Designlet: VMware Cloud on AWS SDDC Connectivity With Direct Connect Private VIF

VMware Architecture

Table of contents

Designlet: VMware Cloud on AWS SDDC Connectivity With Direct Connect Private VIF 3
Introduction 3
Summary and Considerations 3
Planning and Implementation 4
Planning 4
Implementation 5

Designlet: VMware Cloud on AWS SDDC Connectivity With Direct Connect Private VIF

Introduction

This document provides you with recommendations and guidelines on how to connect your VMware Cloud on AWS SDDC to a corporate network using Direct Connect (VIF).

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

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

Summary and Considerations

Use Case	For private network connectivity high bandwidth and low-latency. Customer will need to order the Direct Connect (DX) from their preferred AWS Connect partner to an AWS POP in the same region as their SDDC.
Pre-requisites	
General Considerations	
Performance Considerations	DX is generally the best performing connectivity available between a customer's network and AWS. When a customer procures multiple DXs to the same region, they should order them all to the same AWS account. AWS will attempt to provision them on separate infrastructure to provide high availability/diversity.
Cost implications	
Documentation reference	https://docs.vmware.com/en/VMware-Cloud-on-AWS/services/com.vmware.vmc-aws.networking-security/GUID-417EE2F1-0EA5-4808-BDB3-6FF622EECB95.html
Last Updated	December 2022

Planning and Implementation

There are many best practices and recommendations to keep in mind when you are planning to use AWS Direct Connect (DX) Private VIF.

Planning

- If multiple DXs are used, ensure they are provisioned in the same customer AWS account so that AWS is aware they are for redundancy. AWS will attempt to configure them on separate equipment, if available, to provide high availability.
- All the private VIFs will be attached to the same VGW for the SDDC, so any supported AWS private VIF configuration and BGP attributes will work. Traffic can be balanced across multiple circuits in a multi-active scenario, or different traffic can be routed over different circuits. The SDDC does not provide visibility into the routes on the individual DX VIFs, but only sees the resulting aggregate routes learned by the VGW from any DX VIF.
- The management CIDR is always advertised as 3 different component subnets, while with the default behavior, the first 16 compute segments will be advertised based on their configured CIDR. This limit can be increased on request, pending available capacity. Any requests to increase the limit should include a business justification and consider 100 routes the upper limit.
- Route aggregation is available to enable summarization of multiple segments into a single advertised route. Any segments whose IP range lies inside a defined aggregation will not be advertised, but the defined aggregation CIDR will always be advertised. Segments whose IP ranges do not fall inside a defined aggregation will continue to be advertised individually. The total number of aggregations + non-aggregated segments is still subject to the limit of 16 advertised CIDRs. This feature requires SDDC version 1.18 or higher.
- Route filtering is available to suppress advertisement of all IP ranges from segments attached to the compute gateway (CGW). Defined aggregations will be the only advertised CIDRs (in addition to the management CIDRs). This feature requires SDDC version 1.20 or higher.
- Use AS-prepending or BGP community tags to influence the path used for traffic going from the SDDC to the customer network.
- The default ASN assigned to the SDDC is 64512. If this conflicts with an on-prem ASN, you should change it prior to attaching the VIFs. It is recommended to use the same ASN for all on-prem connections (iBGP) and avoid route re-advertisement. All VIFs connected to the same SDDC will use the same ASN.
- It is possible to use a route-based VPN as a backup to a DX by enabling the option on the DX configuration in the SDDC. When enabled this changes routes learned over DX to be preferred over those learned from VPN. This only has an effect for equivalent networks advertised on both DX and VPN. In all cases more specific routes will be preferred, so advertisements should be managed from the on-prem side to ensure symmetrical paths. When the option is disabled (default), equal network routes are preferred over VPN. It is not recommended to use VPN as a backup for a DX where route aggregation has been enabled as the VPN does not support aggregations and will always advertise the more specific routes, causing it to always be the preferred path for those networks.
- The BGP failover default timers are not configurable and use a keepalive timer of 30 seconds, with a hold timer of 90 seconds. This means that if a connection stops responding, BGP will wait 90 seconds before removing routes. Since BGP will negotiate to the lower timer values between peers, on-prem can be configured with as low as a 1 second keepalive timer and 3 second hold timer. Bidirectional Forwarding Detection (BFD) can also be enabled with liveness detection of 300ms and a multiplier of 3 to reduce failure detection time to sub 1 second. Note BFD will negotiate to the slowest values between peers.
- VMC SDDCs interact with the AWS Virtual Private Gateway (VGW) using API and create static routes in the SDDC. This API is polled every 2 minutes, which provides a worst-case convergence time of 6.5 minutes using the default BGP timers without BFD, or 2 minutes with BFD enabled.
- Since the VGW abstracts the multiple DX VIFs from the SDDC, failover of routes between different DX VIFs does not depend upon the API polling. Failover from an active route is at most 4.5 minutes (BGP only, default timers) or ~1s with BFD in use.
- Jumbo frames up to 8900 bytes are supported over DX private VIF. It is recommended to enable them on the DX, the VIF and on the SDDC even if all workload traffic is expected to only use 1500 byte MTU. This will make it easier to take advantage of the higher throughput of jumbo frames for specific traffic in the future. Using jumbo frames can provide significant performance improvements in VMC.

Implementation

Private VIFs are created on the customer's AWS console under the account and region where the DX is provisioned. When creating the VIF, you must specify to create it "**in another AWS account**", and enter the AWS Account ID shown in the **Direct Connect** page for your SDDC:

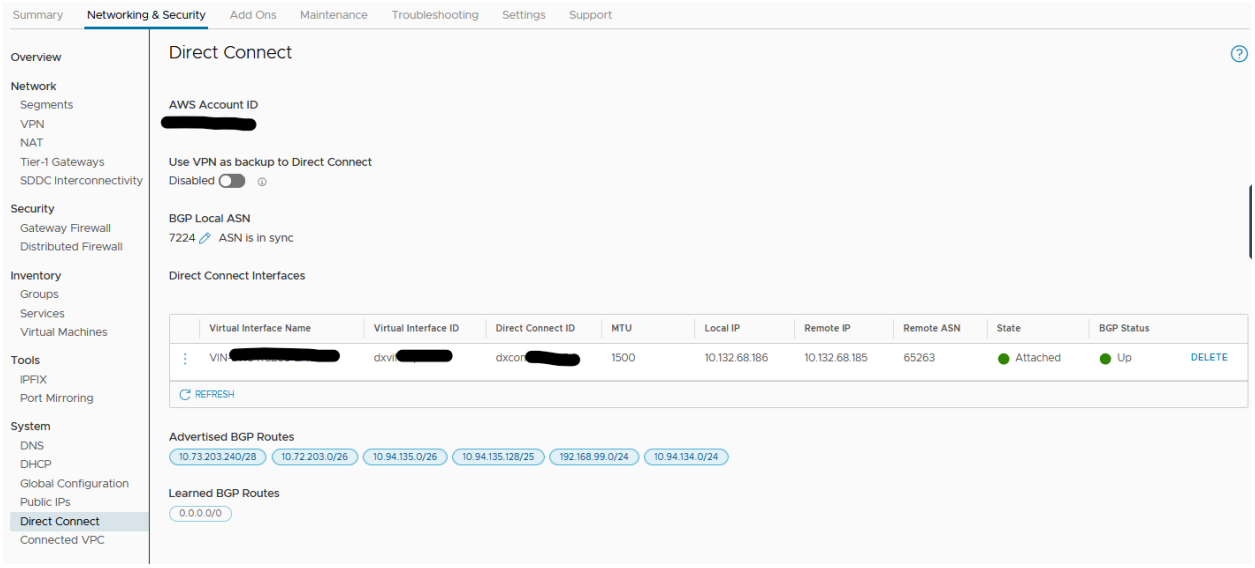

Once the VIF is created, it will appear in the "Direct Connect Interfaces" list in the SDDC.

Note: Until the VIF is attached, it will be visible to all SDDCs in the org, as they all share a common AWS account ID. Once the VIF is attached, it will no longer appear for other SDDCs.

The MTU for the DX is configured on the **Global Configuration** tab. It will prevent configuring a higher MTU than the lowest MTU of any VIF attached the SDDC.