TECHNICAL WHITE PAPER - MAY 2017

MULTI DATA CENTER POOLING WITH NSX

WHITE PAPER

vmware®

Table of Contents

Executive Summary

With the rapid pace of change in business today, the availability of IT resources is increasingly in flux. Because of merger and acquisition activities, or from intentional geo redundancy planning, it is not uncommon for an IT organization to require multiple data center sites.

Yet, too often the network is preventing true workload mobility across locations. This challenge provides new opportunities for Service Providers to introduce additional service offerings to existing customers, grow their business beyond the current customer base, and overall increase revenue.

VMware NSX platform can extend the Service Provider network across administrative domains, sites, and geographies, enabling seamless workload mobility. This means Service Provider tenants now have access to a whole new level of flexibility when planning consolidation activities, upgrade activities, disaster recovery, application level redundancy, and general elasticity. Tenants are able to use the infrastructure (Service Provider hosted) in their various sites as a single pool of resources with no IP reconfiguration required.

Service Providers using the advanced network virtualization of NSX can now transparently interconnect data center capabilities between and within data center environments, whereby they can automatically distribute tenant networking for disaster recovery and business continuity purposes. Doing so greatly accelerates time to recovery, network and system survivability, and drives greater operational continuity for the tenants and providers alike.

Service Providers with a need to extend their services to clients with workloads in other geos (especially where there might be regulations restricting the export of those workloads across borders) can, without building out new data centers, lease capacity with vCloud Air, IBM SoftLayer, and in the near future, VMware Cloud on AWS. These resources can be seamlessly pooled and centrally managed with a single pane of glass.

Multi data center pooling creates a unified, seamless, and resilient pool of networking and security infrastructure to run applications across multiple data centers and to the cloud. In the same way, apps can be deployed in any location and connect to resources located across sites.

Based upon the customer requirements, multi data center pooling can be technically implemented in several different configurations.

Multi Data Center Pooling Use Cases

Service Providers may have multiple data centers or wish to implement multiple data centers for the following reasons:

- Disaster Avoidance/Disaster Recovery-as-a-Service offering
- Organic data center expansion - Data center growth
- Inorganic data center expansion - Acquisition/merger of other Service Provider data centers
- Intra- and Inter-Data Center Workload Migration
- Migration between clouds
- Multiple data centers in dispersed geographical locations to support tenants in those regions

Multi Data Center Pooling Solutions

There are several options to implement multi data center pooling with NSX.

NSX with vSphere Metro Storage Cluster (vMSC)

The use case for this design includes data centers that are close together within a metropolitan or campus area. This multi data center pooling solution has a 10ms RTT latency for storage, or 5ms RTT when using vSAN. In this configuration, there is only one vCenter Server. The cluster(s) are stretched across the sites and share the same (synchronously replicated) datastore, which requires the low storage latency.

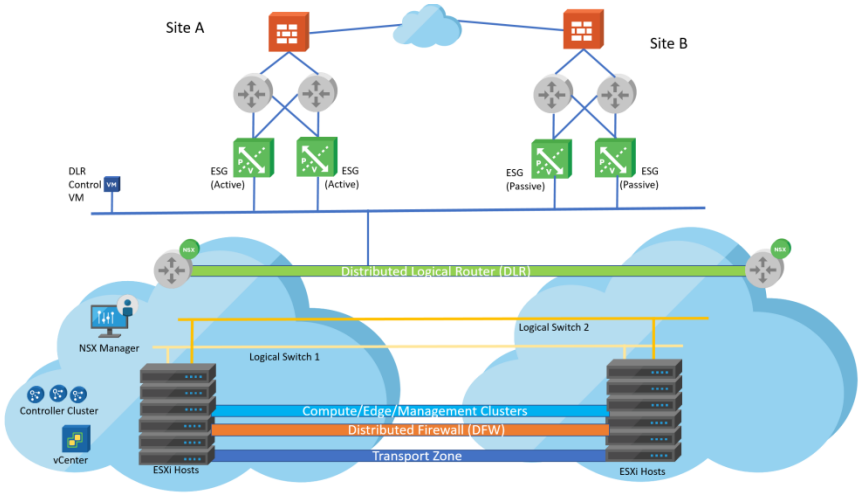

Figure 1: NSX with vSphere Metro Storage Cluster (vMSC)

VMware vMSC infrastructures are implemented with the goal of reaping the same benefits that high-availability clusters provide to a local site, but in a geographically dispersed model with two data centers in different locations.

At its core, a VMware vMSC infrastructure is a stretched cluster. The architecture is built on the idea of extending what is defined as “local” in terms of network and storage. This enables these subsystems to span geographies, presenting a single and common base infrastructure set of resources to the vSphere cluster at both sites. It stretches network and storage between sites.

The primary benefit of a stretched-cluster model is that it enables fully active and workload-balanced data centers to be used to their full potential. They acquire the capability to migrate virtual machines between sites with VMware vSphere®, vMotion®, and vSphere Storage vMotion, enabling on-demand and nonintrusive mobility of workloads. The capability of a stretched cluster to provide this active balancing of resources should always be the primary design and implementation goal. Although often associated with disaster recovery, VMware vMSC infrastructures are not recommended as primary solutions for pure disaster recovery.

Stretched cluster solutions offer the following benefits:

- Workload mobility
- Cross-Site automated load balancing
- Disaster avoidance
- Enhanced downtime avoidance

NSX enhances a vSphere Metro Storage Cluster (vMSC) by providing a faithful reproduction of network and security services used by workload VMs across all sites, without requiring physical L2 extensions between the sites. This ensures continued network operations and consistent security posture of the VMs without a need to reconfigure them or the network, even as they move from one site to another.

Service Providers leveraging NSX with vMSC can meet and exceed their availability and performance SLAs by live migrating customer workloads between sites to accommodate planned updates and outages. Also, workloads can dynamically move between sites when resources are constrained. All this occurs without reconfiguration of the workload or networks, and without affecting the network availability and security posture of the workloads before, during, or after the move.

NSX and Separate vSphere Clusters

This solution does not use shared storage and removes the 10ms requirement enforced by vMSC. Latency requirements are on the NSX control plane and vMotion. Maximum sustained latency for vMotion is 150ms RTT. This configuration uses a single vCenter server. However, the vSphere clusters are local to each site. This means that HA and DRS features cannot be used across sites. NSX distributed logical routers and distributed firewall policies straddle all sites as the single NSX/vCenter domain spans all sites.

Service Providers leveraging NSX with separate vSphere clusters per site can meet and exceed their availability and performance SLAs by scaling customer workloads across sites. Customer workloads can be deployed, moved, or even load-balanced across sites closest to the tenants to improve application response performance.

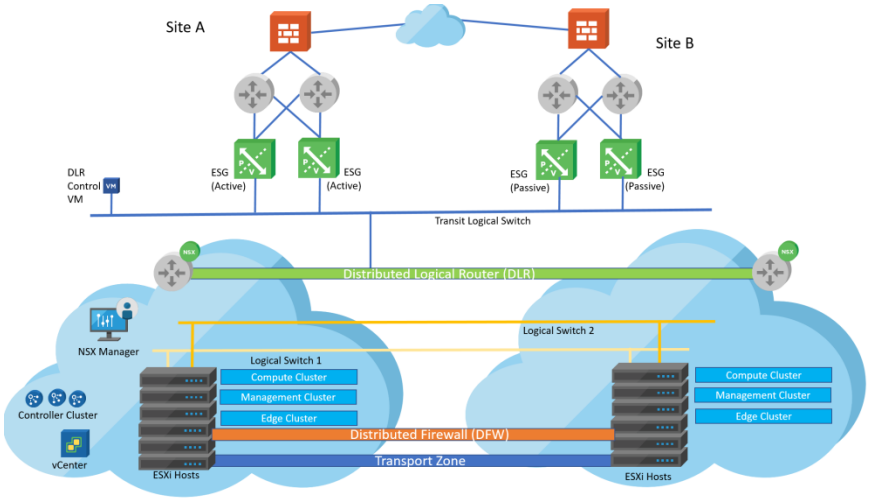

Figure 2: NSX and Separate vSphere Clusters

Cross-VC NSX

This solution provides the ability to span logical networks and security across multiple vCenter Domains and geographic locations. Each location has a vCenter Server and NSX Manager and utilizes Long Distance vMotion to seamlessly move workloads between locations. The latency requirement for the NSX Control plane and vMotion is 150 MS RTT.

Cross-VC NSX allows for Service Providers to create NSX logical networking and common security support across multiple vCenters. Logical switches (LS), distributed logical routers (DLR), and distributed firewalls (DFW) can now be deployed across multiple vCenter domains. These Cross-VC NSX objects are called universal objects. The universal objects are similar to distributed logical switches, routers, and firewalls, except they have global or universal scope, meaning they can span multiple vCenter instances. With Cross-VC NSX functionality, in addition to the prior local-scope single vCenter objects, users can implement Universal Logical Switches (ULS), Universal Distributed Logical Routers (UDLR), and Universal DFW (UDFW) across a multi-vCenter environment that can be within a single data center or across multiple data centers.

Cross-VC NSX introduces the following benefits for Service Providers:

- Logical networking and security across vCenter boundaries/sites
- Consistent security policies across vCenter boundaries/sites
- Enhanced NSX Multi-Site and Disaster Recovery
- No need of physical L2 span for Cross-VC, Long Distance vMotion, workload migration
- vCenter server no longer a mobility or scale boundary

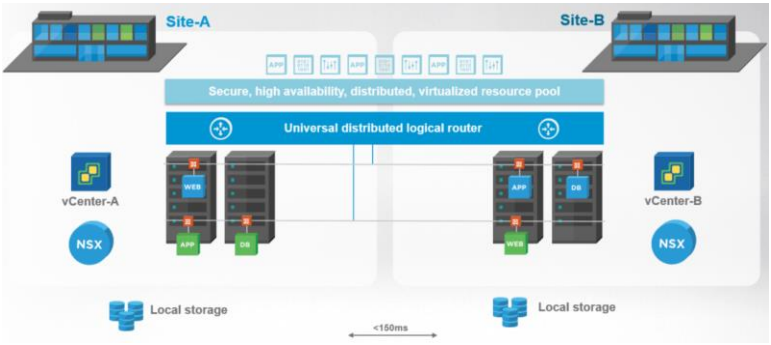

Figure 3: Cross-VC NSX

Layer 2 VPN (L2VPN)

The layer 2 VPN solution provides a simple layer 2 extension across sites. Connectivity between data centers can be VLAN to VLAN, VXLAN to VXLAN, or VLAN to VXLAN. The solution consists of having a vCenter and NSX Manager in the primary site. The secondary site may have vCenter and NSX manager to manage local resources. The latency requirement for vMotion is 150 MS RTT.

One of the primary use cases for L2VPN with Service Providers is “cloud bursting”. With cloud bursting, Service Provider tenants can horizontally scale their on-premise applications into the cloud when demand spikes, without the need or worry of having to procure, set up, configure, and maintain additional hardware (servers, networking equipment, and storage) on-premise. This effectively creates a **Hybrid Cloud** solution.

Another primary use case is Cloud/Data center migration. With Cloud migration, tenants migrate their on-premise workload into the cloud host by a vCloud Air Service Provider. Now more than ever, tenants have compelling reasons for migrating some or all of their workloads to the cloud. These include:

- Migrating their Internet facing applications to the cloud, to decommission their DMZ, thus minimizing the scope and complexity of their networking and security configuration
- Reducing the CapEx and OpEx associated with hardware acquisition, deployment, configuration, and support
- Reducing the CapEx and OpEx associated with maintaining a physical data center

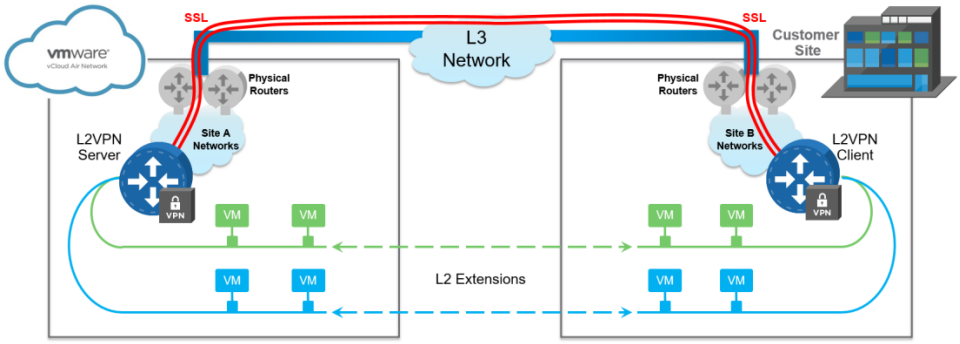

Figure 4: Layer 2 VPN (L2VPN)

Multi Data Center Pooling Solutions Comparison

	Stretched Cluster	Separate Clusters Cross-vC NSX	L2VPN
Scope	Metro	Geo	Global
Latency (Max)	10ms (5ms with VSAN)	150ms	Any
Network	1600 Bytes MTU >=10Gbps	1600 Bytes MTU >=1Gbps	Any <1Gbps
Operations	Common Admin Domain	Common Admin Domain	Common or Separate Admin Domain
Storage	Metro Storage required	Independent Storage	Independent Storage
Features	• Seamless pooling across DCs • Leverage vSphere HA & DRS with consistent networking and security across sites	• VM mobility across sites • Resource Pooling • Consistent networking and security across sites • Cross-VC NSX to provide an enhanced DR solution	• NSX at one or both ends • L2 extension to cloud

Summary

With NSX, Service Providers can now create logical networking and security constructs that span multiple vCenter domains and multiple sites. This capability allows Service Providers to enhance their current offerings, create new offerings, and thus increase revenue.

This capability also allows Service Providers to implement Active-Active data centers, which allows for the following:

- Workload Mobility (live/cold migration) – Tenant workloads can be moved between sites to accommodate disaster avoidance, planned, and unplanned outages.
- Resource Pooling – Tenant workloads can be dynamically or manually moved to accommodate better resource utilization. Resources are no longer isolated based on vCenter boundaries, and idle capacity within another vCenter domain/site can be leveraged for better overall resource utilization.
- Unified Logical Networking and Security policy – Service Providers can create consistent networking and security policies across vCenter domains/sites. Manual replication of networking and security policies or investment in proprietary hardware is no longer a requirement when considering multi data center pooling.
- Disaster Recovery – With networks and security spanning multiple sites, tenant applications can recover in the recovery site and retain their network (IP) configuration. Tenants can also failover/migrate their on-premise applications to a Service Provider data center without a need to change their application configuration, and thus minimize downtime.

Service Providers can also easily extend their presence into other geographies by leveraging capacity in vCloud Air, IBM SoftLayer, and soon, VMware Cloud on AWS, all without having to build and maintain additional data centers.

-

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com