Designlet: VMware Cloud on AWS Connected VPC to Native AWS

VMware Integrations

Table of contents

Designlet: VMware Cloud on AWS Connected VPC to Native AWS

Introduction

This document provides you with recommendations and guidelines on how to use VMware Cloud on AWS SDDC connectivity to native AWS with the Connected VPC.

There are several prerequisites and considerations that you must be aware of before you start using this connectivity.

All 2-node and higher SDDCs must connect to an AWS account and VPC at time of creation. This connection can be optionally deferred up to 2 weeks for 1-node SDDCs. The purpose of this connection is to allow the use of native AWS services, such as S3, RDS, Elastic Load Balancer, etc. directly from VMs running in your SDDC.

Refer to the table below for a summary of the use case, considerations, and other details to see what is possible with the Connected VPC.

Summary and Considerations

Use Case	Access native AWS services and resources directly from VMs running in a VMware Cloud on AWS SDDC through a high-bandwidth, low-latency connection that provide no-cost same-AZ data transfer. (Cross-AZ connectivity within the VPC is also supported, but subject to data transfer charges).
Pre-requisites	
General Considerations	S3 connectivity through the Connected VPC requires an S3 Endpoint to be deployed in the Connected VPC and configured on the main route table. Alternatively, S3 access can be toggled on or off on a per-SDDC basis. Access to S3 buckets in remote regions, or when using a non region-specific URL (for SDDCs that are not located in US-East-1) follows the SDDC's routing for its connection regardless of the S3 Service setting of the SDDC. This is the case for all S3 access when the S3 Service is disabled. The Connected VPC path is only available for management appliances or workload VMs running in the SDDC and connected to a routed network segment. VMs on Layer-2 extended networks (including those with HCX MON enabled) are not able to use the Connected VPC, as their default gateway resides in the on-premises network, and so they follow that gateway's routing tables (in the case of HCX MON-enabled networks, it is due to the lack of a return path from the Connected VPC). In addition, it is not possible to route traffic through the SDDC to the Connected VPC. If any on-prem or remote devices need to reach the Connected VPC, they must make use of native AWS connectivity (such as TGW, AWS VPN or Direct Connect) and bypass the SDDC.
Performance Considerations	
Cost Implications	There is no charge for traffic passing through the Connected VPC when the destination is in the same AZ as the SDDC. When the destination is in a different AZ, traffic is subject to cross-AZ charges of $0.01US per GB in and out (equivalent of $0.02US per GB). Any charges for the services being accessed (e.g. S3) still apply.
Documentation Reference	Accessing AWS Services on docs.vmware.com Using ENI on https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html
Last Updated	December 2021

Planning and Implementation

There are a few key requirements and best practices to keep in mind when you are planning to access native AWS services from your VMware Cloud on AWS SDDC using the Connected VPC.

At the time of deployment, 17 Elastic Network Interfaces (ENIs)are created. Each ENI is assigned an IP address from the selected subnet, and labelled with a Description of "VMware VMC Interface DO NOT USE - *_SDDC ID_##*, where SDDC_ID is the ID of the SDDC, and ## is a number from 0 to 16. (See the documentation reference above for information about ENIs). An ENI is attached to each ESXi host in the SDDC's management cluster (Cluster-1) and has a status of "In-Use". They are not necessarily assigned in numerical order. One of the In-Use ENIs will have a secondary IP address assigned to it. This is the host running the active default Edge, and this secondary IP is used for SNAT of traffic to S3. Any additional ENIs are marked "Available" and are reserved for future host adds or replacements. This also means that 18 IP addresses from the selected subnet are consumed initially.

Netwo...	Subnet ID	VPC ID	Availability Zo...	Security gr...	Description	Status	Primary priva...	Secondary private IPv4
eni-0b3b...	subnet-0dc...	vpc-014...	us-west-2a	default	VMware VMC Interface DO NOT USE - _81f...	Available	192.168.254.105	–
eni-0411...	subnet-0dc...	vpc-014...	us-west-2a	default	VMware VMC Interface DO NOT USE - _81f...	Available	192.168.254.93	–
eni-04bee...	subnet-0dc...	vpc-014...	us-west-2a	default	VMware VMC Interface DO NOT USE - _81f...	Available	192.168.254.89	–
eni-03d5...	subnet-0dc...	vpc-014...	us-west-2a	default	VMware VMC Interface DO NOT USE - _81f...	Available	192.168.254.113	–
eni-0b06...	subnet-0dc...	vpc-014...	us-west-2a	default	VMware VMC Interface DO NOT USE - _81f...	Available	192.168.254.94	–
eni-00fec...	subnet-0dc...	vpc-014...	us-west-2a	default	VMware VMC Interface DO NOT USE - _81f...	Available	192.168.254.126	–
eni-0b9ef...	subnet-0dc...	vpc-014...	us-west-2a	default	VMware VMC Interface DO NOT USE - _81f...	Available	192.168.254.83	–
eni-0a145...	subnet-0dc...	vpc-014...	us-west-2a	default	VMware VMC Interface DO NOT USE - _81f...	Available	192.168.254.115	–
eni-0365...	subnet-0dc...	vpc-014...	us-west-2a	default	VMware VMC Interface DO NOT USE - _81f...	Available	192.168.254.76	–
eni-0c0ad...	subnet-0dc...	vpc-014...	us-west-2a	default	VMware VMC Interface DO NOT USE - _81f...	Available	192.168.254.102	–
eni-0554...	subnet-0dc...	vpc-014...	us-west-2a	default	VMware VMC Interface DO NOT USE - _81f...	Available	192.168.254.124	–
eni-013d...	subnet-0dc...	vpc-014...	us-west-2a	default	VMware VMC Interface DO NOT USE - _81f...	Available	192.168.254.69	–
eni-0ad42...	subnet-0dc...	vpc-014...	us-west-2a	default	VMware VMC Interface DO NOT USE - _81f...	Available	192.168.254.80	–
eni-0997...	subnet-0dc...	vpc-014...	us-west-2a	default	VPC Endpoint Interface vpce-0f2951803a...	In-use	192.168.254.92	–
eni-03b5...	subnet-0dc...	vpc-014...	us-west-2a	default	VMware VMC Interface DO NOT USE - _81f...	In-use	192.168.254.121	192.168.254.114

If the management cluster grows beyond 17 hosts (for example due to unaddressed storage growth, VMware may add hosts to prevent a disk-full situation), additional ENIs are created, and IP addresses assigned from the subnet. For this reason it's best practice to ensure that there are at least 33 IPs available in the subnet for the SDDC, and at least 33 available ENIs in the region (See the AWS VPC quotas page for more details. The current default limit is 5,000 ENIs per region in an AWS account).

In addition, the main route table must have sufficient capacity for all the SDDC's routes - one for the management CIDR, plus one for every routed segment in the SDDC. Since the default route table limit is 100 routes, you may need to request an increase from AWS if you require routes for your native AWS traffic and/or have a large number of routed segments in the SDDC.

Planning

- The primary CIDR of the VPC is added to the SDDC's route table, making the Connected VPC the preferred path for traffic destined for addresses in that CIDR.
- Any secondary CIDRs added to the VPC are ignored by the SDDC, and traffic to that CIDR follows the SDDC's route table to reach the destination. If this network is reachable through on-prem or another path, it's important that the return path follows that same path. Subnets in secondary CIDRs should not use the main route table if there is communication required to the SDDC, as they cannot use the ENI path to the Connected VPC from the SDDC so they should not use the ENI as a return path.
- Resources in the Connected VPC must be in subnets that are associated with the main route table. This is because the SDDC only updates the main route table with the return routes for segments in the SDDC. These routes are updated whenever segments are added or removed, but also need to point to the active Edge instance's ENI, which can change in case of an Edge HA failover or vMotion, and occurs during SDDC Maintenance activities.
- Multi-Edge Traffic Groups can be used to reach resources within the Connected VPC to scale North-South traffic throughput by directing traffic from specific IP addresses defined in a Traffic Group to a separate Edge VM. Traffic to S3 over the Connected VPC always uses the default Edge, as it is be Source NATted to the Edge's IP in the Connected VPC's subnet.
- The ENIs created in the Connected VPC are always assigned to the default VPC Security Group. To preserve security separation between accounts, the SDDC does not have permissions to manage Security Groups in the customer's AWS account, nor to assign alternative a custom Security Group to the ENI.
- To connect to a service that is not running in your VPC, you need to make it reachable from an IP in your VPC using an Interface Endpoint (also called PrivateLink), or using a load balancer that is connected to the a subnet in your VPC. The

exception to this is S3, which requires a Gateway Endpoint to be associated with the main route table of the VPC. VMC has special support for S3, which can be enabled or disabled using the toggle on the Networking & Security -> Connected VPC under Service Access. The default is for S3 to be enabled, which sends all traffic to local S3 buckets over the Connected VPC. If an S3 Endpoint has not been created in the VPC and associated with the main route table, then S3 buckets in the SDDC's region will not be accessible.

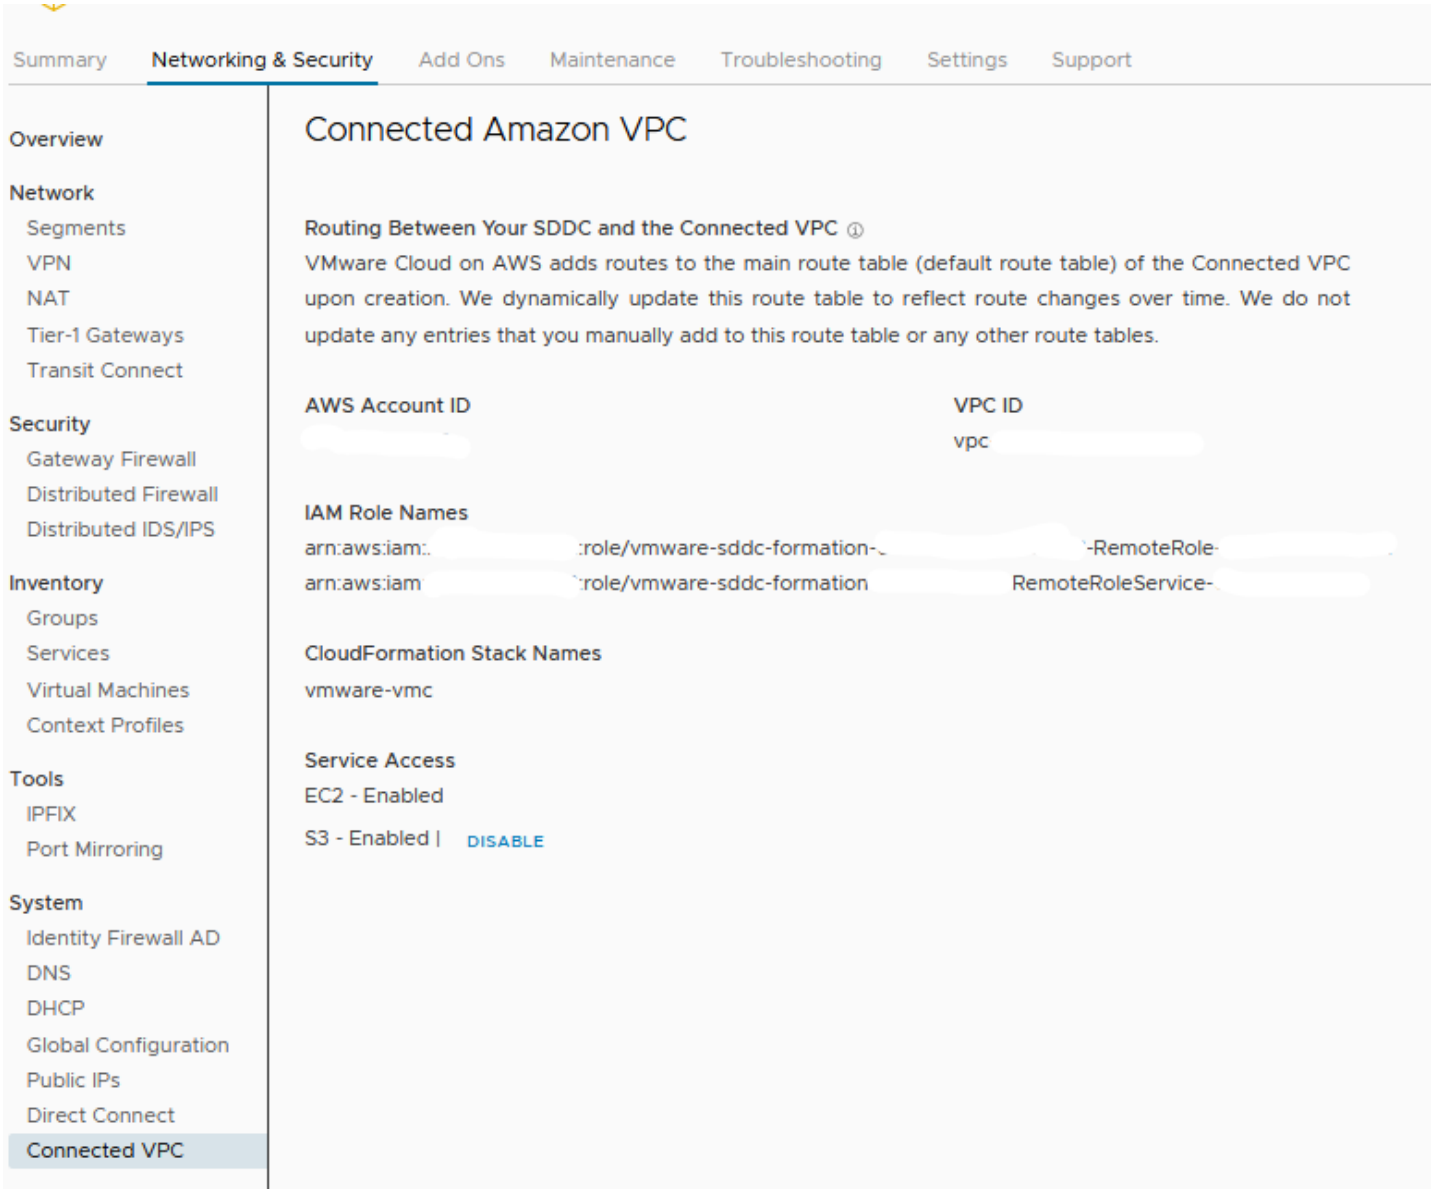

The diagram below is an example of the Connected VPC's main route table, where we can see 3 routes pointing to an ENI, for each of the 2 networks in the SDDC and the management CIDR, as well as an S3 endpoint (vpce-#). Note the destination for the S3 endpoint is a prefix-list (pl), which is a list of IP CIDRs, in this case one that is managed by AWS, and refers to the S3 networks for the Oregon region.

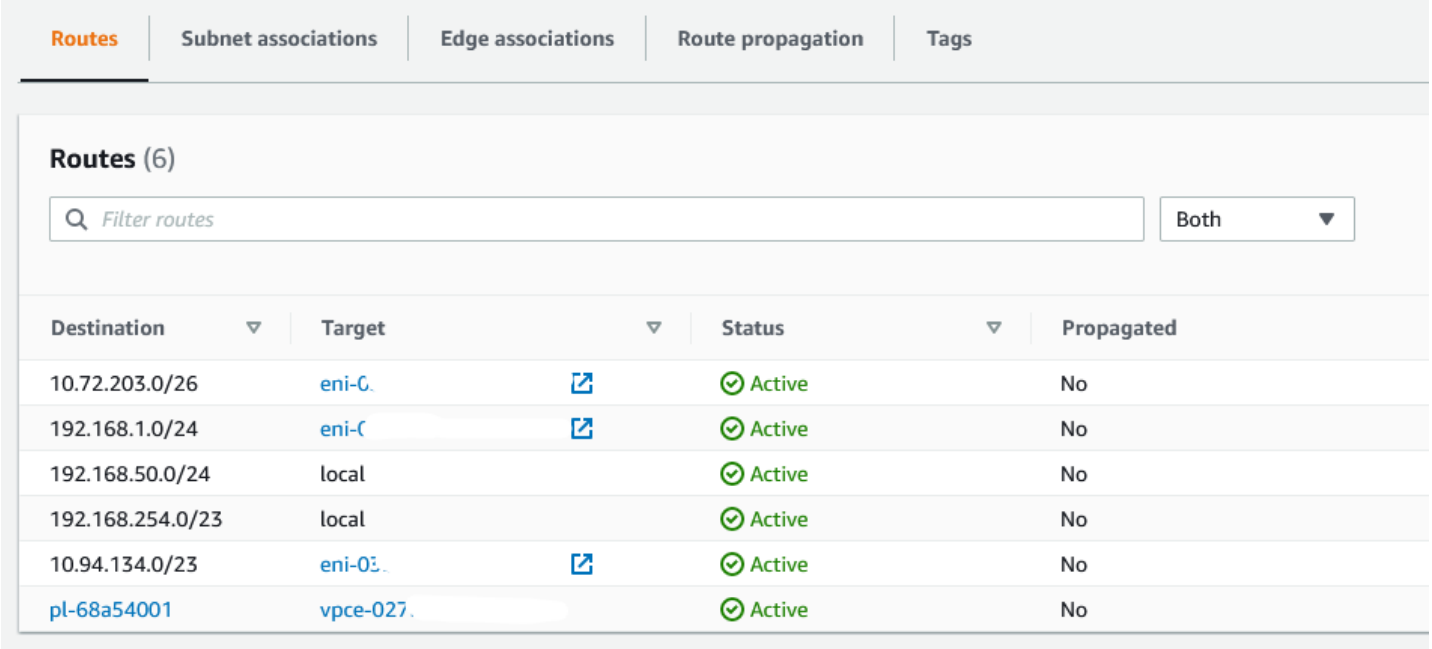

Implementation

When an SDDC is connected to a customer AWS account (which is required for any production SDDC when it is deployed), the customer must select a VPC and a subnet. The AZ of the subnet determines which AZ the SDDC is deployed in. Always refer to the AZ of the subnet as shown in your AWS account, or use the AZ ID rather than AZ name, as AWS rotates the AZ name for each AWS account, but the AZ ID is consistent for all customers.

When configuring network access, the path from a VM in the SDDC to an EC2 instance in the AWS Connected VPC goes through the following network controls:

1. Distributed Firewall (DFW) for the VM's NIC

2. Compute Gateway Firewall, using the VPC Interface (also called the Services Uplink)

3. VPC default Security Group (or if customer has manually assigned a different Security Group to the ENI). Note that because the Security Group is assigned to the ENI, the directions are relative to the SDDC, not the VPC. e.g. Out = Traffic leaving the SDDC towards the AWS VPC, In = Traffic leaving the AWS VPC towards the SDDC. By default, Security Groups allow all traffic to go outbound, and no traffic inbound. Therefore if you have connections being initiated from the AWS VPC towards a VM in the SDDC, you must modify the default Security Group to allow that traffic. Note Security Groups are stateful, so traffic must be defined in the direction it originates, but allow bi-directional traffic flow between those endpoints once the flow is established.

4. The AWS network ACL on the VPC. By default, the ACL allows all traffic. Note ACLs are not stateful, so traffic must be allowed in both directions for return traffic flows.

5. The Security Group attached to the EC2 instance. The rules in this Security Group are applied from the perspective of the EC2 instance, where Out = Traffic leaving the EC2 instance, and In = Traffic coming into the EC2 instance. If the default Security Group is applied to the EC2 instance, then required traffic must be allowed both In and Out (one direction to allow traffic to the SDDC, and the other direction to allow traffic to the EC2 instance.)

6. The firewall on the guest OS - if your guest OS has a firewall running such as iptables or Windows Advanced Firewall, it needs to be configured to allow the desired traffic.

When performing a traceroute to verify connectivity, first ensure that all of the security controls listed above allow the traceroute traffic, which could be UDP or ICMP, depending upon the traceroute command used by your guest OS. From a VM in the SDDC, after the default gateway, you first see a 169.254.x.1 IP, which represents the transfer between the T1 and T0 routers in the SDDC. Then, there is an IP from the management CIDR. For traffic going over the Connected VPC link it this address has the format A.B.C+1.81 for a management CIDR of /23 (e.g. for a management CIDR of 10.2.0.0/23, it is 10.2.1.81), A.B.C+10.129 for a CIDR of size /20, and A.B.168.1 for a CIDR of size /16. If the traffic is being sent over another path, you see a different IP from the management CIDR in that hop.

This path can also be verified and visualized using the NSX Traceflow utility, which is available from the NSX Manager UI (requires SDDC 1.16 or later):

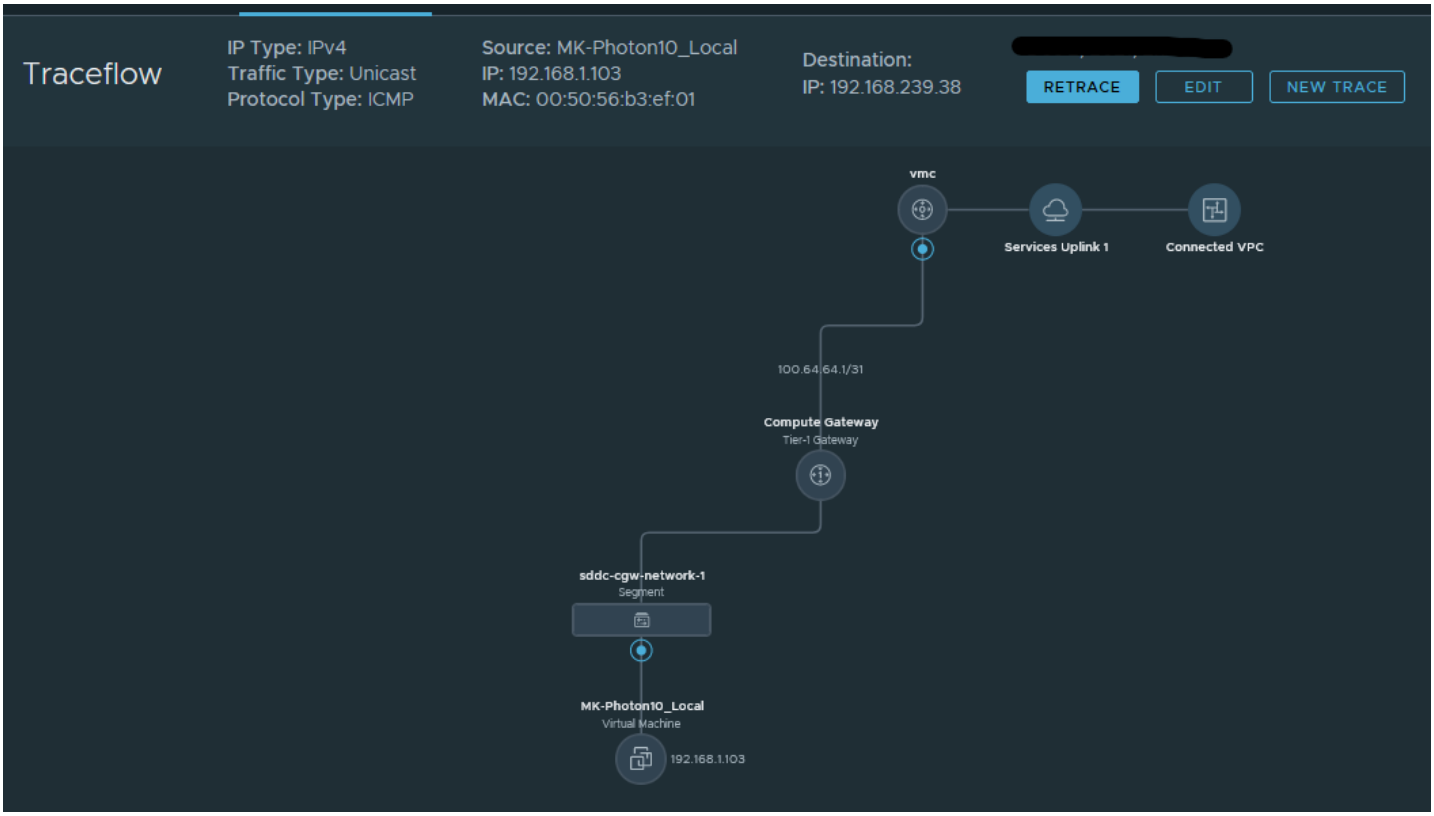

Once an SDDC is connected to a VPC, there is currently no user-serviceable method to change that VPC or subnet. However, you can contact support to request VMC engineering assistance with reconnecting the SDDC to a new AWS account/VPC/subnet.

Author and Contributors

Michael Kolos, Product Solutions Architect, VMware