Designlet: VMware Cloud on AWS SDDC Connectivity With IPSec VPN

VMware Architecture

Table of contents

Designlet: VMware Cloud on AWS SDDC Connectivity With IPSec VPN

Introduction

This document provides you with recommendations and guidelines on how to connect your VMware Cloud on AWS SDDC using IPSec VPN.

There are several prerequisites and considerations that you must be aware of before you start configuring the network.

Refer to the table below for a summary of the use case, considerations, and other details to see if IPSec VPN meets your requirements.

Summary and Considerations

Use Case	When a customer requires connectivity to an SDDC and does not have an AWS Direct Connect (DX) in the desired region, but has reliable Internet. Performance requirements should be no greater than 5-6 Gbps peak total in both directions, with some tolerance for latency.
Pre-requisites	
General Considerations	
Performance Considerations	
Cost implications	Internet egress from the SDDC is billed at $0.05/GB. No additional charge for the VPN connections themselves.
Documentation reference	https://docs.vmware.com/en/VMware-Cloud-on-AWS/services/com.vmware.vmc-aws-networking-security/GUID-92F6C09E-8E74-430E-8F79-C2E5B2150ADA.html
Last Updated	Oct. 2022

Planning and Implementation

There are many best practices and recommendations to keep in mind when you are planning to connect your VMware Cloud on AWS SDDC using IPSec VPN.

Planning

- Use Route-based VPN rather than policy-based VPN.
- Use a GCM-based Cipher such as AES GCM 128, with no Digest Algorithm (GCM ciphers include Checksum) for both Tunnel and IKE.
- Enable Perfect Forward Security to mitigate against Replay Attacks.
- Enable Diffie Hellman Group 14 or higher. This provides a minimum 2048-bit modulus for key exchange.
- Leave the Remote Private IP blank, as long as the on-prem VPN endpoint is using the IP Address specified as the “Remote Public IP” for the local address used for IKE Authentication.
- Use IKEv2. IKEv1 is considered a legacy protocol, and IKEv2 provides several improvements and extensions.
- Use a /30 subnet in the 169.254.32.0-169.254.100.255 range for the BGP IP. Use the first host IP for the BGP Local IP, and the second one for the BGP Remote IP. Reverse the IPs between local and remote on the on-prem device.
- Ensure your on-prem ASN is different from the ASN used on the SDDC, thus using eBGP. The ASN is a global setting on the SDDC, used for all VPN connections. Note that it is unrelated to the ASN used for an AWS Direct Connect (DX), and they both run separate BGP instances which are not connected.
- If multiple VPNs are configured to different endpoints using different BGP ASNs, eBGP dictates that prefixes will be re-advertised from one VPN to the other, which can cause the SDDC to become a transit hub for traffic. This behavior can be influence by the use of the "no-export" BGP community, which will be honored by the SDDC if set by the remote endpoint. Note that BGP prefixes will not be re-advertised between VPN and DX connections regardless of their ASNs, since they do not share the same BGP instance in the SDDC and are not BGP peers.
- Use the exact same IPSec configuration parameters on both sides of the IPSec tunnel. If settings are mismatched, tunnel creation will fail.
- Ensure the ISAKMP/IKE SA lifetime is set to 86400 seconds, Phase 2 SA lifetime is 3600 seconds, and data-based lifetime is disabled as is any idle timeout.
- When using IKEv1, ensure the following settings:
    - Use ISAKMP Main Mode (not aggressive mode)
    - IPsec mode is set to Tunnel
    - Phase 2 mode is set to ESP.
- Use TCP MSS Clamping only when required. In most cases guest OS should use Path MTU Discovery (PMTUD) to automatically determine the optimal packet size that can pass through the VPN without being fragmented. ICMP fragmentation required (Type 3, code 4) messages must be permitted end-to-end between the workloads. TCP MSS Clamping can help when it is not possible to allow ICMP between endpoints, or if there is a guest OS that does not support PMTUD. If needed, first configure the VPN without TCP MSS Clamping enabled to check the maximum path size by using a ping with the “don’t fragment” (DF) bit set. Test subsequent pings with increased packet sizes until pings fail. Set the MSS clamping value to the largest packet that responds with the DF bit set.
- If a policy-based VPN is required, all of the recommendations above still apply, with the addition that you must ensure that the exact same networks are selected on each side, with remote and local swapped. Use as few networks as possible, as the number of tunnels that must be created is equal to the number of local and remote networks multiplied together. You must plan to keep this number below 100.
- Policy-based VPNs cannot be used with another connection when the routes that are to go over the tunnel are also advertised over the other connection. This includes advertising the default route. This is because policy-based VPN is not routed using traditional methods, so routing rules such as more-specific-route-wins do not apply. The traffic intended to go over the policy VPN must be directed to the interface the policy VPN is running on (Internet Interface for public IP, DX for private IP) in order to match the policy and be sent via the policy-based VPN.

Implementation

VPNs are configured using the VMware Cloud console, under an SDDC.

The following screenshot shows the home page of the VMware Cloud/SDDC console:

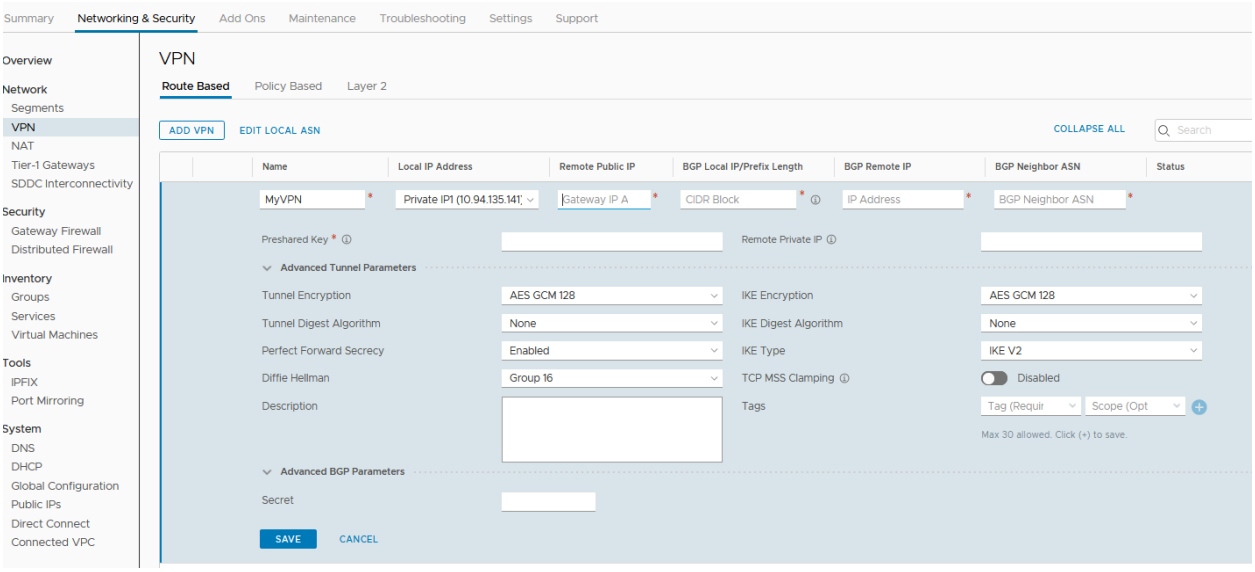

Go to the **Networking & Security** page, and **VPN** is found under the **Network** section on the left-hand menu. There are separate tabs for Route-based, Policy-based, and Layer 2 VPNs.

You can also find the public IP used for the VPN endpoint on the **Overview** page of the **Network & Security** tab, at the very top of the central diagram.

To create a new VPN connection, follow the steps below:

1. Click the **Add VPN** box in the top left and complete the settings to align with your on-prem endpoint.
2. Ensure the Preshared key is identical on both sides. If you are using BGP authentication, enter the secret and ensure that it matches on both sides.