VMware Cloud on AWS: A Technical Overview

VMware Architecture

Table of contents

VMware Cloud on AWS: A Technical Overview

Introduction

Overview

VMware Cloud on AWS (VMC on AWS) is a managed cloud offering that provides dedicated VMware vSphere-based Software Defined Data Centers (SDDC) that are hosted within AWS facilities.

Prior to getting started with the service, it is important to understand a few points.

- The service utilizes facilities and hardware that are owned and managed by AWS.
- The service provides dedicated, private cloud environments in the form of an SDDC.
- Hosts of an SDDC are dedicated to that SDDC. These hosts are exclusive to the SDDC until such a time that they are removed from the SDDC (either manually removed by the customer or replaced due to host failure), at which time they are released into the overall pool of AWS available capacity.
- SDDCs are deployed within a VMware-owned AWS account, not a customer-owned AWS account.
- SDDCs have high-speed access to native AWS services hosted within a separate customer-owned AWS account.
- Native AWS services are billed to the customer-owned AWS account and are not managed by VMware.

The VMware Cloud Organization (Org)

The VMware Cloud Organization (Org) may be thought of as a top-level construct which owns 1 or more cloud services. Within VMware Cloud on AWS, an Org will be the top-level container for SDDCs.

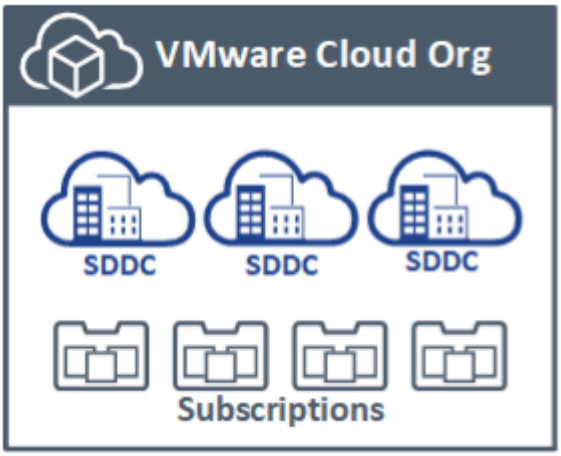

In addition to containing SDDCs, an Org will also contain any host subscriptions which might have been created. These subscriptions are fixed on the following attributes:

- AWS Region
- VMware Cloud Org
- Host instance type

Amazon Web Services Account

A major benefit of VMC on AWS is its ability to provide high-speed, direct access to AWS services. That said, it is required that all customers maintain a dedicated AWS account which will be used to access and manage these services. If you are unsure of how to create an AWS account, then please refer to the AWS documentation for more information on the process.

A few important points on the AWS account:

- The account is necessary in order to provide an SDDC with access to AWS services.
- It is required in order to deploy a production SDDC.
- The account is owned by the customer, not by VMware.
- Billing for the account is handled directly by AWS, not by VMware

AWS Regional Availability

Please refer to the roadmap for a list of AWS regions which offer VMC on AWS.

Software Defined Data Center (SDDC)

The Software Defined Data Center (SDDC) is a collection of bare-metal hosts which are installed with a standard set of VMware software. It is important to understand that each SDDC is running atop dedicated hardware, and that billing for an SDDC is based on the number of hosts dedicated to the SDDC. It is **not** based on the number of VMs running within the SDDC.

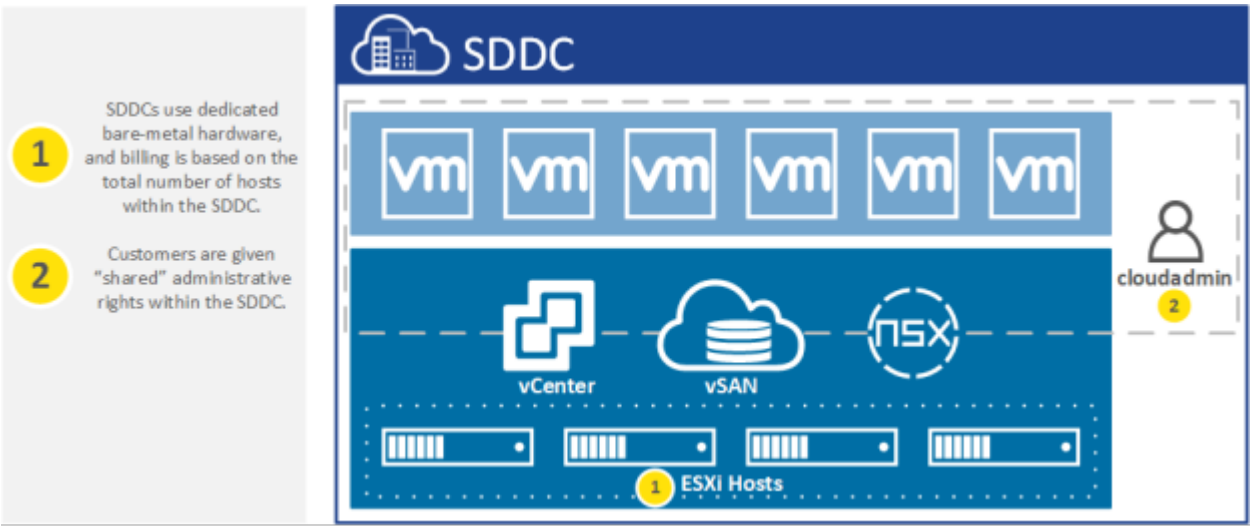

Since VMware Cloud on AWS is a managed service, full admin-level access to the SDDC is not permitted. This restriction is in place to prevent customers from modifying the infrastructure of the SDDC. Instead, customers are given a role that allows them to fully manage workloads which they have deployed within the SDDC. Normally, this permissions model does not impact day-to-day use of the service; however, it's important to keep in mind if you are planning on integrating tools directly with infrastructure components such as vCenter. If the integration you are planning requires admin rights, then it may not function properly. It is recommended to review the documentation for your specific application in order to understand its permission and access requirements.

Key Takeaways of an SDDC:

- An SDDC is deployed on dedicated, bare-metal hosts.
- An SDDC is Deployed with standard components (ESXi, vCenter, NSX, and vSAN).
- Billing is based on the number of hosts within the SDDC, not on the number of VMs.
- Users have the ability to manage their workloads but have limited access to vCenter, vSAN, and NSX.

Integration with AWS Services

Each SDDC is provided direct access to AWS services via a connection to a customer-owned AWS account. This connection is established as part of the SDDC provisioning process and is performed using one of two methods:

1. By choosing an AWS account which has already been connected as part of a previous SDDC deployment, or
2. By creating a new connection to an AWS account.

The term “connected” simply means that the customer has granted permissions for the VMware to enable routing between an SDDC and a VPC within the customer-owned AWS account. These permissions are granted via IAM roles which are created within the connected account using a CloudFormation template. It is important to note that the person who is performing the account connection process must have sufficient permissions (e.g. admin rights) within the AWS account in order to execute this CloudFormation template.

Once a connection is established to the AWS account, it then becomes possible to configure a cross-link between an SDDC and a VPC within that account. The cross-link itself is made up of a series of cross-account ENIs which are attached to a subnet within the VPC. It is these ENIs that provide the hosts of an SDDC with a network forwarding path to resources within the VPC.

The cross-linking process is automated at the time of SDDC deployment. No manual work is required to perform this configuration.

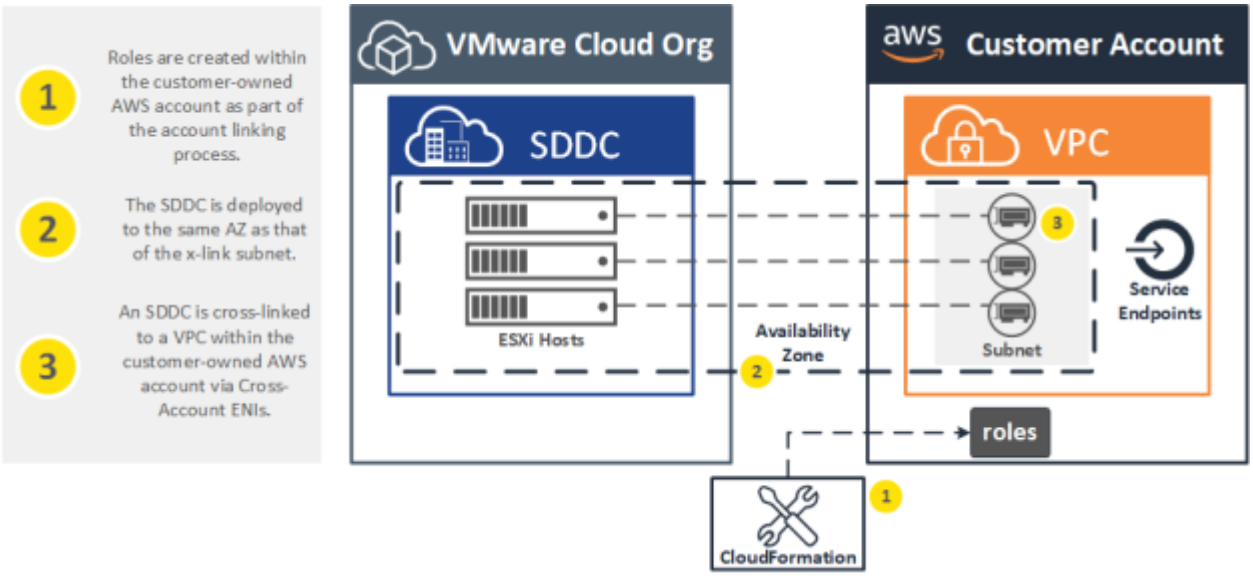

Key Takeaways

- AWS services are managed through a customer-owned AWS account.
- Account linking enables VMware to cross-link SDDCs into a customer-owned AWS account.
- Account linking is performed when the user executes a CloudFormation template within their AWS account.
- The CloudFormation template creates roles which enable VMware to manage SDDC cross-linking.
- The hosts within the SDDC are connected to the customer-owned VPC via cross-account ENIs.

Networking

SDDC Network Architecture

VMware utilizes NSX-T to build a logical overlay network on top of the hosts of the SDDC. There are 2 tiers of routing within the SDDC. At the top tier is the tier-0 edge router, which acts as the north-south gateway for the entire SDDC. Below that are the tier-1 routers (the MGW and CGW), which act as the north-south gateways for their respective networks.

There are 2 layers of firewalls in the SDDC. The first layer is provided by the gateway firewalls, which are designed to protect the north-south border of the SDDC. The gateway firewalls have a "default deny" policy and are implemented at the MGW (for the management network) and at the tier-0 edge (for the compute network). The second layer of firewalls are provided by the distributed firewall. The distributed firewall is enforced at the vNIC level of every VM within the compute network and is designed to enable filtering both north-south and east-west. The distributed firewall has a "default permit" policy which effectively disables the distributed firewall unless the SDDC administrator specifically creates "deny" rules.

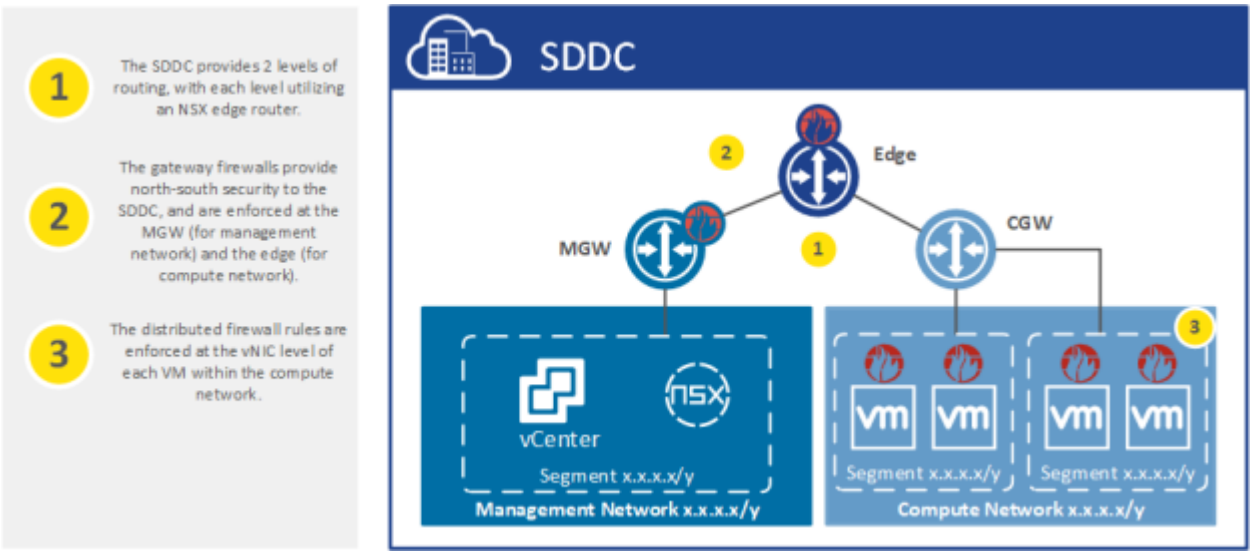

External Connectivity

In the majority of setups, customers wish to maintain some permanent means of direct connectivity between the SDDC and their on-premises environment. IPSec VPN is a common means of accomplishing this. IPSec VPN provides secure connectivity to the private IP address ranges of the SDDC and is implemented with a tunnel to the edge router.

For customers who want a high-speed private connection into their SDDC, VMC on AWS supports the attachment of Direct Connect. As with all AWS services, Direct Connect will be provisioned within the customer-owned AWS account. It is important to note that SDDCs may utilize ***existing*** Direct Connect services. There is no need to provision a dedicated Direct Connect for the SDDC. Utilizing a Direct Connect is the simple matter of provisioning a new Private VIF and then allocated it to the VMware-owned AWS account which is associated to the parent Org of the SDDC. This account is visible from the Direct Connect interface of the SDDC within the VMC console.

In cases where neither IPSec VPN nor Direct Connect are available, the vCenter server appliance within the SDDC may be accessed directly via its public IP address. Per the default firewall policy of the SDDC, vCenter will not be accessible unless connectivity is explicitly permitted.

IP Administration

All SDDCs will be cross-linked to a VPC within the customer-owned AWS account but may also be connected to other networks (such as an on-premises environment). In order to ensure that the SDDC can communicate with other interconnect networks, it is vital that IP addressing be properly planned. IP ranges should be unique and non-overlapping between the SDDC and any networks to which it will be connected. As such, one of the most critical pieces of the design process is proper planning of IP address usage.

Though it is not required, it is a good practice to allocate IP address space in large contiguous chunks. The following table provides a sample IP Administration plan.

Supernet	Subnet level 1	Subnet2 level 2	Description
10.1.0.0/19			on-premises networks
10.1.32.0/19			AWS native
10.1.32.0/19	10.1.32.0/22		AWS Services VPC
10.1.32.0/19	10.1.32.0/22	10.1.32.0/26	AWS Services VPC SDDC x-link
10.1.64.0/19			SDDC1
10.1.64.0/19	10.1.64.0/20		SDDC1 Management
10.1.64.0/19	10.1.80.0/20		SDDC1 Compute
10.1.64.0/19	10.1.80.0/20	10.1.80.0/24	SDDC1 Compute Servers

Workload On-Boarding

Customers will generally populate an SDDC with workloads in one of two ways:

1. Deploying the SDDC as a greenfield environment
2. Migrating existing workloads into the SDDC from another environment

A greenfield deployment operates on the notion that the SDDC is a completely new environment. In this scenario, all workloads will be created from scratch and the SDDC will utilize freshly allocated IP address ranges.

Workload migration scenarios are the most common means of populating an SDDC. In order to address data center evacuation scenarios, VMware has developed a service called HCX which enables transparent workload migration between sites. HCX has been specifically designed to address complex migration scenarios and includes features such as:

- Migration scheduling - enables migrations to be scheduled off-hours.
- WAN optimization and data de-duplication - greatly reduces the time and network bandwidth required to perform migrations.
- Layer-2 network extension - enables workloads to migrate without requiring IP address changes.