IPSec VPN for VMware Cloud on AWS

VMware General

Table of contents

IPSec VPN for VMware Cloud on AWS 3

IPSec Overview 3

IPSec VPN Architecture 4

General Design Considerations for IPSec VPN 5

IPSec Connectivity Options for VMware Cloud on AWS 6

- Policy-based IPSec VPN 6
- Route-based IPsec VPN 6

Use Cases 8

- Securing On-prem and VMware Cloud on AWS SDDC Connectivity Using Policy-based IPSec VPN When On-prem Gateway Does Not Support Route-based IPSec VPN 8
- Point-to-Point Connectivity Using Policy-based IPsec VPN 8
- Securing Connectivity Between On-Premises and VMware Cloud on AWS SDDC with Route-based IPSec VPN 8
- Implementing Direct Connect as a Primary VPN and Route-based IPsec VPN as a Backup 8
- Providing Device and Link Redundancy Using Route-based IPSec VPN 9
- Access to AWS VPCs and SDDC Networks with Secured Multipoint Route-based IPSec VPN Connections 10
- Hub and Spoke with On-prem Connectivity on Route-based IPSec VPN 10
- On-prem and VMware Cloud on AWS SDDC with Direct Connect in a Hub and Spoke IPSec 11

IPSec VPN for VMware Cloud on AWS

IPSec Overview

Internet Protocol security (IPSec) is a secure network protocol suite that authenticates and encrypts the data to provide secure and encrypted communication between two endpoints of a Virtual Private Network (VPN). IPSec uses cryptographic security services to protect communication over Internet Protocol (IP) networks. IPSec supports network-level peer authentication, data origin authentication, data integrity, data confidentiality (encryption), and replay protection. IPSec protects any application traffic over an IP network. Applications are automatically secured by IPSec at the IP layer.

IPSec VPN Architecture

As shown in Figure 1, the high level reference architecture shows how to secure the communication between on-premises and VMware Cloud on AWS SDDC running over an internet connection, using IPSec VPN.

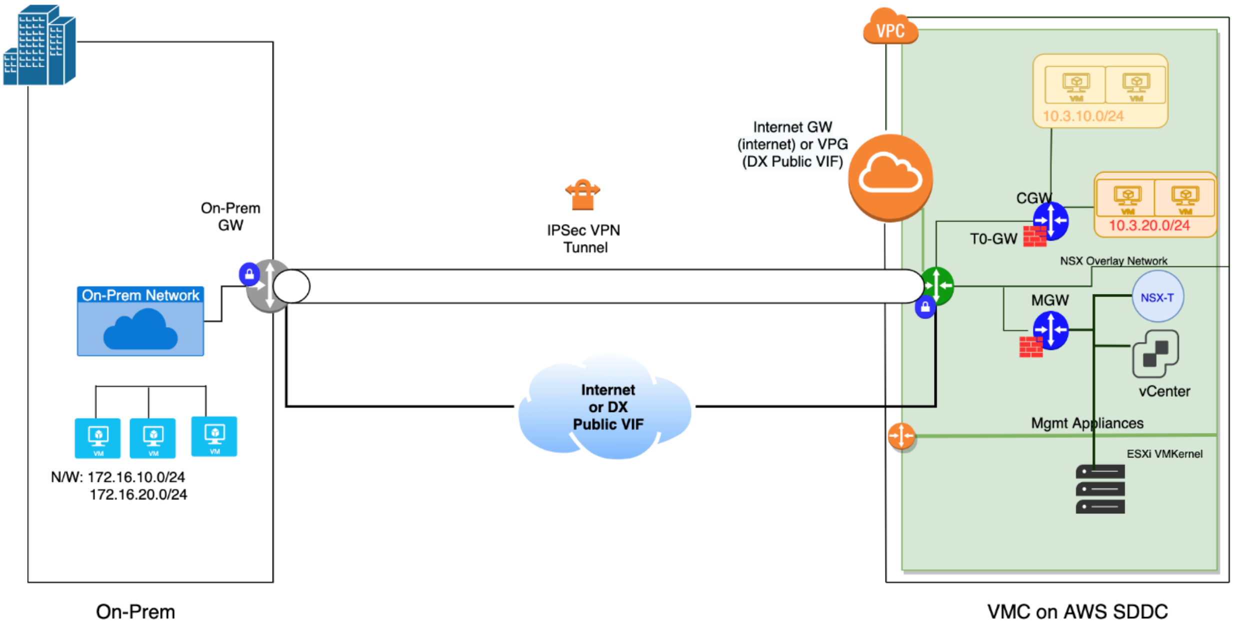

Figure 1 High Level IPSec VPN Reference Architecture

Traffic flows between the on-premises network and VMware Cloud on AWS using an IPSec VPN tunnel. The reference architecture for IPSec VPN consists of the following major components, regardless of the type of IPSec VPN being used.

- On-premises: A private local-area network running within an organization.
- IPSec VPN Connection: It consists of two VPN endpoints providing the VPN initiation and termination (the on-prem VPN GW and the VMware Cloud on AWS T0 Gateway on each sites).
- Network Connectivity: Public internet connection or Direct Connect Public VIF is required.
- On-premises Gateway: It is the initiator of the VPN connection using source as its own public IP and destination as the VMware Cloud on AWS public IP. It also provides connectivity to internet or external network.
- VMware Cloud on AWS T0-Router: It is the endpoint for IPSec tunnel over which the tunnel is terminated.
- VPC: This is the connected VPC to VMware Cloud on AWS SDDC.
- CGW (Compute Gateway): Provides connectivity to the workloads using logical segments of NSX-T.
- MGW (Management Gateway): It provides connectivity to VMware vCenter and NSX-T devices.

General Design Considerations for IPSec VPN

Before you decide to choose IPSec VPN for your on-prem to VMware Cloud on AWS SDDC implementation, note the following design considerations:

Device Selection

The architecture of IPSec VPN implementation depends heavily on the selection of devices and software to provide IPSec services within existing infrastructure. Proper selection of devices used to implement IPSec VPN is key for successful IPSec implementation. For instance, if you plan to use firewall itself as IPSec VPN gateway for your on-premises datacenter, then too much load on such integrated device might impact the functionality of one component over another (IPSec being computationally intensive may impact functionality of firewall in this particular case).

Device Placement

Due to the layered approach towards securing enterprise networks, IPSec gateway placement is a challenging task. Hence, careful placements of devices used to implement IPSec VPN is also very important. If your on-premises VPN gateway is behind a firewall, you must configure that firewall to forward IPsec protocol traffic.

Availability

Availability of VPN gateway is also an important point to consider while architecting implementation of IPSec VPN. You may want to protect the on-premises VPN gateway with a highly available instance or providing fault tolerance for VPN gateway. On the VMware Cloud side of SDDC, you may want to have availability across regions or zones based on your requirement and cost considerations.

Latency

Latency also plays an important role for communicating between on-premise and VMware Cloud on AWS SDDC as it depends on the internet connectivity and bandwidth offered by the service provider. For example, you may want to reconsider hosting an application on VMware Cloud on AWS SDDC and providing accessibility to it to the on-premise users if it is business critical as higher latency may result in slow response from the application. On the other hand, you may want to have disaster recovery for on-premises workloads to the cloud in case if disaster strikes to on-premises and this can be achieved using IPSec VPNs.

Tunnel Parameters Selection

Selecting the tunnel parameters like protocol, encryption algorithm, DH Group, hashing algorithm etc. is also an important key as both the endpoints may not support the recommended values, for e.g. VMware recommends the protocol value of IKE as v2, however, the remote VPN endpoint may only support IKE v1.

IPSec Connectivity Options for VMware Cloud on AWS

The choice of connectivity between the SDDCs depends on various factors such as location of the datacenter, application requirement, CAPEX/OPEX etc. After considering these factors, you may want to opt for a Direct Connect Private VIF connection (DX) or Internet-based connectivity.

DX is a dedicated connectivity that runs from your SDDC to AWS DCs. DX has low latency and high bandwidth as compared to the Internet-based connectivity (economical than DX). Many customers today use DX as a primary and internet as a secondary connectivity option.

With internet connectivity option, there is always a risk of "Man In The Middle" attack and IPSec VPN is one of the options to avoid such attacks. Typically DX is used for accessibility of business critical applications across the sites whereas VPN based connectivity over internet is suitable for use cases such as disaster recovery or non-business critical application across the sites where latency is not a requirement.

VMware Cloud on AWS offers two options for IPSec VPNs - policy-based IPSec VPN and route-based IPSec VPN.

Policy-based IPSec VPN

Policy-based VPNs help by encrypting and encapsulating traffic flowing through an interface according to a defined policy. Policy is applied to a specific interface and only interested traffic that crosses this interface gets evaluated by the policy. In context of SDDC, policy is applied to the uplink of NSX-T Tier0 Router. As part of configuration, the SDDC is configured with a default route which points to its upstream Internet Gateway (IGW). This means that all traffic which is non-local to the SDDC will be sent out through the uplink and will be matched against any IPSec VPN policies which are applied. If the traffic matches a policy, then it is encapsulated and sent through a VPN. Otherwise, it is routed normally to the IGW.

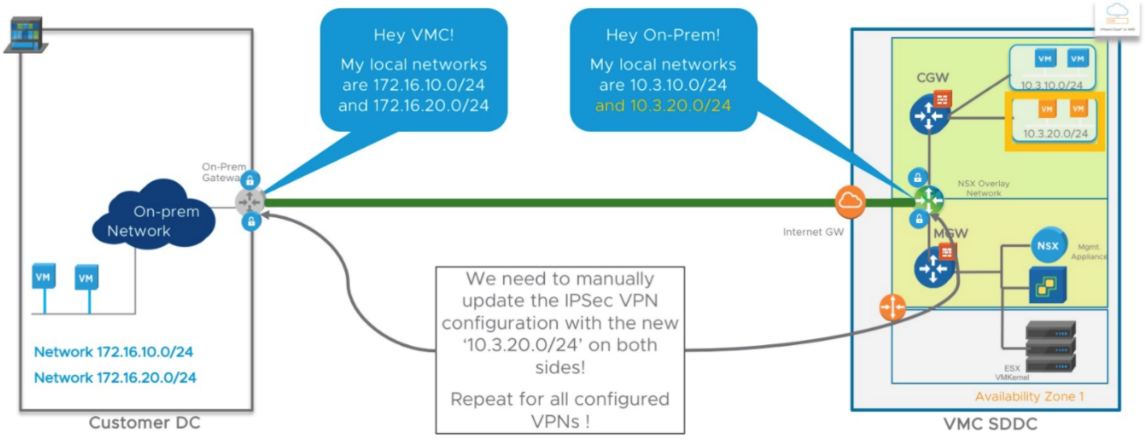
Hey VMC!
My local networks are 172.16.10.0/24 and 172.16.20.0/24
Hey On-Prem!
My local networks are 10.3.10.0/24 and 10.3.20.0/24
On-Prem Gatewa
On-prem Network
VM
VM
Network 172.16.10.0/24
Network 172.16.20.0/24
Customer DC
Internet GW
We need to manually update the IPSec VPN configuration with the new '10.3.20.0/24' on both sides!
Repeat for all configured VPNs !
CGW
VM
VM
10.3.10.0/24
VM
VM
10.3.20.0/24
NSX Overlay Network
MGW
NSX
Mgmt. Appliance
ESX VMKernel
Availability Zone 1
VMC SDDC

Figure 2 - Policy-based IPsec VPN

Route-based IPsec VPN

A route-based VPN creates a virtual interface and traffic hitting this interface is encrypted and decrypted according to the phase 1 and phase 2 IPSec settings. Compared to policy-based IPSec VPNs, route-based IPSec VPNs need extra configuration as they support dynamic routing over tunnel interface. With a route-based VPN, a dynamic routing protocol (such as BGP) is configured within the tunnel and routes are exchanged dynamically. Once the route-based VPN tunnel is established, there is simply no need to update it, even after you create additional network segments in the Cloud or on remote sites.

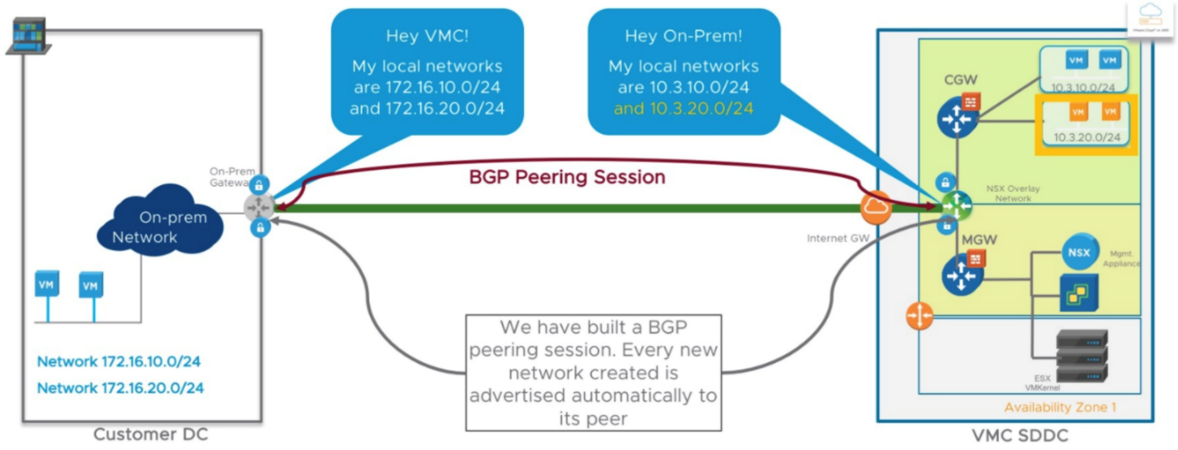

Figure 3 - Route-based IPsec VPN

Use Cases

Securing On-prem and VMware Cloud on AWS SDDC Connectivity Using Policy-based IPSec VPN When On-prem Gateway Does Not Support Route-based IPSec VPN

There may be a situation where the remote device (on-premises gateway in this case) does not support route-based IPSec VPN. The connectivity with VMware Cloud on AWS SDDC can still be established with policy-based IPSec VPN.

Design Considerations & Caveats:

- Policy-based IPSec VPN does not support dynamic routing and it forms Security Associations (SAs) in response to "interesting" traffic matching policy (eventually tears down the SAs in the absence of such traffic).
- You must plan carefully what routes you advertise to the VMware Cloud on AWS SDDC.

Point-to-Point Connectivity Using Policy-based IPsec VPN

Policy-based IPSec VPN is one of the solutions for a point-to-point connectivity (example remote office/branch office) that does not require connectivity to other sites.

Design Considerations & Caveats:

- Point to multipoint or hub and spoke kind of topologies are not supported with policy-based IPSec VPN.
- You must ensure bandwidth and latency requirements are met (not too high latency and low bandwidth) especially in remote office/branch office topologies.
- If any routes are learned via alternative uplinks (i.e., Direct Connect, cross-linked VPC, other route-based VPNs) which encompass networks configured on the policy-based VPN, then the traffic will not pass through the internet uplink and will not hit the policy. In such situation, you must carefully plan what routes you advertise to the SDDC if you intend to use policy-based VPN in conjunction with Direct Connect or route-based VPN.

Securing Connectivity Between On-Premises and VMware Cloud on AWS SDDC with Route-based IPSec VPN

Implementing Direct Connect as a Primary VPN and Route-based IPSec VPN as a Backup

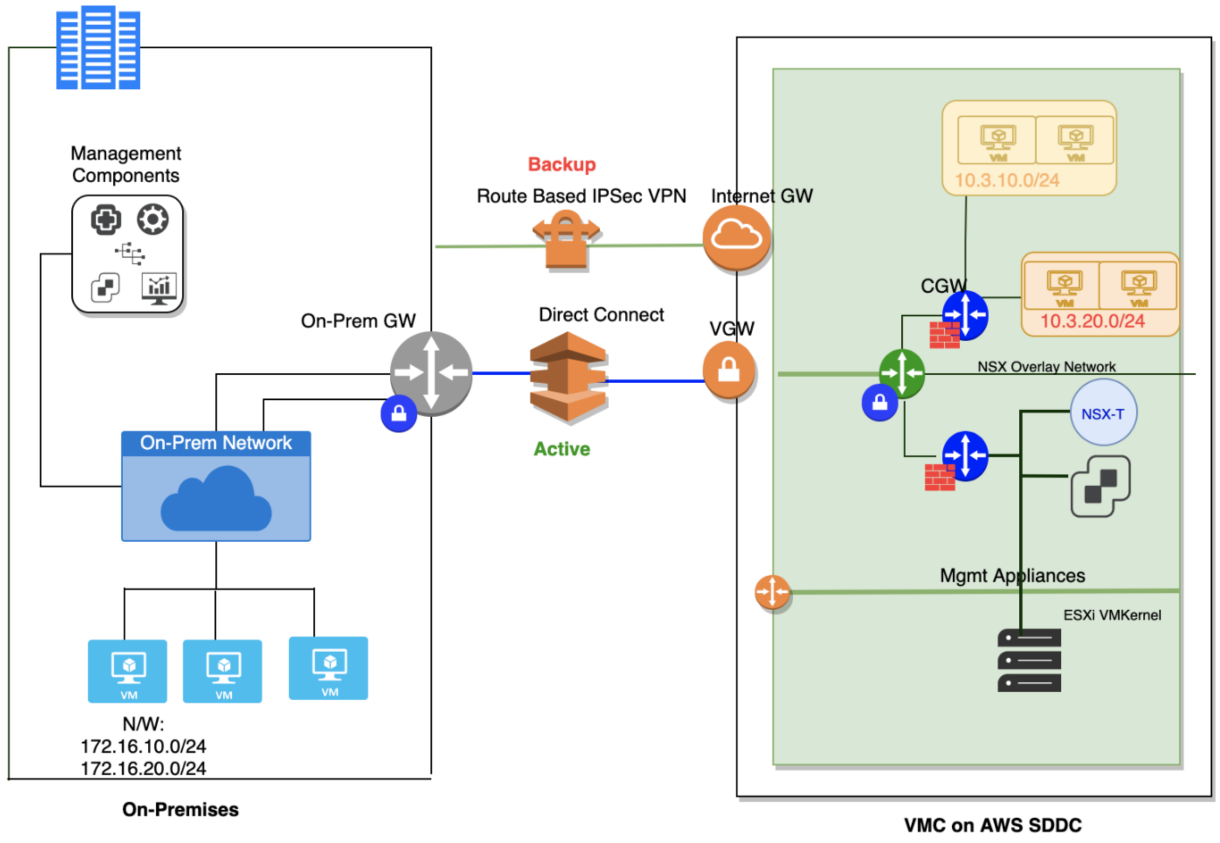

Figure 4 - DX and VPN Connectivity

This use-case is suited for customers trying to save cost by using internet with route-based IPSec VPN running as backup to Direct Connect.

Design Considerations & Caveats

- ESXi management and vMotion traffic will always go over DX, hence this use-case provides backup for management appliance traffic (E.g: vCenter, NSX Managers, etc.) and compute/workload traffic.
- By default, in VMware Cloud, when the same network is learnt over both DX private VIF and route-based IPSec VPN, routes learnt over IPSec VPN has better administrative distance than that of DX private VIF. As a result of this, routes learnt over IPSec VPN are preferred over DX private VIF.
- When the option of route-based IPSec VPN as backup is enabled, traffic egressing the VMware Cloud on AWS SDDC will prefer the DX private VIF path over route-based IPSec VPN. However, due to asymmetric routing, return path from on-prem may not follow the same path. Hence, for traffic egressing the on-prem SDDC towards the VMware Cloud on AWS SDDC, you must prefer the routes learnt via DX over routes learnt via Route-based IPSec VPN Connection.
- You may want to advertise the more specific routes over DX and less specific routes with route-based IPSec VPN connection as backup.
- If you have 2 DX private VIFs, there will be two different DX circuits one of which can be used as backup. In such scenario, route-based IPSec VPN as backup is not very useful.

Providing Device and Link Redundancy Using Route-based IPSec VPN

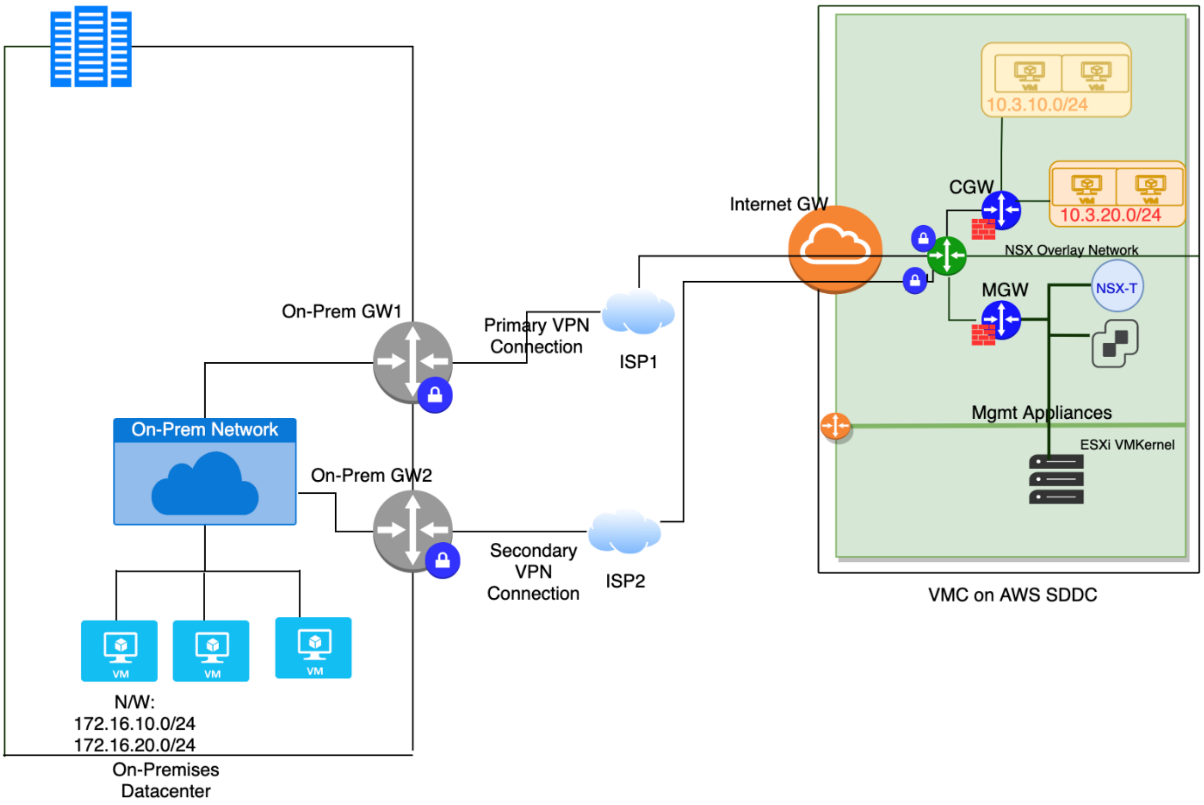

Figure 7 – Link and device redundancy using route-based IPsec VPN

With tunnel resiliency, multiple route based IPSec VPN Tunnels can be built providing better resiliency between the two VPN connections. This feature is useful when an on-prem site has multiple connectivity options with different ISPs and on-prem VPN initiators. As shown in Figure 7, the on-prem sites has two redundant VPN connections initiated using two different VPN gateways with two different ISP links. These connections are terminated on a VMware Cloud SDDC with multiple VPN connection option configured from the UI.

Design Considerations & Caveats

- VMware Cloud on AWS supports multiple VPN tunnels, provided the remote tunnels IPs are different.
- Traffic will not be load balanced in this use case. This use case is mainly focused towards redundant VPN initiators using different ISP links.
- Manual intervention may be required to switchover between the tunnels.
- Switchover time depends on IKED to detect tunnel down.
- For multiple tunnels, careful subnet for the BGP local IP and Remote IP is desired keeping in mind the reserved IP ranges. Click on the "i" icon in BGP remote IP configuration to know more about the reserved IP ranges.
- This use-case can provide link and device redundancy for Route based IPSec VPN.

Access to AWS VPCs and SDDC Networks with Secured Multipoint Route-based IPSec VPN Connections

The difference between the following two deployments is that the deployment indicated in Figure 8 is with all spoke connections over route-based IPSec VPN whereas the deployment indicated in Figure 9 is with one of the on-premises to VMware Cloud on AWS SDDC connectivity with Direct Connect. Please refer to the design considerations section for both the topologies.

Hub and Spoke with On-prem Connectivity on Route-based IPSec VPN

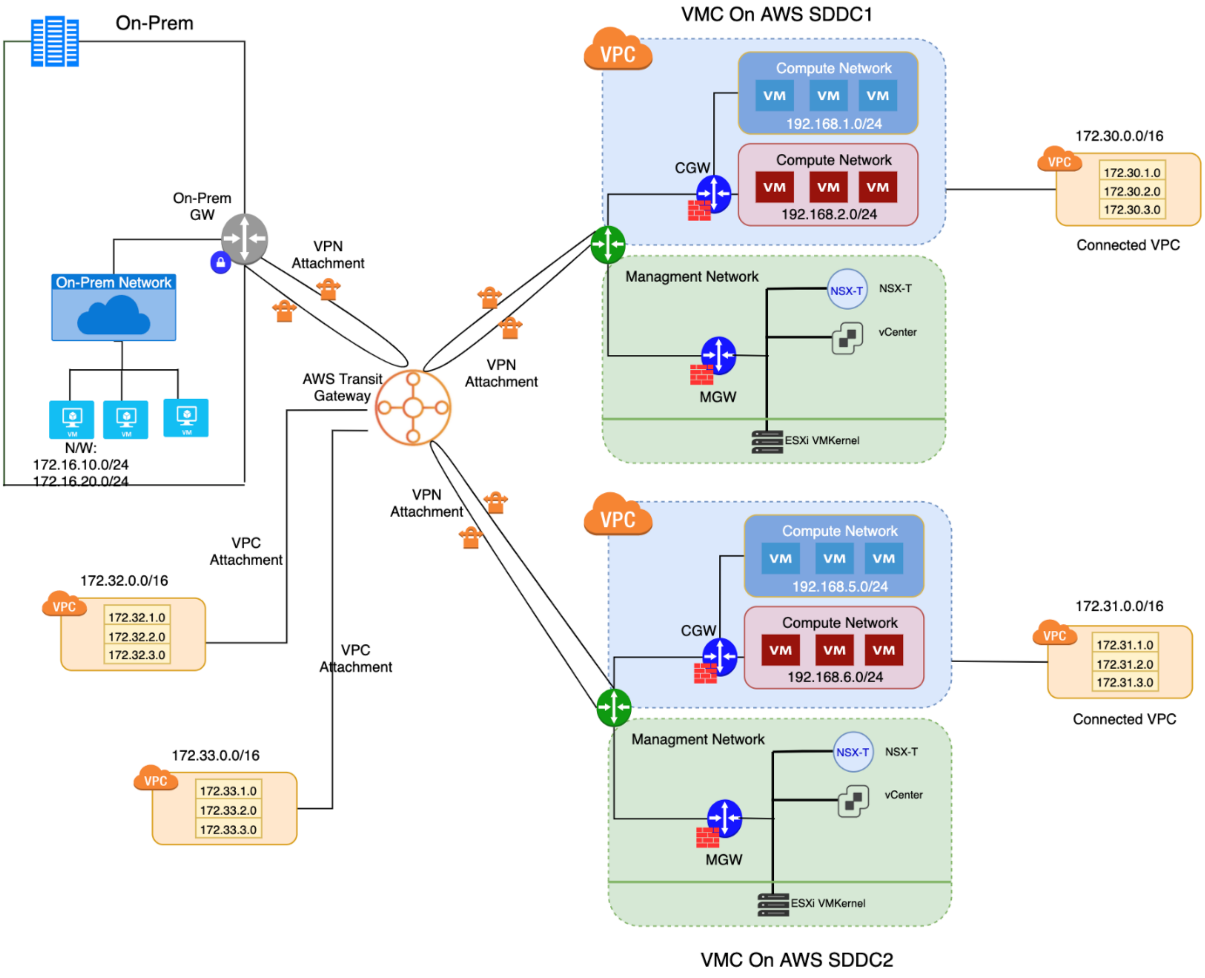

Figure 8 - Hub and Spoke route-based IPSec VPN

This use-case can provide any-to-any connectivity between the on-prem, VMware Cloud SDDCs and native AWS VPCs using AWS Transit Gateway (TGW). Route based IPSec VPNs can be leveraged on top of this to secure multipoint connections between VMware Cloud and AWS and and on-prem.

This use-case not only provides communication between all the sites and native AWS VPCs but can also provide an option to migrate workloads from on-prem to any of the VMware Cloud SDDCs and vice-versa. With 2 VPN tunnels per VPN attachment, this use-case can also provide resiliency for each of the spoke connections from TGW to on-prem or VMware Cloud on AWS SDDC.

Design Considerations and Caveats

- AWS Transit Gateway (TGW) is required to achieve hub and spoke connectivity.
- TGW would incur cost and has to be carefully considered while deploying this use-case.
- TGW has one route table, however you can control the route tables for each of the transit gateway attachments on the TGW.
- Since, this is any-to-any communication, it is important to make use of Compute Gateway firewall rules on VMware Cloud Console and Security Groups in AWS console to restrict or allow traffic between the spoke connections.
- For multiple tunnels, careful subnet for the BGP local IP and Remote IP is desired keeping in mind the reserved IP ranges. Click on the "i" icon in BGP remote IP configuration to know more about the reserved IP ranges.
- Using this topology, the expectation is that default route is learned from TGW tand all traffic routes through i. The caveat is that all internet bound traffic will flow in this direction, which is something you need to account for.

On-prem and VMware Cloud on AWS SDDC with Direct Connect in a Hub and Spoke IPSec

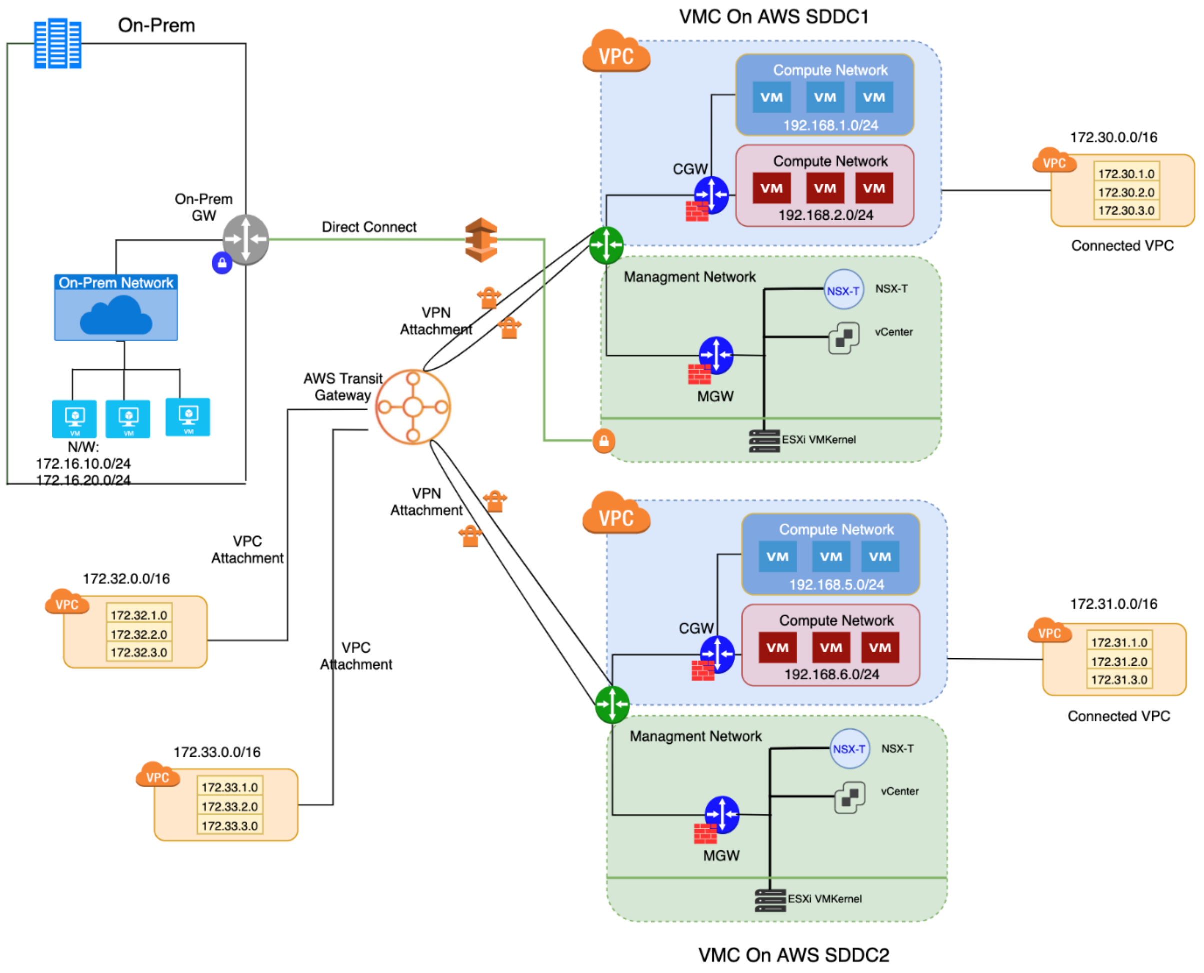

Figure 9 Hub and Spoke route-based IPSec VPN with Direct Connect

As shown in Figure 9, this use-case has Direct Connect (DX) VIF setup between on-prem VMware Cloud on AWS SDDC1. Following are additional design considerations for this use-case.

Design Considerations

- Direct connect as shown in Figure 8, exists between on-prem and VMware Cloud on AWS SDDC1. Since, the other spoke connections are not using direct connect, due to latency constraints, the communication between on-prem to SDDC2 might be slower as compared to communication between on-prem and SDDC1.
- As mentioned under DX and VPN connectivity section, careful consideration for asymmetric routing has to be accounted for while implementing this use-case.
- Using this topology, a design consideration is where you may choose to inject default route. This would again impact internet traffic. It is probably also likely that you have some connectivity between on-prem and the TGW, which introduces the possibility of asymmetrical routes.