

Using vRealize Network Insight

VMware vRealize Network Insight 4.0



vmware®

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

Copyright © 2019 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

1	About vRealize Network Insight User Guide	8
2	Getting Started	9
	Introduction	9
	Homepage	11
	Navigation	12
	Settings	13
3	Configuration	15
	Accounts and Data Sources	15
	Adding a Data Source	16
	Adding an AWS Data Source	17
	VMware Cloud on AWS	20
	Adding an Infoblox DNS Data Source	23
	Adding a NetFlow Collector	25
	Sampled Flows	25
	F5 BIG-IP	25
	NSX-T	27
	Cisco ACI	32
	Adding vRealize Log Insight as a Data Source	34
	Edge Data Collection	36
	Data Management	37
	IP Properties and Subnets	38
	Physical IP and DNS Mapping	38
	Physical Subnets and VLANs	38
	East-West IPs	39
	North-South IPs	39
	Working with Events	40
	Working with System Events	40
	Working with User-Defined Events	42
	Working with Platform Health Events	42
	User Management	43
	Add New User	43
	Assign Administrator Role	43
	Syslog Configuration	44
	LDAP	44
	Configuring Mail Server	46
	Support for Simple Network Management Protocol (SNMP)	47

[Define Timeout Interval for a User](#) 48

4 Capacity 49

[Planning to Scale up the Platform Cluster](#) 50

[Planning to Scale up the Proxy Server](#) 51

5 Clusters 53

[Creating Clusters](#) 53

[Expanding Clusters](#) 54

6 Entity Pages 55

[Virtual Machine Topology](#) 56

[Hosts](#) 56

[VXLAN](#) 57

[VLAN](#) 57

[NSX Manager](#) 57

[Edge Data Collection](#) 58

7 System Dashboards 60

[NI System Dashboard](#) 60

[NI Platform Dashboard](#) 62

[NI Collector Dashboard](#) 63

[Data Source Dashboard](#) 64

[PCI-Compliance Dashboard](#) 65

[Export as PDF](#) 67

[Flow Analytics Dashboard](#) 68

[Analytics - Outlier Detection](#) 70

[How to Detect the Outlier VMs](#) 71

[Analytics: Static and Dynamic Thresholds](#) 72

[Configure Thresholds and Alerts](#) 72

[View the Threshold Configuration Page](#) 74

8 Pins and Pinboards 76

[Pins](#) 76

[Types of Pins](#) 76

[Pinboards](#) 77

[Sharing and Collaboration of Pinboards](#) 81

[To Set A Pinboard as the Home Page](#) 83

[To Duplicate a Pinboard](#) 84

9 Network and Security 85

- Network Visibility 85
 - VM-VM Path 85
 - Monitoring Various States of BGP 93
 - Path to Internet 93
- Security 94
 - Cross vCenter NSX 94
 - Palo Alto Networks 95
 - Cisco ASA Firewall 98
 - Check Point Firewall 100
 - Security Groups 103
 - Policy-Based VPN 104
 - NSX Distributed Firewall Inactive Rules 105

10 Operations 106

- Working with Events 106
 - Working with System Events 106
 - Working with User-Defined Events 108
 - Working with Platform Health Events 109
- Notifications 109
 - Search-based Notifications 109
 - Event Notification Email 110
 - Event Notifications 110

11 Flows 113

- Enabling IPFIX Configuration 113
 - IPFIX Configuration on VDS and DVPG 113
 - VMware NSX IPFIX Configuration 115
 - Adding a NetFlow Collector 116
- Flow Support for Physical Servers 117
 - Configuring a NetFlow Collector in a Physical Device 117
 - Enriching Flows and IP Endpoints 122
 - Search for Physical to Physical Flows 123
- View Blocked and Protected Flows 123
- Network Address Translation (NAT) 124
 - NAT Flow Support - Examples 125
- VMC Flows 126
- Create VPC Flow Log 127

12 Micro-Segmentation 129

- Micro-Segmentation Dashboard 129
- Analyzing the Application 130

- [Viewing Micro-Segmentation And Flow Data in Donut View](#) 130
 - [View Micro-segmentation And Flow Data in Grid View](#) 132
 - [VMC: Planning and Micro-Segmentation](#) 133
 - [Create an Application Manually](#) 134
 - [Creating Tiers for Physical IPs](#) 135
 - [VMC: Planning and Micro-Segmentation](#) 135
 - [Recommended Firewall Rules](#) 136
 - [Exporting Rules](#) 138
- 13 Search** 142
 - [Search Queries](#) 143
 - [Advanced Queries](#) 148
 - [Time Control](#) 151
 - [Search Results](#) 151
 - [Filters](#) 152
 - [vCenter Tags](#) 153
- 14 Upgrade** 156
 - [Online Upgrade of Product](#) 156
 - [Single-Click Offline Upgrade](#) 156
- 15 vRealize Network Insight Disaster Recovery** 158
 - [Sample Disaster Recovery Scenario](#) 159
- 16 Troubleshooting** 162
 - [Audit Logs](#) 162
 - [View the Audit Logs](#) 163
 - [Health](#) 163
 - [Support Tunnel](#) 164
 - [Online Upgrade of Product](#) 164
 - [Single-Click Offline Upgrade](#) 164
 - [View Node Details](#) 165
 - [Capacity](#) 166
 - [Disk Management](#) 167
 - [Creating Clusters](#) 167
 - [Expanding Clusters](#) 168
 - [Creating Support Bundle](#) 169
 - [Migrating Data Sources](#) 169
 - [About Page](#) 170
 - [Customer Experience Improvement Program](#) 171
 - [Common Data Source Errors](#) 171

[Unable to Enable DFW IPFIX](#) 172

17 [Application Migration Using vRealize Network Insight](#) 175

[How do I obtain the CSP Refresh Token for NSX Manager](#) 176

[How Do I Obtain vCenter Credentials](#) 180

[Compute Gateway Firewall Rule](#) 183

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

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.

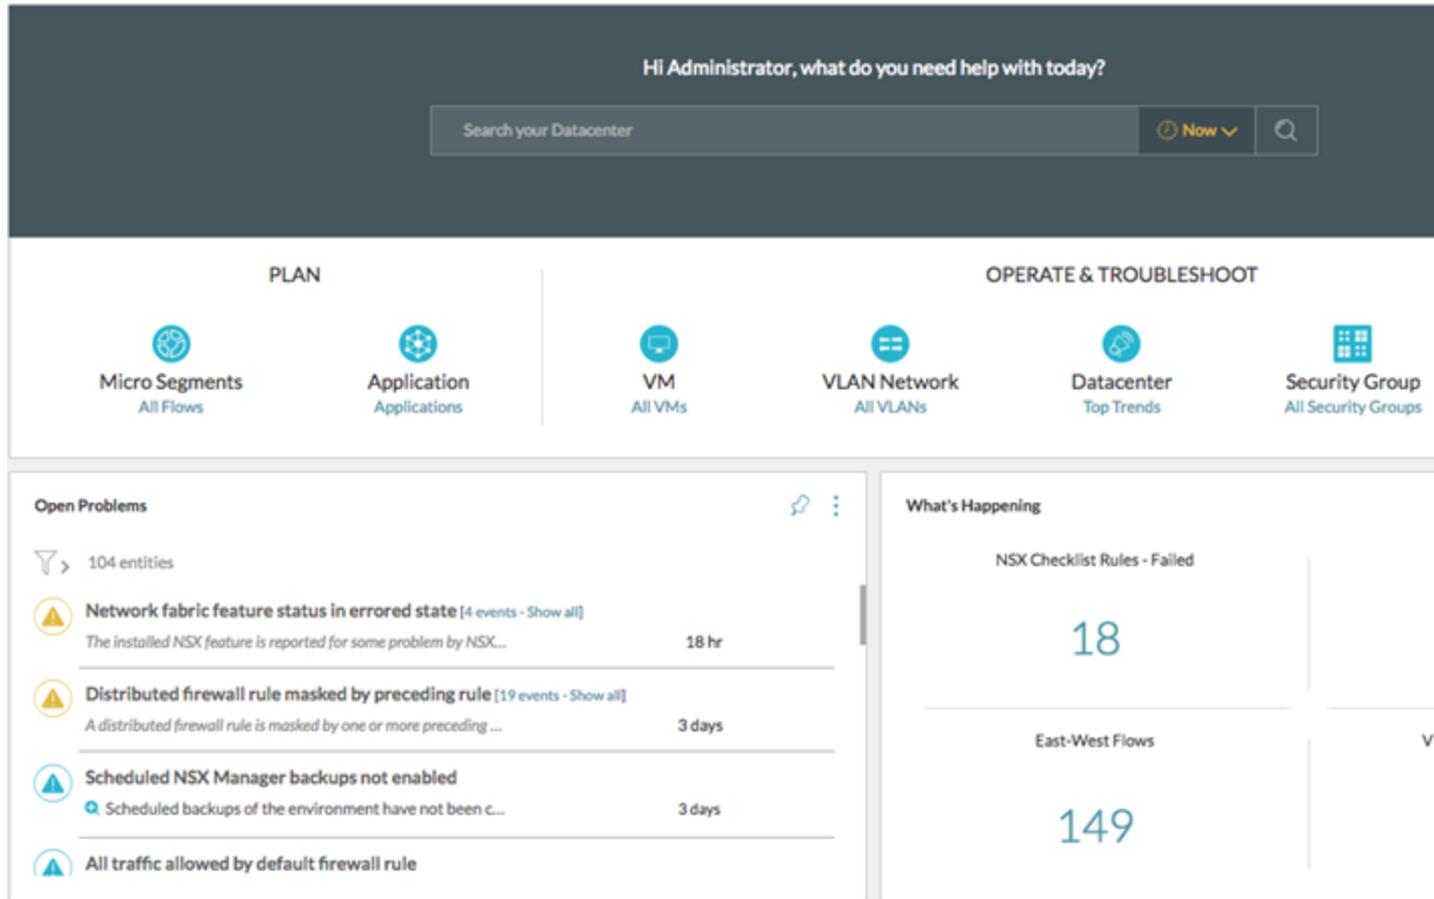
Entities	Description
	Host
	Problem
	NSX Firewall
	Virtual Machine
	vSphere Distributed Switch

Entities	Description
	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
	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:



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.

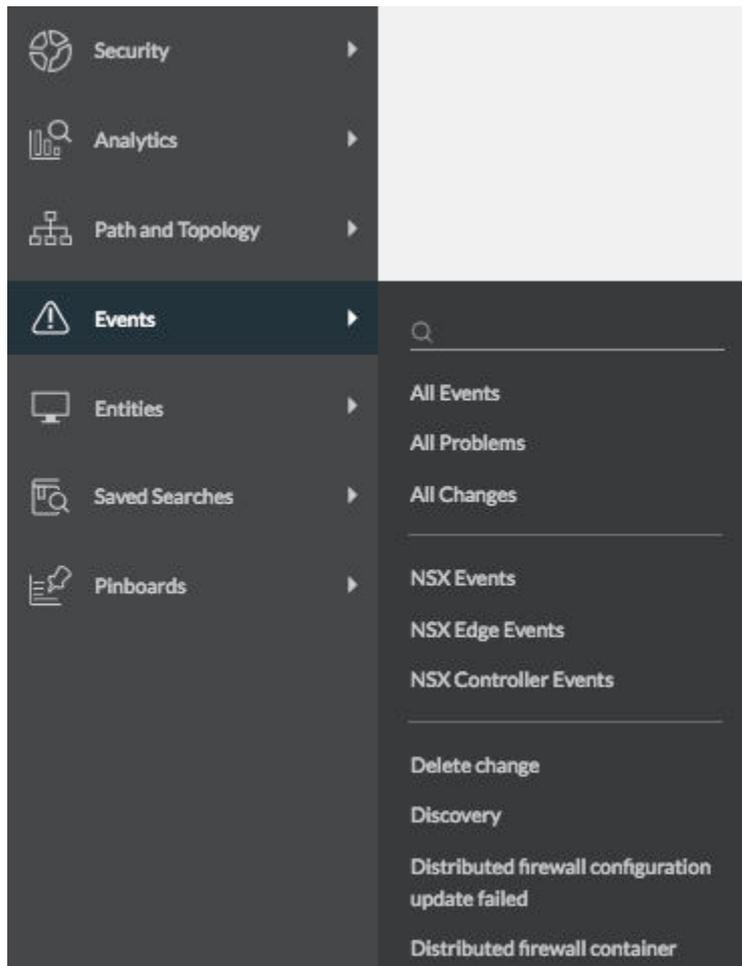
Navigation

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

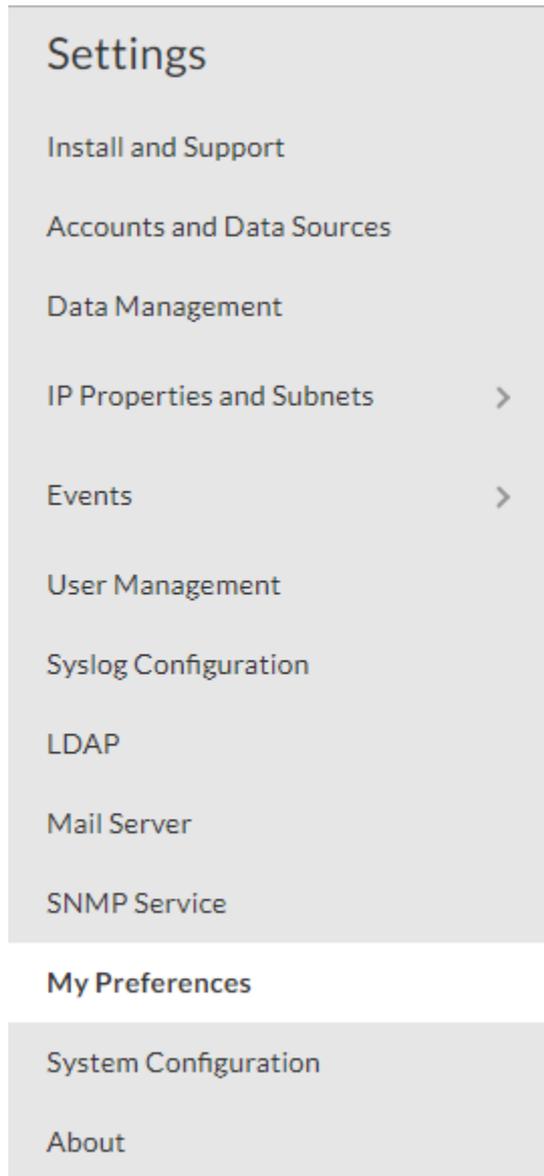
The **Settings** page provides controls to manage data providers, users, and notifications.

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:



Configuration

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 \(SNMP\)](#)
- [Define Timeout Interval for a User](#)

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[™] Chassis 4500 and Cisco[™] 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.

Properties	Description
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 Support** 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.com.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)	<code>us-east-2</code>
US East (N. Virginia)	<code>us-east-1</code>
US West (N. California)	<code>us-west-1</code>
US West (Oregon)	<code>us-west-2</code>
Asia Pacific (Mumbai)	<code>ap-south-1</code>
Asia Pacific (Seoul)	<code>ap-northeast-2</code>
Asia Pacific (Singapore)	<code>ap-southeast-1</code>
Asia Pacific (Sydney)	<code>ap-southeast-2</code>
Asia Pacific (Tokyo)	<code>ap-northeast-1</code>
Canada (Central)	<code>ca-central-1</code>
EU (Frankfurt)	<code>eu-central-1</code>
EU (Ireland)	<code>eu-west-1</code>
EU (London)	<code>eu-west-2</code>
South America (São Paulo)	<code>sa-east-1</code>
Gov Cloud	<code>us-gov-west-1</code>
China (Beijing)	<code>cn-north-1</code>

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.

AWS: Geo-Blocking Support

As the geo-blocking policy is strictly implemented on the corporate firewall, the AWS API calls are limited to specific AWS regions. vRealize Network Insight supports the geo-blocking policy for the AWS environments.

To enable the geo-blocking policy in vRealize Network Insight:

Procedure

- 1 On the **Add AWS Data Source** page, enter the AWS access and the secret keys. Click **Validate**.
- 2 Select **Allow access to specific AWS regions only**. Select the **AWS regions** from the list to enable the automatic collection from the regions. If this option is not selected, then the automatic collection does not happen.
- 3 Click **Submit**.

VMware Cloud on AWS

vRealize Network Insight supports VMware Cloud on AWS (VMC) for the enterprise license users only. You can add VMC vCenter or VMC NSX Manager as a data source.

Adding VMC vCenter as a Data Source

You can add VMC vCenter as a data source.

To add VMC vCenter as a data source:

Note To add VMC vCenter as a data source, you need to have the Cloud Administrator privileges.

Procedure

- 1 To obtain the credentials to add VMC NSX Manager as a data source, perform the following steps:
 - a Log in to your VMWare Cloud Services console.

- b Click **VMware Cloud on AWS** under **My Services**.
 - c Click the name of the desired SDDC.
 - d In the **Settings** tab, copy the **vCenter FQDN** in the **vCenter FQDN** tab. From the **Default vCenter User Account** tab, copy the user credentials.
- 2 In the **Accounts and Data Source** page under **Settings**, click **Add Source**.
- 3 Under **Public Clouds**, click **VMWare Cloud on AWS (vCenter)**.
- 4 In the **Add a new VMware vCenter Account or Source of VMware Cloud on AWS** page, provide the following information. See [How Do I Obtain vCenter Credentials](#).
 - Select the Collector VM.
 - Provide the vCenter FQDN that you have retrieved from the VMware Cloud Services..
 - Provide the user credentials that you retrieved from the VMware Cloud Services.
- 5 Click **Validate**.
- 6 Enter **Nickname** and **Notes** (if any) for the data source and click **Submit**

Adding VMware Cloud on AWS NSX Manager as a Data Source

You can add VMC vCenter as a data source.

Procedure

- 1 To obtain the credentials to add VMware Cloud on AWS NSX Manager as a data source, perform the following steps:
 - a Log in to your VMWare Cloud Services console.
 - b Click **VMware Cloud on AWS** under **My Services**.
 - c Click the name of the desired SDDC.
 - d From the **Support** tab, note the IP address of the NSX Manager.
 - e Click the user profile name on the top right corner of the page.
 - f Click the **API Tokens** tab and copy the Refresh Token.
- 2 In the **Accounts and Data Source** page under **Settings**, click **Add Source**.
- 3 Under **Public Clouds**, click **VMWare Cloud on AWS (NSX Manager)**.
- 4 In the **Add a new VMC NSX Manager Account or Source of VMware Cloud on AWS** page, provide the following information. See [How do I obtain the CSP Refresh Token for NSX Manager](#).
 - Select the corresponding vCenter.
 - The collector is automatically selected based on the selection of the vCenter. VMC NSX Manager has to be added to the same collector VM as that of the corresponding vCenter.
 - Provide the IP address and the CSP Refresh Token.

- Provide the user credentials.
- 5 Click **Validate**.
 - 6 If you want to collect IPFIX flows for DFW, select **Enable DFW IPFIX**.

Note The error messages pop up if the following conditions are not met:

- To enable DFW IPFIX, you need to have the Cloud Administrator privileges.
 - VMC NSX Manager allows only four collectors to be added to its DFW IPFIX collector profile. See also [Unable to Enable DFW IPFIX](#).
-

- 7 Enter **Nickname** and **Notes** (if any) for the data source and click **Submit**

VMC for AWS Entities

Here are the entities related to VMC NSX Manager:

- VMC NSX Manager Data Source
- VMC NSX Manager
- VMC Firewall
- VMC Firewall Rule
- VMC Segment
- VMC Policy-Based VPN
- VMC Group

Some sample search queries related to the VMC entities are:

- VMs where L2 Network = '' (L2 Network -> VMC Segment)
- VMC Policy Based VPN where Tier0 = ''
- VMC Policy Based VPN where Local Network = '' (Local Network of Policy Based VPN Rule)
- VMC Policy Based VPN where Remote Network = '' (Remote Network of Policy Based VPN Rule)
- VMC Group where Translated VM = ''
- VM where VMC Group = ''

Note

- Currently, vRealize Network Insight does not support the following search query:

```
'vmc group where members = <IP/VM/MEMBERSHIP_CRITERIA/CHILD_GROUPS/IPSETS>'
```

- VMC does not support child groups or IPSETS. Hence all the searches like vmc firewall rule where Indirect _____ = '' or vmc group where Indirect _____ = '' are disabled.
-

VMware Cloud for AWS (VMC) Deployment

vRealize Network Insight supports the following deployment model:

- Collector deployed in VMC:
 - a In this deployment model, the collector is deployed as a workload in Compute Gateway in VMC. The platform is deployed in the SDDC on-premises version.
 - b The firewall rules of Management Gateway allow communication to VMC vCenter and VMC NSX Manager over HTTPS.
 - c The collector communicates to the platform using the existing communication mechanisms over VPN or Direct Connect.

The prerequisites for the above deployment model are:

- There should be a Management Gateway firewall rule to allow the vRealize Network Insight collector to invoke vCenter and NSX Manager APIs over HTTPS (443).
- There should be a Compute Gateway rule within the gateway firewall to allow the collector to communicate with the on-premises Platform or the SaaS platform.

Note

- For a single node SDDC in VMC, set the CPU resource reservation for the proxy VM to 1251 MHz. Currently, the proxy OVA delivered as a part of the release has the resource reservation set to 2048 MHz. After importing the OVA in the SDDC vCenter, modify the settings of the proxy VM to use the maximum allowed CPU reservation of 1251 MHz.
-

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.
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 Infoblox 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.

Keyword	Sample Search Query	Description
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 a NetFlow Collector

Procedure

- 1 In the **Settings** page, click **Accounts and Data Sources**.
- 2 Click **Add Source**.
- 3 Under **Flows**, 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.
- 4 Enter **Nickname** and **Notes** as required. Click **Submit**.

Sampled Flows

vRealize Network Insight supports the addition of switches or the routers that support only sampled flow (sFlow) as data sources.

Currently, vRealize Network Insight does not show metrics for the sFlows.

You can search for the sFlows by using the `flow type = 'sFlow'` query.

Adding a Sampled Flow Collector

Procedure

- 1 In the **Settings** page, click **Accounts and Data Sources**.
- 2 Click **Add Source**.
- 3 Under **Flows**, click **Physical Flow Collector**. sFlows are accepted only on the physical collector.
- 4 Enter **Nickname** and **Notes** as required. Click **Submit**.

F5 BIG-IP

vRealize Network Insight supports F5 BIG-IP 12.1.2 and the further versions in the routed mode only. The features like VM-VM path, high availability, VRFs, Routes, Router Interfaces, Switch Ports, Port Channels,

Switch Port metrics, VRF Dashboard, Switch Dashboard and Router dashboard are supported. For search on the F5 BIG-IP entities, use the query string `F5 BIG-IP Data Source`. vRealize Network Insight does not support LLDP neighbors or the neighboring devices in the VM-VM path.

To add F5 BIG-IP as a data source:

Prerequisites

- The user should have at least a Guest role.
- On the F5 BIG-IP UI console, ensure that the following conditions are met:
 - a The terminal access must not be in the disabled mode.
 - b For Partition Access, select ALL.
 - c Enable password authentication for SSH as follows:

Note

- Use root or the administrator role privilege for changing the SSHD configuration.
 - Do not use the root user privilege while adding F5 BIG-IP data source in vRealize Network Insight .
 - Root user does not have HTTP access. The root user privilege is used for the administrative purpose.
-

```
[root@bigip:Active] config # tmsh
root@bigip(Active)(/Common)(tmsh)# edit sys sshd

## Adding the following configuration ##

modify sshd {
  include "
  ChallengeResponseAuthentication no
  PasswordAuthentication yes"
}
#####
Save changes? (y/n/e) y
root@bigip(Active)(/Common)(tmsh)#
root@bigip(Active)(/Common)(tmsh)# save sys config

root@bigip(Active)(/Common)(tmsh)# show running-config sys sshd
sys sshd {
  include "
  ChallengeResponseAuthentication no
  PasswordAuthentication yes"
}
```

Procedure

- 1 On the Settings page, click **Accounts and Data Sources**.
- 2 Click **Add Source**.

3 Under **Routers & Switches** select **F5 BIG-IP**.

4 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.

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

6 If you enable SNMP for the data collection, select **SNMP Version**.

- a If you select 2c, then enter the associated community string.
- b If you select 3, then enter the following:
 - Username
 - Context Name
 - Authentication Type

Note Ensure that you have enabled SNMP by providing the SNMP Access details in the F5 BIG-IP UI console.

7 Provide the **Nickname** and **Notes** as required. Click **Submit**.

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

- NSX-T 2.0, 2.2, and 2.3 versions are supported.
- vRealize Network Insight supports only the NSX-T setups in which vCenter manages the ESXi hosts. Ensure that vCenter is added as Compute Manager in NSX-T.

Note Compute Managers should be added as data sources in vRealize Network Insight before adding NSX-T as a data source.

- vRealize Network Insight supports NSGroups, NSX-T Firewall Rules, IPSets, NSX-T Logical Ports, NSX-T Logical Switches, and NSX-T distributed firewall IPFIX flows.

- 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.
- Ensure that there are no logical switches in the exclusion list in the Distributed Firewall (DFW). If there are any logical switches in this list, then the flows are not reported for any VMs attached to these logical switches.

To add an NSX-T Manager:

- 1 On the **Accounts and Data Source** page under **Settings**, click **Add Source**.
- 2 Under **VMware Manager** in the **Select an Account or Data Type** page, select **VMware NSX-T Manager**.
- 3 Provide the user credentials.

Note If IPFIX is not required, the user must be a local user with the audit level permissions. But if IPFIX is required then the user must have one of the following audit level permissions: **enterprise_admin**, **network_engineer**, or **security_engineer**.

- 4 Select **Enable IPFIX** to update the IPFIX settings on NSX-T. By selecting this option, vRealize Network Insight receives DFW IPFIX flows from NSX-T. For more information on enabling IPFIX, see [Enabling VMware NSX-T DFW IPFIX](#).

Note DFW IPFIX is not supported in the Standard Edition of NSX-T.

Examples for Queries

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

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 instance of vRealize Network Insight including the details on whether it is a system-created or a user-created switch.
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.

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.
NSX-T Transport Zone	Lists the VLAN and the overlay transport zone and the respective details associated with it including the type of the transport node. Note vRealize Network Insight does not support KVM as a data source.
NSX-T Router	Lists the TIER 1 and TIER 0 routers. Click the router shown in the results to view more details associated with it including the NSX-T Edge Cluster and the HA mode.

Support for NSX-T Metrics

The following table displays the vRealize Network Insight entities that support the NSX-T metrics currently and the widgets that display these metrics on the corresponding entity dashboards.

Entities	Widgets on the Entity Dashboard	Supported NSX-T Metrics
Logical Switch	Logical Switch Packet Metrics Logical Switch Byte Metrics	Multicast and Broadcast Rx Multicast and Broadcast Tx Unicast Rx Unicast Tx Dropped Rx Dropped Tx Rx Packets (Total) Tx Packets (Total)
Logical Port	Logical Port Packet Metrics Logical Port Byte Metrics	Multicast and Broadcast Rx Multicast and Broadcast Tx Unicast Rx Unicast Tx Rx Packets (Total) Tx Packets (Total)

Entities	Widgets on the Entity Dashboard	Supported NSX-T Metrics
Router Interface	Router Interface Metrics	Rx Packets Tx Packets Dropped Rx Packets Dropped Tx Packets Rx Bytes Tx Bytes
Firewall Rule	Firewall Rule Metrics	Hit Count Flow Bytes Flow Packets

Here are some sample queries for NSX-T Metrics:

- `nsx-t logical switch where Rx Packet Drops > 0`

This query lists all the logical switches where the count of the dropped received packets is greater than 0.

- `nsx-t logical port where Tx Packet Drops > 0`

This query lists all the logical ports where the count of the dropped transmitted packets is greater than 0.

- `top 10 nsx-t firewall rules order by Connection count`

This query lists the top 10 firewall rules based on the connection count(Hit Count).

Support for NAT in NSX-T

Currently, vRealize Network Insight supports SNAT, DNAT, reflexive rules in the flows and the VM to VM Path.

To obtain all the NAT rules in NSX-T, use the `NSX-T Edge NAT Rule` query. To obtain all the NAT rules in both NSX-V and NSX-T, use the `NAT Rules` query.

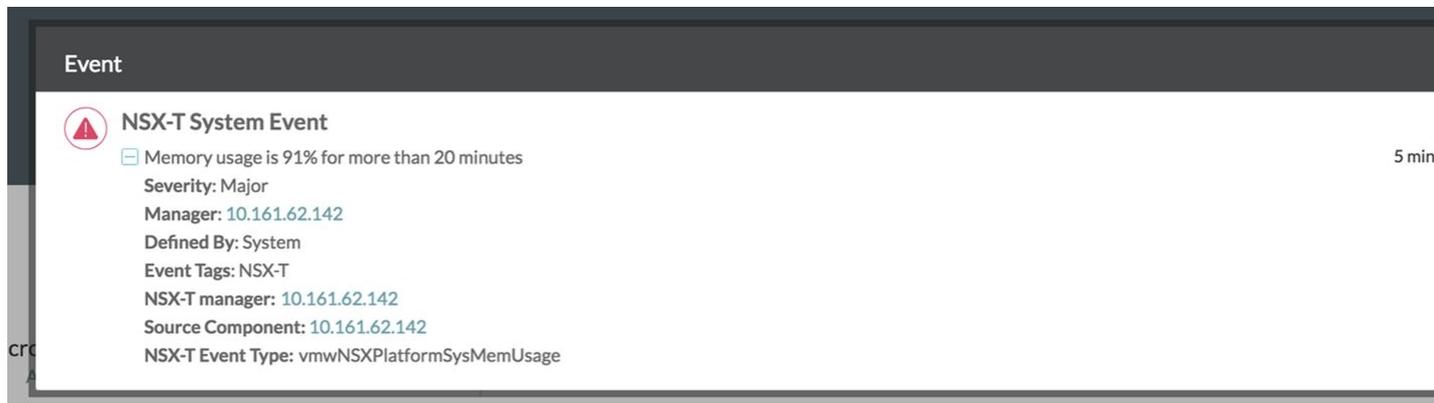
Considerations

- For the VM-VM path with the NSX-T logical routers where NAT service is enabled, vRealize Network Insight wont show the NSX-T edge firewall rules correctly for such a path.
- The NAT rules that are configured on the uplink interface of the VMware NSX-T Tier Router only are processed by the VM to VM path. If NAT is configured on any NSX-T Tier router, then it is expected that there are NAT rules for all the VMs attached to the router else the VM to VM path and the path to Internet does not work. Instead, it displays a missing rule message.
- vRealize Network Insight supports the nested NAT hierarchy.
- NSX-V and NSX-T based edges only are supported.
- vRealize Network Insight supports the edges and the tier routers with NAT-defined uplinks.

- vRealize Network Insight supports SNAT rules with range. However, DNAT must be one-to-one mapping between the destination and translated IP addresses (Parity with NSX-V).
- vRealize Network Insight does not support the following use cases:
 - a In NSX-T, NAT rules can be applied at the service level. For example, in NSX-T, L4 ports set is a type of service and the associated protocols can be TCP or UDP. So in the VM-VM path, the service level details are not supported.
 - b Any port level translation is not supported.
 - c The SNAT match destination address and the DNAT match source address are not supported. Use the SNAT match destination address as the destination IP address when you specify the SNAT rule. Use the DNAT match source address as the source IP address when you specify the DNAT rule. For example, if there is a destination IP address mentioned in the SNAT rule, vRealize Network Insight applies the SNAT rule irrespective of whether the packet has the destination address as the destination IP address.
 - d NSX-T Edge firewall has implications for the data path when enabled with the NAT service on the same logical router. If a flow matches both NAT and Edge firewall, the NAT lookup result takes precedence over firewall. So the firewall is not applied to that flow. If the flow matches only a firewall rule, then the firewall lookup result is honored for that flow.

NSX-T System Events

vRealize Network Insight enables you to view the NSX-T system events for NSX-T version 2.3 onwards.



The source component is the source that emits the event. It can be a manager, controller, or an edge.

The following NSX-T system events are generated by vRealize Network Insight:

- NSX-T MTU Mismatch

Note The mismatch is between the uplinks of T0 and the corresponding interfaces of the uplink device.

- NSX-T Edge Node has no manager connectivity
- NSX-T Edge Node has no controller connectivity

- NSX-T Edge Node's controller connectivity degraded
- Routing Advertisement disabled for Tier-1 router
- No Uplink Connectivity for Tier-1 router

Note The Tier-1 router is not connected to Tier-0.

Cisco ACI

You can add Cisco ACI as a data source. Currently, vRealize Network Insight supports Cisco ACI 3.2. This feature is available only for the enterprise license users.

To add Cisco ACI as a data source, the user should have access to all the tenants and read-only privilege. Here are the steps for the Cisco ACI data source addition:

Procedure

- 1 In the **Accounts and Data Source** page under **Settings**, click **Add Source**.
 - 2 Under **Others**, click **Cisco ACI**.
 - 3 In the **Add a new Cisco ACI Account or Source** page, provide the following information:
 - Select the Collector VM.
 - Provide the IP address of any of the APIC controllers in the cluster.
-
- Note** You do not have to add the individual switches in the ACI fabric.
-
- Provide the user credentials.
 - vRealize Network Insight collects the metric data over SNMP from the individual switches. To enable this task, select **Use SNMP**.
- 4 Click **Validate**.
 - 5 Enter **Nickname** and **Notes** (if any) for the data source and click **Submit**

Cisco ACI Entities

Here is a list of some of the Cisco ACI entities on which you can perform a search:

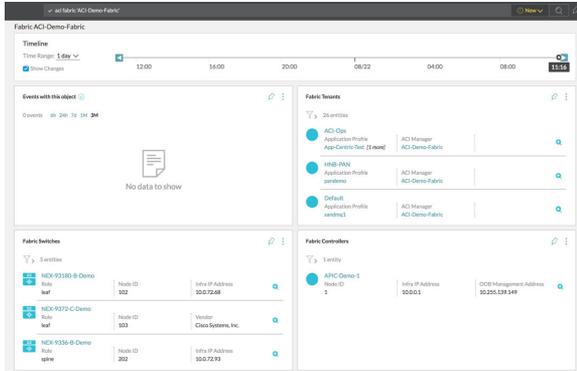
Note The entities are prefixed by `aci`.

- `aci application profile`
- `aci bridge domain`
- `aci endpoint group`
- `aci fabric`
- `aci switch`

- aci tenant

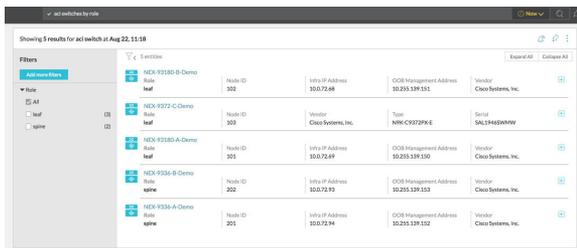
Here are some sample search queries:

- aci fabric 'ACI-Demo-Fabric': This query retrieves information on the tenants, switches, and controllers in the ACI fabric.

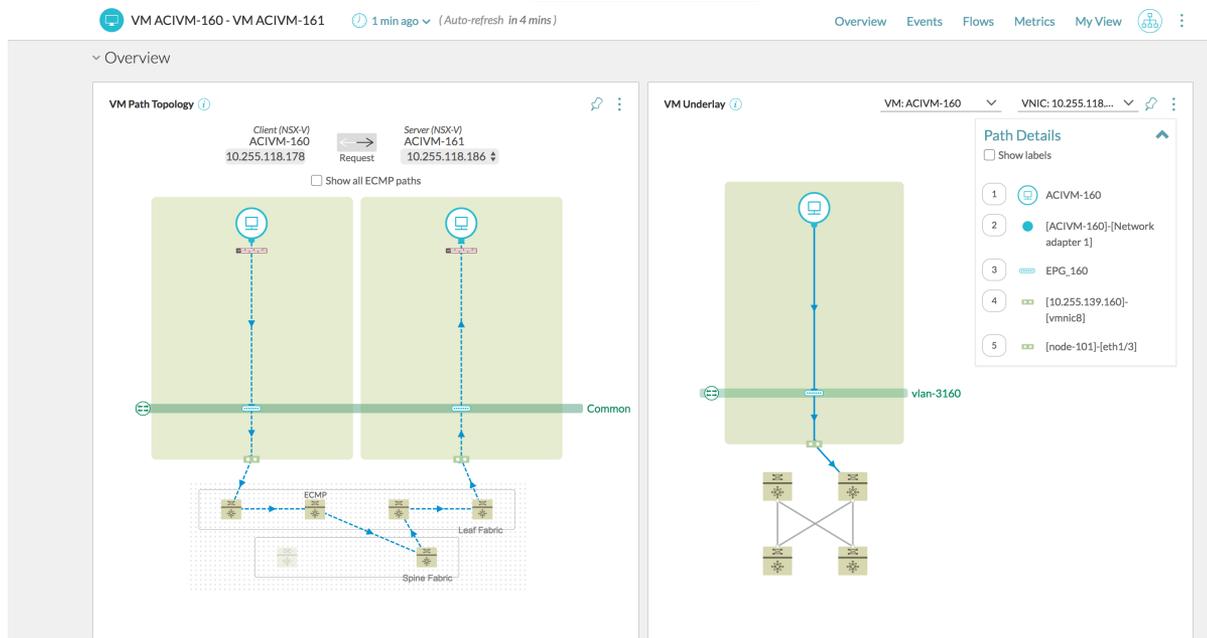


- aci switches by role: This query retrieves information on the various leaf switches or the spine switches in the ACI fabric.

From the switch list, click a switch name to get more details about it.



- aci endpoint group: This query retrieves a list of the endpoint groups with the associated VMs, bridge domains, and VRFs.
- aci application profile 'Production': This query retrieves the application profile of Production with the contained endpoint groups and VMs.
- VMware VM 'ACIVM-160' to VMware VM 'ACIVM-161': This query shows the VM-VM path between the two VMs.



Adding vRealize Log Insight as a Data Source

vRealize Log Insight collects NSX logs dynamically when an NSX event occurs. However, vRealize Network Insight collects data from NSX every 10 minutes. So, adding vRealize Log Insight in vRealize Network Insight enables you to get event information faster, rather than waiting for it.

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

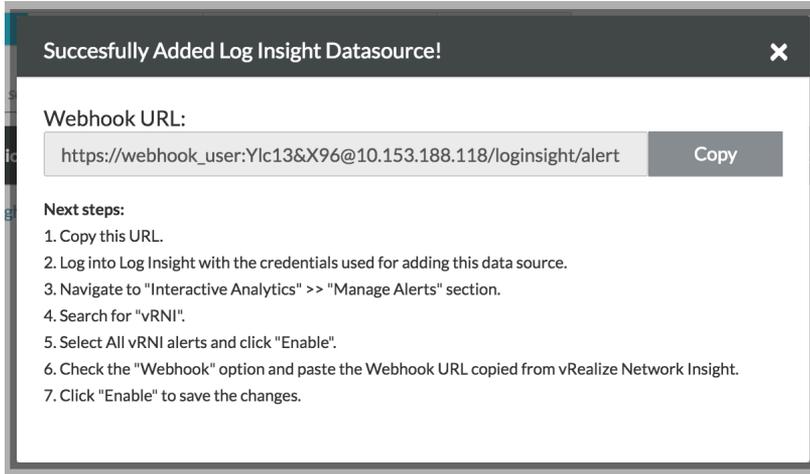
Prerequisites

- 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.

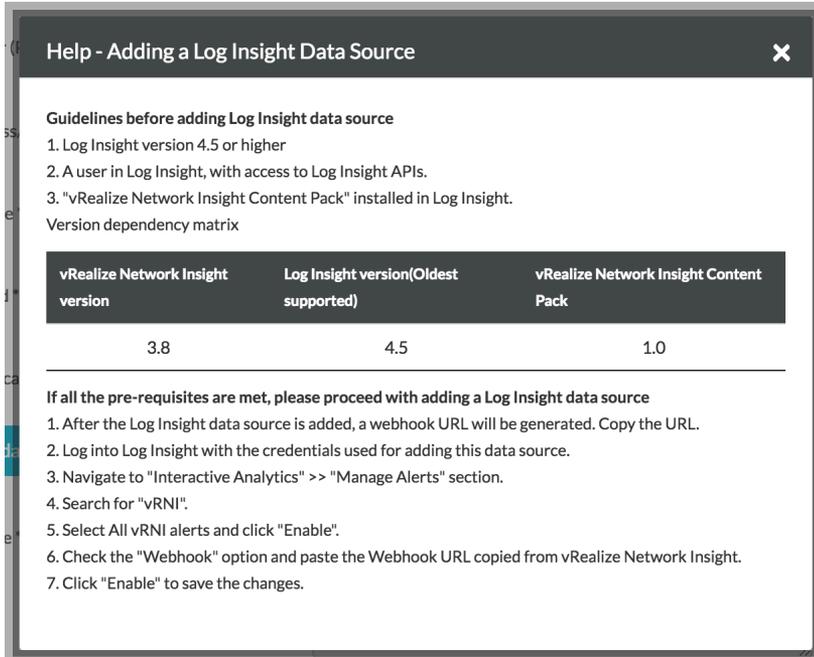


Note The Webhook URL, which is generated after the addition of the data source, is used in vRealize Log Insight.

- 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.



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

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.

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	<ul style="list-style-type: none"> ■ To add NSX Manager as a data source, you have to be a super user, an enterprise administrator, an auditor, or an NSX security administrator. ■ An enterprise administrator, a super user, an NSX security administrator, or an auditor can run the NSX Central CLI commands required by vRealize Network Insight. <p>Note An NSX network administrator cannot add NSX Manager as a data source.</p>
NSX 6.2 and the further releases before NSX 6.4	<ul style="list-style-type: none"> ■ The user should be an administrator to enable the edge data population. ■ An auditor, a super user, or an NSX security administrator can run the NSX Central CLI commands required by vRealize Network Insight. ■ 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
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: <ul style="list-style-type: none"> ■ Critical ■ Moderate ■ Warning ■ Info
Type	This field specifies if the event denotes a Problem or a Change . <p>Note All the events of type Problem are logged into syslog.</p>
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: <ul style="list-style-type: none"> ■ 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.
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**.

Define Timeout Interval for a User

By default, the user session timeout is set to 15 minutes. You can modify this value according to your preference.

Procedure

- 1 On the **Settings** page, click **System Configuration**.

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

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

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
 - Assume that on January 1st (today), the datacenter has 2000 VMs (with flows) across many vCenters.
 - Assume that in March, the number of VMs grows to 3100.
 - 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.
 - Assume that in December, the number of VMs grows to 18100 (with flows).

The deployment model for this scenario is as follows:

- On January 1, deploy a single platform node with the MEDIUM brick size.
 - In March, scale up the platform node to the LARGE brick size.
 - In June, scale out the platform, convert to a three node platform cluster by adding new Platform nodes to the existing Platform.
 - In December, the user needs a four node platform cluster.
- Scenario 2
 - Assume that on January 1st (today), the datacenter has 7000 VMs (with flows) across many vCenters.
 - Assume that in June, the number of VMs grows to 15000 (with flows).
 - Assume that in December, the number of VMs grows to 24000 (with flows).

The deployment model for this scenario is as follows:

- 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

- Scenario 1: 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

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

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

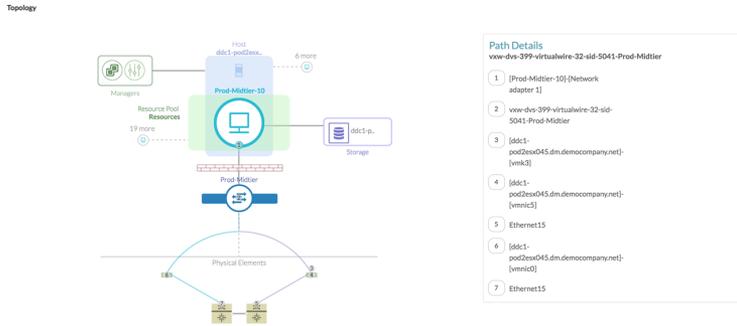
The property pins display important attributes in a two-column layout. Some property pins might also display only a singular attribute value. An example of the property pin is the **VM Properties** pin. The **VM Property** pin displays the properties of a VM, such as operating system, IP address, default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

This chapter includes the following topics:

- [Virtual Machine Topology](#)
- [Hosts](#)
- [VXLAN](#)
- [VLAN](#)
- [NSX Manager](#)

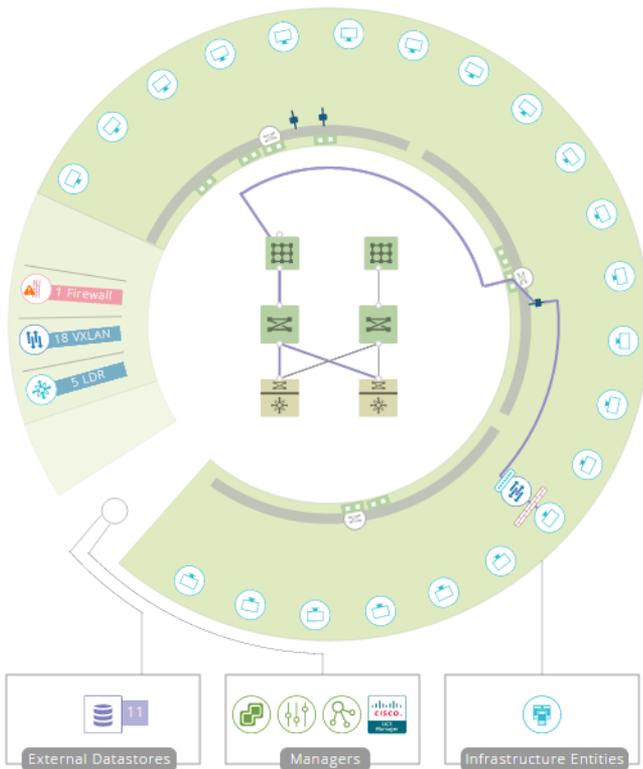
Virtual Machine Topology

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



Hosts

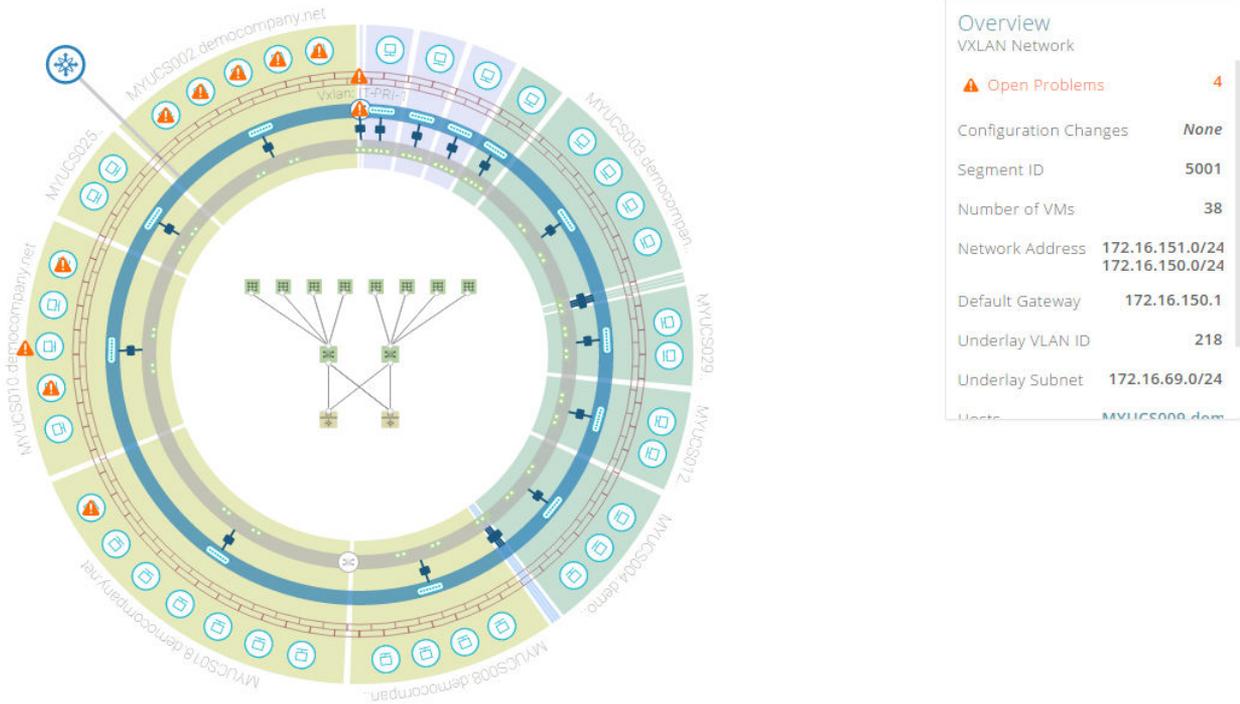
The host topology shows how VMs of a particular host are connected to the virtual and physical components of your data center and also how the host itself is connected with your data center.



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:



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

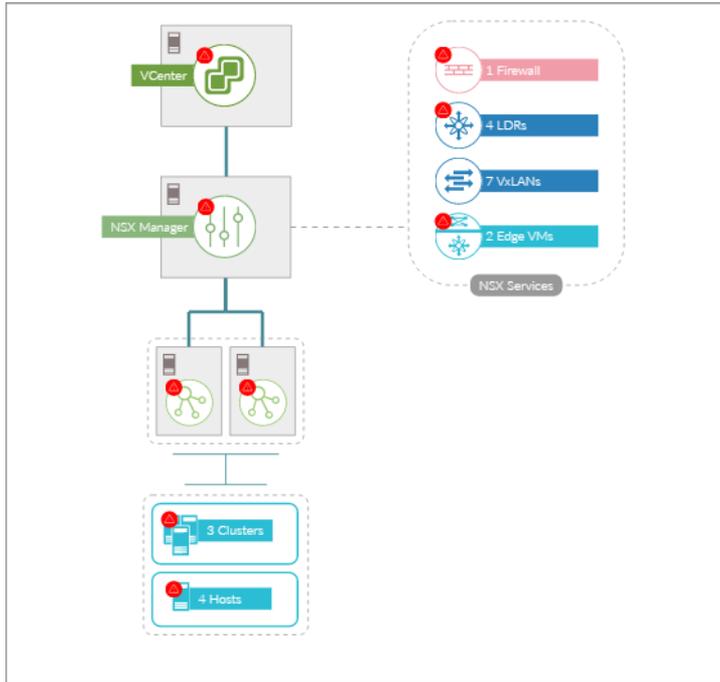
VLAN

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.

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	<ul style="list-style-type: none"> ■ To add NSX Manager as a data source, you have to be a super user, an enterprise administrator, an auditor, or an NSX security administrator. ■ An enterprise administrator, a super user, an NSX security administrator, or an auditor can run the NSX Central CLI commands required by vRealize Network Insight. <p>Note An NSX network administrator cannot add NSX Manager as a data source.</p>
NSX 6.2 and the further releases before NSX 6.4	<ul style="list-style-type: none"> ■ The user should be an administrator to enable the edge data population. ■ An auditor, a super user, or an NSX security administrator can run the NSX Central CLI commands required by vRealize Network Insight. ■ 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

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](#)
- [PCI-Compliance Dashboard](#)
- [Flow Analytics Dashboard](#)
- [Analytics - Outlier Detection](#)
- [Analytics: Static and Dynamic Thresholds](#)

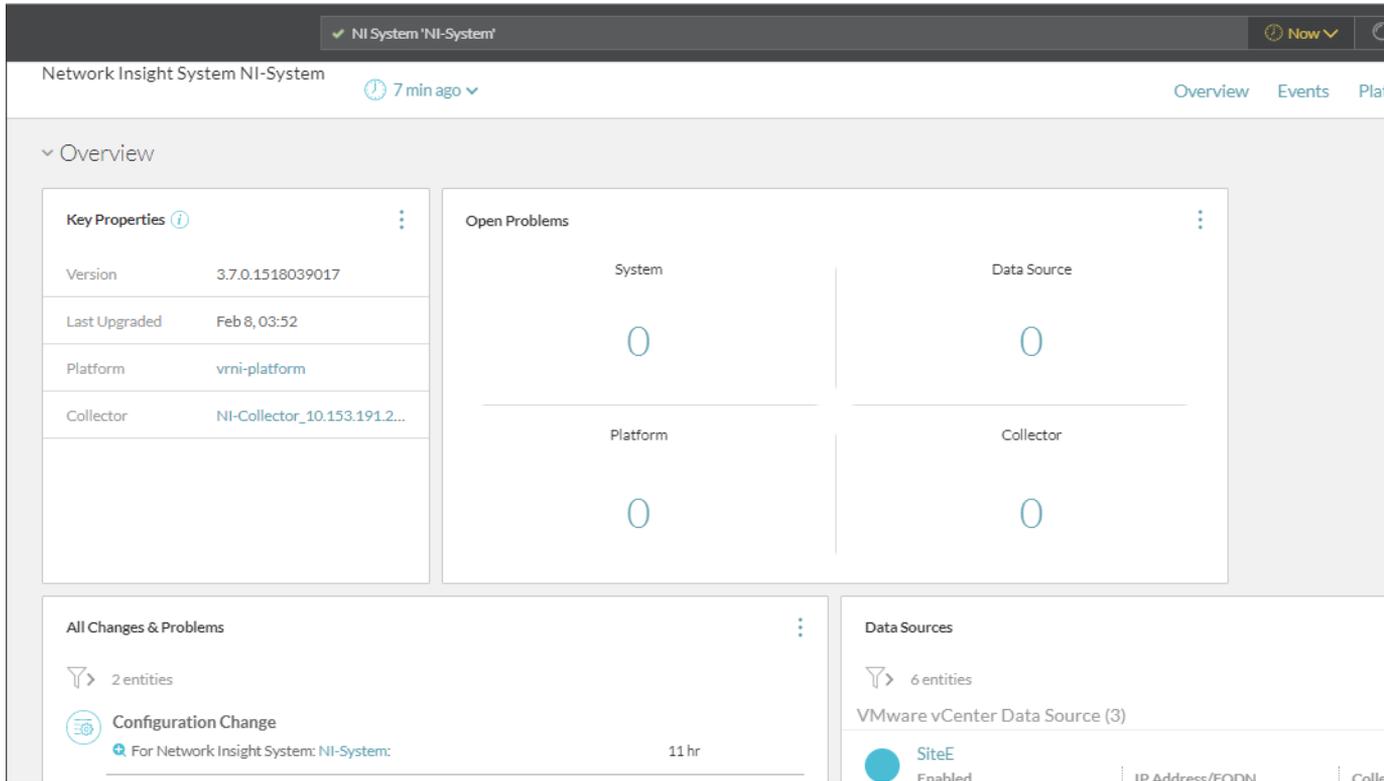
NI System Dashboard

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

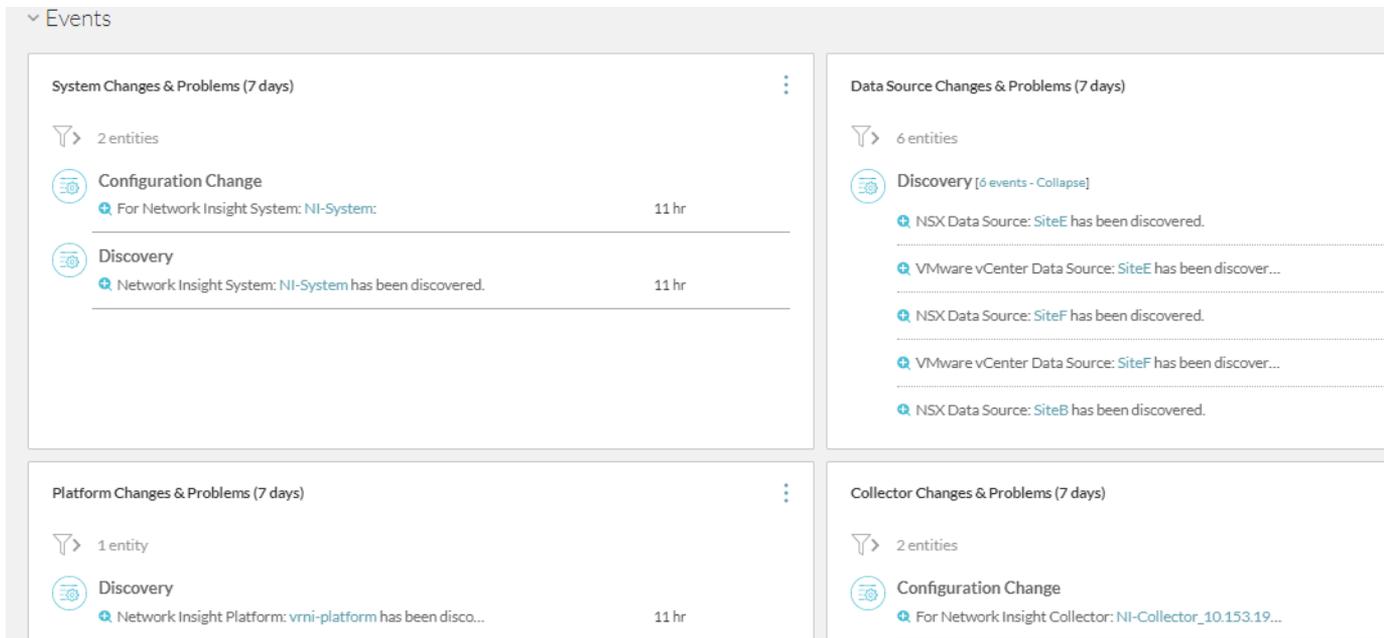
- On the **Install and Support** page, click **View Details** next to **Overview**. The NI System dashboard appears.
- Provide `NI-System` as the search query to view the NI System dashboard.

The NI System dashboard is divided into three sections:

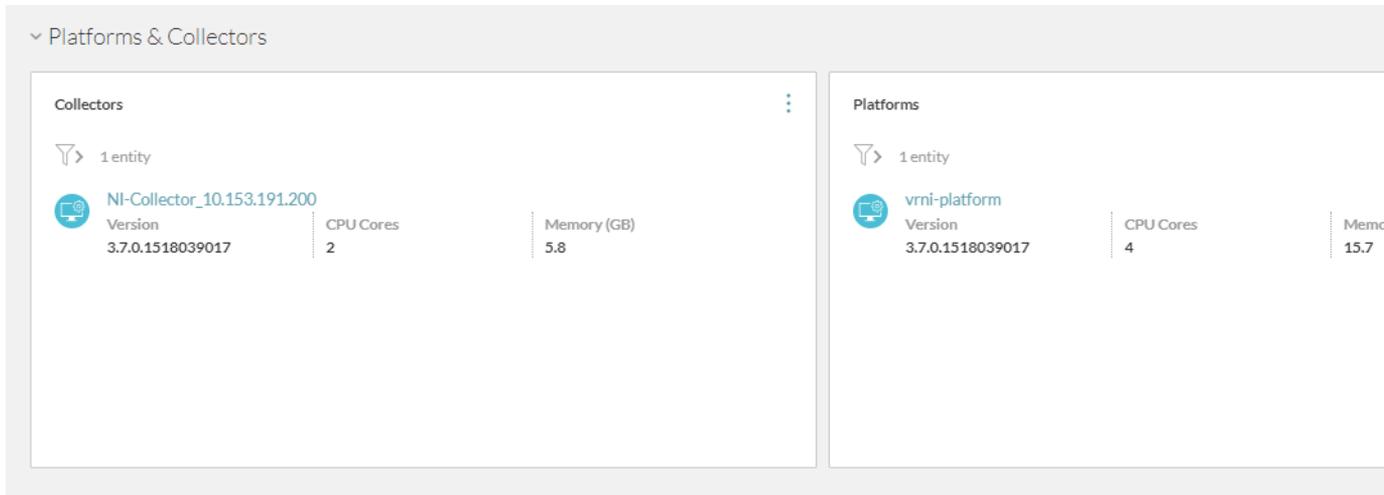
- **Overview:** This section consists of information on the key properties, the data sources, the problems that are open, and all the changes and the problems related to the system. View the details of each data source by clicking it.



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

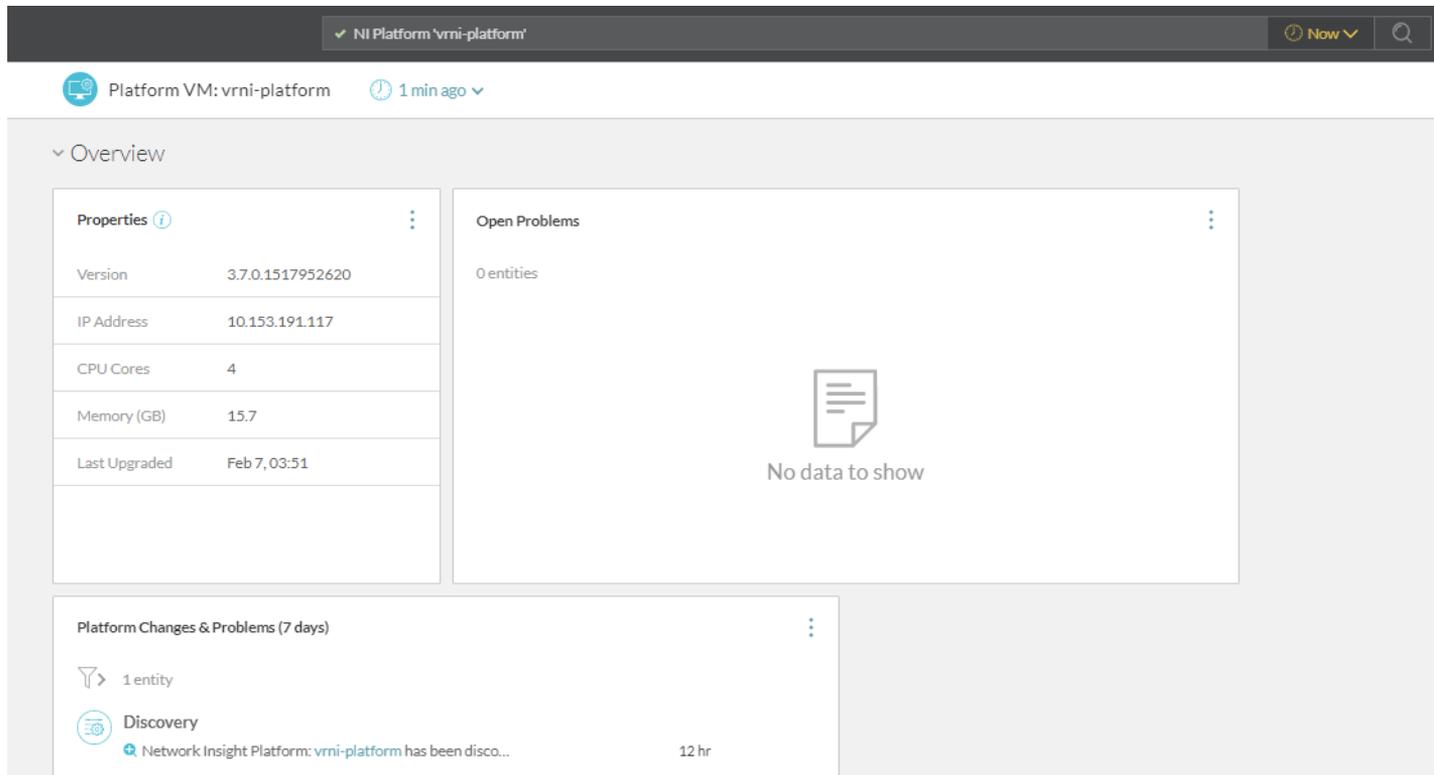


- **Platforms and Collectors:** This section lists all the platforms and the collectors associated with the system. To view more details about any platform or collector, click it.



NI Platform Dashboard

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



- **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.

The screenshot shows the NI Collector Dashboard for the collector node 'NI-Collector_10.153.191.118'. The dashboard is divided into several sections:

- Properties Table:**

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
- Open Problems (Data Sources):** A card showing 0 open problems.
- VMware vCenter Data Source (3):** A list of 6 entities:

Site	Enabled	IP Address/FQDN	Collector
SiteE	Yes	10.197.17.173	NI-Collector_10.153.191.118
SiteF	Yes	10.197.17.137	NI-Collector_10.153.191.118
SiteB	Yes	10.197.17.137	NI-Collector_10.153.191.118
- Events:**
 - Collector Changes & Problems (7 days):** 4 entities, including Configuration Change [3 events - Collapse].
 - Data Source Changes & Problems (7 days):** 7 entities, including 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.

The screenshot displays the configuration page for a VMware vCenter Data Source named 'SiteF'. The interface includes a top navigation bar with the data source name and a 'Now' dropdown menu. Below this, the 'Overview' section is expanded, showing three main panels:

- Properties:** A table listing various configuration options:

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:** A panel indicating '0 entities' and 'No data to show' with a document icon.
- Changes & Problems (7 days):** A panel showing '1 entity' under the 'Discovery' category, with a specific event: 'VMware vCenter Data Source: SiteF has been discover...' occurring '11 hr' ago.

- **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.

PCI-Compliance Dashboard

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

To access the PCI-Compliance feature

- 1 In the navigation panel on the left of the Homepage, select **Security > PCI Compliance**.
- 2 The PCI Compliance window appears. Select the required scope, the corresponding entity, and the duration for which you require the data. Click **Assess**.
- 3 The PCI-Compliance dashboard appears.

PCI-Compliance Dashboard Features

The screenshot displays the PCI Compliance dashboard interface. At the top, there's a header with 'pci compliance of Cluster 'HasS-Cluster-1'', 'PCI Compliance', 'Select Scope', and 'Now'. Navigation tabs include 'Scope', 'Network flows', 'Firewall rules', 'Security changes', and 'My View'. The main content area is titled 'Scope' and contains several panels:

- PCI sections covered:** A list of 7 PCI requirements for a VMware NSX environment, such as 'Section 1.1.1 - A formal process for approving and testing all network connections...'.
- Virtual machines in scope (Section 1.3.1):** A table listing 122 entities. Visible entries include:

Entity	Outgoing Rules	Incoming Rules	Manager
PAN8-Cluster2-ESX-2	Default Rule [2 more]	behind-http-allow... [4 more]	10.197.17.51
VMware vRealize Network Insight Platform	Default Rule [2 more]	behind-http-allow... [4 more]	10.197.17.51
CP_vCenter 6.5	Default Rule [3 more]	behind-http-allow... [5 more]	SG-a1
- Security groups of virtual machines in scope (Section 1.3.1):** A table showing 1 entity:

Entity	Members	Direct Children
SG-a1 (10.197.17.52)	SECURITY_TAG-a1	SECURITY_TAG-a1
- Virtual Machine by security groups (Section 1.3.1):** A table showing 122 entities:

Security Groups	Count of VM
SG-a1	2
- Virtual machine by security tags (Section 1.3.1):** A table showing 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**.

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:

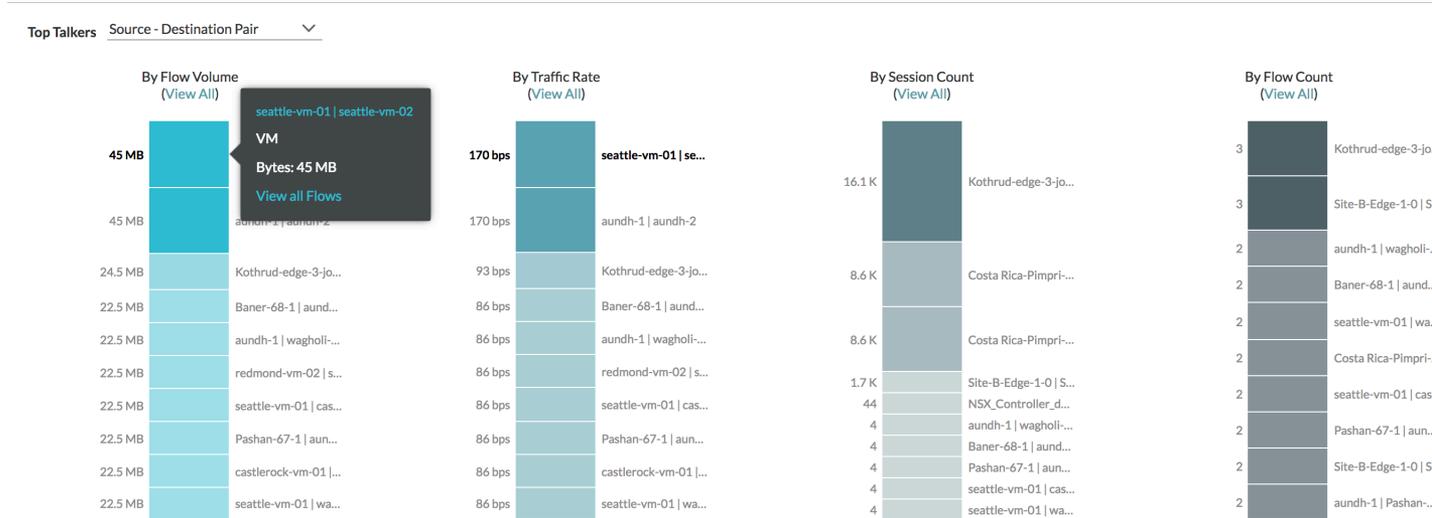
- Top Talkers
- What's New
- Outliers

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



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 Source-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.

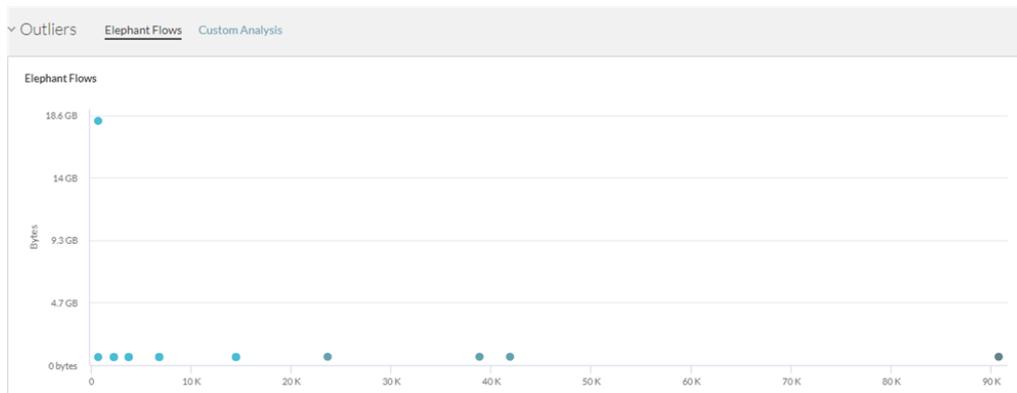
- **New Internal/E-W Services Accessed:** Lists the services that are exposed and accessed by the machines within a data center
- **New Services with Blocked Flows:** Lists services that have blocked flows. This section is populated only for IPFIX.
- **New Firewall Rule Hits:** Lists the new firewall 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	<p>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.</p> <p>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.</p> <p>If you select Security Group, provide the name of the Security Group.</p> <p>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.</p> <p>You can view the micro segmentation for the selected configuration by clicking View Micro-Segments.</p>
Detection Type	Currently, Network Insight enables you to detect the outlier in the system.
Metric	<p>The detection is based on this flow metric. You can select the following options:</p> <ul style="list-style-type: none"> ■ 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.
Traffic Type	You can select Internet , East-West , or All based on the requirement.

Field	Description
Destination Ports	<p>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</p> <hr/> <p>Note The current limit for the number of ports is 20.</p>
Sensitivity	<p>It is a measure of the sensitivity of the detection and reporting that you require. The default value is Medium.</p>
Preview	<p>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.</p>

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.

Analytics: Static and Dynamic Thresholds

vRealize Network Insight enables you to set and configure thresholds and receive alerts based on aberrations in the behavior of the entities. You can configure two types of thresholds:

- **Static Threshold:** If a particular metric value goes beyond or below the configured value, then a static-threshold-based alert is generated.
- **Dynamic Threshold:** If the threshold is determined by the system based on the analysis of the historical data, an alert is generated in case this threshold is violated. The data is analyzed for a period of 7 days before any alert is generated. The process of creating a baseline is restricted to 21 days of the historic data and the older metric values are not considered to create a baseline for the new metric values.

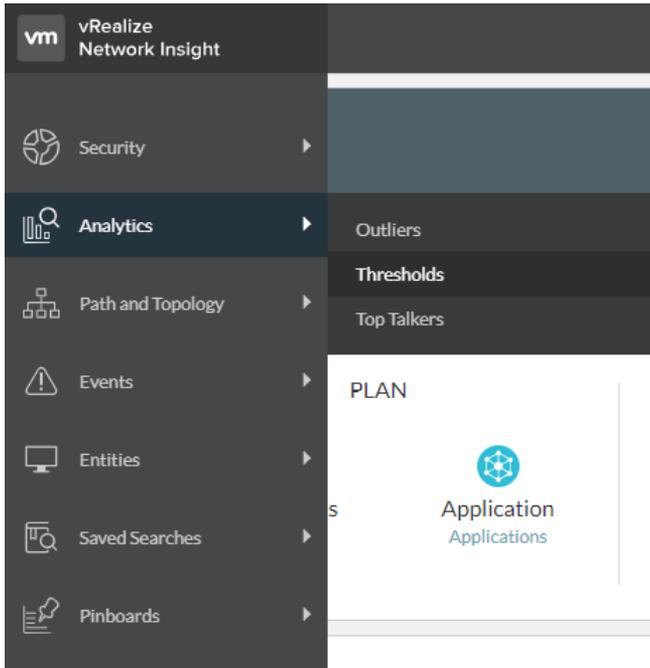
The alert is generated immediately after a threshold is violated.

Configure Thresholds and Alerts

To configure the analytics-associated thresholds and alerts:

Procedure

- 1 On the left navigation panel, click **Analytics**. Click **Thresholds**.



- 2 Click **Add** to add a threshold configuration.
- 3 Provide a unique name for the configuration.
- 4 Select the scope for the threshold that you want to set. Define the set of entities that you want to analyze. Currently, the entities that are available for analysis are VMs and Flows. The scope is based on the search query system. Based on your requirements, create a query from the available suggestions.
- 5 Create a condition for which the alert has to be created. Based on this condition, the system decides if the threshold is violated or not. Select the metric for which you want to create the condition. The default metric is `network traffic rate`. Select the grouping of the entity and the value for which you are checking the threshold. You can set a threshold on a cumulative metric by aggregating the data over a group of entities.

To configure the static threshold, select either of the following threshold conditions from the list:

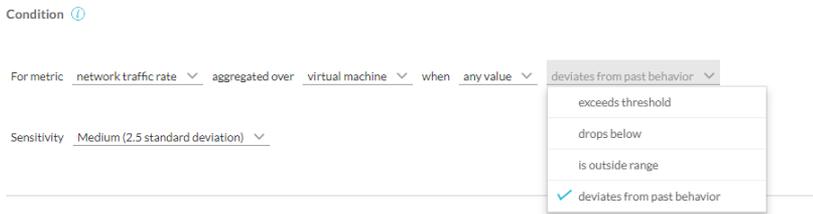
- exceeds threshold
- drops below
- is outside range

When you enter the `Upper Bound` or the `Lower Bound` (if there is range) for `network traffic rate` or `total traffic` or any other metric, ensure that you enter the value in the specified metrics for that particular text box. The following conversion values are for your reference:

- 1 Kbps= 1000 bps

- 1 Mbps= 1000 kbps
- 1 Gbps = 1000 mbps
- 1 KB=1024 B
- 1 MB=1024 KB
- 1 GB = 1024 MB

6 To configure the dynamic threshold, select **deviates** from the past behavior. Select the sensitivity based on your requirement of reporting.



7 When you set the threshold, you can view the associated graph at the top of the page. The pink bar denotes the VMs or the flows violating the threshold. You can view the list of the entities that have violated thresholds and the entities that are within the thresholds in the system.

8 Configure the notifications or alerts by setting the following properties:

- Severity
- Email frequency
- Send notification emails to:

Note Select **Send SNMP Trap** if you have configured SNMP traps on your system.

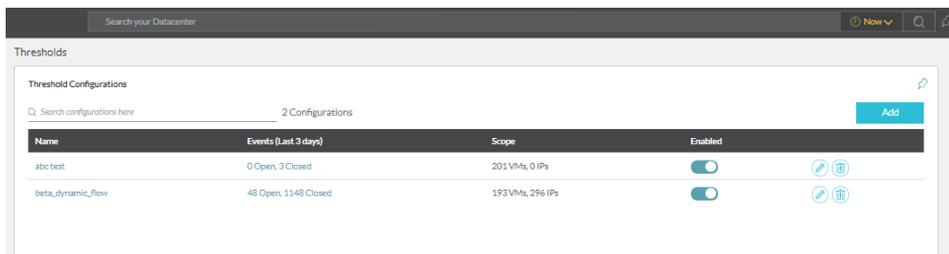
9 Click **Submit** to create the threshold configuration.

View the Threshold Configuration Page

Once you have added a threshold configuration, you can view its details on the **Threshold Configuration** page.

Procedure

1 On the left navigation panel, click **Analytics**. Click **Thresholds**.

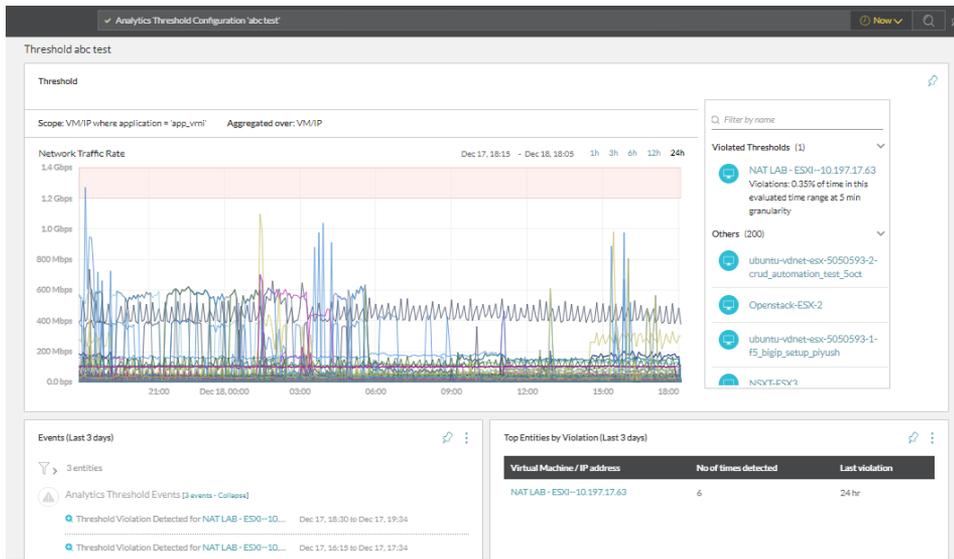


2 The following details for a threshold configuration are provided:

- Name
- Events
- Scope

If the configuration is disabled, then the alert for the violation of that particular threshold is not generated. You can also search for any particular threshold configuration on this page.

3 Click the desired threshold configuration from the list to view the dashboard for that particular configuration.



You can view the following widgets on the dashboard:

- Graph: The threshold graph helps you detect the entities that have violated the thresholds.
- Events: This widget provides the list of events that have been generated for violated thresholds for the last three days.
- Top Entities by Violation: This widget lets you know the top entities that have been the cause of aberrations for the last three days.

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.

This chapter includes the following topics:

- [Pins](#)
- [Pinboards](#)

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

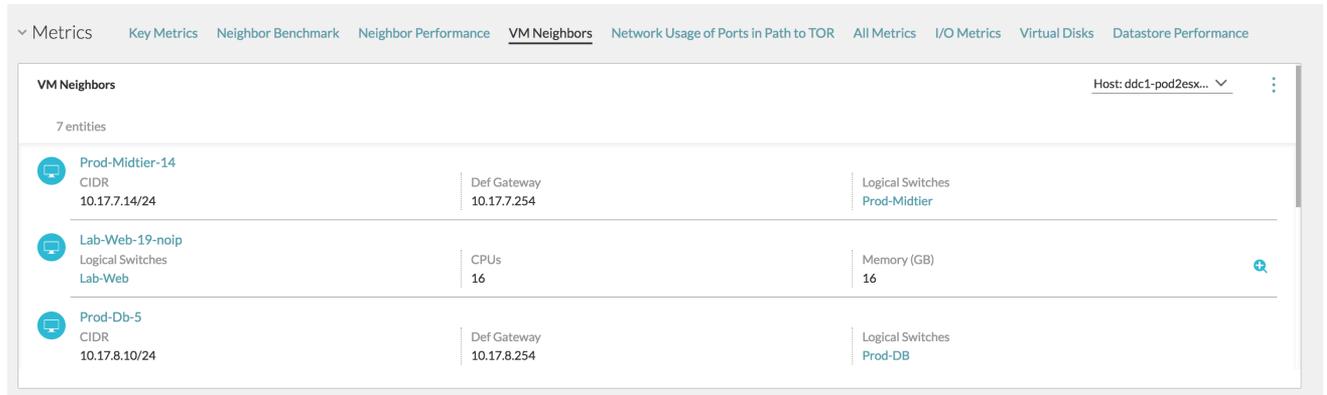
Most of the pins that are available in the software can be categorized into the following:

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.

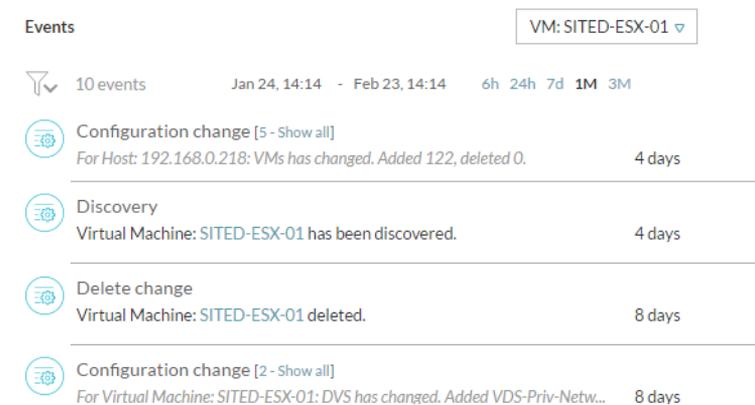


Event View List Pins

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.



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

You can pin any widget from any page on a pinboard to make it easier to access and share data.

To Create a Pinboard

The screenshot displays the 'PLAN' section of the vRealize Network Insight interface. At the top, there are three main categories: 'Micro Segments' (All Flows), 'Application' (Applications), and 'VM' (All VMs). Below these is a section for 'Open Problems' with a search filter showing 125 entities. The list of problems includes:

- Test1**: Event Search: vmware vm returned 8530 results (3 days)
- vms where Power State = 'POWEREDON'**: Event Search: vms where power state = 'poweredon' r... (3 days)
- vms where Power State = 'POWEREDOFF'**: Event Search: vms where power state = 'poweredoff' r... (3 days)
- Create Test Event - 0**

A 'Pin Options' pop-up window is visible on the right side of the screen. It contains a search bar labeled 'Search for pinboard', a 'Recently Modified' section with a list of pinboards: 'Default Pinboard', 'testing', and 'New board 1', and a 'Create New Pinboard' button at the bottom.

- 1 Click the pin icon on the widget that you want to pin.
- 2 Click **Create New Pinboard** in the pop-up window.

Note

- If you have not created any pinboard yet, you can select **Default Pinboard** from the **Recently Modified** list.

Note The Default Pinboard provides the look and feel of a typical pinboard to the first-time user. It helps the user to get familiar with the layout and features of a pinboard. It cannot be shared or deleted. You can copy pins from the default pinboard to any custom pinboard.

- The maximum number of entries that you can see in the Recently Modified list is 15.
- The maximum number of pinboards that you can create across all the users is 500.

Note The total number of pinboards include the custom pinboards, shared pinboards, and the default pinboards.

- The maximum number of pins per pin board is 20.

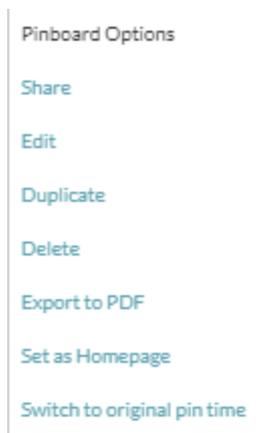
- 3 In the **Create pinboard** window, enter the name and the description for the new pinboard. Click **Create and Pin**.

Note

- The name of the pinboard has to be unique across the system.
 - The maximum number of characters allowed for the pinboard name is 100. You can use only letters, numbers, and spaces in the name of the pinboard.
-

- 4 The **Pinboard created** message appears. Click **Share Now** to share the pinboard immediately.
- 5 To pin the widget to an existing pinboard, select the pinboard under **Recently Modified** and click **Pin**. The message **Your Pin has been added** with the link to the respective pinboard appears.

To Access the Pinboard Options



Click **More Options** on the topmost right corner of a pinboard to access the **Pinboard Options**.

Note You can see all the pinboard options only if you have created the pinboard or if you have shared with any other user with the **View and Edit** permissions. Any other user can only see **Export to PDF** and **Switch to original pin time** options.

You can perform the following actions on the pinboard:

- You can share the pinboard with any other existing Network Insight user.
- You can edit the name of the pinboard and the pin on the pinboard.
- You can rearrange the pins on a pinboard. Their positions are persisted.
- Click **Delete** to delete that particular pinboard.
- Click **Export to PDF** to export the information on the pinboard as a PDF report. For more details, see [Export as PDF](#).
- To view the data on the pin at the time it was pinned, click **Switch to original pin time**. This feature enables you to view the data for each pin at the time it was created.

To Work with the Timeline Slider for a Pinboard

vRealize Network Insight supports a timeline slider on pinboards. To view the pinboard data for any desired time, you can use the timeline slider. When a pinboard loads, it loads all the pins for the current time (**Now**).

To View the Pinboard Library

If you are an admin user, you can see the **My Pinboards** tab and the **All Pinboards** tab in the Pinboard library as shown in the following image. If you are a member user, you can see a list of the pinboards in the Pinboard library.

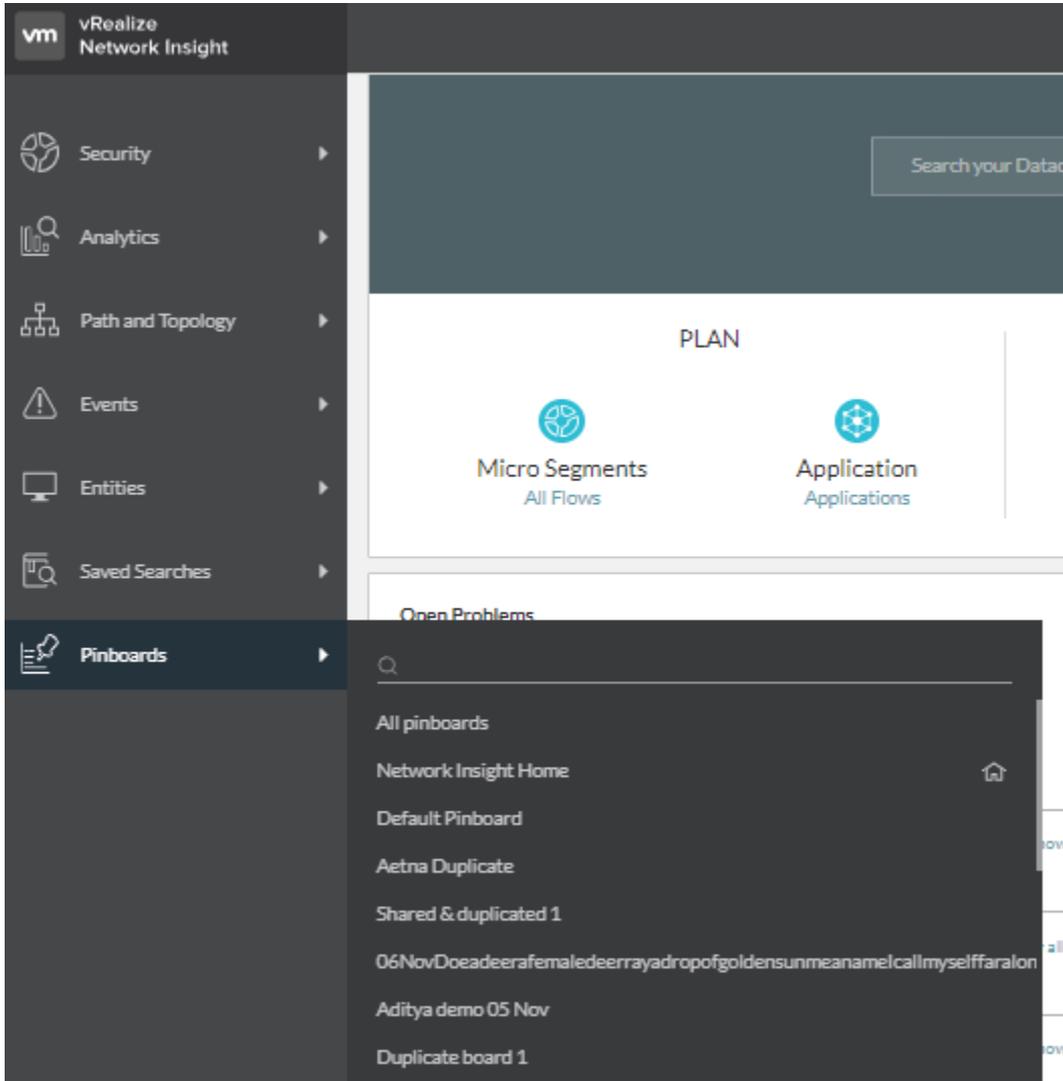
 Pinboards

Search for pinboards 17 pinboards

Pinboard name	Last modified	Owner	Shared	Actions
Network Insight Home 	--	--	--	
Default Pinboard	81 days	Guest 1	Not shared	
Aetna Duplicate	24 days	Guest 1	Not shared	  
Shared & duplicated 1	30 days	Guest 1	5 others	  

- 1 On the left navigation bar of the home page, click **Pinboards**.
- 2 Click **All Pinboards** to view all the pinboards in the system.
- 3 You can view the list of the existing pinboards in the navigation bar. The list has the same items as that of the **My Pinboards** tab in the pinboard library. The last modified pinboard appears at the top of the list. Click the pinboard that you want to view.

Note It takes some time for the pinboard to appear in this list after it is created.



4 You can also perform a search for a pinboard in the library.

To Copy a Pin

- 1 Click the pin icon on the widget.
- 2 Select the pinboard to which you want to copy the pin.
- 3 Click **Add**.

Sharing and Collaboration of Pinboards

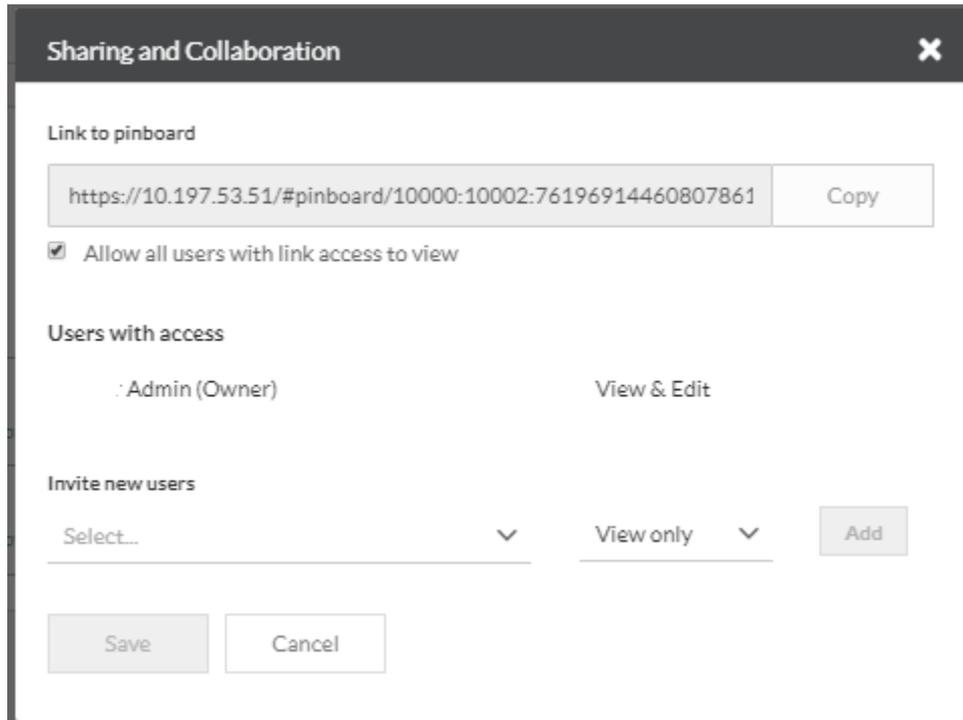
You can share the pinboards that you create with other users. An admin user can view and delete any pinboard. The following are the features of sharing and collaboration of pinboards:

If you have created a pinboard, you can view, edit, or delete it irrespective of you being an admin or a member user.

Table 8-1.

Pinboard Owner	Shared With	Privilege	Possible Action
Admin	Admin	View and Edit	View, Edit, Delete
	Admin	View only	View, Delete
	Member	View and Edit	View, Edit
	Member	View only	View
Member	Admin	View and Edit	View, Edit, Delete
	Admin	View only	View, Delete
	Member	View and Edit	View, Edit
	Member	View only	View

Note If a pinboard must be deleted and the user who created is not available, the admin user can delete it.



To share a pinboard:

Procedure

- 1 Click **More Options** on the pinboard that you want to share.
- 2 Click **Share**.
- 3 You can also share a pinboard from the **Pinboard Library** by clicking the share icon under **Actions**.

- 4 By default, the link sharing is enabled. You can share the link of a pinboard with any user who is logged in.
- 5 You can add the users with whom you want to share the pinboard. You can specify the privileges such as `view` and `view and edit` to a particular user.

Note The user who has only the view privilege cannot share the pinboard with any other user.

- 6 Click **Save** to save the share and collaboration changes that you have made.
- 7 You can view the sharing and collaboration information for any pinboard through either of the following options.
 - In the **Pinboard Library**, you can view the sharing information in the **Shared** column for a particular pinboard.
 - Click the pin icon on the widget. Point to any of the pinboards listed under **Recently Modified** to see the details regarding the owner and with whom it has been shared.

To Set A Pinboard as the Home Page

You can set a pinboard of your choice as your default home page.

Procedure

- 1 Navigate to the desired pinboard that you want to set as the home page.
- 2 Click **Pinboard Options**. Click **Set as Home Page**.

This particular pinboard is set as the home page.

Note Once you set a pinboard as the home page, the **Set as Home Page** option on that pinboard is disabled.

- 3 You can also set a particular pinboard as the default home page from the **My Preferences** page under **Settings**.
- 4 If you want to view the previous home page, then click **Network Insight Home** under **Pinboards** on the left navigation panel. The message **Do you want to set Network Insight Home as Homepage?** pops up. If you want to revert back to the default home page, click **Set Homepage**. Click **Dismiss** to close the message.

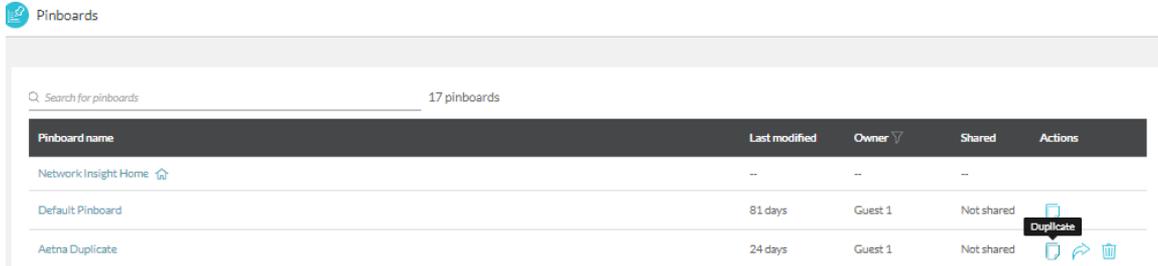
Note

- If you delete a pinboard that you have set as your home page, then the default home page is reset to **Network Insight Home** page. If you are the owner of the pinboard that you are deleting, a message pops up for confirmation for deletion.
 - If another user has set the pinboard that you have created as the home page, when you delete it, the home page reverts back to **Network Insight Home** automatically for that user.
-

To Duplicate a Pinboard

Procedure

- 1 Click the duplicate icon under **Actions** for the particular pinboard in the list in the pinboard library.



- 2 A pop-up comes up where you have to enter the name of the pinboard. The description is same as that of the original pinboard. Click **Duplicate**.

Note The name of the pinboard is mandatory. The **Duplicate** button is not enabled until you enter the name.

- 3 If you are trying to duplicate a pinboard that is shared, then you can opt to retain the source pinboard users and permissions. Select **Keep source pinboard users and permissions** if you want to retain them.

Note If the pinboard that you want to duplicate is shared with you with read-only access, you will not see the **Keep source pinboard users and permissions** option.

The user who duplicates a pinboard becomes the owner of the new pinboard.

Network and Security

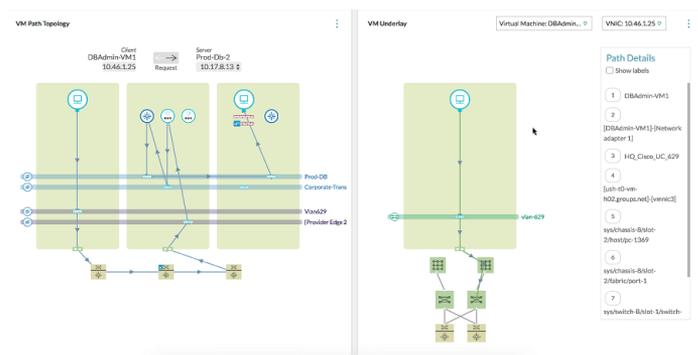
This chapter includes the following topics:

- 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.
-

AWS VM-VM Path

The VM-VM path for AWS provides the path visibility between the on-premises VMs and the AWS EC2 instances.

Currently, vRealize Network Insight supports the following scenarios:

- AWS intra-VPC VM-VM path: This scenario involves the communication between the VMs of the same subnet or different subnets in a particular VPC.
- AWS inter-VPC VM-VM path through the peering connection: This scenario involves the communication between the VM of one VPC to the VM of another VPC through a peering connection.
- AWS VM to Internet: The VM in a VPC communicates to Internet through the Internet Gateway.
- AWS VM to the data center VM through AWS VPN connection: In this scenario, the VM in a VPC communicates to the VM in a data center through the AWS VPN connection. vRealize Network Insight supports VMC SDDC and NSX-V and NSX-T data centers for this scenario.

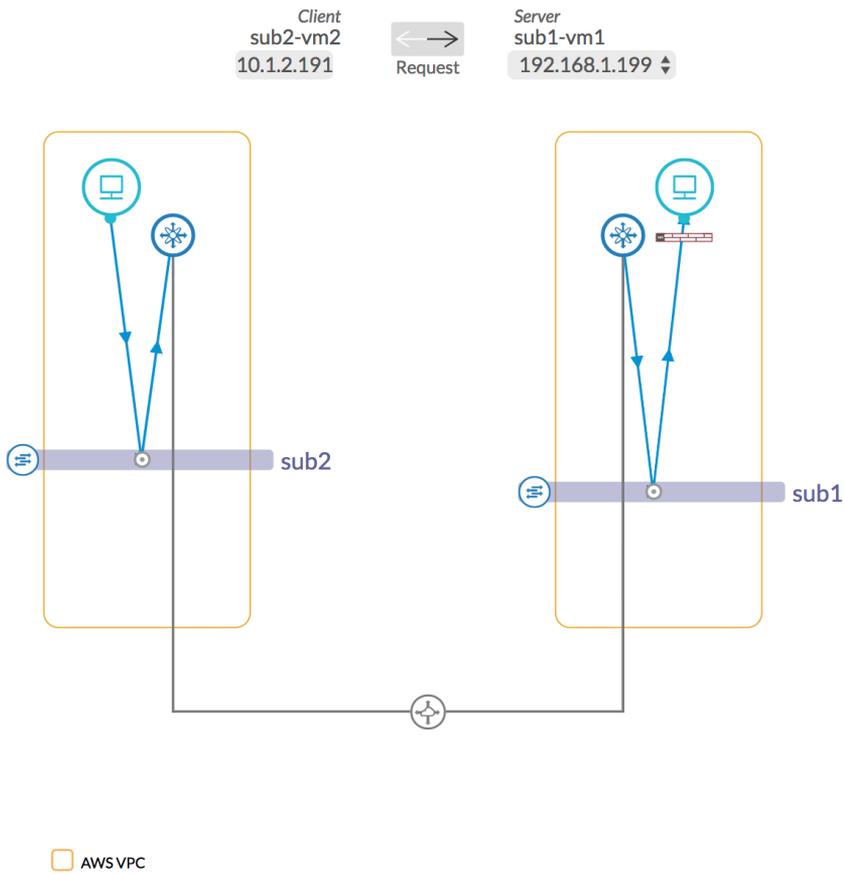
Note

- The hybrid path topology to NSX-T and NSX-V data centers works only when the NSX-T and NSX-V edge routers are configured with a public IP address.
 - vRealize Network Insight does not support the VM underlay topology for AWS.
-

Note

An example of the AWS VM-VM path for the AWS inter-VPC VM-VM path through the peering connection is as follows:

VM Path Topology i



You can view the properties of the peering connection by pointing to its icon in the VM-VM path.

You can search the following entities concerning the AWS VM-VM path:

- AWS Subnet
- AWS Route Table
- AWS Virtual Private Gateway
- AWS Internet Gateway
- AWS VPN Connection
- AWS VPC Peering Connection

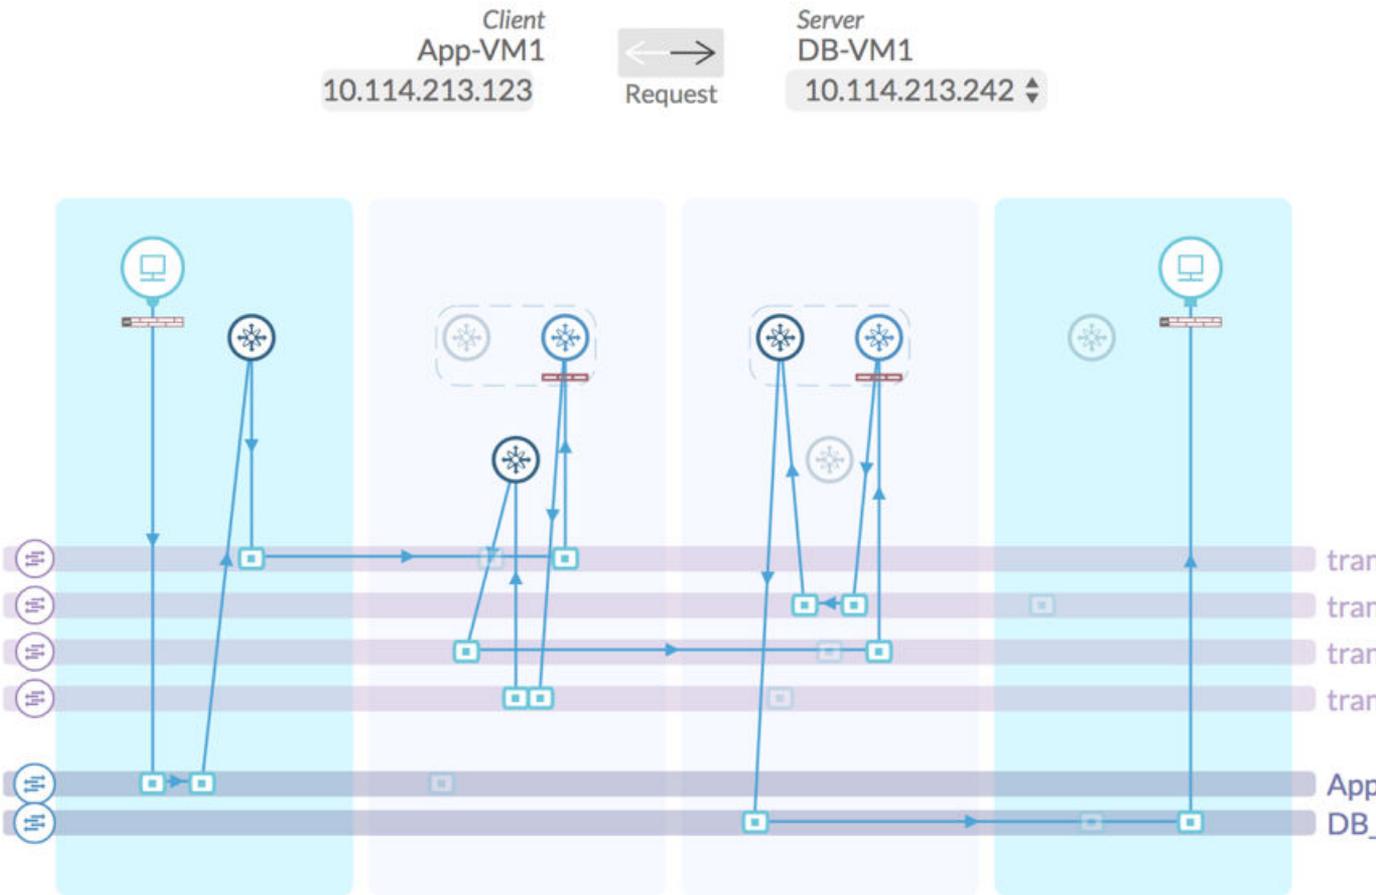
NSX-T

An example for VM-VM path for NSX-T is as follows:



Overview

VM Path Topology i



The blue color represents the host node and the grey color represents the edge node. The icons used in the VM path topology are listed on the right side of the screen along with the labels under Path Details. The distributed routers are shown in the same color irrespective of their tiers. The color of the service router in the topology diagram changes as per the associated tier. All the tier 1 components are shown at the same level and all the tier 0 components are shown at a different level. In NSX-T, the edge firewalls are depicted in the diagram.

To plan security for the NSX-T network, you can select the scope as **NSXT Layer2 Network** and use the following query:

```
plan NSX-T Layer2 Network '<NAME_OF_NSX_T_LOGICAL_SEGMENT>'
```

You can also obtain the same result by performing the following steps:

- Select **Security** from the Navigation side bar.
- Select **NSX-T Layer2 Network** as the scope from the drop-down menu.

Note

- NSX-T related entities such as **NSX-T L2 Network** and **Tags** are available in the scope. You can use these NSX-T related entities in planning, micro-segmentation, and application definition.
 - In The **Group by** drop-down menu, **NSX-T Security Group** is a part of **Security Tag** and **Logical Segment** is part of **VLAN/VXLAN**.
-

VMC: VM-VM Path

vRealize Network Insight supports the following hybrid paths in VMC:

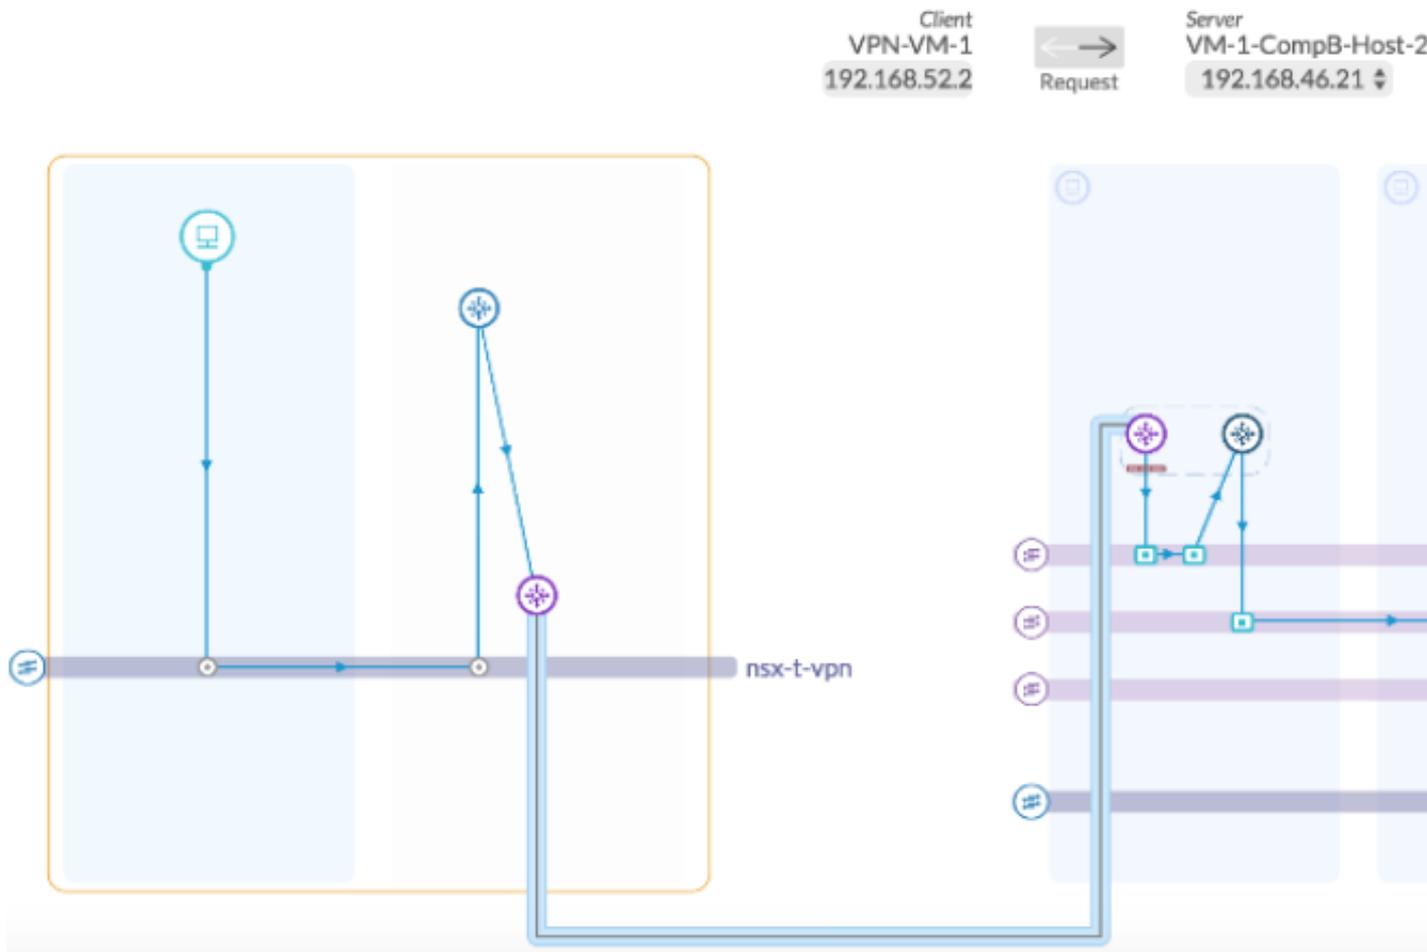
- VMC and VMC
- VMC and NSX-T
- VMC and NSX-V
- VMC and AWS
- Intra VMC

For all VMs present in VMC, the underlay information is shown only until the segment on which the VM lies because the underlying physical elements of the network are abstracted out by VMC and no visibility is present at that level.

A sample VMC and NSX-T VM-VM path is as follows:

NSX-T NAT: Should go through the tunnel and Tier-1 NAT

VM Path Topology ?



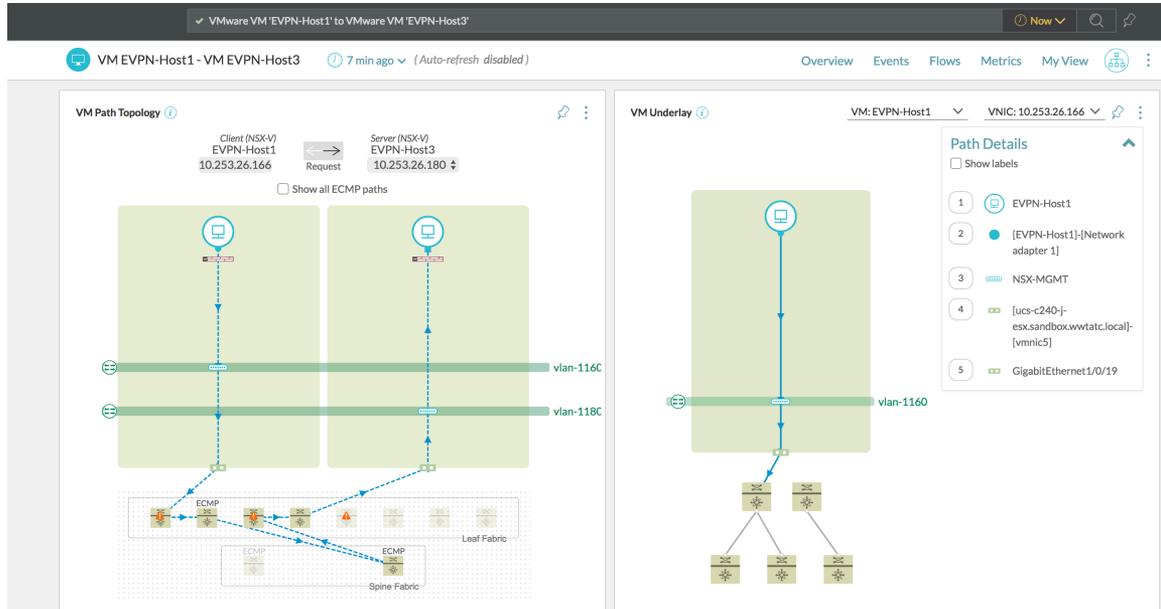
The dark blue line depicts the tunnel.

Support for the Cisco BGP-EVPN Mode

vRealize Network Insight supports the fabric of Cisco 9000 switches configured in the Cisco BGP-EVPN configuration mode for the Enterprise edition only. vRealize Network Insight does not support the switch models other than Cisco Nexus 9000 with the Cisco BGP-EVPN configuration.

Each Cisco Nexus 9000 switch that is a part of the fabric is individually added as a data source. To view all the spine or leaf switches in the fabric, use the switches where role is set query.

A sample VM-VM path for the Cisco BGP-EVPN mode is as follows:

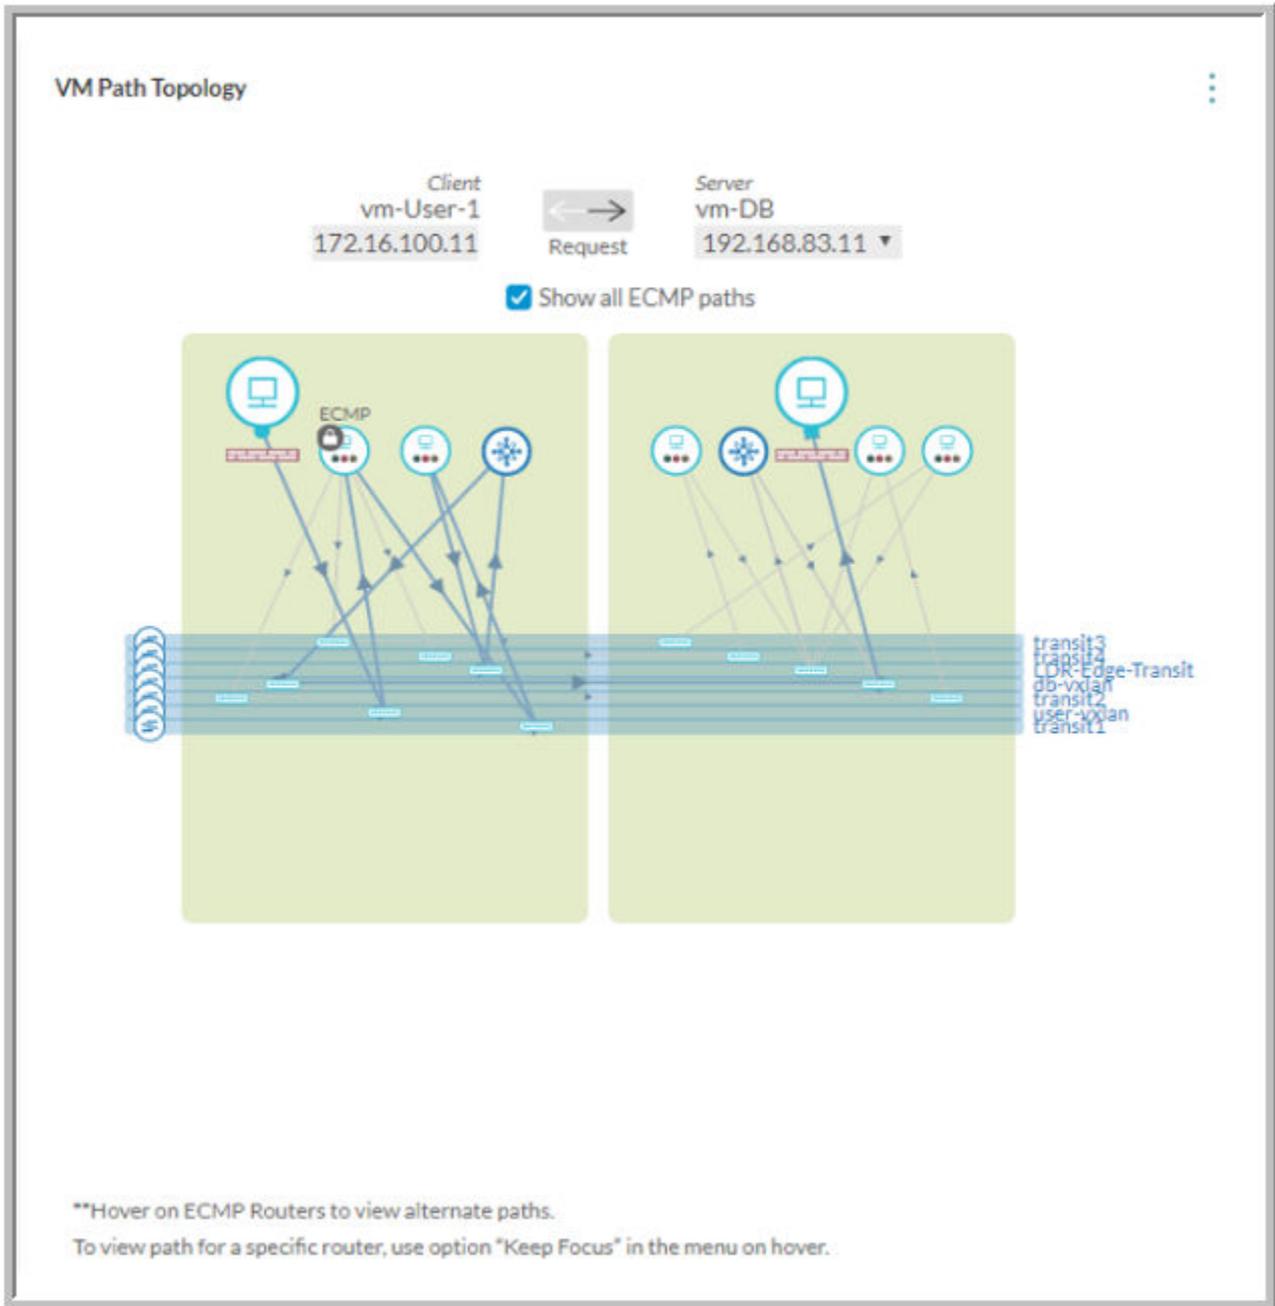


Support for Equal-Cost Multi-Path (ECMP) Route

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



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.

Monitoring Various States of BGP

vRealize Network Insight supports the monitoring of the states of BGP. You can view the neighbors of BGP for an NSX edge or a logical router.

Procedure

- 1 Enter `Routers` in the search bar.
- 2 To view the results of particular NSX manager, filter by selecting `NSX Manager` from the left panel,
- 3 Expand the particular router from list to see the details.
- 4 You can view the following information under **BGP neighbours**.
 - IP Address
 - Remote AS
 - Weight
 - Keep Alive Time
 - Hold Down Time
 - Status

Note

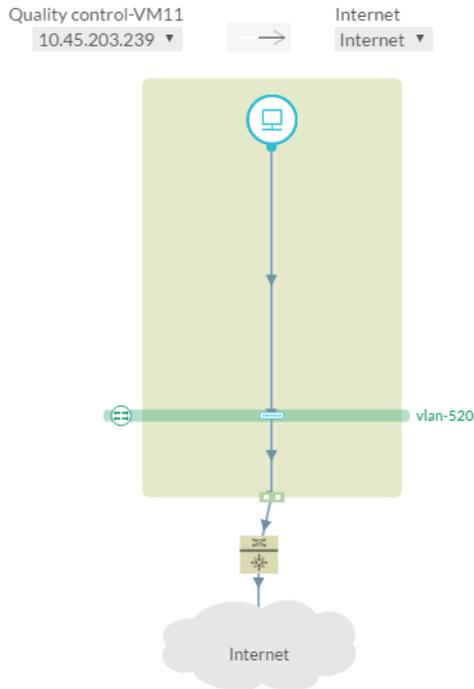
- If the information regarding the neighbors is not fetched, then the `Status` is shown as `Unknown`.
 - If the `Status` is not `Established.up`, then the `One or more BGP neighbours are not in established state` event is raised for that edge. You can also view this event when you search for problems.
-

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

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 the Palo Alto Panorama firewall.

Note vRealize Network Insight does not support the Palo Alto Panorama integration with multiple NSX managers.

To add the Palo Alto Panorama in vRealize Network Insight, the Palo Alto Networks user must have **admin role** with XML API access. Do the following steps to add an admin role for XML API.

- 1 Select **Panorama > Admin Roles**.
- 2 Click **Add** to add a new admin role.
- 3 The Admin Role Profile window opens.
- 4 Enter the name to the role and select **Panorama**.
- 5 Click the **Web UI** tab and disable all entries.
- 6 Click the **XML API** tab and disable all entries, except **configuration** and **Operational Requests**.
- 7 Click **OK** to close the window.
The new admin role appears in the list.
- 8 Click **Commit**.
- 9 Assign this role to an administrator account or create a new user and assign this role to the new user.

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 = <>

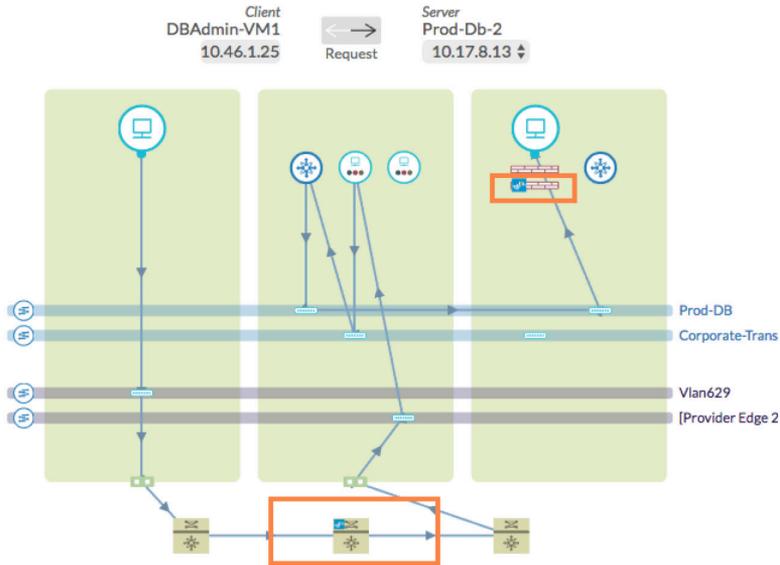
- Query: You can perform a query for all the Palo Alto entities that are supported by vRealize Network Insight. All the entities are prefixed by Palo Alto. Some of the queries are as follows:

Table 9-1.

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 = <>
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

Timeline

Time Range: 1 day 12:00 16:00 20:00 10/26 04:00 8:28

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 connect... 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 L... 65 days
- Service VM's status mismatched between Panorama and NSX Manager**

Palo Alto Virtual Devices

16 entities

IP	Manager	Device Family	Name
10.16.128.13	10.16.128.200	VM	10.16.128.13
10.16.128.14	10.16.128.200	VM	10.16.128.14
10.8.201.3	10.16.128.200	VM	10.8.201.3

Palo Alto Physical Devices

1 entity

IP	Manager	Device Family	Name
10.16.21.2	10.16.128.200	PA-5060	10.16.21.2

Clusters prepared for Palo Alto services

4 entities

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

Hosts deployed with Palo Alto Device

17 entities

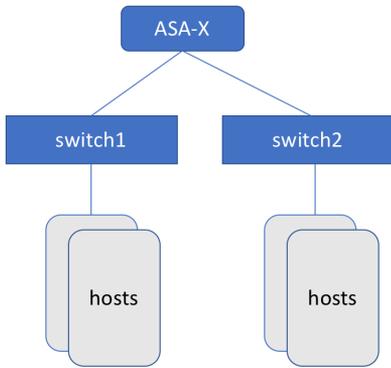
Host	Number of VMs	Power State	Version
ddc1-pod2esx046.dm.democompany.net	3	On	6.0.0
ddc1-pod2esx040.dm.democompany.net	3	On	6.0.0
ddc1-pod2esx037.dm.democompany.net	2	On	6.0.0

Cisco ASA Firewall

vRealize Network Insight supports Cisco ASA firewall.

The features for Cisco ASA firewall are as follows:

- 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:



- Cisco ASA access rules of only Extended type are supported. Other access rule types like Standard, WebType, EtherType, and so on are not supported.
- The Cisco ASA firewall in the VM-to-VM path does not display applicable access rules if the firewall is configured in the Transparent mode.

To add Cisco ASA as data source:

Procedure

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

Table 9-2.

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 this data source. The user should have the enable mode privilege to set terminal length to 0, and for switching security context.
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**.

Example

You can perform a query for all the Cisco ASA entities that are supported by vRealize Network Insight.

Table 9-3.

Entities in Cisco ASA	Keywords	Sample Queries
Security Context	ASA Firewall ASA Security Context	<code>asa firewall where access group = <></code>
Access Rule	ASA Access Rule	<code>asa access rule where source ip = <></code> <code>asa access rule where destination ip = '192.168.2.2'</code> <code>asa access rule where port = <></code> <code>asa access rule where interface = <></code>
Access Group	ASA Access Group	<code>asa access group where interface = <></code>
Network Object / Network Object Group	ASA Network Object ASA Network Object Group	<code>asa network object where ip address = <></code> <code>asa network object group where ip address = <></code>
Service Object / Service Object Group	ASA Service Object ASA Service Object Group	<code>asa service object where port = <></code> <code>asa service where protocol = <></code> <code>asa service object group</code>

Check Point Firewall

vRealize Network Insight supports the following Check Point Management Servers:

- Check Point Security Manager (SmartCenter)
- Check Point Multi-Domain Manager (MDS / Provider-1)

If Check Point MDS is added as data-source, vRealize Network Insight fetches data from all the user-defined domains and the global domain.

vRealize Network Insight uses Check Point public Web API for fetching the data from the Check Point management server. If the VSX gateway is attached to the management server, we use SSH-based CLI commands to fetch the VSX-managed Virtual System VS routing table to support display of the VS gateway in the VM-VM path.

vRealize Network Insight requires read-only privileges for the Web-API access for fetching most of the Check Point data. There are few exceptions as follows:

- If a non-VSX physical gateway is attached to the management server, the user should have read-write access privileges for the Web API. This is required to fetch the gateway routes for using the `run script` Web API for the VM-VM path computation.
- If a VSX gateway is attached to the management server, the user should have the SSH access with the same password. In addition, the user should have access to the CLI command `vsx_util view_vs_conf`. This command is used to fetch the VSX gateway routes for the VM-VM path computation.

- For MDS server IP as data-source, the user should have the Web API access to all domains including the MDS domain and the global domain. It is required to fetch rules, policy packages and other data from all the domains.

Note

- vRealize Network Insight supports Check Point firewall version R80 and further versions.
 - For the VM-VM path, vRealize Network Insight does not support the VSX cluster containing the virtual switch and the virtual router.
-

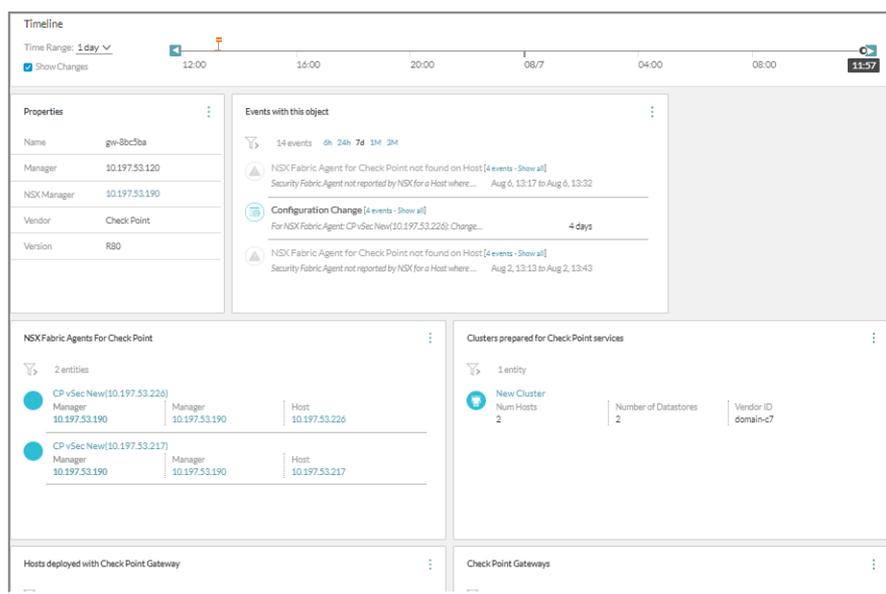
You can perform a query for all the Check Point entities that are supported by vRealize Network Insight. All the entities are prefixed by Check Point. Some of the queries for Check Point are as follows:

Table 9-4.

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 = <>
Domain	Check Point Domain	check point domain where ip address = <> check point policy where domain = <> check point access layer where domain = <>
Gateways and Gateway Cluster	Check Point Gateway Check Point Gateway Cluster	Check Point Gateway Cluster where Policy Package = <>

Entities in Check Point	Keywords	Queries
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)

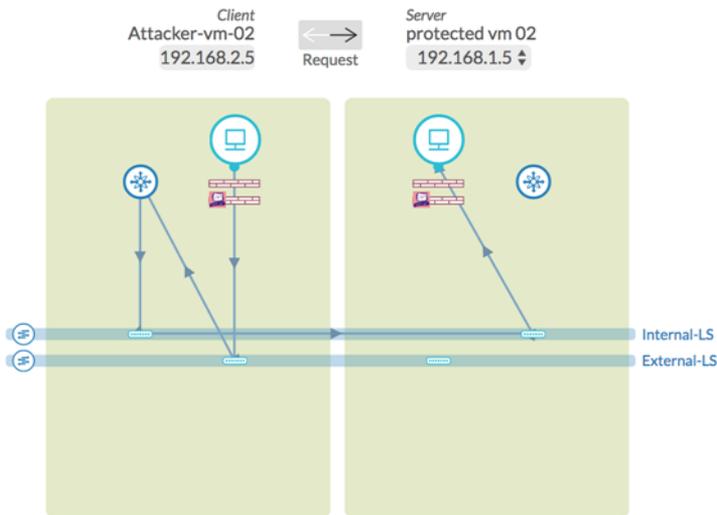
A sample Check Point Manager dashboard is shown as follows:



In a VM-VM topology diagram, you can see the Check Point Service VMs on a host to signify the Check Point rules applied on the particular traffic. The VSX-managed Virtual System (VS) gateway can be seen in the VM-VM path as a physical gateway. The list of applicable Check Point policies is displayed when you click the gateway icon.

Note For the VM-VM path, vRealize Network Insight does not support the VSX cluster containing Virtual Switch and Virtual Router.

VM Path Topology



Here are some scenarios for which the system events are generated for Check Point:

- The NSX fabric agent is not found on the ESX for the Check Point gateway.
- The Check Point service VM is not found.
- The Check point gateway sic status is not communicating.
- The discovery and update events features for the Check Point entities like address range, networks, policies, groups, policy package, service, service group, and so on

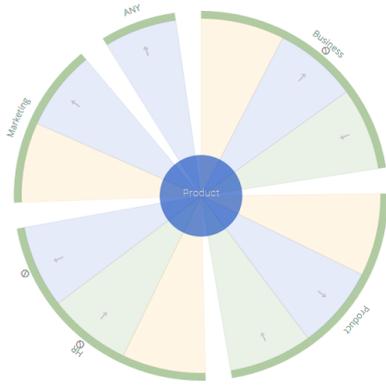
Security Groups

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:

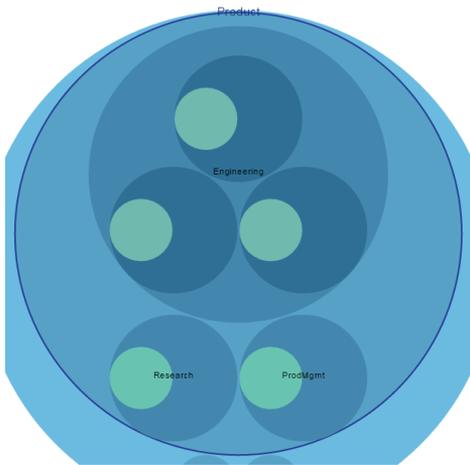
Firewall View

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).



Policy-Based VPN

vRealize Network Insight supports policy-based VPN in VMC, NSX-T, and NSX-V. The following scenarios are supported for the policy-based VPN:

- VPN tunnel between the VMC public IP address and NSX-V/NSX-T/AWS public IP address
- VPN tunnel from the VMC public IP address and the corporate firewall public IP address to a 1:1 NAT between the corporate firewall public IP address and the internal NSX Edge

Note vRealize Network Insight does not support the scenario of the VPN tunnel from the VMC ending on a corporate firewall and no NAT configured with the internal NSX edge.

Policy-Based VPN Entities

vRealize Network Insight fetches data for the L3 `VPN Session` entity which is the actual VPN configured in the data center.

Here are the search terms for the policy-based VPN entities:

Table 9-5.

Search terms	Description
Policy based VPN	All policy-based VPN sessions for VMC, NSX-V, and NSX-T
VMC Policy based VPN	VMC policy-based VPN sessions
NSX-T Policy based VPN	NSX-T policy-based VPN sessions
NSX Policy based VPN	NSX policy-based VPN sessions

NSX Distributed Firewall Inactive Rules

vRealize Network Insight supports the visibility of the NSX distributed firewall rules for which there have been no flows for some time. These rules are known as inactive rules. Such rules use memory heap and can cause security issues. To monitor these inactive rules, vRealize Network Insight provides the following two widgets in the **Security** dashboard:

Note To view the Security dashboard, enter **Security** in the search bar.

- **Unused NSX Firewall Rule:** This widget lists all the NSX firewall rules where no flow is reported on the given time. You can also use the following search query to retrieve these rules:

```
nsx firewall rule where flow is not set
```

Note Ensure that you have enabled NSX Distributed Firewall IPFIX for the specified time.

This chapter includes the following topics:

- [Working with Events](#)
- [Notifications](#)

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 10-1.

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: <ul style="list-style-type: none">■ Critical■ Moderate■ Warning■ Info
Type	This field specifies if the event denotes a Problem or a Change . Note All the events of type Problem are logged into syslog.
Entities	This field specifies that the event is configured to either include or exclude entities for event generation. By default, the value is ALL.

Column	Description
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 10-2.

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: <ul style="list-style-type: none"> ■ 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.

Field	Description
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 10-3.

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.

Flows

This chapter includes the following topics:

- [Enabling IPFIX Configuration](#)
- [Flow Support for Physical Servers](#)
- [View Blocked and Protected Flows](#)
- [Network Address Translation \(NAT\)](#)
- [VMC Flows](#)
- [Create VPC Flow Log](#)

Enabling IPFIX Configuration

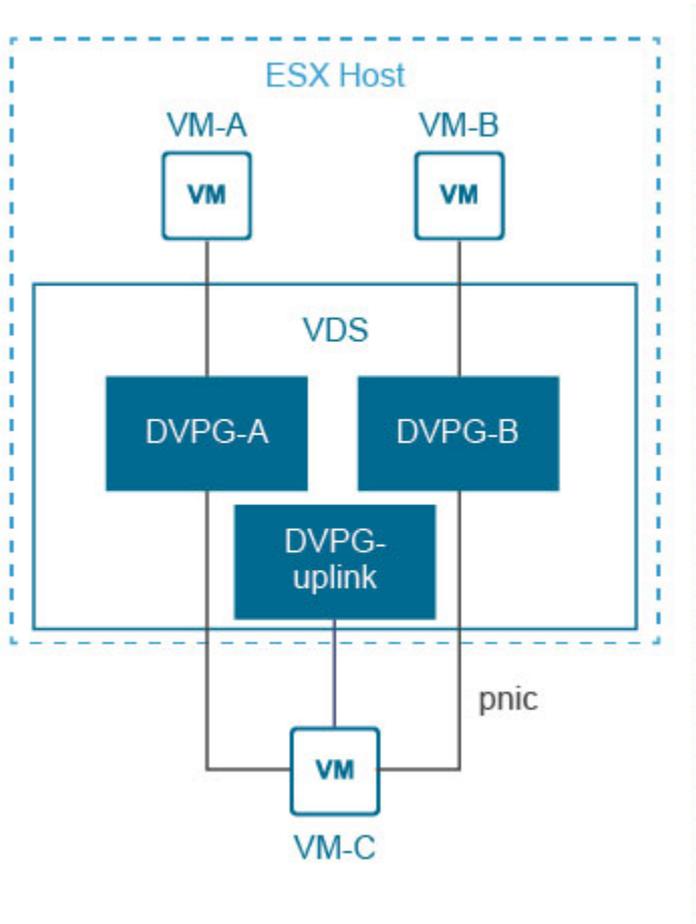
IPFIX is an IETF protocol for exporting flow information.

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.

Enabling VMware NSX-T DFW IPFIX

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

Prerequisites

- Ensure that you have any one of the following privileges:
 - enterprise_admin
 - network_engineer

- security_engineer
- Ensure that the Distributed (DFW) firewall is enabled.
- Ensure that priority 0 is available for the Network Insight IPFIX profile. If there is another IPFIX profile with priority 0, then you have to change it to some other value.

Procedure

- ◆ Select **Enable IPFIX** when adding or editing an NSX-T Manager data source.

What to do next

After you enable IPFIX, vRealize Network Insight creates its own Network Insight Collector profile and Network Insight IPFIX profile on NSX-T. Ensure that you do not modify any of these profiles.

After enabling IPFIX on NSX-T, if the flows are not seen in vRealize Network Insight, then the following events may occur:

- Network Insight Collector Profile is not registered in the NSX-T Manager.
- Network Insight IPFIX Profile is not registered in the NSX-T Manager.
- Network Insight IPFIX Profile port number has changed.
- Network Insight Collector Profile does not match in the Network Insight IPFIX profile in the NSX-T Manager.

Note To resolve all the above issues, enable NSX-T IPFIX again.

- Network Insight IPFIX Profile priority is not zero in the NSX-T Manager.
To resolve this issue, log into NSX-T Manager and set the priority of Network Insight IPFIX Profile to zero.
- Network Insight Collector IP cannot be added in existing Network Insight Collector Profile in the NSX-T Manager.
Delete one of the collectors from the Network Insight Collector Profile in the NSX-T Manager and re-enable NSX-T IPFIX from data source page.
- Distributed Firewall is disabled in NSX-T Manager.
Log into NSX-T Manager and enable the DFW firewall.

Note All the logical switches present in NSX-T are appended in the IPFIX profile within 10-15 minutes.

Adding a NetFlow Collector

Procedure

- 1 In the **Settings** page, click **Accounts and Data Sources**.
- 2 Click **Add Source**.

- 3 Under **Flows**, 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.
- 4 Enter **Nickname** and **Notes** as required. Click **Submit**.

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

Example

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:

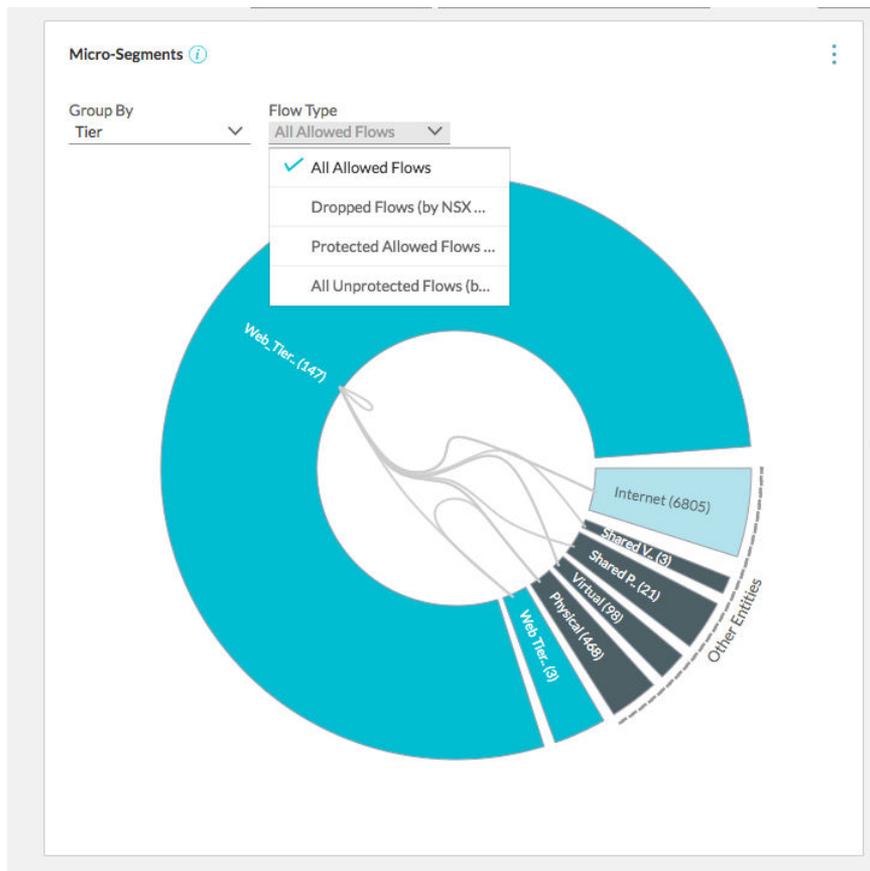
- All Allowed Flows: This option is selected by default. To see all the flows for which the action in the firewall rules is set to **Allowed**, select this option.
- Dropped Flows: This option helps to detect the dropped flows and planning the security in a better way.
- All Protected Flows: This option helps to detect all the flows which have a rule other than of the type any(source) any(dest) any(service) 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(service) 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)

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

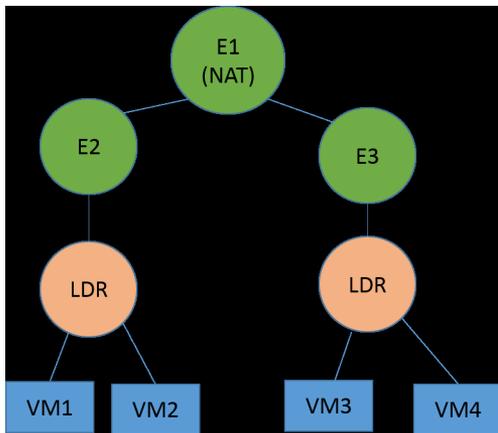
- Currently, vRealize Network Insight supports SNAT, DNAT, reflexive rules in the flows and the VM to VM Path for the NSX-V and NSX-T edges only.
- To obtain all the NAT rules in NSX-T, use the NSX-T Edge NAT Rule query. To obtain all the NAT rules in both NSX-V and NSX-T, use the NAT Rules query.

- Only the NAT rules that are configured on the uplink interface of the VMware NSX-T Tier router are processed by the VM to VM path. If NAT is configured on any NSX-T Tier router, then it is expected that there are NAT rules for all the VMs attached to the router else the VM to VM path and the path to Internet does not work. Instead, it displays a missing rule message.
- vRealize Network Insight supports the nested NAT hierarchy.
- vRealize Network Insight supports the edges and the tier routers with NAT-defined uplinks.
- vRealize Network Insight supports SNAT rules with range. However, DNAT must be one-to-one mapping between the destination and translated IP addresses (Parity with NSX-V).
- vRealize Network Insight does not support the following use cases:
 - a In NSX-T, NAT rules can be applied at the service level. For example, in NSX-T, L4 ports set is a type of service and the associated protocols can be TCP or UDP. So in the VM-VM path, the service level details are not supported.
 - b Any port level translation is not supported.
 - c The SNAT match destination address and the DNAT match source address are not supported. Use the SNAT match destination address as the destination IP address when you specify the SNAT rule. Use the DNAT match source address as the source IP address when you specify the DNAT rule. For example, if there is a destination IP address mentioned in the SNAT rule, vRealize Network Insight applies the SNAT rule irrespective of whether the packet has the destination address as the destination IP address.

NAT Flow Support - Examples

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

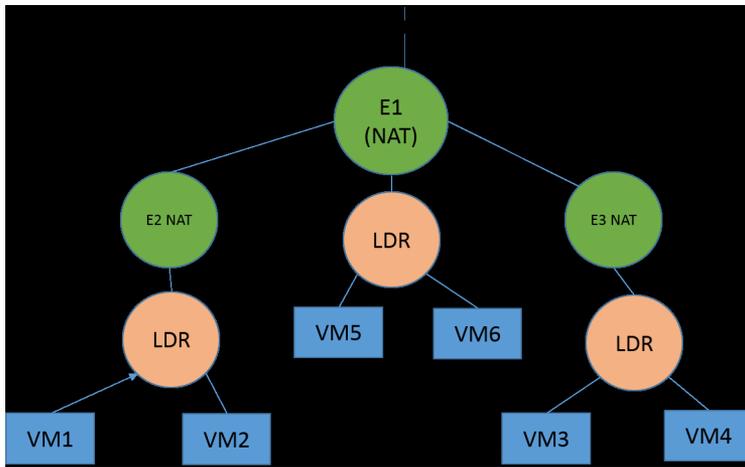
Example 1



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)

VMC Flows

If you have enabled IPFIX on the data source in the **Settings** page, you can view the flow count and the last collection time.

You can search for any particular flow and get the details associated with the entities. For example, you can view the policy segment and the policy group information in Source L2 Network and Source Security Group respectively. You can also view the policy firewall rule attached to flow.

vRealize Network Insight supports the hybrid flows over the VPN. The flow information is enriched with the source and the destination entities.

10 Create an IAM role.

- a In the **VPC Flow Logs is requesting permission to use resources in your account** page, in the **IAM Role**, select **Create a new IAM Role**.
- b In the **Role Name** text box, enter a role name.
 - a Click **Allow**.

11 On the **Create flow log** page, in the **IAM role** drop-down, select the role you created.

12 Click **Create**

Flow log starts publishing on the selected log group. For more information about VPC Flow Log, see the AWS documentation at <https://docs.aws.amazon.com/vpc/latest/userguide/flow-logs.html#create-flow-log>.

Micro-Segmentation

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](#)
- [Analyzing the Application](#)
- [VMC: Planning and Micro-Segmentation](#)
- [Create an Application Manually](#)
- [VMC: Planning and 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.

Analyzing the Application

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

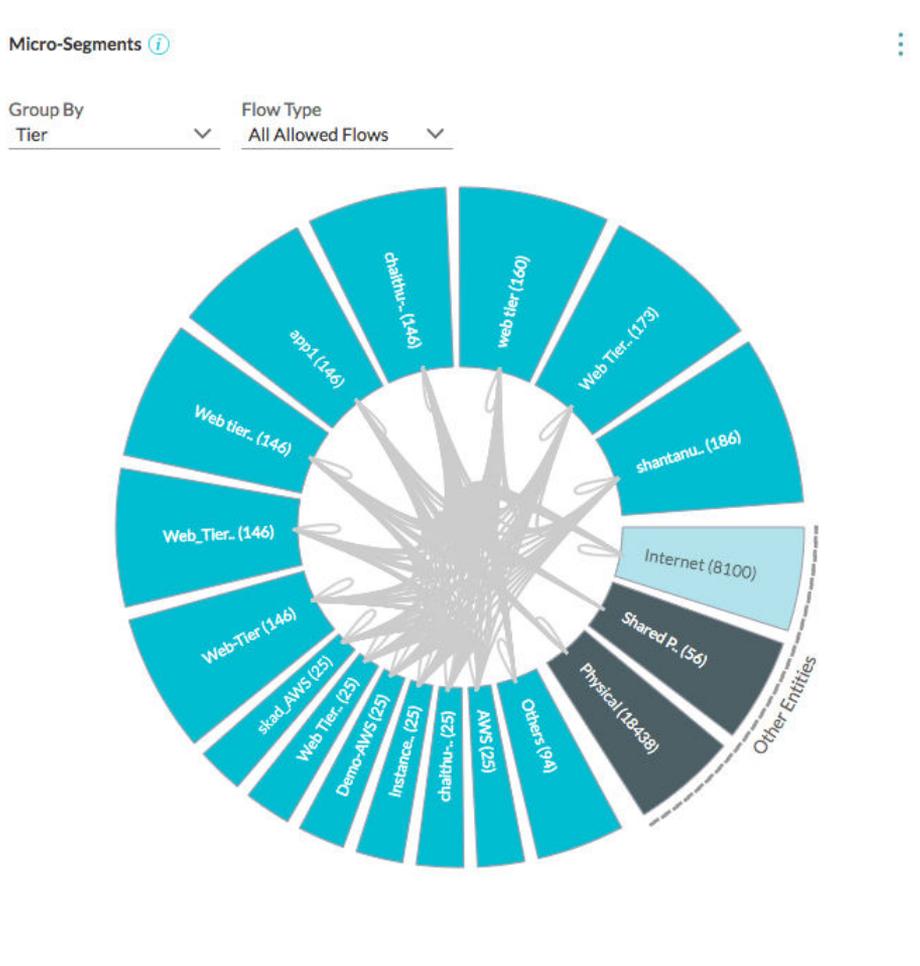
In vRealize Network Insight, a flow is a 4-tuple. It includes:

- Source IP
- Destination IP
- Destination port
- Protocol

You can view the data in two formats: Donut View and the Grid View

Viewing Micro-Segmentation And Flow Data in Donut View

In the Donut view, 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.



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.



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.
-

View Micro-segmentation And Flow Data in Grid View

vRealize Network Insight enables you to see the communication between objects in a tabular or grid view.

Procedure

- 1 Navigate to **Security > Plan Security** and click the grid view  icon.

- 2 Select a value for **Group By** option, for example **VMs, Application, Security Groups** to see the corresponding details in the tabular format.

Field Name	Description
Source Object	Name of the source
Destination Object	Name of the destination
Related Flows	Count of communication or flows between the source and destination Click on the count value to see the related flow details.
Sum of Bytes	Aggregated number of bytes between all the flows
Max of Traffic Rate	Maximum traffic rate observed among all related flows
Count of Sessions	Number of active sessions for the particular flow

Note

- You can click on each column header to sort the data in ascending or descending order.
- You can hide the field from the table view, click the more icon next on the field header and deselect the field name.

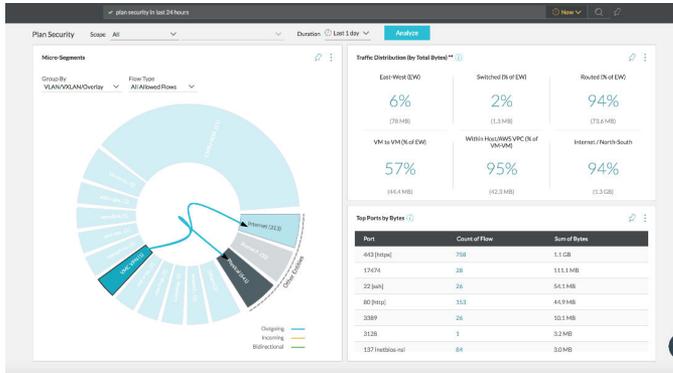
- 3 In addition, you can perform several actions on the grid view page.

- In the filters pane on the left side of the screen, you can perform the following actions:
 - Select an individual source or destination to filter the flows that are related to selected source or destination object.
 - Select the firewall action to see the allowed flows or the dropped flows.
 - Select the protection status to see the flow status.
- Click **Add more filters** to add additional filters.
- To export the tabular data in a CSV format, click the more option on top of the table, and select **Export as CSV**.

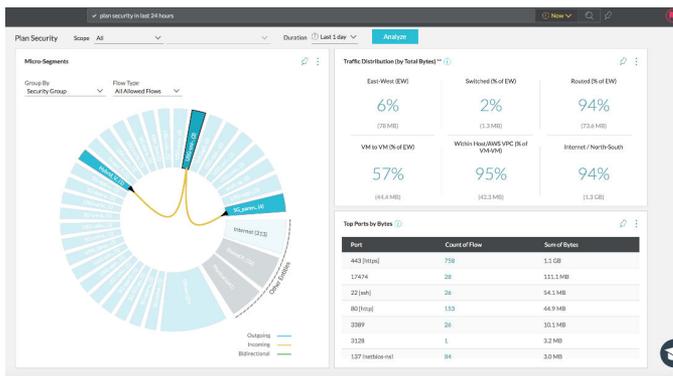
VMC: Planning and Micro-Segmentation

You can plan for a specific VMC segment by selecting **VMC Segment** as the scope in the **Plan Security** page.

For the policy segments, use the `VLAN/VXLAN/Overl`ay clause in the group.



For the policy groups, use the Security Group clause in the group.



Create an Application Manually

To add an application:

Procedure

- 1 In the Search box, type application, and press Enter.
- 2 Click the **Saved Applications** tab and 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.

Creating Tiers for Physical IPs

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](#).

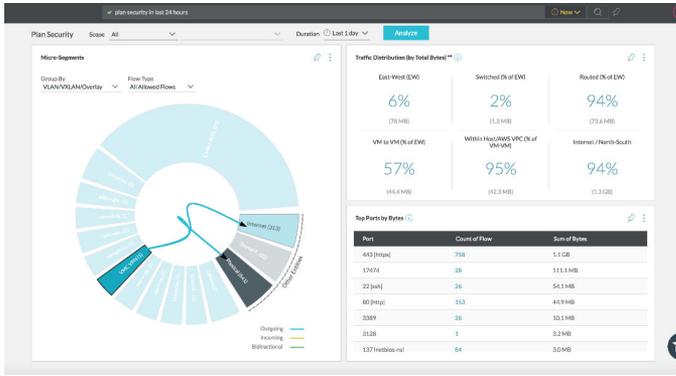
The enriched DNS, Subnet, VLAN information can be used in specifying tiers as follows :

- 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

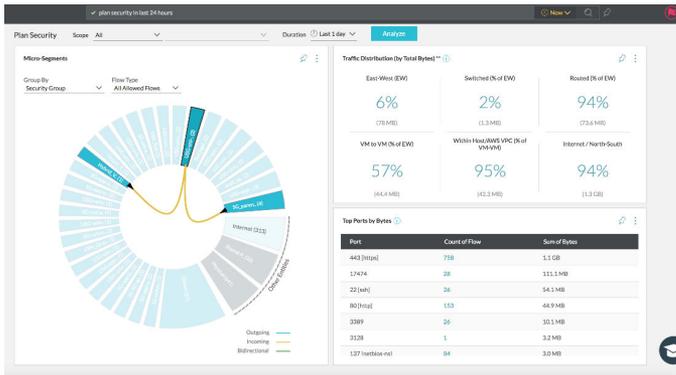
VMC: Planning and Micro-Segmentation

You can plan for a specific VMC segment by selecting **VMC Segment** as the scope in the **Plan Security** page.

For the policy segments, use the VLAN/VXLAN/OverlAy clause in the group.



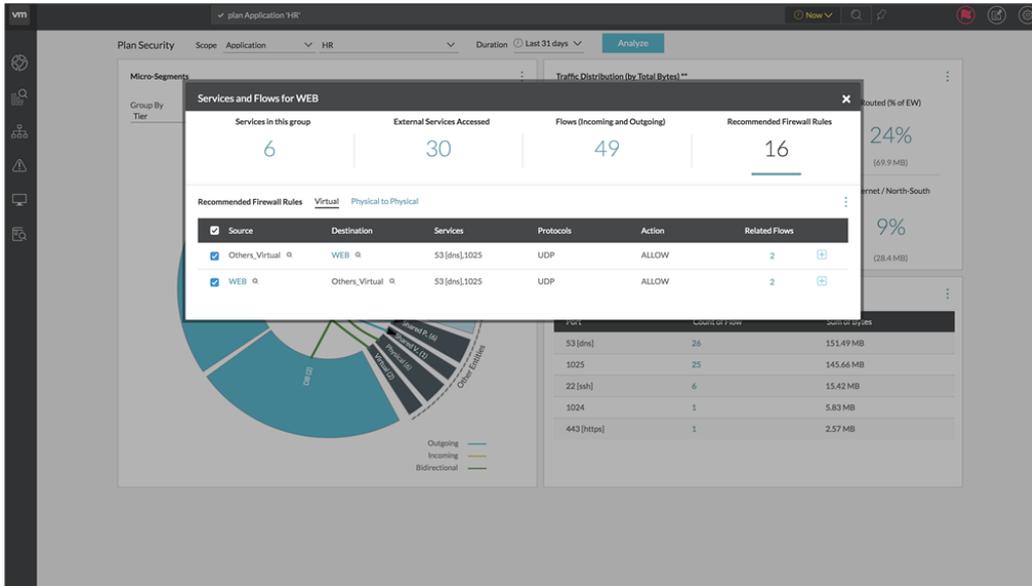
For the policy groups, use the Security Group clause in the group.



Recommended Firewall Rules

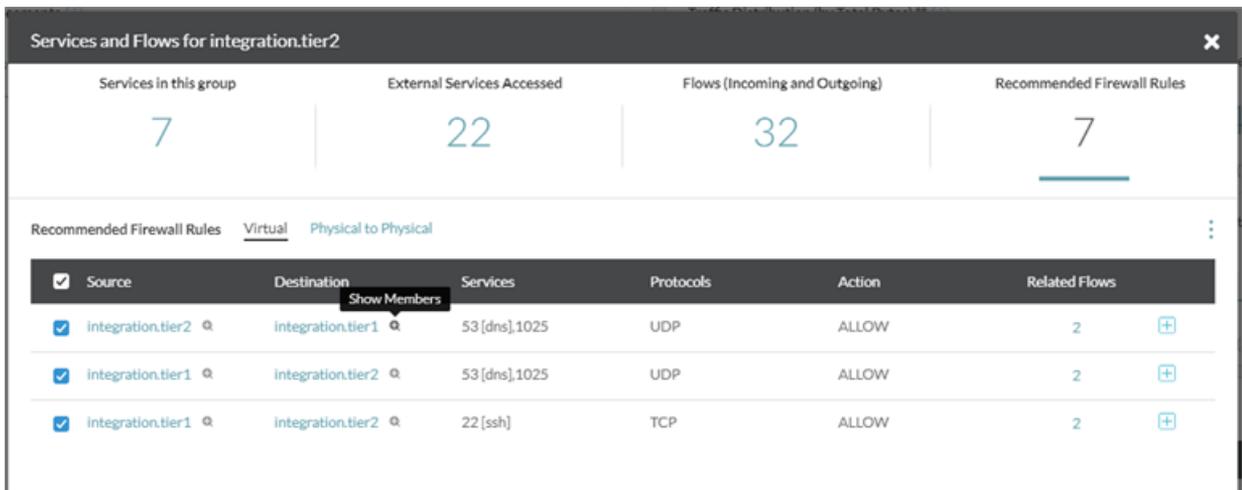
On the **Plan security** page, when you click the wedge or the edge in the topology diagram, you can view the list of the services and flows for that particular segment. Click **Recommended Firewall Rules** to view the rules defined on it. The members of the source or the destination are listed under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



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.



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

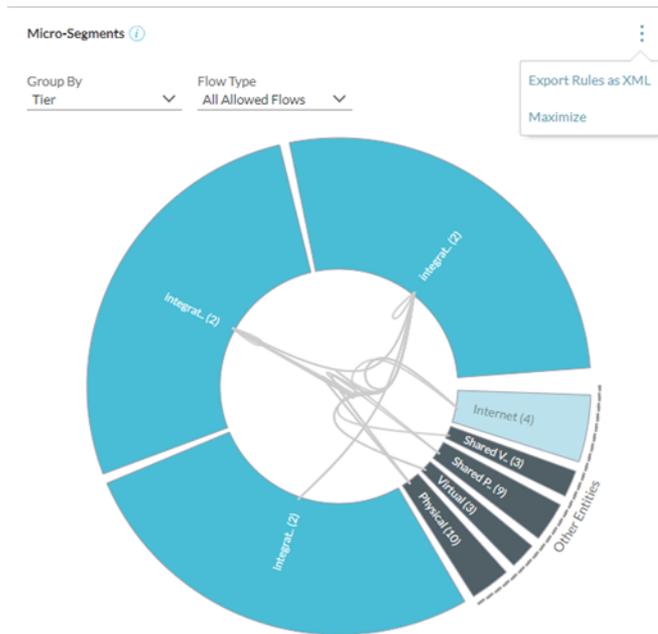
- **Related Flows:** Click the number of the related flows to see the list of flows with the corresponding flow information.
- **View Applied Firewall Rules:** Click the + sign next to the **Related Flows** column to view the applied firewall rules corresponding to the similar sets of flows.

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-0thers_Internet.xml
- SG-0ther.xml

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

NSX DFW Universal Artifacts

It is easy to manage objects in universal security groups across the various vCenter and NSX deployments. vRealize Network Insight supports the generation and the import of the universal artifacts for the Application and Tier groups only. With the universal security groups, it becomes easy to deploy and manage the firewall rules easily in the cross vCenter scenarios. Ensure that you import the universal artifacts on the primary NSX manager. You can manage the membership of the universal security group only through the primary NSX manager.

A universal security group can consist of:

- Other universal groups
- Universal IP sets
- Universal Security Tag

When you export the rules as XML, in addition to the NSX manager specific folders, a universal folder is created which consists of the NSX DFW universal artifacts. The corresponding universal security groups, universal IP sets, universal security tags, and universal DFW firewall rules are created after importing the NSX DFW universal artifacts.

Note

- The universal security tag is supported in only active-standby mode.
 - The universal IP set is supported in both active-active and active-standby modes.
-

You can create universal IP set or universal security tag based on your requirement. If you create the universal security tag, then you can map the application VM to the security tag. Else, the universal IP set is used.

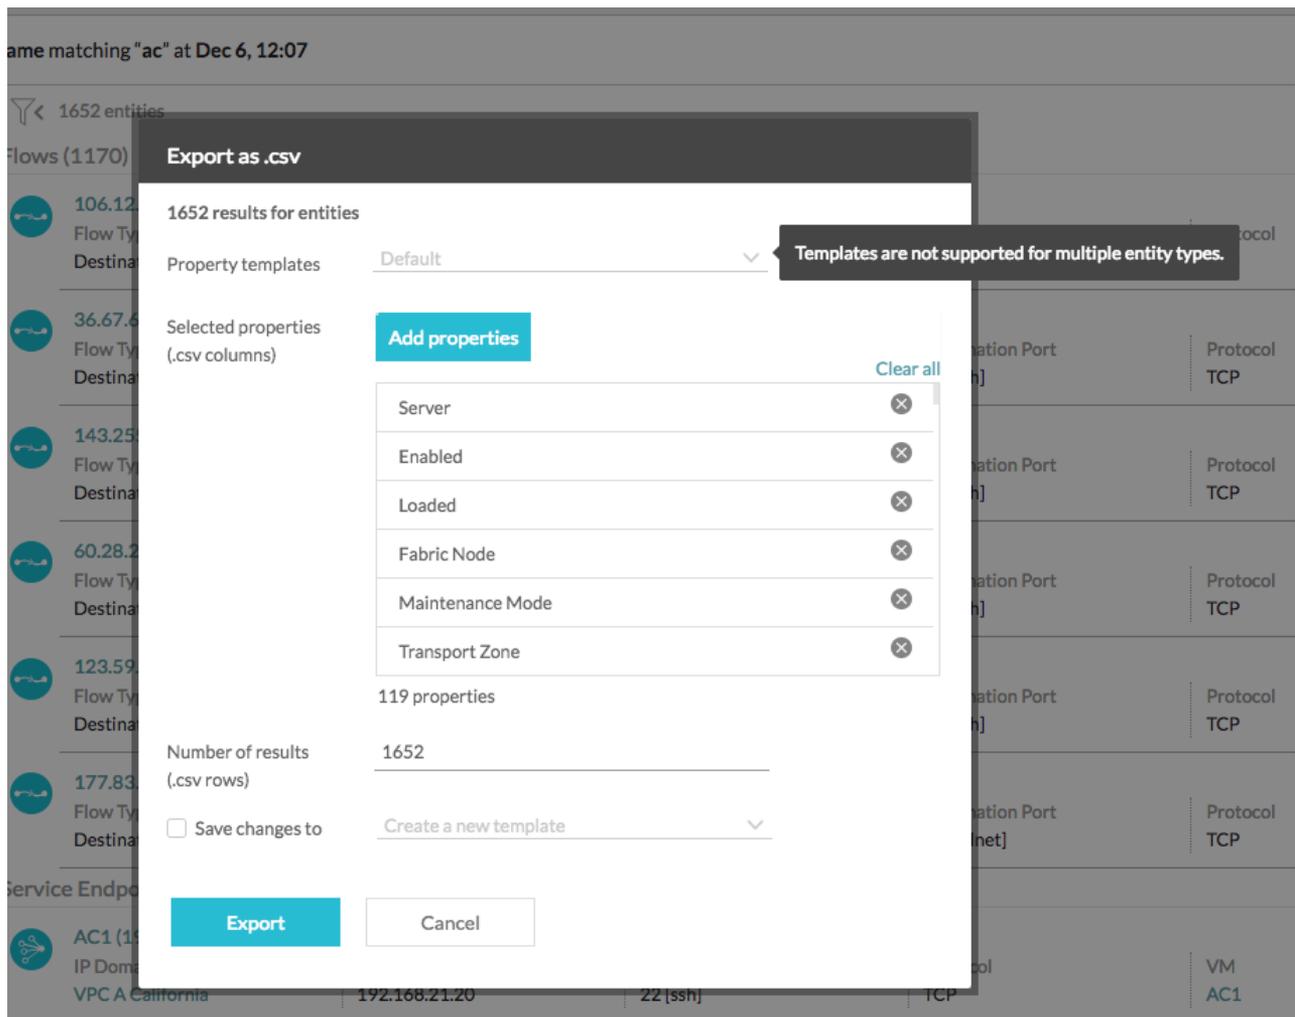
You can use the following flags in the import tool:

Table 12-1.

Flag Name	Description
-uni	To import artifacts from the universal folder.
-utag	To import the universal artifacts with the universal security tags in the membership of the universal security groups.
-log	To create rules in which logging is enabled.
	Note This flag is not specific to universal option.

Save the Configuration for CSV Export as Property Template

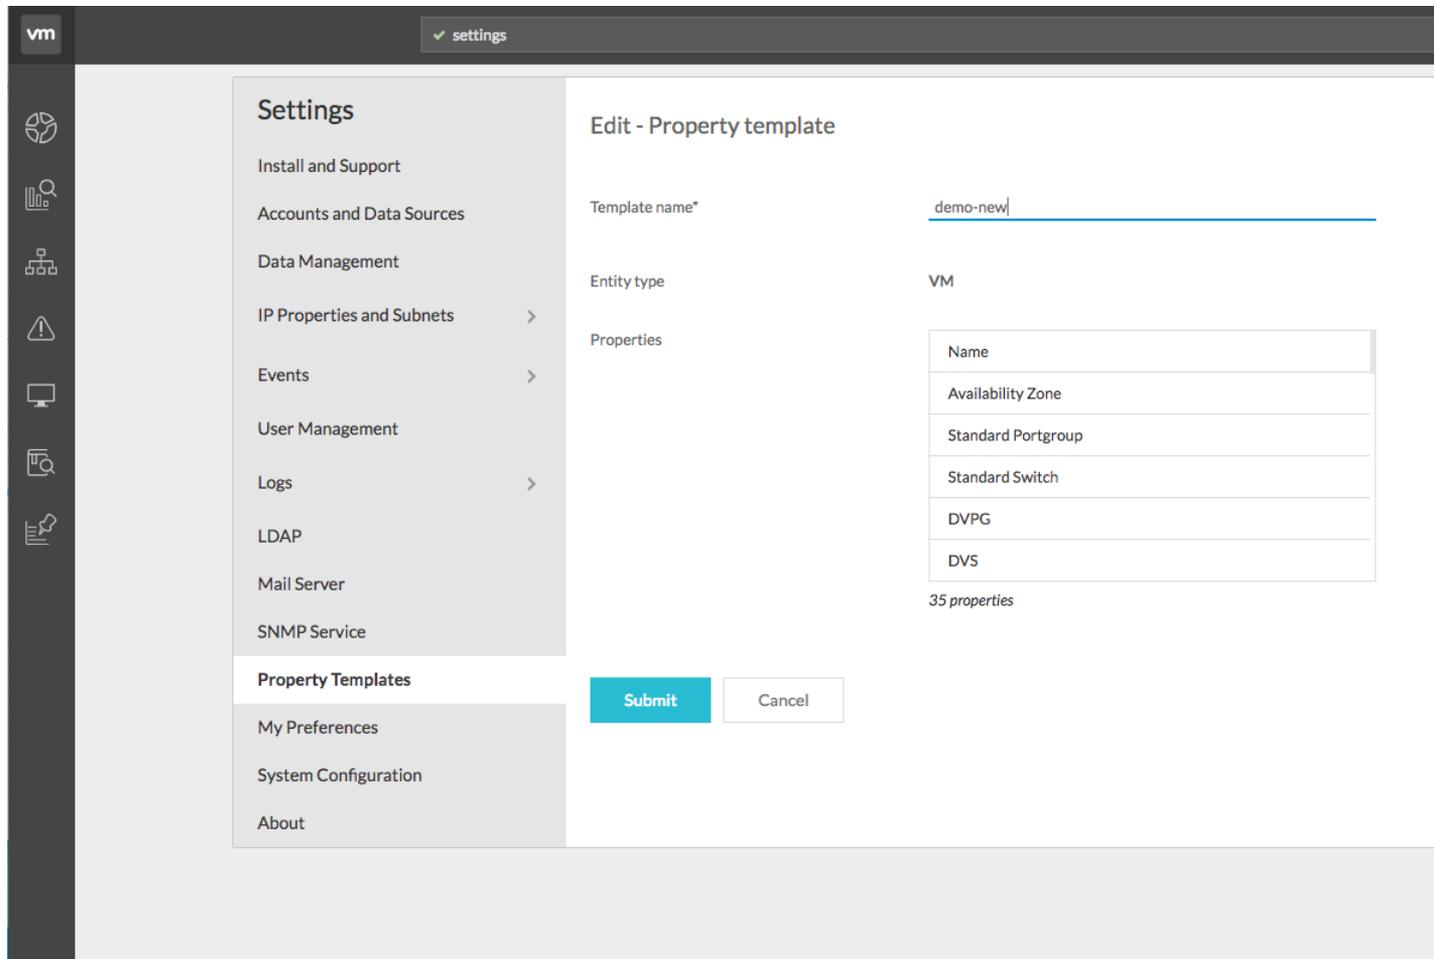
While exporting the data from widgets in the CSV files, you can save the combination of the properties (or columns) that you want to export in the property templates. These property templates are enabled for the CSV export when the results belong to a single entity type. If you search with a keyword that lists the multiple entity types, they cannot save the combination of the properties in the property templates.



When you open the CSV export modal, you will see the default property selections for the search results (based on the entity type). You can change this list of the selected properties and save the new configuration for your future reference. Alternatively, you can also load or open a pre-saved property template from the **Templates** section on the CSV export modal. When you change the value, you will see the selected properties for the selected property template.

Once you make changes to the selected properties for the export, you can create a property template from the CSV export modal or edit an existing property template. This template is of the same entity type as that of the current search results.

You can view the list of the existing property templates in the system by navigating to the **Settings -> Property Templates** page. The list on the **Property Templates** page shows the existing templates with details such as entity type, last updated, and number of properties. You can edit or delete property templates from the **Property Templates** page. You can edit the property template except change its name.



Search

vRealize Network Insight provides a robust search for all the entities in your environment.

Here are some of the terms that can help you with the search feature in vRealize Network Insight:

- **Entities:** A data center consists of physical and logical building blocks such as host, virtual machine, switch, router, NSX Manager and so on. The instances of these blocks are entities.
- **Property:** An entity consists of multiple properties. A property can be either a configuration property or a metric property.
 - a **Configuration Property:** An entity can be described by its configuration properties. A configuration property can be either integer or real value or a string or a boolean value.
 - Name, CPU cores, and operating system for virtual machines
 - Name, and number of virtual machines for hosts
 - b **Metric Property:** Any property which measures a particular characteristic of an entity is a metric property. The values of metric properties are captured at regular intervals of time. CPU usage, memory usage, and network usage for virtual machines are some examples of metric properties.
- **Aggregate Functions:** They can be used in the search queries to compute the total number of instances of a particular entity type or maximum property of an entity. vRealize Network Insight supports following aggregation functions.
 - a sum
 - b max
 - c min
 - d avg

When you search for entities, the software displays the entities that match your search query on the **Results** page.

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:

- [Search Queries](#)

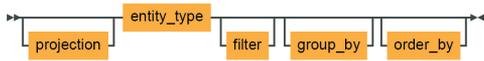
- [Advanced Queries](#)
- [Time Control](#)
- [Search Results](#)
- [Filters](#)
- [vCenter Tags](#)

Search Queries

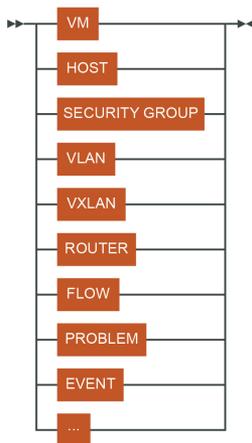
Search queries can be divided into following categories:

1 Structured Queries

A structured query consists of the following components:



- **Entity Type:** An entity type represents the type of object that we want to search. And it can be either in a singular form or in a plural form. The entity type is mandatory in a structured query.



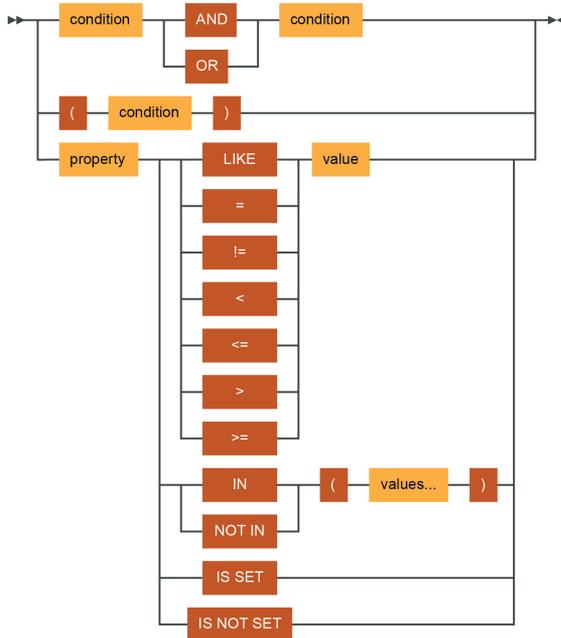
Here are some examples:

- 1 Virtual machines
- 2 Hosts
- 3 Flows
- 4 MTU Mismatch Events
- 5 Problems

- **Filters:** The syntax for filter is as follows:



The syntax for condition is as follows:



A filter clause can be used to filter search results. The condition in a filter clause consists of property, comparison operator, and value. The conditions can be combined with logical operators to form complex conditions. Here is a list of the operators that you can use:

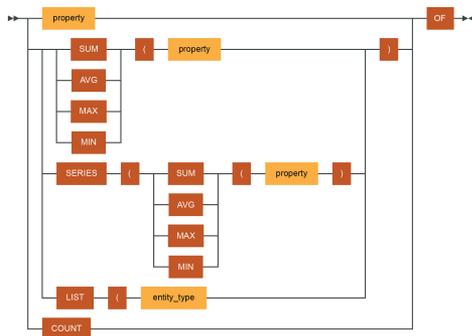
Table 13-1.

Operator	Examples
=	flows where source ip address = '10.16.240.0/24' flows where flow type = 'Source is VM'
!=	vms where ip address != '10.17.0.0/16'
>	vms where memory > 4096 mb
<	vms where cpu usage rate < 70%
>=	vms where memory >= 4096 mb
<=	vms where cpu usage rate <= 70%
like	vms where name like 'app'
not like	vms where name not like 'app'
in	flows where port in (22, 23, 80, 443)
not in	flows where port not in (22, 23, 80, 443)
is set	vms where firewall rule is set
is not set	vms where firewall rule is not set
()	flows where (src tier = 'App' and destination tier = 'DB') OR (destination tier = 'App' and source tier = 'DB')

Operator	Examples
and	flows where src tier = 'App' and destinationtier = 'DB'
or	flows where flow type = 'Source is VMKNIC' or flow type = 'Destination is VMKNIC'

- Projections:** A projection clause in a query decides what fields must be displayed from the filtered entities. This is an optional clause. If the projection clause is not specified, then the default set of fields is shown in the search results. A projection clause can contain any one of the following items:

- Property
- Count
- List
- Aggregation
- Series



- Property:** When entities are searched by an entity type, default set of properties are shown in the search results. Using projections, we can select the fields that should appear in the search results. For example, `os of vms` lists all virtual machines with `OS` property in the search results.

Here are some more examples:

- cpu cores of vms
- source ip address of flows

If a metric property is used, a graph is displayed for each entity with the metric property as y-axis and time as x-axis.

- Count:** The count query can be used to compute the number of objects of an entity type. Here are some examples:

- count of vms
- count of hosts
- count of flows

- 3 **List:** A list operator is helpful if the filter condition cannot be applied on the entity that you fetch.

For example:

```
List(host) of vms where memory <= 2gb
```

This query fetches list of hosts, whereas the filter condition is applied on virtual machines. Here are some more examples:

- List(ip address)of vms where cpu cores = 1

- 4 **Aggregate functions:** An aggregate function allows you to calculate a single value from a numerical config or metric property. The search query language supports the following aggregate functions:

- max
- sum
- min
- avg

Here are some examples:

- sum(memory) of hosts
- sum(memory), sum(cpu cores) of vms
- sum(bytes) of flows

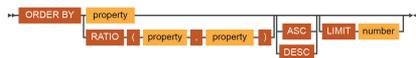
- 5 **Series:** A series operator is used to perform aggregation on the metric properties. For example:

```
series(avg(cpu usage)) of vms where cpu cores = 4
```

This query displays graph containing average cpu usage of all virtual machines with 4 cpu cores. Here are some examples:

- series(sum(network usage)) of vms where name like 'app'
- series(sum(memory usage)) of vms where name like 'db'
- series(avg(cpu usage)), series(avg(memory usage)) of vms

- **Ordering:** The search results can be sorted using order by clause. Only one field is allowed in order by clause. Results are sorted in descending order by default.



Here are some examples:

- 1 vms order by cpu cores
- 2 vms order by cpu cores asc
- 3 flows order by bytes

The `limit` clause can be used to limit the number of results. This must be preceded by the `order by` clause. For example:

```
vms order by memory limit 5
```

- **Grouping:** The entities can be grouped by a property. When entities are grouped by a property, by default, the number of results in each group are shown. By adding a projection, `sum/max/min` of any property can be computed. Adding `order by` clause sorts the results. If `order by` or `projection` clause is present in a query, then the aggregation function must be present.



```
sum(bytes) of flows group by dest vm
```

This query is valid as the query has aggregation function in the projection clause. A query such as `bytes of flows group by dest vm` is invalid as there is no aggregation function in the projection clause.

Here are some examples:

- 1 `vms group by host`
- 2 `sum (bytes) of flows group by dest vm order by sum(bytes)`

2 Entity Queries



- a **Search by entity type:** All entities of an entity type can be listed by searching the entity type.

Examples: `vms`, `hosts`, `flows`, `nsx managers`

- b **Search by entity name**

- Search by full name: If the full name of an entity is known, it can be searched by enclosing the name in single quotes.

Examples: `'prod-68-1'`, `'app1-72-1'`

- Search by partial name: Search by a single word or multiple words fetches all the entities matching the input words.

Examples: `prod`, `app1`

Note If input contains keywords or entity types, then it may be processed as a search query.

- Search by entity type and name: If both the name and the type of an entity are known, it can be searched by querying entity type and entity name together.

Example: The search query `'vm app1'` returns all VMs containing `app1`.

3 Planning Queries

These queries can be used to plan the security of the data center by analyzing flows.

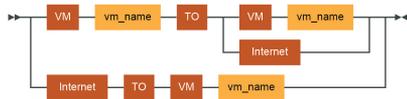


Examples:

- a `plan securitygroup1`
- b `plan host1`
- c `plan security`

4 Path Queries

These queries can be used to show the path between two VMs or the path from VM to Internet.



Examples:

- a `Vm 'vm1' to Vm 'vm2'`
- b `VM 'vm1' to Internet`

Note

- The search queries are not case-sensitive.
- The entity types or the configuration properties can have synonyms. For example, entity type 'virtual machine' has synonym 'vm'.

Advanced Queries

Here are some examples for advanced queries:

Flow Queries for Communication Patterns

- Total Traffic across data centers or sites (DCI link use)


```
sum(bytes) of flows where ( Dst Manager = 'abc' AND src manager = 'cba') OR ( Dst Manager = 'cba' AND src manager = 'abc')
```
- Total VTEP traffic
 - `sum(bytes) of flows where Flow Type = 'Src is VTEP' or flow type = 'Dst is VTEP'` VTEP traffic grouped by VMKNIC
 - `sum(bytes) of flows where Flow Type = 'Src is VTEP' or Flow Type = 'Dst is VTEP'` group by ip
- Other Management Traffic


```
flows where Flow Type = 'Source is VMKNIC' or Flow Type = 'Destination is VMKNIC'
```

Flow Queries for Aggregation and Grouping

- Total Internet traffic by Source VM
`sum(bytes) of flows where Flow Type = 'Internet' group by src vm`
- Top ports by total bytes
`sum(bytes) of flow group by port order by sum(bytes)`
- Top subnet pairs by routed traffic volume
`sum(bytes) of flow where Flow Type = 'Routed' group by Source Subnet Network, destination subnet network order by sum(bytes)`
- Total VM by total pair bytes
`sum(bytes) of flows group by src vm , dest vm order by sum(bytes)`
- Top Server VM/Port by total bytes
`sum(bytes) of flows group by dest vm , port order by sum(bytes)`

Flow Queries for Capacity Estimation and Sizing

- Total bytes of all vm-internet/internet-vm traffic grouped by ESX (Palo Alto Service VM sizing)
`sum(bytes) of flows where flow type = 'internet' and (flow type = ' src is vm ' OR flow type = 'destination is vm ') group by host order by sum(bytes)`
- Aggregated traffic series for matching flows (Palo Alto Service VM sizing)
`series(sum(byte rate)) of flows where host = 'ddc1-pod2esx012.dm.democompany.net' and (Flow Type = 'Source is VM' OR flow type = 'Destination is VM')`

Useful Queries for Application

- VMs in a given application
`VM where application = 'CRM'`
- Routed Flows from a given application
`Flows where source application = CRM and Flow Type = 'Routed'`
- Flows between two tiers (one-way)
`Flows where src tier = 'App' and Destination Tier = 'DB'`
- Flows between two tiers (one-way)
`Flows where (src tier = 'App' and destination Tier = 'DB') OR (destination tier = 'App' and source tier = 'DB')`

Useful Queries for VM and ESX

- Properties of Prod -Midtier-1 VM (MAC, IP, host, and so on)

CPU Usage Rate, Network Rate, Memory Usage Rate, mac address, ip , vxlan , host of vm
'Quality control-VM26'

- Network segments having the highest VM count
vm group by l2 network
- Datastores have highest VM count
vm group by datastore
- Hosts by vSphere version
host group by version
- Hosts by vSphere Builds
host group by OS
- All VMs on all host/blade slotted in a particular UCS chassis (Nested Query)
vm where host in (host where Blade like 'sys/chassis-1')

Useful Queries: Routes

- VNIs by Primary controller
vxlan group by Primary Controller
- Routes for Provider edge 3
routes where vrf = 'Provider Edge 3'
- Routes of DMZ DLR
NextHop Router of routes where VRF = 'LDR-DMZ'
- Routes having the given router as next hop
routes where NextHop Router = 'California-Edge'

Useful Queries: Firewall Rules

- Firewall rules between two VMs
firewall rules from 'Prod-Midtier-1' to 'Prod-Db-1'
- Rules with have ANY source
firewall rules where Service Any = true
- VMs for a given rule
vm where Firewall Rule = 'Prod MidTier to Prod DB - DBService '
- Firewall rules where any port is allowed
firewall rule where action = allow and service any = true
- Flows hitting a particular firewall rule

```
flows where firewall rule = 'Admin to Prod and Lab - SSH'
```

- Denied flows in the system

```
flows where firewall action = deny
```

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

Showing 12 results for Vxlan with filter Num VMs > 0 at

Filters

Add more filters

VM Count

- All (5)
- 1 (5)
- 2 (2)
- 3 (2)

NSX Manager

Scope

12 entities

Entity Name	Number of VMs	NSX Manager	Scope	Segment ID	Network Address
Siteb-Aundh-LS	3	10.197.17.114	Global	5006	192.168.23.0/24
Siteb_P-seattle-vxlan	3	10.197.17.229	Global	5000	172.17.1.0/24
Siteb_P-redmond-vxlan	2	10.197.17.229	Global	5001	172.17.2.0/24
Siteb-Wagholi-LS	2	10.197.17.114	Global	5005	192.168.26.0/24
Siteb-pashan-ls-1	2	10.197.17.114	Global	5002	192.168.24.0/24
Siteb_P-transit-vxlan-2	2	10.197.17.229	Global	5005	172.17.6.0/24
Siteb_P-transit-vxlan-1	2	10.197.17.229	Global	5004	172.17.5.0/24
Siteb-Transit-LS-1	1	10.197.17.114	Global	5003	192.168.21.0/24

Filters

Filters

Add more filters

Default Gateway

- All
- 192.168.23.10 (1)

NSX Manager

Scope

VM Count

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.

vCenter Tags

vRealize Network Insight provides vCenter tags for search and planning.

You can perform a search of VMs based on the vCenter tags and custom attributes. For example, you can use the following query for search by using tags:

```
vm where tag = '{keyname}:{value}'
```

Every tag belongs to a category. In the above example, the keyname is the category to which the tag belongs and value is the name of the tag.

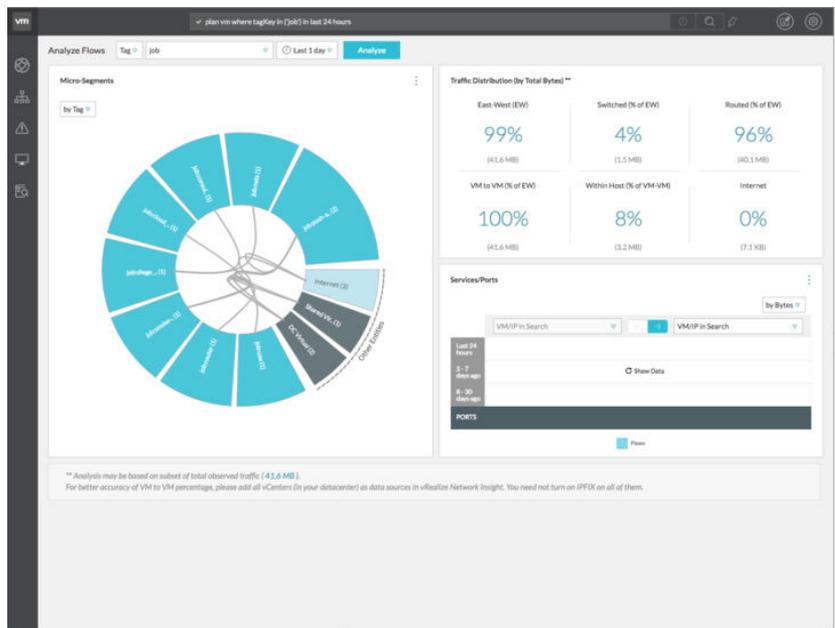
You can also provide an alternate name to a VM by using vCenter tags or custom attributes by using the name key. This alternate name is shown as the other_names property. It is also possible to search and make path queries using the alternate name.

For example, the following queries are supported:

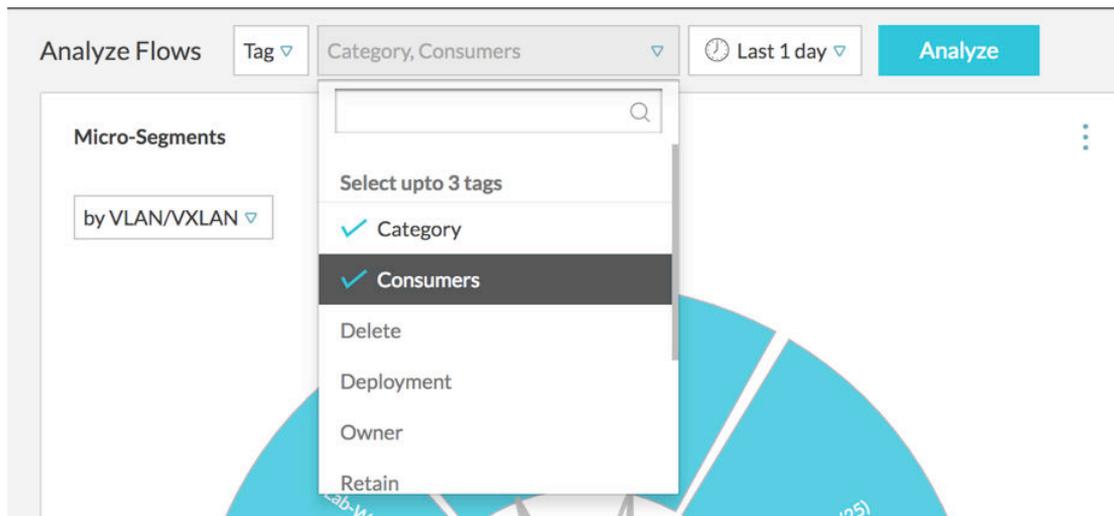
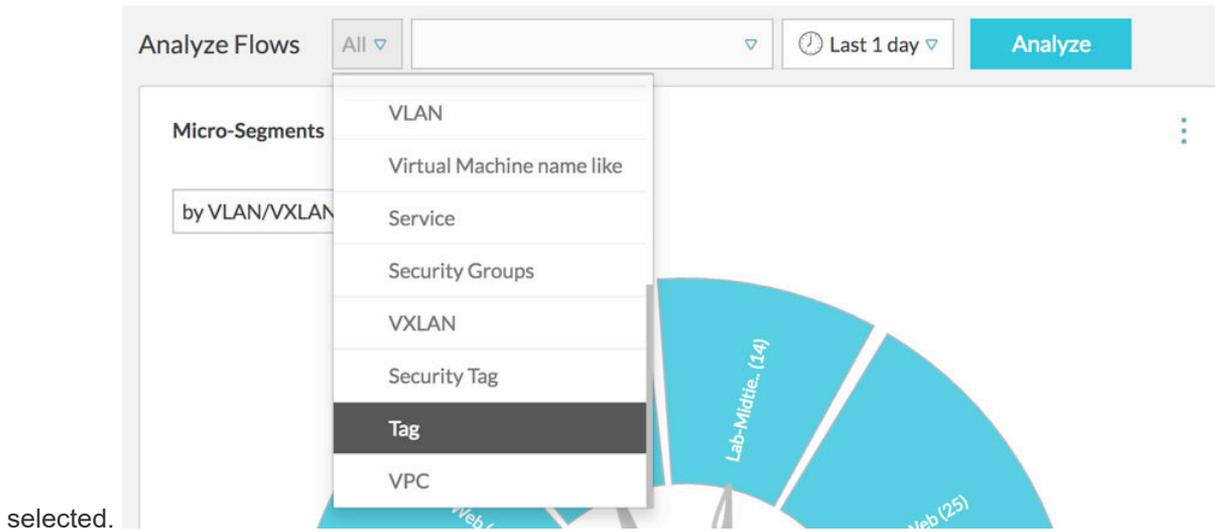
```
vm "other-name-1"
vm "other-name-1" to vm "other-name-2"
```

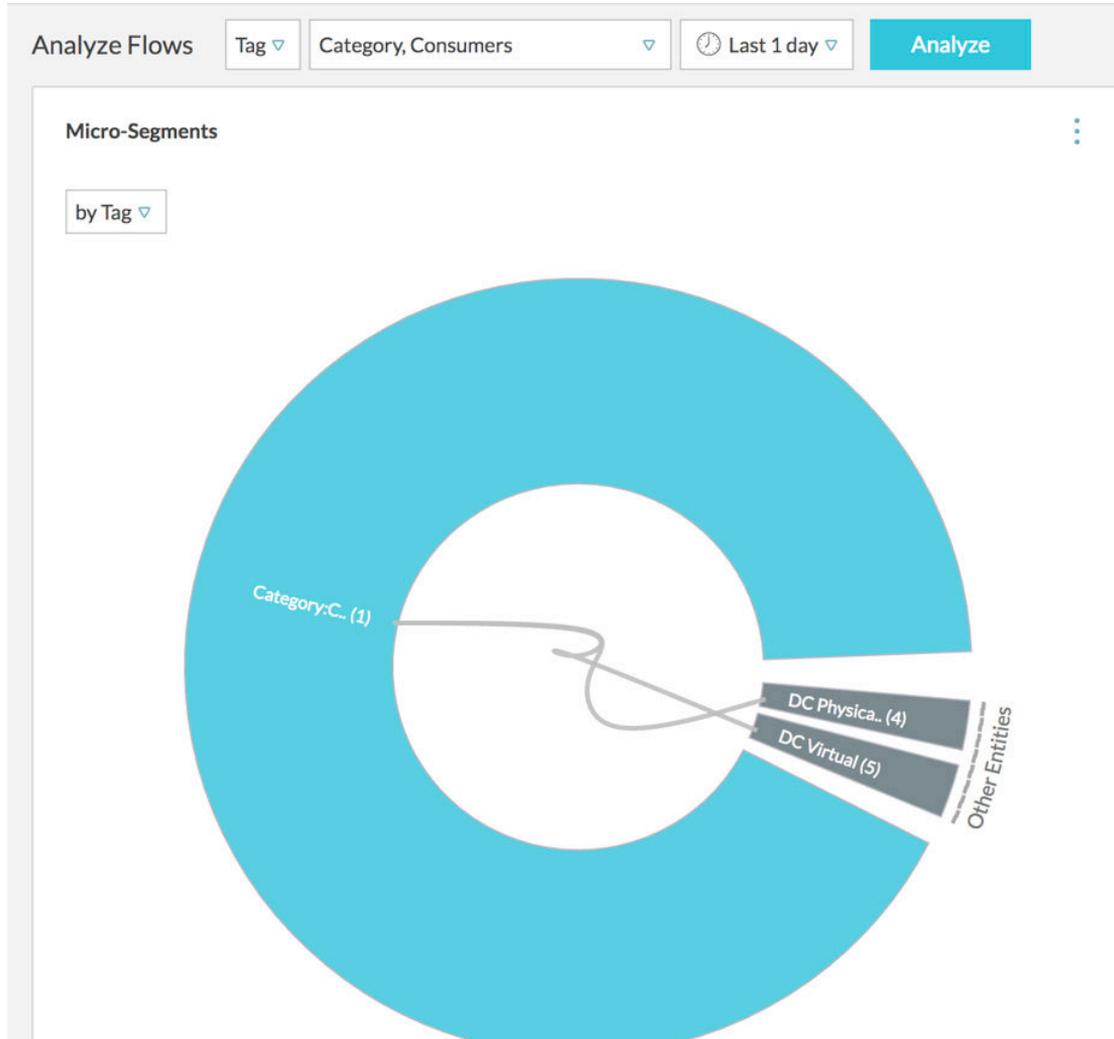
In this example, other-name-1 and other-name-2 are custom attributes with the name key or tags belonging to the name category.

You can also analyze the flows in the network by using the vCenter tags as shown in the figure.



To use the vCenter tags, select the **Tag** option from the **Analyze Flows** drop-down list. You can also select up to three tags at this level. After you select the tag, click **Analyze**. In **Group by Criteria**, **Tag** is





Upgrade

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, 3.8->3.9, and so on.

Ensure that the following disk space requirements are met before going ahead with the upgrade:

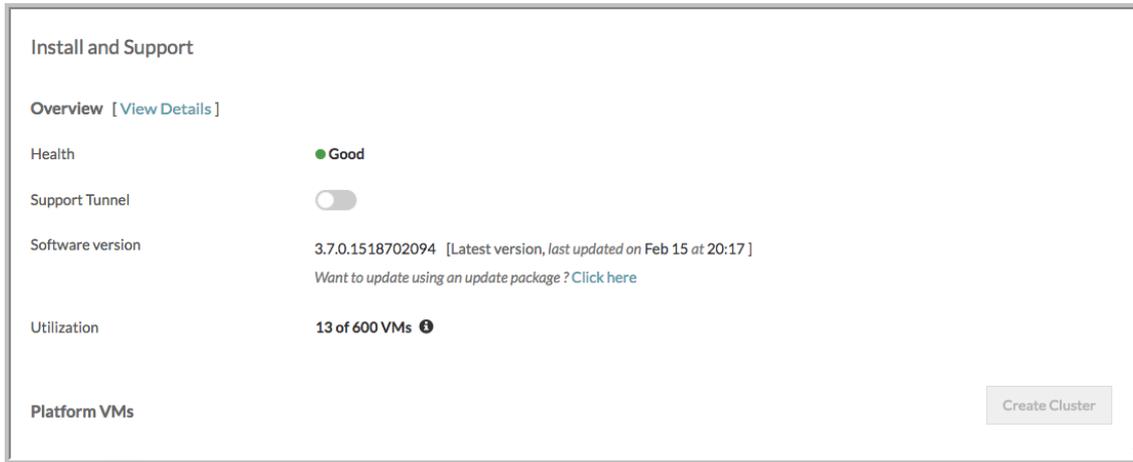
- The following requirements are for both platform and proxy server:
 - /tmp - 6 GB
 - /home - 2 GB

- /- 12 GB (Only for the Platform1 node)

Note This requirement is only for the platform.

Procedure

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



- 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, the upgrade is not triggered.
 - Once the upgrade begins, if a node becomes inactive, the upgrade process does not continue. 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.
 - Once the platforms are upgraded, you can resume your normal vRealize Network Insight operations even though the collector upgrade happens in parallel. Until the upgrade process is completely over, the Node Version Mismatch detected message is shown in the **Install and Support** page.
-
- 5 Upon the completion of upgrade process, all platforms and the collectors nodes are upgraded.

vRealize Network Insight Disaster Recovery

15

VMware Site Recovery Manager (SRM) is a disaster recovery automation software that provides policy-based management, non-disruptive testing, and automated orchestration. vRealize Network Insight supports SRM 8.1 and the further versions. To protect your vRealize Network Insight, SRM automates every aspect of executing a disaster recovery plan to accelerate recovery and eliminate the risks involved when using a manual process.

For information about installing, upgrading, and configuring SRM, see [VMware Site Recovery Manager Documentation](#).

The prerequisites for the disaster recovery operation for vRealize Network Insight are as follows:

- Ensure that you have installed and configured vSphere Replication.
- SRM should be deployed and configured on both the protected and the recovery sites.
- Ensure that the site pairing is configured properly from within the SRM UI before proceeding with the creation of the recovery plan and other components.
- VMware vSphere Replication should be enabled for each of the protected nodes of the vRNI setup in context. While enabling VMware vSphere Replication, provide sufficient RPO considering the vRealize Network Insight node size & the usage so that minimum data loss is expected to incur during a disaster. For more information on replication, see [VMware vSphere Replication Documentation](#).
- Ensure that you create a separate protection group for vRealize Network Insight. For small and non-distributed deployments, ensure that all the VMs are in the same protection group. For distributed deployments, it is recommended that you place all the platforms in a single protection group so that it is easy to recover. You can place the collectors in different protection groups.
- Create a recovery plan and add the protection groups containing vRealize Network Insight VMs to this plan. Ensure that the protection group containing the platform nodes get the higher precedence. In the recovery plan, ensure that the primary platform node is placed in a higher priority group than the other platform nodes.
- Currently, any type of IPv4 customization with SRM is not supported

It is recommended that you migrate or recover vRealize Network Insight VMs to an identical network configuration. Also as per the SRM recommendation, you can perform test run periodically to ensure the existing plan works with underlying infrastructure and the configured RPO limit.

- Migrate or recover vRealize Network Insight VMs to an identical network configuration.

If the recovery site is configured to have the same network configuration as the protected site and a mapping is created between the identical networks, configure all replicated vRealize Network Insight virtual machines to be started with the same IPs, because these VMs are the protected nodes. The recovered system will become operational after the planned migration or disaster recovery has finished successfully.

- Do not specify any IP customization for a recovery plan when the recovery site does not have the same network as that of the protected site. In this scenario, SRM is used for the recovery of the appliance VMs. For configuring network post recovery, manually assign the network settings as follows:
 - 1 Run the `change-network-settings` command simultaneously on all the platform nodes.
 - 2 Run the `update-IP-change` command on the nodes on Platform1, Platform2 and Platform3 consecutively.
 - 3 Run `vrni-proxy set-platform --ip-or-fqdn <with-updated-ip-of-Platform1>` on the collector node.
 - 4 Check the service status. If some of the services on the platform nodes are not running, reboot the nodes in the recommended order.

Note For more information on the commands mentioned above, see *vRealize Network Insight Command Line Reference Guide*.

This chapter includes the following topics:

- [Sample Disaster Recovery Scenario](#)

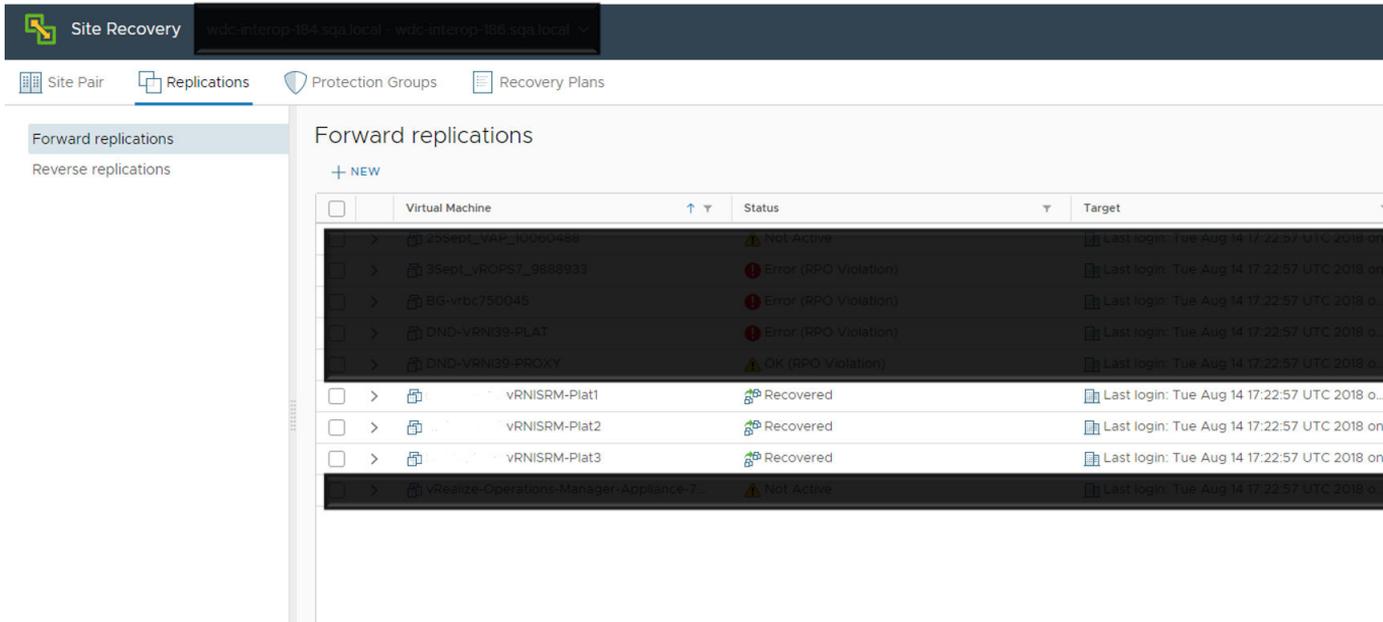
Sample Disaster Recovery Scenario

Here are the steps for a sample scenario for vRealize Network Insight Disaster Recovery (DR):

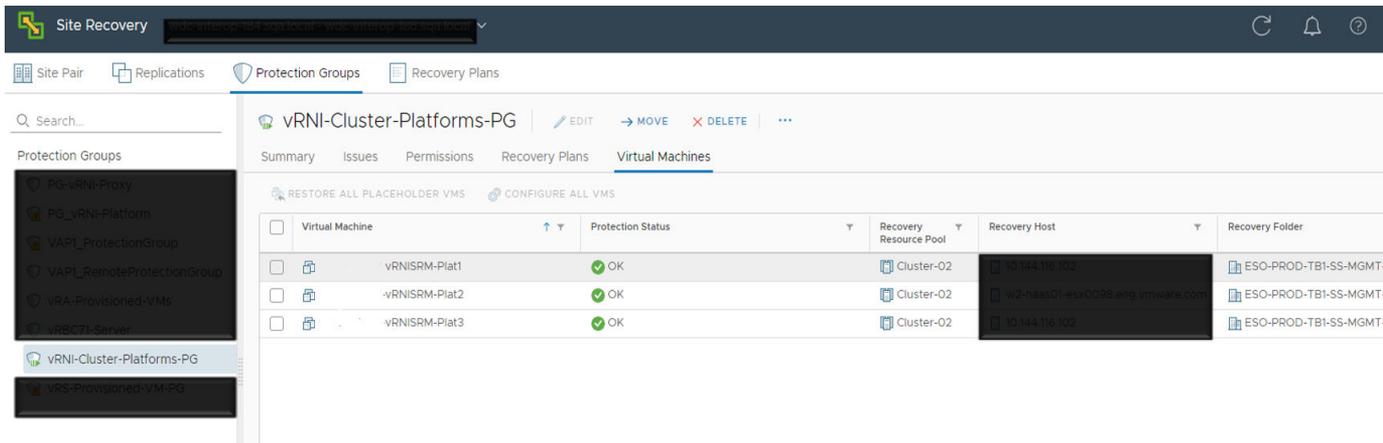
Procedure

- 1 Ensure that SRM is configured and up in both the protected and the recovery sites.

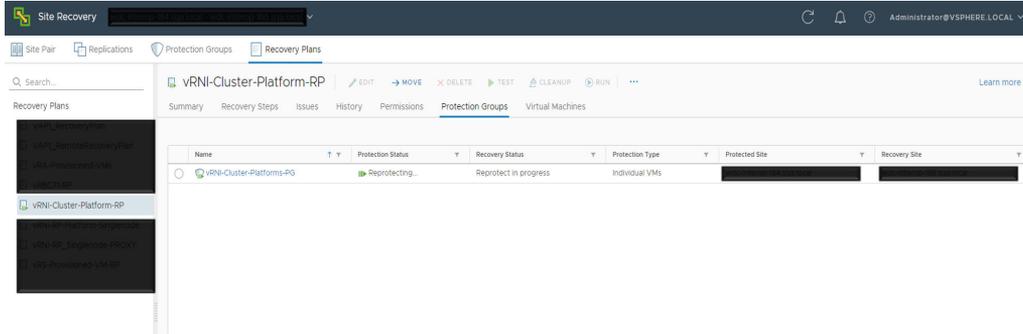
- Configure replication for each of the vRealize Network Insight nodes that are to be protected. While configuring the replication, provide adequate Recovery Point Objectives (RPO) time for the vRealize Network Insight instance. For example, if it is a vRealize Network Insight deployment with a single platform and collector nodes (medium size), then RPO of 45 minutes is good. But if it is a cluster with nodes having bricks of large size, then the adequate RPO should be provided. The snapshot interval configuration is specific to the user environment and requirement.



- Create protection group. Include the VMs that you want to protect under a specific protection group.



4 Create the recovery plan where you include the respective protection groups.



- 5 Perform test recovery. This is to ensure that your recovery plan works as expected.
- 6 SRM recommends that users perform planned migration at regular intervals to validate the integrity of the existingDR plan.
- 7 Suppose the recovery site has a network configuration that forces the vRealize Network Insight VMs to come up with the new IPs. Recover the vRealize Network Insight VMs with a recovery plan that assumes no network change for the recovered VMs. Once the recovery of the VMs is reported as a success in vRealize Network Insight, assign new IP addresses manually to the vRealize Network Insight nodes, apply new certificates, and re-initialize the cluster.
- 8 As IPv4 customization with SRM is not supported currently, as a work around you can perform DR with vRealize Network Insight assuming as if there is no network change.

To manually assign the network settings:

- a Run the `change-network-settings` command simultaneously on all the platform nodes.
- b Run the `update-IP-change` command on the nodes on Platform1, Platform2 and Platform3 consecutively.
- c Run `vrni-proxy set-platform --ip-or-fqdn <with-updated-ip-of-Platform1>` on the collector node.
- d Check the service status. If some of the services on the platform nodes are not running, reboot the nodes in the recommended order.

Troubleshooting

This chapter includes the following topics:

- [Audit Logs](#)
- [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](#)
- [Unable to Enable DFW IPFIX](#)

Audit Logs

Audit logs capture administrative actions carried out in the system. These are regular CRUD operations as well as login and logout events. The administrative actions carried out via UI, CLI or API are logged.

The audit logs capture the actions from API, UI, and CLI.

Features

- The audit log feature is always on.
- vRealize Network Insight supports the UTC format in the audit logs.
- The audit log is integrated with the syslog. You can configure the syslog collector to collect all the audit logs.

- You can export all the audit log data in a CSV file.

View the Audit Logs

Procedure

- 1 On the **Settings** page, click **Audit Logs** under **Logs**.
- 2 The following details are shown on the **Audit Logs** page:

Table 16-1.

Information	Description
Date & Time	Timestamp of the actual action performed.
IP Address	IP address of the client from which the connection is established such as the CLI or the browser.
User Name	User who is performing the action.
Object Type	Object on which the action is being performed.
Operation	Different actions that the user performs on the object.
Object Identifier	Unique identifier for that particular object on which the action is being performed.
Response	Indicator for success or failure of the operation
Details	Details of the settings that have been changed such as the nickname or a property.

- 3 To allow the collection of information when the user logs in through a browser or CLI, enable **Allow collection of Personally Identifiable Information**. This option is disabled by default.

Note The IP Address and the User Name columns are blank if this option is disabled.

- 4 Click **Export as CSV** to export the audit log data in the CSV format.

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, 3.8->3.9, and so on.

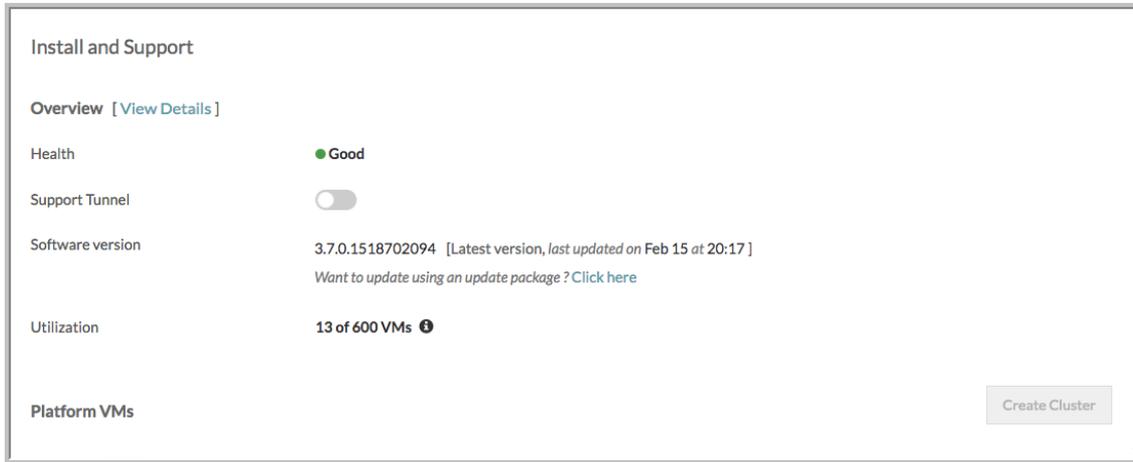
Ensure that the following disk space requirements are met before going ahead with the upgrade:

- The following requirements are for both platform and proxy server:
 - /tmp - 6 GB
 - /home - 2 GB
- /- 12 GB (Only for the Platform1 node)

Note This requirement is only for the platform.

Procedure

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



- 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, the upgrade is not triggered.
- Once the upgrade begins, if a node becomes inactive, the upgrade process does not continue. 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.
- Once the platforms are upgraded, you can resume your normal vRealize Network Insight operations even though the collector upgrade happens in parallel. Until the upgrade process is completely over, the Node Version Mismatch detected message is shown in the **Install and Support** page.

- 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
 Critical: Disk Utilization is high 	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.

- 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:

- 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.

- 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

- 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.
- In the **Edit Collector (Proxy) VM** page, you can assign a nickname to the proxy VM.
- 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.
- The Edit account or source page appears. Ensure that you fill the following information:

Table 16-2.

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
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 16-3.

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.
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 Text	Cause	Resolution
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.

Unable to Enable DFW IPFIX

vRealize Network Insight does not allow you to enable DFW IPFIX.

Problem

While adding a policy manager or source of VMware Cloud on AWS, when you attempt to enable DFW IPFIX, you might see the following error messages:

- No New collectors can be added.

- Provided user does not have the required role. Only users with the following role can enable IPFIX: Cloud Administrator.

Cause

- VMware Cloud on AWS supports only four collectors to its DFW IPFIX collector profile. So, when the existing profile already has four collectors, you see the

No New collectors can be added

message.

The screenshot shows the 'Add a New Policy Manager Account or Source of VMware Cloud on AWS' configuration page. On the left is a 'Settings' sidebar with categories like 'Accounts and Data Sources', 'Data Management', 'IP Properties and Subnets', 'Events', 'User Management', 'Logs', 'LDAP', 'Mail Server', 'SNMP Service', 'Property Templates', 'My Preferences', 'System Configuration', and 'About'. The main content area contains the following fields and controls:

- VCenter ***: vcenter.sddc-35-162-64-191.vmwarevmc.com (VC VMC P...)
- Collector (Proxy) VM ***: NI-Collector_10.153.189.42(Available Capacity: 951 VMs)
Tip: Want to increase capacity of your collector? [Click here](#)
- IP Address/FQDN ***: nsxManager.sddc-35-162-64-191.vmwarevmc.com
- CSP Refresh Token ***: 6f60efe1-6d45-448f-b3d5-76e7e15c92bb
- Validate** button: Validation Successful
- Enable DFW IPFIX**
Selecting this option will enable distributed firewall to send IPFIX flow record to the collector
- Warning message: **No new collectors can be added.**
- Nickname ***: (empty field)
- Notes**: Optional
- Submit** and **Cancel** buttons at the bottom.

- The user does not have the write permission. Only users with **Cloud Administrator** role can perform the write operation on the VMware Cloud on AWS policy manager.

Settings

- Install and Support
- Accounts and Data Sources
- Data Management
- IP Properties and Subnets >
- Events >
- User Management
- Logs >
- LDAP
- Mail Server
- SNMP Service
- Property Templates
- My Preferences
- System Configuration
- About

Edit Account or Source

VCenter * ⓘ vcenter.sddc-34-218-191-237.vmwarevmc.com (VC VMC ...)

Collector (Proxy) VM * NI-Collector_10.153.189.42(Available Capacity: 951 VMs)

Tip: Want to increase capacity of your collector? [Click here](#)

IP Address/FQDN * nsxManager.sddc-34-218-191-237.vmwarevmc.com

CSP Refresh Token * ⓘ 232add00-f35e-4d7d-af61-d6c06aa1d9c2

✔ Validation Successful

Enable DFW IPFIX

Selecting this option will enable distributed firewall to send IPFIX flow record to the collector

! Provided user does not have the required role. Only users with the following role can enable IPFIX: Enterprise Administrator, Cloud Administrator.

Nickname * POLCY VMC M5P2

Notes Optional

Solution

- ◆ To add a new collector, you must:
 - Delete an existing collector, or
 - Create a new profile, or
- ◆ To avoid or fix the user role issue, perform one of the following steps:
 - Assign the **Cloud Administrator** role to the user, or
 - Log in as user with **Cloud Administrator** role.

Application Migration Using vRealize Network Insight

17

Using vRealize Network Insight, you can assess your on-premise environment for application migration to VMware Cloud on AWS or AWS.

Steps	Procedure	references
Step 1	Setting up your Environment	<ul style="list-style-type: none">■ Accept the End User License Agreement (EULA).<ol style="list-style-type: none">a Create a VMware user account or log in to the VMware account.b Update the registration form. New users receive an email to activate their account.c Accept VMware terms and EULA.■ Download the OVA files<ol style="list-style-type: none">a Log in to the VMware Product Download page at https://my.vmware.com/group/vmware/homeb Search for vRealize Network Insight.c Download the latest vRealize Network Insight platform and proxy OVA files.■ Prepare for installation.<ol style="list-style-type: none">a Verify the System Recommendations and Requirements.b Verify the Supported products and versions.
Step 2	Deployment	<ol style="list-style-type: none">1 Deploy the vRealize Network Insight platform OVA file.2 Activate the License.3 Generate a shared secret4 Deploy the vRealize Network Insight Proxy OVA file..5 Deploy vRealize Network Insight in VMware Cloud on AWS (VMC).
Step 3	Data Source Addition	<ol style="list-style-type: none">1 Log in to vRealize Network Insight by using Default Login Credentials.2 Add vCenter Server.3 Add on-prem NSX Manager & underlay switches. For procedure details, see Adding Data Sources.4 Add Add VMC vCenter.5 Add VMC NSX Manager.
Step 4	Model Application	<ul style="list-style-type: none">■ Analyze application dependencies<ol style="list-style-type: none">a Create an Applicationb Creating Tiers for Physical IPsc Analyzing the Applicationd VMC: Planning and Micro-Segmentation■ Analyze recommended firewall rules■ Perform search■ Create and use Pinboards

This chapter includes the following topics:

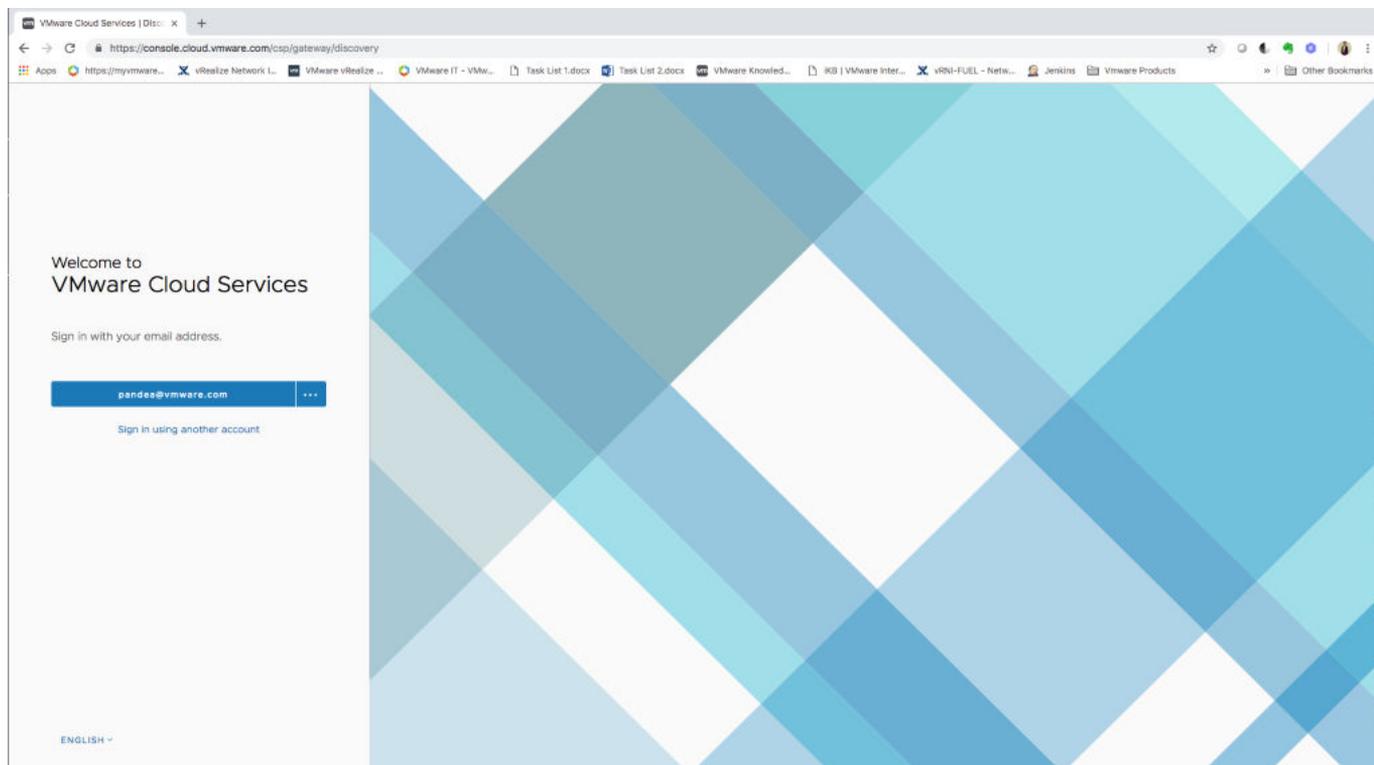
- [How do I obtain the CSP Refresh Token for NSX Manager](#)
- [How Do I Obtain vCenter Credentials](#)
- [Compute Gateway Firewall Rule](#)

How do I obtain the CSP Refresh Token for NSX Manager

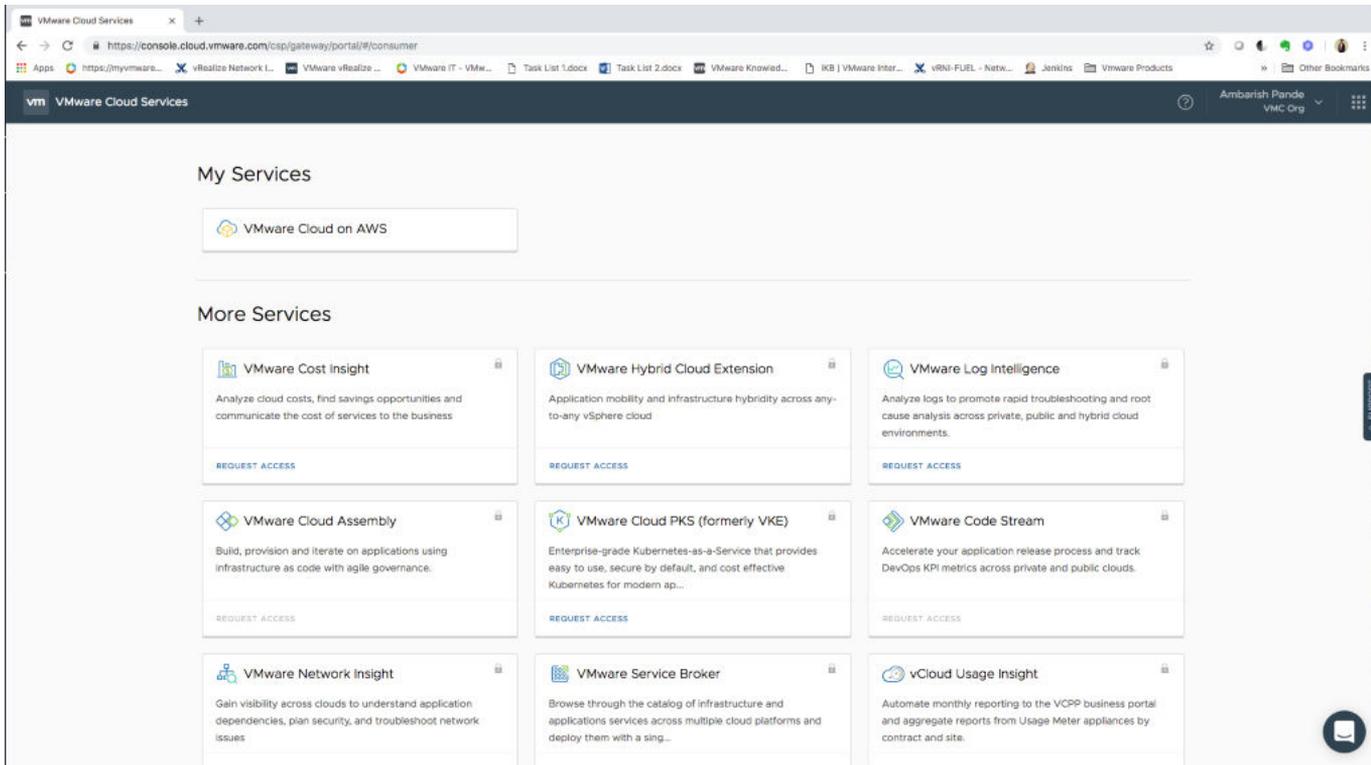
For adding a VMC NSX Manager as a Data Source in to vRealize Network Insight, you need a refresh token.

Procedure

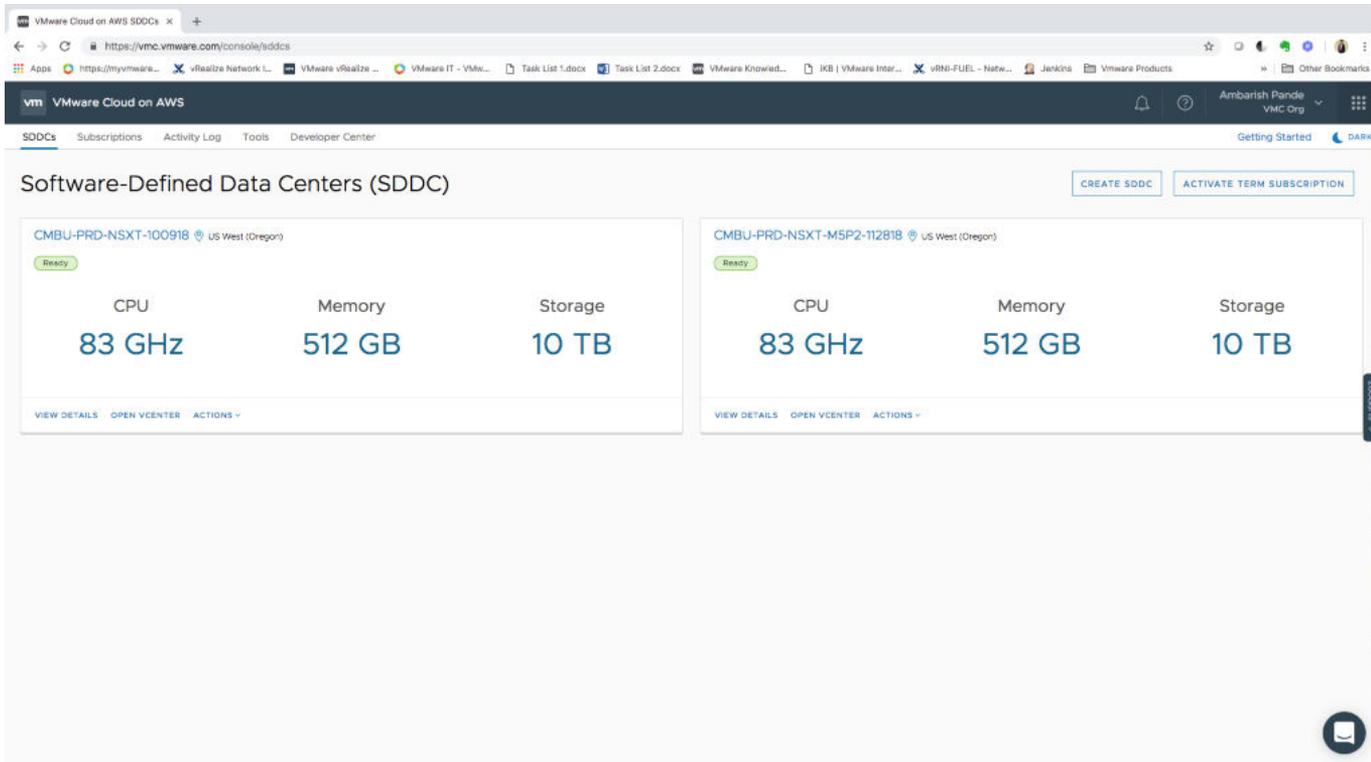
- 1 Log in to the VMware Cloud services console.



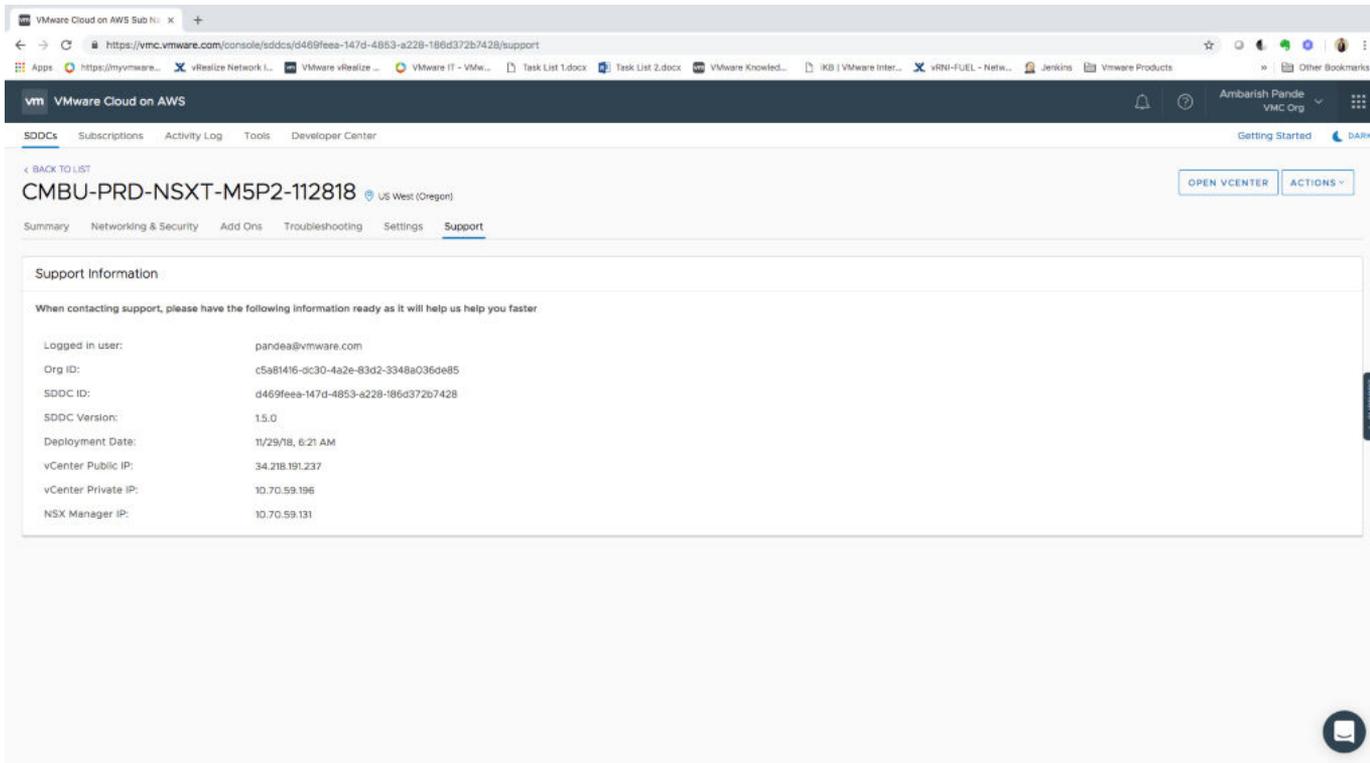
2 Under My Services, click VMware Cloud on AWS.



3 Select the desired Software-Defined Data Center (SDDC).

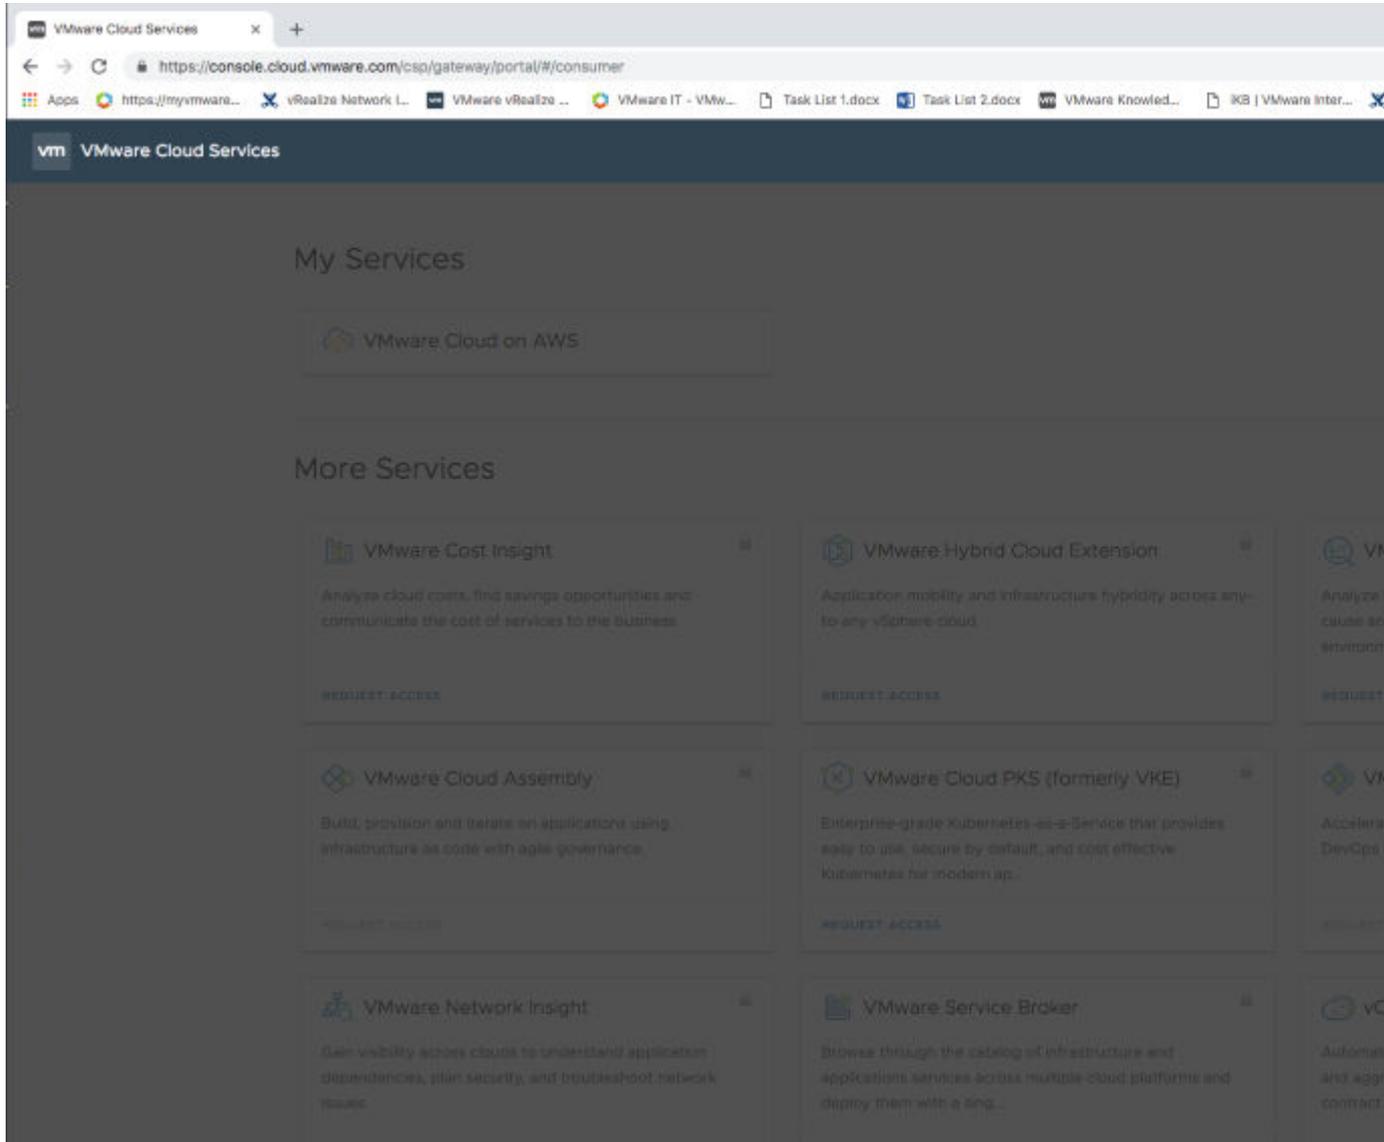


4 Click the **Support** tab.

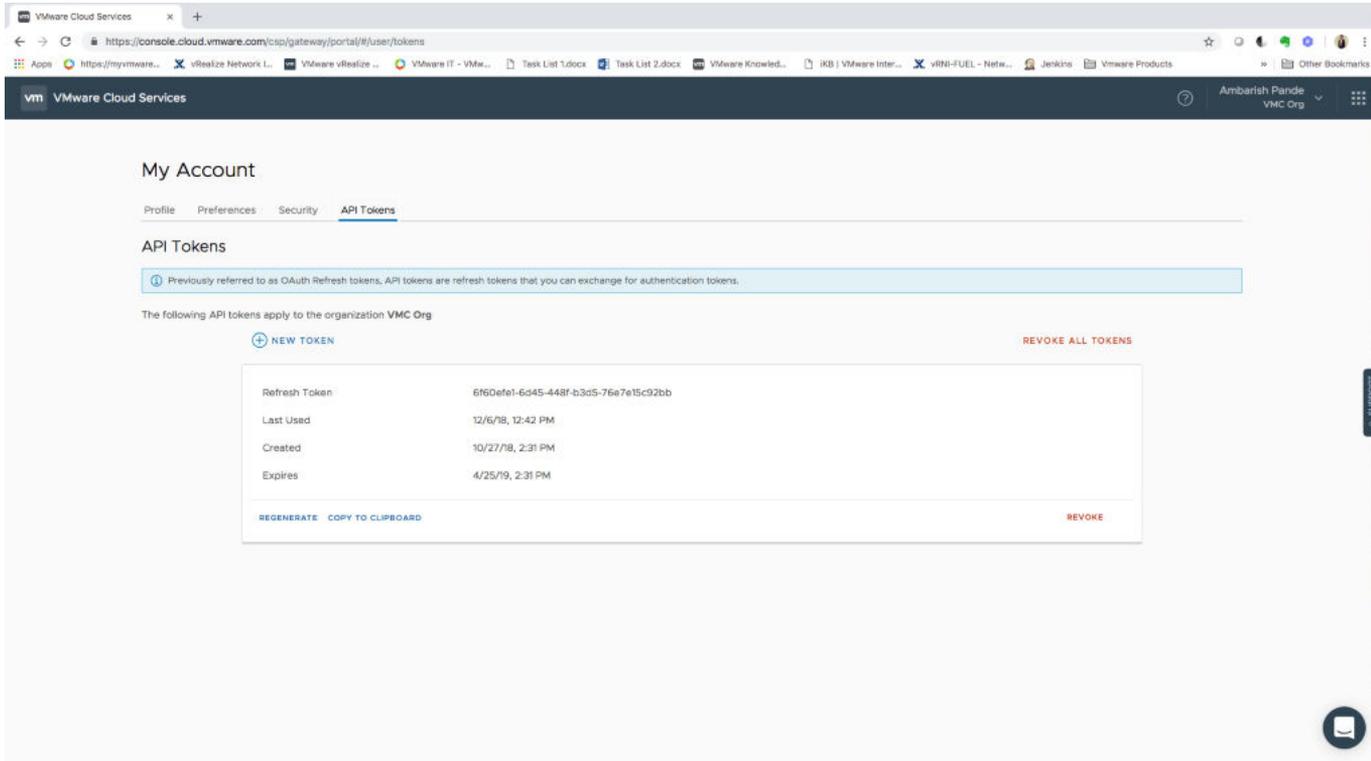


- 5 Make a note of the NSX Manager IP address.
- 6 Click on the organization name on the top banner.

Note Ensure that the organization resides in the selected SDDC.



7 On the **API Tokens** tab, copy the Refresh Token.



The Refresh Token is valid for Six months. vRealize Network Insight does not track the lifecycle of the token.

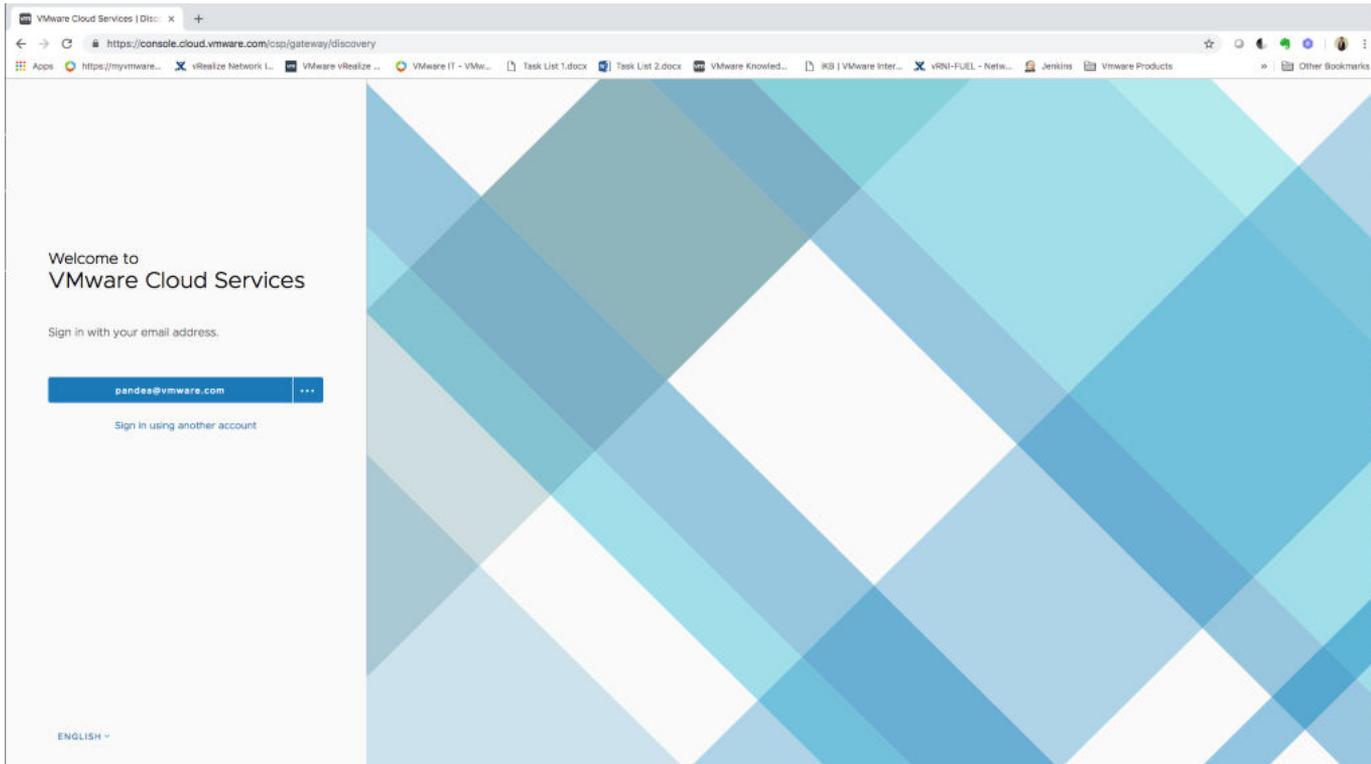
You can use this token for authenticating all VMware Cloud on AWS SDDCs on the organization.

How Do I Obtain vCenter Credentials

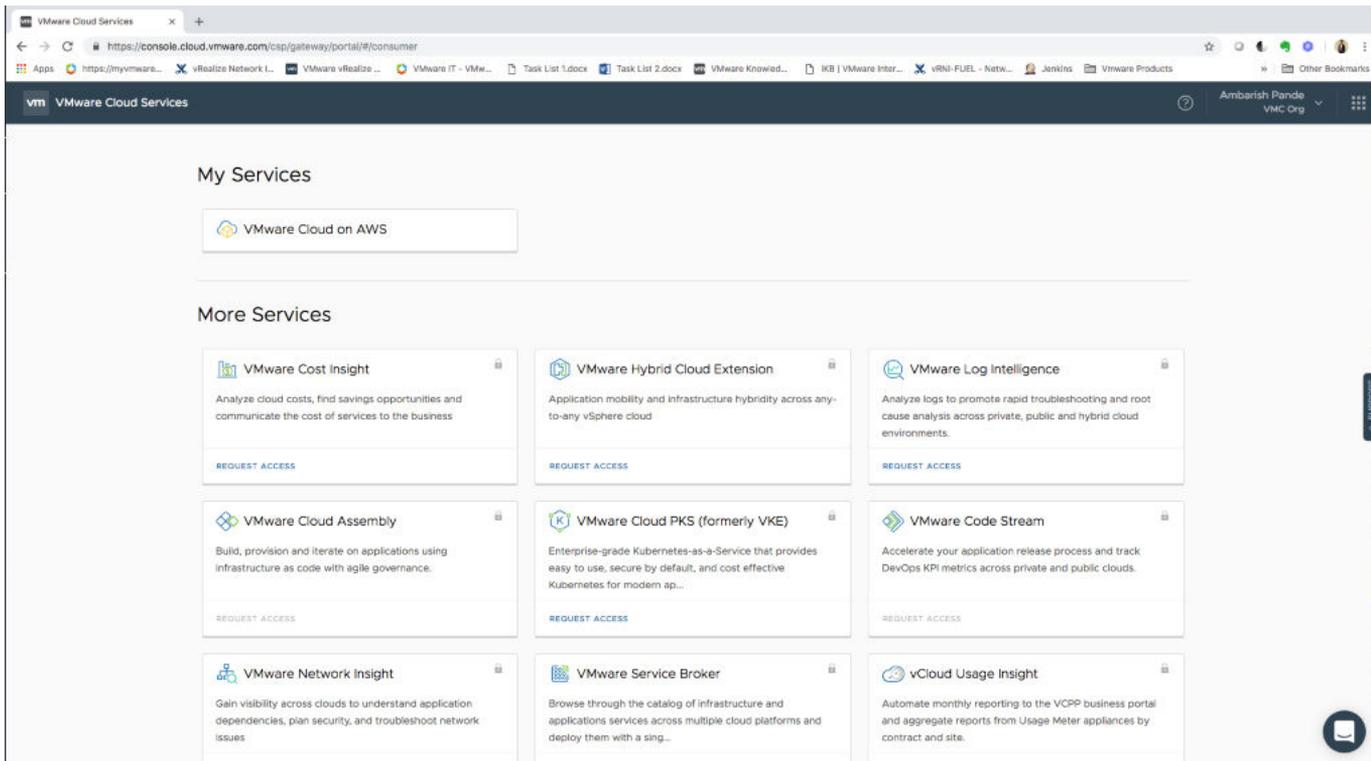
To add a vCenter datasource into vRealize Network Insight, you need the vCenter credentials.

Procedure

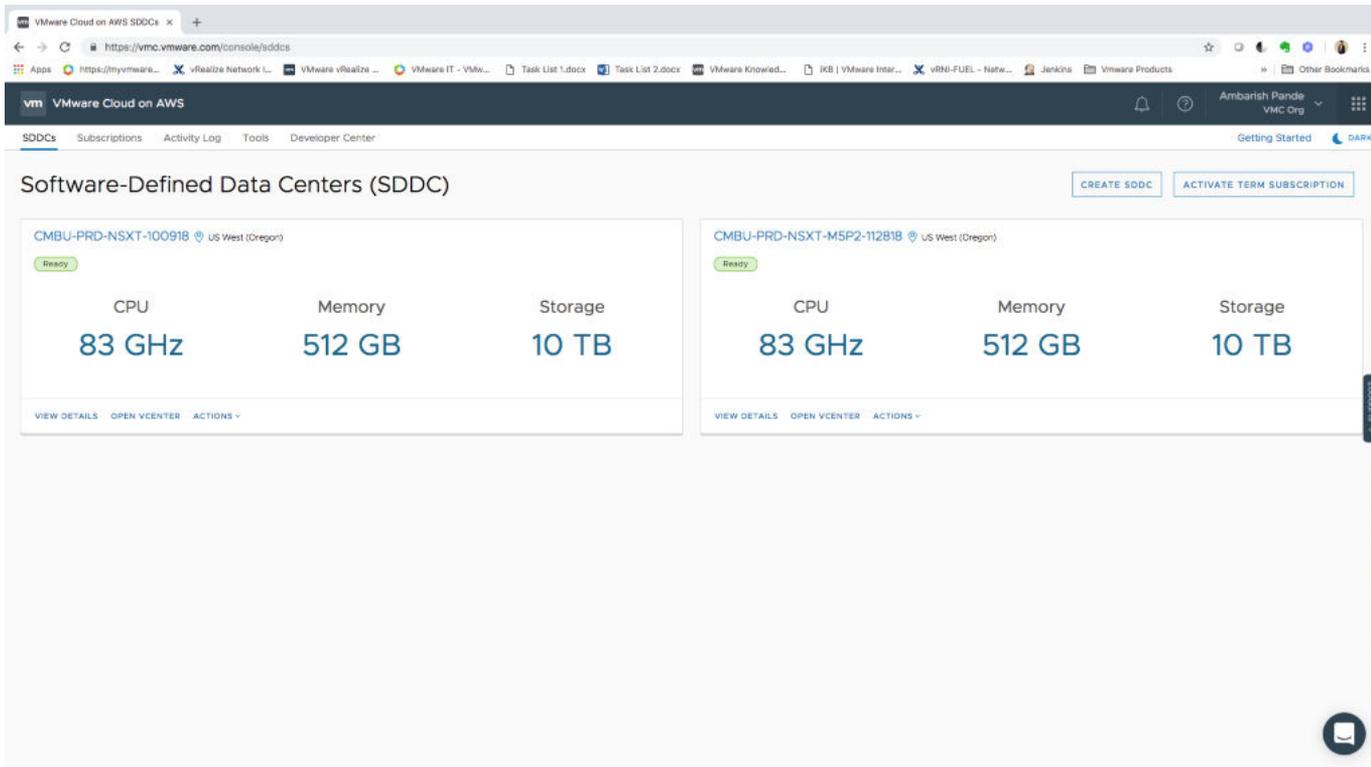
- 1 Log in to the VMware Cloud services console.



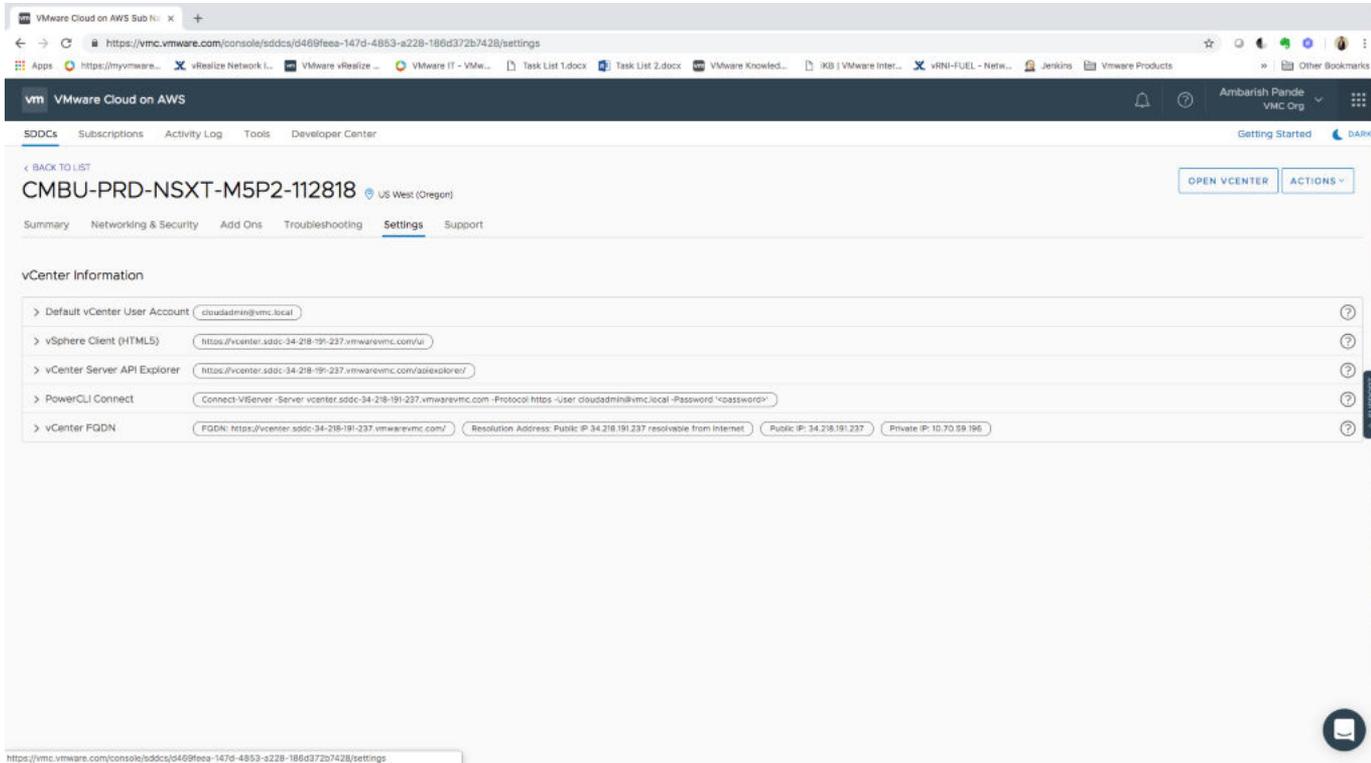
- 2 Under My Services, click VMware Cloud on AWS.



3 Select the desired Software-Defined Data Center (SDDC).



4 Click on the **Settings** tab.

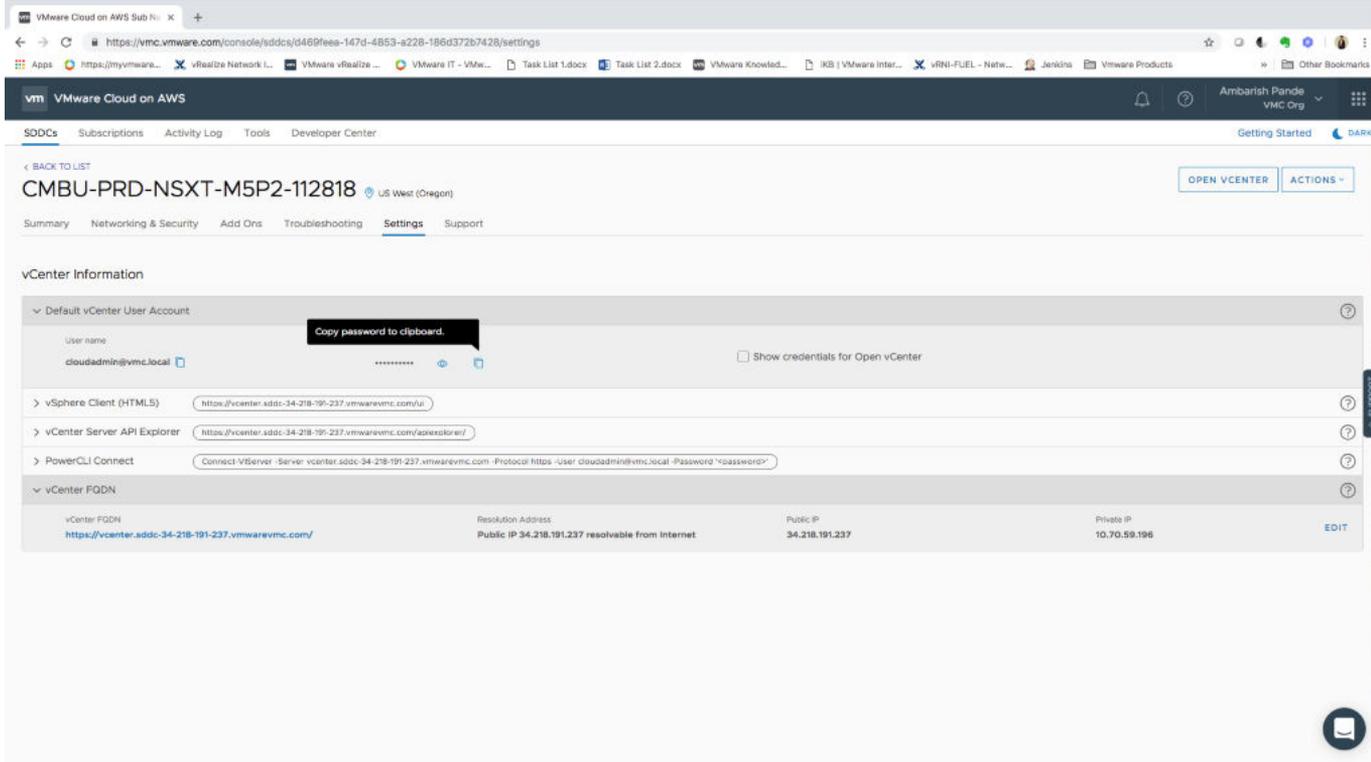


5 Expand vCenter FQDN.

Make a note of the VCenter FQDN details.

6 Expand Default vCenter User account to get the user name and password.

Copy the password and make a note of the user name.



Compute Gateway Firewall Rule

When communicating with the vRealize Network Insight platform; the collector requires HTTPS port 443 to be open for outgoing traffic.

Following VMware hosted URLs are accessed by the collector through the firewall:

- *.vmwareidentity.com
- gaz.csp-vidm-prod.com
- *.vmware.com
- *.ni-onsaas.com

Additionally, NTP and DNS traffic should be allowed for the correct functioning of the vRealize Network Insight or vRealize Network Insight collector.

Create a firewall rule with the following details:

- Name: An appropriate descriptive name
- Source: The name of the VMware Cloud on AWS Group containing the collector IP address.

- Destination: Select **ANY**
- Services – Select **HTTPS, DNS, DNS-UDP, NTP, ICMP**
- Action – **Allow**
- Applied To – **Internet Interface**
- Logging – Enable logging, if required.