# MULTIPLE REGION MULTI-VPC CONNECTIVITY

*"How do I connect multiple VPCs in different AWS Regions?"*

## Overview

Amazon Virtual Private Cloud (Amazon VPC) offers a comprehensive set of virtual networking capabilities that provide AWS customers with many options for designing and implementing networks on the AWS cloud. With Amazon VPC, you can provision logically isolated virtual networks to host your AWS resources. You can create multiple VPCs within the same region or in different regions, in the same account or in different accounts. This is useful for customers who require multiple VPCs for security, billing, regulatory, or other purposes, and want to integrate AWS resources between their VPCs more easily. More often than not, these different VPCs need to communicate privately and securely with one another for sharing data or applications.

This document provides AWS customers with high-level connectivity options for multiple VPCs that reside in different AWS Regions. It includes best practices and guidance, and outlines the most commonly used cross-region VPC network configurations. For guidance on connecting VPCs in the same AWS Region, see Single Region Multi-VPC Connectivity. The following sections address key considerations and recommendations for connecting VPCs in different regions, and assume basic knowledge of Amazon VPC, VPN connections, HA remote network connectivity, network addressing, subnetting, routing, Amazon Elastic Compute Cloud (Amazon EC2), and AWS Direct Connect.

## General Best Practices

When connecting multiple VPCs in different AWS Regions, there are some universal network-design principles to consider:

- Ensure that your VPC network ranges (CIDR blocks) do not overlap.
- Make sure the solution you choose is able to scale according to your current and future VPC connectivity needs.
- Ensure you implement a highly available (HA) design with no single point of failure.
- Consider your data-transfer needs, as this will affect the solution you choose. Some solutions proposed below may prove to be more expensive than others based on the amount of data transferred.
- Use network equipment that supports IPsec VPN tunnels and Border Gateway Protocol (BGP), when applicable.
- Connect only those VPCs that really need to communicate with each other.

## Application on AWS

The following sections offer prescriptive advice for connecting VPCs in different AWS Regions using either AWS-managed networks or non-AWS networks, such as the internet or a customer's existing network backbone. Customers who have already established AWS Direct Connect or VPN connections from on-premises networks to their VPCs typically prefer to reuse existing connections, while customers without existing network infrastructure usually opt for AWS-managed networks. If you design your VPC in adherence to networking best practices, you can easily change from one configuration to another, so select the option that makes most sense for your current networking needs.

### Routing Over AWS Networks

AWS provides high bandwidth, global network infrastructure to support your regional networking needs. AWS Regions are connected to multiple Internet Service Providers (ISPs) as well as to a private global network backbone, which provides lower cost and more consistent cross-region network latency when compared with the public internet. AWS customers with small on-premises network footprints, limited regional network connectivity, or who simply want to leverage AWS networks often choose one of the following options to connect VPCs in different regions.

## Inter-region VPC Peering

This approach leverages inter-region VPC peering connections to encrypt and route traffic between VPCs in different AWS Regions. A VPC peering connection uses the existing infrastructure of a VPC and private IP addresses; it is neither a gateway nor a VPN connection, and does not rely on a separate piece of physical hardware. There is no single point of failure for communication or a bandwidth bottleneck.

### Configuration Details

Inter-region VPC peering connections allow secure communication between VPC resources in different AWS Regions. All network traffic between regions is encrypted, stays on the AWS global network backbone, and never traverses the public internet, thereby reducing threat vectors, such as common exploits and DDoS attacks. VPC peering is appropriate for many scenarios, for example, to provide VPCs full access to each other's resources or to provide a set of VPCs partial access to resources in a central VPC. You can configure peering connections to provide access to part of a CIDR block or to an entire CIDR block of the peer VPC.

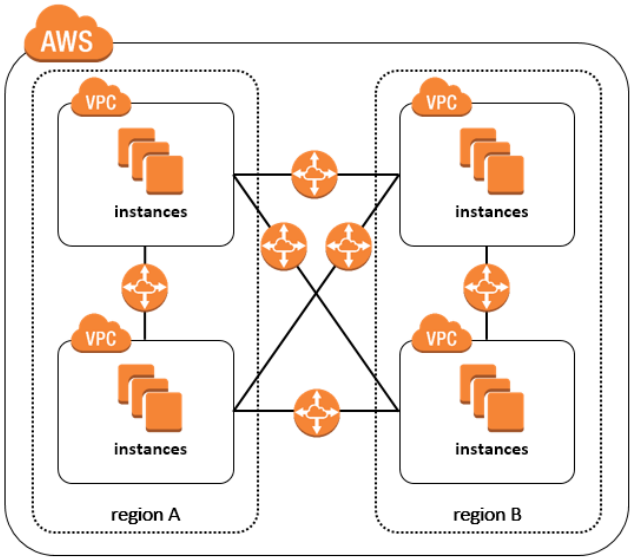


### Considerations

Inter-region VPC peering is available in specific AWS Regions only (see the Amazon VPC Peering Guide for current availability). VPC peering does not support transitive routing, so if you require many-to-many connections, use a fully meshed configuration to allow communication between multiple VPCs. You can peer VPCs with overlapping CIDR blocks to the same VPC, such as a central VPC, but it will require specific adjustments to your subnet route tables (see Configurations with Specific Routes for guidance). Keep in mind that a peering relationship does not allow you to extend these other VPC connection types: VPN or AWS Direct Connect connections to a corporate network; internet connections through an internet gateway; a VPC endpoint to an AWS service; a ClassicLink connection (see Invalid VPC Peering Connection Configurations for detailed information).

## Software VPN Appliances

This approach leverages the Amazon VPC capability to create VPN tunnels between EC2 instances in order to route traffic between VPCs in different AWS Regions. This option uses customer- or AWS Partner Network (APN) member-managed software VPN appliances, and is best suited to customers who want to manage both ends of VPN connections using their preferred VPN software provider. This design optimizes cross-region network transfer costs; however, it requires customers to design and manage their own HA configuration for EC2 network instances.

### Configuration Details

This design uses Elastic IP addresses and VPC internet gateways to facilitate communication between EC2-based software VPN appliances. Although EC2 instances are configured with public IP addresses, network traffic between AWS Regions traverse the AWS global network backbone by default. AWS Marketplace provides multiple third-party and open source options[1] for implementing software VPN appliances. Some of these provide advanced network automation for launching, managing, and configuring software-based VPN appliances to connect multiple VPCs and fully meshed VPN networks.

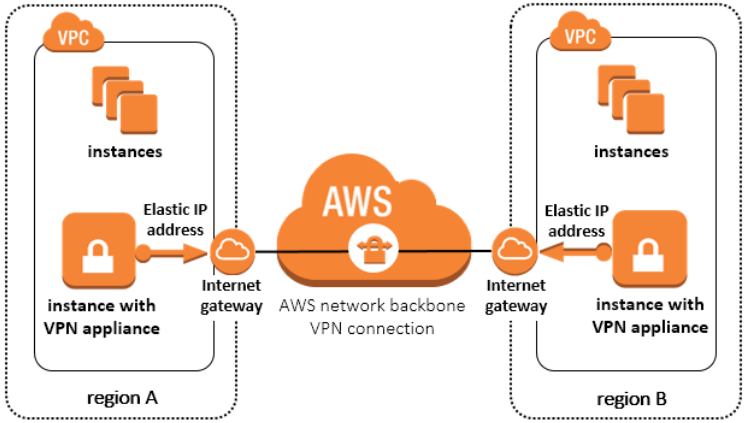


### Considerations

Although network traffic is routed over the AWS global network backbone by default, use of AWS private network infrastructure is provided on a best effort basis and network connectivity

[1] These include products from Aviatrix, Cisco, Fortinet, OpenSWAN, OpenVPN, Palo Alto Networks, Riverbed, Sophos, and Vyatta.

will fail over to AWS ISP networks in the unlikely event that private network connectivity between AWS Regions is not available.[2] You can implement AWS Marketplace products that provide additional network control and security features, such as enhanced monitoring, network-protocol-aware firewall rules, or universal threat management capabilities. To implement this design, you must run network appliances in each VPC, which results in additional EC2 and, potentially, third-party license charges. These EC2 instances can also introduce a single point of failure into the network architecture, and a potential network bottleneck, so be sure to choose a VPN appliance instance size that will meet cross-region network routing requirements. Finally, leverage Auto Recovery for EC2 or other network monitoring and recovery options to decrease the time to recover failed VPN appliances.

## Transit VPC

This approach creates a resilient, highly available transit network in a hub-and-spoke design that supports multiple VPCs and on-premises networks. It uses a dedicated transit VPC(s) hosting EC2-based VPN appliances to route traffic between VPC virtual private gateways (VGWs) in multiple AWS Regions and, optionally, on-premises networks. This option is best suited to customers who want to leverage AWS-provided, automated HA network connectivity features and also optimize their investments in third-party product licensing. Additionally, this design enables customers to control cross-region network traffic using AWS and third-party network security products.

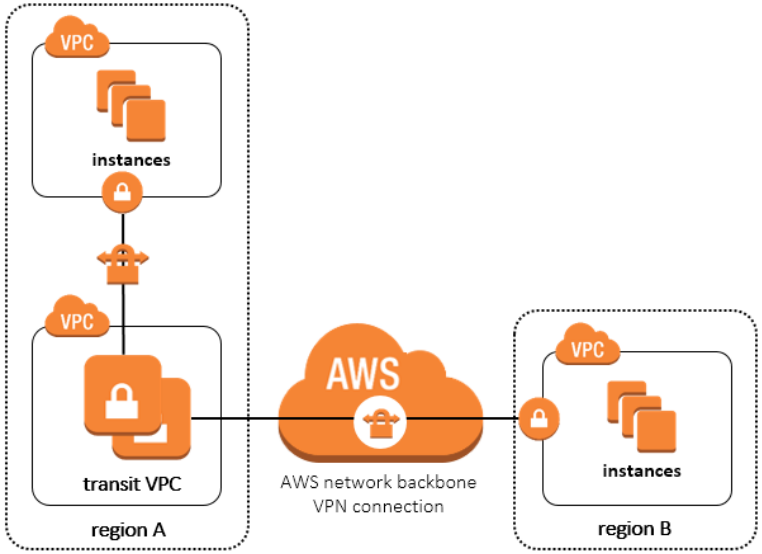


*Single-Region Transit VPC*

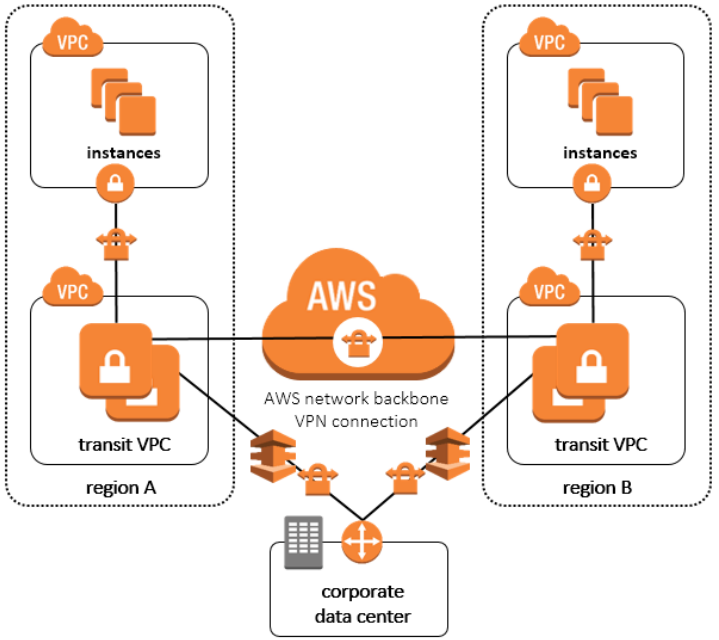


*Regional Transit VPCs Connected to Corporate Network*

### Configuration Details

This design pattern creates dynamically routed VPN connections between spoke VPC VGWs and VPN appliances in one or more transit VPCs (and optionally to on-premises network equipment). The best practice for making a transit network highly available and scalable is to use dynamically routed VPN connections to take advantage of VGW native HA capabilities and reduce network single points of failure. Note that in the diagrams above, all communication with the VPN appliances uses transit VPC internet gateways and Elastic IP addresses. Additionally, AWS highly recommends the use of Auto Recovery for EC2 or Auto Scaling for automatic recovery of failed EC2-based VPN instances.

[2] Note that, for simplicity, the diagrams in this section depict use of AWS backbone network connections and do not depict network failure scenarios where traffic might briefly be routed over public ISP networks.

### Considerations

You can use a transit VPC not only to provide direct network routing between VPCs and on-premises networks, but also 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. However, this approach requires you to configure and manage the EC2-based VPN instances deployed in the transit VPC. We highly recommend using virtual network appliances available in the AWS Marketplace[3] to significantly reduce the level of effort to establish and maintain these VPN connections. As with the previous approach, network traffic is routed over the AWS global network backbone by default, but will fail over to AWS ISP networks in the unlikely event that AWS private network infrastructure is unavailable. This design will result in additional 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.

## Routing Over Non-AWS Networks

Many AWS customers have invested heavily in on-premises networks, and use VPN connections or AWS Direct Connect to connect to AWS. These customers often prefer to leverage their existing infrastructure investments to route VPC traffic across regions, either by establishing new VPN connections or using their own corporate network backbone. AWS customers in this scenario often start by connecting on-premises networks to the closest AWS Region, and then look to expand to additional AWS Regions.

## Internet-Based VPN

This approach uses internet-based VPN connections and existing on-premises infrastructure in a single region to route VPC traffic between regions. This option is best suited for customers who do not have existing network infrastructure investment in the remote region and want to leverage their existing VPN investments to connect VPCs.

### Configuration Details

This design pattern creates a network hub to route traffic between spoke VPCs in different AWS Regions. Intra-region network connectivity is established using either a VPN connection or AWS Direct Connect (both options are portrayed in the diagram), and cross-region connectivity is established with internet-based VPN connections.

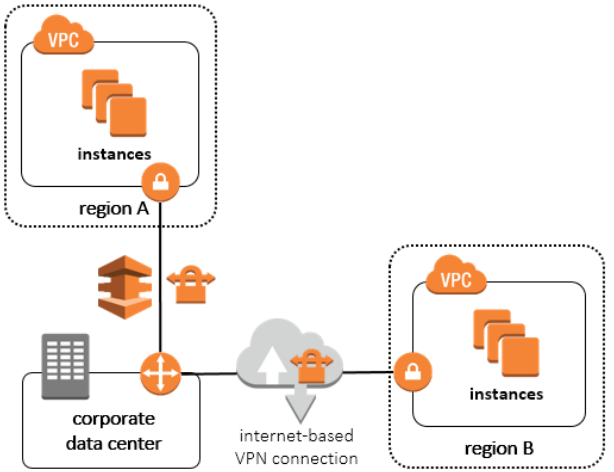


### Considerations

This design leverages existing on-premises network equipment and wide area network (WAN) connections to route network traffic between VPCs. It establishes cross-region VPC connectivity using existing internet connectivity and allows you to apply additional on-premises network monitoring and controls to inter-VPC traffic. Note that the cross-region network traffic is subject to the specific latency, variability, and available bandwidth of the existing internet connections.

## Corporate Network Backbone

This approach routes cross-region VPC traffic over existing corporate networks. This option is best suited for customers who can leverage existing global infrastructure and networks, such as a multinational company with network connectivity between data centers in multiple locations. It is also appropriate for customers who want to ensure their VPC traffic never traverses a public network.

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

**Configuration Details**

As in the previous approach, this design also establishes intra-region network connectivity to existing network infrastructure using either a VPN connection or AWS Direct Connect. Cross-region connectivity leverages existing wide area network (WAN) connections by propagating VPC network advertisements within the company's internal network. Finally, network routing in each VPC must be configured (through either dynamic BGP advertisements or static routes) to route cross-region traffic over the existing network backbone.

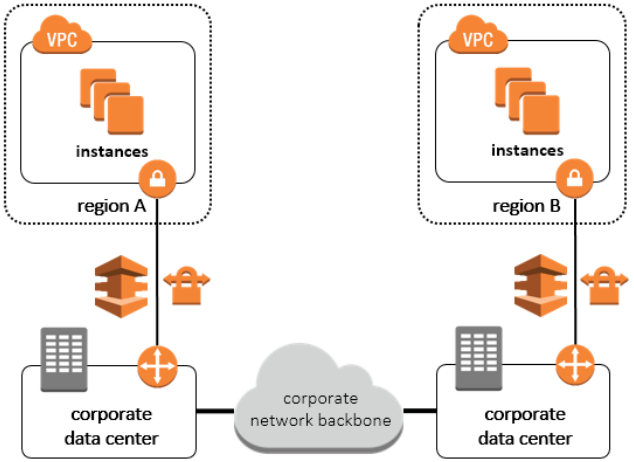


**Considerations**

This design leverages existing on-premises network equipment and wide area network (WAN) connections to route network traffic between VPCs. It also allows you to apply additional on-premises network monitoring and controls to inter-VPC traffic. Note that the cross-region network traffic is subject to the specific latency, variability, and available bandwidth of the existing private network connections.

## Resources

| | |
|---|---|
| Single Region Multi-VPC Connectivity | https://d0.awsstatic.com/aws-answers/AWS_Single_Region_Multi_VPC_Connectivity.pdf |
| AWS Global Transit Network | https://d1.awsstatic.com/aws-answers/aws-global-transit-network.pdf |
| Single Data Center HA Network Connectivity | https://d0.awsstatic.com/aws-answers/AWS_Single_Data_Center_HA_Network_Connectivity.pdf |
| Multiple Data Center HA Network Connectivity | https://d0.awsstatic.com/aws-answers/AWS_Multiple_Data_Center_HA_Network_Connectivity.pdf |
| Multiple-VPC VPN Connection Sharing | https://d0.awsstatic.com/aws-answers/AWS_Multiple_VPC_VPN_Connection_Sharing.pdf |
| AWS Direct Connect Product Website | https://aws.amazon.com/directconnect |
| VPC Product Website | http://aws.amazon.com/vpc/ |
| Amazon VPC Documentation | https://aws.amazon.com/documentation/vpc/ |