Architecting IPv6 networks on AWS

Alexandra Huides

Principal Network Specialist Solutions Architect
AWS

The WHYs

IPv6 adoption on AWS

WHY ADOPT IPv6 ON AWS?

Improve network scalability

Start building experience

Minimize NAT (public & private)

Simplify global connectivity

Improve network scalability

SIMPLY MORE ADDRESSES

NO MORE SUBNETTING CHALLENGES

IPV6-ONLY DEPLOYMENTS SUPPORTED

Start building experience

EASY TO DEPLOY & TEST

BUILD BACWARDS COMPATIBILITY WITH IPV4

ADDRESS WHAT BRINGS VALUE

Minimize NAT (public & private)

NO NEED FOR PUBLIC NAT

NO NEED FOR PRIVATE NAT

IMPROVED VISIBILITY & SECURITY

Simplify global connectivity

NO MORE OVERLAPPING IPs

INTEGRATE MERGERS AND ACQUISITIONS

SUMMARIZATION AND EFFICIENT ROUTING

Approaches

IPv6 adoption on AWS

IPv6 adoption approaches

- Outside in (Edge first)
- Inside out (Internal first)

Outside in (Edge first)

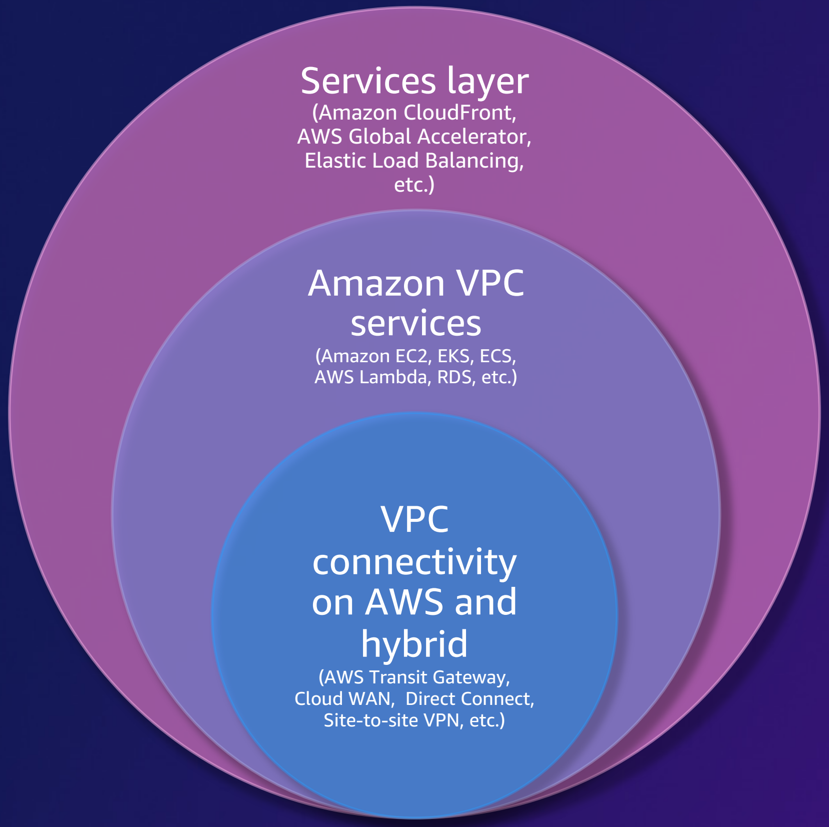

Services layer
(Amazon CloudFront, AWS Global Accelerator, Elastic Load Balancing, etc.)

Amazon VPC services
(Amazon EC2, EKS, ECS, AWS Lambda, RDS, etc.)

VPC connectivity on AWS and hybrid
(AWS Transit Gateway, Cloud WAN, Direct Connect, Site-to-site VPN, etc.)

Outside in (Edge first)

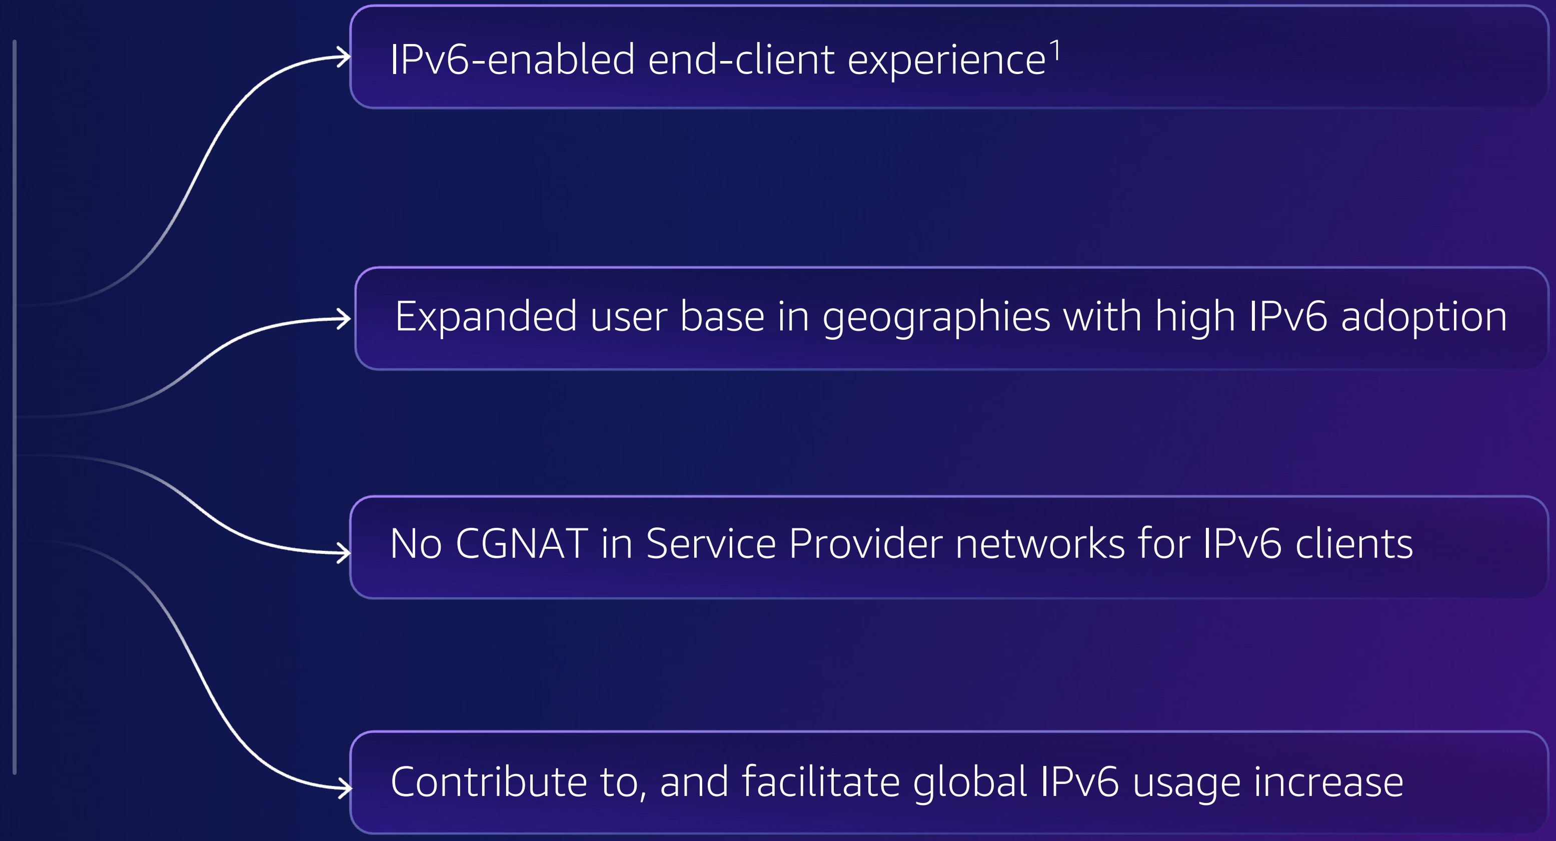

IPv6-enabled end-client experience[1]

Expanded user base in geographies with high IPv6 adoption

No CGNAT in Service Provider networks for IPv6 clients

Contribute to, and facilitate global IPv6 usage increase

[1] Check https://stats.labs.apnic.net/v6perf/XA

Outside in (Edge first)

Aroundhome

We have enabled IPv6 on our load balancers (ALB) and CloudFront distributions so customers can already reach our services through IPv6. It turned out to be a very smooth process without any hiccups.

Within a short amount of time we were able to report nearly 40 percent of our customer traffic to be IPv6

Hendrik Bergunde, Team Lead Technology - Aroundhome

Read more

Inside out (Internal first)

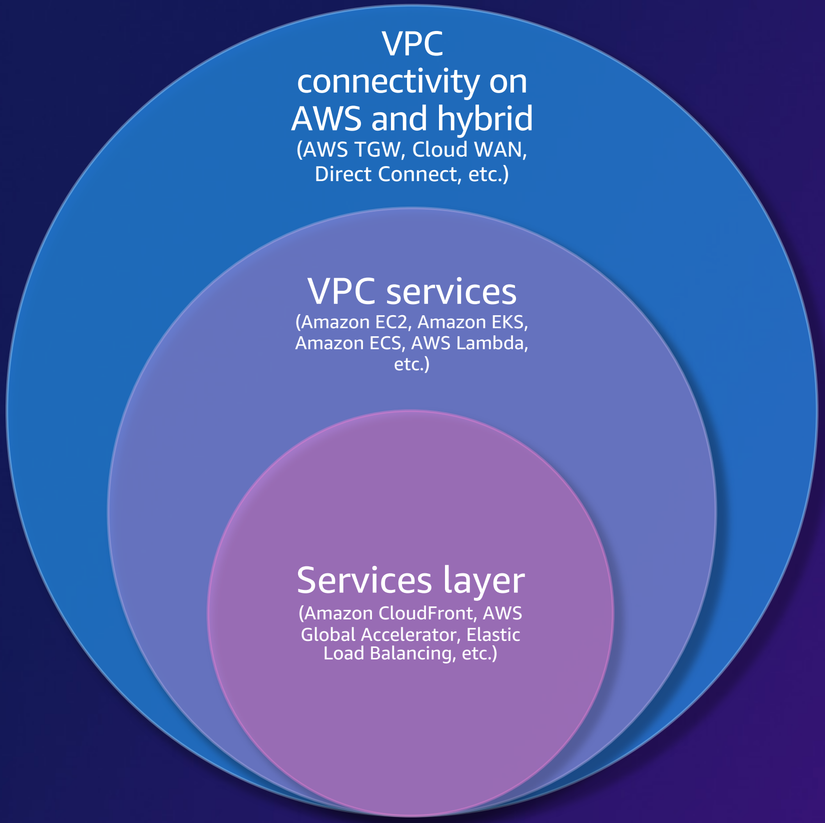
VPC connectivity on AWS and hybrid
(AWS TGW, Cloud WAN, Direct Connect, etc.)

VPC services
(Amazon EC2, Amazon EKS, Amazon ECS, AWS Lambda, etc.)

Services layer
(Amazon CloudFront, AWS Global Accelerator, Elastic Load Balancing, etc.)

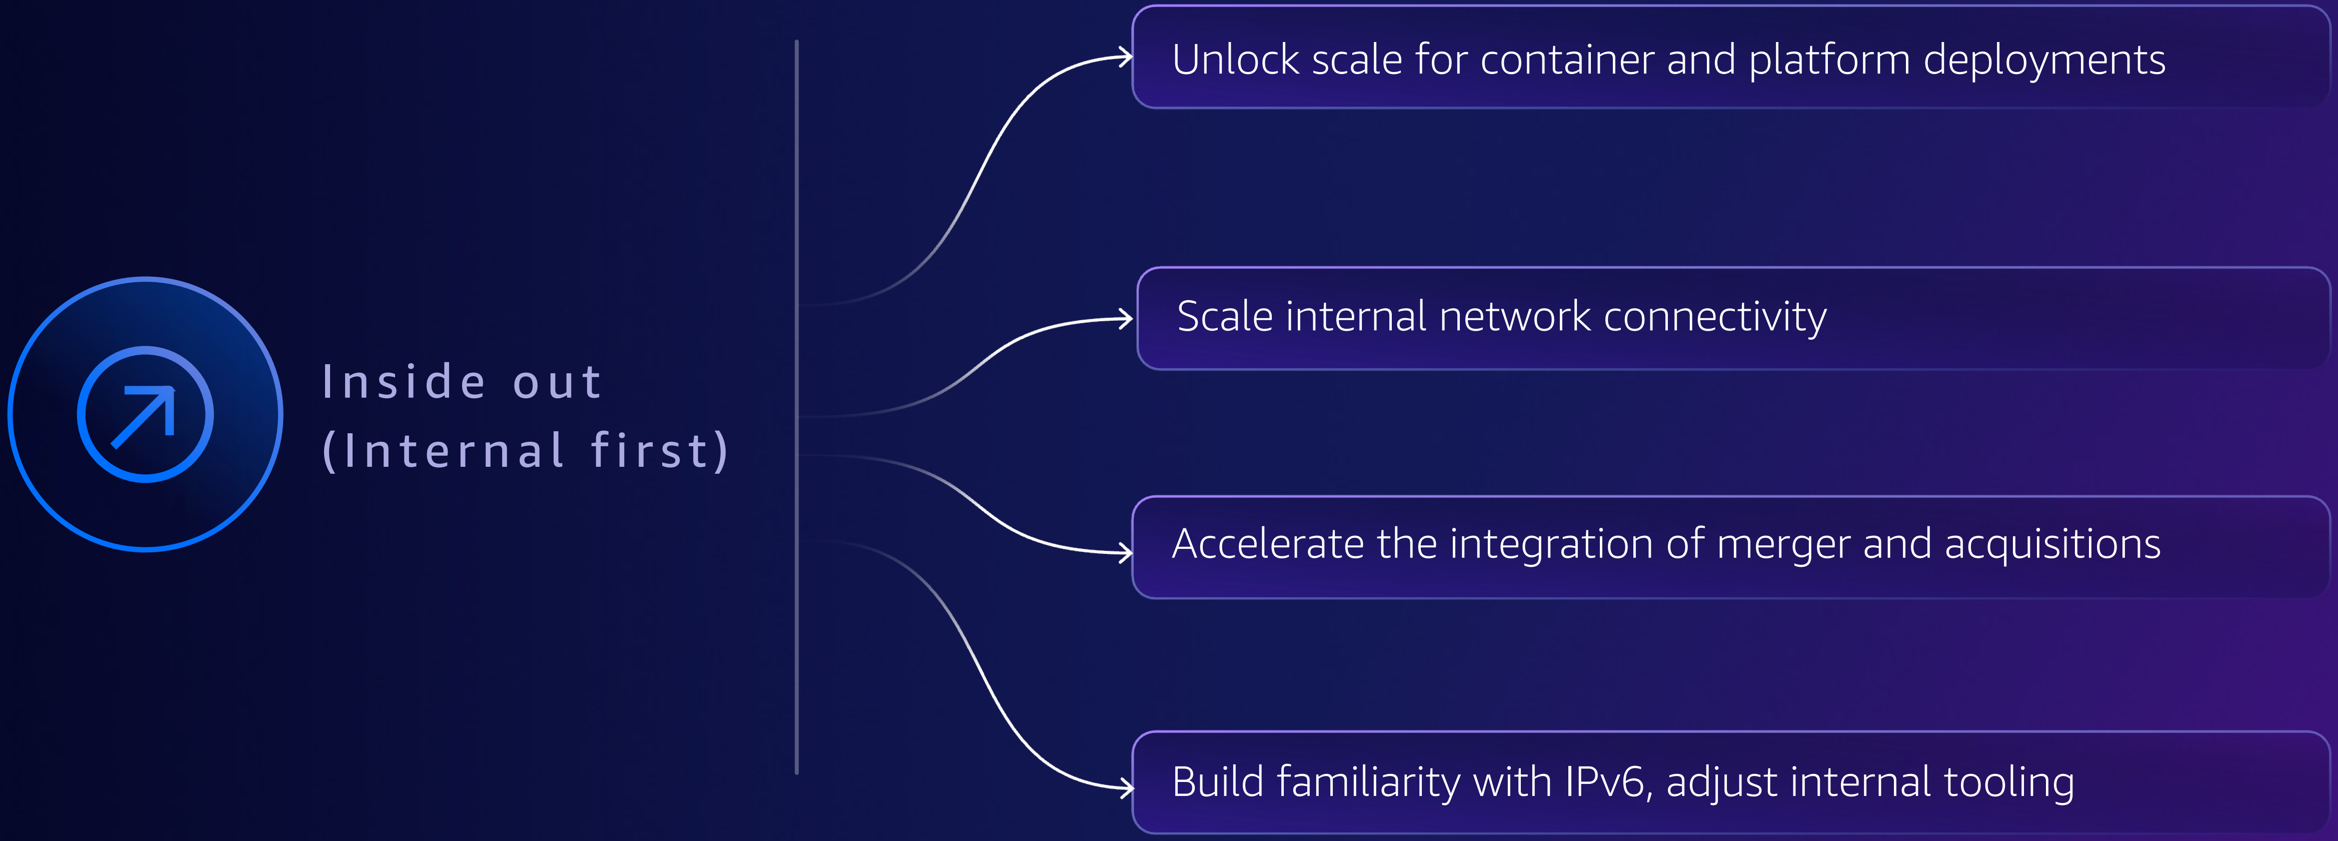

Inside out
(Internal first)

- Unlock scale for container and platform deployments
- Scale internal network connectivity
- Accelerate the integration of merger and acquisitions
- Build familiarity with IPv6, adjust internal tooling

Inside out
(Internal first)

NETFLIX

“IPv6 adoption in the internal network enabled the full IP reachability Netflix needed across the thousands of VPCs without the need for Network Address Translation. Also, the Egress-only Internet Gateway helped maintain the private subnets security posture.

Enabling IPv6 across the Netflix streaming platform in AWS enabled continued hyperscale growth, scalability and innovation.”

Donavan Fritz, Senior Network SRE - Netflix

Read more

IPv6 adoption on AWS

More customer stories

infor

IPv6 adoption on AWS

Outside in and Inside out are complementary approaches!

Focus areas

IPv6 adoption on AWS

IPv6 adoption focus areas

Network

Operating system

App code

Services & tools

Where to start

IPv6 adoption on AWS

IPv6 adoption
where to start

IPv6 addressing plan

Dual stack Amazon VPC

IPv6 addressing plan

Dual stack Amazon VPC

Amazon-provided GUA (VPC-level)

IPv6 addressing plan

Dual stack Amazon VPC

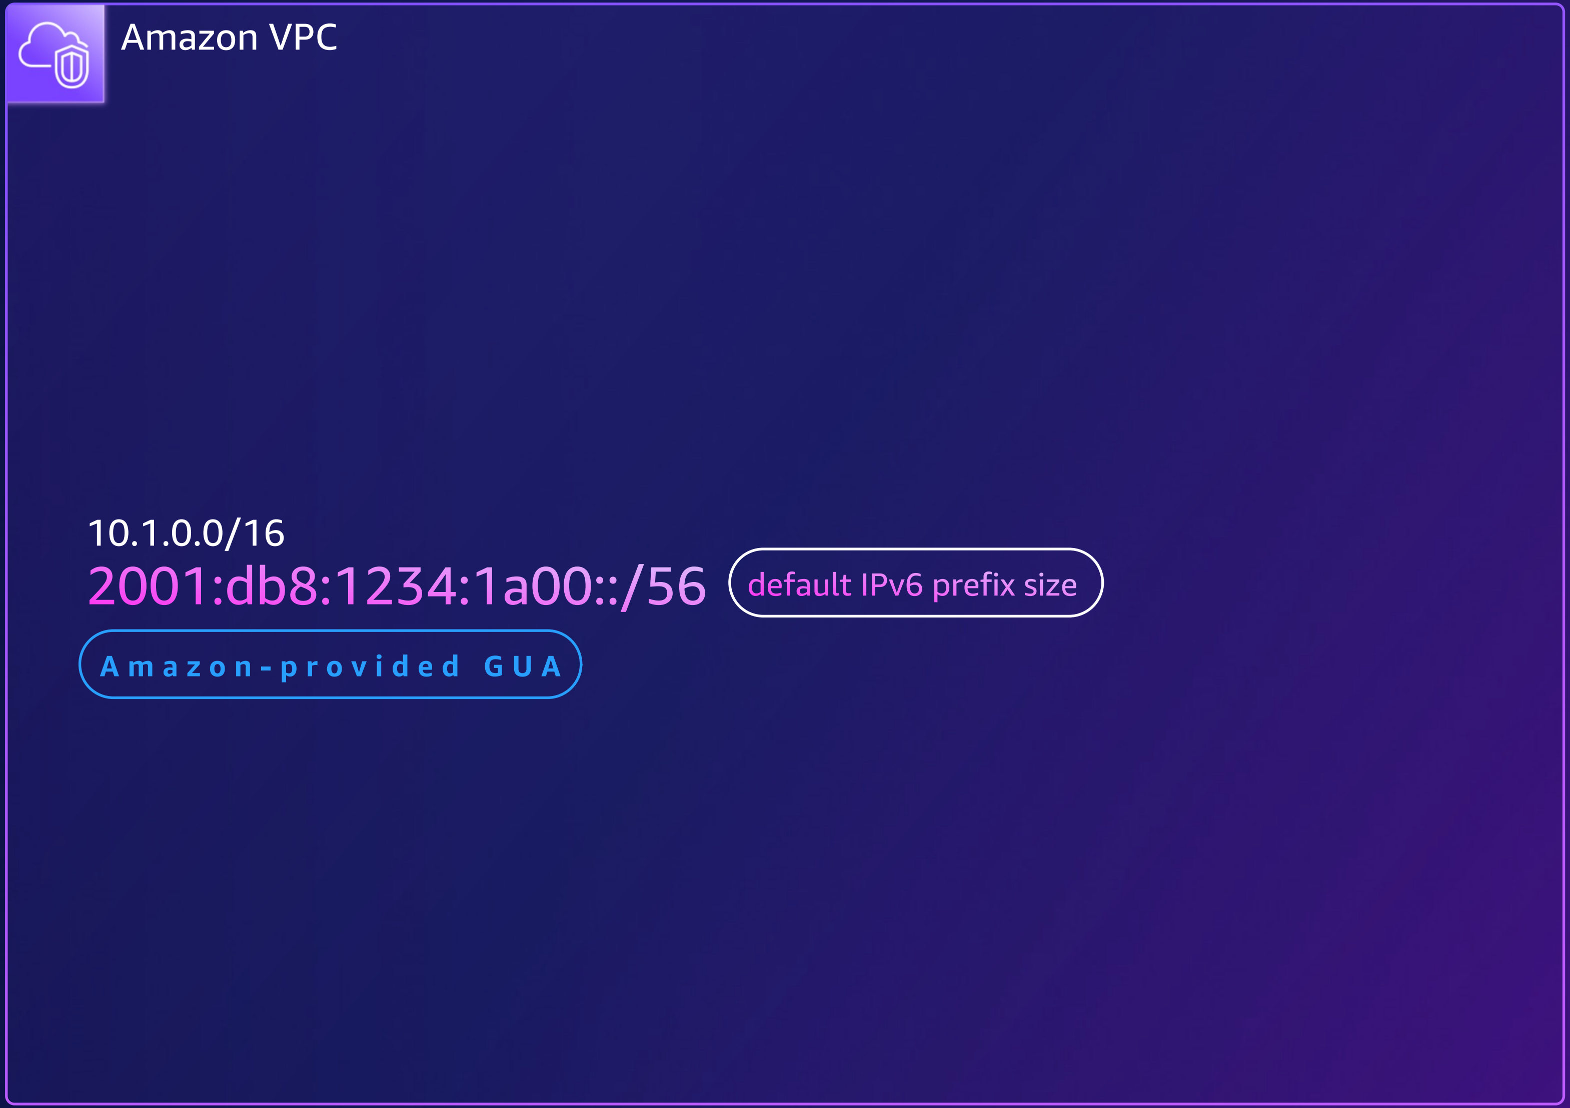

Amazon VPC

10.1.0.0/16

2001:db8:1234:1a00::/56

default IPv6 prefix size

Amazon-provided GUA

Dual stack Amazon VPC

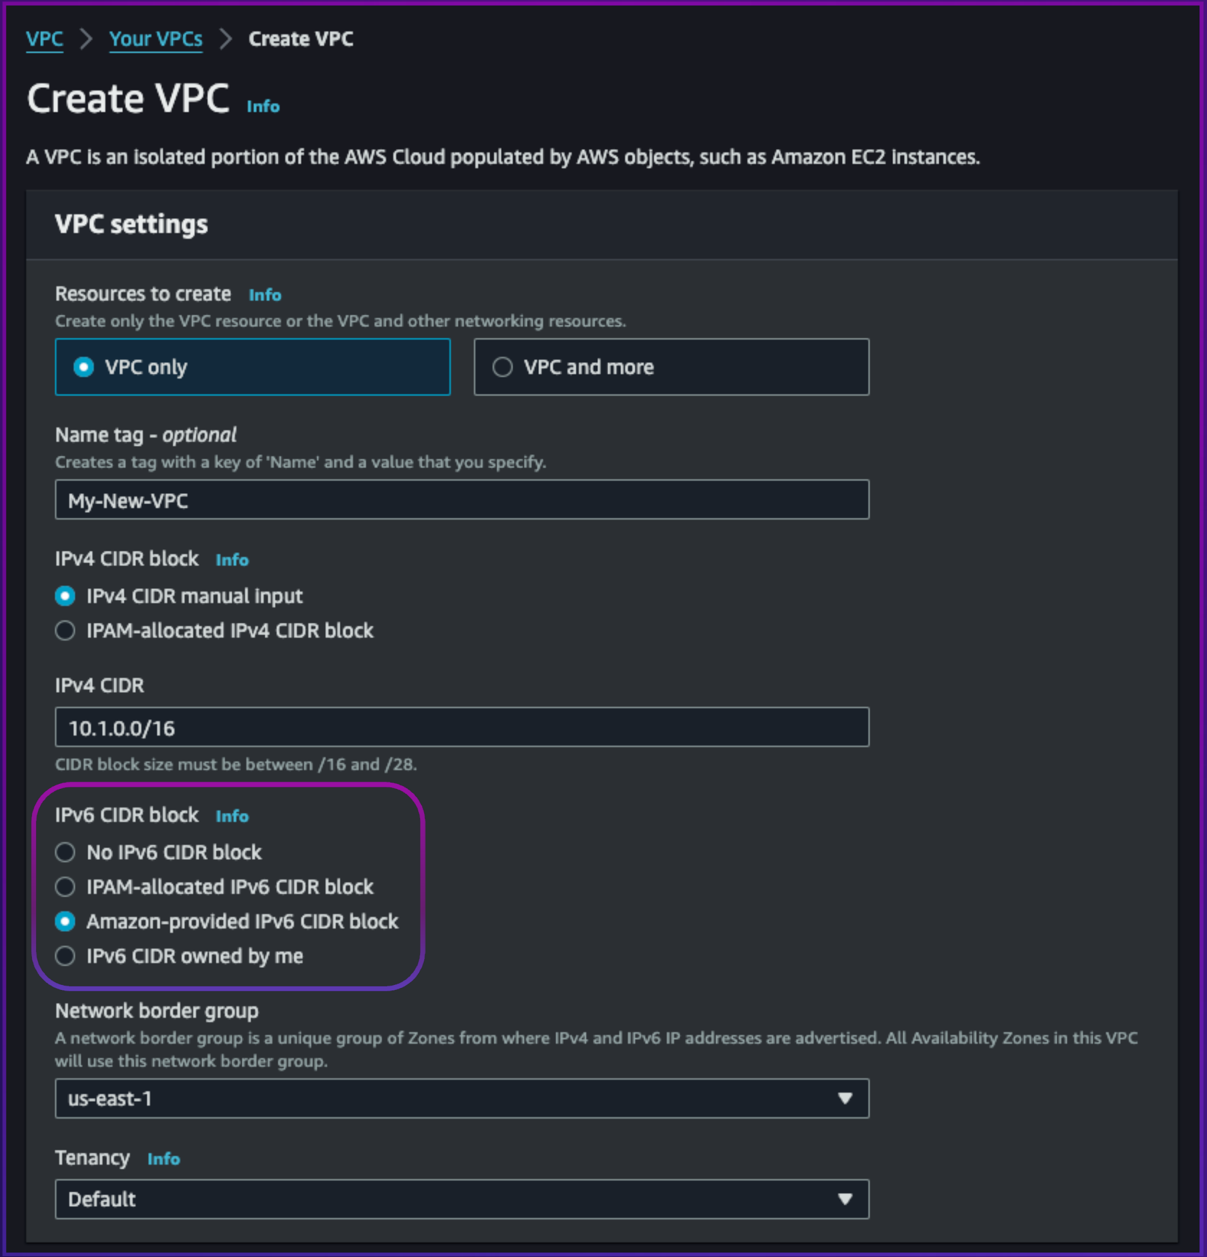

VPC > Your VPCs > Create VPC

Create VPC Info

A VPC is an isolated portion of the AWS Cloud populated by AWS objects, such as Amazon EC2 instances.

VPC settings

Resources to create Info
Create only the VPC resource or the VPC and other networking resources.

● VPC only ○ VPC and more

Name tag - *optional*
Creates a tag with a key of 'Name' and a value that you specify.

My-New-VPC

IPv4 CIDR block Info

● IPv4 CIDR manual input
○ IPAM-allocated IPv4 CIDR block

IPv4 CIDR

10.1.0.0/16

CIDR block size must be between /16 and /28.

IPv6 CIDR block Info

○ No IPv6 CIDR block
○ IPAM-allocated IPv6 CIDR block
● Amazon-provided IPv6 CIDR block
○ IPv6 CIDR owned by me

Network border group
A network border group is a unique group of Zones from where IPv4 and IPv6 IP addresses are advertised. All Availability Zones in this VPC will use this network border group.

us-east-1 ▼

Tenancy Info

Default ▼

Dual stack Amazon VPC

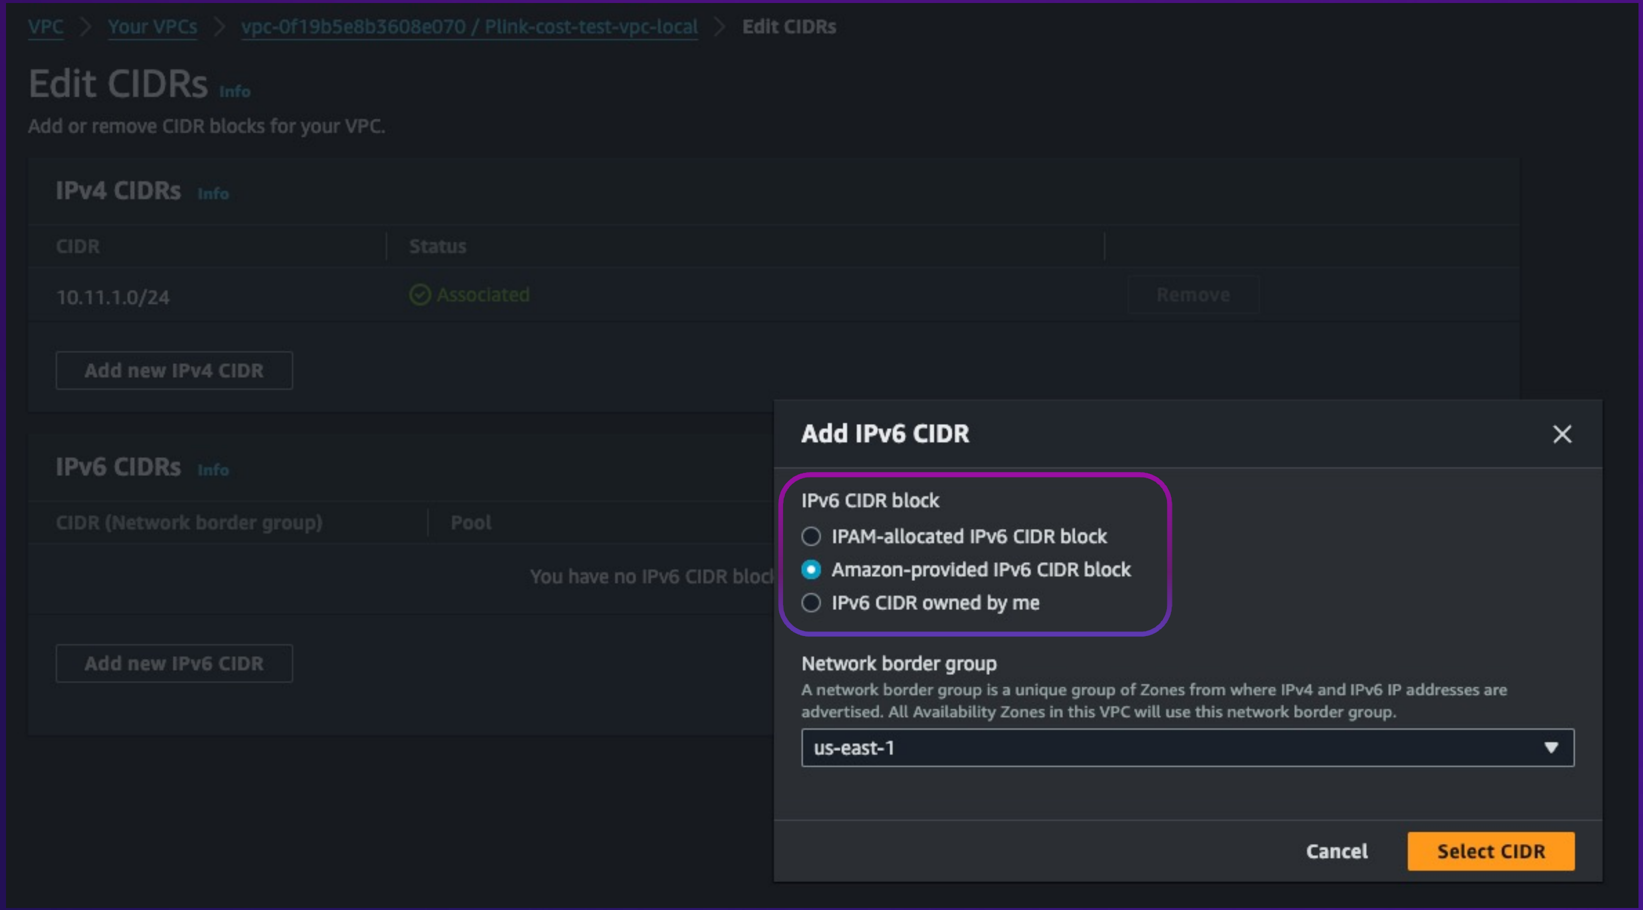

VPC > Your VPCs > vpc-0f19b5e8b3608e070 / Plink-cost-test-vpc-local > Edit CIDRs

Edit CIDRs

Add or remove CIDR blocks for your VPC.

IPv4 CIDRs

CIDR	Status	
10.11.1.0/24	Associated	Remove

IPv6 CIDRs

CIDR (Network border group)	Pool

You have no IPv6 CIDR bloc

Add IPv6 CIDR

IPv6 CIDR block

○ IPAM-allocated IPv6 CIDR block
● Amazon-provided IPv6 CIDR block
○ IPv6 CIDR owned by me

Network border group

A network border group is a unique group of Zones from where IPv4 and IPv6 IP addresses are advertised. All Availability Zones in this VPC will use this network border group.

Network border group: us-east-1

Dual stack Amazon VPC

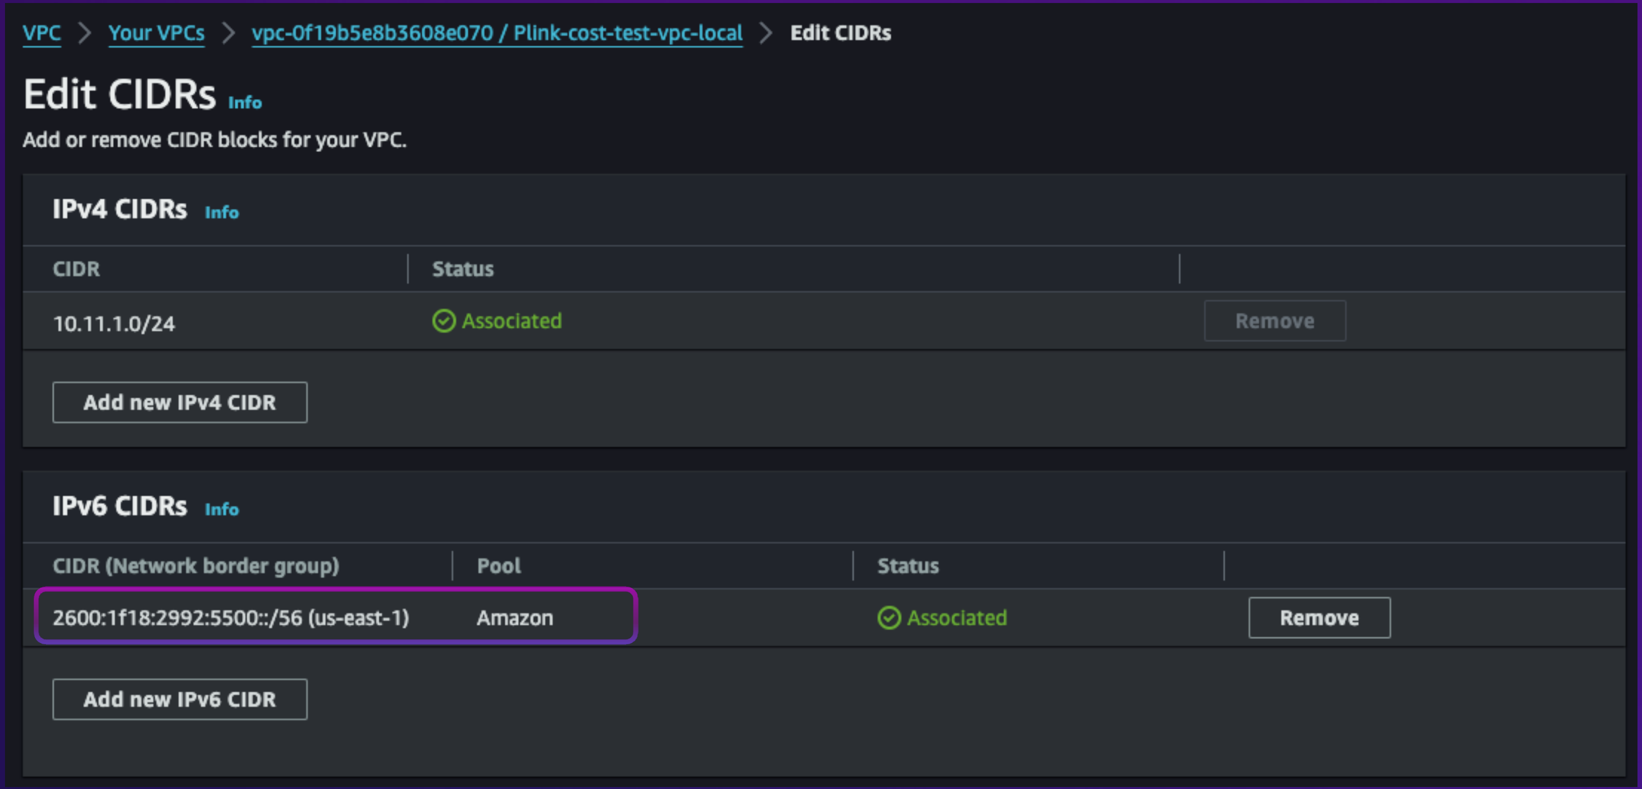

VPC > Your VPCs > vpc-0f19b5e8b3608e070 / Plink-cost-test-vpc-local > Edit CIDRs

Edit CIDRs Info

Add or remove CIDR blocks for your VPC.

IPv4 CIDRs Info

CIDR	Status	
10.11.1.0/24	✓ Associated	Remove

Add new IPv4 CIDR

IPv6 CIDRs Info

CIDR (Network border group)	Pool	Status	
2600:1f18:2992:5500::/56 (us-east-1)	Amazon	✓ Associated	Remove

Add new IPv6 CIDR

Dual stack Amazon VPC

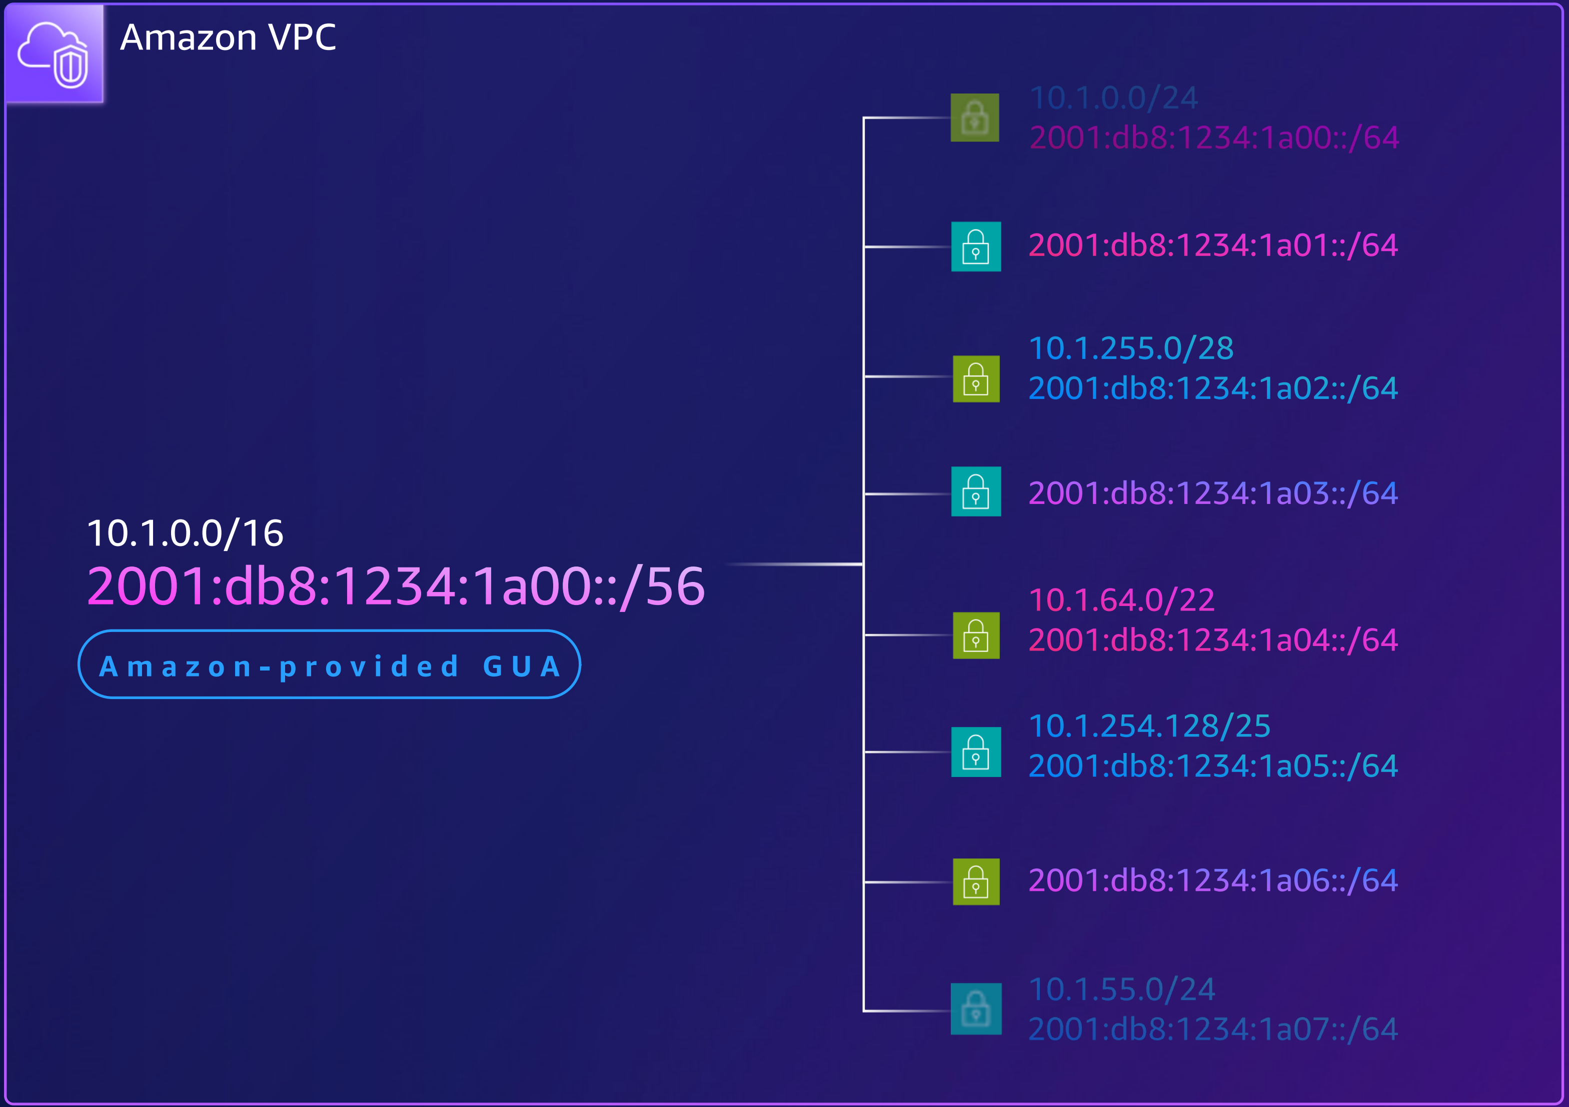

Dual stack Amazon VPC

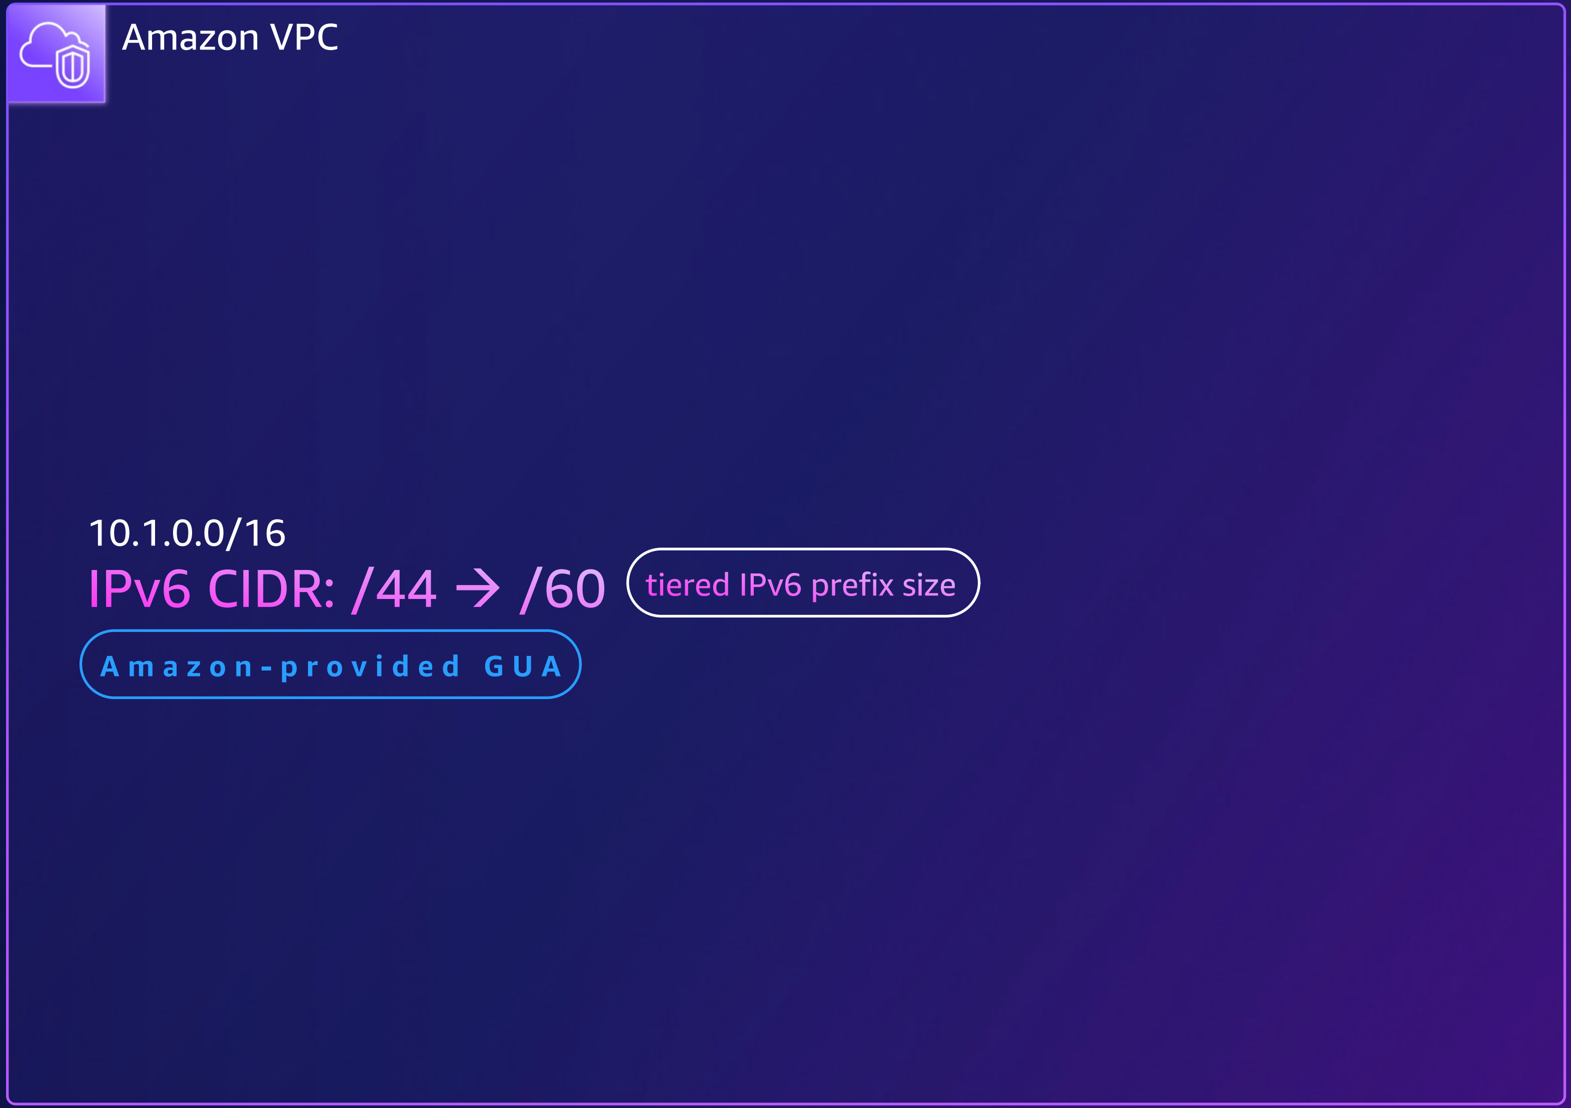

Amazon VPC

10.1.0.0/16

IPv6 CIDR: /44 → /60

tiered IPv6 prefix size

Amazon-provided GUA

Dual stack Amazon VPC

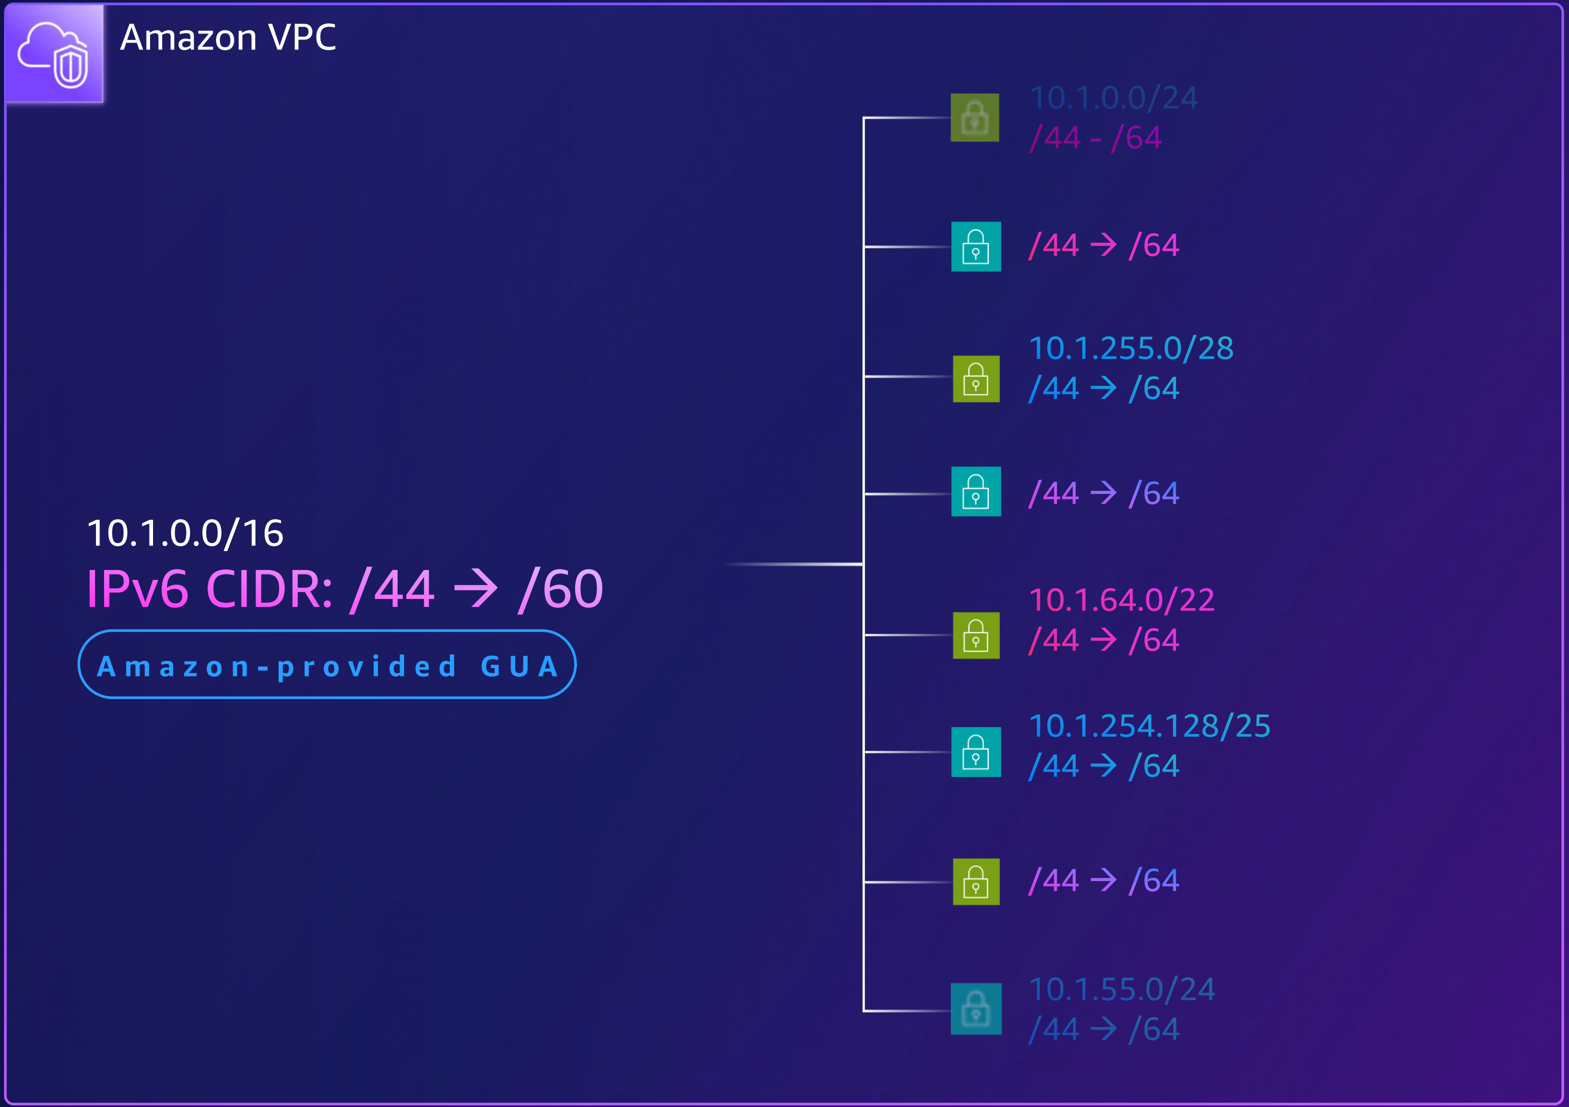

Amazon VPC

10.1.0.0/16

IPv6 CIDR: /44 → /60

Amazon-provided GUA

10.1.0.0/24
/44 - /64

/44 → /64

10.1.255.0/28
/44 → /64

/44 → /64

10.1.64.0/22
/44 → /64

10.1.254.128/25
/44 → /64

/44 → /64

10.1.55.0/24
/44 → /64

Dual stack Amazon VPC

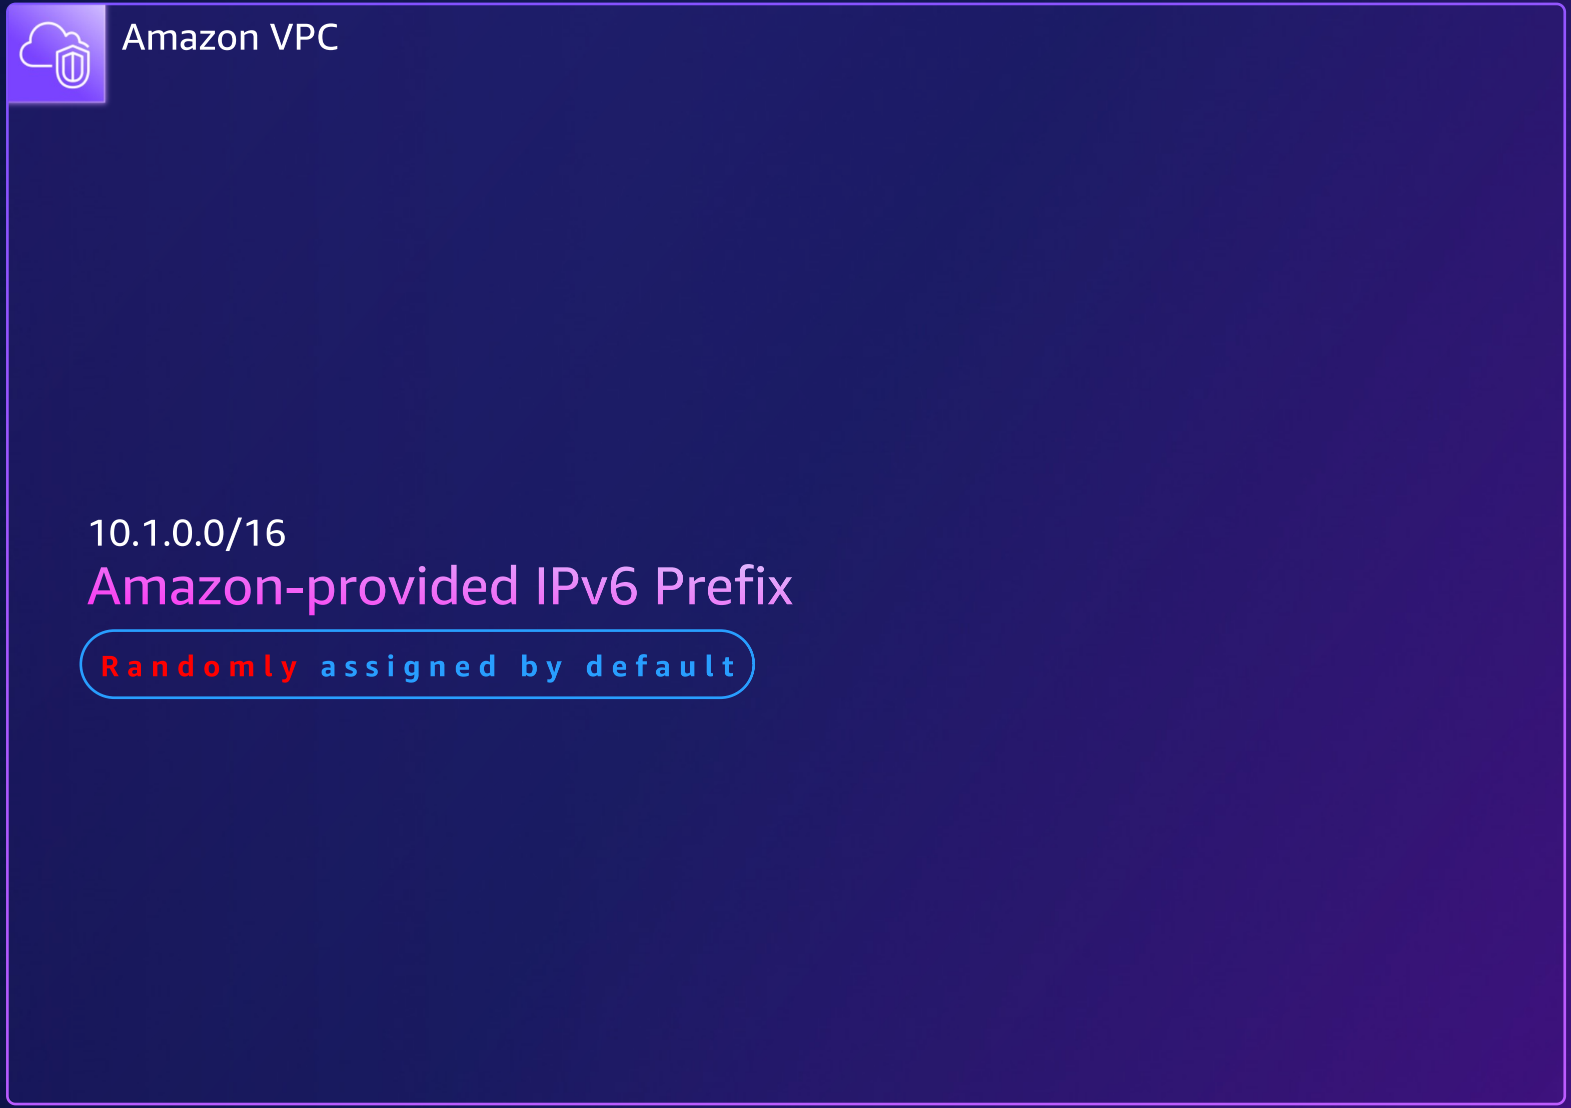

Amazon VPC

10.1.0.0/16

Amazon-provided IPv6 Prefix

Randomly assigned by default

IPv6 addressing plan

Dual stack Amazon VPC

Amazon-provided GUA (VPC-level)

Amazon-provided contiguous IPv6 GUA prefixes

NEW

Amazon VPC IP Address Manager

Amazon-provided contiguous IPv6 prefixes

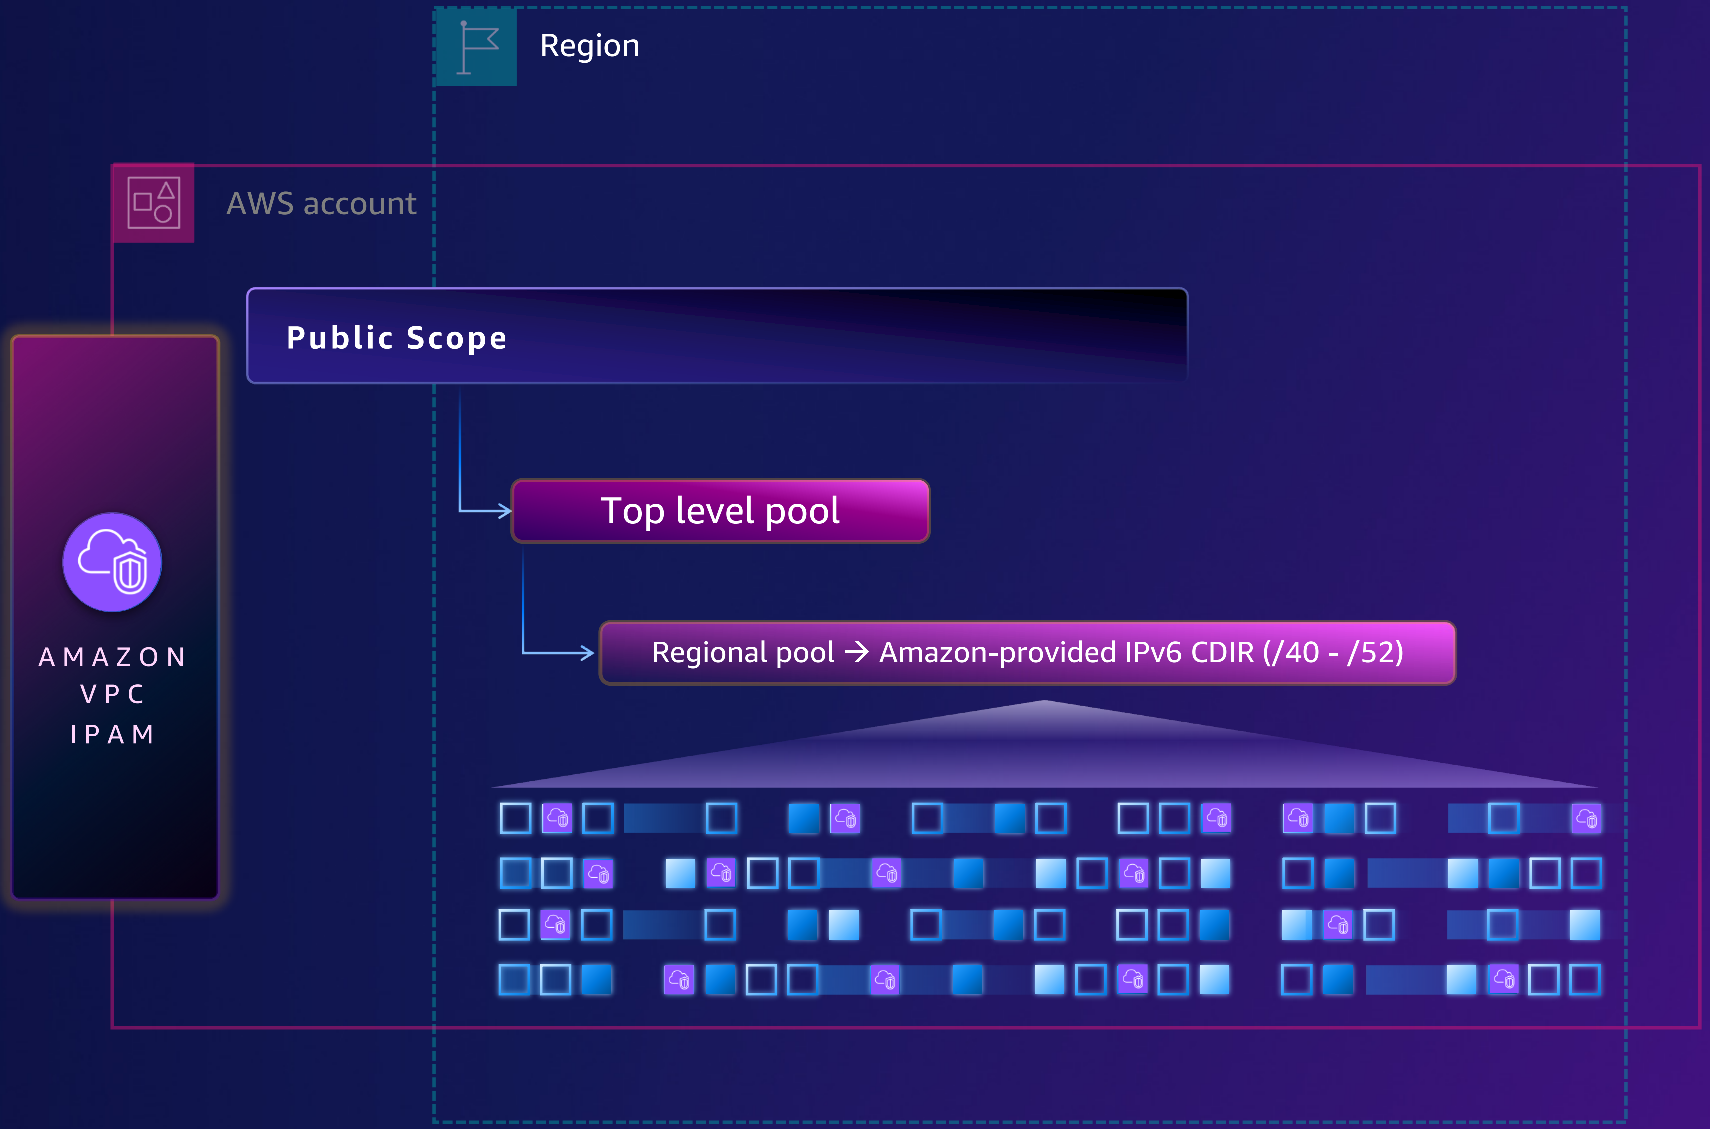

Amazon VPC IP Address Manager

Free tier

For IP management in a single AWS Region and account

Amazon-provided contiguous IPv6 blocks per Region and account

Advanced tier

For IP management across two or more AWS Regions and accounts

Amazon-provided contiguous IPv6 blocks across multiple Regions and accounts

Dual stack Amazon VPC

IPv6 addressing plan

- Amazon-provided GUA (VPC-level)
- Amazon-provided contiguous IPv6 GUA prefixes
- Bring your own IPv6 (BYOIPv6) GUA prefixes

IPv6 addressing plan

BYOIPv6

In Amazon EC2

You can bring each address range to one AWS Region at a time

You cannot share your IP address range with other accounts

You can control if CIDRs in a pool can be publicly advertisable or not

The most specific IPv6 address range that you can bring is **/48** for CIDRs that are publicly advertisable and **/56** for CIDRs that are not publicly advertisable

With VPC IPAM

You can bring each address range to an IPAM top level Pool, and further split it across multiple Regional pools

You can share your IP address range with other accounts

You can control if CIDRs in a pool can be publicly advertisable or not

The most specific IPv6 address range that you can bring is **/48** for CIDRs that are publicly advertisable and **/60** for CIDRs that are not publicly advertisable

IPv6 addressing plan

BYOIPv6

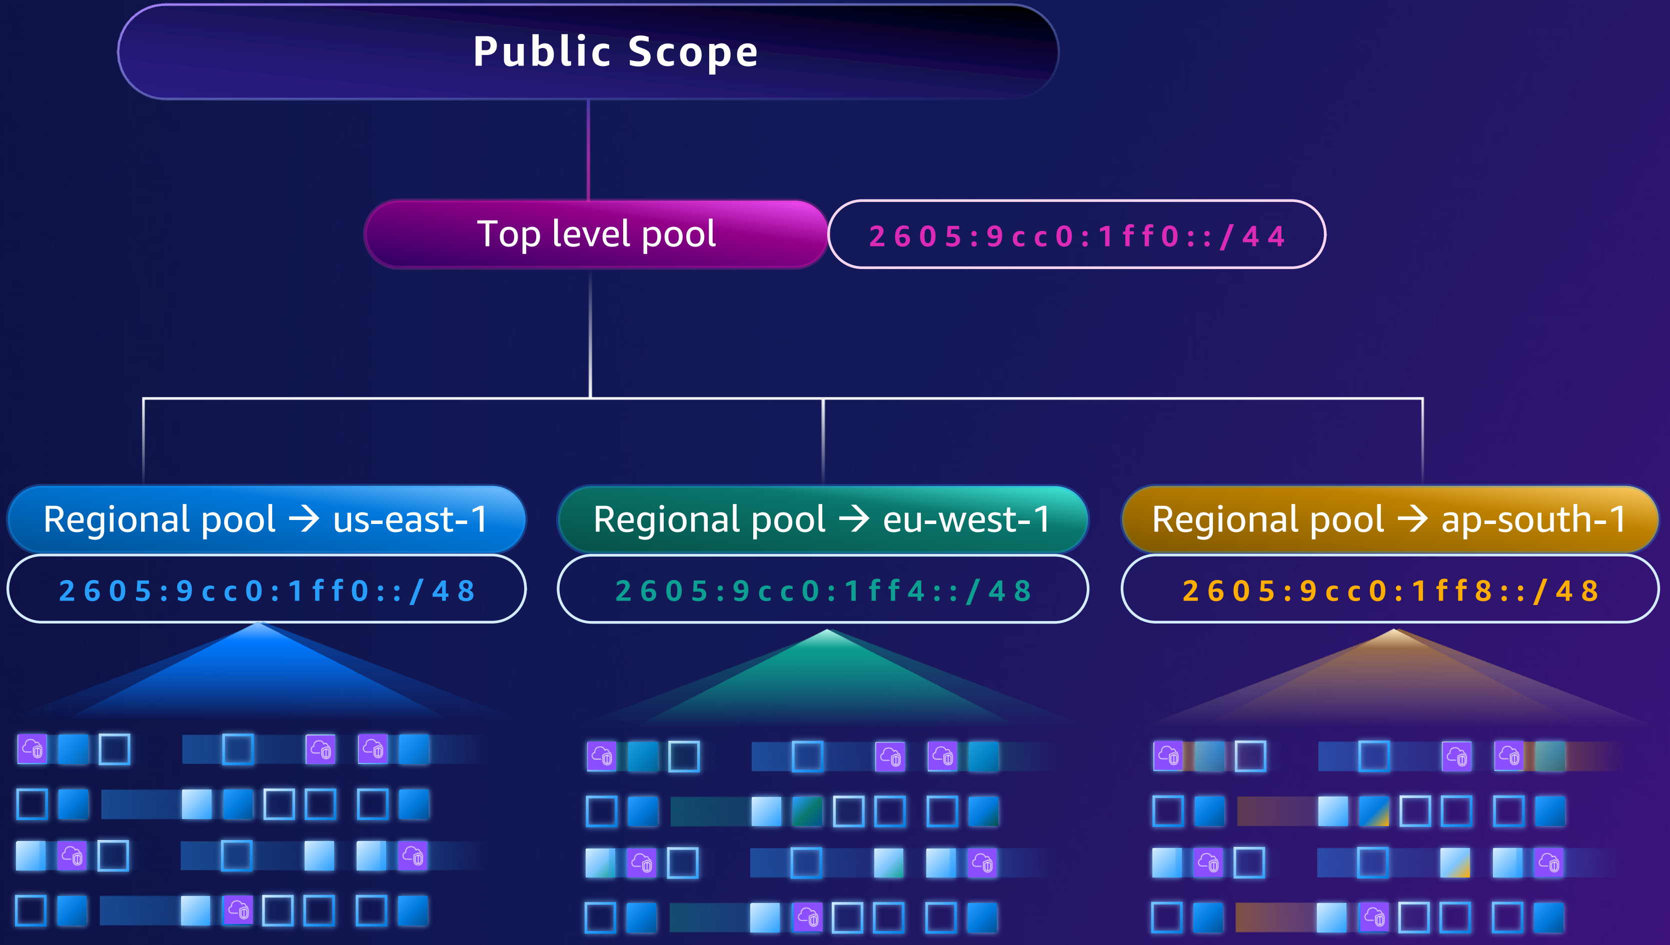

Note: The prefix used for this presentation is an Amazon-owned IPv6 prefix, and prefix sizes are examples.

IPv6 address planning summary

	Provisioning	Globally Unique	Internet advertisement	Internet Connectivity	NAT66 / NPTv6	Summarization capabilities	Considerations
Amazon-provided IPv6 GUA (VPC-level)	Directly at the VPC level	Yes	AWS advertised	Native on AWS	Not Required	No	Not recommended for large scale deployments (many VPCs)
Amazon-provided contiguous IPv6 prefixes	Amazon VPC IPAM free or advanced tiers	Yes	AWS advertised	Native on AWS	Not Required	Yes, for all VPCs created from the same IPAM Pool	Facilitates growth on AWS, doesn't require you to own IPv6 addresses
BYOIPv6	Amazon EC2 or Amazon VPC IPAM	Yes	Configurable	Native on AWS if advertised from AWS	Not Required	Yes, for all VPCs created from the same BYOIP pool	Facilitates growth on AWS, requires you to own IPv6 addresses, and prove ownership through the BYOIPv6 process.
				On-premises if advertised from on-premises	Not Required		

IPv6 design

Dual stack Amazon VPC

Dual stack Amazon VPC

IPv6 design

Dual stack VPC

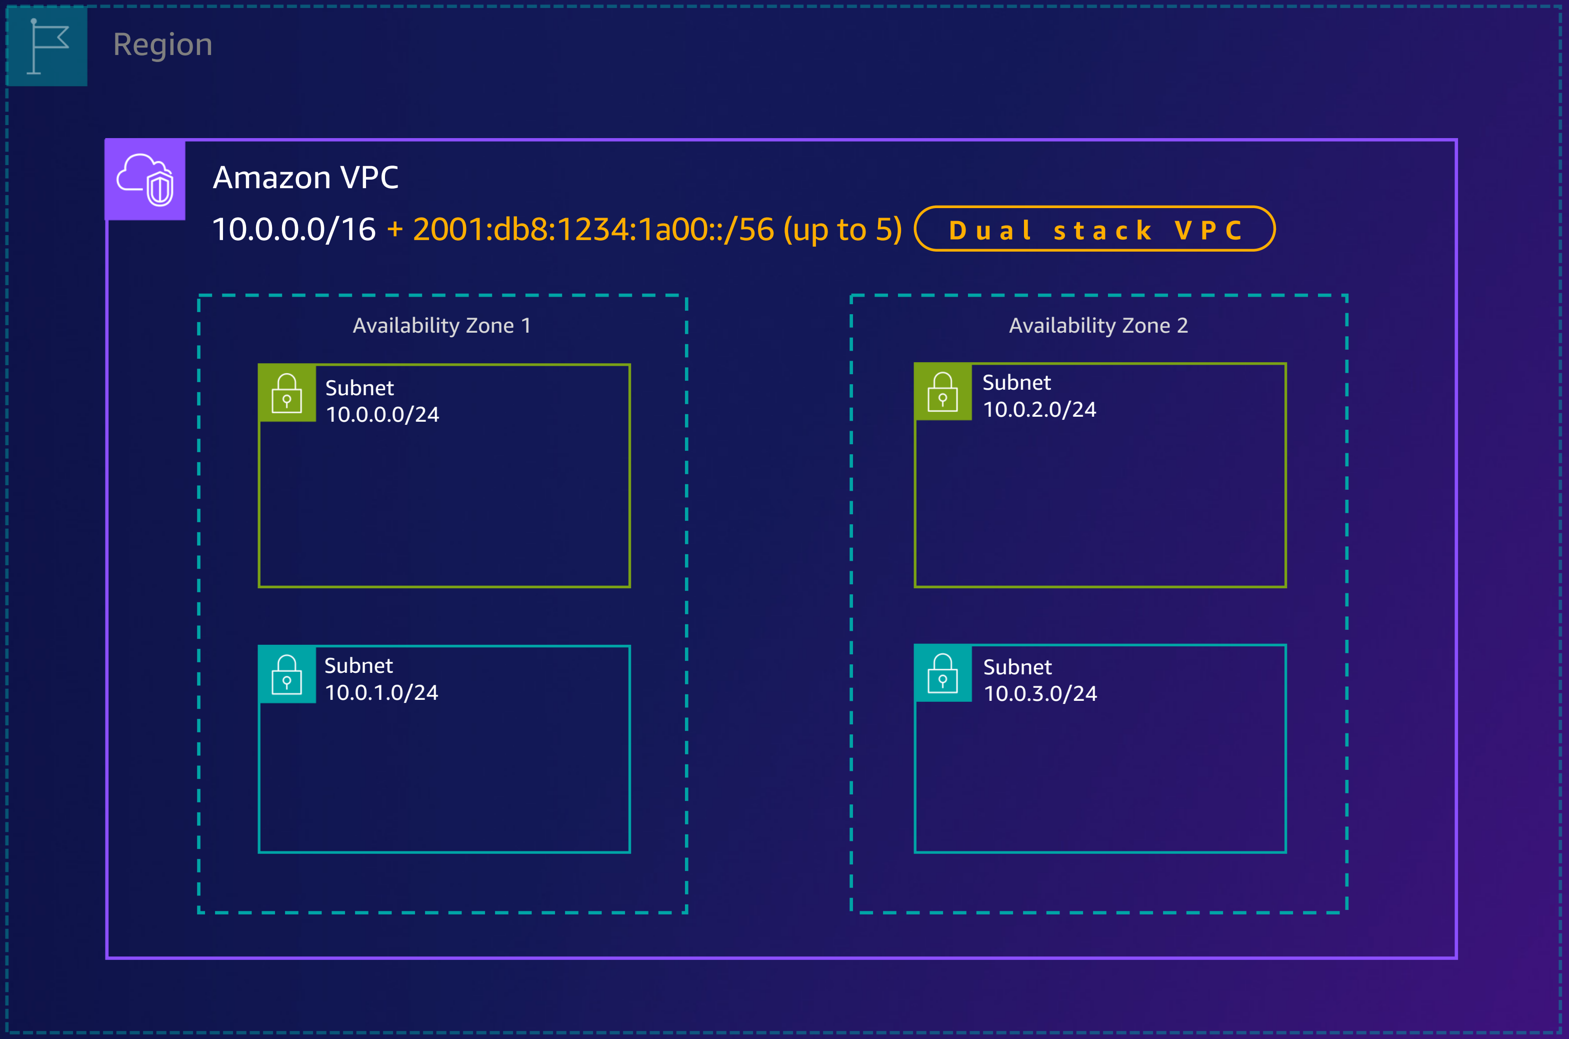

Dual stack Amazon VPC

IPv6 design

- Dual stack VPC
- VPC routing

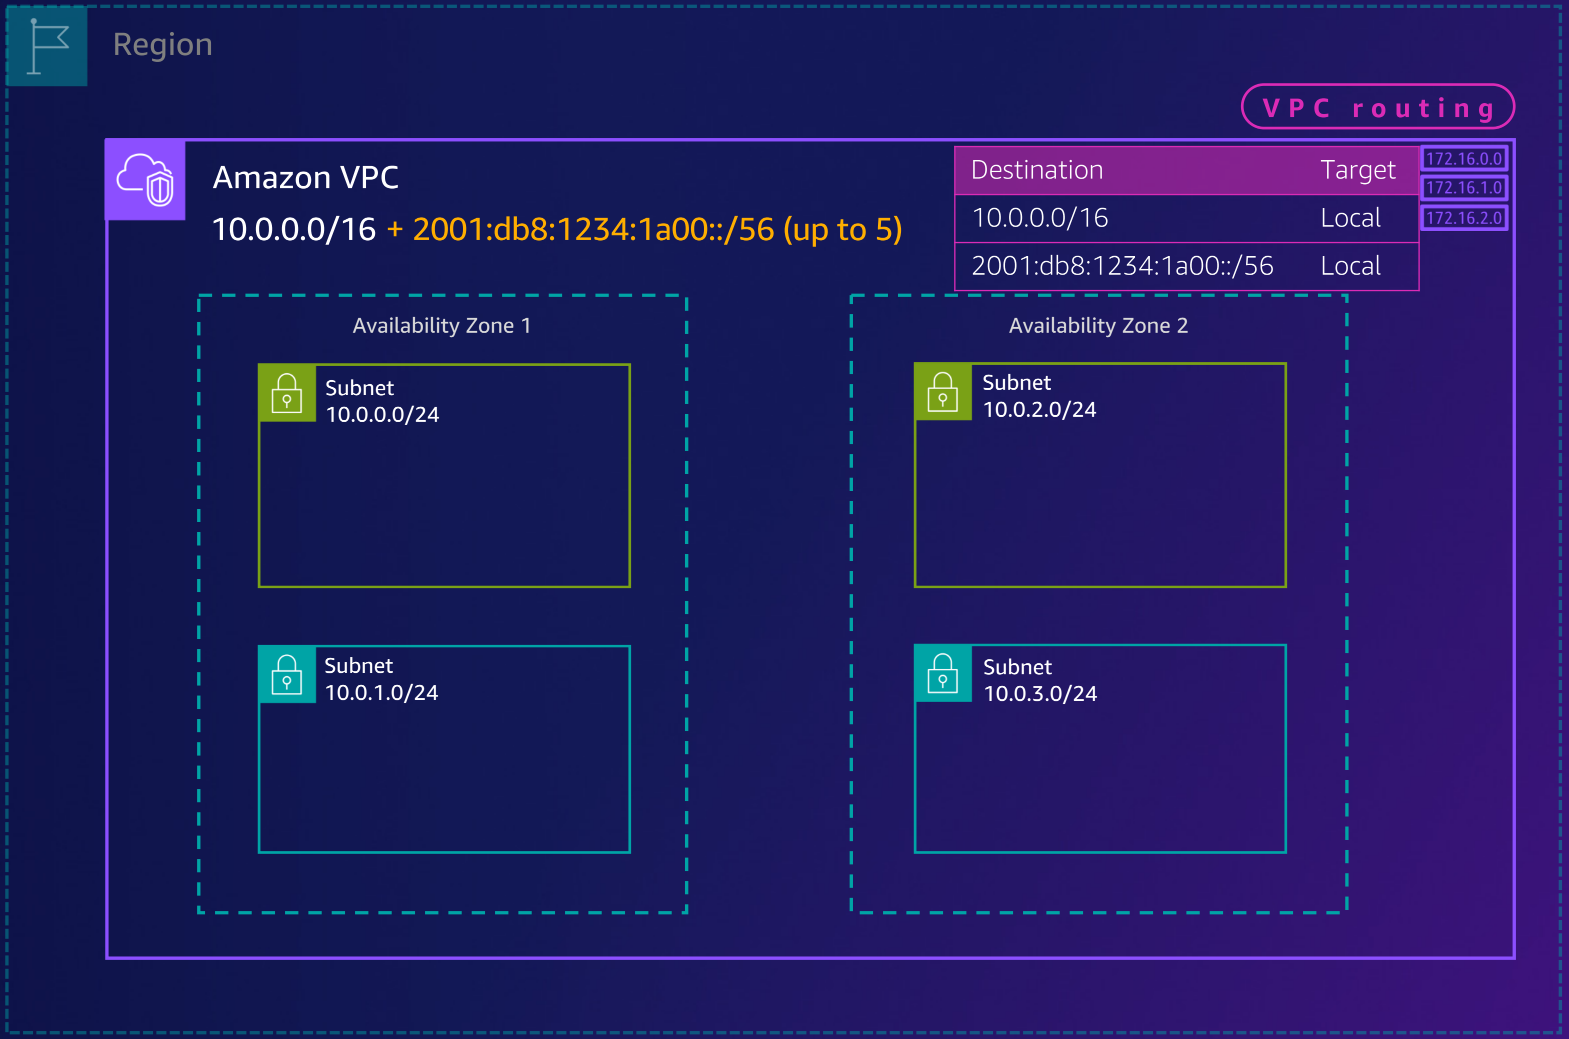

VPC routing

Region

Amazon VPC

10.0.0.0/16 + 2001:db8:1234:1a00::/56 (up to 5)

Destination	Target
10.0.0.0/16	Local
2001:db8:1234:1a00::/56	Local

Dual stack Amazon VPC

IPv6 design

Dual stack VPC

VPC routing

VPC DNS

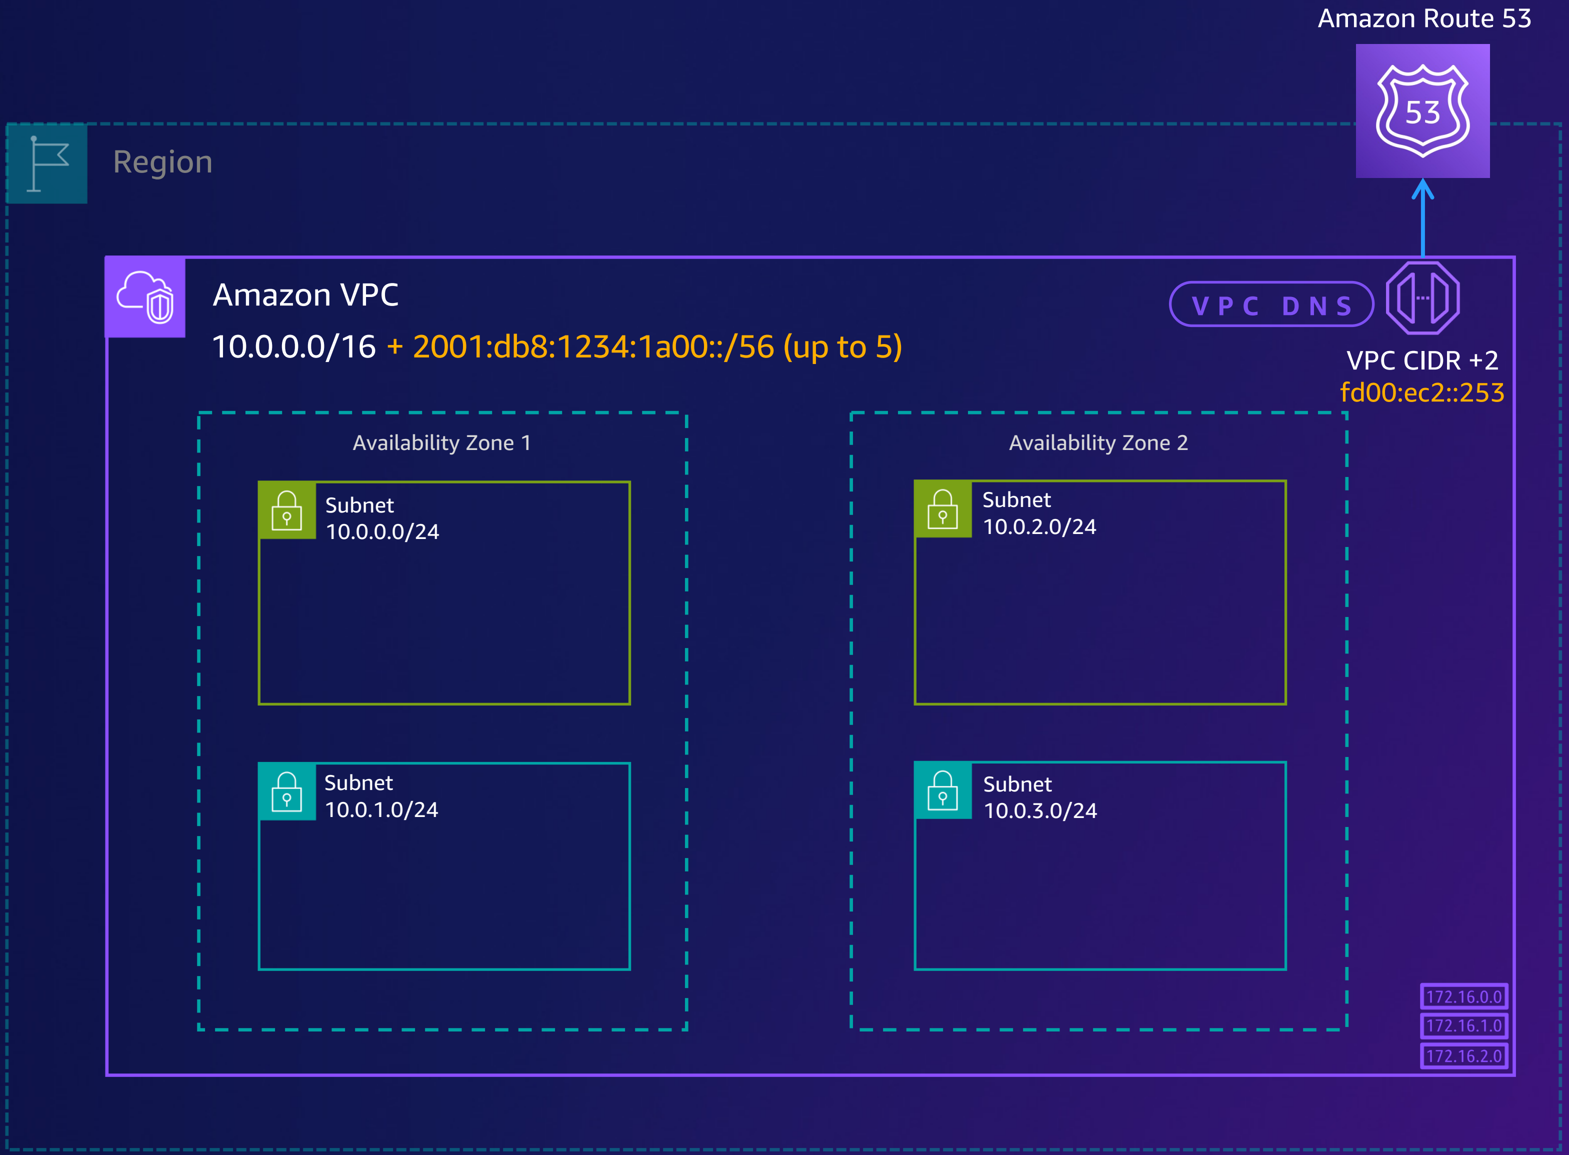

Amazon Route 53

Region

Amazon VPC

10.0.0.0/16 + 2001:db8:1234:1a00::/56 (up to 5)

VPC DNS

VPC CIDR +2

fd00:ec2::253

Availability Zone 1

Subnet 10.0.0.0/24

Subnet 10.0.1.0/24

Availability Zone 2

Subnet 10.0.2.0/24

Subnet 10.0.3.0/24

Dual stack Amazon VPC

IPv6 design

Dual stack VPC

VPC routing

VPC DNS

VPC Subnets

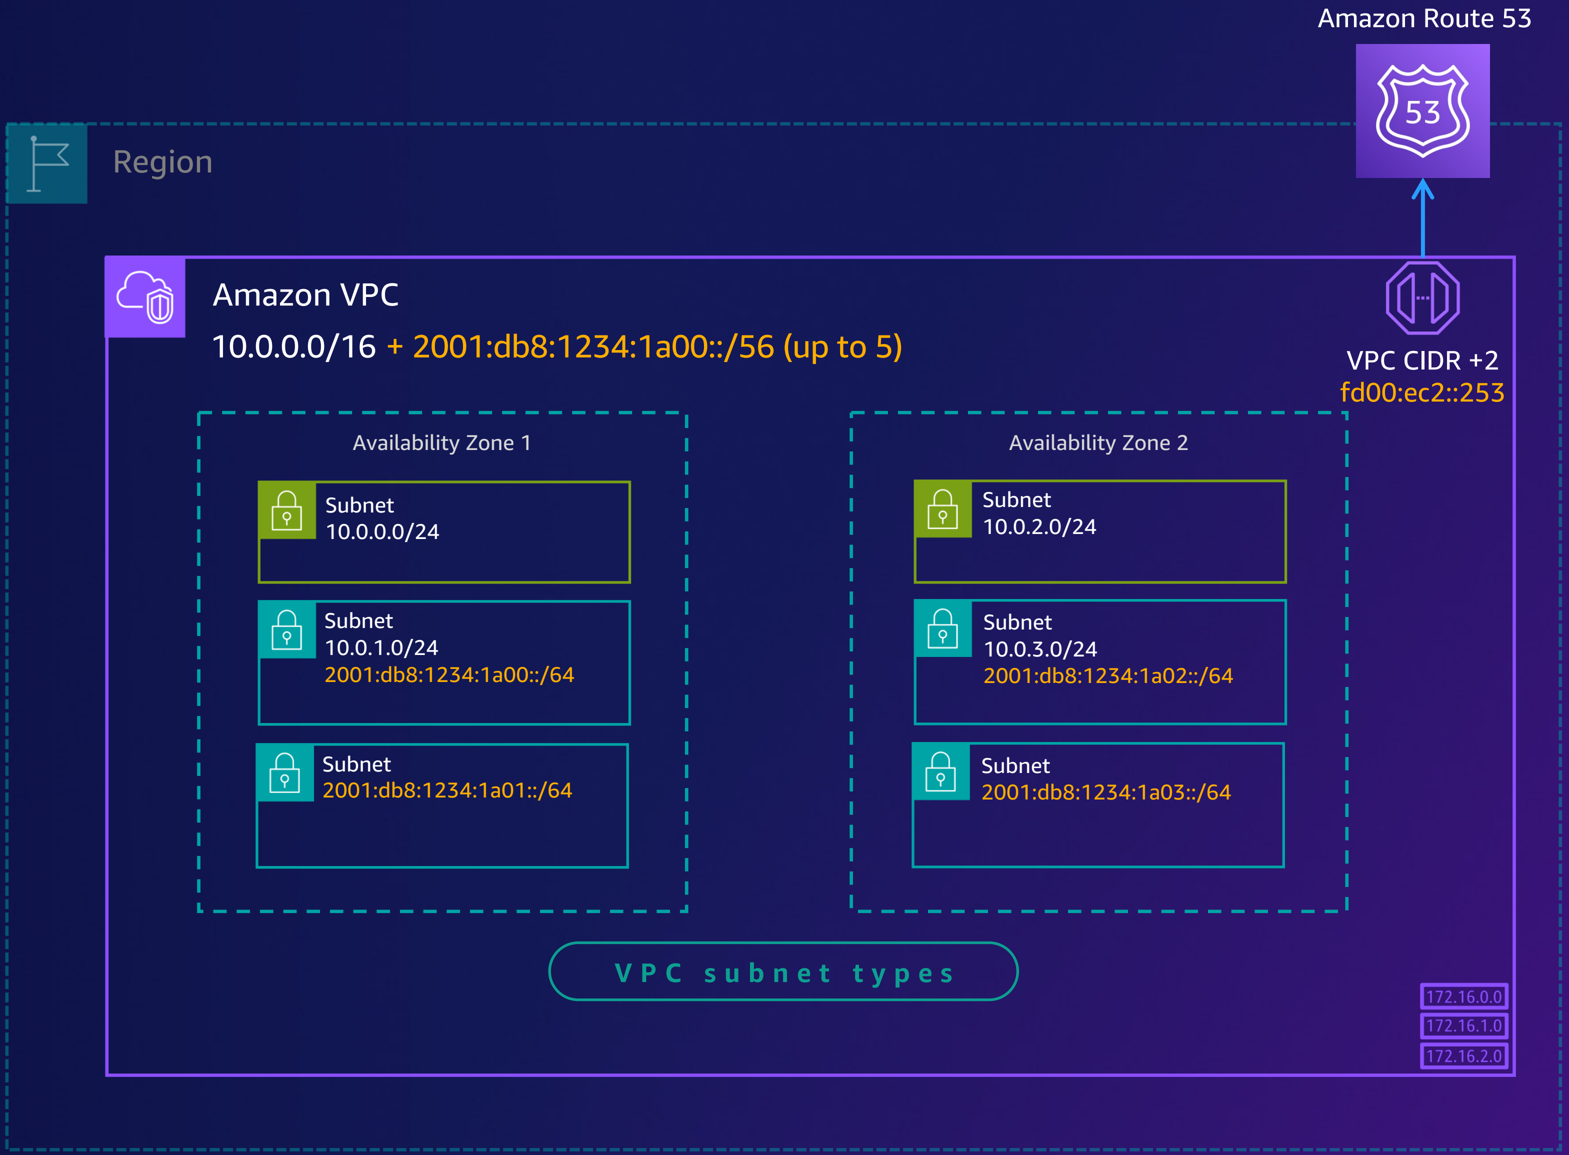

IPv6 support for

Amazon compute services

IPv6 support for Amazon Compute Services[1]

- Amazon Elastic Compute Cloud (EC2)
- Amazon Elastic Kubernetes Service (EKS)
- Amazon Elastic Container Service (ECS)
- AWS Lambda
- Amazon LightSail

[1] more services at: https://docs.aws.amazon.com/vpc/latest/userguide/aws-ipv6-support.html

IPv6 support for Amazon Compute Services

Amazon EC2

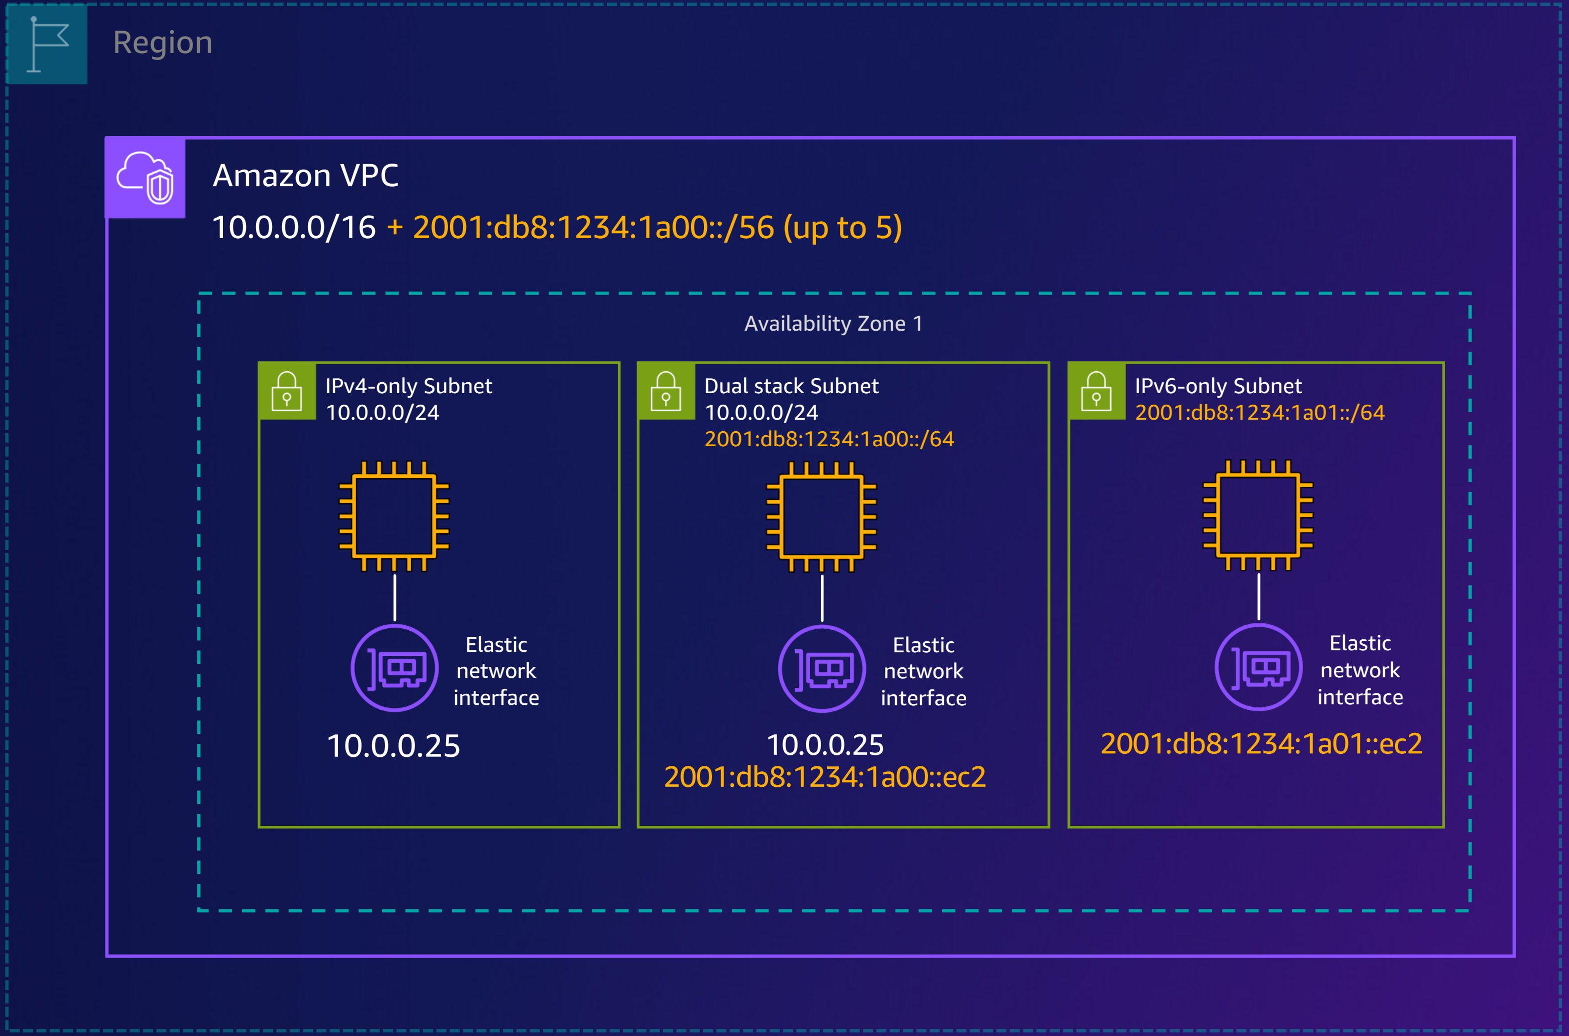

IPv6 support for Amazon Compute Services

Amazon EC2

NEW FOR NITRO INSTANCES

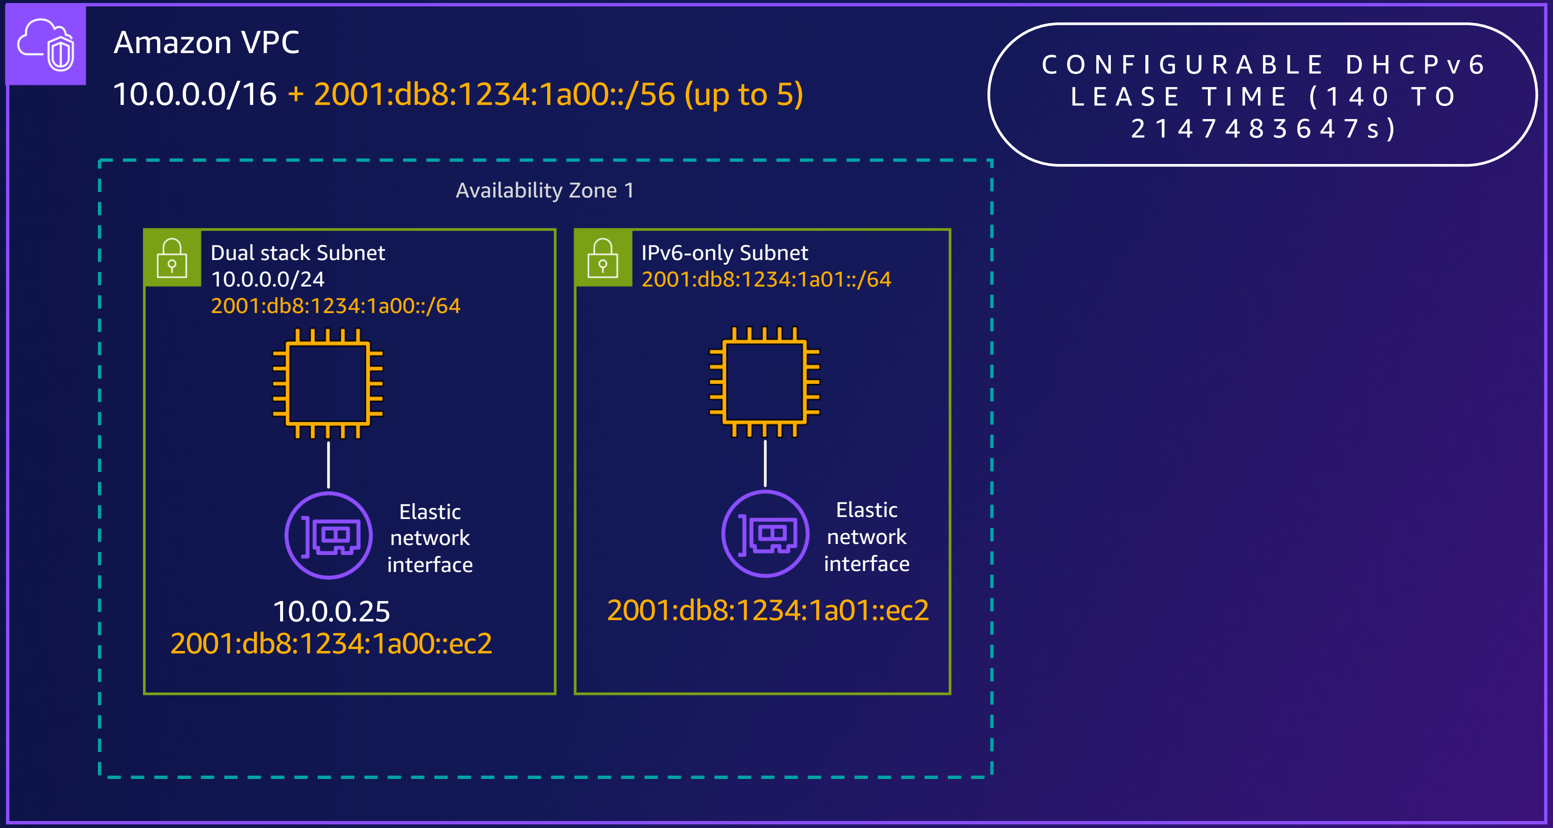

CONFIGURABLE DHCPv6 LEASE TIME (140 TO 2147483647s)

IPv6 support for Amazon Compute Services

Amazon EKS

Dual stack ingress Load Balancer Controller integration

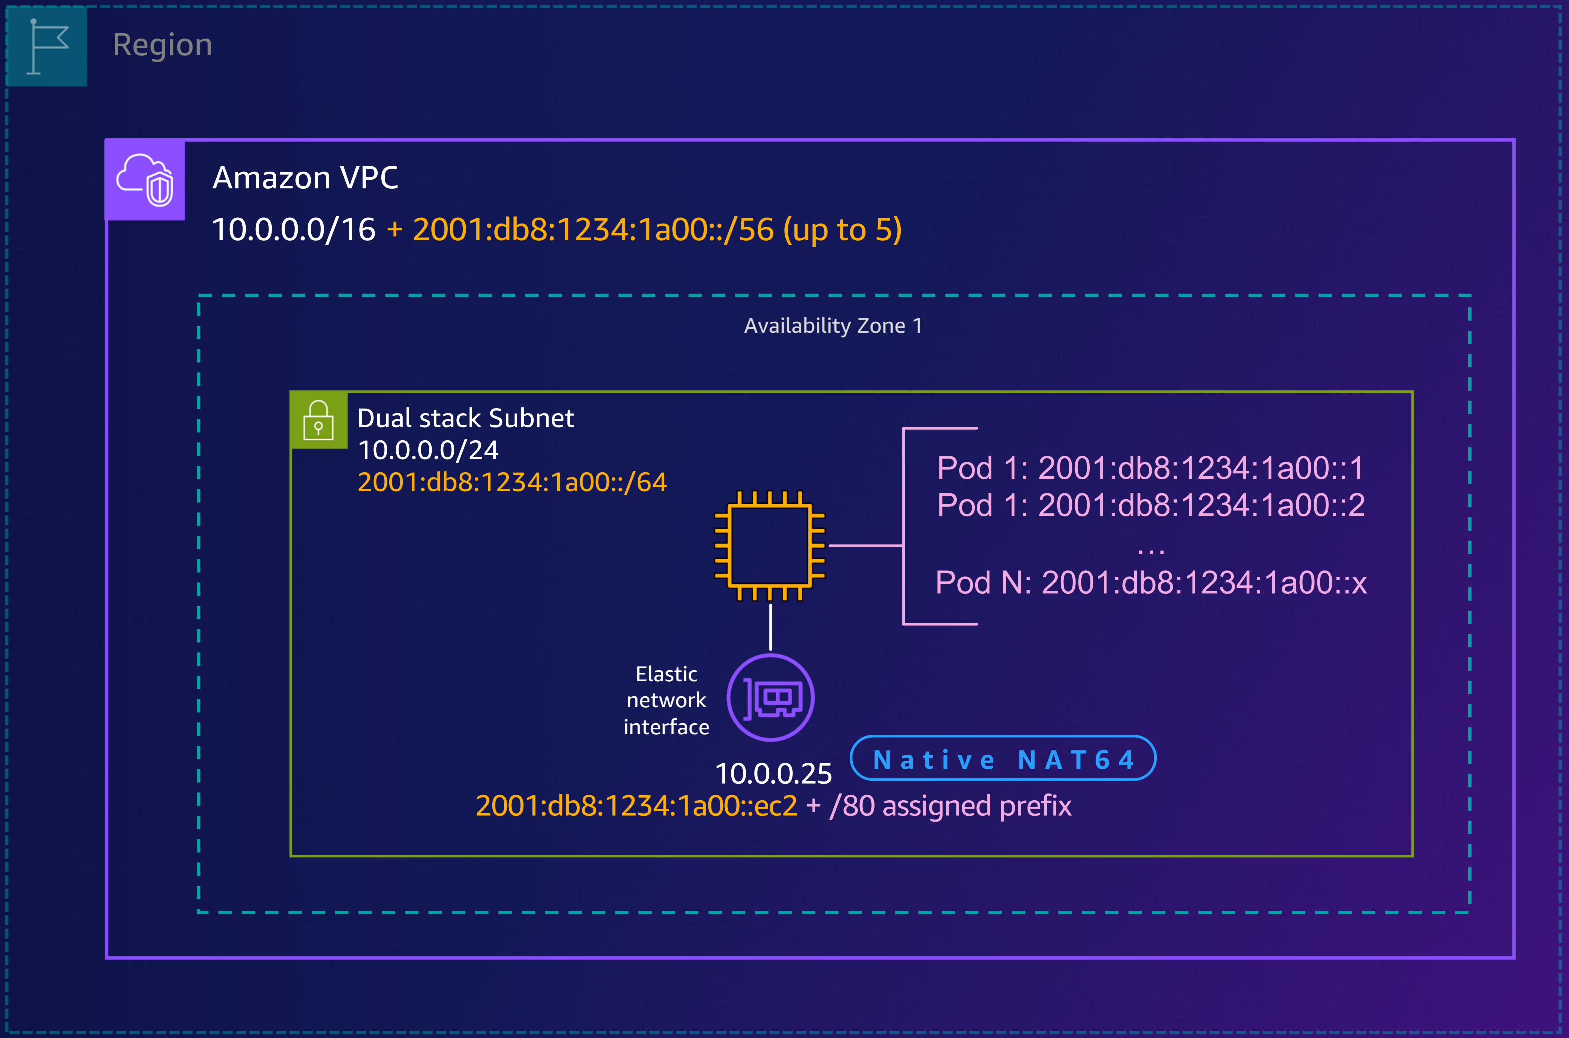

Region
Amazon VPC
10.0.0.0/16 + 2001:db8:1234:1a00::/56 (up to 5)
Availability Zone 1
Dual stack Subnet
10.0.0.0/24
2001:db8:1234:1a00::/64
Pod 1: 2001:db8:1234:1a00::1
Pod 1: 2001:db8:1234:1a00::2
...
Pod N: 2001:db8:1234:1a00::x
Elastic network interface
Native NAT64
10.0.0.25
2001:db8:1234:1a00::ec2 + /80 assigned prefix

IPv6 support for Amazon Compute Services

Amazon ECS

Supported in AWSVPC mode for both EC2 and Fargate

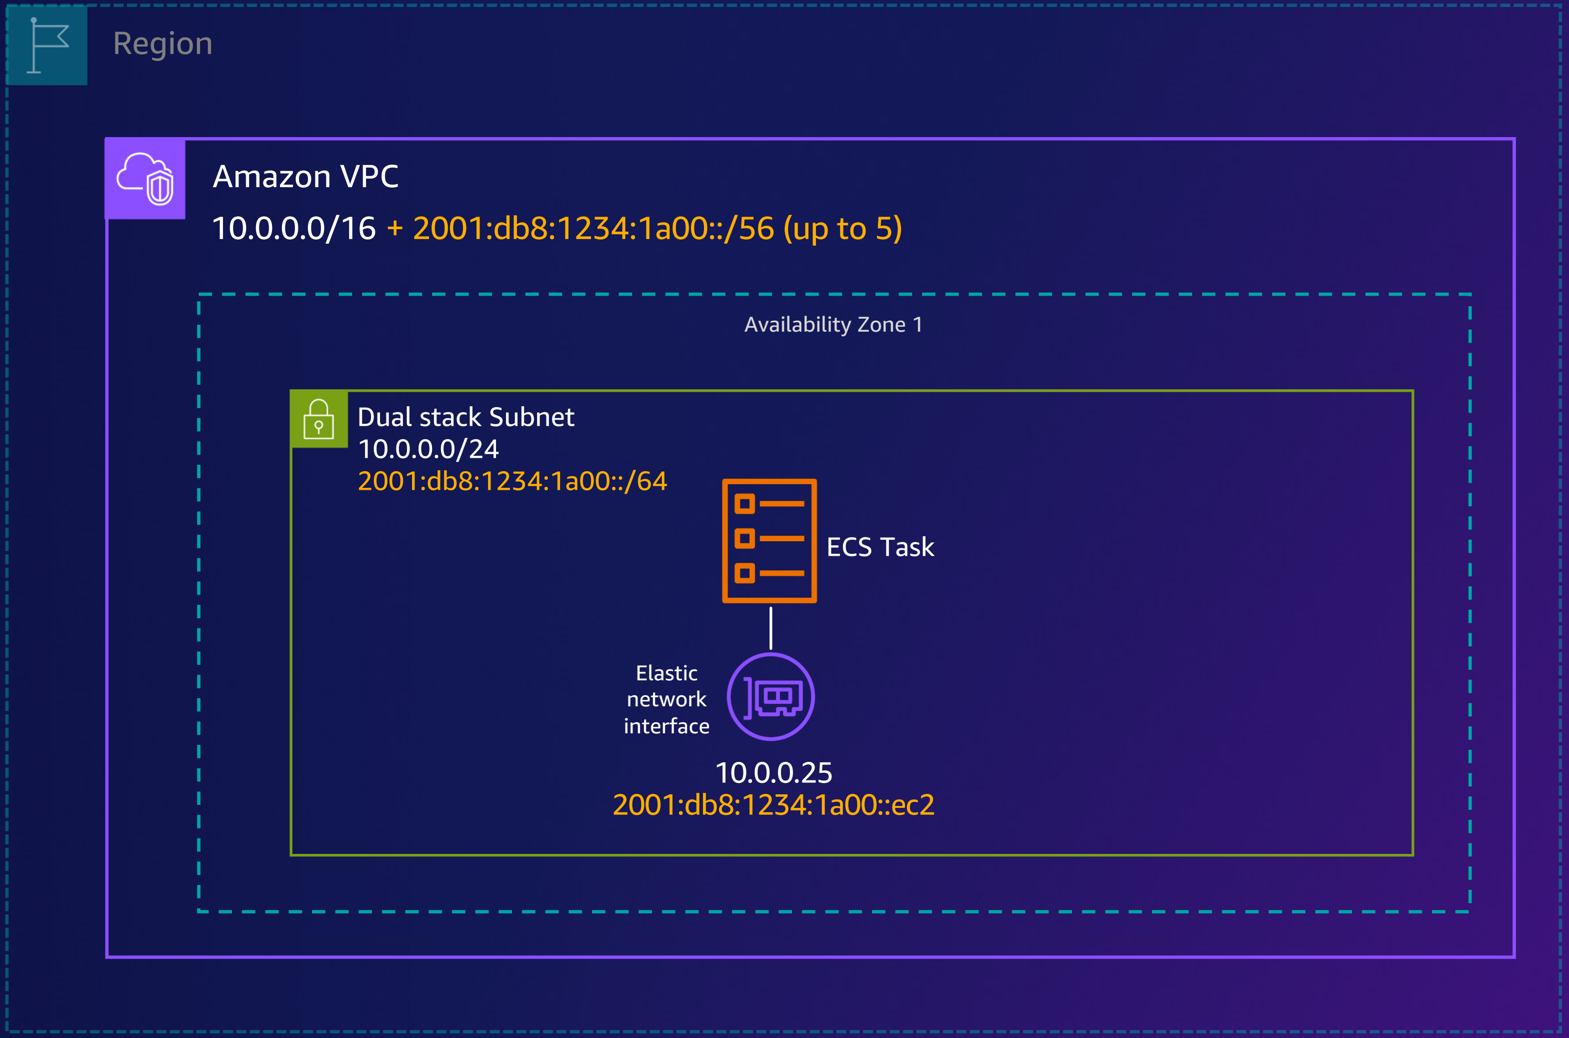

IPv6 support for Amazon Compute Services

AWS Lambda

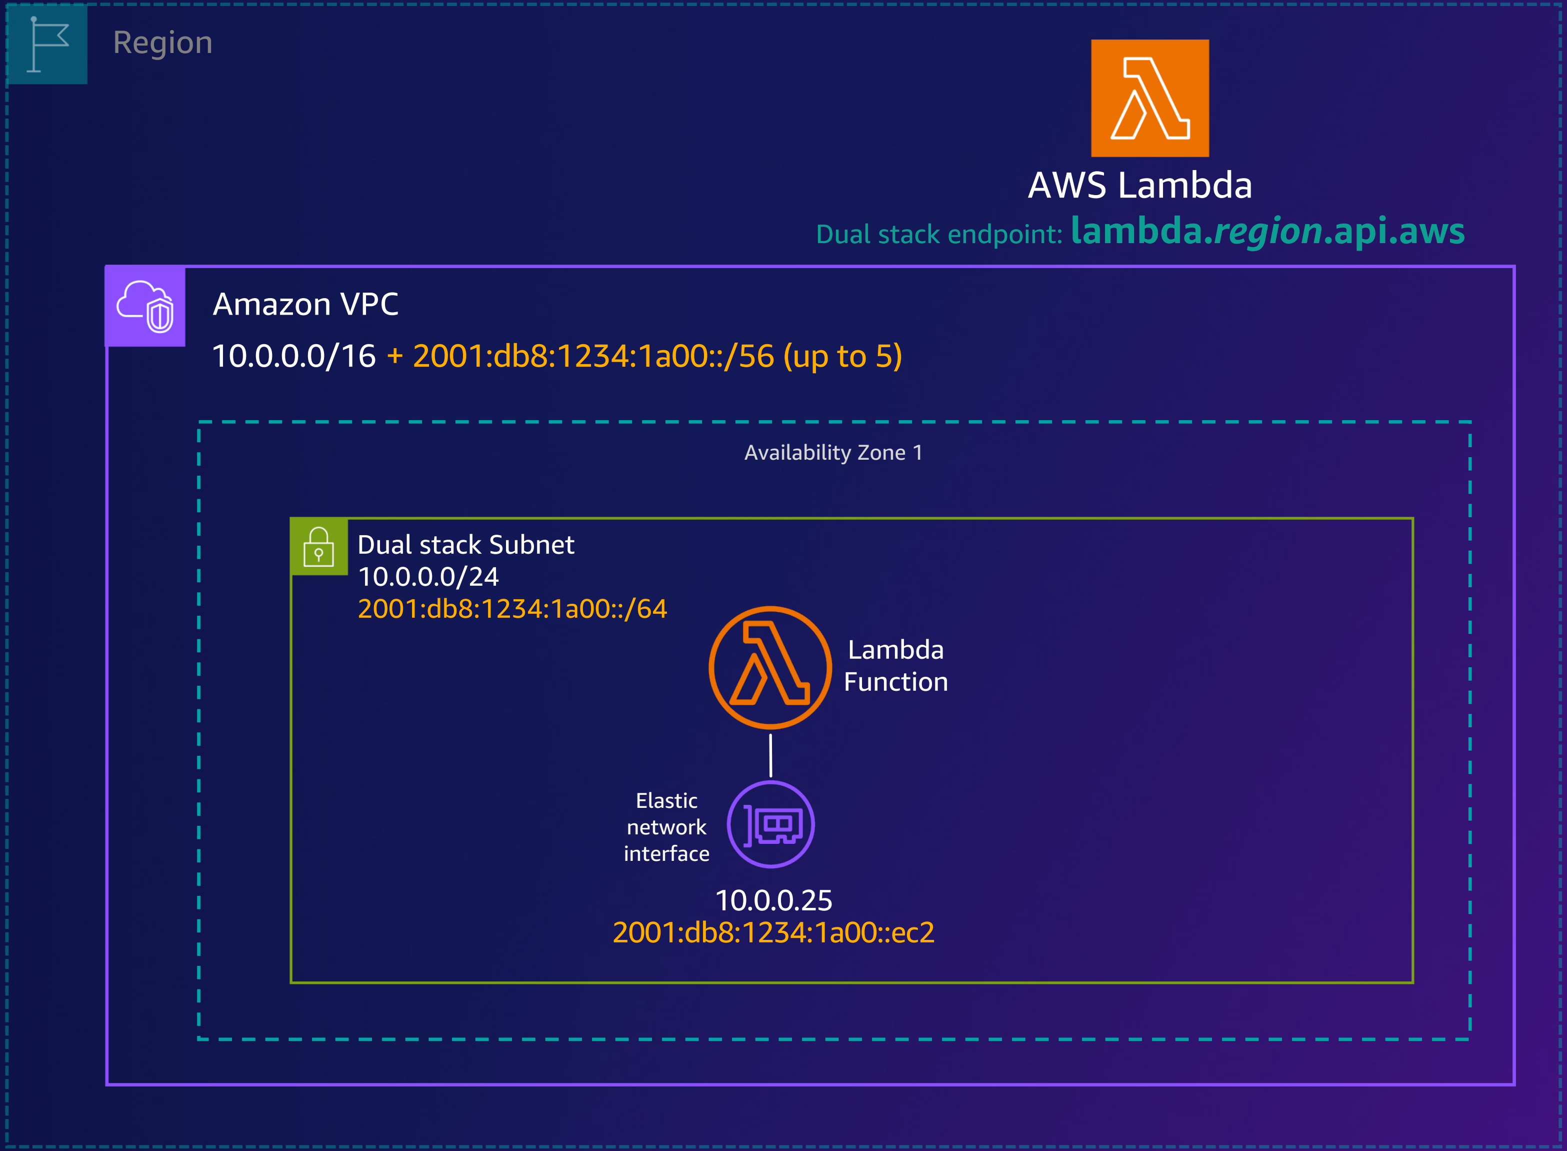

IPv6 support for Amazon Compute Services

Amazon Lightsail

Region
Amazon Lightsail

The easiest way to get started with Amazon Web Services (AWS) - Build applications and websites quickly, with bundled pricing and pre-configured cloud resources

Includes everything you need to launch your project quickly:

- instances (virtual private servers),
- container services,
- managed databases,
- content delivery network (CDN) distributions,
- load balancers,
- SSD-based block storage,
- static IP addresses,
- DNS management of registered domains, and resource snapshots (backups)

IPv6 support for Amazon Compute Services

Amazon Lightsail

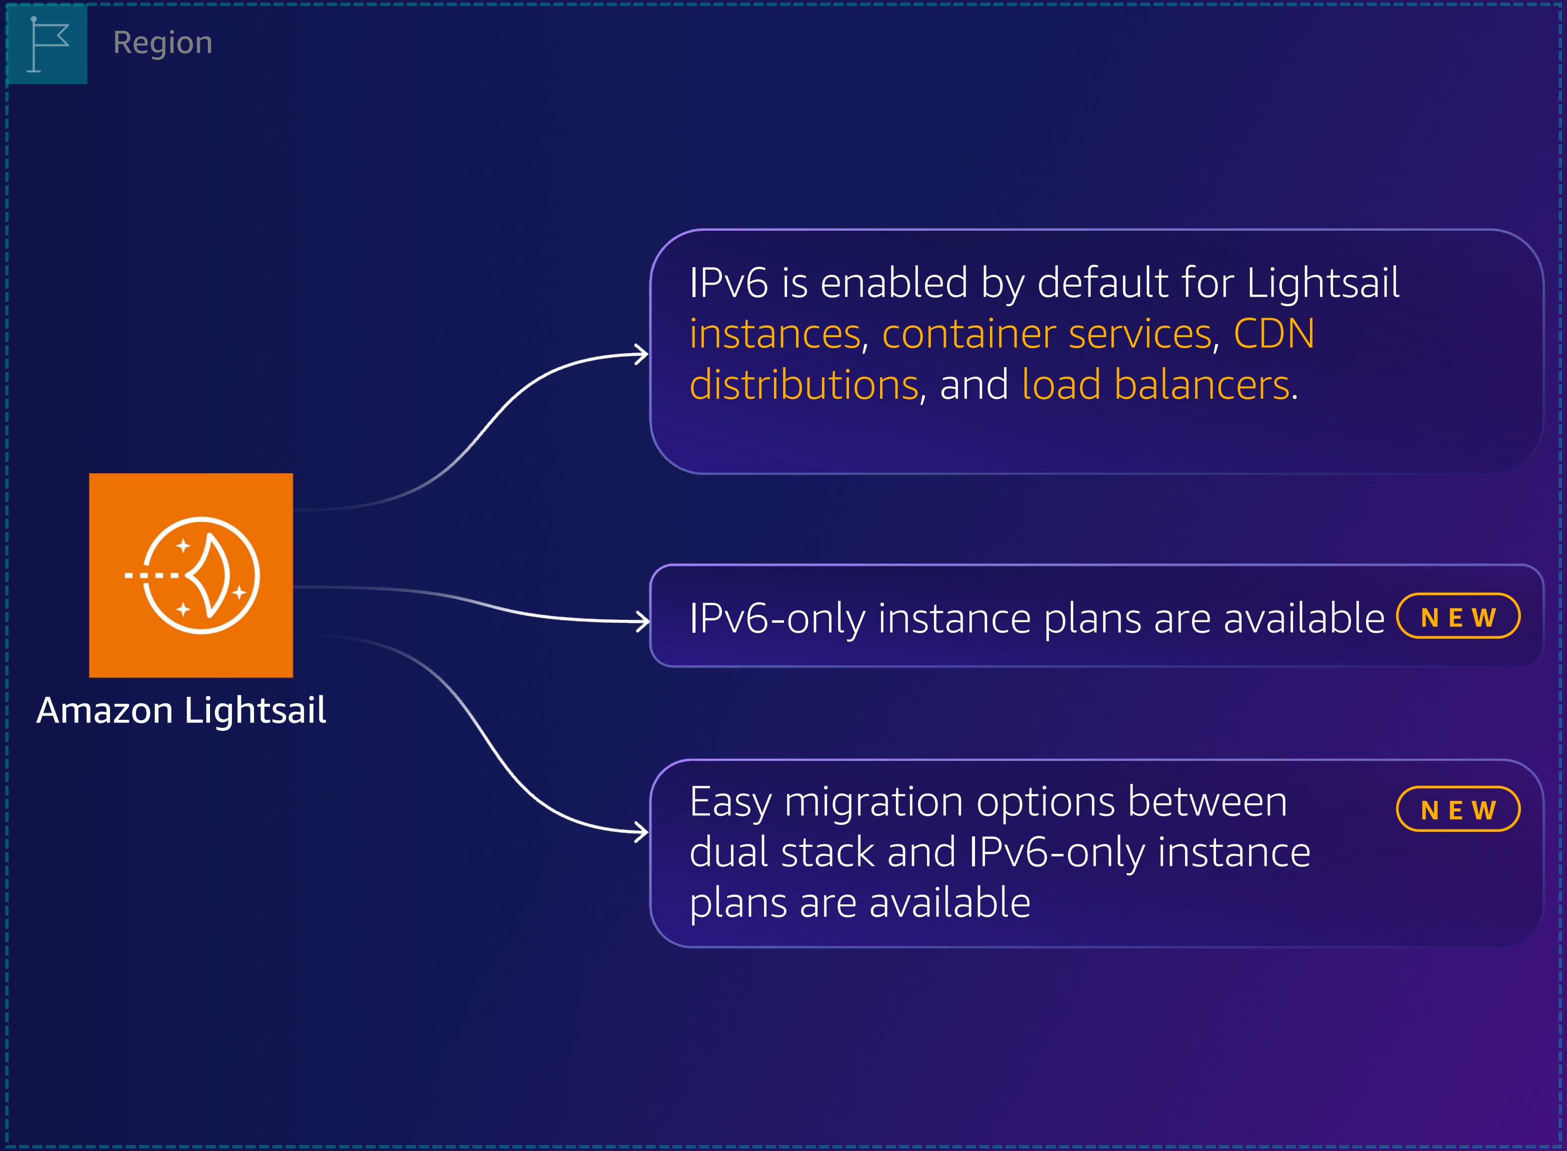

IPv6 is enabled by default for Lightsail instances, container services, CDN distributions, and load balancers.

IPv6-only instance plans are available NEW

Easy migration options between dual stack and IPv6-only instance plans are available NEW

AWS native

IPv6 backwards compatibility with IPv4

Dual stack Amazon VPC

IPv6 backwards compatibility

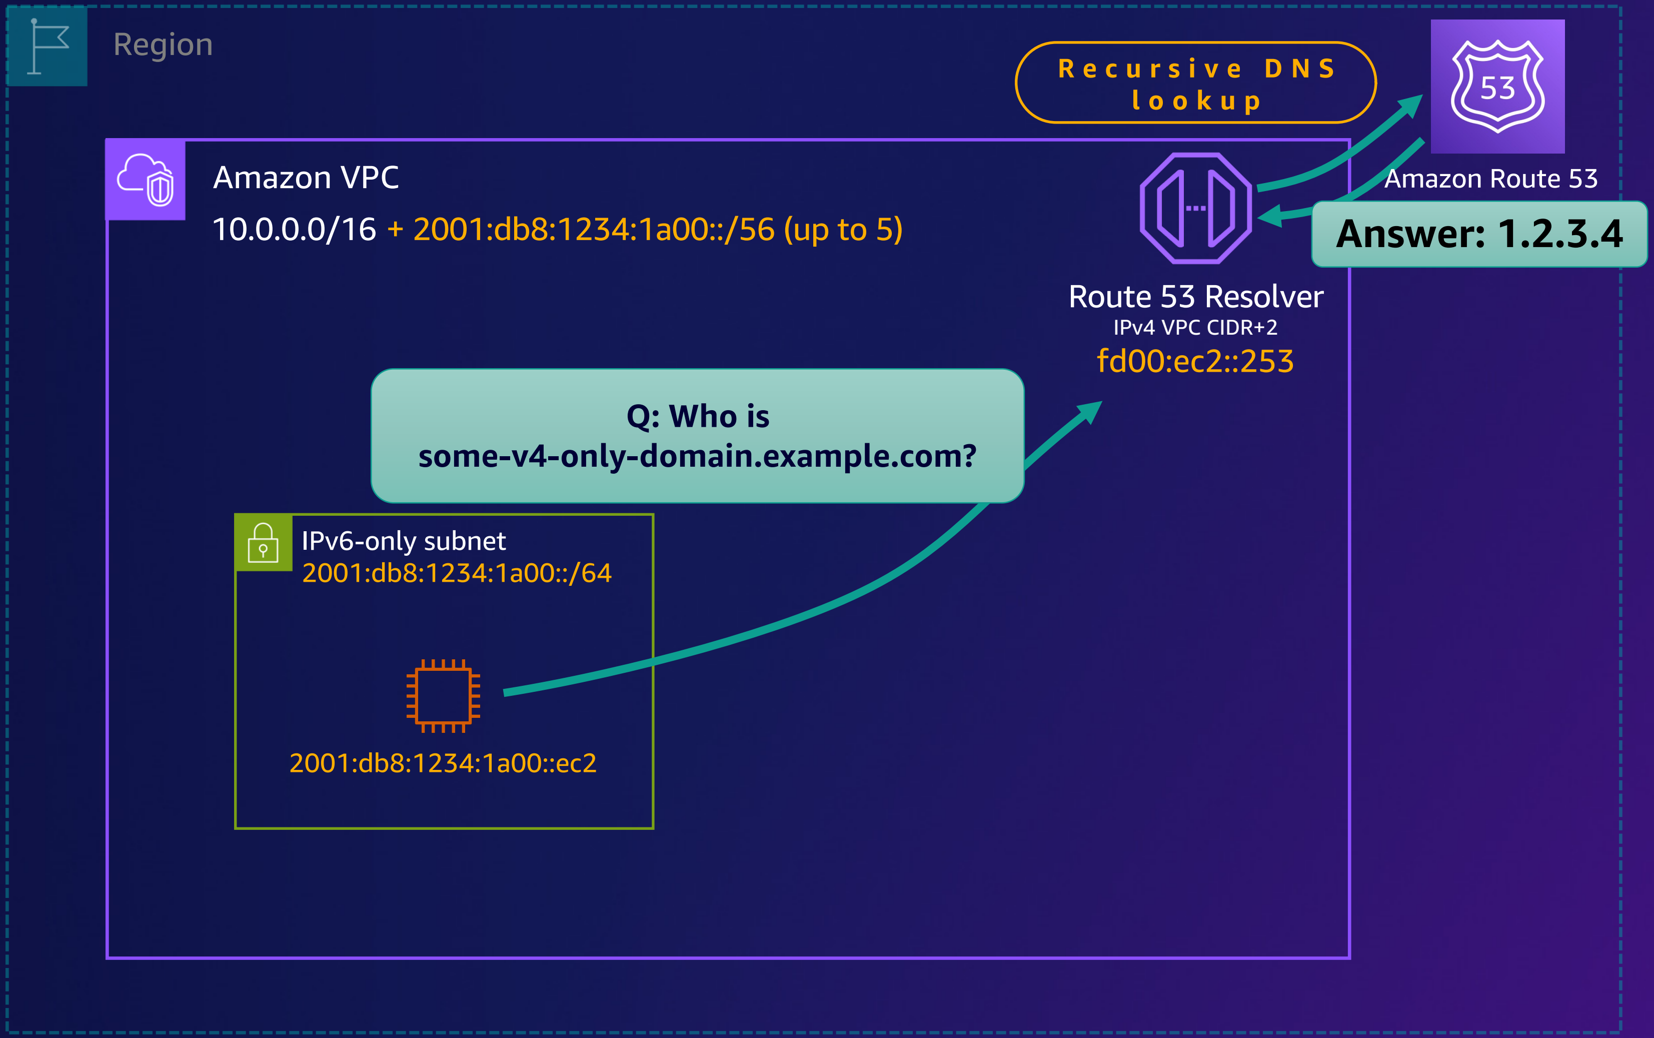

Dual stack Amazon VPC

IPv6 backwards compatibility

DNS64

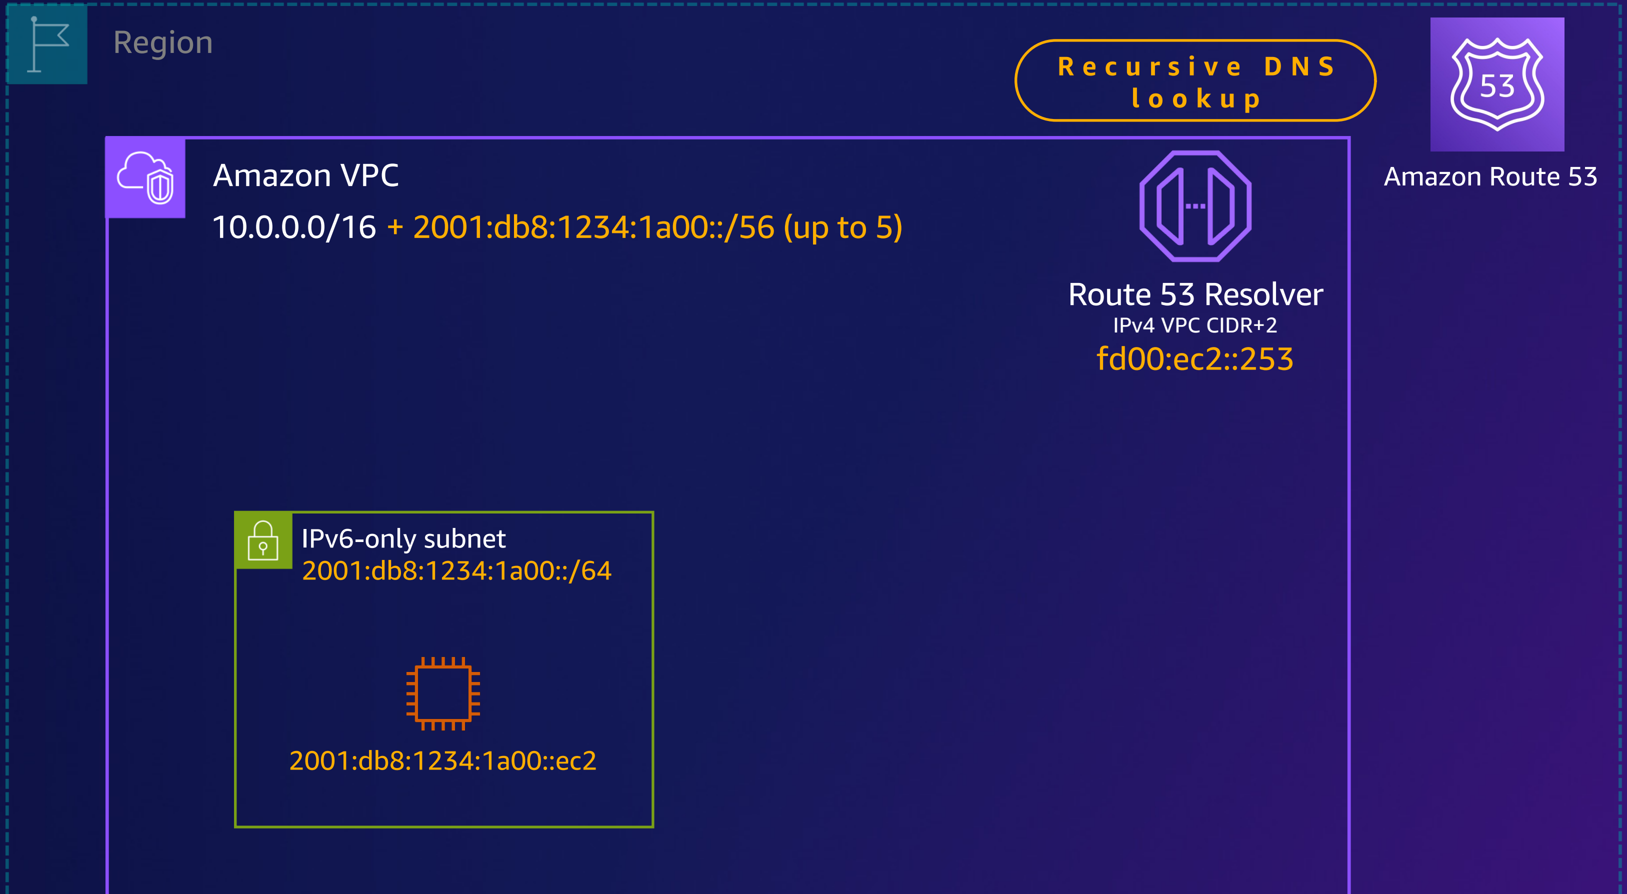

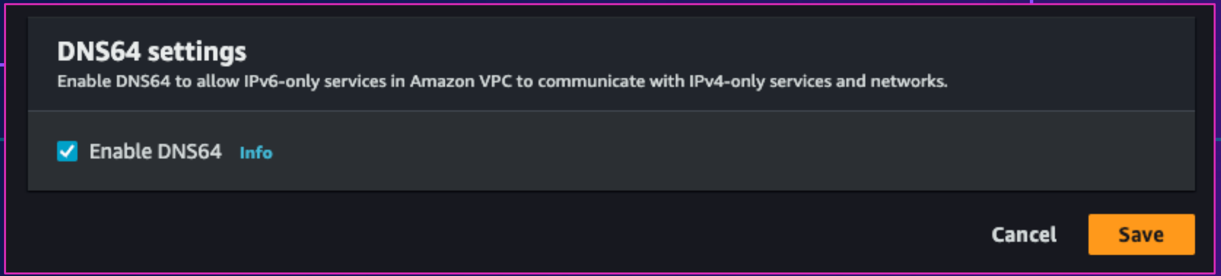

DNS64 settings

Enable DNS64 to allow IPv6-only services in Amazon VPC to communicate with IPv4-only services and networks.

☑ Enable DNS64 Info

Dual stack Amazon VPC

IPv6 backwards compatibility

DNS64

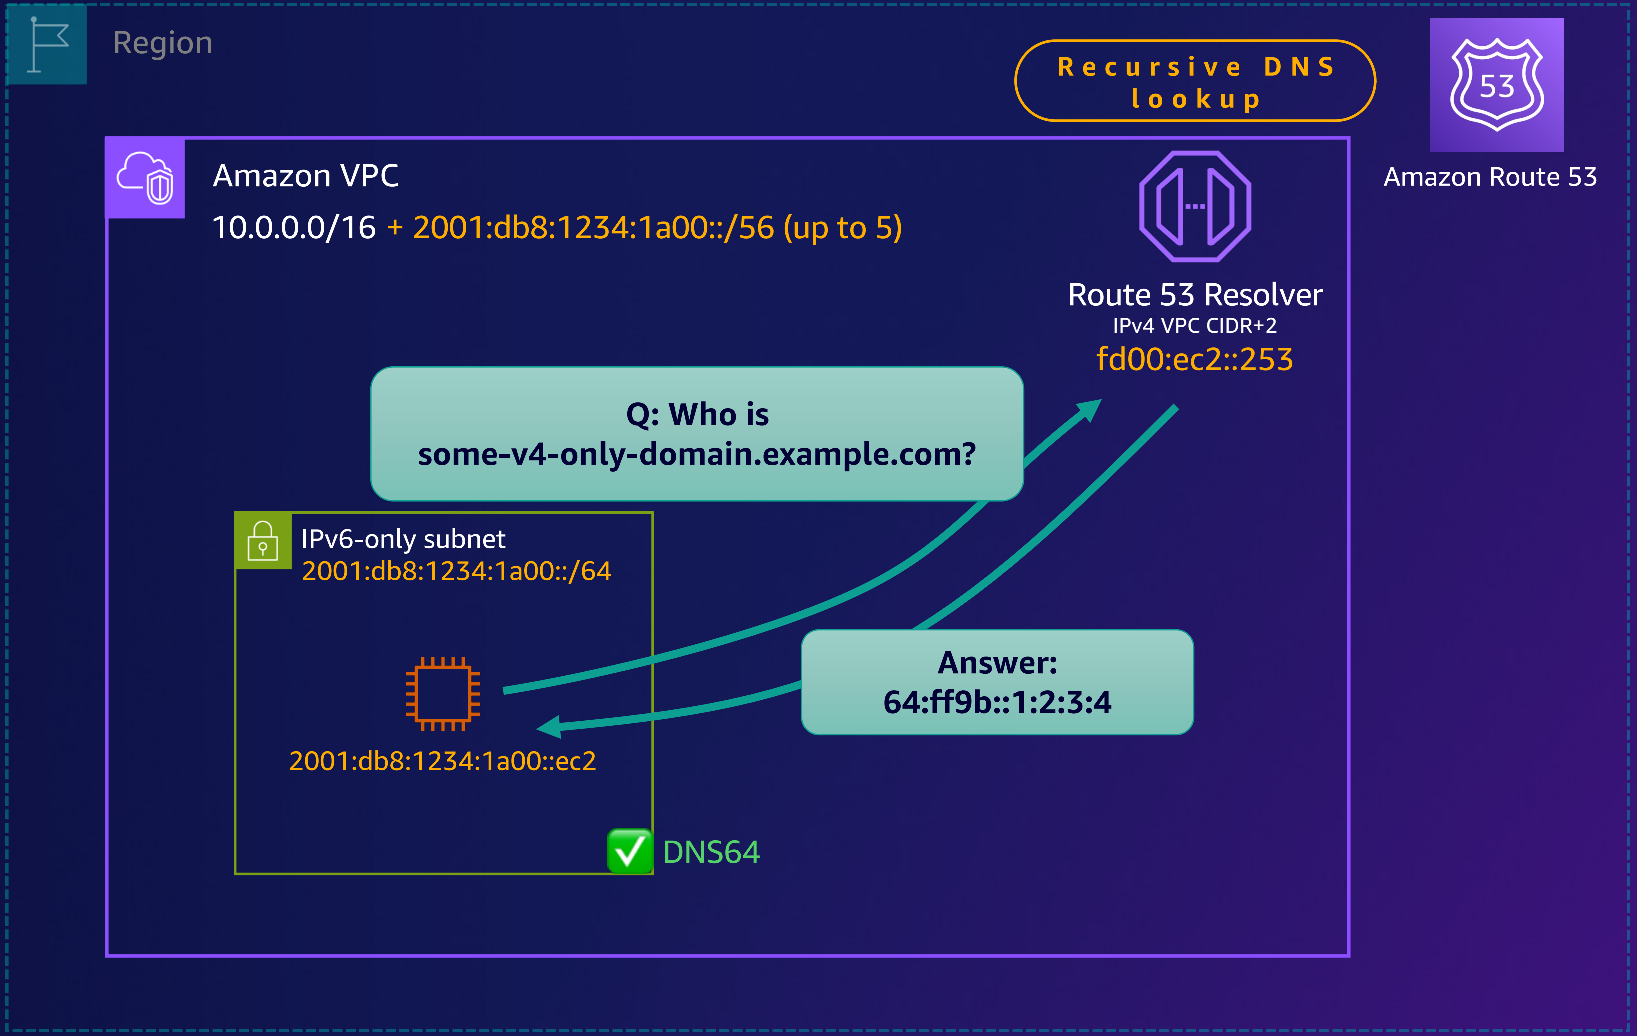

Dual stack Amazon VPC

IPv6 backwards compatibility

DNS64 + NAT64

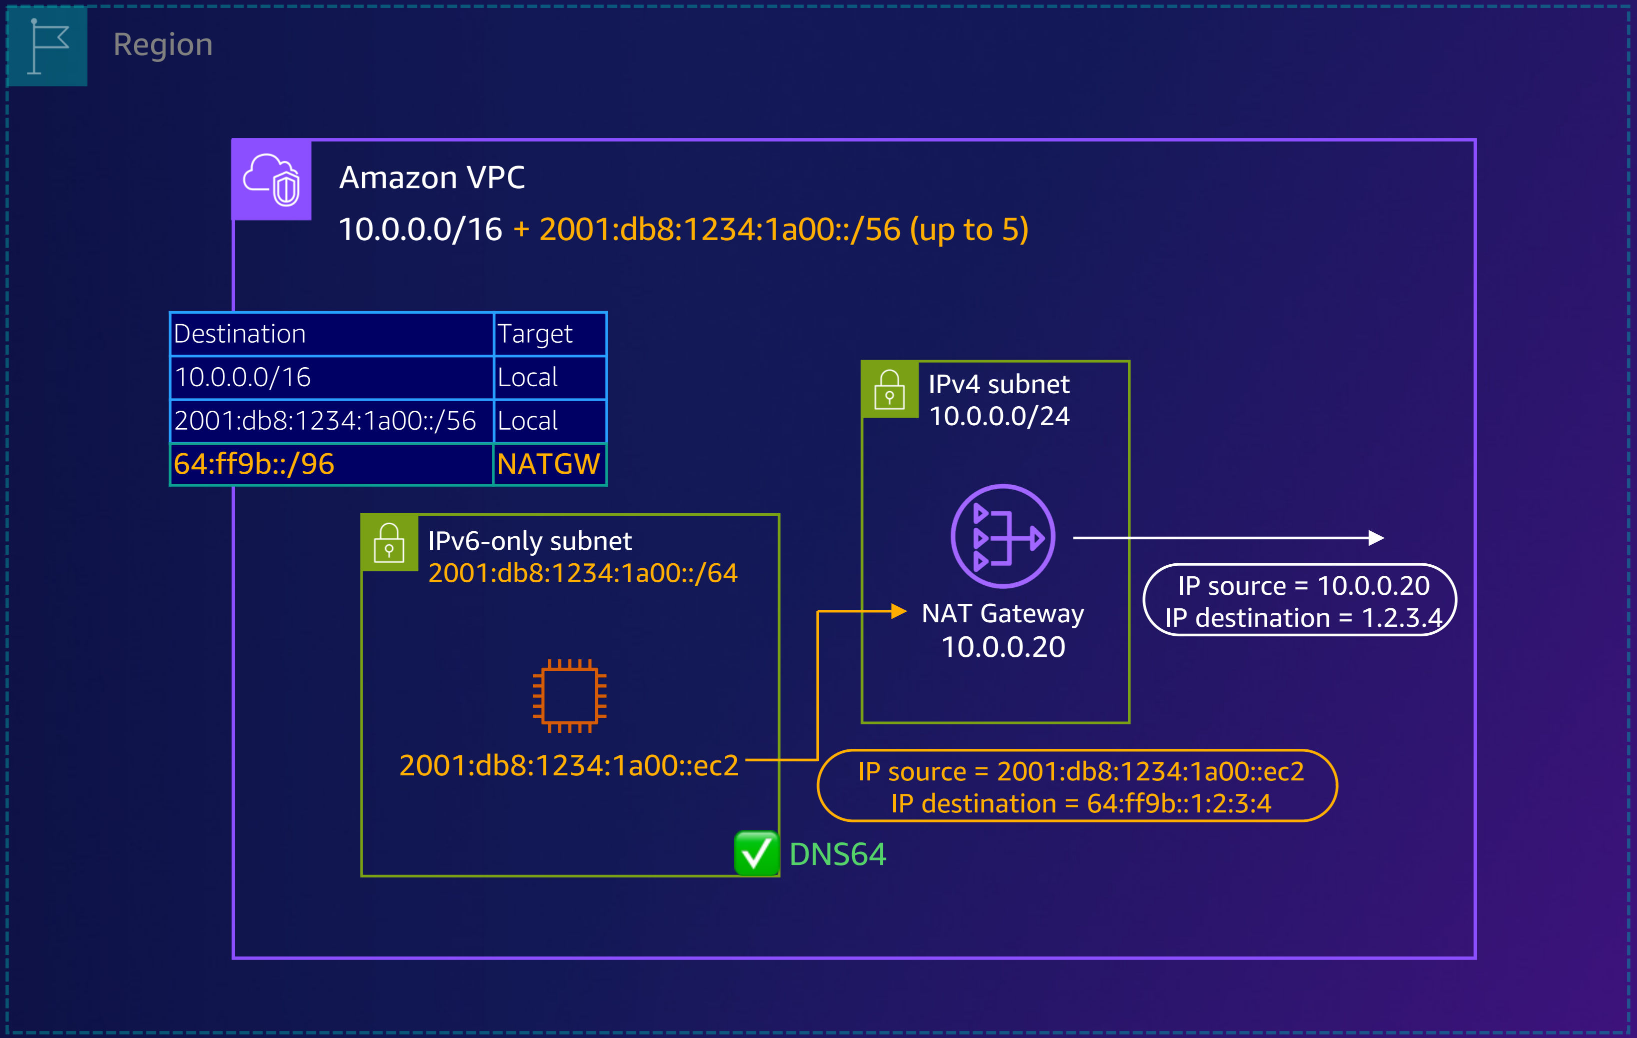

Destination	Target
10.0.0.0/16	Local
2001:db8:1234:1a00::/56	Local
64:ff9b::/96	NATGW

Connectivity

Dual stack Amazon VPC

Dual stack VPC connectivity

Internet connectivity

Dual stack Amazon VPC

Internet connectivity

Public subnets

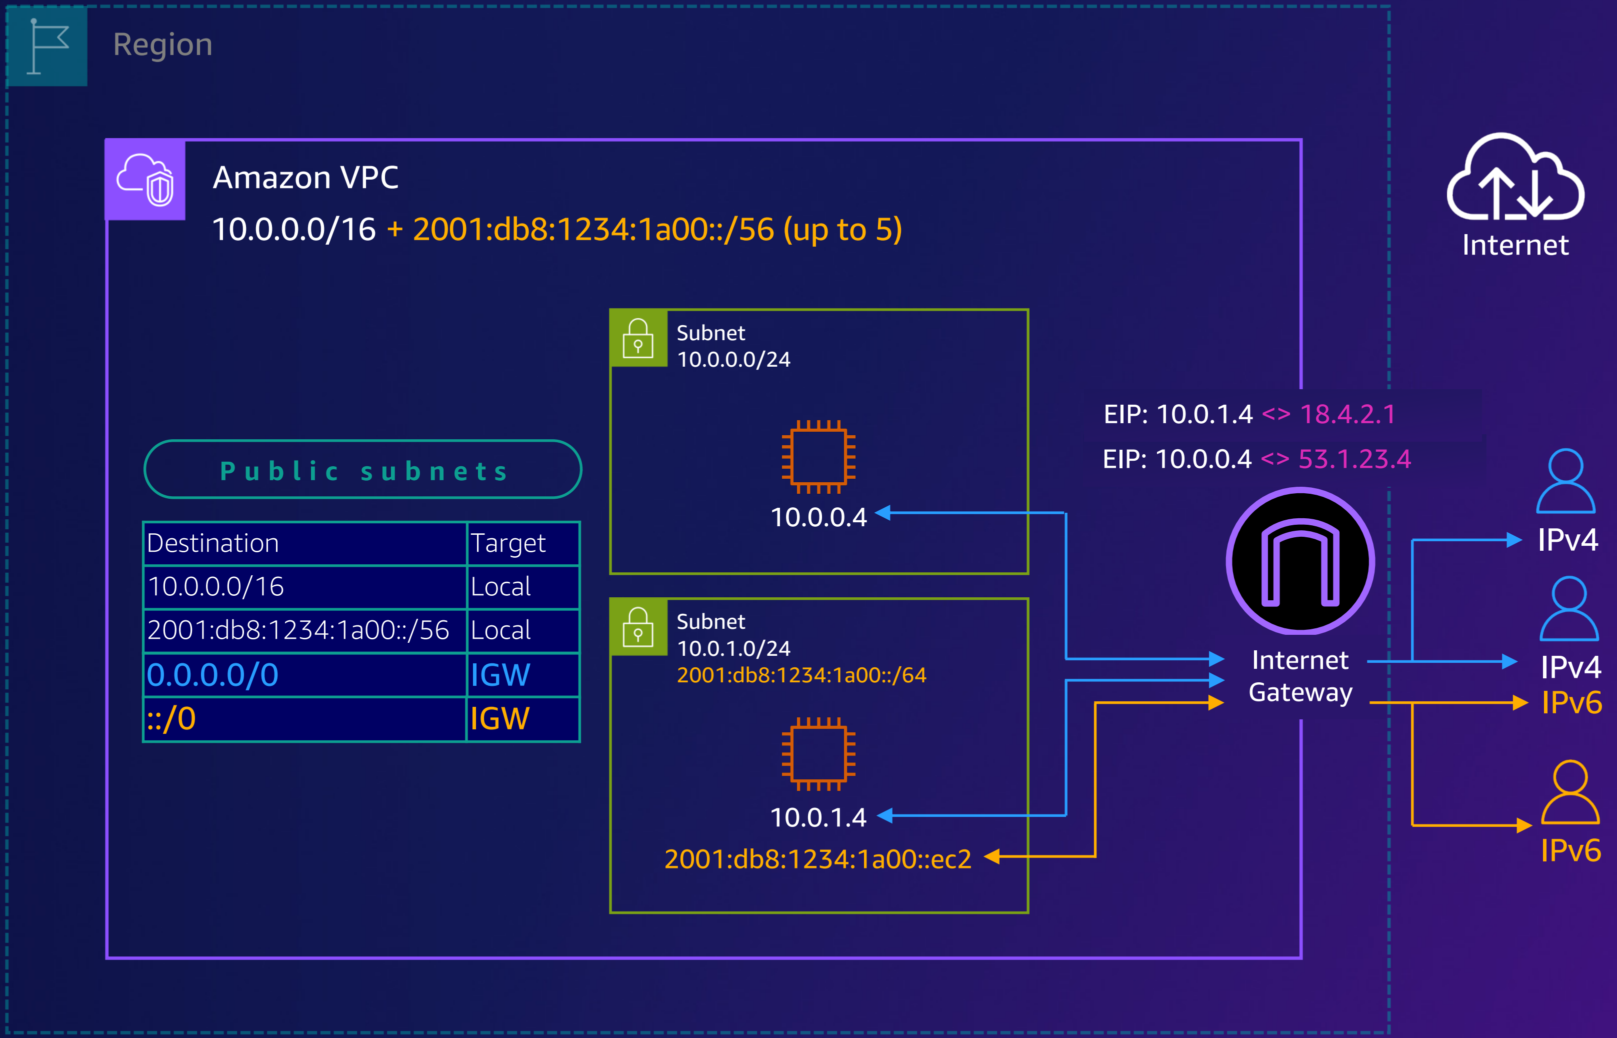

Destination	Target
10.0.0.0/16	Local
2001:db8:1234:1a00::/56	Local
0.0.0.0/0	IGW
::/0	IGW

Dual stack Amazon VPC

Internet connectivity

Public subnets

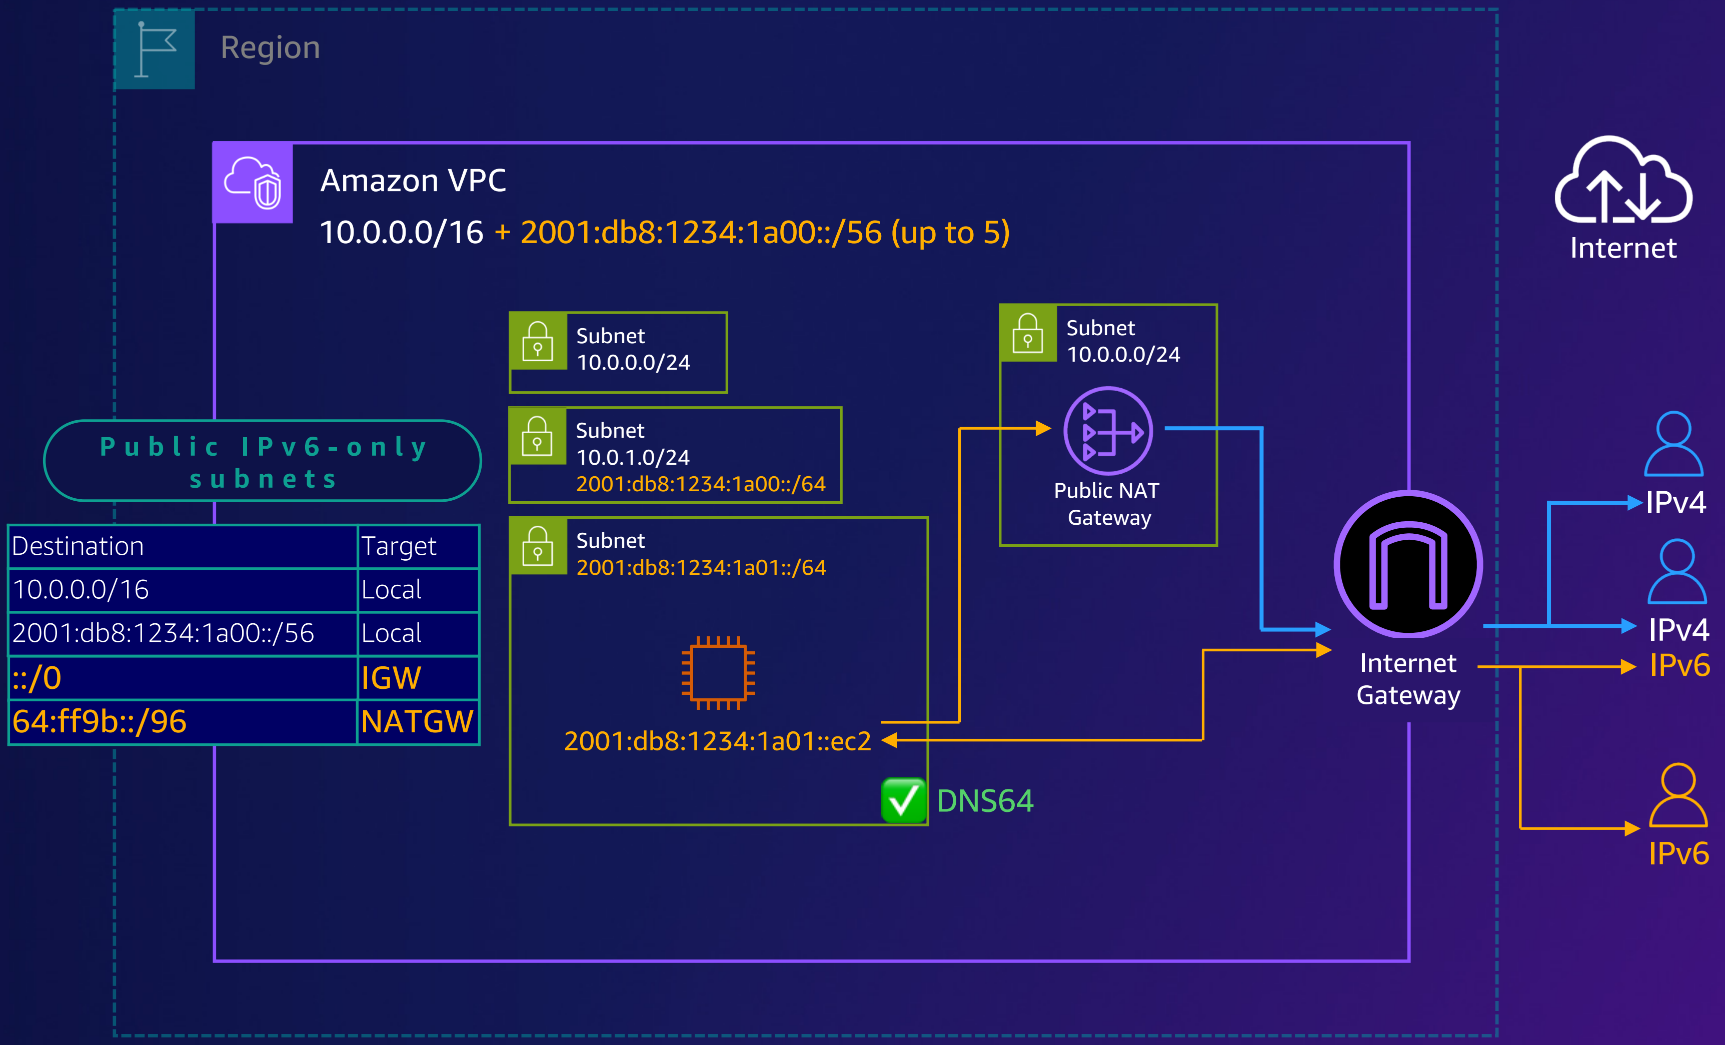

Destination	Target
10.0.0.0/16	Local
2001:db8:1234:1a00::/56	Local
::/0	IGW
64:ff9b::/96	NATGW

Dual stack Amazon VPC

Internet connectivity

Public subnets

Private subnets

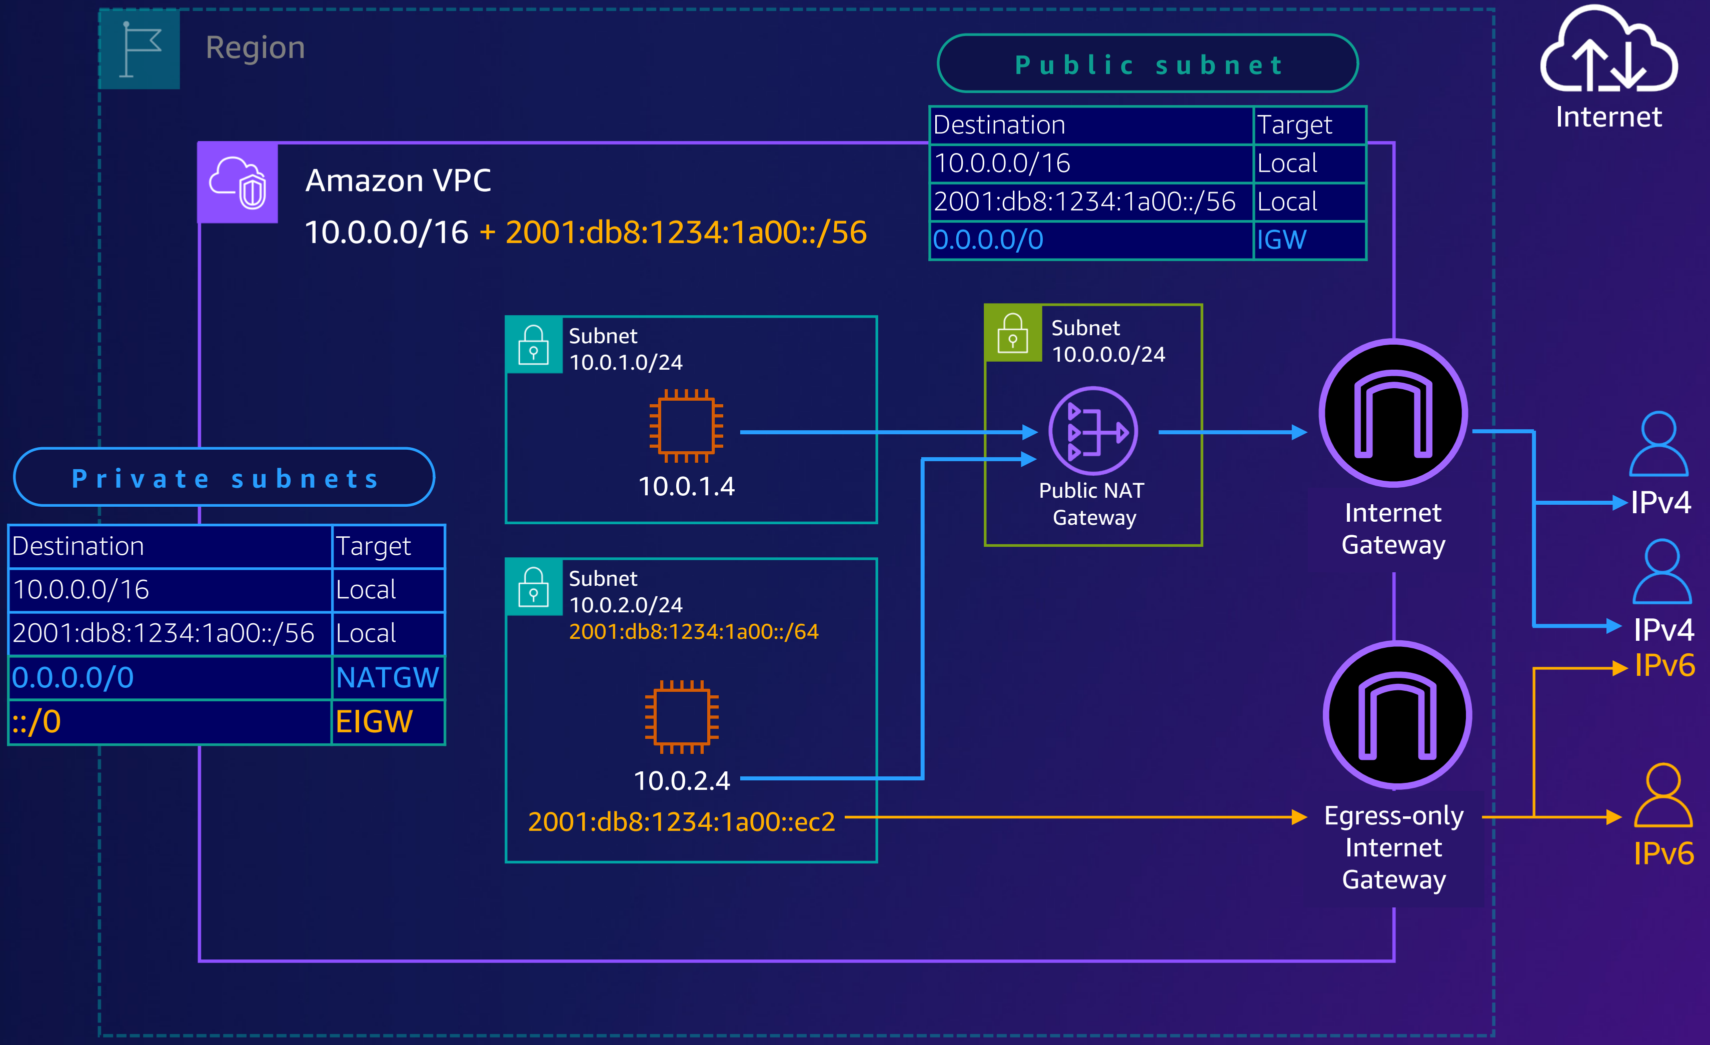

Public subnet

Destination	Target
10.0.0.0/16	Local
2001:db8:1234:1a00::/56	Local
0.0.0.0/0	IGW

Private subnets

Destination	Target
10.0.0.0/16	Local
2001:db8:1234:1a00::/56	Local
0.0.0.0/0	NATGW
::/0	EIGW

Dual stack Amazon VPC

Internet connectivity

Public subnets

Private subnets

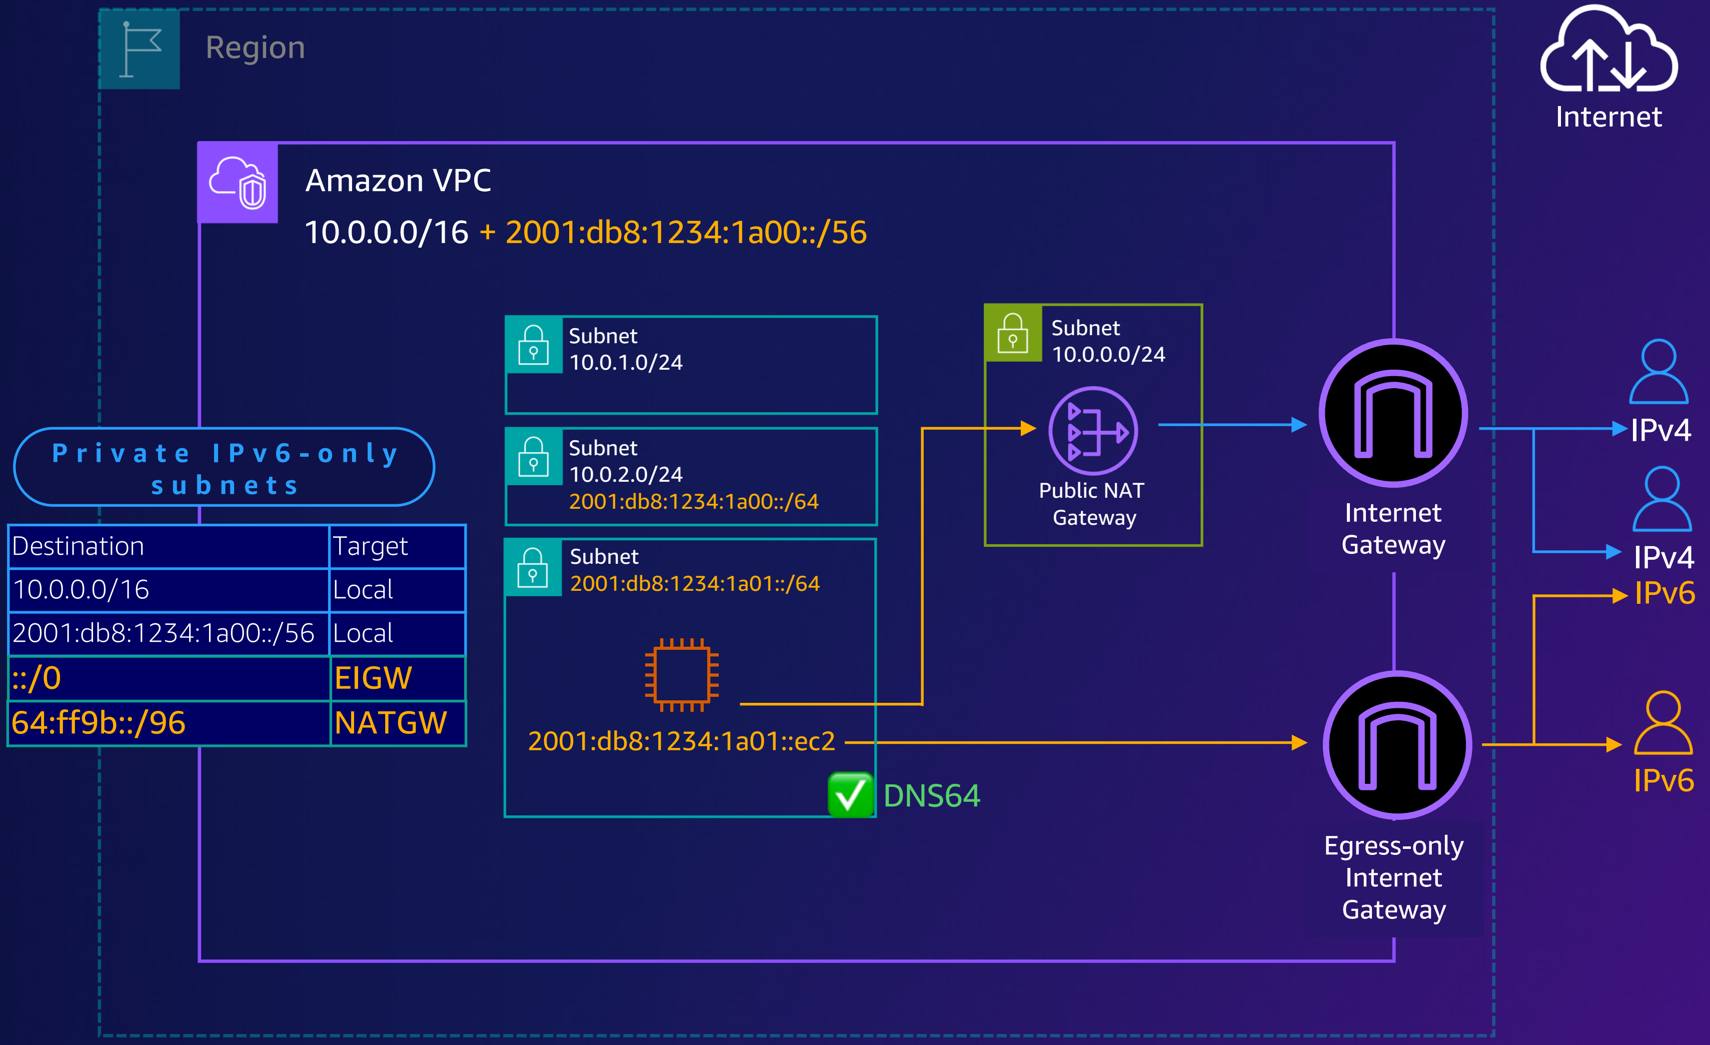

Destination	Target
10.0.0.0/16	Local
2001:db8:1234:1a00::/56	Local
::/0	EIGW
64:ff9b::/96	NATGW

Dual stack VPC connectivity

Internet connectivity

VPC to VPC connectivity

Dual stack Amazon VPC

VPC to VPC connectivity

VPC Peering

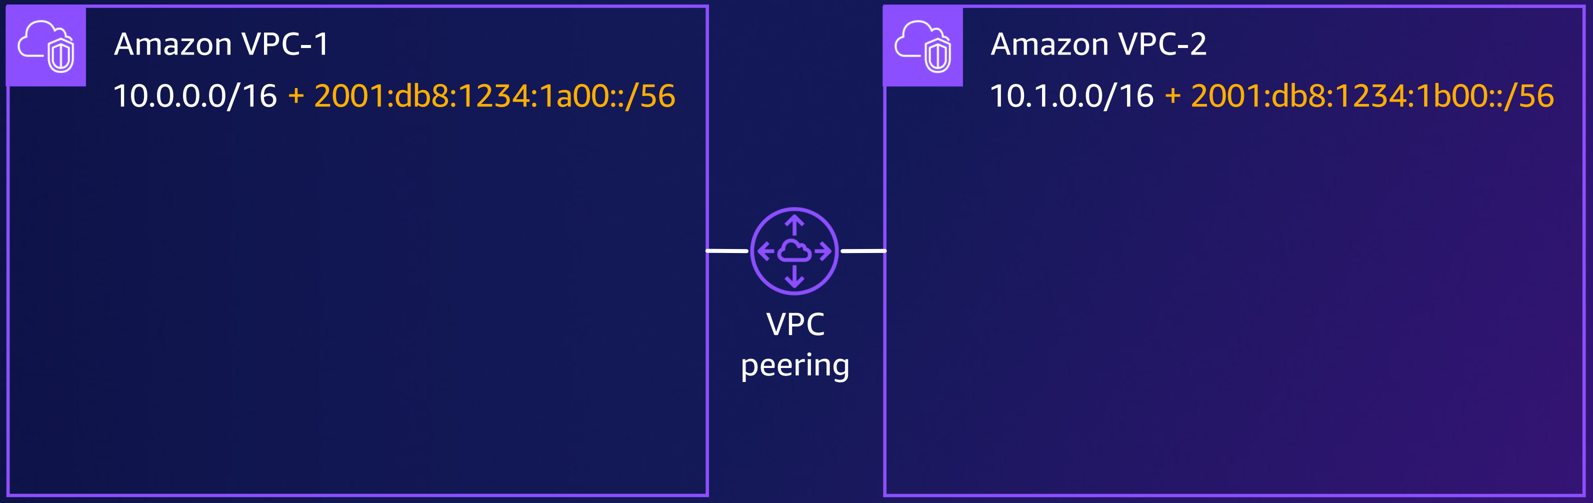
Amazon VPC-1
10.0.0.0/16 + 2001:db8:1234:1a00::/56
VPC
peering
Amazon VPC-2
10.1.0.0/16 + 2001:db8:1234:1b00::/56

VPC-1 Route Table(s)

Destination	Target
10.0.0.0/16	Local
2001:db8:1234:1a00::/56	Local
10.1.0.0/16	PCX-ID
2001:db8:1234:1b00::/56	PCX-ID

VPC-2 Route Table(s)

Destination	Target
10.1.0.0/16	Local
2001:db8:1234:1b00::/56	Local
10.0.0.0/16	PCX-ID
2001:db8:1234:1a00::/56	PCX-ID

Dual stack Amazon VPC

VPC to VPC connectivity

VPC Peering

AWS Transit Gateway

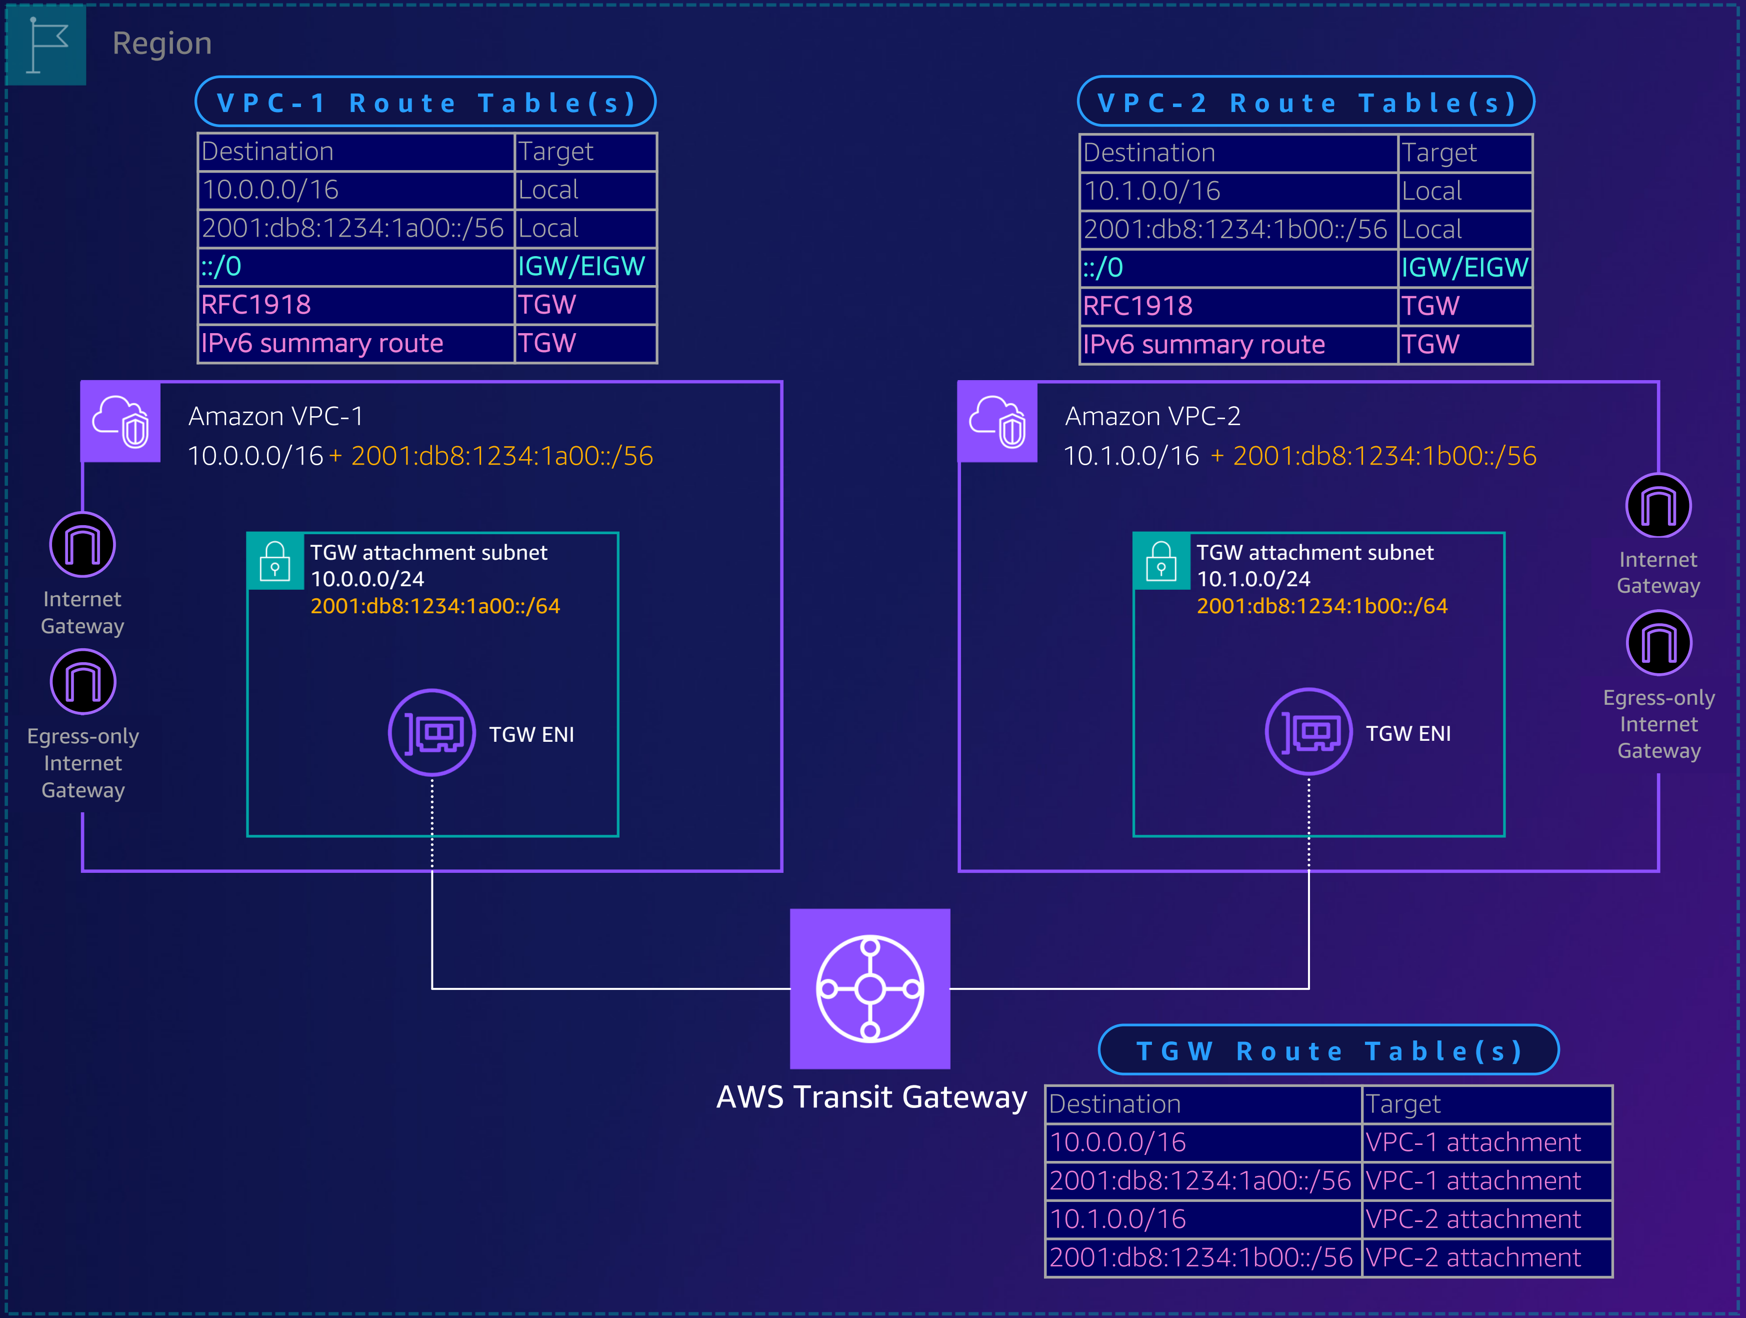

VPC-1 Route Table(s)

Destination	Target
10.0.0.0/16	Local
2001:db8:1234:1a00::/56	Local
::/0	IGW/EIGW
RFC1918	TGW
IPv6 summary route	TGW

VPC-2 Route Table(s)

Destination	Target
10.1.0.0/16	Local
2001:db8:1234:1b00::/56	Local
::/0	IGW/EIGW
RFC1918	TGW
IPv6 summary route	TGW

TGW Route Table(s)

Destination	Target
10.0.0.0/16	VPC-1 attachment
2001:db8:1234:1a00::/56	VPC-1 attachment
10.1.0.0/16	VPC-2 attachment
2001:db8:1234:1b00::/56	VPC-2 attachment

Dual stack Amazon VPC

VPC to VPC connectivity

VPC Peering

AWS Transit Gateway

AWS Cloud WAN

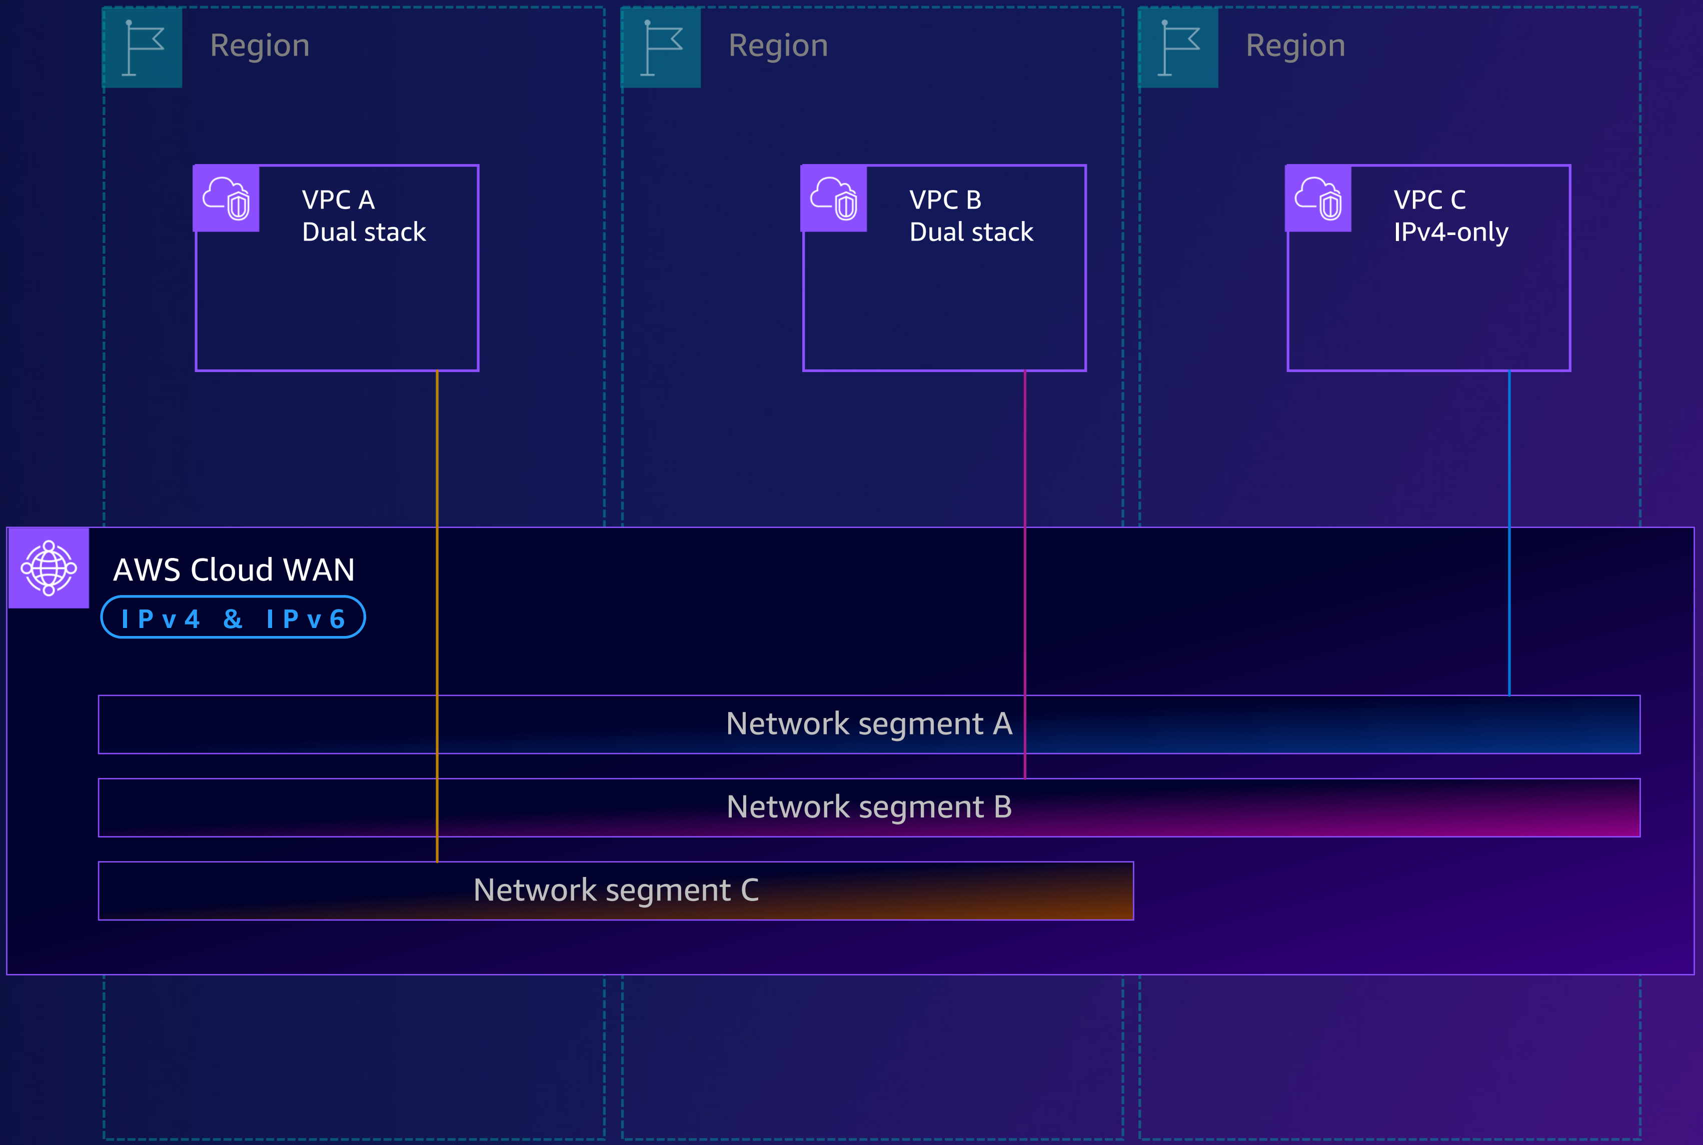

Dual stack VPC connectivity

Internet connectivity

VPC to VPC connectivity

Hybrid connectivity

Dual stack Amazon VPC

Hybrid connectivity

AWS Direct Connect

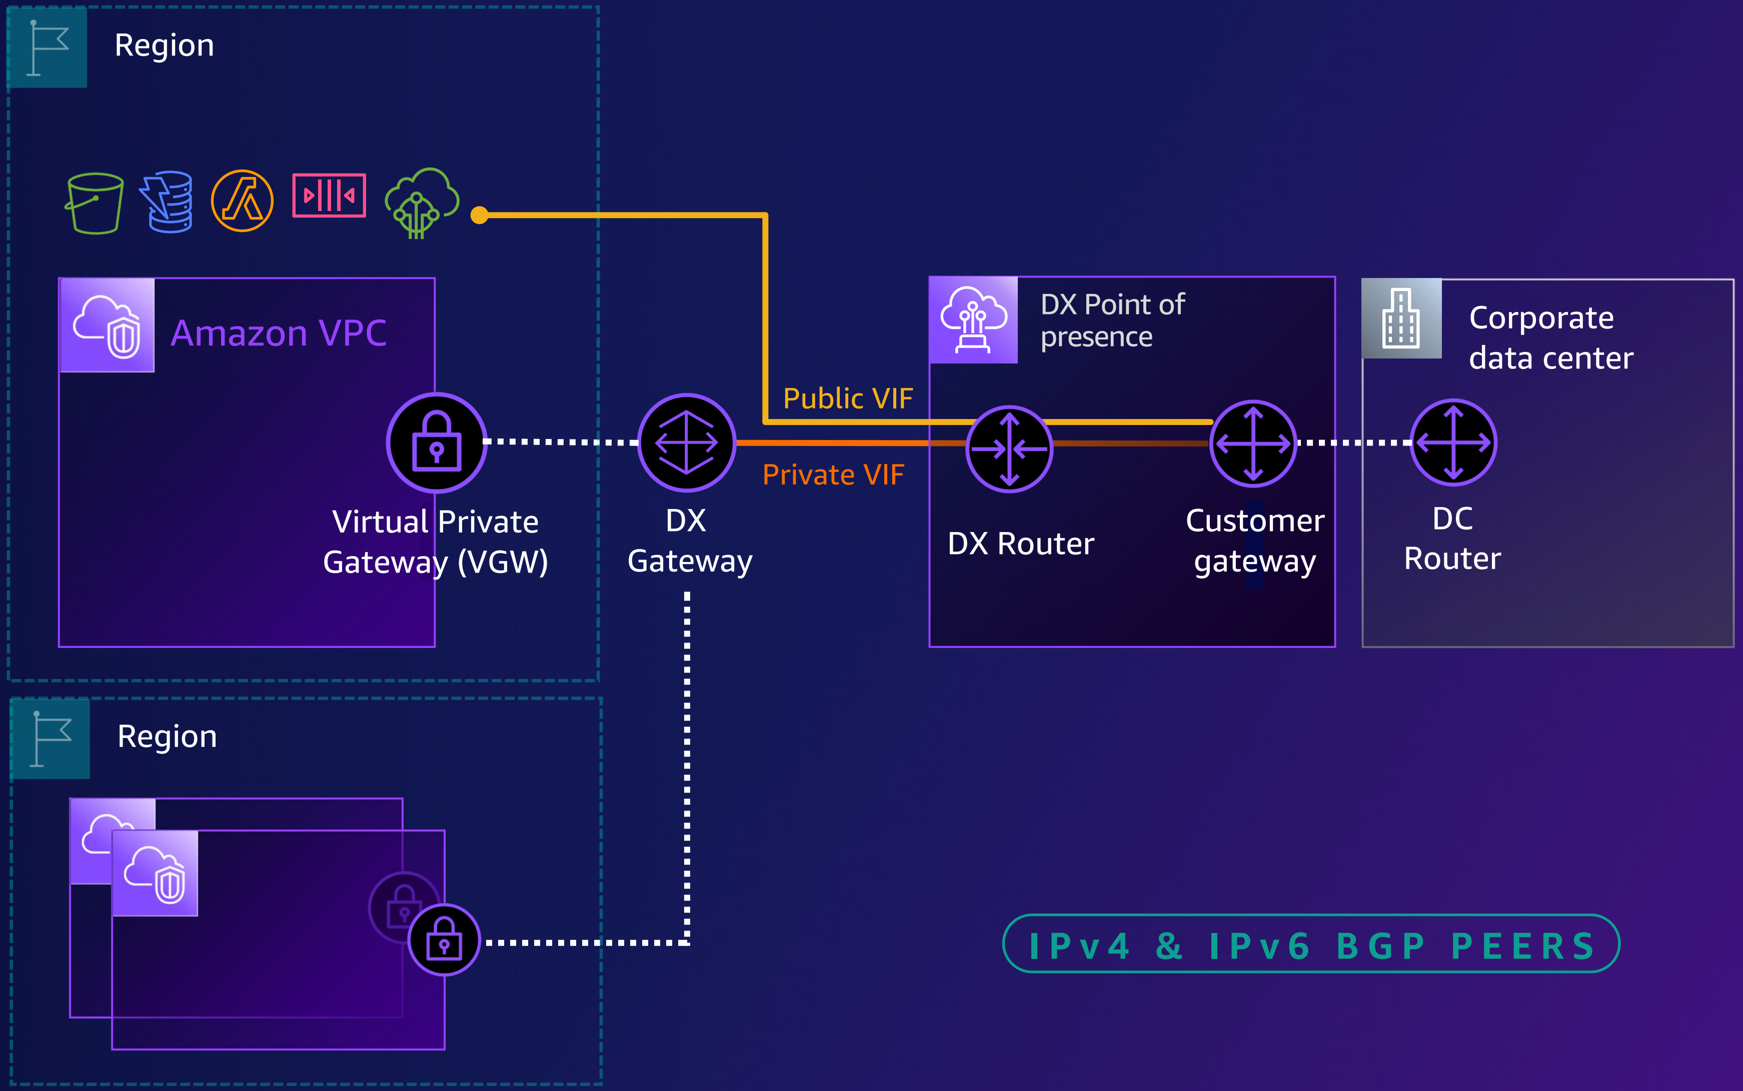

IPv4 & IPv6 BGP PEERS

Dual stack Amazon VPC

Hybrid connectivity

AWS Direct Connect

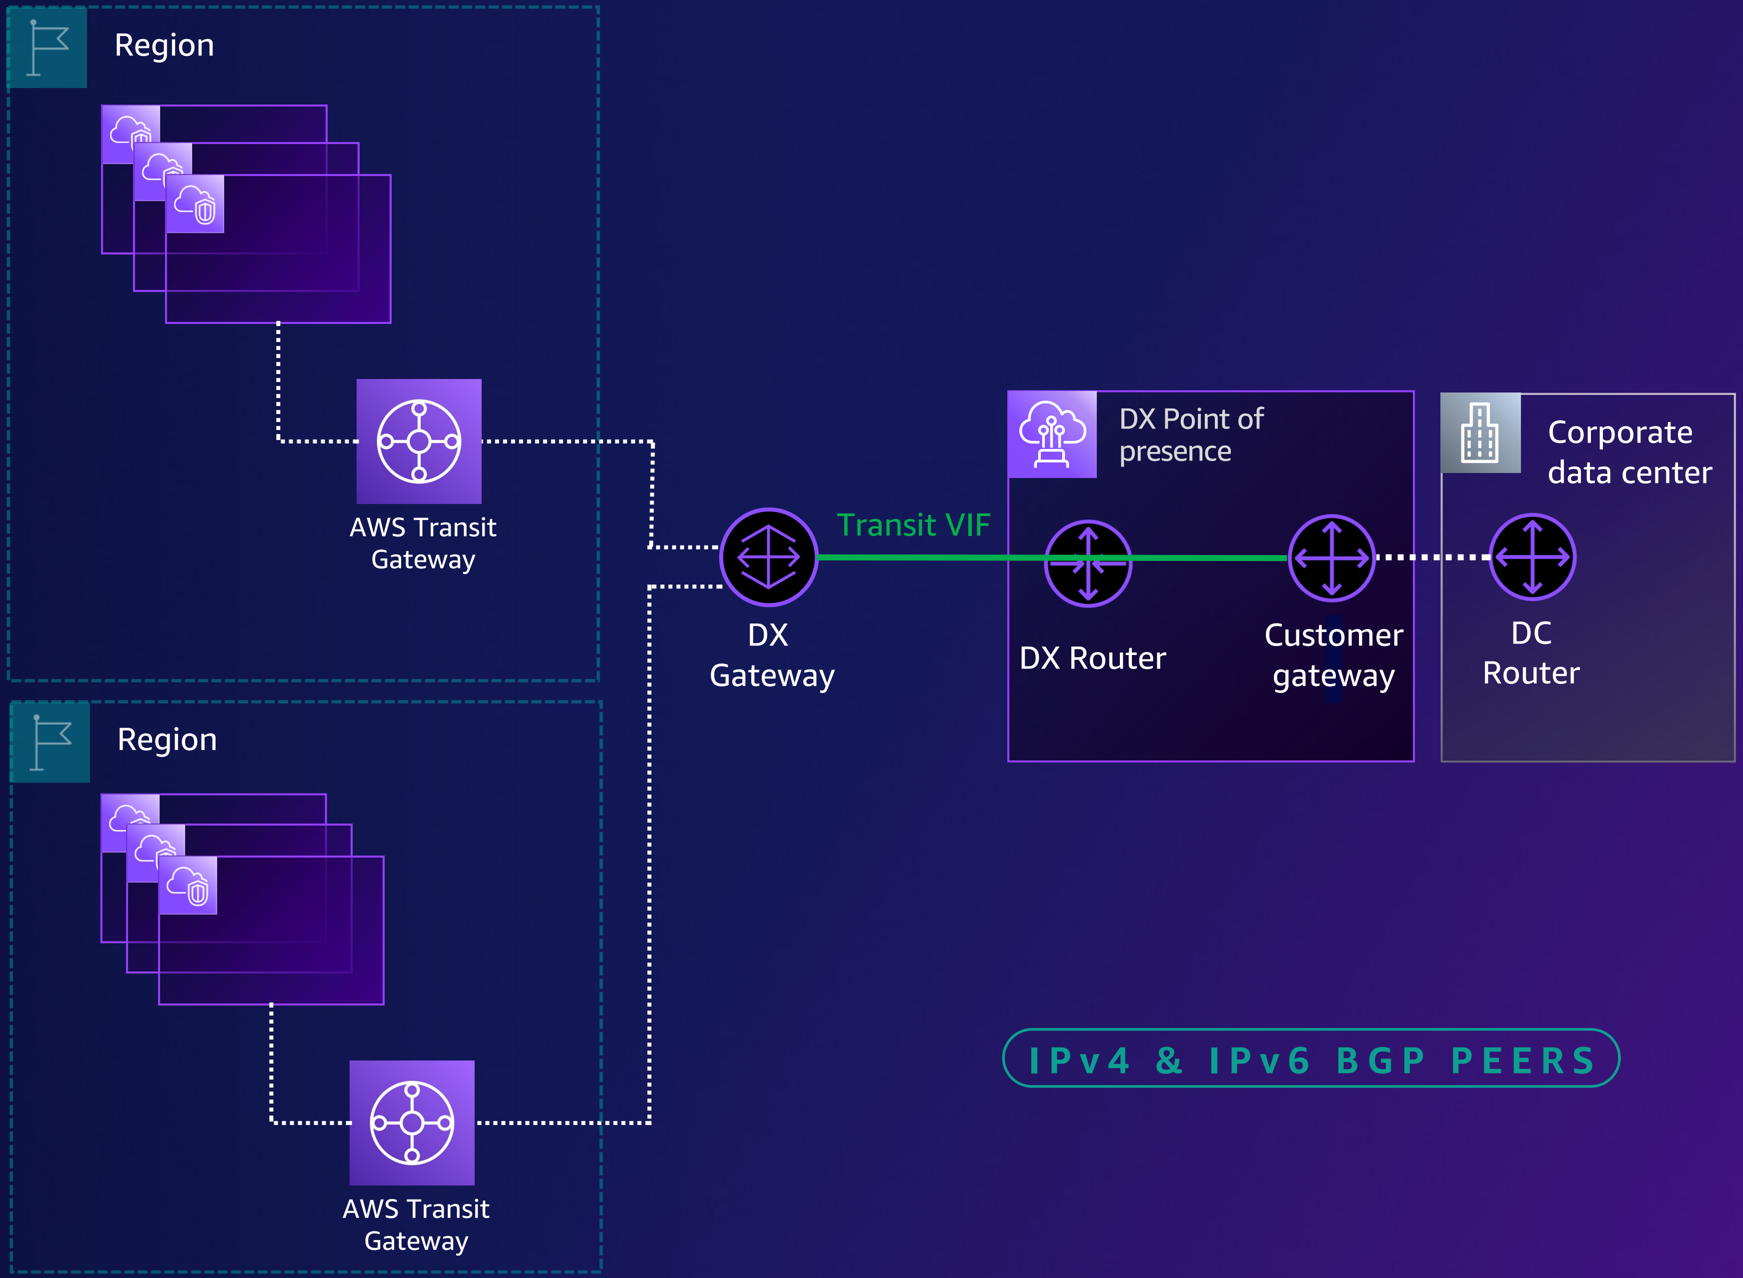

IPv4 & IPv6 BGP PEERS

Dual stack Amazon VPC

Hybrid connectivity

AWS Direct Connect

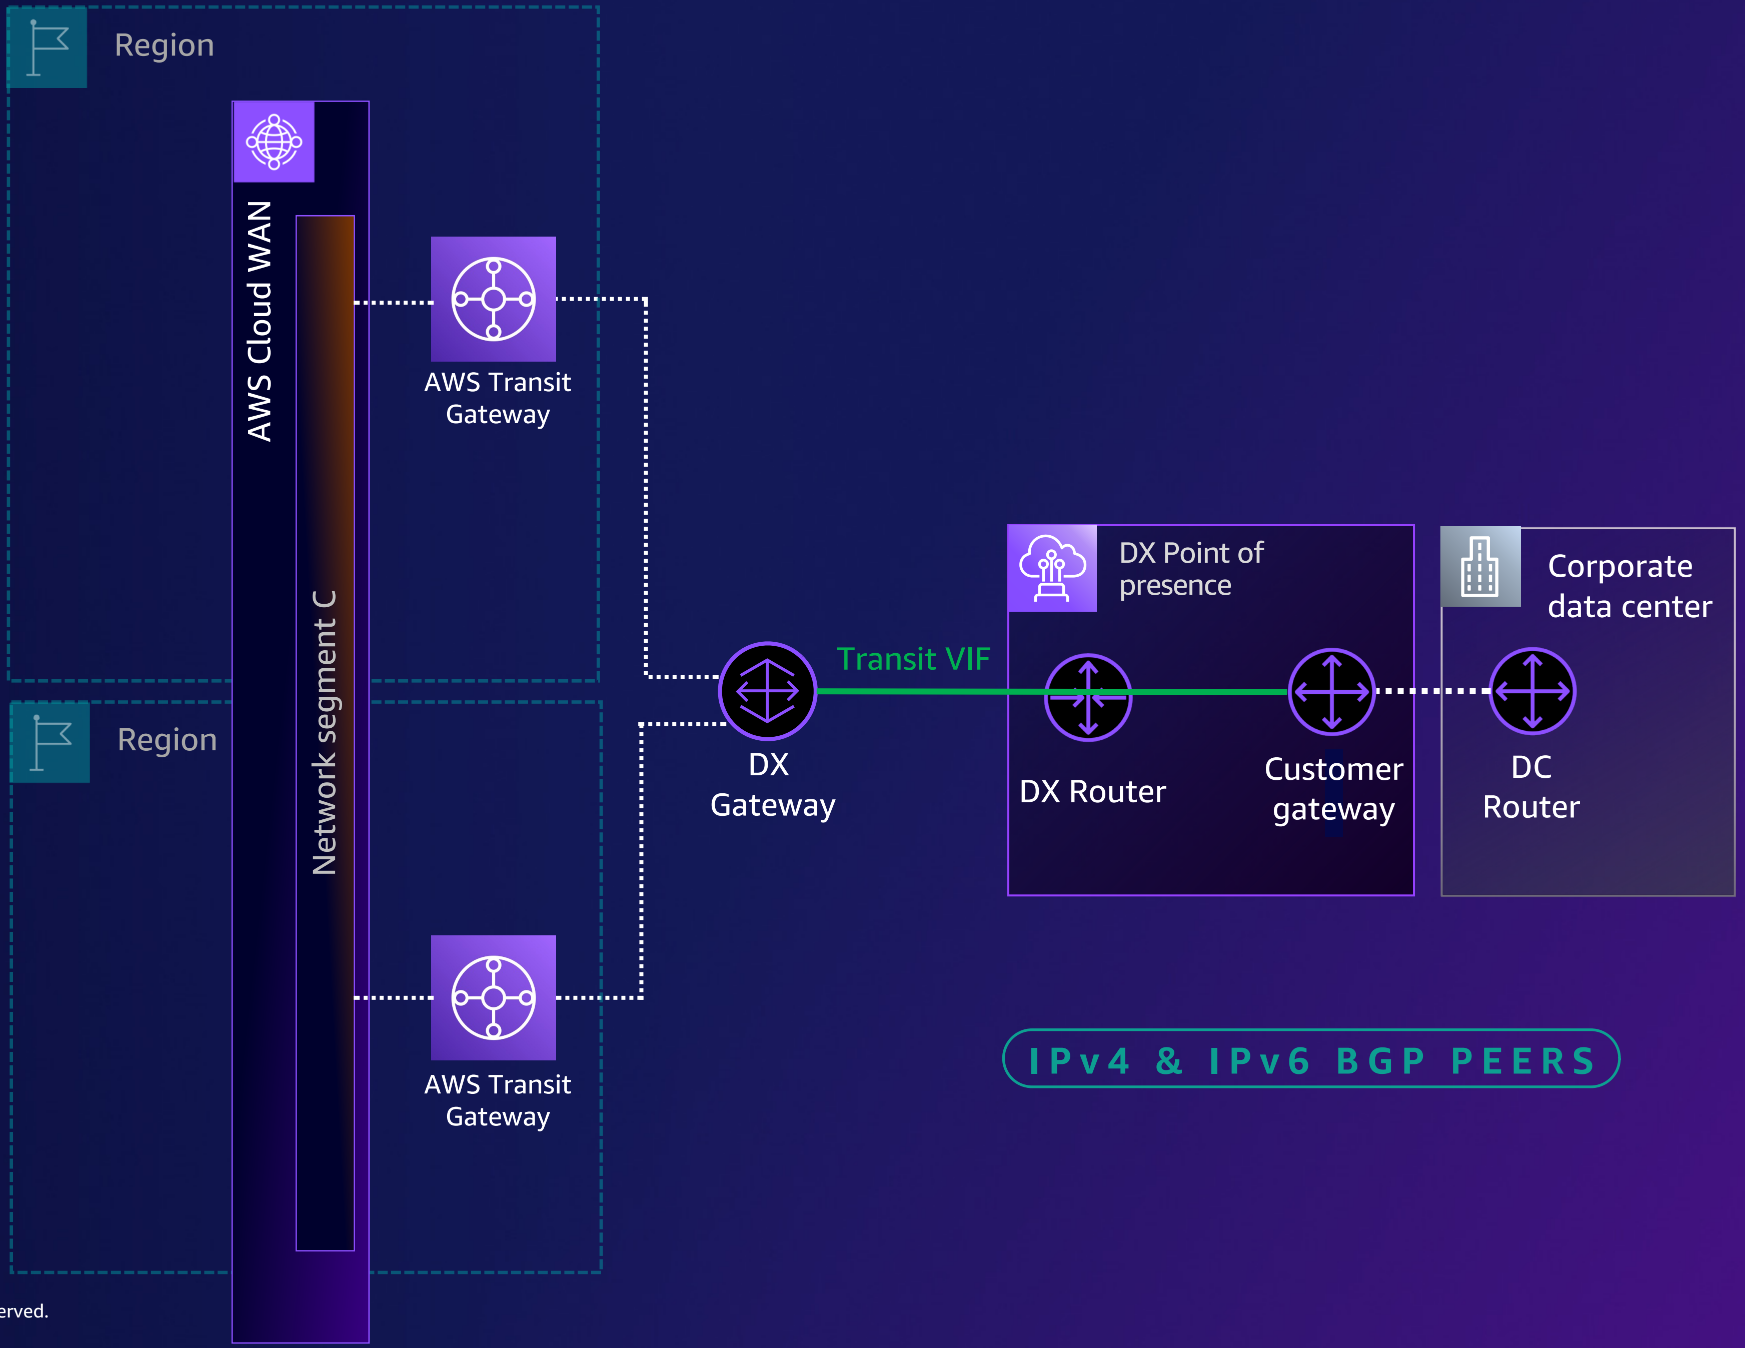

IPv4 & IPv6 BGP PEERS

Dual stack Amazon VPC

Hybrid connectivity

AWS Direct Connect

AWS Site-to-Site VPN

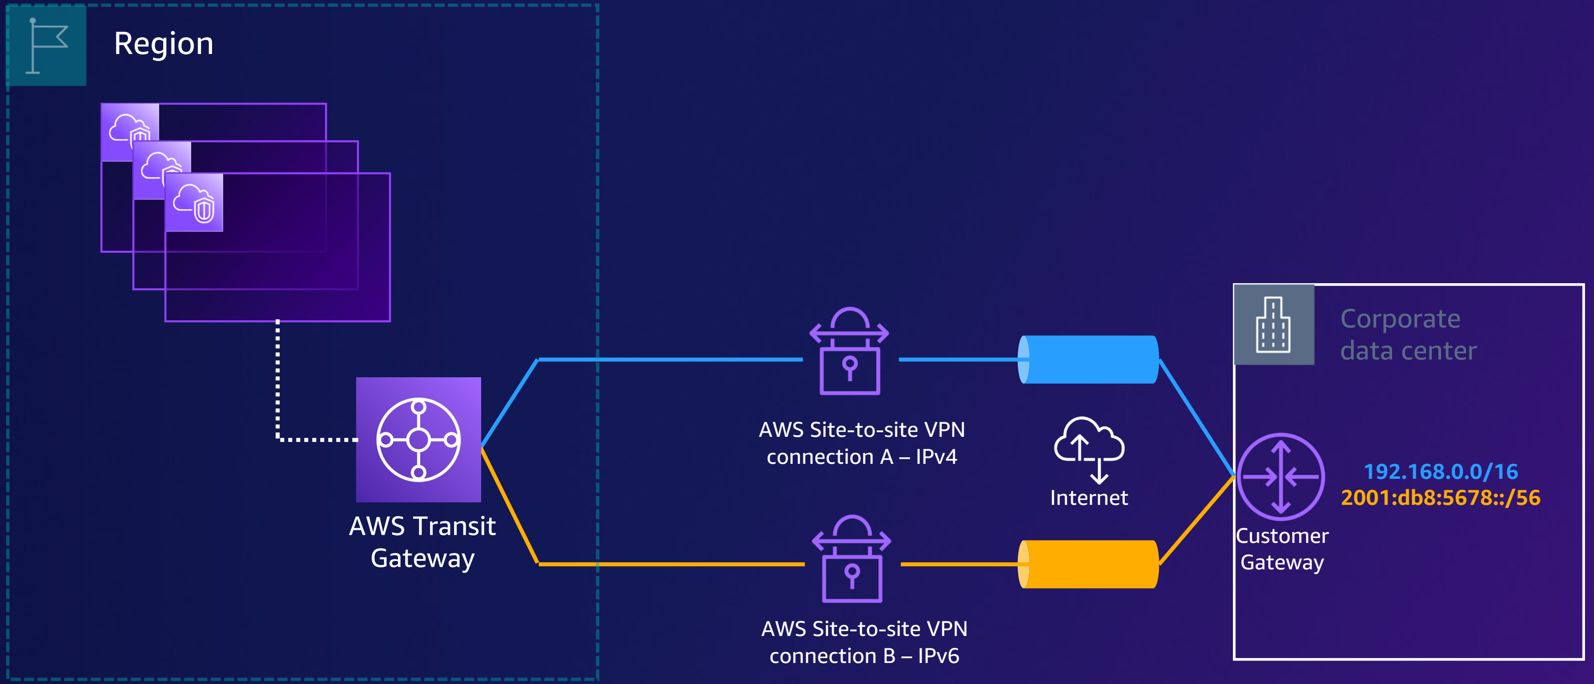

/30 IPv4 inner Tunnel IPs
/126 IPv6 inner Tunnel IPs

Dual stack Amazon VPC

Hybrid connectivity

AWS Direct Connect

AWS Site-to-Site VPN

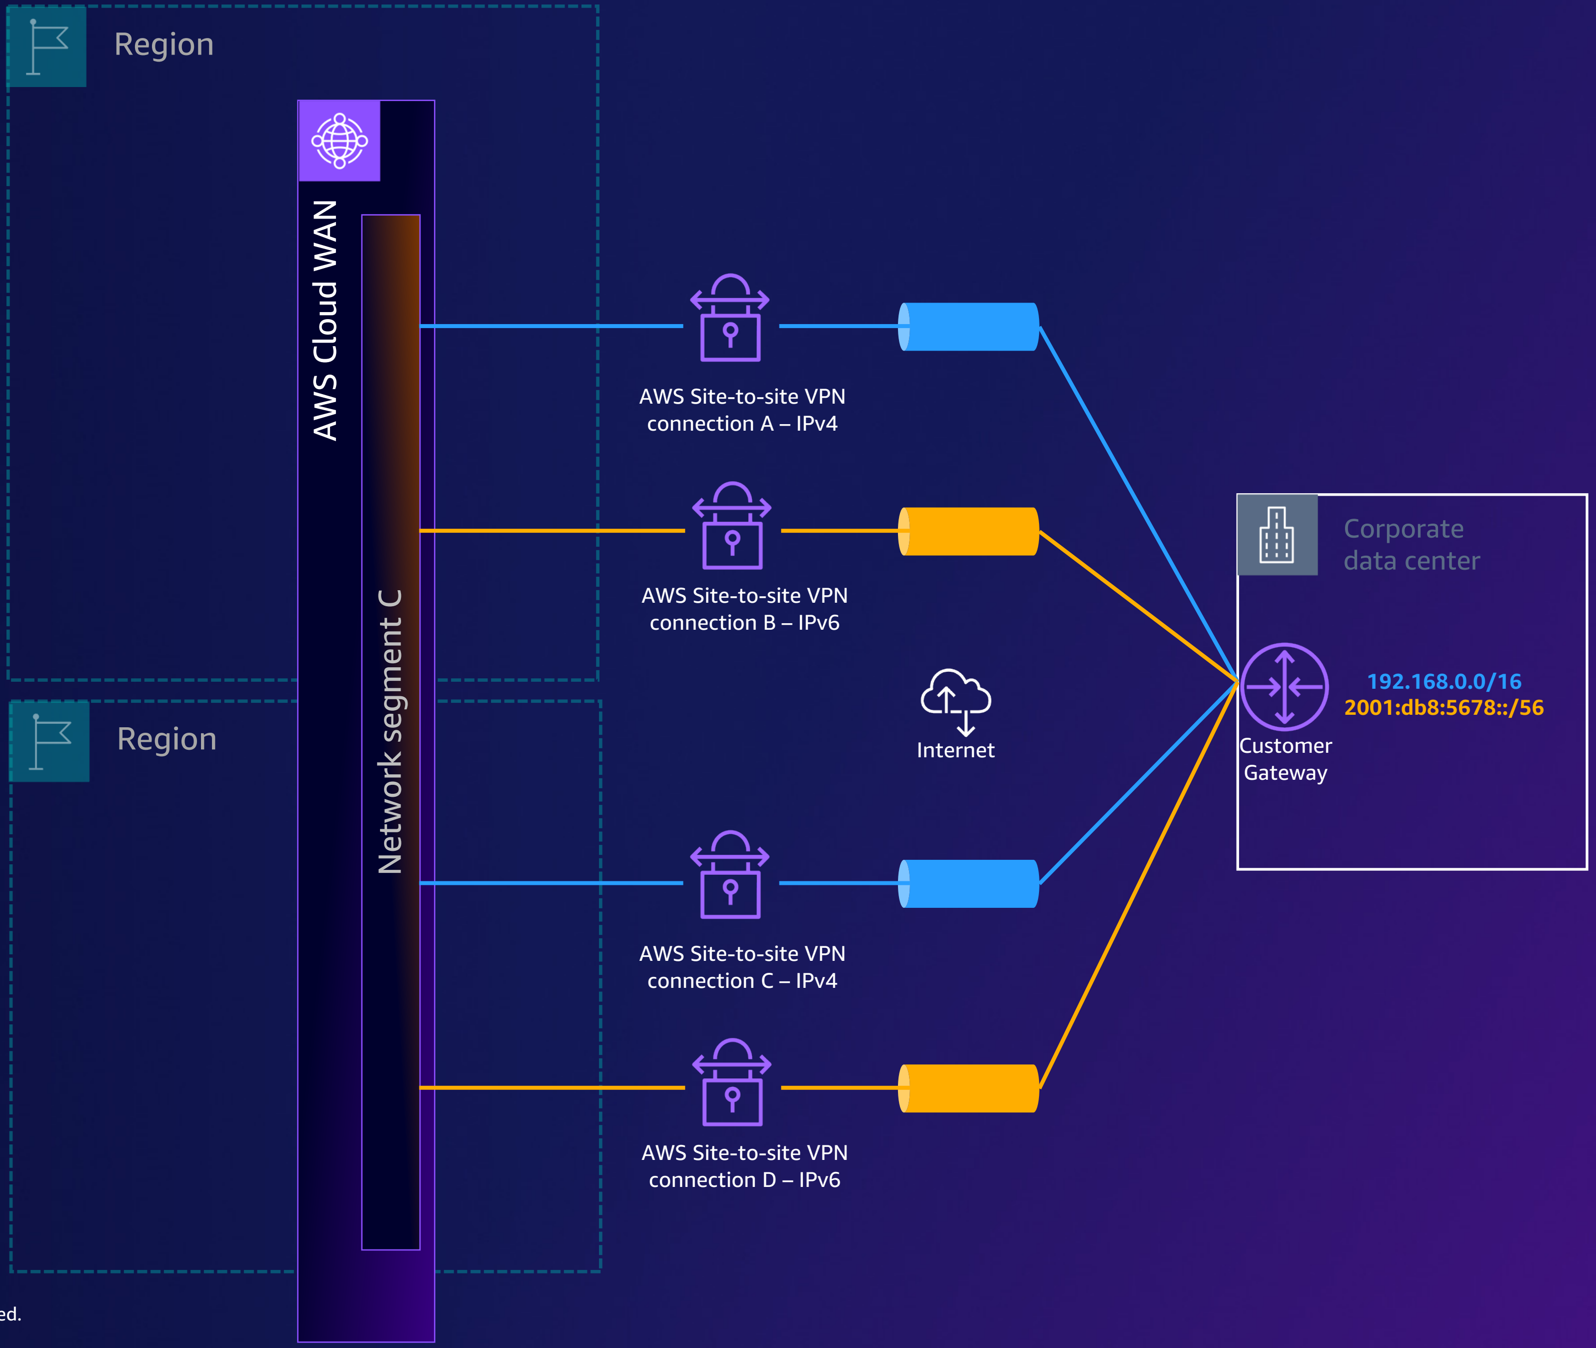
Region
AWS Cloud WAN
Network segment C
AWS Site-to-site VPN connection A – IPv4
AWS Site-to-site VPN connection B – IPv6
Internet
AWS Site-to-site VPN connection C – IPv4
AWS Site-to-site VPN connection D – IPv6
Region
Corporate data center
192.168.0.0/16
2001:db8:5678::/56
Customer Gateway

Scalable IPv6 connectivity with

AWS Application networking

IPv6 for AWS Application Networking

Elastic Load Balancing

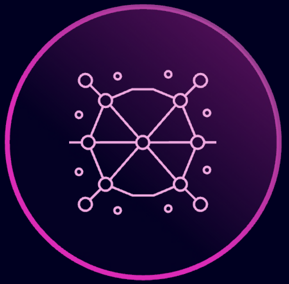

IPv6 for application networking

Elastic Load Balancing

Application Load Balancer

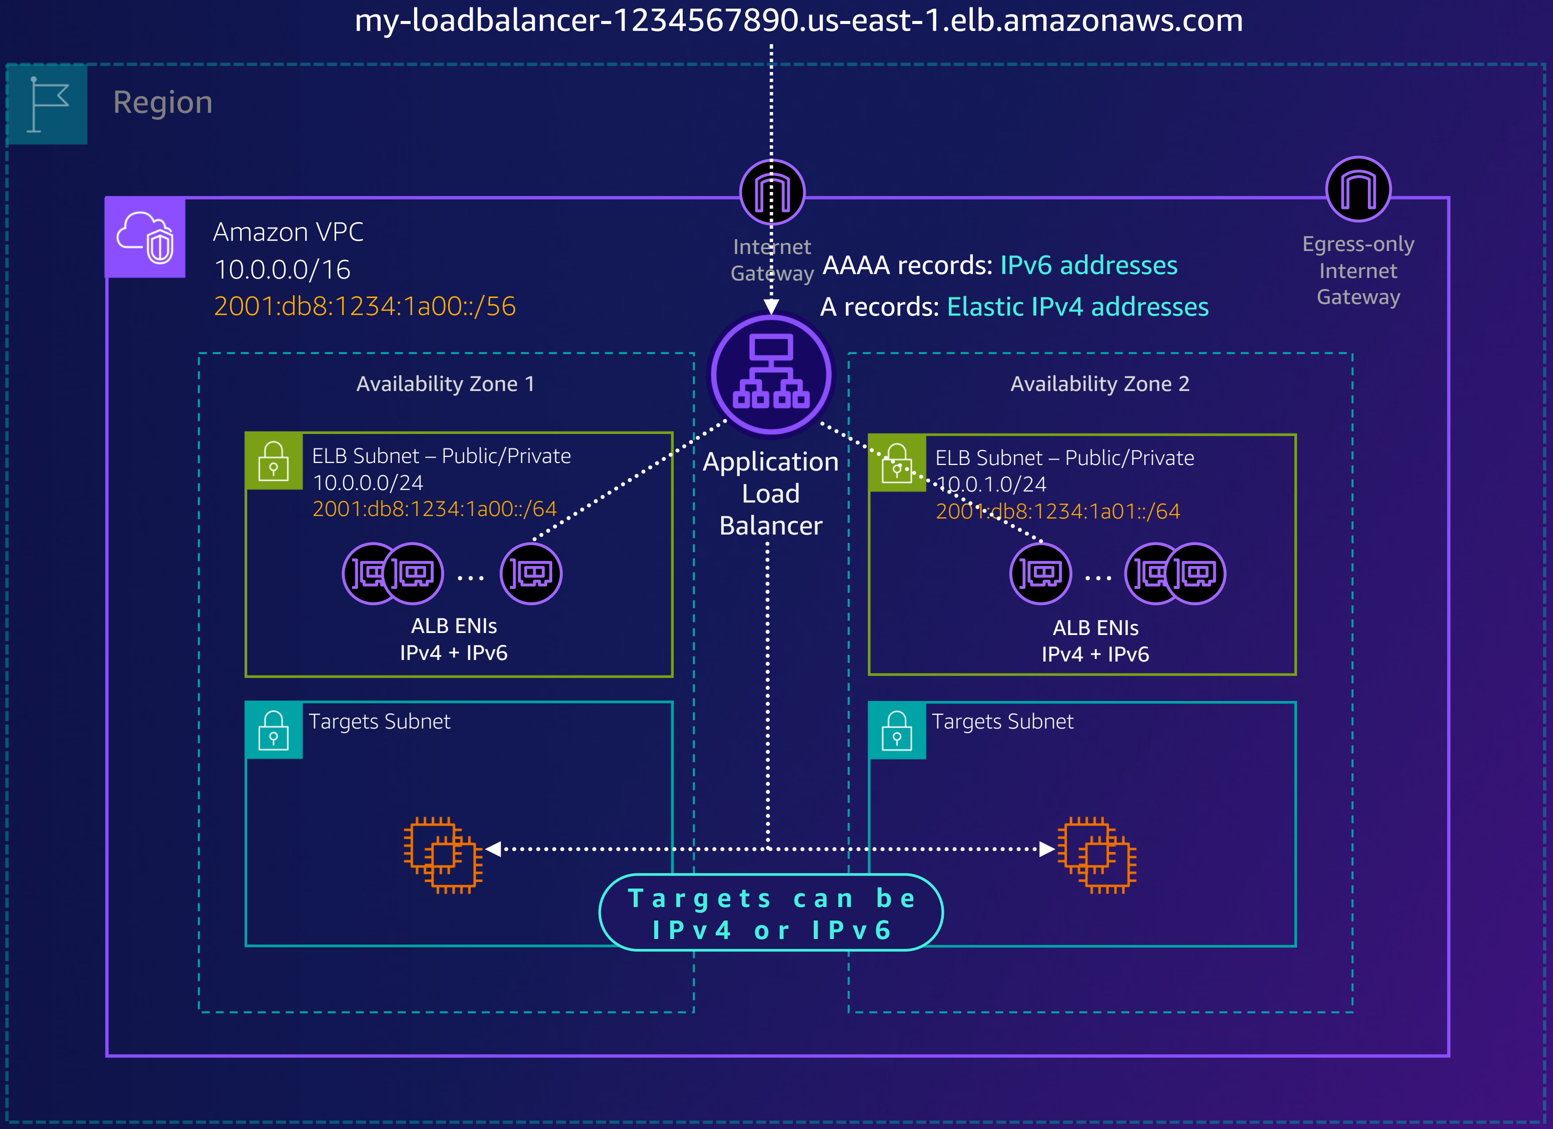

my-loadbalancer-1234567890.us-east-1.elb.amazonaws.com

AAAA records: IPv6 addresses

A records: Elastic IPv4 addresses

NEW

IPv6 for application networking

Elastic Load Balancing

IPv6-ONLY INTERNET-FACING ALB

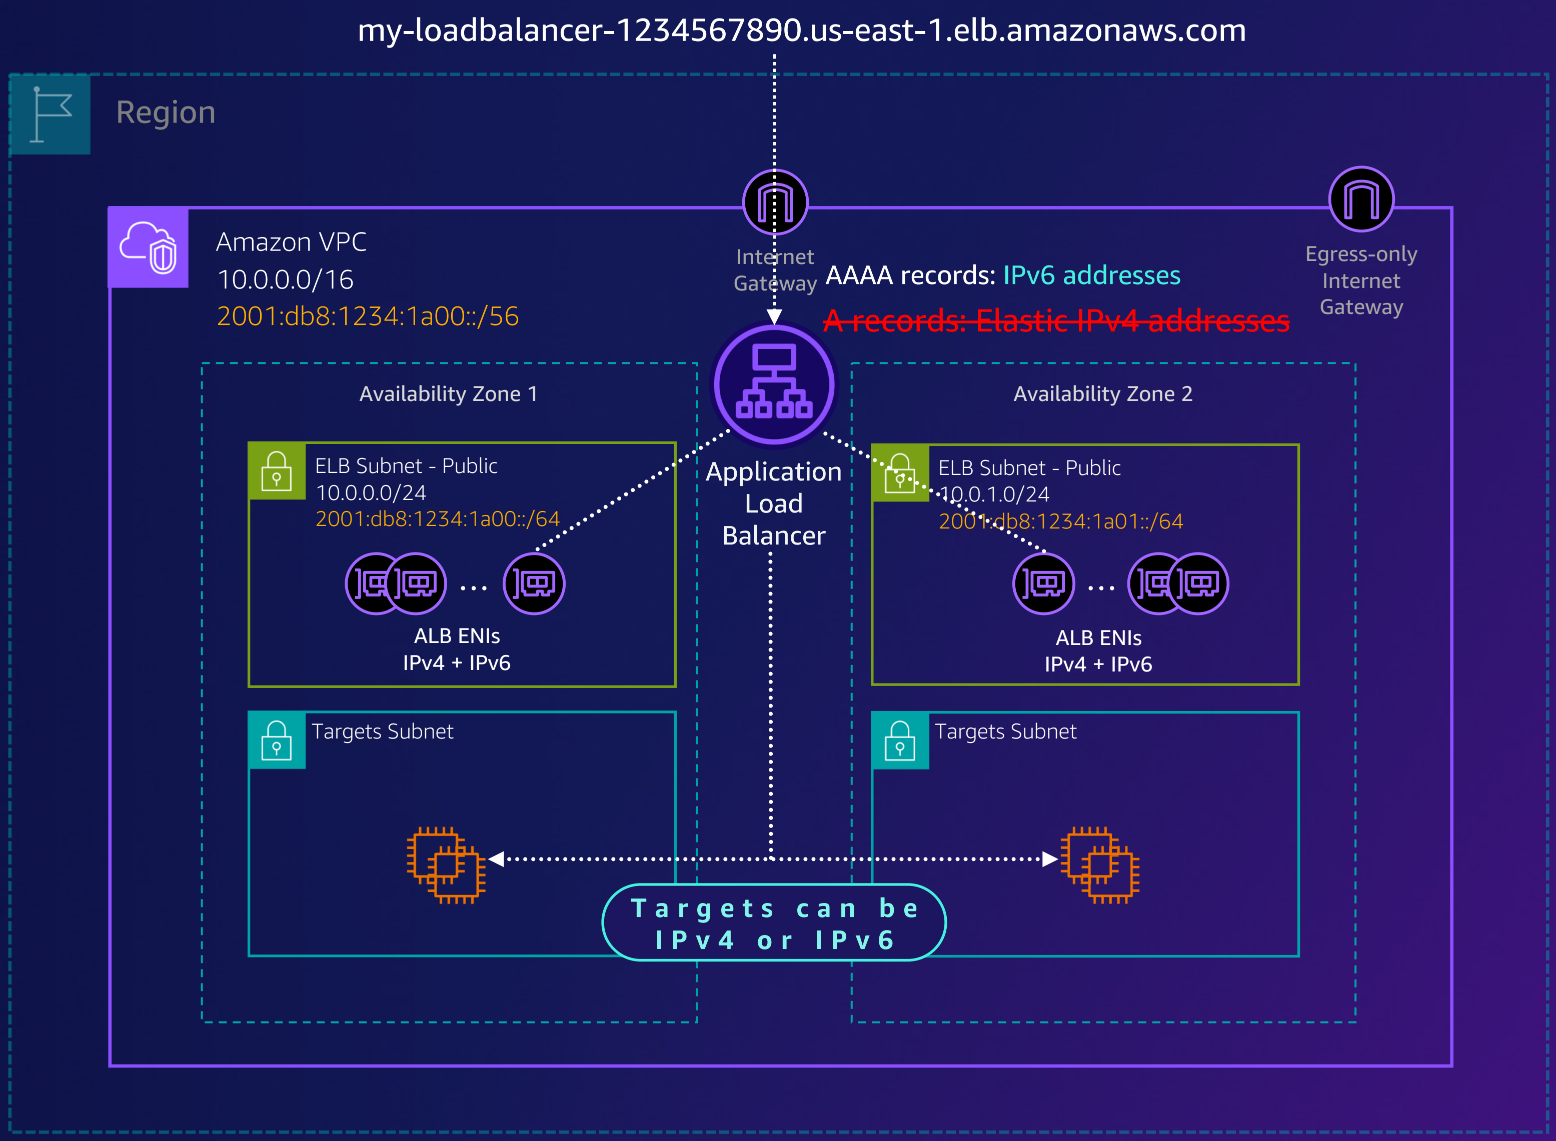

IPv6 for application networking

Elastic Load Balancing

Application Load Balancer

Network Load Balancer

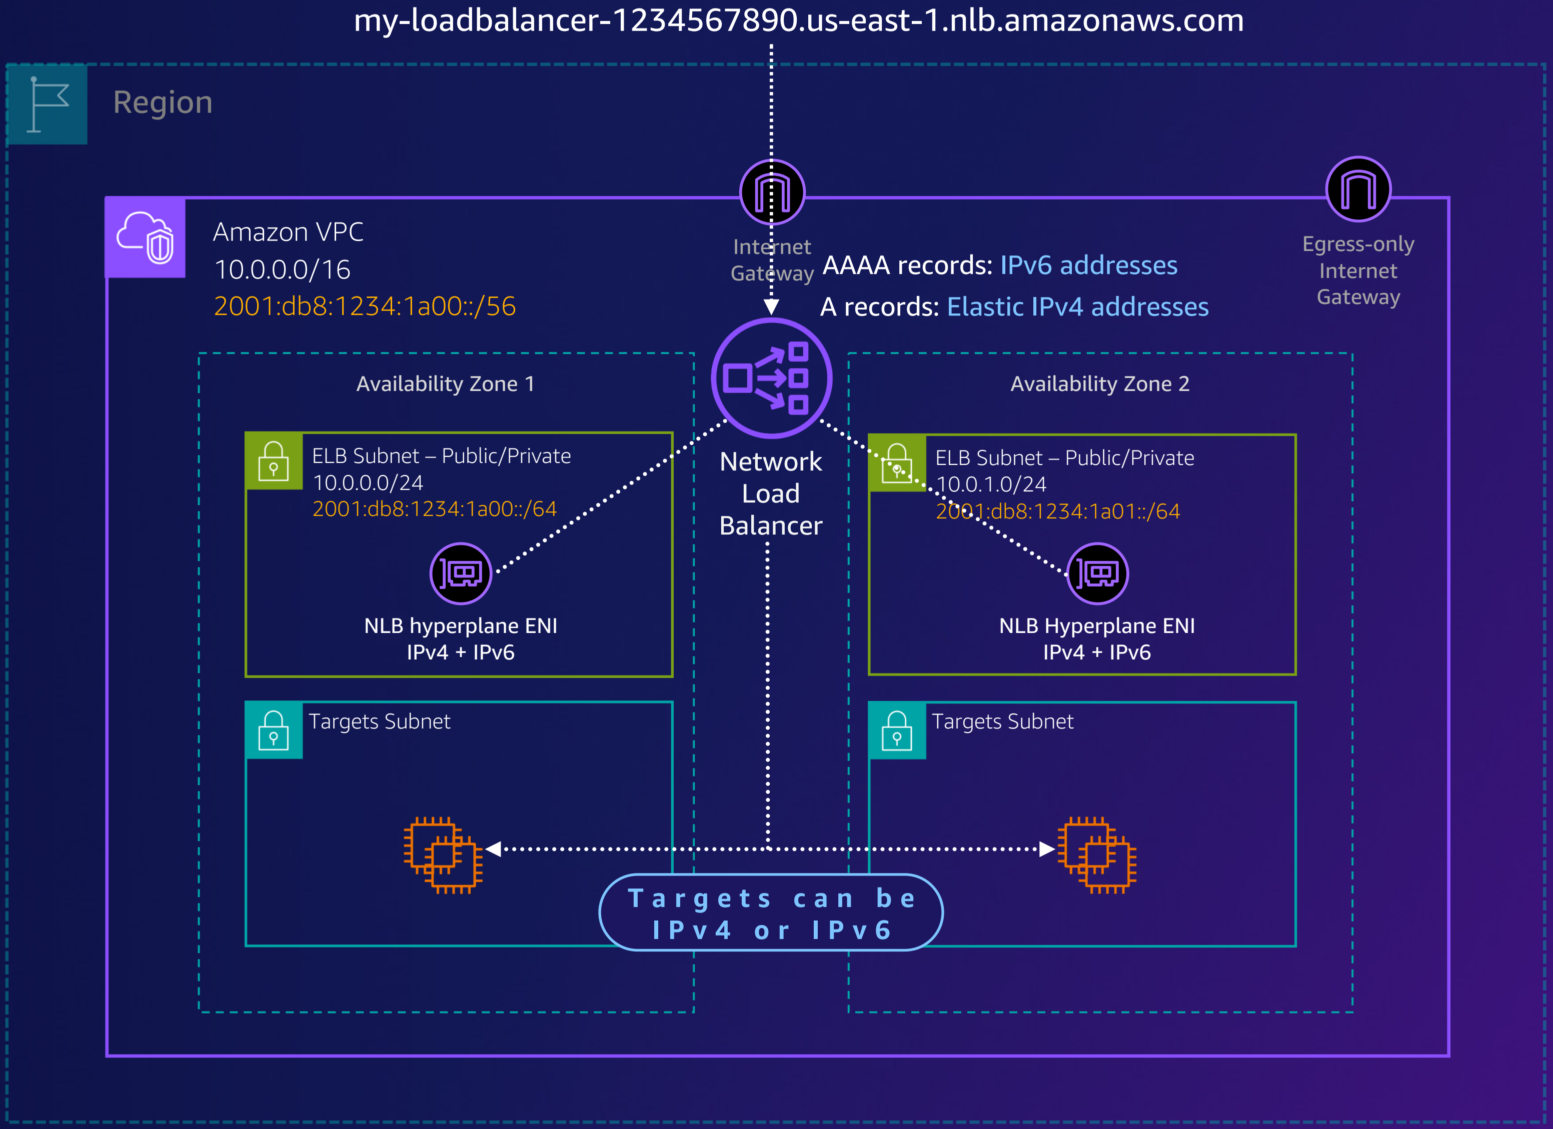

IPv6 for application networking

Elastic Load Balancing

Application Load Balancer

Network Load Balancer

Gateway Load Balancer

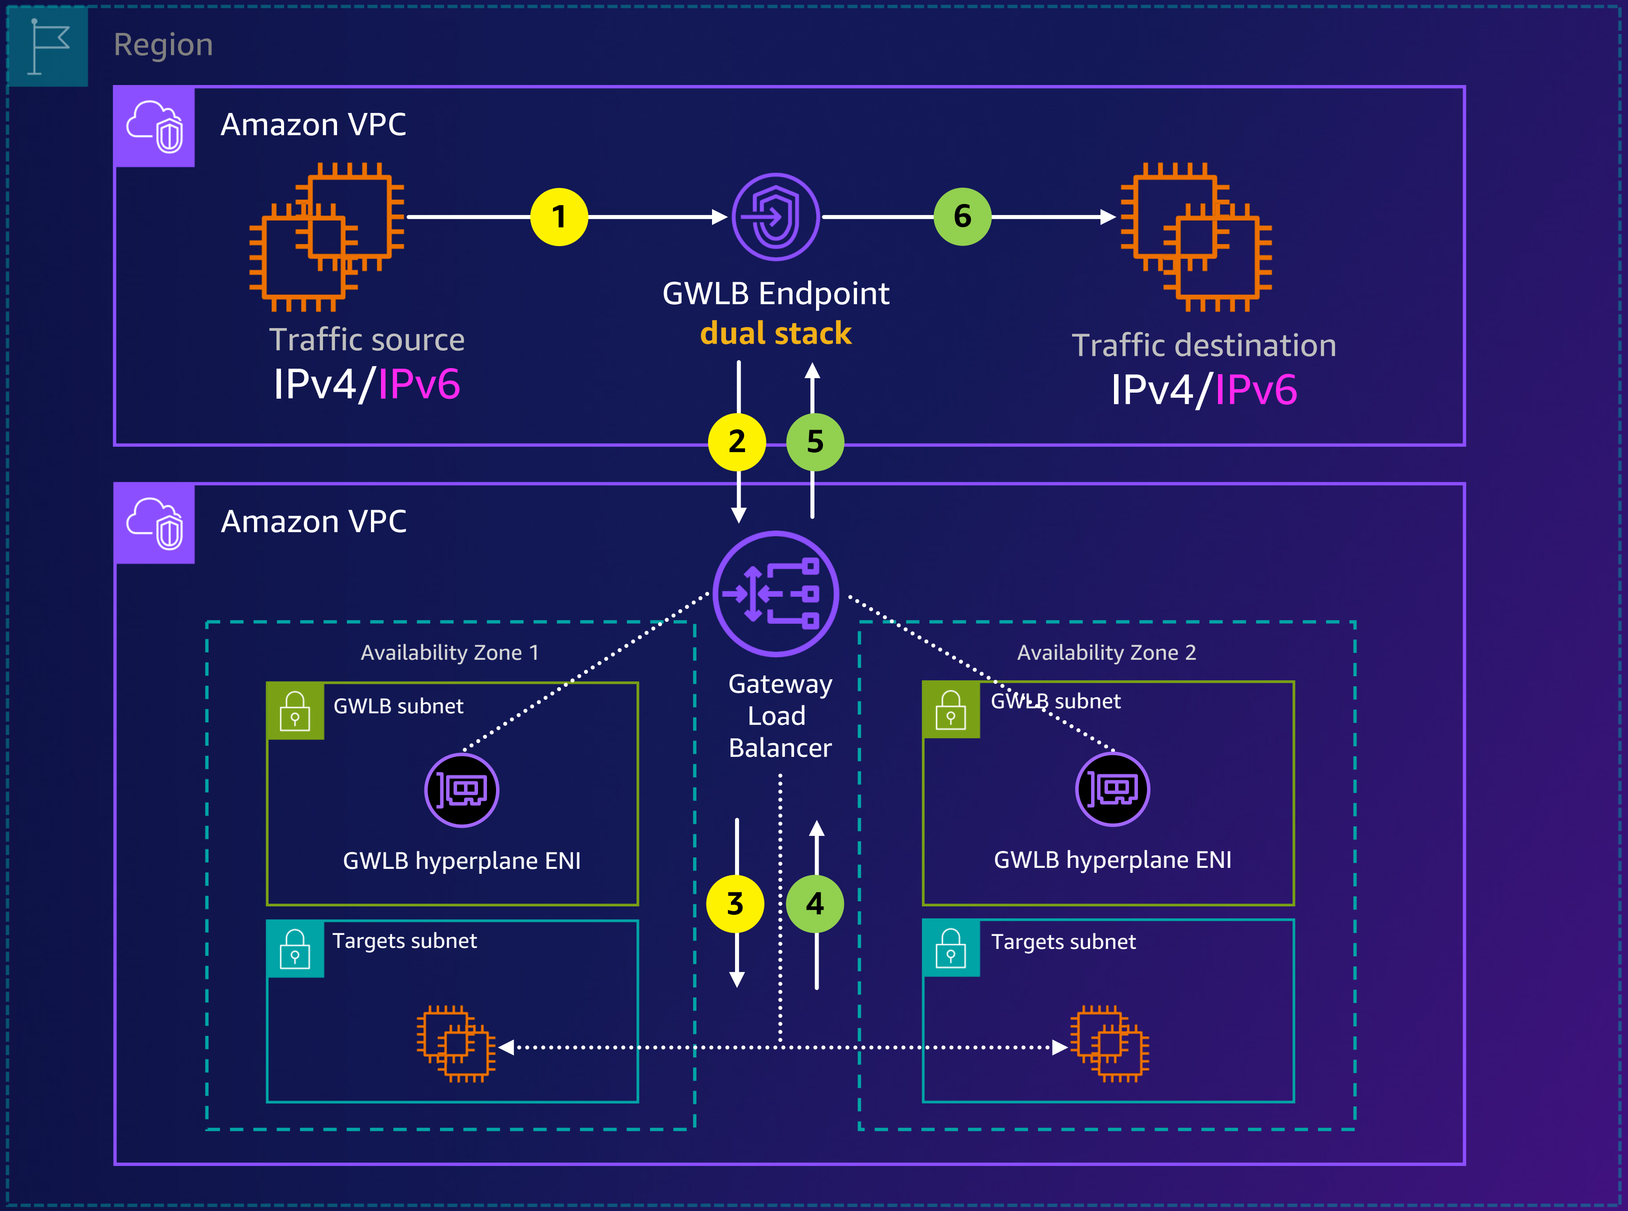

IPv6 for AWS Application Networking

Elastic Load Balancing

Amazon VPC Lattice

IPv6 for application networking

Amazon VPC Lattice

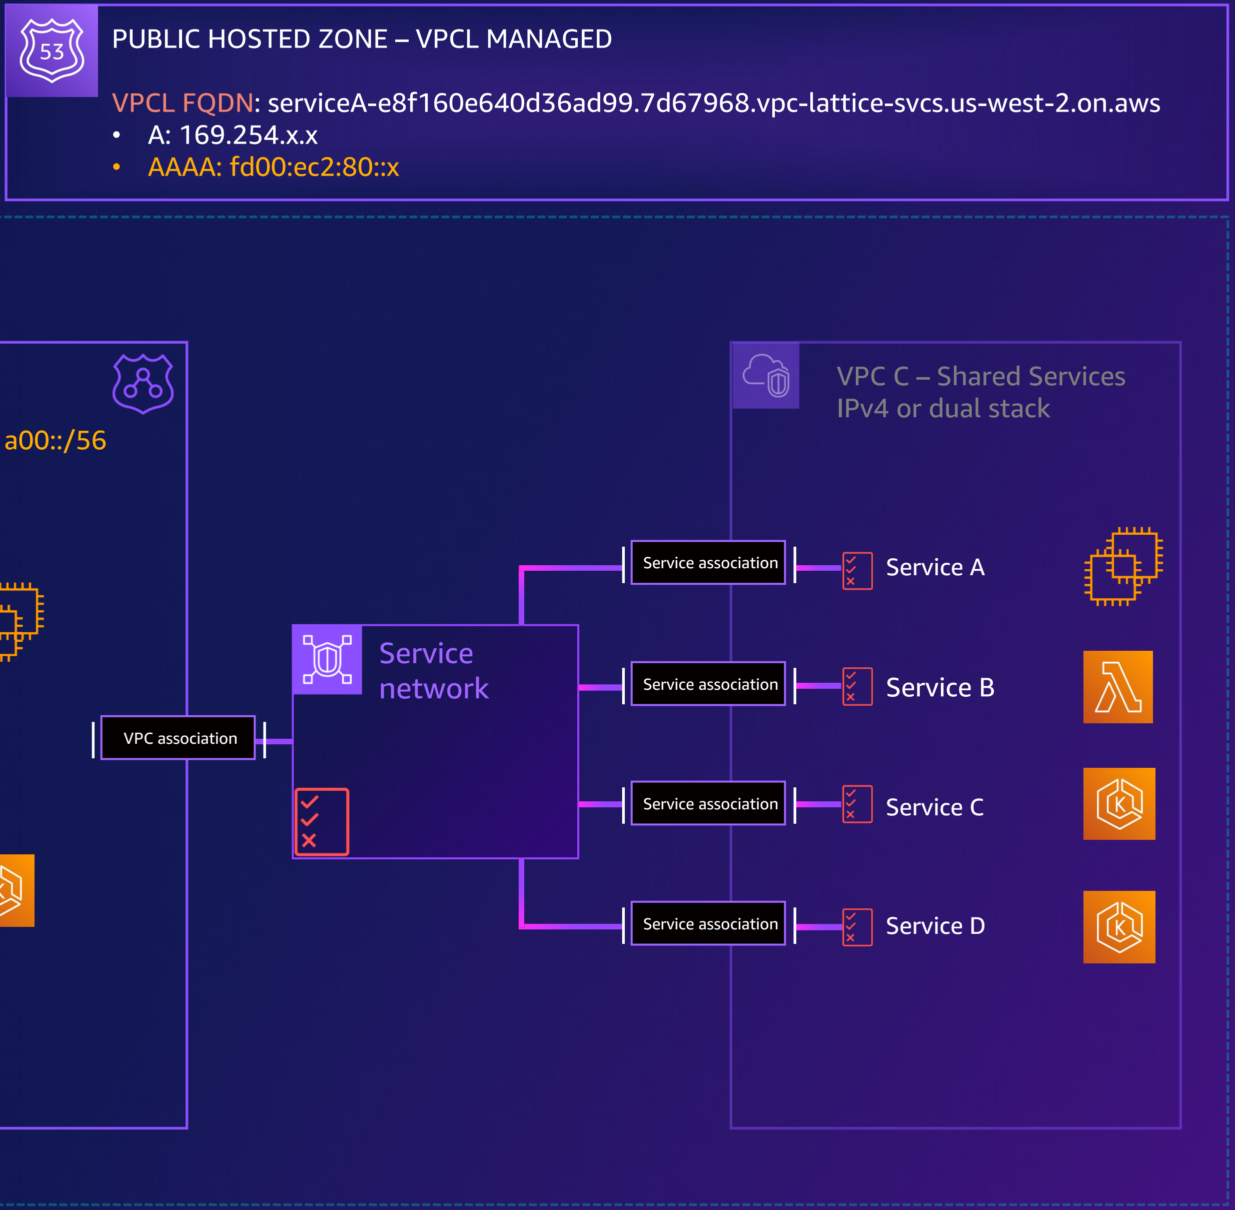

IPv6 for AWS Application Networking

Elastic Load Balancing

Amazon VPC Lattice

AWS PrivateLink

IPv6 for application networking

AWS PrivateLink

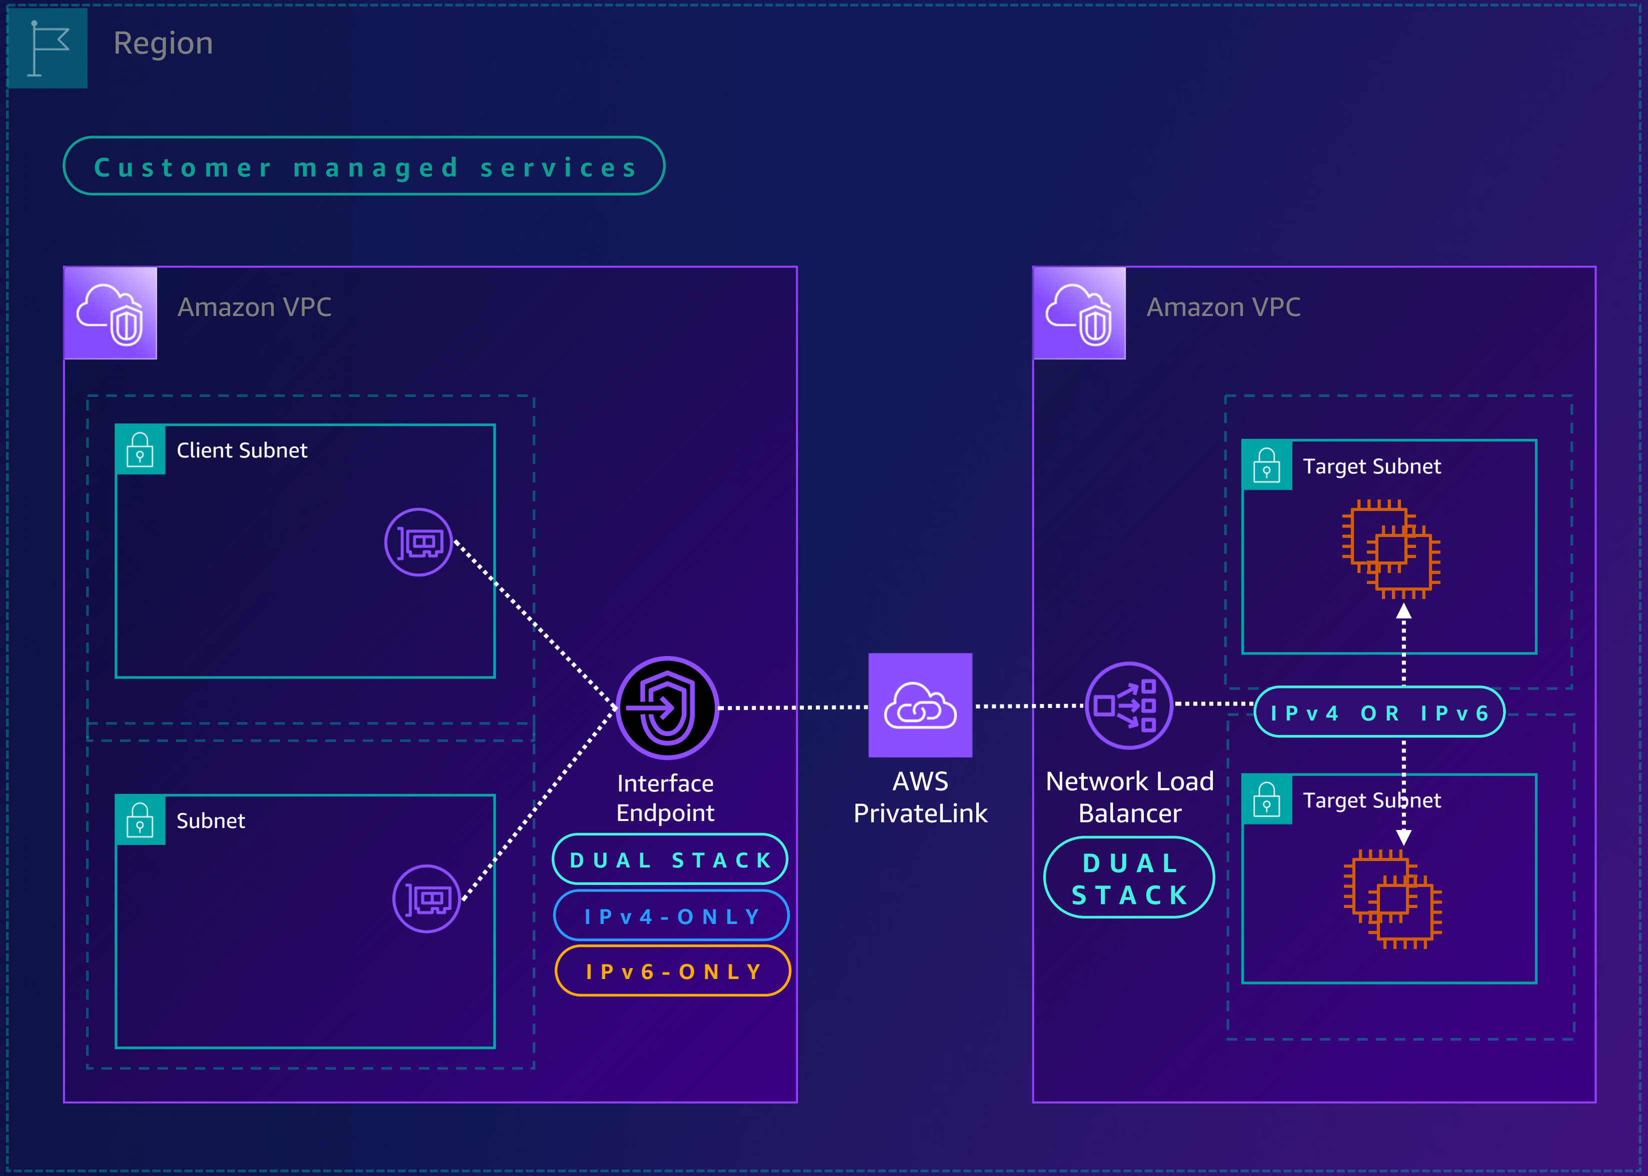

IPv6 on AWS

Secure connectivity

Secure IPv6 connectivity on AWS

VPC Network Access Control Lists

NATIVE IPv4 & IPv6

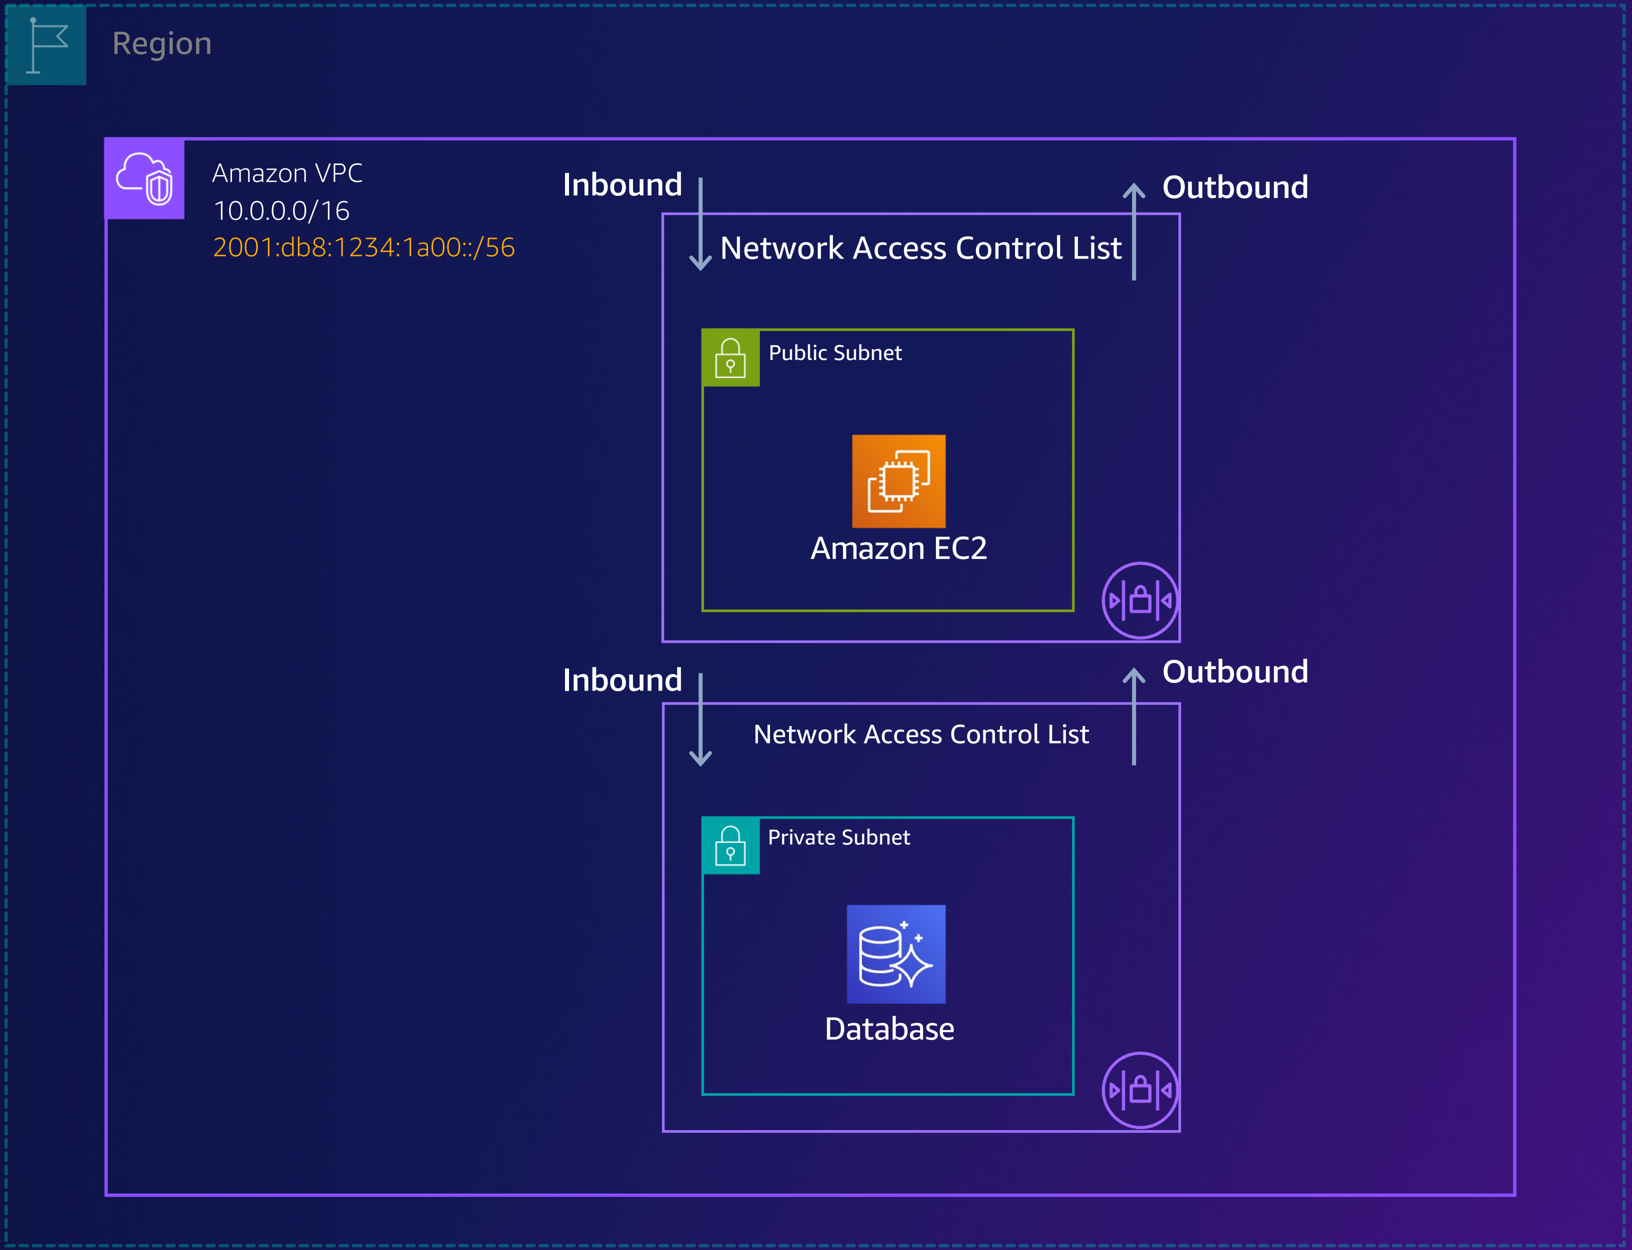

Secure IPv6 connectivity on AWS

VPC Security Groups

NATIVE IPv4 & IPv6

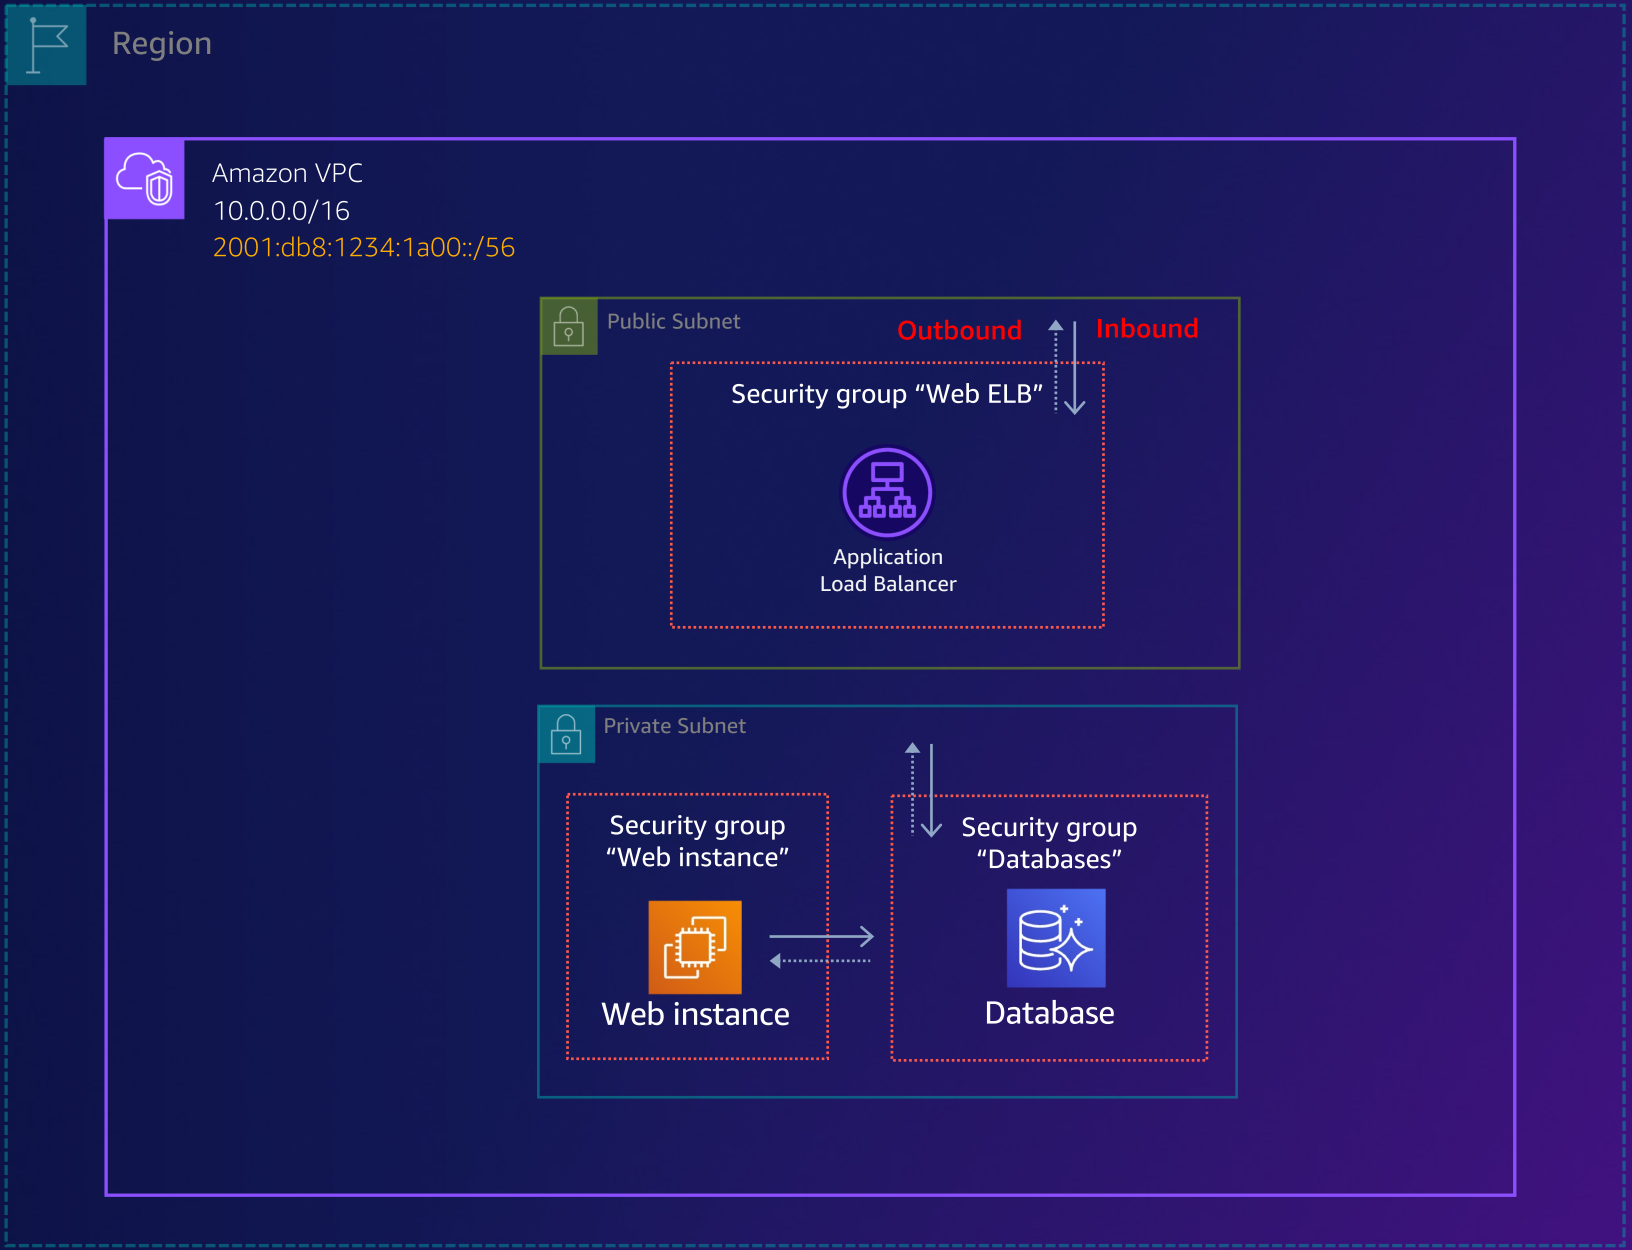

Secure IPv6 connectivity on AWS

AWS Network Firewall

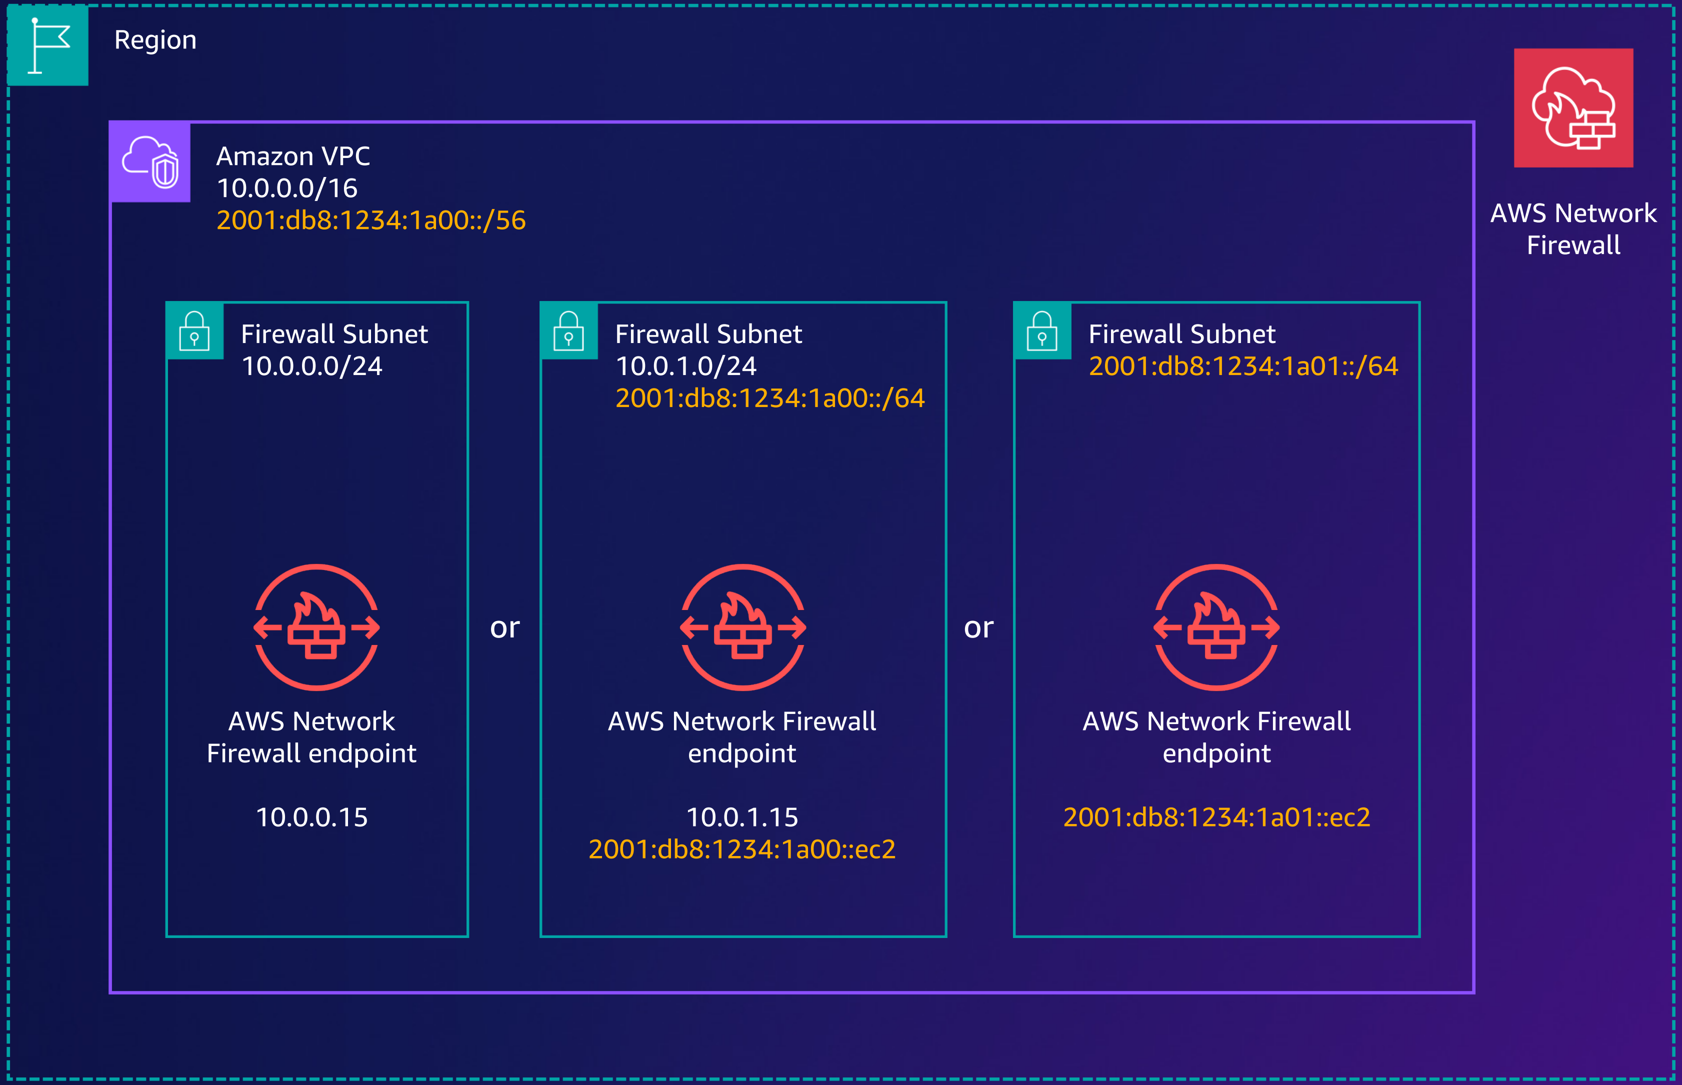

Scalable global IPv6 edge connectivity

Scalable IPv6 edge connectivity

Amazon CloudFront

Scalable IPv6 edge connectivity

Amazon CloudFront

DUAL STACK BY DEFAULT

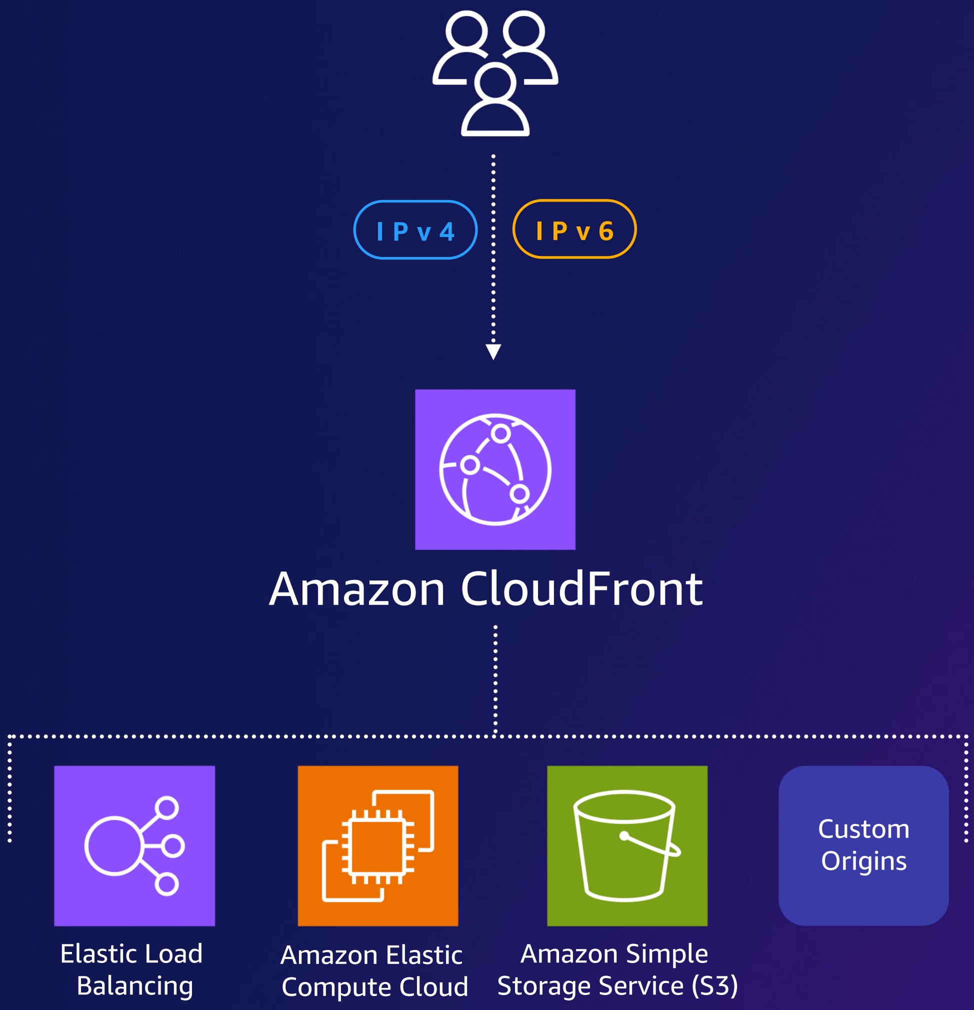
IPv4
IPv6
Amazon CloudFront
Elastic Load Balancing
Amazon Elastic Compute Cloud
Amazon Simple Storage Service (S3)
Custom Origins

Scalable IPv6 edge connectivity

Amazon CloudFront

AWS Global Accelerator

Scalable IPv6 edge connectivity

AWS Global Accelerator

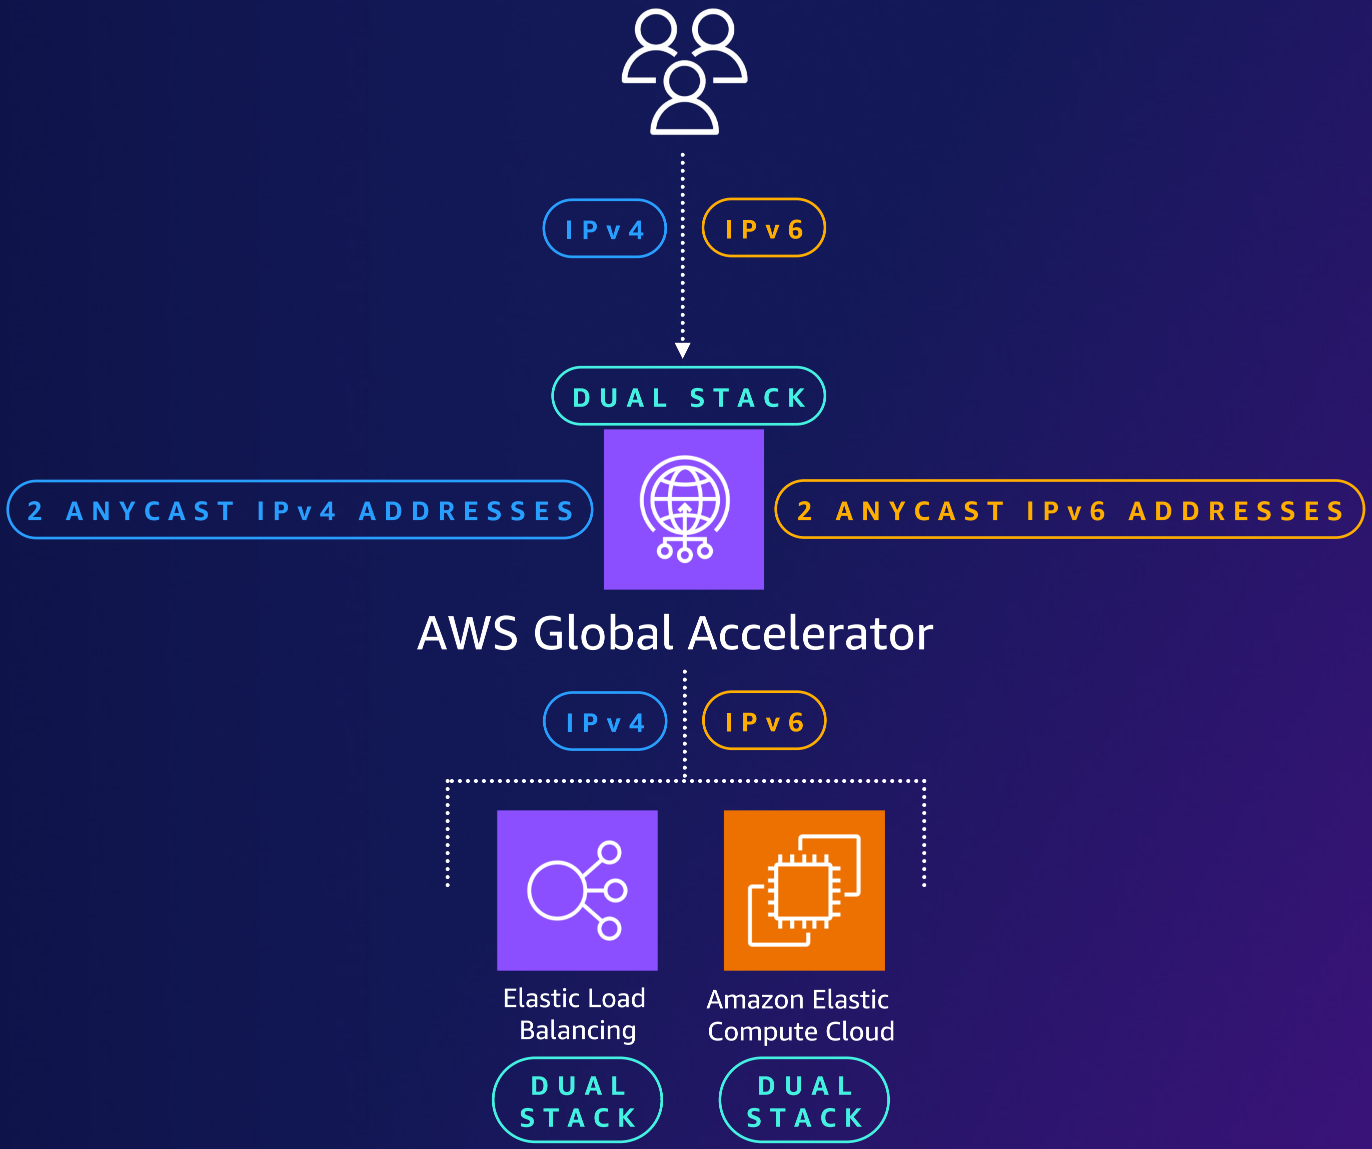

Lessons learned

IPv6 has been around for a long time	It's similar enough to IPv4 to lack adoption incentives and, at the same time, incompatible with IPv4
	Work needs to go into addressing adoption friction
	IPv6 has been around for a long time, and some lost trust on the way
Full switch to IPv6 is determined by the trailing adopters	We all still need to continue supporting IPv4
	Dual stack deployments are considered more complex but ensure backwards compatibility
	Use IPv6-only where it makes sense and allows you to scale beyond IPv4 capabilities
Starting with the business case gets creates traction	Helping leadership understand, in business terms, why IPv6 is needed is critical for resource allocation
	IPv4 exhaustion is seen as a problem for very large scale networks.
	NAT44 is ubiquitous, and also expensive.
	Addressing the fear of unknown drives progress
Creating a POC/small blast radius deployment creates confidence	Discover what works and what doesn't, and what you need to progress
	There will probably be flows that break, or unexpected application behaviors - find them early
	Help AWS work backwards, by identifying the critical services that need IPv6 support prioritized
IP Address Management is critical for scalability	IPv6 allocation is for multiple regions and environments facilitates simplified routing
	Work through the "VLSM mentality" in IPv6 address allocation
	You have full flexibility in advertising or not BYOIPv6 addresses on AWS - avoid 1-way-door decisions
Find your supporters	Success is usually driven through getting buy-in from platform teams, or shared services
	Finding the lowest hanging fruits that come with IPv6 adoption helps show the benefits
	Avoid analysis paralysis, and start TODAY

IPv6 on AWS
Service compatibility matrix

IPv6 on AWS

Start now

All resources

Thank you!