



ADMINISTRATOR GUIDE

Network Performance Monitor

Version 2024.4



© 2024 SolarWinds Worldwide, LLC. All rights reserved.

This document may not be reproduced by any means nor modified, decompiled, disassembled, published or distributed, in whole or in part, or translated to any electronic medium or other means without the prior written consent of SolarWinds. All right, title, and interest in and to the software, services, and documentation are and shall remain the exclusive property of SolarWinds, its affiliates, and/or its respective licensors.

SOLARWINDS DISCLAIMS ALL WARRANTIES, CONDITIONS, OR OTHER TERMS, EXPRESS OR IMPLIED, STATUTORY OR OTHERWISE, ON THE DOCUMENTATION, INCLUDING WITHOUT LIMITATION NONINFRINGEMENT, ACCURACY, COMPLETENESS, OR USEFULNESS OF ANY INFORMATION CONTAINED HEREIN. IN NO EVENT SHALL SOLARWINDS, ITS SUPPLIERS, NOR ITS LICENSORS BE LIABLE FOR ANY DAMAGES, WHETHER ARISING IN TORT, CONTRACT OR ANY OTHER LEGAL THEORY, EVEN IF SOLARWINDS HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

The SolarWinds, SolarWinds & Design, Orion, and THWACK trademarks are the exclusive property of SolarWinds Worldwide, LLC or its affiliates, are registered with the U.S. Patent and Trademark Office, and may be registered or pending registration in other countries. All other SolarWinds trademarks, service marks, and logos may be common law marks or are registered or pending registration. All other trademarks mentioned herein are used for identification purposes only and are trademarks of (and may be registered trademarks) of their respective companies.

Table of Contents

Network Performance Monitor Administrator Guide	7
SolarWinds Platform features in NPM	8
Pre-installation hints	8
Common features	8
NPM licensing model	11
License NPM with other SolarWinds products	11
Manage and activate your NPM license	13
Manage interfaces in NPM	14
View interface status, interface health, and details about downtime in NPM	14
Display physical layout of interfaces on graphical stencils (Device View)	17
Detect and predict duplex mismatches in NPM	21
Edit interface properties in NPM	22
Suspend collecting data for interfaces in NPM	23
Delete interfaces in NPM	25
Remotely manage monitored interfaces in NPM	27
Troubleshoot nodes and interfaces that are Unknown	28
Observe real-time data for nodes and interfaces on charts	31
Monitor capacity usage trends on the network and forecast capacity issues in NPM	33
Forecast capacity for nodes, interfaces, or volumes in NPM	34
Change capacity forecasting settings globally in NPM	35
Customize capacity forecasting settings for single nodes, interfaces, or volumes in NPM	36
Discover your network paths	39
Key features of NetPath™	39
How does NetPath™ work?	39
NetPath requirements for NPM	40
Create a NetPath service in NPM	44
Create a NetPath probe in NPM	46
View a network path in NPM	48
Troubleshoot a NetPath service with external path data in NPM	52

Troubleshoot my network with NetPath data in NPM	54
SolarWinds Platform integration with NetPath	57
Monitor Cisco ACI devices in NPM	59
Requirements	59
Add ACI devices and enable ACI polling	60
Enable polling for Cisco ACI on a monitored node	61
View health scores for ACI members	61
View health score and status history in PerfStack	62
View your SDN infrastructure on Intelligent Maps	62
Create ACI-specific alerts and reports	65
Monitor ASA firewalls with NPM	70
Set up monitoring Cisco ASA firewalls in NPM	70
Understand ASA platform health in NPM	74
Monitor VPN tunnels on ASA firewalls in NPM	80
Review access lists on ASA firewalls in NPM and NCM	82
Monitor Cisco Nexus devices in NPM	83
Set up monitoring for Cisco Nexus devices in NPM	84
Access Nexus-specific information in NPM	89
Monitor Cisco SwitchStack in NPM	91
View stack members and rings	91
View the health of stack members	92
Cisco SwitchStack events	92
Out-of-the-box alerts for SwitchStack	93
Network Insight for F5 BIG-IP load balancers in NPM	95
Set up Network Insight for F5® BIG-IP® load balancers in NPM	95
Monitor services delivered by F5® BIG-IP® load balancers in NPM	98
Status of F5 devices in NPM	101
F5 high availability in NPM	104
F5 health monitors in NPM	105
Events, alerts, and reports for Network Insight for F5® BIG-IP® load balancers in NPM	106
Take an F5 pool member out of rotation in NPM	107
Network Insight for Palo Alto - monitor Palo Alto firewalls with NPM	110

Requirements	110
How to monitor Palo Alto devices	110
View Site-to-Site tunnels on a Palo Alto firewall	112
View the GlobalProtect VPN subview	113
PerfStack	114
Intelligent Maps	114
Troubleshooting	115
Monitor wireless networks in NPM	116
View wireless data in NPM	116
Monitor Meraki wireless infrastructure in NPM	117
Monitor Arista Wireless Manager infrastructure	123
Monitor Aruba Central wireless infrastructure	126
Monitor Extreme Networks Cloud IQ wireless infrastructure	131
Monitor Juniper Mist wireless infrastructure	134
Monitor Ruckus One wireless infrastructure	139
Monitor SD-WAN for Aruba EdgeConnect orchestrators (formerly Silver Peak) with SolarWinds Observability Self-Hosted	143
Monitor SD-WAN for Fortinet FortiManager orchestrators with SolarWinds Observability Self-Hosted	149
Monitor SD-WAN for Meraki organizations with SolarWinds Observability Self-Hosted	155
Monitor SD-WAN for Prisma orchestrators with SolarWinds Observability Self-Hosted	165
Monitor SD-WAN for VeloCloud orchestrators with SolarWinds Observability Self-Hosted	171
Monitor SD-WAN for Viptela orchestrators with SolarWinds Observability Self-Hosted	177
Create wireless heat maps for NPM	183
Display wireless heat maps for NPM in the SolarWinds Platform Web Console	189
Monitor EnergyWise devices in NPM	190
Add the EnergyWise Summary View to the SolarWinds Platform Web Console menu bar in NPM	190
Temporarily reset the current power level of a monitored EnergyWise interface in NPM	190
Monitor Azure V-Nets, V-Net Gateways and Site-to-Site Connections with NPM	192
Add an Azure cloud account for monitoring	192
Display the overview of monitored cloud environment	195
Monitor Virtual Network Gateways on Microsoft Azure clouds	196

Monitor Site-to-Site Connections on Microsoft Azure clouds	197
Configure Azure settings relevant for NPM	198
Set NPM thresholds	199
Create custom monitors in NPM	200
Management Information Base (MIB) in the NPM	201
Monitor custom statistics based on OIDs with Universal Device Pollers in the NPM	201
Manage unique devices on the network with NPM	218
Troubleshoot NPM issues with Performance Analysis dashboards	229
Troubleshoot intermittent network slowdowns with NPM	229
Troubleshoot slow resources in a branch office with NPM	233

Network Performance Monitor Administrator Guide

Welcome to the SolarWinds Network Performance Monitor (NPM) Administrator Guide.

This guide provides an overview of product features and related technologies. In addition, it contains recommendations on best practices, tutorials for getting started with advanced features, and troubleshooting information for common situations.

For information about planning, installing and getting started with NPM, see the [NPM Getting Started Guide](#).

SolarWinds Platform features in NPM

The [SolarWinds Platform](#) is the core of the SolarWinds IT Management Portfolio. It provides a stable and scalable architecture that includes data collection, processing, storage, and presentation. The SolarWinds Platform provides [common features](#), such as user accounts and groups, views, dashboards, reporting, alerting, and more that you can use across all [SolarWinds Platform products](#) and access from the SolarWinds Platform Web Console.

Pre-installation hints

Before you install your SolarWinds Platform products, review the following details:

SolarWinds Platform requirements

[Hardware, software, and port requirements](#) for the SolarWinds Platform server and SolarWinds Orion database.

Licensing

[Licensing differs among SolarWinds Platform products. Activate, add, upgrade or assign licenses](#) with the License Manager in the SolarWinds Platform Web Console.

Installation or upgrade

Use the [SolarWinds Orion Installer](#) to easily install or upgrade multiple SolarWinds Platform products simultaneously.

While installing your SolarWinds Platform products, you might need to [configure SSL for the SolarWinds Platform Web Console](#) or [enable FIPS](#).

Common features

The following features are available in SolarWinds Platform products.

Learn SolarWinds Platform basics

[Log in to your SolarWinds Platform product in a web browser](#) and [meet the SolarWinds Platform Web Console](#).

Review [Events](#), [syslogs](#), or [SNMP traps](#) to know what's going on.

[Get alerts](#) about issues in your environment.

Generate [reports](#) to present the status of the monitored environment.

Review [Performance Analysis dashboards](#), also known as PerfStack™.

[Create, edit, and maintain SolarWinds Platform Web Console user accounts](#) - set user rights, reset passwords, limit access to network segments, and enable authentication with Active Directory.

[View monitored objects on maps in the SolarWinds Platform Web Console](#) - view automatically generated Intelligent Maps as a subview, display objects with their location specified in the OpenStreet format in a widget, or create maps the Network Atlas tool and display them in the SolarWinds Platform Web Console.

Add devices for monitoring and manage monitored devices

Specify which devices to monitor and the information you need, then select the way you get this information. See [Discover and add devices](#).

[Add single nodes](#), [use Active Directory domain controllers to add nodes](#), or [discover devices](#) on your network automatically.

Available polling methods include ICMP, WMI, SNMP, or [agents](#) deployed on Windows, Linux, and UNIX devices.

[Manage monitored devices](#) - edit properties, set the polling method for monitored devices, toggle monitoring on and off, or mute alerts for nodes.

Customize your SolarWinds Platform Web Console

[Customize SolarWinds Platform Web Console](#) - customize dashboards, colors, logo, views, widgets and charts. Learn how to limit what objects users see on views, or specify what you want to see on views for specific device types.

[Create custom properties](#) - create custom fields to associate with monitored network objects and display custom information for monitored devices.

[Create groups and dependencies](#) - organize how monitored data is presented in the SolarWinds Platform Web Console. Set up dependencies to better represent the relationships between network objects and account for constraints on the network.

[Set thresholds](#) - specify thresholds for monitored metrics. Customize general thresholds or use baselines.

Monitor additional metrics and devices

Monitor [hardware health](#) - get insight into hardware issues on the network. Monitor hardware health based on hardware sensors, such as fan status, power supply status, or temperature.

Monitor [virtual environments](#) - monitor your virtual networks (VMware® ESX and ESXi servers, VMware vCenter®) in the SolarWinds Platform Web Console.

[Quality of Experience](#) - use packet analysis sensors to see packet-level traffic information about key devices and applications on your network.

Gain view into details provided by [Cisco UCS devices](#).

Expand the SolarWinds Platform functionality or scale your deployment

Use [SolarWinds High Availability](#) (HA) to provide failover protection for your SolarWinds Platform server and additional polling engines to reduce data loss.

[Integrate a SolarWinds Platform product with the ServiceNow or SolarWinds Service Desk trouble ticketing systems.](#)

Do you need to scale your deployment? See [Scalability Engine Guidelines](#).

Review the [tips for optimizing your deployment](#).


Balance the load on polling engines by [specifying nodes to be polled by individual polling engines](#).

[Manage Additional Polling Engines.](#)

[Troubleshoot your SolarWinds Platform database.](#)

NPM licensing model

The NPM license is based on a number of items to monitor. Each license tier number provides the maximum limit of nodes, interfaces, and volumes to manage and monitor.


 SolarWinds Platform products support both perpetual licenses and subscription licenses. See [License types](#) in the SolarWinds Platform help for details.

NPM is licensed according to the largest number of the following types of monitored network elements:

- **Nodes:** any devices being monitored, such as routers, switches, virtual and physical servers, access points, and modems.
- **Interfaces:** any single points of network traffic, such as switch ports, physical interfaces, virtual interfaces, sub-interfaces, and VLANs.
- **Volumes:** any logical disks being monitored.

NPM has the following available license levels:

License	Number of monitored elements
SL100	Up to 100 nodes, 100 interfaces, and 100 volumes (300 elements in total).
SL250	Up to 250 nodes, 250 interfaces, and 250 volumes (750 elements in total).
SL500	Up to 500 nodes, 500 interfaces, and 500 volumes (1500 elements in total).
SL2000	Up to 2000 nodes, 2000 interfaces, and 2000 volumes (6000 elements in total).
SLX	Virtually unlimited number of elements. With the default polling interval, one polling engine can monitor a maximum of 12,000 elements (the sum of nodes, interfaces, and volumes). To monitor over 12,000 elements, use additional polling engines (APEs) . Each APE requires a license.

 Database size increases with the addition of monitored elements.

License NPM with other SolarWinds products

Your NPM license interacts additively with your other SolarWinds licenses.

For example, if you have an NPM SL500 (500 nodes and 500 volumes) installed with SAM AL150, you can monitor:

- 650 nodes (500 NPM nodes + 150 SAM nodes)
- 650 volumes (matching the node count)
- 500 interfaces monitored with SNMP
- 150 component monitors
- An unlimited number of interfaces polled using WMI

To verify the number of consumed and available component monitors in your license, access the NPM License Summary.

1. Log on to the SolarWinds Web Console with an administrator account.
2. Click Settings > All Settings.
3. Click License Details in the Details section.

MAIN ORION SERVER DETAILS	
Orion	
Module Name	Orion Platform
Version	2017.1
Service Pack	None
Nodes currently monitored	20
Total nodes in license	unlimited
Volumes currently monitored	0
Total volumes in license	unlimited
Total HA Pools in use	0
Total HA Pools in License	0
NPM	
License	i 25 day(s) left in evaluation
Product Name	Network Performance Monitor
Version	12.1
Service Pack	None
Current number of interfaces	54
Allowed number of interfaces	unlimited

Review the Orion and NPM details. You can see the current number of monitored nodes, volumes and interfaces and the number of the total number of elements allowed by your license.

Manage and activate your NPM license

During installation, you will be prompted to activate your NPM license. You need the license key located in the SolarWinds Customer Portal. For more information on licensing NPM, see the [web-based License Manager information](#).

Manage interfaces in NPM


Manage your environment up to the interfaces level.

- You specify interfaces to monitor when adding nodes for monitoring, either [after you have run the Discovery](#) or when [adding single nodes](#).
- If you cannot see an interface on a node, go to the node details view, click [List Resources](#) in the Management widget, and ensure that the interface is selected for monitoring.

View interface status, interface health, and details about downtime in NPM

The downtime information is useful, for example, for SLA providers who want to prove specific times of interface or port unavailability.

Check the [health of interfaces](#) in the Health Summary widget and get more details on the Node Details - Interfaces subview.

 In some areas, an interface being down does not directly impact Internet or intranet connectivity.


Interface status

Interface status is polled via SNMP. Starting with 2022.4, configured [thresholds for interfaces](#) (Interface Errors and Discards, Interface Percent Utilization) are also reflected in the interface status.

View interface downtime information

1. Log in to the SolarWinds Platform Web Console.
2. Navigate to the Node Details view.
3. Review the **Interface Downtime** widget.

By default, the widget shows the interface status in the last 24 hours, each hour represented as a block in color. This widget is also available on Interface Details views.

 To display downtime for all monitored interfaces on a node, add the Interface Downtime widget on node view.

4. To see a detailed view of a problematic section, position your cursor on the graph.

Change the time period

By default, the Interface Downtime widget displays downtime data for the last 24 hours, one block representing 1 hour. You can display any time frame within the stored history.

1. Go to the Interface Downtime widget, and click Edit.
2. Select Custom in the Downtime Period list, and specify the Beginning and End dates and times.
3. When displaying longer time periods, you might need to change the time frame represented by one block. Select Custom in Display Settings, and provide a time period represented by one block.
4. Click Submit.

Set the retention period for interface downtime history

By default, interface status history is stored in the database for 7 days.

1. Log in to the SolarWinds Platform Web Console as an administrator.
2. Click Settings > All Settings in the menu bar.
3. In the Thresholds & Polling grouping, click Polling Settings.
4. Scroll down to Database Settings, and enter a time to retain interface status history in the database in the Downtime History Retention field. Enter a value in days, from 7 to 60 days.

Disable interface downtime monitoring

Monitoring interface downtime can affect the performance of your NPM. To decrease the load, disable interface downtime monitoring. For periods where interface downtime was not monitored, the Interface Downtime widget shows gray blocks.


1. Log in to the SolarWinds Platform Web Console as an administrator.
2. Click Settings > All Settings in the menu bar.
3. In the Thresholds & Polling grouping, click Polling Settings.

4. Clear the Enable Downtime Monitoring box in the Network grouping.
5. Click Submit.

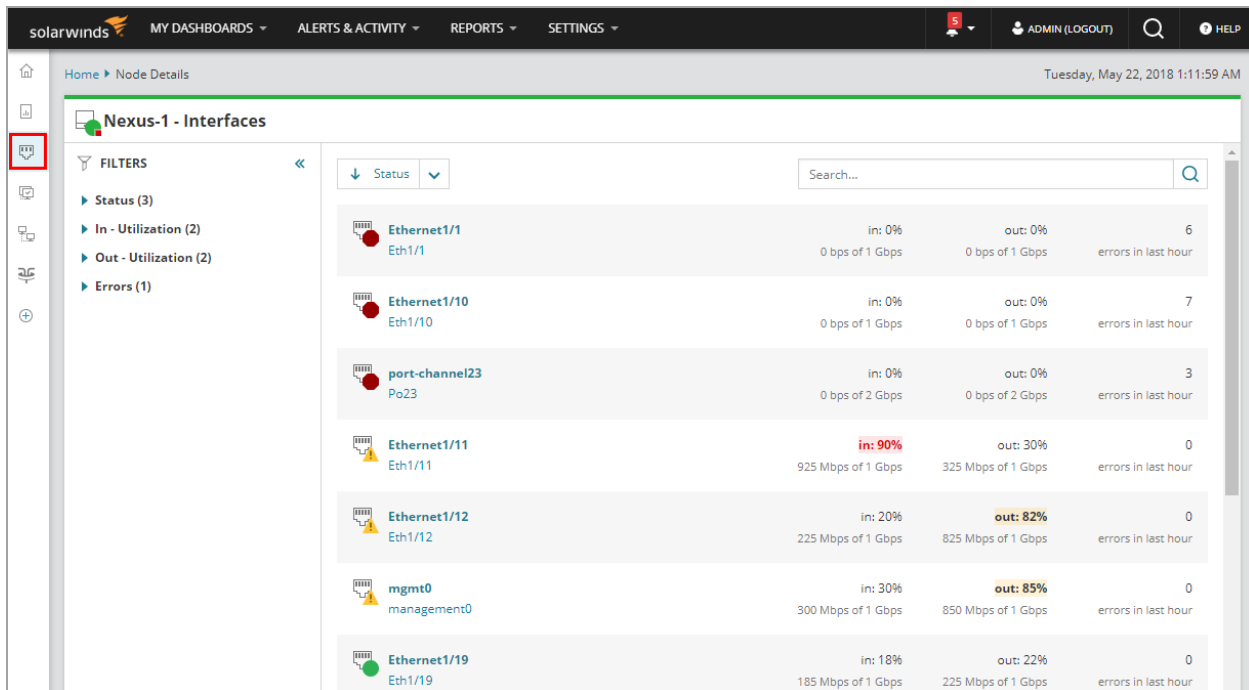
NPM does not monitor downtime for the interface any more. The Interface Downtime widget displays the message "Downtime monitoring is disabled. To enable it, go to Polling Settings."

View interface health

1. In the SolarWinds Platform Web Console, go to the Node Details view.
2. On the Summary subview, review the Health Summary widget to get an overview of the status of interfaces on the node.



3. On the Health Summary widget, click a status group number to go to the Interfaces subview, filtered by the selected group status. The Interfaces subview lists monitored interfaces on the selected device, including relevant details, such as in and out utilization, or any errors that happened in the last hour.



Display physical layout of interfaces on graphical stencils (Device View)

You can visually display interface status and utilization on the device layout for the following rack-mountable devices:

- Cisco Catalyst 2960
- Cisco Catalyst 3750

 You need to have [hardware health monitoring enabled on the device](#).

- EX 2200 Juniper switches
- EX 3300 Juniper switches

The physical layout is available on Node Details views, as the Device View tab.

1. In the SolarWinds Platform Web Console, go to the node details view for a node that is monitored using SNMP.
2. Click the Device View subview. Here, you can:
 - See the status of interfaces on the physical layout
 - Find out more information on individual interfaces
 - See interfaces on all switches in a stack
 - Display In- and Out-utilization of individual interfaces
 - Add unknown interfaces for monitoring

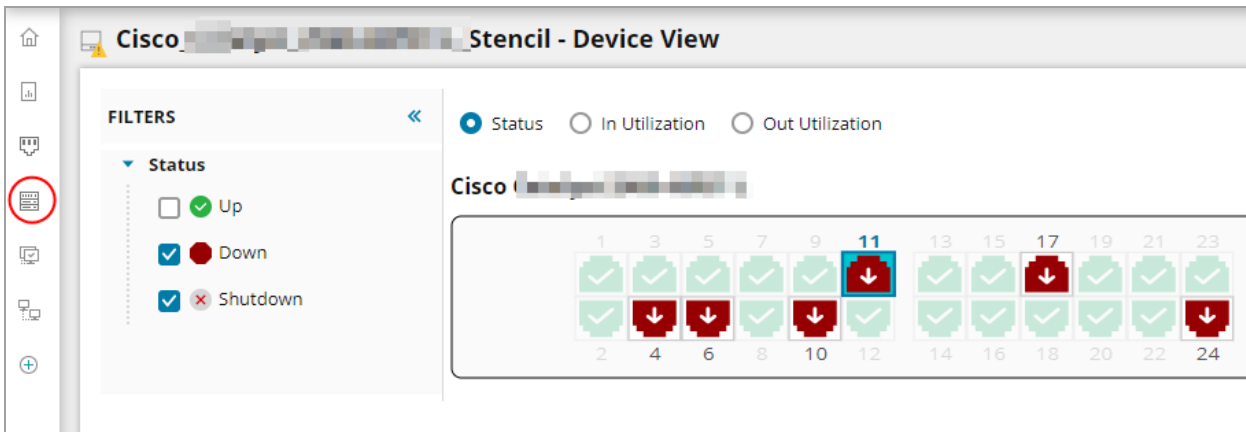
Highlight interfaces with a certain status or interfaces on a VLAN

- In the Status filter, select interface statuses to display. Filtered out interfaces become transparent.
- In the VLAN filter, select one or more VLANs to highlight interfaces that belong to the selection.

The VLAN filter only displays the first 10 options. If there are more options available, click Show all and make your selection in the Options Picker pop-up.

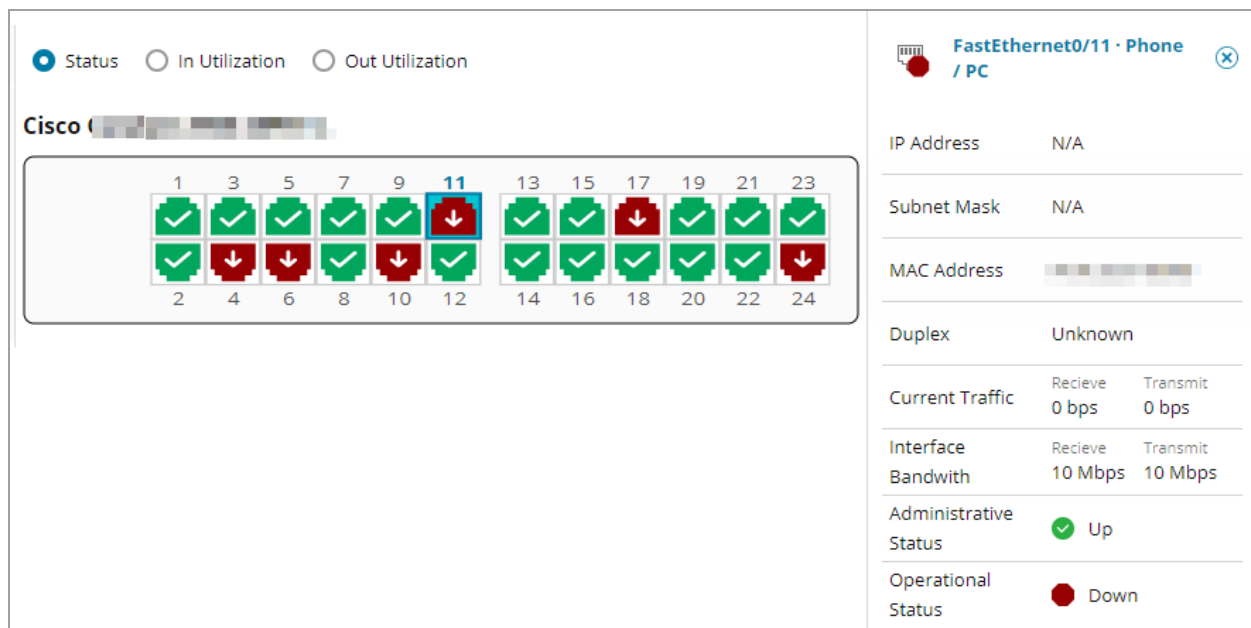
You can combine the filters.

The filters are included in the browser URL. When sharing URL links, you can change the filters by simply adjusting them in the URL.



Display status details for an interface

- Click an interface to display a panel with details on the right.
- Click the interface name in the panel to open the Interface Details view.



Add unknown interface for monitoring

Unknown interfaces are displayed as black on the Device View.

If you have Node Management rights, you can add the interface for monitoring.

1. On the Device View, click a black interface icon.
2. In the side panel, click Start Monitoring Unknown Interface. This opens the List Resource view for the device.
3. Select the box for the interface and submit your changes.

The interface status will be reflected in the Device View.




Display In and Out Utilization of interfaces

Click option buttons above the device stencil to display In or Out Utilization of interfaces.

Click an interface box to display more details about the interface.

When you click an unmonitored interface (black port icon) and have Node Management rights, you can add the interface for monitoring. Click the Start Monitoring Unknown Interface link in the panel.



Display Device View for switch stacks

If you have multiple switches in a stack, their layouts are displayed in the order they are configured in the stack.

Device View
Device View

FILTERS

◻ Status

- ✘ Shutdown 88
- ✔ Up 148
- Down 22
- ⚠ Warning 2

◻ VLAN

- 1 - det 40
- 50 - N 2
- 51 - N 1
- 99 - te 2
- 101 - l 1
- 103 - c 1
- 300 - v 3
- 400 - v 2
- 704 - l 1
- 1954 - b 1

[show all \(12\)](#)

● Status
○ In Utilization
○ Out Utilization

Cisco Catalyst 37xx Stack

WS-C3750G-12S-S (Switch #1) - not supported

WS-C3750G-12S-S (Switch #2) - not supported

WS-C3750G-48TS-S (Switch #3)

1	3	5	7	9	11	13	15
✘	✔	✘	✘	✔	✔	✘	✘
2	4	6	8	10	12	14	16
✘	✘	✘	✘	✔	✔	✘	✘

WS-C3750G-48PS-S (Switch #4)

1	3	5	7	9	11	13	15
✔	✘	✔	✔	✘	✔	✔	✔
2	4	6	8	10	12	14	16
✘	✘	✘	✘	✘	✘	✘	✘

Administrator Guide: Network Performance Monitor

page 20

Detect and predict duplex mismatches in NPM

One of the most common causes of performance issues on 10/100 or 100/1000 Mbit Ethernet links occurs when one port on the link operates at half-duplex while the other port operates at full-duplex.

1. Log into the SolarWinds Platform Web Console.
2. Go to the node details view for the parent node of the interface you want to check for duplex problems.
3. Consult the Possible Duplex Mismatches widget. If there are no errors, the widget is hidden.

The widget lists all duplex interfaces on the node, the percentage of transmit and receive errors, and the neighboring node and interface. If the neighboring interface or node is not monitored, the appropriate columns are empty.

The last column displays the duplex mode issue - Mismatch, or Unknown.

Duplex Mismatch

To be able to detect duplex mismatches, your nodes need to meet the following requirements:

- The nodes must be monitored.
- The nodes must be in the up state during the discovery.
- The nodes must support topology and be interconnected.
- Duplex of both devices must be identified as full or half.

The widget shows all duplex mismatches, not only 100% duplex mismatches. These are reported on by the Duplex Mismatch alert.

Possible Duplex Mismatch

If at least one of the link interfaces has the duplex mode defined as half or full, the widget helps you identify possible mismatches.

Possible duplex mismatches are visible in the duplex mode column as the Unknown duplex mode. They are identified in the following cases:

- If the switch port reports more than 0.5% receive or transmit errors.
- If the switch port reports CRC errors.
- If the switch port reports Late Collision errors.

How do I resolve mismatches?

To resolve a duplex mismatch, make sure your hardware is working, and unify the duplex mode configuration on neighboring interfaces.

Troubleshoot duplex mismatches

The Possible Duplex Mismatches does not display on Node Details view

If the widget does not display on the node details view, there might be a performance issue due to the amount of interfaces and topology connections. Check the following logs for mismatch information:

```
C:\ProgramData\SolarWinds\Logs\Orion\OrionWeb.log
```

```
C:\ProgramData\SolarWinds\InformationService\v3.0\Orion.InformationService.log
```

The Possible Duplex Mismatches widget does not display percentage of errors

Possible causes:

- No statistical data for these interfaces.
- A performance issue connected with getting statistic information for the widget.

Edit interface properties in NPM

1. Log in to the SolarWinds Platform Web Console as an administrator.
2. Click Settings > Manage Nodes.
3. Locate the parent node of the interface you want to manage, and expand the parent node.
4. Select the interface, and click Edit Properties.
5. Make your changes:

Edit the interface name

Adjust the interface name.

- In interface names, aliases, or descriptions, use only the following recommended characters:
a-z A-Z 0-9 space , . - _ () /
- Do not use \ | : * ? , or angle brackets (< or >). Angle brackets and any strings contained within angle brackets are removed during polling, as bracketed text may be incorrectly parsed as web markup tags.
- If you change an interface name, alerts created before the name change will still display the old name. Alerts created after the name change will display the new interface name.

To display the interface as unplugged rather than down, select Display Interface as Unplugged.

Designate bandwidth for the interface

Default transmit and receive bandwidths are 1000 Mb/s. If a device does not report its bandwidth, or the interface bandwidth is constrained by other network devices, specify a custom bandwidth that reflects the performance of the interface.

Select Custom Bandwidth, and provide values for Transmit and Receive Bandwidth, in Mb/s.

Change polling interval

Edit how often NPM polls the interface status and performance data.

Interface Status Polling is the interval in seconds between the status checks on the selected interface. By default, interface status is checked every 120 seconds.

Collect Statistics is the interval in minutes on which performance statistics for the interface are determined. By default, it is every 9 minutes.

Custom properties and dependencies

Provide values for custom properties for the interface, and edit dependencies. See [Creating Custom Properties](#) and [Network object dependencies](#).

Customize alerting thresholds for the interface

You can customize thresholds whose reaching triggers alerts for individual interfaces. You can change alerting thresholds for the following metrics on the interface:

- Received /Transmit Interface Errors and Discards
- Receive/Transmit Interface Utilization

To customize a threshold, select **Override Orion General Thresholds** next to the metric, and provide values for **Warning** and **Critical** Thresholds.

Starting with NPM 2022.4 or Hybrid Cloud Observability 2022.4, you can specify sustained thresholds to specify how long the condition must be true for the threshold to be exceeded. See [Customize thresholds for single objects in the SolarWinds Platform](#).

6. Click Submit.

The interface properties in change according to your updates.

Suspend collecting data for interfaces in NPM

Monitored interfaces are regularly polled for operational status, and collected statistics are displayed in the SolarWinds Platform Web Console.

Maintenance mode

To temporarily stop collecting data or triggering alerts for interfaces, put the interface or [the parent node into a maintenance mode](#).

1. Go to Manage Nodes, and navigate to the interfaces.
2. Select the interfaces, and select a maintenance mode option:
 - Mute alerts: data for the interface is collected, but alerts do not trigger.
 - Stop collecting data: data for the interface is not collected and alerts do not trigger.
 - Schedule a maintenance period: specify a period of time to stop collecting data or mute alerts for the interface.

The maintenance mode settings change according to your settings. For information about resuming alerts, starting collecting statistics, or editing the scheduled maintenance, see the [section on Maintenance Mode for nodes](#).

Set the interface status as Unpluggable

If you do not want to be notified when an interface is down, you can specify that the interface is Unpluggable. The interface status is reflected in the status of the parent node and in alerts.

1. On the Node Management view, select the interface, and click Edit Properties.
2. Select Display Interface as Unplugged Rather Than Down, and click Submit.

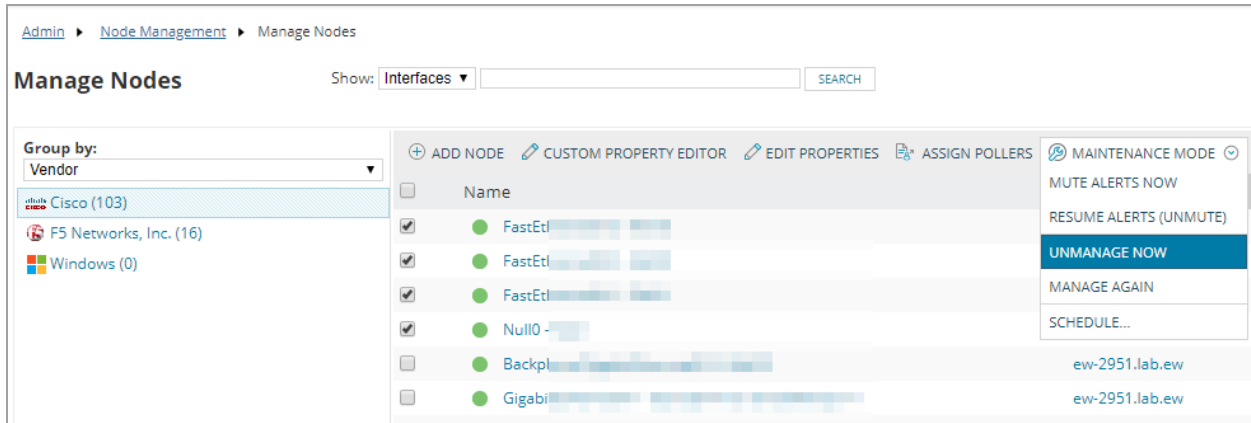
The interface status does not influence the status of the parent node.

Unmanage an interface

If you do not want to poll any data for an interface, unmanage it.

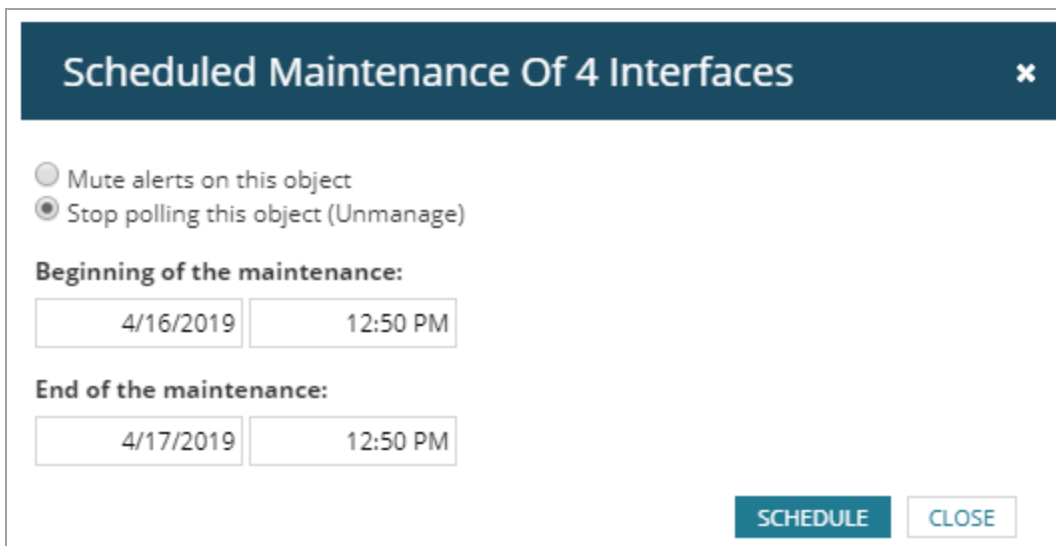
1. Click Settings > Manage Nodes.
2. Ensure that the Show drop-down list is set to Interfaces.
3. Select Interfaces to unmanage.

4. To unmanage interfaces immediately, click Maintenance Mode > Unmanage now.



The selected interface(s) are immediately unmanaged.

5. To schedule a period when the interfaces will be unmanaged:
 - a. Click Maintenance Mode > Schedule.
 - b. Select Stop polling this object (unmanage).
 - c. Specify the Begin and End time for the Interfaces to be unmanaged.
 - d. Click Schedule.



During the specified period of time, no data will be polled for selected interfaces.

Delete interfaces in NPM

To temporarily stop collecting new data for an interface, [unmanage the interface](#) or configure a period when no data will be collected for the interface ([schedule maintenance](#)).

To stop monitoring an interface, remove the interface from monitored resources on the node or delete the interface:

- Remove the interface from resources monitored on the parent node: On the node details view, click List Resources in the Management widget, and clear the interface box.
- If you cannot find the interface in resources monitored on the node, you have probably changed its details and there might be issues with mapping the interface. In this case, [delete and re-add the interface](#).

Delete interfaces from Manage Entities page

1. Click Settings > Manage Nodes.

To view interfaces on a node, click the arrow at the end of a row. Interfaces display in the Related Entities pane.

i To get to Manage Entities from Manage Nodes, click the New Manage Entities page link at the top.

To get back to Manage Nodes, click Commands > Switch Back in the top right corner of the page.

2. Select interfaces to delete, and click Delete.

Selected interfaces are deleted immediately. All references to these interfaces are cleared during the database maintenance (by default at 2:15 a.m.).

The screenshot displays the 'Manage Entities' page in SolarWinds NPM. At the top, there are navigation options: '+ ADD NODE', 'CUSTOM PROPERTY EDITOR', and 'MORE'. Below this is a search bar and a table of nodes. The table has columns for selection, name, and a right-pointing arrow. A red circle highlights the 'X' icon at the end of the second row. To the right, the 'RELATED ENTITIES' pane is open, showing a 'DELETE' button circled in red, with the text '3 item(s) selected.' below it. The pane also has filters for 'Interfaces' and 'Type', and a search bar. The list of related entities includes 'OTHER (2)', 'Console9/0/0', 'NULL0', 'ETHERNET (79)', 'Eth-Trunk1', and 'GigabitEthernet0/0/13'.


Delete interfaces from Manage Nodes

1. Log in to the SolarWinds Platform Web Console, and click Settings > Manage Nodes.
2. In the Show drop-down, select Interfaces.
3. Use the options in the Group by list or the Search field to locate interfaces to delete.
4. Select the interfaces and click Delete in the menu bar.


Selected interfaces are deleted immediately. All references to these interfaces are cleared during the database maintenance (by default at 2:15 a.m).

Remotely manage monitored interfaces in NPM


Using the Node Management utility, you can shut down or enable interfaces, and remotely override configured EnergyWise power settings.

 To manage interfaces remotely, the parent node must have not only a Community String, but also the Read/Write Community String set correctly. See [Edit polling settings](#).

1. Log in to the SolarWinds Platform Web Console as an administrator.
2. Click Settings > Manage Nodes.
3. Expand the parent node of the interface.

 To find the node, use the filter and search tools above the nodes list.

4. Select the interfaces to manage.
5. To shut down the interfaces, click More Actions > Shut Down, and click OK to confirm.
6. To enable the interfaces, click More Actions > Enable.
7. If the selected interface is EnergyWise-enabled, you can override the current power level setting. Click More Actions > Override Power Level, set the power level, and click OK.

 Remote overrides are temporary and reset in accordance with your configured EnergyWise policy for the selected interface. See [Temporarily reset the current power level of a monitored EnergyWise interface in NPM](#).

Troubleshoot nodes and interfaces that are Unknown

This article provides information about the device status in NPM, how it is polled, and what you can do if your nodes or interfaces have the Unknown status.

- [Device status](#)
- [Unknown status](#)
- [No data even though the node is Up](#)

Device status

By default, node status is detected using ICMP. NPM sends a ping request and if the response is not returned, it places the node into the Warning state and fast-polls the device for 120 seconds. If the node still does not respond, you are notified you that the node is Down.

i ICMP only tells you the NPM did not receive a response for the ping request. The device could be down, but there might also be a routing problem, an intermediary device could be down, or something could have blocked the packet on its way to or from the device. See [Get more details about the node](#) in the NPM Getting Started Guide for more details.

Status of sub-elements, such as interfaces and volumes, is detected using SNMP. This is more accurate, because the device tells you that the sub-element is Down.

Unknown status

Unknown status tells you that the node, or the sub-element (interface or volume) cannot be polled for response time and status.

When you add a node to the SolarWinds Platform or remanage a node, its status is Unknown until the data is polled. By default, the interval for polling status is two minutes.

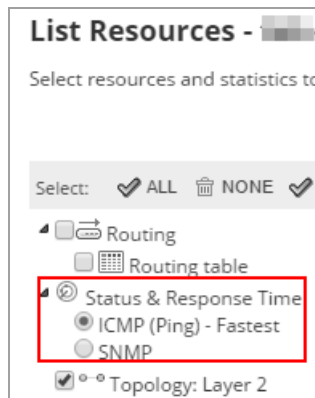
If a node stays in the Unknown status longer than the polling interval, perform the following:

- Ping the device from the polling engine the [device is assigned to](#) using the command line. If you are able to reach the device, status and response time data should also be available for the SolarWinds Platform.
- [Verify that the status and response time are polled through ICMP \(ping\).](#)
- If you are polling the node through SNMP, [verify that the polling method is set correctly.](#)
- [Make sure the status poller is enabled.](#)
- [For Unknown interfaces, re-add the interface for monitoring.](#)

Verify that the status and response time are polled through ICMP

Always use ICMP to poll status and response time when possible. Poll status and response time using SNMP only when ICMP is disabled and SNMP is enabled on the device.

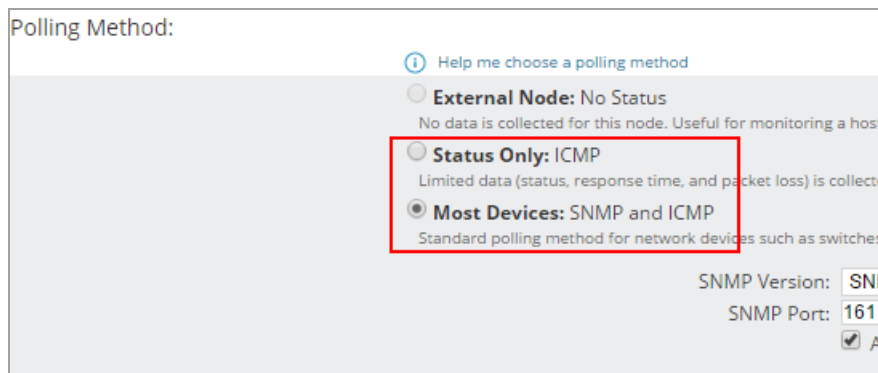
1. On the node details view, click List Resources in the Management resource.
2. In the Status & Response Time option, verify the option. If SNMP is selected, click ICMP (ping), and click Submit. You return to the Node Details view.



After two minutes, the node status is Up.

Verify the polling method for the node

1. On the node details view for the node, click Edit Node in the Management resource.
2. In the Polling Method area, select Most Devices: SNMP and ICMP, provide version, port, and credentials.



3. Click Submit.

After two minutes, the node status is Up. After ten minutes, node statistics are updated.

Verify that the status poller is enabled

1. Log in to the SolarWinds Platform Web Console using an Administrator account.
2. Click Settings > All Settings > Manage Pollers.
3. In the Local Poller Library tab, click SolarWinds Native Poller, and locate the Status & Response Time ICMP poller.

LOCAL POLLER LIBRARY		THWACK COMMUNITY POLLERS	
GROUP BY: Author			
All (37) SolarWinds Native Poller (19) SolarWinds (18)			
<input type="checkbox"/>	Name ▲	Author	Assignments
<input type="checkbox"/>	CPU & Memory	SolarWinds Native Poller	25 Nodes
<input type="checkbox"/>	EnergyWise	SolarWinds Native Poller	2 Nodes
<input type="checkbox"/>	Fibre Channel	SolarWinds Native Poller	1 Node
<input type="checkbox"/>	Hardware Health Sensors	SolarWinds Native Poller	5 Nodes
<input type="checkbox"/>	Multicast Routing	SolarWinds Native Poller	Unassigned
<input type="checkbox"/>	Node Details	SolarWinds Native Poller	33 Nodes
<input type="checkbox"/>	Routing	SolarWinds Native Poller	5 Nodes
<input type="checkbox"/>	Status & Response Time Agent	SolarWinds Native Poller	Unassigned
<input type="checkbox"/>	Status & Response Time ICMP	SolarWinds Native Poller	30 Nodes
<input type="checkbox"/>	Status & Response Time SNMP	SolarWinds Native Poller	Unassigned

4. Click the item in the Assignments column for the poller. A list of nodes where the poller is enabled opens.
5. Search for your node and ensure that the Poller Status is ON.

Node ▲	Scan Result	Poller Status
<input type="checkbox"/>	Multiple matches » Number of other pollers matching this node: 1	ON <input type="checkbox"/>
<input type="checkbox"/>	Multiple matches » Number of other pollers matching this node: 1	ON <input type="checkbox"/>
<input type="checkbox"/>	Multiple matches » Number of other pollers matching this node: 1	ON <input type="checkbox"/>

6. If the poller status is OFF, select the node, and click Enable Poller.

After two minutes, the node status is Up.

Unknown interfaces

For interfaces, the Unknown status might be caused by the following reasons:

- NPM repeatedly did not get results from a status poll.
- The Unknown status was polled on the interface.

- The interface details changed on the device, and NPM is not able to map the interfaces correctly.

i In this case, the interface index for the interfaces in the database has the value -1.

To troubleshoot the issue when NPM is not able to map an interface correctly, re-add the interfaces:

i When you re-add an interface, you lose all historical data for the interface.

1. Log in to the SolarWinds Platform Web Console using an Administrator account.
2. Click Settings > Manage Nodes.
3. Expand the node, select the interface, and click Delete.
4. Select the node, and click List Resources.
5. On List Resources, select the interface, and submit your changes.

The interface is re-added with updated details, and the interface status reflects the polled status.

Data not populating from the device even though the node is Up

If the node is Up and you cannot find data in the SolarWinds Platform Web Console, try the following troubleshooting options:

- Wait ten minutes after you add a device for the SolarWinds Platform to complete the polling cycle, and then refresh your screen.
- Verify that the device was not in [maintenance](#) (unmanaged state) during the requested time period. Click Settings > Manage Nodes, and review the Status for the device.
- [Verify that interfaces or volumes are monitored correctly](#). In Node Management, select the node, and click List resources and confirm the interfaces or volumes are selected.
- Verify that hardware health sensors on the device provide both status and values. If they provide only status, there is no data to be displayed in charts.

Observe real-time data for nodes and interfaces on charts

The following charts support featured data in real time:

- CPU Load & Memory Usage - Real Time Data (Node Details view, Vital Stats)
- Percent Utilization - Real Time Date (Interface Details view)

Observe real-time data on charts

- To see CPU Load & Memory Usage, go to a node details view and click the Vital Stats subview.
- To see Percent Utilization, open an interface details view.

When you open a page with a real-time chart, real-time polling starts automatically. By default, the data is updated every two seconds.

When you leave the page, real-time polling stops in 120 seconds.

i Real-time charts show data for a constant time frame of ten minutes. You cannot change the time frame.

I cannot see the real-time chart on my node/interface details page

If there is no data to display, real-time charts are not displayed.

- Open the node in PerfStack. If you can see real-time data there, you should be able to see real-time data on the real-time chart.
- Make sure the node/interface is monitored by NPM for a few minutes.

Real-Time Data charts do not display immediately when you add the node/interface because the identification of real-time polling capabilities takes some time, usually it is around 5-10 minutes.

i If there is data to display, an administrator might have removed the chart from the view. You can [add the widget](#).

I cannot see both metrics on the real-time chart

If your devices do not support both charted metrics, real-time charts only display the data for the supported metric.

I want the charts to update less frequently.

Change the refresh time:


1. Click Settings > All Settings, and then click Web Console Settings.
2. Scroll down to Modern Chart Settings, and adjust the Real time charts refresh time.

Monitor capacity usage trends on the network and forecast capacity issues in NPM

Capacity forecasting is available for the following metrics of monitored nodes, interfaces, and volumes monitored:


- CPU utilization on nodes
- Memory usage on nodes
- Space usage on volumes
- Receive (in) utilization on interfaces
- Transmit (out) utilization on interfaces

Capacity usage trends are calculated based on historical data. By default, the longest time period taken into account for calculating the capacity forecast is 180 days.

 The more historical data up to 180 days are available, the more precise is the calculated forecast.

Forecast calculation methods

- **Peak calculation** forecasts trends using daily maximum values. This method is suitable for important devices and connections where it is important to completely avoid reaching a certain usage level (threshold).
- **Average calculation** forecasts trends using daily average values. This method is suitable for non-critical network devices or connections where short periods exceeding the threshold level are acceptable.

 By default, the forecast calculation method is set globally for all monitored objects. You can also customize the method for individual objects (nodes, interfaces, or volumes).

Requirements

Capacity forecasting is available for nodes, interfaces, and volumes that meet the following requirements:

- The nodes, interfaces, and volumes must be managed in NPM.
- You need to have enough historical data in the database. By default, 7 days of data are required.

Forecast capacity for nodes, interfaces, or volumes in NPM

Consult graphs or tables to see usage trends of devices on your network, and find out when the capacity of the devices will be fully used.

Locate pending capacity problems

Consult the Top XX Capacity Problems widget to see a list of objects whose usage trend is rising.

If the widget is not in a view, [add it](#).

View capacity usage trends and forecast in graphs

To see a graphical display of capacity usage trends, go to the details view for the node, volume, or interface, and consult the forecast chart:

- CPU Capacity Forecast Chart
- Memory Capacity Forecast Chart
- Storage Capacity Forecast Chart
- Interface Utilization Receive Forecast Chart
- Interface Utilization Transmit Forecast Chart

View capacity usage trends and forecast in tables

For a brief overview of usage trends for a node, volume, or interface, go to the details view for the object, and consult the widget:

- Node Capacity provides an overview of both CPU load and percent memory usage in the past 6 months, a forecast when the warning and critical thresholds will be exceeded, and when the resource will be fully used.
- Volume Capacity provides an overview of volumes capacity usage in the past 6 months, a forecast when the warning and critical thresholds will be exceeded, and when the volume capacity will be fully used.

Forecasts in this widget are calculated using the default method (peak or average) specified for the widget.

Add capacity forecasting widgets

Capacity forecasting widgets display only on views for which they are relevant. For example, interface utilization widgets can only be added on interface detail views.

1. Log in to the SolarWinds Platform Web Console and go to the view where you want to add the widget.
2. Click the pencil icon on the upper left.
3. Click Add widgets on the upper right.
4. Type "forecast" or "capacity" into the Search field.
5. Drag and drop the widget on the page where you want it to be, including in a new column.
6. When complete, click Done Adding Widgets, and then Done Editing. The view is now populated with the widgets you selected.


Change capacity forecasting settings globally in NPM

Capacity forecasting settings include the forecast calculation method and thresholds for the metrics. By default, the settings are set globally.

See [Customize capacity forecasting settings for single nodes, interfaces, or volumes in NPM](#).

Change calculation method and thresholds for nodes or volumes

1. Click Settings > All Settings, and select Orion Thresholds in the Thresholds & Polling section.

 If you are in a capacity forecasting widget, click Edit, and click Orion General Thresholds.

2. Specify values for Critical Level and Warning for the metrics:
 - AVG CPU Load for CPU usage on nodes
 - Disk Usage for volume capacity usage
 - Percent Memory Used for memory usage on nodes
3. For each metric, select the calculation method.
 - Calculate exhaustion using average daily values
 - Calculate exhaustion using peak daily values
4. Click Submit.

You have changed the method and thresholds for calculating capacity forecast for monitored nodes and volumes.

Change calculation method and thresholds for interfaces


1. Click Settings > All Settings, and select NPM Thresholds in the Thresholds & Polling section.
2. Go to the Interface Percent Utilization section, define the Critical and Warning threshold values for the metric.
3. Select the calculation method:
 - Calculate exhaustion using average daily values
 - Calculate exhaustion using peak daily values
4. Click Submit.

You have changed the method and thresholds for calculating capacity forecast for monitored interfaces.

Customize capacity forecasting settings for single nodes, interfaces, or volumes in NPM

You can set different forecast calculation methods and thresholds for individual nodes and volumes.


For interfaces, the calculation method is set globally, and you can customize only the thresholds.

 Set warning and critical thresholds for critical nodes, interfaces, or volumes to lower percentages, so that you have enough time to take measures before capacity issues occur.

Customize capacity forecasting thresholds and calculation methods for nodes

1. Log in to the SolarWinds Platform Web Console as an administrator.
2. Open the Edit Properties page for the node.

Go to Settings > Manage Nodes, select the node, and click Edit Properties.

 If you are in a capacity forecasting resource, click Edit, and click the link to the node's Edit Properties page.

3. On the Edit Properties page, scroll down to Alerting Thresholds.

4. Select **Override Orion General Thresholds for CPU Load or Memory Usage**, and define the Warning and Critical threshold levels.
5. Select the method for calculating trends:
 - Calculate exhaustion using average daily values
 - Calculate exhaustion using peak daily values

i If you want to use baseline thresholds, click **Use Dynamic Baseline Thresholds**. See [Orion baseline data calculation](#).

6. Click **Submit**.

You have changed the method and thresholds for calculating capacity forecast for the node.

Customize capacity forecasting settings for interfaces

1. Log in to the SolarWinds Platform Web Console as an administrator.
2. Open the **Edit Properties** page for the interface.

Go to **Settings > Manage Nodes**. Expand the parent node, select the interface, and click **Edit Properties**.

i If you are in an interface capacity forecasting resource, click **Edit**, and click the link to the interface's **Edit Properties** page.

3. On the **Edit Properties** page, scroll down to **Alerting Thresholds**.
4. Select **Override Orion General Thresholds for Receive Interface Utilization or Transmit Interface Utilization**, and customize the Warning and Critical threshold levels.


i If you want to use baseline thresholds, click **Use Dynamic Baseline Thresholds**. See [Orion baseline data calculation](#).

5. Click **Submit**.

You have changed the thresholds for calculating capacity forecast for the interface.

Customize capacity forecasting settings for volumes

1. Log in to the SolarWinds Platform Web Console as an administrator.
2. Go to **Settings > Manage Nodes**.
3. Select the volume, and click **Edit Properties**.

 To find the volume, locate the node, and click the + sign to display interfaces and volumes on the node.


4. Select Override Orion Capacity Thresholds for Percent Disk Usage.
5. Customize the Warning and Critical threshold levels.
6. Select the appropriate method for calculating trends:
 - Use Average values
 - Use Peak values

7. Click Submit.

You have changed the method and thresholds for calculating capacity forecast for the volume.

Discover your network paths

NetPath™ is a feature that helps you identify network problems faster by automatically creating a map of the problem area, with a wide variety of supporting information. NetPath™ displays the performance details of devices inside and outside of your network.



Key features of NetPath™

- NetPath™ discovers the node-by-node network path.
- NetPath™ quantifies the performance of each link and node along the path.
- NetPath™ isolates the node or connection that is decreasing end-to-end performance.
- If the issue is external, NetPath™ identifies the name of the company that owns the node and displays their contact information.
- If the issue is internal, NetPath™ incorporates data from NPM, NCM, and NTA about your on-premises gear.

How does NetPath™ work?

NetPath™ uses distributed monitoring and path analysis to discover how applications are delivered through the network to your users. To use NetPath™:

1. You deploy agents on Windows computers that act as synthetic users. The agents use advanced probing to discover and test the network path that traffic takes to any network endpoint, such as your local file print server, your website, or external websites.

2. After discovering the path and quantifying the performance of each node and connection, NetPath enriches the picture with additional data about Internet nodes. If you are monitoring non-Internet nodes, NetPath™ incorporates that data too.
3. The result is a clear end-to-end map of how applications are delivered to your users, including your network, the network of your provider, and any other networks you depend on.

NetPath™ answers the following questions:


- How well is my network delivering applications to my users?
- Are the paths to key applications or users down?
- Where is the network problem and who is responsible for it?

NetPath requirements for NPM

Probe computer

Probes are the source of network paths, and the paths are discovered by probes.

You [create a probe](#) on a source computer, which must meet the following requirements:

Type	Requirements
Operating system (64-bit only)	Windows Server 2019
	Windows Server 2016
	Windows Server 2012 R2
	Windows Server 2012
	Windows Server 2008 R2 SP1
	Windows 10 Professional and Enterprise
	Windows 8.1
	Windows 8
	Windows 7
 Windows 10 Home edition is not supported.	
CPU cores	2 CPU cores for 20 paths +1 CPU core per 10 additional paths

Type	Requirements
Hard drive space	1 GB
RAM	2 GB

Ports and other firewall settings

Open the following ports on your firewall for network connectivity used by NetPath™. Review other firewall settings.

Ports

i You may also need to open other ports:

- [NPM ports](#) for communication between polling engines.
- [Agent ports](#) when deploying probes on remote machines using agents.

Port	Pro- to- col	Service or Process	Direc- tion	Sour- ce	Desti- nation	Descrip- tion
User configured	TCP	SolarWinds Agent or JobEngineWorker	Outgoing	NetPath™ probe	Endpoint service	Any ports of the monitored services that are assigned to the probe. Used by the NetPath™ probe to discover service status.
43 443	TCP	SolarWinds.Business- LayerHost (Main server only)	Outgoing	Main polling engine	BGP data providers and announcements, such as: <ul style="list-style-type: none"> • http://whois.arin.net/ui/ • https://stat.ripe.net/ 	Used by NetPath™ to query BGP information about the discovered IP addresses.

Other firewall settings

Setting	Protocol	Service or Process	Direction	Source	Destination	Description
Allow ICMP type 11 (ICMP Time Exceeded)	ICMP	SolarWinds Agent or JobEngineWorker	Incoming	Networking devices along your path	NetPath™ probe	Used by the NetPath™ probe to discover network paths.

Database storage

When calculating the size requirements in SQL Server for NetPath™, you must account for the probing interval and the complexity of the network path from the probe to the monitored service. The complexity of the path is divided into three groups:

- Internal: services with fewer than 10 nodes between the probe and the monitored service.
- Intermediate: multiple paths ending in a single endpoint node. Examples are github.com, linked.com, and visualstudio.com.
- Complex: multiple paths (over 20) ending in multiple endpoint nodes. Examples are google.com and yahoo.com.

This table provides an estimate in megabytes (MB) of the amount of storage consumed by SQL Server over a 30-day period (the default retention time) when monitoring a single service.

Interval (in minutes)	Internal (in MB)	Intermediate (in MB)	Complex (in MB)
1	520	1105	1615
2	325	645	1145
3	200	445	915
4	170	350	750
5	135	265	480
10	80	175	470

Example storage requirement calculation

Your monitoring setup contains the following:

- Five internal monitors with a one-minute interval.
- Three intermediate monitors with a five-minute interval.
- Four complex monitors with a ten-minute interval.

The total storage requirement for SQL Server can be calculated as:

$(5 \times 520) + (3 \times 265) + (4 \times 470) = 5275$ MB over a 30-day time period.

Cloud environment

When you place a probe in a public cloud, consider the following additional requirements:

Provider	Requirements
Amazon	<ul style="list-style-type: none"> • Security group must be enabled on instances that host NetPath™ probes to allow inbound ICMP packets. • Probing services that host on Amazon Web Services (AWS) instances within the same cloud networks may not work.
Azure	<ul style="list-style-type: none"> • Private Internet Protocol (PIP) must be enabled on instances that host NetPath™ probes. • Probing may work within VNET, but may not work if the path crosses the Azure Load Balancer.

Scalability

The scalability of NetPath™ depends on the complexity of the paths you are monitoring, and the interval at which you are monitoring them.

In most network environments:

- You can add up to 100 paths per polling engine.
- You can add up to 1,000 paths per SolarWinds Platform instance.

- You can add 10 - 20 paths per probe.

NetPath™ calculates the recommended path count based on the performance of each probe, and displays it each time you deploy a new path to the probe.

Create New Service

Service Details — Assign Probe

Assign Probe

Probe

(8 of 20 paths assigned) ▼

CREATE NEW PROBE

< BACK CREATE CANCEL

Create a NetPath service in NPM

A service is the destination to which you are mapping. It represents an application, and SolarWinds recommends deploying a service for the most important applications that your users rely on. This can be any TCP-based network service, such as salesforce.com, Microsoft Exchange, Office365, or a file server.

NetPath™ services are monitored by probes. Orion automatically installs a probe on each polling engine, and you can install a probe on any Windows computer. No other software is required on the path.

Create a new service

1. Click My Dashboards > Network > NetPath Services.
2. Click Create New Service.

3. Enter the service details of the target destination of your network path. The service must be TCP-based.
 - a. Enter a host name or IP address and port.

i SolarWinds recommends using the same information that your users access the application by. For example, if they access your internal site by a host name rather than an IP address, enter the host name in NetPath™. That way NetPath™ gets the same service as your users.

- b. Enter the probing interval in minutes.

i SolarWinds recommends starting with a 10-minute interval. See the Probing interval section below to learn how to adjust the probing interval.

- c. Click Next.
4. Select an existing probe from the list, or [Create a NetPath probe in NPM](#) to use a new source.
5. Click Create.

Probing interval

This value determines how often and how long information is polled from the network path. If the value is too low, NetPath™ does not complete the probe and the network path may not show all routes. If the value is too high, the information may not update as frequently as you like.

- If you probe more frequently, the data updates quicker but accuracy is lost. If this happens, NetPath™ identifies it as an issue on the probe displayed in the graph.
- If you probe less frequently, the data updates more slowly but the accuracy of the data increases.

SolarWinds recommends starting with a probing interval of 10 minutes, which is appropriate for most paths. You can adjust the value from there to suit your needs.


Is your network path internal? Does it contain fewer than 10 nodes? If so, you can decrease the interval for more frequent data updates.

Is your network path external and does it contain internet connections? Does it contain more than 10 nodes? If so, you can increase the interval for less load strain on the SolarWinds Platform server, your nodes, and the network. A larger value also saves storage space by writing less NetPath™ data to the database.

Create a NetPath probe in NPM


NetPath™ services are monitored by probes. Orion automatically installs a probe on each polling engine, and you can install a probe on any Windows computer. No other software is required on the path.

A probe is the source you are testing from. It is always the start of the path. Think of a probe as a representative of a user. SolarWinds recommends deploying probes where you have users, for example at each of your office locations.

 The probe must be a Windows computer.

Create a probe

You can create a probe when you [create a service](#), or while assigning an additional probe after you create the service:

1. Click My Dashboards > Network > NetPath Services.
2. Click  next to an entry in the NetPath Services list.
3. Click Create New Probe.
4. Enter the required information on the Create New Probe window.

 Enter the credentials that can be used to log in to the computer and install the software.

5. Click Create.
6. Select the probe from the list.
7. Click Assign.

Assign additional probes

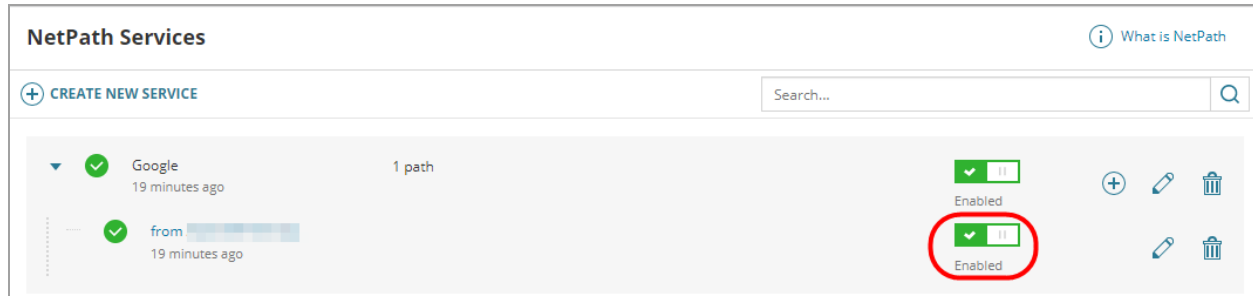
Click  next to an entry in the NetPath™ Services list to assign another probe to the service.

Disable a probe

1. Go to NetPath Services. Click My Dashboards > Network > NetPath Services.
2. Expand the service, and click the switch to disable the probe.

If there is no probe enabled for a service, the service is disabled, too. No data for the service is

polled.



Delete a probe

1. Go to NetPath Services. Click My Dashboards > Network > NetPath Services.
2. Expand the service, and click the waste basket icon to delete the probe.

Deleting the last probe for a service permanently deletes the service.

Troubleshoot probes

If you are creating a probe on an existing SolarWinds Platform Agent, you must enter the primary polling IP address used by Orion for that device.

Check the probe status

If you have other issues with probe deployment, you can check the probe status.

Probes are listed in the Manage Agents section of Agent Management. The NetPath™ probe relies on the Agent infrastructure built into Orion and used for things like QoE and SAM Agents. NetPath™ is an additional plugin in this agent framework.

1. Click Settings > All Settings.
2. Under Node & Group Management, click Manage Agents.
3. Locate the probe in the Agent/Node list by its host name, and select it.
4. Verify the Agent Status is Running, and that the Connection Status is Connected.
5. Click More Actions > View installed agent plugins.
6. Verify the NetPath™ Agent Plugin is installed.


You can also click Edit Settings to change the configuration of the probe, or Delete to remove it.

View a network path in NPM

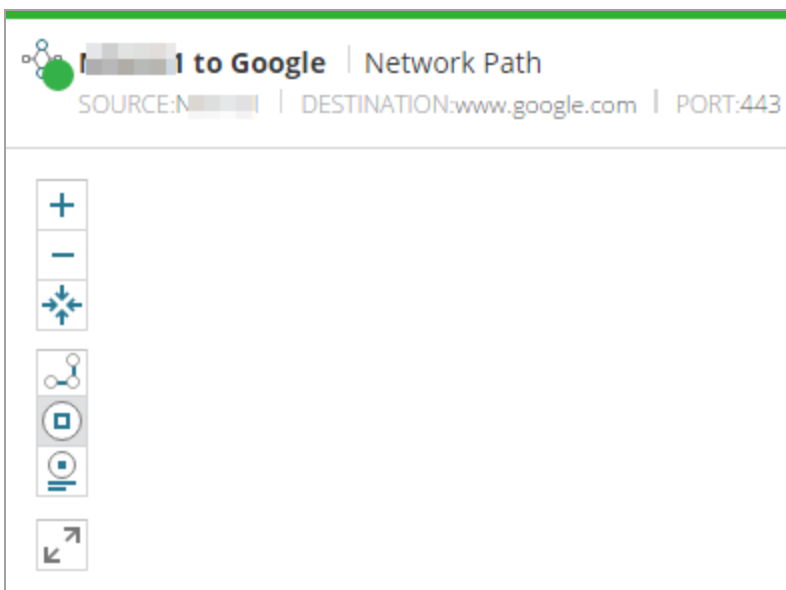
1. Click My Dashboards > Network > NetPath Services. This view displays a list of created network services.
2. Expand a service, and click one of the associated probes to see the network path from that probe to the expanded service.

Path layout

The source is on the left and the destination is on the right. The network path is everything in the middle.

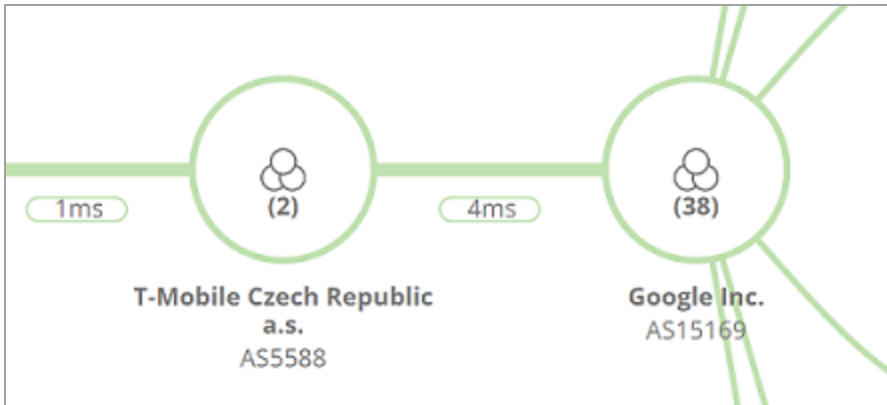


Use the controls in the upper left to change the zoom, detail levels of the path, and the amount of information displayed. You can also use your mouse to pan and zoom.



Objects in the network path include nodes, connections, and interfaces. Point to an object for a summary, and click it for details.

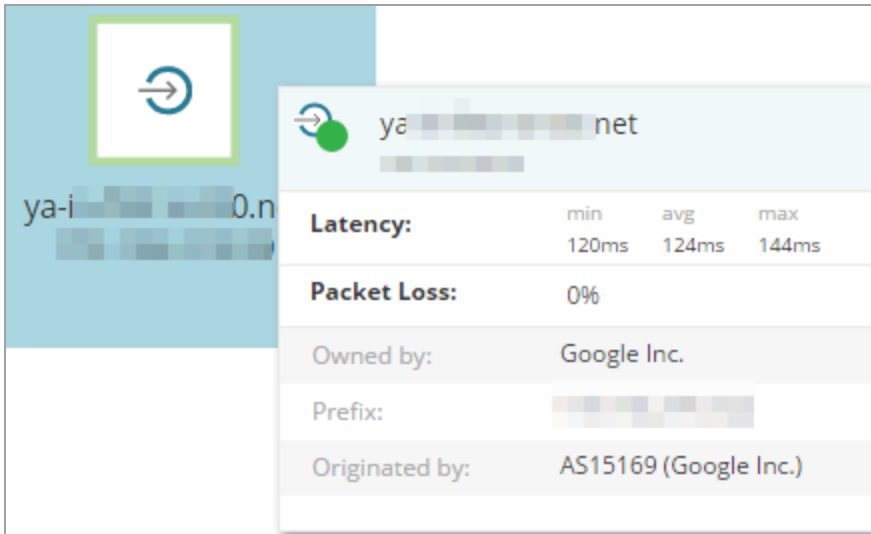
NetPath™ groups nodes into networks represented as larger circles. In the example below, the path goes through two (2) nodes in T-Mobile's network and 38 (38) nodes in Google's network.



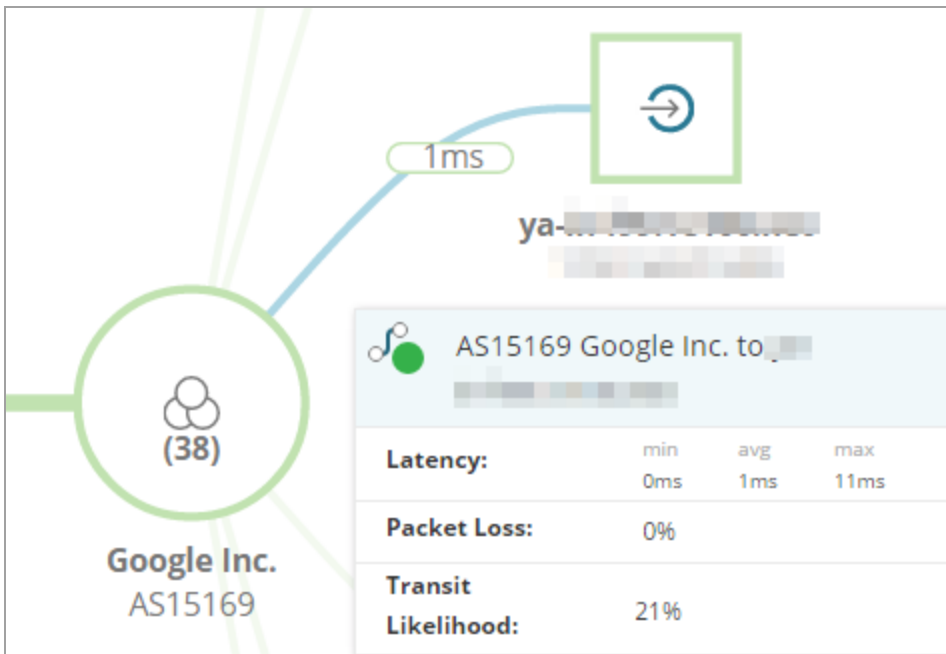
Click the network to show the nodes that comprise it, and click the X on the Expanded filter to collapse it.

The screenshot shows the 'NPM-01 to Google | Network Path' interface. At the top, it displays 'SOURCE:NPM-01 | DESTINATION:www.google.com | PORT:443'. Below this, there is a vertical stack of three icons: a plus sign (+), a minus sign (-), and a starburst icon. To the right of these icons, the text 'EXPANDED:' is followed by a button that says 'Google Inc. (40)' with a close (X) icon on the right side of the button.


The node information is cumulative from the source to that node. When you point to or click a node, the displayed metrics answer the question, “what is the performance between the source, along the path, up to this node?”



A connection between nodes shows latency and packet loss between its two nodes. When you point to or click a link, the displayed metrics answer the question, “what is the performance of this specific link?”



A dotted line illustrates a broken connection to a host that is unreachable. This means that traffic reached the green node, is destined for the endpoint connected with the dashed line, but does not make it.




Use the green, yellow, and red color coding to identify the nodes and connections that may be performing poorly and affecting the end-to-end connection. If you confirm that a service provider is responsible for the outage, you can contact them to resolve the issue.

Path history

The chart on the bottom shows metrics for the end-to-end performance. Select an interval to see the network path and its performance that resulted in that end-to-end performance.

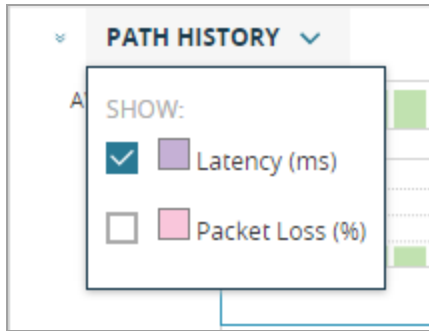
Think of this as your network time machine. You can compare performance metrics from today or a previous time.



Available actions in the path history

- Click a bar in the chart to load the network path from that date and time.
- Click the single arrows, or press the Left and Right Arrow keys, to move one interval at a time.
- Click the double arrows to move to the beginning or the end of the displayed history window.
- Drag the bottom slider to change the history window.


- Click Path History to show or hide Latency and Packet Loss in the chart.



Troubleshoot a NetPath service with external path data in NPM

You can use NetPath™ to diagnose a slow connection to an external service. This example uses amazon.com.

1. Click My Dashboards > Network > NetPath Services.
2. Expand the service that your users reported as slow or unreachable.
3. Click the probe from the office or location that reported the issue.
4. Under Path History, locate the date and time for when your users reported the issue. Here, there is a yellow warning entry at 5:09 p.m. on April 20.

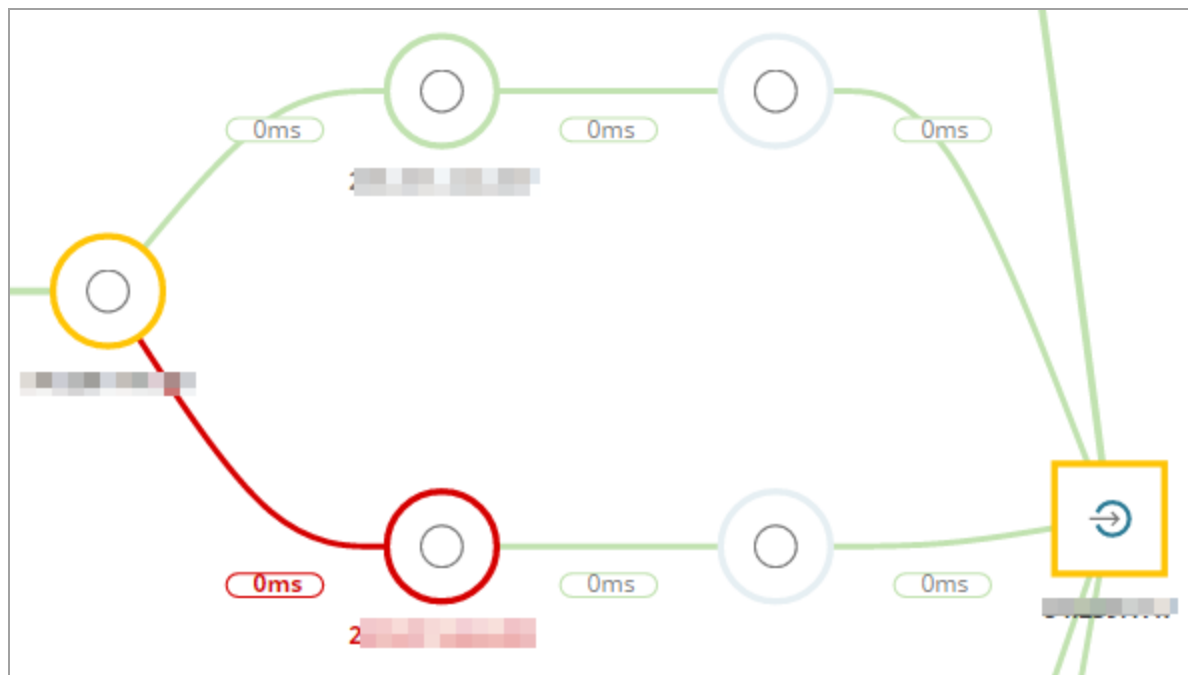


5. Click the yellow bar at 5:09 p.m. in the chart.

6. The problem is in Amazon's network. Click the red Amazon node to expand it.

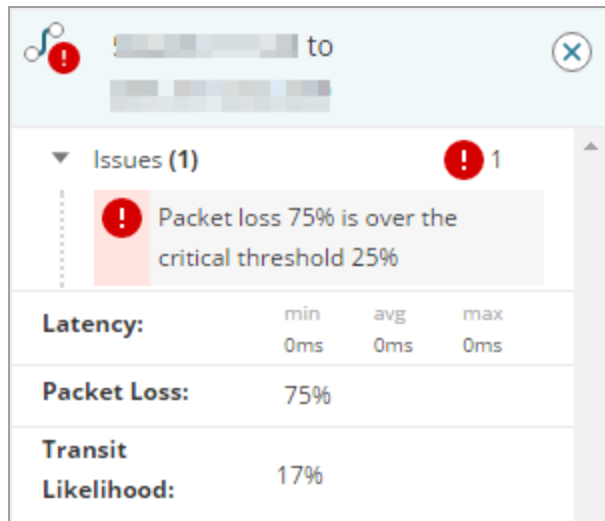


7. Although Amazon's network is large and complex, you should investigate the red and yellow areas.



8. Click the red connection between the two nodes to open the inspector panel.

- Expand the Issues section to see that packet loss is over the critical threshold, and that it is 17% likely that transit passed through this link.




- Click the red 205.251.244.209 node to open the inspector panel.
- Use the phone number or email address to contact the service provider and report the issue. Present the following information to resolve the issue:
 - IP addresses of the nodes in question (54.239.111.33 and 205.251.244.209 in this case)
 - Date, time, and duration of the performance issue
 - Latency and packet loss information

Troubleshoot my network with NetPath data in NPM

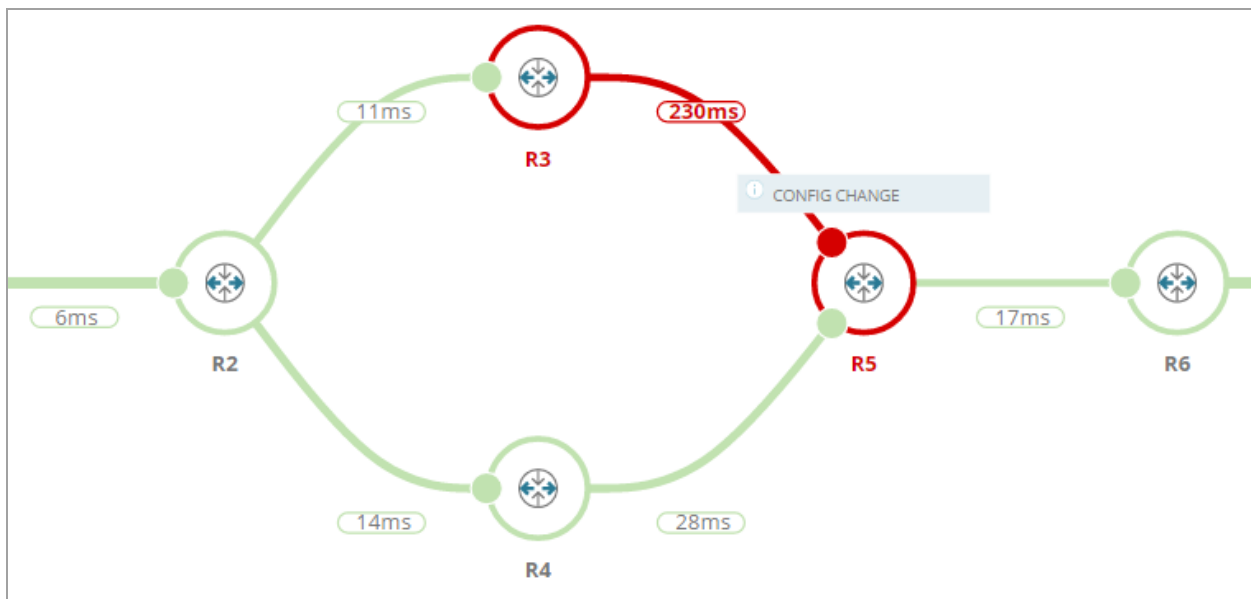
You can use NetPath™ to diagnose a slow connection caused by your internal network. This example shows a node that stopped working properly after a change to its config file.

- Click My Dashboards > Network > NetPath Services.
- Expand the service that your users reported as slow or unreachable.
- Click the probe from the office or location that reported the issue.

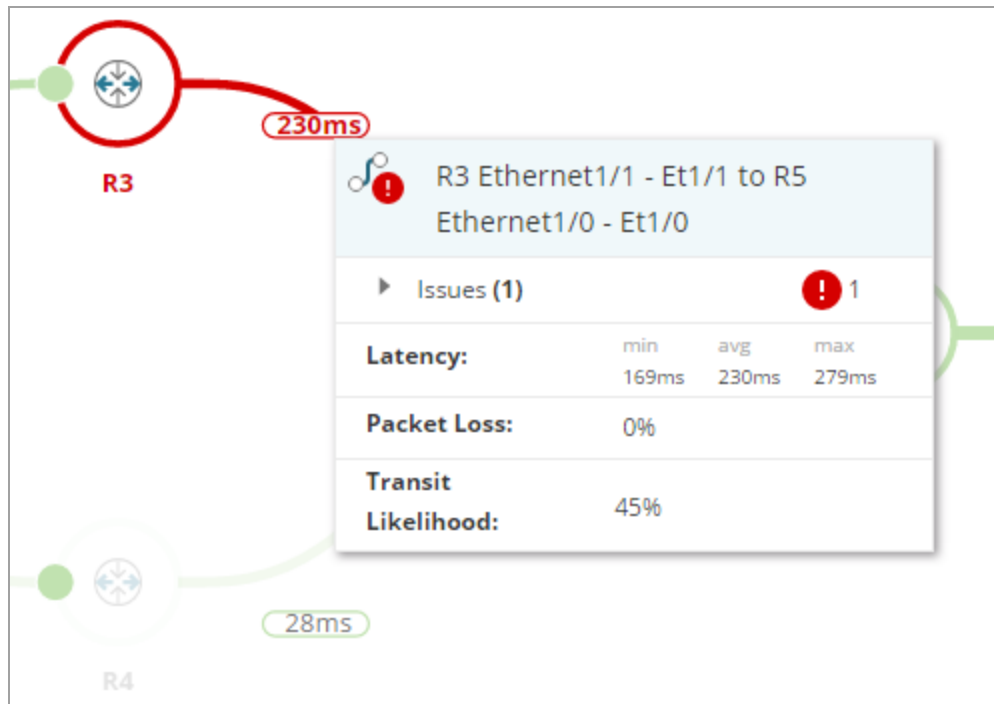
- Under Path History, locate the date and time for when your users reported the issue. Here, there is a red critical entry at 3:26 p.m. on April 14.



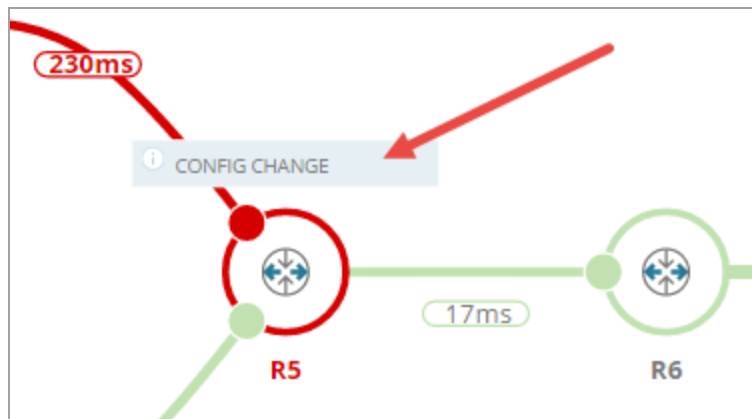
- Click the red bar at 3:26 p.m. in the chart.
- The problem is in the internal network. There is a high latency between nodes R3 and R5.



- Point to the red connection between the two nodes to see that Transit Likelihood is 45%. This means that just under half of your users are likely to experience the problem.



- NCM is installed, so the [SolarWinds Platform integration with NetPath](#) displays information about a config change to node R5. Click the Config Change notification.



9. In the config comparison window, scroll down until you see the highlighted change.

72	!	!
73	interface Ethernet1/0	interface Ethernet1/0
74	ip address 1	ip address
75	ip flow ingress	ip flow ingress
76	ip flow egress	ip flow egress
77	duplex half	duplex half
78		traffic-shape rate 1000000 25000 25000 1000
79	!	!

10. A new command was added on line 78 for `interface Ethernet1/0`. This is the problem. Note the change, and close the config comparison window.
11. Use NCM to revert the config file, or log in to the device and remove the incorrect configuration.

SolarWinds Platform integration with NetPath

NPM integration


NetPath™ is a feature of NPM, and by default displays NPM data and issues.

On the internal portion of the network path, you can:

- See NPM data such as CPU, RAM, interface utilization, and more included in the graph.
- Click a monitored device and go to its Node Details page.
- Click an unmonitored device and add it to Orion to see more data.

NTA integration

NetPath™ uses data from NPM to display information about your internal nodes on the network path, such as bandwidth used for the interface. But what is using that bandwidth?

 NetPath™ and NTA integration requires NTA 4.2 or later.


If you are exporting flow data from those nodes and monitoring it with NTA, NetPath™ displays additional information to identify what is using the most ingress and egress bandwidth.

Click the node or interface in the network path to open the inspector panel, where you can:

- View the top three conversations.
- Select ingress or egress.
- Click a conversation name to view details about that conversation.

NCM integration

NetPath™ displays additional information about NCM nodes with backed-up config files. If traffic through an NCM node was affected after a config change, NetPath™ notifies you that the two events may be correlated.

 NetPath™ and NCM integration requires NCM.

NetPath™ highlights config-related issues on the path, and provides quick access to the configuration data for nodes on the path.

Click the node in the network path to open the inspector panel, where you can:


- Click Commands > View Current to see the config for the device.
- Click Commands > Compare to see two configs side by side for comparison.

Monitor Cisco ACI devices in NPM

Thanks to the broad coverage of Cisco OIDs, you can poll many statistics for hardware components that make up ACI, such as Nexus leaf and spine switches.


You can enable API polling on ACI devices to monitor the following components of your SDN environment:

- Tenants
- Application profiles
- Endpoint groups
- Spine and leaf switches

 To get the best coverage of your ACI environment, enable the API polling on one of your APICs and add leaf and spine switches to NPM as SNMP nodes. To find out how to configure SNMP in APICs, see the Cisco document [Configuring SNMP in APIC](https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/mib/guide/b_Cisco_ACI_MIB_Quick_Reference/b_Cisco_ACI_MIB_Quick_Reference_chapter_01.html) (© 2020 Cisco and/or its affiliates, obtained from https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/mib/guide/b_Cisco_ACI_MIB_Quick_Reference/b_Cisco_ACI_MIB_Quick_Reference_chapter_01.html on February 26, 2020).

To monitor ACI-specific information, complete the following steps:

- [Add an APIC node to NPM](#) for monitoring, or [enable ACI monitoring on an APIC node that is already monitored](#) with NPM.

 To add more ACI devices at the same time, see [Discover your network with the Discovery Wizard](#), and then [enable polling for Cisco ACI](#) on one of the APICs.

- [View members and their health scores on the device.](#)
- [View health score history in PerfStack](#)
- [View ACI environment on Intelligent Maps](#)

Requirements

Requirement	Details
ACI credentials	Cisco API Rest credentials for collecting health scores on ACI entities. REST API must be accessible from the main or additional polling engine server (depending on the polling engine used to poll the node)

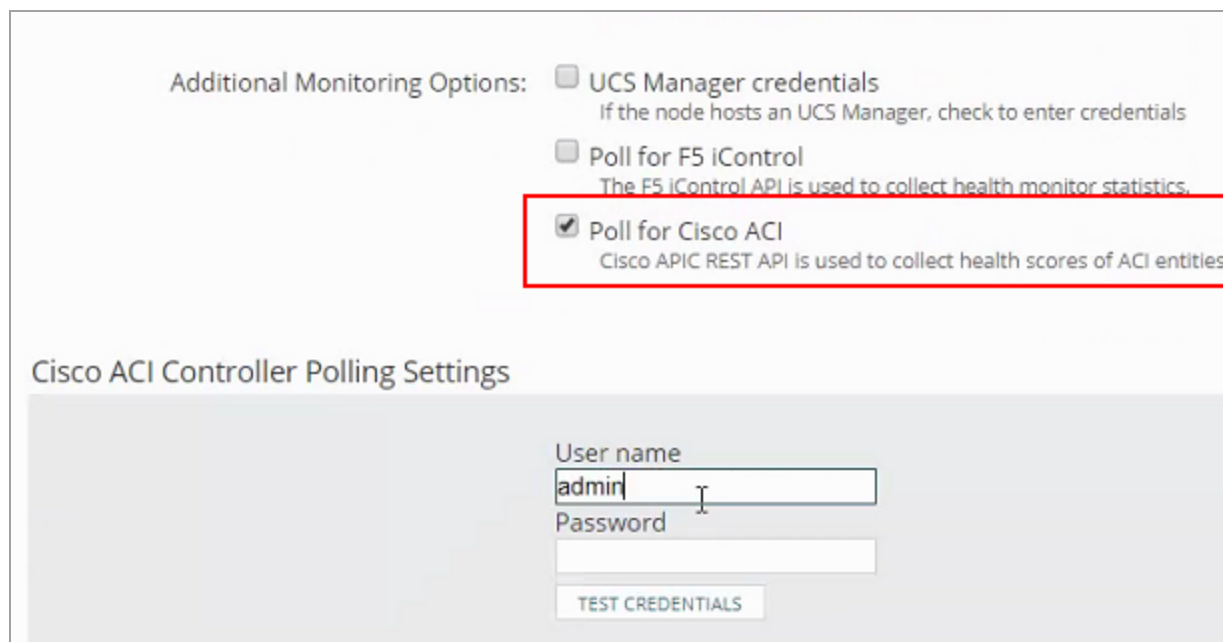
Add ACI devices and enable ACI polling

[Add ACI devices](#) for monitoring. If you have multiple APIC nodes in your ACI system, you can monitor all APIC nodes with NPM to collect health and performance data.

To collect health scores for your ACI environment, enable ACI polling on one of the APIC nodes. Each APIC has a full view of the ACI environment, and enabling ACI polling on multiple nodes thus results in polling and storing redundant information.

i You need Node Management Rights. See [Define what users can access and do](#).

1. Select an APIC to poll for ACI details.
2. Click Settings > All Settings, and click Add Node in the Getting Started grouping.
3. Enter the IP address for the device.
4. Select Most Devices: SNMP and ICMP as the polling method.
5. Scroll down to Additional Monitoring Options, and select Poll for Cisco ACI.
6. Provide the credentials for accessing the Cisco APIC REST API on the device, and click Test Credentials.



Additional Monitoring Options:

- UCS Manager credentials
If the node hosts an UCS Manager, check to enter credentials
- Poll for F5 iControl
The F5 iControl API is used to collect health monitor statistics.
- Poll for Cisco ACI
Cisco APIC REST API is used to collect health scores of ACI entities

Cisco ACI Controller Polling Settings

User name
admin

Password

TEST CREDENTIALS

7. Complete the Add Node wizard.

NPM now polls health scores for devices linked to the APIC, and calculates the ACI status based on the polled health score.

Enable polling for Cisco ACI on a monitored node

When the Cisco ACI devices are already monitored in NPM, make sure polling for ACI is enabled only on one of the APICs. Consider enabling ACI polling on the node with the least load or with the shortest response time.

Polling for Cisco ACI is used to collect health scores for ACI entities.

1. Click Settings > Manage Nodes.
2. Select the node, and click Edit Properties.
3. [Enable Cisco ACI polling](#).
4. Click Submit.

Now you can see the Members subview with health score information for tenants and blades on the monitored device.

View health scores for ACI members


1. In the SolarWinds Platform Web Console, go to the ACI node details view.
2. Click the Members subview.


The Members subview displays the following items:

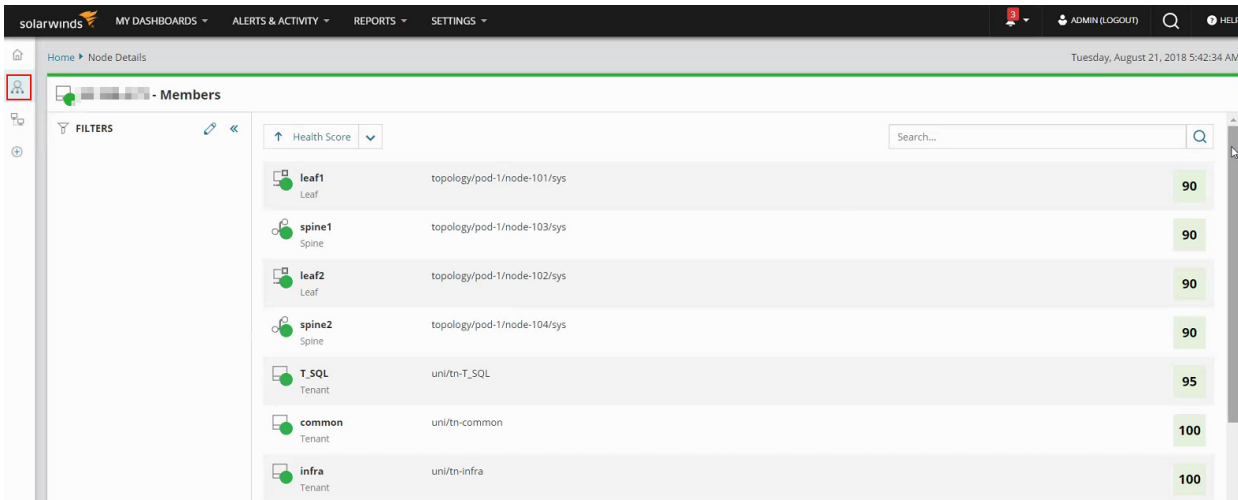
- Name
- Distinguished name
- Member type together with the type-specific icon
- Health score

ACI status uses thresholds defined on the APIC. The following table lists the default thresholds.

0-50	critical status
51-90	warning status
91-100	up

 You can automatically display APIC members on the node on [Intelligent Maps](#). Click the Map icon below in the navigation bar below the Members subview.

 Use a part of the distinguished name to see child entities on a certain level (tenants or application profiles).



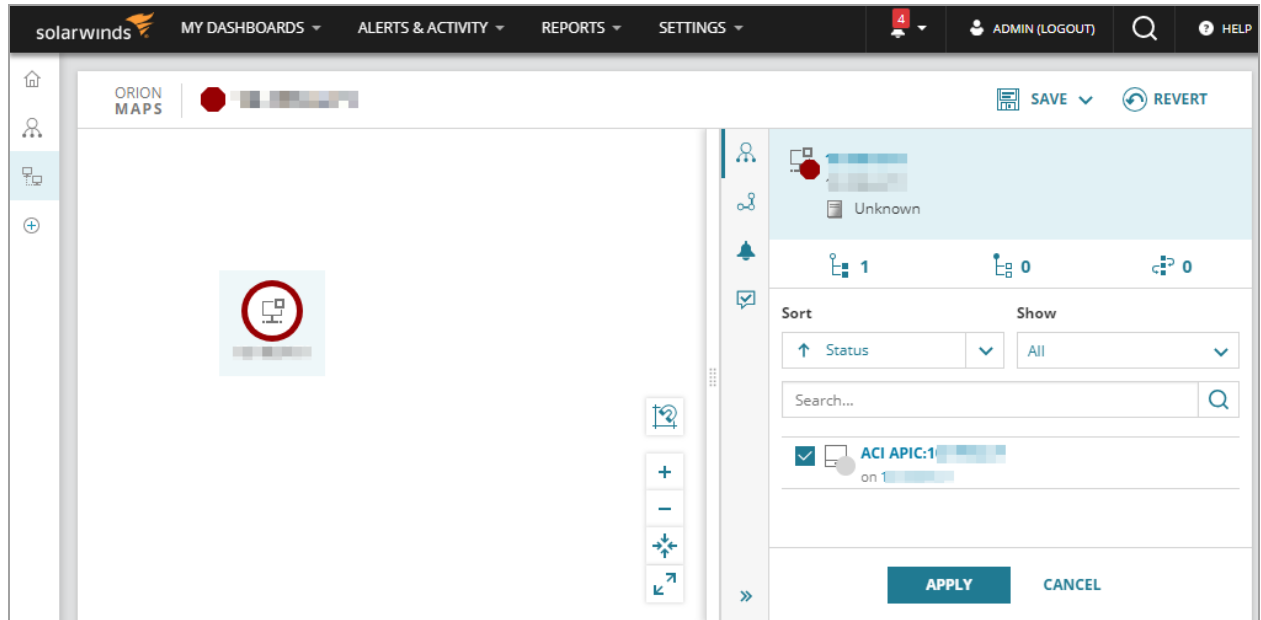
View health score and status history in PerfStack

Click a health score to open this metric as a new [Performance Analysis Dashboard](#). In the PerfStack, you can see the health score and status of the ACI entity, including historical data.

View your SDN infrastructure on Intelligent Maps

1. Go to the Node details for APIC view, and click the Map subview. The map shows only the node.

- In the Inspector panel on the left, select the ACI APIC for the node and click Apply to add it to the map.

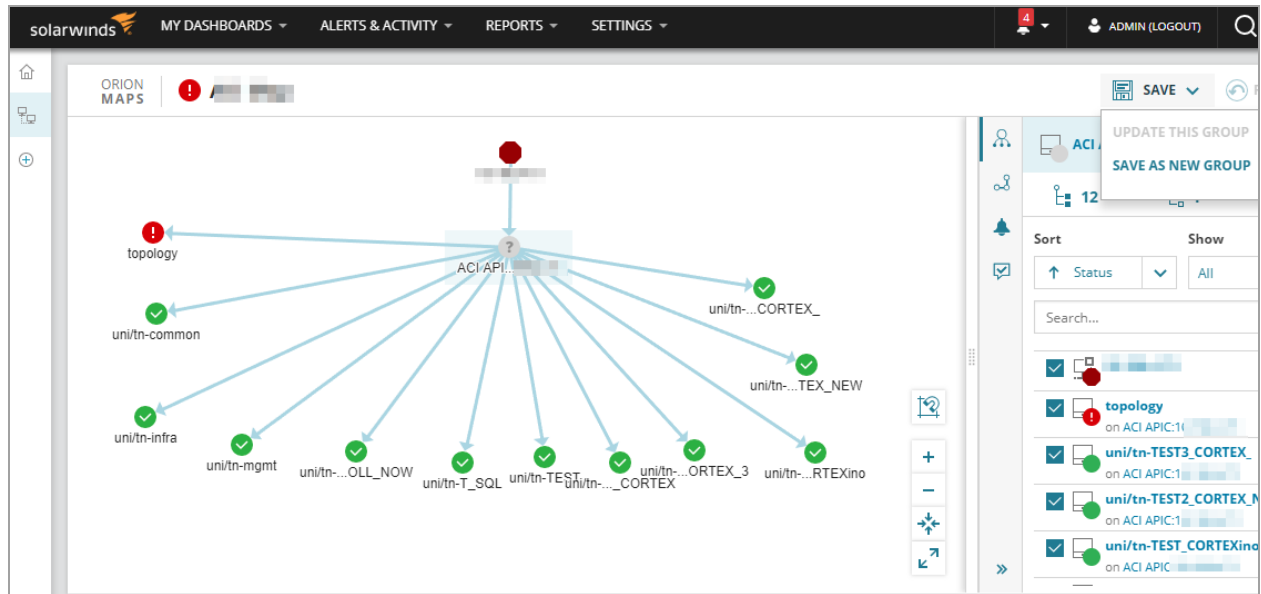


The APIC is added to the map. It represents the gate to the logical layer on your ACI.

- Select the ACI APIC on the map. The Inspector panel populates with all members on the APIC.

4. Select the members in the Inspector panel, and click Apply to add them to the map. The members appear on the map.

i You can now select entities or connections to display more details in the Inspector panel. Use the buttons on the map to apply different layouts, zoom in/out, center the map or extend it to the full screen. See [Intelligent Maps](#) for more details.



5. To keep the map, save the objects as a group.
 - a. Click Save > Save As New Group.
 - b. Provide a name for the group.

The group is created and the Group Details view opens.

💡 You can access the map at any time:

- a. Click My Dashboards > Groups, and then click the group name in the All Groups widget.
- b. When on the Group Details view, click the Map subview.

Create ACI-specific alerts and reports

There are no out-of-the-box alerts and reports specific for ACI. You can [configure custom notifications](#) based on ACI events and [custom reports](#) showing ACI-relevant statistics.

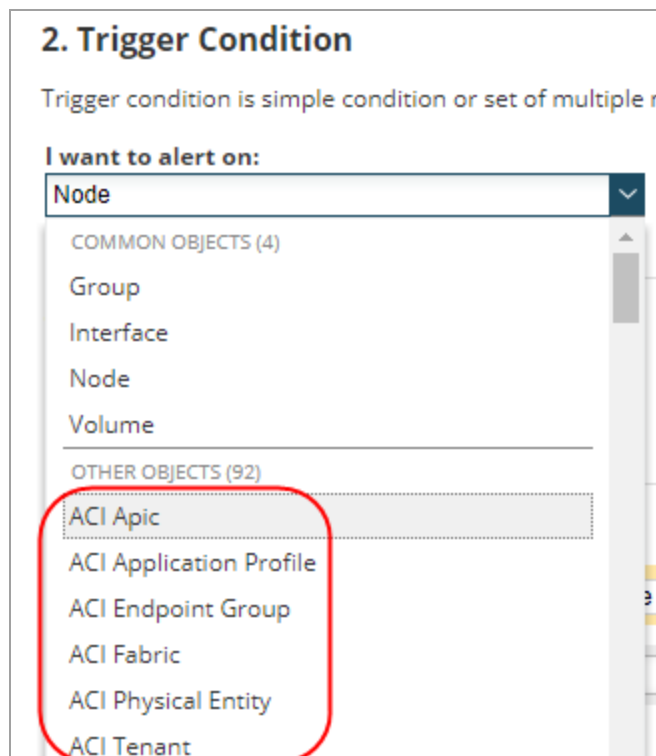
Review the following customization examples:

- [Create an alert to be notified when the health score on an ACI tenant is less than 90%](#)
- [Copy and adjust an alert to be notified when there are polling issues on an APIC](#)
- [Create a report listing ACI tenants with health scores less than or equal to 90%](#)

Alert on health score values lower than 90% on an ACI tenant

Create a new alert to notify you when the health score on an ACI tenant drops below 90%.

1. Select Alerts & Activity > Alerts, and click Manage Alerts.
2. Click Add New Alert.
3. On Trigger Condition, select the ACI entity you want to be alerted about. In this case, select ACI Tenant.



4. In the actual trigger condition, ensure that ACI Tenant is selected in the first drop-down.

- In the neighboring boxes, select Health Score, "is less than or equal to" and enter the percentage when you want to be alerted.

Add New Alert - "ACI Tenant Health Score less than 90"

PROPERTIES
TRIGGER CONDITION
RESET CONDITION
TIME OF DAY
TRIGGER ACTIONS
RESET ACTIONS
SUMMARY

2. Trigger Condition

Trigger condition is simple condition or set of multiple nested conditions which must be met before the alert is triggered. [»Learn more](#)

I want to alert on:

ACI Tenant
▼

The scope of alert: ⓘ

All objects in my environment ([Show List](#))

 Only following set of objects

The actual trigger condition:

Trigger alert when
All child conditions must be satisfied (AND)
▼

⋮

ACI Tenant
▼

Health Score
▼

is less than or equal to
▼

90
%

🗑️

- [Specify the trigger action](#) and complete the wizard.

After the health score on the selected entity falls below the specified percentage, you will be alerted using the method you specified.

Alert on ACI polling issues on an APIC

In case of polling issues, the status of monitored ACI APICs changes to unknown. The following steps describe how you can [duplicate and edit an out-of-the box alert](#) to be alerted if an APIC status changes to unknown.

- Select Alerts & Activity > Alerts, and click Manage Alerts.
- Find the Node is down alert, select it and click Duplicate and Edit.
- On Properties, adjust the alert name.
- On Trigger Condition, change the entity in Trigger condition to ACI APIC.

5. Change the trigger condition to say Alert me when the ACI APIC status is equal to Unknown, and click Next.

2. Trigger Condition

Trigger condition is simple condition or set of multiple nested conditions which must be met before the alert is triggered. [*Learn more](#)

I want to alert on:

The scope of alert: ⓘ

All objects in my environment ([Show List](#))
 Only following set of objects

The actual trigger condition:

Trigger alert when

⋮	<input type="text" value="ACI Apic"/>	⌵	<input type="text" value="Status"/>	⌵	<input type="text" value="is equal to"/>	⌵	<input type="text" value="Unknown"/>
---	---------------------------------------	---	-------------------------------------	---	--	---	--------------------------------------

6. [Adjust the trigger action](#) and complete the wizard.

When an ACI APIC is in Unknown status, you will be notified by the configured trigger action.

Create a report to list all ACI Tenants with health score 90% and less

1. In the SolarWinds Platform Web Console, click Reports >Manage Reports.
2. Click Create new report.
3. Select Custom Table.

- In Select objects to report on, select ACI Tenant, and add a condition that says Health Score is less or equal than 90.

Add Content

1. Resource: Custom Table (Reporting)

2. Select objects you are going to report on...

Selection method:

Dynamic Query Builder Use DQB for selections such as 'All cisco n

Basic Selector Advanced Selector

I want to report on ACI Tenant

Where All child conditions must be satisfied (AND)

- Health Score is less than or equal to 90

+ Add Simple Condition

Selection Name: ACI Tenant with health scores less than 90

- Click Add Column, select the columns for the report, and click Add Columns.

Add Column

Available columns: Search in columns ...

ORION OBJECT: ACI Tenant

GROUP BY: [No Grouping]

ACI Tenant

- ACI Tenant Statistics
- Orion Site

- Database column name
- Description
- Details Url
- Display Name
- Distinguished Name
- Health Score
- Name
- Status

- Back on Edit Resource, provide the title for the report and click Submit.

Edit Resource: Custom Table for Datasource 1 i You can change this after

Title:

Subtitle:

Table layout: [Edit column widths](#)

NAME 🗑	HEALTH SCORE 🗑	STATUS 🗑	DISPLAY NAME 🗑
▶ Advanced	▶ Advanced	▶ Advanced	▶ Advanced

SUBMIT
PREVIEW RESOURCE


You have created a customized report. When you click Reports and click the report name, it will display a list of ACI tenants with health score values that are less or equal to 90%.

Monitor ASA firewalls with NPM

Network Insight for Cisco® ASA automates the monitoring and management of your Cisco ASA infrastructure to provide visibility and help ensure service availability.

Ensure that services dependent on your firewall are available:

- [Monitor VPN tunnels](#): to guarantee the connectivity between sites. Monitor the tunnel status, bandwidth usage, and information about completed phases. View user sessions on remote access tunnels.
- [Monitor firewall high availability health and readiness](#): detect failovers, and keep track of ASA high availability status.

 To have the complete visibility into the health and performance of your firewall infrastructure, and to automate operational activities, such as optimizing your Access lists (ACL), install Network Configuration Manager.

Out-of-the-box alerts

- Failover on ASA node
- High Availability on ASA Node is not up
- VPN Site-to-Site tunnel down
- Connections in use exceeding threshold on ASA node

Out-of-the-box reports

- VPN Site-to-Site Tunnel History - Last 30 Days
- VPN Remote Access Tunnel History - Last 30 Days

Next steps


- [Set up monitoring Cisco ASA firewalls in NPM](#)

Set up monitoring Cisco ASA firewalls in NPM

Data for monitoring Cisco® ASA firewalls is polled by a combination of SNMP and CLI polling. To get accurate ASA-specific information, add the firewall device to NPM as a node, and provide CLI credentials.

What does CLI polling provide?


Enable CLI polling to receive additional ASA-specific details, and to display accurate information for your Cisco ASA devices.

 For example, when polling Site-to-Site VPN tunnels, CLI polling helps filter data polled through SNMP, and display only relevant results. Without CLI polling, you might see failed access attempts from outside as failed tunnels.

Information polled by CLI

- Security level and standby IP address for interfaces
- Number of failed connections per minute on the ASA
- High availability details:
 - Configuration sync state
 - Connection sync state
 - Standby state
 - High availability mode
 - Last failover date and time
 - System HA type and system HA role
 - Peer interfaces
- Firewall mode, serial number, and contexts on the ASA device
- All configured Site-to-Site tunnels on the ASA, including inactive tunnels
- For Site-to-Site tunnels, local IP address, local host name, remote IP address, and remote host name

Requirements

Requirement	Details
Cisco ASA version	Cisco ASA 8.2 and later <div style="border: 1px solid #ccc; padding: 5px; background-color: #e0f2f1;"> Cisco ASA Services Modules are not supported.</div>
ASA user account	Credentials for logging into the ASA device
Enable password	Credentials for polling CLI details. Without this password, you can access the ASA, but cannot poll it.
SSH port	By default, port 22. Open an SSH port for accessing and polling ASA devices through SSH.

Add ASA firewalls using CLI credentials

i You need Node Management Rights. See [Define what users can access and do](#).

1. Click Settings > Manage Nodes, and then click Add Node.
2. Enter the IP address for the device.
3. Select Most Devices: SNMP and ICMP as the polling method, and enter SNMP credentials.
4. Choose resources and add pollers if necessary. You can keep the default settings.
5. On the Change Properties screen, enable CLI monitoring:
 - a. Scroll down to CLI Polling Settings.
 - b. Select Enable CLI Polling, [enter the credentials](#), and click Test.

i Enter a user name and password for logging into the ASA or Nexus device.

If you have configured a security password for CLI polling on the device, provide it in Enable password. Without the Enable Password, CLI polling does not work.

CLI Polling Settings: Enable CLI Polling » [Learn More](#)
Enable polling for Cisco Nexus, Cisco ASA or Palo Alto

Username:

Password:

Enable Password:

SSH Port:

Use Keyboard Interactive Authentication:

6. To use a specific device template, select it. [Device templates](#) are sets of commands you can execute on a device. See [NCM Getting Started Guide](#) for more information.
7. Complete the Add Node wizard.

You can now view the polled ASA firewall information in NPM.

Enable CLI polling on monitored devices

To poll firewall-specific data on ASA devices already monitored in NPM, enable CLI polling for ASAs.

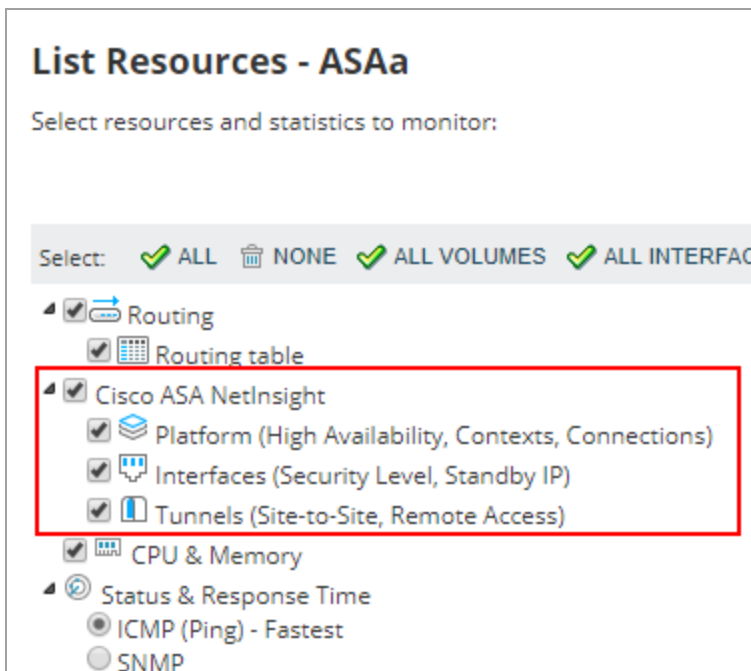
1. On the node details view, click Edit Node in the Management widget.
2. Scroll down to the CLI Polling Settings section.
3. Select Enable CLI Polling, [enter the credentials](#), and click Test.
4. Click Submit.

You can now poll firewall-specific information, such as platform, interface and tunnel details. Pollers for ASA Network Insight are enabled.

Troubleshoot CLI polling

If CLI polling has issues, verify the following:

- You have enabled CLI polling on the ASA device.
- You are using the correct user credentials to log in to the ASA device.
- You are using the correct password for CLI access.
- You have enabled Cisco ASA NetInsight pollers:
 - a. On the node details view, click List Resources in the Management resource.
 - b. Expand Cisco ASA NetInsight, and select to enable the pollers.



- Enable the CLI session trace to extend logging:
 - a. Click Settings > All Settings > CLI Settings in the Product Specific grouping.
 - b. Click Enable Session Tracing, and click Submit.

Review the session trace files located at:

```
%ALLUSERSPROFILE%\Application Data\SolarWinds\Logs\Orion\CLI\Session-Trace
```

Understand ASA platform health in NPM

Understand the health of the Cisco® ASA platform, for example power supplies, ASA high availability status, and other platform-wide health attributes.

1. Log in to the SolarWinds Platform Web Console.
2. On the Summary view, locate your ASA firewall node, and click it to go to the Node Details view.
3. Review the Node Details for ASA - Summary subview.

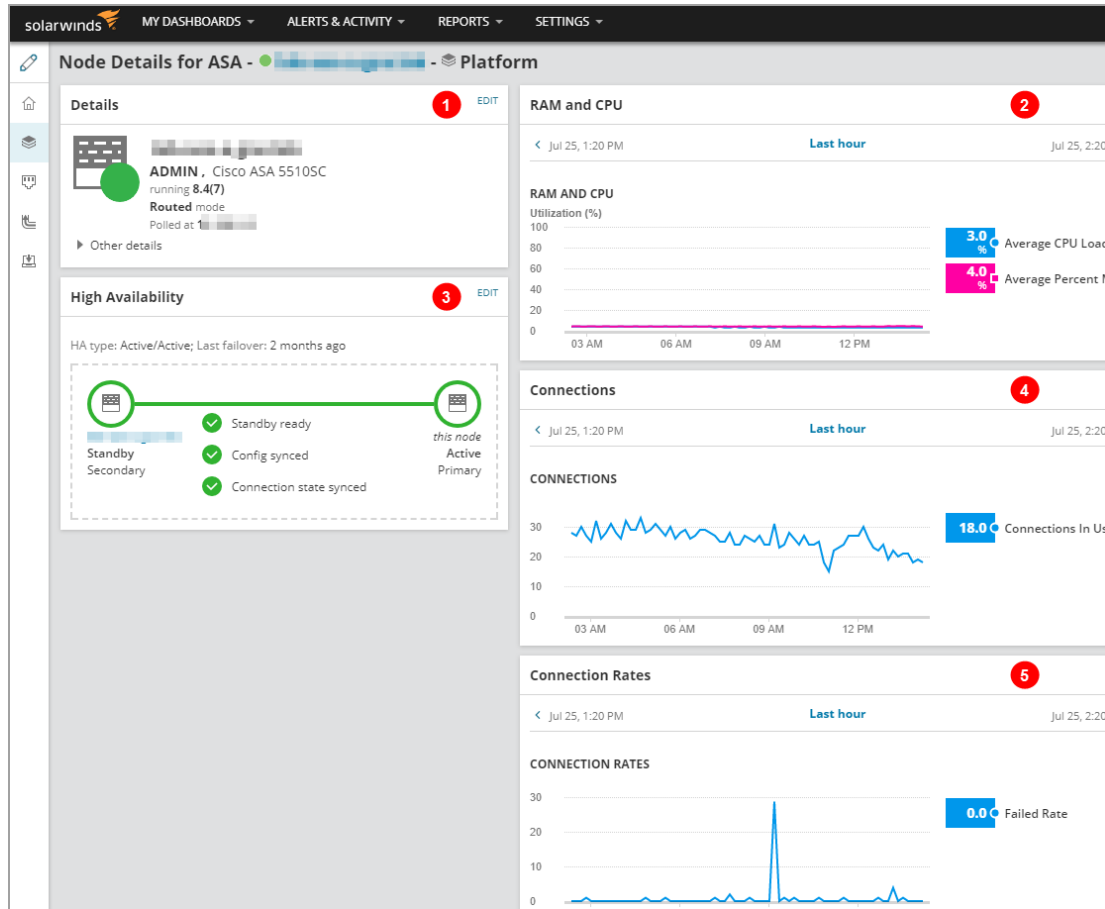
The Summary only displays widgets relevant for the ASA device.

The screenshot shows the 'Node Details for ASA - Summary' page. It features several widgets:

- Details (1):** Shows node information for 'Cisco ASA 5515' running '9.2(4)5'.
- Load Summary (2):** Displays a line chart for 'LOAD SUMMARY' with metrics: Average Percent Memory Us... (18.7%), Average CPU Load (2.0%), and Connections In Use (539.0).
- Management (3):** Lists management actions like 'Edit Node', 'List Resources', 'Pollers', 'Poll Now', 'Rediscover', 'MIB Browser', 'Add New Alert', 'SSH', 'Performance Analyzer', and 'Maintenance Mode'.
- Favorite Site-to-Site VPN (5):** A table showing VPN tunnels with columns for Tunnel, Bandwidth In, Bandwidth Out, Duration, and Other. Two tunnels are marked as 'failed at phase 1'.
- Site-to-Site VPN Health Overview (7):** A pie chart showing the health of VPN tunnels: 1 Up, 2 Down. Total count is 3.
- Platform Summary (4):** Shows 'Hardware Health' as green (7) and 'High Availability' as 'Unknown; See details'.
- Status:** Shows a green bar for 'UP' status.
- Bandwidth (6):** A line chart for 'IN BANDWIDTH' showing 'CAPITA-DMZ : In bps' (14.5 Gbps), 'DMZ1 : In bps' (28.9 Gbps), and 'INSIDE : In bps' (25.4 Gbps).

- 1 Review the node details, such as firmware version, or IP address.
- 2 See the load summary on the device - average percent memory used, average CPU load, and connections in use.
- 3 Click Performance Analyzer to open the Performance Analysis dashboard for the ASA node and view predefined metrics.
- 4 Review the hardware health and high availability status. Click See details to go to the Platform overview, and see more information about High Availability.
Hardware health information is displayed only if it is available on the device.
- 5 See the top 3 Site-to-Site VPN tunnels.
[How do I add tunnels to this resource?](#)
- 6 Review the In and Out bandwidth of favorite interfaces.
[How do I add interfaces here?](#)
- 7 See the basic health overview of monitored Site-to-Site tunnels.

- Click the Platform subview to see more details about the ASA platform health, such as ASA high availability status, RAM and CPU status, connections, and connection rates.



- Review the node details, such as firmware version, or IP address.
- Review the RAM and CPU utilization of the device.
- Review the [node and ASA high availability status](#).
- Review the number of connections in use over a time period.
- Review the number of failed connections over a time period.

What other aspect of the ASA platform are you interested in?


- [Contexts](#)
- [ASA high availability](#)
- [Interfaces](#)

- [Site-to-Site VPN](#)
- [Remote Access VPN](#)
- [Access lists](#)

Monitor contexts

If you have configured contexts on a monitored ASA device, they are listed in the Contexts widget, or resource on the Node Details for ASA - Summary view.

To add a context configured on a monitored ASA device, click the Monitor Node link and [add the context to NPM using CLI credentials](#). NPM provides the same monitoring details as for other ASA nodes.




 Each monitored context requires a node license.

To monitor a context without monitoring the ASA device, [add the context to NPM using CLI credentials](#).

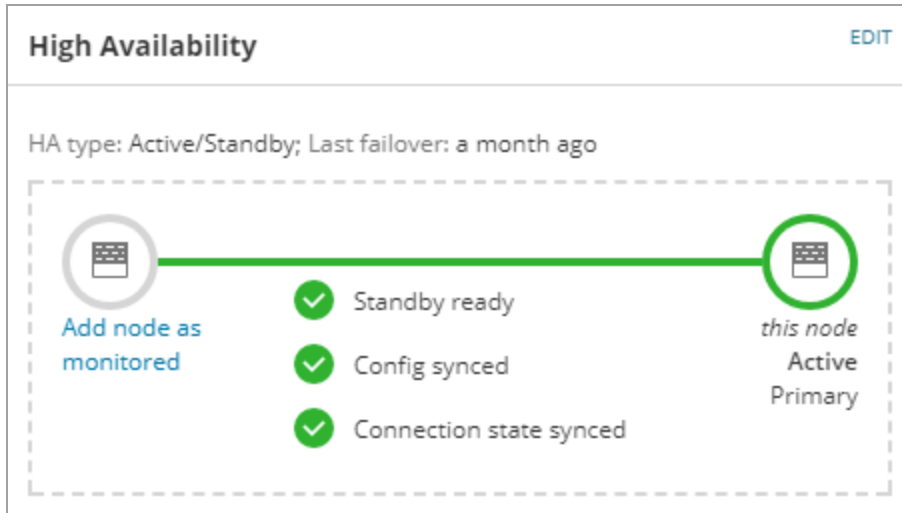
- Monitoring an Administrator context also lists other configured contexts in the widget.
- Monitoring a non-Administrator context only gives you information about the context.

Monitor high availability for Cisco ASA devices

On the Node Details for ASA - Summary, review the high availability information in the Platform Summary resource to help monitor your ASA devices.

Platform Summary		EDIT HELP
HARDWARE HEALTH	 1  6	
HIGH AVAILABILITY	 Unknown; See details	

Click the See details link, and view the High Availability widget on the Platform subview.



ASA node statuses

See the node status options for ASA devices.

The color of the circle indicates the node status.

Icon	Description/Action
	The node is up and running.
	The node's status is Warning. The node did not respond to a ping request and is fast-pollled for 120 seconds.
	The node is not monitored in NPM. For details, see Troubleshoot Unknown nodes . Add the node as monitored to get details about the node, such as the node name.
<div style="border: 1px solid #add8e6; padding: 5px;"> <p>i If the node is monitored with NPM, verify that you configured both an IP address and a stand-by IP address for each active ASA interface so the node can be paired correctly.</p> </div>	
	The node is down. The node did not respond during the fast-poll period of 120 seconds.

Labels next to the icons tell you what type of ASA high availability is configured, and the role of individual nodes:

- Standby/Active
- Primary/Secondary

ASA high availability statuses

NPM polls the following high availability statuses on ASA devices. NPM orders the statuses according to importance with device issues listed first.

- **Standby ready (up, down, or unknown)**

ASA devices (active and standby) see each other and agree that the standby ASA is ready for failover.

- **Configuration state (up, down, or unknown)**

If the Configuration state is synced, both ASA devices report that the configuration is synchronized.

If the Configuration state is not synced, ASA devices report that the configuration is not synchronized. If you have NCM installed, click to see the configuration difference.

- **Connection state sync (up, down, or unknown)**


State - synced means that both ASA devices report that the high availability state is synchronized.

The overall high availability status is indicated by the color of the line:

- Critical status (red): the Standby ready status is down, and the Configuration state and Connection sync are not relevant.
- Warning status (yellow): the Standby ready status is up, and Configuration and Connection states are either down or unknown.
- Up (green): the Standby ready status is up, and the other states are either up or unknown.
- Unknown (gray): the Standby ready status is unknown, and the other statuses are either up or unknown.

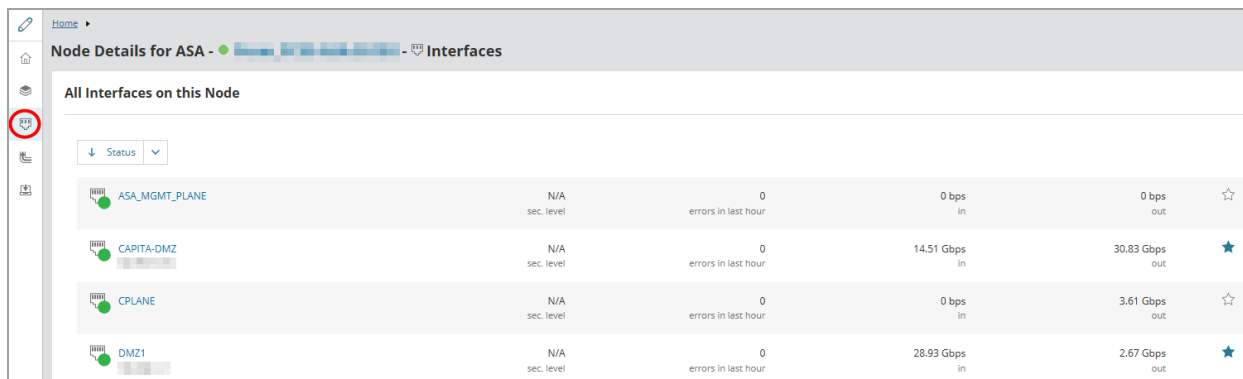
Monitor interfaces

Review the Bandwidth widget, or resource, that shows the traffic going through your favorite interfaces, and then click the Interfaces subview in the navigation bar on the left.

 If the Bandwidth widget is empty, you have no favorite interfaces. [Specify up to three favorite interfaces.](#)

NPM labels interfaces with the `nameif` attribute that reflects the interface function.

To add an interface to widgets on the Summary Page, click the star for the interface.



Interface	sec. level	errors in last hour	in	out	Star
ASA_MGMT_PLANE	N/A	0	0 bps	0 bps	☆
CAPITA-DMZ	N/A	0	14.51 Gbps	30.83 Gbps	★
CPLANE	N/A	0	0 bps	3.61 Gbps	☆
DMZ1	N/A	0	28.93 Gbps	2.67 Gbps	★

Select favorite interfaces and Site-to-Site VPNs for the Summary subview

Specify important interfaces and Site-to-Site VPN tunnels as favorite objects, and keep track of their status directly from the Node Details for ASA - Summary view.

1. For VPN tunnels, click the Site-to-Site VPN subview.
2. For interfaces, click the Interfaces subview.
3. Click the star for objects you want to see on the Summary subview. You can have up to three favorite interfaces and up to three favorite VPN tunnels.

The interfaces with stars are displayed on the Bandwidth widget and VPN tunnels with stars are displayed on the Favorite Site-to-Site VPN resource.

Monitor VPN tunnels on ASA firewalls in NPM

Get basic visibility to your nodes so that you can troubleshoot tunnels with issues.

1. Log in to the SolarWinds Platform Web Console.
2. On the Summary view, locate and click your ASA firewall node to go to the Node Details view.
3. Click the [Site-to-Site VPN](#) or [remote access VPN](#) icon in the subviews menu on the left side of the SolarWinds Platform Web Console.

Site-to-Site VPN

Site-to-Site VPN provides information about office-to-office tunnels.

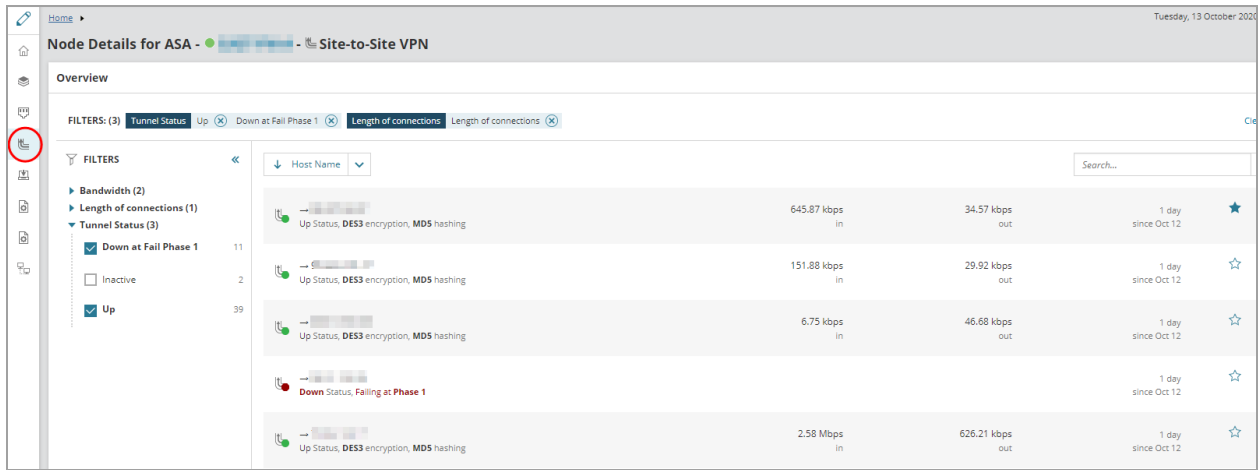
Non-existent or dead tunnels are automatically removed by the Orion Collector Service.

Review the list of Site-to-Site VPN tunnels on the ASA device. Use the search and filter options to find a Site-to-Site VPN tunnel and see more details.

Click the star icon to add a Site-to-Site VPN tunnel to favorites that are featured on the Node Details for ASA - Summary view.

Status information

- If the tunnel is down, see the information about the last phase completed successfully.
- For up tunnels, see the encryption, hashing info, in and out traffic, and the duration of the tunnel.

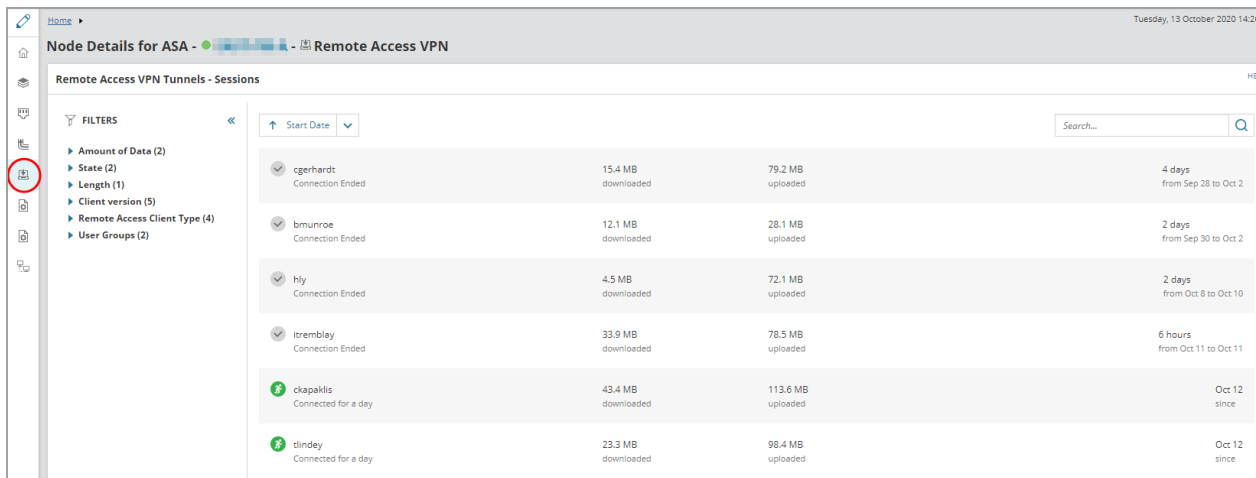


Remote access VPN


On the Remote access VPN subview, you can see a list of remote access tunnels, with the user name and tunnel duration details.

By default, non-existent or dead tunnels are removed after two days.

Search for tunnels, or filter results to find specific tunnels.



Review access lists on ASA firewalls in NPM and NCM

 To monitor access lists and access the ACL subview, you need NCM installed.

See what ACLs are applied to what interfaces and review those ACLs.

1. Log in to the SolarWinds Platform Web Console.
2. On the Summary view, locate and click your ASA firewall node to go to the Node Details view.
3. Click the Access Lists tab.

If you have NCM installed, you can compare the configuration of access lists. Click Compare ACL to go to the NCM widget.

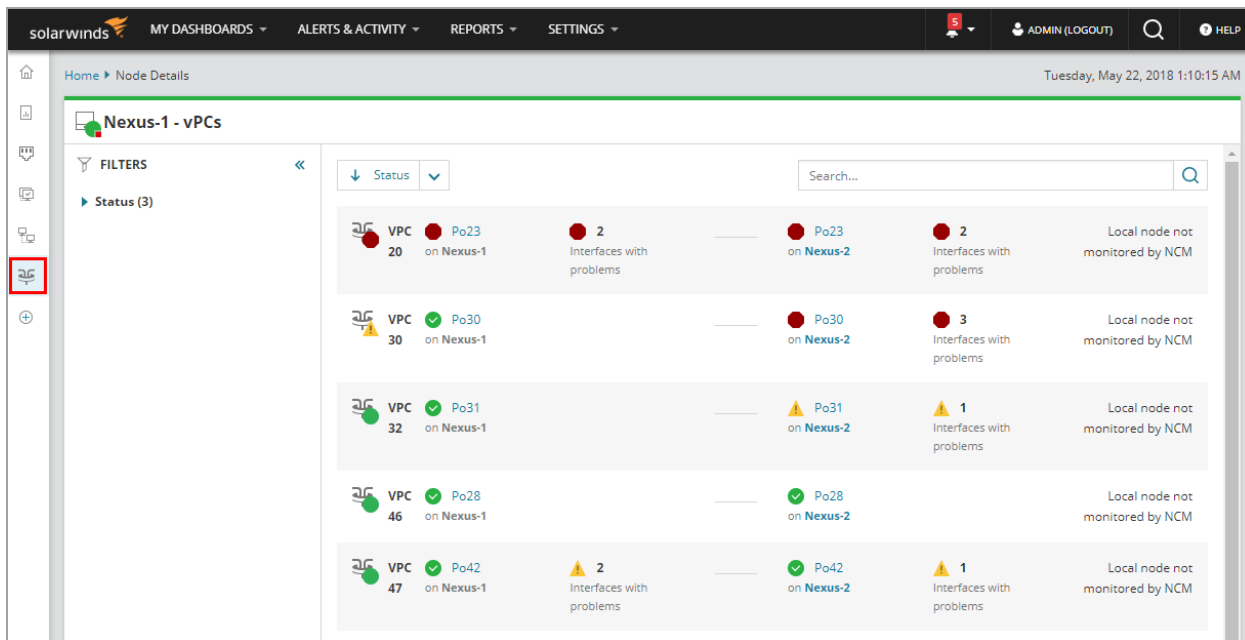
For more details, see [Manage ACLs on Cisco ASA and Nexus devices](#) in NCM documentation.

Monitor Cisco Nexus devices in NPM

With Network Insight for Cisco® Nexus, you can monitor the health and performance of your Nexus devices, view configured virtual Port Channels (vPCs) and their peer vPCs (Switch Virtual Interfaces, SVIs, are supported from NPM 2020.2.6).

The Node Details view for your Nexus devices includes a new Interfaces subview and Nexus-relevant VPCs subview.

To have the complete visibility into the health and performance of your Nexus switches, and to automate operational activities, install [Network Configuration Manager](#).



Out-of-the-box [alerts](#) for Nexus devices

- vPC on Nexus node is not up

Learn more

- [Set up monitoring for Cisco Nexus devices in NPM](#)
- [View information relevant for Nexus devices](#) (interfaces and vPCs)
- [Troubleshoot monitoring for Nexus devices](#)

Set up monitoring for Cisco Nexus devices in NPM


Data for monitoring Cisco® Nexus switches are polled by a combination of SNMP and CLI polling. To get accurate Nexus-specific information, [add the device to NPM as a node, and provide CLI credentials](#).

What does CLI polling provide for Nexus devices?

Enable CLI polling to receive additional details about virtual port channels (vPC), and to display the list of vPCs for your Nexus devices.

Information polled by CLI

- vPCs on the device and their status
- peer vPCs (Switch Virtual Interfaces are supported from NPM 2020.2.6)
- keep-alive links
- related interfaces

 You can provide the CLI credentials in the last step of the Add Node wizard [when adding a Nexus node](#), or on the [Edit Node page](#).

Requirements for Network Insight for Cisco Nexus devices

Requirement	Details
Cisco Nexus versions	5000 Series 7000 Series
Nexus configuration	<p>The vPC feature must be configured.</p> <p>To verify that you have vPCs configured, run the following command on the device:</p> <pre>show vpc brief</pre> <p>If the command doesn't list any vPCs, you need to configure them.</p> <p>NPM doesn't display vPC peer links. If the command lists only vPC peer links, adjust the configuration to include vPCs.</p>
Nexus user account	Credentials for logging into the Nexus device.
Enable password	If you have configured an enable password, you need to provide it.

Requirement	Details
SSH port	By default, port 22. Open an SSH port for accessing and polling Nexus devices through SSH.

Add Nexus devices for monitoring

i You need Node Management Rights. See [Define what users can access and do](#).

1. Click Settings > Manage Nodes, and then click Add Node.
2. Enter the IP address for the device.
3. Select Most Devices: SNMP and ICMP as the polling method, and enter SNMP credentials.
4. Enable [CLI monitoring](#):
 - a. Scroll down to CLI Polling Settings.
 - b. Select Enable CLI Polling, [enter the credentials](#), and click Test.

i Enter a user name and password for logging into the ASA or Nexus device.

If you have configured a security password for CLI polling on the device, provide it in Enable password. Without the Enable Password, CLI polling does not work.

CLI Polling Settings: Enable CLI Polling » [Learn More](#)

Enable polling for Cisco Nexus, Cisco ASA or Palo

Username:

Password: ✕

Enable Password: ✕

SSH Port:

Use Keyboard Interactive Authentication: ▼

5. To use a specific [device template](#), select it. See [NCM Getting Started Guide](#) for more information.
6. Complete the Add Node wizard.

NPM now polls the Nexus-specific information.

Enable CLI polling on monitored devices

To poll vPC data on Nexus devices already monitored in NPM, enable CLI polling.


1. On the node details view, click Edit Node in the Management widget.
2. Scroll down to the CLI Polling Settings section.
3. Select Enable CLI Polling, [enter the credentials](#), and click Test.
4. Click Submit.

You can now poll device-specific information, such as vPCs. The node details view for your Nexus devices includes the vPC subview listing all vPCs with their member ports.

Troubleshoot monitoring Cisco Nexus devices

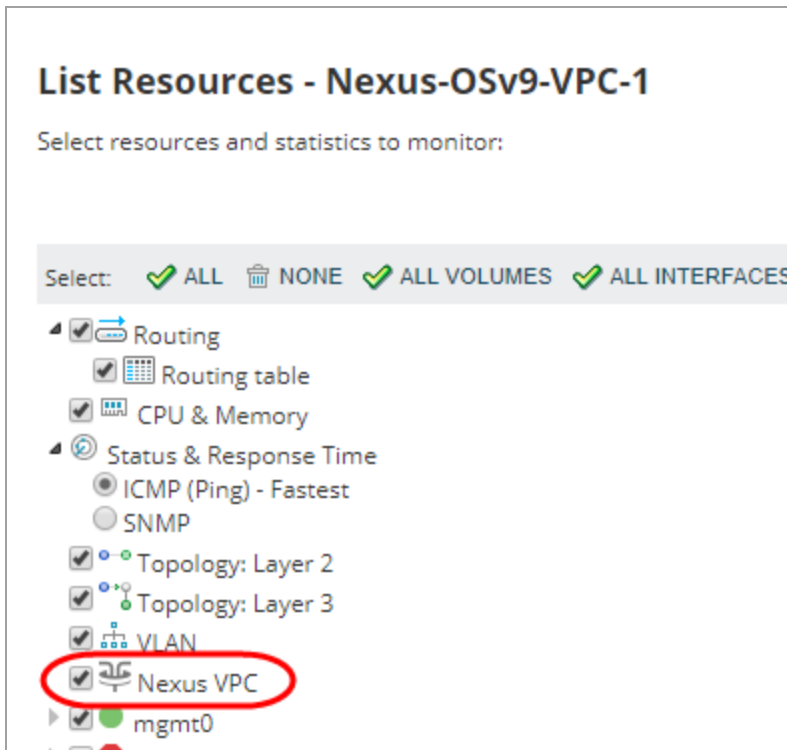
Verify the following:

- The node or interface is monitored in NPM.
If your vPC is configured, NPM can see that it has a peer. There are no details about the peer node. To see the details, [add the node to NPM for monitoring and enable CLI polling](#).

 If you only have NCM installed and cannot see any vPCs, download the running config to get a list of vPCs configured on the device. To get more information about vPCs, install NPM.

- You have enabled CLI polling on the device.
- You are using the correct user credentials to log in to the device.
- You are using the correct password for CLI access.

- You have enabled the Nexus vPC poller:
 - a. On the node details view, click List Resources in the Management resource.
 - b. Ensure Nexus VPC is selected.



Cannot see peer nodes?

If you cannot see the peer node for monitored vPCs, make sure the keep-alive link is alive.

To verify the keep-alive link, run the following command on the device:


```
show vpc peer-keepalive
```

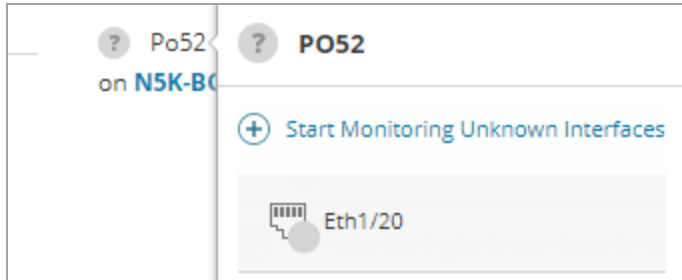
Review the command output. The `vpc keep-alive status` must have the following value:
`peer is alive.`

If the value is different, adjust the configuration on the device.

The interface (port channel) is grey

If an interface is grey and you only see the peer vPC name, click Start monitoring unknown interfaces, and select All interfaces on the List Resources page.

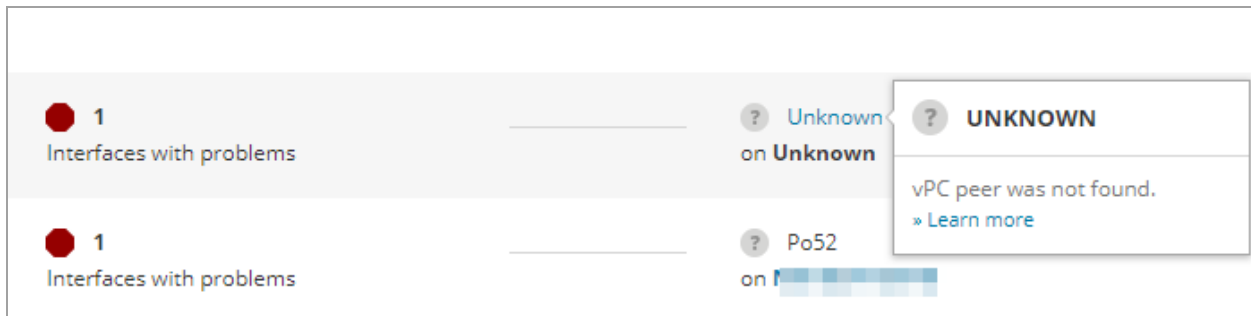
 If you don't want to monitor all interfaces on the device, select at least the interfaces displayed for the vPC on the vPCs subview.



Unknown vPC on the peer

If the vPC on the peer is unknown, the vPC is not configured on the peer node.

Check the configuration on the device, and make sure the vPC is configured correctly.



Extend logging

If the above steps do not solve your issue, enable the CLI session trace to extend logging:

- a. Click Settings > All Settings > CLI Settings in the Product Specific grouping.
- b. Click Enable Session Tracing, and click Submit.


Review the session trace files located at:

```
%ALLUSERSPROFILE%\Application Data\SolarWinds\Logs\Orion\CLI\Session-Trace
```


Access Nexus-specific information in NPM

i To display Nexus-specific details, the device must be monitored in NPM and CLI polling must be enabled. See [Set up monitoring for Cisco Nexus devices in NPM](#).

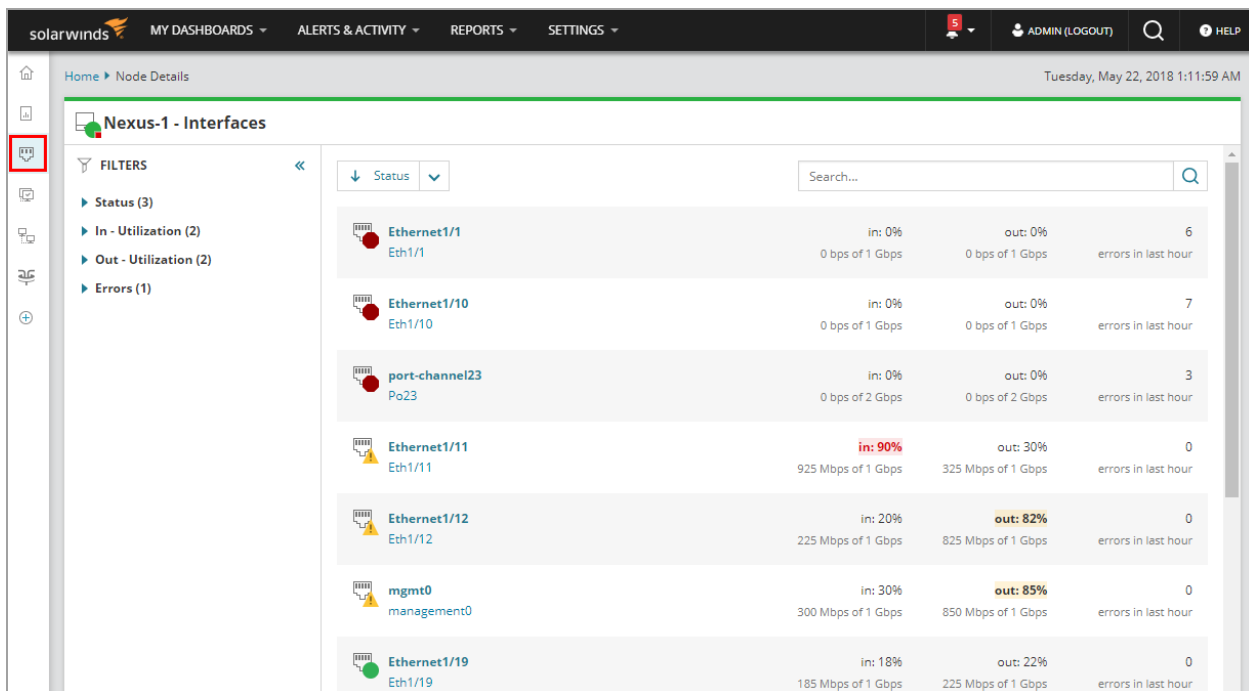
1. In the SolarWinds Platform Web Console, navigate to the Node Details view for a monitored Nexus device.
2. For an overview about the [health of interfaces on the device](#), review the Health Summary widget.



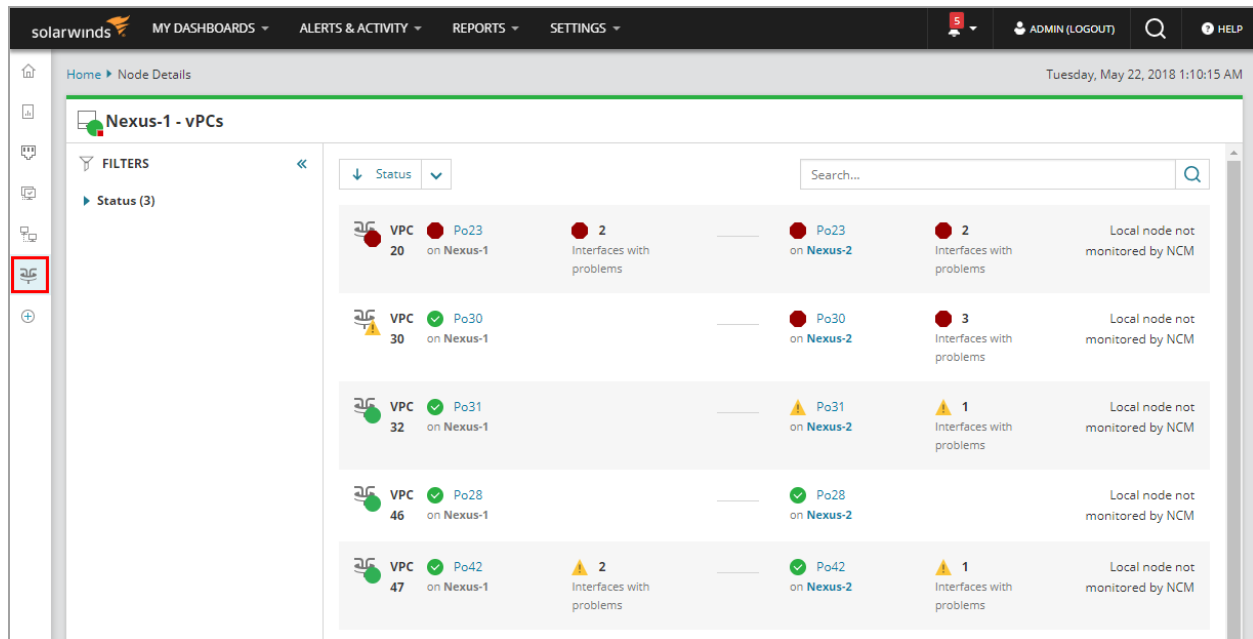
3. For more details about interfaces, click the **Interfaces** subview in the subviews menu on the left. The Interfaces subview lists monitored interfaces on the device, together with their in and out utilization, or errors in the last hour.

i To go to the Interfaces subview, you can also click a status group number on the Health Summary widget. The Interfaces subview opens, filtered by the status you clicked.

4. Use the filter and search options to quickly find more details about any interface on the Nexus node.



- Click the **vPCs** subview to list virtual port channels connected to the monitored Nexus device, together with the interface and the peer vPC details.



Monitor Cisco SwitchStack in NPM

With NPM, you can view the health of individual Cisco® SwitchStack® members, monitor power and data connections between the members, and quickly locate a switch with issues.

Out-of-the-box events and alerts notify you when a member, or a connection between members goes down.

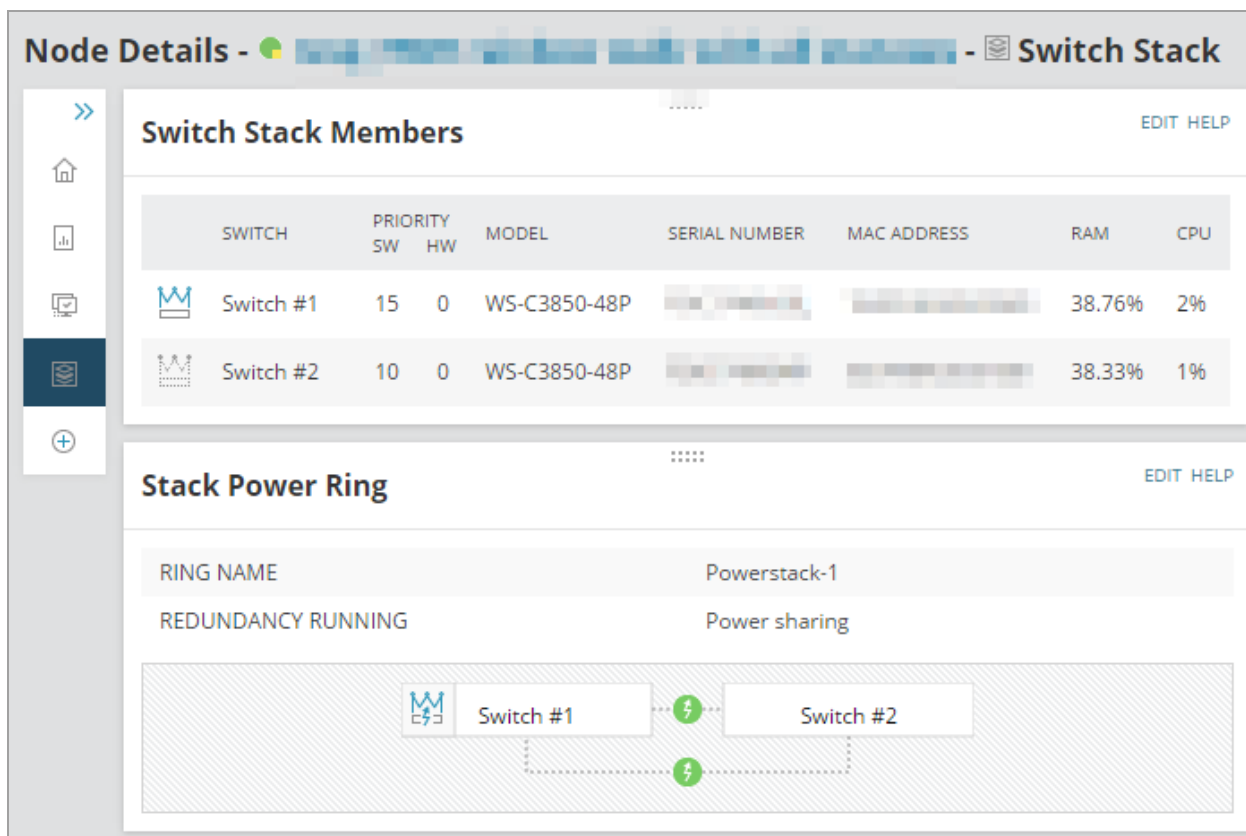
i [Add the Cisco SwitchStack for monitoring](#) as a node. The IP address is always assigned to the master switch (highlighted with a crown icon).

View stack members and rings

When you receive an alert about a SwitchStack problem, go to the SwitchStack node details page, and click the SwitchStack subview.



The subview provides member-specific monitoring with topology maps showing how the data ports and power ports are connected, and information to pinpoint switches with issues.

You can quickly see which switch is having issues, locate it by serial number in the stack, and replace it or resolve the issue.



Node Details - [Node Name] - Switch Stack

Switch Stack Members EDIT HELP

SWITCH	PRIORITY SW	HW	MODEL	SERIAL NUMBER	MAC ADDRESS	RAM	CPU
 Switch #1	15	0	WS-C3850-48P	[REDACTED]	[REDACTED]	38.76%	2%
 Switch #2	10	0	WS-C3850-48P	[REDACTED]	[REDACTED]	38.33%	1%

Stack Power Ring EDIT HELP

RING NAME: Powerstack-1

REDUNDANCY RUNNING: Power sharing

Diagram showing two switches (Switch #1 and Switch #2) connected by power lines (indicated by lightning bolt icons).

View the health of stack members

When you are monitoring hardware health on a Cisco SwitchStack node, you can see the health of individual switches in the stack. The health indicators inform you when the values on a switch are near the safe limits, or when they reach the critical stage.

1. Log in to the SolarWinds Platform Web Console, and go to the SwitchStack node details page, and click the Network subview.
2. Consult the Current Hardware Health resource.
3. Expand a switch in the stack to display hardware health monitors.

Current Hardware Health		MANAGE SENSORS	EDIT	HELP
DEVICE NAME		STATUS	VALUE	
Switch 1		●	3	
▶ Fan	●			
▶ Power Supply	●			
▶ Temperature	●			
Switch 2		●	3	

i The item in the Status column describes the number of sensors monitored on the switch, grouped by the status of the sensor.

You can now troubleshoot the SwitchStack member that is experiencing issues.

See also [Monitor hardware health](#).

Cisco SwitchStack events

Events for Cisco SwitchStack include messages about the following issues and changes:

- Stack ring redundancy loss
- Stack ring failure
- Members being added or removed
- Member number changes
- Master switch changes
- Power redundancy loss
- Power capacity change

Out-of-the-box alerts for SwitchStack


Out-of-the-box SwitchStack alerts inform you about the following items and more:

- SwitchStack Master Changed
- SwitchStack Data Ring Broken
- SwitchStack Member Number Changed
- SwitchStack Power Redundancy Lost


 Not all out-of-the-box alerts are turned on by default.

Create alerts based on SwitchStack events

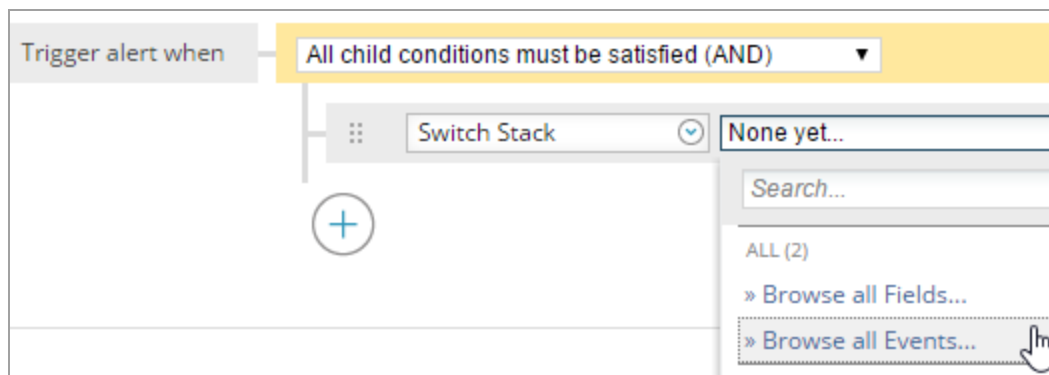
You can [configure additional notifications](#) based on SwitchStack events. For example, you can specify that when a stack ring fails, you want to receive an email with details.

 Out-of-the-box alerts cover the most frequent issues. Review available alerts and [duplicate and edit the alerts](#) if you only need small adjustments.

1. Select Alerts & Activity > Alerts, and click Manage Alerts.
2. Click Add New Alert.
3. On Trigger Condition, select the SwitchStack item you want to alert on.

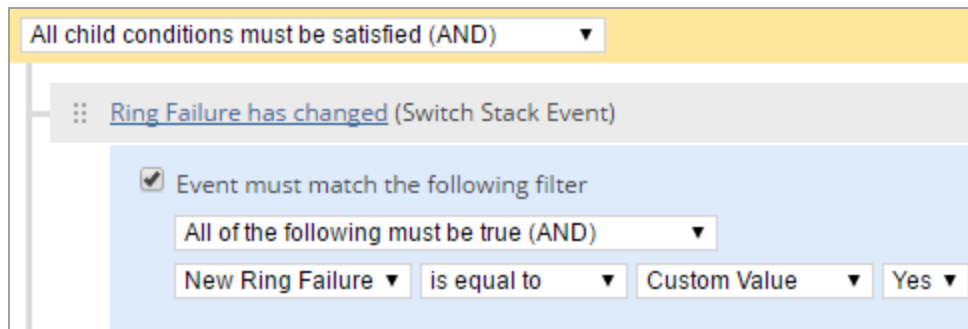


4. In Trigger alert when, select a condition, click the arrow in the second box next to the selected SwitchStack object, and select Browse all events.



5. Select the event you want to alert on and if necessary, complete the trigger condition.

For example, if you want to be notified about a SwitchStack ring failure, select the Ring Failure event, select Event must match the filter, and then select New Ring Failure is equal to Yes.



6. [Specify the trigger action](#) and complete the wizard.

After the trigger condition occurs, you will be notified about it both by the event and the trigger action you specified.

Network Insight for F5 BIG-IP load balancers in NPM

Network Insight provides comprehensive monitoring for the F5® BIG-IP® family of load balancers, giving you the insight you need to keep your most important services running smoothly. Use Network Insight to:

- Monitor the health and performance of all components of application delivery including WideIPs, virtual servers, pool members, and more.
- Identify the components that are contributing to slowness, service outages, or any service that could be affected by an infrastructure problem.
- Visualize your entire application delivery environment and get an instant status of a service or device. Click on any status indicator to see additional details about that component or to show relationships.
- Graphically display relationships and component status. Easily view the relationships from the service through the traffic managers, virtual servers, pools, and pool members along with a detailed status of each component.

i The performance statistics you can monitor on F5® BIG-IP® devices include device status and availability, CPU and memory performance statistics, and interface performance details.

Set up Network Insight for F5® BIG-IP® load balancers in NPM

To monitor the servers and connections in your load balancing environment, make sure your F5 devices meet the following requirements, add the F5 devices for monitoring, and enable F5 iControl.

Requirements

Requirement	Details
supported modules	F5 Local Traffic Managers (LTMs) BIG-IP DNS (formerly called Global Traffic Managers or GTMs)
SNMP	used to poll everything except for health monitors TMOS version 11.2 and later (including 12.0)

Requirement	Details
iControl by F5	used to poll health monitors and to enable and disable the rotation of pool members TMOS version 11.6 and later <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> i The iControl account used in NPM must be a local account on the F5 device. You cannot use Active Directory or TACACS accounts. </div>

Add F5 devices and enable iControl

[Add F5 devices](#) hosting global traffic managers (GTMs) and local traffic managers (LTMs) for monitoring.

i You need Node Management Rights. See [Define what users can access and do](#).

1. Click Settings > All Settings, and click Add Node in the Getting Started grouping.
2. Enter the IP address for the device.
3. Select Most Devices: SNMP and ICMP as the polling method.
4. Enable F5 iControl:
 - a. Scroll down to Additional Monitoring Options, select Poll for F5 iControl, provide the credentials for accessing the iControl API on the F5, and click Test.

i Starting with 2023.2, the credentials are used for retrieving an API token. The token is further used for identification. Using API tokens for authentication improves both security and performance. To disable token-based authentication, see Polling Settings.

F5 iControl Polling Settings Poll for F5 iControl

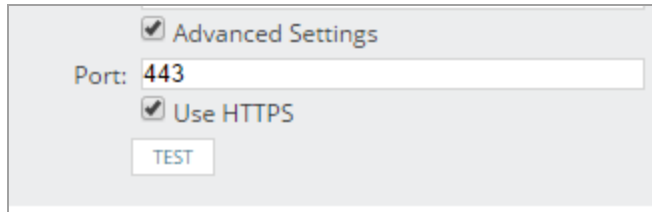
User name:

Password:

Advanced Settings

- b. If iControl does not run on the default port 443, select Advanced Settings, and provide

the port.



Advanced Settings

Port: 443

Use HTTPS

TEST

5. Complete the Add Node wizard.

Both status information and health statistics will be collected on the F5 device, and you can now monitor your load balancing environment:

- [Monitor services delivered by F5® BIG-IP® load balancers in NPM](#)
- [Take an F5 pool member out of rotation in NPM](#)

i See [Discover your network with the Discovery Wizard](#) to add more F5 devices at the same time.

Enable iControl on F5 load balancers

When your F5 devices are already monitored in NPM, make sure iControl is enabled. F5 iControl API is used for collecting health monitor statistics from load balancers, and for enabling and disabling the rotation of pool members.

1. Click Settings > Manage Nodes.
2. Select the node, and click Edit Properties.
3. [Enable F5 iControl](#).
4. Click Submit.

Now you can enable and disable the rotation of pool members and see the health monitors polled on the node.

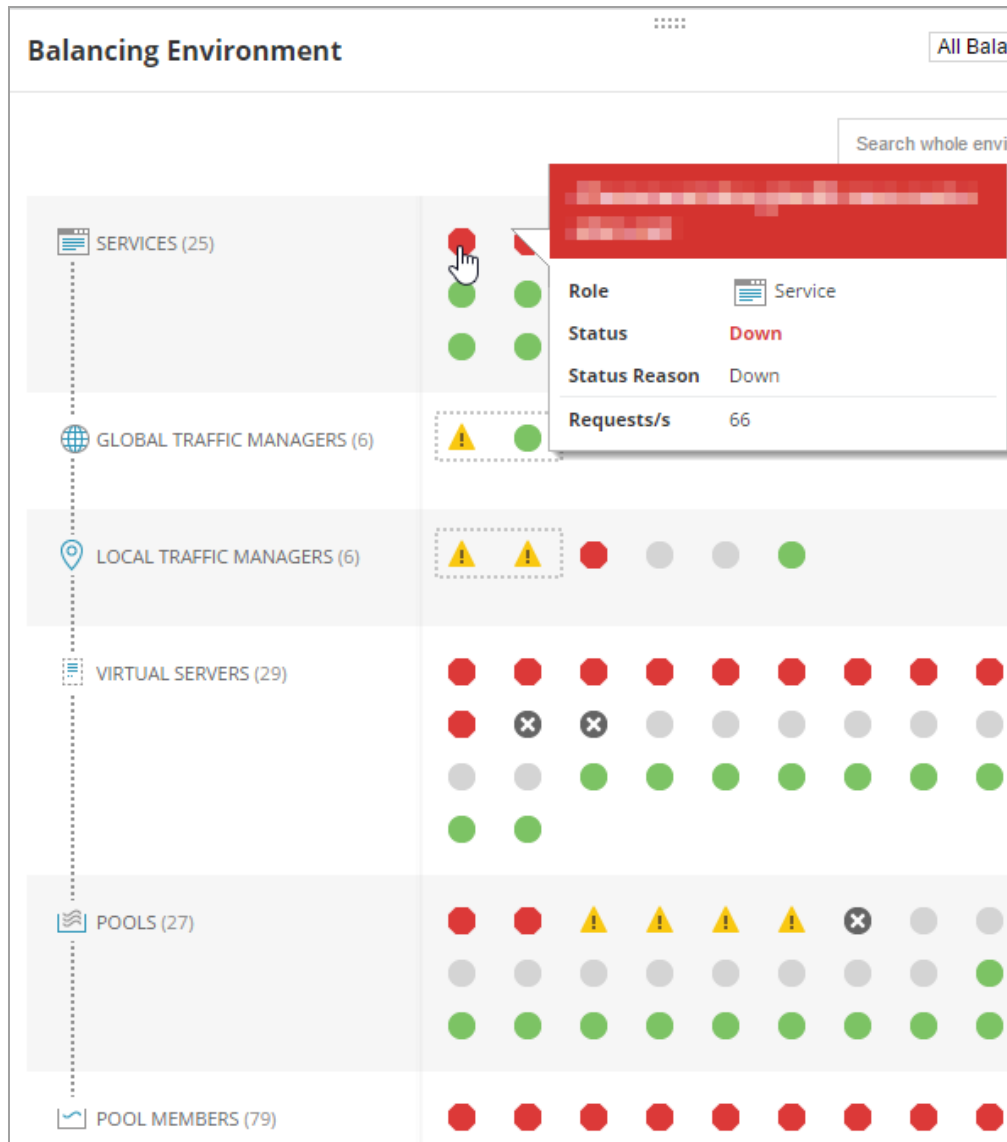
Monitor services delivered by F5® BIG-IP® load balancers in NPM

A load-balanced service is comprised of many components that work together. In the Balancing Environment widget, you can browse all of these components and their relationships and status.

1. Go to My Dashboards > Network > Load Balancing.

The page shows an overview of your load balancing environment.

At the top, you can see your load balanced services. Below Services, there are Global Traffic Managers (GTM) that host the services. The GTMs send users to your Local Traffic Managers (LTM). Your LTMs present virtual servers which are made up of pools, and individual pool members hosting the content.



i Dotted rectangles highlight high availability (H/A) clusters.

2. Point to a service to review the tooltip.
3. To see more detailed information about the component, such as the number of concurrent connections and the load balancing algorithm, click the service and select Display Details Page.

F5 SERVICE DETAILS:

Service Details EDIT HELP

STATUS **Down**

STATUS REASON Down

Balancing Environment EDIT HELP

for [redacted]

Search [redacted]

SERVICES (1)

GLOBAL TRAFFIC MANAGERS (2)

LOCAL TRAFFIC MANAGERS (2)

VIRTUAL SERVERS (2)

POOLS (2)

POOL MEMBERS (5)

Concurrent Connections

TOP 5 VIRTUAL SERVERS BY NUMBER

Apr 29 2016, 5:16 pm

Zoom 1h 12h

CONNECTIONS

50.0 k

40.0 k

30.0 k

20.0 k

10.0 k

30 Apr

18 Apr

Virtual Server

Role Virtual Server

Status **Down**

Status Reason Down

Port [redacted]

IP Address [redacted]

Assigned pool [redacted]

Connections 176

F5 Events

LAST 12 MONTHS

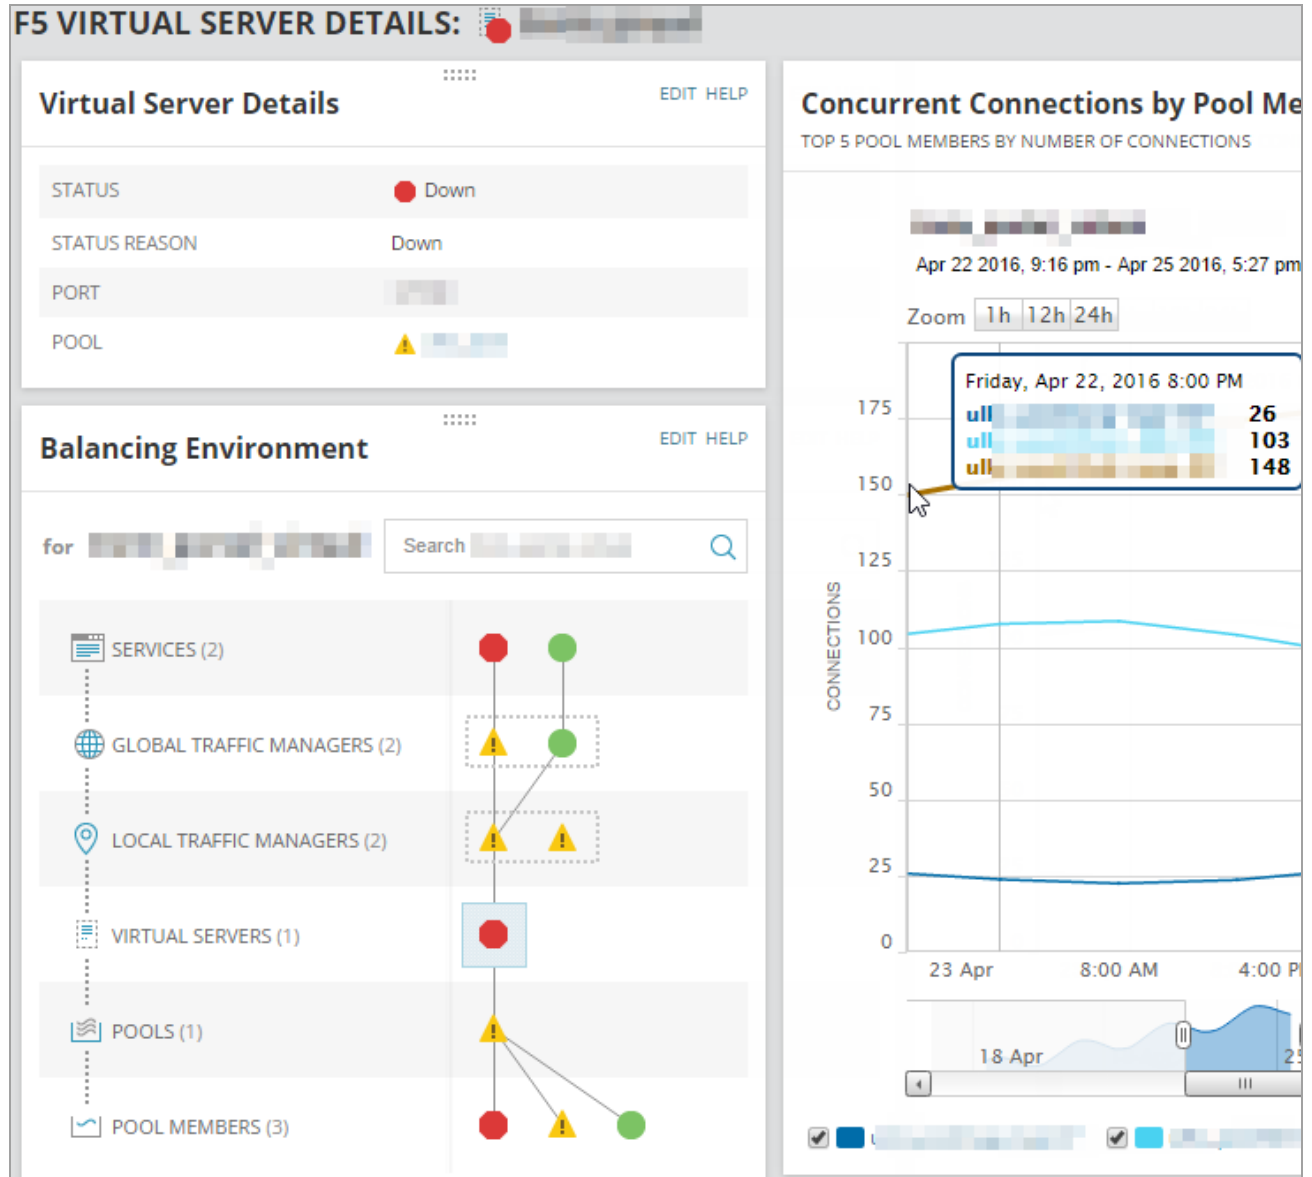
ACTIVITY TIME

Status of F5 GTM service [redacted] changed to Down. 4/25/2016 2:41:44 PM

i The light blue square in the Balancing Environment resource indicates your current position.

- The relational view of the load balancing environment sticks with us on the details page so we can continue to explore around. Select a virtual server, and click Display Details Page. This shows us the number of active connections for each pool member.

We can see the load balancing algorithm and how evenly it is distributing load.



5. Navigate through the load balancing environment to view the health of individual components.
6. Drill in to the components with issues, review the data provided by NPM, such as the status of load balancing components and the reason why they are not up. Use the data to troubleshoot the issues.

Status of F5 devices in NPM











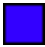

F5 status is information polled directly on the F5 device through SNMP. NPM also polls the status reason from the F5 device and displays the reason in the element's tooltip and on the details pages.

The status for GTM and LTM modules is calculated. LTM status is calculated based on virtual server, GTM status is calculated on the WideIP (service).

i F5 status is not the same as the node status. Both node statistics and F5 statistics are polled through SNMP, but from separate parts of the MIB tree. That's why a node can be up (Orion statistics), but the corresponding load balancing component is down (F5 statistics).







F5 device status mapping to Orion status

We mapped the status icons used for F5 devices to SolarWinds Platform icons. For explanation of F5 statuses, see [F5 support help online](#).

F5 status	Orion status
 Available	 Up
 Unavailable	 Warning
 Down	 Down
Disabled	Unmanaged
 (error)	 (disabled on F5)
 (unlicensed)	 (unreachable)
 Unknown	 Unknown

F5 status in Orion

The table explains what the Orion statuses mean for individual components in the load balancing hierarchy. Status is usually polled on devices, but for some components, such as GTMs and LTMs, it is calculated based on polled values for their child objects.

Status	Load balancing component						
	Service	GTM	LTM	Virtual Server (VS)	Pool	Member	F5 Server
	Reported on the device	All services assigned to the GTM are up	All virtual servers are up	F5 device reports the VS as up	F5 device reports the pool as up	F5 device reports the pool as up	F5 device reports the server as up
	Reported on the device	At least one service is not up	At least one virtual server is not up	Unavailable based on connection limit	No members are currently available	Unavailable based on the connection limit	Unavailable based on the connection limit
	Reported on the device	All services are not up	All virtual servers are not up	Associated objects marked the VS as unavailable. User action necessary	All members are unavailable	The parent F5 server is down or the monitor on the member marked it as down	Down based on monitor
	-	-	-	Unmanaged: Disabled on the F5	-	Unmanaged: Disabled on the F5	Unmanaged: Disabled on the F5
	-	-	-	Unreachable: Parent LTM is down	Unreachable: Parent LTM is down	Unreachable: Parent LTM is down	-
	Reported on the device	GTM is added but not polled yet	LTM is recognized by GTM, but not managed in Orion	Unknown	Unknown	No monitors assigned	No monitors assigned

F5 high availability in NPM

High availability (H/A) is configured on the device level. It does not matter whether you have a GTM or LTM installed on the device, the module is covered by H/A. Devices are connected in traffic groups. If one device fails, another device in the group handles its requests. Devices in a traffic group synchronize the configuration. The configuration is reflected by the synchronization status.

In NPM, we poll the failover and synchronization status.

Devices in one traffic group are connected by dotted rectangles on the Balancing Environment resources. Display the tooltip to see details about the H/A failover and synchronization status.

The screenshot displays the 'Balancing Environment' dashboard. On the left, there are four main categories: SERVICES (25), GLOBAL TRAFFIC MANAGERS (6), LOCAL TRAFFIC MANAGERS (6), and VIRTUAL SERVERS (29). Each category is represented by a vertical dotted line. To the right of these categories is a grid of colored circles representing individual resources. A tooltip is shown over one of the resources, providing the following information:

Role	Global Traffic Manager
Status	Warning
Status Reason	One or more Services are not Up.
IP Address	[Redacted]
Hosting Node	[Redacted]
H/A Status	Standby and In Sync
Requests/s	518

You can see the H/A statuses in tooltips, and on the LTM or GTM detail views. The GTM or LTM details resource shows the H/A status and synchronization status. In the High Availability resource, you can check the details about other members of the traffic group.

Global Traffic Manager Details
EDIT HELP

STATUS	Warning
STATUS REASON	One or more Services are not Up.
POLLING IP ADDRESS	[Redacted]
HOSTING ORION NODE	[Progress bar]
H/A STATUS	Standby
H/A SYNC STATUS	In Sync

High Availability
EDIT HELP

NAME	POLLING IP	H/A STATUS	SYNC STATUS
[Redacted]	[Redacted]	Active	In Sync

F5 health monitors in NPM

To monitor the health of your load balancing environment, NPM polls health monitors on your F5 servers (nodes), and on F5 pool members. Health monitors run periodic tests for network service availability, such as ICMP, HTTP, IMAP, or MSSQL.

To get the health statistics, [F5 iControl API must be enabled](#).

i F5 health monitors are not related to hardware health. The status of an element is based on health monitors polled by F5 iControl API.

Go to a pool member or an F5 server details page to review the health monitors widget.

Pool Member Health Monitors			
MONITOR NAME	TYPE	PORT	STATUS REASON
icmp	ICMP	0	-
sw_web_http	HTTP	80	Unable to connect. No successful responses received before deadline.

i Health monitors require at least one pool member to be up. If no pool members are up, the LTM, the virtual server, and the pool will all be marked as down. Drill down into the pool member to see why it is down.

Events, alerts, and reports for Network Insight for F5® BIG-IP® load balancers in NPM

Each F5 details page includes the F5 Events resource that displays events relevant for the object. Click an event to go to the details page for the object with issues and review the situation.

F5 Events	
ACTIVITY	TIME
Pool [redacted] on F5 device [redacted] is Down.	4/21/2016 4:02:53 PM
F5 server [redacted] is Unknown.	4/21/2016 3:58:33 PM
F5 server [redacted] is Unknown.	4/21/2016 3:58:24 PM
F5 pool [redacted] has 0% of active servers. The alert triggers when less than 30% servers are active.	4/21/2016 1:50:44 PM

Load balancing events include:

- A component status changes to down
The components include virtual IP, Pool, Server, Wide IP, GTM, or servers.
- Health probe status changes up and down
- H/A peer status or synchronization change

- Server is taken out and placed in rotation
- Concurrent connections per pool member exceed a threshold

Out-of-the-box alerts for F5 load balancers

Out-of-the-box alerts cover the most critical issues in your F5 load balancing environment. For example, alerts warn you if the status of your F5 service changes or if a server goes down.

Out-of-the-box reports

NPM includes several out-of-the-box reports for F5 that you can use to view trends, establish baselines, or identify potential issues, such as:

- Average LTM Connections over the last 30 days
- Average service availability over the last 30 days
- Average service resolutions per second over the last 30 days

Take an F5 pool member out of rotation in NPM

When you need to perform maintenance on one of the pool members providing a service, take the server out of rotation so that you can perform maintenance without impacting end users.

Taking server out of rotation means you put the pool member in maintenance mode.

F5 devices support Disabled and Forced Offline modes. NPM uses the Disabled maintenance mode.


i Taking a pool member out of rotation requires that you have [enabled F5 iControl on the device](#).

Why shouldn't I start maintenance immediately after I take a pool member out of rotation?

When you put a pool in maintenance mode, there are still users connected to the server. Disabling the server only disables brand new connections.

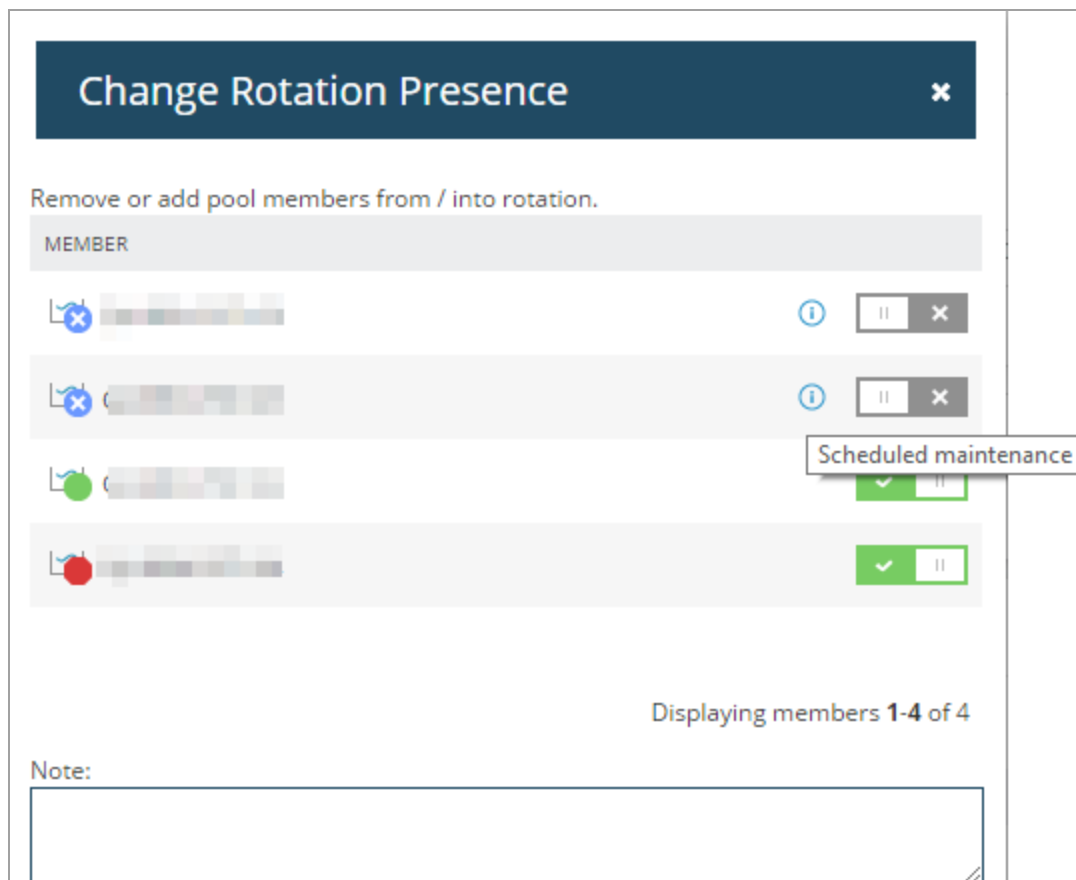
The maintenance mode only changes how the LTM handles incoming requests.

- New users are not sent to the server while the servers is in maintenance mode.
- In the Disabled mode, new connections with existing sessions are not affected. Users who open a new TCP session but were previously using the server, will continue to be sent to this server.
- Existing connections are not affected. Users with an open TCP session with the server will continue to use it.

 SolarWinds recommends that you wait until the existing connections end or time out not to impact the connected users.










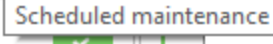



Take a pool member out of rotation

1. Click My Dashboards > Network > Load Balancing, and locate the parent pool of the pool member.
2. Click the parent pool, and click Display Details Page.
3. On the Pool Details view, find the Pool Members resource, and click Change Rotation Presence.
4. Click the green check mark icon next to the pool member to remove it from rotation, and click Submit.




Change Rotation Presence ✕

Remove or add pool members from / into rotation.

MEMBER	Info	Pause	Close
 [blurred]			
 [blurred]			
 [blurred]			
 [blurred]			

Displaying members 1-4 of 4

Note:

 Add a reason for taking the pool member out of rotation in the Note field. An info icon will appear next to the pool member, and your note will be displayed as a tooltip when you hover over the info icon.

The pool is removed from rotation now. To prevent user impact, watch the connection count for the pool member. It should decline over time as existing users finish their sessions and no new users are added. After the connection count has become low, you can begin maintenance.

Network Insight for Palo Alto - monitor Palo Alto firewalls with NPM

Network Insight for Palo Alto firewalls automates the monitoring and management of your Palo Alto infrastructure to provide visibility and help ensure service availability.

For Palo Alto firewalls, you'll find the following subviews:

- [Site-to-Site VPNs](#): Review names of tunnels, status, failure reason message, IN/OUT transferred data, encryption, hashing, virtual system, security zone, and local and peer IP addresses
- [GlobalProtect VPNs](#): Review username, public IP, virtual IP, time of the connection, status, and transferred data

Requirements

Requirement	Detail
REST API access	You need valid credentials to poll Palo Alto devices through REST API.
Log at Session End	To poll Global Protect throughput data, Log Settings in the Action tab for Security Policy Rule must be configured to Log at Session End.
Palo Alto account permissions	To access the statistics, you need an account with the administrator role and the following permissions set for the XML API: <ul style="list-style-type: none"> • Log • Operational Requests

How to monitor Palo Alto devices

To access the Site-to-Site VPNs and Global Protect VPN subviews, [add the device to NPM](#) for monitoring or [enable polling for Palo Alto](#) data on devices already monitored in NPM.

To access the device, NPM calls the device and requests a REST API key, also known as "session key". NPM then accesses the device through the key, not through the credentials directly.

A change of credentials on the firewall can thus lead to issues because NPM needs not only the credentials, but also the session key to poll the data. When you change credentials on the firewall, go to the Edit Node page and change credentials in NPM.

i For Palo Alto devices, NPM provides the Site-to-Site tunnel down out-of-the-box-alert. You can [configure custom notifications](#) based on Palo Alto events and [custom reports](#) showing statistics relevant for Palo Alto devices.

Add Palo Alto devices and enable Palo Alto polling

[Add Palo Alto devices](#) for monitoring.

i You need Node Management Rights. See [Define what users can access and do](#).

1. Click Settings > All Settings, and click Add Node in the Getting Started grouping.
2. Enter the IP address for the device.
3. Select Most Devices: SNMP and ICMP as the polling method.

i If you only need information about tunnels, you can choose ICMP here. If you also want to see data about the Palo Alto node itself, such as traffic, CPU, or memory, select SNMP.

4. Scroll down to Additional Monitoring Options, and select Poll for Palo Alto.
5. Provide the credentials for accessing the Palo Alto device and click Test Credentials.

If the certificate is [not validated successfully](#), for example if it is a self-signed certificate, use a different certificate, or click Accept Certificate to approve the certificate.

Additional Monitoring Options:

- UCS Manager credentials
If the node hosts an UCS Manager, ch
- Poll for F5 iControl
The F5 iControl API is used to collect
- Poll for Cisco ACI
Cisco APIC REST API is used to collect
- Poll for Palo Alto
Palo Alto REST API is used to collect S

Palo Alto Polling Settings

User name:

Password:

6. Complete the Add Node wizard.

NPM now polls Palo Alto details, and you can access the Palo Alto subviews for the device. The polling frequency is the Default Node Statistics Poll Interval and is 10 minutes by default.

Enable polling for Palo Alto on a monitored node

Palo Alto firewalls are polled using REST API to collect Site-to-Site and GlobalProtect VPN information.

1. Click Settings > Manage Nodes.
2. Select the node, and click Edit Properties.
3. Enable Palo Alto polling:
 - a. Scroll down to Additional Monitoring Options, and select Poll for Palo Alto.
 - b. Provide the credentials for accessing the Palo Alto device and click Test Credentials.
4. Click Submit.

NPM now polls Palo Alto details, and you can access the Palo Alto subviews for the device.


View Site-to-Site tunnels on a Palo Alto firewall

By default, NPM polls data every ten minutes using several different polls. At first, NPM polls a list of all entities, and then NPM polls additional data for each entity.

NPM polls static properties, and calculates bandwidth consumption based on the total amount of bytes coming through the firewall.

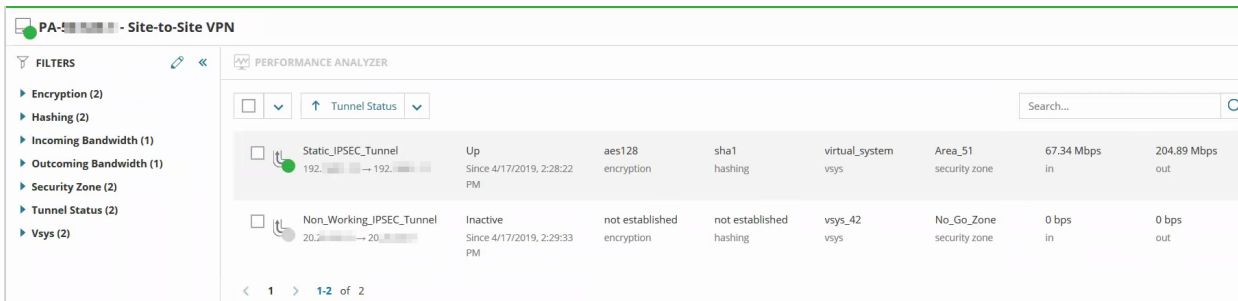
If connectivity issues happen during polling, partial details can be displayed.

The following table captures how statuses polled on Palo Alto devices are displayed in Orion:

Palo Alto status	Orion status	Value
Active	Up	1
Inactive	Down	2
	<div style="border: 1px solid #ccc; padding: 5px; display: inline-block;">  If the tunnel is down in NPM, it's inactive in Palo Alto. </div>	
Init	Inactive	24
Other statuses	Unknown	0

Tunnels are constructed between two firewalls. If NPM discovers that it is monitoring the same tunnel from another firewall, just another side of it, and the node is also monitored with NPM through REST API, you can click the underlined IP address to access the node details view.

To get historical data for monitored tunnels, select up to three tunnels, and click Performance Analyzer to open the tunnels in PerfStack.



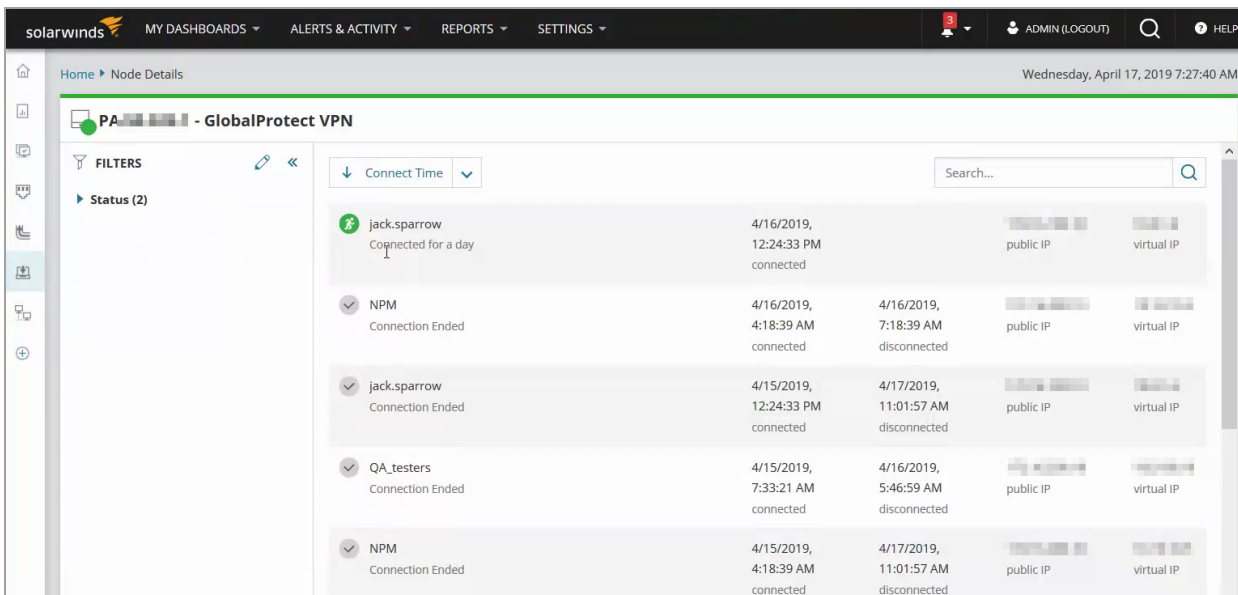
View the GlobalProtect VPN subview

To display a list of active sessions on your Palo Alto firewall, go to the node details view, and click the Global Protect subview.

The subview displays a record for each session. When the session ends, you can see the end time for it. When a session is renewed, it is displayed as a new session in the Global Protect list.

NPM keeps the history of Global Protect sessions for two days. When a session is connected, you can see the following details:

- Public IP provided by the service provider
- Virtual IP provided by the firewall



PerfStack

In Performance Analysis dashboards, you can search for Site-to-Site entities and display the following metrics for them:

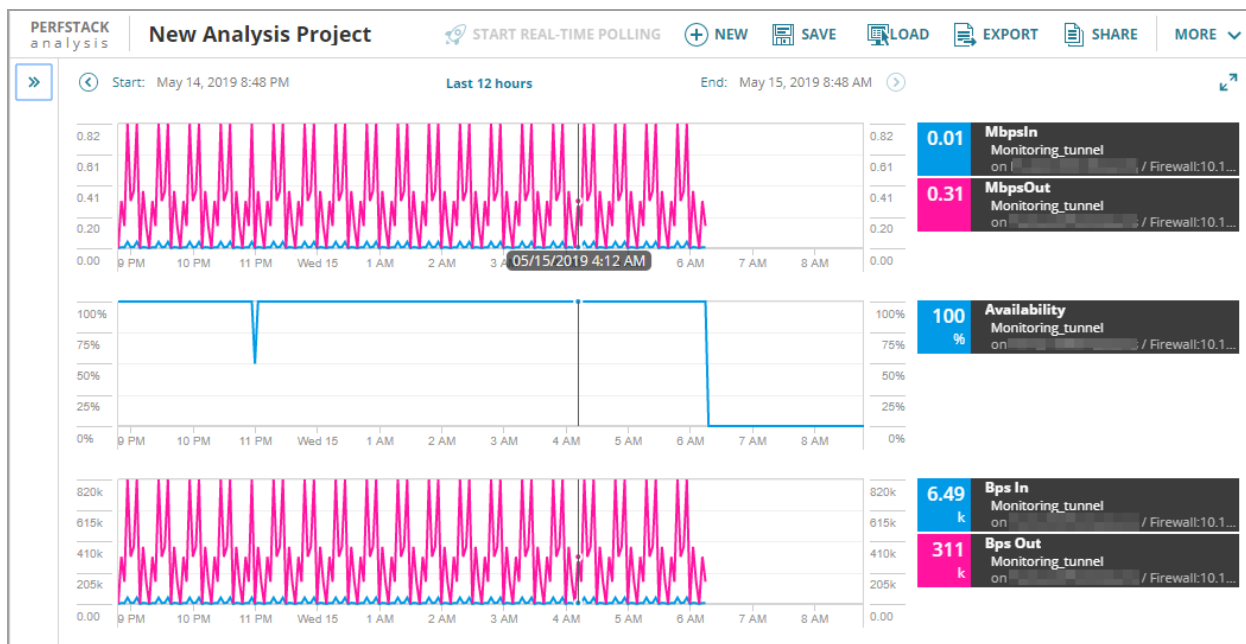
- Availability
- InBps
- OutBps

Access Palo Alto data in PerfStack

1. Go to PerfStack by clicking My Dashboards > Home> Performance Analysis.
2. Search for Palo Alto entities and metrics.

Access Palo Alto data in PerfStack from Palo Alto views

1. Go to the node details view for your Palo Alto firewall.
2. Select up to three tunnels in the Site-to-Site or GlobalProtect VPN subview, and click Performance Analyzer.



Intelligent Maps

You can display Palo Alto firewalls on Intelligent Maps.

Troubleshooting

Failure reason is not accurate

When a Site-to-Site tunnel is down, you can see the reason. This information is polled from traffic logs, based on the last message indicating failure. Whereas this can help you understand your issue, there might be issues. These issues might be caused by the asynchronous nature of polling, NPM might get various information, not related to the current issue.

My Global Protect session is displayed as "ended" although it is active

When a session connecting remotely to the office fails, NPM informs you that it ended and records the end time. If you renew the session, it appears in the Global Protect VPNs subview as a new session.

Monitor wireless networks in NPM


SolarWinds Network Performance Monitor can monitor any 802.11 IEEE-compliant autonomous access point (AP) or wireless controller, and provide details about access points (AP), wireless clients, wireless controllers, thin APs, and rogue APs.

NPM automatically recognizes your wireless APs and controllers as wireless devices when they are added to the SolarWinds Platform database. See [Discover and add network devices](#).

The wireless interfaces are not found during the discovery process. When a wireless device is added, an automatic inventory search is performed. Each wireless interface found is added to the database and the polling begins.

View wireless data in NPM


The Wireless Summary view consists of three tabs - Access Points, Clients, and Rogues. The tabs list monitored wireless access points (APs), clients connected to each AP, and rogue access points.

 You can display the coverage of your wireless access points or the location of connected clients in a map. See [Creating Wireless Heat Maps](#) and [Viewing the location of clients in wireless heat maps](#).

Access point details include the AP name, IP address, device type, SSID, channels used, and the number of clients currently connected.

Client details include client name, SSID, IP Address, MAC Address, Received Signal Strength Indication (RSSI), time connected, data rate, bytes received, and bytes transmitted.

Rogue access point details include MAC Address, SSID, and the number of channels.

 The following IPv6 statistics are currently not monitored:

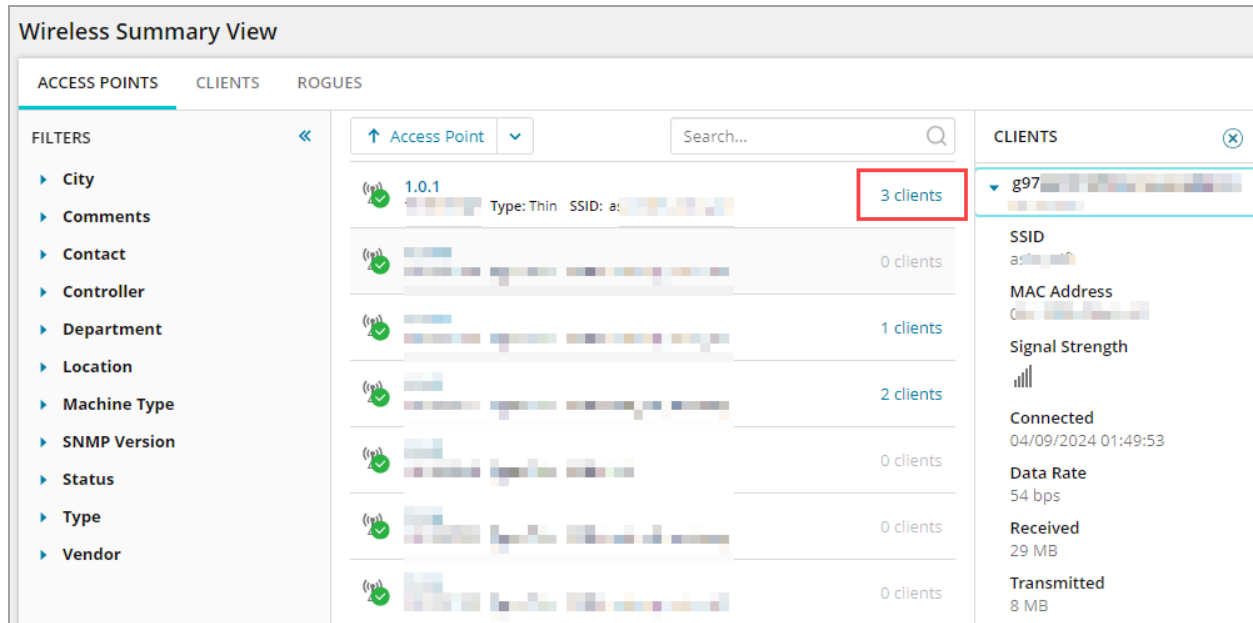
- Connections between wireless users and access points
- Connections between thin access points and controllers

To view wireless access points and clients:

1. Log in to the SolarWinds Platform Web Console.
2. Navigate to Wireless Summary View through My Dashboards > Wireless in the Network menu.
3. To list monitored access points, clients, or rogues, click the corresponding tab.

i To locate an access point, client, or rogue, use the filters, ordering options, or the search box in appropriate tab.

- To see clients currently connected to an access point, locate it in Access Points and click the clients link at the end of the access point row. Clients will be listed in the right part of the page. Expand clients to see more detail, such as SSID, MAC Address, Received or Transmitted data.



The screenshot shows the 'Wireless Summary View' interface. It has three tabs: 'ACCESS POINTS', 'CLIENTS', and 'ROGUES'. The 'ACCESS POINTS' tab is active. On the left, there is a 'FILTERS' sidebar with categories like City, Comments, Contact, etc. The main area displays a table of access points. The first row is selected, and a red box highlights the '3 clients' link at the end of the row. On the right, the 'CLIENTS' tab is active, showing details for a client with ID 'g97'. The details include SSID, MAC Address, Signal Strength, Connected time (04/09/2024 01:49:53), Data Rate (54 bps), Received (29 MB), and Transmitted (8 MB).

- To display the details view for an access point, click the access point. The node details view is specific for the selected device. See [Specify views for device types](#).

Monitor Meraki wireless infrastructure in NPM

[Add Meraki organizations to NPM](#) as external nodes, poll information from the cloud and [display polled data](#) in the SolarWinds Platform Web Console.

What does NPM monitor for Meraki infrastructure?

Meraki infrastructure provides centralized management of end devices, such as wireless devices, switches, or security appliances, as a cloud service. Physical devices act as thin Access Points (APs), managed by the cloud system. Physical devices are installed at physical locations and assigned to customers.

In NPM, Meraki objects are mapped to wireless entities.

Meraki Model	How do you see Meraki objects in Orion?
Organization	Controller
Access Point	Thin Access Point
Client	Client

Metrics with limited support

- **SSID information:** SolarWinds Observability Self-Hosted only polls SSID information for access points to which clients are connected.

Requirements


- Meraki account with administrative privileges
- Enabled access to the Cisco Meraki Dashboard API and generated API key. For details, search for "Cisco Meraki Dashboard API" at <https://documentation.meraki.com> (© 2015 Cisco Systems, Inc, available at <https://documentation.meraki.com>, obtained on February 1, 2017.)

Add Meraki organization to NPM

To monitor Meraki infrastructure with NPM, add the Meraki organization to the SolarWinds Platform database as an external node.

 Each Meraki organization monitored with NPM uses a node license.

1. Log in to the SolarWinds Platform Web Console as an administrator.
2. Click Settings > Manage Nodes, and click Add a Node.
3. In Polling Method, select Orchestrators: API.
4. Under Orchestrators, select Meraki Devices.

 The Polling Hostname or IP Address is disabled and `api.meraki.com` is used as the default host name for Meraki networks. If your Meraki organization has a different domain suffix, such as `api.meraki.cn` or `api.meraki.ca`, go to Advanced configuration, and change the MerakiCloudName.

- a. Log in to Advanced configuration. See [Access the Advanced Configuration](#).
- b. Search for MerakiCloudName on the Global tab.
- c. Change the MerakiCloudName to your organization and save your changes.

5. Provide the API Key you generated in the Cisco Meraki Dashboard.
6. If you have multiple organizations registered, click Get Organization List, and select the organization. If you have one registered organization, it is selected by default.

Meraki Devices: API
API based polling from Meraki Orchestrator.

API Key:

Polling uses the global HTTP proxy settings. [Configure your HTTP proxy](#)

Enable polling for Wireless devices

Enable polling for SD-WAN metrics

i Polling Meraki organizations uses the [global HTTP proxy settings](#). To change the defaults, click the [Configure your HTTP proxy](#) link.

By default both wireless information and SD-WAN metrics are polled. If you do not want to poll one of the items, clear the appropriate box. See [Monitor SD-WAN for Meraki organizations with SolarWinds Observability Self-Hosted](#).

7. Review and adjust the device properties.
 - a. Review your API key, organization, and proxy settings.
 - b. To edit how often the node status, or monitored statistics are updated, change the values in the Polling area.

Node Status Polling: seconds

Collect Statistics Every: minutes

Polling Engine: ● NPM-01 (Primary)

i For critical nodes, you may need to poll status information or collect statistics more frequently than the default polling intervals.

- c. Enter values for custom properties for the node.

The Custom Properties area is empty if you have not defined any custom properties for the monitored nodes. See "[Add custom properties to nodes](#)" in the SolarWinds Platform Administrator Guide.

- d. To adjust when the status of the node changes to Warning or Critical, edit alerting thresholds for the metric. Select the Override box and set thresholds for the node.

Alerting Thresholds

Percent Packet Loss		<input type="checkbox"/> Override Orion General Thresholds
⚠	Warning:	greater than or equal to 30 %
❗	Critical:	greater than or equal to 50 %
Response Time		<input type="checkbox"/> Override Orion General Thresholds
⚠	Warning:	greater than or equal to 500 ms
❗	Critical:	greater than or equal to 1000 ms

8. Click OK, Add Node.

The Meraki organization is now monitored as a wireless controller node. After the first poll, you can see the data from the device in the SolarWinds Platform Web Console.

On the Manage Nodes view, click the added node to see the node details in the Wireless Controller view, or drill down into thin access points listed on the view.

Monitor Meraki organizations

Click My Dashboards > Network > Wireless to see monitored Meraki organizations in the Wireless Overview and to view monitored controllers, access points, and clients.

Wireless Summary View

SHOW: Access Points SEARCH

GROUP BY: [No grouping]	<table border="1" style="width: 100%; border-collapse: collapse; text-align: center;"> <thead> <tr> <th style="width: 20%;">Access Point</th> <th style="width: 20%;">IP Address</th> <th style="width: 10%;">Type</th> <th style="width: 10%;">SSIDs</th> <th style="width: 10%;">Channels</th> <th style="width: 10%;">Clients</th> </tr> </thead> <tbody> <tr> <td> [blurred]</td> <td>[blurred]</td> <td>Thin</td> <td>N/A</td> <td></td> <td>0</td> </tr> <tr> <td> [blurred]</td> <td>[blurred]</td> <td>Thin</td> <td>N/A</td> <td></td> <td>0</td> </tr> </tbody> </table> <p style="text-align: center; font-size: small;"> ⏪ ⏩ Page 1 of 1 ⏪ ⏩ 20 </p>	Access Point	IP Address	Type	SSIDs	Channels	Clients	[blurred]	[blurred]	Thin	N/A		0	[blurred]	[blurred]	Thin	N/A		0
Access Point	IP Address	Type	SSIDs	Channels	Clients														
[blurred]	[blurred]	Thin	N/A		0														
[blurred]	[blurred]	Thin	N/A		0														

i To find access points on a controller, select Controllers in the Group by list.

Click a wireless access point to see details of the access point, such as the controller details, thin access point details, active wireless clients, and wireless clients connected in the past.

Wireless Thin AP - ● [blurred]

Controller Details EDIT HELP

IP ADDRESS	[blurred]
LAST POLL	1/18/2017 6:31:23 AM
THIN ACCESS POINTS	2
ROGUE ACCESS POINTS	0

Thin Access Point Details

IP ADDRESS	MAC	SSID	CURRENT CHANNEL
[blurred]	[blurred]	N/A	

Active Wireless Clients

No clients currently connected to wireless device.

Wireless Client History

TODAY

No clients connected to wireless device for selected period.

Click the controller name to see the wireless controller details, such as node details or list of thin access points.

Wireless Controller - [Green Status Icon]

Average Response Time & Packet Loss THRESHOLDS EDIT HELP

Avg Resp Time: 0 ms (Gauge scale: 0-2500 ms)

Packet Loss: 0% (Gauge scale: 0-100%)

Management EDIT HELP

NODE

- Edit Node
- Poll Now
- Rediscover
- Add New Alert
- SSH
- Maintenance Mode

Node Details EDIT HELP

NODE STATUS	[Green Status Icon]	Node is Up.
POLLING IP ADDRESS	[Redacted]	
DYNAMIC IP		Yes
MACHINE TYPE	[Green Status Icon]	Meraki Networks Cloud WLC
NODE CATEGORY		Other
DNS		dashboard.meraki.com

List of Thin Access Points

ACCESS POINT	IP AD
[Redacted]	[Redacted]
[Redacted]	[Redacted]

Average Response Time & P

Current Percent Utilization

STATUS	INTERFACE
[Redacted]	[Redacted]

Disk Volumes

Edit Meraki organizations

[Edit Meraki polling details, custom properties or alerting thresholds](#) in the same way as other nodes in NPM.

i When you add a Meraki organization to NPM, you can no longer change the polling method from the UI.

Monitor Arista Wireless Manager infrastructure

Add Arista WM orchestrator nodes for monitoring, monitor performance metrics, access points, and SSIDs for the wireless controller.

What does NPM monitor for Arista WM infrastructure?

Monitored entity	Metrics
Wireless Controller	<ul style="list-style-type: none"> Controller name IP address
Access Point	<ul style="list-style-type: none"> Access point name IP address SSID Channels (interface property) MAC (interface property) Radio type (interface property)
Client	<ul style="list-style-type: none"> Client name SSID IP Address MAC Signal strength Connected Data rate Data transmitted/received
Rogue Access Point	<ul style="list-style-type: none"> SSID MAC Signal strength

Requirements

- Arista WM account with an API key ID, API key value, and base URL. The API key needs to have service privileges with Wi-Fi access and WIPS enabled. See [Manage service privileges for users](#).

Limitations

- Performance metrics, such as response time or availability are polled via REST API authorization requests. As a result, the response times are longer than for other nodes, and thus using default thresholds might result in marking the node status as critical. To address this, [customize the Response Time thresholds when adding the nodes](#).

Request frequency and limits

Arista WM has a limit of 10 API calls per second. This limit is per wireless manager. As a result, all incoming calls to a wireless manager, regardless of the customer or user making the call will be throttled to allow 10 API calls per second.

API calls used for polling wireless access points

Nr.	API request	Purpose
1	/session	Login
2	/locations	Status poller
3	/manageddevices/aps	Statistics collection poller - gets APs, interfaces, and SSIDs
4	/aps	Statistics collection poller - gets rogue APs
5	/clients	Statistics collection poller - gets information on clients.

Before you begin

Prepare your Arista WM customer and client details. You will need this information when adding your Arista WM infrastructure for monitoring.

Log in to the Arista WM management console and gather the following details:

- API key ID
- API key ID value
- Base URL

See [Arista WM Getting Started](#) for details about where to find the details.

Add Arista WM wireless infrastructure to NPM

1. Log in to the SolarWinds Platform Web Console as an administrator.
2. Click Settings > Manage Nodes and click Add a Node.

3. In Polling Method, select Orchestrators: API.

i The Polling Hostname or IP Address is disabled and arista.com is used as the default host name.

4. Select Arista Wireless Manager Devices and provide the required information:

- API key ID
- API key ID value
- Base URL



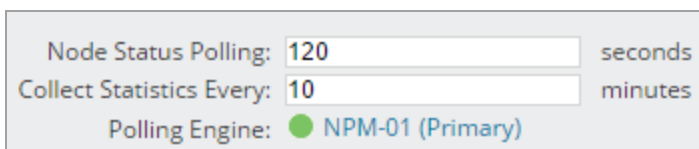
5. Click Get Locations and select the locations to monitor.

- To monitor a particular location, select the location in the dropdown.
- To monitor devices in all locations, select All Locations.

i Polling Arista WM uses the [global HTTP proxy settings](#). To change the defaults, click the Configure your HTTP proxy link.

6. Review and adjust the device properties.

- Review your Arista WM user and client details, as well as proxy settings.
- To edit how often the node status or monitored statistics are updated, change the values in the Polling area.








i By default, Node Status is polled every 120 seconds and requires two requests - an authentication request and the status request. Statistics are polled every 10 minutes and require one authentication request with three entity requests.

- Enter values for custom properties for the node.

The Custom Properties area is empty if you have not defined any custom properties for the monitored nodes. See "[Add custom properties to nodes](#)" in the SolarWinds Platform Administrator Guide.

- d. To adjust when the status of the node changes to Warning or Critical, edit alerting thresholds for the metric. Select the Override box and set thresholds for the node.

 Response times for Arista WM nodes might be 1-2 seconds. Consider adjusting the thresholds accordingly.

Alerting Thresholds	
Percent Packet Loss	<input type="checkbox"/> Override Orion General Thresholds
 Warning:	greater than or equal to 30 %
 Critical:	greater than or equal to 50 %
Response Time	<input type="checkbox"/> Override Orion General Thresholds
 Warning:	greater than or equal to 500 ms
 Critical:	greater than or equal to 1000 ms

7. Click OK, Add Node.

The Arista WM is now added as a node. After the first poll, you can see the data from the device in the SolarWinds Platform Web Console.

On the Manage Nodes view, click the added node to see the node details in the Wireless Controller view, or drill down into thin access points listed on the view. Check Node Details to verify that the Machine Type of the new device is AristaWM Cloud Orchestrator.

Monitor Aruba Central wireless infrastructure

Add Aruba Central orchestrator nodes for monitoring, monitor performance metrics, access points, and SSIDs for the wireless controller.

What does NPM monitor for Aruba Central infrastructure?

Monitored entity	Metrics
Wireless Controller	<ul style="list-style-type: none"> • Controller name • IP address • Utilization of each interface (interface state, name, and availability)

Monitored entity	Metrics
Access Point	<ul style="list-style-type: none"> • Access point name • IP address • Clients • SSID • Channels (interface property) • MAC (interface property) • Radio type (interface property)
Client	<ul style="list-style-type: none"> • Client name • SSID • IP Address • MAC • Signal Strength • Connected

i These metrics are polled via a batch request. Polling additional metrics would require additional requests and might result in exceeding the API requests limit. See [details on Aruba Central Polling request](#).

Requirements

- Aruba Central account with at least the Aruba Central View Only role permissions. See [Aruba Central User Roles...](#)

Limitations

- Performance metrics, such as response time or availability are polled via REST API authorization requests. As a result, the response times are longer than for other nodes, and thus using default thresholds might result in marking the node status as critical. To address this, [customize the Response Time thresholds when adding the nodes](#).
- **Utilization metrics are not collected** - Utilization requests are too expensive request wise and could easily drain the daily limit of 5000 requests, and thus are not collected.
- When the daily limit of requests is reached, non-status polling stops. You can see an event and an error in the Aruba Central log.
- If you have SSO enabled, create the token for accessing Aruba Central's REST API using the offline approach. Then, refresh the token API used in your automated workflows. See [Obtaining Token Using Offline Token Mechanism](#) in Aruba documentation.

Request frequency and limits

Requests	Value	Notes
Aruba Central daily requests	5,000 requests per day	Reset at midnight GMT
Status Poll request quota	free (0 requests used)	Authentication call is used. You can keep the default 120 seconds interval.
Statistics collection requests	4 requests per poll	Provided that no pagination occurs. If you have pagination set up, each page counts as a request. For example, if there are 150 rogue APs in total, with the requests configured to hold a maximum of 100 items, polling rogue APs would count as 2 requests.
Recommended statistics collection interval	10 minutes by default	Requests * polls per hour * hours per day = number of requests per day $4*6*24=576$ requests per day (provided that no pagination occurs).

API calls used for polling wireless access points

API request	Purpose
/oauth2/authorize/central/api/login	Login and status poller
/oauth2/authorize/central/api	Obtaining an authorization code
/oauth2/token	Acquiring or refreshing the access token
/monitoring/v2/aps	Statistics collection poller - gets information on APs and wireless interfaces.
/monitoring/v1/clients/wireless	Statistics collection poller - gets information on clients.
/rapids/v1/rogue_aps	Statistics collection poller - gets information on AP rogues
/monitoring/v2/bssids	Statistics collection poller - gets information on SSIDs
/configuration/v2/groups	Adding or editing node's "Get Groups"

Before you begin

Prepare your Aruba Central customer and client details. You will need this information when adding your Aruba Central infrastructure for monitoring.

Log in to the Aruba Central management console and gather the following details:

- Customer ID
- Username and password
- Client ID
- Client Secret
- Base URL

See [Aruba Central API](#) for details about where to find the details.

Add Aruba Central wireless infrastructure to NPM

1. Log in to the SolarWinds Platform Web Console as an administrator.
2. Click Settings > Manage Nodes, and click Add a Node.
3. In Polling Method, select Orchestrators: API.

i The Polling Hostname or IP Address is disabled and arubanetworks.com is used as the default host name.

4. Select Aruba Central Devices and provide the required information. Log in to the Aruba Management Console to find out the following details:
 - Customer ID
 - Username
 - Password
 - Client ID
 - Client Secret
 - Base URL

These details are necessary to load groups.

5. If you have multiple groups in Aruba Central, click Get Groups, and select a group. If you have only one group, it is selected by default.

i Polling Aruba Central uses the [global HTTP proxy settings](#). To change the defaults, click the Configure your HTTP proxy link.

6. Review and adjust the device properties.

- a. Review your Aruba Central user and client details, as well as proxy settings.
- b. To edit how often the node status, or monitored statistics are updated, change the values in the Polling area.

Node Status Polling:	<input type="text" value="120"/>	seconds
Collect Statistics Every:	<input type="text" value="10"/>	minutes
Polling Engine:	● NPM-01 (Primary)	

i Aruba Central supports 5000 requests a day. This limit is relevant for the value in Collect Statistics Every field. A single poll might require up to 4 requests, be aware that if you poll statistics more frequently than the default settings or if NPM is not the only piece of software using the requests, you might get gaps in data when the daily limit is exceeded. See [Aruba Central API Gateway](#) for details about the Aruba API calls per day limit.

- c. Enter values for custom properties for the node.

The Custom Properties area is empty if you have not defined any custom properties for the monitored nodes. See "[Add custom properties to nodes](#)" in the SolarWinds Platform Administrator Guide.

- d. To adjust when the status of the node changes to Warning or Critical, edit alerting thresholds for the metric. Select the Override box and set thresholds for the node.

💡 Response times for Aruba Central nodes might be 1-2 seconds. Consider adjusting the thresholds accordingly.

Alerting Thresholds	
Percent Packet Loss	<input type="checkbox"/> Override Orion General Thresholds
⚠ Warning:	greater than or equal to 30 %
🚨 Critical:	greater than or equal to 50 %
Response Time	<input type="checkbox"/> Override Orion General Thresholds
⚠ Warning:	greater than or equal to 500 ms
🚨 Critical:	greater than or equal to 1000 ms

7. Click OK, Add Node.

The Aruba Central is now added as a node. After the first poll, you can see the data from the device in the SolarWinds Platform Web Console.

On the Manage Nodes view, click the added node to see the node details in the Wireless Controller view, or drill down into thin access points listed on the view. Check Node Details to verify that the Machine Type of the new device is ArubaCentral Cloud Orchestrator.

Monitor Extreme Networks Cloud IQ wireless infrastructure

Add Extreme Networks Cloud IQ orchestrator nodes for monitoring wireless access points, wireless interfaces, and clients via API polling.

What does NPM monitor for Extreme Networks Cloud IQ infrastructure?

Monitored entity	Metrics
Controller	<ul style="list-style-type: none"> IP address Name
Access Point	<ul style="list-style-type: none"> Availability IP address Name Status
Interface	<ul style="list-style-type: none"> Current channel InBytes MAC Address OutBytes Radio Type SSID
Client	<ul style="list-style-type: none"> Channel IP address MAC address Name Signal strength SSID

Requirements

- Extreme Networks Cloud IQ account with at least read-only permissions.

Limitations

- Not all wireless utilization data is collected (only InBytes and OutBytes for wireless interfaces are collected and processed).
- Rogue AP data is not collected.
- Use the default polling frequency (10 minutes)
- The status poller uses API authorization requests to determine whether a node is responsive. API calls can take more time, and as a result, response time thresholds are exceeded. The nodes might appear to be exceeding the warning or critical status without showing any signs of connectivity problems (0% packet loss). Consider [adjusting the response time thresholds](#) for the nodes to better represent the average response times.
- Poll Now is not supported.
- Discovery is not supported. Use the Add Node wizard instead.

Request frequency and limits

Extreme Networks support up to 7500 API requests for each customer. See [Default Quota](#) in Extreme Networks documentation.

API calls used for polling wireless access points

API request	Purpose
1 \auth\apitoken\info	Status polling
2 \login	Authentication with username & password in order to obtain a Bearer Token.
3 \devices\<ID>\interfaces\wifi	Statistics Collection poller - gets the utilization of wireless interfaces.
4 \devices\radio-information?deviceIds=<ID>	Statistics Collection poller - gets information on AP data. Uses one request per one access point
5 \clients\active	Statistics Collection poller - gets client data.
6 \devices\radio-information	Statistics Collection poller - gets data on wireless interfaces.

Before you begin

Prepare the Extreme Networks Cloud IQ credentials - username and password for the orchestrator node.

Add Extreme Networks Cloud IQ wireless infrastructure to NPM

1. Log in to the SolarWinds Platform Web Console as an administrator.
2. Click Settings > Manage Nodes and click Add a Node.
3. In Polling Method, select Orchestrators: API.

i The Polling Hostname or IP Address is disabled and `api.extremecloudiq.com` is used as the default host name.

4. Select ExtremeCloud IQ Devices and provide the Username and Password for the orchestrator node.

i The polling uses the [global HTTP proxy settings](#). To change the defaults, click the [Configure your HTTP proxy link](#).

5. Review and adjust the device properties.
 - a. Review the user details and proxy settings.
 - b. To edit how often the node status, or monitored statistics are updated, change the values in the Polling area.

Node Status Polling:	<input type="text" value="120"/>	seconds
Collect Statistics Every:	<input type="text" value="10"/>	minutes
Polling Engine:	● NPM-01 (Primary)	

- c. Enter values for custom properties for the node.

The Custom Properties area is empty if you have not defined any custom properties for the monitored nodes. See "[Add custom properties to nodes](#)" in the SolarWinds Platform Administrator Guide.

- d.

To adjust when the status of the node changes to Warning or Critical, edit alerting thresholds for the metric. Select the Override box and set thresholds for the node.

The status poller uses API calls with authentication. This can take longer than if the default ICMP polling was used. Consider adjusting the response time thresholds for the nodes to better represent the average response times.

Node Thresholds

CPU Load Override SolarWinds Platform General Thresholds

Warning: greater than or equal to 80 %
 Critical: greater than or equal to 90 %
 Capacity Trending Calculate exhaustion using average daily values

Memory Usage Override SolarWinds Platform General Thresholds

Warning: greater than or equal to 80 %
 Critical: greater than or equal to 90 %
 Capacity Trending Calculate exhaustion using average daily values

Response Time Override SolarWinds Platform General Thresholds

Warning: Greater than ms For

Critical: Greater than ms For

6. Click OK, Add Node.

The Extreme Networks Cloud IQ device is now added as a node. After the first poll, you can see the data from the device in the SolarWinds Platform Web Console.

On the Manage Nodes view, click the added node to see the node details in the Wireless Controller view, or drill down into thin access points listed on the view. Check Node Details to verify that the Machine Type of the new device is ExtremeCloud IQ Cloud Orchestrator.

Monitor Juniper Mist wireless infrastructure

Add Juniper Mist orchestrator nodes for monitoring access points, interfaces, rogue access points, and clients for the wireless controller.

What does NPM monitor for Juniper Mist infrastructure?

Monitored entity	Metrics
Controller	<ul style="list-style-type: none"> IP address Name

Monitored entity	Metrics
Access Point	<ul style="list-style-type: none"> • Availability • InBps • InBytes • InPackets • IP address • Name • OutBytes • OutBps • OutPackets • Status
Interface	<ul style="list-style-type: none"> • Channel • InBytes • InPackets • MAC Address • OutBytes • OutPackets • Radio Type
Client	<ul style="list-style-type: none"> • InBps • InDataRate • InTotalBytes • InTotalPackets • IP address • Name • OutBps • OutDataRate • OutTotalBytes • OutTotalPackets • Signal strength • SSID
Rogue Access Point	<ul style="list-style-type: none"> • CurrentChannel • MAC Address • Signal strength • SSID

Requirements

You need your the following details for monitoring Juniper Mist infrastructure:

- Juniper Mist Hostname, such as `api.ac2.mist.com`. See [API Endpoints and Global Regions for Mist Juniper Networks](#).
- Authentication data for the used authentication method:
 - API token and Organization ID. See [Find your Organization ID for Mist Juniper Networks](#).
 - Username, password, and Organization ID. See [Find your Organization ID for Mist Juniper Networks](#).
- Site: Select from a list of detected items

Limitations

- The status poller uses API authorization requests to determine whether a node is responsive. API calls can take more time, and as a result, response time thresholds are exceeded. The nodes might appear to be exceeding the warning or critical status without showing any signs of connectivity problems (0% packet loss). Consider [adjusting the response time thresholds](#) for the nodes to better represent the average response times.
- Poll Now is not supported.
- Discovery is not supported. Add nodes using the Add Node wizard.

API calls used for polling wireless access points

	API request	Purpose
1	<code>\api\v1\self</code>	Authentication via token, status polling
2	<code>\api\v1\login</code>	Authentication via username and password.
3	<code>\api\v1\orgs\<OrgID>\sites</code>	Get Site List in the Add/Edit Node wizard.
4	<code>\api\v1\sites\<SiteID>\stats\devices\</code>	Statistics collection poller - gets information on APs and wireless interfaces.
5	<code>\api\v1\sites\<SiteID>\stats\clients</code>	Statistics collection poller - gets information on clients.
6	<code>\api\v1\sites\<SiteID>\insights\rogues</code>	Statistics collection poller - gets information on rogue APs.

Before you begin

Prepare your Juniper Mist customer and client details. See [Requirements](#).

Add Juniper Mist wireless infrastructure to NPM

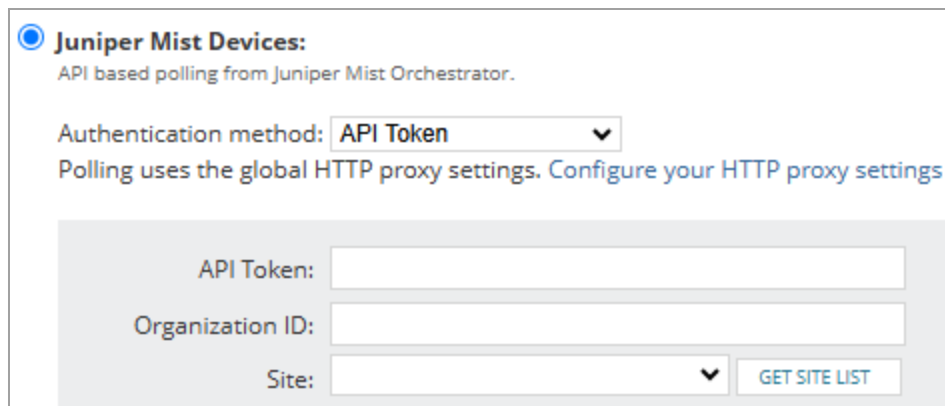
1. Log in to the SolarWinds Platform Web Console as an administrator.
2. Click Settings > Manage Nodes and click Add a Node.
3. Provide the Juniper Mist hostname. See [Requirements](#).
4. In Polling Method, select Orchestrators: API.
5. Select Juniper Mist Devices, select the authentication method, and provide the details.

API Token:

- API Token
- Organization ID

Username & Password:

- E-mail:
- Password:
- Organization ID



The screenshot shows a configuration form for Juniper Mist Devices. At the top, there is a radio button selected for "Juniper Mist Devices:" with the subtext "API based polling from Juniper Mist Orchestrator." Below this, the "Authentication method:" is set to "API Token" in a dropdown menu. A link "Configure your HTTP proxy settings" is provided, with a note that "Polling uses the global HTTP proxy settings." The form contains three input fields: "API Token:", "Organization ID:", and "Site:" (with a dropdown arrow). A "GET SITE LIST" button is located to the right of the "Site:" field.

6. Click Get Site List and select the site to monitor.

i Polling Juniper Mist devices uses the [global HTTP proxy settings](#). To change the defaults, click the Configure your HTTP proxy link.

7. Review and adjust the device properties.

- a. Review your Juniper Mist details and proxy settings.
- b. To edit how often the node status, or monitored statistics are updated, change the values in the Polling area.

Node Status Polling: seconds


Collect Statistics Every: minutes

Polling Engine: ● NPM-01 (Primary)

- c. Enter values for custom properties for the node.

The Custom Properties area is empty if you have not defined any custom properties for the monitored nodes. See "[Add custom properties to nodes](#)" in the SolarWinds Platform Administrator Guide.

To adjust when the status of the node changes to Warning or Critical, edit alerting thresholds for the metric. Select the Override box and set thresholds for the node.

 The status poller uses API calls with authentication. This can take longer than if the default ICMP polling was used. Consider adjusting the response time thresholds for the nodes to better represent the average response times.

- d.

Node Thresholds

Override SolarWinds Platform General Thresholds

CPU Load

Warning: greater than or equal to 80 %

Critical: greater than or equal to 90 %

Capacity Trending: Calculate exhaustion using average daily values

Memory Usage

Override SolarWinds Platform General Thresholds

Warning: greater than or equal to 80 %

Critical: greater than or equal to 90 %

Capacity Trending: Calculate exhaustion using average daily values

Response Time

Override SolarWinds Platform General Thresholds

Warning: Greater than ms For

Critical: Greater than ms For

8. Click OK, Add Node.

The Juniper Mist device is now added as a node. After the first poll, you can see the data from the device in the SolarWinds Platform Web Console.

On the Manage Nodes view, click the added node to see the node details in the Wireless Controller view, or drill down into thin access points listed on the view. Check Node Details to verify that the Machine Type of the new device is Juniper Mist Orchestrator.

Monitor Ruckus One wireless infrastructure

Add Ruckus One orchestrator nodes for monitoring access points with interfaces, rogue access points, and clients for the wireless controller.

What does NPM monitor for Ruckus One infrastructure?

Monitored entity	Metrics
Wireless Controller	<ul style="list-style-type: none"> Controller name IP address
Access Point	<ul style="list-style-type: none"> Access point name IP address SSID Channels (interface property) MAC (interface property) Radio type (interface property) Clients count Bandwidth utilization
Client	<ul style="list-style-type: none"> Client name SSID IP Address MAC Signal strength Connected Data transmitted/received
Rogue Access Point	<ul style="list-style-type: none"> SSID MAC

Requirements

- Ruckus One account with at least read-only permissions. See [Understanding Administrator Roles](#).

Limitations

- Performance metrics, such as response time or availability, are polled via REST API authorization requests. As a result, the response times are longer than for other nodes, and thus using default thresholds might result in marking the node status as critical. To address this, [customize the Response Time thresholds when adding the nodes](#).

Request frequency and limits

As of June 2024, the APIs used below have no rate limiting defined. See [Ruckus One APIs](#).

API calls used for polling wireless access points

API request	Purpose
1 /oauth2/token/{Tenant}	Login
2 /tenants/self	Status poller (when all venues were selected)
3 /venues/{VenueIdentifier}	Status poller (when a specific venue was selected)
4 /venues/aps	Statistics collection poller - gets information on APs.
5 /venues/aps/ {AccessPointIdentifier}/radioSettings	Statistics collection poller - gets information on wireless interfaces.
6 /venues	Statistics collection poller - gets a list of venues to get rogue APs when rogue APIs in all venues should be obtained.
7 /venues/ {VenueIdentifier}/rogueAps/query	Statistics collection poller - gets information on rogue APs.
8 /clients	Statistics collection poller - gets information on clients.

Before you begin

Prepare your Ruckus One customer and client details. You will need this information when adding your Ruckus One infrastructure for monitoring.

Log in to the Ruckus One management console and gather the following details:

- Region
- Tenant ID

- Client ID
- Client Secret

See [Ruckus One API](#) for details about where to find the details.

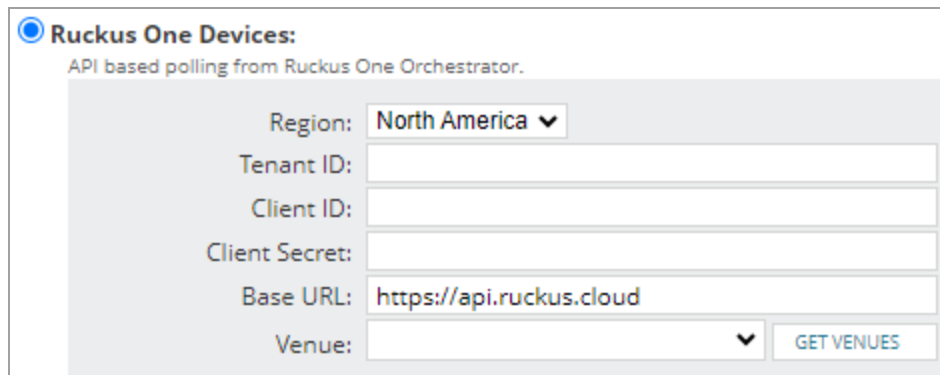
Add Ruckus One wireless infrastructure to NPM

1. Log in to the SolarWinds Platform Web Console as an administrator.
2. Click Settings > Manage Nodes and click Add a Node.
3. In Polling Method, select Orchestrators: API.

i The Polling Hostname or IP Address is disabled and Ruckus.com is used as the default host name.

4. Select Ruckus One Devices and provide the following details:
 - Region - select an option from the dropdown
 - Tenant ID
 - Client ID
 - Client Secret
 - Base URL - filled in automatically, based on the selected region.

When you fill in details, the Get Venues button becomes enabled.



5. Click Get Venues and select a venue.
 - To monitor a particular venue, select it in the drop-down.
 - To monitor devices in all venues, select All venues.

i Polling Ruckus One uses the [global HTTP proxy settings](#). To change the defaults, click the Configure your HTTP proxy link.

6. Review and adjust the device properties.

- a. Review your Ruckus One user details, client details, and proxy settings.
- b. To edit how often the node status, or monitored statistics are updated, change the values in the Polling area.

Node Status Polling:	<input type="text" value="120"/>	seconds
Collect Statistics Every:	<input type="text" value="10"/>	minutes
Polling Engine:	● NPM-01 (Primary)	

i By default, Node Status is polled every 120 seconds and requires two requests - an authentication request and the status request. Statistics are polled every 10 minutes and require one authentication request and the following requests.


- 1 request to get all APs and 1 request per AP to get radios
- 1 request to get all rogue APs for a selected venue and 1 additional request if All venues are selected
- 1 request to get all clients and 1 request per client to get additional details.

- c. Enter values for custom properties for the node.

The Custom Properties area is empty if you have not defined any custom properties for the monitored nodes. See "[Add custom properties to nodes](#)" in the SolarWinds Platform Administrator Guide.

- d. To adjust when the status of the node changes to Warning or Critical, edit alerting thresholds for the metric. Select the Override box and set thresholds for the node.

💡 Response times for Ruckus One nodes might be 1-2 seconds. Consider adjusting the thresholds accordingly.

Alerting Thresholds		
Percent Packet Loss		<input type="checkbox"/> Override Orion General Thresholds
	Warning:	greater than or equal to 30 %
	Critical:	greater than or equal to 50 %
Response Time		<input type="checkbox"/> Override Orion General Thresholds
	Warning:	greater than or equal to 500 ms
	Critical:	greater than or equal to 1000 ms


7. Click OK, Add Node.

The Ruckus One device is now added as a node. After the first poll, you can see the data from the device in the SolarWinds Platform Web Console.

On the Manage Nodes view, click the added node to see the node details in the Wireless Controller view, or drill down into thin access points listed on the view. Check Node Details to verify that the Machine Type of the new device is RuckusOne Cloud Orchestrator.

Monitor SD-WAN for Aruba EdgeConnect orchestrators (formerly Silver Peak) with SolarWinds Observability Self-Hosted

Starting with 2024.2, you can enable SD-WAN monitoring on your Aruba EdgeConnect orchestrators.

 If you are upgrading from a previous version, be aware that in Orchestrator version 9.3, all Aruba EdgeConnect Orchestrator REST API endpoint definitions have changed. See the [NPM 2024.2 release notes](#) for more information. Aruba EdgeConnect was formerly Silver Peak.

For monitored Aruba nodes, you can see orchestrator information. You can use alerts and reports relevant for SD-WAN monitoring.

Monitor SD-WAN edge devices

Add edge devices managed by a monitored Aruba EdgeConnect orchestrator to get further details via Aruba API.

- You can display general details, such as edge device name, model, serial number, status, or type.
- You can display all IP addresses for a specific Aruba SD-WAN device and properly match IPs received in NetFlow data to the node.
- You can monitor uplinks (WAN interfaces).
- You can monitor VPN tunnels - the list of tunnels, their status, and further statistics.

VPN tunnel names are created as follows:

```
{Source Interface Name} → {Peer System IP / Device ID}:{Peer Interface IP}
```

Monitor Aruba EdgeConnect orchestrators

- You can display status and response time, polled via REST (Aruba API).
- You can display the list of all devices managed by an orchestrator and add them in a simplified discovery as API-only nodes.

Monitored properties

SD-WAN metrics monitored for Aruba devices

Aruba devices are polled via API from the Aruba EdgeConnect orchestrator.

- Orchestrator Inventory and Status
- General Edge device Info
- Edge device uplinks (WAN interfaces)
- VPN tunnels

Interface polling

- Status
- Rediscovery
- Statistics

For details, see [Monitor Aruba SD-WAN interfaces](#).

Rate limits

An Aruba EdgeConnect orchestrator is limited to 250 requests/5 seconds for a polling engine. For details, see [Aruba API polling limit was exceeded](#).

Add new SD-WAN devices for monitoring in SolarWinds Observability Self-Hosted


To monitor SD-WAN, add the Aruba EdgeConnect orchestrator as a node, and then add edge devices.

 Each monitored Aruba EdgeConnect orchestrator uses a node license.

Add Aruba EdgeConnect orchestrator

1. Log in to the SolarWinds Platform Web Console as an administrator.
2. Click Settings > Manage Nodes, and then click Add a Node.
3. In Polling Method, select Orchestrators: API.
4. Under Orchestrators, select Aruba EdgeConnect Devices.
5. Select an authentication method:

- For API Token, provide the API token.
- For Username & Password, provide your credentials.

 You need an orchestrator user account with read-only access to get data using Aruba API.


6. Review and adjust the device properties.

- Review your credentials and proxy settings.
- To edit how often the node status, or monitored statistics are updated, change the values in the Polling area.

Node Status Polling: seconds

Collect Statistics Every: minutes

Polling Engine: ● NPM-01 (Primary)



 For critical nodes, you may need to poll status information or collect statistics more frequently than the default polling intervals.

- Enter values for custom properties for the node.

The Custom Properties area is empty if you have not defined any custom properties for monitored nodes. See "[Add custom properties to nodes](#)" in the SolarWinds Platform Administrator Guide.

- To adjust when the status of the node changes to Warning or Critical, edit alerting thresholds for the metric. Select the Override box and set thresholds for the node.

Alerting Thresholds

Percent Packet Loss		<input type="checkbox"/> Override Orion General Thresholds
	Warning:	greater than or equal to 30 %
	Critical:	greater than or equal to 50 %
Response Time		<input type="checkbox"/> Override Orion General Thresholds
	Warning:	greater than or equal to 500 ms
	Critical:	greater than or equal to 1000 ms

7. Click OK, Add Node.

The Aruba EdgeConnect orchestrator is now monitored as an SD-WAN orchestrator. Add connected Aruba devices to complete the SD-WAN configuration.

Add Aruba devices connected to the orchestrator as API-only nodes

When you have added the orchestrator for monitoring, you need to add edge devices you want to monitor.

1. In SolarWinds Platform Web Console, click Settings > Manage Nodes.
2. On the Manage Nodes view, click the added orchestrator node.
3. On the SD-WAN Orchestrator Details view, click Discover Aruba Devices in the Management widget.
4. Select that you want to monitor Aruba devices as API-only nodes and click Continue. Network Sonar Wizard will be launched automatically.
5. Complete the wizard by clicking Discover on the last tab. The wizard discovers the devices based on Aruba API
6. Discovered devices are listed in the Network Sonar Results wizard. Complete the wizard to add devices for monitoring.

When you finish the wizard, go to the SD-WAN Orchestrator Details page (Settings > Manage Nodes > click the orchestrator node). After the next poll, the page will display data not only for the orchestrator, but also for monitored edge devices.

Configure SD-WAN monitoring on devices already monitored with SolarWinds Platform

If you monitored SD-WAN edge devices in using SNMP, you can keep monitoring it via SNMP but you will not get any SD-WAN-related details. See [API-based monitoring for Aruba SD-WAN devices](#).

To get SD-WAN details for the orchestrator and its edge devices, start monitoring them via API:

1. Remove the node from monitoring. See [Delete devices from monitoring](#).
2. Re-add the node and select Orchestrator: API as the polling method. See [Add new SD-WAN devices for monitoring in SolarWinds Observability Self-Hosted](#).

Monitor SD-WAN for Aruba devices

When you enable SD-WAN polling for an Aruba EdgeConnect orchestrator node and click it, the SD-WAN Orchestrator Summary page opens.

By default, it includes widgets you can use to manage the device, view the device details, active alerts, latest events, or AppStack for the device.

Edge Devices

This widget lists device names, IP addresses, models, serial numbers and network IDs of edge devices paired with the orchestrator.

Orchestrator Inventory

This widget displays a list of Aruba devices connected to managed orchestrators. Devices managed by SolarWinds Observability Self-Hosted are marked in the Managed by Platform column.

To add unmanaged devices, click the Discover Aruba Devices button and [add them for monitoring](#).

SD-WAN Map

This widget is available on Orchestrator views. It displays connections between monitored devices on the network. For directly connected devices, you can also see used interfaces.

VPN Tunnels

On SD-WAN Orchestrator views, this widget displays VPN tunnels and their metrics for all edges paired with the orchestrator.

On an edge Node Details view, the widget displays VPN tunnels and their metrics for the edge.

WAN UpLinks

On SD-WAN Orchestrator views, this widget displays WAN uplinks for all edges paired with the orchestrator.

On an edge Node Details view, the widget displays WAN uplinks for the edge.

Monitor Aruba SD-WAN interfaces

Interfaces on Aruba devices are polled via API. Interfaces are discovered during the API discovery. Interface Type is always reported as "Ethernet".

Interface polling intervals use the polling intervals set for nodes. As a result, to force polling interface details, select the parent node in Node Management and click Poll Now. The Poll Now option for individual interfaces is not available.


Threshold limitations for API-polled interfaces

Error Thresholds are not supported because Error polling is not implemented.

Utilization Thresholds are supported but require the bandwidth setup. To make utilization thresholds work, set bandwidth on interfaces manually. Utilization thresholds use interface speed and polling interface speed is not implemented.

What is polled....

When?	What is polled?	What is NOT polled?
Rediscovery	<ul style="list-style-type: none"> • interface index • ifName • MTU • physical address 	<ul style="list-style-type: none"> • interface speed • duplex mode • interface name • interface alias
Status polls	<ul style="list-style-type: none"> • interface index • ifName • admin status • oper status 	<ul style="list-style-type: none"> • interface name • interface last change
Statistics	<ul style="list-style-type: none"> • interface availability • sent and received bytes • sent and received packets 	<ul style="list-style-type: none"> • other information on interface traffic • interface errors

 To get information on interface utilization, set the bandwidth of the interface manually.

Monitor SD-WAN for Fortinet FortiManager orchestrators with SolarWinds Observability Self-Hosted

Starting with NPM 2024.4, you can enable SD-WAN monitoring on your Fortinet FortiManager orchestrators.

For monitored Fortinet FortiManager nodes, you can see orchestrator information. You can use alerts and reports relevant for SD-WAN monitoring.

Monitor SD-WAN edge devices

Add edge devices managed by a monitored Fortinet FortiManager orchestrator to get further details via Fortinet FortiManager API.

- You can display general details, such as edge device name, model, serial number, status, or type.
- You can display all IP addresses for a FortiManager SD-WAN device and properly match IPs received in NetFlow data to the node.
- You can monitor uplinks (WAN interfaces).

Monitor Fortinet FortiManager orchestrators

- You can display status and response time, polled via REST (Fortinet FortiManager API).
- You can display the list of all devices managed by an orchestrator and add them in a simplified discovery as API-only nodes.

Monitored properties

SD-WAN metrics monitored for Fortinet FortiManager devices

Fortinet FortiManager devices are polled via API from the orchestrator.

- Orchestrator Inventory and Status
- General edge device Info
- Edge device uplinks (WAN interfaces)

Interface polling

- Status
- Rediscovery
- Statistics

For details, see [Monitor Fortinet FortiManager SD-WAN interfaces](#).

Rate limits

A Fortinet FortiManager orchestrator is limited to 250 requests/5 seconds for a polling engine.

Session limit

Four active sessions are supported per user for an idle timeout of 15 minutes. As a result, Fortinet FortiManager API polling may fail after a couple of session refresh.

The SolarWinds Platform needs two active sessions for regular and declarative polling. SolarWinds recommends that you do not use the account used by the SolarWinds Platform account for other purposes not to exceed the session limitation.

Requirements and recommendations for polling

- Use dedicated Fortinet credentials for SolarWinds Platform polling.
- Poll all Fortinet FortiManager devices using one polling engine.
- Do not re-run the discovery for 30 minutes after you complete adding the nodes.

Add new SD-WAN devices for monitoring in SolarWinds Observability Self-Hosted

To monitor SD-WAN, add the Fortinet FortiManager orchestrator as a node, and then add edge devices.

 Each monitored Fortinet FortiManager orchestrator uses a node license.

Add Fortinet FortiManager orchestrator

1. Log in to the SolarWinds Platform Web Console as an administrator.
2. Click Settings > Manage Nodes, and then click Add a Node.
3. In Polling Method, select Orchestrators: API.
4. Under Orchestrators, select Fortinet FortiManager Devices.
5. Enter your Fortinet FortiManager username, password, and client ID.

6. Review and adjust the device properties.

- a. Review your credentials and proxy settings.
- b. To edit how often the node status, or monitored statistics are updated, change the values in the Polling area.




Node Status Polling:	<input type="text" value="120"/>	seconds
Collect Statistics Every:	<input type="text" value="10"/>	minutes
Polling Engine:	● NPM-01 (Primary)	

i For critical nodes, you may need to poll status information or collect statistics more frequently than the default polling intervals.

- c. Enter values for custom properties for the node.

The Custom Properties area is empty if you have not defined any custom properties for monitored nodes. See "[Add custom properties to nodes](#)" in the SolarWinds Platform Administrator Guide.

- d. To adjust when the status of the node changes to Warning or Critical, edit alerting thresholds for the metric. Select the Override box and set thresholds for the node.

Alerting Thresholds		
Percent Packet Loss		<input type="checkbox"/> Override Orion General Thresholds
	Warning:	greater than or equal to 30 %
	Critical:	greater than or equal to 50 %
Response Time		<input type="checkbox"/> Override Orion General Thresholds
	Warning:	greater than or equal to 500 ms
	Critical:	greater than or equal to 1000 ms

7. Click OK, Add Node.

The Fortinet FortiManager orchestrator is now monitored as an SD-WAN orchestrator. Add connected Fortinet FortiManager devices to complete the SD-WAN configuration.

Add devices connected to the orchestrator as API-only nodes

When you have added the orchestrator for monitoring, you need to add edge devices you want to monitor.

1. In SolarWinds Platform Web Console, click Settings > Manage Nodes.
2. On the Manage Nodes view, click the added orchestrator node.

3. On the SD-WAN Orchestrator Details view, click Discover Fortinet FortiManager Devices in the Management widget.
4. Select that you want to monitor Fortinet FortiManager devices as API-only nodes and click Continue. Network Sonar Wizard will be launched automatically.
5. Complete the wizard by clicking Discover on the last tab. The wizard discovers the devices based on the Fortinet FortiManager API.
6. Discovered devices are listed in the Network Sonar Results wizard. Complete the wizard to add devices for monitoring.

When you finish the wizard, go to the SD-WAN Orchestrator Details page (Settings > Manage Nodes > click the orchestrator node). After the next poll, the page will display data not only for the orchestrator, but also for monitored edge devices.

Configure SD-WAN monitoring on devices already monitored with SolarWinds Platform

If you monitored SD-WAN edge devices using SNMP, you can keep monitoring it via SNMP but you will not get any SD-WAN-related details.

To get SD-WAN details for the orchestrator and its edge devices, start monitoring them via API:

1. Remove the node from monitoring. See [Delete devices from monitoring](#).
2. Re-add the node and select Orchestrator: API as the polling method. See [Add new SD-WAN devices for monitoring in SolarWinds Observability Self-Hosted](#).

Monitor SD-WAN for Fortinet FortiManager devices

When you enable SD-WAN polling for a Fortinet FortiManager orchestrator node and click it, the SD-WAN Orchestrator Summary page opens.

By default, it includes widgets you can use to manage the device, view the device details, active alerts, latest events, or AppStack for the device.

Edge Devices

This widget lists device names, IP addresses, models, serial numbers and network IDs of edge devices paired with the orchestrator.

Orchestrator Inventory

This widget displays a list of Fortinet FortiManager devices connected to managed orchestrators. Devices managed by SolarWinds Observability Self-Hosted are marked in the Managed by Platform column.

To add unmanaged devices, click the Discover Fortinet FortiManager Devices button and [add them for monitoring](#).

SD-WAN Map

This widget is available on Orchestrator views. It displays connections between monitored devices on the network. For directly connected devices, you can also see used interfaces.

WAN UpLinks

On SD-WAN Orchestrator views, this widget displays WAN uplinks for all edges paired with the orchestrator.

On an edge Node Details view, the widget displays WAN uplinks for the edge.

Monitor Fortinet FortiManager SD-WAN interfaces

Interfaces on Fortinet FortiManager devices are polled via API. Interfaces are discovered during the API discovery. Interface Type is always reported as "Ethernet".

Interface polling intervals use the polling intervals set for nodes. As a result, to force polling interface details, select the parent node in Node Management and click Poll Now. The Poll Now option for individual interfaces is not available.

Threshold limitations for API-polled interfaces

Utilization Thresholds are supported but require the bandwidth setup. To make utilization thresholds work, set bandwidth on interfaces manually.

What is polled....

When?	What is polled?	What is NOT polled?
Rediscovery	<ul style="list-style-type: none"> • interface index • ifName • MTU • physical address • interface speed • interface name • duplex mode • interface alias 	
Status	<ul style="list-style-type: none"> • interface index • ifName • interface name • admin status • oper status 	<ul style="list-style-type: none"> • interface last change
Statistics	<ul style="list-style-type: none"> • interface availability • sent and received bytes • sent and received packets • interface errors 	<ul style="list-style-type: none"> • other information on interface traffic

i To get information on interface utilization, set the bandwidth of the interface manually.

Monitor SD-WAN for Meraki organizations with SolarWinds Observability Self-Hosted

When you enable SD-WAN monitoring on your Meraki orchestrators (dashboards), you can see both WLC and cloud orchestrator information for monitored Meraki nodes. You can use alerts and reports relevant for SD-WAN monitoring.

Monitor Meraki nodes as Wireless Controllers

- You can display access point statuses polled via API.
- Clients are polled for the entire network using one API request instead of multiple requests for each device.
- The API request includes SSID information.

Monitor SD-WAN Edge devices

Add edge devices managed by a monitored Meraki cloud orchestrator to get further details via Meraki API.

- You can display general details, such as edge device name, model, serial number, status, or type.
- You can display all IP addresses for a specific Meraki SD-WAN device and properly match IPs received in NetFlow data to the node.
- You can display CDP and LLDP data from Meraki SD-WAN devices and build connections between them.
- You can display VLANs and Port Maps data from Meraki SD-WAN devices.
- You can display High Availability details in the Warm Spare widget, such as HA role, or HA mode.
- You can monitor uplinks (WAN interfaces), VPN tunnels, VPN tunnel metrics, and performance score.

Monitor cloud orchestrators

- You can display status and response time, polled via REST (Meraki API).
- You can display the list of all devices managed by a cloud orchestrator and add these devices in a simplified discovery.

Monitor devices protected by Warm Spare

You can monitor SD-WAN edge devices protected by Meraki Warm Spare to be highly available.



If you are using Warm Spare, do not use shared LAN ports and make sure each device from the Warm Spare has a dedicated IP address for SNMP polling.

Monitor VPN tunnels for Meraki SD-WAN devices

VPN tunnel names are created as follows:

```
{Source Interface Name} → {Peer Network Name}:{Peer Interface Name}
```

Review the supported Meraki series

Supported devices must have SNMP SysObjectIDs from the following range:

```
1.3.6.1.4.1.29671.2.100 - 1.3.6.1.4.1.29671.2.200
```

Only the following series of Meraki devices are supported for SD-WAN monitoring.

- Meraki MX series
- VMX100
- Meraki Z-Series (Teleworker): Z3, Z3C

Add new SD-WAN devices for monitoring in SolarWinds Observability Self-Hosted

To monitor SD-WAN, add the Meraki orchestrator to the SolarWinds Platform database as an external node, and then add edge devices.

 Each monitored Meraki orchestrator uses a node license.

Add Meraki orchestrator

1. Log in to the SolarWinds Platform Web Console as an administrator.
2. Click Settings > Manage Nodes, and click Add a Node.
3. In Polling Method, select Orchestrators: API.
4. Under Orchestrators, select Meraki Devices.

i The Polling Hostname or IP Address is disabled and `api.meraki.com` is used as the default host name for Meraki networks. If your Meraki organization has a different domain suffix, such as `api.meraki.cn` or `api.meraki.ca`, go to Advanced configuration, and change the MerakiCloudName.

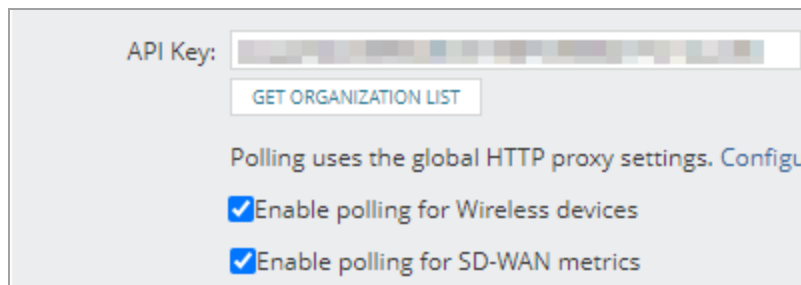
- a. Log in to Advanced configuration. See [Access the Advanced Configuration](#).
- b. Search for MerakiCloudName on the Global tab.
- c. Change the MerakiCloudName to your organization and save your changes.

5. Provide the API Key you generated in the Cisco Meraki Dashboard.

6. Make sure that all boxes are selected:

- Enable Topology polling... - this option
- Enable polling for Wireless devices - this option ensures the basic Meraki monitoring, as described in [Monitor Meraki wireless infrastructure in NPM](#).
- Enable polling for SD-WAN metrics - this option adds SD-WAN metrics.

7. If you have multiple organizations registered, click Get Organization List, and select the organization. If you have one registered organization, it is selected by default.



API Key:

Polling uses the global HTTP proxy settings. [Configure](#)

Enable polling for Wireless devices


Enable polling for SD-WAN metrics

i Polling Meraki dashboards uses the [global HTTP proxy settings](#). To change the defaults, click the Configure proxy link.

8. Review and adjust the device properties.

- a. Review your API key, organization, and proxy settings.
- b. To edit how often the node status, or monitored statistics are updated, change the values in the Polling area.

Node Status Polling:	<input type="text" value="120"/>	seconds
Collect Statistics Every:	<input type="text" value="10"/>	minutes
Polling Engine:	● NPM-01 (Primary)	

 For critical nodes, you may need to poll status information or collect statistics more frequently than the default polling intervals.

- c. Enter values for custom properties for the node.

The Custom Properties area is empty if you have not defined any custom properties for the monitored nodes. See "[Add custom properties to nodes](#)" in the SolarWinds Platform Administrator Guide.

- d. To adjust when the status of the node changes to Warning or Critical, edit alerting thresholds for the metric. Select the Override box and set thresholds for the node.

Alerting Thresholds		
Percent Packet Loss		<input type="checkbox"/> Override Orion General Thresholds
	Warning:	greater than or equal to 30 %
	Critical:	greater than or equal to 50 %
Response Time		<input type="checkbox"/> Override Orion General Thresholds
	Warning:	greater than or equal to 500 ms
	Critical:	greater than or equal to 1000 ms

9. Click OK, Add Node.

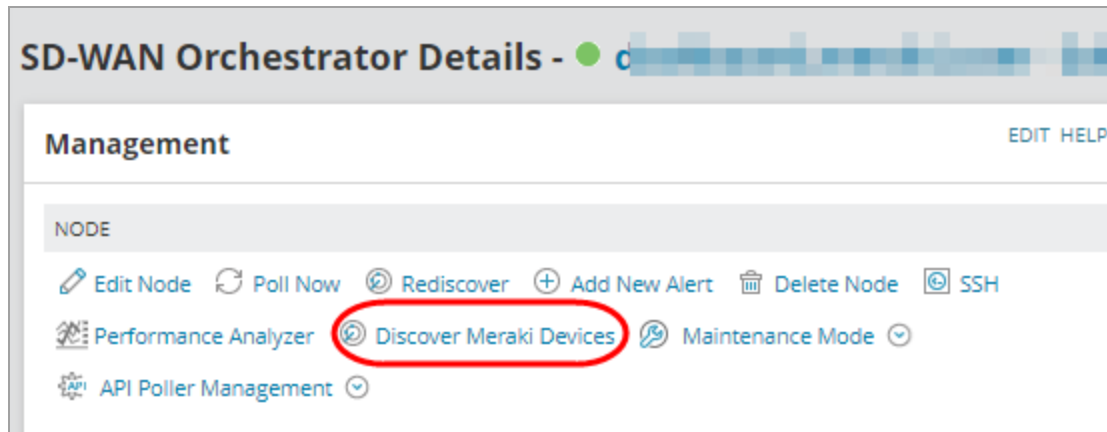
The Meraki orchestrator is now monitored as a wireless controller node and as a SD-WAN orchestrator. After the first poll, you can see the wireless data from the device in the SolarWinds Platform Web Console. Add connected Meraki devices to complete the SD-WAN configuration.

Add Meraki devices connected to the orchestrator

When you have added the Meraki orchestrator for monitoring, you need to add edge devices you want to monitor.

1. In SolarWinds Platform Web Console, click Settings > Manage Nodes.
2. On the Manage Nodes view, click the added Meraki orchestrator node.

3. On the SD-WAN Orchestrator Details view, click Discover Meraki Devices in the Management widget.

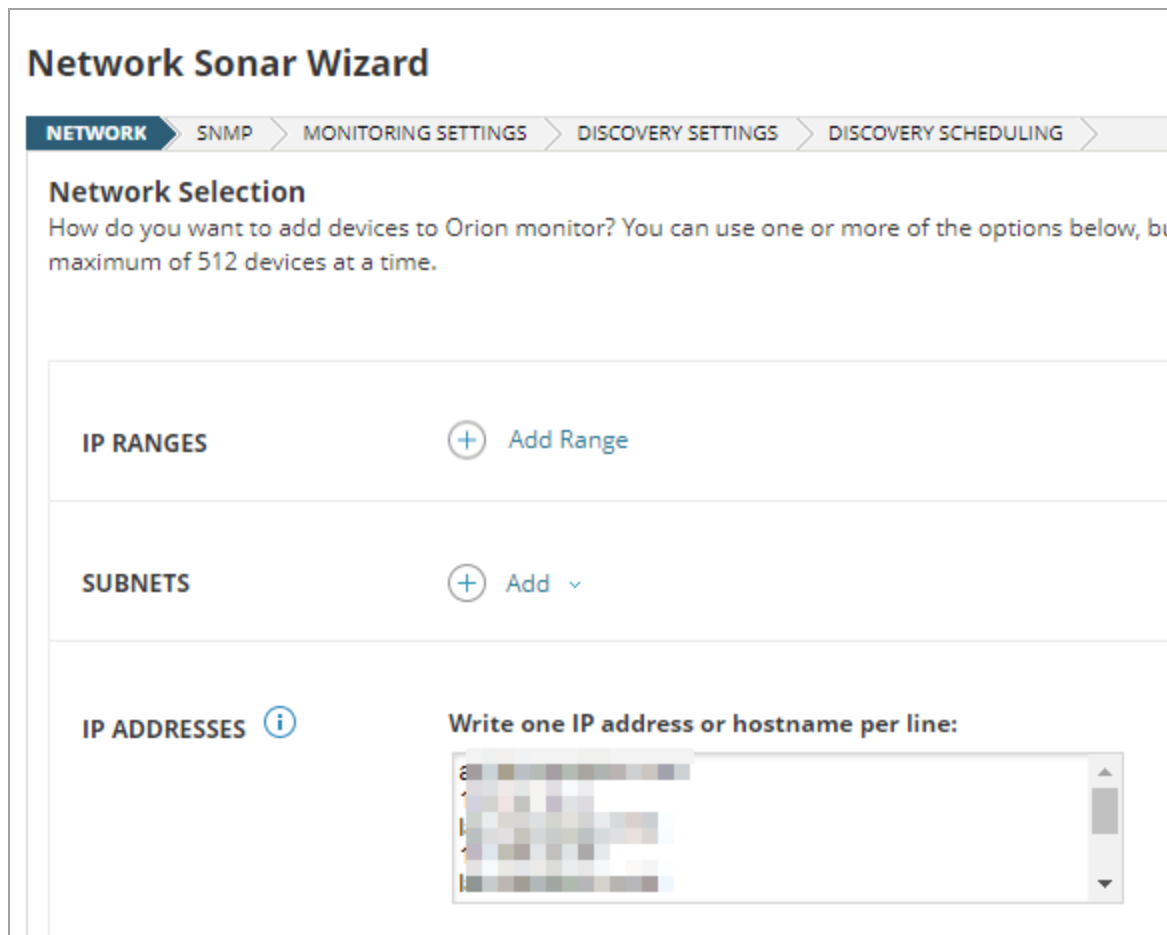


How do I add SD-WAN edge devices for monitoring?

When you click Discover Meraki Devices button on the SD-WAN Orchestrator Details page, SolarWinds Observability Self-Hosted requests Meraki device hostnames and IP addresses from the Meraki Orchestrator.

A simplified version of Network Sonar Wizard opens and the details from the orchestrator are automatically listed in the IP Addresses box.

4. In Network Sonar Discovery, review the hostnames and IP addresses received from the orchestrator. Include only the devices you want to monitor and click Next.



5. On SNMP, make sure [SNMP credentials](#) for the devices are listed and click Next. If appropriate credentials are not listed, click Add New Credential, and define a new set.
6. Complete the wizard by clicking Discover on the last tab. The wizard searches your network for the hostnames/IP addresses.
7. Discovered devices are listed in the Network Sonar Results wizard. Complete the wizard to add the devices for monitoring.

When you finish the wizard, go to the SD-WAN Orchestrator Details page (Settings > Manage Nodes > click the orchestrator node). After the next poll, the page will display data not only for the orchestrator, but also for monitored edge devices.

Configure SD-WAN monitoring on devices already monitored with SolarWinds Platform

If you monitored SD-WAN nodes in a previous version, using either SNMP or Meraki API and upgraded to SolarWinds Observability Self-Hosted, enable SD-WAN polling.

1. In SolarWinds Platform Web Console, edit the node properties.
2. Click Settings > Manage Nodes, select the node and click Edit properties.
3. On the Node Detail page, click Edit Node in the Management widget.
4. On Edit Properties, make sure Meraki Devices is selected as the polling method. See [Add Meraki Orchestrator](#).
5. Select the Enable polling for SD-WAN metrics box and submit your changes. When you go to the Node Details view, the SD-WAN Orchestrator Details page opens by default. The Wireless Summary Details view for the device is available as a subview on the left-hand side of the view.
6. [Add SD-WAN edge devices](#).

After the first poll, you can see SD-WAN data in the widgets.

Monitor SD-WAN for Meraki devices

When you enable SD-WAN polling for Meraki devices, and click the node, the SD-WAN Orchestrator Details page opens.

By default, it includes widgets you can use to manage the device, view the device details, active alerts, latest events, or AppStack for the device.

The following devices display SD-WAN related details.

Edge Devices

This widget lists device names, IP addresses, models, serial numbers and network IDs.

Edge Device Performance and Edge Device Status

These charts display performance and status of individual edge devices in time.

[Network Topology](#)

This widget displays existing physical connections between monitored devices.

Orchestrator Inventory

This widget displays a list of Meraki devices connected to managed orchestrators. Devices managed by SolarWinds Observability Self-Hosted are marked in the Managed by Platform column..

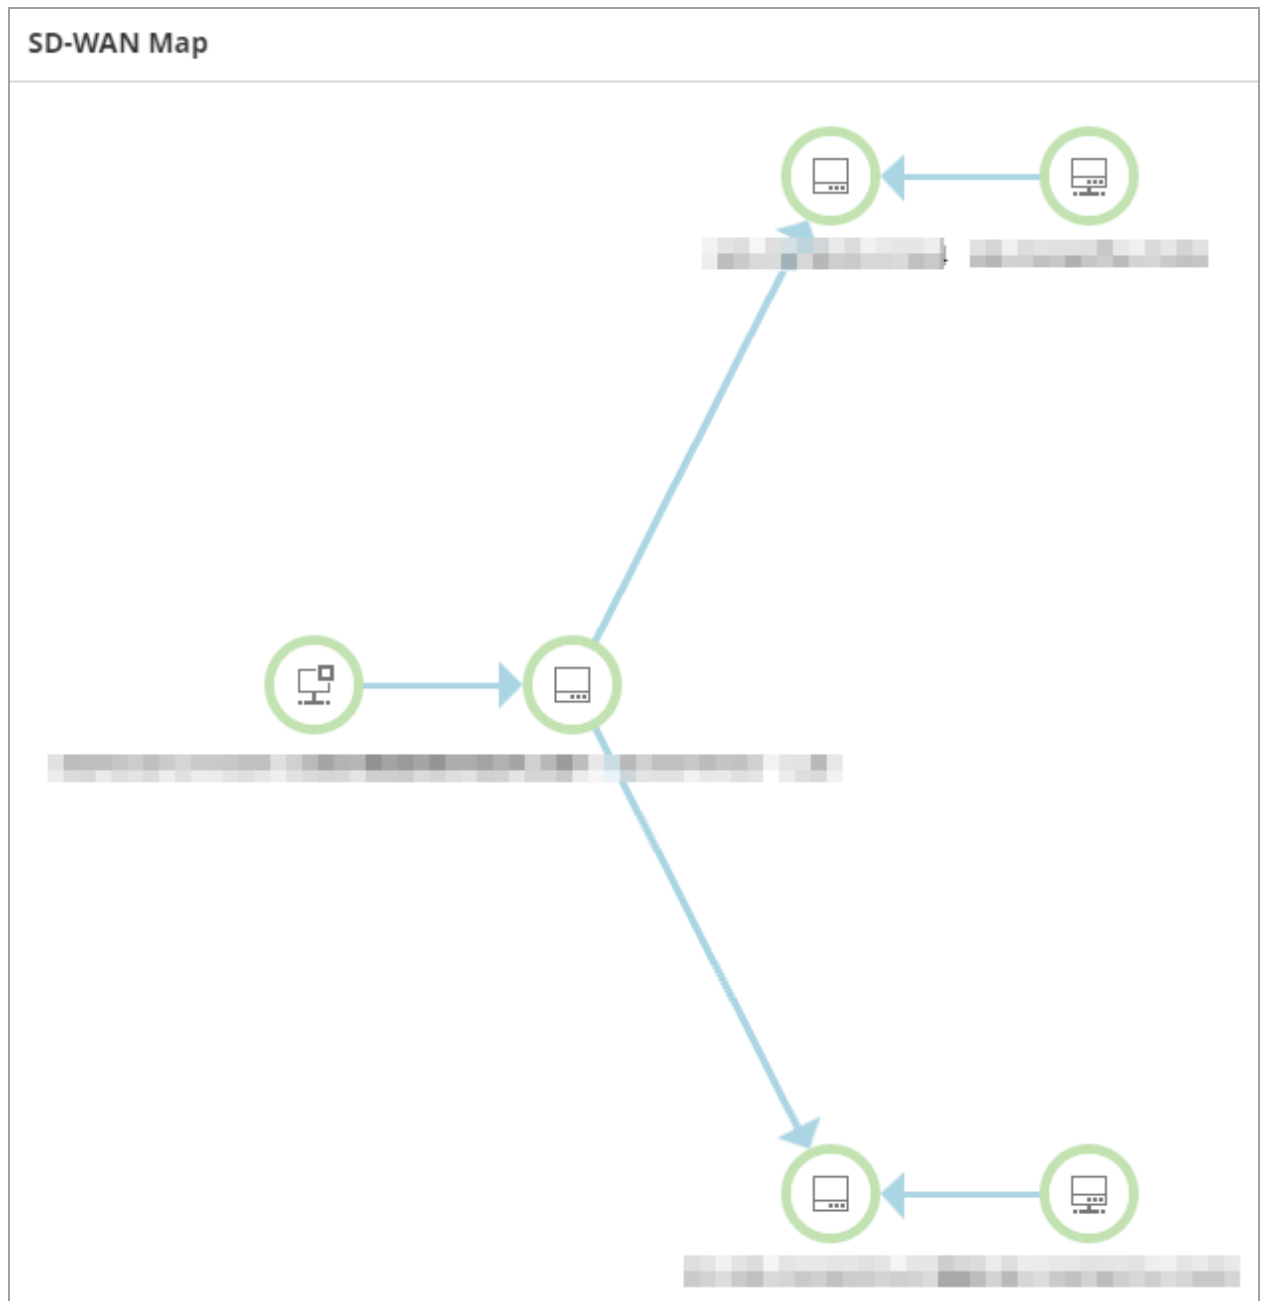
To add unmanaged devices, click the Discover Meraki Devices button and [add them for monitoring](#).

Orchestrator Inventory

DEVICE NAME	STATUS	IP ADDRESS	MODEL	PRODUCT TYPE	SERIAL NO	NETWORK ID	MANAGED BY ORCHESTRATOR
[blurred]	Online	[blurred]	VMX-M	appliance	[blurred]	[blurred]	
[blurred]	Online	[blurred]	MR52	wireless	[blurred]	[blurred]	
[blurred]	Online	[blurred]	MX84	appliance	[blurred]	[blurred]	
[blurred]	Online	[blurred]	MS220-48FP	switch	[blurred]	[blurred]	
[blurred]	Online	[blurred]	MR18	wireless	[blurred]	[blurred]	
[blurred]	Online	[blurred]	MR34	wireless	[blurred]	[blurred]	
[blurred]	Online	[blurred]	MX64	appliance	[blurred]	[blurred]	
[blurred]	Online	[blurred]	MS220-8P	switch	[blurred]	[blurred]	

SD-WAN Map

This widget displays a map of physical and logical connections for the orchestrator. See [Intelligent Maps](#).



Top Tunnels

On SD-WAN Orchestrator views, the widget displays the metric data (jitter, latency, or packet loss) for top VPN tunnels based on the metric. The widget shows top tunnels on all edges monitored by the orchestrator.

On an edge Node Details view, the widget displays the metric data (jitter, latency, or packet loss) for VPN tunnels on the node.


By default, the widgets display a maximum of 10 tunnels. To change the number of tunnels displayed, complete the following steps:

1. Using an account with administrative privileges, click Settings > All Settings and then click Web Console Settings.
2. In Classic Chart Settings, provide the appropriate number of tunnels to be displayed as the value for Maximum number of data series displayed on chart and submit your changes.

VPN Connections

On SD-WAN Orchestrator views, this widget displays VPN tunnels for all edges paired with the orchestrator.

On an edge Node Details view, the widget displays VPN tunnels for the edge.

 This widget is not displayed by default. See [Add widgets to SolarWinds Platform views](#).

VPN Tunnels

This widget displays VPN tunnels configured in the organization networks.

WAN UpLinks

This widget displays a list of uplinks, showing status of individual uplinks, interface names, local and public IP Addresses.

Warm Spare

This widget displays details about devices in the High Availability mode.

Monitor SD-WAN for Prisma orchestrators with SolarWinds Observability Self-Hosted

Starting with 2024.1, you can enable SD-WAN monitoring on your Prisma (formerly CloudGenix) orchestrators.

For monitored Prisma nodes, you can see orchestrator information. You can use alerts and reports relevant for SD-WAN monitoring.

Monitor SD-WAN edge devices

Add edge devices managed by a monitored Prisma orchestrator to get further details via Prisma API.

- You can display general details, such as edge device name, model, serial number, status, or type.
- You can display all IP addresses for a specific Prisma SD-WAN device and properly match IPs received in NetFlow data to the node.
- You can monitor uplinks (WAN interfaces).
- You can monitor VPN tunnels - the list of tunnels and their status. No further statistics are monitored for tunnels.

VPN tunnel names are created as follows:

```
{Source Interface Name} → {Peer SysName}:{Peer Interface Name}
```

Monitor Prisma orchestrators

- You can display status and response time, polled via REST (Prisma API).
- You can display the list of all devices managed by an orchestrator and add them in a simplified discovery as SNMP or API-only nodes.

Review the requirements for Prisma devices

Supported edge devices must have one of the following SNMP SysObjectID:

1.3.6.1.4.1.50114.11.1.10 or 1.3.6.1.4.1.50114.11.1.11

For Prisma devices, the following SD-WAN metrics are monitored:

- Orchestrator Inventory and Status
- General Edge device Info
- Edge device uplinks (WAN interfaces)
- VPN tunnels

Additional metrics for API-only nodes

- Details
- Status

SNMP-polled metrics for SNMP nodes

- Details
- ICMP and SNMP status and response time
- SNMP uptime

SNMP-polled metrics for interfaces of SNMP nodes

- Details
- Status / Availability
- Errors & Discards
- Traffic
- Percent Utilization (regular and real time)


See [SNMP or API-based monitoring for Prisma SD-WAN devices](#) for more details.

Rate limits

A Prisma Orchestrator is limited to 2,000 requests/minute. By default, SolarWinds Observability Self-Hosted is limited to 20 requests/second. For details, see [Prisma API polling limit was exceeded](#).

Add new SD-WAN devices for monitoring in SolarWinds Observability Self-Hosted

To monitor SD-WAN, add the Prisma orchestrator as a node, and then add edge devices.

 Each monitored Prisma orchestrator uses a node license.

Add Prisma orchestrator

1. Log in to the SolarWinds Platform Web Console as an administrator.
2. Click Settings > Manage Nodes, and then click Add a Node.
3. In Polling Method, select Orchestrators: API.
4. Under Orchestrators, select Prisma Devices.

5. Type your Prisma credentials into Client ID, Client Secret, and TSG ID fields and test the credentials.

i These details are generated when you create a service account for your Prisma tenant. See [Add a Service Account...](#) in PaloAlto documentation.

6. Review and adjust the device properties.
 - a. Review your credentials and proxy settings.
 - b. To edit how often the node status, or monitored statistics are updated, change the values in the Polling area.

Node Status Polling:	<input style="width: 100%;" type="text" value="120"/>	seconds
Collect Statistics Every:	<input style="width: 100%;" type="text" value="10"/>	minutes
Polling Engine:	● NPM-01 (Primary)	

i For critical nodes, you may need to poll status information or collect statistics more frequently than the default polling intervals.

- c. Enter values for custom properties for the node.

The Custom Properties area is empty if you have not defined any custom properties for monitored nodes. See "[Add custom properties to nodes](#)" in the SolarWinds Platform Administrator Guide.

- d. To adjust when the status of the node changes to Warning or Critical, edit alerting thresholds for the metric. Select the Override box and set thresholds for the node.

Alerting Thresholds	
Percent Packet Loss	<input type="checkbox"/> Override Orion General Thresholds
⚠ Warning: greater than or equal to 30 %	
! Critical: greater than or equal to 50 %	
Response Time	<input type="checkbox"/> Override Orion General Thresholds
⚠ Warning: greater than or equal to 500 ms	
! Critical: greater than or equal to 1000 ms	

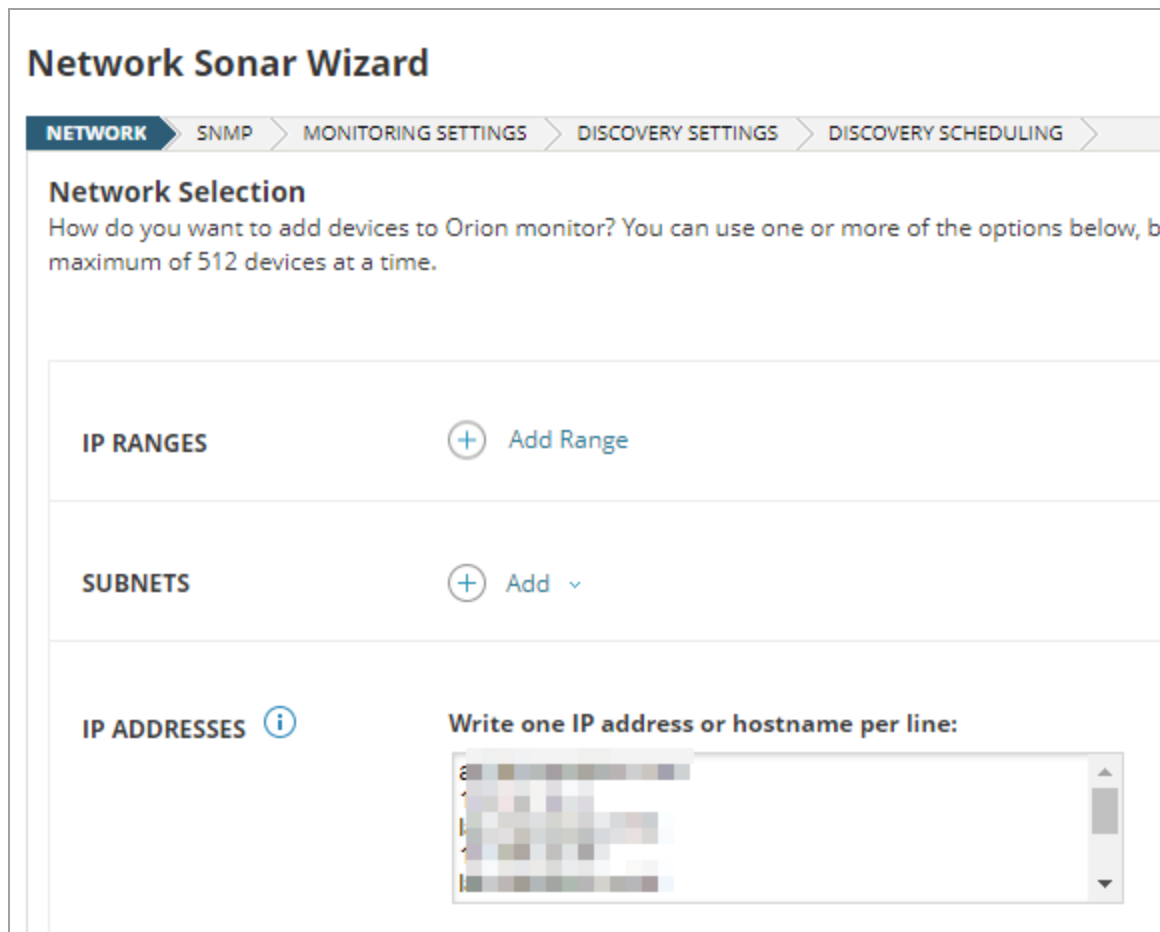
7. Click OK, Add Node.

The Prisma orchestrator is now monitored as an SD-WAN orchestrator. Add connected Prisma devices to complete the SD-WAN configuration.

Add Prisma devices connected to the orchestrator as SNMP nodes

When you have added the orchestrator for monitoring, you need to add edge devices you want to monitor.

1. In SolarWinds Platform Web Console, click Settings > Manage Nodes.
2. On the Manage Nodes view, click the added orchestrator node.
3. On the SD-WAN Orchestrator Details view, click Discover Prisma Devices in the Management widget.
4. Select that you want to monitor Prisma devices as SNMP nodes and click Continue. Network Sonar Wizard will be launched automatically.
5. In Network Sonar Discovery, review the hostnames and IP addresses received from the orchestrator. Include only the devices you want to monitor and click Next.



6. On SNMP, make sure [SNMP credentials](#) for the devices are listed and click Next. If appropriate credentials are not listed, click Add New Credential, and define a new set.
7. Complete the wizard by clicking Discover on the last tab. The wizard searches your network for the hostnames/IP addresses.
8. Discovered devices are listed in the Network Sonar Results wizard. Complete the wizard to add the devices for monitoring.

When you finish the wizard, go to the SD-WAN Orchestrator Details page (Settings > Manage Nodes > click the orchestrator node). After the next poll, the page will display data not only for the orchestrator, but also for monitored edge devices.

Add Prisma devices connected to the orchestrator as API-only nodes

When you have added the orchestrator for monitoring, you need to add edge devices you want to monitor.

1. In SolarWinds Platform Web Console, click Settings > Manage Nodes.
2. On the Manage Nodes view, click the added orchestrator node.
3. On the SD-WAN Orchestrator Details view, click Discover Prisma Devices in the Management widget.
4. Select that you want to monitor Prisma devices as API-only nodes and click Continue. Network Sonar Wizard will be launched automatically.
5. Complete the wizard by clicking Discover on the last tab. The wizard discovers the devices based on Prisma API
6. Discovered devices are listed in the Network Sonar Results wizard. Complete the wizard to add devices for monitoring.

When you finish the wizard, go to the SD-WAN Orchestrator Details page (Settings > Manage Nodes > click the orchestrator node). After the next poll, the page will display data not only for the orchestrator, but also for monitored edge devices.

Configure SD-WAN monitoring on edge devices already monitored with SolarWinds Platform

If you monitored SD-WAN edge devices in a previous version using SNMP and upgraded to SolarWinds Observability Self-Hosted, [add the Prisma orchestrator](#) to automatically pair them with the orchestrator.

Monitor SD-WAN for Prisma devices

When you enable SD-WAN polling for a Prisma orchestrator node and click it, the SD-WAN Orchestrator Summary page opens.

By default, it includes widgets you can use to manage the device, view the device details, active alerts, latest events, or AppStack for the device.

Edge Devices

This widget lists device names, IP addresses, models, serial numbers and network IDs of edge devices paired with the orchestrator.

Orchestrator Inventory

This widget displays a list of Prisma devices connected to managed orchestrators. Devices managed by SolarWinds Observability Self-Hosted are marked in the Managed by Platform column.

To add unmanaged devices, click the Discover Prisma Devices button and [add them for monitoring](#).


SD-WAN Map

This widget is available on Orchestrator views. It displays connections between monitored devices on the network. For directly connected devices, you can also see used interfaces.

VPN Connections

On SD-WAN Orchestrator views, this widget displays VPN tunnels for all edges paired with the orchestrator.

On an edge Node Details view, the widget displays VPN tunnels for the edge.

 This widget is not displayed by default. See [Add widgets to SolarWinds Platform views](#).

VPN Tunnels

On SD-WAN Orchestrator views, this widget displays VPN tunnels and their metrics for all edges paired with the orchestrator.

On an edge Node Details view, the widget displays VPN tunnels and their metrics for the edge.

WAN UpLinks

On SD-WAN Orchestrator views, this widget displays WAN uplinks for all edges paired with the orchestrator.

On an edge Node Details view, the widget displays WAN uplinks for the edge.

Monitor SD-WAN for VeloCloud orchestrators with SolarWinds Observability Self-Hosted

Starting with 2022.4, you can enable SD-WAN monitoring on your VeloCloud orchestrators.

For monitored VeloCloud nodes, you can see orchestrator information. You can use alerts and reports relevant for SD-WAN monitoring.

Monitor SD-WAN edge devices

Add edge devices managed by a monitored VeloCloud orchestrator to get further details via VeloCloud API.

- You can display general details, such as edge device name, model, serial number, status, or type.
- You can display all IP addresses for a specific VeloCloud SD-WAN device and properly match IPs received in NetFlow data to the node.
- You can monitor uplinks (WAN interfaces).

Monitor VeloCloud orchestrators

- You can display status and response time, polled via REST (VeloCloud API).
- You can display the list of all devices managed by an orchestrator and add these devices in a simplified discovery.

Monitor VPN tunnels for VeloCloud SD-WAN devices

VPN tunnel names are created as follows:

```
{Source Interface Name} → {Peer Device Name}:{Peer Link Name}
```

Review the requirements for VeloCloud devices

Supported edge devices must have the following SNMP SysObjectID: 1.3.6.1.4.1.45346

For VeloCloud devices, the following SD-WAN metrics are monitored:

- Orchestrator Inventory and Status
- General Edge device Info
- Edge device Uplinks (WAN interfaces)
- VPN Tunnels
- VLANs
- High Availability (Warm Spare widget)
- VPN tunnel metrics

SNMP-polled metrics for nodes

- Details via Device Studio
- ICMP and SNMP Status and Response time
- SNMP uptime
- CPU & Memory (regular and real time)
- Load average

SNMP-polled metrics for interfaces

- Details
- Status / Availability
- Errors & Discards
- Traffic
- Percent Utilization (regular and real time)

Rate limits

A VeloCloud Orchestrator is limited to 500 requests every five seconds per each type of user (Operator, Partner, and Customer). SolarWinds Observability Self-Hosted doesn't send more requests than this limit but if other requests are sent for the same organization, the limit might be exceeded. For details, see [VeloCloud API polling limit was exceeded](#).


Add new SD-WAN devices for monitoring in SolarWinds Observability Self-Hosted

To monitor SD-WAN, add the VeloCloud orchestrator to the SolarWinds Platform database as a node, and then add edge devices.

 Each monitored VeloCloud orchestrator uses a node license.

Add VeloCloud orchestrator

1. Log in to the SolarWinds Platform Web Console as an administrator.
2. Click Settings > Manage Nodes, and then click Add a Node.
3. Provide the VeloCloud orchestrator hostname.

 Enter a hostname. IP addresses are not supported at the moment.

4. In Polling Method, select Orchestrators: API.
5. Under Orchestrators, select VeloCloud Devices.

6. Select an authentication method:

- For API Token, provide the API token.
- For Username & Password, provide your credentials.

7. Review and adjust the device properties.

- Review your credentials and proxy settings.
- To edit how often the node status, or monitored statistics are updated, change the values in the Polling area.



Node Status Polling:	<input type="text" value="120"/>	seconds
Collect Statistics Every:	<input type="text" value="10"/>	minutes
Polling Engine:	● NPM-01 (Primary)	

i For critical nodes, you may need to poll status information or collect statistics more frequently than the default polling intervals.

- Enter values for custom properties for the node.

The Custom Properties area is empty if you have not defined any custom properties for monitored nodes. See ["Add custom properties to nodes"](#) in the SolarWinds Platform Administrator Guide.

- To adjust when the status of the node changes to Warning or Critical, edit alerting thresholds for the metric. Select the Override box and set thresholds for the node.

Alerting Thresholds		
Percent Packet Loss		<input type="checkbox"/> Override Orion General Thresholds
	Warning:	greater than or equal to 30 %
	Critical:	greater than or equal to 50 %
Response Time		<input type="checkbox"/> Override Orion General Thresholds
	Warning:	greater than or equal to 500 ms
	Critical:	greater than or equal to 1000 ms

8. Click OK, Add Node.

The VeloCloud orchestrator is now monitored as an SD-WAN orchestrator. Add connected VeloCloud devices to complete the SD-WAN configuration.

Add VeloCloud devices connected to the orchestrator

When you have added the orchestrator for monitoring, you need to add edge devices you want to monitor.

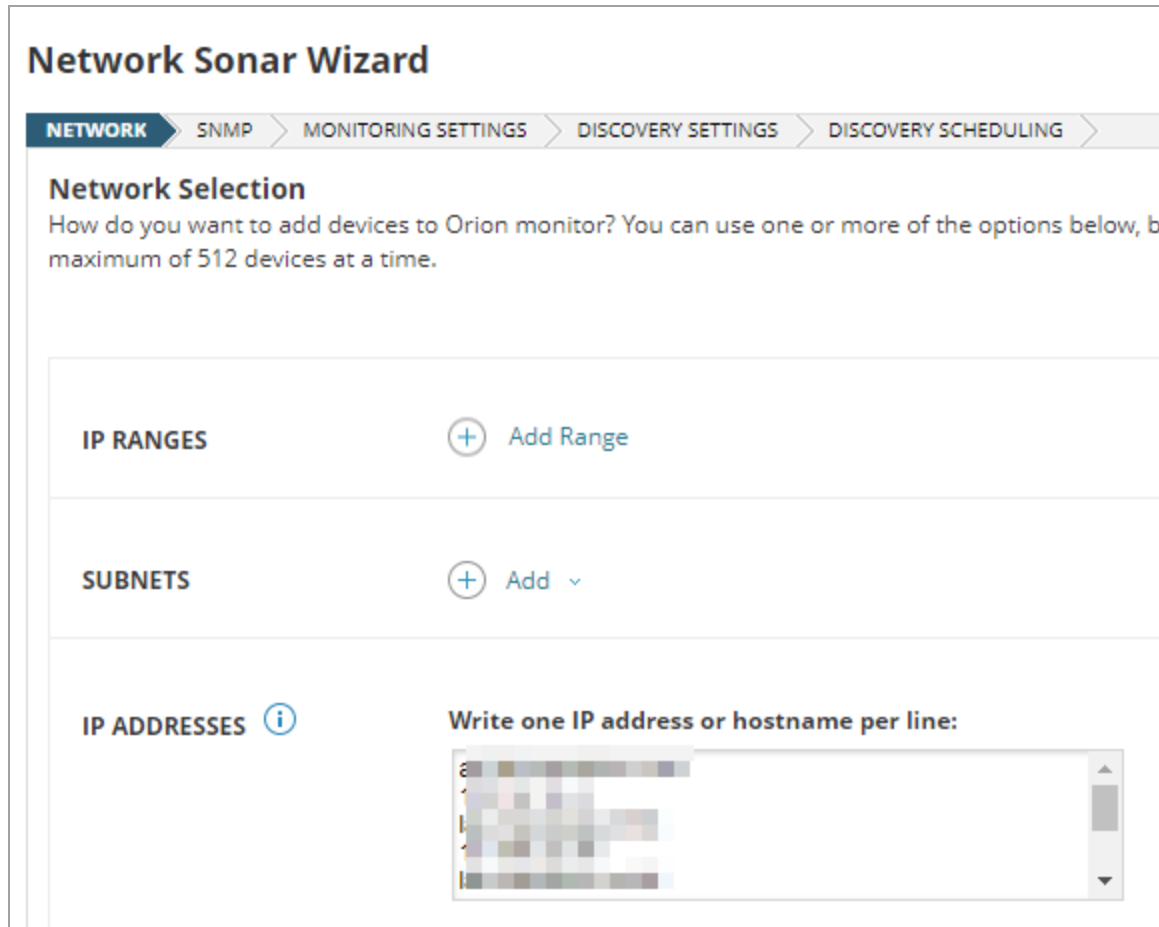
1. In SolarWinds Platform Web Console, click Settings > Manage Nodes.
2. On the Manage Nodes view, click the added orchestrator node.
3. On the SD-WAN Orchestrator Details view, click Discover VeloCloud Devices in the Management widget.

 **How do I add SD-WAN edge devices for monitoring?**

When you click the Discover VeloCloud Devices button on the SD-WAN Orchestrator Details page, SolarWinds Observability Self-Hosted requests VeloCloud device hostnames and IP addresses from the VeloCloud Orchestrator.

A simplified version of Network Sonar Wizard opens and the details from the orchestrator are automatically listed in the IP Addresses box.

4. In Network Sonar Discovery, review the hostnames and IP addresses received from the orchestrator. Include only the devices you want to monitor and click Next.



5. On SNMP, make sure [SNMP credentials](#) for the devices are listed and click Next. If appropriate credentials are not listed, click Add New Credential, and define a new set.

6. Complete the wizard by clicking Discover on the last tab. The wizard searches your network for the hostnames/IP addresses.
7. Discovered devices are listed in the Network Sonar Results wizard. Complete the wizard to add the devices for monitoring.

When you finish the wizard, go to the SD-WAN Orchestrator Details page (Settings > Manage Nodes > click the orchestrator node). After the next poll, the page will display data not only for the orchestrator, but also for monitored edge devices.

Configure SD-WAN monitoring on edge devices already monitored with SolarWinds Platform

If you monitored SD-WAN edge devices in a previous version using SNMP and upgraded to SolarWinds Observability Self-Hosted, [add the VeloCloud orchestrator](#) to automatically pair them with the orchestrator.

Monitor SD-WAN for VeloCloud devices

When you enable SD-WAN polling for a VeloCloud orchestrator node and click it, the SD-WAN Orchestrator Summary page opens.

By default, it includes widgets you can use to manage the device, view the device details, active alerts, latest events, or AppStack for the device.

Edge Devices

This widget lists device names, IP addresses, models, serial numbers and network IDs of edge devices paired with the orchestrator.

List of VLANs on Node

This widget is available on edge Node Details view, in the Network tab. It lists Virtual Local Area Networks (VLANs) on the node.

Orchestrator Inventory

This widget displays a list of VeloCloud devices connected to managed orchestrators. Devices managed by SolarWinds Observability Self-Hosted are marked in the Managed by Platform column.

To add unmanaged devices, click the Discover VeloCloud Devices button and [add them for monitoring](#).

SD-WAN Map

This widget is available on Orchestrator views. It displays connections between monitored devices on the network. For directly connected devices, you can also see used interfaces.

Top Tunnels

On SD-WAN Orchestrator views, the widget displays the metric data (jitter, latency, or packet loss) for top VPN tunnels based on the metric. The widget shows top tunnels on all edges monitored by the orchestrator.

On an edge Node Details view, the widget displays the metric data (jitter, latency, or packet loss) for VPN tunnels on the node.


By default, the widgets display a maximum of 10 tunnels. To change the number of tunnels displayed, complete the following steps:

1. Using an account with administrative privileges, click Settings > All Settings and then click Web Console Settings.
2. In Classic Chart Settings, provide the appropriate number of tunnels to be displayed as the value for Maximum number of data series displayed on chart and submit your changes.

VPN Connections

On SD-WAN Orchestrator views, this widget displays VPN tunnels for all edges paired with the orchestrator.

On an edge Node Details view, the widget displays VPN tunnels for the edge.

 This widget is not displayed by default. See [Add widgets to SolarWinds Platform views](#).

VPN Tunnels

On SD-WAN Orchestrator views, this widget displays VPN tunnels and their metrics for all edges paired with the orchestrator.

On an edge Node Details view, the widget displays VPN tunnels and their metrics for the edge.

i To see VPN tunnel details for VeloCloud orchestrators, you need to have polling for VPN tunnels enabled. When you install 2023.1, the option is enabled by default. When you upgrade to 2023.1 from an earlier version, enable the option manually.

1. Click Settings > All Settings, and then click Polling Settings.
2. Scroll down to VeloCloud settings, enable the Poll VPN Tunnels box, and submit the changes.

Warm Spare

On SD-WAN Orchestrator views, this widget displays High Availability information on all edges paired with the orchestrator.

On an edge Node Details view, the widget displays High Availability information.

High Availability (HA) details displayed in the widget include the device name, status, serial number, HA role, HA mode, and peer device name.

i SD-WAN edges can be installed as a single standalone device or paired with another edge to provide High Availability (HA) support. However, the HA configuration is only for wired WAN connections. Edges configured in HA mode are mirror images of each other and they show up on the Orchestrator as a single Edge.

WAN UpLinks

On SD-WAN Orchestrator views, this widget displays WAN uplinks for all edges paired with the orchestrator.

On an edge Node Details view, the widget displays WAN uplinks for the edge.

Monitor SD-WAN for Viptela orchestrators with SolarWinds Observability Self-Hosted

Starting with 2023.3, you can enable SD-WAN monitoring on your Viptela orchestrators.

For monitored Viptela nodes, you can see orchestrator information. You can use alerts and reports relevant for SD-WAN monitoring.

Monitor SD-WAN edge devices

Add edge devices managed by a monitored Viptela orchestrator to get further details via Viptela API.

- You can display general details, such as edge device name, model, serial number, status, or type.

- You can display all IP addresses for a specific Viptela SD-WAN device and properly match IPs received in NetFlow data to the node.
- You can monitor uplinks (WAN interfaces).

Monitor Viptela orchestrators

- You can display status and response time, polled via REST (Viptela API).
- You can display the list of all devices managed by an orchestrator and add vEdge, cEdge, or vSmart devices in a simplified discovery.

Monitor VPN tunnels for Viptela SD-WAN devices

VPN tunnel names are created as follows:

```
{Source Interface Name} → {Peer System IP / Device ID}:{Peer Interface IP}
```

Review the requirements for Viptela devices

Supported edge devices must have the following SNMP SysObjectID: 1.3.6.1.4.1.41916

For Viptela devices, the following SD-WAN metrics are monitored:

- Orchestrator Inventory and Status
- General Edge device Info
- Edge device uplinks (WAN interfaces)
- VPN tunnels
- VPN tunnel metrics

SNMP-polled metrics for nodes

- Details via Device Studio
- ICMP and SNMP Status and Response time
- SNMP uptime

SNMP-polled metrics for interfaces


- Details
- Status / Availability
- Errors & Discards
- Traffic
- Percent Utilization (regular and real time)

Rate limits

A Viptela Orchestrator is limited to 100 requests/second. By default, SolarWinds Observability Self-Hosted is limited to 50 requests/second. For details, see [Viptela API polling limit was exceeded](#).


Add new SD-WAN devices for monitoring in SolarWinds Observability Self-Hosted

To monitor SD-WAN, add the Viptela orchestrator as a node, and then add edge devices.

 Each monitored Viptela orchestrator uses a node license.


Add Viptela orchestrator

1. Log in to the SolarWinds Platform Web Console as an administrator.
2. Click Settings > Manage Nodes, and then click Add a Node.
3. Provide the Viptela orchestrator hostname or IP address.
4. In Polling Method, select Orchestrators: API.
5. Under Orchestrators, select Viptela Devices.
6. Type your Viptela credentials into Username and Password fields and test the credentials.

 On your Cisco vManage orchestrator device, you might have a self-signed certificate installed by default. If the certificate is not signed by a certification authority, SolarWinds Observability Self-Hosted will display a warning about an invalid SSL certificate. Although you have the option to accept the certificate, SolarWinds recommends that you install a valid SSL certificate for the SD-WAN orchestrator.

7. Review and adjust the device properties.
 - a. Review your credentials and proxy settings.
 - b. To edit how often the node status, or monitored statistics are updated, change the values in the Polling area.

Node Status Polling:	<input type="text" value="120"/>	seconds
Collect Statistics Every:	<input type="text" value="10"/>	minutes
Polling Engine:	● NPM-01 (Primary)	

 For critical nodes, you may need to poll status information or collect statistics more frequently than the default polling intervals.

- c. Enter values for custom properties for the node.

The Custom Properties area is empty if you have not defined any custom properties for monitored nodes. See "[Add custom properties to nodes](#)" in the SolarWinds Platform Administrator Guide.

- d. To adjust when the status of the node changes to Warning or Critical, edit alerting thresholds for the metric. Select the Override box and set thresholds for the node.

Alerting Thresholds	
Percent Packet Loss	<input type="checkbox"/> Override Orion General Thresholds
 Warning:	greater than or equal to 30 %
 Critical:	greater than or equal to 50 %
Response Time	<input type="checkbox"/> Override Orion General Thresholds
 Warning:	greater than or equal to 500 ms
 Critical:	greater than or equal to 1000 ms

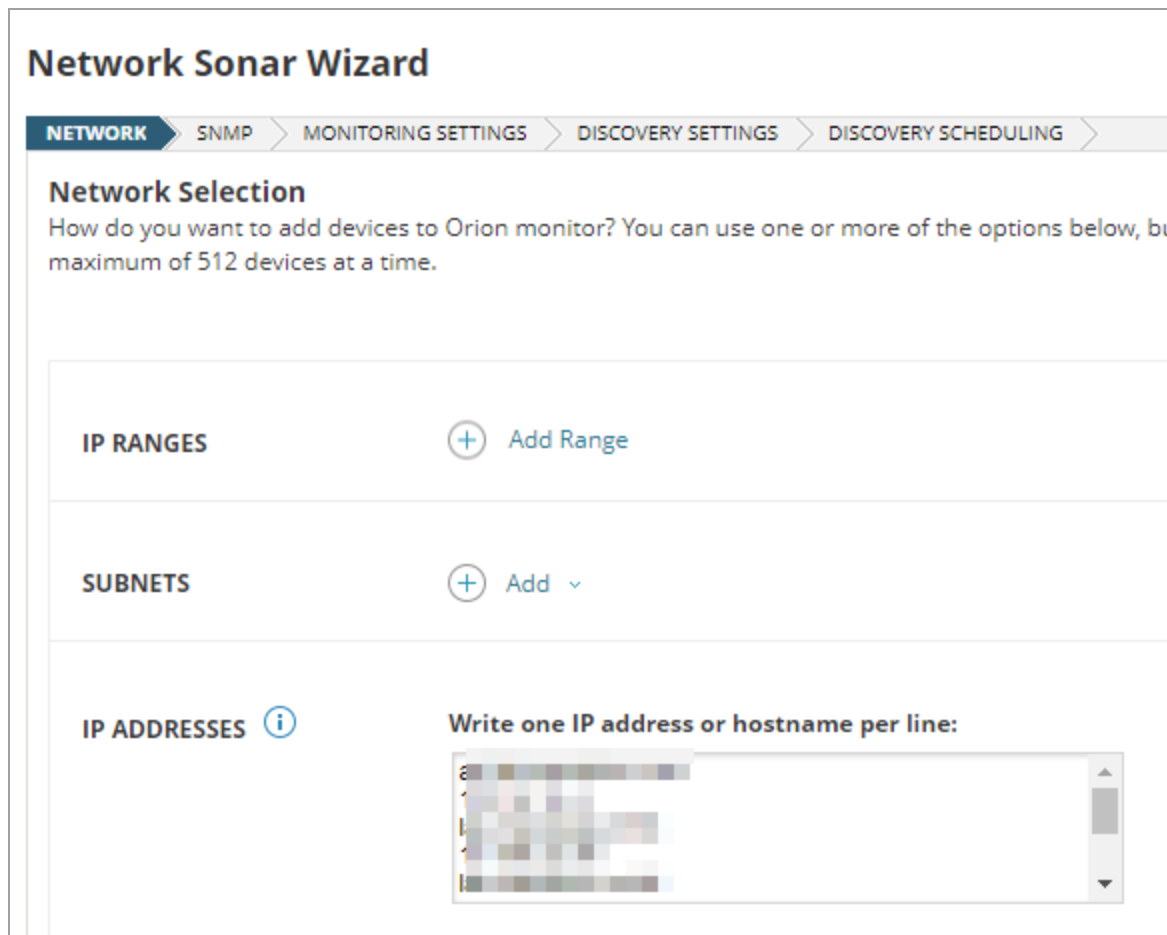
8. Click OK, Add Node.

The Viptela orchestrator is now monitored as an SD-WAN orchestrator. Add connected Viptela devices to complete the SD-WAN configuration.

Add Viptela devices connected to the orchestrator

When you have added the orchestrator for monitoring, you need to add edge devices you want to monitor.

1. In SolarWinds Platform Web Console, click Settings > Manage Nodes.
2. On the Manage Nodes view, click the added orchestrator node.
3. On the SD-WAN Orchestrator Details view, click Discover Viptela Devices in the Management widget.
4. Select what you want to monitor on Viptela devices and click Continue. Network Sonar Wizard will be launched automatically.
 - vEdge
 - cEdge
 - vSmart
5. In Network Sonar Discovery, review the hostnames and IP addresses received from the orchestrator. Include only the devices you want to monitor and click Next.



6. On SNMP, make sure [SNMP credentials](#) for the devices are listed and click Next. If appropriate credentials are not listed, click Add New Credential, and define a new set.
7. Complete the wizard by clicking Discover on the last tab. The wizard searches your network for the hostnames/IP addresses.
8. Discovered devices are listed in the Network Sonar Results wizard. Complete the wizard to add the devices for monitoring.

When you finish the wizard, go to the SD-WAN Orchestrator Details page (Settings > Manage Nodes > click the orchestrator node). After the next poll, the page will display data not only for the orchestrator, but also for monitored edge devices.

Configure SD-WAN monitoring on edge devices already monitored with SolarWinds Platform

If you monitored SD-WAN edge devices in a previous version using SNMP and upgraded to SolarWinds Observability Self-Hosted, [add the Viptela orchestrator](#) to automatically pair them with the orchestrator.

Monitor SD-WAN for Viptela devices

When you enable SD-WAN polling for a Viptela orchestrator node and click it, the SD-WAN Orchestrator Summary page opens.

By default, it includes widgets you can use to manage the device, view the device details, active alerts, latest events, or AppStack for the device.

Edge Devices

This widget lists device names, IP addresses, models, serial numbers and network IDs of edge devices paired with the orchestrator.

Orchestrator Inventory

This widget displays a list of Viptela devices connected to managed orchestrators. Devices managed by SolarWinds Observability Self-Hosted are marked in the Managed by Platform column.

To add unmanaged devices, click the Discover Viptela Devices button and [add them for monitoring](#).

SD-WAN Map

This widget is available on Orchestrator views. It displays connections between monitored devices on the network. For directly connected devices, you can also see used interfaces.

Top Tunnels

On SD-WAN Orchestrator views, the widget displays the metric data (jitter, latency, or packet loss) for top VPN tunnels based on the metric. The widget shows top tunnels on all edges monitored by the orchestrator.

On an edge Node Details view, the widget displays the metric data (jitter, latency, or packet loss) for VPN tunnels on the node.

By default, the widgets display a maximum of 10 tunnels. To change the number of tunnels displayed, complete the following steps:

1. Using an account with administrative privileges, click Settings > All Settings and then click Web Console Settings.
2. In Classic Chart Settings, provide the appropriate number of tunnels to be displayed as the value for Maximum number of data series displayed on chart and submit your changes.

VPN Connections

On SD-WAN Orchestrator views, this widget displays VPN tunnels for all edges paired with the orchestrator.

On an edge Node Details view, the widget displays VPN tunnels for the edge.

i This widget is not displayed by default. See [Add widgets to SolarWinds Platform views](#).

VPN Tunnels

On SD-WAN Orchestrator views, this widget displays VPN tunnels and their metrics for all edges paired with the orchestrator.

On an edge Node Details view, the widget displays VPN tunnels and their metrics for the edge.

WAN UpLinks

On SD-WAN Orchestrator views, this widget displays WAN uplinks for all edges paired with the orchestrator.

On an edge Node Details view, the widget displays WAN uplinks for the edge.

Create wireless heat maps for NPM

Wireless heat maps help you visualize wireless signal coverage on a building floor plan.

i Wireless heat maps are only supported for Cisco wireless controllers. The wireless controllers you want to see on wireless heat maps must be managed in NPM.

Before you begin

- Obtain an image of the wireless coverage area, such as a floor plan.
- Find at least one measurement of the distance between two points on the image, such as the length of a conference room.
- Choose the physical location of access points to accurately place them on the map.

To create wireless heat maps:

1. Start Network Atlas in your SolarWinds program folder. See [Start the Network Atlas](#).
2. On the Welcome to Orion Network Atlas page, click Wireless Heat Map in the Create New section.
3. Enter a name for the new map.
4. [Set a floor plan as the background for visualizing wireless coverage in NPM](#).
5. [Set the wireless heat map scale in NPM](#).

6. [Add wireless access points for NPM wireless heat maps.](#)
To display connected wireless clients on the heat map, you need to add at least four access points, or three access points and a signal sample.
7. Optional: [Improve the accuracy of NPM wireless heat maps by taking samples of the signal strength on real devices.](#)
8. Click Generate to display wireless signal coverage.

See also [Display wireless heat maps for NPM in the SolarWinds Platform Web Console.](#)


Disable the wireless heat map poller in NPM

The wireless heat map poller collects information about the signal strength on monitored access points. By default, this poller is disabled on your devices because of performance issues.


Network Atlas enables the wireless heat map poller on wireless controllers used in wireless maps because the information collected by the poller is required for including access points into wireless heat maps.

When do I need to disable the wireless heat map poller?

If you experience performance issues when working with wireless heat maps, disable the wireless heat map poller on the devices.

 Disabling the poller resolves performance issues, but your wireless heat maps will no longer be updated. The SolarWinds Platform Web Console resources and the Network Atlas will both display the last status generated before you disabled the wireless heat map poller.

1. Log in to the SolarWinds Platform Web Console as an administrator.
2. Click Settings > All Settings in the menu bar.
3. In the Node & Group Management grouping, click Manage Pollers.
4. Locate the wireless heat map poller in the pollers table, and click the item in the Assignments column, such as 1 Node. Clicking the assignments link opens the Assign Wireless Heat Map to Nodes view.
5. Select all nodes for which you want to disable the poller, and then click Off: Disable Poller in the table title.

 Clicking the grey Off icon for the nodes in the Poller Status column disables the poller for the nodes. The icon will turn to green On, and the poller will be disabled.


Set a floor plan as the background for visualizing wireless coverage in NPM

The floor plan should reflect the real dispositions of the office or buildings on the map, so that you can correctly position the wireless access points and reflect the wireless signal coverage on your map.

Requirements:

The floor plan must be a graphic file in one of the following graphics formats:

- Graphics Interchange Format (.gif, non-animated)
- Tagged Image File Format (.tiff)
- Joint Photographic Experts Group (.jpg)
- Microsoft Windows Bitmap (.bmp)
- Portable Network Graphics (.png)

 To ensure the readability of wireless heat maps, use black and white images.


To set a background for wireless heat maps:

1. Create the wireless heat map in the Network Atlas.
2. Click Background Image on the Home ribbon.
3. Navigate to the floor plan image, select the image, and click Open.

The floor plan will appear as the background for your heat map.


Set the wireless heat map scale in NPM

The correct scale is necessary for an accurate display of the wireless coverage provided by your wireless access point.

 You can use online maps, such as the full version of Google Maps, to measure your office building. Locate the building on Google Maps, right-click one wall, and measure the distance to the other wall of the building.

Requirements

- You have already inserted a background image for your wireless heat map (a floor plan).
- You know the distance of two objects displayed on the background image.

 To minimize error, set the scale for the longest distance possible, such as the building or floor length.

To set the map scale:

1. Create the wireless heat map in the Network Atlas.
2. Click Set Scale in the Home ribbon. A blue line segment with squares as end points will appear in the plan.
3. Drag endpoints of the segment to the objects on the map whose distance you know.
4. Fill in the distance between the endpoints into the appropriate field, and select the units (feet or meters).

Example: In floor plans, you usually know the dimensions of individual rooms. Drag and drop the line segment endpoints so that the endpoints are located on the opposite walls, and fill in the width of the room.

5. Click Set Scale to apply the scale to the wireless heat map.

Add wireless access points for NPM wireless heat maps

To generate a wireless heat map, add wireless access points used by client devices into the map.

Requirements

- The wireless LAN controllers must already be managed in your SolarWinds Platform product. See [Monitor wireless networks in NPM](#).
- Only Cisco controllers are supported.
- The wireless heat map poller must be enabled on the wireless LAN controllers that you use in the map.
- To display connected wireless clients on a heat map, add at least four access points, or three access points and a [signal sample](#) on the map.

To add wireless access points:

1. Create a wireless heat map in the Network Atlas.
2. Go to the navigation tree on the left of the Network Atlas main screen.
3. Locate the wireless access points that you want to add to the map.



To find access points on a node, navigate to Orion Objects > vendor name, such as Cisco > appropriate node > Wireless Access Points.

4. Drag the access points to their location on the map.

The selected access points will appear on the map. You can now generate the map.




To make the map more accurate, take signal samples.

Improve the accuracy of NPM wireless heat maps by taking samples of the signal strength on real devices

Wireless heat maps display the ideal wireless signal coverage, they do not count with physical obstacles, such as office walls. To make wireless heat maps more real, measure the signal strength on real devices, such as cell phones, laptops, or tablets connected to your wireless network. The measured values are stored as signal samples and used for calculating the signal coverage on wireless heat maps.

Signal samples represent the signal strength measured in a specified location.

 Take signal samples in places where you expect the signal to be blocked by walls or other obstacles, or in places where the signal strength does not correspond with your heat map.


Take signal samples with cell phones, because polling the signal is usually faster for them.

Simple signal samples

Take a wireless device, walk it to a certain location, and take a signal sample there. Then, walk the device to another location, and take another signal sample. This procedure is called "walking edition" because it requires you to walk through the office.

Multiple signal samples

If you have multiple devices connected to your wireless access points, take multiple signal samples at once (called "sitting edition" because you can do it sitting at your desk).

 Signal samples stay in the map and influence the calculation of wireless heat maps even after the client moves from its position. When you move access points in a map, the signal samples might not be accurate any more. Delete obsolete signal samples, and add new ones.

Requirements

- You need to have a [wireless heat map created](#) and open in the Network Atlas.
- You need to have [wireless access points added](#) into the map.
- You need to have clients, such as cellular phones, tablets, laptops, connected to the access points positioned in your wireless heat maps.

Take simple signal samples

1. Click Take Signal Sample in the Home ribbon. The Signal Sample wizard will display on the right side of the Network Atlas screen as a tab.

2. Walk your device to the location where you want to measure the wireless signal strength and click Next.
3. Select the wireless client (cellular phone, laptop, or tablet) in the drop-down list, and click Next.
4. Drag the client into its current location on the map, and click Next. Network Atlas will start measuring the wireless signal strength in the spot. It can take a few minutes, depending on the device.
5. To add another signal sample, click Repeat, walk the device to a new location, and repeat steps 3 - 4.
6. To apply the measured signal strength to the heat map, click Generate Map.
7. Network Atlas will regenerate the map. Click Close to hide the Signal Sample wizard tab.

Take multiple signal samples at the same time

1. Click Take Signal Sample in the Home ribbon. The Signal Sample wizard will display on the right side of the Network Atlas screen as a tab.
2. Click Use Multiple Devices to Take Signal Samples.
3. Drag the clients to their positions on the wireless heat map, and click Next.

- If there are too many devices, use the search box to find the devices you want to use for creating signal samples.
- Measuring the wireless signal strength can take a few minutes.
- If the signal measuring fails, you can either repeat the measurement for the device, or restart the wizard.

4. Network Atlas will automatically regenerate the map according to the defined signal samples. Click Close to hide the Signal Sample wizard tab.

Troubleshoot NPM wireless heat maps

If your wireless signal coverage on your wireless heat maps is not as expected, you can take the following troubleshooting measures.

- Make sure that the map scale you have entered is precise.
- Make sure that your access points are located correctly.
- Verify that signal samples are up-to-date.
- The signal samples stay in the map even after the device you measured the signal strength on moves away. If you change the position of your access points, or the dispositions of your office, the signal samples might not be accurate and could affect the calculated wireless heat map.


- Delete obsolete signal samples.

To delete a signal sample, open the wireless heat map in the Network Atlas, select the signal sample, and press the Delete key.

- Add new signal samples. See [Improve the accuracy of NPM wireless heat maps by taking samples of the signal strength on real devices](#).

Display wireless heat maps for NPM in the SolarWinds Platform Web Console

1. [Create](#) the wireless heat map in the Network Atlas.
2. Log in to the SolarWinds Platform Web Console.
3. To open a wireless heat map, use one of the following options:
 - Go to the All Wireless Heat Maps resource, and click the thumbnail for the map. The map will open in the Wireless Heat Map view that includes all resources specific for wireless heat maps.

 By default, the All Wireless Heat Maps resource is available on the NPM Summary view.

- Add the Map resource on the view, click Edit, select the map in the list, and click Submit.

Change the time and frequency for regenerating the map

By default, the wireless heat map is regenerated once a day, and the information about clients connected to wireless access points is collected every 5 minutes.

1. Click Settings > All Settings.
2. In the Thresholds & Polling grouping, click Polling Settings.
3. Scroll down to the Wireless Heat Map.
4. Adjust the time when wireless heat maps should be regenerated in Map Generation Start Time.
5. Specify how often the information about clients connected to wireless access points should be collected in Default Client Signal Strength Poll Interval.
6. Click Submit.

Monitor EnergyWise devices in NPM

EnergyWise is a Cisco technology developed to help you cut enterprise energy costs, address environmental concerns, and adhere to government directives around green technologies. By enabling the energy-saving features of EnergyWise-capable devices, you can run business-critical systems in a fully powered state while allowing less critical devices on Power over Ethernet (PoE) ports to power down or drop into standby during off-peak hours.

In the SolarWinds Platform Web Console, you can consult the EnergyWise Summary view and related widgets to help you monitor the energy expended on your network and the energy savings provided by EnergyWise-enabled devices.

- Fully upgrade the IOS of all EnergyWise-enabled devices on your network. For more information, consult your device documentation or www.cisco.com.
- If the EnergyWise Summary view does not display in the SolarWinds Platform Web Console menu bar, see [Add the EnergyWise Summary View to the SolarWinds Platform Web Console menu bar in NPM](#).

Add the EnergyWise Summary View to the SolarWinds Platform Web Console menu bar in NPM

1. Log in to the SolarWinds Platform Web Console as an administrator.
2. Click My Dashboards > Configure.
3. Click Edit beneath the menu bar to which you want to add the EnergyWise Summary view.
4. Drag the EnergyWise button from the Available items list on the left to the correct location in the Selected items list on the right.

i Selected items display from left to right in the selected menu bar as they are listed from top to bottom.


5. Click Submit.

Temporarily reset the current power level of a monitored EnergyWise interface in NPM


Any change made to the power level of a monitored EnergyWise entity is only effective until the next scheduled application of a defined recurrence policy.

To remotely reset the current power level of an interface, the parent node must have not only Community String, but also Read/Write Community String set correctly. See [Edit polling settings](#).

Policies are configured either manually on the monitored device itself or with a configuration management utility, such as SolarWinds NCM. See www.solarwinds.com.

 Some Cisco IOS versions report EnergyWise levels as values 1–11 instead of 0–10. In NPM 10.1.2 and later versions, the levels are automatically corrected. IOS's on some devices are not affected by this issue.

1. Log in to the SolarWinds Platform Web Console as an administrator.
2. Click Settings > Manage Nodes.
3. Locate the device to edit:
 - Use the search tool above the node list to search your database for the parent node of the EnergyWise interface entity you want to reset.
 - Select a Group By option, and click the group including the parent node of the EnergyWise interface entity you want to reset.
4. Expand the parent node, and select the interface entity.
5. Click More Actions > Override Power Level.
6. Select a power level, and click OK.

 To reset the current power level, you can also go to the Interface Details view, and click Set Power Level in the EnergyWise Interface Details resource.

Monitor Azure V-Nets, V-Net Gateways and Site-to-Site Connections with NPM


Starting with NPM 2020.2, you can monitor traffic on Azure Network Gateways on monitored Microsoft Azure cloud accounts.

Supported gateways: Displayed on the [Virtual Network Gateways view](#).

- VPN gateways
- ExpressRoute gateways (only ExpressRoute gateways, no metrics for ExpressRoute connections are supported)

Supported connections to Azure Network gateways: Displayed on [Site-to-Site Connections view](#).

- Site-to-Site connections

 ExpressRoute connections are not supported and you will not be able to see them in NPM.

Add an Azure cloud account for monitoring

Before you begin

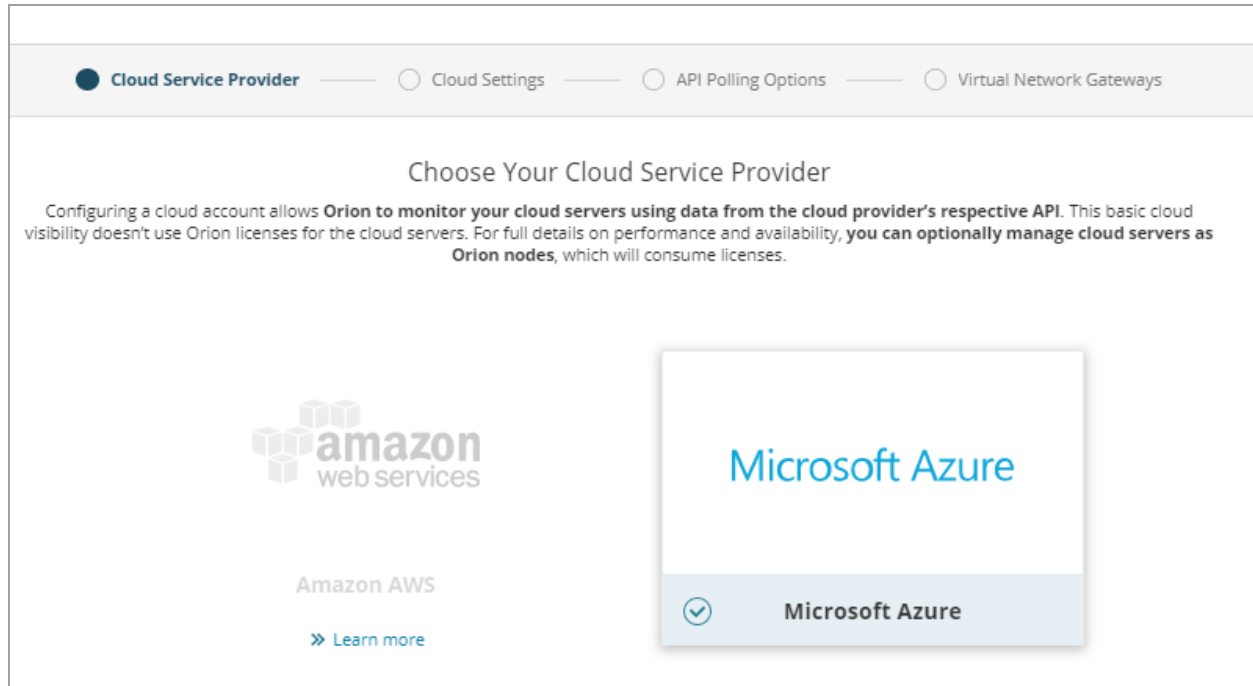
- Ensure that you have [configured your Microsoft Azure cloud account](#) to allow integration with the SolarWinds Platform.
- [Prepare credentials](#) necessary for configuring the SolarWinds Platform to access your Azure cloud.
- Check that Virtual Network Gateways to monitor are connected to Virtual Networks directly. Connections via Virtual WAN are not supported.
- To monitor the Azure cloud, you must use the main polling engine.

For more details, see [Monitor cloud instances and VMs](#) in the SolarWinds Platform documentation.

Use the Add Cloud Account to add an Azure account to NPM for monitoring. The wizard displays steps relevant for installed SolarWinds Platform products. See [Add the first cloud account](#) for details.

1. In the SolarWinds Platform Web Console, click My Dashboards > Home > Cloud.
2. On the Getting Started widget, click Monitor My Cloud instances and complete the wizard.

3. On Cloud Service Provider, select Microsoft Azure, and click Monitor Azure.



4. Provide credentials for the Azure Cloud and click Continue.

Azure Cloud Settings

Cloud Account Display Name

CREDENTIALS
[How do I find these credentials?](#)

Subscription ID

Tenant / Directory ID

Client / Application ID

Application Secret Key

Credential Description Optional

You can use this field to help you remember which Subscription ID was used

[Are you behind a proxy?](#)

TEST CONNECTION

AUTO MONITORING

Automatically monitor any future resources launched

- On API Polling Options, click Continue.

API Polling Options (My Azure Cloud)

Use these options to control how your free Azure Monitor API calls are used.

Enable Azure Monitor API Polling ?

Azure Resource Polling Frequency

5 minutes

Alerting

You can define alerts to be notified when you are approaching and/or have exceeded your free API call limit. Out-of-the-box alerts have been provided for you.

[Manage Alerts](#)

- On Virtual Network Gateways, enable Virtual Network Gateways polling and click Finish.

Virtual Network Gateways (My Azure Cloud)

Virtual Network Gateways statistics are collected if **Enable Azure Monitor API Polling** on **API Polling Options** step is turned on

Enable Virtual Network Gateways polling

Click My Dashboards > Home > Cloud. On the Cloud - Cloud Summary view, you will see two new subviews relevant for NPM - Virtual Network Gateways and Site-to-Site Connections. If you don't have NPM installed, the subviews are empty.

Display the overview of monitored cloud environment

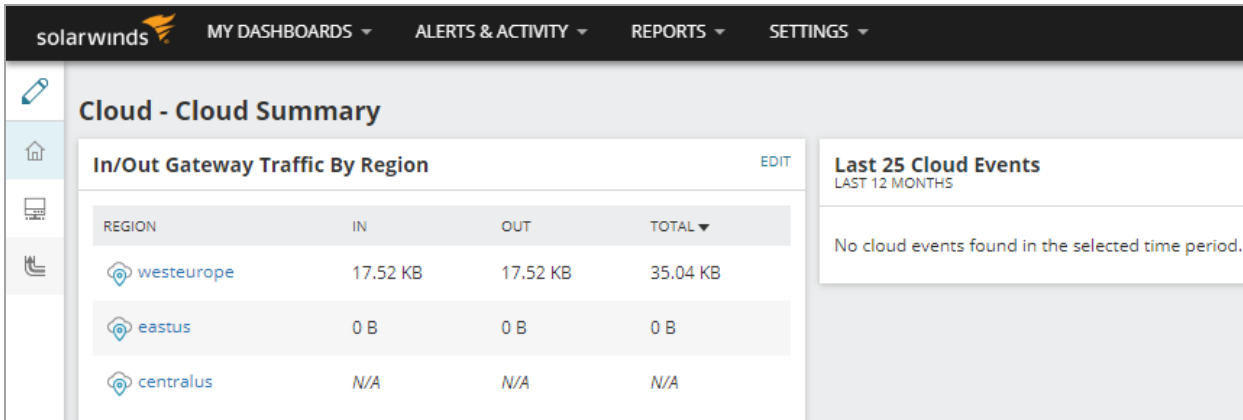
Click My Dashboards > Home > Cloud.

The Cloud Summary view displays summary widgets that give you an overview of monitored cloud accounts.

By default, the Cloud Summary view displays the following widgets:

- The **In/Out Gateway Traffic by Region** widget lists regions with in and out traffic. Click a region to go to the list of gateways in the selected region.
- The **Last 25 Cloud Events** widget lists the last 25 events that occurred on monitored cloud entities. Hover over the event to display a popup with more information or click the Event message to go to the cloud entity details view.

If you have multiple SolarWinds Platform products installed, the Cloud Summary can include additional widgets and additional subviews. See [Explore cloud instances and VMs on the Cloud Summary page](#) in the SolarWinds Platform documentation.



The screenshot shows the SolarWinds interface with a navigation bar at the top containing 'MY DASHBOARDS', 'ALERTS & ACTIVITY', 'REPORTS', and 'SETTINGS'. The main content area is titled 'Cloud - Cloud Summary' and contains two widgets. The first widget, 'In/Out Gateway Traffic By Region', displays a table with the following data:

REGION	IN	OUT	TOTAL
westeurope	17.52 KB	17.52 KB	35.04 KB
eastus	0 B	0 B	0 B
centralus	N/A	N/A	N/A

The second widget, 'Last 25 Cloud Events', shows 'LAST 12 MONTHS' and a message: 'No cloud events found in the selected time period.'

Monitor Virtual Network Gateways on Microsoft Azure clouds

Click My Dashboards > Home > Cloud, and then click Virtual Network Gateways subview.

The Azure Network Gateways view lists gateways available on monitored Azure cloud accounts, with the following details:

- Gateway name, status, type, and location
- Provision state and status
- Name of the related virtual network,
- IP address of the gateway
- Incoming and outgoing tunnel traffic (last polled values)
- Number of site-to-site (S2S) connections on the gateway

Use the search and filtering options to find a gateway.

Click a gateway name to open the gateway in PerfStack and display In, out, and average availability.

You can select up to three gateways and click Open in PerfStack to compare their availability.

Click S2S connections to see site-to-site connections for the gateway on the Site-to-Site Connections view.

Connection Name	Provision State / Status	VNet	Public IP	In	Out	Connections
VPN_express	Failed / Down	Express_route_VNet	N/A	---	---	193
Express_route_Test	Succeeded / Up	Express_route_VNet	Public IP	---	---	192
KarelGateway	Succeeded / Up	DirewolfTest	Public IP	17.4 6 KB	17.4 6 KB	191
VNet2GW	Succeeded / Up	VNet-demo-2	Public IP	0 B	0 B	190

Monitor Site-to-Site Connections on Microsoft Azure clouds

Click My Dashboards > Home > Cloud, and then click the Site-to-Site Connections subview.

Use the search and filter options to find a connection.

Click a connection name to open it in PerfStack and display average in and out bps.

You can select up to three connections and click Open in PerfStack to display average in and out bps for the connections in PerfStack.


Connection Name	Encryption	Hashing	In	Out	VNet	VPN Gateway
third_connection	N/A	N/A	---	---	VNet-demo-1	VNet1GW
test_not_running	N/A	N/A	---	---	VNet-demo-1	VNet1GW
NPM_test_1	N/A	N/A	---	---	VNet-demo-1	VNet1GW

Configure Azure settings relevant for NPM

Edit your Azure cloud details, such as account credentials, or API polling options.

Review the number of API requests you have already used this month to get cloud details to NPM.

1. Click Settings > All Settings, and then click Cloud Infrastructure Monitoring Settings in the Product Specific Settings section.
2. Click Manage Cloud Accounts.
3. Select your cloud account and click Edit Properties.
 - On the General tab, you can adjust your cloud credentials, or enable/disable adding cloud resources on the account added in the future.
 - On the API Polling Options tab, you can enable/disable API polling, change the polling frequency, or review an estimation how many requests all accounts in your cloud subscription used up to get polling details to NPM.

 The SolarWinds Platform measures requests initiated by your Orion deployment. The number covers requests by all accounts in your subscription.

4. Save your changes.

Disable Virtual Network Gateways Polling

To stop collecting data from your Azure account, disable polling from Virtual Network Gateways. Disable this polling to limit the number of API requests and thus decrease your Azure subscription costs.


1. Click Settings > All Settings, and then click Cloud Infrastructure Monitoring Settings in the Product Specific Settings section.
2. Click Manage Cloud Accounts.
3. Select your cloud account and click Virtual Network Gateway Settings.
4. Click the green slider to disable polling and thus stop sending API requests to your Azure account. The slider turns grey and you do not poll Virtual Network Gateways anymore.

Set NPM thresholds

SolarWinds NPM thresholds are relevant for nodes and interfaces. They include Cisco Buffer Misses, Interface Errors and Discards, Interface Percent Utilization, and Flapping Routes.

- When a metric reaches the specified Critical Level threshold on a node or interface, the node or interface will be displayed as bold red in resources and reports.
- When a metric reaches the specified Warning Level thresholds on a node, the node or interface will be highlighted in red in appropriate resources and reports.
- Flapping Routes use different colors when the thresholds are exceeded: red for the error threshold and yellow for the warning threshold.

To change thresholds globally, for all monitored objects, complete the following steps.

 You can also change node or interface thresholds for individual nodes or interfaces. See [Edit node properties](#) or [Edit interface properties in NPM](#).

1. Log in to the SolarWinds Platform Web Console using an account with Administrator Rights.
2. Click Settings > All Settings in the menu bar.
3. In the Thresholds & Polling grouping, click NPM Thresholds.
4. Provide the values for Critical Level and Warning Level for the selected metrics.
See [Network Performance Monitor Thresholds](#) for more information.
5. For the Interface Percent Utilization metric, specify if you want to use average or peak daily values in calculations for capacity forecasting.
6. Click Submit.

See also [Define UnDP Warning and Critical thresholds in NPM](#).

Create custom monitors in NPM

With NPM, you can extend monitoring to non-standard devices, using [object identifiers \(OIDs\) organized in management information bases \(MIBs\)](#).

NPM provides the following advanced monitoring options:

- [Device Studio pollers](#): Create pollers for certain technologies directly in the SolarWinds Platform Web Console.

What is a poller?

Statistics monitored on your devices are specified by pollers. Pollers hold information about a monitored property, how to get the current value for the property, and where and how to display the retrieved data.

What do you need custom pollers for?

- To monitor a specific metric which is not monitored out-of-the box.
- To monitor special equipment.
- To monitor objects although the number of monitored objects exceeds a poller's capacity limitation.

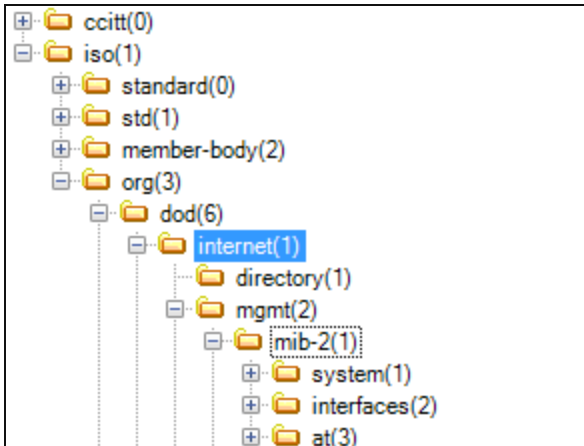
- [Universal Device Pollers](#): If there is a specific metric that is not monitored out-of-the box, or if you have special equipment you need to monitor, create a custom poller based on a specific object identifier (OID) and transform polled results into a resource in the SolarWinds Platform Web Console.

Review the comparison of UnDP and Device Studio pollers to determine which poller to use.

undp	device studio poller
Can poll only one OID.	Can poll multiple OIDs for a given technology.
Cannot perform logical operations or transformations on the polled data.	Can perform logical operations or transformations on the polled data.
Polled values are displayed in dedicated resources.	Polled values are displayed in existing resources.

Management Information Base (MIB) in the NPM

Management Information Base (MIB) is a structure that describes all objects a device can report on, such as CPU, fan, or temperature. MIB contains the name, datatype, and the object identifier (OID). MIB is a hierarchical structure, displayed as a navigation tree. Every entry in the MIB tree is a value for a specific component on a specific device.



Each entry in the tree is followed by a number in parenthesis. Each entry in the tree can be specified using the sequence of numbers, such as 1.3.6.1 (iso.org.dod.internet). The unique numerical value is the OID.

For more information, see [Management Information Base \(MIB\) for the Simple Network Management Protocol \(SNMP\)](#).

Monitor custom statistics based on OIDs with Universal Device Pollers in the NPM


SolarWinds Universal Device Poller (UnDP) is a customization feature of NPM. With UnDP, you can create custom monitors for almost any statistic provided by SNMP based on its Management Information Base (MIB) and object identifier (OID).

With Universal Device Poller, you can monitor:

- Interface traffic
- CPU temperature
- Addressing errors
- UPS battery status
- Current connections to a website

Before you start configuring UnDPs

- Consult your vendor documentation, and find out which OID you want to monitor.
- Create a list of nodes that you want to poll the custom statistic on.


 UnDPs do not collect data from Orion Failover Engine or Hot Standby Engines. If a NPM server fails, data collection for any Universal Device Pollers stops on the server.

UnDPs are tied to the polling engine on which they are hosted. If you move a monitored node from one polling engine to another, you must also move the UnDP poller.

Define a custom statistic to monitor in the NPM

Statistics monitored on your devices are specified by pollers. Pollers hold information about a monitored property, how to get the current value for the property, and where and how to display the retrieved data.

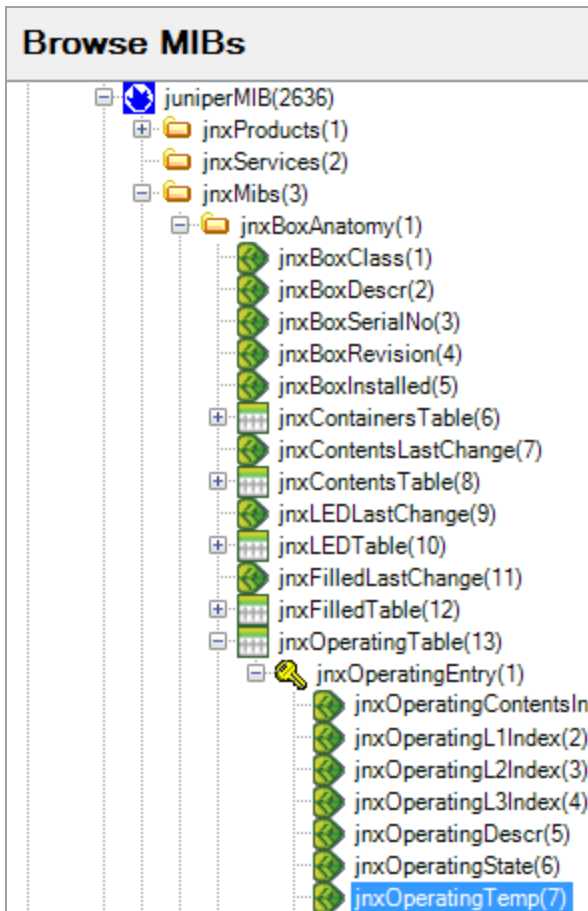
Defining a custom statistic for monitoring means creating a UnDP poller.

 Before you begin, consult your vendor documentation, and find out which OID you want to monitor.

1. Start the Universal Device Poller application, for example by clicking Start > SolarWinds Orion > Universal Device Poller.
2. If prompted, [download and install the MIB database](#).
3. Click New Universal Device Poller.

4. Specify the OID:

- a. Click Browse MIB Tree, and click Search MIBs in the upper-right corner.
- b. Select a Search By option, enter a string, and click Search.
- c. Select the OID, and click Select.



- If you know the OID, fill it in.
- If you know approximately where in the MIB tree you can find the OID, click Browse MIB Tree, navigate in the MIB tree to the OID, and click Select.

5. Test the selected OID against a device. Select a node, and click Test. See [Troubleshooting failed tests](#) if the test fails.

JUNIPER-MIB:jnxOperatingTemp
OID: 1.3.6.1.4.1.2636.3.1.13.1.7

Select a device:

Vendor

Juniper Networks, Inc.

Test Successful

Row ID	Test Result
1.1.0.0	0
2.1.0.0	30
2.2.0.0	30

6. On the Define Your UnDP screen, edit the suggested Name and Description. The poller name is populated automatically. The name is required and cannot contain spaces.


OID:

Enter an OID or browse the MIB tree.


Name:

Description:

7. To customize the value type, SNMP Get type, polling type or interval, click Advanced Options, and change the defaults:
 - a. Select the expected format of values in MIB Value Type.
 - For Rate or Counter, provide a Unit and Time Frame.
 - For Raw Value, select a display Format for the polled raw values .
 - For Raw Value > Enumerated, click Map Values to provide strings corresponding to the values returned by the poller.
 - b. Select SNMP Get Type, and decide whether the poller should poll nodes or interfaces.
 - c. Specify the Polling Interval in minutes. Use values between 1 and 600.

 If you want to use the poller in a transformation, make sure that all pollers in the transformation have the same Polling Interval.


8. Keep default settings for Status (Enabled) and Keep Historical Data (Yes). With these options enabled, you can see the trend of polled values in SolarWinds Platform Web Console views.
9. Specify the Group to which you want to add the poller, and click Next.

 To create a new group, type a name for the group into the Group box.

10. Select devices to poll the statistic, click Test, and then click Next.

 Custom OIDs often work only for identical nodes.

11. If the selected OID is a table, specify labels for the rows in the table.

 **Label the Rows in Your Table**
Select a method for labeling each of the rows

Use interface names from Orion


Use labels from a table column

Use a custom label

12. Select the SolarWinds Platform Web Console views that can display the poller as a chart, gauge, or table, and click Finish.

Include the following on Orion views:			
Chart	Gauge	Table	Orion View Name
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Orion Summary Home
<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Node Details - Summary
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Node Details - Vital Stats
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	Current Top 10 Lists

The new poller is added to All Defined Pollers and will be polled on the selected nodes or interfaces. You can now add Universal Device Poller resources showing the polled values to SolarWinds Platform Web Console views.

- 
- To view the poller status on maps, create a network map, add the poller into the map, and add the map on a view. See [View UnDP status on Network Atlas maps in NPM](#).
 - To check that your UnDP pollers are properly configured, start Orion Diagnostics in your SolarWinds Orion > Documentation and Support program folder, right-click a UnDP, and select Run Tests.

Troubleshooting failed tests

If the test fails on a node or interface, make sure that the following settings are correct:

- Verify that the test node is being polled using the correct community string. See [Edit node properties](#).
- Does the device support the polled MIB or OID? See the vendor documentation to confirm the MIBs supported by your device.
- Can your SolarWinds Platform server access the device? Make sure that the device is responding to both ICMP and SNMP requests.

Select NPM nodes or interfaces to poll a custom statistic


When you have created a UnDP poller, specify the devices (nodes or interfaces) to monitor the statistic.

Before you begin, make sure the UnDP poller is created and enabled. See [Define a custom statistic to monitor in the NPM](#).


1. Start the Universal Device Poller application, for example by clicking Start > SolarWinds Orion > Universal Device Poller.
2. Click Assign Pollers.
3. Navigate the poller tree, select the pollers you want to assign, and click Next.

By default, there are two poller groups:

- Example - all predefined out-of-the box UnDP pollers.
- Default Group - all user-defined UnDPs if they are not assigned to any other group.

 Selecting a poller group selects all pollers in the group. If you do not want to assign all pollers, clear the pollers that you do not want to assign.

4. Expand the node tree down to the interface level, and select the elements to apply the pollers.



- Interfaces are not displayed unless you are assigning an interface poller.
- Selecting a node automatically assigns a selected interface poller to all interfaces on the node. Clear boxes for interfaces that should not be assigned to the poller.

5. Click Test to see current results of the selected pollers on the selected nodes or interfaces. If the test fails, see [Troubleshooting failed tests](#).
6. After you have completed your poller assignments, click Finish.

Transform poller results in the SolarWinds Platform

Values polled by a custom poller are often better understood after a calculation transforms the value to a different format.

For example, if a poller returns temperature values in Celsius, you might want to see the values in Fahrenheit.

i Pollers that you use in a transformation must be assigned to the nodes to poll for values that will be transformed.

1. Start the Universal Device Poller application, for example by clicking Start > SolarWinds Orion > Universal Device Poller.
2. Click Transform Results, and click Next to acknowledge examples of transformations.
3. Type the name and description for the transformation, and click Next. Names must be unique.

Names are required. Any spaces in the name are removed.

Descriptions are optional but might be helpful in identifying the type of information generated by the transformation.

You can also change other default settings:

- a. Select Yes in the Keep Historical Data section. You will be able to view the transformed poller data in charts and gauges in the SolarWinds Platform Web Console.
- b. Select Enabled as the Status if you want your transformation to begin calculating results immediately.

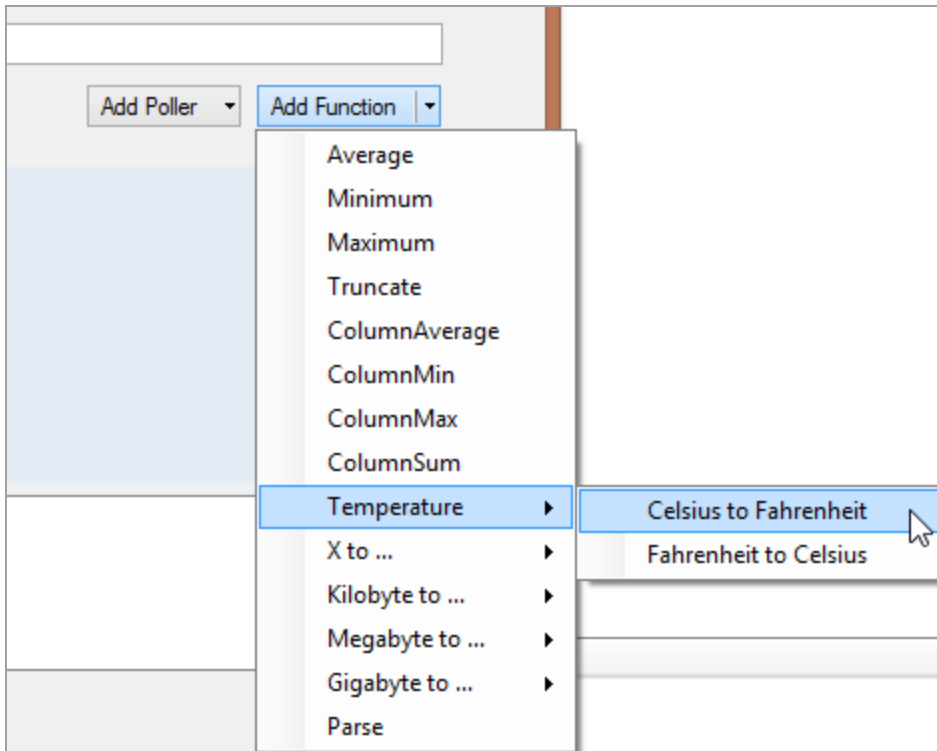
i If you select Disabled, the transformation will not transform polled data.

- c. In the Group field, select a group where you want to add the transformation. To add a group, provide the new group name.
- d. Optional: provide a polling interval.


! Make sure all pollers in the transformation use the same polling interval.

4. Provide the formula for calculating the transformation.

- a. Click Add Function, and select a function.



- b. Click within the bracket, click Add Poller, and select the poller you want to transform.



- Separate pollers with commas. The following example averages the results of three pollers:

```
avg({poller1},{poller2},{poller3})
```

- Use standard mathematical operations:

```
{poller1}+{poller2}
```

- Use the mathematical constants e and π , as `E()` and `PI()`, respectively.
- Nest formulas. The following example returns the average of two poller comparisons:

```
avg(min({poller1},{poller2}),max({poller3},{poller4}))
```



5. Test the transformation on a device, and click Next.

Troubleshooting failed transformation tests


If the test fails, verify the following items:

- Is your formula correct? Ensure that all braces are balanced, that there are no unnecessary spaces, and that all pollers return the same type of values.
- Are you using the correct community string for the node that is being polled for the test? For more information about providing community strings, see [Edit node properties](#).
- Does the device support the polled MIB or OID? See the documentation supplied by the device vendor to confirm supported MIBs for your device.
- Can you access the device from the SolarWinds Platform server? Confirm that the device is responding to both ICMP and SNMP requests.


6. Select nodes for the transformation, and click Test.

-  • Interfaces are not displayed unless your poller transformation operates on an interface poller.

7. If the transformation output is a table, select labels for the rows in the table, and click Next.
8. Select SolarWinds Platform Web Console views where you want to include the transformed values as a chart or table, and click Finish.

-  Click Preview to see how your poller resource will display in the selected SolarWinds Platform Web Console view.

The new transformation is added to All Defined Pollers and applied on the selected nodes or interfaces. You can add a Universal Device Poller resource to display transformed values in the SolarWinds Platform Web Console views.

-  If the transformation combines data from other pollers, make sure that it is assigned to the same node or interface as the pollers used for the transformation and that it has the same polling interval.

Create pollers by duplicating and adjusting pollers in NPM

When creating similar pollers, consider copying a poller and modifying it.

1. Start the Universal Device Poller application, for example by clicking Start > SolarWinds Orion > Universal Device Poller.
2. In the All Defined Pollers pane, locate the poller that you want to duplicate.

i To confirm that you have selected the appropriate poller, view the poller properties in the main Universal Device Poller window.

3. Right-click the poller, and select Duplicate Poller.
4. Change the Name of the poller.
5. Adjust the poller settings. See [Define a custom statistic to monitor in the NPM](#).

Import UnDP pollers to NPM

You can import custom UnDP pollers exported from UnDPs installed with earlier SolarWinds Platform versions.

i You cannot import device-specific MIBs into the SolarWinds MIB Database, but you can import UnDP pollers based on OIDs from device-specific MIBs. Import a poller and assign it to nodes or interfaces in your environment.

1. Start the Universal Device Poller application, for example by clicking Start > SolarWinds Orion > Universal Device Poller.
2. Click File > Import Universal Device Pollers.
3. For each poller you want to import, complete the following steps:
 - a. Click Open, and locate the poller.
 - b. Select the poller, and click Open.
4. Select the pollers to import from the list on the left, and click Import. Selected pollers will move to the pane on the right.

i

- To select multiple pollers, hold down SHIFT or CTRL, and click the pollers you want.
- To remove a poller from the Selected Pollers list, select the poller and click Remove.
- To collapse all folders and see just the group names, hold down SHIFT, and then click – next to any of the group names.

5. Click OK.
6. To begin polling, enable the poller.
 - a. Select the imported poller in the All Defined Pollers pane of the Universal Device Poller window.
 - b. Click Edit Properties.

- c. Confirm that the poller Status is Enabled, and click Finish.

 If Disabled, the poller will not collect data until you enable it.


7. Specify nodes or interfaces to be polled by the imported poller. See [Select NPM nodes or interfaces to poll a custom statistic](#).

When the imported poller is enabled and assigned to the devices, the poller begins collecting statistics. To view the statistics, log in to the SolarWinds Platform Web Console, go to a view for the node or interface to which the poller is assigned, and consult the poller resource. See [View Universal Device Poller statistics in NPM](#).

Export UnDP pollers from NPM

If you want to use your custom UnDPs in later NPM versions or on different polling engines, you need to export them first.

1. Start the Universal Device Poller application, for example by clicking Start > SolarWinds Orion > Universal Device Poller.
2. Click File > Export Universal Device Pollers.
3. In the Pollers pane on the left, navigate to the pollers that you want to export.

-  • To select all pollers in a group, select the group.
• To select multiple pollers, hold down SHIFT or CTRL and click the pollers to export.

4. Select the pollers, and click Export. Pollers will move to the Selected Pollers pane.

 To remove a poller from the list of pollers for export, select the poller and click Remove.

5. Click Save.
6. Navigate to the location where you want to export the selected pollers, provide a File name, and click Save.

Selected pollers will now be stored as a `.UnDP` file in the specified location. You can use the `.UnDP` file to import the pollers on another polling engine.

Temporarily suspend collecting statistics for pollers in NPM

When you assign a poller to nodes or interfaces, it starts collecting statistics on the selected elements. If you want to suspend data collection for a poller without deleting it, disable the poller.

1. Start the Universal Device Poller application, for example by clicking Start > SolarWinds Orion > Universal Device Poller.
2. In the All Defined Pollers pane, navigate to the poller you want to disable.



To confirm that you have selected the appropriate poller, view the poller properties in the main Orion Universal Device Poller window.

3. Select the poller, and click Edit Properties.
4. Set Status to Disabled, and click Finish.

The poller will now still be available in the Universal Device Poller application, but will not collect any statistics.

Define UnDP Warning and Critical thresholds in NPM

If values polled by UnDPs on a device reach a certain level (critical or warning threshold), the UnDP on the device is highlighted in the SolarWinds Platform Web Console.



To get notified about exceeding a threshold in an email, configure an alert.

To see pollers with exceeded thresholds in a map, see [View UnDP status on Network Atlas maps in NPM](#).

1. Log in to the SolarWinds Platform Web Console as an administrator.
2. Click Settings > All Settings in the menu bar.
3. In the Thresholds & Polling grouping, click Custom Poller Thresholds.
4. Select a poller.
5. Select whether the expected polled value is a Text or a Number.

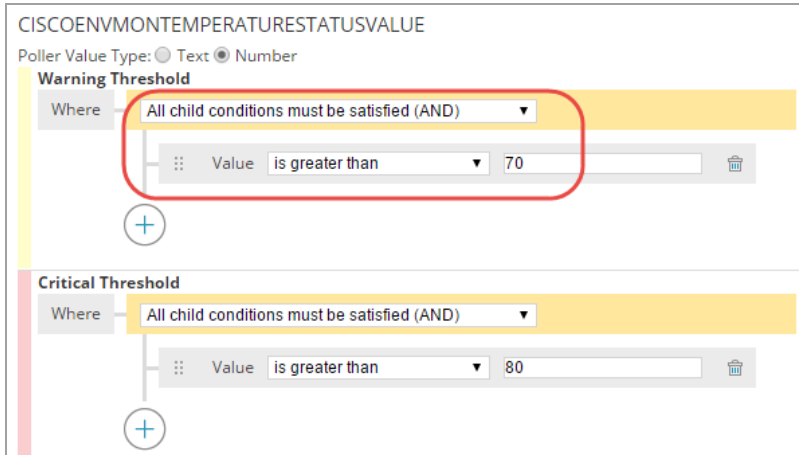


The Poller Value Type determines how the polled value will be interpreted. It also influences the set of possible comparison functions.

- For the Number type, available values include `is greater than` or `less than`.
- For the Text type, available values include for example `contains`.

6. Build conditions to define both Warning and Critical Thresholds:

- a. Select whether All Child Conditions Must Be Satisfied (AND) or if only At Least One Child Condition Must Be Satisfied (OR).
- b. Select a comparison relation, and provide a threshold value on which the comparison is based.



- c. Click + to add additional conditions, as required, to define the poller threshold.

7. After configuring all thresholds, click Submit.

If a value reported by the device belongs to the range defined by the Warning Threshold, pollers in maps will be yellow.

If a value reported by the device belongs to the range defined by the Critical Threshold, pollers in maps will be red.

View Universal Device Poller statistics in NPM

If you want to see a poller results in the SolarWinds Platform Web Console, you need to define which widgets, or resources, should be displayed on which views.

Prerequisites


The poller must be enabled, and assigned to the devices.



Set the poller to collect historical statistics. Without historical data, SolarWinds Platform Web Console widgets will only display the last polled value, and you cannot add charts with the poller results to the SolarWinds Platform Web Console.

Define widgets with UnDP results for SolarWinds Platform Web Console views

1. Start the Universal Device Poller application, for example by clicking Start > SolarWinds Orion > Universal Device Poller.
2. In the All Defined Pollers pane, select the poller whose results you want to add as a SolarWinds Platform Web Console resource.
3. Right-click the poller, and click Web Display.
4. Confirm that Yes is selected, and select the types of poller widgets that you want to display on individual SolarWinds Platform Web Console views.

 Click Preview to see what the poller widget will look like in the SolarWinds Platform Web Console view.

5. Make sure Do Not Show This Poller If It Is Not Assigned is selected. It ensures that the custom poller widget appears only on views for nodes or interfaces that have the custom poller assigned to them and enabled.
6. Click Finish.


When you log in to the SolarWinds Platform Web Console, the selected widgets with poller data will appear on selected views for nodes or interfaces that have the poller assigned to them and enabled.

See also [View UnDP status on Network Atlas maps in NPM](#).

View UnDP status on Network Atlas maps in NPM

In the SolarWinds Platform Web Console network maps, you can see when a Universal Device Poller on a device returns values that exceed the warning and critical thresholds.

1. [Create a Universal Device Poller](#) in the UnDP application.
2. Assign the poller to nodes.
3. [Define warning or critical thresholds](#) specifying when you want the pollers to be highlighted.
4. Create a network map in the Network Atlas, drag the UnDPs into it, and save the map.

 To add a UnDP on a map, start the Network Atlas, navigate to a node on which the UnDP is enabled (Vendor > Node Name > Custom Node Poller), and drag the poller into the map.

5. Log into the SolarWinds Platform Web Console, go to the map view.
6. Locate the Map resource (or add it if not available), click Edit and select your map.

You can now see UnDPs for your nodes in the SolarWinds Platform Web Console map. When the polled UnDP values exceed the warning threshold, the UnDP icon turns yellow on the map. After reaching the critical threshold, the icon turns red.

Assign Universal Device Pollers to devices monitored in NPM

NPM provides both a selection of predefined pollers and the Universal Device Poller (UnDP) utility for defining your own pollers to monitor specific aspects of your network devices.

i UnDPs are SNMP-based. You can only assign them to nodes polled through SNMP.

If you do not see a poller that meets your monitoring needs, use the Universal Device Poller to create a poller. See [Monitor custom statistics based on OIDs with Universal Device Pollers in the NPM](#).

1. Log in to the SolarWinds Platform Web Console as an administrator.
2. Click Settings > Manage Nodes.
3. Select the node, interface, or volume you want to assign Universal Device Pollers to.

i The list only includes nodes polled through SNMP. If you cannot see a node, [check the polling method, and change it to SNMP](#).

See [Add a single node](#) for information about adding nodes, interfaces, or volumes for monitoring.

4. Click Assign Pollers in the Node Management toolbar.
5. Expand the poller group, and select pollers to assign.
6. Click Submit, and click OK to confirm the assignment.

NPM polls data specified by the poller on the node, interface, or volume.

Update the SolarWinds MIB Database for the SolarWinds Platform

SolarWinds maintains a MIB database that serves as a repository for the OIDs used to monitor a wide variety of network devices. The MIB database is updated regularly.

When you are creating a UnDP poller and cannot find an OID in the MIB tree, update the MIB database.

- [Download MIB database from the SolarWinds Platform Web Console](#) (Orion Platform 2020.2 and later)
- [Download MIB database from the Customer Portal](#) (offline environments)
- [Learn more about MIBs](#)

Download the MIB Database from the SolarWinds Platform Web Console


You can check the status of your MIBs database and download the latest MIBs database directly from the SolarWinds Platform Web Console.

1. Log in to your SolarWinds Platform Web Console.
2. Click Settings > All Settings > MIBs Management in the Details section.


The page informs you about the MIBs database installed on your main polling engine and about the latest available version.

3. If an updated version of the MIBs database is available, click Download.

This downloads the latest version of the MIBs database as a MIBs.msi file.

 Keep the Show notifications setting enabled to be informed about an updated MIBs database version on relevant places - in the MIB Browser, in the Device Studio, and in the Universal Device Poller tool. The message includes a link to MIBs Management so that you can quickly download and update the MIBs database.

MIBs Management

 <p>MIBs database is out-of-date Checked on 10.196.4.87</p>	<p>1.1.0.169 Installed</p>	<p>1.1.0.170 Available</p>	<input type="button" value="Download"/>
--	--------------------------------	--------------------------------	---

A Management Information Base (MIB) is a hierarchical structure that describes statistics a device ca

Settings

Show notifications in related Orion locations

4. After the download completes, run MIBs.msi installer and complete the installation wizard.
The Installer informs you when Orion services need to be restarted and restarts them if necessary.
5. Run the .msi file on all deployed polling engines - on the main polling server, HA backup server, and any additional polling engines.

Download the MIB database from the Customer Portal

When your SolarWinds Platform is not connected to the Internet, download the MIB database from the Customer Portal.

Download the MIB database

1. Log in to the Customer Portal (<https://customerportal.solarwinds.com/>) using your SolarWinds Customer ID and Password.
2. In the navigation bar, click Downloads > Orion MIB Database.
3. On the MIB Database page, click **Download As MSI**.
4. If you are using Internet Explorer and it prompts you to add the SolarWinds downloads site <http://solarwinds.s3.amazonaws.com>, add the site to your trusted sites.
5. Install the MIB database. Follow instructions for [2020.2 and later](#) (offline environment).

Install the MIB database in an offline environment

1. If your SolarWinds Platform is in an offline environment, download the MSI version of the MIBs database and transfer the file to your scalability engines.
2. Run the MSI on all deployed polling engines - on the main polling engine server, HA backup server, and any additional polling engines.

When finished, you can check the MIBs database version in the MIBs Management view of the SolarWinds Platform Web Console.

Learn more about MIBs

What is a MIB?

See [What is a MIB, OID, and how they are used](#).

Check an OID in the SolarWinds MIBs database

1. Go to the Details view for a device.
2. In the Management widget, click the MIB Browser button.

The online MIB Browser opens. You can navigate through the SolarWinds MIB database and check whether an OID is available.

Add a device MIB

If you have a specific device MIB, you can have it added to the SolarWinds MIB database. See [Add MIBs to the SolarWinds MIB database](#).

Manage unique devices on the network with NPM

If you have devices on your network that SolarWinds does not recognize for polling, you can either edit an existing poller to suit your device needs, or create a poller specifically tailored to your device.

SolarWinds Platform polls values based on OIDs from the SolarWinds MIB database. There can be OIDs you might want to poll, which are not polled by SolarWinds Platform by default. If these OIDs are in the SolarWinds MIB database, you can create either an UnDP, or use Device Studio to poll for that value, and add support for vendors and technologies that are not natively supported by SolarWinds Platform.


i SolarWinds Platform products poll devices based on OIDs according to the device vendor's MIB. These OIDs must be included in the SolarWinds MIB database. When you create custom pollers, you select OIDs from the SolarWinds MIB database. To poll an OID which is not in the SolarWinds MIB database, define it manually. See [Define object identifiers \(OIDs\) that do not exist in the SolarWinds MIB database](#).


With Device Studio pollers you can:

- Poll devices that do not support any of the OIDs polled for by SolarWinds pollers.
- Poll devices that return incorrect data when polled by SolarWinds pollers.
- Override polled values to display custom static values.

Device Studio technologies in NPM

Device Studio supports a number of technologies. Each technology has a defined set of properties that you can monitor on your devices. The technology you select defines how the polled data are processed, stored, and presented.

technology	usage
CPU & Memory	<p>CPU & Memory is used for collecting data about the CPU and memory load of single processor systems.</p> <p>It provides data to resources related to CPU and memory, such as Average CPU Load & Memory Utilization, Min/Max/Average of Average CPU Load, or Top CPUs by Percent Load.</p> <p>To use this technology, specify a single OID that reports a value from 0 to 100.</p> <p>For example, if a natively polled OID returns incorrect CPU load values, search for an OID that returns a possible value. In the case of CPU load, the load can vary between 0% and 100%, so you must look for an OID that returns a value between 0 and 100.</p> <div data-bbox="370 961 1513 1066" style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p> To determine the OID, consult your device vendor, or carry out a search for an OID that reports the correct value for your device.</p> </div>
Multi CPU & Memory	<p>Multi CPU & Memory provides data to the same resources for multiprocessor systems as the CPU & Memory technology provides for single processor systems.</p> <p>For example, if a natively polled OID returns incorrect CPU load values, search for an OID that returns possible values. In the case of CPU load, the load can vary between 0% and 100% on each CPU core, so you must look for an OID that returns a table of values between 0 and 100, where each row corresponds to a CPU core.</p>
Node Details	<p>Node Details provides data for the Node Details resource, and can be used for devices that are not supported out of the box.</p> <p>To use this technology, specify custom OIDs to poll for Vendor, Machine Type, Software Version, and other data. You can also define custom text to be used instead of the polled value.</p>

 Pollers using other polling technologies, such as VLAN and VRF, are also displayed in the Manage Pollers view. However, it is not possible to create pollers using these technologies in Device Studio.

Data sources used in Device Studio

By creating Device Studio pollers, you can define custom polling definitions in a way that allows you to view the defined set of pollers and the data polled by them as fully integrated entities in the SolarWinds Platform Web Console, including charts, alerts, and reports.


You can define a set of polled data, and then associate these data points with monitored nodes.

The data source you use for polling devices can be:

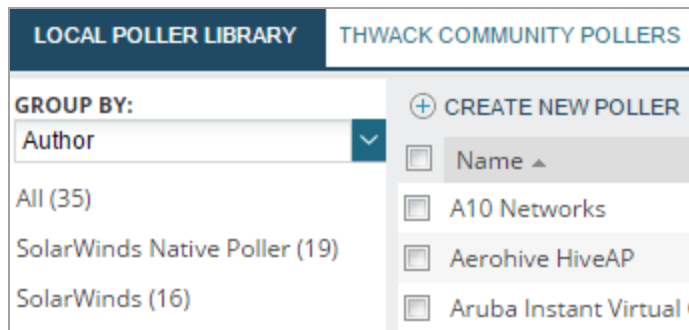
- A polled value or values reported by a device on an OID.
- A calculated value that results from the transformation of polled values.
- A fixed value in the form of a constant number or text. This value is not polled. For example, you can specify the software version of your device as 15.

Create pollers in Device Studio for NPM

To poll unique devices or technologies not supported by default, create a custom poller.

 Reduce the number of Unknown nodes by creating a custom poller.

1. Click Settings > All Settings, and in the Node & Group management grouping, click Manage Pollers.
2. Click Create New Poller.



3. Select a polling technology, type the Poller Package Name, select a test node, and click Next.

i When you are creating the poller, the test node is polled to provide a preview of the results returned by the poller.

Define General Parameters

Tell us what type of data (or technology) you want to poll, select a test node to verify realistic results, and

Technology *: **i** CPU & Memory

Poller Package Name *: Average CPU and Memory Utilization

Test Node *: **i** ● resp_HWH rainbow walk with all statuses [Change test node...](#)

- On the Specify Data Source tab, select a metric you want to define, and click Define Data Source.

Specify Data Source for Pollers

Select "Define Data Source" to choose an object

[Edit Data Source](#) [Poll Current Value](#)

Poller	
<input type="checkbox"/>	Used Memory
<input checked="" type="checkbox"/>	Free Memory
<input type="checkbox"/>	CPU Load

- On the Pick Object Identifier screen, type the OID, or search the MIB database. For information about manually defining OIDs, see [Define object identifiers \(OIDs\) that do not exist in the SolarWinds MIB database for NPM](#).

Browse SolarWinds MIB Database

Available OIDs: 1.3.6.1.2.1.2.1

1 result for "1.3.6.1.2.1.2.1" in name, description, MIB fields:

<input checked="" type="checkbox"/>	ifNumber(1)
-------------------------------------	-------------

Selected OID


ifNumber on ● [Progress Bar]

122

Description:
The number of network interfaces (regardless of their current state) present on this system.


OID: 1.3.6.1.2.1.2.1

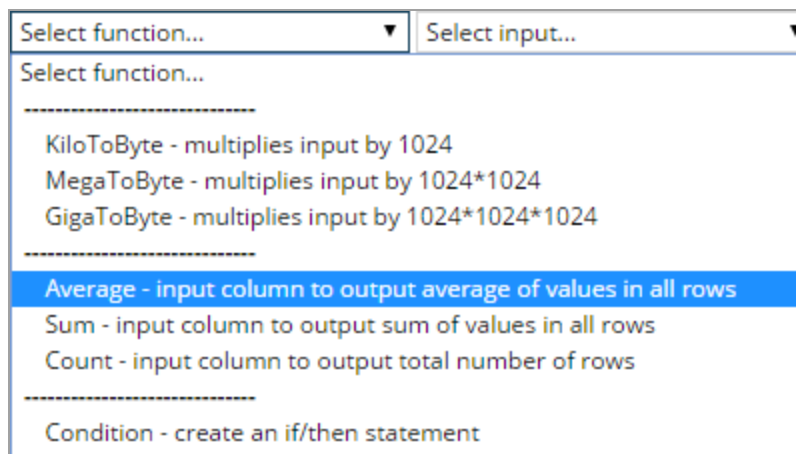
- If necessary, click Add Calculated Value to transform the multiple returned values into a single value, or select a different OID.

 Transforming multiple values to a single value is useful if, for example, the device returns CPU usage as a table of four values (with one value for each CPU core), but you want to use a single value for CPU usage. In this case, you can use the Average function to convert the table of values into a single value.

For more information, see [What is a formula in NPM?](#)

- In the Create a Calculated Value screen, select a function, select an input from the lists, and click Test. You can also define a constant value, for example, if you are creating a CPU and memory poller, and the device you want to poll only supports CPU values.

 Continuing with the previous example, to create an average value out of the four reported values, select the Average function and specify the input values.



Select function... Select input...

Select function...

KiloToByte - multiplies input by 1024

MegaToByte - multiplies input by 1024*1024

GigaToByte - multiplies input by 1024*1024*1024

Average - input column to output average of values in all rows

Sum - input column to output sum of values in all rows

Count - input column to output total number of rows

Condition - create an if/then statement

For more information, see [Formulas used for transforming Device Studio poller results in NPM.](#)

- After testing whether the value is as expected, click Yes, the Data Source Is Reasonable.
- To automatically test the poller on newly added nodes, select Automatically poll nodes during network discovery, and click Next. The test determines whether the Device Studio poller can be assigned to the newly added node.

Network Discovery Settings

- Automatically poll nodes during network discovery, add node,**
This poller will be enabled on nodes where OIDs are polled succes

10. On the Summary tab, review the poller package settings, and click Submit.

Review Your Poller Package Settings

Technology: CPU & Memory

Poller Package Name: Average CPU and Memory Utilization

Description:


Tags:

Author:

The poller is now available in the list of pollers, and you can assign it to nodes.

Define object identifiers (OIDs) that do not exist in the SolarWinds MIB database for NPM

1. On the Pick Object Identifier screen, select the check box under Manually Define Object Identifier (OID).

 Not sure how to get here? See [Create pollers](#) for more details.

2. Type the name and OID.
3. Select the SNMP get type. See [What is the SNMP Get Type in NPM?](#) for more information.
4. Click Poll Current Value From Test Node.

What is the SNMP Get Type in NPM?

The SNMP Get type defines the type of query you have to run to retrieve the appropriate information. You can retrieve scalar values by using either GET or GET NEXT, and you can retrieve values from a particular column in a table value by using GET TABLE.

See [Orion SNMP Get Statements](#) for more details.

 For table records, only the first five values are returned.

What is a formula in NPM?

Values polled by a custom poller are often better understood after a calculation transforms the value to a different format. For example, if a poller returns values in MB, you might want to work with the values presented in GB. The calculations and transformations that are used to manipulate poller results are called formulas.

Two types of values or data sources are available:

- Scalar: one value
- Tabular: column of values

When a new data source is created, the name is generated automatically according to the syntax:

<Property name>Formula<Number>

For example: UsedMemoryFormula1

Formulas used for transforming Device Studio poller results in NPM

formula	description
KiloToByte	Multiplies input by 1024
MegaToByte	Multiplies input by 1024 x 1024
GigaToByte	Multiplies input by 1024 x 1024 x 1024
Average	Returns the average of values from the input columns
Sum	Returns the sum of values from the input columns
Count	Returns the total number of input columns
Condition	Creates an if/then statement
Truncate	Rounds the input decimal number up or down to an integer
Length	Returns the number of characters in the input string
Replace	Replaces the content in the string
IndexOf	Returns the position in the string
SubString	Defines the section of the string of interest

The formulas are divided into three main groups.

type of formula	description
Transformations	Transform data between different units. For example, transform megabytes to bytes.
Aggregations	Transform the values from the input table columns to scalar values. For example, transform the values from the input columns into the average of values.

type of formula	description
Conditions	Transform values according to a logical formula according to the following syntax: if(logical formula), (action to perform if formula is true), (action to perform if formula is false)

Example syntax

SubString

The SubString(,,) calculation takes the following syntax:

```
SubString ([formula], index start, length)
```

For example, if your input is "test", the output will be "es" if you use the following calculation:

```
SubString ([UsedMemoryFormula], 1, 2)
```

As another example, if your input is "test", the output will be "st" if you use the following calculation:

```
SubString ([UsedMemoryFormula1], 2, 2)
```

Replace

The Replace(,,) calculation takes the following syntax:

```
Replace([formula], search string, replacement string)
```

For example, if your input is "test", the output will be "resr" if you use the following calculation:

```
Replace ([UsedMemoryFormula1], "t", "r")
```

Use Regex formulas for transforming poller results

When you define a Regex formula, use the following syntax:

```
Regex([variable], "regular expression")
```

Examples of correct formulas include:

- `Regex([description], "^[a-zA-Z]*[^\s]*")`
- `Regex([description], "(V.[^\s]*)")`
- `Regex([description], "(T.*")`
- `Regex([description], "(C.[^\s]+)")`

Limitations of Regex formulas

When you define a Regex formula, the input string from the test device is interpreted up until the nearest `\r` (new line) character.

The following methods of defining Regex formulas are not supported:

- A backslash sequence for special characters such as the following: (,), {, }, .
- Grouping regular functions such as the following: \w, \W, \s, \S.
- Defining multiple conditions in square brackets such as the following: [^ , -].

Test Device Studio pollers in NPM

A Device Studio poller may not always be seamlessly supported by the device it is tested on. For example, errors occur if the OID the Device Studio poller polls for is not supported by the device, or if the returned value is not of the expected data type defined by the Device Studio poller.

To get the Device Studio poller working in your environment, try the following:

- Test the Device Studio poller on a different node.
- If the device you use for testing is not fully compatible with the Device Studio poller, upgrading the firmware of your test device might help.
- Modify the Device Studio poller to suit the devices you have. For example, you can modify the OID that is used to poll the device.



- Modifying Device Studio pollers this way requires familiarity with the MIB database structure.
- Some of the pollers provided by SolarWinds cannot be modified with Device Studio. You can only modify the poller definition of these pollers in a text editor.

Monitor devices with NPM using THWACK community pollers

Apart from creating your own Device Studio pollers, you can also import pollers provided by contributors of the [THWACK community](#).

The THWACK community pollers are available in the SolarWinds Platform Web Console under Manage Pollers > THWACK Community Pollers. The list is updated automatically every 30 minutes, and it contains the device pollers that have been made available on THWACK, under Network Performance Monitor > NPM Content Exchange > Device Pollers > Documents.

You can group the available pollers according to tags, author, or technology. Click the name of a device poller to view the description of the poller.

To verify whether a poller suits your specific device, test the poller before importing it.


Test THWACK Device pollers

1. Select the THWACK community poller from the list, and click Test Device Poller.
2. Type your THWACK credentials, and click Submit.
3. Select an SNMP node for testing, and click Test Poller.

After the test is finished, you can directly assign the device poller to the test node.

Import Device pollers from THWACK

1. Select the THWACK community poller from the list, and click Import Device Poller.
2. Type your THWACK user credentials, and click Submit.
3. After the import is finished, the poller will be available in the Local Poller Library, and you can assign it to a device. For more information, see [Assign Device Studio pollers to monitored devices in NPM](#).


 If the poller was already imported earlier, you can either overwrite the existing poller, or create a new one.

Import THWACK community pollers to an environment without Internet connection


The THWACK community pollers are only updated automatically if you have a working Internet connection. To import THWACK community pollers to an environment that does not have an Internet connection, download the pollers from a computer which can access the Internet, save them to a portable drive or a USB drive, and import them manually.

Export Device Studio pollers to the THWACK community

1. On the Manage Pollers screen, click the Local Poller Library tab, and select a poller.

 You can export Device Studio pollers that you created, but you cannot export pollers that are provided by SolarWinds.

2. Click Export, and select Export to Thwack.
3. Type your THWACK user credentials, and click Submit.

 If you already logged in to THWACK from the SolarWinds Platform Web Console during the same session, you do not have to enter your credentials again, and the Device Studio poller will be exported immediately.

The Device Studio poller will be available on THWACK, in the Network Performance Monitor > NPM Content Exchange > Device Pollers > Documents section.

Why can't I connect to THWACK from NPM ?

Your SolarWinds Platform server must be able to open internet connections to connect to THWACK.

If the connection is blocked by a firewall or a proxy:

- The list of shared pollers cannot be retrieved from THWACK
- Any operation that relies on communication with THWACK, such as the upload or download of a poller, will fail.

Check your firewall and [proxy settings](#) to make sure that your SolarWinds Platform server can connect to the internet.

Assign Device Studio pollers to monitored devices in NPM

Specify devices on which you want to poll the statistics defined by the poller.

1. On the Manage Pollers page, select a poller, and click Assign.
2. Select the node you want to assign the poller to.
3. If the node has not been scanned yet, click Scan Now.
4. If the scan result is a match or a multiple match, select the node, and click Enable Poller.

 You can only scan SNMP nodes whose status is Up.

Scan monitored objects in NPM to verify if the OIDs match

When a monitored node is scanned, the OIDs of the monitored node and the OIDs specified in the poller are compared to see if they match.

These scenarios are possible:

- If the OIDs do not match, the scan returns a result indicating the mismatch, and the poller cannot be assigned to the monitored node. See [Test Device Studio pollers in NPM](#).
- If the OIDs match, and there is no other poller supporting the specific technology, then the poller is automatically enabled on the node.
- If the OIDs match, but there is already another poller for the technology, the new poller is not enabled. You can enable the poller manually. See [Assign Device Studio pollers to monitored devices in NPM](#).


Troubleshoot NPM issues with Performance Analysis dashboards

With complex networks consisting of cloud, hybrid IT, virtualization, storage area networks, and so on, multi-faceted IT issues can be difficult to pinpoint and diagnose. When an issue surfaces, for example a badly performing application or server, the investigation can take significant time to locate the core issue. The problem could be in storage, network connectivity, user access, or a mix of resources and configurations.

To investigate the issue, create troubleshooting projects with the Performance Analysis (PerfStack™) dashboard that visually correlate historical and real-time data from multiple SolarWinds products and entity types in a single view.

With Performance Analysis dashboards, you can do the following:

- Compare and analyze multiple metric types in a single view, including status, events, and statistics.
- Compare and analyze metrics for multiple entities in a single view, including, nodes, interfaces, volumes, applications, and more.
- Correlate data from across the SolarWinds Platform on a single, shared time line.
- Visualize hybrid data for on-premises, cloud, and everything in between.
- Share a troubleshooting project with your teams and experts to review historical data for an issue.

 PerfStack is designed for expert users to quickly sift through data. If you need help to select relevant metrics or more automated troubleshooting, consider using [AppStack](#) or [NetPath™](#).

Troubleshoot intermittent network slowdowns with NPM

This topic provides an example of how you can troubleshoot an issue where a router drops or delays packets. Bandwidth issues are already ruled out.

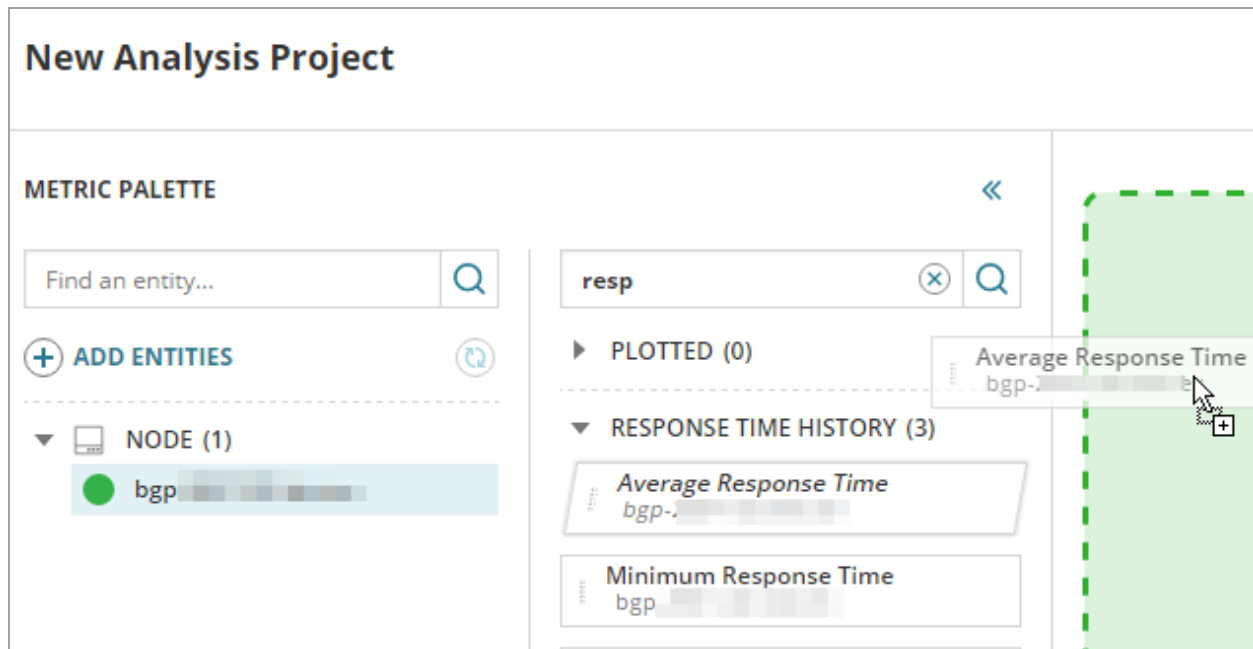
This example describes creating a Performance Analysis dashboard (PerfStack™) that correlates the following metrics to investigate the cause of the network slowdown:

- Average Response Time (ms) and Percent Loss to assess the symptoms of the problem: how fast is the network getting your traffic to the server and how much of your traffic is lost.
- Average CPU Load and Average Percent Memory Used to verify whether the system-wide resources are not overloaded and thus causing the issue.


- Buffer Misses due to No RAM and Total Number of Buffer Misses to see if high RAM usage impacted how the router allocates RAM for traffic forwarding.
- Small, Medium, Big, Large, and Huge Buffer Misses to investigate which size packets are causing the generic buffer misses you found.

To troubleshoot the network slowdown:

1. Click My Dashboards > Home > Performance Analysis.
2. Click Add Entities, add the node to the Metric Palette, and select the node.
3. Expand the Response Time History metric group, and drag Average Response Time and Percent Loss to a new chart.




The graph shows that both average response time and percent packet loss are low. The response time ranges from 2 to 7 ms.




4. To investigate the load, add Average CPU Load and Average Percent Memory metrics to a new graph. Add the memory load in percent to be able to compare the two metrics.

In this example, the CPU load is low (about 5%). The used Average Percent Memory is 18.7%, and so it is not the cause of the issue. The issue might be caused by the allocation of buffer for sending packets.



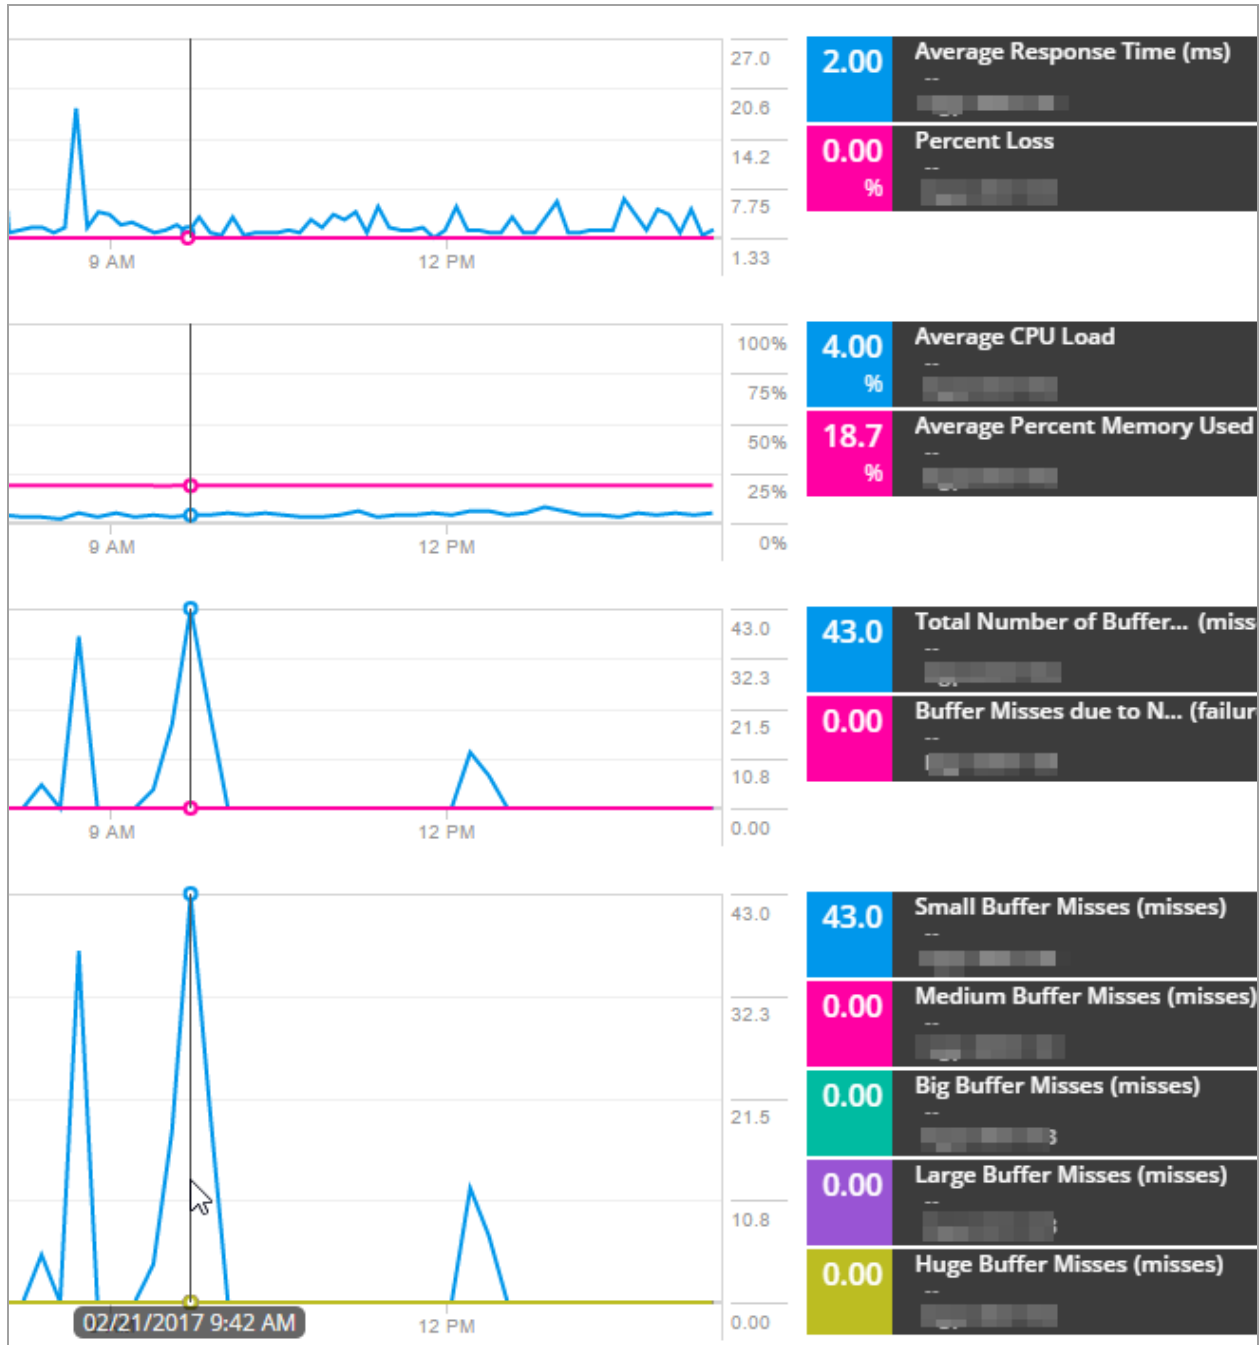
5. Investigate whether it is a buffer issue:
 - a. Add the Total Number of Buffer Misses to a new graph. In this scenario, there are spikes in the Total Number of Buffer Misses graph.
 - b. To find out whether they are caused by a lack of RAM, add Buffer Misses due to No RAM to the Buffer Misses chart.

In this example, Buffer Misses due to no RAM are zero, so insufficient RAM is not the cause of the spikes in buffer misses.



6. To investigate the buffer misses, add metrics for all types of buffer misses into a new graph. The following example shows the size of the packets that are being dropped.

In this scenario, all buffer misses values are zero except for small buffer misses. Small buffer misses are thus the cause of the issue. To resolve the issue, increase the buffer size for small buffers, or configure the router to have a larger boundary.



- To share the dashboard, simply share the URL. The URL contains all data to recreate the view.
- To keep the featured metrics in the dashboard for later use, click Save and enter a name for the Performance Analysis project. You can click Load to view the project again.

Troubleshoot slow resources in a branch office with NPM

This topic provides an example of how you can troubleshoot an issue where employees in a branch office are complaining about slow resources.

Use Performance Analysis (PerfStack™) to correlate metrics, find out the cause of the issue, and resolve the issue.

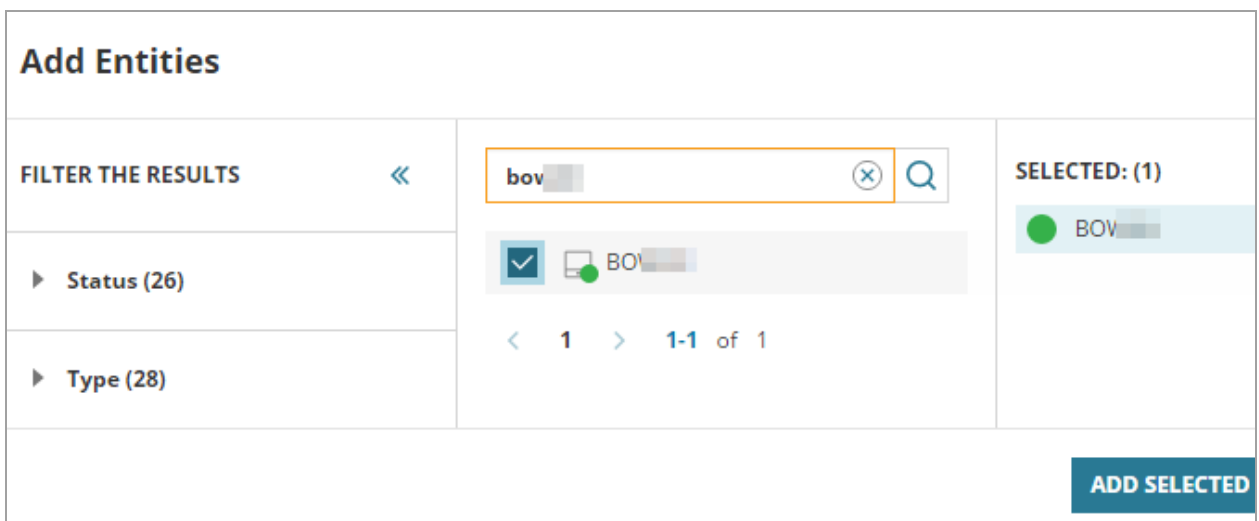
This type of problem in this environment is often due to problems on the WAN interface. Use PerfStack to analyze the interface.

In this example, the Performance Analysis project includes the following metrics:

- Status to know whether the interface is up or whether the status is changing.
- Average Transmit bps to investigate the amount of traffic.
- Transmit Percent Utilization to investigate how close the interface is to being fully saturated.
- Transmit Discards and Percent Discards to investigate discarded data.

Start your analysis with the Branch Office WAN router (BOWAN).

1. Click My Dashboards > Home > Performance Analysis.
2. Locate the interface:
 - a. In the Metric Palette, click Add Entities.
 - b. Enter the router name into the entity filter, select the router, and click Add Selected Items.

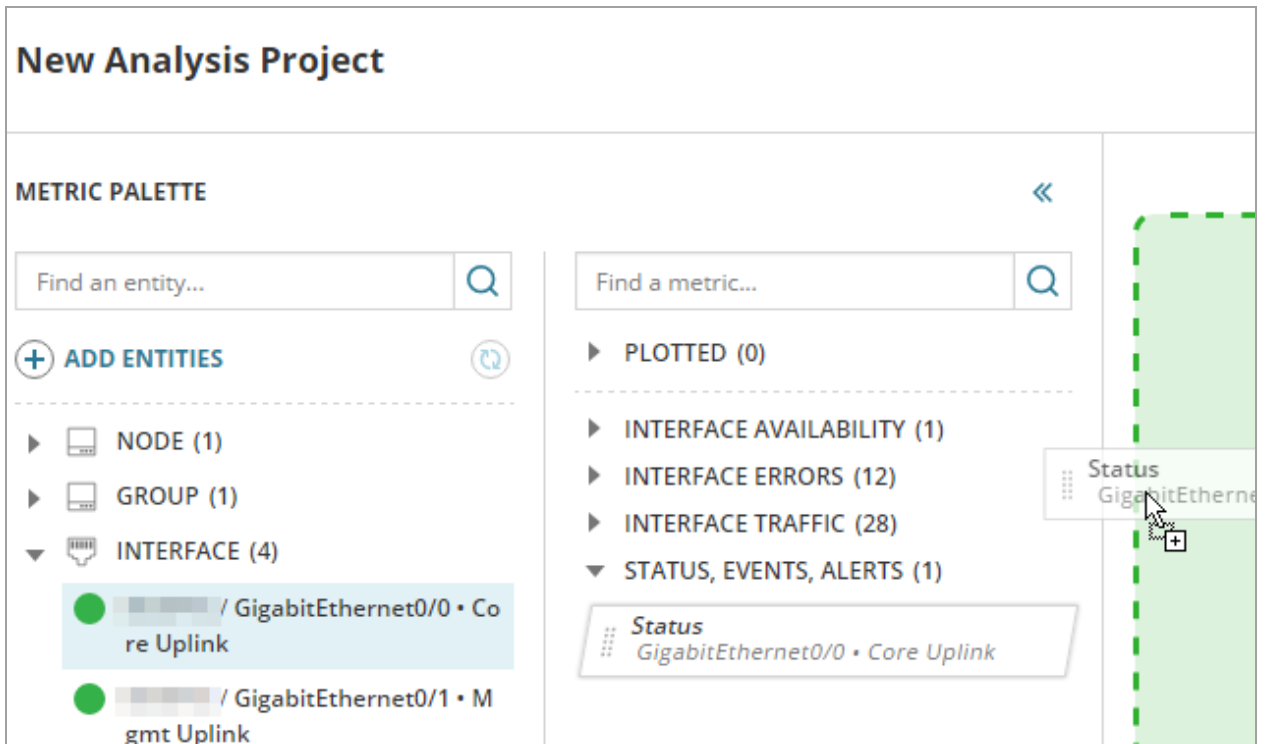


- c. Expand nodes, position the mouse over the router, and click the Add Related Entities button. Related entities, such as groups and interfaces, display in the Metric Palette.




3. Add status information to the dashboard:

- a. Select the Core Uplink interface in the Metric Palette.
- b. Expand Status, Event, Alerts in Metrics and drag Status into the chart area.



4. Investigate the utilization of the interface:


- a. Drag Average Transmit bps into a new chart. The traffic is not very high, it is around 9 Mbps.

 To filter available metrics, type a part of the metric name into the field, for example transmit.

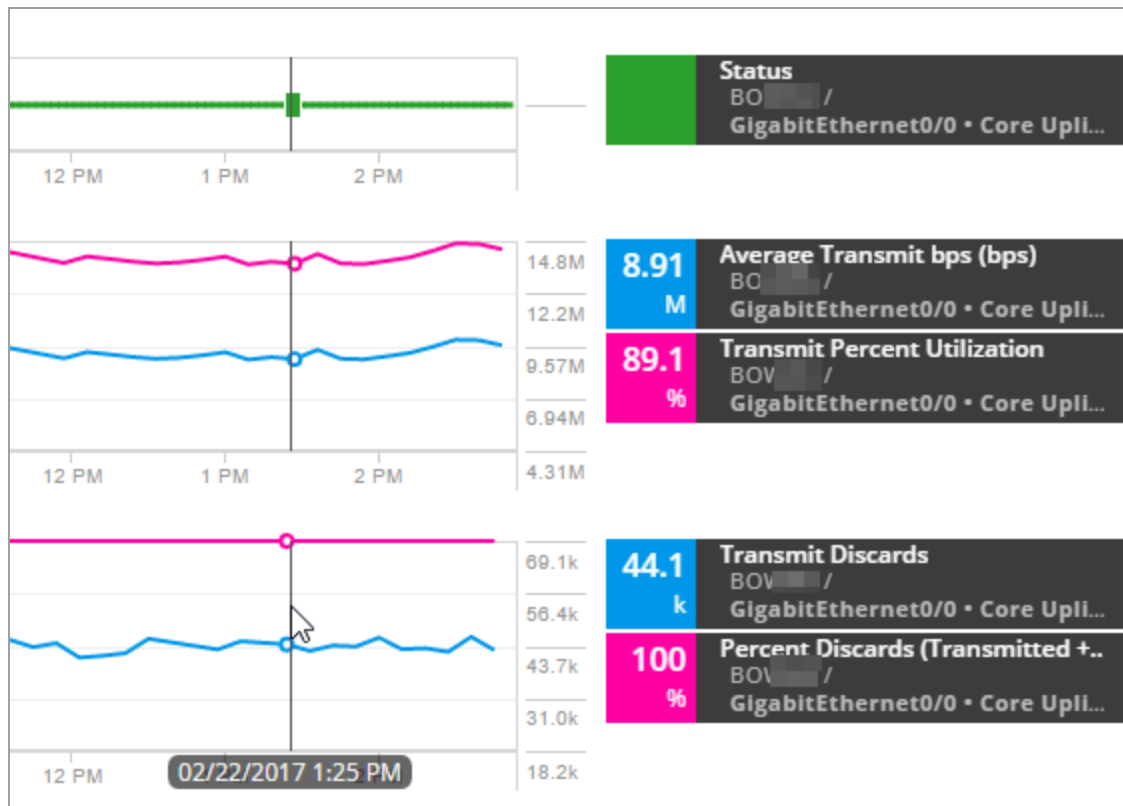
- b. Drag Transmit Percent Utilization into the chart. The percent utilization is high, it is

around 95%.

The two metrics indicate that 9 Mbps is approximately 95% of the interface utilization.



- To understand if high utilization is resulting in discarded frames, drag Transmit Discards and Percent Discards into a new chart. The interface discards 100% of the traffic. Discards can be caused by hardware or configuration issues.



6. To verify the configuration and acting state on the port, log in to the device and check the configuration. For example, go to the BOWAN details view, and click telnet in the Node Details resource.

In this example, you discover that the interface is configured to run at 10 Mb. However, it is a Gigabit interface and it should be running at 1,000 Mbps. Correct the configuration to resolve the issue.

- To share the dashboard, simply share the URL. The URL contains all data to recreate the view.
- To keep the featured metrics in the dashboard for later use, click Save and enter a name for the Performance Analysis project. You can click Load to view the project again.