NOZOMI NETWORKS



Guardian
User Guide

N2OS v26.6.0 — 2026-11-05

Legal notices

Information about the Nozomi Networks copyright and use of third-party software in the Nozomi Networks product suite.

Copyright

Copyright © 2013-2026, Nozomi Networks. All rights reserved. Nozomi Networks believes the information it furnishes to be accurate and reliable. However, Nozomi Networks assumes no responsibility for the use of this information, nor any infringement of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under any patent, copyright, or other intellectual property right of Nozomi Networks except as specifically described by applicable user licenses. Nozomi Networks reserves the right to change specifications at any time without notice.

Third Party Software

Nozomi Networks uses third-party software, the usage of which is governed by the applicable license agreements from each of the software vendors. Additional details about used third-party software can be found at https://security.nozominetworks.com/licenses.

Contents

Chapter 1. Introduction 9
Guardian overview 11
Chapter 2. Sensors 13
List 16
Map 18
Graph 20
Do a force update on sensors 21
Allow a sensor 22
Export a list of sensors 23
Upload a map 24
Arc sensor configuration 25
Deployment settings 25
Configure an Arc sensor 33
Manage the Arc dependencies of a sensor 37
Chapter 3. Alerts 39
Standard mode 42
Expert mode 43
Alert deduplication 44
View alerts in standard mode 45
View alerts in expert mode 45
Closing alerts 46
Actions menu 49
Open the Actions menu in standard mode 51
Open the Actions menu in expert mode 51
Configure an alert 52
Acknowledge an alert 53
Close an alert 53
Download a trace 56
Download a malicious file 56
Edit a note for an alert 57
View a diff from an alert 59
Navigate from an alert 59
Edit a playbook associated with an alert 60
Details page 61
Chapter 4. Assets 65
List 67
Diagram 69
Configure an asset 70

Generate a PDF report of an asset....71
Navigate from an asset....72
View more details for an asset....73
Details page....74
Overview....76
Sessions....77
Alerts....78
Software....79
Vulnerabilities....80
Variables....81
Network graph....82
Physical graph....83
Hardware components....84
Chapter 5. Queries....87
Data sources....91
Basic operators....94
Commands....96
Nodes-specific commands reference....102
Link-events-specific commands reference....104
Functions....105
Examples....109
Pie chart....109
Column chart....109
Where with multiple conditions in OR....110
Bucket and history....111
Join....111
Compute the availability history....112
Complex field types....113
Chapter 6. Smart Polling in Guardian....117
Overview....120
Plan creation....122
Actions menu....141
Add a node from the Network page....158
Log level customization....159
Smart Polling results....160
Activity Log....163
Chapter 7. Arc in Guardian....165
Architecture....168
Arc in Guardian....169
Deployment....170
Advanced....171
Execution details....173
Deployment settings....173

Upstream ... 175
Endpoint ... 176
Network sensor ... 178
Other ... 180
Node points ... 182
Dependencies ... 183
Apply an Arc dependency action to Arc sensors ... 185

Chapter 8. Network ... 187

Nodes ... 190
Configure a node ... 191
Show alerts for a node ... 192
Show requested traces for a node ... 192
Request a trace for a node ... 193
Manage learning for a node ... 195
Navigate from a node ... 196
Add a node to a Smart Polling plan ... 197
Links ... 200
Link events ... 200
Link availability ... 201
Configure a link ... 203
Show alerts for a link ... 204
Show requested traces for a link ... 204
Request a trace for a link ... 204
Show events for a link ... 206
Show captured URLs for a link ... 206
Manage learning for a link ... 207
Navigate from a link ... 208
Physical links ... 210
Sessions ... 211
Show requested traces for a session ... 211
Request a trace for a session ... 212
Navigate from a session ... 215
Traffic ... 216
Details page ... 217

Chapter 9. Graph ... 219

Network ... 222
Assets ... 225
Physical ... 229
Graph controls ... 233
Layout ... 233
Magic wand ... 235

Chapter 10. Process ... 237

List ... 240

Protocol connections....241
Settings....242
Configure a variable....243
View the details of a variable....246
Favorite a variable....248
Navigate from a variable....249
Chapter 11. Reports....251
Dashboard....254
Management....255
Generated....258
Scheduled....259
Settings....260
Create a report....261
Generate a report....263
Download a report....265
Delete a report....265
Add a folder....266
Edit a folder....267
Delete a folder....268
Import a schema....268
Export a schema....269
Upload a custom logo....270
Configure SMTP settings....272
Filter a report globally....274
Add a widget to a report....277
Filter a report with a widget....279
Chapter 12. Assertions....281
Assertion operators....285
Save an assertion....286
Customize assertion name and description....288
Edit an assertion....290
Configure an assertion....291
Configure an assertion on links....292
Configure an assertion on variables....294
Chapter 13. Time machine....297
Time machine diffs....301
Load a snapshot....303
Request a diff....304
Reload the diff operation progress....305
View a diff from an alert....306
Chapter 14. Vulnerabilities....307
Assets....310

List 311
Stats 312
Details page 313
Chapter 15. Administration 315
Administration page 317
Chapter 16. Personal settings 319
Other actions 322
Change your password 323
Edit an OpenAPI key 324
Generate an OpenAPI key 325
Clear your personal settings 327
Request a continuous trace 328
Request a custom trace 329
Show requested traces 331
Download a requested trace 332
Zone filters 333
Glossary 335



Chapter 1. Introduction



Guardian overview

Guardian is the main Nozomi Networks sensor.

Asset discovery

Guardian gives you the ability to automatically track your *industrial control systems (ICS)*, *operational technology (OT)* and *Internet of Things (IoT)*/*Industrial Internet of Things (IIoT)* assets.

- Highly accurate asset inventory of all communicating devices
- Extensive node information including name, type, serial number, firmware version and components
- Actionable risk assessment insights including security and reliability alerts, missing patches and vulnerabilities

Network visualization

Guardian gives you instant visibility of your entire network. This lets you:

- Have instant awareness of your *OT*/*IoT* networks and normal activity patterns
- Access key data such as traffic throughput, *transmission control protocol (TCP)* connections, and protocols
- Use intuitive dashboards and reports with macro and micro views, plus filtering and grouping

Automated vulnerability assessment

Guardian lets you quickly identify which *ICS*, *OT* and *IoT* devices are vulnerable. This provides:

- Efficient prioritization and remediation
- A faster response with vulnerability dashboards, drill-downs and reports
- Based on the U.S. government's *National Vulnerability Database (NVD)* for standardized naming, description and scoring

Continuously monitor your networks and automation systems. Guardian gives you:

- The ability to continuously monitor all your assets, network communications and supported protocols
- Easy access to summarized *ICS*, *OT* and *IoT* risk information
- The ability to highlight potential reliability issues, such as unusual process values

Anomaly-based detection

Guardian builds a baseline of your environment and uses that knowledge to detect threats such as transferred malware, suspicious communications, unwanted operations, or changes to the network.



Chapter 2. Sensors



*The **Sensors** page lets you view all of the sensors that you have in your system.*

The **Sensors** page has these three tabs:

- List
- Map
- Graph

List

*The **List** page lets you view all of the sensors that you have in your system.*

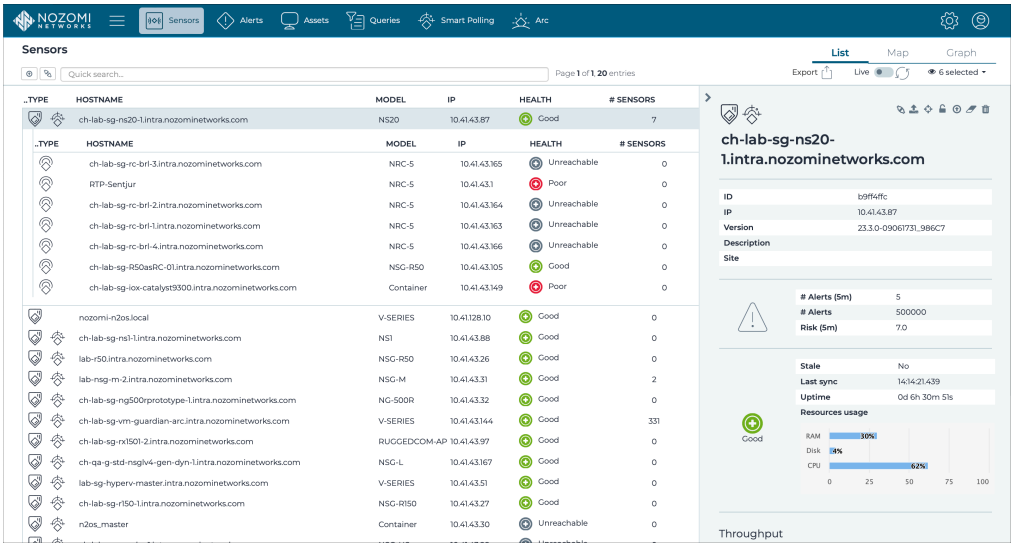


Figure 1. Sensors list

Force update

The force update icon lets you do a force update on the selected sensors.

Allow/disallow

The allow/disallow icon lets you allow or disallow sensors.

Quick search

The quick search field lets you easily do a search on the current page.

Export

The **Export** icon lets you export the current list in either *comma-separated value (CSV)* or Microsoft Excel format.

Download Arc

The **Download Arc** dropdown lets you select, and download, the applicable Arc package for your *operating system (OS)* and architecture.

Live

The **Live** toggle lets you change live view on, or off. When live mode is on, the page will refresh periodically.

Refresh

The **Refresh** icon lets you immediately refresh the current view.

Column selection

The columns selection icon lets you choose which columns to show or hide.

Information pane

The information pane to the right of the list of sensors shows additional information for the selected sensor.

It also lets you do these actions shown below.

Table 1. Sensors list actions

Allow/Disallow sensor	After allowing a sensor, this icon shows: Synchronized data coming from the sensor become part of the Environment of the *Central Management Console (CMC)*. Alerts coming from the sensor can be seen in the **Alerts** section.
Focus on sensor	Allows to filter out only the sensor chosen data, such as Alerts and Environment.
Go to sensor	Connect to a remote sensor directly from the *CMC*. Select this to open a new browser tab to the sensor selected login page. The action is hidden if the *CMC* isn't configured to allow this type of communication between sensors and *CMC*.
Place in map	This action is used to place the sensor on the map.
Toggle version lock	When locked, the sensor will not automatically update its software.
Force update	Even if it is locked, the sensor will automatically update its software, with the version installed on the *CMC*.
Clear sensor data on this machine	Clear all synchronized data at the *CMC* received from the selected sensor. Use this in combination with the clearing of the data on the sensor, and you will be able to restart the synchronization between the sensor and the *CMC* from an empty state
Delete sensor	Clear all data received from the selected sensor and delete it from the list. If the sensor tries to sync with the *CMC* again, it shows as disallowed in the list.

Map

*You can use the **Map** page to upload, and view, a map of the sensors in your environment.*

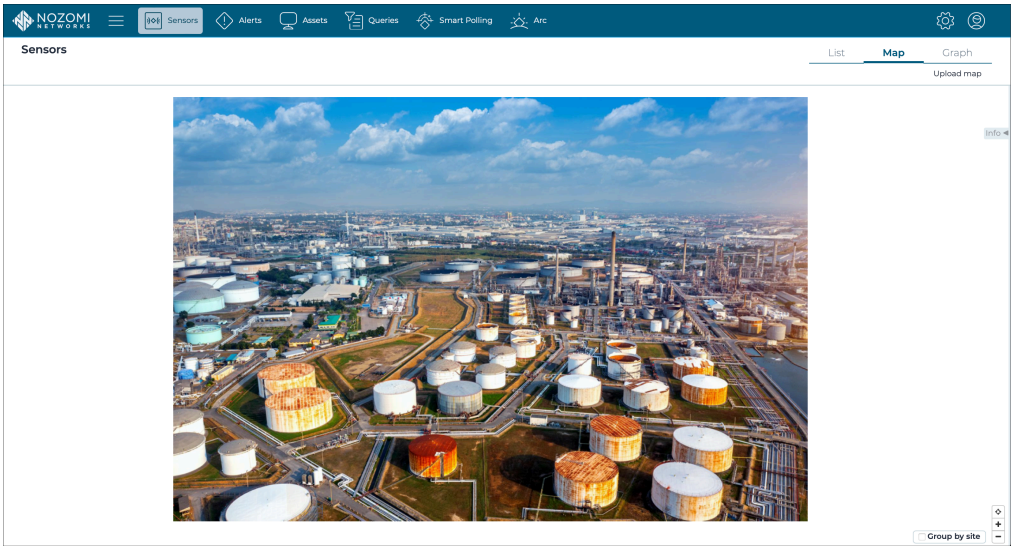


Figure 2. Sensors map

Info(rmation) pane

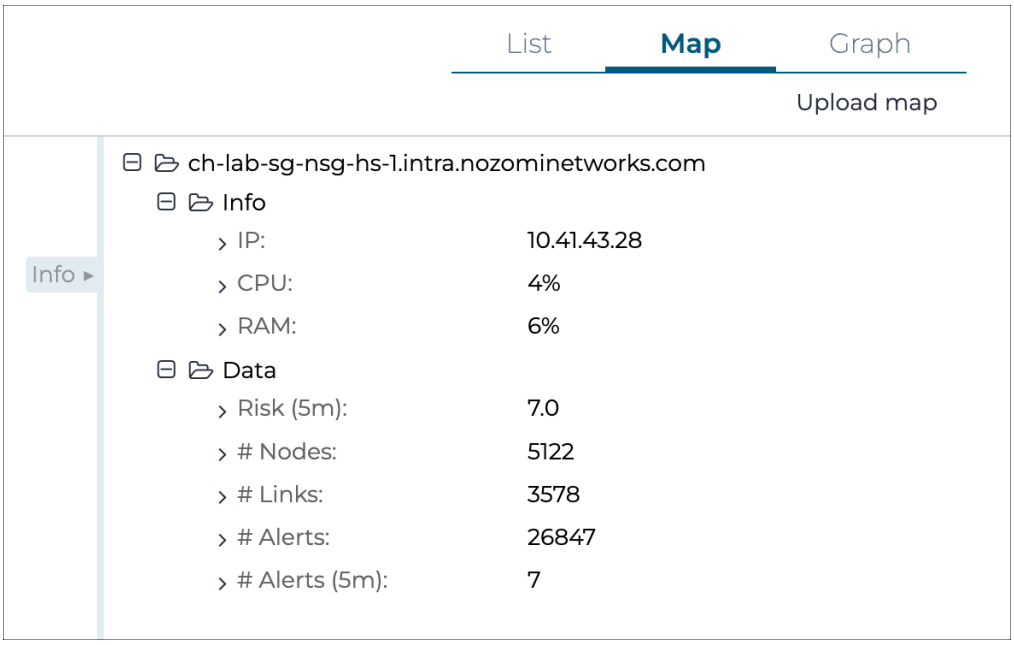


Figure 3. Info(rmation) pane

The **Info** pane lets you view information for the related sensor. The *identifier (ID)* of each sensor is used in the map to help you identify it. The marker color of the sensor relates to the risk of its alerts.

In the map view, a red indicator to the right of the sensor's *ID* shows the number of the alerts in the last five minutes. This indicator only shows if there are some alerts for the sensor. If the alerts in last 5 minutes increase, the sensor marker will blink for one minute.

If the site of the sensor has been specified in  **> System > General**, it is possible to enable the **Group by site** option, in the bottom right corner of the map view. The sensors with the same site will be grouped to deliver a simpler view of a complex installation.



Note:

The sensors map is also available as a widget.

Graph

*The **Graph** page shows a graphical view of all the sensors in you environment.*



Figure 4. Sensors graph

Do a force update on sensors

If you have disabled auto updates, you can use the force update icon to do a manual update.

Procedure

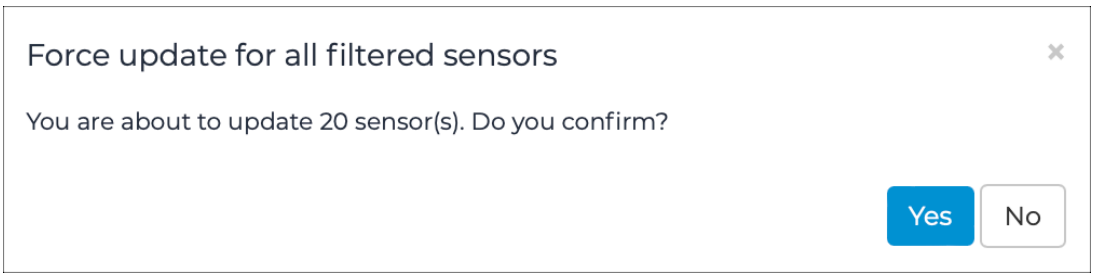
1. In the top navigation bar, select **Sensors**.

   Result: The **Sensors** page opens.

2. In the top left section, select .

   Result: A dialog shows.

3. To confirm, select **Yes**.



Allow a sensor

The allow icon lets you give permission for a new sensor to connect to Central Management Console (CMC).

Procedure

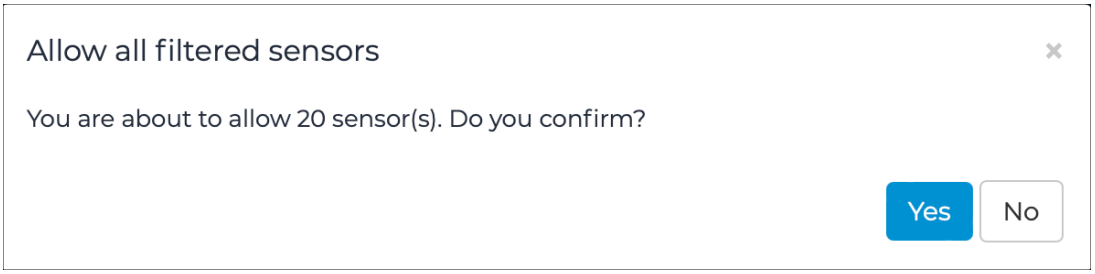
1. In the top navigation bar, select **Sensors**.

   Result: The **Sensors** page opens.

2. In the top left section, select the icon.

   Result: A dialog shows.

3. To confirm, select **Yes**.



Export a list of sensors

You can export a list of the sensors from the current view.

Procedure

1. In the top navigation bar, select **Sensors**.

   Result: The **Sensors** page opens.

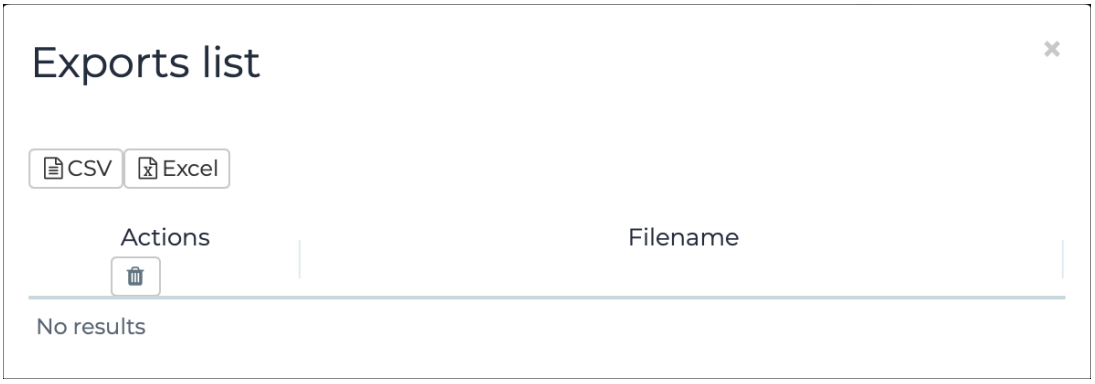
2. In the top right section, select **Export**.

   Result: A dialog shows.

3. Choose an export format:

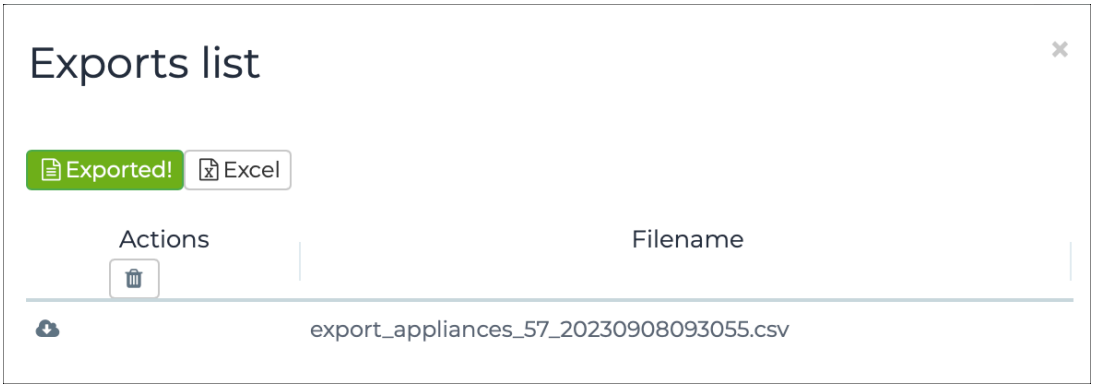
   Choose from:

   - Select **CSV** to create a *CSV* file
   - Select **Excel** to create a Microsoft Excel file



   Result: The file is created and a `Exported` message shows.

4. To download the file, in the bottom-left of the dialog, select



Results

The exported file has been downloaded.

Upload a map

*The **Map** page lets you upload a map of your sensors in your environment.*

Procedure

1. In the top navigation bar, select **Sensors**.

   Result: The **Sensors** page opens.

2. In the top right section, select **Map**.

   Result: The **Map** page opens.

3. In the top right section, select **Upload map.**
4. Select the image file that you want to upload.

Results

The map is uploaded.

Arc sensor configuration

Deployment settings

Configure Arc deployment parameters that will be included in downloaded Arc bundles. These settings define how Arc will behave when deployed from Guardian.

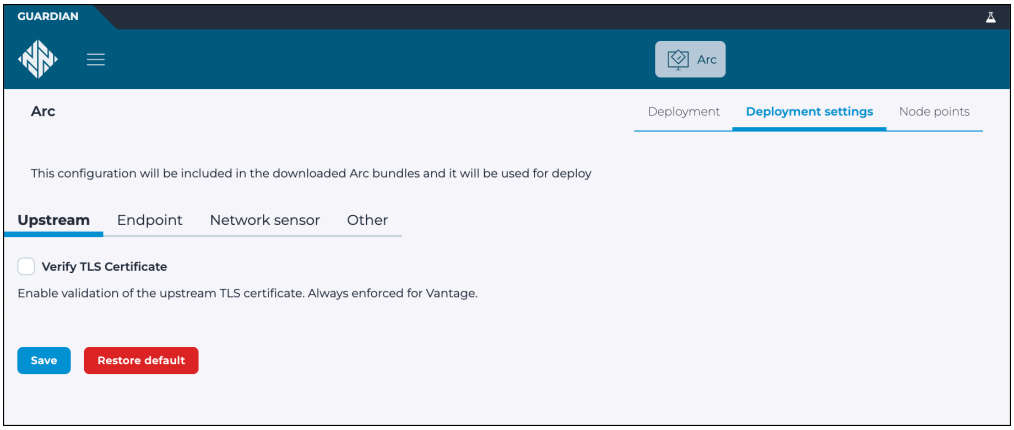


Figure 5. Deployment settings page

The **Deployment settings** page has these tabs:

- Upstream (on page 175)
- Endpoint (on page 176)
- Network sensor (on page 178)
- Other (on page 31)

Restore default

Once the settings have been saved, you can use this button to restore the default configuration.

Upstream

*The **Upstream** page lets you configure the settings for your Arc deployment.*

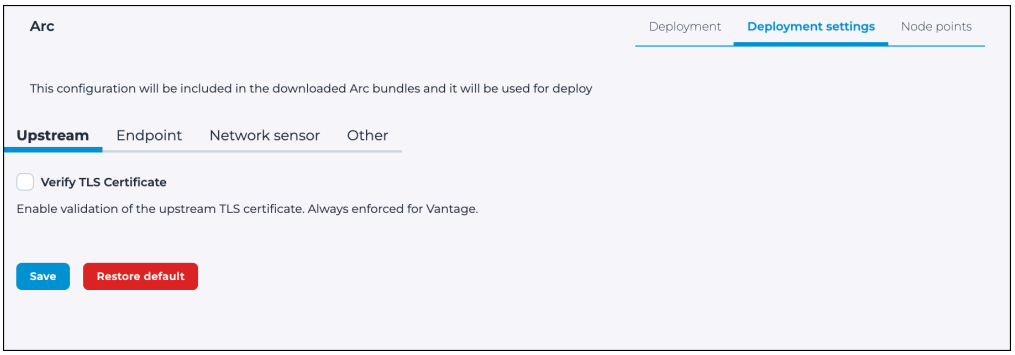


Figure 6. Upstream page

Verify TLS Certificate

Select this option to enforce verification of the upstream *transport layer security (TLS)* certificate when connecting to Guardian. When using Vantage, *TLS* certificate validation is always enforced, regardless of this setting.

Restore default

Once the settings have been saved, you can use this button to restore the default configuration.

Endpoint

*The **Endpoint** page lets you configure the settings for your Arc deployment.*

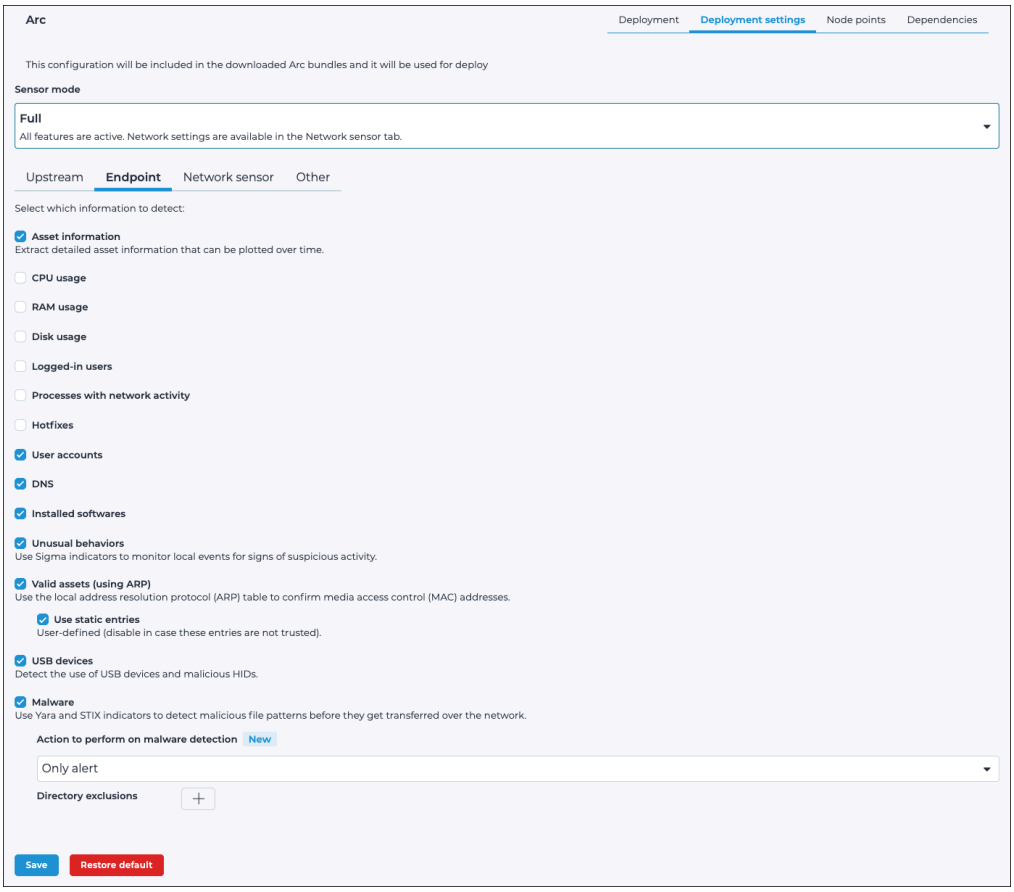


Figure 7. Endpoint page

Asset information

Enable **Asset information** to collect detailed information about the endpoint. Arc collects the following data:

- **CPU usage**: The current processor load of the endpoint.
- **RAM usage**: The amount of memory currently in use on the endpoint.
- **Disk usage**: The current disk space usage on the endpoint.
- **Logged-in users**: The users currently logged in to the endpoint.
- **Processes with network activity**: The running processes that have active network connections, including the local and remote addresses and ports.
- **Hotfixes**: The security patches and updates installed on the endpoint (Windows only).
- **User accounts**: The local user accounts configured on the endpoint.
- **DNS**: The *domain name server (DNS)* servers and configuration of the endpoint.
- **Installed software**: The software packages installed on the endpoint.

Unusual behaviors (Windows only)

This lets you enable/disable Sigma rules for local behavior analysis.

Valid assets (using ARP)

Enable this option to use the local *address resolution protocol (ARP)* table to confirm *media access control (MAC)* addresses for connected assets.

Enable **Use static entries** to include user-defined static *ARP* table entries in the validation. Only enable this if the static entries are trusted.

USB devices (Windows only)

This lets you enable/disable *universal serial bus (USB)* detections.

Malware (Windows only)

The **Malware** checkbox lets you enable protection mode and select from the **Action to perform on malware detection** dropdown. This lets you set the action that Arc will take when it finds a malicious file.

You can choose from these options:

- **Only alert**: Receive an alert with no further action
- **Quarantine**: Move the malicious file to the `quarantine` folder, which is located in the Arc installation folder
- **Delete**: Immediately delete the malicious file. Once deleted, the files cannot be recovered

You can also use the **Directory exclusions** feature to exclude specific directories from malware scanning. Select + to open a file browser and select a directory. You can add as many directories as needed. Arc skips the selected directories when scanning for malware.

Restore default

Once the settings have been saved, you can use this button to restore the default configuration.

Network sensor

*The **Network sensor** page lets you configure the settings for your Arc deployment.*

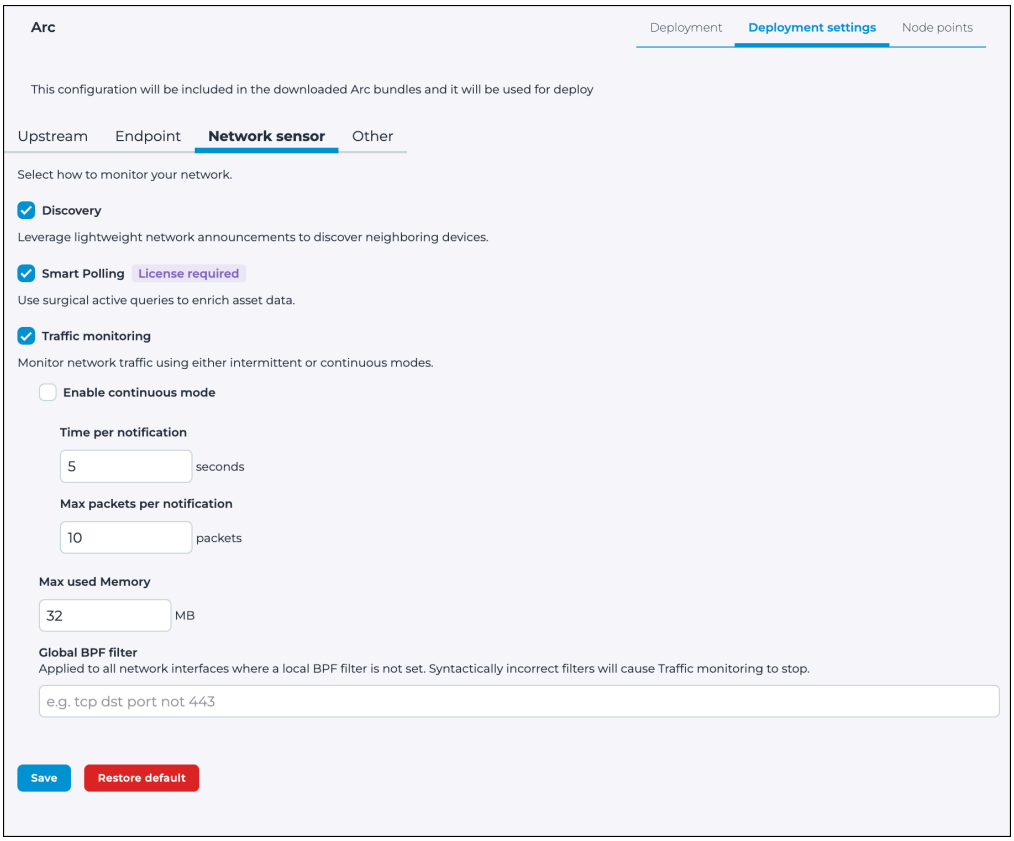


Figure 8. Endpoint page

Discovery

When enabled, this sends out unsolicited lightweight network announcements to discover neighboring nodes.

Discovery uses lightweight protocol-specific broadcast messages to identify network devices. These messages trigger a response from the devices, which includes identity information. The process is repeated at predefined intervals. At each interval, the sensor will identify the suitable network interfaces and send broadcast messages through them to discover devices on each subnetwork connected to the sensor.

Smart Polling

This lets you enable/disable the execution of Smart Polling strategies from Arc. When enabled, this sends out Smart Polling queries following remote requests coming from Guardian to poll assets that Arc can reach, or assets that have been identified with Discovery.

Note:

Smart Polling requires that a Smart Polling license is enabled upstream.

To force Smart Polling from a specific Arc sensor, even when Guardian was the first to monitor a node, you can use a *command-line interface (CLI)* command such as: `vi node 192.168.1.1 capture_device arc[1e6a174c]` In this example, `192.168.1.1` is an *internet protocol (IP)* address of a node you want to poll from a specific Arc sensor. `1e6a174c` are the first eight characters of the Arc sensor *ID* that you want to poll the node with. To find that sensor *ID*, you can select the Arc sensor from the **Sensors** page of your Guardian and read the **ID** field in the right pane. To reset the behavior, you can set the `capture_device` back to the value of the Guardian interface.

Traffic monitoring

When enabled, this checkbox lets you enable/disable traffic monitoring.

Enable continuous mode

This checkbox lets you enable/disable continuous mode. For more details, see **Continuous mode**.

Arc uses two different methods for traffic monitoring:

- Intermittent mode
- Continuous mode

Intermittent mode: This is the default mode, the traffic is monitored, or sniffed, for a duration of 10 seconds at each notify. The purpose of this limitation is to preserve the resources of the host machine, which prevents excessive memory, or *central processing unit (CPU)*, spikes. You can configure these options:

Continuous mode: This mode sniffs traffic continuously from the host's network interface controllers. Depending on the amount of sniffed traffic, continuous mode might utilize more *CPU* and memory on the host. As the traffic is processed upstream, the performance of the remote endpoint is also affected. You can configure:

- **Time [s] per notification**
- **Max packets per notification**
- **Max used Memory (MB)**: this value can be tuned to allow more or less traffic buffering in case the traffic to process exceeds the Arc and network capacity to send it out

Global BPF filter

This field lets you set a Global BPF filter to apply to all the network interfaces. Filters that are applied to single interfaces will take precedence over the global one.

Network interface

This dropdown lets you select a network interface to configure. Each network interface can then be enabled, and be tuned with a monitoring filter.

If you add, remove, or edit the network interfaces on the host, Arc does not automatically add it to the list of sniffing interfaces. For example, if you add a new network card, to enable Arc to use it, you should stop Arc, and then start it again.

Restore default

Once the settings have been saved, you can use this button to restore the default configuration.

Other

*The **Other** page lets you configure the settings for your Arc deployment.*

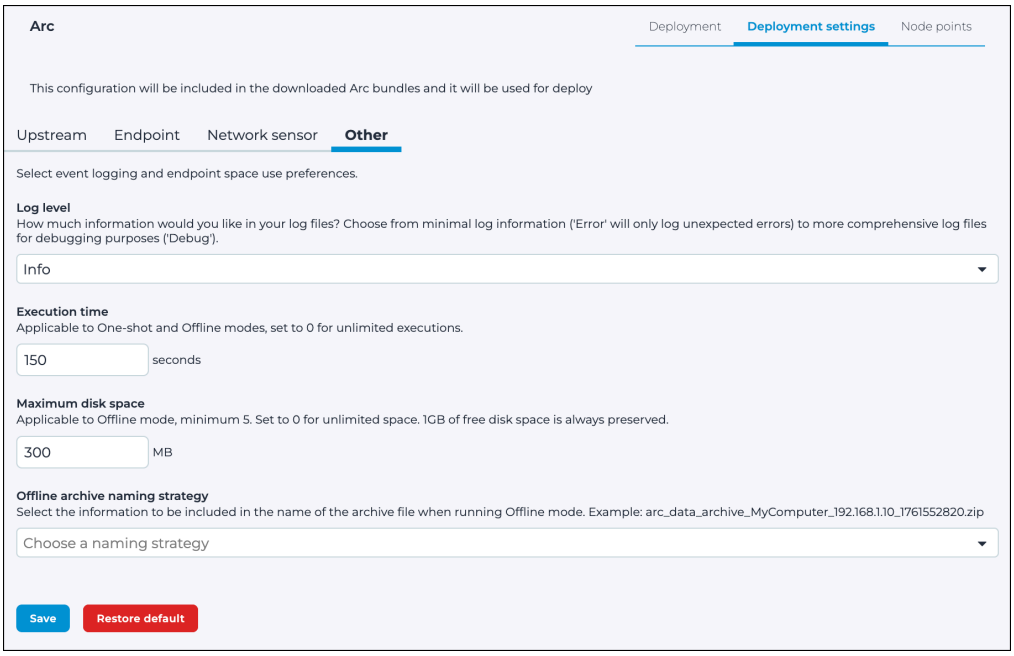


Figure 9. Other page

Log level

This dropdown lets you select the verbosity level for the log files. The options are:

- Debug
- Info
- Warning
- Error

The logging system options have an increasing level of verbosity, from the least verbose to the most verbose. Error < Warning < Info < Debug.

- Error: Creates a minimalistic log, only unexpected errors are logged
- Warning: Creates extra errors that might show on some *OS*s, but that are generally considered as acceptable
- Info: Logs relevant successful events, it shows the program's progress (recommended)
- Debug: Logs extra events that are normally useful for debugging purposes. Given its verbosity it is best to activate it only when debugging activities are involved

Note:

This setting applies to:

- `arc.log.txt`
- `arc_service.log.txt`
- `sp.log.txt`

It doesn't apply to `setup.log`, which covers Arc's initialization process, before the configuration has been loaded and applied.

For more details, see Log files (on page).

Execution time

This field lets you set the time that Arc will run to collect data. This is applicable for One-shot and Offline modes.

Note:
When this is set to 0, the execution time is interpreted as infinite.

Maximum disk space

This field lets you control the maximum amount of disk space in *megabyte (MB)* that will be used for Offline mode.

Offline archive naming strategy

This lets you select the information that will be included in the name of archive files created while running Offline mode.

Restore default

Once the settings have been saved, you can use this button to restore the default configuration.

Configure an Arc sensor

*It is possible to configure an individual Arc sensor directly from the **Sensors** details page for the related sensor.*

About this task

For more details on each configuration setting, see:

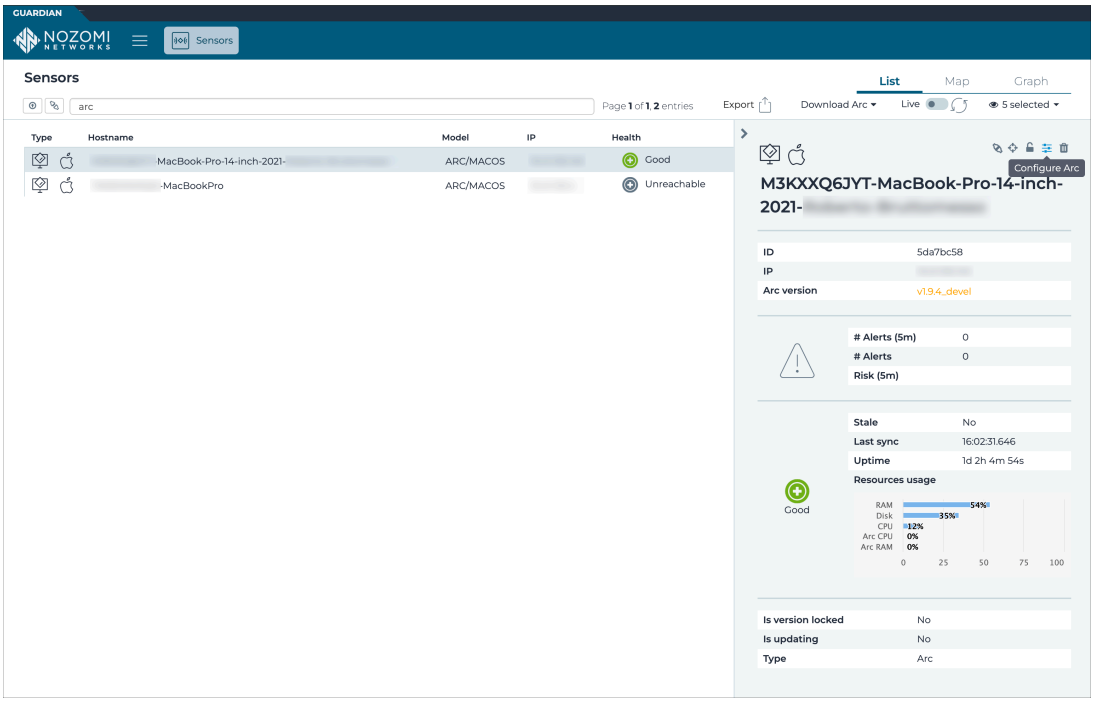
- Upstream (on page 175)
- Endpoint (on page 176)
- Network sensor (on page 178)
- Other (on page 31)

Procedure

1. In the top navigation bar, select **Sensors**.

   Result: The **Sensors** page opens.

2. From the list, select the applicable Arc sensor.



3. In the top right section, select the icon.

   Result: A dialog shows.

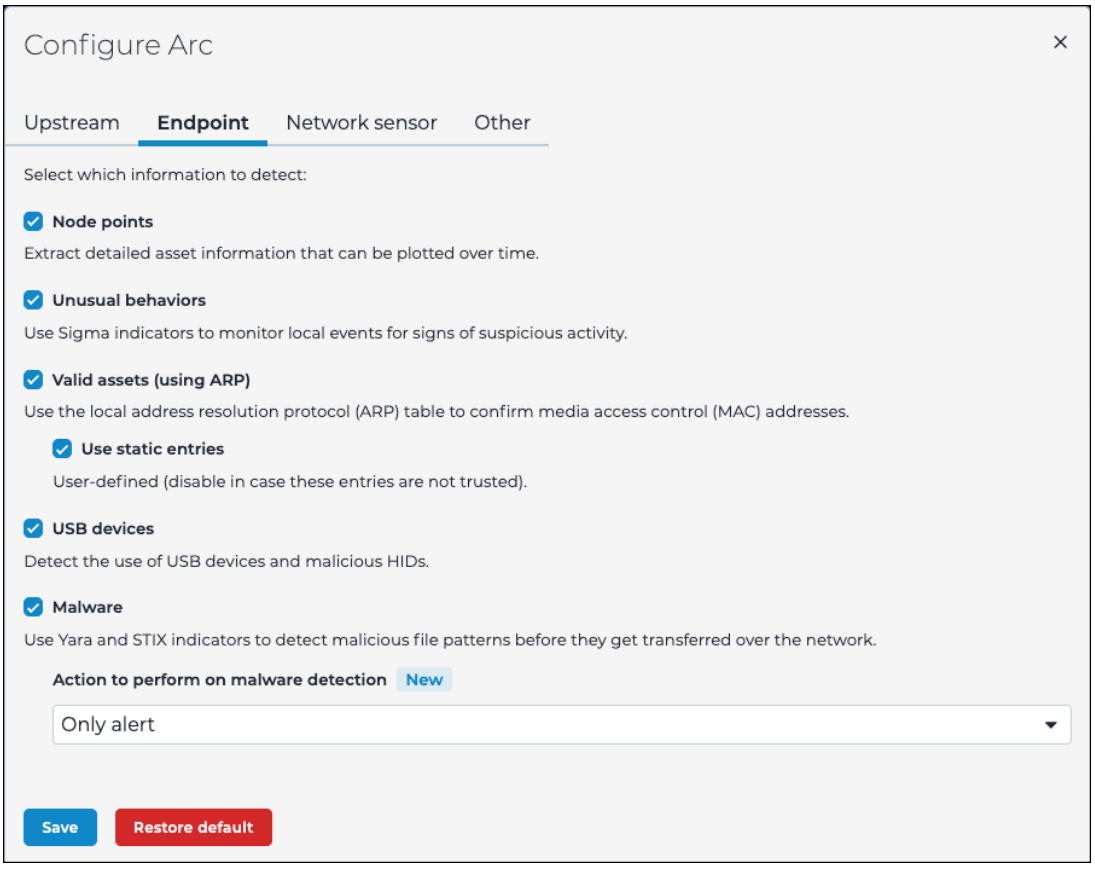
4. In the **Upstream** page, choose the settings that you want.



5. Select **Save**.
6. Select **Endpoint**.

   Result: The **Endpoint** page shows.

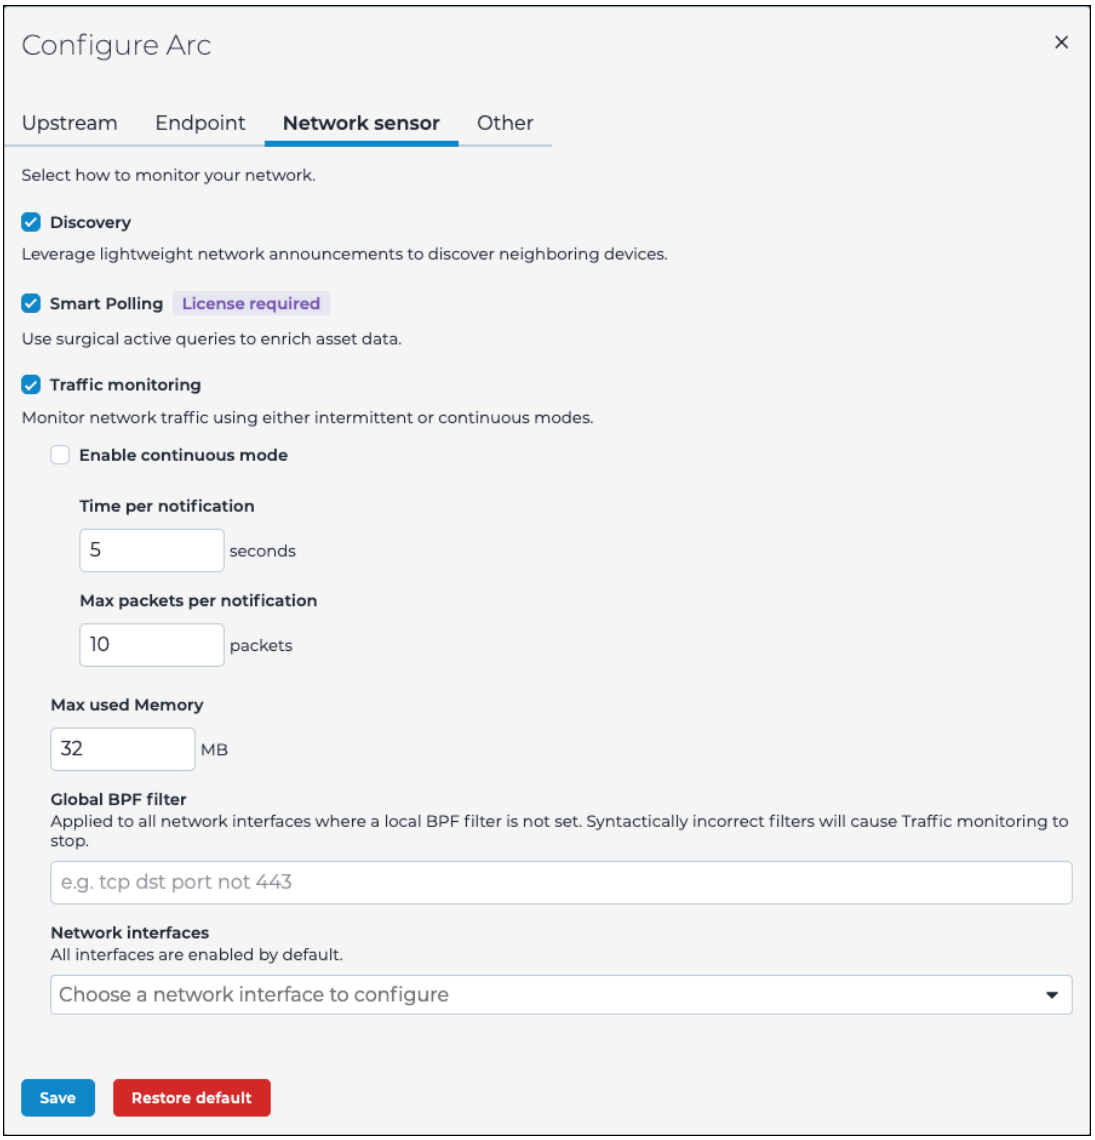
7. In the **Endpoint** page, choose the settings that you want.



8. Select **Save**.
9. Select **Network sensor**.

   Result: The **Network sensor** page shows.

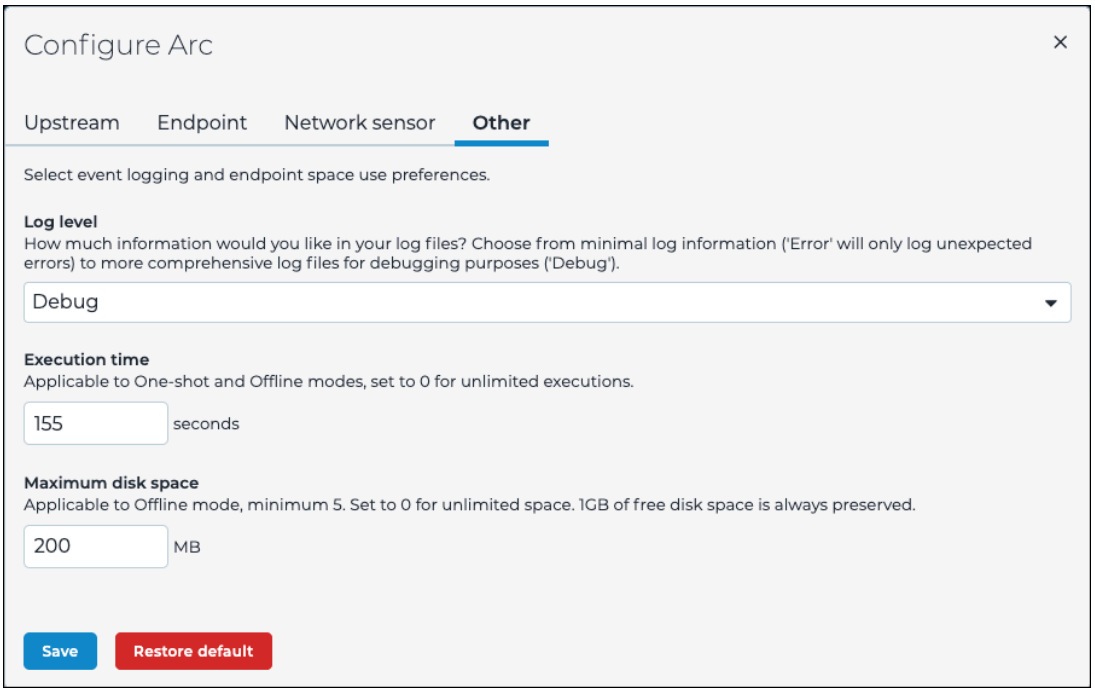
10. In the **Network sensor** page, choose the settings that you want.



11. Select **Save**.
12. Select **Other**.

    Result: The **Other** page shows.

13. In the **Other** page, choose the settings that you want.



14. Select **Save**.

Results

The Arc sensor has been configured.

Manage the Arc dependencies of a sensor

You can install, uninstall, update, enable, or disable the dependencies of an Arc sensor from the Sensors page.

Before you begin

The sensor must be running Arc v2.10.0, or later.

About this task

The Manage Arc dependencies dialog shows the dependencies that the sensor can use, with the installed version or state of each one. The dialog also shows which version is available to install or update. For more details on the dependencies, see Dependencies (on page 183).

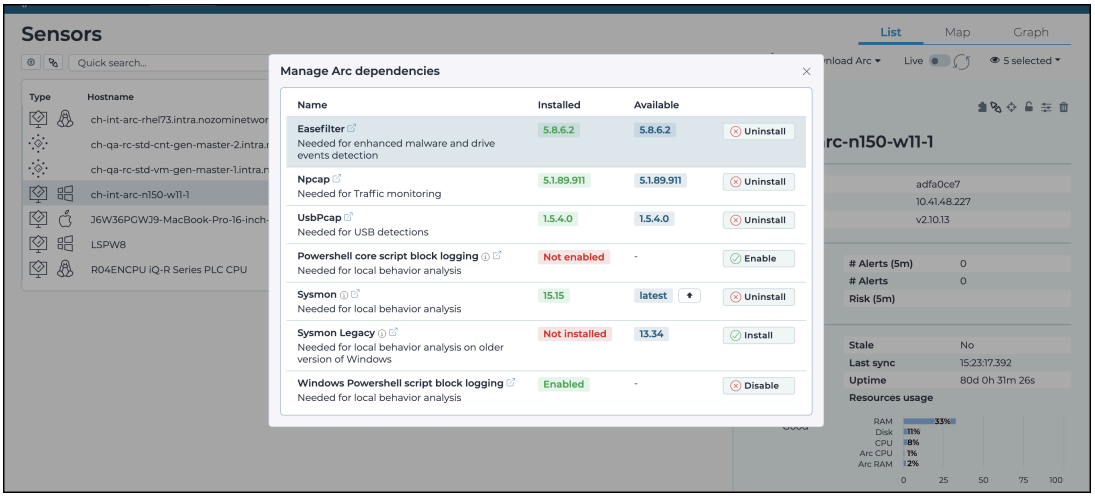
Procedure

1. In the top navigation bar, select **Sensors**.

   Result: The **Sensors** page opens.

2. From the list, select the applicable Arc sensor.
3. In the details pane, select the Manage Arc dependencies icon.

   Result: The Manage Arc dependencies dialog shows.



4. For the applicable dependency, select the action to apply.

   You can select **Install**, **Uninstall**, **Enable**, or **Disable**. To update a dependency, select the up arrow next to the available version.

   Result: A confirmation dialog shows.

5. **Optional:** Select the **Reboot the sensor if the operation requires it** checkbox.

   > **Note:**
   > If this operation requires a reboot, the sensor restarts automatically to complete it.

6. Select **Confirm**.

   Result: The **Request queued** result shows. The sensor runs the action the next time it connects.

Results

The action is queued on the sensor.

Related information

Dependencies (on page 183)

Apply an Arc dependency action to Arc sensors (on page 185)

Chapter 3. Alerts



*The **Alerts** page shows all the latest alerts in the system. It lets you view the alerts in different modes, and carry out actions on the alerts.*

Important:

To perform actions on alerts, the user must belong to a group with admin permissions. Non-admin users can access alerts only if at least one of the groups that they belong to has alerts permission enabled.



Figure 10. Alerts page menu

The top right section of the **Alerts** page has two icons that let you change between these two options:

- Standard mode (on page 42)
- Expert mode (on page 43)

Standard mode

You can view alerts in standard mode to give you an overview of the latest anomalies.

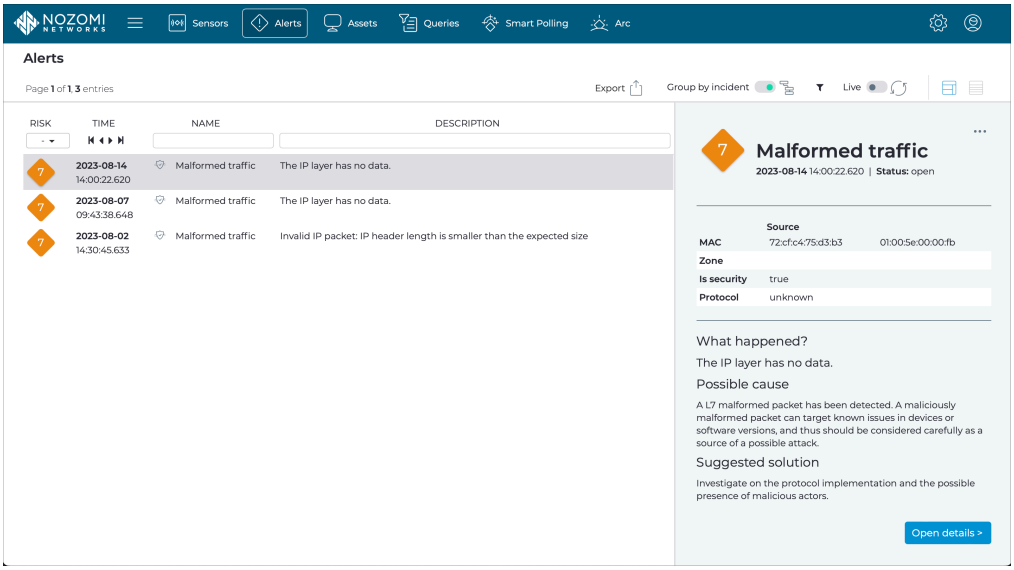


Figure 11. Standard mode

Risk
This shows the risk associated to each alert or incident.

Time
The time related to each event.

Name
The name category of the event.

Description
This shows a detailed description of the related event.

Analysis
If you can select a row, this pane will show a more detailed analysis of the alert.

Expert mode

You can view alerts in expert mode to give you a detailed view of the alerts in the system. This lets you filter, sort, and analyze the information in detail.

Expert mode shows a comprehensive table layout, with details on the alerts and incidents listed, which include:

- Addresses
- Labels
- The roles of the involved nodes, zones, protocol, and ports used in the involved transactions, and more

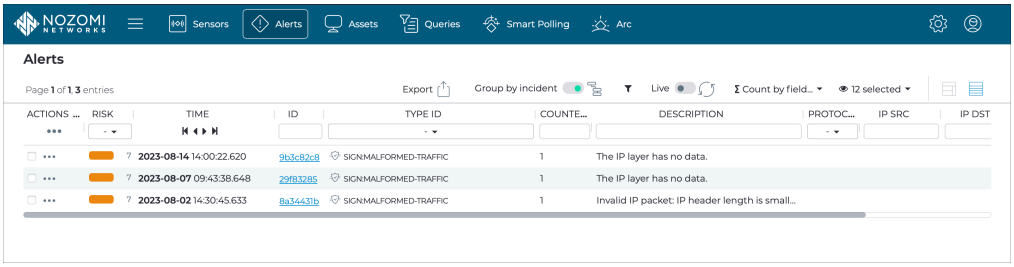


Figure 12. Expert mode

Export

The **Export** icon lets you export the current list in either *CSV* or Microsoft Excel format.

Group by incident

The **Group by incident** icon lets you group alerts by incident. This will show incidents and hides all the alerts that belong to them.

Filter

The filter icon opens a list of items that you let you filter the results.

Live

The **Live** toggle lets you change live view on, or off. When live mode is on, the page will refresh periodically.

Refresh

The **Refresh** icon lets you immediately refresh the current view.

Count by field

The **∑ Count by field** dropdown lets you select a data field on which to group and count the alerts.

Column selection

The columns selection icon lets you choose which columns to show or hide.

Alert deduplication

Alert deduplication groups repeated alerts into one alert entry when the alert type and key values match. It helps you reduce noise and review repeated activity in context.

The system compares alert-specific key values to decide whether two alerts match.

Deduplication can help you focus on unique issues while still preserving recurrence history.

How deduplication works

When deduplication is enabled and a new alert matches an existing alert, the system updates the existing alert.

In expert view, the grouped alert shows these column values:

- **Creation time**: The timestamp of the first time this alert was triggered.
- **Time**: The timestamp of the most recent trigger.
- **Counter**: The total number of times the alert has been triggered.
- **Occurrence timeline**: The timestamps of up to 100 of the most recent occurrences.

If a matching alert occurs after you close a deduplicated alert, the alert reopens automatically.

Deduplication key

Each alert type defines its own deduplication key fields. Only alerts with matching key values are grouped together.

The alert type ID is always part of the deduplication key. Common fields also include source node, destination node, and protocol-specific attributes.

View alerts in standard mode

You can view alerts in standard mode to give you an overview of the latest anomalies.

Procedure

1. In the top navigation bar, select **Alerts**.

   Result: The **Alerts** page opens.

2. In the top right corner, select the standard mode icon.

   Result: The Standard mode (on page 42) view opens.

View alerts in expert mode

You can view alerts in expert mode to give you a detailed view of the alerts in the system. This lets you filter, sort, and analyze the information in detail.

Procedure

1. In the top navigation bar, select **Alerts**.

   Result: The **Alerts** page opens.

2. In the top right corner, select the expert mode icon.

   Result: The Expert mode (on page 43) view opens.

Closing alerts

When you close an alert, or incident, a dialog lets you select a reason, and specify the learning process.

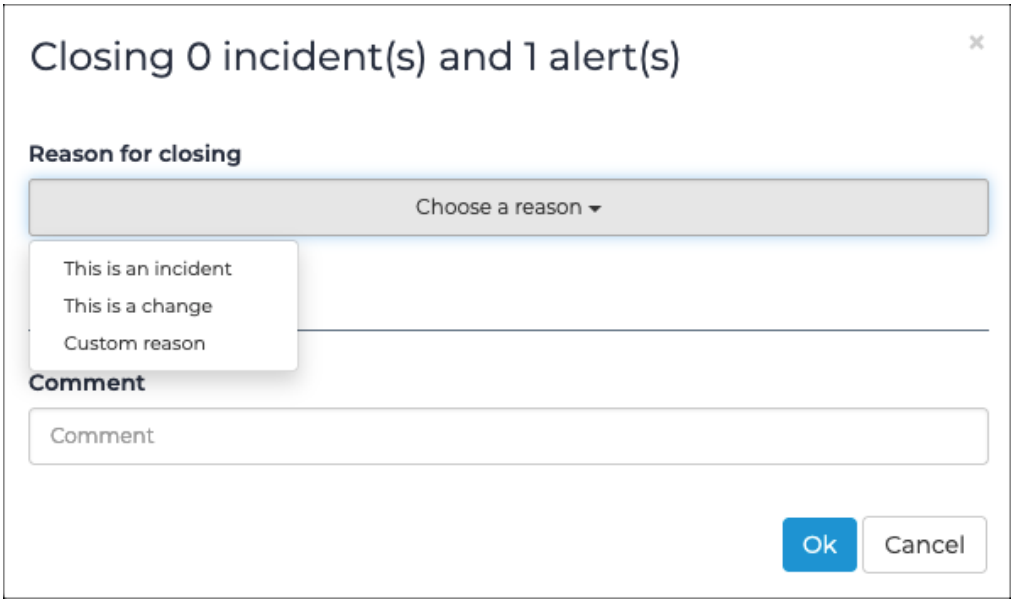


Figure 13. Alerts closing dialog

The **Reason for closing** dropdown has these options:

- **This is a change**: If the cause of the alert is an intended change to the network, such as:
  - A new computer being attached
  - New communication between two nodes that were not previously communicating

  Guardian can learn the change that has been detected as part of the environment baseline. When you close an alert in this way, the *intrusion detection system (IDS)* is instructed to learn the related objects. For example, when a `VI:NEW-NODE` alert is closed as a change, Guardian registers that the corresponding node is part of the environment and will not raise subsequent `VI:NEW-NODE` alerts about the same node.
- **This is a change**: If the cause of the alert is an intended change to the network, such as a new computer being attached, or a new communication between two nodes that were not talking before, the change detected by Guardian can be learned as part of the environment baseline. When closing an alert in this way, the IDS is instructed to learn the corresponding objects. For example, when a VI:NEW-NODE alert is closed as a change, Guardian registers that the corresponding node is part of the environment and will not raise subsequent VI:NEW-NODE alerts about the same node.

- **This is an incident**: If the cause of the alert is a configuration error, an attack, a malfunctioning device, or other security incident, the change is not learned as part of the environment baseline. When closing an alert in this way, the IDS is instructed to delete the corresponding objects. For example, a new node entering the network for the first time causes a VI:NEW-NODE alert. If an alert closes as an incident, reference to the new node is deleted. The VI:NEW-NODE alert is raised again in subsequent communication involving the same node.
- **Custom reason**: This lets you write a custom reason for closing an alert. You can enter a text string as the closing reason, with a request to apply one of the two described behaviors.

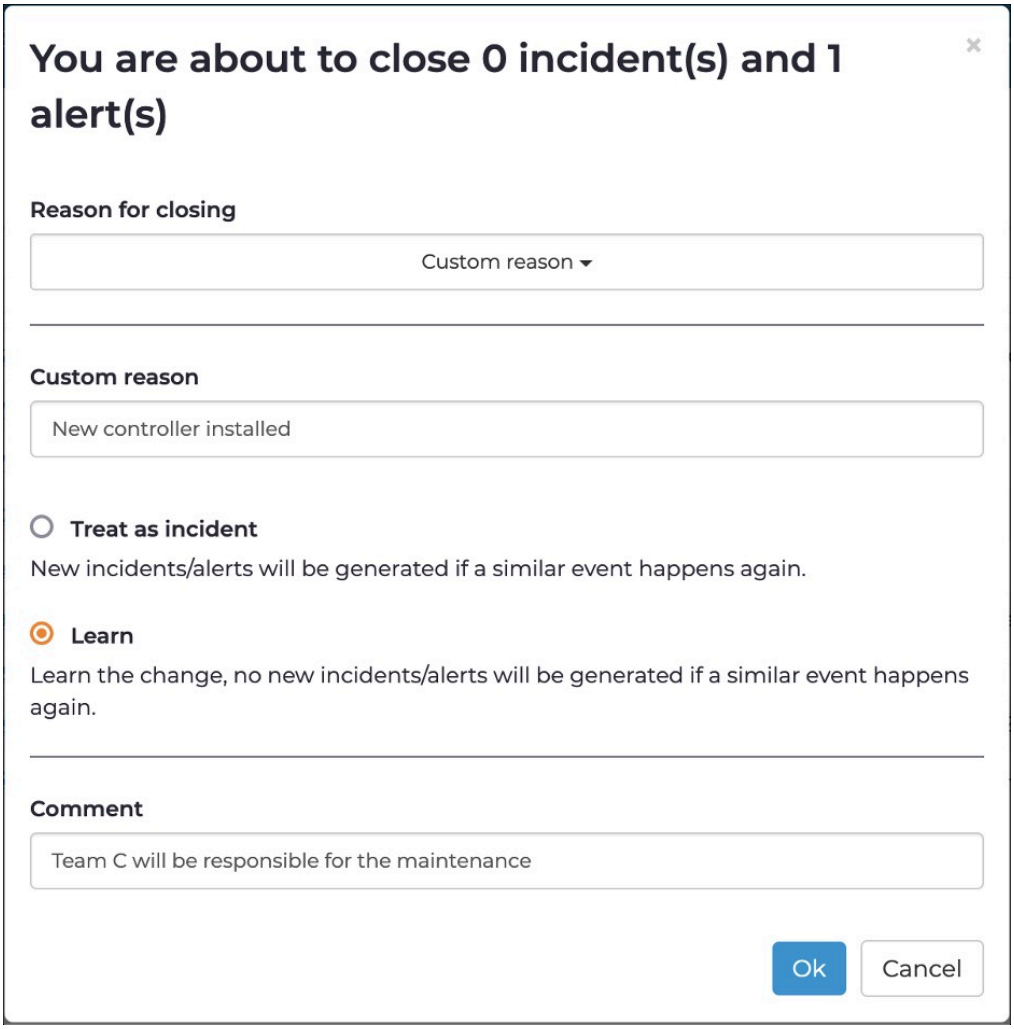


Figure 14. Closing alert for custom reason with comment

You can add a comment so that it shows in the **alert audit** log.

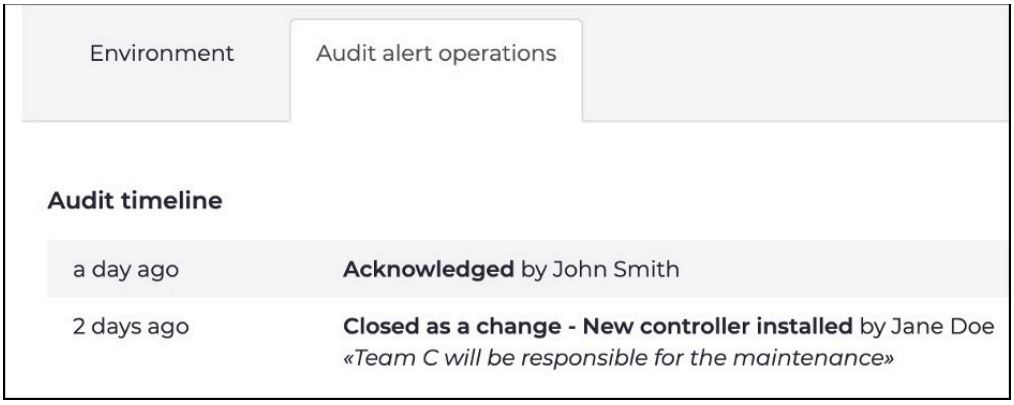


Figure 15. Audit alert operations

Actions menu

*The **Actions** menu gives you access to all the actions that you can do for the related alert.*

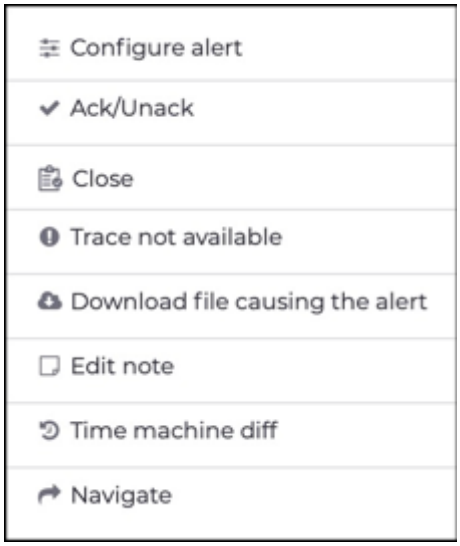


Figure 16. Actions menu

Note:

The options available in the **Actions** menu can change. The options will depend on:

- The type of alert
- The state of the sensor
- Whether the sensor is a Guardian or a *CMC*

Configure alert

You can use the **Configure alert** option to create a new alert rule for future events that are similar to the current one.

Ack/Unack

Once an alert or incident shows, you can mark it as acknowledged. You can also change the status back to unacknowledged again.

Close

Once an alert or incident has been addressed, you can mark it as closed, and choose the type of learning operation to perform.

Download trace

If a trace is available, you can choose to download it. The trace contains the packet that triggered the alert, along with an extract of the same session before and after that packet. Traces might be unavailable if the appliance is under stress. For detections that require multiple packets, such as **Multiple login failures**, the trace might not contain enough traffic to reproduce the alert. Incidents do not have an associated trace.

Download file causing the alert

Once a sensor has detected a malicious file, it is possible to download it for analysis. After you select this option, a dialog shows to warn the user that the file has been identified as malicious, or unwanted. To download the file, the user must acknowledge that they will do so at their own risk. In a *CMC*, this option is only available after the applicable file has been requested, (see below).

Edit note

Once an alert or incident shows, you can write a note for it, or edit an existing one.

Time machine diff

It is possible to open a time machine diff which corresponds to the time of the alert, or incident.

Navigate

Alerts and incidents have related nodes, links, vulnerabilities, or sessions. The **Actions** menu lets you navigate to these links.

Open the Actions menu in standard mode

*You can open the **Actions** menu to give you access to additional operations. When you are in standard mode, there are two different options available to open the menu.*

Procedure

1. View alerts in standard mode (on page 45).
2. Choose a method to open the **Actions** menu.

   Choose from:

   - Select the applicable alert, go to the top right corner and select the ••• icon
   - Select the **Open details** button. In the detailed view window, go to the top left corner and select the ••• icon

   Result: The **Actions** menu opens.

Open the Actions menu in expert mode

*You can open the **Actions** menu to give you access to additional operations.*

Procedure

1. View alerts in expert mode (on page 45).
2. Choose a method to open the **Actions** menu.

   Choose from:

   - To the left of the applicable alert, select the ••• icon
   - Select the link in the **ID** column. In the detailed view window, go to the top left corner and select the ••• icon

   Result: The **Actions** menu opens.

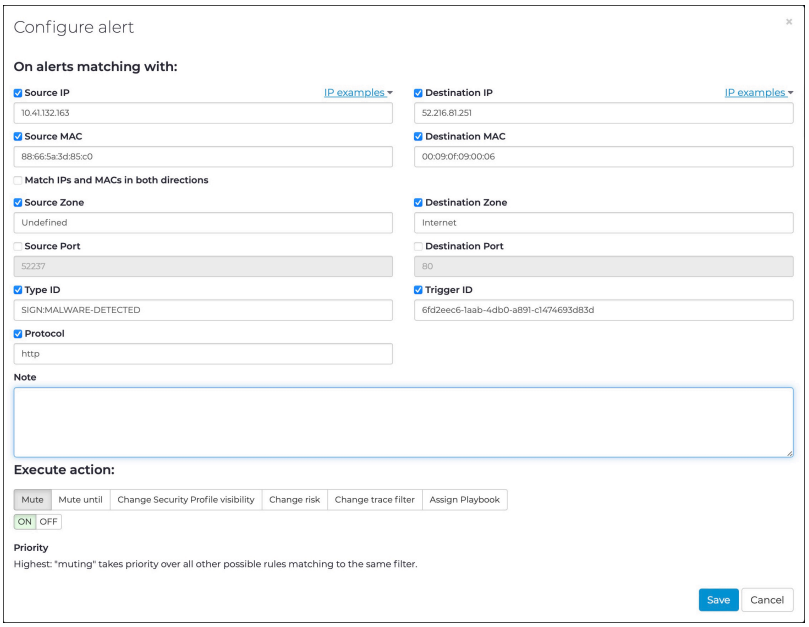
Configure an alert

*You can use the **Configure alert** option to create a new alert rule for future events that are similar to the current one.*

Procedure

1. Open the **Actions** menu for the selected alert(s) with one of the these options:
   - Open the Actions menu in standard mode (on page 51)
   - Open the Actions menu in expert mode (on page 51)
2. Select **Configure alert**.

   Result: A dialog shows.



3. In each of the fields, enter the necessary details.
4. Select **Save**.

Results

The alert has been configured.

Acknowledge an alert

Once an alert or incident shows, you can mark it as acknowledged. You can also change the status back to unacknowledged again.

Procedure

1. Open the **Actions** menu for the applicable file with one of the these options:
   - Open the Actions menu in standard mode (on page 51)
   - Open the Actions menu in expert mode (on page 51)
2. Select **Ack/Unack**.

   Result: After a few seconds, the tick symbol changes from grey to green.
3. If necessary, click the **Ack/Unack** option again to cancel the acknowledgment.

   Result: After a few seconds, the tick symbol changes from green to grey.

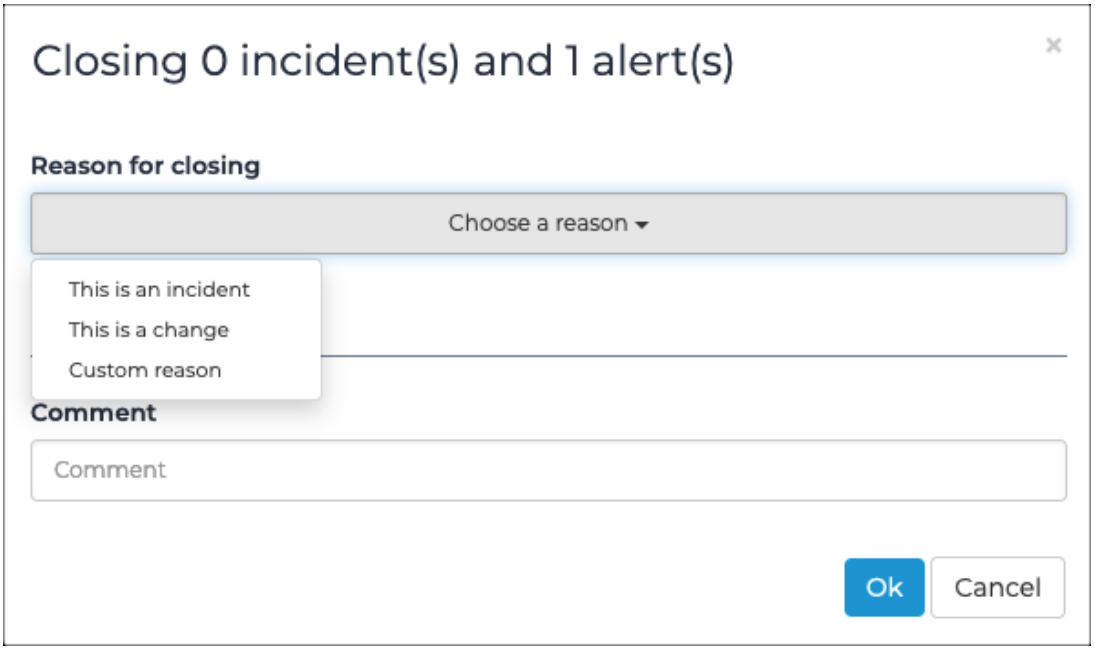
Close an alert

Once an alert or incident shows, you can mark it as closed, and choose the type of learning operation to perform.

Procedure

1. Open the **Actions** menu for the applicable file with one of the these options:
   - Open the Actions menu in standard mode (on page 51)
   - Open the Actions menu in expert mode (on page 51)
2. Select **Close**.

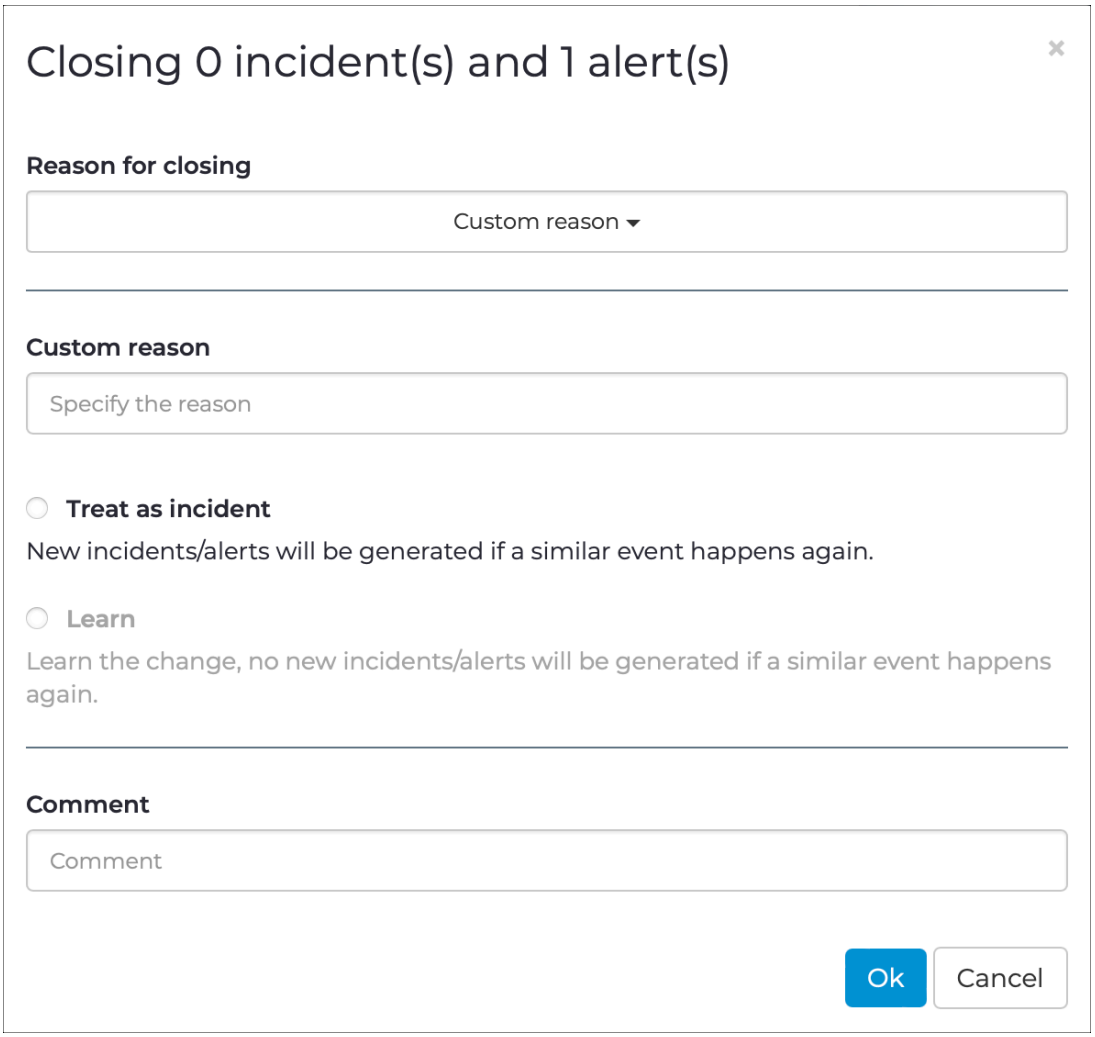
   Result: A dialog shows.



3. In the **Reason for closing** dropdown, choose one of these options:

   Choose from:
   - **This is a security incident**
   - **Custom reason**

4. If you chose **Custom reason**, the dialog will extend to show more options.



5. If you chose **Custom reason**, do the steps that follow.
   a. In the **Custom reason** field, enter the details as necessary.
   b. Choose from one of these options:
      - **Treat as incident**
      - **Learn**
6. If necessary, enter a comment in the **Comment** field.
7. Select **Ok**.

Download a trace

If a trace is available, you can choose to download it. The trace contains the packet that triggered the alert, along with an extract of the same session before and after that packet. Traces might be unavailable if the appliance is under stress.

Procedure

1. Open the **Actions** menu for the applicable file with one of the these options:
   - Open the Actions menu in standard mode (on page 51)
   - Open the Actions menu in expert mode (on page 51)
2. Select **Download trace**.

   Result: The file downloads to your downloads folder.

Download a malicious file

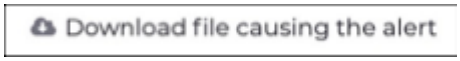
Once a sensor has detected a malicious file, it is possible to download it for analysis.

About this task

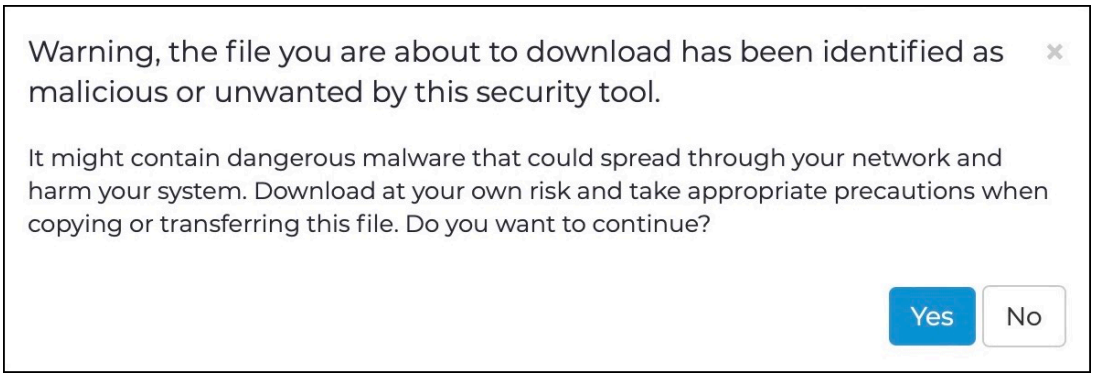
Only qualified personnel should download malicious files. Before you do this procedure, make sure that you have the correct permissions.

Procedure

1. Open the **Actions** menu for the applicable file with one of the these options:
   - Open the Actions menu in standard mode (on page 51)
   - Open the Actions menu in expert mode (on page 51)
2. Select **Download file causing the alert**.



Result: A dialog shows.



3. If you want to proceed, select **Yes**.

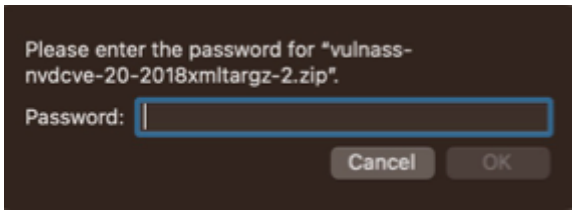
> **Important:**
> Nozomi Networks recommends that only qualified personnel download malicious, or unwanted, files. Download these files at your own risk.

Result: The file downloads to your downloads folder and a password dialog that contains a randomized, single-use password shows.



4. Copy the password, and select **Ok**.
5. Go to the folder where the file was downloaded to.
6. Double-click the *ZIP* file to open it.

Result: A dialog that prompts you to enter a password shows.



7. Paste the password into the password field and select **OK**.

Result: You can now access the file.

Edit a note for an alert

Once an alert or incident shows, you can write a note for it, or edit an existing one.

Procedure

1. Open the **Actions** menu for the applicable file with one of the these options:
   - Open the Actions menu in standard mode (on page 51)
   - Open the Actions menu in expert mode (on page 51)
2. Select **Edit note**.

Result: A note is opened with the identification of the related alert, or incident.

3. In the text field, write a note.

> **Note:**
> There is a character limit of 1000 characters.

4. Select **Save**.

Results

The note has been edited.

View a diff from an alert

This automatic feature will use the previous and subsequent snapshots according to the time of the alert.

Procedure

1. In the top navigation bar, select **Alerts**.

   Result: The **Alerts** page opens.

2. Choose a method to open the actions menu.

   Choose from:

   - In the table, select the hyperlink to open the details page. Select **Actions**
   - In the table, select the ••• icon

3. Select **Time machine diff**.

   Result: The time diff shows.

4. To see more details on the right side of the graph, select the applicable node or link.

Navigate from an alert

Alerts and incidents have related nodes, links, vulnerabilities, or sessions. The actions menu lets you navigate to these links.

Procedure

1. Open the **Actions** menu for the applicable file with one of the these options:

   - Open the Actions menu in standard mode (on page 51)
   - Open the Actions menu in expert mode (on page 51)

2. Click the **Navigate** option.

   Result: A dialog shows.

3. Select the desired link.

Edit a playbook associated with an alert

A playbook associated with an alert can be modified. A change you make on the playbook only affects the playbook related to that specific alert.

About this task

You can edit a playbook associated with an alert from the **Playbook** tab in the **Alerts** page.

Procedure

1. In the top navigation bar, select **Alerts**.

   Result: The **Alerts** page opens.

2. Find the alert on which to edit the assigned playbook.
3. Select the alert, then select **Playbook** at the bottom of the screen.
4. To edit the playbook, select **Edit**.

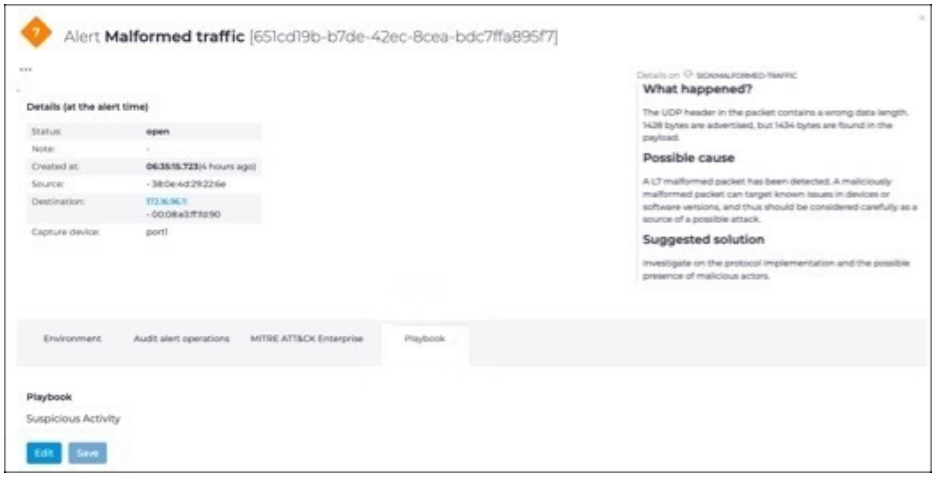


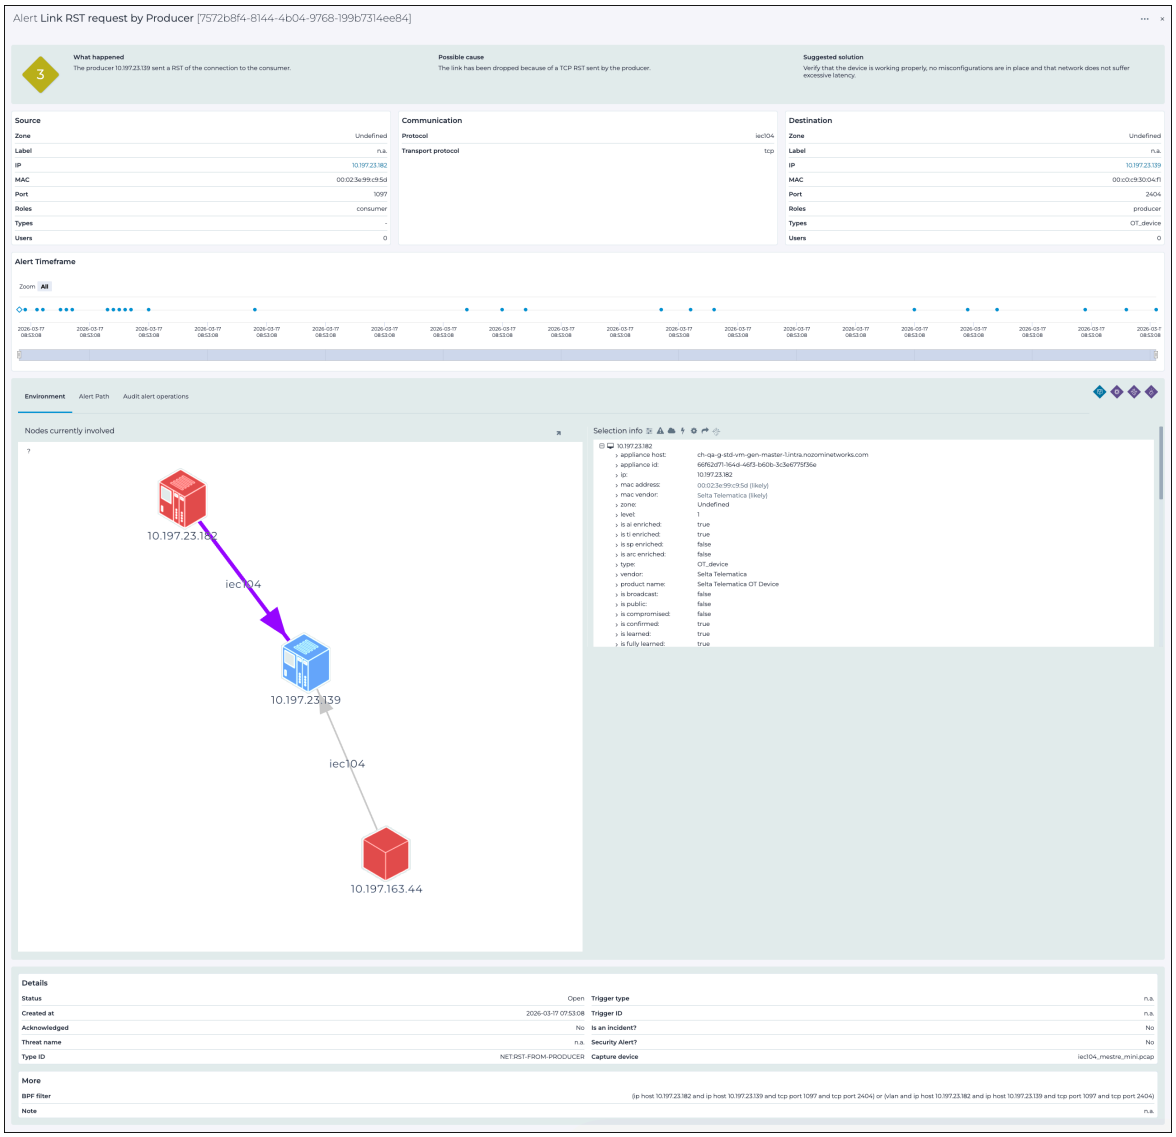
Figure 17. Playbook tab

5. Edit the playbook as necessary.
6. Select **Save**.

> **Note:**
> The playbook template from which the alert playbook was generated remains unchanged, as do all other alert playbooks generated from the same playbook template.

Details page

*The **Alerts** details page provides a comprehensive view of a triggered alert, including cause, context, and recommended response. It displays communication paths, affected nodes, and relevant MITRE ATT&CK tactics. Use this page to investigate incidents and track alert handling activity.*



What happened

Shows a summary and a list of key details for the related alert.

Possible cause

Shows a possible cause of the alert.

Suggested solution

Our recommendation on the best way to solve the problem.

Alert Timeframe

When alert deduplication is enabled, the system groups repeated alerts of the same type that share identical deduplication key values, instead of creating a new alert entry each time. The **Alert Timeframe** section shows the occurrence history for that grouped alert.

The section includes the following fields:

- **Creation time**: The timestamp of the first time this alert was triggered.
- **Time**: The timestamp of the most recent trigger.
- **Counter**: The total number of times the alert has been triggered.
- **Occurrence timeline**: The timestamps of up to 100 of the most recent occurrences.

Use the occurrence timeline to understand recurrence patterns over time. Timestamps that are close together can indicate a burst of repeated activity, while larger gaps can indicate periodic or intermittent behavior. The timeline shows only the most recent 100 occurrences, so older occurrences will not appear.

This information is also accessible through queries. If alert deduplication is disabled, this section is not shown.

Environment

Shows a graph with the logical connections for the related nodes. The **Selection info** section shows all the applicable information for the alert. It also gives you access to the actions menu.

Alert Path

The **Alert Path** view displays the physical connection path between source and destination devices involved in a triggered alert. This path is traced at the cable level, including intermediate switches and any other devices connected along the route. The alert path shows with a red line.

This view helps users assess the potential impact of response actions, such as disabling a specific switch port to isolate a device. By identifying additional assets that share the same physical infrastructure, operators can evaluate containment strategies and take targeted action directly from the alert interface.

Audit alert operations

The **Audit alert operations** section shows a timeline of user or system actions taken on a specific alert. This section helps you track how an alert is handled throughout its lifecycle.

It records these types of activity:

- Acknowledgement of the alert
- Escalation of the alert to an incident
- Changes to the alert's severity or classification
- Notes or comments that were added
- Automatic responses or triggered workflows, if the system is integrated with them

MITRE ATT&CK Enterprise

The **MITRE ATT&CK for ICS** section helps security analysts understand how the current alert maps to known tactics and techniques used by threat actors in *ICS* environments.

This section shows which tactics from the MITRE ATT&CK for ICS framework are relevant to the alert. The tactics represent different stages of an attack lifecycle, such as:

- Initial Access
- Execution
- Persistence
- Evasion
- Discovery
- Lateral Movement
- Collection
- Command and Control
- Inhibit Response Function
- Impair Process Control
- Impact

Each tactic can display one or more techniques observed in the alert. This mapping helps analysts determine the attacker's goals and possible next steps.



Chapter 4. Assets



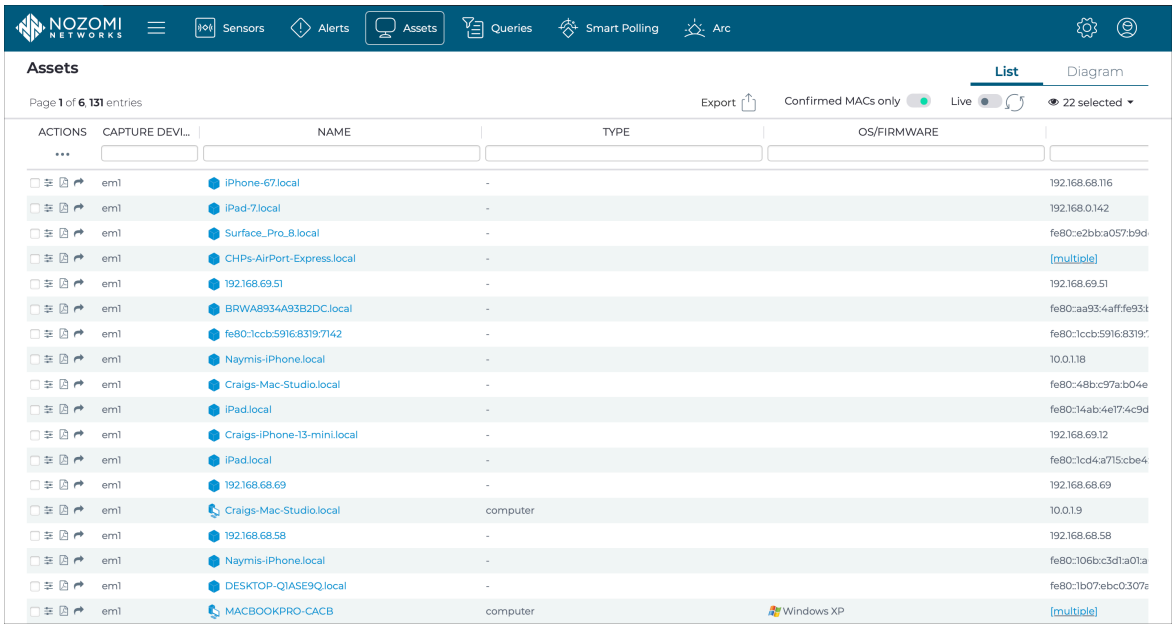
*The **Assets** page shows all the physical components and systems in the local network environment and their associated details. It also lets you perform actions on those assets. Depending on the nodes and components involved, assets can range from a simple personal computer to an operational technology (OT) device.*

The top right section of the **Assets** page has these two tabs:

- List (on page 67)
- Diagram (on page 69)

List

*The **List** page shows all the assets in table format.*



Export

The **Export** icon lets you export the current list in either *CSV* or Microsoft Excel format.

Confirmed MACs only

The **Confirmed MACs only** toggle lets you select only assets that have a confirmed *MAC* address.

Live

The **Live** toggle lets you change live view on, or off. When live mode is on, the page will refresh periodically.

Refresh

The **Refresh** icon lets you immediately refresh the current view.

Column selection

The columns selection  icon lets you choose which columns to show or hide.

Diagram

*The **Diagram** page uses the Purdue model format to display the assets. Assets are shown in separate rows, according to their level.*

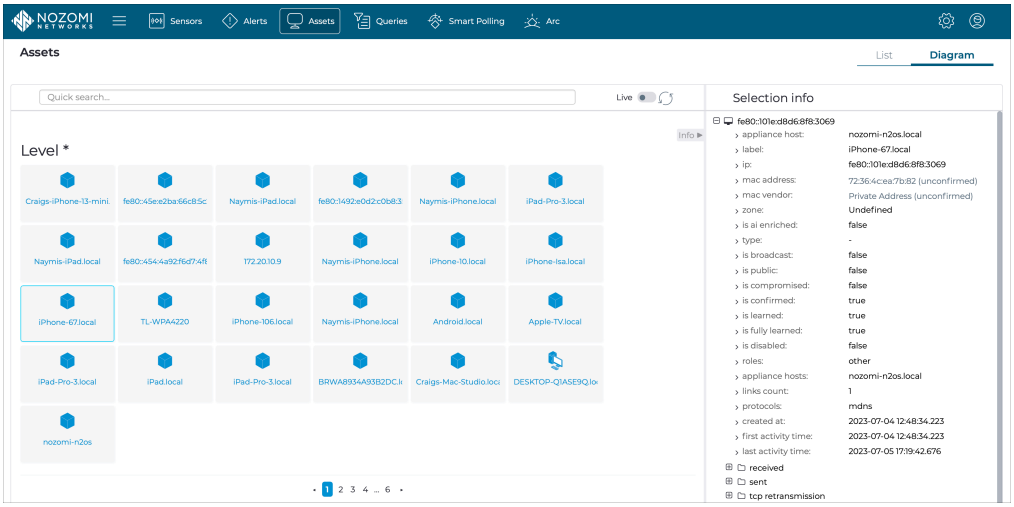


Figure 18. Diagram page

Search bar

The search bar lets you search for a specific item.

Live

The live toggle lets you immediately refresh the graph.

Refresh

The refresh icon lets you immediately refresh the graph.

Levels section

The levels section shows an icon for each asset, and shows on which level of the Purdue model it is.

Selection info

When you select the link below an asset's icon, the **Selection info** pane shows more details for the selected asset.

Configure an asset

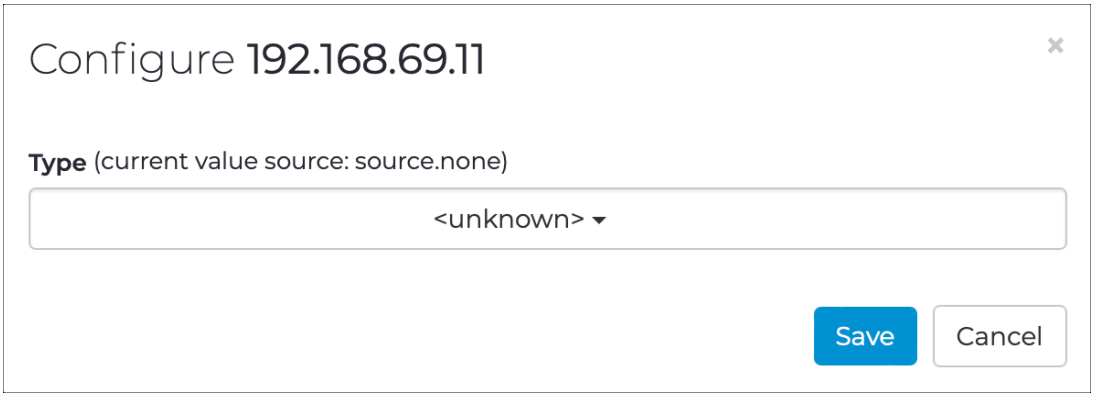
*The **Lists** page lets you configure assets.*

Procedure

1. In the top navigation bar, select **Assets**.

   Result: The **Assets** page opens.

2. In the Actions column, to the left of the applicable asset, select the icon.



3. From the dropdown, select the applicable type.
4. Select **Save**.

Results

The asset has been configured.

Generate a PDF report of an asset

*The **Lists** page lets you generate a report in portable document format (PDF).*

Procedure

1. In the top navigation bar, select **Assets**.

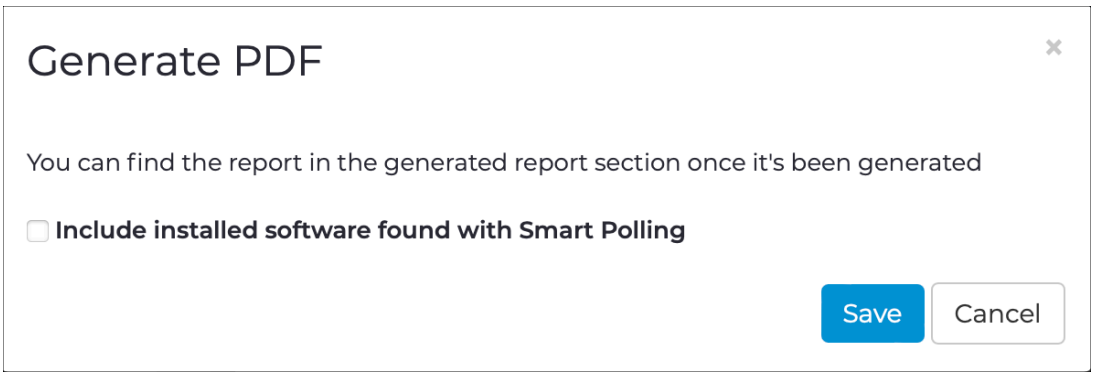
   Result: The **Assets** page opens.

2. In the Actions column, to the left of the applicable asset, select the icon.

   Result: A dialog shows.

3. **Optional:**

   If necessary, select the **Include installed software found with Smart Polling** checkbox.



4. Select **Save**.

   Result: The *portable document format (PDF)* file generates in the background. When it is ready, you can view it on the **Reports** page.

5. To view the report, go to **Reports > Generated**.

Navigate from an asset

*The **Lists** page lets you use hyperlinks to navigate to entities that are related to an asset.*

Procedure

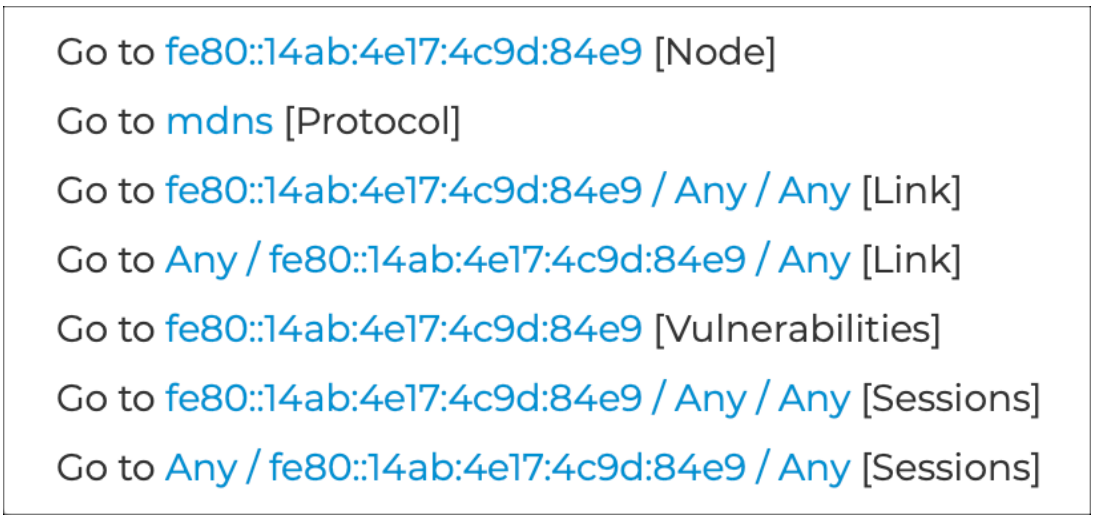
1. In the top navigation bar, select **Assets**.

   Result: The **Assets** page opens.

2. In the Actions column, to the left of the applicable asset, select the icon.

   Result: A list of related entities shows.

3. Select the hyperlink that you want to navigate to.



Results

The entity shows in the applicable page.

View more details for an asset

*Both the **Lists** and **Diagrams** pages lets you view more details for a specific asset.*

Procedure

1. In the top navigation bar, select **Assets**.

   Result: The **Assets** page opens.

2. In the top right section, select either **List** or **Diagram**
3. Choose a method to view more details for an asset:

   Choose from:

   - If you chose the **List** page, in the **NAME** column for the applicable asset, select the hyperlink
   - If you chose the **Diagram** page, below the applicable asset, select the hyperlink

Results

The Details page (on page 74) for the asset shows.

Details page

The details page lets you view more detailed information for an assets.

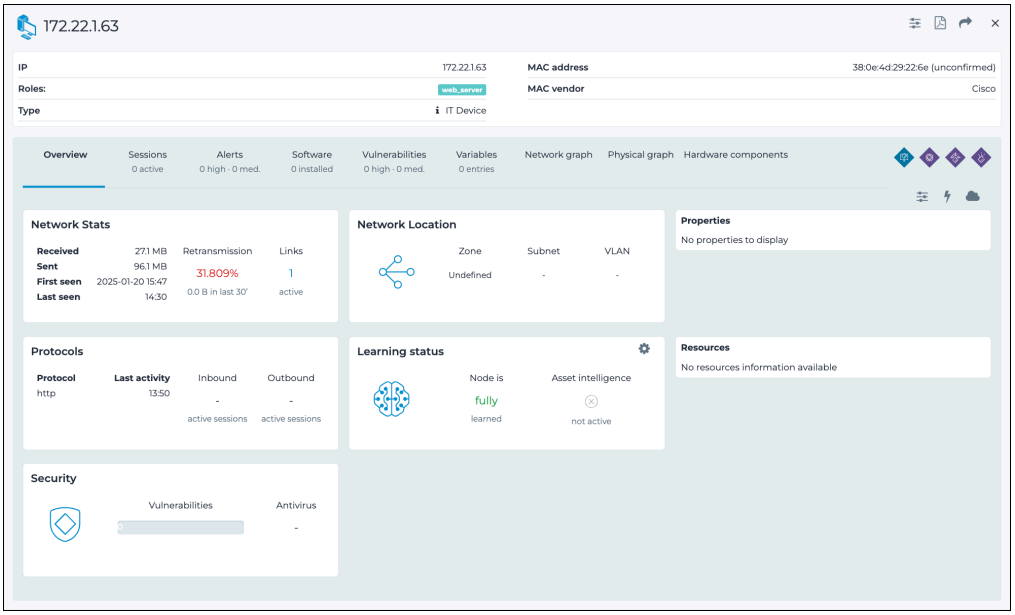


Figure 19. Asset details page

This details page shows more details for an asset.

The top section of the screen contains generic data. You can hover your mouse over the information icon to display the source, granularity and confidence of the corresponding piece of data. Data includes:

- *IP* (address)
- Roles
- Vendor
- *MAC* address
- *MAC* vendor

The details window has these tabs:

- Overview (on page 76)
- Sessions (on page 77)
- Alerts (on page 78)
- Software (on page 79)
- Vulnerabilities (on page 80)
- Variables (on page 81)
- Network graph (on page 82)
- Physical graph (on page 83)
- Hardware components (on page 84)

Information icon

When you hover over the icon, you can see information for:

- Source
- Granularity
- Confidence

Source

Information source	Description
manual	Information that is manually added from the configuration
imported data	Imported information
passive detection	Information from deep packet inspection
asset-kb	Information from Asset Intelligence
smart-polling	Information from Smart Polling

Granularity

Level of detailed information	Description
manual-or-import	Information manually added or imported
complete	Detailed information that has been extracted
partial	Detailed, but not complete information
generic	A family/generic value is found, but it is not detailed
unknown	Unknown

Confidence

Level of confidence in information	Description
manual-or-import	Information manually added or imported, with the highest level of confidence at this level
high	High level of confidence
good	Good level of confidence
low	Low level of confidence
unknown	Unknown confidence

Overview

*The **Overview** page shows a general overview of information for items such as network statistics and location, protocols, and learning status for the related assets.*

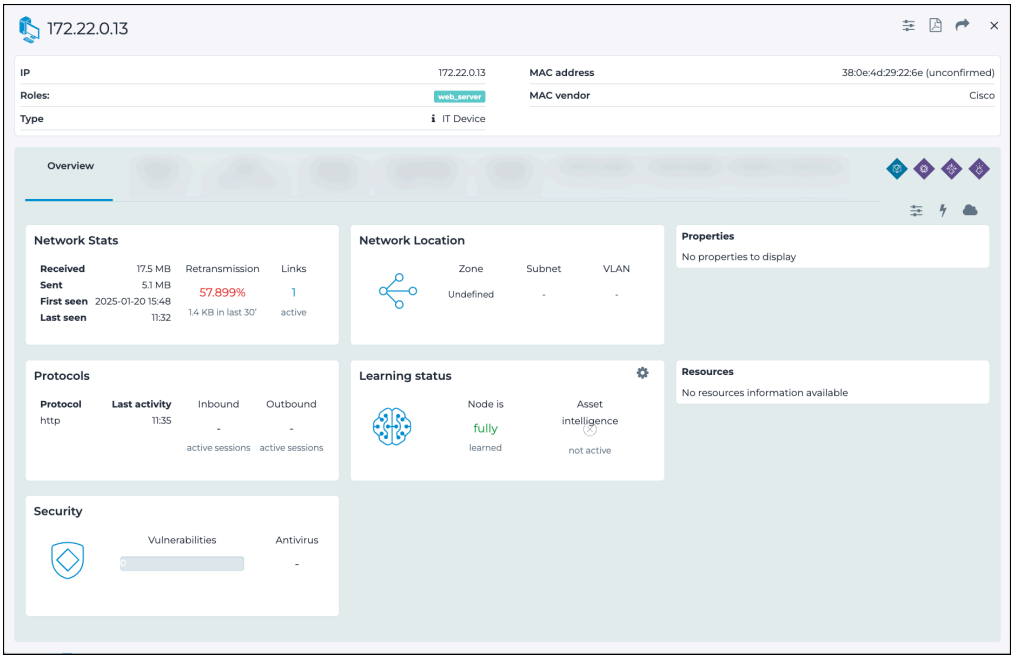


Figure 20. Overview tab

Network Stats

This shows useful statistics for the network activity for the related device.

Network Location

This shows information for the location for the related device on the network.

Properties

This shows additional information for the related device.

Protocols

This shows information for the different protocols that the related device uses.

Learning status

This shows the learning status of the related asset.

Security

This shows the security status of the related asset.

Sessions

*The **Sessions** page shows detailed information for communication sessions between devices.*

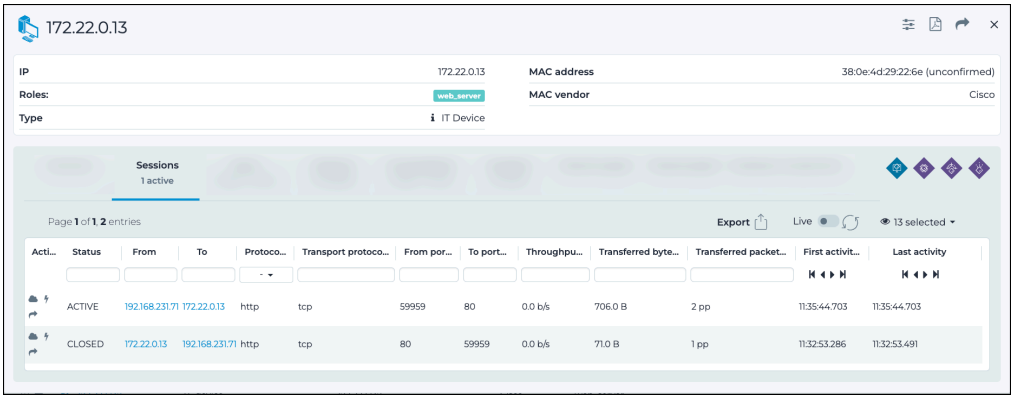


Figure 21. Sessions tab

Export

The **Export** icon lets you export the current list in either *CSV* or Microsoft Excel format.

Live / refresh

The **Live** icon lets you change live view on, or off. When live mode is on, the page will refresh approximately every five seconds.

Column selection

The columns selection icon lets you choose which columns to show or hide.

Alerts

*The **Alerts** page shows detailed information for all the alerts that have been raised for the related assets.*

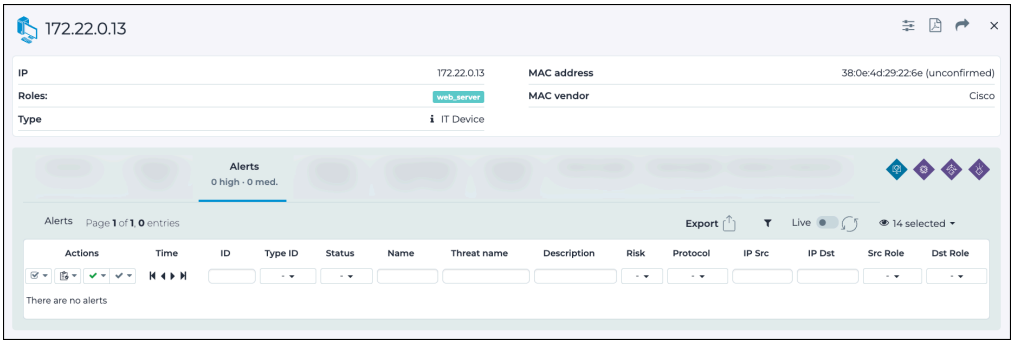


Figure 22. Alerts tab

Export

The **Export** icon lets you export the current list in either *CSV* or Microsoft Excel format.

Live / refresh

The **Live** icon lets you change live view on, or off. When live mode is on, the page will refresh approximately every five seconds.

Column selection

The columns selection icon lets you choose which columns to show or hide.

Software

*The **Software** page shows a list of software applications that are installed on the related assets.*

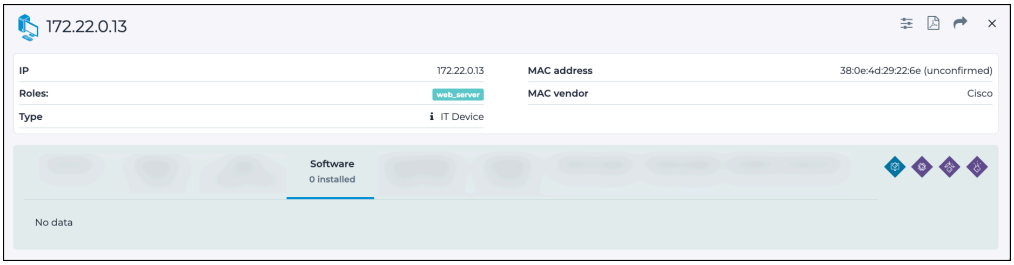


Figure 23. Software tab

Vulnerabilities

*The **Vulnerabilities** page shows a list of vulnerabilities that are present on the related asset.*

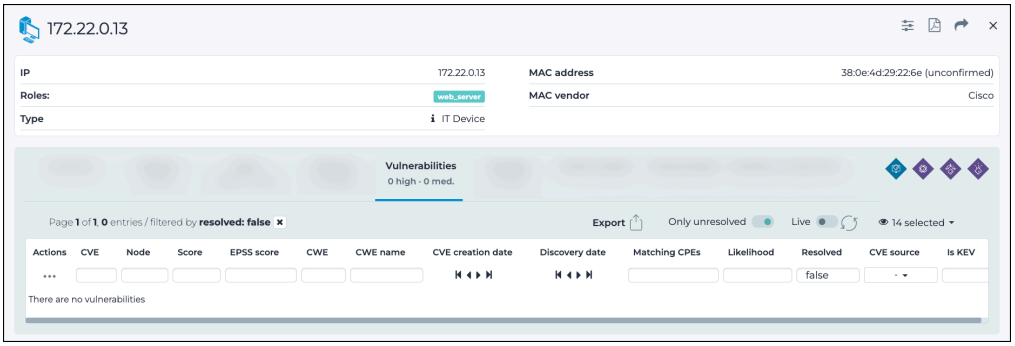


Figure 24. Vulnerabilities tab

Export

The **Export** icon lets you export the current list in either *CSV* or Microsoft Excel format.

Only unresolved

The **Only unresolved** toggle lets you filter the column to only show unresolved vulnerabilities.

Live / refresh

The **Live** icon lets you change live view on, or off. When live mode is on, the page will refresh approximately every five seconds.

Column selection

The columns selection icon lets you choose which columns to show or hide.

Variables

*The **Variables** page shows detailed information for the variables hosted by the related asset.*

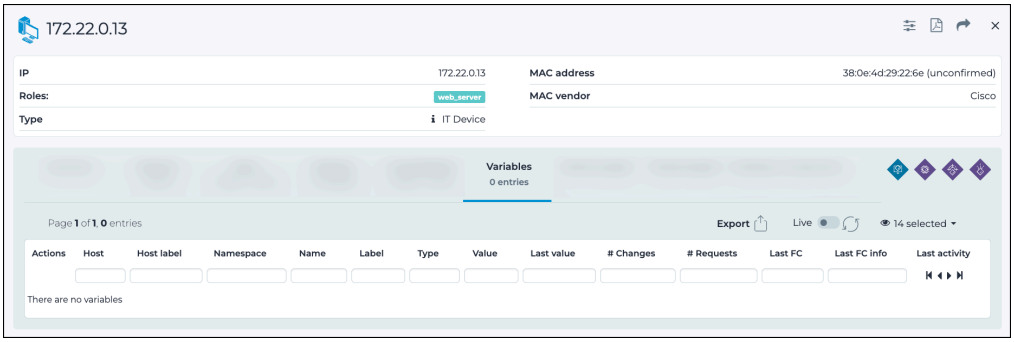


Figure 25. Variables tab

Export

The **Export** icon lets you export the current list in either *CSV* or Microsoft Excel format.

Live / refresh

The **Live** icon lets you change live view on, or off. When live mode is on, the page will refresh approximately every five seconds.

Column selection

The columns selection icon lets you choose which columns to show or hide.

Network graph

*The **Network graph** page shows how an asset communicates across the network, using real traffic data. It is ideal for uncovering lateral movement or validating expected behavior.*

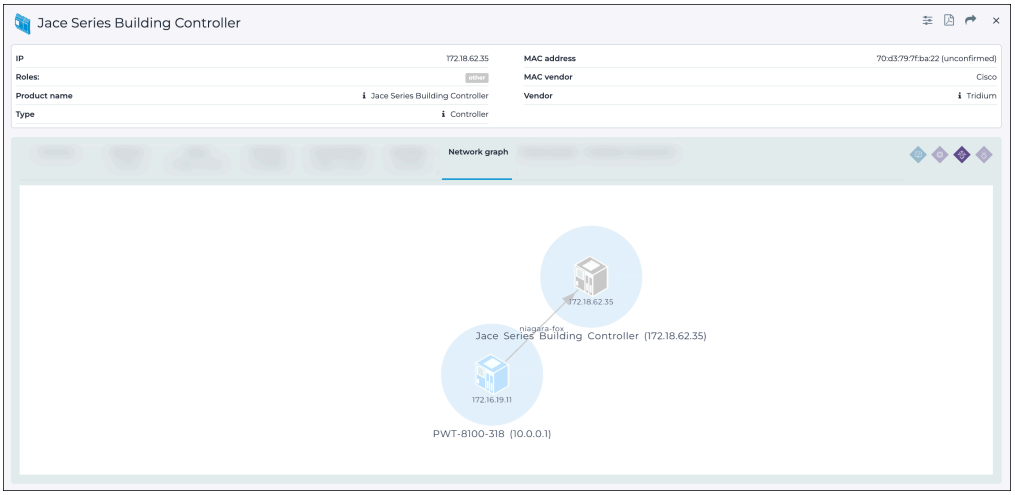


Figure 26. Network graph tab

The **Network graph** tab displays the asset's logical connections, based on observed network traffic. It shows which devices the asset communicates with, using lines between nodes to represent live or historical communication paths.

This view is useful for analyzing device behavior, identifying communication patterns, and detecting potential lateral movement. While the **Physical graph** focuses on physical connections, the **Network graph** provides visibility into traffic-based interactions, offering a complementary perspective on the asset's network role.

Physical graph

*The **Physical graph** page shows the device's physical link to its switch, including port-level data. It is ideal for tracing, troubleshooting, or safely disconnecting assets.*

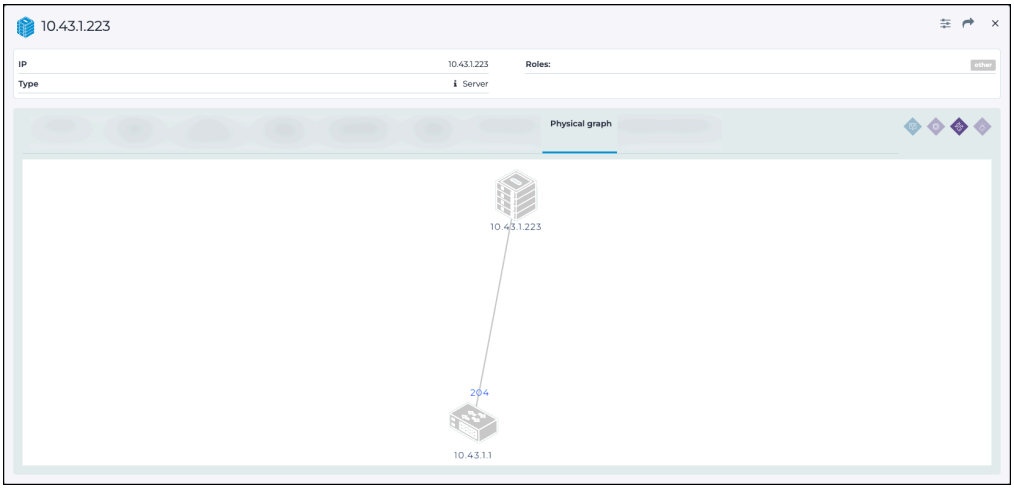


Figure 27. Physical graph tab

The **Physical graph** has a cable-level view that shows how a specific device is physically connected to the network. It shows the direct connection between the asset and its switch, including the exact port in use. This information provides clear insight into the physical location and wiring of the device.

This view supports operational tasks such as device tracing, maintenance, and isolation. Users can confirm port assignments or prepare to disconnect a device without referencing external network diagrams. The graph delivers actionable, asset-level visibility into physical connectivity.

Hardware components

*The **Hardware Components** page shows a list of hardware modules installed on the same chassis as the selected controller, providing details on each component for better visibility and management of the device's physical configuration.*

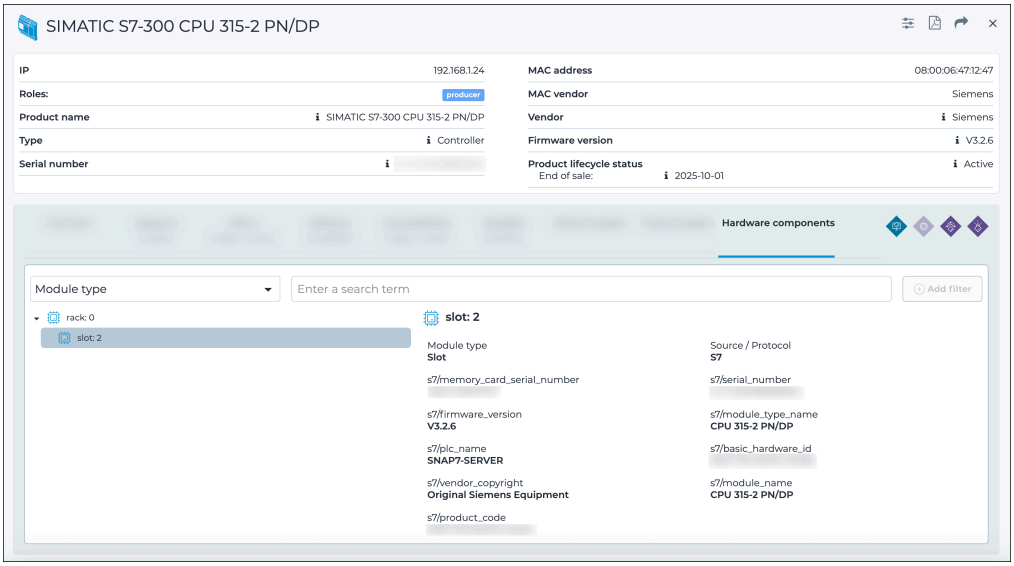


Figure 28. Hardware components tab

Choose

This dropdown lets you select from a list of attributes.

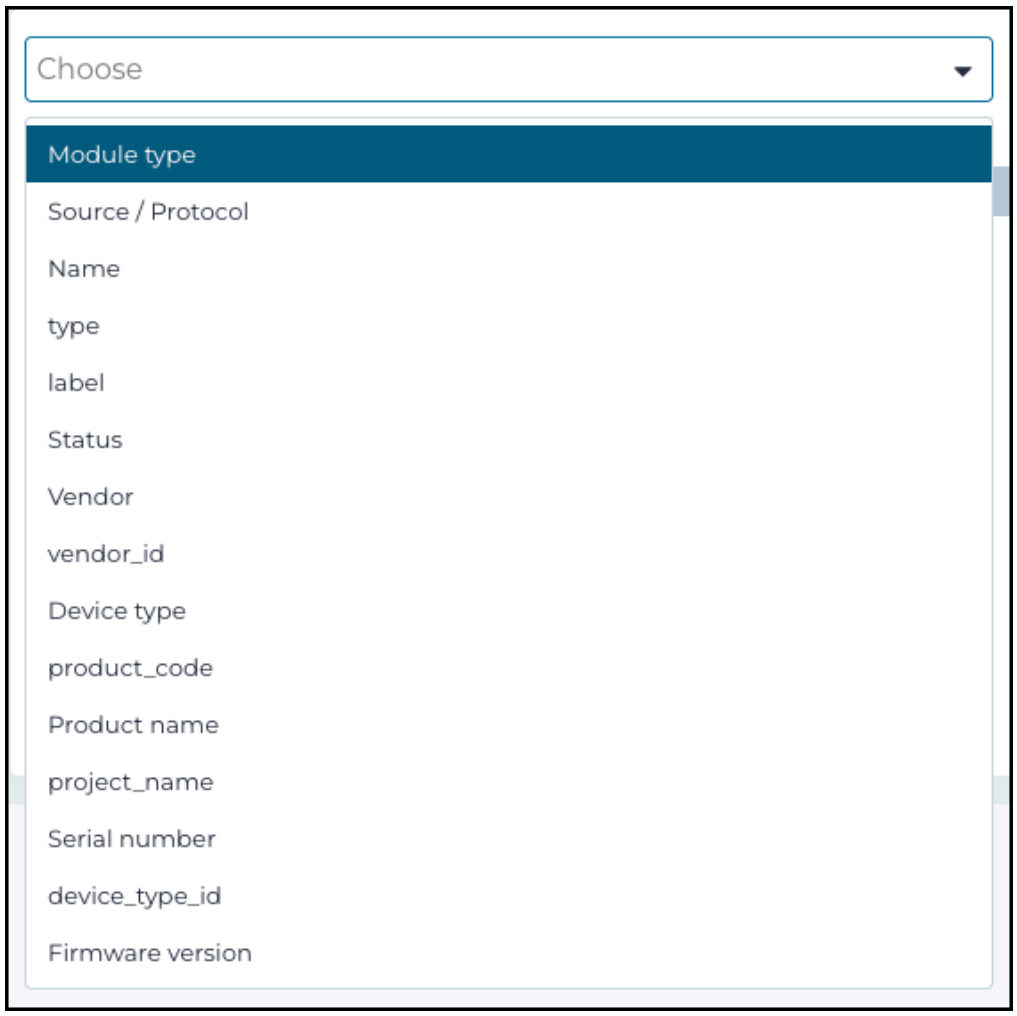


Figure 29. Choose dropdown

Search bar

The search bar lets you enter a term to filter the list.



Figure 30. Search bar

Add filter

Select this button to apply the filter in the search bar.



Figure 31. Add filter button

Chapter 5. Queries



You can use the Nozomi Networks Query Language (N2QL) syntax to create complex data processes to obtain, filter, and analyze lists of information from the Nozomi Networks software.

In *Nozomi Networks Query Language (N2QL)*, queries consist of:

- Data sources (on page 91)
- Commands (on page 96)
- Functions (on page 105)

Data sources

Queries start by calling a data source. For example:

```
nodes | sort received.bytes desc | head
```

This query will show, in table format, the first 10 nodes that received the most bytes. If you add the `pie` command at the end of the query, the results will show in a pie chart format, where each slice has `node id` as the label and the `received.bytes` field as data.
For example:

```
nodes | sort received.bytes desc | head | pie ip received.bytes
```

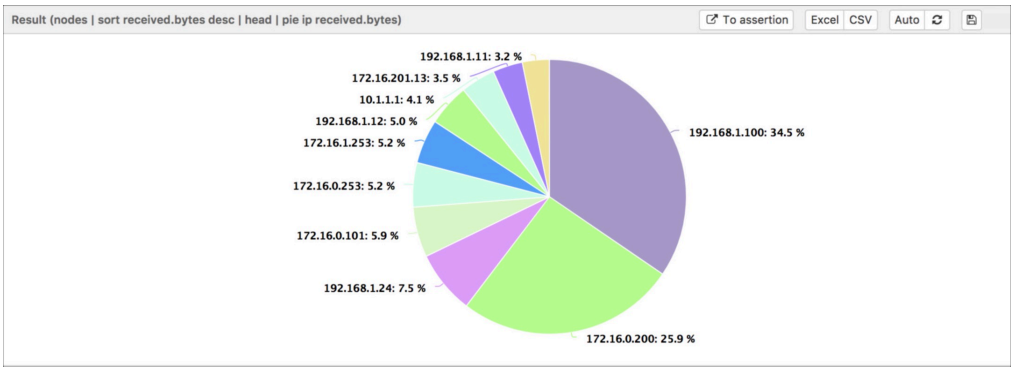


Figure 32. Queries example

Functions

You might not achieved your desired result just using queries. Consequently, query syntax supports functions. With functions, you can apply calculations to the fields and use the results as a new temporary field. For example, the query:

```
nodes | sort sum(sent.bytes,received.bytes) desc | column ip
 sum(sent.bytes,received.bytes)
```

uses the `sum` function to `sort` on the aggregated parameters, which produces a chart with the columns representing the sum of the `sent` and `received` bytes.

Prefix

The `$` is a prefix that changes the interpretation of the right hand side (rhs) of a `where` clause. By default, the rhs is interpreted as a string. With the `$` prefix, the interpretation of the rhs changes to a field name.

For example, in a query such as:

```
nodes | where id == 17.179.252.2
```

the right side of the `==` is expected to be a constant. If you create a query such as:

```
nodes | where id == id
```

the query tries to match all of the nodes having `id` equal to the string `id`.
If, however, you use the `$`, the second field is interpreted as a field, not a constant:

```
nodes | where id == $id
```

and returns the full list of records.

Data sources

These are the available data sources with which you can start a query.

alerts	Raised events
appliances	Downstream connected sensors synchronizing data to this, local one
assertions	Assertions saved by the users. An assertion represents an automatic check against other query sources
assets	Identified assets. Assets represent a local (private), physical system to care about, and can be composed of one or more Nodes. Broadcast nodes, grouped nodes, internet nodes, and similar cannot be Assets accordingly
audit_log	System's log for important operational events, e.g., login, backup creation, etc.
captured_files	Files reconstructed for analysis
captured_logs	Logs captured passively over the network
captured_urls	URLs and other protocol calls captured over the network. Access to files, requests to DNS, requested URLs and other are available in this query source
cpe_items	CPE maps definitions
cve_files	CVE definitions
dhcp_leases	IP to Mac bindings due to the presence of DHCP
function_codes	Protocols' function codes used in the environment
health_log	System's Health-related events, e.g. high resource utilization or hardware-related issues or events
link_events	Events that can occur on a Link, like it being available or not
links	Identified links, defined as directional one-to-one associations with a single protocol (i.e. source, destination, protocol)
microsoft_hotfixes	Microsoft hotfix information

node_cpe_changes	Common Platform Enumeration changes identified over known nodes. On the event of update of a CPE (on hardware, operating system and software versions), an entry in this query source is created to keep track of software updates or better detection of software
node_cpes	Common Platform Enumeration identified on nodes (hardware, operating system and software versions)
node_cves	Common Vulnerability Exposures: vulnerabilities associated to identified nodes' CPEs
node_points	Data points extracted over time, via Smart Polling or via Arc, from monitored Nodes
node_points_last	node_points last samples per each included data point
nodes	Identified nodes, where a node is an L2 or L3 (and above) entity able to speak some protocol
packet_rules	Packet rules definitions
protocol_connections	Identified protocol handhsakes/connections needed to decode process variables
report_files	Generated report files available for consultation
report_folders	Generated report folders
sessions	Sessions with recent network actvity. A Session is a specific application-level connection between nodes. A Link can hold one or more Session at a given time
sessions_history	Archived sessions
sigma_rules	Sigma rules definitions
sp_executions	Executions of Smart Polling plans
sp_node_executions	Results of Smart Polling plans executions per node
stix_indicators	STIX definitions
subnets	Identified network subnets
threat_models	Threat Modeling definitions
trace_requests	Trace requests in processing

variable_history	Process variables' history of values
variables	Identified process variables
yara_rules	YARA rules definitions
zone_links	A list of protocols exchanged by the defined zones
zones	Defined network zones

Basic operators

Operator	`|` (pipe, AND logical operator)
Description	Add a where clause with a logical AND, append it using the pipe character (`|`). For example, the query below returns links that are from 192.168.254.0/24 `AND` going to 172.217.168.0/24.
Example	**links	where from in_subnet? 192.168.254.0/24	where to in_subnet? 172.217.168.0/24**

Operator	`OR`
Description	To add a where clause with a logical OR, append it using the OR operator. For example, the query below returns links with either the http `OR` the https protocols.
Example	**links	where protocol == http OR protocol == https**

Operator	`!` (exclamation point, NOT logical operator)
Description	Put an exclamation point (`!`) before a term to negate it. For example, the query below returns links that do NOT (`!`) belong to 192.168.254.0/24.
Example	**nodes	where ip !in_subnet? 192.168.254.0/24	count**

Operator	`->`
Description	To change a column name, select it and use the `->` operator followed by the new name. It is worth noting that specific suffixes are parsed and used to visualize the column content differently. For example: • `_time` data is shown in a timestamp format (1647590986549 becomes 2022-03-18 09:09:46.549) • `_bytes` adds KB or MB, as applicable (50 becomes 50.0 B) • `_percent` adds a percentage sign (50 becomes 50%) • `_speed` adds a throughput speed in Mb/s (189915 becomes 1.8 Mb/s) • `_date` converts numbers into a date format (2022-06-22 15:43:31.297 becomes 2022-06-2214:24:09.280 becomes 2022-06-24 (current day)) • `_packets` adds pp after the number of packets (50 becomes 50 pp)
Example 1	`nodes	select created_at created_at->my_integer	where my_integer > 946684800000`
Example 2	`nodes	select created_at->my_creation_time`

Example 3	`nodes | select tcp_retransmission.bytes->my_retrans_bytes`

Operators	==, =, <, >, <=, and >=
Description	Queries support the mathematical operators listed above.

Operator	`"` (Quotation marks)
Description	Use quotation marks (`"`) to specify an empty string. Consider these two cases where this technique is useful: • Finding non-empty values. Example 1 below returns assets where the `os` field is not blank. • Specifying that a value in the query is a string (if its type is ambiguous). Example 2 below tells concat to treat the `"--"` parameter as a fixed string to use rather than as a field from the **alerts** table.
Example 1	`assets | where os != ""`
Example 2	`alerts | select concat(id_src,"--",id_dst)`

Operator	`in?`
Description	`in?` is only used with arrays; the field `type` must be an array. The query looks for the text strings you specify using `in?` and returns arrays that match one of them. The example below uses `in?` to find any node having `computer` or `printer` as elements in the array.
Example	`assets | where type in? ["computer","printer_scanner"]`

Operator	`include?`
Description	The query looks for the text string you specify using `include?` and returns strings that match it. The example below uses `include?` to find assets where the **os** field contains the string **Win**.
Example	`assets | where os include? Win`

Commands

Syntax	`select <field1> <field2> ... <fieldN>`
Parameters	• the list of field(s) to output
Description	The select command takes all the input items and outputs them with only the selected fields

Syntax	`exclude <field1> <field2> ... <fieldN>`
Parameters	• the list of field(s) to remove from the output
Description	The exclude command takes all the input items and outputs them without the specified field(s)

Syntax	**where <field> <==|!=|<|>|<=|>=|in?|!in?|include?|!include?|exclude?| start_with?|end_with?|in_subnet?> <value>**
Parameters	• field: the name of the field to which the operator will be applied • operator • value: the value used for the comparison. It can be a number, a string, or other data type. Advanced operators can use other data types, such as: ◦ a list (using JSON syntax) when using the **`in?`** operator, for example: `nodes | where ip in? ["172.18.41.44"]` ◦ a list (using JSON syntax) when using the **`include?`**, **`!include?`**, **`exclude?`**, **`start_with?`**, **`!start_with?`**, **`end_with?`**, or **`!end_with?`** operators. The condition matches if *any* element in the list matches; for negated operators, it matches only if *none* of the elements match. An empty list always returns no results for positive operators and all results for negated operators. For example: `assets | where name include? ["Switch","Router"]` ◦ another property when using the '$' symbol, for example: `nodes | where ip != $id`
Description	The where command will send to the output only the items which fulfill the specified criterion, many clauses can be concatenated using the boolean **OR** operator

Example	• `assets | where vendor !in? ["VMware","Cisco"]` • `nodes | where roles include? "consumer" OR zone == "office"` • `nodes | where label != "" | where label !include? ".net"` • `nodes | where ip in_subnet? "192.168.1.0/24"` • `<value>` can also be another `<field>`, as in: `links | where from_zone == $to_zone | select from_zone to_zone` • Match nodes whose name contains any of several substrings: `nodes | where name include? ["Switch","Router","Firewall"]` • Exclude nodes whose name starts with any of the given prefixes: `nodes | where name !start_with? ["test","lab","dev"]` • Find assets whose type ends with any of the given suffixes: `assets | where type end_with? ["controller","gateway"]` • Exclude assets whose vendor name contains any of the listed strings: `assets | where vendor exclude? ["unknown","generic"]`

Syntax	`sort <field> [asc|desc]`
Parameters	• field: the field used for sorting • asc|desc: the sorting direction
Description	The sort command will sort all the items according to the field and the direction specified, it automatically understands if the field is a number or a string

Syntax	`group_by <field> [ [avg|sum] [field2] ]`
Parameters	• field: the field used for grouping • avg|sum: if specified, the relative operation will be applied on field2
Description	The group_by command will output a grouping of the items using the field value. By default the output will be the count of the occurrences of distinct values. If an operator and a **field2** are specified, the output will be the average or the sum of the **field2** values

Syntax	`head [count]`
Parameters	• count: the number of items to output
Description	The head command will take the first **count** items, if **count** is not specified the default is 10

Syntax	`uniq [<field1> <field2> ... <fieldN>]`
Parameters	• an optional list of fields on which to calculate the uniqueness
Description	The uniq command will remove from the output the duplicated items

Syntax	`expand <field>`
Parameters	• field: the field containing the list of values to be expanded
Description	The expand command will take the list of values contained in **field** and for each of them it will duplicate the original item substituting the original **field** value with the current value of the iteration

Syntax	`expand_recursive <field>`
Parameters	• field: the field to be recursively expanded
Description	The expand_recursive command will recursively parse the content of **field**, expanding each array or json structure until a scalar value is found. It generates a new row for each array element or json field. For each new row, it duplicates the original item substituting the original **field** value with the current value of the iteration and adding a new field that represents the current iteration path from the root

Syntax	`sub <field>`
Parameters	• field: the field containing the list of objects
Description	The sub command will output the items contained in **field**

Syntax	`count`
Parameters	
Description	The count command outputs the number of items

Syntax	`pie <label_field> <value_field>`
Parameters	• label_field: the field used for each slice label • value_field: the field used for the value of the slice, must be a numeric field

Description	The pie command will output a pie chart according to the specified parameters

Syntax	`column <label_field> <value_field ...>`
Parameters	• label_field: the field used for each column label • value_field: one or more field used for the values of the columns
Description	The column command will output a histogram; for each label a group of columns is displayed with the value from the specified value_field(s). The variant `column_colored_by_label` returns bars of different colors depending on their labels.

Syntax	`history <count_field> <time_field>`
Parameters	• count_field: the field used to draw the Y value • time_field: the field used to draw the X points of the time series
Description	The history command will draw a chart representing an historic series of values

Syntax	`distance <id_field> <distance_field>`
Parameters	• id_field: the field used to tag the resulting distances. • distance_field: the field on which distances are computed among entries.
Description	The distance command calculates a series of distances (that is, differences) from the original series of `distance_field`. Each distance value is calculated as the difference between a value and its subsequent occurrence, and tagged using the **id_field**. For example, assuming we're working with an **id** and a **time** field, entering **alerts | distance id time** returns a table where each distance entry is characterised by the **from_id**, **to_id**, and **time_distance** fields that represent time differences between the selected alerts.

Syntax	`bucket <field> <range>`
Parameters	• field: the field on which the buckets are calculated • range: the range of tolerance in which values are grouped

Description	The bucket command will group data in different buckets, different records will be put in the same bucket when the values fall in the same multiple of <range>

Syntax	`join <other_source> <field> <other_source_field>`
Parameters	• other_source: the name of the other data source • field: the field of the original source used to match the object to join • other_source_field: the field of the other data source used to match the object to join
Description	The join command will take two records and will join them in one record when <field> and <other_source_field> have the same value

Syntax	`gauge <field> [min] [max]`
Parameters	• field: the value to draw • min: the minimum value to put on the gauge scale • max: the maximum value to put on the gauge scale
Description	The gauge command will take a value and represent it in a graphical way

Syntax	`value <field>`
Parameters	• field: the value to draw
Description	The value command will take a value and represent it in a textual way

Syntax	`reduce <field> [sum|avg]`
Parameters	• field: the field on which the reduction will be performed • sum or avg: the reduce operation to perform, it is sum if not specified
Description	The reduce command will take a series of values and calculate a single value

Syntax	size()
Parameters	• field: the field to calculate the size of

Description	If the field is an array, then the size function returns the number of entries in the array. If the field contains a string, then the size function returns the number of characters in the string. **Note**: The size function may only be used on the following data sources: alerts, assets, captured_files, links, nodes, packet_rules, sessions, stix_indicators, subnets, variables, yara_rules, zones, and zone_links.
Example:	`assets | where size(ip) > 1`

Nodes-specific commands reference

Syntax	`where_node <field> < ==|!=|<|>|<=|>=|in?|include?|exclude?| start_with?|end_with? > <value>`
Parameters	• field: the name of the field to which the operator will be applied • operator • value: the value used for the comparison. It can be a number, a string or a list (using JSON syntax), the query engine will understand the semantics.
Description	The where_node command will send to the output only the items which fulfill the specified criterion, many clauses can be concatenated using the boolean **OR** operator. The where_node command is similar to the where command, but the output will also include all the nodes that are communicating directly with the result of the search. **Note**: This command is only applicable to the nodes table.

Syntax	`where_link <field> < ==|!=|<|>|<=|>=|in?|include?|exclude?| start_with?|end_with? > <value>`
Parameters	• field: the name of the links table's field to which the operator will be applied. • operator • value: the value used for the comparison. It can be a number, a string or a list (using JSON syntax) the query engine will understand the semantics.
Description	The where_link command will send to the output only the nodes which are connected by a link fulfilling the specified criterion. Many clauses can be concatenated using the boolean **OR** operator. **Note**: This command is only applicable to the nodes table.

Syntax	`graph [node_label:<node_field>]`
	`[node_perspective:<perspective_name>]`
	`[link_perspective:<perspective_name>]`

Parameters	• node_label: add a label to the node, the label will be the content of the specified node field • node_perspective: apply the specified node perspective to the resulting graph. Valid node perspective values are: ◦ roles ◦ zones ◦ transferred_bytes ◦ not_learned ◦ public_nodes ◦ reputation ◦ appliance_host • link_perspective: apply the specified link perspective to the resulting graph. Valid link perspectives are: ◦ transferred_bytes ◦ tcp_firewalled ◦ tcp_handshaked_connections ◦ tcp_connection_attempts ◦ tcp_retransmitted_bytes ◦ throughput ◦ interzones ◦ not_learned
Description	The graph command renders a network graph by taking some nodes as input.

Link-events-specific commands reference

Syntax	`availability`
Parameters	
Description	The availability command computes the percentage of time a link is UP. The computation is based on the link events UP and DOWN that are seen for the link.

Syntax	`availability_history <range>`
Parameters	• range: the temporal window in milliseconds to use to group the link events
Description	The availability_history command computes the percentage of time a link is UP by grouping the link events into many buckets. Each bucket will include the events of the temporal window specified by the range parameter.

Syntax	`availability_history_month <months_back> <range>`
Parameters	• months_back: number of months to go back in regards to the current month to group the link events • range: the temporal window in seconds to use to group the link events
Description	The availability_history command computes the percentage of time a link is UP by grouping the link events into many buckets. Each bucket will include the events of the temporal window specified by the range and months parameters.

Functions

Functions are always used in conjunction with other commands, such as `select`. In the following examples, functions are shown in **bold**:

- Combining functions with `select: nodes | select id type` **`color`**`(type)`
- Combining functions with `where: nodes | where` **`size`**`(label) > 10`
- Combining functions with `group_by: nodes | group_by` **`size`**`(protocols)`

Here is the complete list of functions:

Syntax	`abs(<field>)`
Parameters	• the field on which to calculate the absolute value
Description	The abs function returns the absolute value of the field

Syntax	`bitwise_and(<numeric_field>,<mask>)`
Parameters	• numeric_field: the numeric field on which apply the mask • mask: a number that will be interpreted as a bit mask
Description	The bitwise_and function calculates the bitwise & operator between the numeric_field and the mask entered by the user

Syntax	`cast(<field>,<type>)`
Parameters	• field: the field to cast • type: one of: integer (4 bytes integer), bigint (8 bytes integer), real (4 bytes floating point number), boolean, text, json
Description	The cast function converts a field to the specified type

Syntax	`coalesce(<field1>,<field2>,...)`
Parameters	• a list of fields or string literals in the format "<chars>"
Description	The coalesce function will output the first value that is not null

Syntax	`color(<field>)`
Parameters	• field: the field on which to calculate the color
Description	The color function generates a color in the rgb hex format from a value

Note	Only available for nodes, links, variables and function_codes

Syntax	`concat(<field1>,<field2>,...)`
Parameters	• a list of fields or string literals in the format "<chars>"
Description	The concat function will output the concatenation of the input fields or values

Syntax	`date(<time>)`
Parameters	• time defined as unix epoch
Description	The date function returns a date from a raw time

Syntax	`day_hour(<time_field>)`
Parameters	• time_field: the field representing a time
Description	The day_hour function returns the hour of the day plus the sensor's local time offset from UTC, i.e. a value in the range 0 through 23. Be careful when accounting for daylight saving time. Use **day_hour_utc** when absolute precision is desired

Syntax	`day_hour_utc(<time_field>)`
Parameters	• time_field: the field representing a time
Description	The **day_hour_utc** function returns the hour of the day expressed in UTC for the current time field, i.e. a value in the range 0 through 23

Syntax	`days_ago(<time_field>)`
Parameters	• time_field: the field representing a time
Description	The days_ago function returns the amount of days passed between the current time and the time field value

Syntax	`dist(<field1>,<field2>)`
Parameters	• the two fields to compute the distance on

Description	The dist function returns the distance between field1 and field2, which is the absolute value of their difference

Syntax	`div(<field1>,<field2>)`
Parameters	• field1 and field2: the two field to divide
Description	The div function will calculate the division field1/field2

Syntax	`hours_ago(<time_field>)`
Parameters	• time_field: the field representing a time
Description	The hours_ago function returns the amount of hours passed between the current time and the time field value

Syntax	`is_empty(field) == true | false`
Parameters	• field: the field to check to evaluate whether it is empty or not
Description	The is_empty command takes a field as input and returns only the entries that are either empty / not empty.
Example	**nodes | where is_empty(label) == false**

Syntax	`is_recent(<time_field>)`
Parameters	• time_field: the field representing a time
Description	The is_recent function takes a time field and returns true if the time is not farther than 30 minutes

Syntax	`minutes_ago(<time_field>)`
Parameters	• time_field: the field representing a time
Description	The minutes_ago function returns the amount of minutes passed between the current time and the time field value

Syntax	`mult(<field1>,<field2>,...)`
Parameters	• a list of fields to multiply

Description	The mult function returns the product of the fields passed as arguments

Syntax	`round(<field>,[precision])`
Parameters	• field: the numeric field to round • precision: the number of decimal places
Description	The round function takes a number and outputs the rounded value

Syntax	`seconds_ago(<time_field>)`
Parameters	• time_field: the field representing a time
Description	The seconds_ago function returns the amount of seconds passed between the current time and the time field value

Syntax	`split(<field>,<splitter>,<index>)`
Parameters	• field: the field to split • splitter: the character used to separate the string and produce the tokens • index: the 0 based index of the token to output
Description	The split function takes a string, separates it and outputs the token at the <index> position

Syntax	`sum(<field>,...)`
Parameters	• a list of fields to sum
Description	The sum function returns the sum of the fields passed as arguments

Syntax	`to_time(<field>)`
Parameters	• field: a time field, or a string literal in ISO-8601 format, e.g. "2024-01-15T10:30:00Z"
Description	The to_time function converts a time field, or an ISO-8601 datetime string, into the numeric timestamp (epoch milliseconds) used by time fields. If the datetime string does not specify a UTC offset, it is assumed to be UTC

Examples

Pie chart

An example on how to create a pie chart to understand the media access control (MAC) vendor distribution in a network.

We choose nodes as our query source and we start to group the nodes by mac_vendor:

```
nodes | group_by mac_vendor
```

We can see the list of the vendors in our network associated with the occurrences count. To better understand our data we can use the sort command, so the query becomes:

```
nodes | group_by mac_vendor | sort count desc
```

In the last step we use the pie command to draw the chart with the mac_vendor as a label and the count as the value.

```
nodes | group_by mac_vendor | sort count desc | pie mac_vendor count
```

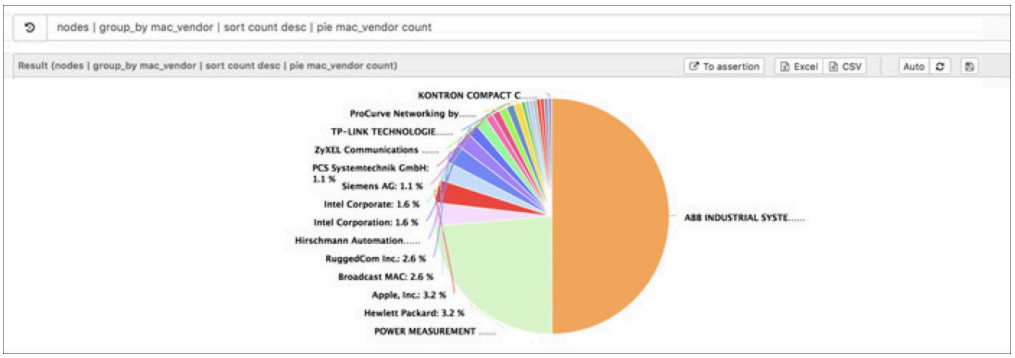


Figure 33. Pie chart example

Column chart

An example on how to create a column chart with the top nodes by traffic.

To start, you need to get the nodes and select the:

- `id`
- `sent.bytes`
- `received.bytes`
- `sent.bytes`
- `received.bytes`

To calculate the sum , you need to use the `sum` function. The query is:

```
nodes | select id sent.bytes received.bytes
 sum(sent.bytes,received.bytes)
```

When you execute this query, the sum field has a very long name. You can rename it to be more comfortable with these commands:

```
nodes | select id sent.bytes received.bytes
 sum(sent.bytes,received.bytes)->sum
```

To obtain the top nodes by traffic, you can sort and take the first 10:

```
nodes | select id sent.bytes received.bytes
 sum(sent.bytes,received.bytes)->sum | sort sum desc | head 10
```

Finally, to display the data in a graphical way, you can use the `column` command:

```
nodes | select id sent.bytes received.bytes
 sum(sent.bytes,received.bytes)->sum | sort sum desc | head 10 | column
 id sum sent_bytes received_bytes
```

Note:
You can access an inner field of a complex type with the dot syntax, in the example the dot syntax is used on the fields `sent` and `received` to access their bytes sub field.

Note:
After accessing a field with the dot syntax, it will gain a new name to avoid ambiguity; the dot is replaced by an underscore. In the example `sent.bytes` become `sent_bytes`

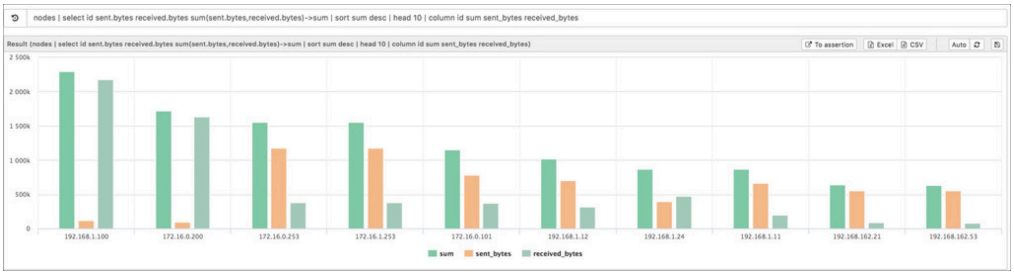


Figure 34. Column chart example

Where with multiple conditions in OR

An example of a query to get all the nodes with a specific role, in particular all the nodes which are web or domain name server (DNS) servers.

With the `where` command, you can separate many conditions with `OR`

Note:
Because the roles field contains a list of values, you can use the `include?` operator to check if a value was contained in the list.

```
nodes | where roles include? web_server OR roles include? dns_server |
 select id roles
```



Figure 35. Where with multiple conditions in OR example

Bucket and history

An example of a query to calculate the distribution of link events towards an internet protocol (IP) address.

You can filter all the `link_events` with `id_dst` equal to `192.168.1.11` After this you can sort by time, this is a very important step because bucket and history depend on how the data are sorted.

Then you can use `bucket` to group the data by time. The final step is to use the `history` command to draw a chart, we pass `count` as a value for the Y axis and `time` for the X axis.

The `history` command is particularly suited for displaying a big amount of data, in the image below we can see that there are many hours of data to analyze.

```
link_events | where id_dst == 192.168.1.11 | sort time asc | bucket time
 36000 | history count time
```

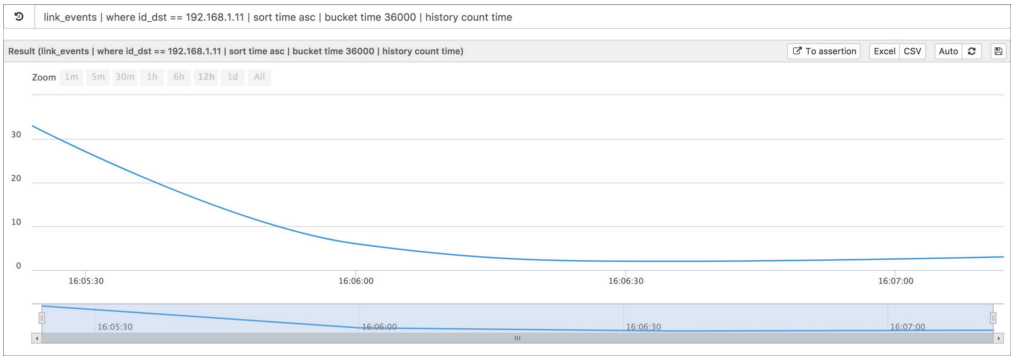


Figure 36. Bucket and history example

Join

An example query to join two data sources to obtain a new data source with more information. In particular, how to list the links with the labels for the source and destination nodes.

You can match the `from` field of the links with the `id` field of the nodes to ask for the links, and join them with the nodes:

```
links | join nodes from id
```

After executing the query above you will get all the links fields, plus a new field called `joined_node_from_id`, it contains the node which satisfies the `link.from == node.id` condition. You can use the dot syntax to access the sub fields of `joined_node_from_id`

Because we also want to get the labels for the `to` field of the `links` you add another `join` and exclude the empty labels of the node referred by `to` to get more interesting data:

```
links | join nodes from id | join nodes to id | where
 joined_node_to_id.label != ""
```

This will obtain a huge amount of data, which is difficult to understand. To only get the relevant information, you can use a `select`:

```
links | join nodes from id | join nodes to id | where
 joined_node_to_id.label != "" | select from joined_node_from_id.label to
 joined_node_to_id.label protocol
```

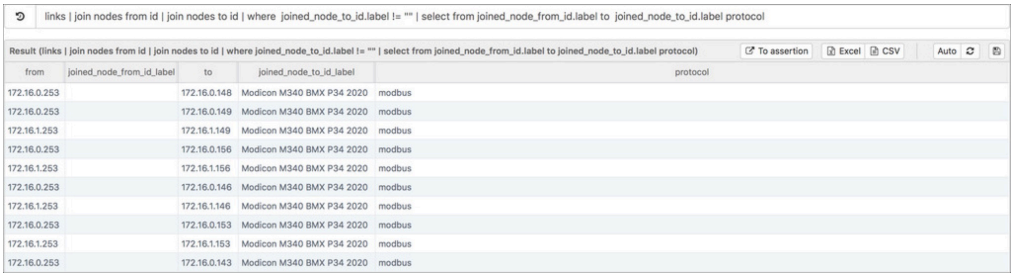


Figure 37. Join example

Compute the availability history

An example query to compute the availability history for a link.

In order to achieve a reliable availability, it is recommended to enable the **Track availability** feature on the desired link.

Start from the `link_events` data source, filtered by source and destination ip in order to precisely identify the target link. Consider also filtering by protocol to achieve a higher degree of precision.

```
link_events | where id_src == 10.254.3.9 | where id_dst == 172.31.50.2
```

The next step is to sort the events by ascending time of creation. Without this step the `availability_history` might produce meaningless results, such as negative values. Then, to compute the `availability_history` with a bucket of 1 minute (60000 milliseconds), you can complete query is as follows:

```
link_events | where id_src == 10.254.3.9 | where id_dst == 172.31.50.2 |
 sort time asc | availability_history 60000
```

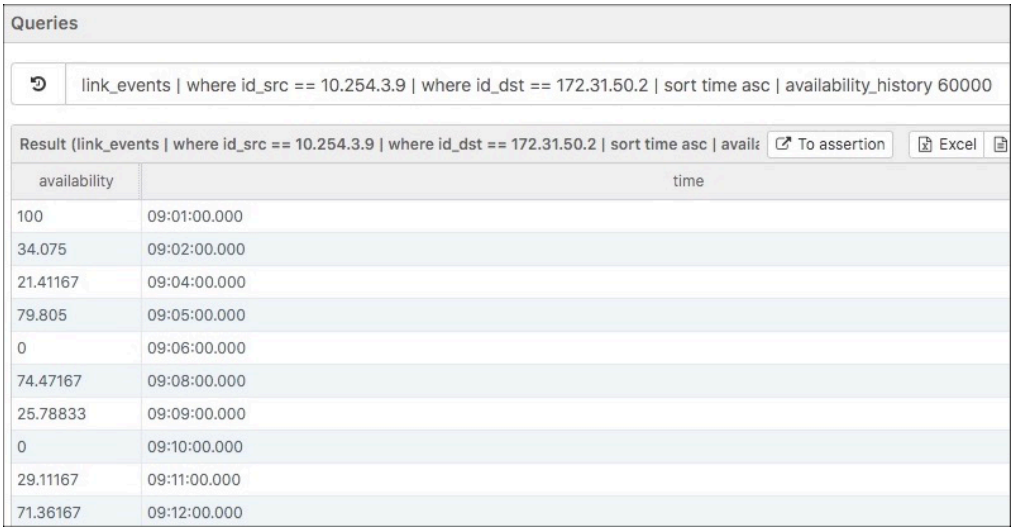


Figure 38. Availability history example

Note:
By default, `link_events` generation is disabled. To enable it, you can use the configuration rule described in **Configure links**.

Complex field types

Single scalar values

To query single scalar values, apply the commands that are explained in this section.

Objects

Objects show in braces: {object}

```
{
"source": "ARP",
"likelihood": 1,
"likelihood_level": "confirmed"
}
```

An example on how to query only `confirmed` *MAC* addresses.

> **Note:**
> Possible values are:
> - `confirmed`
> - `likely`
> - `not confirmed`

Since `mac_address:info` is an object, you can access subfields like `mac_address:info.likelihood_level` to apply the `where` condition:

```
nodes | select mac_address:info mac_address:info.likelihood_level | where
 mac_address:info.likelihood_level == confirmed
```

Since N2OS 24.1 is possible to access complex objects with a different syntax that is compatible with Vantage, using the `/` operator, the query specified above becomes:

```
nodes | select mac_address:info/likelihood_level | where
 mac_address:info.likelihood_level == "confirmed"
```

Note that also the `"confirmed"` literal can now be quoted and the query can be executed in Vantage without any change.

Arrays

> **Note:**
> For example, a `parent` in the alerts table.

Arrays show in braces: {array}

```
[
"5b867836-2b41-4c15-ab6f-4ae5f0251e30"
]
```

An example on how to only query alerts that have a parent incident, with a known incident id with the value: `d36d0`

Since the `parents` field is an array, you can use `expand` first to get an entry for each parent, then apply your condition:

```
alerts | expand parents | where expanded_parents include? d36d0
```

Object arrays

> **Note:**
> For example, `function_codes` in the links table.

Object arrays are a combination of the above examples. Therefore, they show an object included in a [{..},{..},..] :

```
[
{
"name": "M-SEARCH",
"is_learned": true,
"is_fully_learned": true
}
]
```

An example on how to query learned function codes.

Since `function_codes` is an object array, you can use `expand` first, to get an entry for each function code, then use the `.` operator (`function_code.is_learned`) to apply your `where` condition:

```
links | select from to protocol function_codes | expand function_codes |
 where expanded_function_codes.is_learned == true
```



Chapter 6. Smart Polling in Guardian



*The **Smart Polling** page lets you view and manage Smart Polling.*

Smart Polling has these tabs:

- Overview (on page 120)
- Activity Log (on page 163)

> **Note:**
> Arc sensors execute Smart Polling when Guardian requests it. Enable execution in the Arc local UI under **Execution options**.

Active sensor

Smart Polling on Guardian selects the sensor or sensors for execution based on the `capture_devices` field of the polled node. Smart Polling will engage the corresponding sensors. When at least one successful execution occurs from a sensor, Smart Polling remembers to use that sensor in future executions. If a future execution fails, or Smart Polling records no executions for more than a week, it reverts to the sensors in the `capture_devices` field.

To add a sensor to the applicable sensors for a node, you can use the Guardian *CLI* to modify the `capture_devices` field. For example, to add an Arc to the applicable sensors for a node, use `vi node 192.168.1.1 capture_device arc[1e6a174c]`, where `192.168.1.1` is the node *ID* and the value in brackets is the first eight characters of the Arc sensor's *ID*. The manually added sensor is then added to the node's `capture_devices` list, and may be chosen according to the algorithm described above. If a node is only visible to Arc sensors but you want to use Guardian instead, use `vi node 192.168.1.1 capture_device guardian`.

The *OS* of the sensor that polls the node decides which network interface to use based on the routing table. All Smart Polling strategies are based on the *IP*, which means that all routing decisions are delegated to the *OS*.

Overview

Learn how Smart Polling plans are displayed, managed, and executed in Guardian. It covers plan statuses, how to add or search for plans, and how to interpret result data. Icons and table fields provide insight into device polling outcomes and plan performance.

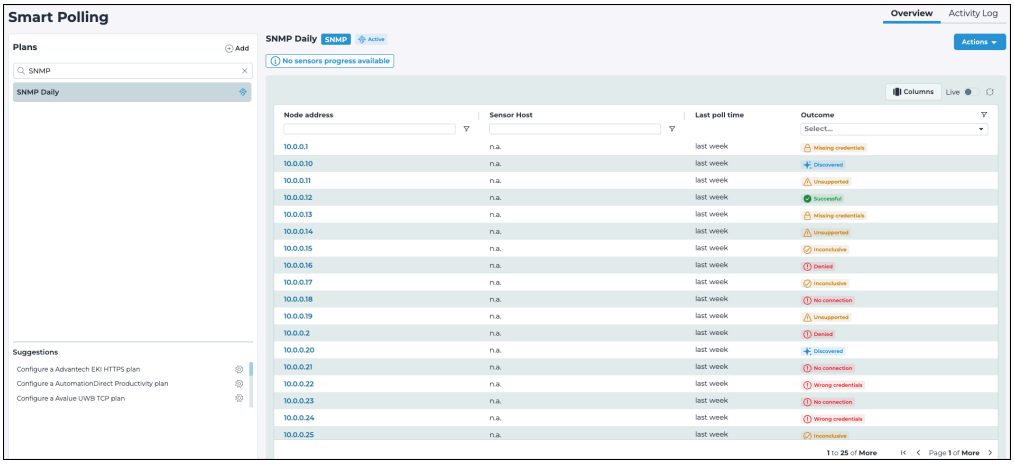


Figure 39. Populated Plans page

Plans

The **Plans** section shows a list of all the Smart Polling plans that have been added for your organization. The icons show the current status of the plan:

- : Active
- : Paused

Add button

The ⊕ **Add** icon lets you add a new Smart Polling plan.

Search field

The search field lets you easily find a specific Smart Polling plan.

Results

After a Smart Polling plan runs, its results show on the plan pages. When you select a specific plan, the table shows the:

- Status of the plan through the `ACTIVE` | `INACTIVE` indicator
- Result of the last polling of each node polled within the last two weeks

The table columns show:

- **Node address**: The *IP* address of the node that was detected
- **Sensor Host**: The name of the sensor

- **Last poll time**: The time since the last execution of the plan for that node
- **Outcome**: The result of the last execution of the plan, which can be:
  - `Successful`: Smart Polling was able to poll the host and enrich the corresponding asset
  - `No connection`: Verify that the device is up, its services are running, that the sensor has a route to that host and no firewall is blocking the connection
  - `Missing credentials`: Add the credentials to the **Credentials Manager**
  - `Wrong credentials`: Update the credentials in the **Credentials Manager**
  - `Inconclusive`: This device is unsupported or only partly supported by Smart Polling
  - `Unsupported`: This strategy is not supported by the sensor that tried to poll the device
  - `Queued`: The sensor has not yet picked up the node for execution
  - `Denied`: The node lies in a subnet that Smart Polling does not have the permission to reach

For more details about the additional information that Smart Polling can provide, see Smart Polling results (on page 160).

Plan creation

Learn how to create a Smart Polling plan in Guardian by choosing a strategy, defining parameters, adding targets and credentials, and scheduling when data collection runs.

Select the ⊕ **Add** icon to enter the plan creation wizard.

Strategy

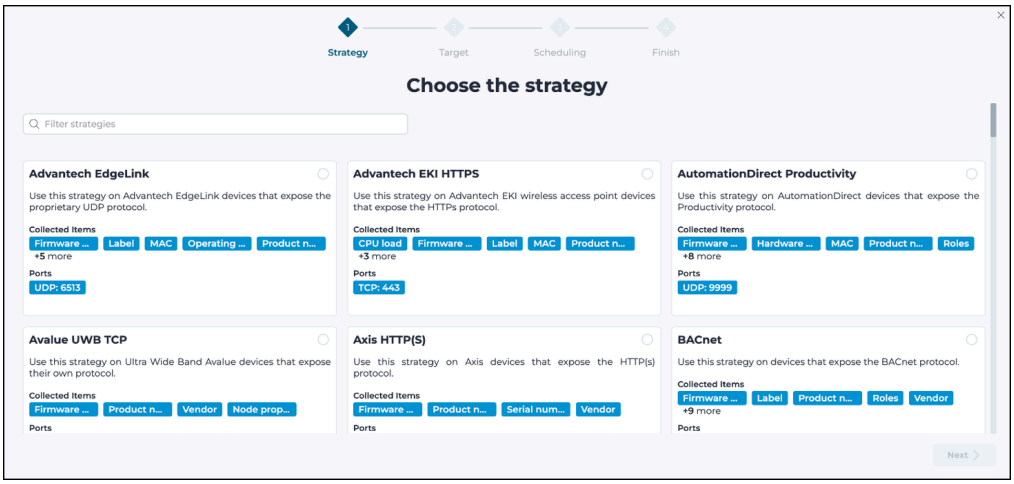


Figure 40. Choose the strategy page

Filter strategies: This field lets you enter text to easily find a specific strategy.

All available strategies show on the page with a summary. The summary section shows the:

- Name of the strategy
- A description of when it is useful to use the strategy
- The information that is collected
- Ports used

Parameters

Customize the plan with parameters used by the selected strategy, such as the port to use. Not all strategies accept parameters, in which case this step is skipped.

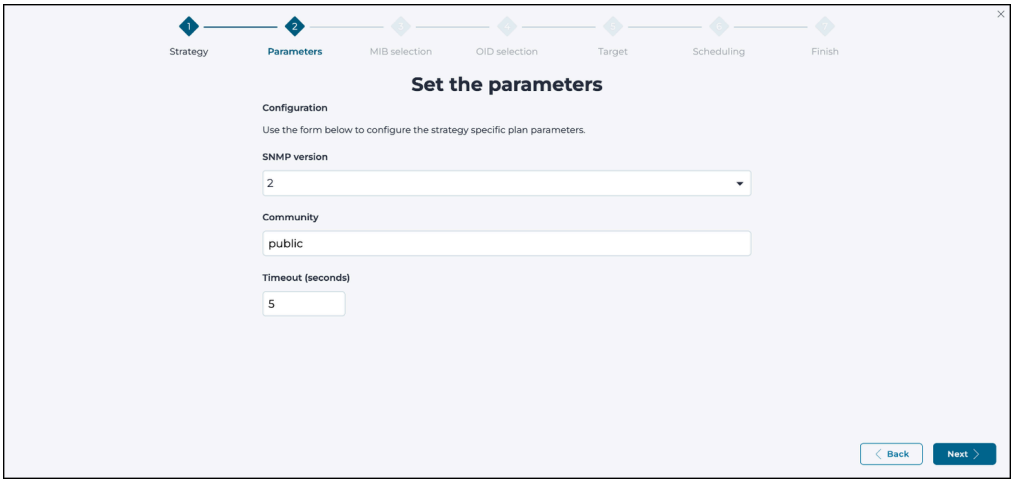


Figure 41. Parameters page

Choose the MIB files

This step shows only for *simple network management protocol (SNMP)* plans. The **MIB selection** page lists the *management information base (MIB)* files that Smart Polling can use. A *MIB* file defines the *object identifier (OID)*s of an *SNMP* device: their names, their types, and where they are in the tree. You choose the *OID*s to poll from these files on the next page.

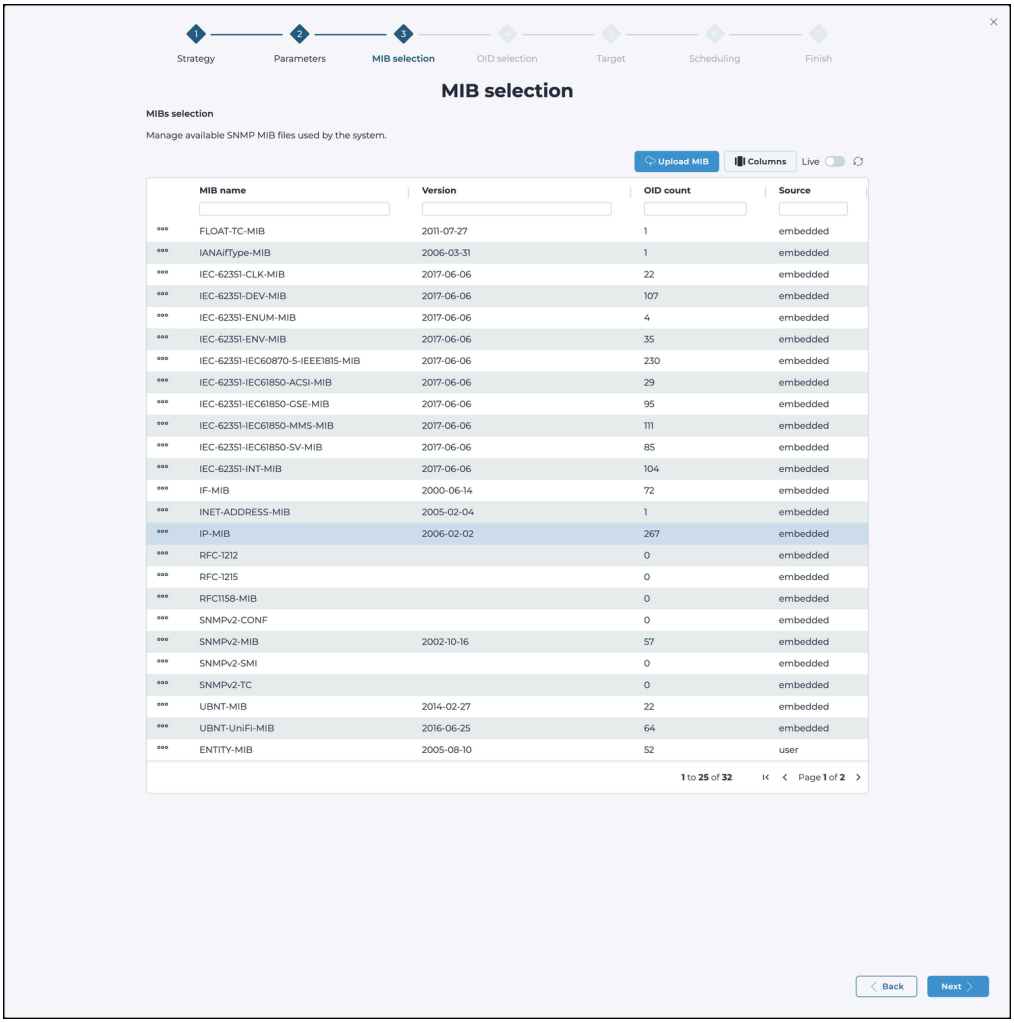


Figure 42. Available MIB files

For each *MIB* file, the page shows the:

- **MIB name**
- **Version**
- **OID count**: The number of *OID*s in the file
- **Source**: Where the file came from. The value is `embedded` for files that are included with the product, and `user` for files uploaded by users

Above the list, you can use the:

- **Upload MIB** button: Add your own *MIB* files. For the procedure, see Upload a MIB file (on page 136)
- **Columns** button: Choose the columns that show in the list
- **Live** toggle: Turn the automatic refresh of the list on or off
- Refresh icon: Get the *MIB* files that are on the sensor at this moment, and update the list

You can also enter text in the box under a column heading to filter the list.

Each row has a menu with a **Delete** option. You cannot delete the *MIB* files that are included with the product.

The Guardian stores the files that you upload, so the same files are available when you create or edit any *SNMP* plan.

Choose how to use OID values

This step also shows only for *SNMP* plans. On the **OID selection** page, browse the *OID*s from the *MIB* files, and choose how to use each one: collect it as a node point, walk it on every poll, or map its value to a node attribute.

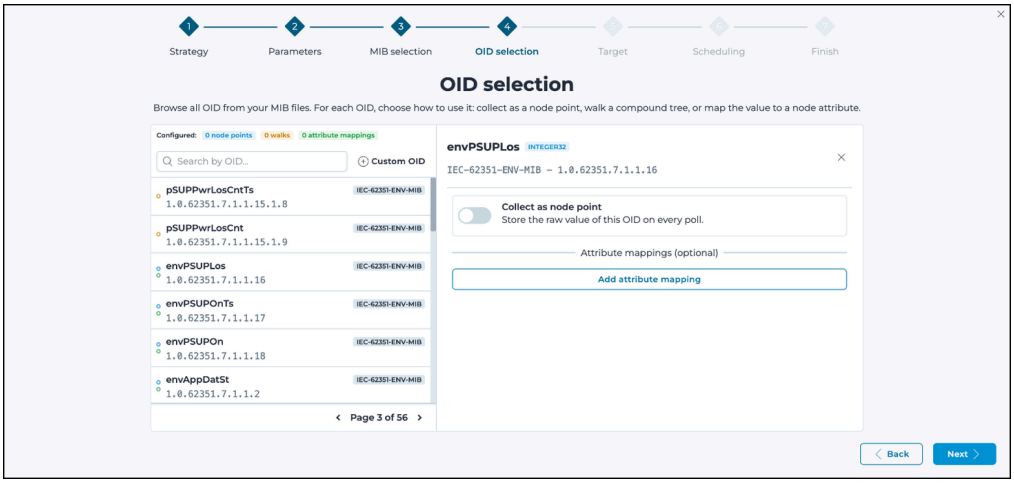


Figure 43. Selected OID and its options

The left side of the page shows:

- **Configured**: A summary of how many node points, walks, and attribute mappings you have configured
- **Search by OID**: A field that filters the list as you type. You can enter the number of an *OID*
- **Custom OID**: A button that you can use to add an *OID* that is not in a *MIB* file. For the procedure, see Add a custom OID to the list (on page 139)
- A paged list of the *OID*s. Each entry shows the name of the *OID*, the *MIB* file that it is from, and its number

The coloured dots next to each *OID* show the ways that you can use it, in the same colours as the **Configured** summary: blue for a node point of a single-value *OID*, orange for a walk, and green for an attribute or custom field mapping. A dot is filled when you have enabled that use.

Select an *OID* in the list to show its details on the right: the name, the type, and the *MIB* file and number that it is from. When you have not selected an *OID*, the panel shows **Select an OID to configure its mapping**.

The toggle in the panel turns the collection of the *OID* on and off. Its label depends on the type of the *OID*:

- For an *OID* that holds a single value: **Collect as node point**. When the toggle is on, Smart Polling stores the raw value of the *OID* on every poll
- For an *OID* that represents a *MIB* node or a table: **Walk this OID on every poll**. Select the toggle to traverse all child *OID*s on each poll; every value that the walk finds becomes a separate node point

Under **Attribute mappings (optional)**, you can map the value of the *OID* to one or more node attributes. The mappings are independent of the selection toggle. For the procedure, see Map an OID to a node attribute (on page 137).

Note:

Smart Polling walks at most 200 custom *OID*s in one polling session.

Target

Some of the strategies that are available require you to enter the related credentials. Therefore, there are two possible versions of the **Target** page that can show.

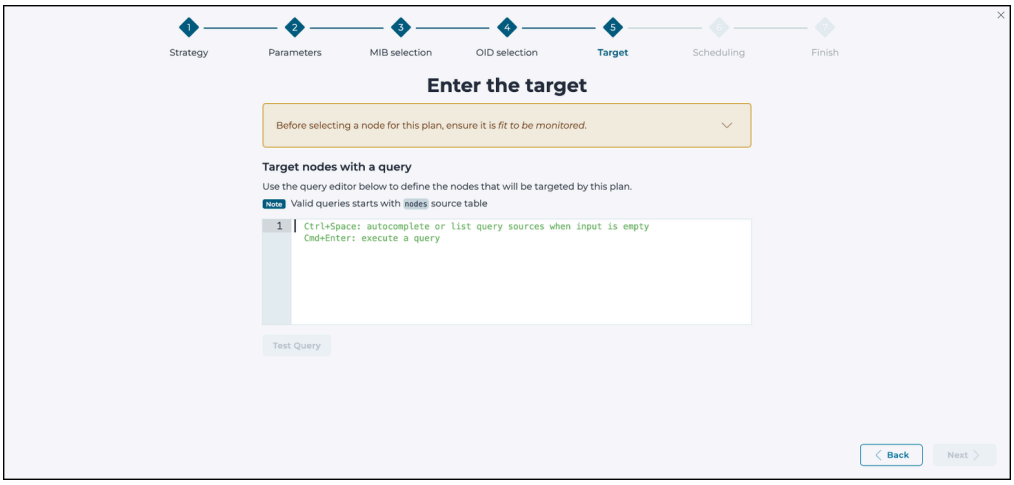


Figure 44. Query version of the target page

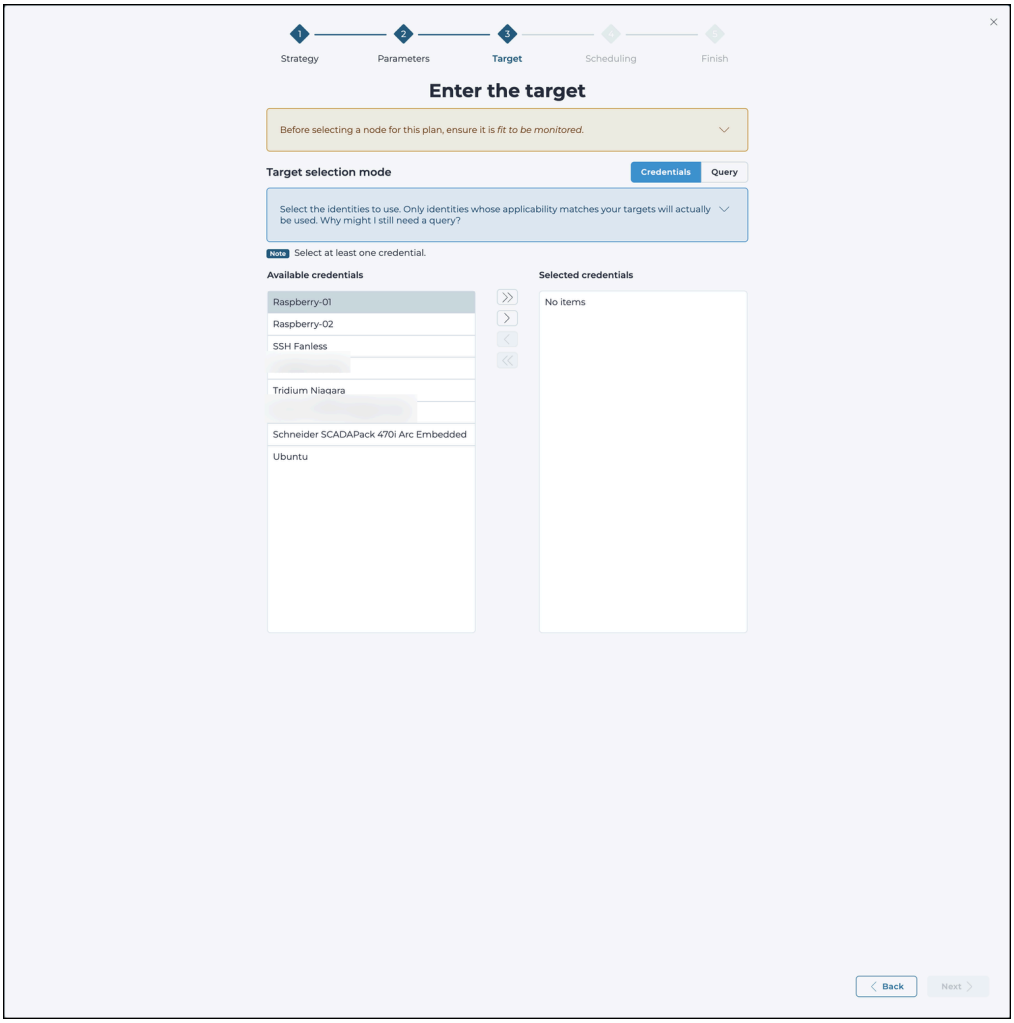


Figure 45. Credentials version of the target page

For strategies that use credentials, the page also has a **Target selection mode** switch with the options **Credentials** and **Query**. Select **Query** to target nodes with a query instead. If no credentials are available for the strategy, the page tells you to add one in **Administration > Credentials Manager**.

For strategies that require credentials, you must select one, or more, of the available credentials that show in the panel on the left. To move the credentials to the panel on the right, you can use the:

- > icon to select one credential
- >> icon to select all credentials

Scheduling

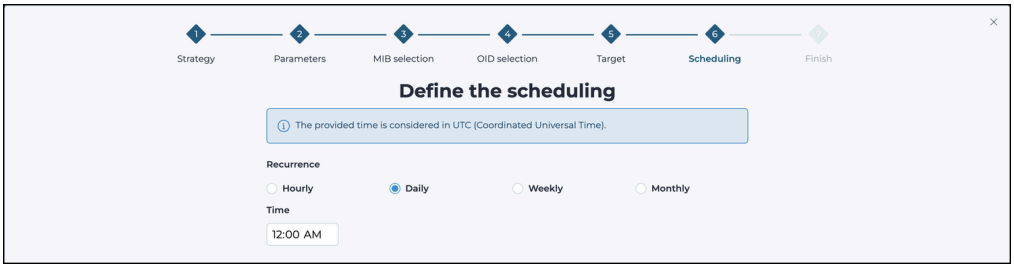


Figure 46. Scheduling page

Under **Recurrence**, you must select the granularity of the scheduling, choosing among:

- **Hourly**
- **Daily**
- **Weekly**
- **Monthly**

Time: You then must set the time that the Smart Polling plan will be executed. The time is set in .

Finish

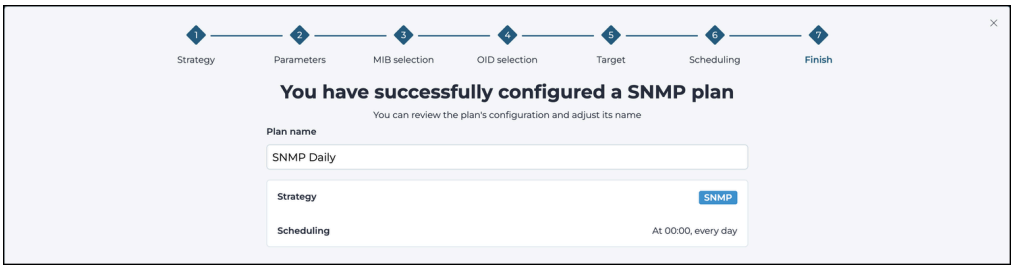


Figure 47. Finish page

On the **Finish** page, you can edit the **Plan name**.

Add a Smart Polling plan that requires credentials

Create a Smart Polling plan in Guardian that uses a strategy that requires credentials. Choose specific or all credentials, configure a schedule, and complete the setup.

Before you begin

Make sure that the correct credentials for the Smart Polling plan you are about to add are in the **Credentials Manager**.

About this task

Use this process to create a Smart Polling plan that targets systems requiring authentication. This method lets Guardian access and collect data from sources that need credentials. You can configure the schedule of the plan to suit your operational needs.

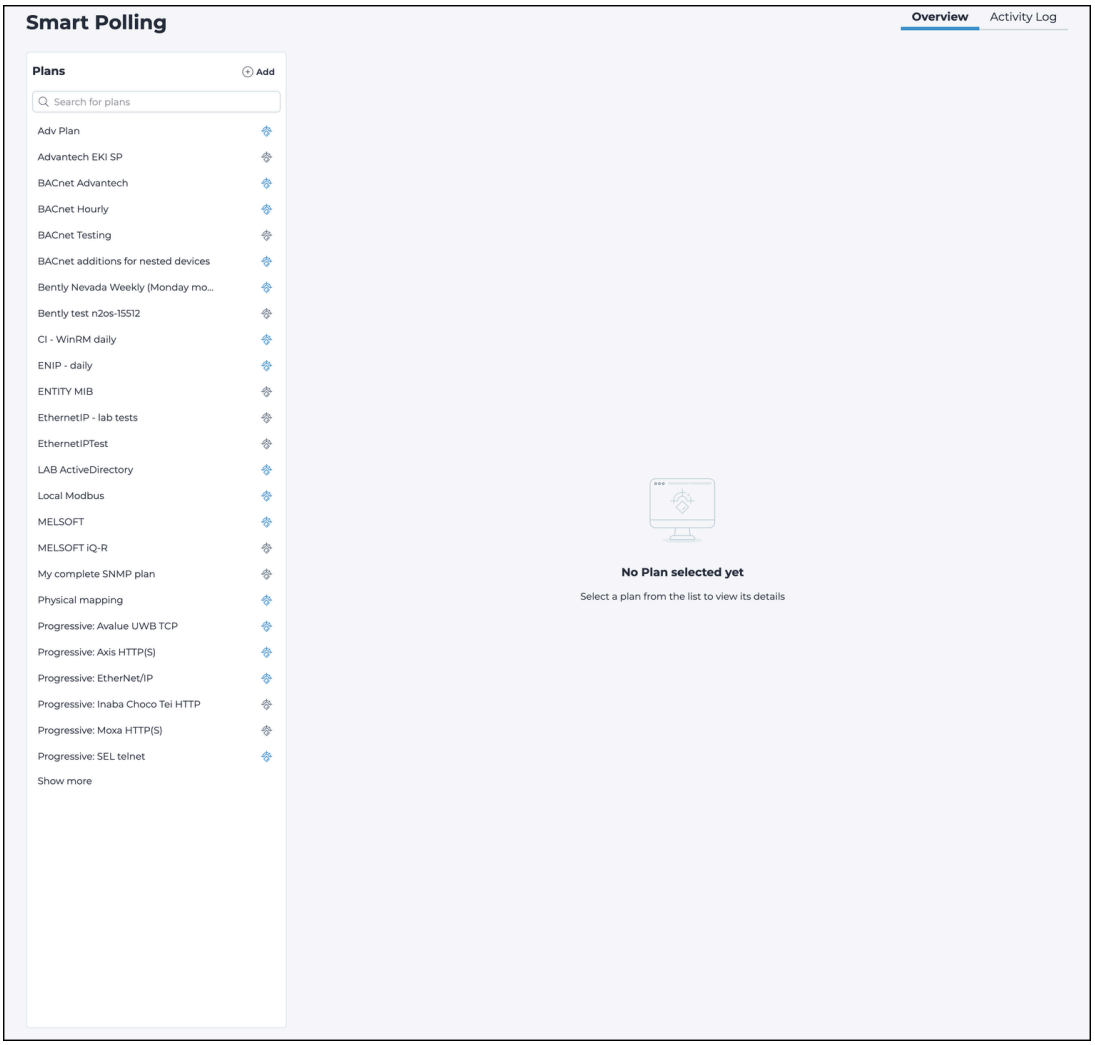
Procedure

1. In the top navigation bar, select ☰ > **Smart Polling**.

   Result: The **Smart Polling** page opens.

2. If the welcome page shows, select **Go to the Smart Polling page**.

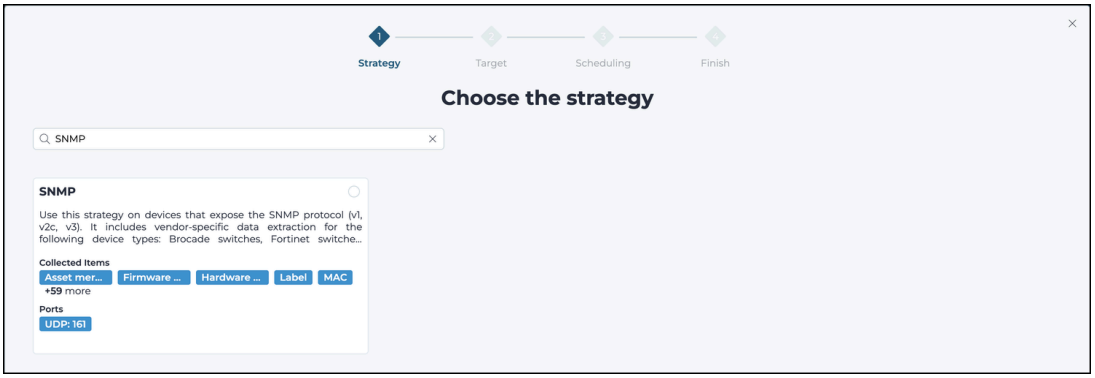
3. On the **Smart Polling** page, select the ⊕ **Add** icon.



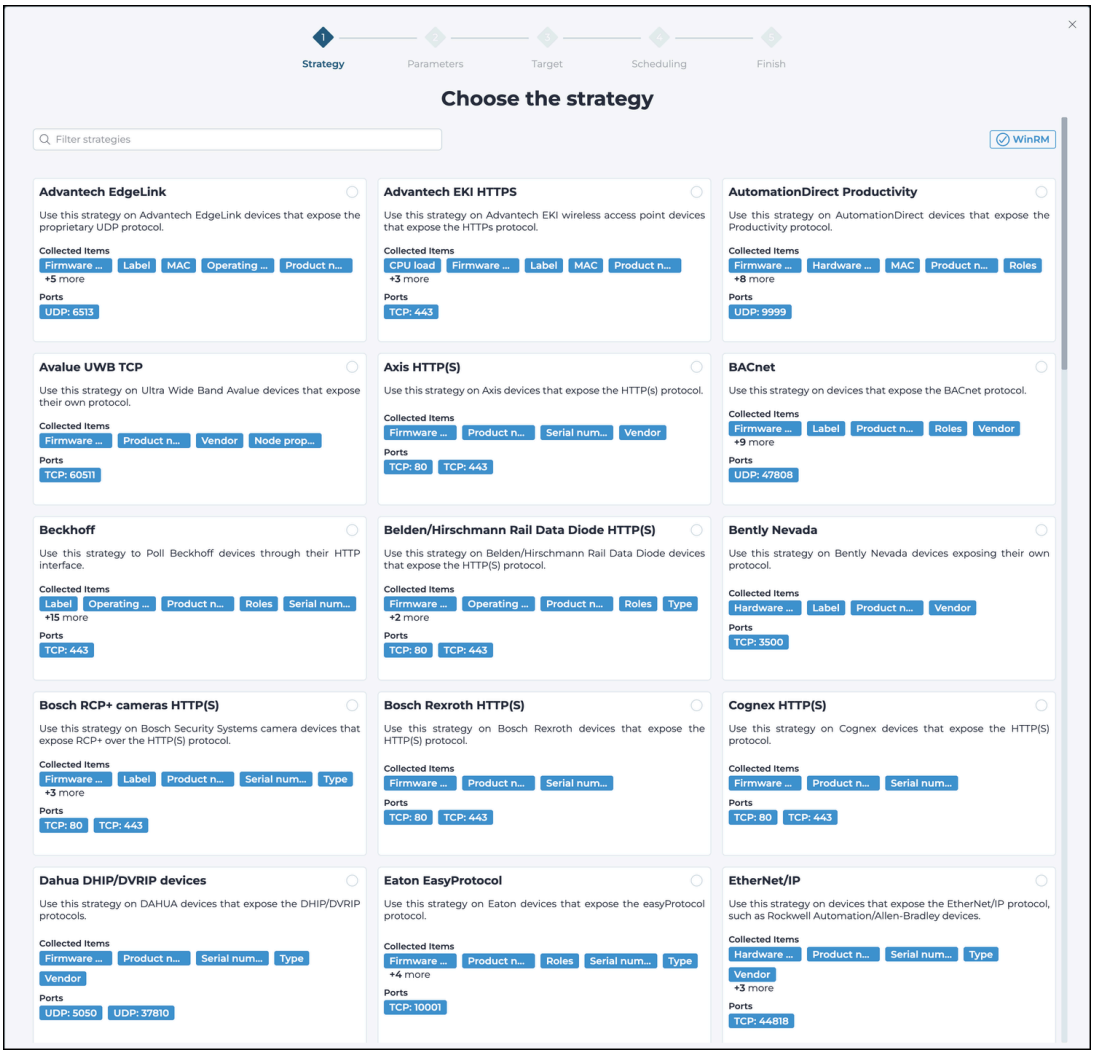
Result: The **Choose the Strategy** page opens.

4. **Optional:**

   To find the correct strategy, enter text in the **Filter strategies** field.



5. In the top right corner of the relevant strategy, select the radio button.



6. Select **Next**.

   Result: Either the **Parameters** page or the **Target** page will show.

7. **Optional:** If the **Parameters** page shows, enter the necessary details.
8. Select **Next**.

   Result: The **Target** page will show.
9. From the **Available credentials** panel on the left, select the applicable credential.
10. Choose an option:

    Choose from:

    - To select one credential, select the > icon.
    - To select all the available credentials, select the >> icon.
11. Select **Next**.

    Result: The **Scheduling** page will show.
12. On the **Scheduling** page, under **Recurrence**, select the granularity of the scheduling:

    Choose from:

    - **Hourly**
    - **Daily**
    - **Weekly**
    - **Monthly**
13. Set the time at which the Smart Polling plan will be executed. The time is set in .
14. Select **Next**.

    Result: The **Finish** page will show.
15. **Optional:** In the **Plan name** field, edit the name of the plan.
16. Select **Confirm**.
17. Select **Done**.

Results

The Smart Polling plan has been added.

Add a Smart Polling plan that does not require credentials

Create a Smart Polling plan in Guardian for targets that do not require credentials. Define the query, configure the schedule, and complete the setup.

About this task

Use this process to create a Smart Polling plan for devices or systems that do not require authentication. This method lets Guardian collect data from non-credentialed targets using a query-based selection. You can configure the schedule of the plan to suit your operational needs.

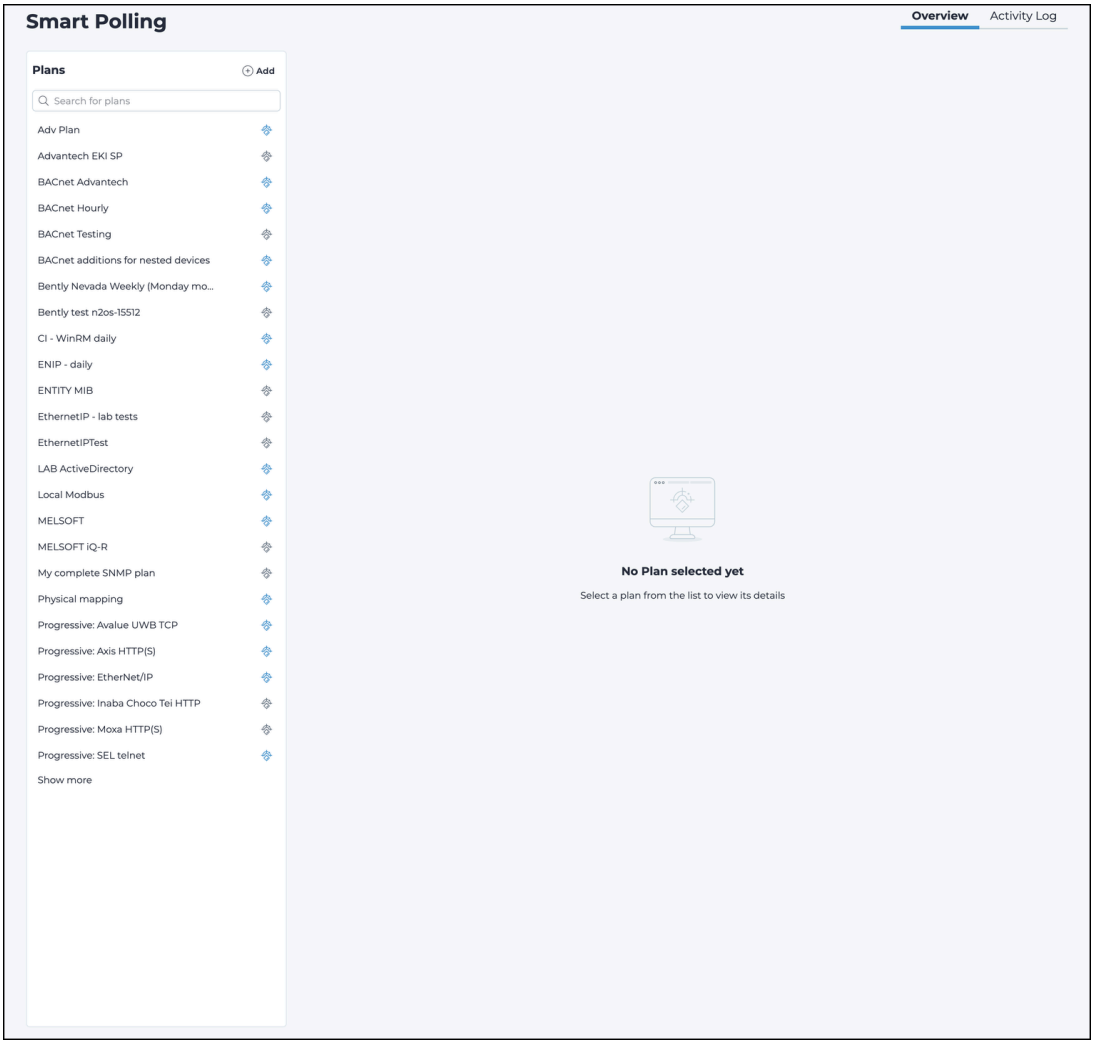
Procedure

1. In the top navigation bar, select ☰ > **Smart Polling**.

   Result: The **Smart Polling** page opens.

2. If the welcome page shows, select **Go to the Smart Polling page**.

3. On the **Smart Polling** page, select the ⊕ **Add** icon.



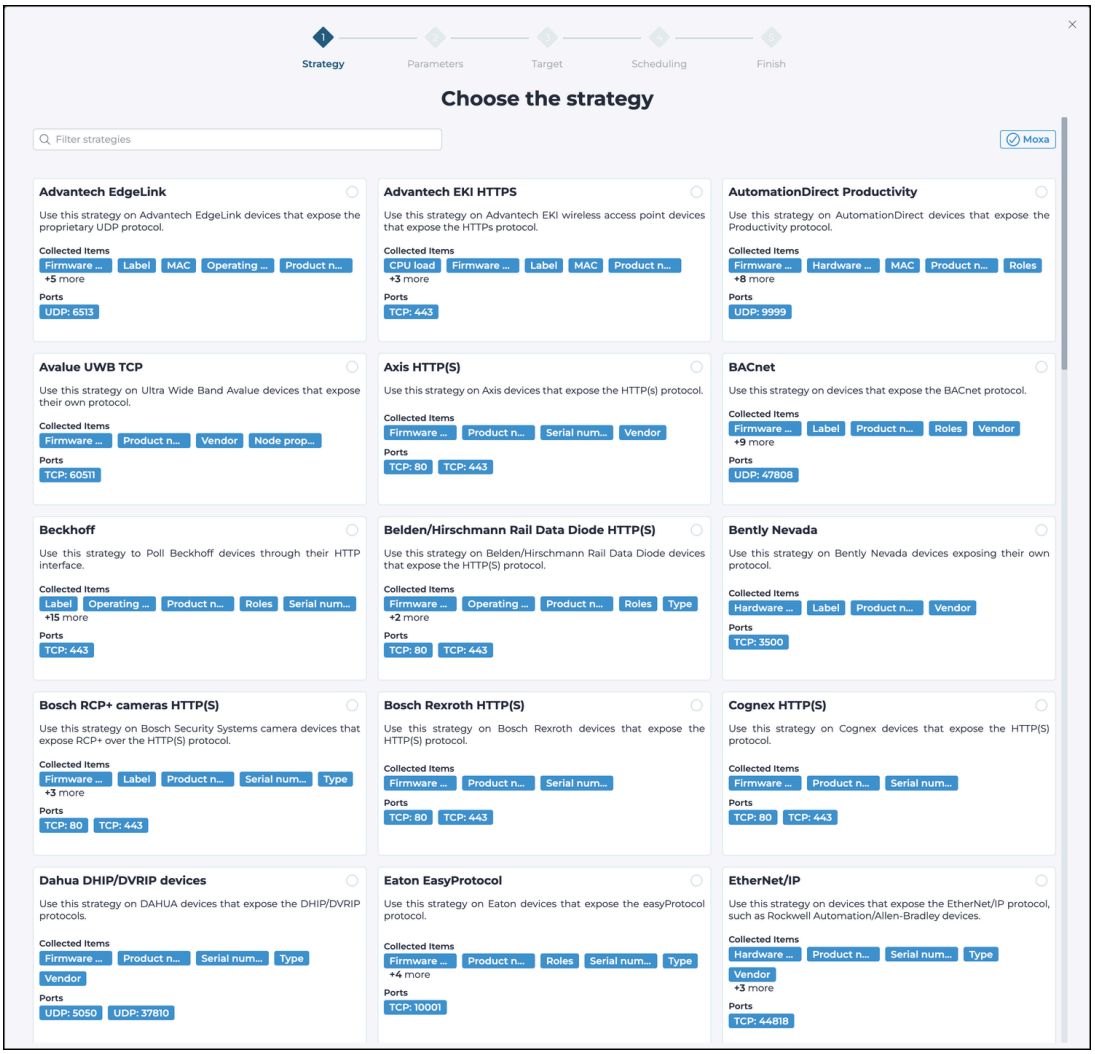
Result: The **Choose the Strategy** page opens.

4. **Optional:**

   To find the correct strategy, enter text in the **Filter strategies** field.



5. In the top right corner of the relevant strategy, select the radio button.



6. Select **Next**.

   Result: Either the **Parameters** page or the **Target** page will show.

7. **Optional:** If the **Parameters** page shows, enter the necessary details.

> **Note:**
> Smart Polling proposes a safe default for each of these parameters. Unless you have specific requirements, we recommend that you use these parameters.

8. Select **Next**.

   Result: The **Target** page will show.

9. In the **Target nodes with a query** section, enter a query.
   For example, `nodes | where ip in_subnet? "192.168.45.0/24"`



10. Select **Test query**.
    If successful a `Success` indicator and the amount of nodes found will show and the **Next** button will be enabled.
11. Select **Next**.

    Result: The **Scheduling** page will show.
12. On the **Scheduling** page, under **Recurrence**, select the granularity of the scheduling:

    Choose from:
    - **Hourly**
    - **Daily**
    - **Weekly**
    - **Monthly**
13. Set the time at which the Smart Polling plan will be executed. The time is set in .
14. Select **Next**.

    Result: The **Finish** page will show.
15. **Optional:** In the **Plan name** field, edit the name of the plan.
16. Select **Confirm**.
17. Select **Done**.

Results

The Smart Polling plan has been added.

Upload a MIB file

Upload a management information base file to make its objects available when you create a plan that uses the Simple Network Management Protocol.

About this task

When you create an *SNMP* plan, you choose the *OIDs* to poll from the *MIB* files that are available. If the *OIDs* that you need are not in an available *MIB* file, you can upload your own.

> **Note:**
> If the file depends on other *MIB* files, you must upload the other files first. The upload fails if a newer version of the file is already stored.

Each file can be up to 10 MB. The file can have any extension, as long as its content is formatted as a *MIB* file.

Procedure

1. In the plan creation wizard for an *SNMP* plan, go to the **MIB selection** page.
2. Select **Upload MIB**.

   Result: The **Upload MIB files** drawer shows.
3. Add the *MIB* files to upload.

   Select the upload area to open a file selector, or drag and drop the files onto it.

   You can add more than one file at a time.

   Result: The files are uploaded, and a message shows how many *OIDs* were imported.

Results

The *OIDs* from the uploaded files are available on the **OID selection** page.

The Guardian stores the uploaded files. When you create or edit another *SNMP* plan, you do not have to upload them again.

Map an OID to a node attribute

Map the value of an object identifier to a node attribute, so that Smart Polling uses the value it collects to update that attribute on the node.

About this task

When you create an *SNMP* plan, you can map the value of an *OID* to a node attribute. You can map an *OID* to one or more of the following attributes:

- Firmware version
- Label
- Product name
- Serial number
- Vendor
- Any custom field with 'string' type

Procedure

1. In the plan creation wizard for an *SNMP* plan, go to the **OID selection** page.
2. Find the *OID* in the list.

   Use the arrows below the list to go to other pages, or enter text in the **Search by OID** field. You can search by the number of the *OID*.
3. Select the *OID*.

   Result: The details of the *OID* show on the right.
4. Under **Attribute mappings**, select **Add attribute mapping**.

   Result: The **Mapping #** box shows.
5. From the **Map this OID to** list, select the node attribute.

   Custom fields show in the list as **Custom field:** followed by the name of the field.

   If you added a mapping, but did not select a target attribute, the button to go to the next step is disabled.
6. **Optional:** You can map a single *OID* to multiple attributes. To do so, select **Add attribute mapping** again, and repeat the previous step.

Results

The number of attribute mappings in the **Configured** summary at the top of the page increases, and the dot for the mapping is filled.

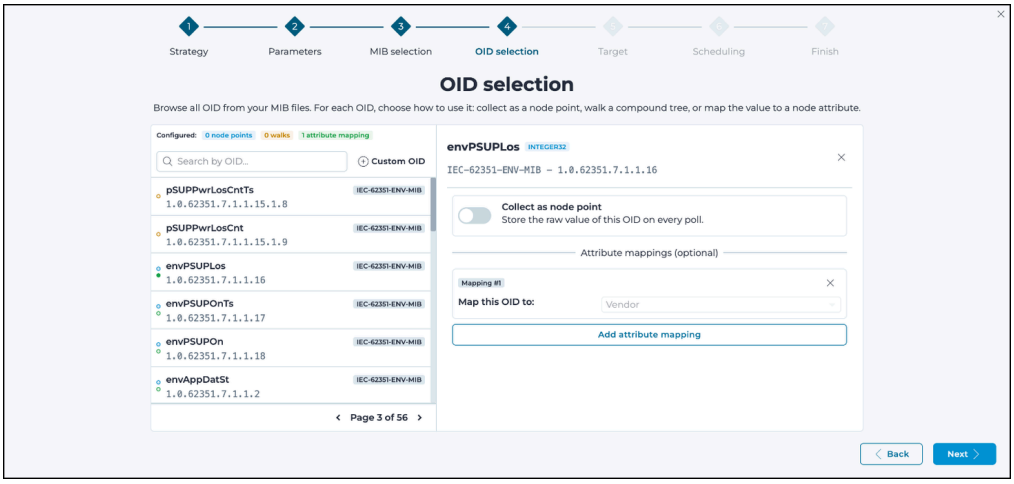


Figure 48. Attribute mapping on the OID selection page

What to do next

To remove a mapping, select the cross symbol on the mapping.

Add a custom OID to the list

Add a custom object identifier that is not in any available management information base file, so that you can select it when you create a plan.

About this task

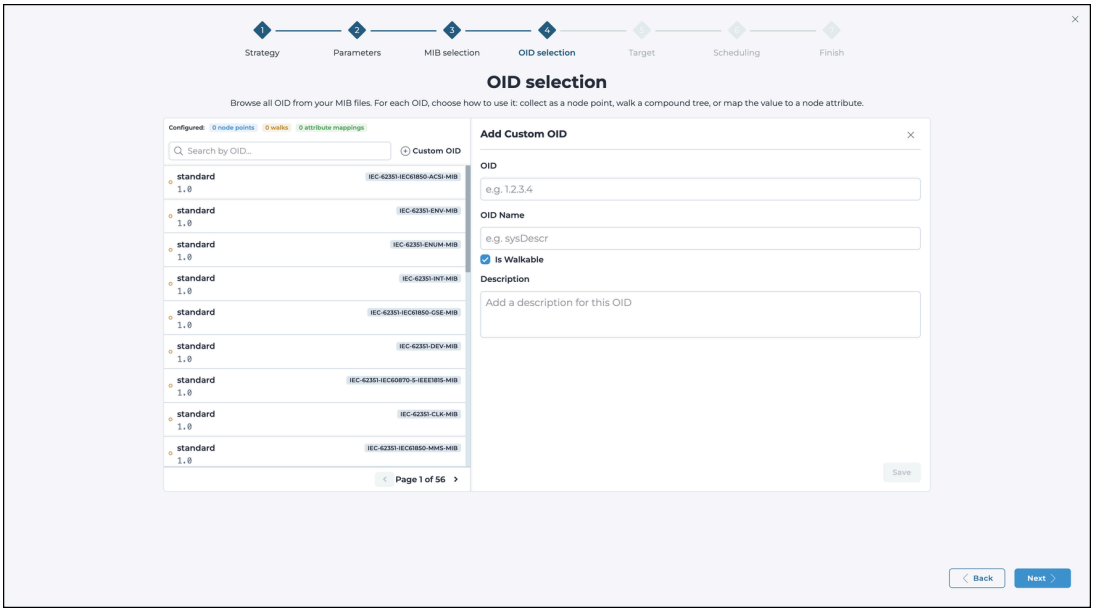
If an *OID* is not in any available *MIB* file, you can define it yourself as a custom *OID*.

Procedure

1. In the plan creation wizard for an *SNMP* plan, go to the **OID selection** page.
2. Select **Custom OID**.

   Result: The **Add Custom OID** form shows.

3. In the **OID** field, enter the *OID* as numbers separated by dots, for example `1.3.6.1.2.1.47.1.1.1`.



4. In the **OID Name** field, enter a name for the *OID*, for example `entPhysicalTable`.
5. **Optional:** Select **Is Walkable** to switch it off if the *OID* holds a single value. This option is on by default. A walkable *OID* represents a *MIB* node or a table. When it is walked, every value that is found is stored as a separate node point.
6. **Optional:** If **Is Walkable** is off, from the **Base Type** list, select the data type that the *SNMP* server returns for the *OID*: string or integer.
7. **Optional:** Enter a description for the *OID*.
8. Select **Save**.

   Result: The custom *OID* is added.

What to do next

When you select a custom *OID* in the list, the *OID* editor opens. The **Edit** and **Delete** icons show on the top right of the editor.

To change a custom *OID*, select it in the list, and then select the **Edit** icon.

To delete a custom *OID*, select it in the list, select the **Delete** icon, and then confirm in the **Delete Custom OID** dialog. You cannot undo this action.

Actions menu

Pause a Smart Polling plan

Temporarily suspend a Smart Polling plan without deleting it. Use this option to stop polling activity while preserving the plan configuration.

About this task

When you pause a Smart Polling plan, it temporarily disables its polling activity without removing its configuration. This is useful when you want to stop monitoring for maintenance, troubleshooting, or scheduling reasons, but plan to resume later.

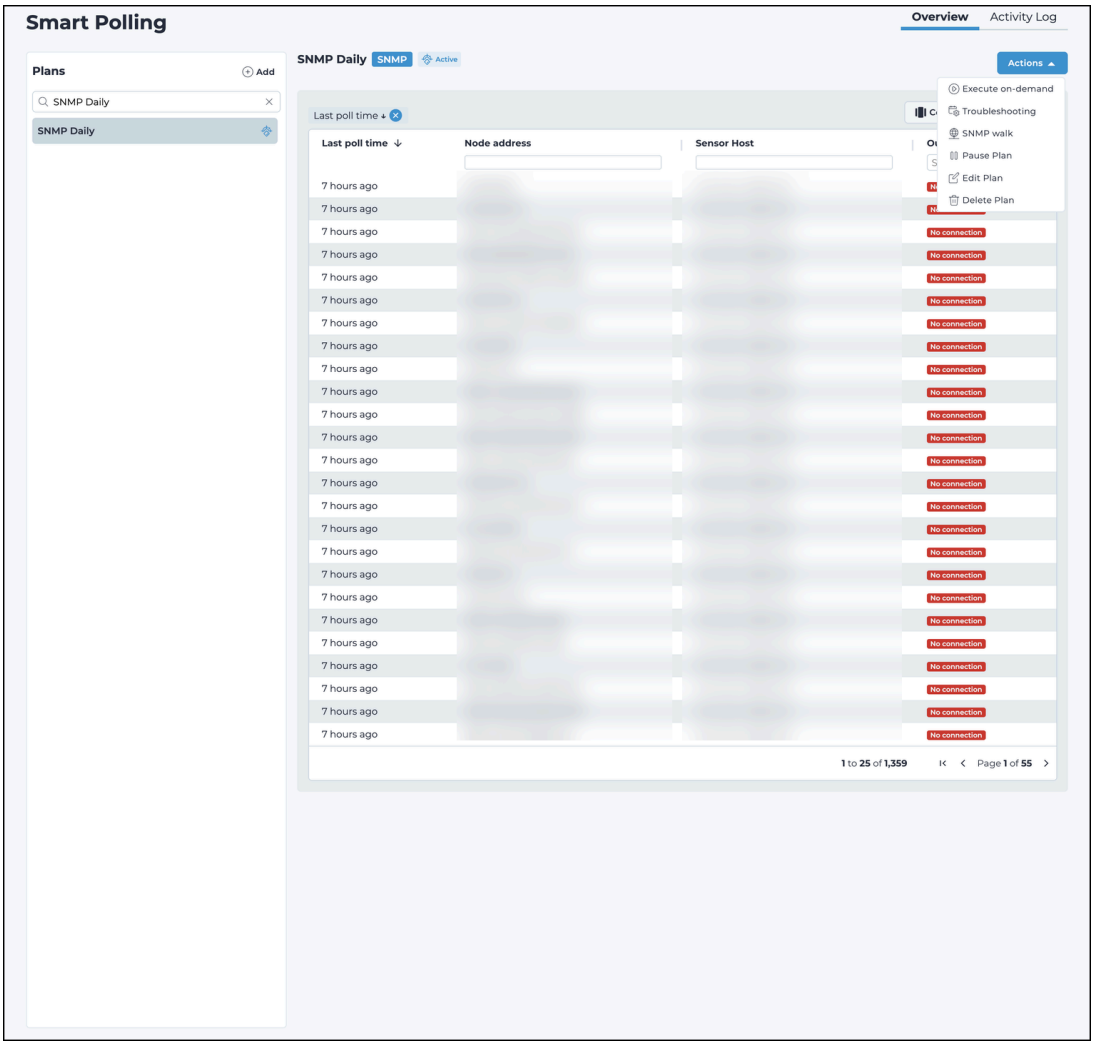
Procedure

1. In the top navigation bar, select ☰ > **Smart Polling**.

   Result: The **Smart Polling** page opens.

2. From the list of Smart Polling plans on the left, select the applicable plan.

3. In the top right corner, select **Actions**.



4. Select **Pause Plan**.

Results

The Smart Polling plan has been paused.

Resume a Smart Polling plan

Reactivate a previously paused Smart Polling plan to resume automated monitoring. This action restarts polling using the plan's saved configuration.

About this task

When you resume a Smart Polling plan, it restarts monitoring activity that was previously paused. This is useful after scheduled maintenance, issue resolution, or other intentional interruptions in polling. The plan resumes with its original configuration intact.

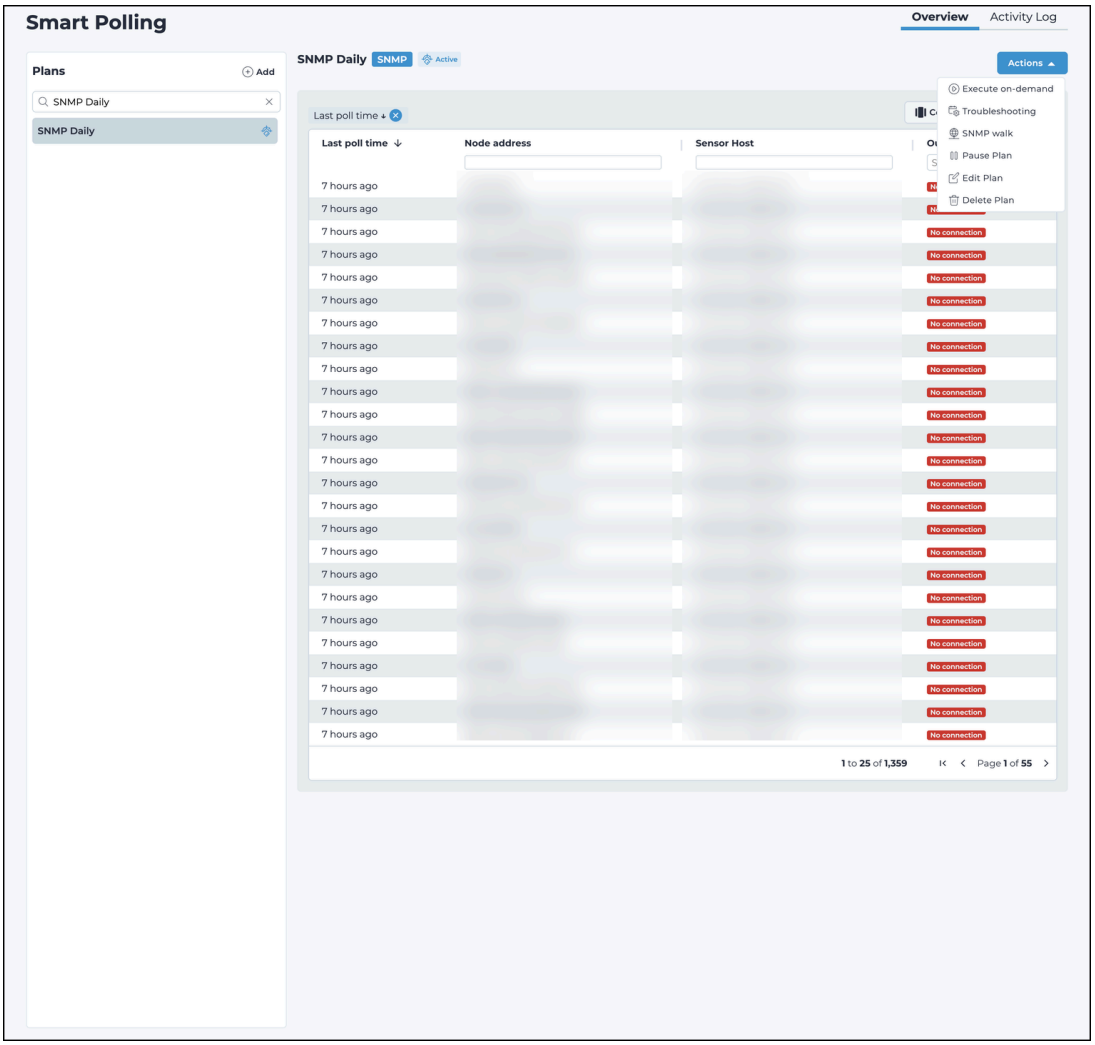
Procedure

1. In the top navigation bar, select ☰ > **Smart Polling**.

   Result: The **Smart Polling** page opens.

2. From the list of Smart Polling plans on the left, select the applicable plan.

3. In the top right corner, select **Actions**.



4. Select **Resume Plan**.

Results

The Smart Polling plan has been resumed.

Edit a Smart Polling plan

Modify an existing Smart Polling plan to adjust its configuration or update associated settings. Editing a plan lets you fine-tune monitoring behavior.

About this task

Editing a Smart Polling plan lets you change its targets, credentials, polling intervals, or other configurations. This is useful when your monitoring requirements change or you need to correct or refine existing settings.

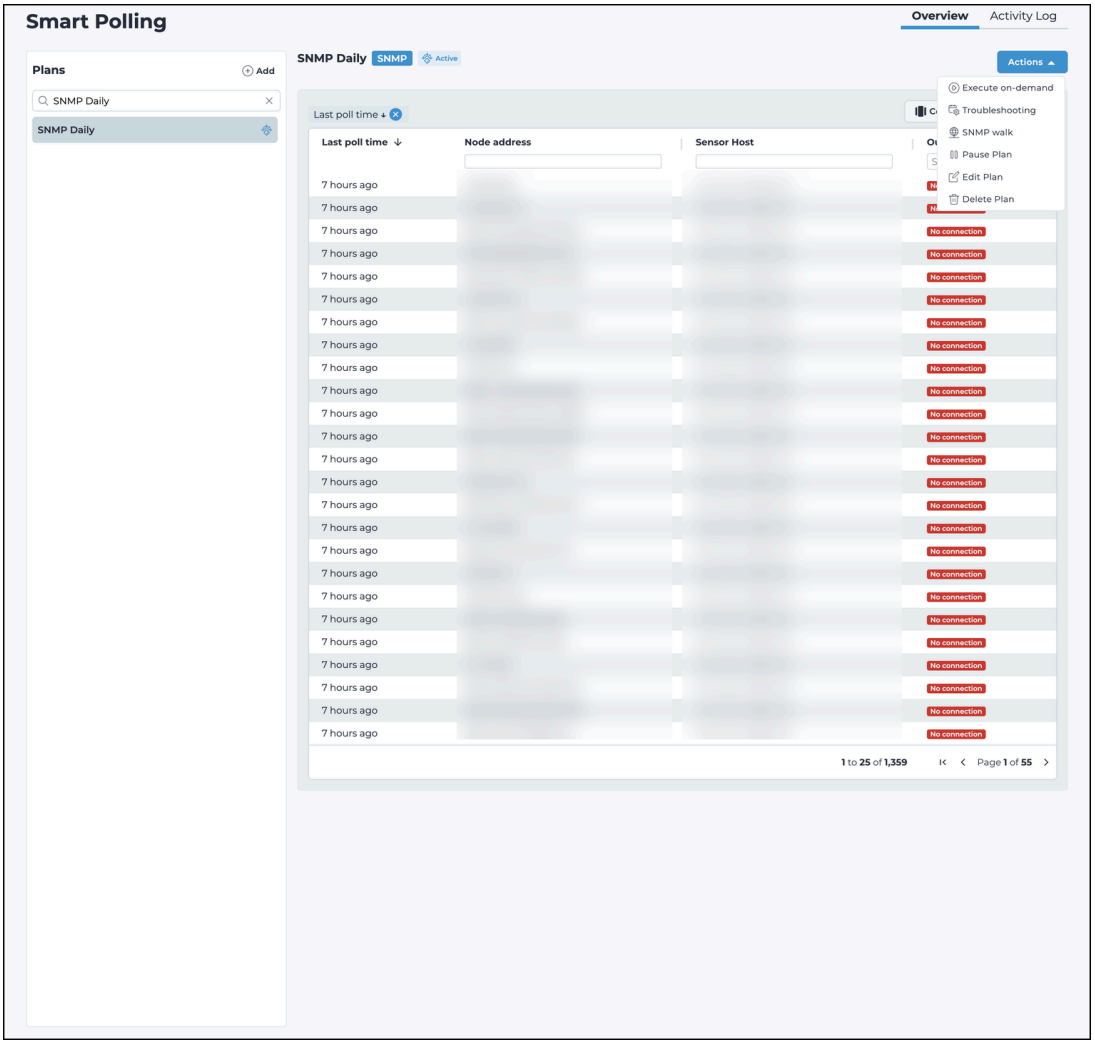
Procedure

1. In the top navigation bar, select ☰ **> Smart Polling**.

   Result: The **Smart Polling** page opens.

2. From the list of Smart Polling plans on the left, select the applicable plan.

3. In the top right corner, select **Actions**.



4. Select **Edit Plan**.

   Result: The plan editing wizard opens in a modal window.

5. Use the applicable procedure to edit the plan:

   Choose from:

   - Add a Smart Polling plan that requires credentials (on page 128)
   - Add a Smart Polling plan that does not require credentials (on page 132)

   You can then confirm your edited details on the **Scheduling** and **Finish** pages.

Delete a Smart Polling plan

Delete an existing Smart Polling plan to remove obsolete or unneeded monitoring configurations. This action is permanent and cannot be undone.

About this task

Deleting a Smart Polling plan permanently removes its configuration and stops all associated monitoring activity. Use this option when a plan is no longer relevant or needed. Be sure to confirm the deletion, as this action cannot be undone.

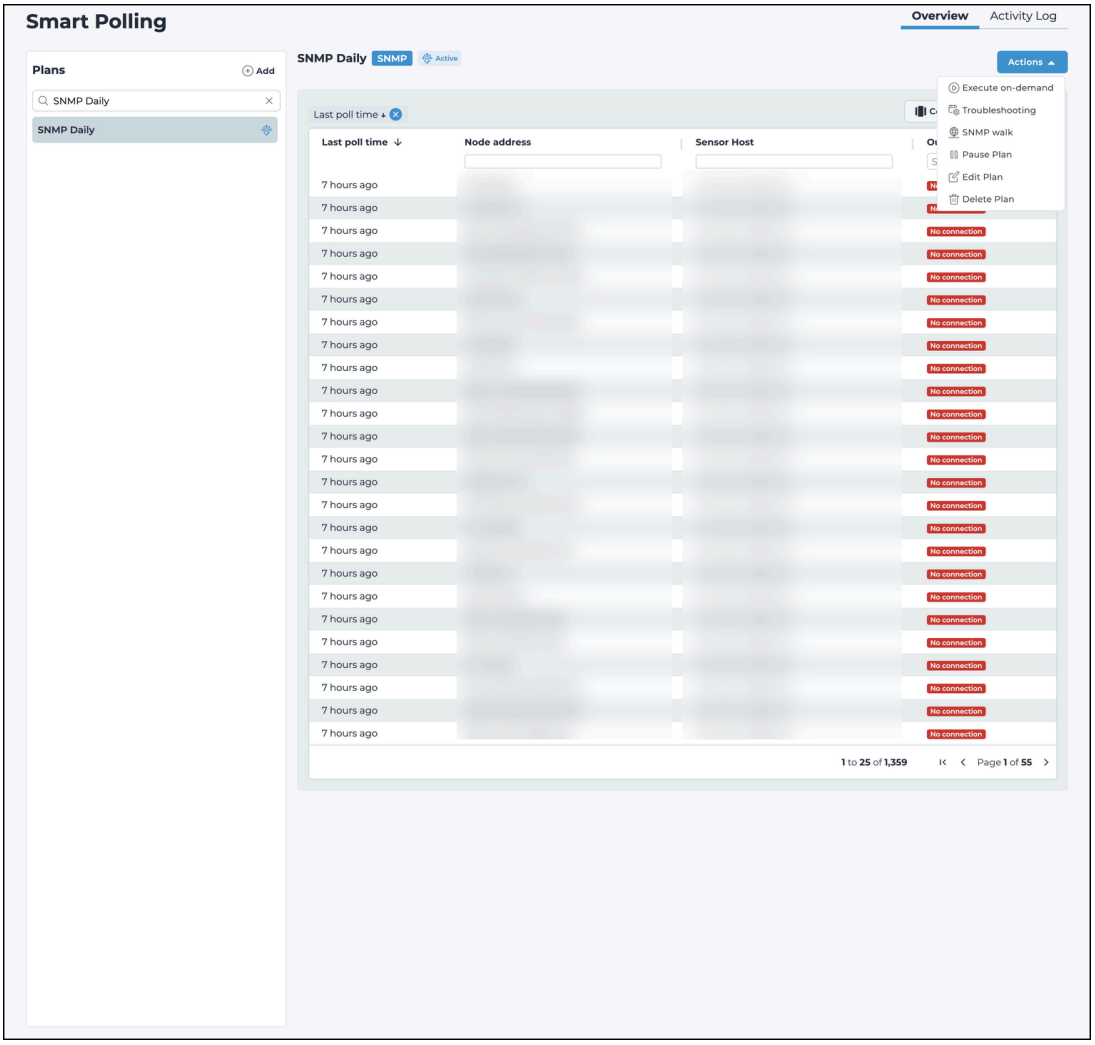
Procedure

1. In the top navigation bar, select ☰ **> Smart Polling**.

   Result: The **Smart Polling** page opens.

2. From the list of Smart Polling plans on the left, select the applicable plan.

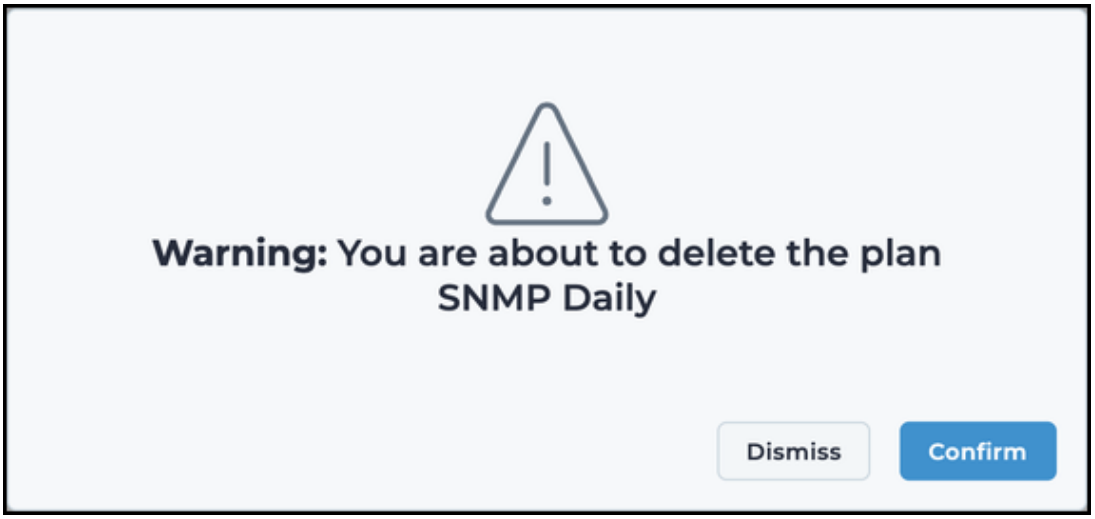
3. In the top right corner, select **Actions**.



4. Select **Delete Plan**.

   Result: A dialog shows

5. Select **Confirm**.



Results

The Smart Polling plan has been deleted.

Execute a Smart Polling plan on demand

Run the selected Smart Polling plan once, independently of its schedule. The execution is performed immediately and does not interfere with the plan schedule.

About this task

Using this action, you can perform an execution of the selected plan that runs independently of the chosen schedule. The execution is performed immediately, and it does not interfere with the schedule. For instance, if a plan is scheduled to be executed every 24 hours, and the on-demand execution action is run 8 hours after the scheduled execution, the next execution will still run 24 hours after the previous scheduled execution. You can run a plan on demand even when the plan is paused.

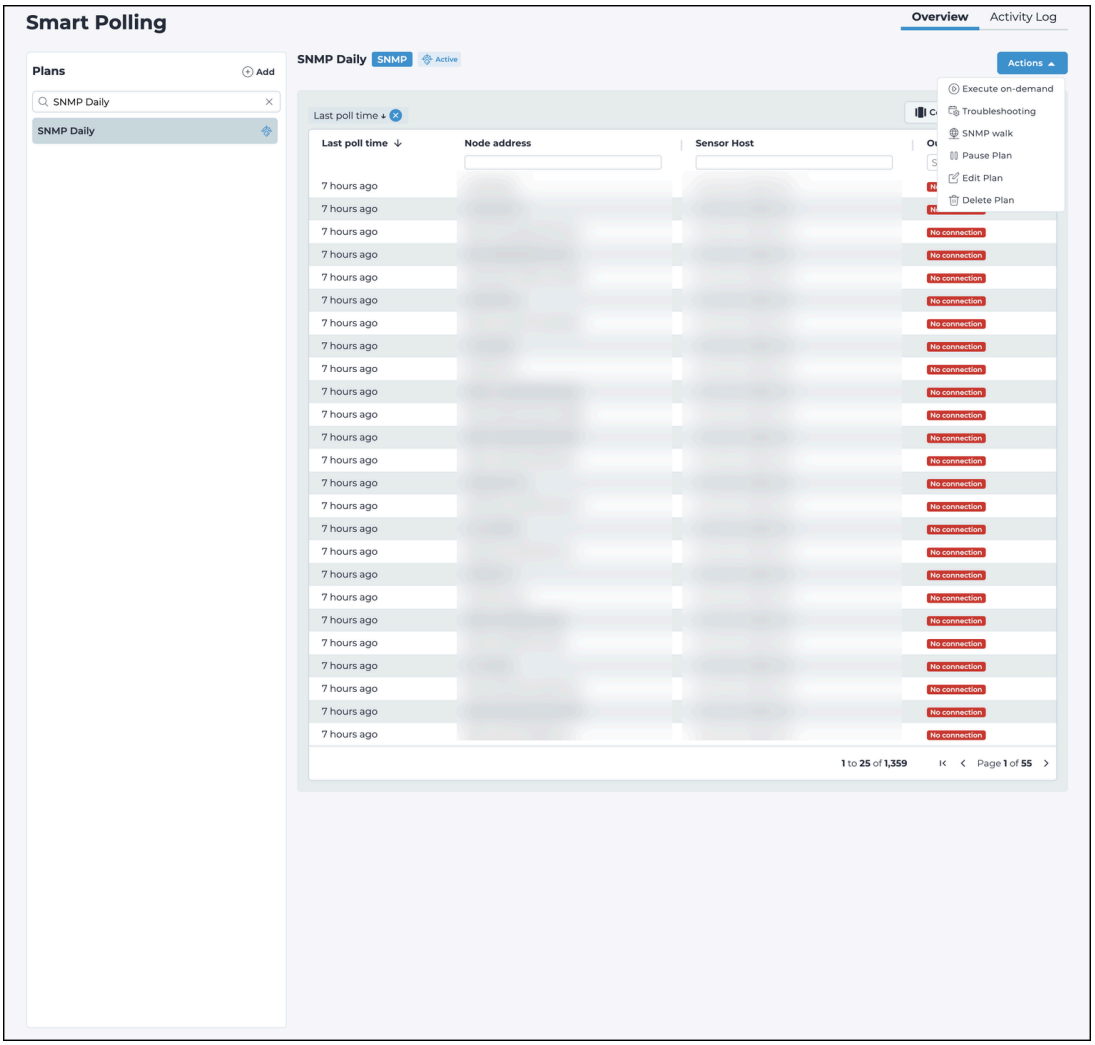
Procedure

1. In the top navigation bar, select ☰ > **Smart Polling**.

   Result: The **Smart Polling** page opens.

2. From the list of Smart Polling plans on the left, select the applicable plan.

3. In the top right corner, select **Actions**.



4. Select **Execute on-demand**.

   Result: A dialog shows.

5. Select **Confirm**.

Results

The Smart Polling plan runs once, immediately.

Troubleshooting

Use the Smart Polling troubleshooting tools to understand and preview the behavior of a plan, so you can maximize its success.

Smart Polling provides two troubleshooting tools. Select **Troubleshooting** in the **Actions** menu to open a modal window where the tools show:

- Run a dry run (on page 151)
- Preview a plan schedule (on page 154)

Run a dry run

Perform a single polling of a node without recording data. Use this tool to test a plan against a specific IP address without changing the node data or the activity log.

About this task

In this tool, you can choose an *IP* address to poll with the selected plan. The result of the polling will not be recorded in the activity log, and the node data will not be updated.

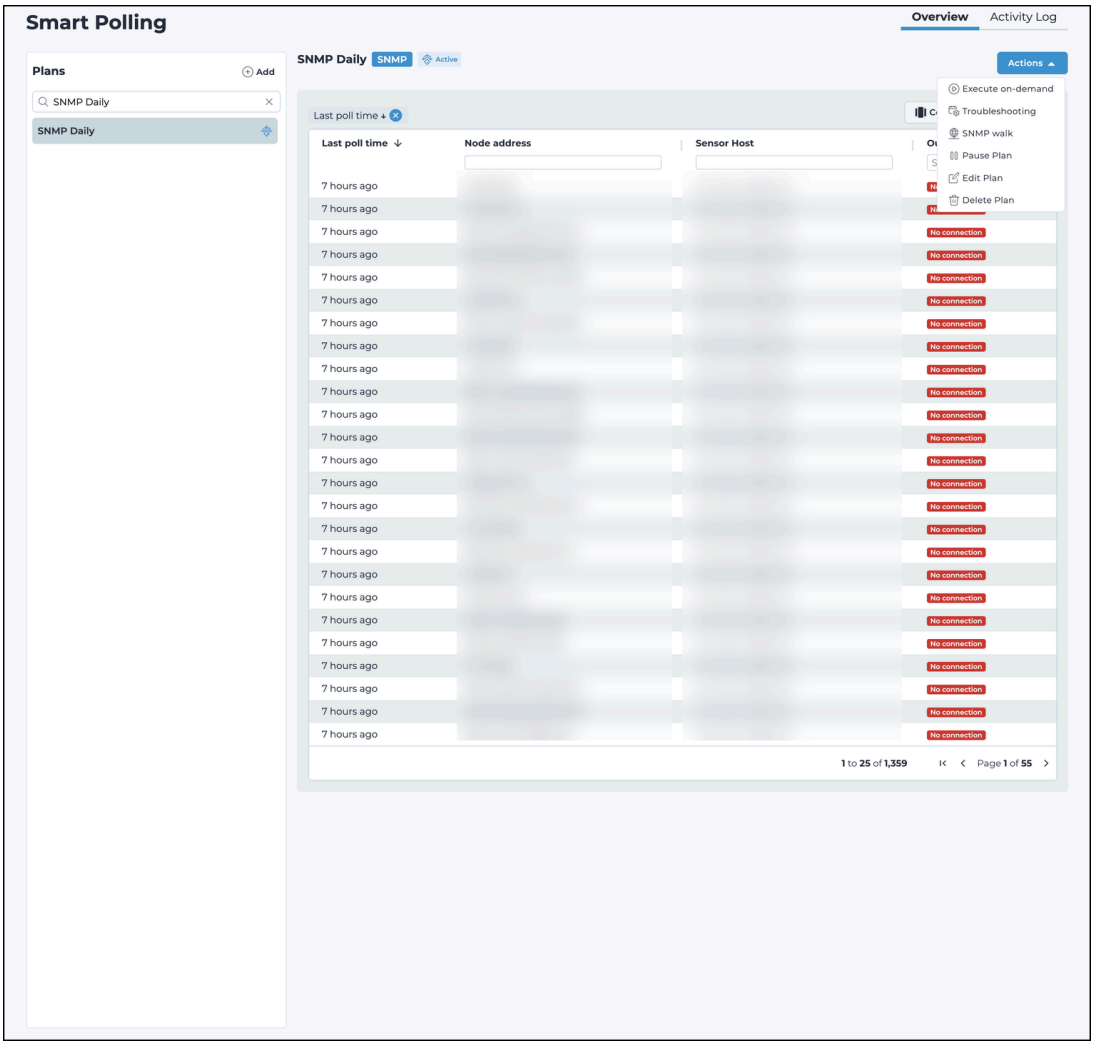
Procedure

1. In the top navigation bar, select ☰ > **Smart Polling**.

   Result: The **Smart Polling** page opens.

2. From the list of Smart Polling plans on the left, select the applicable plan.

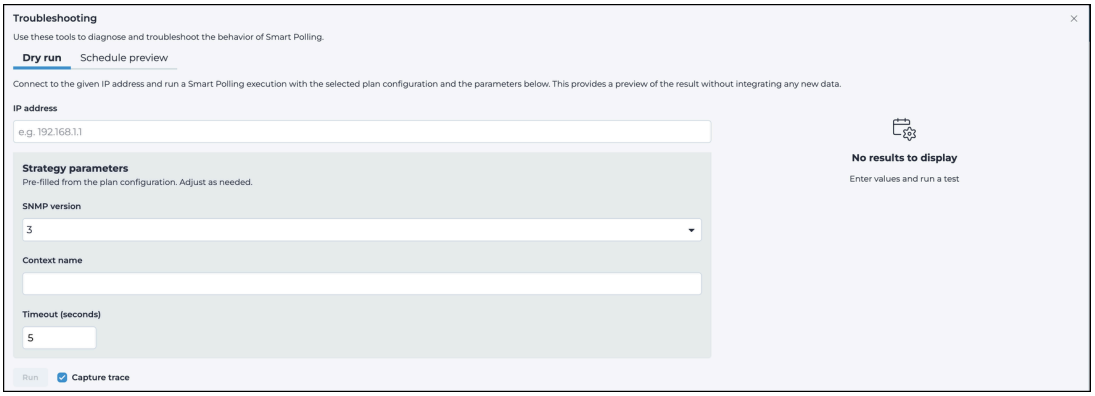
3. In the top right corner, select **Actions**.



4. Select **Troubleshooting**.

   Result: A dialog shows.

5. Enter the *IP* address of the device to poll in the **IP address** input.



6. **Optional:** Change the parameters used in this execution.
   The changes to these parameters will not be saved in the plan configuration.
7. **Optional:** Select the **Capture trace** box to download a *packet capture (pcap)* file for this poll.
8. Select **Run**.

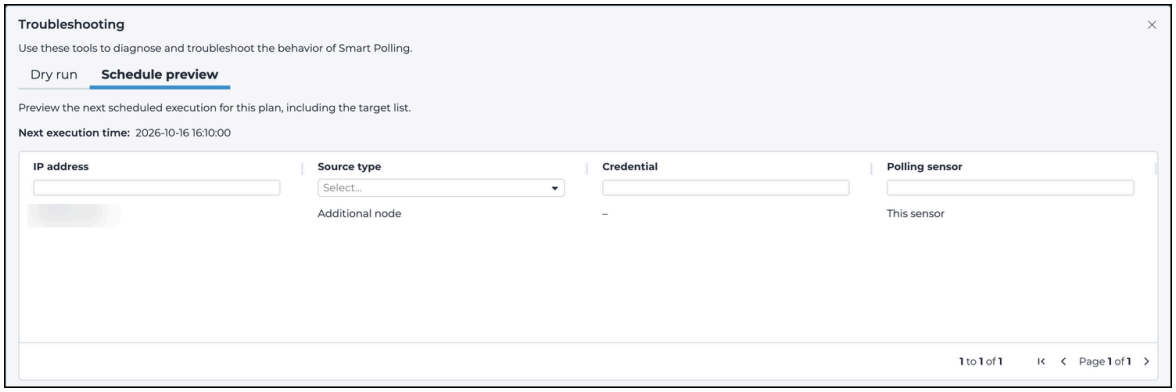
   Result: The node is polled and the results pane updates.
9. On the right side, inspect the results.
10. **Optional:** Select **Download trace** to get the capture file.

Preview a plan schedule

Get a preview of the next execution of a Smart Polling plan. The preview shows the time of the next scheduled execution, if any, and the list of nodes that will be polled.

About this task

This tool does not perform any action. It only gives diagnostic information. The page shows the time of the next scheduled execution, if any, and the list of nodes that will be polled.



Column label	Description
IP address	The address of the polled device
Source type	The reason why this node is selected
Credential	The name of the credential used, if any
Polling sensor	The hostname of the sensor that will do the polling

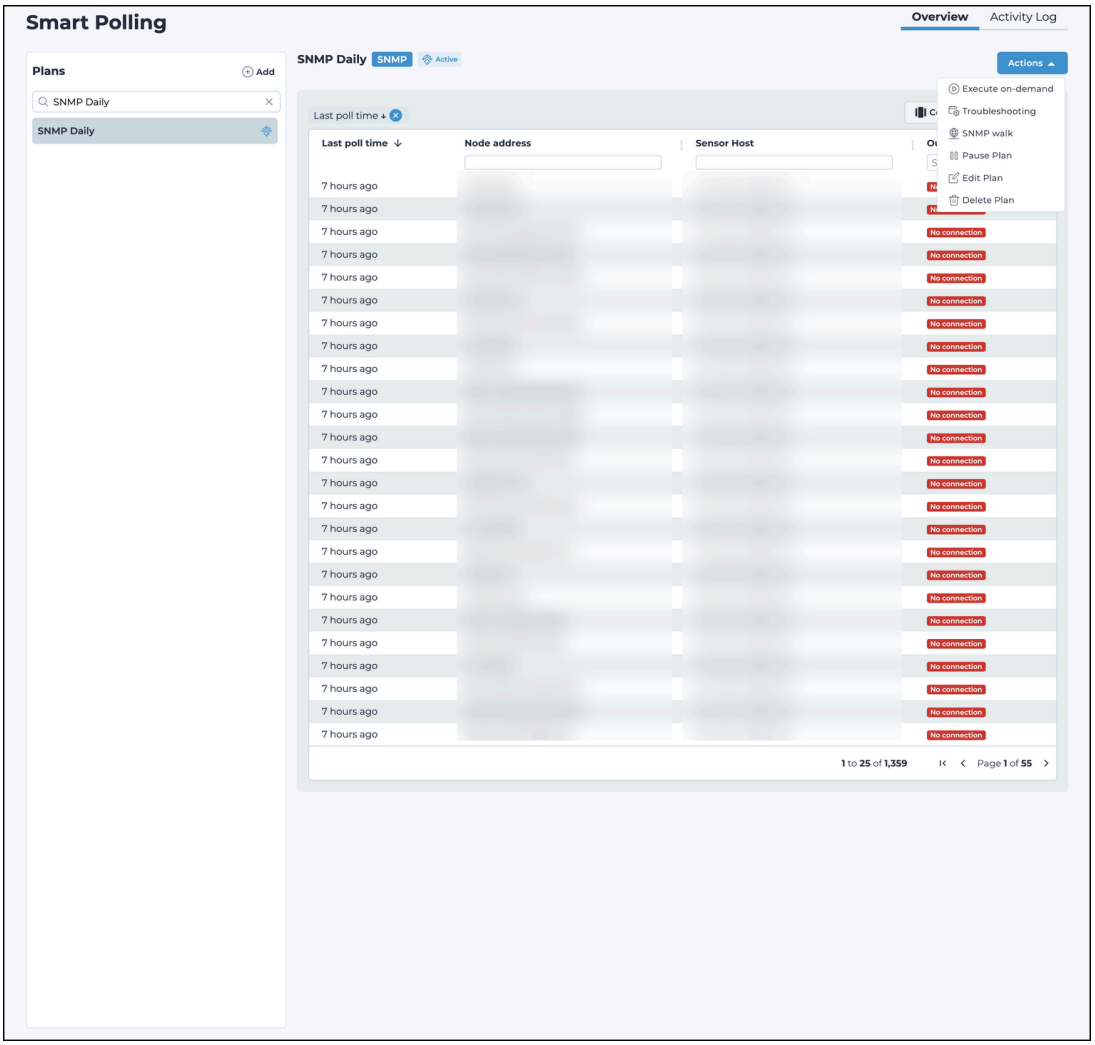
Procedure

1. In the top navigation bar, select ☰ > **Smart Polling**.

   Result: The **Smart Polling** page opens.

2. From the list of Smart Polling plans on the left, select the applicable plan.

3. In the top right corner, select **Actions**.



4. Select **Troubleshooting**.

   Result: A dialog shows.

5. From the tab bar, choose **Schedule preview**.

   Result: The tool shows the next execution details and the list of nodes.

Run an SNMP walk

Run a full or partial SNMP walk on a monitored device to get detailed information about it, and download the result in textual format.

About this task

For *SNMP* plans, this action lets you perform a full or partial walk on the monitored device. You can download the result in textual format.

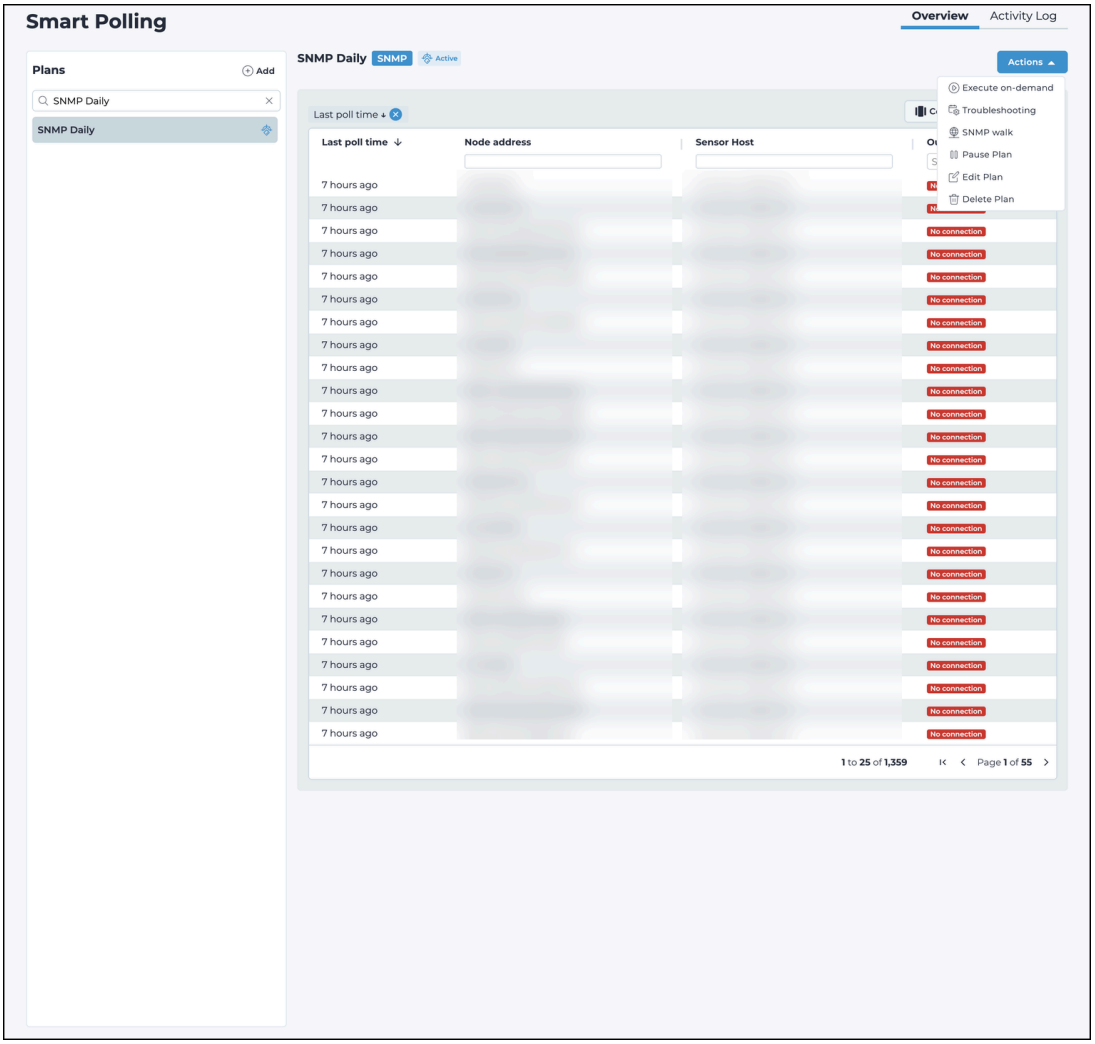
Procedure

1. In the top navigation bar, select ☰ > **Smart Polling**.

   Result: The **Smart Polling** page opens.

2. From the list of Smart Polling plans on the left, select the applicable *SNMP* plan.

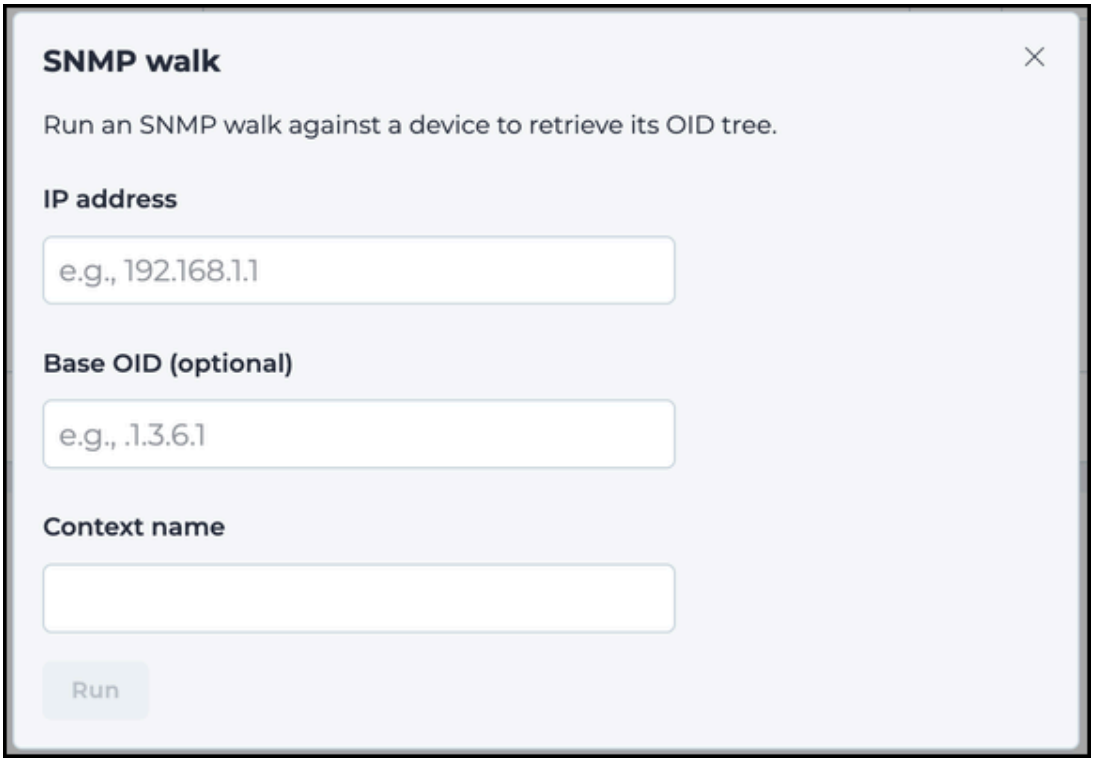
3. In the top right corner, select **Actions**.



4. Select *SNMP* **walk**.

   Result: A dialog shows.

5. Enter the *IP* address of the applicable node.



6. **Optional:** Enter a base *OID*.
   By default, the root *OID* is used.
7. **Optional:** Change the community or context string.
   By default, the setting of the plan is used.
8. Select **Run**.

   Result: The node is polled.

9. When the results are available, select **Download** to get the textual file containing the results.

Add a node from the Network page

Once you have created a Smart Polling plan, you can add arbitrary nodes to the target from the Network page. The Smart Polling plan will poll these nodes even if the plan's query does not return them.

Procedure

1. In the top navigation bar, select icon > **Network**.

   Result: The **Network** page opens.

2. In the top right, select **Nodes**.

   Result: The **Nodes** page opens.

3. To the left of the name of the applicable plan, select the icon.

   Result: A dialog shows.

4. From the **Select an existing plan to add the node to** dropdown, select the plan you want to add the selected node to.

   > **Note:**
   > Fields that are not modified in this dialog are automatically populated with the plan-configured values.

   Result: A dialog shows.

5. Edit the field(s) as necessary to change the parameter(s) of the existing plan.
6. Select **Add**.

Results

The node(s) has (have) been added to the Smart Polling plan.

Log level customization

Smart Polling logs self-diagnostic information about its operations and activities during execution.

When Smart Polling logs self-diagnostic information, the logs are collected in these files:

- `/data/log/n2os/n2ossp.log`
- `/data/log/n2os/n2os_strategist.log`

To change the level of detail in the logs, you can run this command in the *CLI*:

```
conf.user configure sp log_level <LEVEL>
```

The change applies immediately. You don't need to edit any files or restart any services. where `<LEVEL>` is one of the following values (in increasing order of verbosity):

- `FATAL`
- `ERROR`
- `WARN`
- `INFO`
- `DEBUG`



Note:
The default value is `INFO`.



Note:
The values are ordered from the most to the least important, and the configured level is the minimum to be printed. For example:

- `ERROR` will print log lines for both `ERROR` and `FATAL` messages
- `FATAL` will print log lines only for `FATAL` messages

To configure Smart Polling to only produce `ERROR` and `FATAL` messages, use this *CLI* rule:

```
conf.user configure sp log_level ERROR
```

Smart Polling results

Smart Polling retrieves details from nodes to enrich asset data. The enriched information also appears in other places, such as the installed software, hotfixes, and vulnerabilities of an asset.

After a Smart Polling plan runs successfully, some or all of the information on the details page of an asset is enriched. To show the details of an asset, select it on the Assets page.

Overview

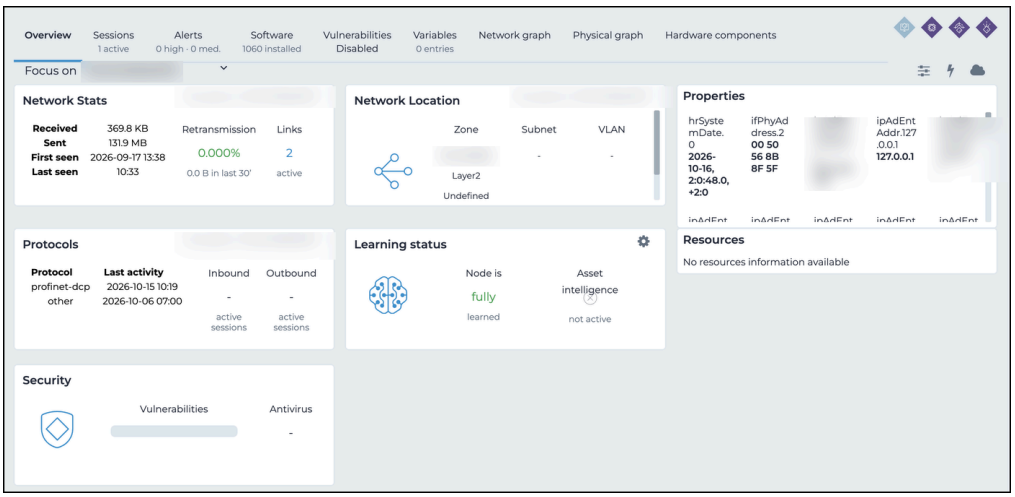


Figure 49. Assets page overview

On the **Overview** tab, Smart Polling enriches these widgets:

- **Resources**: Shows the CPU, RAM, and disk node points that Smart Polling collects from the node. The widget shows only when Smart Polling data is available for the asset
- **Security**: Shows the antivirus status that Smart Polling reports, next to the distribution of the vulnerabilities of the asset

Software tab

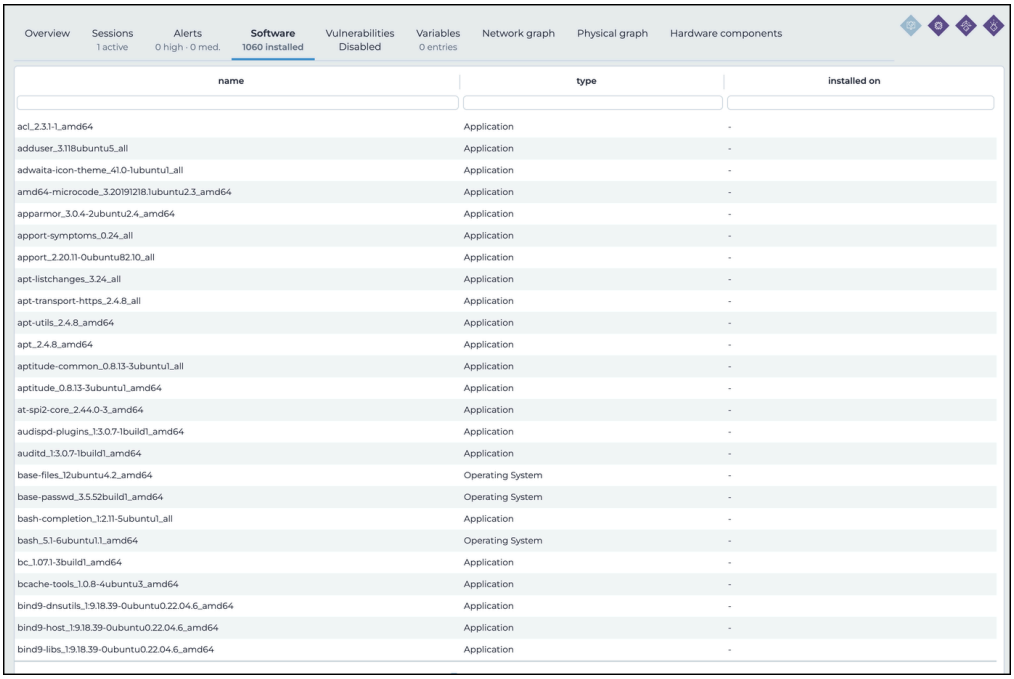


Figure 50. Software tab of an asset

The **Software** tab lists the software that is installed on the asset. Smart Polling fills it with the latest installed software that it retrieves from the node. This applies to the *secure shell (SSH)*, *Windows Remote Management (WinRM)*, and *Windows Management Instrumentation (WMI)* strategies.

Hotfixes tabs

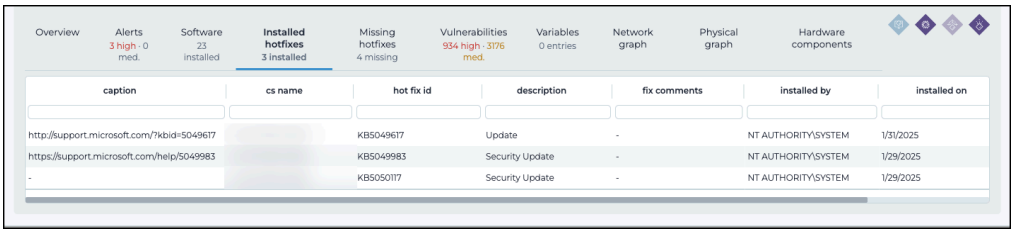


Figure 51. Installed hotfixes tab of an asset

The **Installed hotfixes** tab lists the hotfixes that are installed on the asset. Smart Polling fills it with the latest hotfixes that it retrieves from the node. This applies to the *WinRM* and *WMI* strategies. The tab shows for Windows assets, and for other assets only when they have hotfix data.

The **Missing hotfixes** tab lists the hotfixes that are missing on the asset. Smart Polling fills it with the hotfix data that it retrieves from the node.

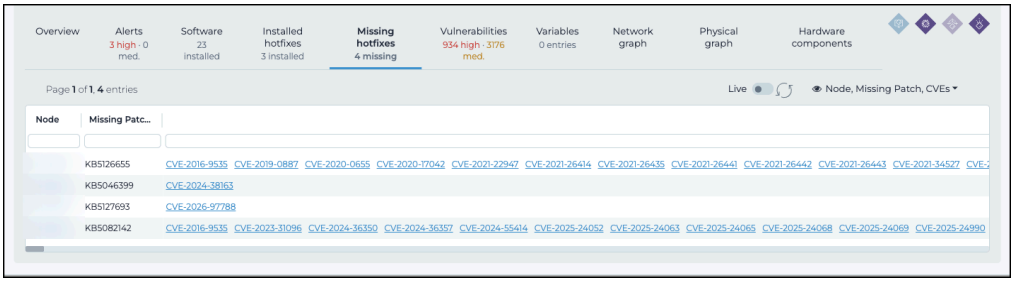


Figure 52. Missing hotfixes tab of an asset

Vulnerabilities tab

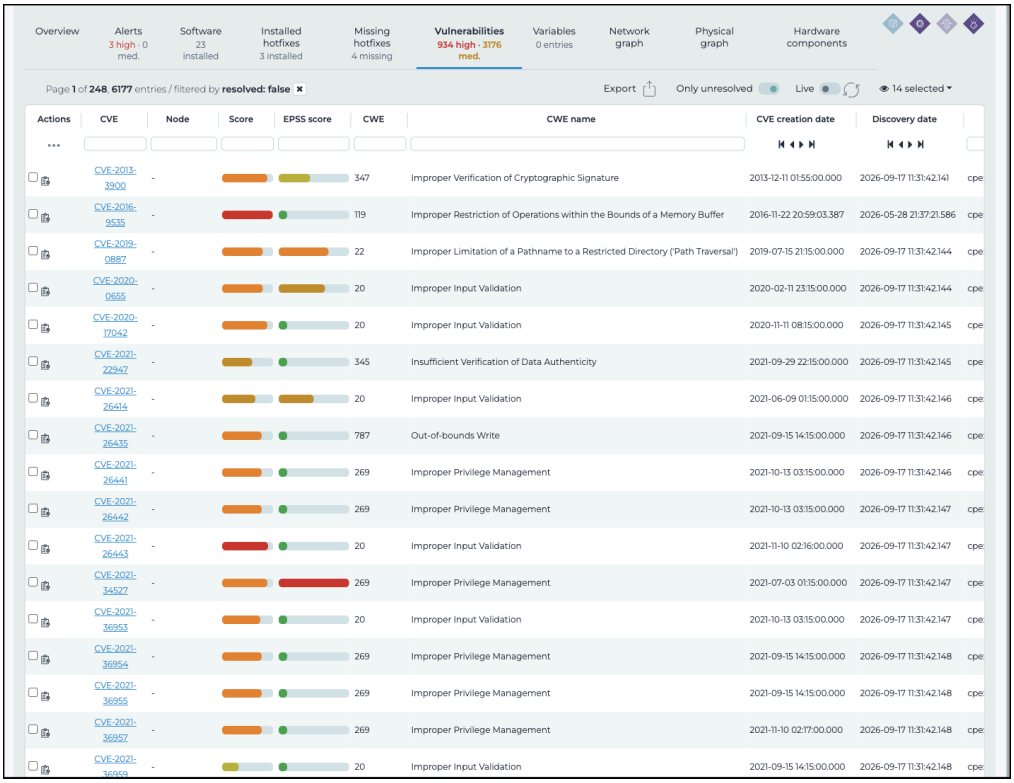


Figure 53. Assets page Vulnerabilities tab

The **Vulnerabilities** tab lists the vulnerabilities of the asset. The **Matching CPEs** column shows the *Common Platform Enumeration (CPE)*s that match each vulnerability. Smart Polling provides the information that Guardian uses to identify them.

Physical graph

In addition to the tabs above, Smart Polling provides the necessary information to populate the physical graph, which shows on the **Physical graph** tab of the asset. To do this, you need to use the *SNMP* strategy to poll network switches or routers.

Activity Log

*Use the **Activity Log** page in Guardian to inspect the history of Smart Polling activities. Each record in the table represents an execution of a Smart Polling plan on a specific node.*

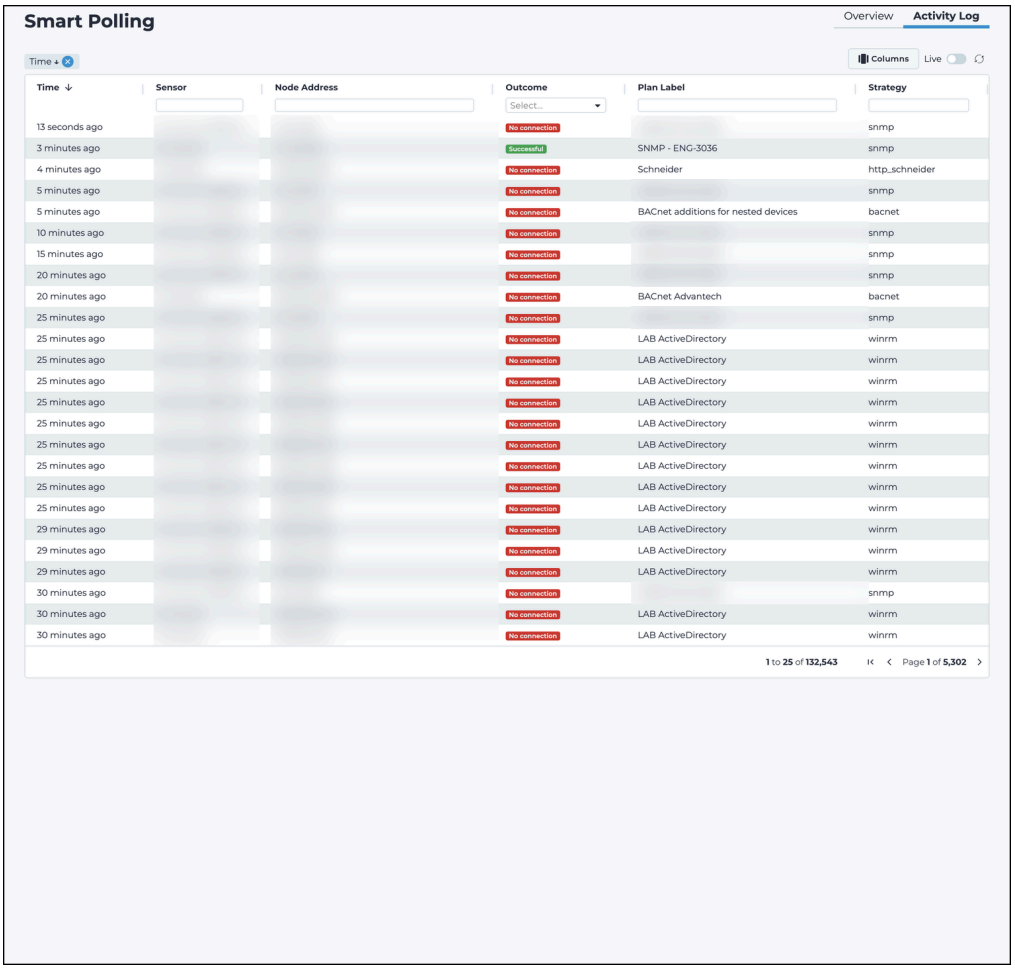


Figure 54. Activity Log page

Node executions

The **Activity Log** page shows the node executions. These are the individual executions of Smart Polling on a specific node with a configured plan. The table gives a view that you can filter and sort.

You can double-click a row in the table to open a sidebar with additional information. This includes a more granular report of the activity.

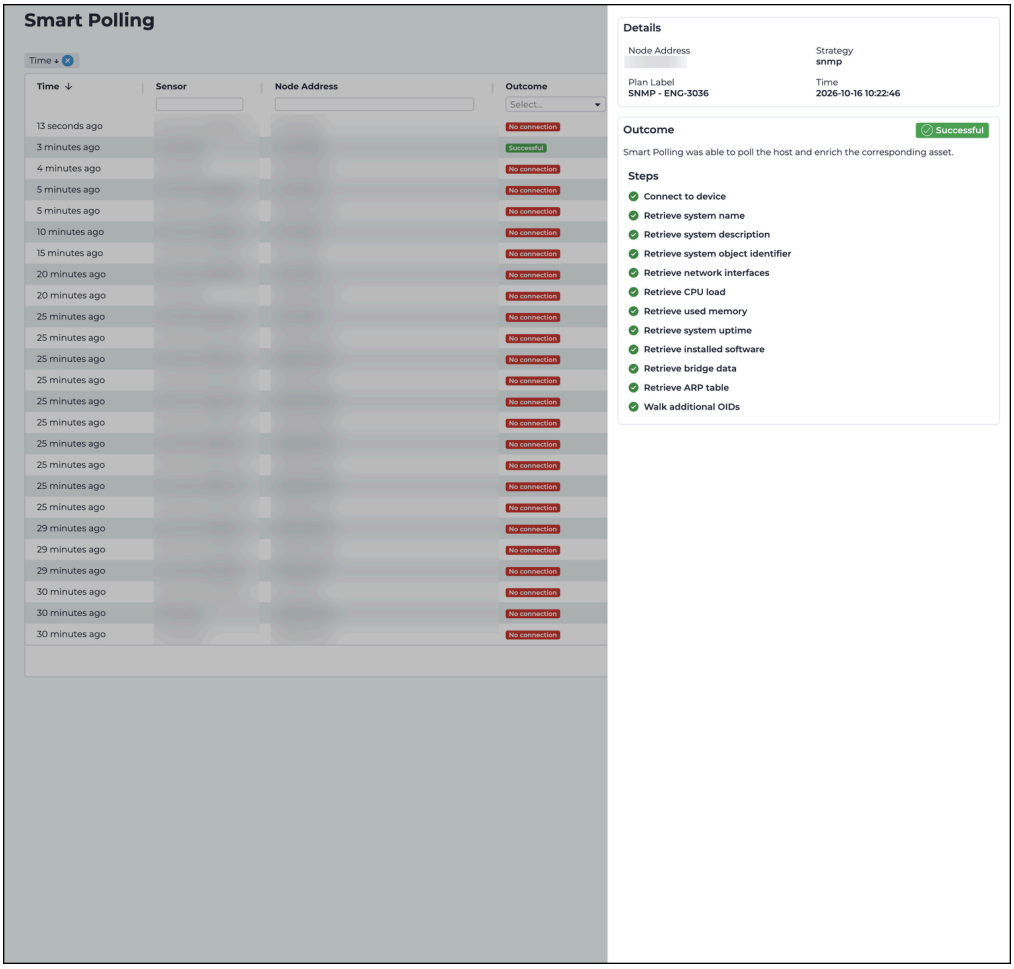


Figure 55. Activity Log sidebar

Column descriptions

Column label	Description
Time	The time the sensor connected to the monitored node.
Sensor	The hostname of the sensor performing the poll. This can be the Guardian sensor or a connected Arc. When the poll runs on the Guardian itself, the value is **This sensor**.
Node address	The address of the polled node.
Outcome	The result of the polling.
Plan label	The label of the executed plan.
Strategy	The strategy of the executed plan.

Chapter 7. Arc overview



Arc™ *is a host-based sensor that detects and defends against malicious or compromised endpoints, and insider attacks. You can use Arc sensors to aggregate data for analysis and reports, either on-premises, or in the Vantage cloud.*

General

When detecting cyberthreats, identifying vulnerabilities, or analyzing anomalies in your processes, it is critical to have as much detailed network and system information as possible. More accurate and timely access to data leads to better diagnostics and a faster time to repair.

Arc gives you enhanced endpoint data collection and asset visibility for your networks. This enhanced visibility gives you more:

- Vulnerability assessment capabilities
- Endpoint protection
- Traffic analysis capabilities
- Accurate diagnostics of in-progress threats and anomalies

Arc lets you easily identify compromised hosts that have:

- Malware
- Rogue applications
- Unauthorized *USB* devices
- Suspicious user activity

Operating Systems (OS)

Arc sensors are endpoint executables that run on hosts on these *OSs*:

- Microsoft Windows
- Linux
- Apple macOS

The data that is collected can be sent to either Guardian or Vantage.

Use cases and deployment scenarios

Arc lets you:

- Incorporate air-gapped devices into the analysis and reporting system
- Gain deeper intelligence or insight on critical endpoint devices
- Continuously monitor and protect endpoints
- Automatically deploy sensors across thousands of devices
- Use a low-impact process to scan air-gapped networks
- Deploy with *mobile device management (MDM)* solutions

Continuous monitoring

Because the Arc sensor is on the host, it can monitor traffic continuously, even when the device is not sending or receiving traffic.

User-specific activity monitoring

With more access to endpoint data, Arc lets you connect network traffic and anomalies with specific users. This helps to identify potential insider threats and makes corrective actions both easier and quicker.

Local behavioral analysis (Sigma rules)

Sigma is a common open-source standard that lets you analyze log files to identify malicious events. They are not necessarily related to network artifacts, and as such, would not be detected without residing on a machine. Nozomi Networks Labs curates all the Sigma rules that are loaded into Arc. A *Threat Intelligence (TI)* active license is needed to receive curated rules from the upstream Nozomi Networks endpoint.

Temporary deployment

It is not necessary to keep the Arc executable on a host after you have collected information. This means that you can remove it after data has been collected to conserve host resources, and maintain a clean host environment.

Threat prevention

Starting with v2.0, Arc enforces actions to prevent threats in addition to monitoring.

Architecture

It is important to understand the different architecture possibilities that are available with Arc.

You can connect Arc:

- To Guardian
- To Vantage

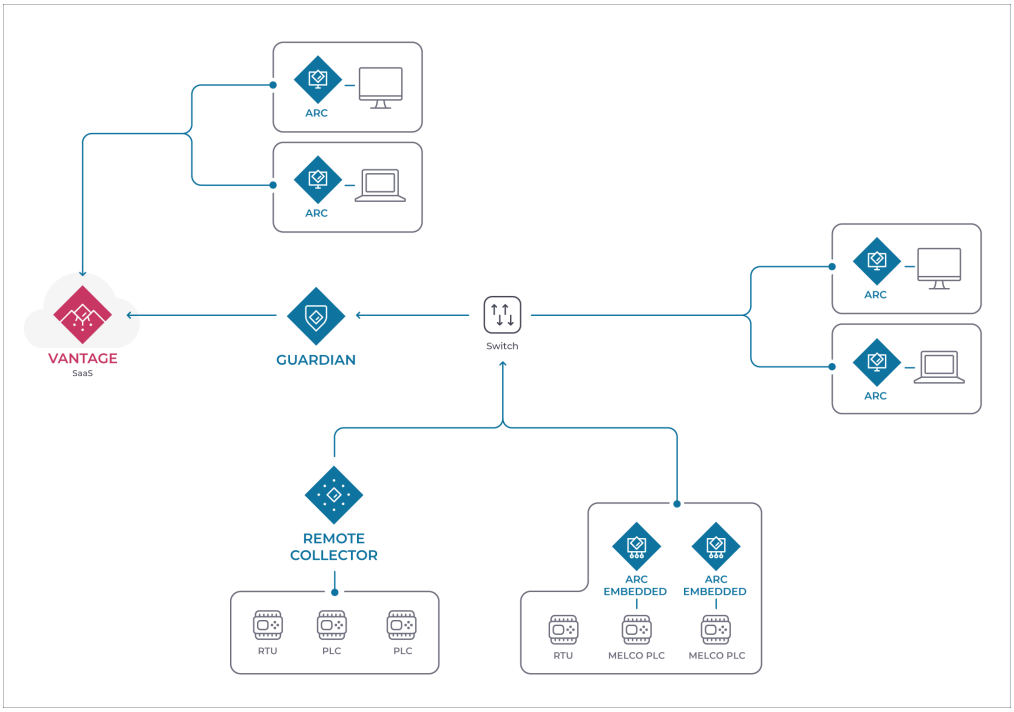


Figure 56. Arc architecture example

Arc in Guardian

*The **Arc** button in the Guardian Web UI lets you access the different pages for Arc.*



Figure 57. Arc button in Guardian Web UI

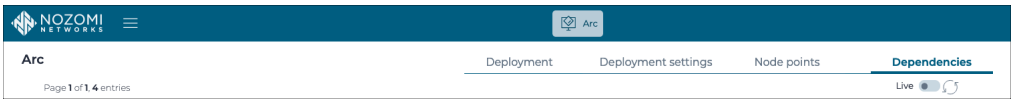


Figure 58. Arc button in Guardian Web UI (not connected to Vantage)

When you select **Arc** in the Guardian Web *user interface (UI)*, you get access to these pages:

- **Deployment**
- **Deployment settings**
- **Node points**
- **Dependencies**

As of *Nozomi Networks Operating System (N2OS)* v26.6.0, the **Dependencies** page is shown for all Guardians. When an upstream appliance manages the dependencies, the dependency files are synchronized from it and cannot be uploaded on the Guardian.

Configure an Arc sensor

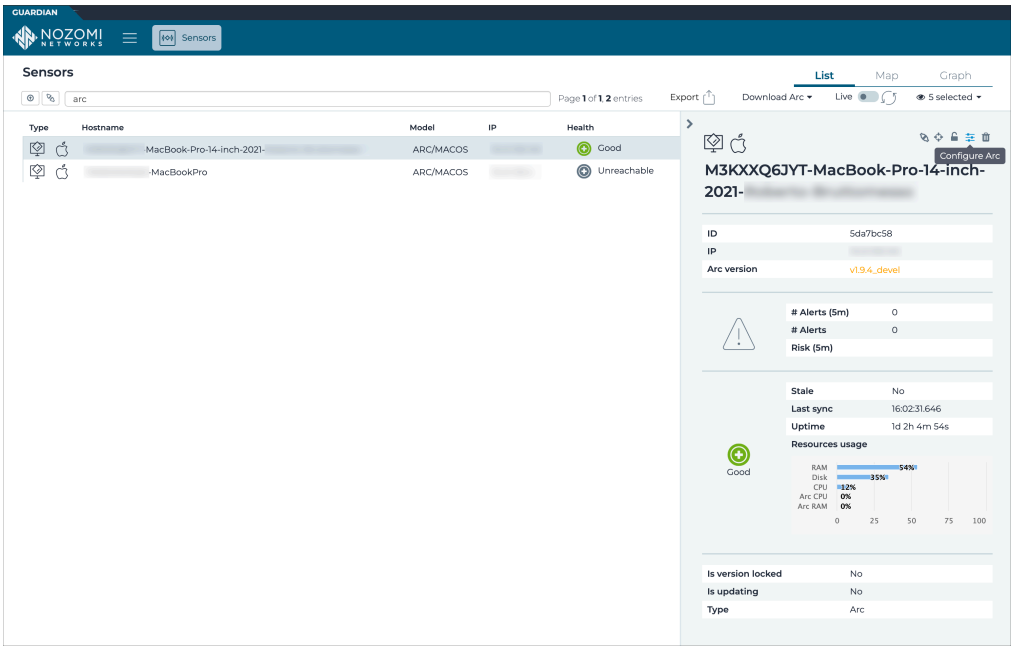


Figure 59. Configure an Arc sensor

You can configure an individual Arc sensor directly from Guardian. To do this, you can select the applicable Arc sensor from the **Sensors** list, and select the icon.

Deployment

*The **Deployment** page shows a table of all the devices available for Arc deployment.*

The table only shows machines which have an *OS* that matches one that Arc supports.

As Guardian detects the installed *OS*, the correct Arc package will be automatically deployed.

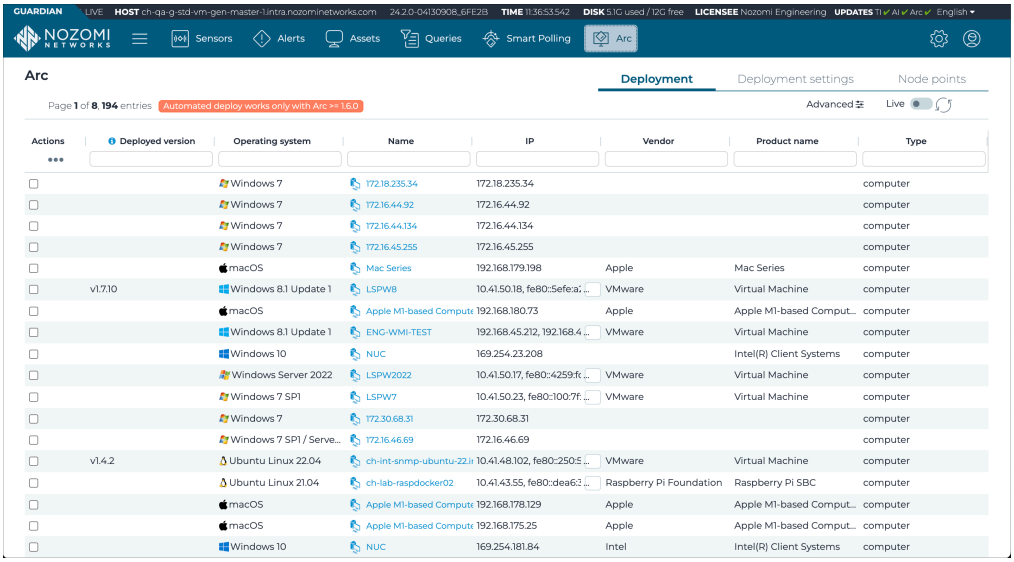


Figure 60. Deployment page

Advanced

The **Advanced** button lets you access the **Advanced** page. For more details, see Advanced (on page 171).

Execution details

The **Execution details** lets you access the **Activity Log**. For more details, see Execution details (on page 173).

Live toggle

The **Live** toggle lets you change live view on, or off. When live mode is on, the page will refresh periodically.

Refresh

The icon lets you immediately refresh the current view.

Actions

The **ACTIONS** column has a checkbox for each row in the table. This lets you select multiple nodes before you then apply an action to them.

The **ACTIONS** menu icon ••• gives you access to these options:

- Select all in current page
- Select none in current page
- Invert selection in current page
- Deploy Service mode: this installs Arc in Service mode for the selected devices
- Remove Service mode: this removes the Arc previously installed in Service mode for the selected devices

Operating System

The **OPERATING SYSTEM** column shows the *OS* for each of the Arc sensors in the table. The field at the top of the column lets you use the *OS* to filter the table.

IP

The **IP** column shows the *IP* for each of the Arc sensors in the table. The field at the top of the column lets you use the *IP* to filter the table.

Vendor

The **VENDOR** column shows the vendor name for each of the Arc sensors in the table. The field at the top of the column lets you use the vendor name to filter the table.

Product name

The **PRODUCT NAME** column shows the product name for each of the Arc sensors in the table. The field at the top of the column lets you use the product name to filter the table.

Type

The **TYPE** column shows the device type for each of the Arc sensors in the table. The field at the top of the column lets you use the device type to filter the table.

Advanced

The **Advanced** *page lets you interact with nodes that have no operating system (OS) detected, or do not show on the same page in the table.*

The default table view only shows nodes that have had their *OS* detected. Also, if you select multiple nodes, actions will only be applied to a single page of nodes. To overcome these limitations, you can use the **Advanced** button to go to the **Advanced** page. This will let you interact with a:

- Set of nodes that cannot be shown on a single page
- Set of nodes that have no *OS* detected



Figure 61. Advanced page

Strategy

Automatic: This selection will use the *OS* that has been detected on the node to automatically choose a deployment strategy. You can select multiple nodes that have a different *OS*. This strategy will ignore a host if it has no *OS*.

SSH (Linux): This selection will force the *SSH* strategy, regardless of the *OS*, and deploy the correct Arc package for Linux.

Query

This field lets you create and execute queries on the nodes. This lets you filter and selectively install packages.

Timeout (seconds)

The **Timeout** dropdown lets you set the amount of time that Arc will try to communicate with a host machine before it skips it and goes to the next one.

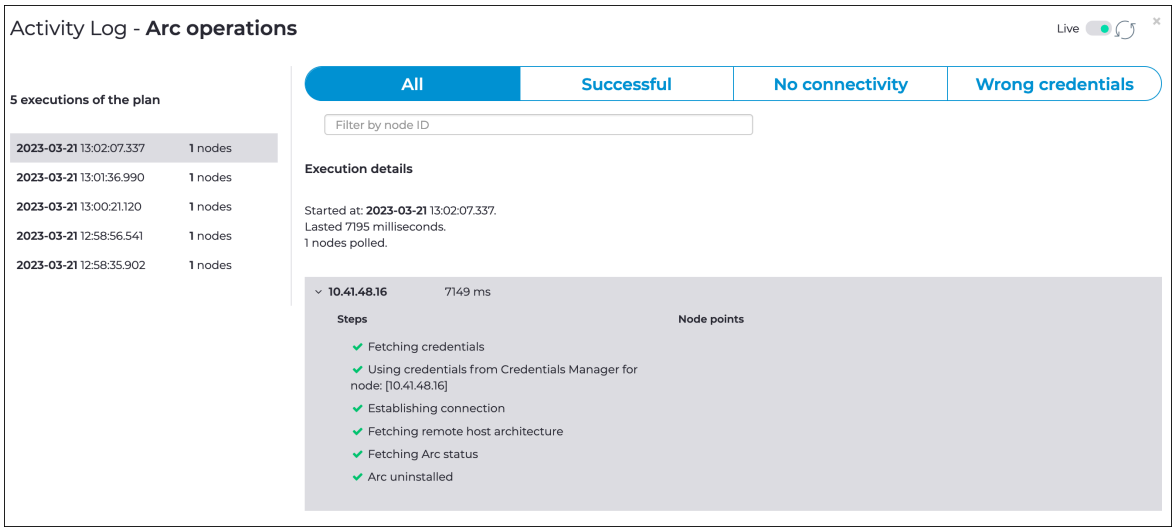
Execution details

*The **Execution details** button gives you access to the **Activity Log**.*

The **Activity Log** lets you troubleshoot the results of the executed deployments. When you select an execution on the left side of the page, you can analyze the selection.

You can use the **Filter by node ID** to focus on a single issue, such as:

- Credential missing, or
- Wrong credentials



Live toggle

The **Live** toggle lets you change live view on, or off. When live mode is on, the page will refresh periodically.

Refresh

The ↻ icon lets you immediately refresh the current view.

Deployment settings

Configure Arc deployment parameters that will be included in downloaded Arc bundles. These settings define how Arc will behave when deployed from Guardian.

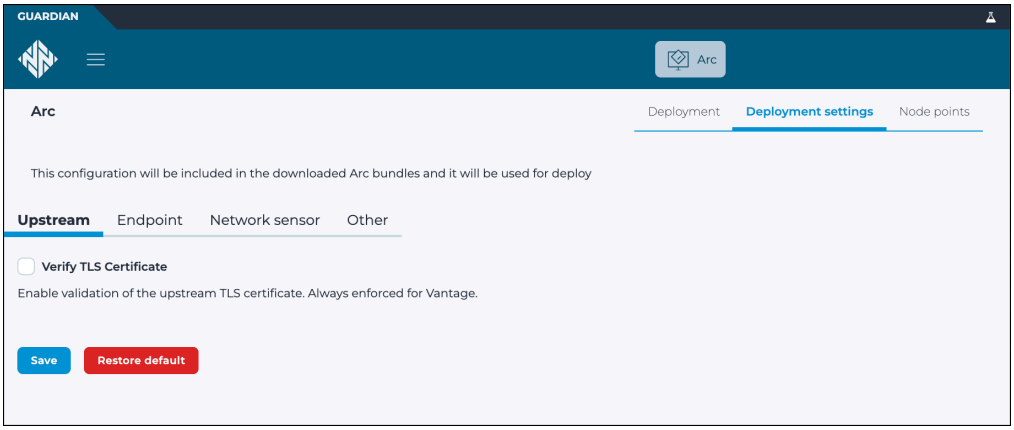


Figure 62. Deployment settings page

The **Deployment settings** page has these tabs:

- Upstream (on page 175)
- Endpoint (on page 176)
- Network sensor (on page 178)
- Other (on page 31)

Restore default

Once the settings have been saved, you can use this button to restore the default configuration.

Upstream

*The **Upstream** page lets you configure the settings for your Arc deployment.*

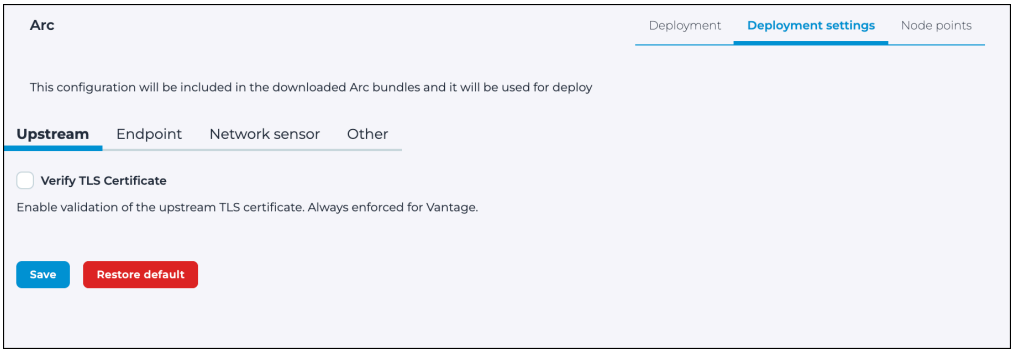


Figure 63. Upstream page

Verify TLS Certificate

Select this option to enforce verification of the upstream *TLS* certificate when connecting to Guardian. When using Vantage, *TLS* certificate validation is always enforced, regardless of this setting.

Restore default

Once the settings have been saved, you can use this button to restore the default configuration.

Endpoint

*The **Endpoint** page lets you configure the settings for your Arc deployment.*

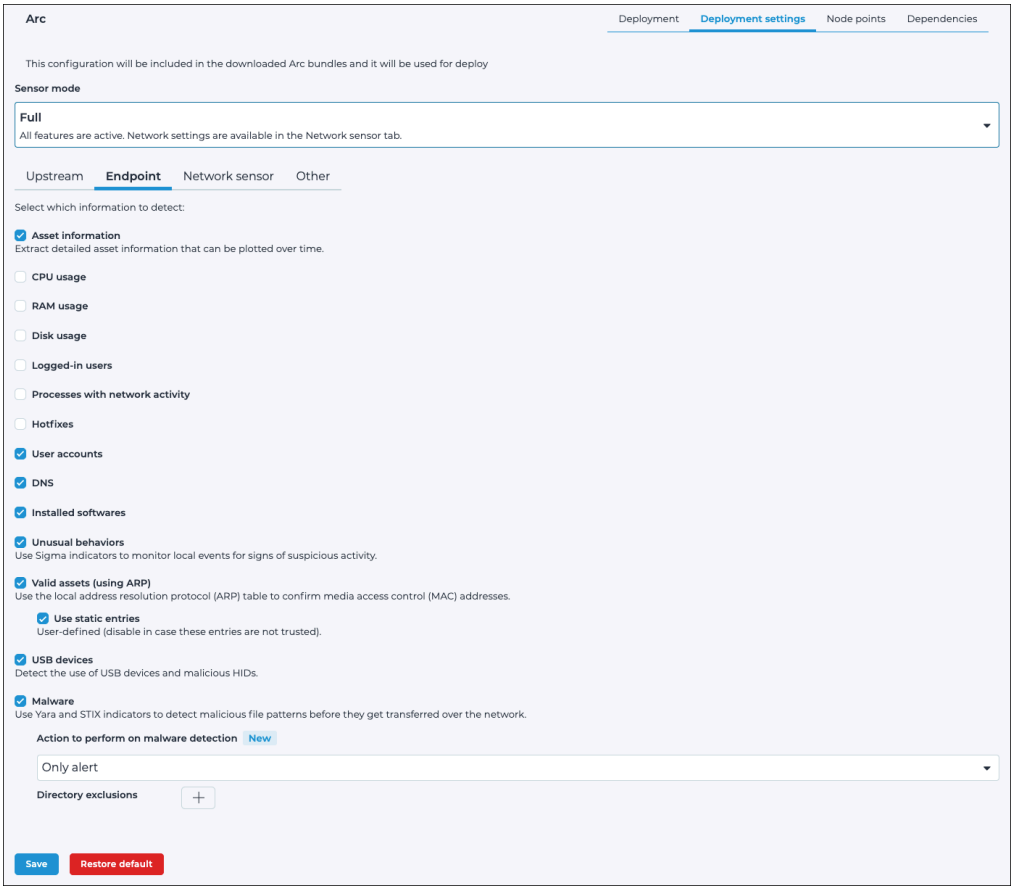


Figure 64. Endpoint page

Asset information

Enable **Asset information** to collect detailed information about the endpoint. Arc collects the following data:

- **CPU usage**: The current processor load of the endpoint.
- **RAM usage**: The amount of memory currently in use on the endpoint.
- **Disk usage**: The current disk space usage on the endpoint.
- **Logged-in users**: The users currently logged in to the endpoint.
- **Processes with network activity**: The running processes that have active network connections, including the local and remote addresses and ports.
- **Hotfixes**: The security patches and updates installed on the endpoint (Windows only).
- **User accounts**: The local user accounts configured on the endpoint.
- **DNS**: The *DNS* servers and configuration of the endpoint.
- **Installed software**: The software packages installed on the endpoint.

Unusual behaviors (Windows only)

This lets you enable/disable Sigma rules for local behavior analysis.

Valid assets (using ARP)

Enable this option to use the local *ARP* table to confirm *MAC* addresses for connected assets.

Enable **Use static entries** to include user-defined static *ARP* table entries in the validation. Only enable this if the static entries are trusted.

USB devices (Windows only)

This lets you enable/disable *USB* detections.

Malware (Windows only)

The **Malware** checkbox lets you enable protection mode and select from the **Action to perform on malware detection** dropdown. This lets you set the action that Arc will take when it finds a malicious file.

You can choose from these options:

- **Only alert**: Receive an alert with no further action
- **Quarantine**: Move the malicious file to the `quarantine` folder, which is located in the Arc installation folder
- **Delete**: Immediately delete the malicious file. Once deleted, the files cannot be recovered

You can also use the **Directory exclusions** feature to exclude specific directories from malware scanning. Select + to open a file browser and select a directory. You can add as many directories as needed. Arc skips the selected directories when scanning for malware.

Restore default

Once the settings have been saved, you can use this button to restore the default configuration.

Network sensor

*The **Network sensor** page lets you configure the settings for your Arc deployment.*

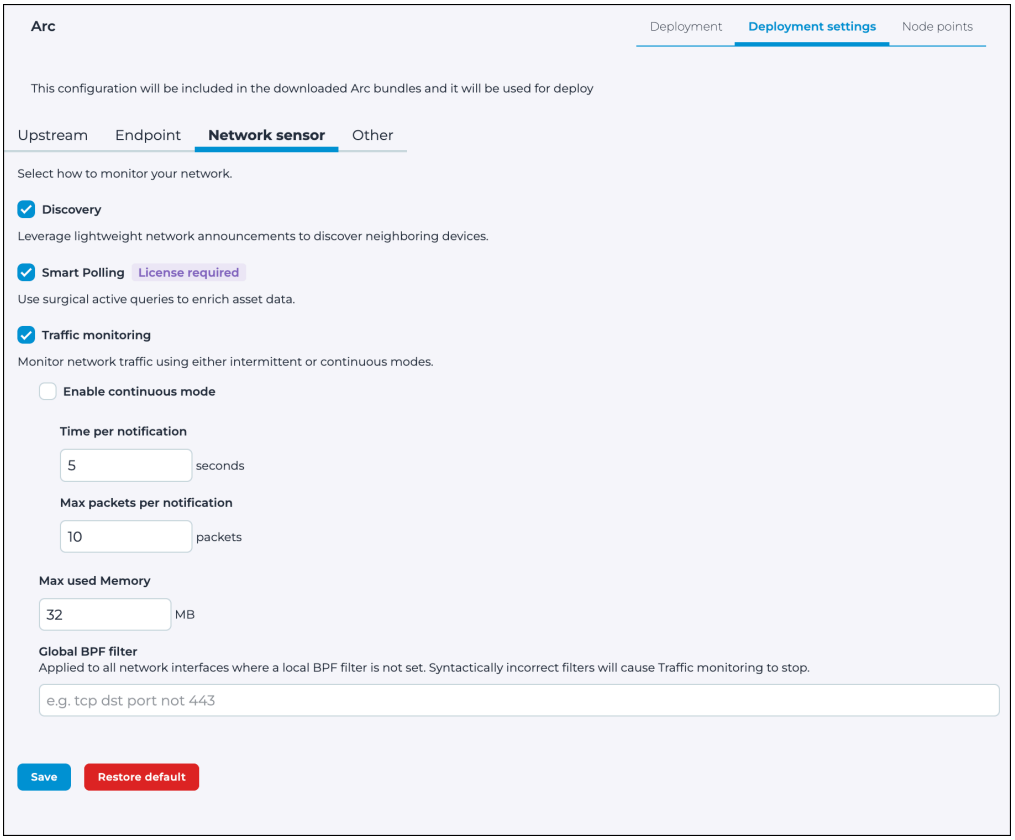


Figure 65. Endpoint page

Discovery

When enabled, this sends out unsolicited lightweight network announcements to discover neighboring nodes.

Discovery uses lightweight protocol-specific broadcast messages to identify network devices. These messages trigger a response from the devices, which includes identity information. The process is repeated at predefined intervals. At each interval, the sensor will identify the suitable network interfaces and send broadcast messages through them to discover devices on each subnetwork connected to the sensor.

Smart Polling

This lets you enable/disable the execution of Smart Polling strategies from Arc. When enabled, this sends out Smart Polling queries following remote requests coming from Guardian to poll assets that Arc can reach, or assets that have been identified with Discovery.

Note:

Smart Polling requires that a Smart Polling license is enabled upstream.

To force Smart Polling from a specific Arc sensor, even when Guardian was the first to monitor a node, you can use a *CLI* command such as: `vi node 192.168.1.1 capture_device arc[1e6a174c]` In this example, `192.168.1.1` is an *IP* address of a node you want to poll from a specific Arc sensor. `1e6a174c` are the first eight characters of the Arc sensor *ID* that you want to poll the node with. To find that sensor *ID*, you can select the Arc sensor from the **Sensors** page of your Guardian and read the **ID** field in the right pane. To reset the behavior, you can set the `capture_device` back to the value of the Guardian interface.

Traffic monitoring

When enabled, this checkbox lets you enable/disable traffic monitoring.

Enable continuous mode

This checkbox lets you enable/disable continuous mode. For more details, see **Continuous mode**.

Arc uses two different methods for traffic monitoring:

- Intermittent mode
- Continuous mode

Intermittent mode: This is the default mode, the traffic is monitored, or sniffed, for a duration of 10 seconds at each notify. The purpose of this limitation is to preserve the resources of the host machine, which prevents excessive memory, or *CPU*, spikes. You can configure these options:

Continuous mode: This mode sniffs traffic continuously from the host's network interface controllers. Depending on the amount of sniffed traffic, continuous mode might utilize more *CPU* and memory on the host. As the traffic is processed upstream, the performance of the remote endpoint is also affected. You can configure:

- **Time [s] per notification**
- **Max packets per notification**
- **Max used Memory (MB)**: this value can be tuned to allow more or less traffic buffering in case the traffic to process exceeds the Arc and network capacity to send it out

Global BPF filter

This field lets you set a Global BPF filter to apply to all the network interfaces. Filters that are applied to single interfaces will take precedence over the global one.

Network interface

This dropdown lets you select a network interface to configure. Each network interface can then be enabled, and be tuned with a monitoring filter.

If you add, remove, or edit the network interfaces on the host, Arc does not automatically add it to the list of sniffing interfaces. For example, if you add a new network card, to enable Arc to use it, you should stop Arc, and then start it again.

Restore default

Once the settings have been saved, you can use this button to restore the default configuration.

Other

*The **Other** page lets you configure the settings for your Arc deployment.*

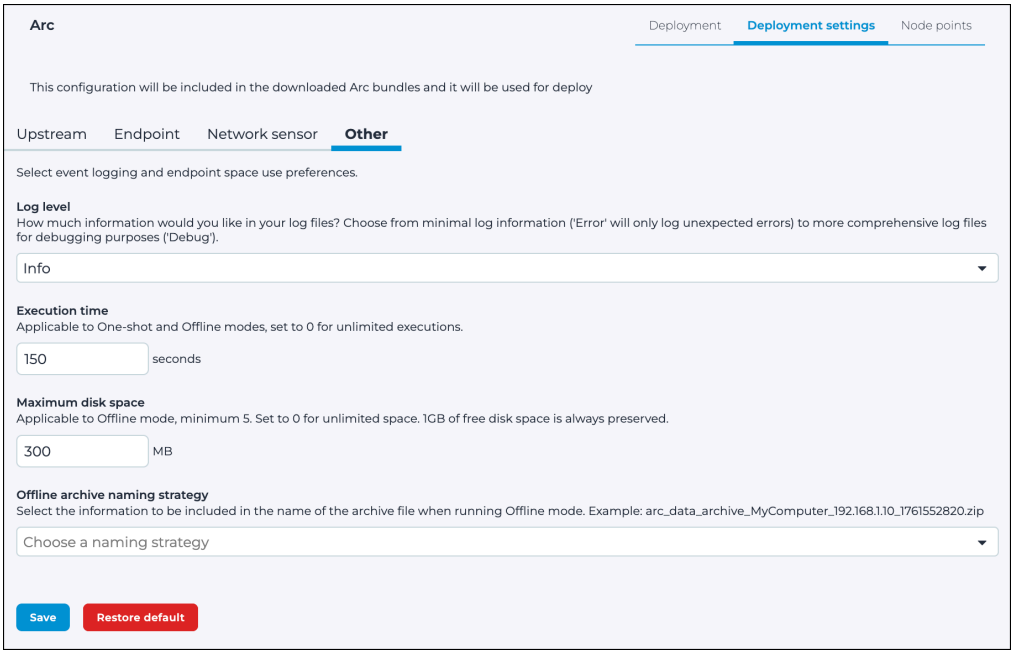


Figure 66. Other page

Log level

This dropdown lets you select the verbosity level for the log files. The options are:

- Debug
- Info
- Warning
- Error

The logging system options have an increasing level of verbosity, from the least verbose to the most verbose. Error < Warning < Info < Debug.

- Error: Creates a minimalistic log, only unexpected errors are logged
- Warning: Creates extra errors that might show on some *OS*s, but that are generally considered as acceptable
- Info: Logs relevant successful events, it shows the program's progress (recommended)
- Debug: Logs extra events that are normally useful for debugging purposes. Given its verbosity it is best to activate it only when debugging activities are involved

Note:

This setting applies to:

- `arc.log.txt`
- `arc_service.log.txt`
- `sp.log.txt`

It doesn't apply to `setup.log`, which covers Arc's initialization process, before the configuration has been loaded and applied.

For more details, see Log files (on page).

Execution time

This field lets you set the time that Arc will run to collect data. This is applicable for One-shot and Offline modes.

Note:
When this is set to 0, the execution time is interpreted as infinite.

Maximum disk space

This field lets you control the maximum amount of disk space in *MB* that will be used for Offline mode.

Offline archive naming strategy

This lets you select the information that will be included in the name of archive files created while running Offline mode.

Restore default

Once the settings have been saved, you can use this button to restore the default configuration.

Node points

*The **Node points** page shows data points that are collected over time, and represent the state of the target machine.*

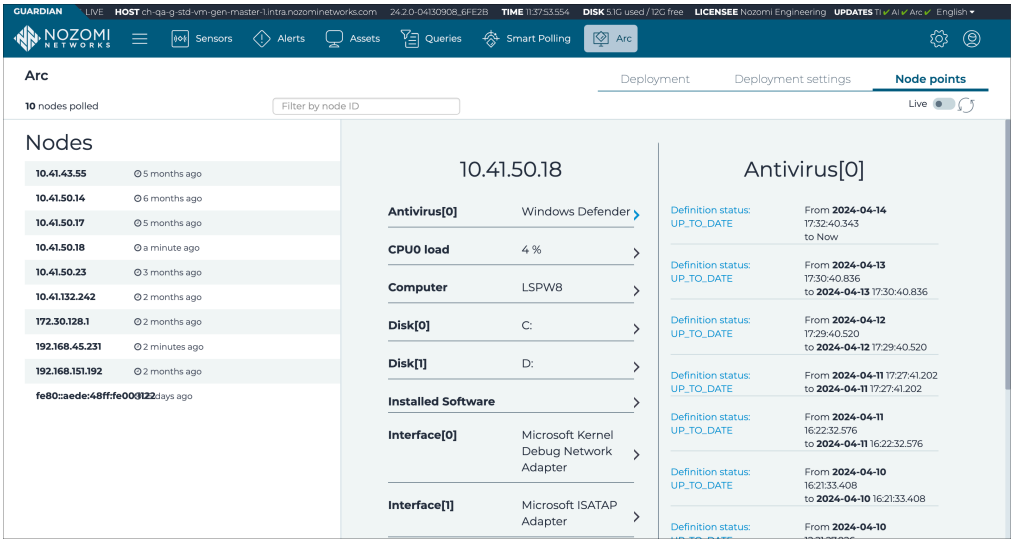


Figure 67. Node points page in Guardian

Node points count

This shows the number of the nodes polled.

Filter by node ID

This field lets you use the node *ID* to filter the nodes.

Live toggle

The **Live** toggle lets you change live view on, or off. When live mode is on, the page will refresh periodically.

Refresh

The  icon lets you immediately refresh the current view.

Nodes

The list of nodes that show at least one node point.

Dependencies

*The **Dependencies** page shows the dependencies that the connected Arc sensors can use, and lets you manage their dependency files and apply actions to multiple sensors at once.*

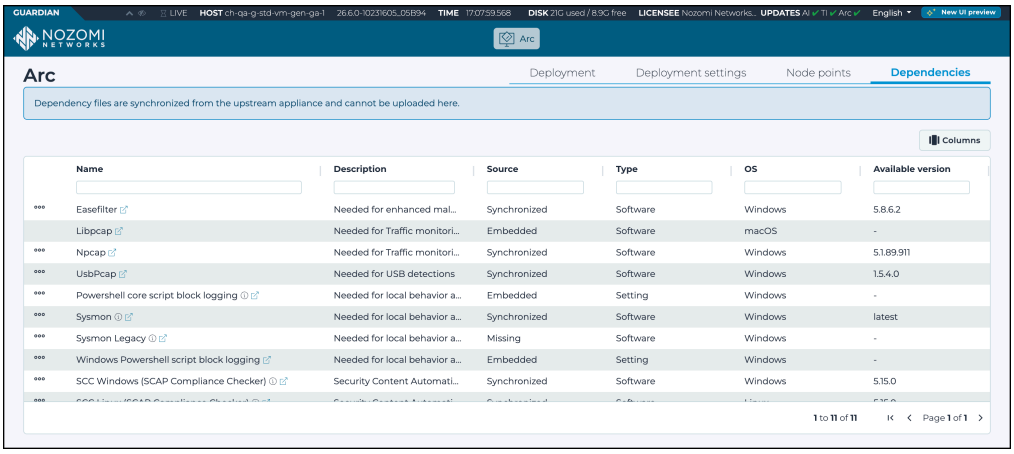


Figure 68. Dependencies page in Guardian

The table shows one row for each dependency that Arc sensors can use on the appliance. You can use the filter under each column header to narrow the list. The table has these columns:

- **Name**: The dependency's name. Some dependencies also show an information icon with more details, an external link to more information, or both
- **Description**: Why the dependency is needed
- **Source**: Where the dependency file comes from
- **Type**: **Software** for dependencies that are installed on the sensor, or **Setting** for dependencies that are enabled or disabled on the sensor
- **OS**: The operating system that the dependency applies to
- **Available version**: The version of the dependency file that is available. Some dependencies show **latest** instead of a version number, and a dash (-) shows when no version is available

The **Source** column shows one of these values:

- **Embedded**: The dependency ships with Arc
- **Synchronized**: The dependency file is synchronized from the upstream appliance
- **Uploaded**: The dependency file is uploaded on this appliance
- **Missing**: No dependency file is held for the dependency; you must upload it

Note:

When an upstream appliance manages the dependencies, the dependency files are synchronized from it and cannot be uploaded on this appliance. The **Upload dependency file** action isn't available.

When no upstream appliance manages the dependencies, the **Available version** column shows a dash (-) and the **Download dependency file** action is unavailable. You use the **Upload dependency file** action to upload the files that the dependencies need.

For each dependency, the actions menu shows the actions that apply to it:

- **Download dependency file** shows for all dependencies except Sysmon and Sysmon Legacy, which Nozomi Networks doesn't distribute due to licensing limitations. It's unavailable when no dependency file is held. The Sysmon file is downloaded directly from Microsoft when a connection is available; otherwise you upload it manually. Sysmon Legacy always needs a manual upload, because Microsoft doesn't distribute download links for older versions of Sysmon
- **Upload dependency file** shows only when no upstream appliance manages the dependencies
- **Install on Arc sensors**, **Update on Arc sensors**, and **Uninstall from Arc sensors** show for software dependencies
- **Enable on Arc sensors** and **Disable on Arc sensors** show for setting dependencies

Note:

Install and Update are unavailable when no dependency file is held for the dependency.

The actions that apply to Arc sensors open a dialog where you select the applicable sensors. The dialog shows the host, IP address, model, and Arc version of each sensor, together with the state of the dependency on it.

Related information

Apply an Arc dependency action to Arc sensors (on page 185)

Apply an Arc dependency action to Arc sensors

You can apply an Arc dependency action to one or more Arc sensors from the Dependencies page of the Arc section.

Before you begin

The sensor must be running Arc v2.10.0, or later.

About this task

The available actions depend on the type of the dependency and on whether an upstream appliance manages the dependency files. For more details, see Dependencies (on page 183).

Procedure

1. In the top navigation bar, select **Arc**.
2. Select the **Dependencies** tab.
3. For the applicable dependency, open the actions menu and select the action to apply, for example **Install on Arc sensors**.

   Result: A dialog shows, listing the Arc sensors that the dependency applies to.
4. Select the applicable Arc sensors.
   Sensors that have a pending request for the same dependency, or that require a reboot, aren't selectable.
5. Select **Confirm**.

   Result: A confirmation message shows, with the number of selected Arc sensors.
6. **Optional:** Select the **Reboot the sensor if the operation requires it** checkbox.

   > **Note:**
   > If this operation requires a reboot, the sensor restarts automatically to complete it.
7. Select **Confirm**.

   Result: The **Request queued** result shows on how many sensors the request was queued. Each sensor runs the action the next time it connects. The result also shows how many sensors already had the same request pending, and on how many sensors the request couldn't be queued.
8. Select **Close**.

Results

The action is queued on the selected Arc sensors.

Related information

Dependencies (on page 183)

Chapter 8. Network



*The **Network** page gives you access to multiple pages that show nodes, links, sessions in tabular format and graphs and traffic in a graphical format.*

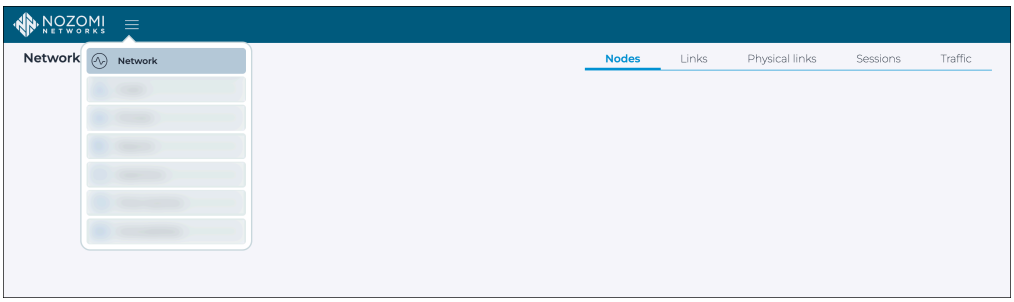


Figure 69. Network page

The **Network** page has these tabs:

- Nodes (on page 190)
- Links (on page 200)
- Physical links (on page 210)
- Sessions (on page 211)
- Traffic (on page 216)

Nodes

*The **Nodes** page lets you view all the network nodes in your environment and perform actions on them.*

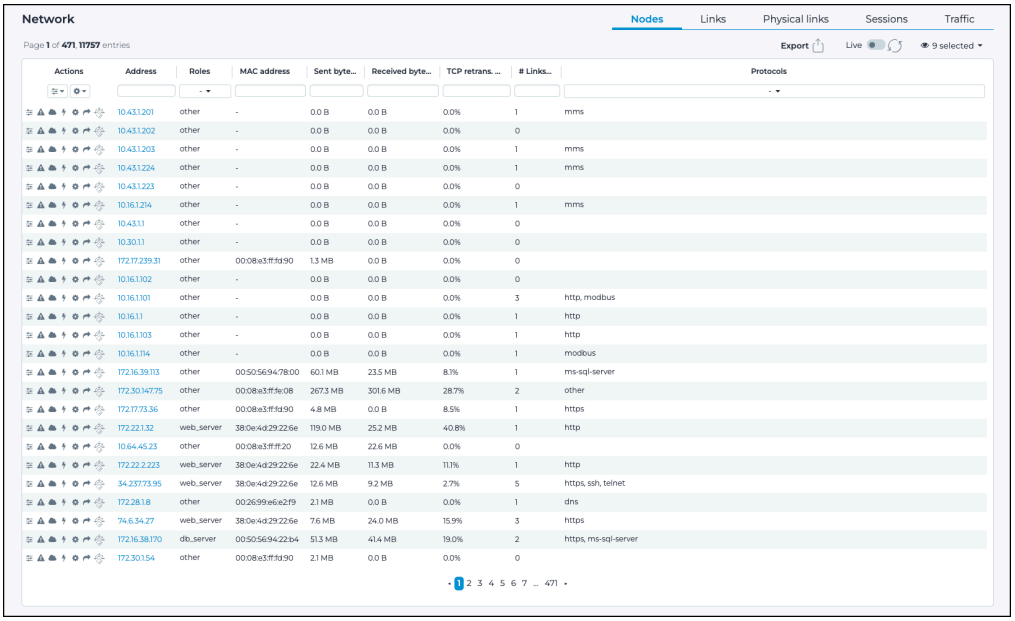


Figure 70. Nodes page

Export

The **Export** icon lets you export the current list in either *CSV* or Microsoft Excel format.

Live / refresh

The **Live** icon lets you change live view on, or off. When live mode is on, the page will refresh approximately every five seconds.

Column selection

The columns selection icon lets you choose which columns to show or hide.

Configure a node

*The **Nodes** page lets you configure nodes.*

Procedure

1. In the top navigation bar, select icon > **Network**.

   Result: The **Network** page opens.

2. Select **Nodes**.

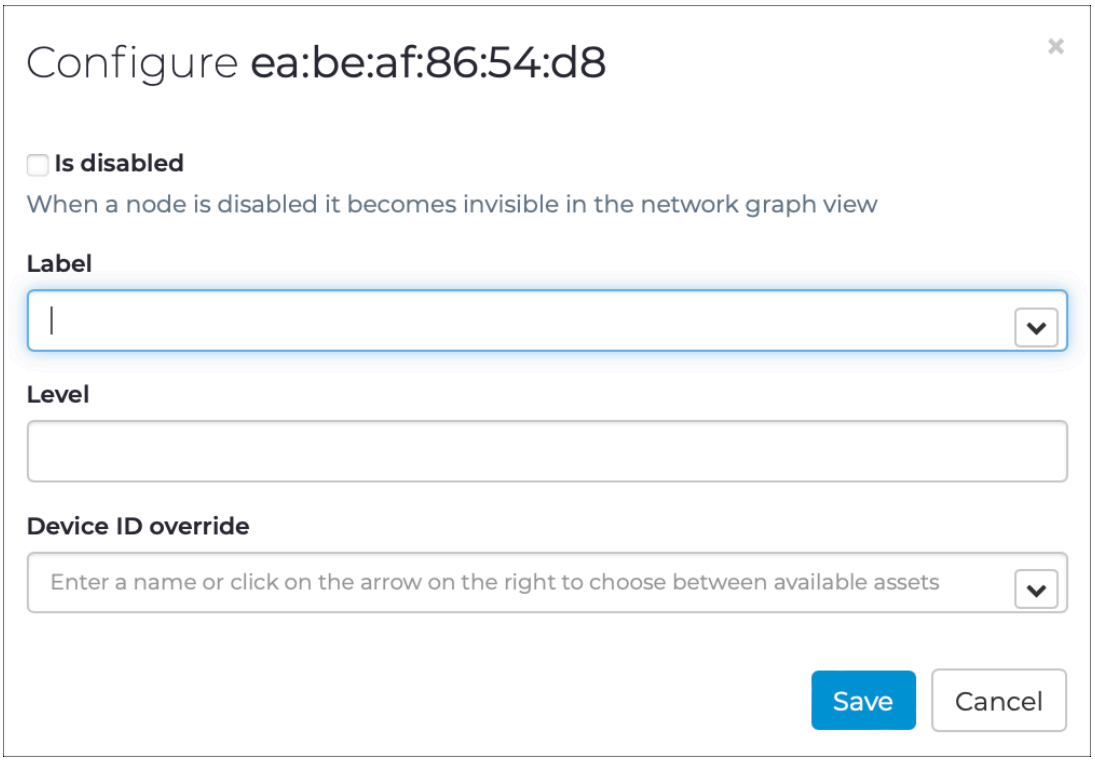
   Result: The **Nodes** page opens.

3. To the left of the applicable node, select the configure icon.

   Result: A dialog shows.

4. **Optional:**

   To make the node(s) invisible in network graph view, select **Is disabled**.



5. From the **Label** dropdown, select an asset and assign the node to it.
6. In the **Level** field, enter a level in accordance with the Purdue model classification.
7. From the **Device ID override** dropdown, remove or re-assign Device ID to overwrite the automatically assigned **Device ID**.
8. Select **Save**.

Results

The node has been configured.

Show alerts for a node

*The **Nodes** page lets you show alerts for a selected node.*

Procedure

1. In the top navigation bar, select icon > **Network**.

   Result: The **Network** page opens.

2. Select **Nodes**.

   Result: The **Nodes** page opens.

3. To the left of the applicable node, select the icon.

   Result: A list of all the requested traces for the node shows.

Show requested traces for a node

*The **Nodes** page lets you show requested traces for a selected node.*

Procedure

1. In the top navigation bar, select icon > **Network**.

   Result: The **Network** page opens.

2. Select **Nodes**.

   Result: The **Nodes** page opens.

3. To the left of the applicable node, select the icon.

   Result: A list of all the requested traces for the node shows.

Request a trace for a node

*The **Nodes** page lets you request a trace for a selected node.*

Procedure

1. In the top navigation bar, select ☰ **icon > Network**.

   Result: The **Network** page opens.

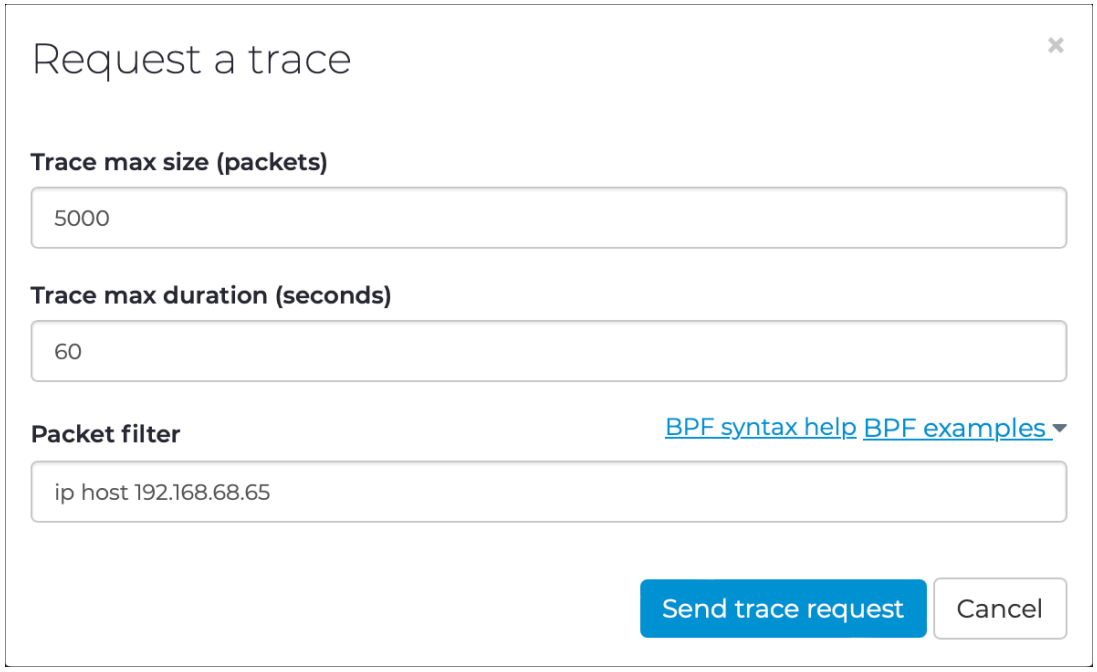
2. Select **Nodes**.

   Result: The **Nodes** page opens.

3. To the left of the applicable node, select the ⚡ icon.

   Result: A dialog shows.

4. To set the maximum packet size, in the **Trace max size (packets)** field, enter a value.



> **Note:**
> The default size is 5000 packets.

5. To set the maximum duration of the trace, in the **Trace max duration (seconds)** field, enter a value.

> **Note:**
> The default value is 60 seconds.

6. **Note:**
   The **Packet filter** field is automatically populated with a *Berkeley Packet Filter (BPF)* that captures the packets to/from the selected node, but you can customize this.

   If necessary, customize this field.

   Note:
   You can select **BPF syntax help** to show more information on *BPF* syntax.

   Note:
   You can select **BPF examples** to see some examples.

7. Select **Send trace request**.

Results

The trace has been requested.

Manage learning for a node

*The **Nodes** page lets you manage the learning for a selected node.*

About this task

The **Manage Learning** dialog lets you learn and delete the entire node and its individual details, such as *IP* or *MAC* address.

Note:

The icon adjacent to the node shows its status. When the node details are:

- A green icon shows when the node details have been entirely learned
- An orange icon shows when the node details are only partially learned
- A red icon shows when the node details have not been learned

Individual details have either a green or red icon, depending on whether they are learned or not. When you learn, or delete a node, all of its details are affected in the same way. When you learn, or delete, an individual detail, only that detail's learning status changes.

Procedure

1. In the top navigation bar, select icon > **Network**.

   Result: The **Network** page opens.

2. Select **Nodes**.

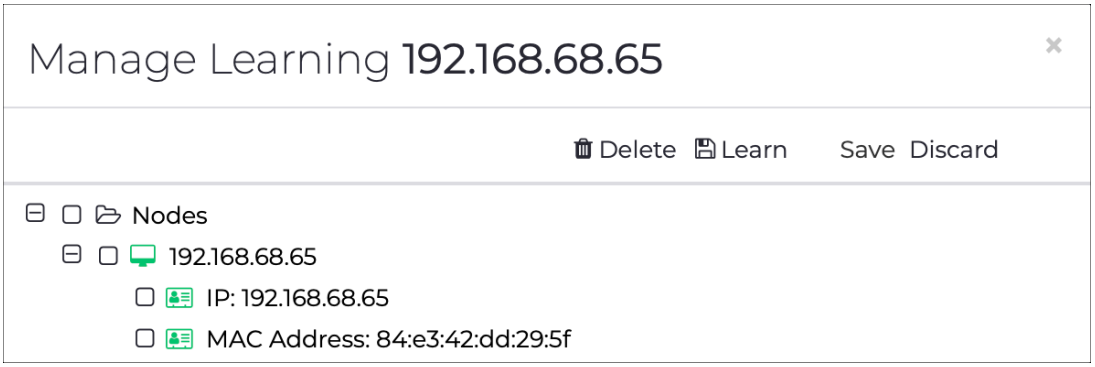
   Result: The **Nodes** page opens.

3. To the left of the applicable node, select the icon.

   Result: A dialog shows.

4. **Optional:**

   Select an item(s) to delete.



   Choose from:

   - To delete the *IP* address, select the checkbox to the left.
   - To delete the *MAC* address, select the checkbox to the left.
   - To delete the node, and both the *IP* and *MAC* addresses, select the checkbox to the left.

   a. Select **Delete** to delete the selected item(s).

5. **Optional:** Select an item(s) to learn.

   Choose from:

   - To learn the *IP* address, select the checkbox to the left.
   - To learn the *MAC* address, select the checkbox to the left.
   - To learn the node, and both the *IP* and *MAC* addresses, select the checkbox to the left.

   a. Select **Learn** to learn the selected item(s).

6. Select **Save**.

Navigate from a node

*The **Nodes** page lets you navigate to related entities from a selected node.*

Procedure

1. In the top navigation bar, select **icon > Network**.

   Result: The **Network** page opens.

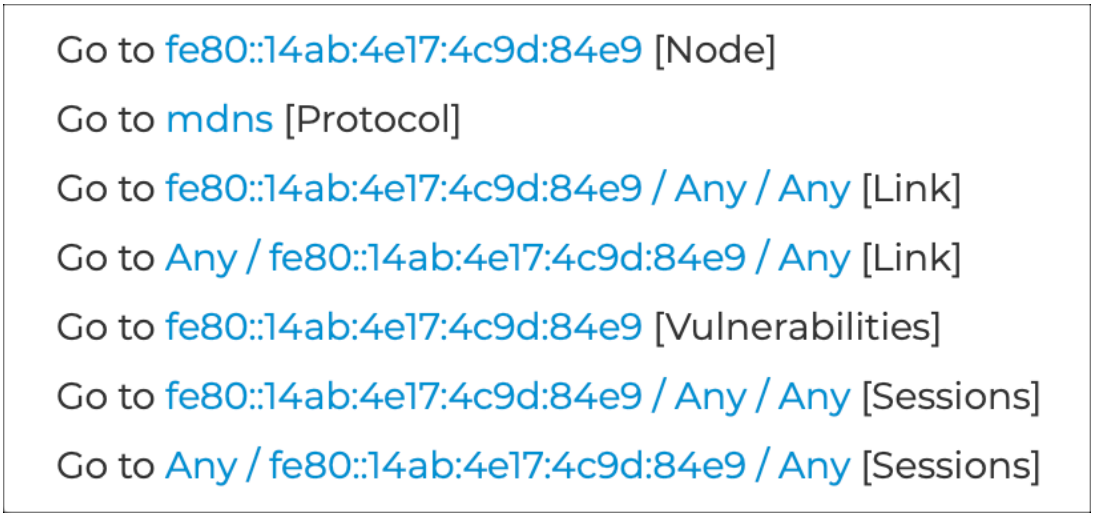
2. Select **Nodes**.

   Result: The **Nodes** page opens.

3. In the Actions column, to the left of the applicable node, select the icon.

   Result: A list of related entities shows.

4. Select the hyperlink that you want to navigate to.



Results

The entity shows in the applicable page.

Add a node to a Smart Polling plan

*The **Nodes** page lets you add a node to an existing Smart Polling plan.*

About this task

The Smart Polling icon only shows if Smart Polling is present.

Procedure

1. In the top navigation bar, select **icon > Network**.

   Result: The **Network** page opens.

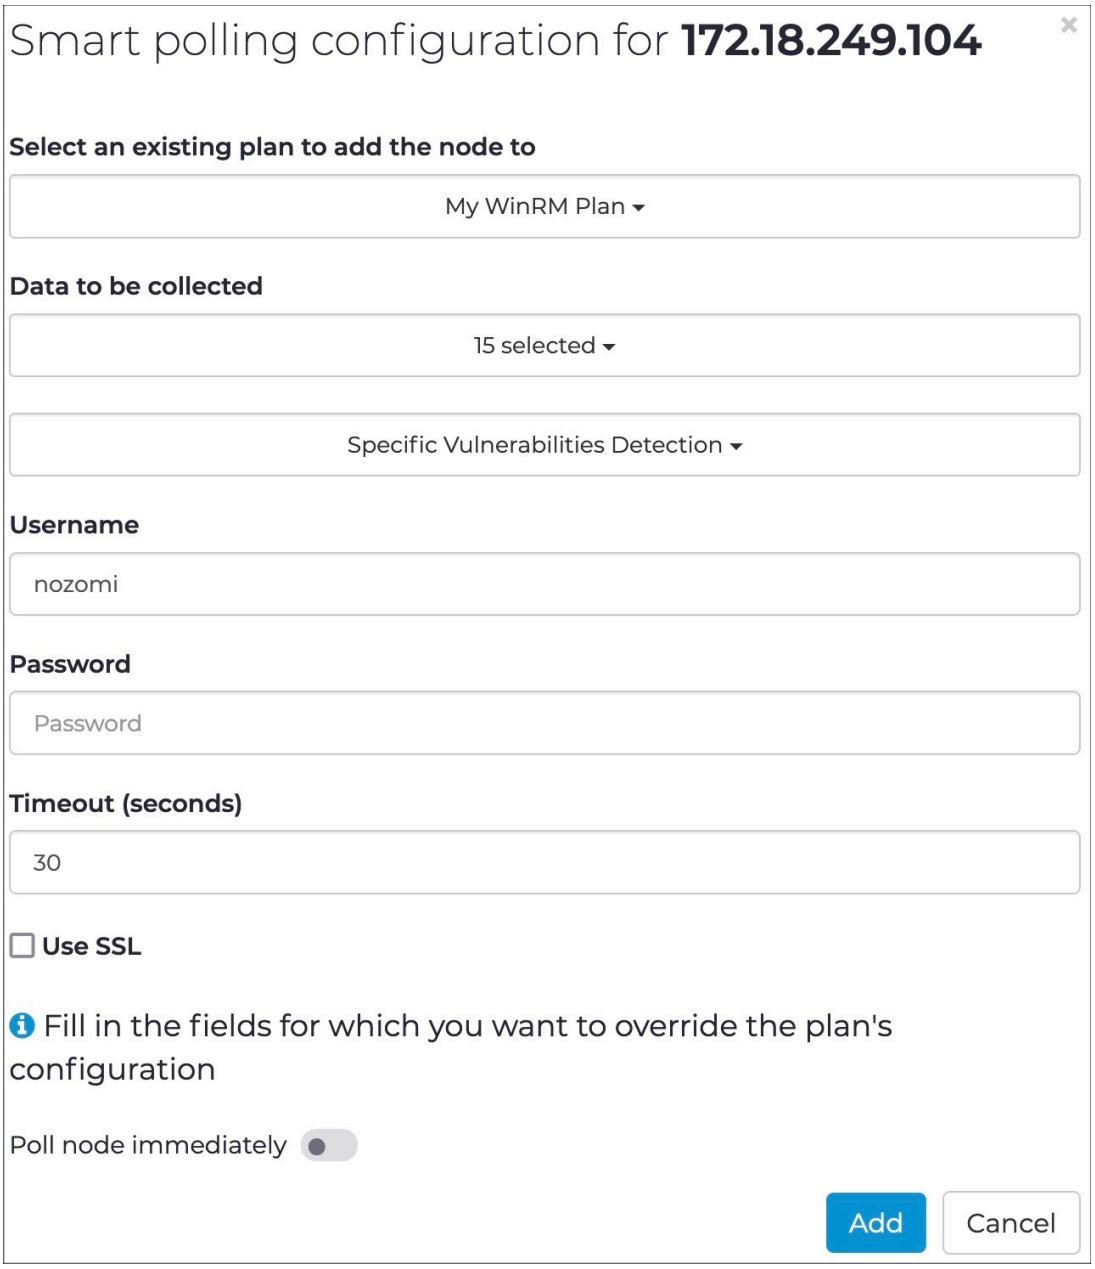
2. Select **Nodes**.

   Result: The **Nodes** page opens.

3. To add a node to a plan with an optionally different configuration from the plan's original one, in the Actions column, to the left of the applicable node, select the icon.

   Result: A dialog shows.

4. From the **Select an existing plan to add the node to** dropdown, select an existing plan to that you would like to add the node to.



5. **Optional:** In the **Timeout (seconds)** field, enter a value to override the value in the current plan.
6. **Optional:** To poll the node immediately, set the **Poll node immediately** toggle to on.

> **Note:**
> If you do not set the **Poll node immediately** toggle to on, the node will be polled at the next execution of the selected plan.

Links

*The **Links** page shows all the links in your environment.*

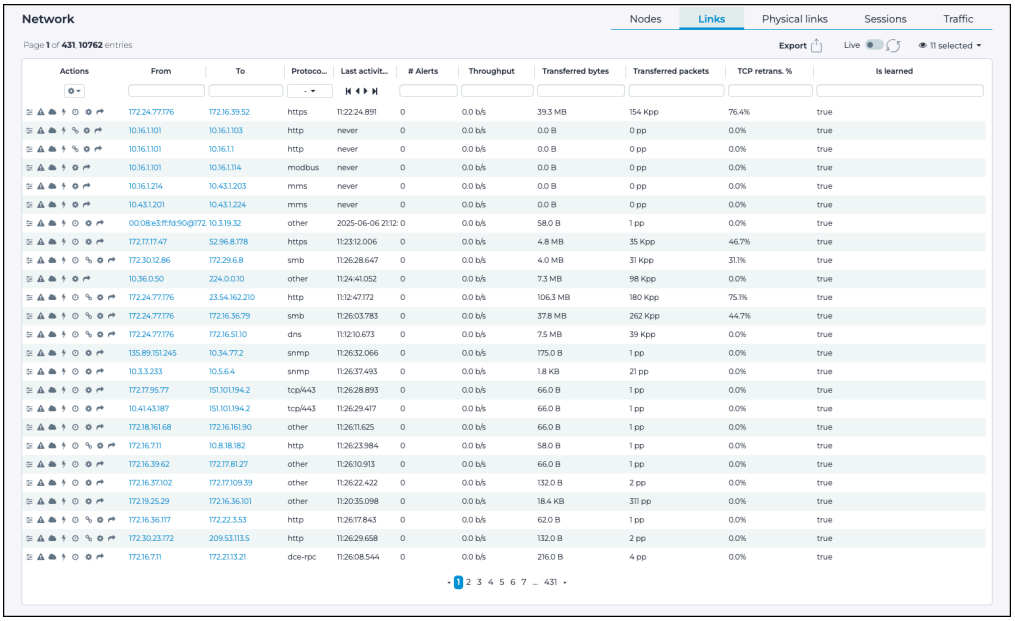


Figure 71. Links page

Export

The **Export** icon lets you export the current list in either *CSV* or Microsoft Excel format.

Live / refresh

The **Live** icon lets you change live view on, or off. When live mode is on, the page will refresh approximately every five seconds.

Column selection

The columns selection icon lets you choose which columns to show or hide.

Link events

Link events are shown in a graphical format.

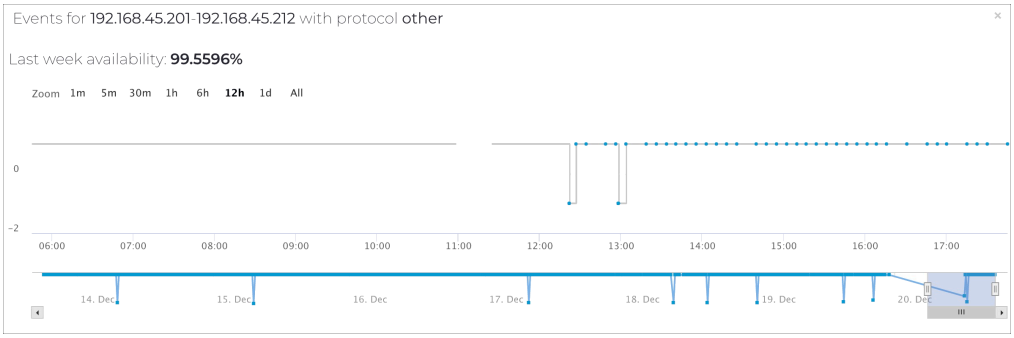


Figure 72. Link events

Availability

Link availability is based on UP and DOWN events.

Time span

You can use the time span control to view only the events in the specified time range. The available options are:

- Zoom
- 1m(inute)
- 5m
- 30m
- 1h(our)
- 6h
- 12h
- 1d(ay)
- All

Graphical history

A point with a value of 1 represents an UP event, a value -1 represents a DOWN event.

Link availability

A history of events is stored for each link.

Two events are of particular interest for computing availability:

- UP - This occurs when an activity is detected on an inactive link
- DOWN - This occurs when an active link stops its activity

Each event has a timestamp to track the precise moment of its occurrence.

Guardian computes the total downtime of a link by taking into consideration the history of events within a finite time window. Then, it sums the time spans of all events starting with a DOWN event and ending with an UP event. All links are considered active by default. Therefore, the availability of the link is 100% minus the percentage of the total downtime.

Track availability

The **Track availability** feature allows an accurate computation of availability. It enables the monitoring of activity on a link at regular intervals, generating extra UP and DOWN events, depending on the detected activity on both sides of the link during the last interval.

We recommend that you select a value greater than the expected link polling time to avoid checks that are too frequent and are likely to produce spurious DOWN events.

link_events generation is disabled by default. To enable it, see the configuration rule shown in **Configuration**.

Configure a link

*The **Links** page lets you configure links.*

Procedure

1. In the top navigation bar, select icon > **Network**.

   Result: The **Network** page opens.

2. Select **Links**.

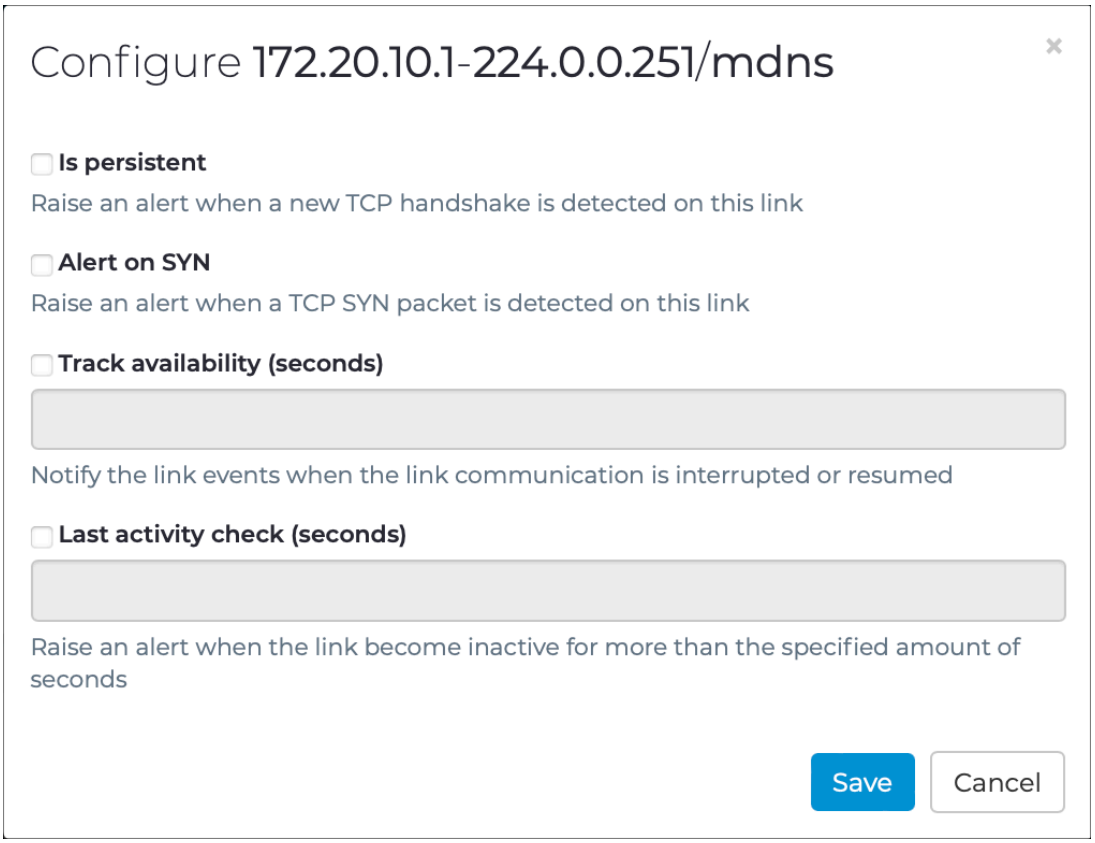
   Result: The **Links** page opens.

3. To the left of the applicable link, select the configure icon.

   Result: A dialog shows.

4. **Optional:**

   To raise an alert when a new *TCP* handshake is detected on the link, select **Is persistent**.



5. **Optional:** To raise an alert when a *TCP* SYN packet is detected on the link, select **Alert on SYN**.
6. **Optional:** To notify the link events when the link communication is interrupted or resumed., select **Track availability (seconds)** and enter a value.

7. **Optional:** To raise an alert when the link becomes inactive for more than the specified number of seconds, select **Last activity check (seconds)** and enter a value.
8. Select **Save**.

Results

The link has been configured.

Show alerts for a link

*The **Links** page lets you show alerts for a selected link.*

Procedure

1. In the top navigation bar, select **icon > Network**.

   Result: The **Network** page opens.

2. Select **Links**.

   Result: The **Links** page opens.

3. To the left of the applicable link, select the icon.

   Result: A list of all the requested traces for the link shows.

Show requested traces for a link

*The **Links** page lets you show requested traces for a selected link.*

Procedure

1. In the top navigation bar, select **icon > Network**.

   Result: The **Network** page opens.

2. Select **Links**.

   Result: The **Links** page opens.

3. To the left of the applicable link, select the icon.

   Result: A list of all the requested traces for the link shows.

Request a trace for a link

*The **Links** page lets you request a trace for a selected link.*

Procedure

1. In the top navigation bar, select icon > **Network**.

   Result: The **Network** page opens.

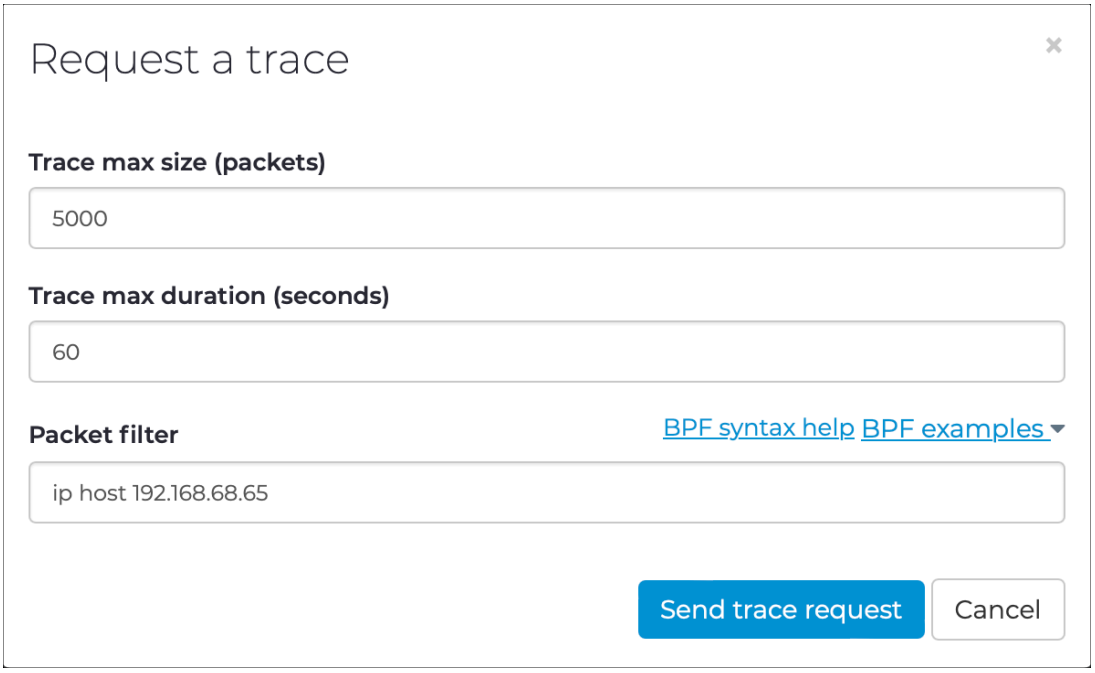
2. Select **Links**.

   Result: The **Links** page opens.

3. To the left of the applicable link, select the icon.

   Result: A dialog shows.

4. To set the maximum packet size, in the **Trace max size (packets)** field, enter a value.



> **Note:**
> The default size is 5000 packets.

5. To set the maximum duration of the trace, in the **Trace max duration (seconds)** field, enter a value.

> **Note:**
> The default value is 60 seconds.

6. **Note:**
   The **Packet filter** field is automatically populated with a *BPF* that captures the packets to/from the selected link, but you can customize this.

   If necessary, customize this field.

   Note:
   You can select **BPF syntax help** to show more information on *BPF* syntax.

   Note:
   You can select **BPF examples** to see some examples.

7. Select **Send trace request**.

Results

The trace has been requested.

Show events for a link

*The **Links** page lets you show requested traces for a selected link.*

Procedure

1. In the top navigation bar, select **icon > Network**.

   Result: The **Network** page opens.

2. Select **Links**.

   Result: The **Links** page opens.

3. To the left of the applicable link, select the icon.

   Result: A list of all the events for the link shows.

Show captured URLs for a link

*The **Links** page lets you show captured URLs for a selected link.*

Procedure

1. In the top navigation bar, select icon > **Network**.

   Result: The **Network** page opens.

2. Select **Links**.

   Result: The **Links** page opens.

3. To the left of the applicable link, select the icon.

   Result: A list of all the captured *uniform resource locator (URL)*s from the analyzed traffic for the link shows.

Manage learning for a link

*The **Links** page lets you manage the learning for a selected node.*

About this task

The **Manage Learning** dialog lets you learn and delete the entire node and its individual details, such as *IP* or *MAC* address.

Note:

The icon adjacent to the node shows its status. When the node details are:

- A green icon shows when the node details have been entirely learned
- An orange icon shows when the node details are only partially learned
- A red icon shows when the node details have not been learned

Individual details have either a green or red icon, depending on whether they are learned or not. When you learn, or delete a node, all of its details are affected in the same way. When you learn, or delete, an individual detail, only that detail's learning status changes.

Procedure

1. In the top navigation bar, select icon > **Network**.

   Result: The **Network** page opens.

2. Select **Links**.

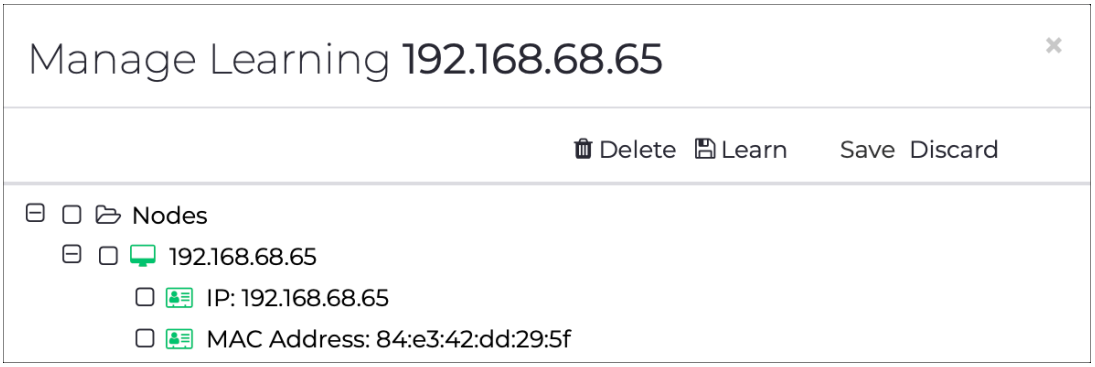
   Result: The **Links** page opens.

3. To the left of the applicable node, select the icon.

   Result: A dialog shows.

4. **Optional:**

   If necessary, select one, or more, items and select **Delete**.



5. **Optional:** If necessary, select one, or more, items and select **Learn**.
6. Select **Save**.

Navigate from a link

*The **Links** page lets you navigate to related entities from a selected node.*

Procedure

1. In the top navigation bar, select **icon > Network**.

   Result: The **Network** page opens.

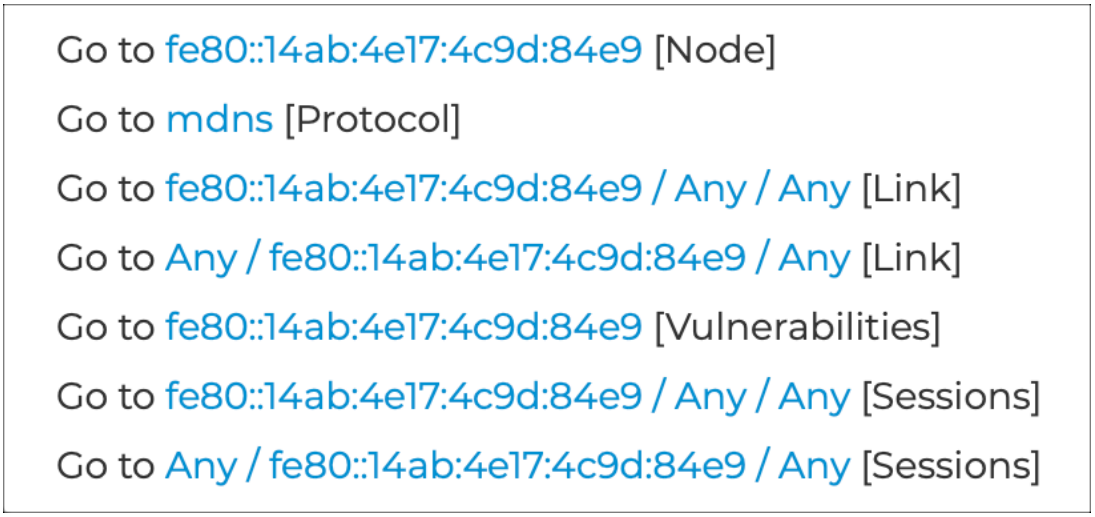
2. Select **Links**.

   Result: The **Links** page opens.

3. In the Actions column, to the left of the applicable node, select the icon.

   Result: A list of related entities shows.

4. Select the hyperlink that you want to navigate to.



Results

The entity shows in the applicable page.

Physical links

*The **Physical links** page lets you view all the physical links in your environment and perform actions on them. Smart Polling captures and displays the physical connections between the nodes in the physical graph. For switches, the graph also shows the physical port that maps each connection.*

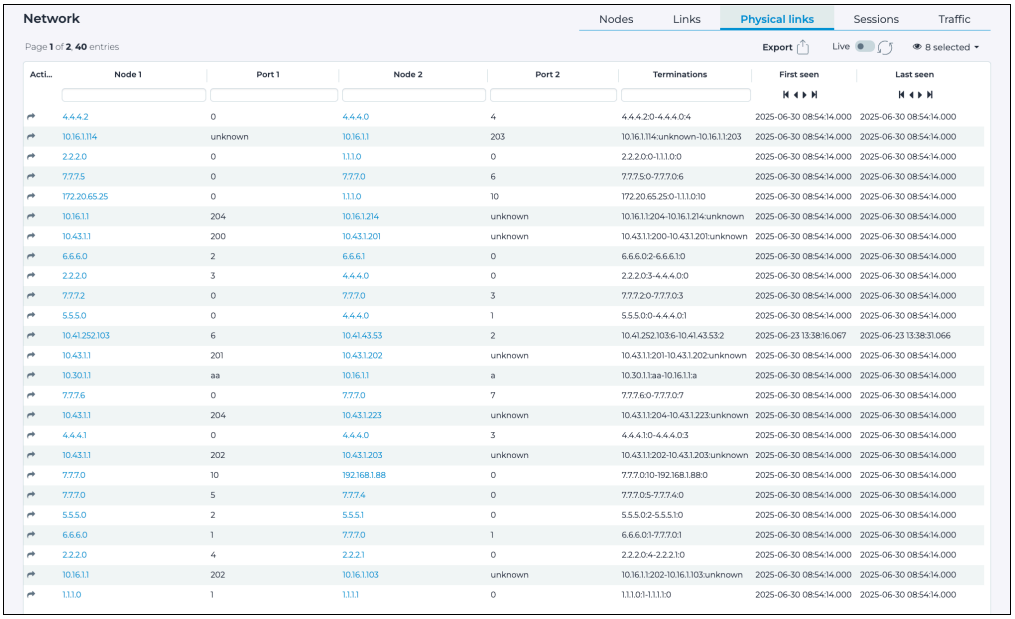


Figure 73. Physical links page

Export

The **Export** icon lets you export the current list in either *CSV* or Microsoft Excel format.

Live / refresh

The **Live** icon lets you change live view on, or off. When live mode is on, the page will refresh approximately every five seconds.

Column selection

The columns selection icon lets you choose which columns to show or hide.

Sessions

A session is established at a certain point in time, and later turned down. An established communication session might involve more than one message in each direction.

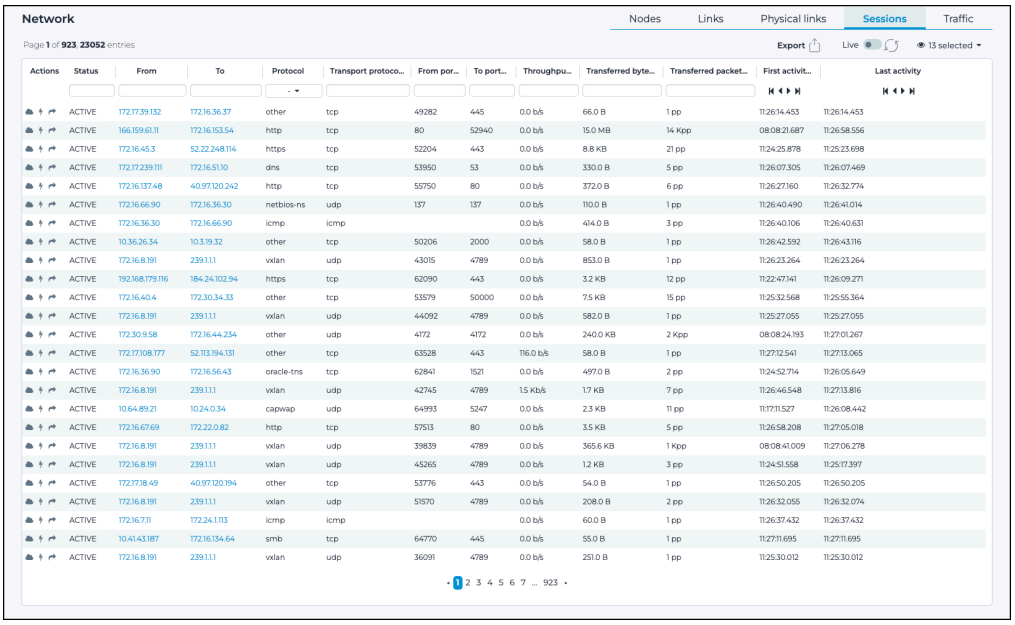


Figure 74. Sessions page

Export

The **Export** icon lets you export the current list in either *CSV* or Microsoft Excel format.

Live / refresh

The **Live** icon lets you change live view on, or off. When live mode is on, the page will refresh approximately every five seconds.

Column selection

The columns selection icon lets you choose which columns to show or hide.

Show requested traces for a session

*The **Sessions** page lets you show requested traces for a selected node.*

Procedure

1. In the top navigation bar, select icon > **Network**.

   Result: The **Network** page opens.

2. Select **Sessions**.

   Result: The **Sessions** page opens.

3. To the left of the applicable node, select the icon.

   Result: A list of all the requested traces for the session shows.

Request a trace for a session

*The **Sessions** page lets you request a trace for a selected node.*

Procedure

1. In the top navigation bar, select icon > **Network**.

   Result: The **Network** page opens.

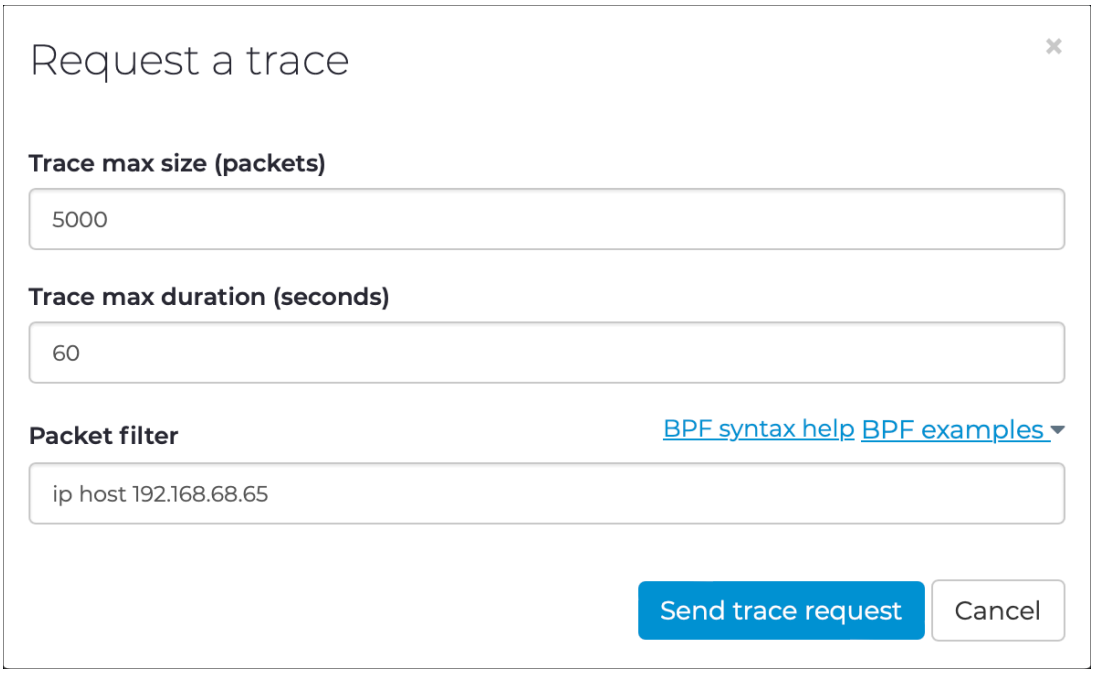
2. Select **Sessions**.

   Result: The **Sessions** page opens.

3. To the left of the applicable node, select the icon.

   Result: A dialog shows.

4. To set the maximum packet size, in the **Trace max size (packets)** field, enter a value.



> **Note:**
> The default size is 5000 packets.

5. To set the maximum duration of the trace, in the **Trace max duration (seconds)** field, enter a value.

> **Note:**
> The default value is 60 seconds.

6. **Note:**
   The **Packet filter** field is automatically populated with a *BPF* that captures the packets to/from the selected node, but you can customize this.

   If necessary, customize this field.

   Note:
   You can select **BPF syntax help** to show more information on *BPF* syntax.

   Note:
   You can select **BPF examples** to see some examples.

7. Select **Send trace request**.

Results

The trace has been requested.

Navigate from a session

*The **Sessions** page lets you navigate to related entities from a selected node.*

Procedure

1. In the top navigation bar, select icon > **Network**.

   Result: The **Network** page opens.

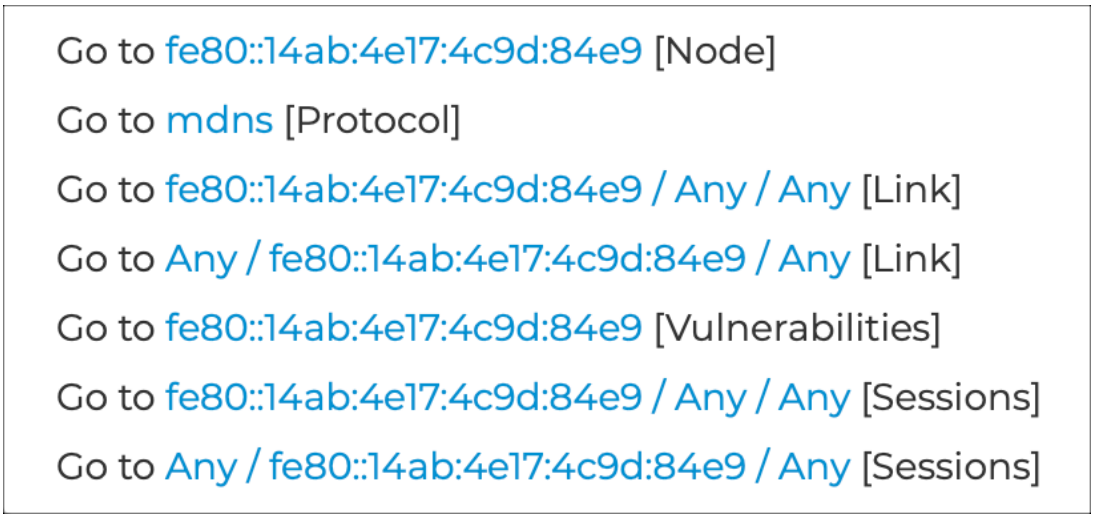
2. Select **Sessions**.

   Result: The **Sessions** page opens.

3. In the Actions column, to the left of the applicable node, select the icon.

   Result: A list of related entities shows.

4. Select the hyperlink that you want to navigate to.

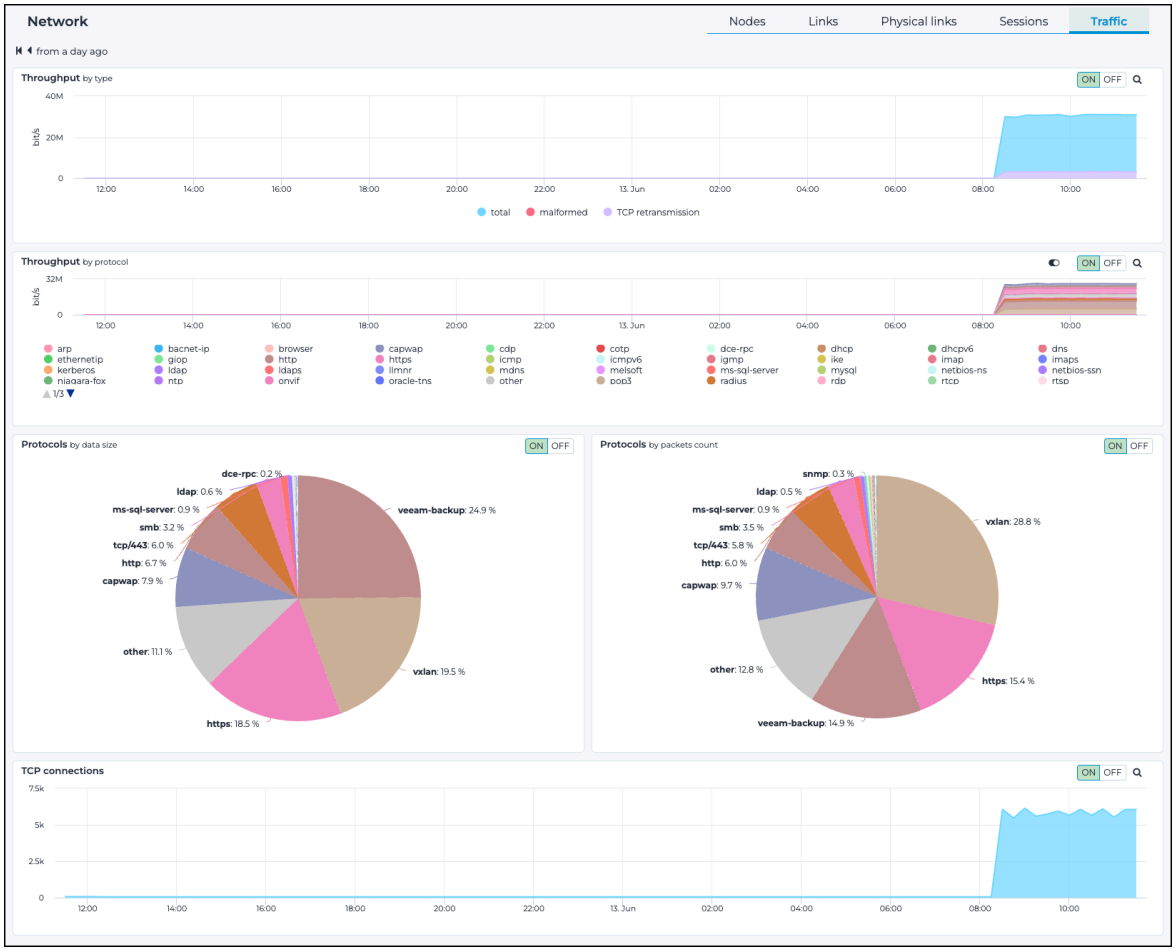


Results

The entity shows in the applicable page.

Traffic

*The **Traffic** page shows charts with information about throughput, protocols, and open transmission control protocol (TCP) connections.*



Throughput by type

This section shows traffic by macro category.

Throughput by protocol

This section shows traffic by protocol.

Protocols by data size

This section shows the proportion of packets sent by protocol, in pie chart format.

Protocols by packets count

This section shows the proportion of traffic generated by protocol, in pie chart format.

TCP connections

This section shows the number of open *TCP* connections.

Details page

The details page lets you view more detailed information for a network.

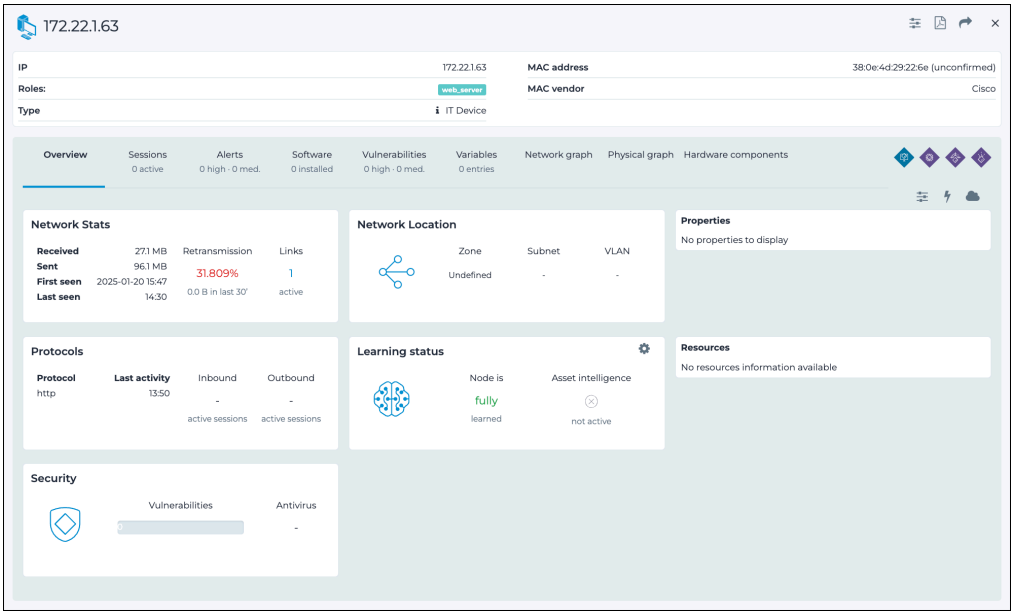


Figure 75. Network details page

This details page shows more details for a network.

The top section of the screen contains generic data. You can hover your mouse over the information icon to display the source, granularity and confidence of the corresponding piece of data. Data includes:

- *IP* (address)
- Roles
- Vendor
- *MAC* address
- *MAC* vendor

The details window has these tabs:

- Overview (on page 76)
- Sessions (on page 77)
- Alerts (on page 78)
- Software (on page 79)
- Vulnerabilities (on page 80)
- Variables (on page 81)
- Network graph (on page 82)
- Physical graph (on page 83)
- Hardware components (on page 84)



Chapter 9. Graph



*The **Graph** pages gives a graphical representation of the nodes in the environment.*

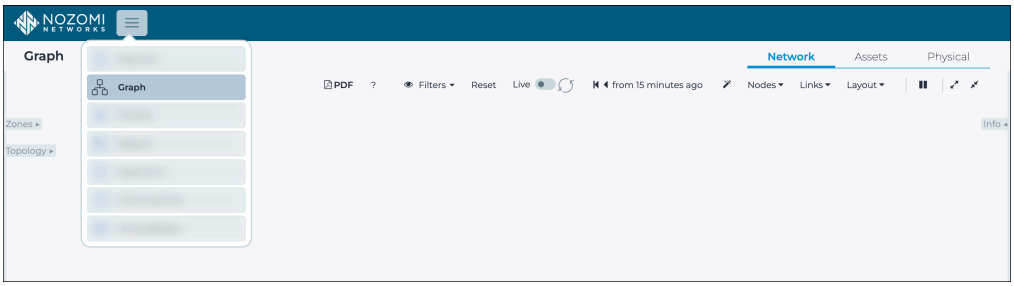


Figure 76. Graph page

The **Graph** page has these tabs:

- Network (on page 222)
- Assets (on page 225)
- Physical (on page 229)

Network

*The **Network** page displays a live graph of connected assets, links, and filters in your network. You can apply filters, adjust the layout, and select node and link perspectives. Interactive options let you customize the visualization and export reports.*

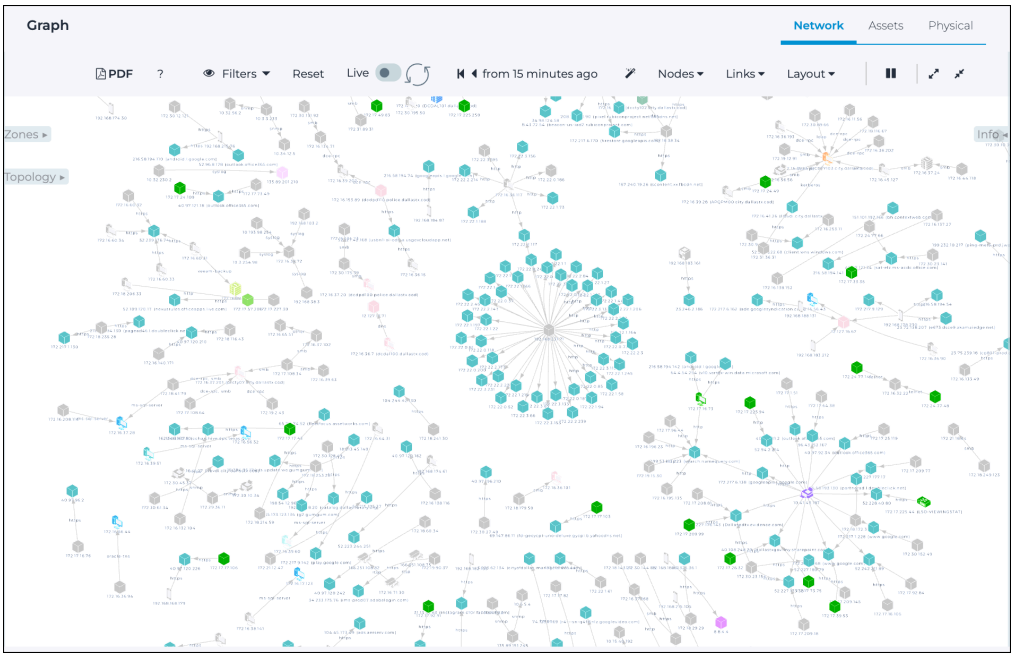


Figure 77. Network page

Generate PDF report

The generate PDF icon lets you create a *PDF* report of the selected item, which will show in the **Reports** page when it is ready.

Filters

This indicates active graph filtering, when present. Filters can be from the filter bar (see R and S below), or activated from the zone/topology graph when you select a link/node in the zone/topology graphs. Once a filter is enabled with a value, the graph is automatically updated. If more than one filter is enabled, then a logical and criteria is applied. Only nodes that satisfy all of the specified filters are shown.

Note:

If a node passes the filters, then all of the directly connected nodes are shown in the graph. For example if a specific *IP* filter is used, then the specified node is shown along with all the nodes connected to it.

Reset

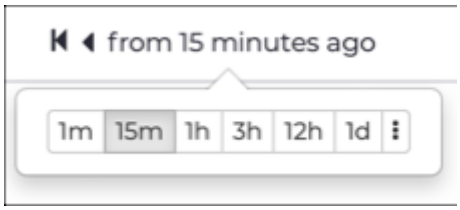
This resets customizations and reloads the data.

Live / refresh

The **Live** icon lets you change live view on, or off. When live mode is on, the page will refresh approximately every five seconds.

Time

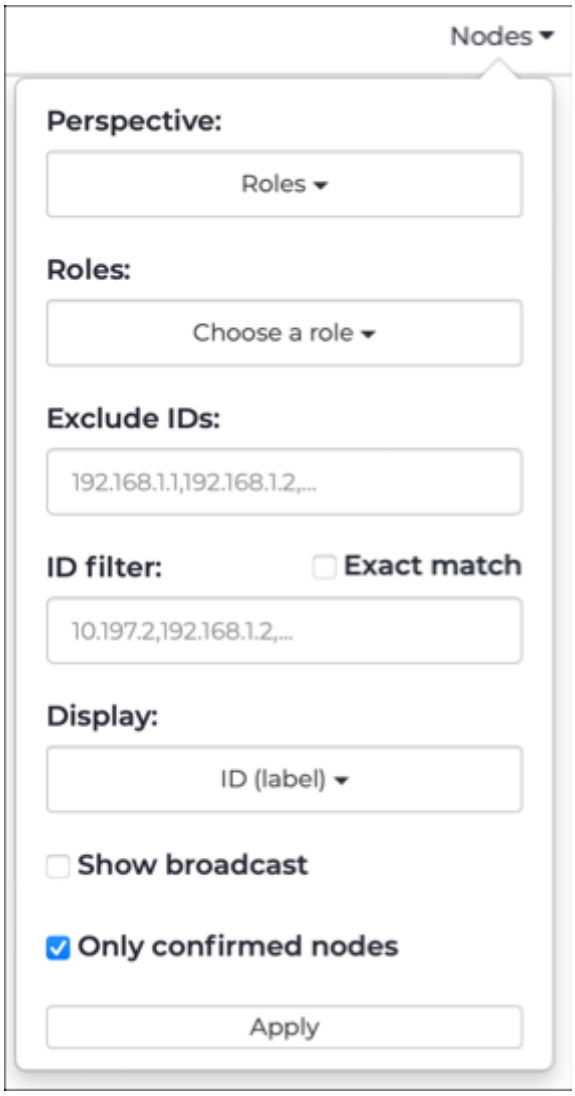
These icons let you select an activity time range.



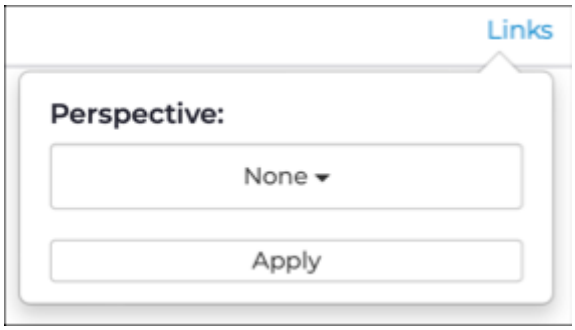
For more details, Magic wand (on page 235).

Nodes

This dropdown lets you select node visualization configuration options.



Links



This dropdown lets you configure visualization options.

Layout

This dropdown lets you select a layout for the graph. For more details, see Layout (on page 233).

Play-pause

The play-pause icon lets you pause, or restart the motion of the graph.

Increase-Decrease icon size

The increase and decrease icons lets you change the size of the icons in the graph.

Assets

*The **Assets** page shows linked assets as graph vertices when available, and defaults to node display when no asset is associated. This view provides a contextual representation of asset relationships within the network topology.*

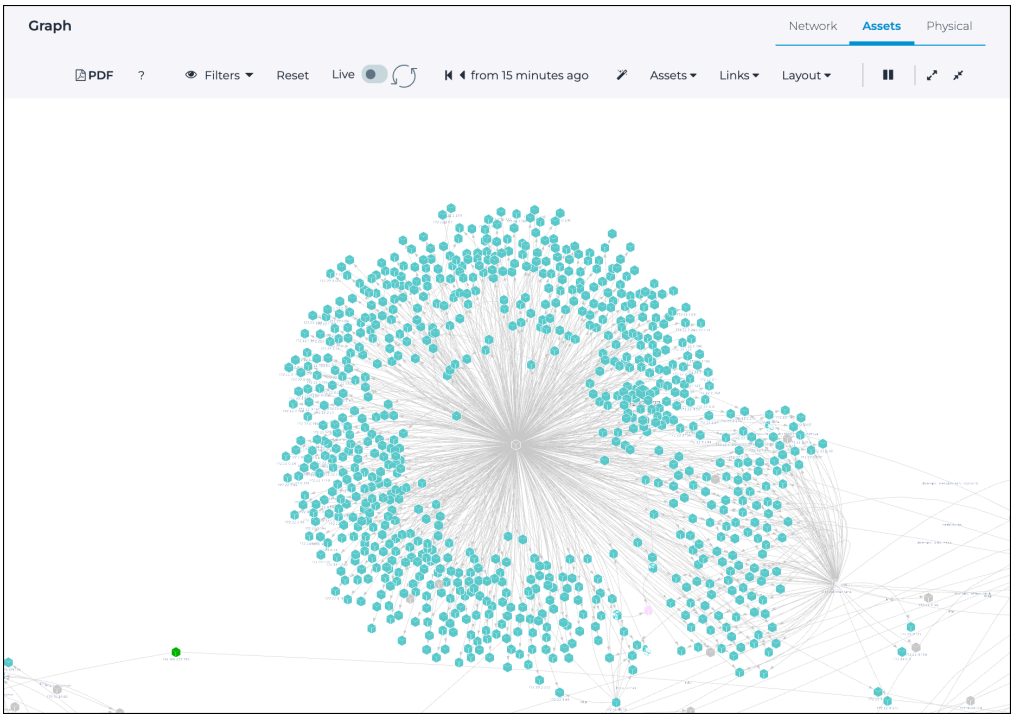


Figure 78. Assets page

Overview

The asset graph groups nodes by asset when possible, unlike the standard network graph. For example, if an asset includes two nodes, the network graph displays two separate icons, one for each node, while the asset graph displays a single icon that represents the asset.

In the case that a node does not belong to an asset then the asset graph shows the node itself.

Generate PDF report

The generate PDF icon lets you create a *PDF* report of the selected item, which will show in the **Reports** page when it is ready.

Filters

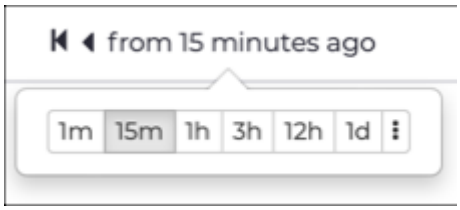
This indicates active graph filtering, when present. Filters can be from the filter bar (see R and S below), or activated from the zone/topology graph when you select a link/node in the zone/topology graphs. Once a filter is enabled with a value, the graph is automatically updated. If more than one filter is enabled, then a logical and criteria is applied. Only nodes that satisfy all of the specified filters are shown.

Note:
If a node passes the filters, then all of the directly connected nodes are shown in the graph. For example if a specific *IP* filter is used, then the specified node is shown along with all the nodes connected to it.

Reset
This resets customizations and reloads the data.

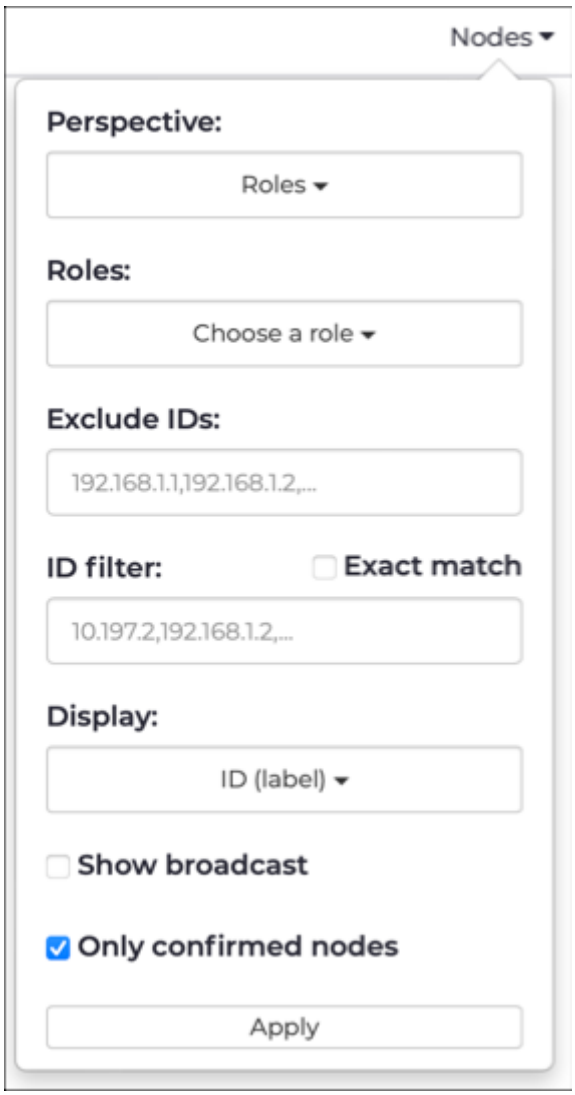
Live / refresh
The **Live** icon lets you change live view on, or off. When live mode is on, the page will refresh approximately every five seconds.

Time
These icons let you select an activity time range.

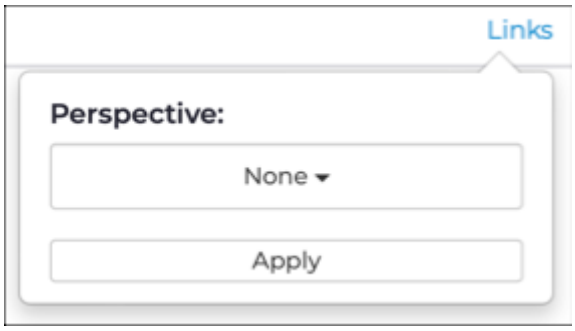


For more details, Magic wand (on page 235).

Nodes
This dropdown lets you select node visualization configuration options.



Links



This dropdown lets you configure visualization options.

Layout

This dropdown lets you select a layout for the graph. For more details, see Layout (on page 233).

Play-pause

The play-pause icon lets you pause, or restart the motion of the graph.

Increase-Decrease icon size

The increase and decrease icons lets you change the size of the icons in the graph.

Physical

*The **Physical** page displays a real-time interactive map of all physical device connections in the network. Nodes are shown only if a physical link has been detected for them. The graph, shows how devices connect at the cable level. It is a live, visual way to understand your physical network layout down to the port.*

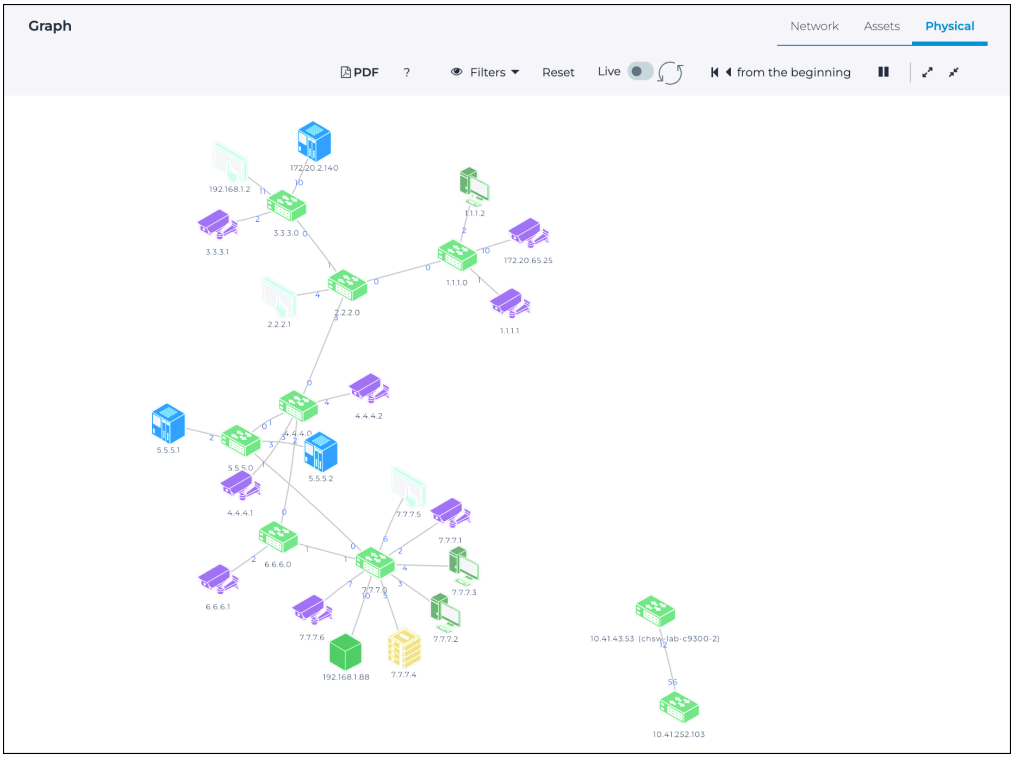


Figure 79. Physical page

Overview

The Physical graph view provides a visual representation of cable-level connections between devices in the network. It uses data collected through Smart Polling and system integrations to display physical links in a dynamic topology graph that reflects actual hardware wiring, including switch port details.

This view supplements logical graphs by showing real-time or historical physical connectivity between assets. Devices are positioned based on link data, with visual indicators that show port assignments and device types. Use the Physical graph to examine network structure, trace physical paths, and identify connectivity changes.

Generate PDF report

The generate PDF icon lets you create a *PDF* report of the selected item, which will show in the **Reports** page when it is ready.

Filters

This indicates active graph filtering, when present. Filters can be from the filter bar (see R and S below), or activated from the zone/topology graph when you select a link/node in the zone/topology graphs. Once a filter is enabled with a value, the graph is automatically updated. If more than one filter is enabled, then a logical and criteria is applied. Only nodes that satisfy all of the specified filters are shown.

> **Note:**
> If a node passes the filters, then all of the directly connected nodes are shown in the graph. For example if a specific *IP* filter is used, then the specified node is shown along with all the nodes connected to it.

Reset

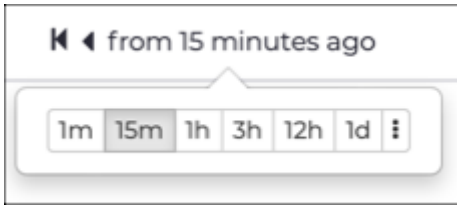
This resets customizations and reloads the data.

Live / refresh

The **Live** icon lets you change live view on, or off. When live mode is on, the page will refresh approximately every five seconds.

Time

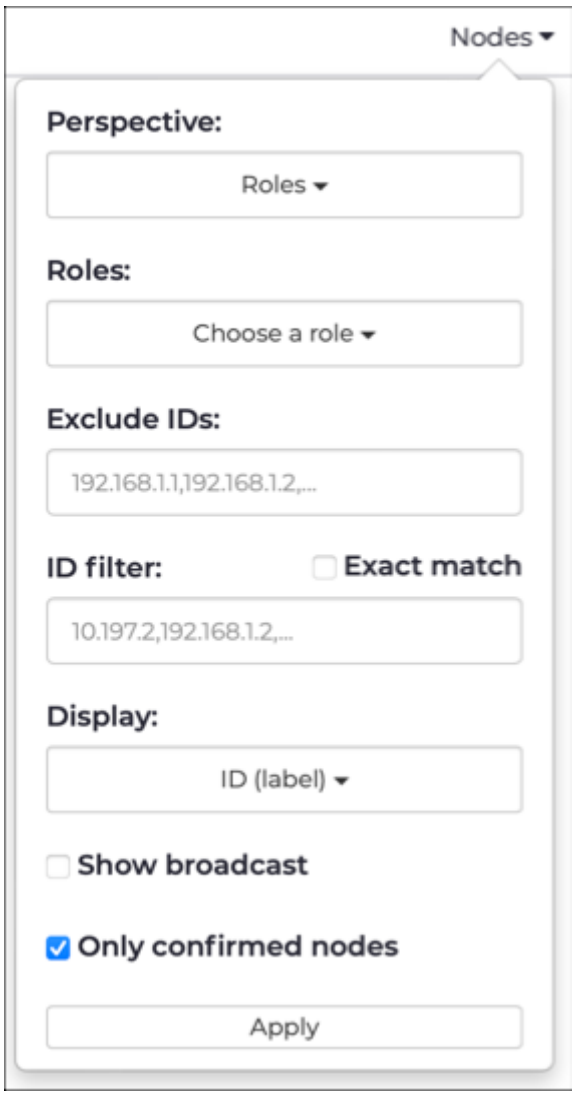
These icons let you select an activity time range.



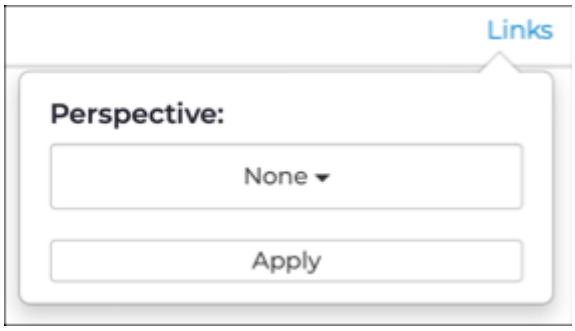
For more details, Magic wand (on page 235).

Nodes

This dropdown lets you select node visualization configuration options.



Links



This dropdown lets you configure visualization options.

Layout

This dropdown lets you select a layout for the graph. For more details, see Layout (on page 233).

Play-pause

The play-pause  icon lets you pause, or restart the motion of the graph.

Increase-Decrease icon size

The increase  and decrease  icons lets you change the size of the icons in the graph.

Graph controls

A list of the different control options in the graph view.

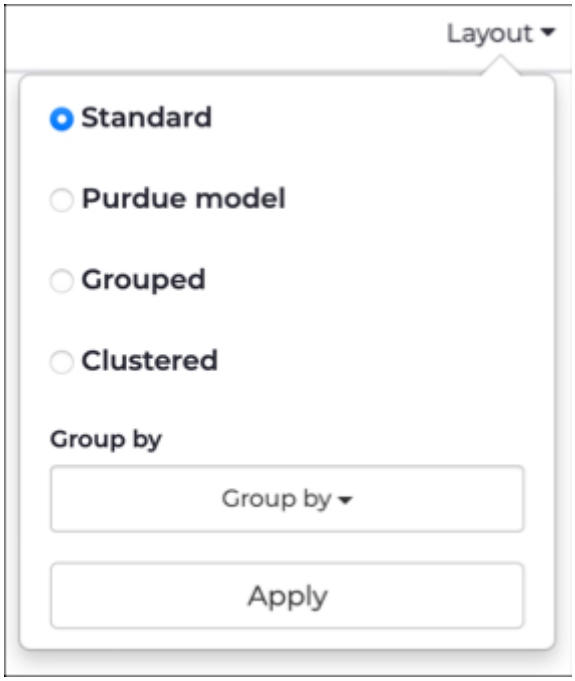
Move	Click and drag anywhere in the graph other than on a node.
Zoom (mode 1)	With the cursor positioned inside the graph window, scroll the mouse forward or backward. Zoom is centered on the mouse position.
Zoom (mode 2)	With the cursor positioned inside the graph window, press Z on the keyboard and move the mouse up or down. Zoom is centered on the mouse position.
Increase icon and text size.	Select the  icon.
Decrease icon and text size.	Select the  icon.
View detailed information for a node or link in the information pane.	Select a node or link with a single click
Show a new window with additional information for a node	Select a node with a double click
Show a node or link	Move your mouse over the node or link
Show a node or link, and the elements directly connected to it	Click and hold your mouse button on a node or link without releasing it

Layout

*The **Layout** dropdown lets you select a layout for the graph.*

The options are:

- Standard (on page 234)
- Purdue model (on page 234)
- Grouped (on page 234)
- Clustered (on page 235)



Standard

This is the default layout. The type of visualization depends on the criteria defined in Group by (on page 235):

- Group by (on page 235) not defined: All of the nodes and links are shown
- Group by (on page 235) defined: Nodes that belong to the same group (based on the defined criteria) are collapsed into a single node

Purdue model

Nodes are arranged in separate rows, according to their level. You can distinguish the levels and isolate potential communication problems that cross two or more levels.

Grouped

Nodes are grouped according to the criteria defined in Group_by. The graph is visualized as follows:

- Group by (on page 235) not defined: All nodes and links are shown.
- Group by (on page 235) defined: Nodes that belong to the same group are shown and are placed inside a circle that represents the group. Links between nodes within the same group are shown. However, links between groups are replaced with lines that connect the circles

Clustered

Nodes are clustered according to the criteria below. Once nodes are clustered, a single circle represents the node cluster. Upon zoom-in, the circle expands and the internal nodes display. A cluster may contain multiple subclusters. This layout is useful when visualizing large graphs because it provides an overview of the graph, along with sufficient details.

Nodes are clustered depending on the values defined by Group_by:

- Group by (on page 235) not defined: Nodes are clustered based on connections. Nodes with a large number of links act as a cluster center with neighboring nodes assigned to the same cluster
- Group by (on page 235) defined: At the highest level, a cluster is created for each group. Inside each high level cluster are subclusters created around nodes with a high number of links. For example, if Group_by=Zones, then a cluster is created for each zone, and inside each zone other subclusters may be created around nodes with a high number of links

Group by

This dropdown lets you define the group used for Standard, Grouped, and Clustered layouts. Nodes with the chosen property, such as zone or subnet, are assigned to the same group. The group displays depending on the selected layout.

The **Group by** dropdown lets you select from:

- Asset
- Cluster
- Level
- Roles
- Subnet
- Type
- Zone
- Site
- Host

Magic wand

The magic wand icon opens a wizard to help filter the graph and view only the desired information to help you reduce visualized data from large graphs.

The graph wizard gives hints to help you improve the graph performance. Settings that are annotated with an orange exclamation point are considered suboptimal. Settings annotated with green thumbs are considered helpful.

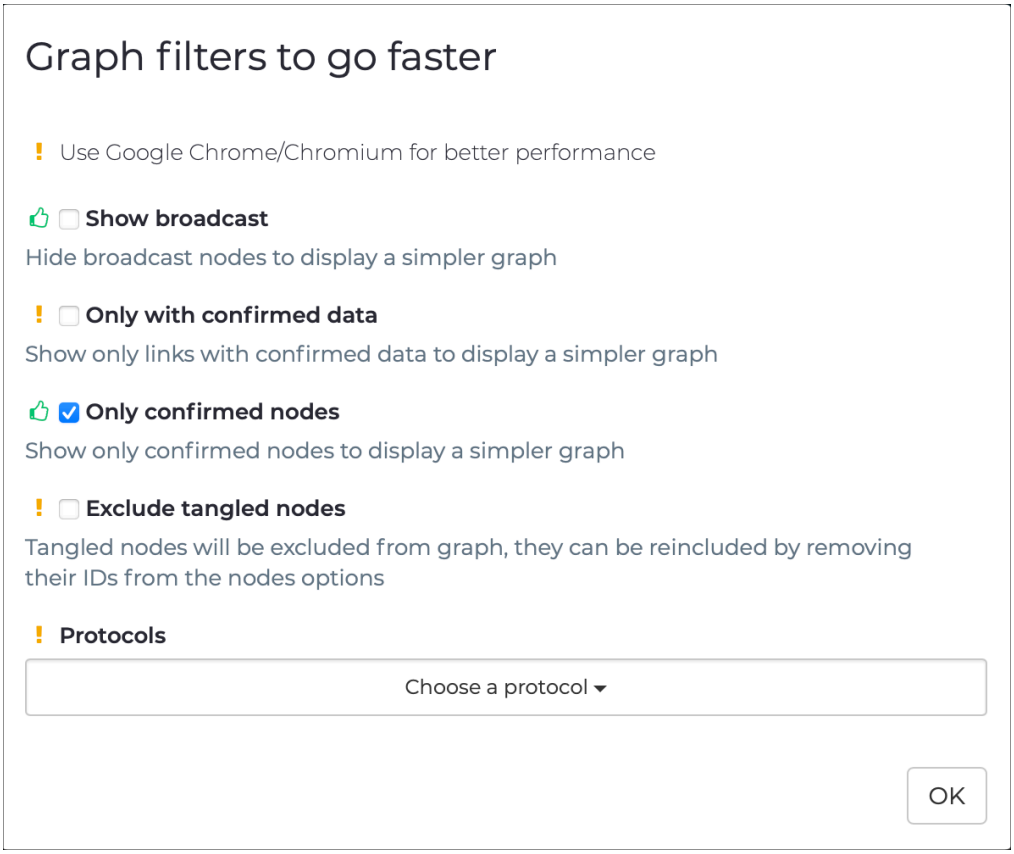


Figure 80. Magic wand dialog

Show broadcast

Broadcast addresses are not actual network nodes in that no asset is bound to a broadcast address. They are used to represent communications that a node uses towards an entire subnet. Removing broadcast nodes reduces the complexity of a graph.

Only with confirmed data

You can select this to hide unconfirmed links to reduce the complexity of an entangled graph.

Only confirmed nodes

You can select this to hide unconfirmed nodes to reduce the size of a large graph.

Exclude tangled nodes

To improve the readability of the graph, you can select this to exclude nodes whose connections cause the graph to be too complex.

Protocols

This dropdown lets you filter nodes and edges to show only those items that are communicating with the related protocol(s).

Chapter 10. Process



Process is a set of repeatable functions that a business does to deliver a core value.

Process includes:

- Repeatable tasks
- Data collection
- Resource control in accordance with business policies

Variables model communication between operational devices as they participate in the industrial process.

Individual values within operational devices are represented as variables, and Guardian tracks them over time in **Process**.

The **Process** page has these tabs:

- List (on page 240)
- Protocol connections (on page 241)
- Settings (on page 242)

List

*The **List** page shows detailed information about variables in your environment.*

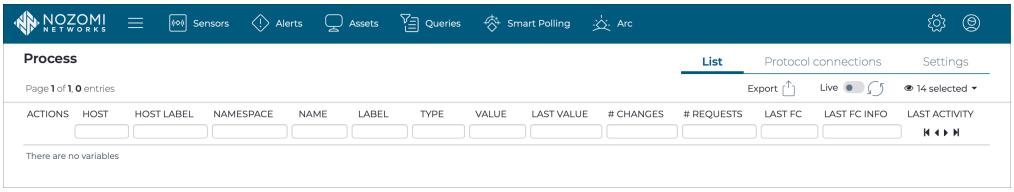


Figure 81. List page

Export

The **Export** icon lets you export the current list in either *CSV* or Microsoft Excel format.

Live / refresh

The **Live** icon lets you change live view on, or off. When live mode is on, the page will refresh approximately every five seconds.

Column selection

The columns selection icon lets you choose which columns to show or hide.

Protocol connections

*The **Protocol connections** page lets you configure the additional parameters necessary to extract information contained inside certain protocols.*

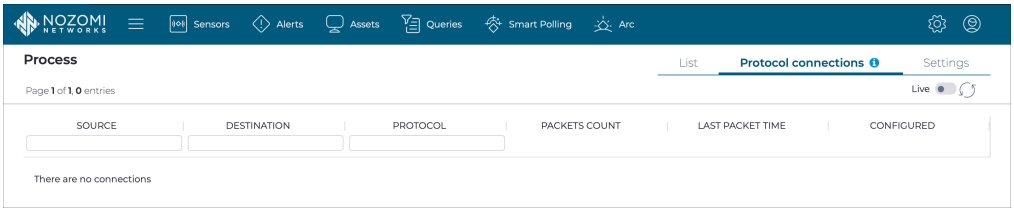


Figure 82. Protocol connections page

Live / refresh

The **Live** icon lets you change live view on, or off. When live mode is on, the page will refresh approximately every five seconds.

Settings

*The **Settings** page lets you configure which strategy to use to extract variables inside the process that are being communicated over different protocols.*

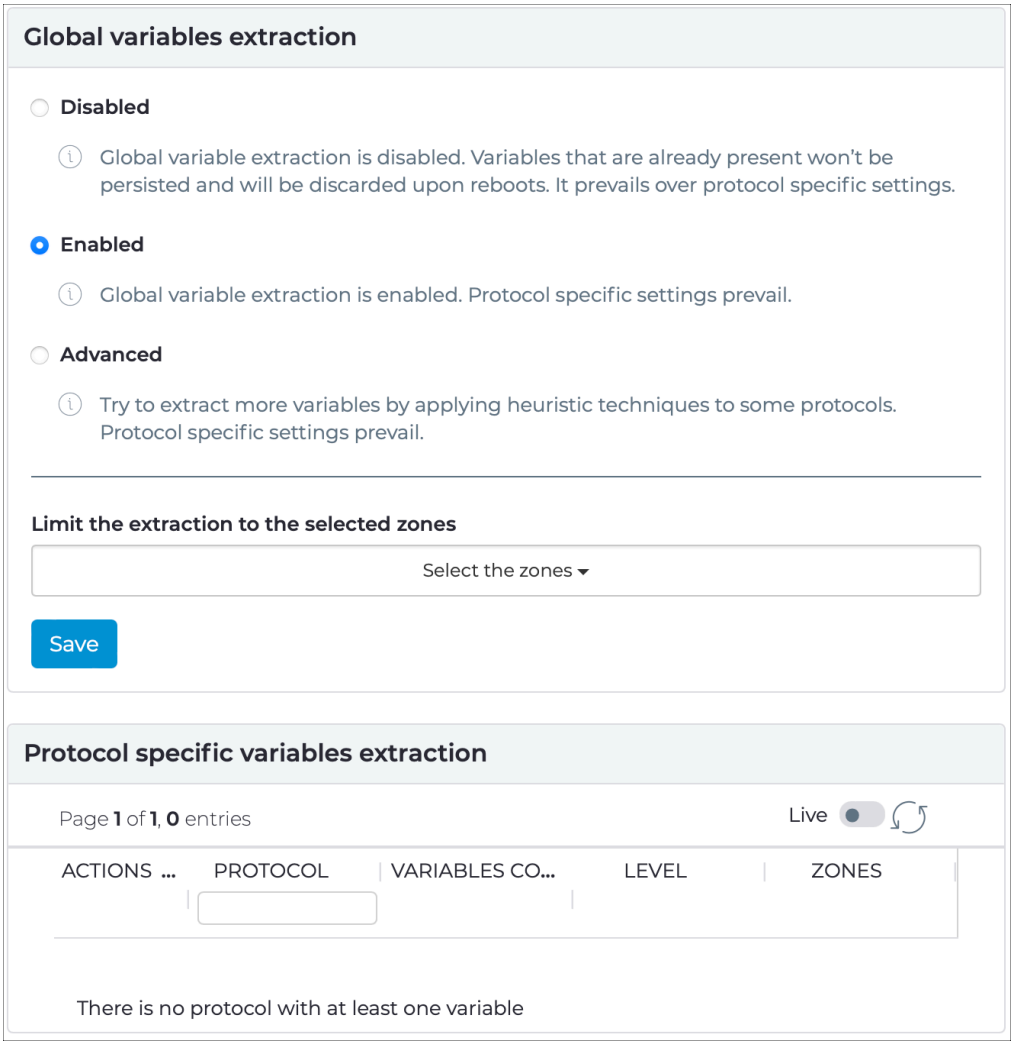


Figure 83. Process variable extraction tuning

Disabled

This checkbox lets you disable global variables extraction.

Enabled

This checkbox lets you enable global variables extraction.

Advanced

This checkbox lets you try to apply heuristic techniques to some protocols to extract more variables. Protocol specific settings prevail.

Limit the extraction to the selected zones

This dropdown lets you Limit the extraction to the selected zones.

Configure a variable

*The **List** page lets you navigate to related nodes, links, vulnerabilities, or sessions.*

Procedure

1. In the top navigation bar, select icon > **Process**.

   Result: The **Process** page opens.

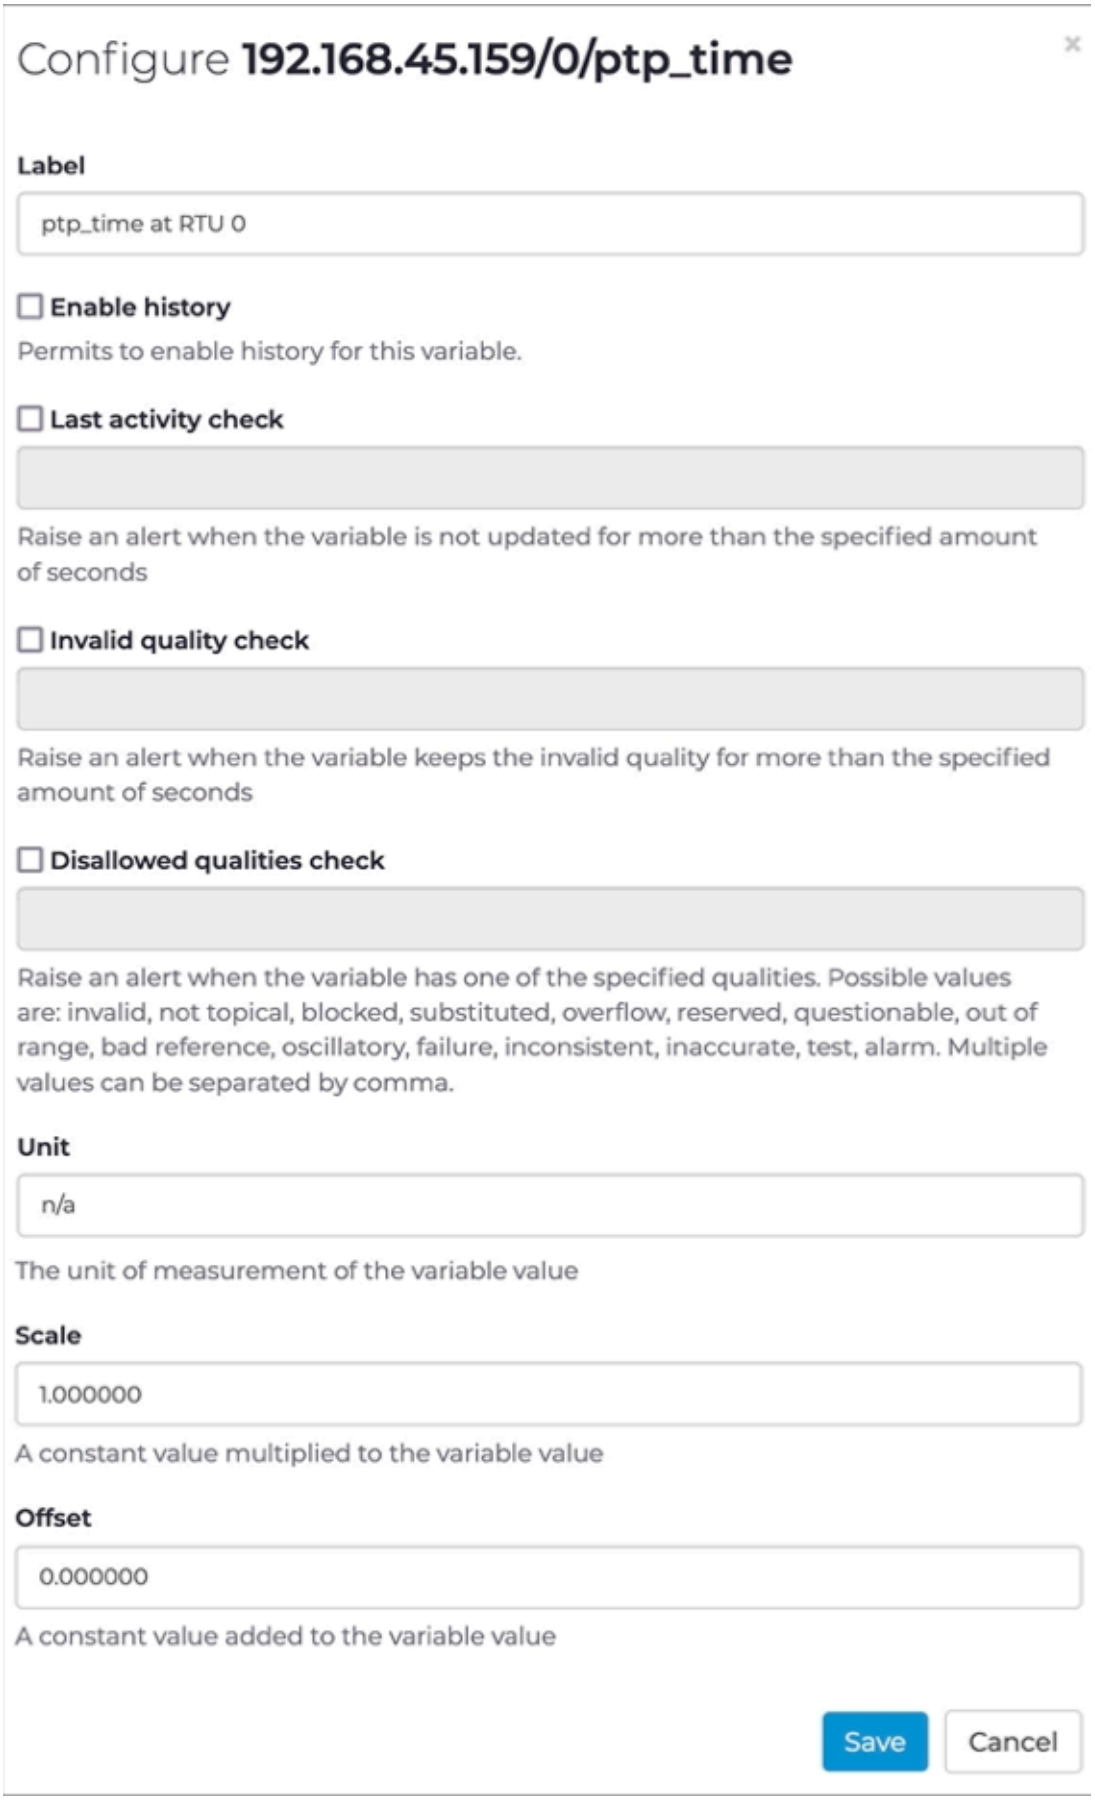
2. Select **List**.

   Result: The **List** page opens.

3. To the left of the applicable *variable*, select the configure icon.

   Result: A dialog shows.

4. In the **Label** field, enter a label for the *Variable*.



5. **Optional:** If necessary, select the **Enable history**.

> **Note:**
> This lets you enable *Variable* history.

6. **Optional:** If necessary, select the **Last activity check**.

> **Note:**
> This triggers an alert when the *Variable* is not updated for more than the specified amount of seconds.

7. **Optional:** If necessary, select the **Invalid quality check**.

> **Note:**
> This triggers an alert when the *Variable* keeps the invalid quality for more than the specified amount of seconds.

8. **Optional:** If necessary, select the **Disallowed qualities check**.

> **Note:**
> This triggers an alert when the *Variable* has one of the listed qualities. You can use a comma to separate values.

9. In the **Unit** field, enter a value.
10. In the **Scale** field, enter a value.
11. In the **Offset** field, enter a value.

Results

The *Variable* has been configured.

View the details of a variable

*The **List** page lets you view the details of a variable.*

Procedure

1. In the top navigation bar, select icon > **Process**.

   Result: The **Process** page opens.

2. Select **List**.

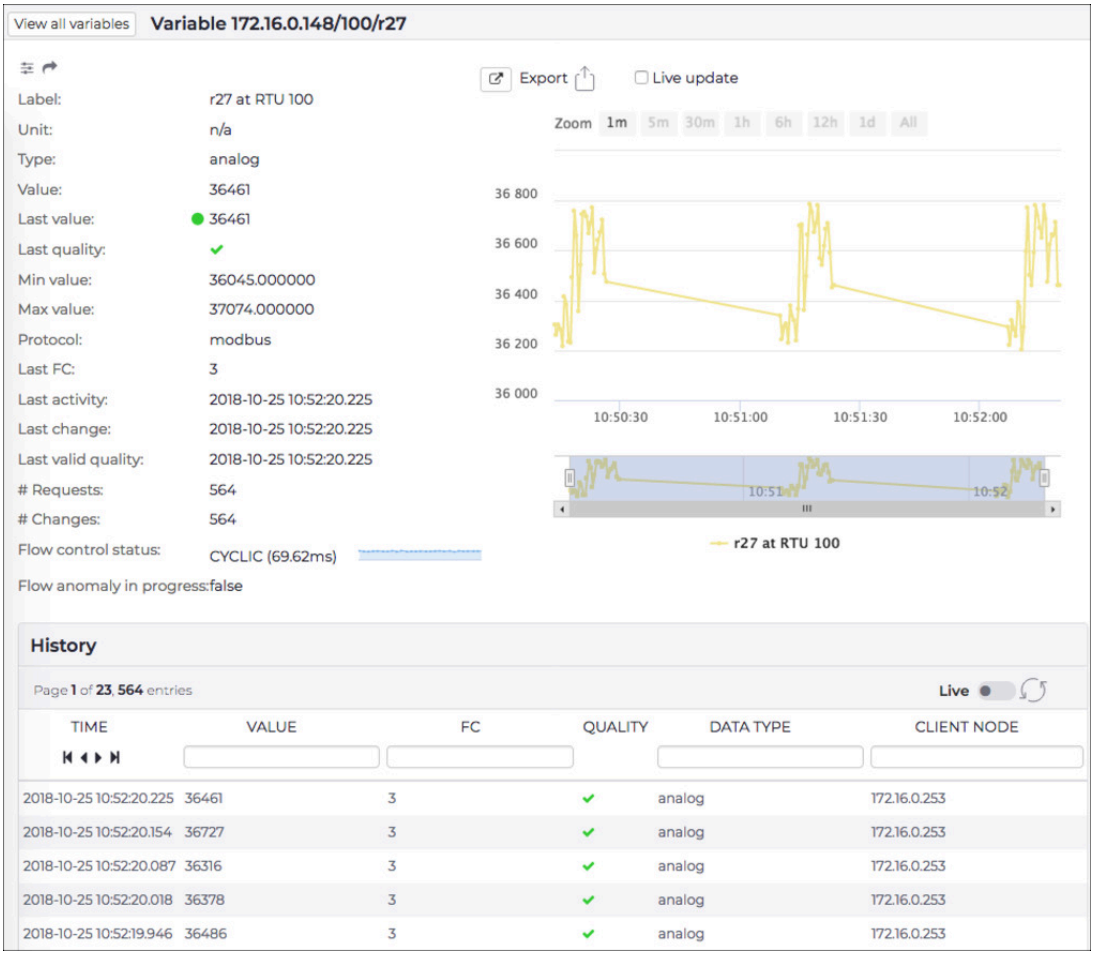
   Result: The **List** page opens.

3. To the left of the applicable *Variable*, select the search icon.

   Result: The Variables details window opens.

4. **Optional:**

   To open the chart in another window, select the icon.



5. **Optional:** To export the results, select **Export**
6. **Optional:** To update the chart to real-time, select **Live update**.

Favorite a variable

*The **List** page lets you add a Variable to the favorite variables.*

Procedure

1. In the top navigation bar, select icon > **Process**.

   Result: The **Process** page opens.

2. Select **List**.

   Result: The **List** page opens.

3. To the left of the applicable *Variable*, select the search icon.

   Result: The *Variable* is added to the Favorite variables list.

Navigate from a variable

*The **List** page lets you navigate to related nodes, links, vulnerabilities, or sessions.*

Procedure

1. In the top navigation bar, select icon > **Process**.

   Result: The **Process** page opens.

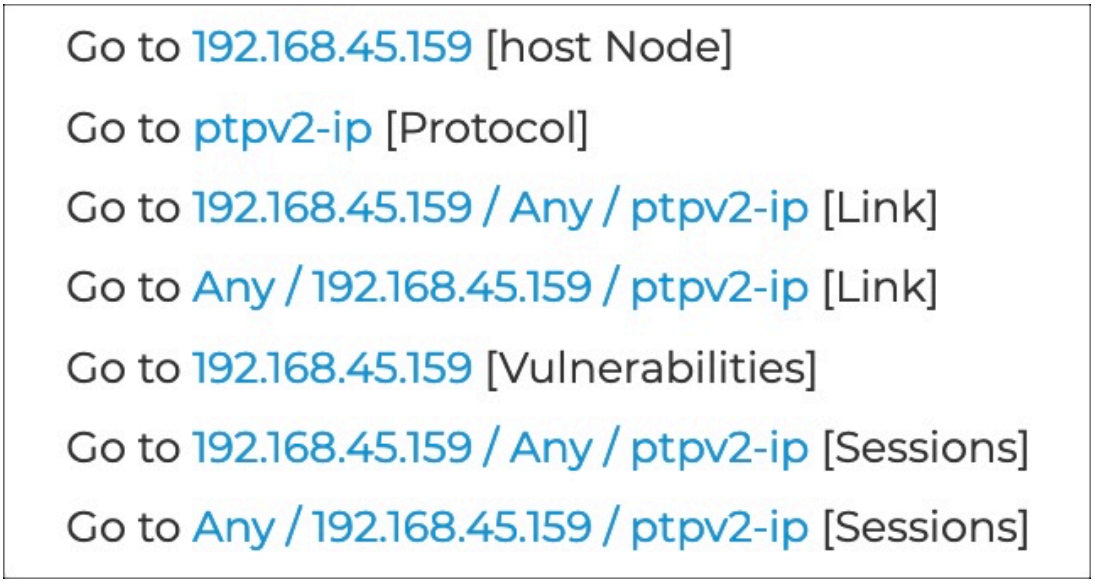
2. Select **List**.

   Result: The **List** page opens.

3. To the left of the applicable variable, select the navigate icon.

   Result: A dialog shows.

4. Select the desired link.





Chapter 11. Reports



*The **Reports** page lets you manage, generate, schedule and view reports.*

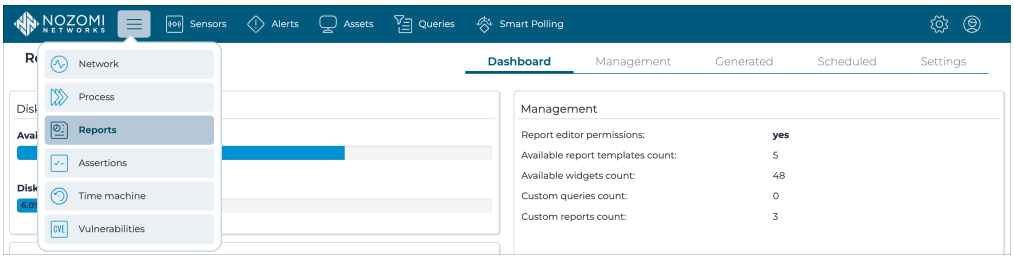


Figure 84. Reports page

The **Reports** page has these tabs:

- Dashboard (on page 254)
- Management (on page 255)
- Generated (on page 258)
- Scheduled (on page 259)
- Settings (on page 260)

Dashboard

The Dashboard page shows an overview of information related to reports, which includes disk availability, report settings, generated reports, report management, and scheduled reports.

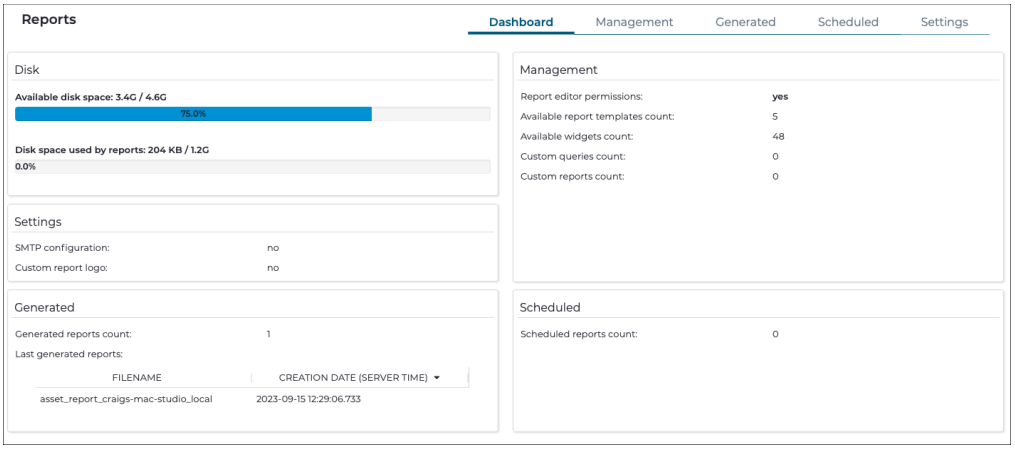


Figure 85. Dashboard page

Disk

This section shows:

- Available disk space
- Disk space used by reports

Management

This section shows a summary of information from the Management (on page 255) page.

Settings

This section shows a summary of information from the Settings (on page 260) page.

Generated

This section shows a summary of information from the Generated (on page 258) page.

Scheduled

This section shows a summary of information from the Scheduled (on page 259) page.

Management

*The **Management** page lets you manage all your reports.*

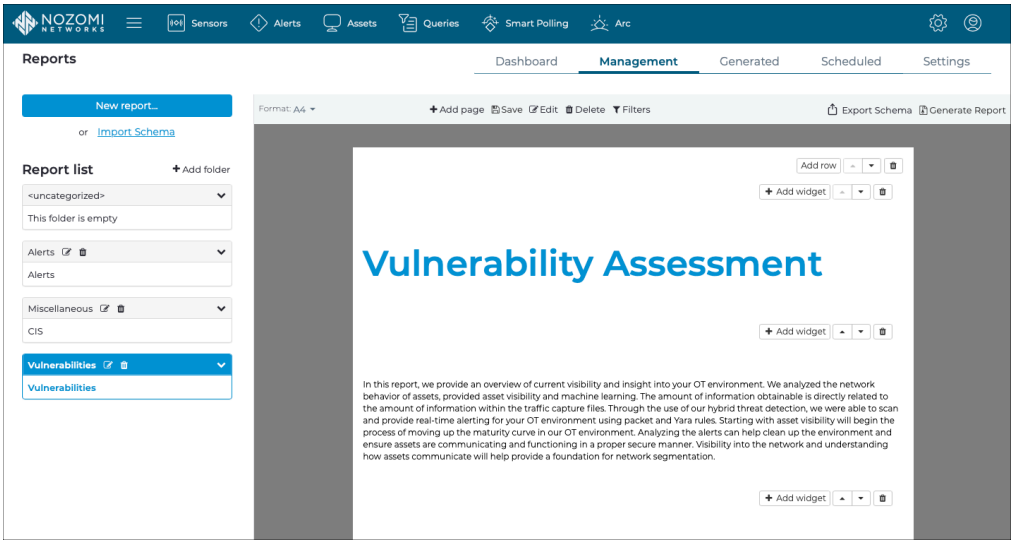


Figure 86. Management page

New report

This button lets you Create a report (on page 261).

Import schema

This button lets you Import a schema (on page 268).

Report list

This section shows a list of created and saved reports. You can add folders (on page 266) to group the reports in.

Add page

This button lets you add a page to the bottom of the current report.

Save

This button lets you save the changes to the current report.

Edit

This button lets you edit the current report.

Delete

This button lets you delete the current report.

Filters

This button lets you filter the contents of the current report.

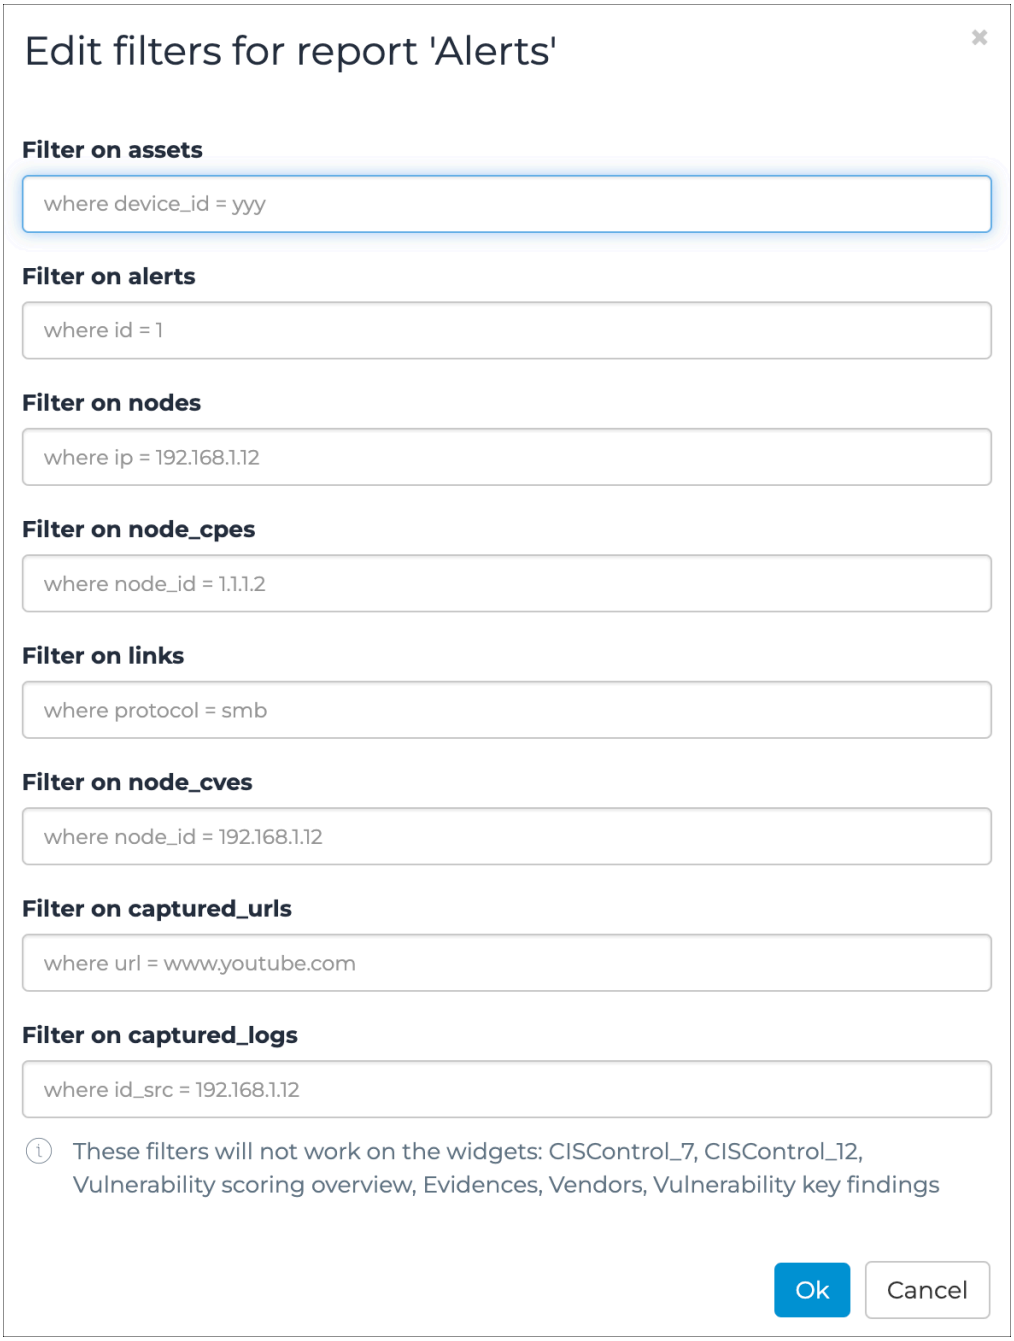


Figure 87. Filter dialog

Export schema

This button lets you export a schema (on page 269) for the current report in *JavaScript Object Notation (JSON)* format.

Generate report

The generate report  icon lets you generate a report (on page 263) in one of these formats:

- *PDF*
- *CSV*
- Microsoft Excel

Generated

*The **Generated** page lets you view, download, edit, and delete generated reports.*

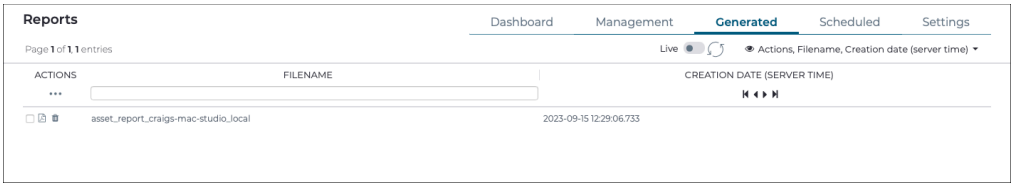


Figure 88. Generated page

Live / refresh

The **Live** icon lets you change live view on, or off. When live mode is on, the page will refresh approximately every five seconds.

Column selection

The columns selection icon lets you choose which columns to show or hide.

Scheduled

*The **Scheduled** page lets you view, download, edit, and delete scheduled reports.*

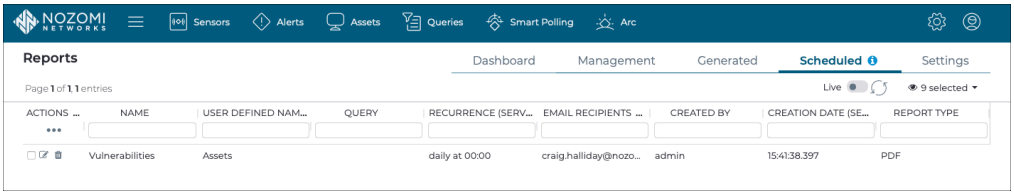


Figure 89. Scheduled page

Live / refresh

The **Live** icon lets you change live view on, or off. When live mode is on, the page will refresh approximately every five seconds.

Column selection

The columns selection icon lets you choose which columns to show or hide.

Settings

*The **Settings** page lets you change report settings, upload custom logos, and configure simple mail transfer protocol (SMTP) settings.*

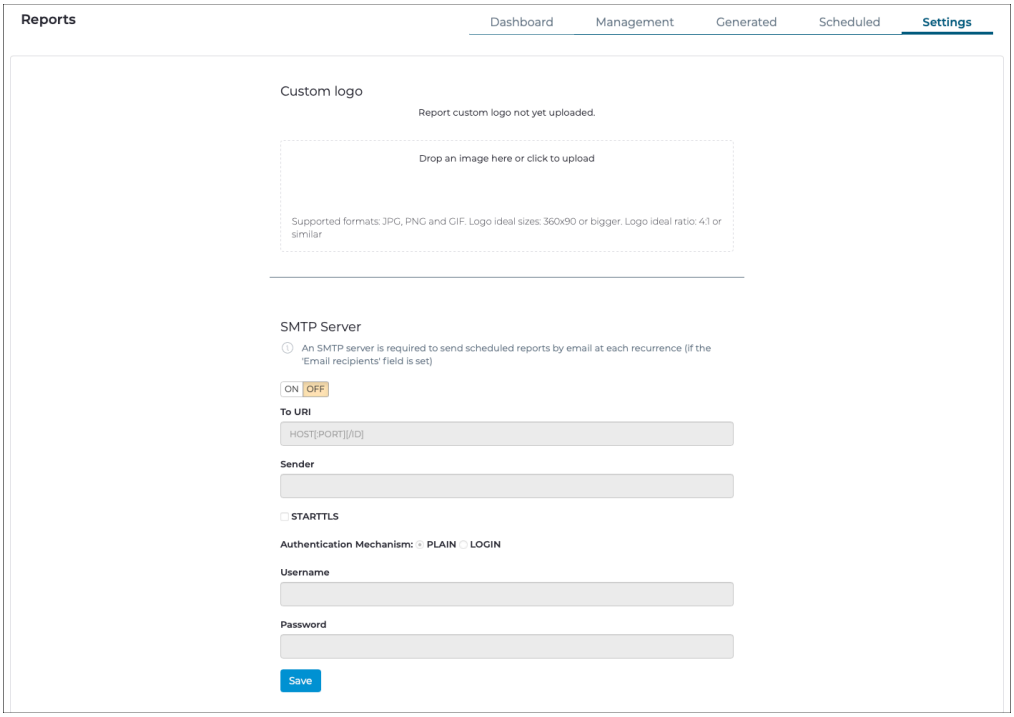


Figure 90. Settings page

Custom logo

This section lets you upload a custom logo (on page 270) that will show in your reports.

SMTP server

If you want to send emails that contain scheduled reports, you need to configure (on page 272) the *simple mail transfer protocol (SMTP)* server settings.

Create a report

*The **Management** page lets you create a new report.*

Procedure

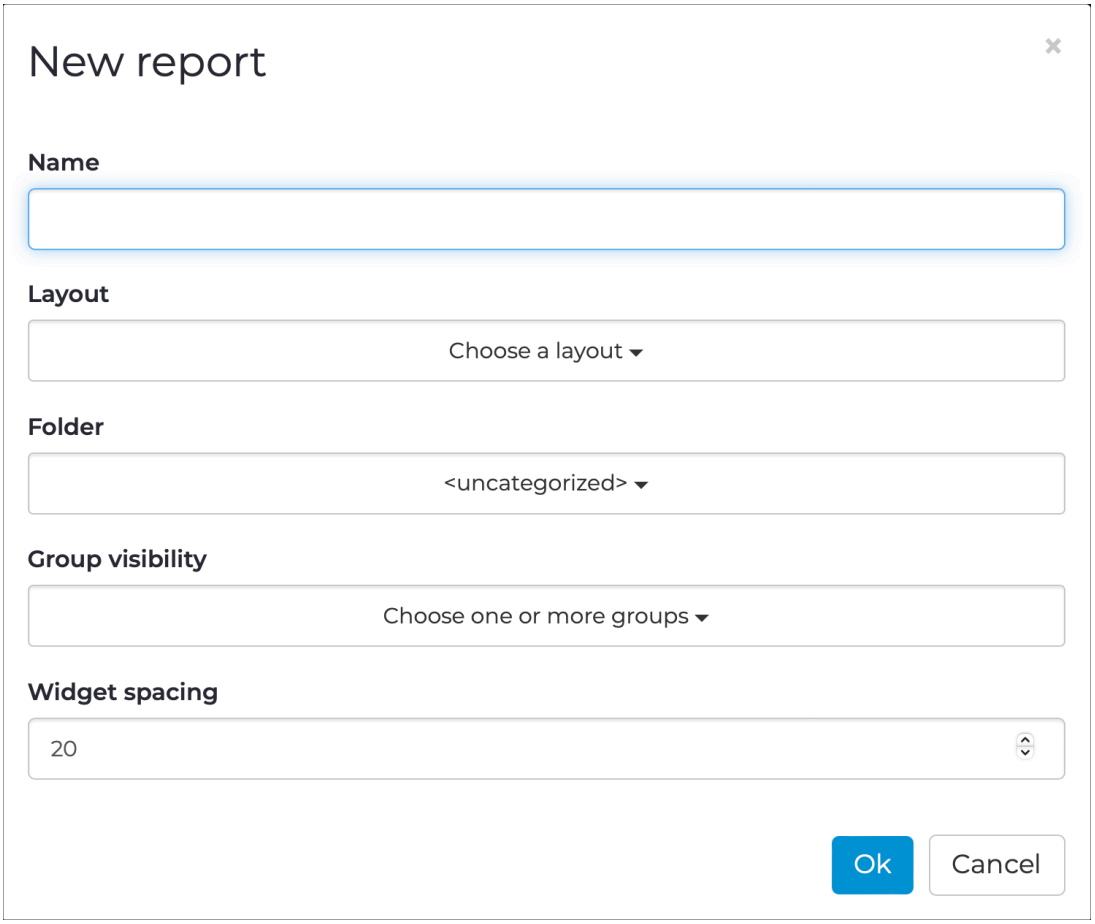
1. In the top navigation bar, select icon > **Reports**.

   Result: The **Reports** page opens.

2. Select **Management**.
3. In the section on the left, select **New report**.

   Result: A dialog shows.

4. In the **Name** field, enter a name for the report.



5. From the **Layout** dropdown, select a layout for the report.
6. From the **Folder** dropdown, select a folder for the report.
7. From the **Group visibility** dropdown, select the group(s) that will be able to view the report.

8. From the **Widget spacing** dropdown, enter a value.
9. Select **Ok**.

Results

The report has been created.

Generate a report

You can generate both scheduled, or on-demand, reports, in multiple file formats.

Procedure

1. In the top navigation bar, select **icon > Reports**.

   Result: The **Reports** page opens.

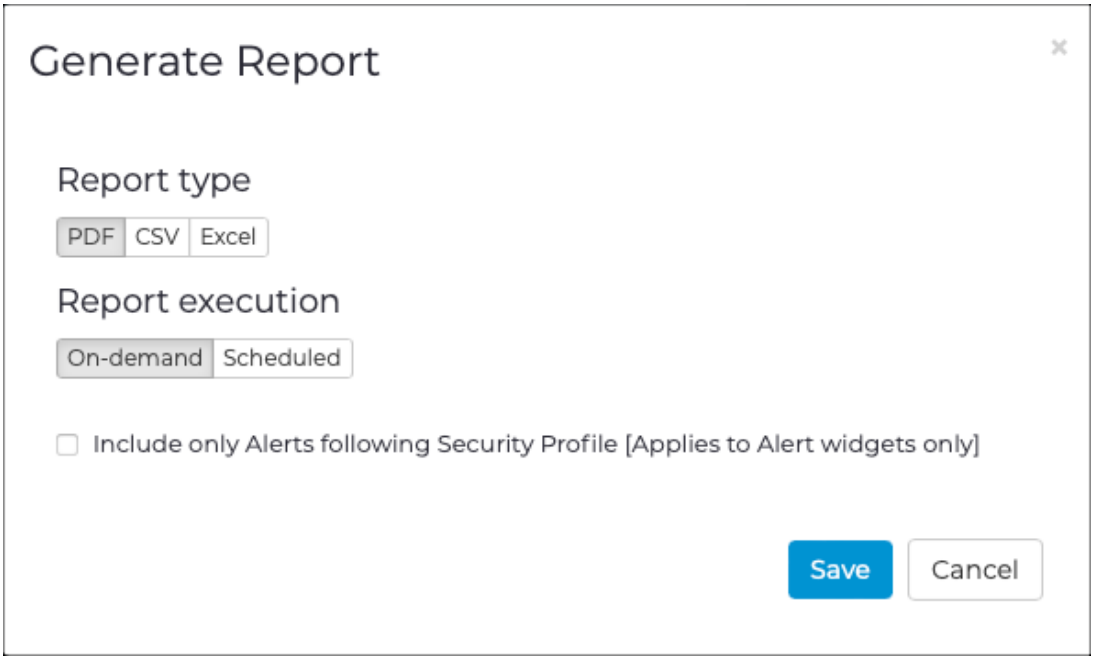
2. Select **Management**.
3. In the section on the left, select the report that you want to generate.
4. In the top right, select **Generate Report**.

   Result: A dialog shows.

5. In the **Report type** section, choose a format for the report:

   Choose from:

   - *PDF*
   - *CSV*
   - Excel



6. In the **Report execution** section, choose the type of execution:

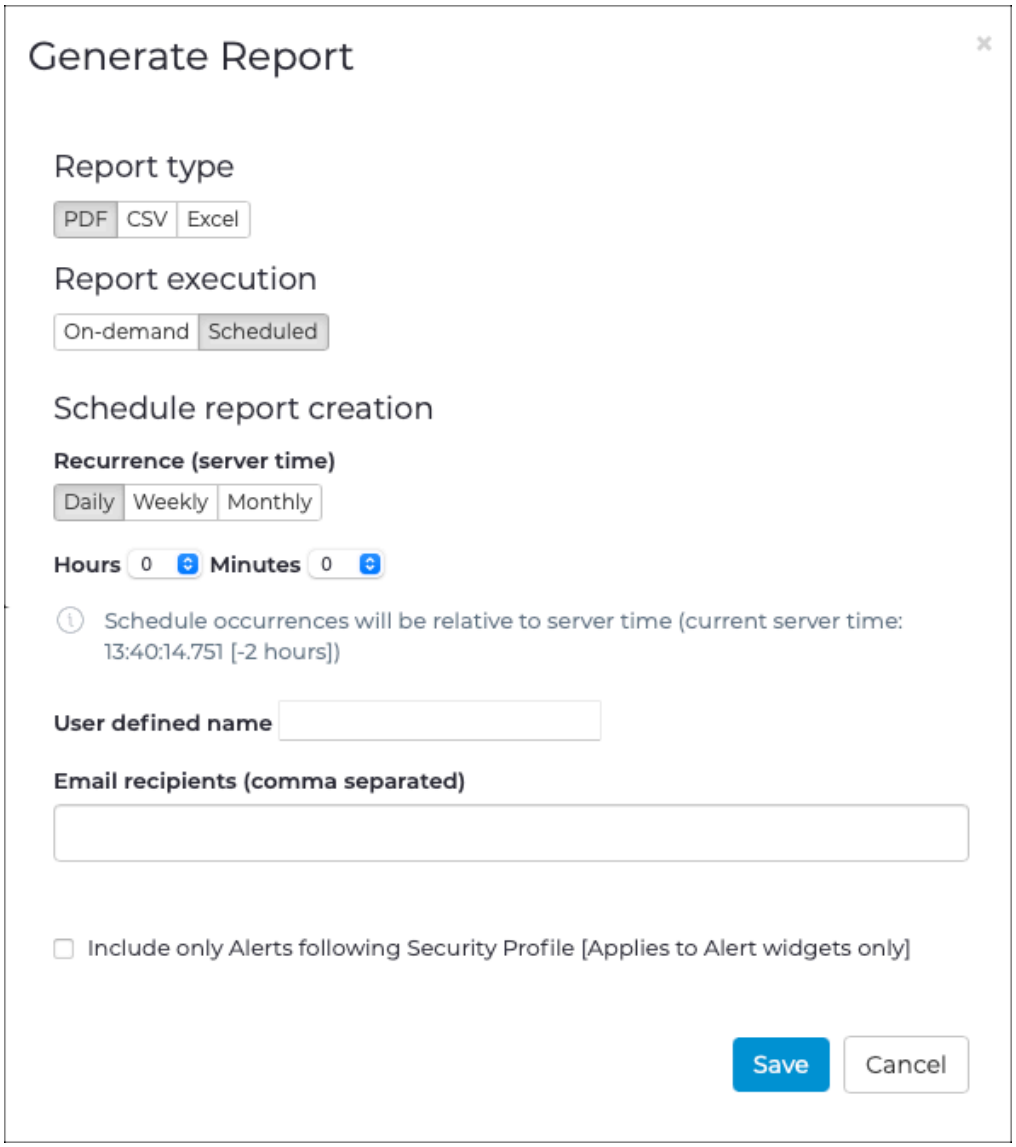
   Choose from:

   - On-demand
   - Scheduled

7. If you chose **On-demand**, select **Save**.

   Result: The report starts to generate. When the generation is complete, the report will show in the Generated (on page 258) page.

8. If you chose **Scheduled**, do the steps below.
   a. In the **Recurrence (server time)** section, enter the settings that you want.



   b. **Optional:** In the **User defined name** field, enter a name for the report.

   c. **Optional:** In the **Email recipients (comma separated)** field, enter the email addresses of the people that you would like to receive the reports.

   d. **Optional:** If necessary, select the **Include only Alerts following Security Profile [Applies to widgets only]** checkbox.

Results

The report has been generated, or scheduled, as applicable.

Download a report

*The **Generated** page lets you download reports.*

Procedure

1. In the top navigation bar, select icon > **Reports**.

   Result: The **Reports** page opens.

2. Select **Generated**.
3. To the left of the applicable report, select the download icon.

   Result: The download starts.

Results

The report has been downloaded to your downloads folder.

Delete a report

*The **Generated** page lets you delete generated reports.*

Procedure

1. In the top navigation bar, select icon > **Reports**.

   Result: The **Reports** page opens.

2. Select **Generated**.
3. To the left of the applicable report, select the download icon.

Results

The report has been deleted.

Add a folder

*The **Management** page lets you add folders for you to organize your reports.*

Procedure

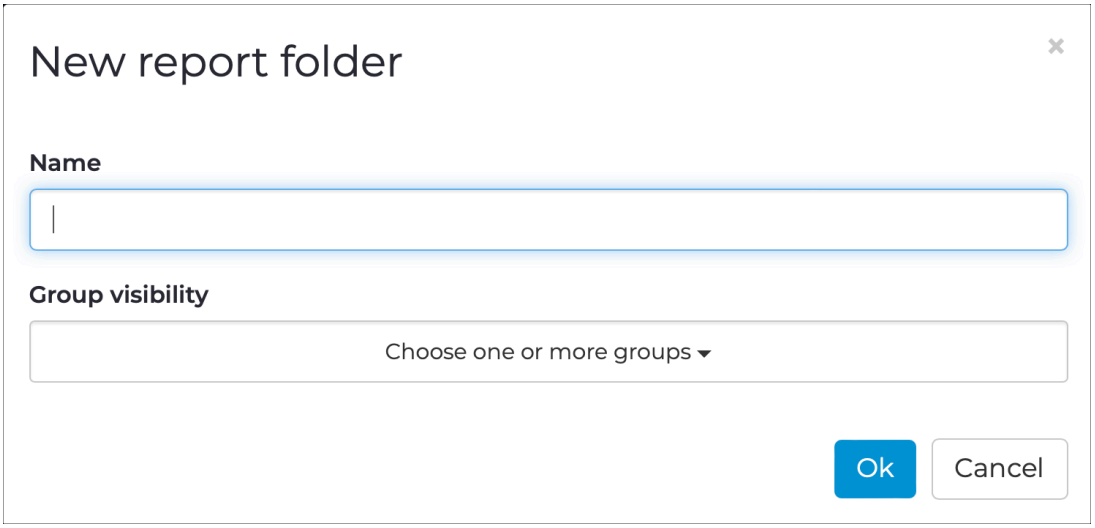
1. In the top navigation bar, select icon > **Reports**.

   Result: The **Reports** page opens.

2. Select **Management**.
3. In the **Reports list** section on the left, select **Add folder**.

   Result: A dialog shows.

4. In the **Name** field, enter a name for the folder.



5. From the **Group visibility** dropdown, select the group(s) that will be able to view the reports.
6. Select **Ok**.

Results

The folder has been added.

Edit a folder

*The **Management** page lets you delete reports.*

Procedure

1. In the top navigation bar, select ☰ **icon > Reports**.

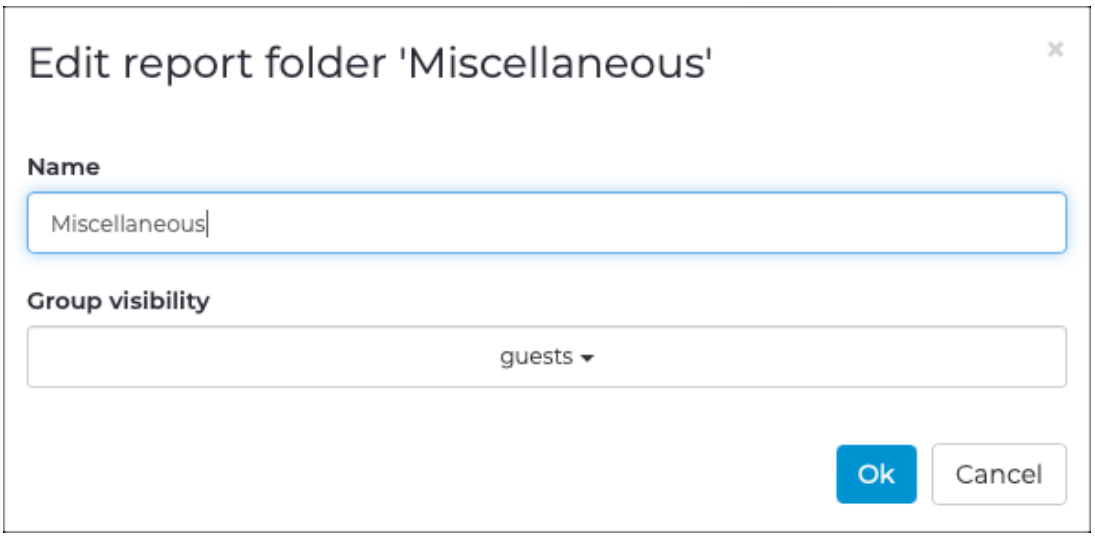
   Result: The **Reports** page opens.

2. Select **Management**.
3. In the section on the left, to the right of the applicable folder's name, select the download ✎ icon.

   Result: A dialog shows.

4. **Optional:**

   If necessary, in the **Name** filed, edit the name of the folder.



5. **Optional:** If necessary, in the **Group visibility** filed, edit the visibility of the folder.
6. Select **Ok**.

Results

The report has been edited.

Delete a folder

*The **Management** page lets you delete reports.*

Procedure

1. In the top navigation bar, select icon > **Reports**.

   Result: The **Reports** page opens.

2. Select **Management**.
3. In the section on the left, to the right of the applicable folder's name, select the download icon.

Results

The report has been deleted.

Import a schema

*The **Management** page lets you import a schema that has previously been exported.*

Procedure

1. In the top navigation bar, select icon > **Reports**.

   Result: The **Reports** page opens.

2. Select **Management**.
3. In the section on the left, select **Import Schema**.
4. Select the schema to import.

Results

The schema has been imported.

Export a schema

*The **Management** page lets you export a schema.*

Procedure

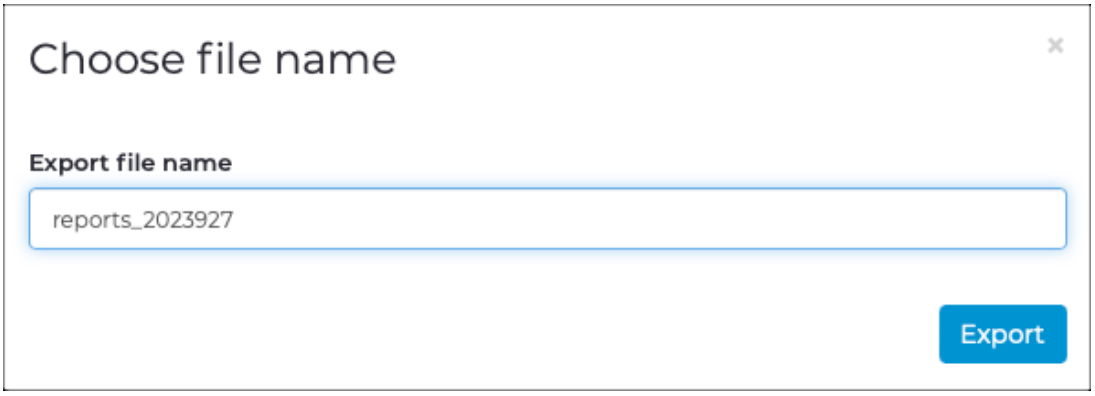
1. In the top navigation bar, select ☰ **icon > Reports**.

   Result: The **Reports** page opens.

2. Select **Management**.
3. In the top right, select **Export Schema**.

   Result: A dialog shows.

4. In the **Export file name** field, enter a name for the schema.



5. Select **Export**.

Results

The schema has been exported in *JSON* format.

Upload a custom logo

You can add a custom logo that will show in your reports.

Procedure

1. In the top navigation bar, select icon > **Reports**.

   Result: The **Reports** page opens.

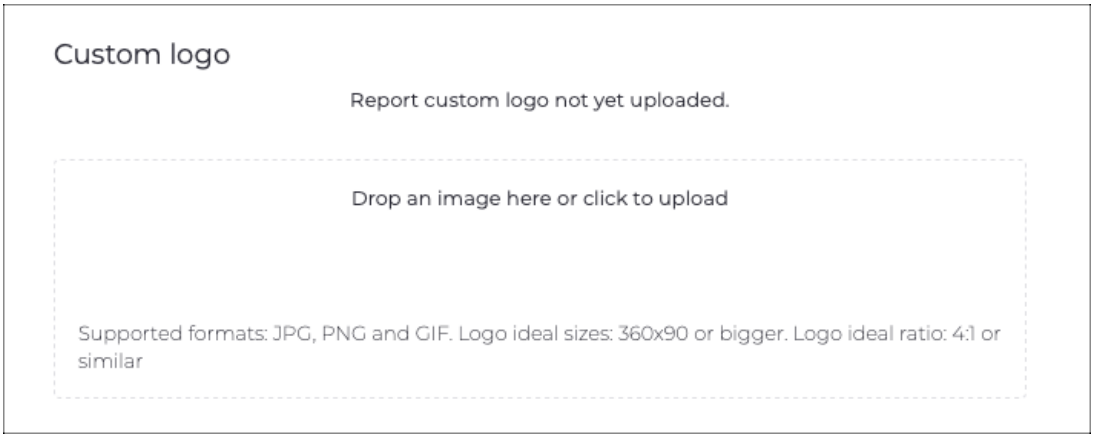
2. Select **Settings**.

   Result: The Settings (on page 260) page opens.

3. Choose a method to upload a custom logo:

   Choose from:
   - Drag your image file into the **Drop an image here or click to upload** field
   - Click in the **Drop an image here or click to upload** field

4. If you chose the second method, select the correct file to upload.



Note:
Logos should be 360x90 pixels or bigger, with an ideal aspect ratio of 4:1, or similar.

Note:
Supported formats are:
- *graphics interchange format (GIF)*
- *joint photographic experts group (JPEG)*
- *portable network graphics (PNG)*

5. Wait for the file to upload.
6. Select **Save**.

Results

Your custom logo will now be added to your reports.

Configure SMTP settings

If you want to send emails that contain scheduled reports, you need to configure the simple mail transfer protocol (SMTP) server settings.

Procedure

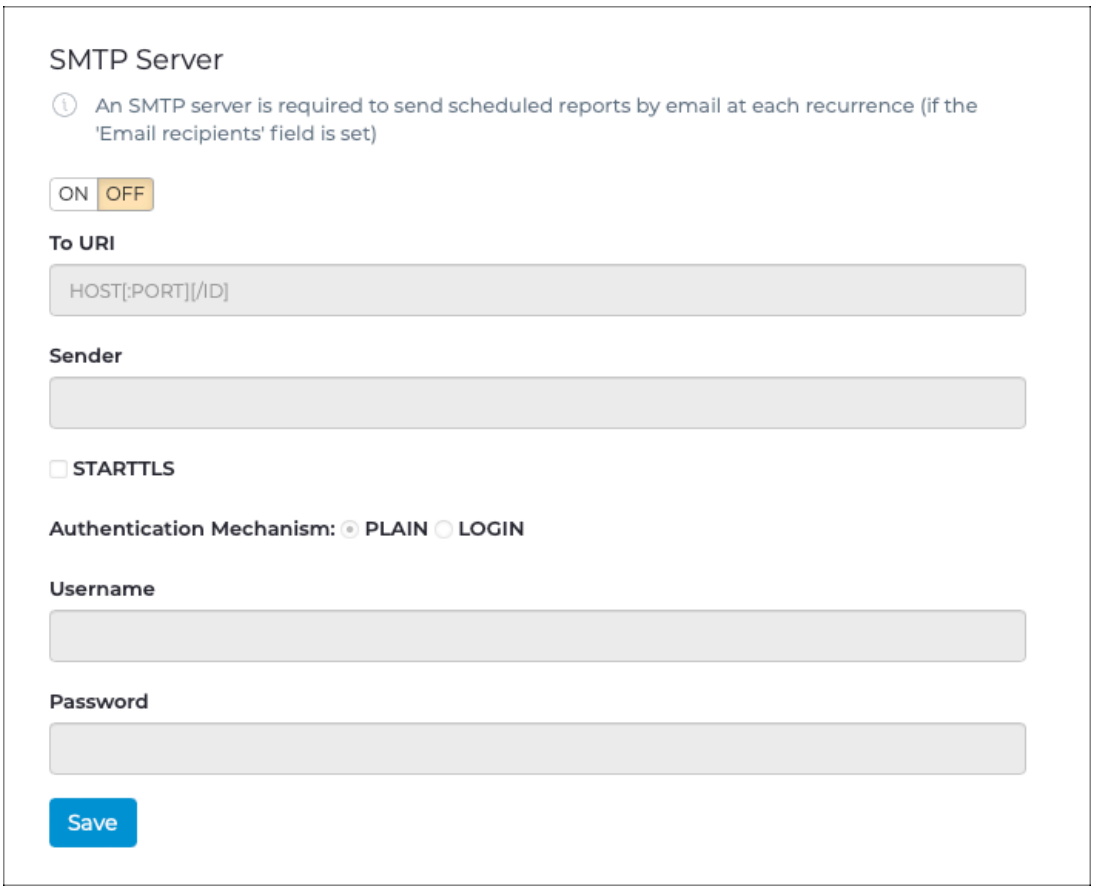
1. In the top navigation bar, select icon > **Reports**.

   Result: The **Reports** page opens.

2. Select **Settings**.

   Result: The Settings (on page 260) page opens.

3. In the **SMTP Server** section, set the toggle to **ON**.



4. In the **To URI** field, enter the host URI information. For example, `HOST[:PORT][/ID]`
5. In the **Sender** field, enter the sender identification information.

6. To use encryption, select the **STARTTLS** checkbox.

> **Note:**
> If you do not select this option, reports will be sent without encryption.

7. To start the authentication process, choose an Authentication Mechanism:

Choose from:

- **PLAIN**
- **LOGIN**

> **Note:**
> The default setting is **PLAIN**.

8. If you chose **LOGIN**, enter your credentials.
    a. In the **Username** field, enter your username.
    b. In the **Password** field, enter your password.
9. Select **Save**.

Results

Emails for scheduled reports that have email recipients will now be sent at the next scheduled occurrence.

Filter a report globally

You can filter a report globally, which is the default filter. This lets you use a specific category to apply filters to the entire report.

Procedure

1. In the top navigation bar, select icon > **Reports**.

   Result: The **Reports** page opens.

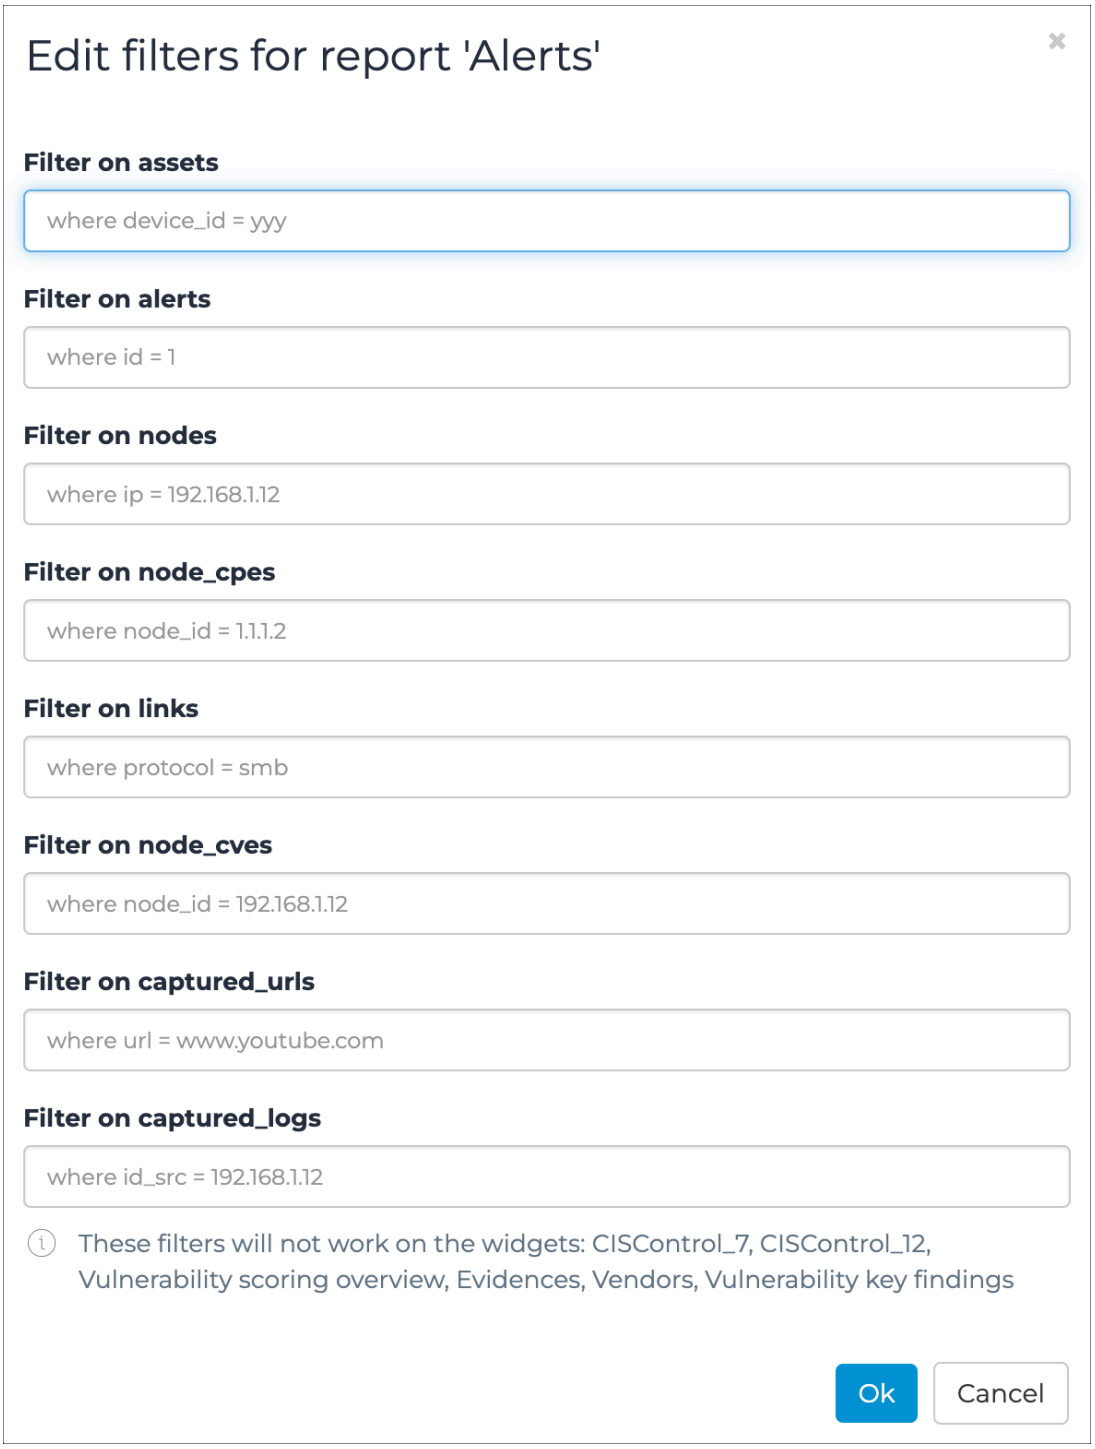
2. Select **Management**.

   Result: The Management (on page 255) page opens.

3. In the top section, select **Filters**.

   Result: A dialog shows.

4. Select the category on which to filter, then enter your filter query in the related field.



5. Select **Ok**.

> **Note:**
> At the bottom of the dialog is a list of widgets on which filters will not work.

Results

The filter(s) has (have) been applied to the report.

Add a widget to a report

*The **Management** page lets you add widgets to reports.*

Procedure

1. In the top navigation bar, select ☰ **icon > Reports**.

   Result: The **Reports** page opens.

2. Select **Management**.

   Result: The Management (on page 255) page opens.

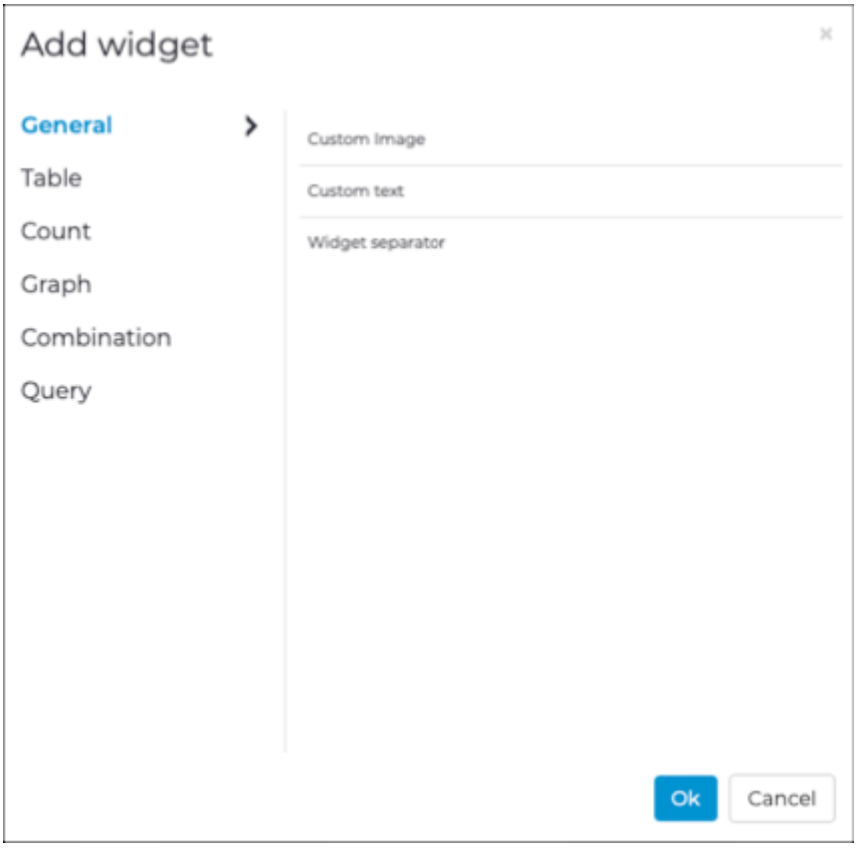
3. In the **Reports list** section on the left, select the applicable report.

   Result: The report opens.

4. On the right side of the report, choose the section that you want to add the widget to. Select **Add widget**.

   Result: A dialog shows.

5. In the left pane, choose the type of widget that you want to add.



6. From the list, select the widget that you want.

Results

The widget has been added to the report.

Filter a report with a widget

You can use widgets to filter a report.

Procedure

1. In the top navigation bar, select ☰ **icon > Reports**.

   Result: The **Reports** page opens.

2. Select **Management**.

   Result: The Management (on page 255) page opens.

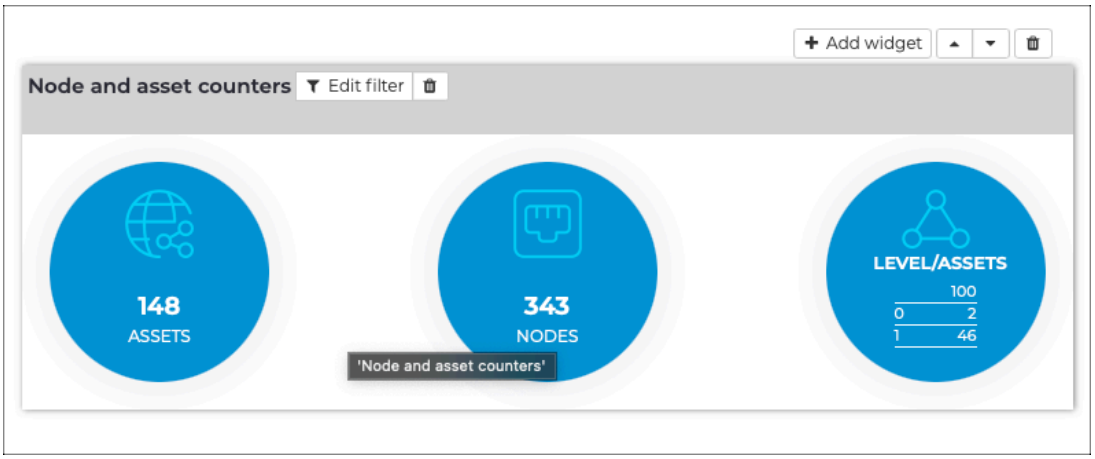
3. In the **Reports list** section on the left, select the applicable report.

   Result: The report opens.

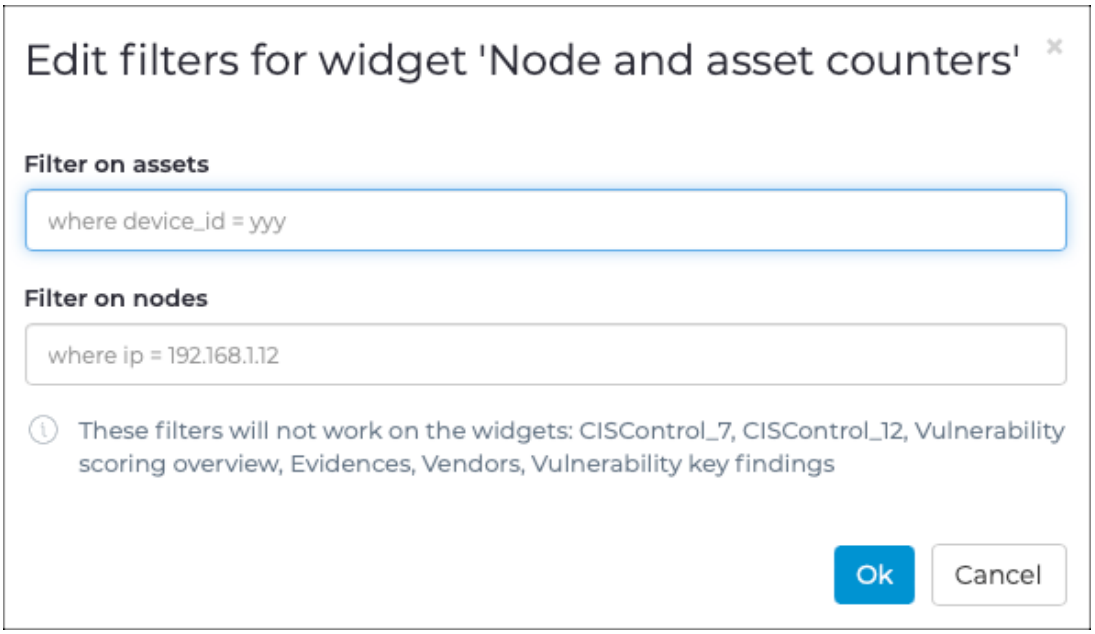
4. Find the widget that you would like to use and hover your mouse over it.

   Result: At the top of the widget, more buttons show.

5. Select **Edit filter**.



6. Select the category on which to filter, then enter your filter query in the related field.



7. Select **Ok**.

> **Note:**
> At the bottom of the dialog is a list of widgets on which filters will not work.

Results

The filter(s) has (have) been applied to the report.

Chapter 12. Assertions



*The **Assertions** page shows all the assertions and lets you configure them.*

A valid assertion is a normal query with a special command appended at the end. Assertions can be saved in a specific order and can be continuously executed in the system.

Queries are based on the *N2QL*, which you can use to ensure that certain conditions are met on the observed system. An assertion is typically either an empty value, or a specific value. When an unexpected value appears, or when the value is different than the expected, the system alerts the user.

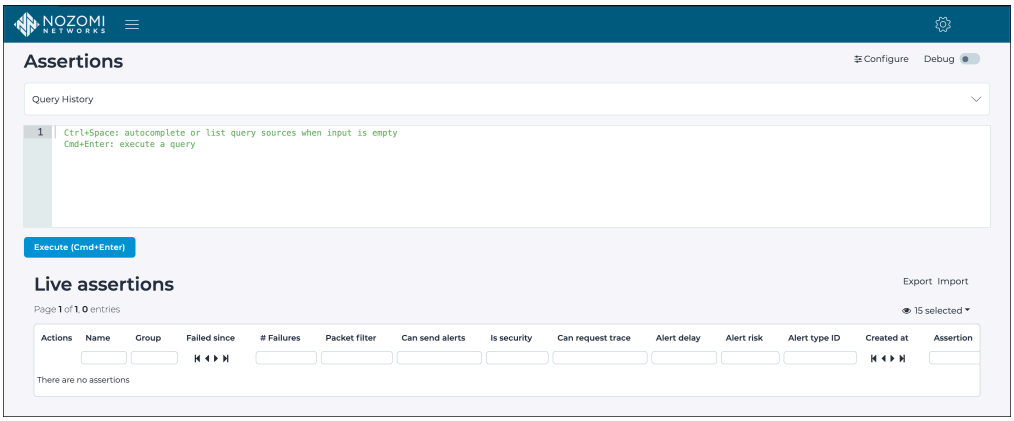


Figure 91. Assertions page

Configure

This lets you configure the execution interval in seconds.

History button

This button shows a history of the previous queries that have been entered in the query field.

Query field

This field is where you enter your query.

Debug button

Because assertions with logical operators and brackets can quickly become complex, the debug icon decomposes the query, and executes each part to show intermediate results.



Figure 92. Complex assertion being debugged

Export

This lets you export assertion groups in *JSON* format.

Import

This lets you import assertion groups in *JSON* format.

Assertion operators

Table 2. Assertion operators

Operator	Description
`assert_all <field> <op> <value>`	The assertion is satisfied when each element in the query result set matches the given condition.
`assert_any <field> <op> <value>`	The assertion is satisfied when at least one element in the query result set matches the given condition.
`assert_empty`	The assertion is satisfied when the query returns an empty result set.
`assert_not_empty`	The assertion is satisfied when the query returns a non-empty result set.

Save an assertion

You can save assertions to have them continuously executed in the system.

Procedure

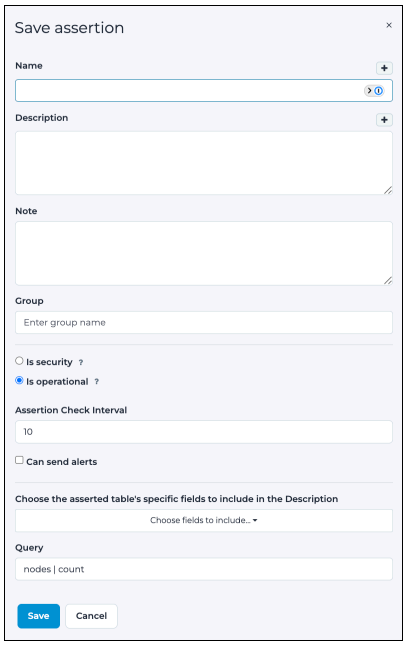
1. In the top navigation bar, select ☰ **icon > Assertions**.

   Result: The **Assertions** page opens.

2. In the query field, enter a query.
3. Select **Enter**.
4. To save the assertion, select **Save**.

   Result: A dialog shows.

5. In the **Name** field, enter a name for the query.



Note:
You can use field placeholders to include dynamic values from your query results. For more details, see Customize assertion name and description (on page 288).

6. In the **Description** field, enter a description.

   > **Note:**
   > You can use field placeholders to include dynamic values from your query results. For more details, see Customize assertion name and description (on page 288).

7. In the **Group name** field, enter a group name for the assertion.
8. Choose from one of these options:

   Choose from:
   - **Is security ?**
   - **Is operational ?**

9. In the **Assertion Check Interval** field, choose the interval in seconds at which the assertion will be rechecked.

   > **Note:**
   > You can select an interval between 10 seconds and 1 day.

10. **Optional:** If you want the assertion to trigger an alert, select the **Can send alerts** checkbox.
11. From the **Choose the asserted table's specific fields to include in the Description** dropdown, select the fields to include in the assertion description.
12. In the **Query** field, enter the assertion query.
13. Select **Save**.

    Result: The saved assertion will be listed at the bottom of the page with a green or red color to indicate the result.

Customize assertion name and description

You can customize assertion names and descriptions with dynamic field values using string interpolation.

About this task

When saving an assertion, you can use field placeholders in the name and description fields. These placeholders are replaced with actual values from your query results when alerts are generated. This lets you create more descriptive and context-aware assertion names and alert descriptions.

Procedure

1. In the top navigation bar, select icon > **Assertions**.

   Result: The **Assertions** page opens.

2. In the query field, enter a query.
3. Select **Enter**.
4. Select **Save**.

   Result: A dialog shows.

5. In the **Name** field, enter a name for the assertion. To insert a field value from your query, select the plus icon button next to the field.

   Result: A placeholder `{{FIELD_NAME}}` is inserted into the field with `FIELD_NAME` highlighted. Replace it with the actual field name from your query.

   > **Tip:**
   > Field names must match the fields returned by your query. For example, if your query returns `host`, `name`, or `ip` fields, you can use `{{host}}`, `{{name}}`, or `{{ip}}`.

6. In the **Description** field, enter a description. You can also use field placeholders here by selecting the plus icon button.
   Example: `The sensor {{host}} with model {{model}} has been offline for more than 15 minutes`
7. Complete the remaining fields in the dialog and select **Save**.

Results

When the assertion generates an alert, the placeholders in the name and description are replaced with the actual values from the query results.

Example

If you save an assertion with the following configuration:

- Name: `Sensor offline: {{host}}`
- Description: `The sensor {{host}} with model {{model}} has not synchronized in the last 15 minutes`
- Query: `sensors | where minutes_ago(last_sync) > 15 | assert_empty`

When a sensor with `host` = `nozomi-guardian` and `model` = `V-SERIES` fails the assertion, the generated alert will display:

- Name: `Sensor offline: nozomi-guardian`
- Description: `The sensor nozomi-guardian with model V-SERIES has not synchronized in the last 15 minutes`

Edit an assertion

This lets you edit the details for an existing assertion.

Procedure

1. In the top navigation bar, select icon > **Assertions**.

   Result: The **Assertions** page opens.

2. In the query field, enter a query.
3. To execute the query, elect **Enter**.

> **Note:**
> You can use the logical operators && (and) and || (or) to combine multiple assertions. Round brackets ( ) change the logical grouping as in a mathematical expression.

4. **Optional:** To the right of the query field, select the debug icon.

> **Note:**
> Because assertions with logical operators and brackets can quickly become complex, the debug icon decomposes the query, and executes each part to show intermediate results.



Result: The debug section shows.

Configure an assertion

This lets you configure the execution interval of the assertions, in seconds.

Procedure

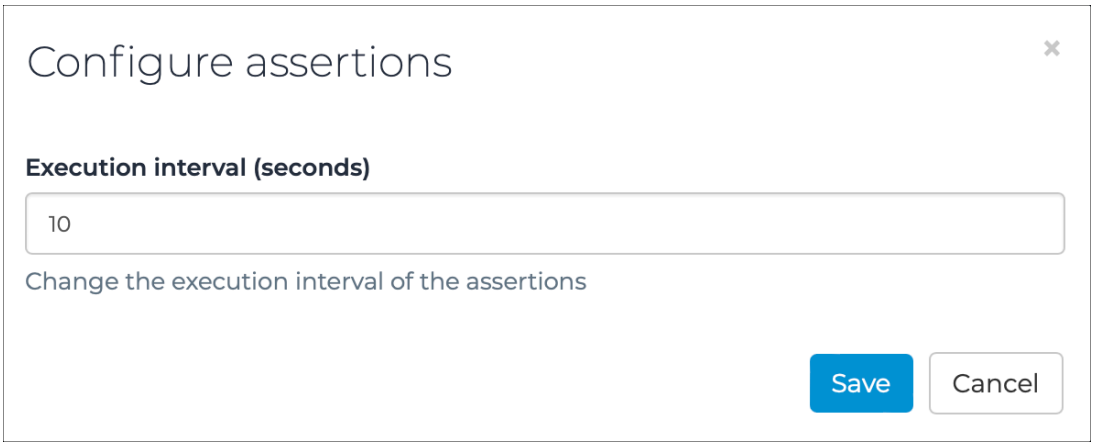
1. In the top navigation bar, select icon > **Assertions**.

   Result: The **Assertions** page opens.

2. In the top right, select **Configure**.

   Result: A dialog opens.

3. In the **Execution interval (seconds)** field, enter an interval.



4. Select **Save**.

Configure an assertion on links

This lets you configure the scope of the assertion for the related links.

Procedure

1. In the top navigation bar, select icon > **Network**.

   Result: The **Network** page opens.

2. Select **Links**.

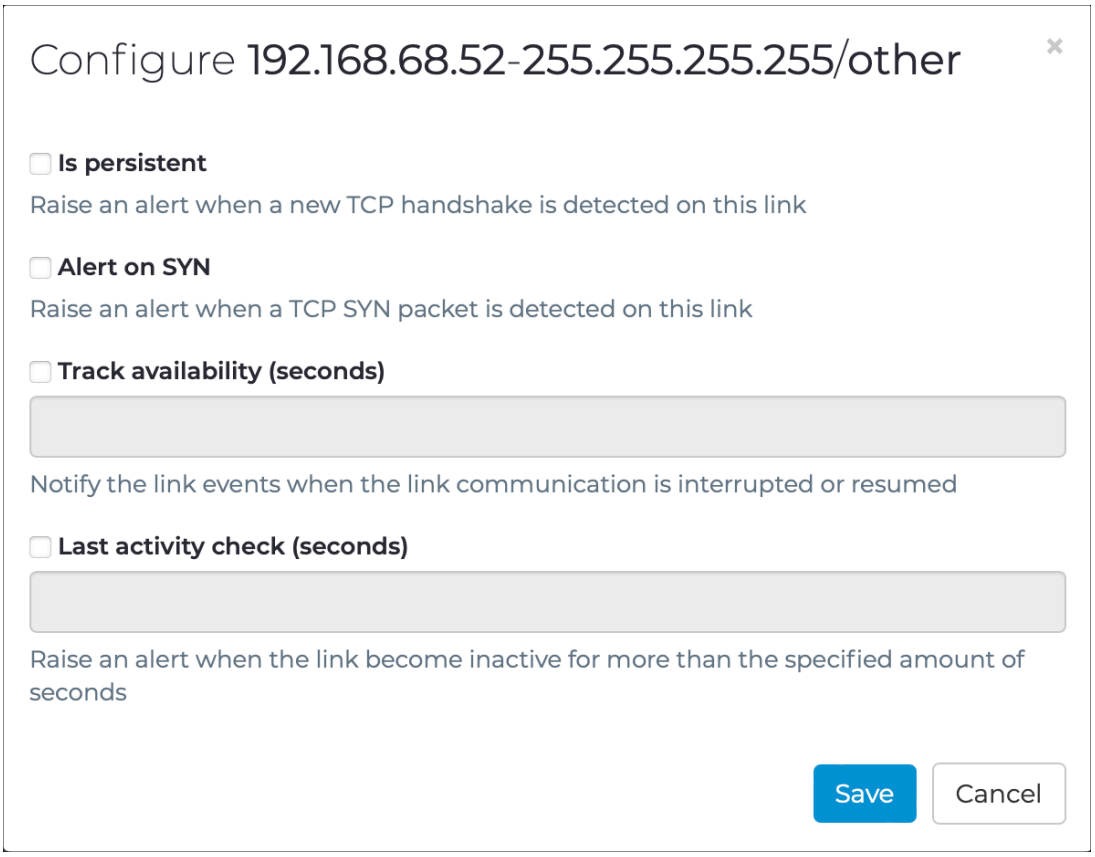
   Result: The **Links** page opens.

3. To the left of the applicable link, select the configure icon.

   Result: A dialog opens.

4. **Optional:**

   Select **Is persistent**.



Note:
When selected, this check raises a new alert whenever a *TCP* handshake is successfully completed on the link.

5. **Optional:** Select **Alert on SYN**.

> **Note:**
> When selected, this check raises a new alert whenever a client sends a TCP SYN on the link.

6. **Optional:** Select **Track availability (seconds)**.

> **Note:**
> When selected, a link is considered non-functioning if it is unresponsive for the specified time.

7. **Optional:** Select **Last activity check (seconds)**.

> **Note:**
> When selected, this check raises an alert when the link is not receiving data for more than the specified time.

8. Select **Save**.

Results

The assertion has been configured.

Configure an assertion on variables

This lets you configure the scope of the assertion for the related variables.

Procedure

1. In the top navigation bar, select icon > **Process**.

   Result: The **Process** page opens.

2. Select **List**.

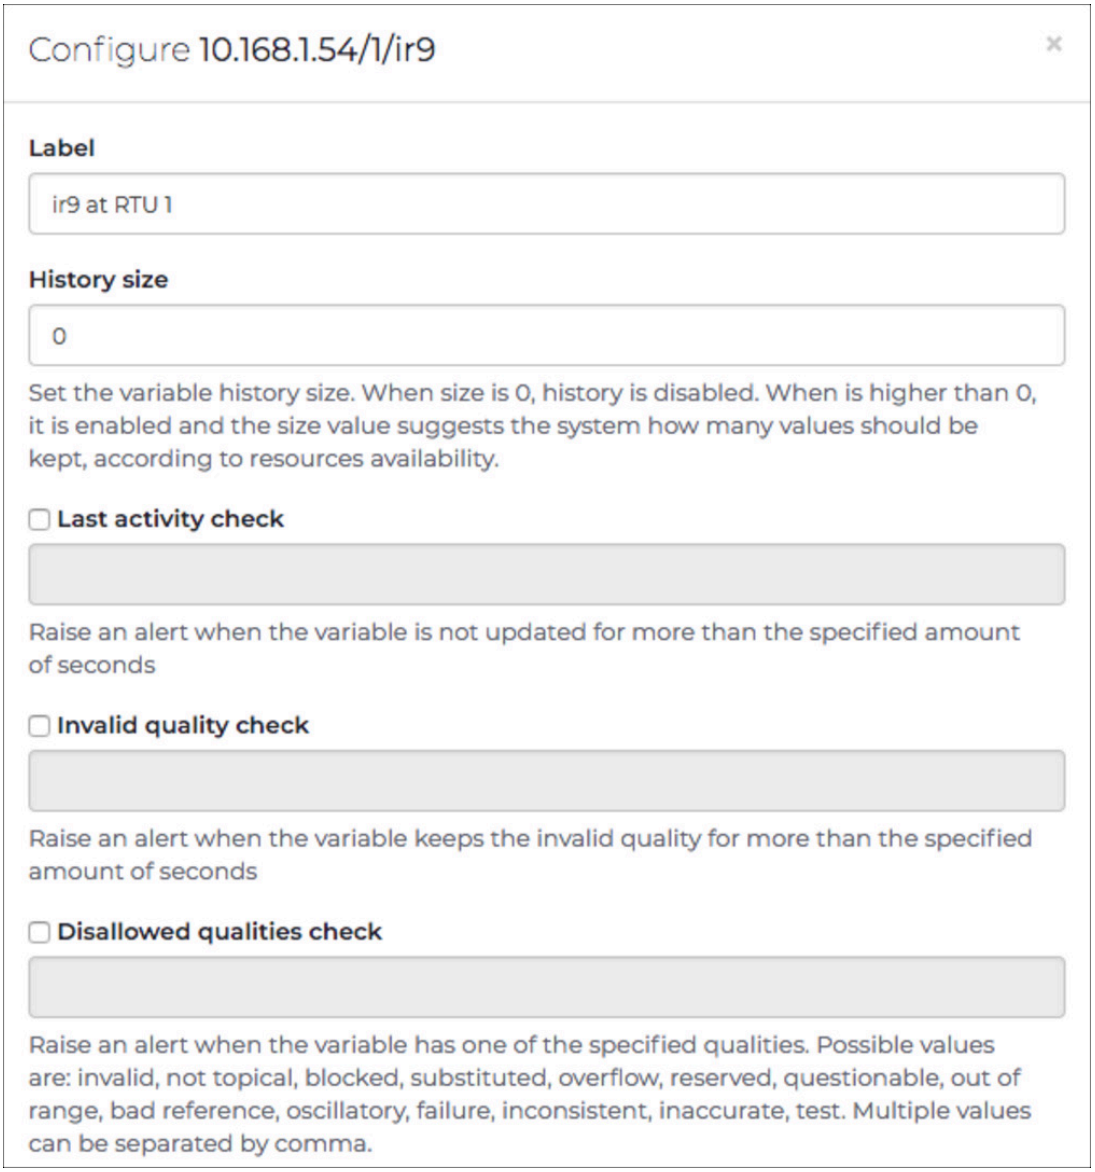
   Result: The **List** page opens.

3. To the left of the applicable link, select the configure icon.

   Result: A dialog opens.

4. **Optional:**

   In the **Label** field, enter a label for the assertion.



5. **Optional:** In the **History size** field, enter a value.

> **Note:**
> This sets the *Variable* history size. When the size is 0, history is disabled. When it is higher than 0, it is enabled, and the size value suggests how many values that the system should keep, depending on the available resources.

6. **Optional:** In the **Last activity check** field, enter a value.

> **Note:**
> When selected, this check raises an alert when the *Variable* is either not measured or is changed for more than the specified number of seconds.

7. **Optional:** In the **Invalid quality check** field, enter a value.

> **Note:**
> When selected, this check raises an alert when the *Variable* maintains an invalid quality for more than the specified amount of seconds.

8. **Optional:** In the **Disallowed quality check** field, enter a value.

> **Note:**
> When selected, this check raises an alert when the *Variable* gains one of the specified qualities.

9. Select **Save**.

Results

The assertion has been configured.

Chapter 13. Time machine



Time machine lets you load a snapshot, which is a previously saved state, and go back in time, to analyze the data from a past situation. You can load a single snapshot and use the platform as usual or load two snapshots and compare the user interface to highlight changes.

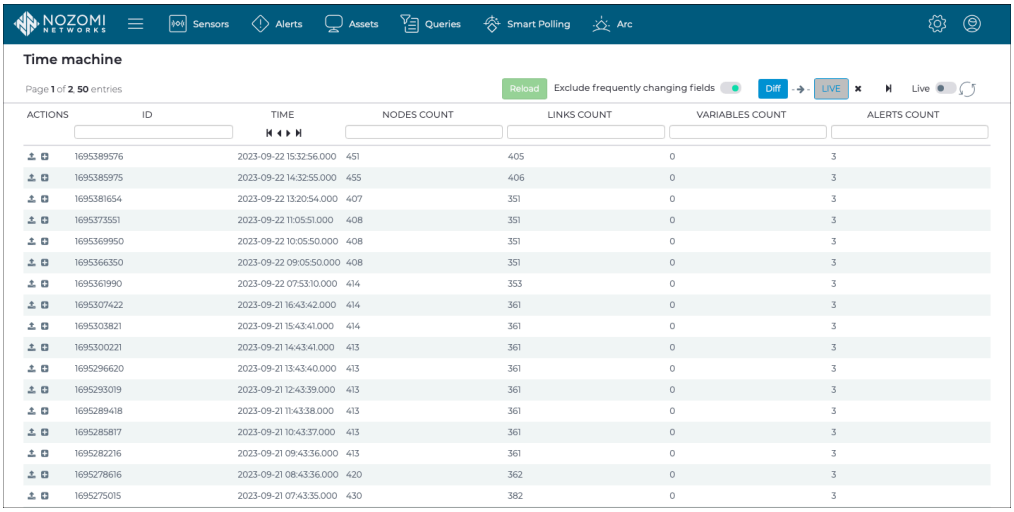


Figure 93. Time machine page

Reload

If you initiate a diff and then navigate to another area of the software, when you return to this page you might not see the progress of the diff. The **Reload** button lets you either:

- Reload the progress bar (if the diff is still in progress)
- View the results (if the process has finished)

Exclude frequently changing fields

You can select the **Exclude frequently changing fields** toggle to on to exclude the fields that change frequently. This excludes the diff all fields that are affected from normal traffic handling. Examples of frequently changing fields are the number of bytes received / sent, and the last activity time.

Diff

The **Diff** button starts the diff process. The system evaluates the diff baseline - target files and estimates how *CPU* / memory intensive the diff operation is going to be. If there is not enough free memory at the moment, the diff will be aborted with the appropriate message. If the diff is estimated to take more than a few minutes, a warning will show and a confirmation dialog will show.

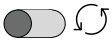
Live

The **LIVE** button lets you choose the current view as part of the diff.

Step forward

The step forward icon lets you go back to the present time view after you have viewed a snapshot.

Live / refresh

The **Live**  icon lets you change live view on, or off. When live mode is on, the page will refresh approximately every five seconds.

Load snapshot

The load snapshot  icon lets you select a historical snapshot.

Select for diff

The select for diff  icon to choose the related snapshot as part of the diff.

Time machine diffs

You can use time machine to view diffs between historical snapshots, or an historical snapshot and the current live view.

Snapshots

When you view a snapshot, the *graphical user interface (GUI)* will change to gray.



The graph view and the use of color let you quickly see the nodes, or links, that have been added, removed, or changed:

- Added items are green
- Removed items are red
- Changed items are blue

Diff views

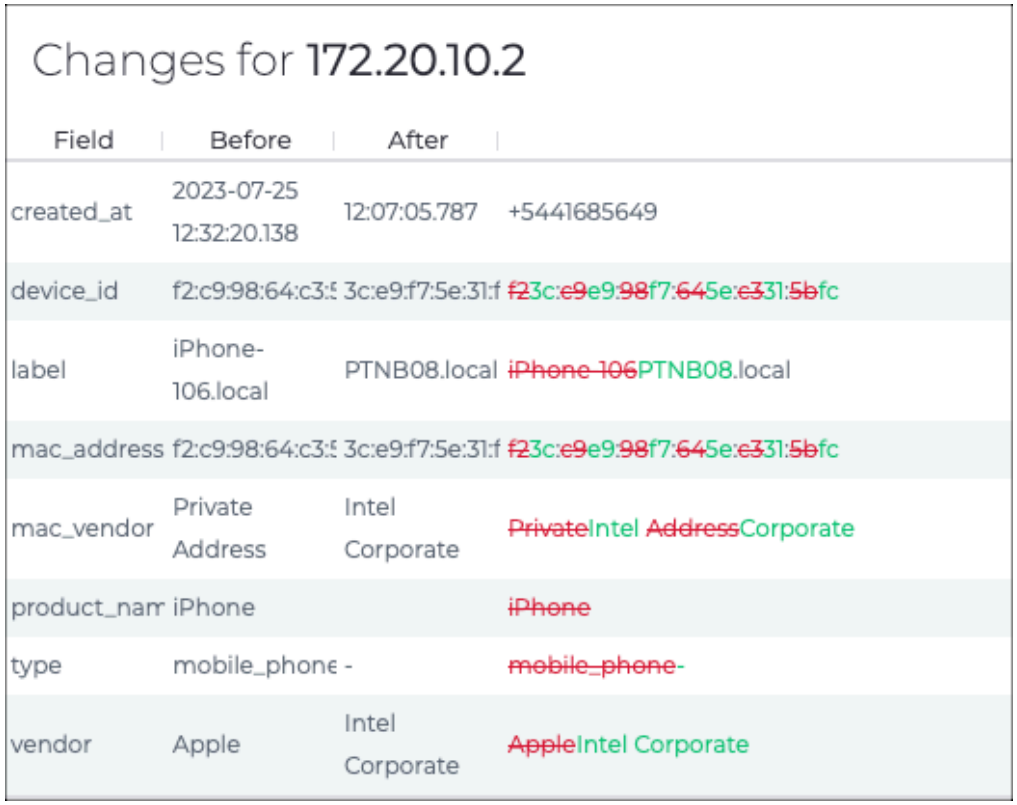


Figure 94. Table view

In the graph view, you can select a link or node to see more information in the pane on the right side.

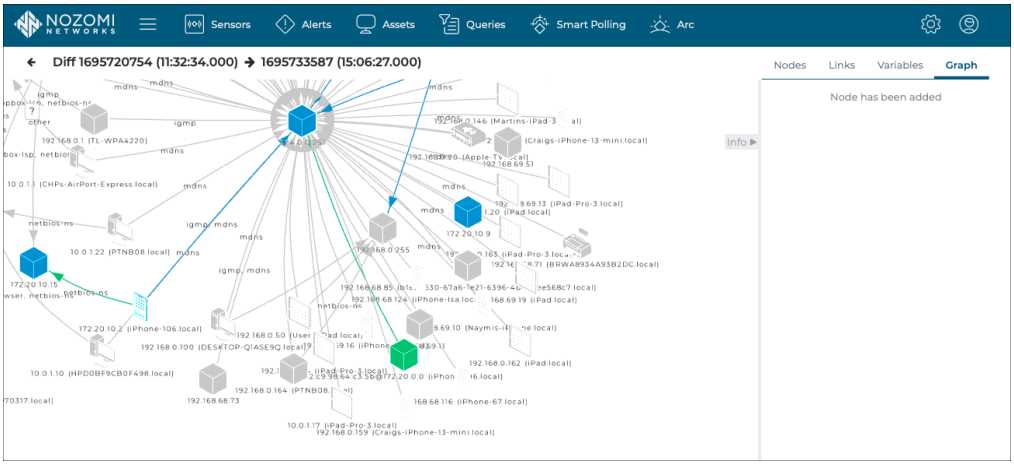


Figure 95. Graph view

Load a snapshot

To create a diff, you must first load a snapshot.

Procedure

1. In the top navigation bar, select icon > **Time machine**.

   Result: The **Time machine** page opens.

2. To the left of the applicable snapshot, select the load snapshot icon.

   Result: The user interface turns gray to highlight that you are viewing a static snapshot.



3. To return to the present time view, select the icon.

Request a diff

The time machine lets you create a diff between to historical snapshots, or a snapshot and the current view.

Procedure

1. In the top navigation bar, select icon > **Time machine**.

   Result: The **Time machine** page opens.

2. To select the baseline of the diff, to the left of the applicable snapshot, select the select for diff icon.

   Result: At the top of the table, the details of the snapshot are loaded next to the **Diff** icon.

3. Choose a target for the diff:

   Choose from:

   - To select another snapshot from the past, to the left of the applicable snapshot, select the load snapshot icon
   - To select the current, live environment, at the top of the table, select the **LIVE** icon

4. **Optional:** Select the **Exclude frequently changing fields** toggle to on.
5. Now that the baseline and the target have been set, select the **Diff** icon.

   Result: The system evaluates the diff baseline / target files and estimates how *CPU* / memory intensive the diff operation is going to be. If there is not enough free memory at the moment, the diff will be aborted with the appropriate message. If the diff is estimated to take more than a few minutes, a warning will show and a confirmation dialog will show.

6. As soon as the diff operation starts, a dialog shows the progress. To stop the diff operation, select **Abort**.

Results

After the diff has been computed, the diff results show.

Reload the diff operation progress

Diff operations can take a long time to conclude. While an operation is in progress, you can continue to use the software as normal.

Procedure

1. In the top navigation bar, select icon > **Time machine**.

   Result: The **Time machine** page opens.

2. At the top of the table, select **Reload**.

   Result: If the diff operation is still in progress, the progress dialog shows. If the operation has finished, the diff results show.

View a diff from an alert

This automatic feature will use the previous and subsequent snapshots according to the time of the alert.

Procedure

1. In the top navigation bar, select **Alerts**.

   Result: The **Alerts** page opens.

2. Choose a method to open the actions menu.

   Choose from:

   - In the table, select the hyperlink to open the details page. Select **Actions**
   - In the table, select the ••• icon

3. Select **Time machine diff**.

   Result: The time diff shows.

4. To see more details on the right side of the graph, select the applicable node or link.

Chapter 14. Vulnerabilities



The Nozomi Networks software continuously discovers vulnerabilities in monitored assets.

The Nozomi Networks software continuously discover vulnerabilities. To do this, it matches the *CPE* of a device with the National Vulnerability Database, and other data sources.

The Vulnerabilities page has these tabs:

- Assets (on page 310)
- List (on page 311)
- Stats (on page 312)

Assets

*The **Assets** page shows a list of assets with known vulnerabilities, along with a summary of the severity of the vulnerability.*

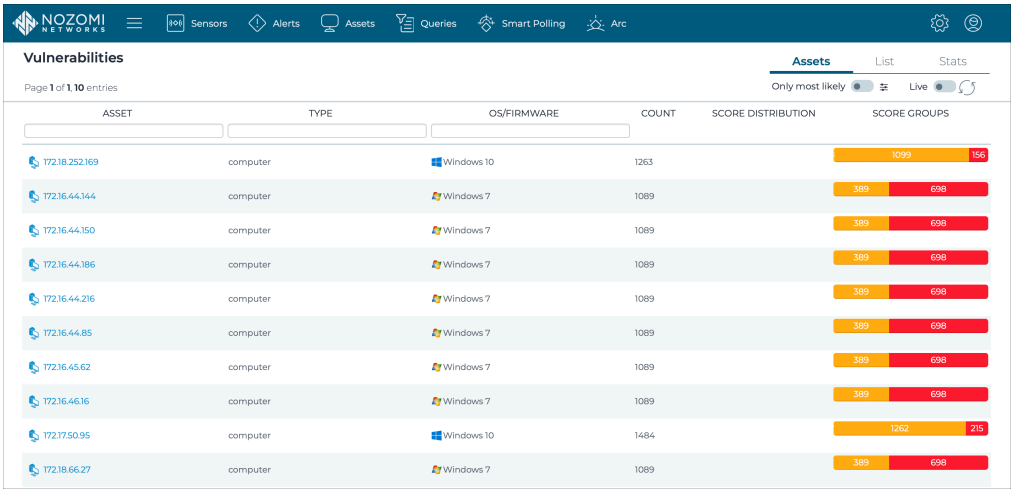


Figure 96. Assets page

Only most likely

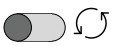
This toggle lets you filter the view to show only the assets the match the criteria that you have set for likelihood threshold.

Likelihood threshold settings

Likelihood threshold is a value between 0.1 and 1.0 where 1.0 represents the maximum likelihood of the *Common Vulnerabilities and Exposures (CVE)* to be present. Likelihood is the confidence of the software's correct assignment of a *CPE* to the hardware of the monitored asset. The higher the likelihood, the higher the software's confidence that the vulnerabilities assigned to an asset are in fact are relevant to that asset. **Likelihood threshold** is the minimum likelihood a vulnerability needs in order for it to be shown in this page when the **Only most likely** toggle is set to on. As a guideline, we suggest that you use:

- 0.8 for a high level of confidence
- 0.5 for a medium level of confidence
- 0.3 for a low level of confidence

Live / refresh

The **Live**  icon lets you change live view on, or off. When live mode is on, the page will refresh approximately every five seconds.

List

*The **List** page shows a comprehensive list of vulnerabilities in the environment. This lets you perform global, in-depth analysis.*

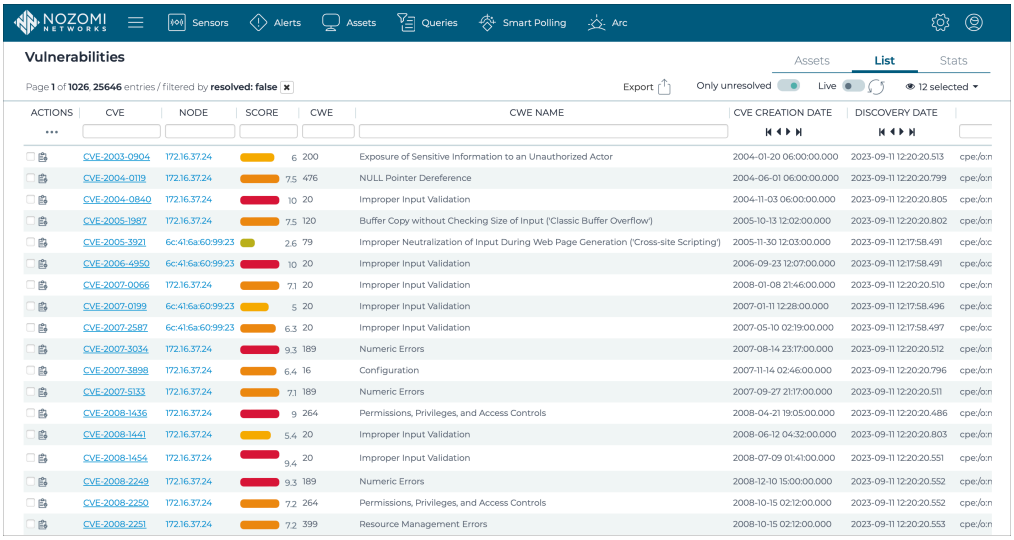


Figure 97. List page

Export

The **Export** icon lets you export the current list in either *CSV* or Microsoft Excel format.

Only resolved

This lets you show only `Unresolved` vulnerabilities. Vulnerability status options are:

- `Unresolved`
- `Mitigated`
- `Accepted`

`Mitigated` and `Accepted` lead to a resolution status that equals true.

Live / refresh

The **Live** icon lets you change live view on, or off. When live mode is on, the page will refresh approximately every five seconds.

Column selection

The columns selection icon lets you choose which columns to show or hide.

Stats

*The **Stats** page shows high level information in a graphical format that shows the top common platform enumerations (CPEs), common vulnerabilities and exposures (CVEs), and common weakness enumerations (CWEs).*

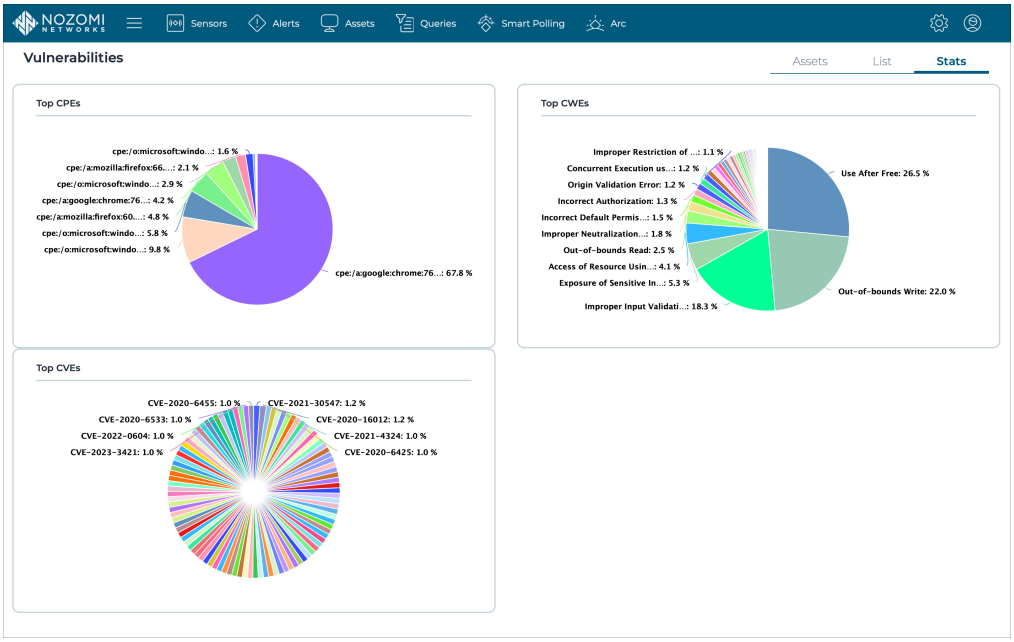


Figure 98. Stats page

Top CPEs

This section shows the:

- Title of vulnerability
- Percentage of the total vulnerabilities
- Actual count of that vulnerability

Top CWEs

This section shows the:

- Title of vulnerability
- Percentage of the total vulnerabilities
- Actual count of that vulnerability

Top CVEs

This section shows the:

- Title of vulnerability
- Date of vulnerability type
- Percentage of the total vulnerabilities
- Actual count of that vulnerability

Details page

The details page shows you all the details for the related vulnerability.

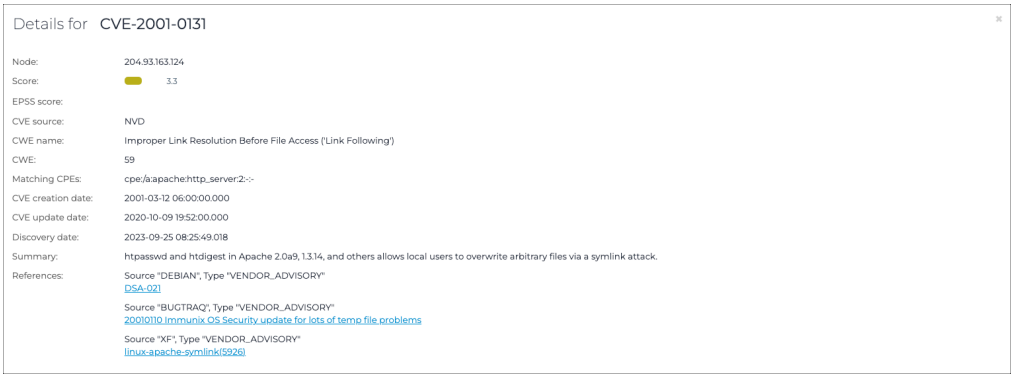


Figure 99. Details page



Chapter 15. Administration



Administration page

The administration page lets a user with administrator privileges configure settings and do other tasks.

For more details, see the **Guardian - Administrator Guide**.



Chapter 16. Personal settings



The personal settings page lets you do actions that are specific to your profile.

General

You can select the icon to access the personal settings menu.

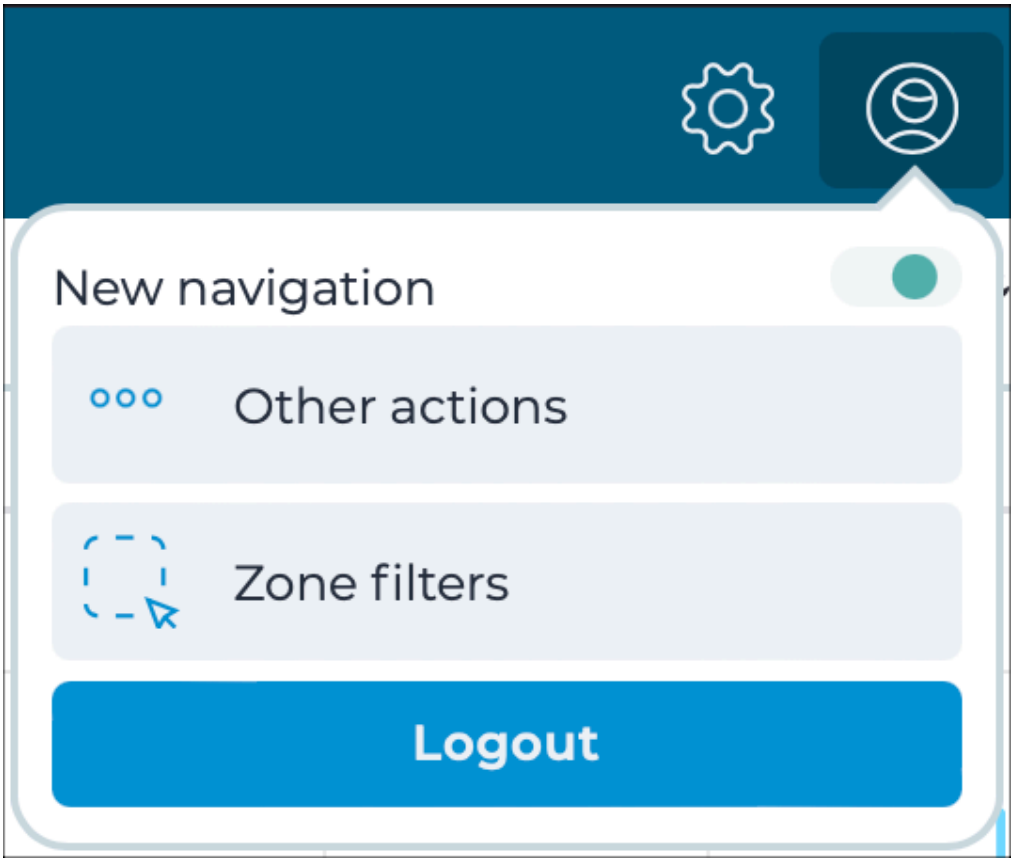


Figure 100. Personal settings menu

The personal settings page has these two sections:

- Other actions (on page 322)
- Zone filters (on page 333)

Logout

A button in the personal settings menu lets you log out of the software.

Other actions

*The **Other actions** page provides access to essential user settings, including options to clear personal settings, manage trace requests, and customize the UI theme with Auto, Light, or Dark modes.*

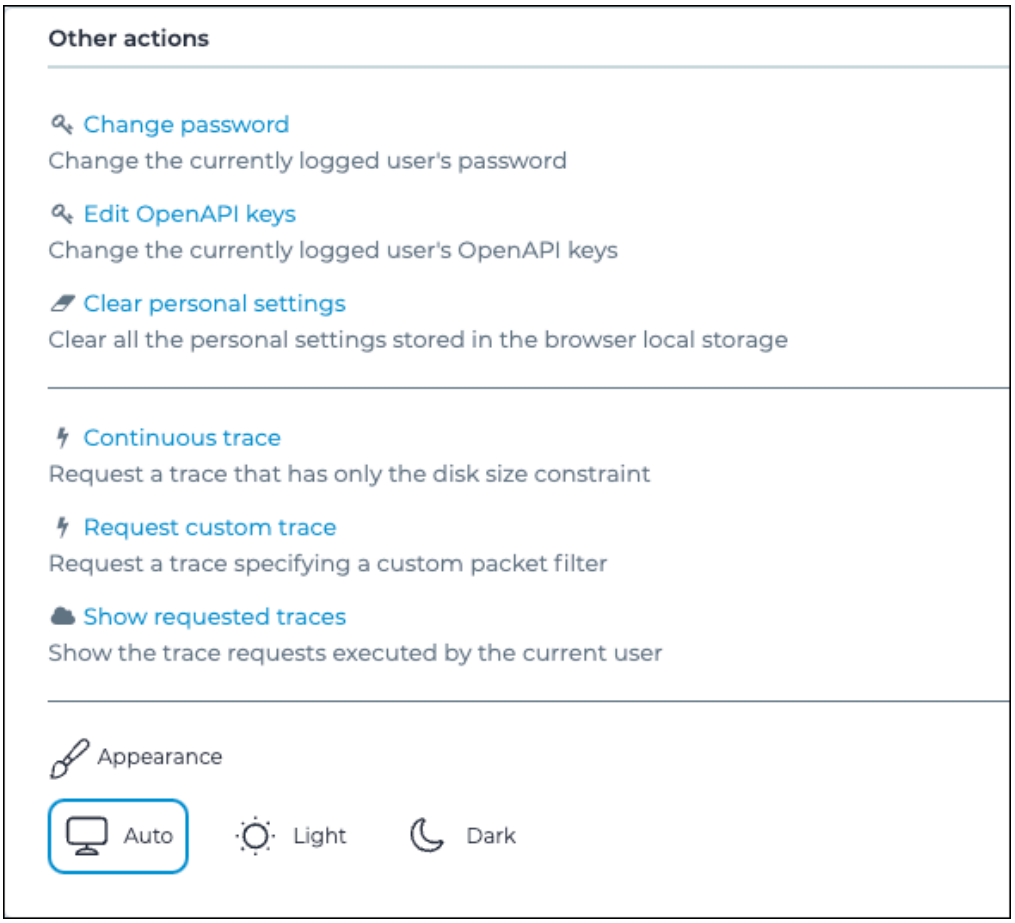


Figure 101. Other actions page

The other actions page lets you access these options:

- Change your password (on page 323)
- Edit an OpenAPI key (on page 324)
- Clear your personal settings (on page 327)
- Request a continuous trace (on page 328)
- Request a custom trace (on page 329)
- Show requested traces (on page 331)

The **Appearance** section lets you customize the visual appearance of the *UI*. You can choose one of these themes:

- Auto: Automatically adjusts the theme based on system settings
- Light: Uses a bright interface for better visibility in well-lit environments
- Dark: Uses a dark interface for reduced eye strain in low-light environments

Change your password

The profile settings menu lets you change your password.

Procedure

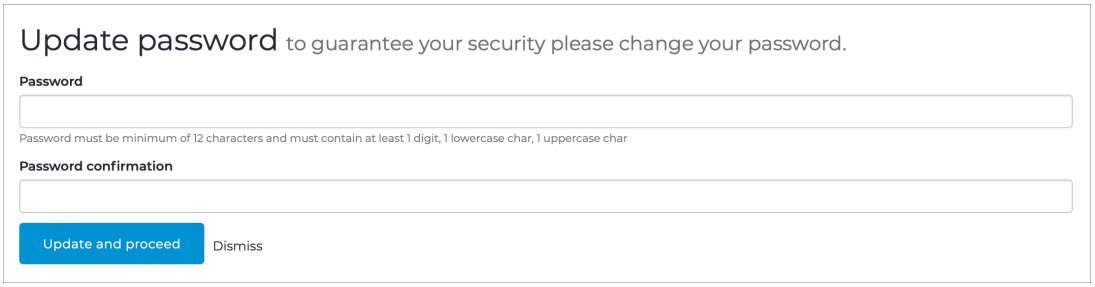
1. In the top navigation bar, select 

   Result: A menu shows.

2. Select **Other actions**.
3. Select **Change password**.

   Result: A dialog shows.

4. In the **Password** field, enter a new password.



5. In the **Password confirmation** field, enter the password again.
6. Select **Update and proceed**.

Results

Your password has been changed.

Edit an OpenAPI key

The profile settings menu lets you edit your OpenAPI keys.

Procedure

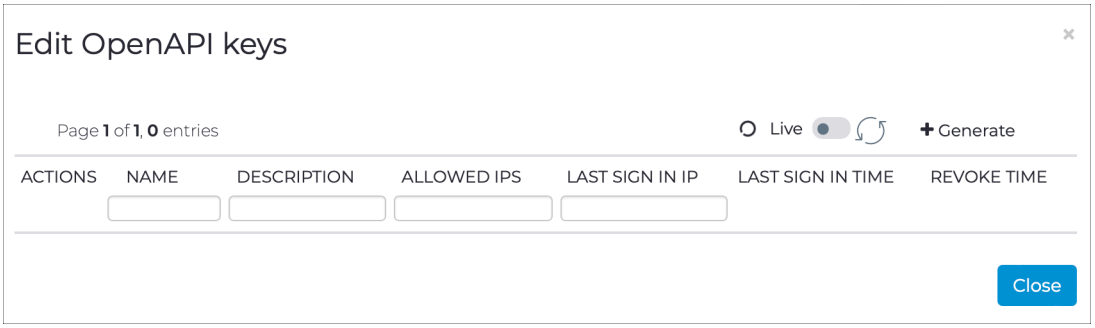
1. In the top navigation bar, select 

   Result: A menu shows.

2. Select **Other actions**.
3. Select **Edit OpenAPI keys**.

   Result: A dialog shows.

4. Edit the OpenAPI keys as necessary.



5. Select **Close**.

Results

Your OpenAPI keys have been edited.

Generate an OpenAPI key

The profile settings menu lets you generate an OpenAPI key.

Procedure

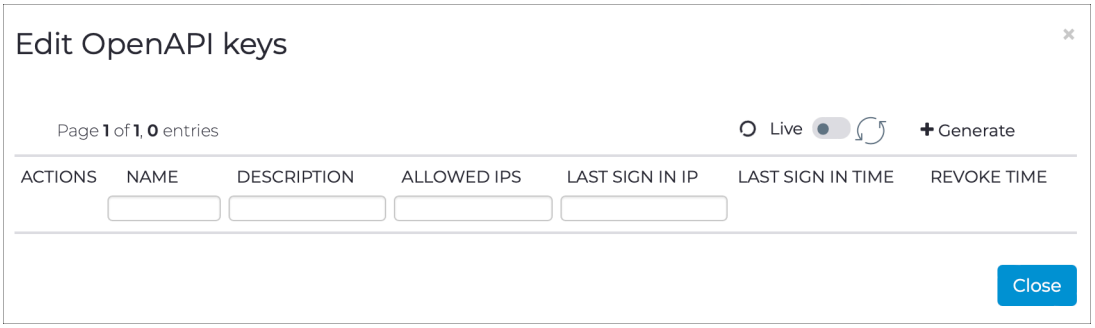
1. In the top navigation bar, select 

   Result: A menu shows.

2. Select **Other actions**.
3. Select **Edit OpenAPI keys**.

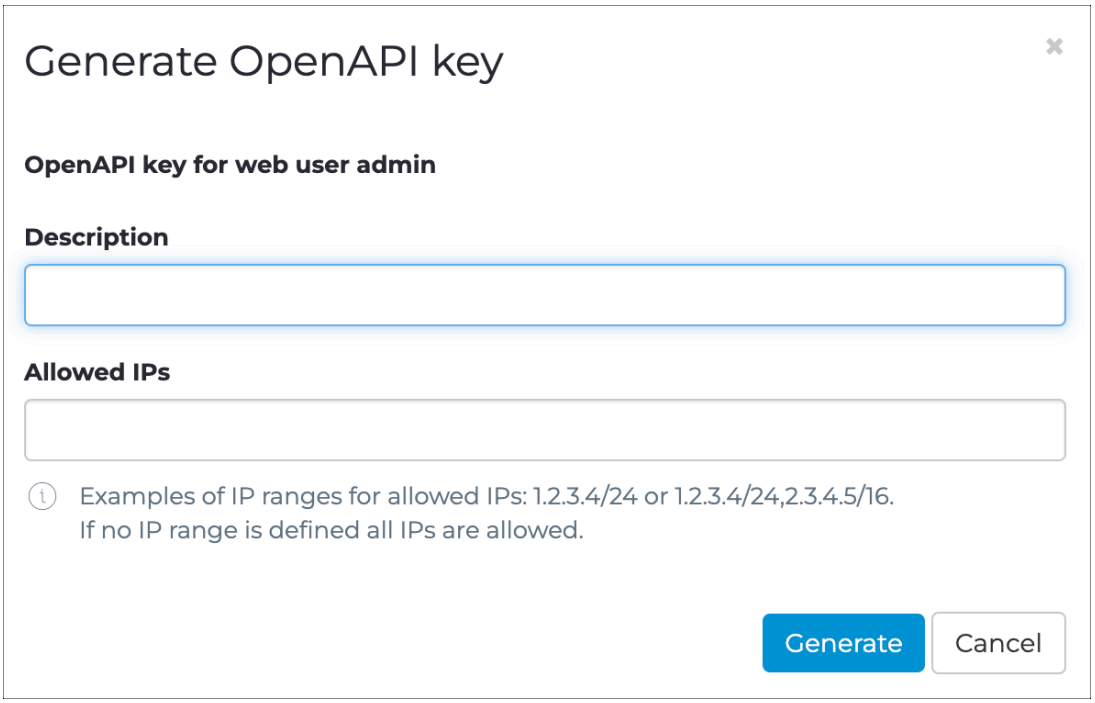
   Result: A dialog shows.

4. In the top right, select **+ Generate**.



   Result: A dialog shows.

5. In the **Description** field, enter a description for the OpenAPI key.



6. In the **Allowed IPs** field, enter the details of the allowed *IP* addresses.
7. Select **Generate**.

Results

Your OpenAPI key has been generated.

Clear your personal settings

The profile settings menu lets you clear your personal settings that are stored in the local browser local.

Procedure

1. In the top navigation bar, select 

   Result: A menu shows.

2. Select **Other actions**.
3. Select **Clear personal settings**.

   Result: A dialog shows.

4. Select **OK**.



Results

Your personal settings have been cleared from the local browser.

Request a continuous trace

The profile settings menu lets you request a trace that has only the disk size constraint.

About this task

The maximum count of packets, and the run time, constrain regular traces. For example, being constrained to capture at most 1000 packets in no more than one minute. A continuous trace does not have such limits. It collects packets without constraints, as long as the disk has free space to store the captured traffic. The packets are then saved as individual files that are 100 *MB* in size.

Procedure

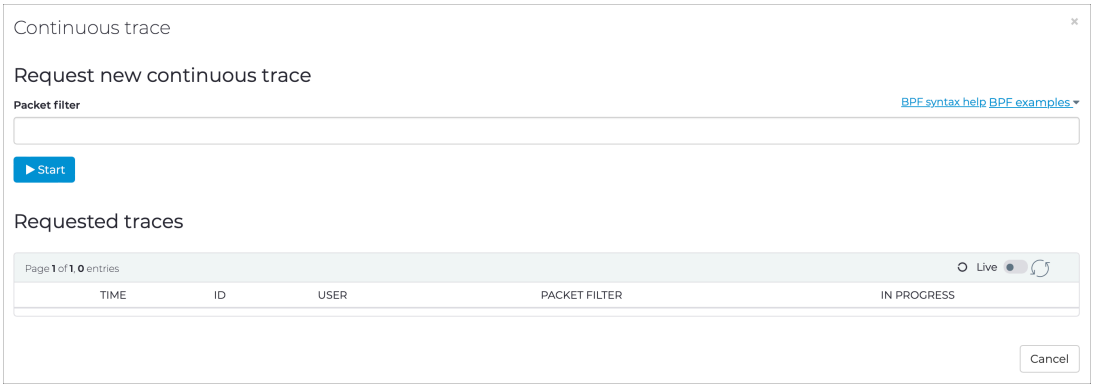
1. In the top navigation bar, select 

   Result: A menu shows.

2. Select **Other actions**.
3. Select **Continuous trace**.

   Result: A dialog shows.

4. In the **Packet filter** field, enter a *BPF* filter.



5. Select **Start**.

Results

The trace has been requested.

Request a custom trace

The profile settings menu lets you request a trace specifying a custom packet filter.

Procedure

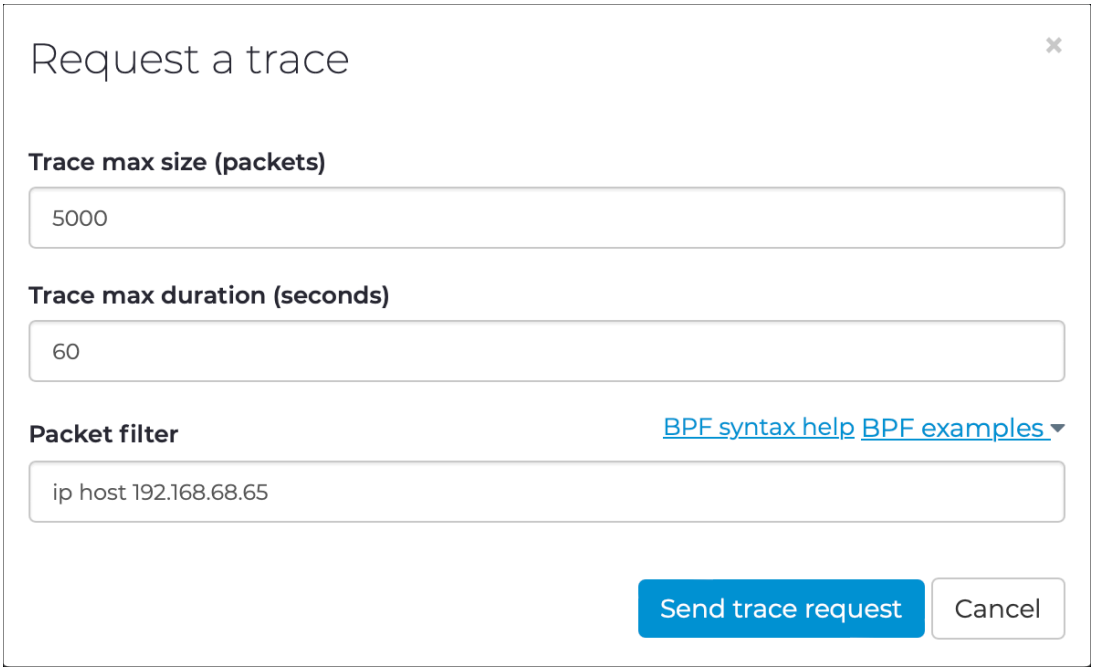
1. In the top navigation bar, select 

   Result: A menu shows.

2. Select **Other actions**.
3. Select **Request custom trace**.

   Result: A dialog shows.

4. To set the maximum packet size, in the **Trace max size (packets)** field, enter a value.



> **Note:**
> The default size is 5000 packets.

5. To set the maximum duration of the trace, in the **Trace max duration (seconds)** field, enter a value.

> **Note:**
> The default value is 60 seconds.

6. **Note:**
   The **Packet filter** field is automatically populated with a *BPF* that captures the packets to/from the selected node, but you can customize this.

   If necessary, customize this field.

   Note:
   You can select **BPF syntax help** to show more information on *BPF* syntax.

   Note:
   You can select **BPF examples** to see some examples.

7. Select **Send trace request**.

Results

The trace has been requested.

Show requested traces

The profile settings menu lets you show all the traces that you have requested.

Procedure

1. In the top navigation bar, select 

   Result: A menu shows.

2. Select **Other actions**.
3. Select **Show requested traces**.

Results

All your requested traces show.

Download a requested trace

Once you have requested a trace, and it is ready, you can download it.

Procedure

1. In the top navigation bar, select 

   Result: A menu shows.

2. Select **Other actions**.
3. Select **Show requested traces**.

   Result: All your requested traces show.

4. For the applicable trace, select **Download trace**.



Results

The trace has been downloaded.

Zone filters

The personal settings menu lets you filter zones.

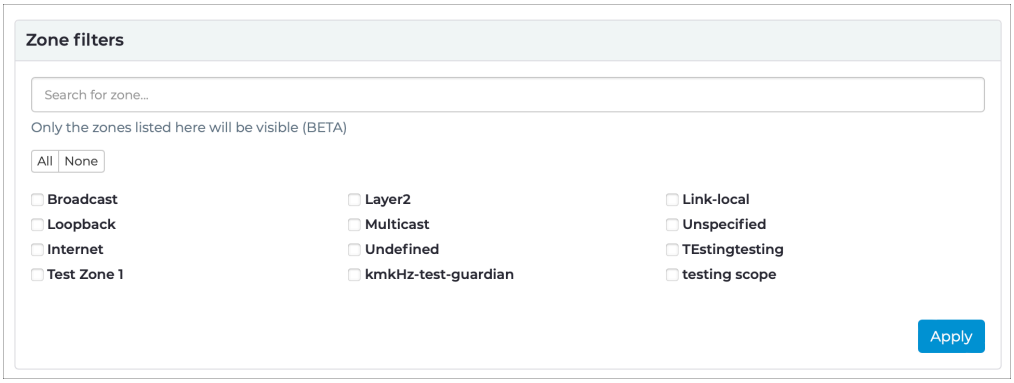


Figure 102. Zone filters page



Glossary



Address Resolution Protocol

ARP is a communication protocol that is used to discover the link layer address, such as a *MAC* address, that is associated with a given internet layer address. This is typically an IPv4 address.

Amazon Machine Image

An AMI is a type of virtual appliance that is used to create a virtual machine for the Amazon Elastic Compute Cloud (EC2), and is the basic unit of deployment for services that use EC2 for delivery.

Amazon Web Services

AWS is a subsidiary of the Amazon company that provides on-demand cloud computing platforms governments, businesses, and individuals on a pay-as-you-go basis.

Application Programming Interface

An API is a software interface that lets two or more computer programs communicate with each other.

Assertion Consumer Service

An ACS is a version of the SAML standard that is used to exchange authentication and authorization identities between security domains.

Asset Intelligence™

Asset Intelligence is a continuously expanding database of modeling asset behavior used by N2OS to enrich asset information, and improve overall visibility, asset management, and security, independent of monitored network data.

Berkeley Packet Filter

The BPF is a technology that is used in some computer operating systems for programs that need to analyze network traffic. A BPF provides a raw interface to data link layers, permitting raw link-layer packets to be sent and received.

Central Management Console

The Central Management Console (CMC) is a Nozomi Networks product that has been designed to support complex deployments that cannot be addressed with a single sensor. A central design principle behind the CMC is the unified experience, that lets you access information in a similar way as on the sensor.

Central Processing Unit

The main, or central, processor that executes instructions in a computer program.

Certificate Authority

A certificate, or certification authority (CA) is an organization that stores, signs, and issues digital certificates. In cryptography, a digital certificate certifies the ownership of a public key by the named subject of the certificate.

Classless Inter-Domain Routing

CIDR is a method for IP routing and for allocating IP addresses.

Command-line interface

A command-line processor uses a command-line interface (CLI) as text input commands. It lets you invoke executables and provide information for the actions that you want them to do. It also lets you set parameters for the environment.

Comma-separated Value

A CSV file is a text file that uses a comma to separate values.

Common Event Format

CEF is a text-based log file format that is used for event logging and information sharing between different security devices and software applications.

Common Platform Enumeration

CPE is a structured naming scheme for information technology (IT) systems, software, and packages. CPE is based on the generic syntax for Uniform Resource Identifiers (URI) and includes a formal name format, a method for checking names against a system, and a description format for binding text and tests to a name.

Common Vulnerabilities and Exposures

CVEs give a reference method information-security vulnerabilities and exposures that are known to the public. The United States' National Cybersecurity FFRDC maintains the system.

Common Weakness Enumeration

CWE is a category system for software and hardware weaknesses and vulnerabilities. It is a community project with the aim to understand flaws in software and hardware and create automated tools that can be used to identify, fix, and prevent those flaws.

Configuration file

A CFG file is a configuration, or config, file. They are files that are used to configure the parameters and initial settings for a computer program.

cURL

cURL is a command-line tool and library used to transfer data to or from a server. It supports many protocols, including HTTP, HTTPS, FTP, and more. In API contexts, curl is commonly used to test and interact with endpoints by sending requests and receiving responses.

Domain Name Server

The DNS is a distributed naming system for computers, services, and other resources on the Internet, or other types of Internet Protocol (IP) networks.

ESXi

VMware ESXi (formerly ESX) is an enterprise-class, type-1 hypervisor developed by VMware for deploying and serving virtual computers. As a type-1 hypervisor, ESXi is not a software application that is installed on an operating system (OS). Instead, it includes and integrates vital OS components, such as a kernel.

Extensible Markup Language

XML is a markup language and file format for the storage and transmission of data. It defines a set of rules for encoding documents in a format that is both human-readable and machine-readable.

Federal Information Processing Standards

FIPS are publicly announced standards developed by the National Institute of Standards and Technology for use in computer systems by non-military American government agencies and government contractors.

File Allocation Table

FAT is a file system architecture used in computers for managing disk space. It maintains a table to track the allocation of files on a disk, supporting efficient data storage, access, and management

File Transfer Protocol

FTP is a standard communication protocol that is used for the transfer of computer files from a server to a client on a computer network. FTP is built on a client–server model architecture that uses separate control and data connections between the client and the server.

Fully qualified domain name

An FQDN is a complete and specific domain name that specifies the exact location in the hierarchy of the Domain Name System (DNS). It includes all higher-level domains, typically consisting of a host name and domain name, and ends in a top-level domain.

Garbage Collector

The Garbage Collector is a system process that discards nodes, assets, and links that are no longer useful, saving system resources.

Gigabit per second

Gigabit per second (Gb/s) is a unit of data transfer rate equal to: 1,000 Megabits per second.

Gigabyte

The gigabyte is a multiple of the unit byte for digital information. One gigabyte is one billion bytes.

Graphical User Interface

A GUI is an interface that lets humans interact with electronic devices through graphical icons.

Graphics Interchange Format

GIF is a bitmap image format that is widely used on the internet.

High Availability

High Availability is a mode that permits the CMC to replicate its own data on another CMC.

Host-based intrusion-detection system

HIDS is an internal Nozomi Networks solution that uses sensors to detect changes to the basic firmware image, and record the change.

Hypertext Markup Language

Hypertext Markup Language (HTML) is the standard markup language used to structure and display content on the web. It defines the structure of web pages using elements represented by tags.

Hypertext Transfer Protocol

HTTP is an application layer protocol in the Internet protocol suite model for distributed, collaborative, hypermedia information systems. HTTP is the foundation of data communication for the World Wide Web, where hypertext documents include hyperlinks to other resources that the user can easily access, for example by a mouse click or by tapping the screen in a web browser.

Hypertext Transfer Protocol Secure

HTTPS is an extension of the Hypertext Transfer Protocol (HTTP). It is used for secure communication over a computer network, and is widely used on the Internet. In HTTPS, the communication protocol is encrypted using Transport Layer Security (TLS) or, formerly, Secure Sockets Layer (SSL). The protocol is therefore also referred to as HTTP over TLS, or HTTP over SSL.

Identifier

A label that identifies the related item.

Identity Provider

An IdP is a system entity that creates, maintains, and manages identity information. It also provides authentication services to applications within a federation, or a distributed network.

Industrial Control Systems

An ICS is an electronic control system and related instrumentation that is used to control industrial processes.

Industrial Internet of Things

The IIoT is a name for interconnected devices, sensors, instruments, which are networked together with industrial applications. This connectivity allows for analysis and data collection, which can facilitate improvements in efficiency and productivity.

Information Object Addresses

An Information Object Address (IOA) is a unique identifier used in IEC 60870-5-101 and IEC 60870-5-104 protocols to specify individual data points, such as sensors or actuators, within a telecontrol system. The IOA format can vary and is configurable, typically appearing as either a 1-byte–2-byte or a 1-byte–1-byte–1-byte structure, depending on the system requirements.

Information Technology

IT is the use of computers to process, create, store, and exchange data and information.

Internet Control Message Protocol

ICMP is a supporting protocol in the internet protocol suite. Network devices use it to send error messages and operational information to indicate success or failure when communicating with another IP address. ICMP differs from transport protocols such as TCP and UDP in that it is not typically used to exchange data between systems.

Internet of Things

The IoT describes devices that connect and exchange information through the internet or other communication devices.

Internet Protocol

An Internet Protocol address, or IP address, identifies a node in a computer network that uses the Internet Protocol to communicate. The IP label is numerical.

Intrusion Detection System

An intrusion detection system (IDS), which can also be known as an intrusion prevention system (IPS) is a software application, or a device, that monitors a computer network, or system, for malicious activity or policy violations. Such intrusion activities, or violations, are typically reported either to a system administrator, or collected centrally by a security information and event management (SIEM) system.

IPv6

Internet Protocol version 6 is the most recent version of the Internet Protocol, designed to replace IPv4. It uses 128-bit numerical addresses to identify devices, allowing for a vastly larger address space.

JavaScript Object Notation

JSON is an open standard file format for data interchange. It uses human-readable text to store and transmit data objects, which consist of attribute–value pairs and arrays.

Joint Photographic Experts Group

JPEG, or JPG, is a method of lossy compression that is used for digital images. The degree of compression can be adjusted, allowing a selectable tradeoff between storage size and image quality.

Lightweight Directory Access Protocol

LDAP is an open, vendor-neutral, industry standard application protocol that lets you access and maintain distributed directory information services over an internet protocol (IP) network.

Lightweight Directory Access Protocol Secure

LDAP over SSL or Secure LDAP is the secure version of LDAP.

Link Layer Discovery Protocol

LLDP is a vendor-neutral link layer protocol used by network devices for advertising identity, capabilities, and neighbors on a local area network.

Managed Security Service Provider

In computing, MSSP are network security services that have been outsourced to a service provider. A company that provides this type of service is known as an MSSP.

Management Information Base

An MIB is a collection of definitions that specify the properties of the managed resources within a network device (like routers or switches) and how they can be accessed using SNMP (Simple Network Management Protocol).

Media Access Control

A MAC address is a unique identifier for a network interface controller (NIC). It is used as a network address in network segment communications. A common use is in most IEEE 802 networking technologies, such as Bluetooth, Ethernet, and Wi-Fi. MAC addresses are most commonly assigned by device manufacturers and are also referred to as a hardware address, or physical address. A MAC address normally includes a manufacturer's organizationally unique identifier (OUI). It can be stored in hardware, such as the card's read-only memory, or by a firmware mechanism.

Megabyte

The megabyte is a multiple of the unit byte for digital information. One megabyte is one million bytes.

Message Queuing Telemetry Transport

Message Queuing Telemetry Transport is a lightweight, publish-subscribe network protocol designed for constrained devices and low-bandwidth, high-latency, or unreliable networks. It is commonly used in Internet of Things (IoT) applications.

Mobile Device Management

This is the administration of mobile devices, such as tablet computers, laptops, and smartphones. It is often implemented with a third-party product with features for specific vendors of mobile devices.

National Institute of Standards and Technology

NIST is an agency of the United States Department of Commerce. NIST's mission is to promote American innovation and industrial competitiveness.

National Vulnerability Database

The National Vulnerability Database (NVD) is the U.S. government repository of standards-based vulnerability management data represented using the Security Content Automation Protocol (SCAP). This data enables automation of vulnerability management, security measurement, and compliance. NVD includes databases of security checklists, security related software flaws, misconfigurations, product names, and impact metrics.

Network Address Translation

NAT is a method of mapping an internet protocol (IP) address space into another one. This is done by modifying network address information in the IP header of packets while in transit across a traffic routing device.

Network-Attached Storage

NAS is a dedicated file storage system that provides local-area network (LAN) users with centralized, shared access to data through standard protocols such as NFS or SMB.

Network Interface Controller

A network interface controller (NIC), sometimes known as a network interface card, is a computer hardware component that lets a computer connect to a computer network.

Network Time Protocol

The NTP is a networking protocol to synchronize clocks between computer systems over variable-latency, packet-switched data networks.

Newline Delimited JSON

NDJSON is a data format where each line is a valid JSON value, used for streaming or storing structured records efficiently.

Nozomi Networks Operating System

N2OS is the operating system that the core suite of Nozomi Networks products runs on.

Nozomi Networks Query Language (N2QL)

N2QL is the language used in queries in Nozomi Networks software.

Object Identifier

In computing, OIDs are an identifier mechanism which is standardized by the International Telecommunication Union (ITU) and ISO/IEC for the naming of an object, concept, or other item with a globally-unambiguous persistent name.

Open Virtual Appliance

An OVA file is an open virtualization format (OVF) directory that is saved as an archive using the .tar archiving format. It contains files for distribution of software that runs on a virtual machine. An OVA package contains a .ovf descriptor file, certificate files, an optional .mf file along with other related files.

Operating System

An operating system is computer system software that is used to manage computer hardware, software resources, and provide common services for computer programs.

Operational Technology

OT is the software and hardware that controls and/or monitors industrial assets, devices and processes.

Packet Capture

A pcap is an application programming interface (API) that captures live network packet data from the OSI model (layers 2-7).

Packet Capture Next Generation

A pcapNg is the latest version of a pcap file, an application programming interface (API) that captures live network packet data from the OSI model (layers 2-7).

Portable Document Format

PDF is a Adobe file format that is used to present documents. It is independent of operating systems (OS), application software, hardware.

Portable Network Graphics

PNG is a raster graphics file format that supports lossless data compression.PNG was developed as an improved, non-patented replacement for graphics interchange format (GIF).

Privacy-Enhanced Mail

PEM is a standard file format that is used to store and send cryptographic keys, certificates, and other data. It is based on a set of 1993 IETF standards.

Programmable Logic Controller

A PLC is a ruggedized, industrial computer used in industrial and manufacturing processes.

Protected Extensible Authentication Protocol

PEAP is a protocol that encloses the Extensible Authentication Protocol (EAP) within an encrypted and authenticated Transport Layer Security (TLS) tunnel.

Random-access Memory

Computer memory that can be read and changed in any order. It is typically used to store machine code or working data.

Remote Terminal Unit

An RTU is a microprocessor-controlled electronic device that acts as an interface between a SCADA (supervisory control and data acquisition) system, or distributed control system, to a physical object. It transmits telemetry data to a master system, and uses messages from the master supervisory system to control connected objects.

Representational State Transfer

Representational State Transfer (REST) is an architectural style for designing networked applications. It uses stateless, client-server communication via standard HTTP methods (GET, POST, PUT, DELETE) to access and manipulate web resources represented in formats like JSON or XML.

Seamless Message Protocol

Seamless Message Protocol is a communication protocol used primarily in industrial automation for device-level communication. It enables data exchange between controllers and devices over Ethernet networks.

Secure Copy Protocol

SCP is a protocol for the secure transfer of computer files between a local host and a remote host, or between two remote hosts. It is based on the secure shell (SSH) protocol.

Secure Shell

A cryptographic network protocol that let you operate network services securely over an unsecured network. It is commonly used for command-line execution and remote login applications.

Secure Sockets Layer

A secure sockets layer ensures secure communication between a client computer and a server.

Security Assertion Markup Language

SAML is an open standard, XML-based markup language for security assertions. It allows for the exchange of authentication and authorization data different parties such as a service provider and an identity provider.

Security Information and Event Management

SIEM is a field within the computer security industry, where software products and services combine security event management (SEM) and security information management (SIM). SIEMs provide real-time analysis of security alerts.

Server Message Block

Is a communication protocol which provides shared access to files and printers across nodes on a network of systems. It also provides an authenticated interprocess communication (IPC) mechanism.

Simple File Transfer Protocol

SFTP was proposed as an unsecured file transfer protocol with a level of complexity intermediate between TFTP and FTP. It was never widely accepted on the internet.

Simple Mail Transfer Protocol

SMTP is an internet standard communication protocol that is used for the transmission of email. Mail servers and other message transfer agents use SMTP to send and receive mail messages.

Simple Network Management Protocol

SNMP is an Internet Standard protocol for the collection and organization of information about managed devices on IP networks. It also lets you modify that information to change device behavior. Typical devices that support SNMP are: printers, workstations, cable modems, switches, routers, and servers.

Simple Text Oriented Messaging Protocol

STOMP is a simple text-based protocol, for working with message-oriented middleware (MOM). It provides an interoperable wire format that allows STOMP clients to talk with any message broker supporting the protocol.

Single Sign-on

SSO is an authentication method that lets users log in to one or more related, but independent, software systems.

Spanning Tree Protocol

Spanning Tree Protocol is a network protocol that prevents loops in Ethernet networks by creating a spanning tree topology. It selectively blocks redundant paths and ensures a loop-free logical topology.

Structured Threat Information Expression

STIX™ is a language and serialization format for the exchange of cyber threat intelligence (CTI). STIX is free and open source.

Supervisory control and data acquisition

SCADA is a control system architecture which has computers, networked data communications and graphical user interfaces for high-level supervision of processes and machines. It also covers sensors and other devices, such as programmable logic controllers (PLC), which interface with process plant or machinery.

Text-based User Interface

In computing, a text-based (or terminal) user interfaces (TUI) is a retronym that describes a type of user interface (UI). These were common as an early method of human–computer interaction, before the more modern graphical user interfaces (GUIs) were introduced. Similar to GUIs, they might use the entire screen area and accept mouse and other inputs.

Thick Provisioning (Static)

Thick provisioning, also known as fat provisioning, is a method of allocating a fixed amount of physical storage to a virtual disk at creation. The full storage space is reserved immediately, regardless of how much data is actually stored.

Thin Provisioning (Dynamic)

Thin provisioning is a method of allocating storage space on a just-in-time basis as data is written, rather than reserving the full amount upfront. This approach makes more efficient use of available storage capacity.

Threat Intelligence™

Nozomi Networks **Threat Intelligence™** feature monitors ongoing OT and IoT threat and vulnerability intelligence to improve malware anomaly detection. This includes managing packet rules, YARA rules, STIX indicators, Sigma rules, and vulnerabilities. **Threat Intelligence™** allows new content to be added, edited, and deleted, and existing content to be enabled or disabled.

Transmission Control Protocol

One of the main protocols of the Internet protocol suite.

Transport Layer Security

TLS is a cryptographic protocol that provides communications security over a computer network. The protocol is widely used in applications such as: HTTPS, voice over IP, instant messaging, and email.

Uniform Resource Identifier

A URI is a unique string of characters used to identify a logical or physical resource on the internet or local network.

Uniform Resource Locator

An URL is a reference to a resource on the web that gives its location on a computer network and a mechanism to retrieving it.

Uninterruptible Power Supply

A UPS is an electric power system that provides continuous power. When the main input power source fails, an automated backup system continues to supply power.

Universally unique identifier

A UUID is a 128-bit label that is used for information in computer systems. When a UUID is generated with standard methods, they are, for all practical purposes, unique. Their uniqueness is not dependent on an authority, or a centralized registry. While it is not impossible for the UUID to be duplicated, the possibility is generally considered to be so small, as to be negligible. The term globally unique identifier (GUID) is also used in some, mostly Microsoft, systems.

Universal Plug and Play

UPnP is a network protocol that enables devices to automatically discover and communicate with each other using broadcast messages. While it facilitates easy device identification and connectivity, UPnP lacks robust authentication, making it vulnerable to unauthorized access in cybersecurity contexts.

Universal Serial Bus

Universal Serial Bus (USB) is a standard that sets specifications for protocols, connectors, and cables for communication and connection between computers and peripheral devices.

User Datagram Protocol

UDP is a lightweight, connectionless communication protocol used for fast, time-sensitive data transmission, such as video streaming, online gaming, and VoIP. UDP prioritizes speed and low latency over guaranteed delivery or error correction.

User Interface

An interface that lets humans interact with machines.

Variable

In the context of control systems, a variable can refer to process values that change over time. These can be temperature, speed, pressure etc.

Virtual DOM

A virtual DOM, or vdom, is a lightweight JavaScript representation of the Document Object Model (DOM). It is used in declarative web frameworks such as Elm, React, and Vue.js. It enables the updating of the virtual DOM is comparatively faster than updating the actual DOM.

Virtual Hard Disk

VHD is a file format that represents a virtual hard disk drive (HDD). They can contain what is found on a physical HDD, such as disk partitions and a file system, which in turn can contain files and folders. They are normally used as the hard disk of a virtual machine (VM). They are the native file format for Microsoft's hypervisor (virtual machine system), Hyper-V.

Virtual Local Area Network

A VLAN is a broadcast domain that is isolated and partitioned in a computer network at the data link layer (OSI layer 2).

Virtual Machine

A VM is the emulation or virtualization of a computer system. VMs are based on computer architectures and provide the functionality of a physical computer.

Windows Management Instrumentation

WMI is the Microsoft implementation of Web-Based Enterprise Management (WBEM), an industry initiative to develop a standard technology for accessing management information in an enterprise environment. WMI uses the Common Information Model (CIM) industry standard to represent systems, applications, networks, devices, and other managed components. It allows querying and management of Windows systems.

Windows Remote Management

Is a Microsoft implementation of Web Services-Management in Windows which lets systems access or exchange management information across a common network. You can utilize the built-in command line tool, or scripting objects, WinRM can be used with remote computers that may have baseboard management controllers (BMCs) to acquire data.

YARA

YARA is a tool designed to help identify and classify malware samples by describing them in terms of textual or binary patterns. YARA rules are used to detect files, processes, or network traffic that match known malware signatures.

Zettabyte File System

A file system and logical volume manager that combines storage pooling with data integrity features. N2OS uses ZFS for the `/data` file system.

ZIP

An archive file format that supports lossless data compression. The format can use a number of different compression algorithms, but DEFLATE is the most common one. A ZIP file can contain one or more compressed files or directories.