

Ranjit Singh Thakurratan

Learning VMware NSX

Second Edition

Next-generation network administration skills unveiled



www.hellodigi.ir

Packt >

Learning VMware NSX

Second Edition

Next-generation network administration skills unveiled

Ranjit Singh Thakurratan



BIRMINGHAM - MUMBAI

Learning VMware NSX

Second Edition

Copyright © 2017 Packt Publishing

All rights reserved. No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, without the prior written permission of the publisher, except in the case of brief quotations embedded in critical articles or reviews.

Every effort has been made in the preparation of this book to ensure the accuracy of the information presented. However, the information contained in this book is sold without warranty, either express or implied. Neither the author, nor Packt Publishing, and its dealers and distributors will be held liable for any damages caused or alleged to be caused directly or indirectly by this book.

Packt Publishing has endeavored to provide trademark information about all of the companies and products mentioned in this book by the appropriate use of capitals. However, Packt Publishing cannot guarantee the accuracy of this information.

First published: November 2015

Second edition: August 2017

Production reference: 1210817

Published by Packt Publishing Ltd.
Livery Place
35 Livery Street
Birmingham

B3 2PB, UK.

ISBN 978-1-78839-898-5

www.packtpub.com

Credits

<p>Author Ranjit Singh Thakurratan</p>	<p>Copy Editor Safis Editing</p>
<p>Reviewer Renjith Menon</p>	<p>Project Coordinator Kinjal Bari</p>
<p>Commissioning Editor Vijin Boricha</p>	<p>Proofreader Safis Editing</p>

Acquisition Editor

Prachi Bisht

Indexer

Mariammal Chettiyar

Content Development Editor

Trusha Shriyan

Graphics

Kirk D'Penha

Technical Editor

Sayali Thanekar

Production Coordinator

Shantanu Zagade

About the Author

Ranjit Singh Thakurratan is a four year VMware vExpert (2013-2017) and works as a principal chief architect at DellEMC. Ranjit holds a master's degree in information technology—infrastructure assurance and an engineering degree in computer science, and has over ten years of hands-on IT experience. He has presented at numerous VMUG UserCon conferences held at Boston, Washington DC, New York, Denver, and Dallas. He also runs a technology blog RJApproves.com and can be reached via his Twitter handle [@RJAPPROVES](https://twitter.com/RJAPPROVES). Apart from technology, Ranjit is also interested in astrophysics, animal welfare, and open source projects.

About the Reviewer

Renjith Menon is an IT systems engineer with over 12+ years of experience working in an environment covering IT systems administration, systems operations, systems health monitoring, systems upgrades, and deploying new software and services. He possesses diversified experience covering implementation of new systems and troubleshooting of software components along with maintenance and configuring of systems.

At work, on a daily basis, he deals with virtualization technologies, backup and replication technologies, most Microsoft server roles, and much more. He likes tech blogging and learning new technologies. He started his blogging journey back in 2007, and recently, he also started his personal blog.

You can follow him on Twitter at [@vcrenjith](#).

www.PacktPub.com

For support files and downloads related to your book, please visit www.PacktPub.com.

Did you know that Packt offers eBook versions of every book published, with PDF and ePub files available? You can upgrade to the eBook version at www.PacktPub.com and as a print book customer, you are entitled to a discount on the eBook copy. Get in touch with us at service@packtpub.com for more details.

At www.PacktPub.com, you can also read a collection of free technical articles, sign up for a range of free newsletters and receive exclusive discounts and offers on Packt books and eBooks.



<https://www.packtpub.com/mapt>

Get the most in-demand software skills with Mapt. Mapt gives you full access to all Packt books and video courses, as well as industry-leading tools to help you plan your personal development and advance your career.

Why subscribe?

- Fully searchable across every book published by Packt
- Copy and paste, print, and bookmark content
- On demand and accessible via a web browser

Customer Feedback

Thanks for purchasing this Packt book. At Packt, quality is at the heart of our editorial process. To help us improve, please leave us an honest review on this book's Amazon page at <https://www.amazon.com/dp/178839898X>.

If you'd like to join our team of regular reviewers, you can email us at customerreviews@packtpub.com. We award our regular reviewers with free eBooks and videos in exchange for their valuable feedback. Help us be relentless in improving our products!

This book would not have been possible without the endless support of my dear wife Monica, our pup Shyla, my lovely parents, and my brothers.

Table of Contents

Preface

- What this book covers
- What you need for this book
- Who this book is for
- Conventions
- Reader feedback
- Customer support
 - Downloading the color images of this book
 - Errata
 - Piracy
 - Questions

1. Introduction to Network Virtualization

- Introducing network virtualization
- Concepts of network virtualization
- Introducing the NSX-V network virtualization platform
- NSX features and services
- NSX configuration maximums
- Summary

2. NSX Core Components

- Introduction to NSX core components
- NSX manager
- NSX controller clusters
- VXLAN architecture overview
 - A sample packet flow
- Transport zones
- NSX Edge Services Gateway
- Distributed firewalls
- Cross-vCenter NSX
- Summary

3. NSX Installation and Configuration

- Preparing your environment
- Downloading and deploying NSX Manager
- Overview of the NSX Manager interface
- Configuring NSX Manager
- Managing NSX using the vSphere web client
- Deploying the control plane (Controller Virtual Machines)
- Deploying the data plane

Summary

4. NSX Functional Services

Primary and secondary NSX managers

Benefits of cross-vCenter NSX

Configuring VXLAN

Assigning a segment ID pool

Transport zones

Logical switching

L2 bridges

Deploying an NSX Edge logical router

Edge services gateway

Logical firewalls

Summary

5. Edge Services Gateway

DNS and DHCP services

DHCP service

DNS service

Routing

Configuring OSPF on Edge services gateway

Configuring logical distributed router OSPF

Configuring BGP

Configuring route redistribution

Logical Edge load balancers

Virtual private networks

SSL VPN-Plus

SSL VPN-Plus network access mode

IPSEC VPN

L2 VPN

More Edge services and configurations

Adding a sub-interface

Force sync NSX Edge with NSX Manager

Configuring remote syslog servers

Redeploying an NSX Edge

Summary

6. Service Composer

Service Composer

Security groups

Security policies

Security group and security policy mapping

Network extensibility

Summary

7. Monitoring

Endpoint Monitoring

Flow monitoring

Traceflow

Summary

8. Managing NSX

NSX Manager settings

 Date and time settings

 Syslog servers

 DNS servers

 Technical support logs

 SSL certificates

Backup and restore

 NSX Manager backup

NSX Manager domain registration

 Configuring SNMP traps

Controller cluster operations

Summary

9. Conclusion

Preface

Networking is one of the most important and critical components of any IT architecture. Architects always want to lay the foundation and solve networking before compute and storage is looked at. This is why I have always felt the need to understand networking better and understand how things work.

Network virtualization technology gave me a behind the scenes, hands-on look at how networking works and the concepts that made it happen. The ability to deploy virtual appliances such as switches, load balancers, and routers and examine their functionality was very appealing. The pace at which I learned network virtualization gave me the confidence to grow in a field that I thought I lacked significant knowledge in. Throughout my journey, I understood that network virtualization not only made me think like a network engineer but also made me apply networking concepts very creatively to a virtualized environment, and this made it all the more appealing. What was complex before now became extremely simple.

The journey started in 2015 when I went around presenting at multiple VMUG UserCon sessions all around the country. My presentation was Getting started with VMware NSX - basics and best practices. My aim here was not to talk about what NSX can do but to talk about how easy it was to get started and to squash some common misconceptions about NSX. I wasn't sure if this was a topic worth talking about and wasn't expecting a large crowd. My time slot to present was right after lunch, which wasn't very appealing.

I was quickly proved wrong. In each and every city I presented, all of my sessions were completely full. People were eager to see how to get started with NSX and in fact, preferred this presentation over any NSX presentations happening at the same conference. This is when the first edition of Learning VMware NSX was born. Following the tremendous reception and feedback, we decided to have this second edition to ensure continuity and address changes in the network virtualization technology.

The aim of the book is to connect to that day-to-day administrator and that network engineer and make it easy for them to understand NSX. The book explains the basics and covers the deployment of various features of network virtualization in simple, clear language and with screenshots to allow you to visualize the workflow as you read.

I hope you enjoy working with this second of edition Learning VMware NSX, and that it

helps you learn how to use and understand NSX and network virtualization. We are constantly looking for feedback and advice, so feel free to reach out to us by all means necessary.

What this book covers

[Chapter 1](#), Introduction to Network Virtualization, gets you started with an introduction to network virtualization and an overview of its concepts.

[Chapter 2](#), NSX Core Concepts, talks about all the different components of NSX and how they work together.

[Chapter 3](#), NSX Installation and Configuration, covers deploying and configuring NSX.

[Chapter 4](#), NSX Functional Services, discusses the deployment and configuration of different NSX services such as logical switching, L2 bridging, and Edge gateway services.

[Chapter 5](#), Edge Services Gateway, goes deeper into the services offered by Edge gateway and looks at deploying and configuring them.

[Chapter 6](#), Service Composer, discusses different NSX security policies, because one of the most important capabilities of NSX is its security features.

[Chapter 7](#), Monitoring, looks at enabling the monitoring of our environment using NSX.

[Chapter 8](#), Managing NSX, talks about NSX administrative tasks such as backup and restore along with NSX manager settings.

[Chapter 9](#), Conclusion, concludes the second edition of the Learning VMware NSX series and provides additional reference links and author contact information.

What you need for this book

Although you can dive right into this book, I recommend setting up a modest home lab of three servers running VMware ESXi and vCenter. You are also encouraged to spend time exploring the hands-on labs offered for free by VMware. The specific NSX labs that will help you greatly are HOL-1703- SDC-1 and HOL-1703- USE-2. The labs help you get started with NSX without having to worry about the intricacies of having to set it up.

You can get to hands-on labs by visiting <http://labs.hol.vmware.com> and searching for the two labs I mentioned previously.

Who this book is for

The book is for anyone who is interested in learning more about software-defined network virtualization tools. System administrators, network administrators, solution engineers, sales engineers, and solution architects are some of those who will find this book very educational.

Conventions

In this book, you will find a number of text styles that distinguish between different kinds of information. Here are some examples of these styles and an explanation of their meaning. Code words in text, database table names, folder names, filenames, file extensions, pathnames, dummy URLs, user input, and Twitter handles are shown as follows: "Log in to the NSX manager with the username `admin` and the password that was set during deployment time."

A block of code is set as follows:

```
acl is_foo hdr_dom(host) -i foo
acl is_bar hdr_dom(host) -i bar
use_backend pool_1 if is_foo
use_backend pool_2 if is_bar
```

Any command-line input or output is written as follows:

```
[root@host:~] esxcli software vib list | grep esx
```

New terms and **important words** are shown in bold. Words that you see on the screen, for example, in menus or dialog boxes, appear in the text like this: "You will see the Deploy OVF Template screen."



Warnings or important notes appear like this.



Tips and tricks appear like this.

Reader feedback

Feedback from our readers is always welcome. Let us know what you think about this book-what you liked or disliked. Reader feedback is important for us as it helps us develop titles that you will really get the most out of. To send us general feedback, simply email feedback@packtpub.com, and mention the book's title in the subject of your message. If there is a topic that you have expertise in and you are interested in either writing or contributing to a book, see our author guide at www.packtpub.com/authors.

Customer support

Now that you are the proud owner of a Packt book, we have a number of things to help you to get the most from your purchase.

Downloading the color images of this book

We also provide you with a PDF file that has color images of the screenshots/diagrams used in this book. The color images will help you better understand the changes in the output. You can download this file from https://www.packtpub.com/sites/default/files/downloads/LearningVMwareNSXSecondEdition_ColorImages.pdf.

Errata

Although we have taken every care to ensure the accuracy of our content, mistakes do happen. If you find a mistake in one of our books-maybe a mistake in the text or the code-we would be grateful if you could report this to us. By doing so, you can save other readers from frustration and help us improve subsequent versions of this book. If you find any errata, please report them by visiting <http://www.packtpub.com/submit-errata>, selecting your book, clicking on the Errata Submission Form link, and entering the details of your errata. Once your errata are verified, your submission will be accepted and the errata will be uploaded to our website or added to any list of existing errata under the Errata section of that title. To view the previously submitted errata, go to <https://www.packtpub.com/books/content/support> and enter the name of the book in the search field. The required information will appear under the Errata section.

Piracy

Piracy of copyrighted material on the internet is an ongoing problem across all media. At Packt, we take the protection of our copyright and licenses very seriously. If you come across any illegal copies of our works in any form on the internet, please provide us with the location address or website name immediately so that we can pursue a remedy. Please contact us at copyright@packtpub.com with a link to the suspected pirated material. We appreciate your help in protecting our authors and our ability to bring you valuable content.

Questions

If you have a problem with any aspect of this book, you can contact us at questions@packtpub.com, and we will do our best to address the problem.

Introduction to Network Virtualization

This chapter begins with a brief introduction to network virtualization, followed by an overview of its concepts. We then introduce VMware's NSX-V network virtualization solution that allows you to deploy and manage your own software-defined networking stack. We will go over all the features and services of NSX, followed by its configuration maximums. By the end of this chapter, you will have a thorough understanding of the concepts of network virtualization, and NSX-V as a network virtualization solution.

In this chapter, we will cover:

- Introducing network virtualization
- Concepts of network virtualization
- Introducing the NSX-V network virtualization platform
- NSX features and services
- NSX configuration maximums
- Summary

Introducing network virtualization

Today's datacenter demands are a paradigm shift from what they were a decade ago. As the cloud consumption model is being rapidly adopted across the industry, the need for on-demand provisioning of compute, storage, and networking resources is greater than ever. One of the biggest contributing factors to enable the cloud consumption model is **server virtualization**.

Server virtualization has enabled fast consumption of compute resources along with add-on functionality and services. Snapshots, clones, and templates are all now easier than ever with server virtualization.

If you have worked in a datacenter, you would agree that networking is always challenging to work with. Once the networking design is established, any changes that need to be made are always challenging because of a lack of flexibility due to increasing complexity and demands on the environment. While compute and storage have rapidly improved in their speed of deployment and consumption, networking continues to remain a challenge in today's environments, where simple tasks such as creating a new VLAN are becoming increasingly complex and time consuming.



A metaphor: Today's networking is similar to building roads and highways in a city. Once you have the highways and roads established, it is not easy to expand them, or simply remove and replace them, without affecting traffic. You always have to think ahead and build to facilitate future growth and flexibility. Similarly, traditional networks in a datacenter have to be built to handle future growth and should be flexible enough to allow for changes as they happen.

Network virtualization is the virtualization of network resources using software and networking hardware that enables faster provisioning and deployment of networking resources. Network virtualization lays the foundation for software-defined networking, which allows instant deployment of services to be offered to the consumers. Services such as Edge gateways, VPN, DHCP, DNS, and load balancers can be instantly provisioned and deployed because of the software aspect of network virtualization. The networking hardware allows for physical connectivity, while the software is where all the network logic resides allowing for a feature-rich network service offering.

Network virtualization allows for consumption of simplified logical networking devices

and services that are completely abstracted from the complexities of the underlying physical network. Lastly, network virtualization is key for a **software-defined data center (SDDC)**.

Concepts of network virtualization

Now that we have defined what network virtualization is about, let's go over some of the key concepts of network virtualization and software-defined networking:

- **Decoupling:** An important concept of network virtualization is the decoupling of software and the networking hardware. The software works independently of the networking hardware that physically interconnects the infrastructure. Any networking hardware that can inter-op with the software is always going to enhance the functionality, but it is not necessary. Remember that your throughput on the wire will be always limited by your network hardware performance.
- **Control plane:** The decoupling of software and networking hardware allows you to control your network better because all the logic resides in the software. This control aspect of your network is called the control plane. The control plane provides the means to configure, monitor, troubleshoot, and also allow automation against the network.
- **Data plane:** The networking hardware forms the data plane where all the data is forwarded from source to destination. The management of data resides in the control plane; however, the data plane consists of all the networking hardware whose primary function is to forward traffic over the wire from source to destination. The data plane holds all the forwarding tables that are constantly updated by the control plane. This also prevents any traffic interruptions if there is a loss of the control plane, because the networking hardware, which constitutes the data plane, will continue to function without interruptions.
- **Application Programming Interface (API):** The API is one of the important aspects of a virtualized network and allows for true software-defined networking by instantly changing the network behavior. With the API, you can now instantly deploy rich network services in your existing network. Network services such as Edge gateway, VPN, Firewall, and load balancers can all be deployed on the fly by means of an API.

Introducing the NSX-V network virtualization platform

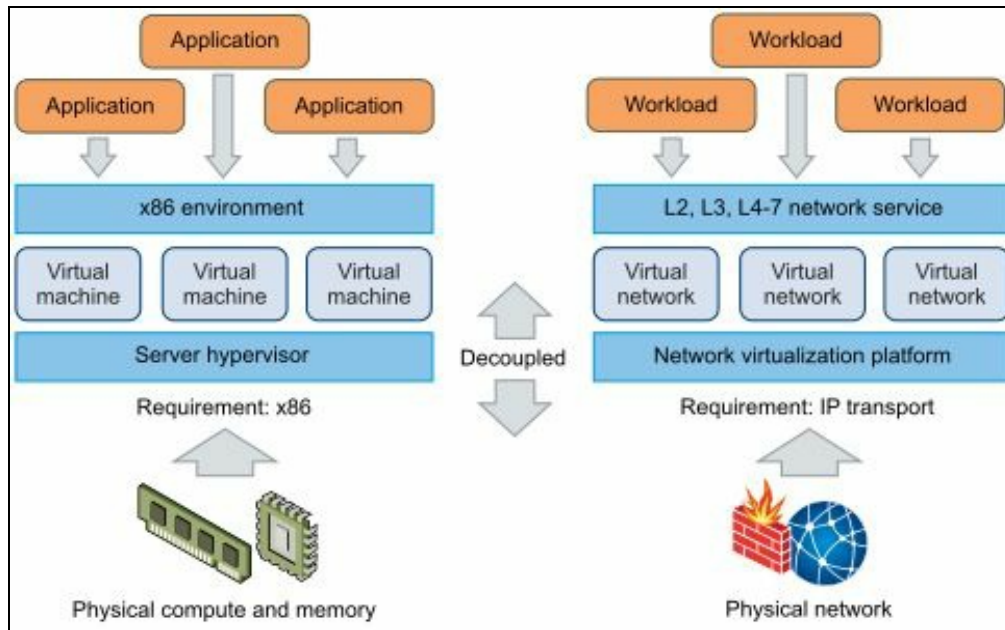
VMware NSX-V is a network virtualization platform that allows for software-defined networks and is a critical component of software-defined datacenter architecture. VMware's NSX-V software abstracts the underlying physical network by introducing a software layer that makes it easy to consume network resources by creating multiple virtual networks. NSX-V also allows for deploying multiple logical network services on top of the abstracted layer.



VMware acquired NSX from Nicira in July, 2012. Nicira's NSX was primarily being used for network virtualization in a Xen-based hypervisor.

VMware now has two flavors of NSX: NSX-V, and NSX-MH. NSX-V is NSX for a VMware-based hypervisor while **NSX-Multi Hypervisor (NSX-MH)** is for OpenStack environments. The two versions have many similarities but also are dissimilar in some aspects. This book focuses on the **NSX-VMware (NSX-V)** version of NSX only. NSX-V will be referred to as NSX for the rest of the book.

The following figure represents the software abstraction of a physical network and networking hardware by NSX. This is synonymous with how the VMware vSphere hypervisor achieves software abstraction of CPU, memory, and storage, making it possible for the creation of multiple virtual machines:



Just as the vSphere hypervisor allows you to create, delete, snapshot, and monitor a virtual machine, NSX allows you to programmatically create, delete, snapshot, and monitor a virtual network. NSX can be deployed on your current physical network infrastructure, and does not require you to upgrade your existing infrastructure.

Lastly, NSX deployment is non-disruptive to your existing network and traffic. It can seamlessly be deployed on top of your existing infrastructure, and consumption of its services can take place in conjunction with your traditional network.

NSX features and services

Before we get started with NSX, it is important to understand some of its features and services.



NSX 6.2 is the current NSX version as of this writing.

Some NSX features are listed as follows. We will discuss these features in great detail in the following chapters:

- **Logical switching:** NSX allows the ability to create L2 and L3 logical switching that enables workload isolation and separation of IP address space between logical networks. NSX can create logical broadcast domains in the virtual space that prevent the need to create any logical networks on the physical switches. This means you are no longer limited to 4096 physical broadcast domains (VLANs).
- **NSX gateway services:** The Edge gateway services interconnect your logical networks with your physical networks. This means a virtual machine connected to a logical network can send and receive traffic directly to your physical network through the gateway.
- **Logical routing:** Multiple virtual broadcast domains (logical networks) can be created using NSX. As multiple virtual machines subscribe to these domains, it becomes important to be able to route traffic from one logical switch to another. Logical routing helps achieve this by routing traffic between logical switches, or even between a logical switch and public networks. Logical routing can be extended to perform east-west routing that saves unnecessary network hops, increasing network efficiency. Logical routers can also provide north-south connectivity allowing access to workloads living in the physical networks. Logical routers also help avoid hairpinning of traffic, thereby increasing network efficiency.



East-west traffic is traffic between virtual machines within a datacenter. In the current context, this typically will be traffic between logical switches in a VMware environment.

North-south traffic is traffic moving in and out of your datacenter. This is any traffic that either enters your datacenter or leaves your datacenter.

- **Logical firewall:** NSX allows you the option of a distributed logical firewall or an Edge firewall for use within your software-defined networking architecture. A distributed logical firewall allows you to build rules based on attributes that include not just IP addresses and VLANs, but also virtual machine names and vCenter objects. The Edge gateway features a firewall service that can be used to impose security and access restrictions on north-south traffic.
- **Extensibility:** There are third-party VMware partner solutions to integrate directly into the NSX platform that allow a vendor choice in multiple service offerings. There are many VMware partners who offer solutions such as traffic monitoring, IDS, and application firewall services that can integrate directly into NSX. This enhances management and end user experience by having one management system to work with.

The features listed earlier enable NSX to offer a wide variety of services that can be consumed in your infrastructure. These services can be deployed and configured by the NSX API as well. Some of the NSX services are listed as follows:

- **Load balancer:** NSX Edge offers a variety of services and the logical load balancer is one of them. The logical load balancer distributes incoming requests among multiple servers to allow for load distribution while abstracting this functionality from end users. The logical load balancer can also be used as a **high availability (HA)** mechanism to ensure your application has the most uptime.
- **Virtual private networks (VPN):** The NSX Edge offers the VPN service that allows you to provision secure encrypted connectivity for end users to your applications and workloads. Edge VPN service offers SSL-VPN plus it allows for user access and IPSEC site-to-site connectivity, which enables two sites to be interconnected securely.
- **Dynamic Host Configuration Protocol (DHCP):** NSX Edge offers DHCP services that allow IP address pooling, and also static IP assignments. An administrator can now rely on the DHCP service to manage all IP addresses in your environment, rather than having to maintain a separate DHCP service. The DHCP service can also relay DHCP requests to your existing DHCP server as well. The NSX Edge DHCP service can relay any DHCP requests generated from your virtual machines to a pre-existing physical or virtual DHCP server, without any interruptions.
- **Domain name system (DNS):** NSX Edge offers a DNS relay service that can relay any DNS requests to an external DNS server.
- **Service composer:** The service composer allows you to allocate network and multiple security services to security groups. Virtual machines that are part of these

security groups are automatically allocated the services.

- **Data security:** NSX data security provides visibility into sensitive data, ensures data protection, and reports back on any compliance violations. A data security scan on designated virtual machines allows NSX to analyze and report back on any violations based on the security policy that applies to these virtual machines.

Other NSX features include cross-vCenter networking and security, which allow you to manage multiple vCenter NSX environments using a primary NSX manager. This not only allows centralized management, but also extends one or more services and features across multiple vCenter environments. We will talk more about cross vCenter networking in the upcoming chapters.

NSX configuration maximums

Let's have a look at what the NSX configuration maximums are. VMware has not published an official document, so the following limits listed were gathered by reviewing NSX documentation and online research. Some websites that contributed include www.vmguru.com.

Some of these limits are hard limits while most of them are soft limits, beyond which VMware does not support such configurations. For example, if you exceed the number of concurrent connections per Edge gateway, it will affect your gateway's performance, but won't cause it to halt or reject new connections. The hard limit verses soft limit documentation is not explicitly published, but VMware NSX support can clarify if needed. The chances are that you will scale out your environment before reaching these maximums.

The maximums for NSX follow.



NSX 6.2 is the current NSX version as of this writing. Configuration maximums can differ based software release. Always refer to the most up-to-date documentation to ensure accuracy.

The following table shows the limits for **NSX – vCenter Maximums**:

Description	Limit
vCenters	1
NSX Managers	1
DRS clusters	12
NSX controllers	3

Hosts per cluster	32
Hosts per Transport Zone	256

A **Transport Zone** defines the scope of a logical switch and can span one or more vSphere clusters. We will discuss this in greater depth in the upcoming chapters.

The following table shows the limits for **Switching Maximums**:

Description	Limit
Logical switches	10,000
Logical switch ports	50,000
Bridges per distributed logical router	500

The following table shows the limits for **Distributed Logical Firewall Maximums**:

Description	Limit
Rules per NSX Manager	100,000
Rules per VM	1,000
Rules per host	10,000
Concurrent connections per host	2,000,000

Security groups per NSX Manager	10,000
---------------------------------	--------

The following table shows the limits for **Distributed Logical Router (DLR) Maximums**:

Description	Limit
DLRs per host	1,000
DLR per NSX Manager	1,200
Interfaces per DLR	999
Uplink interfaces per DLR	8
Active routes per DLR	2,000
Active routes per NSX Manager	12,000
OSPF adjacencies per DLR	10
BGP peers per DLR	10



Open Shortest Path First (OSPF) and Border Gateway Protocol (BGP) are routing protocols.

The following table shows the limits for **NSX Edge Services Gateway (ESG) Maximums**:

--	--

Description	Limit
Total number of Edge service gateways per NSX Manager	2,000
Interfaces per ESG (internal, uplink or trunk)	10
Sub-interfaces on a trunk	200
NAT rules per ESG	2,000
Static routes per ESG	2,048

The following table shows the limits for **Edge Services Gateway Compact Maximums**:

Description	Limit
OSPF routes per ESG	20,000
OSPF adjacencies per ESG	10
BGP peers per ESG	10
BGP routes per ESG	20,000
Total routes per ESG	20,000
Concurrent connections per ESG	64,000

--	--

The following table shows the limits for **Edge Services Gateway Large Maximums**:

Description	Limit
OSPF routes per ESG	50,000
OSPF adjacencies per ESG	20
BGP peers per ESG	20
BGP routes per ESG	50,000
Total routes per ESG	50,000
Concurrent connections per ESG	1,000,000

The following table shows the limits for **Edge Services Gateway X-Large Maximums**:

Description	Limit
OSPF routes per ESG	100,000
OSPF adjacencies per ESG	40
BGP peers per ESG	50

BGP routes per ESG	250,000
Total routes per ESG	250,000
Concurrent connections per ESG	1,000,000

The following table shows the limits for **Edge Services Gateway Quad-Large Maximums**:

Description	Limit
OSPF routes per ESG	100,000
OSPF adjacencies per ESG	40
BGP peers per ESG	50
BGP routes per ESG	250,000
Total routes per ESG	250,000
Concurrent connections per ESG	1,000,000

The following table shows the limits for **Edge Services Gateway Overall Maximums**:

Description	Limit
Load balancer VIPs	64

Load balancer pools	64
Load balancer servers per pool	32
Firewall rules per ESG	2,000

The following table shows the limits for **DHCP, VPN Service Maximums**:

Description	Limit
DHCP pools per Edge service gateway (all Sizes)	20,000
Number of IPSEC tunnels per Edge gateway - Compact	512
Number of IPSEC tunnels per Edge gateway - Large	1600
Number of IPSEC tunnels per Edge gateway - X-Large	4096
Number of IPSEC tunnels per Edge gateway - Quad-Large	6000
SSL VPN number of concurrent connections (compact/large/x-large/quad-large)	50/100/100/1000

The following table shows the limits for **Multi-vCenter NSX Supported Features**:

Description	Limit

Logical switch	Yes
L2 bridges	No
Logical distributed router	Yes
Distributed firewall	Yes
Edge services	No
IP security groups	Yes

Summary

We started this chapter with an introduction to network virtualization and software-defined networking. We discussed the concepts of network virtualization and introduced VMware's NSX network virtualization platform. We then discussed different NSX features and services, including logical switching, logical routing, Edge gateway services, extensibility, service composer, and data security. We also briefly discussed the multi-vCenter NSX feature. We ended the chapter with configuration maximums for NSX. In [Chapter 2](#), NSX Core Components, we will look at the different components of NSX and VXLAN.

NSX Core Components

This chapter begins with a brief introduction to NSX core components followed by a detailed discussion of these core components. We will go over the three different control planes and how each of these NSX core components fits into that architecture. Next we go over the VXLAN architecture and transport zones that allow us to create and extend overlay networks across multiple clusters. We will also look at NSX Edge and the distributed firewall in greater detail and take a look at the newest NSX feature: multi-vCenter or cross-vCenter NSX deployment. By the end of this chapter, you will have a thorough understanding of NSX core components and also their functional inter-dependencies.

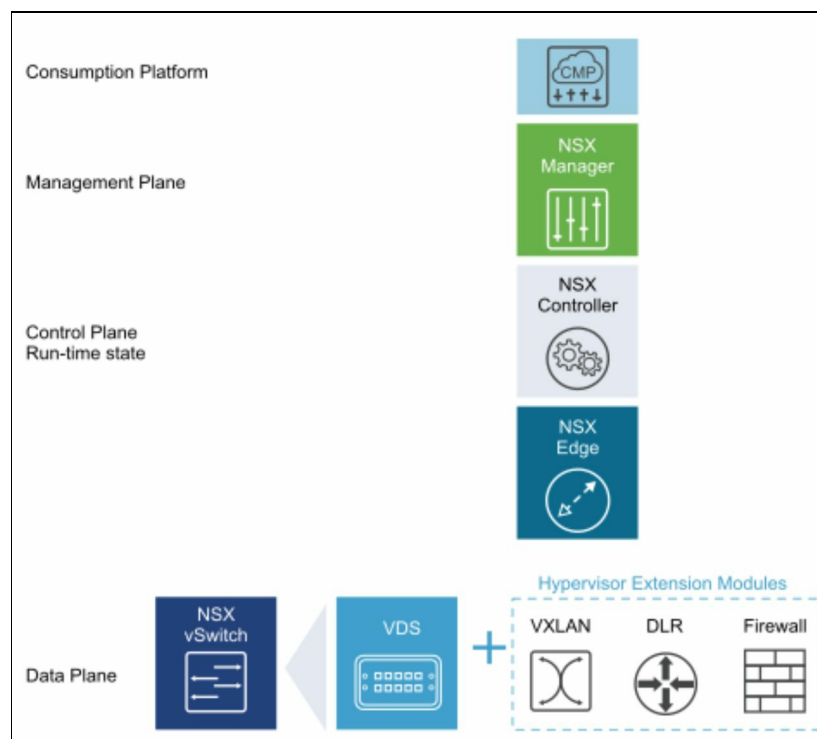
In this chapter, we will cover:

- Introduction to NSX core components
- NSX manager
- NSX controller clusters
- VXLAN architecture overview
- Transport zones
- NSX Edge
- Distributed firewall
- Cross-vCenter NSX

Introduction to NSX core components

The foundational core components of NSX are divided across three different planes. The core components of a NSX deployment consist of an NSX manager, controller clusters, and hypervisor kernel modules. Each of these is crucial for your NSX deployment; however, they are decoupled to a certain extent allowing for resiliency during failure of multiple components. For example, if your controller clusters fail, your virtual machines will still be able to communicate with each other without any network disruption. You have to always ensure that NSX components are always deployed in a clustered environment so they are protected by vSphere HA.

The high-level architecture of NSX primarily describes three different planes wherein each of these core components fits in. They are the **Management Plane**, the **Control Plane**, and the **Data Plane**. The following figure represents how the three planes are interlinked with each other. The management plane is how an end user interacts with NSX as a centralized access point while the data plane consists of north-south or east-west traffic:



(Courtesy VMware)

Notice the consumption platform. The consumption platform allows NSX to interface with multiple cloud management platforms that an organization puts in place so end

users can use NSX without having to access the core NSX manager. NSX has REST APIs that can enable rich integration with any cloud management platform:



Out-of-the-box, NSX integrates into VMware vRealize Automation, vCloud Director, and OpenStack with the Neutron plug-in for NSX.

- **Management plane:** The management plane primarily consists of the NSX manager. The NSX manager is a centralized network management component and primarily allows for a single management point. It also provides the REST API that a user can use to perform all NSX functions and actions. During the deployment phase, the management plane is established when the NSX appliance is deployed and configured. This management plane directly interacts with the control plane and also the data plane. The NSX manager is then managed via the vSphere web client and CLI. The NSX manager is configured to interact with vSphere and ESXi and, once configured, all of the NSX components are then configured and managed via the vSphere web GUI.

It is important to note that the preceding holds true even for a cross-vCenter environment. In a cross-vCenter environment, there are a primary and a secondary NSX manager. The primary manager is responsible for all the universal components such as universal logical switches and universal firewall rules. The secondary manager is responsible for components that are deployed locally to its vCenter.



There is a 1:1 relationship between an NSX manager and a vCenter. There can be up to seven secondary NSX managers associated with a primary NSX manager in a cross-vCenter environment.

- **Control plane:** The control plane consists of the NSX controller that manages the state of virtual networks. NSX Controllers also enable overlay networks (VXLAN) that are multicast-free, making it easier to create new VXLAN networks without having to enable multicast functionality on the physical switches. The controllers also keep track of all information about the virtual machines, hosts, and VXLAN networks and can perform ARP suppression as well. No data passes through the control plane and the lack of controllers does not affect network functionality between virtual machines.

Overlay networks and VXLANs can be used interchangeably. They both





represent L2 over L3 virtual networks. To enable high-availability, controller nodes are deployed in a cluster of three members and failure of the controller nodes does not impact any data-plane traffic.

- **Data plane:** The NSX data plane primarily consists of the NSX logical switch. The NSX logical switch is part of the vSphere distributed switch and is created when a VXLAN network is created. The Logical switch and other NSX services are enabled at the hypervisor kernel level after the installation of the hypervisor kernel modules (VIBs). This logical switch is key in enabling overlay networks that are able to encapsulate and send traffic over existing physical networks. It also allows for gateway devices that allow L2 bridging between virtual and physical workloads.

The data plane receives its updates from the control plane as hypervisors maintain local virtual machine and VXLAN (Logical switch) mapping tables well. The loss of data plane will cause the loss of the overlay (VXLAN) network as virtual machines that are part of a NSX logical switch will not be able to send and receive data.

NSX manager

The NSX manager allows us to create, configure, and manage NSX components in an environment. The NSX manager provides a graphical user interface and also the REST APIs that allow you to interact with various NSX components. NSX manager is a virtual machine that you can download as an OVA and deploy it on any ESX host managed by vCenter.



The NSX version at the time of writing is 6.3 and only supports 1:1 vCenter connectivity.

NSX manager, once deployed and configured, can then deploy controller clusters and prepare the ESXi host, which involves installing various vSphere installation bundles (VIB) that allow for network virtualization features such as VXLAN, logical switching, and logical routing. The NSX manager can also deploy and configure Edge gateway appliances and its services.

Because NSX manager is deployed as a single virtual machine, it relies on VMware's HA functionality to ensure its availability. There is no NSX manager clustering available at the time of writing. It is important to note that loss of NSX manager will cause loss of management and API access, but does not disrupt virtual machine connectivity.

Finally, the NSX manager's configuration UI allows an administrator to collect log bundles and also back up the NSX configuration.



Always deploy NSX manager from the OVA template, which creates a unique UUID. In a cross-vCenter environment, each NSX manager needs to have its own unique UUID.

NSX controller clusters

NSX controller provides a control plane functionality to distribute logical routing and VXLAN network information to the underlying hypervisor. Controllers are deployed as virtual appliances and should be deployed in the same vCenter NSX manager is connected to. In a production environment it is recommended to deploy a minimum of three controllers. For better availability and scalability we need to ensure DRS anti-affinity rules are configured to deploy controllers on a separate ESXI host. Control plane to management and data plane traffic is secured by certificate-based authentication.

It is important to note that controller nodes employ a scale-out mechanism where by each controller node uses a **slicing** mechanism that divides the workload equally across all the nodes. This renders all controller nodes active at all times. If one controller node fails then the other nodes are reassigned the tasks that were owned by the failed node to ensure operational status. VMware NSX controller uses a Paxos-based algorithm within the NSX controller cluster. The controller removes dependency on multicast routing/PIM in the physical network and also suppresses broadcast traffic in VXLAN networks.



NSX version 6.3 only supports up to three controller nodes. Ensure that the controller clusters are deployed on a storage system that has a peak write latency of less than 300 ms and a mean latency of less than 100 ms. Slow disks can cause a controller to become unstable and can cause downtime.

VXLAN architecture overview

One of the most important functions of NSX is enabling virtual networks. These virtual networks or overlay networks have become very popular due to that fact that they can leverage an existing network infrastructure without the need to modify it in any way. The decoupling of logical networks from the physical infrastructure allows users to scale rapidly. Overlay networks or VXLANs were developed by a host of vendors that include Arista, Cisco, Citrix, Red Hat, and Broadcom. This allows the VXLAN standard to be implemented by multiple vendors due to this joint effort in developing its architecture.

VXLAN is a layer 2 over layer 3 tunneling protocol that allows for logical network segments to extend on routable networks. This is achieved by encapsulating the **Ethernet frame** with additional UDP, IP, and VXLAN headers. Consequently, this increases the size of the packet by 50 bytes. Hence, VMware recommends increasing the MTU size to a minimum of 1,600 bytes for all interfaces in the physical infrastructure and any associated vSwitches.

When a virtual machine generates traffic meant for another virtual machine on the same virtual network, the hosts these source and destination virtual machines run on are called **VXLAN tunnel end points (VTEP)**. VTEPs are configured as separate VMKernel interfaces on the hosts. The outer IP header block in the VXLAN frame contains the source and the destination IP addresses that contain the source hypervisor and the destination hypervisor. When a packet leaves the source virtual machine, it is encapsulated at the source hypervisor and sent to the target hypervisor. The target hypervisor, upon receiving this packet, decapsulates the Ethernet frame and forwards it to the destination virtual machine.

Once the ESXI host is prepared from NSX manager we need to configure VTEP. NSX supports multiple VXLAN vmknics per host for uplink load balancing features. In addition to this, Guest VLAN tagging is also supported.

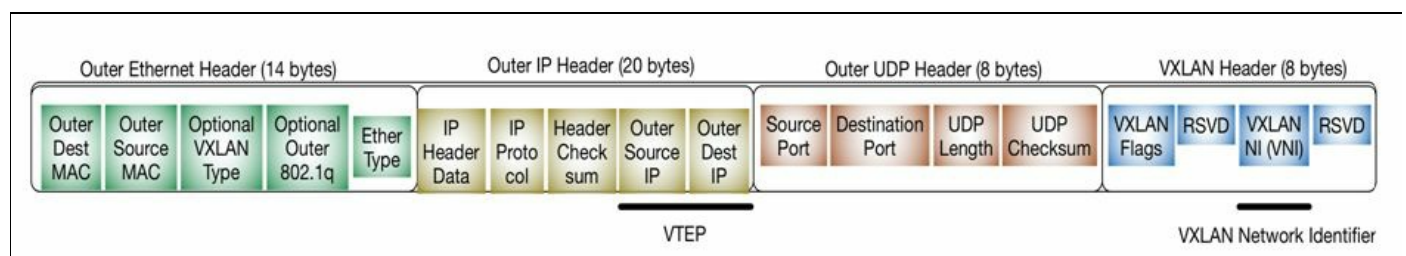
A sample packet flow

When a virtual machine generates traffic **broadcast, unknown unicast, or multicast (BUM)** meant for another virtual machine on the same **virtual network identifier (VNI)** on a different host, the outer IP header block in the VXLAN frame contains the source and the destination IP addresses that contain the source hypervisor and the destination hypervisor. When a packet leaves the source virtual machine, it is encapsulated at the source hypervisor and sent to the target hypervisor. The target hypervisor, upon receiving this packet, de-encapsulates the Ethernet frame and forwards it to the destination virtual machine. Control plane modes play a crucial factor in optimizing the VXLAN traffic depending on the control plane modes selected for the logical switch/transport scope:

- Unicast
- Hybrid
- Multicast

By default, a logical switch inherits its replication mode from the transport zone. However, we can set this on a per logical switch basis. The segment ID is needed for multicast and hybrid mode.

The following is a representation of the VXLAN encapsulated packet showing the VXLAN headers:



As indicated in the preceding figure, the outer IP header identifies the source and the destination VTEPs. The VXLAN header also has the VNI, a 24-bit unique network identifier. This allows for scaling virtual networks beyond the 4094 VLAN limitation placed by the physical switches. Two virtual machines that are part of the same virtual network will have the same virtual network identifier, similar to how two machines on the same VLAN share the same VLAN ID.

Transport zones

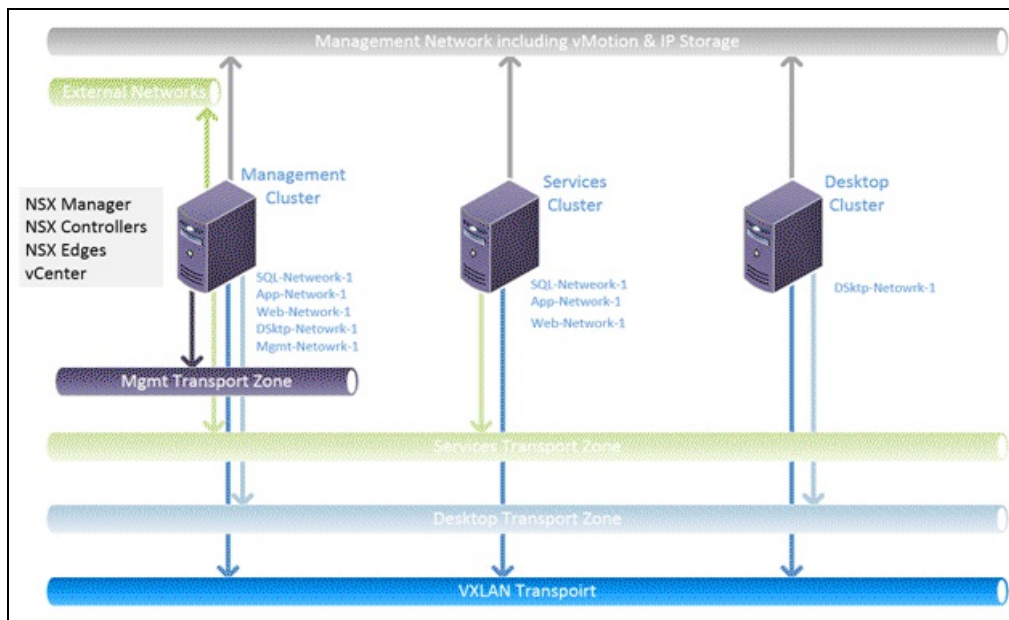
A group of ESXi hosts that are able to communicate with one another over the physical network by means of VTEPs is said to be in the same transport zone. A transport zone defines the extension of a logical switch across multiple ESXi clusters that span across multiple virtual distributed switches.

A typical environment has more than one virtual distributed switch that spans across multiple hosts. A transport zone enables a logical switch to extend across multiple virtual distributed switches and any ESXi hosts that are part of this transport zone can have virtual machines as part of that logical network. A logical switch is always created as part of a transport zone and ESXi hosts can participate in them.

The following figure shows a transport zone that defines the extension of a logical switch across multiple virtual distributed switches:



A universal transport zone allows a logical switch (in this case a universal logical switch) to span multiple hosts across multiple vCenters. A universal transport zone is always created by the primary NSX server and is synchronized to the secondary NSX managers.



Courtesy <http://dailyhypervisor.com/>

NSX Edge Services Gateway

The NSX **Edge Services Gateway (ESG)** offers a feature-rich set of services that include NAT, routing, firewall, load balancing, L2/L3 VPN, and DHCP/DNS relay. NSX API allows each of these services to be deployed, configured, and consumed on-demand. You can install the NSX Edge as an ESG or as a DLR.



The number of Edge appliances including ESGs and DLRs is limited to 250 on a host.

The Edge Services Gateway is deployed as a virtual machine from the NSX manager, which is accessed using the vSphere web client. Four different form factors are offered for different-sized environments. It is important that you factor in enough resources for the appropriate Edge Services Gateway when building your environment.

Edge Services Gateway can be deployed in different sizes. The following are the available size options for Edge Services Gateway appliances:

- **X-Large:** The X-large form factor is suitable for high-performance firewalls, load balancer, and routing, or a combination of multiple services. When an X-large form factor is selected, the Edge Services Gateway will be deployed with 6 vCPU and 8 GB of RAM.
- **Quad-Large:** The Quad-large form factor is ideal for a high-performance firewall. It will be deployed with 4 vCPU and 1 GB of RAM.
- **Large:** The Large form factor is suitable for medium performance routing and firewalls. It is recommended that in production you start with the Large form factor. The Large Edge Services Gateway is deployed with 2 vCPU and 1 GB of RAM.
- **Compact:** The compact form factor is suitable for DHCP and DNS replay functions. It is deployed with 1 vCPU and 512 MB of RAM.

Once deployed, the ESG or the DLR gateway appliance can be upgraded by using the API or the UI. The upgrade action will incur an outage. We will look at upgrading the form factor in more detail in the following chapters. Edge gateway services can also be deployed in an active/standby mode to ensure high availability and resiliency. A heart beat network between the Edge appliances ensures state replication and uptime. If the active gateway goes down and after the declared dead time passes, the standby Edge

appliance takes over.



The default declared dead time is 15 seconds and can be reduced to 6 seconds.


Let's look at some Edge services:

- **Network address translation:** The NSX Edge supports both source and destination NAT allowing NAT for all traffic flowing through the Edge appliance. If the Edge appliance supports more than 100 virtual machines, it is recommended that a Quad instance be deployed to allow for high performance translation.
- **Routing:** The NSX Edge allows for centralized routing that allows logical networks deployed in the NSX domain to be routed to the external physical network. The Edge supports multiple routing protocols including OSPF, iBGP, and eBGP. The Edge also supports static routing.
- **Load balancing:** The NSX Edge also offers the load balancing functionality, which allows load balancing of traffic between the virtual machines. The load balancer supports different balancing mechanisms including IP Hash, least connections, URI based, and round-robin.
- **Firewall:** NSX Edge provides stateful firewall functionality that is ideal for north-south traffic flowing between the physical and the virtual workloads behind the Edge gateway. The Edge firewall can be deployed alongside the hypervisor kernel-based distributed firewall that is primarily used to enforce security policies between workloads in the same logical network.
- **L2/L3VPN:** The Edge also provides L2 and L3 VPN, which makes it possible to extend L2 domains between two sites. IPSEC site-to-site connectivity between two NSX Edges or other VPN termination devices can also be set up.
- **DHCP/DNS relay:** NSX Edge also offers DHCP and DNS relay functions allowing you to offload these services to the Edge gateway. Edge only supports DNS relay functionality and can forward any DNS requests to the DNS server. The Edge gateway can be configured as a DHCP server to provide and manage IP addresses, default gateways, DNS servers, and search domain information for workloads connected to the logical networks.

Distributed firewalls

NSX provides L2-L4 stateful firewall services by means of a distributed firewall that runs in the ESXi hypervisor kernel. Because the firewall is a function of the ESXi kernel it provides massive throughput and performs at near line rate. When the ESXi host is initially prepared by NSX, the distributed firewall service is installed in the kernel by deploying the kernel VIB—**VMware internetworking service insertion platform (VSIP)**. VSIP is responsible for monitoring and enforcing security policies on all the traffic flowing through the data plane. The **distributed firewall (DFW)** throughput and performance scales horizontally as more ESXi hosts are added.

DFW instances are associated to each vNIC and every vNIC requires one DFW instance. A virtual machine with 2 vNICs has two DFW instances associated with it, each monitoring its own vNIC and applying security policies to it. DFW is ideally deployed to protect virtual to virtual or virtual to physical traffic. This makes DFW very effective in protecting east-west traffic between workloads that are part of the same logical network. DFW policies can also be used to restrict traffic between virtual machines and external networks because it is applied at the vNIC of the virtual machine. Any virtual machine that does not require the firewall protection can be added to the exclusion list:



Courtesy <http://humairahmed.com/>

DFW fully supports **vMotion** and the rules applied to a virtual machine always follow the virtual machine. This means any manual or automated vMotion triggered by DRS

does not cause any disruption in its protection status.

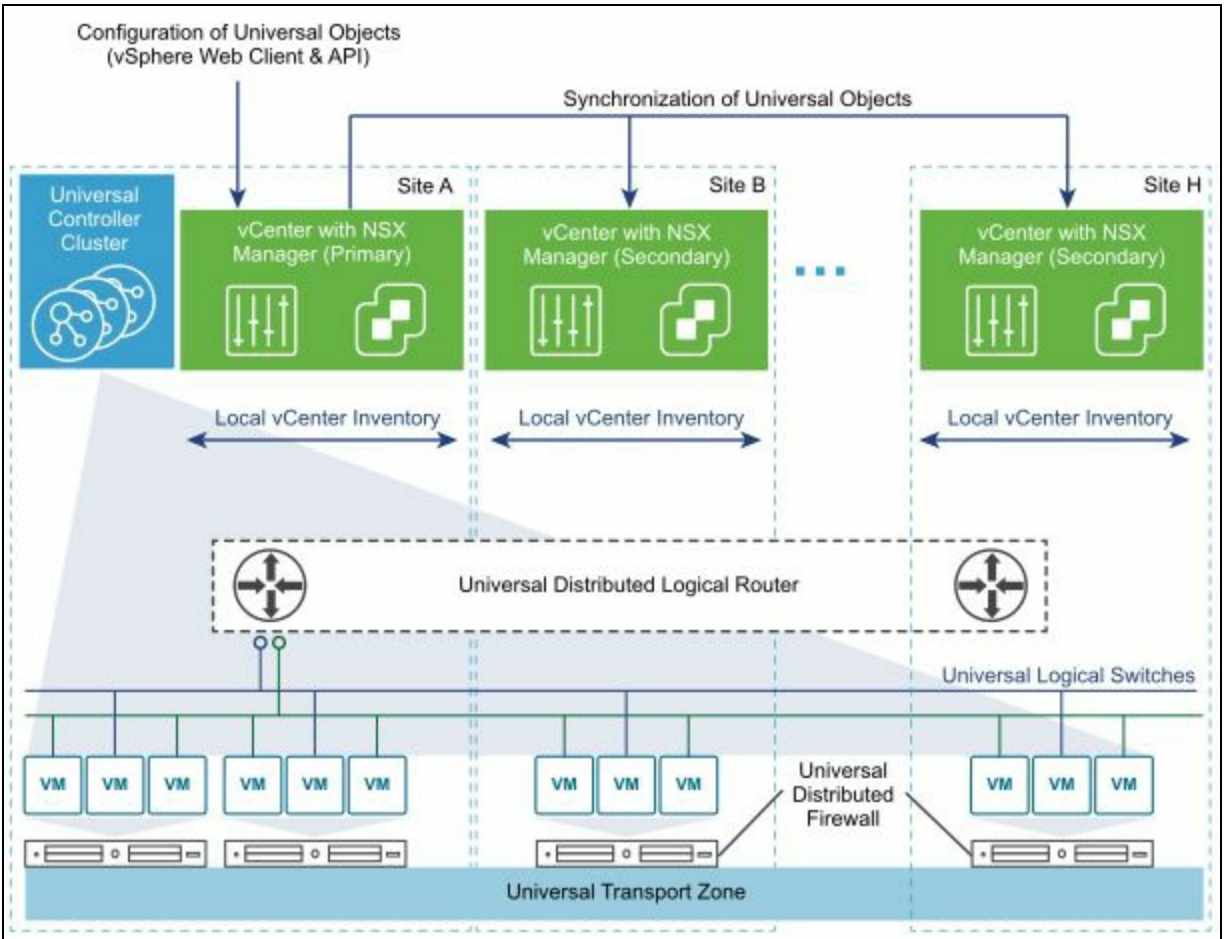
The VSIP kernel module also adds **spoofguard** and traffic redirection functionalities as well. The spoofguard function maintains a VM name and IP address mapping table and prevents IP spoofing. Spoofguard is disabled by default and needs to be manually enabled per logical switch or virtual distributed switch port group. Traffic redirection allows traffic to be redirected to a third-party appliance that can do enhanced monitoring if needed. This allows third-party vendors to interface with DFW directly and offer custom services as needed. We will discuss this in more detail in the following chapters.

Cross-vCenter NSX

Beginning from NSX 6.2, you can manage multiple vCenter NSX environments using the cross-vCenter functionality. This allows you to manage multiple vCenter NSX environments from a single primary NSX manager. This allows ease of management and also enables many new functions that include extending networks and other features such as distributed logical routing. Cross-vCenter NSX deployment also allows for centralized management and also eases disaster recovery architectures.

In a cross-vCenter deployment, multiple vCenters are all paired with their own NSX manager per vCenter. One NSX manager is assigned the primary while other NSX managers become secondary. This primary NSX manager can now deploy a universal controller cluster that provides the control plane. Unlike a standalone vCenter-NSX deployment, secondary NSX managers do not deploy their own controller clusters.

The primary NSX manager also creates objects whose scope is universal. This means these objects extend to all the secondary NSX managers. These universal objects are synchronized across all the secondary NSX managers and can be edited and changed by the primary NSX manager only. This does not prevent you from creating local objects on each of the NSX managers:



Pic courtesy—VMWare

Similar to local NSX objects, primary NSX manager can create global objects such as universal transport zone, universal logical switches, universal distributed routers, universal firewall rules, and universal security objects. There can only be one universal transport zone in a cross-vCenter NSX environment; after it gets created, it gets synchronized across all the secondary NSX managers. When a logical switch is created inside a universal transport zone, it becomes a universal logical switch that spans the layer 2 network across vCenters. All traffic is routed using the universal logical router and any traffic that needs to be routed between a universal logical switch and a logical switch (local scope) requires an Edge Services Gateway. We will discuss these features in upcoming chapters.



You can create objects that are local to that specific vCenter NSX environment in both primary and secondary NSX managers. You can have a maximum of seven secondary NSX managers and one primary NSX manager.

Summary

We began this chapter with a brief introduction to NSX core components and looked at management, control, and the data plane. Next we discussed the NSX manager and NSX controller clusters. This was followed by a VXLAN architecture overview discussion where we looked at the VXLAN packet. We then discussed transport zones and NSX Edge gateway services. The chapter ended with NSX Distributed firewall services and also an overview of Cross-vCenter NSX deployment.

In the [Chapter 3](#), NSX Installation and Configuration, we will start deploying our NSX environment and deploy all of its components starting with the NSX manager. We will configure the NSX manager using its user interface and build our control plane by deploying controller clusters.

NSX Installation and Configuration

This chapter describes the step-by-step installation of NSX and its configuration. We will begin by getting your environment ready for NSX and then go over downloading and deploying NSX. We will then go over the NSX Manager management interface and configure it with your vCenter and license it. We will spend time getting to know the NSX user interface in the vSphere web client and build our control plane by deploying controller clusters. This is followed by our data plane preparation, which involves preparing ESXi clusters for network virtualization.

In this chapter, we will cover:

- Preparing your environment
- Downloading and deploying NSX Manager
- Overview of the NSX Manager interface
- Configuring NSX Manager
- Managing NSX using the vSphere web client
- Deploying the control plane (Controller Virtual Machines)
- Deploying the data plane

Preparing your environment

Before installing NSX, it is important to understand its requirements. NSX Manager and its related components require a considerable amount of resources and planning ahead is very important. The following table lists the minimum resource requirements for NSX Manager and its related components:

Component	CPU	Memory	Disk Space
NSX Manager	4 vCPU	16 GB	60 GB
NSX Controller	4 vCPU	4 GB	20 GB
NSX Edge	1 vCPU (Compact) 2 vCPU (Large) 4 vCPU (Quad Large) 6 vCPU (X-Large)	512 MB(Compact) 1GB (Large) 2GB (Quad Large) 8GB (X-Large)	Compact, Large, Quad Large: 1 disk 584MB + 1 disk 512MB XLarge: 1 disk 584MB + 1 disk 2GB + 1 disk 256MB
Guest Introspection	2 vCPU	1GB	4GB

You also need to have vCenter 6.0 or later installed in your environment and with each server running ESXi version 6.0 or newer. NSX also requires a range of ports to be allowed in your network. We will need TCP port 80 and 443 open for vSphere communication and NSX REST API functionality. We also need TCP ports 1234, 5671,

and 22 for host to controller cluster communication, the rabbit MQ message bus, and SSH console access, respectively.

The following is a list of ports that must be open for NSX to operate flawlessly:

Source	Target	Port	Protocol	Purpose	Sensitive	TLS	Authenti
Client PC	NSX Manager	443	TCP	NSX Manager Administrative Interface	No	Yes	PAM Authentic
Client PC	NSX Manager	80	TCP	NSX Manager VIB Access	No	No	PAM Authentic
ESXi Host	vCenter Server	80	TCP	ESXi Host Preparation	No	No	-
vCenter Server	ESXi Host	80	TCP	ESXi Host Preparation	No	No	-
ESXi Host	NSX Manager	5671	TCP	RabbitMQ	No	Yes	Rabbit M user/pass
ESXi Host	NSX Controller	1234	TCP	User World Agent Connection	No	Yes	-
NSX Controller	NSX Controller	2878, 2888, 3888	TCP	Controller Cluster - State Sync	No	Yes	IPsec
NSX	NSX			Inter-Controller			

Controller	Controller	7777	TCP	RPC Port	No	Yes	IPsec
NSX Controller	NSX Controller	30865	TCP	Controller Cluster - State Sync	No	Yes	IPsec
NSX Controller	NTP Time Server	123	TCP	NTP client connection	No	Yes	No Authentic
NSX Manager	NSX Controller	443	TCP	Controller to Manager Communication	No	Yes	User/Pass
NSX Manager	vCenter Server	443	TCP	TCP vSphere Web Access	No	Yes	-
NSX Manager	vCenter Server	902	TCP	vSphere Web Access	No	Yes	-
NSX Manager	ESXi Host	443	TCP	Management and provisioning connection	No	Yes	-
NSX Manager	ESXi Host	902	TCP	Management and provisioning connection	No	Yes	-
NSX Manager	DNS Server	53	TCP	DNS client connection	No	No	-

NSX Manager	Syslog Server	514	TCP	Syslog connection	No	Yes	-
NSX Manager	NTP Time Server	123	TCP	NTP client connection	No	Yes	-
vCenter Server	NSX Manager	80	TCP	TCP Host Preparation	No	Yes	-
REST Client	NSX Manager	443	TCP	NSX Manager REST API	No	Yes	User/Pass
NSX Controller	NTP Time Server	123	UDP	NTP client connection	No	Yes	No Authentic
NSX Manager	DNS Server	53	UDP	DNS client connection	No	No	-
NSX Manager	Syslog Server	514	UDP	Syslog connection	No	Yes	-
NSX Manager	NTP Time Server	123	UDP	NTP client connection	No	Yes	-
VXLAN Tunnel End Point (VTEP)	VXLAN Tunnel End Point (VTEP)	8472 or 4789*	UDP	Transport network encapsulation between VTEPs	No	Yes	-

ESXi Host	ESXi Host	6999	UDP	ARP on VLAN LIFs	No	Yes	-
ESXi Host	NSX Manager	8301, 8302	UDP	DVS Sync	No	Yes	-
NSX Manager	ESXi Host	8301, 8302	UDP	DVS Sync	No	Yes	-

You will also need virtual distributed switches in your environment, which is the foundation for VXLAN logical segments.



Remember not to upgrade VMware tools on each NSX appliance as specific functionality is tied into each version of these tools. Upgrading without official VMware guidance can potentially break your NSX deployment.

Downloading and deploying NSX Manager

NSX Manager is an appliance and it can be downloaded online from VMware's website. The download is available in the OVA format and can be then deployed in your environment. Before discussing the deployment of NSX, let's go over some considerations for deploying NSX.

VMware recommends that NSX Manager be deployed on a separate management cluster that will be separate from the compute cluster where your software-defined networks will be deployed. This decouples the management plane from the compute plane and allows for higher availability of your management systems. The management cluster will also run the vCenter server and NSX controllers are deployed on the compute cluster. The controller cluster virtual machines should be deployed on the management network and should be able to reach the vCenter server and hypervisors. It is also important to recall that vCenter and NSX have a 1:1 relationship with one NSX Manager only connected to one vCenter.



NSX Manager deploys controller instances on the clusters managed by the NSX registered vCenter. You cannot deploy controller instances to a cluster that is not registered by the NSX Manager.

It is important to also size the environment appropriately to allow for easy resource additions if needed. NSX Manager also does not have built-in HA functionality and relies on ESXi's HA and DRS features to avoid downtime and resource contention. A minimum of three nodes is recommended for a management cluster to run NSX Manager. The management cluster should also be completely separate from any IP addressing space that is chosen to run compute production instances (instances such as your web, database, and application footprints), thereby keeping the management and compute planes isolated.




Always deploy an NSX Manager from the OVA to ensure that a unique UUID is created for every deployment. This is necessary for a cross-vCenter environment.

Once the NSX Manager OVA file is downloaded, we will proceed to import it to our vSphere cluster as follows:

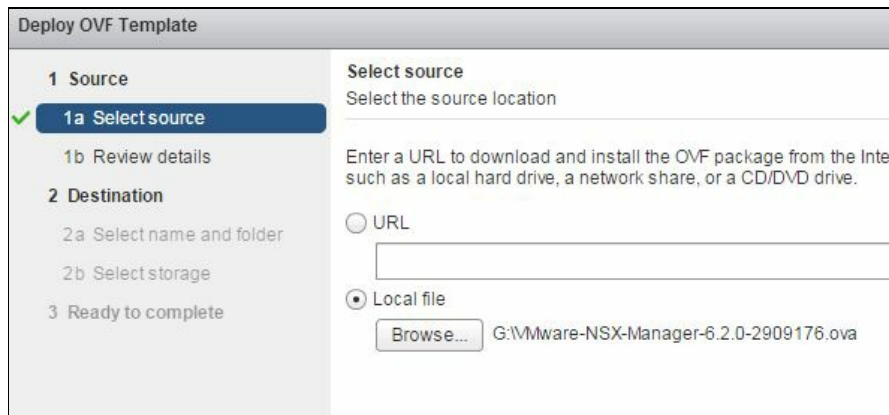


Deployment can be done using the Windows-based vSphere client or vSphere web client. Because NSX can only be managed using the web client, we will deploy using the web client, which is preferred.

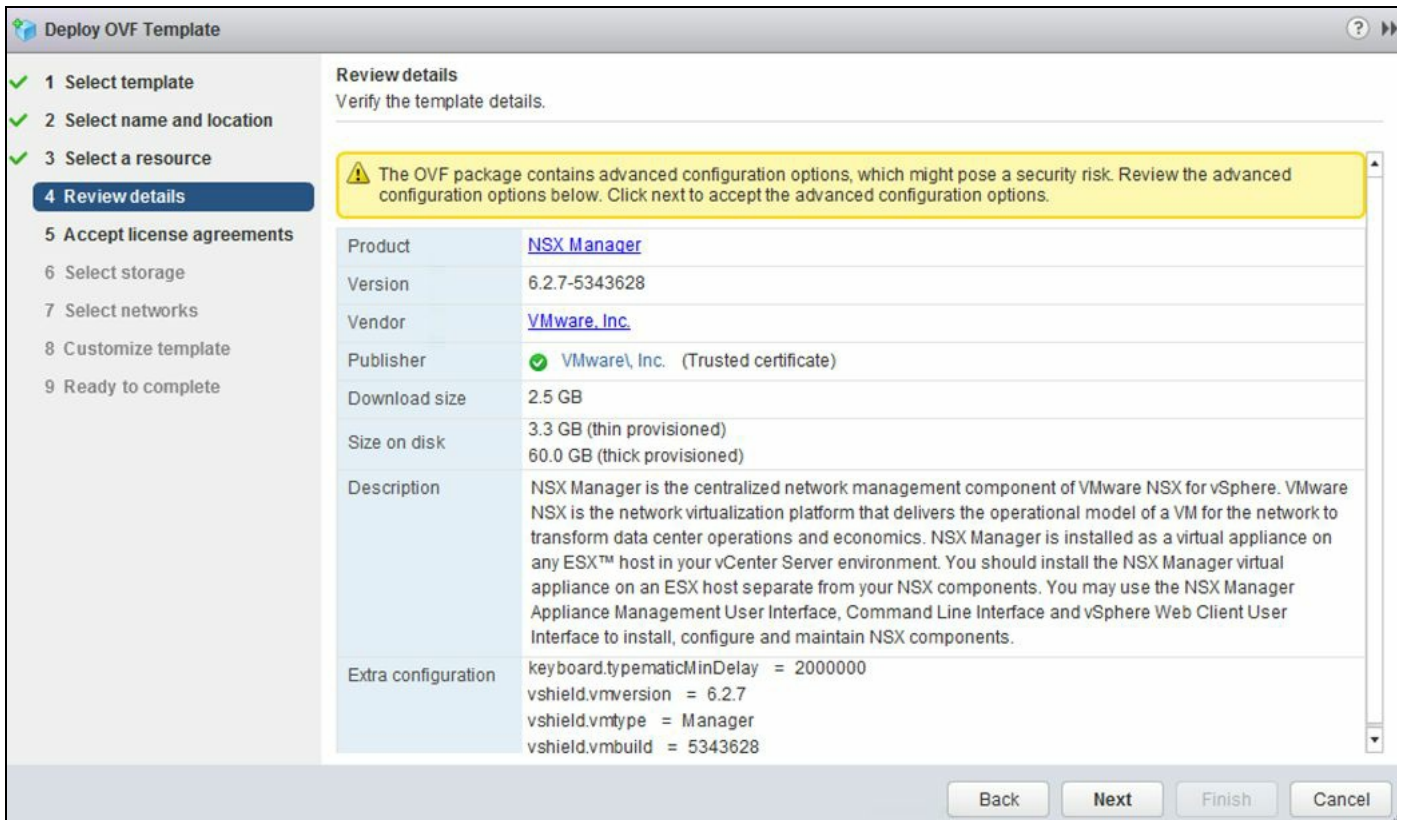
1. Get to the OVF deployment screen by clicking on Hypervisors and Clusters. Expand the vCenter, right-click the hypervisor, and select Deploy OVF Template...:



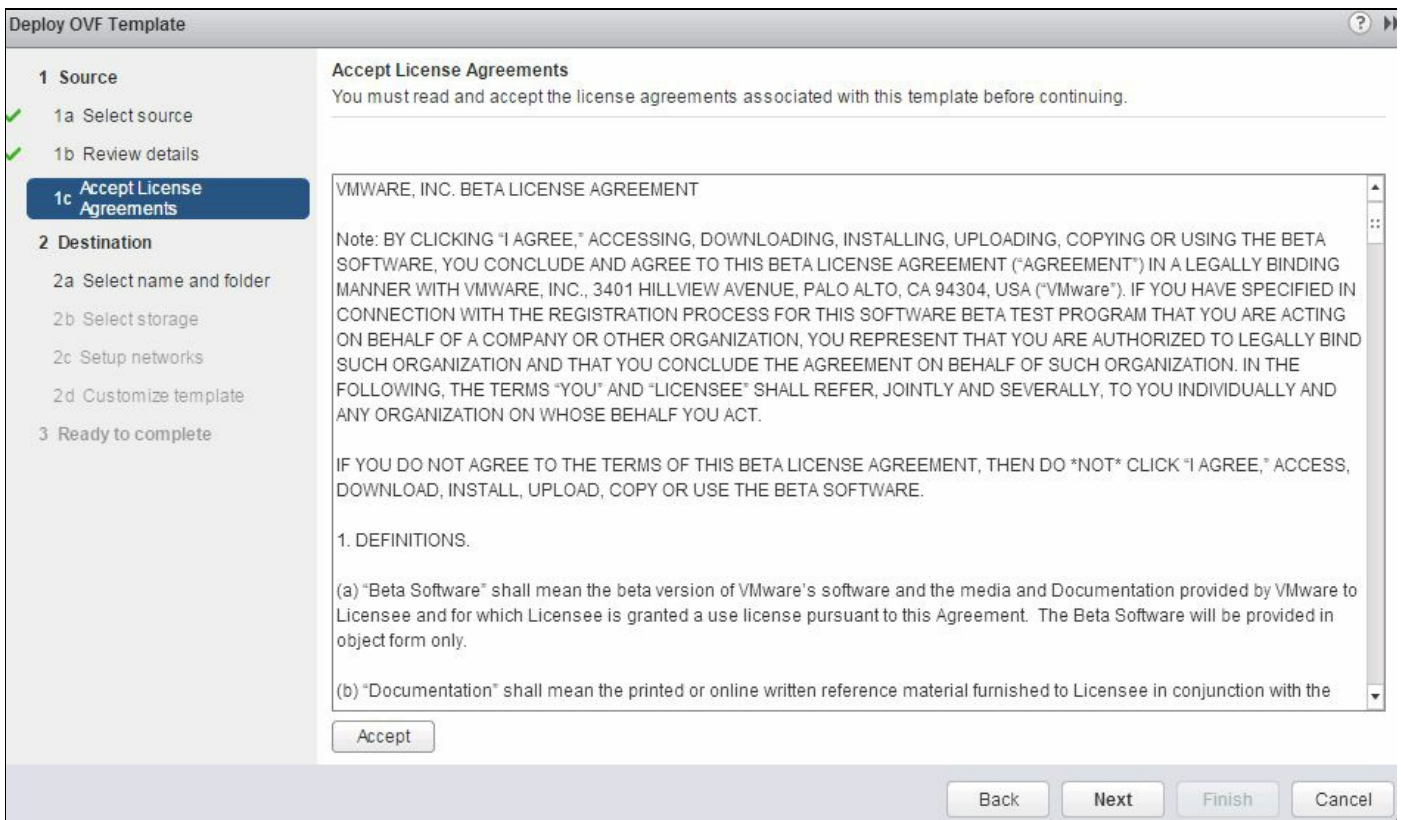
2. You will see the Deploy OVF Template... screen:



3. Select the OVF/OVA package that you have downloaded and click Next. You will see the details of the NSX OVA. Accept the extra configurations by clicking on the checkbox and click Next:

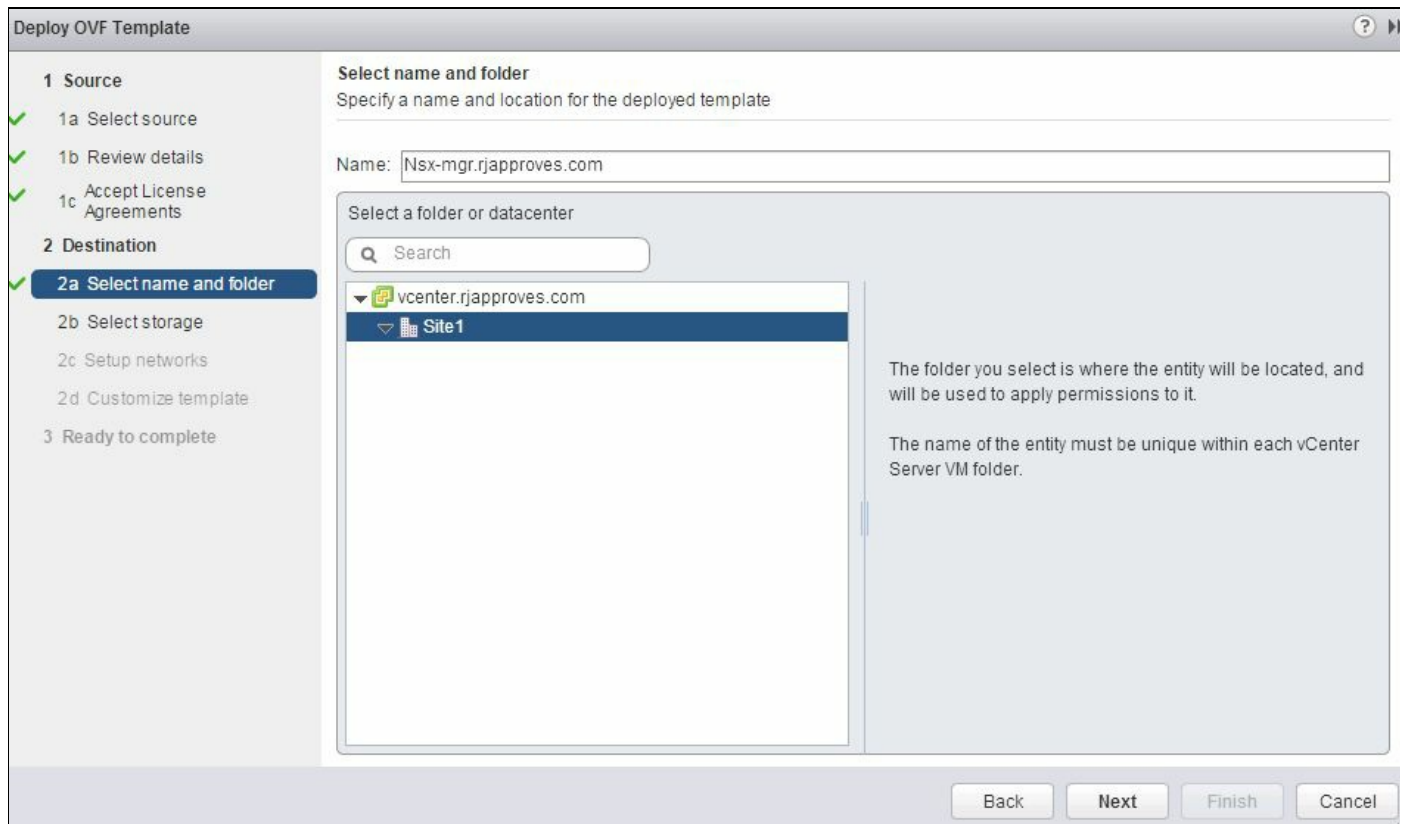


4. Accept the EULA by clicking the Accept button and then click Next:



5. Name your NSX Manager instance and select a datacenter or a folder you want it

to be deployed to. Click Next to proceed to the next screen:



6. Next, select the storage location and the appropriate policy. The default disk format is Thick Provision Lazy Zeroed, but I am going to deploy this appliance in a thin disk format to conserve space in my lab. The best practice is to deploy either Thick Provision Lazy Zeroed or Thick Provision Eager Zeroed:

Deploy OVF Template

1 Source

- 1a Select source
- 1b Review details
- 1c Accept License Agreements

2 Destination

- 2a Select name and folder
- 2b Select storage**
- 2c Setup networks
- 2d Customize template

3 Ready to complete

Select storage
Select location to store the files for the deployed template

Select virtual disk format: **Thick Provision Lazy Zeroed**

VM Storage Policy: **Datastore Default** ⓘ

The following datastores are accessible from the destination resource that you selected. Select the destination datastore for the virtual machine configuration files and all of the virtual disks.

Name	Capacity	Provisioned	Free	Type	Storage DRS
Local-hlu001	465.50 GB	540.33 GB	192.23 GB	VMFS	

Back Next Finish Cancel

7. Select the network onto which this NSX Manager will be deployed. Typically this will be a management network where NSX Manager is able to talk to the vCenter it will connect to. Click Next:

Deploy OVF Template

1 Source

- ✓ 1a Select source
- ✓ 1b Review details
- ✓ 1c Accept License Agreements

2 Destination

- ✓ 2a Select name and folder
- ✓ 2b Select storage
- ✓ 2c Setup networks
- 2d Customize template

3 Ready to complete

Setup networks
Configure the networks the deployed template should use

Source	Destination	Configuration
VSMgmt	VM Network	✓

IP protocol: IPv4 IP allocation: Static - Manual ⓘ

Source: VSMgmt - Description
This network provides connectivity to this virtual machine.

Destination: VM Network - Protocol settings
No configuration needed for this network

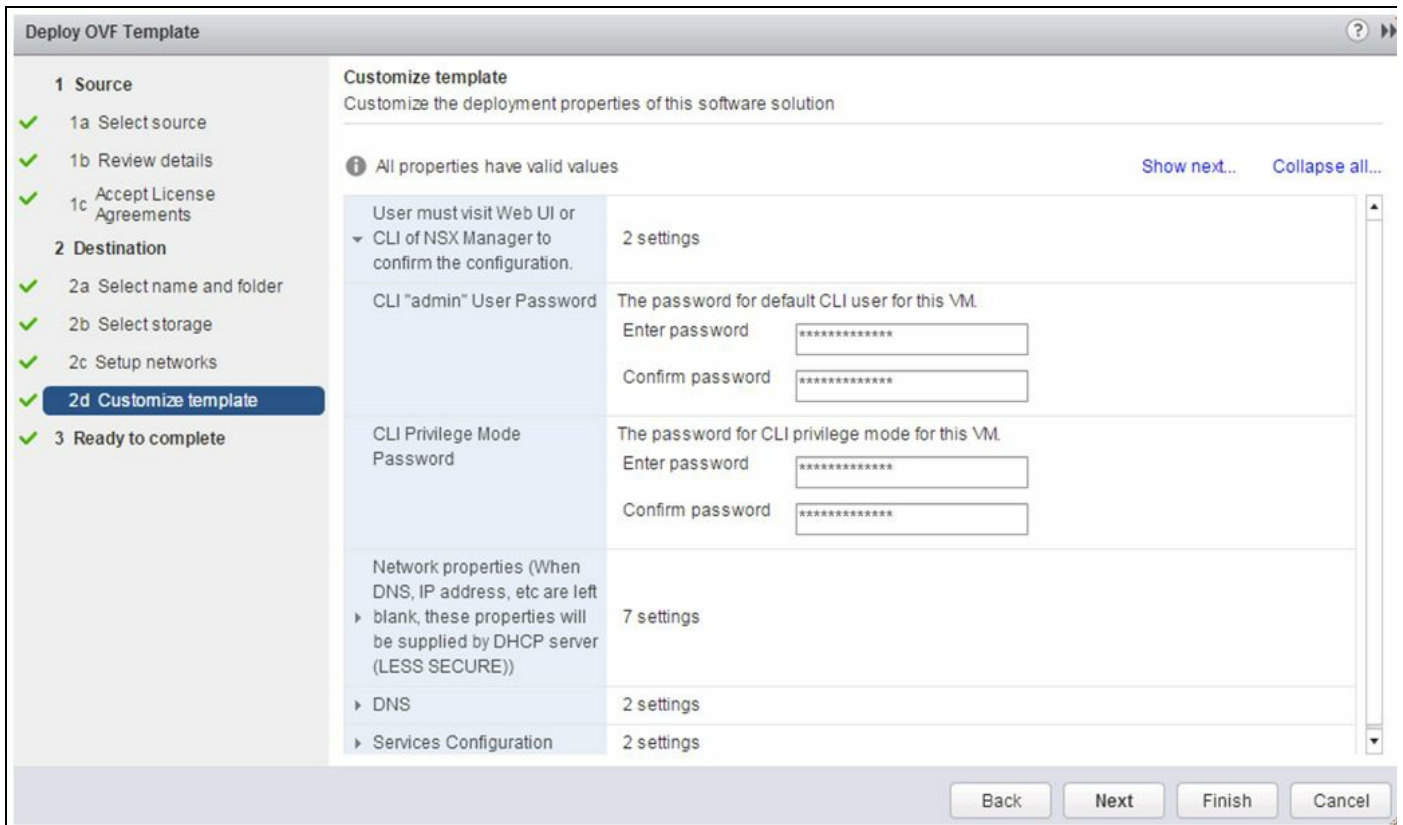
Back Next Finish Cancel

8. Next you will have to input all the values you'll configure this appliance with. This includes CLI admin and privileged mode passwords, IP address assignments, and DNS and NTP settings as well. There is an option to enable SSH mode; however, in production environments this is not a best practice unless you have taken appropriate measures to secure SSH access. Leaving the IP section blank will invoke a DHCP request to your DHCP server.

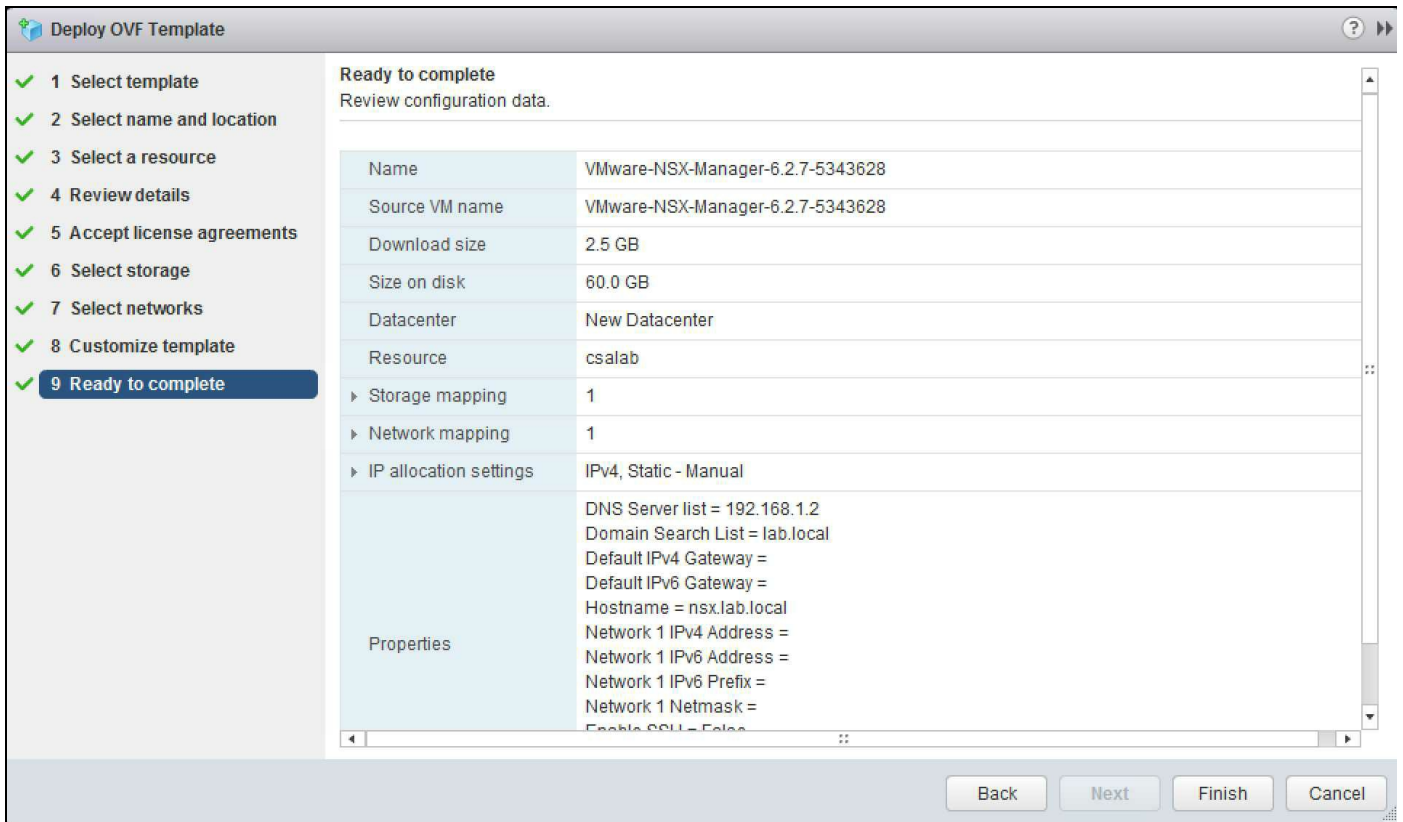


Ensure that forward and reverse DNS lookup entries are in place.


9. Click Next once you have completed all the values appropriately:



10. Review the summary page and click Finish:



11. You will now see your NSX Manager being deployed in the Tasks pane:



Overview of the NSX Manager interface

Once NSX Manager is successfully deployed you should be able to see it in your inventory:

1. Click on the NSX Manager virtual machine and review the Summary tab:

The screenshot displays the vSphere Summary page for a virtual machine named **nsx.lab.local**. The interface includes a navigation bar with tabs: Getting Started, Summary (selected), Monitor, Configure, Permissions, Snapshots, Datastores, Networks, and Update Manager.

System Information:

- Guest OS: Other Linux (64-bit)
- Compatibility: ESXi 5.0 and later (VM version 8)
- VMware Tools: Running, version:2147483647 (Guest Managed)
- DNS Name: nsx.lab.local
- IP Addresses: (None listed)
- Host: esx3.csalab.local

Resource Usage:

- CPU USAGE: 130.00 MHz
- MEMORY USAGE: 163.00 MB
- STORAGE USAGE: 47.07 GB

VM Hardware:

- CPU: 4 CPU(s), 130 MHz used
- Memory: 8192 MB, 163 MB memory active
- Hard disk 1: 60.00 GB
- Network adapter 1: vxw-dvs-84-virtualwire-1-sid-5000-Lab-vxlan (connected)
- CD/DVD drive 1: Disconnected
- Floppy drive 1: Disconnected
- Video card: 4.00 MB
- Other: Additional Hardware
- Compatibility: ESXi 5.0 and later (VM version 8)

NSX Activity Monitoring:

- Virtual Machine Data Collection: Disabled
- Global Data Collection: Enabled

Security Group Membership:

Name	Description
This list is empty.	

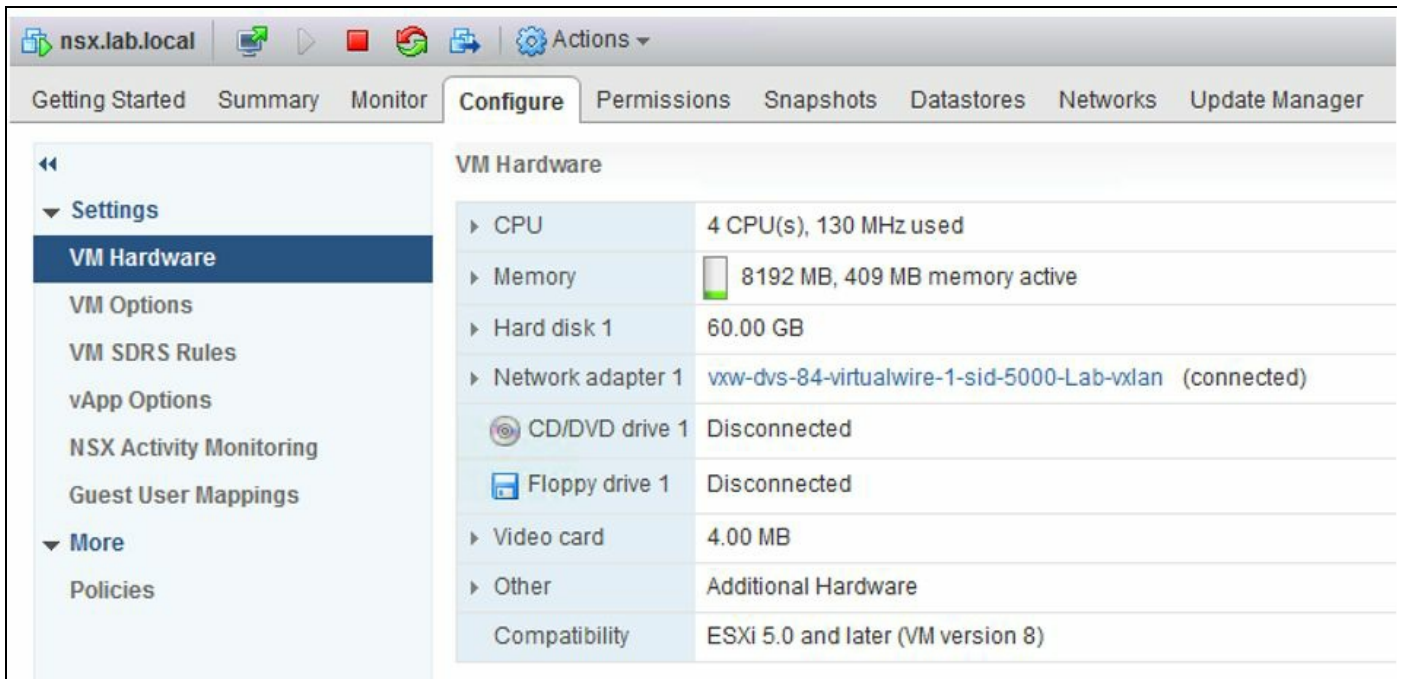
Security Tags:

Security tags can not be applied to a service virtual machine.

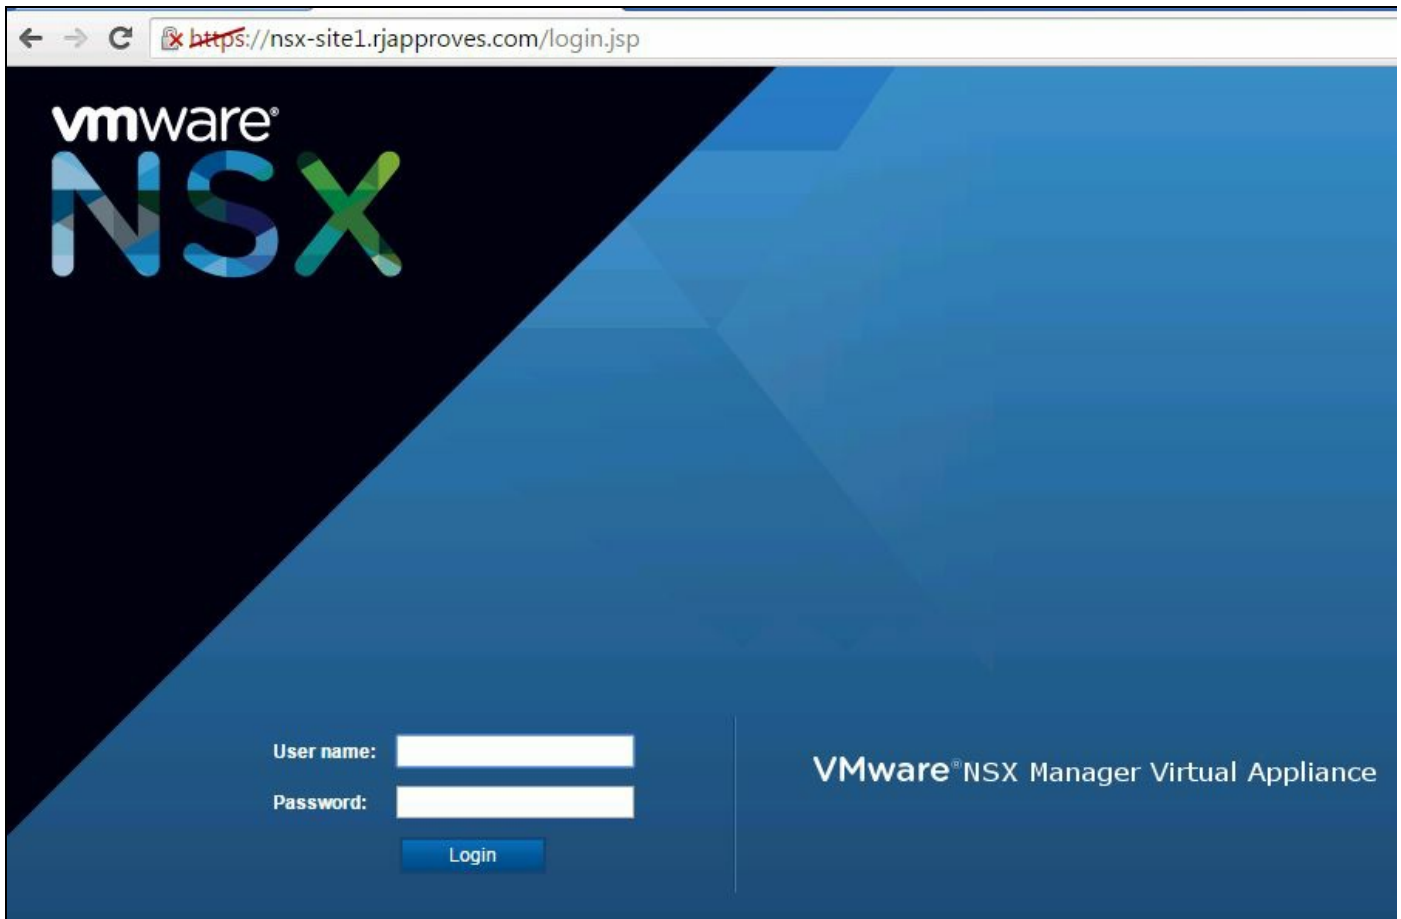
Tags:

Assigned Tag	Category	Description
This list is empty.		

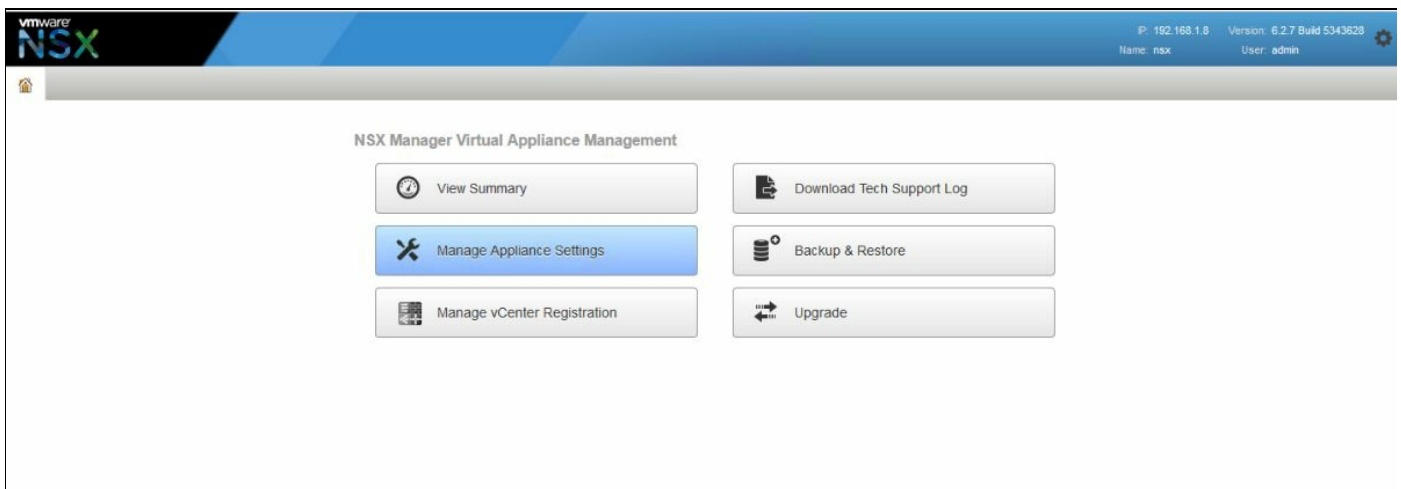
2. You should be able to see that the NSX virtual machine adds notes to the summary page. You will also see that this virtual machine is taking up about 8 GB of disk space initially.
3. Let's click on the Configure tab to review its default hardware requirements:



4. The NSX Manager virtual machine has been allocated four CPUs and about 16 GB of RAM. The hard disk is set to be 60 GB in size.
5. Let's proceed to power on this virtual machine. Right-click the NSX Manager virtual machine and navigate to Power | Power On.
6. You will now see that NSX is powered on and VMware guest tools have been started.
7. Let's access the NSX Manager from the browser and review its interface and options. Open up a new browser or a new tab and enter the FQDN or the IP address of the NSX Manager. The NSX Manager web interface is accessed over SSL:



8. Log in to the NSX Manager with the username `admin` and the password that was set during deployment time. Once logged in you will see the splash page, which allows you to configure your appliance:



9. The View Summary view gives you a summary of the current state of the NSX Manager appliance. This includes resource consumptions and a service status:

Upgrade the system virtual hardware to at least 16 GB RAM.

NSX Manager Virtual Appliance

DNS Name: nsx.lab.local
 IP Address: 192.168.1.8
 Version: 6.2.7 Build 5343628
 Uptime: 3 days, 6 hours, 55 minutes
 Current Time: Wednesday, 07 June 2017 10:39:46 PM CDT

Resource Usage:

- CPU:** Free: 1721 MHz, Used: 140 MHz, Capacity: 1881 MHz
- MEMORY:** Free: 4496 MB, Used: 3477 MB, Capacity: 7973 MB
- STORAGE:** Free: 57G, Used: 21G, Capacity: 78G

Common components

Name	Version	Status	
vPostgres		Running	Stop
RabbitMQ		Running	Stop

System-level components

Name	Version	Status	
SSH Service		Stopped	Start

NSX Management Components

Name	Version	Status	
NSX Universal Synchronization Service		Stopped	Start
NSX Management Service	6.2.7 Build 5343628	Running	Stop

10. To enable SSH access to the manager, simply click on the Start button and the service will be started.

11. The Manage view allows you to manage and configure NSX Manager:

SETTINGS

- General
- Network
- SSL Certificates
- Backups & Restore
- Upgrade

COMPONENTS

- NSX Management Service

Time Settings [Unconfigure NTP Servers] [Edit]

Specify NTP server below. For SSO configuration to work correctly it is required that the time on this virtual appliance and NTP server should be in sync. It is recommended to use the same NTP server used by the SSO server.

NTP Server: 0.north-america.pool.ntp.org
 Timezone: US/Central
 Date/Time: 06/07/2017 22:40:01

Syslog Server [Edit]

You can specify the IP address or name of the syslog server that can be resolved using the above mentioned DNS Server(s).

Syslog Server: _____
 Port: _____
 Protocol: _____

Locale [Edit]

Below is the current locale information.

Locale: en-US

12. Some of the settings you can configure include configuring a syslog server, the network, SSL certificates, and backup and restore; you can also upgrade the appliance. We will learn more about these in the upcoming chapters.

The NSX management service allows you to configure NSX Manager with a Virtual Center and a lookup service:

The screenshot shows the VMware NSX Manager interface. At the top, the VMware NSX logo is on the left, and the IP address (192.168.1.8), Version (6.2.7 Build 5343628), Name (nsx), and User (admin) are on the right. Below the header, there are tabs for 'Summary' and 'Manage'. A left sidebar contains 'SETTINGS' (General, Network, SSL Certificates, Backups & Restore, Upgrade) and 'COMPONENTS' (NSX Management Service). The main content area is divided into two sections: 'Lookup Service URL' and 'vCenter Server'. The 'Lookup Service URL' section includes a table with fields for 'Lookup Service URL' (https://vcenter.lab.local:443/lookupservice/sdk), 'SSO Administrator User Name' (administrator@vsphere.local), and 'Status' (Connected). The 'vCenter Server' section includes a table with fields for 'vCenter Server' (vcenter.lab.local), 'vCenter User Name' (administrator@vsphere.local), and 'Status' (Connected - Last successful inventory update was on Thu, 08 Jun 2017 02:48:12 GMT).

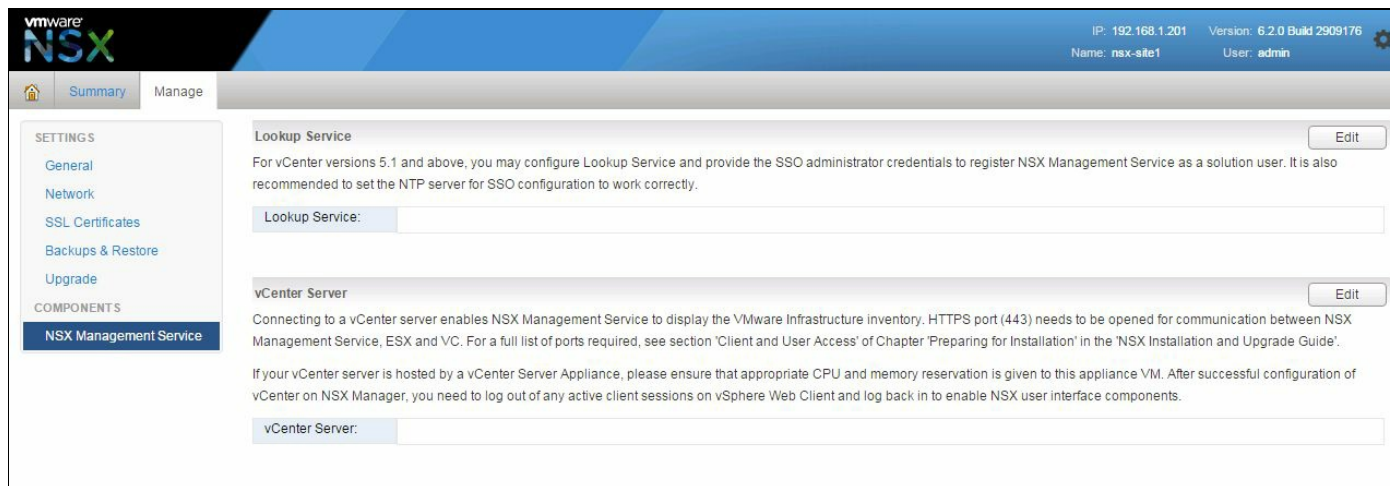
13. Lastly, to log out of your NSX Manager, click on the gear symbol in the upper-right corner and click on the Logout option:



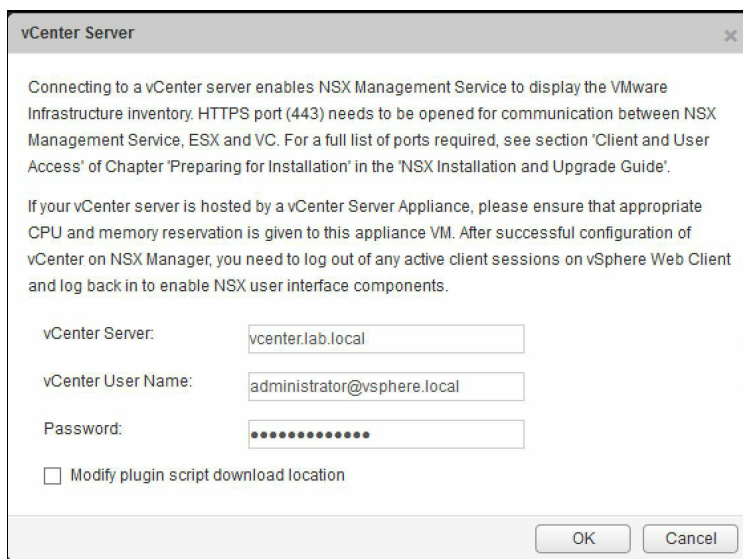
Configuring NSX Manager

Now that we have deployed NSX Manager and taken a quick overview of it, let's go ahead and configure this manager with our vCenter server:

1. Once you are logged into NSX Manager, click on Manage. Under the Components section, click on NSX Management Service:



2. Under the vCenter Server section click Edit. You will see the following screen:



3. Enter your vCenter Server FQDN and a vCenter User Name and Password. Click OK when done.

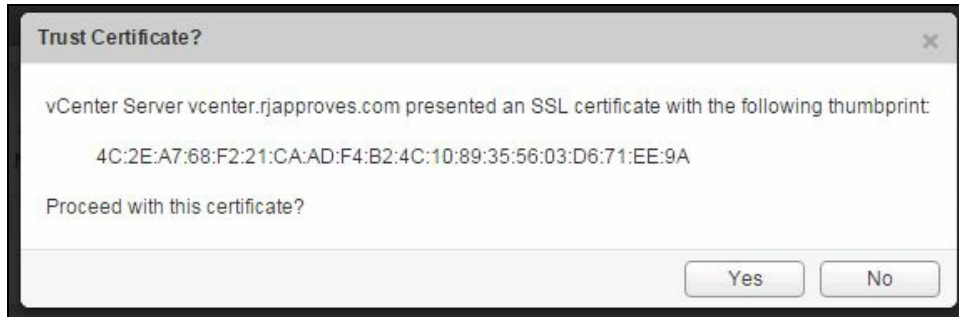
It is a best practice to create a custom user for NSX Manager. You have



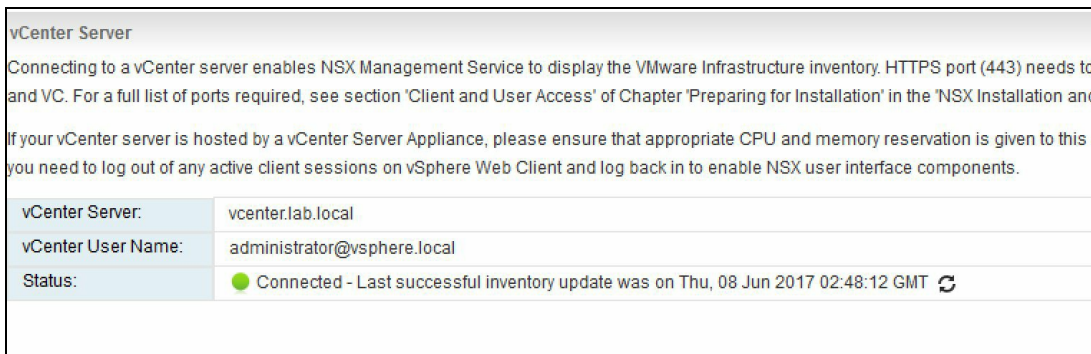


roles built in such as Enterprise administrator, NSX administrator, and Security administrator.

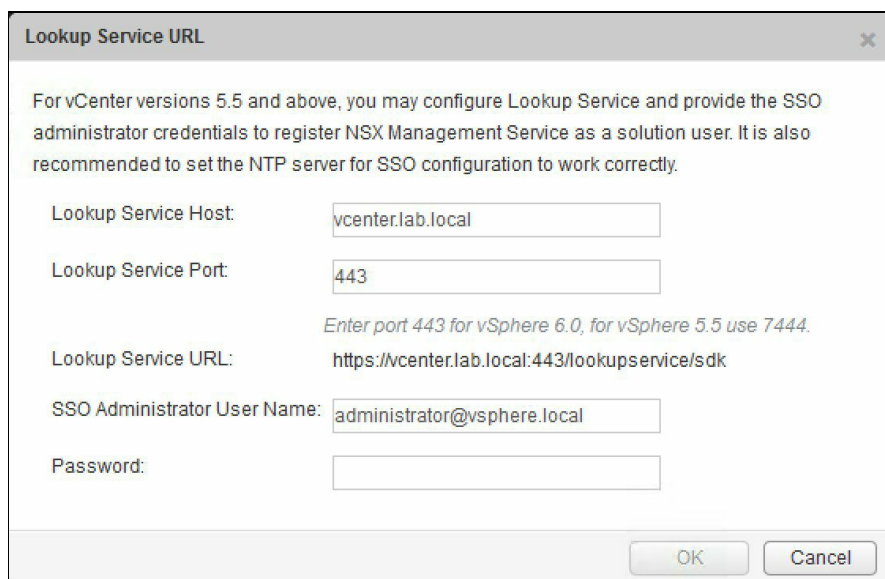
4. Accept the SSL certificate by clicking Yes:




5. Once connected you will see a Connected status:



6. To use SSO on NSX Manager you should have the SSO service installed on the vCenter. Click on Edit in the Lookup Service section and fill out your SSO service credentials. Accept the SSL certificate prompts:



7. Once connected you will see the status update as follows:

Lookup Service URL	
For vCenter versions 5.5 and above, you may configure Lookup Service and provide the SSO server for SSO configuration to work correctly.	
Lookup Service URL:	https://vcenter.lab.local:443/lookupservice/sdk
SSO Administrator User Name:	administrator@vsphere.local
Status:	● Connected 

8. We have now configured our NSX Manager with a vCenter server.


Managing NSX using the vSphere web client

NSX can only be managed using the vSphere web client. When NSX is configured with a vCenter, it installs a plugin for the web client. There is no management functionality available in the classic windows client.



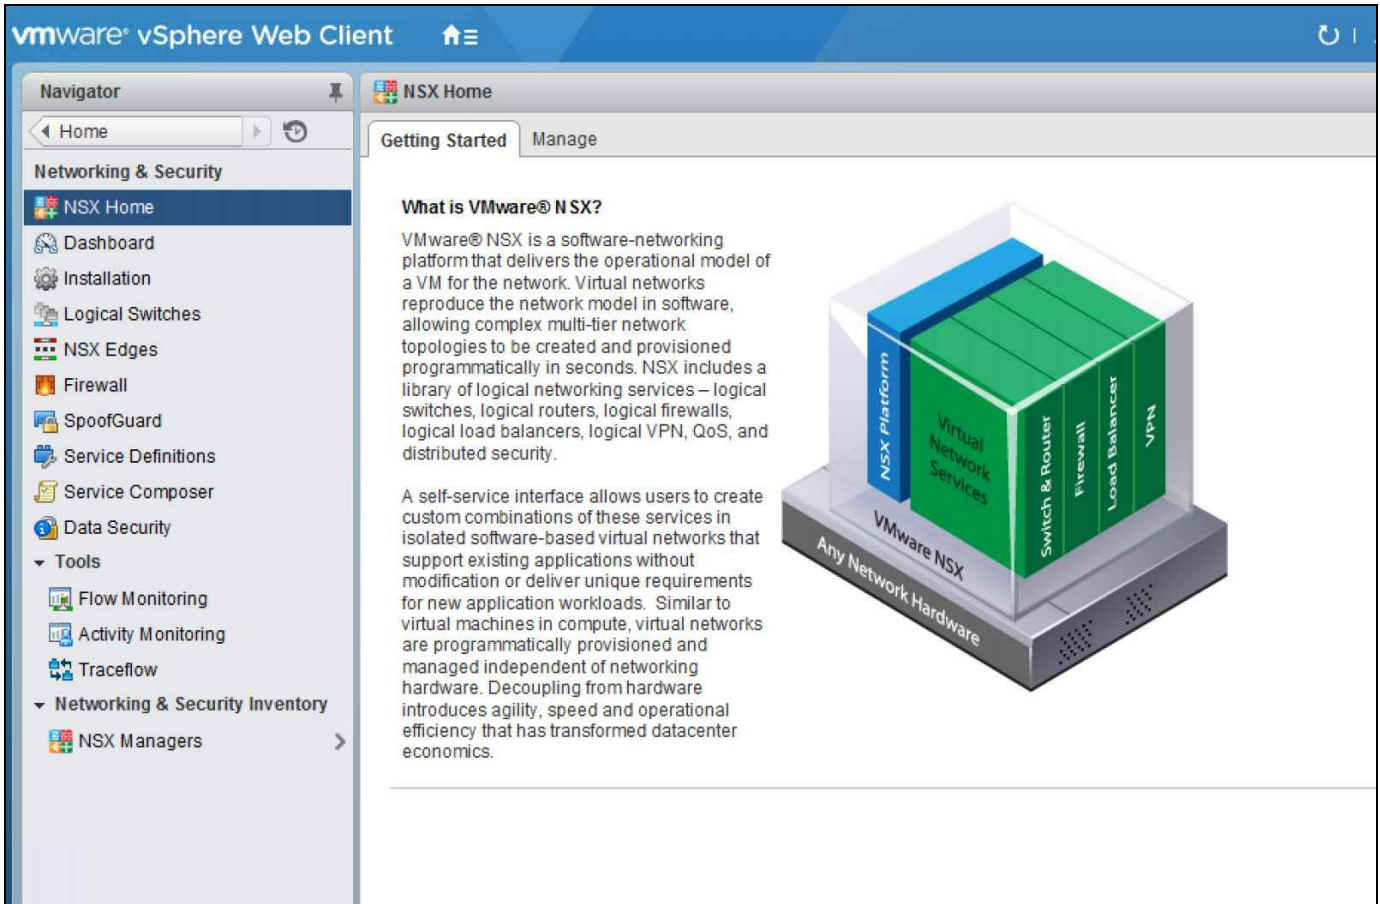
If you are already logged in, you will need to log out and log back in to the vSphere web client after about five minutes for the plugin to show up.

Once logged in to your vSphere web client you will see a new icon called Networking & Security. You will see this on the left side navigation pane as well:

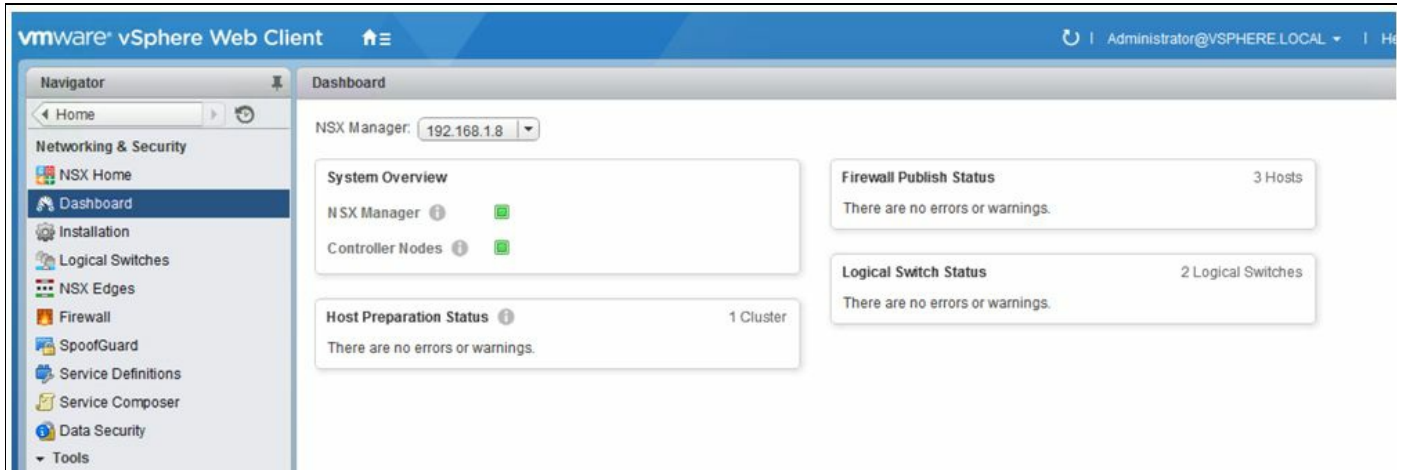


Now perform the following steps to generate a license:


1. Click on the Networking & Security icon or menu option and you should see the following screen:



2. There is a Dashboard view that gives overall system status:



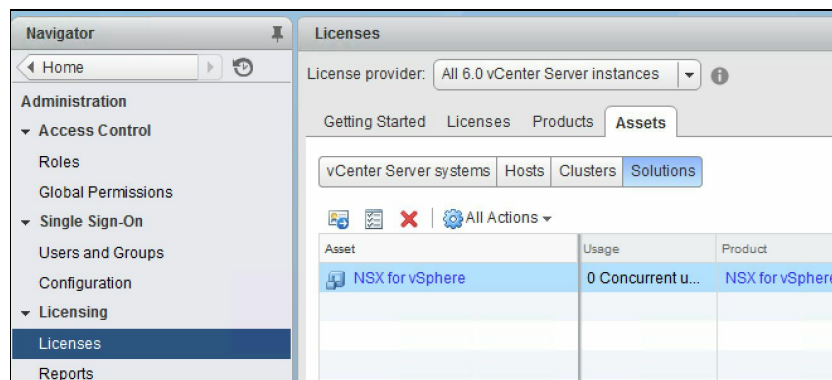
3. There are many menu items, but let's click on NSX Managers under Networking & Security Inventory. This shows us the NSX Manager that is registered:



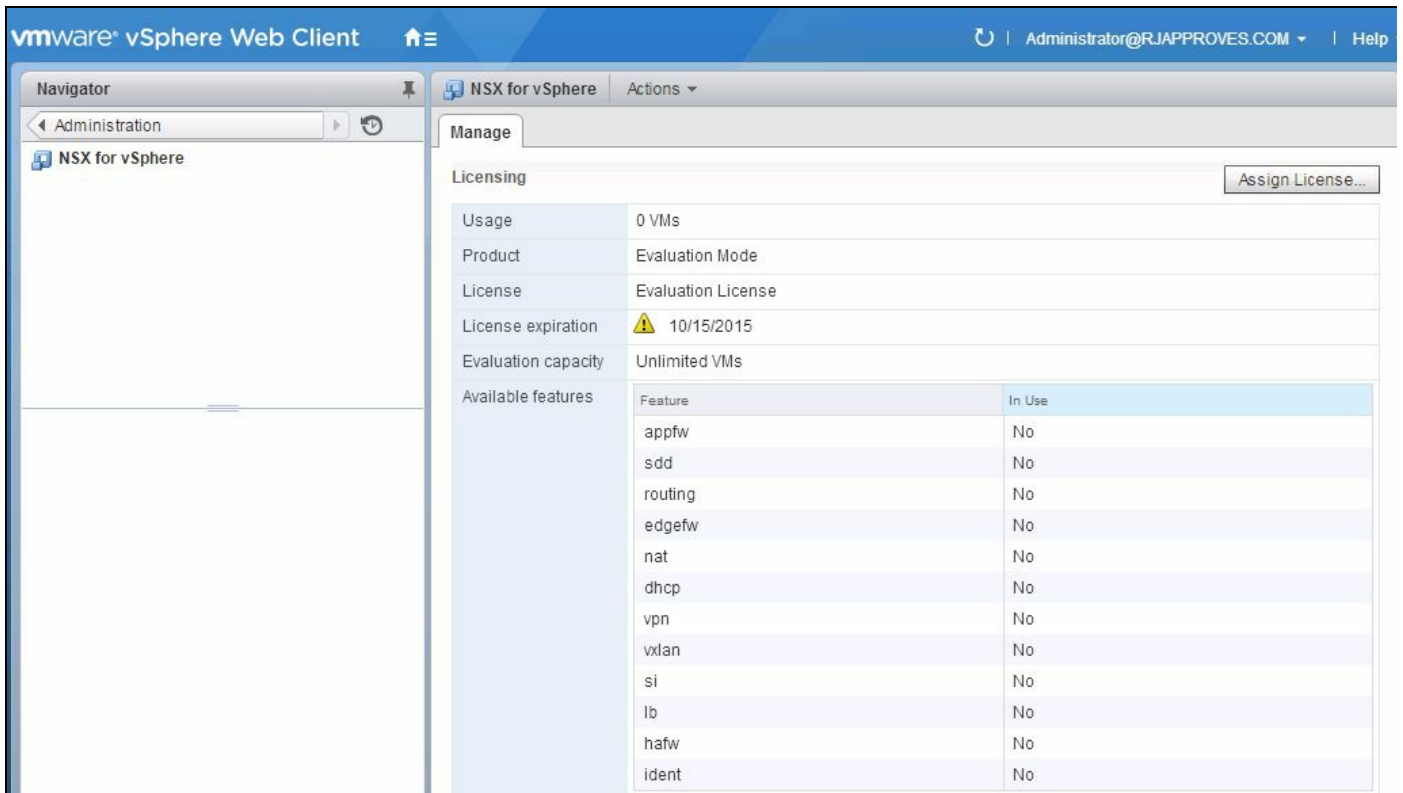
4. Clicking on NSX Managers allows you to view the current licensing status and manage and monitor NSX Manager:



5. Before we begin setting up the control plane and prepare our clusters for Network virtualization, let's add the appropriate licensing for NSX. NSX by default allows up to 60 days of license-free usage.
6. Log in to the web client and click on Administration | Licenses | Assets | Solutions:



7. Click on the NSX for vSphere | Assign License... button:



8. In the License view, click the + symbol to open up the licensing window:



9. Copy and paste or type in your license key and click Next.
10. Click Finish when ready. Make sure you select NSX license from the licensing options.

Deploying the control plane (Controller Virtual Machines)

The control plane consists of the controller nodes that eliminate the need for multicast support from the physical network infrastructure to manage VXLAN-based logical switches. A minimum of three controller virtual machines is required for high availability.

To deploy controller nodes, go to Networking & Security | Installation, and under the Management tab click the + symbol in the NSX Controller nodes section:



VMware recommends that you add three controllers for scale and redundancy. Controller cluster virtual machines employ a slicing mechanism, that is, all of them are active at any given time. If a controller virtual machine fails, the tasks failover to the other active peers. This is done by means of a master controller that is elected by means of a majority vote. To maintain this majority it is always recommended to deploy controller cluster virtual machines in odd numbers. Ensure there is no resource contention in the cluster while deploying the controller nodes. You cannot modify resources for the controller nodes.

A screenshot of the 'Add Controller' dialog box in VMware NSX Manager. The dialog contains several fields for configuration: NSX Manager (192.168.1.201), Datacenter (Site1), Cluster/Resource Pool (Site1), Datastore (Local-hlu001), Host (esx2.rjapproves.com), Folder (Discovered virtual machi...), Connected To (VM Network), IP Pool (empty), Password (empty), and Confirm password (empty). A red arrow points to the IP Pool field with the text 'This field is required'. The dialog has 'OK' and 'Cancel' buttons at the bottom.

Perform the following steps to add a controller:

1. Select the appropriate NSX Manager, Datacenter, Cluster/Resource Pool, Datastore, desired host if any, folder if any, and the network port group as well. The network port group can be a logical switch, port group, or a distributed port

group. The controller node must be reachable by NSX Manager and the vSphere hosts it is to communicate with. Usually controller nodes are part of the management network.

2. Click Select to define the IP pool, which is a range of pre-determined IP addresses. These IP's should be able to reach the vCenter and the management network including the compute hypervisors. The controller node, when deployed, gets an IP allocated from this pool of IP's:



3. Click on New IP Pool... and fill in the values appropriately. The IP range needs to be added and sized appropriately to accommodate the number of controller nodes you plan to deploy:

Add Static IP Pool

Name: * Controller-pool

Gateway: * 192.168.1.1
A gateway can be any IPv4 or IPv6 address.

Prefix Length: * 24

Primary DNS: 192.168.1.50

Secondary DNS: 8.8.8.8

DNS Suffix: rjapproves.com

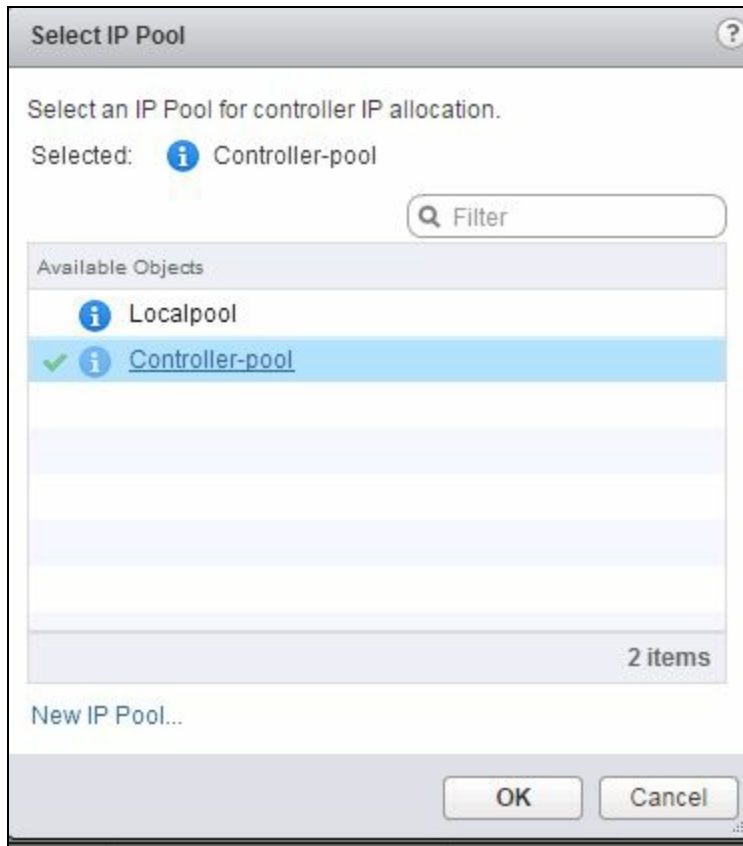
Static IP Pool: * 192.168.1.210-192.168.1.211

for example 192.168.1.2-192.168.1.100 or
abcd:87:87::10-abcd:87:87::20

OK Cancel

4. Click OK when complete and you will see the pool show up in the list of available pools.

5. Click a pool to select it and press OK:




6. Next, enter the password and click OK.

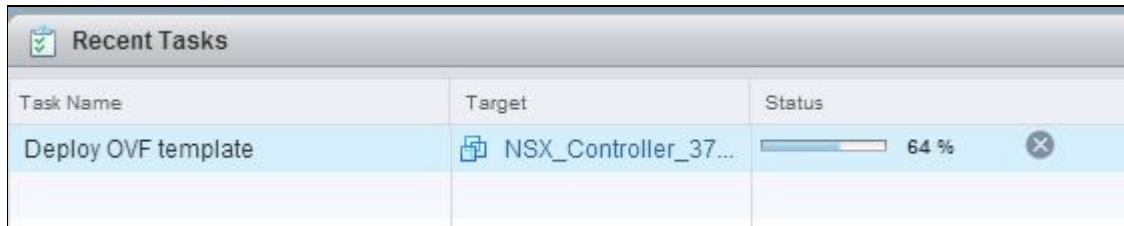
The password must be eight characters and must follow three of the following four rules:

- At least one upper-case letter
- At least one lower-case letter
- At last one number
- At least one special character

7. When done, you will now see that the controller node is being deployed with the Deploying status shown:

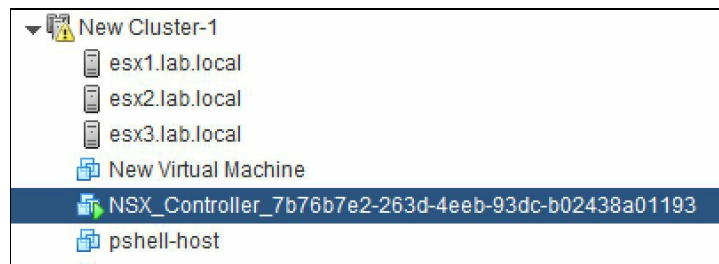
NSX Controller nodes				
Controller IP Address	ID	Status	Software Version	NSX Manager
192.168.1.210	controller-2	Deploying	6.2.44354	 192.168.1.201

8. You can also monitor the controller node progress in the Recent Tasks pane:



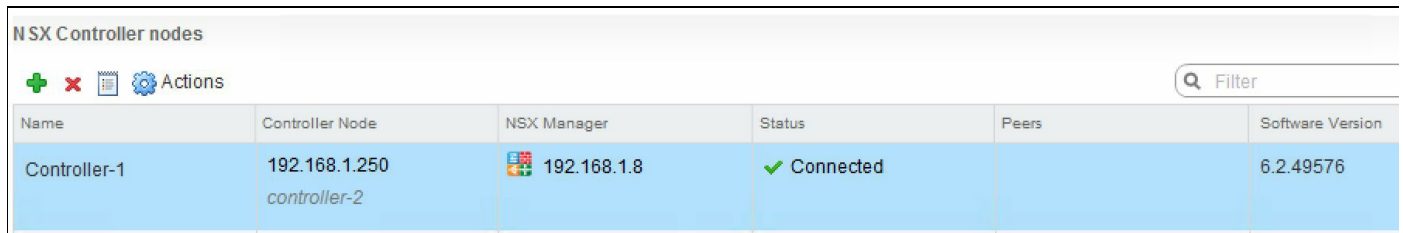
Task Name	Target	Status
Deploy OVF template	NSX_Controller_37...	64 %


9. Once it is deployed you will see the NSX controller in the host as shown in the following screenshot. NSX Manager will continue to power on the controller node once the deployment is complete:



New Cluster-1
esx1.lab.local
esx2.lab.local
esx3.lab.local
New Virtual Machine
NSX_Controller_7b76b7e2-263d-4eeb-93dc-b02438a01193
pshell-host

10. The controller status in NSX Manager gets updated once the controller node is powered on and ready to use:

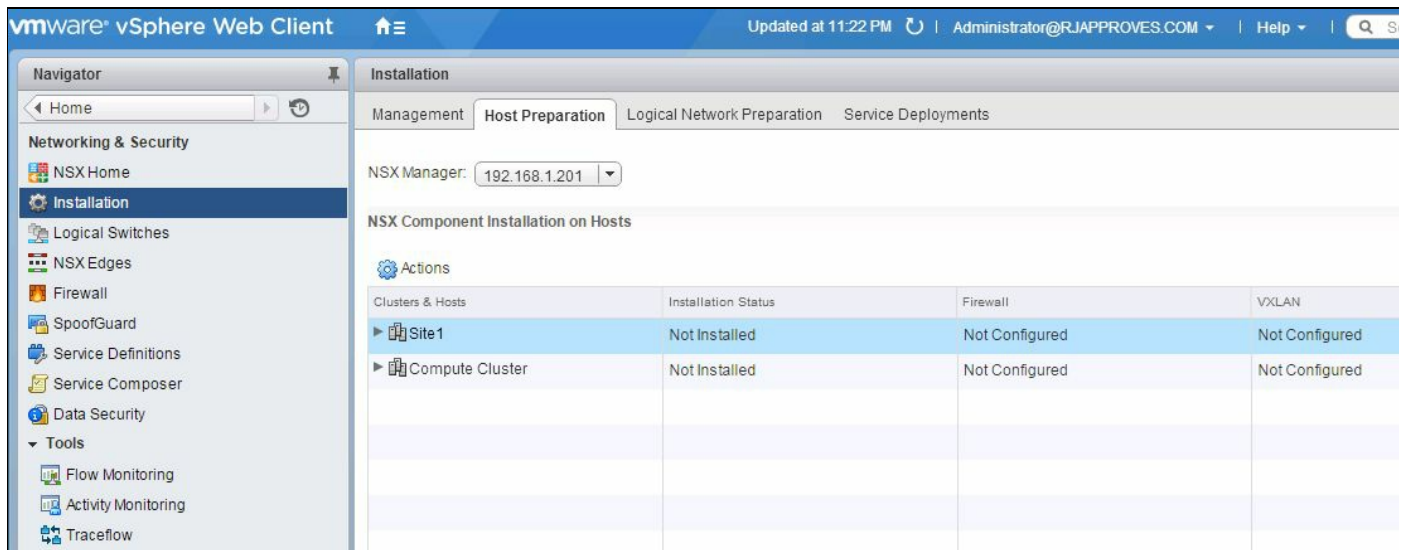


Name	Controller Node	NSX Manager	Status	Peers	Software Version
Controller-1	192.168.1.250 <i>controller-2</i>	 192.168.1.8	✓ Connected		6.2.49576

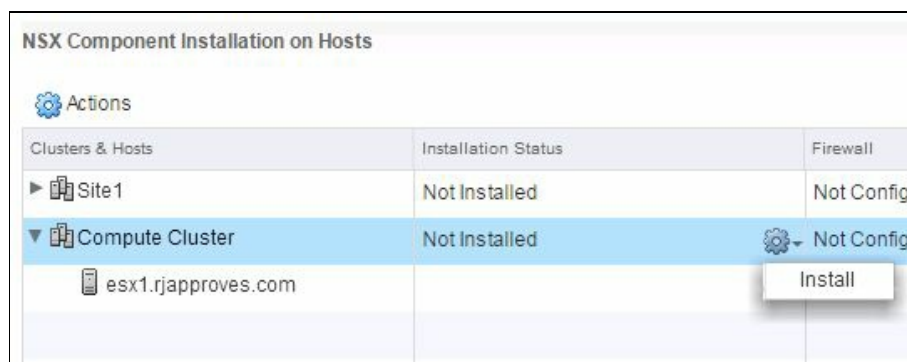
Deploying the data plane

Deploying the data plane involves preparing the ESXi hosts that enable them to participate in network virtualization. To prepare your environment for network virtualization, network components must be installed on a per-cluster basis. Any new hypervisors added to this cluster will automatically be prepared for network virtualization. After the network virtualization software (installed as a VIB) is installed, the logical firewall is enabled on that cluster. One important prerequisite is to ensure all hosts in the cluster are part of the distributed virtual switch:

1. To prepare a cluster go to the Networking & Security | Installation | Host Preparation tab:



2. Pick your cluster that you want to prepare for NSX component installation on the host. Hover over the Installation Status column and you will find the Install option:



3. Click Install. Click Yes in the prompt to proceed. You will now see the Installing

status.



Although not necessary, in rare instances or a reinstall a ESXi host reboot will be required.

- Once the installation is complete you will see the status updated and the distributed firewall enabled.
- We will learn more about configuring VXLAN in [Chapter 4](#), NSX Functional Services:

NSX Component Installation on Hosts			
Actions			
Clusters & Hosts	Installation Status	Firewall	VXLAN
▼ New Cluster-1	✓ 6.2.7.5343628	✓ Enabled	✓ Configured
esx3.lab.local	✓ 6.2.7.5343628	✓ Enabled	
esx1.lab.local	✓ 6.2.7.5343628	✓ Enabled	
esx2.lab.local	✓ 6.2.7.5343628	✓ Enabled	

- All hosts in the Compute cluster will have three VIBs installed. The VIBs installed prepare the hosts in the cluster for logical routing, VXLAN, and the distributed firewall. You can log in to a host and run the following command to find the VIBs installed:

```
[root@host:~] esxcli software vib list | grep esx
```

```
esx-vsip 6.0.0-0.0.5325006  
areCertified 2017-05-03  
esx-vxlan 6.0.0-0.0.5325006  
areCertified 2017-05-03
```


Summary

We began this chapter by preparing our environment and understanding the different limitations of NSX in order to begin its deployment. We then went ahead and deployed NSX Manager using the vSphere web client. Remember that NSX cannot be managed by traditional Windows clients and can only be managed by the vSphere web client. We then had a quick overview of NSX Manager to familiarize ourselves with the management interface functionality. Next we went ahead and configured our NSX Manager with our vCenter server. This allows us to go ahead and start configuring NSX services for our vCenter. We managed our NSX Manager using the vSphere web client and continued to deploy our control plane. Lastly, we prepared our hosts for network virtualization by deploying the data plane that by default enables the distributed firewall service on all ESXi hosts in a cluster.

In the [Chapter 4](#), NSX Functional Services, we will go over the step-by-step deployment and configuration of multiple NSX features and services and configure logical switches, transport zones, and other NSX functions.

NSX Functional Services

In this chapter, we will go over the step-by-step deployment and configuration of multiple NSX features and services. In [Chapter 3, NSX Installation and Configuration](#), we got your environment ready by downloading and deploying NSX. We also went ahead and configured the NSX manager, built our control plane, and deployed the controller clusters, followed by preparing our hosts for network virtualization. This now leaves us with an environment where we can go ahead and configure primary and secondary NSX managers, transport zones, logical switching, L2 bridges, logical routers, and a firewall, and begin to look at an Edge services gateway as well.

In this chapter, we will cover the following topics:

- Primary and secondary NSX managers
- Transport zones
- Logical switching
- L2 bridges
- Edge Services Gateway
- Logical firewall

Primary and secondary NSX managers

In the previous chapters, we have discussed cross-vCenter NSX deployment. We will now look at configuring primary and secondary NSX managers in a multi-vCenter NSX deployment.



NSX 6.2 and later allows you to manage multiple vCenter NSX environments from a single primary NSX Manager.

Before we delve more deeply into primary and secondary NSX manager configuration, let's briefly go over the benefits of cross-vCenter NSX deployment.

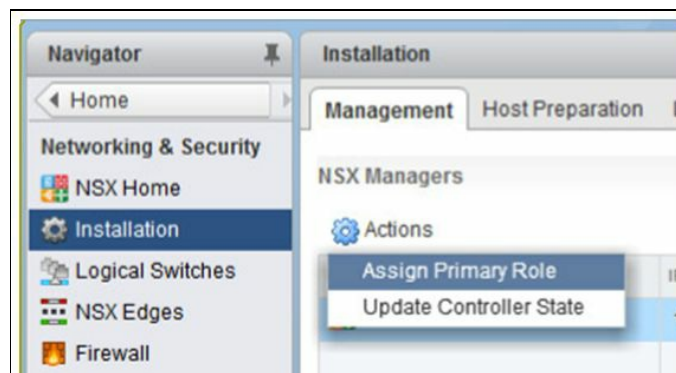
Benefits of cross-vCenter NSX

Cross-vCenter NSX environments can be managed centrally. This means that organizations with multiple vCenter instances and multiple NSX instances can manage them all with one central primary NSX instance. With a cross-vCenter environment, you can increase the span of your logical VXLAN networks. This means the same logical VXLAN network is now available across multiple vCenters. Virtual machines on multiple vCenters can now be part of the same logical VXLAN network without any additional network changes. This is powerful.

Cross-vCenter NSX environments also allow you to centralize security and policy management. Firewall rules applied on a universal logical network (a logical network that spans multiple vCenters) is enforced across all virtual machines across all vCenters. Even features such as vMotion between vCenters become easy, with no network changes.

There can only be one primary NSX manager, and up to seven secondary NSX managers. You can select one primary NSX manager, following which you can start creating universal objects and deploying universal controller clusters as well. The universal controller cluster will provide the control plane for the cross-vCenter NSX environment. Remember that in a cross-vCenter environment, the secondary NSX managers do not have their own controller clusters.

Before assigning a primary role to a NSX manager, ensure that it has the controller clusters deployed and the clusters prepared and configured:

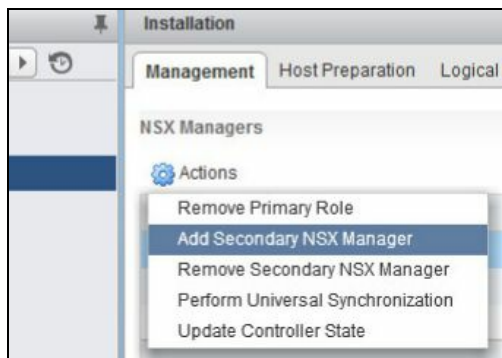


The commands shown in the screenshot assign the primary role to the NSX manager.

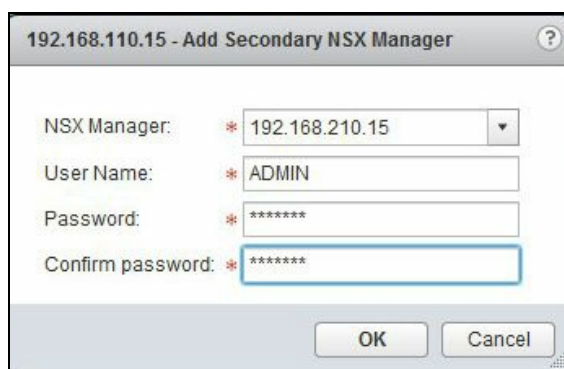
On the secondary NSX manager, make sure that no controller clusters are deployed in

the NSX manager that is going to be assigned the secondary role:

1. Log in to the vCenter linked to the primary NSX manager.
2. Go to Home | Networking & Security | Installation and select the Management tab.
3. Click on the primary NSX manager, then Actions | Add Secondary NSX Manager:



4. Enter the IP address, User Name, and Password of the secondary NSX manager and click OK:



5. Check the certificate, and upon registration you now have a secondary NSX manager configured.

Make sure both vCenters and NSX managers are able to communicate with each other. Make sure the following ports are opened between the vCenter environments:

Source	Target	Port	Protocol	Purpose
VXLAN Tunnel End Point	VXLAN Tunnel End Point	8472 (default before NSX 6.2.3) or 4789 (default in new installs of NSX 6.2.3 and	UDP	Transport network encapsulation

(VTEP)	(VTEP)	later)		between VTEPs
Primary NSX manager	Secondary NSX manager	443	TCP	Cross-vCenter NSX universal sync service
Primary NSX manager	vCenter Server	443	TCP	vSphere API
Secondary NSX manager	vCenter server	443	TCP	vSphere API
Primary NSX Manager	NSX universal controller cluster	443	TCP	NSX controller REST API
Secondary NSX manager	NSX universal controller cluster	443	TCP	NSX controller REST API
ESXi host	NSX universal controller cluster	1234	TCP	NSX control plane protocol
ESXi host	Primary NSX manager	5671	TCP	RabbitMQ
ESXi host	Secondary NSX manager	5671	TCP	RabbitMQ

Now that we have the primary and secondary managers configured, we can proceed to create universal objects in the primary NSX manager, which spans the multi-site cross-

vCenter environment.

We will discuss universal objects in the following sections as we discuss different aspects of NSX.

Configuring VXLAN

Before we jump into configuring transport zones, let's look at configuring VXLAN, which is part of preparing a cluster. VXLAN allows for layer 2 logical switching across hosts and can span multiple layer 3 domains. VXLANs are configured on a per-cluster basis, where each cluster that is to participate in NSX is mapped to a **virtual distributed switch (VDS)**. All hosts in that cluster are then enabled for logical switches.



If you only plan to use the distributed firewall feature of NSX, you do not need to configure VXLAN transport layer.

When you configure VXLAN networking, you must provide a vSphere distributed switch, a VLAN ID, an MTU size, an IP addressing mechanism (DHCP or IP pool), and a NIC teaming policy.

Jumbo frames are necessary, so the MTU for each switch must be set to 1550 or higher. By default, it is set to 1600.

To configure VXLAN, follow these steps:

1. Navigate to Home | Networking & Security | Installation and select the Host Preparation tab.
2. Click Not Configured in the VXLAN column:

NSX Component Installation on Hosts			
Actions			
Clusters & Hosts	Installation Status	Firewall	VXLAN
▼ New Cluster-1	✓ 6.3.2.5672532	✓ Enabled	Not Configured
esx3.lab.local	✓ 6.3.2.5672532	✓ Enabled	
esx1.lab.local	✓ 6.3.2.5672532	✓ Enabled	
esx2.lab.local	✓ 6.3.2.5672532	✓ Enabled	

3. Select a distributed virtual switch and a VLAN ID, if there is one. The MTU needs to be 1600, and ensure jumbo frames are enabled on the physical network. Select the desired NIC teaming policy:



The number of VTEPs is not editable in the UI. The VTEP number is set to match the number of dvUplinks on the vSphere distributed switch being prepared.

New Cluster-1 - Configure VXLAN Networking

Switch: * dvSwitch

VLAN: * 0

MTU: * 1600

VMKNic IP Addressing: * Use DHCP
 Use IP Pool New IP Pool...

VMKNic Teaming Policy: * Fail Over

VTEP: * 1

OK Cancel

- When you select the New IP Pool, a pop-up appears that allows you to configure an IP pool for your VXLAN IP addresses:

Add Static IP Pool

Name: * vxlan

Gateway: * 192.168.1.1
A gateway can be any IPv4 or IPv6 address.

Prefix Length: * 24

Primary DNS: 192.168.1.2

Secondary DNS:

DNS Suffix: lab.local

Static IP Pool: * 192.168.1.5-192.168.1.10


for example 192.168.1.2-192.168.1.100 or
abcd:87:87::10-abcd:87:87::20

OK Cancel

- Click OK when done.
- You will now see that the VXLAN is configured:

Actions			
Clusters & Hosts	Installation Status	Firewall	VXLAN
▼ New Cluster-1	✓ 6.3.2.5672532	✓ Enabled	✓ Configured
esx3.lab.local	✓ 6.3.2.5672532	✓ Enabled	⚙
esx1.lab.local	✓ 6.3.2.5672532	✓ Enabled	
esx2.lab.local	✓ 6.3.2.5672532	✓ Enabled	

7. You will also see a new distributed port group appear:



8. You will also see the VXLAN configuration status and the VMKnic on a per-host basis in the Logical Network Preparation tab:

Installation

Management Host Preparation **Logical Network Preparation** Service Deployments

NSX Manager: 192.168.1.8

VXLAN Transport Segment ID Transport Zones

VXLAN Port: 4789 Change

Clusters & Hosts	Configuration Status	Switch	VLAN	MTU	VMKnic IP Addressing	Teaming Policy	VTEP
▼ New Cluster-1	✓ Unconfigure	dvSwitch	0	1600	IP Pool	Fail Over	1
esx1.lab.local	✓ Ready				✓ vmk1: 192.168.1.18		
esx2.lab.local	✓ Ready				✓ vmk1: 192.168.1.19		
esx3.lab.local	✓ Ready				✓ vmk1: 192.168.1.17		

9. Clicking on Unconfigure will unconfigure the VXLAN from the cluster:

Clusters & Hosts	Configuration Status	Switch	VLAN	MTU	VMKnic IP Addressing
▼ New Cluster-1	✓ Unconfigure	dvSwitch	0	1600	IP Pool
esx1.lab.local	✓ Ready				✓ vmk1: 192.168.1.18
esx2.lab.local	✓ Ready				✓ vmk1: 192.168.1.19
esx3.lab.local	✓ Ready				✓ vmk1: 192.168.1.17

New Cluster-1 - Unconfigure cluster

Do you want unconfigure VXLAN on the selected Cluster 'New Cluster-1'?

In the simplest terms, a VXLAN VTEP is an additional IP added as a VMKernel address to a host. A VTEP is effectively a host with a VMKernel IP address assigned when the host is added to a transport zone.

Assigning a segment ID pool

Now that we have our cluster configured and the VXLAN functionality set up, it is time to set up the VXLAN segments. VXLAN segments are built between the VXLAN tunnel end points, which are hypervisors. Each VXLAN tunnel has a segment ID (VNI), and you must specify a segment ID pool for each NSX manager. All traffic will be bound to its segment ID, which allows for isolation.

For environments without a NSX controller, you must also add a multicast address range to spread traffic across your network. Specifying a range avoids overloading a single multicast address.

A range is specified when creating a VXLAN segment ID. Remember that segment ID has a 1:1 ratio with a logical network. The segment ID range controls the number of logical switches that can be created. Ensure that you pick the range that allows you to easily create logical switches as needed. You can always increase it if you ever max out.



You can have up to 16 million potential segment IDs. You must, however, not configure more than 10,000 VNIs in a single vCenter because vCenter limits the number of dvPortgroups to 10,000.

In a cross-vCenter deployment, make sure your segment IDs (VNIs) do not overlap. In a single NSX manager/vCenter environment, non-overlapping segment IDs are enforced; however, in a cross-vCenter environment, we have to do this manually.

A transport zone not using a controller cluster (unicast) will need to use the multicast or hybrid replication mode, in which case you must add a multicast address range. We will learn a bit more about transport zones in the next section.



Do not use `239.0.0.0/24` or `239.128.0.0/24` as the multicast address range because these networks are used for local subnet control, causing the physical switches to flood all traffic that uses these addresses.

When VXLAN multicast and hybrid replication modes are configured, a copy of multicast traffic is delivered only to hosts that have sent IGMP join messages. To avoid the physical network flooding all multicast traffic to all hosts, we need to make sure that the underlying physical switch is correctly configured with IGMP snooping and an

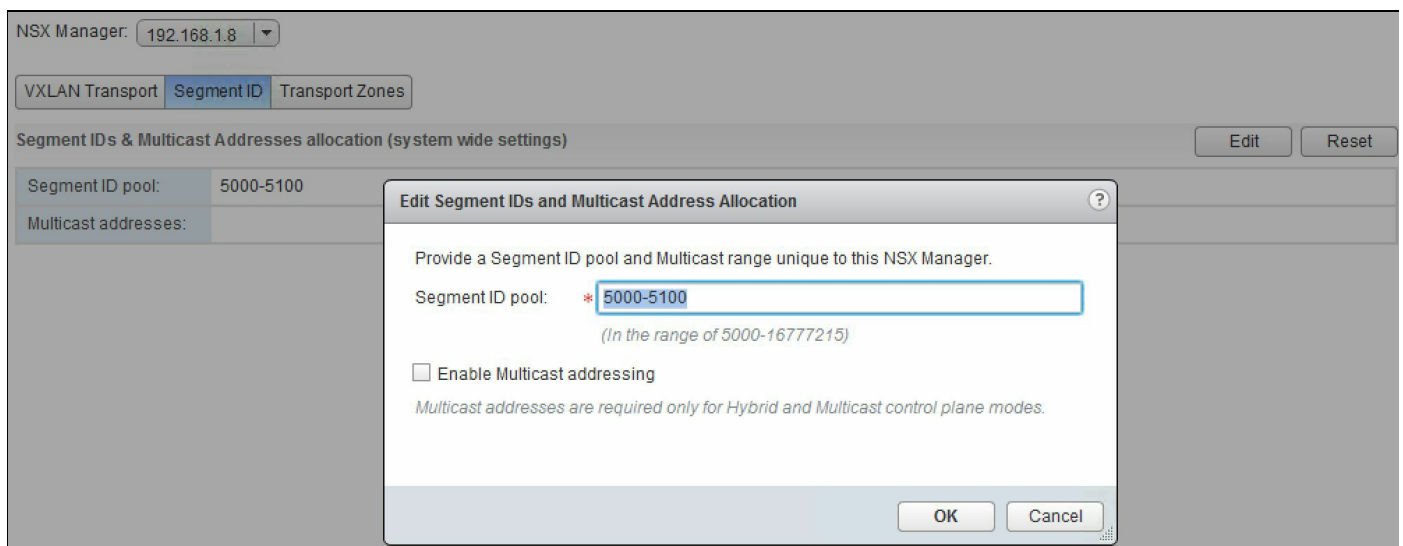
IGMP querier in network segments that carry VTEP traffic.



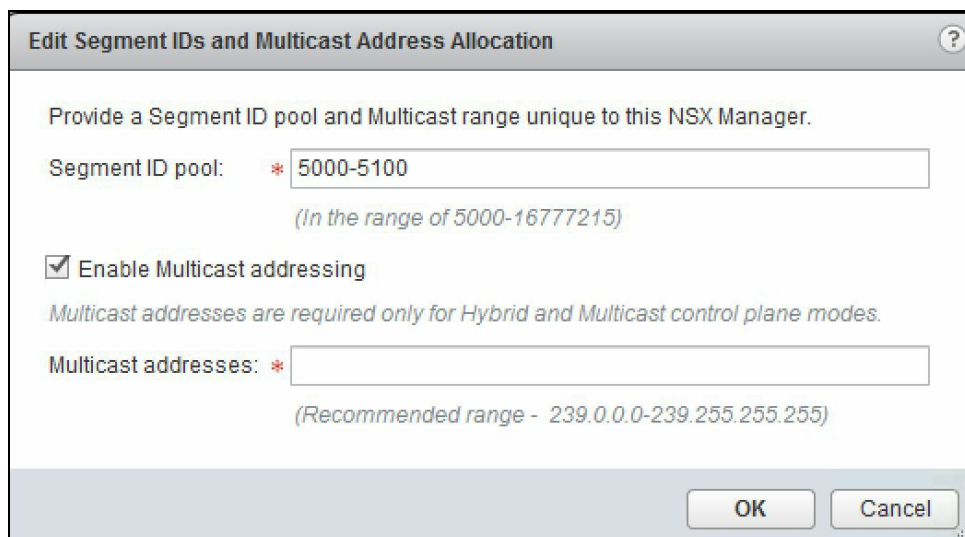
The recommended multicast address range starts at 239.0.1.0/24 and excludes 239.128.0.0/24.

To configure the segment ID, follow these steps:

1. Navigate to Home | Networking & Security | Installation and select the Logical Network Preparation tab.
2. Click Segment ID | Edit:



3. Multicast addressing is optional. If you are not deploying any controllers and prefer multicast addressing, select this option and enter a multicast address range:



4. Click OK to complete.

Transport zones

We have briefly looked at transport zones in the previous chapters. A transport zone controls the domain of a logical switch among the hosts. In other words, it controls which hosts a logical switch can reach. A transport zone is configured on a per-cluster basis and can span multiple clusters in a vCenter. A universal transport zone can span multiple clusters across multiple vCenters. A transport zone dictates which host and, by extension, which virtual machines are allowed to participate in a particular network. In a typical environment, there can be more than one transport zone that is mapped to a host or to a cluster. However, a logical switch can only belong to one transport zone.

If a virtual machine belongs to a different transport zone, you will not be able to directly communicate with that virtual machine. This means that a vNIC is limited to spanning within the bounds of a transport zone. A virtual machine, however, can have multiple vNICs, each belonging to a different transport zone.

In a cross-vCenter NSX environment, you can create a universal transport zone that includes clusters from any vCenter in the entire environment, thereby extending your logical network. However, you can only create one universal transport zone.

A universal transport zone is created by the primary NSX manager and is synchronized across all the secondary NSX managers. A universal logical switch associated with a universal transport zone can extend to one or more vSphere clusters across multiple vCenters. There can only be one universal transport zone.

To add a transport zone, follow these steps:

1. Go to Home | Networking & Security | Installation | Logical Network Preparation tab | Transport Zones.
2. Click on the + icon to add a new transport zone. To add a universal transport zone in a cross-vCenter NSX environment, you have to select the primary NSX manager:

Mark this object for Universal Synchronization

Name:

Description:

Replication mode:

Multicast
Multicast on Physical network used for VXLAN control plane.

Unicast
VXLAN control plane handled by NSX Controller Cluster.

Hybrid
Optimized Unicast mode. Offloads local traffic replication to physical network.

Select clusters that will be part of the Transport Zone

Name	NSX vSwitch	Status
<input checked="" type="checkbox"/> Compute Cluster	Site1-DvSwitch	<input checked="" type="checkbox"/> Normal

OK Cancel

3. Enabling Mark this object for Universal Synchronization allows this to be a universal transport zone.

4. Name the transport zone appropriately.

5. Select a Replication mode:

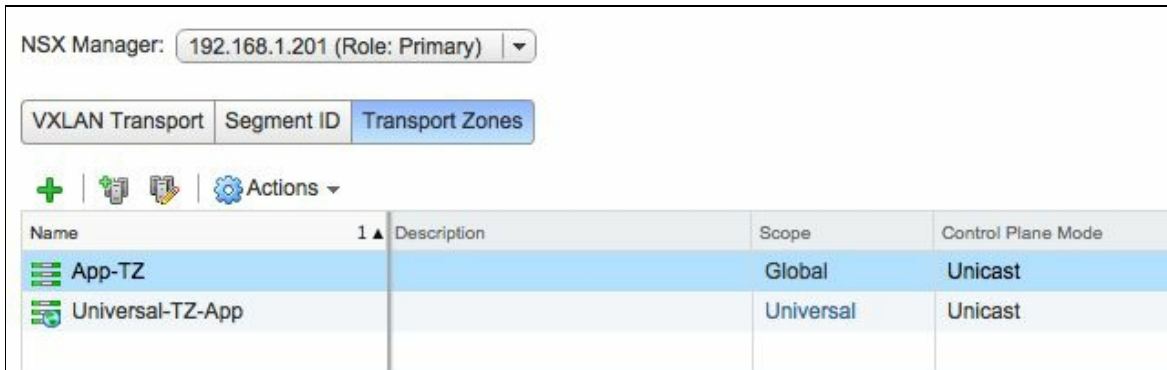
- Multicast: Instead of deploying controller clusters, you can use multicast IP addresses in the physical network for your control plane in this mode. This mode requires PIM/IGMP configured in your physical network.
- Unicast: The control plane is managed by NSX by means of the controller clusters. No changes are needed in the physical network.
- Hybrid: Hybrid mode offloads any local traffic replication to the physical network by means of multicast. This requires IGMP snooping to be configured in your physical network.



Remember to watch out for overlapping multicast addresses in a cross-vCenter NSX environment. The recommended multicast address range starts at 239.0.1.0/24 and excludes 239.128.0.0/24. 239.128.0.0/24 should not be used as this range is used for local subnet control causing the physical

switch to flood all traffic using that address range.

6. Select the clusters you want this transport zone to span to. A transport zone will remain local to the NSX Manager it was created in. A universal transport zone will span to all NSX environments in a cross-vCenter NSX deployment. Click OK when done:



Name	Description	Scope	Control Plane Mode
App-TZ		Global	Unicast
Universal-TZ-App		Universal	Unicast

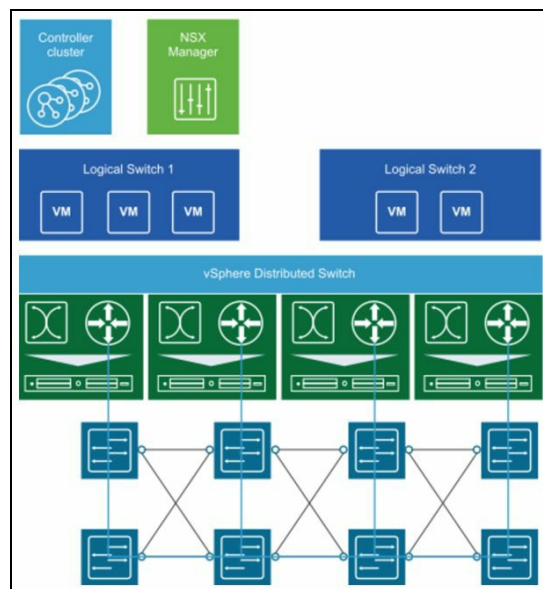
7. The transport zone is now created. Notice the scope and the icon difference between a Global and a Universal transport zone.

You can double-click on a transport zone to manage its settings. Alternatively, you can use the Actions drop-down to change the settings. You can also add or remove any clusters from the transport zone membership as required.

Logical switching

Logical switches have a similar functionality to physical switches—they permit the isolation of your applications and tenants for security and for other purposes. A logical switch creates a broadcast domain to allow the isolation of virtual machines. A logical switch is mapped to a unique VXLAN, which encapsulates the virtual machine traffic and carries it over the physical IP network.

Using logical switches, you can create VLANs in a similar way to a physical switch that spans large compute clusters. This now allows you to migrate your virtual machines using vMotion without any limitations of the physical network. This also allows you to have more control of your virtual network without having to alter or modify the physical network constructs. The logical switch is mapped to a VXLAN that encapsulates the traffic over the physical network. We discussed VXLAN and its packet architecture in [Chapter 2, NSX Core Components](#):



Picture courtesy of VMware

Logical switches are local to a virtual center NSX deployment; however, you can create a universal logical switch that can span multiple vCenters and effectively can extend your VXLAN across two different sites.

The NSX controller maintains the information about all virtual machines, hosts, logical switches, and its VXLANs.

Here are the prerequisites for creating a logical switch:


- vSphere distributed switches must be configured. You cannot deploy logical switches on standard switches.
- NSX controllers must be deployed.
- Your compute host clusters must be prepared/configured in NSX and ready to go.
- VXLAN must be configured.
- A transport zone and a segment ID pool must be configured.

Let's also look at some prerequisites for creating a universal logical switch for a multi-vCenter NSX environment:

- vSphere distributed switches must be configured on both sides
- NSX manager and NSX controllers must be deployed on both vCenters
- On both sites, hosts must be prepared and a VXLAN must be configured
- A primary NSX manager must be assigned
- A universal segment ID pool and a universal transport zone must be configured

To configure a logical switch, open up your vSphere web client and navigate to Home | Networking & Security | Logical Switches.

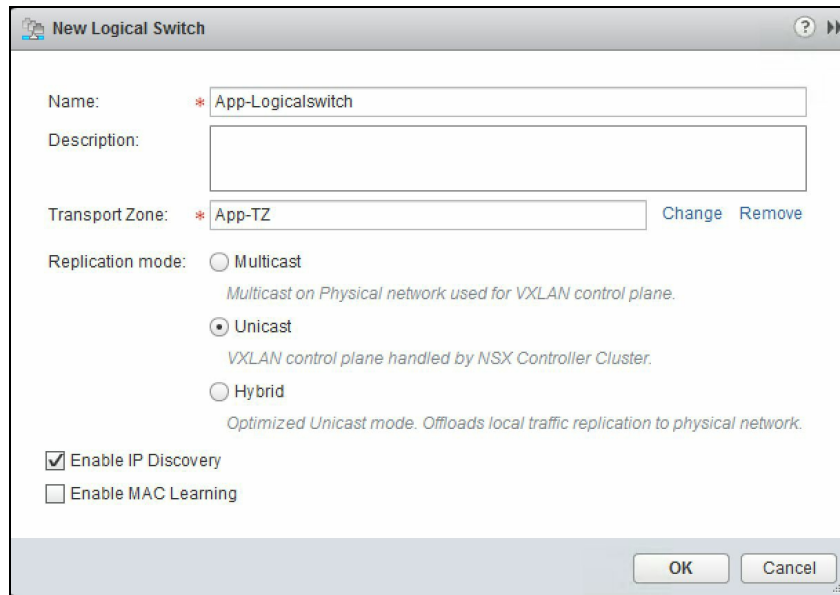
From the NSX Manager drop-down menu, select the NSX manager where you want to set up your logical switch. To create a universal logical switch, you must select the primary NSX manager. If enhanced vCenter linked mode is configured, you will be able to select and manage multiple NSX managers from any vCenter; however, one NSX manager can only be associated with one vCenter:



Click the + sign to add a new logical switch:



Due to lab restrictions, I will be switching between VMware HOL and my custom-built lab. This may cause you to see different names in the screenshots. The detailed explanation should help avoid any confusion.



1. Enter a Name for your logical switch.
2. Enter a Description (optional).
3. Select a Transport Zone. If you select a universal transport zone, this will create a universal logical switch. In this example, I selected a transport zone that is local to my vCenter.
4. By default, the logical switch inherits the control plane replication mode set for the transport zone. You can change that by selecting one of the available modes. For a universal logical switch, ensure that your multicast addresses do not conflict with any other multicast addresses across any NSX manager in the environment.
5. Enable IP discovery is enabled by default and allows for ARP suppression between VMs connected to the same logical switch. There should not really be any reason to disable this (optional).
6. Select enable MAC learning if your virtual machines have multiple MAC addresses or use virtual NICs that are trunking VLANs. This setting builds a VLAN/MAC pairing table on each vNIC.

You have now created your logical switch:


Logical Switches			
NSX Manager: 192.168.1.8			
Segment ID	Name	Status	Transport Zone
5000	App-Logicals witch	Normal	App-TZ

Now that we have created our logical switch, let's add some virtual machines to this

switch. Adding virtual machines to this switch will allow them to communicate with other instances that are part of the logical switch, similarly to how traditional networking works. You can connect virtual machines to a logical switch or to a universal logical switch:

1. Select the logical switch that you have created:

Segment ID	Name	Status	Transport Zone
5000	App-Logicswitch	Normal	App-TZ

2. Click on the  button to add virtual machines to this logical switch:

App-Logicswitch - Add Virtual Machines

1 Select Virtual Machines

Select Virtual Machines

Select VMs to connect to this network

Select from the available objects list on the left and move it to the selected objects list by double clicking the object or using the arrow key.

Available Objects (6 items): testvm, New Virtual Machine, pshell-host, Win2012-Linkedclone, Win2012-base, CentOS

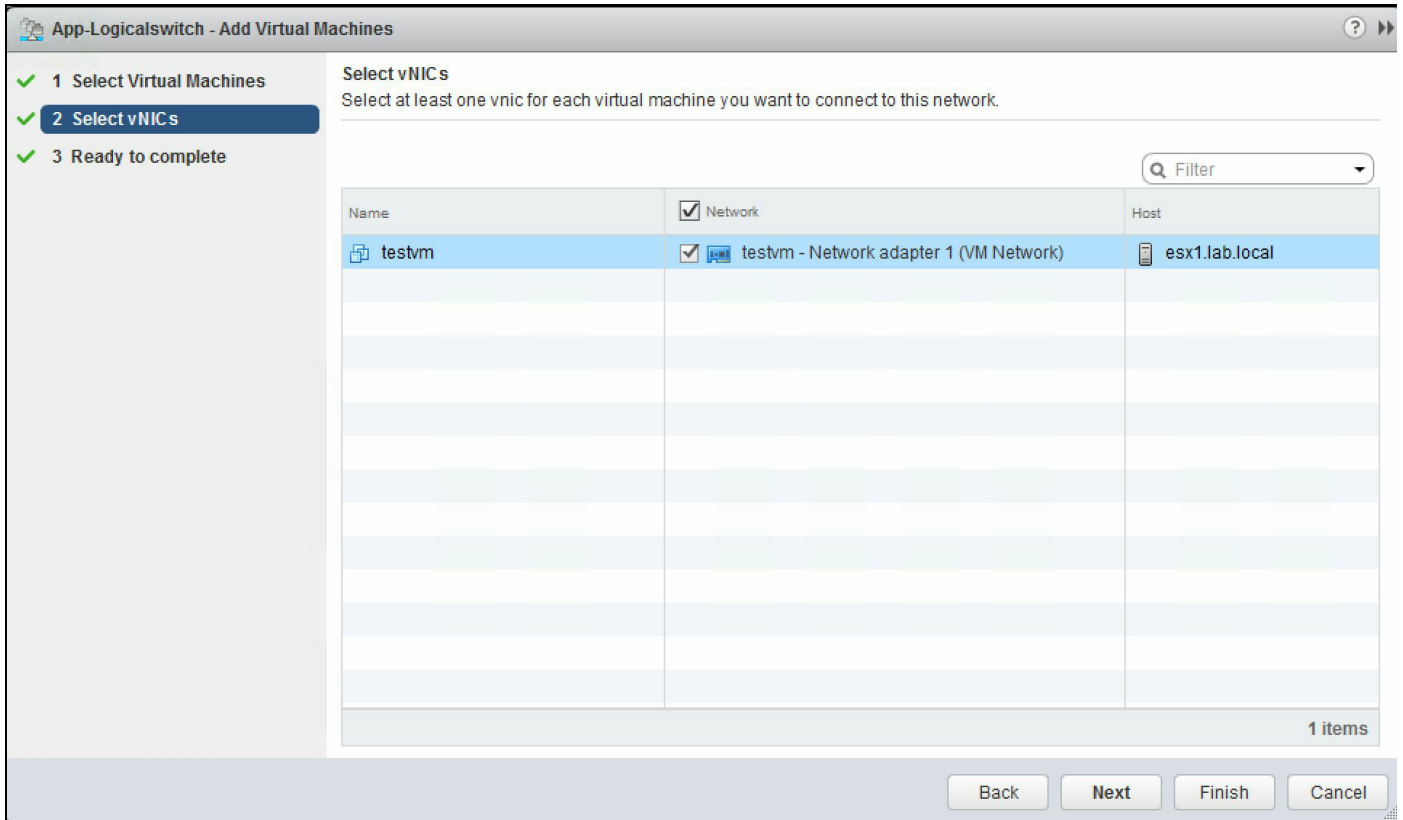
Selected Objects (1 item): testvm

Buttons: Back, Next, Finish, Cancel

3. Select the virtual machines from the list and click on Next. Notice that your virtual machines managed by the vCenter are listed here. For a multi-vCenter

environment, you need to add virtual machines from each vCenter by selecting the appropriate NSX manager.

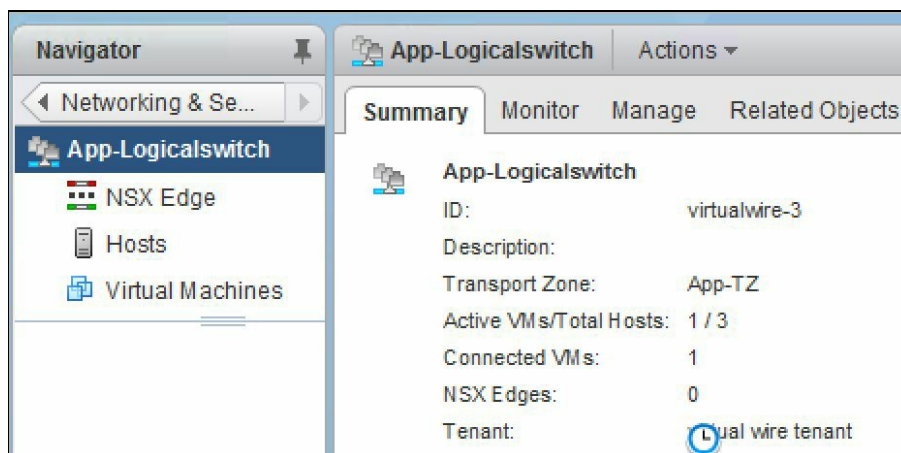
4. Select the vNIC for the virtual machine that needs to be part of the logical switch network. You may have a scenario where a virtual machine has more than one vNIC that is part of multiple logical switches. Click on Next when done and click Finish:



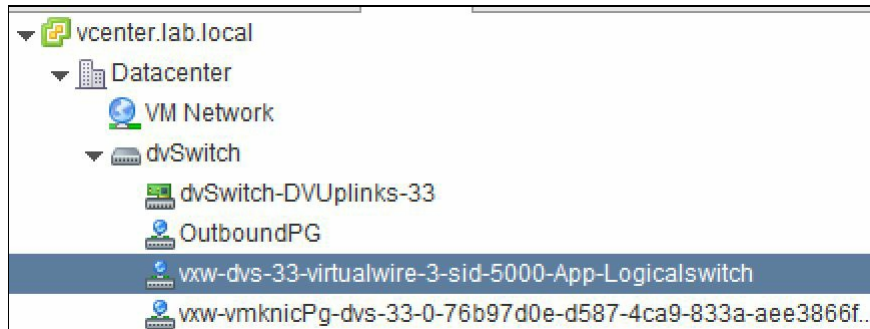
You have now connected a virtual machine to this logical switch. You can also associate a virtual machine with a logical switch by simply editing its settings and changing its vNIC association in vCenter.

You can review all the virtual machines connected to a logical switch by double-clicking the logical switch.

Go back to the logical switches view and double-click the logical switch we created. You will see a logical switch overview. You will notice that the virtual machines count is 1. The number of hosts this logical switch is configured on is also listed:

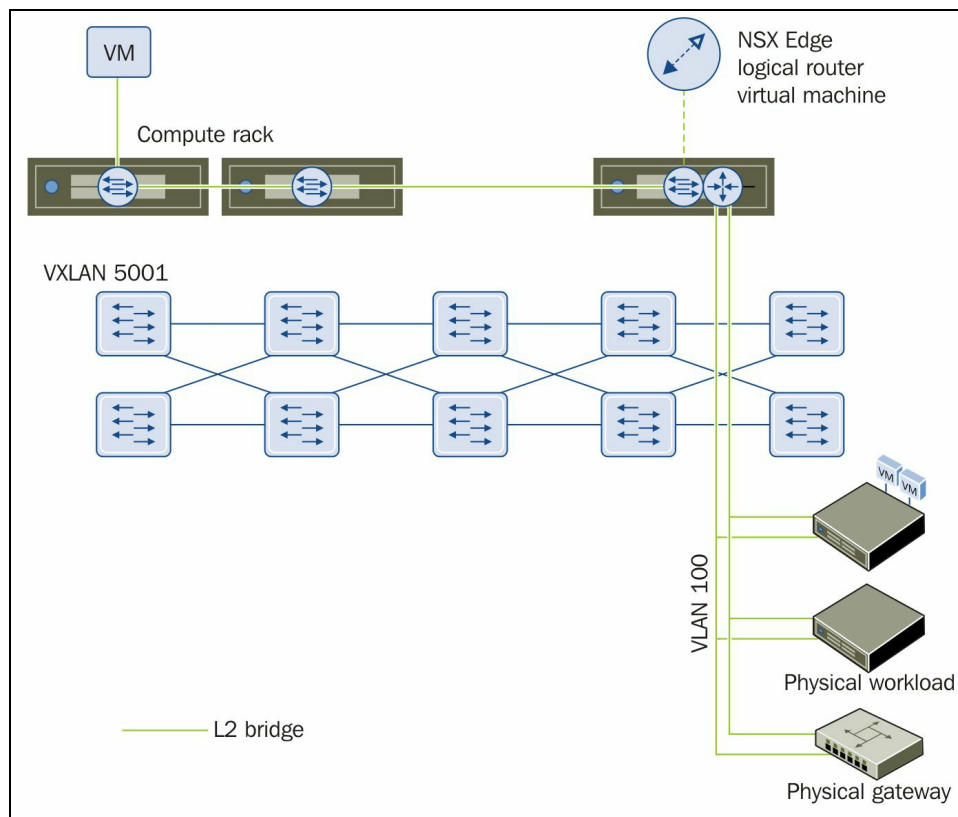


You will also see a new distributed port group has been created, which is the logical switch. You can connect virtual machines directly to this distributed port group as well:



L2 bridges

A logical switch can be connected to a physical switch VLAN by means of an L2 bridge. This allows you to extend your virtual logical networks to access existing physical networks by bridging the logical VXLAN with the physical VLAN. This L2 bridging is accomplished by means of an NSX Edge logical router that maps to a single physical VLAN on the physical network. However, L2 bridges should not be used to connect two different physical VLANs or two different logical switches. You also cannot use a universal logical router to configure bridging, and a bridge cannot be added to a universal logical switch. This means that in a multi-vCenter NSX environment you cannot extend a logical switch to a physical VLAN at another datacenter by means of L2 bridging:



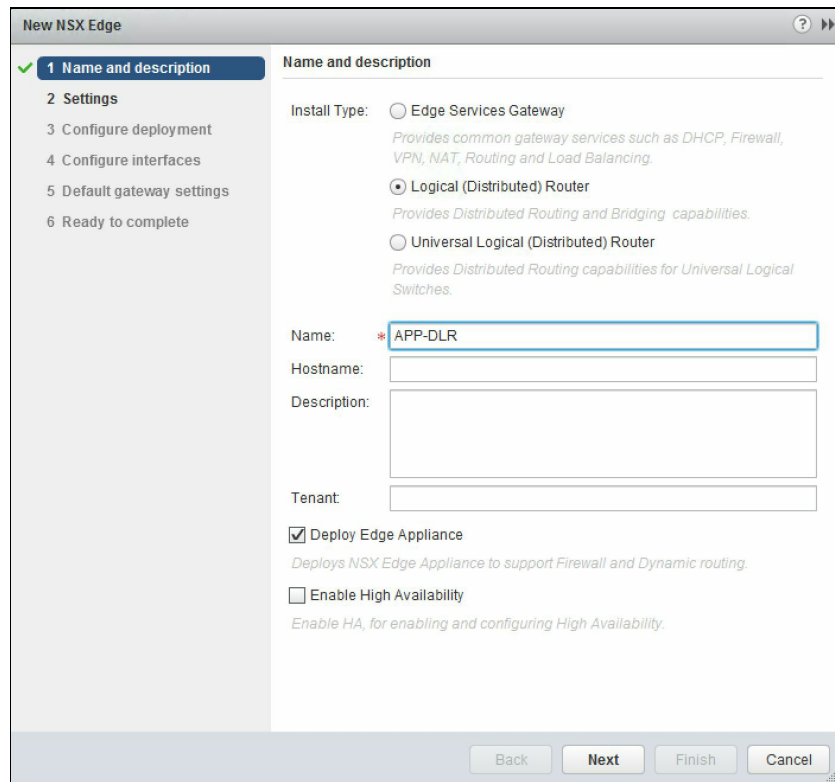
If **high availability (HA)** is configured for the NSX Edge router, the host on which the secondary Edge appliance is living must also be connected to the physical VLAN to allow for seamless bridge failover during a primary Edge failure event.

To add an L2 bridge, an NSX Edge logical router must be deployed.

Deploying an NSX Edge logical router

To deploy an NSX Edge logical router, let's perform the following set of steps:

1. Go to Home | Networking & Security | NSX Edge.
2. Select the appropriate NSX manager and click the + icon.
3. Select Logical (Distributed) Router. You cannot use a universal logical router to configure bridging:



The screenshot shows the 'New NSX Edge' configuration wizard. The left sidebar lists the steps: 1 Name and description (selected), 2 Settings, 3 Configure deployment, 4 Configure interfaces, 5 Default gateway settings, and 6 Ready to complete. The main area is titled 'Name and description' and contains the following options and fields:

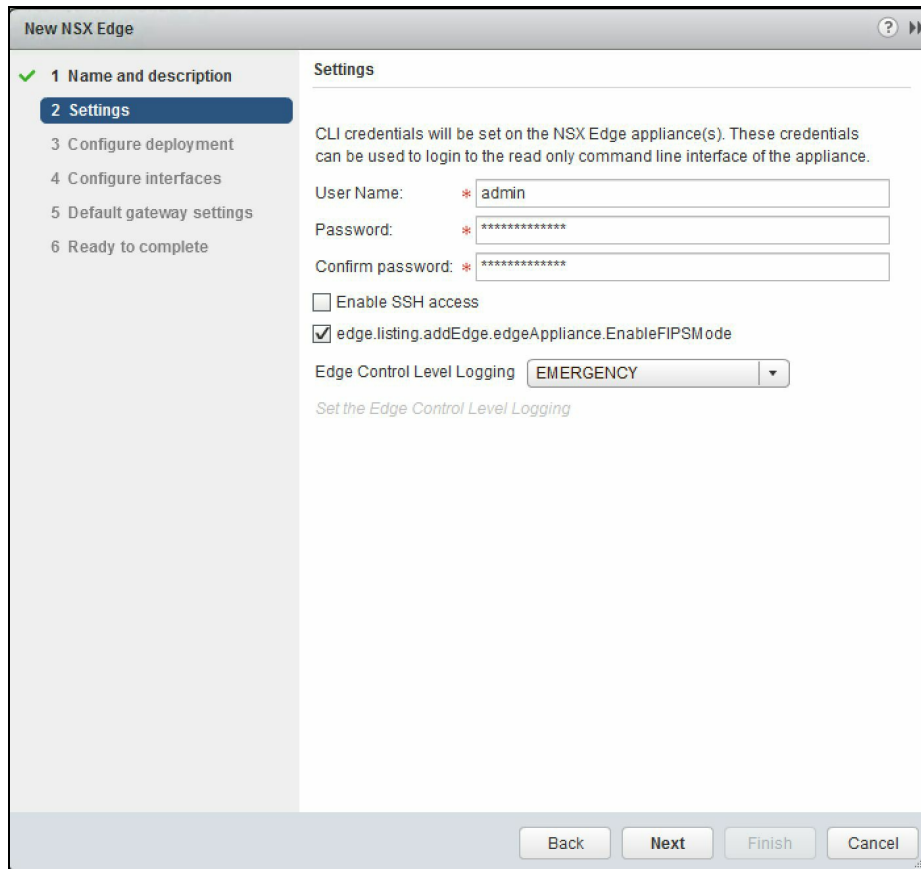
- Install Type:** Three radio button options:
 - Edge Services Gateway: Provides common gateway services such as DHCP, Firewall, VPN, NAT, Routing and Load Balancing.
 - Logical (Distributed) Router: Provides Distributed Routing and Bridging capabilities.
 - Universal Logical (Distributed) Router: Provides Distributed Routing capabilities for Universal Logical Switches.
- Name:** * APP-DLR (text input field)
- Hostname:** (text input field)
- Description:** (text input field)
- Tenant:** (text input field)
- Deploy Edge Appliance: Deploys NSX Edge Appliance to support Firewall and Dynamic routing.
- Enable High Availability: Enable HA, for enabling and configuring High Availability.

At the bottom, there are four buttons: Back, Next, Finish, and Cancel.

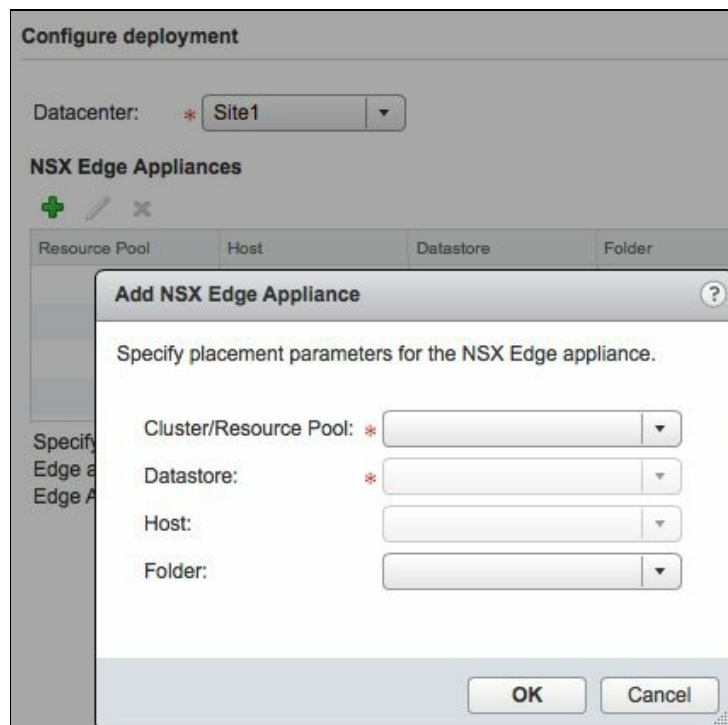
4. Check Enable High Availability to ensure uptime during a downtime scenario. Ensure the other hosts participating in NSX Edge HA have physical VLAN connectivity to allow for L2 bridging. Click Next.
5. Enter the `admin` password and set the logging level.
6. Click Next:



Enabling **Federal Information Processing Standard (FIPS)** is optional, and by default is disabled. When you enable FIPS mode, any secure communication to or from the NSX Edge uses cryptographic algorithms or protocols that are allowed by FIPS.

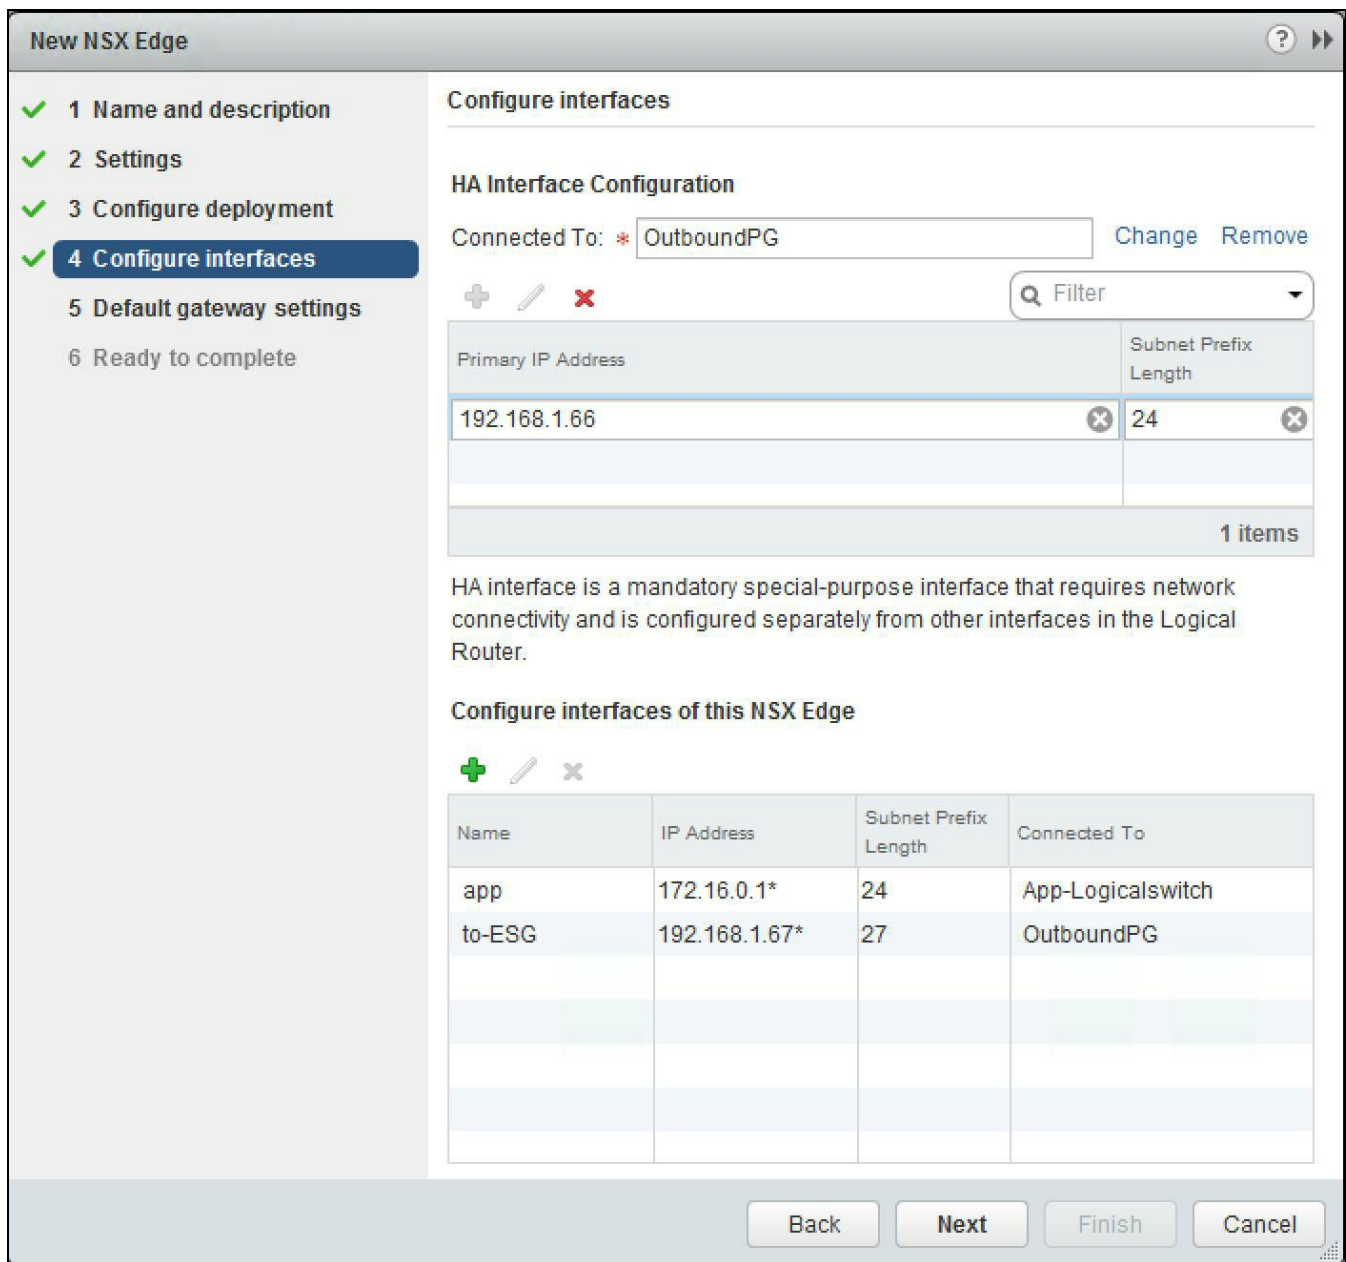


7. Select the Datacenter and click the + icon to specify the resource pool and datastore this appliance needs to be deployed on:



8. Next, we configure the HA interface configuration and the uplink interface for the

NSX Edge. The uplink interface should be connected to a physical VLAN-backed distributed port group to allow for L2 bridging. Alternatively, it can also be connected to an Edge gateway to route traffic for any north-south communication:



9. To configure interfaces, click on the + sign. Select the uplink for the Edge and enter the gateway IP for your network.

10. Click Next:

Edit Interface ?

Name: *

Type: Internal Uplink

Connected To: * [Change](#) [Remove](#)

Connectivity Status: Connected Disconnected

Configure subnets

+ ✎ ✖

Primary IP Address	Subnet Prefix Length
172.16.0.1	24

1 items

MTU:

11. Configure the default gateway and click Next:

New NSX Edge

- ✓ 1 Name and description
- ✓ 2 Settings
- ✓ 3 Configure deployment
- ✓ 4 Configure interfaces
- 5 Default gateway settings**
- 6 Ready to complete

Default gateway settings

Configure Default Gateway

vNIC: * app

Gateway IP: * 192.168.1.1

MTU: 1500

Admin Distance: 1

Back Next Finish Cancel

12. Review the summary and click Finish.
13. Monitor the progress in the Tasks pane:

NSX Edges

NSX Manager: 192.168.1.8 (Role: Primary)

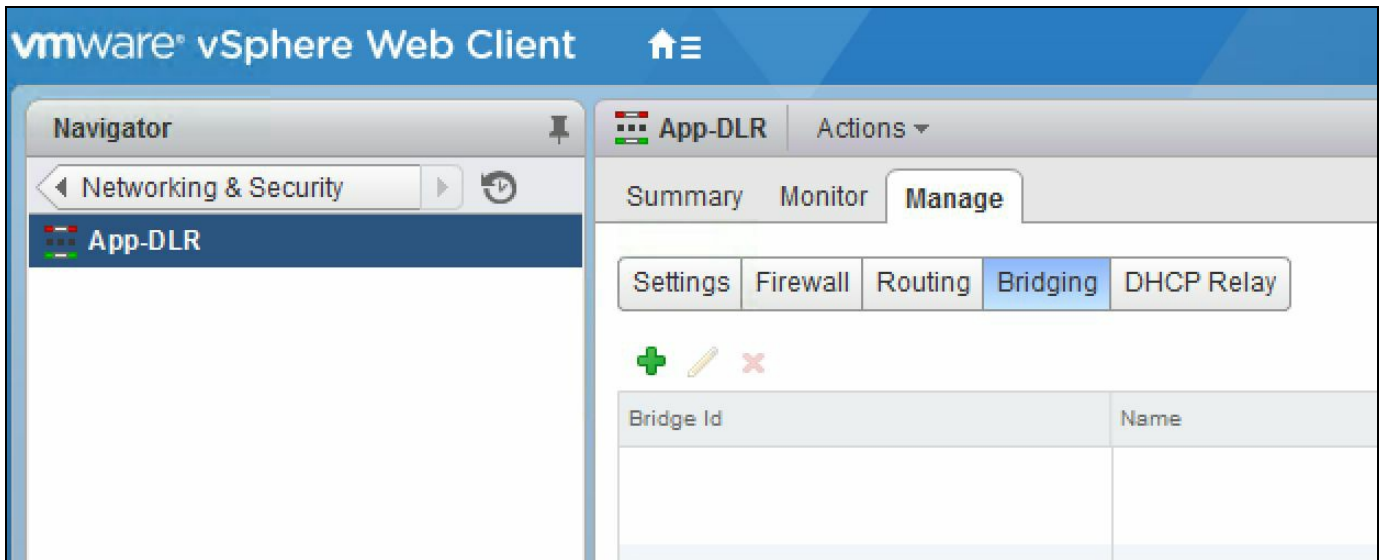
+ | Actions

1 Installing 4 Failed


Id	Name	Type	Tenant
edge-5	App-DLR	Logical Router	Default

Now that the NSX Edge-distributed logical router is deployed, let's configure this Edge logical router and enable the L2 bridging functionality:

1. Select the NSX Edge logical router and double-click on it.
2. Click on the Manage | Bridging | + icon:

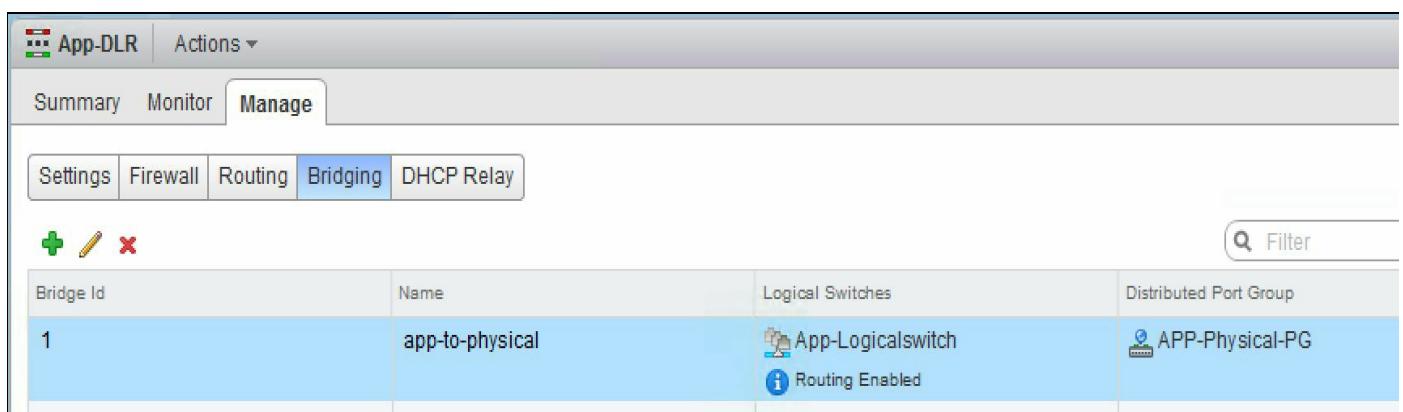


3. Name the bridge and select a Logical Switch and the Distributed Port Group to enable L2 bridging between them. Click OK:



The distributed port group must be a VLAN tagged port group. You cannot connect your Edge to a flat network distributed port group.

4. Click Publish Changes. Your L2 bridge is now set up:



Edge services gateway

In the previous section, we looked at deploying an NSX Edge distributed logical router. Let's look at adding an NSX Edge services gateway. You can always deploy multiple NSX Edge services gateway virtual appliances. Each appliance is configured with interfaces, and an Edge appliance can have up to 10 virtual interfaces (uplink and internal). The internal interfaces act as a gateway to all virtual machines in a port group, while the uplink interfaces can be connected to the outgoing network, which can be your physical network.

Only the enterprise administrator role, which allows for NSX operations and security management, can deploy an Edge services gateway:

1. Go to Home | Networking & Security | NSX Edges and click the + icon.
2. Select the Edge Services Gateway and enter a unique name for the appliance. The Edge services gateway will be deployed in high availability mode if Enable High Availability is selected. Click Next:

New NSX Edge ? >>

1 Name and description

2 Settings

3 Configure deployment

4 Configure interfaces

5 Default gateway settings

6 Firewall and HA

7 Ready to complete

Name and description

Install Type: Edge Services Gateway
Provides common gateway services such as DHCP, Firewall, VPN, NAT, Routing and Load Balancing.

Logical (Distributed) Router
Provides Distributed Routing and Bridging capabilities.

Universal Logical (Distributed) Router
Provides Distributed Routing capabilities for Universal Logical Switches.

Name: *

Hostname:

Description:

Tenant:

Deploy NSX Edge
Select this option to create a new NSX Edge in deployed mode. Appliance and interface configuration is mandatory to deploy the NSX Edge.

Enable High Availability
Enable HA, for enabling and configuring High Availability.

3. Enter the desired admin password. Enable SSH access is optional and only recommended for troubleshooting purposes. Enable auto rule generation is checked by default and allows the automatic creation of firewall rules, NAT, and routing configuration, which control traffic for some Edge services. Disabling this will cause you to manually add these rules and configurations. Click Next.
4. Select the appliance size depending on the size of your environment. You can also upgrade the appliance size after its deployment using the Convert to option.

Assign it to a resource pool or a cluster and to a datastore. Click Next:

New NSX Edge

- ✓ 1 Name and description
- ✓ 2 Settings
- 3 Configure deployment**
- 4 Configure interfaces
- 5 Default gateway settings
- 6 Firewall and HA
- 7 Ready to complete

Configure deployment

Datcenter: *

Appliance Size: Compact
 Large
 X-Large
 Quad Large

NSX Edge Appliances

+ ✎ ✕

Resource Pool	Host	Datastore	Folder
New Cluster-1		nfs-datastore	

Specifying a resource pool and datastore is mandatory for configuring the NSX Edge appliance.

5. Click + to configure interfaces to this Edge appliance. An Edge appliance must have at least one internal interface if configured for HA to work. An interface can have multiple overlapping subnets.

An interface can be configured with both IPv4 and IPv6 addresses and can have multiple non-overlapping subnets. If an interface has more than one IP address, you can select the primary IP address. Only one primary IP address is allowed per interface and the Edge uses that IP as the source address for locally generated traffic. Optionally, you can specify a MAC address for each IP you enter. You can also set the MTU if needed. You can select Enable Proxy ARP if you want the Edge gateway to answer ARP requests intended for the virtual machines. Enabling Send ICMP Redirect conveys the routing information to hosts:

Add NSX Edge Interface

vNIC#: 0

Name: * app-internal

Type: Internal Uplink

Connected To: App-Logicswitch Change Remove

Connectivity Status: Connected Disconnected

+ ✎ ✕

Primary IP Address	Secondary IP Addresses	Subnet Prefix Length
192.168.0.1 ✕		24 ✕

1 items

Comma separated lists of Secondary IP Addresses. Example: 1.1.1.1,1.1.1.2,1.1.1.3

MAC Addresses:

You can specify a MAC address or leave it blank for auto generation. In case of HA, two different MAC addresses are required.

MTU:

Options: Enable Proxy ARP Send ICMP Redirect

Reverse Path Filter:

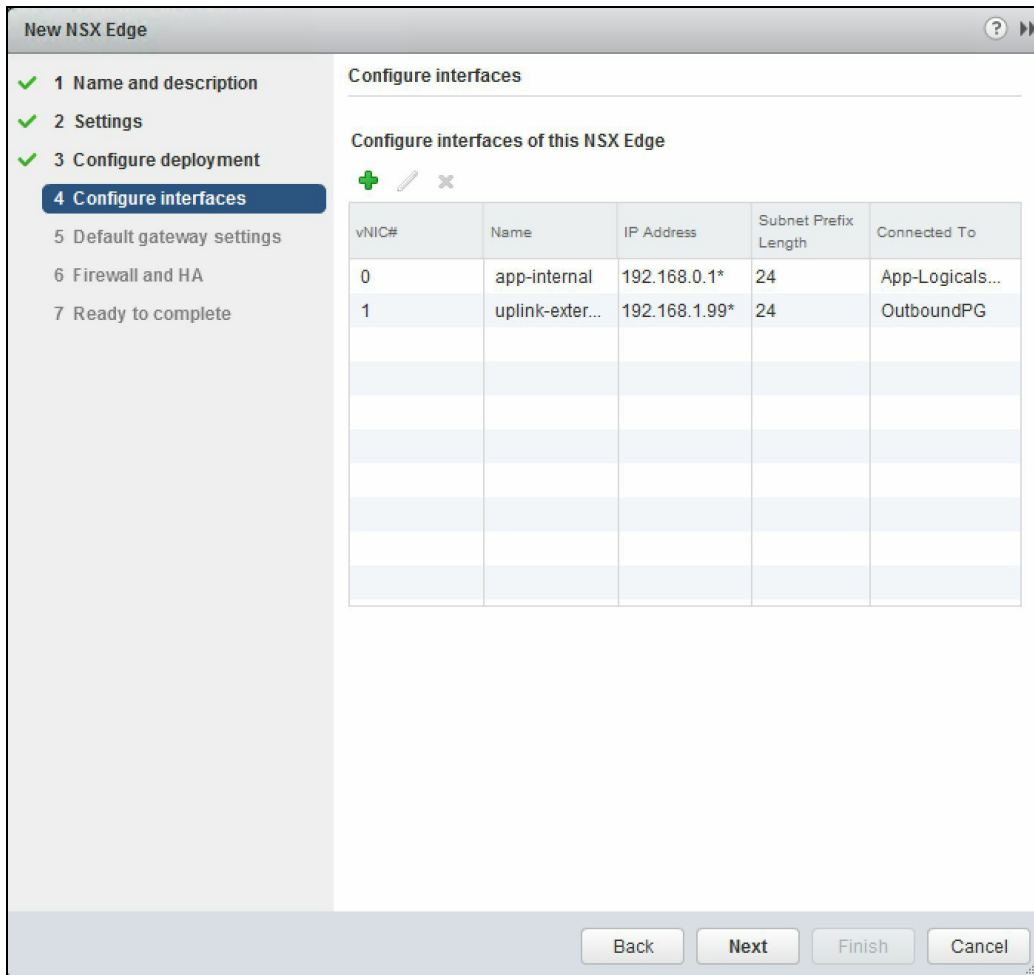
Fence Parameters:

Example: ethernet0.filter1.param1=1

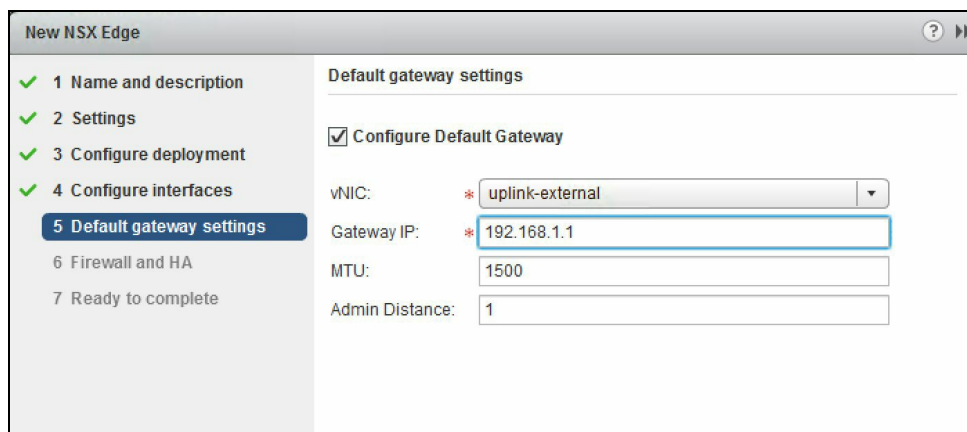
OK Cancel

Reverse Path Filter, when set to Enabled, verifies the reachability of the source address in packets being forwarded. In this mode, the packet must be received on the interface that the router would use to forward the return packet. It is enabled by default. In loose mode, the source address must appear in the routing table.

Configure fence parameters if you want to reuse IP and MAC addresses across different fenced environments. Click OK:

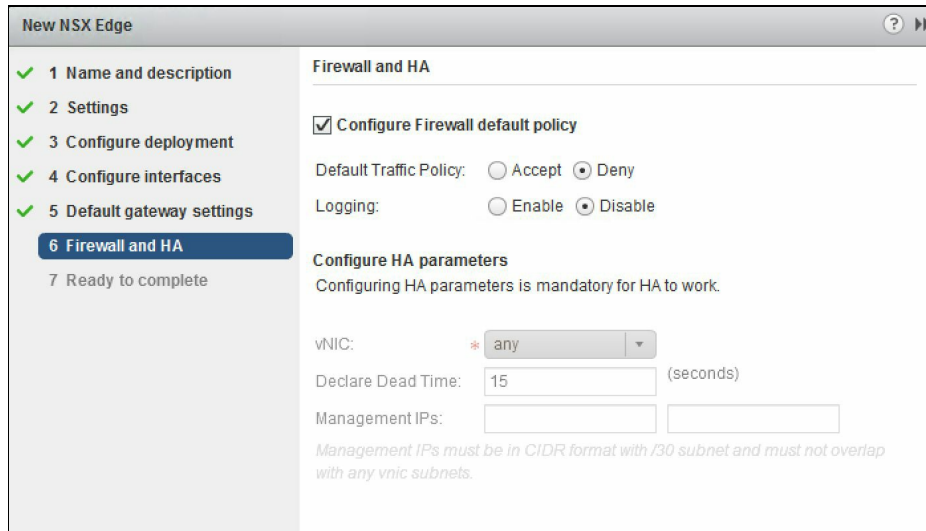


6. Click Next to enter the default gateway IP information. Click Next:



7. Select Configure Firewall default policy to change it to allow traffic that is legitimate. By default, the firewall will deny all traffic. You can enable logging, and logs are stored on the Edge appliance itself. If HA is enabled, it automatically chooses an internal interface and assigns it link-local IP addresses. You can also set the HA heartbeat timeout as well, which determines the time interval after which a failure of the Edge appliance is declared and HA action is initiated. Click

Next:



New NSX Edge

- 1 Name and description
- 2 Settings
- 3 Configure deployment
- 4 Configure interfaces
- 5 Default gateway settings
- 6 Firewall and HA**
- 7 Ready to complete

Firewall and HA

Configure Firewall default policy

Default Traffic Policy: Accept Deny

Logging: Enable Disable

Configure HA parameters

Configuring HA parameters is mandatory for HA to work.

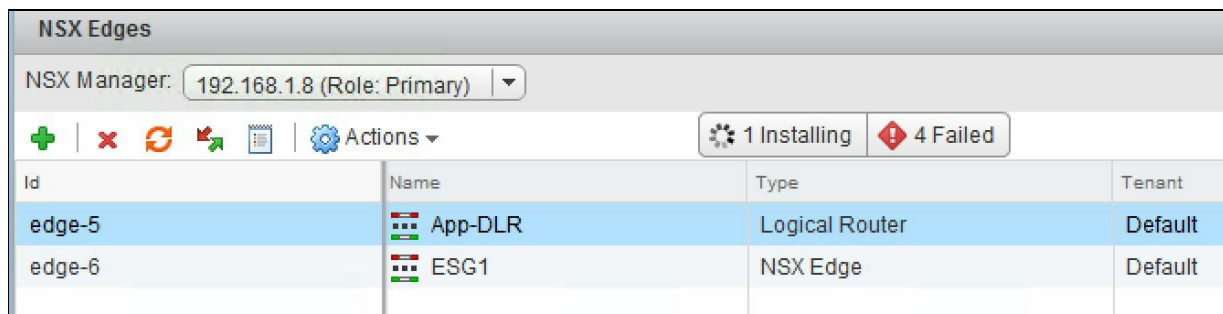
vNIC: * any

Declare Dead Time: 15 (seconds)

Management IPs:

Management IPs must be in CIDR format with /30 subnet and must not overlap with any vnic subnets.

8. Click Finish, and you will see an Edge appliance being deployed:



NSX Edges

NSX Manager: 192.168.1.8 (Role: Primary)

+ | x | ↻ | ↶ | ⚙ Actions

⚙ 1 Installing | ❌ 4 Failed

Id	Name	Type	Tenant
edge-5	App-DLR	Logical Router	Default
edge-6	ESG1	NSX Edge	Default

We will look at configuring Edge appliances with more features in the upcoming sections and chapters.

Logical firewalls

Logical firewalls are of two types: distributed firewall and Edge firewall. A distributed firewall is ideally deployed to protect any east-west traffic, while an Edge firewall protects any north-south traffic.

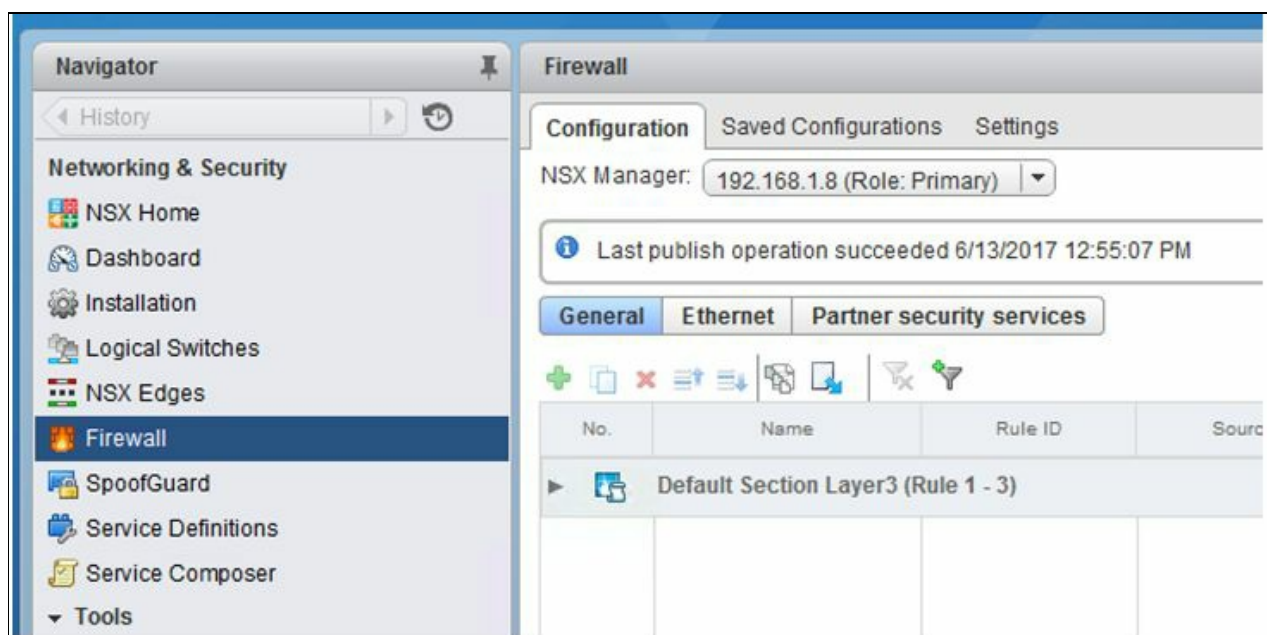


Server-to-server traffic is considered east-west, while client-server traffic is known as north-south.


The firewall rules UI allows you to add sections to separate firewall rules. Both L2 and L3 rules can have multiple sections that can be managed accordingly. For cross-vCenter environments, you must create a universal section before you can add the universal rules, and you must manage the universal rules from the primary NSX manager.

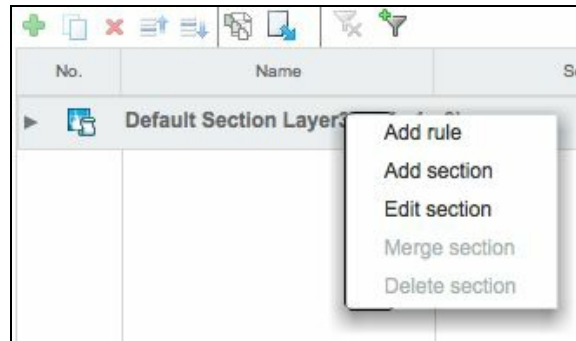
To add a firewall rule, follow these steps:

1. Go to Networking & Security | Firewall.
2. Ensure that you have selected the NSX Manager where you want to configure the rules. In a cross-vCenter environment, select the primary NSX manager to add universal firewall rules:

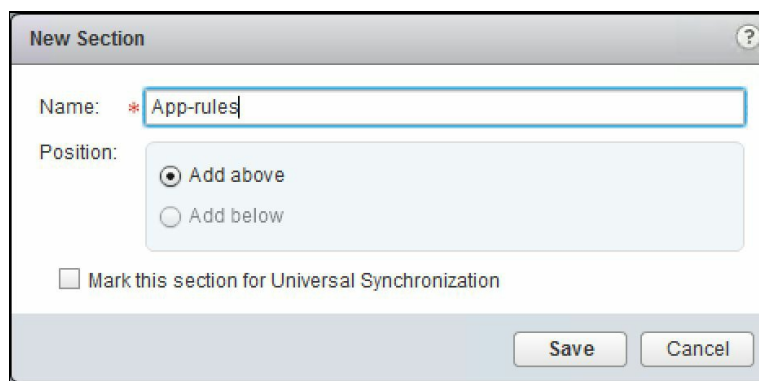


3. The General tab allows you to add L3 rules, while the Ethernet tab allows you to add L2 rules.

4. Right-click Default Section Layer3 (Rule 1 - 3) and click Add section to add a new section. You can even scroll to the right to find the  icon to add a new section. Sections allow us to group similar rules that apply to a particular task. For example, all rules that apply to an application could be part of the APP section:



5. Name the section appropriately:




6. When Mark this section for Universal Synchronization is enabled, it automatically marks this section as a universal section. You will only see this option when configuring rules on the primary NSX manager. Leaving this unchecked will only apply this section to the local domain. Click OK when complete.
7. Click Publish Changes.

Now that the firewall section has been created, let's go ahead and add some rules to it. Firewall rules are applied to multiple objects at the source and destination levels for each rule. There are many vCenter objects that the rules can be applied to, and some of them are clusters, virtual machines, vNIC, and IP addresses. Before we go ahead and configure some local firewall rules, make sure your virtual machines have VMware tools installed on them.

To add local firewall rules, follow these steps:

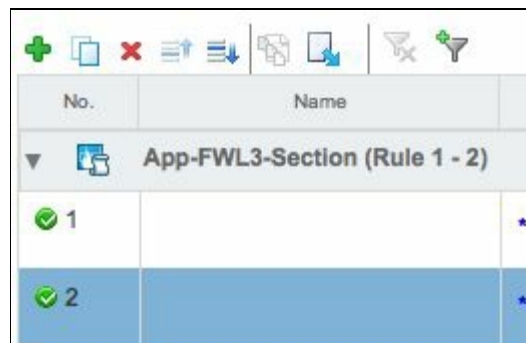
1. Go to Networking & Security | Firewall.
2. Remember that the General tab allows you to add L3 rules, while the Ethernet tab allows you to configure L2 rules. Click the + icon in the section where you want to add a new rule:



3. A new rule shows up with any rules under each section:





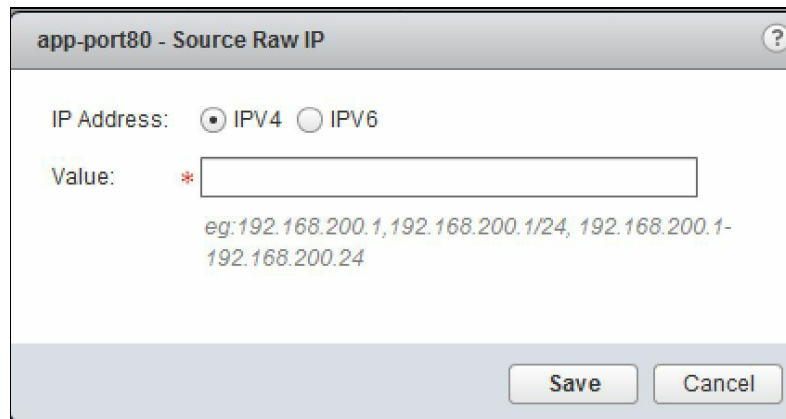
4. If you have multiple rules in a section and want to reorder the rules, highlight a rule and click the up or the down arrow icons to reorder it. The first matched rule is enforced in the firewall, so ordering your rules is important:




5. Point to the name of the rule and click the edit icon to name it:

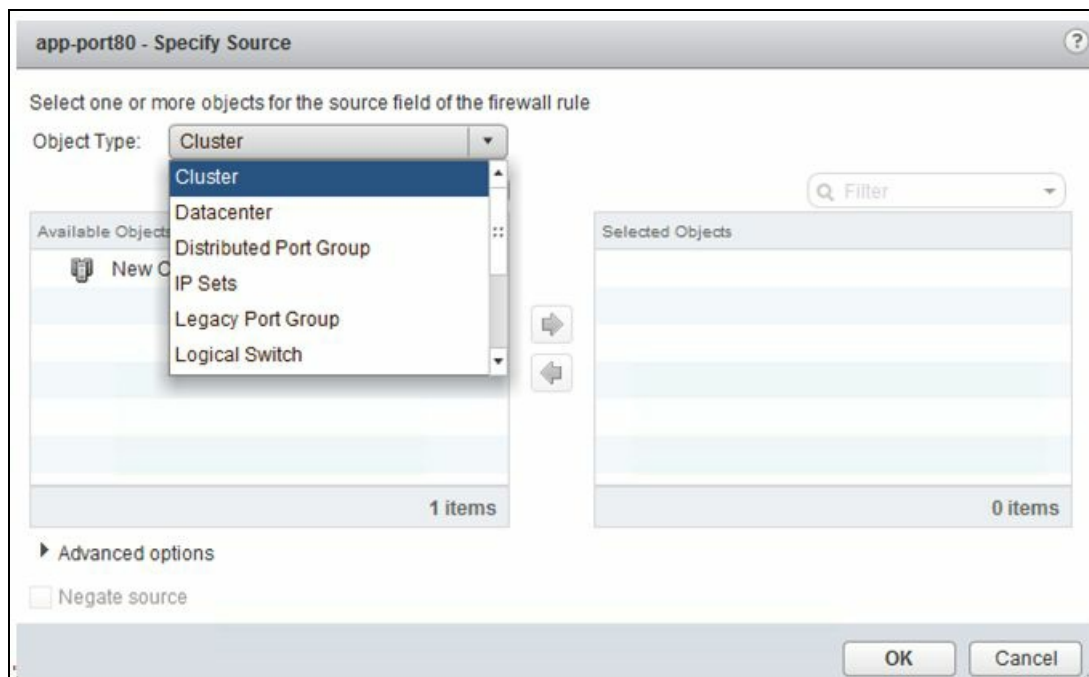


- Point to the source section, and you will see two icons. Click  to specify an IP address as source, and the  icon is to specify an object. You can specify both IPv4 and IPv6 addresses. You can have a rule with both an IP address and an object together.
- Click on the IP icon to add an IP address, then click Save:



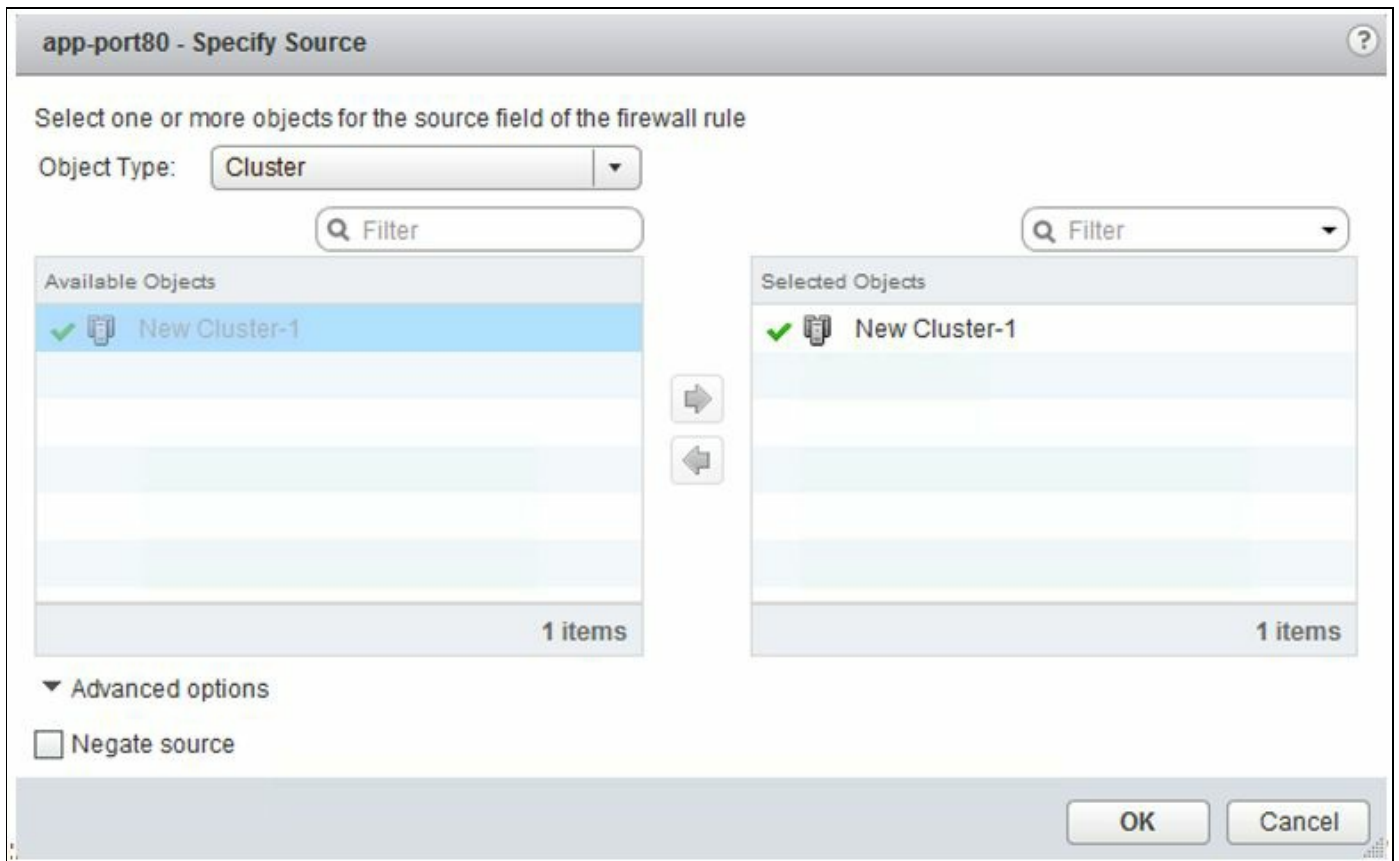
- To add an object, click on the  icon.






- Select the appropriate object from the Object Type drop-down. Select the available objects and click the arrow to include them. Click OK when done:



- The Advanced options link supplies the option to Negate Source. By default it's

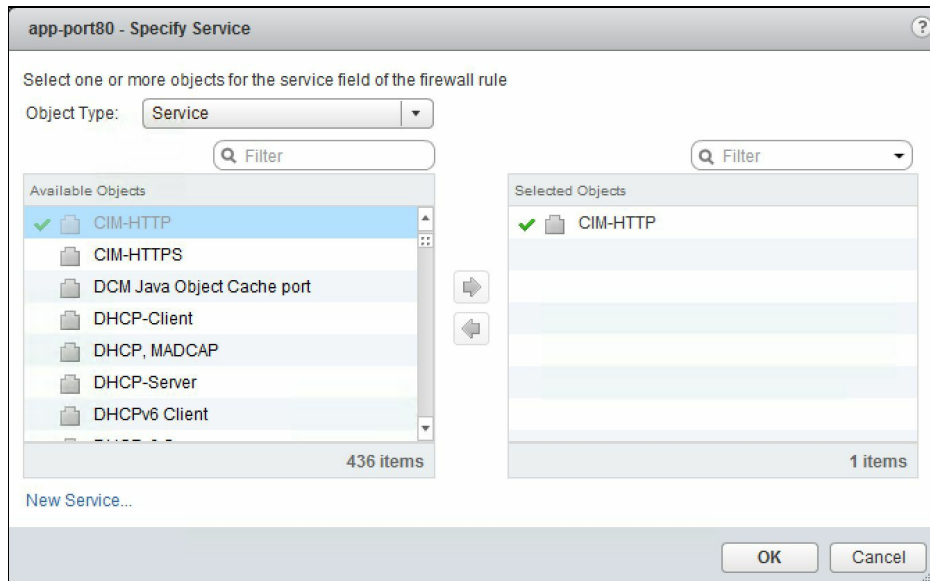
unchecked; if it's checked, the rule will apply to all source traffic except the ones you have selected here. If unchecked, the rule will apply to all traffic coming from the source set in the rule:



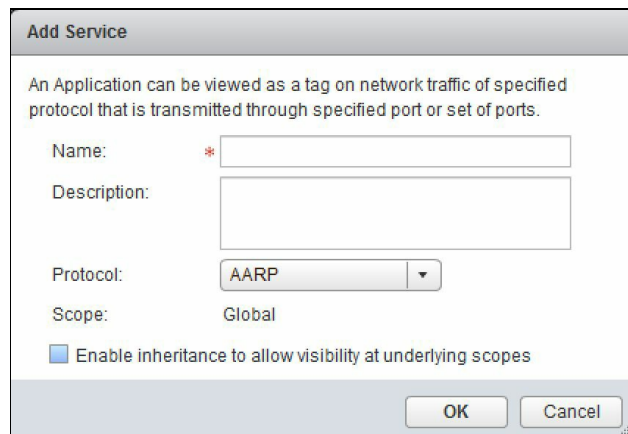
11. Similar to the source section, point to the destination section and you will see two icons. Click  to specify an IP address as source and click the  icon to specify an object. You can specify both IPv4 and IPv6 addresses. You can have a rule with both an IP address and an object together. The advanced option in destination has the Negate destination rule; if selected, this will apply the rule to traffic going to all destinations except the listed destination.
12. Point to the service section to specify the kind of service you want the rule to allow traffic to. Click on the  icon to define your port or the  icon to select from one of the pre-defined services. You can have both the port and a service in a rule.
13. Click on the  icon to define the service and click Save:



14. Click on the  icon to define a service or a service group:



You can also click on **New Service Group** and define your own service group. Once defined, this group is automatically added to the list of service groups. Click **OK**:



15. Point to the Action section to define the rule behavior. Click on the edit icon to define the action:

The screenshot shows a dialog box titled "app-port80 - Edit Action". It has the following fields and controls:

- Action: A dropdown menu with "Allow" selected.
- Direction: A dropdown menu with "In/Out" selected.
- Packet Type: A dropdown menu with "Any" selected.
- Tag: An empty text input field.
- Log: Two radio buttons, "Log" and "Do not log". "Do not log" is selected.
- Comments: A large empty text area.
- Buttons: "Save" and "Cancel" buttons at the bottom right.

16. The Applied To section identifies the scope at which the rule is applicable. Click on the edit icon to define the scope. You can select specific Edge gateways or apply this rule on all the Edge gateways. Click OK:


The screenshot shows a dialog box titled "Specify Applied To". It contains the following elements:


- Text: "Specify containers on which this rule will be applied."
- Checkboxes:
 - Apply this rule on all clusters on which Distributed Firewall is installed.
 - Apply this rule on all the Edge gateways.
- Text: "(For Edges with version 6.1.0 and higher)"
- Text: "Select one or more objects for the applied to field of the firewall rule"
- Search filters: Two "Filter" input fields with search icons.
- Available Objects: A list with one item, "ESG-1", which is highlighted in blue. Below the list is a counter "1 items".
- Selected Objects: An empty list. Below the list is a counter "0 items".
- Buttons: "OK" and "Cancel" buttons at the bottom right.

17. When done, click Publish Changes and the rule will be published.

In a cross-vCenter NSX environment, rules created in the universal firewall section are

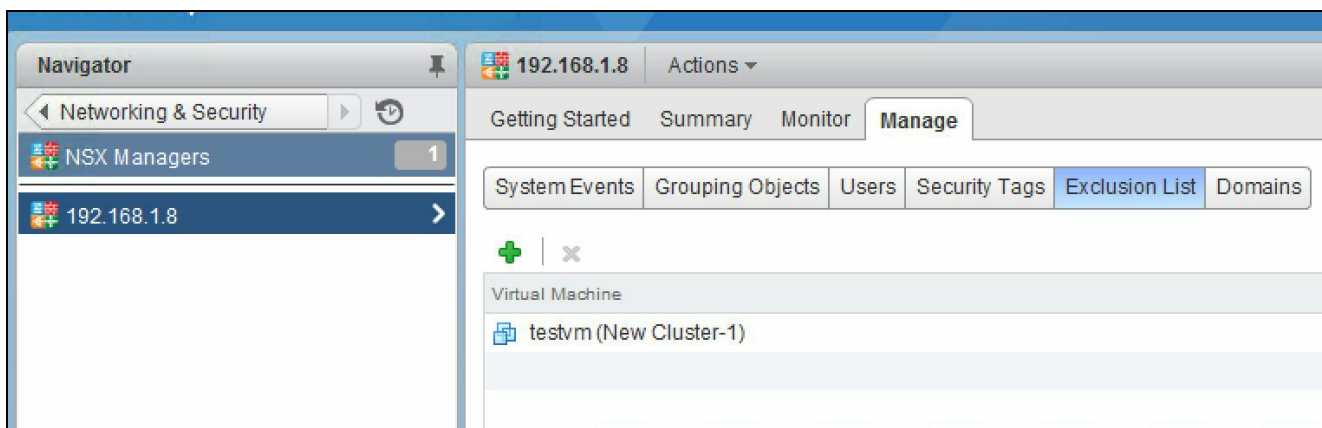
replicated to the secondary NSX managers. To create a universal firewall rule, follow these steps:

1. The universal firewall section is represented by the  icon. Click on the + sign in this section to add a universal rule.
2. Follow the previous procedure to add a rule, and click Publish Changes when done.

To delete a firewall rule, simply select the rule by clicking on it and click the  icon. Once done, you have to Publish Changes for the deletion to take effect.

You can exclude certain virtual machines from being protected by the distributed firewall. NSX managers, NSX controllers, and Edge appliances are automatically excluded. VMware recommends that the service exclusion list contain the vCenter server, the SQL server for vCenter, and any virtual machines that are running in promiscuous mode. To add virtual machines to the exclusion list, follow these steps:

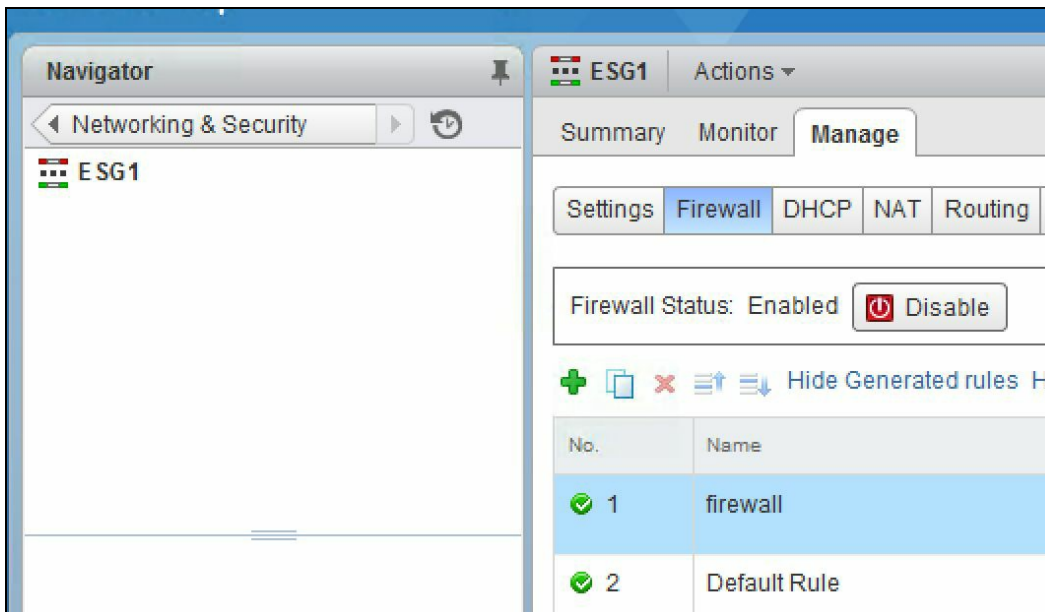
1. Go to Networking & Security | NSX Managers | Manage | Exclusion List.
2. Click the + icon and select a virtual machine to add it to the Exclusion List.
3. Click OK when done. The virtual machine is no longer protected by the distributed firewall:



NSX Edge services gateway also allows firewall protection that protects the north-south traffic in your environment. Let's look at configuring the Edge firewall. The default Edge firewall policy blocks all traffic. To access the Edge firewall, you must have an Edge gateway services appliance deployed and configured:

1. Go to Networking & Security | NSX Edges. Double-click on the Edge appliance listed.

2. Click on the Manage tab | Firewall:



3. You should see the rules, including Default Rule.

4. You can change the rule behavior by pointing to the Action row and editing it to any action that you prefer.

The process to add a new rule is similar to the logical firewall described earlier. You click the + symbol and continue to name the rule. Any new rules added are marked as the user type. Once you have configured your rule, click Publish Changes to publish and apply the rule. You can also change the order/priority of the rule by clicking the move up or move down arrows.

Summary

We began this chapter by discussing and deploying primary and secondary NSX managers, followed by learning how to configure VXLANs and segment IDs. We then configured our transport zones, which define the scope of logical switches in your environment. A logical switch can only be part of one transport zone, and there can only be one universal transport zone in a cross-vCenter NSX environment. We then looked at creating logical switches and L2 bridges. L2 bridges connect your logical networks to your physical networks for seamless network integration and extension. We looked at deploying a distributed logical router to enable L2 bridging. We also got started with the Edge services gateway by deploying it. We protected our environment by deploying a logical firewall and configuring rules in it. We finished the chapter by configuring the Edge services gateway firewall service, which is similar to the logical firewall configuration. Ideally, the Edge services gateway firewall protects north-south traffic, while the logical firewall protects east-west traffic.

In [Chapter 5](#), Edge Services Gateway, we will look at the NSX Edge gateway services and deploy and configure multiple Edge services.

Edge Services Gateway

We were introduced to NSX Edge services gateway in the previous chapter. We will continue to explore the deployment and configuration of different services that the Edge services gateway has to offer. We will look at configuring OSPF and BGP routing using the Edge services gateway. We will then walk through the deployment of logical Edge load balancers. The Edge services gateway also offers the ability to set up virtual private networks that enable secure access to your environment. We will look at configuring virtual private networks followed by configuring DNS and DHCP services. We will finish the chapter by looking at some more Edge services gateway configurations.

In this chapter, we will cover:

- DNS and DHCP services
- Routing
- Logical Edge load balancers
- Virtual Private Networks
- More Edge services and configurations

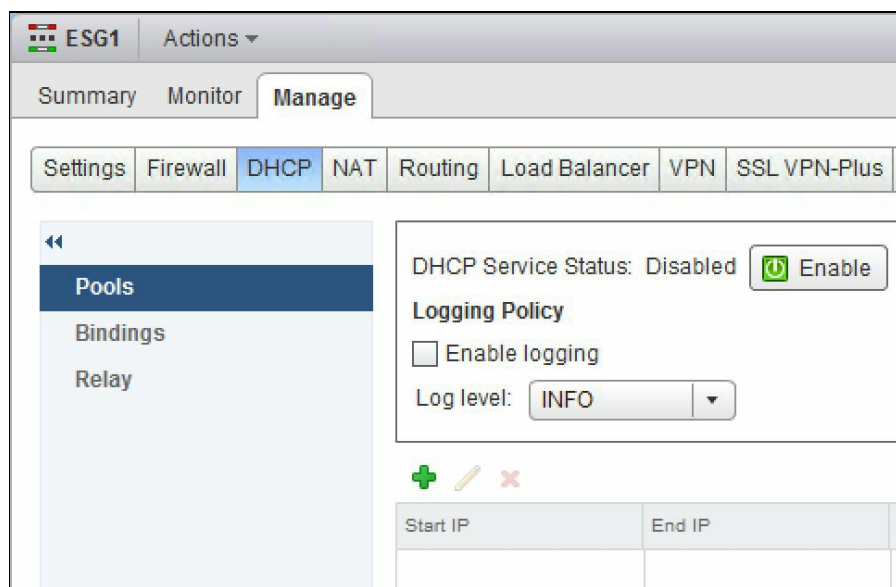
DNS and DHCP services

NSX Edge services gateway offers DNS and DHCP services. Let's look at configuring these services in the NSX Edge gateway.

DHCP service

NSX allows for one-to-one static NAT IP address allocation and IP address pooling. When the DHCP service is set up, it listens to any DHCP discovery requests on the internal interfaces and responds:

1. Go to Home | Networking & Security | NSX Edges, and double click an NSX Edge, and navigate to Manage | DHCP:



2. Click the + icon to add a new pool:

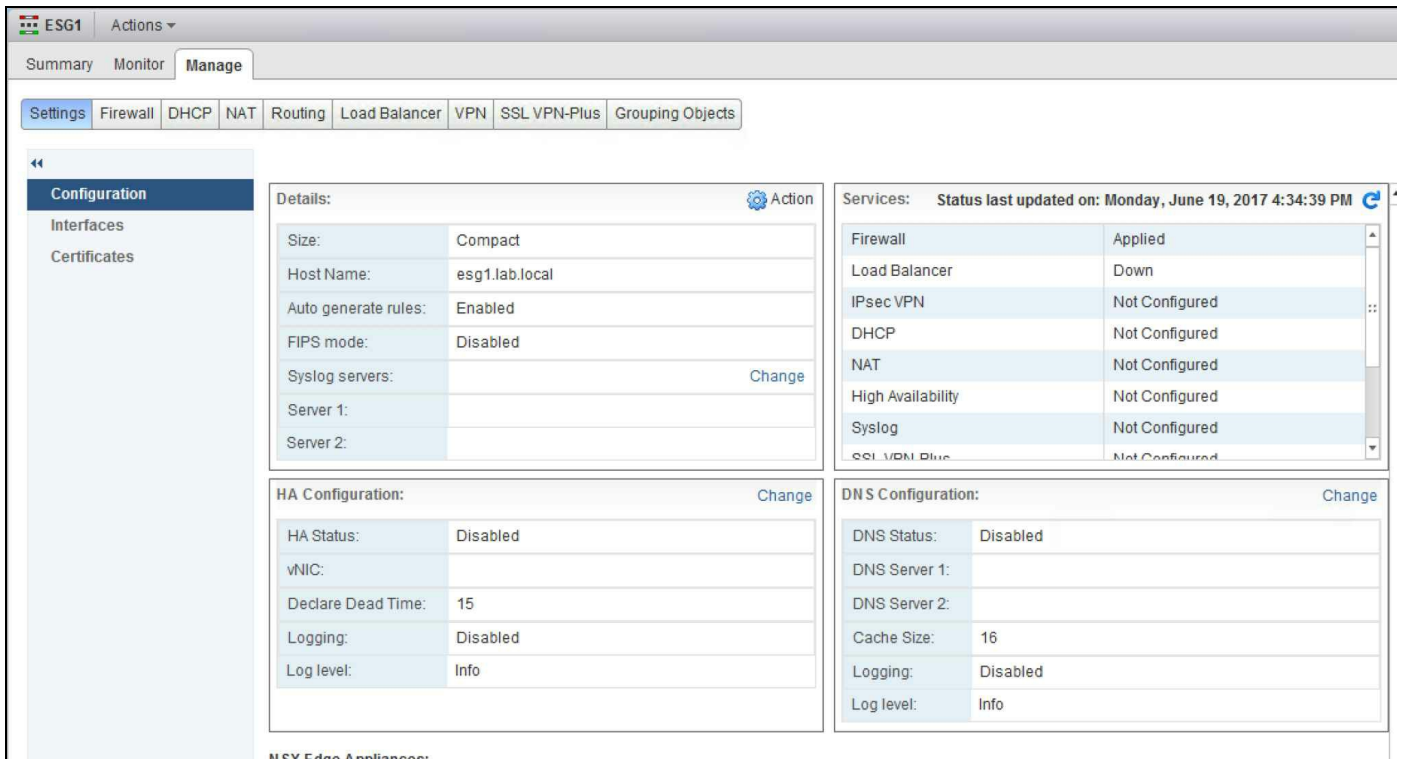
The screenshot shows a dialog box titled "Add DHCP Pool". It features two unchecked checkboxes at the top: "Auto Configure DNS" and "Lease Never Expires". Below these are several input fields: "Start IP:" and "End IP:" (both marked with a red asterisk), "Domain Name:", "Primary Name Server:", "Secondary Name server:", "Default Gateway:", and "Subnet Mask:". The "Lease Time:" field is set to "86400" with "(seconds)" next to it. At the bottom right, there are "OK" and "Cancel" buttons.

3. Enabling Auto Configure DNS allows you to configure the DNS for every DHCP binding automatically.
4. Enabling Lease Never Expires binds the IP to the MAC address forever.
5. Enter the Start IP and the End IP. Fill in the rest of the form with applicable values and click OK when done.
6. Now that a pool has been added, let's go ahead and enable the DHCP service.
7. Click on Enable and click Publish Changes. You can also enable logging by selecting Enable logging and setting the appropriate log level.

DNS service

NSX Edge can be configured with external DNS servers and can relay name resolution requests:

1. Go to Home | Networking & Security | NSX Edges and double click on Edge | Manage | Settings | Configuration:



2. In the DNS Configuration section, click Change:



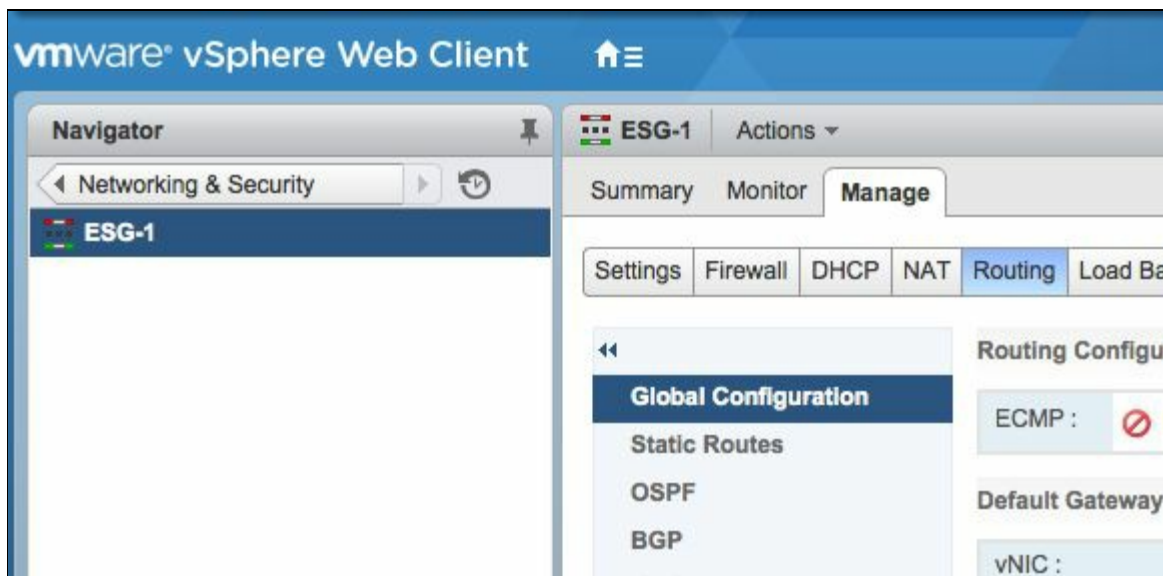
3. Select Enable DNS service to enable the service. Enter the external DNS server names and the cache size if required. Enable logging if needed and click OK when done.

Routing


In [Chapter 4](#), NSX Functional Services, we looked at deploying a logical router and the NSX Edge services gateway. We will now look at how to enable routing services using the NSX Edge services gateway. These services allow for more customized routing within your environment to suit your needs. You can configure the default gateway for the router and **equal cost multi-path (ECMP)** routing that allows for highly available deployment of multiple Edge gateways to prevent bottlenecks. You can even configure dynamic routing that updates routing tables with real-time logical network changes. All these can be set by editing the global configuration of an Edge services gateway.

Once the Edge services gateway is deployed, follow these steps for global configuration:

1. Go to Home | Networking & Configuration | NSX Edges.
2. Double click the Edge device that needs to be configured.
3. Go to Manage | Routing | Global Configuration:



4. To enable ECMP routing, click Enable. ECMP allows the next-hop packet to be forwarded to a single destination over multiple best paths that can be added statically or dynamically using routing protocols such as OSPF and BGP. These multiple paths are added as comma separated values when defining the static routes.
5. To add a default gateway, click Edit under the Default Gateway section:



6. Select an interface that will be the outgoing interface for the next hop. Set a Gateway IP. The Admin Distance is a metric used to choose which route to take when multiple routes are available for a network. The lower the metric, the higher the priority. The value ranges from 1 to 255. The default ranges are: Connected (0), Static (1), External BGP (20), OSPF Intra-Area (30), OSPF Inter-Area (110), and Internal BGP (200).
7. Click OK.



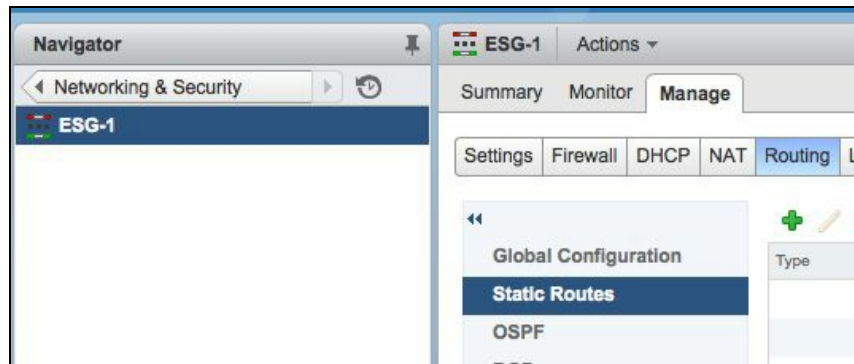
Remember that in dynamic routing, routing protocols running in routers continuously exchange network status updates between each other as broadcast or multicast.

To configure dynamic routing, follow these steps:

1. Click Edit in the dynamic routing configuration section.
2. Router ID shows the first uplink of the Edge appliance that is used to push the routes to the kernel for dynamic routing.
3. Select Enable Logging to enable the logs.
4. Click Publish Changes to save.

To set a static route in the Edge gateway, do the following:

1. Go to Home | Networking & Services | NSX Edges, and double click the Edge appliance, and then navigate to Manage | Routing tab.
2. Select Static Routes on the left menu:



3. Click the + icon:

4. Enter the Network in CIDR notation. For example, 192.168.1.0/24.
5. Enter the Next Hop information. This is an IP address that the Edge is able to reach directly. If ECMP is configured, then multiple next hop IP addresses can be mentioned with a comma.
6. Select an interface on which this static route will apply.
7. Select the appropriate Admin Distance.
8. Click OK.
9. Click Publish Changes when done.

Configuring OSPF on Edge services gateway

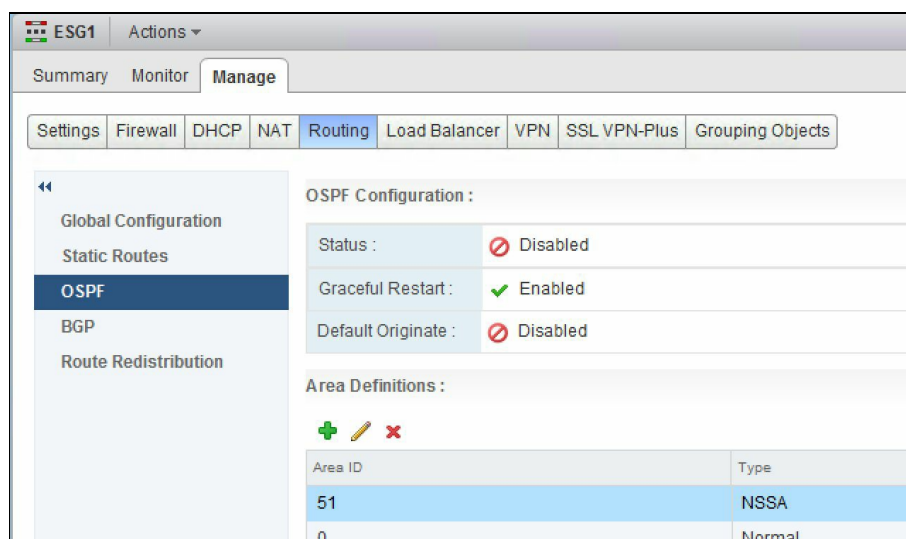
The Edge services gateway supports OSPF, BGP, and IS-IS routing protocols. An Edge services gateway can learn and advertise routes when configured with the OSPF routing protocol. When there are routes of equal cost, OSPF provides dynamic load balancing between these routes. The routing table size can become a challenge and an OSPF network limits the size of these routing tables by dividing the network into routing areas to optimize traffic flow. An area is identified by an area ID and is comprised of routers, links, and a logical collection of OSPF networks that have the same area identification.



Open Shortest Path First (OSPF) is a routing protocol which uses a link state routing algorithm and operates within a single autonomous system.

Before we begin, ensure that a Router ID is configured. As we learned earlier, a Router ID is simply an uplink interface for the Edge services gateway that connects to the external peer. Perform the following set of steps to configure a Router ID:

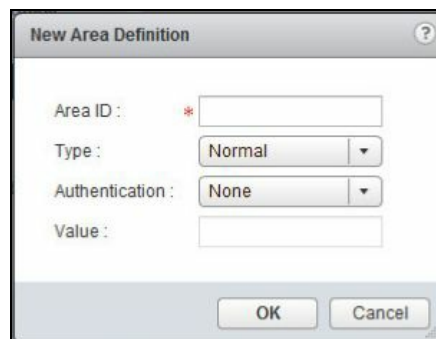
1. Log in to NSX and navigate to Home | Networking & Security | NSX Edges, and double click on the Edge services gateway, and then go to the Manage | Routing | OSPF:



2. Click Edit to Enable OSPF configuration:

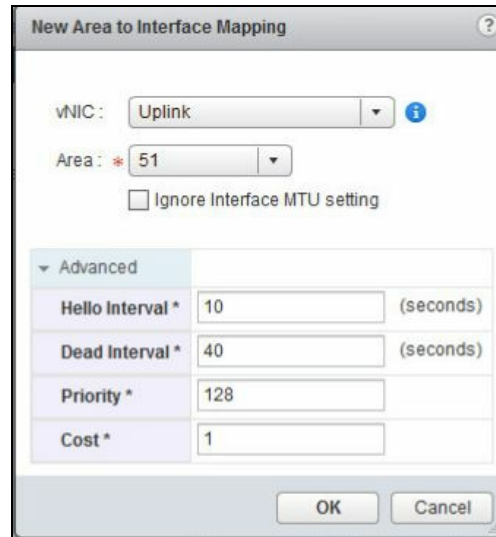


3. Click Enable OSPF. Enable Graceful Restart allows for non-stop packet forwarding even if the OSPF process is being restarted. This helps in non-disruptive packet routing.
4. If you want your Edge services gateway to advertise itself as a default gateway to its peers, click OK when done.
5. We need to configure OSPF areas. By default, Area 51 is configured, which can be deleted if needed.
6. Click the + icon to add a new area:



7. Type in an Area ID—it can be an IP address or a decimal number. If your area contains a specific network segment, you can type in `192.168.1.0` to easily identify that area. Next, select the Type of area—Normal or NSSA. **NSSA** stands for **Not-So-Stubby Area**. NSSA prevents the flooding of external autonomous system link state advertisements by relying on the default routes to external destinations. NSSAs are typically placed at the Edge of an OSPF routing domain.
8. Next, select the type of Authentication. This is optional, but once set you need to ensure that all routers in the area have the same authentication type set with the same password. For MD5 authentication, both receiving and transmitting routers must have the same key configured. Click OK when done.

9. Now that we have an area defined, we will proceed to map an interface to that area. In the New Area to Interface Mapping, click the + icon:



The screenshot shows a dialog box titled "New Area to Interface Mapping". It contains the following fields and options:

- vNIC: Uplink (dropdown menu)
- Area: 51 (dropdown menu)
- Ignore Interface MTU setting
- Advanced section (expanded):
 - Hello Interval *: 10 (seconds)
 - Dead Interval *: 40 (seconds)
 - Priority *: 128
 - Cost *: 1
- Buttons: OK, Cancel

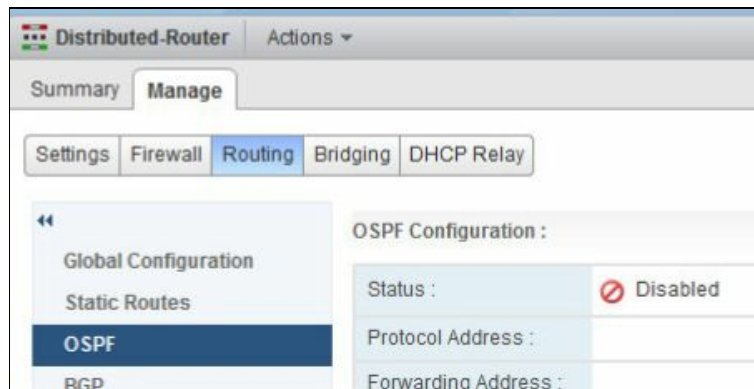
10. Select the interface and the Area ID to map. There are optional Advanced settings that can be changed. The Hello Interval is the default interval between the hello packets sent to the interface. The Dead Interval is the time out interval before the router declares a neighbor down. The Priority is the default routing priority for the interface. Finally, the Cost of the interface displays the default overhead required to send packets over that interface. You can define the cost based on bandwidth on an interface so that the Edge can determine what the lowest cost to route traffic across multiple interfaces is. The higher the bandwidth on an interface, the lower the cost.
11. Click OK when done.
12. Click Publish Changes when done.

Configuring logical distributed router OSPF

Configuring OSPF on the logical distributed router enables VM connectivity across logical routers. It also forms the bridge between the logical router and the Edge services gateway. Configuring OSPF on a logical distributed router is similar in many ways to configuring OSPF on an Edge services gateway.

Ensure that a Router ID is configured. Perform the following set of steps:

1. Go to Home | Networking & Security | NSX Edges, and double click on your logical distributed router, and go to Manage | Routing | OSPF:



2. Click Edit under OSPF Configuration:



3. Click Enable OSPF. In the Protocol Address, enter the IP address that the OSPF

protocol will use to form adjacencies with its peers. The Forwarding Address is then used by the distributed router (in the hypervisor) to forward packets. Click OK.

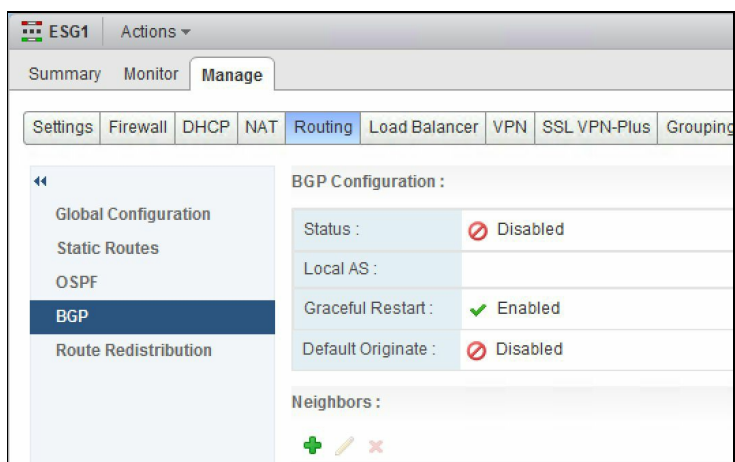
4. Configuring the OSPF areas and area to interface mapping is similar to the configuration described in the Edge services gateway.
5. Once done, click Publish Changes.

Configuring BGP

Let's look at configuring the border gateway protocol. BGP makes core routing decisions and includes a table of IP networks.

The BGP is an exterior gateway protocol designed to exchange routing information among **autonomous systems (AS)** on the internet. BGP is relevant to network administrators of large organizations that connect to two or more ISPs, as well as to internet service providers who connect to other network providers. If you are the administrator of a small corporate network or an end user, then you probably don't need to know about BGP:

1. Go to Home | Networking & Security | NSX Edge, and double click an NSX Edge, and go to Manage | Routing | BGP:



2. Click Edit under BGP configuration:



3. Click Enable BGP. Click Enable Graceful Restart to continue to forward packets

even if the BGP process restarts. Enable Default Originate allows this Edge to advertise itself as a gateway to its peers. Type the Router ID in the Local AS. This is what the router uses to advertise itself to routers in other autonomous systems.

4. In the Neighbors field, click the + icon:

The 'New Neighbour' dialog box contains the following fields and sections:

- IP Address: *
- Remote AS: *
- Weight: 60
- Keep Alive Time: 60 (Seconds)
- Hold Down Time: 180 (Seconds)
- (BGP Keep alive timer value needs to be one third of hold down timer)*
- Password:
- BGP Filters: A section with a search bar and a table.

Direction	Action	Network	IP Prefix GE	IP Prefix LE

0 items

Buttons: OK, Cancel

5. Type the IP Address of the neighbor that is the peer of this router.

6. Next, type in the Remote AS. The Weight determines the priority of the route.

7. The Keep Alive Time field determines how often the keep alive messages are sent to the peers. Hold Down Time determines the time-out duration before it declares a peer dead.

8. If authentication is required, set a password in the Password section, but ensure that each of the BGP neighbors is configured with the same password.

9. To filter any routes, click the + icon in the BGP Filters:

The 'New BGP Filter' dialog box contains the following fields:

- Direction: *
- Action: *
- Network: *
- Network should be specified as 'ANY' or in CIDR format e.g. 192.169.1.0/24.*
- IP Prefix GE
- Expected IP Prefix GE (0-32)*
- IP Prefix LE
- Expected IP Prefix LE (0-32)*

Buttons: OK, Cancel

10. Set the Direction for filtering traffic and the Action to perform. Next, type the Network in the CIDR format that you want to filter traffic to or from the neighbor. Next, enter the IP Prefixes GE and IP Prefixes LE and click OK.
11. Click Publish Changes when done.

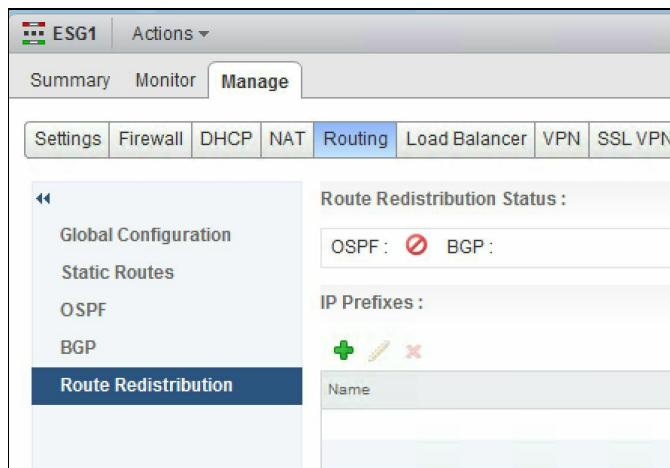


Intermediate System to Intermediate System (IS-IS) is no longer supported with NSX 6.3.x.

Configuring route redistribution

In an environment where there are multiple routing protocols being used, route redistribution enables cross-protocol route sharing:

1. Go to Home | Networking & Security | NSX Edges, and double click an NSX Edge, and go to Manage | Routing | Route Redistribution:



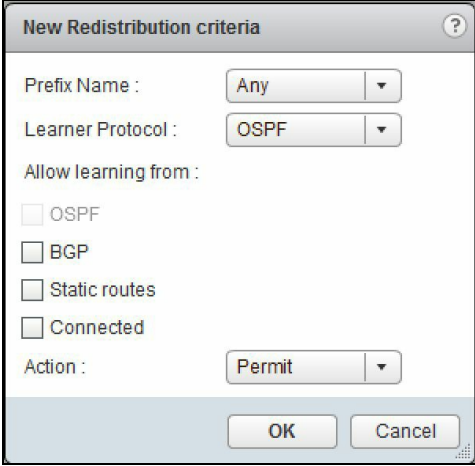
2. Click Edit to select the protocols to which you want to enable route redistribution. Only the configured protocols will show:



3. Click the + icon to add an IP prefix:



4. Next you need to specify the redistribution criteria for the IP prefix entered. Click the + icon in the route redistribution table:



The image shows a dialog box titled "New Redistribution criteria". It contains the following fields and options:

- Prefix Name : Any (dropdown menu)
- Learner Protocol : OSPF (dropdown menu)
- Allow learning from :
 - OSPF
 - BGP
 - Static routes
 - Connected
- Action : Permit (dropdown menu)

At the bottom of the dialog box are two buttons: "OK" and "Cancel".

5. Select the appropriate Prefix Name from the drop-down menu.
6. Next, specify the Learner Protocol that will learn routes from other protocols.
7. The Allow learning from field allows you to specify the protocols to learn the routes from.
8. Click OK when done.
9. Click Publish Changes when done.

Logical Edge load balancers

Load balancers allow network traffic to be balanced across multiple servers to increase performance, and they also allow the high availability of services. This distribution of an incoming service among multiple servers is transparent to the end users, which makes deploying load balancers a critical component of any environment.



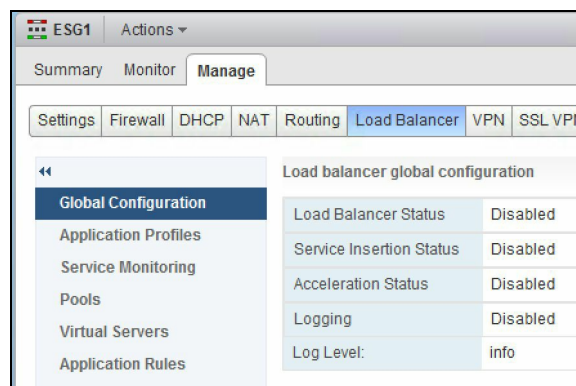
A good use-case for load balancers for further reading can be found at: <http://cloudmaniac.net/load-balance-vmware-psc-with-nsx/>

The Edge services gateway offers logical Edge load balancers that allow you to utilize the load balancing functionality and distribute incoming traffic across multiple virtual machine instances.

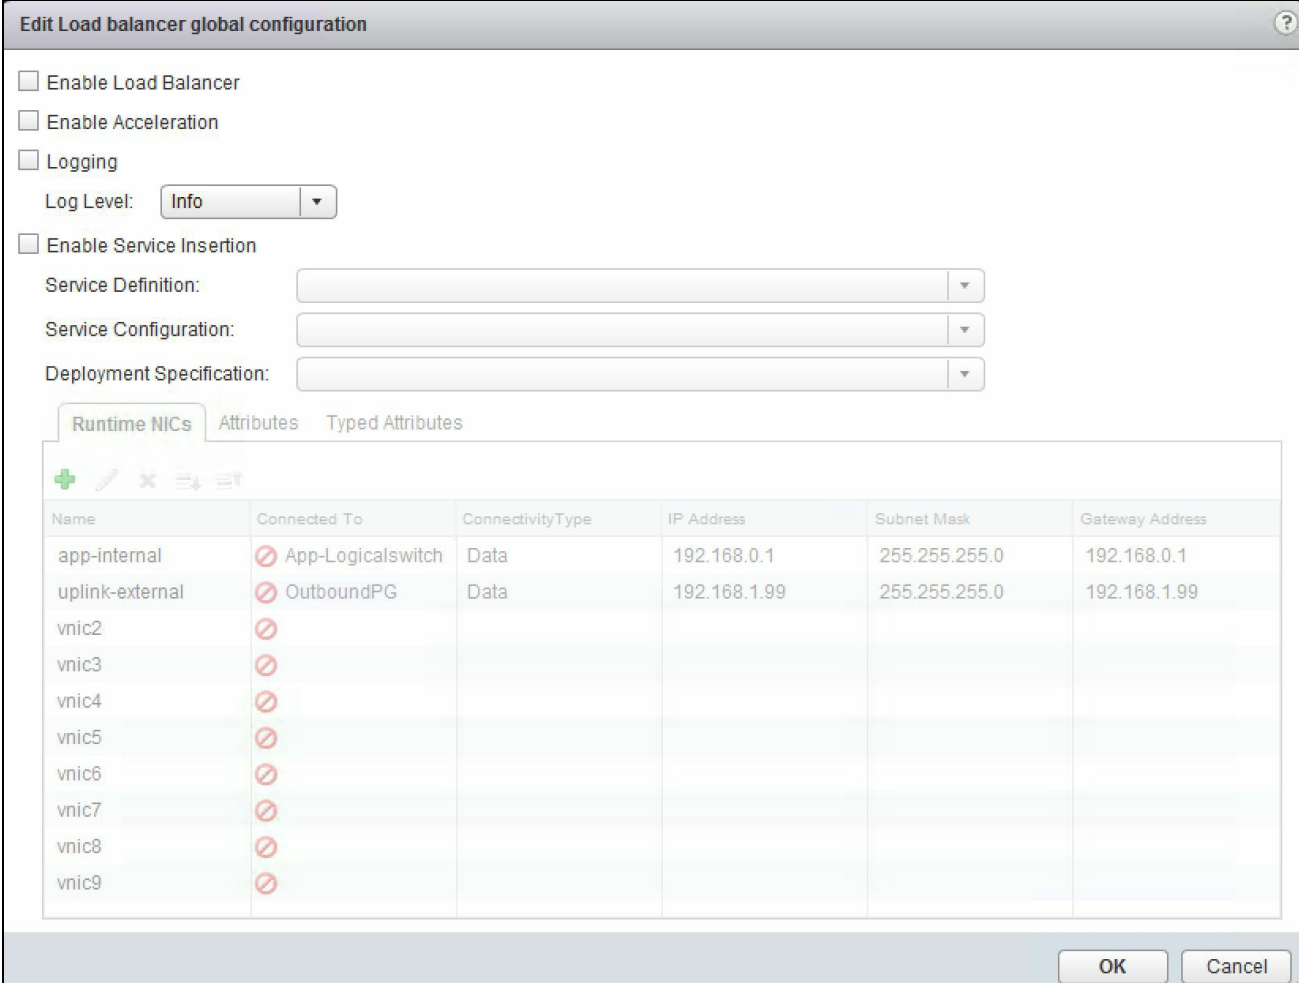
An Edge services gateway instance must be deployed in order to enable the load balancer service. There are multiple steps required when configuring a load balancer service, and it begins with enabling the service, followed by configuring the application profile to define the behavior based on the traffic type. Once these are defined, you will create a service monitor to enable a health check of the services behind the load balancer. This will stop you from sending traffic to a dead node. You can then create a server pool that has the list of servers participating in the load balancer and a virtual node that will receive all the traffic and distribute it among the pool based on the policies set.

To configure the load balancer service, do the following:

1. Go to Home | Networking & Security | NSX Edges, and double click on an Edge gateway, and go to Manage | Load Balancer | Global Configuration:



2. Click Edit to enable the load balancer service:



The screenshot shows the 'Edit Load balancer global configuration' dialog box. The 'Enable Load Balancer' checkbox is checked. The 'Log Level' is set to 'Info'. The 'Runtime NICs' tab is active, showing a table of network interfaces.

Name	Connected To	ConnectivityType	IP Address	Subnet Mask	Gateway Address
app-internal	App-Logicals witch	Data	192.168.0.1	255.255.255.0	192.168.0.1
uplink-external	OutboundPG	Data	192.168.1.99	255.255.255.0	192.168.1.99
vnic2					
vnic3					
vnic4					
vnic5					
vnic6					
vnic7					
vnic8					
vnic9					

Select Enable Load Balancer to enable the load balancer service.

Enable Acceleration enables the Edge load balancer to use the faster L4 load balancer engine rather than the L7 engine.

Layer 4 load balancer takes routing decisions based on IPs and TCP or UDP ports. It has a packet view of the traffic exchanged between the client and a server and takes decisions packet by packet. The layer 4 connection is established between a client and a server.

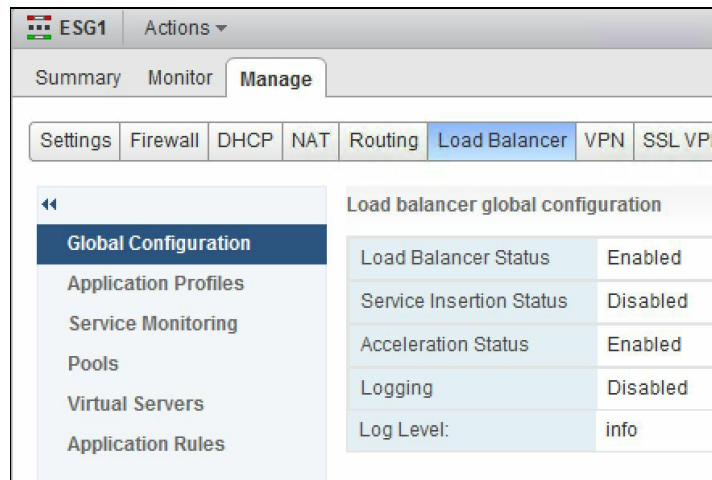
A layer 7 load balancer takes routing decisions based on IPs, TCP, or UDP ports or other information it can get from the application protocol (mainly HTTP). The layer 7 load balancer acts as a proxy and maintains two TCP connections: one with the client and one with the server.

Logging allows you to specify the level of logging. Increasing the logging level

will cause a lot of data to be collected on the Edge appliance. In such instances, a syslog server is recommended as per best practice.

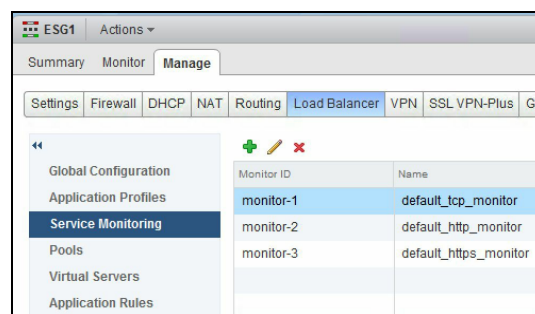
Enable Service Insertion allows the Edge appliance to work with a third-party load balancer directly.

3. Click OK to enable the load balancer service:



Now that the load balancer service is enabled, we will proceed to create a service monitor to monitor and define health check parameters. This service monitor is associated with a pool of servers that will be serving the traffic behind the load balancer:

1. Go to Home | Networking & Security | NSX Edges, and double click on an Edge appliance, and then go to Manage | Load Balancer | Service Monitoring:



You will notice the default monitors in place. You can either edit them or remove them if not needed.

2. Click on the + icon to add a New Service Monitor:

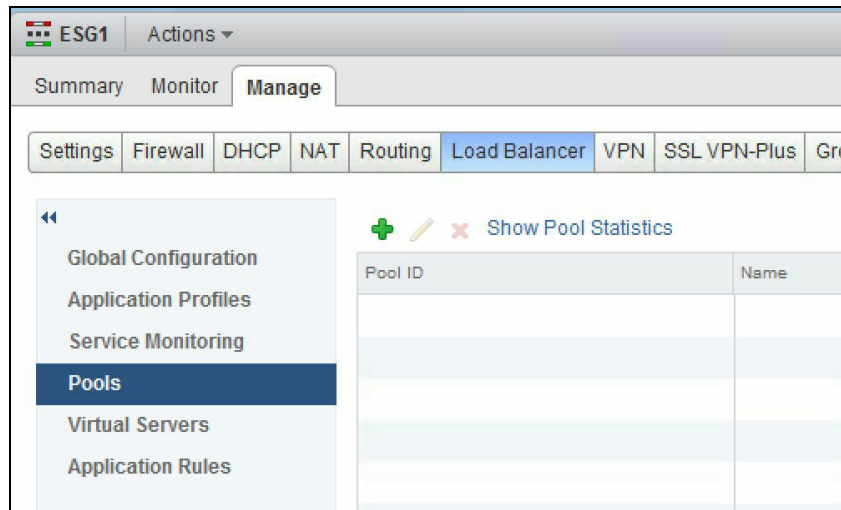
The image shows a 'New Service Monitor' dialog box with the following fields and values:

- Name: * (empty)
- Interval: 10 (seconds)
- Timeout: 15 (seconds)
- Max Retries: 3
- Type: HTTP
- Expected: (empty)
- Method: GET
- URL: /
- Send: (empty)
- Receive: (empty)
- Extension: (empty text area)

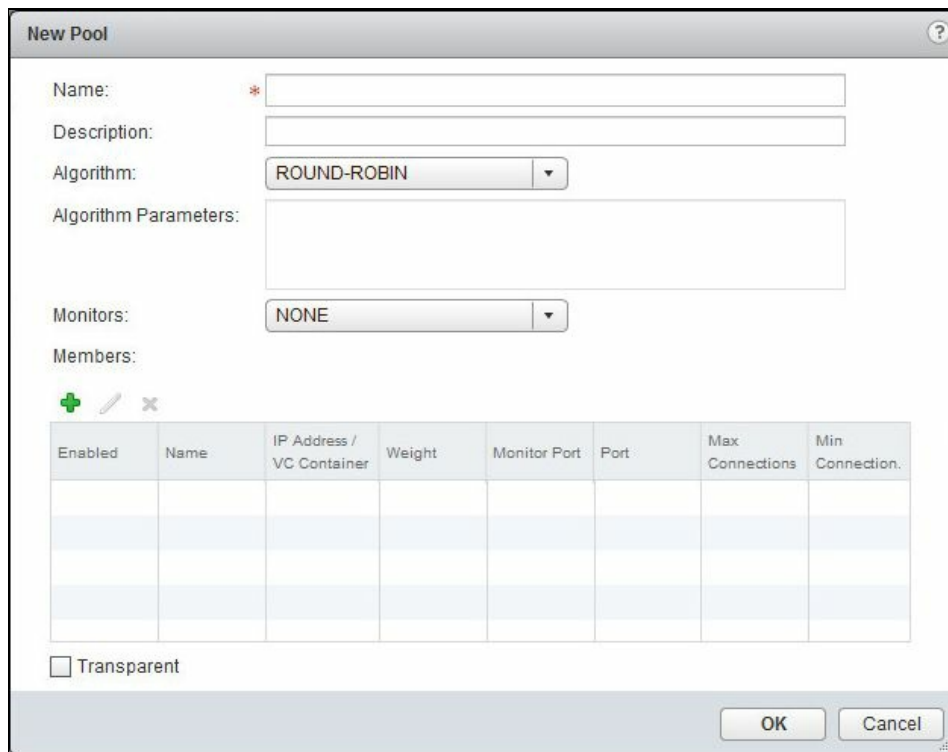
3. Enter a Name for the service, followed by the Interval frequency to ping the service.
4. Enter the Timeout for the maximum time in seconds to receive a response, followed by the maximum allowed retries before declaring a failure.
5. The Type field defines the way the service is checked for the HTTP and HTTPS types.
6. Enter the expected string in the Expected section. This is what the monitor should expect when it checks for the HTTP or the HTTPS service.
7. Select Method and the base URL to test. If the method is POST instead of GET, type the data in the Send field to send to the server. When the Expected string matches, the monitor continues to also match the Receive string.
8. The Extension allows you to enter advanced monitoring parameters, such as key:value pairs. This is optional.
9. Click OK when done.

We will now proceed to create a server pool so we can associate it with the service monitor:

1. Click on Pools in the Load Balancer section of the Edge appliance:



2. Click on the + icon to create a New Pool:



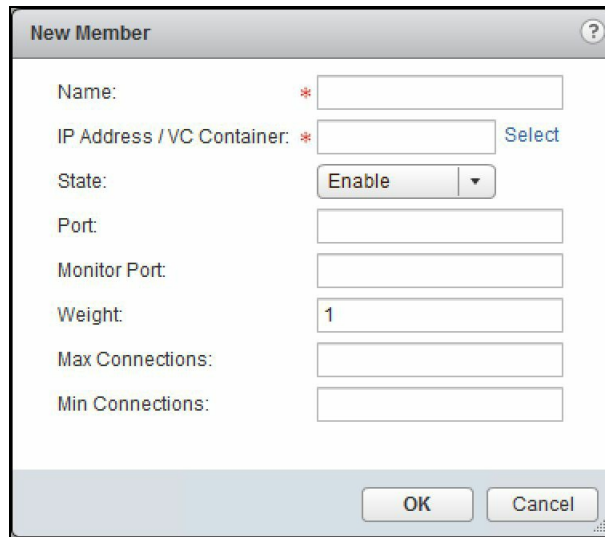
3. Enter the Name of the pool. Select the appropriate Algorithm for the service. This depends on what service you are working with to load balance traffic. The four options are:

- Round-Robin: Each server is used in turn according to the priority or weight assigned to it.
- IP_Hash: Selects a server based on the hash of the source and destination IP address of each packet.
- Least_Conn: Directs new connections to the server in the pool that has the least

connections active.

- URI: The URI is hashed and divided by the total weight of all the servers in the pool. The result is used to decide which server in the pool will take the request.

4. Select the Monitors that apply to the pool from the drop-down menu.
5. Add members to the pool by clicking the + icon:



6. Type in a Name for the member followed by the IP address or a virtual center object such as a cluster.
7. Set the State of the member. Choose between the Enable, Disable, and Drain options.



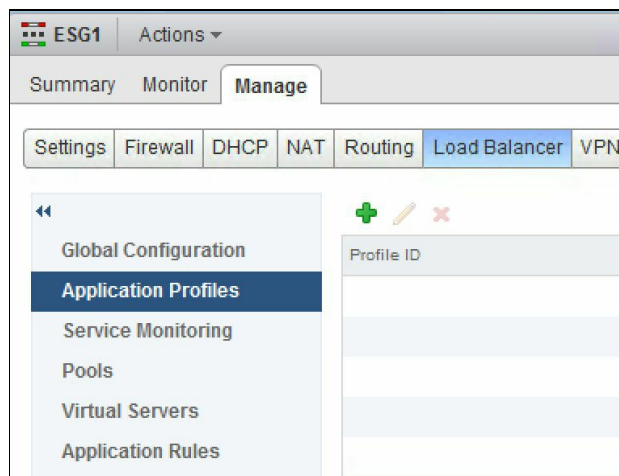
The Drain option forces the server to shut down gracefully for maintenance. This setting also removes the server from the load balancing pool, but it will still be used for existing and new connections from clients with persistent session to that server.

8. Type in the Port member where the traffic is to be sent. The Monitor Port is the port where the member receives the health check pings. The Weight determines how much traffic this member can handle.
9. The Max Connections and Min Connections allow you to manage traffic and number of connections appropriately. Click OK.

10. The Transparent option allows the backend servers in the pool to see the source IP of the request. Transparent is disabled by default and the backend servers only see the traffic coming in from the internal load balancer IP.
11. Click OK when done.

Before we create a virtual server to map to the pool, we have to define an application profile that defines the behavior of a particular type of network traffic. When traffic is received, the virtual server processes the traffic based on the values defined in the profile. This allows for greater control over managing your network traffic:

1. Select Application Profiles:



2. Click the + icon:

New Profile

Name:

Type:

Enable SSL Passthrough

HTTP Redirect URL:

Persistence:

Cookie Name:

Mode:

Expires in (Seconds):

Insert X-Forwarded-For HTTP header

Enable Pool Side SSL

Virtual Server Certifica... Pool Certificates

Service Certificates CA Certificates CRL

Configure Service Certificate

	Common Name	Issuer	Validity
<input checked="" type="radio"/>	vcenter.lab.local	CA	Mon Jul 11 2016 - Mon Jul
<input type="radio"/>	VSM_SOLUTION_1100	VSM_SOLUTION_1100	Tue May 2 2017 - Thu Apr
<input type="radio"/>	VSM_SOLUTION_1100	VSM_SOLUTION_1100	Tue May 2 2017 - Thu Apr
<input type="radio"/>	VSM_SOLUTION_7d3b	VSM_SOLUTION_7d3b	Tue May 2 2017 - Thu Apr
<input checked="" type="radio"/>	VSM_SOLUTION_7d3b	VSM_SOLUTION_7d3b	Tue May 2 2017 - Thu Apr

Cipher:

Client Authentication:

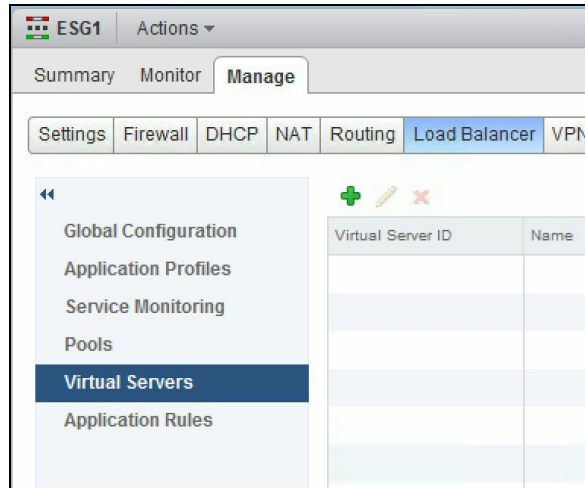
OK Cancel

3. Type a Name for the profile and the Type of traffic. If you want to redirect HTTP traffic, enter the URL it needs to be redirected to.
4. Specify the Persistence that applies to the profile. Persistence tracks and stores any session data. There are different persistence methods supported for different types of traffic.
5. Selecting HTTPS allows you to terminal SSL certificates at the load balancer or even configure an SSL to pass through to your backend pool servers.
6. Select any Cipher algorithms that are negotiated during the SSL/TLS handshake that apply.
7. Click OK when done.

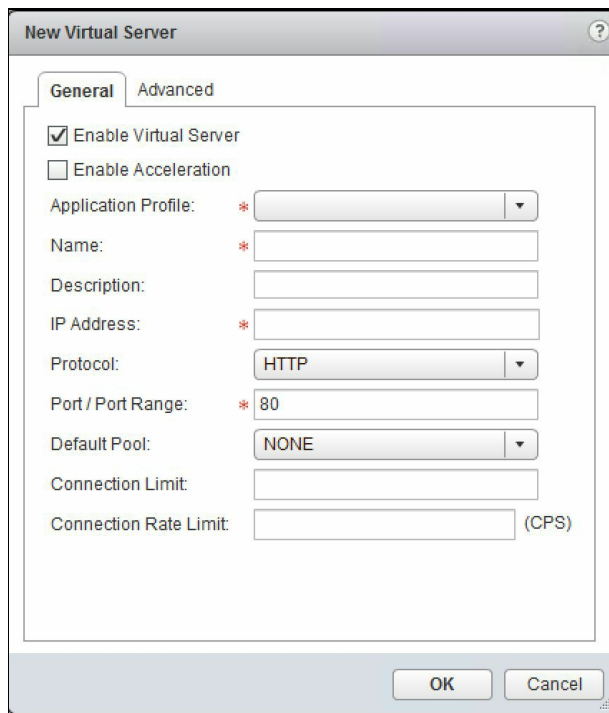
Now that we have the application profile created, let's create a virtual server and associate it with the pool. Once this is done, external traffic can be directed to the

virtual server IP that, in turn, distributes the traffic across the pool members based on the algorithm we have defined:

1. Select the Virtual Servers and click the + icon to add a new virtual server:



2. You will see the New Virtual Server window pop out:



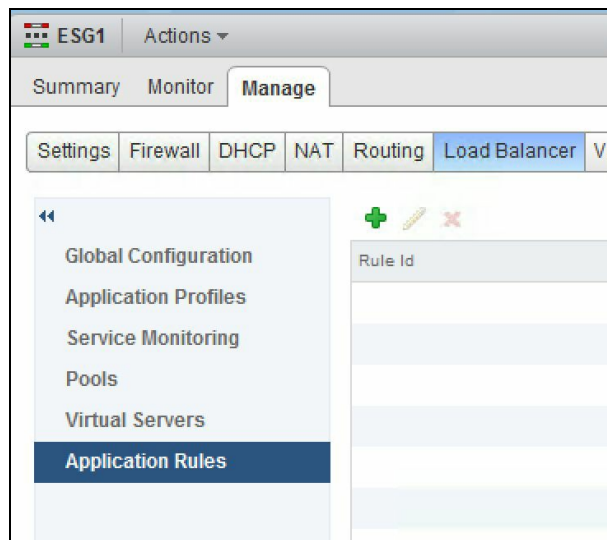
Check Enable Acceleration for the NSX Edge load balancer to use the faster L4 load balancer engine rather than the L7 load balancer engine.

3. Select the Application Profile for your virtual server.

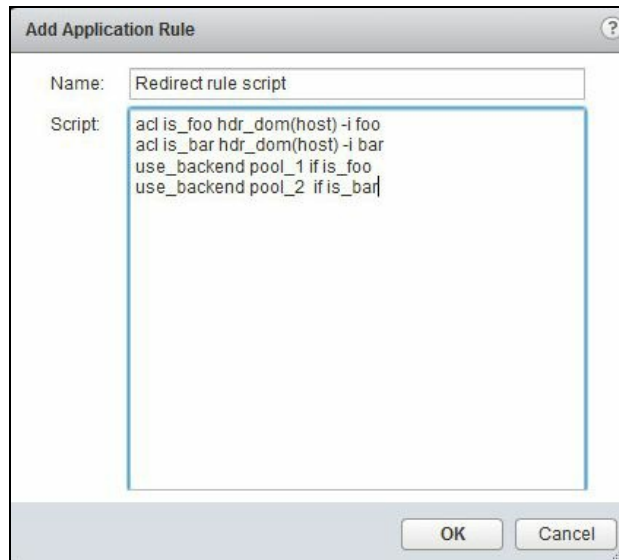
4. Type the Name for the virtual server.
5. Enter the IP Address of the virtual server. This is the IP address the load balancer will be receiving all the external traffic from.
6. Select the Protocol the virtual server will handle and the port to receive the traffic on.
7. Select the Default Pool and set the Connection Limit if applicable.
8. Click OK when done.

Let's now look at Application Rules to understand their applications and configuration. An application rule allows you to specify rules and logic to manage your traffic to make intelligent redirection. You can use an application rule to directly manipulate and manage IP application traffic. This becomes critical in fine-tuning your traffic to the application you are running behind the perimeter:

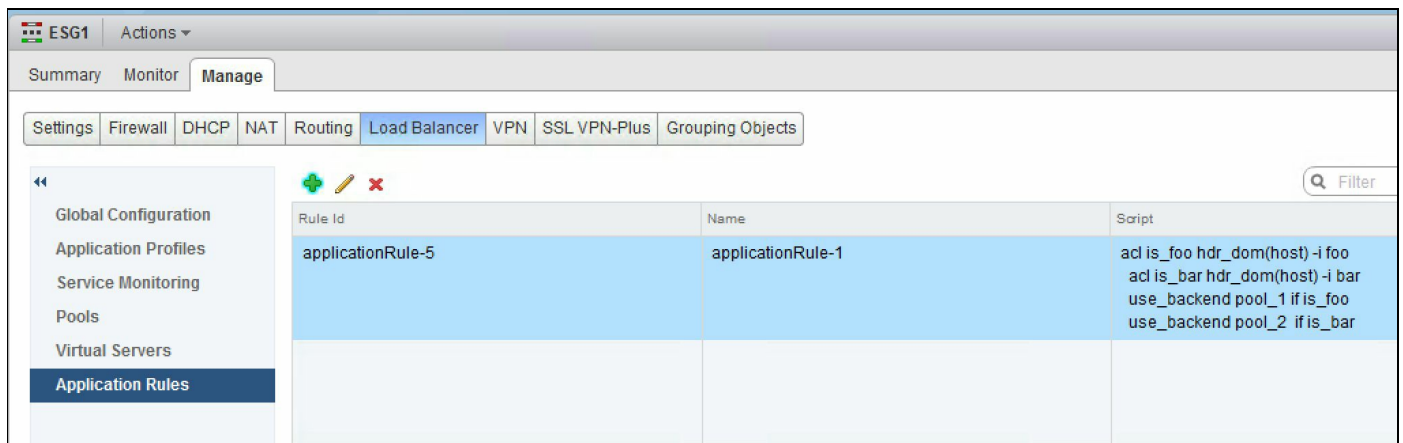
1. Log in to the vSphere web client and go to Home | Networking & Security | NSX Edges.
2. Double click on an NSX Edge and navigate to the Manage | Load Balancer tab:



3. Click Application Rules and click the Add icon:



4. Type the Name and a Script for the rule. Click OK when done:



A good example of an application rule script is shown as follows. This script directs requests to a specific load balancer pool according to a domain name. The following rule directs requests to `foo.com` to `pool_1`, and requests to `bar.com` to `pool_2`.

```

acl is_foo hdr_dom(host) -i foo
acl is_bar hdr_dom(host) -i bar
use_backend pool_1 if is_foo
use_backend pool_2 if is_bar

```

You can also look into more rule syntax examples at VMware's NSX online documentation at:

<http://pubs.vmware.com/nsx-63/topic/com.vmware.nsx.admin.doc/GUID-A5779D43-AC0F-4407-AF4A-0C1622394452.html>

Virtual private networks

Virtual private networks (VPNs) allow you to securely connect a remote device or a remote site to your corporate infrastructure. NSX Edge supports three types of VPN connectivity. SSL VPN-Plus allows remote users to access corporate applications securely. IP-SEC VPN offers site-to-site connectivity and L2 VPN allows you to extend your data center by allowing virtual networks to span across data centers securely.

SSL VPN-Plus

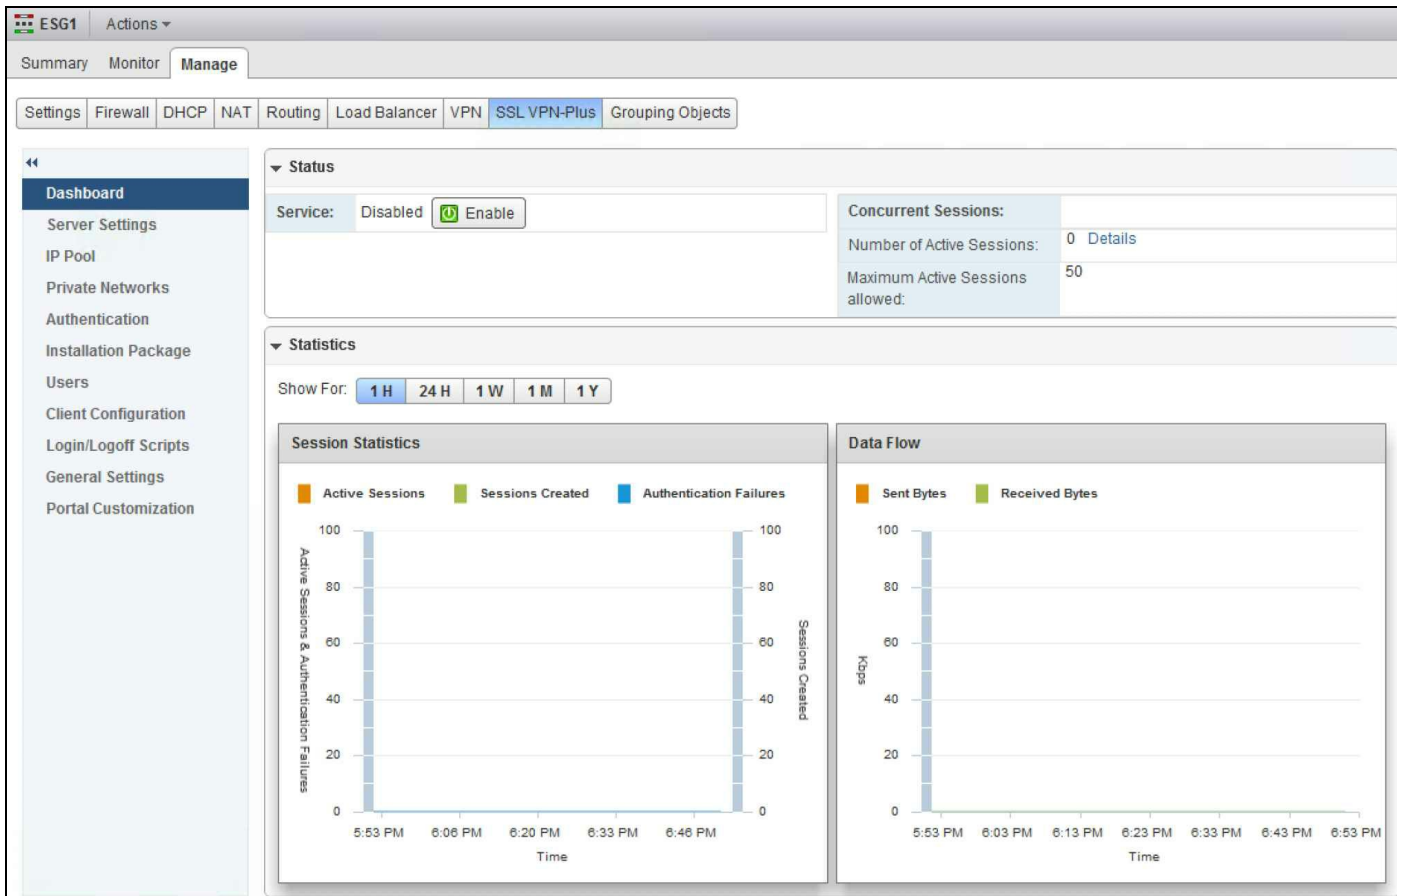
SSL VPN-Plus allows remote users to securely access applications and servers in a private network. There are two modes in which SSL VPN-Plus can be configured: network access mode and web access mode.

In the network access mode, a remote user can access the internal private network securely. This is done by virtue of a VPN client that the remote user downloads and installs on their operating system. In web access mode, the remote user is able to access the private networks without any VPN client software.

SSL VPN-Plus network access mode

Before we begin configuring for the network access mode, make sure that the SSL VPN gateway is accessible externally over port 443. We will systematically configure the network access mode by following multiple steps:

1. Go to Home | Networking & Security | NSX Edges, and double click on the Edge gateway services device, and then navigate to the Manage | SSL VPN-Plus tab:



2. Click on Server Settings and click the Change button to select the IP on which the server will respond:

The 'Change Server Settings' dialog box is shown with the following configuration:

- IPv4 Address: 192.168.1.99 (Primary)
- IPv6 Address: None
- Port: 443
- Cipher List: AES128-SHA
- Server Certificate: Use Default Certificate

Common Name	Issued To	Valid Dates
		7/11/2016 2:59:57 PM - 7/6/2026 3...
		7/11/2016 2:59:57 PM - 7/6/2026 3...
vcenter.lab.local	vcenter.lab.local	7/11/2016 3:09:43 PM - 7/6/2026 3...
VSM SOLUTION 11005fad-bd:	VSM SOLUTION 11005fad-bd:	5/2/2017 7:08:28 PM - 4/8/2117 7:...

Buttons: OK, Cancel

3. Select the IPv4 Address or IPv6 Address, the desired Port number, and the encryption method. If you have a certificate installed, select that if needed.

4. Click OK when done.
5. We will now continue to add an IP pool to provide an IP address to the remote user when a VPN connection is established.

6. On the IP Pools tab, click the + icon:

The screenshot shows a dialog box titled "Add Static IP Pool". It contains the following fields and options:

- IP Range:** Two text boxes separated by "To".
- Netmask:** A text box.
- Gateway:** A text box with the note "This will add an IP address in na0 interface."
- Description:** A large text area.
- Status:** Radio buttons for "Enabled" (selected) and "Disabled".
- Advanced:** A section header.
- Primary DNS:** A text box.
- Secondary DNS:** A text box.
- DNS Suffix:** A text box with the example "Example.eng.vmware.com" below it.
- WINS Server:** A text box.

At the bottom right, there are "OK" and "Cancel" buttons.

7. Type the IP Range, Netmask, and the Gateway that is typically the external interface of the NSX Edge gateway.
8. Type in a Description for the pool and then set the status to Enabled to enable the pool. You may also add custom DNS settings and provide a WINS server if needed.
9. Click OK when done.
10. We will now add the private network the remote VPN user will be able to access. Click on the Private Network tab on the left and click on the + icon:

11. Enter a Network in the CIDR format and a Description. Specify if the traffic should traverse the tunnel or bypass the tunnel and be sent directly to the private server.
12. A tunnel here indicates the SSL VPN-Plus-enabled Edge gateway. If you choose to Send Traffic over the tunnel, leave the Enable TCP Optimization checked to optimize the internet speed, followed by specifying the port numbers for which the traffic will be optimized. Traffic for ports not listed will not be optimized.
13. For multiple ports, you have to create multiple private networks pointing to the same subnet with a different port each time.
14. Select Enabled for Status.
15. Click OK when done.
16. We will now configure an authentication mechanism for users who will be able to access the VPN. The Edge services gateway SSL VPN-Plus supports external authentication mechanisms such as Active Directory, LDAP, Radius, and RSA. Click on the Authentication tab and click the + icon:

Add Authentication Server

Authentication Server Type: AD

Enable SSL:

IP Address: *

Port: * 389

Timeout: 10 Sec(s)

Status: Enabled Disabled

Advanced:

Search base: *

Bind DN: *

Bind Password:

Retype Bind Password:

Login Attribute Name: * sAMAccountName

Search Filter: * objectClass=*

Use this server for secondary authentication

Terminate Session if authentication fails

OK Cancel

17. Select the appropriate Authentication Server Type and fill out the information. Wherever applicable, ensure that you Enable SSL to encrypt all traffic.
18. The maximum Timeout for a VPN connection authentication is 3 minutes and is non-configurable.

We will now create an installation package for the end user. This installation package contains the VPN software client required to make the connection.

1. Click on the Installation Package tab and click the + icon:

The screenshot shows a dialog box titled "Add Installation Package". At the top, there is a "Profile Name" field with a red asterisk indicating it is required. Below this is a table with two columns: "Gateway" and "Port". The "Port" column contains the value "443". To the right of the table are "OK" and "Cancel" buttons. Below the table, there is a section "Create installation packages for:" with three checkboxes: "Windows" (checked), "Linux", and "Mac". Below that is a "Description" text area. The "Status" section has two radio buttons: "Enabled" (selected) and "Disabled". The "Installation Parameters for Windows:" section contains several checkboxes: "Start client on logon", "Allow remember password", "Enable silent mode installation", "Hide SSL client network adapter", "Hide client system tray icon", "Create desktop icon" (checked), "Enable silent mode operation", and "Server security certificate validation". At the bottom of the dialog are "OK" and "Cancel" buttons.

2. Type a Profile Name for the installation package. Type the FQDN or the external IP address for the Edge gateway server. This is the IP the client will connect to. If you need to bind additional Edge uplink interfaces, click the + icon and add them.
3. Select the operating system for which the package needs to be created. By default, the package is created for Windows.
4. Select Enabled to advertise and display the installation package on the installation package page.
5. Customize the installation package by choosing your preferred parameter for Windows.
6. Click OK when done.
7. Let's now add a remote user. Click on Users and click the + icon:



Some useful LDAP configuration options are listed at the below URL:
<http://pubs.vmware.com/nsx-63/topic/com.vmware.nsx.admin.doc/GUID-A5A4B31A-4A38-4CF0-8776-FA40B52544D0.html>

Add User ?

User ID: *

Password: *

Re-type Password: *

First Name:

Last Name:

Description:

Password Details:

Password never expires:

Allow change password:

Change password on next login:

Status: Enabled Disabled

OK Cancel

8. Fill in the form appropriately and click OK when done.

9. Now that the setup is done, let's enable the SSL VPN-Plus service. Click on Dashboard and click Enable and answer the prompt to enable the service. Once enabled, open a browser and access the Edge gateway services over HTTPS. Log in using the username you created in Step 7 to download the VPN client. Log in to the VPN client based on the user and authentication mechanism applicable:

Summary Monitor **Manage**

Settings Firewall DHCP NAT Routing Load Balancer VPN **SSL VPN**

«

Dashboard

Server Settings

IP Pool

Private Networks

Authentication

Service enabled successfully!

▼ Status

Service: Enabled

IPSEC VPN

The NSX Edge service gateway supports site-to-site IPSEC VPN that allows you to connect an NSX Edge services gateway-backed network to another device at the remote site. NSX Edge can establish secure tunnels with remote sites to allow secure traffic flow between sites. The number of tunnels an Edge gateway can establish depends on the size of the Edge gateway deployed.

A compact Edge gateway can create a maximum of 512 tunnels. A large Edge gateway can create a maximum of 1600 tunnels, while a quad-large can handle a maximum of 4096 tunnels. An X-Large Edge gateway can handle up to 6000 tunnels.

NSX supports AES (AES128-CBC), AES256 (AES256-CBC), Triple DES (3DES192-CBC), DH-2 (Diffie-Hellman group 2), DH-5 (Diffie-Hellman group 5), and AES-GCM (AES128-GCM) IPSEC VPN algorithms.

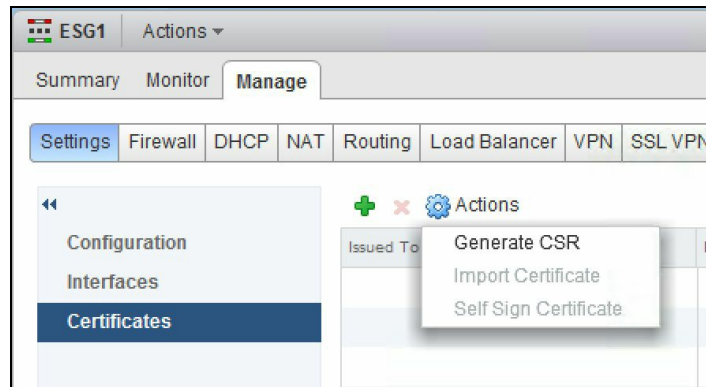
Before we begin configuring our IPSEC VPN, ensure dynamic routing is disabled on the Edge uplink to allow specific routes defined for any VPN traffic.



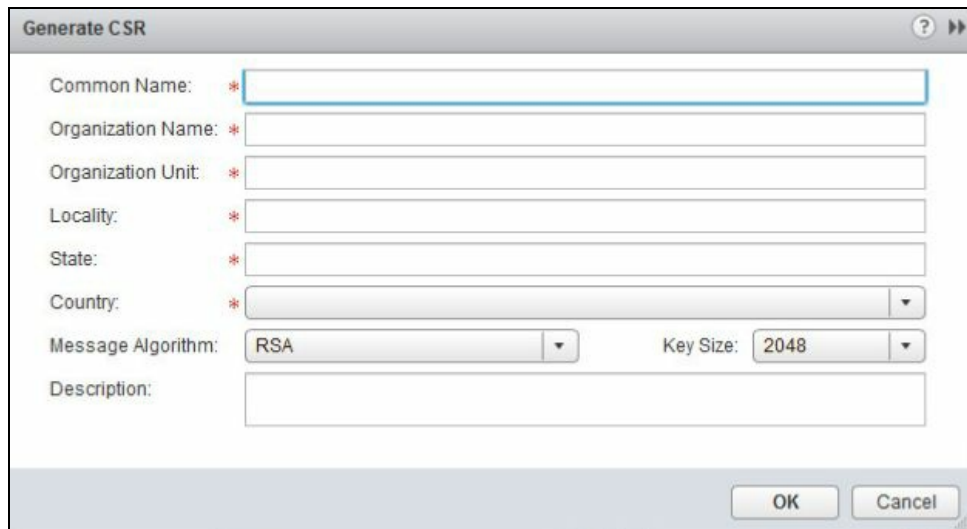
Having dynamic routing enabled causes routes to be updated as the router learns of new routes, which can cause traffic disruption in an IPSEC VPN setup.

Let's begin by generating a certificate to enable certificate authentication. You can import a CA-signed certificate or use Open-SSL to generate a CA-signed certificate. Self-signed certificates cannot be used with an IPSEC VPN. They can only be used with load balancers and SSL VPNs. Perform the following set of steps to generate a certificate to enable certificate authentication:

1. Go to Home | Networking & Security | NSX Edges, and double click on an Edge appliance, and navigate to the Manage tab | Settings | Certificates:

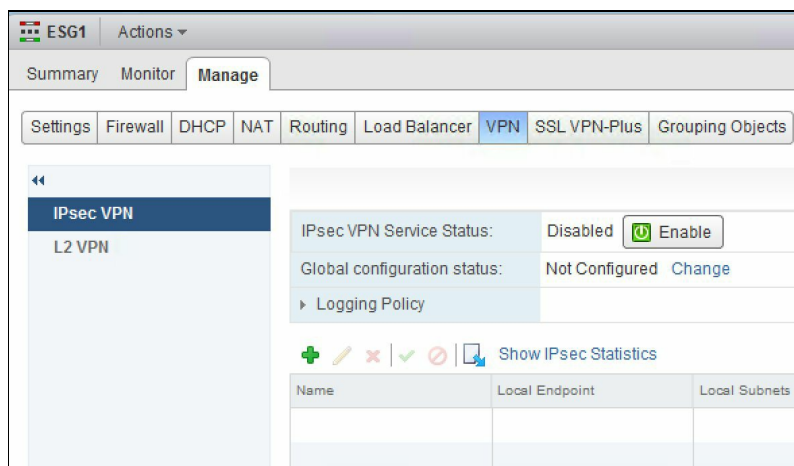


2. Click Actions and click Generate CSR. This generates your CSR:

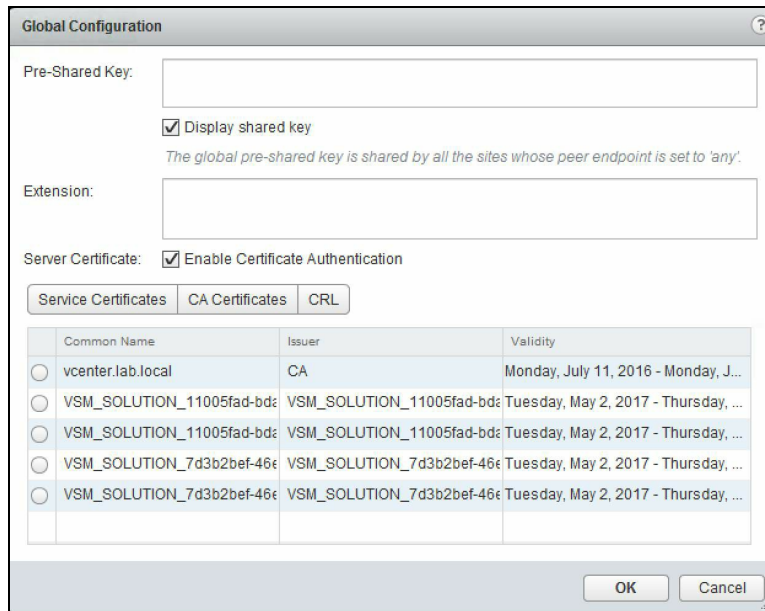


3. Fill out the appropriate details and click OK.

4. We will now set the Global configuration status. Click on the VPN tab and click on IPsec VPN:



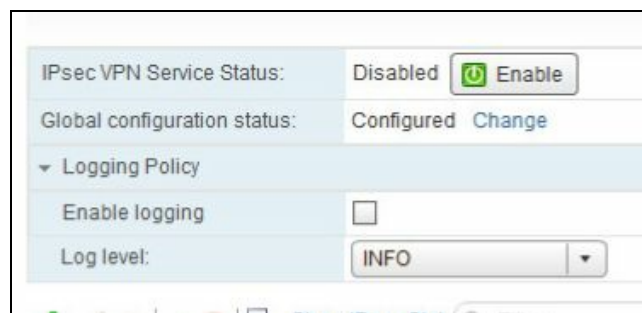
5. Click on Change beside the Global configuration status:



6. Enter the global Pre-Shared Key that is shared by all the sites whose peer endpoints are set to any.
7. Enable Certificate Authentication and select the appropriate certificate. Click OK when done.
8. Next, we enable logging for our IPsec VPN. Expand the logging policy and check Enable logging. Set the appropriate logging level. Increasing the logging level increases the amount of data stored on the Edge appliance and can negatively impact performance:



A best practice is to configure a syslog server so that all logs can be exported to it and not stored locally at the Edge services gateway appliance.



9. Next, we configure the IPsec VPN parameters. Click the + icon:

10. Type the Name of the tunnel and enter the Edge gateway IP address in the Local Id field. This will be the Peer ID on the remote site.
11. Type the IP address of the Local Endpoint, which is the IP address of your Edge gateway.
12. Next, type the Local Subnets that are shared between two sites.
13. Enter the Peer ID to uniquely identify the peer site. This ID must be the common name in the peer's certificate for any peers using certificate authentication. Ideally, it is best practice to stick with the IP address as the peer ID.
14. In the Peer Endpoint field, type the IP address of the peer site. Next, type the internal IP address of the Peer Subnets.
15. Next, select the appropriate Encryption Algorithm. If you require anonymous sites to connect to your VPN, enter the Pre-Shared Key to allow such capability. Clicking Display shared key displays the key on the peer site.

16. Next, select the cryptography scheme that allows the NSX Edge and the peer site to establish a shared secret over an insecure channel.
17. In the Extension field, type one of the following:
 - `securelocaltrafficbyip=IPAddress`: This redirects Edge's local traffic through the IPSEC tunnel. This is the value by default.
 - `passthroughSubnets=PeerSubnetIPAddress`: This allows for overlapping subnets on both sides.
18. Click OK when done. NSX Edge now creates a tunnel between the local subnet and the peer subnet.
19. Click Enable and Publish Changes when done.

L2 VPN

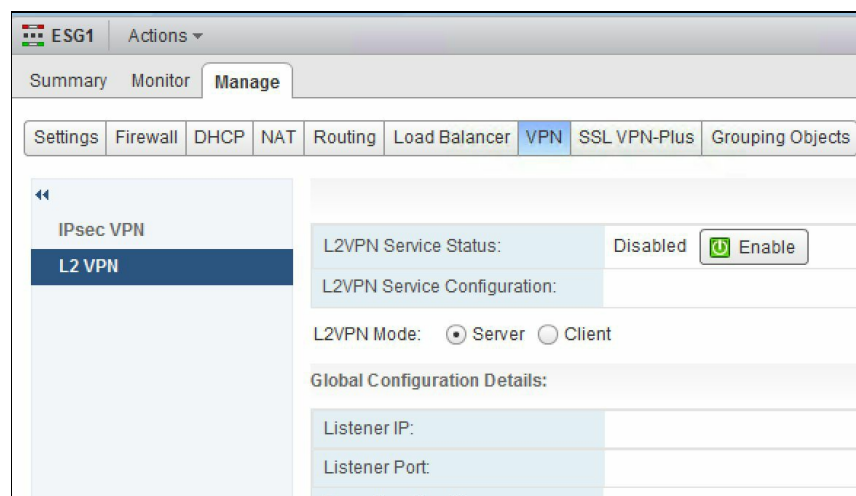
An L2 VPN allows you to stretch multiple logical networks across multiple sites. The networks can be both traditional VLANs and VXLANs. In such a deployment, a virtual machine can move between sites without a change in its IP address.

An L2 VPN is deployed as a client and server where the destination Edge is the **server** and the source Edge is the **client**. Both the client and the server learn the MAC addresses of both local and remote sites. For any sites that are not backed by an NSX environment, a standalone NSX Edge gateway can be deployed.

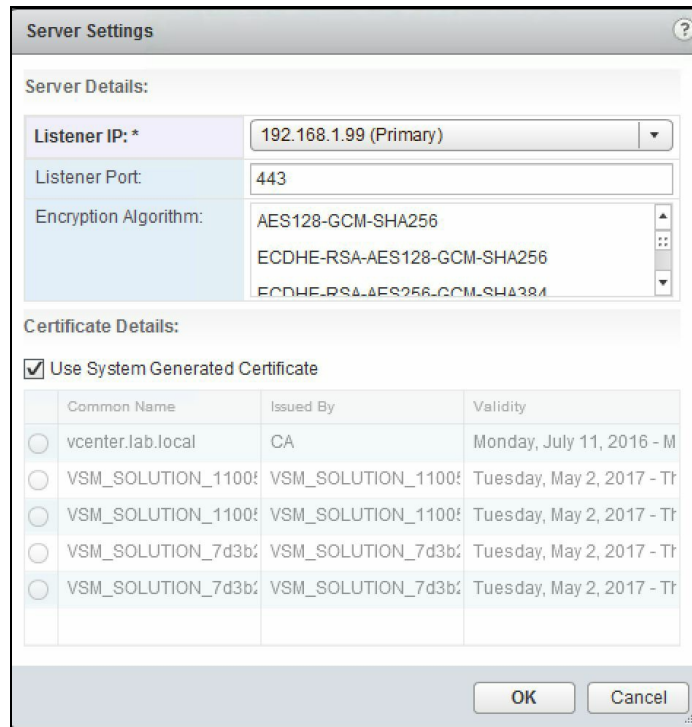
Before we begin configuring L2 VPN, ensure that a sub-interface is added to a trunk interface of the NSX Edge. You can learn more on adding a sub-interface in the Edge configuration section of this chapter.

To begin configuration for an L2 VPN, follow these steps:

1. Go to Home | Networking & Security | NSX Edges, and double click the Edge appliance, and navigate to Manage | VPN | L2 VPN:



2. To configure the L2 VPN server that is the destination Edge, select the L2 VPN Mode to Server and click on the Change button:



3. In the Listener IP, select the primary or the secondary IP of the Edge. Change the port if needed and select the appropriate Encryption Algorithm.
4. Select the certificate if configured and click OK when done. You can always use a system-generated certificate if needed.
5. Next, we add a peer site. Under the Site Configuration Details tab, click the + icon:

6. Enter a unique Name for the peer site and a Description if necessary. Enter the User ID and Password for the peer site for authentication.
7. Click on Select Sub Interfaces and select the appropriate interfaces. These are the interfaces that are stretched with the client.
8. If the default gateway for virtual machines is the same across both sites, type in the gateway IP in the Egress Optimization Gateway Address section. This will allow local routing and increase performance.
9. Enabling the Enable Unstretched Networks field allows you to identify network subnets that you do not wish to extend between two sites.
10. Make sure the top Enable Peer Site tab is checked and then click OK when done.

Let's go ahead and configure the L2 VPN client. The L2 VPN client is the source that NSX Edge initiates communication with alongside the destination Edge (L2 VPN server):

The screenshot displays the management interface for an ESG1 device, specifically the VPN configuration page. The top navigation bar includes 'Summary', 'Monitor', and 'Manage'. Below this, a secondary navigation bar lists various services: 'Settings', 'Firewall', 'DHCP', 'NAT', 'Routing', 'Load Balancer', 'VPN', 'SSL VPN-Plus', and 'Grouping Objects'. The 'VPN' tab is selected. On the left, a sidebar shows 'IPsec VPN' and 'L2 VPN', with 'L2 VPN' being the active selection. The main configuration area shows the following settings:

- L2VPN Service Status:** Disabled, with an 'Enable' button.
- L2VPN Service Configuration:** A field with a 'Delete Configuration' button.
- L2VPN Mode:** Radio buttons for 'Server' and 'Client', with 'Client' selected.
- Global Configuration Details:** A section with a 'Change' button, containing several input fields:
 - Server Address:
 - Server Port:
 - Encryption Algorithm:
 - Stretched Interface:
 - Egress Optimization Gateway Address:
 - Unstretched Networks:
 - User Id:
 - CA Certificate:
 - Proxy Settings:
- Fetch Status:** A button at the bottom right of the configuration area.

1. In the L2 VPN mode, select Client and click the Change button:

The screenshot shows a 'Client Settings' dialog box with two tabs: 'Client Details' and 'Advanced'. The 'Client Details' tab is active. It contains the following fields and options:

- Server Address: ***: A text input field.
- Server Port:**: A text input field containing '443'.
- Encryption Algorithm:**: A dropdown menu with 'AES128-GCM-SHA256' selected and 'ECDHE-RSA-AES128-GCM-SHA256' visible below it.
- Stretched Interfaces:**: A text input field with a 'Select Sub Interfaces' button to its right.
- Egress Optimization Gateway Address:**: A text input field with a hint: 'Example: Comma separated list of IP address' and 'Ex:191.1.1.1,192.1.1.1'.
- Unstretched Networks:**: A checkbox labeled 'Enable Unstretched Networks' which is currently unchecked, with a text input field below it. A hint below the field reads: 'Example: Comma separated list of networks' and 'Ex:192.168.10.0/24,192.168.20.0/24'.
- User Details:** A section containing three text input fields: 'User Id: *', 'Password: *', and 'Confirm Password: *'.

At the bottom right of the dialog box are 'OK' and 'Cancel' buttons.

2. Enter the Server Address, which is the L2 VPN server IP or FQDN. Select the same Encryption Algorithm as set on the server and select the appropriate Stretched Interfaces.
3. Type in the Egress Optimization Gateway Address and then type the user credentials.
4. If the client NSX Edge does not have a direct internet connection, it can reach the server NSX Edge over a proxy server, which can be set in the advanced settings.
5. Click OK and Publish Changes.
6. Enable the L2 VPN service on the client; this will establish the L2 VPN connectivity between the sites.

More Edge services and configurations

In this section, we will look at a few configuration steps for some common actions that you will perform on the Edge services gateway. In a production environment, you will often perform these configurations either during initial setup or after.

We will be looking at adding a sub-interface, a force-sync NSX Edge with NSX Manager, configuring remote syslog servers, and redeploying an NSX Edge appliance.

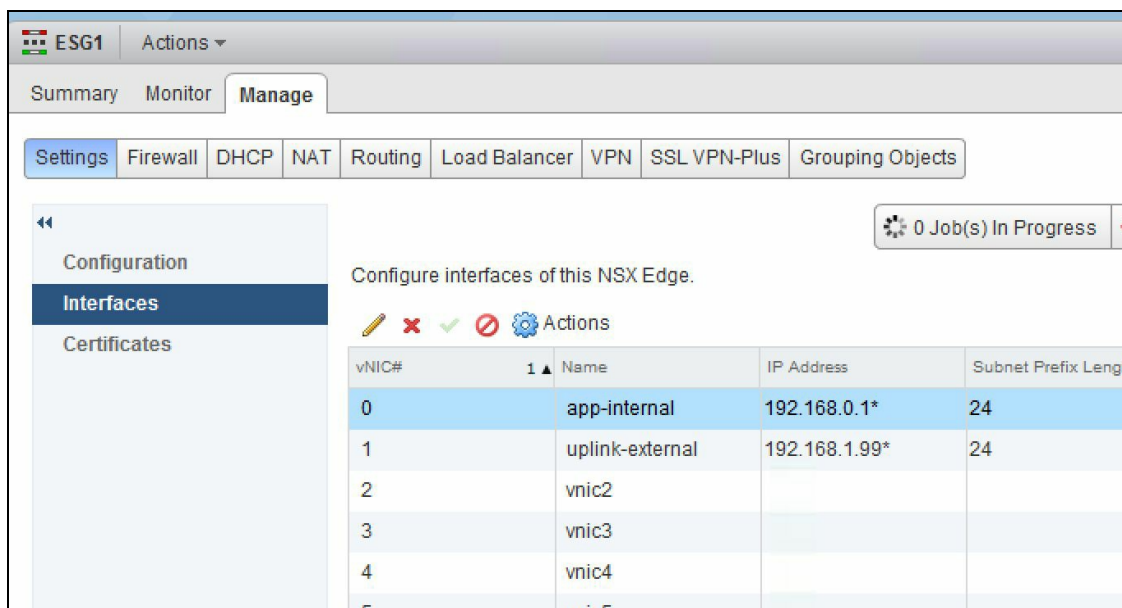
Adding a sub-interface

An NSX Edge services gateway can have up to ten internal, external (physical), or trunk interfaces, while an Edge distributed router can have up to eight uplink interfaces and up to a thousand internal (sub) interfaces.

A sub-interface, or an internal interface, is a logical interface that is created and mapped to the physical interface. Sub-interfaces are simply a division of a physical interface into multiple logical interfaces. This logical interface uses the parent physical interface to move data. Remember that you cannot use sub-interfaces for HA because a heartbeat needs to traverse a physical port from one hypervisor to another between the Edge appliances.

To configure a sub-interface, do the following:

1. Go to Home | Networking & Security | NSX Edges, and double click an Edge appliance, and navigate to Manage | Settings | Interfaces:



2. Select the Edit icon to edit an interface:

Edit NSX Edge Interface
?

vNIC#: 2

Name: * vnic2

Type: Trunk

Connected To: VM Network Change Remove

Connectivity Status: Connected Disconnected

Sub Interfaces:

+ ✎ ✕
Filter

vNIC#	Name	Network	VLAN / VNI	Tunnel ID	Status

* indicates new SubInterface

MAC Addresses:

You can specify a MAC address or leave it blank for auto generation. In case of HA, two different MAC addresses are required.

MTU:

OK
Cancel

3. Enter a Name for the interface. Select Trunk from the Type drop-down menu.
4. Click Select to select a port group to connect to this interface. Click the + sign to add a sub-interface to this trunk interface:

Add Sub Interface

Enable Sub-Interface

Name: *

Tunnel Id: *
Values from 1-4093

Backing Type: Network ▾

Network: Select Remove

Configure Subnets:

+ ✎ ✕

Primary IP Address	Secondary IP Address	Subnet Prefix Length

MTU:
Values from 68-9000

Options: Enable Send Redirect
Reverse Path Filter: Enabled ▾

OK Cancel

5. Enter a Name and a Tunnel ID as applicable in your environment.
6. In the Backing Type, select a VLAN for a VLAN network and type in the VLAN ID that this sub-interface should use.
7. For a Network type, select a logical switch or a distributed port group so that NSX can use the VLAN ID for that switch or port group and use it in the trunk.
8. Selecting None creates an internal sub-interface that is used to route packets in a stretched network and an unstretched network.
9. Click + to configure subnets. A sub-interface can have one primary IP and multiple secondary IPs. Enable Send Redirect allows you to convey routing information to hosts. Click OK to return to interface configuration.
10. Enter a MAC address. Sub-interfaces do not support HA, so you may leave this blank to allow for auto-generation.
11. Select the appropriate MTU with a minimum value of 1600.
12. Click OK when done.

Force sync NSX Edge with NSX Manager

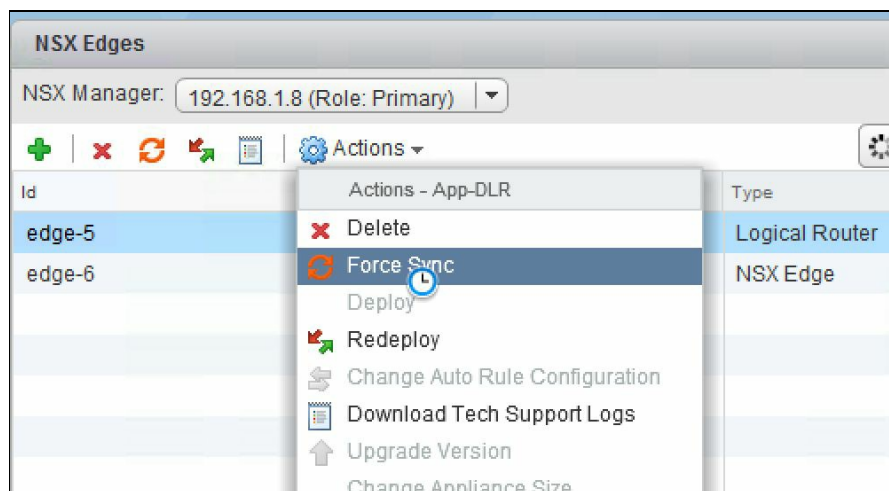
The NSX Manager holds all the configuration data for each NSX Edge and keeps check of the synchronization via regular polling between the NSX Manager and the Edge services gateway. A good example of an Edge appliance going out of sync is low disk space, where no more new configuration data can be written to the appliance, causing it to go out of sync. Any Edge services gateway that is out of sync can be resynchronized by issuing a force sync request.

Force sync is a feature that synchronizes the Edge configuration from the NSX Manager to all of its components in an environment. A synchronization action is initiated from the NSX Manager to the NSX Edge that refreshes and reloads the Edge configuration. Initiating a force sync action causes minor north-south traffic interruption. East-west traffic remains unaffected. The interruption is due to a reboot of the Edge appliances.

In a cross-vCenter environment, you have to first apply force sync on the NSX Edge at the primary NSX Manager site before applying it on the secondary sites.

To apply force sync, follow these steps:

1. Go to Home | Networking & Security | NSX Edges.
2. Select an Edge device, click on Actions and select Force Sync:

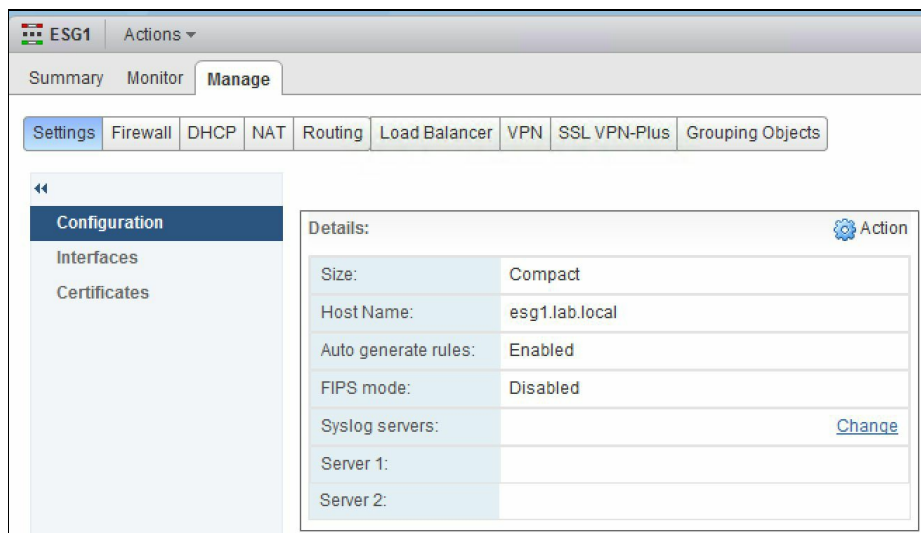


3. Click Yes when the prompt appears to enforce sync action.

Configuring remote syslog servers

VMware recommends configuring syslog servers to avoid log flooding on the Edge appliances. When logging is enabled, logs are stored locally on the Edge appliance and consume space. If left unchecked, this can have a performance impact on the Edge appliance and can also result in the Edge appliance stopping due to lack of disk space.

1. Go to Home | Networking & Security | NSX Edge, and double click an Edge appliance, and navigate to Manage | Settings | Configuration:



2. Click Change in the Syslog servers row:




3. Enter the syslog server's IP address and select the Protocol as applicable. Click OK when done.

Redeploying an NSX Edge

Occasionally you will notice that your Edge appliance is not performing as expected. If force sync does not resolve the issue, then a redeploy is necessary. Redeploying an Edge appliance essentially redeploys the Edge services gateway and is a disruptive action.

In a multi-vCenter environment, you are required to first redeploy the NSX Edge appliance on the primary side before re-deploying the secondary side. It is necessary to redeploy both sides in a multi-vCenter environment:

1. Go to Home | Networking & Security | NSX Edges, and select Edge, and navigate to Actions | Redeploy:



2. Click Yes on the prompt to complete the redeployment process.

Summary

In this chapter, we looked at how to configure different supported routing protocols on the NSX Edge. We started the chapter by looking at configuring DNS and DHCP services. Next, we looked at configuring OSPF and BGP routing protocols. We also looked at configuring logical load balancers that allow us to implement a load balancer for our applications. We then looked at configuring VPNs on our Edge appliance so that remote users can connect securely. Edge supports SSL-VPN, IPSEC VPN, and L2 VPN.

In the [Chapter 6](#), Service Composer, we will look at NSX's security capabilities, including security composer and security groups. We will also learn about mapping security policies to virtual machines.

Service Composer

In this chapter, we will look at the security aspects of NSX. One of the primary use cases of NSX is its ability to offer a comprehensive toolset for data security. We will look at what security composer is and how it can be used to assign security services to your applications. We will look at creating security groups and policies and mapping a security policy to a group. We will next look at data security and install a data security policy. We will then learn about network extensibility and how to integrate third-party services with NSX.

In this chapter, we will cover:

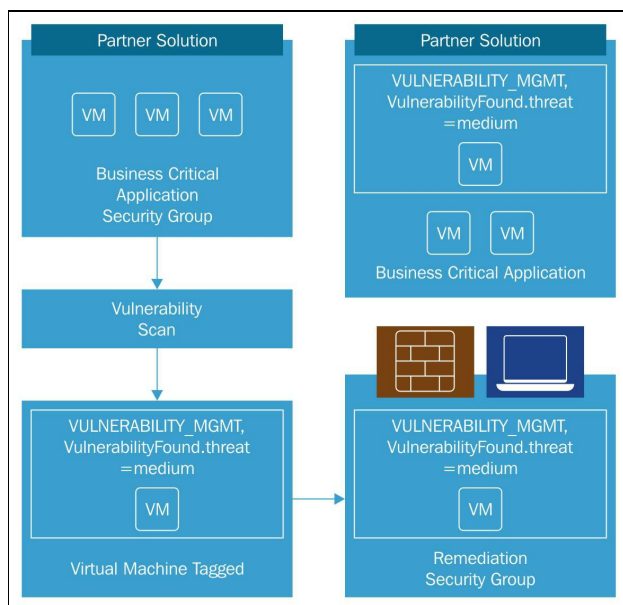
- Service Composer
- Network extensibility

Service Composer

Service Composer provides an administrator with the ability to define a scalable and tiered security policy independent of the underlying infrastructure or routed topology. This is the feature with the NSX platform that allows security to scale and allows for security policies that are enforced at a unit level, protecting virtual to physical or physical to virtual communications and allowing event-driven security actions.

Service Composer consists of security groups and security policies that allow you to provision security services to your virtual machines. Service Composer in effect has mappings between security groups, policies, and virtual machines.

Security groups are a collection of instances that you want to protect. You can group your virtual machines to be part of a security group or can have vCenter objects as part of a security group. You can have a security group that consists of other security groups. You can even define a security group to have instances that have security tags:



Courtesy—VMware

Service Composer helps you consume security services with ease.

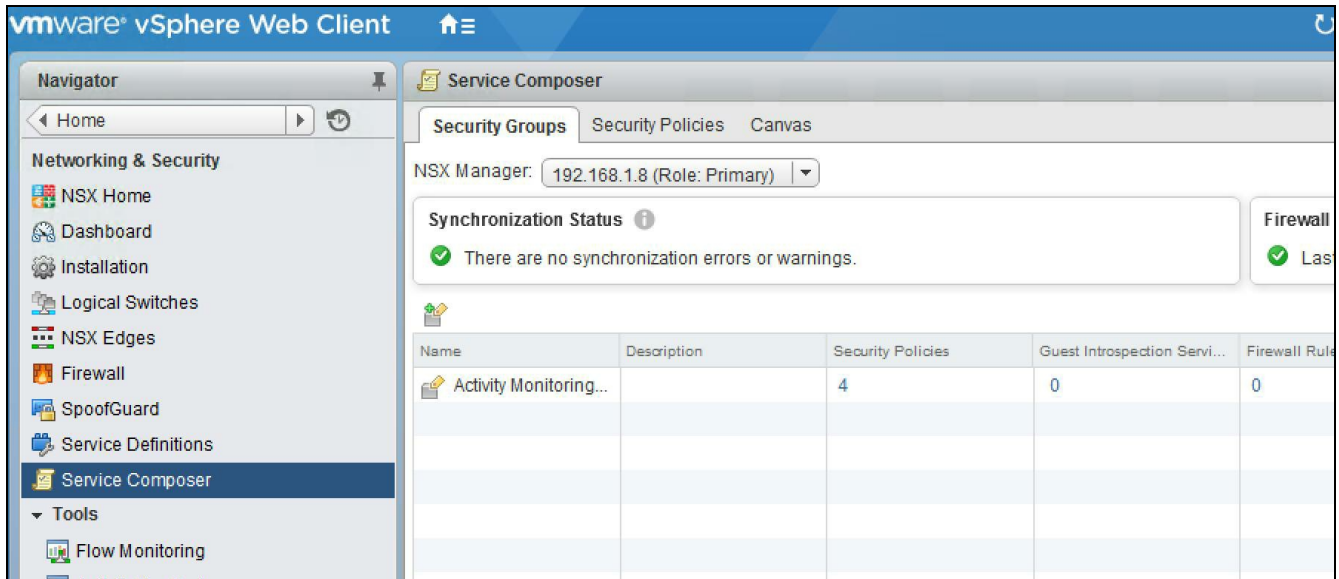
Security policies are collections of security services and their service configurations. Firewall rules, endpoint services, and network introspection services are all part of the services in a security policy. You map security policies to security groups. When a security policy is mapped to a security group, the policy applies to all the virtual

machines that are part of that security group.

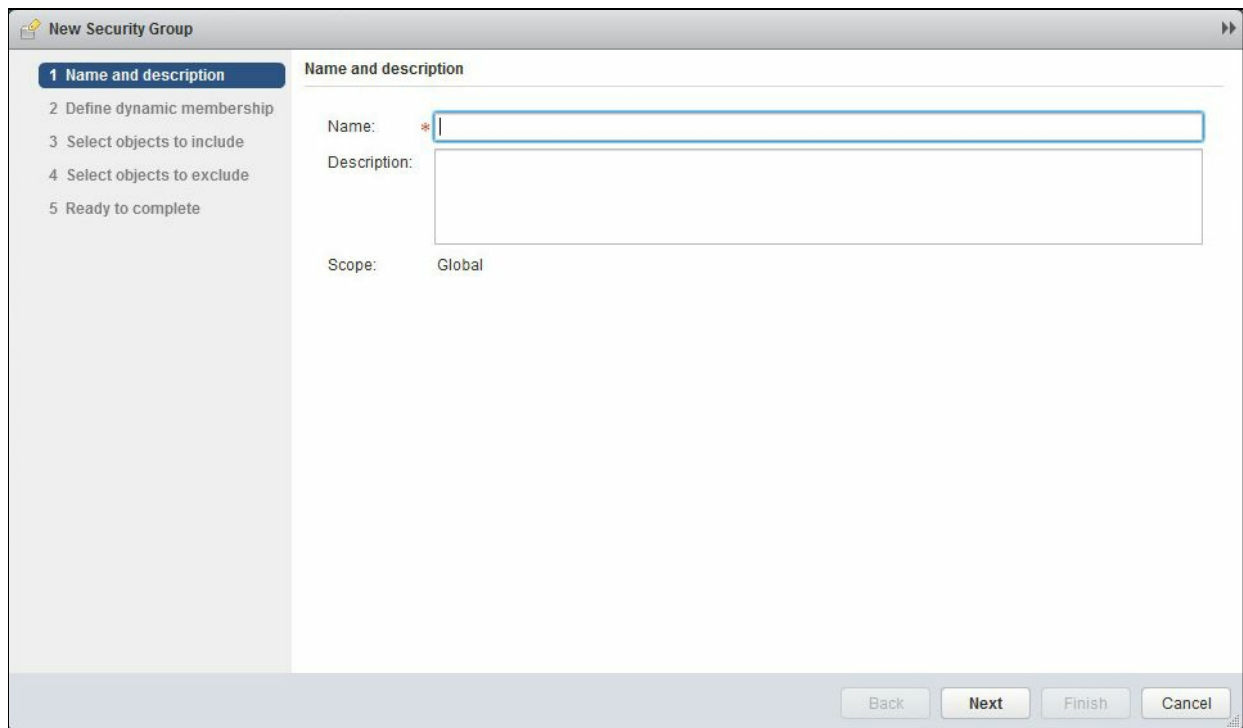
Security groups

We will now look at creating a security group in Service Composer:

1. Go to Home | Networking & Security | Service Composer:



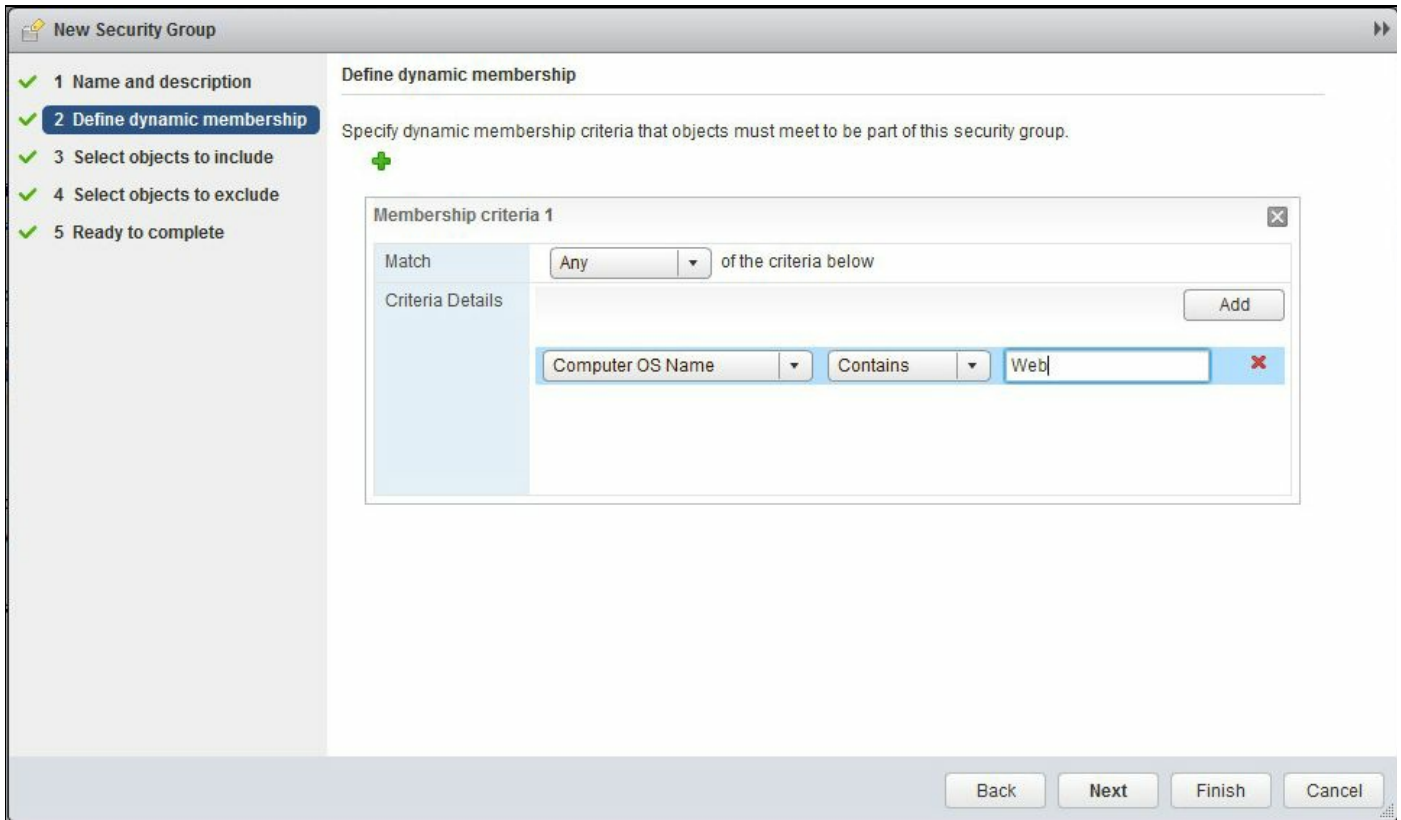
2. Click on the Security Groups tab and click the + icon to add a new security group:



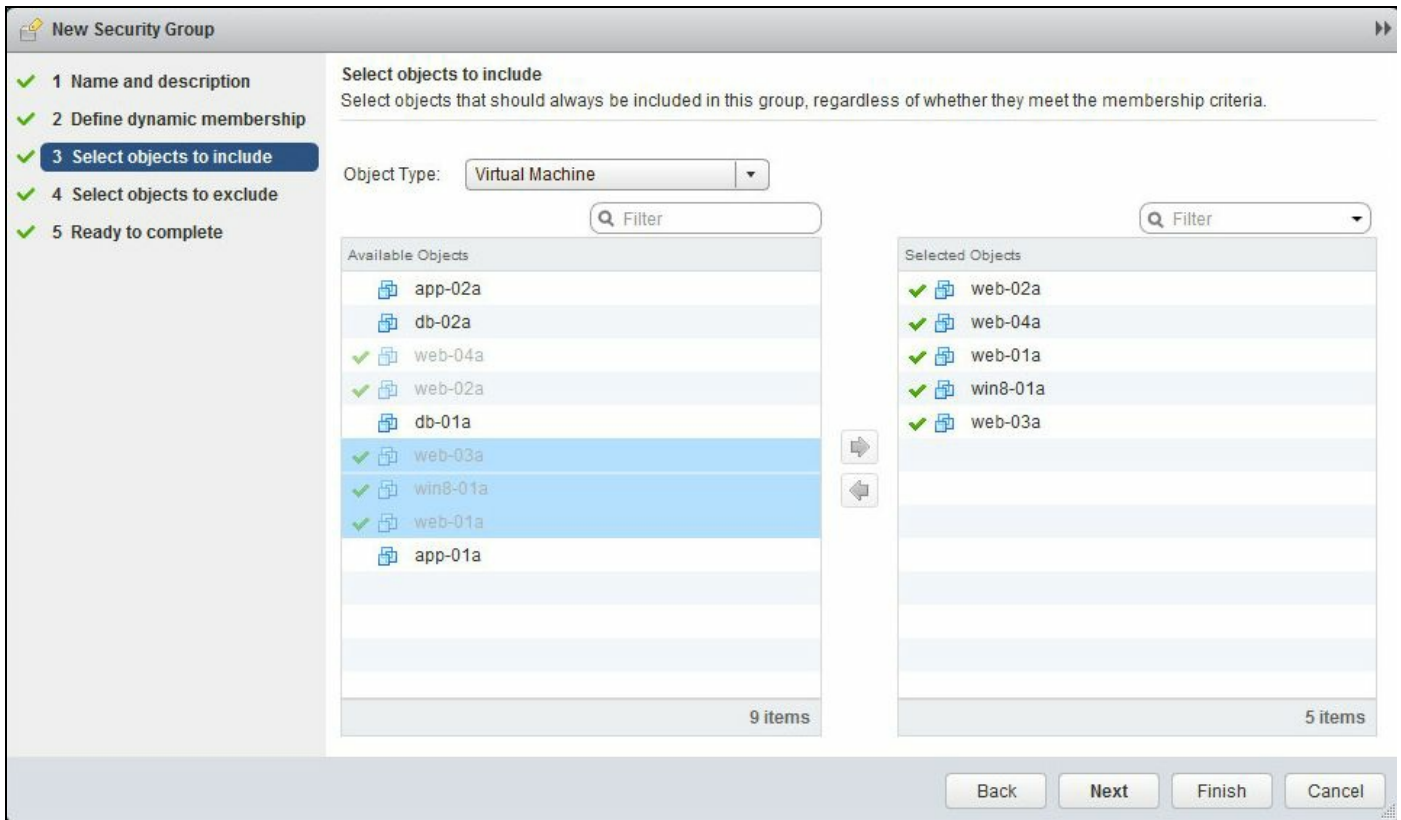
3. Enter the name of the security group. It's a good idea to make the name of the group

similar to the function of the security group. For example, if the group contains instances that belong to your application then name the group `SG-APP`. Click Next.

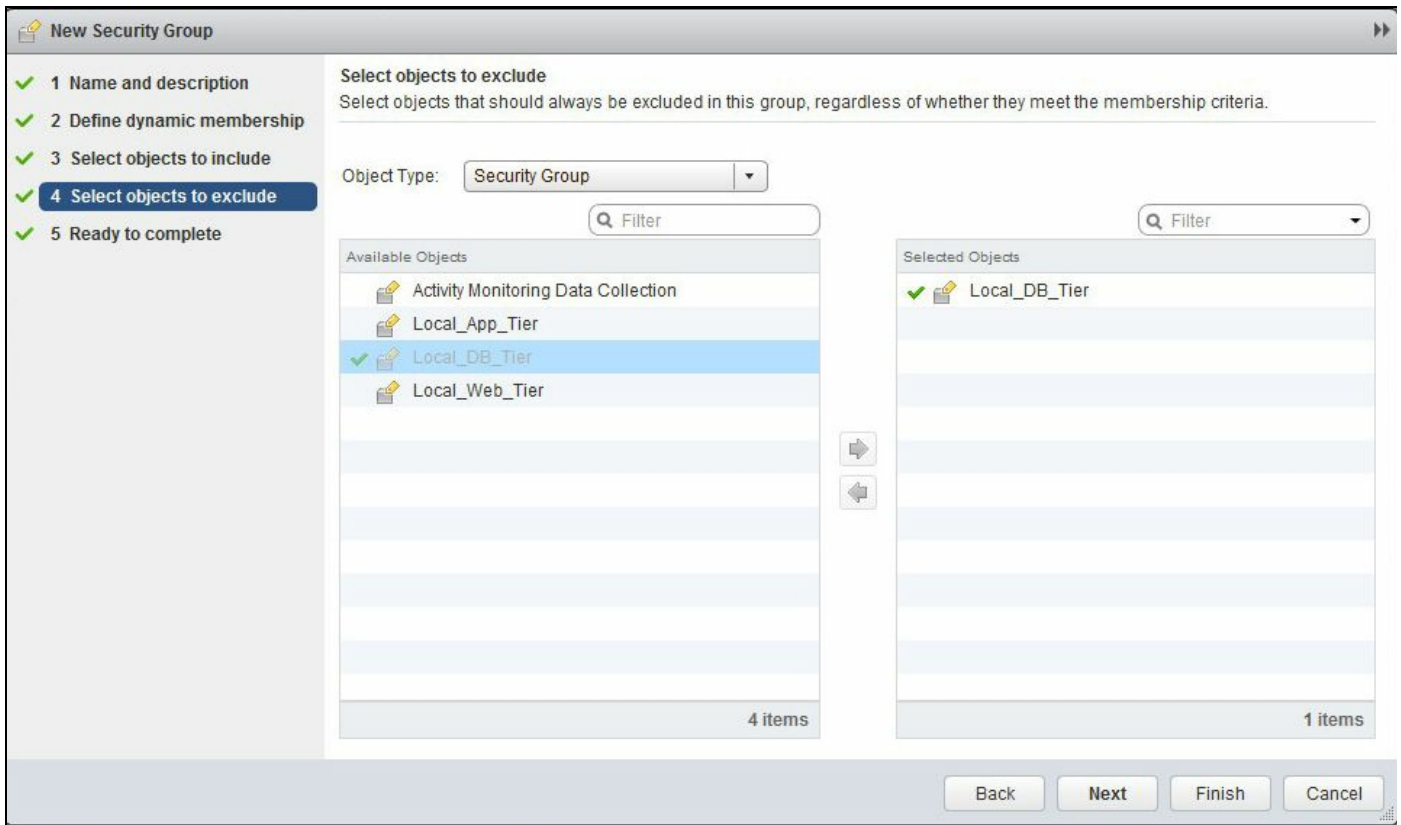
4. Define the criteria that any instance or an object should meet to be added to the security group. Remember that this is a dynamic action so any object that satisfies the criteria will be part of the group:



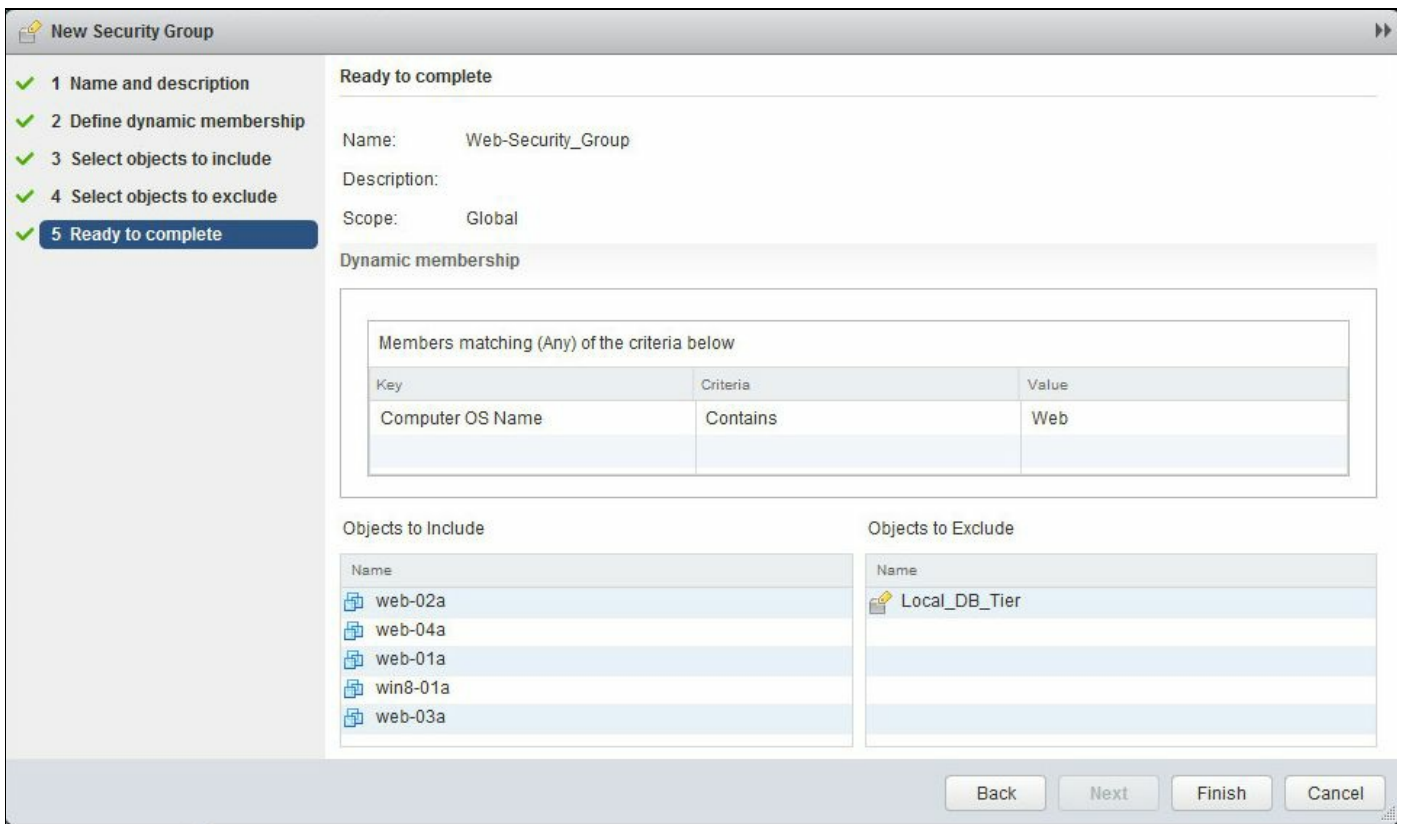
5. Clicking the + sign allows you to add more criteria. Within criteria you can click the Add button to add more criteria. For example, you may want to have the security group applied to all virtual machines where the name contains `web`. Click Next when done.
6. Select the types of object that you want to always include in the group. This includes a range of objects including virtual machines and even other security groups. The selected objects will always be part of the group. Click Next:




7. Select the objects that you want to exclude from the group. For example, you may want to exclude specific virtual machines that need to not have a security group applied to it such as a DMZ web server. Click Next:



8. Review the summary and click Finish to create the security group:



You have now successfully created a security group:



We will now create a security policy and associate it with this security group.

Security policies

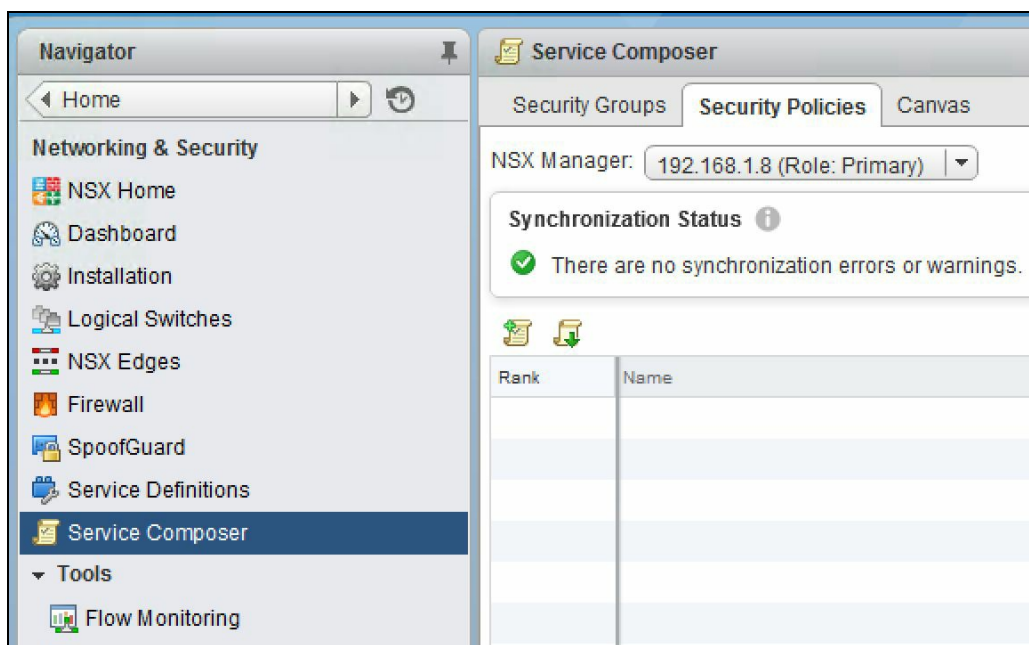
Security policies are sets of rules that apply to a virtual machine, network, or firewall services. Security policies are reusable rulesets that can be applied to security groups. Security policies express three types of rulesets:

- **Endpoint services:** Guest-based services such as anti-virus solutions and vulnerability management
- **Firewall rules:** Distributed Firewall policies
- **Network introspection services:** Network services such as intrusion detection systems and encryption

These rules are applied to all objects and virtual machines that are part of a security group to which this policy is associated. The order of the policies applied depends on the assigned weight. The weight of a policy determines the order in the list of policy rules that apply.

You need to ensure that you have some introspection services installed, such as a distributed firewall that has been installed and configured. We looked at enabling distributed firewalls in [Chapter 3, NSX Installation and Configuration](#):

1. Go to Home | Networking & Security | Service Composer:



2. Click on the Security Policies tab.

3. Click on the + icon on the left to add a new security policy:

The screenshot shows the 'New Security Policy' configuration window. The left sidebar indicates the current step is '1 Name and description'. The main form fields are: 'Name' (App Security Policy), 'Description' (empty text area), 'Inherit security policy' (checkbox), and 'Parent policy' (dropdown menu). An 'Advanced options' section is visible below. The bottom right corner contains 'Back', 'Next', 'Finish', and 'Cancel' buttons.

4. Enter a Name and a Description for the policy. Select Inherit security policy to allow the policy to inherit rules from another policy. This allows you to better manage policies and reduce unnecessary duplication of rules.
5. Under Advanced options, set the weight accordingly. A policy with the highest weight is given the first precedence. By default, a new policy has the highest weight. Click Next when done.

6. Under the Guest Introspection Services, click the + icon to add a new guest introspection service:

Add Guest Introspection Service

Name: App-vms

Description:

Action: Apply Block

Service Type: Anti Virus

Service Name: VMware Data Security

Service Profile: None

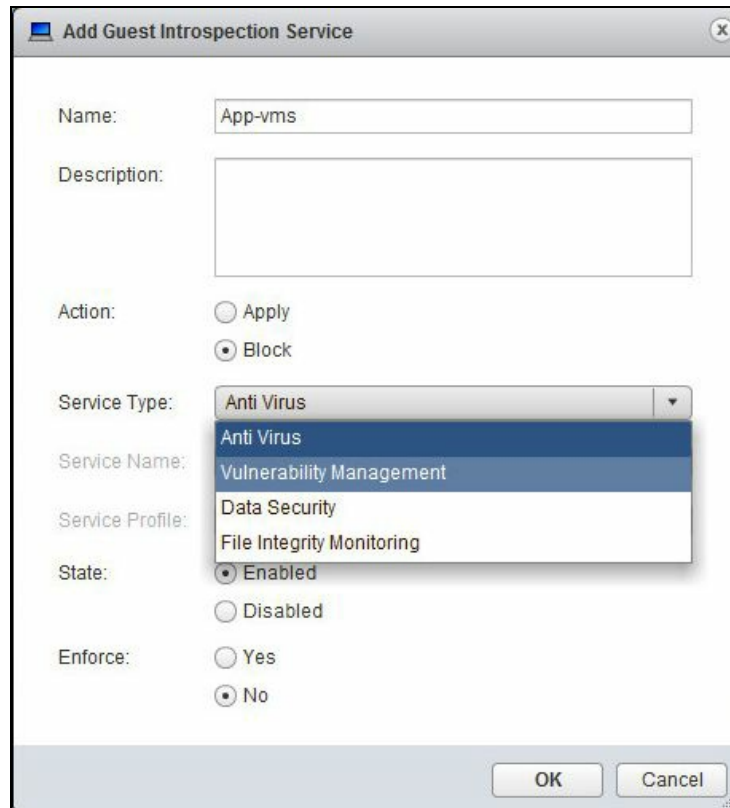
State: Enabled Disabled

Enforce: Yes No

OK Cancel

7. Enter a Name and a Description for the service. Enter the default action for the service:

- If you choose the Apply action ensure that VMware Data Security is configured and in place.
- If you choose to Block the service, select the Service Type:

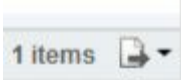


8. Select the Service Name applicable for guest introspection. You will see more service names offered as provided by a third-party vendor who integrates with NSX. We will look at integration of third-party services later.

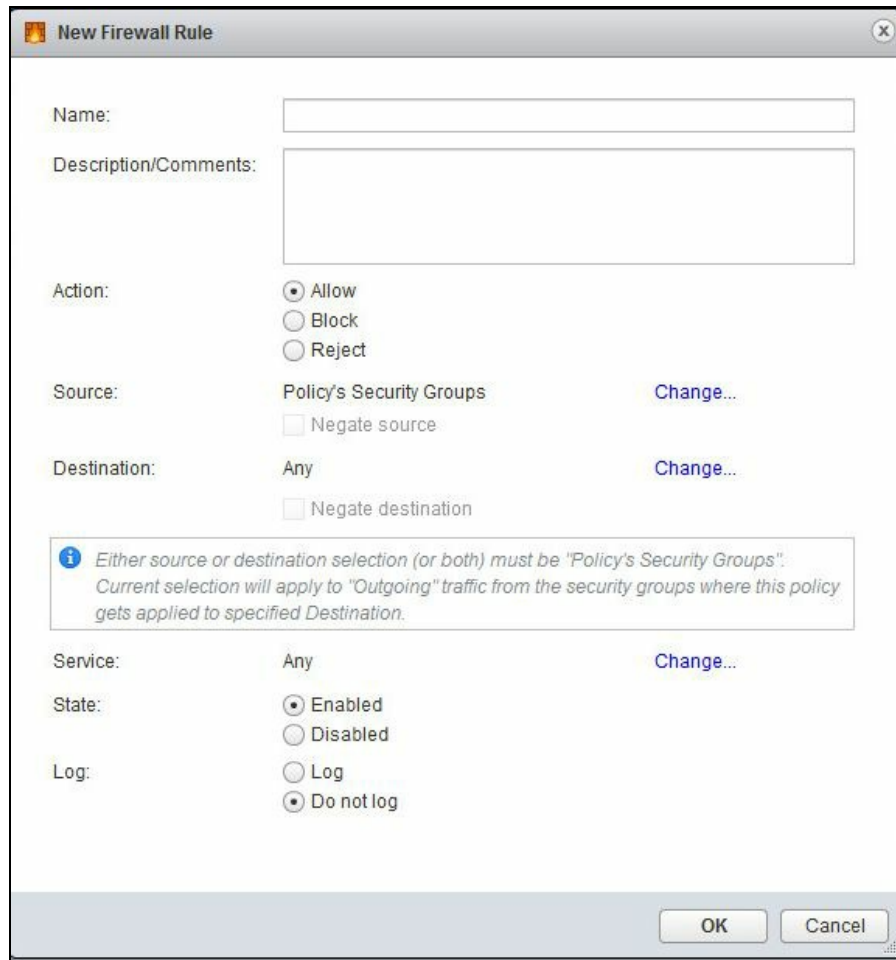


Make sure the required partner services have been registered with NSX manager for Service Type and Service Name to populate. We will look at setting up partner services later in this chapter. To learn more about registering partner services, refer to the following URL: <http://pubs.vmware.com/nsx-63/topic/com.vmware.nsx.admin.doc/GUID-EA477D96-E2D3-488B-90AA-2F19B4AE327D.html>.

Select a State option to enable or disable a defined service.

9. Select a value for Enforce. Enforcing allows you to prevent any changes to this policy. If another policy inherits from this policy, then Enforce forces this policy to be applied before the new policy is applied.
10. Click OK and click Next when done. You may also export services by clicking the  icon.

11. Click the + icon to create a new firewall rule:



12. Enter a Name and a Description for the rule and set an appropriate action.
13. Click the appropriate Change... to select the source for which this policy applies. You can specify a specific security group that this policy applies to.
14. Click the appropriate Change... to select the destination security group for this policy.
15. Clicking Negate destination allows the source to access all security groups except the one you have selected.
16. Click the appropriate Change... to select the services or service groups this rule will apply to. You can also create custom services if needed.
17. Select the state for this rule and the log setting, and click Ok. The rules defined here show up in the firewall table.



VMware recommends that you do not edit Service Composer rules that are listed in the firewall table as the rules are managed by Service Composer. Doing so will require you to re-synchronize Service Composer rules with firewall rules by selecting the Synchronize firewall rules option from the Action menu in the Security Policies tab.

18. Click OK when done and click Next.
19. Click the + symbol to add a network introspection service. To define the network introspection service you need to have specific service providers (vendor appliances) integrated in your environment. We will look at third-party integration in the upcoming sections:

The screenshot shows a configuration dialog box for a network policy. It contains the following fields and options:

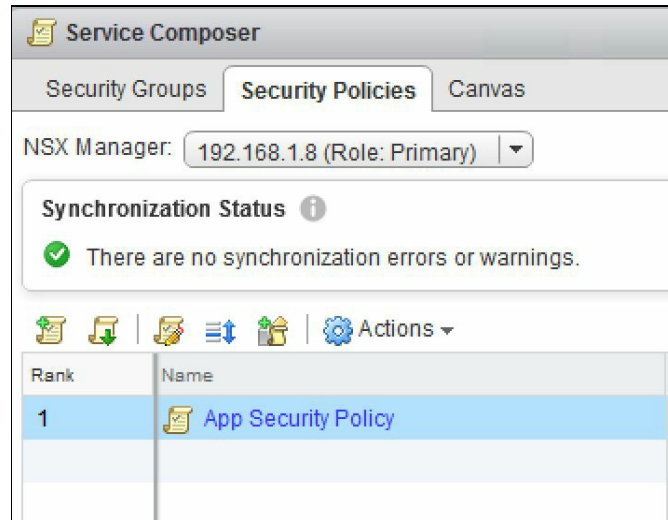
- Name:** A text input field.
- Description:** A larger text input field.
- Action:** Two radio buttons: Redirect to service and Do not redirect.
- Service Name:** A dropdown menu.
- Profile:** A dropdown menu.
- Source:** A dropdown menu set to "Policy's Security Groups" with a "Change..." link. Below it is a checkbox for "Negate source".
- Destination:** A dropdown menu set to "Any" with a "Change..." link. Below it is a checkbox for "Negate destination".
- Service:** A dropdown menu set to "Any" with a "Change..." link.
- State:** Two radio buttons: Enabled and Disabled.
- Log:** Two radio buttons: Log and Do not log.

At the bottom right are "OK" and "Cancel" buttons. A blue information icon is located above a text box containing the following message:

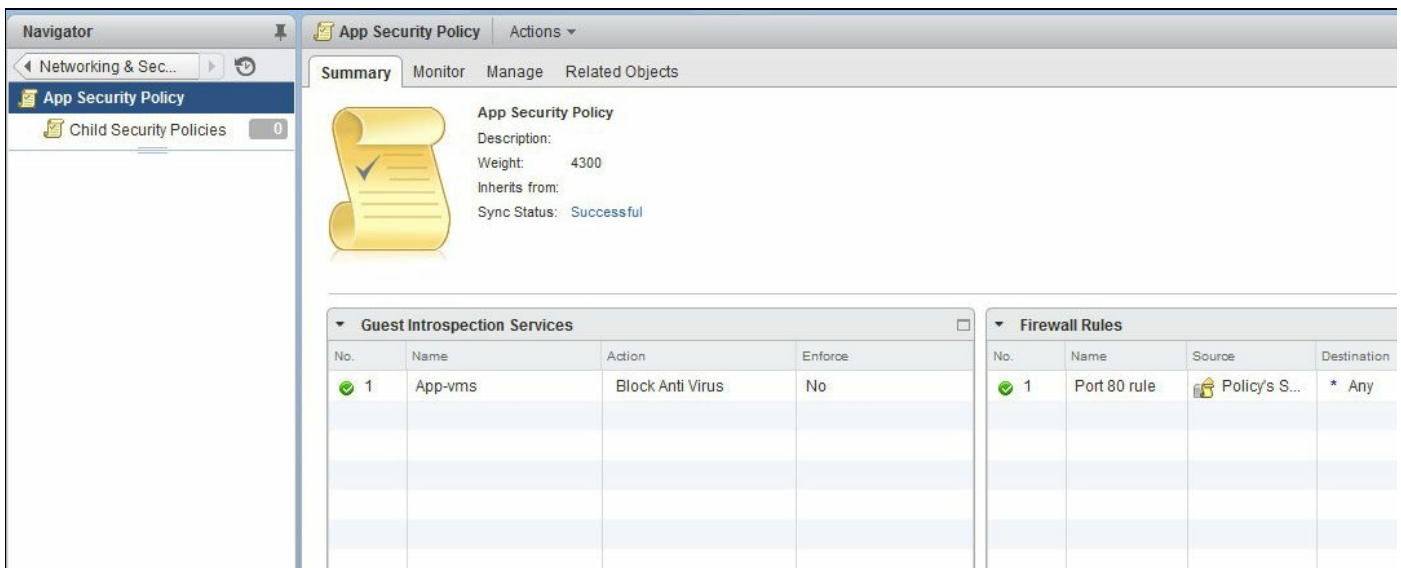
Either source or destination selection (or both) must be "Policy's Security Groups". Current selection will apply to "Outgoing" traffic from the security groups where this policy gets applied to specified Destination.

20. Enter a Name and the Description for the policy. Select Redirect to service to redirect to the third-party service.
21. Select the Service Name and the Profile of the network introspection service provider/vendor.
22. Click the appropriate Change... to select the Source and Destination security groups for the policy.
23. Click the appropriate Change... to select a specific Service that you want to add.
24. Select the State and the Log option and click Ok.
25. Click Next when done.
26. Review the summary and click Finish.

27. The policy is now created and added:





You can click on the policy to review it:



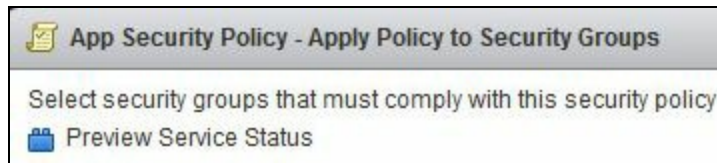
Security group and security policy mapping

Now that we have a security group and a security policy configured, we need to map the security policy to a security group. All the objects that are part of the security group get its security policy rules applied to them. Objects can be dynamically added to the security group based on the criteria they meet:


1. Go to the Home | Networking & Security | Service Composer | Security Policies tab.
2. Select a security policy and click on the  icon to associate the policy to a security group. Alternatively, you may also right-click on the policy and select Apply Policy:



3. Select the security group you want this policy associated with. You can also click on Preview Service Status to identify any services that will not be applied to the group:

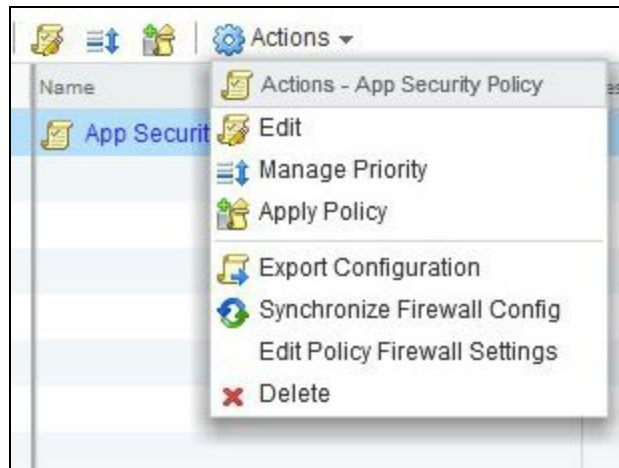


4. Click on the Preview Service Status.
5. For example, the following screenshot shows that VMware Data Security is not installed on the host:



6. Click OK when done. We now have a security policy associated with a security group.

There are actions that can be performed on a security policy:



Manage Priority lets you manage the order of the policy by placing it over or under any other rules. Export Configuration lets you export the policy configuration. Apply Policy maps the policy to a security group.

Synchronize Firewall Config allows you to synchronize the firewall configuration. Firewall rules in the security policy show up in the firewall configuration section. VMware does not recommend you edit the composer rules in the firewall section. Always edit rules in the security policy and click Synchronize Firewall Config to apply changes appropriately.

Go to the Home | Networking & Security | Firewall section. You will see your composer rules with the associated group listed:

Navigator

- Home
- Networking & Security
 - NSX Home
 - Dashboard
 - Installation
 - Logical Switches
 - NSX Edges
 - Firewall**
 - SpoofGuard
 - Service Definitions
 - Service Composer
- Tools
 - Flow Monitoring
 - Activity Monitoring
 - Endpoint Monitoring
 - Traceflow
- Networking & Security Inventory
 - NSX Managers

Firewall

Configuration Saved Configurations Settings

NSX Manager: 192.168.1.8 (Role: Primary)

Last publish operation succeeded 6/13/2017 1:56:28 PM

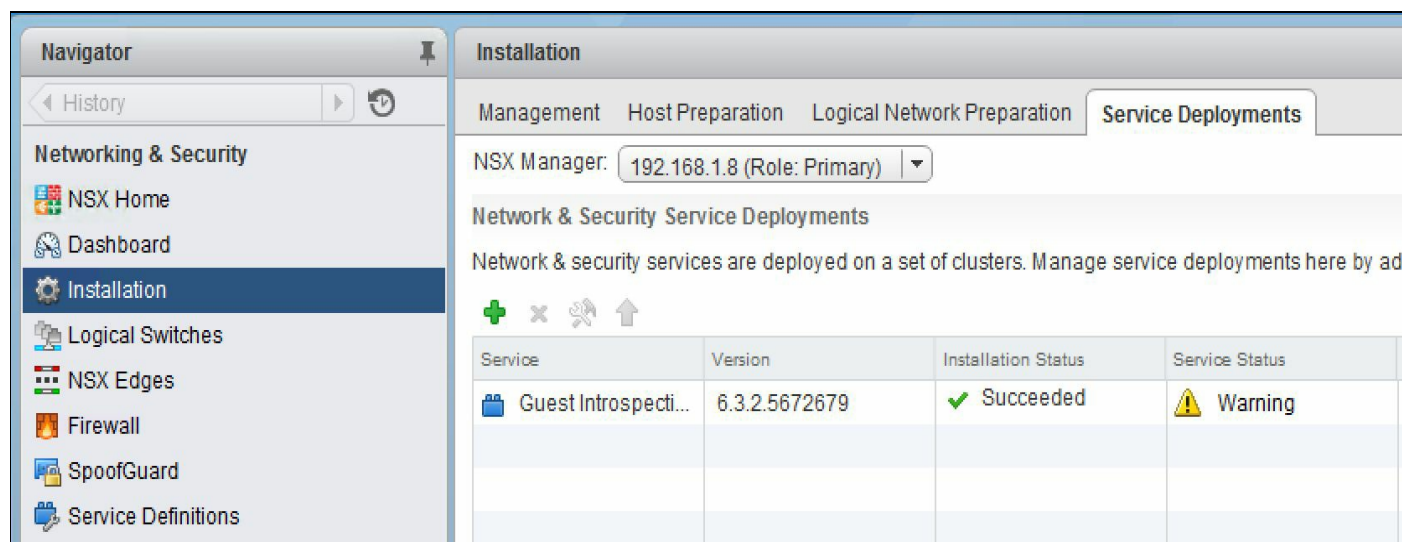
General Ethernet Partner security services

No.	Name	Rule ID	Source	Destination	Service	Action	Applied To
App-rules (Rule 1)							
1	app-port80	1006	New Clust...	* any	CIM-HTTP	Allow	Distributed Fire...
Default Section Layer3 (Rule 2 - 4)							
2	Default Rule NDP	1003	* any	* any	IPv6-ICMP ... IPv6-ICMP ...	Allow	Distributed Fire...
3	Default Rule DHCP	1002	* any	* any	DHCP-Cli... DHCP-Ser...	Allow	Distributed Fire...
4	Default Rule	1001	* any	* any	* any	Allow	Distributed Fire...

Network extensibility

We will briefly look at network extensibility, a vast and a fairly complex topic that covers multiple vendors and deployment architectures. NSX has the ability to integrate with multiple vendors to provide a rich network feature set and also a single pane of glass to manage multiple physical network devices.

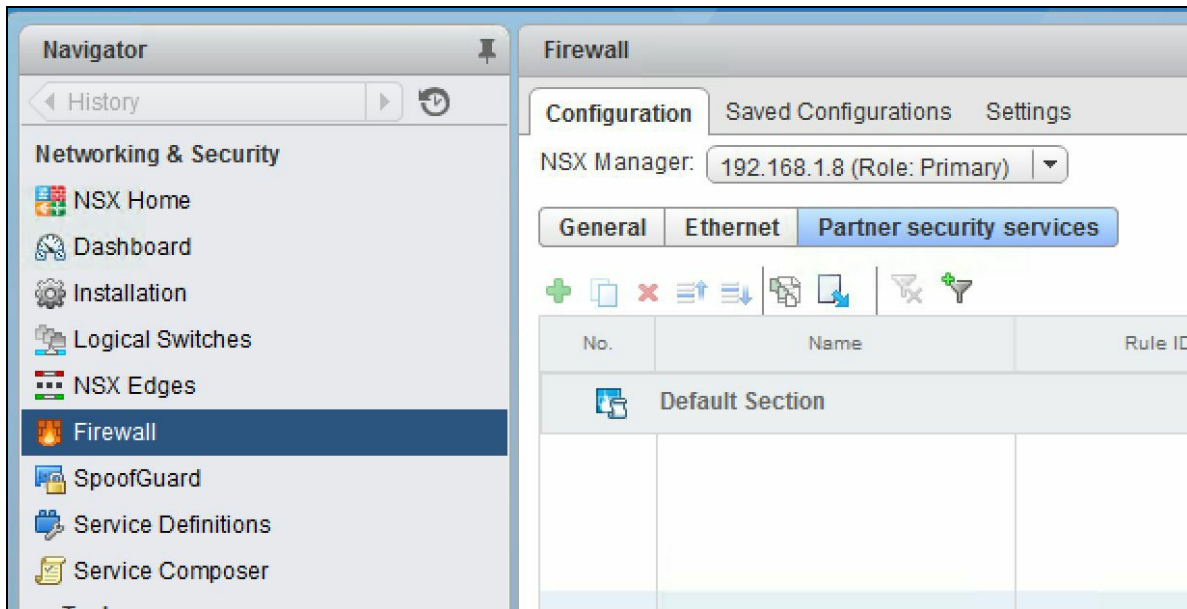
A third-party service provider (vendor) must register with the NSX manager using the specified login credentials. Once registered, we will deploy the partner service that allows NSX to offer the third-party service:



Service	Version	Installation Status	Service Status
Guest Introspecti...	6.3.2.5672679	✓ Succeeded	⚠ Warning

We have seen saw earlier, while creating security policies, the Redirect to service option, which allows traffic to be redirected to these third-party services. You may also redirect traffic by adding firewall rules.

Click the Partner security services tab in the firewall section:



During creation of the new rule, select Action as Redirect to allow for traffic redirection to the third-party vendor.



To learn more about Partner security services, visit the following URLs:

<http://pubs.vmware.com/nsx-63/topic/com.vmware.nsx.admin.doc/GUID-EA477D96-E2D3-488B-90AA-2F19B4AE327D.html>

<http://pubs.vmware.com/nsx-63/topic/com.vmware.nsx.admin.doc/GUID-C3FD9AA7-96AB-48DA-9C60-0F764EACE783.html>

Summary

One of the most important factors for a successful enterprise is compliance and security. We started this chapter by looking at Service Composer. We created service groups with dynamic groupings of virtual machines and vCenter objects. We also created a service policy and associated it with a service group. NSX can integrate with partner services and we looked at a very generic integration process. We briefly discussed security policy and firewall rule redirection as well.

In [Chapter 7, Monitoring](#), we will look at monitoring our environment using the tools offered by NSX. We will look at activity monitoring, flow monitoring, and trace flow, among others.

Monitoring

In this chapter, we will be looking at monitoring our environment and the different tools NSX has to offer. We will begin the chapter looking at Endpoint Monitoring, which allows visibility into applications to verify if security policies are being enforced appropriately. We will also look at how to enable data collection on one or more virtual machines. We will then discuss flow monitoring and the traceflow feature.

In this chapter, we will cover:

- Endpoint Monitoring
- Flow monitoring
- Traceflow

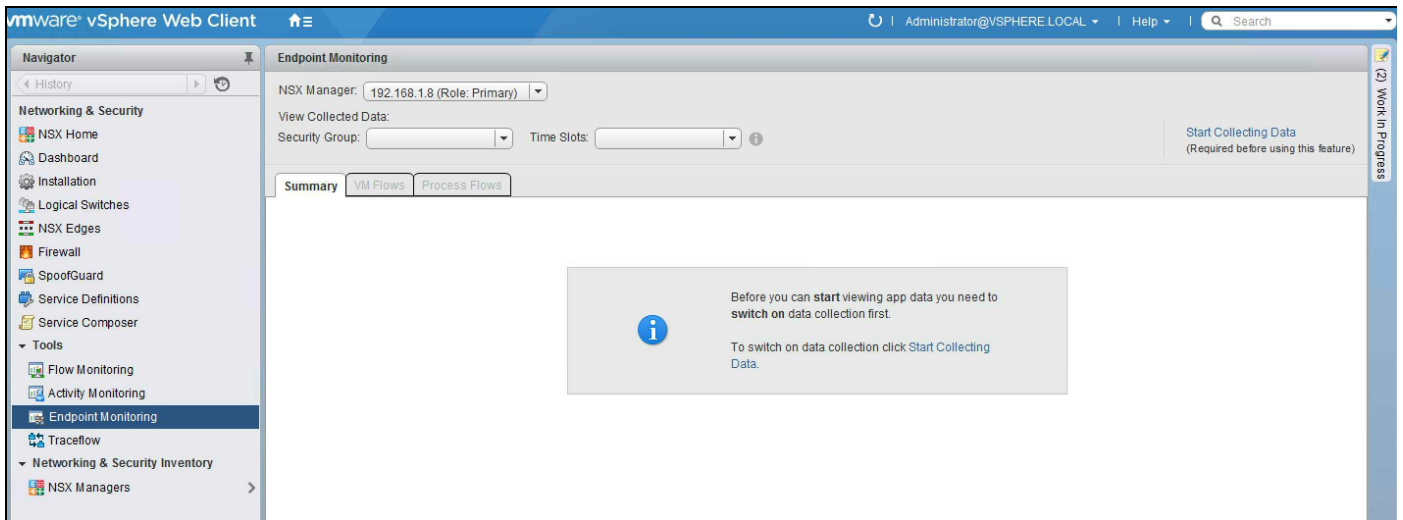
Endpoint Monitoring



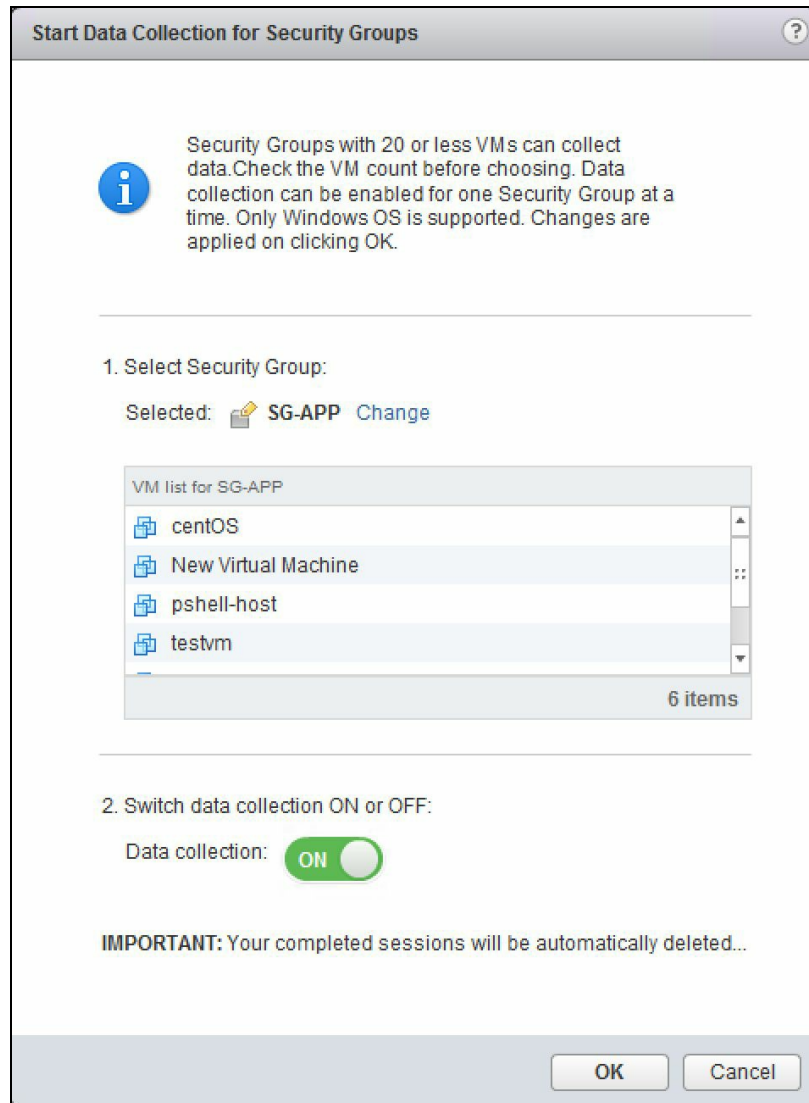
Remember that, as of NSX 6.3.0, the NSX Activity Monitoring feature has been deprecated. You can continue to use this feature at your discretion, but be aware that it will be removed from NSX in a future release. VMware recommends you use Endpoint Monitoring instead of Activity Monitoring. (Courtesy VMware)

Endpoint Monitor provides insight and visibility into applications running within an operating system to ensure that security policies are being enforced correctly. For example, Activity Monitor can help if you have a misconfigured web server that is receiving traffic on HTTP instead of HTTPS as you intended. You can now run reports to monitor in-bound and out-bound traffic to the machines managed by the vCenter.

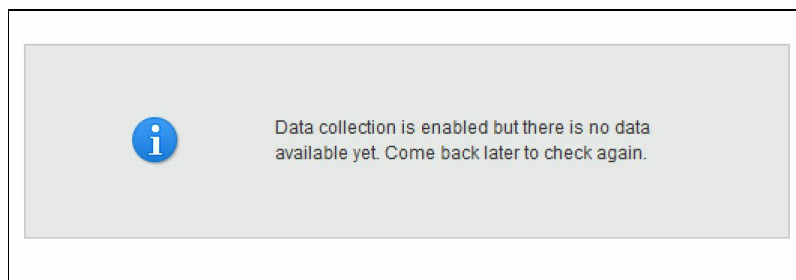
Endpoint Monitoring require guest introspection to be installed. On virtual machines, you will need to the install guest introspection driver, which is part of the VMware tools installation. A full installation typically installs the driver onto the operating system. The purpose of the guest introspection driver (VMCI driver) is to detect all the applications running on the operating system and to send this information to the guest introspection appliance:



1. Click on Start Collecting Data.
2. Click on Select Security Group and flip the Switch data collection ON or OFF option to ON.
3. Click OK when done:




Data collection is now enabled and will take some time for NSX to gather all virtual machine activity:



Once the data collection is complete, the summary screen displays the details of NSX Manager, the security group, and the time slot of the collected data. The first box shows the total number of running virtual machines and the total number of processes that are generating traffic. When you click on the virtual machine, the action is redirected to the VM Flows tab. Clicking the number of processes generating traffic takes you to the

Process Flows tab.

Once data is gathered, remember to select Stop Collecting Data by clicking on the right side link. This will release the resources back to NSX Manager and avoid unnecessary data collection:



Flow monitoring

NSX Flow monitoring is a feature that allows detailed traffic monitoring to and from protected virtual machines. Flow monitoring can uniquely identify different machines and different services that are exchanging data and when enabled can identify which machines are exchanging data over specific applications. Flow monitoring also allows live monitoring of TCP and UDP connections and can be used as an effective forensic tool.

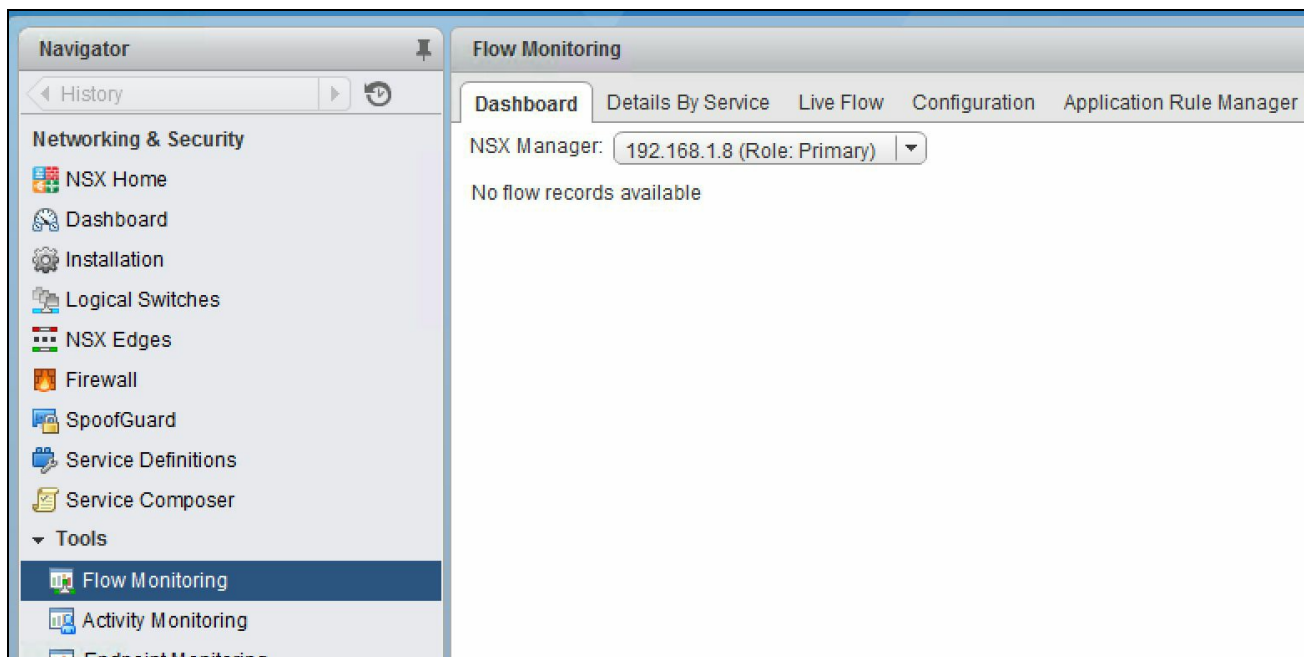


Flow monitoring can only be turned on for NSX deployments where a firewall is enabled.

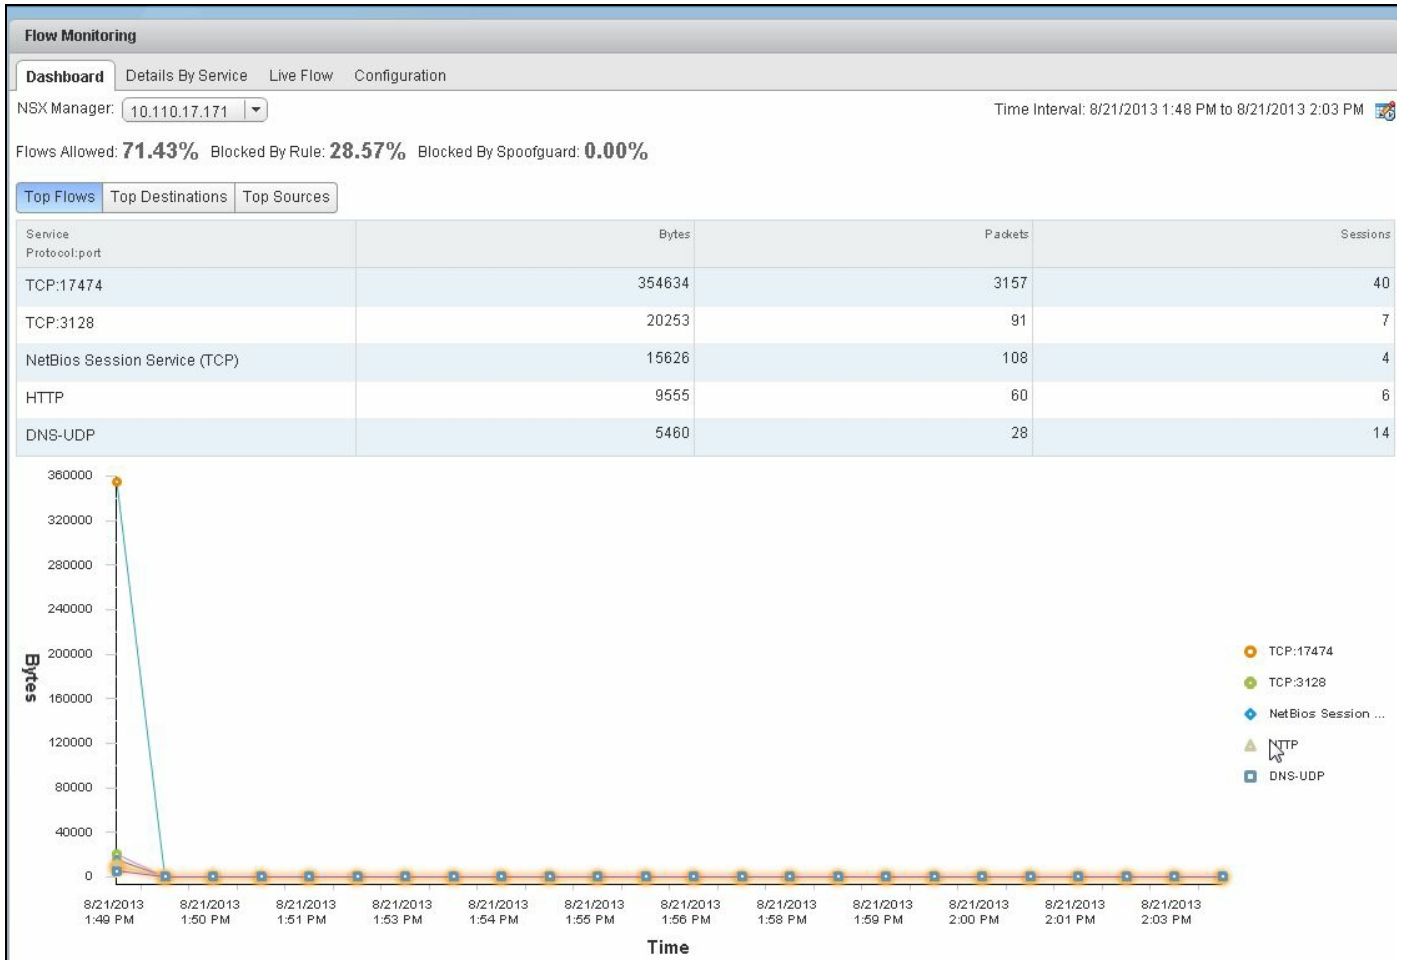
Flow monitoring data can be polled to a set interval and then analyzed. The default period is 24 hours and the minimum is one hour while the maximum data collection interval is two weeks. Keep an eye on the disk space being consumed by NSX Manager as the polling interval is set.

To view the flow monitoring data, follow these steps:


1. Log in to your vCenter web client and navigate to Networking & Security | Flow Monitoring:



2. Click Flow Monitoring under the Dashboard tab:



Courtesy - VMware

3. To change the time interval for the flow, click on the  icon on the right:

Change Time Interval

Last 15 minutes
 Last 1 hour
 Last 12 hours
 Last 24 hours
 Last 1 week
 Last 2 weeks
 From: 10/28/2015 11:53 To: 10/28/2015 12:08

OK Cancel

4. You will see the Flows Allowed, Blocked By Rule, and Blocked By Spoofguard metrics:

NSX Manager:	192.168.110.15
Flows Allowed:	100.00%
Blocked By Rule:	0.00%
Blocked By Spoofguard:	0.00%

5. The Top Flows tab shows the total incoming and outgoing traffic over the specified period of time:

Service Protocol:port	Bytes	Packets
DNS-UDP	12.76 KB	131
DNS-UDP	9.35 KB	123
NetBios Session Service (TCP)	6.16 KB	48
ICMP:echo-reply	5.97 KB	102
IPv6-ICMP	2.20 KB	28

6. The Top Destinations tab shows the incoming traffic per destination while the Top Sources tab shows the specified outgoing traffic per source:

Destination	Incoming Traffic (Bytes)	Packets
192.168.110.10	15.51 KB	171
fd53::11	12.97 KB	134
8.8.8.8	5.74 KB	98
win8-01a (fe80::d1fc:198e:f37e:aa37)	904	11
00:50:56:01:20:a5	598	13

7. The Top Sources tab, as in the following screenshot, shows the outgoing traffic per source:

Source	Outgoing Traffic (Bytes)	Packets
win8-01a (192.168.100.222)	21.64 KB	275
win8-01a (fe80::d1fc:198e:f37e:aa37)	13.74 KB	145
fe80::1	744	7
00:50:56:01:20:a5	598	13
00:50:56:ae:09:87	598	13

8. The Details By Service tab shows the allowed and blocked flows including the number of sessions for each type of flow. You can click on a service to view the traffic flow and the rules that apply. You can also choose to edit a rule by clicking Edit Rule in the Actions column or you can add a rule by clicking Add Rule in the Actions column:


Type	Service	Bytes	Session
UDP	DNS-UDP	12.76 KB	17
UDP	DNS-UDP	9.35 KB	16
TCP	NetBios Session Service (TCP)	6.16 KB	2
OTHER	ICMP:echo-reply	5.97 KB	2
OTHER	IPv6-ICMP:0	2.20 KB	28

One of the most interesting features is the ability to see a live flow for a selected vNIC. You can monitor all live TCP and UDP connections to a vNIC using the live flow feature. To do so, perform the following set of steps:

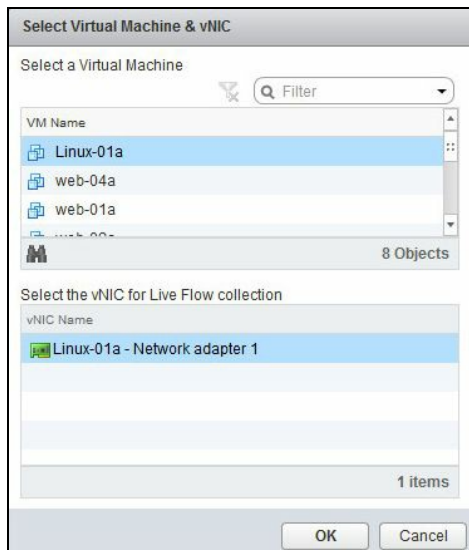


You can monitor a maximum of two vNICs per host and a maximum of five vNICs per vCenter.

1. Log in to the vCenter web client and go to Networking & Security | Flow Monitoring.
2. Click on the Live Flow tab in the dashboard:



3. Click Browse to select a vNIC and click Start when done:



The refresh rate can be set accordingly.

4. Click Stop when done:

NSX Manager: 192.168.110.15

Live Flow will be shown for the selected vNIC. Please select a vNIC and press start to s

vNIC: win8-01a - Network adapter 1 Browse Start Stop

Refresh Rate: 5 Seconds

RuleId	Direction	Flow Type	Protocol	Source IP	Source Port	Des
1001	OUT	Active	UDP	fe80::d1fc:198e:f3	49440	fd53::
1001	OUT	Active	UDP	192.168.100.222	58835	192.1
1001	OUT	Active	UDP	fe80::d1fc:198e:f3	58933	fd53::
1001	OUT	Inactive	UDP	fe80::d1fc:198e:f3	61723	fd53::
1001	OUT	Active	UDP	fe80::d1fc:198e:f3	52444	fd53::
1001	OUT	Active	UDP	192.168.100.222	58933	192.1
1001	OUT	Active	UDP	fe80::d1fc:198e:f3	61870	fd53::
1001	OUT	Inactive	UDP	192.168.100.222	61723	192.1
1001	OUT	Active	UDP	192.168.100.222	52444	192.1



Using Live flow increases NSX Manager's resource consumption so this feature should be used sparingly.

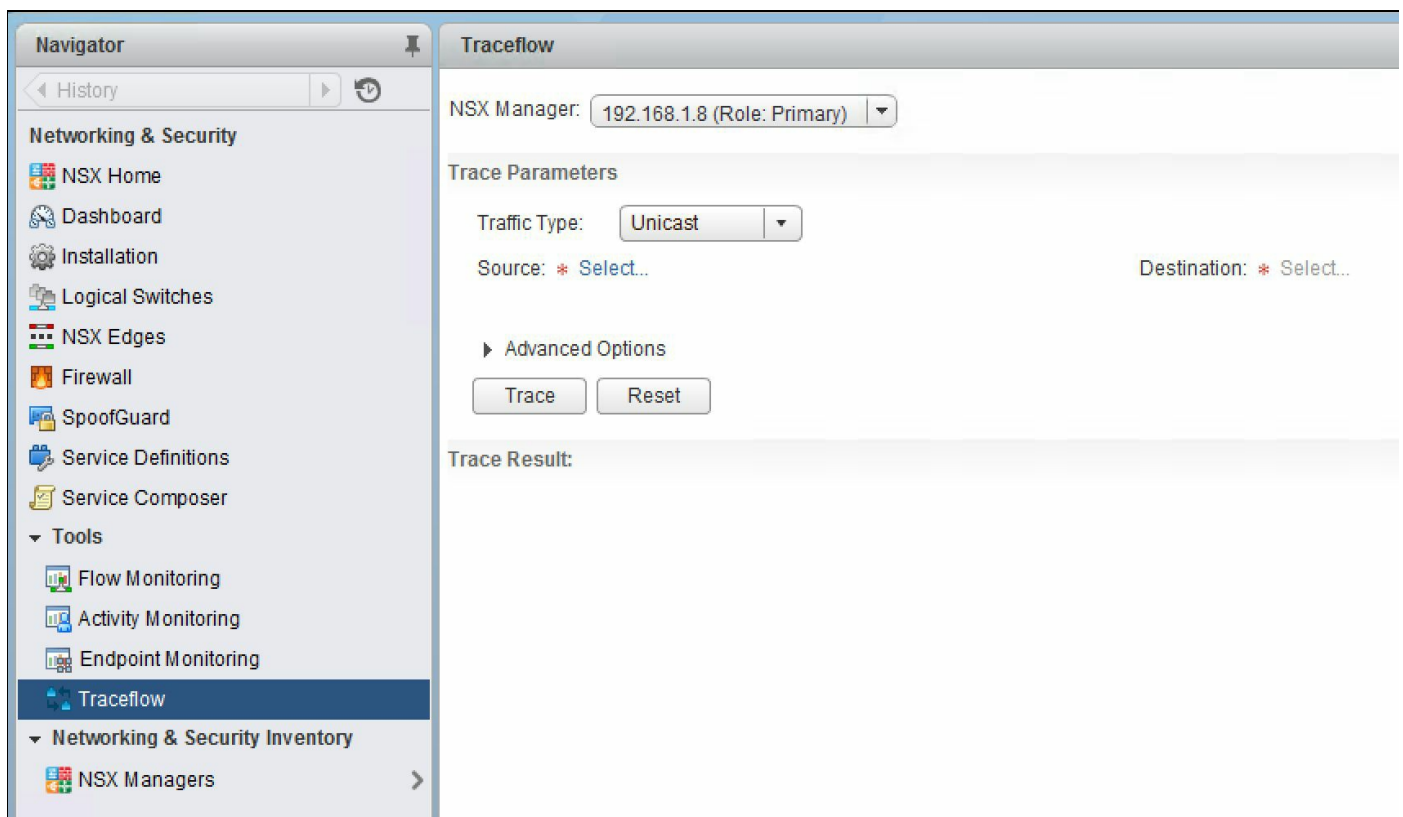
Traceflow

Traceflow is an interesting tool and was built to allow administrators to seamlessly troubleshoot their virtual network environment by tracing a packet flow, in a similar way to the legacy Packet Tracer application. Traceflow allows you to inject a packet into the network and monitor its flow across the network. This flow allows you to monitor your network and identify issues such as bottlenecks or disruptions.

Traceflow allows you to construct your own packets with custom headers and packet sizes. The target destination of this test packet can be a NSX-managed overlay network or underlay network devices such as a host or a logical router. The source will always be a vNIC from a VM. These packets are injected in the virtual distributed switch and support the layer 2 unicast, multicast and broadcast and layer 3 unicast traffic types.

To use traceflow, follow these steps:

1. Log in to your vCenter web client and go to Networking & Security | Traceflow:




2. Select the appropriate Traffic Type. You can choose from Unicast, L2 Multicast, and L2 Broadcast.

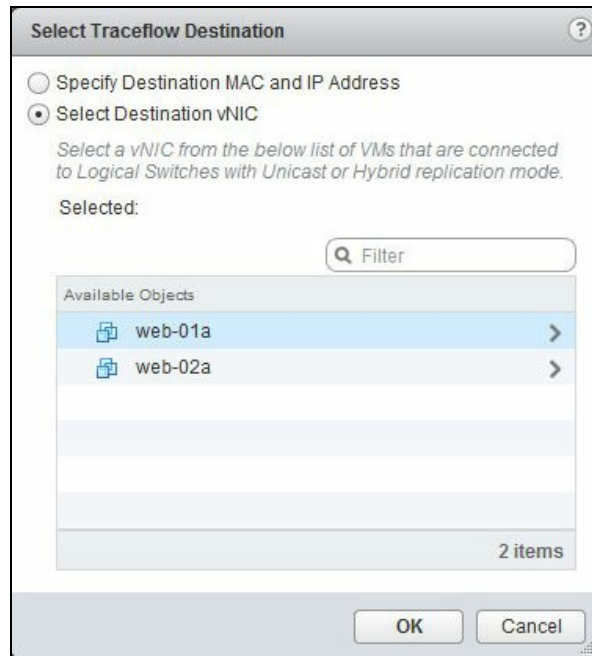
3. Select the source VM vNIC where the packet will be sent from:



4. You can click on the right arrow to select the vNIC:



5. Select the Traceflow Destination:




6. If L2 Broadcast traceflow is selected, enter the Subnet Prefix Length:



7. For an L2 Multicast traceflow, the multicast group addresses are entered:



8. You may leave the Advanced Options at the default unless your network tests require a certain scenario to be created:



9. Click Trace when done:

Trace Result: Traceflow delivered observation(s) reported

1 Delivered

Sequence	Observation Type	Host	Component Type	Component Name
0	Injected	esx-03a.corp.local	vNIC	vNIC
1	Received	esx-03a.corp.local	Firewall	Firewall
2	Forwarded	esx-03a.corp.local	Firewall	Firewall
3	Forwarded	esx-03a.corp.local	Physical	esx-03a.corp.local
3	Forwarded	esx-03a.corp.local	Physical	esx-03a.corp.local
4	Received	esx-05a.corp.local	Physical	esx-05a.corp.local
4	Received	esx-05a.corp.local	Physical	esx-05a.corp.local
4	Received	esx-05a.corp.local	Physical	esx-05a.corp.local

As you can see, the packet was successfully delivered and the entire trace is visible:

4	Received	esx-01a.corp.local	Physical	esx-01a.corp.local
4	Received	esx-01a.corp.local	Physical	esx-01a.corp.local
5	Received	esx-01a.corp.local	Firewall	Firewall
6	Forwarded	esx-01a.corp.local	Firewall	Firewall
7	Delivered	esx-01a.corp.local	vNIC	vNIC

Summary

We started this chapter with a look at three interesting tools that NSX offers to allow us to efficiently and thoroughly monitor and analyze our environment. Activity Monitoring is a tool that has been replaced by Endpoint Monitor, which allows us visibility into applications. You will also need to enable guest introspection in VMware tools as well. We then looked at flow monitoring. We ended the chapter by looking at the traceflow tool to identify network issues and allow us end-to-end packet trace visibility.

In [Chapter 8](#), Managing NSX, we will look at the operational management of our NSX environment. We will look at backing up NSX Manager, setting up syslog servers, controller cluster actions, and many more.

Managing NSX

In this chapter, we will be looking at some of the most common operations to manage our NSX environment. We will begin the chapter by looking at NSX Manager actions, including backing up and restoring your NSX environment, management settings including, setting up syslog servers, and setting up NSX with a Windows domain. We will also look at controller cluster actions, including changing its passwords, downloading tech support logs, and also setting up a syslog server. We will also look at checking the communication channel health for communications between NSX Manager and its associated components.

In this chapter, we will cover:

- NSX Manager settings
- Backup and restore
- NSX Manager domain registration
- Controller cluster actions
- Summary

NSX Manager settings

There are several settings in the NSX Manager that can be edited to suit your environment. These settings can also be edited using the NSX CLI; however, we will learn to update these settings using the UI.

Date and time settings

You can change the date and time settings as needed.



An NSX Manager reboot is needed after any date/time change is made.

You need to perform the following steps:

1. Log in to the NSX Manager virtual machine appliance.
2. Click on Appliance Management | Manage appliance settings:

Time Settings	
Specify NTP server below. For SSO configuration to work correctly use the same NTP server used by the SSO server.	
NTP Server	192.168.110.1
Timezone	UTC
Date/Time	11/12/2015 03:29:39

3. Next to Time Settings, click Edit:

Time Settings

Specify NTP server below. For SSO configuration to work correctly it is required that the time on this virtual appliance and NTP server should be in sync. It is recommended to use the same NTP server used by the SSO server.

NTP Server: 192.168.110.1

Timezone: UTC

Date/Time: 11/12/2015 03:29:39

OK Cancel

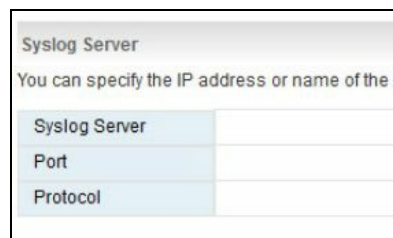
4. Click OK when done.
5. Reboot the appliance to apply the changes.

Syslog servers

Having a syslog server is recommended. Configuring NSX Manager with a remote syslog server enables you to collect, view, and save all log files to a central location. This enables you to store logs for compliance purposes; when you are using a tool such as VMware vRealize Log insight, it enables you to create alarms and use the built-in search engine to review logs.

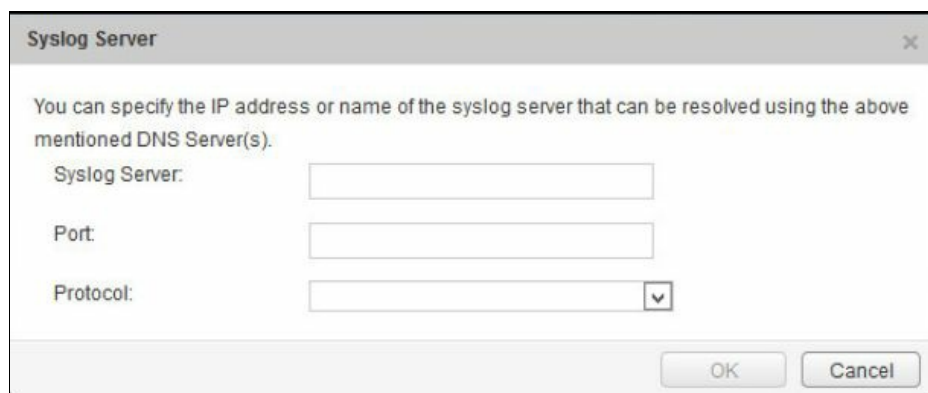
The steps to configure NSX Manager to start forwarding logs to a syslog server are as follows:

1. Log in to the NSX Manager virtual appliance.
2. Click on Manage Appliance Settings | General.
3. Locate the Syslog Server and click Edit next to it:



The screenshot shows a dialog box titled "Syslog Server". Below the title, there is a subtitle: "You can specify the IP address or name of the s". The dialog contains three input fields: "Syslog Server", "Port", and "Protocol".

4. Enter Syslog Server values as applicable and click OK:



The screenshot shows the "Syslog Server" dialog box with the following content:

You can specify the IP address or name of the syslog server that can be resolved using the above mentioned DNS Server(s).

Syslog Server:

Port:

Protocol:

OK Cancel

DNS servers

It is important to have DNS configured in your environment and that NSX Manager is configured to use DNS servers. The steps to configure the NSX Manager to use DNS servers are as follows:

1. Log in to the NSX Manager virtual appliance.
2. Click on Manage Appliance Settings | Network:

DNS Servers	
To resolve all objects referenced using a hostname, you must provide one or more DNS servers common to vCenter, ESX hosts and other vSphere components (If primary or secondary server is removed, the next available dns server in the line would assume the responsibility).	
IPv4 DNS Servers	
Primary Server	192.168.110.10
Secondary Server	
IPv6 DNS Servers	
Primary Server	
Secondary Server	
Search Domains	corp.local


3. Locate the DNS Servers and click Edit next to it:


The screenshot shows a dialog box titled "DNS Servers" with a close button (X) in the top right corner. The text inside reads: "To resolve all objects referenced using a hostname, you must provide one or more DNS servers common to vCenter, ESX hosts and other vSphere components (If primary or secondary server is removed, the next available dns server in the line would assume the responsibility)." Below this text are three sections: "IPv4 DNS Servers" with "Primary Server:" (192.168.110.10) and "Secondary Server:" (empty); "IPv6 DNS Servers" with "Primary Server:" (empty) and "Secondary Server:" (empty); and "Domain Search List" with "Search Domains:" (corp.local). At the bottom right are "OK" and "Cancel" buttons.

4. Enter DNS Server values and click OK when done.

Technical support logs

During NSX troubleshooting, you will be required to download technical support logs in order to review them and submit them to VMware technical support. Perform the following steps:

1. Log in to the NSX Manager virtual appliance.
2. Click on Manage Appliance Settings.
3. In the right corner of the screen, click the  icon:




4. Click Download Tech Support Log. NSX Manager now prepares the log bundle; when ready click Save to download the log bundle to your desktop. The log bundle is compressed and downloaded as a `.gz` file.

SSL certificates

NSX Manager generates a self-signed certificate during the initial install; however, you can also configure it to use a CA signed certificate to authenticate its identity and secure communication. VMware recommends that you generate the certificate using NSX Manager's Generate Certificate option. The process of registering a certificate includes generating a certificate, signing it with a CA, and then importing the signed certificate.

To generate a **certificate signing request (CSR)**:

1. Log in to the NSX Manager virtual appliance.
2. Click on Manage Appliance Settings | SSL Certificates:



3. Click Generate CSR to generate the SSL certificate:

4. Complete the form appropriately.
5. Click OK when done.

Now that the CSR is generated, download the certificate by clicking Download CSR.

You will now send this to your CA to get it signed. Use the Import option to upload the signed certificates that you receive from the CA and save your signed certificates.

Backup and restore

Backups are critical for an NSX environment that allows you to restore them appropriately during a system failure. Apart from vCenter, you can also perform backup operations on the NSX Manager, controller clusters, NSX Edge, firewall rules, and Service Composer. All these can be backed up and restored individually.

NSX Manager backup

NSX Manager can be backed up on-demand or on a scheduled basis. NSX Manager backups are saved over any FTP or SFTP that the NSX Manager can access.



NSX Manager restore is supported for the same NSX Manager version as the backup.

Perform the following set of steps to back up and restore data:

1. Log in to NSX Manager via web access.
2. Click on Backup & Restore:

Backups & Restore	
FTP Server Settings:	
Scheduling:	OFF
Exclude:	
Backup History	
File Name	

3. Next to FTP Server Settings, click Change to specify the backup location:

Backup Location	
IP/Host name:	<input type="text"/>
Transfer Protocol:	<input type="text"/>
Port:	<input type="text"/> Invalid Port number
User name:	<input type="text"/>
Password:	<input type="text"/>
Backup Directory:	<input type="text"/>
Filename Prefix:	<input type="text"/>
Pass Phrase:	<input type="text"/>
OK Cancel	

4. Fill in the appropriate fields and click OK.
5. Clicking Backup starts an on-demand backup:

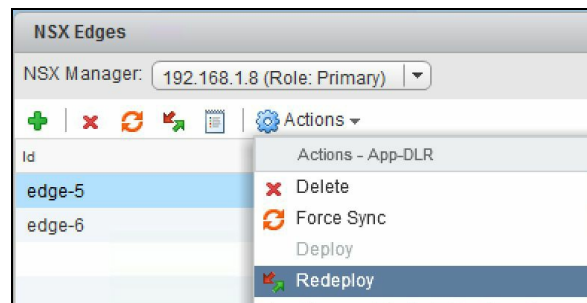



6. To schedule a backup, click Change near the Schedule option to set a schedule:

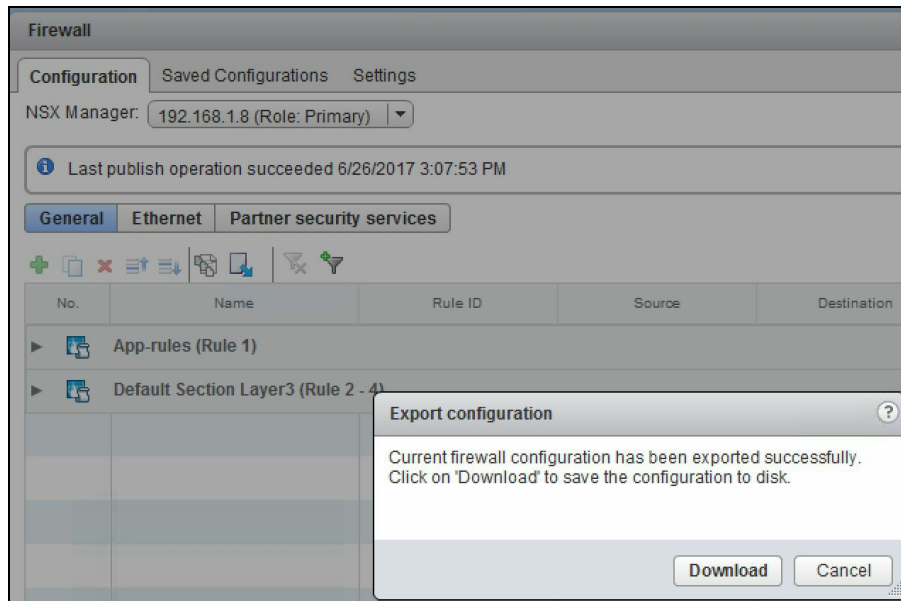
A dialog box titled 'Create or Schedule Backup'. It contains four dropdown menus: 'Backup Frequency', 'Day of week', 'Hour of day', and 'Minute'. A 'Schedule' button is located at the bottom right of the dialog.

In an event of a failure that requires a restore operation, redeploy a fresh NSX Manager of the same backup version, configure the FTP server settings, and use the Restore option to restore a backup.

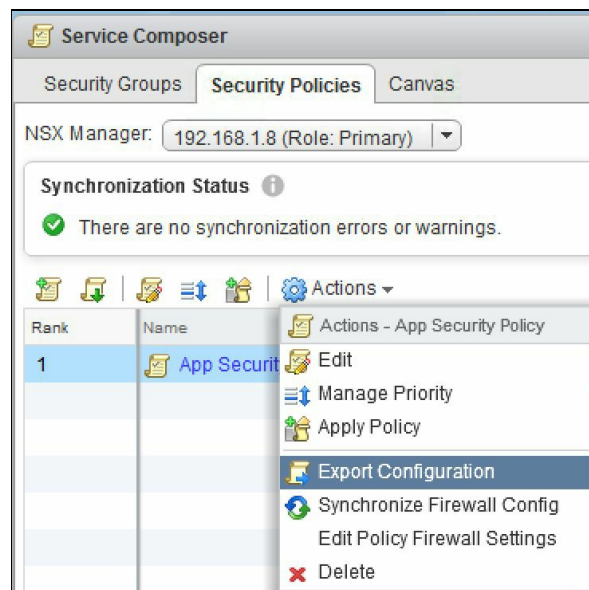
As part of NSX Manager backup, the NSX Edge appliance configuration is also backed up. Edge appliances include logical routers and also any Edge services gateway appliances. Inaccessible or failed NSX Edge appliances can be redeployed by using the Redeploy action:



The firewall configuration can be exported using the  Export icon in the firewall section. This exports all the firewall configuration and rules in XML format to be downloaded and saved at a safe location. This also contains any NSX Service Composer rules:



Any Service Composer configuration can be backed up by exporting it. This is done using the Export Configuration option, found under Actions in a security policy that you pick:



The configuration is downloaded to a selected location on your desktop and can then be moved to a safe location. The downloaded configuration can also be imported into another NSX Manager.

To import a security policy, click the  icon:

Import Configuration

1 Select configuration file

2 Ready to complete

Select configuration file
Select the Service Composer configuration to be imported

Select Configuration: Browse...

Suffix:

All the imported object names will be suffixed with this string

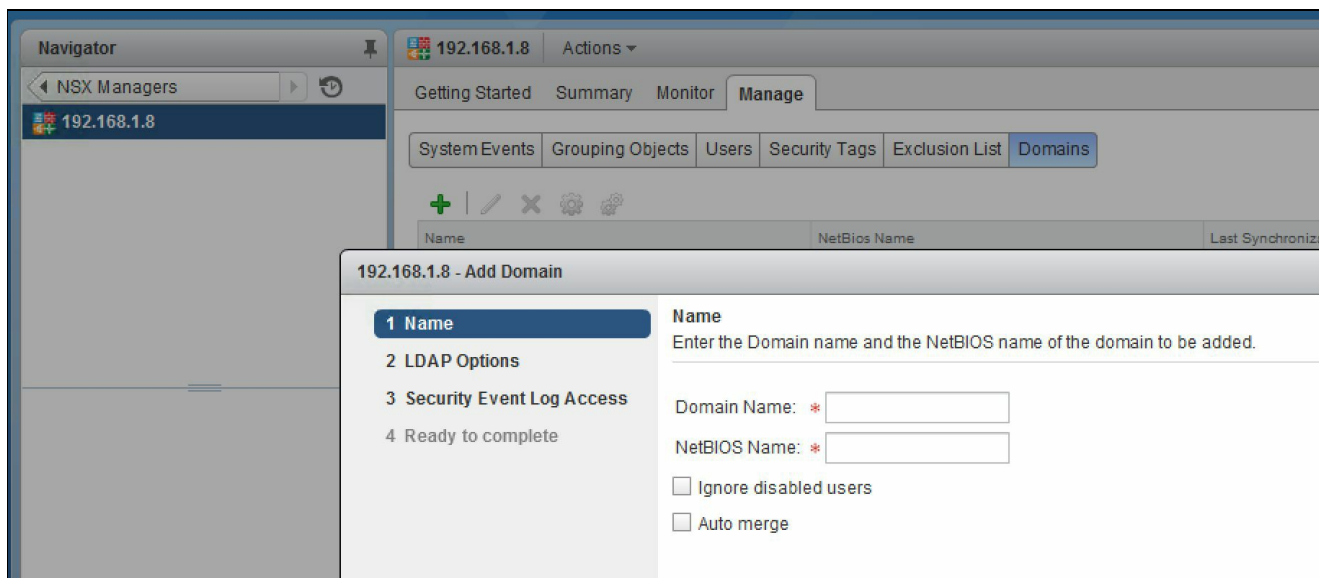
Back Next Finish Cancel

NSX Manager domain registration

It is recommended to register your NSX deployment with your corporate domain. This allows NSX to map user and group information and also the associations that allow you to create security bindings based on these relationships.

Register a domain to your NSX Manager with the following procedure:

1. Log in to the vSphere webclient and navigate to Home | Networking & Security | NSX Managers. Select your NSX Manager, and go to the Manage | Domain tab:



2. Click the + sign to add a new domain.
3. Enter a fully qualified Domain Name.
4. Specify the domain controller that the domain needs to be synchronized with. Select the protocol.
5. Enter the user credentials for the account and click Next.
6. Select Use Domain Credentials if you want to authenticate using the LDAP server.
7. You can additionally select Yes or No for the security event log segment. This step is used by Active Directory Event Log Scraper, which allows you to configure the identity firewall, where firewall rules are tied to your Active Director users. This logging allows NSX to detect when a user has logged in to a specific machine so those security policies can be applied. In the Security Event Log Access page,

select either CIFS or WMI for Connection Method to access security event logs on the specified AD server and change the port number if required:

192.168.1.8 - Add Domain

- ✓ 1 Name
- ✓ 2 LDAP Options
- 3 Security Event Log Access**
- 4 Ready to complete

Security Event Log Access

Specify the options that affect access to the Security Event logs on specified server in the domain. If required, specify the user name and password of an alternate domain account for log access.

Do you have the Security Event Log available for Server 192.168.1.2?

Yes No

Server: * 192.168.1.2

Connection Method: CIFS

Port: * 445

Use Domain Credentials

User Name: *

Password: *

Back Next Finish Cancel

8. Click Next and click Finish when done.

By default, all domains are synchronized automatically every three hours. Under NSX Managers | Name columns | Manage tab, click on the appropriate synchronization technique. While delta synchronization synchronizes only changed AD objects, a full synchronization does a complete refresh of all AD objects.

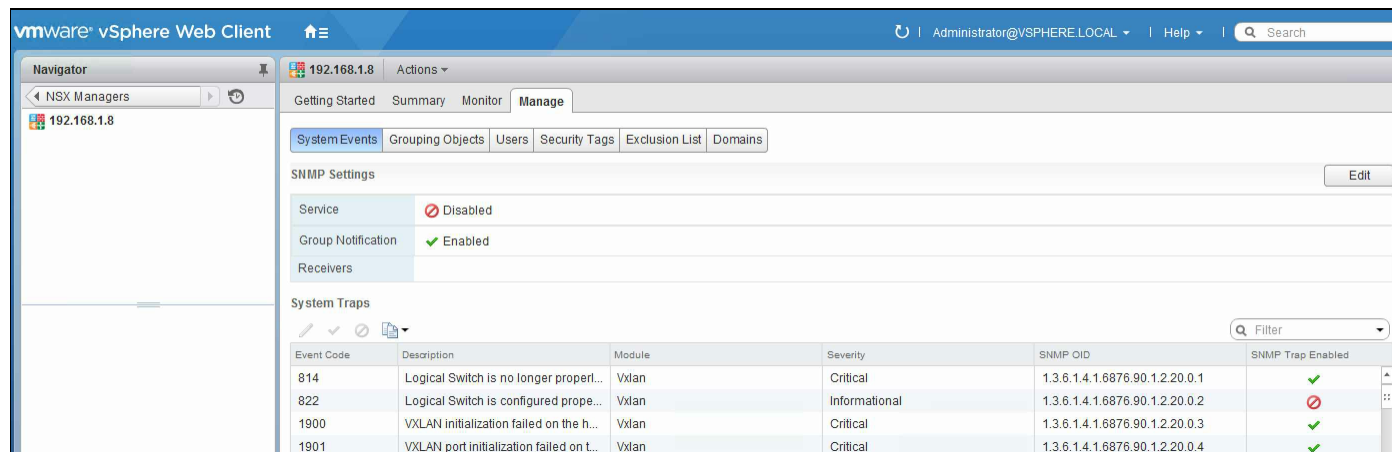
Configuring SNMP traps

Simple network management protocol (SNMP) traps are alert messages sent from a remote SNMP-enabled device to a collector. You can configure the SNMP agent to forward SNMP traps.

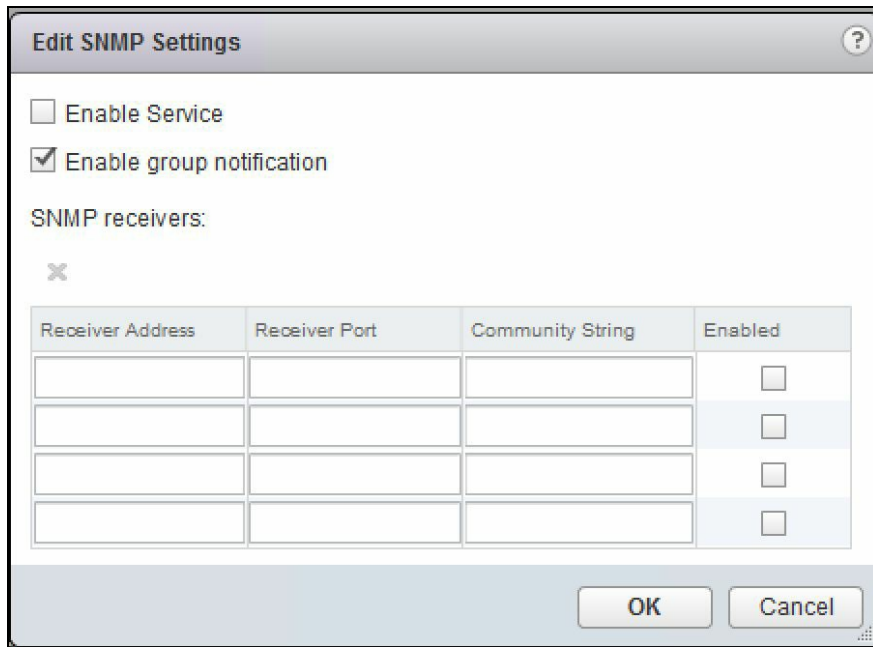
By default, the SNMP trap mechanism is disabled. When the SNMP trap is enabled, only critical and high severity notifications are sent to the SNMP manager.

To enable SNMP traps:

1. Log in to the vSphere web client, navigate to Home | Networking & Security | Network & Security Inventory | NSX Manager.
2. Select NSX Manager IP Address.
3. Click on Manage | System Events:



4. Click Edit and configure the requested information:




5. Checking Enable group notification allows system events to be sent out as part of groups, which reduces the load on the SNMP server.
6. You can configure up to four receivers for traps to be sent out to.
7. Click OK when done.

Controller cluster operations

There are certain operations that can be performed on a controller cluster. You should always change the default password for controller clusters to ensure data security.

To change the passwords on a controller cluster:

1. Log in to the vSphere web client and navigate to Home | Networking & Security | Installation.
2. The Management tab shows the list of controllers. Pick one whose password needs to be changed. It is recommended that all three controllers have different passwords.
3. Click Actions | Change Controller Cluster Password:



4. Enter a new password and click OK when done.

The control plane remains unaffected in the event a single NSX controller fails. However, VMware recommends redeploying the entire cluster and using the Update Controller State mechanism to synchronize the state of the controller cluster.




The Update Controller State option causes logical routers to be redeployed and VXlan to be resynchronized.

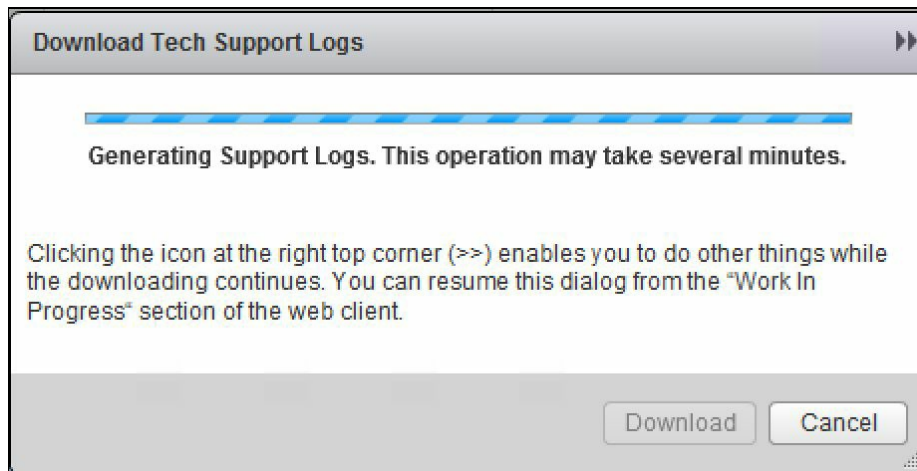
The following steps are used to deploy a fresh controller cluster:

1. Log in to the vSphere web client and navigate to Home | Networking & Security | Installation | Management.
2. Delete the NSX controllers by clicking on each one and then the **X** Delete icon.
3. Deploy a new set of NSX controllers using the standard NSX controller deployment steps.

- Once deployed, go to the NSX Manager pane under the Management tab and click on Actions | Update Controller State:



- You can also download Tech support logs from a specific controller. To do that select a controller and click on  to begin the log collection. You will see the following as soon as you click on the icon:



- Once the support log generation is done, download it to your desired location on your system.
- Configuring a syslog server for NSX controllers is done via an API call. Once configured, NSX Manager sends all audit logs and system events to this syslog server. Configuring the syslog server via the API is the only supported way as of this version.
- To enable syslog on NSX controller, use the following API, which adds a controller syslog exporter and configures it on a controller node:

```
Request
POST https://NSXManager-IP/api/2.0/vdn/controller/{controller-id}/syslog
Request Body:
<controllerSyslogServer>
<syslogServer>SYSLOG SERVER IP</syslogServer>
<port>514</port>
```

```
<protocol>UDP</protocol>
<level>INFO</level>
</controllerSyslogServer>
```

VMware recommends the protocol remain UDP.

9. To query the syslog server, use the following API:

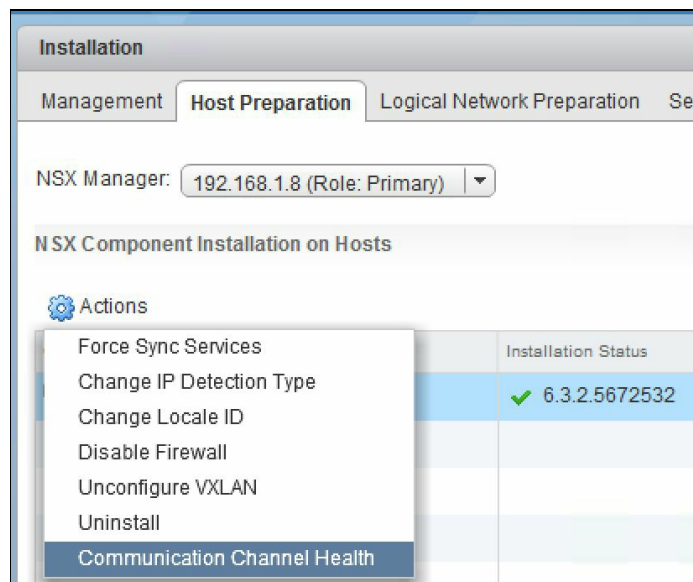
```
Request
GET https://NSXManager-IP/api/2.0/vdn/controller/{controller-id}/syslog
```

If you wish to delete the syslog configuration, use the following API:

```
Request
DELETE https://NSXManager-IP/api/2.0/vdn/controller/{controller-id}/syslog
```

NSX constantly checks on communication health between NSX Manager and all of its associated and deployed components. To check on the status of the communication channel health:

1. Log in to the vSphere web client and navigate to Home | Networking & Security | Installation | Host Preparation.
2. Select a cluster and expand it. Select any host and click Actions and then Communication Channel Health:



The Communication Channel Health information is displayed:

Summary

We started this chapter by looking at NSX Manager settings. We looked at many settings including changing the date and time, configuring a syslog server, configuring DNS servers, gathering tech support logs, and SSL certificates configuration. We then looked at backing up and restoring our NSX Manager, and configuring our NSX Manager with an active directory. NSX Controller operations described how to change NSX controller passwords, which is a best practice, and also talked about how to download tech support logs for controllers and configure them with a syslog server. Finally, we looked at reviewing communication channel health between NSX Manager and all of its components.

Conclusion

By now you should have a pretty good understanding about NSX and should feel comfortable in using it. In this book, we went over the basics and step-by-step deployment of NSX and its features. It is important to understand that this is not the end, but in fact the beginning of the NSX ecosystem. There is a lot more to learn and understand about NSX, but the preceding eight chapters will give you a head start.

There are lots of other articles about NSX that are worth reading along with this book. In this conclusion chapter, I will list these recommended articles; I suggest you go through them to get a better understanding of NSX features and use cases:

- First there is always VMware's NSX documentation. NSX documentation is very well written and has a lot of information that will further help you. More specifically, you should read up on the administration guide. The administration guide is available at:

<https://docs.vmware.com/en/VMware-NSX-for-vSphere/6.3/com.vmware.nsx.admin.doc/GUID-B5C70003-8194-4EC3-AB36-54C848508818.html>

- You can also read the NSX design guide. Although this was written for version 6.1, its concepts are still valid and give you a pretty good idea about some industry best practices and design considerations when deploying NSX components, features, and services. You can get the guide at:

<https://www.vmware.com/files/pdf/products/nsx/vmw-nsx-network-virtualization-design-guide.pdf>

- It's easy enough to deploy NSX in a greenfield environment, but when it comes to brownfield environments this technical whitepaper can be quite helpful in getting you started:

<https://www.vmware.com/files/pdf/products/nsx/VMware-NSX-Brownfield-Design-and-Deployment-Guide.pdf>

- There are a lot of vendors who have provided technical white papers that are worth a read. Palo Alto networks has been at the forefront by enabling its PAN firewall to work and inter-op with NSX. Here is a good technical white paper on this integration:

<https://www.vmware.com/files/pdf/products/nsx/NSX-Palo-Alto-Networks-WP.pdf>

- Then VMware NSX and F5 design guide is another such document and supplies a lot of insight into deployment and design ideas:

<https://f5.com/Portals/1/Premium/Architectures/RA-VMware-NSX-Design-Guide.pdf>

- To learn more about some cross-vCenter and multi-site solutions, the following blog is quite helpful:

<https://blogs.vmware.com/networkvirtualization/2016/03/cross-vc-nsx-multi-site-solutions.html/>

If you don't have a lab ready to go you can always start learning NSX instantly by using VMware's hands-on labs. Just go to <http://labs.hol.vmware.com/HOL/catalogs/> and explore the different NSX labs that are deployed instantly and can be accessed by any machine that has connectivity and a web browser. You will need to register to access the lab; registration is free and open. Each lab comes with lab exercise guide and I recommend going over the labs in this order:

- **HOL-1703-SDC-1 - VMware NSX: Introduction and Feature Tour**—This gives you a basic hands-on introduction to NSX. Doing the exercises in order lets you deploy and test all NSX functionality.
- **HOL-1703-USE-2 - VMware NSX: Distributed Firewall with Micro-Segmentation**—This gives you a deeper look into NSX including deploying and configuring some of its advanced features.
- **HOL-1703-USE-3 - VMware NSX: Operations and visibility**—Learn all the operational features and how to monitor an NSX environment.
- **HOL-1703-SDC-4 - VMware NSX: Installation and configuration**—Learn to install and configure NSX with this virtual lab!

All the preceding links will be posted on the book's website at <http://www.rjapproves.com/learningVMwareNSX/>.

You can also always reach me via my blog site www.RJApproves.com or my Twitter account [@RJApproves](https://twitter.com/RJApproves) to discuss use cases and any issues that you run into. You can connect with me professionally at LinkedIn at www.linkedin.com/in/rjapproves.

I hope you enjoyed this book and that you will continue on your path to consuming the rich features and services offered by NSX. I wish you all the best!