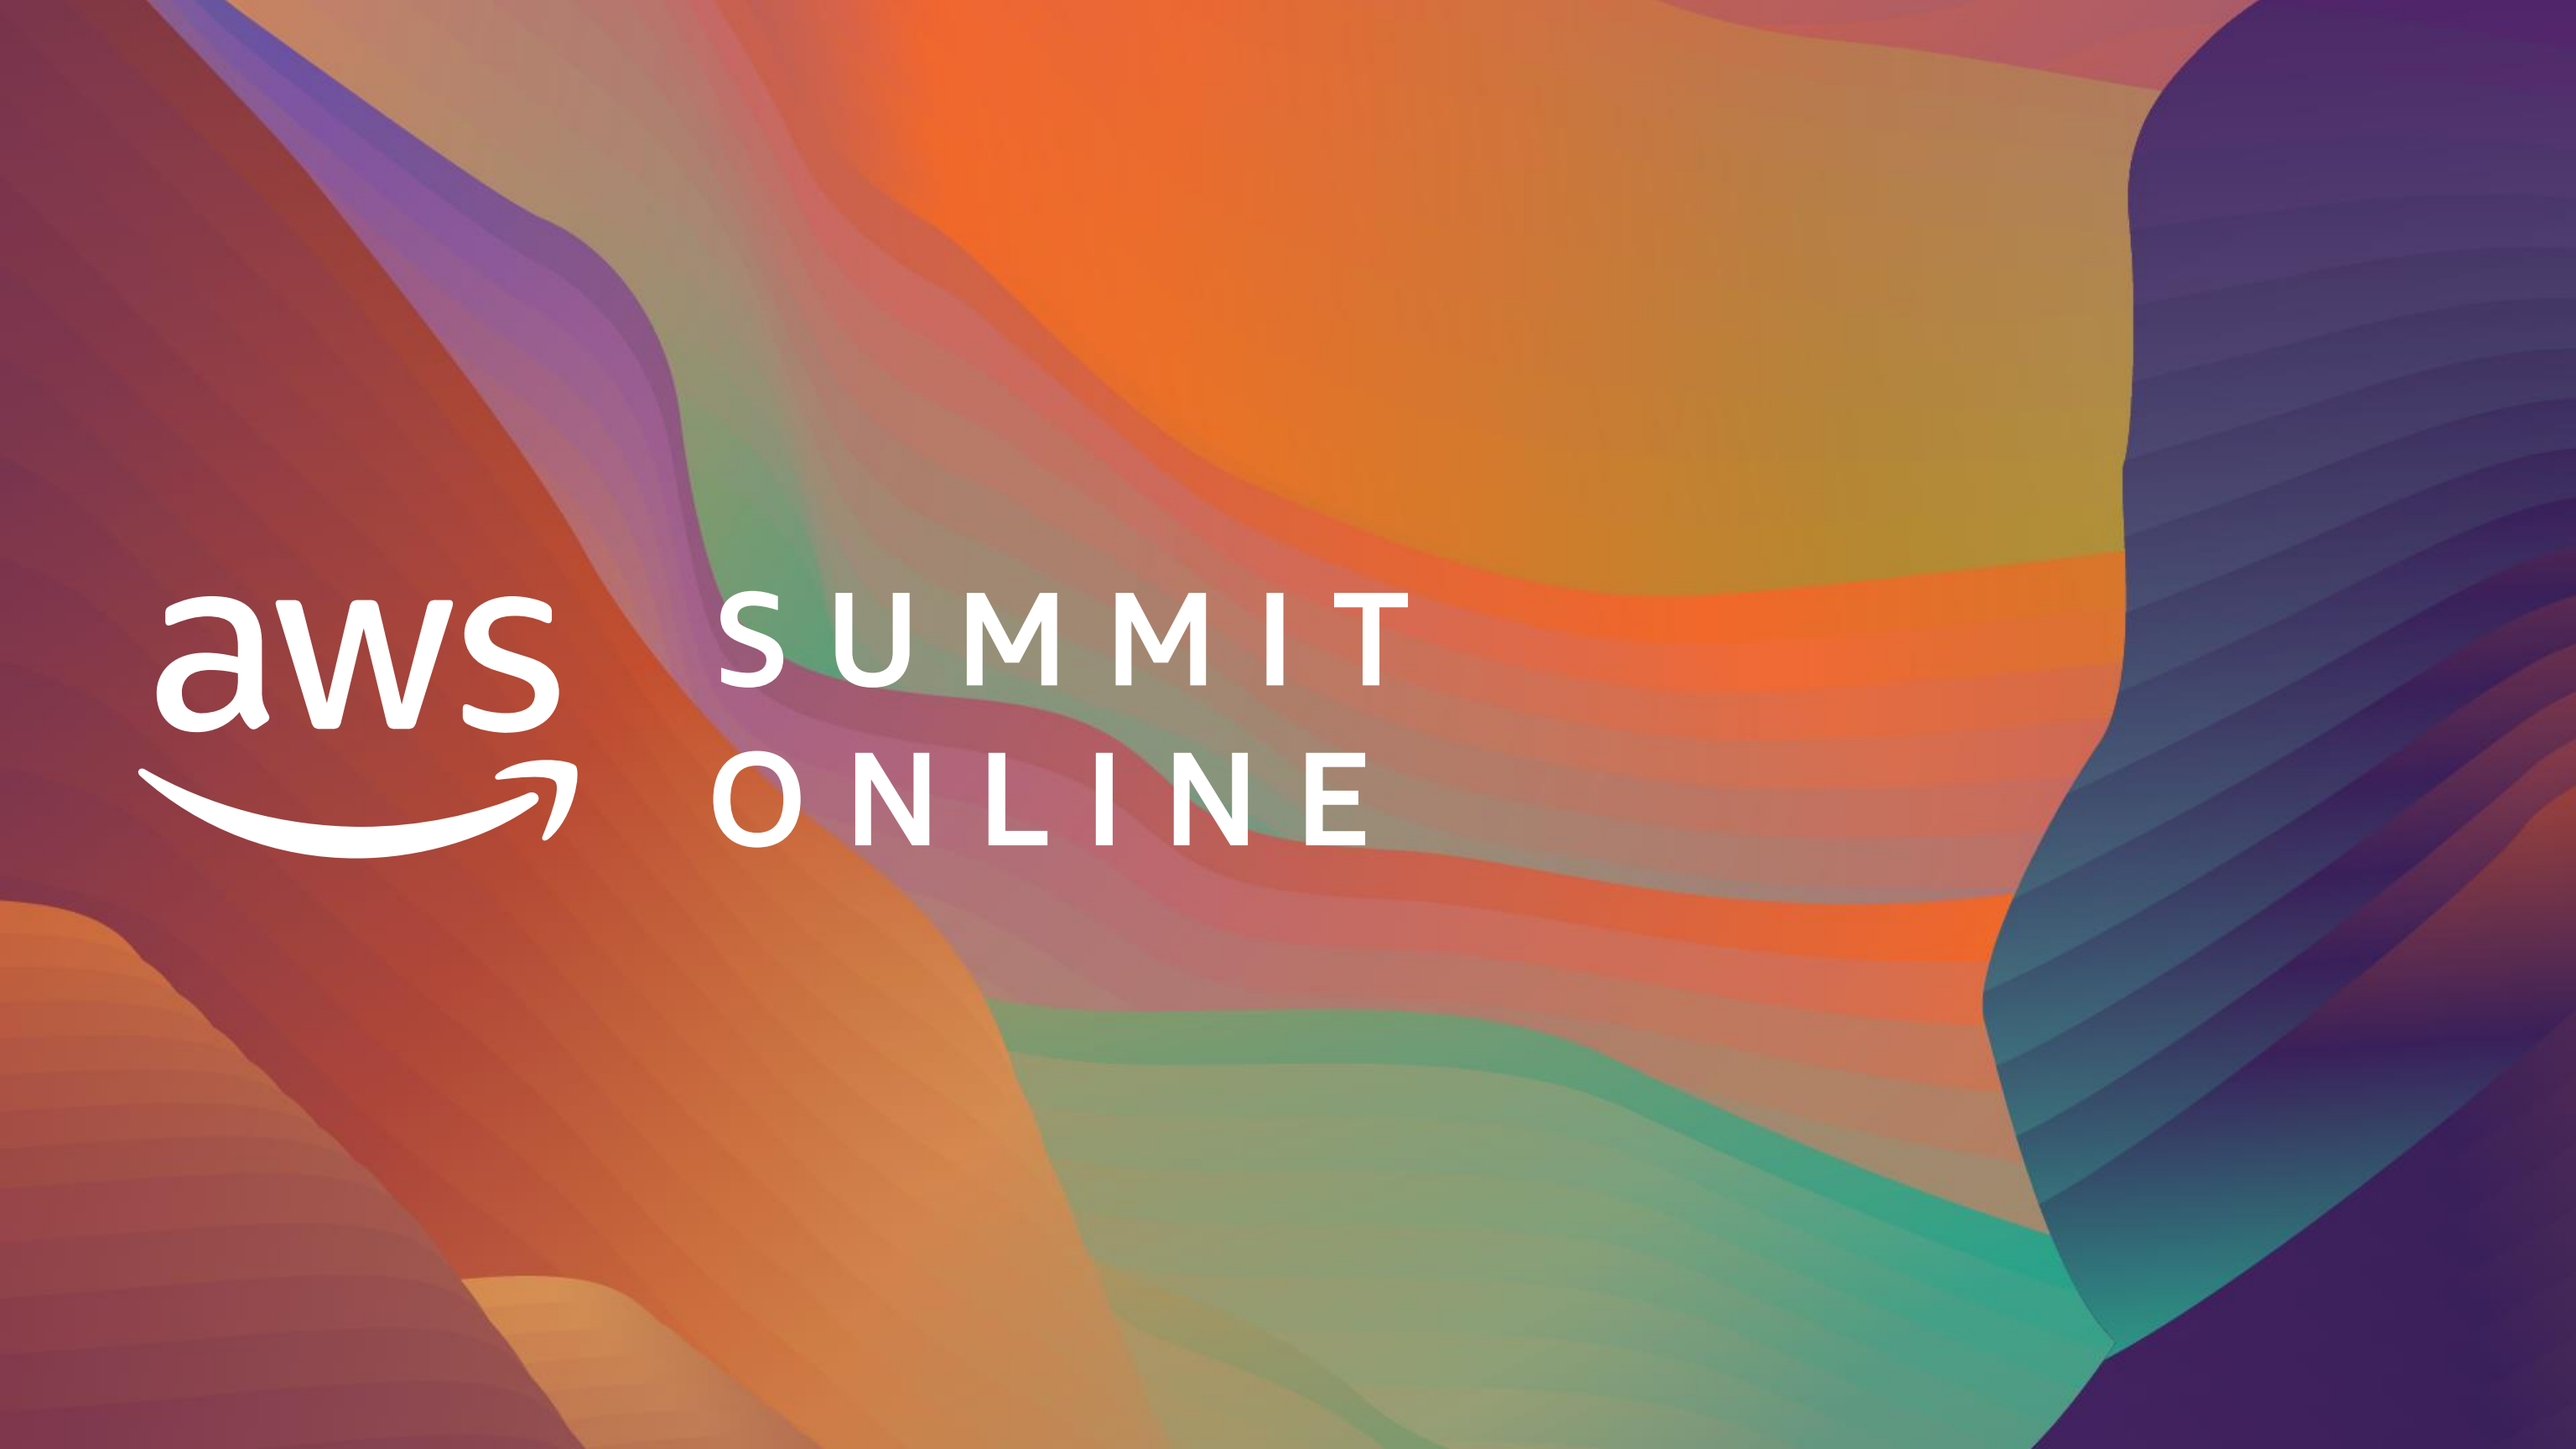
aws
SUMMIT
ONLINE

Advanced VPC connectivity patterns

Brett Looney

Global Solutions Architect
Amazon Web Services

aws SUMMIT ONLINE

Agenda

Connecting to AWS (AWS Direct Connect)

Routing within AWS (VPC peering and AWS Transit Gateway)

Sharing services in AWS (AWS PrivateLink)

DNS (Amazon Route 53 Resolver)

Connecting to AWS

aws SUMMIT ONLINE

AWS Direct Connect

Private connections to VPCs

The original way: Private VIF

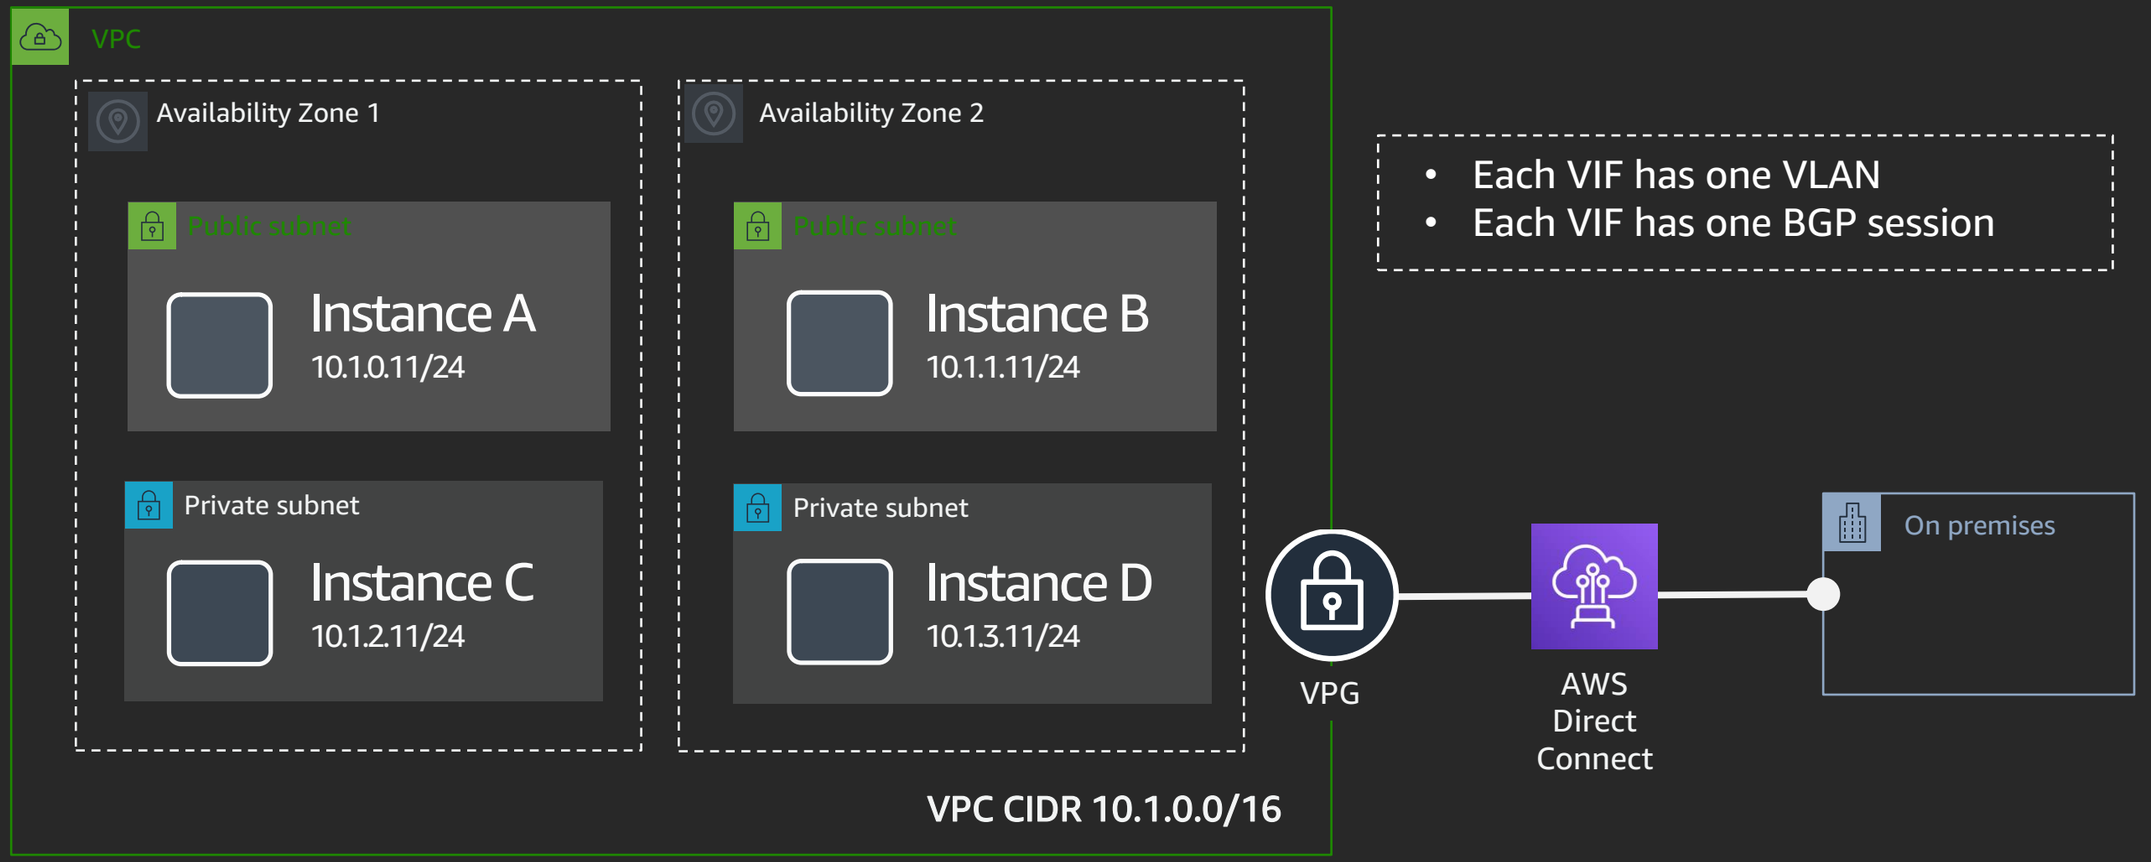

Side note: AWS Site-to-Site VPN

Managed private connections to VPCs

Primary or backup connection

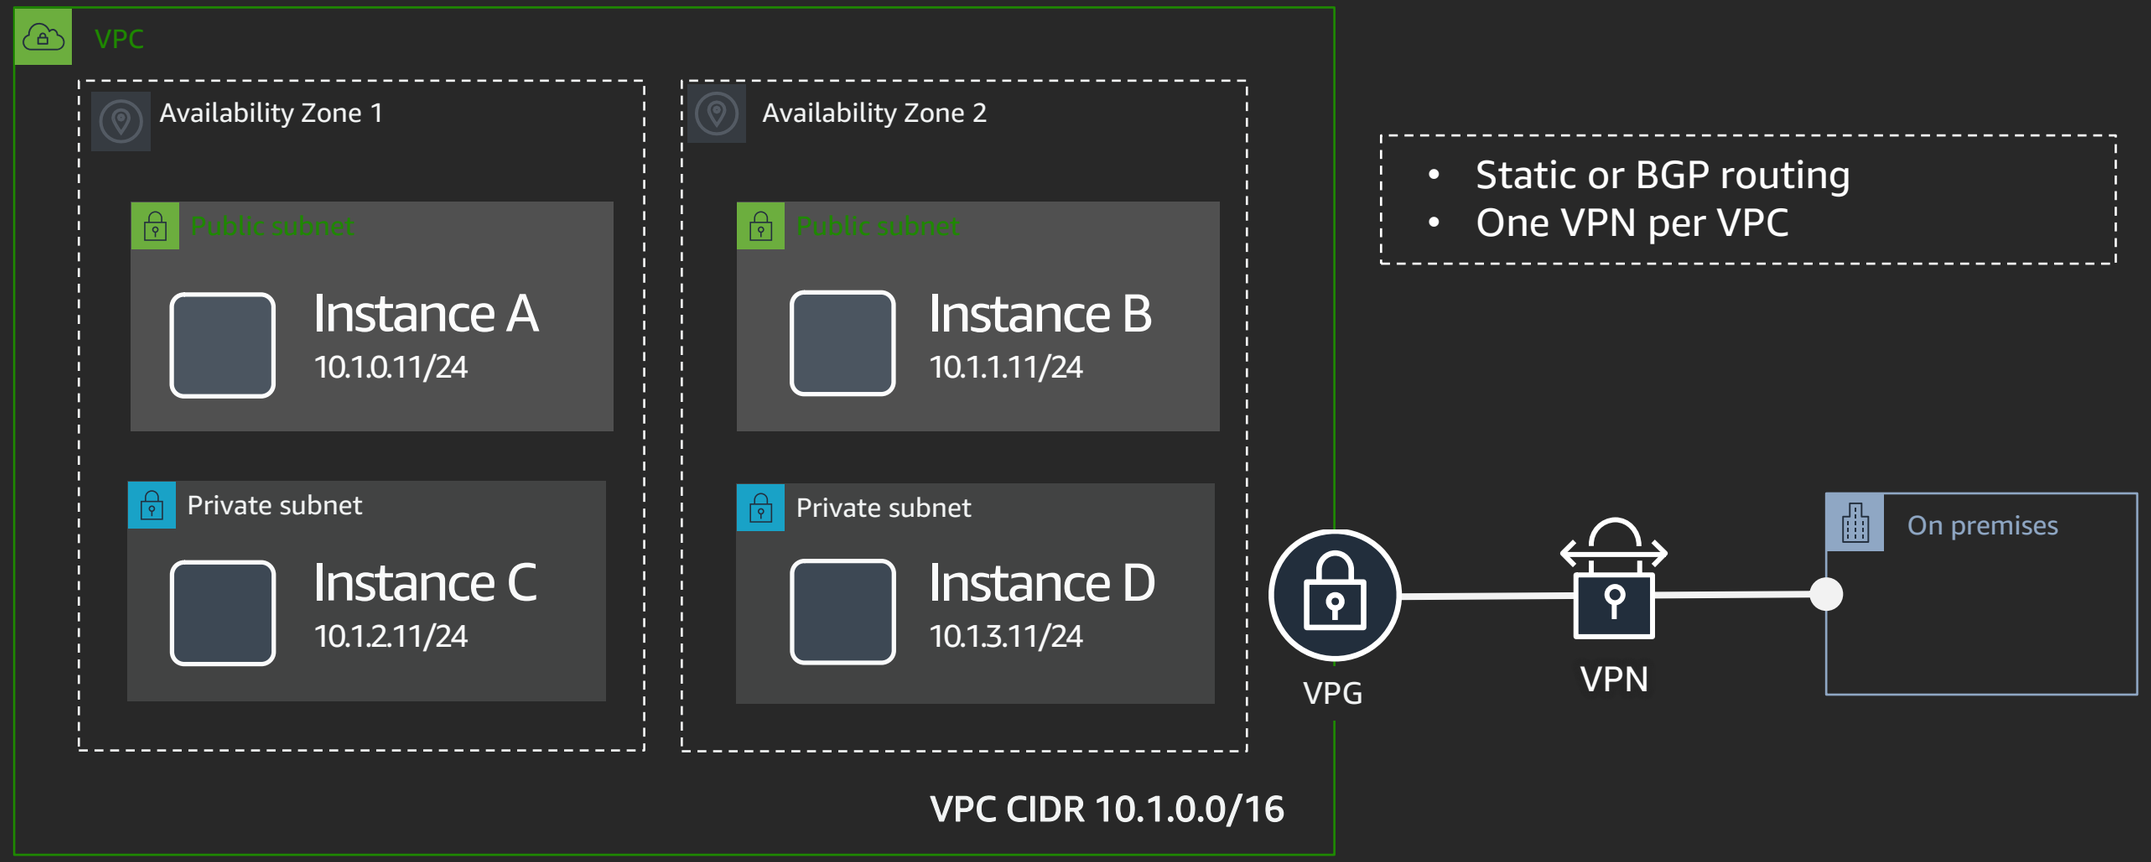

AWS Direct Connect

Private connections to VPCs

The newer way: AWS Direct Connect gateways

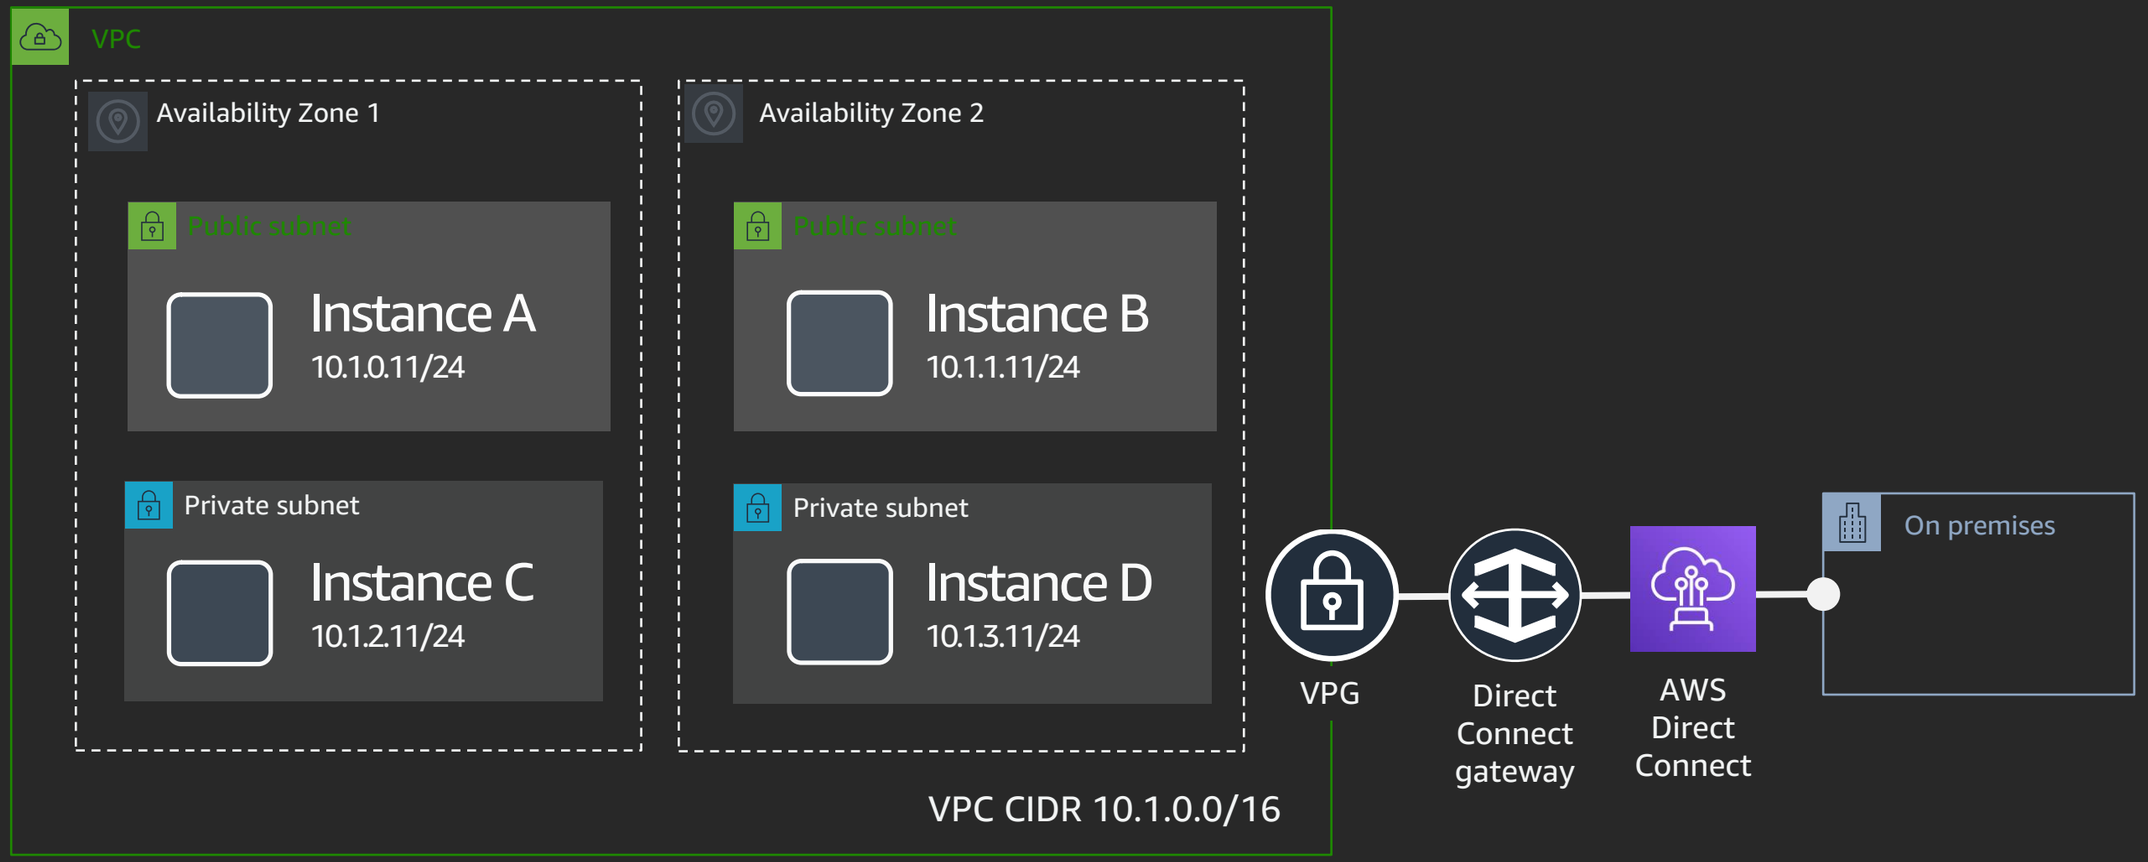

AWS Direct Connect with Direct Connect gateway

Private connections to VPCs

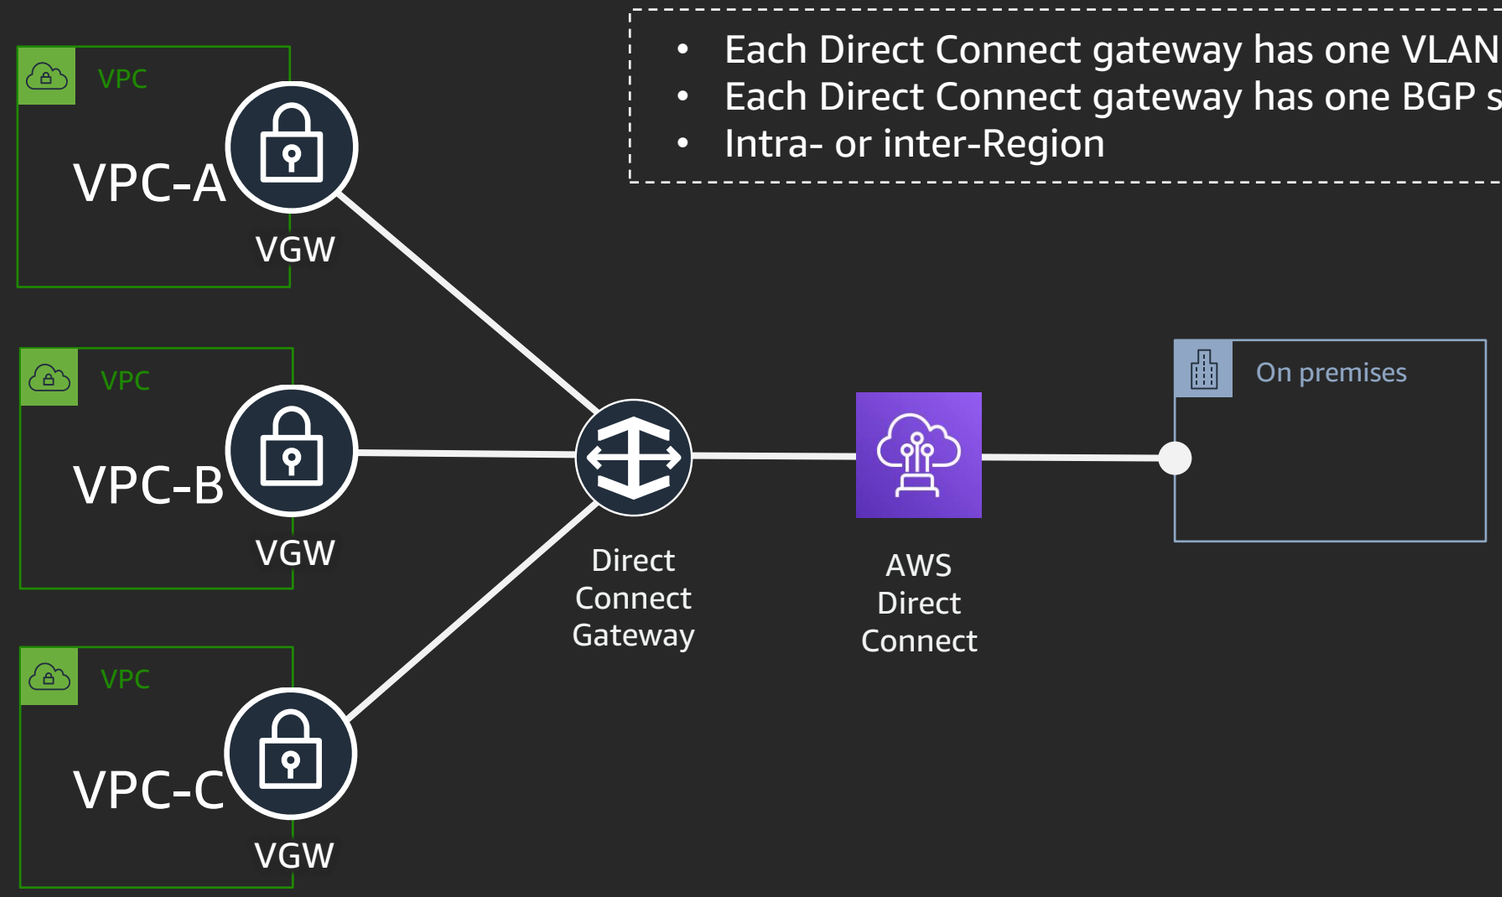

Side note: Customer premises equipment

Or: What do you need to do?

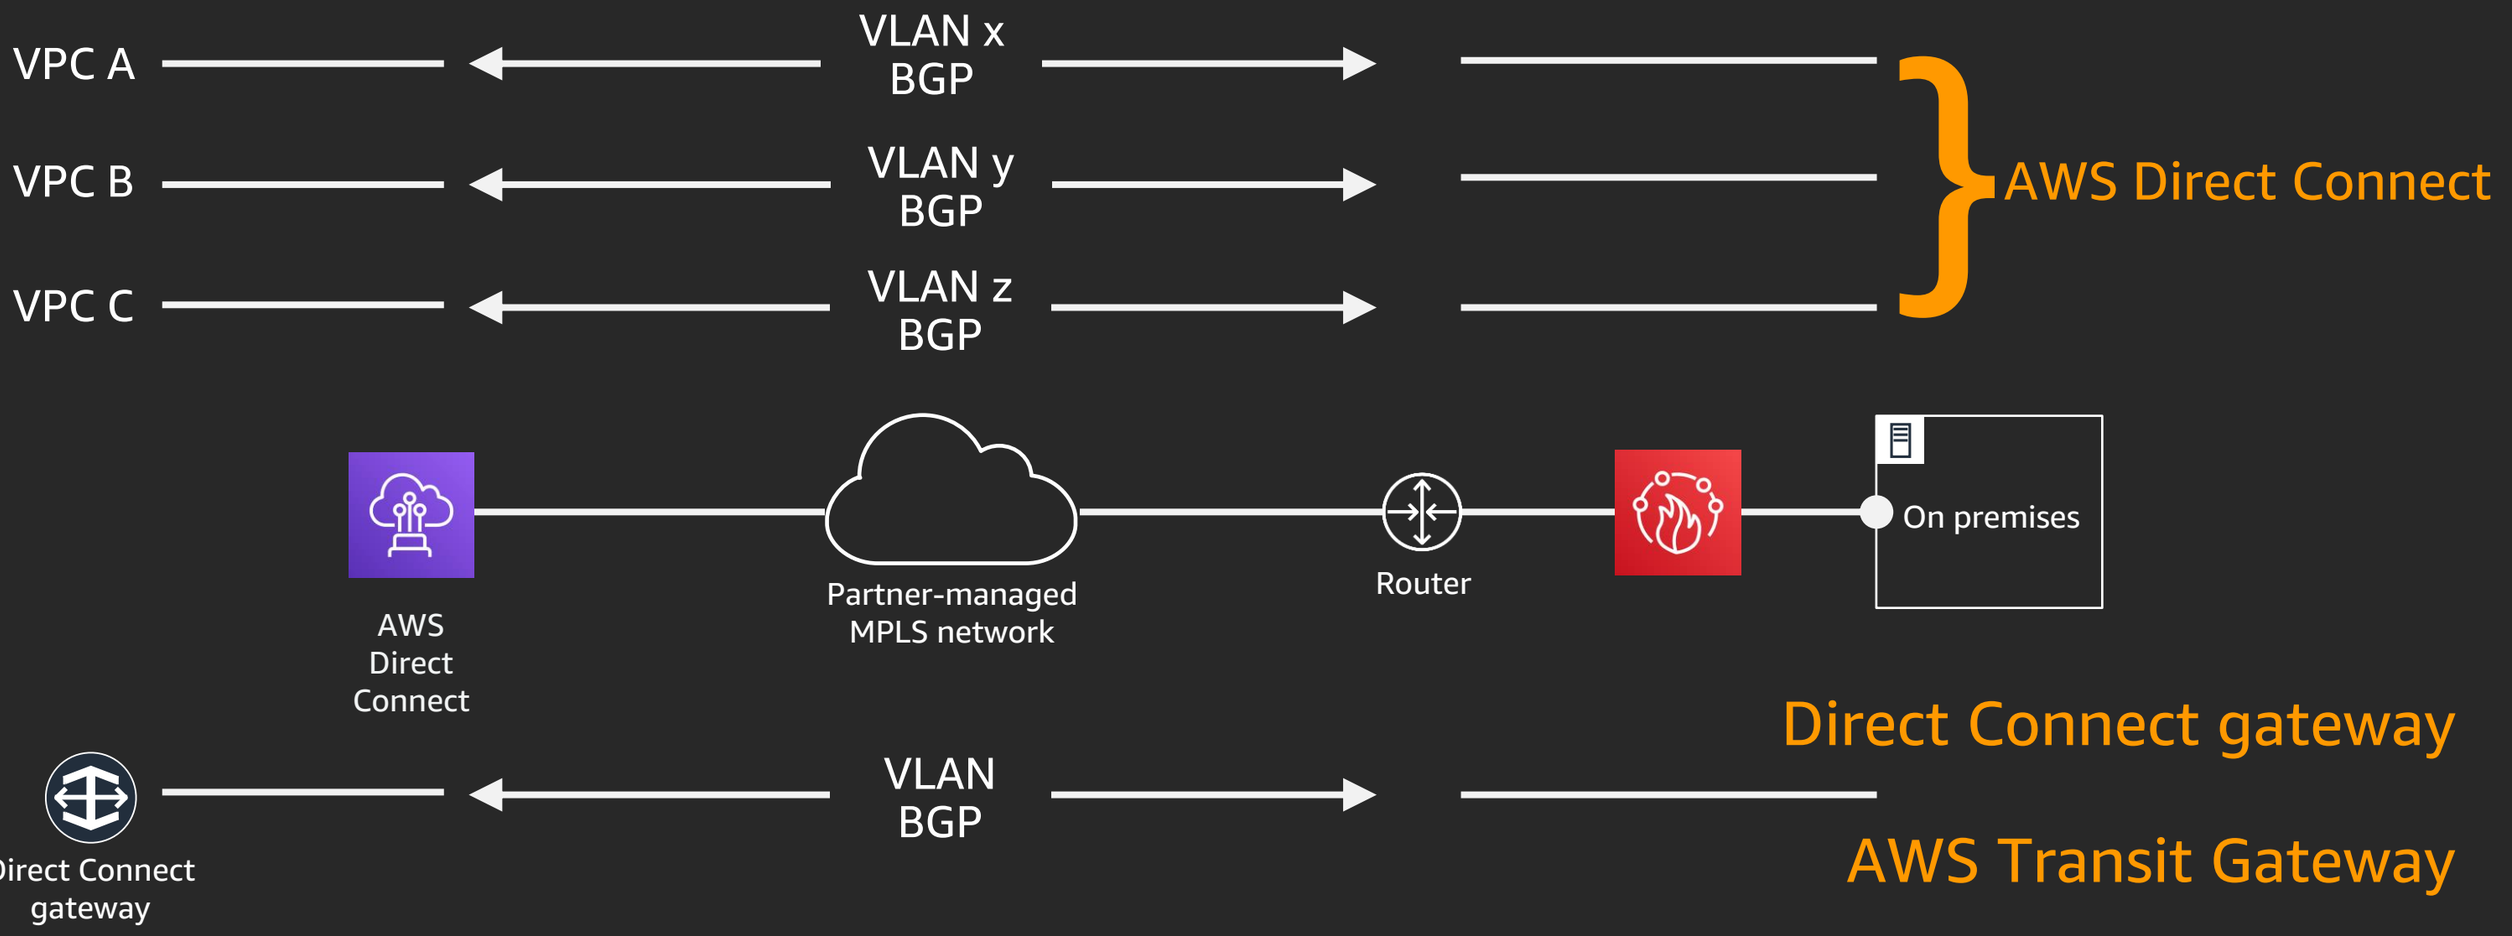

AWS Direct Connect

Public connections

Why have a public connection to AWS over a private network?

- Reduce congestion on existing internet link
- Deliver guaranteed bandwidth
- Fewer latency issues (jitter)
- Encrypt traffic to VPC over AWS Direct Connect using VPN

Why not have a public connection to AWS over a private network?

- Requires additional routed connections (VLANs and BGP)
- Requires customer-side firewall
- Not just AWS services, all AWS customer IPs are advertised
- May require customer-side NAT

AWS Direct Connect

Public connections

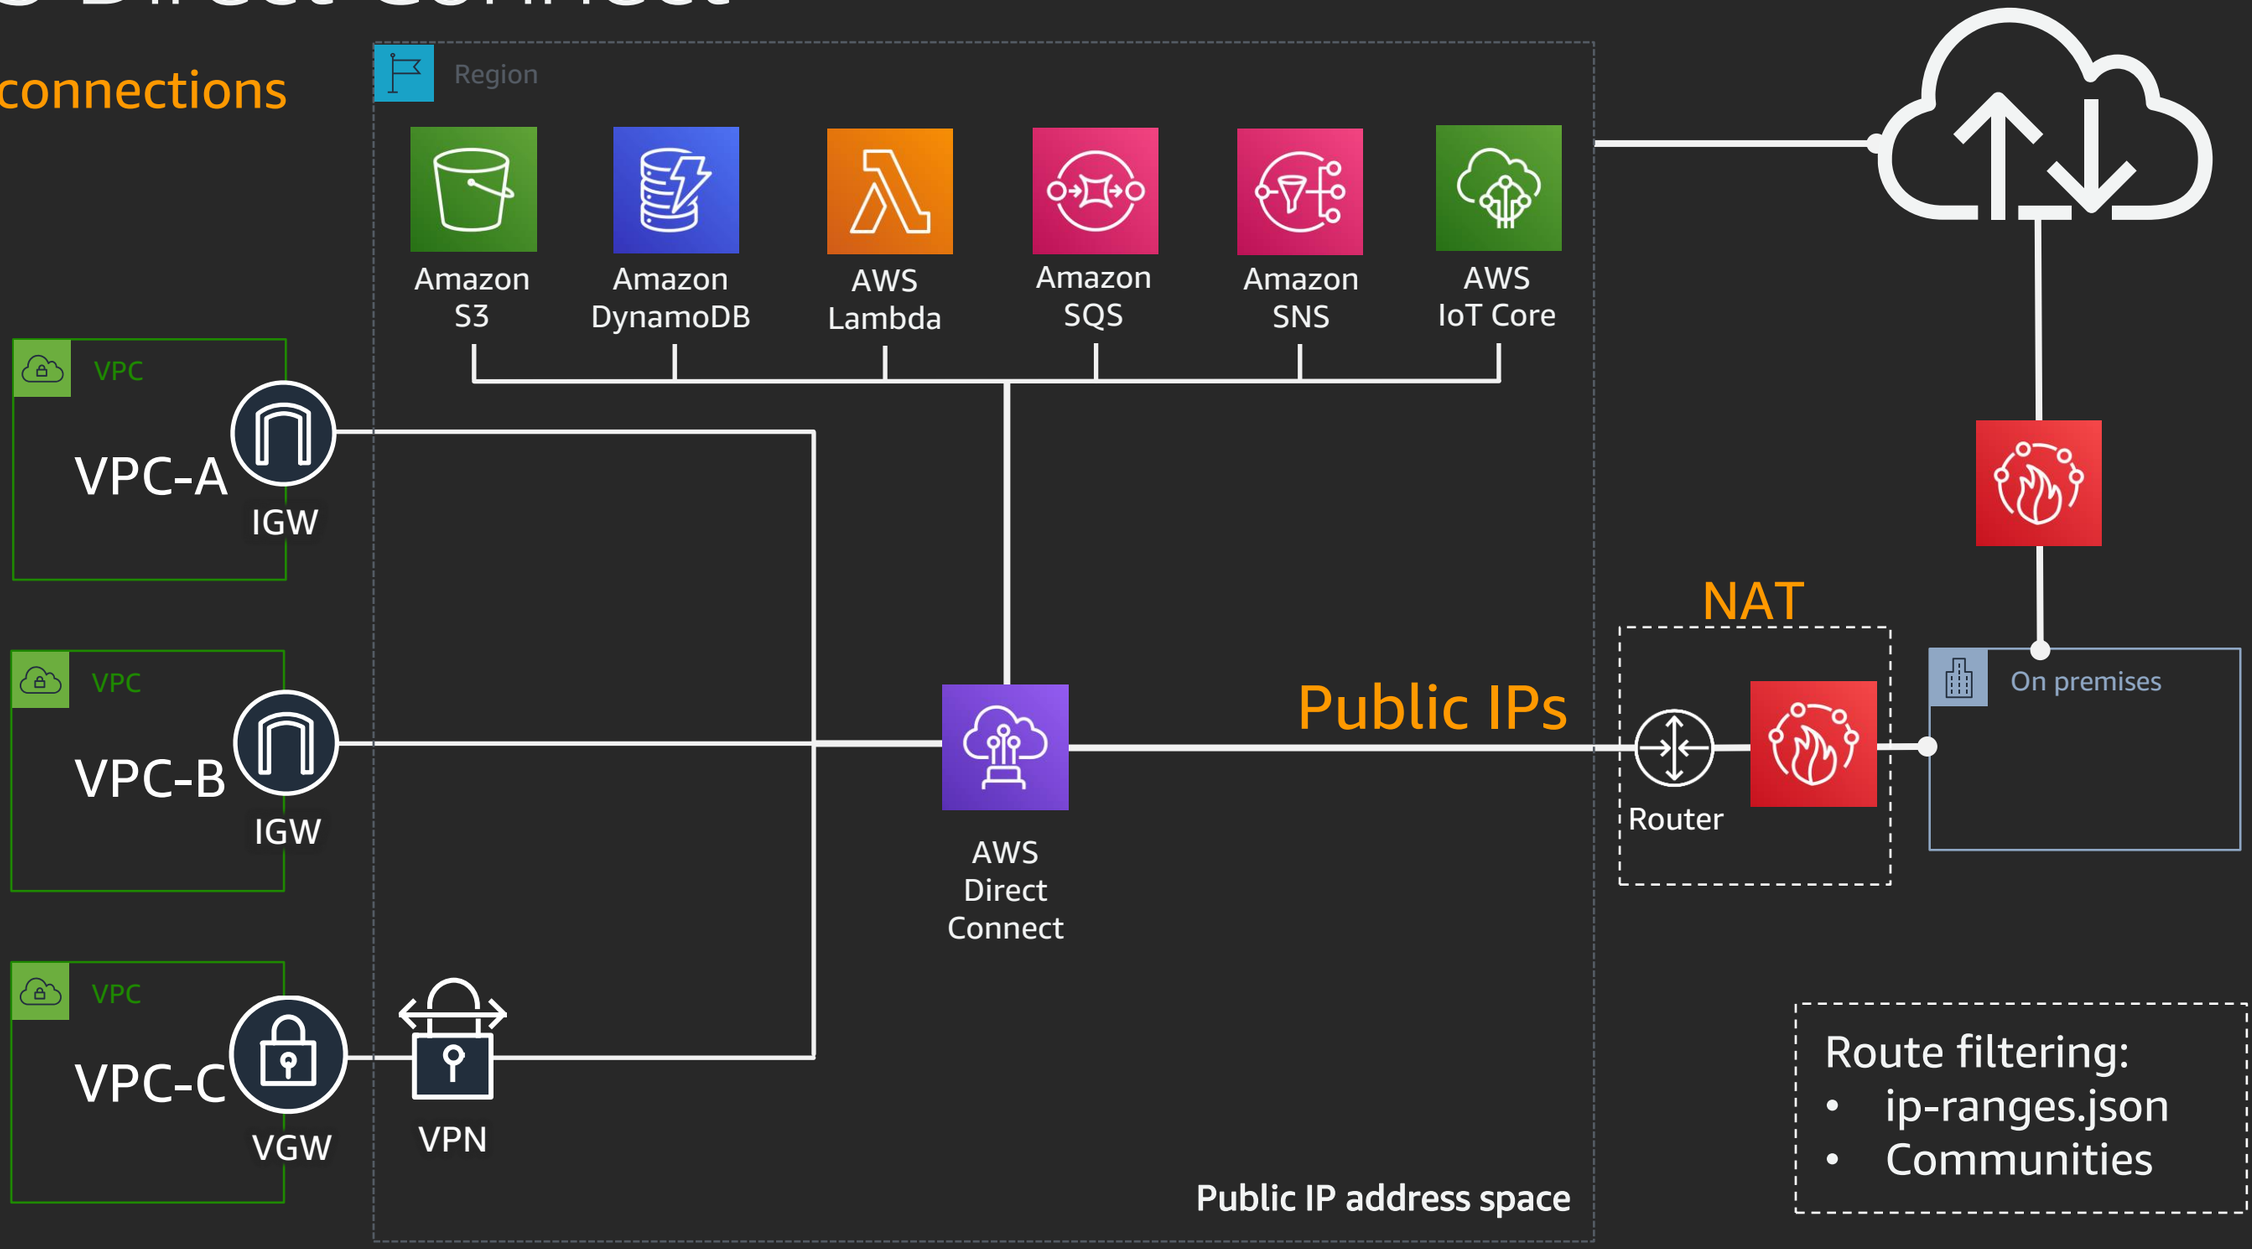

Route filtering:

- ip-ranges.json
- Communities

Routing within AWS

aws SUMMIT ONLINE

VPC peering

Or: Start simply, and then …

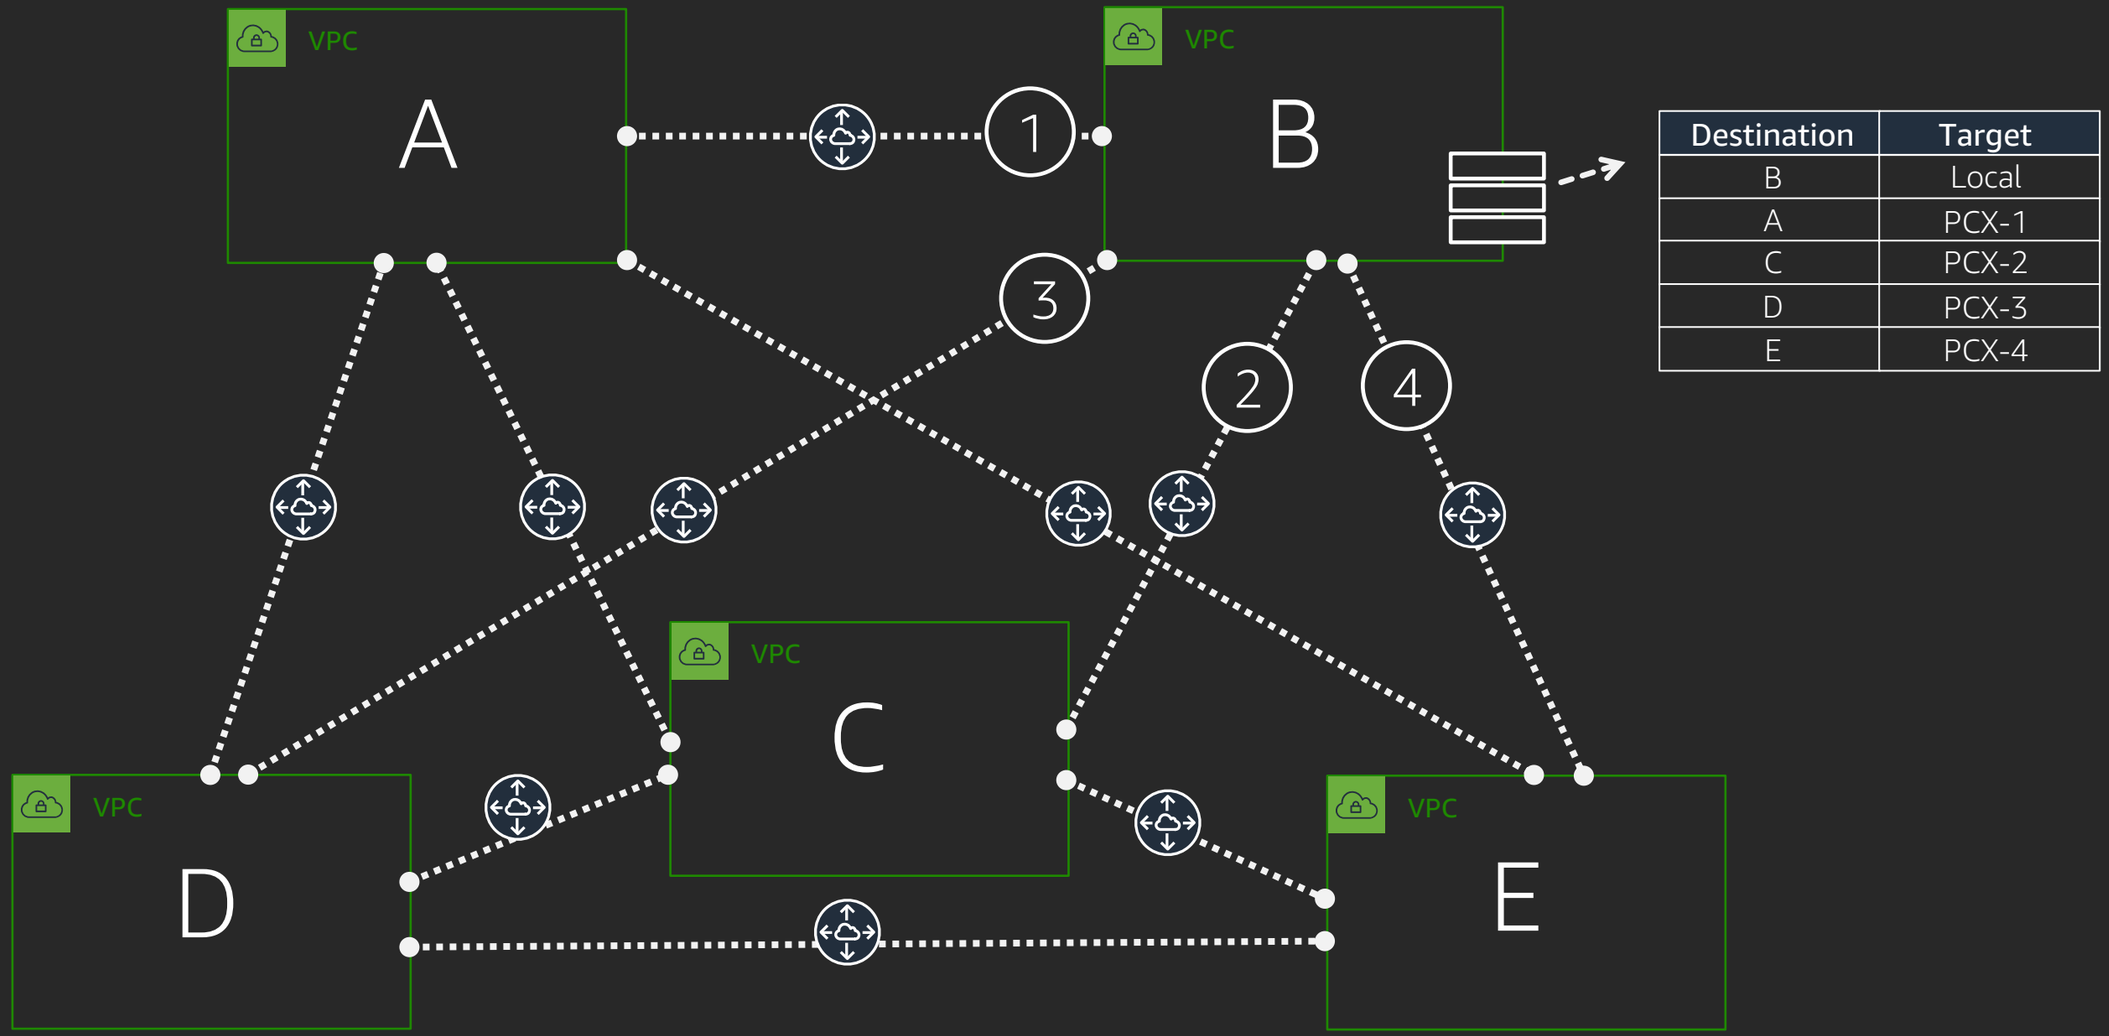

Destination	Target
B	Local
A	PCX-1
C	PCX-2
D	PCX-3
E	PCX-4

VPC peering

Q: How many peering connections do I need for a full mesh?

$$\frac{n(n-1)}{2}$$

VPC peering

Q: How many peering connections do I need for a full mesh?

$$\frac{10(10-1)}{2} = 45$$

VPC peering

Q: How many peering connections do I need for a full mesh?

$$\frac{100(100-1)}{2} = 4,500$$

VPC peering

Why is this a problem?

Static routes per Amazon VPC route table

100

Amazon VPC peering connections per Amazon VPC

125

AWS Transit Gateway

Cloud-scale routing

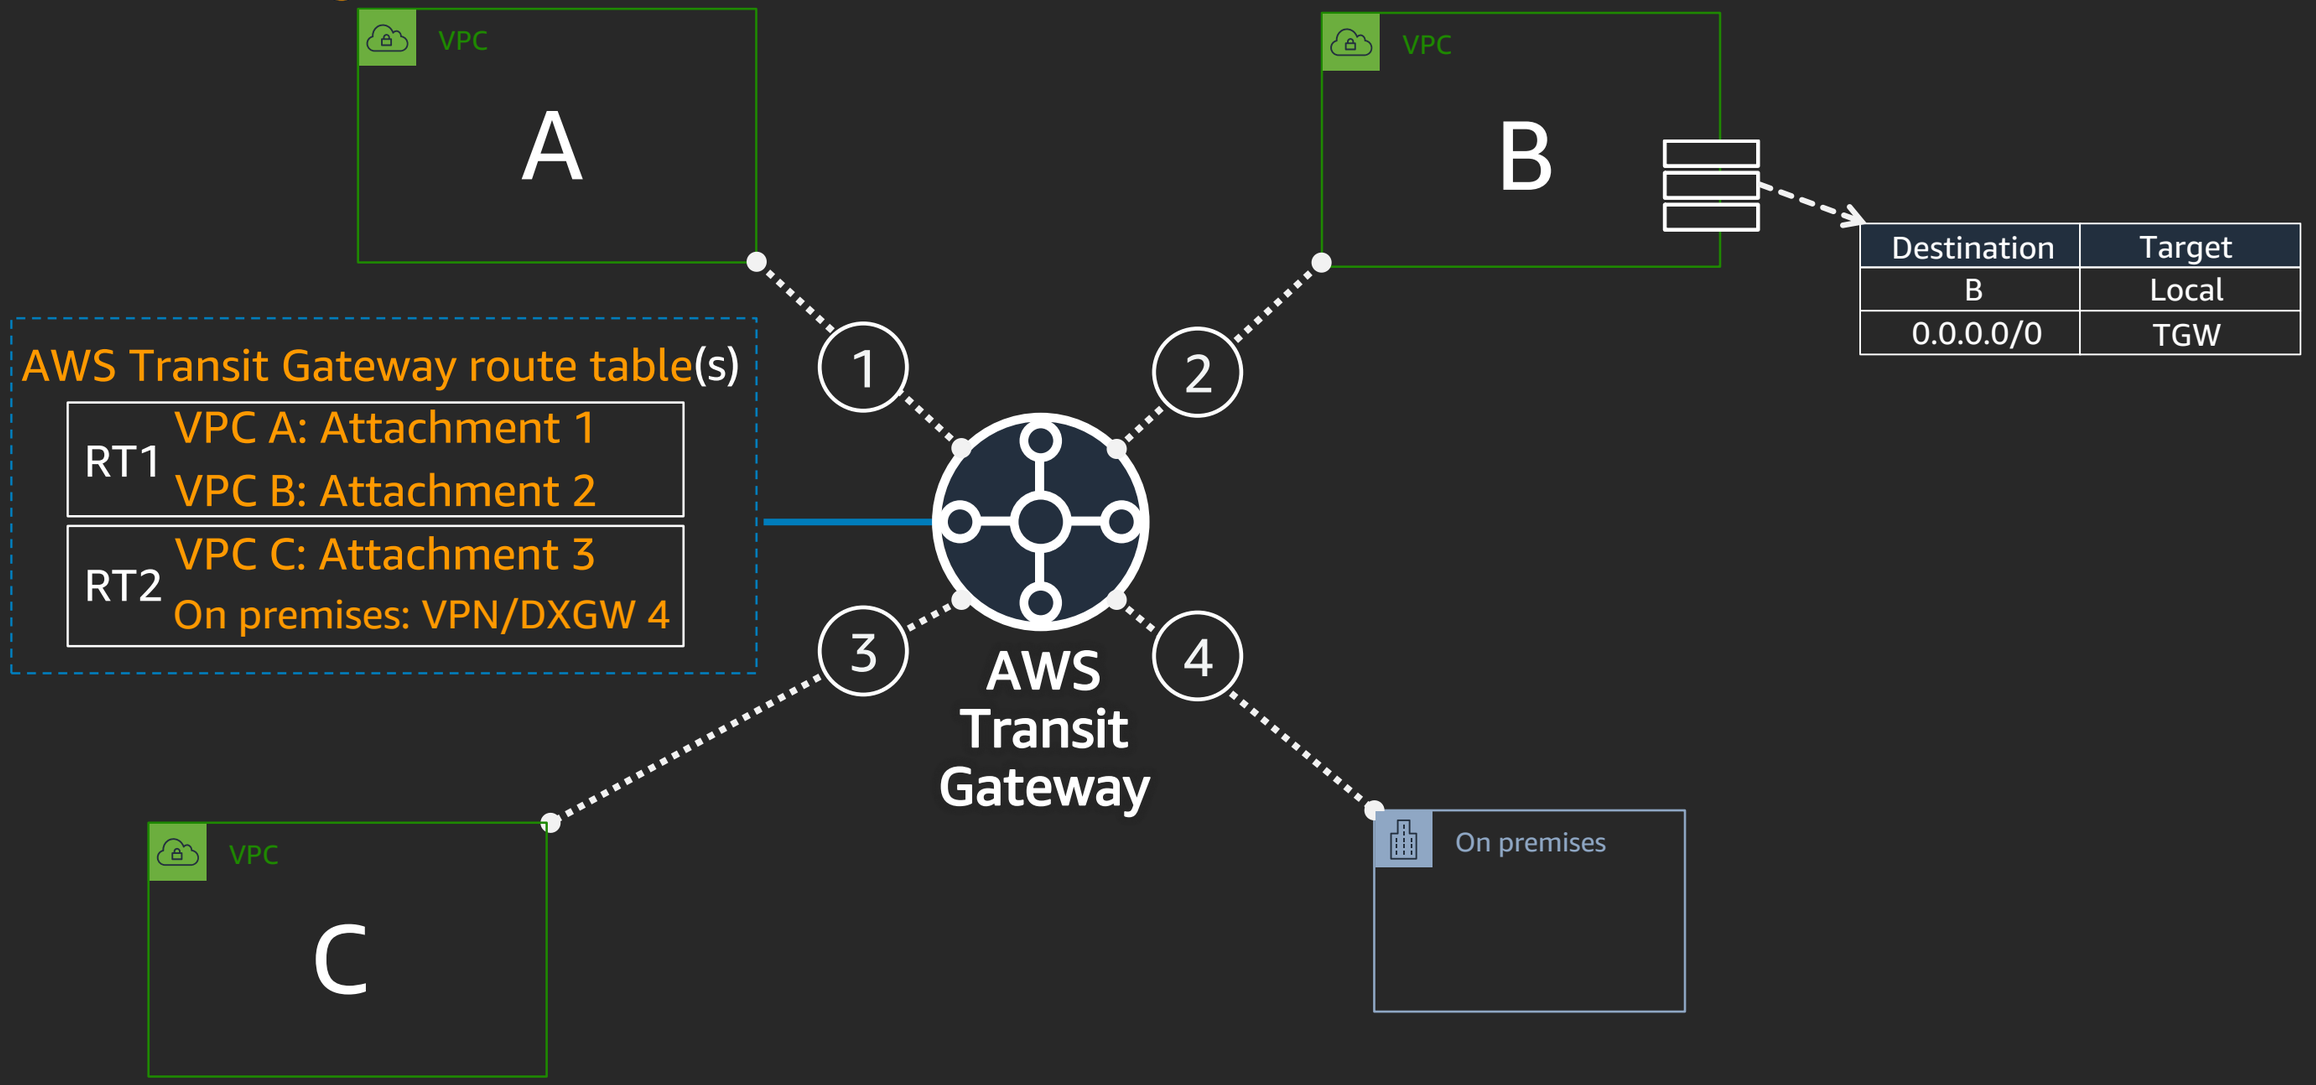

AWS Transit Gateway terms

Attachment

The connection from an Amazon VPC, VPN, or AWS Direct Connect to AWS Transit Gateway

Association

The route table used to route packets coming from an attachment (from an Amazon VPC and VPN)

Propagation

The route table where the attachment's routes are installed

AWS Transit Gateway example

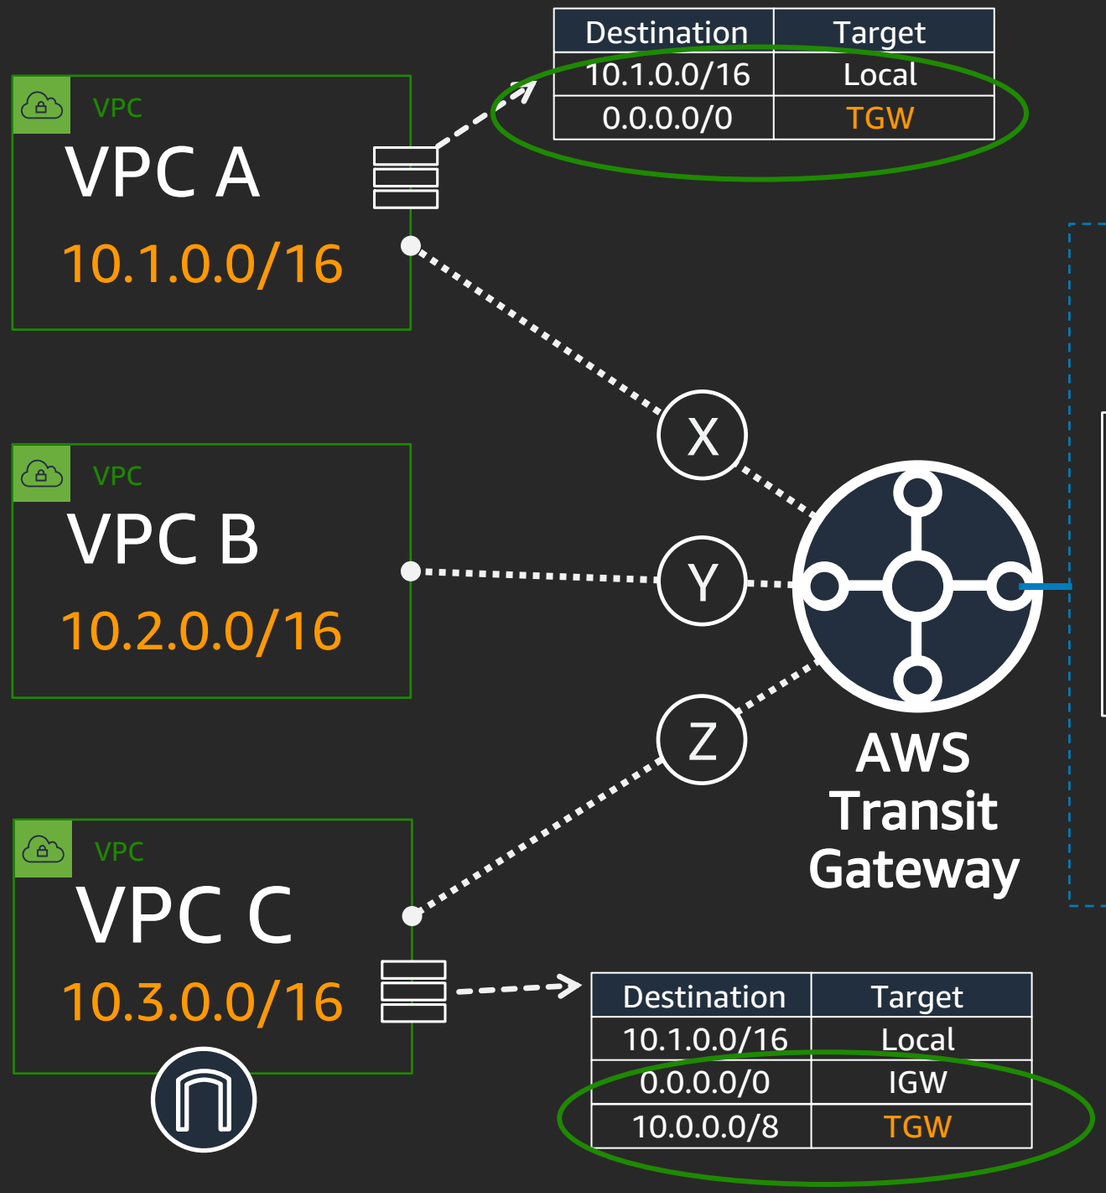

AWS Transit Gateway route table(s)

	Associations	Propagations	Routes
RT1	VPC A from X	VPC A from X	10.1.0.0/16 via X
	VPC B from Y	VPC B from Y	10.2.0.0/16 via Y
	VPC C from Z	VPC C from Z	10.3.0.0/16 via Z

With propagation turned off, you can still statically configure routes

AWS Transit Gateway

Other advantages

Can create multiple AWS Transit Gateway route tables

- Use these to create isolated sets of VPCs
- For example:
  - Production VPCs can use AWS Direct Connect and communicate with each other and on premises
  - Dev and test VPCs can use AWS Direct Connect and communicate with each other and on premises
  - Production, dev, and test cannot communicate with each other

---

VPCs with overlapping CIDR blocks can use the same AWS Transit Gateway

- But they must communicate through NAT
- AWS Transit Gateway does not fix fundamental routing problems

Sharing services in AWS

aws SUMMIT ONLINE

AWS PrivateLink

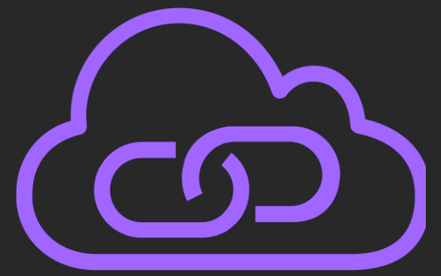

What is it?

Reach resources in another VPC, AWS services, and on premises

- Eliminate the exposure of data to the public internet
- Without peering or routing
- Resources appear as a local IP

Bonus

- IP ranges in each VPC can overlap
- AWS PrivateLink performs double-sided NAT for you

AWS PrivateLink quick overview

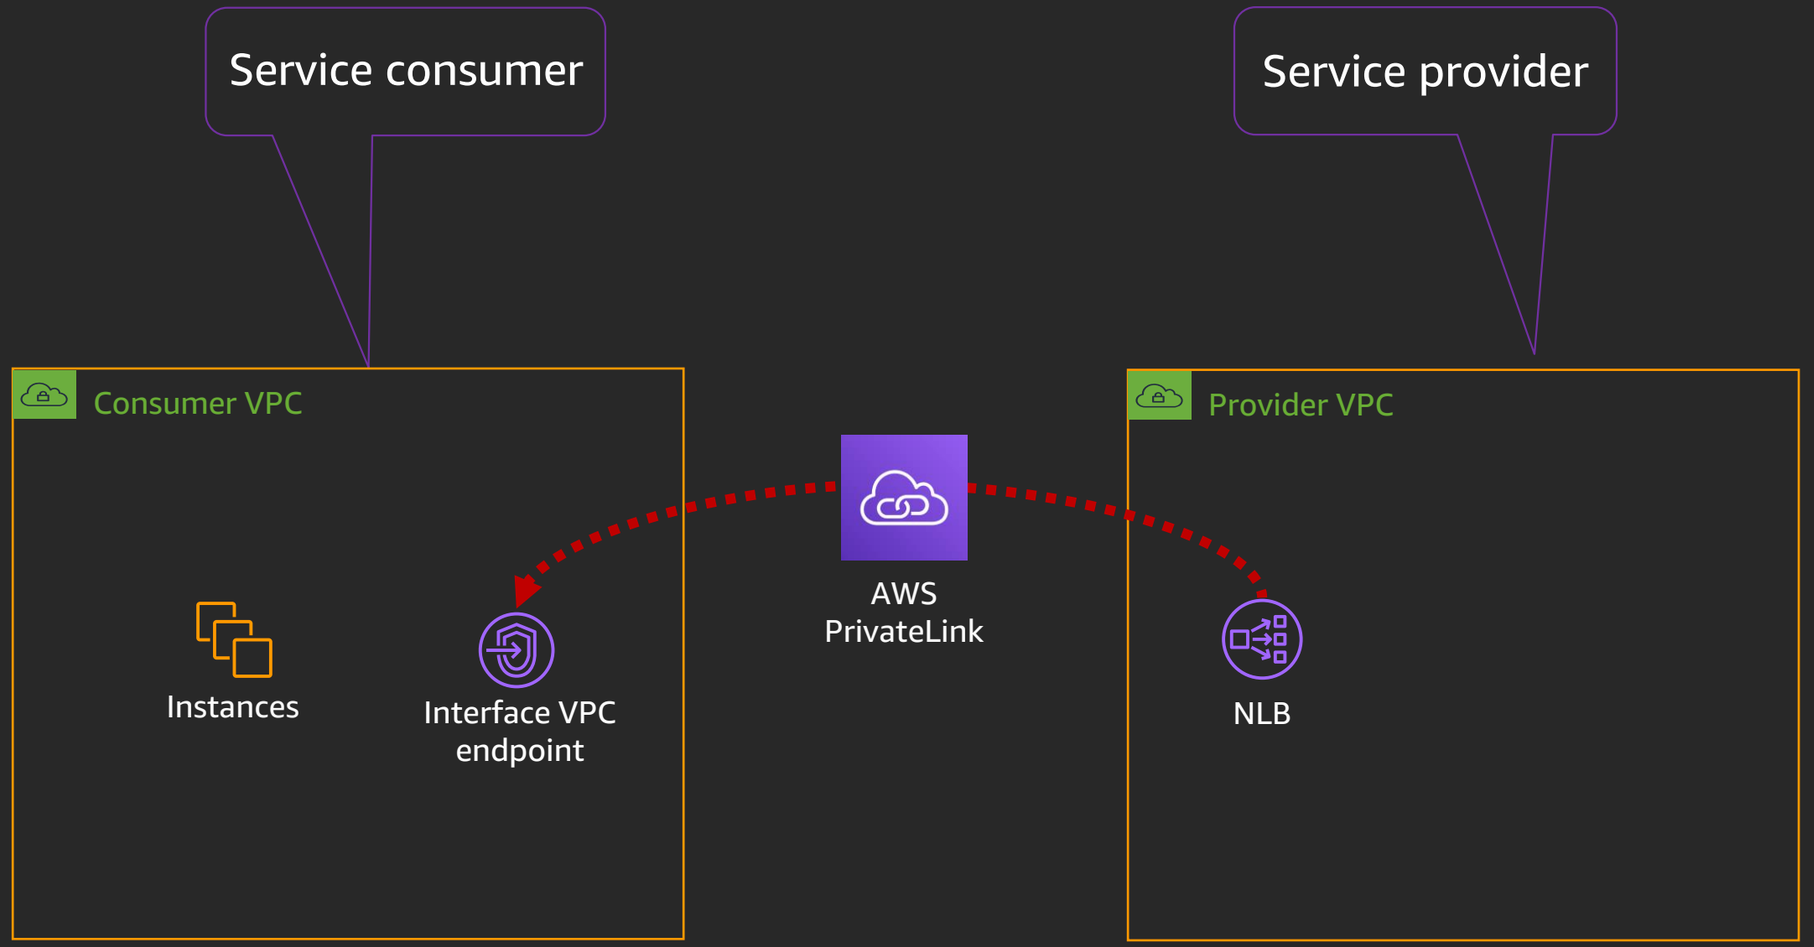

AWS PrivateLink interface endpoints – AWS services

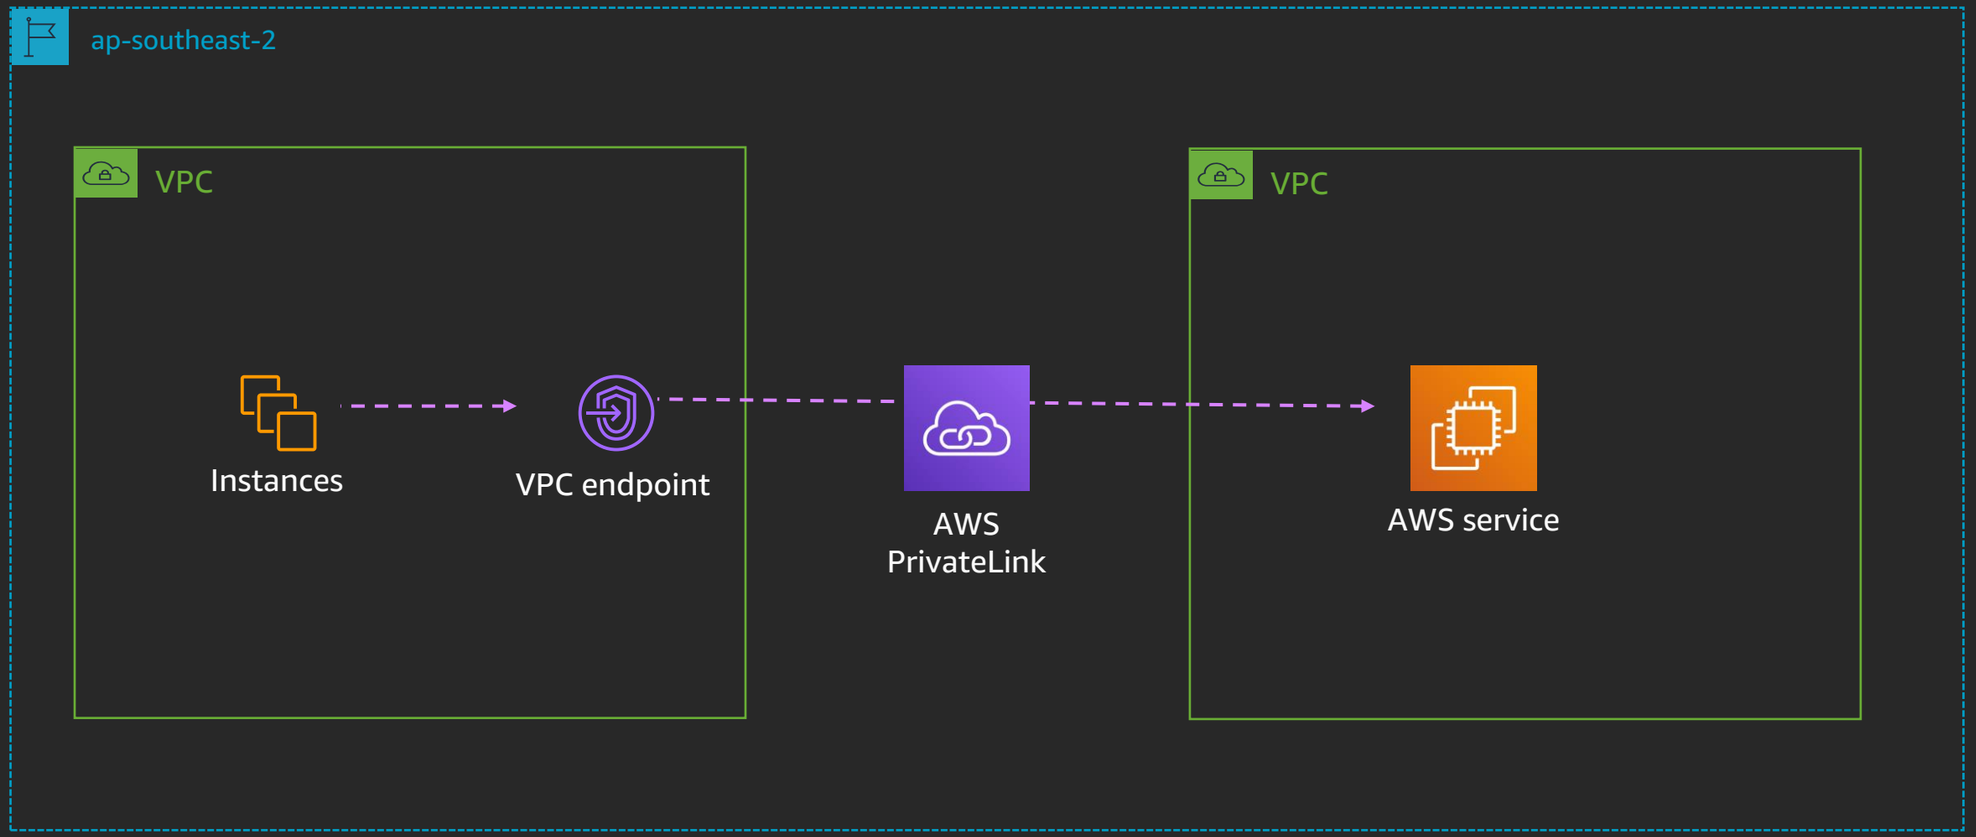

AWS PrivateLink interface endpoints

Endpoint services and SaaS

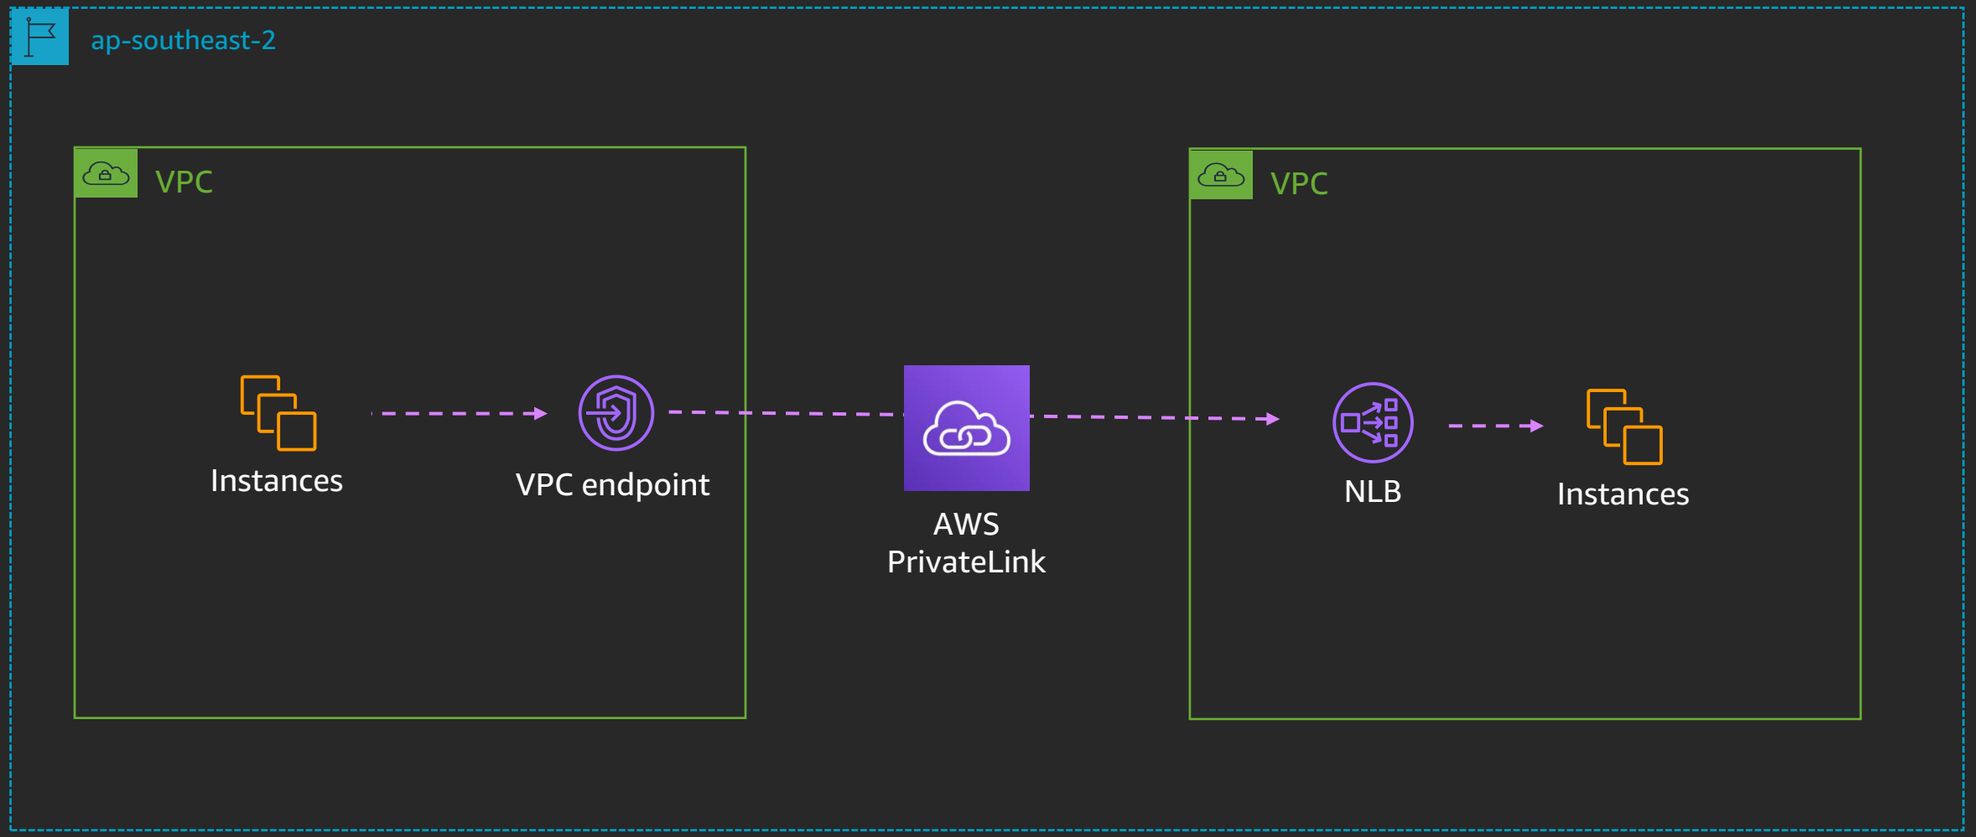

VPC endpoints and ENIs

- A VPC endpoint is a collection of ENIs spanning subnets
- Within a subnet, a VPCE is represented as an ENI
  - At most one elastic network interface per AZ
  - An ENI is used to connect to a AWS PrivateLink enabled service

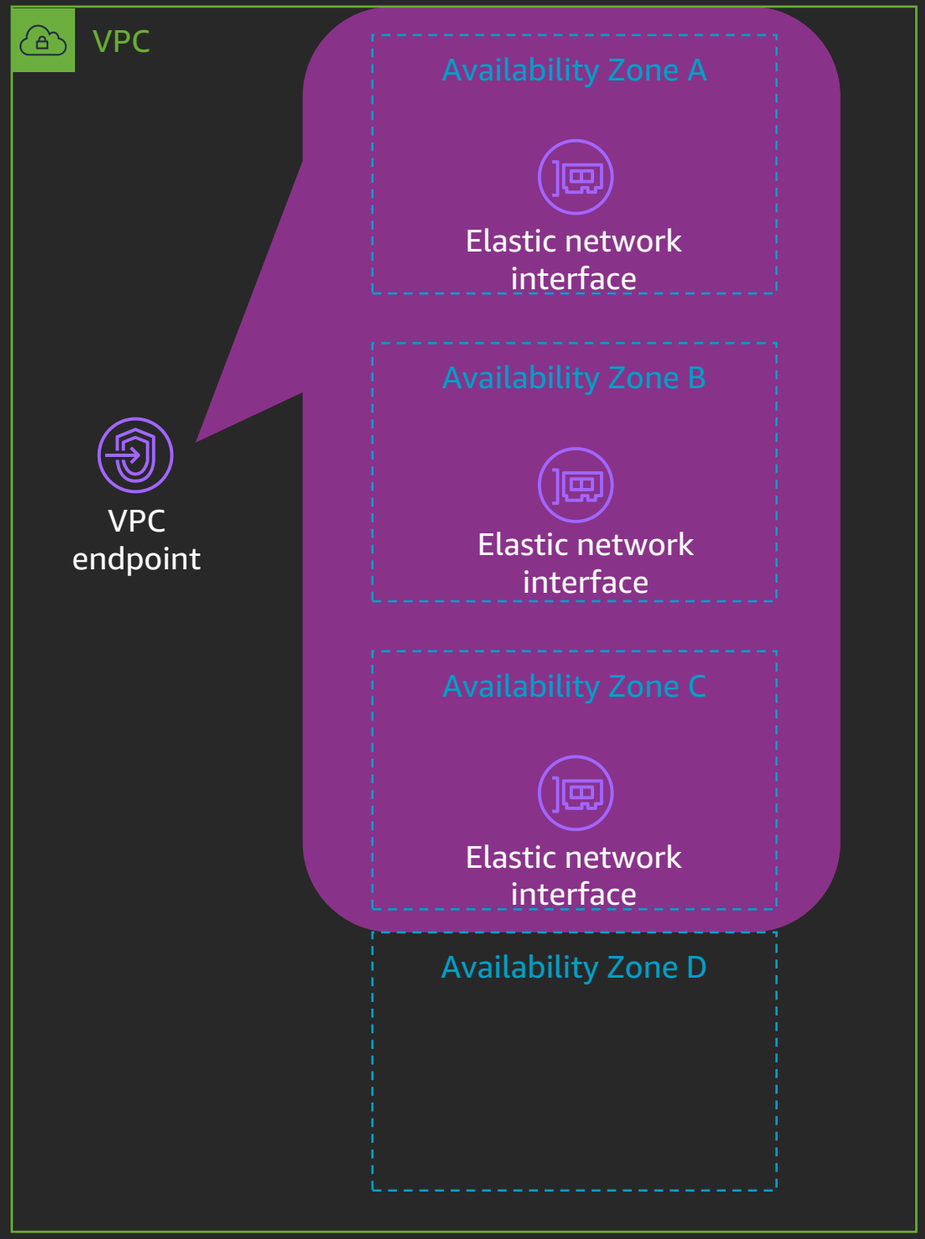

On-premises service consumers

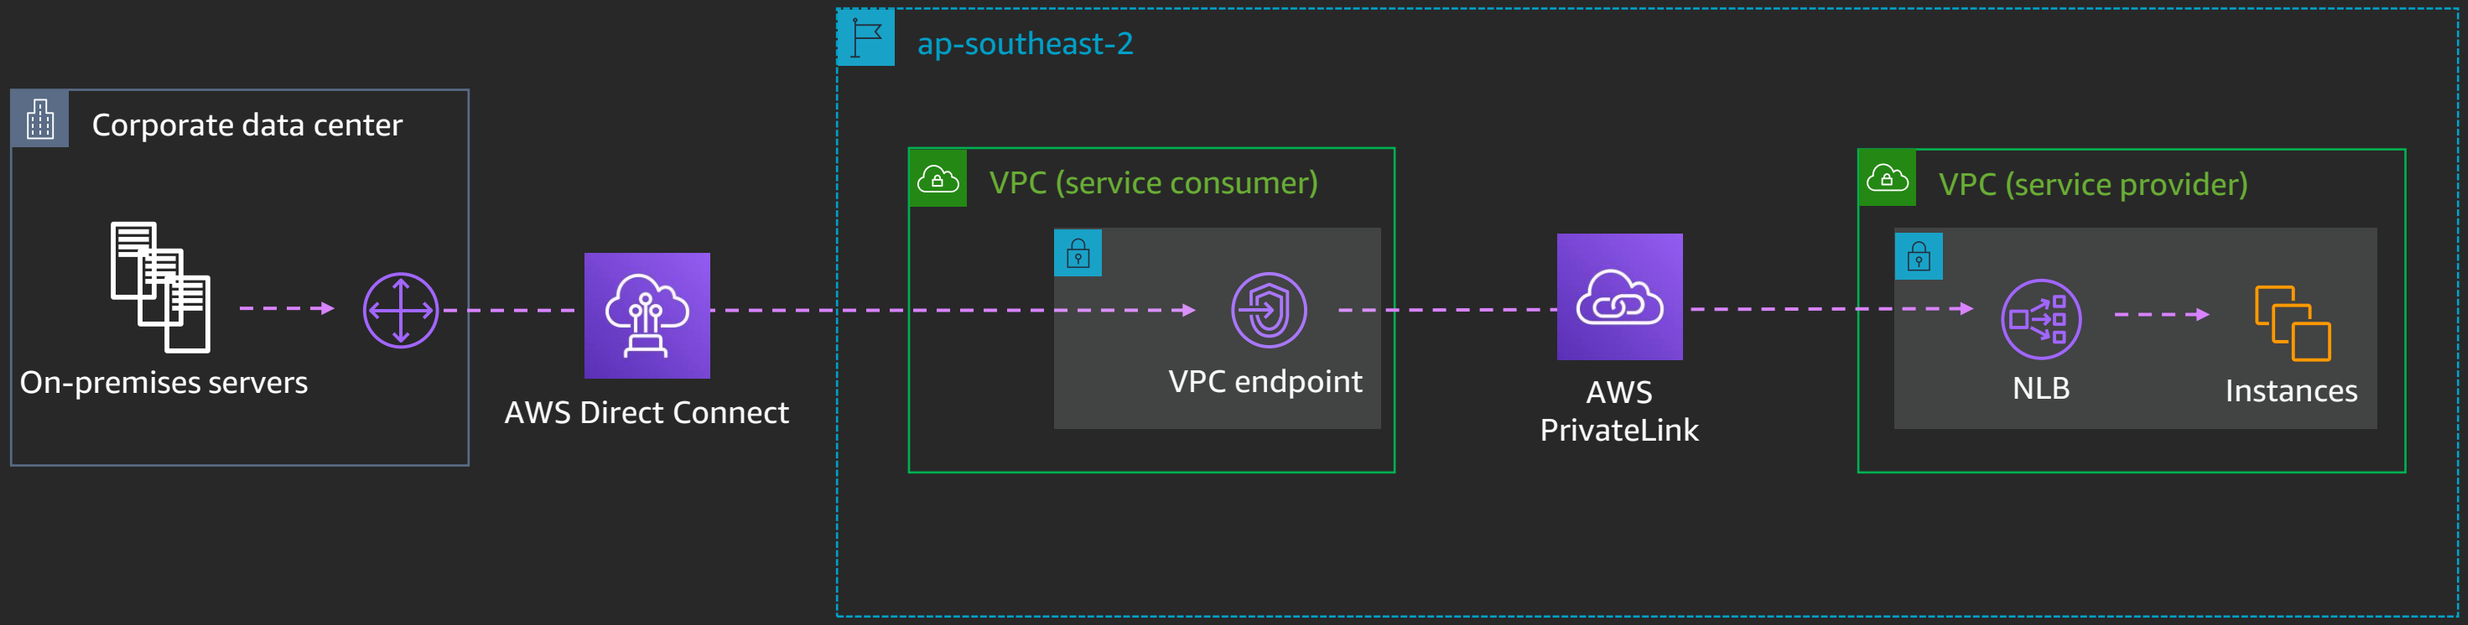

On-premises service providers

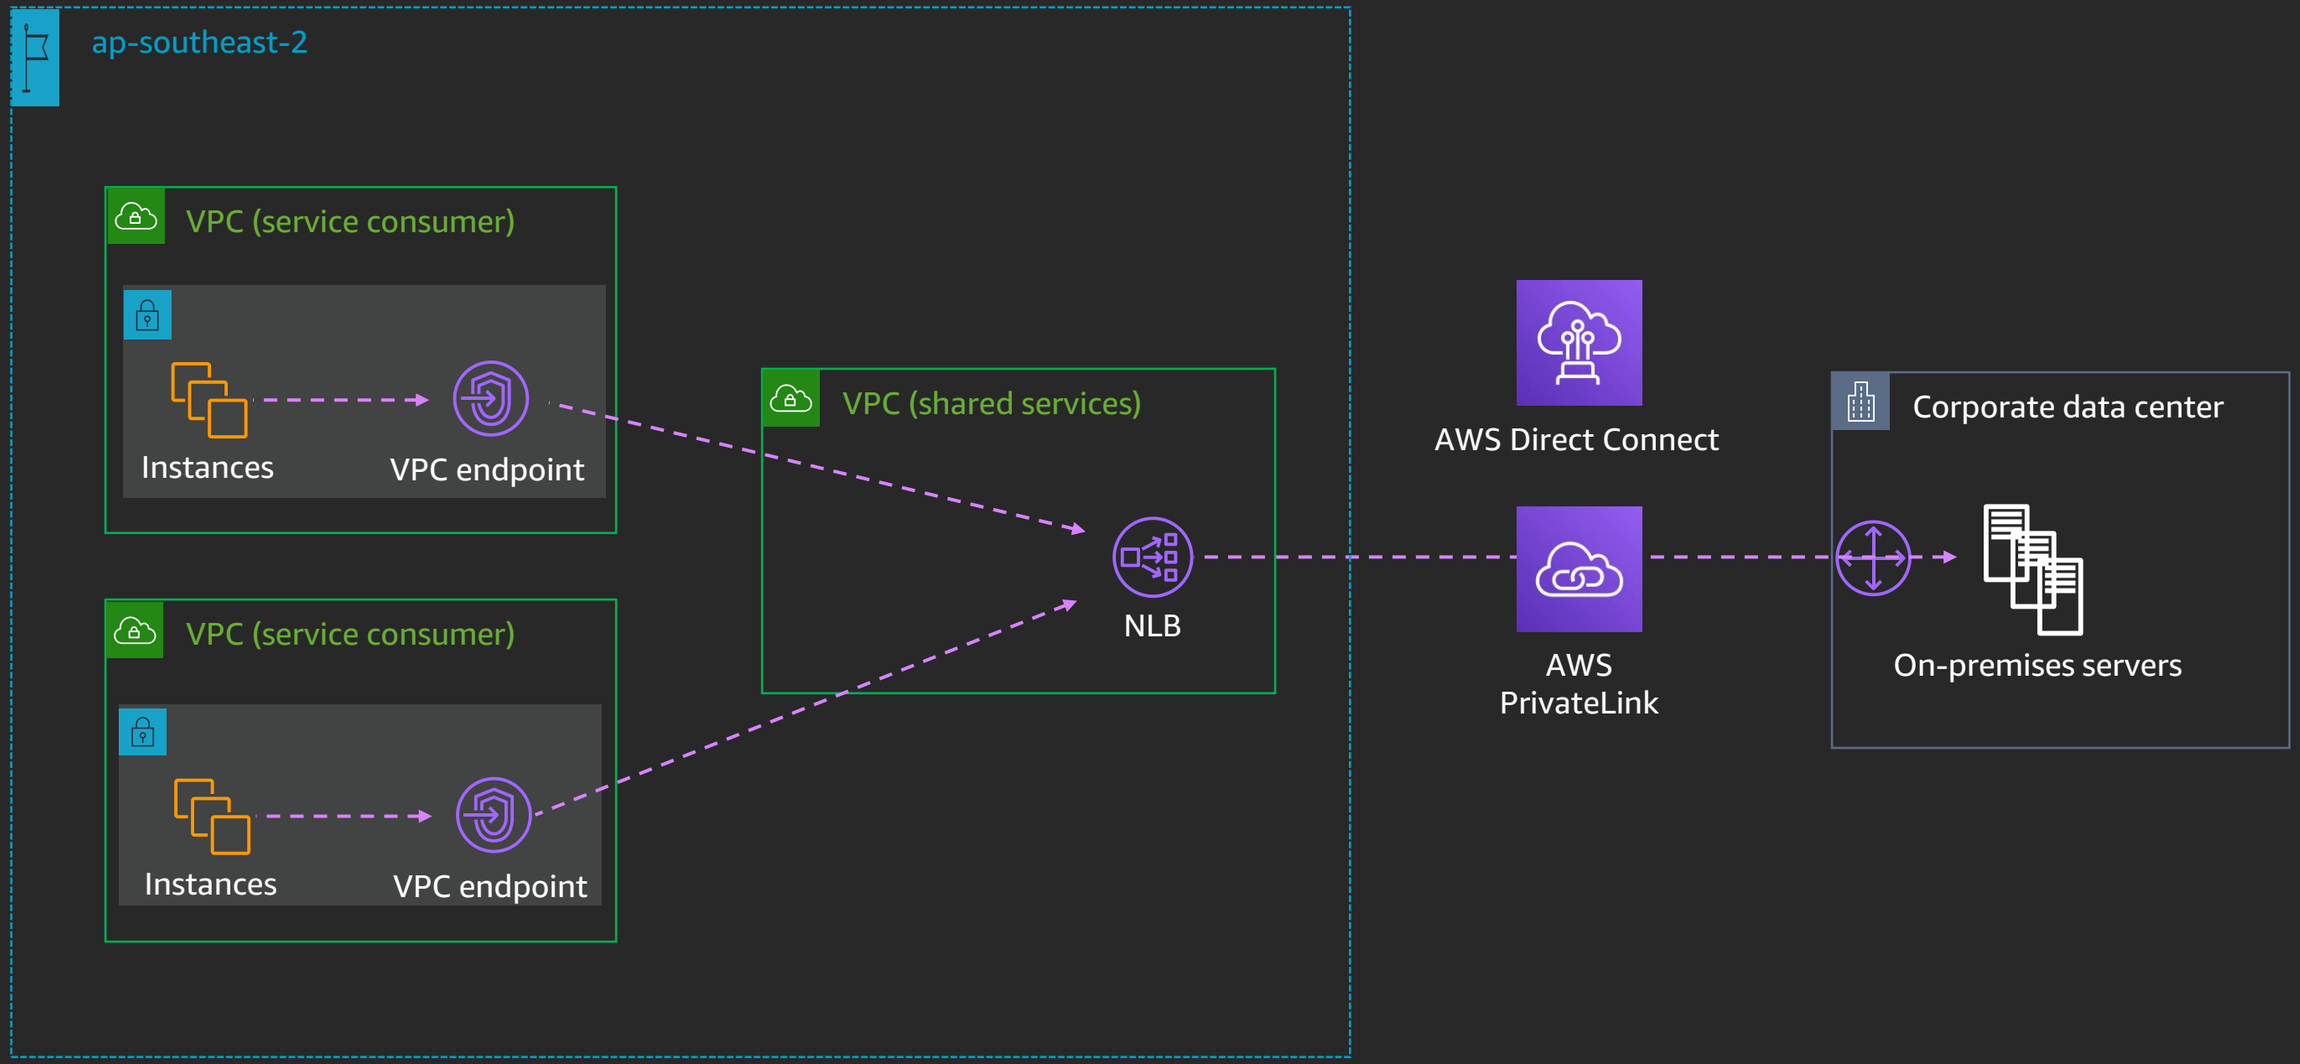

Cross-Region connectivity to cervices

Note: Avoid inter-Region dependencies

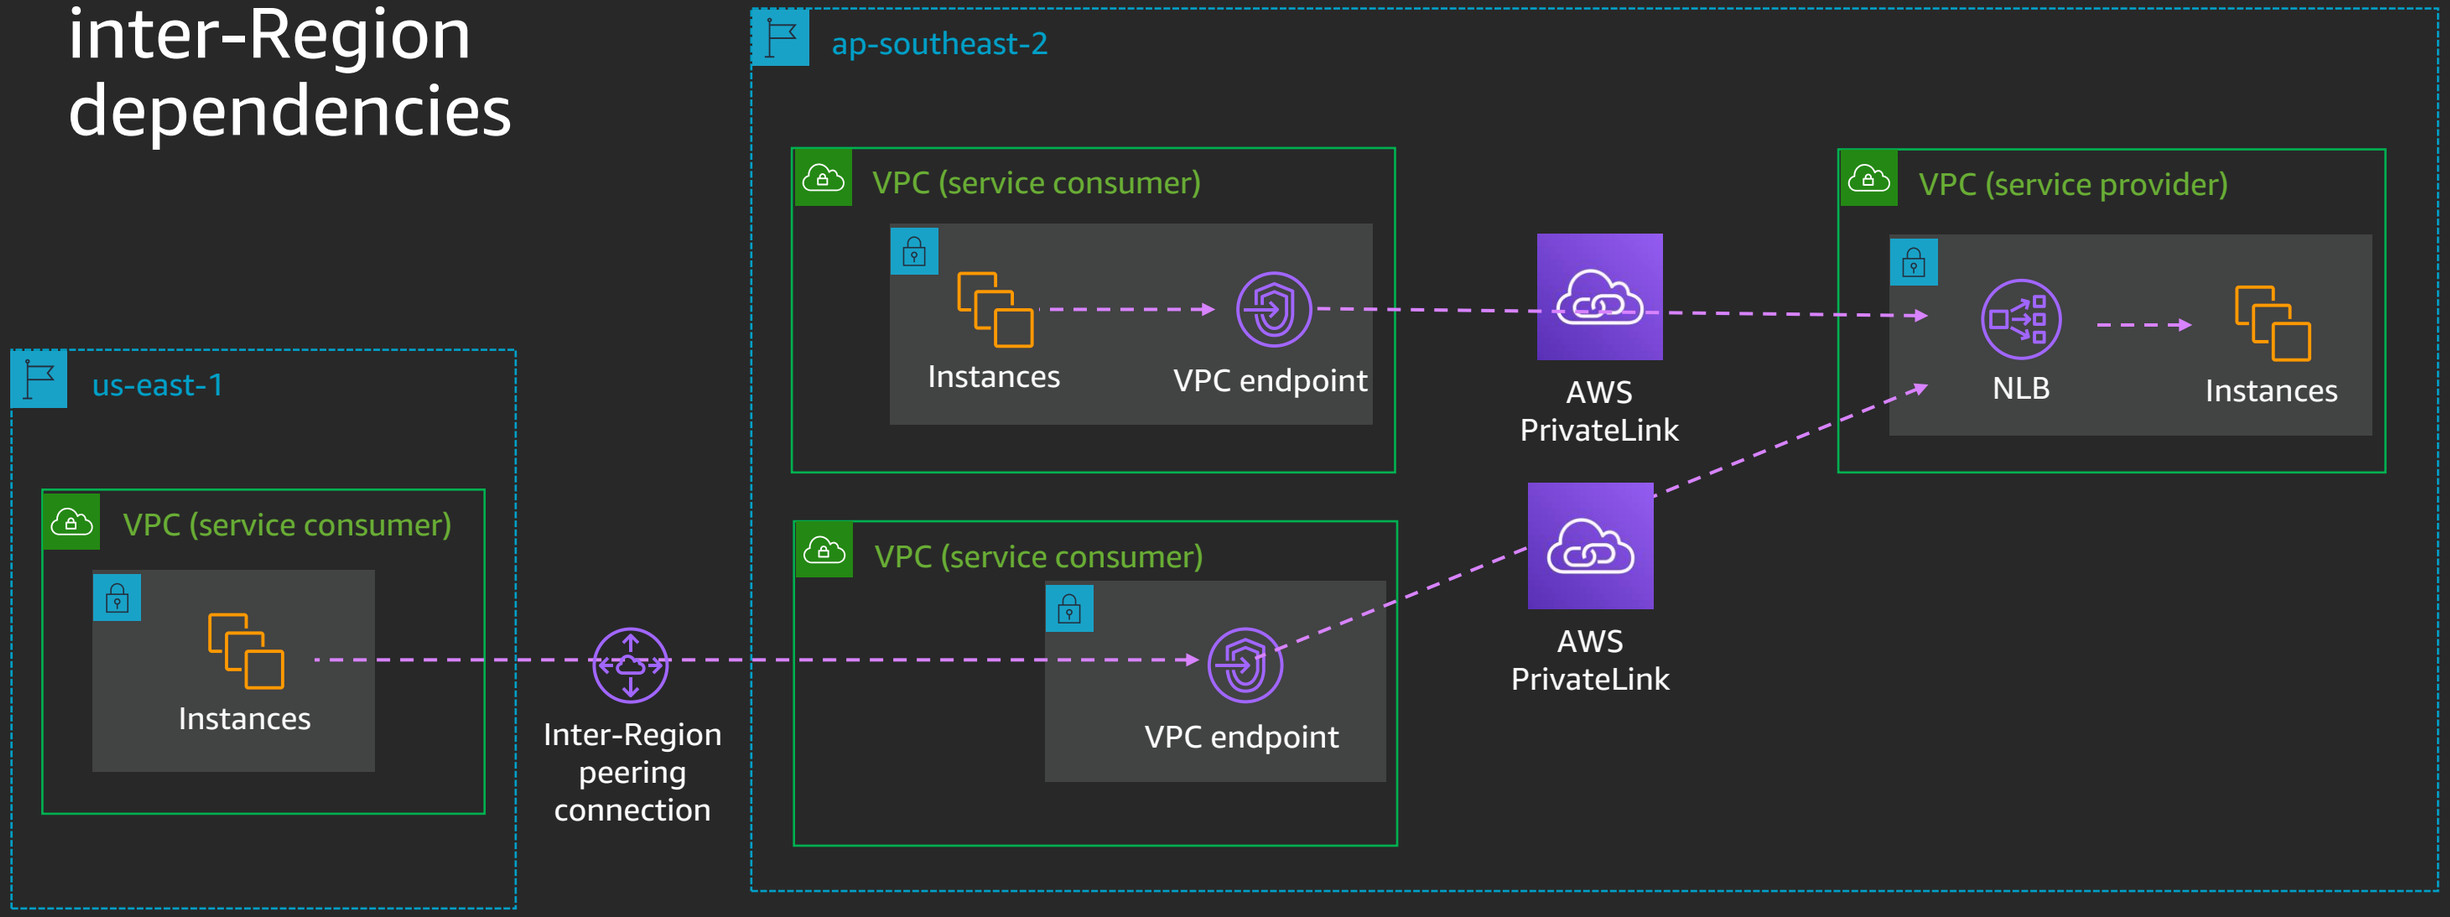

Presenting services in another Region

Note: Avoid inter-Region dependencies

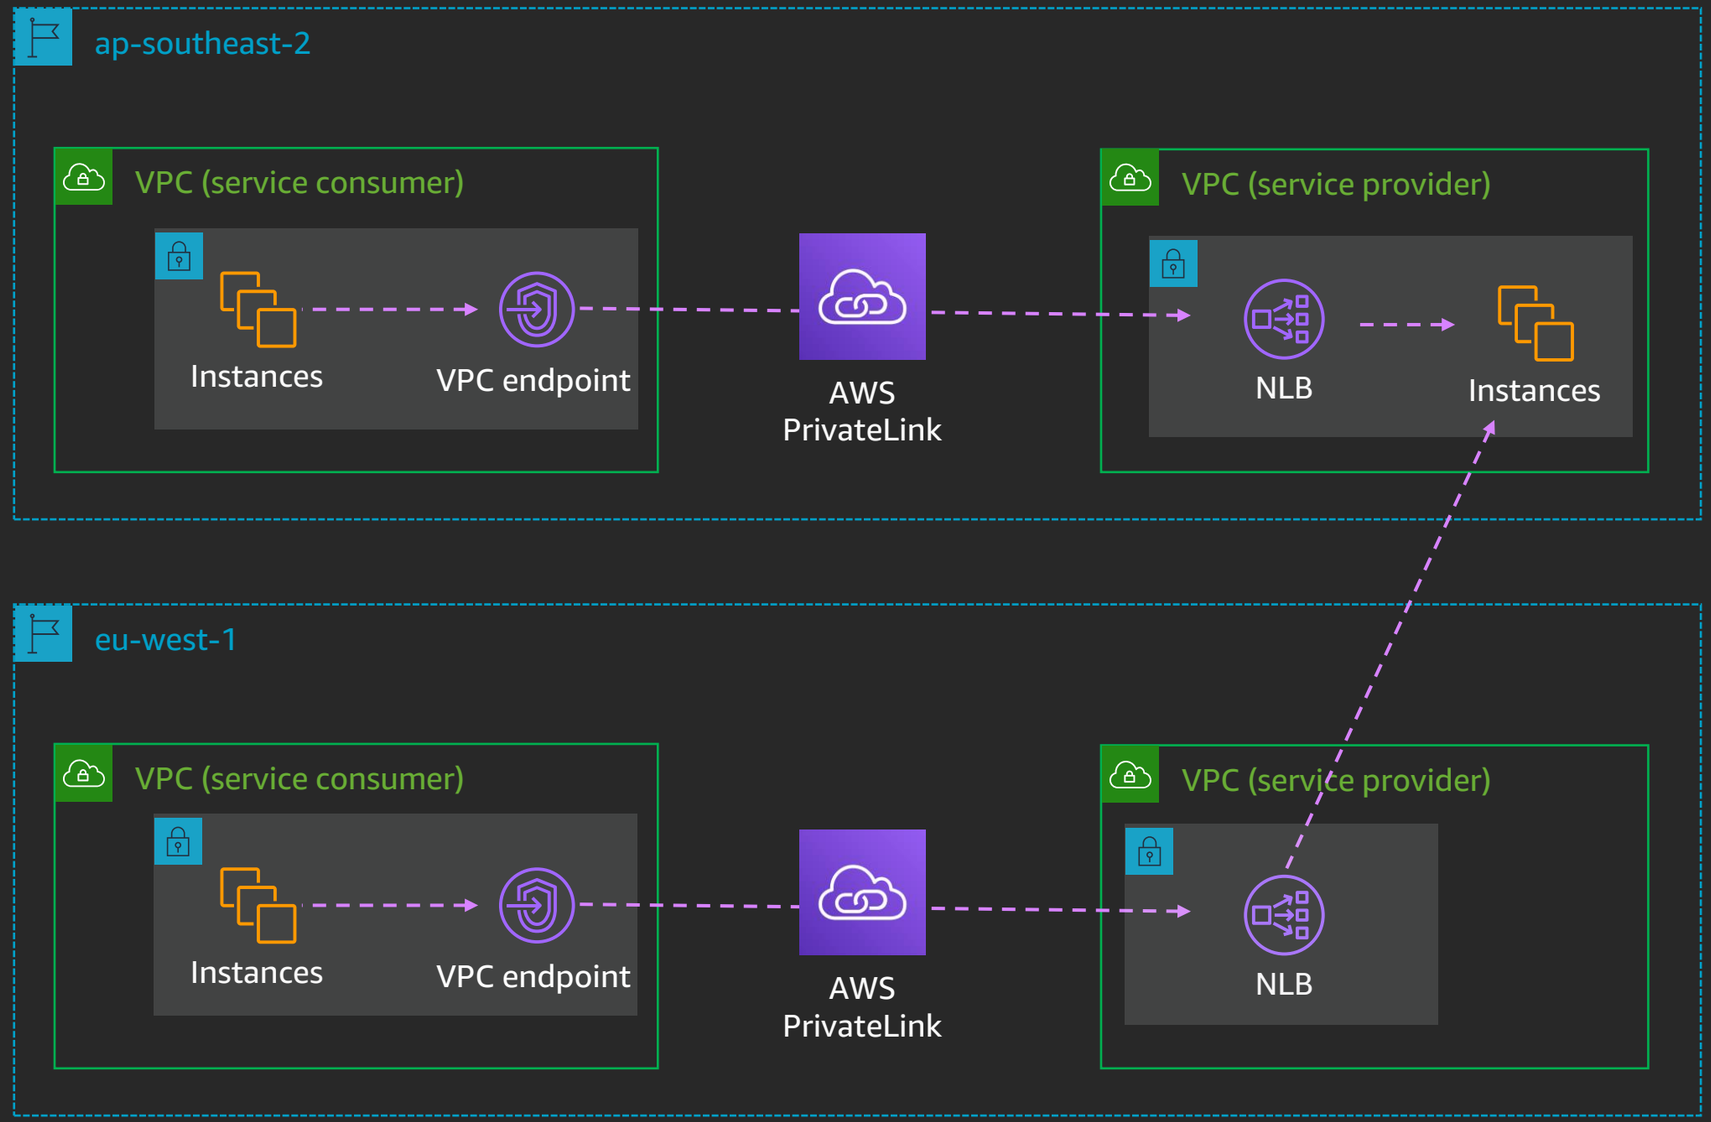

Shared VPC services

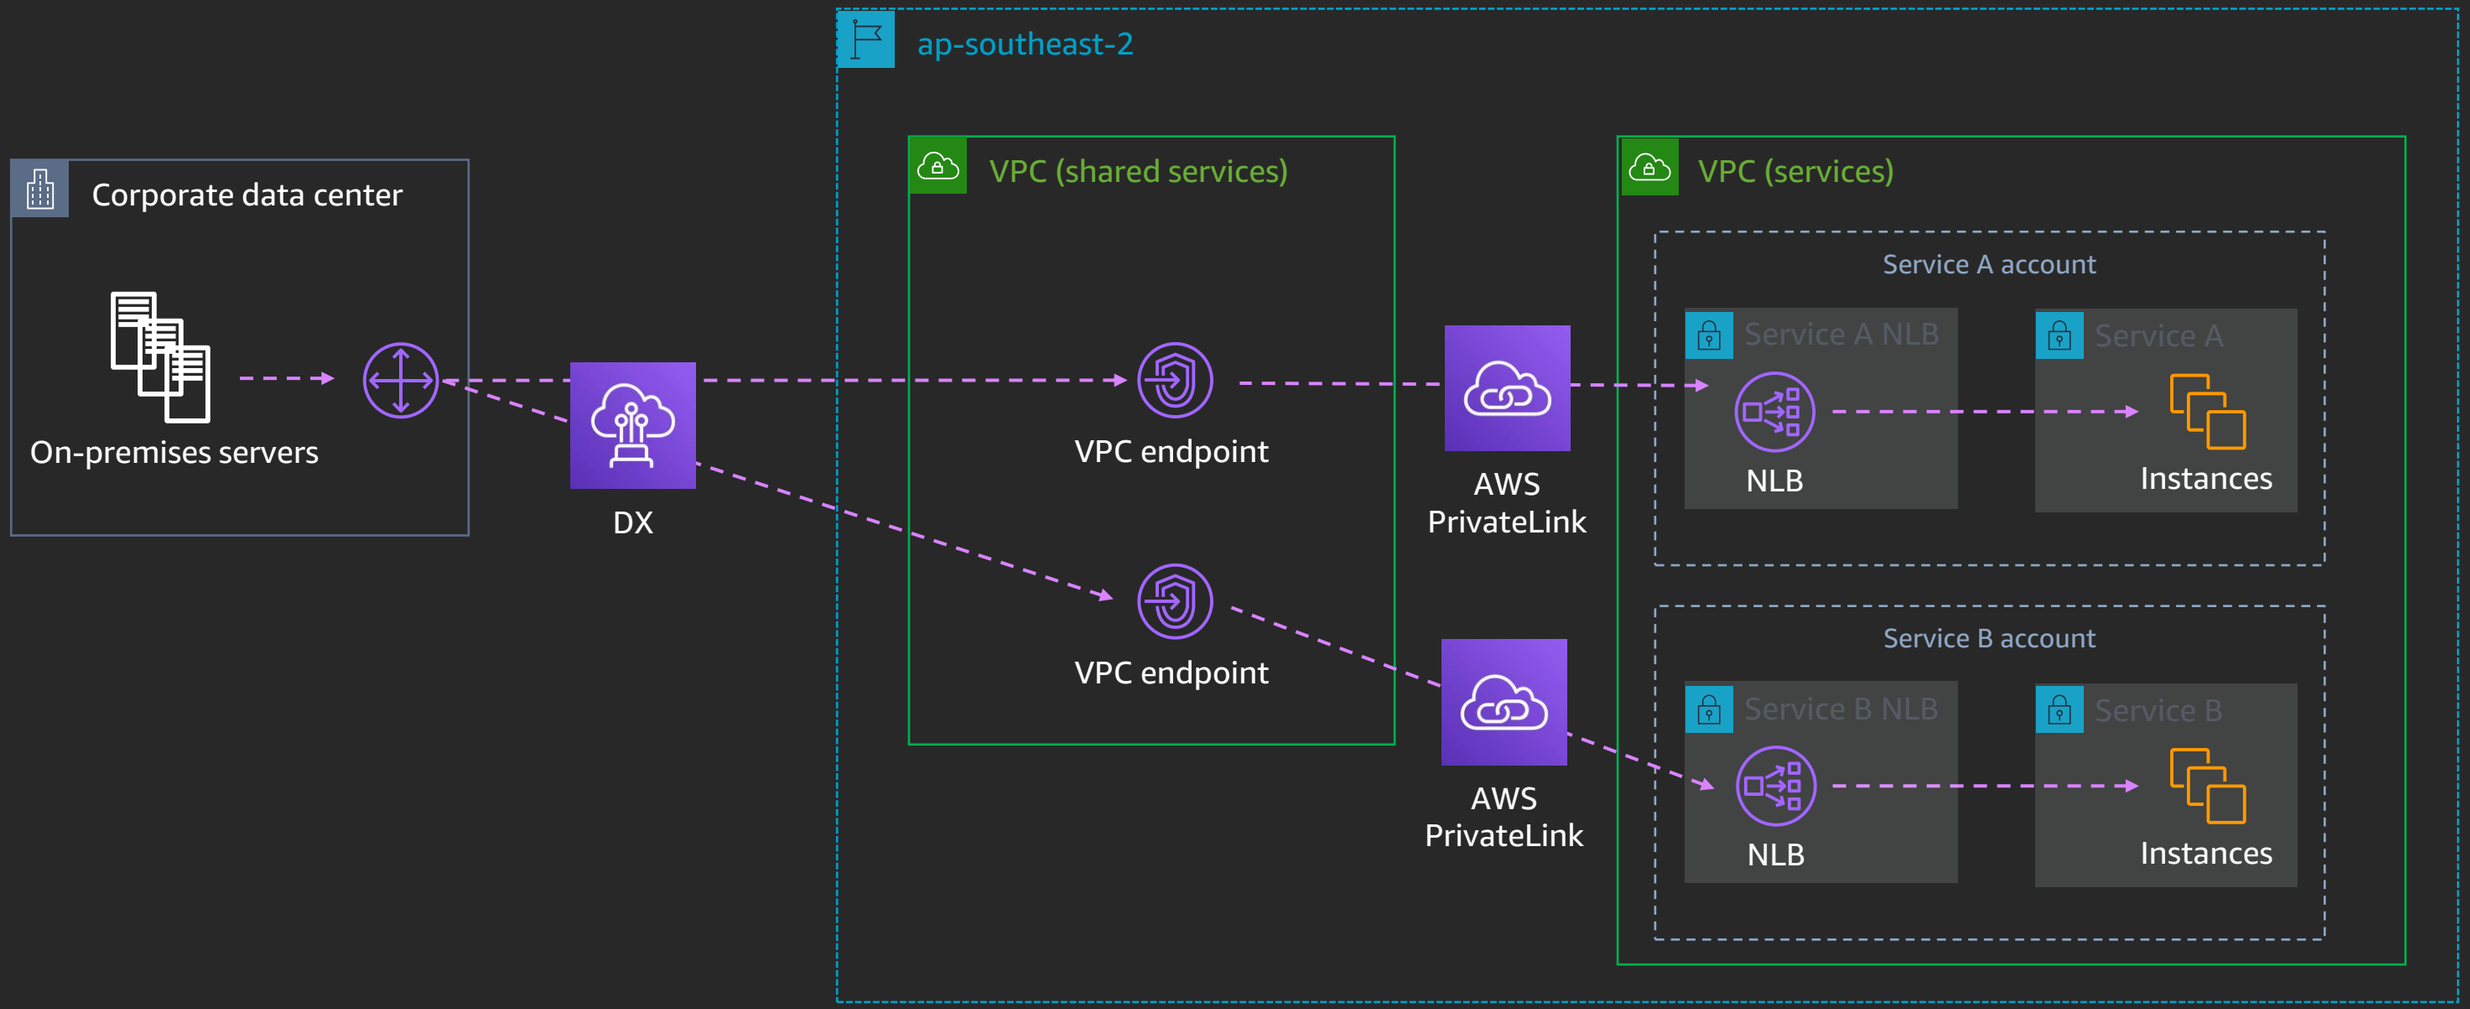

AWS PrivateLink

- Use at least two ENIs per VPCE
- Consider DNS infrastructure to meet your needs
- Ensure that service provider NLB has an elastic network interface in each Availability Zone
  - Cross-zone load balancing if don’t have service in each Availability Zone
- Avoid building inter-Region dependencies

DNS

aws SUMMIT ONLINE

Amazon Route 53 in VPC

DNS in VPC known as:

- AmazonProvidedDNS
- VPC Resolver
- +2 Resolver
- .2 Resolver
- EC2 DNS resolver

The EC2 DNS Resolver needed an official name

Amazon Route 53 Resolver

Route 53 Resolver – VPC view

- Recursive DNS server
- +2 IPs from VPC CIDR
- Built-in redundancy

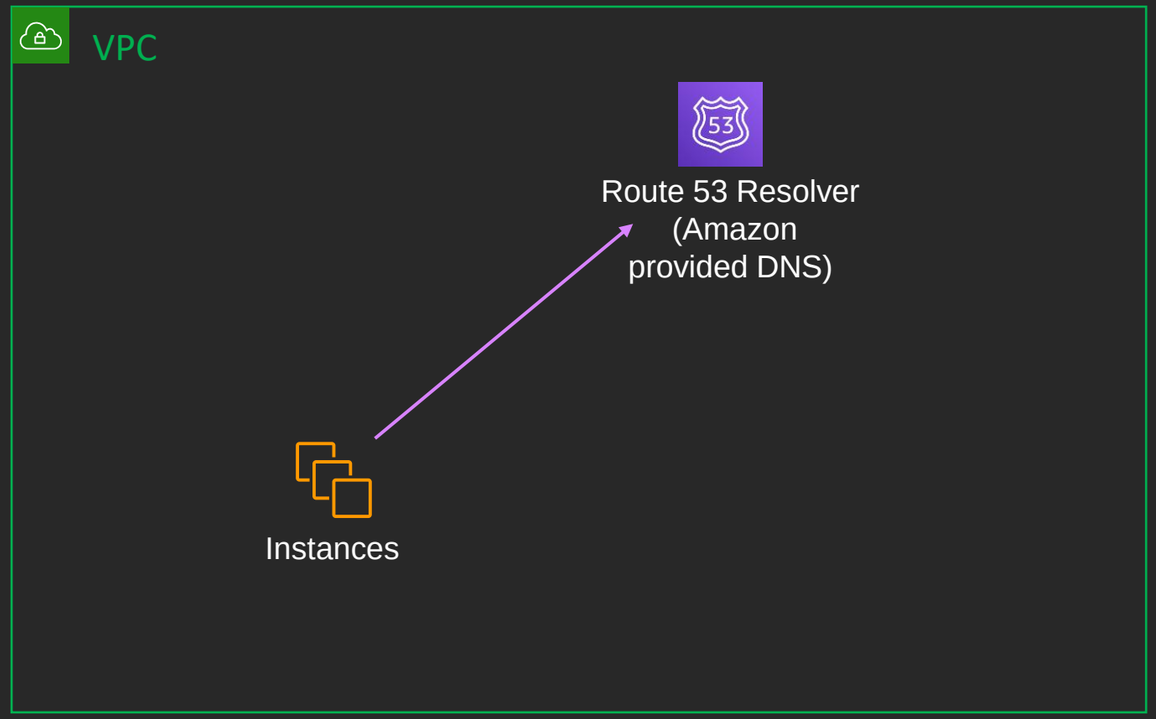

Route 53 Resolver

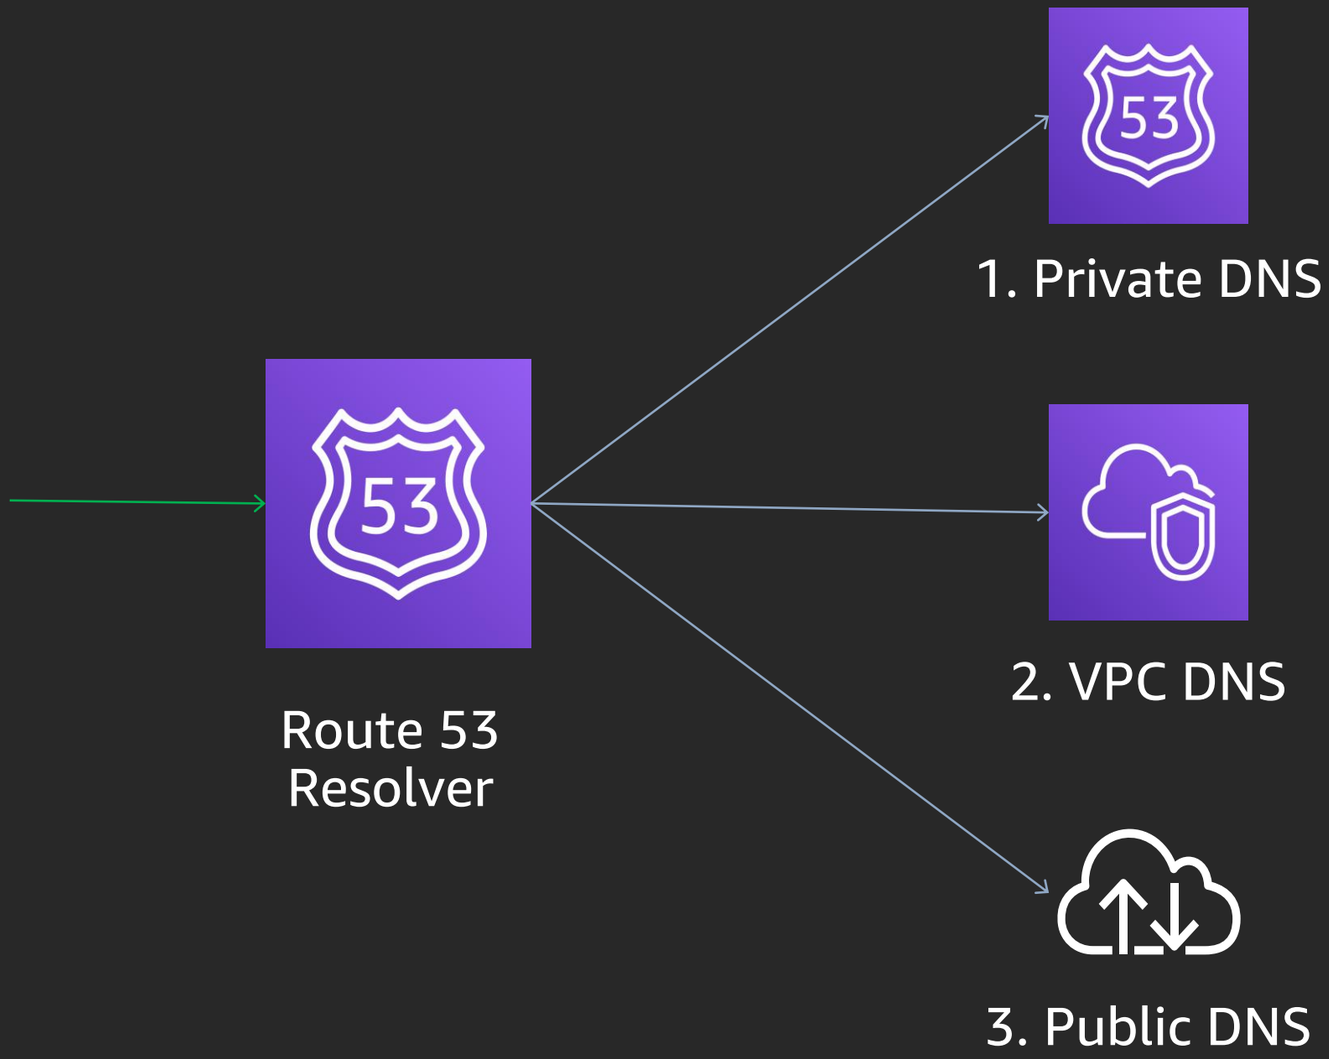

Route 53 Resolver endpoints

- Inbound endpoint:
  - Share VPC DNS view
- Outbound endpoint:
  - Share corporate DNS view
- Built-in redundancy
- Nomenclature:
  - One “endpoint” == multiple ENIs
- 10,000 QPS per elastic network interface

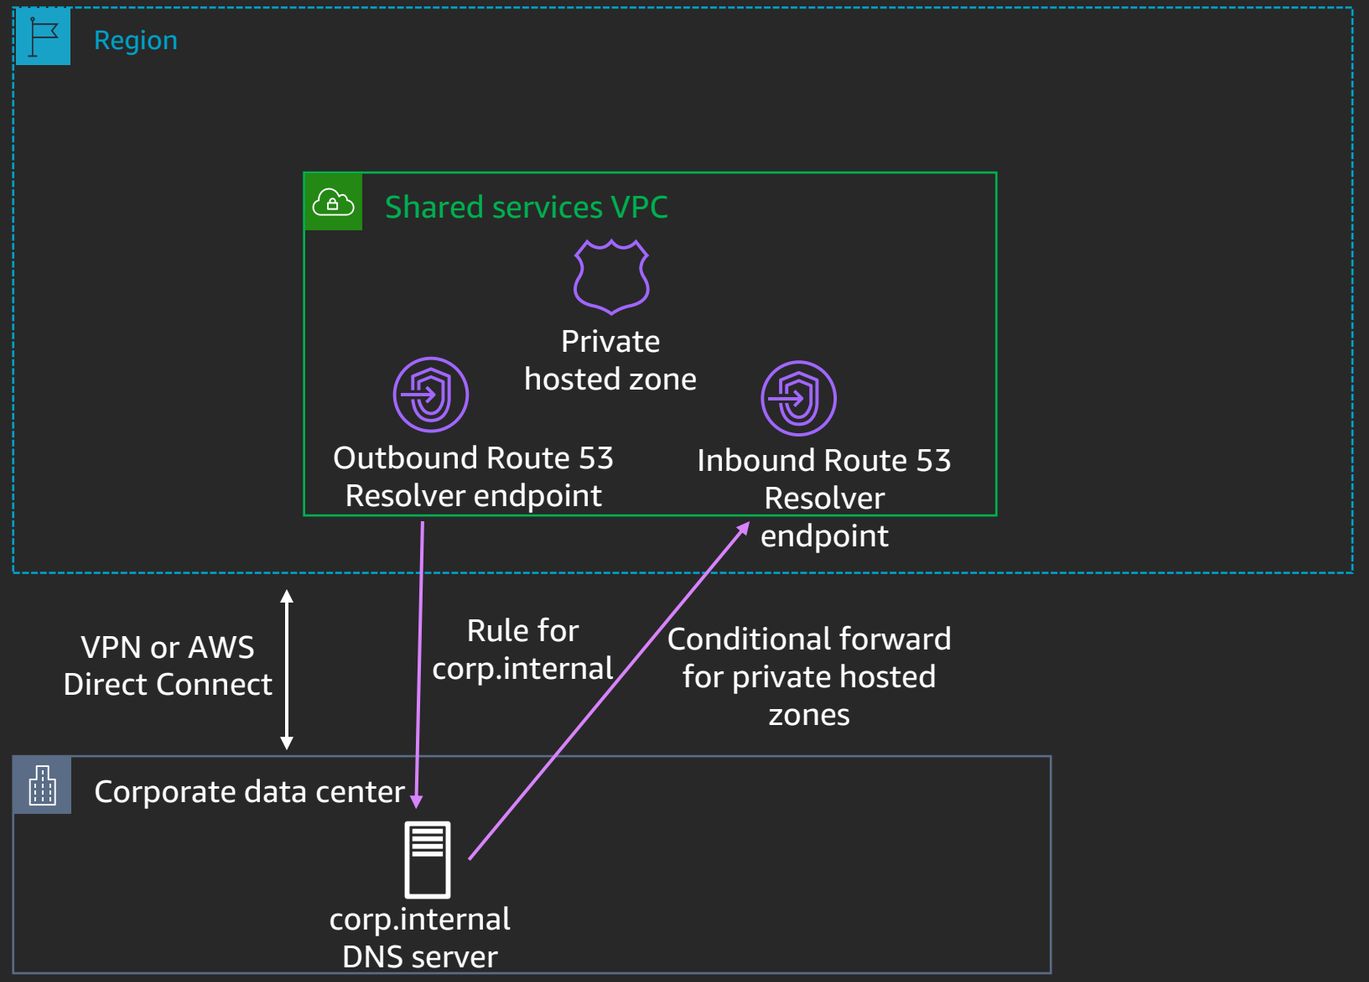

Route 53 Resolver rules

Configure how Route 53 Resolver makes queries

Two types: FORWARD and SYSTEM

Route 53 Resolver rules

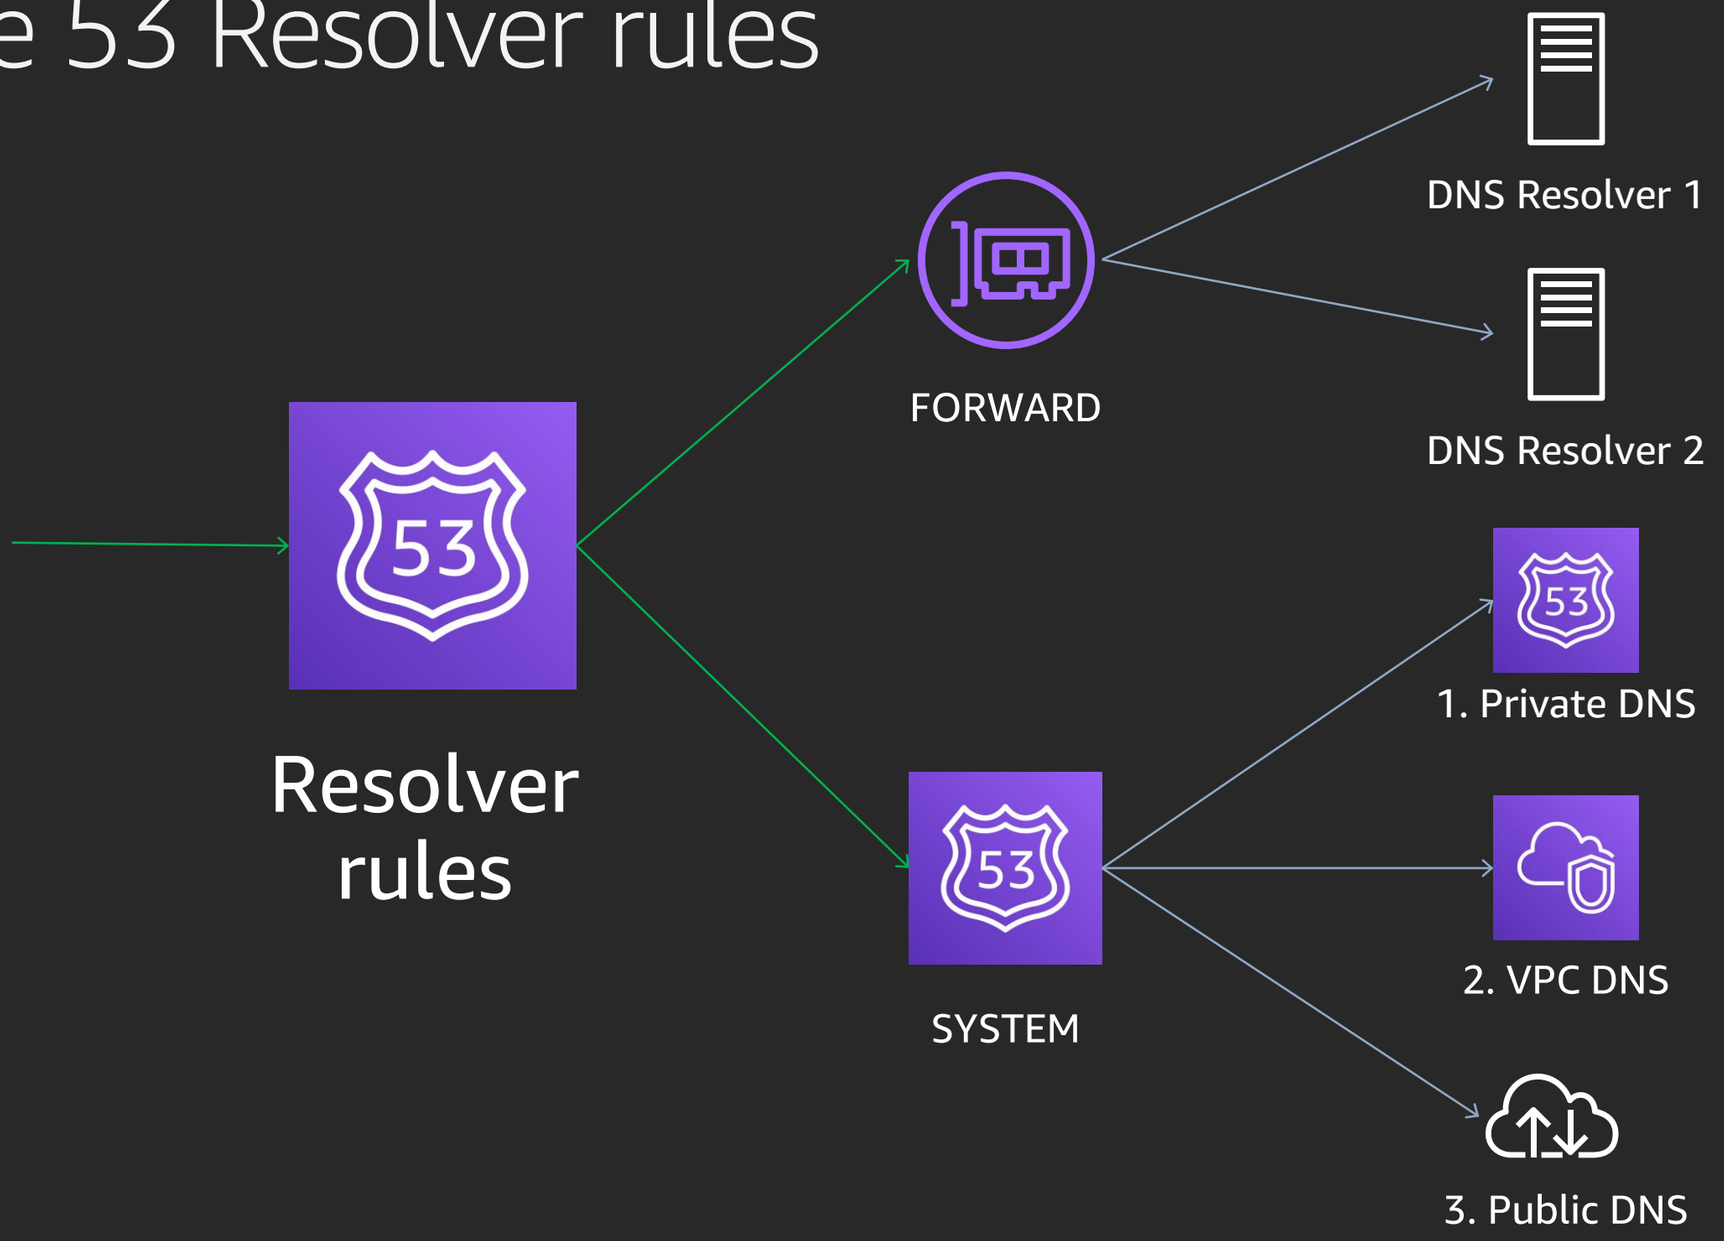

Sharing resolver VPC endpoints

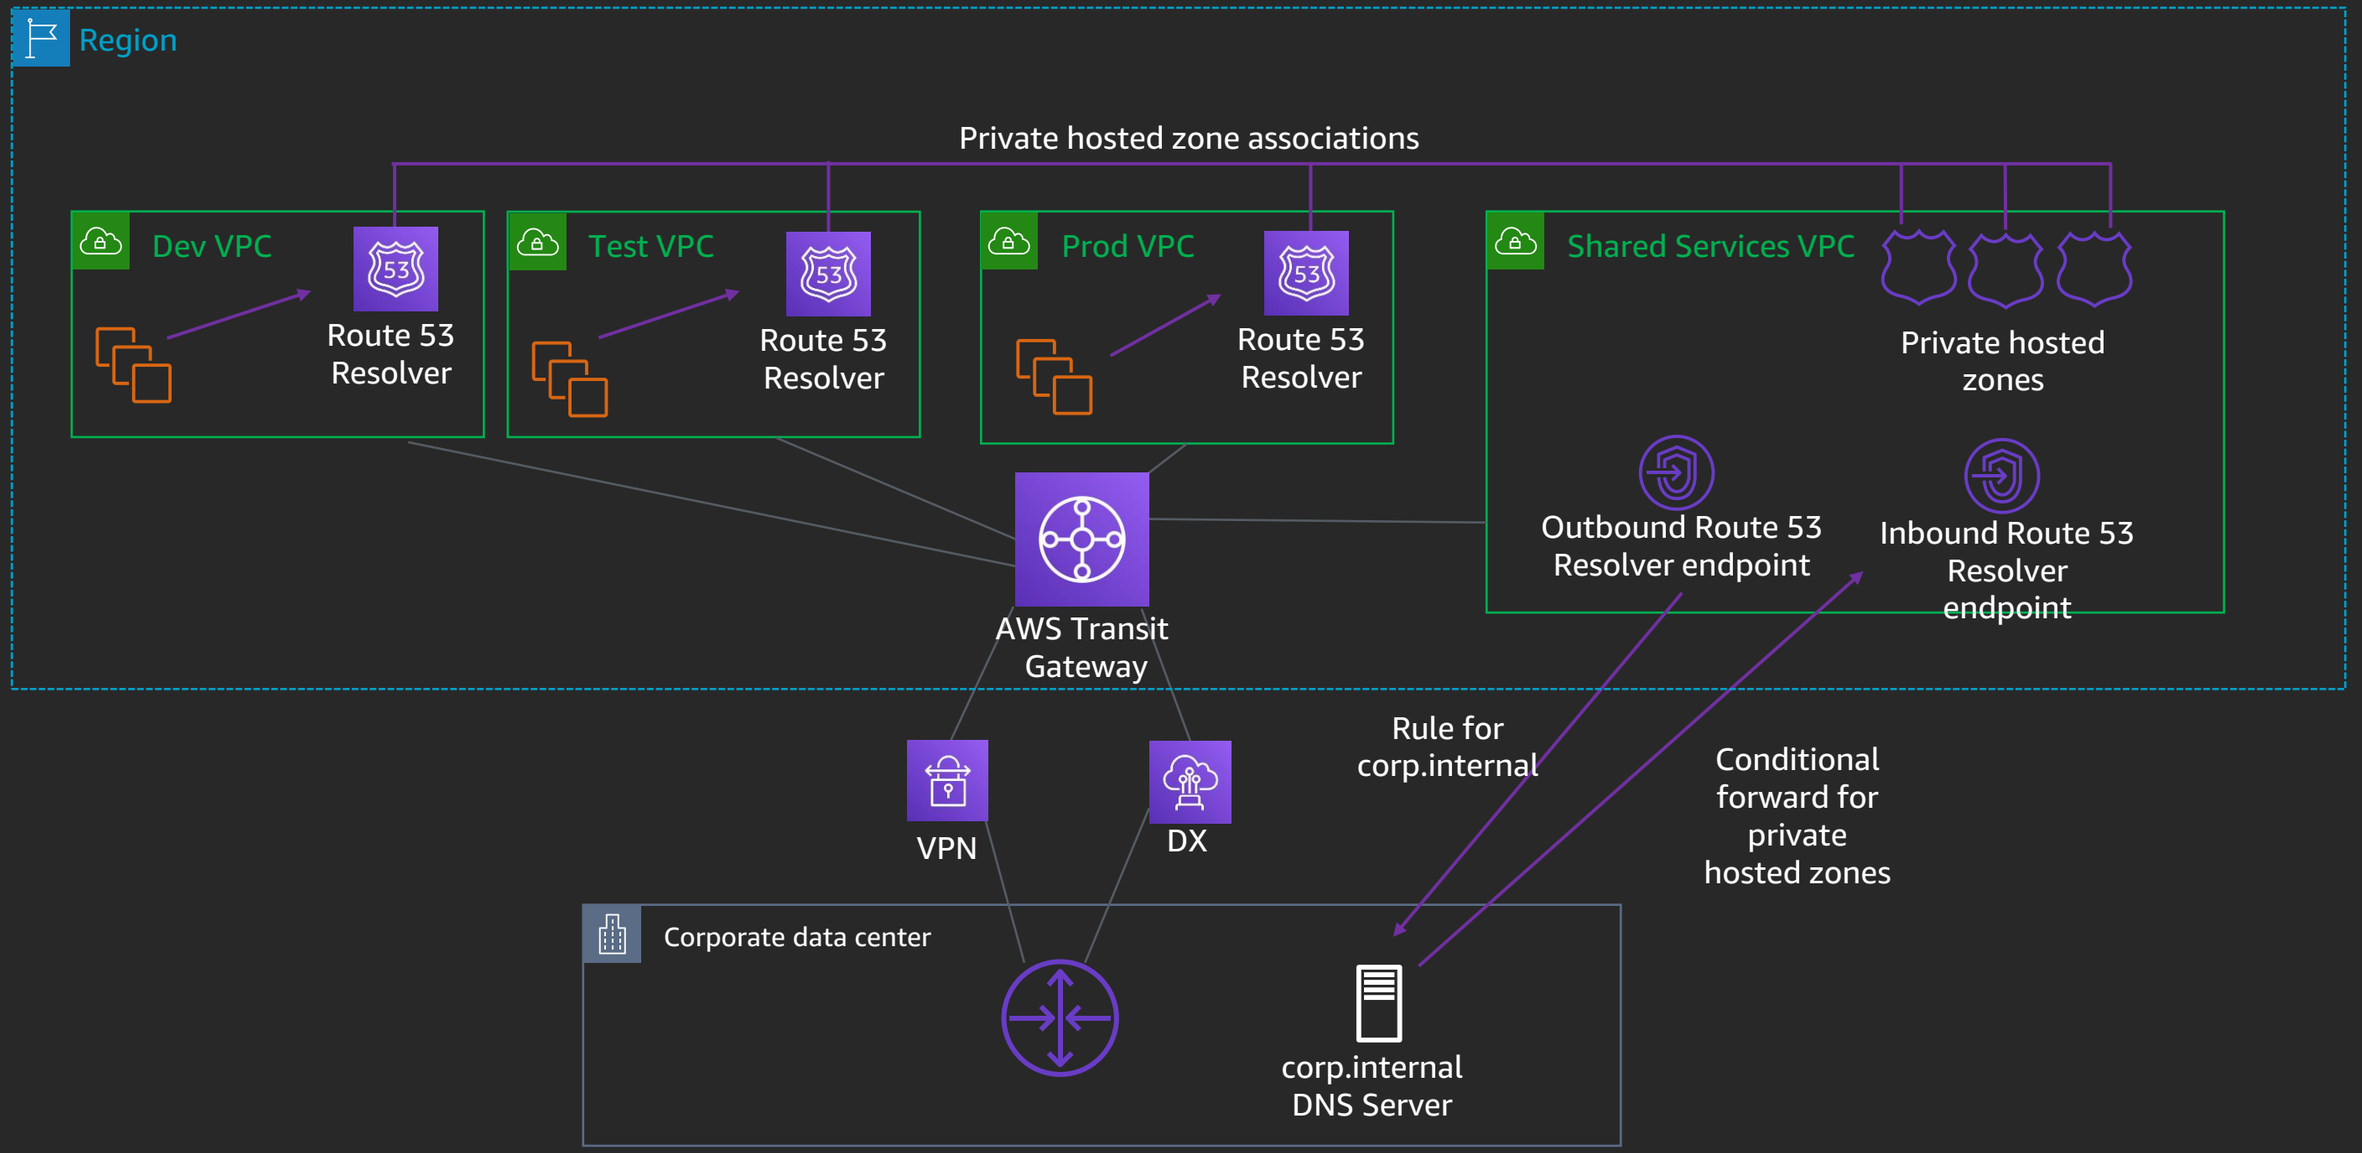

Active Directory hybrid DNS

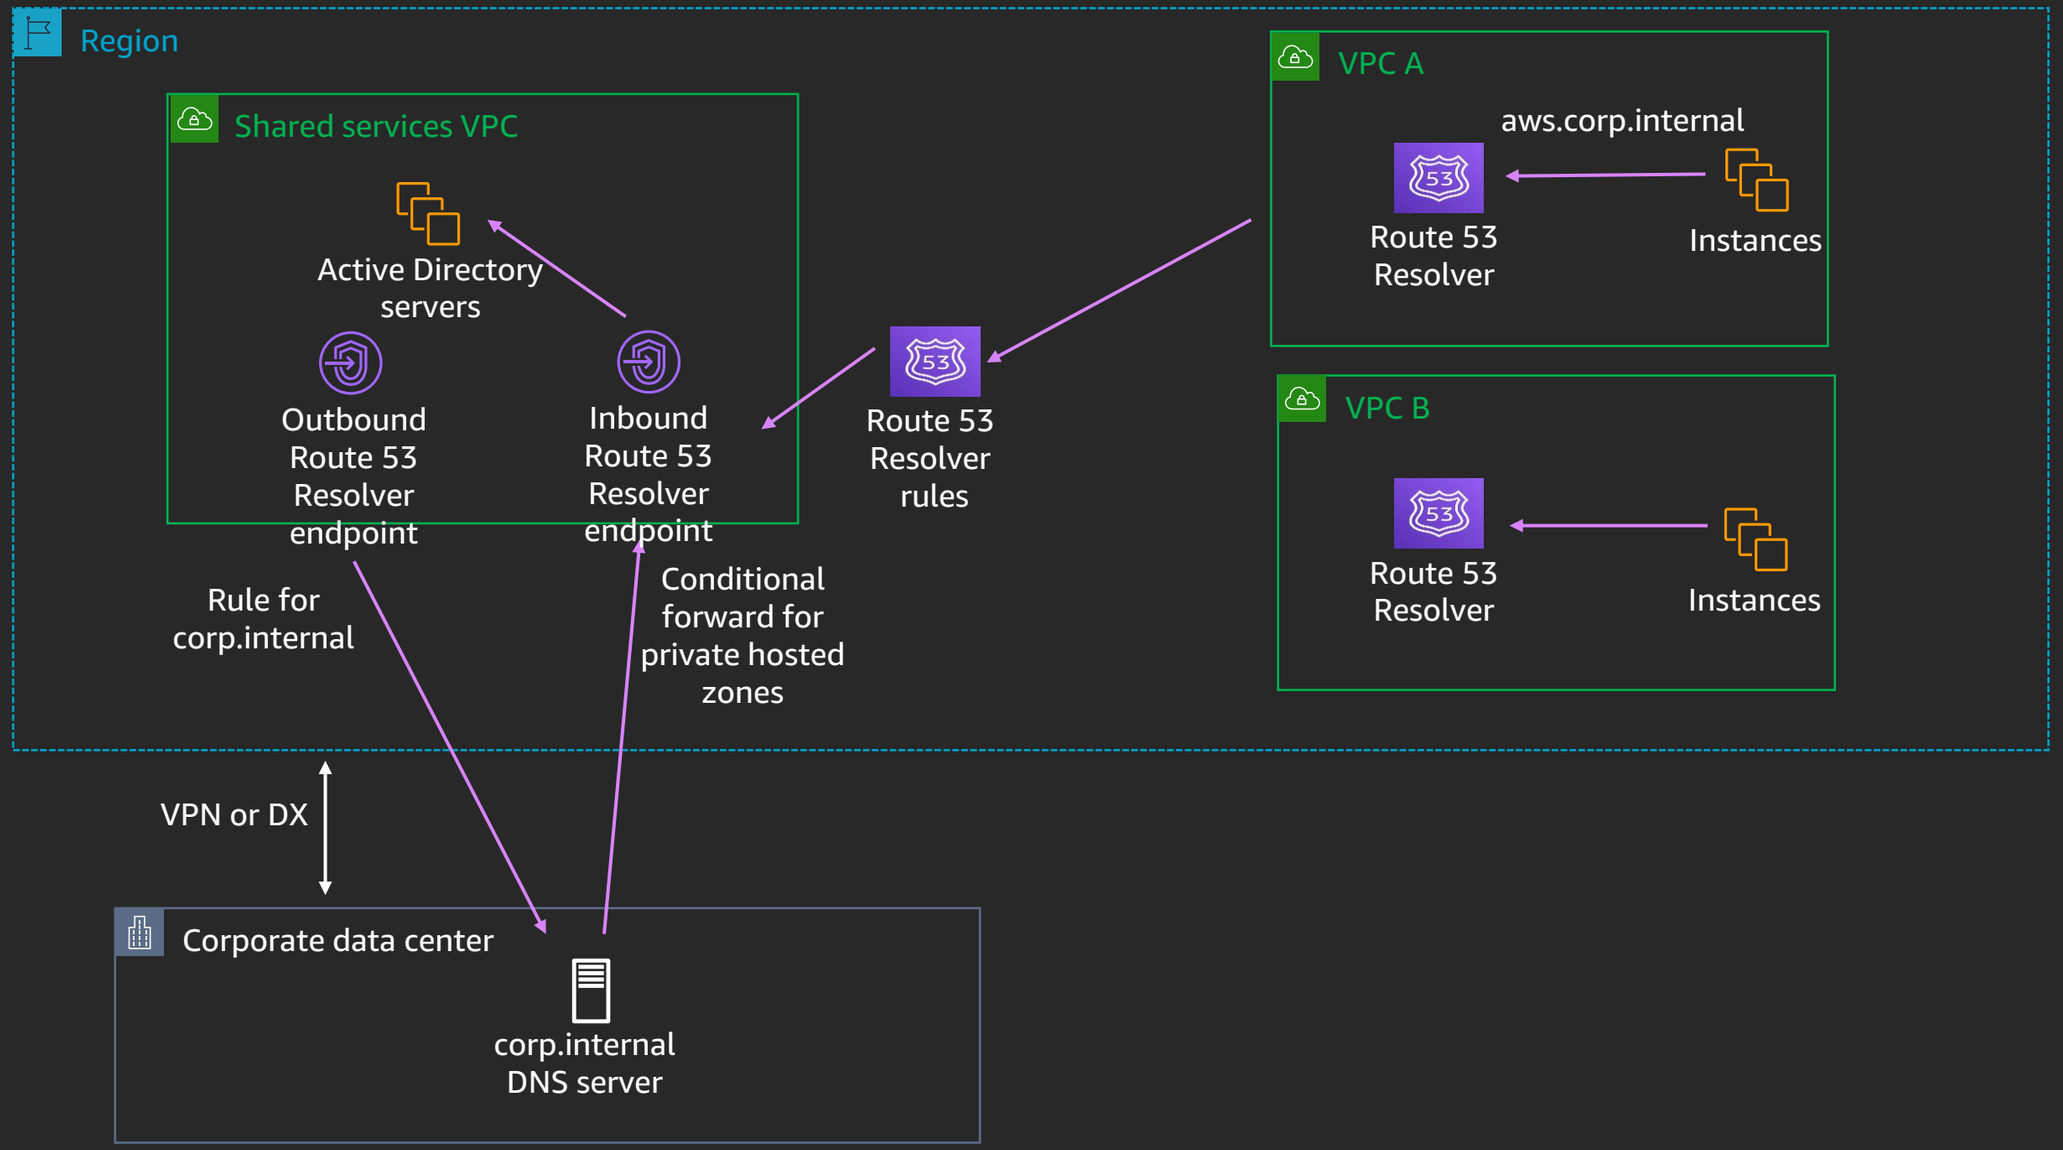

Route 53 best practices

- Within a VPC use the “.2” Route 53 Resolver
- Always use resolver endpoint ENIs in multiple Availability Zones
- Use conditional forwarding for on premises
- Avoid A records to VPCE ENIs
  - Alias record or CNAME
- Avoid pointing outbound endpoints at inbound endpoints
  - Limit: 10,000 QPS per elastic network interface
- Set CloudWatch alarms on resolver endpoints approaching QPS limits

Route 53 best practices

Inbound endpoints

- Use a retrying DNS resolver on premises
- Specify your IPs

Outbound endpoints

- Use forwarding sparingly
- Maintain fixed IPs as targets

Key takeaways

- AWS PrivateLink endpoints are highly available
- Amazon Route 53 is highly available and fault tolerant
- AWS PrivateLink and Amazon Route 53 enable you to create novel data flows

---

- AWS Direct Connect gateway makes changing VPCs easy
- Public VIFs are useful, but you are connected to everything
- AWS Transit Gateway is better than VPC peering for 99% of use cases

Bonus slides: AWS Outposts

AWS Outposts connectivity

- AWS Direct Connect or AWS VPN
- 1, 10, 40, or 100-Gb/s networking
  - 2x connections to local network
  - Each connection can be a single link or LAG
  - Two local VLANs (service and local data connectivity)
    - /30 or /31 required for local traffic (private IPs are okay)
    - /26 required for service VLAN (public or NATable range required)
- Local gateway (connects to your on-premises network)
  - Needs additional /26 or larger for 1:1 NAT to local subnet
- VPC is stretched to AWS Outposts using a subnet (same as Availability Zone)
- Cannot use AWS Outposts as transit to a Region

Thank you!

Brett Looney

brettski@amazon.com

aws SUMMIT ONLINE