Virtual IPS Administration Guide

Revision A

McAfee Network Security Platform 8.3

For Private, Public, and Hybrid Clouds

COPYRIGHT

TRADEMARK ATTRIBUTIONS

LICENSE INFORMATION

License Agreement

Contents

Preface **5**
About this guide . . . 5
Conventions . . . 5
Find product documentation . . . 5

1 About Cloud and Data Center Security **7**

Network Security Platform Data Center Security

2 Virtual IPS Sensor deployment on VMware ESX and KVM **11**
Virtual IPS Sensors - advantages . . . 13
Virtual IPS Sensor models . . . 13
Requirements for deploying Virtual IPS Sensors . . . 14
Considerations . . . 15
Supported modes for Virtual IPS Sensors . . . 15
Features not supported by Virtual IPS Sensors . . . 15
Deploying Virtual IPS Sensors on VMware ESX Server . . . 16
Install Virtual IPS Sensors on VMware ESX . . . 16
Deployment scenarios for Virtual IPS Sensors . . . 30
Verify the deployment . . . 67
Deployment of Virtual IPS Sensors on KVM . . . 67
Access KVM . . . 68
Install the Virtual IPS Sensor on KVM . . . 68
Troubleshooting scenarios . . . 80
Uninstall the Virtual IPS Sensor from KVM . . . 81
Add the Virtual IPS Sensor in the Manager . . . 82
Manage Virtual IPS Sensor licenses . . . 83
Generate the Virtual IPS Sensor License Compliance report . . . 86

3 IPS for virtual networks using VMware NSX **89**
Securing virtual networks with Open Security Controller . . . 89
Security challenges in an SDDC . . . 90
How OSC secures virtual networks? . . . 91
Advantages of Open Security Controller . . . 93
Virtual IPS Sensors deployed through Open Security Controller . . . 93
Deploying next generation IPS service to a virtual network . . . 94
Terminologies . . . 95
Components involved in IPS service . . . 97
High-level steps to implement a security service . . . 99
How the IPS service works . . . 101
Requirements for deploying IPS service . . . 105
Considerations . . . 106
Prepare a VMware ESXi host for NSX . . . 108
Define an IP pool for virtual security appliances . . . 110
Define virtualization connectors . . . 114

Define manager connectors 116
Manage software images for security appliances 119
Manage distributed appliances 124
Jobs and tasks 136
Create a security group in VMware NSX 140
Create a security policy in VMware NSX 144
Apply a security policy to a security group in VMware NSX 147
Configure Virtual Security System to fail-close or fail-open 149
Manager functions regarding IPS service deployment 153
FAQs regarding IPS service 162

Index **165**

Preface

Contents

- *About this guide*
- *Find product documentation*

About this guide

This information describes the guide's target audience, the typographical conventions and icons used in this guide, and how the guide is organized.

Conventions

This guide uses these typographical conventions and icons.

Italic	Title of a book, chapter, or topic; a new term; emphasis
Bold	Text that is emphasized
`Monospace`	Commands and other text that the user types; a code sample; a displayed message
Narrow Bold	Words from the product interface like options, menus, buttons, and dialog boxes
Hypertext blue	A link to a topic or to an external website
	**Note:** Extra information to emphasize a point, remind the reader of something, or provide an alternative method
	**Tip:** Best practice information
	**Caution:** Important advice to protect your computer system, software installation, network, business, or data
	**Warning:** Critical advice to prevent bodily harm when using a hardware product

Find product documentation

On the **ServicePortal**, you can find information about a released product, including product documentation, technical articles, and more.

Task

1. Go to the **ServicePortal** at https://support.mcafee.com and click the **Knowledge Center** tab.
2. In the **Knowledge Base** pane under **Content Source**, click **Product Documentation**.
3. Select a product and version, then click **Search** to display a list of documents.

1 About Cloud and Data Center Security

Cloud computing and virtualization have begun to play a huge role in the IT infrastructure and scalability of an organization's business. However, while the terms virtualization and cloud are used almost interchangeably, it is important to note the distinction that this guide makes between the two terms. Virtualization is the technology that separates the physical infrastructure to create dedicated resources. It separates compute environments from the actual physical hardware thereby allowing servers, workstations, storage, and other systems to be independent of the hardware layer. Virtualization is the foundation on which Cloud computing is based and Cloud computing is the delivery of shared computing resources as a service or on-demand through the Internet. For better readability, it might simply be referred to as the "cloud" in this guide.

Early adoption of the cloud can provide organizations with an opportunity to transform their business models and gain a competitive edge. Cloud environments are generally the most cost-effective solutions compared to conventional on-premise environments both in the short and the long term. Measurable benefits such as lower costs, greater agility, and better resource utilization have spurred initial adoption. Cloud computing is comprised of two specific deployment models which are described here:

- **Public cloud** refers to cloud infrastructure that is managed by a business, academic or government organization or a combination of these. All infrastructure and data reside on the premises of the cloud provider. Public cloud models have all your data and applications residing on the cloud provider's servers. The primary advantage of the public cloud model is the ease to deploy. Any individual or business can sign up for advanced services on the cloud with just a credit card. Obtaining such services without the public cloud would require expensive and time-consuming resources to be set up within your organization network.
- **Data center** refers to cloud computing that delivers the same benefits of public cloud such as scalability and self-service provisioning but through a proprietary architecture. However, while public cloud deployments serve multiple clients or organizations, data centers (they might sometimes be known as private clouds) serve only one organization. Data center deployments keep data and applications within your control so that they do not have to be stored and operated from a third-party organization's infrastructure. From this description, it would seem that security is not a challenge in a data center deployment. However, because data center deployments require new deployments for resource pooling and elastic scalability, it is prone to security challenges that must be anticipated and planned for.

McAfee Network Security Platform is a full featured next-generation IPS solution ready for the unique demands of cloud environments. It is an intelligent security solution that discovers and blocks sophisticated threats in the network with unmatched speed, accuracy, and simplicity. Combined with network virtualization and security platforms. Network Security Platform delivers best-in-class enterprise security against sophisticated attacks on virtual infrastructures. You are able to deploy it as a standalone Virtual IPS Sensor to monitor both east-west and north-south traffic or as a service that is orchestrated across a software defined data center.

The following type of solutions are available:

- Standalone Virtual IPS Sensor, a virtual instance of the physical IPS Sensor, can be deployed on hypervisors and used to monitor traffic between virtual machines.
- Cluster solution which comprises several Virtual IPS Sensors that are clustered in a single appliance and orchestrated in a data center using a tool such as Open Security Controller. The cluster can be used to deploy IPS as a service in environments that use network virtualization software such VMware NSX.

Standalone and distributed IPS appliances

In a standalone Virtual IPS deployment, any number of Virtual IPS Sensor can be installed per hypervisor. Each Virtual IPS Sensor is managed separately through the Network Security Manager (Manager). You can have different policies configured for each Virtual IPS Sensor. Maintenance and troubleshooting of each Virtual IPS Sensor is also carried out individually.

In the distributed solution, a logical container contains several Virtual IPS Sensors within itself. The container is known as Virtual Security System and, unlike Virtual IPS Sensors, the Virtual Security System appears in the Manager as a single device with several instances of the Virtual IPS Sensors. The Virtual Security System is orchestrated using Open Security Controller and is managed through VMware NSX. When a security policy is applied to a Virtual Security System, all instances of the Virtual IPS Sensor within it are updated with the same policy. Each Virtual IPS Sensor is configured similarly but functions independently of the other and provides security to the specific host it is deployed on. Maintenance and troubleshooting of a Virtual Security System involves managing one device in the Manager. For example, when you apply a configuration update to one Virtual Security System appliance, all Virtual IPS Sensor instances contained in it are updated with the same configuration. Such centralized management of several instances of the Virtual IPS Sensor proves beneficial in deploying a scalable security solution across your data center.

In subsequent chapters, we will look at both these deployment models and the various environments in which they can each be deployed.

This section describes the various models and the virtualization environments in which they can be deployed.

Virtual IPS Sensor model	Supported virtualization environment	Type of solution
IPS-VM600	VMware ESX and KVM	Standalone
IPS-VM100	VMware ESX and KVM	Standalone
IPS-VM100-VSS	VMware NSX	Distributed

Network Security Platform Data Center Security

Chapter 2 *Virtual IPS Sensor deployment on VMware ESX and KVM*
Chapter 3 *IPS for virtual networks using VMware NSX*

2 Virtual IPS Sensor deployment on VMware ESX and KVM

Enterprises are moving towards virtual IT infrastructures such as private and public cloud, virtual data centers for servers, and virtual machines for clients. Security requirements for a virtual network might vary vastly when compared to physical networks. For example, monitoring of peer-to-peer traffic and access control in a virtual network have their own challenges. Based on the network architecture and security requirements, virtual security products are required to protect virtual IT infrastructures. Even for physical networks, virtual security products can bring in savings in terms of cost and space.

A Virtual IPS Sensor is McAfee's virtual next-generation IPS product. It is a virtual instance of the NS-series Sensor software, which you can install as a virtual appliance on

- VMware ESX server
- Kernel-based Virtual Machine (KVM)

You do not require the Sensor hardware to deploy a Virtual IPS Sensor. Though primarily designed to protect virtual networks, you can deploy a Virtual IPS Sensor to protect physical networks as well.

Similar to a physical Sensor, you use a Manager to configure and manage Virtual IPS Sensors. The Manager can be installed on a physical server or on a virtual machine. Also, you can use the same Manager to manage both virtual and physical Sensors including heterogeneous Sensor environments.

A Virtual IPS Sensor supports several features that are supported by a physical Sensor. Except for the fact that you deploy Virtual IPS Sensors in a virtual environment, the process of configuring and managing them is similar to that of physical Sensors. Virtual IPS Sensors also function similar to their physical counterparts when it comes to protecting your networks. With the added advantage of being a virtual instance, you can deploy a Virtual IPS Sensor to protect various network architectures. Some of the common scenarios are covered in this document.

You install a Virtual IPS Sensor in a virtualization environment. You then deploy this Virtual IPS Sensor to inspect traffic between:

- Virtual machines (VMs) on the virtualization environments
- East-west traffic between VMs in virtual environments
- North-south traffic inside and outside the perimeter

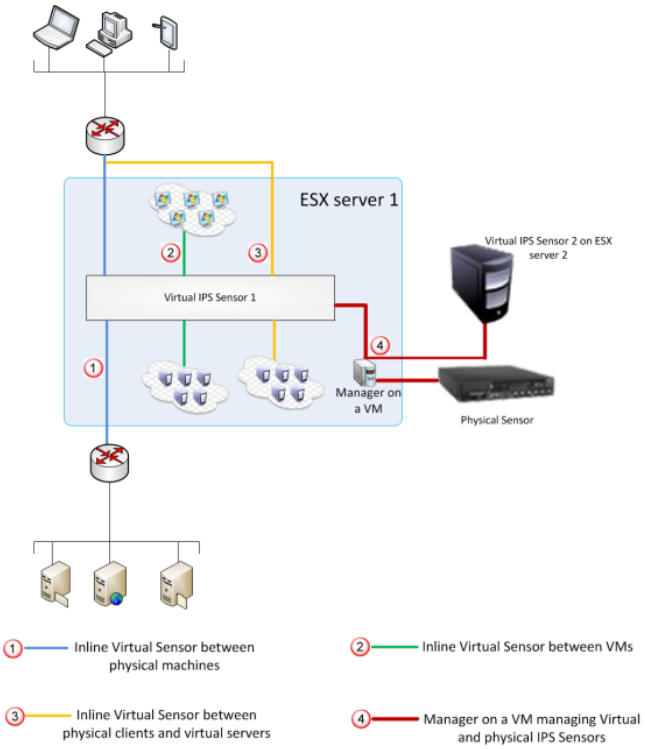

Figure 2-1 Virtual IPS Sensor deployment

To use the information in this document, familiarity with the following might be required:

- For VMware ESX Server: Administration of VMware ESXi hosts including virtual networks within VMware ESXi hosts.
- For KVM: Administration of KVM hypervisor including virtual network configuration.
- Management of guest virtual machines.
- Installation, configuration, and management of Network Security Sensor and Manager.

Contents

- *Virtual IPS Sensors - advantages*
- *Virtual IPS Sensor models*
- *Requirements for deploying Virtual IPS Sensors*
- *Considerations*
- *Deploying Virtual IPS Sensors on VMware ESX Server*
- *Deployment of Virtual IPS Sensors on KVM*
- *Add the Virtual IPS Sensor in the Manager*
- *Manage Virtual IPS Sensor licenses*
- *Generate the Virtual IPS Sensor License Compliance report*

Virtual IPS Sensors - advantages

The following are the advantages of Virtual IPS Sensors:

- As with any virtual technology, Virtual IPS Sensors too bring in savings in terms of space, cost of the appliance, maintenance costs, power to run the appliances, power for air-conditioning, and so on.
- Virtual IPS Sensors enable inspection of traffic that never leave a virtualization environment. Also, you can implement security policies, (Exploit, DoS, Advanced Malware policies) on this traffic.
- Virtual IPS Sensors are very easy and quick to deploy. Therefore, they can scale up to any rapid expansions of your virtual network.
- You can deploy Virtual IPS Sensors without any physical access to the virtualization environment.
- Though protecting virtual networks has its unique challenges, there is no compromise on security with respect to Virtual IPS Sensors. It protects networks similar to a physical Sensor, including features such as application identification and visualization, Firewall policies, Advanced Malware policies and so on. Also, just like physical Sensors, a Virtual IPS Sensor can integrate and communicate with other McAfee products.
- Based on factors such as the size of your network and network architecture, it is possible to use Virtual IPS Sensors to protect both your virtual and physical networks.
- You can use the same Manager as a single-point-of-control for managing your physical and Virtual IPS Sensors.
  - Provides the flexibility of using virtual, physical, or a mixture of virtual and physical Sensors to protect your networks.
  - Managing Sensors is simplified and centralized.
  - You can define the security policies to both virtual and physical networks at one place. So, the task of defining and implementing security policies for the virtual networks still rests with your security experts.
  - Provides a consolidated view of the threat information from both virtual and physical Sensors.
  - Enables consistent and consolidated report generation for all your networks.

  This Manager can also be installed on a VM.
- Virtual IPS Sensors are self-contained with respect to protecting your networks. That is, there is no requirement for any third-party applications to protect virtual or physical networks.
- Virtualization of virtual ports is available. That is, you can virtualize the interfaces of a physical Sensor by creating sub-interfaces based on VLAN or CIDR. Following the same procedure, you can create sub-interfaces of the monitoring ports of a Virtual IPS Sensor.

Virtual IPS Sensor models

The table describes the available Virtual IPS Sensor models.

Model	Maximum Sensor throughput	Number of monitoring ports	Management port	Response port	Logical CPU Cores	Memory	Storage
IPS-VM100	100 Mbps	4	1	1	3	4 GB minimum required.	8 GB
IPS-VM600	600 Mbps	6	1	1	4	6 GB minimum required.	8 GB

The kind of traffic being inspected and the features that you enable are some of the primary factors that affect the throughput of a Sensor. For these details and other capacity values for Virtual IPS Sensors, see the *Virtual IPS Sensor capacity by model number* section in the *McAfee Network Security Platform Best Practices Guide.*

Requirements for deploying Virtual IPS Sensors

This section discusses the server requirements for deploying Virtual IPS Sensors on VMware ESX server and KVM.

Server requirements

Component	Minimum requirement
Virtualization software	• VMware ESX 5.0 and later OR • KVM 2.0.0

Hardware	• CPU:
	  • Intel Xeon processor 5000 series or higher
	  • Minimum required processor speed is 1 GHz; recommended is 2 GHz or more
	  • VM100-IPS: 3 logical CPU cores
	  • VM600-IPS: 4 logical CPU cores
	• Memory:
	  • VM100-IPS: 4 GB
	  • VM600-IPS: 6 GB
	• Ethernet ports or bridges:
	  • VM100-IPS: 4 Ethernet ports or bridges
	  • VM600-IPS: 6 Ethernet ports or bridges

Other requirements

- Network Security Manager 8.1.x or later for VMware ESX installed on a virtual or physical machine.
- Network Security Manager 8.3.x or later for KVM installed on a virtual or physical machine.
- You require one license per Virtual IPS Sensor. The license is also specific to the Virtual IPS Sensor model you purchased. Make sure you have secured the required number of licenses from McAfee.
- You must exclude the Virtual IPS Sensor from VMware Distributed Resource Scheduler (DRS).

- Do not install VMware tools for a Virtual IPS Sensor.
- For optimal and predictable performance follow these guidelines.
  - You must configure each Virtual IPS Sensor VM to execute on as many cores as required for the Sensor model by assigning CPU affinity to the VM. For example, an IPS-VM100 VM must be affinitized to 3 logical cores.
  - CPU, memory, and disk resources must be reserved for the Sensor.

Considerations

Review this section and its sub-sections before you deploy a Virtual IPS Sensor.

- Based on how you deploy the Virtual IPS Sensor, you might have to re-configure some of the vSwitches on your VMware ESXi host.
- Currently, For VMware ESX, Virtual IPS Sensor deployment involves standard vSwitches and distributed vSwitches.
- You need an Active Fail-Open kit to deploy Sensor monitoring ports in fail-open mode. This is applicable to scenarios where the Virtual IPS Sensor is between two physical network devices. For inline inspection of traffic to virtual machines, only fail-closed mode is applicable.
- Currently, failover-deployment of Virtual IPS Sensors is not supported.

Under test conditions, it is observed that it takes less than a minute for a Virtual IPS Sensor to restart and become fully operational again. This factor mitigates the risk of significant network disruption due to inline Virtual IPS Sensor monitoring ports going down. It is recommended to deploy **Active Failopen Kit** to avoid network disruption due to a restart of the Virtual IPS Sensor.

Supported modes for Virtual IPS Sensors

You can deploy a Virtual IPS Sensor in the following modes:

- SPAN mode
- Inline fail-closed mode
- For inline fail-open mode, you must use an external Active Fail-Open Bypass Kit

Tap mode is not applicable to Virtual IPS Sensors.

Features not supported by Virtual IPS Sensors

The following are the list of features that are not supported by Virtual IPS Sensors.

Table 2-1 Features not supported

Feature name
TAP mode
SSL decryption
Sensor failover
Cloud threat detection
Traffic prioritization with application content
Jumbo frame parsing

Deploying Virtual IPS Sensors on VMware ESX Server

To deploy Virtual IPS Sensors, you must first install the Virtual IPS Sensors and establish trust with a Manager. After trust is established, you can deploy the Virtual IPS Sensor for protecting your networks. How you configure the Virtual IPS Sensor, however, depends on the network architecture and your security needs.

The following is a high-level procedure that you can consider to install and subsequently deploy a Virtual IPS Sensor:

1 Verify your VMware ESX server meets the hardware and software requirements. See Requirements for deploying Virtual IPS Sensors on page 14.

2 Install the Virtual IPS Sensors and establish a trusted communication channel between the Virtual IPS Sensors and a Manager. See Install Virtual IPS Sensors on VMware ESX on page 16.

3 Determine how you want to deploy the Virtual IPS Sensor and configure it accordingly. See Deploying Virtual IPS Sensors on VMware ESX Server on page 16.

Install Virtual IPS Sensors on VMware ESX

Before you can deploy a Virtual IPS Sensor to protect your network, you must install the Virtual IPS Sensor and establish trust between the Virtual IPS Sensor and the Manager.

The following are the high-level steps to install a Virtual IPS Sensor:

1 Identify the network to which you in which you want to place the Virtual IPS Sensor and the Manager. Accordingly, identify the vSwitch and the port group for the management port and the Manager. You can use the default switch port group of vSwitch0, which is VM Network, to connect the Sensor management port. However, if required, you can also create a vSwitch to connect the management port. To create a vSwitch for management port connectivity, see Create a standard vSwitch for the management port on page 16.

2 For every Virtual IPS Sensor that you plan to deploy, import the required licenses in the Manager. See Manage Virtual IPS Sensor licenses on page 83.

3 Add the Virtual IPS Sensor in the Manager. See Add the Virtual IPS Sensor in the Manager on page 82.

4 Install the Virtual IPS Sensor and establish trust with the Manager. See Install Virtual IPS Sensors on VMware ESX on page 16

Create a standard vSwitch for the management port

Before you begin

To create a vSwitch, you might be required to connect an additional physical NIC on the VMware ESX.

After you install the Virtual IPS Sensor, you need a standard vSwitch to which you connect the management port of the Virtual IPS Sensor. If you have installed the Manager on the same VMware ESX, you can connect it to this switch as well. When you create a standard vSwitch, VMware ESX creates a default port group for this vSwitch.

Task

1 Connect an additional physical NIC on the VMware ESX to the adjacent physical switch.

This is required for the vSwitch that you are creating.

2 Log on to the VMware ESX as the root user in VMware vSphere Web Client.

3 In the vSphere **Home** tab, select **Hosts and Clusters.**

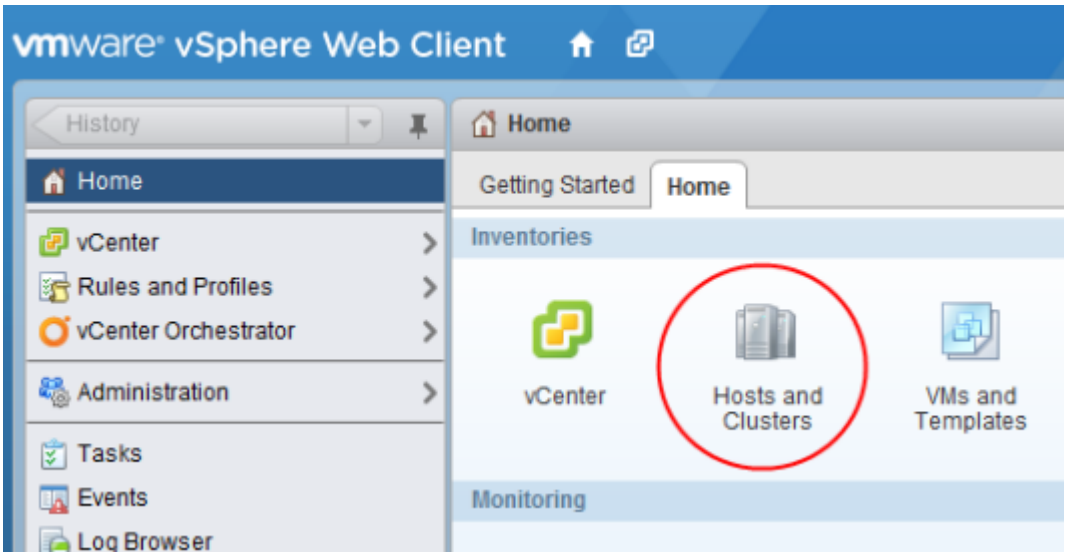

4 Select the required VMware ESX server and select **Manage** | **Networking** | **Virtual switches.**

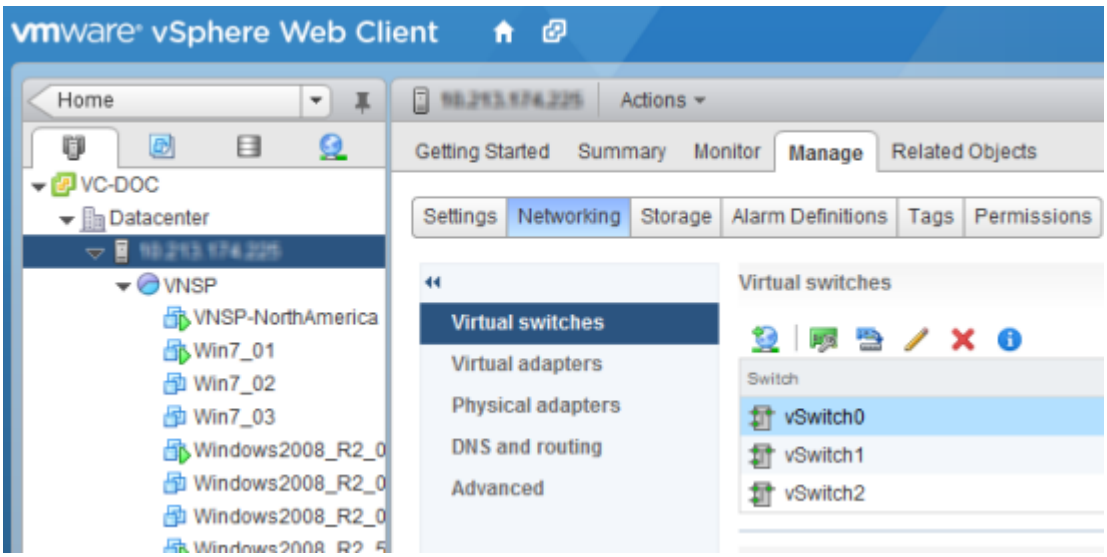

5 Click on the **Add host networking** icon.

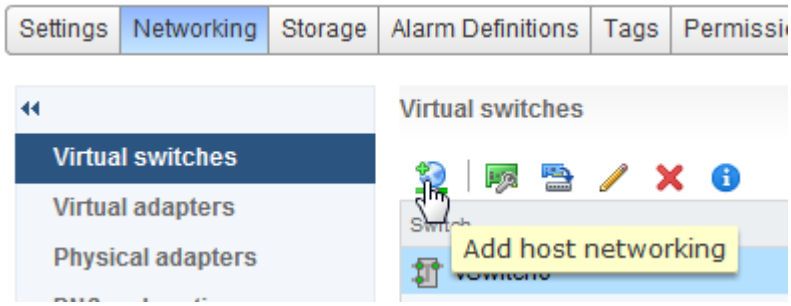

6 For **Select connection type**, as an example select **Physical Network Adapter** and click **Next**.

Selecting the connection type depends on your network design and requirements. For example, if the Manager is installed on a physical machine, then you must select physical network adapter. If not the Sensor and the Manager cannot communicate. If the Manager is installed on a virtual machine and you plan to use the same vSwitch for both the Sensor management port and the Manager, then you might choose **Virtual Machine Port Group for a Standard Switch**. For the scenarios explained in this document, select **Physical Network Adapter**.

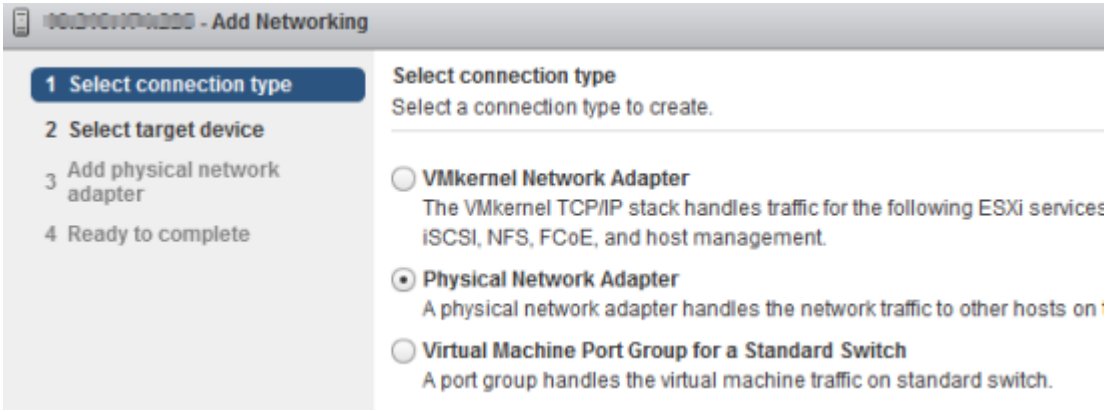

7 For **Select target device**, select **New standard switch**, the required number of ports, and then click **Next**.

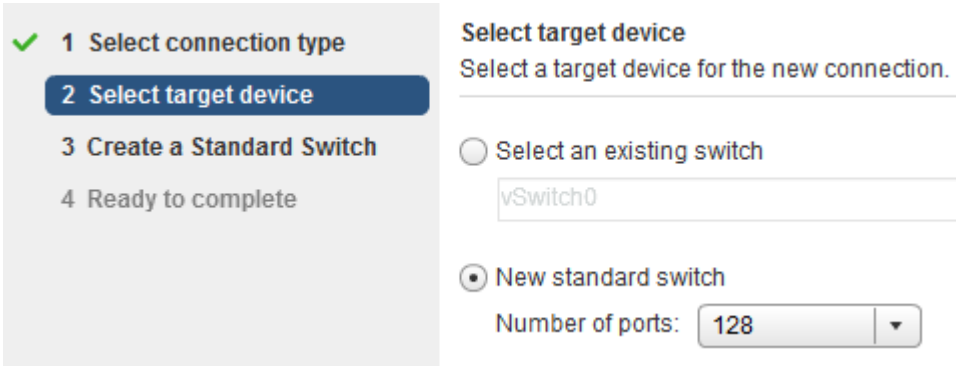

8 For **Create a Standard Switch** step, click **Add adapters** icon.

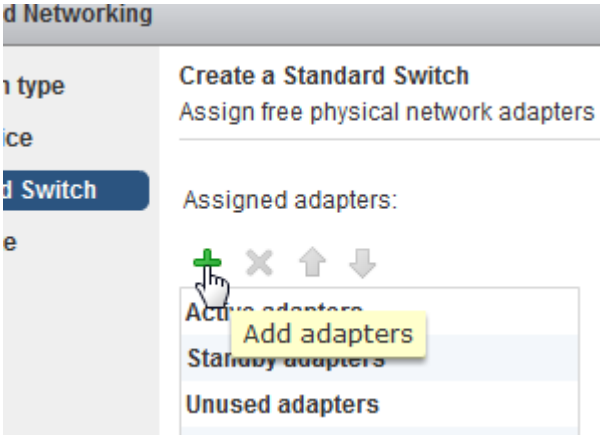

9 In the **Add Physical Adapters to the Switch** dialog box, select the required network adapter and click **OK**.

Make sure that a physical NIC corresponding to the network adapter you selected is connected to the network.

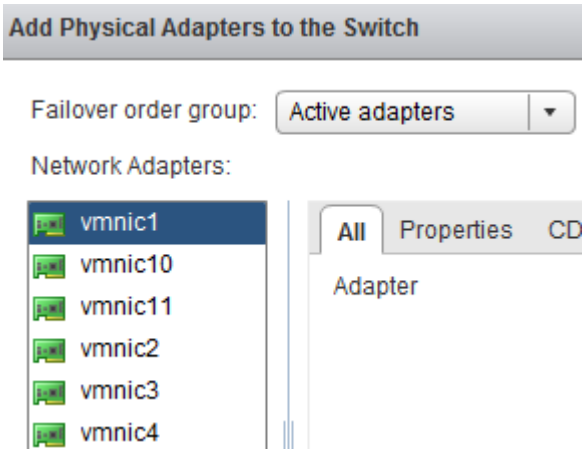

10 Verify the properties of the selected adapter and click **Next** and then select **Finish.**

The vSwitch that you created is listed in the **Virtual switches** section.

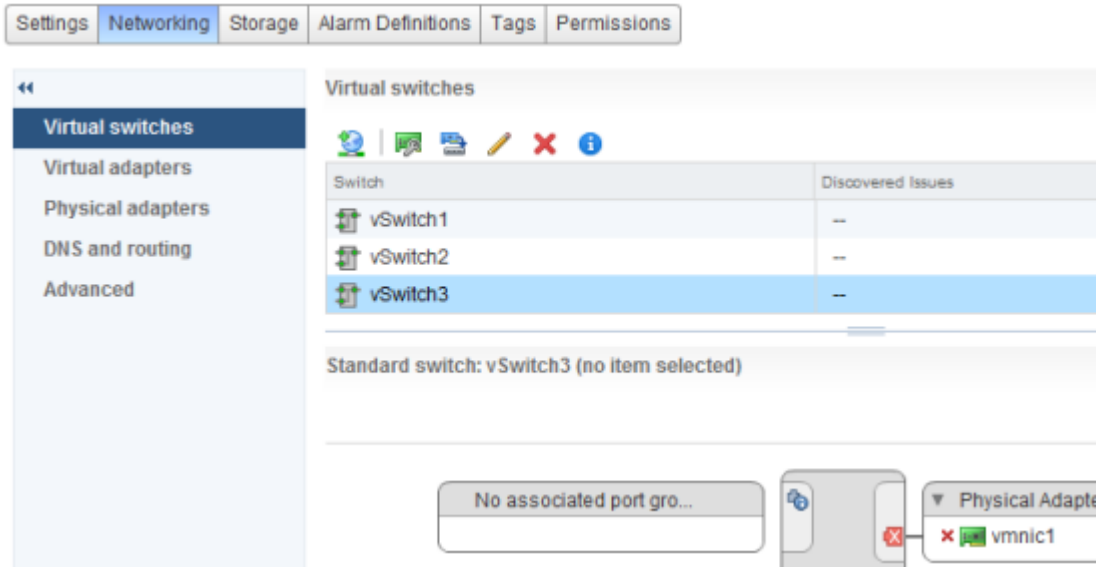

11 Add the required switch port groups for the vSwitch that you created; select that vSwitch and then click on the **Add host networking** icon.

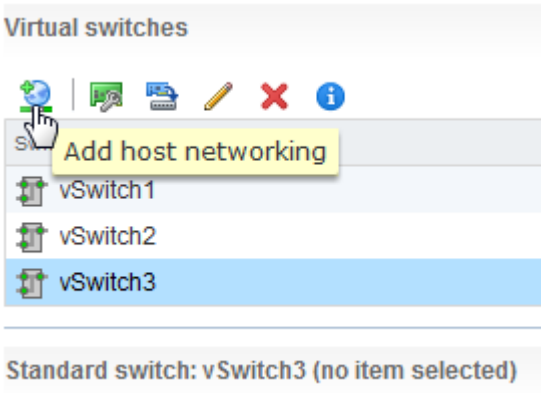

12 In the **Select connection type** step, select **Virtual Machine Port Group for a Standard Switch** and then click **Next.**

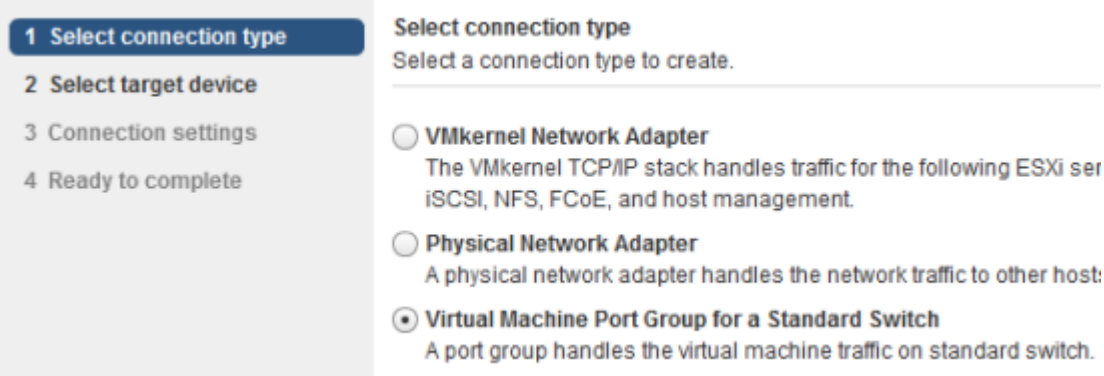

13 In the **Select target device** step, select **Select an existing standard switch** and make sure the vSwitch that you created is selected. Then click **Next.**

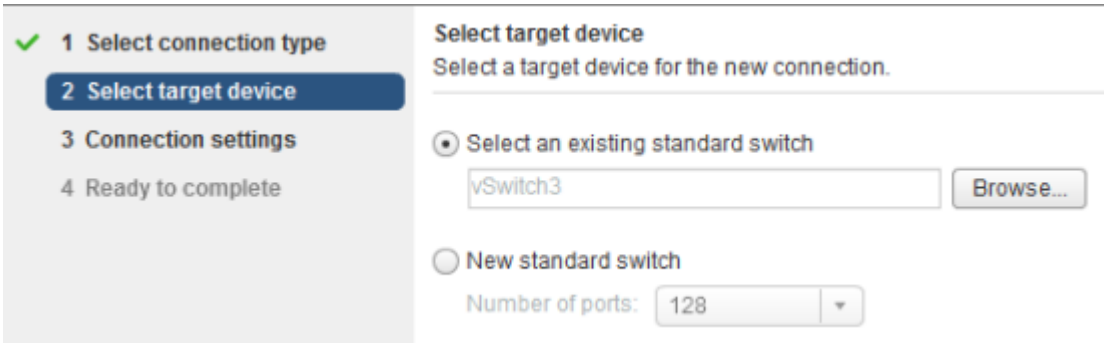

14 In the **Network Label** field, enter the required name for the default port group that the wizard creates for the switch.

You can modify **Network Label** later. For easier management, you can name it as *Management Port Group*.

15 Optionally, in **VLAN ID** field, select **All (4095)** and select **Next.**

If required, you can specify the required VLAN ID as per your network configuration.

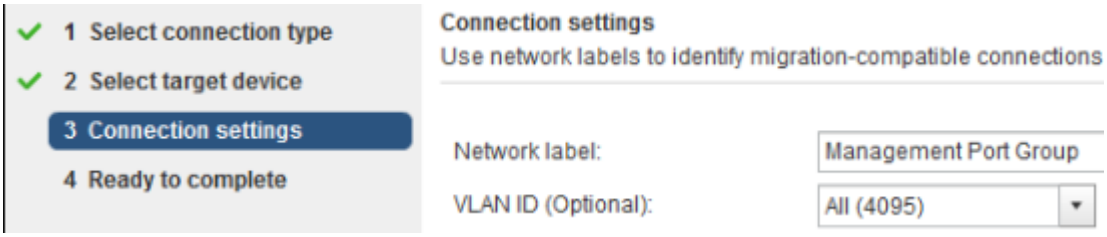

16 Review the details displayed in the **Ready to complete** step and click **Finish.**

This vSwitch is now listed with the switch port group that you created.

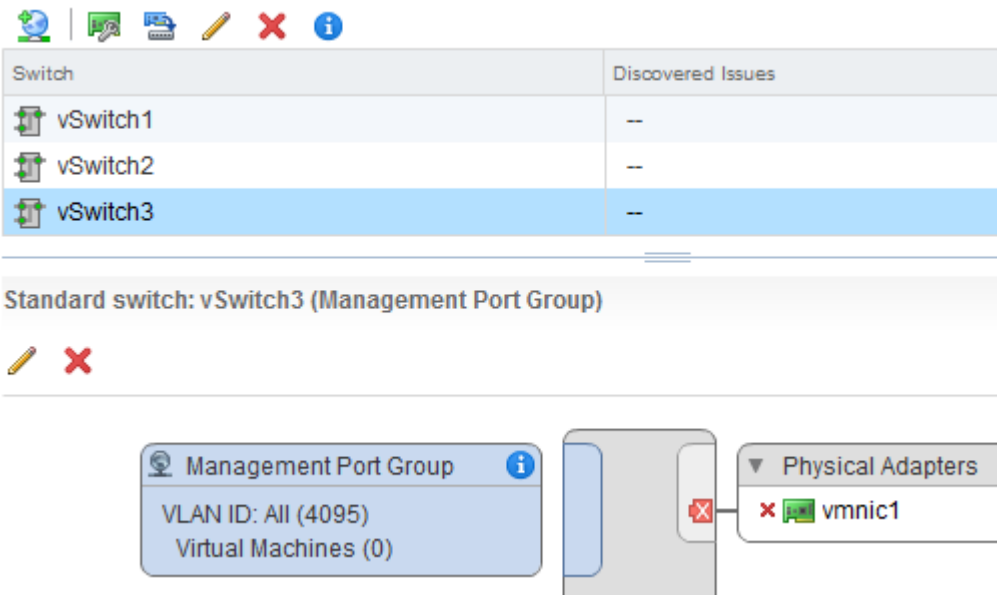

17 Move the mouse over the physical adapter and click on it.

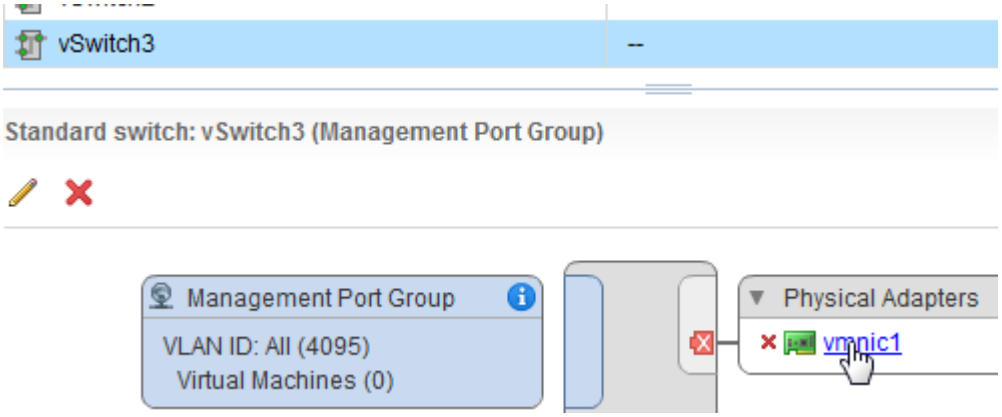

18 Click on the **Edit adapter speed** icon.

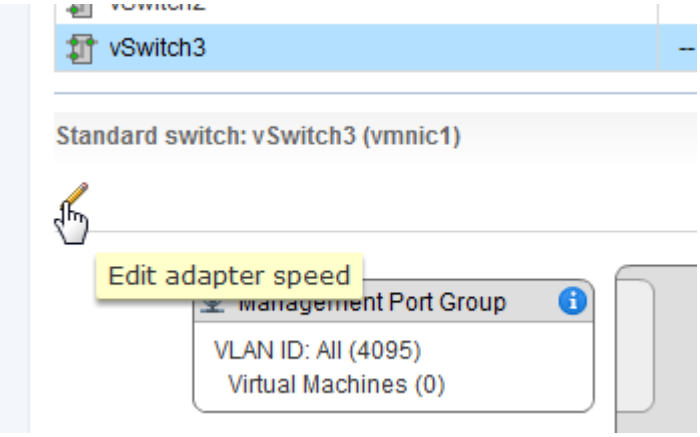

19 Verify if the **Configured Speed, Duplex** is set to **Auto negotiate.**

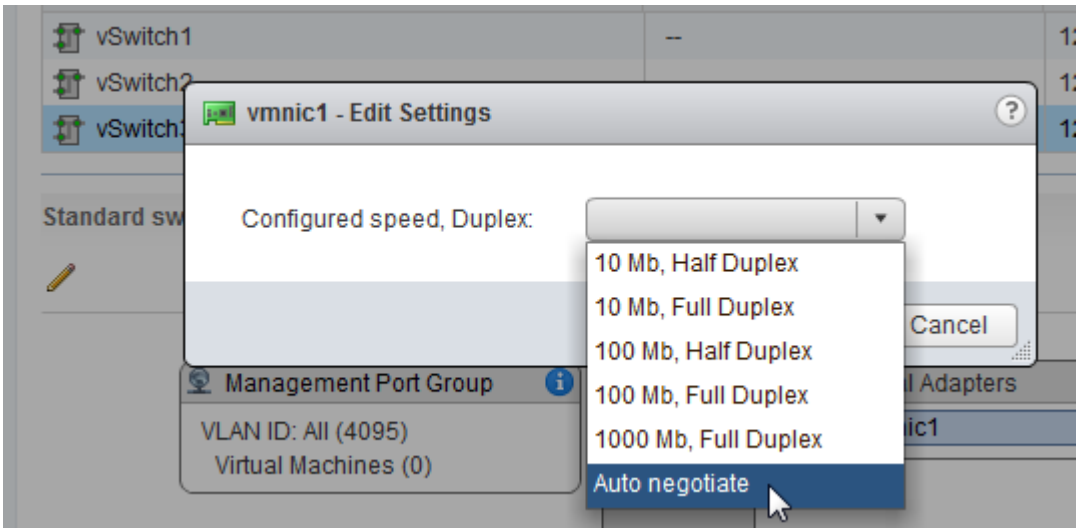

For other property values, you can leave them with the default values.

Install the Virtual IPS Sensor

Before you begin

- The Virtual IPS Sensor installation file is a .ova file. Make sure this file is accessible from your client machine.
- Standard vSwitches and switch port groups are available for the Sensor management port and monitoring ports that you plan to use. Consider that you are installing IPS-VM100, which has one management port, one response port, and 4 monitoring ports. Currently you plan to deploy ports 1-2 in inline mode between 2 vSwtiches. You

do not plan to use ports 3 and 4 for now. For this example, make sure you have the following:

- Standard vSwitch with switch port group for the management port. The Virtual IPS Sensor must be able to communicate with the Manager through this switch port group.
- Two standard vSwitches with switch port groups for monitoring port 1 and 2. That is, the Virtual IPS Sensor will act as a bridge between these two vSwitches with port pair 1-2 inline between these two vSwitches.
- Different dummy switch port groups for the Sensor ports that you do not plan to use now — response port and monitoring ports 3 and 4.

- You have added the Virtual IPS Sensor in the Manager.

Task

1 Log on to the ESX as the root user in VMware vSphere Web Client.

2 In the vSphere **Home** tab, select **Hosts and Clusters**.

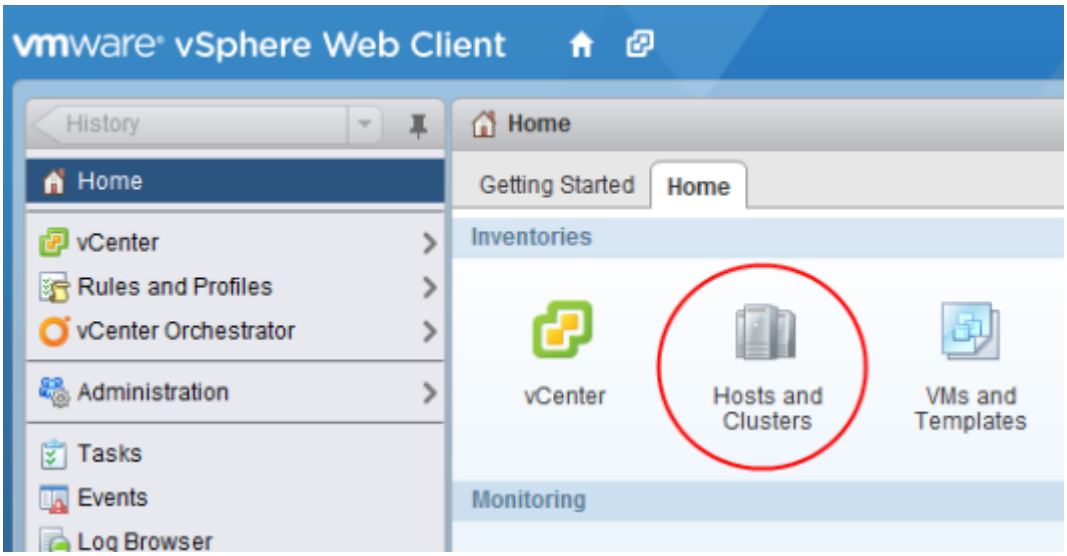

3 Navigate to the required node such as a resource pool, right-click, and select **Deploy OVF Template.**

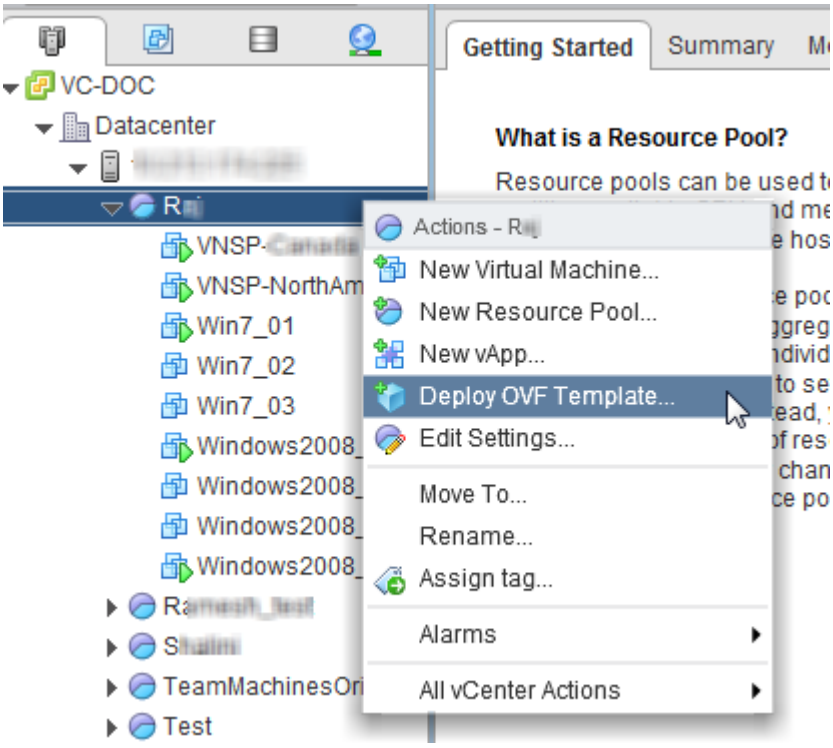

4 Click **Browse** and locate the .ova file.

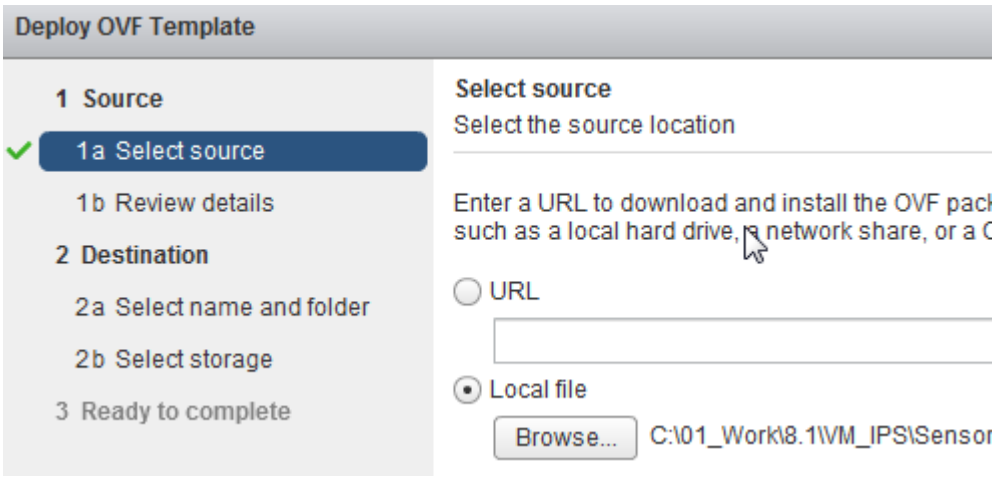

5 Review the details and click **Next**.

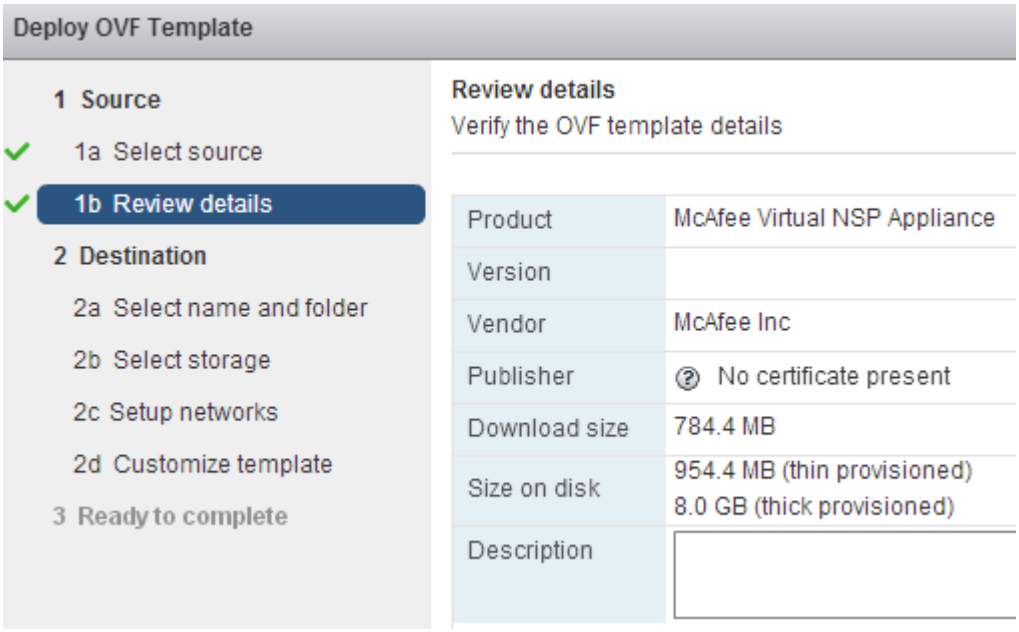

6 In the **Name** field, enter a name for the Virtual IPS Sensor and also select the corresponding datacenter.

Preferably enter the same name that you entered when adding the Virtual IPS Sensor in the Manager.

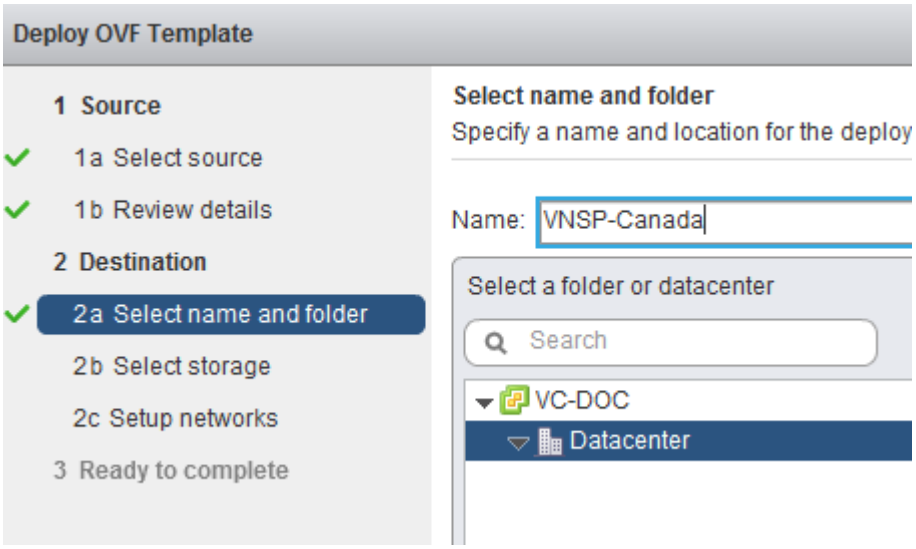

7 From the **Select virtual disk format** list, select **Thin Provision.**

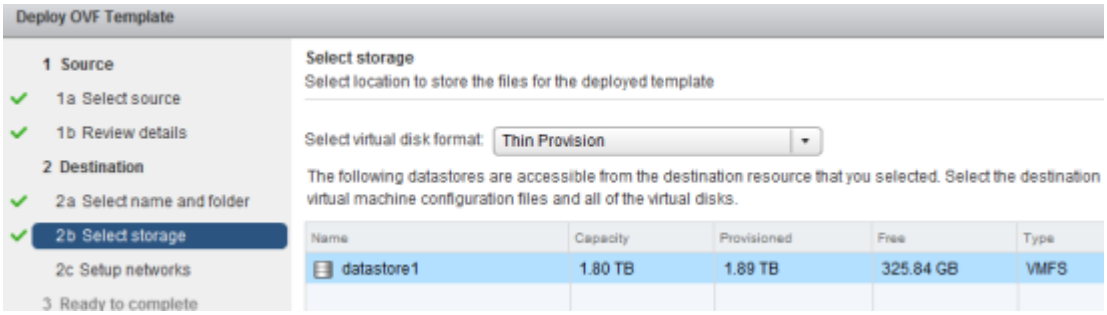

8 In the **Setup networks** section, select the switch port groups for the corresponding Sensor ports.

- For example scenario 2 for example, assign VNSP Client Port for monitoring port 1 and VNSP Server Port for monitoring port 2.
- Assign temporary, non-functional switch port groups to the unused Sensor ports, that is the response port and ports 3 and 4 (for scenario 2). You select a functional switch port group for the response port when you configure the Virtual IPS Sensor for IDS (SPAN).

> You must never assign the same port group to peer monitoring ports. For example, monitoring ports 3 and 4 must not be assigned the same port group. If you do, it results in a loop within the ESX.

- For the management port, assign the port group belonging to the vSwitch that you created for the Sensor management port. This switch port group must enable communication with the Manager server.

> Within the same Virtual IPS Sensor, no **Source** (monitoring port or the response port) should have the same **Destination** (switch port group) as that of the Sensor management port.

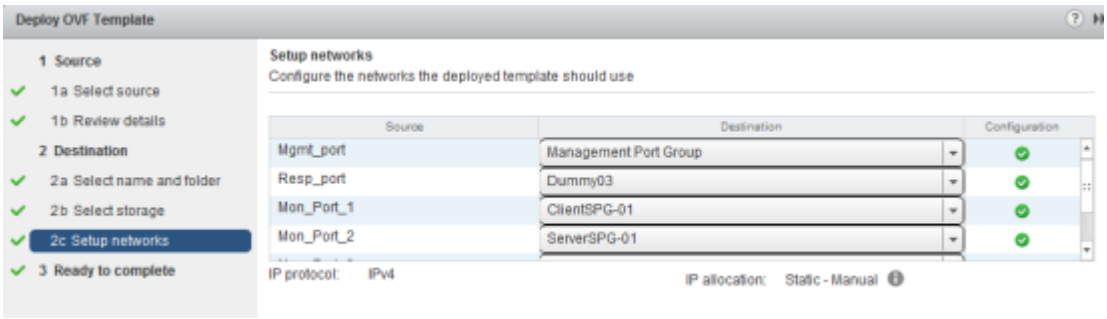

9 In the **Customize template** page, specify the Sensor setup details.

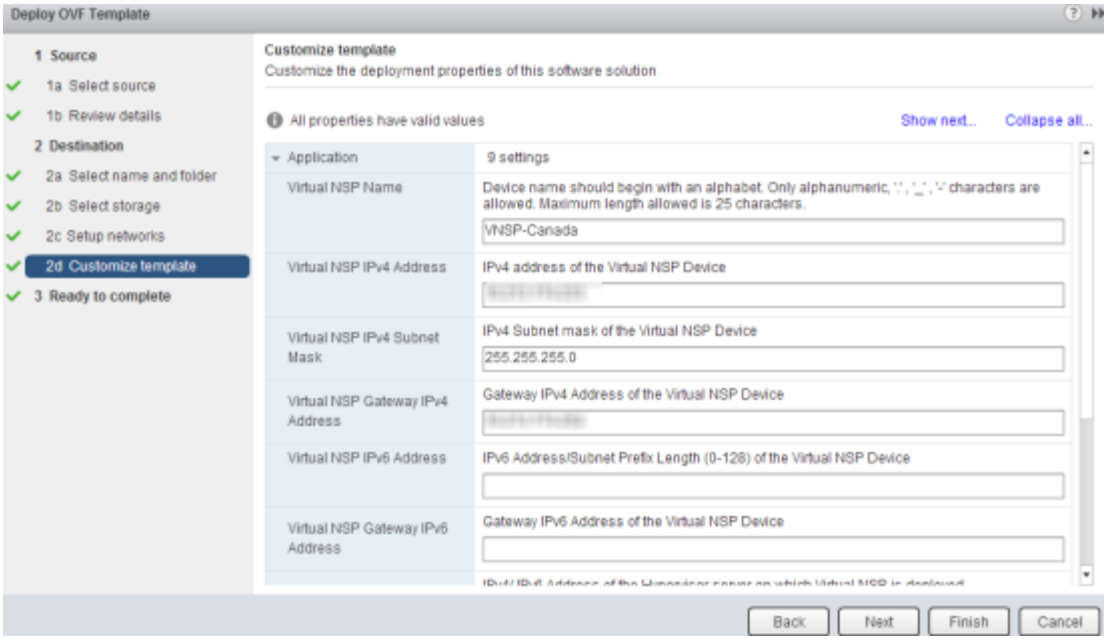

a Enter the same Sensor name that you specified in the Manager.

b Optionally, enter the IPv4 address for the Sensor.

You can specify IPv4, IPv6, or both type of IP addresses to the Sensor.

c If you had specified an IPv4 address, specify the subnet mask for the IPv4 address that you provided.

d If you had specified an IPv4 address, specify the default gateway for the IPv4 address.

This is mandatory if the Sensor needs to communicate outside its network. For example, the Manager could be on a different subnet.

e Optionally, specify an IPv6 address to the Sensor.

f If you had specified an IPv6 address, specify the default gateway for the IPv6 address.

g Specify the IPv4 or IPv6 address of the hypervisor such as VMware ESX server on which you are deploying the Virtual IPS Sensor.

h Specify the Manager's primary IPv4 or IPv6 address.

To specify the Manager's secondary IP address, use the `set manager secondary ip` command in the Sensor CLI after the Sensor is installed.

i Specify the shared secret key and also confirm it by re-entering.

j Click **Next**.

10 Review the configuration that you specified, select **Power on after deployment**, and then click **Finish**.

Click **Back** and make changes, if required. Note that the Sensor setup details that you entered are listed under **Properties**.

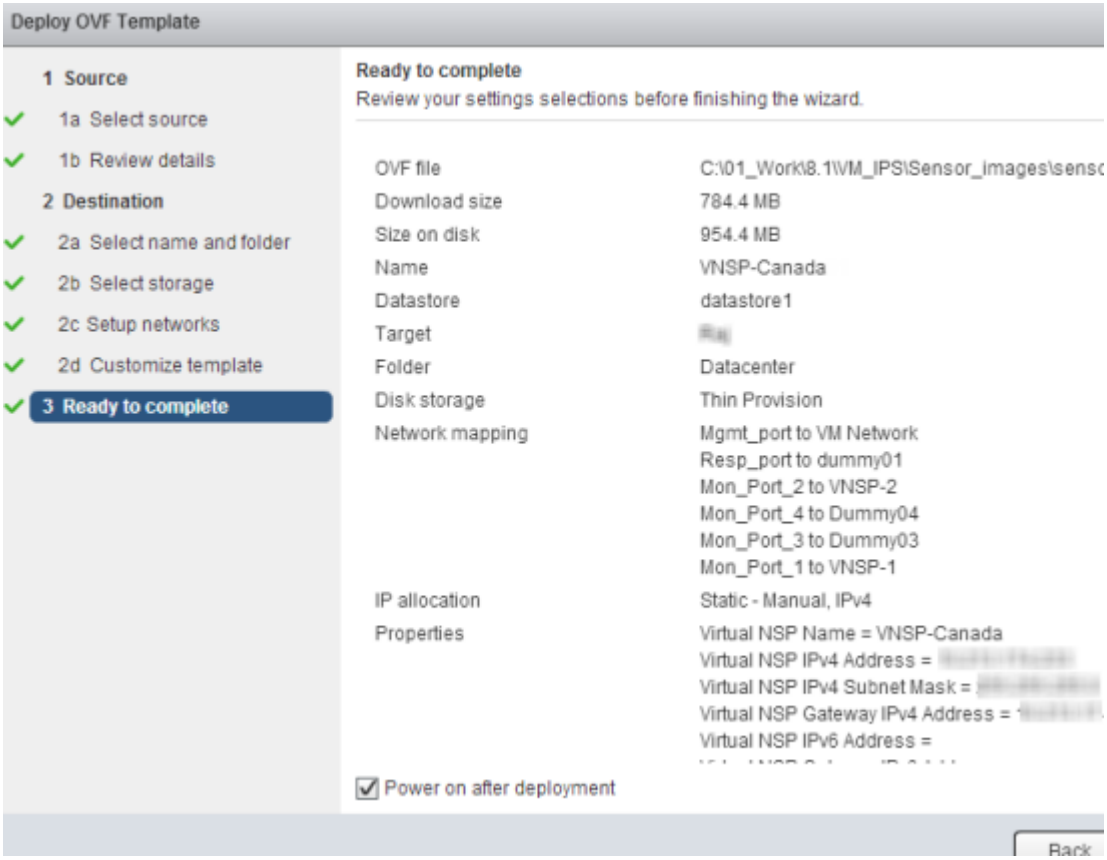

11 After the Virtual IPS Sensor is installed, open an SSH client session to logon to the Sensor.

Alternatively, you can click **Launch Console** in the vSphere Web Client.

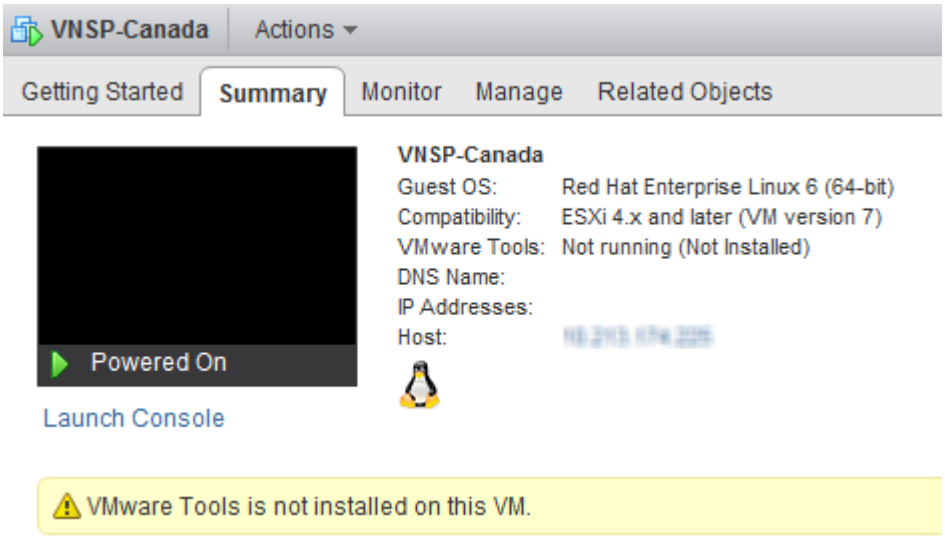

12 In the Sensor CLI, enter admin and admin123 as the login name and password respectively.

13 Use the `status` CLI command to check if trust is established with the Manager and if signature set is present in the Sensor.

If the signature set is not present, you can deploy the signature set from the **Deploy Pending Changes** page of the Manager.

Deployment scenarios for Virtual IPS Sensors

There are subsections that describe some scenarios for you to understand various deployment options. These scenarios are examples used for the sake of explanation; they might not exactly match real or typical network architectures. You can use the scenarios to determine the Virtual IPS Sensor deployment process for your network. As a VMware ESX administrator, your must identify a process that requires the least amount of changes to your network architecture and configuration.

The following is a high-level procedure that you can consider to deploy a Virtual IPS Sensor:

1 Determine how you want to deploy the Virtual IPS Sensor. You can review the scenarios discussed in the subsections. When deciding the deployment type, you can factor in:

- The kind of protection that your network requires. For example, whether you need to place monitoring ports in inline mode against SPAN mode.
- How to deploy the Virtual IPS Sensor with the least changes to your VMware ESX configuration.

2 Follow the procedural information provided for the corresponding scenario.

a Evaluate your VMware ESX deployment and identify the vSwitches that need to be modified and those that need to be created.

b Evaluate the port groups that can be used and the ones that need to be created.

c Identify the vSwitches and port groups for the Sensor monitoring ports.

d Identify the vSwitches and port groups for the clients and servers that you need to protect.

3 Verify if the deployment is functioning as expected. See Verify the deployment on page 67.

Scenario 1: Inspection of traffic between virtual machines in SPAN mode

This scenario involves using a SPAN monitoring port to inspect traffic between virtual machines on the same VMware ESX. Similar to the SPAN mode deployment of a physical Sensor, the SPAN mode deployment of a Virtual IPS Sensor is also simple and non-intrusive.

Scenario description before Virtual IPS Sensor deployment

- The servers are installed on guest VMs on the VMware ESX.
- These servers are connected to a standard vSwitch, *vSwitch0*.
- vSwitch0 has a physical adapter, which is connected to networks outside the VMware ESX.

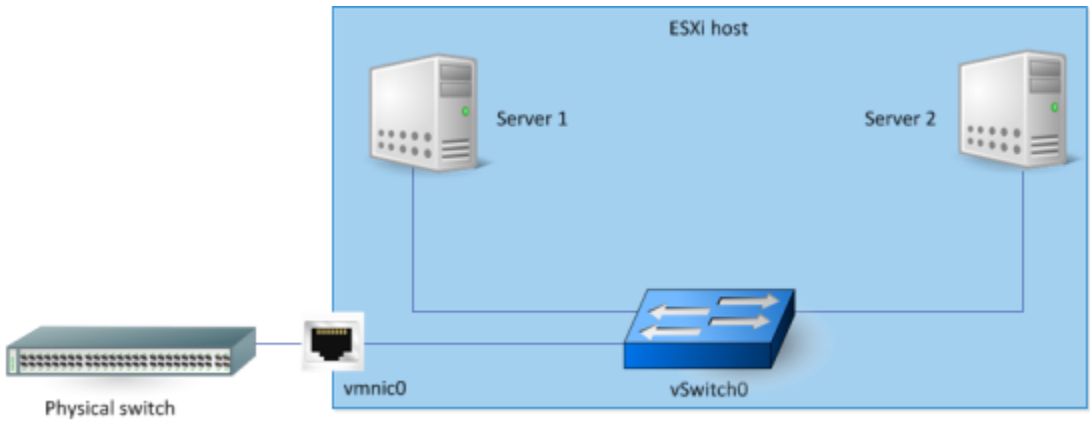

Figure 2-2 Scenario before Virtual IPS Sensor deployment

Scenario description after Virtual IPS Sensor deployment

- In vSwitch0 create a new switch port group set in promiscuous mode.
- The Virtual IPS Sensor is deployed on the VMware ESX.
- Consider that the Manager is installed on a VM connected to vSwitch1.
- In this scenario, the Manager is connected to the management port of the Virtual IPS Sensor through vSwitch1. This vSwitch1 has a physical adapter vminc1. So, you can access the Manager and the Sensor from outside the VMware ESX.
- Monitoring port 1 of the Virtual IPS Sensor is in SPAN mode. This is connected to the promiscuous switch port group in vSwitch0. Therefore, a copy of all the packets of vSwitch0 are sent to port 1 for intrusion detection.

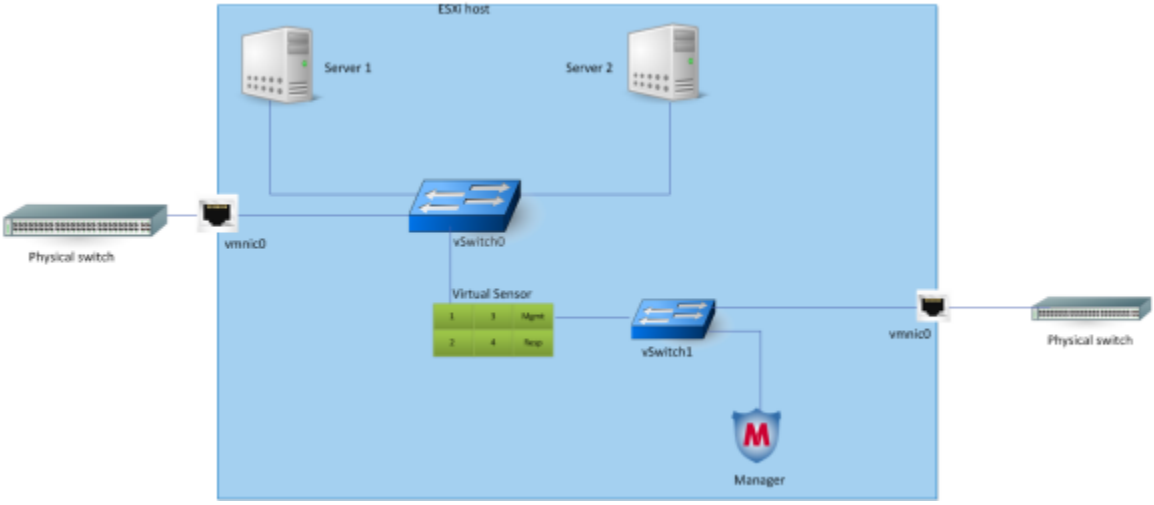

Figure 2-3 Scenario after Virtual IPS Sensor deployment

Scenario 1: High-level steps for Virtual IPS Sensor deployment

This section assumes the following for deploying the Virtual IPS Sensor for scenario 1.

- The VMware ESX server meets the requirements as discussed in Requirements for deploying Virtual IPS Sensors on page 14.
- You have the privileges on the VMware ESX server to add and modify vSwitches and port groups.
- You have installed the Virtual IPS Sensor and established trust with the Manager successfully. As an example in this scenario, the management port is connected to vSwitch1.
- As an example, this section uses the IPS-VM100 Virtual IPS Sensor to explain the deployment.
- This scenario involves only a Sensor monitoring port deployed in SPAN mode.
- This section uses only the vSphere Client for configurations on the VMware ESX.

Task

1 Modify vSwitch0 to create a switch port group in promiscuous mode.

Refer to Modify an existing standard vSwitch for a monitoring port on page 53. Subsequently, you assign this switch port group to the SPAN port.

2 Modify vSwitch0 to create a switch port group.

Subsequently, you assign this switch port group to the Sensor response port.

3 Configure the SPAN port on the Virtual IPS Sensor in the Manager.

a Click the **Devices** tab.

b Select the domain from the **Domain** drop-down list.

c In the left pane, click the **Devices** tab.

d Select the device from the **Device** drop-down list.

e Select **Setup** | **Physical Ports**.

f Double-click on monitoring port 1 and then from the **Mode** drop-down, select **SPAN or Hub**.

g Click **OK**.

h Click **Save** in the **Monitoring Port Details** panel.

i Select **Deploy Pending Changes** and in the **Deploy Pending Changes** page select **Configuration & Signature Set** for the required Virtual IPS Sensor and click **Update**.

4 Assign the switch port groups you created in steps 1 and 2 to the Sensor SPAN port and the response port respectively.

See Specify the switch port groups for monitoring ports on page 59.

5 Verify if you have deployed the Virtual IPS Sensor correctly and whether it is inspecting traffic.

Refer to Verify the deployment on page 67.

Scenario 2: Inline inspection of traffic between virtual machines

This scenario involves inspecting the traffic between virtual machines on the same VMware ESX.

Scenario description before Virtual IPS Sensor deployment

- The clients and servers belong to the same subnet (10.10.10.x).
- The clients and servers are connected to different virtual machine port groups within the same standard vSwitch (vSwitch0).
- For the sake of this discussion, assume that the clients and servers have no access from outside the VMware ESX. That is, there is no physical NIC associated with vSwitch0.

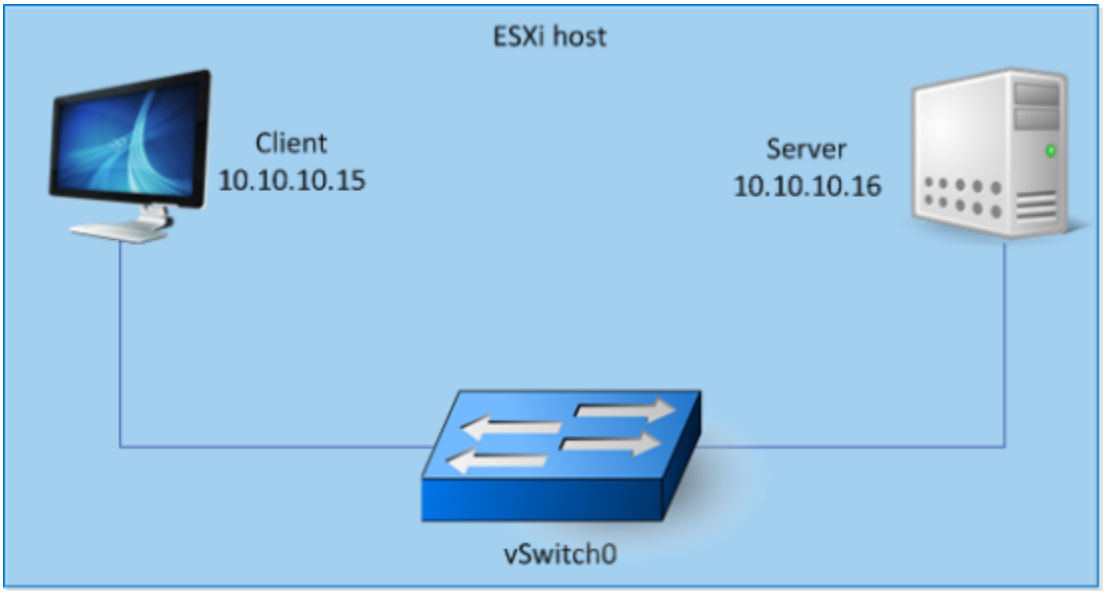

Figure 2-4 Scenario before Virtual IPS Sensor deployment

Scenario description after Virtual IPS Sensor deployment

- Two more standard vSwitches (vSwitch1 and vSwitch2) are now added.
- The Virtual IPS Sensor is deployed on the VMware ESX.
- In this scenario, the Manager is installed on a VM connected to vSwitch2.
- In this scenario, the Manager is connected to the management port of the Virtual IPS Sensor through vSwitch2. This vSwitch2 has a physical adapter vminc0. So, you can access the Manager and the Sensor from outside the VMware ESX.
- The monitoring port pair 1-2 of the Virtual IPS Sensor is inline between the client and server.

- The client and the monitoring port 1 are connected to two different port groups within vSwitch0. The port group to which the monitoring port is connected is set to promiscuous mode.
- Similarly, the server and monitoring port 2 are connected to two different port groups within vSwitch1. Any traffic from the client to the server is inspected by the monitoring port pair 1-2.

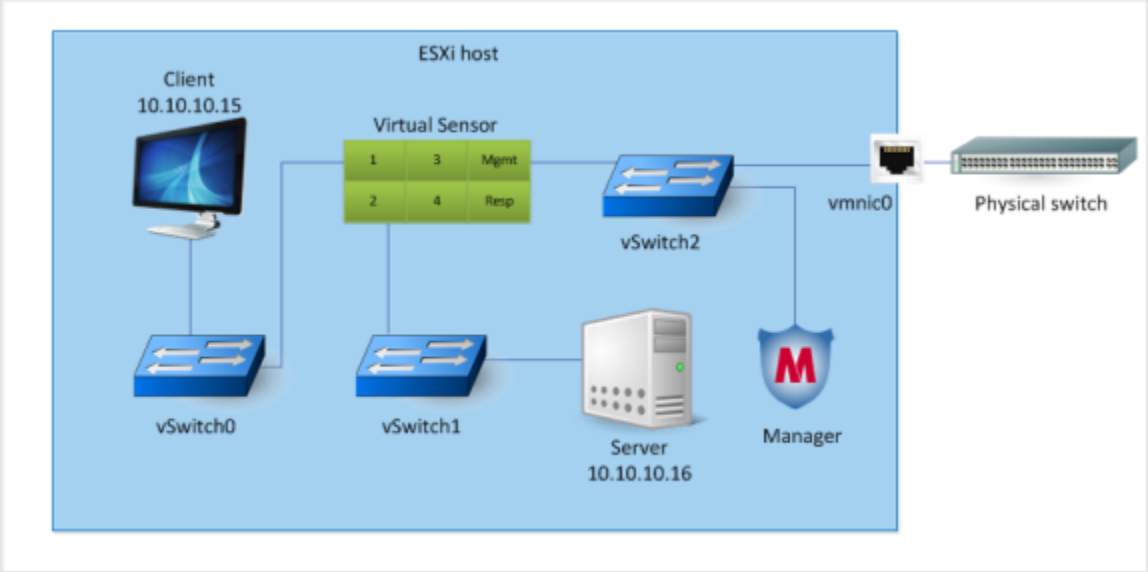

Figure 2-5 Scenario after Virtual IPS Sensor deployment

Scenario 2: High-level steps for Virtual IPS Sensor deployment

This section assumes the following for deploying the Virtual IPS Sensor for scenario 2.

- The VMware ESX server meets the requirements as discussed in Requirements for deploying Virtual IPS Sensors on page 14.
- You have the privileges on the VMware ESX server to add and modify vSwitches and port groups.
- You have installed the Virtual IPS Sensor and established trust with the Manager successfully. As an example in this scenario, the management port is connected to vSwitch2.
- As an example, this section uses the IPS-VM100 Virtual IPS Sensor to explain the deployment.
- This scenario involves only a Sensor monitoring port pair deployed in inline fail-closed mode.
- This section uses only the vSphere Client for configurations on the VMware ESX.

Task

1 Create vSwitch1 for connecting Sensor monitoring port 2 and the 10.10.10.16 server.

Refer to Create a standard vSwitch for a monitoring port on page 43.

2 Modify vSwitch0 to connect monitoring port 1.

Refer to Modify an existing standard vSwitch for a monitoring port on page 53.

3 Assign the corresponding switch port group (promiscuous mode) that you created in step 1 to monitoring port 2.

See Specify the switch port groups for monitoring ports on page 59.

4 Change the switch port group for the 10.10.10.16 server such that it is now connected to vSwitch1.

5 Assign the corresponding switch port group (promiscuous mode) that you created in step 2 to monitoring port 1.

See Specify the switch port groups for monitoring ports on page 59.

6 Verify if you have deployed the Virtual IPS Sensor correctly and whether it is inspecting traffic.

Refer to Verify the deployment on page 67.

Scenario 3: Inspection of traffic to virtual servers

This scenario involves inspecting the traffic going to and coming out of virtual servers installed on a VMware ESX. In this deployment, the Sensor monitoring port acts as a gateway to the protected servers.

Scenario description before Virtual IPS Sensor deployment

- The servers are installed on guest VMs on the VMware ESX.
- These servers are connected to a standard vSwitch, vSwitch0.
- vSwitch0 has a physical adapter, which is connected to networks outside the VMware ESX.

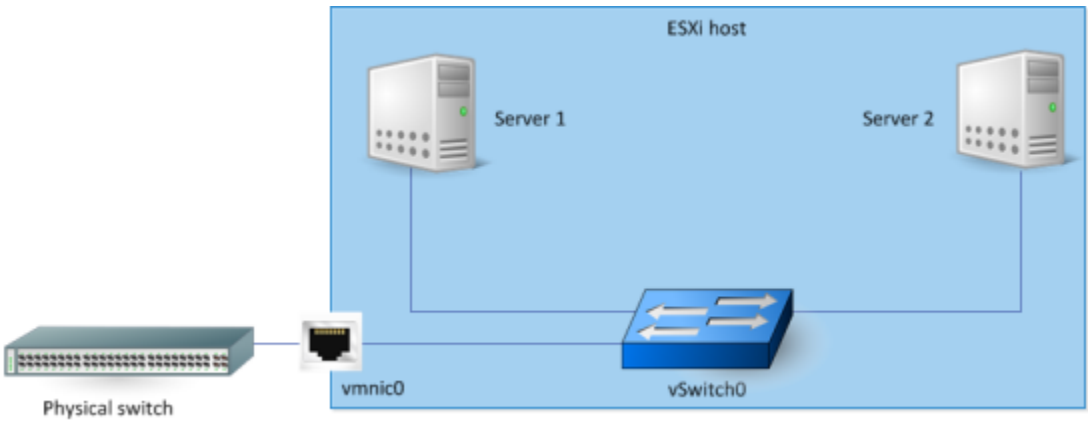

Figure 2-6 Scenario before Virtual IPS Sensor deployment

Scenario description after Virtual IPS Sensor deployment

- Two more standard vSwitches (vSwitch1 and vSwitch2) are now added.
- The Virtual IPS Sensor is deployed on the VMware ESX.
- The Manager is installed on a VM connected to vSwitch2.
- The management port of the Virtual IPS Sensor is connected to vSwitch2. This virtual switch has a physical adapter. So, you can access the Manager and the Sensor from outside the VMware ESX.
- The monitoring port pair 1-2 of the Virtual IPS Sensor is inline between external network through vmnic0 and the server farm on the VMware ESX.

- The servers and the monitoring port 1 are connected to two different port groups in vSwitch0. The port group to which the monitoring port is connected is set to promiscuous mode.
- Monitoring port 2 is connected to vSwitch1, which is in turn connected to external network through vmnic0. Therefore, any traffic to the servers from the outside network is inspected by the port pair 1-2.

Note that the monitoring port 1 is connected to a promiscuous switch port group on vSwitch0. Therefore, the Sensor will inspect traffic between Server 1 and Server 2 as well though it is not inline. Effectively this is as if monitoring port 1 is in SPAN mode. To avoid the Sensor from inspecting the traffic between the servers, define ACLs on the Sensor accordingly.

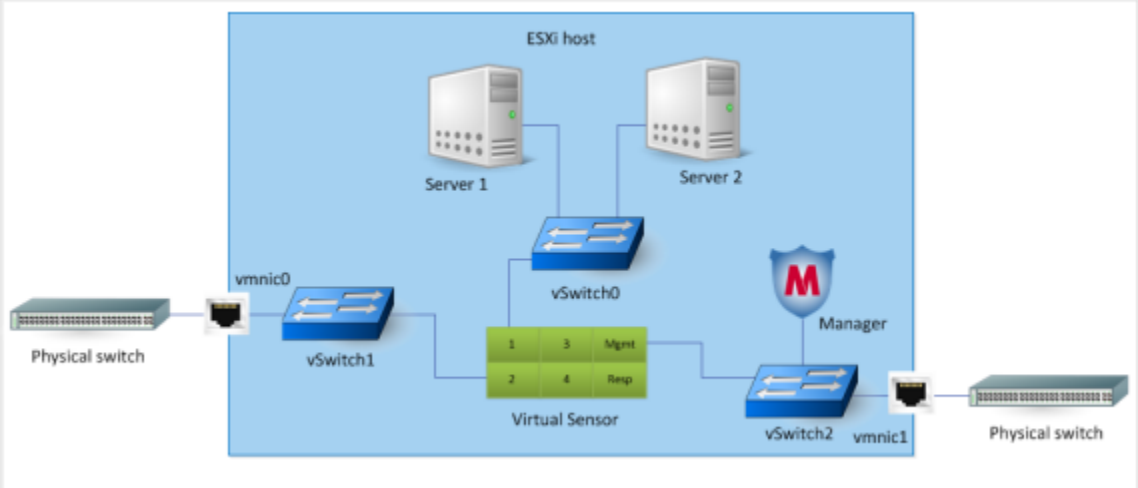

Figure 2-7 Scenario after Virtual IPS Sensor deployment

Scenario 3: High-level steps for Virtual IPS Sensor deployment

This section assumes the following for deploying the Virtual IPS Sensor for scenario 3.

- The VMware ESX server meets the requirements as discussed in Requirements for deploying Virtual IPS Sensors on page 14.
- You have the privileges on the VMware ESX server to add and modify vSwitches and port groups.
- You have installed the Virtual IPS Sensor and established trust with the Manager successfully. As an example in this scenario, the management port is connected to vSwitch2.
- As an example, this section uses the IPS-VM100 Virtual IPS Sensor to explain the deployment.
- This scenario involves only a Sensor monitoring port pair deployed in inline fail-closed mode.
- This section uses only the vSphere Client for configurations on the VMware ESX.

Task

1 Create a standard vSwitch for connecting Sensor monitoring port 2 with the external network through vmnic0. For this scenario, consider it is vSwitch1.

Refer to Create a standard vSwitch for a monitoring port on page 43.

Make sure this vSwitch has a physical adapter and that this is connected to the corresponding external switch.

2 Modify vSwitch0 to connect monitoring port 1 and the virtual servers.

For this scenario, you need not change the switch port group for the virtual servers. Create a new switch port group for monitoring port 1. Refer to Modify an existing standard vSwitch for a monitoring port on page 53.

3 Assign the corresponding switch port group (promiscuous mode) to the monitoring ports.

See Specify the switch port groups for monitoring ports on page 59.

- The switch port group that you created in step 1 (vSwitch1) must be assigned to monitoring port 2.
- The switch port group that you created in step 2 (vSwitch0) must be assigned to monitoring port 1.

4 Verify if you have deployed the Virtual IPS Sensor correctly and whether it is inspecting traffic.

Refer to Verify the deployment on page 67.

Scenario 4: Inspection of traffic between physical machines

In addition to virtual networks, you can use the Virtual IPS Sensor to protect physical networks as well. This is typically used in very large networks, where space might be a constraint to install multiple physical Sensors.

Scenario description before Virtual IPS Sensor deployment

- The clients and servers are connected to different physical switches.
- For the sake of simplicity, assume that the client and servers are on the same VLAN.
- The two switches are connected through their trunk ports. Since all the machines are assumed to be on the same VLAN, no routing is required.

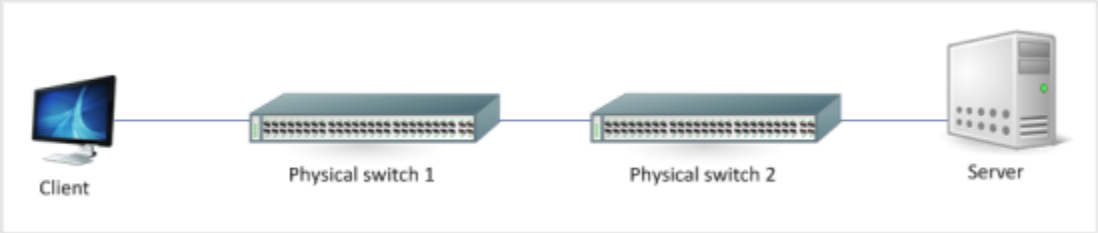

Figure 2-8 Scenario before Virtual IPS Sensor deployment

Scenario description after Virtual IPS Sensor deployment

- The trunk ports of the two switches are connected two physical NICs on the VMware ESX.
- These NICs are connected to two different vSwitches, which are in turn connected to monitoring port pair 1-2.

- vSwitch2 is used to connect the management port, which in turn is connected to a Manager on a physical machine.
- The monitoring port pair 1-2 of the Virtual IPS Sensor is now inline between the client and the server.

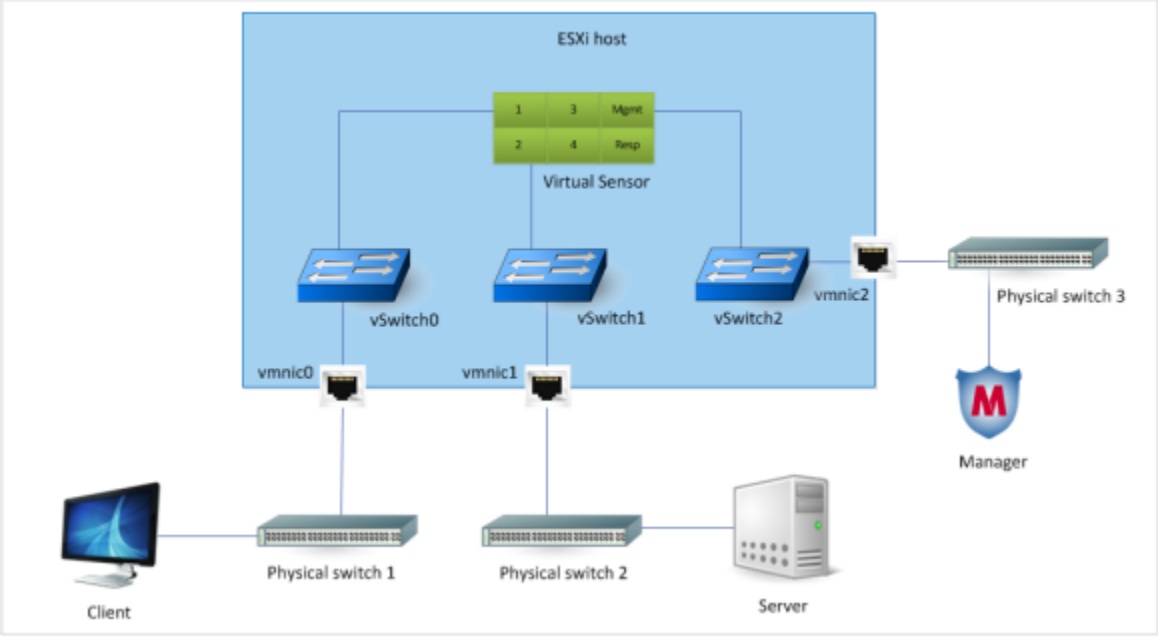

Figure 2-9 Scenario after Virtual IPS Sensor deployment

Scenario 4: High-level steps for Virtual IPS Sensor deployment

This section assumes the following for deploying the Virtual IPS Sensor for scenario 4.

- The VMware ESX server meets the requirements as discussed in Requirements for deploying Virtual IPS Sensors on page 14.
- You have the privileges on the VMware ESX server to add and modify vSwitches and port groups.
- You have installed the Virtual IPS Sensor and established trust with the Manager successfully. As an example in this scenario, the management port is connected to vSwitch2.
- As an example, this section uses the IPS-VM100 Virtual IPS Sensor to explain the deployment.
- This scenario involves only a Sensor monitoring port pair deployed in inline fail-closed mode.
- This section uses only the vSphere Client for configurations on the VMware ESX.

Task

1 Connect the trunk port of Physical switch 1 and Physical switch 2 to two different physical NICs on the VMware ESX.

2 Create two standard vSwitches for connecting Sensor monitoring ports 1 and 2.

Refer to Create a standard vSwitch for a monitoring port on page 43. Both these switches need physical adapters. For this scenario, vSwitch0 must be assigned a physical adapter that is connected to Physical switch 1. Similarly, vSwitch1 must be assigned a physical adapter that is connected to Physical switch 2.

3 Create a switch port group in vSwitch0, which corresponds to the trunk port on Physical switch 1.

a Log on to the VMware ESX as the root user in VMware vSphere Web Client.

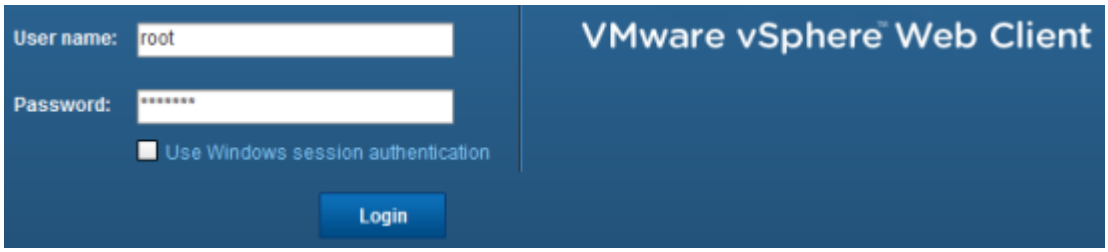

b In the vSphere **Home** tab, select **Hosts and Clusters.**

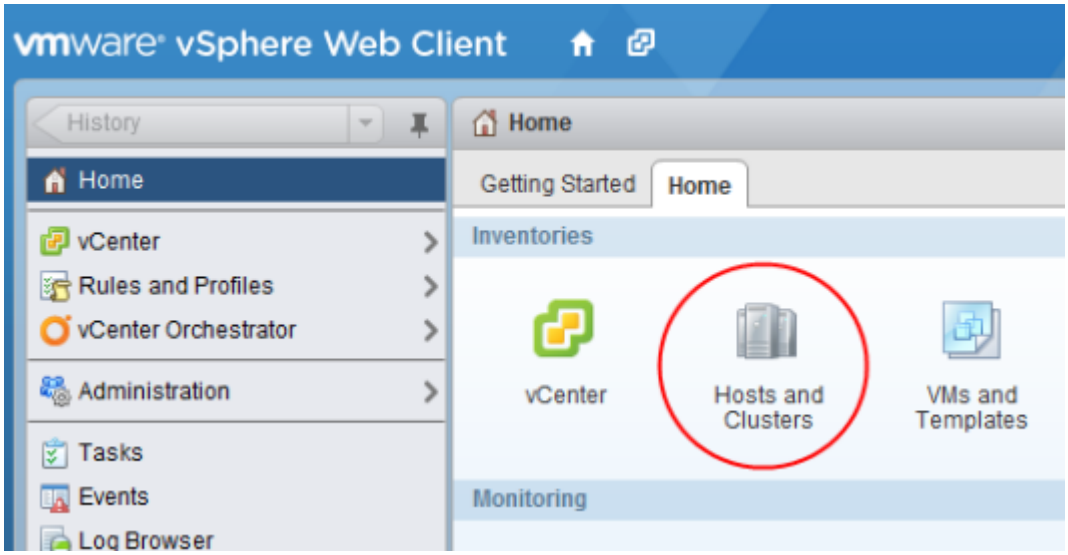

c Select the required VMware ESX server and select **Manage** | **Networking** | **Virtual switches** | **vSwitch0.**

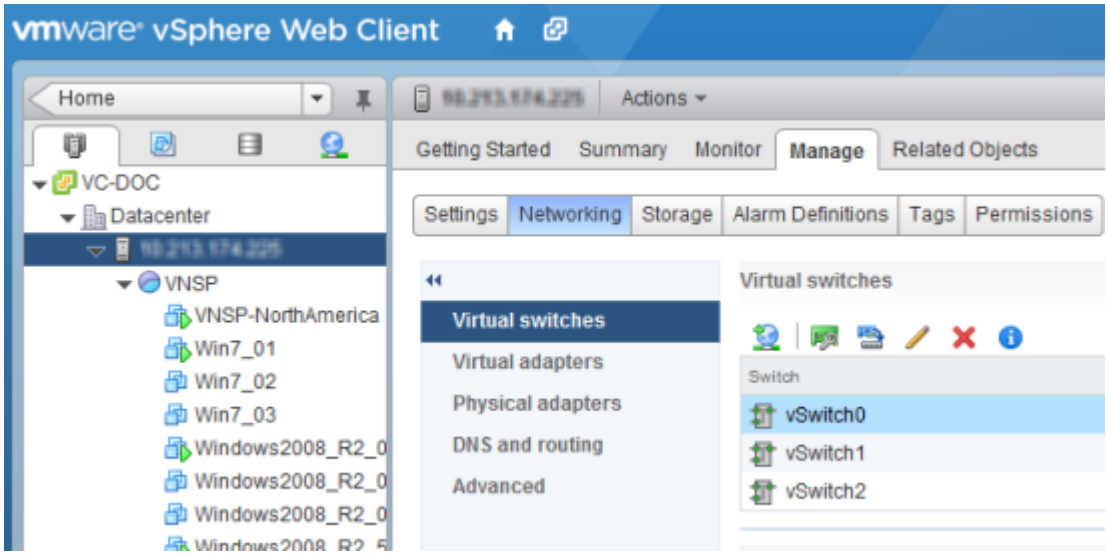

d Click on **Add host networking** icon for vSwitch0.

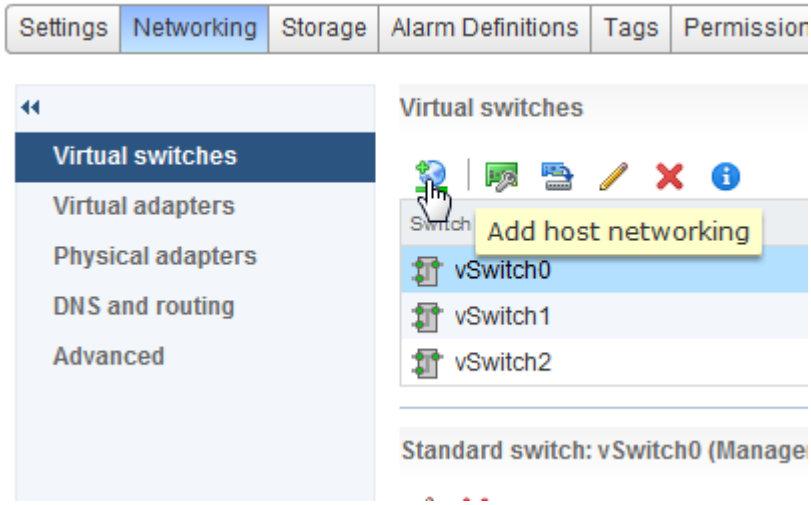

e For **Select target device**, select **Select an existing standard switch** and make sure vSwitch0 is selected.

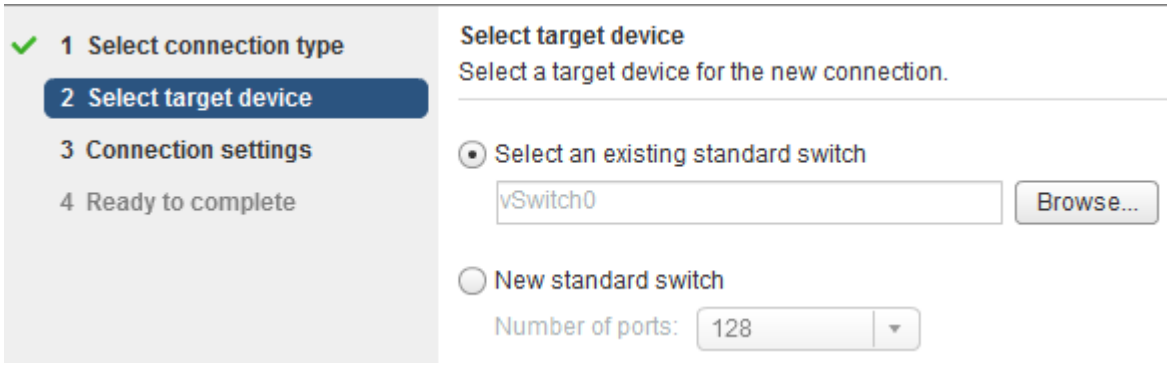

f In the **Network Label** field, enter a name.

For example, enter *Physical Client Port*.

g In the **VLAN ID (Optional)** field, select **All (4095)** because this switch port group corresponds to the trunk port of a physical switch.

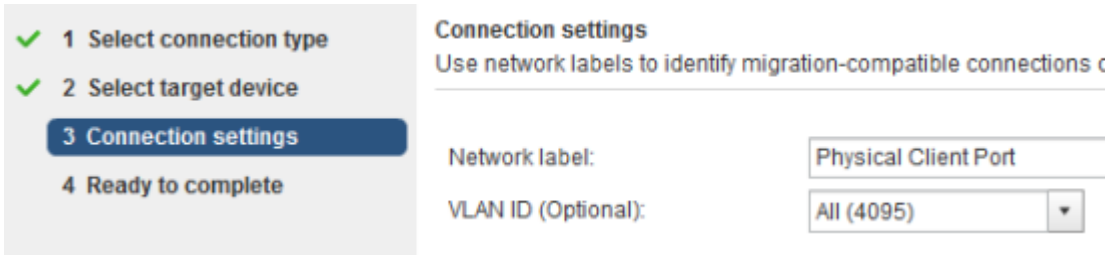

h Click **Next** and then **Finish**.

i Under **Standard switch: vSwitch0 (VM Network)**, click on the switch port group that you created.

In this example, it is Physical Client Port.

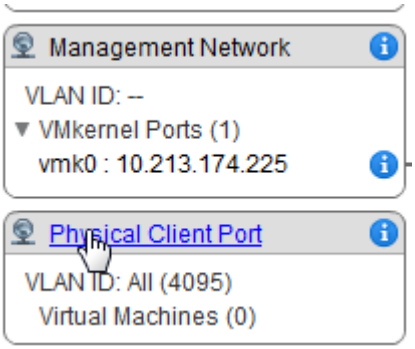

j With the switch port group selected, click on the **Edit Settings** icon for the switch port group.

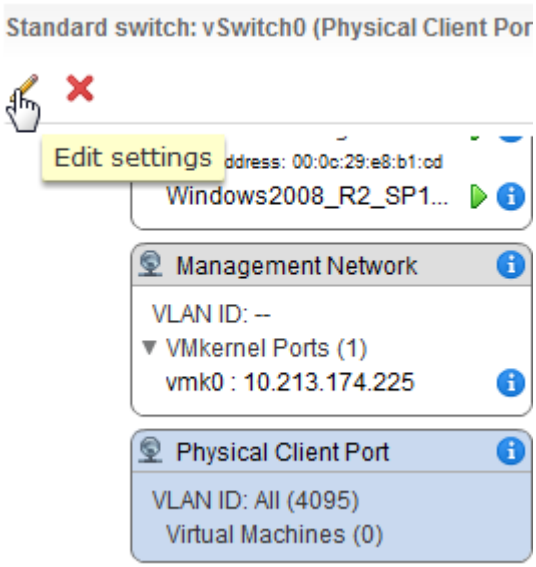

k In the **Edit Settings** dialog, select the **Security** tab and make sure the fields are set with the following values and click **OK**.

- **Promiscuous mode —** Accept. This is set to accept because traffic related to all the hosts connected to Physical switch 1 is involved.
- **MAC Address Changes —** Reject.
- **Forged Transmits —** Accept.

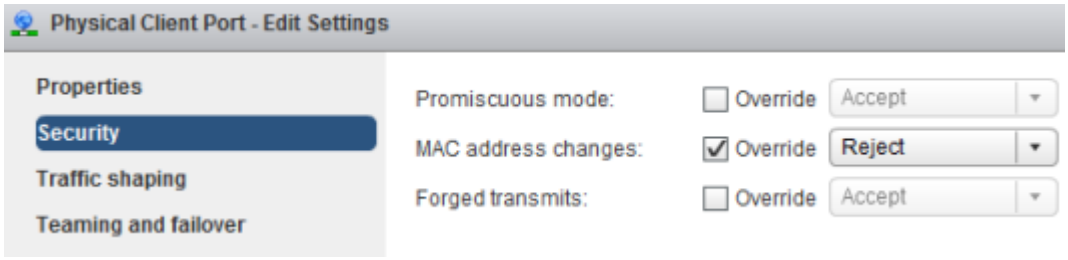

4 Use the previous step to create a similar switch port group in vSwitch1.

This corresponds to the trunk port on Physical switch 2.

5 Assign the switch port group that you created in step 3 to monitoring port 1.

See Specify the switch port groups for monitoring ports on page 59.

6 Assign the switch port group that you created in step 4 to monitoring port 2.

See Specify the switch port groups for monitoring ports on page 59.

Make sure the monitoring port 1 is connected to the corresponding port group (VNSP100-PG-Port1) on vSwitch0 and the monitoring port 2 is connected to the corresponding port group (VNSP100-PG-Port2) on vSwitch1. Both these port groups must have their VLAN ID as All (4095). This is required since the monitoring ports are connected to trunk ports of the physical switches.

7 Verify if you have deployed the Virtual IPS Sensor correctly and whether it is inspecting traffic.

Refer to Verify the deployment on page 67.

Create a standard vSwitch for a monitoring port

Before you begin

If you require external access to VMs connected to this switch, you will be required to connect an additional physical NIC on the VMware ESX to a physical switch.

You connect monitoring ports to standard vSwitches. When you create a standard vSwitch, VMware ESX creates a default port group for this vSwitch. Each monitoring port in a Virtual IPS Sensor must be connected to different port groups.

Task

1 Optionally, connect an additional physical NIC on the VMware ESX to the adjacent physical switch.

In scenario 4, for example, you must connect an additional NIC to the corresponding physical switch.

2 Log on to the VMware ESX as the root user in VMware vSphere Web Client.

3 In the vSphere **Home** tab, select **Hosts and Clusters.**

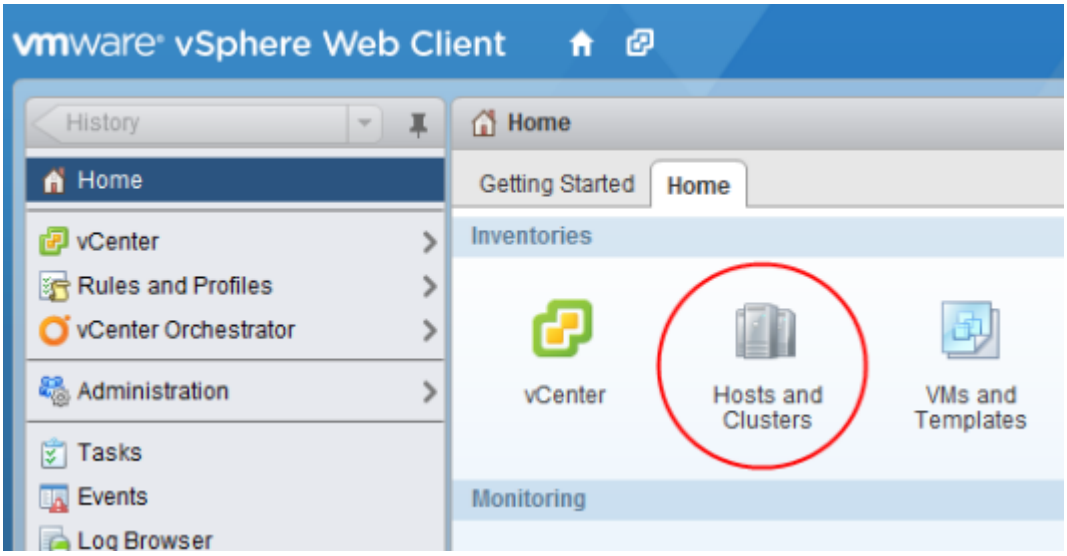

4 Select the required VMware ESX server and select **Manage** | **Networking** | **Virtual switches.**

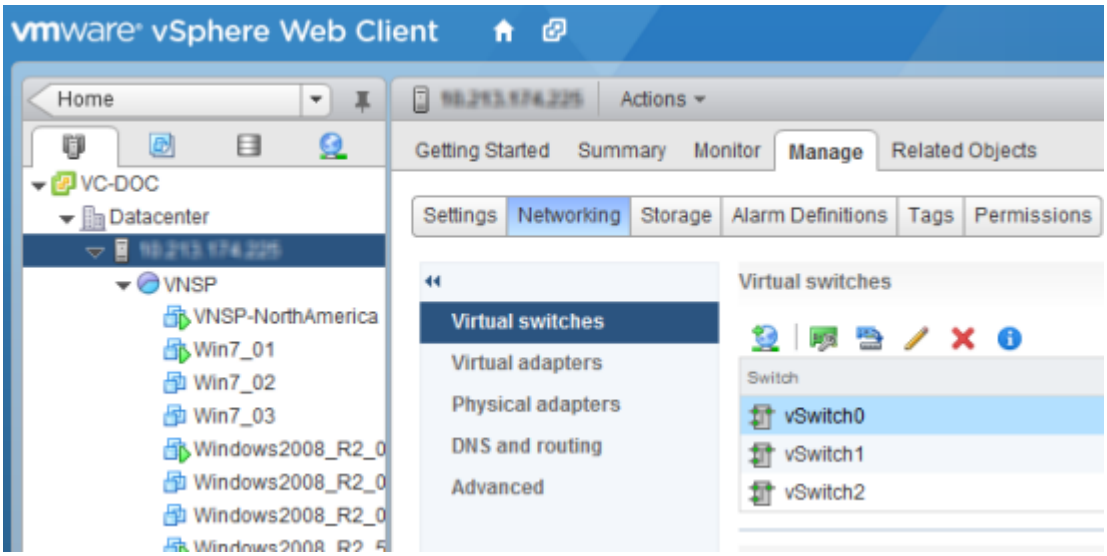

5 Click on the **Add host networking** icon.

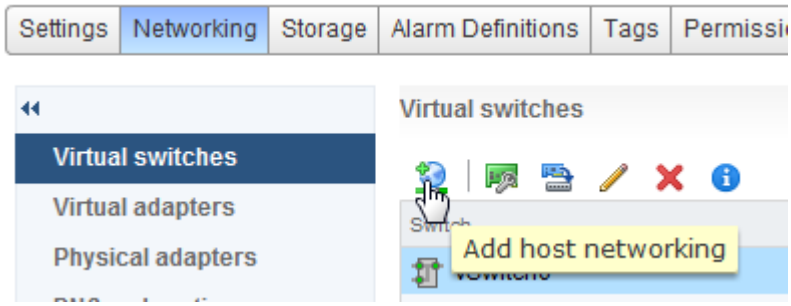

6 In the **Select connection type** section, select the required connection type and click and click **Next**.

Selecting the connection type depends on your network design and requirements. If the VMs that will be connected to this switch do not need access outside the VMware ESX then you might select **Virtual Machine Port Group for a Standard Switch**. However, for requirements as in scenario 4, it is mandatory to select **Physical Network Adapter**. Consider that you now select **Virtual Machine Port Group for a Standard Switch**.

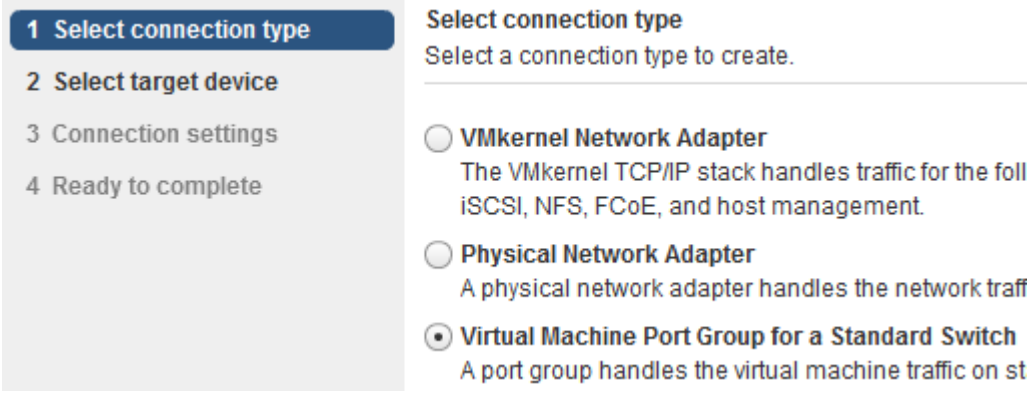

7 For **Select target device**, select **New standard switch**, the required number of ports, and then click **Next**.

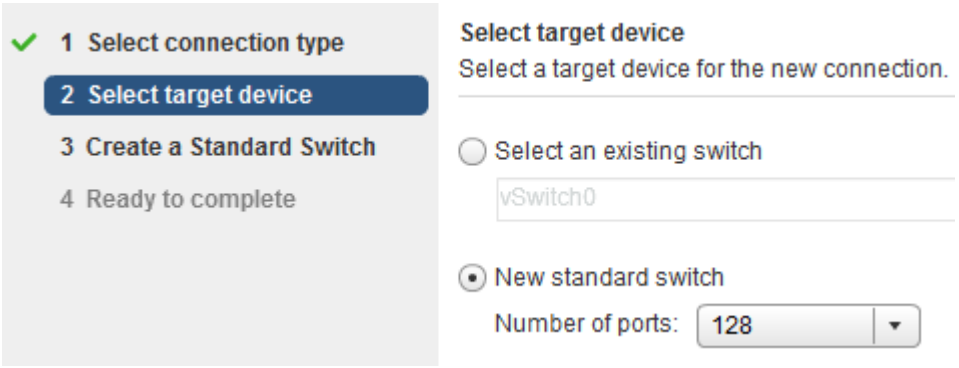

8 Based on your network requirements, click on the **Add adapters** icon and select the corresponding physical network adapter. Then click **Next**.

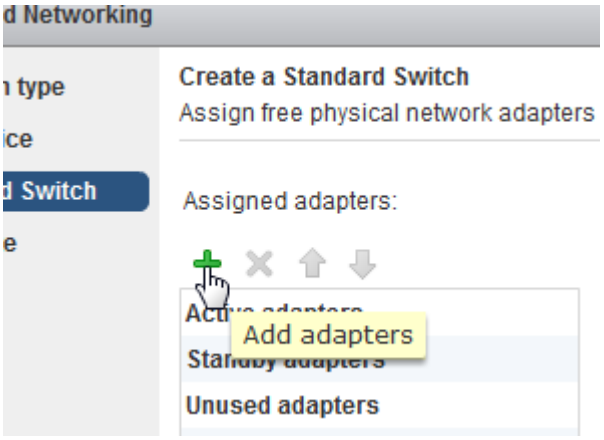

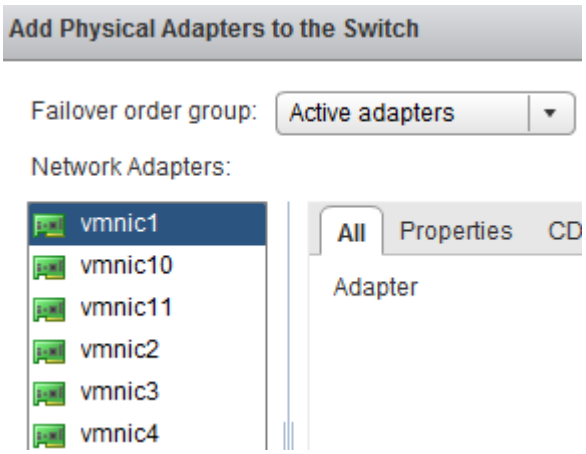

If you select an adapter, make sure that a physical NIC corresponding to the network adapter you selected is connected to the network.

9 In the **Network Label** field, enter the required name for the default port group that the wizard creates for the switch.

You can modify **Network Label** even later. For easier management, name this as *VNSP100-PG-Port1* for example.

10 In **VLAN ID**, select **All (4095)** and select **Next**.

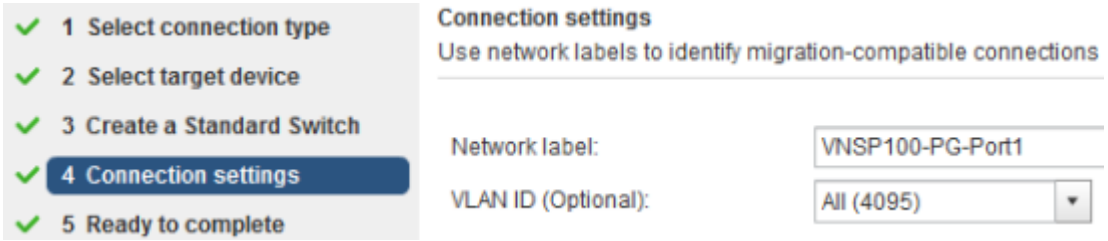

11 Click **Finish.**

This vSwitch is now listed under **Virtual Switches** in the **Networking** tab.

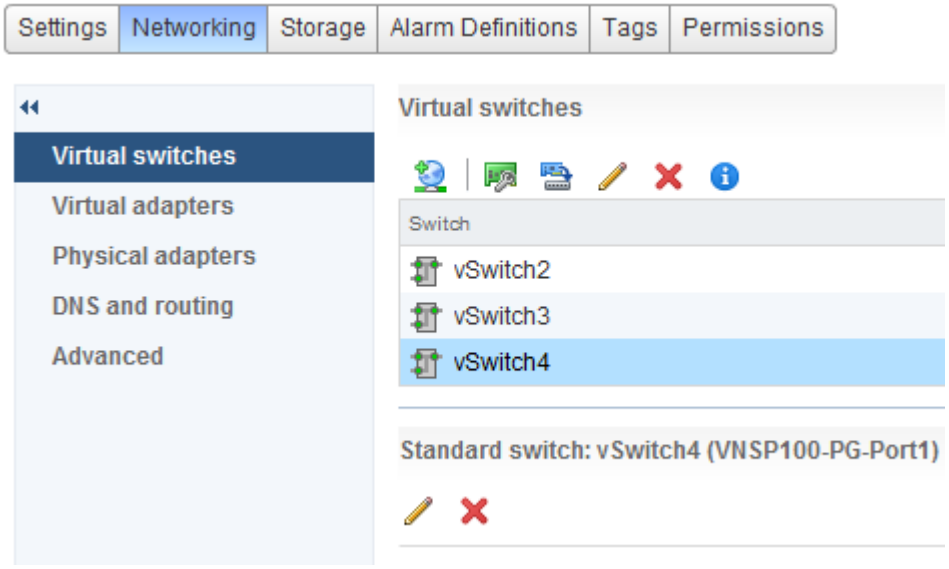

12 Select the vSwitch that you created, move the mouse over its physical adapter, and click on it.

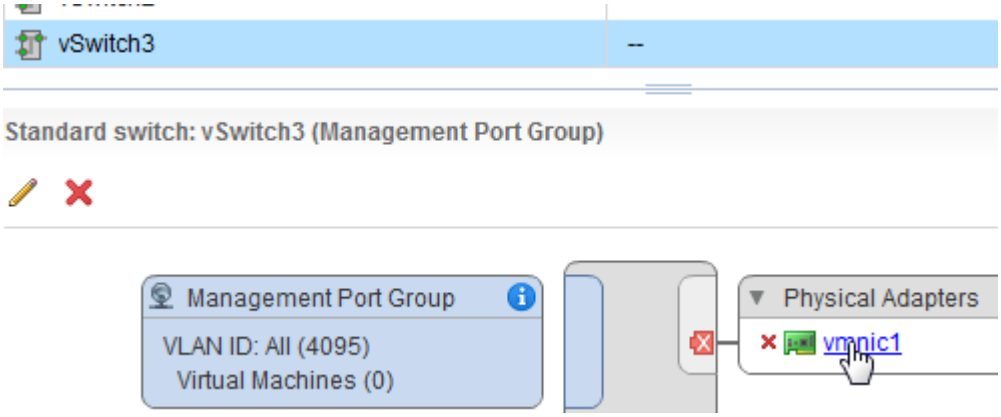

13 Click on the **Edit adapter speed** icon.

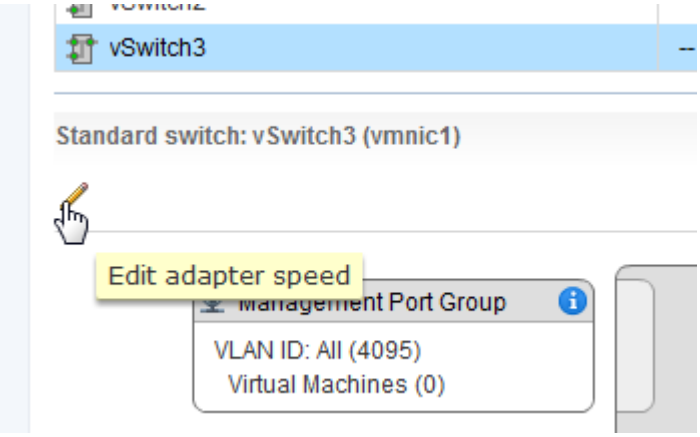

14 Verify if the **Configured Speed, Duplex** is set to **Auto negotiate.**

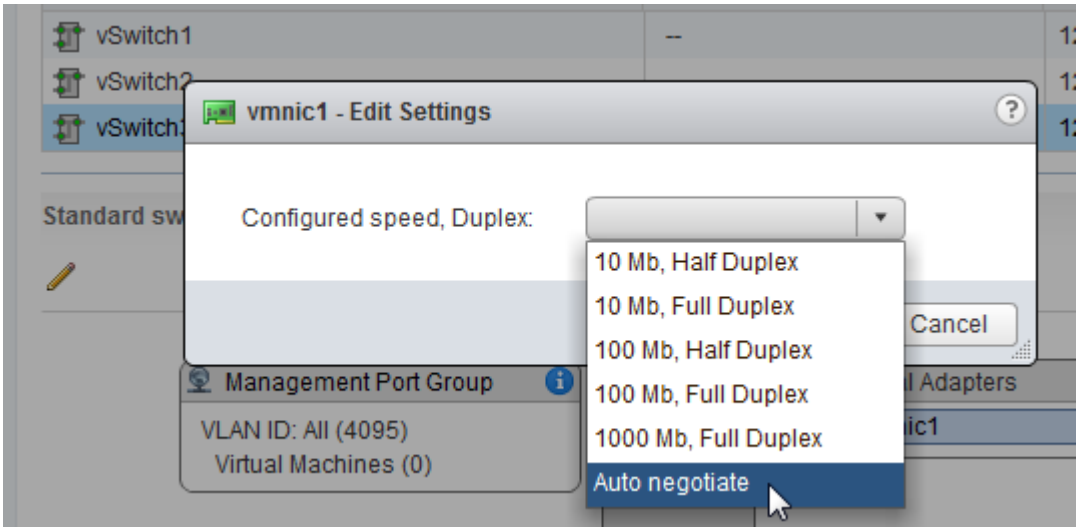

For other property values, you can leave them with the default values.

15 Modify the security properties of the vSwitch.

a Select the vSwitch and click on the **Edit settings** icon.

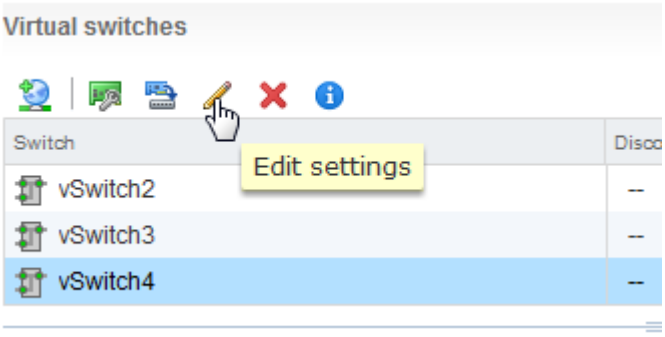

b Select **Security** and make sure the fields are set to the values mentioned and then click **OK**.

- **Promiscuous mode —** Reject.
- **MAC Address Changes —** Reject.
- **Forged Transmits —** Accept.

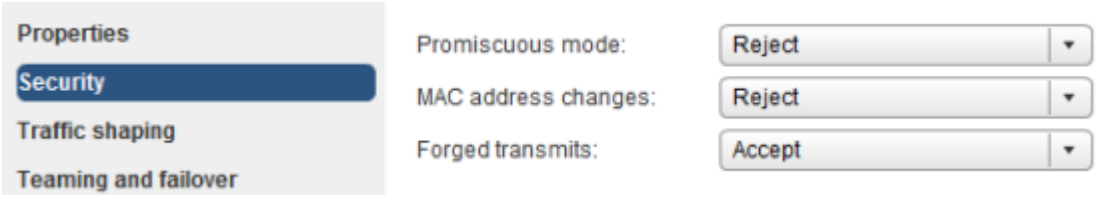

16 Modify the security settings of the default port group.

It is assumed that you will use this port group to connect the monitoring port of a Sensor.

a Move the mouse over the default port group and click on it.

The switch port group is now selected.

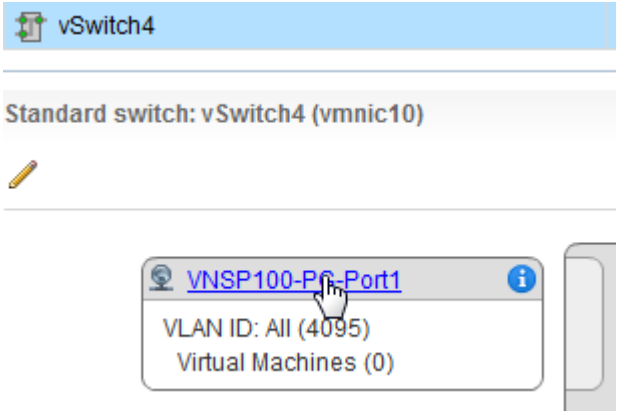

b Click on the **Edit settings** icon for the switch port group.

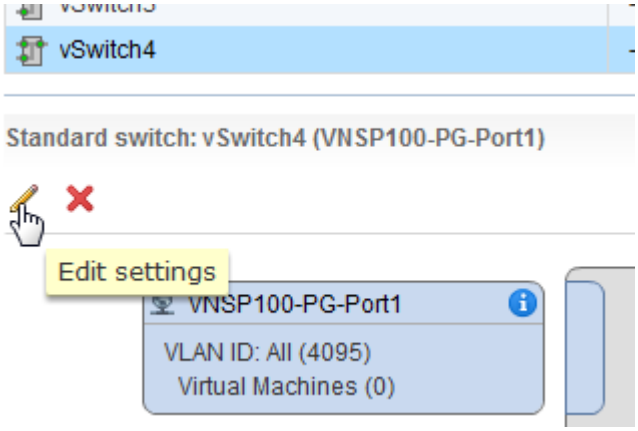

c Click **Properties** and modify the Network Label if required.

d Make sure **VLAN ID** is set to **All (4095)**.

This switch port group must receive all VLAN traffic similar to a trunk port.

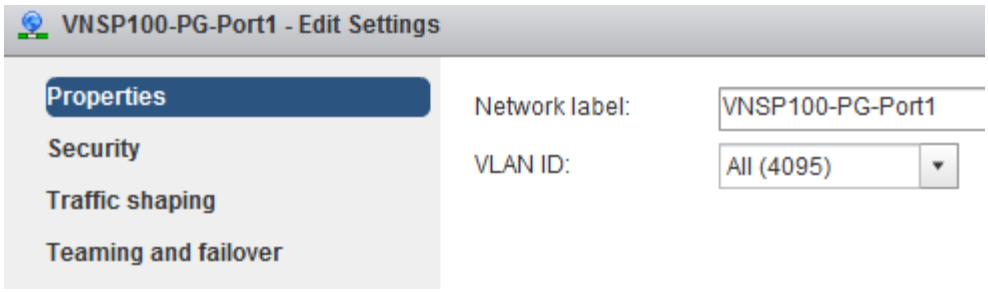

e Click **Security**, select the **Override** check box next to **Promiscuous Mode** and then select **Accept** from the drop-down.

This is mandatory for the port group that you will use for any Sensor monitoring port.

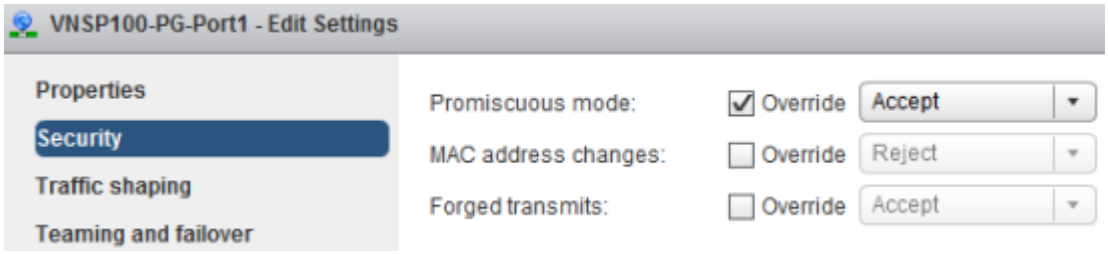

17 Create a new port group for the other VMs in this vSwitch.

For example, in scenario 2, this is the port group that you will use for the 10.10.10.16 server.

Skip this step for scenario 4.

a Select the corresponding vSwitch and click on the **Add host networking** icon.

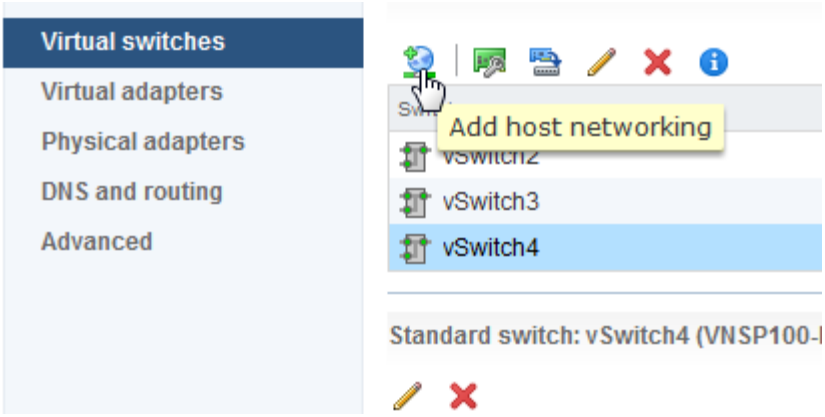

b In the **Select connection type** step, select **Virtual Machine Port Group for a Standard Switch** and then click **Next.**

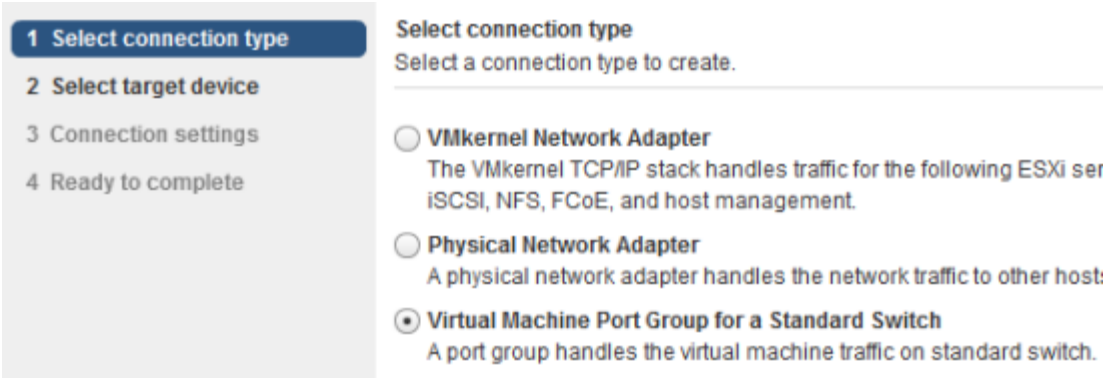

c In the **Select target device** step, select **Select an existing standard switch** and make sure the vSwitch that you created is selected. Then click **Next.**

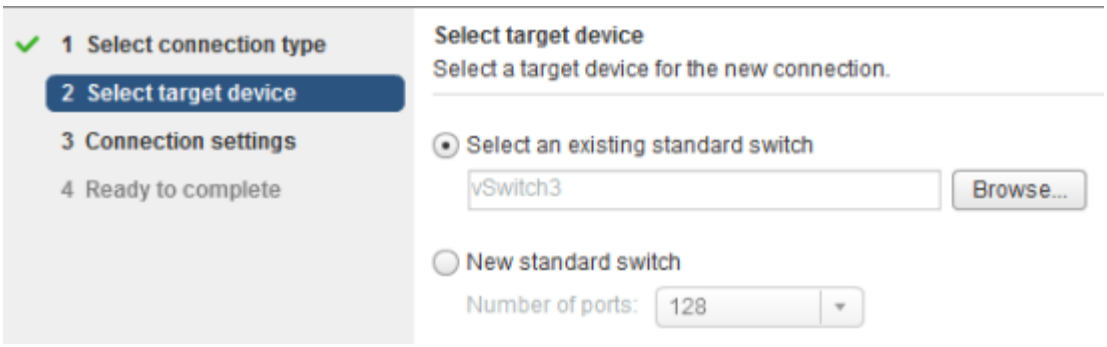

d In the **Network Label** field, enter the required name for the default port group that the wizard creates for the switch.

You can modify **Network Label** later. For easier management, you can name it as *Server Port*.

e In the **VLAN ID (Optional)** field, you can specify the required VLAN. For scenario 1, for example, select **None (0)** and click **Next** and then **Finish**.

None (0) means that the traffic is not tagged with a VLAN.

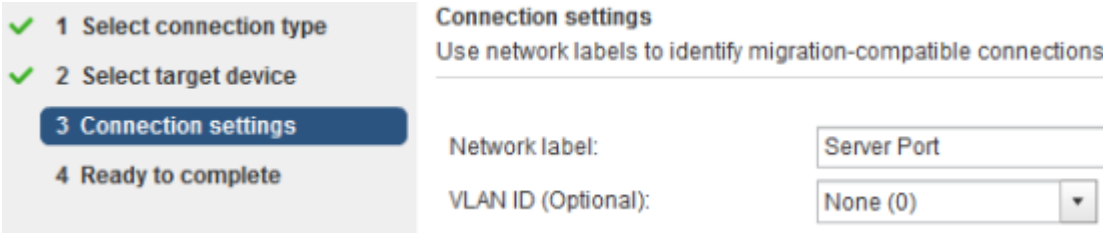

f Move the mouse over the switch port group and click on it.

The switch port group is now selected.

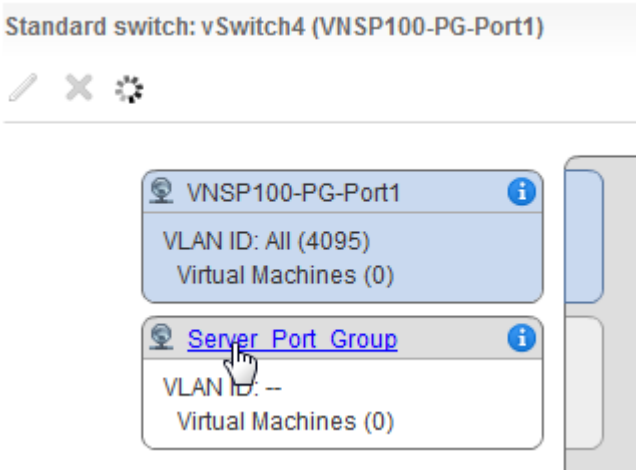

g Click on the **Edit settings** icon for the switch port group.

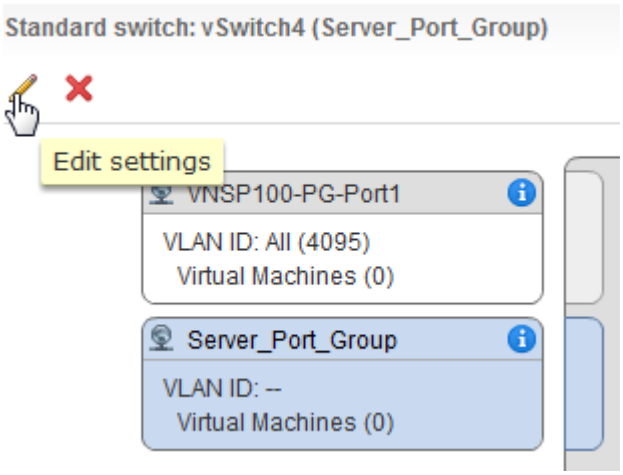

h In the **Security** tab, make sure the fields are set with the following values and click **OK**.

- **Promiscuous mode —** Reject.
- **MAC Address Changes —** Reject.
- **Forged Transmits —** Accept.

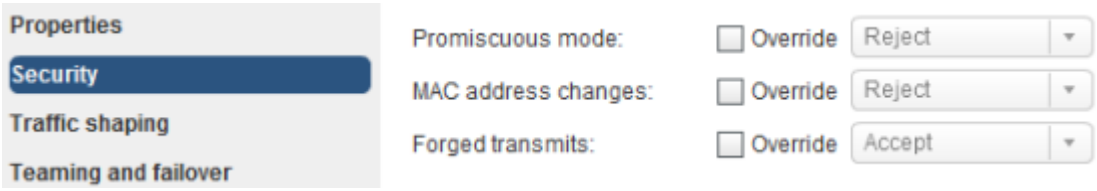

18 Click **OK** to close the **Edit settings** dialog.

Modify an existing standard vSwitch for a monitoring port

If you plan to connect a monitoring port to an existing vSwitch, you might have to modify some of the configuration. For example, in scenario 2, you need to modify virtual switch 1 to connect it to monitoring port 1.

Task

1 Log on to the VMware ESX as the root user in VMware vSphere Web Client.

2 In the vSphere **Home** tab, select **Hosts and Clusters.**

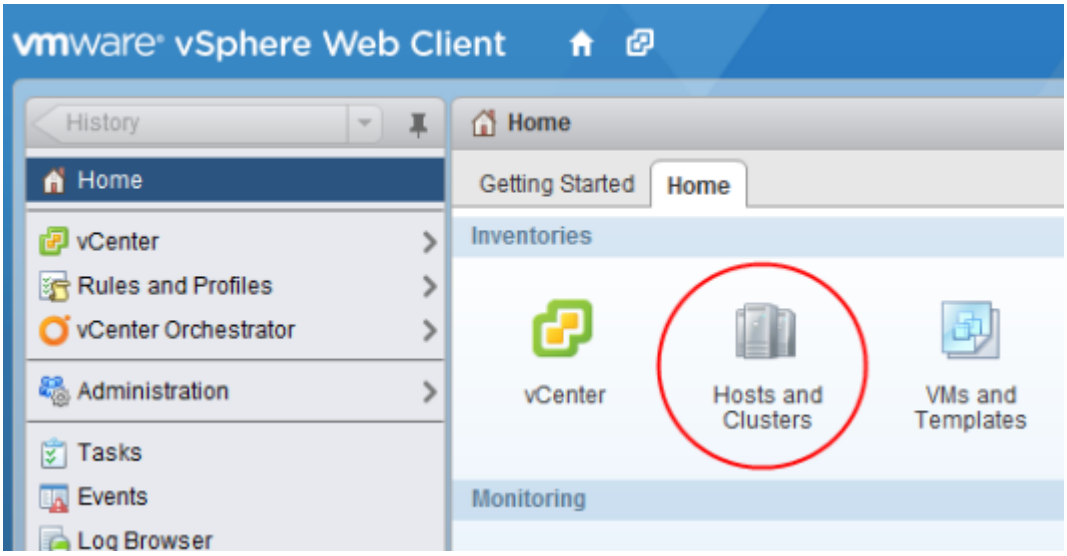

If the page takes time to render, click on the page-refresh icon at the top.

3 Select the required VMware ESX server and select **Manage** | **Networking** | **Virtual switches.**

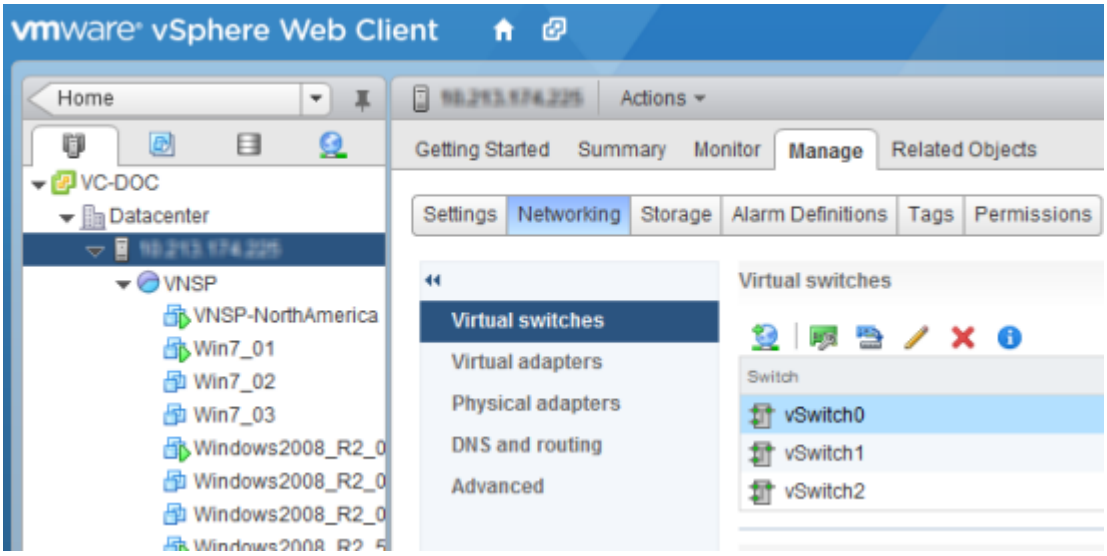

4 Select the required vSwitch and click on its **Edit settings** icon.

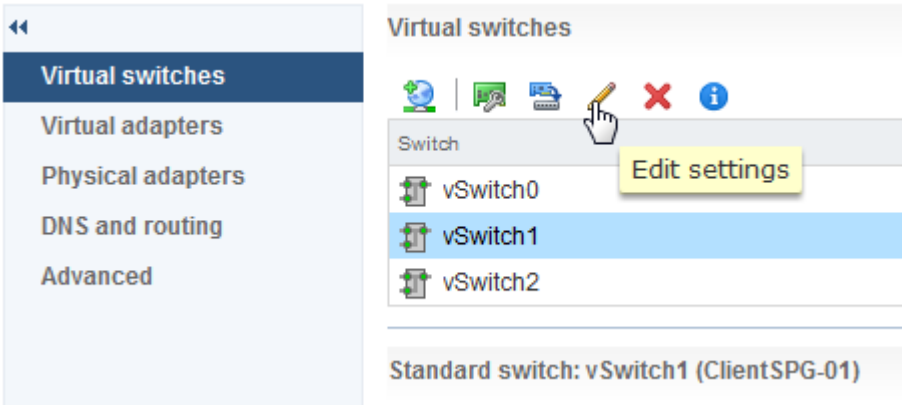

5 Select **Security** and make sure the fields are set to the values mentioned and then click **OK**.

- **Promiscuous mode —** Reject.
- **MAC Address Changes —** Reject.
- **Forged Transmits —** Accept.

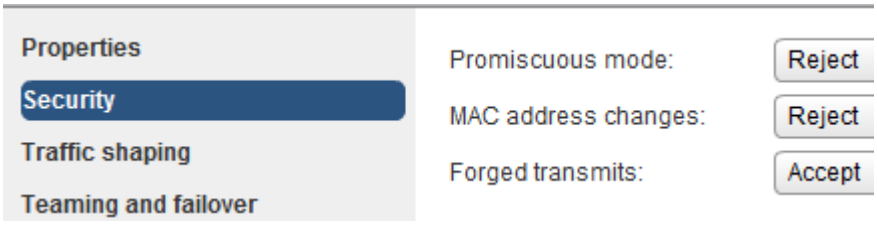

6 Modify the port group that you are using for the VMs in this switch.

For example, in scenario 2, this is the port group that you use for the 10.10.10.15 client.

This step might be relevant only for scenario 2.

a Move the mouse over the required port group and click on it.

The switch port group is now selected.

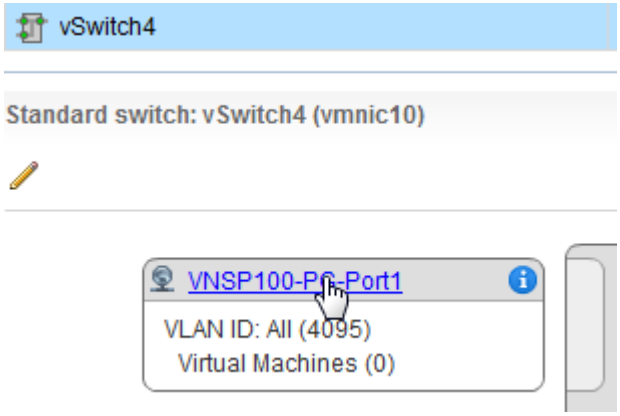

b Click on the **Edit settings** icon for the switch port group.

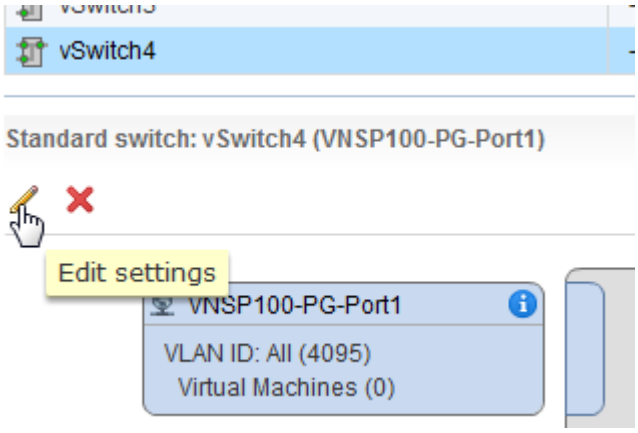

c Click **Properties** and modify the Network Label if required.

For example, change it to Client Port.

d In the VLAN ID (Optional) field, you can specify the required VLAN. For the scenarios in this section, select None (0).

e Click **Security** and make sure the fields are set with the following values and click **OK**.

- **Promiscuous mode —** Reject.
- **MAC Address Changes —** Reject.
- **Forged Transmits —** Accept.

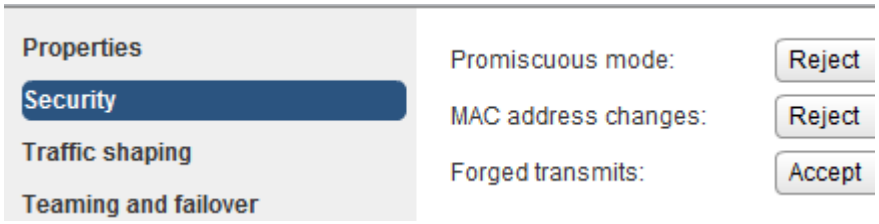

7 Create a port group to connect the monitoring port of a Sensor.

a Select the corresponding vSwitch and click on the **Add host networking** icon.

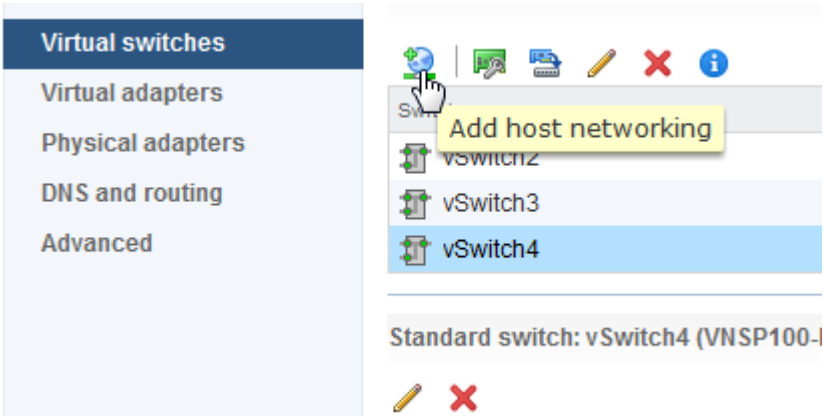

b In the **Select connection type** step, select **Virtual Machine Port Group for a Standard Switch** and then click **Next**.

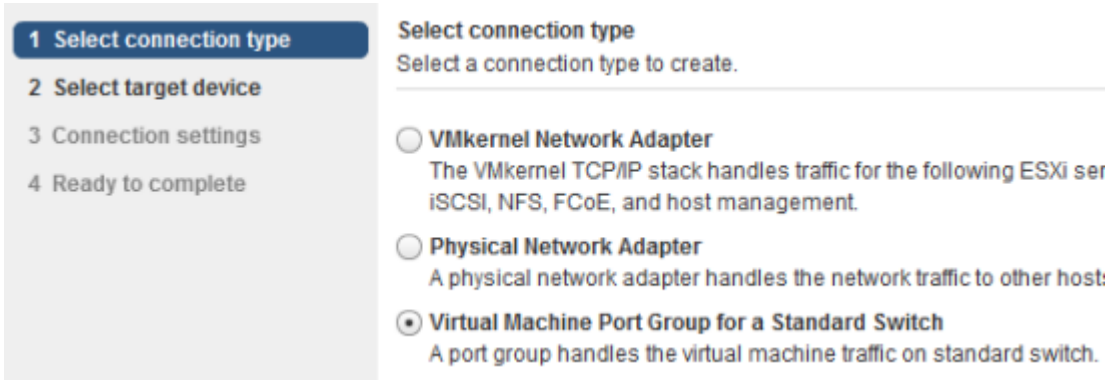

c In the **Select target device** step, select **Select an existing standard switch** and make sure the vSwitch that you created is selected. Then click **Next**.

d In the **Network Label** field, enter the required name

For example, enter VNSP Client Port.

e Make sure VLAN ID is set to All (4095), click **Next** and then **Finish**.

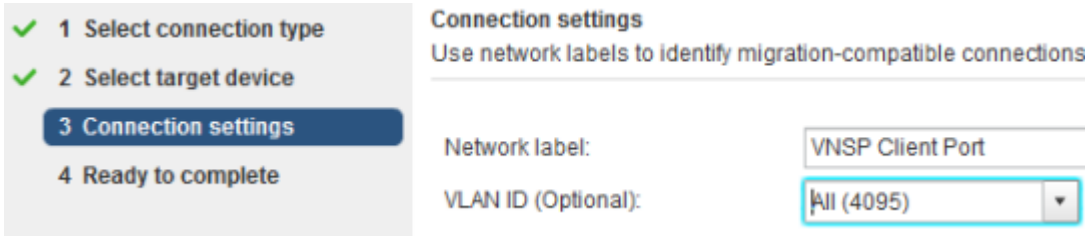

f Move the mouse over the switch port group and click on it.

The switch port group is now selected.

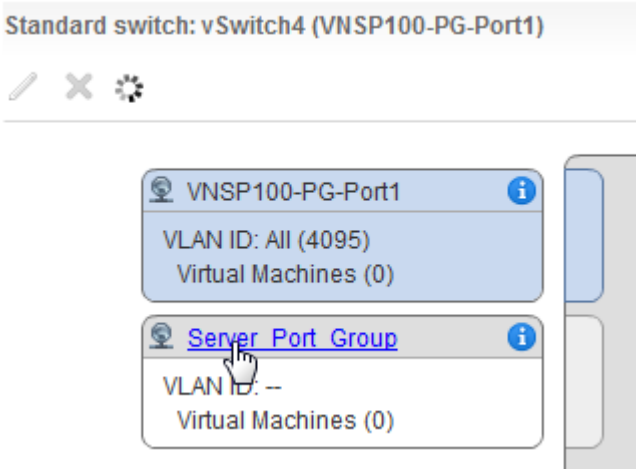

g Click on the **Edit settings** icon for the switch port group.

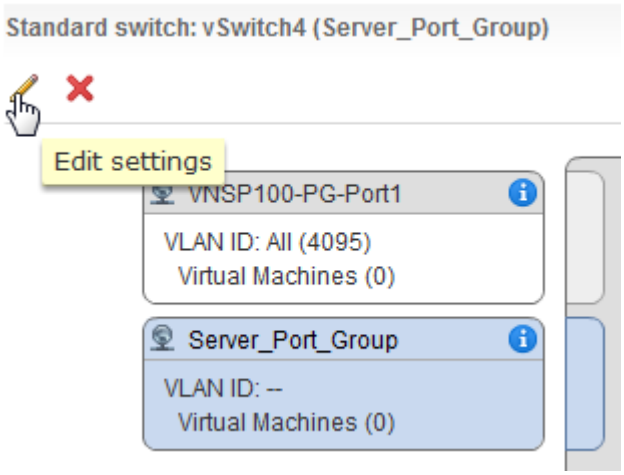

h In the Security tab, select **Override** next to **Promiscuous mode** and then select **Accept** from the drop-down. Click **OK** when done.

This is mandatory for the port group that you will use for any Sensor monitoring port.

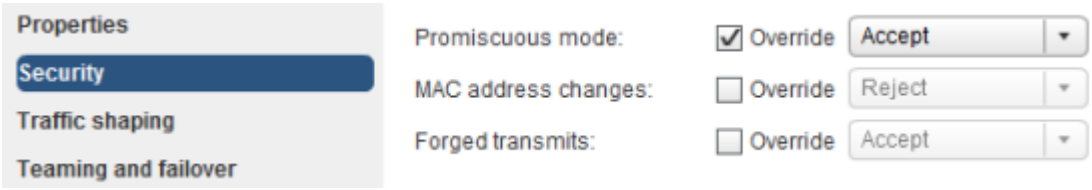

Specify the switch port groups for monitoring ports

When you install the Virtual IPS Sensor, you select the switch port group for the required Sensor ports. No two port must be connected to the same switch port group to prevent loopback. As a precaution, the monitoring and response ports are disconnected by default.

As a good practice, assess the vSwitches and the switch port groups that you would require. Then you can create them in your VMware ESX before you begin installing the Virtual IPS Sensor. You can also create the dummy switch port groups for the ports that you do not plan to deploy.

The first two network adapters correspond to the management port and the response port respectively. The remaining adapters correspond to the monitoring ports. The following graphic shows the mapping between network adapters and the Sensor ports for IPS-VM100.

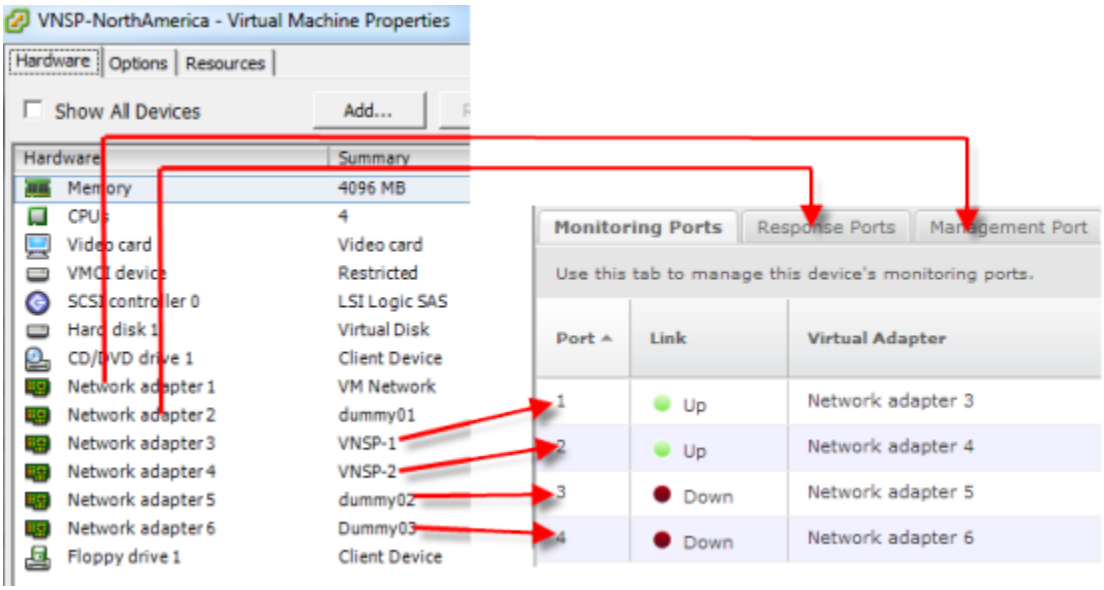

Figure 2-10 Network adapter to Sensor port mapping

The left side is the Virtual Machine Properties dialog that you can access in the vSphere client for a Virtual IPS Sensor. To access the Virtual Machine Properties dialog, select the Virtual IPS Sensor in the vSphere client and then select **Edit virtual machine settings**. The right side of the graphic shows the **Physical Ports** page in the Manager for a Virtual IPS Sensor.

In case of IPS-VM600, network adapters 3 through 8 correspond to the monitoring ports in the same order. Network adapters 1 and 2 correspond to the management port and the response port respectively.

Task

1 Log on to the VMware ESX as the root user in VMware vSphere Web Client.

2 In the vSphere **Home** tab, select **Hosts and Clusters.**

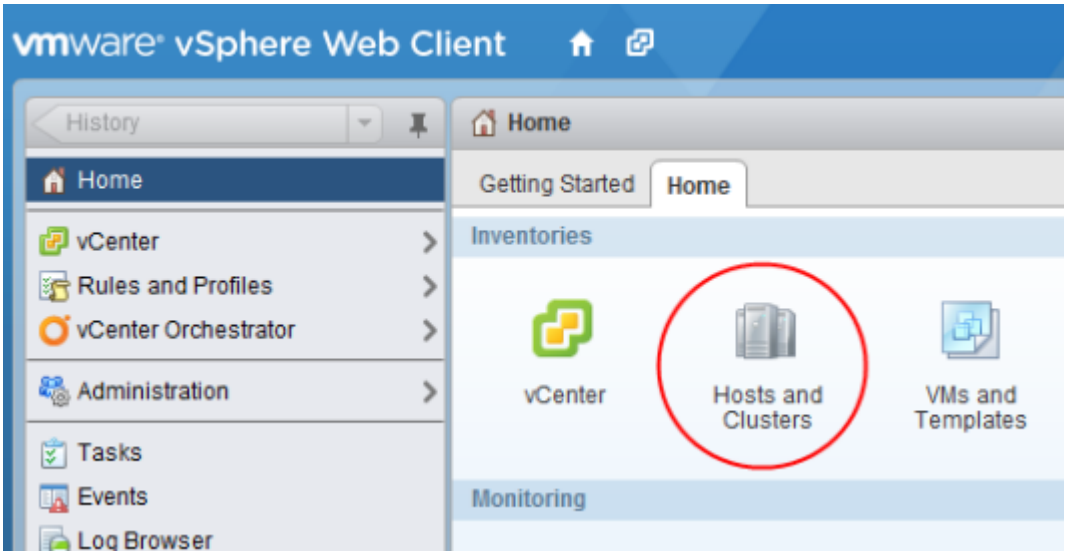

3 Under the corresponding VMware ESX, select the required Virtual IPS Sensor.

You can view the network adapters corresponding to the Sensor ports in the VM Hardware section. The current status of these adapters is also displayed.

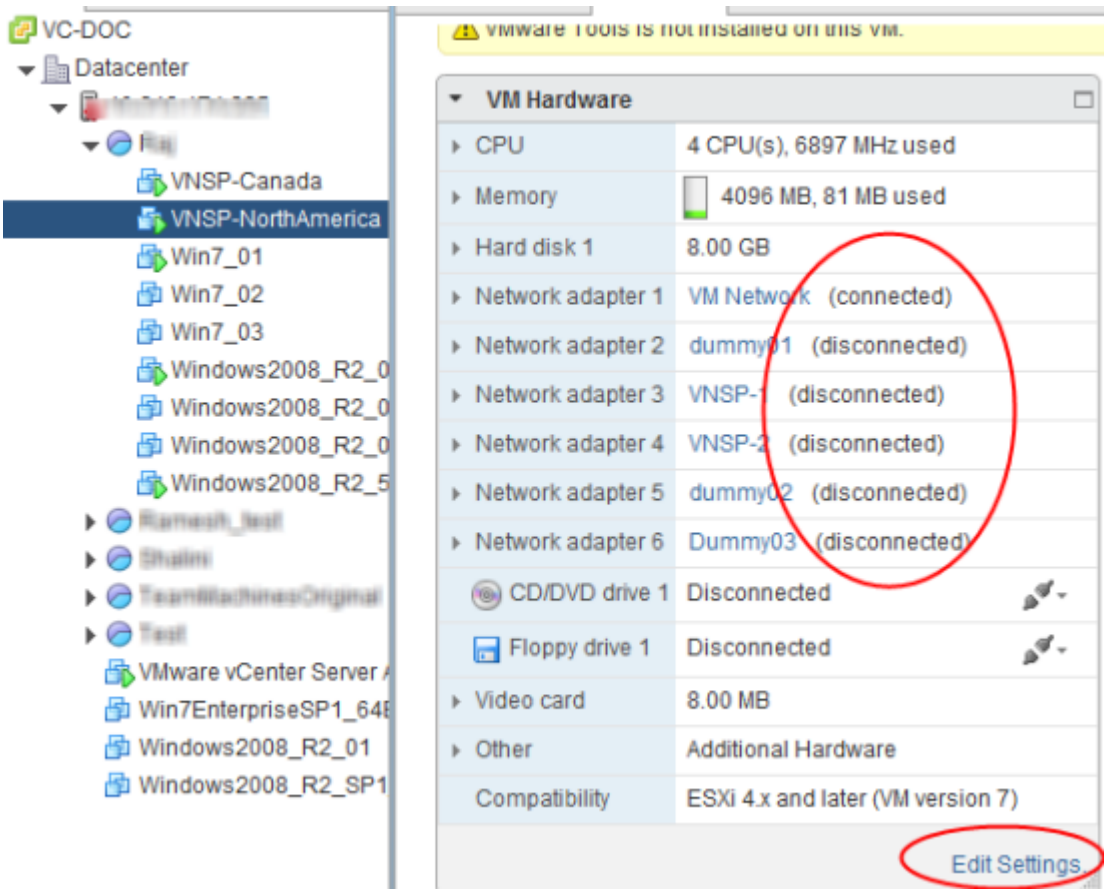

4 Click **Edit Settings** in the **VM Hardware** section.

5 In the **Edit Settings** dialog, select **Connected** check box for the required network adapters (Sensor ports) and click **OK**.

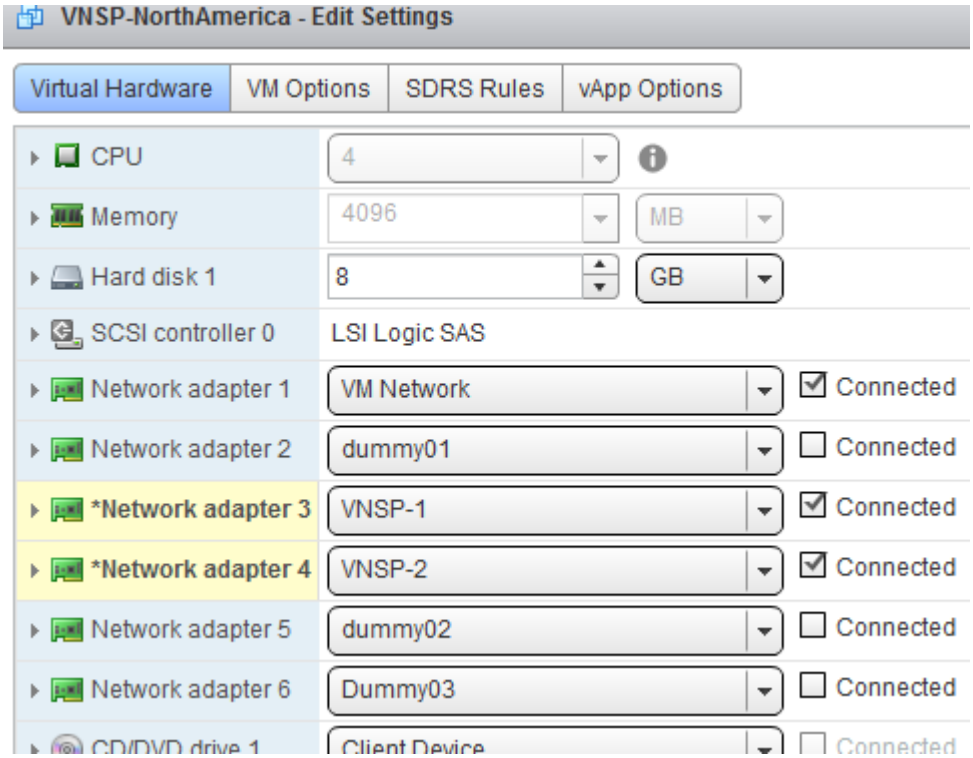

In the **VM Hardware** section, you can verify that these network adapters are now connected.

Deploying Virtual IPS Sensor monitoring ports in inline fail-open mode

Before you begin

Before you deploy monitoring ports in inline fail-open mode, make sure the Sensor is functioning as expected with the same ports in inline fail-closed mode.

If a Virtual IPS Sensor receives traffic from a physical network device, then you can deploy an inline pair in inline fail-open mode. Consider a scenario, wherein the Virtual IPS Sensor inspects traffic between virtual machines as shown (scenario 4).

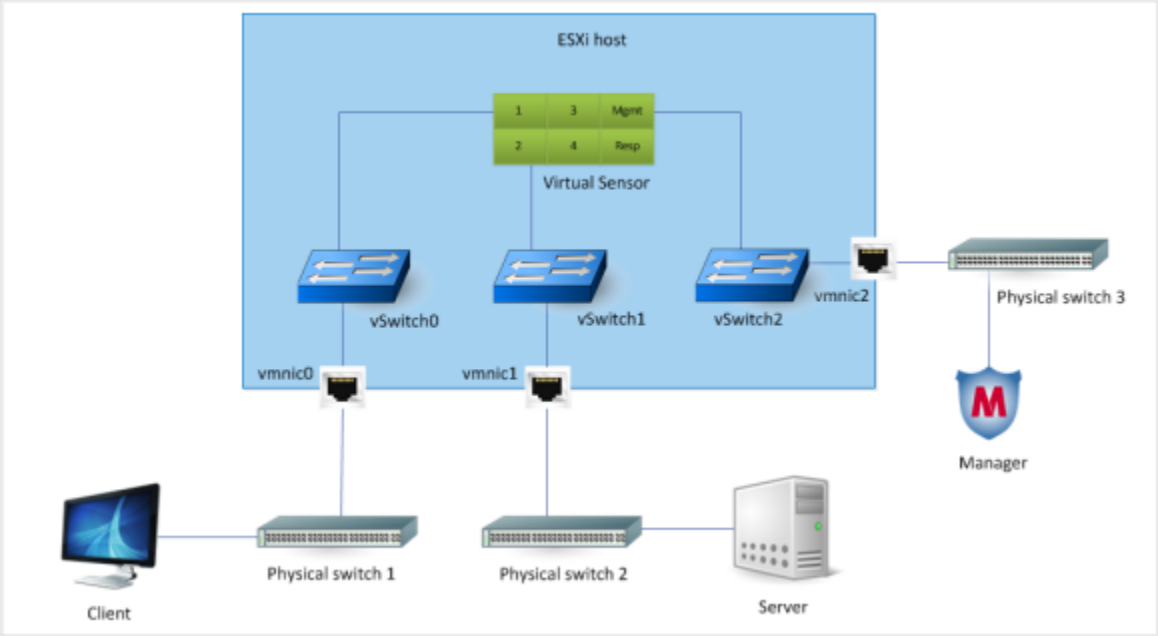

Figure 2-11 Scenario, wherein the Virtual IPS Sensor inspects traffic between virtual machines

The monitoring port pair 1-2 is inline between physical switches 1 and 2. By default, inline monitoring ports of a Virtual IPS Sensor are in the fail-closed mode. The network between the client and the server is broken under any of the following conditions:

- Link failure in either port 1 or 2.
- Power or application failure in the Virtual IPS Sensor.
- Either vSwitch0 or vSwitch1 is down.
- Link failure in either vmnic0 or vmnic1.
- The VMware ESX server is down.

To mitigate the risk of network breakdown due to these conditions, you can deploy the monitoring port 1-2 in fail-open mode. Fail-open operation for the monitoring ports of a Virtual IPS Sensor, require the use of an external copper bypass switch.

- Only McAfee-certified 10/100/1000 external Copper active fail-open bypass kits are supported.
- Installing the active fail-open bypass kit involves a brief network downtime.

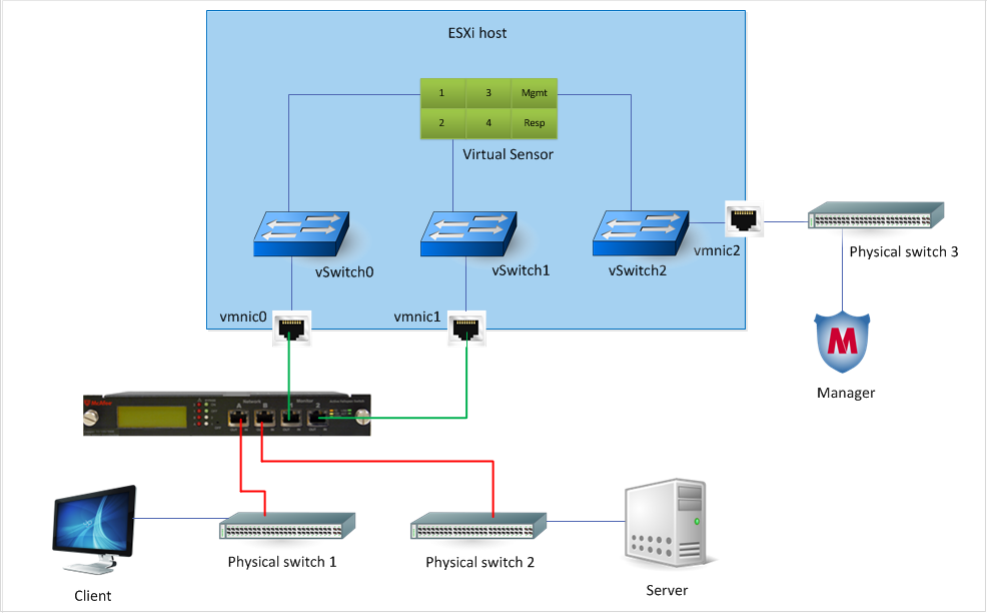

Figure 2-12 Scenario, wherein the active fail-open bypass kit is installed

Task

1 Install the 10/100/1000 external Copper active fail-open bypass kit and power it on (with dual power sources).

Refer to *McAfee Network Security Platform Copper Active Fail-open Bypass Kit Guide* for information.

2 Disconnect the trunk port of Physical switch 1 from vmnic0, and connect the trunk port to the port marked A in the bypass kit.

3 Similarly, disconnect the trunk port of Physical switch 2 from vmnic1, and connect the trunk port to the port marked B in the bypass kit.

4 Connect the port marked 1 to vmnic0 and the port marked 2 to vmnic1.

5 In the Manager, set the port in inline fail-open active mode.

a Click the **Devices** tab.

b Select the domain from the **Domain** drop-down list.

c In the left pane, click the **Devices** tab.

d Select the device from the **Device** drop-down list.

e Select **Setup** | **Physical Ports**.

f Double-click on the required port and select **Inline Fail Open – Active** in the **Mode** field.

g Confirm and then click **Save**.

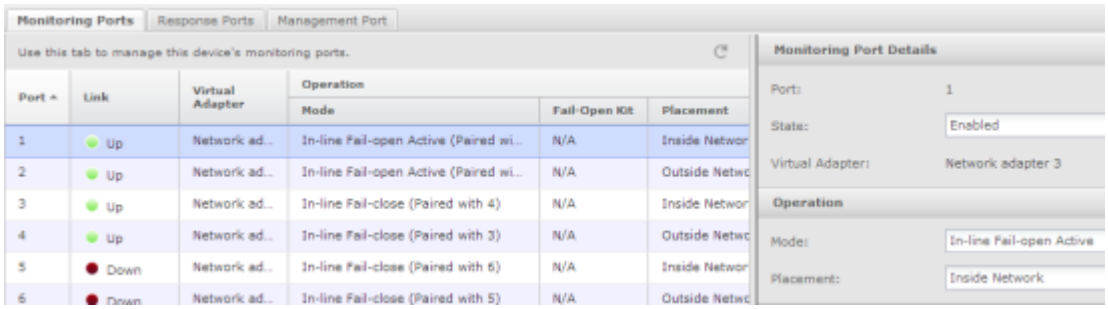

- In the Sensor CLI, run the `show intfport <port number>` command and verify that **Fail-Open Switch** shows **PRESENT** and **Fail-Open Port** shows **INLINE**.
- When the Sensor is operating, the switch is *on* and routes all traffic directly through the Virtual IPS Sensor.

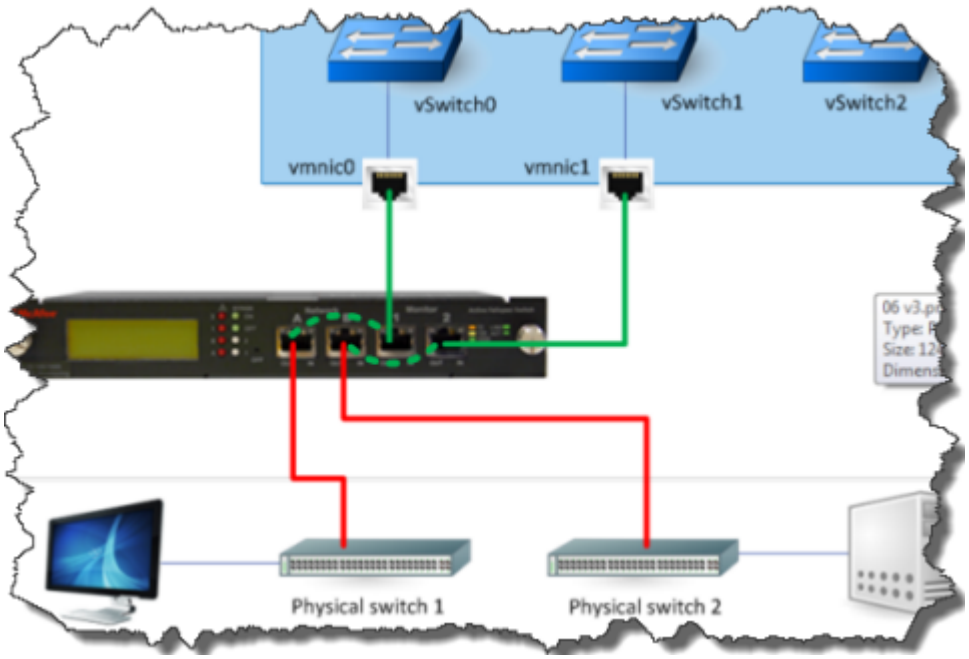

Figure 2-13 Routing when Sensor is operating

- When the Sensor fails, the switch automatically shifts the Virtual IPS Sensor to a bypass state: in-line traffic continues to flow through the network link, but is no longer routed through the Sensor.

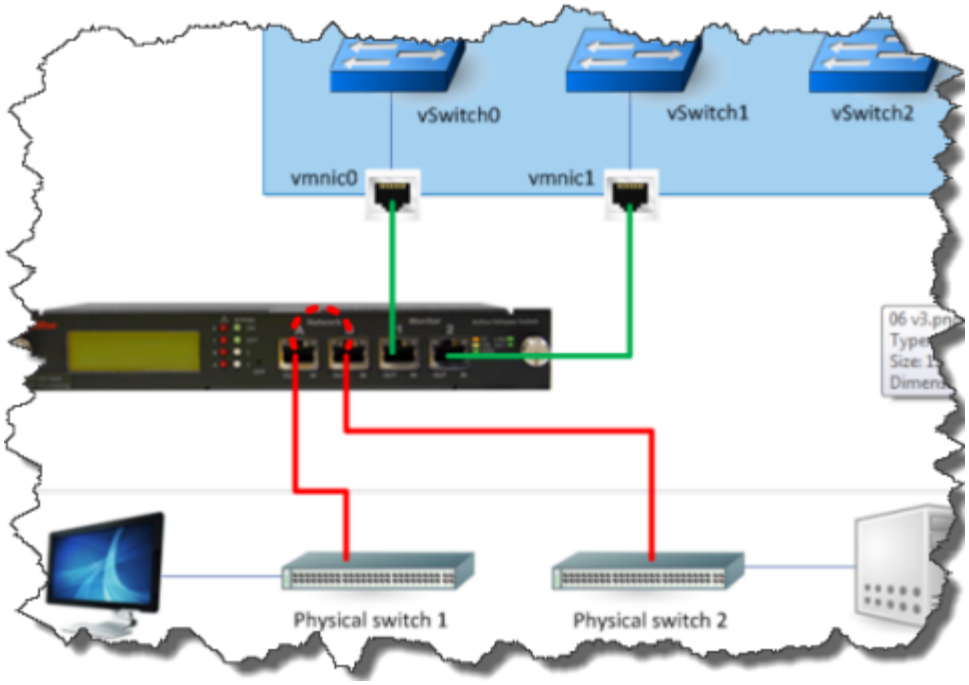

Figure 2-14 Routing when Sensor fails

- When a monitoring port goes down, its status in the Manager is shown as *unknown*. A critical-fault message is also generated. This message is cleared when the monitoring port pair is back inline.

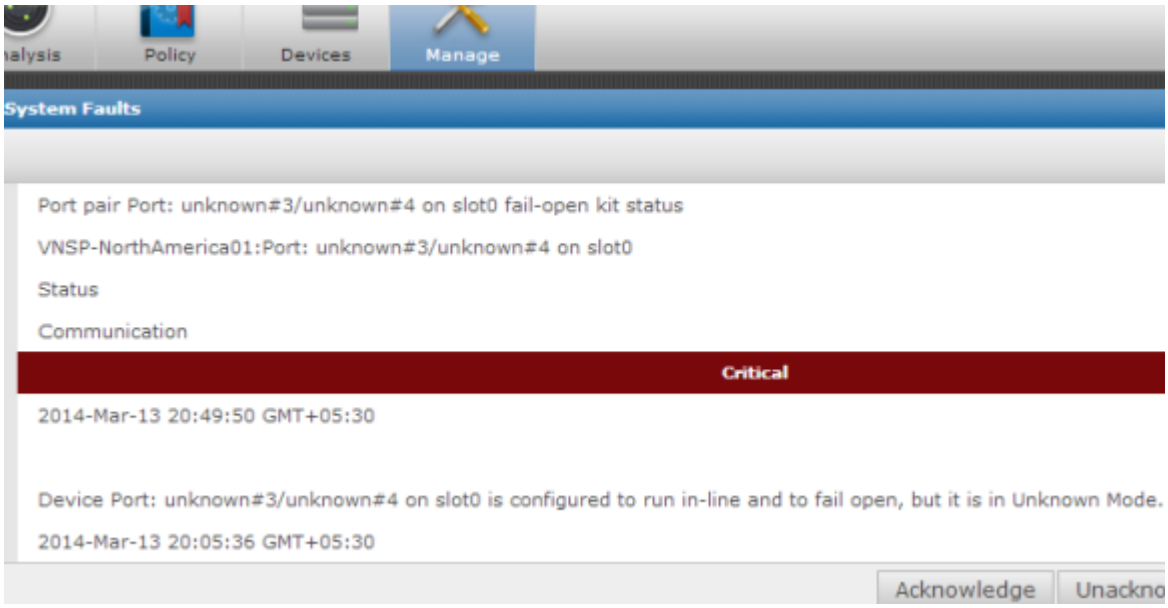

- If the monitoring port is down or if the Sensor goes into layer 2 bypass mode, the fail-open bypass kit turns *on* and the traffic bypasses the Sensor. If you run the `show intfport <port number>` command, the **Fail-Open Port** field displays **BYPASS**.

- If the Sensor layer 2 bypass mode, the **Fail-Open Port** field, displays **LAYER2_BYPASS**.
- Once the Sensor resumes normal operation, the bypass switch returns to the *off* state, again enabling in-line monitoring.

For the details of how the external Copper active fail-open bypass kit works, see the *10/100/1000 Copper Active Fail-open Bypass Kit Guide.*

6 You can also configure the bypass switch to operate in tap mode.
- See *10/100/1000 Copper Active Fail-open Bypass Kit Guide* for information.
- There is no specific configuration in the **Physical Ports** page in the Manager for tap mode. The configuration is only in the bypass switch.
- When the bypass switch is in tap mode, and you run `show intfport <port number>` command, the **Fail-Open Port** field displays **TAP**.

Verify the deployment

You can use the following information to verify if you have deployed the Virtual IPS Sensor correctly and if it is inspecting traffic.

Task

1 Make sure the Sensor management port and the Manager server and reachable to each other.

2 On the Sensor CLI, use the status and show commands to see if trust is established, the channels are up, and if the Sensor is in good health.

3 On the Manager Dashboard, check the **System Health** monitor to verify if the Sensor is active.

4 In the Manager, select **Devices** | **<Admin Domain Name>** | **Devices** | **<Device name>** | **Setup** | **Physical Ports** and check if the monitoring ports and up.

5 Verify if the client and server are reachable from one another.

6 Send a sample attack from the client to the server, for example root.exe, and check if an alert is raised in the Attack Log with the correct details.

7 After you deploy a Virtual IPS Sensor, the process of configuring and managing it is similar to that of a physical Sensor. Therefore, refer to the corresponding section for information. For example, the procedures related to virtualization of monitoring ports is similar to that of physical ports. For information on how to create sub-interfaces out of the monitoring ports of a Virtual IPS Sensor, see Chapter *How to understand virtualization* , *McAfee Network Security Platform IPS Administration Guide.*

Deployment of Virtual IPS Sensors on KVM

To deploy Virtual IPS Sensors on KVM, make sure you have installed KVM on a Linux operating system such as Ubuntu 14.04. Through the Linux operating system, install the Virtual IPS Sensor and establish trust with a Manager. Once you have established trust, the Virtual IPS Sensor is ready to be configured to protect your network. How you configure the Virtual IPS Sensor, however, depends on the network architecture and your security needs.

The following is a high-level procedure that you can consider to install and subsequently deploy a Virtual IPS Sensor:

Task

1 Verify your KVM server meets the hardware and software requirements. See Requirements for deploying Virtual IPS Sensors on page 14.

2 Install the Virtual IPS Sensor and establish a trusted communication channel between the Virtual IPS Sensor and a Manager.

3 Determine how you want to deploy the Virtual IPS Sensor and configure it accordingly. There are several method to deploy the Virtual IPS Sensor on KVM. In this document, we provide you some samples.

   You will be provided with the QCOW image which is the Virtual IPS Sensor image.

Tasks

- *Access KVM* on page 68

Access KVM

This section describes the steps that can be used to access KVM hypervisor on your Ubuntu 14.04 server.

Task

1 Install Ubuntu 14.04 on the server and configure the required network interfaces.

It is recommended that you configure atleast two network interfaces.

2 Launch Putty and login to the Ubuntu 14.04 server using its IP address.

3 Install KVM on the Ubuntu 14.04 server. For more information, refer to the *KVM Installation documentation*.

4 Once you have installed KVM, you can begin deploying your Virtual IPS Sensor using:

- Virtual Machine Manager console - enables you to configure and manage your virtualization host, networking, and storage resources. It allows you to create and deploy virtual machines and services on private clouds that you manage.

5 You can check the status of your Virtual IPS Sensor by using the `virsh` command line. The `virsh` command line interface tool is used for managing guest virtual machines and the hypervisor.

- `virsh list` - lists the Virtual IPS Sensors that are installed on KVM
- `virsh console <vmips name>` - allows you to access the Virtual IPS Sensor

Install the Virtual IPS Sensor on KVM

Before you can deploy a Virtual IPS Sensor to protect your network, you must install the Virtual IPS Sensor on KVM, and establish trust between the Virtual IPS Sensor and the Manager.

The following are the high-level steps to install a Virtual IPS Sensor:

1 Identify the network on which you want to place the Virtual IPS Sensor and the Manager.

2 Install the Virtual IPS Sensor and establish trust with the Manager.

3 Add the Virtual IPS Sensor in the Manager. See Add the Virtual IPS Sensor in the Manager on page 82.

4 For every Virtual IPS Sensor that you plan to deploy, import the required licenses in the Manager. See Manage Virtual IPS Sensor licenses on page 83.

Sample 1: Install a Virtual IPS Sensor through the command prompt

Before you begin

Before you begin to deploy the Virtual IPS Sensor through the KVM command line interface (CLI), decide whether you want use Open vSwitches or Linux bridges for traffic flow.

You can choose to deploy the Virtual IPS Sensor through the CLI in KVM using either Open vSwitches or Linux bridges depending on your requirements. Depending on which Sensor model you have decided to use, you will require

- Six Open vSwitches or Linux bridges to connect the six interfaces of the IPS-VM100
- Eight Open vSwitches or Linux bridges to connect the eight interfaces of the IPS-VM600

These devices are used as:

- One management bridge – Used for communication between the KVM management interface and the Virtual IPS Sensor management port
- One response bridge – Used by the response port of the Virtual IPS Sensor to take response actions during an attack (for example, a TCP reset is sent to an attacker IP address through the response port which closes the session with the attacker and the target)
- Monitoring bridges – Used by the Virtual IPS Sensor to inspect traffic. These ports can be deployed in line or in SPAN mode. Each monitoring port requires one bridge or Open vSwitch — the IPS-VM100 has four monitoring ports and the IPS-VM600 has six monitoring ports.

Let us illustrate this sample deployment mechanism through a IPS-VM100 Sensor.

Task

1 Begin by creating bridges or switches.

In this illustration, we will name the bridges and switches br1 through br6.

- If you are using Open vSwitch, use the `ovs-vsctl` command.

```
ovs-vsctl add-br br1
ovs-vsctl add-br br2
..
..
ovs-vsctl add-br br6
```

- If you are using a Linux bridge, use the `brctl` command.

```
brctl addbr br1
brctl addbr br2
..
..
brctl addbr br6
```

2 Create an XML file using the sample provided in the section, Sample XML file on page 70.

3 Place the XML file in the a convenient location, in the Linux server, where it can be accessed and modified.

4 Edit the XML file by `vi default-vm100.xml` command.

For this illustration, we have chosen to create bridges instead of Open vSwitches. The modified XML file resembles the same below.

5 Replace the bridge IDs with the bridge IDs you provided while creating the bridges.

6 Create the virtual machine.

General syntax

```
virsh define <XML file name>
```

Sample syntax

```
virsh start default-vm100.xml
```

If the operation is successful, you see a message that reads `Domain docvmips100 defined from default-vm100.xml`.

7 Start the Virtual IPS Sensor image up.

General syntax

```
virsh start <instance name>
```

Sample syntax

```
virsh start docvmips100
```

If this is successful, you see a message that reads `docvmips600 started`.

8 Initialize the Virtual IPS Sensor service.

General syntax

```
virsh console <instance name>
```

Sample syntax

```
virsh console docvmips100
```

If this is successful, you see a message that reads `Connected to domain docvmips100`.

9 Use the default credentials to log on to the Virtual IPS Sensor.

Default credentials to all Sensors are:

Username: `admin`

Password: `admin123`

The Virtual IPS Sensor setup begins.

Sample XML file

A sample of the XML file required for your deployment is provided below. We consider deployment of an IPS-VM100 for which the parameters considered in the illustration are:

- Number of CPUs required is 3
- Memory require is 4 GB

- Six bridges are required for which customizable bridge IDs
  - `vmips_mgmt_br` is the Management bridge
  - `resp_br` is the Response bridge
  - `mon_br3` through `mon_br6` are Monitoring bridges
- NIC type is virtio
- Hard disk emulation is IDE
- Path of the file along with the filename is `/home/doc/sensorsw_vm100_83714.qcow`

```
<domain id="2" type="kvm">
 <name>vmips</name>
 <memory unit="KiB">4194304</memory>
  <currentMemory unit="KiB">4194304</currentMemory>
  <vcpu placement="static">3</vcpu>
  <resource>
   <partition>/machine</partition>
  </resource>
  <os>
   <type arch="x86_64" machine="pc-1.0">hvm</type>
   <boot dev="hd" />
  </os>
  <features>
   <acpi />
   <apic />
   <pae />
  </features>
  <clock offset="utc" />
  <on_poweroff>destroy</on_poweroff>
  <on_reboot>restart</on_reboot>
  <on_crash>restart</on_crash>
  <devices>
   <emulator>/usr/bin/kvm</emulator>
   <disk device="disk" type="file">
    <driver name="qemu" type="qcow2" cache='none'/>
    <source file="/home/doc/sensorsw_vm100_83714.qcow" size="10G"/>
    <target bus="ide" dev="hda" />
    <alias name="ide0-0-0" />
    <address bus="0" controller="0" target="0" type="drive" unit="0" />
   </disk>
   <controller index="0" type="scsi">
    <alias name="scsi0" />
    <address bus="0x00" domain="0x0000" function="0x0" slot="0x07" type="pci" />
   </controller>
   <controller index="0" type="ide">
    <alias name="ide0" />
    <address bus="0x00" domain="0x0000" function="0x1" slot="0x01" type="pci" />
   </controller>
   <controller index="0" type="usb">
    <alias name="usb0" />
    <address bus="0x00" domain="0x0000" function="0x2" slot="0x01" type="pci" />
   </controller>
   <controller index="0" model="pci-root" type="pci">
    <alias name="pci0" />
   </controller>

   <interface type="bridge">
    <source bridge="vmips_mgmt_br" />
    <virtualport type="openvswitch" />
    <target dev="vnet0" />
    <model type="virtio" />
    <alias name="net0" />
   </interface>

   <interface type="bridge">
    <source bridge="mon_br1" />
    <virtualport type="openvswitch"/>
```

```
    <target dev="vnet1" />
    <model type="virtio" />
    <alias name="net1" />
   </interface>

   <interface type="bridge">
    <source bridge="resp_br" />
    <virtualport type="openvswitch" />
    <target dev="vnet2" />
    <model type="virtio" />
    <alias name="net2" />
   </interface>

   <interface type="bridge">
    <source bridge="mon_br2" />
    <virtualport type="openvswitch" />
    <target dev="vnet3" />
    <model type="virtio" />
    <alias name="net3" />
   </interface>

   <interface type="bridge">
    <source bridge="mon_br3" />
    <virtualport type="openvswitch" />
    <target dev="vnet4" />
    <model type="virtio" />
    <alias name="net4" />
   </interface>

   <interface type="bridge">
    <source bridge="mon_br4" />
    <virtualport type="openvswitch" />
    <target dev="vnet5" />
    <model type="virtio" />
    <alias name="net5" />
   </interface>

   <serial type="pty">
    <source path="/dev/pts/5" />
    <target port="0" />
    <alias name="serial0" />
   </serial>

   <console tty="/dev/pts/5" type="pty">
    <source path="/dev/pts/5" />
    <target port="0" type="serial" />
    <alias name="serial0" />
   </console>

   <input bus="ps2" type="mouse" />
    <graphics autoport="yes" listen="127.0.0.1" port="5900" type="vnc">
     <listen address="127.0.0.1" type="address" />
    </graphics>
    <video>
     <model heads="1" type="cirrus" vram="9216" />
     <alias name="video0" />
     <address bus="0x00" domain="0x0000" function="0x0" slot="0x02" type="pci" />
    </video>
    <memballoon model="virtio">
     <alias name="balloon0" />
     <address bus="0x00" domain="0x0000" function="0x0" slot="0x05" type="pci" />
    </memballoon>
  </devices>
  <seclabel model="apparmor" relabel="yes" type="dynamic">
   <label>libvirt-f70ba969-610e-6035-abec-8ec3f48cb389</label>
   <imagelabel>libvirt-f70ba969-610e-6035-abec-8ec3f48cb389</imagelabel>
  </seclabel>
</domain>
```

Sample 2: Install a Virtual IPS Sensor through the KVM user interface

Another way to install the Virtual IPS Sensor on KVM is using the KVM user-interface.

Task

1 Log on to the Linux server user interface using the IP address and the credentials.

2 Go to the Virtual Machine Manager.

If you have any instances running, you see these instances listed in this window.

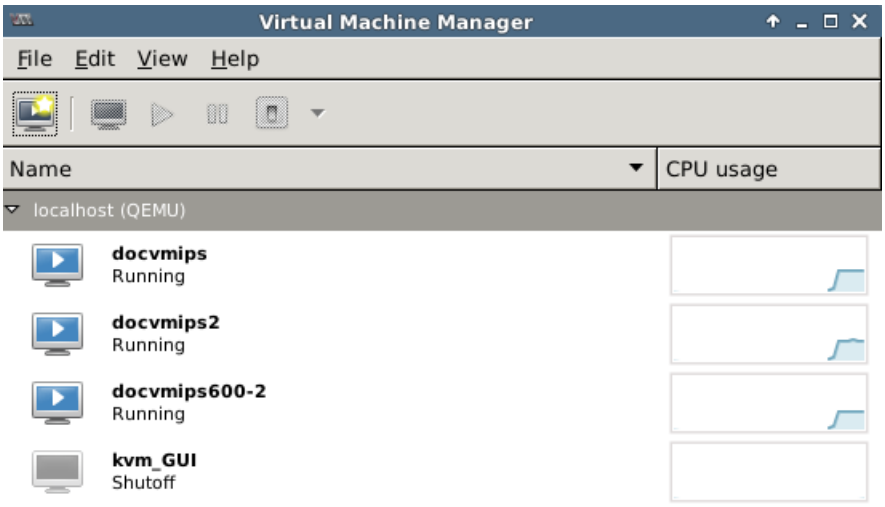

3 Click the icon to create a new virtual machine.

The wizard window appears.

4 Enter a name for the Virtual IPS Sensor, choose the appropriate option to find the installer image from the list, and click **Forward**.

Since we have the image in another location, we select **Import existing disk image**.

Clicking **Forward** brings you to step 2 where you must select:

- the path where the image file is located
- type of OS

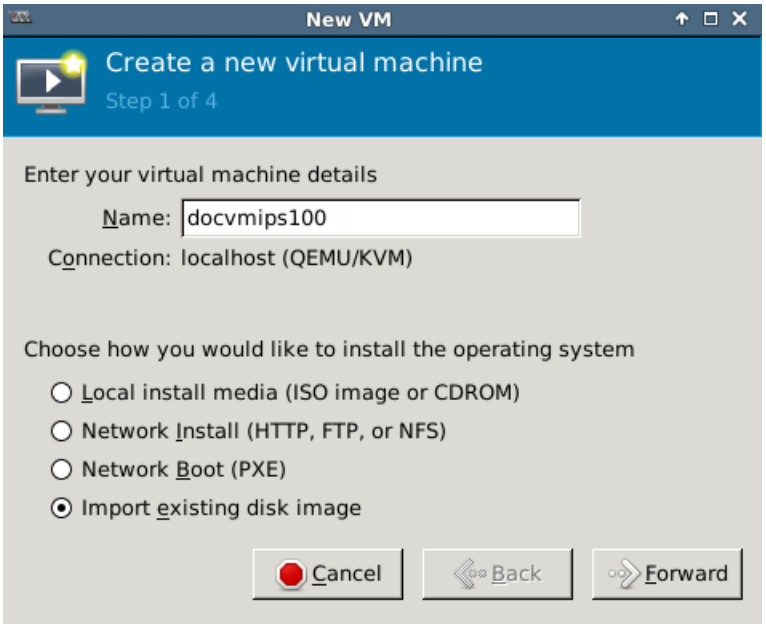

5 From the **OS type** drop-down list, select Linux.

6 From the **Version** drop-down list, select Generic kernel.

7 Browse to the location where the Virtual IPS Sensor image is located and select the QCOW image.

We recommend that you place the software image in a folder other than the root folder.

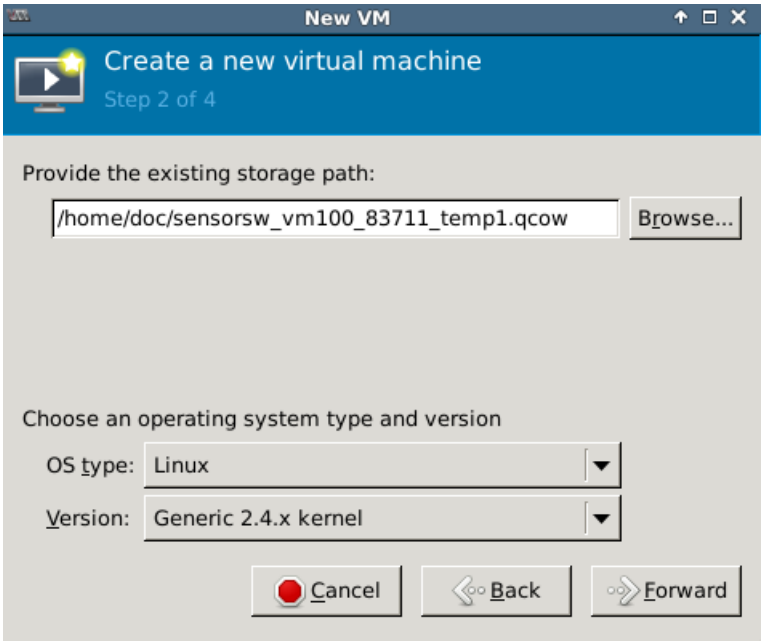

8 Click **Forward**.

We come to step 3 in the deployment where you must select the memory and CPU requirements for the Virtual IPS Sensor.

9 Manually enter the **Memory (RAM)** required and number of **CPUs** required.

For memory and CPU requirements of each Virtual IPS Sensor model refer the section, Requirements for deploying Virtual IPS Sensors on page 14

10 Click **Forward**.

We come to step 4 of the wizard where you opt whether to customize the configuration before installation. You can also review your settings in the previous steps.

11 Select the **Customize configuration before install** checkbox.

12 Expand the **Advanced options** tab.

a Click the **Virtual network** drop-down and select **Specify shared device name**.

This option is meant to specify a shared device which facilitates communication with the Virtual IPS Sensor management port. The **Bridge name** field appears.

b Enter `vr1` which is the shared management bridge to which all Virtual IPS Sensor management interfaces connect with in this illustration.

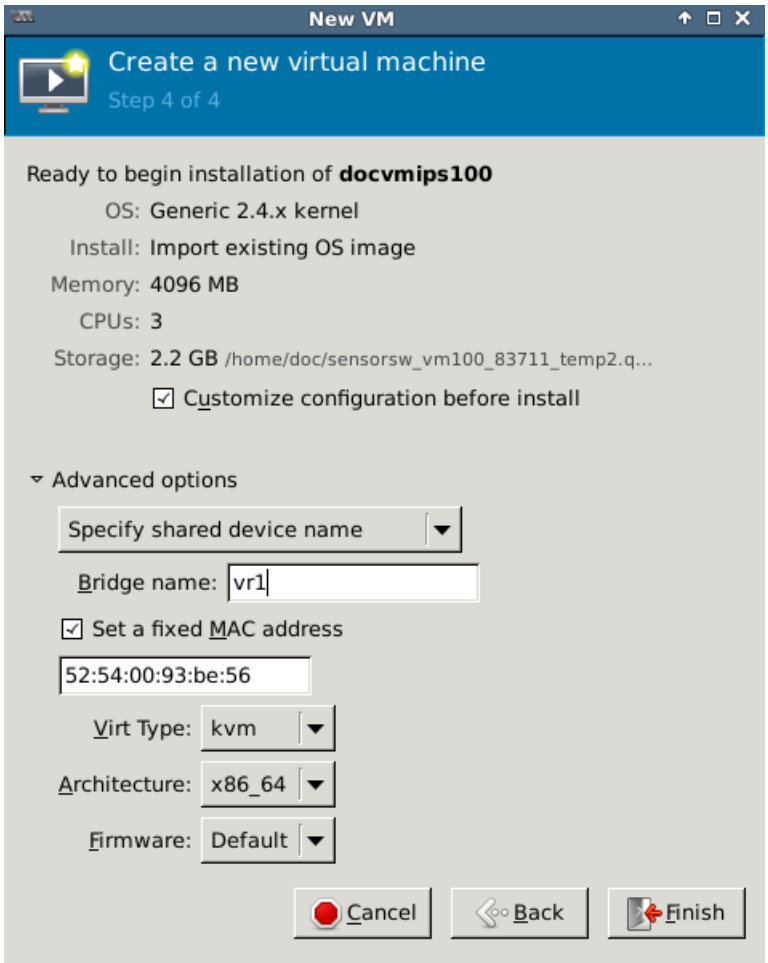

c Make sure the rest of the settings are configured as shown in the list below. For this illustration, these were default settings.

1 Select the **Set a fixed MAC address** checkbox and make sure a MAC address is seen.

2 From the **Virt Type** drop-down list make sure kvm is selected.

3 From the **Architecture** drop-down list, make sure x86_64 is selected.

4 From the **Firmware** drop-down list, make sure Default is selected.

d Click **Finish**.

You are routed to the next step in the deployment where you can review you entire configuration tab-by-tab. You will also need to verify some of the settings in these tabs.

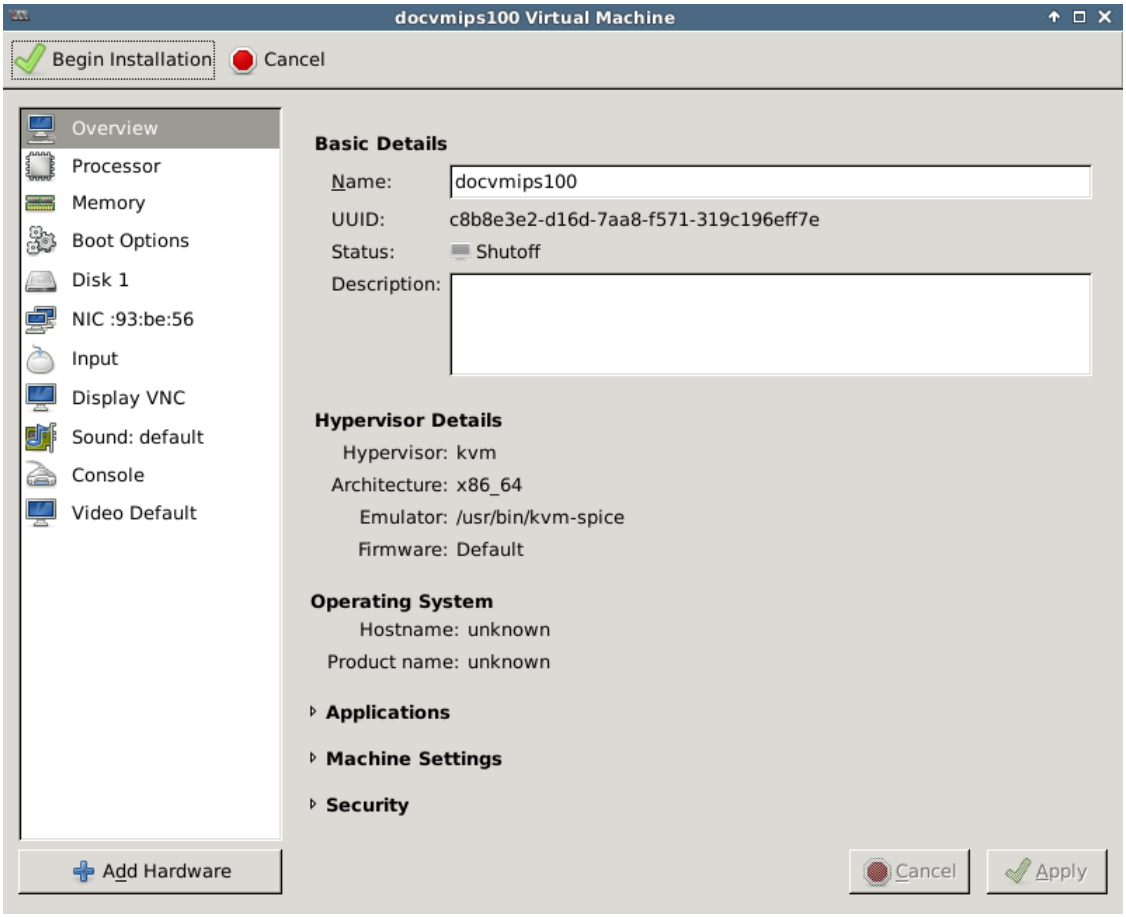

13 To verify and modify appropriate settings:

a Click the **Disk 1** tab.

1 Expand the **Advanced** options tab.

2 From the **Disk bus** drop-down list, select IDE since other formats are not supported.

3 From the **Storage format** drop-down list, select qcow2.

4 Click **Apply** to confirm your changes.

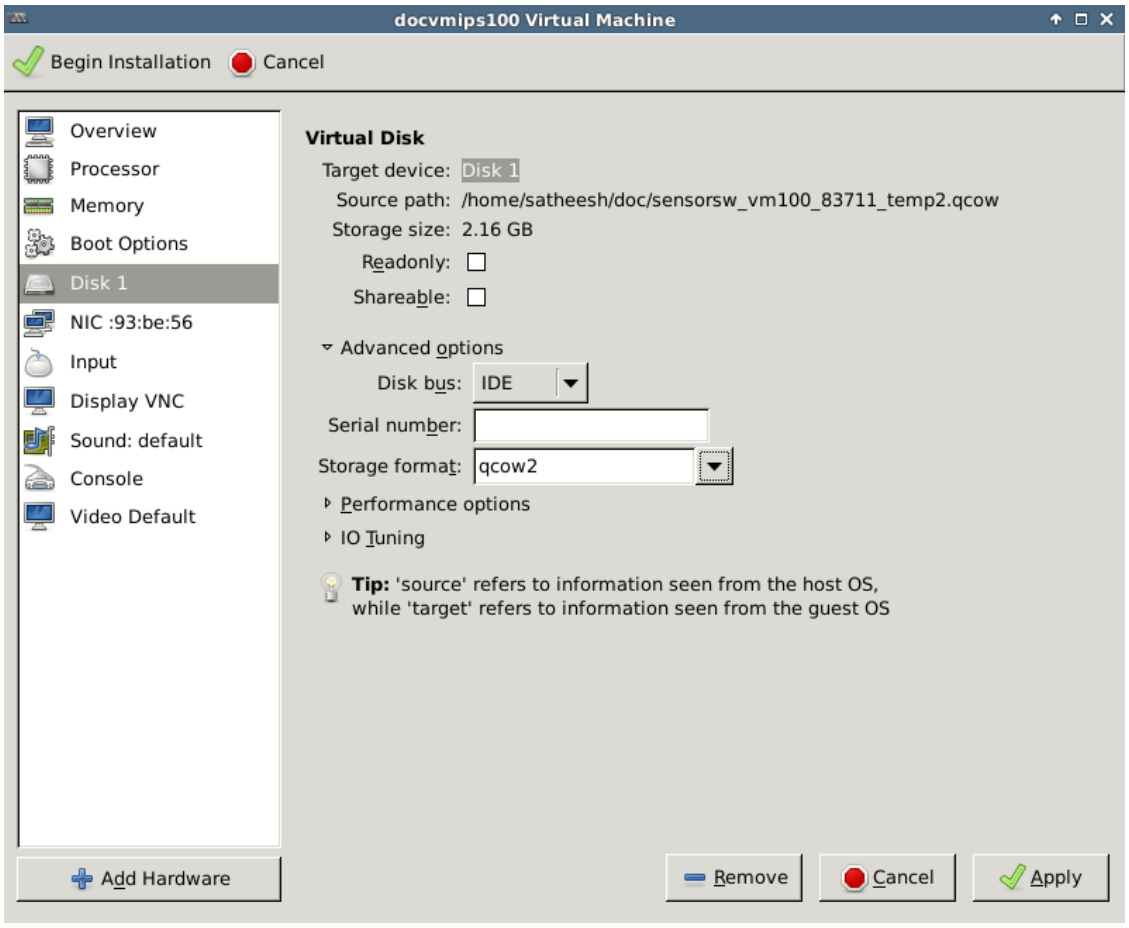

b Click the **NIC** tab.

1 From the **Device model** drop-down list, select virtio.

You can also select e1000 but we recommend selecting virtio for best results.

2 Click **Apply** to confirm your changes.

c Click **Add Hardware**.

The **Add New Virtual Hardware** wizard appears. We require additional NICs (beyond the management NIC) for a Virtual IPS Sensor to function normally.

1 Click on the **Network** tab.

2 From the **Host device** drop-down list, select **Specify a shared device name**.

The **Bridge name** field appears.

3 Enter `vr2` which is the first bridge that we are creating manually.

4 From the **Device model** drop-down, select virtio.

5 Click **Finish**.

A new NIC with a MAC address appears in the menu on the left. The MAC address in the menu refers to that of the Virtual IPS Sensor interface.

d Repeat step c to create another four NICs since we are about to deploy a IPS-VM100.

If you want to deploy a IPS-VM600, create seven NICs in addition to the first NIC that appears in the menu. The difference between the number NICs required for each model is because the IPS-VM100 has four monitoring ports whereas the IPS-VM600 has six monitoring ports.

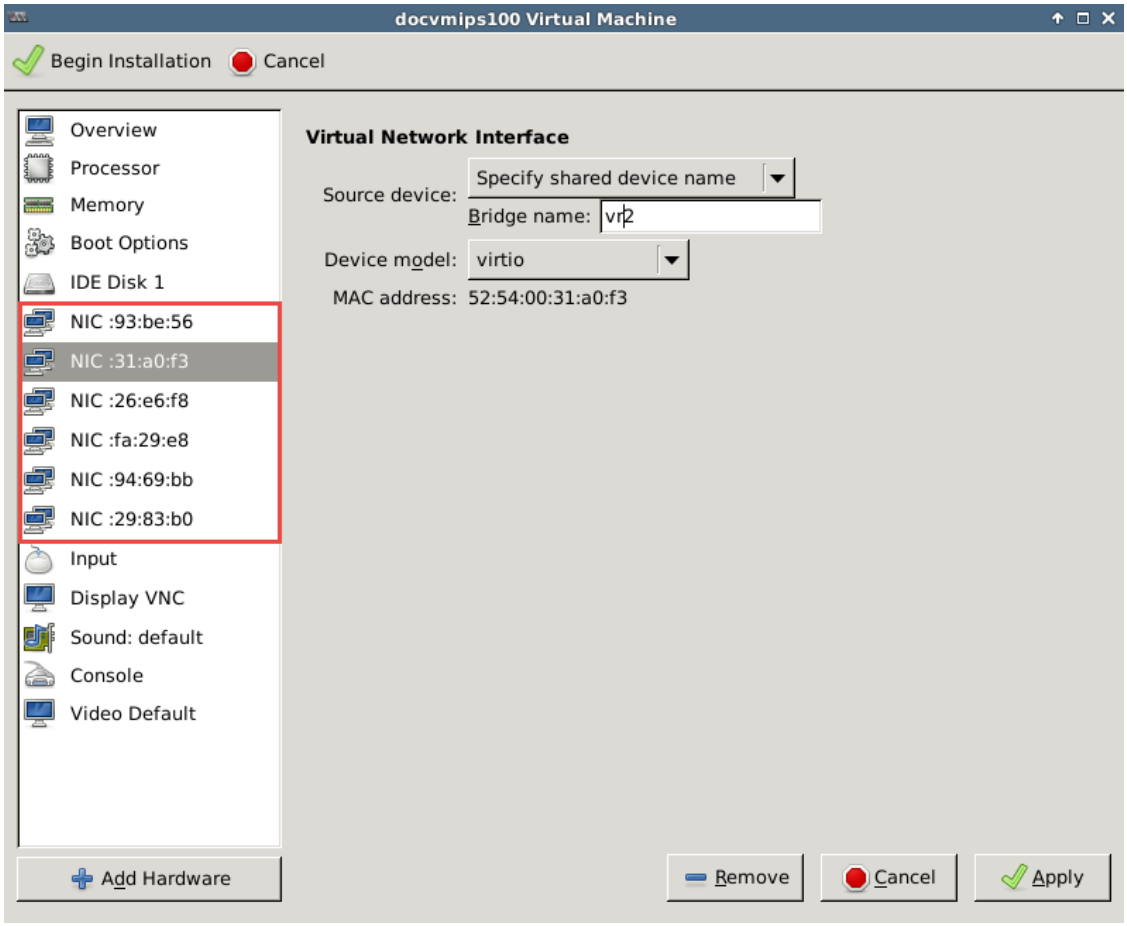

14 Click **Begin Installation**.

Creation of the Virtual IPS Sensor virtual machine begins. This process take a few minutes. After installation of the Virtual IPS Sensor, you are routed to the login prompt for the Virtual IPS Sensor CLI which displays the name of the instance you chose in step 1 of the wizard.

15 Provide default credentials for the Virtual IPS Sensor and set it up like you would any other Network Security Platform IPS Sensor.

After you complete setup of the Virtual IPS Sensor, you are able to assign it to a Manager. Managing the Virtual IPS Sensor through the Manager is the same as managing any other Sensor.

Sample 3: Install a Virtual IPS Sensor using a command

Before you begin

Before you begin to deploy the Virtual IPS Sensor through the KVM command line interface (CLI), you must decide whether you want use Open vSwitches or Linux bridges for traffic flow.

You can choose to deploy the Virtual IPS Sensor through the CLI in KVM using either Open vSwitches or Linux bridges depending on your requirements. Depending on which Sensor model you have decided to use, you will require

- Six Open vSwitches or Linux bridges to connect the six interfaces of the IPS-VM100
- Eight Open vSwitches or Linux bridges to connect the eight interfaces of the IPS-VM600

These devices are used as:

- One management bridge – Used for communication between the KVM management interface and the Virtual IPS Sensor management port
- One response bridge – Used by the response port of the Virtual IPS Sensor to take response actions during an attack (for example, a TCP reset is sent to an attacker IP address through the response port which closes the session with the attacker and the target)
- Monitoring bridges – Used by the Virtual IPS Sensor to inspect traffic. These ports can be deployed in line or in SPAN mode. Each monitoring port requires one bridge or Open vSwitch — the IPS-VM100 has four monitoring ports and the IPS-VM600 has six monitoring ports.

These steps illustrate a sample deployment mechanism of a IPS-VM600 Sensor using a command in the KVM command line interface.

Task

1 Begin by creating bridges or switches.

In this illustration, we will name the bridges br1 through br8.

- If you are using Open vSwitch, use the `ovs-vsctl` command.

```
ovs-vsctl add-br br1
ovs-vsctl add-br br2
..
..
ovs-vsctl add-br br8
```

- If you are using a Linux bridge, use the `brctl` command.

```
brctl addbr br1
brctl addbr br2
..
..
brctl addbr br8
```

2 Open the string provided in the block below, in a notepad editor such as Notepad++, to install your Virtual IPS Sensor.

Strings for both Virtual IPS Sensor models, IPS-VM100 and IPS-VM600, are provided.

IPS-VM100

```
virt-install --name=<instance_name> --ram=4096 --arch=x86_64 --vcpus=3 --os-type=linux --i
mport --disk path=/home/doc/
sensorsw_vm100_83714_temp2.qcow,bus=ide,size=6,format=qcow2 --network
bridge:mgmt_br,model=virtio  --network bridge:mon_br1,model=virtio --network
bridge:resp_br,model=virtio --network bridge:mon_br2,model=virtio --network
bridge:mon_br3,model=virtio --network bridge:mon_br4,model=virtio
```

IPS-VM600

```
virt-install --name=<instance_name> --ram=6144 --arch=x86_64 --vcpus=4 --os-type=linux --i
mport --disk path=/home/doc/
sensorsw_vm600_83714_temp1.qcow,bus=ide,size=8,format=qcow2 --network
bridge:mgmt_br,model=virtio  --network bridge:mon_br1,model=virtio --network
bridge:resp_br,model=virtio --network bridge:mon_br2,model=virtio --network
bridge:mon_br3,model=virtio --network bridge:mon_br4,model=virtio --network
bridge:mon_br5,model=virtio --network bridge:mon_br6,model=virtio
```

3 Edit the parameters mentioned below depending on your setup.

- Instance name: This is the name of the Virtual IPS Sensor instance also referred to in the Sensor CLI as Sensor name. It must be unique since it will be the same name that reflects in the Manager
- Bridge IDs: These must match the bridge that you provided while creating the bridges using the `brctl addbr` command in KVM
- Disk path: This is the path of the QCOW file within the Linux file system

4 Log on to the Linux CLI using your credentials.

5 Copy the string from the notepad editor and paste it in the CLI.

If this is successful, installation of the Virtual IPS Sensor begins. After the installation is complete the Virtual IPS Sensor comes to the `setme login` prompt. Setting up the Virtual IPS Sensor from this point is the same as setting up any other Sensor.

Troubleshooting scenarios

While deploying the Virtual IPS Sensor you need make sure that vital parameters required to bring it up are mentioned as per requirements. The parameters along with their requirements that are critical for the deployment are:

- CPU: IPS-VM100 - 3 | IPS-VM600 - 4
- Memory (RAM): IPS-VM100 - 4096 MB (4 GB) | IPS-VM600 - 6144 MB (6 GB)
- NIC type: virtio
- Number of bridges: IPS-VM100 - 6 | IPS-VM600 - 8

Errors you will see if any of these parameters are incorrect

When any of the parameters above are configured incorrectly, the command prompt displays errors which enable you to correct the problem. The errors along with the messages that you receive are mentioned in this section. All screenshots in this section use an IPS-VM100 to illustrate. Similar errors messages appear if you attempt to deploy the IPS-VM600 with incorrect parameters.

Incorrect parameter	Error
Number of CPUs required is provided incorrectly	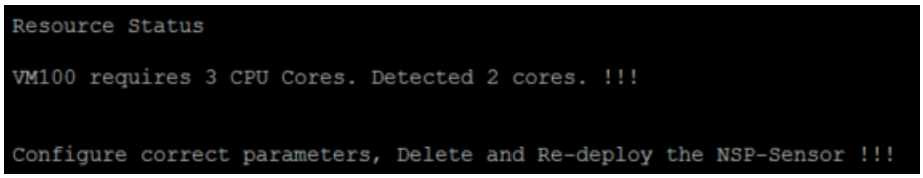
Memory allocated is other than the required memory	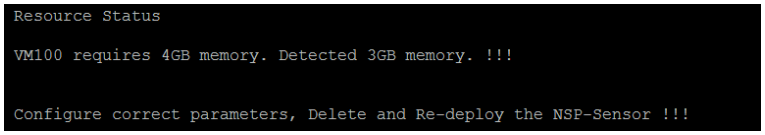
NICs created are not virtio type of NICs	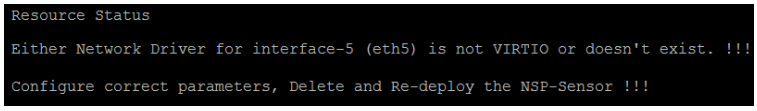
Number of NICs created is fewer than the required number	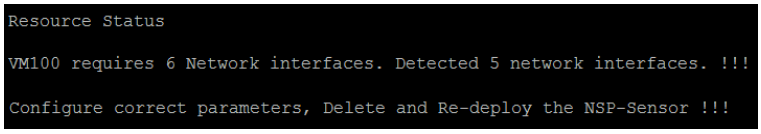

Uninstall the Virtual IPS Sensor from KVM

This section describes the steps required to uninstall your Virtual IPS Sensor from KVM.

Sample 1: Uninstall a Virtual IPS Sensor through the command prompt

To uninstall the Virtual IPS Sensor from KVM, you must undefine the domain that deployed for the instance and then remove the instance itself.

Task

1. Log on to the Linux server through a hyperterminal client such as PuTTY.
2. Run the `virsh undefine <XML file name>` command.

   If the operation is successful, you see `Domain <XML file name> is undefined`.
3. Run the `virsh destroy <instance name>` command.

   If the operation is successful, you see `Domain <instance name> is destroyed`.

   You can now reuse the bridges to deploy another instance of a Virtual IPS Sensor.

Sample 2: Uninstall a Virtual IPS Sensor through the KVM user interface

Another way to remove the Virtual IPS Sensor in KVM is through the KVM user interface.

Task

1 Log on to the Linux server user interface using the IP address and the credentials.

2 Go to the Virtual Machine Manager.

All instances you are running are listed in this window.

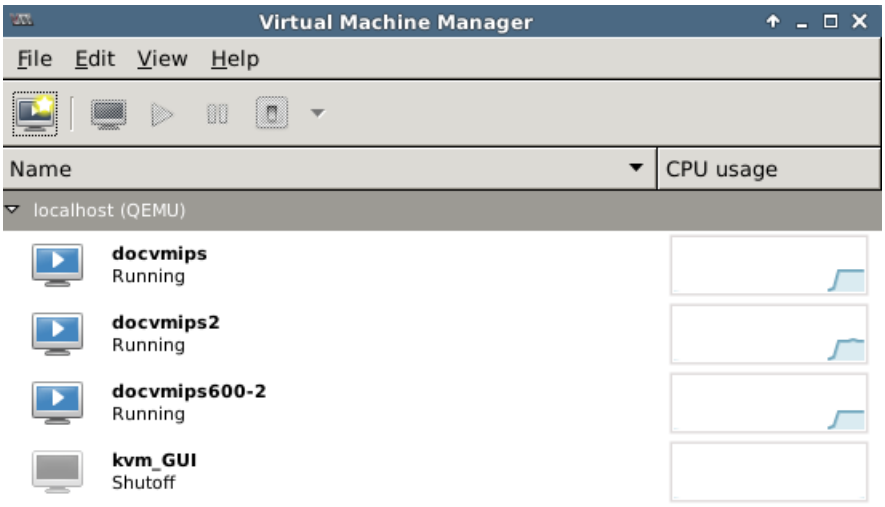

3 Right-click the instance that you want to remove and click Delete.

The Delete option is greyed out until a virtual machine is shut down. To shut down the Virtual IPS Sensor, log on to the Sensor CLI and use the `shutdown` command.

After you click Delete, you are prompted to choose whether you want to delete the instance from the list or from the disk.

4 Select either of the options and remove the Virtual IPS Sensor.

Add the Virtual IPS Sensor in the Manager

Before you begin

- For VMware ESX, you have a Manager version 8.1.7.x or later either on a virtual or a physical machine.
- For KVM, you have a Manager version 8.3.7.x or later either on a virtual or a physical machine.

You can add the Virtual IPS Sensor in the Manager and specify a shared secret key. You can use a Manager installed on a virtual machine or on a physical server.

Task

1 In the Manager select **Devices** | **<Admin Domain Name>** | **Global** | **Add and Remove Devices.**

2 Click **New.**

3 Specify at least the **Device Name**, **Device Type**, and **Shared Secret**.

- Select **IPS Sensor** as the **Device Type**.
- You must use the same **Device Name** and Shared Secret when you deploy the Virtual IPS Sensor.

4 Click **Save**.

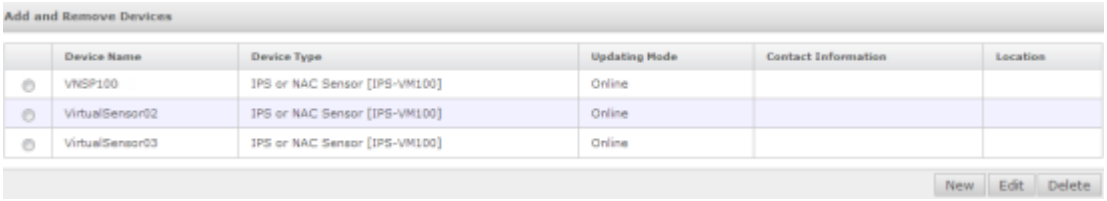

Figure 2-15 Virtual IPS Sensors added in a Manager

For more information on deleting a configured Sensor, see topic *Delete a Sensor from the Manager*, section *How to replace a Sensor*, chapter *Troubleshooting* in the *McAfee Network Security Platform IPS Administration Guide.*

Manage Virtual IPS Sensor licenses

Before you begin

- The license file that you received from McAfee is accessible from your Manager client.
- You have access to the root-admin domain in the Manager.

Virtual IPS Sensors are licensed per installation. You need to secure a license per Virtual IPS Sensor managed by a Manager. Also, licenses are model-specific. Licenses apply to all Virtual IPS Sensor models. McAfee provides license files that you can import into the Manager. Each license file will contain the number of supported Virtual IPS Sensors.

One license file might contain one or more licenses. Before you deploy Virtual IPS Sensors, it is a good practice to first import the required licenses into the Manager. You can import and delete Virtual IPS Sensor license files from the Virtual IPS Sensor Licenses page. This page is available only in the root admin domain. So, even for the Virtual IPS Sensors managed by child domains, the license file is imported only in the root admin domain.

The Manager periodically checks if there are enough licenses imported already. If there are not enough licenses, then this information is displayed in the `License Summary` section of the Virtual IPS Sensor Licenses page. Also, a critical system health fault message is raised for the Manager. The Manager periodically checks the number of licenses against the installed number of Virtual IPS Sensors and raises this fault message accordingly.

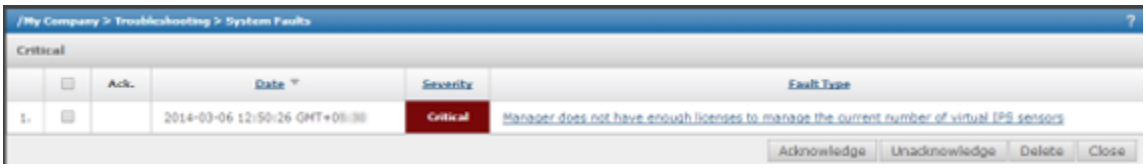

Figure 2-16 Virtual IPS Sensor license-non-compliance fault message

Notes

- There is no option currently to export license files from a Manager. So, make sure you safely retain the license files that you received from McAfee.
- You cannot import the Virtual IPS Sensor licenses in the Central Manager.
- In case of MDR, you must import the license file in the currently active Manager. This license file is pushed to the peer as part of data synchronization.
- To increase the number of licenses, simply import additional license files provided by McAfee.
- If you want to transfer a part of the licenses to a different Manager, you must contact McAfee Support.
- If you delete a license file, and if the number of licenses is less than the number of Virtual IPS Sensors managed, then the **License Summary** information is displayed accordingly.
- Consider you want to move a Virtual IPS Sensor to a different Manager Server. Delete the license file in the source Manager and then import that license file in the target Manager. Then, uninstall the Sensor from the source Manager and then install it with the target Manager.
- If you enable GTI participation, then the following information is sent to McAfee Global Threat Intelligence (McAfee GTI):
  - Number of Virtual IPS Sensor licenses required to be compliant.
  - Number of currently managed Virtual IPS Sensors.
  - Total number of licenses currently imported.

  You can verify this by clicking **Show Me What I'm Sending** in the **Global Threat Intelligence** page **(Manager | <Admin Domain Name> | Integration | GTI).**
- If you want to generate a report on your compliance with respect to Virtual IPS Sensor licenses, you can generate the **Virtual IPS Sensor License Compliance** report.

Task

1 In the Manager select **Manager | <Admin Domain Name> | Setup | Virtual IPS Sensor Licenses.**

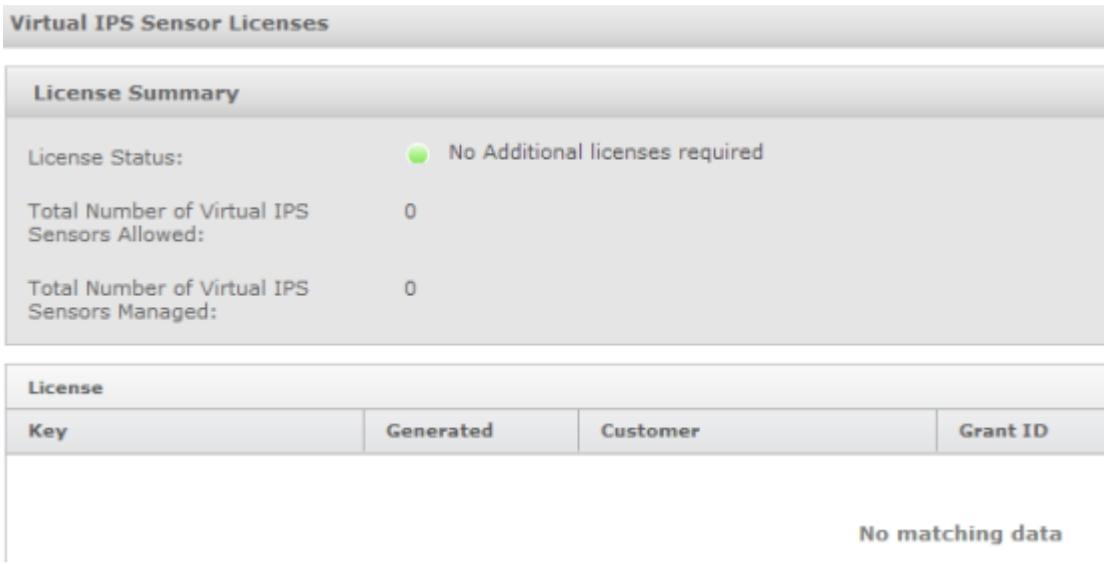

2 To import licenses, click **Add**, browse to the .zip or .jar file containing the licenses, and then click **OK**.

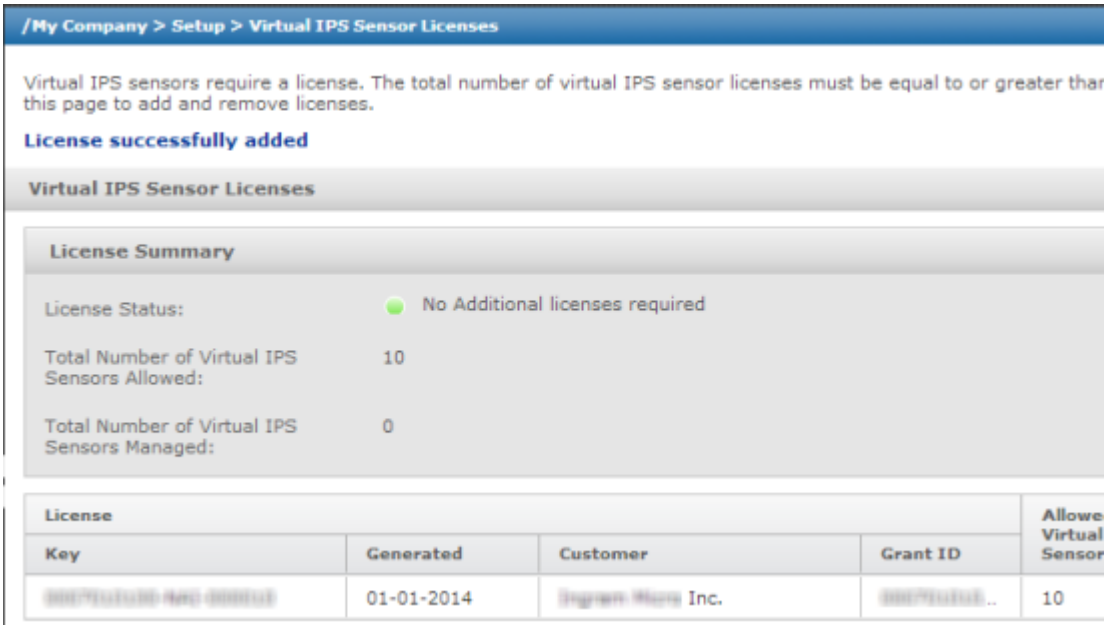

3 To delete a license file from the Manager, select the required license file and click **Remove**.

4 Click **OK** to confirm deletion.

If deleting the license file results in non-compliance, an alert message is displayed and you must click **OK** to proceed.

Option	Description
License Status	Indicates the current compliance status.
Total Number of Virtual IPS Sensors Allowed	This number quantifies the number of Virtual IPS Sensors that can be under this Manager.
Total Number of Virtual IPS Sensors Managed	Indicates the number of Virtual IPS Sensors currently managed by the Manager.
License Key	Displays the license key of the license file.
Generated	The date when the license file was generated.
Customer	The customer for whom the license file was generated.
Grant ID	The McAfee Grant ID of the corresponding customer.
Allowed Virtual IPS Sensors	This corresponds to the number of licenses per license file.
Imported Time	The timestamp of when the license file was imported into the Manager.
Imported By	The user who imported the license file.
Comment	Enables you to add your comment per license file that is imported. Double-click in the **Comment** field and enter your comment. Click outside this field and the your comment is automatically saved.
Add	Click to import a license file.
Remove	Select a license file and then click **Remove** to delete a license file from the Manager.

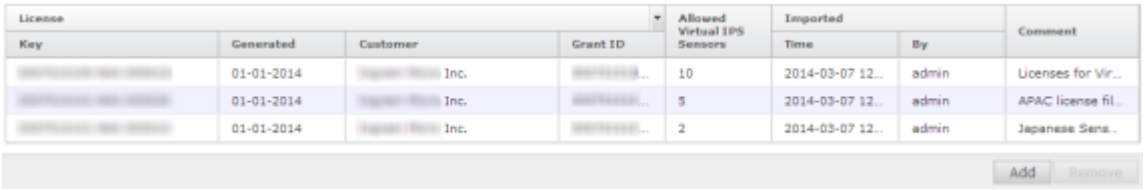

Figure 2-17 Details of the imported license files

Generate the Virtual IPS Sensor License Compliance report

If you want to generate a report on your compliance with respect to Virtual IPS Sensor licenses, you generate the Virtual IPS Sensor License Compliance report. This report summarizes the current Virtual IPS Sensor license compliance. In addition, it also lists the Virtual IPS Sensors currently managed by the Manager.

Task

1 In the Manager select **Manager** | **<Admin Domain Name>** | **Reporting** | **Configuration Reports** | **Virtual IPS Sensor License Compliance.**

2 Select the required option from the **Output Format** list and click **Submit.The Virtual IPS Sensor License Compliance Report** is displayed

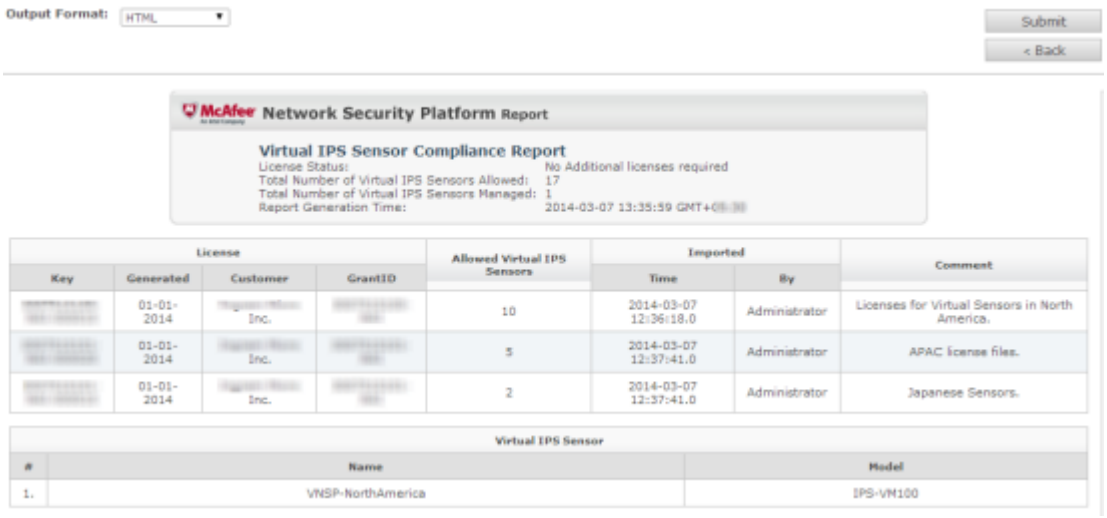

Figure 2-18 Virtual IPS Sensor License Compliance Report

Option	Description
License Status	Indicates the current compliance status.
Total Number of Virtual IPS Sensors Allowed	This number quantifies the number of Virtual IPS Sensors that can be under this Manager.
Total Number of Virtual IPS Sensors Managed	Indicates the number of Virtual IPS Sensors currently managed by the Manager.
Report Generation Time	Time when you generated the report.
License Key	Displays the license key of the license file.
Generated	The date when the license file was generated.
Customer	The customer for whom the license file was generated.
Grant ID	The McAfee Grant ID of the corresponding customer.
Allowed Virtual IPS Sensors	This corresponds to the number of licenses per license file.
Imported Time	The timestamp of when the license file was imported into the Manager.
Imported By	The user who imported the license file.
Comment	The user comment that was entered in the **Virtual IPS Sensor Licenses** page.
Virtual IPS Sensor Name	Names of the Virtual IPS Sensors managed under the admin domain.
Virtual IPS Sensor Model	The model of the Virtual IPS Sensors.

3 IPS for virtual networks using VMware NSX

If your Manager version is later than 8.3.7.44, you can integrate McAfee® Network Security Platform with Open Security Controller (OSC) to provide next-generation IPS service for software-defined datacenters (SDDC).

This section provides the following information:

- An overview of OSC and how it collaborates with McAfee Network Security Platform and NSX to provide IPS service to SDDCs.
- Detailed procedures to deploy next-generation IPS for virtual networks using OSC.

  To deploy IPS service to SDDCs, you configure OSC, Manager, and NSX. To configure NSX, you need vCenter web client.

This section assumes the following:

- You have installed and set up OSC virtual appliance. For information, see the latest *Open Security Controller Product Guide*.
- For information on installing and configuring VMware NSX, refer to VMware documentation.
- For information on installing and configuring the virtual infrastructure, refer to the corresponding documentation. For example, for VMware virtual infrastructure, refer to VMware documentation.

Limited information on how to set up and configure a VMware-based SDDC is provided in *OSC Product Guide*. However, this information is designed only for setting up test environments. For production environments and details of VMware functionalities, you must refer to VMware documentation.

Contents

- *Securing virtual networks with Open Security Controller*
- *Deploying next generation IPS service to a virtual network*

Securing virtual networks with Open Security Controller

OSC is a centralized platform to enable software-defined security for software-defined datacenters (SDDC). OSC provides a common set of management services, acting as a broker between the security solutions and the virtual infrastructure. You can use OSC to provide service such as next-generation IPS to virtual infrastructures.

OSC integrates with a hypervisor and an networking provider to provide security solutions as a service to your virtual networks. Using OSC as a liaison between the security service and its associated components, and the virtualization providers, you are able to provide security services for virtual networks. Currently, you are able to integrate with McAfee® Network Security Platform (Network Security Platform).

To illustrate this, consider a virtual environment that uses VMware vCenter as its hypervisor and VMware NSX as its SDN controller to deploy security services on virtual infrastructure.

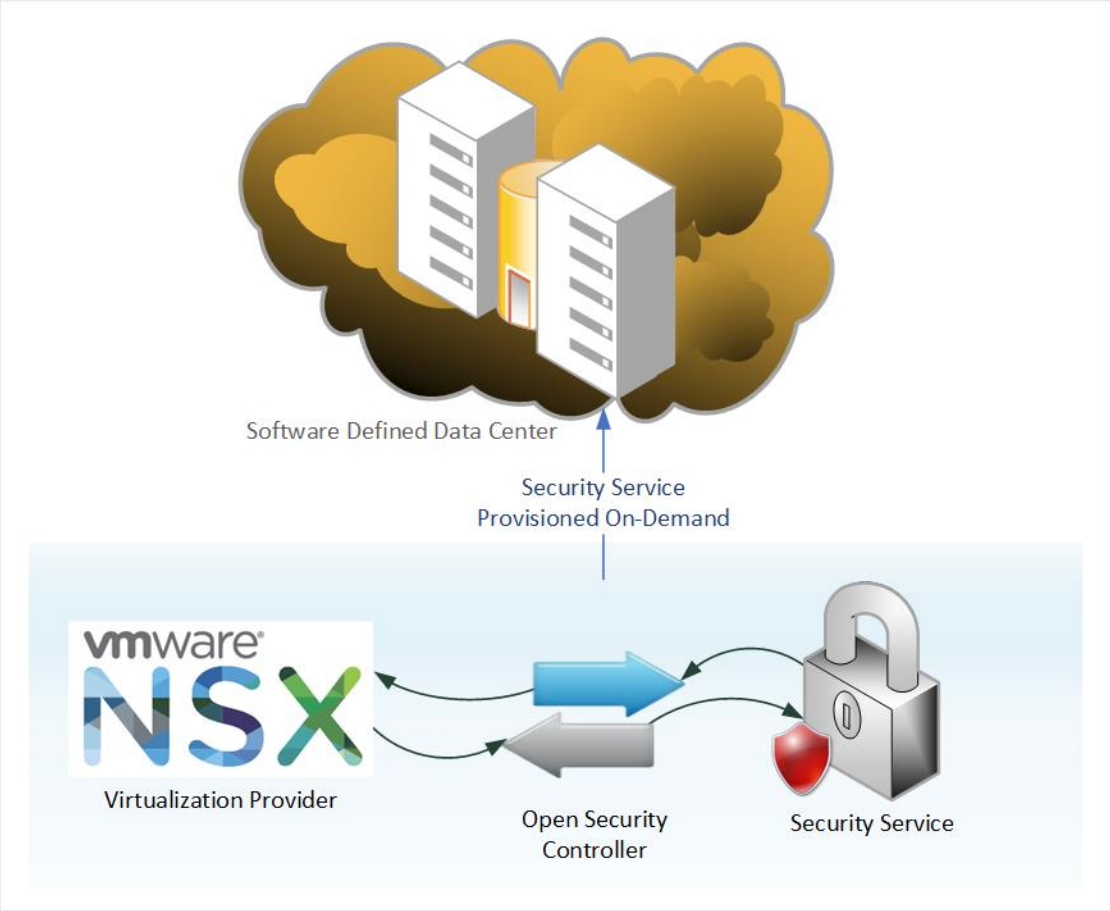

Figure 3-1 OSC solution overview

OSC is a virtual appliance that you install on a VMware ESXi host. It provides a Java-based web application for configuration and management. You can deploy OSC on existing virtual infrastructure without any configuration changes to those virtual networks.

Contents

- *Security challenges in an SDDC*
- *How OSC secures virtual networks?*
- *Advantages of Open Security Controller*
- *Virtual IPS Sensors deployed through Open Security Controller*

Security challenges in an SDDC

Consider a large-scale SDDC consisting of hundreds of hosts aggregated under multiple clusters. Virtualization provides flexibility and agility to its users, wherein they can spin up virtual machines (VMs). Users can spin up isolated logical networks as easily as one can spin up VMs. All these

possibilities require no changes in the physical networking configuration. When multiple users spin up new networks and move working VMs across physical boxes in such a large-scale data center, security is threatened.

To match with the capabilities of virtualization solutions, OSC can seamlessly, non-intrusively, and non-disruptively integrate security services with existing virtualized environments. This enables network security services to keep pace with the speed, agility, and scalability of virtualization features and solutions.

How OSC secures virtual networks?

To understand how OSC can orchestrate security for virtual networks, consider a VMware-based SDDC as illustrated here. For the sake of explanation, this SDDC is shown to contain only a few VMware ESXi hosts.

Two VMware ESXi hosts are clustered together. A few Windows VMs are connected to a distributed vSwitch that spans across these two VMware ESXi hosts. VMware vCenter and NSX are installed on a third VMware ESXi host in the same data center but outside the cluster. With vMotion, you can move the VMs between host-1 and host-2.

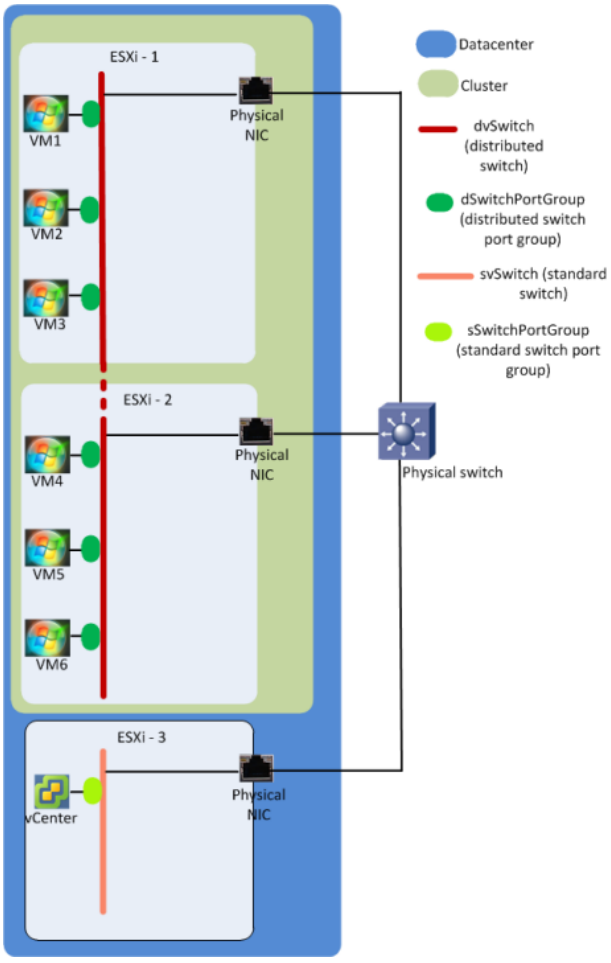

Figure 3-2 An example SDDC

Consider a data center in which you want to implement a security service function (security service) using OSC. This illustration shows next-generation IPS.

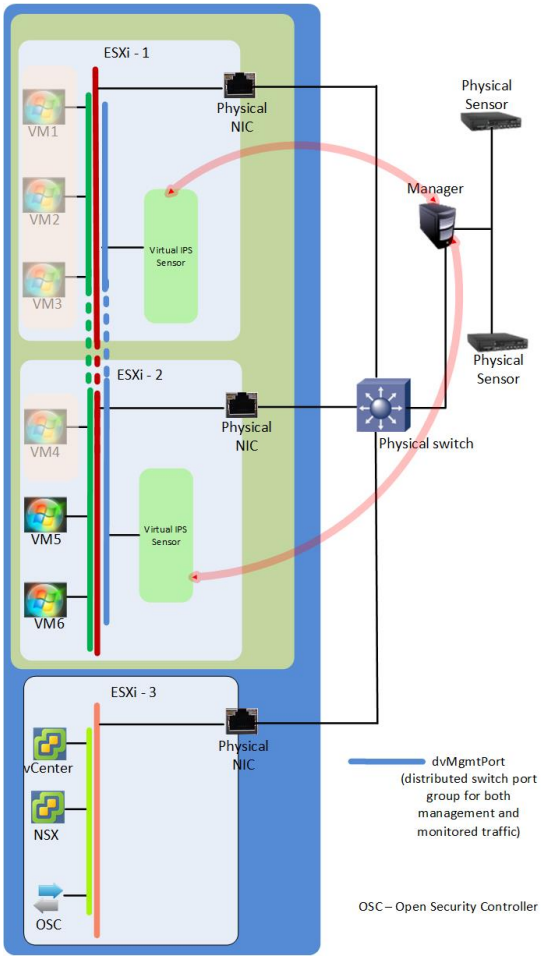

Figure 3-3 SDDC with IPS service through OSC

To deploy a security service, some of the generic tasks you need to complete are.

1. Install the OSC virtual appliance in a VMware ESXi host outside the cluster in which you want to deploy the security service.
2. You deploy the security service to an ESXi cluster such that when you select a VMware ESXi cluster, the security service is available to all VMs on the corresponding VMware ESXi hosts. In this example, there is only one cluster. OSC collaborates with vCenter and NSX such that virtual security is automatically deployed in every host in the selected cluster.
3. You select a security manager for managing these security appliances. Consider that you are using an existing security manager, which is managing other security appliances in your network. Only minimal user intervention is required to install these security appliances. Discovery and establishment of trust with the corresponding security manager is automatic.
4. In NSX, you create security groups containing the objects you want to protect in the cluster. For example, you can select the cluster itself or the distributed switch port group. Then the security service is available for all VMs corresponding to the selected object. In this example, VMs 1 through 4 are selected for IPS service. So traffic related to these VMs is subjected to next-generation IPS analysis. In effect, if VM1 communicates with VM2, traffic does not exit the host. However, the security service inspects such traffic.

Notes:

- Even if you migrate one of the protected VMs to a different host, the same security service is automatically provided to that VM.
- If you add a host to the cluster, you must make sure that network virtualization components are installed on the host as part of host preparation through NSX. Then, an instance of the security service is automatically installed on that host. Trust is also automatically established, by way of exchange or a password or a certificate key, between this instance of the security service and the corresponding security manager.
- Consider that you select the distributed switch port group (dSwitchPortGroup) as the object to be protected. Then any new VM added to this switch port group is automatically subjected to IPS (or other security service).

Advantages of Open Security Controller

- OSC facilitates simple, seamless, non-intrusive, and non-disruptive security service integration with an existing virtualized environment.
- The best-in-class security solutions available to your physical networks are available to your virtual networks as a software-based service.
- Regarding securing virtual networks, OSC can cope with the flexibility, scalability, and agility of virtualization solutions. After you deploy a security service, your virtual networks are protected with minimal user-intervention as they scale up.
- When you deploy an IPS service, you can use your current security manager and those security policies for the SDDC.
- You do not need to change your physical or virtual network architecture to provide a security service to your virtual networks.
- Provides visibility of intra-VM traffic (east-to-west) for security.
- Under test conditions, OSC did not impact functionality or performance of virtualization solutions.

Virtual IPS Sensors deployed through Open Security Controller

When you implement IPS service using OSC, it involves the deployment of a type of Virtual IPS Sensors. These Virtual IPS Sensors are referred to as Virtual Security Systems. A Virtual Security System is to some extent similar to the regular Virtual IPS Sensors, such as the IPS-VM100 Sensor.

Table 3-1 Differences between Virtual IPS Sensors and Virtual Security Systems

Virtual IPS Sensors	Virtual Security Systems
You manually install Virtual IPS Sensor and establish trust with the Manager. Sensor shared secret key is required to establish trust.	Sensor deployment and trust establishment is automatic. OSC orchestrates the deployment and trust establishment. Sensor shared secret key is not relevant. Post deployment of Virtual Security System instances, if necessary, you can re-establish trust between the Manager and a specific Virtual Security System instance from the OSC web application.
Each Virtual IPS Sensor is an individual entity.	A Virtual Security System is a container object for its instances. That is, if your VMware ESXi cluster consists of two VMware ESXi hosts, a Virtual IPS Sensor is deployed in each host. These two Virtual IPS Sensors are the instances of one Virtual Security System. So, the Virtual Security System is the logical container of these two Virtual IPS Sensors. The Virtual IPS Sensors (or instances of a Virtual Security System) are the appliances, which inspect traffic. Each Virtual Security System instance is an individual entity. However, these instances are managed as one Virtual Security System. This is similar to how a Manager manages a Sensor failover pair. Therefore, the same security policies and other configuration are applied to the instances of a Virtual Security System.
The scope of Virtual IPS Sensors is limited to the host in which it is deployed.	Virtual Security System caters to the entire cluster. You can upgrade or downgrade the software version of the Virtual Security System instances in the cluster. Also, when you deploy the changes from the Manager, the changes are deployed to the Virtual Security System instances in the cluster.
Virtual IPS Sensors work with standard virtual switches and standard switch port groups only.	Virtual Security System instances work with distributed virtual switches and distributed switch port groups only.
Virtualization features such as vMotion and DRS are not supported.	Next-generation IPS service is available to a VM regardless of where it is located in the data center. Also, the IPS configuration applied on that VM does not change because the VM is now in a different VMware ESXi host. DRS has no impact on how a Virtual Security System function.
Inline and SPAN are supported.	Only inline is supported.
Inline deployment might require changes to your standard virtual switches. So, it might involve a brief network downtime.	No changes to distributed switches, port groups, or physical network configuration are required.
Depending on the model, a virtual management port, a virtual response port, and a number of virtual monitoring ports are visible in the Manager.	The virtual ports are not exposed.

Deploying next generation IPS service to a virtual network

In the current release, OSC supports only next-generation IPS to be deployed as a service to your virtual networks. When you deploy the IPS service, the traffic from and to the protected VMs are subjected to inline inspection by a Virtual IPS Sensor. If a VM is the source of malicious traffic, the Virtual IPS Sensor takes the configured response action.

One Virtual IPS Sensor is installed per VMware ESXi host. The Manager that you specified in the Security Manager Connector manages all these Virtual IPS Sensors. These Virtual IPS Sensors are configured similarly but function independently. That is, these Virtual IPS Sensors provide IPS to their respective VMware ESXi hosts but implement the same IPS and other IPS configuration.

To deploy the Virtual IPS Sensors, OSC integrates with NSX. This integration also ensures that the relevant traffic is routed through the Virtual IPS Sensor for inspection.

The IPS service makes available all relevant next-generation IPS features for your dynamic virtual networks. Deploying the IPS service is non-intrusive and non-disruptive even though the Virtual IPS Sensors are deployed in inline mode only. Scaling up or changing your virtual networks do not warrant any kind user-intervention to your IPS service deployment. Also, any change to the IPS configuration is automatically applied to all Virtual IPS Sensors. OSC does not take any action directly but orchestrates the actions by its integration with VMware ESX, vCenter, Manager, and the Virtual IPS Sensors.

Contents

- *Terminologies*
- *Components involved in IPS service*
- *High-level steps to implement a security service*
- *How the IPS service works*
- *Requirements for deploying IPS service*
- *Considerations*
- *Prepare a VMware ESXi host for NSX*
- *Define an IP pool for virtual security appliances*
- *Define virtualization connectors*
- *Define manager connectors*
- *Manage software images for security appliances*
- *Manage distributed appliances*
- *Jobs and tasks*
- *Create a security group in VMware NSX*
- *Create a security policy in VMware NSX*
- *Apply a security policy to a security group in VMware NSX*
- *Configure Virtual Security System to fail-close or fail-open*
- *Manager functions regarding IPS service deployment*
- *FAQs regarding IPS service*

Terminologies

To configure OSC to act as a broker, you first define the building blocks. You then use these building blocks to configure OSC so that it can act as a broker between the virtualization provider and security solutions.

- Virtualization connector — In this building block, you define the virtualization provider entities. You must confirm that the virtualization provider is accessible to OSC.
  - For VMware, you define the IP address and administrator logon credentials for NSX and vCenter.
- Appliance instances — The virtual security appliances, which intercept the traffic from the VMs.

  For IPS, Virtual IPS Sensors are the security appliance instances, which are referred to as *Virtual Security System instances*.
- Security service manager connector (Manager connector) — In this building block, you define the management console for managing the security appliances. For IPS, you define the Manager IP address and the root admin logon credentials, which will manage the Virtual IPS Sensors installed in the hosts.

- Security service function (security service) — This component refers to the security service you intend to deploy such as next-generation IPS. You can use the **Service Function Catalog** page to upload corresponding software images for further deployment through OSC.
- Distributed appliances — A distributed appliance, associates the security solution and the virtualization solution. That is, you define a distributed appliance using the virtualization connectors and security manager connector as building blocks.

  In a distributed appliance, you specify the following:

  - One security manager connector.
  - The model and version of the security appliance.

    For IPS, this is the version and model of the Virtual IPS Sensors, which are later deployed in the VMware ESXi hosts.
  - For each virtualization connector, you must select a Manager admin domain. The security appliances are managed under the specified admin domain. In the case of IPS, if you select *My Company* (root admin domain), all Virtual IPS Sensors are managed under *My Company* in the Manager.

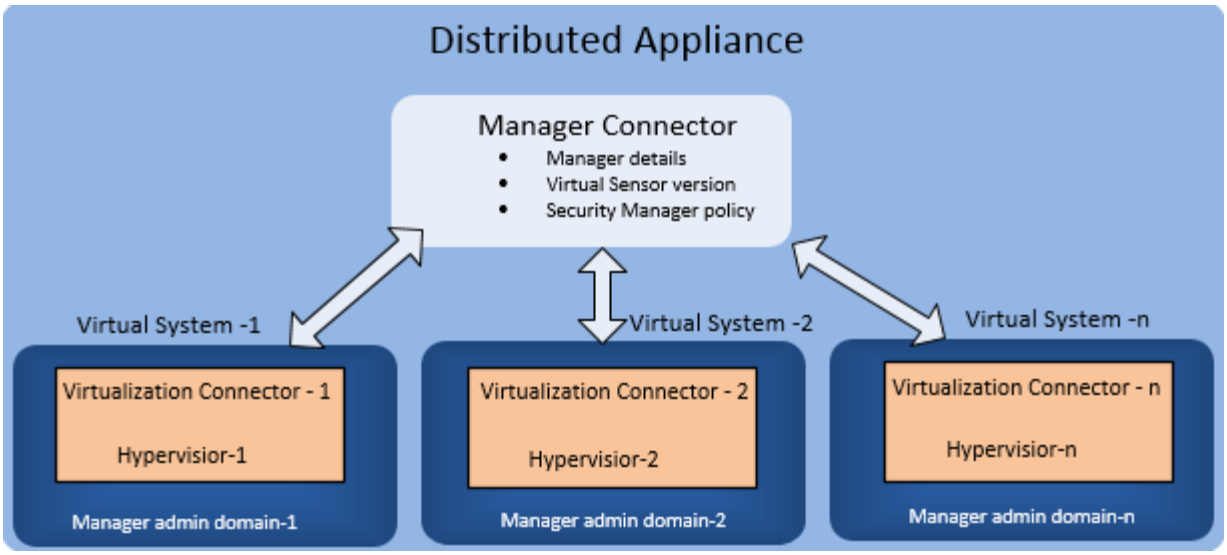

- Virtual system: A virtualization connector associated with a manager domain is a virtual system.

  The most common example of a virtual system is the Virtual Security System for IPS. A Virtual Security System is the logical container object for all deployed virtual security service functions or Virtual Security System instances.

- OSC agents: Security services deployed through OSC have the following agents.
  - Control Path Agent: This agent is responsible for communication between the security services and the security manager.
  - Data Path Agent: This agent makes sure the traffic from the VMs are routed through the security service for inspection in case of VMware.
- Job and tasks: Some of the actions that you perform in OSC are treated as jobs and tasks. The high-level action is treated as a job. For example, synchronizing a distributed appliance is a job. A job might consist of a number of tasks. That is, a job can be broken down into tasks. For example, if synchronizing a distributed appliance is the job, checking the manager connector and validating existing NSX components are some of the tasks. When all tasks are completed successfully, the corresponding job is complete.

  Jobs and tasks enable you to easily track and troubleshoot your actions in OSC. When you trigger a job, the state, status, start time, completed time, and so on are displayed for the job as well as its component tasks.

Components involved in IPS service

The following are the components directly related to IPS service provided through OSC:

- Manager: This is the Network Security Manager which manages all the Virtual IPS Sensors deployed through OSC. You provide the IP address and admin logon credentials when you define the Security Manager Connector. When you deploy the IPS Service, OSC ensures that the installed Virtual IPS Sensors automatically establish trust with the specified Manager.

  

  There is no shared secret key configuration required on Virtual IPS Sensors and Manager.

  Once trust is established, these Sensors and the Manager communicate similar to the regular Sensors. Currently, you cannot use Managers that are part of an Manager Disaster Recovery (MDR) pair to manage the Virtual IPS Sensors deployed through OSC.
- Virtual IPS Sensor deployed through OSC: The Virtual IPS Sensor deployed through OSC for the most part are similar to the regular Virtual IPS Sensors. However, these Sensor images are slightly modified to exchange communication with OSC. Therefore, you cannot install the regular Virtual IPS Sensors through OSC; you must use the specific Virtual IPS Sensor images for OSC. Once these Sensors are deployed and trust established with the Manager, these Sensors function similar to the regular Virtual IPS Sensors.

  

  The Virtual IPS Sensors deployed through OSC support all the relevant features similar to the regular Virtual IPS Sensors.
- VMware ESX Agent: In vCenter, the Virtual IPS Sensors deployed through OSC are grouped under a resource pool called VMware ESX Agents. This resource pool is created under the cluster for which you deploy IPS service. It is recommended that you do not attempt to change the settings for these Sensors grouped in the VMware ESX Agent resource pool.
- Security Group in NSX: This is a group of VMware objects for which the IPS service is provided. For example, if you include a distributed port group, the Virtual IPS Sensor monitors traffic related to all the VMs connected to this distributed port group. Similarly, if you select a cluster, the Virtual IPS Sensor monitors traffic related to all the VMs in the cluster. You define this security group through vCenter.
- Security Policy in NSX: You define this policy in NSX using the vCenter web client. You define the required security services in a security policy. Then you apply this policy to a security group in NSX (described above).

A security policy has three sections:

- Guest introspection services - You define the VM-level security services such as anti-virus, vulnerability management, data security, and so on in this section. These features are native to NSX and not provided through OSC. So, this section is not relevant for IPS service.
- Network Introspection Services - This section relates to services such as IPS and IDS through external solutions. So, you define the details related to IPS service in this section and apply it to the required security group.

In the security policy, you specify the following:

- You must select the required distributed appliance from the security name list. The distributed appliances that you defined in OSC are listed under service name in vCenter.
- Security profile. These are the IPS policies that are currently defined in the corresponding Manager. When you create a distributed appliance, OSC collects the IPS policy names from the Manager specified in the Security Manager Connector. OSC provides this list of policies to NSX, which is listed under Profile in vCenter. So, the IPS policies in the Manager are the same as that of security profiles in NSX/vCenter.
- Source and destination of traffic based on which the selected security profile (IPS policy) must be applied. This option is to enable you to specify the inbound or outbound direction.

You should select the same profile (IPS policy) for both inbound and outbound or specify different policies.

- Policy groups in the Manager: The policy group in the Manager is the collection of next-generation IPS policies defined in the Manager. For example, a policy group contains definitions for exploit attacks, recon attacks, DoS attacks, and advanced malware. When you create a distributed appliance, OSC provides these policy groups from the Manager to NSX. Then, vCenter displays this list under security profile when you configure network introspection services in the security policy in NSX.

High-level steps to implement a security service

Any security service through OSC is a collaboration between vCenter, NSX, the security service manager, and the virtual security service appliances that are orchestrated by OSC. To illustrate what this means, consider an SDDC as shown here. Assume that you want to provide a security service such as IPS to VMs 1 through 4. vCenter and NSX are up and running on VMware ESXi-3.

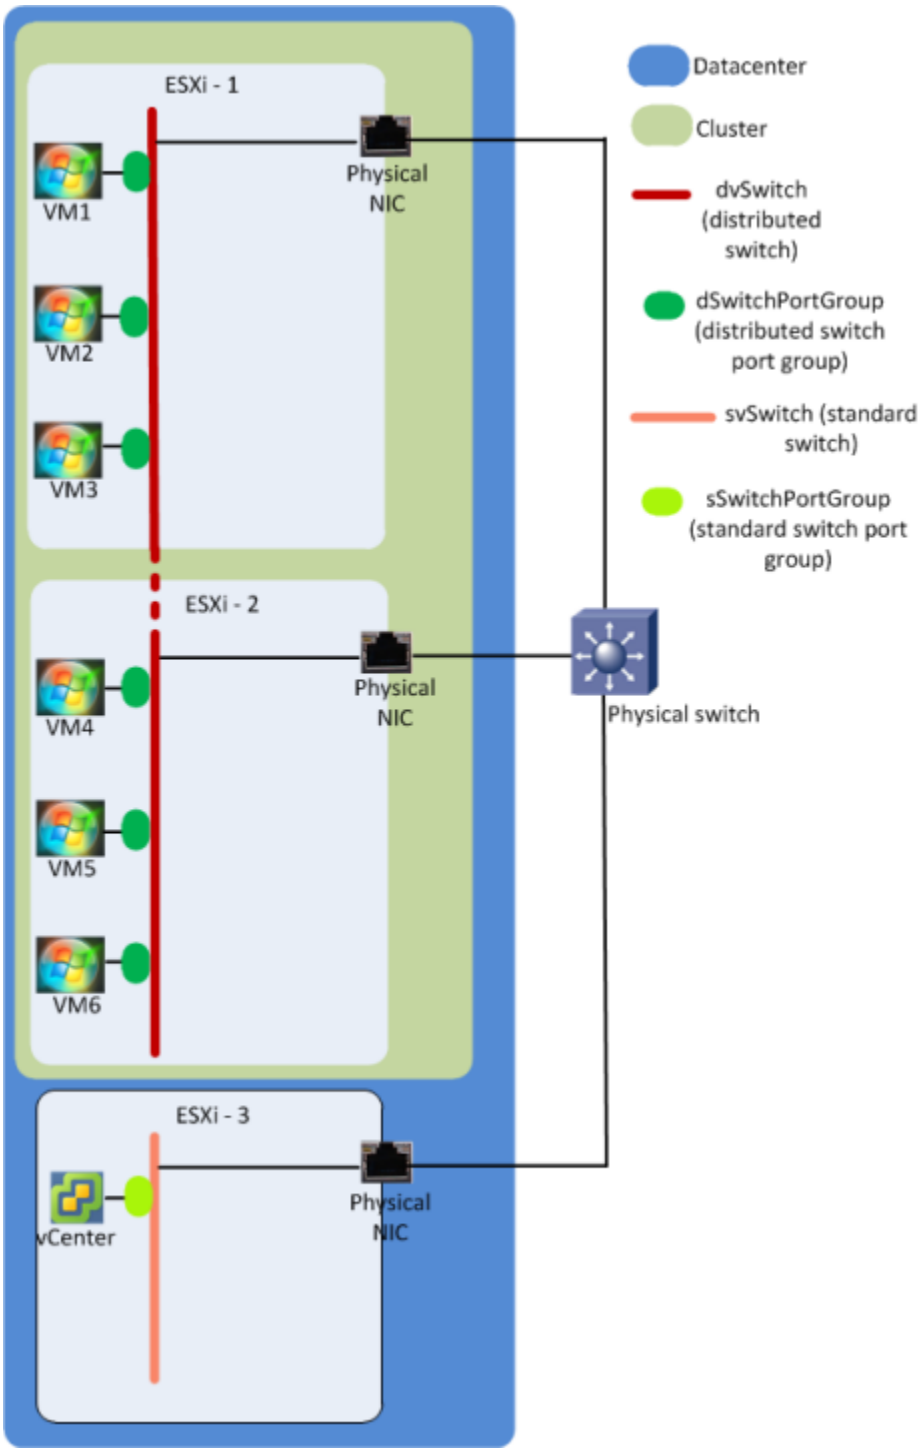

Figure 3-4 An example of SDDC

The following are the high-level steps for configuring the security service using OSC. There are multiple orders by which you can complete this configuration. The steps provided here are in the recommended order.

1 When you deploy the security service to a cluster, OSC collaborates with NSX to install a virtual security service appliance in each host of that cluster. Each of these virtual security service appliances needs a management port IP address. In this example, there are two hosts in the cluster which means that you need two IP addresses for the management ports of these appliances.

   You define a pool of IP addresses in the NSX Manager. For each IP address pool, you also define the required network details such as the default gateway and DNS server IP addresses.

   The number of IP addresses in an IP address pool depends on the number of VMware ESXi hosts that you plan to include in the corresponding cluster. You must also factor in the VMware ESXi hosts that you might add to the cluster in future. When you add VMware ESXi hosts to the cluster after deploying a security service, NSX automatically installs an instance of the virtual security service appliance in those VMware ESXi hosts. NSX needs an IP address to assign to these virtual appliances. Consider that the cluster for which you plan to provide the security service currently has two VMware ESXi hosts. However, you plan to include 5 more VMware ESXi hosts later on. Though two IP addresses are currently sufficient, you need 5 more for the VMware ESXi hosts that you plan to add.

2 Complete the following configuration in OSC.

   a Create a virtualization connector by providing the IP address and logon credentials for NSX and vCenter. You can create as many virtualization connectors as you require. For example, you want to provide the same security service to multiple clusters managed by different vCenters. Then create virtualization connector for each vCenter. The same security policies are applied to all hosts in all these clusters when you implement the security service.

      In this example, one virtualization connector is sufficient since there is only one cluster managed by one vCenter. See Define virtualization connectors on page 114.

   b Create a manager connector by providing its IP address and admin logon credentials. See Define manager connectors on page 116.

   c Define the virtual security service appliance in OSC and import the required security service function images into OSC. See Change the software version of security appliances on page 121.

   d Create a distributed appliance using the virtualization connectors, manager connector, and the required virtual security service image from the previous 3 steps. When you associate a virtualization connector with a security manager domain, a Virtualization System is created in OSC. See Manage distributed appliances on page 124.

   e Deploy the Virtualization Systems that you created in the previous step in the relevant clusters. If there are multiple clusters that you want to protect, you must deploy the Virtualization System separately for each cluster. Then, OSC collaborates with the corresponding vCenter and NSX to deploy Virtual IPS Sensors in each VMware ESXi host in the specified cluster. See Deploy virtual systems on page 127.

3 Complete the following configuration in NSX using vCenter web client.

   a NSX needs to know the IP addresses of the VMs to be protected for it to route the traffic for the network introspection service (that is, to the Virtual Security System instances). For NSX to know the IP addresses, VMware tools must be running on the VMs. If VMware tools is not running, you must include the IP addresses of such VMs in a security group. Before you create the security group, create an IP set object containing the IP addresses of VMs on which VMware tools is not running.

   b In the **Networking & Security** tab of vCenter, create a security group and add the VMs that you want to protect. In our example, you include VMs 1 through 4 in the security group. See Create a security group in VMware NSX on page 140.

   c Create a security policy and in the **Network Introspection Services** step, select the corresponding distributed appliance and the security service policy for both inbound and outbound traffic.

   The distributed appliances are listed as **Service Name** and the security service policies are listed as **Profiles** in vCenter. See Create a security policy in VMware NSX on page 144.

   Apply the security policy that you created in the previous to the policy group created in step 1. See Apply a security policy to a security group in VMware NSX on page 147.

4 Complete the following in the security service manager.

   a Log on to the security service manager and verify whether the Virtual Security System is listed in the appropriate device list.

   b Check the status of the Virtual Security System in the security service manager. Also, in case of IPS, verify whether a signature set is present. If not, take appropriate measures to provide one. For example, in Network Security Platform, you must deploy pending changes to the Virtual Security System from the Devices tab. The pending changes are automatically updated to all individual Virtual IPS Sensors of that Virtual Security System.

   c If necessary, log on to the CLI of the virtual security service appliance and view the configuration.

5 To verify successful deployment, send sample attack traffic from one of the protected VMs and check if an alert is displayed in the security service manager.

6 By default, Virtual Security System instances are deployed in fail-open mode. You can configure a Virtual Security System to fail-close if necessary. See Configure Virtual Security System to fail-close or fail-open on page 149.

How the IPS service works

When you deploy a Virtual System from OSC, NSX creates one Virtual IPS Sensor per host in the specified cluster. This Virtual IPS Sensor provides IPS service for the protected VMs in that host.

1 When you create a security manager connector in OSC (step 2-2 in High-level steps to implement a security service on page 99), OSC gathers the policy group names from the Manager and provides them to NSX. This is how the Manager policy groups are available for you to select when you add a network introspection service in a security policy in NSX.

2 As described in step 2-3 in High-level steps to implement a security service on page 99, you import the required Virtual IPS Sensor software images into OSC. OSC provides these image files to NSX to install the Virtual IPS Sensors in the VMware ESXi hosts.

3 When you create the Virtualization System, you select the Manager admin domain. OSC uses this information to align the Virtual IPS Sensors under the selected admin domains in the Manager.

4 When you deploy the Virtualization System, OSC orchestrates the installation of the Virtual IPS Sensors in the corresponding VMware ESXi hosts through NSX. The following actions are also completed automatically.

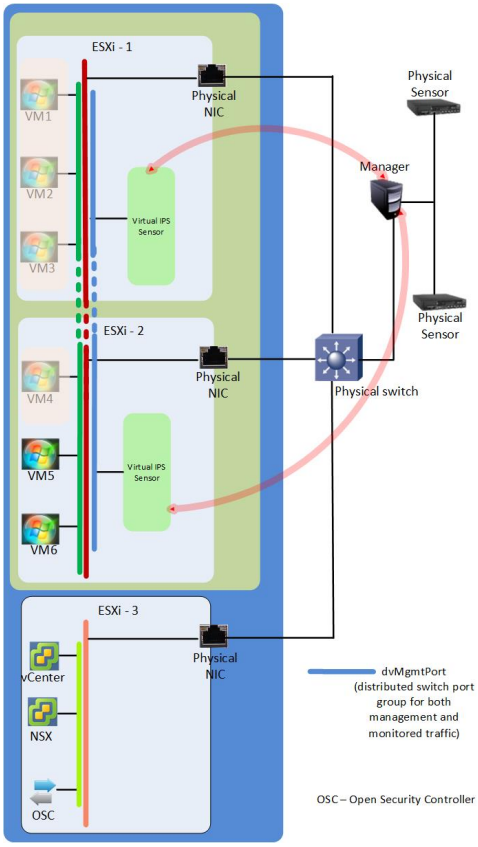

Figure 3-6 SDDC with IPS service through OSC

- OSC creates the corresponding Virtual Security Sensor objects in the Manager and under the selected admin domain. Recall that a Virtual Security Sensor is the logical container in the Manager for the deployed Virtual IPS Sensors. Virtual Security Sensor is comparable to the *Sensor failover* object created in the Manager when you pair up Sensors for failover. Virtual Security Sensor in the Manager is the equivalent of Virtual Systems in a Distributed Appliance in OSC.
- The Virtual IPS Sensors automatically establish trust with the Manager specified in the Security Manager Connector.
- The internal monitoring ports are automatically configured for operation.
- When you create or delete a new policy group in the Manager, OSC automatically updates NSX so the changes reflect in NSX as well.
- When you add an VMware ESXi host to the cluster, NSX automatically installs the Virtual IPS Sensor on this new VMware ESXi host and aligns the Virtual IPS Sensor under the corresponding admin domain in the Manager.
- If you change the IPS configuration, you must deploy the changes to the Virtual Security Sensor. OSC makes sure that these changes are deployed in all Virtual IPS Sensors of that Virtual Security Sensor.

In our example, VMs 1 to 4 are included in the security group of NSX. That is, these are VMs for which IPS service is provided currently. If VM1 communicates with VM2, NSX directs this traffic to pass through the inline Virtual IPS Sensor on VMware ESXi-1. Any traffic originating from or destined to VMs 1 to 3 pass through this Virtual IPS Sensor.

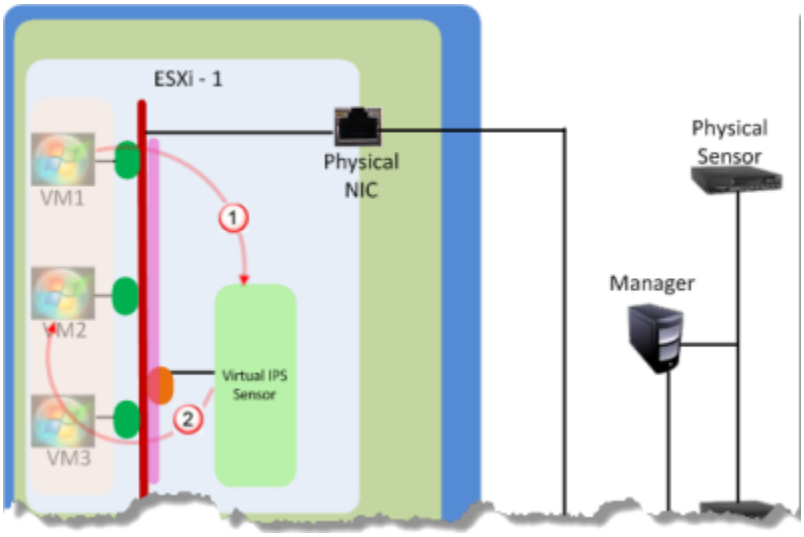

Consider VM1 communicates with VM4, which is on ESXi-2. The Virtual IPS Sensor on VMware ESXi-1 applies the outbound policy group on this traffic before it exits out of the host. When it reaches VMware ESXi-2, the Virtual IPS Sensor on this host applies the inbound policy group before forwarding it. So, if both these Virtual IPS Sensors detect an attack in this traffic, they raise separate alerts in the Manager. The corresponding Virtual IPS Sensors take the required response action as configured in the policies.

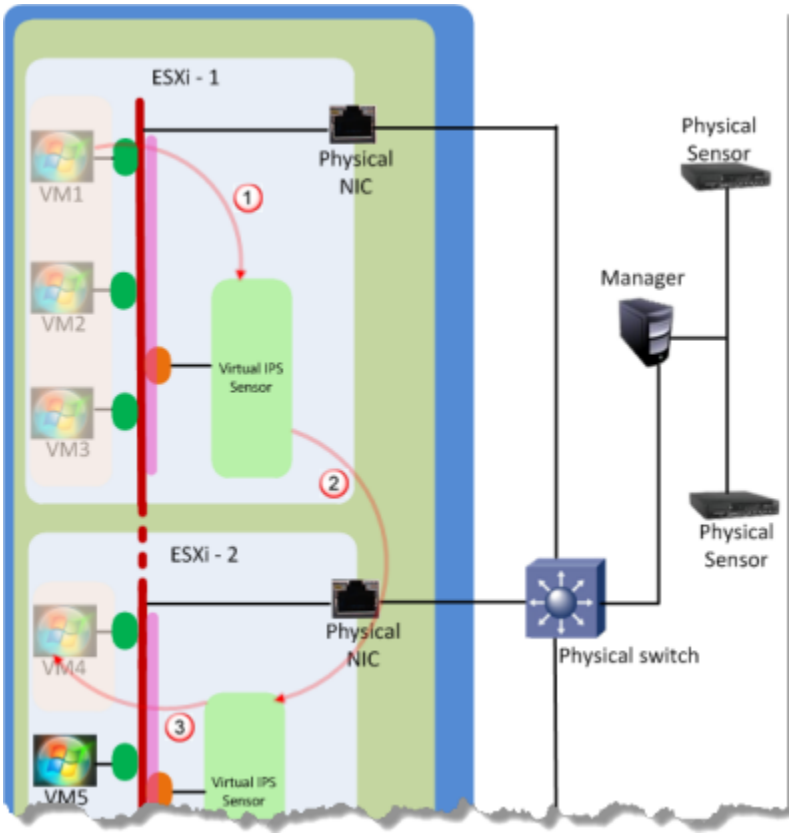

The Virtual IPS Sensor does not receive the traffic if two VMs not belonging to the security group communicate with each other.

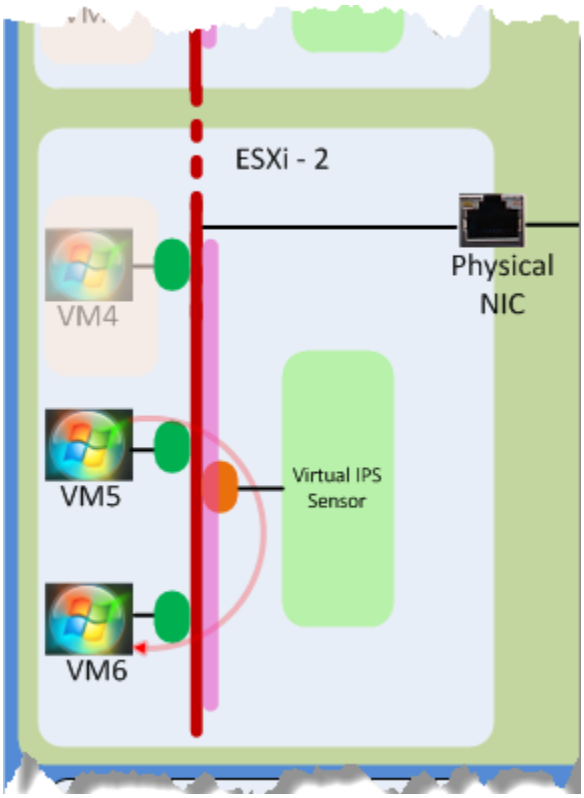

Requirements for deploying IPS service

Make sure that you meet the following requirements before you deploy IPS service using OSC.

- The hosts must be running on VMware ESXi version 6.0 or later.
- You can deploy IPS service only at a cluster level. So, make sure that you have *clustered* the hosts for which you want to provide IPS service. You must create a cluster even if you want to deploy IPS service to only one host.
- You have VMware vCenter version 6.0 or later installed (corresponding to the VMware ESXi version).
- You have VMware NSX version 6.2.4 or later installed.
- All hosts in the clusters have network virtualization components installed (Host Preparation).
- OSC requires distributed virtual switches of version 5.5. Standard virtual switches are not supported.
- If the corresponding clusters contain more than one VMware ESXi host, you must set up an NFS datastore for those clusters. In case of clusters with more than one VMware ESXi hosts, the Virtual IPS Sensors are installed only in NFS datastores. If a cluster contains only one VMware ESXi host, a VMFS datastore will suffice.
- You have Manager 8.3.7.44 or later installed.
- You have imported the Virtual IPS Sensor software images into OSC.

- OSC can communicate with the vCenter, NSX, and the Manager.
- VMware tools must be running on the protected VMs. This enables the NSX (NetX) APIs to redirect traffic to the Virtual IPS Sensor on the host. If VMware tools are not installed on the VMs, an alternative is to create a grouping object based on IP sets in NSX. Then, you must add this grouping object in the security group.

Considerations

Review the following before you deploy IPS service using OSC.

- Because OSC currently supports only IPv4 addressing, the Manager, IP pool for Virtual Security Systems, NSX, and vCenter Server must all have an IPv4 address. Also, the products with which you plan to integrate the Virtual Security System instances must also have an IPv4 address.
- MDR is not supported currently. So, you cannot use a Manager, which is part of an MDR.
- Only inline mode is supported. You configure inline fail-open and fail-close modes not on the Sensor but through an NSX mechanism.
- Failover of Virtual Security System instances is not supported.
- IPS-VM100-VSS Virtual Security System model is similar to the IPS-VM100 Virtual IPS Sensor with respect to Sensor performance and capacity values. See the *Network Security Platform Best Practices Guide* for performance and capacity values of IPS-VM100.

IPS features supported by Virtual Security System

The following are the list of IPS features supported by Virtual Security Systems in this release:

- Inline fail-open and fail-close are through an NSX mechanism
- To apply the policies, you must use Policy Groups in the Manager
  - IPS policies
  - Advanced Malware policies
  - Inspection Options policies
  - Connection Limiting policies
  - Firewall policies
- Snort custom attack definitions
- McAfee custom attack definitions
- Protection of Web application servers
- Advanced Traffic Inspection
- Layer 2 passthru mode is supported but implemented differently when compared to physical Sensors
- Layer 7 data collection
- MPLS traffic inspection
- IPv6 traffic inspection
- HTTP response scanning
- Inspection of double VLAN tagged traffic
- Monitoring Sensor performance
- Synchronization of Sensor clock using an NTP server

- Display Sensor CLI audit log events in the Manager
- TACACS+ user in audit logs
- Secure Transfer of Files from Sensor CLI
- Application Identification and Visualization
- Advanced Traffic inspection
- SmartBlocking of attacks including use of IP Reputation to augment SmartBlocking
- Integration with McAfee GTI for IP reputation and file reputation. This includes protection from high-risk hosts.
- Inspection of X-Forwarder-For Header Information. Reputation lookup and quarantine of client IP addresses in the XFF header.
- Stateless Firewall access rules
- Granular access control for CLI commands (for TACACS users)
- Advanced callback detection including Bot Command and Control server activity detection
- Web server protection against DoS attacks
- Integration with McAfee Endpoint Intelligence Agent (McAfee EIA)
- Sensor autorecovery is done through a VM recovery feature.
- Syn cookie
- Latency monitor
- Traffic prioritization
- IP fragmentation flood
- Inspection of tunneled traffic including GRE tunneled traffic
- Passive device profiling
- Attacks using evasion techniques – supported but some combinations may not be detected.

IPS features - not supported

The following are the list of IPS features not supported by Virtual Security Systems in this release:

- SPAN Mode
- TAP Mode
- Sensor failover
- DoS
- DNS DoS protection
- SSL Decryption
- Rate Limiting
- ARP Spoofing protection
- IP Spoofing protection

- Virtualization of monitoring ports with VLAN (VIDS)
- Virtualization of monitoring ports with CIDR (VIDS)
- Remediation
- IPv6 traffic support on the Management port
- Jumbo Frame Parsing
- Monitoring Sensor Performance
- Netflow export to NTBA
- GTI IP Reputation Integration – SmartBlocking
- Packet Capture
- Simulated Blocking
- Cloud Threat Detection
- Network Forensics
- Passive Device Profiling
- Support for 256 SSL certificates
- Integration with Endpoint Intelligence Agent
- Traffic Prioritization with Application Content (Rate Limiting with App ID)
- DNS spoof protection feature (to prevent DDoS attacks)
- Fail-over
- Malware detection of files downloaded using HTTP range request (Split file download)
- Application Visualization per port
- Incident Generator

Prepare a VMware ESXi host for NSX

Before you begin

- You have successfully installed and configured the NSX Manager.
- You have admin privileges in the vCenter Server.

To protect the VMs through OSC and NSX, you must prepare the VMware ESXi host of those VMs. You use the Networking & Security options in the vCenter Server Web Client to prepare an VMware ESXi host.

Task

1 Log on to vCenter Server Web Client with administrator privileges.

2 Select **Home** | **Networking & Security** | **Installation** | **Host Preparation**.

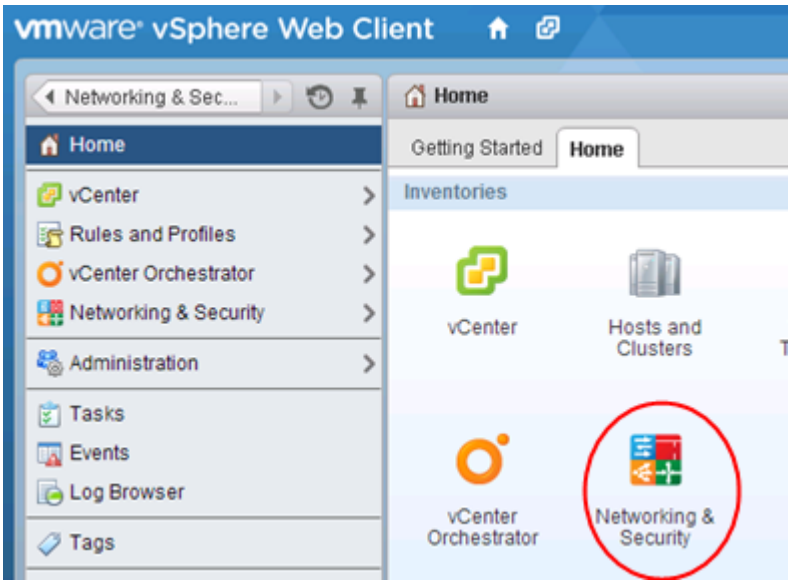

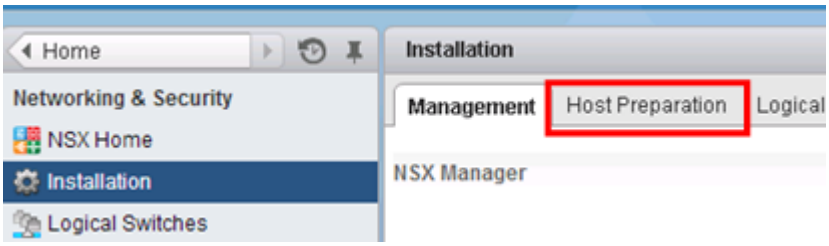

3 Select the required **NSX Manager**.

4 Click **Install** for the required cluster or VMware ESXi host.

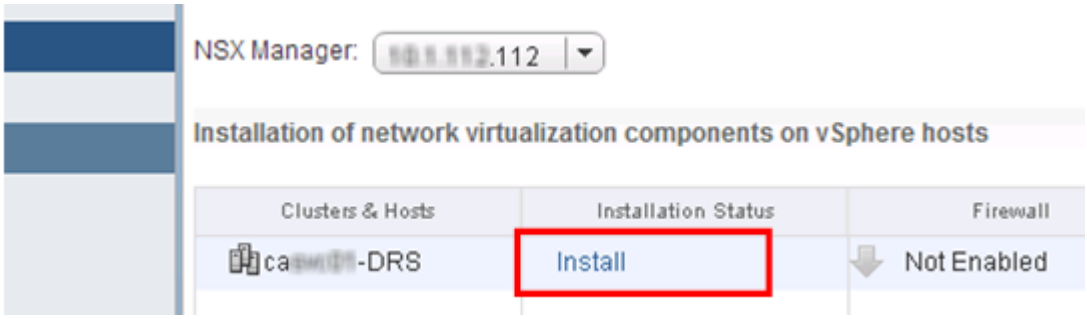

5 Click **Yes** to proceed with the installation.

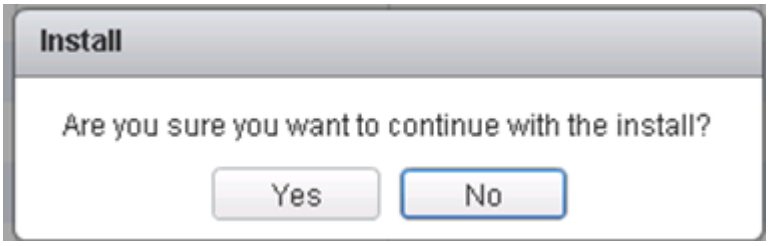

You can monitor the progress of the task.

The installation status column shows if the task completed successfully.

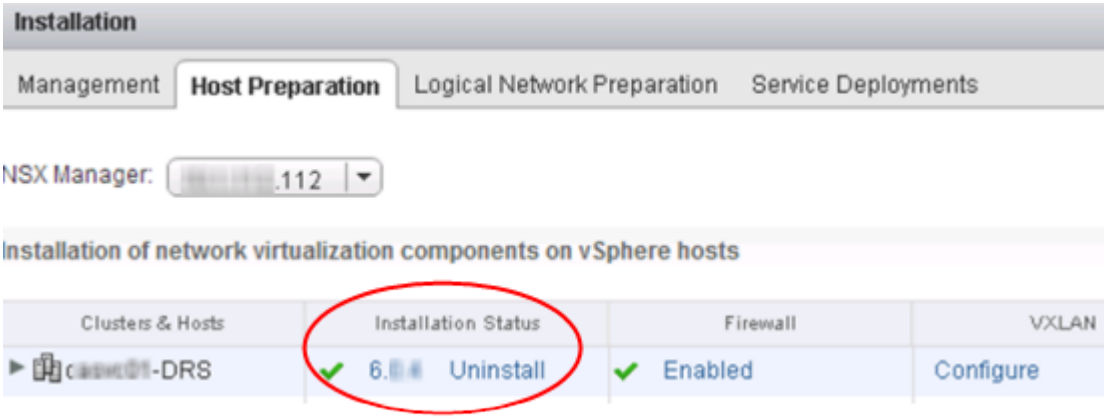

If the host preparation task fails, restart the corresponding VMware ESXi hosts and repeat the task.

Define an IP pool for virtual security appliances

If you have installed NSX, you can define the IP pool for the virtual security appliances. Along with the IP addresses, you also define other network settings such as the default gateway IP address and the subnet mask.

Currently, OSC supports only IPv4 addresses. So, the IP pool must contain IPv4 addresses only.

Task

1 Log on to vSphere Web Client as the root user.

2 In the vSphere **Home** tab, select **Networking & Security.**

3 Select **Networking & Security Inventory | NSX Managers.**

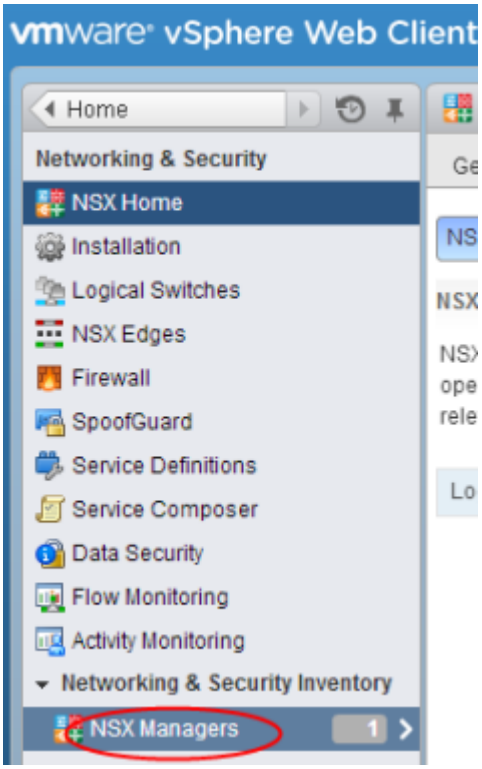

4 Select the NSX Manager in which you want to define the IP pool and then select **Manage** | **Grouping Objects** | **IP Pools.**.

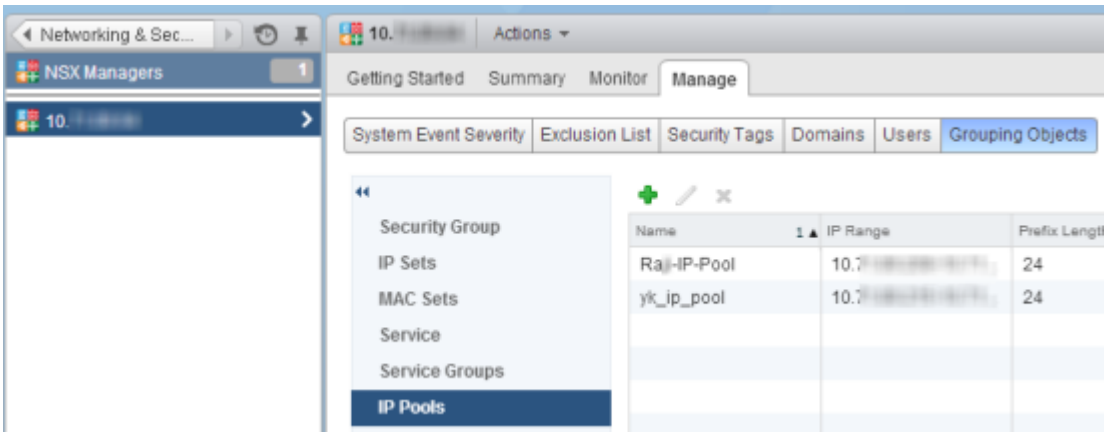

5 Click ✚ to add an IP pool.

6 In the **Add IP Pool** window, enter the details and click **OK**.

Table 3-2 Option definitions

Option	Definition
Name	Enter a relevant name for the IP pool.
Gateway	Enter the IP address of the default gateway for the IP addresses. After you create the IP pool, you cannot modify the default gateway IP address.
Prefix Length	Enter the network prefix length of the IP addresses.
Primary DNS	Enter the IP address of the primary or the preferred DNS server for the IP addresses.
Secondary DNS	Optionally, enter the IP address of the secondary DNS server.
Static IP Pool	Enter the range of valid IPv4 addresses. Make sure there is no IP address clash. That is, the IP addresses in the IP pool must not have been assigned to a network object or already included in a different IP pool.
OK	Click to save the settings and create the IP pool.
Cancel	Click to close the dialog box without saving the changes.

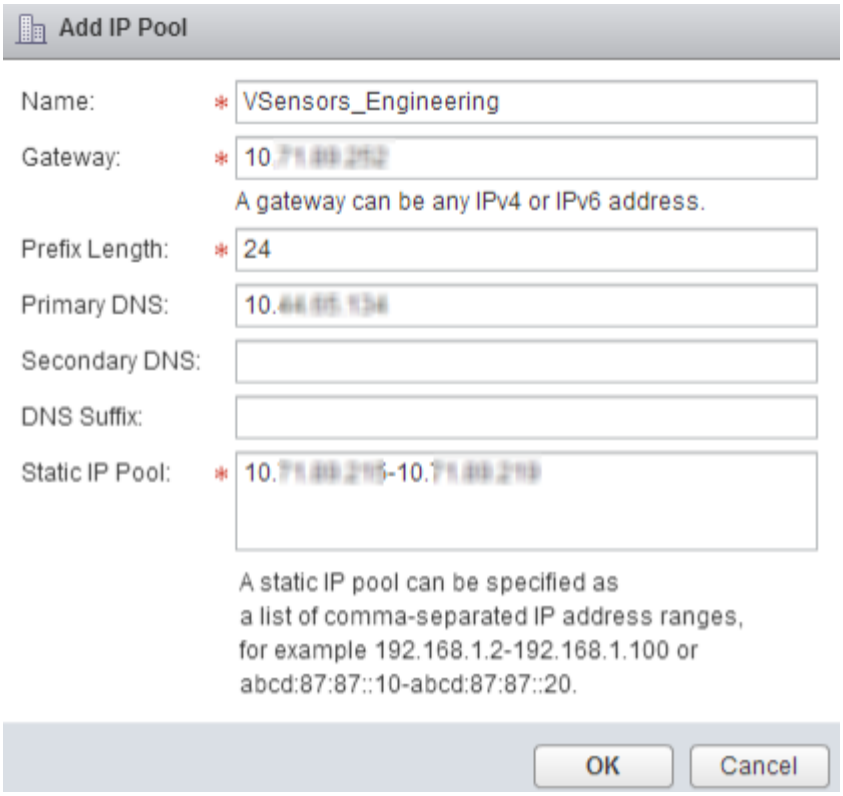

The summary view displays the count of IP addresses in an IP pool and the count of addresses in use.

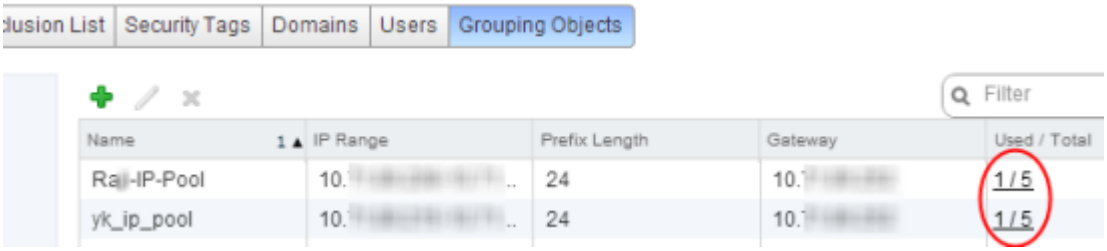

Define virtualization connectors

You are able to define virtualization connectors from the OSC web application.

Task

1 In the OSC web application, select **Setup** | **Virtualization Connectors**.

The **Virtualization Connector** page displays the currently available virtualization connectors.

Figure 3-7 Virtualization Connector page

Table 3-3 Option definitions

Option	Definition
Name	Name of the virtualization connector record.
Type	Virtualization provider which you mention when you create the virtualization connector. An example is VMware.
Controller IP	IP address of the virtual security controller such as VMware NSX.
Provider IP	IP address of the virtualization provider server. Clicking the hyperlink provided for the IP address opens the login screen of the virtualization provider. Certain virtualization providers are configured on specific ports of the server. You are required to enter the port at the end of the URL.

2 Take one of the following actions:

To create a new virtualization connector, click **Add** and enter the options in the **Add Virtualization Connector** dialog.

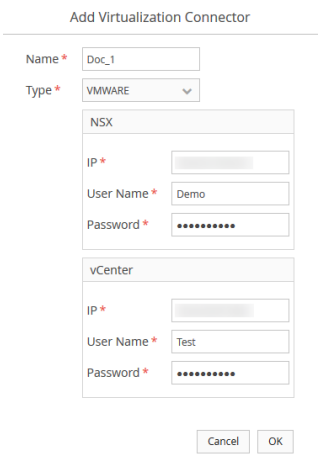

Figure 3-8 VMware virtualization connector

Table 3-4 Option definitions

Option	Definition
Name	Name that enables you easily identify a virtualization connector record.
Type	Virtualization provider from the list of currently supported providers. You are provided with following option: • VMware
Cancel	Closes the dialog without saving the changes.
OK	Closes the dialog box with the changes saved to the OSC database. A warning displays if OSC is unable to connect to virtualization provider using the IP address and credentials. You can still create the virtualization connector. However, if you use this virtualization connector in a distributed appliance, you cannot delete the distributed appliance or virtualization connector record. If you are using NSX and delete a virtualization connector, OSC deletes the related data from NSX. So, if OSC is unable to log on to the NSX defined in the virtualization connector, the task of deleting the virtualization connector fails.
	VMWare
NSX	
IP	IPv4 address of VMware NSX Manager Virtual Appliance.
User Name	Logon name of an admin user.
Password	Corresponding password.
vCenter	
IP	IPv4 address of VMware vCenter with which the NSX management service is connected.
User Name	Root admin user name of the vCenter.
Password	Corresponding password.

3 To edit a virtualization connector record, select the record and click **Edit**.

You cannot change the **Type**. After you complete making the changes, click **OK** to save the changes.

McAfee highly recommends that you do not change the IP addresses of the virtualization connector servers (the SDN and the virtualization provider) after deployment of virtual security system instances.

4 To delete a virtualization connector record, select the record and click **Delete**.

If the virtualization connector you want to delete is used in a distributed appliance, you must first delete the distributed appliance record. To delete a distributed appliances record, see Deleting a distributed appliance for VMware on page 133.

Define manager connectors

You are able to configure manager connectors from the OSC web application.

Task

1 In the OSC web application, select **Setup** | **Manager Connectors**.

The **Manager Connector** page appears, displaying currently available manager connectors.

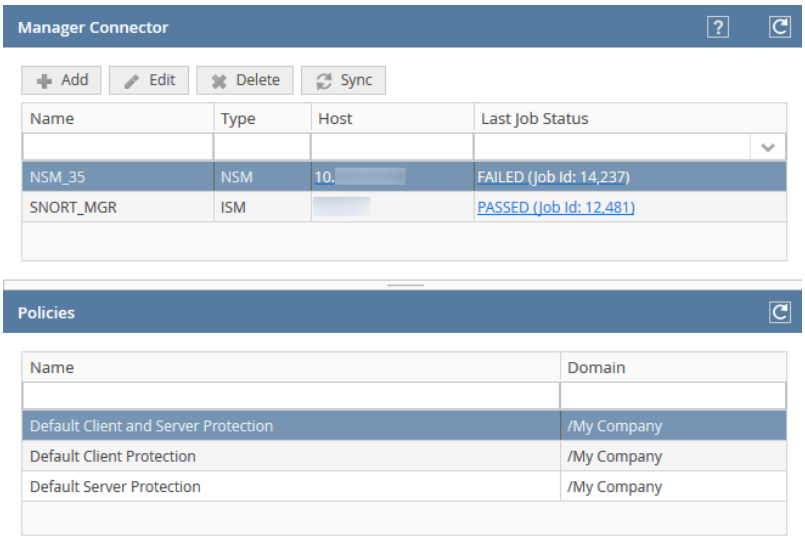

Figure 3-10 Manager Connector page

Table 3-5 Option definitions

Option	Definition
	Manager Connector
Name	Name of the manager connector record.
Type	Security service manager type which you select when you create the virtualization connector, such as Network Security Manager or SMC.
Host	IP address of the security service manager server. Clicking the hyperlink provided for the IP address opens the login screen of the security service manager. Certain security service managers are configured on specific ports of the server. You are required to enter the port at the end of the URL.
Last Job Status	Status of the most recent job. Clicking the hyperlink provided for the job ID, routes you directly to the **Jobs** page with tasks for the specific job displayed in the **Tasks** pane.
	Policies
Name	Name of the policy in the security service manager. Initially all default policies from the security service manager are linked with OSC. Different default policies are linked depending on the security service manager it is deployed in.
Domain	Name of the domain from which the policies are linked in the security service manager

2 Take the appropriate action:

To create a new manager connector, click **Add** and enter the options in the **Add Manager Connector** dialog.

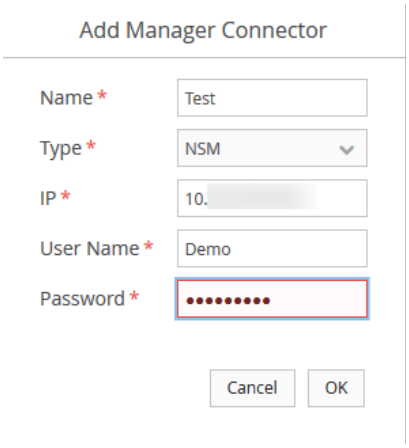

Figure 3-11 Add manager connector

Table 3-6 Option definitions

Option	Definition
Name	Enter a name, which enables you easily identify a manager connector record.
Type	Based on the type of security service, select the manager connector type. For IPS service, select NSM which refers to the Network Security Manager.
Cancel	Click to close the dialog without saving the changes.
OK	Click to close the dialog box with the changes saved to the OSC database. When you click **OK**, all the admin domains and policy groups available in the Manager are sent to OSC.
	NSM
IP	Enter the IPv4 address of the Manager server.
User Name	Enter the logon name for the Manager. This logon name must have **Super User** role assigned to it. The default logon name with **Super User** role is *admin*.
Password	Enter the corresponding password.
	SMC
IP	Enter the IPv4 address of the Management Server. There might be multiple IPv4 addresses configured to reach the Management Server. However, you must enter the one that is configured for listening.
API Key	A 24-character alphanumeric authentication key that is randomly generated when you configure the API Client in the Management Client. OSC uses this key to communicate with the SMC API.

3 To edit a manager connector record, select the record and click **Edit**.

You cannot change the **Type**. After you complete making the changes, click **OK** to save the changes. A job is started and the job number is displayed at the right-side bottom of the **Manager Connector** page. You can monitor the progress of this job in the **Jobs** page.

McAfee highly recommends that you do not change the IP address of the security manager after deployment of the virtual security system instances.

4 To delete a manager connector record, select the record and click **Delete**.

If the manager connector you want to delete is used in a distributed appliance, you must disassociate the manager connector from the distributed appliances before you delete the manager connector. Alternatively, you can delete the distributed appliance record. To delete a distributed appliances record, see Deleting a distributed appliance for VMware on page 133.

5 To synchronize any changes with the manager connectors, click **Sync**.

When there are any policy updates made in the security appliance manager, you can trigger manual synchronization to update the changes in the manager connector. A new information pop-up appears notifying you that the job has begun along with the job number. After the synchronization is complete, the **Last Job Status** for that instance changes from **RUNNING** to **PASSED** or **FAILED** depending on the result.

Manage software images for security appliances

To provide a security service, OSC orchestrates the installation of the corresponding virtual security appliance on designated hosts. For this, it is essential to follow this sequence of steps for the reasons explained beside the step:

1 Import relevant software images into OSC – Software images imported to OSC will be used to deploy security service on designated assets which can be located on any of the supported virtual environments.

2 Create a distributed appliance, in OSC – This is where you specify the model and software version combination for that security service.

- For security service function such as IPS, you select the virtual security system (Virtual IPS Sensor) model and software version, manager connector, and virtualization connector.

When you later deploy either of the security services, the SDN controller installs the corresponding model and software version of the security appliance in each designated instance.

You use the **Service Function Catalog** page in OSC to import software images for security appliances. When you import a particular software image, the model and version information are automatically populated. You can import multiple software image versions for each virtual security appliance model. You can then upgrade or downgrade the software image for deployed appliances.

Define service function catalog

Before you begin

- You have admin rights in the OSC web application.
- The .zip file containing the virtual security appliance software image is available from your endpoint.

You use the **Service Function Catalog** page to manage the software for security service deployments. For example, you can manage the Virtual IPS Sensor software images for IPS.

Managing the security function catalog consists of just one step.

- Import the required software images corresponding to that model. Using the example of IPS, IPS-VM100-VSS model and version of that model is added by default. You can several software versions of the same appliance model. For example, you can add 8.1 and 8.2 software versions of IPS-VM100-VSS in the security function catalog.

After you add software images to the security function catalog, you are able to select a specific image for each virtualization system at the time when you create distributed appliances. For existing distributed appliances, you are able to modify the choice of software image for that security service. You are also able to also downgrade the software image in the same manner.

The **Service Function Catalog** page provides information about the existing security service deployments. You can view information or delete the security service deployments.

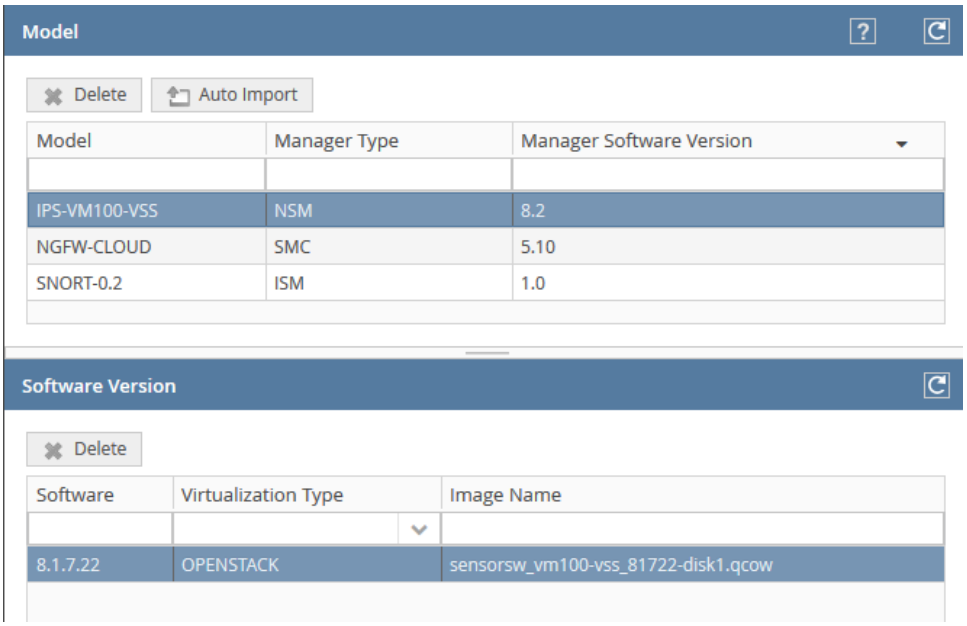

Figure 3-12 Security service function catalog

Option	Definition
	Model
Model	Displays the name of the security service Manager
Manager Type	Displays the type of Manager
Manager Software Version	Displays the Manager software version installed
	Software Version
Software	Displays the software version of the security service installed
Virtualization Type	Displays the name of the type of virtualization deployed
Image Name	Name of the security service image installed

Task

1 In the OSC web application, select **Setup** | **Service Function Catalog**.

2 Click **Auto Import** in the **Model** section.

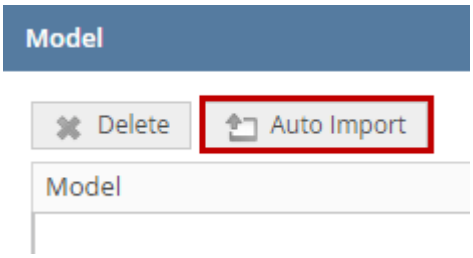

Figure 3-13 Auto Import option to upload software image files

The **Auto Import Appliance Software Version** pop-up appears.

3 In this pop-up, click **Choose File** and select the zipped virtual security system image file.

4 Click **OK** to begin uploading this file.

You cannot edit any records for a model.

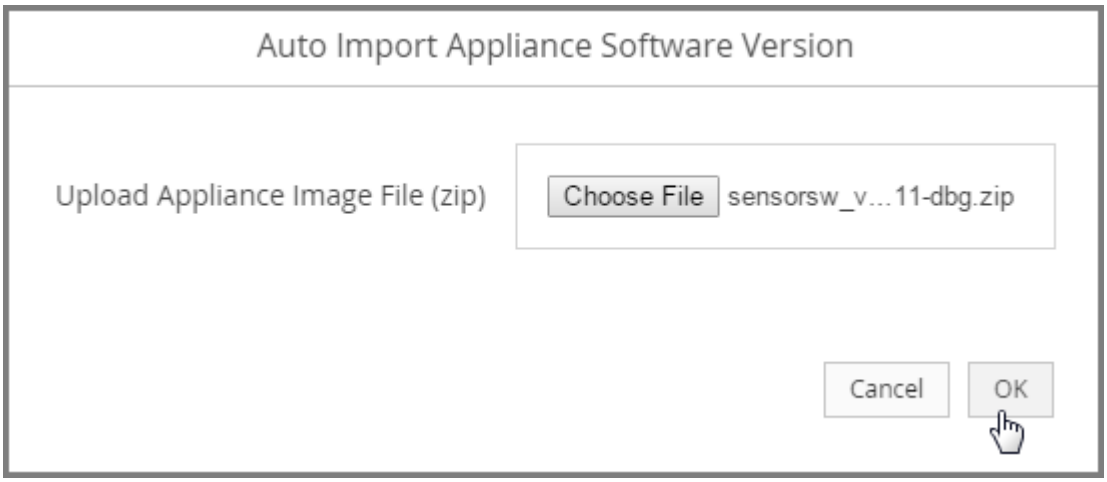

Figure 3-14 Auto Import File Select Pop-up

A progress bar appears providing the status of the file upload. At the end of the upload, it validates the image file before applying it.

Change the software version of security appliances

Before you begin

- Using the security function catalog, you have successfully imported the required software image in OSC web application. Importantly, make sure that you specified the software version of the image correctly. See Define service function catalog on page 119.
- You have the required access to deploy security services in NSX.

Before or after deploying a security service, you might want to change the software version of the corresponding security appliances. For example, after you deploy IPS service, you might want to upgrade the software image of the virtual security system instance. Alternatively, you can downgrade to an earlier version as well.

After you import the required software images in the security function catalog of OSC, you can change the image to the required version.

- After you change the version of security appliances, you must resolve installation status of the service deployment in NSX. NSX then deletes the current security appliances from the datastore and installs the version that you selected. It is evident that security appliances are not actually upgraded or downgraded but replaced with new instances installed with a selected software version. You have to check the **Jobs/Tasks** status to make sure that the process is complete without failures.
- In case of some security service functions such as IPS, the security service manager is unaware that the existing instances of the virtual security system are deleted and new ones are installed. The security service manager considers that the version of the virtual security systems has changed and trust is re-established. The name of virtual security system and the instances remain unchanged in the security service manager. The IP addresses assigned to the deployed security appliances also remain unaffected.
- After you change the software version, you must resolve the installation status of the corresponding service deployment in NSX.
- Until the new instances of the security appliances are fully up and running, the security service is suspended.
- Because security appliances are installed and not upgraded or downgraded, you can switch to a different version regardless of the current functional state of the appliances.

Task

1 In the OSC web application, select **Setup** | **Distributed Appliances**.

2 Select the required distributed appliance record and click **Edit**.

3 In the **Edit Distributed Appliance** dialog box, select the required security model-version combination from the **Service Function Definition** drop-down list and click **OK**.

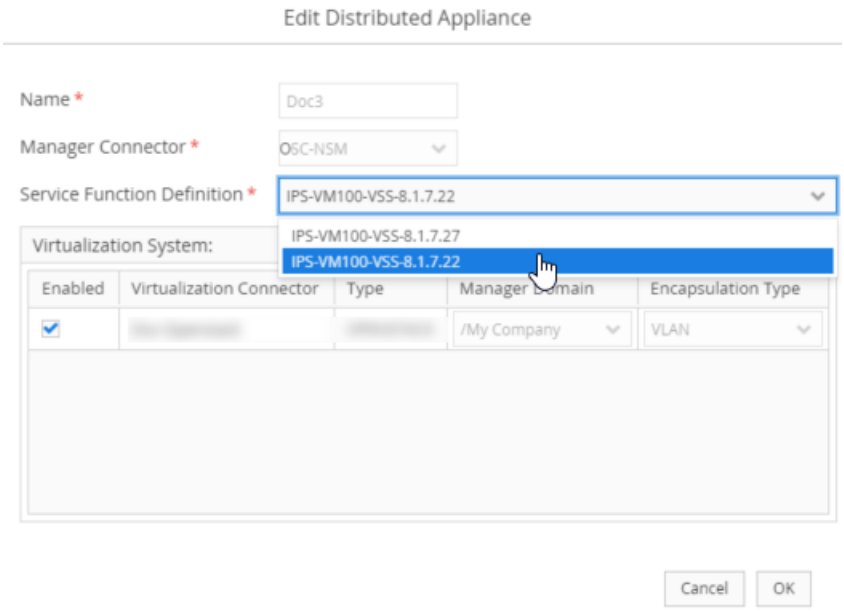

Figure 3-15 Select the required version for the security appliance

4 Review the warning message and click **OK** to proceed with the version change.

> Warning messages disappear from the screen automatically if you move your mouse or click a key on your keyboard. However, they remain on the screen if you do not perform either of these actions.

NSX stops the corresponding virtual security appliances to install the selected version of the security appliances. So the installation status of the virtual security appliances is now *Failed*.

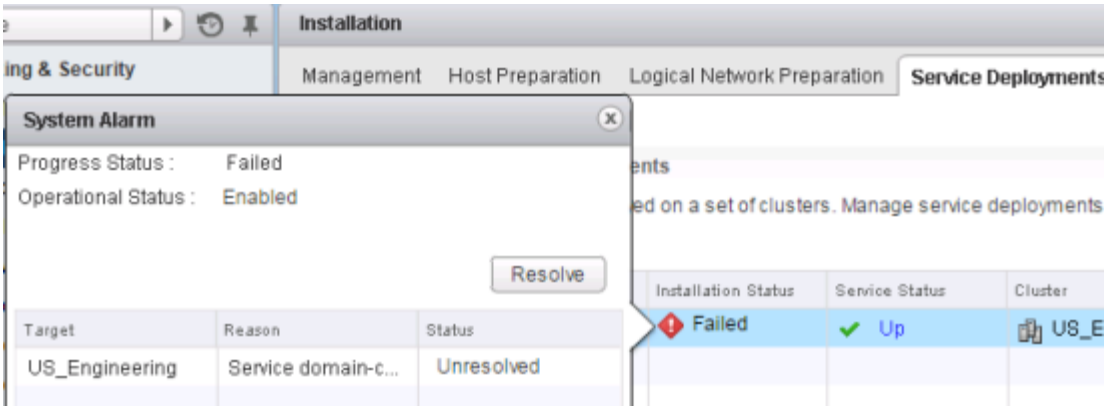

Figure 3-17 Installation status of service deployment

5 Log on to vCenter and resolve the installation status for the corresponding service.

The status on this will be one of four possibilities:

- *Unknown* – No information is available.
- *Down* – Includes an error message about the health of the security service. The most relevant indicators in this context are *Discovery* and *Inspection-Ready*. When the status is Down, it implies that both indicators are *False* and that you must investigate your deployment.
- *Warning* – Implies that the security service discovery was complete, but is not inspection-ready.
- *Up* – Implies that both indicators are positive.

6 After the **Installation Status** turns to *Succeeded* and the **Sevice Status** displays *Up*, deploy the configuration changes to the virtual security system from the Manager.

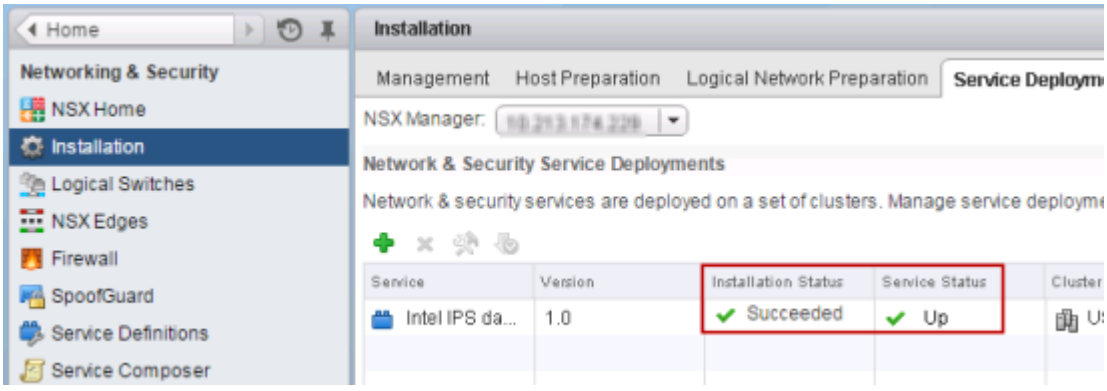

Figure 3-18 Installation status of service deployment

When you deploy the configuration changes, the Manager pushes the signature set to the virtual security system instances.

7 When you deploy configuration changes, the *Propagating Manager File...* job is triggered in OSC. Select **Status** | **Jobs** and make sure that *Propagating Manager File...* job is passed.

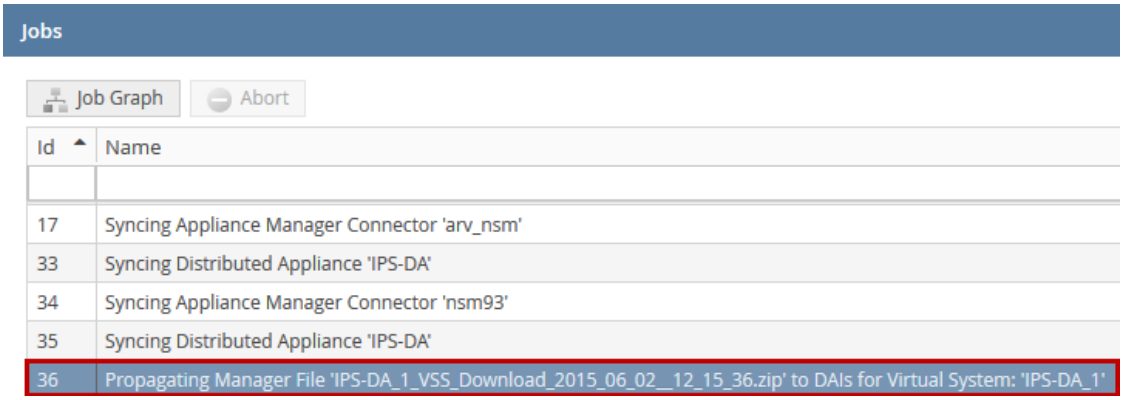

Figure 3-19 Propagating Manager File....

If the *Propagating Manager File....* job fails, deploy the configuration changes again from the Manager.

8 In the security manager, make sure the virtual security system and its member instances are connected and up-to-date.

9 Select **Status** | **Appliance Instances** and make sure the **Discovered** and **Inspection-Ready** for the corresponding appliances are *true*. If not, the security service function is suspended.

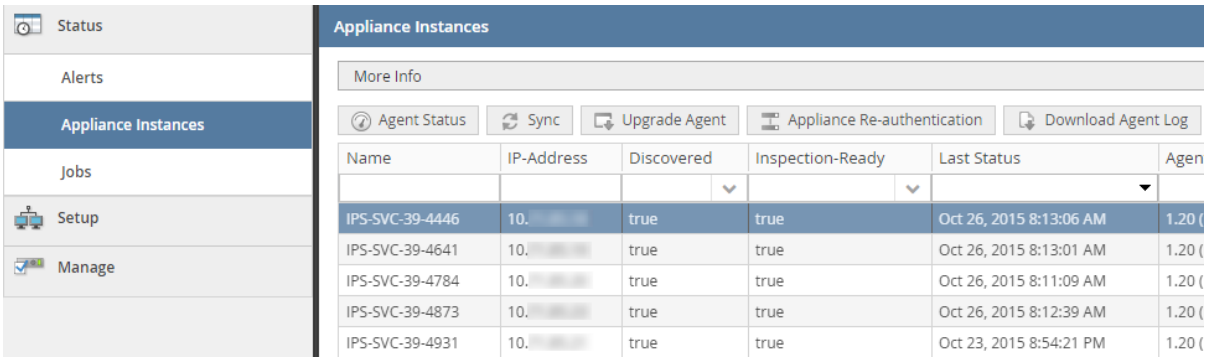

Figure 3-20 Appliance instances status

Manage distributed appliances

Before you begin

- The required virtualization connectors, the manager connector, and the security-appliance image file are available in OSC.
- VMware vCenter, NSX Manager, and the Manager are all up and reachable to OSC.

You manage distributed appliances from the OSC web application.

Task

1 In the OSC web application, select **Setup** | **Distributed Appliances**.

The **Distributed Appliances** page appears, displaying available distributed appliances.

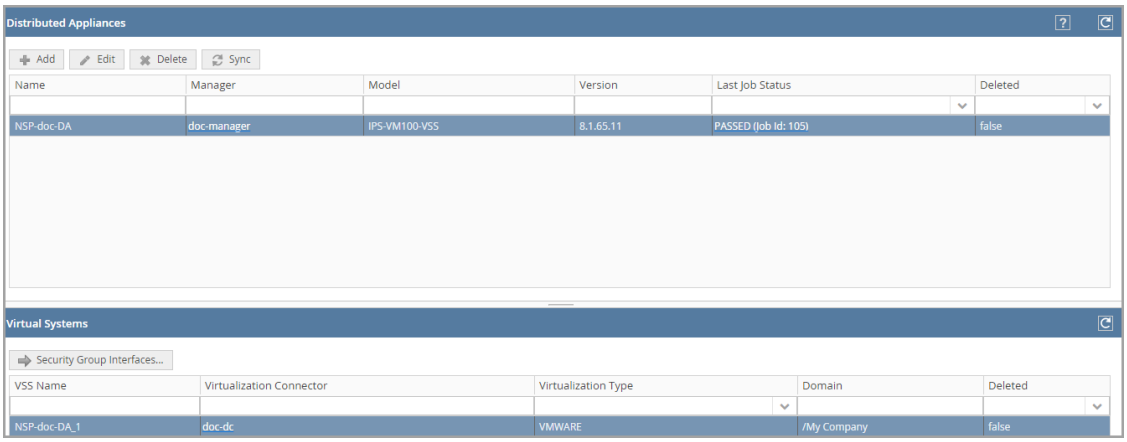

Figure 3-21 Distributed Appliances page

2 To create a new distributed appliance, click **Add** and enter the options in the **Add Distributed Appliance** dialog.

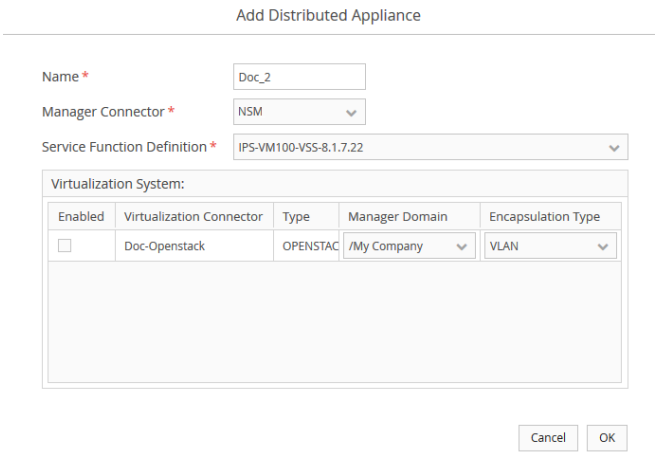

Figure 3-22 Add Distributed Appliance dialog box

Table 3-7 Option definitions

Option	Definition
Name	Enter a relevant name for the distributed appliance record.
Manager Connector	Select the manager connector for the distributed appliance. For IPS, select the manager connector that refers to a Network Security Manager.
Service Function Definition	Select a corresponding security appliance. This list of security appliances is from the security function catalog. So if you are unable to find a specific appliance definition, add it in the appliance catalog.
Enabled	Click to select a particular virtualization connector.

Table 3-7 Option definitions *(continued)*

Option	Definition
Virtualization Connector	This is the list of all added virtualization connectors. You can select multiple virtualization connectors. However, there is only one manager connector per distributed appliance. That is, you can map multiple virtualization connectors with one manager connector. If you select a virtualization connector, which is already selected in a different distributed appliance, that virtualization connector is automatically disassociated from the earlier distributed appliance.
Type	This is the virtualization provider corresponding to a virtualization connector.
Manager Domain	This is the list of admin domains from the security manager which you specified in the corresponding manager connector. • When you select the manager connector, OSC displays the current admin domains from the corresponding security manager. • When you select an admin domain, OSC creates the Virtualization System under this admin domain in the security manager. This is called as VSS in the Manager and is the logical container of the Virtual IPS Sensors installed in each VMware ESXi host. A VSS is similar to a failover Sensor object in the Manager.
Cancel	Click to close the dialog without saving the changes.
OK	Click to close the dialog with the changes saved to the OSC database. When you create a manager connector, all the admin domains and policy groups available in the Manager are sent to OSC. So, when you click **OK** while creating the distributed appliance, OSC gathers the current policy groups from the Manager and provides this list to the NSX Manager. This is how the Network Security Platform policy groups are available in the vCenter as profiles. You select these profiles when you create a security policy in vCenter.

- You cannot edit if the distributed appliance record is deleted. This is indicated by the status mentioned in the **Deleted** column for a record.
- To delete a distributed appliances record, see Deleting a distributed appliance for VMware on page 133.

3 To edit a distributed appliance record, select the record and click **Edit**.

4 A job is started and the job number is displayed at the right-side bottom of the **Distributed Appliances** page. You can monitor the progress of this job in the **Jobs** page.

You cannot change the **Name**, the **Manager Connector**, and the **Manager Domain** options. After you complete making the changes, click **OK** to save the changes.

Tasks

- *Deploy virtual systems* on page 127
- *Maintaining virtual appliance instances* on page 130
- *Deleting a distributed appliance for VMware* on page 133

Deploy virtual systems

Before you begin

- You have created the distributed appliance successfully.
- VMware vCenter, NSX Manager, and the Network Security Manager are all up and can be reached by OSC.
- If the cluster contains more than one VMware ESXi host, you must set up an NFS datastore to deploy the virtual systems. In case of clusters with more than one VMware ESXi host, virtual security service appliances are installed only in NFS datastores. If a cluster contains only one VMware ESXi host, a VMFS datastore will suffice.
- You have prepared the VMware ESXi hosts in the cluster for NSX.
- You have created a distributed switch port group for the virtual security service appliance management ports. Through this switch port group, the virtual security service appliance must be able to communicate with the Network Security Manager, OSC, NSX, and vCenter Server.
- You have created the IP address pool to assign IP addresses for the virtual security service appliance management ports. See Define an IP pool for virtual security appliances on page 110.

When you create a distributed appliance, a virtual system (virtual security system) record is automatically created. The virtual security system is visible in the security service manager, in OSC, and in NSX as a security service. You can then deploy the virtual system as a security service from NSX.

When you deploy a virtual system, OSC collaborates with vCenter, NSX, and the Network Security Manager to deploy the virtual security appliance in all hosts (hypervisors). In the case of IPS service, for example, NSX installs a virtual security system instance (a virtual IPS Sensor) in each VMware ESXi host of the cluster. These virtual security system instances are automatically assigned network details and have established trust with the Network Security Manager.

Task

1 In the OSC web application, select **Setup** | **Distributed Appliances**.

The **Distributed Appliances** page displays the currently available distributed appliances.

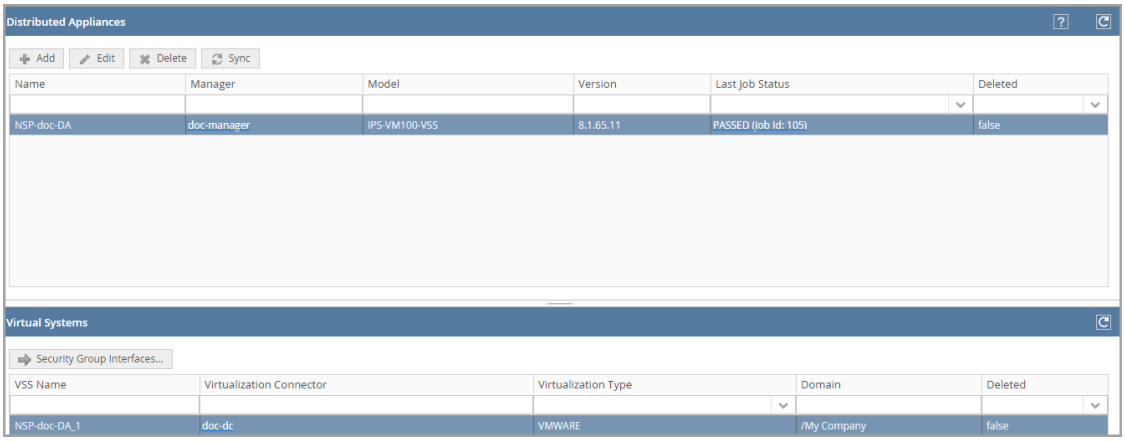

Figure 3-23 Distributed Appliances page

2 Select the required distributed appliance you created and ensure the following.
- The **Last Job Status** shows *Passed*.
- The virtual system is created and listed in the **Virtual Systems** section of the page.

3 In the Network Security Manager, make sure that the virtual system is automatically added.

4 Log on to vSphere Web Client as the root user.

5 In the vSphere **Home** tab, select **Networking & Security** | **Installation** | **Service Deployments**.

Figure 3-24 vSphere home page

6 From the **NSX Manager** list, select the required NSX Manager.

7 Click ✚ to create a service deployment.

The **Deploy Network & Security Services** wizard opens.

8 In the **Select services & schedule** step, select the service named after the distributed appliance you created.

For example, if the distributed appliance you created is *DA_N_America*, the corresponding service is *Intel IPS DA_N_America.*

Figure 3-25 Select services & schedule

9 If you want to deploy the virtual system now, select **Deploy now** and then click **Next**. Else, select the date and time for deployment and then click **Next**.

10 In the **Select clusters** step, select the required data center and the cluster, for which you want to provide IPS service.

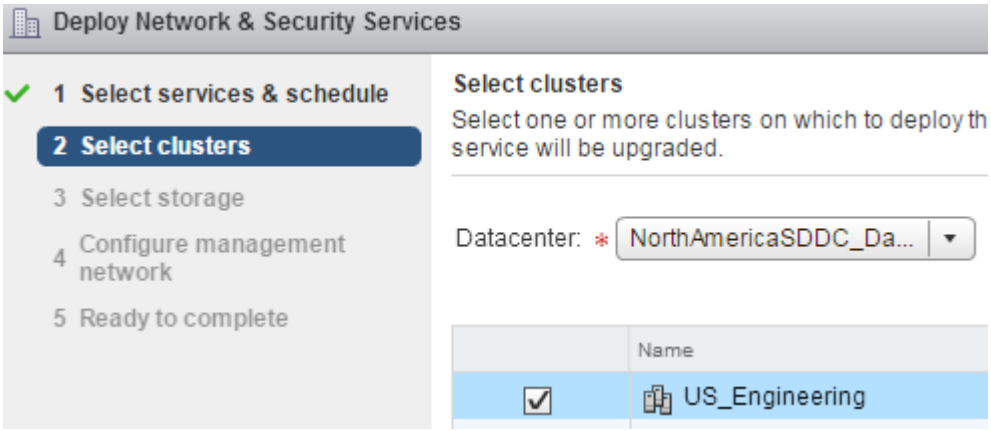

Figure 3-26 Select clusters

11 In the **Select storage** step, select the required datastore.

If the corresponding cluster contains more than one VMware ESXi host, you must select an NFS datastore.

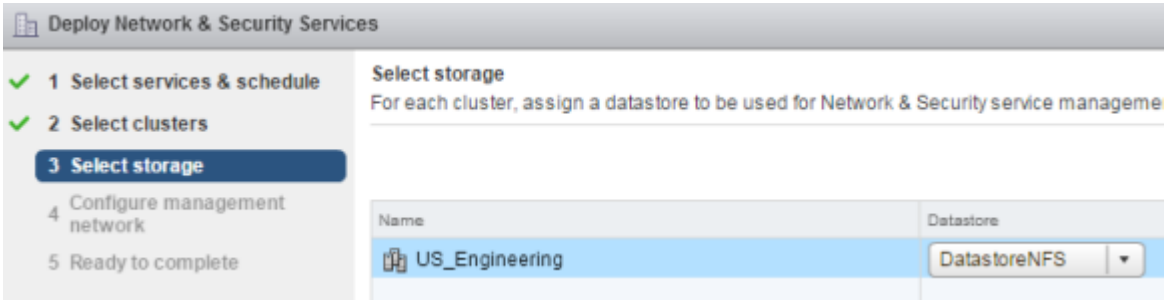

Figure 3-27 Select storage

12 In the **Configure management network** page, the record containing the selected service and cluster is displayed. Complete the following in **Configure management network** step.

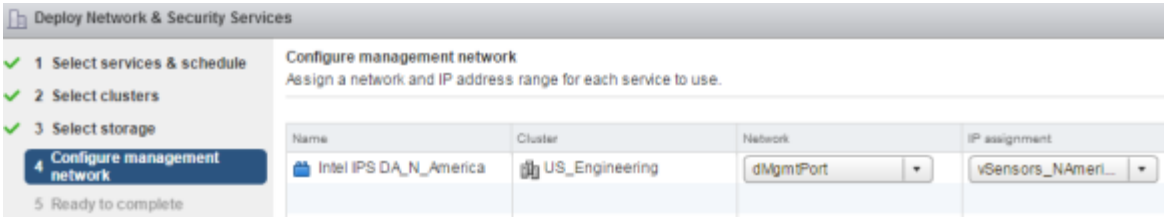

Figure 3-28 Configure management network

a From the **Network** drop-down list, select the distributed switch port group which the virtual security service appliances must use for management data. That is, the virtual security service appliance management port uses the switch port group you select here.

b From the **IP assignment** drop-down list, select the IP address pool which you configured and then click **Next**.

13 In the **Ready to complete** step, review the configuration and click **Finish**.

Depending on the number of VMware ESXi hosts and your network infrastructure, it takes some minutes for the virtual security service appliances to be deployed.

14 Make sure that the **Installation Status** shows up as *Succeeded* and **Service Status** shows up as *Up* in the **Installation** page.

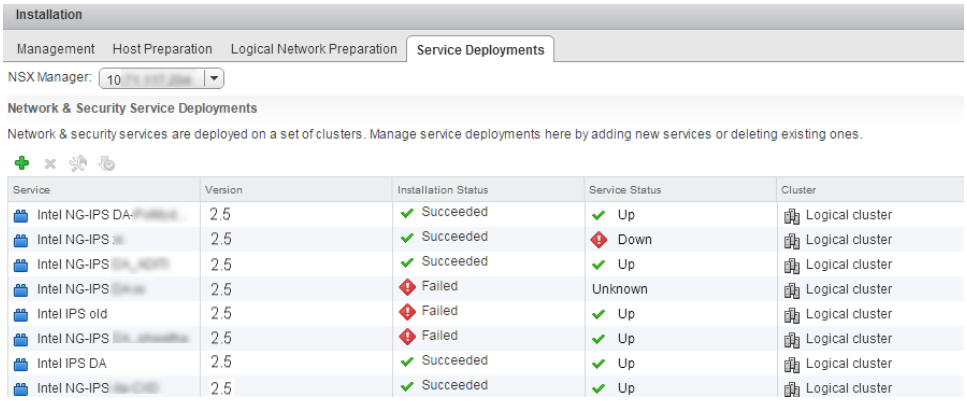

Figure 3-29 Installation page

15 In the Network Security Manager, deploy all changes to make sure the individual virtual security service appliances are updated.

16 Select the virtual system name based on the distributed appliance name and then select the appropriate option to view summary information about it.

Recall that one virtual security service appliance is automatically deployed per VMware ESXi host in the cluster.

17 In OSC, select **Status** | **Appliance Instances**.

The deployed virtual security system instances (security appliances) are listed. Make sure the state of **Discovered** and **Inspection-Ready** are *True*.

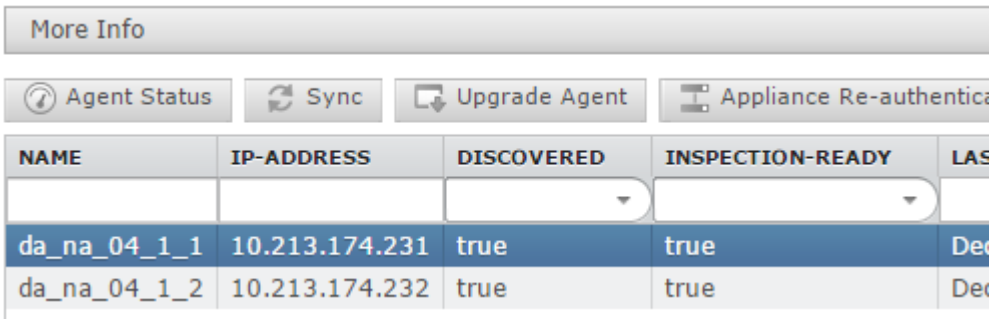

Figure 3-30 Deployed virtual security system instances

Maintaining virtual appliance instances

When you create a distributed appliance, you enable the required virtualization connectors in the distributed appliance. For every enabled virtualization connector, OSC creates a virtual security system (virtual system) by default. For Network Security Platform, this virtual security system in the distributed appliance corresponds to the virtual security system in the Network Security Manager.

A virtual security system is assigned a name by assigning a sequentially increasing number to the name of the distributed appliance. When you deploy this virtual security system on the required cluster, NSX automatically installs a virtual security appliance in each VMware ESXi host of the cluster. In the case of Network Security Platform, a virtual security appliance corresponds to the virtual security system instance, that is the virtual IPS Sensor installed in each VMware ESXi host.

After you deploy the virtual security system, you can view and maintain deployed virtual appliances from the **Appliance Instances** page of OSC.

Task

1 In OSC, select **Status** | **Appliance Instances**.

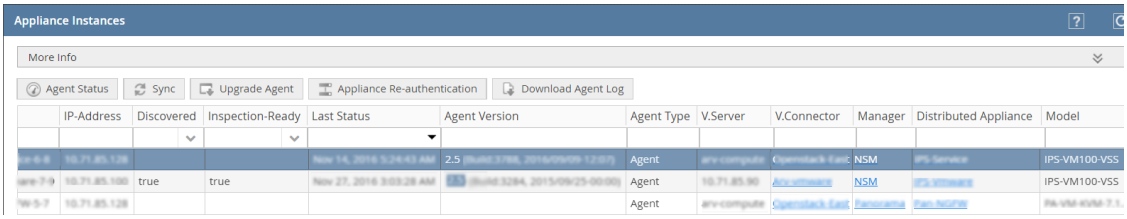

Figure 3-31 Appliance instances

The **Appliance Instances** page lists all the virtual appliances that are currently deployed.

2 Select an appliance instance and click **Agent Status** to view the details of the OSC agents running on the appliance. Most of these details are also visible in the security manager.

- **Refresh** — Refreshes the **Open Security Controller Agent(s) Status Window**.
- **Close** — Closes the **Open Security Controller Agent(s) Status Window**.
- **Name** — Name assigned to the virtual security system.
- **IP**:
  - For VMware — The IP address assigned to the management port of the virtual appliance. This IP address is randomly assigned from the IP pool, which you specified when you deployed the security service.
- **V. Server**— IP address of the virtualization provider or hypervisor server.
- **OSC IP** — The IP address of OSC.
- **Manager IP** — The IP address of the corresponding Manager, which manages the appliance. For Network Security Platform, it is the IP address of the Network Security Manager.
- **Version** — The version of the Control Path Agent (CPA) on the appliance.
- **Agent time** — The time as per CPA's clock.
- **Uptime** — Indicates how long the CPA is up and running.
- **CPA PID** — Unique ID for the CPA.
- **DPA PID** — Unique ID for the Data Path Agent (DPA) running on the appliance.
- **DPA Info** — The version details of the DPA.
- **DPA Stats** — Displays the number of packets received, transmitted, dropped, and so on by the appliance.
- **Discovered** — Displays *true* or *false* which are corresponding values that suggest whether the instance of the security service is discovered or not.
- **Inspection Ready** — Displays *true* or *false* which are corresponding values that suggest whether the instance of the security service is available for deployment and configuration or not.

3 Click **Sync** to manually synchronize changes by selecting the required appliance instance.

The most common scenario in which you will use this is if an automatic update fails. Such an update is initiated if you change network settings of OSC or change the password of the default *agent* user name. OSC agents in the virtual security system instances (appliances) are updated automatically. If the default update fails, use the manual sync option.

When you click **Sync**, a *Syncing Open Security Controller Agent(s)* job is triggered.

4 To upgrade OSC agents in the virtual security system instances (appliances), select an appliance instance and click **Upgrade Agent**.

An *Upgrade Open Security Controller Agent(s)* job is triggered.

5 If the **Discovered** state of an appliance is *false*, select the appliance and click **Appliance Re-authentication** to re-authenticate the appliance with the corresponding manager.

For Network Security Platform, the Virtual Security System instance re-establishes trust with the Network Security Manager when you click **Appliance Re-authentication**.

In some situations, there might be a slight delay before the **Discovered** state turns to *true*. One such example is when you just deploy a security service in NSX. In such cases, wait for the state to change instead of clicking **Appliance Re-authentication** again.

6 Select an appliance instance and click **Download Agent Log** to save log files related to the appliance.

Either use the log files for troubleshooting or share it with your sales partner to investigate further.

7 To filter the displayed records, enter or select a value in the required column headers and press the Enter key.

Table 3-8 Option definitions

Option	Definition
Name	Unique name assigned to each virtual appliance by default. The name consists of numeric ID appended to the virtual security system name. For example, if *DA_N_America_6_7* is the name, *DA_N_America_6* is the name of the virtual security system and 7 is the numeric ID appended to give the virtual appliance a unique name.
IP-Address	IP address assigned to the management port of the virtual appliance. This IP address is randomly assigned from the IP pool, which you specified when you deployed the security service.
Discovered	Indicates if an appliance is discovered by the manager managing that appliance. If this is false for an appliance, click **Appliance Re-authentication** to trigger re-authentication with the corresponding manager. For Network Security Platform, the discovered state is true if the output of the `status` command for the appliance indicates the following: • Trust is established with the Network Security Manager. • Alert and log channels are up.
Inspection-Ready	Indicates that the appliance is ready for traffic. The inspection-ready state in Network Security Platform is *True* if the output of the `status` command indicates the following: • Signature file is present in the virtual security system instance. • System is initialized. • System health is good.
Last Status	Timestamp of when OSC last checked on the virtual appliance. This time is as per the OSC system clock.

Table 3-8 Option definitions *(continued)*

Option	Definition
Agent Version	OSC agent on the virtual security system instance. This agent handles the management traffic between the appliance and the corresponding manager.
V.Server	IP address of the virtualization (VMware ESXi) server on which the virtual appliance is installed.
V.Connector	Unique name of the corresponding virtualization connector that you mentioned when you configured that connector. Clicking the hyperlink provided for the virtualization connector of any of the agents, routes you directly to the list of virtualization connectors in the **Virtualization Connectors** page.
Manager	Unique name of the corresponding manager connector that you mentioned when you configured that connector. Clicking the hyperlink provided for the manager of any of the agents, routes you directly to the list of manager connectors in the **Manager Connectors** page.
Distributed Appliance	Unique name of the corresponding distributed appliance that you mentioned when you configured that appliance. Clicking the hyperlink provided for the distributed appliance of any of the agents, routes you directly to the list of distributed appliances in the **Distributed Appliances** page.
Model	Model number of the security appliance as mentioned in the **Service Function Catalog** page.
Version	Software version of the security appliance as mentioned in the **Service Function Catalog** page.

Deleting a distributed appliance for VMware

Before you begin

1. You have access rights to uninstall service deployments in NSX.
2. The details you provided in the virtualization connectors and manager connector used in the corresponding distributed appliance are valid.

When you successfully delete a distributed appliance, the corresponding virtual security system instances (ESX agents) are deleted. Therefore, it results in the termination of the security service provided by these virtual security system instances.

To delete a distributed appliance, you must first sequentially delete the related objects in NSX as explained in this section.

Task

1 In the vSphere **Home** tab, select **Networking & Security** | **Installation** | **Service Deployments**.

2 Select the corresponding service deployment and click ✖.

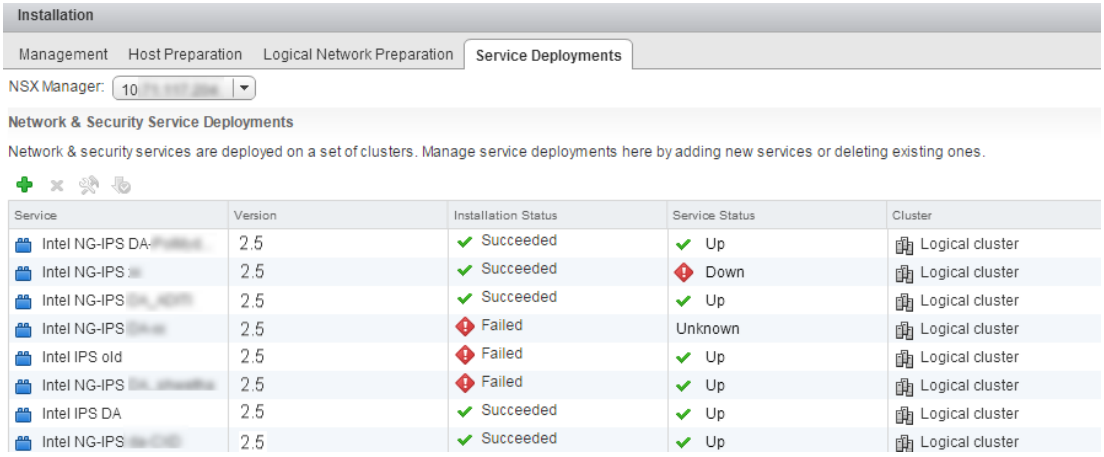

Figure 3-32 Service deployments

3 Select **Delete now** or schedule the deletion and click **OK**.

Depending on your configuration, it might take several minutes to uninstall the service deployment. This process deletes the virtual security system instances (ESX agents), implying the security service is terminated.

4 Select **Service Composer** | **Security Policies** and then select the NSX Manager.

5 Select the security policy used in the deleted Network Security Platform service deployment and click .

The security groups to which you assigned the security policy are displayed in pop-up window.

6 Deselect all the security groups to which you assigned the security policy and click **OK**.

7 In the **Networking & Security** pane, select **Service Definitions**.

8 Select the corresponding service definition and click the edit icon.

You can identify the service definition by the name of the distributed appliance you want to delete.

9 Select **Service Instances**.

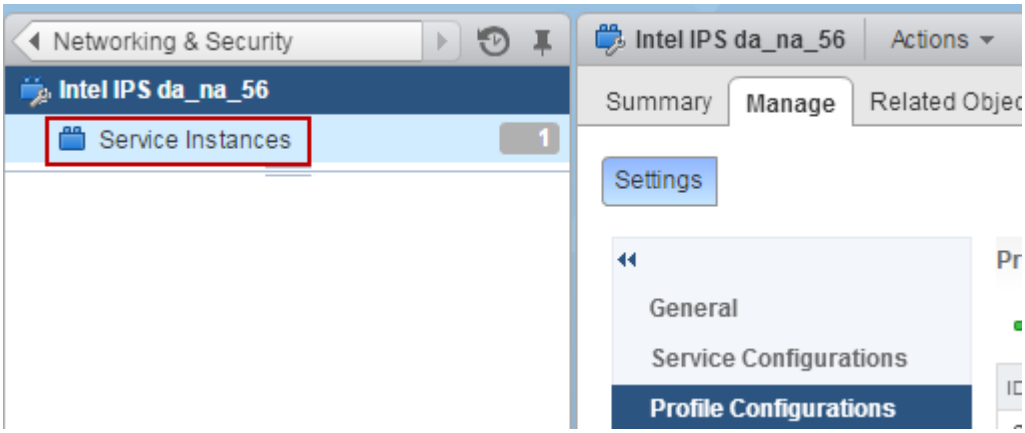

Figure 3-33 Service instances

10 Select the instance, which is displayed.

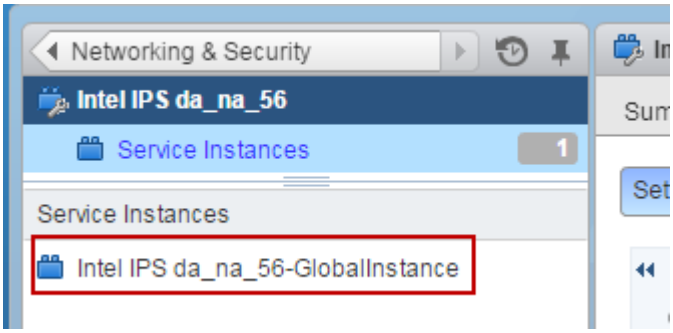

Figure 3-34 Select the instance

The corresponding service profiles are listed on the right side.

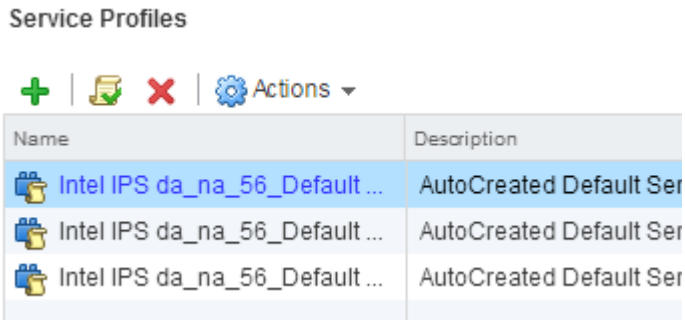

Figure 3-35 Service profiles

11 Select and delete all the service profiles one by one.

12 Select the service instance and delete the service instance.

The corresponding security policy in NSX is automatically deleted.

13 In the **Networking & Security** pane, select **Service Definitions**.

14 Delete the corresponding service definition.

Locate the service definition based on the name of the distributed appliance you want to delete. This completes the deletion of the related objects in NSX.

15 In OSC web application, select **Setup** | **Distributed Appliances**.

16 Select the distributed appliance and click **Delete**.

- Deleting the distributed appliance, deletes the service definition in NSX.
- After the distributed appliance deletes successfully, you can delete the corresponding manager connector and virtualization connector, if required.

Jobs and tasks

Certain actions you perform in OSC are tracked as jobs. When you start a job, it triggers one or more background activities in OSC. These background activities are tracked as tasks of that job. Therefore, a job is completed only when all its tasks are successfully completed.

Jobs and tasks enable easy tracking and troubleshooting. For example, if a job failed, you just have to look at the failed task to locate the stage at which the job failed. If a job is running for a long time, you can troubleshoot by looking at the task at which the processing is stuck.

OSC triggers a job, when you take any of the following actions:

- Create, edit, synchronize, or delete a manager connector.
- Create, edit, synchronize, or delete a distributed appliance.
- Synchronize appliance instances.
- Appliance instance re-authentication.
- Upgrade the software for an appliance instance.
- Modify the password of the default users.

Viewing jobs and tasks

You can view the jobs and the corresponding tasks in the **Jobs** page.

Task

1 In the OSC web application, select **Status** | **Jobs**.

The jobs are listed in the top pane of the page. When you click on a job, the corresponding tasks are listed in the bottom pane.

All time stamps displayed in the **Jobs** page are according to OSC system time. You can use `show clock` command or the **Manage** | **Server** | **Summary** page to check the current date and time on OSC. To change the system time, use the `set time` and `set timezone` commands.

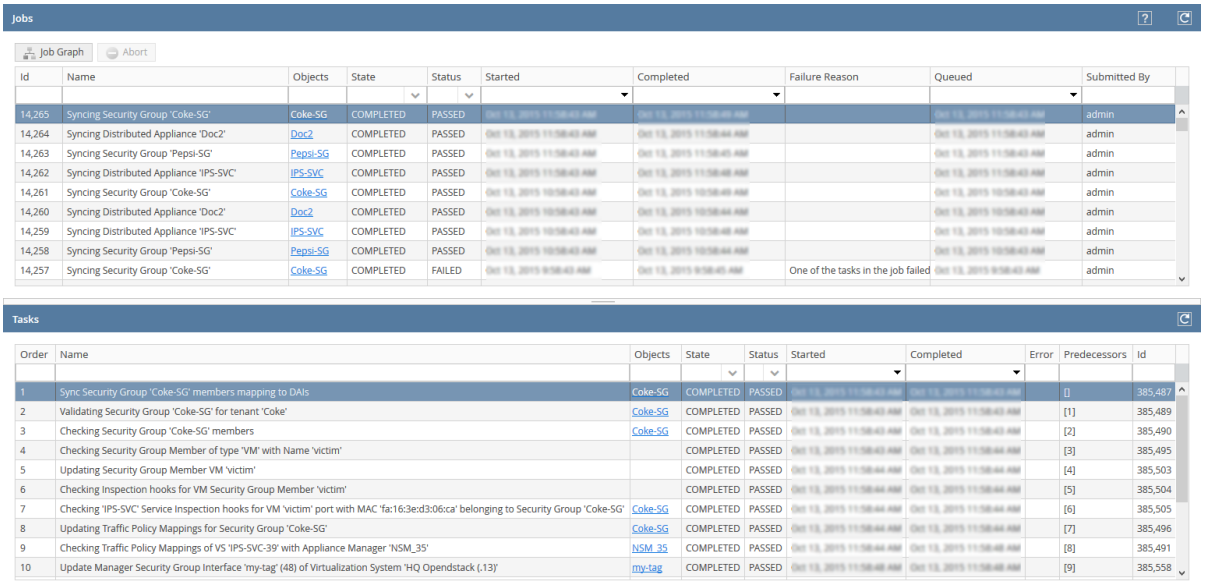

Figure 3-36 Jobs page

Table 3-9 Option definitions in the Jobs pane

Option	Definition
Job Graph	Click to see a graphical representation of the order of tasks for the selected job.
Abort	Click if you do not need a job to be processed any further or remove it from the processing queue. Even when you abort a job, the entries for all tasks of that job are created; task IDs are assigned for all tasks; job ID is assigned for the job.
ID	This is a system-assigned unique ID to each job. This ID is assigned in a sequential order.
Name	This is a system-defined name to a job. The name of the connector is appended at the end of the name. For example, if you delete a distributed appliance named *example_DA,* the name assigned for this job is *Delete Distributed Appliance 'example_DA''*.
Objects	Unique name of the distributed appliance executing that job. Clicking the hyperlink provided routes you directly to the list of distributed appliances in the **Distributed Appliances** page.

Table 3-9 Option definitions in the Jobs pane *(continued)*

Option	Definition
State	Indicates the current state of the job. The following are the possible values: • **NOT_RUNNING** — indicates that OSC started the job but is unable to complete the process. • **QUEUED** — indicates that the job is in the queue to be processed. For example, there might be too many concurrent jobs being processed and the jobs in the queue are processed as soon as system resources are available. • **RUNNING** — indicates that the job is now being processed. • **COMPLETED** — indicates that OSC completed processing the job. However, check the status of the job to see if the job was completed successfully.
Status	Indicates the result for a job. The following are the possible values: • **FAILED** — indicates that one or more the tasks of that job failed. So, OSC is unable to complete the job successfully. • **PASSED** — indicates that all tasks of the job are completed successfully. Hence, the jobs are also completed successfully. • **ABORTED** — A user clicked **Abort** to stop processing the job further.
Started	The time stamp of when the job is started. This is empty for jobs which processing never started. For example, if you aborted jobs, which are queued for processing, this time stamp is empty.
Completed	The time stamp of when the job is completed regardless of the job **Status**. The **Completed** time stamp is available only for those jobs, which are in completed **State**. Regarding **Status**, the completed time stamp is displayed for all failed, passed, and aborted jobs. In case of failed and passed jobs, this time stamp indicates when the last task was completed. In case of aborted jobs, this time stamp indicates when a user aborted the job.
Failure Reason	Displalys the reason for the failure of the job
Queued	The time stamp of when a user started the action. If OSC processes the job as soon as it is triggered, this time stamp is the same as that of the **Started** time stamp.
Submitted By	Name of the user who initiated the job.

When you click **Job Graph**, a graphical representation of the order and dependency of tasks displays. The tasks are color-coded to indicate their status.

- Green indicates the task succeeded.
- Yellow indicates the task is in progress. Not shown if the task is in different state.
- Red indicates the task failed.
- Gray indicates that the task is skipped.
- White indicates that the task is aborted or not yet scheduled.
- Orange indicates that the task is scheduled for execution.
- Light Blue indicates that there are predecessor tasks and only some of them are currently completed.
- Magenta arrow head indicates all tasks (not just its predecessors) before this task must be successfully completed.

- Black arrow head normal indicates allow task to start execution when any of its predecessors succeed.
- Black arrow head empty indicates all predecessor tasks must be in a completed state.

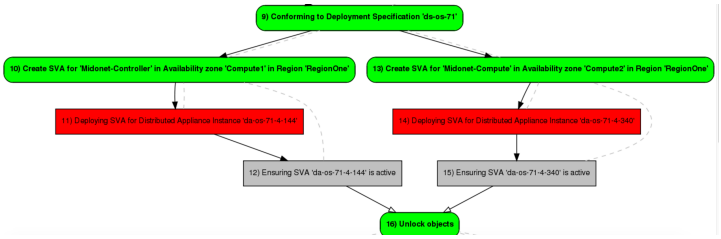

Figure 3-37 Job graph

2 Click a job to view its tasks.

Table 3-10 Option definitions in the Tasks pane

Option	Definition
Order	Indicates the sequence in which the tasks are executed.
Name	This is a system-defined name to a task. The name of the relevant connector is appended at the end for some of the tasks.
Objects	Unique name of the distributed appliance executing that task. Clicking the hyperlink provided routes you directly to the list of distributed appliances in the **Distributed Appliances** page.
State	Indicates the current state of the job. The following are the possible values: • **NOT_RUNNING** — indicates that OSC started the task but is now aborted. • **QUEUED** — indicates that the task is in the queue to be processed. For example, there might be too many concurrent tasks being processed and the tasks in the queue are processed as soon as system resources are available. • **PENDING** — indicates that there are predecessor tasks and only some of them are currently completed. • **RUNNING** — indicates that the task is now being processed. • **COMPLETED** — indicates that OSC completed processing the task. However, check the status of the task to see if it was completed successfully.
Status	Indicates the result for a task. The following are the possible values: • **FAILED** — indicates that the task failed. So, OSC is unable to complete the job successfully. • **SKIPPED** — indicates that OSC skipped this task and proceeded to the next task. When a prerequisite task fails, the current task is skipped. • **PASSED** — indicates that the task is completed successfully. • **ABORTED** — indicates that the task was started but a user clicked **Abort** to stop processing the job further.
Started	The time stamp of when the task is started. This is empty for tasks which never started.

Table 3-10 Option definitions in the Tasks pane *(continued)*

Option	Definition
Completed	The time stamp of when the task is completed regardless of the job **Status**. The **Completed** time stamp is available only for those tasks, which are in completed **State**. Regarding **Status**, the completed time stamp is displayed for all failed, passed, skipped, and aborted tasks. In case of failed and passed tasks, this time stamp indicates when the task was completed. In case of aborted jobs, this time stamp indicates when a user aborted the job. In case of skipped tasks, this time stamp indicates when a task was skipped.
Error	Displays the reason for failed tasks.
Predecessors	Indicates the **Order** number of one or more tasks, which must be completed to complete this task. Click **Job Graph** to see the graphical representation of the order in which tasks are executed.
ID	This is a system-assigned unique ID to each task. This ID is assigned in a sequential order and is unique across jobs.

3 Use the following options to change the display in the jobs and tasks pane.

- The page automatically shows the updated content. If necessary, you can click in the jobs or the tasks pane.
- To change the order of the columns in the jobs or tasks pane, click on a column header and drag it to where you want in the pane.
- Click on a column header to sort the records in the ascending or descending order. For example, click on the **Order** column in the tasks pane to display the records in descending order, that is, the last task is displayed first and the first task is displayed last.
- To filter the records, enter or select the values in the required columns and press the Enter key. All records containing the specified values are listed. For example, enter *distributed appliance* **Name** column and select **Failed** in the **Status** column to view the failed jobs with the **Failure Reason** related to distributed appliances. Follow a similar procedure to filter records in the tasks pane.

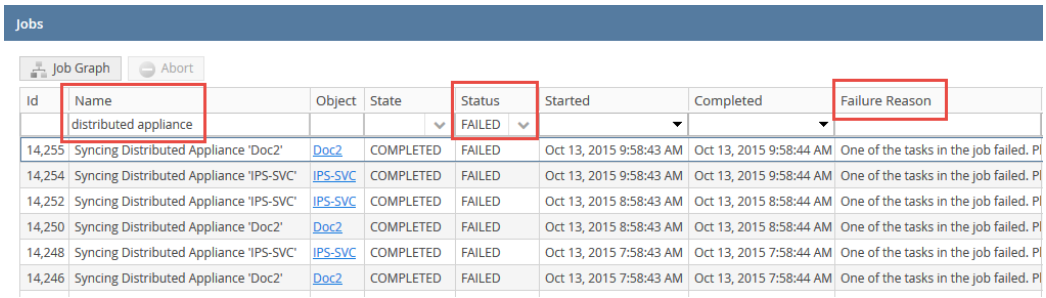

Figure 3-38 Filter records in the Jobs page

- To remove the filters, delete the search strings and selections from the columns and press *Enter* on your keyboard.

Create a security group in VMware NSX

In an NSX Manager, you create a security group and then include the required VMs in that group. Then you can apply an NSX security policy to this security group, the corresponding security service is provided to those VMs included in the security group.

If VMware tools is not running on the VMs, you must include those VMs in the security group and also create an IP set containing the IP addresses of those VMs and include that IP set in the security group.

Task

1 Log on to vSphere Web Client as the root user.

2 In the vSphere **Home** tab, select **Networking & Security.**

Figure 3-39 vSphere home page

3 Select **Service Composer** | **Security Groups.**

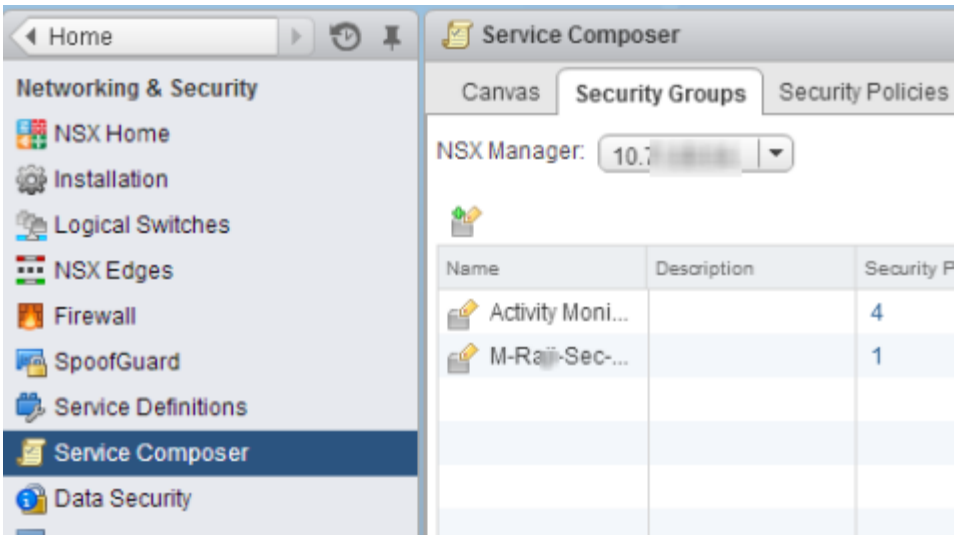

Figure 3-40 Service composer

4 From the **NSX Manager** list, select the NSX Manager in which you want to define the security group.

5 Click  to create a security group.

6 In the **New Security Group** wizard, enter a meaningful name and, if required, a description, then click **Next**.

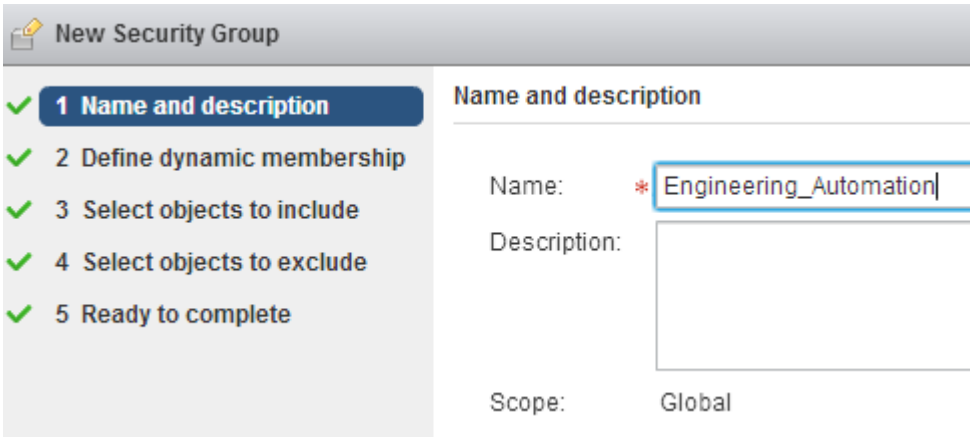

Figure 3-41 New security group

7 Select **Select objects to include**.

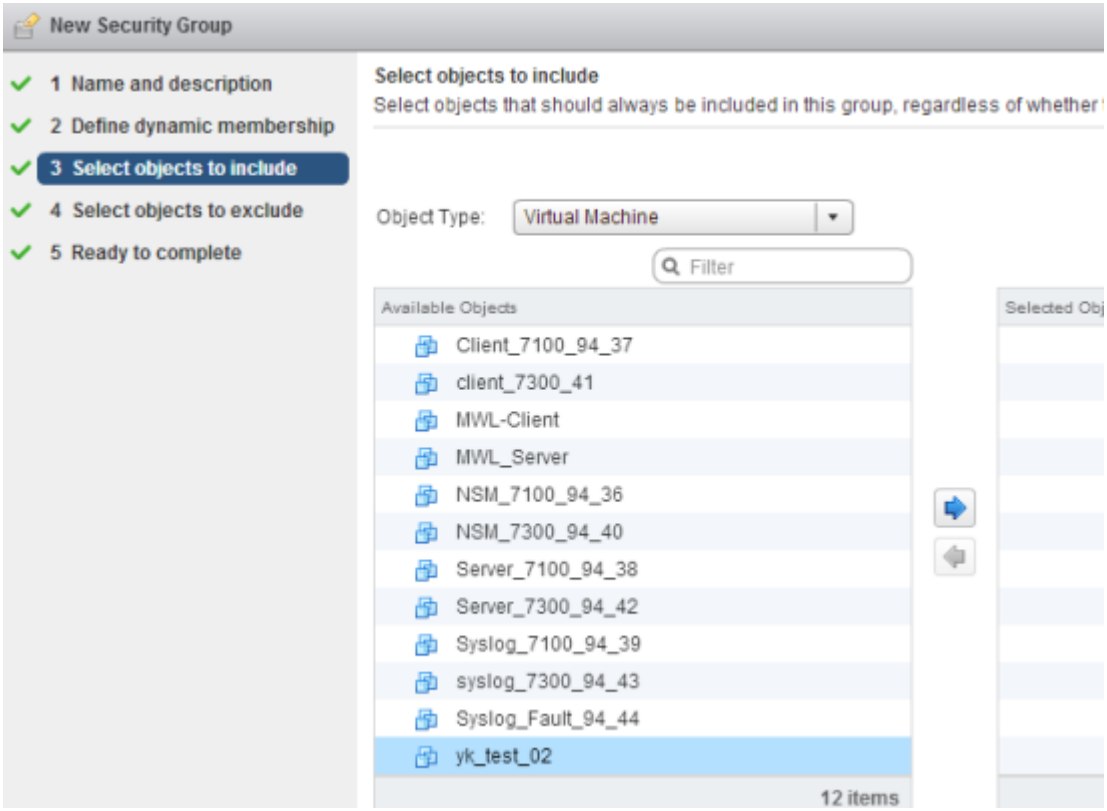

Figure 3-42 Select objects to include

8 From the **Object Type** drop-down list, select the object based on which you want to include VMs.

For example, if you select distributed port group in this list, the distributed port groups currently defined in the data center are listed. When you select a distributed port group, all the VMs connected to this port group are included in the security group.

- You can also base the inclusion on multiple object types. For example, you can select a few distributed switch port groups and some VMs.
- For VMs on which VMware tools is not running, you must include them in the security group. Also, you must create an IP set object and include that IP set in the security group. To include the IP set, select **IP Sets** from the **Object Type** list.

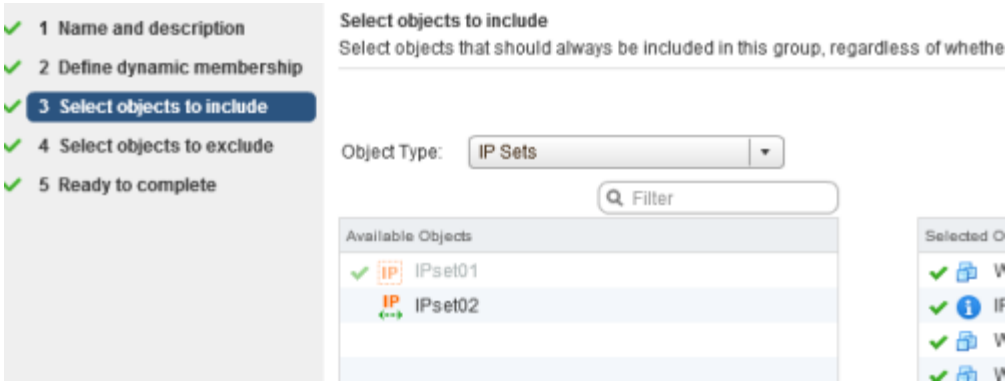

Figure 3-43 Select objects to include

Consider that you selected **Virtual Machine** from the **Object Type** drop-down list.

9 Select the required objects (in our example, VMs) and click  to move them under **Selected Objects**. Then click **Next**.

10 If you want to exclude any VMs, click **Select objects to exclude** in the **New Security Group** wizard.

This is similar to how you included VMs based on objects. For example, you want to include all the VMs connected to a distributed switch port group except for 5 server VMs. Then, you can include the distributed switch port group in step 3 of **New Security Group** wizard and exclude only those 5 VMs in step 4.

11 Click **Ready to complete**, review the objects included and those excluded. Then click **Finish** to create the security group.

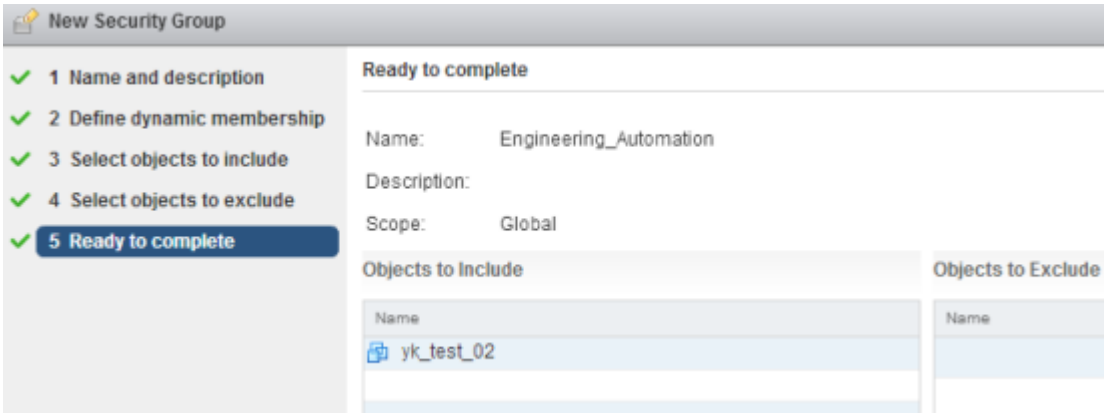

Figure 3-44 Ready to complete

The security group you created is listed in the **Security Groups** tab for the corresponding NSX Manager.

Create a security policy in VMware NSX

In an NSX Manager, you create a security policy, which you can apply to a security group in NSX. This security policy contains the security profile to be applied on the VMs included in that security group.

You must select the Network Security Platform policy group as a security profile in a security policy of NSX. Then, when you apply this security policy to a security group, a Virtual IPS Sensor uses this Network Security Platform policy group to inspect traffic related to the protected VMs.

Task

1 Log on to vSphere Web Client as the root user.

2 In the vSphere **Home** tab, select **Networking & Security.**

Figure 3-45 vSphere home tab

3 Select **Service Composer** | **Security Policies.**

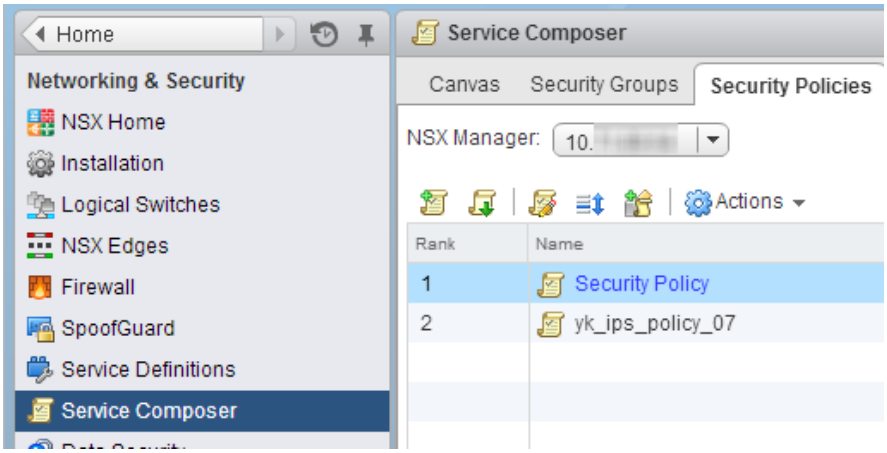

Figure 3-46 Service composer

4 From the **NSX Manager** list, select the NSX Manager in which you want to define the security policy.

5 Click to create a security policy.

6 In the **New Security Policy** wizard, enter a meaningful name and, if required, a description and click **Next**.

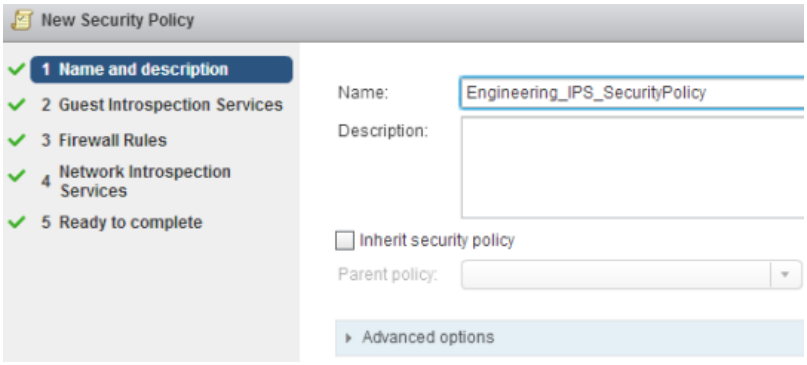

Figure 3-47 New security policy

7 Select **Network Introspection Services** and click ✚ to add a network introspection service.

All products that currently integrate with OSC are network introspection services.

> For the security services to work as expected, you must add two network introspection services. One introspection service is for inbound traffic to the security group. The second one is for the outbound traffic from the security group.

8 In the **Add Network Introspection Service** dialog, enter the required options and click **OK**.

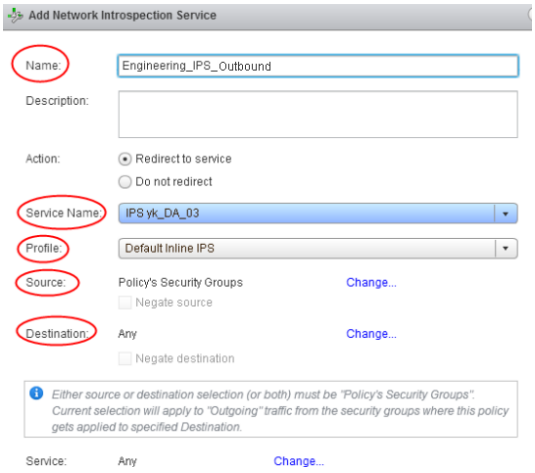

Figure 3-48 Add network introspection service

9 Follow a similar procedure to create the inbound security service.

- You can opt for the same **Profile** or choose another.
- You must select **Any** for **Source** and **Policy's Security Groups** for **Destination**.

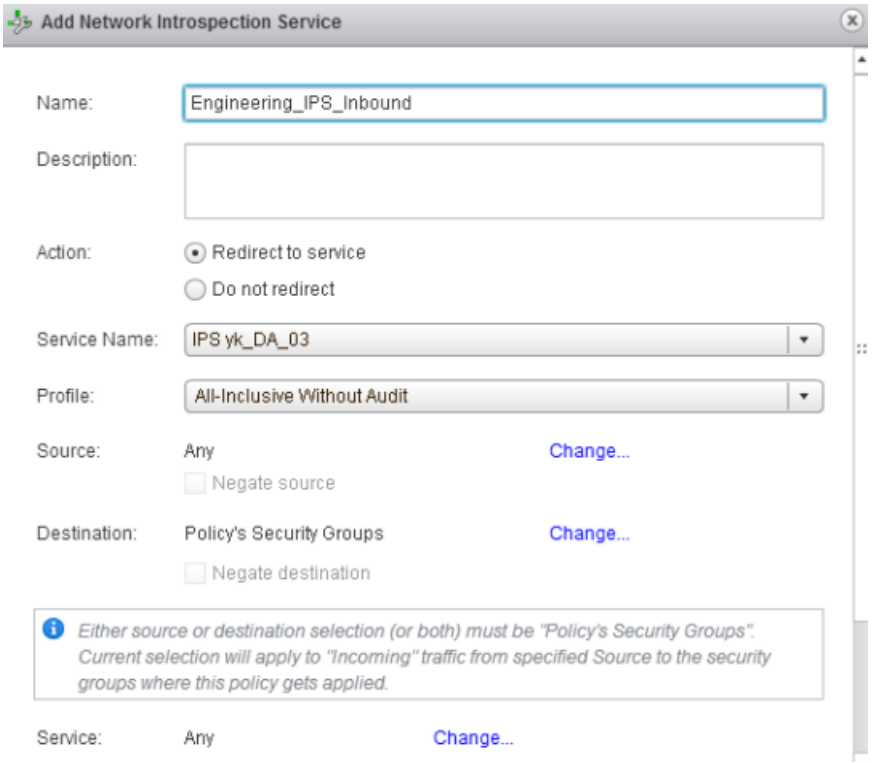

Figure 3-49 Policy's security groups

10 Select **Ready to complete**, review the configuration, and click **Finish** to create the security policy.

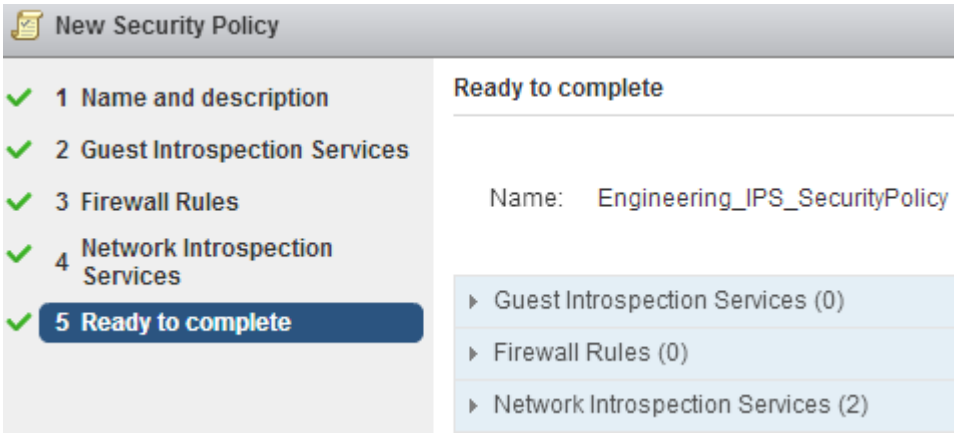

Figure 3-50 Ready to complete

The security policy you created is listed in the **Security Policies** tab for the corresponding NSX Manager.

Apply a security policy to a security group in VMware NSX

Before you begin

Make sure you have created the required security groups and the security policies in the NSX Manager.

In an NSX Manager, you create a security policy, which you can apply to a security group. This creates the association between security groups and security policies.

Task

1 Log on to vSphere Web Client as the root user.

2 In the vSphere **Home** tab, select **Networking & Security.**

Figure 3-51 vSphere home tab

3 Select **Service Composer** | **Security Policies.**

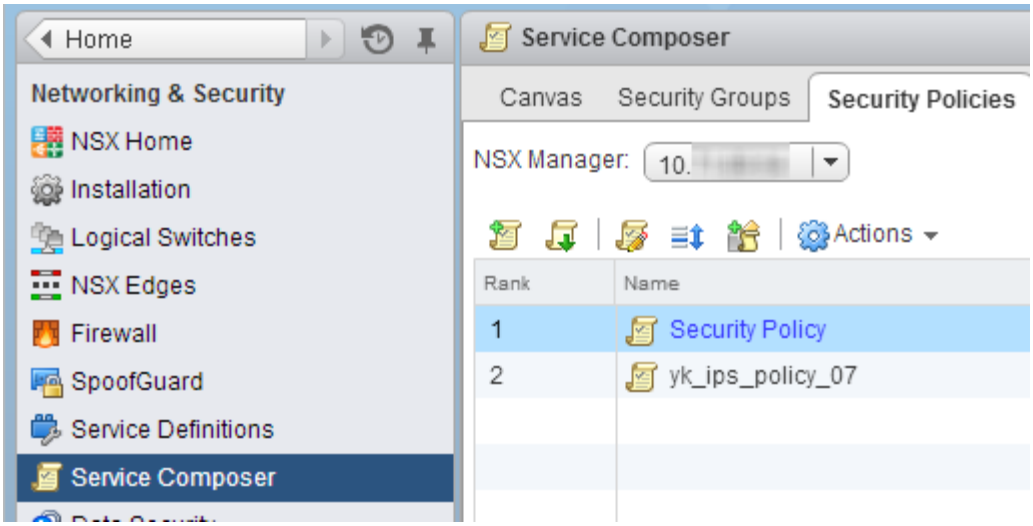

Figure 3-52 Service composer

4 From the **NSX Manager** list, select the corresponding NSX Manager.

5 Select the security policy you want to apply and click 

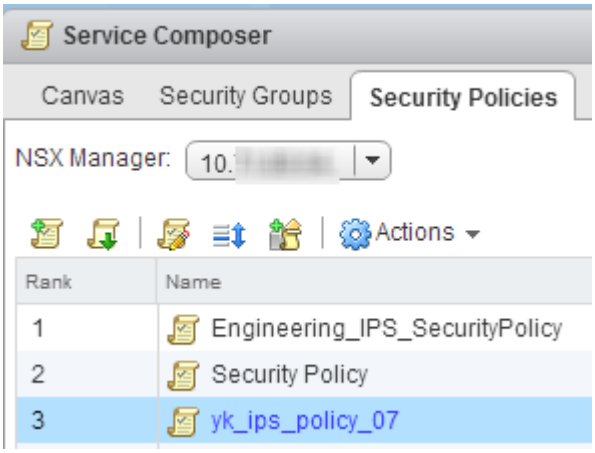

Figure 3-53 NSX manager

6 Select the security groups on which you want to apply the security policy and click **OK**.

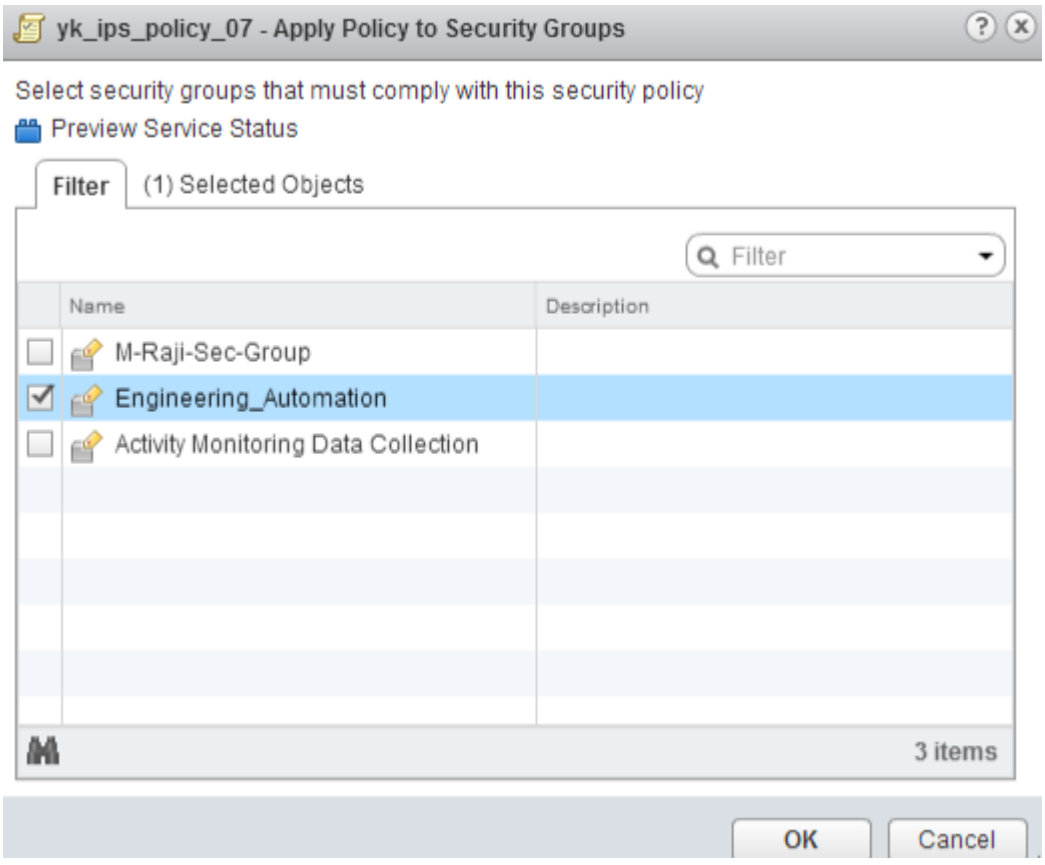

Figure 3-54 Select the security groups

You can view the details in the **Security Policies** tab.

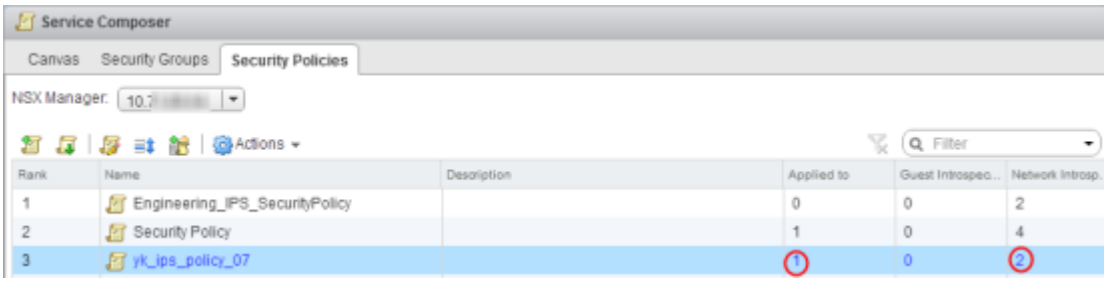

Figure 3-55 Security policies

Configure Virtual Security System to fail-close or fail-open

When you successfully install the security service, the Virtual Security System instances are deployed in fail-open mode by default. You can configure the Virtual Security System instances to run in fail-close or fail-open mode. Similar to any other configuration, the fail-close or fail-open setting of a Virtual Security System applies to all its instances.

The fail-open or fail-close configuration is implemented through an NSX mechanism. The solution uses an attribute in the service definitions of NSX to implement fail-open or fail-close configuration.

Task

1 In the vSphere **Home** tab, select **Networking & Security** | **Service Definitions**.

2 Select the corresponding service definition and click the edit icon.

You can identify the service definition by the name of the distributed appliance.

3 Select **Service Instances**.

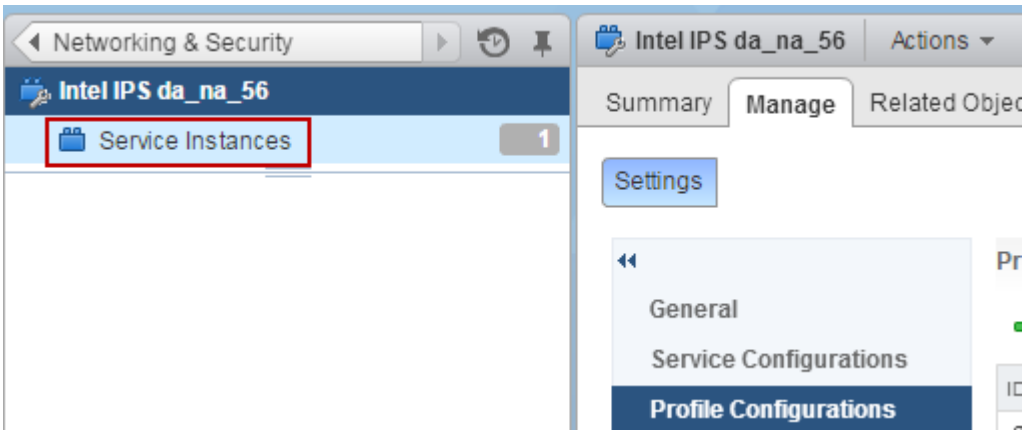

Figure 3-56 Service instances

4 Select the instance, which is displayed.

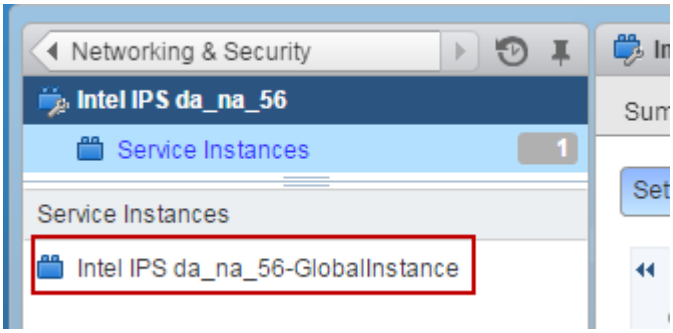

Figure 3-57 Select the instance

The corresponding service profiles are listed on the right side.

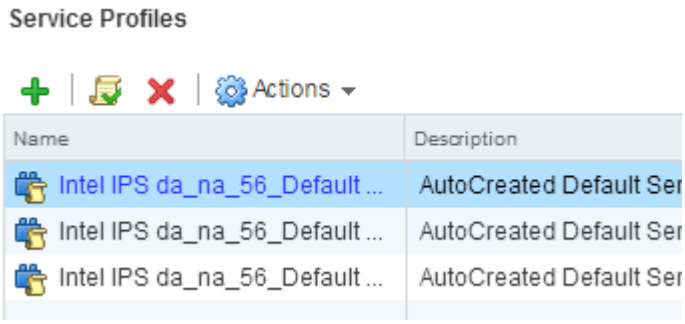

Figure 3-58 Service profiles

5 Double-click on the service profile, which you have used in the network introspection service of the applied security policy.

Consider that you selected the *Default Client and Server Protection (IDS IPS)* profile in the inbound and outbound network introspection service (as shown below).

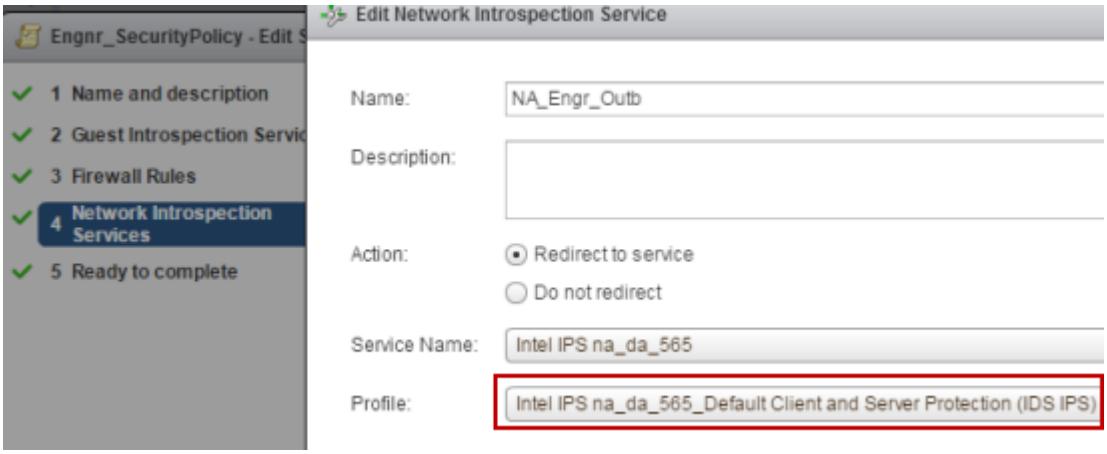

Figure 3-59 Edit network introspection service

Double-click on *Default Client and Server Protection* service profile.

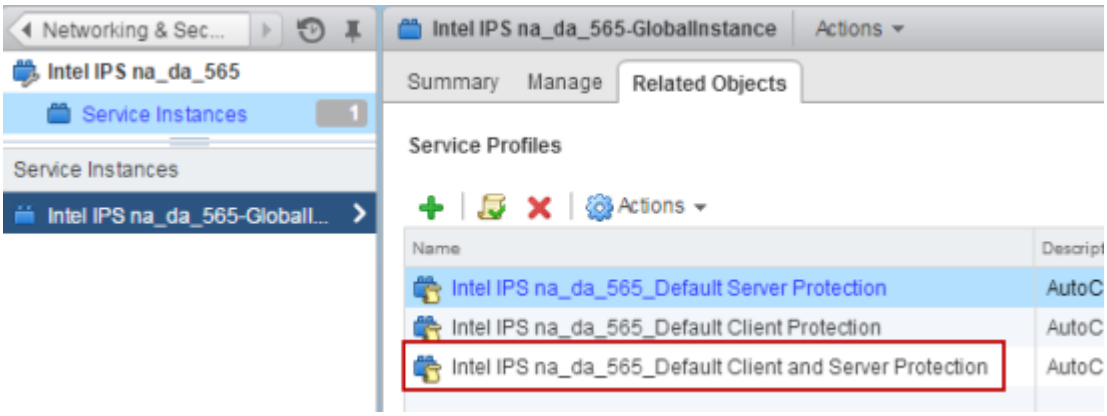

Figure 3-60 Select the service profile

6 Select the **Failure Policy** attribute and click **Edit.**

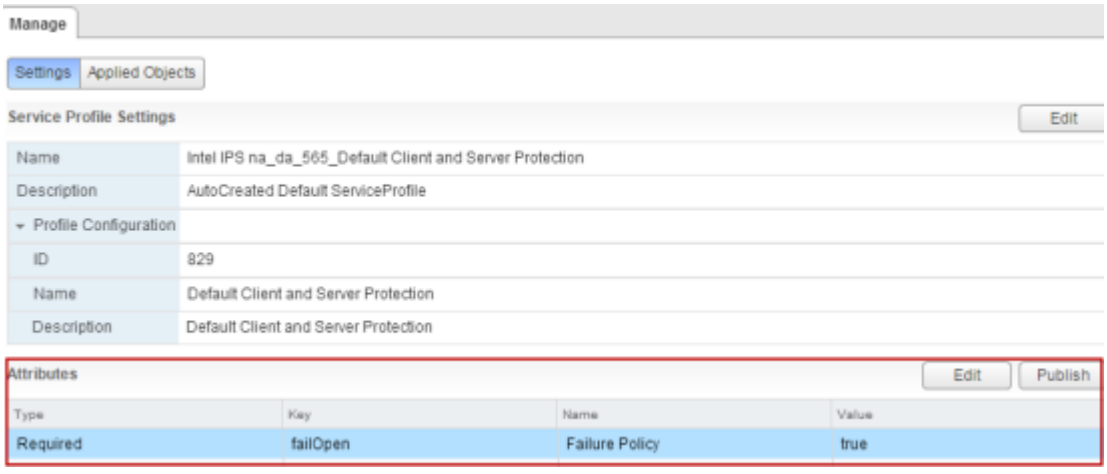

Figure 3-61 Failure policy attribute

7 By default, the **Failure Policy** attribute's fail-open key is set to true. This means the Virtual Security System instance fail-opens in case of a failure. Set the value to false, if you want the Virtual Security System instance to fail-close.

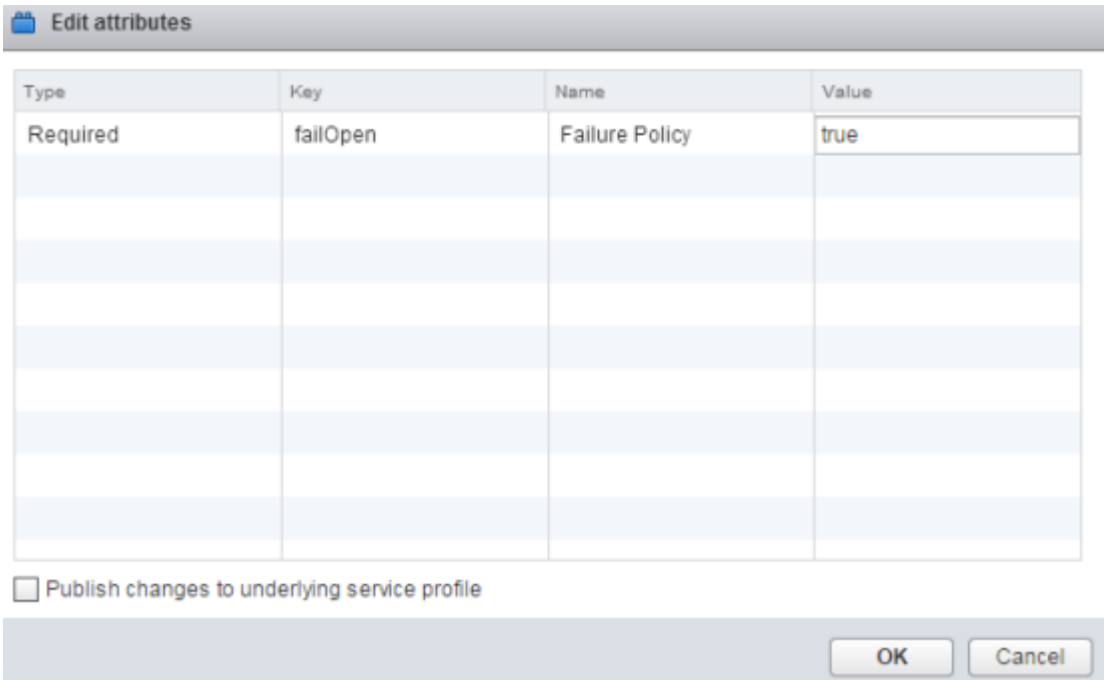

Figure 3-62 Edit attributes

8 In the **Edit attributes** dialog box, select **Publish changes to underlying service profile** and click **OK**.

9 Select the attribute and click **Publish**.

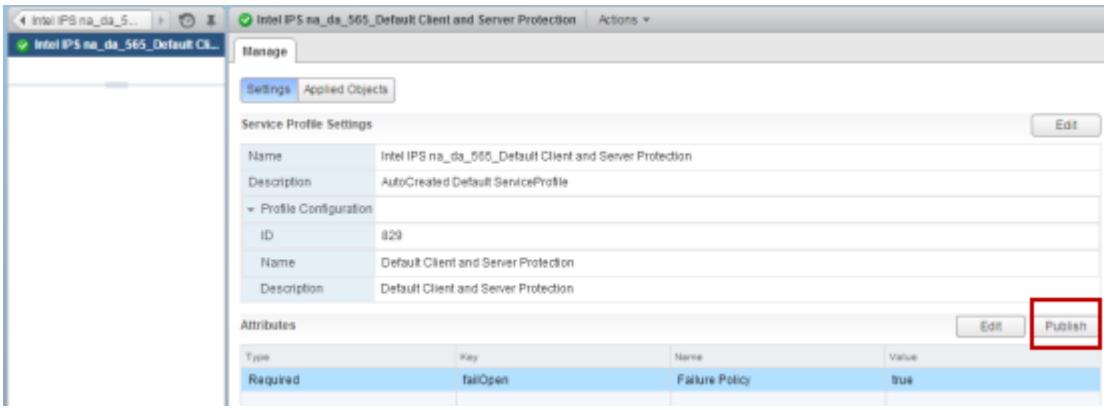

Figure 3-63 Select publish

Repeat the process if you used different profiles for inbound and outbound network introspection services.

Manager functions regarding IPS service deployment

In the context of IPS service through OSC, you use the Manager to perform the following tasks:

- Define the Manager admin domains, which should manage the Virtual Security System and the member instances. You can select the required domain for each **Virtualization System** in a distributed appliance. Then, the virtual security system is created under the domain you selected in the distributed appliance.
- Create a policy group containing the security policies for next-generation IPS. For example, the policy group can contain an IPS policy, and Advanced Malware policy. When you deploy the IPS service, OSC collaborates with vCenter and NSX to ensure the Virtual Security System instances use the selected policy group for traffic inspection.
  - For a Virtual Security System, the policy assignment is only through a policy group. That is, you cannot assign IPS policy and so on separately. To access the **Policy Groups** page, select **Policy** | **<Admin Domain Name>** | **Intrusion Prevention** | **Objects** | **Policy Groups**.
  - For the Virtual Security System, you assign the policy groups that you create in the Manager through the security policies of NSX. That is, you cannot assign the policy group to a Virtual Security System instance or its interface through the Manager.
- When you create a new policy group, OSC makes sure that the security policies in NSX are updated accordingly. The *Syncing Appliance Manager Connector <manager connector name>* job is triggered in OSC.
- **Important note on renaming or deleting policy groups in the Manager:** Recall that the policy groups in the Manager are automatically available as service profiles and profile configurations in NSX. So, when you delete even an unused policy group in the Manager or rename a policy group, you must manually delete the service profile and profile configuration in NSX. Else the subsequent manager-connector sync and distributed-appliance sync jobs fail.

  1 In the **Networking & Security** pane, select **Service Definitions**.

  2 Select the corresponding service definition and click the edit icon.

  3 Select **Service Instances**.

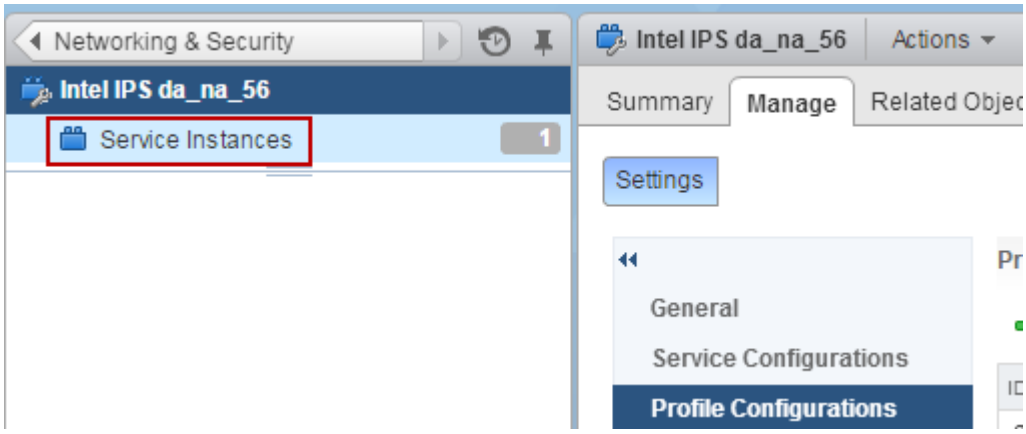

4 Select the instance, which is displayed.

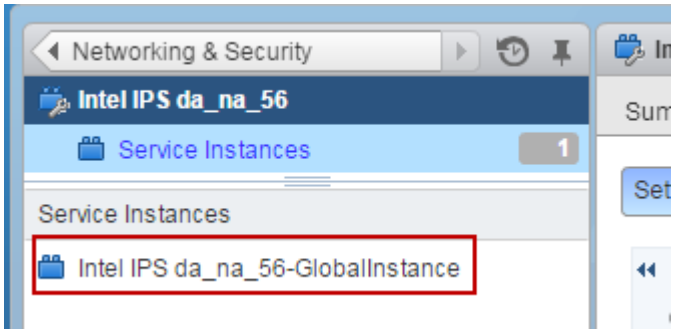

The corresponding service profiles are listed on the right side.

5 Delete the service profile corresponding to the policy group you renamed or deleted in the Manager.

6 In the **Networking & Security** pane, select **Service Definitions**.

7 Select the corresponding service definition and click the edit icon.

8 Delete the profile configuration corresponding to the policy group you renamed or deleted in the Manager.

- View details of the deployed Virtual Security System instances. For example, you can view the VMware ESXi servers on which the Virtual Security System instances are installed, the IP addresses of the Virtual Security System instances, and so on. See View summary details for virtual security systems on page 154.
- Use the Manager to deploy configuration changes, signature sets, and so on for the Virtual Security System instances.
- For information on how to change the policy group assigned to virtual security system instances, see Assign policy groups to virtual security systems on page 156.
- If you want to modify an attack definition, deploy configuration changes from the Manager after you modify the attack definition. For example, if you use the *Default Prevention* policy in the policy group. If you modify any attack definition in *Default Prevention*, deploy the configuration changes to the virtual security system.
- View the alert and other details sent by the Virtual Security System instances.

View summary details for virtual security systems

Before you begin

You have deployed the IPS security service in NSX.

You might want to view the details of virtual security systems and its deployed instances in the Manager.

Task

- In the Manager, select **Devices** | **<Admin Domain Name>** | **Devices** | **<virtual security system name>** | **Summary**.

> The virtual security system is system-defined. The name is derived by adding a sequential number to the distributed appliance name.

The device **Summary** page displays.

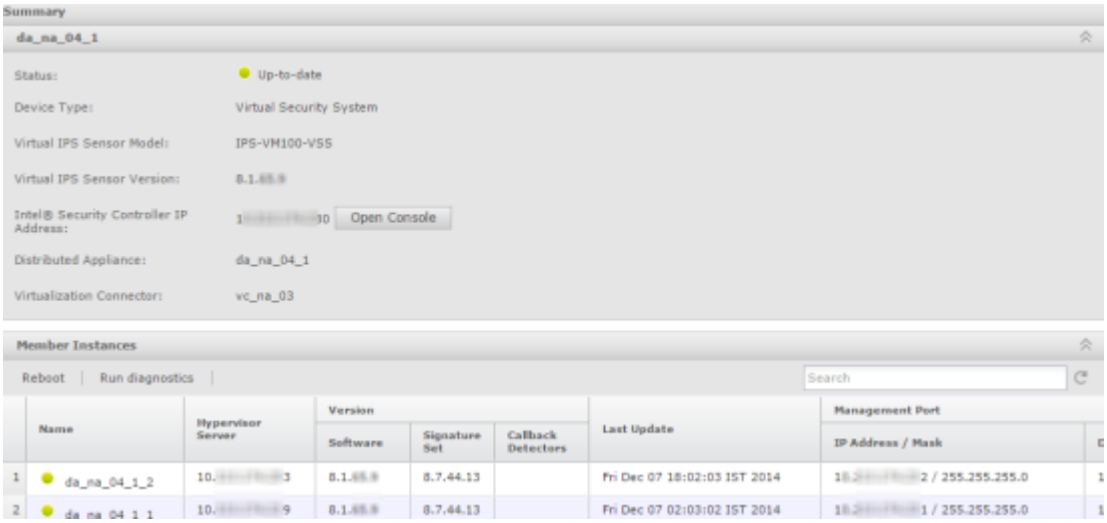

Figure 3-64 Summary details of a virtual security system

Table 3-11 Option definitions

Option	Definition
Status	Indicates whether there are pending changes to be deployed to the virtual security system or if it is up to date. Green indicates that the virtual security system is up to date and blue indicates that there are pending changes to be deployed. The term *virtual security system* collectively refers to all virtual security system instances it contains.
Virtual IPS Sensor Model	Indicates the virtual security system model. An example is IPS-VM100-VSS.
Virtual IPS Sensor Version	Indicates the Sensor image version used to create the virtual security system instances.
Open Security Controller IP Address	The IP address of the associated OSC. Click **Open Console** to access the OSC web application.
Virtualization Connector	When you define the distributed appliance, you select a Manager domain for each enabled virtualization connector. The corresponding virtualization connector for the current admin domain is listed.
Member Instances	A member instance corresponds to a Virtual IPS Sensor installed on a hypervisor. The container object of the member instances is the virtual security system.
Name	This is a system-defined name for a virtual security system instance. The name is derived by adding a sequential number to the virtual security system name.
Hypervisor Server	The IP address of the hypervisor on which the corresponding virtual security system instance is installed. In the case of VMware, this is the IP address of the VMware ESXi hosts in the protected cluster.

Table 3-11 Option definitions *(continued)*

Option	Definition
Version	The software image, signature set, and callback detectors versions on the corresponding virtual security system instance.
Last Update	The time stamp of when pending changes were deployed last.
Management Port	The network settings of virtual security system instance. When you deploy the IPS service, NSX assigns the network settings for each instance based on the IP address pool you defined in NSX. These network settings act as the management port settings for the virtual security system instance.
Last Reboot	The time stamp of when a virtual security system instance was last restarted.
Reboot	To restart a virtual security system instance, select the record and click **Reboot**. The IPS service to the VMs on the corresponding hypervisor is suspended until the virtual security system instance is functional again.
Run diagnostics	When you click, a diagnostic trace for the corresponding virtual security system instance is generated. You can provide this file to McAfee support for analysis. This diagnostic trace file is available on the **Diagnostics Trace** page in the Manager. To access the **Diagnostics Trace** page, select **Troubleshooting** | **Diagnostics Trace**.

Assign policy groups to virtual security systems

Before you begin

You have created the required policy group in the Manager.

You might want to apply a different policy group to the deployed instances of a virtual security system.

Task

1. Log on to vSphere Web Client as the root user.
2. In the vSphere **Home** tab, select **Networking & Security**.
3. Select **Service Composer** | **Security Policies**.
4. From the **NSX Manager** list, select the corresponding NSX Manager.
5. Select the required security policy and click  .
6. Select **Network Introspection Services**.
7. Select the required network introspection service and click the edit icon.

8 From the **Profile** drop-down list, select the required policy group and then click **OK**.

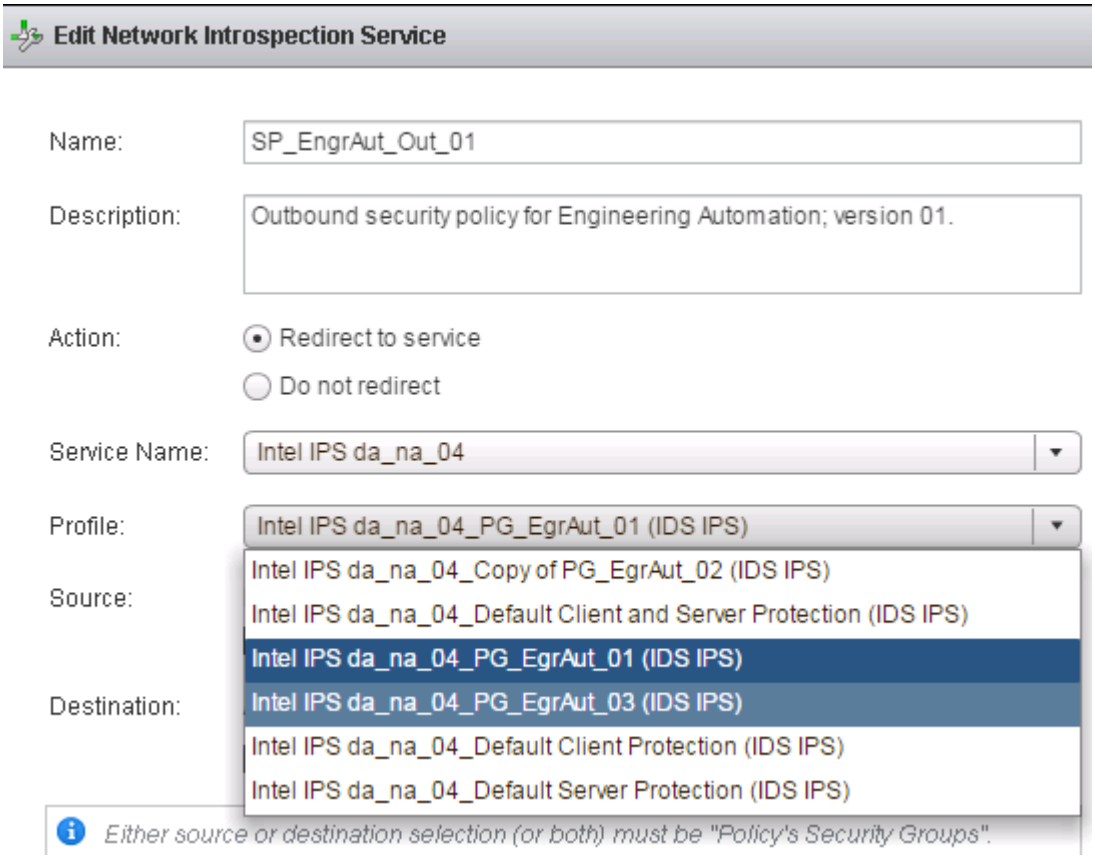

Figure 3-65 Select the policy group in NSX

9 Click **Finish**.

In OSC, the *Syncing Service Profile* job is triggered for every network introspection service that you change. Make sure this job passes successfully.

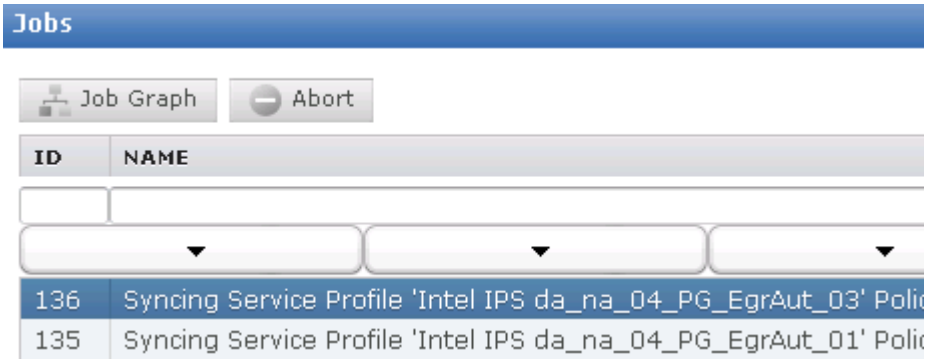

Figure 3-66 Job for synchronizing change in policy group

You do not require to deploy configuration changes from the Manager.

Quarantining endpoints using NSX features

Before you begin

You have the required access to create security groups and security policies in NSX.

In case of alerts detected by Virtual Security System instances, you can use the native Quarantine feature to quarantine the source endpoint of an attack. You can also use the security tags and security policies in NSX to quarantine the source or the target endpoint of an attack.

To quarantine endpoints using NSX features, you tag the source or destination VM in alert from the Attack Log. In NSX, create a security group, which dynamically includes VMs tagged in the Attack Log. To this group, you assign a security policy, which effectively quarantines the tagged VM.

Notes:

- The Quarantine feature of IPS and the security tags of NSX are exclusive to each other.
- In case of security tags, currently you can only quarantine endpoints from the **Attack Log**. You cannot enable quarantine through security tags in the attack definitions.
- Currently, there is no indication in the Attack Log that an endpoint is tagged.
- Currently, to release an endpoint from quarantine, you must manually remove the tag in NSX.

This section uses an example to explain how to quarantine endpoints using security tags and security policies in NSX.

Task

1 In NSX, create a security group to dynamically include tagged VMs.

a In vCenter web client, select **Home** | **Networking & Security** | **Service Composer** | **Security Groups**.

b In the **Service Composer** page, select the NSX Manager from the drop-down list.

c Click  to create a security group.

d Enter a relevant name and description.

e In the **Define dynamic membership** step, select **any** for **Match**.

f In **Criteria Details**, select **Security Tag** and **Equals to** in the first two drop-down.

g In the text box, exactly enter `OSC-Quarantine`.

The Manager assigns `OSC-Quarantine` as the tag. So, you must also exactly enter this string for a tagged VM to be included in this security group.

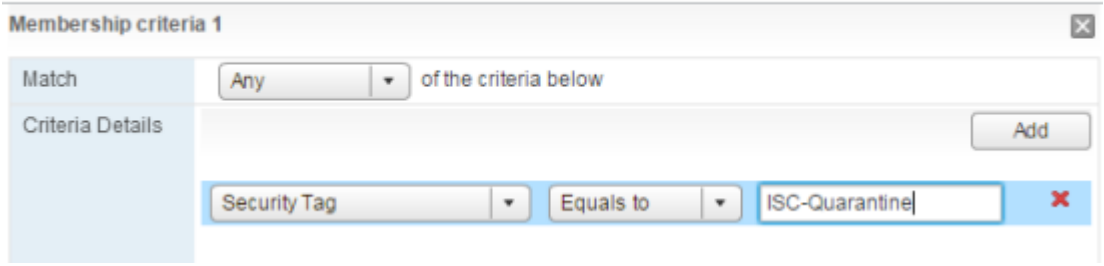

Figure 3-67 Adding criteria to include tagged VMs

h Click **Finish**.

2 In NSX, create a security policy to quarantine VMs tagged as *OSC-Quarantine*.

a In the **Securities Policies** tab of **Service Composer** page, click  to create a security policy.

b Enter a relevant name and description and proceed to step 3, **Firewall Rules**.

c Add a firewall rule to allow traffic to the VM hosting the remediation portal.

For **Source**, select **Policy's Security Groups**. For **Destination** select the security group, which contains the VM hosting the remediation portal.

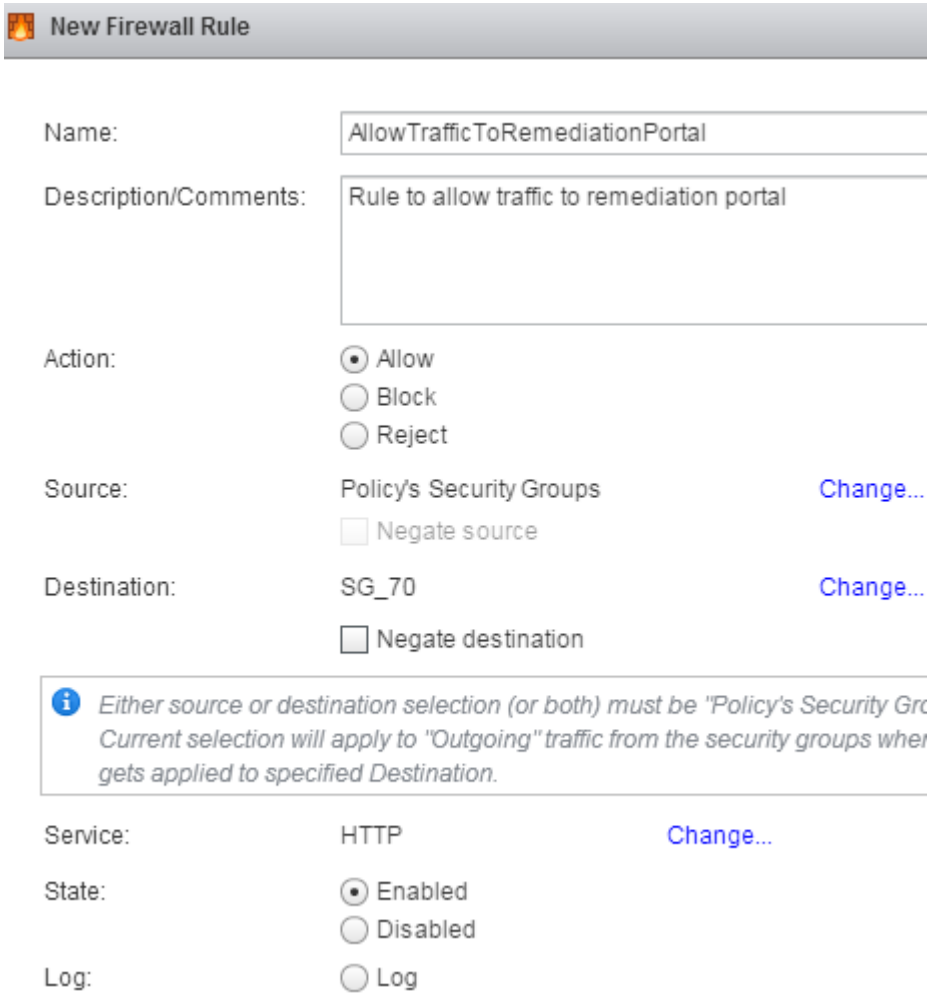

Figure 3-68 Rule to allow traffic to remediation portal

d Add the last rule, which blocks all traffic.

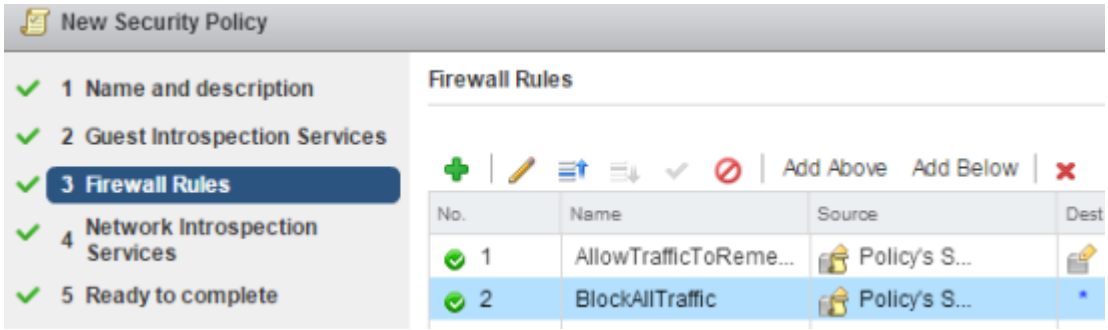

Figure 3-69 Rule to block all traffic

e Click **Finish**.

3 Apply the security policy from the previous step to the security group created in step 1.

a In the **Securities Policies** tab of **Service Composer** page, select the relevant security policy and click 

4 Select the relevant security group and click **OK**.

5 Tag a VM from the Attack Log.

a Select an alert and click **Other Actions**.

b Go to **Tag Endpoint** | **via OSC**.

c Select either the source IP address or the destination IP address.

The **Tag Virtual Endpoint via OSC** pop-up opens.

The endpoint and VSS server is detected by default.

d Select the tag from the **Tag to Assign** drop-down list.

e Click **Tag**.

A successfully tagged message is displayed.

Only the default tags are available for tagging the endpoint.

6 Confirm tagging in NSX.

a In vCenter web client, select **Home** | **Networking & Security** | **Networking & Security Inventory** | **NSX Managers**.

b Select the relevant NSX Manager and then select **Manage** | **Security Tags**.

c Click on the **VM Count** for **OSC-Quarantine**.

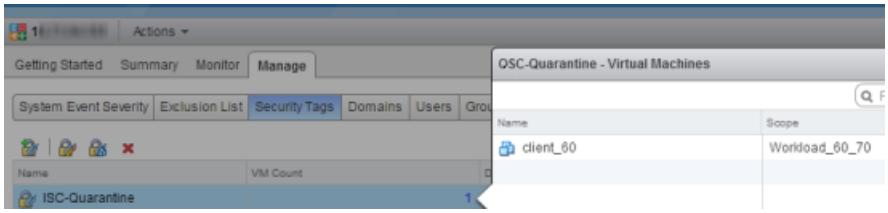

Figure 3-70 VM Count for OSC-Quarantine

You can also check the current virtual machines included in the corresponding security group.

7 If required, release a VM from quarantine.

To release a VM from quarantine, you must remove the *OSC-Quarantine* tag for the VM.

a In vCenter web client, select **Home** | **Networking & Security** | **Networking & Security Inventory** | **NSX Managers**.

b Select the relevant NSX Manager and then select **Manage** | **Security Tags**.

c Select *OSC-Quarantine* and click 

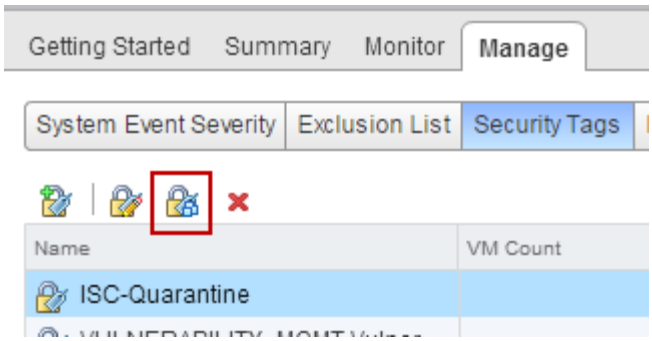

d Deselect the check box for the VMs you want to release from quarantine.

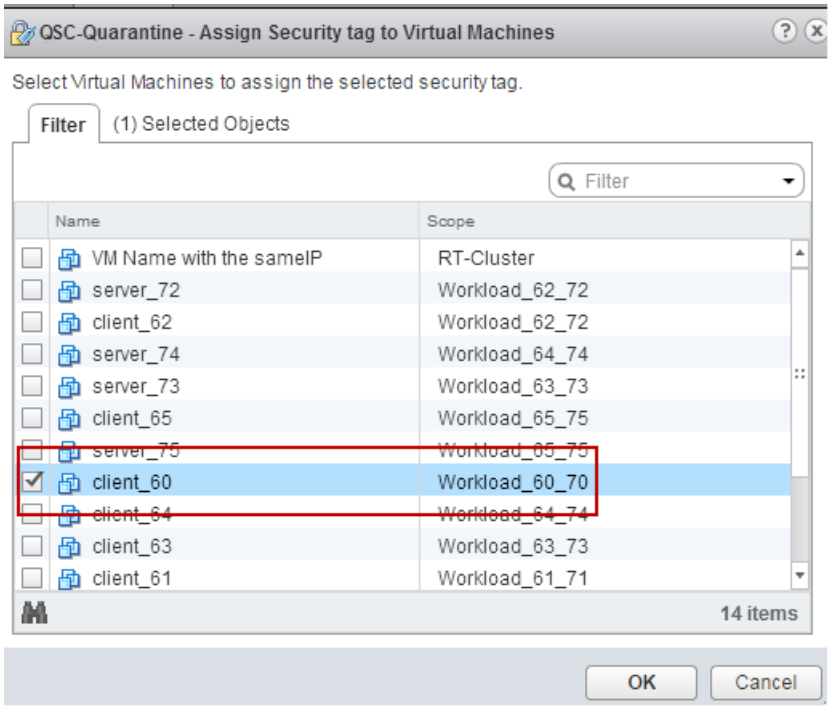

Figure 3-71 Release VMs from quarantine

FAQs regarding IPS service

How to make sure the IPS service is deployed successfully?

1. The installation status is *succeeded* and service status is *up* in the **Installation** page of NSX.
2. In NSX, you have created the security policy with network introspection services for inbound and outbound.
3. In NSX, you have applied the security policy to the security group.
4. The status is *true* for **DISCOVERED** and **INSPECTION-READY** in the **Appliances Instances** page of OSC web application.
5. In the Manager, select **Devices** | **<Admin Domain Name>** | **Devices** | **<Virtual Security System name>** and make sure the Virtual Security System is up-to-date and all the member instances are connected.

I have made sure IPS service is deployed successfully. However, the Sensor is not detecting attack traffic?

- Make sure your deployment meets all the requirements mentioned in Requirements for deploying IPS service on page 105.
- Make sure you have included the protected virtual machines correctly in the security group. See Create a security group in VMware NSX on page 140 .
- Make sure you have configured the security policy as described in Create a security policy in VMware NSX on page 144. Especially, make sure you have selected the correct policy group in the **Profile** list when you add the network introspection service.
- In the Manager, make sure the attack you are testing is enabled in the corresponding policy group.

How do I change a current policy applied on the Virtual Security System instances?

- Consider that you want to apply the same policy group but apply a different IPS policy. Then, select the required IPS policy in the policy group in the Manager.
- To change the applied policy group, do the following:
  1. In NSX, select **Service Composer** | **Security Policies**.
  2. Edit the required security policy and select **Network Introspection Services**.
  3. Edit the inbound or outbound network introspection service to select the required policy group from the **Profile** drop-down list.

  If indicated in the Manager, deploy the configuration changes to the Virtual Security System instances.

I only want to inspect traffic in either inbound or outbound.

You must define network introspection services in NSX for both inbound and outbound. To not inspect outbound traffic for example, apply a policy group with the Default DoS and Reconnaissance Only policy selected for IPS and none for the other policies.

Unable to delete a distributed appliance record in OSC

You must delete the related objects in NSX before you can delete a distributed appliance record. See Deleting a distributed appliance for VMware on page 133.

Manager connector sync and distributed appliance sync jobs fail or **What happens if I rename or delete a policy group in the Manager?**

See *Important note on renaming or deleting policy groups in the Manager* in Manager functions regarding IPS service deployment on page 153.

Index

A

about this guide 5

C

conventions and icons used in this guide 5

D

documentation
- product-specific, finding 5
- typographical conventions and icons 5

M

McAfee ServicePortal, accessing 5

S

ServicePortal, finding product documentation 5

T

technical support, finding product information 5