AWS GLOBAL TRANSIT NETWORK

"How do I build a global transit network on AWS?"

Overview

Amazon Virtual Private Cloud (Amazon VPC) provides customers with the ability to create as many virtual networks as they need, as well as different options for connecting those networks to each other and to non-AWS infrastructure. There are two common strategies for connecting multiple, geographically dispersed VPCs and remote networks: one is to implement a hub-and-spoke network topology that routes all traffic through a network transit hub; the other is to create a meshed network that uses individual connections between all networks. Both approaches can create an efficient and available transit network, each offering specific benefits and tradeoffs for different business needs.

This document addresses key considerations for implementing a global transit network on AWS, and provides general best practices and an overview of common transit network patterns. The following sections assume basic knowledge of AWS Transit Gateway, highly available remote-network connectivity,[1] IPsec VPNs, network addressing, subnetting, and routing.

General Best Practices

When creating transit networks, there are some universal network-design principles to consider. For example, the transit network will become a critical component of your network backbone, so choose network vendor products you are familiar with and comfortable supporting. With this in mind, consider the following AWS remote-connectivity best practices:

- Implement non-overlapping network ranges for your private networks to simplify the ability to route between remote networks. While it is possible to implement NAT rules in the transit network to compensate for overlapping networks, doing so adds additional complexity to the network design.
- Implement measures to ensure your network is highly available, resilient, and scalable. For example, leverage multiple dynamically routed, rather than statically routed, connections between networks to enable automatic failover between available connections, or use systems to monitor and manage network connectivity and availability in real time.
- Leverage a multi-account, multi-VPC strategy to guide your network design and layout, such as those provided by AWS Landing Zone or the AWS Control Tower service.

Application on AWS

The following sections provide high-level design overviews, including associated benefits and considerations, for creating either a hub-and-spoke network or a meshed network to directly route network traffic between global networks both on-premises and in the cloud. Implementing a global transit network virtually can reduce costs associated with colocation transit hubs or physical network gear. A global transit network is applicable to customers with the following use case/requirements:

- AWS resources in geographically dispersed VPCs need access to a wide variety of on-premises or remote infrastructure.
- Customer VPCs are located in different AWS Regions.
- Complex network-routing is required to implement a hybrid network architecture.
- Security or compliance programs require additional network-based monitoring or filtering between resources in different networks (e.g., Network Intrusion Detection Systems or next-generation firewalls).

Transit Networking Options

The following table provides a comparison of key attributes available in the AWS Transit Gateway and transit VPC offerings.

[1] See the *Resources* section for relevant Solution Briefs.

	AWS Transit Gateway	Transit VPC (via partner solution)
Deployment model	AWS native service	Customer-managed instances
Infrastructure used	Virtual private connections	VPN and virtual private gateways
Scalability	High horizontal scale	More difficult to manage at scale
VPC segmentation	Flexible	Varies depending on partner
Product integration(s)	Aviatrix Orchestrator for Transit Gateway; Terraform Provider	Aviatrix, Cisco, Juniper, Riverbed

Hub-and-Spoke Network (Transit VPC)

This approach uses either the AWS Transit Gateway service or host-based VPN appliances running in a dedicated VPC to perform routing between spoke networks through a central hub. Both the Transit Gateway and the transit VPC solution simplify network management and minimize the number of connections required to connect multiple VPCs and remote networks. AWS Transit Gateway avoids the need to route traffic through an Amazon EC2 instance reducing instance cost and bandwidth limitations of instances. On the other hand, Amazon EC2-based transit VPC solutions often provide additional security options, visibility, and edge connectivity options.

Configuration Details with AWS Transit Gateway

This design uses AWS Transit Gateway to provide scalable, high-bandwidth connectivity across VPCs and on-prem connections (including VPN connections). With Transit Gateway, you only have to create and manage an attachment from the central gateway in to each Amazon VPC, on-premises data center, or remote office across your network.

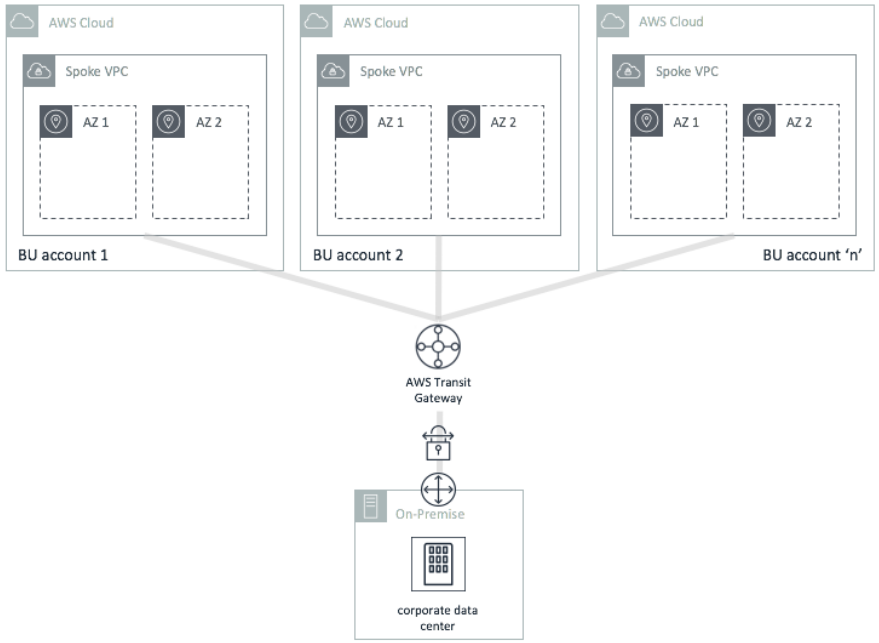

To connect spoke VPCs to the central gateway, use a Transit Gateway attachment. Then, update the VPC route tables to route traffic to other VPCs via the Transit Gateway. You can also segment your network by creating multiple route tables in a Transit Gateway, and associating and propagating Amazon VPCs and VPNs to them. This will allow you to create isolated networks inside a Transit Gateway similar to virtual routing and forwarding (VRFs) in traditional networks.

AWS Transit Gateway is inherently scalable and highly available. Multiple spoke VPCs in the same region can be connected to the same Transit Gateway for routing to other VPCs.

Configuration Details with Host-Based VPN Appliances (Transit VPC)

This design deploys VPN appliances on Amazon Elastic Compute Cloud (Amazon EC2) instances in separate Availability Zones of a transit VPC. We highly recommend leveraging virtual network appliances from the AWS Marketplace[2] to significantly reduce the level of effort to establish and maintain these VPN connections.

[2] For recommended products, search AWS Marketplace for one the following terms: Aviatrix, Cisco CSR 1000V, Fortinet FortiGate, Palo Alto Networks, Sophos UTM, Vyatta

Spoke VPCs are connected to the transit network through dynamically routed VPN connections between their virtual private gateways (VGWs) and the network appliances. This design uses VPN connections from spoke VPCs rather than VPC peering to enable routing between any connected network, including external networks or VPCs in other AWS Regions. This also allows spoke VPC resources to leverage VGW capabilites for routing and failover in order to maintain highly available network connections to the transit VPC network appliances. Remote networks also connect to the transit VPN appliances using redundant, dynamically routed VPN connections. Once connected, leverage dynamic routing protocols to automatically route traffic around potential network failures as well as to propogate network routes to remote networks.

Note that in the diagram to the right, all communication with the VPN appliances (including the VPN connection between the corporate data center and other provider networks and the transit VPC) uses the transit VPC Internet gateway and Elastic IP addresses. In addition to using dynamically routed connections, we highly recommend the use of Auto Recovery for EC2 to protect instances in the transit VPC.

Along with providing direct network routing between VPCs and on-premises networks, this design also enables the transit VPC to implement more complex routing rules, such as network address translation between overlapping network ranges, or to add additional network-level packet filtering or inspection.

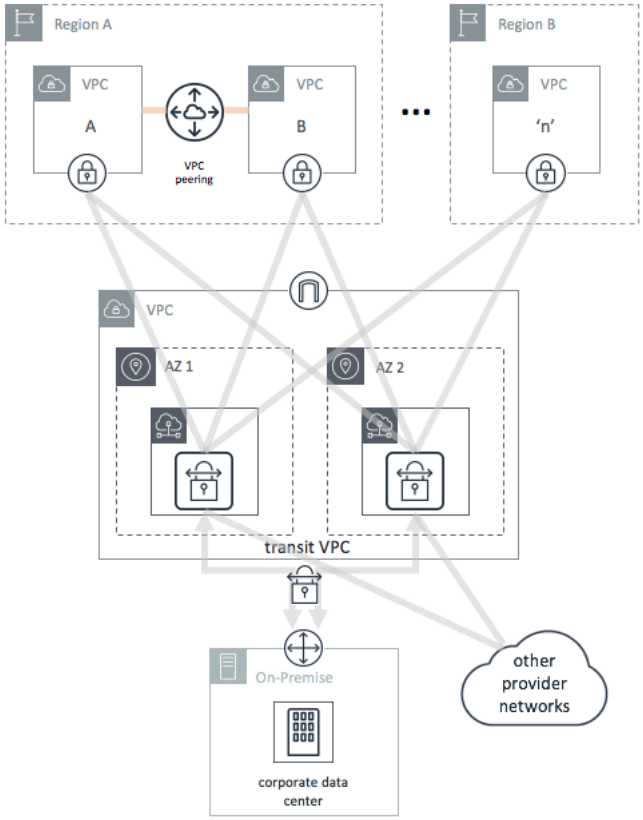

Considerations

This design supports any IP-based connectivity requirements between Amazon VPCs and remote resources with minimal on-premises network changes. It also provides an opportunity to select products available on the AWS Marketplace that integrate seamlessly with AWS-provided VPN connections, without the need to deploy these products into existing data centers. However, this design does require the customer to configure and manage the EC2-based VPN instances deployed in the transit VPC. This will result in additional Amazon EC2 and, potentially, third-party license charges. Also, be aware that this design will generate additional data-transfer charges for traffic traversing the transit VPC: data is charged when it is sent from a spoke VPC to the transit VPC, and again from the transit VPC to the on-premises network.

See the AWS Global Transit Network solution for information on how to deploy fully automated transit VPCs in minutes. The Cisco option actively monitors a customer's environment for specifically tagged VGWs to automatically join to the transit network. Also, it supports VPCs located in multiple AWS regions and in different AWS accounts. AWS Partner Network (APN) Partners Aviatrix and Juniper also provide automated solutions that allow customers to quickly and easily deploy a secure and managed transit VPC network. See the AWS Global Transit Network Resources & FAQ page for more information.

Meshed Network

A fully or partially meshed design uses individual VPN connections between networks without a central hub. This approach reduces the number of hops in the network which can reduce latency and simplify troubleshooting. The implementation and ongoing management of a meshed network can be more complex than the hub-and-spoke approach, but it might be more suitable for companies with large inter-regional data transfer needs, or who have extensive compliance requirements for network logging and monitoring.

Configuration Details

This design deploys VPN appliances on EC2 instances in each VPC which use fully (or partially) meshed point-to-point VPN connections to route network traffic. On-premises VPN devices can also join the transit network by creating individual VPN connections to VPN instances as needed. VPC route tables are configured to route transit network traffic through the VPN appliances to the intended destination network.

We highly recommend leveraging an APN Partner or AWS Marketplace offering that automatically provisions, manages, and monitors the availability of these networking instances and associated VPN connections.

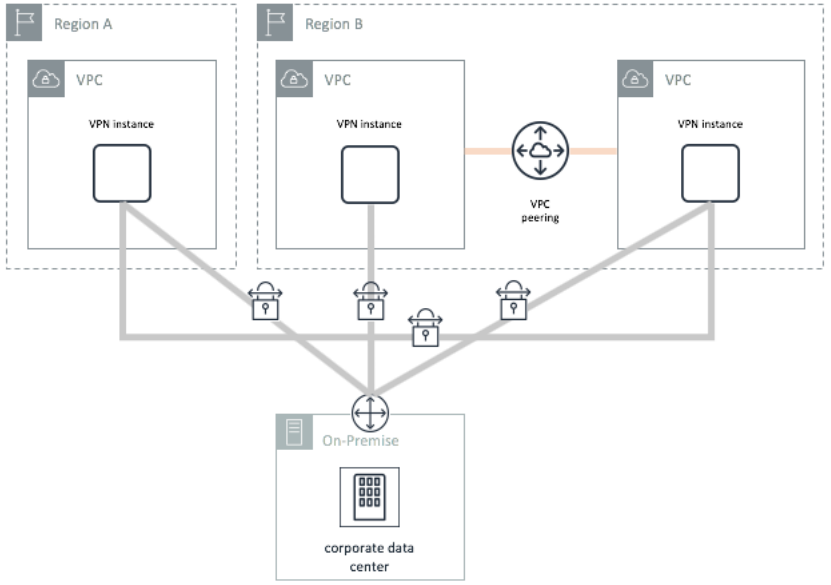

Considerations

This design does not rely on a central hub for routing all transit traffic, which allows it to scale more effectively by sending traffic directly to another VPC. However, it relies on EC2 instances to provide transit network connectivity, which will result in additional EC2 instance costs and introduces a single point of failure between a VPC and remote networks. It also requires more involved instance and VPN provisioning, monitoring, management, and recovery, which is why we recommend using a comprehensive partner offering, such as those offered by Aviatrix (Aviatrix Mesh Network) or Riverbed (SteelConnect) to automate and simplify these processes.

Resources

AWS Global Transit Network	https://aws.amazon.com/solutions/aws-global-transit-network/
AWS Transit Gateway	https://aws.amazon.com/transit-gateway/
Transit Gateways and Transit VPCs: Reference Architectures	https://www.youtube.com/watch?v=ar6sLmJ45xs
Amazon VPC Documentation	https://aws.amazon.com/documentation/vpc/
Multiple-VPC VPN Connection Sharing	https://aws.amazon.com/answers/networking/aws-multiple-vpc-vpn-connection-sharing/
Aviatrix Systems, Inc. Partner page	https://aws.amazon.com/partners/find/partnerdetails/?id=001E000001Dhh2sIAB
Riverbed Partner page	https://aws.amazon.com/partners/find/partnerdetails/?id=001E000000RI134IAB