NSX DESIGN GUIDE FOR SMALL DATA CENTERS

Version 2.0 (December 2016)

NSX DESIGN GUIDE SERIES

Intended Audience

This document is targeted toward virtualization and network architects interested in deploying VMware NSX network virtualization solution in a vSphere environment.

Revision History:

Version	Date Released	Comments
1.0	August 2016	First Release
2.0	December 2016	Added new content and fixed some figures

Executive Summary

VMware NSX is the network virtualization technology that decouples the networking services from the underlying physical infrastructure. VMware NSX allows for a new software based approach to networking that provides the same operational model as a virtual machines (VM). Virtual networks can easily be created, modified, backed-up and deleted within minutes.

By providing the physical networking constructs in software, VMware NSX provides similar benefits as server virtualization did with VMs. Businesses can see the impact in terms of increased efficiency, effective resources utilization, productivity, flexibility, agility and cost savings.

Document Structure

This document will present the audience with the NSX introduction, business use-cases and overview of design in Large and Medium data centers. The beginning of the document serves as a refresher to those who are already familiar with the NSX design and deployment aspects.

The document goes on to present a NSX for small data center, its relevance, and what are the main building blocks of designing NSX in small data centers.

The document talks about popular NSX deployment models in small data centers, gives details around protecting and designing based on the individual NSX components, like NSX ESG and DLR Control VM etc.

Towards the end of the document, it talks about the growth option to take NSX even further and grow it into the medium and large scale deployment.

Introduction

NSX has emerged as the leading software platform to virtualize network and networking services. Many customers have deployed NSX to run their production and non-production workload to get the benefits that comes with virtual networks and software defined network approaches. NSX has been deployed from small to medium to large sizes of data centers to enable a wide-range of use-cases.

There are situations where large enterprises have also deployed NSX in their small data centers islands within the overall large environment. There are also situations where small and medium businesses (SMBs) are deploying NSX with small number of hosts to take advantage of network virtualization. Regardless of the size of the enterprise, small data center is a viable option and relevant for all type of customers, enterprises and businesses.

The NSX Reference Design Guide discusses design aspects to deploy NSX in all data center sizes. This document uses the NSX Reference Design Guide as a baseline and provides additional and/or supportive guidance to successfully run NSX in SMB Data Centers. It is assumed that readers have gone through the concepts and design options discussed in NSX reference design guide.

In addition, readers are highly encouraged to take a look at Software Defined Data Center (SDDC) VMware Validated Design Guide (VVD) that provides most comprehensive and extensively tested blueprint to build and operated SDDC.

NSX Customer Use Cases

NSX has been widely accepted and deployed in production by many customers. Figure 1 lists some of the most important use cases that customers are deploying NSX for.

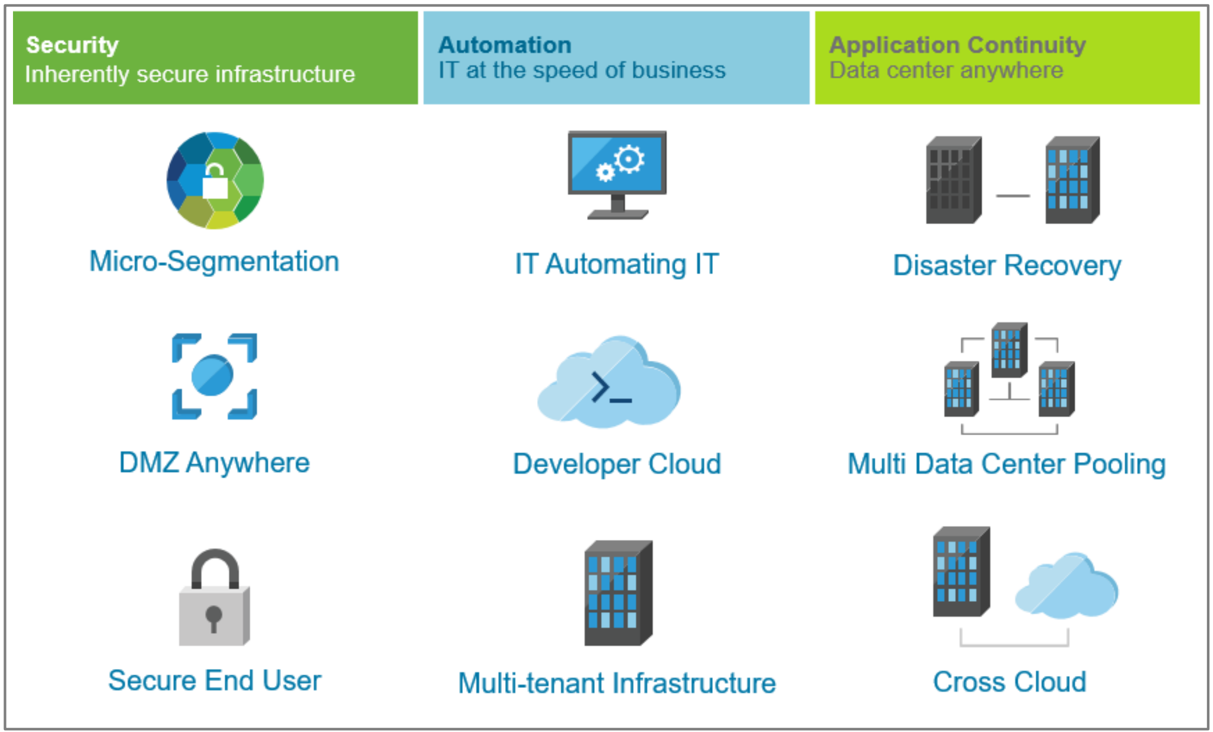

Figure 1

It is important to realize that regardless of the size of the data center, customers can take advantage of these use-cases with NSX. NSX is a software product which is extremely flexible and can be designed to suit various deployment models and use-cases that go beyond Security, Automation and Application Continuity use-cases.

Security

NSX can be used to create a secure infrastructure, which can create a zero-trust security model. Every virtualized workload can be protected with a full stateful firewall engine at a very granular level. Security can be based on constructs such as MAC, IP, ports, vCenter objects, security tags, active directory groups, etc. Intelligent dynamic security grouping can drive the self-adaptive security posture within the infrastructure.

Automation

VMware NSX provides a full RESTful API to consume networking, security and services, which can be used to drive automation within the infrastructure. In small data centers, automation tools like REST API and PowerNSX can be useful to programmatically configure network and security services, or to pull the information from VMware NSX deployments for simple operations tasks.

Application Continuity

NSX provides a way to easily extend networking and security up to eight vCenters either within or across data centers. NSX can extend or stretch L2 and L3 networks across data centers in distributed fashion. NSX also ensure that the security policies are consistent across those stretched networks and hence provide a seamless, distributed and available Network and Security overlay. All of it is done using software based technologies, without requiring expensive hardware.

The next section is an introduction to NSX for vSphere. Readers familiar with this topic can skip the next section and can directly jump to NSX in Small Data Center Use-Cases.

NSX for vSphere Components

vSphere is the foundation for NSX for vSphere (referred to as NSX throughout this document) deployment. It is important to have good understanding of what vSphere and NSX components are involved into the design. For a successful NSX deployment, it is imperative to have a good vSphere deployment in place with proper vSphere clustering, compute, network and storage. For detailed discussions on these topics, the reader can refer to the NSX Reference Design Guide.

Figure 2 shows various layers of NSX for vSphere architecture based on the role being performed by each NSX components. From a very high level, the NSX solution architecture can be seen as divided between management, control and data planes. In the traditional networking model, the control and data plane is combined together. NSX and other software defined networking architectures follow an approach where the data plane is separated from the control plane. This approach provides the advantage of decoupling from hardware dependencies, and allows all networking services to be virtualized following the same operational model that compute and storage virtualization has been providing for years.

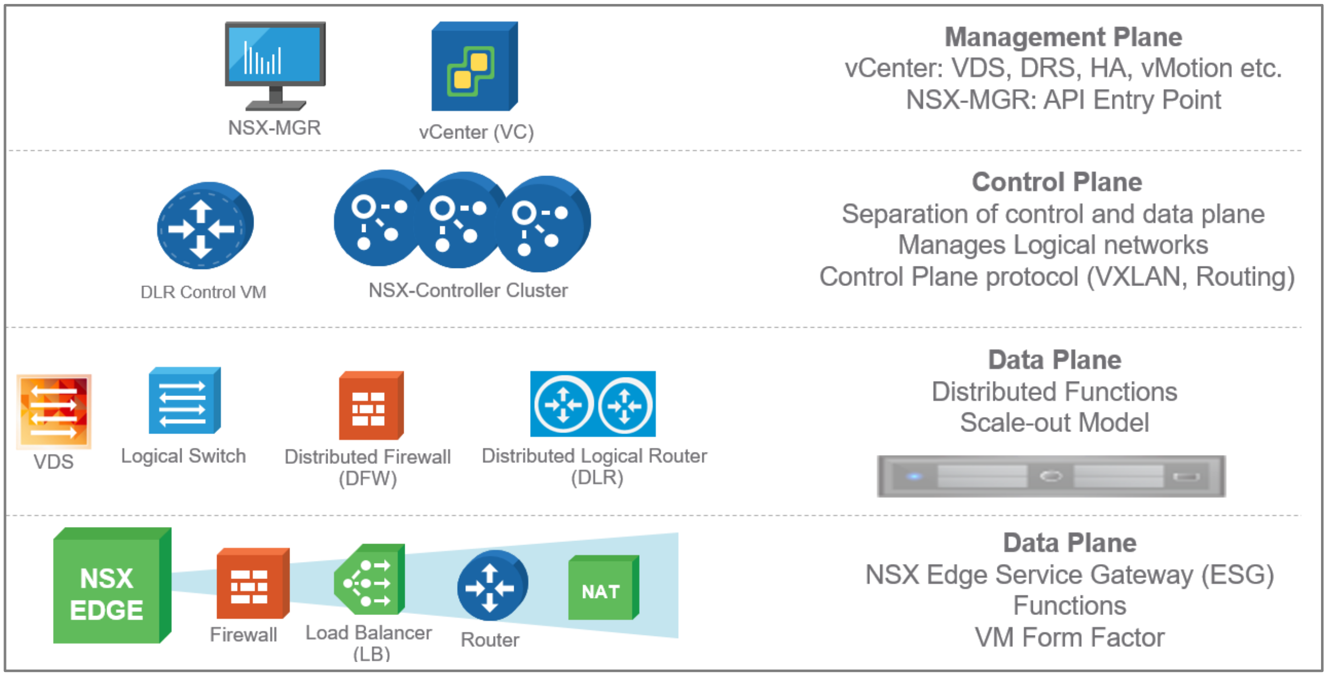

Figure 2

NSX Edge Service Gateway (ESG) is a virtual machine form factor appliance. ESG sits at the edge of the virtual network and can be considered as on-ramp/off-ramp point between physical and virtual network. ESG VM can provide many edge services like Firewall, Load Balancer and Edge Routing etc. The NSX reference design recommends to deploy edge related VMs into their own separate clusters because their compute, network and storage requirements are very specific to their roles and responsibilities.

NSX in Small Data Center Use-Cases

One must understand that Small Data Center (DC) does not mean that it is only relevant for small customers. Many large enterprises deploy NSX with small footprint or small number of ESXi resources in the beginning and then they expand to larger footprint. This could be due to number of different reasons for example budget, staffing or simply because of small scale deployment that they would have in the beginning. The advantage is that even if NSX is deployed in small footprint, it can easily grow into a medium or large size deployment.

On a broad scale, Small Data Center use-cases can be divided based on business function and application that are being deployed.

Functional Level Use Cases

Organizations deploy NSX with small footprint in specific functional areas or groups that they have. For instance

- Disaster recovery and/or avoidance
- Pre-Prod vs Test environments
- Compliance / DMZ
- Business units with their own operational model
- Etc.

Application Level Use Cases

Many customers deploy NSX in small DCs to tackle one or more application level use-cases that they have. For instance

- VDI
- Load Balancer
- Agentless Antivirus (AV)
- Etc.

NSX Advantage for Small Data Centers

Organizations adopt NSX not just because of its technical strength and advantages that they gain while deploying networking services in software. They also get the advantage in terms of its simplicity, ease of use and operational flexibility. Some of these advantages are highlighted here.

Simplicity and Modularity

Small customers like the idea of its simplicity and modularity, where they have peace of mind to grow and add more features as they increase the capacity or user base. They do not need to purchase all the networking hardware upfront with lots of unknown down the road. NSX provides those customers software based networking services that they can spin up anytime they want without incurring additional hardware cost.

Procurement

Customer are also thrilled because all the networking and security services are bundled within the same product and platform, so they do not need to worry about contacting multiple vendors not just for purchase but also for support agreement and licenses procurement and cost. Customer are getting everything with the NSX under one roof.

Ease of Operations

Majority of the customers are already familiar with the operational model vSphere has provided them for years. NSX is seamlessly integrated within the same model. It enhances their operational model and sits nicely on top of it. Hence the learning curve to adopt the new technology is minimal.

The next section discusses the importance of well architected vSphere design. vSphere design (including but not limited to vCenter, Platform Services Controller, vSphere Cluster, VDS, DRS, HA etc.) is the foundation for a successful NSX deployment. These aspects must be reviewed carefully prior to designing NSX based overlay, virtual networks and services.

vSphere Design

The vSphere and vSphere cluster design plays a very important role for NSX deployment. vSphere deployment best practices recommend to create purpose build clusters. Purpose built cluster not only limit the fault domain but also help carve out resources based on the functions that these clusters perform.

In large data centers with many ESXi hosts, the recommendation to have multiple vSphere clusters can easily be justified and achieved. In small data centers (due to availability of few ESXi hosts or constraints like budget) customers can also deploy management, compute and edge workloads into one single vSphere cluster. Such a design would require careful allocation, reservation and monitoring of resources. In an ideal situation, one should avoid to oversubscribe the vCPU and memory resources in small data centers.

NSX in Single vSphere Cluster

The NSX deployment best practices suggest to have separate vSphere clusters for management, edge and compute virtual machines and distributed functions. There are situations and scenarios where customers may not be able to dedicate ESXi host resources to create multiple vSphere clusters for management, edge and compute roles. If a vCenter (VC) deployment consist of a single vSphere cluster where management, edge and compute resources are collapsed into one single cluster, such a DC will be designated as a small DC throughout this document.

This guide discusses considerations and guidance to deploy NSX in small DCs and provide network architects confidence to successfully design, deploy and operate NSX.

The following section provides a high-level overview of NSX vSphere cluster design in large, medium and small data centers.

Large Data Center Cluster Design

A large sized DC could typically have hundreds of ESXi hosts running thousands of virtual machines. The north-south bandwidth requirement is usually more than 10Gig. Figure 3 shows separate management, edge and compute clusters hosting respective workloads/VMs in large DC. It also shows possible distribution of various management and edge VMs on separate hosts within the cluster to avoid single point of failure. The design guidance of such data center is discussed in the NSX Reference Design in details.

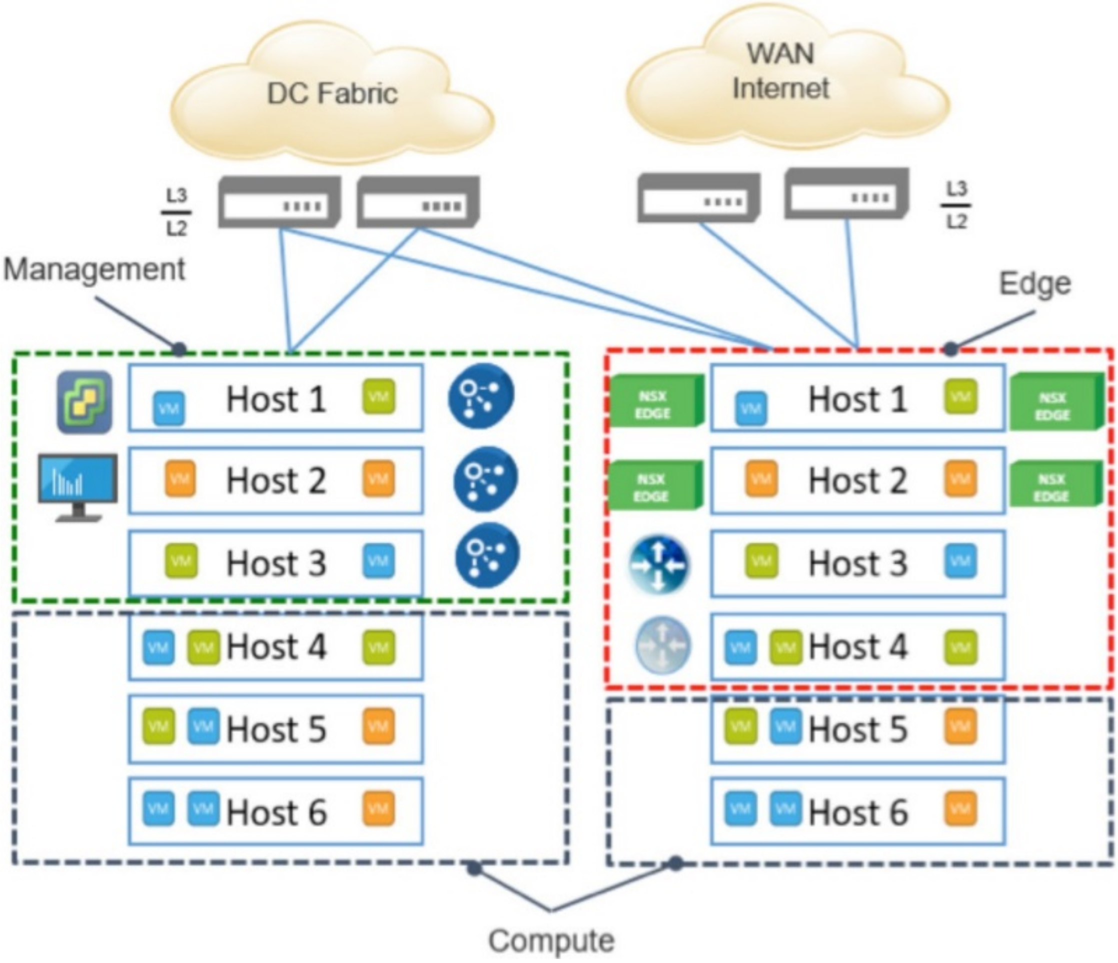

Figure 3

Medium Data Center Cluster Design

A medium sized DC could typically have somewhere between 10-100 ESXi hosts running from hundreds of virtual machines. The north-south bandwidth requirement is usually less than 10Gig.

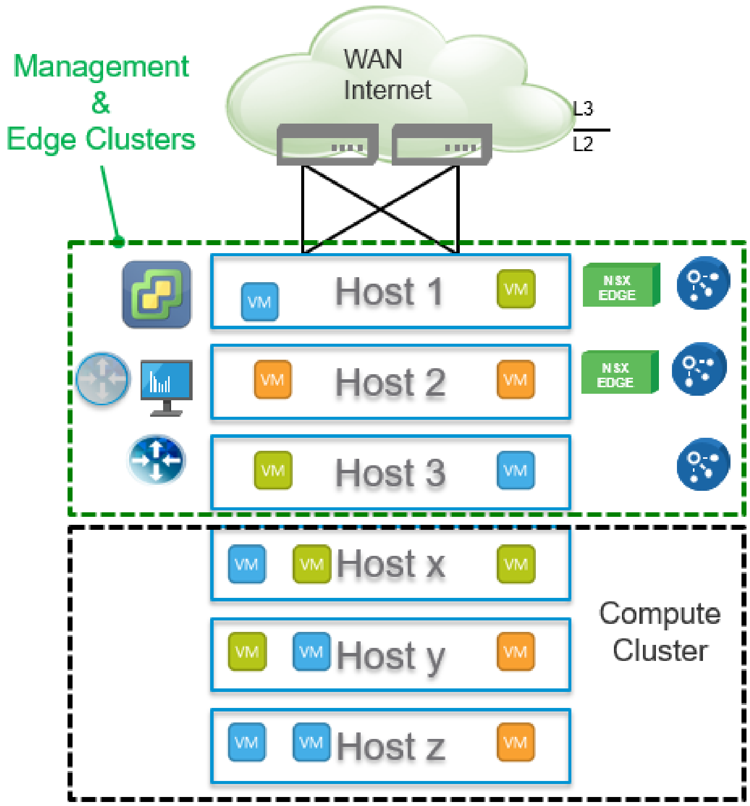

Figure 4

It is recommended to have separate clusters like in the case of a large DC deployment. There are situations where it is not possible or feasible to carve out separate management, edge and compute clusters. Figure 4 shows that in those situations, customers can collapse management and edge resources into a single cluster and can have single or multiple compute clusters for compute workloads and applications. Figure 4 also shows a possible distribution of management and edge VMs to protect again single point of failure. The design guidance of such data center is discussed in the NSX reference guide in details.

Small Data Center Cluster Design

A small size DC would typically have less than 10 ESXi hosts running somewhere between ten to 100 virtual machines including management, edge and compute workloads. The north-south bandwidth requirement is almost always less than 10Gig.

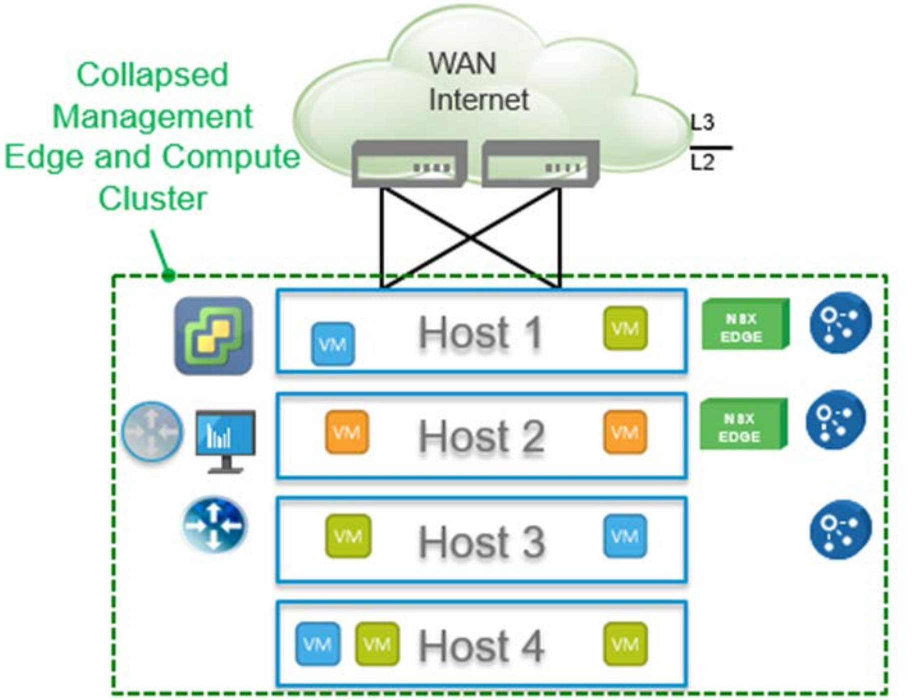

Figure 5

In a small DC case, since the number of ESXi hosts are less, there is no other option but to deploy NSX in one single cluster. Deploying NSX in a single cluster, requires some considerations especially resource reservation that should be carefully monitored and assured. The best part of the solution is that it leverages both vSphere and NSX technologies to provide the assurance and confidence needed for such deployments.

The rest of the document is focused on options and choices that a customer should consider when deploying NSX in small DC situations.

Small DC Deployment Models

NSX is extremely flexible and expandable. It can easily be implemented with many deployment models, variations and options. For small and medium (SMB) data centers customer have mainly deployed NSX in the following deployment models

- Security Focused Deployment Model
- Centralized Edge Deployment Model
- Full Stack Deployment Model

The Security focused deployment model is the most popular and widely deployed model not only in small and medium data centers but in large enterprises as well. This is mainly achieved using the NSX distributed Firewall functionality. This could be the entry point in majority of the NSX deployments.

The centralized edge deployment model is based on the Edge Services Gateway (ESG) VMs. It could be considered in situations where a customer is transitioning from security focused deployment model to full stack model but not ready to change the network MTU (Maximum Transmit Unit) for VXLAN (Virtual Extensible LAN) enablement.

The full stack model is where customer will get full benefits for virtual networking with features like logical networking (VXLAN), distributed routing etc. In this model customers create networking topologies and offer various networking services all in software based overlay. The general recommendation from VMware is to deploy NSX with Full Stack model.

The emphasis here will be on Security Focused and Full stack deployment models. This document will be covering the deployment options that are widely adopted among many production customers. It is important to understand that the order in which these models are listed does not mean a customer must deploy them in the same order.

Security Focused Deployment Model

Security focused deployment is the most popular choice for deployment in small and medium data centers. This model is also the foundation of Micro-Segmentation that allows applications to be segmented based on their security requirements and not based on physical constructs like IP addresses or VLANs.

Design Considerations

Figure 6 shows a logical topology of security focused deployment. The main component in this model is NSX distributed firewall (DFW) which is an unmatched scale-out, distributed and stateful firewall technology embedded right into the hypervisor kernel. DFW deployment is non-disruptive. VXLAN or any changes in physical (like MTU or routing) infrastructure are not required. NSX manager VM is the only VM that needs to be deployed to implement Micro-Segmentation (or DFW) and it does not need NSX Controller or Edges to be deployed.

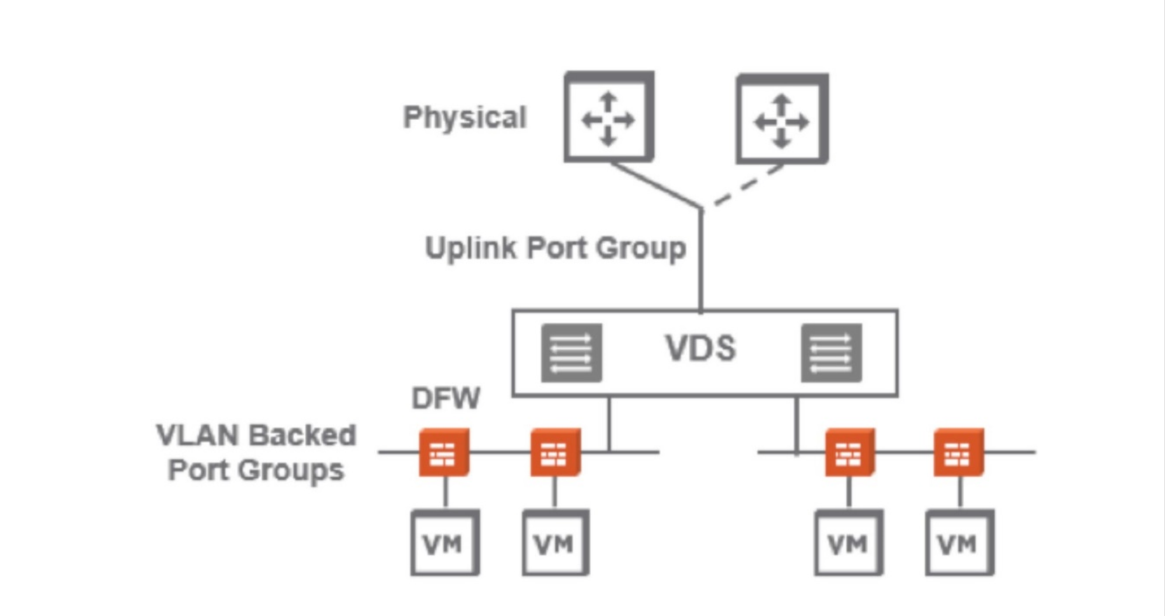

Figure 6

In small DC, management and mompute resources will be collapsed in a single vSphere cluster. Table 1 shows the components and functions required to design this model. Since this model does not require VXLAN, customers can use existing VLAN backed-port groups to implement DFW.

Planes in Single Cluster	Components / VMs
Management Plane	NSX Manager, VC, LogInsight, vRops and other management VMs
Compute	Compute VMs Service VMs
Data Plane	ESXi Kernel Components (NSX Distributed Firewall and vSphere Distributed Switch)

Table 1

Customers can also easily add other security services like Agentless Antivirus (AV) with this model without re-designing or disturbing existing DFW implementation. Agentless AV solution is delivered via NSX Guest Introspection technology. And it requires additional service VMs and integration with NSX security partners in this space.

For additional information refer to NSX Design Guide.

Security Focused Deployment Model with Micro-Segmentation (DFW) Use-Case

The footprint to implement this mode is extremely minimal. Customers can easily deploy it on two ESXi hosts. It is highly recommended to have at least three hosts in production.

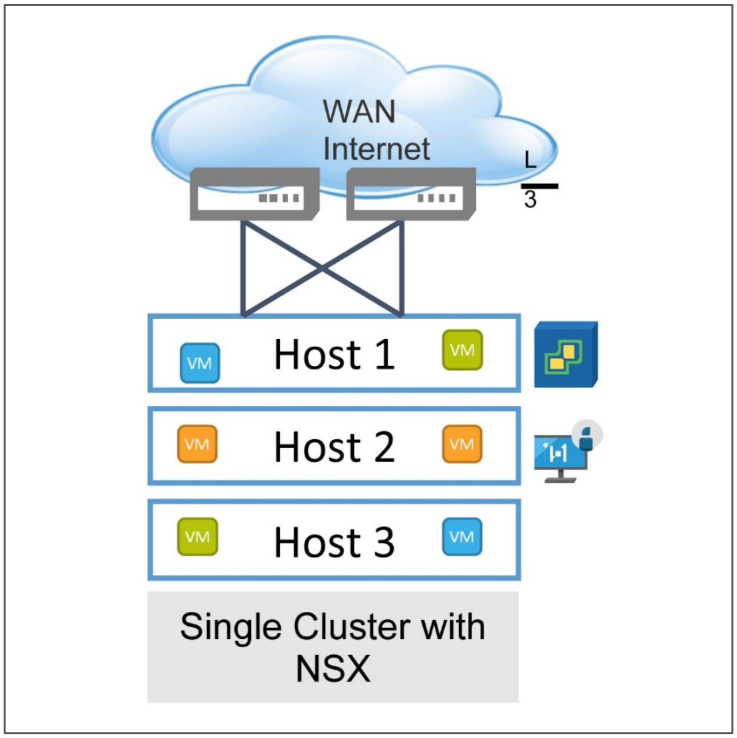

Figure 7

Figure 7 does not show the requirement for compute workload. Customers should size according to their workload requirement to add additional compute hosts in the cluster. It is always advisable to have more hosts with adequate capacity to sustain at least single host failure.

VM Name	vCPU	MEM (GB)	Storage (GB)	VM Count
Tiny vCenter Appliance with Embedded PSC	2	8	116	1
NSX Manager	4	16	60	1
Total	6	24	176	2

Table 2

Table 2 shows an example where a customer deployed this model using only two virtual machines.

Security Focused Deployment Model with Agentless AV Use-Case

This use-case could be an add-on to the DFW use-case or could be deployed on its own without DFW. This model would require NSX Guest Introspection Service VM (GI-SVM) and additional partner Service VM (Partner-SVM) per ESXi hosts. These services VMs (GI-SVM and Partner-SVM) are automatically deployed at the vSphere cluster level. SVM are automatically deployed when a ESXi hosts enters into the cluster and automatically removed when an ESXi host level the cluster.

The SVMs are tied to one particular host when deployed. Hence it is extremely important to note that these SVMs should not be vMotioned or Storage vMotioned to any other hosts.

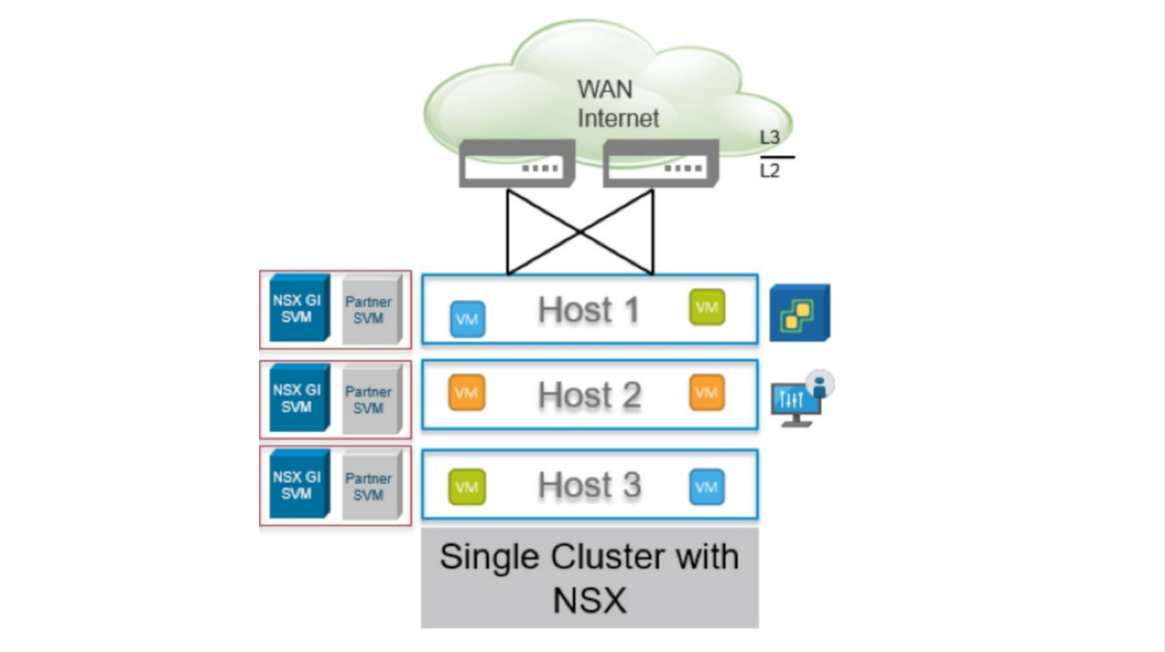

Figure 8

Figure 8 shows the design layout for this use-case in one single cluster. Notice the addition of NSX GI SVM and Partner SVM.

Centralized Edge Deployment Model

NSX Edge Service Gateway (ESG) VM is the core component in this deployment model. ESG VM can be deployed with or without NSX DFW. This deployment model could be adopted as an intermediate step going from security only deployment to NSX full stack deployment (with VXLAN, DRL etc.). This model is not a very popular choice among customers because it does not give full benefit of what customers can get with the NSX full stack deployment.

Design Considerations

NSX Edge VM in the core element in this design. Customers usually pick up this model if they do not have requirement to optimize their East-West routing or do not have much East-West traffic in their data centers.

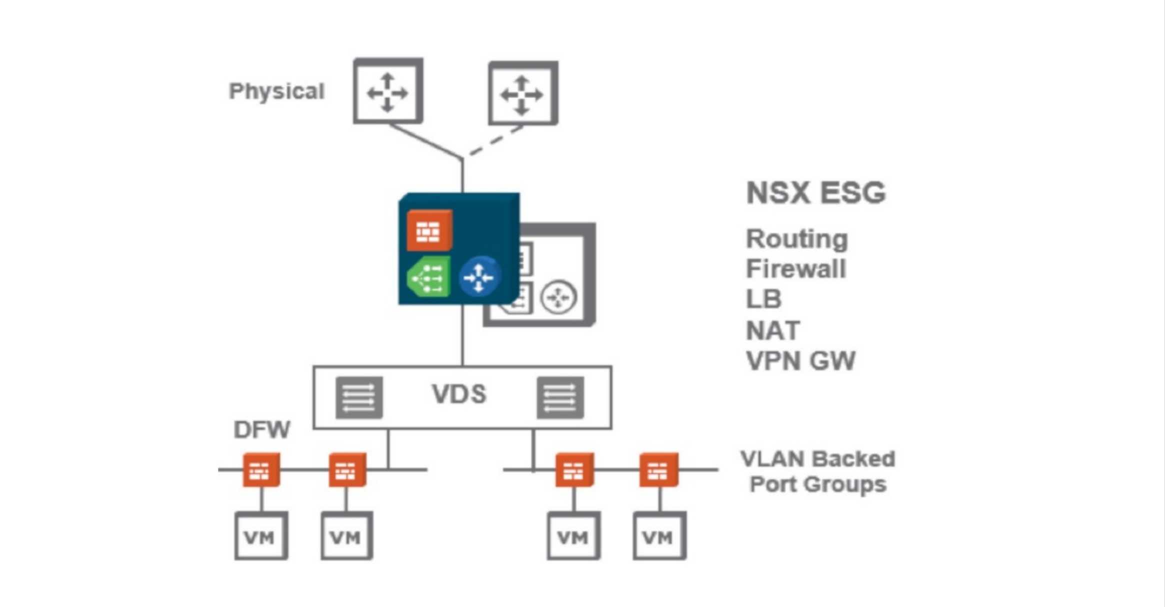

Figure 9

NSX ESG is a multi-function gateway in VM form factor and one can enable multiple services on a single ESG like dynamic/static routing, edge firewall, NAT VPN etc. Figure 9 shows a logical topology where NSX ESG VM is connected to physical device (physical router, switch or firewall) on North side. In small DC since the N-S bandwidth requirement is not very high hence this edge can be deployed in HA mode (active/standby VMs). NSX Edge HA and vSphere technologies (like vSphere HA, DRS etc.) also help improving overall availability of this deployment model.

Depending on the requirement and use-case, customers can also enable ECMP on the centralized ESG to achieve higher bandwidths. In case of ECMP it should be noted that all the Edge stateful services (like NAT, Firewall etc.) will be turned off by default.

Customers can deploy this model with minimal modification to their existing physical infrastructure. There is no MTU change required on the infrastructure and it can be implemented with existing VLAN backed port-groups. This model cannot take advantage of NSX VXLAN or DLR functionalities.

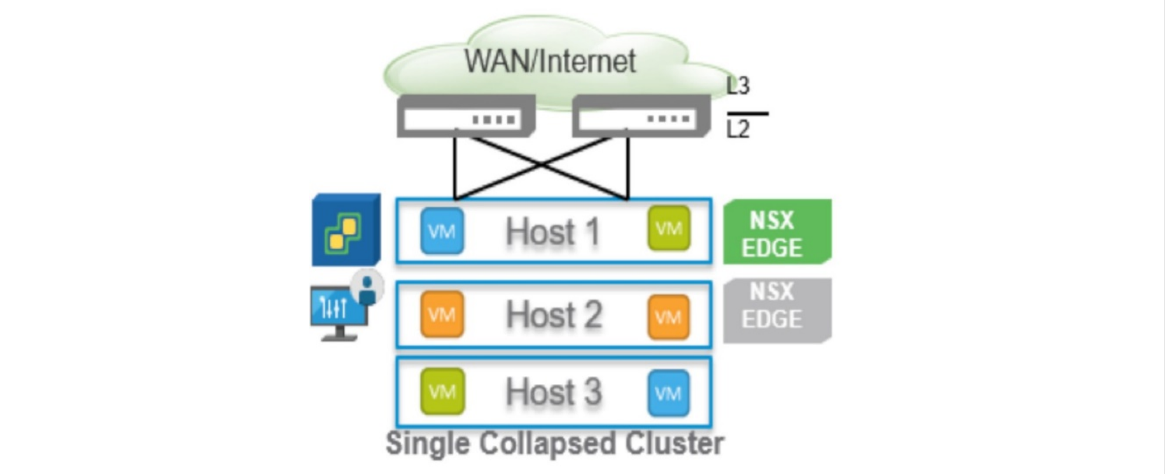

Figure 10

Figure 10 shows a possible placement of VMs on a single vSphere cluster with 3 ESXi hosts. In this layout, the NSX Edge VM is deployed in HA. It is important to make sure that both ESG VMs are not placed on the same host because in case of host failure, both active/standby VMs will go down and it would impact accessibility to applications. It is highly recommended to use DRS with proper anti-affinity rule to provide protection against single host failure.

Centralized Edge Deployment Model with VXLAN Logical Switches

There are situations where customer requirement is to deploy applications behind VXLAN backed-port groups (also known as VXLAN Logical Switch or simple logical switch) to take advantage of seamless L2 extension across racks. Notice that this model does not have a requirement to deploy NSX DLR (for East-West Traffic optimization). With these requirements it is possible to enable VXLAN Logical Switch with centralized edge deployment model.

Deploying application and workload behind logical switch also give customer advantage to easy migration to full stack deployment when they are ready for next phase. This model requires deployment of NSX Controllers and MTU must be at least 1600 byte for the VXLAN transport VLAN. The default gateway for workload/or application VMs will be the internal interface of ESG. When customer is ready to deploy DLR and NSX in full stack mode in future, they can easily switch the default gateway from ESG to DLR with minimal intervention.

In this model all the cluster functions (management, edge and compute) will be residing on the single vSphere cluster. The entire cluster will be prepared for NSX which means that DFW and VXLAN will be available on all hosts. Refer to the “Full Stack Deployment Model” for additional design consideration with Central Edge with VXLAN Logical Switch model.

Full Stack Deployment Model

Full stack model provides full abstraction from underlying hardware by utilizing industry standard VXLAN technology. It allows customers to deploy networking topologies in software and eliminate the dependency on physical hardware and physical hardware constraints. This model not only provides distributed firewall (micro-segmentation) but also optimize east-west routing with NSX logical switch (LS) and distributed logical router (DLR). Customers get the benefit of making security and routing decision closed to the workload.

This design is also popular for disaster avoidance and disaster recovery (DR) scenarios where customers want to recover a portion of their main site to a small data center location.

Design Considerations

This mode requires MTU to be increased to at least 1600 MTU or more on the physical infrastructure. This MTU requirement is only required for the VLAN that will be used for VXLAN transit traffic (or for VTEP to VTEP communication).

Cluster Function	Components
Management Plane	NSX Manager, Controllers, VC, DB Server and other management VMs
Compute	Compute VMs Service VMs
Data Plane East-West	ESXi Kernel Component (VXLAN, DLR, DFW, VDS)
Data Plane North-South	Active/Standby DLR Control VM ESG VM (HA or ECMP Mode)

Table 3

Table 3 shows the placement of various VMs and components in single vSphere cluster based on the functions they are performing.

Deployment Considerations

In this model all the cluster functions (management, edge and compute) will be residing on the single vSphere cluster. The entire cluster will be prepared for NSX which means that DFW and VXLAN will be available on all hosts.

It is mandatory to have at least 3 ESXi hosts for the deployment because full stack mode does require NSX controllers cluster. NSX controller cluster is clustered formed with 3 VMs. These NSX controllers VMs should be running on separate hosts all the time. Anti-affinity rule should be created to make sure they are not running on the same hosts.

It is recommended to at least have 4 ESXi host for a production deployment which would provide resiliency in case of single host failure. Figure 11 shows a possible placement of VMs for this model. In order to reduce the footprint, static routing is used between NSX DLR and NSX edge hence there is no DLR control VM shown in figure 11.

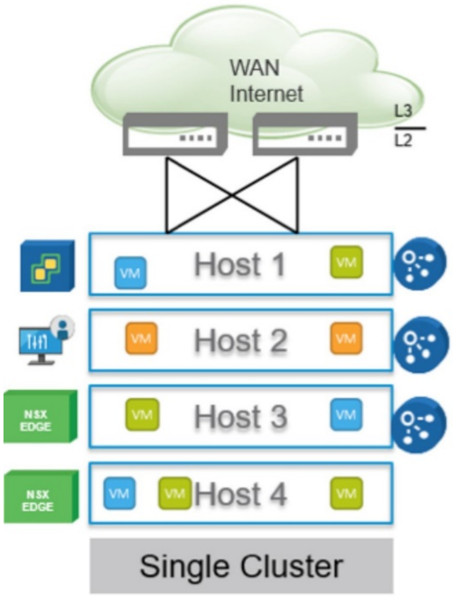

Figure 11

Customers can use Table 4 as a reference to calculate the resources needed to deploy VMs in this model. Notice that this does not include requirement for workload VMs. Workload VM requirement should be calculated and taken into consideration beside the NSX components VMs and is outside the scope of this document.

Dynamic routing (BGP or OSPF) is also supported in this model. Dynamic routing between DLR and ESG will require deployment of DLR Control VM in HA and it will increase number of VMs by 2. ESG VM comes in different form factors. It is recommended to use large form factor for majority of production deployments.

Component	VMs	vCPU	MEM (GB)	Storage (GB)
Tiny vCenter Appliance with Embedded PSC	1	2	8	116
NSX Manager	1	4	16	60
Controllers	3	12 (3x4)	12 (3 x 4)	60 (3x20)
Edge VM (Large)*	2*	4 (2x2)	1 (2 x 0.5)	2 (2 x ~1)
Total	7	22	37	~ 238

*Table 4 (* ESG with High Availability with static routing)*

Individual Component Considerations

The focus of this section is on design and deployment aspects for individual components that are involved in NSX deployment. There are features and resiliency options built into both NSX and vSphere products and one should leverage both to make sure the design is protected. With the proper use of these knobs and features, customers can successfully design their NSX Data Center to sustain failure conditions. Properly sizing and placement of these individual components will also provide confidence to network architect that their environment is stable and meet business SLAs.

vCenter Server

For NSX-v deployment, vCenter (VC) and vSphere Distributed Switch (VDS) are pre-requisite components. This document assumes that vCenter is already deployed with vSphere deployment best practices.

vCenter can be deployed in Tiny, Small, Medium and Large configuration. Table 5 illustrates the number of hosts and VM that are supported per vCenter deployment type. Table 5 also shows a potential mapping between VC deployment option to NSX deployment type.

VC VM Options	Max Number of Hosts	Max Number of VMs	Potential NSX Deployment Type	vCPU	MEM (GB)	Disk (GB) Embedded PSC
Tiny	10	100	Small DC	2	8	116
Small	100	1000	Small DC	4	16	136
Medium	400	4000	Medium DC	8	24	275
Large	1000	10,000	Large DC	16	32	325

Table 5 – VC Deployment Options

For small DC, most likely, a customer would have either Tiny or Small VC VM. Choosing Tiny or Small VC VM size is not a mandatory requirement. A customer might deploy medium or large size VC VM with less than 10 ESXi host. Medium or Large VC VM deployment will result in better performance and easy growth in future.

vSphere License

Designing NSX in a small DC require increased reliance on vSphere High Availability and DRS features. In general, proper resource reservation is a key to make sure both NSX and compute VMs are running at healthy levels. But all or some of those options might not be at one's disposal due to the VC license type. Table 6 lists some of the most popular vSphere license types and the supported features.

vSphere License Options	Notable Features/Limitations
Essential+	Supports up to 3 hosts, standard vSwitch and vSphere HA
Standard	Supports up to 1000, standard vSwitch and vSphere HA
Enterprise+	Supports all enterprise grade DC virtualization features (HA, DRS, VDS etc.)

Table 6 – vSphere License

VDS is one of the main requirements to install and support NSX deployment. VDS is only available with vSphere Enterprise Plus license in legacy option. Only NSX license allows customers to use VDS regardless of their type of vSphere license.

NSX License

There are three different licensing tiers for NSX. Table 7 lists some of the NSX features and NSX licensing requirement.

Features	Standard	Advance	Enterprise
Distributed Routing and Switching (DLR/VXLAN)	✓	✓	✓
NSX ESG (except load balancer)	✓	✓	✓
SW L2 bridging	✓	✓	✓
Distributed Firewall (DFW – Micro-Segmentation)		✓	✓
NSX Edge load balancing		✓	✓
Cross vCenter NSX			✓

Table 7 – NSX license options and corresponding features

Customer can easily compare various licensing option on VMware web site at following URL:

http://www.vmware.com/products/vsphere.html

vCenter Deployment Options

The most popular deployment option for small DC is to deploy embedded PSC with vCenter appliance as shown in Figure 12.

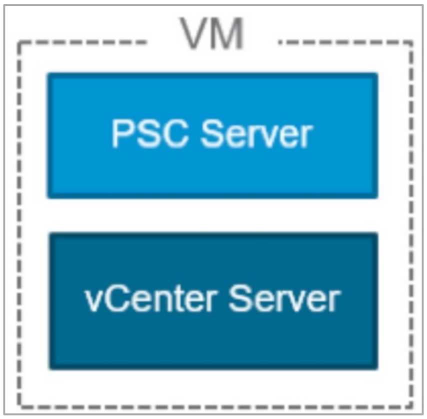

Figure 12

External PSC deployment option is an alternate option which is recommended for medium-large environments with multiple vCenters. This option is shown in Figure 13 and is useful in situation where customer is planning to grow number of virtual machines in future.

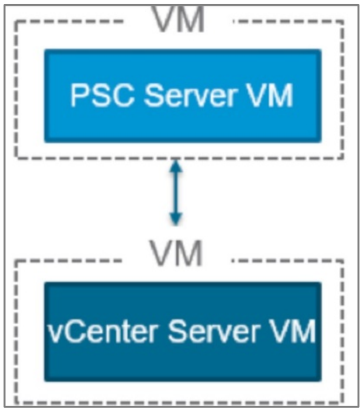

Figure 13

NSX Manager

NSX manager virtual machine deployment is done via an OVA file which can be downloaded from VMware website. Since NSX manager VM is not a locked-down VM hence it is possible to modify its vCPU and MEM values via vCenter UI. It is highly recommended to stick with the default values. Changing NSX manager VM setting may lead to error, disruption and/or poor performance.

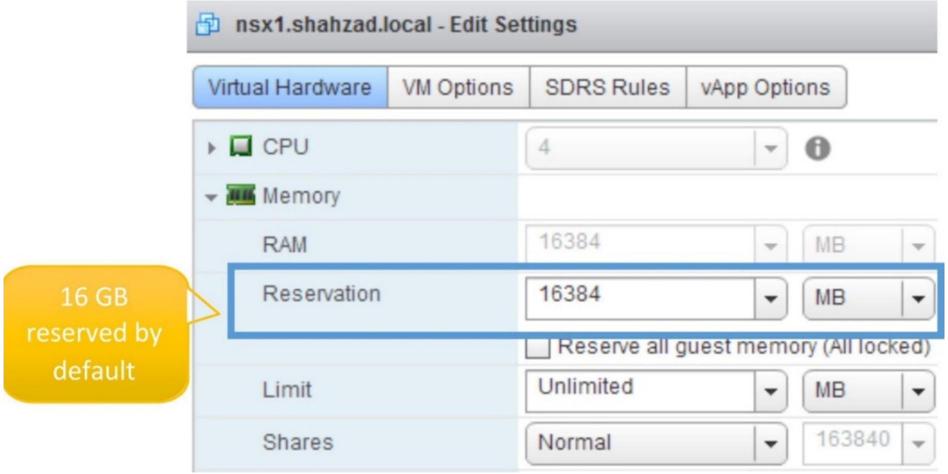

nsx1.shahzad.local - Edit Settings

CPU: 4

Memory

RAM: 16384 MB

Reservation: 16384 MB

□ Reserve all guest memory (All locked)

Limit: Unlimited MB

Shares: Normal 163840

16 GB reserved by default

Figure 14

NSX manager is the management plane component and never participates in the data path. If NSX VM goes out of service for a while, it will not impact any traffic in the network. During unavailability of NSX-MGR, new changes will not be permitted though. Customers should take NSX manager VM backup and also configure NSX manager DB backup in the NSX manager UI. NSX manager DB backup can be used to restore the last save NSX configuration in the event of VM failure.

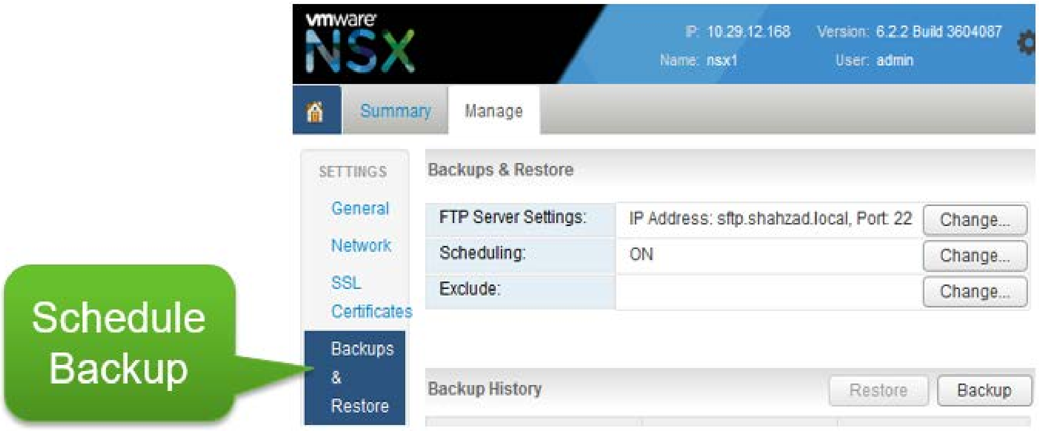

IP: 10.29.12.168 Version: 6.2.2 Build 3604087

Name: nsx1 User: admin

Backups & Restore

FTP Server Settings: IP Address: sftp.shahzad.local, Port: 22

Scheduling: ON

Exclude:

Schedule Backup

Figure 15

NSX Manager “VM Exclusion List”

NSX manager (NSX-MGR) VM exclusion list features makes sure that the VM placed in the list will never get the DFW rules applied to them. The recommendation is to add all the management VMs and most importantly vCenter VM in the list. This is to make sure that customers don’t lock themselves down logging into the management VMs or vCenter VM by accidently configuring a wrong DFW firewall rule. Alternate is to configure some fine grained DFW rules to make sure access to management VMs and vCenter VM is always granted.

Note that NSX VMs like NSX-MGR, DLR Control VM and NSX ESG VMs are automatically part of exclusion list.

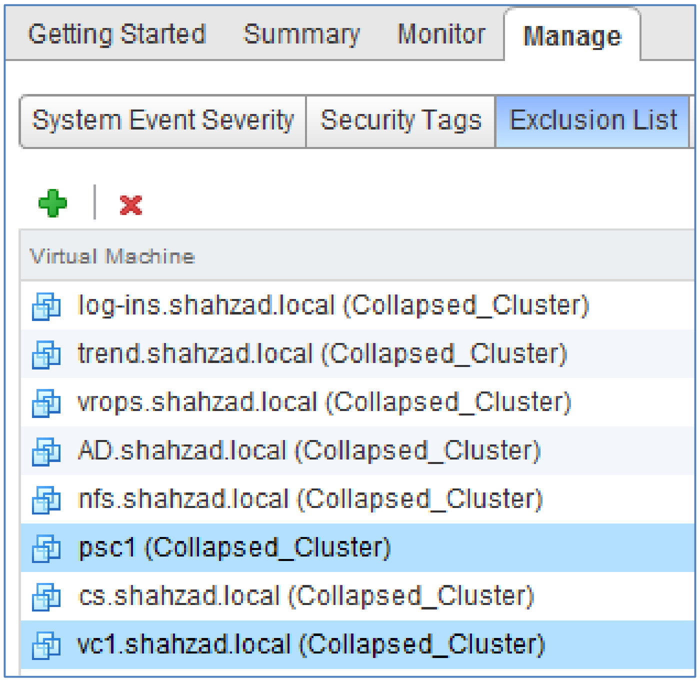

Virtual Machine
log-ins.shahzad.local (Collapsed_Cluster)
trend.shahzad.local (Collapsed_Cluster)
vrops.shahzad.local (Collapsed_Cluster)
AD.shahzad.local (Collapsed_Cluster)
nfs.shahzad.local (Collapsed_Cluster)
psc1 (Collapsed_Cluster)
cs.shahzad.local (Collapsed_Cluster)
vc1.shahzad.local (Collapsed_Cluster)

Figure 16

NSX Controllers

NSX controller is responsible for programming data plane (ESXi hosts) with the logical switch (VXLAN) and distributed routing (DLR) information in the kernel. NSX mandate configuring 3 controller VM in order to provide built-in high availability and load distribution. NSX controller clustering, when used in conjunction with the vSphere high availability and DRS, provides protection in case of failure and make sure the operations are running smoothly.

In order to prevent single host failures, it is recommended to create DRS anti-affinity rules. These rules should be created to make sure each controller VM is residing on a separate host. It is suggested to create "SHOULD" anti-affinity rules. "SHOULD" rules are best effort rules that will put VMs on separate host. In case host is not available, these "SHOULD" rule can also place two VMs on single host. Use of at least 4 host will ensure that cluster can sustain single host failure.

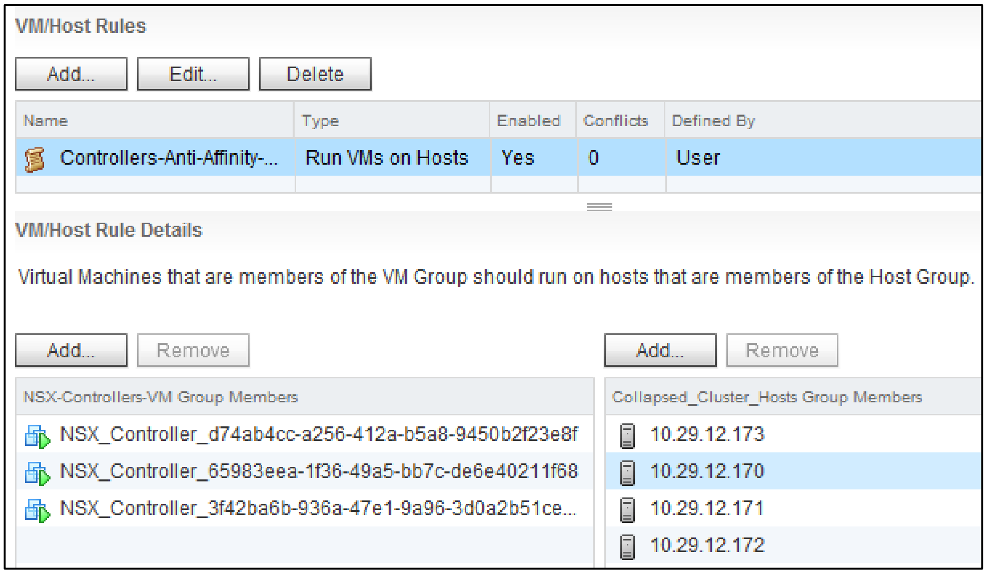

VM/Host Rules

Name	Type	Enabled	Conflicts	Defined By
Controllers-Anti-Affinity-...	Run VMs on Hosts	Yes	0	User

VM/Host Rule Details

Virtual Machines that are members of the VM Group should run on hosts that are members of the Host Group.

NSX-Controllers-VM Group Members	Collapsed_Cluster_Hosts Group Members
NSX_Controller_d74ab4cc-a256-412a-b5a8-9450b2f23e8f	10.29.12.173
NSX_Controller_65983eea-1f36-49a5-bb7c-de6e40211f68	10.29.12.170
NSX_Controller_3f42ba6b-936a-47e1-9a96-3d0a2b51ce...	10.29.12.171
	10.29.12.172

Figure 17 – NSX Controller Affinity Rules

NSX controller VMs are deployed via NSX-manager UI. These VMs are deployed as "locked-down" VM. Which means it is not possible to change the default vCPU and Memory settings for these VMs. These VMs are also, by default, created with memory reservation which cannot be changed via vCenter UI. These setting can be seen in Figure 18.

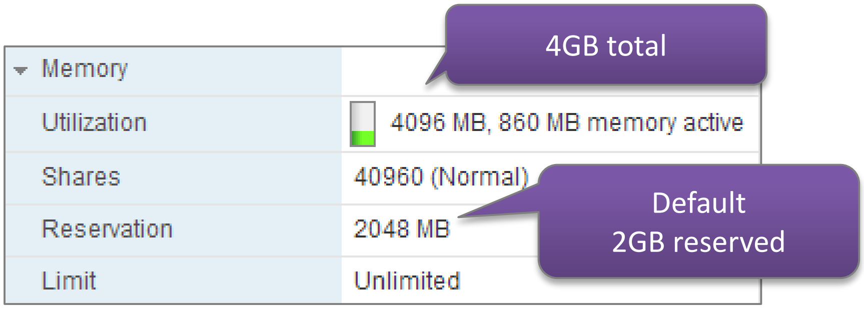

Memory	
Utilization	4096 MB, 860 MB memory active
Shares	40960 (Normal)
Reservation	2048 MB
Limit	Unlimited

Figure 18 – Default NSX Controller VM Settings

In small DC customers can deploy controller VM with confidence due to pre-reserved memory settings. Note that similar to NSX-MGR, NSX controllers are not in the data-path. In case one controller goes down, it will not impact the traffic because other controller VM in the cluster will assume responsibilities for the functions being performed by the outaged VM.

DLR Control VM

This VM is only needed if dynamic routing protocol is required between DLR and ESG. As a good practice, customers are encouraged to use dynamic routing protocol between DLR and ESG. In small data centers, a customer might choose to deploy DLR with the static routing. In that situation, DLR control VM is not needed and some compute capacity can be saved.

Anti-affinity rule is created automatically if DRL control VM is deployed. This is only possible if DRS is enabled on the cluster.

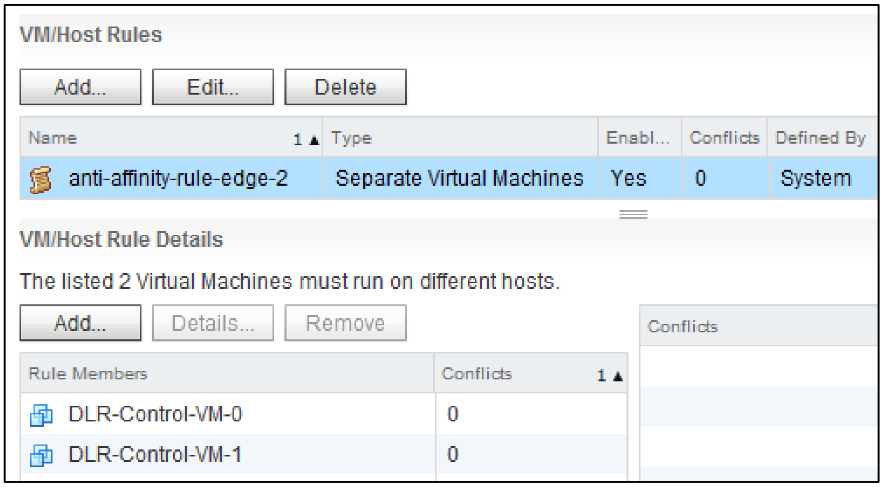

Name 1 ▲	Type	Enabl...	Conflicts	Defined By
anti-affinity-rule-edge-2	Separate Virtual Machines	Yes	0	System

The listed 2 Virtual Machines must run on different hosts.

Rule Members	Conflicts 1 ▲
DLR-Control-VM-0	0
DLR-Control-VM-1	0

Figure 19

These VMs must be deployed in HA mode (active/standby) for production deployment.

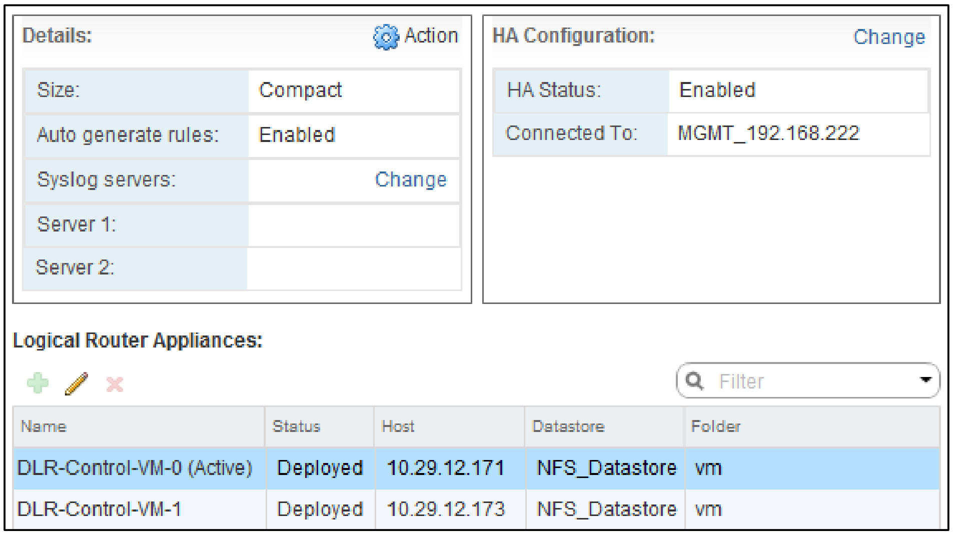

Size:	Compact
Auto generate rules:	Enabled
Syslog servers:	Change
Server 1:	
Server 2:	

HA Status:	Enabled
Connected To:	MGMT_192.168.222

Name	Status	Host	Datastore	Folder
DLR-Control-VM-0 (Active)	Deployed	10.29.12.171	NFS_Datastore	vm
DLR-Control-VM-1	Deployed	10.29.12.173	NFS_Datastore	vm

Figure 20

Notice that these VMs are also locked-down VMs and CPU/MEM modification is disabled.

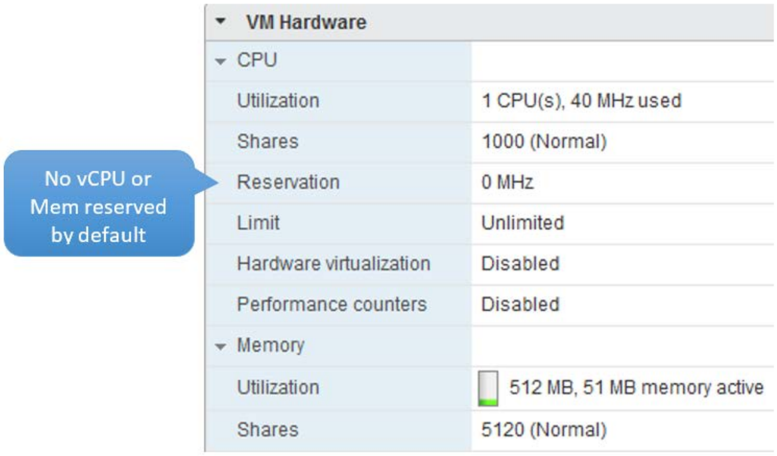

No vCPU or Mem reserved by default

VM Hardware	
CPU	
Utilization	1 CPU(s), 40 MHz used
Shares	1000 (Normal)
Reservation	0 MHz
Limit	Unlimited
Hardware virtualization	Disabled
Performance counters	Disabled
Memory	
Utilization	512 MB, 51 MB memory active
Shares	5120 (Normal)

Figure 21

Edge Services Gateway (ESG)

ESG VM sits at the edge of the NSX virtual network and provides on-ramp/off-ramp capabilities between virtual overlay and physical underlay. ESG VM can be deployed into two different modes

1 - ESG with HA

2 - ESG with ECMP

ESG with HA option will essentially deploy two ESG VMs. One will act as primary and the second will be in standby mode. In small data center ESG with HA is the most preferred option because in majority of the cases, customers would like to also run ESG services like DHCP, NAT, VPN etc. on the same ESG VM. This is not possible with the ESG in ECMP mode.

ESG with ECMP should be deployed when more than 10G north-south bandwidth is required. This is not a very common deployment option in small data centers. In this mode all ESG VMs are active and participate in sending traffic in/out from virtual network to physical underlay. Since all ESG VMs are active in case of ECMP, the traffic might go out from one ESG VM and come back in from another. Due to this fact, it is not allowed to run stateful services on ESG in ECMP mode.

Following decision tree guide customer to decide which option is best for their particular needs and growth forecast.

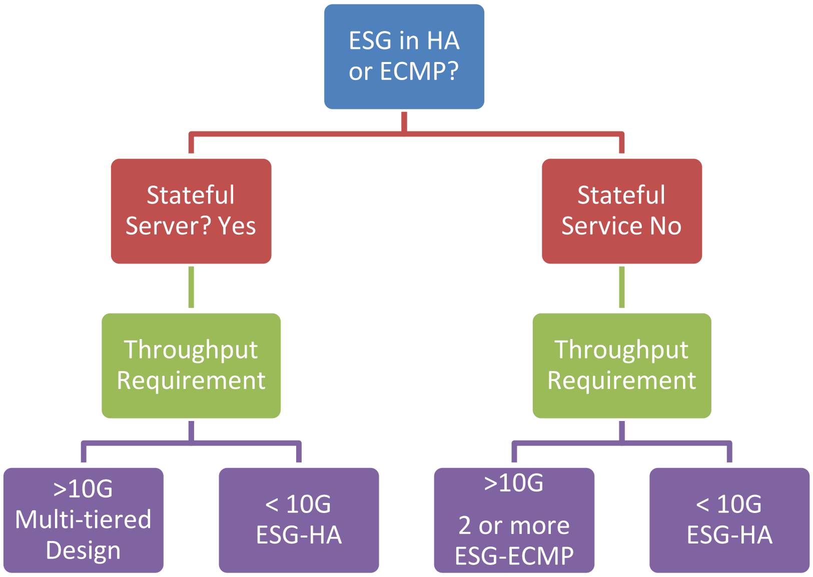

ESG Multi-Tiered Design

In multi-tier ESG design, a customer might deploy ESG in Tier0 and Tier1 fashion. Where Tier0 ESGs are configured in ECMP mode for higher throughput. And Tier1 ESGs, that are close to the actual workload VMs, are configured with HA to provide stateful services.

This design is definitely not very common in small data center. If a design needs higher throughput requirement with stateful services, then it is recommended to follow NSX design guide for medium size or large data center designs.

Note that this document is listing the most commonly asked design choices. Other designs are also possible depending on scale.

General ESG Deployment Considerations For Small Data Centers

ESG VM should be deployed in large form factor. This option is good for majority of the services that a small DC design might require except load balancer. For NSX load balancer, if a design requirement is to have L7 engine turned on then it is highly recommended to deploy ESG in X-large form factor. NSX automatically creates these ESG VMs with the vCPU/MEM reservation that is needed to run these VMs smoothly. Once these are deployed, the UI will not allow to modify these values.

Table 8 shows the ESG VM form factor and its resource requirement.

VM Size	vCPU	Memory(GB)	HD (GB)	Suitable For
Large	2	1	1	All services except L6 LB
X-Large	6	8	2.5	L7 LB

Table 8 – ESG VM Form Factors

ESG Deployed in HA

When ESG is deployed in HA the anti-affinity rules is automatically created if DRS is enabled on the cluster to make sure the active/standby VMs are not hosted on the same ESXi host.

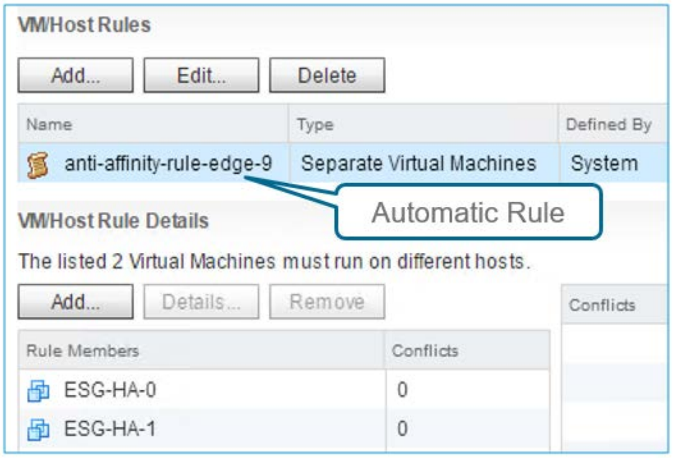

VM/Host Rules

Name	Type	Defined By
anti-affinity-rule-edge-9	Separate Virtual Machines	System

VM/Host Rule Details

The listed 2 Virtual Machines must run on different hosts.

Rule Members	Conflicts
ESG-HA-0	0
ESG-HA-1	0

Conflicts

Figure 22

If the DRS is not enabled on the cluster, then customer should make sure that active and standby VMs are not residing on the same ESXi host at any time. If active and standby ESG are resident on same ESXi host and if that ESXi host loses connectivity or goes down for some reason, then it may cause service disruption. Recommended design configuration with ESG in HA is shown in Figure 23.

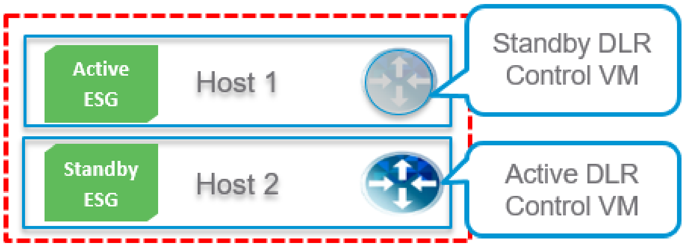

Figure 23

ESG Deployed in ECMP

With ESG in EMCP mode, all ESG VMs are active. They will actively participate in sending traffic. If one ECMP ESG VM goes down, the traffic will still flow through other active ECMP ESG VMs. With ESG in ECMP mode, make sure to disable ESG HA option in the NSXI UI. Enabling both ECMP and ESG HA option will create one extra and unnecessary standby VM per ESG VM. Another aspect is that these Active ECMP VMs should not be deployed on the same ESXi host with the Active DLR control VM. Make sure to create anti-affinity rule so that the distribution and placement of VMs is according to recommended guidelines.

Recommended design configuration is shown in Figure 24.

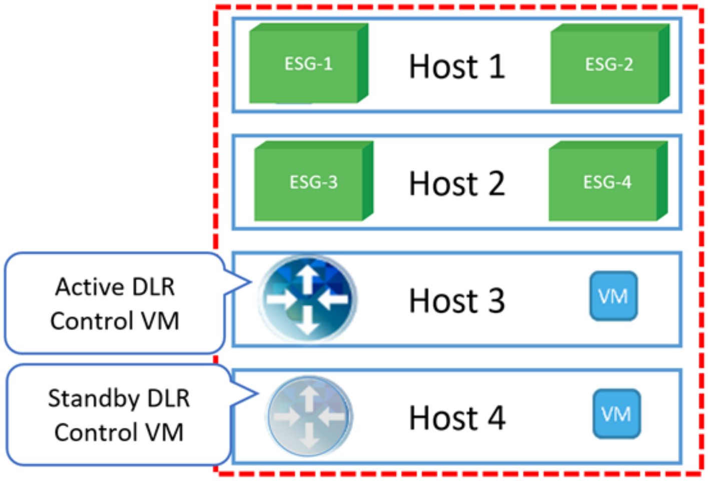

Figure 24

VDS Considerations

VDS is the only supported virtual switch for NSX deployment. vSphere Standard Switch or any other 3rd party vendor is not supported with NSX. VDS does requires vSphere Enterprise Plus license. This requirement is lifted if NSX license is purchased. With NSX license, any vSphere license (vSphere 5.5 U3 or 6.0+) will support VDS. Essentially VDS comes free with NSX license.

In small data center, since there is only one cluster having all the VMs, it is recommended to keep design simple and use single VDS. Multiple port groups like vMotion, management should be created as per the NDX design guide. VTEP port groups and VM kernel port(s) created automatically based on the nic teaming selection.

Recommended VTEP vmknic teaming policy is Route Based on Originating Port (Source-ID). This option provides multiple VXLAN paths for traffic leaving from ESXi host. It will create multiple VTEP ports per host in order to achieve multipathing. This might not be desirable for majority of small DC customers as they might want to keep the design simple from configuration and troubleshooting point of view. If single VTEP per ESXi host is desirable, one can choose NIC fail-over option as teaming policy.

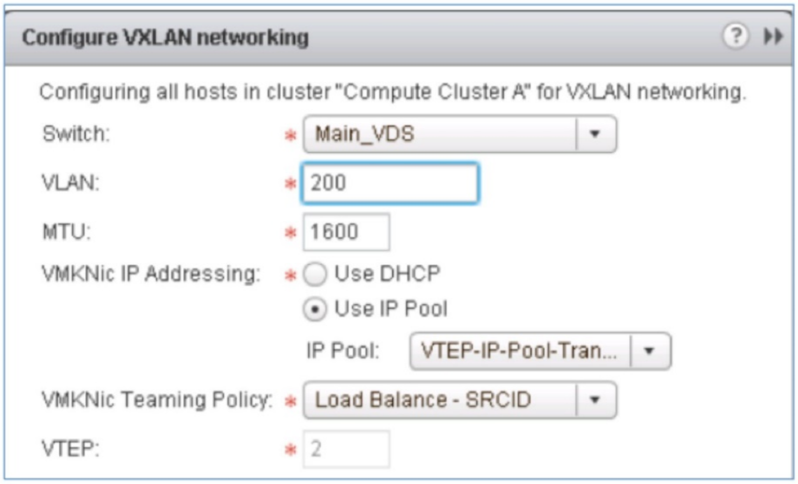

Configure VXLAN networking

Configuring all hosts in cluster "Compute Cluster A" for VXLAN networking.

Switch: * Main_VDS

VLAN: * 200

MTU: * 1600

VMKNic IP Addressing: * ○ Use DHCP

● Use IP Pool

IP Pool: VTEP-IP-Pool-Tran...

VMKNic Teaming Policy: * Load Balance - SRCID

VTEP: * 2

Growing and Expanding NSX Deployment

Growing and expanding NSX is not just one dimensional decision and option. With NSX, it is important to notice that this expansion can be in different areas and directions and not just growth in terms of compute addition. NSX is all in software, it is highly flexible and versatile and provides various options to grow.

Following figure provides some of the growth possibilities and options:

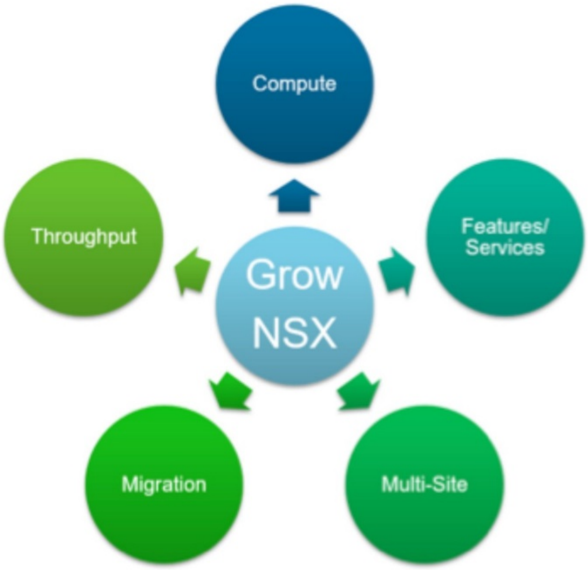

Customers can easily add additional compute capacity to existing or new cluster based on their business needs. New features can be added (like L2 VPN or Load Balancer etc.) to existing deployment without re-architecting the overall design. Customers can also add more bandwidth or throughout to the existing deployment by simply adding more physical links or by adding more ECMP ESG in scale out fashion.

There are customers who have deployed NSX in small data center but quickly seen the potential and advantages of the solution and are growing NSX beyond single site and have deployed in cross-VC with multiple NSX.

Example Growth Scenario

There could be many options and scenarios that can be shown and discussed here but due to the length of the document, only one scenario is discussed here. A customer can deploy NSX Security Focused deployment model as shown in figure below and can easily expand it to NSX full stack deployment model.

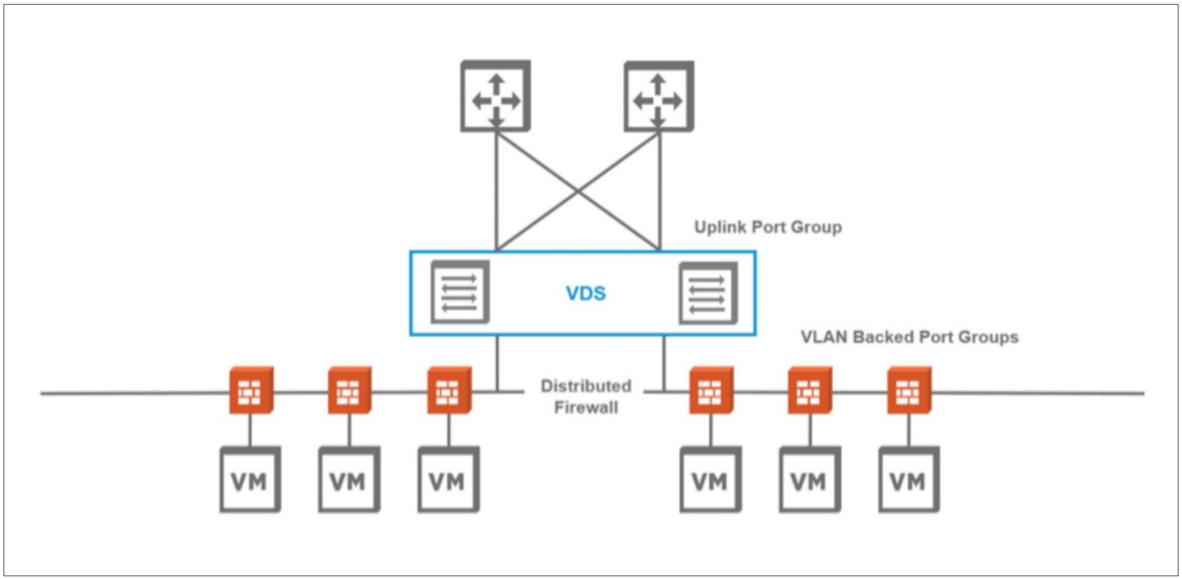

From a very high level, expanding from Security Focused deployment mode to Full Stack will require following

- MTU should be 1600 or more to implement VXLAN based overlay
- IP addresses/subnet will be needed for VTEP and other components like VXLAN Transit logical switch and DRL control VM etc.
- VTEP must be configured on ESXi hosts. Number of VTEP will depend on the NIC teaming type used on the VDS
- If dynamic routing is needed in the virtual network, it would require deployment of DRL control VM in HA mode
- NSX edges (ESG) VMs will be deployed in the either ECMP or HA mode depending on business requirement. This action does require changes on the physical network and peering NSX ESG with the physical routing domain and devices
- The default gateway will be moved to NSX DRL. This does not require re IP on workload VM. This action does require changes on the physical network
- As part of the expansion, 3rd party services can be added for features like Guest Introspection and Network Introspection

Details of above mentioned steps can be found in NSX for vSphere Brownfield Design Guide.

A high level architectural diagram is shown in figure below

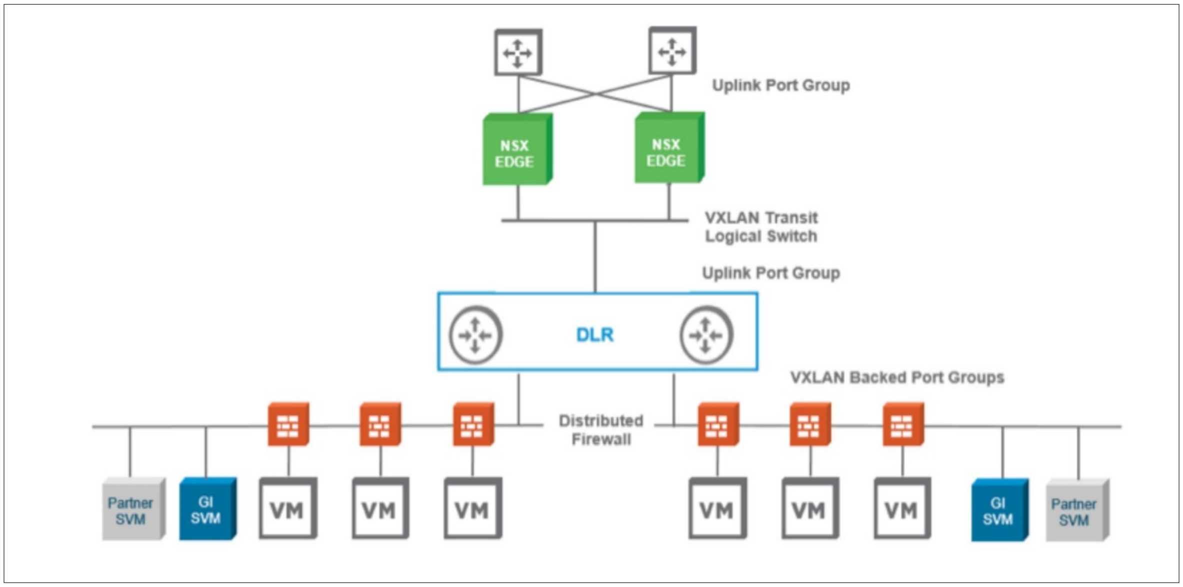

Capacity Planning and Monitoring

It is important to understand that like any other networking design there are consolidation and oversubscription ratios that should be monitored on the ongoing basis. Proper capacity planning and monitoring would avoid the risk of starting with small DC with small workload footprint but then organically growing while not keeping an eye on original scope of Small DC design.

If a customer design NSX in small DC for a certain number of workload, then it is imperative to add resources as the deployment grow. This could be done easily with the tools like VMware vRealize Automation or vRealize Network Insight.

Summary

NSX has been successfully deployed for large number of customers with various datacenter sizes including small data centers. There are also cases where some large enterprises deployed NSX with small footprint in the beginning. In almost all the cases, those deployments then grew to a medium or larger size NSX deployments.

Regardless of the customer size, NSX and vSphere together provides options and features that will make sure that customer deployments are successful and they can get the full benefit of network virtualization technology based on the business problem they are trying to solve.

Appendix A

List of relevant references to design, deploy and manage NSX.

Planning and Design

- NSX for vSphere Network Virtualization Design Guide
- NSX for vSphere Cross-VC Design Guide
- NSX for vSphere Brownfield Design Guide
- Software Defined Data Center (SDDC) VMware Validated Design Guide (VVD)

Implementation

- NSX for vSphere Installation Guides

Operate and Manage

- NSX for vSphere Operations Guide

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