Using vRealize Network Insight

VMware vRealize Network Insight 3.8

You can find the most up-to-date technical documentation on the VMware website at:

https://docs.vmware.com/

If you have comments about this documentation, submit your feedback to

docfeedback@vmware.com

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Contents

1 About vRealize Network Insight User Guide 6

2 Getting Started 7

Introduction 7

Homepage 9

Navigation 10

Settings 12

3 Configuration 13

Accounts and Data Sources 13

Data Management 24

IP Properties and Subnets 25

Working with Events 27

User Management 30

Syslog Configuration 31

LDAP 32

Configuring Mail Server 34

Support for Simple Network Management Protocol (SNMP) 35

User Session Preferences 36

4 Capacity 37

Planning to Scale up the Platform Cluster 38

Planning to Scale up the Proxy Server 39

5 Clusters 41

Creating Clusters 41

Expanding Clusters 42

6 Entity Pages 43

Virtual Machine Topology 44

Hosts 44

VXLAN 45

VLAN 45

NSX Manager 45

7 System Dashboards 48

NI System Dashboard 48

NI Platform Dashboard 50

NI Collector Dashboard 51
Data Source Dashboard 52
Viewing the PCI-Compliance Dashboard 53
Viewing the Flow Analytics Dashboard 56
Analytics - Outlier Detection 58

8 Network and Security 61
Network Visibility 61
Security 64

9 Operations 73
Working with Events 73
Notifications 76
Pins and Pinboards 79

10 Flows 83
Enabling IPFIX Configuration 83
Flow Support for Physical Servers 86
View Blocked and Protected Flows 92
Network Address Translation (NAT) 93
VPC Flow Logs 95

11 Micro-Segmentation 97
Micro-Segmentation Dashboard 97
Micro-Segmentation Planning Topology 98
Application-Centric Micro-Segmentation 100
Recommended Firewall Rules 102

12 Search 105
Simple Search 105
Advanced Search 106
Time Control 107
Search Results 108
Filters 108

13 Upgrade 110
Online Upgrade of Product 110
Single-Click Offline Upgrade 110

14 Troubleshooting 112
Health 112
Support Tunnel 113

Online Upgrade of Product 113
Single-Click Offline Upgrade 113
View Node Details 114
Capacity 115
Disk Management 116
Creating Clusters 116
Expanding Clusters 117
Creating Support Bundle 118
Migrating Data Sources 118
About Page 119
Common Data Source Errors 120

About vRealize Network Insight User Guide 1

The *vRealize Network Insight User Guide* provides information about using vRealize Network Insight.

Intended Audience

This information is intended for administrators or specialists responsible for using vRealize Network Insight. The information is written for experienced virtual machine administrators who are familiar with enterprise management applications and datacenter operations.

VMware Technical Publications Glossary

VMware Technical Publications provides a glossary of terms that might be unfamiliar to you. For definitions of terms as they are used in VMware technical documentation, go to http://www.vmware.com/support/pubs.

Getting Started

2

This chapter includes the following topics:

- Introduction
- Homepage
- Navigation
- Settings

Introduction

VMware vRealize Network Insight delivers intelligent operations for software-defined networking and security. It helps customers build an optimized, highly-available, and secure network infrastructure across multi-cloud environments. It accelerates micro-segmentation planning and deployment, enables visibility across virtual and physical networks, and provides operational views to manage and scale the VMware NSX deployments.

Think of your entire data center as being composed of entities and their relationships. As an example, a virtual machine is an entity, and the virtual machine is part of a Host which is another entity. vRealize Network Insight provides visibility and information on numerous entities that are part of your data center.

Table 2-1.

Entities	Description
	Host
	Problem
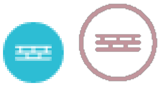	NSX Firewall
	Virtual Machine

Table 2-1. (Continued)

Entities	Description
	vSphere Distributed Switch
	Physical Switch
	Virtual Port Group
	Cisco Fabric Extender
	Logical Switch
	Datastore
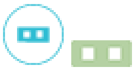	Physical Network Interface Card
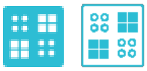	Security Group
	Blade
	Router
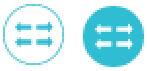	VLAN
	Group of VMs
	Configuration Changes
	Router Interface
	Troubleshoot

Table 2-1. (Continued)

Entities	Description
NAT	Network Access Translation (NAT)
✉	Mail Server

Homepage

The VMware vRealize Network Insight homepage provides you a quick summary of what is happening in your entire data center. It provides you a quick access to the important components of vRealize Network Insight of your data center.

The homepage is divided into the following sections:

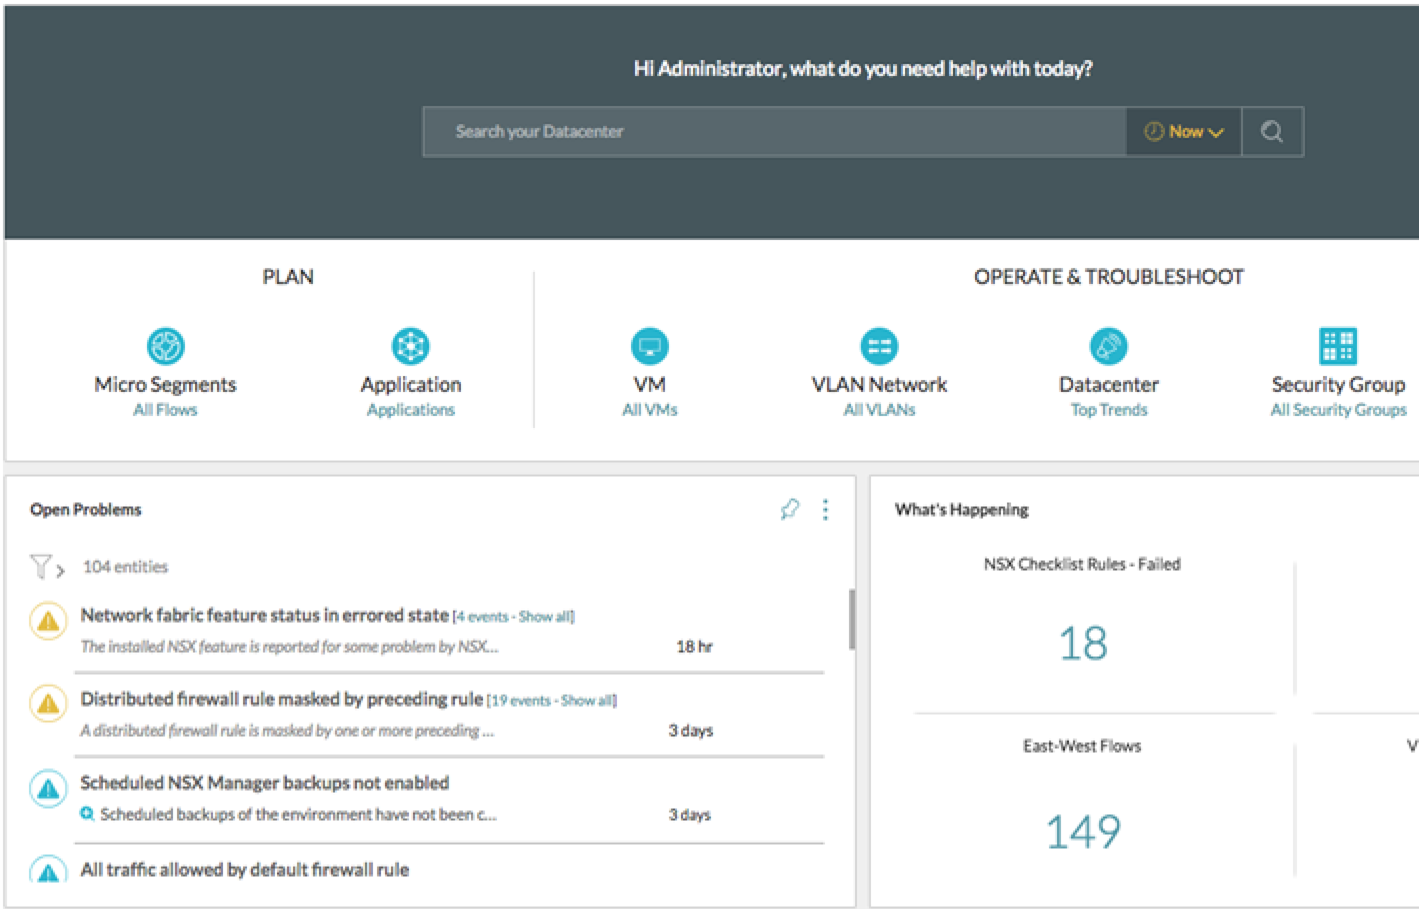

Network fabric feature status in errored state [4 events - Show all]
The installed NSX feature is reported for some problem by NSX... 18 hr

Distributed firewall rule masked by preceding rule [19 events - Show all]
A distributed firewall rule is masked by one or more preceding ... 3 days

Scheduled NSX Manager backups not enabled
Scheduled backups of the environment have not been c... 3 days

All traffic allowed by default firewall rule

Search Bar

The Search bar provides you the ability to search across your data center network (and its corresponding entities). You can use the search bar to search for the entities that are available in your data center. The search bar is available at the top of the homepage.

Based on your requirement, you can perform search as per the following time line options:

- **Presets**: Using this option, you can narrow down your search results for presets such as `last week`, `last 3 days`, `last 24 hours`, `yesterday`, `today`, `last 2 hours`, `last hour`, and now (current time).
- **At**: Using this option, you can narrow down your search results for a particular date and time.
- **Between**: Using this option, you can search for data between a particular time interval.

Plan Section

- **Micro Segments**: You can plan the micro-segmentation of the network based on the flows between all the VMs.
- **Application**: You can define your applications and analyse their flows, and plan their security.

Operate and Troubleshoot Section

The **Operate and Troubleshoot** section provides visibility, metrics, and analytics for the following components:

- Virtual Machine (VM)
- VLAN Network
- Data Center
- NSX Security Group
- VMware NSX

Open Problems

The **Open Problems** section provides a quick glance of the critical events that the platform finds in your data center. All such similar events are grouped. Use **Show All** to view all the events. To view more details of an event, click 🔍 (**View Details**).

What's Happening

The **What's Happening** section provides a quick view of very high-value properties from your data center. To view the property details, click the count of a particular property. This section also contains filters on the left side to filter the events, and expand all and collapse all buttons to view the details of the events.

Navigation

vRealize Network Insight contains a navigation panel on the left that helps users to navigate quickly to the key product features such as Security, Topologies, Entities, Events, and Saved Searches of interest without having to type any search queries.

The Navigation Panel contains the following options:

- Security: Provides you the following options:
  - Plan Security: Allows you to analyze the flows in the environment and helps to plan the micro-segments within the environment. You can select all the entities or select a particular entity and then select the duration to analyze the selected entity.
  - Applications: Allows you to create applications in vRealize Network Insight by using custom search. Once you create an application, you can plan it accordingly.
  - PCI Compliance: The PCI-Compliance dashboard helps in assessing compliance against the PCI requirements only in the NSX environment.
- Path and Topology: Allows you to view any VM to VM path or topology of several entities of the data center.
- Events: Allows you to view the events (changes and problems) in your environment. There is also a list of event types so that you can quickly view a specific type of event.
- Entities: Displays the list of all the different types of entities present in your environment. Click any entity type from the given list to view a list of all the entities of that type. The text box above the entities list can be used to narrow down the list based on text entered.
- Saved Searches: Displays the searches that have been saved previously.

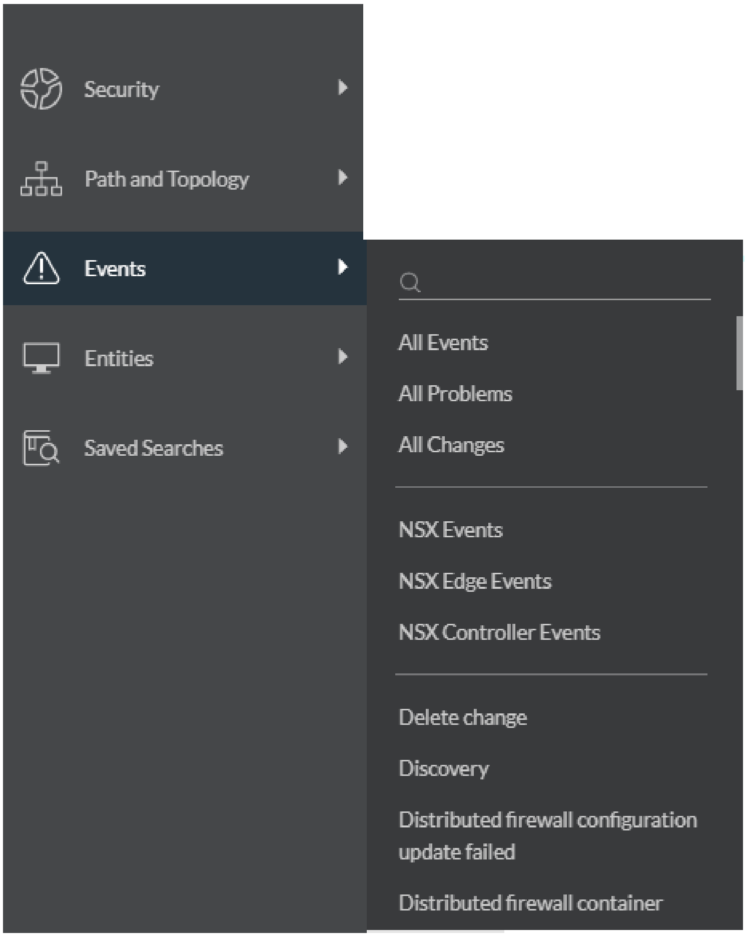

Settings

The **Settings** page provides controls to manage data providers, users, and notifications.

To go to the **Settings** page:

1 On the top-right hand corner in the Home page, click the Profile icon.

2 Click **Settings**. The **Settings** page appears as shown.

You can configure the following on the **Settings** page:

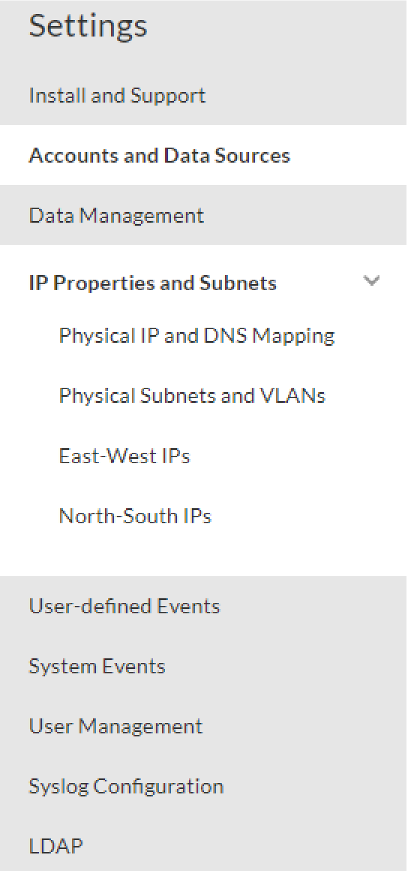

Settings

Install and Support

Accounts and Data Sources

Data Management

IP Properties and Subnets

Physical IP and DNS Mapping

Physical Subnets and VLANs

East-West IPs

North-South IPs

User-defined Events

System Events

User Management

Syslog Configuration

LDAP

Configuration

3

This chapter includes the following topics:

- Accounts and Data Sources
- Data Management
- IP Properties and Subnets
- Working with Events
- User Management
- Syslog Configuration
- LDAP
- Configuring Mail Server
- Support for Simple Network Management Protocol (SNMP)
- User Session Preferences

Accounts and Data Sources

Data sources provide the application the ability to gather data from certain aspects of your data center. These range from your NSX installation to physical devices such as Cisco[TM] Chassis 4500 and Cisco[TM] N5K.

For each added data source, you can view the following information:

- All: Displays all the available data sources.
- With Problems: Displays the data sources where vRealize Network Insight has found a problem.
- With Recommendations: Displays auto generated recommendations from vRealize Network Insight for the data sources that require additional information.
- Disabled: Displays the data sources that have been disabled.

For each data source, you can view the following details:

Table 3-1.

Properties	Description
Type(Nickname)	Displays name of the Data source.
IP Address/FQDN	Displays IP address or FQDN details for the Data Source.
Last Collection	Displays the last collection time on which the data is collected.
Discovered VMs	Displays the number of VMs that have been discovered for that data source.
Collector VM	Displays the name of the collector to which the data source has been added. This column is not visible if all the listed data sources have been added on the same collector. You can view this column only if the data sources are present on different collectors.
Enabled	Indicates if the data source is enabled or not.
Actions	Displays options to edit and delete the data source.

vRealize Network Insight provides the following functions to enable easy access to the information of data sources.

- You can perform search for a data source by its name, its IP address, or by the collector VM name by using the search bar above the column headers.
- You can filter information by different data sources in the **Type(Nickname)** column.
- You can filter information by various collector VMs in the **Collector VM** column.
- The data sources are sorted by their types and nicknames in the alphabetical order.

Note The Discovered VMs column is populated only if the data source is vCenter or AWS source.

Adding a Data Source

To add a Data Source

1. In the **Install and Suport** page under **Settings**, click **Accounts and Data Sources**.
2. Click **Add new source**.
3. Select an account or a source type.
4. Provide the following information:

Table 3-2.

Properties	Description
Collector(Proxy) VM	Select the proxy VM from the drop-down menu.
IP Address/FQDN	Enter the IP Address or the FQDN details.
`Username`	Enter the user name you want to use for a particular data source.
`Password`	Enter the password.

5 After entering the information in the text boxes, click **Validate**.

- When you are adding a VMware vCenter or an AWS data source, if the number of VMs discovered for a specified data source exceeds the capacity of the platform or a proxy node or both, the validation fails. You will not be allowed to add a data source until you increase the brick size of the platform or create a cluster.

  The specified capacity for each brick size with and without flows is as follows:

  Table 3-3.

Brick Size	VMs	State of Flows
Large	6k	Enabled
Large	10k	Disabled
Medium	3k	Enabled
Medium	6k	Disabled

- If the validation is successful, you can add advanced data collection sources for the data source (not all data sources contain this feature). Following advanced data collection sources are available:

  - For VMware vCenter, you can enable NetFlow (IPFIX). For more information on IPFIX, read the Enabling IPFIX configuration on VDS and DVPG section.
  - For VMware NSX Manager, you can enable automatic NSX Edge Population using SSH to allow vRealize Network Insight to collect advanced data. However, for NSX Manager 6.2 and above, use NSX central CLI instead of `ssh`. You can select this option to allow vRealize Network Insight to collect data for NSX Edge directly from NSX Manager using the NSX Central CLI. This feature also requires NSX Manager credentials with System Admin privileges.
  - Many data sources also use SNMP (Simple Network Management Protocol) for richer data collection. For such data sources, select the SNMP version and enter the community string to allow vRealize Network Insight to collect richer data from the data source.

6 Enter the required details in the text boxes for advanced data collection sources.

7 Enter Nickname and Notes (if any) for the data source and click **Submit** to add the data source to the environment.

Adding an AWS Data Source

To add an AWS data source:

Prerequisites

- The custom policy of the AWS account user to add AWS data source is as follows:

  ```
  {
      "Version": "2012-10-17",
      "Statement": [
          {
              "Effect": "Allow",
              "Action": [
                  "iam:ListAccountAliases"
              ],
              "Resource": [
                  "*"
              ]
          },
          {
              "Effect": "Allow",
              "Action": [
                  "ec2:Describe*"
              ],
              "Resource": "*"
          },
          {
              "Action": [
                  "logs:Describe*",
                  "logs:Get*",
                  "logs:TestMetricFilter",
                  "logs:FilterLogEvents"
              ],
              "Effect": "Allow",
              "Resource": "*"
          }
      ]
  }
  ```

- There are a list of URLs which should be accessible from the Collector VM to access AWS. The AWS can be deployed in multiple regions. There are separate URLs associated with different regions. If you are unaware of the region or the service, have a wildcard entry for the URL such as `*.amazonaws.com`.

  Note The wildcard entry does not work for the China region.

  But if you want to give fine-grained access to separate URLs, there are 4 services based on the region:

  Regions except GovCloud and China

  - `ec2.<REGION>.amazonaws.com`
  - `logs.<REGION>.amazonaws.com`
  - `sts.<REGION>.amazonaws.com`

- `iam.amazonaws.com`

GovCloud Region

- `ec2.us-gov-west-1.amazonaws.com`
- `logs.us-gov-west-1.amazonaws.com`
- `sts.us-gov-west-1.amazonaws.com`
- `iam.us-gov.amazonaws.com`

China (Beijing) Region

- `ec2.cn-north-1.amazonaws.con.cn`
- `logs.cn-north-1.amazonaws.com.cn`
- `sts.cn-north-1.amazonaws.com.cn`
- `iam.cn-north-1.amazonaws.com.cn`

You can use any of the following values for `REGION` based on the AWS region:

Region Name	Region
US East (Ohio)	`us-east-2`
US East (N. Virginia)	`us-east-1`
US West (N. California)	`us-west-1`
US West (Oregon)	`us-west-2`
Asia Pacific (Mumbai)	`ap-south-1`
Asia Pacific (Seoul)	`ap-northeast-2`
Asia Pacific (Singapore)	`ap-southeast-1`
Asia Pacific (Sydney)	`ap-southeast-2`
Asia Pacific (Tokyo)	`ap-northeast-1`
Canada (Central)	`ca-central-1`
EU (Frankfurt)	`eu-central-1`
EU (Ireland)	`eu-west-1`
EU (London)	`eu-west-2`
South America (São Paulo)	`sa-east-1`
Gov Cloud	`us-gov-west-1`
China (Beijing)	`cn-north-1`

Procedure

1. Select **Account/Data Sources**. Click **Add Source**.
2. Under **Public Clouds**, click **Amazon Web Services**.

3 Add your AWS account by using Amazon Access Key ID and corresponding Secret Access Key.

Note Your Amazon Access Key ID is a 20-digit string with a corresponding Secret Access Key. For more details, see http://docs.aws.amazon.com/general/latest/gr/managing-aws-access-keys.html.

Note To add AWS Gov Cloud Region as a data source, create an AWS IAM user by using the recommended policy in the AWS account with access to the Gov Cloud region. Use the Access key and the Secret key for the newly created account to add the data source to vRealize Network Insight.

This process might take 15–20 minutes for adding and displaying your account data.

4 After you have validated your AWS account, you can select **Enable Flows data collection** to get deeper insights.

Adding an Infoblox DNS Data Source

Infoblox DNS offers an advanced solution to manage and control DNS. It uses Infoblox Grid to ensure that the DNS is highly available throughout the network. vRealize Network Insight allows the users to add Infoblox Grid as a DNS data provider. The DNS data from Infoblox is used only for enriching the flows where either the source or the destination IP addresses are associated with the physical devices. This feature is available only for the enterprise license.

The Infoblox DNS data co-exists with the DNS data that is imported by using CSV.

If you configure an Infoblox DNS data source on a proxy server, you can configure other data sources also on the same proxy server. You do not need a dedicated proxy server for Infoblox.

Considerations

- vRealize Network Insight supports only single-grid mode for Infoblox in the current release.
- Only A Records are supported in the current release. Shared A Records are not supported currently.
- The DNS enrichment is supported only for the IP addresses that are marked as physical in the current release.
- If there are multiple FQDNs for a single physical IP address, all FQDNs are returned.

Procedure

1 On the **Settings** page, click **Accounts and Data Sources**.

2 Click **Add new source**.

3 Click **Infoblox** under **DNS**.

4 Provide the following information:

Table 3-4.

Properties	Description
`Collector(Proxy) VM`	Select the proxy VM from the drop-down menu.
`IP Address/FQDN`	Enter the IP Address/FQDN of Infoblox Grid Master.

Table 3-4. (Continued)

Properties	Description
`Username`	Enter the user name you want to use for a particular data source.
`Password`	Enter the password.

5 Click **Validate**.

Note Ensure that you have the `API Privilege` to access the Infloblox APIs.

6 Enter **Nickname** and **Notes** (if any) for the data source and click **Submit** to add the Infoblox DNS data source to the environment.

Enriching Flows with the Infoblox DNS Data

vRealize Network Insight supports two sources of DNS information:

- Imported CSV file
- Infoblox DNS

Note If there is a conflict between Infoblox DNS and the CSV file, the information from Infoblox DNS takes precedence.

You can use various search queries to find out more about the source of DNS entries in a flow.

Table 3-5.

Keyword	Sample Search Query	Description
DNS Provider	`Flows where DNS Provider='Infoblox'`	Provides the list of flows in which the DNS data is obtained from Infoblox.
DNS Provider	`Flows where DNS Provider='CSV'`	Provides the list of flows in which the DNS data is obtained from CSV.
Source DNS Provider	`Flows where Source DNS Provider='Infoblox'`	Provides the list of flows in which the DNS provider for the source IP address is Infoblox.
Destination DNS Provider	`Flows where Destination DNS provider='Infoblox'`	Provides the list of flows in which the DNS provider for the destination IP address is Infoblox.

Adding vRealize Log Insight as a Data Source

In the vRealize Network Insight and vRealize Log Insight integration, the alerts generated by vRealize Log Insight are consumed by vRealize Network Insight. Whenever a security group is created or modified, the logs of NSX are sent to vRealize Log Insight which in turn sends an alert. After receiving the alert, vRealize Network Insight polls the NSX Manager on which the security group was created and fetches the corresponding data for the changed security groups. Currently, this integration supports only the security group CRUD-related alerts.

The vRealize Network Insight and vRealize Log Insight integration supports the following versions:

- VMware vRealize Log Insight 4.5 and later versions
- vRealize Network Insight v3.8 and later versions
- VMware NSX Manager v6.2 and later versions

Prerequisites

- You should have installed the vRealize Log Insight server in your environment.
- You should have configured the NSX Manager to send the logs to vRealize Log Insight.
- You should have installed the NSX Content Pack for vRealize Network Insight in vRealize Log Insight.

Procedure

1. Create or reuse a vRealize Log Insight user with access to the APIs of vRealize Log Insight.
2. On the **Install and Support** page, click **Accounts and Data Sources**.
3. Click **Add Source**.
4. Click **Log Insight** under **Log Servers**.

5 On the **Add a New Log Insight Server Account or Source** page, click **Instructions** next to the page title. A pop-up window appears that provides the prerequisites for adding the Log Insight data source and the instructions to enable the Webhook URL on vRealize Log Insight.

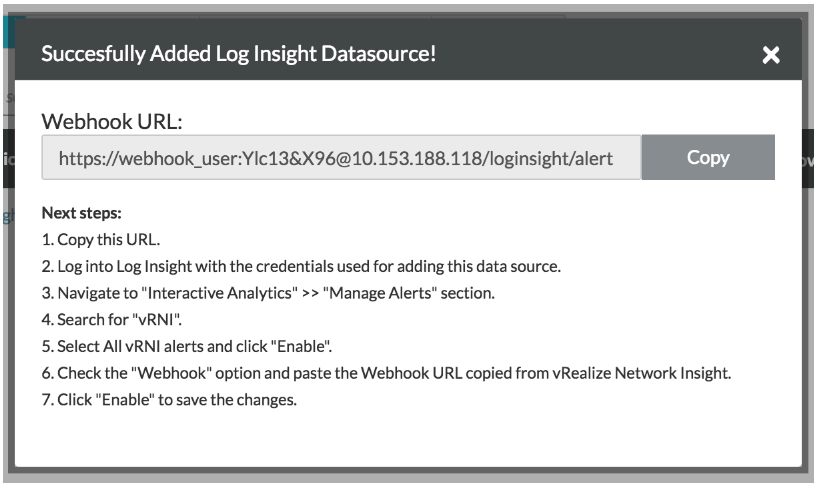

Succesfully Added Log Insight Datasource!

Webhook URL:

https://webhook_user:Ylc13&X96@10.153.188.118/loginsight/alert

Next steps:

1. Copy this URL.
2. Log into Log Insight with the credentials used for adding this data source.
3. Navigate to "Interactive Analytics" >> "Manage Alerts" section.
4. Search for "vRNI".
5. Select All vRNI alerts and click "Enable".
6. Check the "Webhook" option and paste the Webhook URL copied from vRealize Network Insight.
7. Click "Enable" to save the changes.

Note The Webhook URL, which is generated after the addition of the data source, is used in vRealize Log Insight.

6 Enter the required details.

Name	Description
Collector (Proxy) VM	Select the IP address of the data collector that you have deployed for the data collection process.
IP Address / FQDN	Enter the IP address or the FQDN of the data source.
User Name	Enter the user name you want to use for a particular data source.
Password	Enter the password for the data source.
Authentication Provider	Select the respective authentication provider for the credentials that you have provided.

7 After the data source has been created, a pop-up window appears that will provide the Webhook URL and the steps that have to be performed to enable this URL on vRealize Log Insight. Copy the Webhook URL. Log in with the credentials that were used for adding this data source. Enable alerts in the vRealize Log Insight application and configure this Webhook URL. Send Test Alert to ensure that the integration is successful.

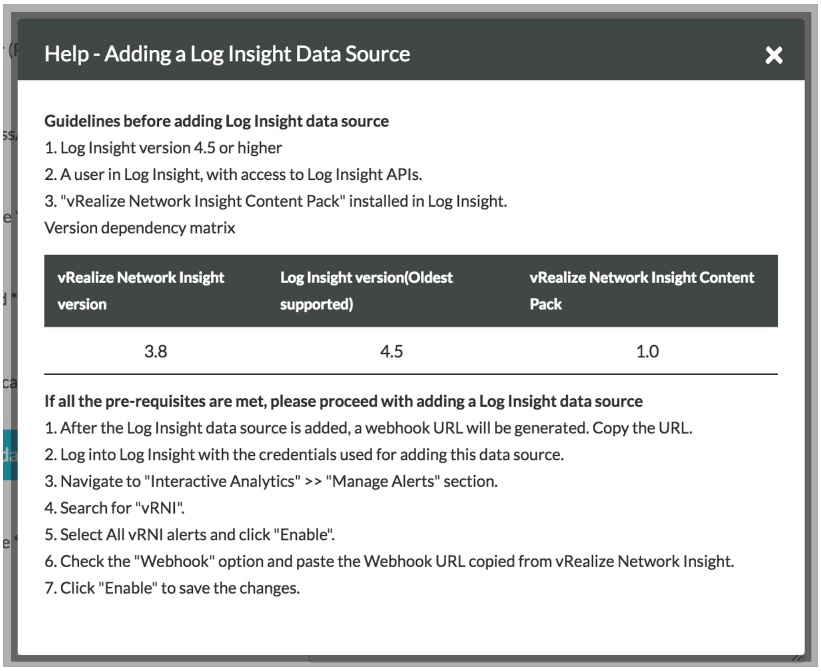

Help - Adding a Log Insight Data Source

Guidelines before adding Log Insight data source
1. Log Insight version 4.5 or higher
2. A user in Log Insight, with access to Log Insight APIs.
3. "vRealize Network Insight Content Pack" installed in Log Insight.
Version dependency matrix

vRealize Network Insight version	Log Insight version(Oldest supported)	vRealize Network Insight Content Pack
3.8	4.5	1.0

If all the pre-requisites are met, please proceed with adding a Log Insight data source
1. After the Log Insight data source is added, a webhook URL will be generated. Copy the URL.
2. Log into Log Insight with the credentials used for adding this data source.
3. Navigate to "Interactive Analytics" >> "Manage Alerts" section.
4. Search for "vRNI".
5. Select All vRNI alerts and click "Enable".
6. Check the "Webhook" option and paste the Webhook URL copied from vRealize Network Insight.
7. Click "Enable" to save the changes.

Note Any alert displayed on the vRealize Log Insight data source in vRealize Network Insight is resolved in an hour.

Add Cisco ASA as a Data Source

You can add Cisco ASA-X series as a data source in vRealize Network Insight.

Note Considerations

- vRealize Network Insight supports only routed mode.
- vRealize Network Insight supports only Cisco ASA-X series.
- Currently, vRealize Network Insight supports Cisco ASA operating system version 9.4.
- vRealize Network Insight does not support the cluster deployment of Cisco ASA.
- vRealize Network Insight does not support the high availability of Cisco ASA.
- vRealize Network Insight does not support Cisco ASA if it is directly connected to the host. A topology that is similar to the following one is supported:

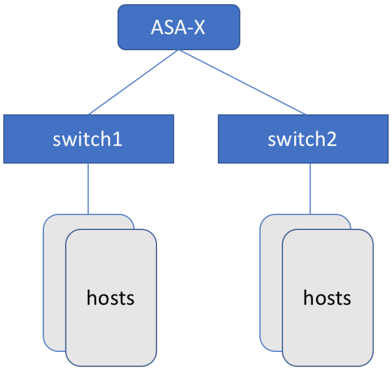

Procedure

1. In the **Install and Suport** page under **Settings**, click **Accounts and Data Sources**.
2. Click **Add new source**.
3. Select **Routers & Switches**.
4. Select **Cisco ASA**.
5. Provide the following information:

   Table 3-6.

Properties	Description
Collector(Proxy) VM	Select the proxy VM from the drop-down menu.
IP Address/FQDN	Enter the IP Address or the FQDN details.
`Username`	Enter the user name you want to use for this data source.
`Password`	Enter the password. Ensure that you enter the same password as the one that you used for the enable mode of Cisco ASA.

6. After entering the information in the text boxes, click **Validate**.

Edge Data Collection

Whenever you add an NSX data source, you can enable the automatic edge data collection. In the previous releases, the edge data collection was done either by NSX Central CLI or Edge-SSH session. From the current release onwards, the edge data collection is done by NSX Central CLI. So no edge data providers are created under NSX Manager.

Note Validation of NSX User Privileges

While adding the NSX data source and enabling the edge population, the NSX user privileges are validated.

- In NSX 6.4 and the further releases, an enterprise admin user can run the NSX Central CLI commands. The user credentials that need to be specified while adding NSX Manager as a data source must be of an enterprise admin or super user.
- In NSX 6.2 and the further releases previous to NSX 6.4, the user should be a super user to enable the edge data population. The user credentials that need to be specified while adding NSX Manager as a data source must be of an enterprise admin or super user.
- Suppose that a user has the enterprise admin privilege in NSX 6.3 and is working on the current release of vRealize Network Insight, an `Insufficient Privileges` error comes up on the **Accounts and Data Sources** page for **VMware NSX Manager**. The error is shown because the user has to be a super user to run the NSX Central CLI commands in NSX 6.3.

Table 3-7.

NSX Version	User
NSX 6.4 and the further releases	■ The user credentials that need to be provided while adding NSX Manager as a data source must be of an enterprise admin or super user. ■ An enterprise admin user can run the NSX Central CLI commands required by vRealize Network Insight.
NSX 6.2 and the further releases before NSX 6.4	■ The user should be a super user to enable the edge data population. ■ The user credentials that need to be provided while adding NSX Manager as a data source must be of an enterprise admin or super user.

Data Management

In vRealize Network Insight, you can specify for how long do you want to retain your data.

Note vRealize Network Insight supports configurable data management on an enterprise license only. In the advanced license edition, the data retention defaults to 1 month.

The data is divided into the following categories:

Table 3-8.

Category	Minimum Value	Maximum Value
Events	1 month	13 months
Entities and Configuration Data	1 month	3 months
Metrics	1 month	13 months

Table 3-8. (Continued)

Category	Minimum Value	Maximum Value
Flows	NA	1 month
Miscellaneous Data	NA	100 GB of additional disk space

Note For all the categories, the minimum value is the default value.

Different policies can be configured and controlled for each category. You can configure the policy as per your requirement.

To configure data management:

1. On the top-right corner of the Home page, click ⚙ and then click **Settings**.
2. In the **Settings** section, click **Data Management**.
3. When you log in for the first time, this page shows the default data.
4. Click the information icon on more information on how data occupies the disk.
5. Click **Change Policy** to change the data retention period for the various categories of data. Once you make the changes, the information is recorded in the database.
6. Click **Submit**.

Note The retention period for low-resolution metrics is longer than the high-resolution metrics.

IP Properties and Subnets

Physical IP and DNS Mapping

To provide the information for the flows between physical devices, you can import the DNS mapping file. The supported formats for the DNS mapping file are the Bind and CSV file format. Ensure that you have placed these files in a single ZIP file.

Note vRealize Network Insight does not support the password-protected ZIP files.

Procedure

1. In the **Settings** page, click **IP Properties and Subnets.**.
2. Click **Physical IP and DNS Mapping**.

3 Click **Upload and Replace** to upload your DNS mapping file. After you select and upload the file, click **Validate**. The number of DNS records is displayed after the validation.

   The **Upload and Replace** operation removes any existing DNS mappings and replaces them with the the mappings that are being imported. The DNS Mapping file consists of the following three fields:

   - Host Name
   - IP Address
   - Domain Name

Physical Subnets and VLANs

You can define a mapping between subnet and a VLAN.

You can use this mapping for the following:

- Enriching the information about the IP entities that are learned from physical to physical flows by adding the source and destination subnets and the Layer2 networks associated with the flow.
- Planning the network topology based on the subnet and VLAN for physical addresses.

Procedure

1 In the **Settings** page, click **IP Properties and Subnets.**.

2 Click **Physical IP and DNS Mapping**.

3 In the **Settings** page, under **IP Properties and Subnets**, click **Physical Subnets and VLANs**.

   This page lists all the subnets and the associated VLAN IDs.

4 Click **Add** to add the subnet and VLAN information.

5 After defining the mapping information, you can only edit the VLAN ID that is associated with the subnet. It is not possible to change to the subnet CIDR associated with the VLAN Id. To edit a subnet associated with the VLAN ID, delete the subnet to be edited and create a subnet VLAN mapping with the required values.

   When the subnet-VLAN mapping information is updated, a new VLAN is created for the specified VLAN ID and the subnet information is associated with this VLAN.

6 To delete the subnet-VLAN ID mapping, click the delete icon.

   Note All VLAN creation, updation, and deletion operations do not happen immediately after the subnet and VLAN mappings are created. It takes some time for the changes to be propagated and the corresponding VLAN being to be created or modified.

East-West IPs

The IPs that are within the range of RFC1918 standard are considered private IPs. The IPs that are outside the RFC1918 are treated as Internet IPs. However, users can specify their East-West IPs (datacenter public IPs) that they want to be treated as non-Internet IPs while tagging flows and micro-segmentation, even if these are outside the private IP address range as defined by RFC1918.

To specify public IPs to be treated as non-internet IPs

1. On the top-right corner of Home page, click the Profile icon, and then click **Settings**.
2. In the Settings section, click **East-West IPs**.
3. In the IP Addresses box, enter specific IPs, or IP ranges, or subnets, which are to be treated as non-internet IPs.
4. Click **Save**. The IP Addresses Saved confirmation message is displayed upon successful saving.

North-South IPs

The IPs that are in the RFC1918 space are categorized as North-South IPs. The users can specify their North-South IPs while tagging flows and micro-segmentation.

To specify North-South IPs:

1. On the top-right corner of Home page, click the Profile icon ⚙, and then click **Settings**.
2. In the Settings section, click **North-South IPs**.
3. In the IP Addresses box, enter specific IPs, or IP ranges, or subnets.
4. Click **Save**. The IP Addresses Saved confirmation message is displayed upon successful saving.

Working with Events

On the **Settings** page, click **Events** to view the various types of events:

- System Events
- User Defined Events
- System Health Events

Working with System Events

The event is defined either by the system or the user. The system events are predefined events.

The system events are listed in the System Events page under Settings. The following fields are specified for each event. You can filter the information based on your requirements in all the following columns except the Event column.

Table 3-9.

Column	Description
Event	This field specifies the name of the event.
Severity	This field specifies the severity of the event. You can set it to the following values: - Critical - Moderate - Warning - Info
Type	This field specifies if the event denotes a problem or a change.
Entities	This field specifies that the event is configured to either include or exclude entities for event generation. By default, the value is `All`.
Notifications	This field specifies the types of notifications that are sent. The notifications can be sent by email or SNMP trap or both.
Enabled	This option is selected if the event is enabled.

When you hover the mouse on each event, you can see More Information. By clicking this option, you can see the description, event tags, and entity type for that event.

You can perform the following tasks on the system events:

- Edit an event
- Perform bulk edit
- Disable an event for a particular entity

Edit System Event

To edit an event:

1 Click the edit icon after the **Enabled** column for a particular event.

2 You can add or remove event tags if required.

3 You can change the severity.

4 Check Include/Exclude entities if you want the event to be enabled or disabled for selected entities.

5 To create inclusion rules:

   a Select **Inclusion List**.

   b Specify the entities which you want to include for the event under **Conditions**.

6 To create exclusion rules:

   a Select **Exclusion List**.

b Specify the entities which you want to exclude for the event under **Conditions**.

Note

- You can create multiple rules in both inclusion and exclusion lists.
- When you select `NSX Manager`, you can add exceptions in both the lists. You can define exception if you want the inclusion or the exclusion rule to hold exception for a particular entity.
- You can also specify `Custom Search` by writing your own query to include or exclude entities.

7 Select **Enable Notifications** if you want to configure when the notifications have to be sent. Specify the email address and the frequency at which you would like to receive the emails.

Prerequisites

Perform a Bulk Edit on an Event

1 In the **System Events** page, when you select multiple events, the options **Enable**, **Disable**, and **Edit** appear above the list.

2 Click **Edit**.

3 In the **Edit** page, you have the following options:

- **Override existing values**: In this option, only the fields that you edit will get overwritten.
- **Add to existing**: In this option, you can add to the existing values such as email addresses and event tags.

4 Click **Submit**.

Disable an Event

1 You can select an event in the **Open Problems** widget in the Homepage. You can also enter **Problems** in the search bar and select an event from the list.

2 Select a particular event and click **Archive**.

3 Select **Disable all events of this type in future for** and select an entity or all entities.

4 Click **Save**.

Note The changes made in severity, tags, or inclusion/exclusion rules will reflect for the future events. The existing events continue to show the old configuration.

Working with User-Defined Events

The user-defined events are based on search.

All the user-defined events are listed on the **User-defined Events** page under **Settings**. The following fields are specified for each event.

Table 3-10.

Field	Description
Name (Search Criteria)	This field specifies the name of the event and the search criteria for the event.
Severity	This field specifies the severity of the alert. You can set it to the following values: ■ Critical ■ Moderate ■ Warning ■ Info
Type	This field specifies if the event denotes a problem or a change.
Notify when	This field specifies when the notification has to be sent.
Created By	This field specifies who created the event.
Enabled	This option is selected if the event is enabled.

You can edit or delete the event. While editing it, you can specify the email address and the frequency of the email notification.

Working with Platform Health Events

The Platform Health Events page is your one-stop page to view all the events that provide details on the overall health of the system. These events might have occurred on a datasource or a node in the infrastructure. You can also view these events through search.

Table 3-11.

Field	Description
Event	This field specifies the name of the event.
Severity	This field specifies the severity of the event. You cannot change the severity of the event.
Type	This field specifies if the event denotes a problem or a change.
Notifications	This field specifies the types of notifications that are sent. The notifications can be sent by email or SNMP trap or both.

User Management

The admin user can add new users and configure memberships and other settings of existing users. The users with membership role of administrator only can view the **User Management** tab.

Add New User

1 In the **Settings** page, click **Create new user**, and provide the required information in the form.

  The form has the following text boxes:

Table 3-12.

Properties	Description
Name	Enter the name of the user.
Email (Login ID)	Enter your email or login ID if any.
Role	Select the role from drop-down list.
Password	Enter the password.
Re-enter new password	Re-enter the password for confirmation.

2 Click **Add User** to save the user information.

Assign Administrator Role

You can assign an administrator role to any LDAP user.

Even if that particular user is not logged in, you can still assign the administrator role to that user. To assign the administrator role:

1 In the **Settings** page, click **User Management**.

2 Click the **LDAP Users** tab.

3 Click **Assign Admin Role**.

4 Provide the login ID of the user to whom you want to assign the administrator role.

5 Click **Add User**.

6 Once you add the user, you can see the login ID in the LDAP Users tab.

7 To change the role, click the edit icon next to the login ID in the LDAP Users tab.

Syslog Configuration

You can configure remote syslog servers for vRealize Network Insight by using the **Syslog Configuration** page.

While every proxy server can potentially have a different remote syslog server, all the platform servers in a cluster use the same remote syslog server.

In the current release, the vRealize Network Insight problem events and platform/proxy server syslogs are sent to the remote syslog server.

Currently, vRealize Network Insight supports only UDP for communication between vRealize Network Insight servers and remote syslog servers. So ensure that your remote syslog servers are configured to accept syslog traffic over UDP.

To configure syslogs:

1 In the **Settings** page, click **Syslog Configuration**. The **Syslog Configuration** page has the configured syslog servers and their mappings to the virtual appliances listed. If you are accessing this page the first time, then the syslog is disabled by default and the list of servers on this page does not appear.

2 To add a syslog server:

   a Click **Add Syslog Server**.

   b Enter IP Address, nickname, and port number of the server. The standard port number for UDP is `514`.

   c To test the configuration, click **Send Test Log**.

   d Click **Submit**.

   e If it is the first server that you have added, then enable syslog at the top of the page.

3 To map the server to platforms and proxies:

   a Click **Edit Mapping**.

   b Select the syslog server for All Platforms and Proxy servers.

   c If you do not want to enable syslog on any proxy server or on the platform, select the **No server** option.

   d Click **Submit**.

   Note After you make the changes, it might take a few minutes for them to be effective.

LDAP

vRealize Network Insight supports the following two types of users:

- User created on vRealize Network Insight Platform VM
- LDAP users

To allow the LDAP users log into vRealize Network Insight, configure the LDAP service in the vRealize Network Insight Platform as follows:

To Enable LDAP-Based User Authentication

1 On the **Settings** page, click **LDAP**.

2 Click **Configure**.

3 On the **Configure LDAP** page, type the appropriate domain, LDAP Host URL, and LDAP credentials in the respective boxes. See the following table for individual field descriptions.

Table 3-13.

Field	Description
Domain	This is typically the last part of the user email address after the '@' sign. Example: For a user logging in as johndoe@example.com, this field is `example.com`
LDAP Host URLs	You can specify multiple LDAP Host URLs separated by commas.
Username	User with the necessary rights to log in using the settings provided.
Password	Password of the user.

You can configure a group and provide a role to the members of that group. To enable this functionality, select **Group based access control**.

a Under **Base DN**, type the Base DN, the point from which the server starts searching for users.

b Under **Group DN**, add groups .

c For each group, select the role of the user as member or administrator from the drop-down menu. If you select the administrator role for a particular group, then all the members of that group have the administrator privilege. Similarly, if you select the member role for a particular group, then all the members of that group have the member privilege. If this option is not selected, then the group setting is used to assign the privileges. But other valid LDAP users who do not belong to the groups that you have added can log in to the product.

d Click **Add more** to add groups in the inclusion list.

To allow access to the users only from the LDAP groups (direct or inherited membership) that you have added, select the **Restrict access to members of the above groups only** check box. To

4 Click **Submit** to configure LDAP.

After the LDAP configuration is successful, a new drop-down menu is available on the login screen where users can select whether they want to log in locally or using their LDAP credentials.

The LDAP credentials are not saved anywhere.

Considerations about Groups and Inheritance

- For the groups that you have added under Group DN, their child groups also can log in using the LDAP credentials.
- Inheritance is not considered for the role assignment. For example, if a user has to be an administrator, the direct group to which the user belongs should be assigned the administrator role. The user belonging to the child group will not have the administrator role.

- Suppose that you have assigned the administrator role to a group and you want to exclude a particular user in that group from the administrator role, perform the following steps:

  a On the **Settings** page, click **User Management**.

  b Under the **LDAP Users** tab, you can see the assigned role of that particular user and also that the role has been inherited from the group.

  c Click the edit icon. Under **Role**, select **Member** from the drop-down menu for that user. In this way, you assign a role directly to the user.

  d Click **Save Changes**.

  e Enter your password to confirm. Click **Authorize**.

- Suppose that you want a user to inherit the role from the group to which the user belongs, then perform the following steps:

  a On the **Settings** page, click **User Management**.

  b Under the LDAP Users tab, you can see the assigned role of that particular user and also that the role has been directly assigned to the user.

  c Click the delete icon to delete that LDAP user.

  d When that particular user logs in, the user inherits the role from the parent group by default.

- While a user is logged in, if someone changes the role of the group to which the user belongs, the new role comes into effect only after the user logs out.

- Suppose that there are some LDAP users who are logged in before an upgrade. After an upgrade, the LDAP users have direct roles and do not inherit from the group.

- Suppose that a user belongs to multiple groups. For example, a user belongs to Group A, Group, B, and Group C. If Group A is assigned the administrator role, and Group B and Group C are assigned the member role, then the user inherits the administrator role.

Configuring Mail Server

To configure mail server:

1 On the top-right corner of Home page, click the **Profile** icon, and then click **Settings**.

2 Click **Mail Server**.

3 Select the SMTP server check box.

4 Enter appropriate values in the boxes.

Table 3-14.

Field	Description
Sender Email	Sender's email address.
SMTP Hostname/IP Address	Hostname or IP address of the SMTP server.

Table 3-14. (Continued)

Field	Description
Encryption	The following encryption options are available: None, TLS, and SSL.
SMTP Port Number	Port number of the SMTP server (default 25).

Optionally, for additional security, select the Authentication checkbox, and enter the user name and password.

Note To verify whether the notification mail is correctly set up, click **Send test Email**.

5 Click **Submit** to complete the configuration.

Support for Simple Network Management Protocol (SNMP)

The product supports the following two versions of SNMP:

1 v2c

2 v3

Configuring SNMP service

1 On the top-right hand corner in the Home page, click the Profile icon, and then click **Settings**.

2 On the **Settings** page, click **SNMP**, and then click **Configure SNMP Service**.

3 On the **Configure SNMP Service** page, in the Version box, select SNMPv2c or SNMPv3 protocol.

Note SNMPv2c protocol does not require authentication. SNMPv3 protocol supports authentication.

4 In the Destination IP Address/FQDN box, enter the IP address of the SNMP agent, or enter the Fully Qualified Domain Name (FQDN).

5 In the Destination Port box, enter **162**.

6 If you select the SNMPv2c protocol, in the Community String box, enter **Public**. If you select the SNMPv3 protocol, in the Username box, enter the name of the user you created in the SNMP agent.

For SNMPv3, additionally,

- Select the **Use Authentication** checkbox.
- Select an authentication protocol, and then enter the password you had set for the particular user in the SNMP agent. Optionally, in the Privacy Protocol and Privacy Phrase boxes, select a privacy protocol and a privacy phrase respectively.

To verify whether the configuration is correctly done, click **Test SNMP trap**, and then find whether the trap has been sent to the SNMP agent.

7 Click **Submit**.

User Session Preferences

You can set the preference for the user session timeout.

Procedure

1 On the **Settings** page, click **Preferences**.

Note The **Preferences** tab is visible only to the `admin user`.

2 Click the edit icon to change your preference for the user session timeout.

Note The default value for the session timeout is `15` minutes.

3 Drag the slider bar to set the timeout value for the session. The value ranges from `15` minutes to `24` hours.

4 You can also view the details on who modified the timeout value and when in the **Last Modified** field.

5 Click **Submit**. The Success message appears to confirm that the updated session duration will be effective from the next login.

Note The new value for the user session timeout will come into effect only after you log out and log in again.

Capacity

4

vRealize Network Insight provides the approximate capacity and load information of a collector node and a platform. This limits-based information helps you to prevent the performance and experience issues later.

Understanding Capacity

There are two kinds of capacity:

- VM capacity: It is defined as the number of discovered VMs that a node or a setup can handle.
- Flow capacity: It is defined as the number of flows that a node or a setup can handle.

The capacity is defined as follows:

- Single platform with one or more proxy nodes: The capacity of a proxy node or the platform is the number of discovered VMs that it can handle without the degradation of performance.
- Cluster setup: The capacity of the platform in a cluster setup is the aggregation of all the capacities of all the platform nodes while the capacity of proxy nodes is considered at the level of an individual node.

Accessing the Capacity Information

You can view **VM Capacity** and **Flow Capacity** on the **Install and Support** page.

For every collector node listed under Collector (Proxy) VMs, only the VM capacity information is provided.

Note When the number of discovered VMs from the data sources across the deployment exceed the capacity of either the system or the collector or both, you will not be allowed to trigger the upgrade.

To view the discovered VMs for a data source:

1 In the **Accounts and Data Sources** page, you can see the number of VMs that have been discovered for a particular data source which is already added and currently active. This column will have a value only if the data source is vCenter or AWS source.

Note The discovered VM count includes placeholder and template VMs. So it can be different from the count of VMs in the product.

This chapter includes the following topics:

- Planning to Scale up the Platform Cluster
- Planning to Scale up the Proxy Server

Planning to Scale up the Platform Cluster

Three or more LARGE platform bricks can be connected together to form a platform cluster.

Note Ensure that you take a backup of the Platform1 node before you create clusters. Refer to VMware best practices to take the backup of virtual machines (like VMware VDP using VADP). Restore the Platform1 node from backup if there is an unrecoverable error while creating the cluster. It is recommended that you use cleanly deployed platform nodes while creating clusters. Redeploy the new platform nodes (p2-pn) before restarting cluster creation process if there is an unrecoverable error.

To decide the required number of platform bricks:

```
Number of bricks needed = Round off to next Integer ((Total number of managed VMs) /
(Capacity of LARGE Platform brick in table above))
```

Scaling up Scenarios for the Platform Cluster

- Scenario 1

  a Assume that on January 1st (today), the datacenter has 2000 VMs (with flows) across many vCenters.

  b Assume that in March, the number of VMs grows to 3100.

  c Assume that in June, the number of VMs grows to 6100 which could be because of the additions of few more vCenters or the expansion of the existing vCenters.

  d Assume that in December, the number of VMs grows to 18100 (with flows).

  The deployment model for this scenario is as follows:

  a On January 1, deploy a single platform node with the MEDIUM brick size.

  b In March, scale up the platform node to the LARGE brick size.

  c In June, scale out the platform, convert to a three node platform cluster by adding new Platform nodes to the existing Platform.

  d In December, the user needs a four node platform cluster.

- Scenario 2

  a Assume that on January 1st (today), the datacenter has 7000 VMs (with flows) across many vCenters.

  b Assume that in June, the number of VMs grows to 15000 (with flows).

  c Assume that in December, the number of VMs grows to 24000 (with flows).

The deployment model for this scenario is as follows:

a On January 1, deploy a three node platform cluster.

b In June or later, as the environment size gets closer to exceeding 18000, the user needs a four node platform cluster.

c In December, as the environment size gets closer to exceeding 24000, the user needs a five node platform cluster.

Planning to Scale up the Proxy Server

The scaling up of the proxy node is independent of the platform nodes in the cluster. Typically, the users install one or more proxy VMs per site. Within a site, the number of proxy VMs needed is a simple function of total number of VMs for which it has to collect data. Refer to the capacity of proxy VMs in the brick size table in the System Requirements section.

You can add a data source (maybe a vCenter or a switch) to exactly one proxy VM.

Scaling up Scenarios for the Proxy Server

- Scenario1: Suppose there are 2000 VMs in a vCenter.

  Install one medium proxy VM. Assign the vCenter to this proxy using the product UI.

- Scenario 2: 1000 VMs in vCenter1 and 2000 VMs in vCenter2 (all of them are in one data center)

  Install one medium Proxy VM. Assign both vCenters to this proxy using the product UI.

- Scenario 3: 1000 VMs in vCenter 1 and 2000 VMs in vCenter2 (all of them are in the same data center)

  Install one medium Proxy VM. Assign both vCenters to this proxy using the product UI.

- Scenario 4: 1000 VMs in vCenter1 (data center1) and 2000 VMs in vCenter2 (data center2)

  Install one medium Proxy VM in each data center. Assign vCenter1 to proxy VM in same data center using Product UI. Assign vCenter2 to Proxy VM in its data center using the product UI.

- Scenario 5: 9,000 VMs in vCenter1 without flows (data center1)

  Install one large proxy brick. Assign this vCenter to this proxy using the product UI.

- Scenario 6: 11,000 VMs in vCenter1 with flows (data center1)

  This scenario is not supported. Maximum number of VMs that can be managed by one proxy VM is 10,000 without flows OR 6,000 with flows. And one vCenter can be added to only one proxy at a time.

- Scenario 7: vCenter1 with 2000 VMs in January, vCenter2 with 5000 VMs in June

  Install one medium Proxy VM in January and assign vCenter1 to it. Install the second large proxy VM in June and assign vCenter2 to it.

Proxy VMs with a Platform Cluster

The number of proxy VMs does not depend on the number of VMs in a platform cluster. All proxy VMs communicate only to the first platform VM (`platform1` in the following example) in a platform cluster. A few example deployment models that are supported are as follows:

- Case 1: A proxy VM connects to a platform cluster.

  The proxy connects to `platform1`.

- Case 2: Many Proxy VMs connect to a platform cluster

  All the proxies are connected to `platform1`. And then `platform1` VM load balances both proxy requests and the data processing to other platform VMs in this cluster internally automatically.

- Case 3: A proxy VM connects to the single platform node deployment
- Case 4: Many proxy VMs connect to one platform node deployment

Clusters

5

This chapter includes the following topics:

- Creating Clusters
- Expanding Clusters

Creating Clusters

You can create clusters from the **Install and Support** page.

Prerequisites

At least two additional platforms are required. The additional platform VMs should be deployed and powered on.

To create cluster

1. Click **Create Cluster** for **Platform VMs**.
2. On the **Create Cluster** page, enter the following information:
   - **IP Address**: Enter the IP address of the new platform that you want to add.
   - **Password**: Enter the support user password of the platform VM. If you have not changed the password yet, then refer the *Default Login Credentials* section in *vRealize Network Insight Installation Guide* for the password.
3. To keep adding more platforms, click **Add more** and enter the IP address and the support user password.
4. Click **Submit**. Click **Yes**.

5 After creating a cluster, the user needs to log in to the product again.

Note

- The **create cluster** option is enabled only when the platform is of large brick size. All platforms should be of large brick to create cluster.
- To receive telemetry data, ensure that you enable telemetry on all the platform nodes.
- To expand clusters, refer the *Expanding a Cluster* section in the *vRealize Network Insight Installation Guide*.

Expanding Clusters

Once the cluster is created, you can expand the cluster by adding more platform nodes to it.

To expand clusters:

Procedure

1. On the **Install and Support** page, click **Expand Cluster** for **Platform VMs**.
2. The IP addresses of the VMs that are part of the cluster already are listed on the Expand Cluster page. To add one or more nodes to the existing cluster, provide the IP address of the node and the support user password.

   Note

   - Currently, vRealize Network Insight supports 10 nodes in an existing cluster. Once the limit is reached, the **Add more** button is disabled.
   - Ensure that all the new nodes are non-provisioned and are reachable through SSH.
   - Ensure that you have taken a backup of the existing platform VMs before you go ahead with the cluster expansion.
3. Click **Submit**.

   The step-by-step progress is displayed.
4. Once the cluster expansion link is completed, a message indicating success is displayed.

   While the cluster expansion is in progress, the application cannot be used for any other operation.

Entity Pages

6

The entity pages provide a comprehensive outlook of the entities that are present in your data center. This information can range from detailed topologies to show relationships with other entities of your data center to detailed metrics about a particular entity.

Each entity page is a collection of pins and each pin shows specific information related to the entity. The information provided is both real time and historical, and an exhaustive list of metrics and properties for the entity.

If you want to visit the Help content, then click **Help** on the top-right corner of the entity page.

Timeline

The timeline provides you the following information:

- The state of the data center at a particular time in the past
- A bird’s eye view of events that were detected across a selected time range

Select the time range of the timeline that you want to view.

To view a particular timeline, select the time range by using the **Time Range** option.

Property Pins

The property pins display important attributes in a two-column layout. Some property pins might also display only a singular attribute value. An example of the property pin is the **VM Properties** pin. The **VM Property** pin displays the properties of a VM, such as operating system, IP address, default gateway, logical switches, CPU, memory, power state, and so on.

This chapter includes the following topics:

- Virtual Machine Topology
- Hosts
- VXLAN
- VLAN
- NSX Manager

Virtual Machine Topology

The virtual machine topology provides a comprehensive view of a singular virtual machine in relation to the rest of your data center.

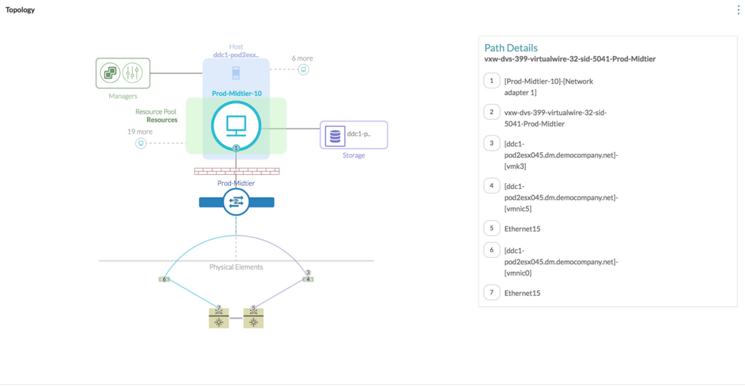

Path Details
vxw-dvs-399-virtualwire-32-sid-5041-Prod-Midtier

1 [Prod-Midtier-10]-[Network adapter 1]

2 vxw-dvs-399-virtualwire-32-sid-5041-Prod-Midtier

3 [ddc1-pod2esx045.dm.democompany.net]-[vmk3]

4 [ddc1-pod2esx045.dm.democompany.net]-[vmnic5]

5 Ethernet15

6 [ddc1-pod2esx045.dm.democompany.net]-[vmnic0]

7 Ethernet15

Hosts

The host topology shows how VMs of a particular host are connected to the virtual and physical components of your data center and also how the host itself is connected with your data center.

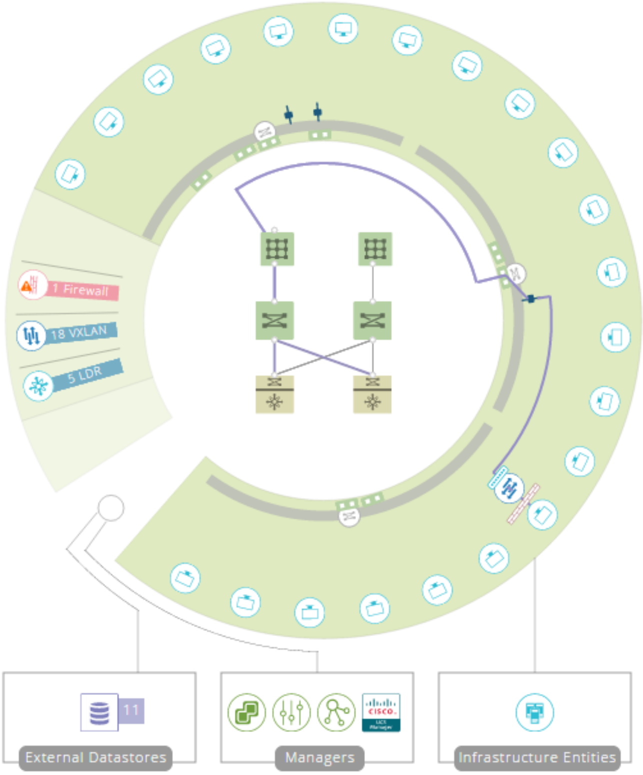

VXLAN

Virtual eXtensible Local Area Network (VXLAN) overlay networking technology is an industry standard that is developed by VMware jointly with the major networking vendors.

The VXLAN topology is an innovative visualization that gives you an overview of the selected VXLAN. The following diagram elucidates the various components that make up the visualization:

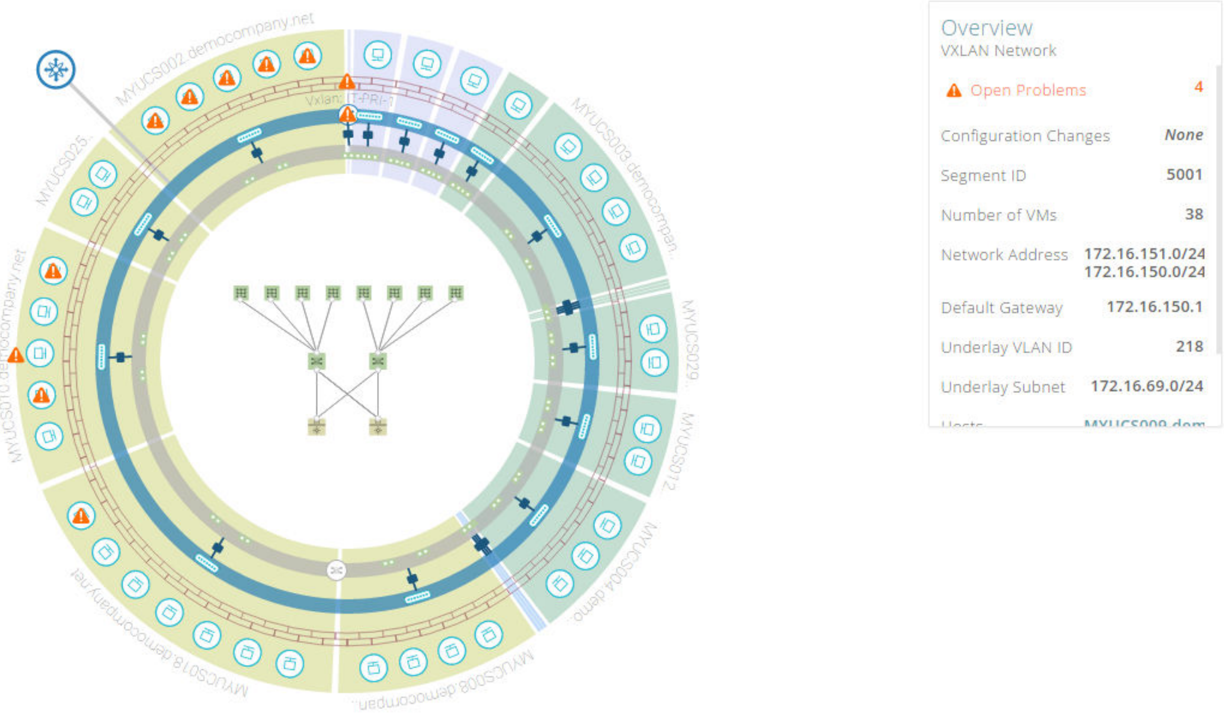

Overview
VXLAN Network

Open Problems	4
Configuration Changes	None
Segment ID	5001
Number of VMs	38
Network Address	172.16.151.0/24 172.16.150.0/24
Default Gateway	172.16.150.1
Underlay VLAN ID	218
Underlay Subnet	172.16.69.0/24
Hosts	MYUCS009.dom

Note Both virtual and physical components can be visualized in this manner.

VLAN

Virtual LANs (VLANs) enable a single physical LAN segment to be further segmented so that groups of ports are isolated from one another as if they were on physically different segments.

The VLAN topology is constructed in a similar manner as the VXLAN topology.

NSX Manager

The NSX Manager topology shows the components that are associated with the NSX Manager.

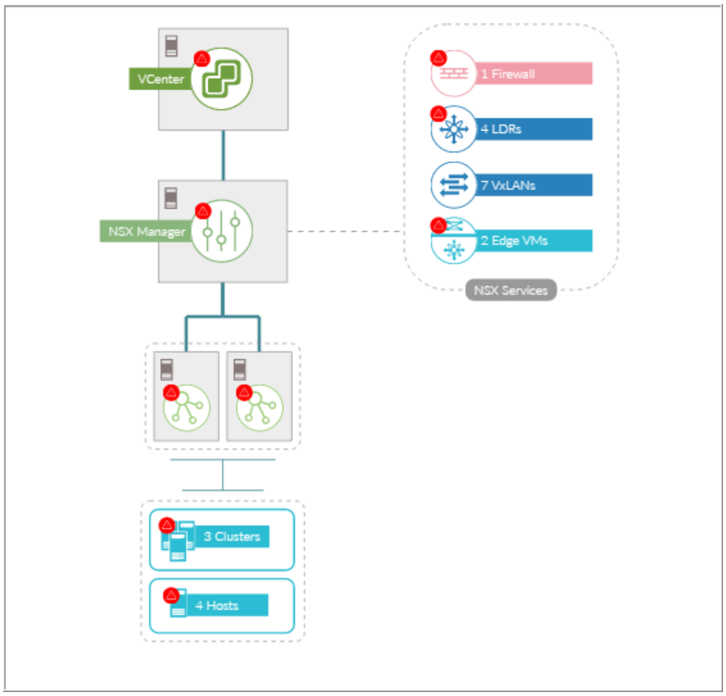

Edge Data Collection

Whenever you add an NSX data source, you can enable the automatic edge data collection. In the previous releases, the edge data collection was done either by NSX Central CLI or Edge-SSH session. From the current release onwards, the edge data collection is done by NSX Central CLI. So no edge data providers are created under NSX Manager.

Note Validation of NSX User Privileges

While adding the NSX data source and enabling the edge population, the NSX user privileges are validated.

- In NSX 6.4 and the further releases, an enterprise admin user can run the NSX Central CLI commands. The user credentials that need to be specified while adding NSX Manager as a data source must be of an enterprise admin or super user.
- In NSX 6.2 and the further releases previous to NSX 6.4, the user should be a super user to enable the edge data population. The user credentials that need to be specified while adding NSX Manager as a data source must be of an enterprise admin or super user.
- Suppose that a user has the enterprise admin privilege in NSX 6.3 and is working on the current release of vRealize Network Insight, an `Insufficient Privileges` error comes up on the **Accounts and Data Sources** page for **VMware NSX Manager**. The error is shown because the user has to be a super user to run the NSX Central CLI commands in NSX 6.3.

Table 6-1.

NSX Version	User
NSX 6.4 and the further releases	■ The user credentials that need to be provided while adding NSX Manager as a data source must be of an enterprise admin or super user. ■ An enterprise admin user can run the NSX Central CLI commands required by vRealize Network Insight.
NSX 6.2 and the further releases before NSX 6.4	■ The user should be a super user to enable the edge data population. ■ The user credentials that need to be provided while adding NSX Manager as a data source must be of an enterprise admin or super user.

System Dashboards

7

The vRealize Network Insight dashboard provides an insight into properties and changes of the system entity over a period of time. It is a snapshot of all the information that is required to analyse your system. There are four types of vRealize Network Insight dashboards:

This chapter includes the following topics:

- NI System Dashboard
- NI Platform Dashboard
- NI Collector Dashboard
- Data Source Dashboard
- Viewing the PCI-Compliance Dashboard
- Viewing the Flow Analytics Dashboard
- Analytics - Outlier Detection

NI System Dashboard

The NI System dashboard provides a snapshot of all the information related to the system. To access the NI System dashboard:

- On the **Install and Support** page, click **View Details** next to **Overview**. The NI System dashboard appears.
- Provide `NI–System` as the search query to view the NI System dashboard.

The NI System dashboard is divided into three sections:

- Overview: This section consists of information on the key properties, the data sources, the problems that are open, and all the changes and the problems related to the system. View the details of each data source by clicking it.

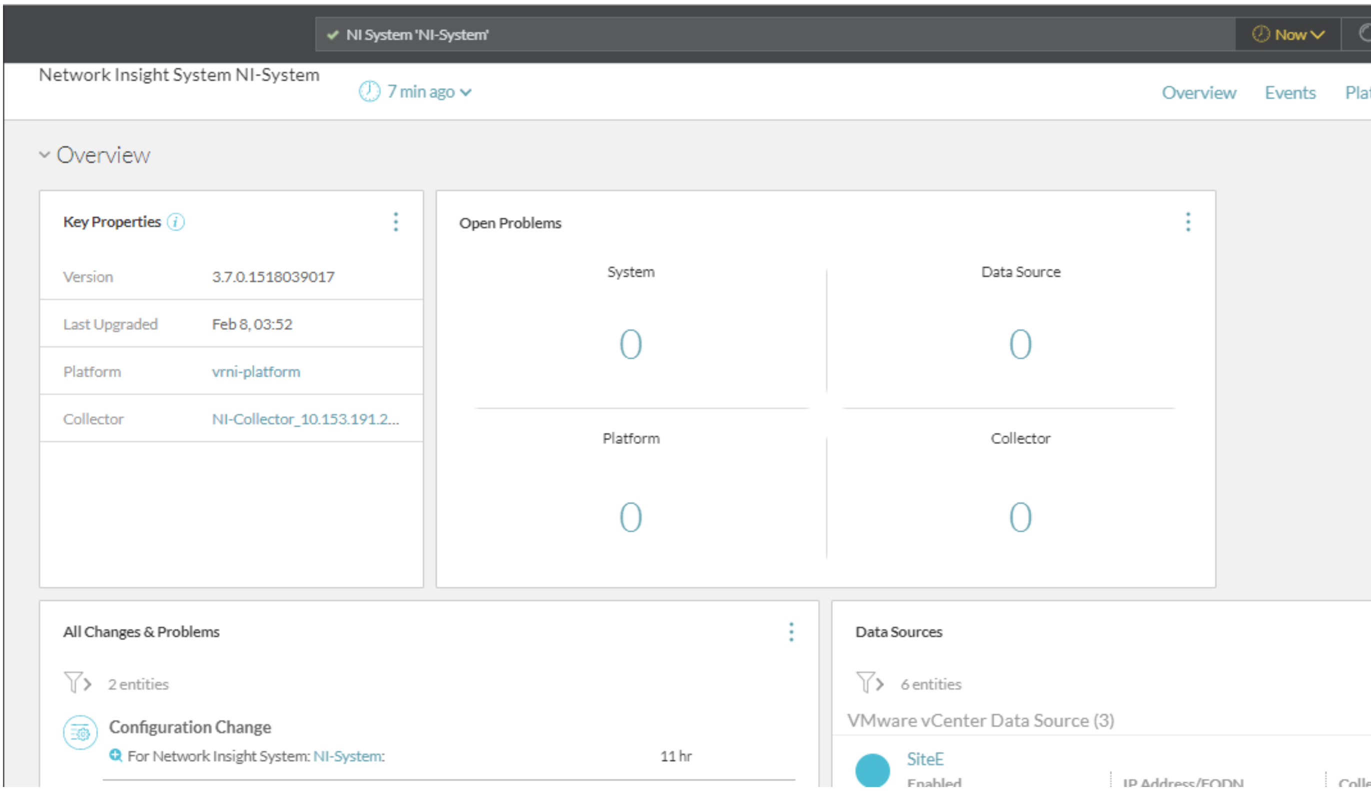

✓ NI System 'NI-System'

Now

Network Insight System NI-System

7 min ago

Overview Events Plat

Overview

Key Properties

Version	3.7.0.1518039017
Last Upgraded	Feb 8, 03:52
Platform	vrni-platform
Collector	NI-Collector_10.153.191.2...

Open Problems

System	Data Source
0	0
Platform	**Collector**
0	0

All Changes & Problems

2 entities

Configuration Change
For Network Insight System: NI-System: 11 hr

Data Sources

6 entities

VMware vCenter Data Source (3)

SiteE
Enabled | IP Address/FQDN | Colle

- Events: This section lists all the problems and changes in the system, data sources, platforms, and the collectors.

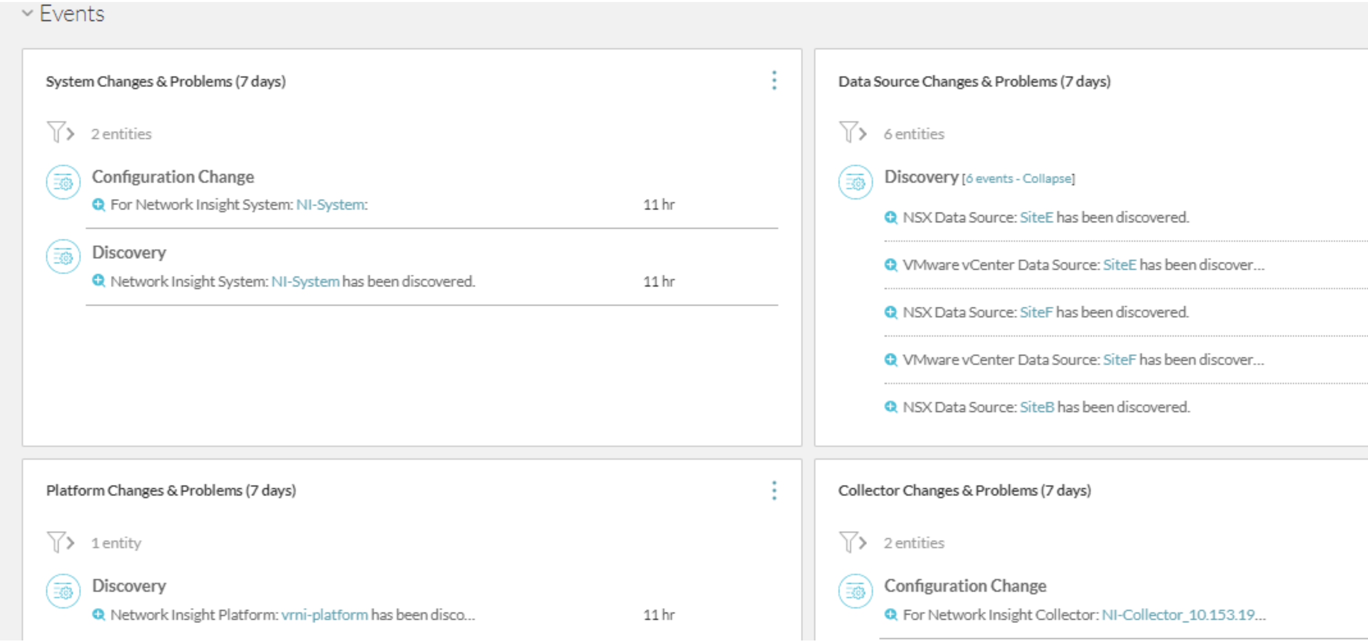

Events

System Changes & Problems (7 days)

2 entities

Configuration Change
For Network Insight System: NI-System: 11 hr

Discovery
Network Insight System: NI-System has been discovered. 11 hr

Data Source Changes & Problems (7 days)

6 entities

Discovery [6 events - Collapse]
NSX Data Source: SiteE has been discovered.
VMware vCenter Data Source: SiteE has been discover...
NSX Data Source: SiteF has been discovered.
VMware vCenter Data Source: SiteF has been discover...
NSX Data Source: SiteB has been discovered.

Platform Changes & Problems (7 days)

1 entity

Discovery
Network Insight Platform: vrni-platform has been disco... 11 hr

Collector Changes & Problems (7 days)

2 entities

Configuration Change
For Network Insight Collector: NI-Collector_10.153.19...

- Platforms and Collectors: This section lists all the platforms and the collectors associated with the system. To view more details about any platform or collector, click it.

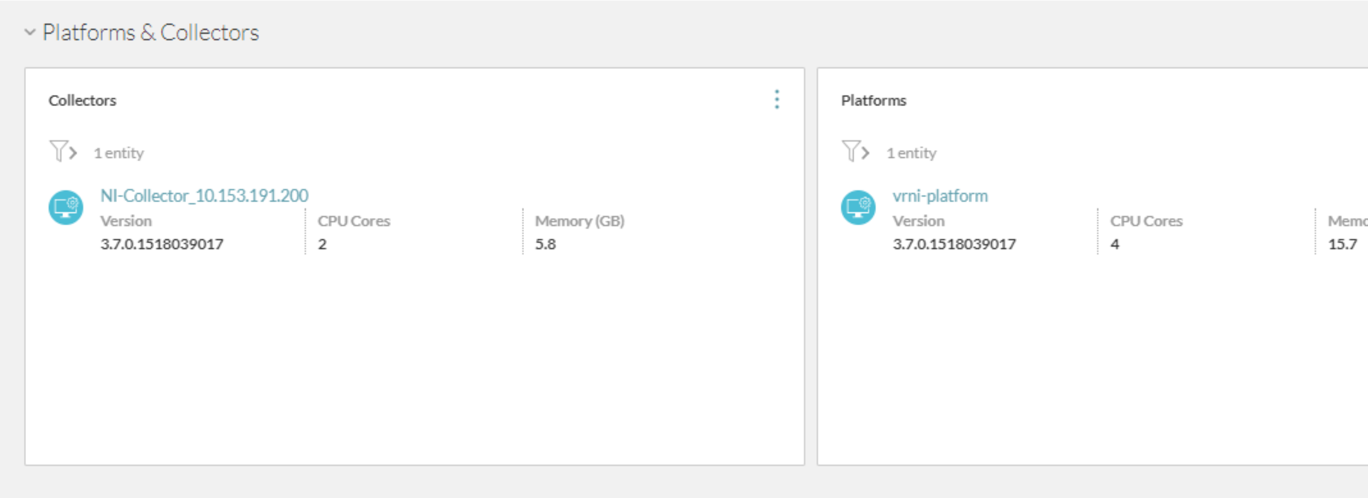

Platforms & Collectors

Collectors

1 entity

NI-Collector_10.153.191.200

Version	CPU Cores	Memory (GB)
3.7.0.1518039017	2	5.8

Platforms

1 entity

vrni-platform

Version	CPU Cores	Memo
3.7.0.1518039017	4	15.7

NI Platform Dashboard

The NI Platform dashboard provides a snapshot of the properties, changes, and problems of a particular platform node.

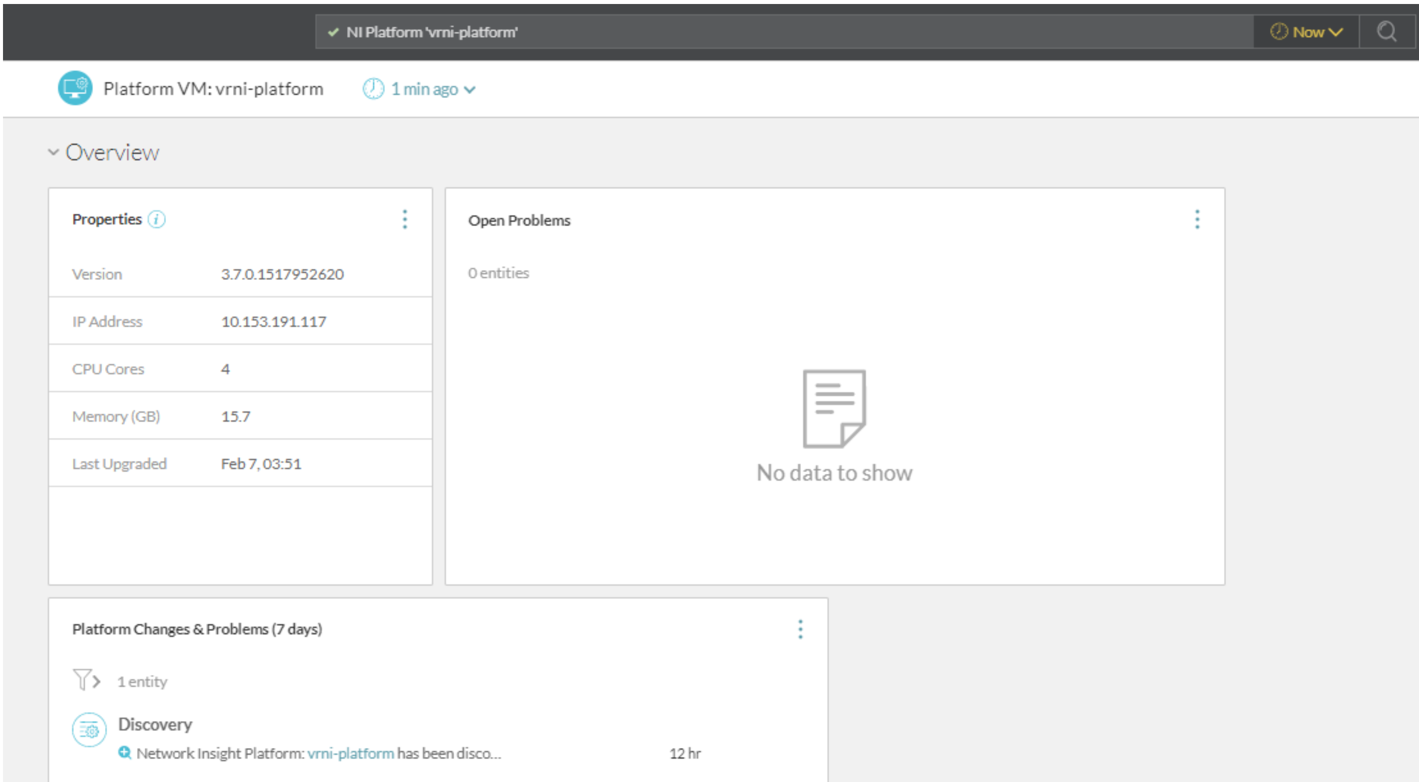

Platform VM: vrni-platform

1 min ago

Overview

Properties

Version	3.7.0.1517952620
IP Address	10.153.191.117
CPU Cores	4
Memory (GB)	15.7
Last Upgraded	Feb 7, 03:51

Open Problems

0 entities

No data to show

Platform Changes & Problems (7 days)

1 entity

Discovery

Network Insight Platform: vrni-platform has been disco... 12 hr

- Properties: This pin provides details of the node such as name, IP address, CPU cores, memory, the last upgraded time, and the version.

- Open Problems: This pin indicates the problems that are associated with the platforms and are open.
- Platform Changes and Problems: This pin highlights the changes and problems that occurred in the platform in the last seven days.
- Metrics: Click Platform Metrics to get the graphical representation of the metrics such as CPU Usage, Memory Usage, and Data Disk Usage.

NI Collector Dashboard

The NI Collector dashboard provides a snapshot of the properties, changes, and problems of a particular collector node.

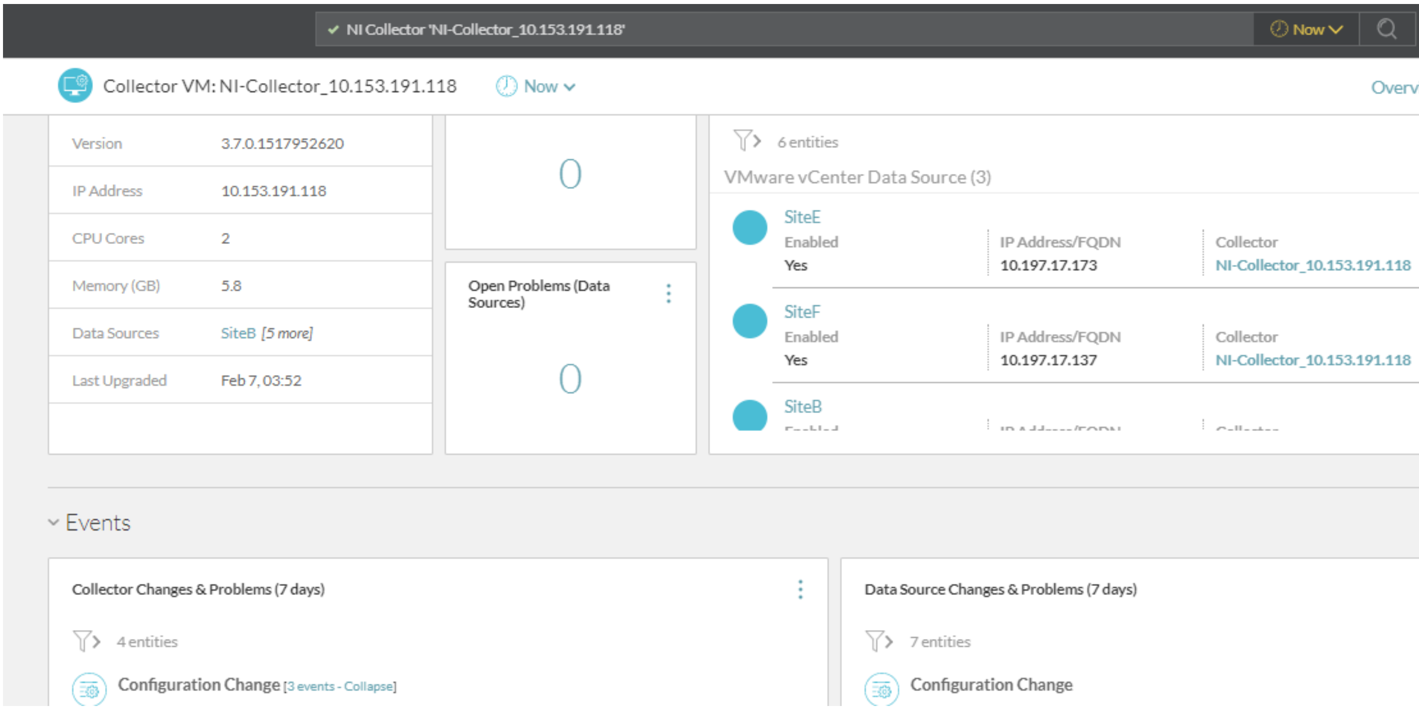

✔ NI Collector 'NI-Collector_10.153.191.118'

Now

Collector VM: NI-Collector_10.153.191.118 Now

Overv

Version	3.7.0.1517952620
IP Address	10.153.191.118
CPU Cores	2
Memory (GB)	5.8
Data Sources	SiteB *[5 more]*
Last Upgraded	Feb 7, 03:52

0

Open Problems (Data Sources)

0

6 entities

VMware vCenter Data Source (3)

SiteE Enabled Yes	IP Address/FQDN 10.197.17.173	Collector NI-Collector_10.153.191.118
SiteF Enabled Yes	IP Address/FQDN 10.197.17.137	Collector NI-Collector_10.153.191.118
SiteB Enabled	IP Address/FQDN	Collector

Events

Collector Changes & Problems (7 days)

4 entities

Configuration Change [3 events - Collapse]

Data Source Changes & Problems (7 days)

7 entities

Configuration Change

- Properties: This pin provides details of the node such as name, IP address, CPU cores, memory, the last upgraded time, and the version. You can get the following information from this pin:
  - IP Address and the version of the node
  - Brick Size
  - The time at which it was rebooted last
  - The number of CPU cores
  - The number of data source that were added on it
  - The size of memory, total disk space used, and the total data disk space used

    Note The values mentioned for the storage are all approximate values.
  - The time including the year when it was last upgraded

  - Netmask
  - The NTP server and the DNS server details
  - Search Domain
  - The total number of VMs running on this particular collector
  - The time when the last activity was done
- Open Problems(Collectors): This pin lists all the open problems related to the collector.
- Open Problems (Data Sources): This pin lists all the open problems related to the data sources.
- Data Source Changes and Problems: This widget highlights the changes that occurred in the data source in the last seven days.
- Data Sources and NetFlow Reporters: This widget provides you the details of the data sources and the NetFlow reporters available in the collector. The number of flows are shown for each NetFlow reporter. For data sources, the number of flows and the discovered VMs are shown.
- Metrics: Click Collector Metrics to get the graphical representation of the metrics such as CPU Usage, Memory Usage, and Data Disk Usage

Data Source Dashboard

The data source dashboard provides a snapshot of the properties, changes, and problems of a particular data source.

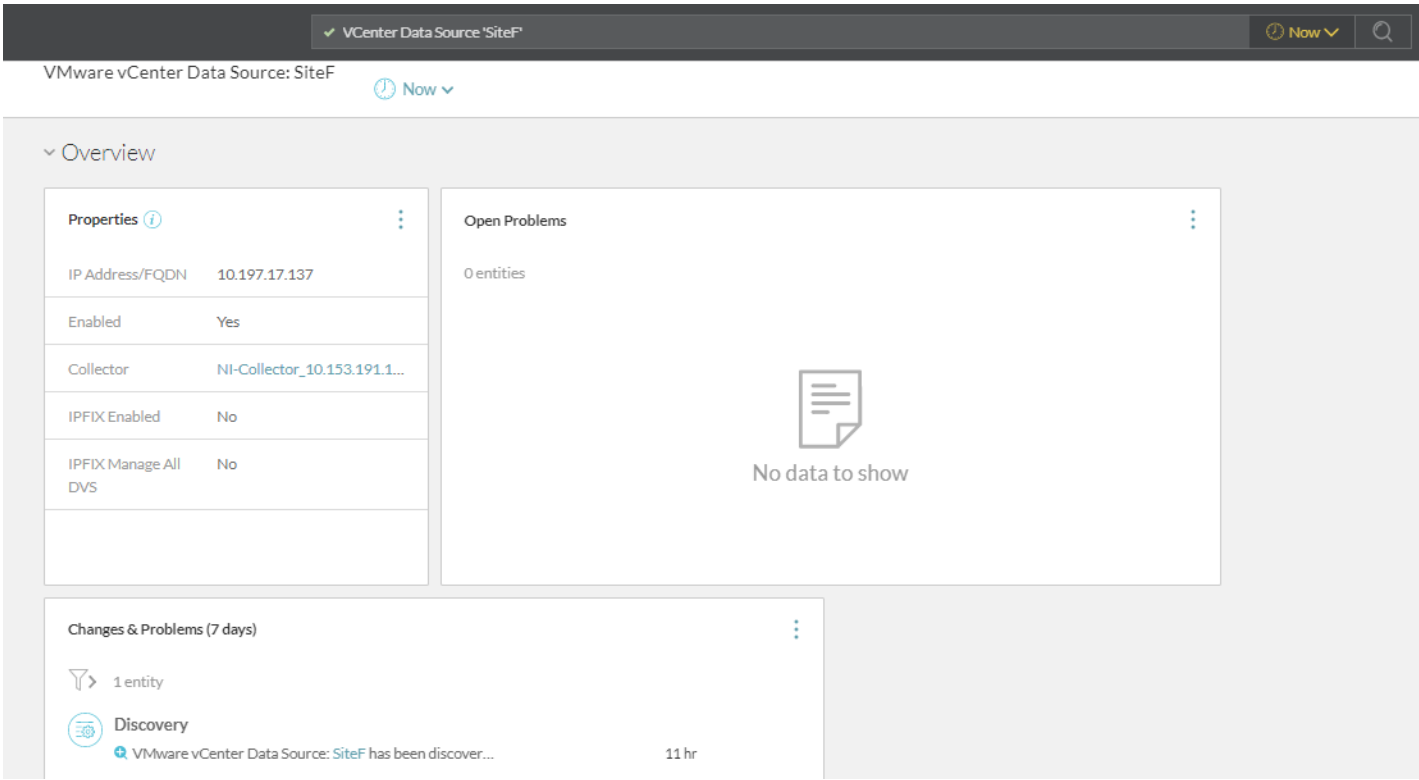

VMware vCenter Data Source: SiteF

Properties

IP Address/FQDN	10.197.17.137
Enabled	Yes
Collector	NI-Collector_10.153.191.1...
IPFIX Enabled	No
IPFIX Manage All DVS	No

Open Problems

0 entities

No data to show

Changes & Problems (7 days)

1 entity

Discovery

VMware vCenter Data Source: SiteF has been discover... 11 hr

- Properties: You can view the following information on this pin:
  - IP Address/FQDN: The IP address or the FQDN of the data source
  - Collector: The name of the collector on which it is added
  - Enabled: Indicates if the data source is enabled or not
  - IPFIX Enabled: Indicates if IPFIX is enabled or not
  - IPFIX Manage All DVS: Indicates if vRealize Network Insight is configured to receive NetFlow from all DVS
  - Discovered VMs: The number of discovered VMs
  - Last collection: The time when the last collection was done
  - Flows: The number of flows
  - Last Flow Received: The time when the last flow was received
- Open Problems: This pin lists all the open problems associated to the data source.
- Changes and Problems: This pin lists all the changes and the problems encountered in a particular data source in the last seven days.

Viewing the PCI-Compliance Dashboard

The PCI-Compliance dashboard is available only for the Enterprise License users.

To access the PCI-Compliance feature

1. In the Homepage, select **Security** > **PCI Compliance**.
2. The PCI Compliance window appears. Select the required scope from the drop-down menu.
3. The PCI-Compliance dashboard appears.

PCI-Compliance Dashboard Features

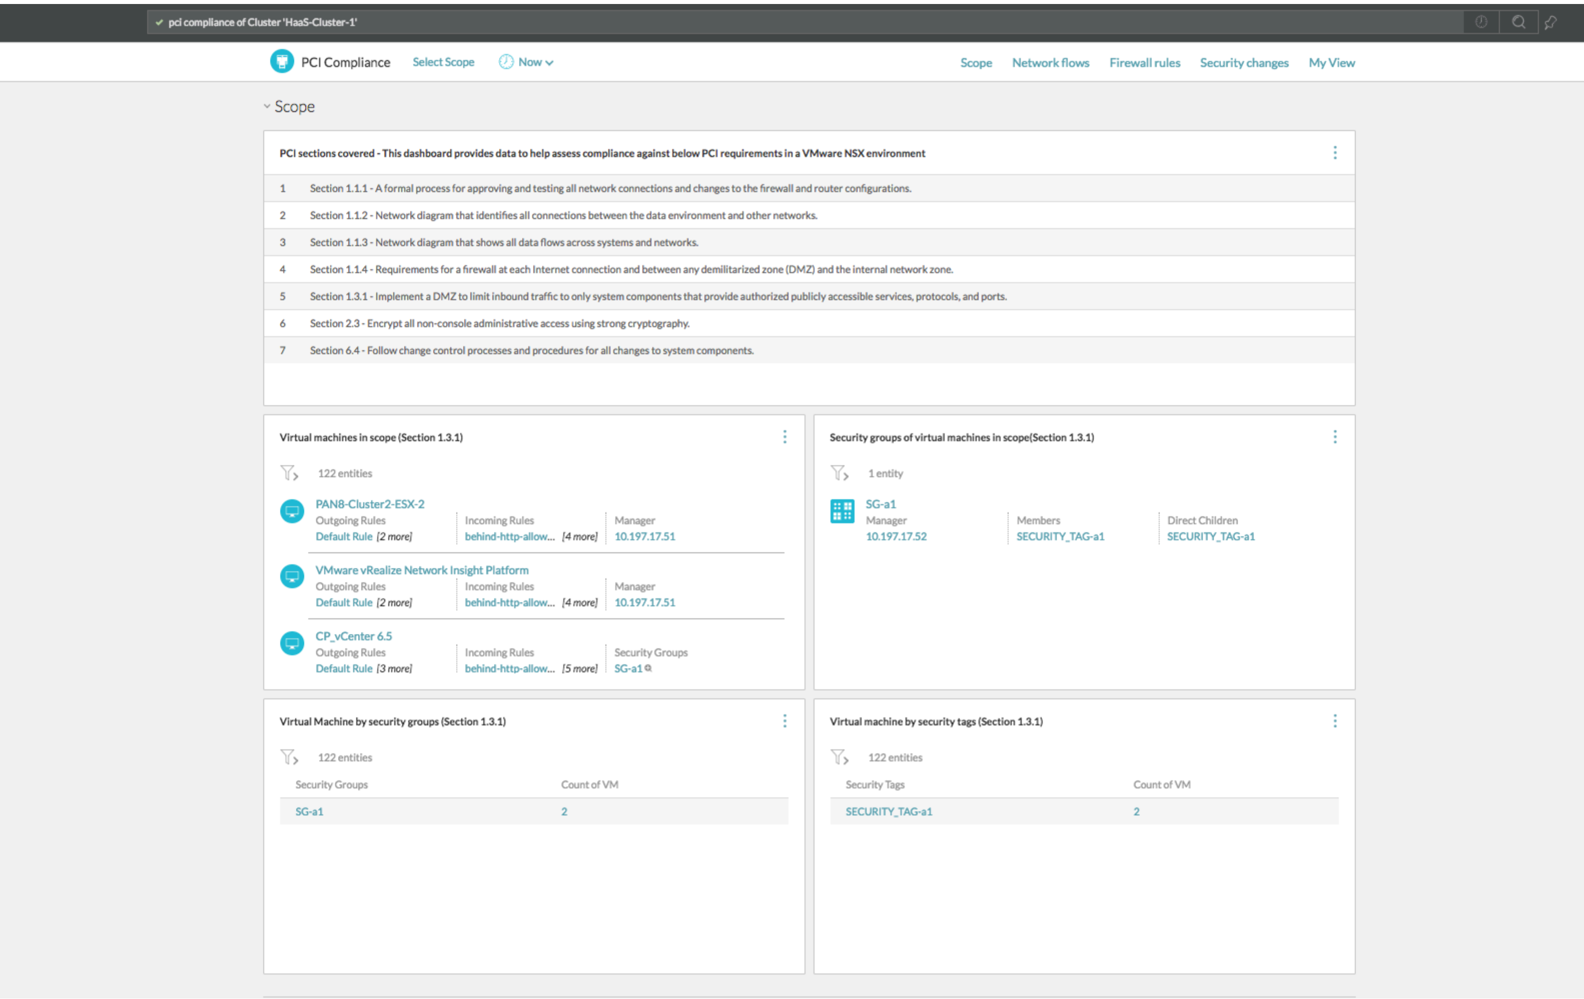

PCI sections covered - This dashboard provides data to help assess compliance against below PCI requirements in a VMware NSX environment

1	Section 1.1.1 - A formal process for approving and testing all network connections and changes to the firewall and router configurations.
2	Section 1.1.2 - Network diagram that identifies all connections between the data environment and other networks.
3	Section 1.1.3 - Network diagram that shows all data flows across systems and networks.
4	Section 1.1.4 - Requirements for a firewall at each Internet connection and between any demilitarized zone (DMZ) and the internal network zone.
5	Section 1.3.1 - Implement a DMZ to limit inbound traffic to only system components that provide authorized publicly accessible services, protocols, and ports.
6	Section 2.3 - Encrypt all non-console administrative access using strong cryptography.
7	Section 6.4 - Follow change control processes and procedures for all changes to system components.

Virtual machines in scope (Section 1.3.1)

122 entities

PAN8-Cluster2-ESX-2

Outgoing Rules	Incoming Rules	Manager
Default Rule [2 more]	behind-http-allow... [4 more]	10.197.17.51

VMware vRealize Network Insight Platform

Outgoing Rules	Incoming Rules	Manager
Default Rule [2 more]	behind-http-allow... [4 more]	10.197.17.51

CP_vCenter 6.5

Outgoing Rules	Incoming Rules	Security Groups
Default Rule [3 more]	behind-http-allow... [5 more]	SG-a1

Security groups of virtual machines in scope(Section 1.3.1)

1 entity

SG-a1

Manager	Members	Direct Children
10.197.17.52	SECURITY_TAG-a1	SECURITY_TAG-a1

Virtual Machine by security groups (Section 1.3.1)

122 entities

Security Groups	Count of VM
SG-a1	2

Virtual machine by security tags (Section 1.3.1)

122 entities

Security Tags	Count of VM
SECURITY_TAG-a1	2

The PCI-Compliance dashboard helps in assessing compliance against the PCI requirements only in the NSX environment. These requirements are mentioned under the first pin in the dashboard. The rest of the pins in the dashboard that provide data for the assessment of these requirements are as follows:

- Network flow diagram: It shows the data flow, firewalls, connections, and other details associated with a network.
- Flows: It lists the flows that you view in the network flow diagram.

- Clear text protocol flows based on the destination port: The traffic that flows on certain ports are in clear text. This pin displays the clear text protocol flows based on a particular destination port.
- Virtual machines in scope: It shows the virtual machines in the scope that you have selected in the query. This pin shows the outgoing rules, incoming rules, and security groups for virtual machines in that scope.
- Security groups of virtual machines: It lists the security groups of the virtual machines.
- Virtual machine count by Security Groups: You can view the list of the virtual machines in a security group by clicking Count in this pin.
- Virtual machine count by Security Tags: You can view the list of virtual machines with security tags by clicking Count in this pin.
- Firewall rules applied on internal traffic : You can view the firewall rules for the traffic between the virtual machines within the selected scope.
- Firewall rules applied on incoming traffic: You can view the firewall rules for the traffic that is coming from a virtual machine outside the scope to the virtual machine within the selected scope.
- Firewall rules applied on outgoing traffic: You can view the firewall rules for the traffic that is going to a virtual machine outside the scope from the virtual machine within the selected scope.
- Security tag membership changes: The changes related to the membership for security tags are shown in this pin.
- Security group membership changes: The changes related to the membership of a security group are shown in this pin.
- Firewall rule changes: The changes related to any firewall rule is listed in this pin.

Note If NSX has nested security groups, then the scope of PCI Compliance should be other than security group.

Export as PDF

vRealize Network Insight enables you to create and export the information on the PCI Compliance dashboard as a PDF report.

Procedure

1 In the PCI Compliance dashboard, click **Export as PDF** on the right top side of the page. The Export to PDF window appears.

2 The Export to PDF window lists all the widgets and their respective properties available on the PCI Compliance dashboard. Select the widgets and the properties that you want to export.

Note

- You have to select at least one property.
- The maximum number of properties that you can select is 20.
- The maximum number of entries in the list view that can be exported is 100.
- Certain widgets do not allow you to select the properties. In such instances, specify only the number of entries.

3 Provide a title for the PDF report.

Note

- The maximum number of characters in the title is 200.
- The maximum number of pages that can be generated in the report are 50.

4 Click **Preview**. You can see the preview of the complete report.

5 Click **Export PDF**.

Viewing the Flow Analytics Dashboard

The Flow Analytics dashboard provides an insight into data centers, devices, and flows. It is a context-based dashboard as it performs analysis based on the entities, flows, and the time range that you select.

To access the Flow Analytics dashboard:

- Search **Flows**.
- Click **Flow Analytics**.
- The Flow Analytics dashboard appears.

The various sections in the Flow Analytics Dashboard are:

- Top Talkers
- What's New
- Outliers

Top Talkers

This section helps you to recognize which entities are talking the most in your environment. You can select different kinds of entities such as Source-Destination pair, VM, Cluster, L2 Network, Subnet. This widget lists the top 10 talkers in the entity category that you select. It helps the customer to plan for network optimization. The metrics that are represented by bars in this widget are as follows:

- By Flow Volume: Indicates the traffic volume.

- By Traffic Rate: Indicates the rate of traffic.
- By Session Count: Indicates the number of sessions.
- By Flow Count: Indicates the number of flows

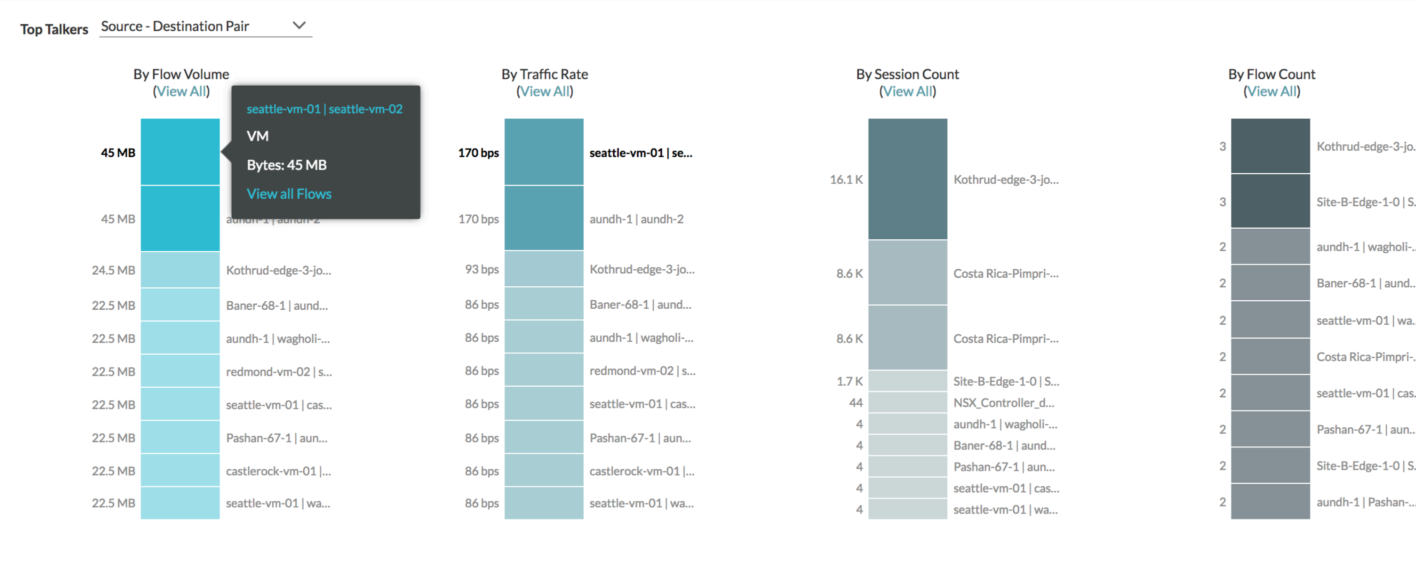
Top Talkers
Source - Destination Pair
By Flow Volume
(View All)
seattle-vm-01 | seattle-vm-02
VM
Bytes: 45 MB
View all Flows
45 MB
45 MB
aundh-1 | aundh-2
24.5 MB Kothrud-edge-3-jo...
22.5 MB Baner-68-1 | aund...
22.5 MB aundh-1 | wagholi-...
22.5 MB redmond-vm-02 | s...
22.5 MB seattle-vm-01 | cas...
22.5 MB Pashan-67-1 | aun...
22.5 MB castlerock-vm-01 |...
22.5 MB seattle-vm-01 | wa...
By Traffic Rate
(View All)
170 bps seattle-vm-01 | se...
170 bps aundh-1 | aundh-2
93 bps Kothrud-edge-3-jo...
86 bps Baner-68-1 | aund...
86 bps aundh-1 | wagholi-...
86 bps redmond-vm-02 | s...
86 bps seattle-vm-01 | cas...
86 bps Pashan-67-1 | aun...
86 bps castlerock-vm-01 |...
86 bps seattle-vm-01 | wa...
By Session Count
(View All)
16.1 K Kothrud-edge-3-jo...
8.6 K Costa Rica-Pimpri-...
8.6 K Costa Rica-Pimpri-...
1.7 K Site-B-Edge-1-0 | S...
44 NSX_Controller_d...
4 aundh-1 | wagholi-...
4 Baner-68-1 | aund...
4 Pashan-67-1 | aun...
4 seattle-vm-01 | cas...
4 seattle-vm-01 | wa...
By Flow Count
(View All)
3 Kothrud-edge-3-jo.
3 Site-B-Edge-1-0 | S.
2 aundh-1 | wagholi-.
2 Baner-68-1 | aund.
2 seattle-vm-01 | wa.
2 Costa Rica-Pimpri-.
2 seattle-vm-01 | cas.
2 Pashan-67-1 | aun.
2 Site-B-Edge-1-0 | S.
2 aundh-1 | Pashan-.

Note

- If a VM appears in one or more metrics, when you point to that VM in a bar, it will also be highlighted in other bars.
- When you click a VM in the metrics bar, the complete list of flows coming to this VM is shown.
- When you select VM as the entity in the Top Talkers list, all the flows related to this VM irrespective of it being the source or destination is shown. If you select Source VM in the list, then only the flows coming from this VM are considered.
- If you are considering the physical flows, you can select either Source IP or Destination IP.
- After you select the Source-Destination pair and point on the metric bar, if you click the link in the tool tip, the corresponding dashboard appears. For example, for a VM in Sourec-Destination pair, the VM-VM path dashboard appears.
- For a flow group view or a flow entity projection or a flows group query, you cannot see the **Flow Analytics** button.

What's New

This section helps you to track what services and entities are discovered in the data center in the selected time range. The sections in this widget are as follows:

- New Virtual Machines Accessing Internet: Lists the new VMs that access Internet.
- New Internet Services Accessed: Lists the new Internet services discovered in the environment.
- New Internal Services Accessed: Lists the new intranet services that are discovered and accessed from the Internet endpoint.

- New Internal/E-W Services Accessed: Lists the services that are exposed and accessed by the machines within a data center
- New Services with Blocked Flows: Lists services that have blocked flows. This section is populated only for IPFIX.
- New Firewall Rule Hits: Lists the new firewall rules that are brought into effect. This section is populated only for IPFIX.

Outliers

This section helps you to track and analyze related data. It consists of the following sections:

- Elephant Flows: This section helps to identify the flows which have small count of sessions and high throughput versus flows which have large count of sessions and small throughput. Typically, the flows with the large session counts and small throughput are also referred as mice flows. The analysis is based on the ratio of bytes to the number of sessions. Each dot in the graph represents multiple flows. When you point to a dot, you can see the list of flows. To view the details of a particular flow, click that flow in the list.
- Custom Analysis: This section allows you to visualize the flow data on 2 dimensions of your choice. It helps in analyzing the data to find the outliers in various ways.

  Note The metrics represented in this section are the approximate values and not the exact values.

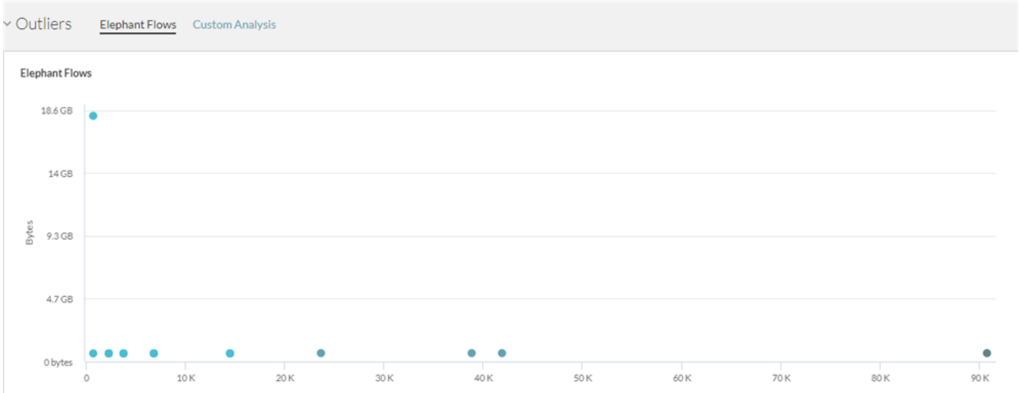

Analytics - Outlier Detection

Network Insight offers outlier detection based on the metrics associated with the flows defined over the VMs and physical IP addresses. These VMs/IPs should have similar traffic patterns so that a classification of a particular VM/IP as an outlier is of value. For example, the VMs, which belong to the same tier of an application, generally perform the same function for the application, such as the VMs of an SQL database serving requests for a web application. For these kind of VMs, the number of requests received, the amount of traffic sent out, the session count, and so on go through a series of similar variations.

Through outlier detection, Network Insight enables you to detect a particular VM which might be experiencing very different traffic pattern compared to other VMs/IPs in the group. For example, if the VM is sending or receiving much higher/lower traffic compared to the rest of the group. It could be because of a wrongly configured load balancer, DDOS attack, and so on. Network Insight classifies such VMs/IPs as outliers. By looking at these outliers, the user easily knows about this unexpected behaviour and takes appropriate actions.

How to Detect the Outlier VMs

Procedure

1 On the sidebar, click **Analytics**. Click **Outlier**.

2 Click **Add** to add a configuration.

3 In the **Analytics/Configure** page, provide the following details for the configuration:

Table 7-1.

Field	Description
Name	Name of the configuration
Scope	Name of the group that defines the VMs and the IPs for which the analysis needs to be done. You can select Application Tier or Security Group as the scope. If you select Application Tier, provide the name of the application and the tier separately. The number of VMs and Physical IPs that are defined for the tier is shown next to the name of the tier. If you select **Security Group**, provide the name of the Security Group. **Note** The current limit for the number of VMs and Physical IPs in a tier is 200. Choose a tier or a security group with VMs and Physical IPs less than this limit. The scope should also contain a minimum of 3 VMs/Physical IPs. You can view the micro segmentation for the selected configuration by clicking **View Micro-Segments**.
Detection Type	Currently, Network Insight enables you to detect the outlier in the system.
Metric	The detection is based on this flow metric. You can select the following options: ■ **Bytes** ■ **Packets** ■ **Sessions** ■ **Traffic Rate**
Traffic Direction	You can select **Outgoing**, **Incoming**, or **Both** as the traffic direction. If you select **Both**, then you can specify Incoming or Outgoing in the preview of the configuration.

Table 7-1. (Continued)

Field	Description
Traffic Type	You can select **Internet**, **East-West**, or All based on the requirement.
Destination Ports	You can either select all ports detected on the flows discovered on the selected scope or manually enter the destination ports of your choice. If you select **All Ports**, the number of the destination ports is shown. If you select **Manually enter ports**, then enter the ports in the autocomplete text box, the analysis would be restricted to only these ports **Note** The current limit for the number of ports is 20.
Sensitivity	It is a measure of the sensitivity of the detection and reporting that you require. The default value is **Medium**.
Preview	This section provides a preview of the particular configuration based on the inputs and parameters that you have provided. Specify the ports and the traffic direction if you have selected Both for Traffic Direction before. You will be able to identify the outlier VM in the graph.

Note

- The outlier is detected by evaluating the data available in last 24 hours.
- You need a continuous flow of IPFIX data to detect the outlier.

4 Click **Submit** to create the analytics configuration.

5 Once the application is created, it is available in the list view of the applications in the Analytics Configurations page. Click that particular application to see a dashboard associated with it.

Network and Security

8

This chapter includes the following topics:

- Network Visibility
- Security

Network Visibility

VM-VM Path

The VM-VM path topology draws a detailed connection that exists between any two virtual machines in your environment.

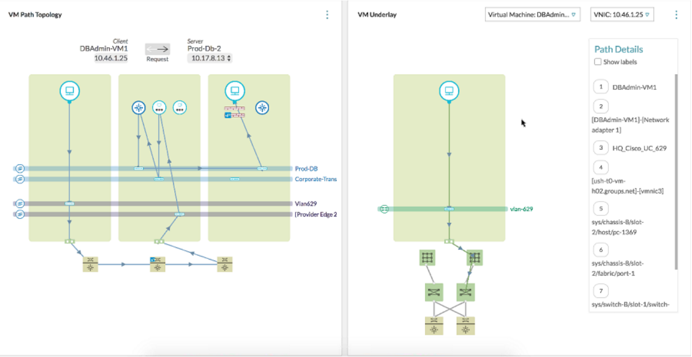

The topology involves both Layer 3 and Layer 2 components. This topology can be viewed using the search query `vm_name_1` to `vm_name_2`. If a path exists, the VM-VM path visualization proceeds to populate all the components that exist between `vm_name_1` to `vm_name_2` and also draws an animated path. If the routers are physical, then they are shown outside the boundary.

In the VM Path topology, if you hover your mouse on any of the routers, edges, or LDRs that are involved in the path, the complete routing or NAT information is shown.

The VM Underlay section that is on the right side of the VM Path topology shows the underlay information of the VMs involved and their connectivity to the top of the rack switches and the ports involved. In the VM underlay section, the components are labeled if you select **Show labels** under **Path Details.** In this section, the drop-down list at the top shows the endpoint VMs and the active VMs at the edges. For each edge VM, the neighboring drop-down list shows the ingress and the egress interface IP addresses. Based on the selection, the underlay path for that particular interface is shown. .

You can also reverse the path direction using the arrows on top of the topology map.

The topology map gives more visibility regarding the ports involved in the VM-VM path. In the **Path Details** section, the name of the actual port channel is shown.

Note

- There is no complete visibility for layer 2 on the physical front. If a packet is traversing from one switch to another, there maybe multiple switches involved. But the topology does not show the switches in the underlay network.

Support for Equal-Cost Multi-Path (ECMP) Route

vRealize Network Insight provides ECMP support in the VM-VM path.

The VM-VM path shows the following information on ECMP:

- The multiple ECMP paths from source to destination
- The routers on which ECMP occurs
- The possible outgoing paths for a given router (VRF)
- The route for the possible path

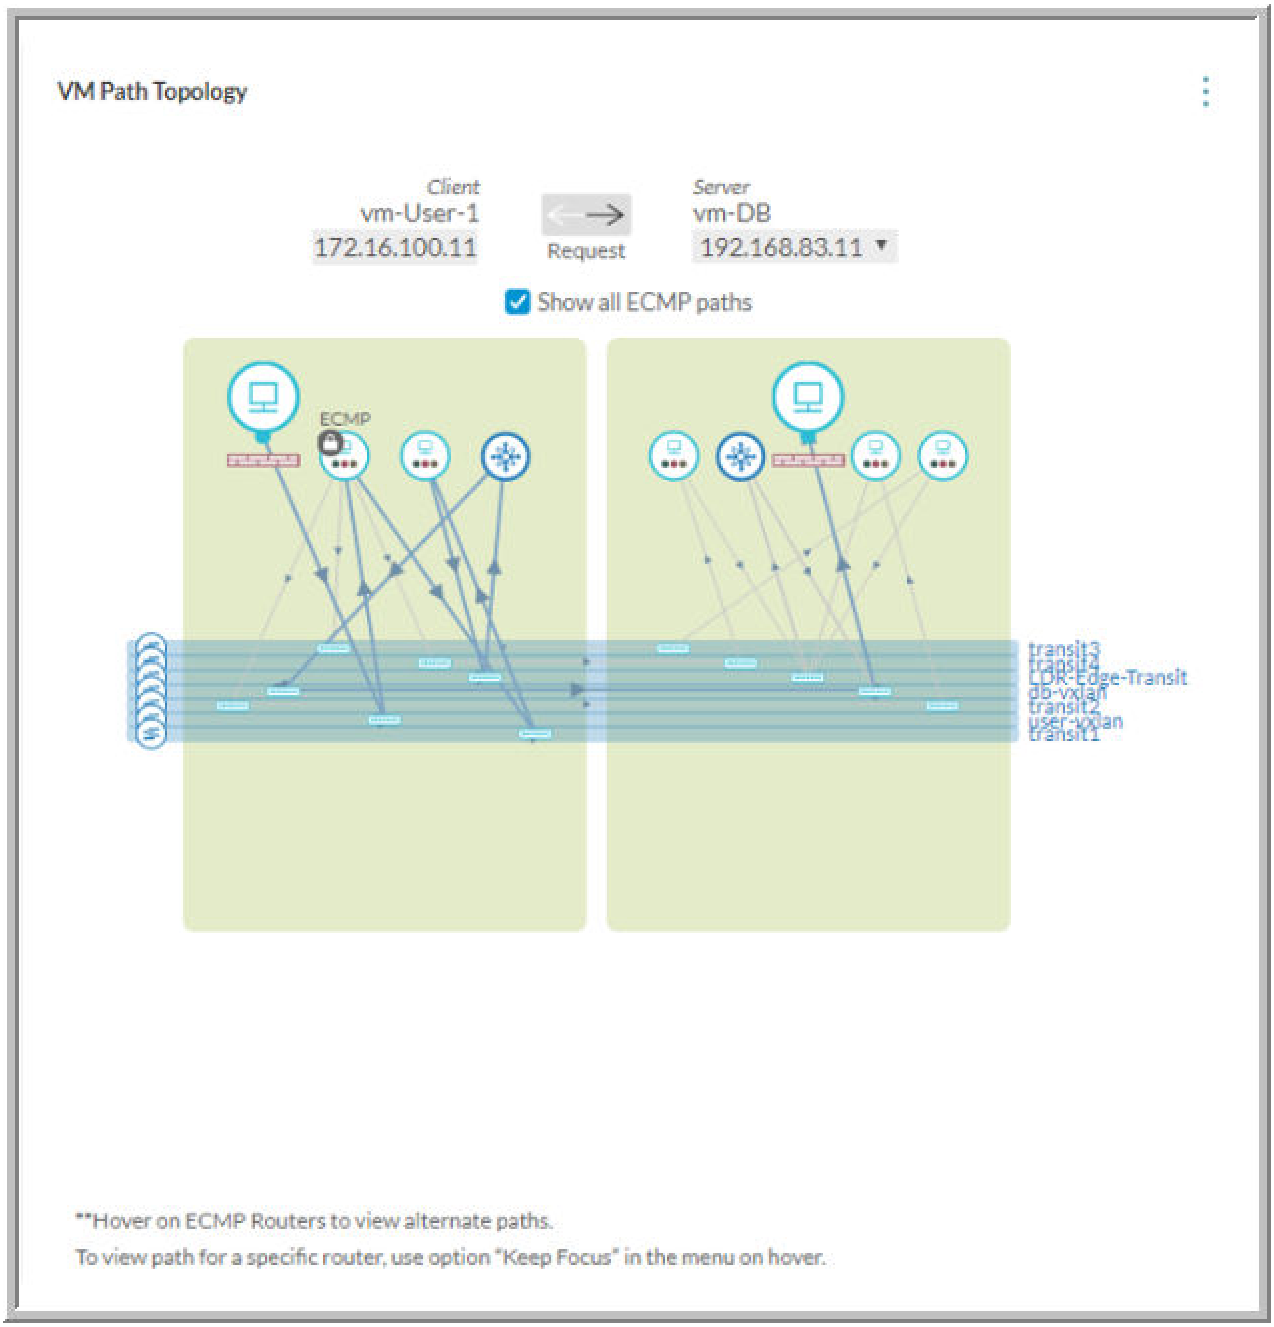

VM Path Topology

Client
vm-User-1
172.16.100.11

Request

Server
vm-DB
192.168.83.11 ▾

☑ Show all ECMP paths

**Hover on ECMP Routers to view alternate paths.
To view path for a specific router, use option "Keep Focus" in the menu on hover.

In the preceding figure, you can see the ECMP-enabled routers. If you point over them, the additional paths are shown. Also, you can create a path by selecting and locking the routers as per your requirement. If you want to view all the ECMP paths between the two VMs, select the **Show all ECMP paths** option in the topology diagram.

If you want to view the path for a particular router, point on the router and click **Keep Focus**. The paths specific to the router is shown.

Support for the L2 Bridges

The L2 or the VLAN bridges create a single broadcast domain from multiple VLANs. In the previous releases, if the VM-VM path involved an L2 bridge between two or more VLANs, the VM-VM path did not work. From this release onwards, vRealize Network Insight supports L2 bridging. Currently, this feature is supported only for the Cisco ASA routers.

Path to Internet

For each virtual machine that is present in your environment, vRealize Network Insight shows you how the VM is connected to the Internet by using an animated path in the **Path to Internet** pin.

The path populates all the components (both virtual and physical) that exist between the virtual machine and the Internet. It draws an animated path that connects each component in a sequence. The path direction can also be reversed by using the arrows situated above the visualization.

Point your mouse pointer to the entity icons to get their addressable names. Click an icon on the path to display a summarized account of its primary attributes. You can also maximize the pin to see the path details.

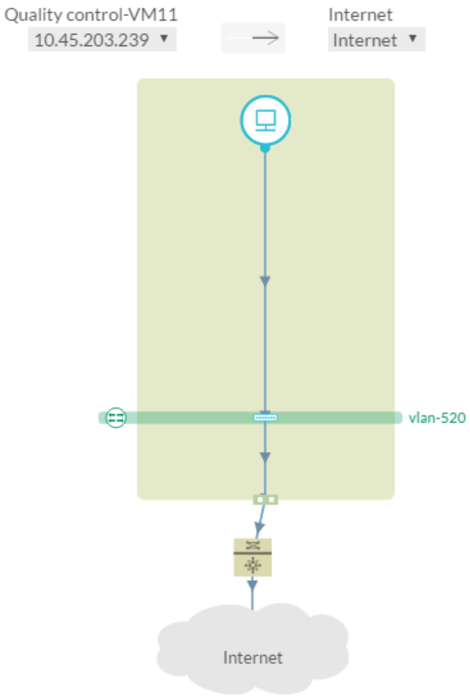

Security

NSX-T

VMware NSX-T is designed to address the emerging application frameworks and architectures that have heterogeneous endpoints and technology stacks. In addition to vSphere, these environments may also include other hypervisors, containers, bare metal, and public clouds. vRealize Network Insight supports NSX-T deployments where the VMs are managed by vCenter.

Considerations

- vRealize Network Insight supports NSGroups, NSX-T Firewall Rules, IPSets, NSX-T Logical Ports, and NSX-T Logical Switches.
- vRealize Network Insight does not support displaying information about routers, edge nodes, underlay and overlay paths of VMs..
- Only NSX-T 2.0 is supported.
- vRealize Network Insight does not support the KVM hosts as well as the individual ESXi servers added to NSX-T.
- vRealize Network Insight supports both NSX-V and NSX-T deployments. When you use NSX in your queries, the results include both NSX-V and NSX-T entities. NSX Manager lists both NSX-V and NSX-T Managers. NSX Security Groups list both NSX-T and NSX-V security groups. If NSX-V or NSX-T is used instead of NSX, then only those entities are displayed. The same logic applies to the entities such as firewall rules, IPSets, and logical switches.

To Add an NSX-T Manager as a Data Source

Here are the prerequisites for adding an NSX-T Manager as a data source:

- Before adding NSX - T, add at least one vCenter which is associated with NSX - T to vRealize Network Insight.
- It is recommended that you add all the vCenters associated with NSX-T as data sources in vRealize Network Insight.

To add an NSX-T Manager:

1. On the **Accounts and Data Source** page under **Settings**, click **Add Source**.
2. Under **VMware Manger** in the **Select an Account or Data Type** page, select **VMware NSX-T Manager**.
3. Provide the user credentials. The user should be a local user with the audit level permissions.

Examples for Queries

Here are some examples for queries related to NSX-T:

Table 8-1.

Queries	Search Results
NSX-T Manager where VC Manager=10.197.53.214	NSX-T Manager where this particular VC Manager has been added as the compute manager
NSX-T Logical Switch	Lists all the NSX-T Logical switches present in the particular instance of vRealize Network Insight
NSX-T Logical Ports where NSX-T Logical Switch = 'DB-Switch'	Lists the NSX-T logical ports belonging to that particular NSX-T logical switch, DB-Switch.

Table 8-1. (Continued)

Queries	Search Results
VMs where NSX-T Security Group = 'Application-Group' Or VMs where NSGroup = ‘Application-Group’	Lists all the VMs in that particular security group, Application-Group.
NSX-T Firewall Rule where Action='ALLOW'	Lists all the NSX-T Firewall Rules which have their action set as ALLOW.
NSX-T Firewall Rule where Destination Security Group = ‘CRM-Group’	Lists the firewall rules where the CRM-Group is the Destination Security Group. The results include both Direct Destination Security Groups and Indirect Destination Security Groups.
NSX-T Firewall Rule where Direct Destination Security Group = ‘CRM-Group’	Lists the firewall rules where the CRM-Group is the Destination Security Group. The results include only the Direct Destination Security Groups.
VMs where NSX-T Logical Port = ‘App_Port-Id-1’	Lists all the VMs which have that particular NSX-T Logical Port.

Cross vCenter NSX

In a cross-vCenter NSX environment, you can have multiple vCenter Servers, each of which must be paired with its own NSX Manager.

One NSX Manager is assigned the role of primary NSX Manager, and the others are assigned the role of secondary NSX Manager. The primary NSX Manager is used to deploy a universal controller cluster that provides the control plane for the cross-vCenter NSX environment. The secondary NSX Managers do not have their own controller clusters. The primary NSX Manager can create universal objects, such as universal logical switches. These objects are synchronized to the secondary NSX Managers by the NSX Universal Synchronization Service. You can view these objects from the secondary NSX Managers, but you cannot edit them there. You must use the primary NSX Manager to manage universal objects. The primary NSX Manager can be used to configure any of the secondary NSX Managers in the environment.

The following Universal objects are supported:

- Universal LDR
- Universal Transport Zone
- Universal Logical Switch
- Universal Firewall Rule
- Universal Security Group
- Universal IPSets
- Universal Service
- Universal Service Groups
- Universal Segment Range

Palo Alto Networks

vRealize Network Insight supports Palo Alto Panorama 8.0.

The Palo Alto Network features that are supported by vRealize Network Insight are as follows:

- Interrelation of Palo Alto and NSX entities: The VM membership of the address and the address group of Palo Alto Networks is computed based on the IP Address to VM mapping. This membership info can be queried as follows:
  - `VM where Address = <>`
  - `Palo Alto address where vm = <>`
  - `VM where Address Group = <>`
  - `Palo Alto address group where vm = <>`
- Query: You can perform a query for all the Palo Alto entities that are supported by vRealize Network Insight. All the entities are prefixed by Palo Alto. Some of the queries are as follows:

Table 8-2.

Entities	Queries

Palo Alto Address	`Palo Alto address where vm = <>`
	`VM where Address = <>`

Palo Alto Address Group	`Palo Alto address group where Translated VMs = <>`
	`VM where address group = <>`

Palo Alto Device	`Palo Alto Device where Version = <>`
	`Palo Alto Device where connected = true`
	`Palo Alto Device where family = 'PA-5060'`

Palo Alto Physical Device	`Palo Alto Physical Device where model = 'PA-5060'`
Palo Alto VM Device	`Palo Alto VM Device where model = 'PA-VM'`

Palo Alto Device Group	`Palo Alto Device Group where device = <>`
	`Palo Alto Device Group where address = <>`
	`Palo Alto Device Group where address group = <>`

Palo Alto Service	`Palo Alto service where Port = <>`
	`Palo Alto service where Protocol = <>`

Palo Alto Service Group	`Palo Alto service group where Member = <>`

Palo Alto Policy	`Palo Alto Policy where Source vm = <> and Destination vm = <>`
	`Palo Alto Policy where Source IP = <> and Destination IP = <>`

Palo Alto firewall	`Palo Alto firewall where Rule = <>`

Table 8-2. (Continued)

Entities	Queries
Palo Alto Zone	`Palo Alto Zone where device = <>`

Palo Alto Virtual System	`Palo Alto Virtual System where Device = <>`
	`Palo Alto Virtual System where Device Group = <>`

Note Other than the queries, you can also use facets to analyze the search results.

- VM to VM Path: As a part of the VM-VM topology, vRealize Network Insight displays the Palo Alto VM Series firewall on the host. The applicable rules are displayed when one clicks the firewall icon. If a firewall device (routing device) of Palo Alto Network is also present in the path, then that device is also displayed. When you click the device icon, you can see the basic information such as a Routing table, Interfaces, and a table containing the applied firewall rules.

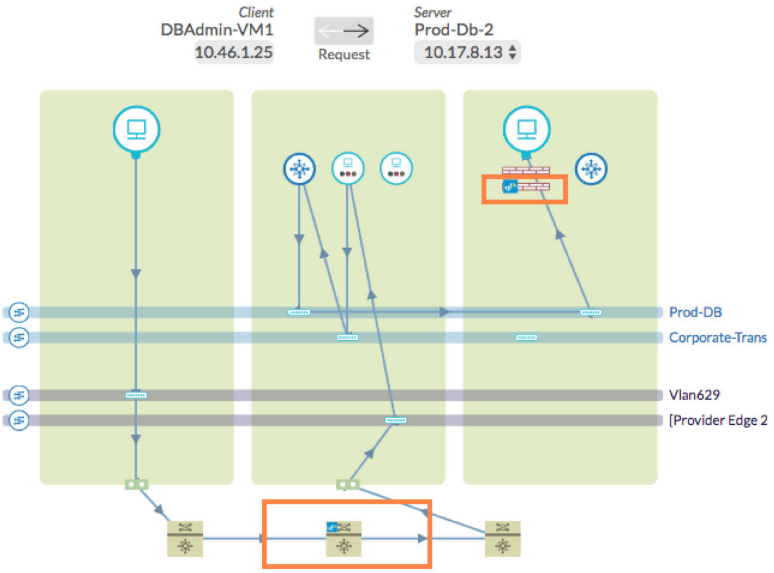

- You can view some system events related to the following scenarios for Palo Alto Networks:
  - Palo Alto device not connected to Panorama (manager)
  - NSX Manager not in registered with Panorama
  - NSX fabric agent not found on the ESX for palo alto device
  - Palo alto device not found on Panorama for NSX fabric agent
  - Out of sync security group membership data

- You can create and register multiple service definitions in Panorama with a given NSX manager. If different ESXi clusters have workloads that require the VM-Series firewall to handle traffic differently, then multiple service definitions are created. Each service definition has an associated device group from which the policies are picked. While displaying the VM-VM path in vRealize Network Insight, the correct set of policies based on the cluster information of the VM should be considered.

A sample Palo Alto Manager dashboard is shown as follows:

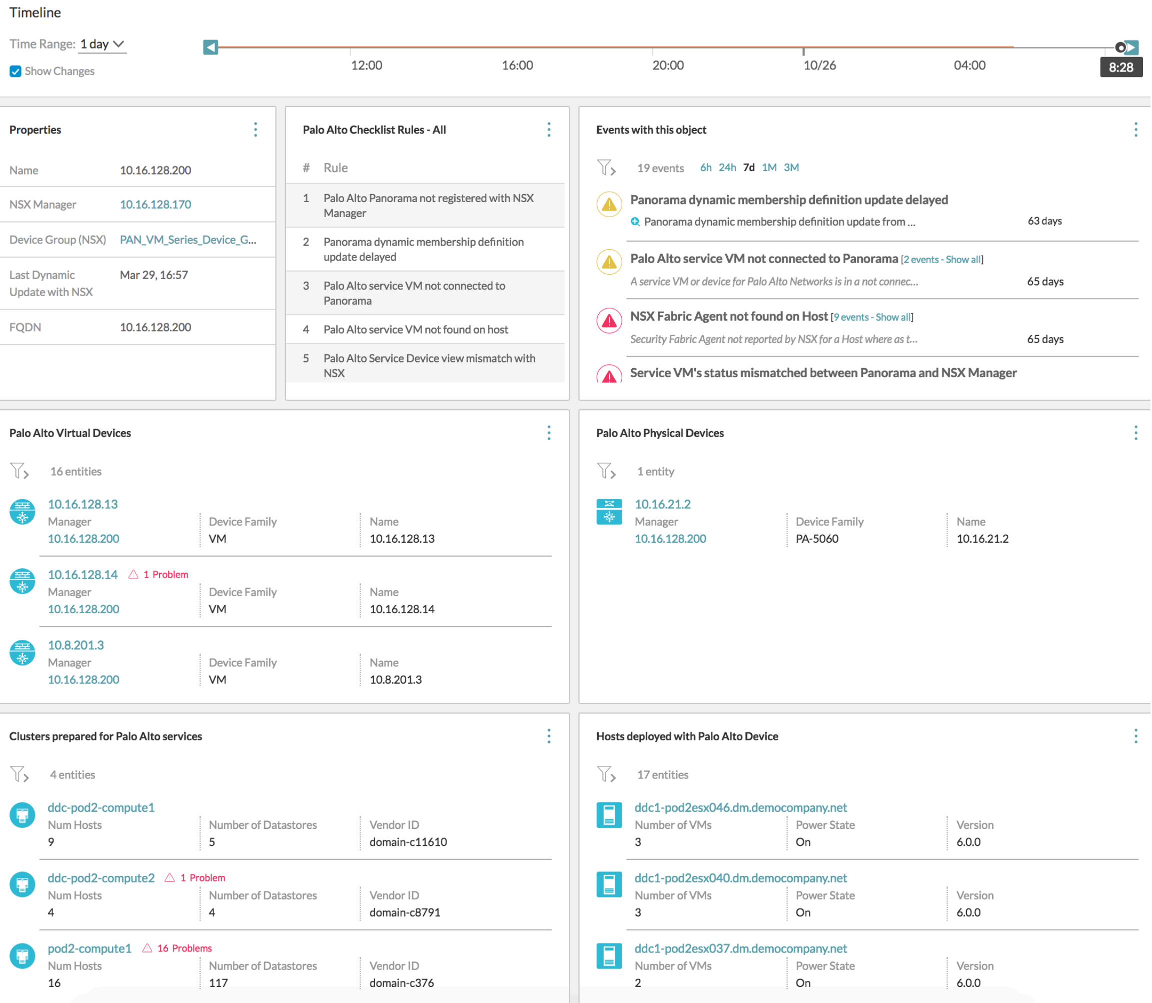

Timeline

Time Range: 1 day

☑ Show Changes

Properties

Name	10.16.128.200
NSX Manager	10.16.128.170
Device Group (NSX)	PAN_VM_Series_Device_G...
Last Dynamic Update with NSX	Mar 29, 16:57
FQDN	10.16.128.200

Palo Alto Checklist Rules - All

#	Rule
1	Palo Alto Panorama not registered with NSX Manager
2	Panorama dynamic membership definition update delayed
3	Palo Alto service VM not connected to Panorama
4	Palo Alto service VM not found on host
5	Palo Alto Service Device view mismatch with NSX

Events with this object

19 events 6h 24h 7d 1M 3M

Panorama dynamic membership definition update delayed
Panorama dynamic membership definition update from ... 63 days

Palo Alto service VM not connected to Panorama [2 events - Show all]
A service VM or device for Palo Alto Networks is in a not connec... 65 days

NSX Fabric Agent not found on Host [9 events - Show all]
Security Fabric Agent not reported by NSX for a Host where as t... 65 days

Service VM's status mismatched between Panorama and NSX Manager

Palo Alto Virtual Devices

16 entities

10.16.128.13 Manager 10.16.128.200	Device Family VM	Name 10.16.128.13
10.16.128.14 △ 1 Problem Manager 10.16.128.200	Device Family VM	Name 10.16.128.14
10.8.201.3 Manager 10.16.128.200	Device Family VM	Name 10.8.201.3

Palo Alto Physical Devices

1 entity

10.16.21.2 Manager 10.16.128.200	Device Family PA-5060	Name 10.16.21.2

Clusters prepared for Palo Alto services

4 entities

ddc-pod2-compute1 Num Hosts 9	Number of Datastores 5	Vendor ID domain-c11610
ddc-pod2-compute2 △ 1 Problem Num Hosts 4	Number of Datastores 4	Vendor ID domain-c8791
pod2-compute1 △ 16 Problems Num Hosts 16	Number of Datastores 117	Vendor ID domain-c376

Hosts deployed with Palo Alto Device

17 entities

ddc1-pod2esx046.dm.democompany.net Number of VMs 3	Power State On	Version 6.0.0
ddc1-pod2esx040.dm.democompany.net Number of VMs 3	Power State On	Version 6.0.0
ddc1-pod2esx037.dm.democompany.net Number of VMs 2	Power State On	Version 6.0.0

Check Point Firewall

vRealize Network Insight currently supports post R80 version of the Check Point firewall.

You can perform a query for all the Checkpoint entities that are supported by vRealize Network Insight. All the entities are prefixed by `Check Point`. Some of the queries for Checkpoint are as follows:

Table 8-3.

Entities in Check Point	Keywords	Queries

IPset	Check Point Address Range
	Check Point Network
		`vm where Address Range = <>`
		`vm where Address Range = <>`
		`Check Point Address Range where Translated VM = <>`

Grouping	Check Point Network Group	`Check Point Network Group where Translated VM = <>`
		`vm where Network Group = <>`

Service/ Service Group	Check Point Service
	Check Point Service Group
		`Check point service where Port = <>`
		`Check point service where protocol = <>`

Access Layer	Check Point Access Layer	`Check Point Policy where Access Layer = <>`

Policy Package	Check Point Policy package	`Check Point Policy where Policy Package = <>`
		`Check Point Policy Package where Rule = <>`

Policy	Check Point Policy	`Check point policy where source ip = <> and Destination IP = <>`
		`Rule where source ip = <> and Destination IP = <> (will display other rules– nsx, redirect along with check point policies in the system)`

Gateways and Gateway Cluster	Check Point Gateway Check Point Gateway Cluster	`Check Point Gateway Cluster where Policy Package = <>`

A sample Check Point Manager dashboard is shown as follows:

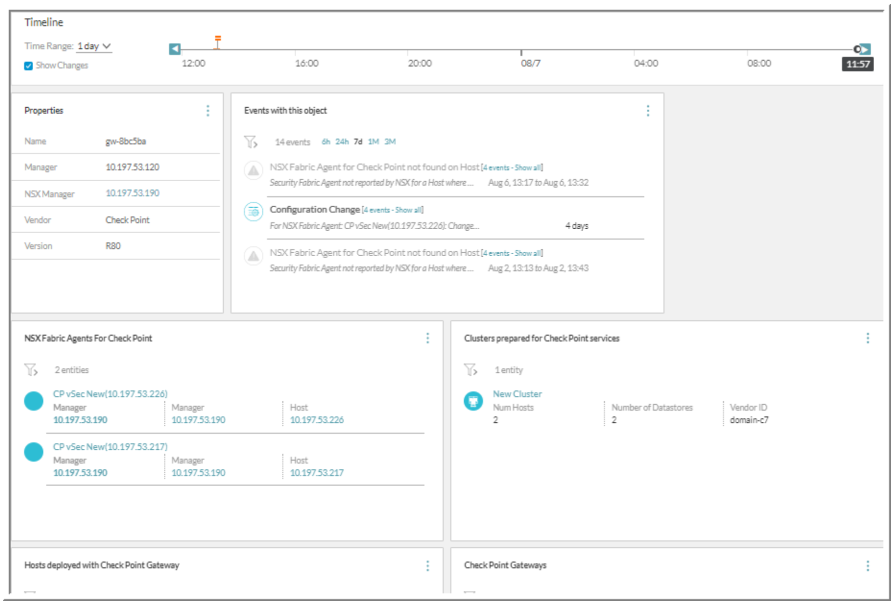

Timeline

Time Range: 1 day

☑ Show Changes

Properties

Name	gw-8bc5ba
Manager	10.197.53.120
NSX Manager	10.197.53.190
Vendor	Check Point
Version	R80

Events with this object

14 events 6h 24h 7d 1M 3M

NSX Fabric Agent for Check Point not found on Host [4 events - Show all]
Security Fabric Agent not reported by NSX for a Host where ... Aug 6, 13:17 to Aug 6, 13:32

Configuration Change [4 events - Show all]
For NSX Fabric Agent: CP vSec New(10.197.53.226): Change... 4 days

NSX Fabric Agent for Check Point not found on Host [4 events - Show all]
Security Fabric Agent not reported by NSX for a Host where ... Aug 2, 13:13 to Aug 2, 13:43

NSX Fabric Agents For Check Point

2 entities

CP vSec New(10.197.53.226)

Manager	Manager	Host
10.197.53.190	10.197.53.190	10.197.53.226

CP vSec New(10.197.53.217)

Manager	Manager	Host
10.197.53.190	10.197.53.190	10.197.53.217

Clusters prepared for Check Point services

1 entity

New Cluster

Num Hosts	Number of Datastores	Vendor ID
2	2	domain-c7

Hosts deployed with Check Point Gateway

Check Point Gateways

Also, in a VM-VM topology, you can see the Check Point Service VMs on a host to signify the Check Point rules applied on particular traffic.

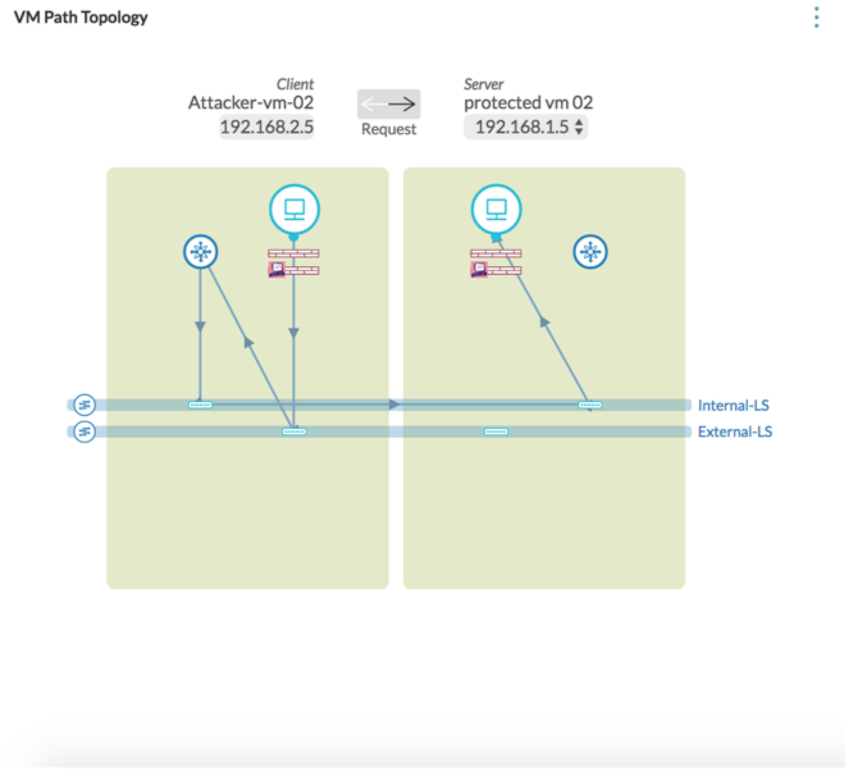

You can view some system events related to the following scenarios for Check Point:

- NSX fabric agent not found on the ESX for check point gateway.
- Check point service vm not found.
- Check point gateway sic status not communicating.

- Discovery and update events for check point entities like address range, networks, policies, groups, policy package, service, service group, and so on.

Security Groups

Security Groups are a set of groups that are managed through a common set of permissions.

The Security Group topology has the following two views:

Firewall View

The Security Group firewall topology displays the relation between the selected Security Group and other Security Groups by showcasing the firewall rules that are applicable between the Security Groups.

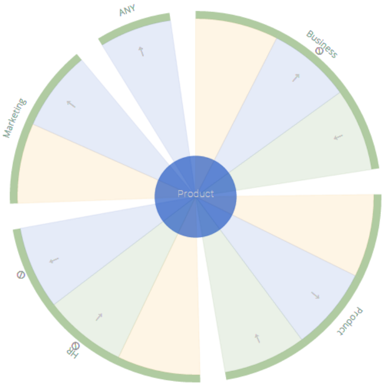

Container View

The Security Group container topology displays how the Security Group is structured with respect to its parent Security Groups or children (Security Groups or other entities).

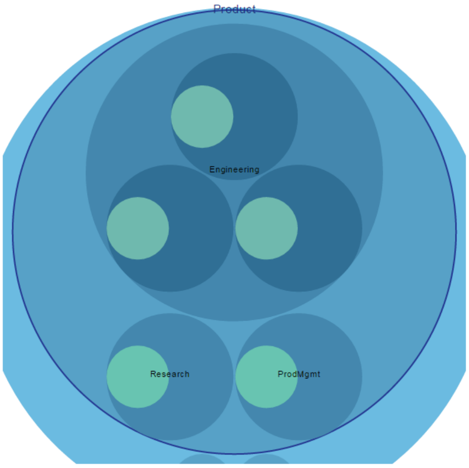

Operations

9

This chapter includes the following topics:

- Working with Events
- Notifications
- Pins and Pinboards

Working with Events

On the **Settings** page, click **Events** to view the various types of events:

- System Events
- User Defined Events
- System Health Events

Working with System Events

The event is defined either by the system or the user. The system events are predefined events.

The system events are listed in the System Events page under Settings. The following fields are specified for each event. You can filter the information based on your requirements in all the following columns except the Event column.

Table 9-1.

Column	Description
Event	This field specifies the name of the event.
Severity	This field specifies the severity of the event. You can set it to the following values: ■ Critical ■ Moderate ■ Warning ■ Info
Type	This field specifies if the event denotes a problem or a change.

Table 9-1. (Continued)

Column	Description
Entities	This field specifies that the event is configured to either include or exclude entities for event generation. By default, the value is `All`.
Notifications	This field specifies the types of notifications that are sent. The notifications can be sent by email or SNMP trap or both.
Enabled	This option is selected if the event is enabled.

When you hover the mouse on each event, you can see More Information. By clicking this option, you can see the description, event tags, and entity type for that event.

You can perform the following tasks on the system events:

- Edit an event
- Perform bulk edit
- Disable an event for a particular entity

Edit System Event

To edit an event:

1. Click the edit icon after the **Enabled** column for a particular event.
2. You can add or remove event tags if required.
3. You can change the severity.
4. Check Include/Exclude entities if you want the event to be enabled or disabled for selected entities.
5. To create inclusion rules:
   - a Select **Inclusion List**.
   - b Specify the entities which you want to include for the event under **Conditions**.
6. To create exclusion rules:
   - a Select **Exclusion List**.
   - b Specify the entities which you want to exclude for the event under **Conditions**.

   Note

   - You can create multiple rules in both inclusion and exclusion lists.
   - When you select `NSX Manager`, you can add exceptions in both the lists. You can define exception if you want the inclusion or the exclusion rule to hold exception for a particular entity.
   - You can also specify `Custom Search` by writing your own query to include or exclude entities.
7. Select **Enable Notifications** if you want to configure when the notifications have to be sent. Specify the email address and the frequency at which you would like to receive the emails.

Prerequisites

Perform a Bulk Edit on an Event

1. In the **System Events** page, when you select multiple events, the options **Enable**, **Disable**, and **Edit** appear above the list.
2. Click **Edit**.
3. In the **Edit** page, you have the following options:
   - **Override existing values**: In this option, only the fields that you edit will get overwritten.
   - **Add to existing**: In this option, you can add to the existing values such as email addresses and event tags.
4. Click **Submit**.

Disable an Event

1. You can select an event in the **Open Problems** widget in the Homepage. You can also enter **Problems** in the search bar and select an event from the list.
2. Select a particular event and click **Archive**.
3. Select **Disable all events of this type in future for** and select an entity or all entities.
4. Click **Save**.

Note The changes made in severity, tags, or inclusion/exclusion rules will reflect for the future events. The existing events continue to show the old configuration.

Working with User-Defined Events

The user-defined events are based on search.

All the user-defined events are listed on the **User-defined Events** page under **Settings**. The following fields are specified for each event.

Table 9-2.

Field	Description
Name (Search Criteria)	This field specifies the name of the event and the search criteria for the event.
Severity	This field specifies the severity of the alert. You can set it to the following values: ■ Critical ■ Moderate ■ Warning ■ Info
Type	This field specifies if the event denotes a problem or a change.

Table 9-2. (Continued)

Field	Description
Notify when	This field specifies when the notification has to be sent.
Created By	This field specifies who created the event.
Enabled	This option is selected if the event is enabled.

You can edit or delete the event. While editing it, you can specify the email address and the frequency of the email notification.

Working with Platform Health Events

The Platform Health Events page is your one-stop page to view all the events that provide details on the overall health of the system. These events might have occurred on a datasource or a node in the infrastructure. You can also view these events through search.

Table 9-3.

Field	Description
Event	This field specifies the name of the event.
Severity	This field specifies the severity of the event. You cannot change the severity of the event.
Type	This field specifies if the event denotes a problem or a change.
Notifications	This field specifies the types of notifications that are sent. The notifications can be sent by email or SNMP trap or both.

Notifications

Search-based Notifications

The search-based notifications can be categorized as follows:

- System-based notification
- User-defined notification

System-based notification parameters are predefined and upon activating notification alert, notification in the form of mails are sent. User-defined notifications are set by users, based on their requirements. You can create email notifications based on your search query. After you run a search, on the Results page, the **Create notification** option is displayed. For each search, you can:

- Select the condition when you want to receive the notifications.
- Define how frequently you want to receive the notifications.
- Enter the email recipients for each notification (by default, your email ID is present in the receiver's list; you can also add multiple email IDs).

For a user-defined search:

- It is mandatory for you to assign a name to the search-based notification.
- It is mandatory to select the severity of a search-based event that is marked as a problem.
- The user-defined events are uniquely identified by the search criteria.
- You can specify the notification frequency as **Immediately** or **As a daily digest**.

You can manage your notifications from the **Settings > Search-based Notifications** page. On the **Search-based Notifications** page, you can view the existing notifications, edit them, activate or deactivate them, and also delete unwanted notifications.

Event Notification Email

The notifications are sent in the form of emails.

To set up notification, users have to first configure the mail server. To know how to configure mail server, see Configuring mail server .

Specifying Notification Events for Emails to be sent

Users can specify events for which mail notifications are to be sent.

To specify events

1. On the **Settings** page, click **Search-based Notifications**, or simply search for any information using the Search box.
2. On the Search-based Notifications page, click the **Create Notification** icon. A notification dialog box is displayed.
3. In the **Receive notification when** box, select the event on the occurrence of which notifications are to be sent.
4. In the **Notify** box, select the frequency at which the notifications are to be sent.
5. If the event is undesirable, select the **Mark it as a problem** check box.
6. Enter the email addresses to which the notifications are to be sent, and then click **Save**.

Note To verify whether the notification mail is correctly set up, click **Send test Email**.

Event Notifications

vRealize Network Insight contains a list of predefined system events (system problems and system changes) for which you can receive automated email notifications every four hours.

You can view the list of notifications on the **Settings > System Notifications** page.

Archiving Problems

Archiving a Problem

1 Click the Show All link (if there is more than one instance of an event) to display all instances of the event.

2 Hover on the instance of the event that you want to archive to display a set of icons, and then click the Archive icon .

3 In the Event specific dialog box

   a Select This event from the You are about to archive list, if you want to archive only this event.

   b Select All events of this type from the You are about to archive list, if you want to archive all events of the same type in the system.

4 Click **Save**.

Viewing all archived events

1 On the Home page, type events in Search box and press **Enter**. A list of events is displayed.

2 On the left hand pane, in the Archived facet, select True checkbox (highlighted in the screenshot below).

You can view all archived events here.

To restore an archived event

1 On the Archived event, click the Archived icon . (See the preceding section on To view an archived event to know how to go to the Archived events page).

2 In the Event specific dialog box

   a Select This event from the You are about to restore from archive list, if you want to restore only this event.

   b Select All events of this type from the You are about to restore from archive list, if you want to restore all similar type of events.

   c Click Save to complete restoring.

Disabling Events

Users can selectively disable events and prevent notifications from being sent in future.

To disable event notification

Method 1

1 On the event, click the **Show All** link (if there is more than one instance of an event) to display all instances of the event.

2 Hover on the instance of the event, whose notification you want to disable. This displays a set of icons, click the Archive icon .

3 In the Event specific dialog box, select the **Disable all events of this type in future** checkbox, and then click **Save**.

Method 2

1. On the top-right corner of **Home** page, click the **Profile** icon, and then click **Settings**.
2. In the **Settings** section, click **Event Notifications** to see a list of all enabled and disabled events.
3. On the enabled event that you want to disable, in the **Enabled** column, click the left-side space of the respective slider.
4. In the **Confirm Action** dialog box, click **Yes**.

Configuring Event Notification Service

Users can enable customer notifications for different events

To set notification services

1. On Settings, go to Event Notification, and click the (edit) icon corresponding to the problem, for which you want to enable e-mail notifications and SNMP.
2. In the Edit System Notification dialog box, enter the email address to which you want the email notification to be sent. In the Email Frequency box, select the time frequency at which you want to receive notifications.
3. Select the Enable SNMP trap for this event checkbox to set SNMP notifications.
4. Click **Save**.
5. Once successfully enabled, the respective mail and SNMP icons appear, as highlighted in the screenshot below.

Pins and Pinboards

All parts of the application are denoted as pins; fundamental units that can be saved and grouped to club data that you think can be useful together and to share them with other members of your team. You can pin a search query and also the pins that are available for an entity.

To add a pin, click the Pin icon. All your saved pins are displayed in Pinboards section which can be invoked by clicking the Pinboard icon in the header.

Pins

The information on each entity page is segregated into pins. All the entity pages are made up of pins and each pin contains a specific bit of information related to the entity.

The pins have the following features:

- You can maximize the view of any pin using the More options ( ) button and also view more information about the pin using the **Help** option.
- Pins can also contain filters so that you can drill down on the data that is displayed on the pin.
- Many pins also contain the Export as CSV option so that you can export the data present in the pin in CSV format. You can select the specific properties and the number of CSV rows you want to export in the dialog that is displayed.

Types of Pins

Most of the pins that are available in the software can be categorized into the following:

Metrics Pins

The metrics pins show important metrics pertaining to the selected entity.

The metrics pin uses the cubism graph to display data by dividing each graph into two bands and transposing the higher value one over another. The higher values hence are shown in darker color and are easier to discern.

You can select the particular metric to display from the drop-down present in the pin header and change the selection of entities to display.

The time range can be modified by either using the range presets or entering in a custom date/time.

An example of the Metrics pin is the VM Metrics pin. This pin displays the network traffic rate, network Tx rate, network Rx rate, and packet drops of the virtual machine.

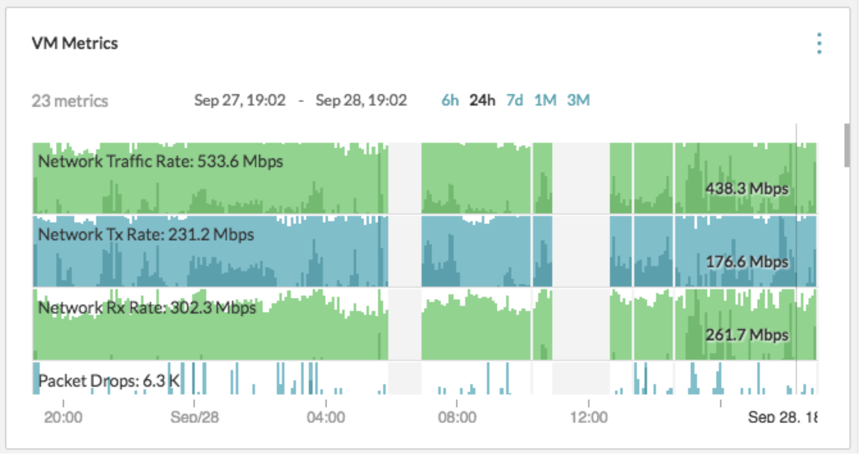

Entity List View Pins

The Entity List View pins display a list of entities that are grouped by a common theme. The list shows important attributes per entity.

You can see more attributes of a particular entity by clicking the magnify icon on the far right. Clicking the entity name takes you to the entity page.

Like other pins, the filter icon houses various facets with which the list can be filtered. An example of the Entry List View pin is the VM Neighbors pin. By default, this pin shows the VMs that are present on the same host. You can also filter VMs by Security Groups, VXLAN, and datastore.

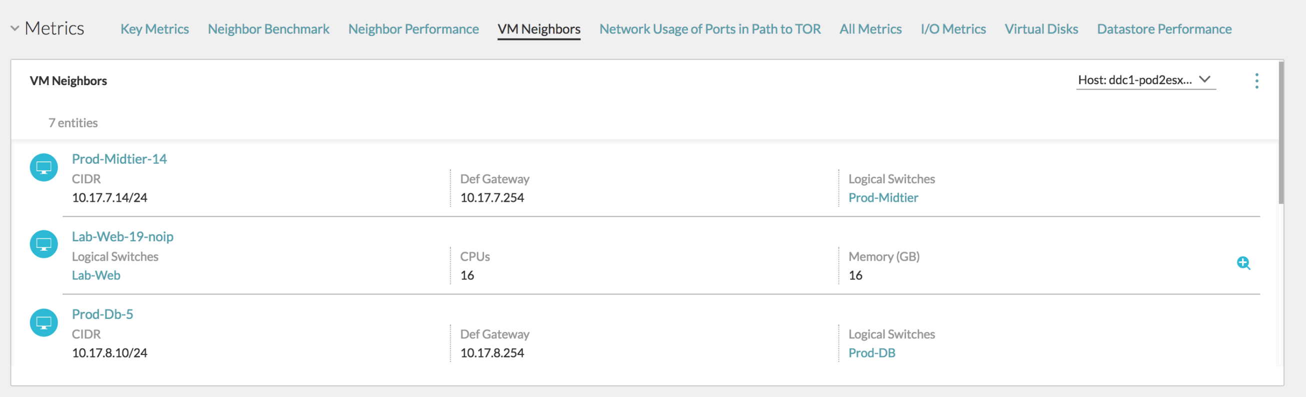

VM Neighbors

Host: ddc1-pod2esx...

7 entities

Prod-Midtier-14
CIDR 10.17.7.14/24 | Def Gateway 10.17.7.254 | Logical Switches Prod-Midtier

Lab-Web-19-noip
Logical Switches Lab-Web | CPUs 16 | Memory (GB) 16

Prod-Db-5
CIDR 10.17.8.10/24 | Def Gateway 10.17.8.254 | Logical Switches Prod-DB

Event View List Pins

The Events List view pins provide a list of events in chronological order for a particular entity or group of entities (that can be selected from the dropdown in the pin header).

You can change how far back in time (from now) should the pin show the events by using the available presets or entering in a custom date/time. Other filter options such as **Event Status** and **Event Type** can be selected by clicking on the filter icon.

In the below image, the events related to VM Prod-db-vm21 and its related entities are displayed. You can click the entity name to view events from other related entities. Using the filter you can filter the events based on their status and their types. An event can be a change or a problem related to an entity.

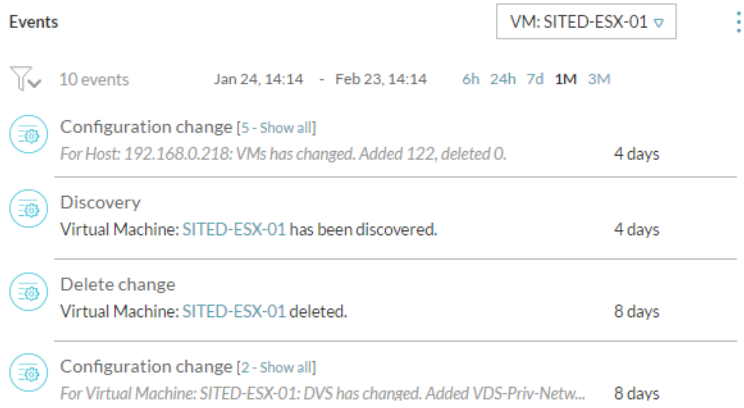

10 events Jan 24, 14:14 - Feb 23, 14:14

Configuration change [5 - Show all]
For Host: 192.168.0.218: VMs has changed. Added 122, deleted 0. 4 days

Discovery
Virtual Machine: SITED-ESX-01 has been discovered. 4 days

Delete change
Virtual Machine: SITED-ESX-01 deleted. 8 days

Configuration change [2 - Show all]
For Virtual Machine: SITED-ESX-01: DVS has changed. Added VDS-Priv-Netw... 8 days

You can search for the events by using the events search query. You can search for open or closed events with queries such as open events or closed events. You can also search for problems with the same modifiers.

Pinboards

Pinboards are how you group pins together. You can pin any widget from any page to make it easier to access various data.

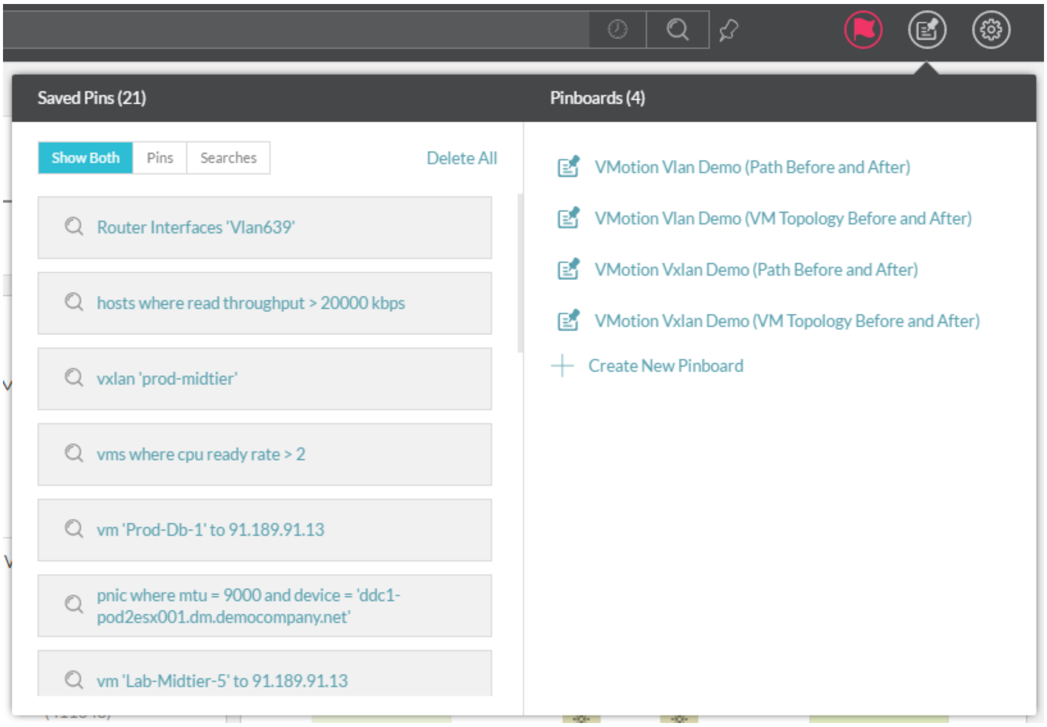

Saved Pins (21)

Pinboards (4)

Router Interfaces 'Vlan639'

hosts where read throughput > 20000 kbps

vxlan 'prod-midtier'

vms where cpu ready rate > 2

vm 'Prod-Db-1' to 91.189.91.13

pnic where mtu = 9000 and device = 'ddc1-pod2esx001.dm.democompany.net'

vm 'Lab-Midtier-5' to 91.189.91.13

VMotion Vlan Demo (Path Before and After)

VMotion Vlan Demo (VM Topology Before and After)

VMotion Vxlan Demo (Path Before and After)

VMotion Vxlan Demo (VM Topology Before and After)

To create a pinboard:

1 Click **Pinboards** and select the pins that you want to add to the pinboard from the **Saved Pins** section.

2 Click **Create New Pinboard**.

3 Enter the name of the pinboard, add a note related to any information that you want to share with others, and enter the email IDs or name of the users with whom you want to share the pinboard and click **Create**.

4 To add a pin to an existing pinboard, after selecting the pin, click **Add** beside an existing pinboard where you want to add the pin.

To share the pinboard link:

1 Click **Share** on the rightmost side of the pinboard.

2 Click **Copy** to copy the link. You can share this link only with the users whom you have added in the "Share Pinboard with" list during the creation of the pinboard.

To view the pinboard by using the time selector:

You can also jump to a pinboard view on a particular date and time by using the time selector.

1 In the time selector just next to **Jump**, select **Custom time** or **Current Time** to view the flow of data for a particular pin.

2 Click **Update**. The flow of data for that particular pin can be seen for the given time.

Flows

10

This chapter includes the following topics:

- Enabling IPFIX Configuration
- Flow Support for Physical Servers
- View Blocked and Protected Flows
- Network Address Translation (NAT)
- VPC Flow Logs

Enabling IPFIX Configuration

IPFIX is an IETF protocol for exporting flow information.

A flow is defined as a set of packets transmitted in a specific timeslot, and sharing the same 5-tuple values - source IP address, source port, destination IP address, destination port, and protocol. The flow information may include properties such as timestamps, packets/bytes count, Input/Output interfaces, TCP Flags, VXLAN ID, Encapsulated flow information, and so on.

IPFIX Configuration on VDS and DVPG

A VDS in vSphere environment can be configured to export flow information using IPFIX. Flow monitoring has to be enabled on all the port groups attached to the VDS. If packets arrive on port X of a VDS and exit from port Y, a corresponding flow record is emitted if flow monitoring is enabled on port Y.

To analyze the complete information of any session, the IPFIX data about packets in both the directions is required. Refer the following diagram where VM-A is connected to DVPG-A and is talking to VM-C. Here DVPG-A will only provide data about the C→A packets, and DVPG-Uplink will provide data about A→C packets. To get the complete information of A's traffic, IPFIX should be enabled on DVPG-A, DVPG-uplink.

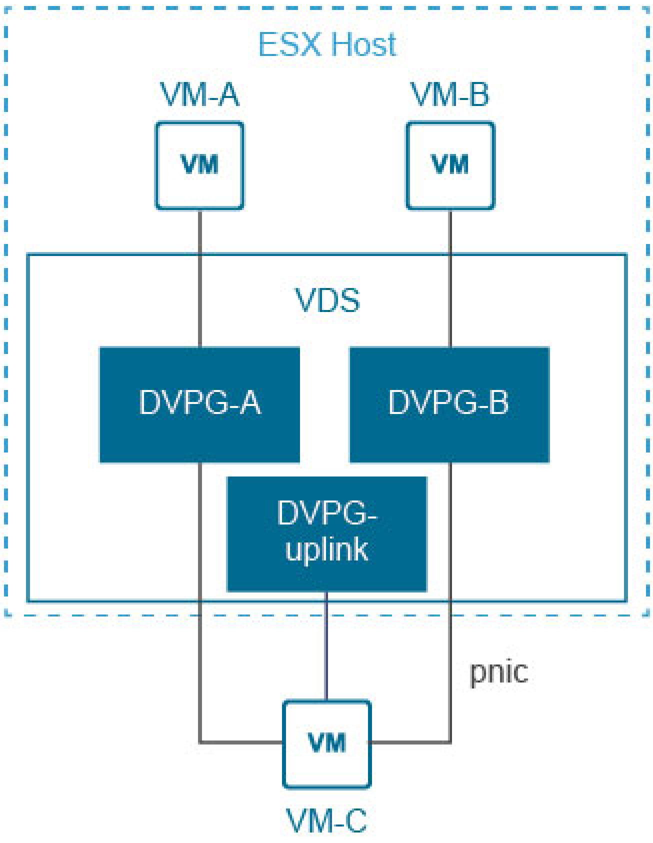

vRealize Network Insight Proxy VM has built-in collector/receiver for IPFIX flow information. You can enable the IPFIX information collection in the vCenter Data Source settings at various levels of granularity.

Enabling IPFIX Configuration on VDS and DVPG

To enable IPFIX information at vCenter level:

Procedure

1 Select **Enable Netflow (IPFIX)** when you are adding vCenter.

2 Select the VDS for which you want to enable IPFIX from the list of available VDS in vCenter.

3 A notification icon is displayed for the VDS where one of the hosts has unsupported version of ESXi. If vRealize Network Insight has detected that IPFIX is already configured for a VDS with some other IP address apart from vRealize Network Insight Proxy VM, then it displays the **Override** button. Click **Override** to view the list of DVPGs under that VDS.

4 The list of available DVPGs for the selected VDS is displayed. All the DVPGs are selected by default. Turn **Manual Selection** on to select specific DVPGs for which you want to enable IPFIX. Select the desired DVPGs and click **Submit**.

   Note The DVPG with a notification icon denotes that it is the uplink DVPG and it has to be selected.

VMware NSX IPFIX Configuration

VMware NSX IPFIX provides network monitoring data similar to that provided by physical devices and gives administrators a clear view of virtual network conditions.

VMware NSX virtualizes the network by allowing the network administrator the ability to decouple the network from physical hardware. This functionality makes it easy to grow and shrink the network as needed and making the network transparent to the applications traversing it.

By using NSX IPFIX in a virtualized network, the network administrators gain visibility into the virtual overlay network. The VXLAN IPFIX reporting using Netflow is enabled on the host uplink. It provides visibility on the VTEP that is encapsulating the packet, and the details of the VM that generated the inter-host traffic on an NSX Logical Switch (VXLAN).

The distributed firewall implements stateful tracking of flows. As these tracked flows go through a set of state changes, IPFIX can be used to export data about the status of that flow.

The tracked events include flow creation, flow denial, flow update, and flow teardown. The denied events are exported as syslogs.

Enabling VMware NSX-V IPFIX

To enable VMware NSX-V IPFIX in vRealize Network Insight:

Prerequisites

- Ensure that you have the security administrator or enterprise administrator credentials.
- It is recommended that you enable VDS IPFIX on all the DVS and DVPGs from which NSX IPFIX data has to be collected. You can enable VDS IPFIX from the details page of the associated vCenter.

Procedure

- ◆ Select **Enable IPFIX** when adding or editing a NSX-V Manager data source.

Adding a NetFlow Collector

Procedure

1 In the **Settings** page, click **Accounts and Data Sources**.

2 Click **Add Source**.

3 In the **Flows** section, click **NetFlow Collector**. The Collector VM that is used for NetFlow is a dedicated collector. It cannot be used for any other data source. If any other data source is also added on the proxy server, it is not available as a NetFlow collector.

Flow Support for Physical Servers

vRealize Network Insight supports the device that sends the NetFlow data of versions v5, v7, and v9. If the DNS Mapping and Subnet-VLAN mapping information is provided, vRealize Network Insight can enrich the NetFlow data with DNS Domains, DNS Host Names, Subnets, and Layer 2 networks. This feature is available for the Enterprise License users only.

To configure NetFlow in vRealize Network Insight, perform the following steps:

1 Adding a NetFlow Collector

2 Configuring a NetFlow Collector in a Physical Device

3 Physical IP and DNS Mapping

4 Physical Subnets and VLANs

Configuring a NetFlow Collector in a Physical Device

To send the NetFlow information to the vRealize Network Insight NetFlow collector, configure the physical device manually. Here are the steps for the configuration in most of the physical devices:

1 Create a flow record.

  The required fields for a flow record are as follows:

  - Mark the following fields as `Match`.
    - `ipv4 protocol`
    - `ipv4 source address`
    - `ipv4 destination address`
    - `transport source-port`
    - `transport destination-port`
    - `interface input`
  - Mark the following fields as `Collect`.
    - `direction`
    - `counter bytes`
    - `counter packets`
    - `timestamp sys-uptime first`
    - `timestamp sys-uptime last`

- Mark the following field as `Match` or `Collect`. If not, skip it.
    - `transport tcp flags`

2 Create a flow exporter.

- Provide vRealize Network Insight NetFlow Proxy IP and Port 2055.

3 Configure the flow cache as follows:

- Active timeout: 30 seconds
- Inactive timeout: 60 seconds

4 Create the flow monitor using the created flow record and flow exporter.

5 Configure the monitor on each interface.

Prerequisites

The sample steps to configure the physical devices are provided in the following sections:

- Cisco 4500
- Cisco Nexus 1000v
- Cisco Nexus 9000

Note The steps may vary from version to version and device to device.

Cisco 4500

1 To create the flow record

```
configure terminal

flow record netflow-original

match ipv4 protocol

match ipv4 source address

match ipv4 destination address

match transport source-port

match transport destination-port

match interface input

collect transport tcp flags

collect counter bytes

collect counter packets

collect timestamp sys-uptime first

collect timestamp sys-uptime last

End
```

2 To create the flow exporter

```
configure terminal

flow exporter e1

destination <PROXY_IP>

transport udp 2055

end
```

3 To create the flow monitor

```
configure terminal

flow monitor m1

record netflow-original

exporter e1

end
```

4 To configure the timeouts

```
configure terminal

cache timeout inactive 30

cache timeout active 60

end
```

5 To configure the flow monitor for each interface on the ingress mode and the egress mode or at least the ingress mode

```
configure terminal

interface <INTERFACE_NAME>

ip flow monitor m1 unicast input

end
```

Cisco Nexus 1000v

1 To configure timeouts

```
configure terminal

Active timeout 60

Inactive timeout 15

end
```

2 To configure the exporter

```
configure terminal
flow exporter <EXPORTER_NAME>
destination <PROXY_IP>
transport udp 2055
source <VSM_IP_OR_SUBNET>
end
```

3 To configure the flow monitor for each interface:

```
configure terminal
flow monitor <MONITOR_NAME>
record netflow-original
exporter <EXPORTER_NAME>
end
```

4 To configure the flow monitor for each interface on the ingress mode and the egress mode or at least the ingress mode

```
configure terminal
port-profile type vethernet <IF_NAME>
ip flow monitor <MONITOR_NAME> input
ip flow monitor <MONITOR_NAME> output
.
.
end
```

Cisco Nexus 9000

Here are some of the sample device commands for Cisco Nexus 9000:

1 To enable the NetFlow feature

```
configure terminal
feature netflow
end
```

2 To create flow record

```
configure terminal
flow record vrni-record
match ipv4 protocol
match ipv4 source address

match ipv4 destination address

match transport source-port

match transport destination-port

match interface input

collect transport tcp flags

collect counter bytes

collect counter packets

collect timestamp sys-uptime first

collect timestamp sys-uptime last

End
```

3 To create flow exporter

```
configure terminal

flow exporter vrni-exporter

destination <PROXY_IP>

transport udp 2055

version 9

source <INTERFACE_NAME>

end
```

4 To create the flow monitor for each interface

```
configure terminal

flow monitor vrni-monitor

record vrni-record

exporter vrni-exporter

end
```

5 To configure timeouts

```
configure terminal

cache timeout inactive 30

cache timeout active 60

end
```

6 To configure the flow monitor for each interface on the ingress mode and the egress mode or at least the ingress mode

```
configure terminal

interface <INTERFACE_NAME>

ip flow monitor vrni–monitor input

end
```

Enriching Flows and IP Endpoints

You can import the DNS mapping and the subnet-VLAN mapping information through the UI.

The flow information is enriched with the following types of information based on the import of the DNS data and the specification of subnet-VLAN mappings.

- Source DNS Domain
- Source DNS Host Name
- Destination DNS Domain
- Destination DNS Host Name
- Source L2 Network
- Source Subnet network
- Destination L2 Network
- Destination Subnet network

The IP Endpoint information is enriched with the following types of information based on the import of the DNS data and the specification of subnet-VLAN mappings.

- DNS Domain
- DNS Host Name
- FQDN
- L2 Network
- Subnet network

For more information on enriching flows through the DNS information, refer Physical IP and DNS Mapping.

For more information on enriching flows through the Subnet-VLAN mapping, refer Physical Subnets and VLANs.

Note

- The DNS mapping and subnet information are enhanced only for the physical IPs. No subnet or DNS mapping information is associated with any virtual NIC.
- The information is enriched only for flows that have been seen by vRNI after this information has been imported.

Search for Physical to Physical Flows

You can search for the physical to physical flows based on the following attributes:

- Source DNS Host
- Destination DNS Host
- Source DNS Domain
- Destination DNS Domain
- Source Subnet Network
- Destination Subnet Network

You can search for Physical-Physical flows based on the following attributes. A few examples of flow search query using the enriched DNS and Subnet-VLAN mapping information are as follows:

```
bytes,Dns Domain,Dns Host,l2 network of flows where flow type = 'Physical-Physical'
```

```
bytes,Dns Domain,Dns Host,l2 network of flows where flow type = 'Source is VM' and
flow type = 'Destination is Physical'
```

```
bytes,Dns Domain,Dns Host,l2 network of flows where flow type = 'Source is Internet'
and flow type = 'Destination is Physical'
```

View Blocked and Protected Flows

The NSX-IPFIX integration enables the visibility of the blocked and protected flows in the system.

The basic filters in the Micro-Segmentation Planning page are as follows:

- All Allowed Flows: This option is selected by default. To see all the flows for which the action in the firewall rules is set to **Alllowed**, select this option.
- Dropped Flows: This option helps to detect the dropped flows and planning the security in a better way.
- All Protected Flows: This option helps to detect all the flows which have a rule other than of the type `any(source) any(dest) any(port) allow` associated with it. Such flows are known as protected flows.
- All Unprotected Flows: This option helps to detect all the flows that have the default rules of the type `any(source) any(dest) any(port) allow`. Such flows are known as unprotected flows.

The firewall rules are visible only for the allowed and unprotected flows.

For example, if you are in the planning phase and you want to see the allowed flows in the system, perform the following steps:

1. On the Micro-Segmentation Planning page, for a particular group, select **All Allowed Flows** from the drop-down menu.

2 Click the dropped flows in the topology diagram to see the corresponding recommended firewall rules.

3 Implement those firewall rules by exporting them into NSX manager.

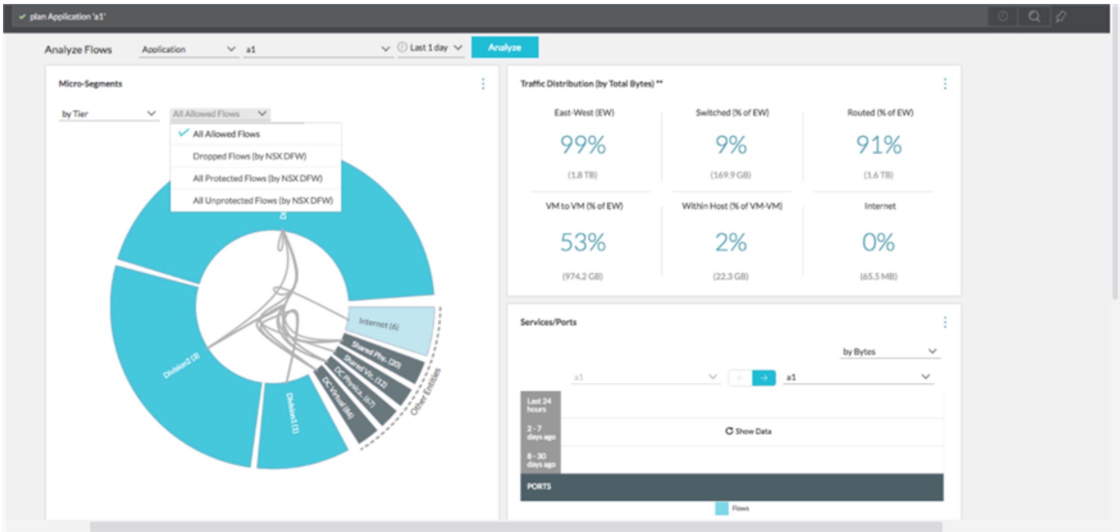

Network Address Translation (NAT)

vRealize Network Insight lists both SNAT and DNAT rules that are configured on the VMware NSX® Edge. The `NAT Rules` query lists all the SNAT and DNAT rules.

The VM to VM path also includes and shows the Edge NAT gateways configured in the path. Only the NAT rules configured on the Uplink interface of the VMware NSX® Edge are processed by the VM to VM path. The nested NAT hierarchy is also supported.

NAT Support

The NAT flow support in vRealize Network Insight is as follows:

- Only NSX-based edges are supported.
- Only edges with defined uplinks are supported.
- Only edges with NAT-defined uplinks are supported.
- The flow of the following NAT domains are reported:
  - Default domain
  - The single child domain of the default NAT domain

NAT Flow Support - Examples

This section consists of few examples for the supported NAT flow in vRealize Network Insight.

Example 1

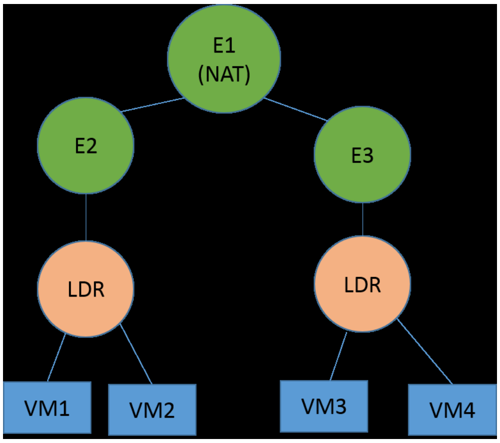

In the above topology, E2, E3, LDRs, VMs (VM1, VM2, VM3, VM4) are part of NAT domain E1. Anything above E1 such as uplink of E1 is part of default NAT domain. The above topology consists of the following:

The flow from VM1 to VM2 and vice versa is reported in vRealize Network Insight. Similarly the flow from VM3 to VM4 and vice versa is reported.

Example 2

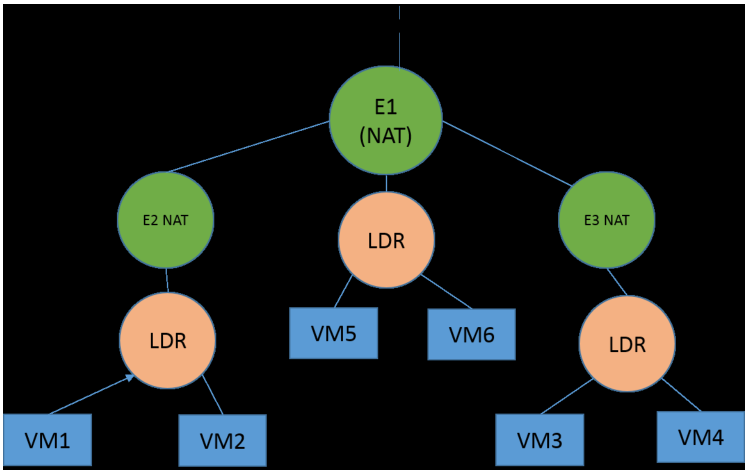

The above topology consists of the following:

- VM1 and VM2 are part of E2 domain.
- VM3 and VM4 are part of E2 domain.
- E2 and E3 NAT domains are child domains of E1 NAT domain.
- E1 is the single child of default NAT domain.
- VM5 and VM6 are part of E1 NAT domain.

In the above topology, the following flows are reported in vRealize Network Insight:

- Flow from VM5 to VM6

- Flow from (VM1, VM2) to (VM3, VM4)

VPC Flow Logs

You can create a flow log for VPC. For more information on creating flow logs at the VPC level, refer the AWS documentation at https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/flow-logs.html#create-flow-log. Once the flow logs are created, you can enable the VPC flow logs in vRealize Network Insight after validating your AWS account.

Recommended Policy for Flow Log Creation

Review your users' IAM policies to ensure that the `iam:PassRole` permission on the IAM role used in the `CreateFlowLogs` call is granted appropriately.

If the user who calls `CreateFlowLogs` does not have the IAM `PassRole` permission, the system returns an access denied message.

The following example demonstrates how you can verify if the correct permission is already assigned to users that have created the VPC Flow Logs in the past 90-days. If you have any additional questions or concerns, contact the AWS Support team.

Here is the sequence of steps to be run using the AWS CLI to determine if users who created the VPC Flow Logs have the correct permissions assigned to them:

1 Check your CloudTrail logs for events related to creating the flow logs, by searching for attributes with a key of `EventName` and a value of `CreateFlowLogs`. In this example, only one `CreateFlowLogs` event is found, and this command was invoked by the user `admin-temp`.

```
% aws cloudtrail lookup-events --lookup-attributes AttributeKey=EventName, \
       AttributeValue=CreateFlowLogs
 {
     "Events": [
         {
             "EventName": "CreateFlowLogs",
             "Resources": [
                 {
                     "ResourceType": "AWS::IAM::Role",
                     "ResourceName": "arn:aws:iam::123456789012:role/flowlogsRole"
                 },
                 {
                     "ResourceType": "AWS::Logs::LogGroup",
                     "ResourceName": "example-flow-logs"
                 },
                 {
                     "ResourceType": "AWS::EC2::FlowLog",
                     "ResourceName": "fl-1a1a1a1a"
                 },
                 {
                     "ResourceName": "vpc-2b2b2b2b"
                 }
```

```
            ],
            "EventId": "1a1a1a1a-ffff-1111-9999-1234567890af",
            "EventTime": 1514764800.0,
            "Username": "admin-temp",
```

2 Audit the permissions assigned to the IAM user `admin-temp`. Specifically, look for the `PolicyNames` assigned to this IAM user. In this example, the policy name assigned is `inline-pass-role-policy`. Using the CLI, review the details of this user policy. Look for the `iam:PassRole` permission. In this example, the user policy does include the `iam:PassRole` permission.

```
% aws iam list-user-policies --user-name admin-temp
{
    "PolicyNames": [
        "inline-pass-role-policy"
    ]
}

% aws iam get-user-policy --user-name admin-temp --policy-name inline-pass-role-policy
{
    "UserName": "admin-temp",
    "PolicyName": "inline-pass-role-policy",
    "PolicyDocument": {
        "Version": "2012-10-17",
        "Statement": [
            {
                "Sid": "VisualEditor0",
                "Action": "iam:PassRole",
                "Effect": "Allow",
                "Resource": "arn:aws:iam::123456789012:role/flowLogsRole"
            }
        ]
    }
}
```

3 Using the credentials of the `admin-temp` IAM user, you can create a new flow log and verify that the flow log is created successfully.

```
% aws ec2 create-flow-logs --deliver-logs-permission-arn \
         arn:aws:iam::123456789012:role/flowlogsRole \
         --log-group-name example-flow-logs --resource-ids vpc-2b2b2b2b \
         --resource-type VPC --traffic-type ALL
```

As this user has the correct permission `iam:PassRole` on the role `arn:aws:iam::123456789012:role/flowLogsRole` it uses in the `CreateFlowLogs` call, this call succeeds.

Refer the AWS Documentation on `IAM Roles for Flow Logs` for more information.

Micro-Segmentation

11

vRealize Network Insight provides planning and recommendations for implementing the micro-segmentation security. It helps the user to manage and scale the VMware NSX deployments quickly and confidently.

This chapter includes the following topics:

- Micro-Segmentation Dashboard
- Micro-Segmentation Planning Topology
- Application-Centric Micro-Segmentation
- Recommended Firewall Rules

Micro-Segmentation Dashboard

You can analyze the flows by selecting scope and segment them accordingly based on entities such as VLAN/VXLAN, Security Groups, Application, Tier, Folder, Subnet, Cluster, virtual machine (VM), Port, Security Tag, Security Group, and IPSet. The micro-segmentation dashboard provides the analysis details with the topology diagram. This dashboard consists of the following sections:

- Micro-Segments: This widget provides the diagram for topology planning. You can select the type of group and flows. Based on your inputs, you can view the corresponding topology planning diagram.
- Traffic Distribution: This widget provides the details of the traffic distribution in bytes.
- Top Ports by Bytes: This widget lists the top 100 ports that record the highest traffic. The metrics for the flow count and the flow volume are provided. You can view the flows for a particular port by clicking the count of flows corresponding to that port.

To access the micro-segmentation dashboard:

Procedure

1 On the navigation panel on the left side of the home page, click **Security>Plan Security**.

2 Select the scope, subscope, and the duration for which you want to plan and analyze. Click **Analyze**.

3 The micro-segmentation dashboard appears.

Micro-Segmentation Planning Topology

The micro-segmentation planning topology shows all the flows that are present in your environment by dividing the flows into segments.

In vRealize Network Insight , a flow is a 4-tuple. It includes:

- Source IP
- Destination IP
- Destination port
- Protocol

The blue lines denote the outgoing flows, the green lines denote the incoming flows, and the yellow lines denote the flows that are bidirectional. You can click any of the segments to view its details.

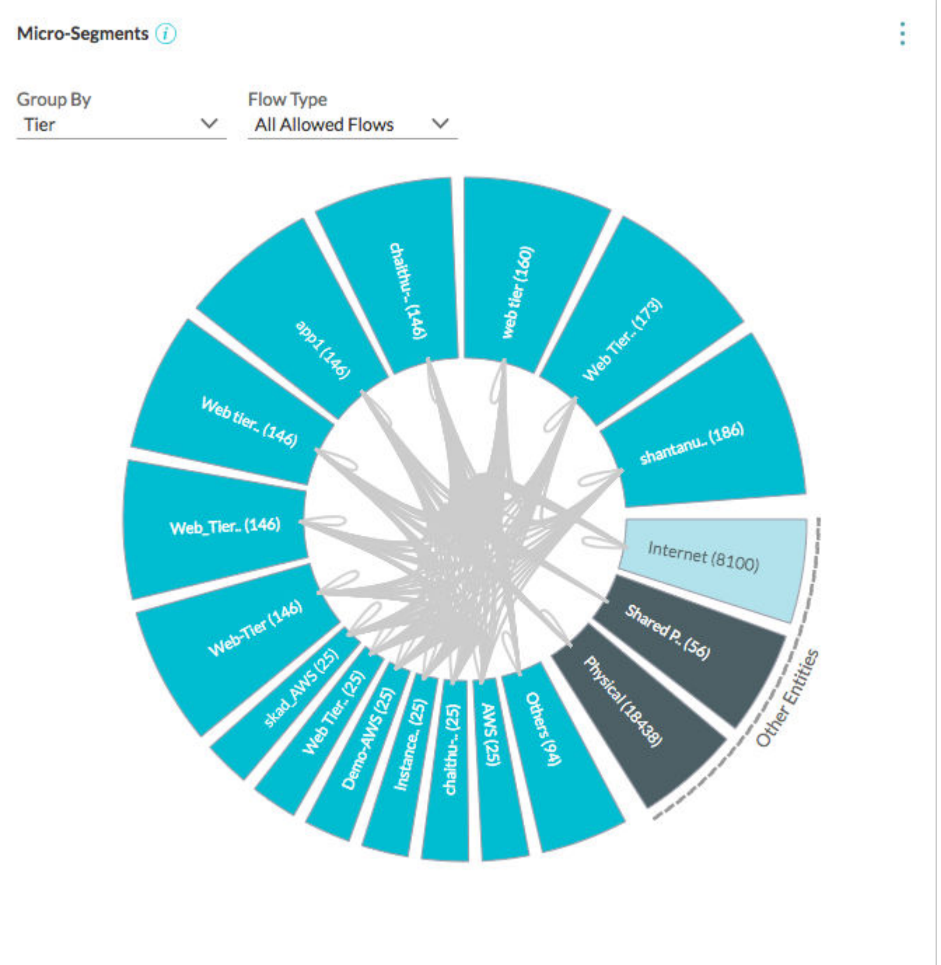

Micro-Segments

Group By: Tier

Flow Type: All Allowed Flows

The VMs that are outside the selected scope are grouped as **Other Entities** in the micro-segmentation planning topology.

You can also analyze the flows by creating subgroups as per Physical, Other Virtual, and Internet categories.

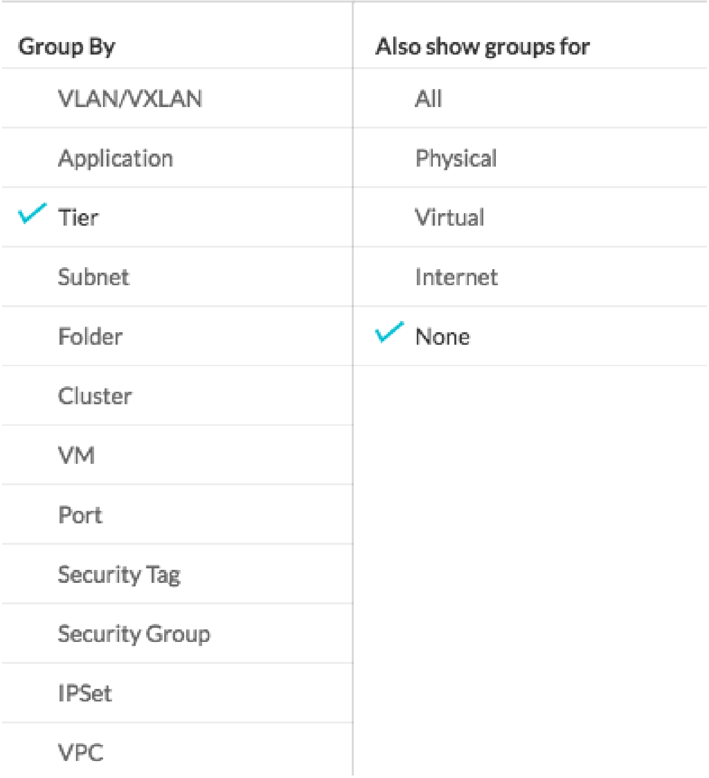

Group By	Also show groups for
VLAN/VXLAN	All
Application	Physical
✓ Tier	Virtual
Subnet	Internet
Folder	✓ None
Cluster	
VM	
Port	
Security Tag	
Security Group	
IPSet	
VPC	

Each group is expanded into a wedge. In the following topology, the wedge for **Physical group** is seen.

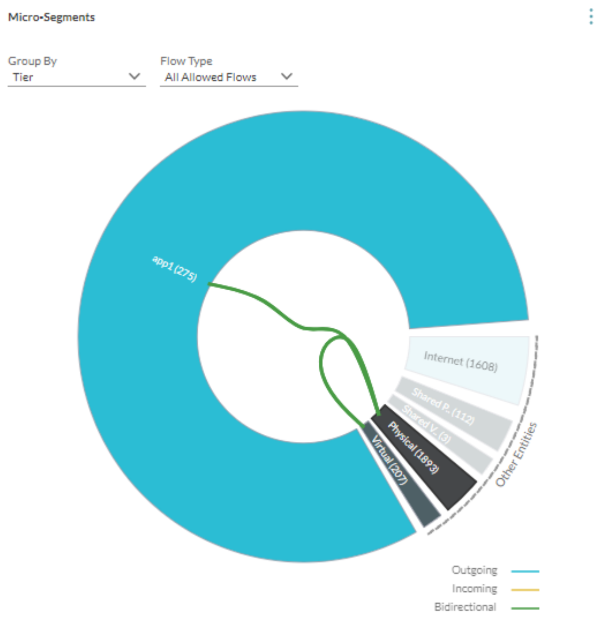

Micro-Segments

Group By: Tier

Flow Type: All Allowed Flows

The Flows pin shows that the flows for different time intervals segregated by ports. You can either view all the flows or view the flows between two entities. You can filter the flows by Allowed and Blocked flows. You can view flows by either Total Bytes or by Allowed Session Count. For the flows that are protected by a firewall, a Protected by Firewall sign is used to denote that the flows in that port that are protected by a firewall.

The planning for a scope such as an entire data center or a cluster selects flows that have VMs or Physical Servers (identified by the Physical IPs) as the source or the destination.

A topology has two distinct zones:

- Internal: This zone includes the VMs or the IP addresses in the scope.
- External: This zone includes the VMs or the IP addresses that are out of scope but talk to the VM or IP addresses in the internal zone. The external zone consists of the following wedges:
  - DC Virtual: It includes the source or the destination data center internal VMs that are talking to VMs or IP addresses in the internal zone and are not hosting any well-known shared services such as LDAP, NTP, and so on.
  - Shared Virtual: It includes the destination data center internal VMs hosting well-known shared services such as LDAP, NTP, and so on to which the VMs or IP addresses in the internal zone are talking.
  - DC Physical: It includes the source or the destination data center internal physical IP addresses that are talking to VMs or IP addresses in the internal zone and are not hosting any well-known shared services like LDAP, NTP, and so on.
  - Shared Physical: It includes the destination data center internal Physical IP addresses hosting well-known shared services such as LDAP, NTP, and so on to which the VMs or IP addresses in the internal zone are talking.
  - Internet: It includes the source or the destination data center external VMs or the physical IP addresses that are talking to the VMs or IP addresses in the internal zone.

Note

- Data center Internal implies RFC 1918 designated IPs by default + any overrides defined in E-W settings.
- Data center External implies non-RFC 1918 designated IPs by default + any overrides defined in N-S settings.

Application-Centric Micro-Segmentation

An application is a collection of tiers. Each tier in an application is a collection of VMs based on the user-defined filter criteria. The applications allow you to create a hierarchical group of VMs and visualize traffic/flows between the tiers of the same application. The traffic/flows can be visualized between applications.

To add an application:

Procedure

1. In the Search box, type application, and press Enter.
2. Click **Add Application**.
3. On the **Add Application** page, in the Application Name box, type a name for the application, which you want to create.
4. In the Tier section, type a name of the tier, which you want to create under Application (parent level). You can create a tier for VMs or physical machines as per requirements.
5. In the Virtual Machines/IP Addresses box, select the appropriate VMs by any of the following conditions:

   VM PROPERTIES

   a. VM Names - Name of the VMs, which you want to group in the tier you are creating
   b. IP Addresses - IP Addresses of the VMs or physical machines, which you want to group in the tier you are creating. The count of the IP addresses is shown at the right side of the text box.
   c. VMs with Service Ports - Service ports of the VMs, which you want to group in the tier you are creating
   d. Custom Search - It is an open search

   VMs IN

   a. Application - Select this option if the VMs are located in any previously created application
   b. Cluster - Select this option if the VMs are located in any cluster
   c. Folders - Select this option if the VMs are located in any folder
   d. VXLAN - Select this option if the VMs are located in any VXLAN
   e. VLAN - Select this option if the VMs are located in any VLAN

   Note For entering multiple values, set apart the individual values by comma.

   Optional: In case, you want to create multiple tiers under one application, click **Add Tier**.
6. Select Analyze Flows to view the flows before you finally add the application. You are able to see the tiers based on VMs or physical addresses accordingly.
7. Click **Save** to create the application.

Plan Applications

While creating an application, you can select **Custom IP Search** from the drop-down list to create tiers for the physical IPs based on the enriched fields. For more information on the enriched fields, refer Enriching Flows and IP Endpoints.

The enriched DNS, Subnet, VLAN information can be used in specifying tiers as follows :

- Web

  `Query: IP Endpoint where Subnet Network = '172.16.101.0/24'`

- App

  `Query: IP Endpoint where Dns Domain = app.example.com`

- DB

  `Query: IP Endpoint where L2 Network = 'vlan-102'`

- Common Services

  `Query: IP Endpoint where Dns Domain = svc.example.com`

Recommended Firewall Rules

On the **Plan security** page, when you click the wedge or the edge in the topology diagram, you can view the list of the services and flows for that particular segment. Click **Recommended Firewall Rules** to view the rules defined on it. The members of the source or the destination are listed under the following types of rules:

- Physical to Physical: This tab lists all the rules associated with the physical and Internet IPs. The rules can be for physical-physical, physical-Internet, Internet-physical, or Internet-Internet entities.
- Virtual: This tab lists all rules where at least one of the endpoints is a VM.

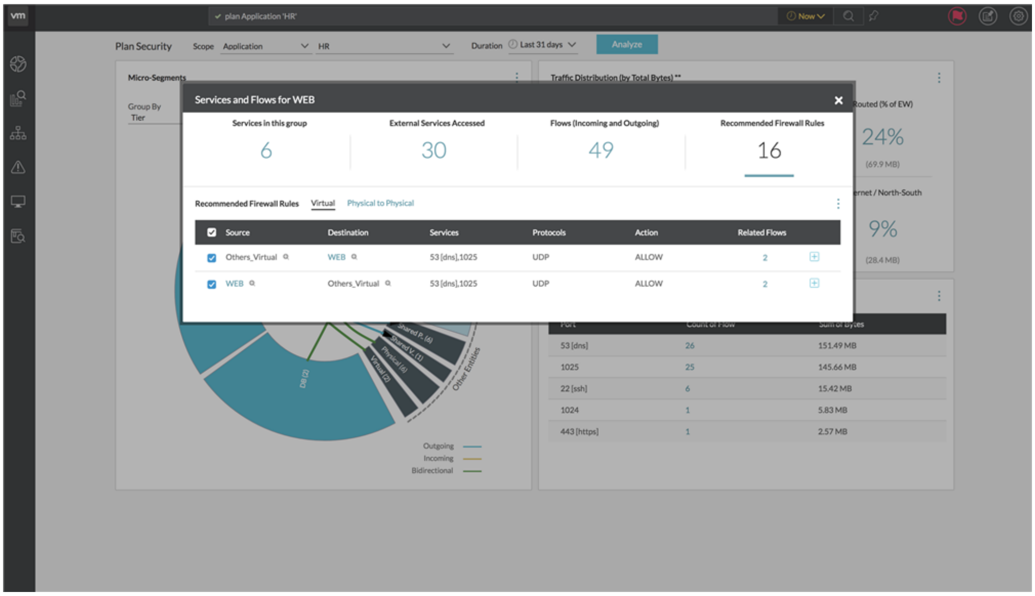

Plan Security Scope: Application HR Duration: Last 31 days Analyze

Services and Flows for WEB

Services in this group: 6
External Services Accessed: 30
Flows (Incoming and Outgoing): 49
Recommended Firewall Rules: 16

Recommended Firewall Rules Virtual Physical to Physical

☑	Source	Destination	Services	Protocols	Action	Related Flows
☑	Others_Virtual	WEB	53 [dns],1025	UDP	ALLOW	2
☑	WEB	Others_Virtual	53 [dns],1025	UDP	ALLOW	2

Port	Count of Flow	Sum of Bytes
53 [dns]	26	151.49 MB
1025	25	145.66 MB
22 [ssh]	6	15.42 MB
1024	1	5.83 MB
443 [https]	1	2.57 MB

For each firewall rule, the following details are available:

- Show members of the group: Click the + sign next to the name of the entity to see the members of the group.

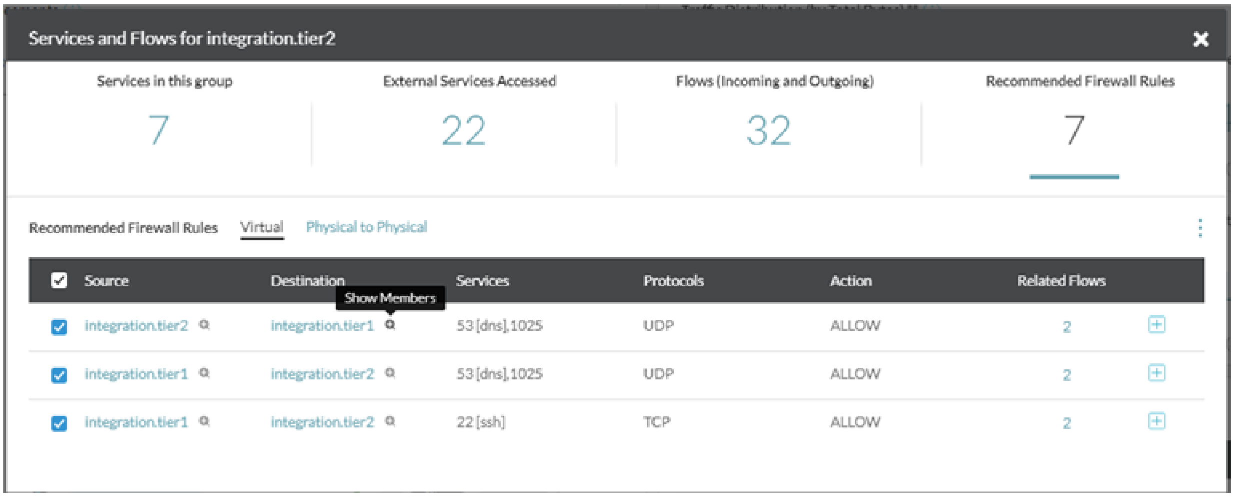

Services and Flows for integration.tier2

Services in this group	External Services Accessed	Flows (Incoming and Outgoing)	Recommended Firewall Rules
7	22	32	7

Recommended Firewall Rules Virtual Physical to Physical

☑	Source	Destination	Services	Protocols	Action	Related Flows	
☑	integration.tier2	integration.tier1	53 [dns],1025	UDP	ALLOW	2	⊞
☑	integration.tier1	integration.tier2	53 [dns],1025	UDP	ALLOW	2	⊞
☑	integration.tier1	integration.tier2	22 [ssh]	TCP	ALLOW	2	⊞

Note

- The members are not shown for the groups belonging to the Internet category.
- If a security group has both virtual and physical IPs, the physical and the Internet IPs are not shown in the list of the members of that particular group.

- Source
- Destination
- Services
- Protocols
- Action
- Related Flows: Click the number of the related flows to see the list of flows with the corresponding flow information.
- View Applied Firewall Rules: Click the + sign next to the **Related Flows** column to view the applied firewall rules corresponding to the similar sets of flows.

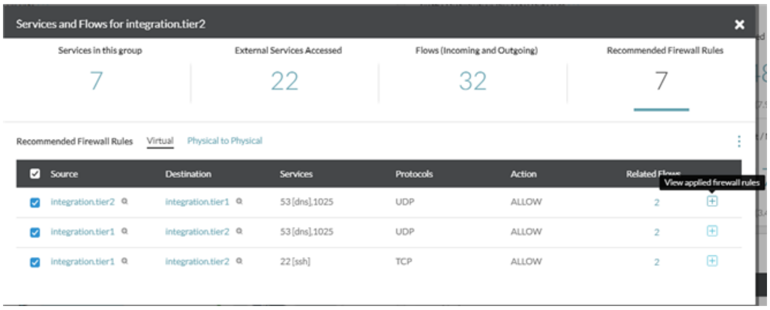

Services and Flows for integration.tier2

Services in this group	External Services Accessed	Flows (Incoming and Outgoing)	Recommended Firewall Rules
7	22	32	7

Recommended Firewall Rules Virtual Physical to Physical

☑	Source	Destination	Services	Protocols	Action	Related Flows	
☑	integration.tier2	integration.tier1	53 [dns],1025	UDP	ALLOW	2	⊞
☑	integration.tier1	integration.tier2	53 [dns],1025	UDP	ALLOW	2	⊞
☑	integration.tier1	integration.tier2	22 [ssh]	TCP	ALLOW	2	⊞

You can export the recommended rules as XML or CSV based on your requirement. Refer to Exporting Rules for more information on these artifacts.

Exporting Rules

You can export all the rules as XML for the entire topology. You can find this menu item in the **Micro-Segmentation Planning** page as follows:

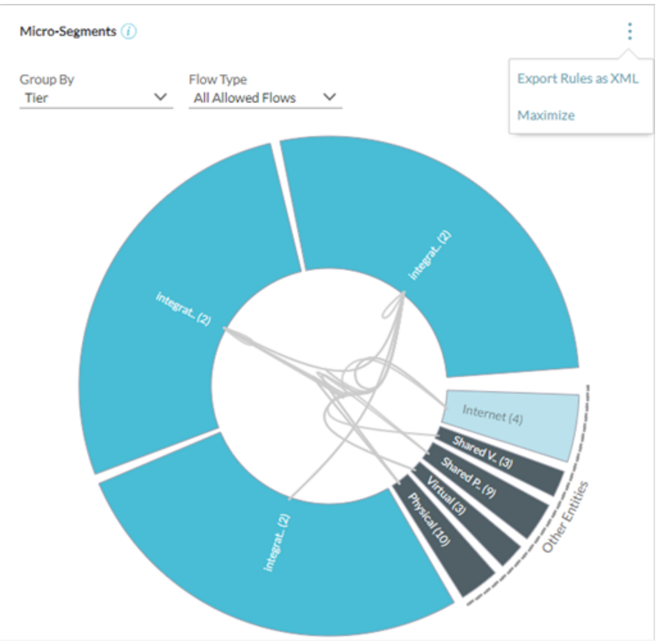

The Export as XML option is available only for the following entities:

- Security Group
- Application Tier

If the planning scope spans a single NSX Manager only, the generated artifacts contain the XML files corresponding to the recommended services and firewall rules. If the planning scope spans multiple NSX managers, the generated artifacts contain the XML files corresponding to the recommended services, IPsets, security groups, and the firewall rules.

The following are the placeholder artifacts for security groups:

- `SG-Others_Internet.xml`
- `SG-Other.xml`

You can export all the rules as XML or CSV for a particular wedge or edge depicted in the topology diagram.

Search

12

vRealize Network Insight provides a robust search for all the entities in your environment. When you search for entities, the software displays the entities that match your search query on the **Results** page.

The search-bar uses natural language to search through various aspects of your SDDC. For each search query, the search bar suggests you the next term that you can use to narrow down your search results. For example, when you enter the term **vm**, the search bar displays a possible list of terms that you can add to your existing term to narrow down your search results. The search bar also validates each search query. A check mark denotes a valid search query and a cross mark denotes an invalid search query. The **Help** page provides examples of currently supported queries.

This chapter includes the following topics:

- Simple Search
- Advanced Search
- Time Control
- Search Results
- Filters

Simple Search

The features of simple search are as follows:

- Entity Types
  - Example: VM, Host, VLAN, VXLAN, and so on
  - Some pre-defined types can be used to search for multiple related entity types.
    - Example: L2 Network represents VLAN, VXLAN, and Native L2 networks.
  - Auto-complete can be used to explore entity types and other prop
- Event Types
  - Example: MTU Mismatch Event, Membership Change Event, and so on

  - Problems or Changes can be used as keywords to search for all problems or change events

VMs	Show all the VMs
VM 'vm1'	Show VM with name 'vm1'
L2 Networks	Show all L2 networks (VLAN, VXLAN, Native)
Problems	Show all problem events
MTU Mismatch Events	Show all MTU Mismatch Events

- Configuration properties
  - To find entities matching a configuration property value
    - Example: <Name>, <IP Address>, <IP Address 1 - IP Address 2>, and so on
    - Different types of properties: String, Numeric, IP, Range, Reference, and so on
    - Reference is used to represent data center as a graph. So VM has a reference to Host, Cluster and so on
- Metric properties
  - To find all entities with an applicable metric
    - Example: CPU Usage Rate, Network Usage Rate, and so on
- Planning
  - This can be used to plan the security of the data center by analyzing the flows
  - Example: Plan Security Group 'SG_All', Plan Host 'Host-1', and so on
  - "Plan Security" can be used to plan security of the entire data center.
- Path
  - This can be used to show the path between two VMs or the path from VM to Internet
    - VM ‘dev1’ to VM ‘db1’
    - VM ‘dev2’ to Internet

Note Auto-complete can be used to explore more entity types and properties.

Advanced Search

The features of advanced search are:

- Filtering
  - Search results can be filtered using the properties that they have.

    Example:
    - VMs where IP Address = 192.168.0.1

    - VMs where CPU Usage Rate > 90%
    - VMs where host = ‘host1’
  - Filters can be combined using the following logical operators:
    - and
    - or
    - not

    Example:

    - VMs where IP Address = 192.168.0.0/16 and Network Rate > 1 Mbps
    - VMs where IP Address = CPU Usage Rate > 90% or Network Rate > 1 Mbps
- Projection
  - Get properties or metrics.
    - IP Address of VMs
    - CPU Usage Rate, CPU Count of VMs
  - 1 Aggregation (SUM, AVG, MAX, MIN) can be used for numeric properties and metrics.
      - AVG(CPU Usage Rate) of VMs
- Sorting
  - Results can be sorted using the **order by** clause.
    - Order: asc or desc (optional)

    Example: VMs order by CPU Usage Rate
  - Limit the # of results

    Example: top 10 VMs order by CPU Usage Rate
- Group By
  - Search results can be grouped by a given property into buckets. By default, groups by results are ordered by count of entities in each bucket.

    Example: VMs group by Host

    This property returns list of hosts with # of VMs on each host
  - Group by results can be sorted by applying aggregation on numeric properties and metrics

SUM(Bytes) of Flows group by Port order by SUM(Bytes)	Returns list of ports ordered by sum of total bytes for all the flows on that port
SUM(CPU Count) of VMs group by Host order by SUM(CPU Count)	Returns list of hosts ordered by sum of vCPUs of all VMs on every host

Time Control

Time-control allows you to run a search query within the context of a selected time or time range. You can select from a list of presets such as last 24 hours, last 3 days, and so on. You can also specify a particular date and time using the **At** option or even a range using the **Between** option.

Search Results

The search results page provides a detailed list of concerned entities that match a particular search. The page itself provides numerous information that ranges from the list of entities, their corresponding properties, and facets to filter the search results to refine your search.

You can also expand or collapse each entry in the search results to view more information about a particular entry. You can also create a notification for each search.

Note You can point to a particular property in the search results and also in the entity pages to view a tool tip containing more information about that property.

The following graphic shows the search results for the VXLANs where num vms > 0 search query for a time from the past.

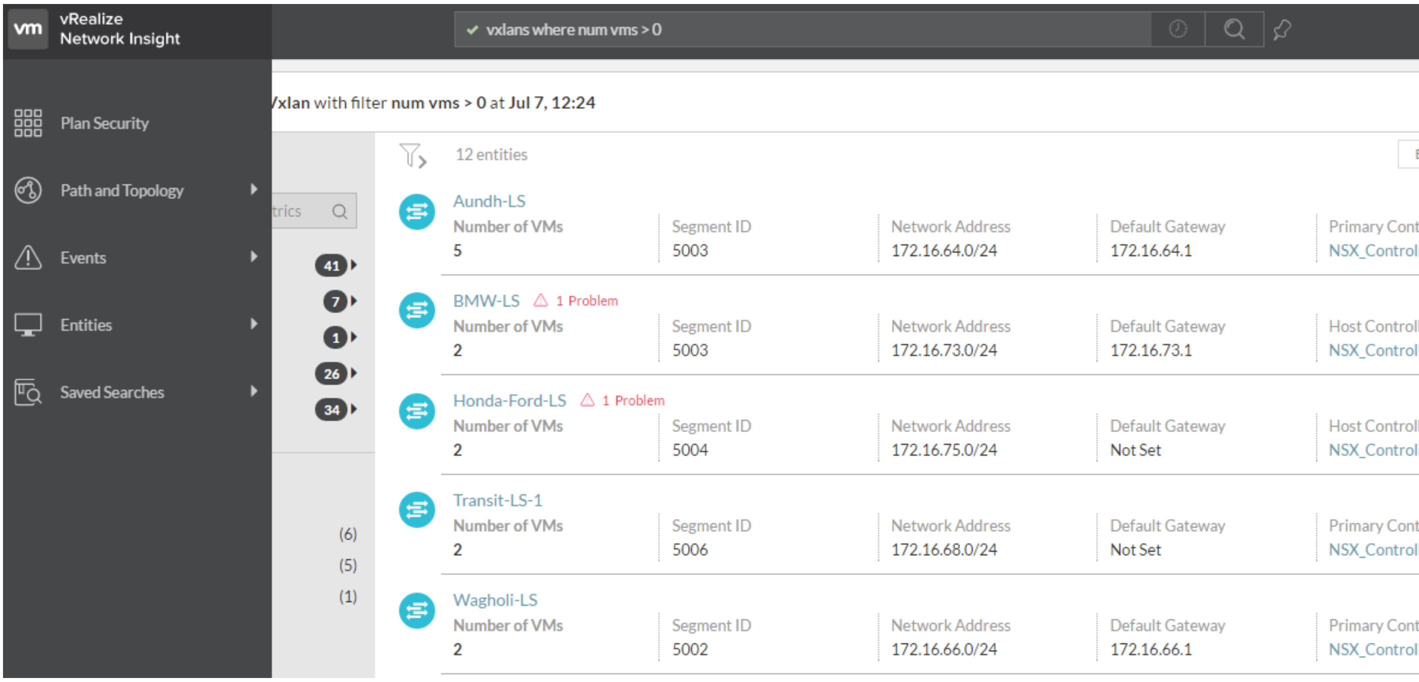

vxlans where num vms > 0

Vxlan with filter num vms > 0 at Jul 7, 12:24

12 entities

	Number of VMs	Segment ID	Network Address	Default Gateway	
Aundh-LS	5	5003	172.16.64.0/24	172.16.64.1	Primary Cont NSX_Control
BMW-LS ⚠ 1 Problem	2	5003	172.16.73.0/24	172.16.73.1	Host Controll NSX_Control
Honda-Ford-LS ⚠ 1 Problem	2	5004	172.16.75.0/24	Not Set	Host Controll NSX_Control
Transit-LS-1	2	5006	172.16.68.0/24	Not Set	Primary Cont NSX_Control
Wagholi-LS	2	5002	172.16.66.0/24	172.16.66.1	Primary Cont NSX_Control

Filters

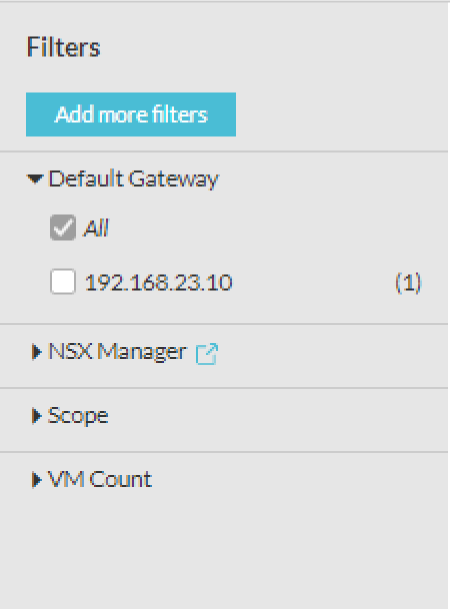

Once you get the search results, click Add more filters on the left pane as per your requirements. You can view a series of filter categories that you can use to narrow down the search results. The number of available filters for each category is mentioned in a small box beside the category. View the available filters for that category (along with a short explanation for each filter) and click to apply that filter. You can also use the filter search box to search for a particular filter and vRealize Network Insight automatically shows the filters that match your search query and you can click to apply that filter. Each filter has several properties to refine the search results. When you select a filter property from one of the filters, then the selected property is highlighted in the search results.

Upgrade

13

This chapter includes the following topics:

- Online Upgrade of Product
- Single-Click Offline Upgrade

Online Upgrade of Product

The **Update** option is available in the **Overview** section on the **Install and Support** page.

The **Update** option lets you know if the latest version of the product is available for an upgrade. A notification message appears in the product, and you can opt to upgrade to the latest version from the UI itself. To upgrade to the latest version:

1. In case a latest version is available, a message appears on the upper-right corner of the browser window.
2. Click **View details** in the notification.
3. You can view the new features, which are available in the new version.
4. Click **Install now** to start the upgrade.

Alternatively,

1. If a newer version is available, the information is displayed in the **Overview** section at the **Update** option.
2. Click **View Details**, to view the new features, which are available in the new version.
3. Click **Install now** to start the upgrade.

Single-Click Offline Upgrade

vRealize Network Insight supports the single-click offline upgrade of the product from Release 3.7 to the future releases such as 3.7->3.8, 3.7->3.9, and so on.

Procedure

1 On the **Install and Support** page, under **Software version**, click **Click here**.

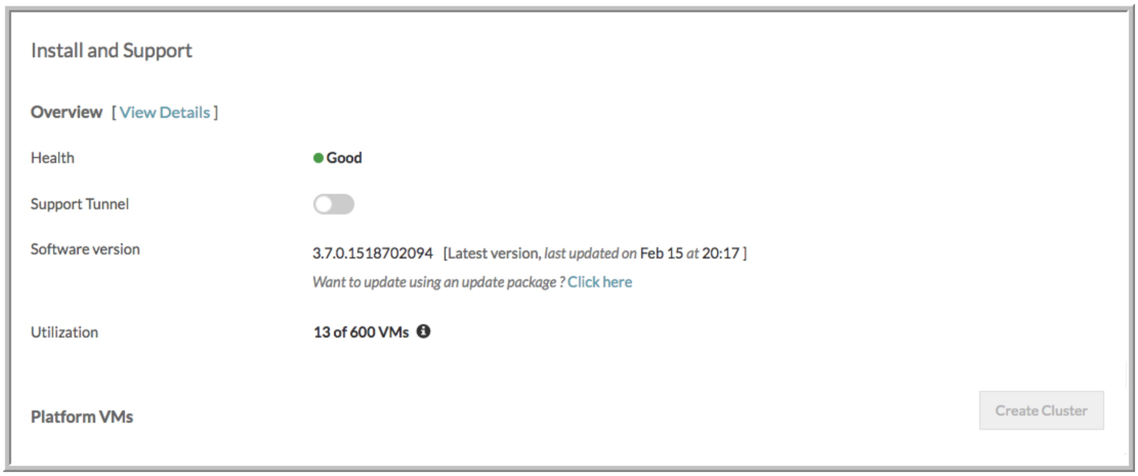

Install and Support

Overview [View Details]

Health ● Good

Support Tunnel

Software version 3.7.0.1518702094 [Latest version, *last updated on* Feb 15 *at* 20:17]
Want to update using an update package ? Click here

Utilization 13 of 600 VMs

Platform VMs

2 Ensure that you have saved the update package to your local disk so that you can upload it on this page. Click **Browse** to select the file and click **Upload**.

3 After the upload is complete, a dialog box with the basic upgrade instructions appears before the upgrade begins. To proceed further, click **Install Now**.

4 Once the upgrade process begins, a pop-up window appears. This pop-up window provides the status of each node.

Note

- Ensure that all the nodes are online before beginning the upgrade. If any node is inactive before the upgrade begins, you will not be allowed to upload the upgrade bundle.
- Once the upgrade begins, if a node becomes inactive, the upgrade process hangs. The upgrade will not resume until the node becomes active again.
- Until the upload of the package happens, the user should take care that the session is not closed. If the session ends, the user has to restart the upload process.
- The Platform 1 becomes the upgrade server here. If Platform1 is offline, then no other node is upgraded.

5 Upon the completion of upgrade process, all platforms and the collectors nodes are upgraded.

Troubleshooting

14

This chapter includes the following topics:

- Health
- Support Tunnel
- Online Upgrade of Product
- Single-Click Offline Upgrade
- View Node Details
- Capacity
- Disk Management
- Creating Clusters
- Expanding Clusters
- Creating Support Bundle
- Migrating Data Sources
- About Page
- Common Data Source Errors

Health

The **Health** indicator is available in the **Overview** section on the **Install and Support** page.

The **Health** indicator turns red if any of the following malfunctioning events occur:

- If proxy stops collecting flow data
- If platform stops processing data due to some reason; for example, insufficient disk space
- If search indexer lags behind, resulting in outdated search result

The overall health indicator displays the number of irregularities, with a Red light on. The individual irregularities are listed with their details, when the number of problems against overall health, is clicked on. In case of normal functioning, the health indicator shines a Green light.

Support Tunnel

The **Support Tunnel** option is available in the **Overview** section on the **Install and Support** page.

The support tunnel allows the vRealize Network Insight engineering team to remotely connect to customer's platform and collector VMs on the SSL secured connection. You have to request the access to support tunnel when the vRNI engineering team needs to access the setup for advanced troubleshooting or debugging.

Note Ensure that the traffic to `support2.ni.vmware.com` on port `443` is allowed.

Online Upgrade of Product

The **Update** option is available in the **Overview** section on the **Install and Support** page.

The **Update** option lets you know if the latest version of the product is available for an upgrade. A notification message appears in the product, and you can opt to upgrade to the latest version from the UI itself. To upgrade to the latest version:

1. In case a latest version is available, a message appears on the upper-right corner of the browser window.
2. Click **View details** in the notification.
3. You can view the new features, which are available in the new version.
4. Click **Install now** to start the upgrade.

Alternatively,

1. If a newer version is available, the information is displayed in the **Overview** section at the **Update** option.
2. Click **View Details**, to view the new features, which are available in the new version.
3. Click **Install now** to start the upgrade.

Single-Click Offline Upgrade

vRealize Network Insight supports the single-click offline upgrade of the product from Release 3.7 to the future releases such as 3.7->3.8, 3.7->3.9, and so on.

Procedure

1 On the **Install and Support** page, under **Software version**, click **Click here**.

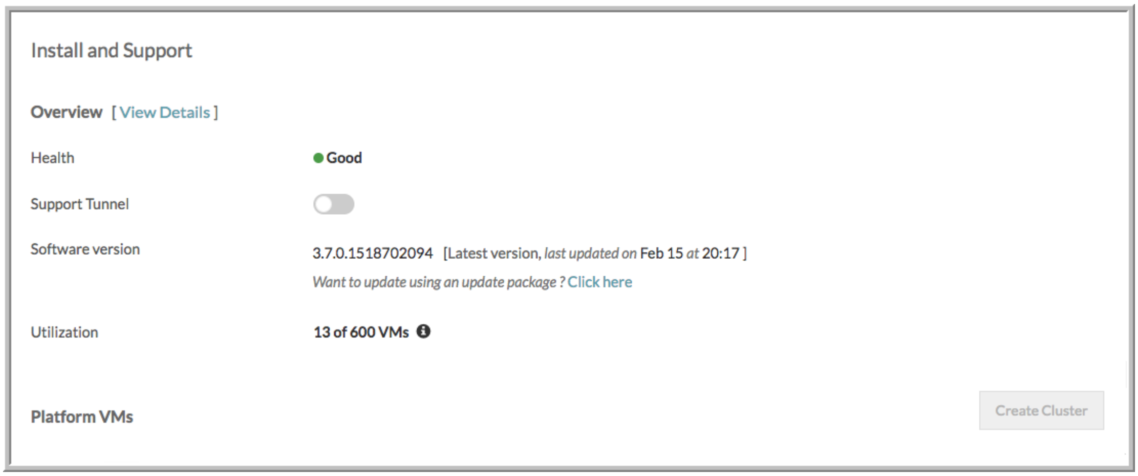

Install and Support

Overview [View Details]

Health ● Good

Support Tunnel

Software version 3.7.0.1518702094 [Latest version, *last updated on* Feb 15 *at* 20:17]

Want to update using an update package ? Click here

Utilization 13 of 600 VMs

Platform VMs

2 Ensure that you have saved the update package to your local disk so that you can upload it on this page. Click **Browse** to select the file and click **Upload**.

3 After the upload is complete, a dialog box with the basic upgrade instructions appears before the upgrade begins. To proceed further, click **Install Now**.

4 Once the upgrade process begins, a pop-up window appears. This pop-up window provides the status of each node.

Note

- Ensure that all the nodes are online before beginning the upgrade. If any node is inactive before the upgrade begins, you will not be allowed to upload the upgrade bundle.
- Once the upgrade begins, if a node becomes inactive, the upgrade process hangs. The upgrade will not resume until the node becomes active again.
- Until the upload of the package happens, the user should take care that the session is not closed. If the session ends, the user has to restart the upload process.
- The Platform 1 becomes the upgrade server here. If Platform1 is offline, then no other node is upgraded.

5 Upon the completion of upgrade process, all platforms and the collectors nodes are upgraded.

View Node Details

You can view the details of each node in a platform or a collector.

Procedure

1 To view the details of a particular platform node, click its name that is listed under **Platform VMs** on the **Install and Support** Page.

   The NI Platform dashboard appears.

2 To view the details of a particular collector node, click its name that is listed under **Collector (Proxy) VMs** on the **Install and Support** page.

   The NI Collector dashboard appears.

Capacity

vRealize Network Insight provides the approximate capacity and load information of a collector node and a platform. This limits-based information helps you to prevent the performance and experience issues later.

Understanding Capacity

There are two kinds of capacity:

- VM capacity: It is defined as the number of discovered VMs that a node or a setup can handle.
- Flow capacity: It is defined as the number of flows that a node or a setup can handle.

The capacity is defined as follows:

- Single platform with one or more proxy nodes: The capacity of a proxy node or the platform is the number of discovered VMs that it can handle without the degradation of performance.
- Cluster setup: The capacity of the platform in a cluster setup is the aggregation of all the capacities of all the platform nodes while the capacity of proxy nodes is considered at the level of an individual node.

Accessing the Capacity Information

You can view **VM Capacity** and **Flow Capacity** on the **Install and Support** page.

For every collector node listed under Collector (Proxy) VMs, only the VM capacity information is provided.

Note When the number of discovered VMs from the data sources across the deployment exceed the capacity of either the system or the collector or both, you will not be allowed to trigger the upgrade.

To view the discovered VMs for a data source:

1 In the **Accounts and Data Sources** page, you can see the number of VMs that have been discovered for a particular data source which is already added and currently active. This column will have a value only if the data source is vCenter or AWS source.

Note The discovered VM count includes placeholder and template VMs. So it can be different from the count of VMs in the product.

Disk Management

If the disk utilization is high for a platform or a collector, an event is triggered to warn the user. Also, a recommendation of how much more disk space needs to be added is provided. You can view the event in the platform or the collector dashboard. The alert is also shown in the corresponding collector or the platform section in the **Install and Support** page.

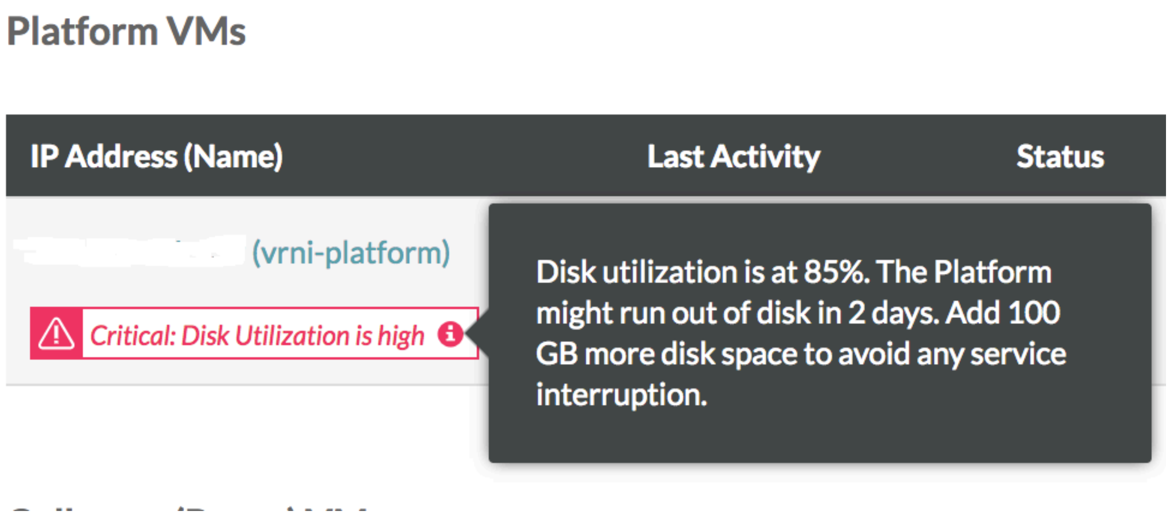

Platform VMs

IP Address (Name)	Last Activity	Status
[redacted] (vrni-platform) ⚠ Critical: Disk Utilization is high ℹ		

Disk utilization is at 85%. The Platform might run out of disk in 2 days. Add 100 GB more disk space to avoid any service interruption.

You can add disks to the nodes by performing the following steps:

Note Do not expand the existing hard disk.

Procedure

1. Log into vCenter through the Web client with sufficient privileges.
2. Right click the node and click **Edit settings**.
3. Add the hard disk as per the recommendation provided in the alert.
4. vRealize Network Insight takes few minutes to detect the appliance and add it to the `/var` partition.

Creating Clusters

You can create clusters from the **Install and Support** page.

Prerequisites

At least two additional platforms are required. The additional platform VMs should be deployed and powered on.

To create cluster

1. Click **Create Cluster** for **Platform VMs**.
2. On the **Create Cluster** page, enter the following information:
   - **IP Address**: Enter the IP address of the new platform that you want to add.
   - **Password**: Enter the support user password of the platform VM. If you have not changed the password yet, then refer the *Default Login Credentials* section in *vRealize Network Insight Installation Guide* for the password.
3. To keep adding more platforms, click **Add more** and enter the IP address and the support user password.
4. Click **Submit**. Click **Yes**.
5. After creating a cluster, the user needs to log in to the product again.

Note

- The **create cluster** option is enabled only when the platform is of large brick size. All platforms should be of large brick to create cluster.
- To receive telemetry data, ensure that you enable telemetry on all the platform nodes.
- To expand clusters, refer the *Expanding a Cluster* section in the *vRealize Network Insight Installation Guide*.

Expanding Clusters

Once the cluster is created, you can expand the cluster by adding more platform nodes to it.

To expand clusters:

Procedure

1. On the **Install and Support** page, click **Expand Cluster** for **Platform VMs**.
2. The IP addresses of the VMs that are part of the cluster already are listed on the Expand Cluster page. To add one or more nodes to the existing cluster, provide the IP address of the node and the support user password.

   Note

   - Currently, vRealize Network Insight supports 10 nodes in an existing cluster. Once the limit is reached, the **Add more** button is disabled.
   - Ensure that all the new nodes are non-provisioned and are reachable through SSH.
   - Ensure that you have taken a backup of the existing platform VMs before you go ahead with the cluster expansion.

3 Click **Submit**.

The step-by-step progress is displayed.

4 Once the cluster expansion link is completed, a message indicating success is displayed.

While the cluster expansion is in progress, the application cannot be used for any other operation.

Creating Support Bundle

To create support bundle:

1 In the **Platform VMs** or Proxy VMs table, in the **Support Bundle** column, click the **Create Support Bundle** icon .

Note Only two support bundles can be present at one given time, so while creating a new one, if there are already two support bundles present, the older one is deleted.

2 Click **Yes** to confirm creation of a new support bundle.

A new support bundle is created displaying data and time as download link. To initiate the download of support bundle, click the link.

Migrating Data Sources

If a proxy VM is down or deleted, you can add a new proxy VM and migrate data source from the old proxy VM to the new proxy VM.

To migrate a data source:

Procedure

1 In the **Install and Support** page, under the **Collector (Proxy) VMs** section, click the edit icon.

If a proxy VM is down, you can see the error message that proxy VM is not available under the same section.

2 In the **Edit Collector (Proxy) VM** page, you can assign a nickname to the proxy VM.

3 The Edit Collector (Proxy) page lists all the data sources added to the proxy. To migrate a data source, click **Migrate** for a particular data source.

4 The Edit account or source page appears. Ensure that you fill the following information:

Table 14-1.

Fields	Description
Collector (Proxy) VM	Name of the new proxy VM to which the data source has to be migrated
IP Address	Pre-filled IP/FQDN address of data source

Table 14-1. (Continued)

Fields	Description
Username	Username for the data source
Password	Password for the data source

5 Click **Validate**. Click **Submit**. The data source is then deleted in the old proxy VM and is added to the new proxy VM.

6 Once the migration is successful, you will see the new proxy VM against the data source in the **Enabled** column in the **Accounts and Data Sources** page.

Note

- If you are migrating vCenter to another proxy VM, then sure that you migrate the corresponding NSX Manager also to the same proxy VM.
- When you migrate NSX Manager to another proxy VM, the child data providers such as NSX Controller and NSX Edge are migrated as well to the new proxy VM.

About Page

This page displays the license details and the product version number that you are currently using.

Add License

vRealize Network Insight supports the addition of multiple licenses.

To add a license:

1 In the About page, click **Add License**.

2 Provide the license key for the **New License Key** field

3 Click **Validate**.

4 Click **Activate**.

5 You can see the list of licenses in the page.

6 You can also delete the license by clicking the delete icon next to the Expiration column. If the license belongs to an Enterprise edition and if it is the last remaining Enterprise edition in the system, then ensure that you have deleted the AWS account before you delete the Enterprise license.

Change License

In the event of expiry of evaluation license, when you log in to the product, a message appears stating that the license has expired and that you need to renew your license. Follow the steps below to change license.

To change license:

1 Click the link contained in the Expiry message to go to the Change License page. Alternatively, in **Settings**, click **About**, and then click **Change License**.

2 In the **Change License** page, in **New License Key**, enter the new license key you received from VMware.

3 Click **Validate**.

4 Click **Activate**.

Customer Experience Improvement Program

This product participates in VMware's Customer Experience Improvement Program (CEIP).

The details regarding the data collected through CEIP and the purposes for which it is used by VMware are set forth at the Trust & Assurance Center at

https://www.vmware.com/solutions/trustvmware/ceip.html

To join or leave the CEIP for this product:

1 In the **About** page, under Customer Experience Improvement Program, click **Modify**.

2 The CEIP window pops up. To join CEIP, check **Enable**. This action activates CEIP and sends data to https://vmware.com.

3 To leave CEIP, uncheck **Enable**.

4 Click **Submit**.

Common Data Source Errors

When you add a data source, you can come across several errors. This table contains the list of common errors with the cause and resolution for each.

Table 14-2.

Error Text	Cause	Resolution
Invalid Response from Data Source	vRealize Network Insight Proxy was unable to process the information received from the Data Source as the information was not in the expected format.	In some data providers this problem is observed intermittently and might go away in the next polling cycle. If it occurs consistently, contact support.
Data Source is not reachable from Proxy VM	Data source IP address on SSH/REST (port 22 or 443) is either not reachable from the vRealize Network Insight Proxy VM or the data source is not responding. This error occurs while adding the data source.	Verify connectivity to the data source from vRealize Network Insight Proxy VM on port 22 or 443. Make sure data source is up and running and the firewall is not blocking connection from vRealize Network Insight Proxy VM to the data source.
No NSX Controller found	An NSX Controller has been selected in the NSX Manager data source page but there is no NSX Controller installed.	Install an NSX Controller on NSX Manager and then select NSX Controller check box on the NSX Manager data source page.

Table 14-2. (Continued)

Error Text	Cause	Resolution
Data source type or version mismatch	Provided data source IP Address/FQDN is not of selected data source type.	Verify that provided data source IP Address/FQDN is of selected data source type and version is supported by vRealize Network Insight
Error connecting to data source	vRealize Network Insight Proxy VM is unable to connect to the data source. This error occurs after adding the data source.	Verify connectivity to the data source from vRealize Network Insight Proxy VM on port 22 or 443. Make sure that the data source is up and running and firewall is not blocking connection from vRealize Network Insight Proxy VM to the data source.
Not found	vRealize Network Insight Proxy VM is not found.	Check if pairing is done between vRealize Network Insight Proxy VM and vRealize Network Insight Platform VM.
Insufficient privileges to enable IPFix	The user who is trying to enable IPFIX in vCenter does not have the following privileges: DVSwitch.Modify; DVPortgroup.Modify	Provide adequate privileges to the user.
IP/FQDN is invalid	The IP/FQDN provided on the data source page is not valid or does not exist.	Provide valid IP/FQDN address.
No data being received	vRealize Network Insight Platform VM is not receiving data from vRealize Network Insight Proxy VM for that data source.	Contact Support.
Invalid credentials	Provided credentials are invalid.	Provide the correct credentials.
Connection string is invalid	The IP/FQDN provided on data source page is not in proper format	Provide valid IP/FQDN address.
Recent data may not be available, due to processing lag	vRealize Network Insight Platform VM is overloaded and lagging behind in processing data.	Contact support.
Request timed out, please try again	Could not complete request in specified time.	Try again. If the issue is not fixed, then contact support.
Failed for unknown reason, please retry or contact support	Request failed for some unknown reason.	Try again. If the issue is not fixed, then contact support.
Password authentication for SSH needs to be enabled on device	SSH login using password is disabled on the device added	Enable password authentication for SSH on the device being added for monitoring.
SNMP connection error	Error connecting to the SNMP port	Verify if SNMP is configured correctly on the target device.