VMware Cloud on AWS

Practitioners Quick Reference for Data Center Extension

WHAT IS VMWARE CLOUD ON AWS?

VMware Cloud on AWS is a unified SDDC platform that integrates vSphere, vSAN, and NSX on top of bare metal hardware from AWS. Utilizing the same trusted, enterprise-grade VMware products and tools, customers can tap into the broad range of AWS services while taking advantage of the functionality, elasticity, and security they expect to gain in the cloud.

Organizations often need to extend their data center to the cloud for a variety of reasons. Today's environment demands that new projects and geographical locations can be rapidly supported, seasonal or unplanned capacity needs met on demand, and hybrid applications deployed to meet evolving needs. At the same time, as-needed environments for testing, development, labs and training also require data centers to be capable of expanding in response to business demands. While extending the data center can help to meet the needs of dynamic, evolving organizations, it can also be complex and time-consuming.

VMware Cloud™ on AWS provides an ideal solution for extending data center to public cloud by offering a production-ready, enterprise-grade vSphere®-based environment on the AWS Cloud that is also delivered by VMware as a cloud service. VMware Cloud on AWS doesn't require re-tooling or re-educating IT teams because the solution delivers consistent vSphere-based infrastructure and operations that are already implemented in existing on-premises data centers. In addition, there is no need to purchase or install hardware or infrastructure software because all lifecycle management is performed by VMware.

Key Considerations when Extending On-Premises Data Center Environments to the Cloud

In order to extend on-premises environments to VMware Cloud on AWS, IT teams need to understand their connectivity options, hybrid operations and resource management.

vmware | aws

INTERCONNECTIVITY OPTIONS: AT A GLANCE

- AWS Direct Connect (DX)
- Route Based IPSEC VPN
- L2VPN
- VMware HCX
- AWS Transit Gateway

Key Consideration 1: Interconnectivity

One of the key aspects to consider while extending data center environments to VMware Cloud on AWS is understanding the various options available for establishing network connectivity between the two environments, depending on each organization's requirements and use-cases. The selection of connectivity options should be based on considerations such as:

- The need for live migration of workloads to and from the cloud; vs. moving power-off VMs
- The ability to migrate workloads without changing their IP addresses
- A requirement for layer 2 connectivity between workloads on-premises and in the cloud

IT teams can leverage AWS Direct Connect for high-bandwidth, low latency private connectivity or use route-based IPSec VPNs. In addition, they can also extend layer 2 networks with L2VPN capabilities. With VMware HCX®, IT teams can also leverage multi-site hybrid interconnect to enable operations between environments on different on-premises vSphere versions and network types. Finally, for any-to-any connectivity between VMware Cloud on AWS Software Defined Data Center (SDDC) and AWS Virtual Private Clouds (VPC), IT teams can leverage AWS Transit Gateways (TGW).

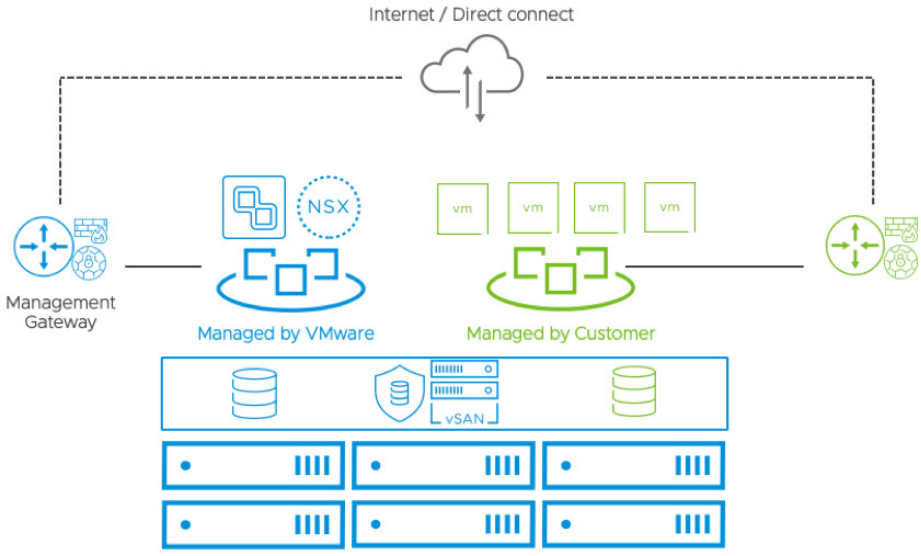

AWS Direct Connect (DX)

- Provides high bandwidth, low latency private connectivity
- Carries all traffic across this link: vMotion®, ESXi™ management, appliance management, workload traffic, etc.
- Offers the ability to learn all on-premises networks and expose SDDC network segments to the on-premises data center
- Takes advantage of VMware NSX® features such as Distributed Firewalling and Micro-Segmentation
- Configuration of a VPN connection over DX allows IT teams to take advantage of live vMotion to move workloads back and forth between the on-premises data center and VMware Cloud on AWS

Route Based IPSEC VPN

- Configure a single VPN tunnel design between on-premises data center and Cloud SDDC
- Configure a virtual tunnel interfaces (VTIs) in the SDDC and run BGP over IPSec VPN to automatically learn networks
- Configure IPSec with private addressing over DX and public addressing over VPN
- View advertised and learned routes as well as traffic statistics

L2VPN

- Extends virtual machine networks and workload traffic
- Extends up to 100 on-premises networks to VMware Cloud on AWS
- Ideal for disaster recovery, cloud bursting, and migration from on-premises to the cloud without the need to re-IP

VMware HCX

- Provides high performance, multi-site interconnectivity capabilities by abstracting infrastructure and allowing interoperation across different versions of vSphere and different network types
- Provides the ability to setup a WAN optimized multi-site IPSec VPN mesh for secure site-to-site connectivity
- Stretches Layer 2 networks and extends data centers between sites
- Performs bulk workload migrations with live bi-directional vMotion and the ability to retain MAC and IP addresses

AWS Transit Gateway

- Eases the creation of any-to-any communications between native AWS VPCs and VMware Cloud on AWS SDDCs
- Connectivity to on-premises can be shared easily with all connecting VMware Cloud on AWS SDDCs and native AWS VPCs

HYBRID CLOUD OPERATIONS: AT A GLANCE

- Hybrid Linked Mode with vCenter Cloud Gateway
- Content Library
- Support for VMware and 3rd party technology solutions for operations and management

Key Consideration 2: Hybrid Cloud Operations

Operational consistency helps simplify the management of hybrid cloud environments. VMware Cloud on AWS offers this consistency because the infrastructure and the operations associated with day-to-day management are both familiar and leverage existing VMware and 3rd party technologies used in on-premises data center environments. The most important capabilities that are available to hybrid clouds built with VMware Cloud on AWS are:

- A single-pane of glass for virtual machine management for on-premises and cloud
- Unified and consistent access to workload software and files
- The ability to use the same operations tools for monitoring, planning, and assessing their workloads

The following VMware Cloud on AWS capabilities are key to ensuring the ease of hybrid cloud operations:

Hybrid Linked Mode

Hybrid Linked Mode is a flexible solution that allows the joint management of both VMware Cloud on AWS and on-premises single sign-on (SSO) domains. Hybrid Linked Mode retains the separation between on-premises and VMware Cloud on AWS permissions to avoid issues if the two environments need to be broken apart. Once Hybrid Linked Mode is established, on-premises workloads can be migrated to VMware Cloud on AWS. Migration works both ways and workloads can be migrated back from VMware Cloud on AWS to on-premises.

Hybrid Linked Mode provides two options:

- A one-way trust from on-premises to VMware Cloud on AWS (e.g. VMware Cloud on AWS trusts the on-premises users) with the option to link and unlink as needed.
- An on-premises appliance called the vCenter® Cloud Gateway brings hybrid connectivity on-premises. This keeps both on-premises and in-cloud vCenters separate and allows single-pane hybridity through an appliance in the data center.

Content Library

A Content Library can be created in the VMware Cloud on AWS SDDC by subscribing to a Content Library that exists in an on-premises data center. Items published in the on-premises library can then be imported into the VMware Cloud on AWS SDDC library. This method works for transferring templates, ISO images, scripts, and other files.

Existing VMware and 3rd party validated technologies

IT teams can continue to use popular VMware and 3rd party validated technologies to manage VMware Cloud on AWS environments in an integrated way. One popular tool is VMware vRealize® Operations™.

vRealize Operations enables IT teams to most efficiently plan their capacity and migration strategy when adopting VMware Cloud on AWS. IT teams can proactively identify upcoming capacity shortfalls and help model future requirements. In addition, they can reclaim capacity and right-size environments to gain the maximum utilization across the footprint. Finally, for workloads that are identified to move to VMware Cloud on AWS, they can discover application dependencies prior to moving the workloads.

1. Plan

- › Determine potential costs and efficiency gains using VMware Cloud on AWS
- › Compare costs and capacity requirements for different deployment scenarios

2. Assess

- › Discover application dependencies
- › Confirm the readiness (e.g., health and compliance) of the destination
- › Model scenarios to ensure available capacity
- › Confirm the move is successful once workloads have been migrated

ELASTICITY AND DYNAMIC RESOURCE MANAGMENT: AT A GLANCE

- Elastic DRS
- Elastic vSAN
- Compute Policies

Key Consideration 3: Elasticity and Dynamic Resource Management

VMware Cloud on AWS allows IT teams to perform on-demand scaling of capacity and policy-driven resource management. This enables workloads to start running in the cloud with minimal upfront cost, and with the confidence that resources will be allocated as needed when workload compute or storage needs grow.

Elastic DRS

Elastic DRS for VMware Cloud on AWS (eDRS) enables automatic scaling of a cluster in response to demand, or lack of demand, by adding or removing hosts automatically based on specific configured policies. The eDRS algorithm runs at short intervals and looks at predefined resource thresholds for CPU, memory, and storage. The thresholds cannot be changed by the user and differ based on the policy configured. While the algorithm runs every 5 minutes, the scaling decisions also take into account trends that are tracked over time. If any of the resources consistently remain above the defined threshold, a scale-up recommendation alert is generated, and a host is added to the cluster. A vCenter Server event is posted to indicate the start, completion, or failure of scaling out on the cluster. Conversely, a scale-down recommendation alert is only generated when all resources are consistently below the threshold, triggering the removal of a host. A vCenter Server event is posted to

EXPERIENCE IT TODAY

Get a feature walkthrough

Try a Hands On Lab: VMware Cloud on AWS Getting Started

RESOURCES

Learn more about our VMware Cloud on AWS service at the VMware Cloud on AWS website

Review the VMware Cloud on AWS Solution Brief and VMware Cloud on AWS Total Cost of Ownership

Watch informative demos, overview videos, webinars and hear from our customers: VMware Cloud on AWS on YouTube

Read our latest VMware Cloud on AWS blogs

Follow us on Twitter @vmwarecloudaws and give us a shout with #VMWonAWS

Get started now with VMware Cloud on AWS: https://cloud.vmware.com/vmc-aws/get-started

Read VMware Cloud on AWS technical documentation

indicate the start, completion, or failure of the scaling in operation on the cluster.

Elastic vSAN

Elastic vSAN combines the enterprise-grade storage capabilities of VMware vSAN with automated provisioning and management of Amazon Elastic Block Store (EBS) volumes.

Elastic vSAN is best suited for workloads with high storage capacity requirements and low to moderate performance requirements. Examples are data warehouses with modest performance requirements where storage capacity consumption grows at a much faster rate than compute capacity needs.

Using a new diskless host type, Elastic vSAN enables customers to specify at cluster creation how much storage they require on a per-host basis. The VMware Cloud on AWS service will then dynamically build the hosts from the component services inside AWS. This flexibility empowers customers to control costs by deploying fewer total hosts where the capacity requirements exceed the compute and memory needs.

Users are advised to refer to the VMware Cloud on AWS Sizer Tool to find the right configuration for their needs.

Compute Policies

Compute policies provide a higher level of abstraction beyond cluster rules to capture customer intent at an SDDC level rather than at a cluster level. A compute policy consists of a capability and one or more vSphere tags. vSphere tags identify the vCenter objects to which a policy applies, whereas the capability describes the intended behavior for the objects identified by these tags. Use the vSphere client Compute Policies editor to create and delete compute policies.There are 5 compute policies that users can take advantage of to improve resource management, availability, performance and software licensing costs.

- **VM-Host Anti-Affinity policy** allows the user to specify anti-affinity relations between a group of VMs and a group of hosts. This can be useful to avoid running general purpose workloads on hosts that are running resource intensive applications to avoid resource contention.
- **VM-VM Anti-Affinity policy** allows users to spread a specific group of virtual machines across multiple hosts. This policy can be useful when users wish to prevent simultaneous failure of those virtual machines in the event that a host fails.
- **VM-VM Affinity policy** allows the user to specify affinity relations between VMs. This policy can be useful when two or more VMs can benefit from placement on the same host to keep latency to a minimum.
- **VM-Host Affinity policy** constrains the placement of tagged VMs on specifically tagged hosts in each cluster, thereby circumventing the need to define rules on a per-cluster basis. This policy can be useful when users wish to optimize software licensing costs for applications workloads.
- **Disable DRS vMotion** allows the user to specify that a virtual machine not be migrated away from the host on which it was powered-on, unless the host is placed into maintenance mode. This policy can be useful for vMotion-sensitive workloads such as large transactional databases and real-time transaction processing applications.

Conclusion

These three considerations will enable IT teams to quickly begin extending the data center footprint with VMware Cloud on AWS using known concepts, tools and processes. Learn more by reviewing the available resources and participating in a Hands-On Lab.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com
Item No: Data Center Extension_VMware Cloud on AWS_Practitioners Quick Reference 5/19