Building a Scalable and Secure Multi-VPC AWS Network Infrastructure

Building a Scalable and Secure Multi-VPC AWS Network Infrastructure: AWS Whitepaper

Table of Contents

Building a Scalable and Secure Multi-VPC AWS Network Infrastructure

Publication date: **April 17, 2024** (*Document history*)

Amazon Web Services (AWS) customers often rely on hundreds of accounts and virtual private clouds (VPCs) to segment their workloads and expand their footprint.This level of scale often creates challenges around resource sharing, inter-VPC connectivity, and on-premises facilities to VPC connectivity.

This whitepaper describes best practices for creating scalable and secure network architectures in a large network using AWS services such as Amazon Virtual Private Cloud (Amazon VPC), AWS Transit Gateway, AWS PrivateLink, AWS Direct Connect, Gateway Load Balancer, AWS Network Firewall, and Amazon Route 53. It demonstrates solutions for managing growing infrastructure—ensuring scalability, high availability, and security while keeping overhead costs low.

Introduction

AWS customers begin by building resources in a single AWS account that represents a management boundary which segments permissions, costs, and services. However, as the customer's organization grows, greater segmentation of services becomes necessary to monitor costs, control access, and provide easier environmental management. A multi-account solution solves these issues by providing specific accounts for IT services and users within an organization. AWS provides several tools to manage and configure this infrastructure, including AWS Control Tower.

AWS Control Tower initial deployment

When you set up your multi-account environment using AWS Control Tower, it creates two Organizational Units (OUs):

- **Security OU** – Within this OU, AWS Control Tower creates two accounts:
- Log Archive
- Audit (This account corresponds to the security Tooling account discussed previously in the guidance.)
- **Sandbox OU** – This OU is the default destination for accounts created within AWS Control Tower. It contains accounts in which your builders can explore and experiment with AWS services, and other tools and services, subject to your team's acceptable use policies.

AWS Control Tower allows you to create, register, and manage additional OUs to expand the initial environment to implement the guidance.

The following diagram shows the OUs initially deployed by AWS Control Tower. You can expand your AWS environment to implement any of the recommended OUs included in the diagram, to meet your requirements.
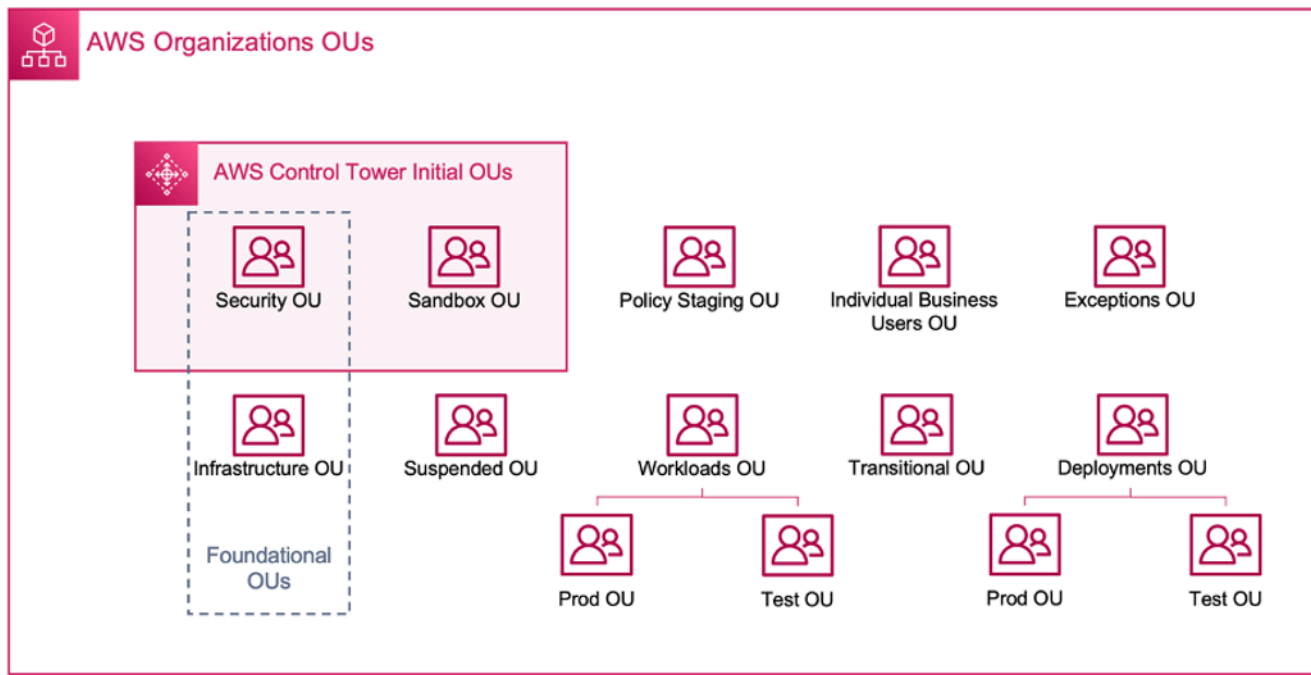
AWS organizational OUs

For further details on multi-account environment using AWS Control Tower, refer to *Appendix E* in the *Organizing Your AWS Environment Using Multiple Accounts* whitepaper.

Most customers begin with a few VPCs to deploy their infrastructure. The number of VPCs a customer creates is usually related to their number of accounts, users, and staged environments (production, development, test, and so on). As cloud usage grows, the number of users, business units, applications, and Regions that a customer interacts with also grow, leading to the creation of new VPCs.

As the number of VPCs grows, cross-VPC management becomes essential for the operation of the customer's cloud network. This whitepaper covers best practices for three specific areas in cross-VPC and hybrid connectivity:

- **Network connectivity** – Interconnecting VPCs and on-premises networks at scale.

- **Network security** – Building centralized egress points for accessing the internet and endpoints such as network address translation (NAT) gateway, VPC endpoints , AWS PrivateLink , AWS Network Firewall , and Gateway Load Balancers.

- **DNS management** – Resolving DNS within the Control Tower and hybrid DNS.

IP address planning and management

In order to build a scalable multi-account multi-VPC network design, IP address planning and management is imperative. A good IP addressing scheme needs to consider your current and future networking needs. Your IP address scheme IP needs to cover your on-prem workloads, your cloud workloads, and should also allow for future expansion (for example, addition of new AWS Regions, business units, and mergers or acquisitions). It should also prevent your teams from inadvertently creating overlapping IP CIDRs. If overlapping IP CIDR is desired such as for isolated or disconnected workloads, this decision needs to be conscious and should account for implications on routing, security and cost. You might also need to consider creating necessary approval processes for such exceptions. A good IP addressing scheme also helps with simplifying your network design and routing configuration.

Key considerations:

- Plan your IP addressing scheme (both public and private IPs) up front and select an IP address management tool to allocate, manage, and track IP address usage across all of your workloads.

- Use hierarchical and summarized IP addressing schemes.

- Plan for consistent IP assignment based on environment, AWS Region, organization, or business unit.

- Designate distinct IP CIDRs (both IPv4 and IPv6) for on-premises and cloud networks.

- Proactively prevent and track overlapping IP CIDRs.

- Size your IP CIDRs appropriately to enable scaling and future growth.

- Enable your workloads for IPv6 or dual-stack compatibility to reduce IP conflicts and address IPv4 space depletion.

You can use Amazon VPC IP Address Manager (IPAM) to simplify planning, tracking, and monitoring both public and private IP addresses for your AWS workloads. IPAM allows you to organize, allocate, monitor, and share IP address space across multiple AWS Regions and AWS accounts. It also helps with automatic allocation of CIDRs to VPCs using specific business rules.

Refer to the [Amazon VPC IP Address Manager Best Practices](#), [Managing IP pools across VPCs and Regions using Amazon VPC IP Address Manager](#), and [IP Address Management for AWS Control Tower](#) blog posts to learn IP addressing best practices and how to use IPAM to manage IP pools across VPCs, AWS Regions, and AWS Control Tower.

Are you Well-Architected?

The AWS Well-Architected Framework helps you understand the pros and cons of the decisions you make when building systems in the cloud. The six pillars of the Framework allow you to learn architectural best practices for designing and operating reliable, secure, efficient, cost-effective, and sustainable systems. Using the AWS Well-Architected Tool, available at no charge in the AWS Management Console, you can review your workloads against these best practices by answering a set of questions for each pillar.

For more expert guidance and best practices for your cloud architecture—reference architecture deployments, diagrams, and whitepapers—refer to the AWS Architecture Center.

VPC to VPC connectivity

Customers can use two different VPC connectivity patterns to set up multi-VPC environments: *many to many*, or *hub and spoke*. In the many-to-many approach, the traffic between each VPC is managed individually between each VPC. In the hub-and-spoke model, all inter-VPC traffic flows through a central resource, which routes traffic based on established rules.

VPC peering

The first way to connect two VPCs is to use VPC peering. In this setup, a connection enables full bidirectional connectivity between the VPCs. This peering connection is used to route traffic between the VPCs. VPCs in different accounts and AWS Regions can also be peered together. All data transfer over a VPC peering connection that stays within an Availability Zone is free. All data transfer over a VPC peering connection that crosses Availability Zones is charged at the standard in-region data transfer rates. If the VPCs are peered across Regions, standard inter-Region data transfer charges will apply.

VPC peering is point-to-point connectivity, and it does not support transitive routing. For example, if you have a VPC peering connection between VPC A and VPC B and between VPC A and VPC C, an instance in VPC B cannot transit through VPC A to reach VPC C. To route packets between VPC B and VPC C, you are required to create a direct VPC peering connection.

At scale, when you have tens or hundreds of VPCs, interconnecting them with peering can result in a mesh of hundreds or thousands of peering connections. A large number of connections can be difficult to manage and scale. For example, if you have 100 VPCs and you want to setup a full mesh peering between them, it will take 4,950 peering connections $[n(n-1)/2]$ where n is the total number of VPCs. There is a maximum limit of 125 active peering connections per VPC.

Network setup using VPC peering

If you are using VPC peering, on-premises connectivity (VPN and/or Direct Connect) must be made to each VPC. Resources in a VPC cannot reach on-premises using the hybrid connectivity of a peered VPC, as shown in the preceding figure.

VPC peering is best used when resources in one VPC must communicate with resources in another VPC, the environment of both VPCs is controlled and secured, and the number of VPCs to be connected is less than 10 (to allow for the individual management of each connection). VPC peering offers the lowest overall cost and highest aggregate performance when compared to other options for inter-VPC connectivity.

AWS Transit Gateway

AWS Transit Gateway provides a hub and spoke design for connecting VPCs and on-premises networks as a fully managed service without requiring you to provision third-party virtual appliances. No VPN overlay is required, and AWS manages high availability and scalability.

Transit Gateway enables customers to connect thousands of VPCs. You can attach all your hybrid connectivity (VPN and Direct Connect connections) to a single gateway, consolidating and controlling your organization's entire AWS routing configuration in one place (refer to the following figure). Transit Gateway controls how traffic is routed among all the connected spoke networks using route tables. This hub-and-spoke model simplifies management and reduces operational

costs because VPCs only connect to the Transit Gateway instance to gain access to the connected networks.
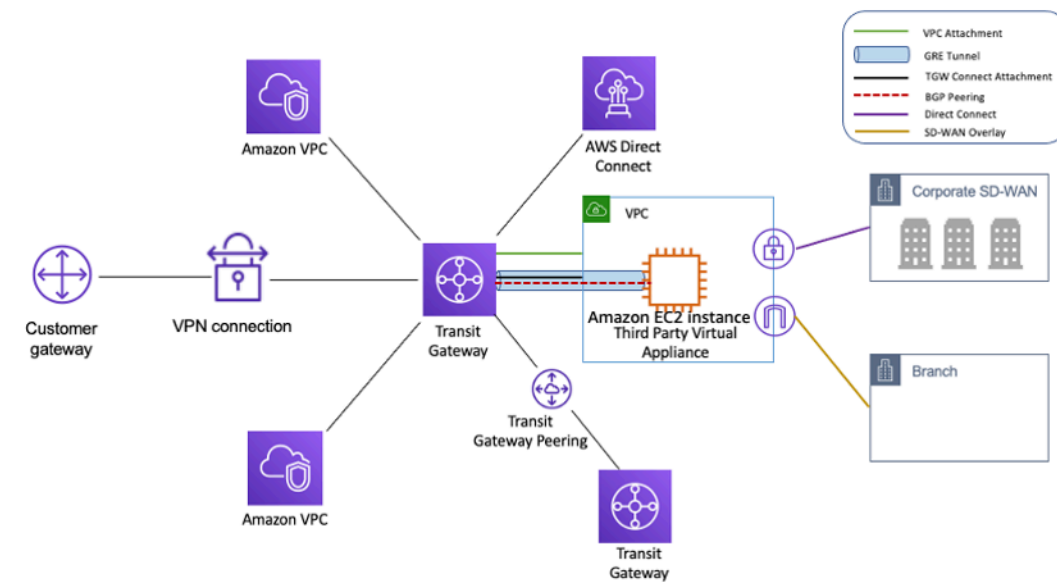
Hub and spoke design with AWS Transit Gateway

Transit Gateway is a Regional resource and can connect thousands of VPCs within the same AWS Region. You can connect multiple gateways over a single Direct Connect connection for hybrid connectivity. Typically, you can use just one Transit Gateway instance connecting all your VPC instances in a given Region, and use Transit Gateway routing tables to isolate them wherever needed. Note that you do not need additional transit gateways for high availability, because transit gateways are highly available by design; for redundancy, use a single gateway in each Region. However, there is a valid case for creating multiple gateways to limit misconfiguration blast radius, segregate control plane operations, and administrative ease-of-use.

With Transit Gateway peering, customers can peer their Transit Gateway instances within same or multiple Regions and route traffic between them. It uses the same underlying infrastructure as VPC peering, and is therefore encrypted. For more information, refer to Building a global network using AWS Transit Gateway Inter-Region peering and AWS Transit Gateway now supports Intra-Region Peering.

Place your organization's Transit Gateway instance in its Network Services account. This enables centralized management by network engineers who manage the Network services account. Use AWS Resource Access Manager (RAM) to share a Transit Gateway instance for connecting VPCs across multiple accounts in your AWS Organization within the same Region. AWS RAM enables you to easily and securely share AWS resources with any AWS account, or within your AWS

Organization. For more information, refer to the [Automating AWS Transit Gateway attachments to a transit gateway in a central account](#) blog post.

Transit Gateway also allows you to establish connectivity between SD-WAN infrastructure and AWS using Transit Gateway Connect. Use a Transit Gateway Connect attachment with Border Gateway Protocol (BGP) for dynamic routing and Generic Routing Encapsulation (GRE) tunnel protocol for high performance, delivering up to 20 Gbps total bandwidth per Connect attachment (up to four Transit Gateway Connect peers per Connect attachment). By using Transit Gateway Connect, you can integrate both on-premises SD-WAN infrastructure or SD-WAN appliances running in the cloud through a VPC attachment or AWS Direct Connect attachment as the underlying transport layer. Refer to [Simplify SD-WAN connectivity with AWS Transit Gateway Connect](#) for reference architectures and detailed configuration.

Transit VPC solution

[Transit VPCs](#) can create connectivity between VPCs by a different means than VPC peering by introducing a hub and spoke design for inter-VPC connectivity. In a transit VPC network, one central VPC (the hub VPC) connects with every other VPC (spoke VPC) through a VPN connection typically leveraging BGP over [IPsec](#). The central VPC contains [Amazon Elastic Compute Cloud](#) (Amazon EC2) instances running software appliances that route incoming traffic to their destinations using the VPN overlay. Transit VPC peering has the following advantages:

- Transitive routing is enabled using the overlay VPN network — allowing for a hub and spoke design.

- When using third-party vendor software on the EC2 instance in the hub transit VPC, vendor functionality around advanced security (layer 7 firewall/Intrusion Prevention System (IPS)/ Intrusion Detection System (IDS)) can be used. If customers are using the same software on-premises, they benefit from a unified operational/monitoring experience.

- The Transit VPC architecture enables connectivity that may be desired in some use-cases. For example, you can connect an AWS GovCloud instance and Commercial Region VPC or a Transit Gateway instance to a Transit VPC and enable inter-VPC connectivity between the two Regions. Evaluate your security and compliance requirements when considering this option. For additional security, you may deploy a centralized inspection model using design patterns described later in this whitepaper.

Transit VPC with virtual appliances

Transit VPC comes with its own challenges, such as higher costs for running third-party vendor virtual appliances on EC2 based on the instance size/family, limited throughput per VPN connection (up to 1.25 Gbps per VPN tunnel), and additional configuration, management and resiliency overhead (customers are responsible for managing the HA and redundancy of EC2 instances running the third-party vendor virtual appliances).

VPC peering vs. Transit VPC vs. Transit Gateway

Table 1 — Connectivity comparison

Criteria	VPC peering	Transit VPC	Transit Gateway	PrivateLink	Cloud WAN	VPC Lattice
Scope	Regional/ Global	Regional	Regional	Regional	Global	Regional
Architecture	Full mesh	VPN-based hub-and-spoke	Attachments-based hub-and-spoke	Provider or Consumer Model	Attachments based, multi-region	App to App connectivity
Scale	125 active Peers/VPC	Depends on virtual router/EC2	5000 attachments per Region	No limits	5000 attachments per core network	500 VPC associations per service

| Criteria	VPC peering	Transit VPC	Transit Gateway	PrivateLink	Cloud WAN	VPC Lattice
Segmentation	Security groups	Customer managed	Transit Gateway route tables	No segmentation	Segments	Servie and service network policies
Latency	Lowest	Extra, due to VPN encryption overhead	Additional Transit Gateway hop	Traffic stays on AWS backbone, customers should test	Uses the same dataplane as Transit Gateway	Traffic stays on AWS backbone, customers should test
Bandwidth limit	Per instance limits, no aggregate limit	Subject to EC2 instance bandwidth limits based on size/family	Up to 100 Gbps (burst)/attachment	10 Gbps per Availability Zone, automatically scales up to 100 Gbps	Up to 100 Gbps (burst)/attachment	10 Gbps per Availability Zone
Visibility	VPC Flow Logs	VPC Flow Logs and CloudWatch Metrics	Transit Gateway Network Manager, VPC Flow Logs, CloudWatch Metrics	CloudWatch Metrics	Network Manager, VPC Flow Logs, CloudWatch Metrics	CloudWatch Access Logs
Security group cross-referencing	Supported	Not supported	Not supported	Not supported	Not supported	Not applicable

Criteria	VPC peering	Transit VPC	Transit Gateway	PrivateLink	Cloud WAN	VPC Lattice
IPv6 support	Supported	Depends on virtual appliance	Supported	Supported	Supported	Supported

AWS PrivateLink

AWS PrivateLink provides private connectivity between VPCs, AWS services, and your on-premises networks without exposing your traffic to the public internet. Interface VPC endpoints, powered by AWS PrivateLink, make it easy to connect to AWS and other services across different accounts and VPCs to significantly simplify your network architecture. This allows customers who may want to privately expose a service/application residing in one VPC (service provider) to other VPCs (consumer) within an AWS Region in a way that only consumer VPCs initiate connections to the service provider VPC. An example of this is the ability for your private applications to access service provider APIs.

To use AWS PrivateLink, create a Network Load Balancer for your application in your VPC, and create a VPC endpoint service configuration pointing to that load balancer. A service consumer then creates an interface endpoint to your service. This creates an elastic network interface (ENI) in the consumer subnet with a private IP address that serves as an entry point for traffic destined for the service. The consumer and service are not required to be in the same VPC. If the VPC is different, the consumer and service provider VPCs can have overlapping IP address ranges. In addition to creating the interface VPC endpoint to access services in other VPCs, you can create interface VPC endpoints to privately access supported AWS services through AWS PrivateLink, as shown in the following figure.

With Application Load Balancer (ALB) as target of NLB, you can now combine ALB advanced routing capabilities with AWS PrivateLink. Refer to Application Load Balancer-type Target Group for Network Load Balancer for reference architectures and detailed configuration.

AWS PrivateLink for connectivity to other VPCs and AWS Services

The choice between Transit Gateway, VPC peering, and AWS PrivateLink is dependent on connectivity.

- **AWS PrivateLink** — Use AWS PrivateLink when you have a client/server set up where you want to allow one or more consumer VPCs unidirectional access to a specific service or set of instances in the service provider VPC or certain AWS services. Only the clients with access in the consumer VPC can initiate a connection to the service in the service provider VPC or AWS service. This is also a good option when client and servers in the two VPCs have overlapping IP addresses because AWS PrivateLink uses ENIs within the client VPC in a manner that ensures that are no IP conflicts with the service provider. You can access AWS PrivateLink endpoints over VPC peering, VPN, Transit Gateway, Cloud WAN, and AWS Direct Connect.
- **VPC peering and Transit Gateway** — Use VPC peering and Transit Gateway when you want to enable layer-3 IP connectivity between VPCs.

Your architecture will contain a mix of these technologies in order to fulfill different use cases. All of these services can be combined and operated with each other. For example, AWS PrivateLink handling API style client-server connectivity, VPC peering for handling direct connectivity requirements where placement groups may still be desired within the Region or inter-Region connectivity is needed, and Transit Gateway to simplify connectivity of VPCs at scale as well as edge consolidation for hybrid connectivity.

VPC sharing

Sharing VPCs is useful when network isolation between teams does not need to be strictly managed by the VPC owner, but the account level users and permissions must be. With Shared VPC, multiple AWS accounts create their application resources (such as Amazon EC2 instances) in shared, centrally managed Amazon VPCs. In this model, the account that owns the VPC (owner) shares one or more subnets with other accounts (participants). After a subnet is shared, the participants can view, create, modify, and delete their application resources in the subnets shared with them. Participants cannot view, modify, or delete resources that belong to other participants or the VPC owner. Security between resources in shared VPCs is managed using security groups, network access control lists (NACLs), or through a firewall between the subnets.

VPC sharing benefits:

- Simplified design — no complexity around inter-VPC connectivity

- Fewer managed VPCs

- Segregation of duties between network teams and application owners

- Better IPv4 address utilization

- Lower costs — no data transfer charges between instances belonging to different accounts within the same Availability Zone

> ⓘ **Note**
>
> When you share a subnet with multiple accounts, your participants should have some level of cooperation since they're sharing IP space and network resources. If necessary, you can choose to share a different subnet for each participant account. One subnet per participant enables network ACL to provide network isolation in addition to security groups.

Most customer architectures will contain multiple VPCs, many of which will be shared with two or more accounts. Transit Gateway and VPC peering can be used to connect the shared VPCs. For example, suppose you have 10 applications. Each application requires its own AWS account. The apps can be categorized into two application portfolios (apps within the same portfolio have similar networking requirements, App 1–5 in 'Marketing' and App 6–10 in 'Sales').

You can have one VPC per application portfolio (two VPCs total), and the VPC is shared with the different application owner accounts within that portfolio. App owners deploy apps into their respective shared VPC (in this case, in the different subnets for network route segmentation and isolation using NACLs). The two shared VPCs are connected via the Transit Gateway. With this setup, you could go from having to connect 10 VPCs to just two, as seen in the following figure.
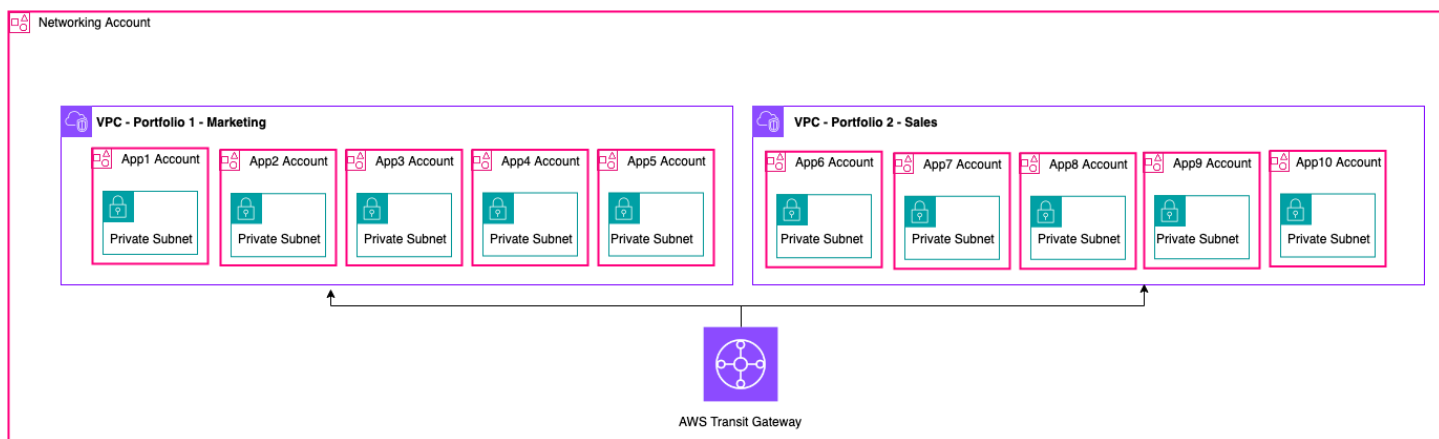
Example setup – shared VPC

> **Note**
>
> VPC sharing participants cannot create all AWS resources in a shared subnet. For more information, refer to the Limitations section in the VPC Sharing documentation.
> For more information about the key considerations and best practices for VPC sharing, refer to the VPC sharing: key considerations and best practices blog post.

Private NAT Gateway

Teams often work independently and they might create a new VPC for a project, which may have overlapping classless inter-domain routing (CIDR) blocks. For integration, they might want to enable communication between networks with overlapping CIDRs, which is not achievable through features such as VPC peering and Transit Gateway. A private NAT gateway can help with this

use case. Private NAT gateway uses a unique private IP address to perform source NAT for the overlapping source IP address, and ELB does the destination NAT for the overlapping destination IP address. You can route traffic from your private NAT gateway to other VPCs or on-premises networks using Transit Gateway or a virtual private gateway.
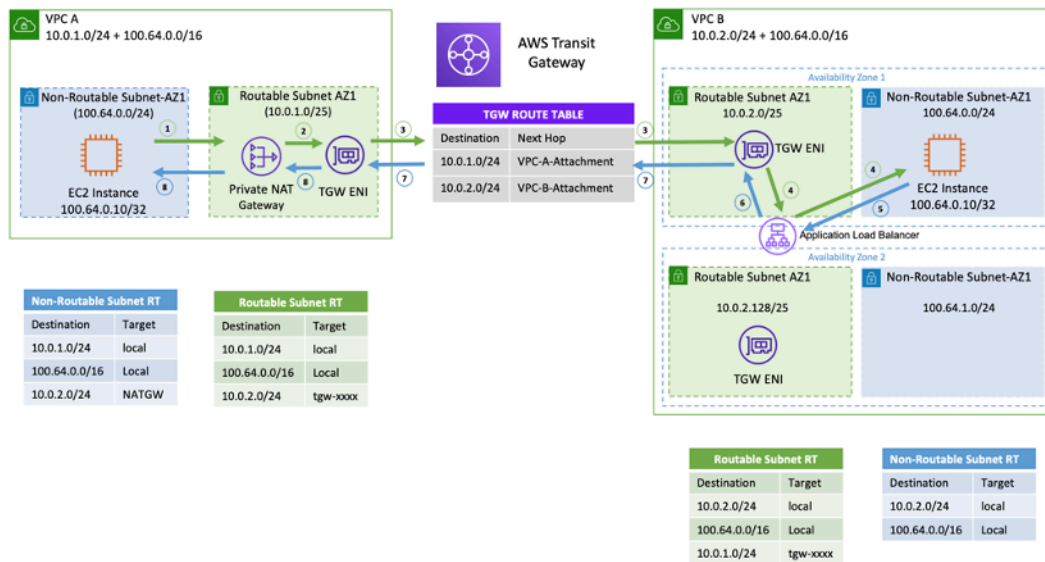
Example setup – Private NAT gateway

The preceding figure shows two non-routable (overlapping CIDRs, `100.64.0.0/16`) subnets in VPC A and B. To establish a connection between them, you can add secondary non-overlapping/ routable CIDRs (routable subnets, `10.0.1.0/24` and `10.0.2.0/24`) to VPC A and B, respectively. The routable CIDRs should be allocated by the network management team responsible for IP allocation. A private NAT gateway is added to the routable subnet in VPC A with an IP address of `10.0.1.125`. The private NAT gateway performs source network address translation on requests from instances in non-routable subnet of VPC A (`100.64.0.10`) as `10.0.1.125`, the ENI of the private NAT gateway. Now traffic can be pointed to a routable IP address assigned to the Application Load Balancer (ALB) in VPC B (`10.0.2.10`), which has a target of `100.64.0.10`. Traffic is routed through Transit Gateway. Return traffic is processed by the private NAT gateway back to the original Amazon EC2 instance requesting the connection.

The private NAT gateway can also be used when your on-premises network restricts access to approved IPs. The on-premises networks of few customers are required by compliance to communicate only with private networks (no IGW) only through a limited contiguous block of approved IPs owned by the customer. Instead of allocating each instance a separate IP from the block, you can run large workloads on AWS VPCs behind each allow-listed IP using private NAT

gateway. For details, refer to the [How to solve Private IP exhaustion with Private NAT Solution](#) blog post.

Example setup – How to use Private NAT gateway to provide approved IPs for on-premises network

AWS Cloud WAN

AWS Cloud WAN is a new way to connect networks together that we were previously able to do with Transit Gateways, VPC Peering, and IPSEC VPN tunnels. Previously you would configure one or more VPCs, connect them together with one of the previous methods, and use IPSEC VPN or AWS Direct Connect to connect to on-premises networks. You would have your network and security posture constructs defined in one place, and your networks in another. Cloud WAN allows you to centralize all of these constructs in a single place. By policy, you can segment your networks to determine who can talk to who, and isolate production traffic via these segments from development or test workloads, or your on premises networks.
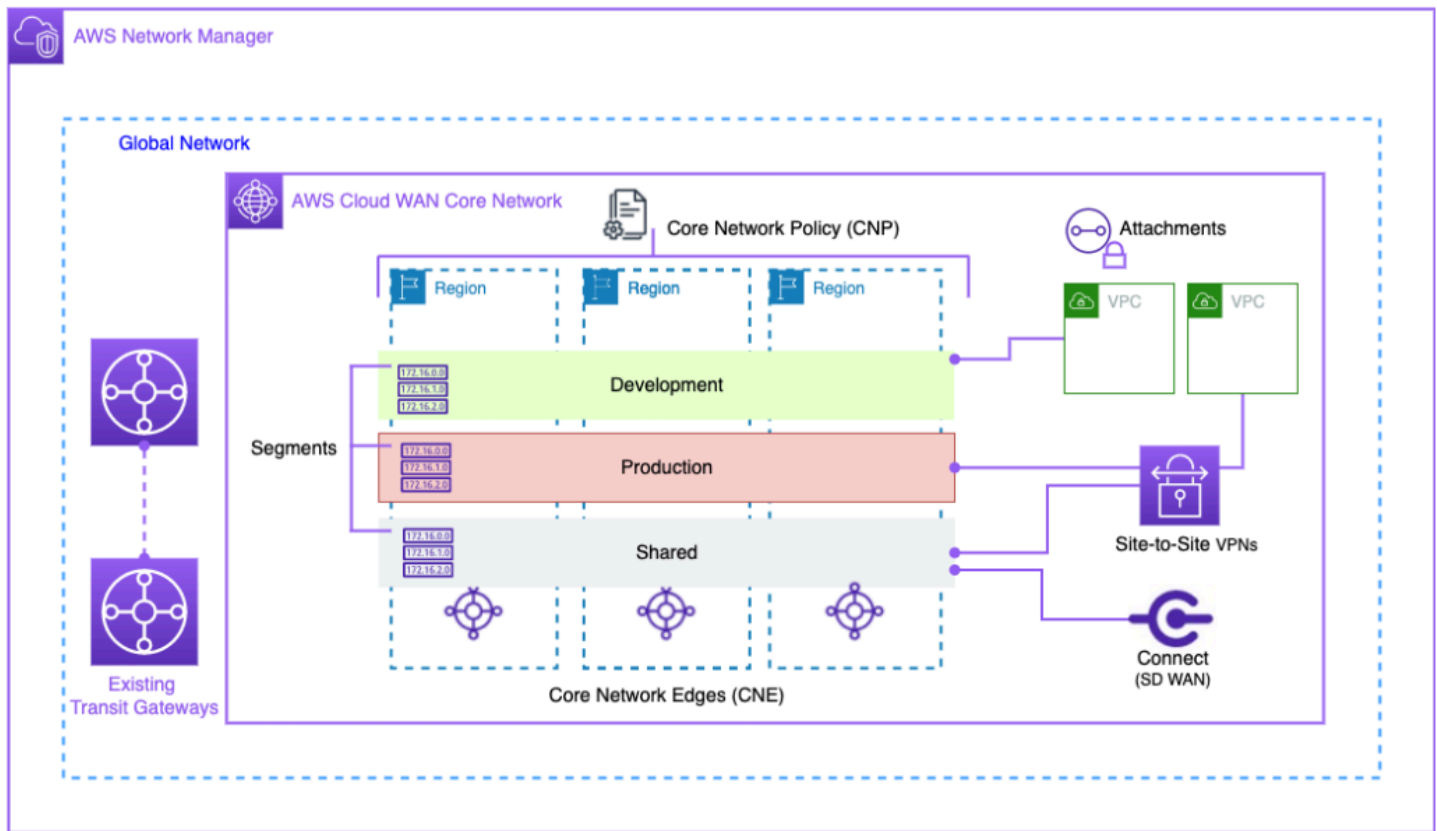
Cloud WAN Block Diagram

Manage your global network through the AWS Network Manager user interface and APIs. The global network is the root-level container for all your network objects; the core network is the part of your global network managed by AWS. A core network policy (CNP) is a single versioned policy document that defines all aspects of your core network. Attachments are any connections or resources you wand to add to your core network. A core network edge (CNE) is a local connection point for attachments that comply with the policy. Network segments are routing domains which, by default, allow communication only within a segment.

To use CloudWAN:

1. In AWS Network Manager, create a global network and associated core network.

2. Create a CNP that defines segments, ASN range, AWS Regions and tags to be used to attach to segments.

3. Apply the network policy.

4. Share the core network with your users, accounts, or organizations using the resource access manager.

5. Create and tag attachments.

6. Update routes in your attached VPCs to include the core network.

Cloud WAN was designed to simplify the process of connecting your AWS infrastructure globally. It allows you to segment traffic with a centralized permissions policy and use your existing infrastructure at your company locations. Cloud WAN also connects your VPCs, SD-WANs, Client VPNs, firewalls, VPNs, and data center resources to connect to Cloud WAN. For more information, see AWS Cloud WAN blog posts.

AWS Cloud WAN enables a unified network connecting cloud and on-premises environments. Organizations use next-gen firewalls (NGFWs) and intrusion prevention systems (IPSs) for security. The AWS Cloud WAN and Transit Gateway migration and interoperability patterns blog post describes architectural patterns for centrally managing and inspecting outbound network traffic in a Cloud WAN network, including single-Region and multi-Region networks, and configures route tables. These architectures ensure data and applications remain safe while maintaining a secure cloud environment.

For more information about Cloud WAN, see the Centralized outbound inspection architecture in AWS Cloud WAN blog post.

Amazon VPC Lattice

Amazon VPC Lattice is a fully managed application networking service that is used to connect, monitor, and secure services across various accounts and virtual private clouds. VPC Lattice helps to interconnect services within a logical boundary, so that you can manage and discover them efficiently.

VPC Lattice components consists of:

- **Service** - This is a unit of application running on an instance, a container, or a Lambda function and consists of listeners, rules and target groups.

- **Service network** - This is the logical boundary that is used to automatically implement service discovery and connectivity and apply common access and observability policies to a collection of services.

- **Auth policies** - IAM resource policies that can be associated with a service network or individual services to support request-level authentication and context-specific authorization.

- **Service Directory** – A centralized view of the services that you own or that have been shared with you through the AWS Resource Access Manager.

VPC Lattice usage steps:

1. Create the service network. The service network usually resides on a network account where a network administrator has full access. The service network can be shared across multiple accounts within an organization. Sharing can be performed on individual services or the entire service account.

2. Attach VPCs to the service network to enable application networking for each VPC, so that different services can start consuming other services that are registered within the network. Security groups are applied to control traffic.

3. Developers define the services, which are populated in the service directory and registered into the service network. VPC Lattice contains the address book of all services configured. Developers can also define routing polices to use blue/green deployments. Security is managed at the service network level where authentication and authorization policies are defined and at the service level where access policies with IAM are implemented.

VPC Lattice communication flows

More details can be found in the [VPC Lattice user guide](#).

Hybrid connectivity

This section focuses on securely connecting your cloud resources with your on-premises data centers. There are three approaches for enabling hybrid connectivity:

- **One-to-one connectivity** — In this setup, a VPN connection and/or Direct Connect private VIF is created for every VPC. This is accomplished by using the virtual private gateway (VGW). This option is great for small numbers of VPCs, but as a customer scales their VPCs, managing hybrid connectivity per VPC can become difficult.

- **Edge consolidation** — In this setup, customers consolidate hybrid IT connectivity for multiple VPCs at a single endpoint. All the VPCs share these hybrid connections. This is accomplished by using AWS Transit Gateway and the AWS Direct Connect gateway.

- **Full mesh hybrid consolidation** — In this setup, customers consolidate connectivity for multiple VPCs at a single endpoint using CloudWAN, built on AWS Transit Gateway. This is a full policy-based approach to networking in one or more AWS accounts, represented in code. At this time, using AWS Direct Connect for edge connectivity requires peering Transit Gateway to CloudWAN.

VPN

There are various ways to set up VPN to AWS:

AWS VPN options

- **Option 1: Consolidate VPN connectivity on Transit Gateway** — This option leverages the Transit Gateway VPN attachment on Transit Gateway. Transit Gateway supports IPsec termination for site-to-site VPN. Customers can create VPN tunnels to the Transit Gateway, and can access the VPCs attached to it.Transit Gateway supports both Static and BGP-based Dynamic VPN connections. Transit Gateway also supports Equal-Cost Multi-Path (ECMP) on VPN attachments. Each VPN connection has a maximum of 1.25-Gbps throughput per tunnel. Enabling ECMP allows you to aggregate throughput across VPN connections, allowing to scale beyond the default maximum limit of 1.25 Gbps. In this option, you pay for Transit Gateway pricing as well as AWS VPN pricing. AWS recommendeds using this option for VPN connectivity. For more information, refer to the Scaling VPN throughput using AWS Transit Gateway blog post.

- **Option 2: Terminate VPN on Amazon EC2 instance** — This option is leveraged by customers in edge cases when they want a particular vendor software feature set (such as Cisco DMVPN or Generic Routing Encapsulation (GRE)), or they want operational consistency across various VPN deployments. You can use the transit VPC design for edge consolidation, but it is important to remember that all the key considerations from the *VPC to VPC connectivity* section for transit VPC are applicable to hybrid VPN connectivity. You are responsible for managing high availability, and you pay for EC2 instance as well as any vendor software licensing and support costs.

- **Option 3: Terminate VPN on a virtual private gateway (VGW)** — This AWS Site-to-Site VPN service option enables a one-to-one connectivity design where you create one VPN connection (consisting of a pair of redundant VPN tunnels) per VPC. This is great way to get started with VPN connectivity into AWS, but as you scale the number of VPCs, managing a growing number of VPN connections can become challenging. Therefore, edge consolidation design leveraging Transit Gateway will eventually be a better option. VPN throughput to a VGW is limited to 1.25 Gbps per tunnel and ECMP load balancing is not supported. From a pricing perspective, you only pay for AWS VPN pricing, there is no charge for running a VGW. For more information, refer to AWS VPN Pricing and AWS VPN on virtual private gateway.

- **Option 4: Terminate VPN connection on client VPN endpoint** — AWS Client VPN is a managed client-based VPN service that enables you to securely access your AWS resources and resources in your on-premises network. With Client VPN, you can access your resources from any location using an OpenVPN or AWS provided VPN client. By setting up a Client VPN endpoint, clients and users can connect to establish a Transport Layer Security (TLS) VPN connection. For more information, refer to the AWS Client VPN documentation.

- **Option 5: Consolidate VPN connection on AWS Cloud WAN** — This option is similar to the first option in this list, but it uses the CloudWAN fabric to programmatically configure VPN connections through the network policy document.

AWS Direct Connect

While VPN over internet is a great option to get started, internet connectivity may not be reliable for production traffic. Because of this unreliability, many customers choose AWS Direct Connect. AWS Direct Connect is a networking service that provides an alternative to using the internet to connect to AWS. Using AWS Direct Connect, data that would have previously been transported over the internet is delivered through a private network connection between your facilities and AWS. In many circumstances, private network connections can reduce costs, increase bandwidth, and provide a more consistent network experience than internet-based connections. There are several ways to use AWS Direct Connect to connect to VPCs:

Ways to connect your on-premises data centers using AWS Direct Connect

- **Option 1: Create a private virtual interface (VIF) to a VGW attached to a VPC** — You can create 50 VIFs per Direct Connect connection, allowing you to connect to a maximum of 50 VPCs (one VIF provides connectivity to one VPC). There is one BGP peering per VPC. Connectivity in this

setup is restricted to the AWS Region that the Direct Connect location is homed to. The one-to-one mapping of VIF to VPC (and lack of global access) makes this the least preferred way to access VPCs in the Landing Zone.

- **Option 2: Create a private VIF to a Direct Connect gateway associated with multiple VGWs (each VGW is attached to a VPC)** — A Direct Connect gateway is a globally available resource. You can create the Direct Connect gateway in any Region and access it from all other Regions, including GovCloud (excluding China). A Direct Connect Gateway can connect to up to 20 VPCs (via VGWs) globally in any AWS account over a single private VIF. This is a great option if a Landing Zone consists of a small number of VPCs (ten or fewer VPCs) and/or you need global access. There is one BGP peering session per Direct Connect Gateway per Direct Connect connection. Direct Connect gateway is only for north/south traffic flow and does not permit VPC-to-VPC connectivity. Refer to Virtual private gateway associations in the AWS Direct Connect documentation for more details. With this option, the connectivity is not restricted to the AWS Region where the Direct Connect location is homed to. AWS Direct Connect gateway is only for north/south traffic flow and does not permit VPC-to-VPC connectivity. An exception to this rule is when a supernet is advertised across two or more VPCs that have their attached VGWs associated with the same AWS Direct Connect gateway and on the same virtual interface. In this case, VPCs can communicate with each other through the AWS Direct Connect endpoint. Refer to AWS Direct Connect gateways documentation for more details.

- **Option 3: Create a transit VIF to a Direct Connect gateway associated with Transit Gateway** — You can associate a Transit Gateway instance to a Direct Connect gateway by using a Transit VIF. AWS Direct Connect now supports connections to Transit Gateway for all port speeds, providing a more cost-effective choice for Transit Gateway users when high-speed connections (greater than 1Gbps) are not required. This enables you to use Direct Connect at speeds of 50, 100, 200, 300, 400, and 500 Mbps connecting to Transit Gateway. Transit VIF allows you to connect your on-premises data center to up to six Transit Gateway instances per AWS Direct Connect gateway (which can connect to thousands of VPCs) across different AWS Regions and AWS accounts over single transit VIF and BGP peering. This is the simplest setup among the options for connecting multiple VPCs at scale, but you should be mindful of the Transit Gateway quotas. One key limit to note is that you can advertise only 200 prefixes from a Transit Gateway to on-premises router over the transit VIF. With the previous options, you pay for Direct Connect pricing. For this option, you also pay for the Transit Gateway attachment and data processing charges. For more information, refer to the Transit Gateway Associations on Direct Connect documentation.

- **Option 4: Create a VPN connection to Transit Gateway over Direct Connect public VIF** — A public VIF allows you to access all AWS public services and endpoints using the public IP addresses. When you create a VPN attachment on a Transit Gateway, you get two public IP

addresses for VPN endpoints at the AWS side. These public IPs are reachable over the public VIF. You can create as many VPN connections to as many Transit Gateway instances as you want over Public VIF. When you create a BGP peering over the public VIF, AWS advertises the entire AWS public IP range to your router. To ensure that you only permit certain traffic (for example, allowing traffic only to the VPN termination endpoints) you are advised to use a firewall on-premises facilities. This option can be used to encrypt your Direct Connect at the network layer.

- **Option 5: Create a VPN connection to Transit Gateway over AWS Direct Connect using Private IP VPN** — Private IP VPN is a feature that provides customers the ability to deploy AWS Site-to-Site VPN connections over Direct Connect using private IP addresses. With this feature, you can encrypt traffic between your on-premises networks and AWS through Direct Connect connections without the need for public IP addresses, thus enhancing security and network privacy at the same time. Private IP VPN is deployed on top of Transit VIFs, so it allows you to use Transit Gateway for centralized management of customers' VPCs and connections to the on-premises networks in a more secured, private, and scalable manner.

- **Option 6: Create GRE tunnels to Transit Gateway over a transit VIF** – The Transit Gateway Connect attachment type supports GRE. With Transit Gateway Connect, SD-WAN infrastructure can be natively connected to AWS without having to set up IPsec VPNs between SD-WAN network virtual appliances and Transit Gateway. The GRE tunnels can be established over a transit VIF, having Transit Gateway Connect as the attachment type, providing higher bandwidth performance compared to a VPN connection. For more information, refer to the Simplify SD-WAN connectivity with AWS Transit Gateway Connect blog post.

The "transit VIF to Direct Connect gateway" option might seem to be the best option because it lets you consolidate all your on-premises connectivity for a given AWS Region at a single point (Transit Gateway) using a single BGP session per Direct Connect connection; however, some of the limits and considerations around this option might lead you to use both private and transit VIFs in conjuction for your Landing Zone connectivity requirements.

The following figure illustrates a sample setup where Transit VIF is used as a default method for connecting to VPCs and a private VIF is used for an edge use case where exceptionally large amounts of data must be transferred from an on-premises Data Center to the media VPC. Private VIF is used to avoid Transit Gateway data processing charges. As a best practice, you should have at least two connections at two different Direct Connect locations for maximum redundancy—a total of four connections. You create one VIF per connection for a total of four private VIFs and four transit VIFs. You can also create a VPN as backup connectivity to AWS Direct Connect connections.

With the "Create GRE tunnels to Transit Gateway over a transit VIF" option, you get the capability to natively connect your SD-WAN infrastructure with AWS. It eliminates the need to setup IPsec VPNs between SD-WAN network virtual appliances and Transit Gateway.

Sample reference architecture for hybrid connectivity

Use the Network Services account for creating Direct Connect resources enabling demarcation of network administrative boundaries. The Direct Connect connections, Direct Connect gateways, and Transit Gateways can all reside in a Network Services account. To share the AWS Direct Connect connectivity with your Landing Zone, simply share the Transit Gateway through AWS RAM with other accounts.

MACsec security on Direct Connect connections

Customers can use MAC Security Standard (MACsec) encryption (IEEE 802.1AE) with their Direct Connect connections for 10 Gbps and 100 Gbps dedicated connections at select locations. With this capability, customers can secure their data on the layer 2 level, and Direct Connect delivers point-to-point encryption. To enable the Direct Connect MACsec feature, ensure that the MACsec pre-requisites are met. Because MACsec protects links on a hop-by-hop basis, your device must have a direct layer 2 adjacency with our Direct Connect device. Your last-mile provider can help you verify that your connection will work with MACsec. For more information refer to Adding MACsec security to AWS Direct Connect connections.

AWS Direct Connect resiliency recommendations

With AWS Direct Connect, customers can achieve highly resilient connectivity into their Amazon VPCs and AWS resources from their on-premises networks. It is best practice that customers connect from multiple data centers to eliminate any single point physical location failures. It is also recommended that, depending on the type of workloads, customers utilize more than one Direct Connect connection for redundancy.

AWS also offers the AWS Direct Connect Resiliency Toolkit, which provides customers with a connection wizard with multiple redundancy models; to help them determine which model works best for their service level agreement (SLA) requirements and design their hybrid connectivity using Direct Connect connections accordingly. For more information, refer to AWS Direct Connect Resiliency Recommendations.

AWS Direct Connect SiteLink

Previously, configuring site-to-site links for your on-premises networks was only possible by using direct circuit buildout through dark fiber or other technologies, IPSEC VPNs, or by using third-party circuit providers with technologies such as MPLS, MetroEthernet, or legacy T1 circuits. With the advent of SiteLink, customers can now enable direct site-to-site connectivity for their on-premises location that terminate at an AWS Direct Connect location. Use your Direct Connect circuit to provide site-to-site connectivity without having to route traffic through your VPCs, bypassing the AWS region completely.

Now, you can create global, reliable, and pay-as-you-go connections between the offices and data centers in your global network by sending data over the fastest path between AWS Direct Connect locations.
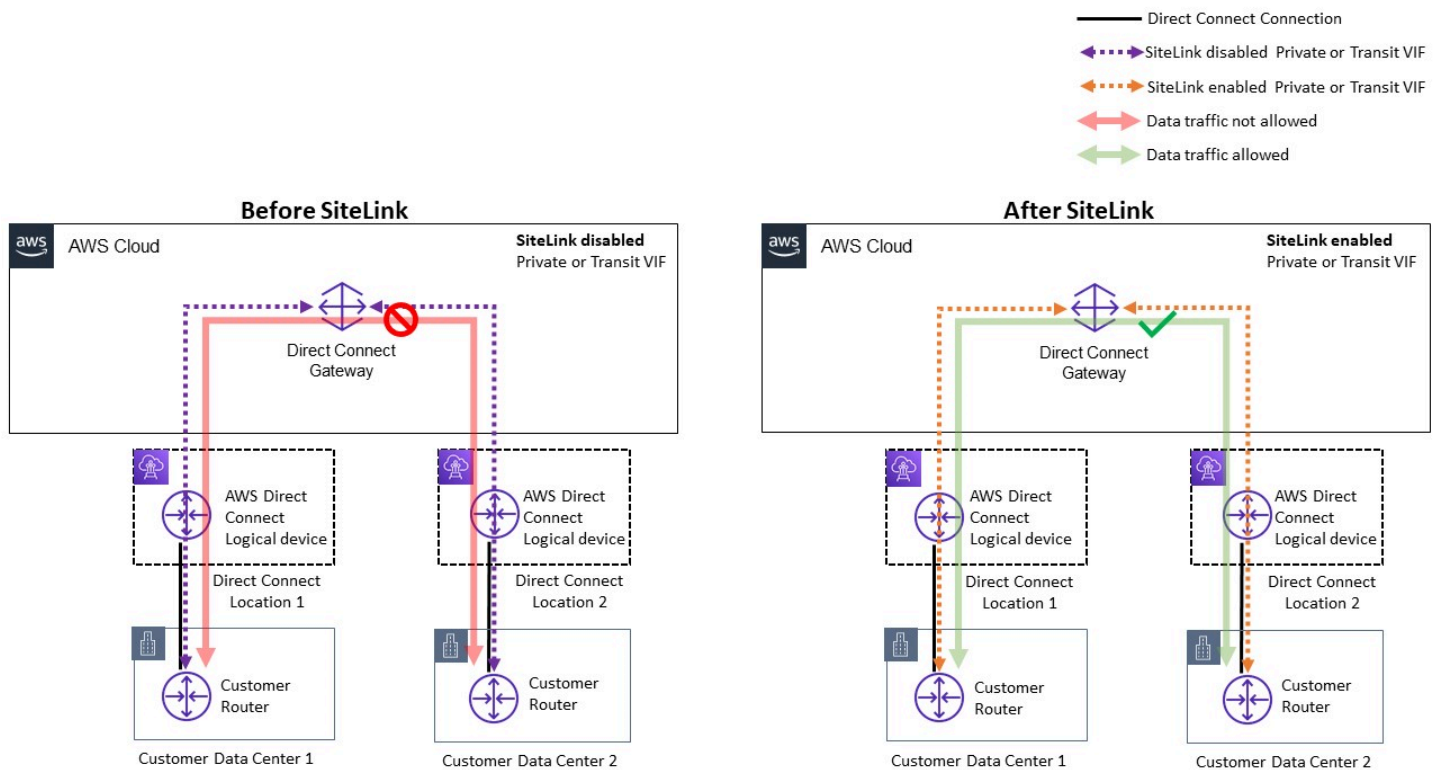
Sample reference architecture for AWS Direct Connect SiteLink

When using SiteLink, you first connect your on-premises networks to AWS at any of over 100 AWS Direct Connect locations worldwide. Then, you create virtual interfaces (VIFs) on those connections and enable SiteLink. Once all VIFs are attached to the same AWS Direct Connect gateway (DXGW), you can start sending data between them. Your data follows the shortest path between AWS Direct Connect locations to its destination, using the fast, secure, and reliable AWS global network. You don't need to have any resources in any AWS Region to use SiteLink.

With SiteLink, the DXGW learns IPv4/IPv6 prefixes from your routers over SiteLink enabled VIFs, runs BGP best path algorithm, updates attributes such as NextHop and AS_Path, and re-advertises these BGP prefixes to the rest of your SiteLink-enabled VIFs associated with that DXGW. If you disable SiteLink on a VIF, the DXGW will not advertise the learned on-premises prefixes over this VIF to the other SiteLink-enabled VIFs. The on-premises prefixes from a SiteLink disabled VIF is only advertised to the DXGW Gateway associations, such as AWS Virtual Private Gateways (VGWs) or Transit Gateway (TGW) instances associated with the DXGW.
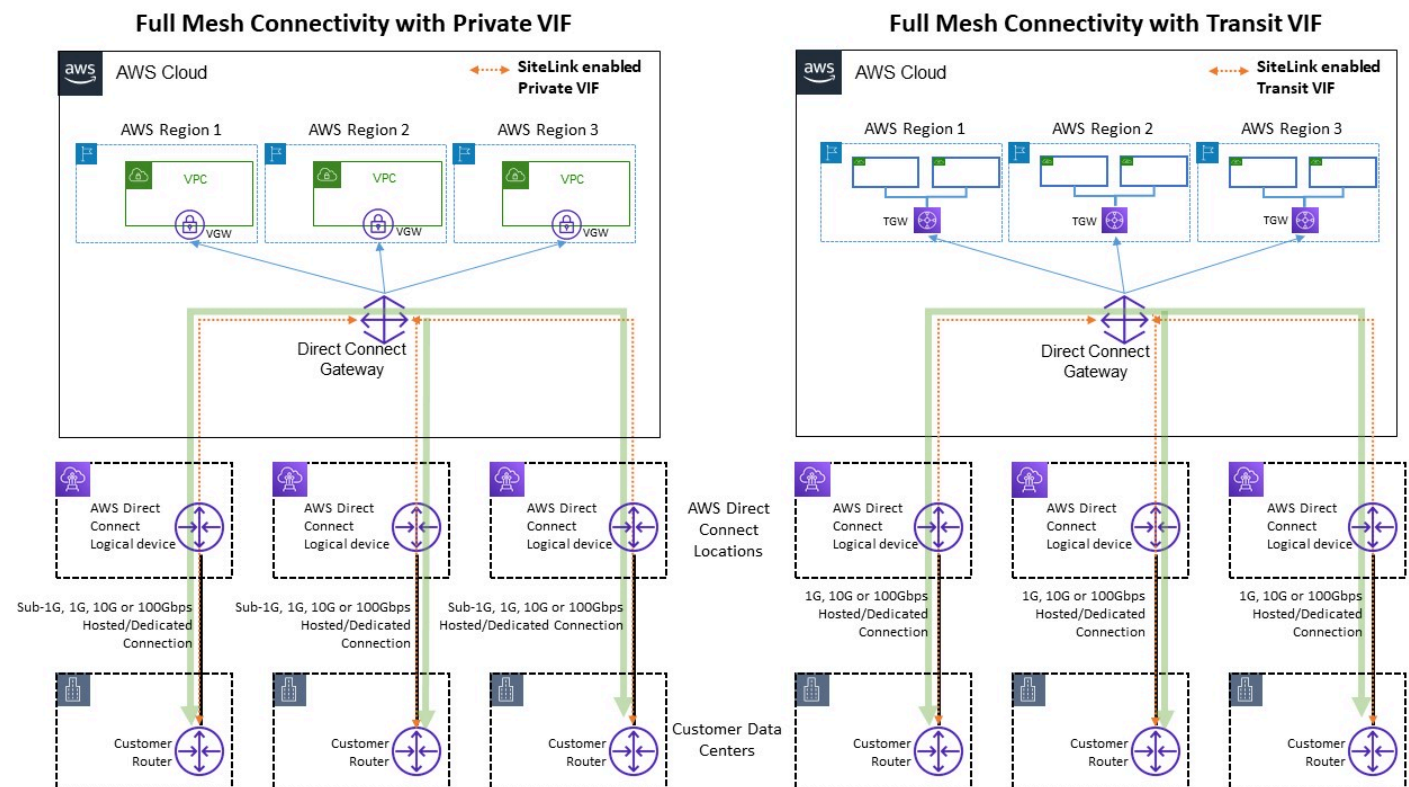
Sitelink allows traffic flows example

SiteLink allows customers to use the AWS global network to function as a primary or secondary/ backup connection between their remote locations, with high bandwidth and low latency, with dynamic routing to control which locations can communicate with each other and with your AWS regional resources.

For more information, refer to [Introducing AWS Direct Connect SiteLink](#).

Centralized egress to internet

As you deploy applications in your multi-account environment, many apps will require outbound-only internet access (for example, downloading libraries, patches, or OS updates). This can be achieved for both IPv4 and IPv6 traffic. For IPv4, this can achieved through network address translation (NAT) in the form of a NAT gateway (recommended), or alternatively, a self-managed NAT instance running on an Amazon EC2 instance, as a means for all egress internet access. Internal applications reside in private subnets, while NAT Gateways and Amazon EC2 NAT instances reside in a public subnet.

AWS recommends that you use NAT gateways because they provide better availability and bandwidth and require less effort on your part to administer. For more information, refer to Compare NAT gateways and NAT instances.

For IPv6 traffic, egress traffic can be configured to leave each VPC through an egress only internet gateway in a decentralized manner or it can be configured to be sent to a centralized VPC using NAT instances or proxy instances. The IPv6 patterns are discussed in Centralized egress for IPv6.

Topics

- Using the NAT gateway for centralized IPv4 egress
- Using the NAT gateway with AWS Network Firewall for centralized IPv4 egress
- Using the NAT gateway and Gateway Load Balancer with Amazon EC2 instances for centralized IPv4 egress
- Centralized egress for IPv6

Using the NAT gateway for centralized IPv4 egress

NAT gateway is a managed network address translation service. Deploying a NAT gateway in every spoke VPC can become cost prohibitive because you pay an hourly charge for every NAT gateway you deploy (refer to Amazon VPC pricing). Centralizing NAT gateways can be a viable option to reduce costs. To centralize, you create a separate egress VPC in the network services account, deploy NAT gateways in the egress VPC, and route all egress traffic from the spoke VPCs to the NAT gateways residing in the egress VPC using Transit Gateway or CloudWAN, as shown in the following figure.

> ⓘ **Note**
>
> When you centralize NAT gateway using Transit Gateway, you pay an extra Transit Gateway data processing charge — compared to the decentralized approach of running a NAT gateway in every VPC. In some edge cases when you send huge amounts of data through NAT gateway from a VPC, keeping the NAT local in the VPC to avoid the Transit Gateway data processing charge might be a more cost-effective option.
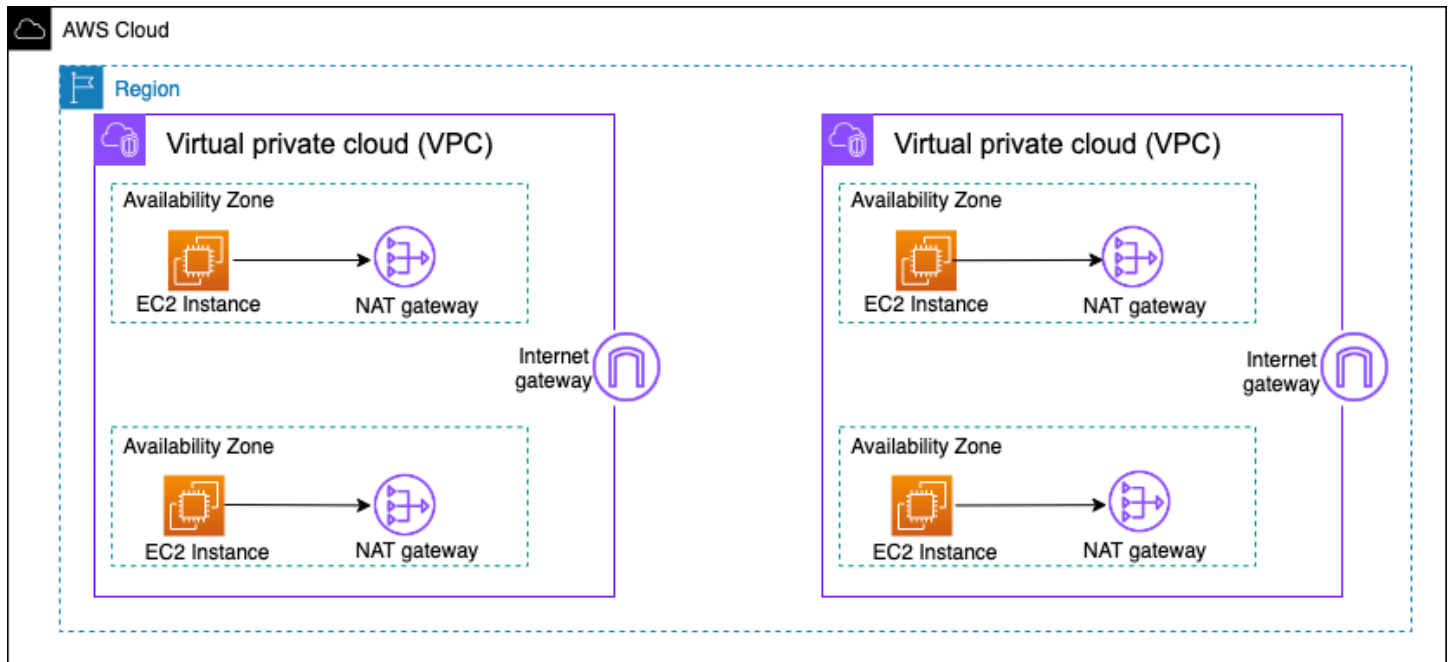
Decentralized high availability NAT gateway architecture

Centralized NAT gateway using Transit Gateway (overview)

Centralized NAT gateway using Transit Gateway (route table design)

In this setup, spoke VPC attachments are associated with Route Table 1 (RT1) and are propagated to Route Table 2 (RT2). There is a [Blackhole](#) route to disallow the two VPCs from communicating with each other. If you want to allow inter-VPC communication, you can remove the `10.0.0.0/8 -> Blackhole` route entry from RT1. This allows them to communicate via the transit gateway. You can also propagate the spoke VPC attachments to RT1 (or alternatively, you can use one route table and associate/propagate everything to that), enabling direct traffic flow between the VPCs using Transit Gateway.

You add a static route in RT1 pointing all traffic to egress VPC. Because of this static route, Transit Gateway sends all internet traffic through its ENIs in the egress VPC. Once in the egress VPC, traffic follows the routes defined in the subnet route table where these Transit Gateway ENIs are present. You add a route in subnet route tables pointing all traffic towards the respective NAT gateway in the same Availability Zone to minimize cross-Availability Zone (AZ) traffic. The NAT gateway subnet route table has internet gateway (IGW) as the next hop. For return traffic to flow back, you must add a static route table entry in the NAT gateway subnet route table pointing all spoke VPC bound traffic to Transit Gateway as the next hop.

High availability

For high availability, you should use more than one NAT gateway (one in each Availability Zone). If a NAT gateway is unavailable, traffic might be dropped in that Availability Zone that is traversing the affected NAT gateway. If one Availability Zone is unavailable, the Transit Gateway endpoint along with NAT gateway in that Availability Zone will fail, and all traffic will flow via the Transit Gateway and NAT gateway endpoints in the other Availability Zone.

Security

You can rely on security groups on the source instances, blackhole routes in the Transit Gateway route tables, and the network ACL of the subnet in which the NAT gateway is located. For example, customers can use ACLs on the NAT Gateway public subnet(s) to allow or block source or destination IP addresses. Alternatively, you can use NAT Gateway with AWS Network Firewall for centralized egress described in the next section to meet this requirement.

Scalability

A single NAT gateway can support up to 55,000 simultaneous connections per assigned IP address to each unique destination. You can request a quota adjustment to allow up to eight assigned IP addresses, allowing for 440,000 simultaneous connections to a single destination IP and port. NAT gateway provides 5 Gbps of bandwidth and automatically scales to 100 Gbps. Transit Gateway generally does not act as a load balancer and would not distribute your traffic evenly across NAT gateways in the multiple Availability Zones. The traffic across the Transit Gateway will stay within an Availability Zone, if possible. If the Amazon EC2 instance initiating traffic is in Availability Zone 1, traffic will flow out of the Transit Gateway elastic network interface in the same Availability Zone 1 in the egress VPC and will flow to the next hop based on that subnet route table that elastic network interface resides in. For a complete list of rules, refer to NAT gateways in the Amazon Virtual Private Cloud documentation.

For more information, refer to the Creating a single internet exit point from multiple VPCs Using AWS Transit Gateway blog post.

Using the NAT gateway with AWS Network Firewall for centralized IPv4 egress

If you want to inspect and filter your outbound traffic, you can incorporate AWS Network Firewall with NAT gateway in your centralized egress architecture. AWS Network Firewall is a managed

service that makes it easy to deploy essential network protections for all of your VPCs. It provides control and visibility to Layer 3-7 network traffic for your entire VPC. You can do URL/domain name, IP address, and content-based outbound traffic filtering to stop possible data loss, help meet compliance requirements, and block known malware communications. AWS Network Firewall supports thousands of rules that can filter out network traffic destined for known bad IP addresses or bad domain names. You can also use Suricata IPS rules as part of the AWS Network Firewall service by importing open-source rulesets or authoring your own Intrusion Prevention System (IPS) rules using Suricata rule syntax. AWS Network Firewall also allows you to import compatible rules sourced from AWS partners.

In the centralized egress architecture with inspection, the AWS Network Firewall endpoint is a default route table target in the transit gateway attachments subnet route table for the egress VPC. Traffic between spoke VPCs and the internet is inspected using AWS Network Firewall as shown in the following diagram.
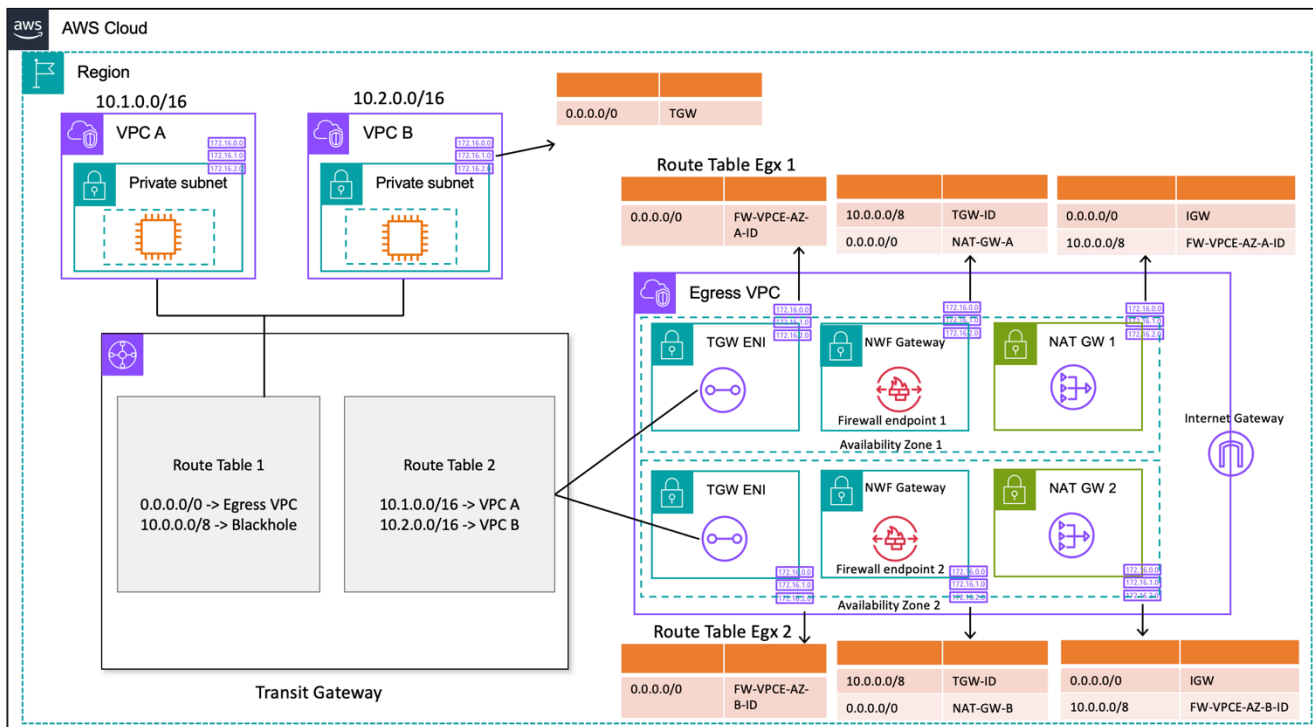
Centralized egress with AWS Network Firewall and NAT gateway (route table design)

For centralized deployment model with Transit Gateway, AWS recommends deploying AWS Network Firewall endpoints in multiple Availability Zones. There should be one firewall endpoint in each Availability Zone the customer is running workloads in, as shown in the preceding diagram. As a best practice, the firewall subnet should not contain any other traffic because AWS Network Firewall is not able to inspect traffic from sources or destinations within a firewall subnet.

Similar to the previous setup, spoke VPC attachments are associated with Route Table 1 (RT1) and are propagated to Route Table 2 (RT2). A Blackhole route is explicitly added to disallow the two VPCs from communicating with each other.

Continue to use a default route in RT1 pointing all traffic to egress VPC. Transit Gateway will forward all traffic flows to one of the two availability zones in the egress VPC. Once traffic reaches one of Transit Gateway ENIs in the egress VPC, you hit a default route which will forward traffic to one of the AWS Network Firewall endpoints in their respective availability zone. AWS Network Firewall will then inspect traffic based on the rules you set before forwarding traffic to the NAT gateway using a default route.

This case doesn't require Transit Gateway appliance mode, because you aren't sending traffic between attachments.

> ### ⓘ Note
>
> AWS Network Firewall does not perform network address translation for you, this function would be handled by the NAT gateway after traffic inspection through the AWS Network Firewall. Ingress routing is not required in this case as return traffic will be forwarded to the NATGW IPs by default.

Because you are using a Transit Gateway, here we can place the firewall prior to NAT gateway. In this model, the firewall can see the source IP behind the Transit Gateway.

If you were doing this in a single VPC, we can use the VPC routing enhancements that allow you to inspect traffic between subnets in the same VPC. For details, refer to the [Deployment models for AWS Network Firewall with VPC routing enhancements](#) blog post.

Scalability

AWS Network Firewall can automatically scale firewall capacity up or down based on traffic load to maintain steady, predictable performance to minimize costs. AWS Network Firewall is designed to support tens of thousands of firewall rules and can scale up to 100 Gbps throughput per Availability Zone.

Key considerations

- Each firewall endpoint can handle about 100 Gbps of traffic, if you require higher burst or sustained throughput, contact [AWS support](#).

- If you choose to create a NAT gateway in your AWS account along with Network Firewall, standard NAT gateway processing and per-hour usage charges are waived on a one-to-one basis with the processing per GB and usage hours charged for your firewall.

- You can also consider distributed firewall endpoints through AWS Firewall Manager without a Transit Gateway.

- Test firewall rules before moving them to production, similar to a network access control list, as order matters.

- Advanced Suricata rules are required for deeper inspection. Network firewall supports encrypted traffic inspection for ingress as well as egress traffic.

- The HOME_NET rule group variable defined the source IP range eligible for processing in the Stateful engine. Using a centralized approached, you must add all additional VPC CIDRs attached to the Transit Gateway to make them eligible for processing. Refer to the Network Firewall documentation for more details on the HOME_NET rule group variable.

- Consider deploying Transit Gateway and egress VPC in a separate Network Services account to segregate access based on delegation of duties; for example, only network administrators can access the Network Services account.

- To simplify deployment and management of AWS Network Firewall in this model, AWS Firewall Manager can be used. Firewall Manager allows you to centrally administer your different firewalls by automatically applying protection you create in the centralized location to multiple accounts. Firewall Manager supports both distributed and centralized deployment models for Network Firewall. To learn more, see the blog post How to deploy AWS Network Firewall by using AWS Firewall Manager.

Using the NAT gateway and Gateway Load Balancer with Amazon EC2 instances for centralized IPv4 egress

Using a software-based virtual appliance (on Amazon EC2) from AWS Marketplace and AWS Partner Network as an exit point is similar to the NAT gateway setup. This option can be used if you want to use the advanced layer 7 firewall/Intrusion Prevention/Detection System (IPS/IDS) and deep packet inspection capabilities of the various vendor offerings.

In the following figure, in addition to NAT gateway, you deploy virtual appliances using EC2 instances behind a Gateway Load Balancer (GWLB). In this setup, the GWLB, Gateway Load Balancer Endpoint (GWLBE), virtual appliances and NAT gateways are deployed in a centralized VPC which is connected to Transit Gateway using VPC attachment. The spoke VPCs are also connected to the

Transit Gateway using a VPC Attachment. Because GWLBEs are a routable target, you can route traffic moving to and from Transit Gateway to the fleet of virtual appliances that are configured as targets behind a GWLB. GWLB acts as a bump-in-the-wire and transparently passes all Layer 3 traffic through third-party virtual appliances, and thus invisible to the source and destination of the traffic. Therefore, this architecture allows you to centrally inspect all of your egress traffic traversing through Transit Gateway.

For more information on how the traffic flows from the applications in the VPCs to the internet and back through this setup, refer to Centralized inspection architecture with AWS Gateway Load Balancer and AWS Transit Gateway.

You can enable appliance mode on Transit Gateway to maintain flow symmetry through virtual appliances. This means the bidirectional traffic is routed through the same appliance and the Availability Zone for the life of the flow. This setting is particularly important for stateful firewalls performing deep packet inspection. Enabling appliance mode removes the need for complex workarounds, such as source network address translation (SNAT), to force traffic to return to the correct appliance to maintain symmetry. Refer to Best practices for deploying Gateway Load Balancer for details.

It is also possible to deploy GWLB endpoints in a distributed manner without Transit Gateway to enable egress inspection. Learn more about this architectural pattern in the Introducing AWS Gateway Load Balancer: Supported architecture patterns blog post.
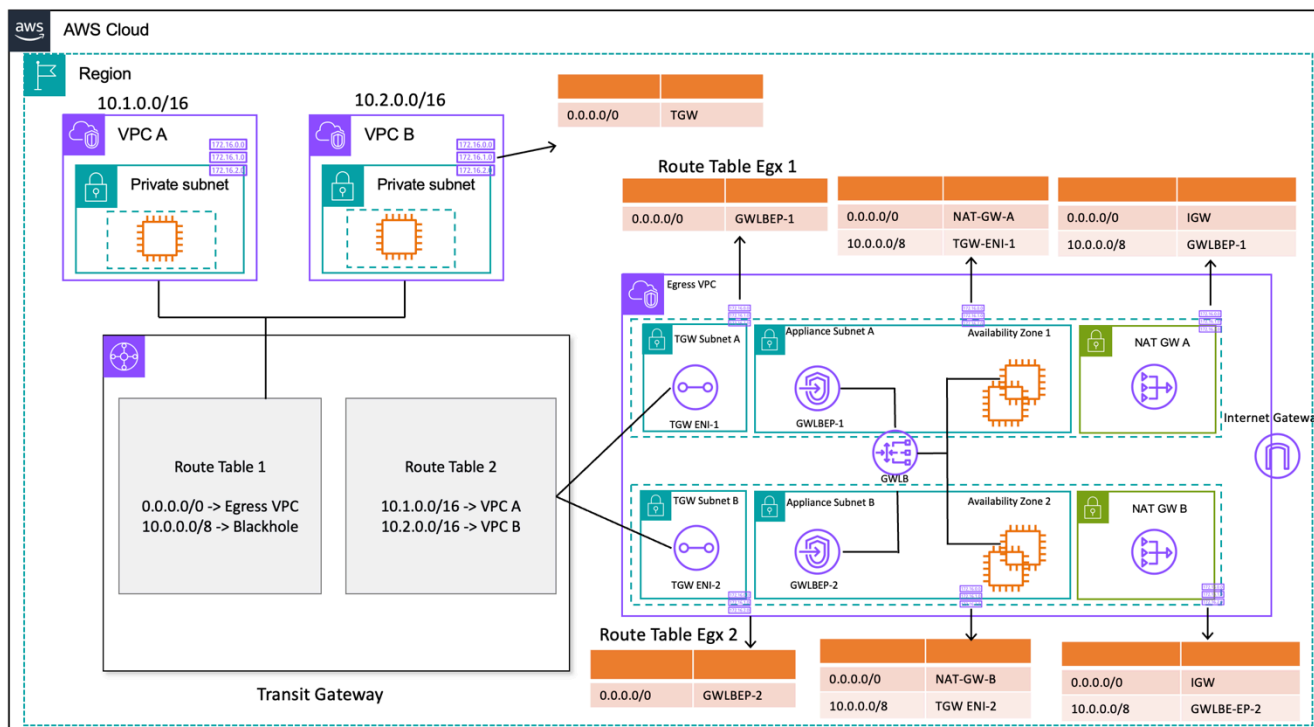
Centralized egress with Gateway Load Balancer and EC2 instance (route table design)

High availability

AWS recommends deploying Gateway Load Balancers and virtual appliances in multiple Availability Zones for higher availability.

Gateway Load Balancer can perform health checks to detect virtual appliance failures. In case of an unhealthy appliance, GWLB reroutes the new flows to healthy appliances. Existing flows always go to the same target regardless of health status of target. This allows for connection draining and accommodate health check failures due to CPU spikes on appliances. For more details, refer to section 4: Understand appliance and Availability Zone failure scenarios in the blog post Best practices for deploying Gateway Load Balancer. Gateway Load Balancer can use auto scaling groups as targets. This benefit takes out the heavy lifting of managing availability and scalability of the appliance fleets.

Advantages

Gateway Load Balancer and Gateway Load Balancer endpoints are powered by AWS PrivateLink, which allows for the exchange of traffic across VPC boundaries securely without the need to traverse the public internet.

Gateway Load Balancer is a managed service that removes undifferentiated heavy lifting of managing, deploying, scaling virtual security appliances so that you can focus on the things that matter. Gateway Load Balancer can expose the stack of firewalls as an endpoint service for customers to subscribe to using the AWS Marketplace. This is called Firewall as a Service (FWaaS); it introduces a simplified deployment and removes the need to rely on BGP and ECMP to distribute traffic across multiple Amazon EC2 instances.

Key considerations

- The appliances need to support Geneve encapsulation protocol to integrate with GWLB.
- Some third-party appliances can support SNAT and overlay routing (two-arm mode) therefore eliminating the need to create NAT gateways for saving costs. However, consult with an AWS partner of your choice before using this mode as this is dependent on vendor support and implementation.
- Take a note of GWLB idle timeout. This can result in connection timeouts on clients. You can tune your timeouts on client, server, firewall, and OS level to avoid this. Refer to *Section 1: Tune TCP*

keep-alive or timeout values to support long-lived TCP flows in the [Best practices for deploying Gateway Load Balancer](#) blog post for more information.

- GWLBE are powered by AWS PrivateLink, so AWS PrivateLink charges will be applicable. You can learn more in the [AWS PrivateLink pricing page](#). If you are using the centralized model with Transit Gateway, the TGW data processing charges will be applicable.

- Consider deploying Transit Gateway and egress VPC in a separate Network Services account to segregate access based on delegation of duties, such as only network administrators can access Network Services Account.

Centralized egress for IPv6

To support IPv6 egress in dual stack deployments that have centralized IPv4 egress, one of two patterns must be chosen:

- Centralized IPv4 egress with decentralized IPv6 egress

- Centralized IPv4 egress and centralized IPv6 egress

In the first pattern, shown in the following diagram, egress-only internet gateways are deployed in each spoke VPC. Egress-only internet gateways are horizontally scaled, redundantly, and highly available gateways that allow outbound communication over IPv6 from instances inside your VPC. They prevent the internet from initiating IPv6 connections with your instances. Egress-only internet gateways have no charge. In this deployment model, IPv6 traffic flows out of the egress-only internet gateways in each VPC and IPv4 traffic flows over the centralized NAT Gateways deployed.

Centralized IPV4 egress and decentralized outbound only IPv6 egress

In the second pattern, shown in the following diagrams, egress IPv6 traffic from your instances is sent to a centralized VPC. This can be accomplished by using IPv6-to-IPv6 Network Prefix Translation (NPTv6) with NAT66 instances and NAT Gateways or by using Proxy Instances and

Network Load Balancer. This pattern is applicable if centralized traffic inspection for outbound traffic is required and it cannot be performed in each spoke VPC.

Centralized IPv6 egress using NAT gateways and NAT66 instances

Centralized IPv4 and IPv6 egress using proxy instances and Network Load Balancer

The IPv6 on AWS whitepaper describes the centralized IPv6 egress patterns. The IPv6 egress patterns are discussed in more detail in the blog Centralized outbound internet traffic for dual stack IPv4 and IPv6 VPCs, along with special considerations, sample solutions, and diagrams.

Centralized network security for VPC-to-VPC and on-premises to VPC traffic

There might be scenarios where a customer wants to implement a layer 3-7 firewall/IPS/IDS within their multi-account environment to inspect traffic flows between VPCs (east-west traffic) or between an on-premises data center and a VPC (north-south traffic). This can be achieved different ways, depending on use case and requirements. For example, you could incorporate the Gateway Load Balancer, Network Firewall, Transit VPC, or use centralized architectures with Transit Gateways. These scenarios are discussed in the following section.

Considerations using a centralized network security inspection model

To reduce costs, you should be selective in what traffic passes via your AWS Network Firewall or Gateway Load Balancer. One way to proceed is to define security zones and inspect traffic between untrusted zones. An untrusted zone can be a remote site managed by a third party, a vendor VPC you don't control/trust, or a sandbox/dev VPC, which has more relaxed security rules compared to rest of your environment. There are four zones in this example:

- **Untrusted Zone** — This is for any traffic coming from the 'VPN to remote untrusted site' or the third-party vendor VPC.

- **Production (Prod) Zone** — This contains traffic from the production VPC and on-premises customer DC.

- **Development (Dev) Zone** —This contains traffic from the two development VPCs.

- **Security (Sec) Zone** — Contains our firewall components Network Firewall or Gateway Load Balancer.

This setup has four security zones, but you might have more. You can use multiple route tables and blackhole routes to achieve security isolation and optimal traffic flow. Choosing the right set of zones is dependent on your overall Landing Zone design strategy (account structure, VPC design). You can have zones to enable isolation between Business Units (BUs), applications, environments, and so on.

If you want to inspect and filter your VPC-to-VPC, inter-zone traffic, and VPC-on-premises traffic, you can incorporate AWS Network Firewall with Transit Gateway in your centralized architecture. By having the hub-and-spoke model of the AWS Transit Gateway, a centralized deployment model can be achieved. The AWS Network Firewall is deployed in a separate security VPC. A separate security VPC provides a simplified and central approach to manage inspection. Such a VPC architecture gives AWS Network Firewall source and destination IP visibility. Both source and destination IPs are preserved. This security VPC consists of two subnets in each Availability Zone; where one subnet is a dedicated to AWS Transit Gateway attachment and the other subnet is dedicated to the firewall endpoint. The subnets in this VPC should only contain AWS Network Firewall endpoints because Network Firewall can't inspect traffic in the same subnets as the endpoints. When you use Network Firewall to centrally inspect traffic, it can perform deep packet inspection (DPI) on ingress traffic. The DPI pattern is expanded upon in the *Centralized Inbound Inspection* section of this paper.

VPC-to-VPC and on-premises to VPC traffic inspection using Transit Gateway and AWS Network Firewall (route table design)

In the centralized architecture with inspection, the Transit Gateway subnets require a separate VPC route table to ensure the traffic is forwarded to firewall endpoint within the same Availability Zone. For the return traffic, a single VPC route table containing a default route towards the Transit Gateway is configured. Traffic is returned to AWS Transit Gateway in the same Availability Zone after it has been inspected by AWS Network Firewall. This is possible due to the appliance mode

feature of the Transit Gateway. The appliance mode feature of the Transit Gateway also helps the AWS Network Firewall to have stateful traffic inspection capability inside the security VPC.

With the appliance mode enabled on a transit gateway, it selects a single network interface using flow hash algorithm for the entire life of the connection. The transit gateway uses the same network interface for the return traffic. This ensures that bidirectional traffic is routed symmetrically—it's routed through the same Availability Zone in the VPC attachment for the life of the flow. For more information on appliance mode, refer to Stateful appliances and appliance mode in the Amazon VPC documentation.

For different deployment options of security VPC with AWS Network Firewall and Transit Gateway, refer to the Deployment models for AWS Network Firewall blog post.

Using Gateway Load Balancer with Transit Gateway for centralized network security

Often times, customers want to incorporate virtual appliances to handle the traffic filtering and to provide security inspection capabilities. In such use cases, they can integrate Gateway Load Balancer, virtual appliances, and Transit Gateway to deploy a centralized architecture for inspecting VPC-to-VPC and VPC-to-on-premises traffic.

Gateway Load Balancer is deployed in a separate security VPC along with the virtual appliances. The virtual appliances that will inspect the traffic are configured as targets behind the Gateway Load Balancer. Because Gateway Load Balancer endpoints are a routable target, customers can route traffic moving to and from Transit Gateway to the fleet of virtual appliances. To ensure flow symmetry, appliance mode is enabled on the Transit Gateway.

Each spoke VPC has a route table that is associated with the Transit Gateway, which has the default route to the Security VPC attachment as the next-hop.

The centralized Security VPC consists of appliance subnets in each Availability Zone; which have the Gateway Load Balancer endpoints and the virtual appliances. It also has subnets for Transit Gateway attachments in each Availability Zone, as shown in the following figure.

For more information on centralized security inspection with Gateway Load Balancer and Transit Gateway, refer to the Centralized inspection architecture with AWS Gateway Load Balancer and AWS Transit Gateway blog post.

VPC-to-VPC and on-premises-to-VPC traffic inspection using Transit Gateway and AWS Gateway Load Balancer (route table design)

Key considerations for AWS Network Firewall and AWS Gateway Load Balancer

- Appliance mode should be enabled on the Transit Gateway when doing east-west inspection.

- You can deploy the same model for inspection of traffic to other AWS Regions using AWS Transit Gateway Inter-Region peering.

- By default, each Gateway Load Balancer deployed in an Availability Zone distributes traffic across the registered targets within the same Availability Zone only. This is called Availability Zone affinity. If you enable cross-zone load balancing, Gateway Load Balancer distributes traffic across all registered and healthy targets in all enabled Availability Zones. If all targets across all Availability Zones are unhealthy, Gateway Load Balancer fails open. Refer to section 4: Understand appliance and Availability Zone failure scenarios in the Best practices for deploying Gateway Load Balancer blog post for more details.

- For multi-Region deployment, AWS recommends that you set up separate inspection VPCs in the respective local Regions to avoid inter-Region dependencies and reduce associated data transfer costs. You should inspect traffic in the local Region instead of centralizing inspection to another Region.

- Cost of running an additional EC2-based high availability (HA) pair in multi-Region deployments can add up. For more information, refer to the [Best practices for deploying Gateway Load Balancer](#) blog post.

AWS Network Firewall vs. Gateway Load Balancer

Table 2 — AWS Network Firewall vs Gateway Load Balancer

Criteria	AWS Network Firewall	Gateway Load Balancer
Use case	Stateful, managed, network firewall with intrusion detection and prevention service capability compatible with Suricata.	Managed service which makes it easy to deploy, scale and manage third-party virtual appliances
Complexity	AWS managed service. AWS handles the scalability and availability of the service.	AWS managed service. AWS will handle the scalability and availability of the the Gateway Load Balancer service. The customer is responsible for managing the scaling and availability of the virtual appliances behind Gateway Load Balancer.
Scale	AWS Network Firewall endpoints are powered by AWS PrivateLink. Network Firewall supports up to 100 Gbps of network traffic per firewall endpoint.	Gateway Load Balancer endpoints support maximum bandwidth of up to 100 Gbps per endpoint
Cost	AWS Network Firewall endpoint cost + Data processing charges	Gateway Load Balancer + Gateway Load Balancer endpoints + virtual appliances + data processing charges

Centralized inbound inspection

Internet-facing applications, by their nature, have a larger attack surface and are exposed to categories of threats most other types of applications don't have to face. Having the necessary protection from attacks on these types of applications, and minimizing the impact surface area, are a core part of any security strategy.

As you deploy applications in your Landing Zone, many apps will be accessed by the users over the public internet (for example, through a Content Delivery Network (CDN), or through a public facing web-application) via a public facing load balancer, API gateway or directly through an internet gateway. You can secure your workloads and applications in this case by using AWS Web Application Firewall (AWS WAF) for Inbound Application Inspection, or alternatively IDS/IPS Inbound Inspection using Gateway Load Balancer or AWS Network Firewall.

As you continue to deploy applications in your Landing Zone, you might have a requirement to inspect inbound internet traffic. You can achieve this in multiple ways, using either distributed, centralized, or combined inspection architectures using Gateway Load Balancer running your third-party firewall appliances or AWS Network Firewall with advance DPI and IDS/IPS capabilities through the use of open source Suricata rules. This section covers both Gateway Load Balancer and AWS Network Firewall in a centralized deployment, using AWS Transit Gateway acting as a central hub for routing traffic.

AWS WAF and AWS Firewall Manager for inspecting inbound traffic from the internet

AWS WAF is a web application firewall that helps protect your web applications or APIs against common web exploits and bots that may affect availability, compromise security, or consume excessive resources. AWS WAF gives you control over how traffic reaches your applications by enabling you to create security rules that control bot traffic and block common attack patterns, such as SQL injection or cross-site scripting (XSS). You can also customize rules that filter out specific traffic patterns.

You can deploy AWS WAF on Amazon CloudFront as part of your CDN solution, the Application Load Balancer that fronts your web servers, Amazon API Gateway for your REST APIs, or AWS AppSync for your GraphQL APIs.

Once you deploy AWS WAF, you can then create your own traffic filter rules using the visual rule builder, code in JSON, managed rules maintained by AWS, or you can subscribe to third-party rules from the AWS Marketplace. These rules can filter out unwanted traffic by evaluating the traffic against the specified patterns. You can further use Amazon CloudWatch for monitoring incoming traffic metrics and logging.

For centralized management across all your accounts and applications in AWS Organizations, you can use AWS Firewall Manager. AWS Firewall Manager is a security management service which allows you to centrally configure and manage firewall rules. As your new applications are created, AWS Firewall Manager makes it easy to bring new applications and resources into compliance by enforcing a common set of security rules.

Using AWS Firewall Manager, you can easily roll out AWS WAF rules for your Application Load Balancers, API Gateway instances, and Amazon CloudFront distributions. AWS Firewall Manager integrates with AWS Managed Rules for AWS WAF, which gives you an easy way to deploy pre-configured, curated AWS WAF rules on your applications. For more information on centrally managing AWS WAF with AWS Firewall Manager, refer to Centrally manage AWS WAF (API v2) and AWS Managed Rules at scale with AWS Firewall Manager.

Centralized inbound traffic inspection using AWS WAF

In the preceding architecture, applications are running on Amazon EC2 instances in multiple availability zones in the private subnets. There is a public-facing Application Load Balancer (ALB) deployed in front of the Amazon EC2 instances, load balancing the requests between different targets. The AWS WAF is associated with the ALB.

Advantages

- With AWS WAF Bot Control, you get visibility and control over common and pervasive bot traffic to your applications.

- With Managed Rules for AWS WAF, you can quickly get started and protect your web application or APIs against common threats. You can select from many rule types, such as those that address issues such as the Open Web Application Security Project (OWASP) Top 10 security risks, threats specific to Content Management Systems (CMS) like WordPress or Joomla, or even emerging Common Vulnerabilities and Exposures (CVE). Managed rules are automatically updated as new issues emerge, so that you can spend more time building applications.

- AWS WAF is a managed service and no appliance is needed for inspection in this architecture. In addition, it provides near real-time logs through Amazon Data Firehose. AWS WAF gives near real-time visibility into your web traffic, which you can use to create new rules or alerts in Amazon CloudWatch.

Key considerations

- This architecture is best suited for HTTP header inspection and distributed inspections, because AWS WAF is integrated on a per-ALB, CloudFront distribution and API Gateway. AWS WAF does not log the request body.

- Traffic going to a second set of ALB (if present) may not get inspected by the same AWS WAF instance; because a new request would be made to the second set of ALB.

Centralized inbound inspection with third-party appliances

In this architectural design pattern, you deploy third-party firewall appliances on Amazon EC2 across multiple availability zones behind an Elastic Load Balancer (ELB) such as an Application/ Network Load Balancer in a separate Inspection VPC.

The Inspection VPC along with other Spoke VPCs are connected together through a Transit Gateway as VPC attachments. The applications in Spoke VPCs are frontend by an internal ELB

which can be either ALB or NLB depending on the application type. The clients over the internet connect to the DNS of the external ELB in the inspection VPC which routes the traffic to one of the Firewall appliances. The Firewall inspects the traffic and then routes the traffic to the Spoke VPC through Transit Gateway using the DNS of the internal ELB as shown in the following figure. For more information regarding inbound security inspection with third-party appliances, refer to the How to integrate third-party firewall appliances into an AWS environment blog post.
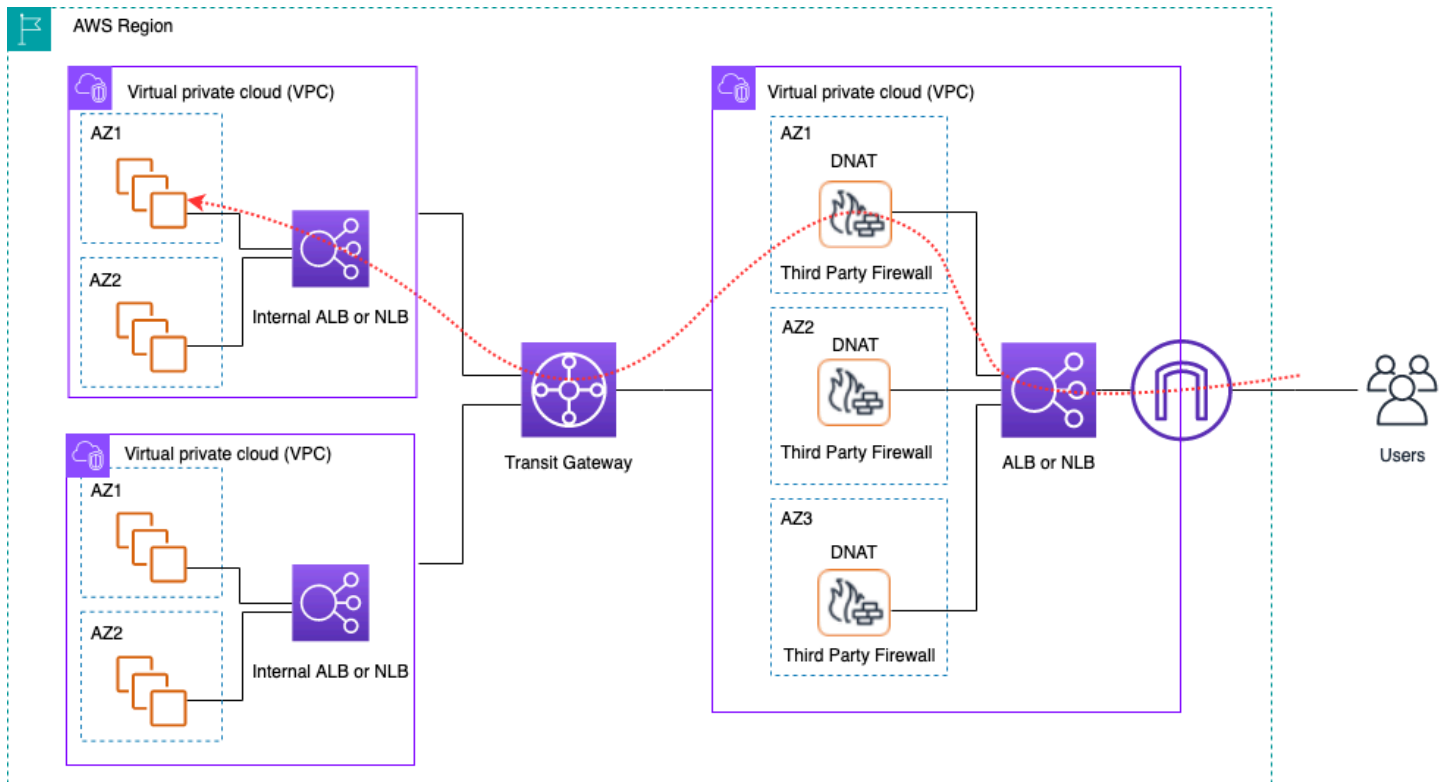
Centralized ingress traffic inspection using third-party appliances and ELB

Advantages

- This architecture can support any application type for inspection and advanced inspection capabilities offered through third-party firewall appliances.

- This pattern supports DNS based routing from firewall appliances to spoke VPCs, which allows the applications in Spoke VPCs to scale independently behind an ELB.

- You can use Auto Scaling with the ELB to scale the firewall appliances in the Inspection VPC.

Key considerations

- You need to deploy multiple firewall appliances across Availability Zones for high availability.

- The firewall needs to be configured with and perform Source NAT in order to maintain flow symmetry, which means the client IP address won't be visible to the application.

- Consider deploying Transit Gateway and Inspection VPC in the Network Services account.

- Additional third-party vendor firewall licensing/support cost. Amazon EC2 charges are dependent on instance type.

Inspecting inbound traffic from the internet using firewall appliances with Gateway Load Balancer

Customers use third-party next-generation firewalls (NGFW) and intrusion prevention systems (IPS) as part of their defense in depth strategy. Traditionally these often are dedicated hardware or software/virtual appliances. You can use Gateway Load Balancer to scale these virtual appliances horizontally to inspect traffic from and to your VPC, as shown in the following figure.
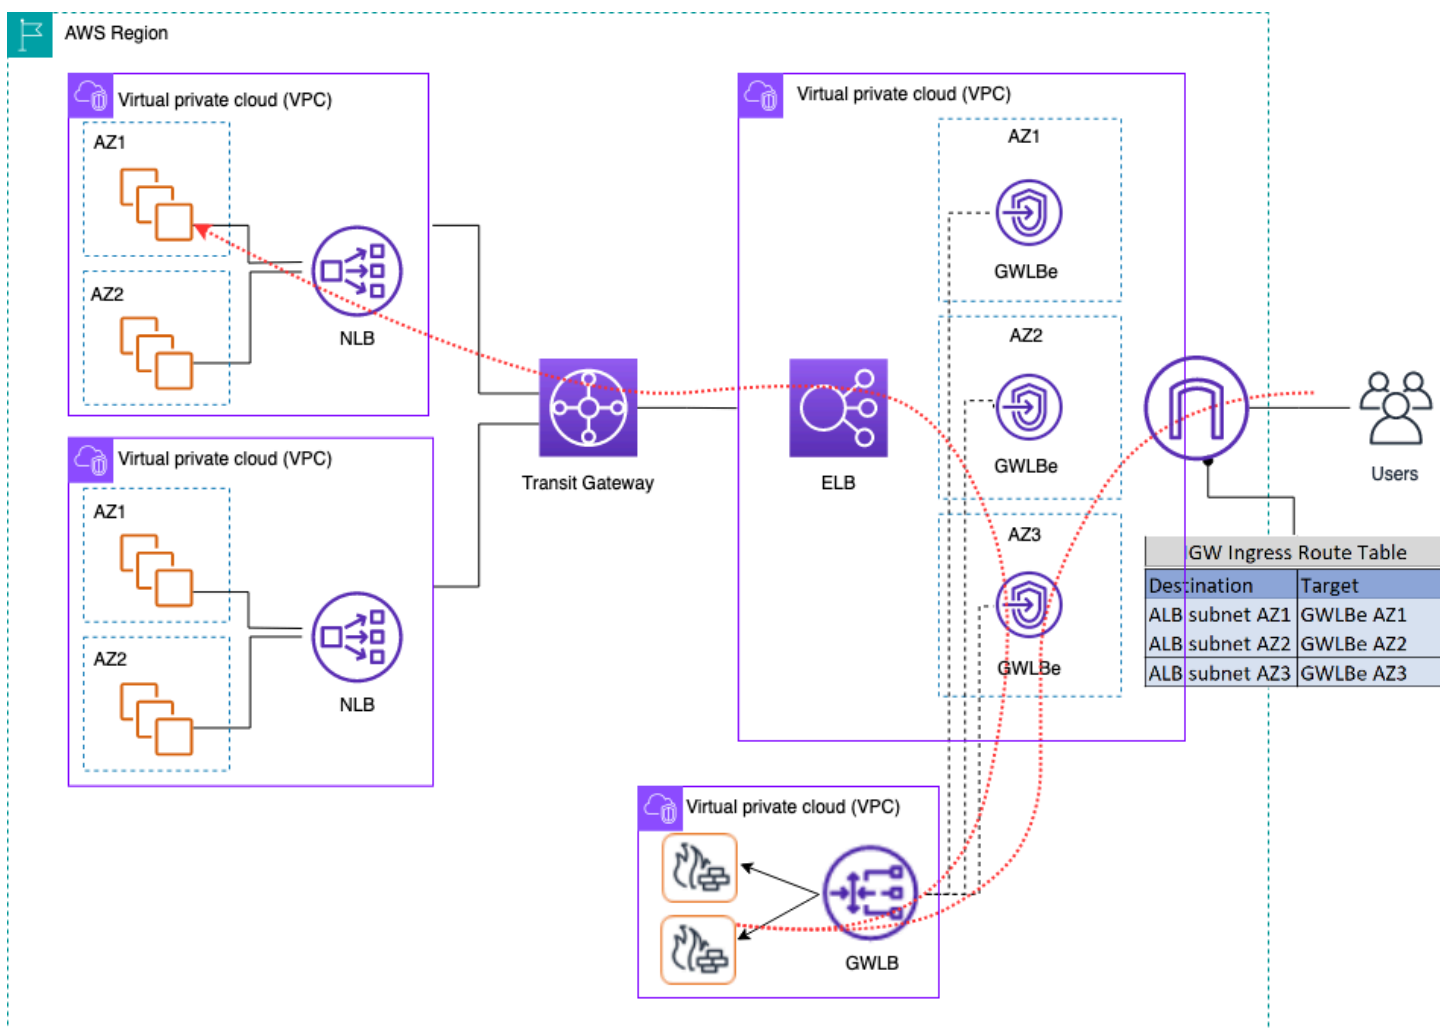
Centralized ingress traffic inspection using firewall appliances with Gateway Load Balancer

In the preceding architecture, Gateway Load Balancer endpoints are deployed into each Availability Zone in a separate edge VPC. The next-generation firewalls, intrusion prevention systems etc. are deployed behind the Gateway Load Balancer in the centralized appliance VPC. This appliance VPC can be in the same AWS account as the spoke VPCs or different AWS account. Virtual appliances can be configured to use Auto Scaling groups and are registered automatically with the Gateway Load Balancer, allowing auto scaling of the security layer.

These virtual appliances can be managed by accessing their management interfaces through an Internet Gateway (IGW) or using a bastion host setup in the appliance VPC.

Using the VPC ingress routing feature, the edge route table is updated to route inbound traffic from internet to firewall appliances behind Gateway Load Balancer. Inspected traffic is routed via Gateway Load Balancer endpoints to target VPC instance. Refer to the Introducing AWS Gateway Load Balancer: Supported architecture patterns blog post for details on various ways to use Gateway Load Balancer.

Using the AWS Network Firewall for centralized ingress

In this architecture, ingress traffic is inspected by AWS Network Firewall before reaching the rest of the VPCs. In this setup, traffic is split among all firewall endpoints deployed in the Edge VPC. You deploy a public subnet between the firewall endpoint and the Transit Gateway subnet. You can use an ALB or NLB, which contain IP targets in your spoke VPCs while handling Auto Scaling for targets behind them.
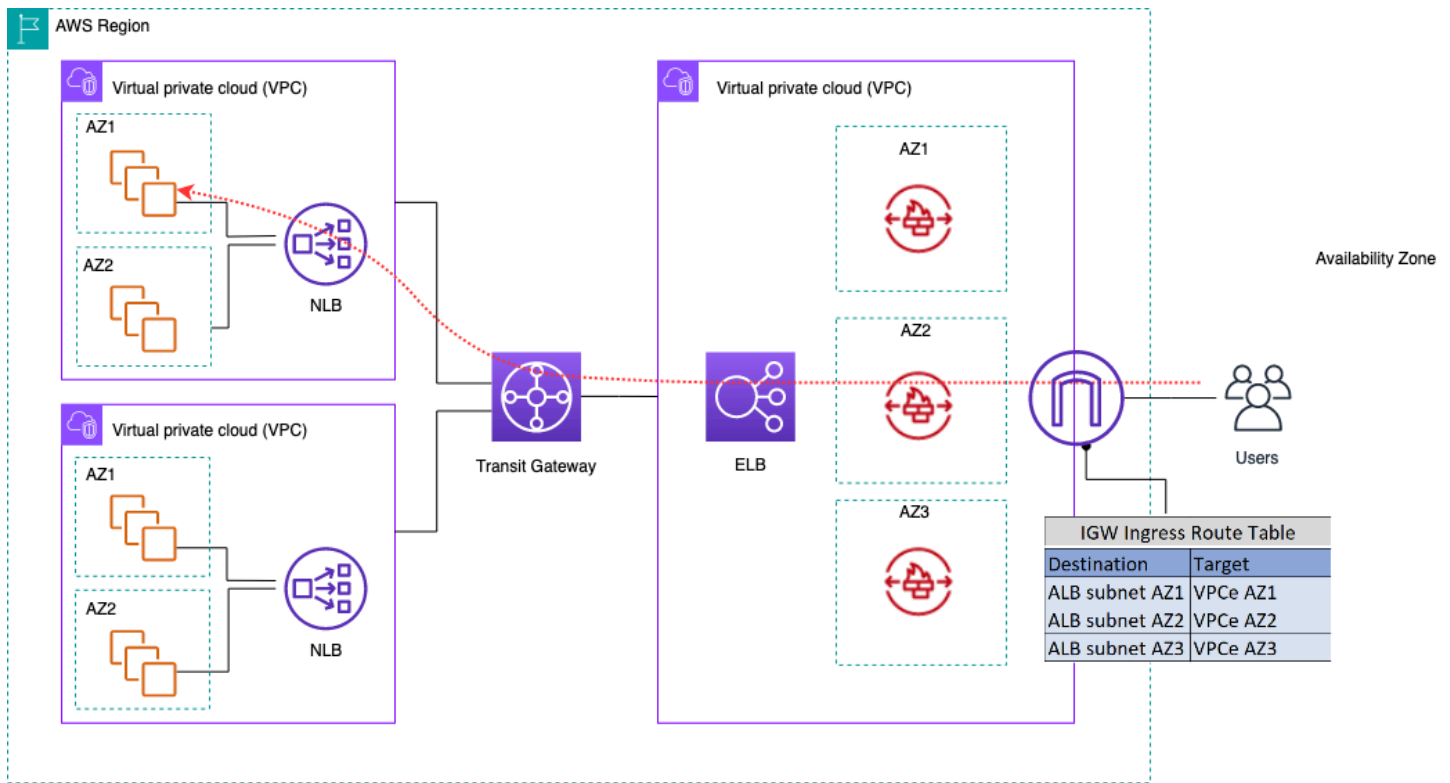
Ingress traffic inspection using AWS Network Firewall

To simplify deployment and management of AWS Network Firewall in this model, AWS Firewall Manager can be used. Firewall Manager allows you to centrally administer your different firewalls by automatically applying protection you create in the centralized location to multiple accounts. Firewall Manager supports both distributed and centralized deployment models for Network Firewall. The blog post How to deploy AWS Network Firewall by using AWS Firewall Manager provides more details on the model.

Deep Packet Inspection (DPI) with AWS Network Firewall

Network Firewall can perform deep packet inspection (DPI) on ingress traffic. Using a Transport Layer Security (TLS) certificate stored in AWS Certificate Manager (ACM), Network Firewall can decrypt packets, perform DPI, and re-encrypt packets. There are a few considerations for setting up DPI with Network Firewall. First, a trusted TLS certificate must be stored in ACM. Second, Network Firewall rules must be configured to correctly send packets for decryption and re-encryption. Refer to the blog post TLS inspection configuration for encrypted traffic and AWS Network Firewall for more details.

Key considerations for AWS Network Firewall in a centralized ingress architecture

- Elastic Load Balancing in Edge VPC can only have IP addresses as target types, not a hostname. In the preceding figure, the targets are the private IPs of the Network Load Balancer in spoke VPCs. Using IP targets behind the ELB in the edge VPC results in the loss of Auto Scaling.

- Consider using AWS Firewall Manager as a single pane of glass for your firewall endpoints.

- This deployment model uses traffic inspection right as it enters the edge VPC, so it has the potential to reduce the overall cost of your inspection architecture.

DNS

When you launch an instance into a VPC, excluding the default VPC, AWS provides the instance with a private DNS hostname (and potentially a public DNS hostname) depending on the DNS attributes you specify for the VPC and if your instance has a public IPv4 address. When the `enableDnsSupport` attribute is set to `true`, you get a DNS resolution within the VPC from Route 53 Resolver (+2 IP offset to the VPC CIDR). By default, Route 53 Resolver answers DNS queries for VPC domain names such as domain names for EC2 instances or Elastic Load Balancing load balancers. With VPC peering, hosts in one VPC can resolve public DNS hostnames to private IP addresses for instances in peered VPCs, provided the option to do so is enabled. The same is applicable for VPCs connected via AWS Transit Gateway. For more information, refer to Enabling DNS Resolution Support for a VPC Peering Connection.

If you want to map your instances to a custom domain name, you can use Amazon Route 53 to create a custom DNS-to-IP-mapping record. An Amazon Route 53 hosted zone is a container that holds information about how you want Amazon Route 53 to respond to DNS queries for a domain and its subdomains. Public Hosted Zones contain DNS information that is resolvable over the public internet while Private Hosted Zones are a specific implementation that only presents information to VPCs that have been attached to the specific private hosted zone. In a Landing Zone setup where you have multiple VPCs or accounts, you can associate a single private hosted zone with multiple VPCs across AWS accounts and across Regions (doable with SDK/CLI/API only). The end hosts in the VPCs use their respective Route 53 Resolver IP (+2 offset the VPC CIDR) as the name server for DNS queries. The Route 53 Resolver in VPC accepts DNS queries only from resources within a VPC.

Hybrid DNS

DNS is a critical component of any infrastructure, hybrid or otherwise, as it provides the hostname-to-IP-address resolution that applications rely on. Customers implementing hybrid environments usually have a DNS resolution system already in place, and they want a DNS solution that works in tandem with their current system. Native Route 53 resolver (+2 offset of the base VPC CIDR) is not reachable from on-premises networks using VPN or AWS Direct Connect. Therefore, when you integrate DNS for the VPCs in an AWS Region with DNS for your network, you need a Route 53 Resolver inbound endpoint (for DNS queries that you are forwarding to your VPCs) and a Route 53 Resolver outbound endpoint (for queries that you are forwarding from your VPCs to your network).

As shown in the following figure, you can configure outbound Resolver endpoints to forward queries it receives from Amazon EC2 instances in your VPCs to DNS servers on your network. To forward selected queries, from a VPC to an on-premises network, create Route 53 Resolver rules that specify the domain names for the DNS queries that you want to forward (such as example.com), and the IP addresses of the DNS resolvers on your network where you want to forward the queries. For inbound queries from on-premises networks to Route 53 hosted zones, DNS servers on your network can forward queries to inbound Resolver endpoints in a specified VPC.
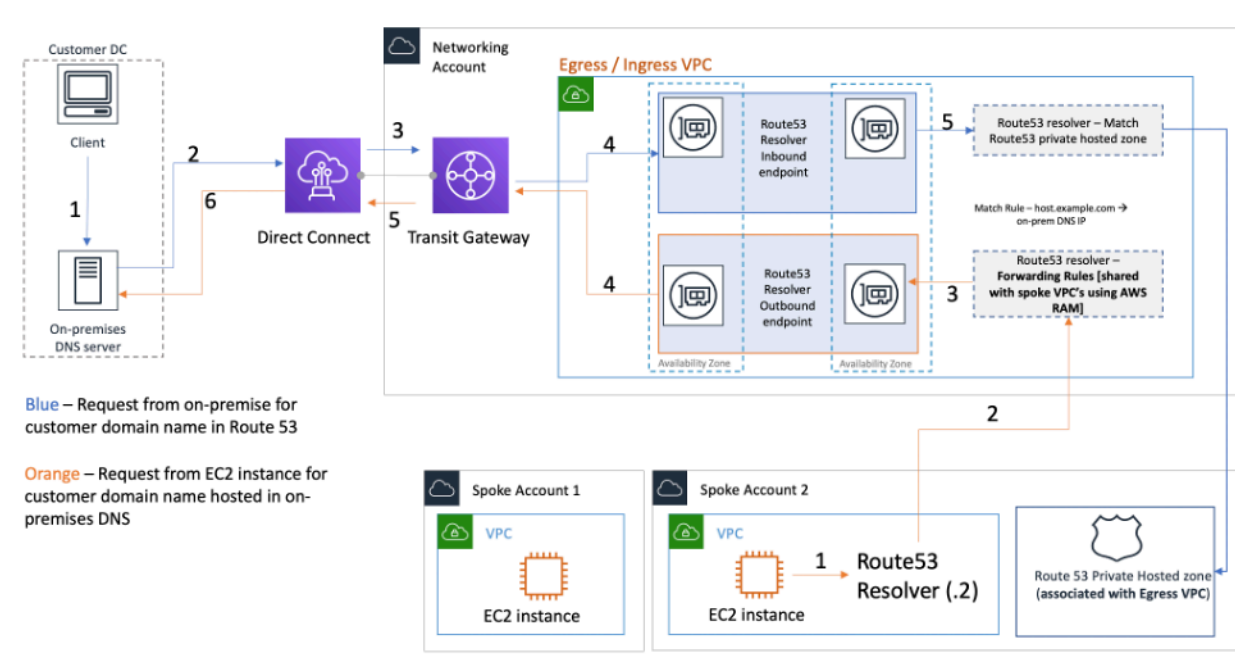

Hybrid DNS resolution using Route 53 Resolver

This allows your on-premises DNS resolvers to easily resolve domain names for AWS resources, such as Amazon EC2 instances or records in a Route 53 private hosted zone associated with that VPC. In addition, Route 53 Resolver endpoints can handle up to approximately 10,000 queries per second per ENI, so can scale to much larger DNS query volume easily. Refer to Best practices for Resolver in the Amazon Route 53 documentation for more details.

It is not recommended that you create Route 53 Resolver endpoints in every VPC of the Landing Zone. Centralize them in a central egress VPC (in the Network services account). This approach allows for better manageability while keeping the costs low (you are charged an hourly fee for each inbound/outbound resolver endpoint you create). You share the centralized inbound and outbound endpoint with the rest of the Landing Zone.

- **Outbound resolution** — Use the Network Services account to write resolver rules (based on which DNS queries will be forwarded to on-premises DNS servers). Using Resource Access Manager (RAM), share these Route 53 Resolver rules with multiple accounts (and associate with VPCs in the accounts). EC2 instances in spoke VPCs can send DNS queries to Route 53 Resolver

and Route 53 Resolver Service will forward these queries to on-premises DNS server via the outbound Route 53 Resolver endpoints in the egress VPC. You don't need to peer spoke VPCs to the egress VPC, or connect them via Transit Gateway. Do not use the IP of the outbound resolver endpoint as the primary DNS in the spoke VPCs. Spoke VPCs should use Route 53 Resolver (to offset of the VPC CIDR) in their VPC.

Centralizing Route 53 Resolver endpoints in ingress/egress VPC

- **Inbound DNS resolution** – Create Route 53 Resolver inbound endpoints in a centralized VPC and associate all the private hosted zones in your Landing Zone with this centralized VPC. For more information, refer to Associating More VPCs with a Private Hosted Zone. Multiple Private Hosted Zones (PHZ) associated with a VPC cannot overlap. As shown in the preceding figure, this association of PHZ with the centralized VPC will enable on-premises servers to resolve DNS for any entry in any private hosted zone (associated with central VPC) using the inbound endpoint in the centralized VPC. For further more information on hybrid DNS setups, refer to Centralized DNS management of hybrid cloud with Amazon Route 53 and AWS Transit Gateway and Hybrid Cloud DNS Options for Amazon VPC.

Route 53 DNS Firewall

Amazon Route 53 Resolver DNS Firewall helps filter and regulate outbound DNS traffic for your VPCs. A primary use of the DNS Firewall is to help prevent data exfiltration of your data by defining domain name allow-lists which allow resources in your VPC to make outbound DNS requests only for the sites your organization trusts. It also gives customers the ability to create *blocklists* for domains they don't want resources inside a VPC to communicate with via DNS. Amazon Route 53 Resolver DNS firewall has the following features:

Customers can create rules to define how DNS queries are answered. The actions that can be defined for the domain names include NODATA, OVERRIDE and NXDOMAIN.

Customers can create alerts for both allow-lists and deny-lists to monitor the rule activity. This can come in handy when customers want to test the rule before moving it to production.

For more information, refer to the How to Get Started with Amazon Route 53 Resolver DNS Firewall for Amazon VPC blog post.

Centralized access to VPC private endpoints

A VPC endpoint allows you to privately connect your VPC to supported AWS services without requiring an internet gateway or a NAT device, VPN connection, or AWS Direct Connect connection. Therefore, your VPC is not exposed to the public internet. Instances in your VPC do not require public IP addresses to communicate with AWS service endpoints with this interface endpoint. Traffic between your VPC and other services does not leave the AWS network backbone. VPC endpoints are virtual devices. They are horizontally scaled, redundant, and highly available VPC components. Two types of endpoints can currently be provisioned: interface endpoints (powered by AWS PrivateLink) and gateway endpoints. Gateway endpoints can be utilized to access Amazon S3 and Amazon DynamoDB services privately. There is no additional charge for using gateway endpoints. Standard charges for data transfer and resource usage apply.

Interface VPC endpoints

An interface endpoint consists of one or more elastic network interfaces with a private IP address that serves as an entry point for traffic destined to a supported AWS service. When you provision an interface endpoint, a cost is incurred for every hour the endpoint is running along with data processing charges. By default, you create an interface endpoint in every VPC from which you want to access the AWS service. This can be cost prohibitive and challenging to manage in the Landing Zone setup where a customer wants to interact with a specific AWS service across multiple VPCs. To avoid this, you can host the interface endpoints in a centralized VPC. All the spoke VPCs will use these centralized endpoints via Transit Gateway.

When you create a VPC endpoint to an AWS service, you can enable private DNS. When enabled, the setting creates an AWS managed Route 53 private hosted zone (PHZ), which enables the resolution of public AWS service endpoint to the private IP of the interface endpoint. The managed PHZ only works within the VPC with the interface endpoint. In our setup, when we want spoke VPCs to be able to resolve VPC endpoint DNS hosted in a centralized VPC, the managed PHZ won't work. To overcome this, disable the option that automatically creates the private DNS when an interface endpoint is created. Next, manually create a Route 53 private hosted zone matching the service endpoint name and add an Alias record with the full AWS service endpoint name pointing to the interface endpoint.

1. Log in to AWS Management Console and navigate to Route 53.

2. Select the private hosted zone and navigate to **Create Record**.

3. Populate the **Record Name** field, select Record Type as **A**, and enable **Alias**.

   Note that some services, such as the [Docker and OCI client endpoints](#) (`dkr.ecr`), require a wildcard alias (*) be used for the **Record Name**.

4. Under **Route Traffic to section**, select the service to which the traffic should be sent and select the region from the dropdown list.

5. Select the appropriate routing policy and enable the option to **Evaluate target health**.

You [associate](#) this private hosted zone with other VPCs within the Landing Zone. This configuration allows the spoke VPCs to resolve the full-service endpoint names to interface endpoints in the centralized VPC.

> ⓘ **Note**
>
> To access the shared private hosted zone, the hosts in the spoke VPCs should use the Route 53 Resolver IP of their VPC. Interface endpoints are also accessible from on-premises networks over VPN and Direct Connect. Use conditional forwarding rules to send all DNS traffic for the full-service endpoint names to Route 53 Resolver inbound endpoints, which will resolve DNS requests according to the private hosted zone.

In the following figure, Transit Gateway enables traffic flow from the spoke VPCs to the centralized interface endpoints. Create VPC Endpoints and the private hosted zone for it in Network Services Account and share it with spoke VPCs in the spoke accounts. For more details on sharing endpoint information with other VPCs, refer to the [Integrating AWS Transit Gateway with AWS PrivateLink and Amazon Route 53 Resolver](#) blog post.

> ⓘ **Note**
>
> A distributed VPC endpoint approach that is, an endpoint per VPC allows you to apply least privilege policies on VPC endpoints. In a centralized approach, you will apply and manage policies for all spoke VPC access on a single endpoint. With growing number of VPCs, the complexity of maintaining least privilege with a single policy document might grow. Single policy document also results in larger blast radius. You are also restricted on the [size of the policy document](#) (20,480 characters).

Centralizing interface VPC endpoints

Cross Region endpoint access

When you want multiple VPCs set up in different Regions that share a common VPC endpoint, use a PHZ, as outlined earlier. Both VPCs in each Region will be associated with the PHZ with the alias to the endpoint. In order to route traffic between VPCs in a multi-Region architecture, the Transit Gateways in each Region need to be peered together. For more information, refer to this blog: Using Route 53 Private Hosted Zones for Cross-account Multi-Region Architectures.

VPCs from different regions can be routed to one another using either Transit Gateways or VPC Peering. Use the following documentation for peering Transit Gateways: Transit gateway peering attachments.

In this example, the Amazon EC2 instance in the VPC `us-west-1` Region will use the PHZ to get the private IP address of the endpoint in the `us-west-2` Region and route the traffic to the `us-west-2` Region VPC over the Transit Gateway peering or VPC peering. Using this architecture, the traffic remains within the AWS network, securely allowing the EC2 instance in `us-west-1` to access the VPC service in `us-west-2` without going over the internet.

Multi-Region VPC endpoints

> **ⓘ Note**
>
> Inter-Region data transfer charges do apply when accessing endpoints across Regions.

Referring to the previous figure, an endpoint service is created in a VPC in the `us-west-2` Region. This endpoint service provides access to an AWS service in that Region. In order for your instances in another Region (such as `us-east-1`) to access the endpoint in the `us-west-2` Region, you need to create an address record in the PHZ with an alias to the desired VPC endpoint.

First, make sure that the VPCs in each Region are associated with the PHZ that you created.

When deploying an endpoint in multiple Availability Zone, the IP address of the endpoint returned from DNS will be from any of the subnets in the Availability Zone that is allocated.

When invoking the endpoint, use the fully qualified domain name (FQDN) that is in the PHZ.

AWS Verified Access

AWS Verified Access delivers secure access to applications in private network without a VPN. It evaluates requests in real time like identity, device, and location. This service grants access based on policy for applications and connecting the users by improving security of the organization. Verified Access provide access to private applications by acting as an identity aware reverse proxy. User identity and device health, if applicable, are performed before routing traffic to application.

The following diagram provides a high-level overview of Verified Access. Users send requests to access an application. Verified Access evaluates the request against the access policy for the group and any application-specific endpoint policies. If access is allowed, the request is sent to the application through the endpoint.
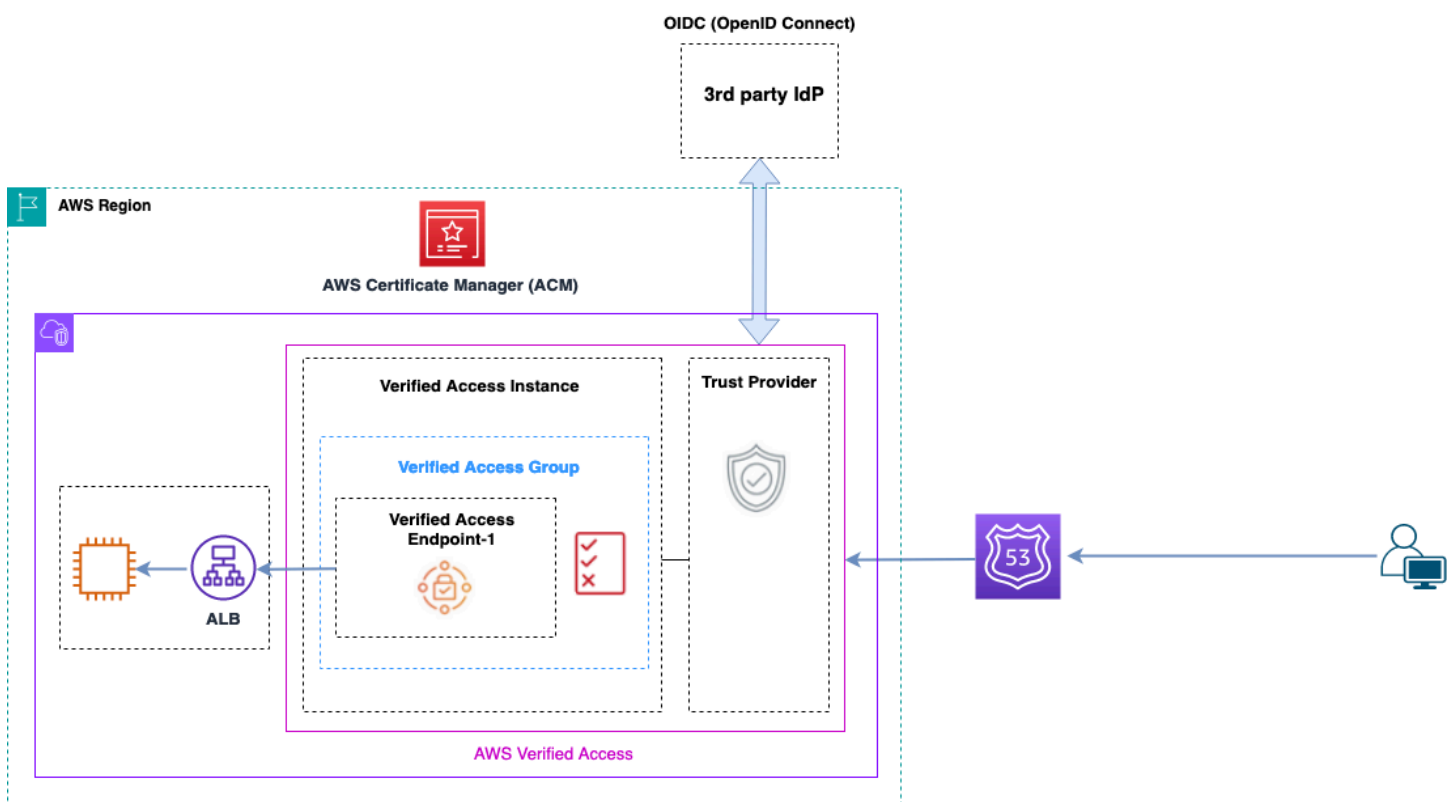
Verified Access overview

The main components in an AWS Verified Access architecture are:

- **Verified Access instances** – An instance evaluates application requests and grants access only when your security requirements are met.
- **Verified Access endpoints** – Each endpoint represents an application. An endpoint can be NLB, ALB or network interface.

- **Verified Access group** – A collection of Verified Access endpoints. We recommend that you group the endpoints for applications with similar security requirements to simplify policy administration.

- **Access policies** – A set of user-defined rules that determine whether to allow or deny access to an application.

- **Trust providers** – Verified Access is a service that facilitates the management of user identities and device security states. It is compatible with both AWS and third-party trust providers, requiring that at least one trust provider be attached to each Verified Access instance. Each of these instances can include a single identity trust provider as well as multiple device trust providers.

- **Trust data** – The security data that your trust provider sends to Verified Access, such as a user's email address or the group they belong to, is evaluated against your access policies each time an application request is received.

More details can be found in [Verified Access blog posts](#).

Conclusion

As you scale your usage of AWS and deploy applications in the AWS Landing Zone, the number of VPCs and networking components increases. This whitepaper explained how you can manage this growing infrastructure ensuring scalability, high availability, and security while keeping costs low. Making the right design decisions when using services such as Transit Gateway, Shared VPC, AWS Direct Connect, VPC endpoints, Gateway Load Balancer, AWS Network Firewall, Amazon Route 53, and third-party software appliances becomes critical. It is important to understand the key considerations of each approach and work backwards from your requirements and analyze as to which option or combination of options fit you best.

Contributors

The following individuals contributed to this document:

- Sohaib Tahir, Solutions Architect, Amazon Web Services
- Shirin Bhambhani, Solutions Architect, Amazon Web Services
- Kunal Pansari, Solutions Architect, Amazon Web Services
- Eric Vasquez, Solutions Architect, Amazon Web Services
- Tushar Jagdale, Solutions Architect, Amazon Web Services
- Ameer Shariff, Solutions Architect, Amazon Web Services
- Glenn Davis, Solutions Architect, Amazon Web Services
- Nick Kniveton, Solutions Architect, Amazon Web Services
- Sidhartha Chauhan, Principal Solutions Architect, Amazon Web Services

Document history

To be notified about updates to this whitepaper, subscribe to the RSS feed.

Change	Description	Date
Major update	Updates throughout whitepaper for changes to CloudWAN, Amazon VPC Lattice, ENA Express, hybrid connectivity, AWS Direct Connect Sitelink, Deep Packet Inspection, and AWS Verified Access.	April 17, 2024
Minor update	Updated diagrams to be more consistent, updated DX connectivity options to include private IP VPN, and numerous minor changes throughout.	July 6, 2023
Minor update	Updated AWS Control Tower information, reflected new throughput limits for various services, updated NAT gateway diagram, updated security section for centralize egress.	April 4, 2023
Minor update	Added section: Cross Region endpoint access.	July 19, 2022
Major update	Updated Transit Gateway section with Transit Gateway Connect, updated Transit VPC section; updated AWS Direct	February 22, 2022

Connect section with MACsec and resiliency recommendations; updated AWS PrivateLink section. Added VPC peering vs. Transit VPC vs. Transit Gateway comparison table; added centralized inbound inspection section; updated centralized network security for VPC-to-VPC and VPC-on-premises to VPC and centralized egress to internet with AWS Network Firewall and Gateway Load Balancer design patterns; added private NAT gateway and Amazon Route 53 DNS Firewall sections.

[Minor update](#)	Updated *Transit Gateway vs VPC peering* section	April 2, 2021
[Whitepaper updated](#)	Corrected text to match the options illustrated in Figure 7	June 10, 2020
[Initial publication](#)	Whitepaper published.	November 15, 2019

Notices

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents current AWS product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers, or licensors. AWS products or services are provided "as is" without warranties, representations, or conditions of any kind, whether express or implied. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.