NOZOMI NETWORKS



Guardian
Administrator 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 **5**
Guardian overview 7
Chapter 2. Administration **9**
Administration page 11
Settings 13
Security control panel 13
Features 28
Users management 35
CLI 68
Dashboards 69
Threat Intelligence 77
Custom fields 83
Discovery 84
Data integration 85
Firewall integration 134
Credentials manager 157
Zone configurations 160
Synchronization settings 164
Alert playbooks 181
Update service contents 187
System 189
General 189
Date and time 190
Updates and licenses 192
Network interfaces 196
Upload traces 206
Export 212
Import 216
Health 229
Audit 233
Data 234
Migration tasks 236
Operations 238
Support 240
Backup and restore 241
Glossary 245



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. Administration



Administration page

The administration page lets a user with administrator privileges configure settings and do other tasks.

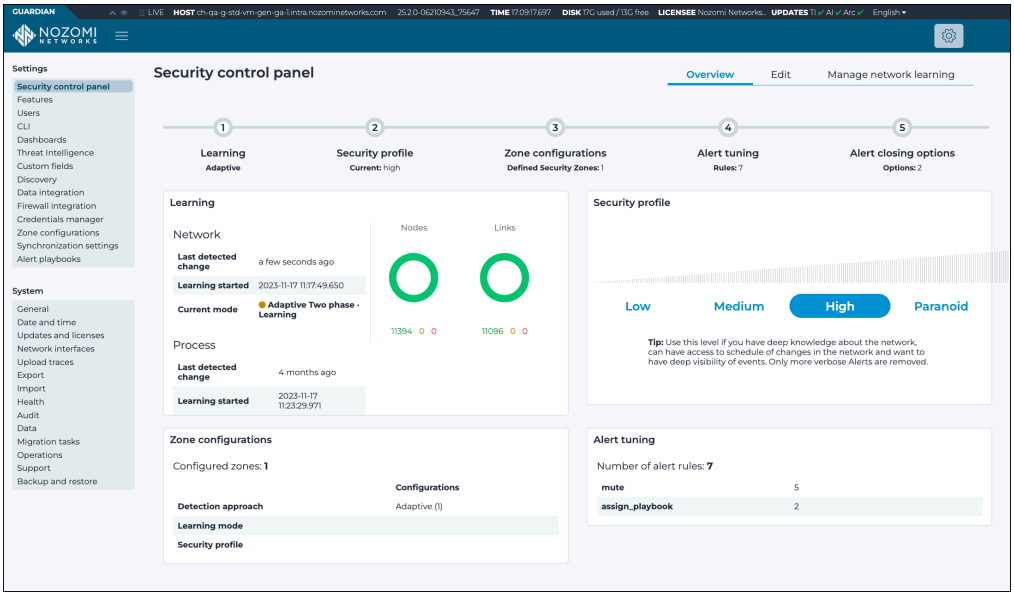


Figure 1. Administration page

Settings

The **Settings** section has these pages:

- Security control panel (on page 13)
- Features (on page 28)
- Users management (on page 35)
- CLI (on page 68)
- Dashboards (on page 69)
- Threat Intelligence (on page 77)
- Custom fields (on page 83)
- Discovery (on page 84)
- Data integration (on page 86)
- Firewall integration (on page 134)
- Credentials manager (on page 157)
- Zone configurations (on page 160)
- Synchronization settings (on page 164)
- Alert playbooks (on page 182)

System

The **System** section has these pages:

- General (on page 189)
- Date and time (on page 190)
- Updates and licenses (on page 192)
- Network interfaces (on page 196)
- Upload traces (on page 206)
- Export (on page 212)
- Import (on page 216)
- Health (on page 229)
- Audit (on page 233)
- Data (on page 234)
- Migration tasks (on page 236)
- Operations (on page 238)
- Support (on page 240)
- Backup and restore (on page 241)

Settings

Security control panel

*The **Security control panel** page shows an overview of the current status of the learning process and lets you configure the features that manage the: learning, security profile, zones, alerts tuning, and alert closing options.*

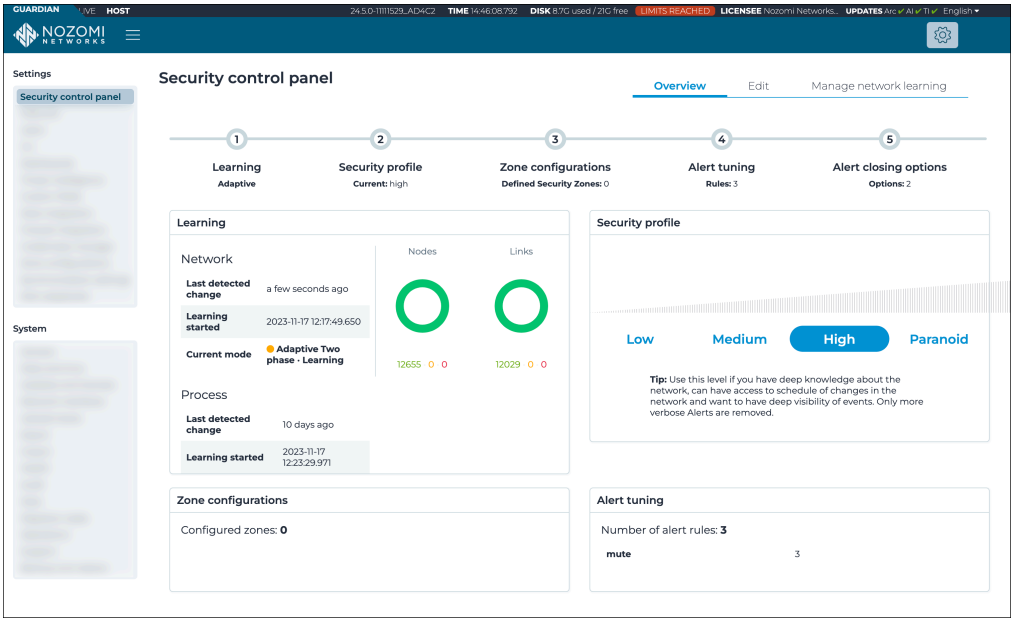


Figure 2. Security control panel page

The **Security control panel** page has these tabs:

- Overview (on page 14)
- Edit (on page 15)
- Manage network learning (on page 25)

Overview

*The **Overview** page shows an overview of the learning status, security profile chosen, zone configurations, and alert tuning rules configured.*

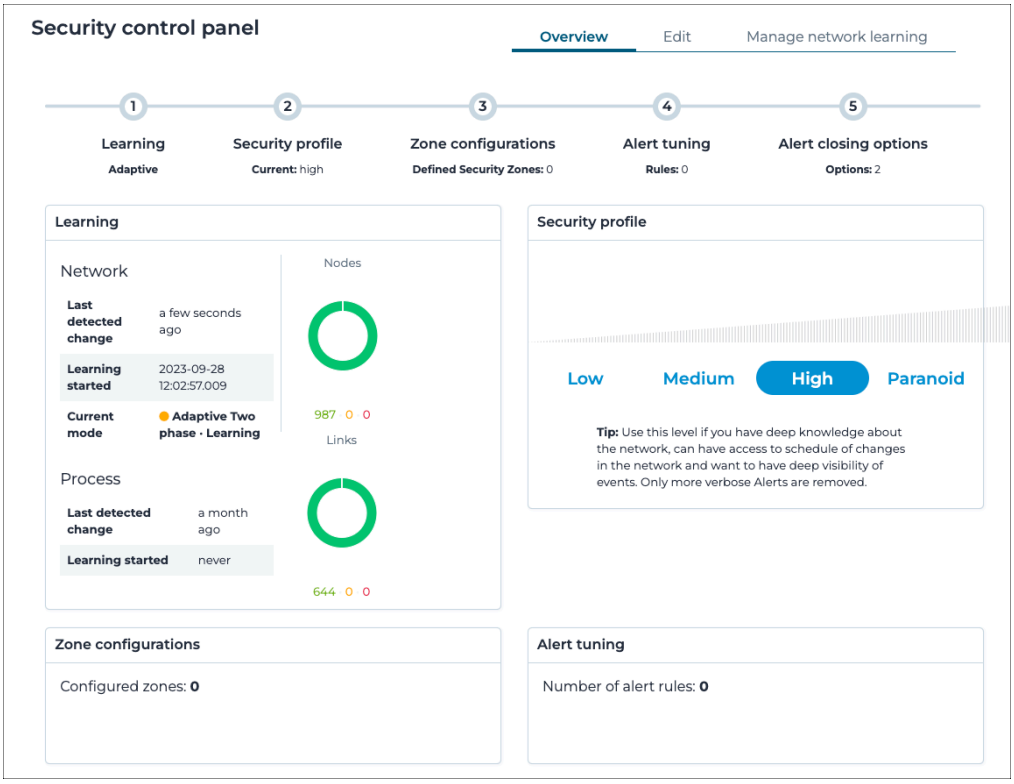


Figure 3. Overview page

The **Overview** page shows an overview for each of these pages:

- Learning (on page 16)
- Security profile (on page 19)
- Zone configurations (on page 20)
- Alert tuning (on page 21)
- Alert closing options (on page 24)

Edit

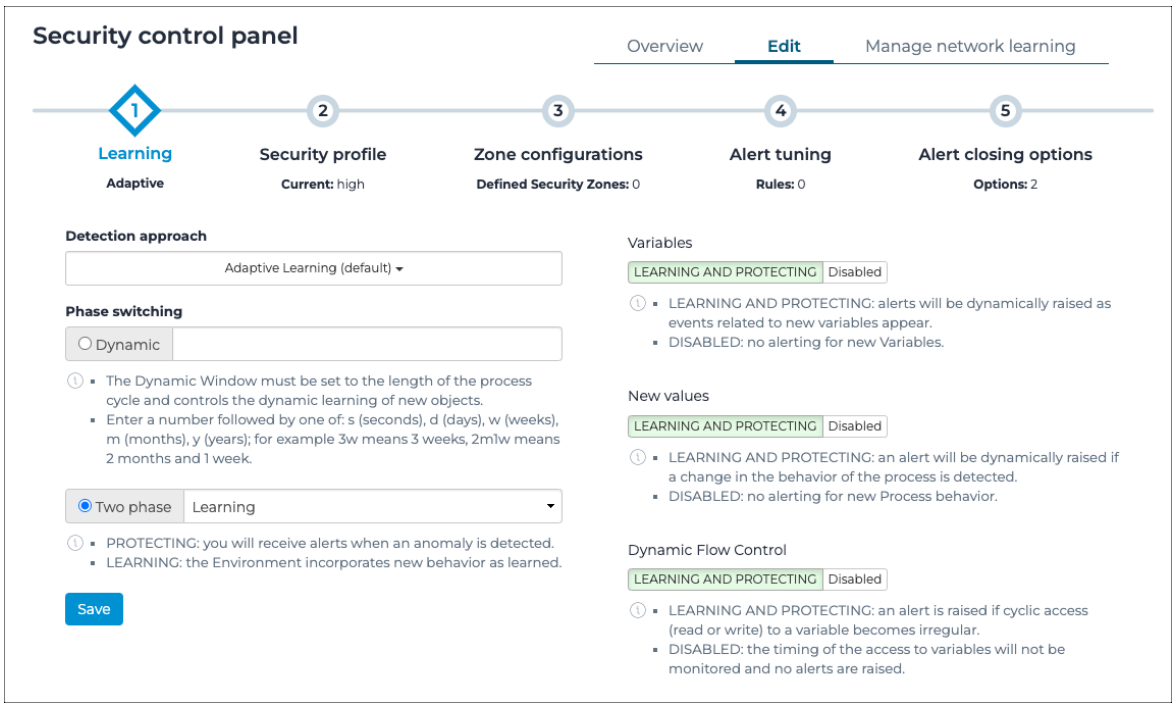
*The **Edit** page lets you configure the security features through simple steps.*

The **Edit** page has these pages:

- Learning (on page 16)
- Security profile (on page 19)
- Zone configurations (on page 20)
- Alert tuning (on page 21)
- Alert closing options (on page 24)

Learning

*The **Learning** page lets you manage how the typical behavior and components of your environment are learned. The software needs to learn the normal processes, patterns and communication of your environment in order for it to be able to detect anomalies.*



Learning

Detection approach

This dropdown lets you choose from these options:

- Adaptive Learning (default)
- Strict

Adaptive Learning uses a less granular and more scalable approach to anomaly detection where deviations are evaluated at a global level, rather than at a single node level. For example, the addition of a device similar to the ones already installed in the learned network will not produce alerts. This holds true for the appearance of a similar communication. Adaptive Learning shows its maximum capabilities when combined with Asset Intelligence.

Strict uses a detailed anomaly-based approach, so that deviations from the baseline will be detected and alerted. This approach is called strict because it requires the learned system to behave like it has behaved during the learning phase, and requires some knowledge of the monitored system in order to be maintained over time.

The engine has two distinct learning goals: the network and the process. For both cases the engine can be in learning and in protection mode, and they can be governed independently.

Network Learning is the learning of:

- Nodes
- Links
- Protocols
- Function Codes for example, commands, that are sent from one node to another

A wide range of parameters is checked in this engine and can be fine-tuned.
Process Learning is the learning of variables, and their behavior. You can use specific checks to fine-tune this learning.

Phase switching

You can set this to one of these options:

- Dynamic
- Two-phase

Dynamic, or Dynamic window, lets you configure the time interval in which an engine considers a change to be learned. Every engine does this kind of evaluation **per node** and **per network segment**.

After this period of time, the learning phase is automatically, and safely, switched to protection mode. This:

- Raises alerts when something is different from the learned baseline
- Adds suspicious components to the environment with the `is learned` attribute set to off, in such a way that an operator can confirm, delete, or take the appropriate action in the Manage network learning (on page 25) page

In this way, stable network nodes and segments become protected automatically. This prevents you from being overwhelmed with alerts due to the premature closing of learning mode.

The **Two phase** dropdown lets you choose between:

- Learning
- Protecting

Learning: in this mode, the environment incorporates new behavior as learned.

Protecting: in this mode, you will receive alerts when an anomaly is detected

Variables

If you set the toggle to **Learning and Protecting**, an alert will be raised if a new *variable* is detected. You can set the toggle to **Disabled**, if you want to prevent this happening. For example, when you are making planned changes to your environment.

New values

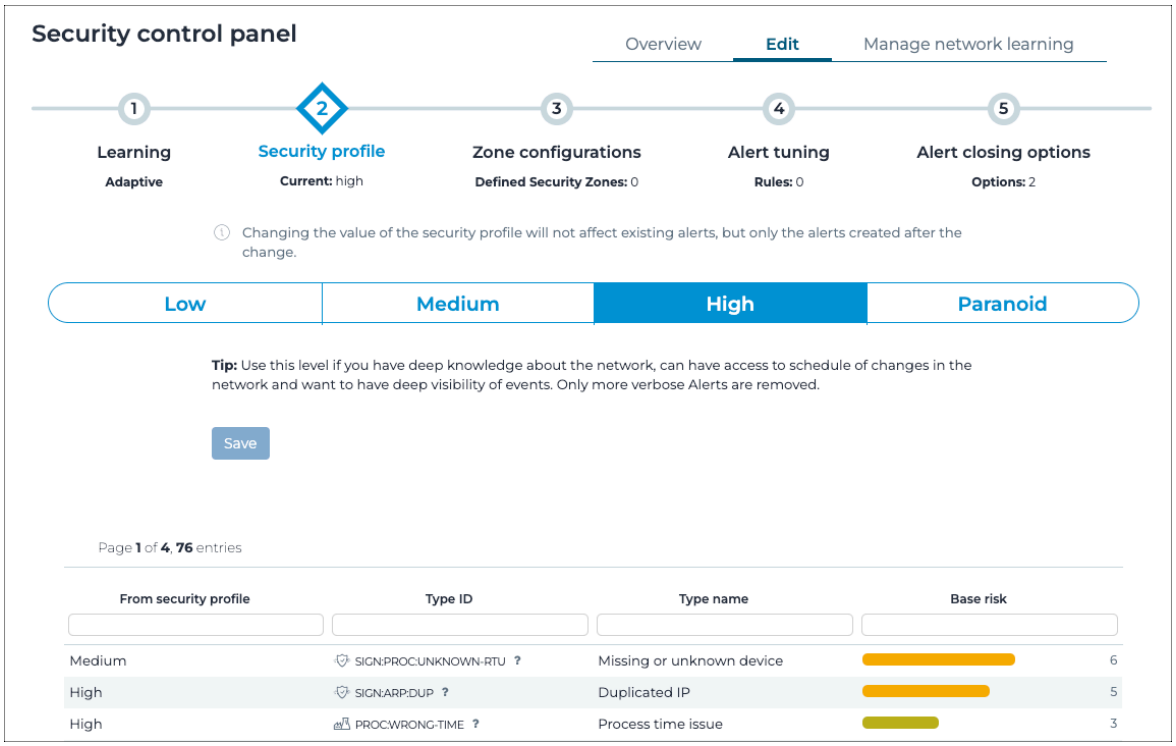
If you set the toggle to **Learning and Protecting**, an alert will be dynamically raised if a change in the behavior of the process is detected. You can set the toggle to **Disabled**, if you want to prevent this happening. For example, when you are making planned changes to your environment.

Dynamic flow control

If you set the toggle to **Learning and Protecting**, an alert will be raised if reading and writing patterns of a *Variable* are detected. You can set the toggle to **Disabled**, if you want to prevent this happening. For example, when you are making planned changes to your environment.

Security profile

*The **Security profile** page lets you change the visibility of alerts based on their type.*



General

When you change the value of the **Security Profile**, it has an immediate affect on newly-generated alerts, and it has no effect on existing alerts. The default setting is **Medium**.

Alerts which are not visible under the current configurations are not stored in the database, unless they are part of an incident. To change this behavior, you can set the option `save_invisible_alerts` to `true`

Low

Use this level if you are starting to get control of the network, you do not have deep information about it, or if you are resource-constrained. This level reduces that amount of alerts.

Medium

Use this level if you have some information about the network and how it operates, and have knowledge of scheduled network changes. In this scenario, it will be useful to help you understand if new elements should be allowed or not.

High

Use this level if you have deep knowledge about the network, can have access to schedule of changes in the network and want to have deep visibility of events. Only more verbose Alerts are removed.

Paranoid

Use this level for maximum visibility, but at the cost of a more time-consuming effort in alert management.

Zone configurations

*You can select this to go to the **Zone configurations** page. You can customize all settings related to the learning engine and the security profile on a per-zone basis.*

For more details, see **Settings > Zone configurations**.

Alert tuning

*The **Alert tuning** page lets you customize the alert behavior. Specifically, you can impose conditions on one, or many, fields to match criteria. This feature can be selectively enabled for specific user groups.*

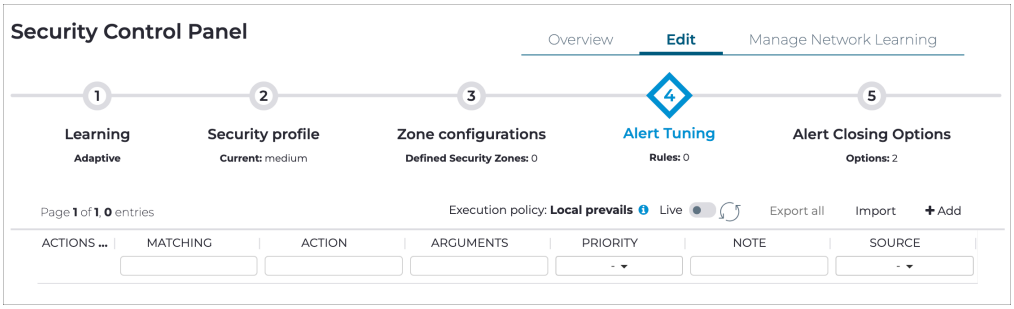


Figure 4. Alert tuning page

Execution policy

As alert rules can be propagated from upstream connections, conflicts between rules are possible. A conflict is detected when multiple rules, performing the same action, match an alert. To deal with these collisions, the execution algorithm takes into consideration the source of the rules. The user can choose three policies:

- **upstream_only**: alert rules are managed in the top *Central Management Console (CMC)*, or with Vantage. Creation and modification are disabled in the lower-level sensors. Only the rules received from upstream are executed
- **upstream_prevails**: in case of conflicts, rules coming from upstream are executed
- **local_prevails**: in case of conflicts, rules created locally are executed

A special case is represented by the **mute** action. Consider the following example: the execution policy is **local_prevails** and a mute rule is received by Guardian from an upstream connection. This rule will be ignored if at least one local rule matches the alert. Conversely, with the execution policy set to **upstream_prevails**, local **mute** will be ignored if at least one rule coming from upstream matches the alert.

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.

Export all

Use this to export the table.

Import

This lets you import alert rules.

Note:

The maximum file size is 2 *gigabyte (GB)*. The supported file type is `.nozomi_alert_rules`.

+ Add
This lets you add and configure an alert.

Alert configuration settings

Source IP	Enter the *internet protocol (IP)* address of the source that you want to filter.
Destination IP	Enter the *IP* address of the destination that you want to filter.
Source MAC	Enter the *media access control (MAC)* address of the source that you want to filter.
Destination MAC	Enter the *MAC* address of the destination that you want to filter.
Match IPs and MACs in both directions	Select this if you want to select all the communications between two nodes (*IP* or *MAC*) independently of their role in the communication (source or destination).
Source Zone	Enter the zone of the source that you want to filter.
Destination Zone	Specify the zone of the destination that you want to filter.
Type ID	The type ID of the alert, this field is precompiled if you create a new modifier from an alert in the **Alerts** page.
Trigger ID	Unique identifier corresponding to the specific condition that has triggered the alert.
Protocol	Enter the protocol that you want to filter.
Note	Enter free-form text that describes details of the alert rule.
Execute action	Select an action to perform on the matched alerts: • **Mute**: Switch ON/OFF: to mute or not the alert • **Mute until**: Specify a date until which the alert will be muted • **Change Security Profile visibility**: Set to **ON** to force the visibility of the selected alert type for any selected profile, or to **OFF** to hide it for any selected profile. Useful for extending or reducing the default provided security profiles as needed • **Change risk**: Set a custom risk value for the alert • **Change trace filter**: Define a custom trace filter to apply to this alert • **Assign playbook**: Define a playbook to be attached to the matching alerts. The playbook to be attached has to be selected from the list of available playbook templates
Priority	Set a custom priority; when multiple rules trigger on an alert, the rule with the highest priority applies. **Normal** is the default value if no selection is made.

Alert closing options

*The **Alert closing options** page lets you customize the closure details of alerts and incidents. When alerts and incidents are closed, the user must choose the reason why the closure happens. There are two default reasons:* **actual incident** *and* **baseline change**.

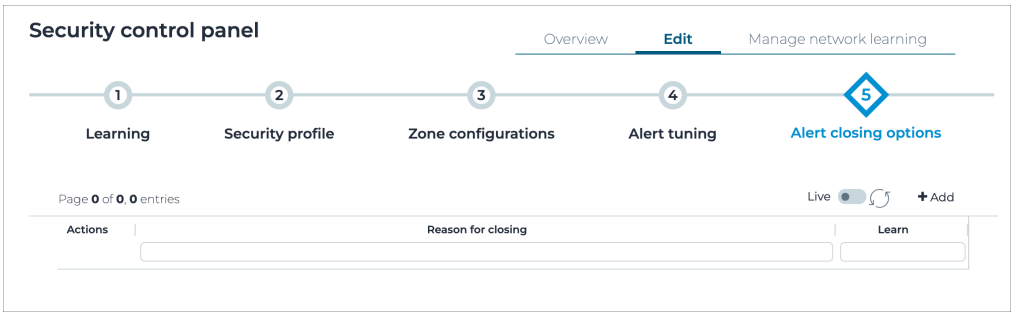
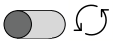


Figure 5. Alert closing options page

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.

Add

This button lets you add a new alert closing option.

Manage network learning

*The **Manage network learning** page lets you review and manage the Network Learning status in detail.*

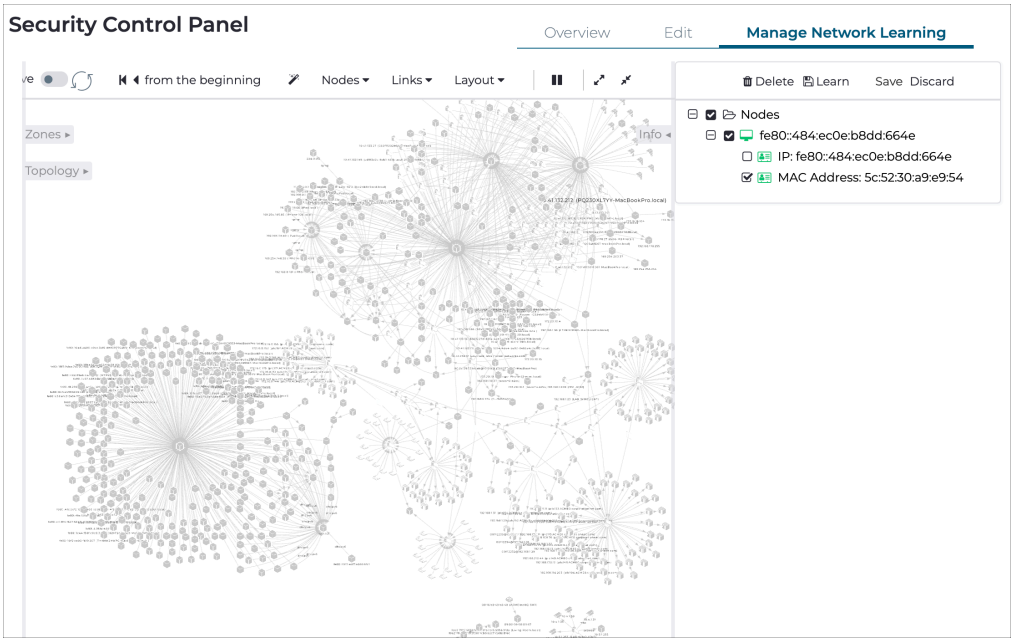


Figure 6. Manage network learning page

The graph is initialized with the node and link not learned perspectives. The items that are not known to the system are highlighted in red or orange. This makes it easy to discover new elements, and take an action on them.

To the right of the graph view, is a panel with additional buttons.

Delete
This lets you delete the selected item(s) from the system.

Learn
This lets you add the selected item(s) in the system.

Save
When the configuration is complete, you can select **Save** to make the changes persistent.

Discard
This lets you discard all unsaved changes to the system.

Learn a protocol

The Manage network learning page lets you select links to let you learn protocols.

Procedure

1. In the top navigation bar, select 

   Result: The administration page opens.

2. In the **Settings** section, select **Security control panel**.

   Result: The **Security control panel** page opens.

3. In the top right section, select **Manage network learning**.

   Result: The **Manage network learning page** opens.

4. Select a red or orange link.

   Result: Information for the link will show in the pane on the right. If the protocol is *supervisory control and data acquisition (SCADA)*, a function code will also show.

5. In the pane on the right, below the selected link, select one, or more, protocol(s).
6. Select **Learn**.

   Result: A save icon will show to the right of the selected item(s) and the **Save** and **Discard** buttons will be highlighted.

7. Select **Save**.

   Result: The protocol(s) will be learned, and it will show green.

Results

If all protocols were learned, the link will show in gray. If only some of the protocols were learned, the link will show in orange.

Learn a node

The Manage network learning page lets you select links to let you learn protocols.

Procedure

1. In the top navigation bar, select 

   Result: The administration page opens.

2. In the **Settings** section, select **Security control panel**.

   Result: The **Security control panel** page opens.

3. In the top right section, select **Manage network learning**.

   Result: The **Manage network learning** page opens.

4. Select a red or orange node.

   Result: Information for the link will show in the pane on the right. For each node, the node id will be shown, and below two children items will show the node *IP* address and the *MAC* address.

5. In the pane on the right, below the selected node, select one, or more, item(s) to be learned.
6. Select **Learn**.

   Note:
   If you select **Delete**, all the items will be selected, and deleted.

   Result: A save icon will show to the right of the selected item(s) and the **Save** and **Discard** buttons will be highlighted.

7. Select **Save**.

   Result: The node(s), and all the selected items will show green.

Results

If all nodes were learned, the node will show in gray. If only some of the protocols were learned, the link will show in orange.

Features

*The **Features** page shows an overview of the current status of system features configuration and lets you fine tune specific values.*

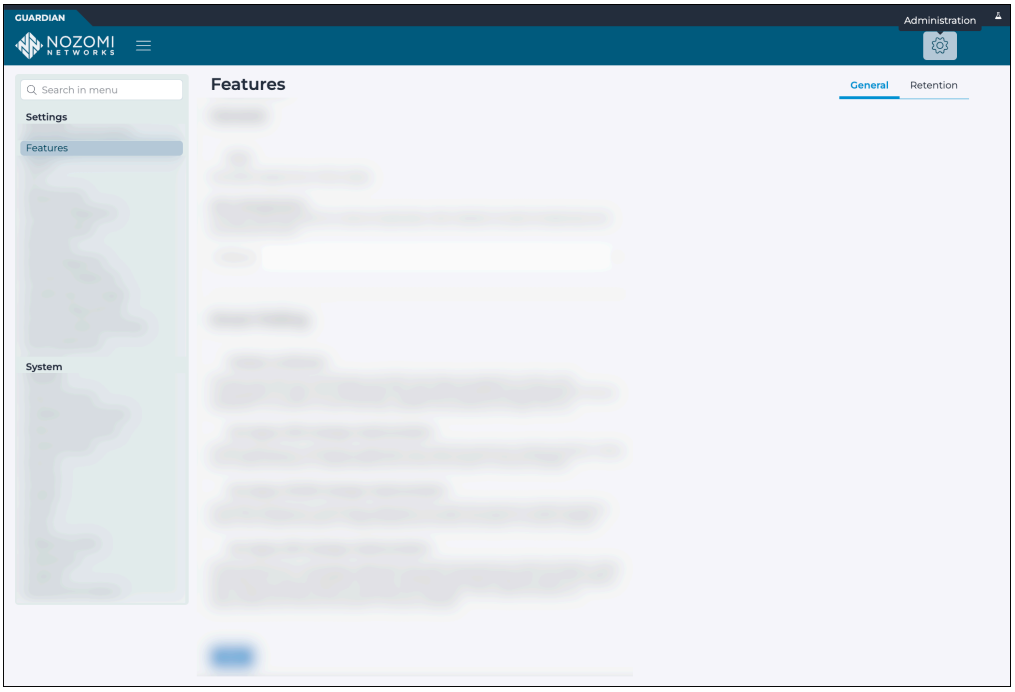


Figure 7. Features page

The **Features** page has these tabs:

- General (on page 29)
- Retention (on page 31)

General

*The **General** page lets you configure the behavior of a sensor.*

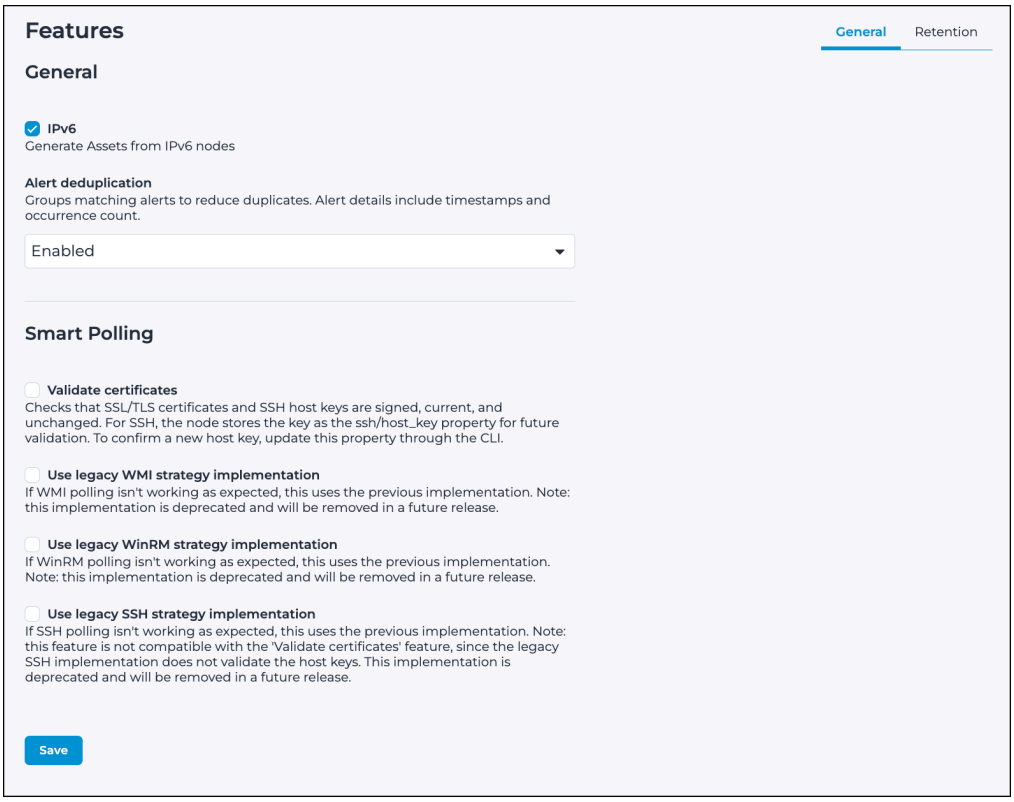


Figure 8. General page

IPv6

This checkbox enables the creation of assets with an *Internet Protocol version 6* address that are detected through passive monitoring.

Alert deduplication

This control lets you turn the deduplication of alerts on or off. When deduplication is enabled, Guardian groups identical events into a single alert instead of creating individual alerts for each event. This simplifies alert management and provides better visibility of anomalies.

By default, this feature is turned off. The default will change to enabled in *Nozomi Networks Operating System (N2OS)* v26.5.0. Leaving the status set to **Default** will result in the feature being turned on automatically when *N2OS* v26.5.0 is installed.

Validate certificates

By default, Smart Polling does not check that the *secure sockets layer (SSL)*/*transport layer security (TLS)* certificate is valid when connecting to an *hypertext transfer protocol secure (HTTPS)* server, nor does it verify that the host key of an *secure shell (SSH)* server is known. You can enable these security checks here.

Smart Polling uses the *IP* address to contact *HTTPS* servers. For certificate verification to succeed, the certificate must contain the server's *IP* address among the valid names.

For *SSH*, the host key is compared against the content of the `ssh/host_key` property of the polled node.

Use legacy WMI strategy implementation

If you find issues with your WMI Smart Polling plans that did not occur before, reverting to the legacy implementation of this strategy may resolve them. Contact technical support if this is the case.

Use legacy WinRM strategy implementation

If you find issues with your *Windows Remote Management (WinRM)* Smart Polling plans that did not occur before, reverting to the legacy implementation of this strategy may resolve them. Contact technical support if this is the case.

Use legacy SSH strategy implementation

If you find issues with your *SSH* Smart Polling plans that did not occur before, reverting to the legacy implementation of this strategy may resolve them. Contact technical support if this is the case.



Note:

The legacy implementation will not verify the *SSH* host key even if **Validate certificates** is turned on.

Retention

*The **Retention** tab lets you configure how much historical data Guardian stores for each data category.*

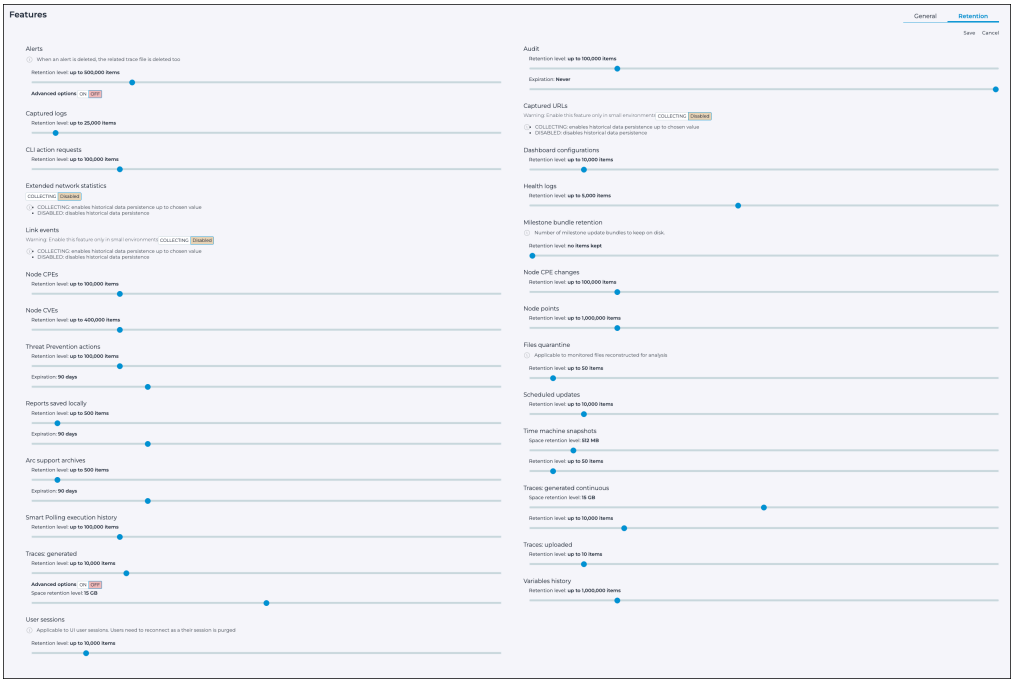


Figure 9. Retention tab

The **Retention** tab displays a set of data category cards in a two-column layout. Each card shows one or more sliders you can adjust to control how much data Guardian stores. Some categories also have a **COLLECTING** / **DISABLED** toggle, and some offer **Advanced options** that let you set per-subcategory limits within the overall total.

Expiration

Sets the maximum age of stored items in days. When an item exceeds this age, the system deletes it. Set to **Never** to keep items indefinitely.

Retention level

Sets the maximum number of items to store. When the limit is reached, the system deletes the oldest items to make room for new ones.

Space retention level

Sets the maximum disk space to use for storing items. When the limit is reached, the system deletes the oldest items.

Discard rate

Sets the maximum rate of incoming items per minute. The system discards new items when they exceed this threshold.

Data categories

The following data categories appear on the **Retention** tab. Categories that require optional features only appear when you have licensed and enabled those features on this sensor.

Data category	Controls	Notes
Left column (top to bottom)		
Alerts	Retention level	Deleting an alert also deletes the related trace file. Advanced options let you set a separate sub-limit for alerts outside the active security profile.
Captured logs	Retention level	
CLI action requests	Retention level	
Extended network statistics	Space retention level	Can be toggled on or off.
Link events	Retention level, Discard rate	Can be toggled on or off. Enable only in small environments.
Node CPE changes	Retention level	
Node points	Retention level	
Reports saved locally	Retention level, Expiration	
Arc support archives	Retention level, Expiration	
Smart Polling execution history	Retention level	
Threat Prevention actions	Retention level, Expiration	Tracks device actions taken from an alert (for example, blocking or unblocking a USB device), including their status.

Data category	Controls	Notes
Traces: generated	Retention level, Space retention level	Advanced options let you set sub-limits for traces generated by high risk alerts, medium risk alerts, low risk alerts, and by user request. The space retention level control is not available on systems where traces use tmpfs-based storage.
User sessions	Retention level	Applies to UI user sessions only. Users must reconnect when their session is purged.
Right column (top to bottom)		
Audit	Retention level, Expiration	
Captured URLs	Retention level, Discard rate	Can be toggled on or off. Enable only in small environments.
Dashboard configurations	Retention level	
Health logs	Retention level	
Milestone bundle retention	Retention level	Controls how many milestone update bundles are kept on disk. Set to 0 to disable retention (no bundles are kept).
Node CPEs	Retention level	
Node CVEs	Retention level	Requires Vulnerability Assessment.
Files quarantine	Retention level	Applies to monitored files reconstructed for analysis.
Scheduled updates	Retention level	
Time machine snapshots	Space retention level, Retention level	
Traces: generated continuous	Space retention level, Retention level	
Traces: uploaded	Retention level	

Data category	Controls	Notes
Variables history	Retention level	

Users management

*The **Users management** page shows all the pages that you need to let you manage authentication and authorization policies for users and groups.*

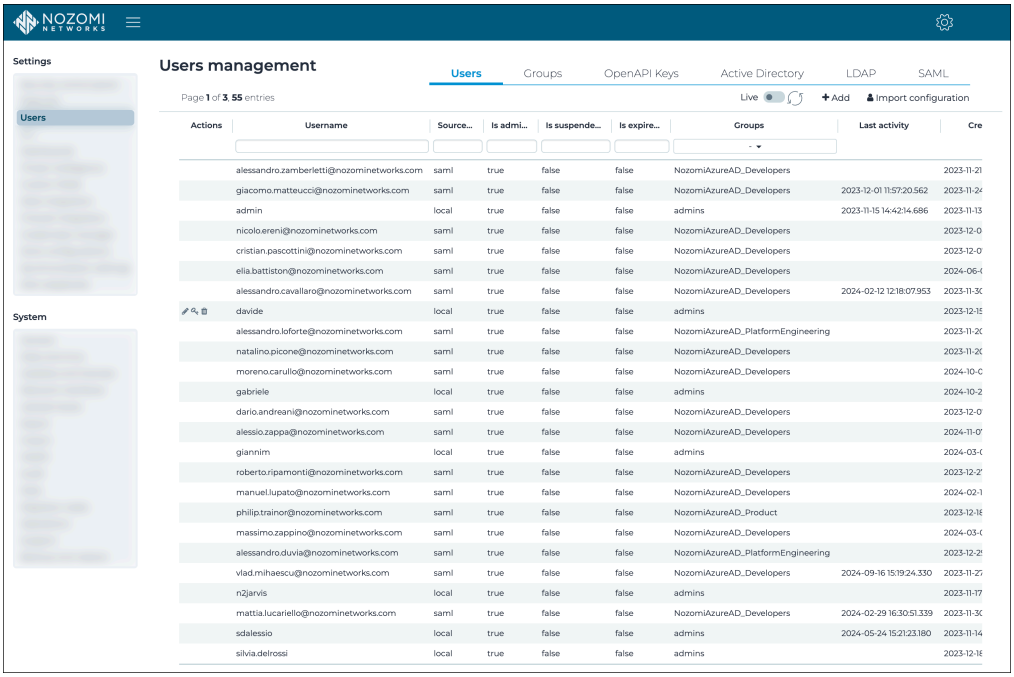


Figure 10. Users management page

General

The **Users management** page has these tabs:

- Users
- Groups
- OpenAPI Keys
- Active Directory
- LDAP
- SAML

There are four different types of user:

- Local users (on page 35)
- Active Directory users (on page 36)
- LDAP users (on page 36)
- SAML users (on page 36)

Local users

Authentication is enforced with a password, and the user is created from the Web *user interface (UI)*.

Active Directory users

Active Directory manages authentication. User properties and groups are imported from the Active Directory. You need to configure Active Directory (on page 57) for it to work correctly.

LDAP users

lightweight directory access protocol (LDAP) manages authentication. User properties and groups are imported from *LDAP*. You need to configure LDAP (on page 60) for it to work correctly.

SAML users

A password is not required as *single sign-on (SSO)* authentication is enforced through an authentication server that uses *security assertion markup language (SAML)*. You can use the Web *UI* to add users (on page 65), or to Import SAML users from a CSV file (on page 41).

Users

*The **Users** page shows a list of users and lets you create and delete users, and change the password and/or username of existing users.*

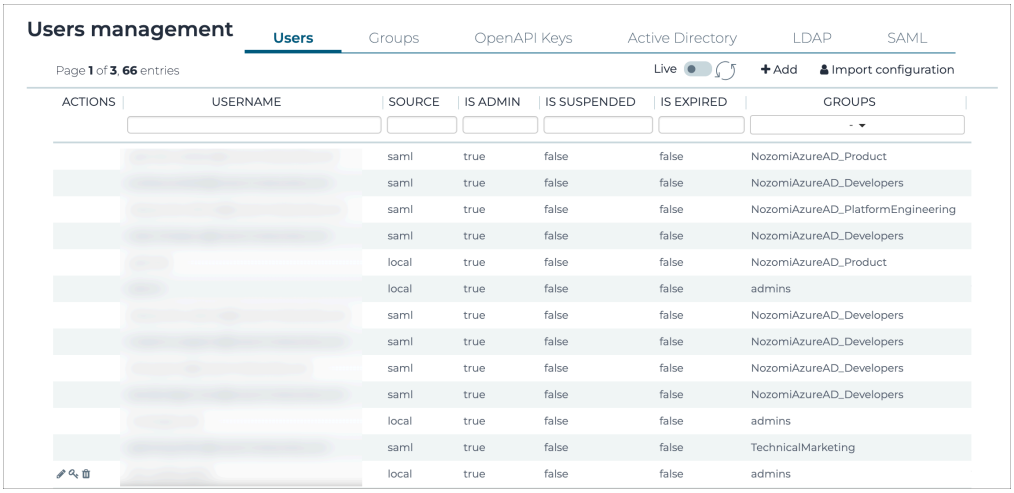


Figure 11. Users page

Add a user

*The **Users** page lets you add a new user.*

About this task

When you add a new user, you will typically select **Local** or **SAML** as the source. You can choose to select a user from the **Active Directory** or from **LDAP**, but you must first make sure that the user exists in the Active Directory or in LDAP. Therefore, in these cases, we recommend that you import these users directly from the Active Directory or from LDAP.

Procedure

1. In the top navigation bar, select 

   Result: The administration page opens.

2. In the **Settings** section, select **Users**.

   Result: The **Users management** page opens.

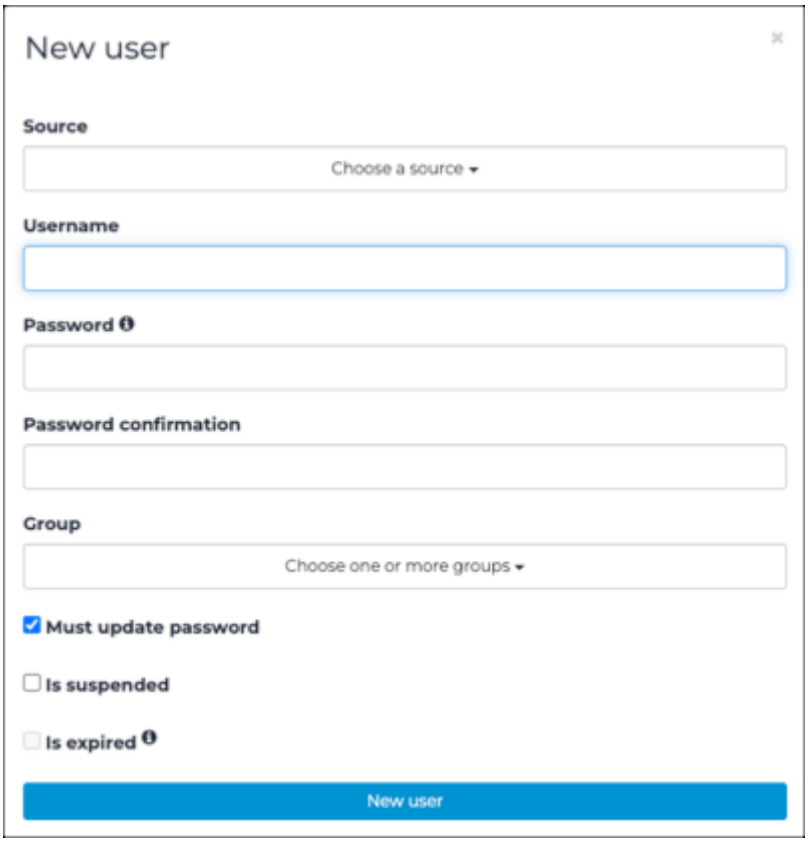
3. In the top right section, select **Users**.

   Result: The **Users** page opens.

4. In the top right section, select **+Add**.

   Result: A dialog shows.

5. From the **Source** dropdown, select a source for the user.



Choose from:

- Local
- Active Directory
- LDAP
- SAML

> **Note:**
> The dialog might change after you make your selection.

6. In the **Username** field, enter a value.
7. In the **Password** field, enter a password.

> **Note:**
> This step is not applicable if you chose **SAML** as the source.

8. In the **Password confirmation** field, enter the password again.

> **Note:**
> This step is not applicable if you chose **SAML** as the source.

9. From the **Group** dropdown, select a group for the user.
10. **Optional:** If necessary, select **Must update password**. This is selected by default.

> **Note:**
> When this is selected, the user will be prompted to update their password the next time they log in.

11. **Optional:** If necessary, select **Is suspended**.

> **Note:**
> When this is selected, the user will not be able to log in.

12. **Optional:** If necessary, select **Is expired**.

> **Note:**
> When this is selected, the user is forced to change their password the first time that they log in after the expiration date.

13. Select **New user**.

Results

The user has been added.

Import SAML users from a CSV file

*The **Users** page lets you import security assertion markup language (SAML) users from a comma-separated value (CSV) file.*

About this task

The template for the *comma-separated value (CSV)* file is three fields for each row, separated by commas:

- The first field defines the user name
- The second field is the **Authentication** group that is associated with the user (typically an **Authentication**-only group)
- The last field includes one or more groups (separated by semicolons) that define additional groups associated with the user (typically used to define allowed features)

Procedure

1. In the top navigation bar, select 

   Result: The administration page opens.

2. In the **Settings** section, select **Users**.

   Result: The **Users management** page opens.

3. In the top right section, select **Users**.

   Result: The **Users** page opens.

4. In the top right section, select **Import configuration**.

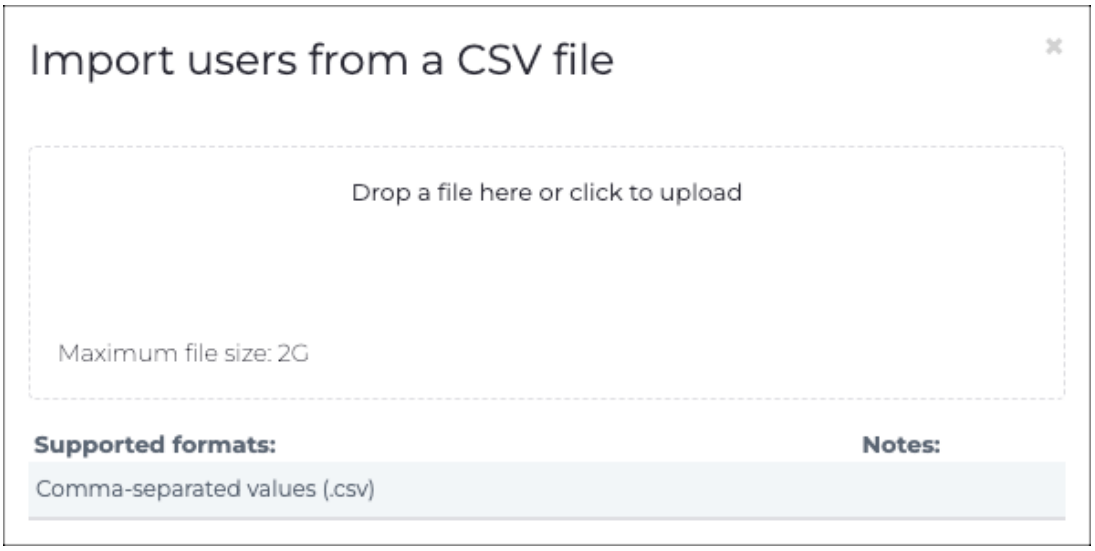
   Result: A dialog shows.

5. Choose a method to upload a file.

   Choose from:

   - Drag your file into the **Drop a file here or click to upload** field
   - Click in the **Drop a file here or click to upload** field

6. If you chose the second method, select the correct file to upload.



Note:
The supported format is *CSV*.
The maximum permitted file size is 2 *GB*.

7. Wait for the file to upload.

Results

Once the import is complete, you will receive a message with the number of users that have been correctly imported.

Edit a local user

*You can use the **Users** page to edit the details for an existing user.*

Procedure

1. In the top navigation bar, select 

   Result: The administration page opens.

2. In the **Settings** section, select **Users**.

   Result: The **Users management** page opens.

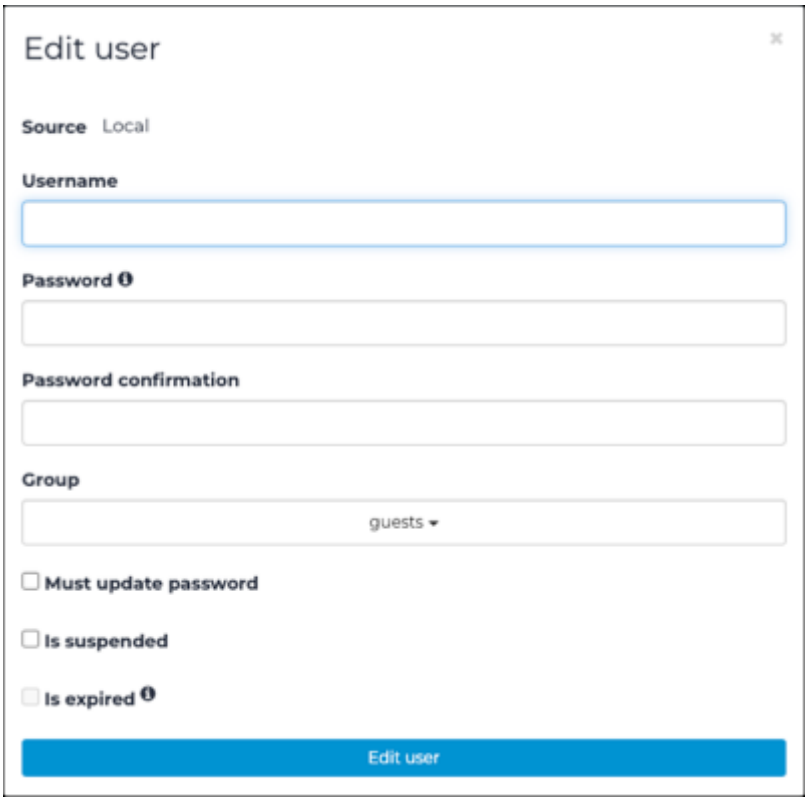
3. In the top right section, select **Users**.

   Result: The **Users** page opens.

4. To the left of the applicable user, select the  icon.

   Result: A dialog shows.

5. Edit the details as necessary.



6. Select **Edit user**.

Results

The details for the local user have been edited.

Add an SSH key for an admin user

You can use secure shell (SSH) public keys to log into the SSH console without typing a password. You must add SSH public keys to the user account to configure SSH password-less authentication.

Procedure

1. In the top navigation bar, select 

   Result: The administration page opens.

2. In the **Settings** section, select **Users**.

   Result: The **Users management** page opens.

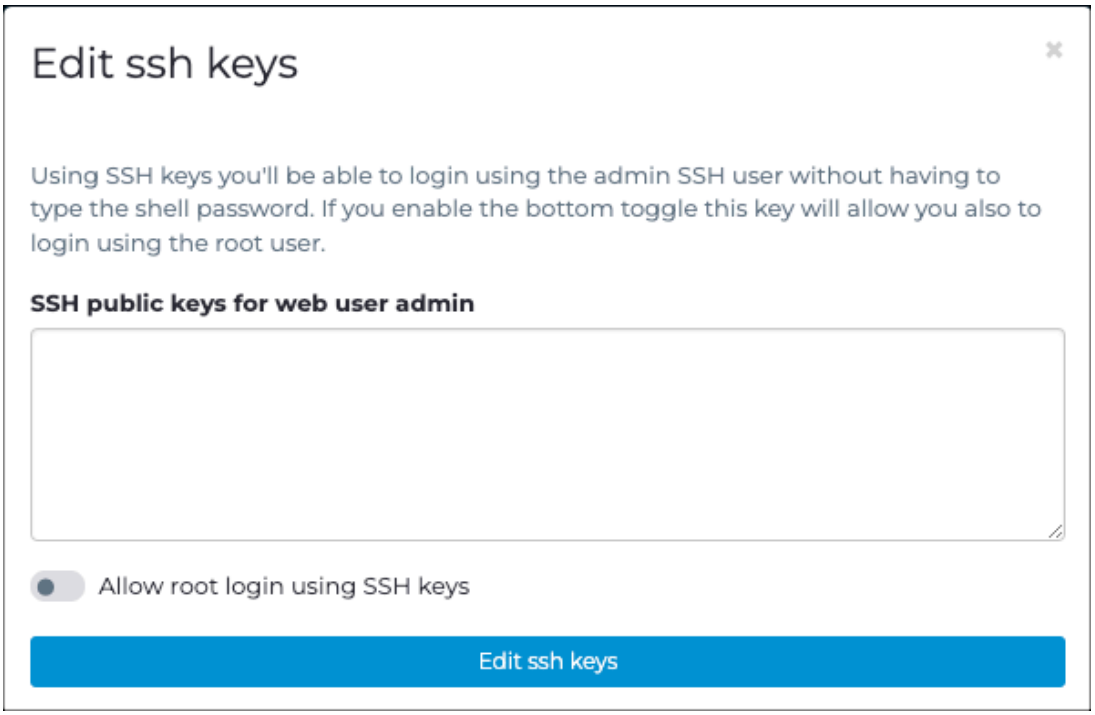
3. In the top right section, select **Users**.

   Result: The **Users** page opens.

4. To the left of the applicable user, select the  icon.

   Result: A dialog shows.

5. In the field, paste the public key in the **SSH public keys for web user admin** field.



> **Note:**
> Every admin user has a key. If you need more than one key, paste one per line. Non-admin users must use a password for *SSH* authentication. When an admin user leaves, the associated *SSH* keys are removed.

> **Note:**
> The pasted key should not contain new lines. The system will not use invalid keys.

6. **Optional:** To enable logging in with the root account, select **Allow root login using SSH keys**.
7. Select **Edit ssh keys**.

Results

The *SSH* public keys are propagated to all the sensors that are directly connected. The default key propagation interval is 30 minutes. You can change this with: `conf.user configure ssh_key_update interval <seconds>` in the **CLI**.

Groups

*The **Groups** page shows a list of all the user groups.*

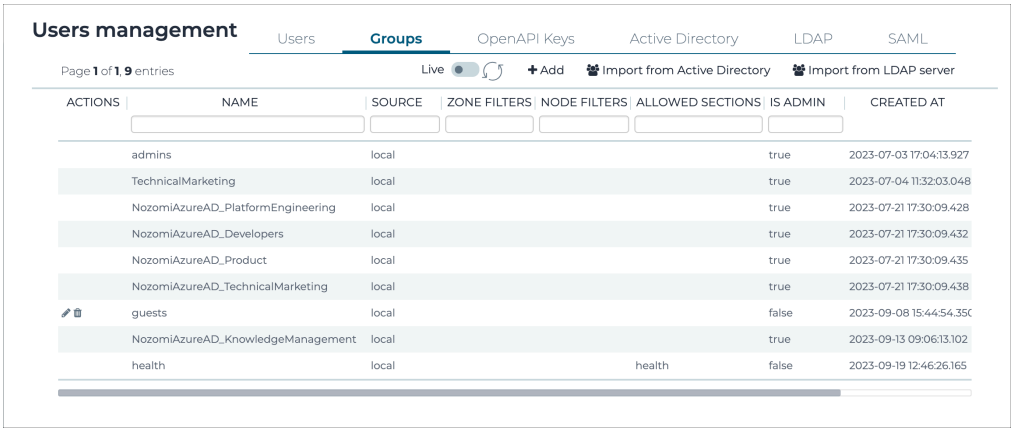


Figure 12. Groups page

Authorization policies

User groups define authorization policies. Each group includes a:

- List of allowed features
- Filter to enable visualization of just specific node subsets

When a user belongs to a group, the user can only:

- Perform the operations that the group allows
- See the nodes that the group node filter defines

A user can belong to several groups and will inherit the authorizations of those groups. When a user belongs to multiple groups, any node that satisfies the filter of any group is visible, and its features are available.

Two group types have predefined authorization policies:

- **Administrators**: All features are available
- **Authentication Only**: Only the authentication feature is available

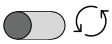
When a group is neither **Administrators** nor **Authentication Only**, the allowed features (sections) can be enabled/disabled individually.

Note:

After a reboot, the local default admin of the Web *UI* will automatically be recreated (within the default admin's user group) if it has been deleted, or if it doesn't exist. This is to make sure that a user cannot mistakenly delete it.

Note:
We recommend to have at least one group without admin privileges. Therefore, after a reboot, if a user group with no admin privileges doesn't exist, a guests user group will be created automatically, without a user in it.

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.

+Add

This lets you add a new group.

Import from Active Directory

This lets you import a group from Active Directory.

Import from LDAP server

This lets you import a group from a *LDAP* server.

Add a local group

*The **Groups** page lets you add new user groups.*

Procedure

1. In the top navigation bar, select ⚙

   Result: The administration page opens.

2. In the **Settings** section, select **Users**.

   Result: The **Users management** page opens.

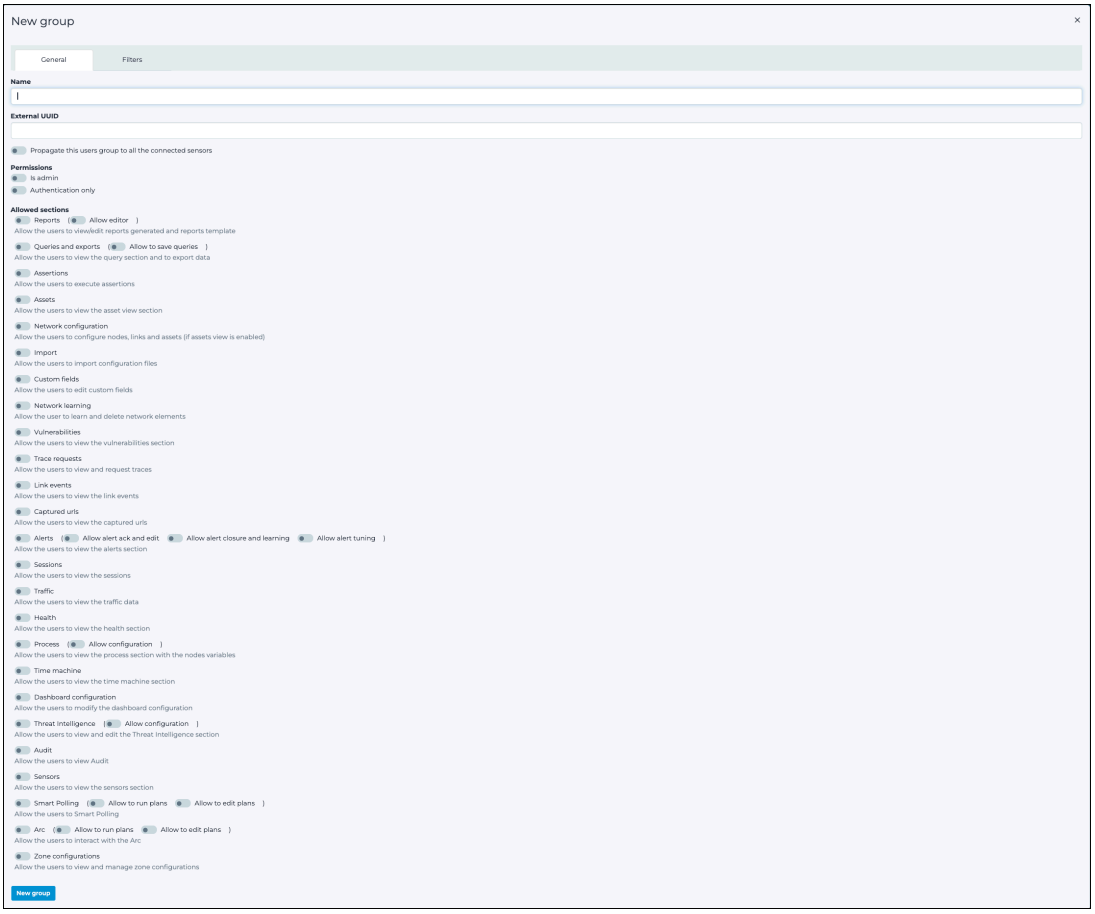
3. In the top right section, select **Groups**.

   Result: The **Groups** page opens.

4. In the top right section, select **+Add**.

   Result: A dialog shows.

5. In the top left, select **General**.



6. In the **Name** field, enter a name for the group.

7. **Optional:** If necessary, in the **External UUID** field, enter a *universally unique identifier (UUID)*.

   Note:
   This is useful should the user group be created through *SAML* integration and the external *identity provider (IdP)* uses an *identifier (ID)* rather than a human-readable group name.

8. If you want the group to propagate to connected sensors, select **Propagate this users group to all the connected sensors**.
9. If the group belongs to a predefined type, select the appropriate one:

   Choose from:
   - Is admin
   - Authentication only

10. If you do not select a predefined group type, continue with the steps below to manually select section(s) that the group can view and interact with.
    a. In the top left, select **Filters**.
    b. From the **Zone filters** dropdown, select one or more zone(s) to define zone filters to limit zone visibility to the users in the group.
    c. In the **Node filters** field, enter a list of subnet addresses in *classless inter-domain routing (CIDR)* format.

       Note:
       The addresses must be comma-separated. This limits the nodes users in a group can view in:
       - Nodes
       - Links
       - Variables
       - List
       - Graph
       - Queries
       - Assertions

11. Select **New group**.

Results

The group has been added.

Edit a group

*You can use the **Group** page to edit the details for an existing group.*

Procedure

1. In the top navigation bar, select ⚙

   Result: The administration page opens.

2. In the **Settings** section, select **Users**.

   Result: The **Users management** page opens.

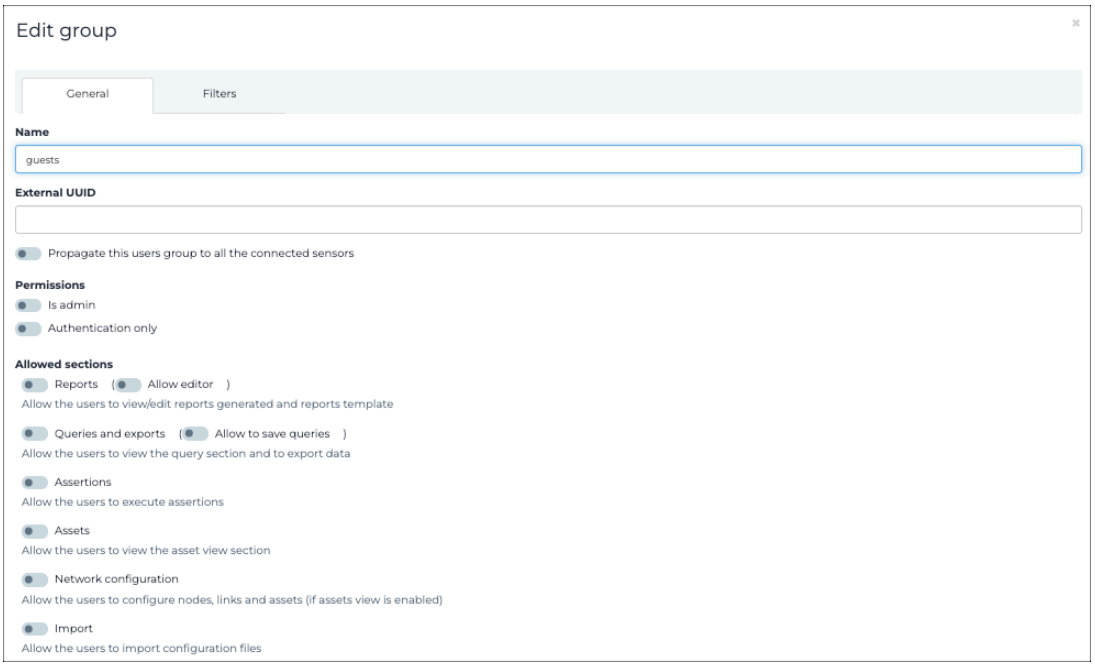
3. In the top right section, select **Groups**.

   Result: The **Groups** page opens.

4. To the left of the applicable group, select the ✎ icon.

   Result: A dialog shows.

5. Edit the details as necessary.



6. Select **Edit group**.

Results

The details for the group have been edited.

Import an Active Directory group

To permit Active Directory users to log into the system, you can import an existing group from an Active Directory infrastructure.

Procedure

1. In the top navigation bar, select ⚙

   Result: The administration page opens.

2. In the **Settings** section, select **Users**.

   Result: The **Users management** page opens.

3. In the top right section, select **Groups**.

   Result: The **Groups** page opens.

4. In the top right, select **Import from Active Directory**.

   Result: A dialog shows.

5. In the **Username** field, enter the Active Directory user logon name.



> **Note:**
> This should be in format `<domainname>\<domainusername>` format.

6. In the **Password** field, enter a password.
7. To retrieve the list of groups, select **Retrieve groups**.

> **Note:**
> You can also select **Filter by group name**, and type the name of the group that you want to retrieve.

8. Filter and select the desired groups to import.
9. **Optional:** To import related groups, for example, parent groups, select **Get also member-of groups**.

10. Select **Import configuration**.

    Result: The list of groups opens.

11. Edit the group permissions as necessary.

> **Note:**
> Active Directory users that belong to this group are automatically assigned to it and inherit all permissions of the configured group.

Results

Users can log into the system with the `<domainname>\<domainusername>` username and their current domain password in the login screen.

Import an LDAP group

To permit existing lightweight directory access protocol (LDAP) users to log into the system, you can import an existing group from an LDAP server.

Procedure

1. In the top navigation bar, select ⚙

   Result: The administration page opens.

2. In the **Settings** section, select **Users**.

   Result: The **Users management** page opens.

3. In the top right section, select **Groups**.

   Result: The **Groups** page opens.

4. In the top right, select **Import from LDAP server**.

   Result: A dialog shows.

5. In the **Username** field, enter a username.



> **Note:**
> This requires an admin user with full *LDAP* server permission.
> The **Username** for the *LDAP* server should be a distinguished name (DN) that follows the *LDAP* standard. For example, `cn=username,cn=group,dc=nozominetworks,dc=com`.

6. In the **Password** field, enter a password.
7. To retrieve the list of groups, select **Retrieve groups**.

> **Note:**
> You can also select **Filter by group name**, and type the name of the group that you want to retrieve.

8. Filter and select the desired groups to import.
9. **Optional:** To import related groups, for example, parent groups, select **Get also member-of groups**.

10. Select **Import configuration**.

    Result: The list of groups opens.

11. Edit the group permissions as necessary.

> **Note:**
> *LDAP* users that belong to this group are automatically assigned to it and inherit all permissions of the configured group.

Results

Users can log into the system with the username and their current domain password in the login screen.

OpenAPI Keys

*The **OpenAPI Keys** page lets you manage the openAPI keys. You can use OpenAPI keys for bearer token authentication instead of basic authentication, which uses a username and password.*

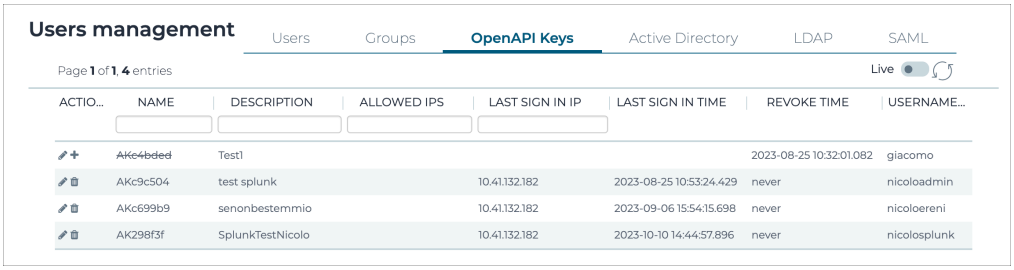


Figure 13. OpenAPI keys page

General

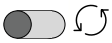
You can only assign OpenAPI keys to local users.

Local users can access their OpenAPI keys in **Personal settings > Other actions > Edit openAPI keys**.

For more details, see the **Guardian User Guide**.

For more details about open *application programming interface (API)* keys, see the **SDK User Manual**.

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.

Active Directory

In addition to local users, you can configure existing Active Directory users for login.

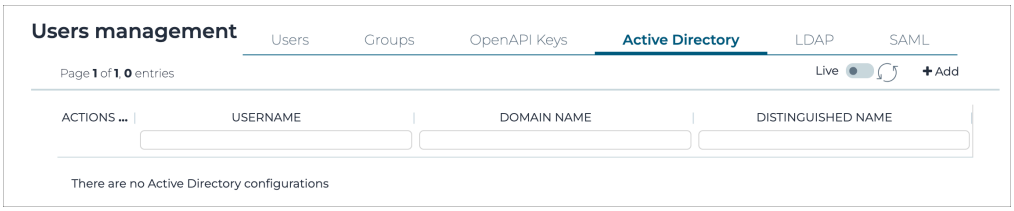


Figure 14. Active Directory page

Active Directory permissions are defined based on the user group. When you set the primary group for a user, that user is excluded from the corresponding group membership in the Active Directory. This is because the Active Directory does not support primary group functionality. It does not query the primary group attribute when building the group membership of a user, therefore the primary group on the Active Directory server is not visible in Guardian.

Configure Active Directory integration

*The **Active Directory** page lets you configure integration.*

Procedure

1. In the top navigation bar, select ⚙

   Result: The administration page opens.

2. In the **Settings** section, select **Users**.

   Result: The **Users management** page opens.

3. In the top right section, select **Active Directory**.

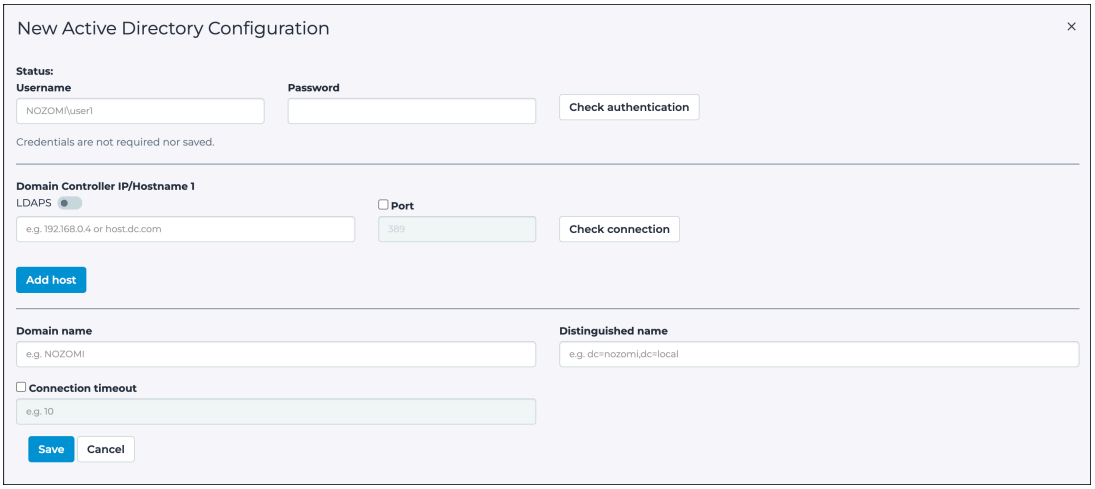
   Result: The **Active Directory** page opens.

4. In the top right, select **+Add**.

   Result: A dialog shows.

5. **Optional:**

   In the **Username** field, enter a username.



6. **Optional:** In the **Password** field, enter a password.

> **Note:**
> In order to connect and integrate into the Active Directory, users must belong to at least one group with read permission on the server. Administrator privileges are not required.

7. **Optional:** To check that Active Directory is running correctly, select **Check authentication**.

8. In the **Domain Controller IP/Hostname 1** section:

   Choose from:

   - To use *LDAP*, leave the **LDAPS** toggle unselected
   - To use *lightweight directory access protocol secure (LDAPS)*, select the **LDAPS** toggle to on.

9. **Optional:** If necessary, and you chose *LDAPS*, select **Verify SSL**.

   > **Note:**
   > By default, the server's *SSL* certificate is not verified.

10. If you chose *LDAP*, in the **Port** field, enter a 389. If you chose *LDAPS*, in the Port field, enter 636.
11. To check that Active Directory is running correctly on the port, select **Check connection**.
12. To add another domain controller *IP* address, select **Add host**.
13. In the **Domain name** field, enter the applicable details for your Active Directory.
14. In the **Distinguished name** field, enter the applicable details for your Active Directory.
15. **Optional:** If necessary, select **Connection timeout** and enter a value (in seconds).
16. Select **Save**.

LDAP

*The **LDAP** page shows a list of all the lightweight directory access protocol (LDAP) configurations.*

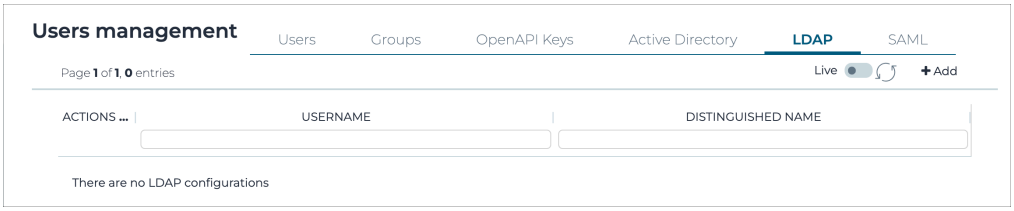


Figure 15. LDAP page

Add LDAP users

You can add existing lightweight directory access protocol (LDAP) users for login. LDAP permissions are defined based on the user group.

Before you begin

Make sure that you have:

- The domain name (i.e., pre-Windows 2000 name), referred to as `<domainname>`
- The domain distinguished name, referred to as `<domainDN>`
- One or more domain controller *IP* addresses, referred to as `<domaincontrollerip>`

About this task

The supported *LDAP* formats are:

- v2
- v3

Procedure

1. In the top navigation bar, select ⚙

   Result: The administration page opens.

2. In the **Settings** section, select **Users**.

   Result: The **Users management** page opens.

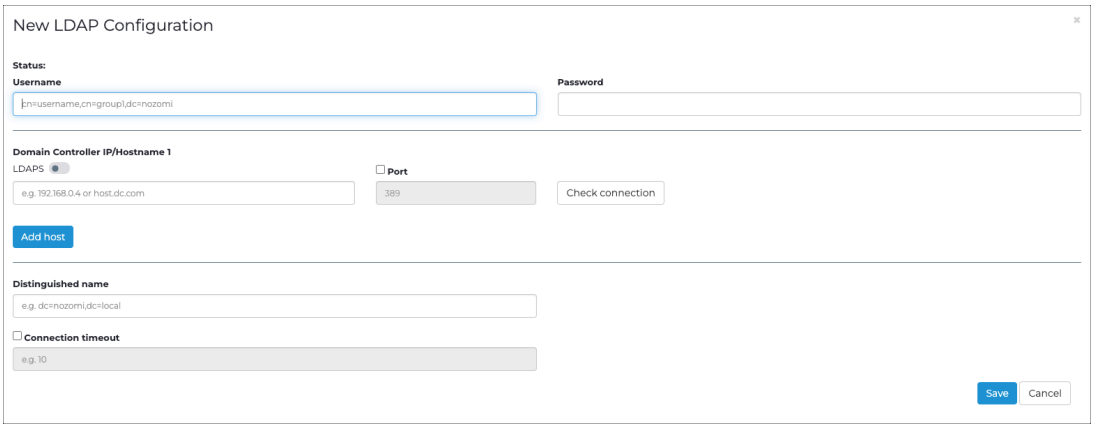
3. In the top right section, select **LDAP**.

   Result: The **LDAP** page opens.

4. In the top right section, select **+Add**.

   Result: A dialog shows.

5. In the **Username** field, enter a username.



> **Note:**
> This requires an admin user with full *LDAP* server permission. The **Username** for the *LDAP* server should be a distinguished name (DN) that follows the *LDAP* standard. For example, `cn=username,cn=group,dc=nozominetworks,dc=com`.

6. In the **Password** field, enter a password.
7. In the **Domain Controller IP/Hostname 1** section:

   Choose from:

   - To use *LDAP*, leave the **LDAPS** toggle unselected
   - To use *LDAPS*, select the **LDAPS** toggle to on.

8. **Optional:** If necessary, and you chose *LDAPS*, select **Verify SSL**.

> **Note:**
> By default, the server's *SSL* certificate is not verified.

9. If you chose *LDAP*, in the **Port** field, enter a 389. If you chose *LDAPS*, in the Port field, enter 636.
10. To check that Active Directory is running correctly on the port, select **Check connection**.
11. To add another domain controller *IP* address, select **Add host**.
12. In the **Distinguished name** field, enter a value.

13. **Optional:** If necessary, select **Connection timeout**, and enter a value in seconds.
14. To save the changes and validate the data, select **Save**.

> **Note:**
> If there are errors, they will show next to the **Status** field.

SAML

Nozomi Networks supports security assertion markup language (SAML) single sign-on (SSO) authentication.

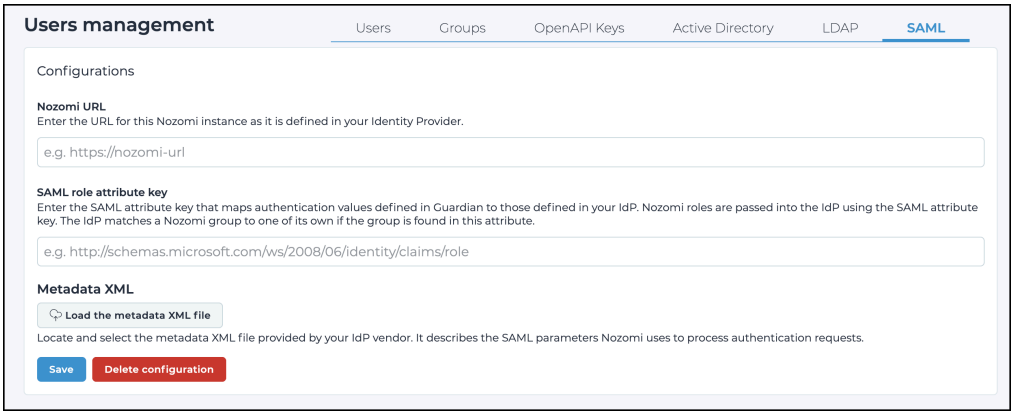


Figure 16. SAML page

General

Your *IdP* must be compatible with *SAML* 2.0.

The *SAML* configuration process is often error-prone. To implement *SAML* integration you should be familiar with:

- The *SAML* protocol
- Your *IdP* software
- The exact details of your specific *IdP* implementation

You can select **Propagate the *SAML* metadata** to push the current metadata to sensors directly connected to this machine, instead of manually downloading and re-uploading the metadata file on each one. This is useful, for example, after renewing this machine's *SSL* certificate, which changes the metadata.

Note:

This option is disabled until this machine is set and enabled as *IdP* with an *IdP uniform resource locator (URL)* specified. Only a *CMC* can be configured as *IdP*.

Note:

Metadata is only propagated to sensors directly connected to this machine. Sensors connected through a mid-level *CMC* do not receive the update.

Additional configuration

Typically, *SAML* with authentication, replies are sent back from the same host that originally received the request. Occasionally, *SAML* requests are chained between different *IdP*s, and replies might come from a different host. By default, the Web *UI* content security rules block these types of replies.

You can use the `csp form-action-urls` configuration key to override this behavior.

To accept replies from an *IdP SSO* target *URL* that differs from the one specified in the *SAML* metadata, you can issue the configuration rule: `conf.user configure csp form-action-urls <additional_url>` in the *command-line interface (CLI)*.

If you need to specify more than one *URL*, you should use spaces to separate them. After this change, you need to run the `service webserver stop` command, in a shell console to apply it.

Clock skew

Occasionally, the *IdP* and Guardian system times can differ. By default, the system accepts requests with up to 60 seconds difference. You can use the `saml clock_drift` configuration key to override this behavior.

To change the value, you can issue `conf.user configure saml clock_drift <allowed_seconds>` in the *CLI*.

After this change, you can run the `service webserver stop` command in a shell console to apply it.

Limitations

The *SAML* logout protocol is not supported.

Configure SAML integration

*The **SAML** page lets you add and configure security assertion markup language (SAML).*

Before you begin

Make sure that you have defined a new application in your *IdP*. This should consist of:

- The *assertion consumer service (ACS) URL* for Nozomi Networks. An *ACS* specifies the `/auth` path such as `https://10.0.1.10/saml/auth`
- The issuer *URL* for your *IdP*, which specifies the `/saml/metadata` path, such as `/saml/metadata`. This value depends on your *IdP*
- The metadata *eXtensible Markup Language (XML)* file that describes the *SAML* parameters of your *IdP*. Before configuring your sensor, download the file from your *IdP* vendor and save it to a location accessible to Nozomi Networks

About this task

Note:

You can configure a mid-level *CMC* as *IdP*. For more details, see **N2OS Configuration > Miscellaneous > Configure identity provider**. You can download the metadata from `https://<address>/idp/saml/metadata`, where `address` is the *IP* address of the *CMC* that is configured as *IdP*.

Procedure

1. In the top navigation bar, select ⚙

   Result: The administration page opens.

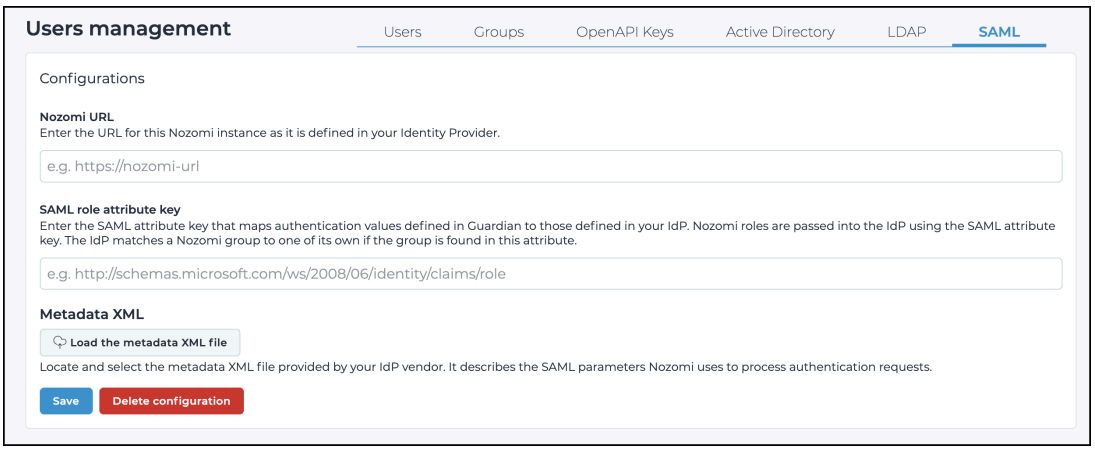
2. In the **Settings** section, select **Users**.

   Result: The **Users management** page opens.

3. In the top right section, select **SAML**.

   Result: The **SAML** page opens.

4. In the **Nozomi URL** field, enter the *URL* for your Nozomi Networks instance.



5. In the **SAML role attribute key** field, enter a string that will be used to map role names between the sensor and your *IdP*.

> **Note:**
> The value in this field is used to compare groups defined in the sensor with those defined in your *IdP*. The nature of this value depends on your *IdP*. (For example, if you are using Microsoft Office 365 as your *IdP*, the value might be `http://schemas.microsoft.com/ws/2008/06/identity/claims/role`

> **Note:**
> After you insert the *URL* in the **IdP URL** field, a **Open metadata page** hyperlink will show. This will open a new page that contains the metadata you can use in the downstream sensor. The metadata will be only be available after you save the configuration.

6. Select **Save**.
7. On the sensor login page, select **Single Sign On**.
8. To test the integration, use the credentials from your *IdP*.

> **Note:**
> For *SAML* to work properly, groups that match *SAML* roles must exist in the system. Groups are found using the role name. For example, if the *SAML* role attribute specifies an **Operator** role, the *IdP* looks for the **Operator** group when authorizing an authenticating user.

Results

SAML has been configured, and the login page shows a new **Single Sign On** button.

CLI

*The **CLI** page lets you change configuration parameters and perform troubleshooting activities.*

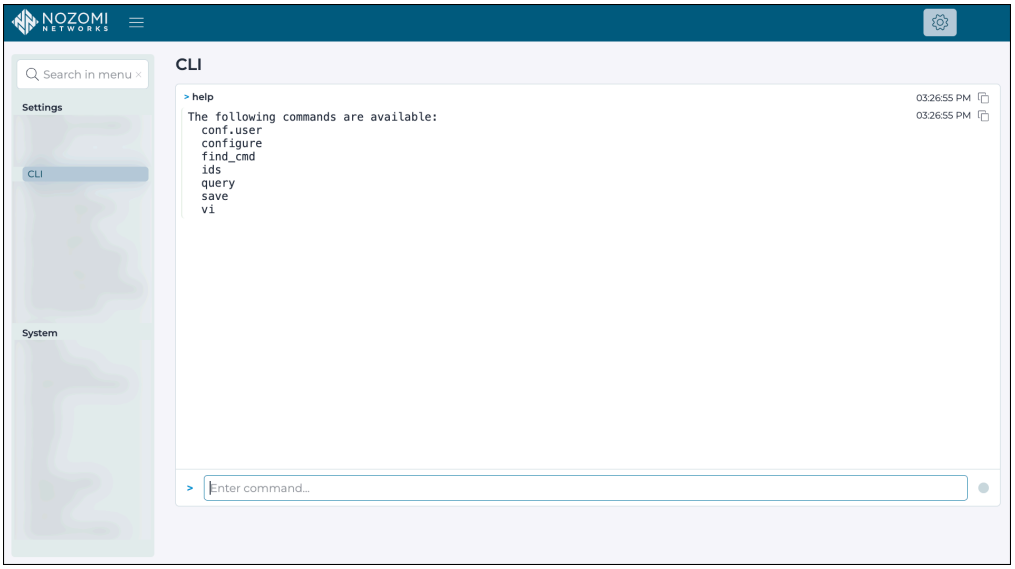


Figure 17. CLI page

If you enter an unknown command, an error will show.

When you select the copy icon, next to the command or command-output, the formatted text will be copied to the clipboard.

When the output takes more than 300 ms to show, a duration will appear next to the timestamp.

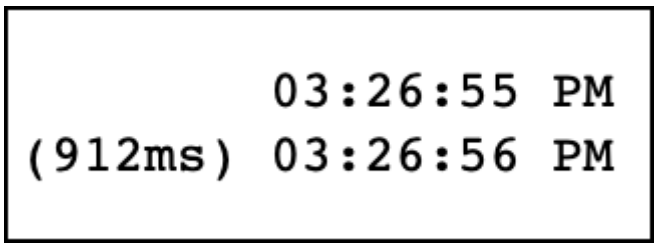


Figure 18. Duration shown next to timestamp

Table 1. Keyboard shortcuts

Shortcut	Action
Up arrow	Shows the previous search from the history
Down arrow	Shows the subsequent search from the history

Dashboards

*The **Dashboards** page lets you create and configure widget-based dashboards that provide information about your network. The dashboards created here will show on the Guardian home page, and give an overview of the monitored environment.*

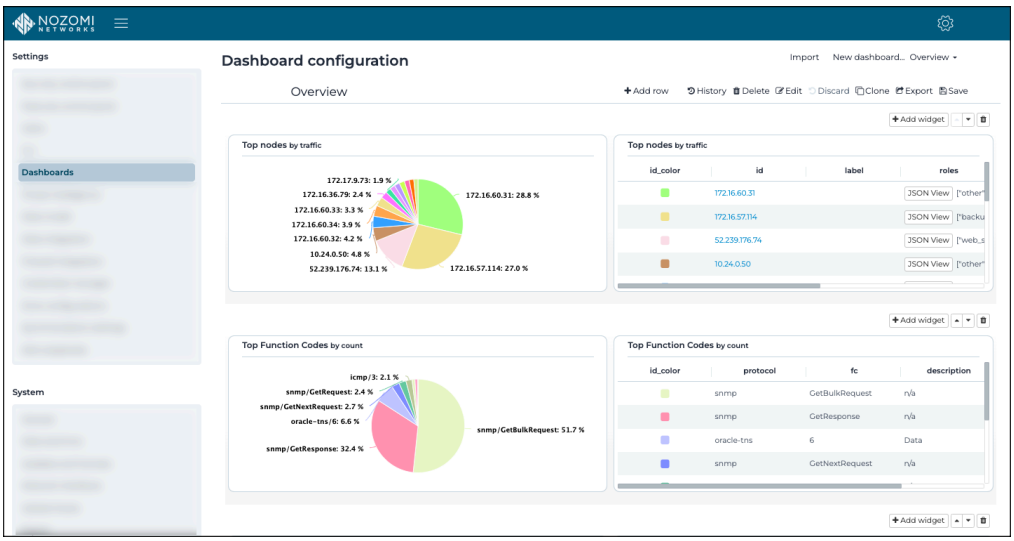


Figure 19. Dashboards page

Table 2. Default widgets

Widget	Description
Environment information	This provides a high level view of your network, in terms of the number of assets, nodes, links, protocols and variables
Asset overview	Assets, grouped by Purdue level
Alert flow over time	Alert risk charted over time
Latest alerts	Latest alerts (the most recent being first)
Failed assertions	List of failed assertions

Import
This lets you Import a dashboard (on page 75).

New dashboard
This lets you Create a dashboard (on page 71).

Dashboard selector
This dropdown lets you select from dashboards that already exist.

+ Add row
This lets you add a new row to the dashboard.

History
This lets you revert to a previous version of the dashboard.

Delete
This lets you delete the current dashboard.

Edit
This lets you edit the name and the group of users who can see this dashboard when they log in to the *CMC*.

Discard
This lets you discard the change that you have made to the dashboard.

Clone
This lets you clone the current dashboard to make a new one.

Export
This lets you export the current dashboard in *JavaScript Object Notation (JSON)* format.

Save
This lets you save the changes.

+ Add widget
This lets you add one, or more, widgets to the dashboard, that you can then configure.

Create a dashboard

*You can use the **Dashboards** page to load a dashboard that has been created previously.*

Procedure

1. In the top navigation bar, select 

   Result: The administration page opens.

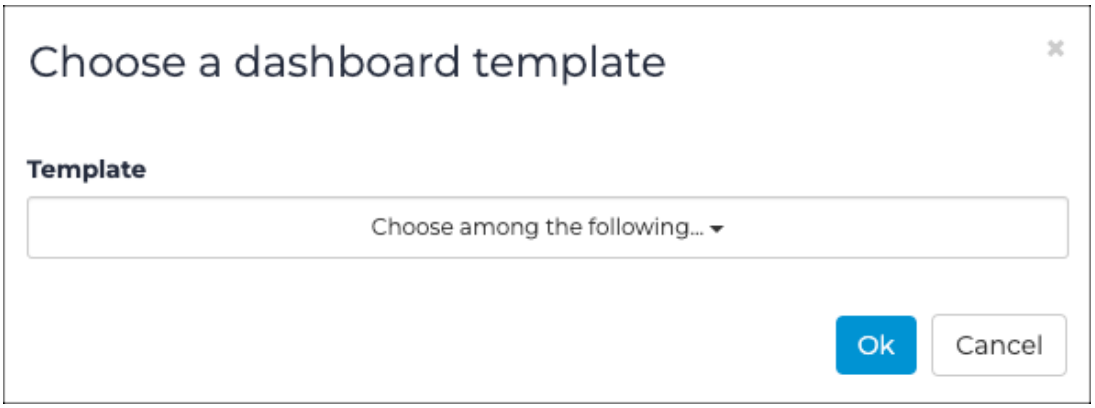
2. In the **Settings** section, select **Dashboards**.

   Result: The **Dashboard configuration** page opens.

3. In the top right section, select **New dashboard**.

   Result: A dialog shows.

4. If you want the new dashboard to be based on a template, from the dropdown, select a dashboard.



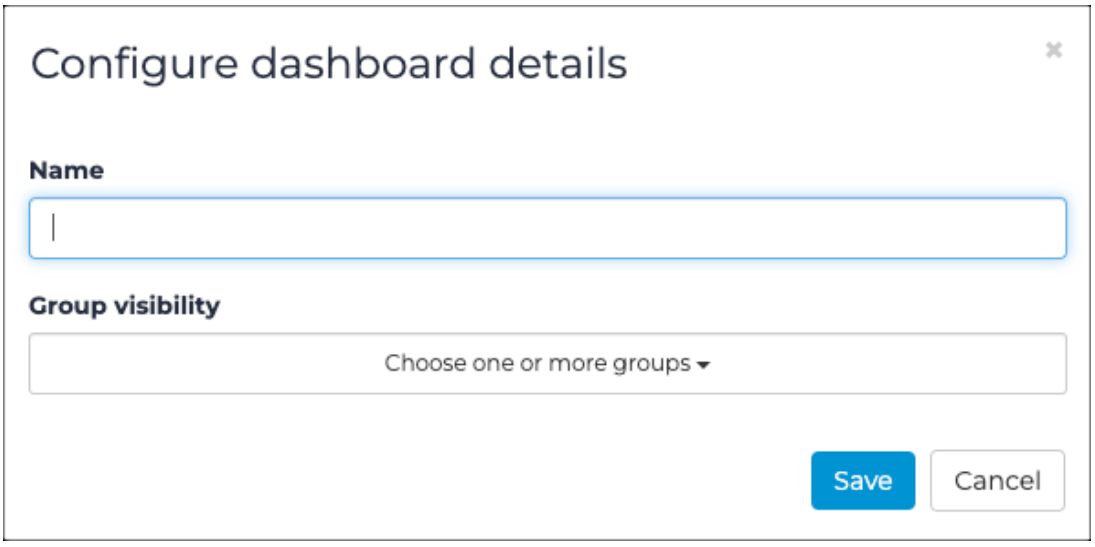
> **Note:**
> If you do not select a template, an empty dashboard will be created.

5. Select **Ok**.

6. In the top right section, select  **Save**.

   Result: A dialog shows.

7. In the **Name** field, enter a name for the dashboard.



8. From the **Group visibility** dropdown, select one or more groups to add the dashboard to.
9. Select **Save**.

Results

The dashboard has been created.

Related information

Configure a dashboard (on page 73)

Configure a dashboard

Once a dashboard has been created, you can configure it.

Procedure

1. Choose a method.

   Choose from:

   - Create a dashboard (on page 71)
   - Select a dashboard (on page 75)
   - On the homepage after you log in. In the top right section of the dashboard, select the  icon.

2. **Optional:** In the top right section, select **+ Add row**.

   Result: A new row shows.

3. In the top right section, select **+ Add widget**.

   Result: A list shows.

4. From the list, select one, or more widgets to add.

   Result: The widgets are added to the row.

5. Select somewhere on the screen other than the list.

   Result: The list closes.

6. **Optional:** If necessary, configure the widget.
   a. Hover your mouse over the widget.
      Result: The configuration toolbar shows.
   b. To increase the width of the widget, select ↔
   

   c. To decrease the width of the widget, select →←
   d. To increase the height of the widget, select ↕
   e. To decrease the height of the widget, select ↕
   f. To adjust the height of all of the widgets in the same row, select ☰
   g. To move the widget to the left, select ◀
   h. To move the widget to the right, select ▶
   i. To move the widget up to the row above, select ▲
   j. To move the widget down to the row below, select ▼
   k. To delete the widget, select 🗑
7. Select **Save**.

Results

The dashboard has been configured.

Select a dashboard

*You can use the **Dashboards** page to load a dashboard that has been created previously.*

Procedure

1. In the top navigation bar, select 

   Result: The administration page opens.

2. In the **Settings** section, select **Dashboards**.

   Result: The **Dashboard configuration** page opens.

3. In the top right section, select **New dashboard**.

   Result: A dialog shows.

4. From the dropdown, select a dashboard.

Results

The applicable dashboard shows.

Import a dashboard

*You can use the **Dashboards** page to import an existing dashboard.*

About this task

You can export dashboards from the **Dashboards** page, or from the **Export data** page to export dashboards in *JSON* format. These files can then be imported into other *CMC*s.

Procedure

1. In the top navigation bar, select 

   Result: The administration page opens.

2. In the **Settings** section, select **Dashboards**.

   Result: The **Dashboard configuration** page opens.

3. In the top right section, select **Import**.

   Result: A dialog shows.

4. Select the *JSON* format dashboard file.

Results

The applicable dashboard has been imported.

Export a dashboard

*You can use the **Dashboards** page to export an existing dashboard.*

About this task

This procedure shows you how to export a dashboard from the **Dashboards** page. To export from the **Export data** page, see **System > Export > Export a content pack**. You can export dashboards from the **Dashboards** page, or from the **Export data** page.

Procedure

1. In the top navigation bar, select 

   Result: The administration page opens.

2. In the **Settings** section, select **Dashboards**.

   Result: The **Dashboard configuration** page opens.

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

Results

The applicable dashboard has been exported in *JSON* format.

Threat Intelligence

*The **Threat Intelligence** page lets you manage packet rules, YARA rules, Sigma rules, structured threat information expression (STIX) indicators and vulnerabilities to provide detailed threat information.*

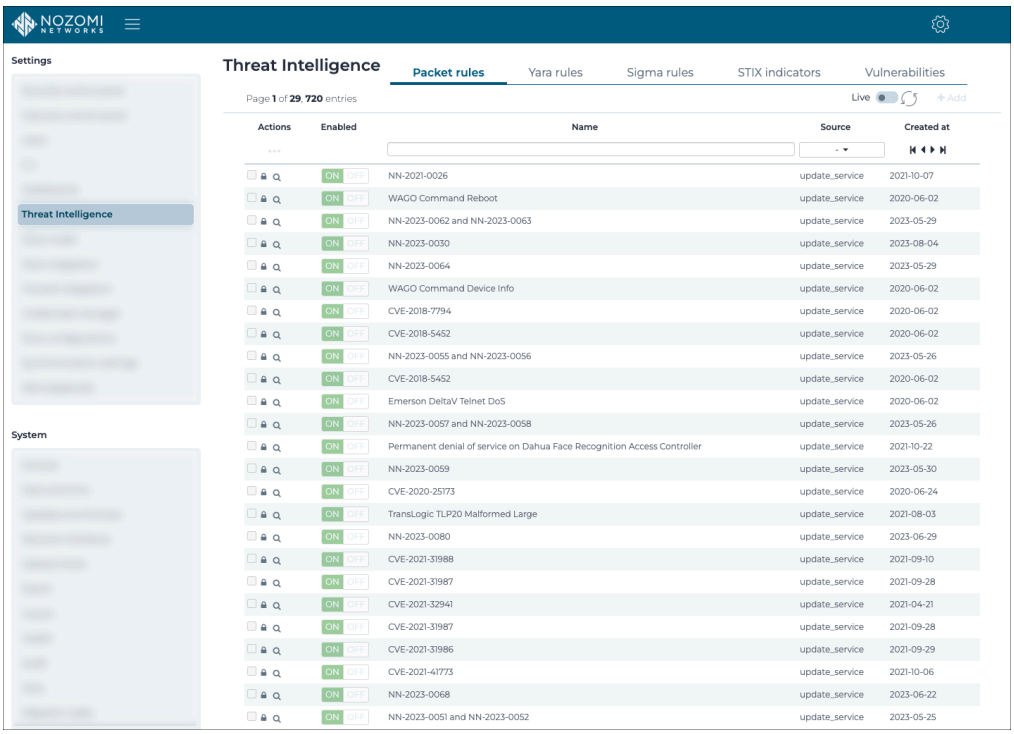


Figure 20. Threat Intelligence page

The **Threat Intelligence** page has these tabs:

- Packet rules (on page 78)
- Yara rules (on page 79)
- Sigma rules (on page 80)
- STIX indicators (on page 81)
- Vulnerabilities (on page 82)

Packet rules

*The **Packet rules** page lets you manage the packet rules for Threat Intelligence.*

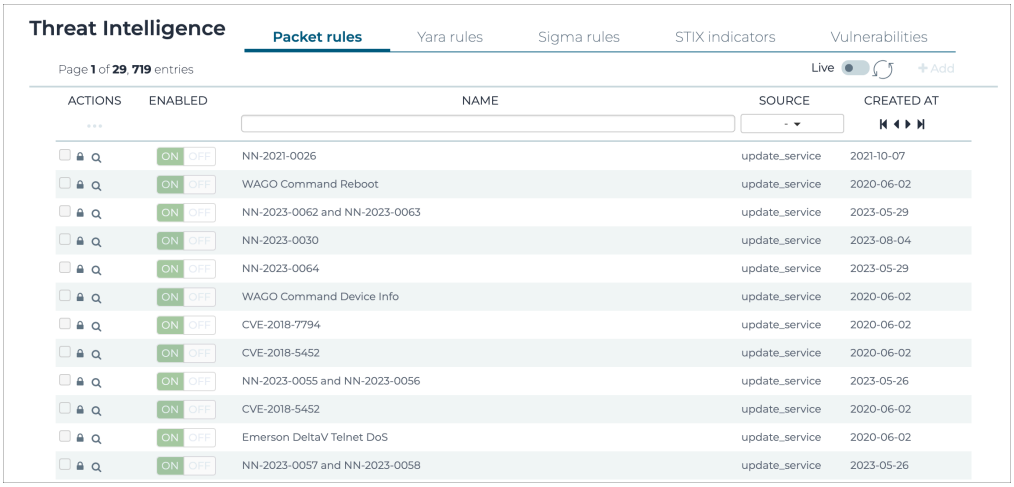


Figure 21. Packet rules page

Packet rules are executed on every packet. They raise an alert of type `SIGN:PACKET-RULE` if a match is found.

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.

Add

This lets you add a new packet rule.

Yara rules

*The **YARA rules** page lets you manage the YARA rules for Threat Intelligence.*

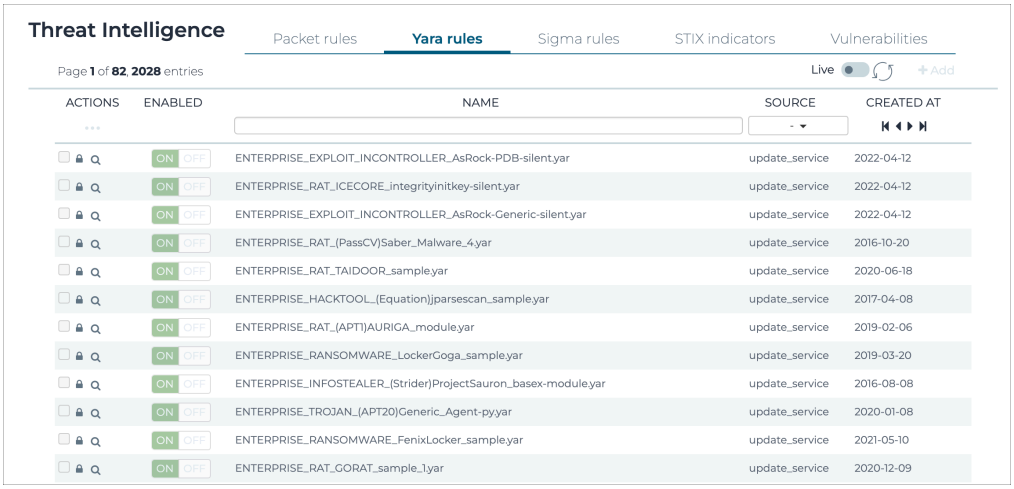


Figure 22. Yara rules page

YARA rules are executed on every file transferred over network protocols such as *hypertext transfer protocol (HTTP)* or *server message block (SMB)*. When a match is found, an alert of type `SIGN:MALWARE-DETECTED` is raised.

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.

Add

This lets you add a new YARA rule.

Sigma rules

*The **Sigma rules** page lets you manage the Sigma rules for Threat Intelligence that are applicable in an Arc deployment.*

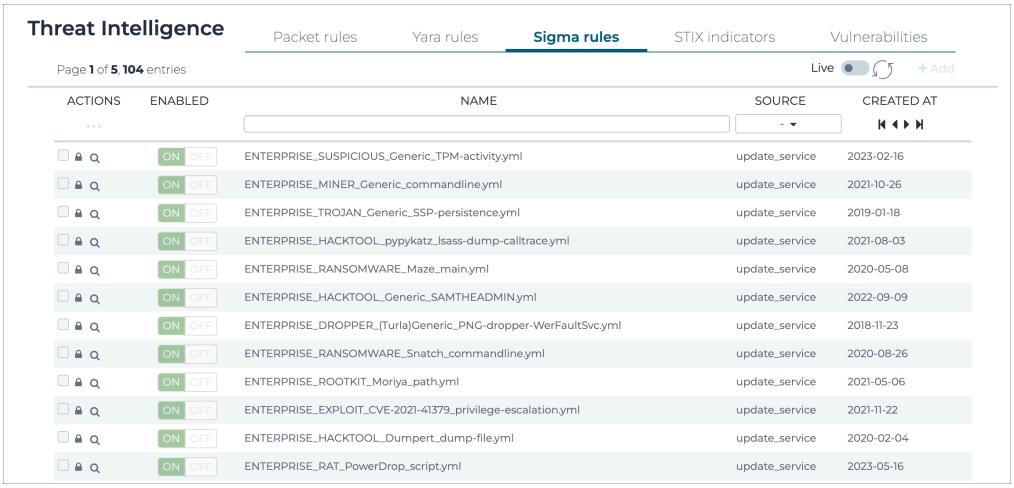


Figure 23. Sigma rules page

Sigma rules are versatile and generic rules written in the Sigma language. Primarily employed in threat detection and *security information and event management (SIEM)* systems, Sigma rules aim to standardize and offer a uniform method for describing log patterns across diverse security devices, applications, and platforms.

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.

Add

This lets you add a new Sigma rule.

STIX indicators

*The **STIX indicators** page lets you manage the structured threat information expression (STIX) indicators for Threat Intelligence.*

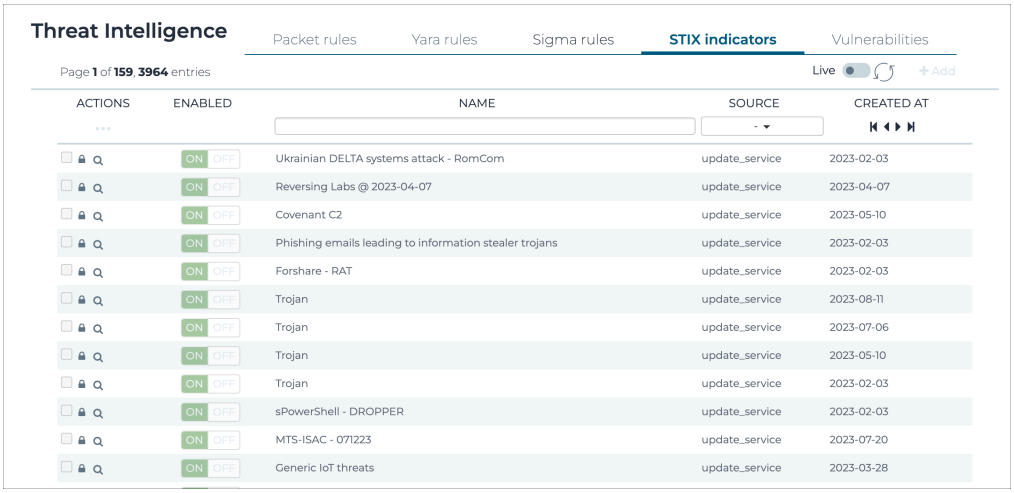


Figure 24. STIX indicators page

Structured Threat Information Expression (STIX) indicators contain information about malicious *IP* addresses, *URL*s, malware signatures, or malicious *domain name server (DNS)* domains. This information enriches existing alerts and raises new ones.

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.

Add

This lets you add a new *STIX* indicator.

Vulnerabilities

*The **Vulnerabilities** page lets you manage the vulnerabilities for Threat Intelligence.*

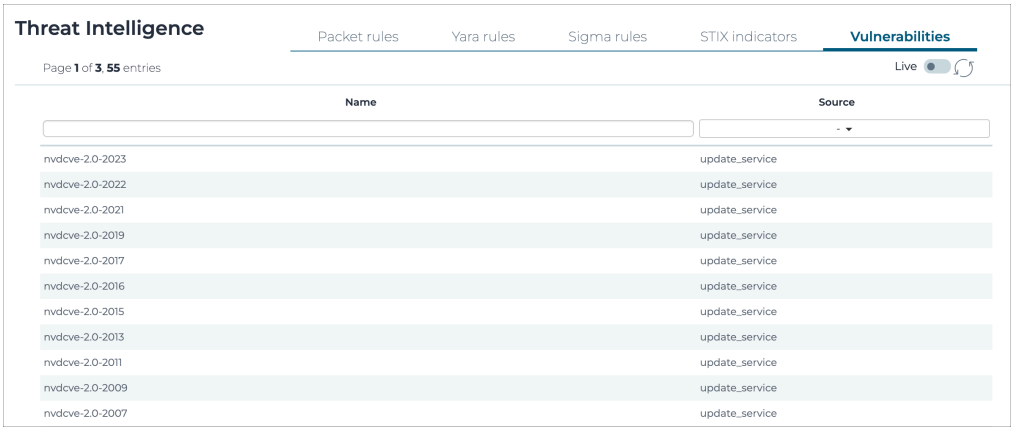


Figure 25. Vulnerabilities page

Vulnerabilities are assigned to each node, depending on the installed hardware, *operating system (OS)*, and the software identified in the traffic. The software uses *Common Vulnerabilities and Exposures (CVE)* , a dictionary that gives definitions for publicly-disclosed cybersecurity vulnerabilities and exposures.

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.

Add

This lets you add a new vulnerability.

Custom fields

*The **Custom fields** page lets you create a custom field to the **Nodes** (All-In-One CMC only) and **Assets** tables. A custom field lets you add information that might be relevant to your organization, and cannot be extracted from network traffic.*

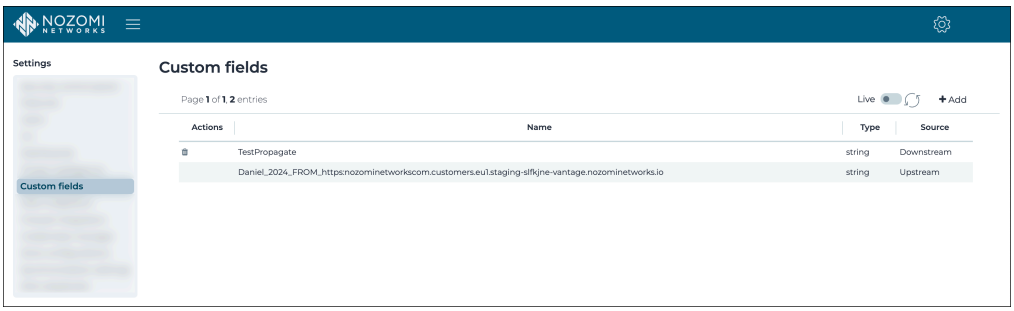


Figure 26. Custom fields page

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.

Add

This lets you add a custom field.

Discovery

Enables lightweight network announcements to detect neighboring devices. Use this option to initiate device discovery within the local network. Discovery is disabled by default.

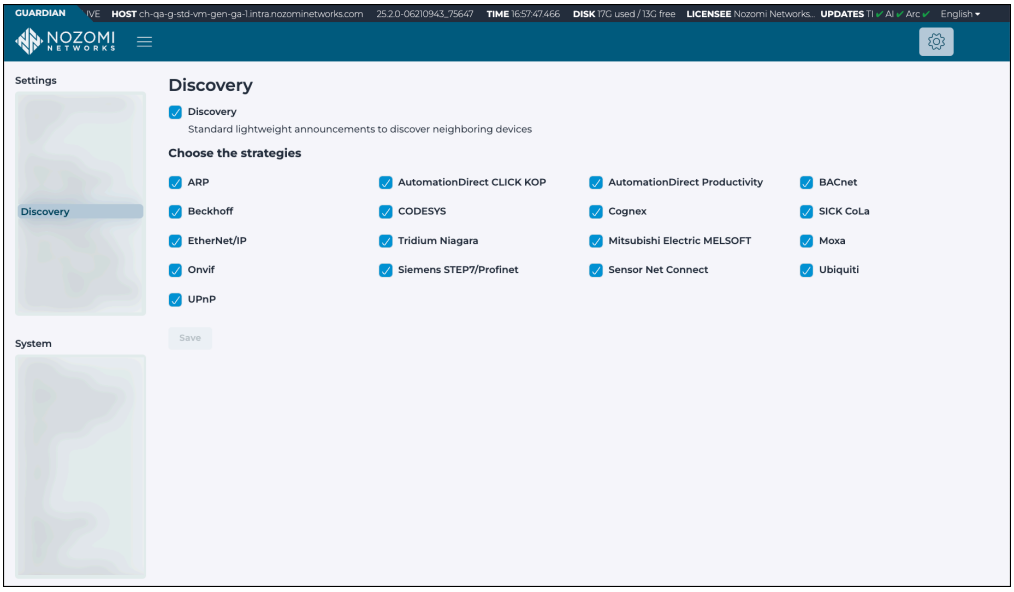


Figure 27. Discovery page

This checkbox lets you enable/disable Discovery. The default setting is disabled. When the checkbox is selected, the list of Discovery strategies is shown and you can enable or disable each strategy individually.

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.

Data integration

Data integration overview

*An overview of **Data integration** in the Nozomi Networks software.*

The Nozomi Networks solution uses third-party platforms to exchange data, in specific formats and methods. After configuring the endpoints of the third-party platforms for data integration, Guardian generates messages for:

- Alerts
- Health
- Audits

Connectivity status and status is checked with the data integration endpoint.
The third-party platforms that the Nozomi Networks solution uses for exchanging data are:

- Google Chronicle (Ingestion API - UDM)
- IBM QRadar (LEEF)
- Common Event Format (CEF)
- ServiceNow
- Tanium
- Splunk - Common Information Model (JSON)
- SMTP forwarding
- SNMP trap
- Syslog Forwarder
- Custom JSON
- Custom CSV
- DNS Reverse Lookup
- Kafka
- Cisco ISE
- External Storage
- Microsoft Endpoint Configuration Manager (WinRM RPC)
- Microsoft Endpoint Configuration Manager (DB)
- NetWitness
- Aruba ClearPass
- CarbonBlack Defense
- Active Directory
- Webhook

Data integration

*The **Data integration** page lets you exchange data with different third-party platforms.*

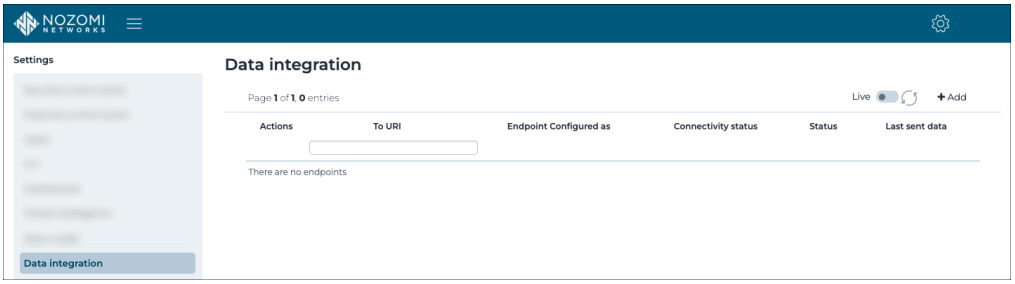


Figure 28. Data integration page

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.

Add

This lets you add a new endpoint.

Google Chronicle (Ingestion API - UDM)

This integration lets you forward alerts, asset information and vulnerabilities to an instance of Google Chronicle.

To configure this integration, you need to upload the Service Account file and to specify the matching customer *ID*.

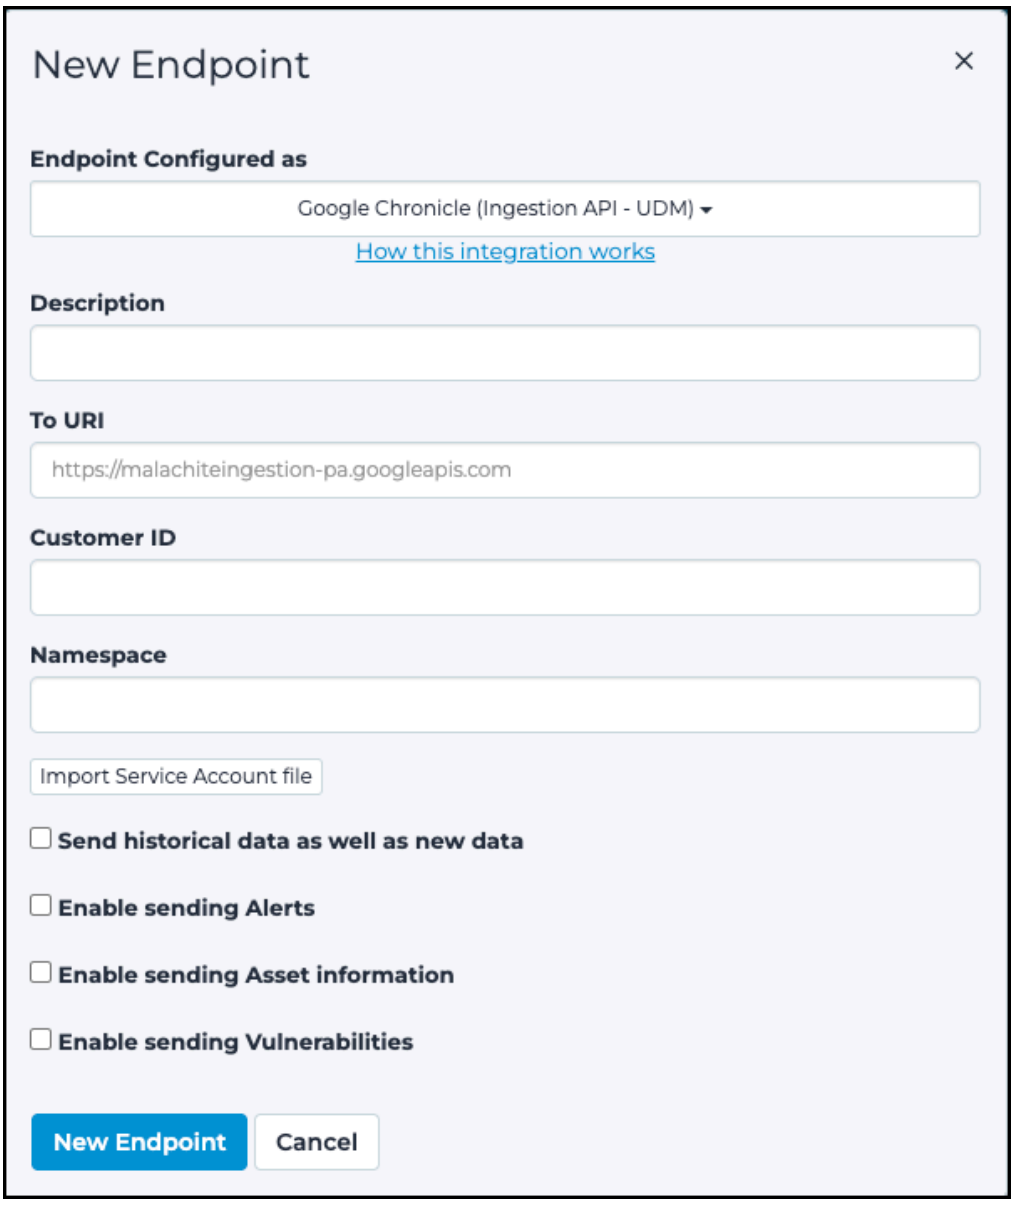


Figure 29. Google Chronicle dialog

IBM QRadar (LEEF)

The IBM QRadar integration lets you send all alerts, and optionally health logs, in LEEF format. You can also send asset information to QRadar beginning with version 2.0.0 of the QRadar App.

You can select **How this integration works** to view additional details.

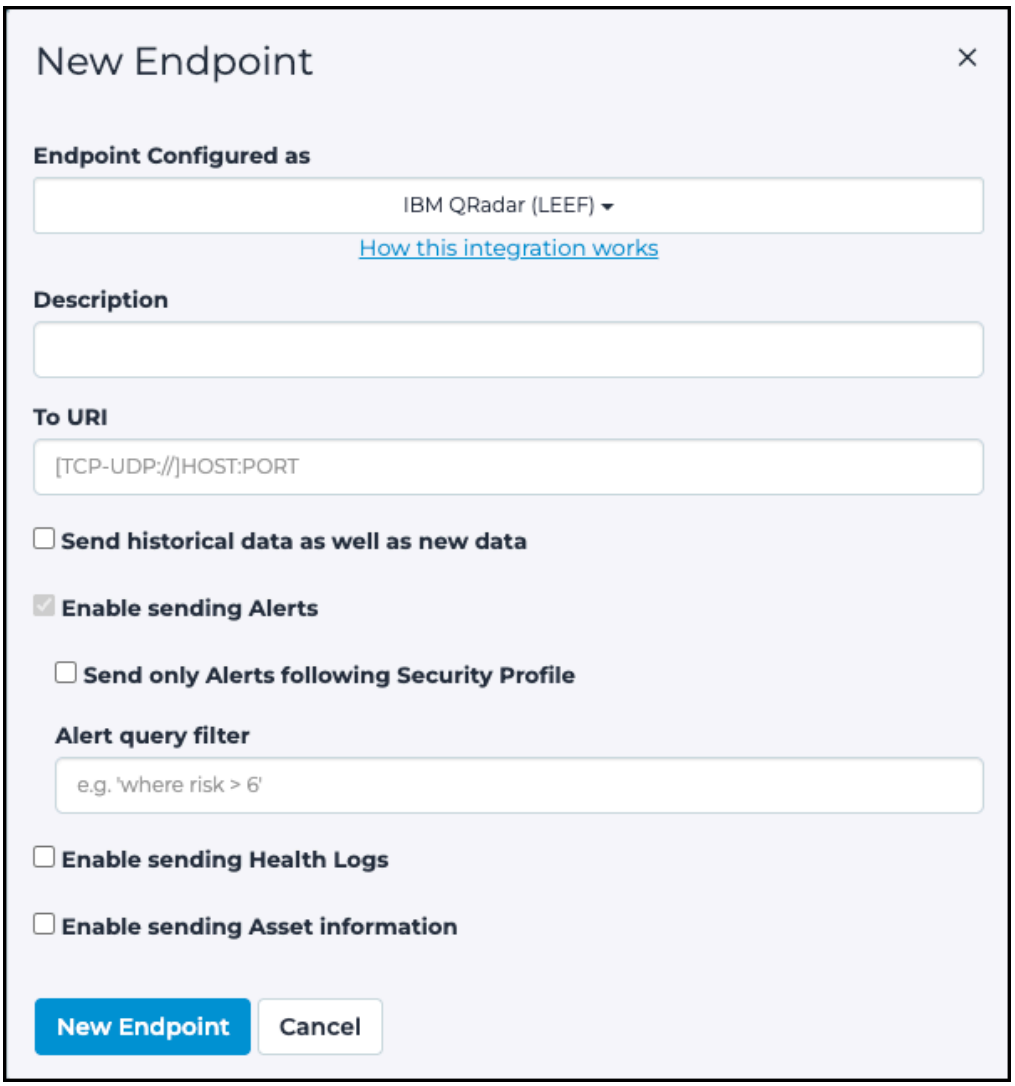


Figure 30. IBM QRadar (LEEF) dialog

Common Event Format (CEF)

*CEF lets you send alerts and health logs in CEF format. You can also enable encryption of the data through the TLS checkbox and check the validity of the CEF server's certificate with the CA-emitted **TLS** certificate checkbox.*

You can select **How this integration works** to view additional details.

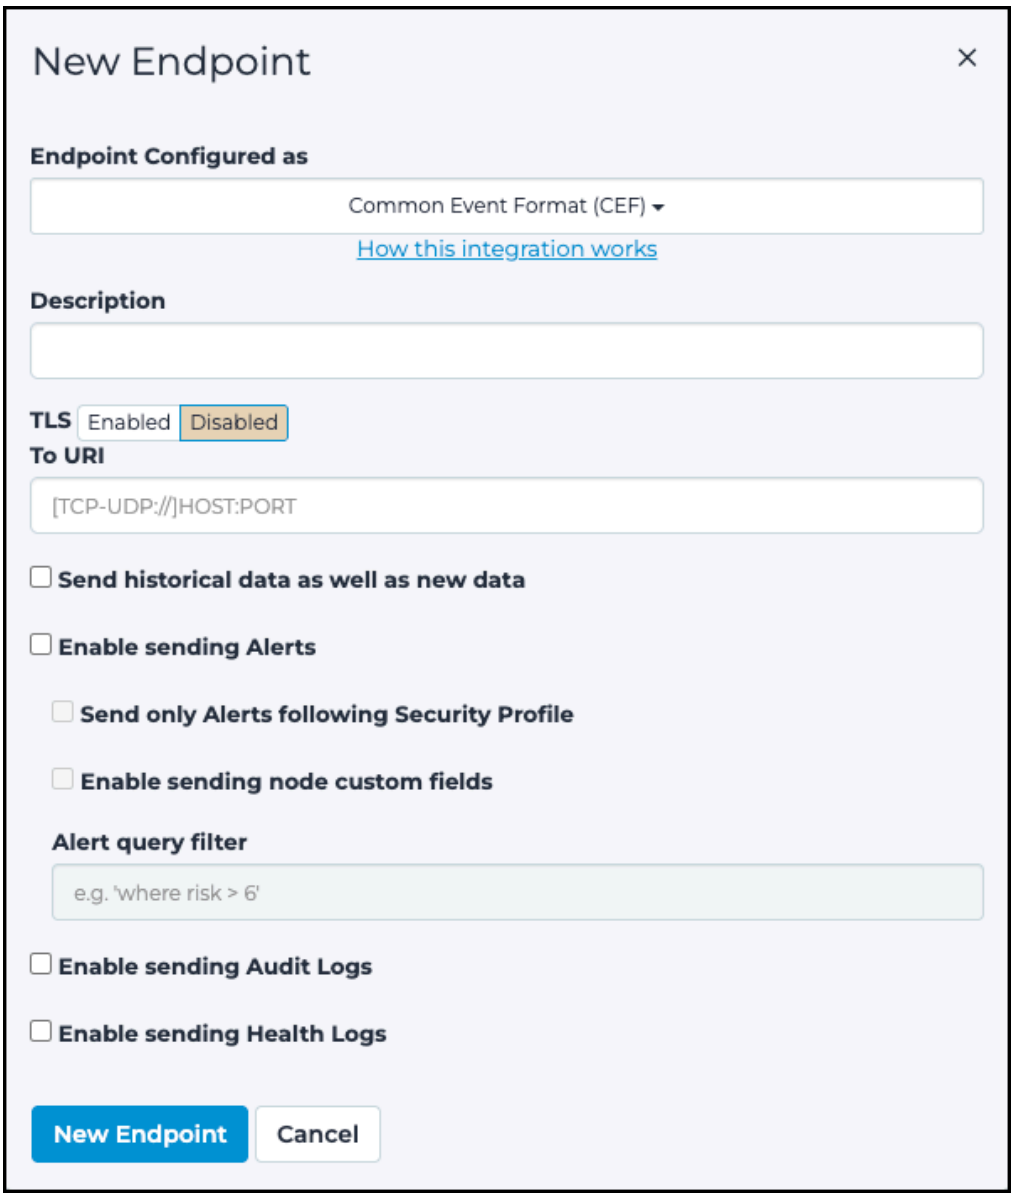


Figure 31. Common Event Format (CEF) dialog

Custom fields

The Nozomi Networks solution has defined custom label fields in our common event format (CEF) implementation. Ensure that your integration recognizes these custom labels and deals with them appropriately.

Field Value	Label Value	Label Sample	Field Sample
cs1	cs1Label	Risk	Risk level for the alert
cs2	cs2Label	IsSecurity	Is this a security alert
cs3	cs3Label	Id	Alert ID (not Alert Type ID) of the alert in the Nozomi system
cs4	cs4Label	Detail	Alert details
cs5	cs5Label	Parents	Parent IDs of the alert if related to others
cs6	cs6Label	n2os_schema	This is the Nozomi Schema version
flexString1	flexString1Label	mitre_attack_techniques	T0843
flexString2	flexString2Label	mitre_attack_tactics	Impair Process Control, Inhibit Response Function, Persistence
flexString3	flexString3Label	Name	Suspicious Activity

The *common event format (CEF)* data integration now sends the name attribute of alerts in the `flexString CEF` field. For example:

```
nozomi-ids.local n2osevents[0]: CEF:0|Nozomi Networks|N2OS|
21.9.0-01051414_C13FC|SIGN:MULTIPLE-UNSUCCESSFUL-LOGINS|Multiple
unsuccessful logins|8|
app=smb
dvc=172.16.193.105
dvchost=nozomi-ids.local
cs1=8.0
cs2=true
cs5=["22114bf0-813c-434c-b4d7-933d2a54b4e1"]
cs6=3 cs1Label=Risk
cs2Label=IsSecurity
cs3Label=Id
```

```
cs5Label=Parents
cs6Label=n2os_schema
flexString1=T0843
flexString1Label=mitre_attack_techniques
flexString2=impair_process_control, inhibit_response_function,
 persistence
flexString2Label=mitre_attack_tactics
flexString3=suspicious_activity
flexString3Label=name
dst=192.168.1.77
dmac=f0:1f:af:f1:40:5c
dpt=445
msg=Multiple unsuccessful logins detected with protocol smb. The
 usernames
'', 'DOMAIN\VCA07_12$' attempted at least 40 connections in 15 seconds
src=192.168.1.227
smac=d8:9e:f3:3a:cb:3a
spt=57280
proto=TCP
start=1651456283700
```

ServiceNow

The ServiceNow integration allows you to receive asset information from a ServiceNow instance and enrich local nodes.

You can select **How this integration works** to view additional details.

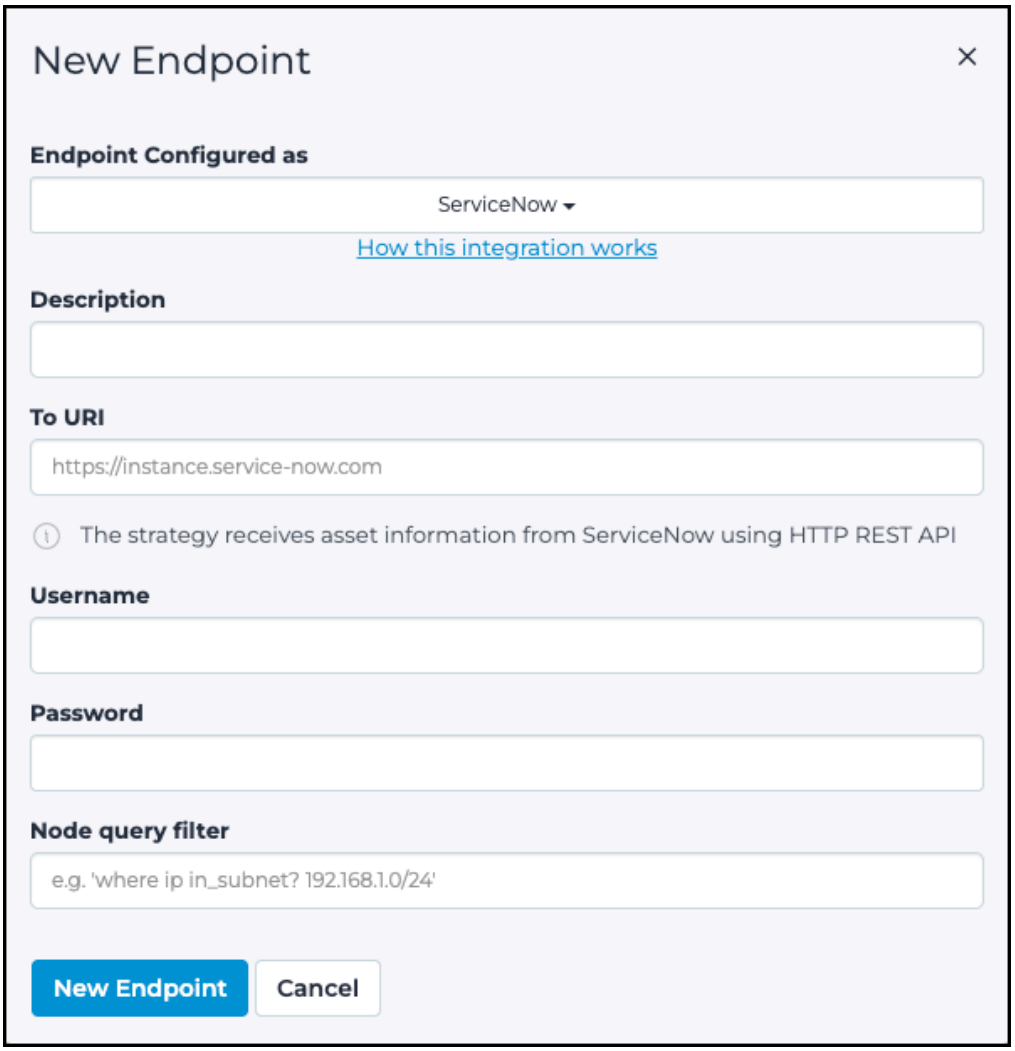


Figure 32. ServiceNow dialog

Tanium

This integration enables seamless forwarding of node and asset data to a Tanium instance, facilitating centralized management and enhanced visibility within your network.

You can select **How this integration works** to view additional details.

Note:
If the Tanium instance does not have a valid signed *HTTPS certificate authority (CA)*, users must add an ! before the *URL*. For example, `!https://192.168.1.1`

Note:
Nodes are sent, not assets.

Note:
Nodes are sent regardless of whether *MAC* addresses are confirmed or not (all nodes).

Note:
If integrating with the *CMC*, use All-In-One mode. This is because multicontext a *CMC* does not have nodes.

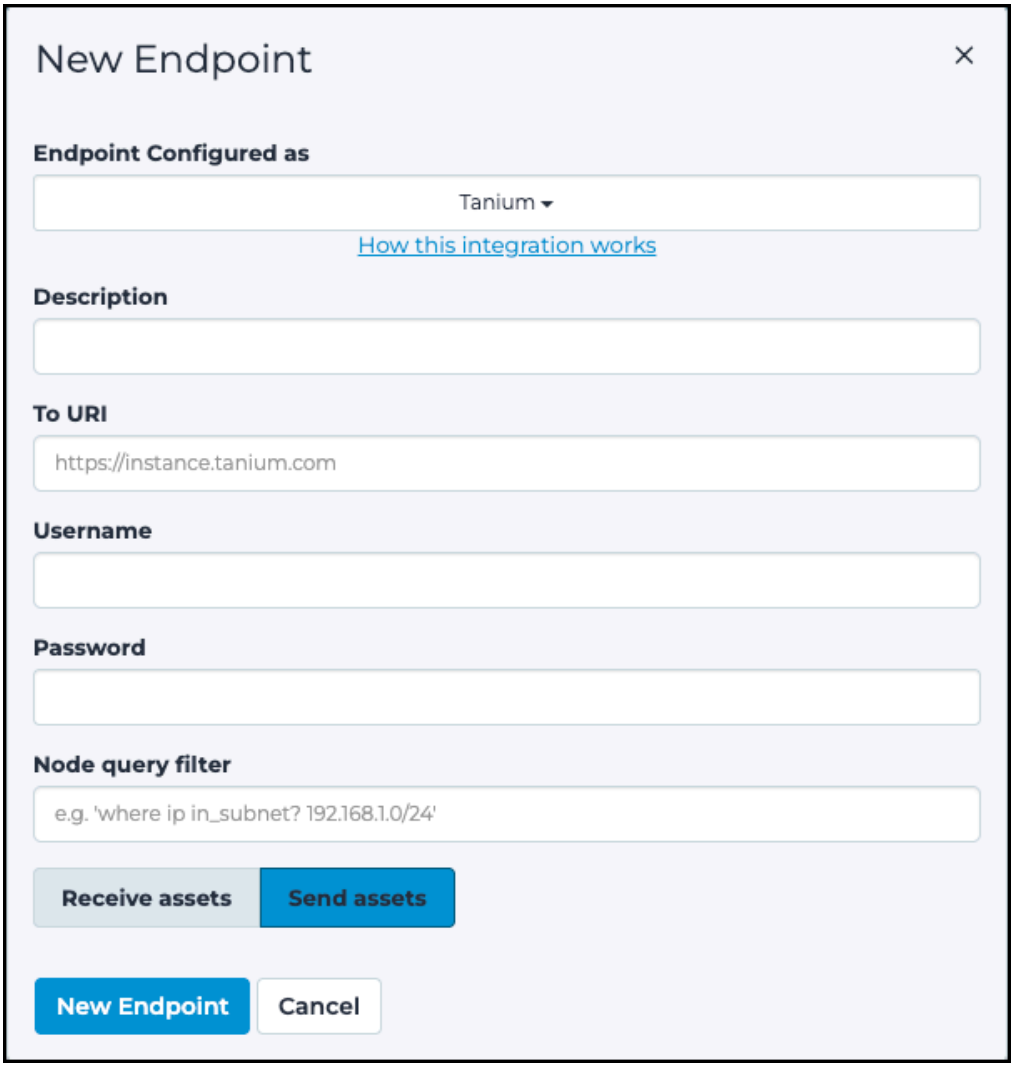


Figure 33. Tanium dialog

Receive assets

The integration can be configured to receive asset data from Tanium. When configuring the integration to receive data, it is important to ensure that only one integration at a time is designated to interact with a specific endpoint. This means that if you have an integration set to receive data from an endpoint, another integration should not be set to send data to that same endpoint. This will help to avoid potential conflicts, or data inconsistencies.

Send assets

The integration can be configured to send asset data to Tanium. When configuring the integration to send data, it is important to ensure that only one integration at a time is designated to interact with a specific endpoint. This means that if you have an integration set to send data to an endpoint, another integration should not be set to receive data to that same endpoint. This will help to avoid potential conflicts, or data inconsistencies.

Splunk - Common Information Model (JSON)

If you need to send alerts to a Splunk - JavaScript Object Notation (JSON) instance, you can use integration. Data are sent in JSON format and you are also able to filter on alerts. You can also send health logs and audit logs.

You can select **How this integration works** to view additional details.

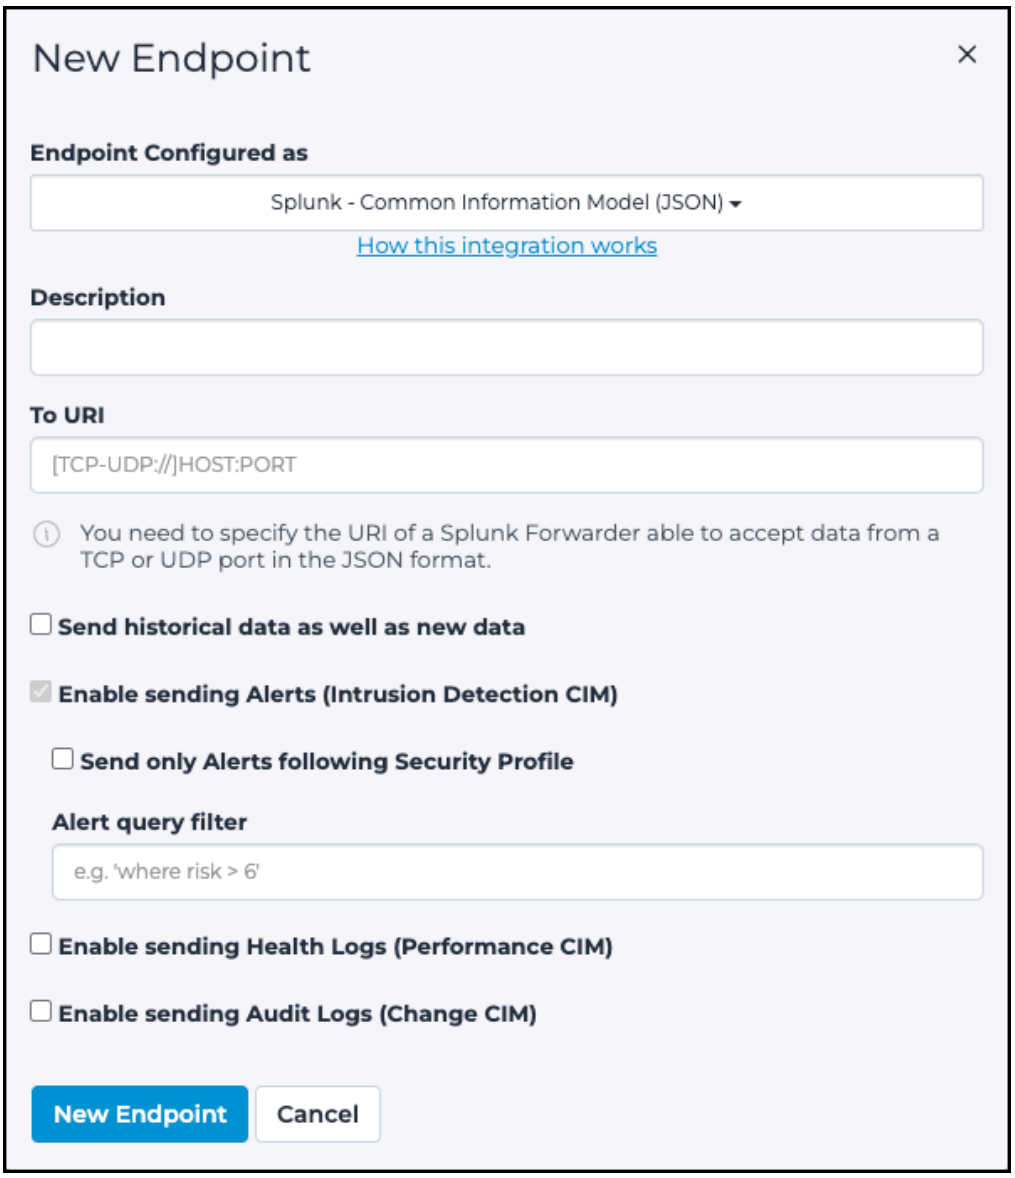


Figure 34. Splunk dialog

SMTP forwarding

To send reports, alerts and/or health logs to an email address, you can configure a simple mail transfer protocol (SMTP) forwarding endpoint. In this case, you are also able to filter alerts.

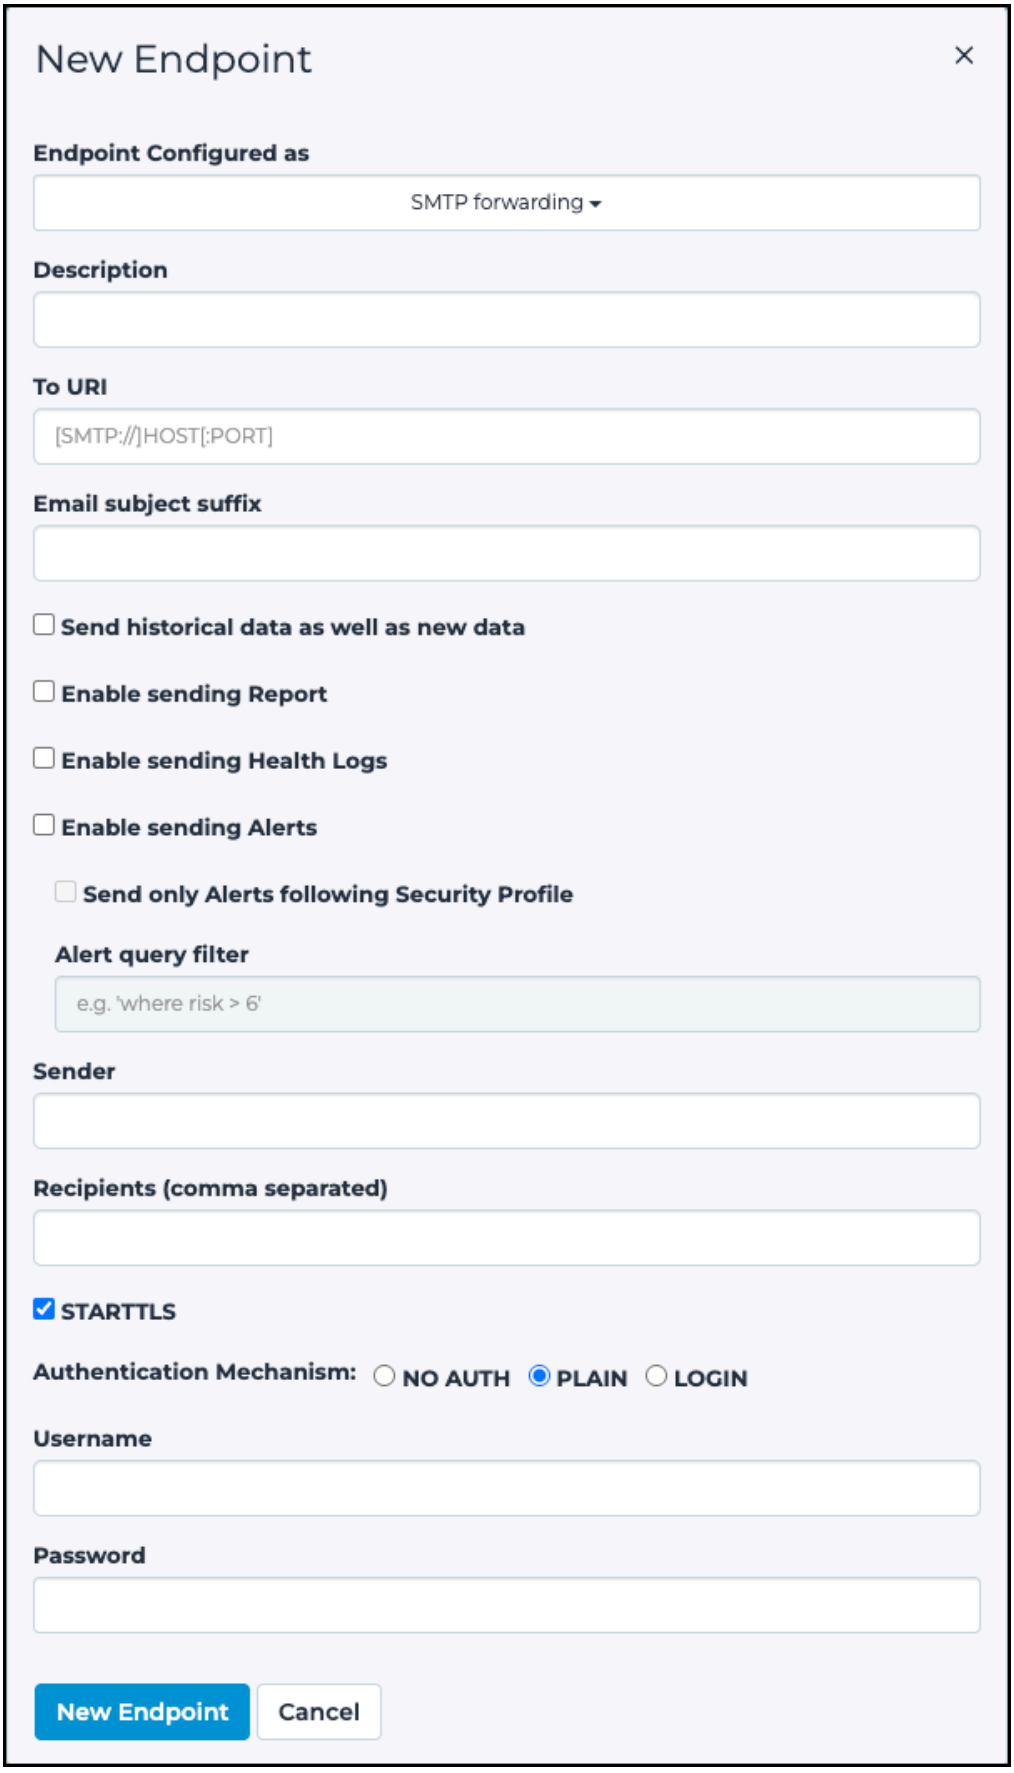


Figure 35. SMTP forwarding dialog

SNMP trap

Use this kind of integration to send alerts through a simple network management protocol (SNMP) trap.

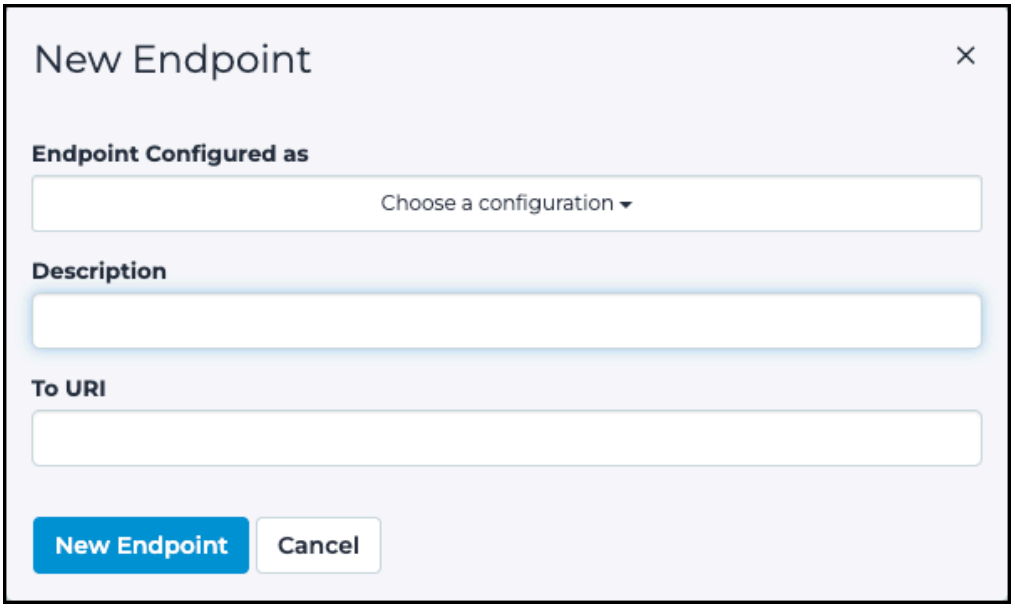


Figure 36. SNMP trap dialog

Syslog Forwarder

Use this type of integration to send syslog events captured from monitored traffic to a syslog endpoint.

It is useful for passively capturing logs and forwarding them to a *SIEM*.

Note:
In order to enable syslog events capture see Enable Syslog capture feature in the Basic configuration section of the Configuration chapter of this manual.

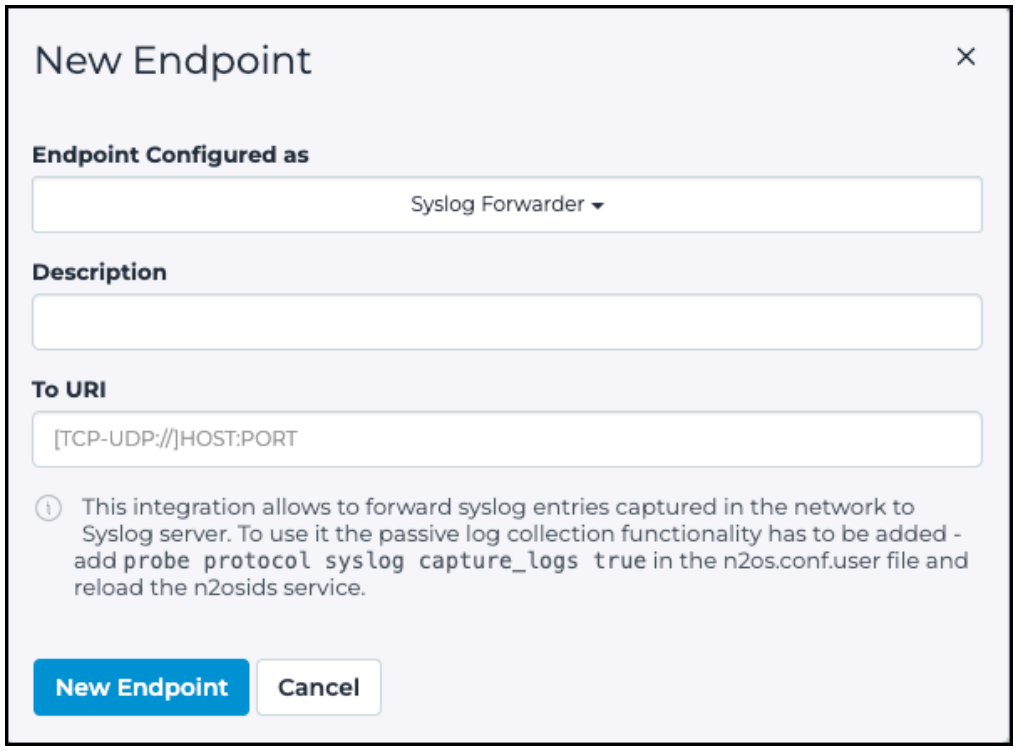


Figure 37. Syslog Forwarder dialog

Custom JSON

This type of integration uses the JavaScript Object Notation (JSON) format to send the results of the related query to a specific uniform resource identifier (URI).

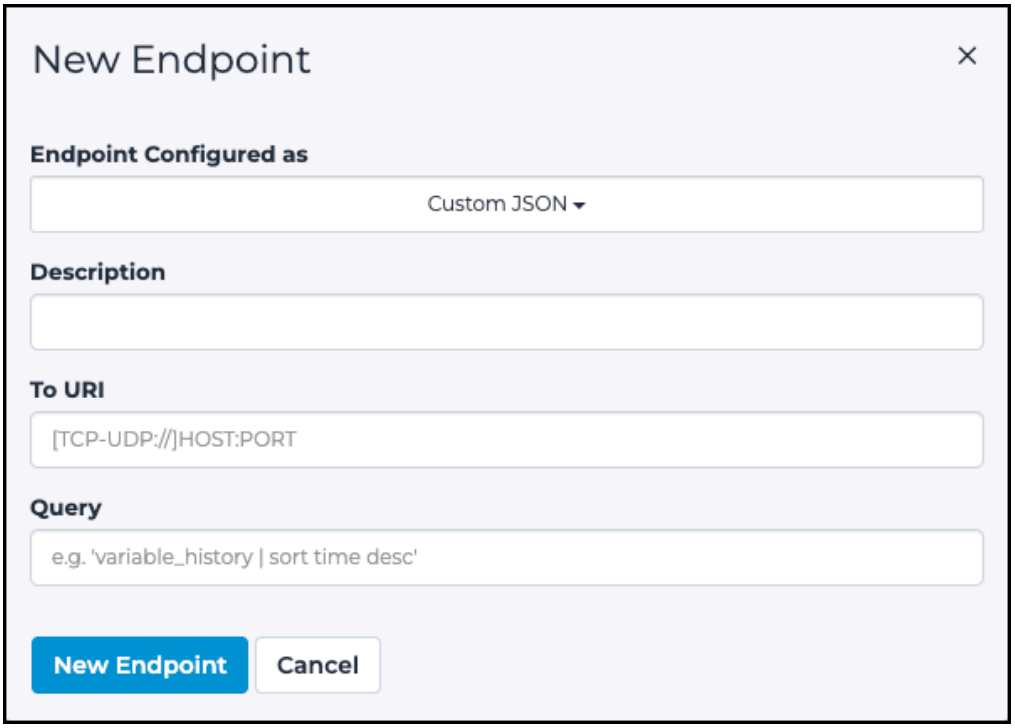


Figure 38. Custom JSON dialog

Custom CSV

This type of integration sends the results of the specified query to a specific URI in comma-separated value (CSV) format.

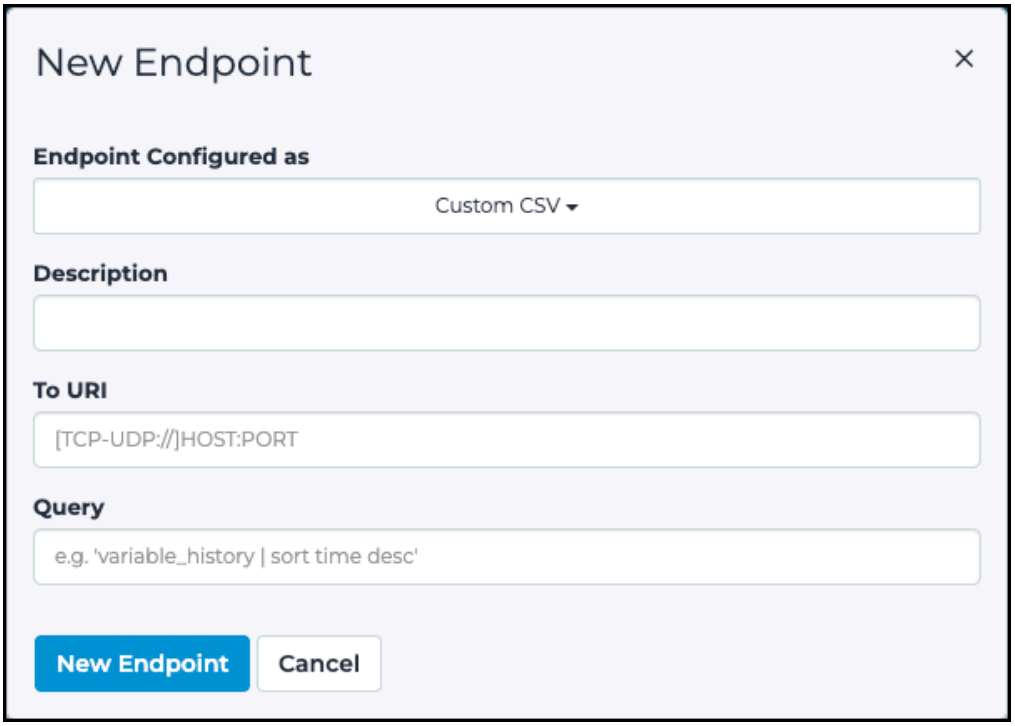


Figure 39. Custom CSV dialog

DNS Reverse Lookup

This integration sends reverse domain name server (DNS) requests for the nodes in the environment and uses the names provided by the DNS as nodes' labels. You can pre-filter the nodes by specifying a query filter.

The strategy runs once a day by default, but you can select **Rerun the strategy on all the data** to run it on demand.

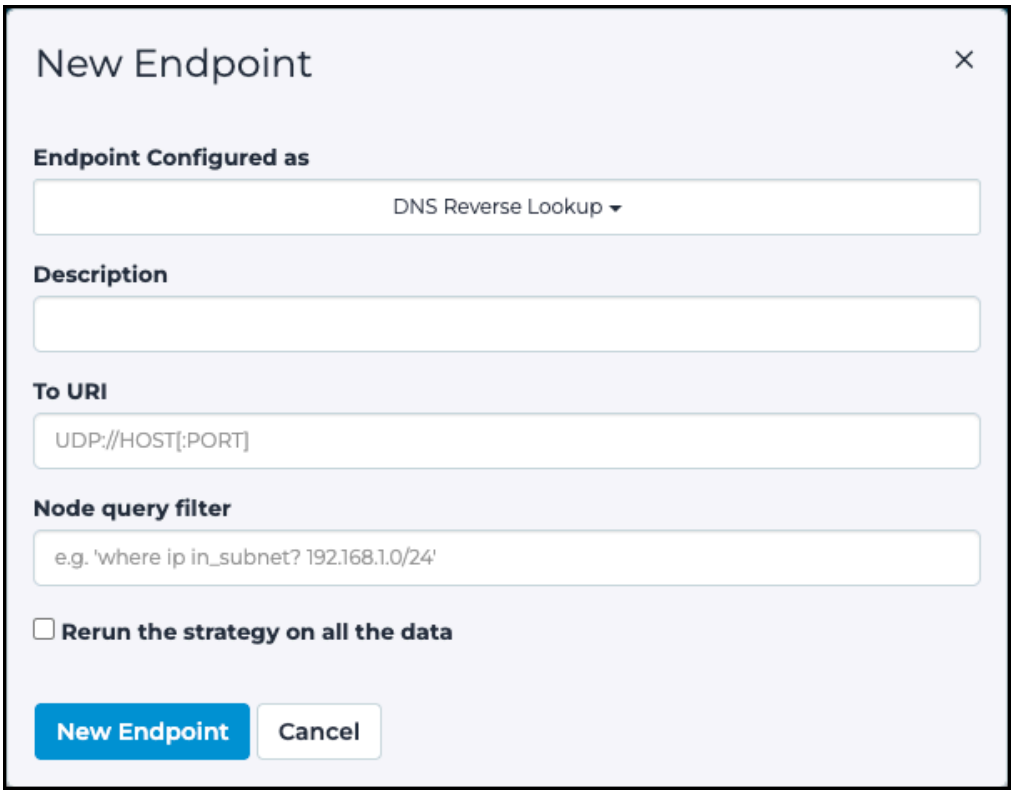


Figure 40. DNS Reserve Lookup dialog

Kafka

The Kafka integration allows you to send the results of custom queries in JavaScript Object Notation (JSON) format to existing topics of a Kafka cluster.

You can select **How this integration works** to view additional details.

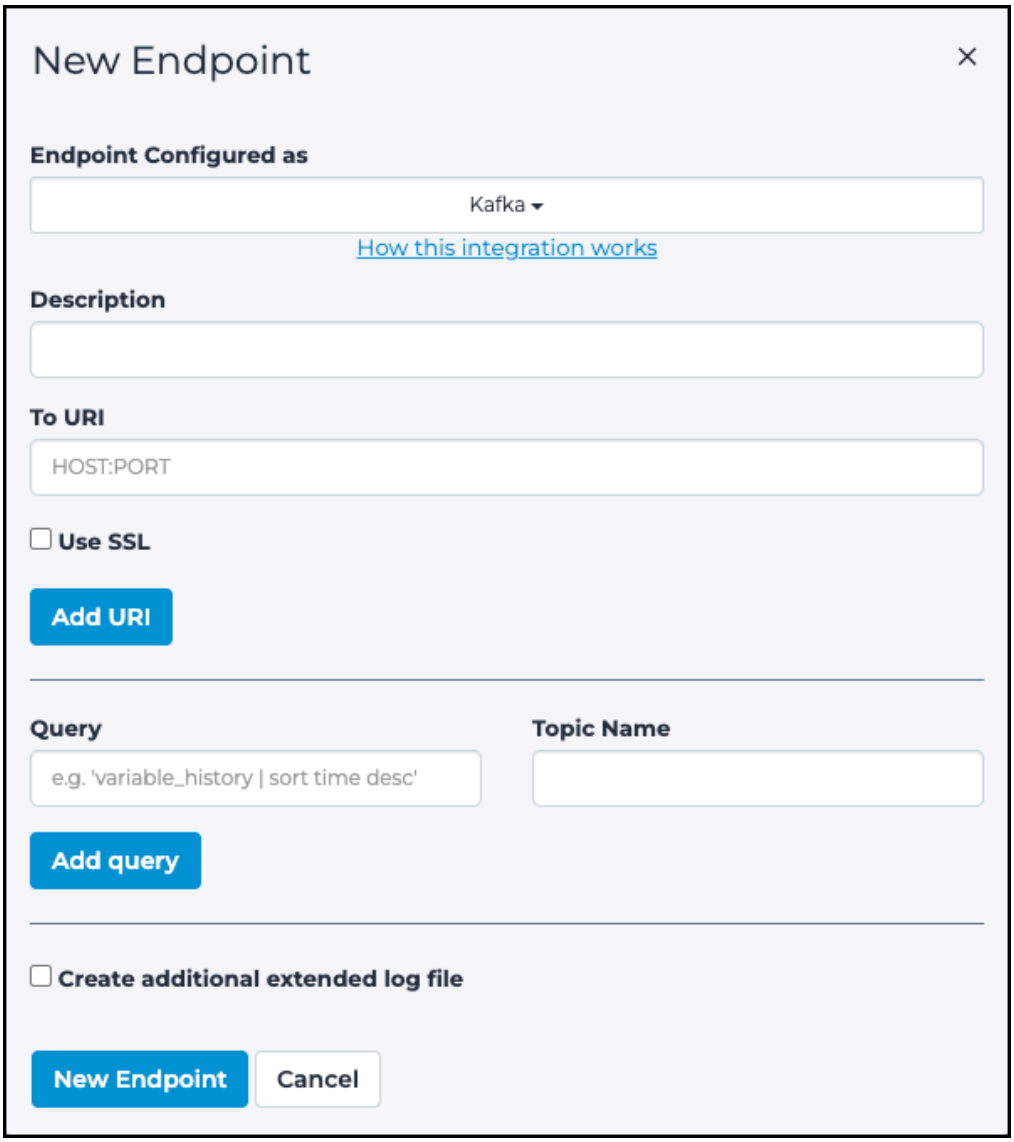


Figure 41. Kafka dialog

Cisco ISE

With the Cisco ISE integration, you can use the simple text-oriented messaging-protocol (STOMP) to send the results of custom node queries to Cisco's ISE asset information, or you can use hypertext transfer protocol (HTTP) to receive asset information from a Cisco ISE instance and enrich local nodes.

You can select **How this integration works** to view additional details about certificate usage and Cisco ISE environment requirements.

Send assets configuration

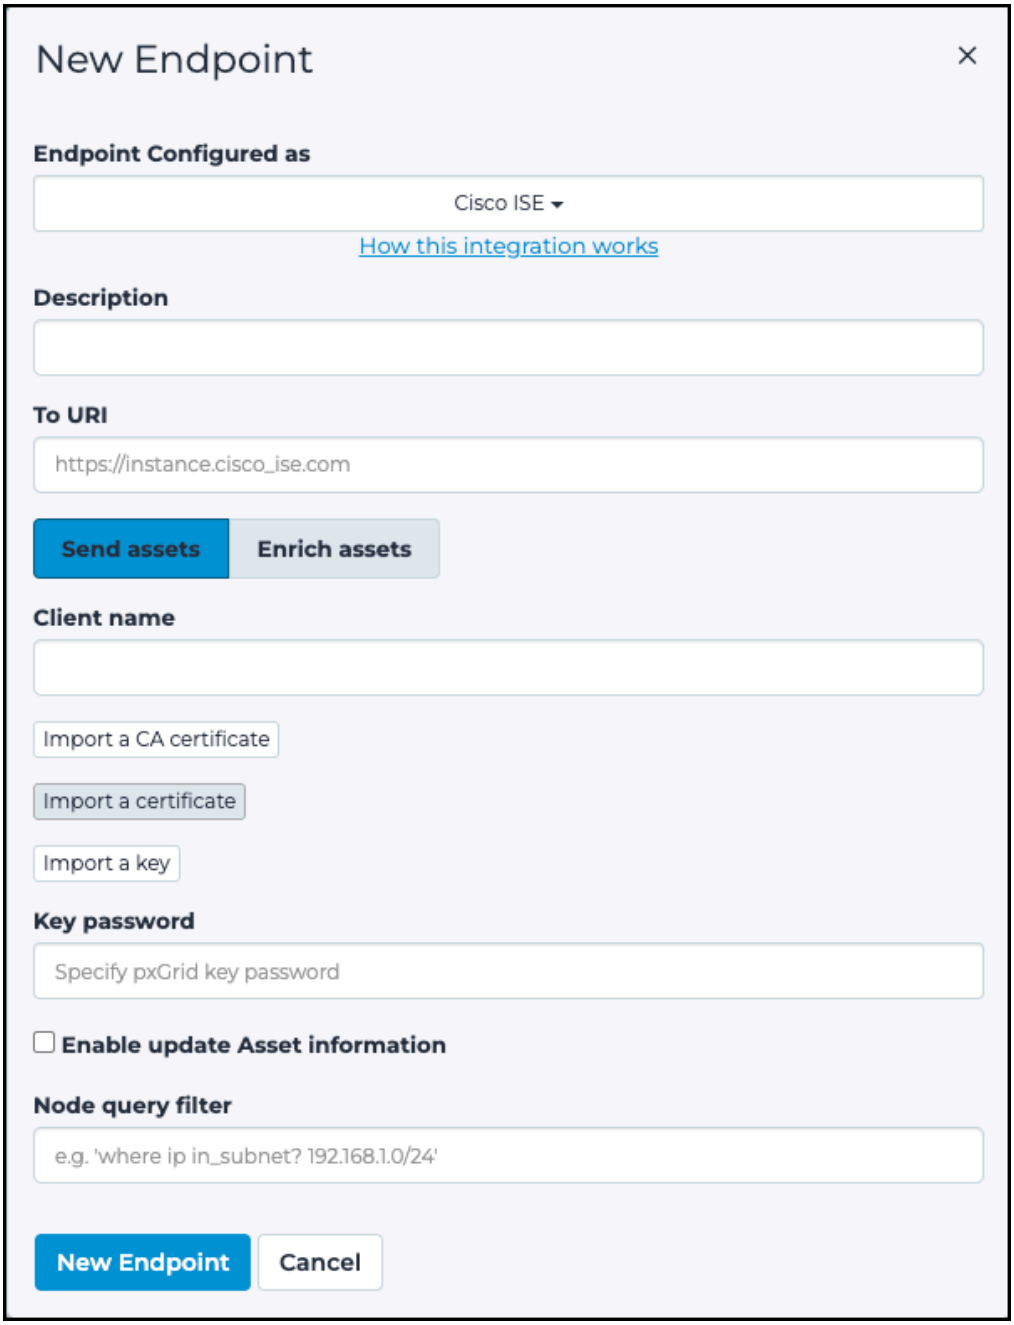


Figure 42. Send assets dialog

To configure Cisco ISE, you need to create these custom string attributes:

- `n2os_change_flag`
- `n2os_operating_system`
- `n2os_product_name`
- `n2os_vendor`
- `n2os_type`
- `n2os_appliance_site`
- `n2os_zone`

You then need to create a new profile and set the required condition `n2os_change_flag` custom attribute equal to `change`.

You then need to modify the existing profiles or, if no profiles are expected to be assigned to assets from n2os, create a new profile. Add the required condition for `n2os_change_flag`.

Note:

Due to a long-standing bug in Cisco's PxGrid *API*, the performance when sending assets is halved, requiring two network calls for each updated record. Nozomi Networks is working with Cisco to address this issue. Cisco has not provided a target date for this bug.

Enrich assets

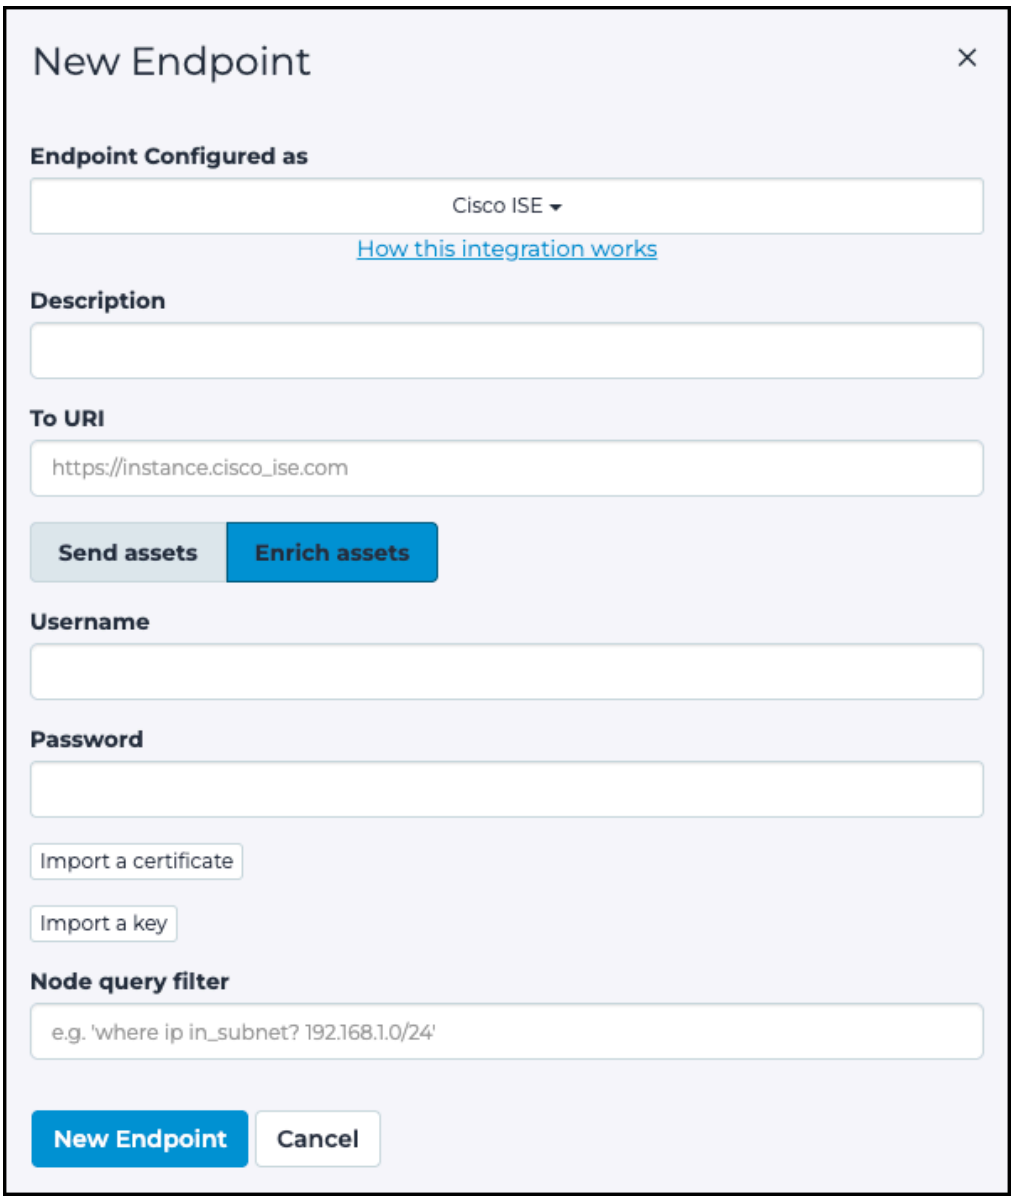


Figure 43. Enrich assets dialog

External Storage

The external storage integration uploads files to an external machine. This enables the external machine to keep remote copies of files that are kept beyond the retention settings.

The file location becomes transparent to the user, who can retrieve them seamlessly from external storage when the files are removed from the local file system. You can also choose a connection protocol for storing the files. Available protocols are *SMB*, *file transfer protocol (FTP)*, and *SSH*.

Important:

Microsoft *OS*s are the only *OS* that support the *SMB* connection protocol.

Compatibility with third-party devices is not guaranteed. These devices might require additional configuration changes, that include:

- Permission changes
- The creation of new network shares
- The creation of new users

Kerberos authentication is not supported.

This functionality is currently only available for trace *packet capture (pcap)* files on Guardian.

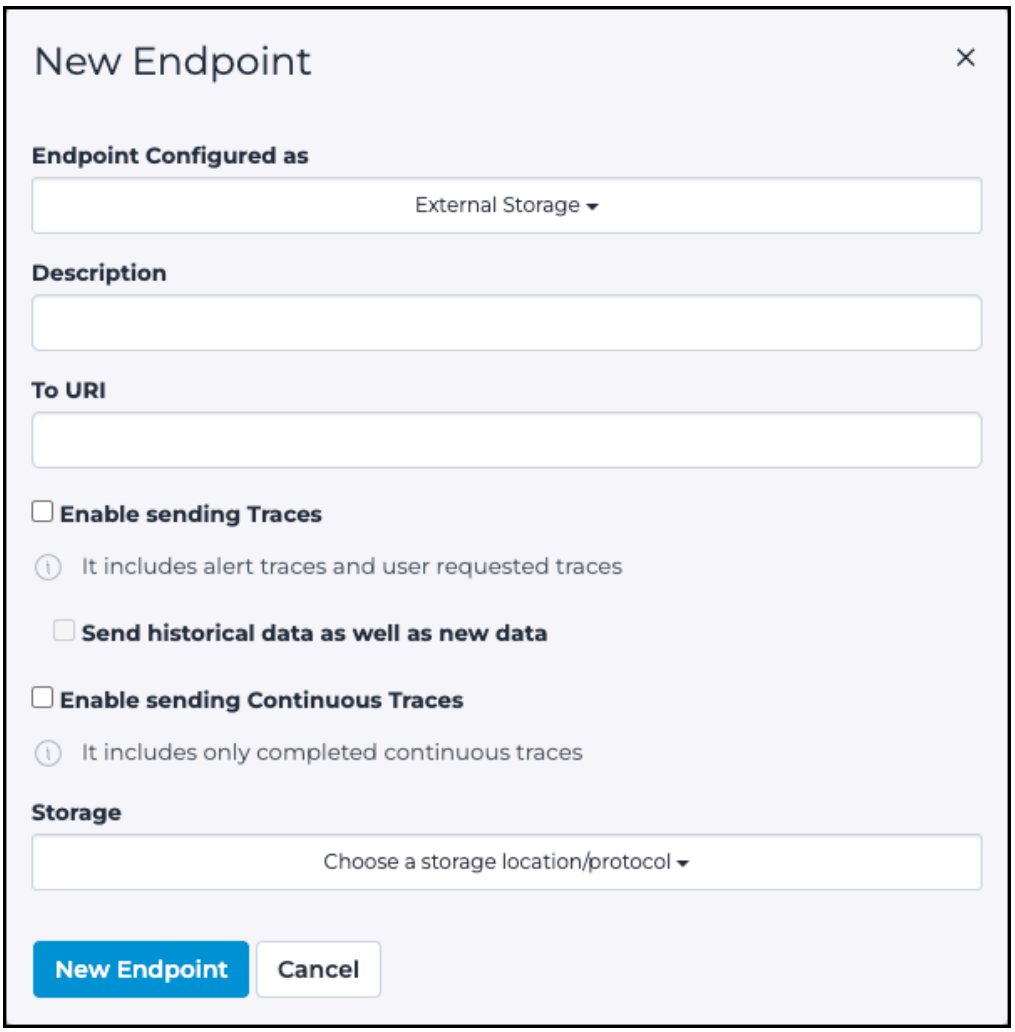


Figure 44. External Storage dialog

Microsoft Endpoint Configuration Manager (WinRM RPC)

With the Windows Remote Management (WinRM) RPC, you can collect information coming from the Microsoft Endpoint Configuration Manager to update Windows nodes.

Collected items:

- **OS information**: Returns *OS* information as version, service pack, build and architecture
- **Hostnames**: Returns host name information to configure the node label
- **Interfaces information**: Returns interface data to populate the node *MAC* address
- **Installed software**: Returns installed software and populates the node *Common Platform Enumeration (CPE)*
- **Hotfixes**: Returns installed software version updates and checks to see if there are node *CVE*s to close



Note:

It is important to filter the strategy nodes because without the filter, the strategy waits for the timeout of non-Microsoft nodes that are not reachable. This significantly decreases the performance of the data integration strategy. If not specified, the port is the default: 5986 if *SSL* is enabled, 5985 otherwise.

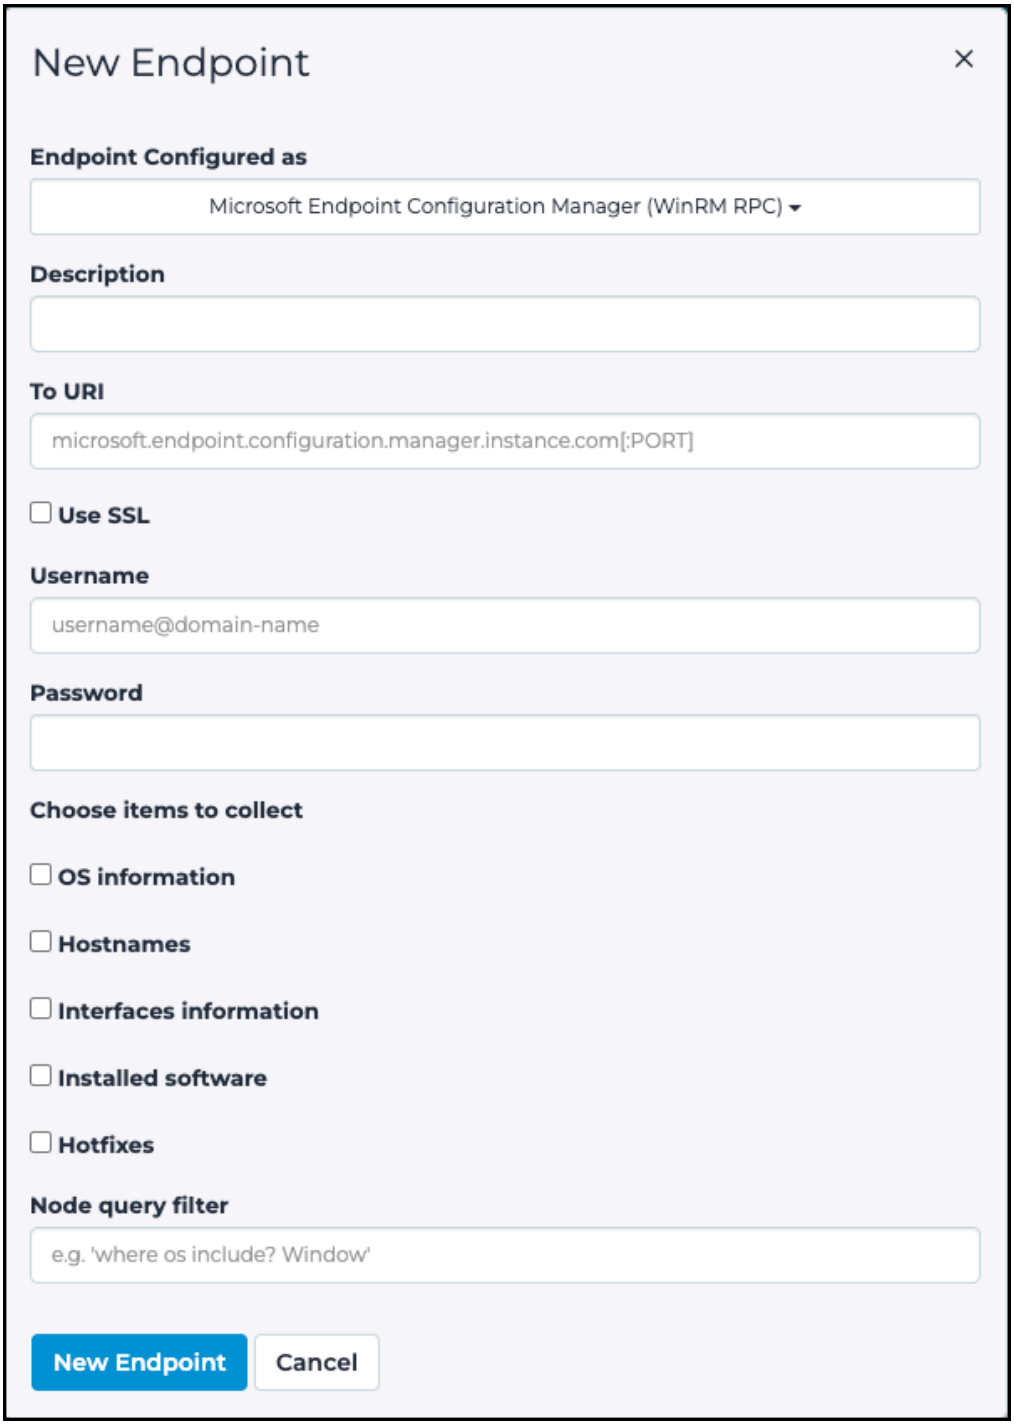


Figure 45. Microsoft Endpoint Configuration Manager (WinRM RPC) dialog

Microsoft Endpoint Configuration Manager (DB)

The goal of this integration is to collect information from the Microsoft Endpoint Configuration Manager DB to update the existing Windows nodes.

Collected items:

- **OS information**: Returns *OS* information as version, service pack, build and architecture
- **Hostnames**: Returns host name information to configure the node label
- **Interfaces information**: Returns interface data to populate the node *MAC* address
- **Installed software**: Returns installed software and populates the node *CPE*
- **Hotfixes**: Returns installed software version updates and checks to see if there are node *CVE*s to close



Note:

It is important to filter the strategy nodes because without the filter, the unreachable non- Microsoft nodes time out, significantly decreases timing of the data integration. If not specified, the port is the default: 5986 if *SSL* is enabled, 5985 otherwise. The database name default value format is CM_[Site code], which may be changed by an administrator.

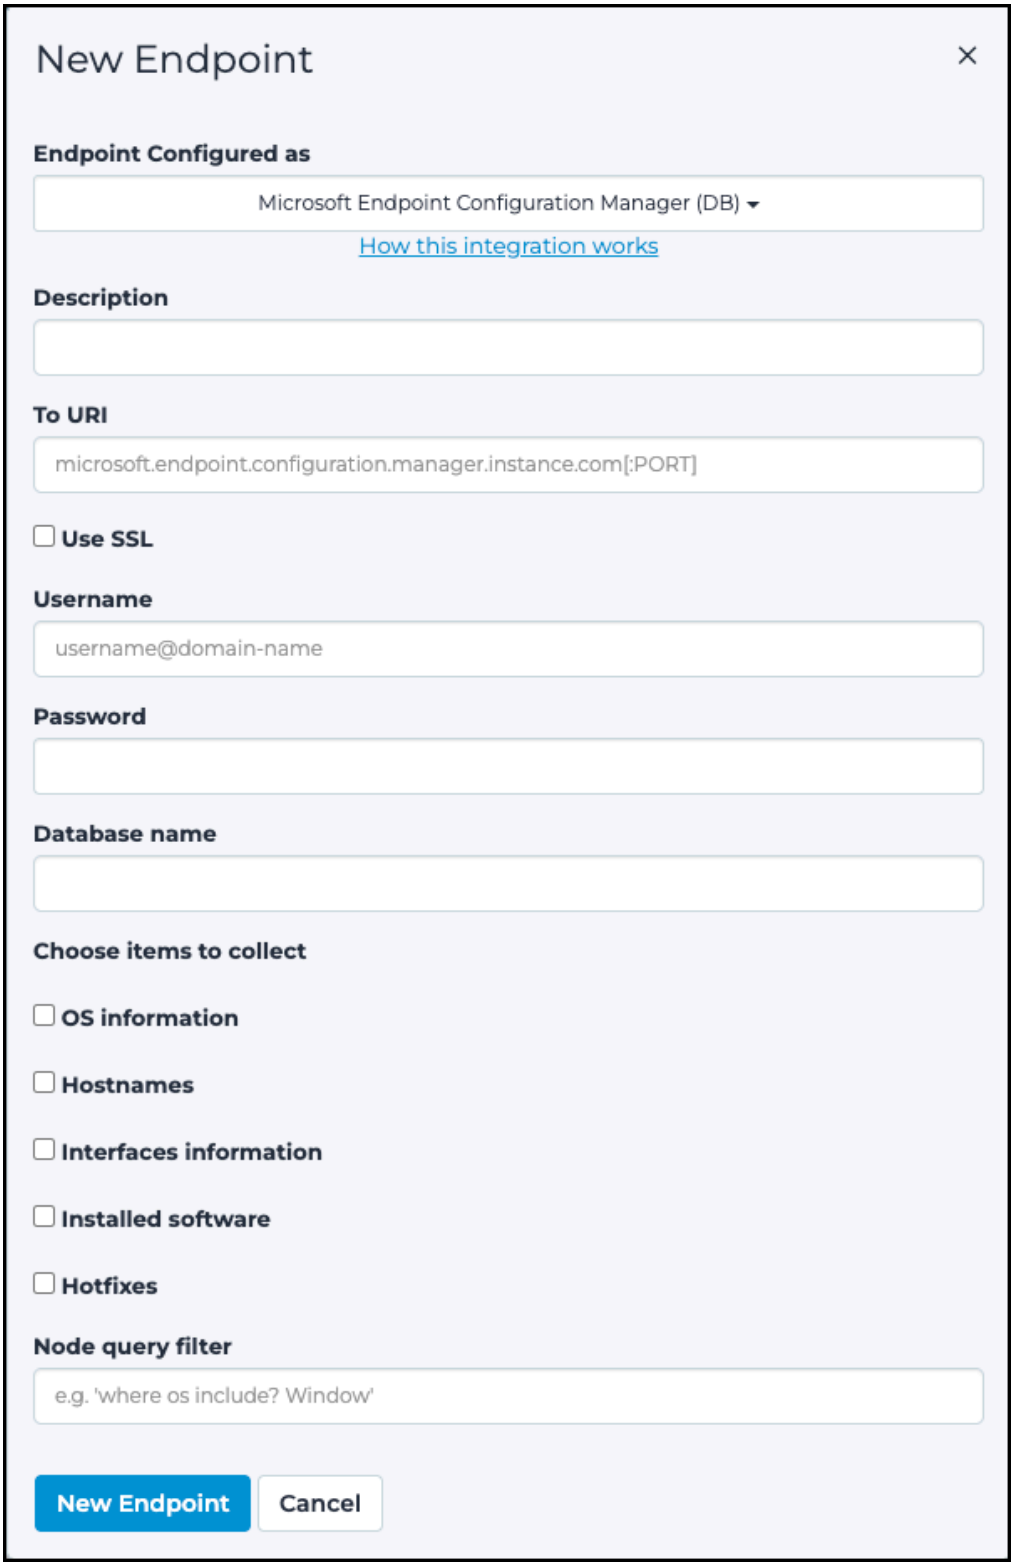


Figure 46. Microsoft Endpoint Configuration Manager (DB) dialog

NetWitness

If you require sending alerts to a NetWitness instance, you can utilize this integration. Data is transmitted in JavaScript Object Notation (JSON) format prepended by the **NOZOMI:** *header. You can also apply filters, and determine whether historical data should be sent or not.*

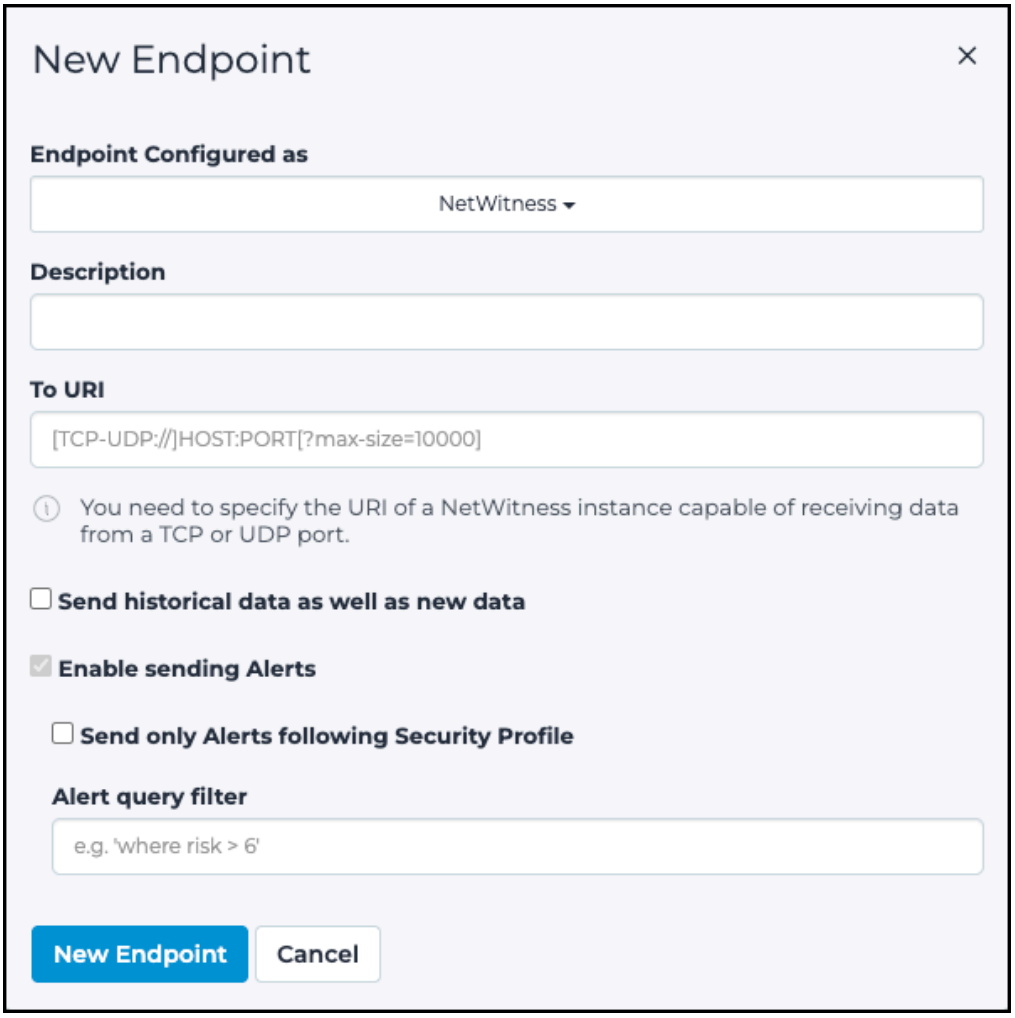


Figure 47. NetWitness dialog

Aruba ClearPass

*This integration lets you send node information to Aruba ClearPass. Only nodes with confirmed media access control (MAC) addresses are sent. You can specify a query filter to pre-filter the nodes. If **Enable update Node information** is not selected, nodes are only sent once.*

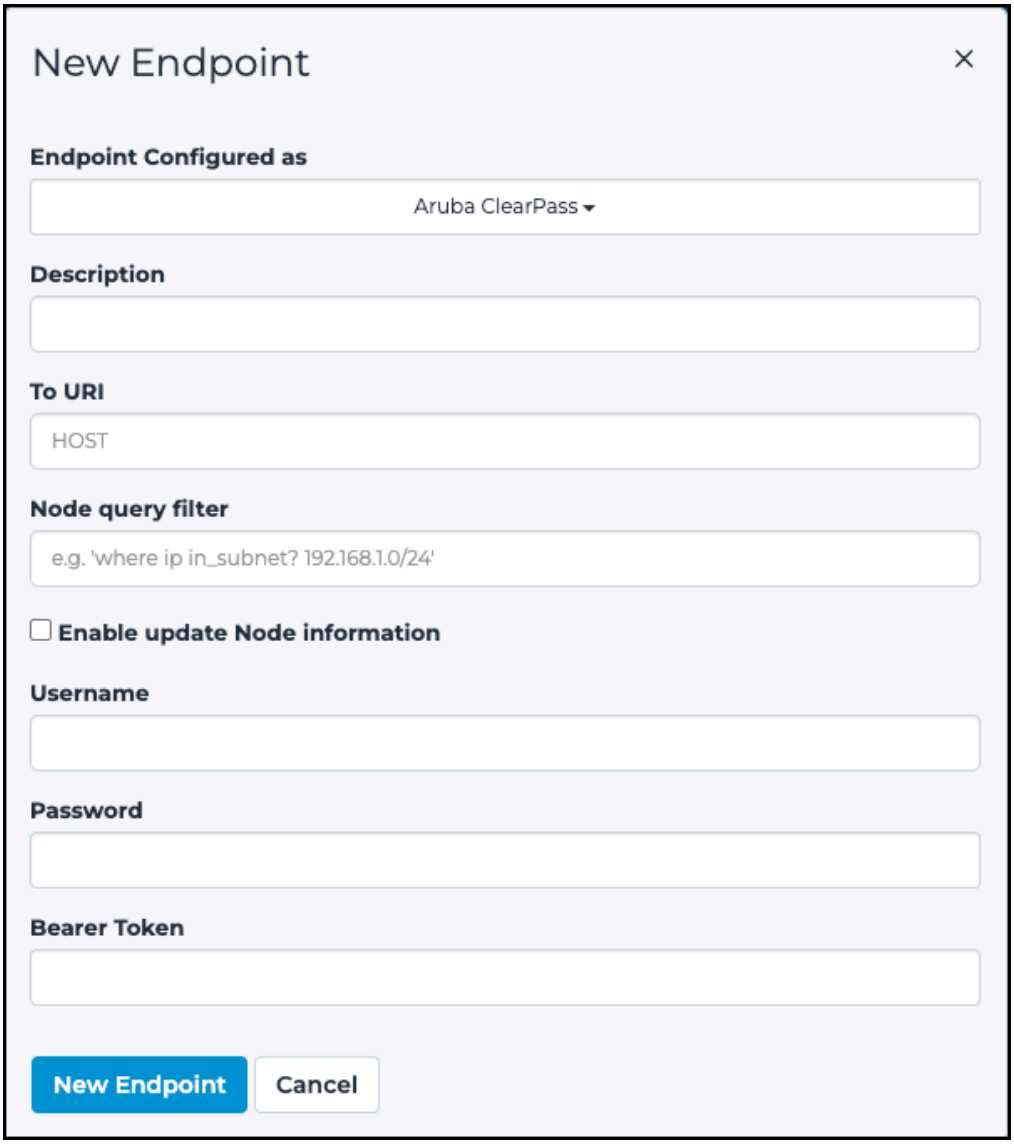


Figure 48. Aruba ClearPass dialog

CarbonBlack Defense

The strategy runs once a day by default. The strategy uses its representational state transfer (REST) application programming interface (API) over hypertext transfer protocol secure (HTTPS) to extract information about nodes from an external CarbonBlack service.

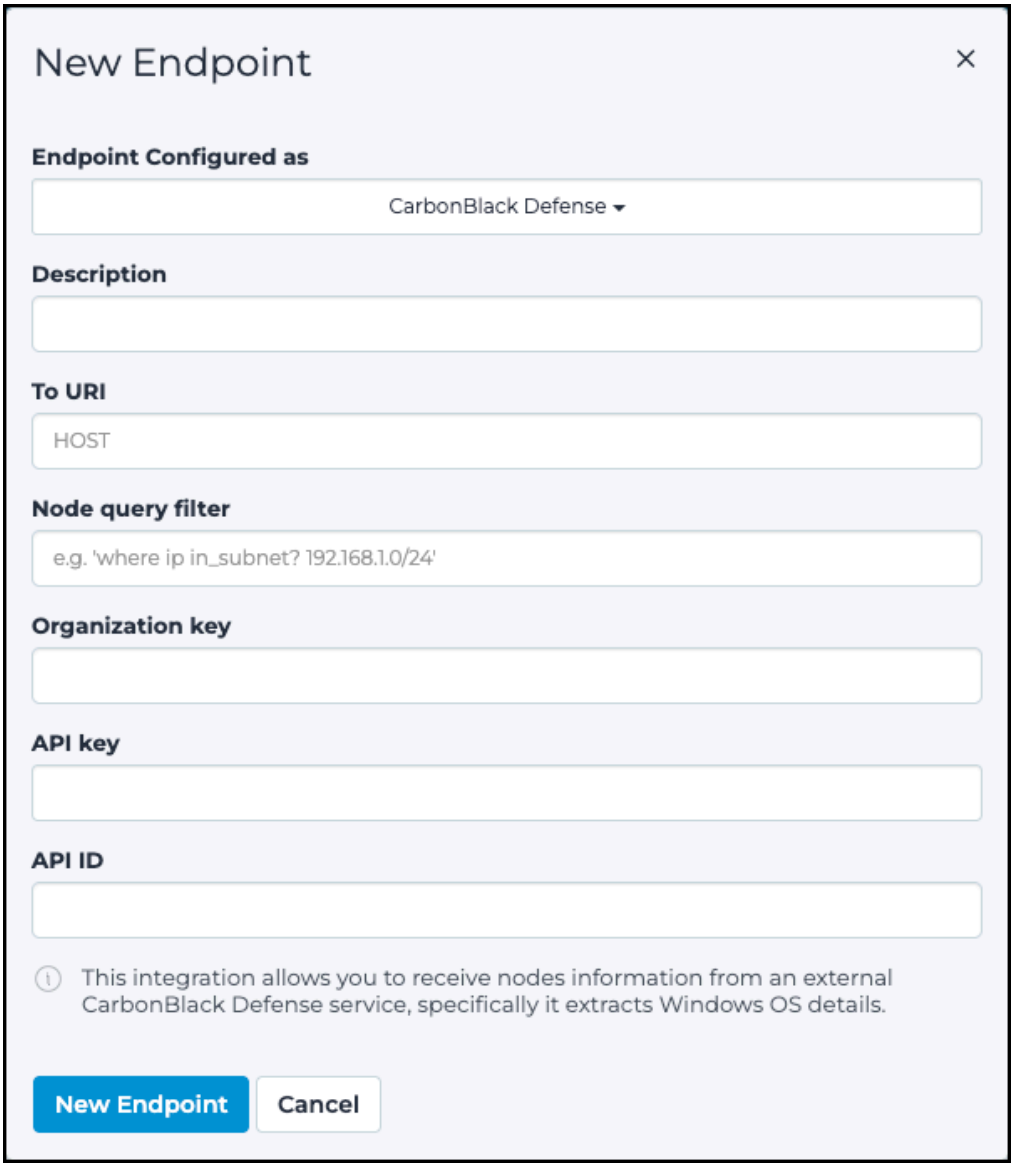


Figure 49. CarbonBlack Defense dialog

Active Directory

The integration allows you to enrich nodes with information of objects that have a `Computer` *class coming from an Active Directory server. Fields being enriched are:* `operating system`, `managedBy` *(inside properties) and* `type`.

We recommend that you also set up DNS Reverse Lookup in order to enrich nodes with their hostname. This is because this integration matches nodes with objects in Active Directory based on the DNS Host Name value found in Active Directory.

If you select TLS to be enabled, the integration will use *LDAPS* instead of *LDAP*.

To use the integration, and search objects in Active Directory, it is necessary to have a user that can perform *LDAP* queries.

Search base refers to the base of the subtree in which the search is to be constrained.

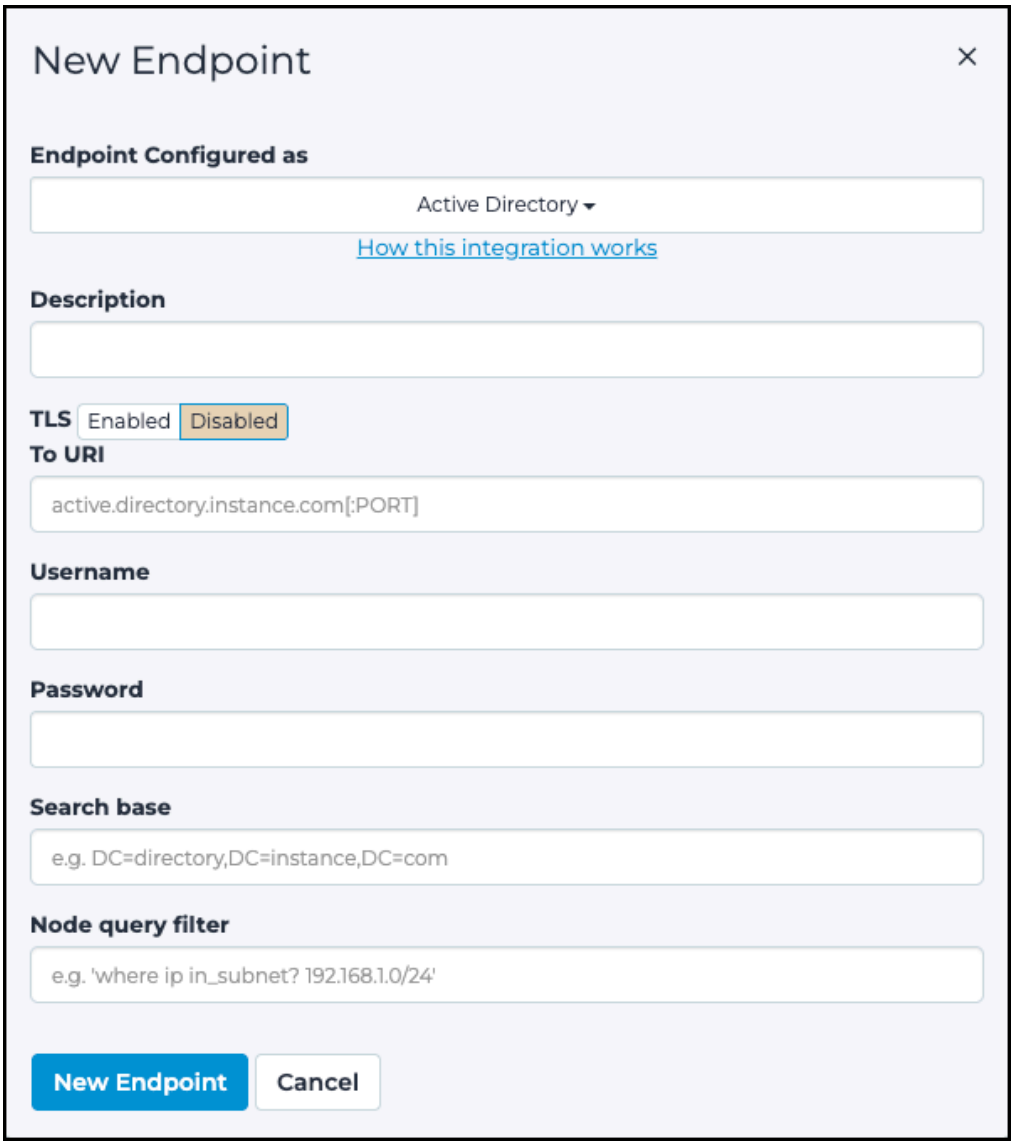


Figure 50. Active Directory dialog

Nozomi syslog data events and syslog messages

For customers implementing syslog, Guardian generates three types of syslog events: alerts, health, and audit.

Note:
As the set of alert messages inside each alert type ID category increases over time, perform searches on alert type IDs, health type IDs, and audit type IDs, rather than on the alert message itself.

Alert events

A description of alert events in common event format (CEF).

Alert events

Alert events should be identified by the alert type ID. There are many alert types in the Nozomi Networks environment.

For a full list of alert types, see **Alerts** in the **Alerts and Incidents - Reference Guide**.

Alert events in *CEF* have the following format, as shown in this example:

```
<137>Oct 17 2019 22:32:23 local-sg-19.x n2osevents[0]: CEF:0|Nozomi

 Networks|N2OS|19.0.3-10142120_A2F44|SIGN:MALWARE-DETECTED|Malware
 detected|
                    9|
                    app=smb
                    dvc=172.16.248.11
                    dvchost=local-sg-19.x
                    cs1=9.0
                    cs2=true
                    cs3=d25c520f-7f79-4820-b5ae-d1b334b05c75
                    cs4={trigger_type: yara_rules, trigger_id:
 MALW_DragonFly2.yar}
                    cs5=["5740a157-08e8-490f-85ad-eef23657e3cb"]
                    cs6=1
                    cs1Label=Risk
                    cs2Label=IsSecurity
                    cs3Label=Id
                    cs4Label=Detail
                    cs5Label=Parents
                    cs6Label=n2os_schema
                    flexString1=T0843
                    flexString1Label=mitre_attack_techniques
                    flexString2=Impair process (etc)
                    flexString2Label=mitre_attack_tactics
                    flexString3=Suspicious Activity
                    flexString3Label=name
                    dst=172.16.0.55
                    dmac=00:0c:29:28:dd:c5
                    dpt=445
                    msg=Suspicious transferring of malware named
 'TemplateAttack_DragonFly_2_0'
                    was detected involving resource '\\172.16.0.55\ADMIN
                    \CVcontrolEngineer.docx' after a 'read' operation
 [rule author: US-CERT
                    Code Analysis Team - improved by Nozomi Networks]
 [yara file name:
                    MALW_DragonFly2.yar]
                    src=172.16.0.253
                    smac=00:04:23:e0:04:1c
                    spt=1148
                    proto=TCP
                    start=1571351543431
```

Note the **highlighted** part of the Alert message. This is the Alert Type *ID*. This should be used as the key for performing searches once Nozomi Networks syslog events have been ingested into the integration platform.

Best practice

Make sure that your parsing logic extracts the appropriate data. If you are integrating with *CEF* messages, a *CEF* parser must be used. Do not use regular expressions. This will ensure the integration integrity in the future. When using the correct parser for the data that is expected, be sure to test different inputs to ensure that data is correctly extracted from the messages.

Health events

A description of health events in common event format (CEF).

Health events

Health events in CEF have the following format, as shown in this example:

```
<131>Oct 10 2019 15:57:48 local-sg-19.x n2osevents[0]: CEF:0|Nozomi
                    Networks|N2OS|19.0.3-10201846_FD825|HEALTH|Health
 problem|0|
                    dvchost=local-sg-19.x
                    cs6=1
                    cs6Label=n2os_schema
                    msg=LINK_DOWN_on_port_em0
```

Note the **highlighted** part of the health message. This is the health type *ID*. This should be used as the key for performing searches once Nozomi Networks syslog events have been ingested into the integration platform.

Best practice

Make sure that your parsing logic extracts the appropriate data. If you are integrating with *CEF* messages, a *CEF* parser must be used. Do not use regular expressions. This will ensure the integration integrity in the future. When using the correct parser for the data that is expected, be sure to test different inputs to ensure that data is correctly extracted from the messages.

Audit events

A description of audit events in common event format (CEF).

Audit events in *CEF* are in this format:

```
<134>Oct 10 2019 16:00:18 local-sg-19.x n2osevents[0]: CEF:0|Nozomi
Networks|N2OS|19.0.3-10201846_FD825|AUDIT:SESSIONS:CREATE|User signed
 in|0|
dvchost=local-sg-19.x
cs1=Mozilla/5.0 (Macintosh; Intel Mac OS X 10.15; rv:69.0) Gecko/20100101
Firefox/69.0
cs6=1
cs1Label=browser
cs6Label=n2os_schema
msg=User signed in
src=172.16.248.1
suser=admin
start=1570723218425
```

Note the **highlighted** part of the audit message. This is the Audit Type ID. This should be used as the key for performing searches once Nozomi Networks syslog events have been ingested into the integration platform.

Best practice

Make sure that your parsing logic extracts the appropriate data. If you are integrating with *CEF* messages, a *CEF* parser must be used. Do not use regular expressions. This will ensure the integration integrity in the future. When using the correct parser for the data that is expected, be sure to test different inputs to ensure that data is correctly extracted from the messages.

Connectivity status and status

A description of Connectivity status and status in common event format (CEF).

Connectivity status and status

Connectivity status represents the connectivity status with the data integration endpoint, which is updated when a new connection is initiated. The value is **OK** if the data integration preliminary connection check returns a success response, otherwise you will see the value of the error / exception retrieved while performing these preliminary checks.

Status represents the state of the last data operation. For example, was the data integration able to send data, or did it receive errors. The value is **OK** if the integration returns a success response, a 200 status code, or similar. Otherwise the system displays the value of the error / exception retrieved while trying to send the data.

Both **Connectivity Status** and **Status** support the **OK** value across all data integrations. If either is not OK, one or more errors have occurred. The value then displays error details, depending on the specific data integration.

Webhook

The Webhook integration pushes Nozomi Networks data, such as nodes, asset information, asset CVEs, or alerts, to any HTTP or HTTPS endpoint that accepts POST requests.

The Webhook integration is the most flexible outbound integration and can connect to custom platforms, middleware, or serverless functions.

Limitations

- Only one entity type can be selected per integration instance; create multiple instances for multiple entity types
- Authentication supports Bearer token or Basic authentication
- Batch size is 50 records per request
- The integration receives no data from the endpoint

Data flow

Attribute	Value
Transport	*HTTPS* or *HTTP* POST
Authentication	Bearer token or Basic authentication
Format	*JSON*, *CSV*, or *newline delimited JSON (NDJSON)*
Batch size	50 records per *HTTP* POST
Retry	Cursor is not advanced on failure; records are retried on the next run

Entities exchanged

Entity	Definition	Trigger	Create versus update
Alert	A security or operational alert.	Incremental by `record_created_at`	New records only; by `record_updated_at` if `send_updates` is enabled.
Node	A network node, such as an *OT* or *information technology (IT)* device.	Incremental by `time`	New records only; daily deduplication if `send_updates` is enabled.

Entity	Definition	Trigger	Create versus update
Asset information	Asset information discovered by Nozomi Networks.	Incremental by `time`	New records only; daily deduplication if `send_updates` is enabled.
Asset CVE	A CVE finding associated with an asset.	Incremental by `time`	New records only.

Entity attribute mapping

Fields sent are the raw appliance query result columns for the selected entity. No field renaming is applied. The *JSON* keys match the appliance column names.

Common alert fields when entity is `alerts`:

Nozomi field	JSON key	Type	Example
`id`	`id`	string	`550e8400-...`
`time`	`time`	integer	`1729409051000`
`type_id`	`type_id`	string	`VI:NEW-FUNC-CODE`
`name`	`name`	string	`New Function Code`
`description`	`description`	string	`New SCADA function code detected`
`risk`	`risk`	float	`3.5`
`ip_src`	`ip_src`	string	`192.168.1.1`
`ip_dst`	`ip_dst`	string	`192.168.1.2`
`is_security`	`is_security`	boolean	`true`
`appliance_host`	`appliance_host`	string	`guardian-01`

Common node fields when entity is `nodes`:

Nozomi field	JSON key	Type	Example
`id`	`id`	string	`<uuid>`
`ip`	`ip`	array	`["192.168.1.1"]`
`mac_address`	`mac_address`	string	`00:09:0f:6d:f1:c1`

Nozomi field	JSON key	Type	Example
label	label	string	plc-01
vendor	vendor	string	Siemens
os	os	string	SIMATIC S7

Examples

Outbound payload in JSON format for alerts:

```
[
  {
    "id": "550e8400-e29b-41d4-a716-446655440000",
    "time": 1729409051000,
    "type_id": "VI:NEW-FUNC-CODE",
    "name": "New Function Code",
    "description": "New SCADA function code detected on PLC",
    "risk": 3.5,
    "ip_src": "192.168.1.1",
    "ip_dst": "192.168.1.2",
    "is_security": false,
    "appliance_host": "guardian-01"
  }
]
```

Outbound payload in CSV format for nodes:

```
id,ip,mac_address,label,vendor
550e8400-...,192.168.1.1,00:09:0f:6d:f1:c1,plc-01,Siemens
```

Outbound payload in NDJSON format for nodes:

```
{"id":"550e8400-...","ip":
["192.168.1.1"],"mac_address":"00:09:0f:6d:f1:c1","label":"plc-01","vendo
r":"Siemens"}
{"id":"550e8401-...","ip":
["192.168.1.2"],"mac_address":"00:09:0f:6d:f1:c2","label":"hmi-01","vendo
r":"Siemens"}
```

Verification

- Check the receiving endpoint's access log for POST requests from the appliance *IP* address
- In the *UI*, check the run history for `webhook` and verify the sent counts.
- Use a service such as `webhook.site` to capture and inspect raw payloads during testing.

Configure the Webhook integration

Configure a Webhook endpoint to push Nozomi Networks data to an HTTP or HTTPS endpoint.

Before you begin

Make sure that you have:

- The target webhook *URL*
- Authentication credentials for the endpoint (Bearer token, or username and password for Basic authentication)

About this task

The following parameters are required.

Parameter	Description
To URI	The webhook *URL*, for example `https://myapp.example.com/nozomi-hook`.
Entity	One of: `alerts`, `nodes`, `assets`, `asset_cves`.

Select one of the following authentication methods.

Authentication method	Description
Bearer token	Sends the token in the `Authorization: Bearer <token>` header.
Basic authentication	Uses a username and password for Basic authentication.

The following parameters are optional.

Parameter	Description	Default
`description`	A description for the endpoint configuration.	None
`format`	Request body format: `json`, `csv`, or `ndjson`.	`json`
`query_filter`	An extra query filter to limit the records sent by the integration.	None
`send_updates`	Re-sends updated records when supported by the selected entity type.	`false`
`page_count`	Number of records per chunk.	300

Procedure

1. In the top navigation bar, select 

   Result: The administration page opens.

2. In the **Settings** section, select **Data integration**.

   Result: The **Data integration** page opens.

3. Select **Add**.

   Result: A dialog shows.

4. From the **Choose a configuration** dropdown, select **Webhook**.
5. Set **Endpoint Configured as** to **Webhook**.
6. Enter a value in **Description**.
7. Enter the target *URL* in **To URI**.
8. Select the authentication method and enter the required credentials:

Method	Action
Bearer token	Enter the token value.
Basic authentication	Enter a username and password.

9. Select the output **Format**:
   - **JSON**
   - **CSV**
   - **NDJSON**
10. Select one **Entity** type:
    - **Enable sending Alerts**
    - **Enable sending Nodes**
    - **Enable sending Asset information**
    - **Enable sending Asset CVEs**

    Note:
    Only one entity type can be enabled per Webhook endpoint. Create a separate endpoint for each additional entity type.

11. **Optional:** Select **Send updates** to re-send updated records when the selected entity type supports it.
12. **Optional:** Enter a **Query filter** to limit the records sent.
13. Select **Edit Endpoint** to save the configuration.

Results

The Webhook endpoint is saved and begins pushing data to the configured *URL* on the next run.

Troubleshooting

The Webhook integration can fail because of authentication issues, permission restrictions, payload size limits, query or data availability mismatches, or format incompatibilities between the appliance and the receiving endpoint. See the following topics for symptom-specific causes and remedies.

HTTP 401

Diagnose and resolve HTTP 401 errors returned by the Webhook integration.

Condition

The Webhook integration returns an *HTTP* 401 error when sending data to the endpoint.

Possible cause

The configured authentication credentials are invalid, expired, or not accepted by the endpoint.

Procedure

1. Verify that the selected authentication method matches the endpoint requirements and re-enter valid credentials.
2. Send a test request to the endpoint with the same credentials to confirm they authenticate successfully.

HTTP 403

Diagnose and resolve HTTP 403 errors returned by the Webhook integration.

Condition

The Webhook integration returns an *HTTP* 403 error when sending data to the endpoint.

Possible cause

The credentials are valid but do not have permission to create or update resources on the endpoint.

Procedure

1. Confirm that the service account has write permissions for the target endpoint path.
2. Update endpoint-side access control rules to allow requests from the appliance.

HTTP 413

Diagnose and resolve HTTP 413 errors returned by the Webhook integration.

Condition

The Webhook integration returns an *HTTP* 413 error when sending data to the endpoint.

Possible cause

The request payload exceeds the maximum size accepted by the receiving endpoint.

Procedure

1. Reduce the `page_count` value in the endpoint configuration to decrease the payload size.
2. If needed, increase the endpoint's allowed request size limit.

No data received

Diagnose and resolve cases where the Webhook integration sends no data.

Condition

The receiving endpoint does not receive any data from the Webhook integration.

Possible cause

The selected entity has no matching records, or the configured query filter excludes all records.

Procedure

1. Confirm that the selected entity type has records available in the *UI*.
2. Temporarily clear the query filter and rerun the integration to verify data flow.

Data is not in the expected format

Diagnose and resolve format mismatches between the Webhook integration and the receiving endpoint.

Condition

The receiving endpoint cannot parse the data sent by the Webhook integration.

Possible cause

The configured output format does not match what the receiving endpoint expects.

Procedure

1. Check the **Format** setting in the endpoint configuration and confirm the receiving endpoint supports that format.
2. Set the format expected by the receiver, then send a new test batch and confirm successful parsing.

Configure data integration

*The **Data integration** page lets you configure one of the available options.*

Procedure

1. In the top navigation bar, select 

   Result: The administration page opens.

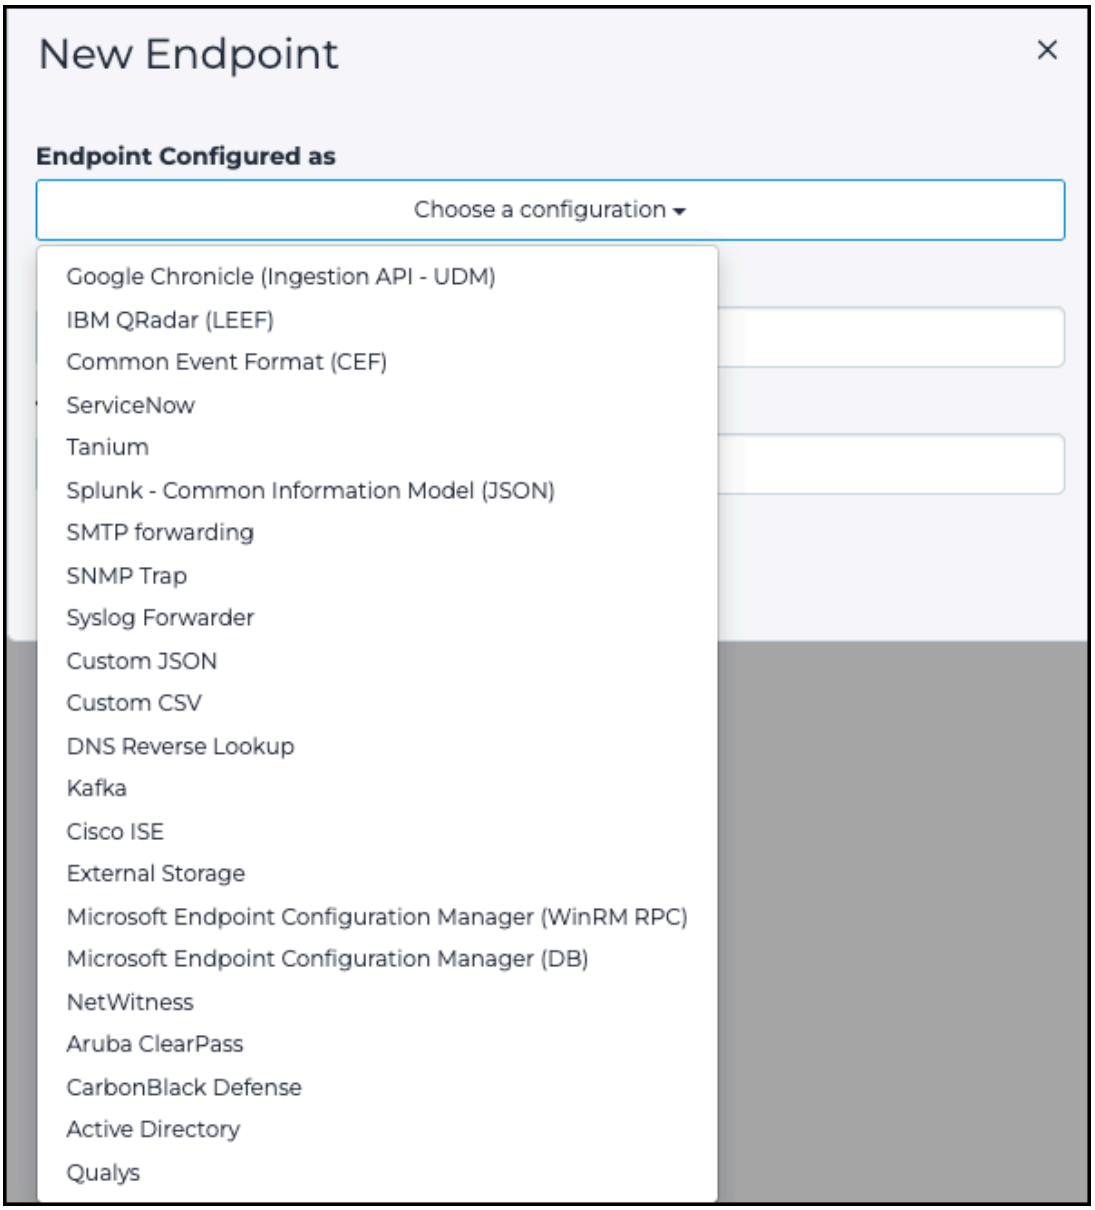
2. In the **Settings** section, select **Data integration**.

   Result: The **Data integration** page opens.

3. In the top right section, select **Add**.

   Result: A dialog shows.

4. From the **Choose a configuration** dropdown, select an option.



5. Enter the details as necessary for the option that you chose.

Firewall integration

*The **Firewall integrations** page shows all the firewall integrations and lets you add new ones.*

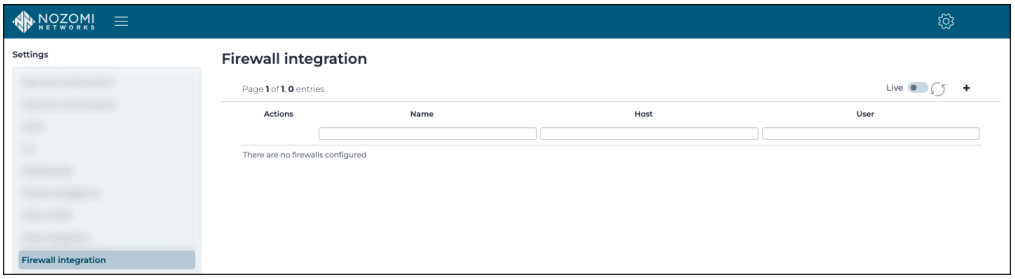


Figure 51. Firewall integrations page

General

Guardian lets you configure firewall integrations. Guardian discovers, identifies, and learns the behavior of assets on your network. Through integration with the firewall, unlearned nodes and links are automatically blocked through block policies. Block policies are not created for nodes and links in the learned state.

Note:

For some firewall integrations, Guardian supports session kill.

After the integration has been set up, policies are produced and inserted in the firewall. The policies are displayed in the **Policies** section.

Firewall rules strategies

For supported integrations, you can choose how Guardian decides when to create firewall rules. Two strategies are available:

- **Block active alerts**: Guardian creates link-blocking policies when specific alerts are triggered. Rules are generated based on real-time detection events, not on node learning state. For more details, see Alert-based firewall rule creation (on page 136).
- **Block unlearned**: Guardian creates block policies for nodes and links that have not yet been learned. This is the default behavior.

Features

Firewall integrations only work when, in the Security control panel (on page 13), the:

- **Detection approach** is set to **Strict**
- **Phase switching** is set to **Protecting**

It does not work when the policy for zones is set to override the **Protecting** and **Strict** mode. In this mode, we can see new nodes, but they are not learned.
If the global learning policy is set to **Adaptive Learning** and **Learning**, and a zone is set to **Adaptive Learning** and **Protecting**, we see new nodes, but they are not learned, however links to new nodes are learned automatically.

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.

Add

The **+** icon lets you add a new firewall integration.

Alert-based firewall rule creation

Guardian can create firewall rules dynamically in response to specific alerts, providing event-driven network enforcement independent of node learning state.

Overview

The **Block active alerts** firewall rules strategy lets Guardian create link-blocking policies when specific alerts are triggered. Unlike the default **Block unlearned** strategy, which acts on node learning state, alert-based rule creation responds to real-time detection events. Guardian creates only link-blocking policies; node-based blocking does not apply in this mode.

Use this strategy when you want enforcement to be driven by detected threats rather than by whether assets have been learned.

Supported alert types

You select which alert types trigger firewall rule creation. The following types are available:

- `SIGN:MITM`
- `SIGN:SCADA-INJECTION`
- `SIGN:PACKET-RULE`
- `SIGN:OUTBOUND-CONNECTIONS`
- `SIGN:MALWARE-DETECTED`
- `SIGN:MULTIPLE-ACCESS-DENIED`
- `SIGN:MULTIPLE-UNSUCCESSFUL-LOGINS`
- `SIGN:NETWORK-SCAN`

Only the selected alert types generate firewall rules.

Severity and risk thresholds for packet rules

For `SIGN:PACKET-RULE` alerts, Guardian applies additional filtering before creating a firewall rule. Guardian creates a rule only when both of the following conditions are met:

- The alert severity is greater than or equal to the configured **Severity threshold**.
- The alert risk score is greater than or equal to the configured **Risk threshold**.

Both thresholds are integers. This lets you avoid generating rules for low-impact events and fine-tune automation based on your risk tolerance.



Note:

Both conditions must be met simultaneously. If either threshold is not reached, no firewall rule is created for that alert.

Related information

Configure Fortinet FortiGate (on page 144)

Configure Palo Alto Networks v10.1 (on page 148)

Configure Barracuda

Configure Guardian firewall integration with the Barracuda firewall.

Before you begin

Make sure that you have administrator privileges.

About this task

Guardian integration supports Barracuda *API* v8.3.

Procedure

1. In the top navigation bar, select ⚙

   Result: The administration page opens.

2. In the **Settings** section, select **Firewall integration**.

   Result: The **Firewall integration** page opens.

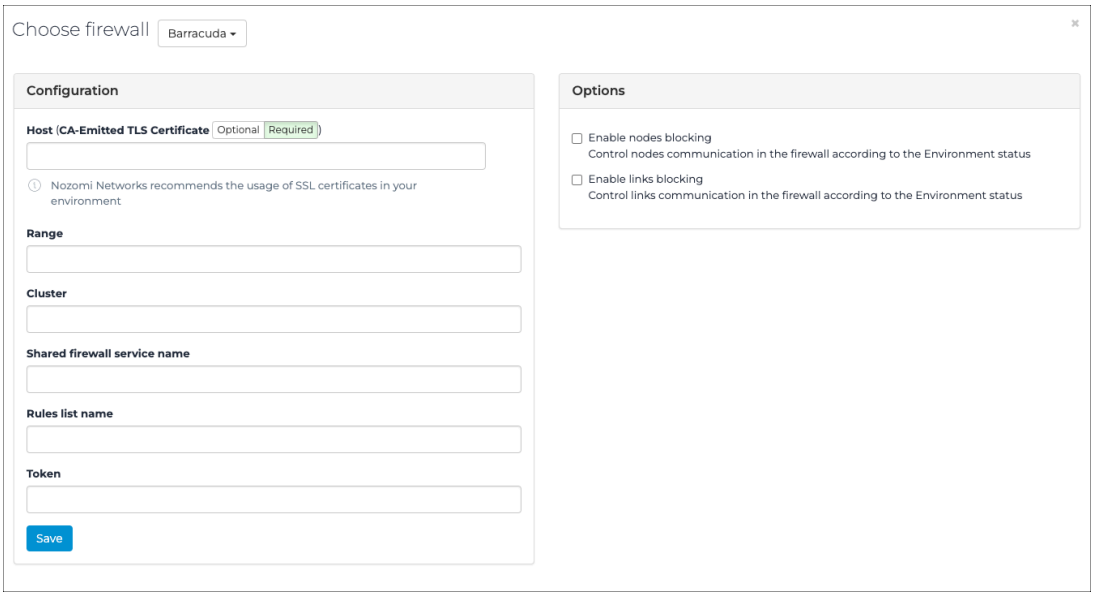
3. In the top right section, select **+**

   Result: A dialog shows.

4. From the **Choose firewall** dropdown, select **Barracuda**.

   Result: A dialog shows.

5. If it is not populated already, in the **Host (CA-Emitted TLS Certificate)** field, enter the host *IP* address.



6. In the **Range** field, enter the *IP* address range.
7. In the **Cluster** field, enter the cluster.
8. In the **Shared firewall service name** field, enter a name.
9. In the **Rules list name** field, enter a name.

10. In the **Token** field, enter the token.
11. **Optional:** If necessary, tune the integration's behavior in the **Options** section.
    a. If necessary, select **Enable nodes blocking**.
    b. If necessary, select **Enable links blocking**.
12. Select **Save**.

Results

The firewall integration has been configured.

Configure Cisco ASA

Configure Guardian firewall integration with the Cisco ASA firewall.

Before you begin

Make sure that you have administrator privileges.

Procedure

1. In the top navigation bar, select ⚙

   Result: The administration page opens.

2. In the **Settings** section, select **Firewall integration**.

   Result: The **Firewall integration** page opens.

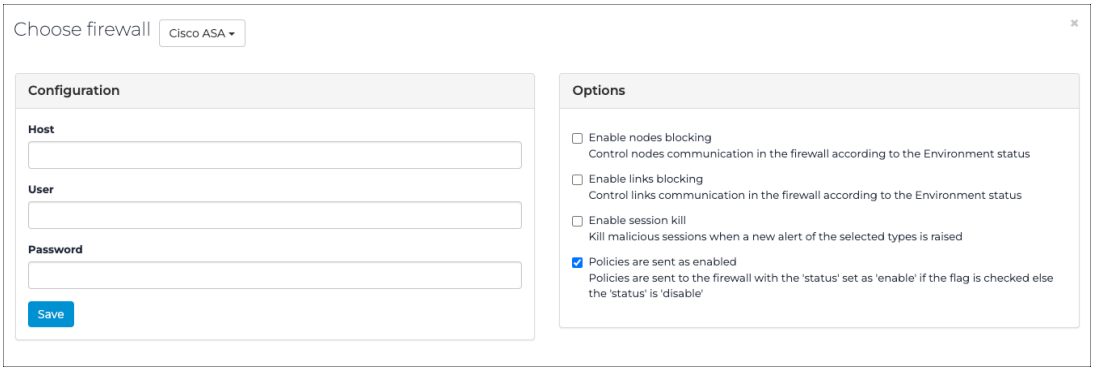
3. In the top right section, select **+**

   Result: A dialog shows.

4. From the **Choose firewall** dropdown, select **Cisco ASA**.

   Result: A dialog shows.

5. If it is not populated already, in the **Host** field, enter the host *IP* address.



6. In the **User** field, enter your user name.
7. In the **Password** field, enter your password.
8. **Optional:** If necessary, tune the integration's behavior in the **Options** section.
   a. If necessary, select **Enable nodes blocking**.
   b. If necessary, select **Enable links blocking**.
   c. If necessary, select **Enable session kill**. Then select the specific alert type(s).
   d. If necessary, select **Policies are sent as enabled**.
9. Select **Save**.

Results

The firewall integration has been configured.

Configure Cisco FTD

Configure Guardian firewall integration with the Cisco FTD firewall.

Before you begin

Make sure that you have administrator privileges.

Procedure

1. In the top navigation bar, select ⚙

   Result: The administration page opens.

2. In the **Settings** section, select **Firewall integration**.

   Result: The **Firewall integration** page opens.

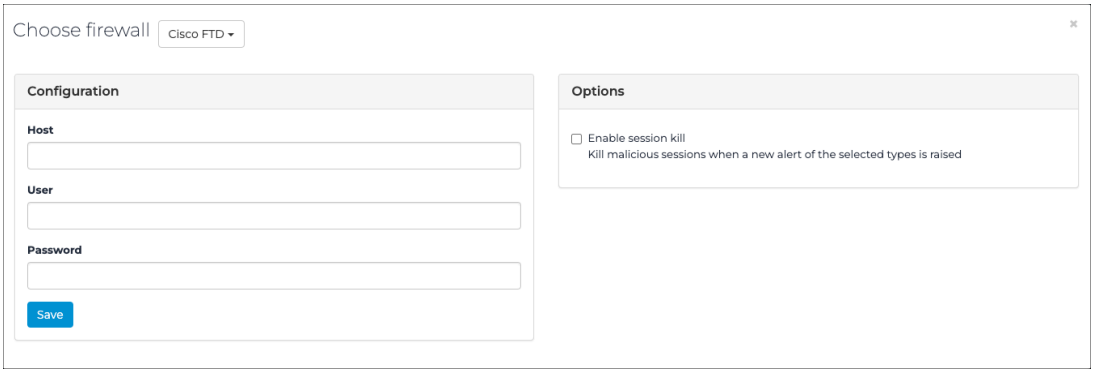
3. In the top right section, select **+**

   Result: A dialog shows.

4. From the **Choose firewall** dropdown, select **Cisco FTD**.

   Result: A dialog shows.

5. If it is not populated already, in the **Host** field, enter the host *IP* address.



6. In the **User** field, enter your user name.
7. In the **Password** field, enter your password.
8. **Optional:** If necessary, in the **Options** section, select **Enable session kill**. Then select the specific alert type(s).in the **Options** section.
9. Select **Save**.

Results

The firewall integration has been configured.

Configure Cisco ISE

Configure Guardian firewall integration with the Cisco ISE firewall.

Before you begin

Make sure that you have administrator privileges.

About this task

The integration between Guardian and Cisco ISE lets Cisco customers to extend network access controls and policy enforcement to their *OT* and *IoT* networks from the Cisco ISE. Guardian uses the pxGrid platform to integrate with Cisco ISE. Along with the client associated with the certificate and the certificate password, you need to upload the identity certificate and the private key. The preferred method of authenticating with the Cisco ISE is to use certificates. Guardian supports authentication with certificates issued by:

- The Cisco ISE internal *CA*
- An external *CA* (third-party certificates)

Procedure

1. In the top navigation bar, select 

   Result: The administration page opens.

2. In the **Settings** section, select **Firewall integration**.

   Result: The **Firewall integration** page opens.

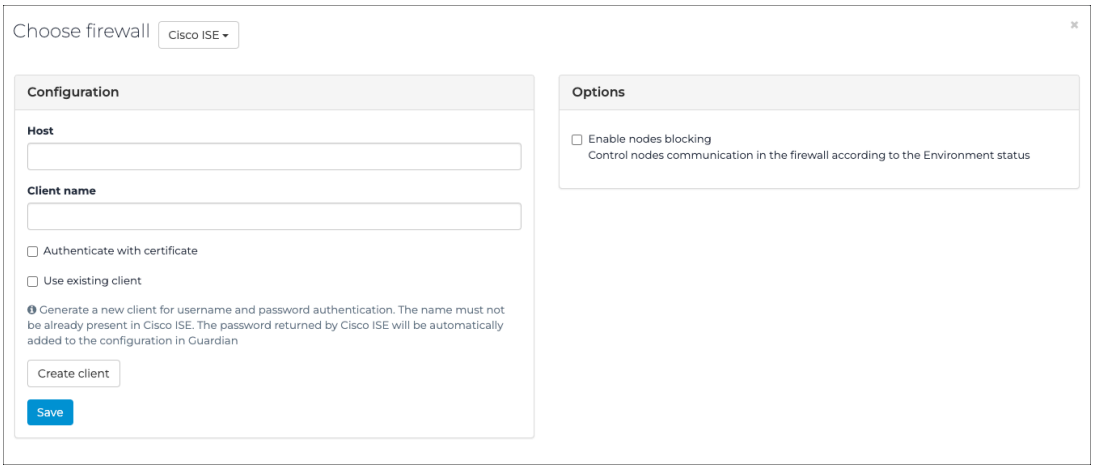
3. In the top right section, select **+**

   Result: A dialog shows.

4. From the **Choose firewall** dropdown, select **Cisco ISE**.

   Result: A dialog shows.

5. If it is not populated already, in the **Host** field, enter the host *IP* address.



> **Note:**
> The host *IP* address is the *IP* address of the Cisco ISE firewall that you are configuring.

6. In the **Client name** field, enter the name of the client.

> **Note:**
> The client name is taken from the Cisco ISE pxGrid Services screen on the Cisco ISE Web *UI*. For more details, see the appropriate Cisco ISE documentation.

7. **Optional:** Authenticate with a Cisco ISE internal *CA* certificate. select **Authenticate with certificate**., then enter the password in the Password field.
    a. Select **Authenticate with certificate**.
    b. In the **Password** field, enter your password.
8. **Optional:** Use a third-party certificate. check the Use third party certificate box, then import the certificate(s), using one of the following methods
    a. Select **Use third party certificate**.
    b. Choose a method to import the certificate:

    Choose from:
    - **Import the CA certificate**
    - **Import the certificate**
    - **Import the key**
9. If you chose, **Import the CA certificate** or **Import the certificate**, continue from step 10 (on page 143) . If you chose **Import the key**, continue from step 11 (on page 143).

10. Import the certificate.
    a. Select **Import the certificate**.
       Result: a dialog shows.
    b. Select the file and import it.
11. Import the key.
    a. Select **Import the key**.
    b. Select the file and import it.
12. **Optional:** If necessary, in the **Options** section, select **Enable nodes blocking**.
13. Select **Save**.

Results

The firewall integration has been configured.

Configure Fortinet FortiGate

Configure Guardian firewall integration with the Fortinet FortiGate firewall.

Before you begin

Make sure that:

- You have administrator privileges
- You have generated the REST *API* access token from the firewall admin Web *UI*
- You have added the Guardian address subnet to trusted hosts

> **Note:**
> The access token needs to have permission to insert, read, and delete entities such as:
>
> - Addresses
> - Addrgroups
> - Routes
> - Sessions
> - Policies

About this task

Guardian integration supports FortiOS versions 6.2, 6.4, 7.0, 7.2, and 7.4.9. This integration uses the REST *API*.

Procedure

1. In the top navigation bar, select ⚙

   Result: The administration page opens.

2. In the **Settings** section, select **Firewall integration**.

   Result: The **Firewall integration** page opens.

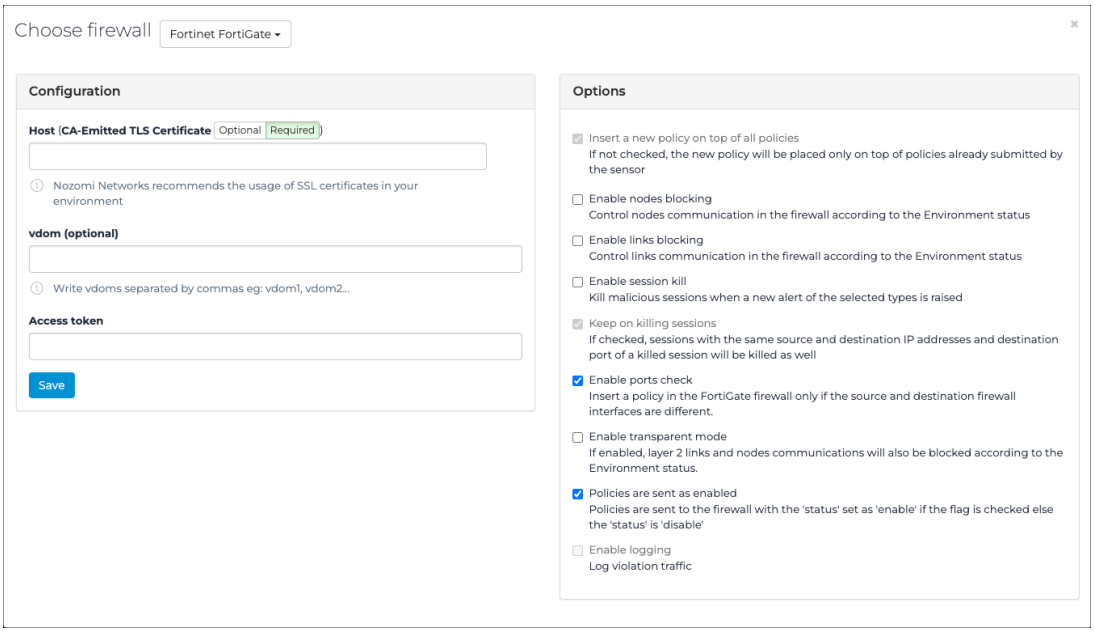
3. In the top right section, select **+**

   Result: A dialog shows.

4. From the **Choose firewall** dropdown, select **Fortinet FortiGate**.

   Result: A dialog shows.

5. If it is not populated already, in the **Host (CA-Emitted TLS Certificate)** field, enter the host *IP* address.



6. **Optional:** In the **vdom (optional)** field, enter one or more *Virtual DOM (vdom)*s. Use a comma to separate multiple entries.
7. In the **Access token** field, enter the access token.

8. **Optional:** If necessary, in the **Options** section, select one or more of these options:
   a. Select **Insert a new policy on top of all policies**.
   b. For **Firewall rules strategy**, select a strategy.

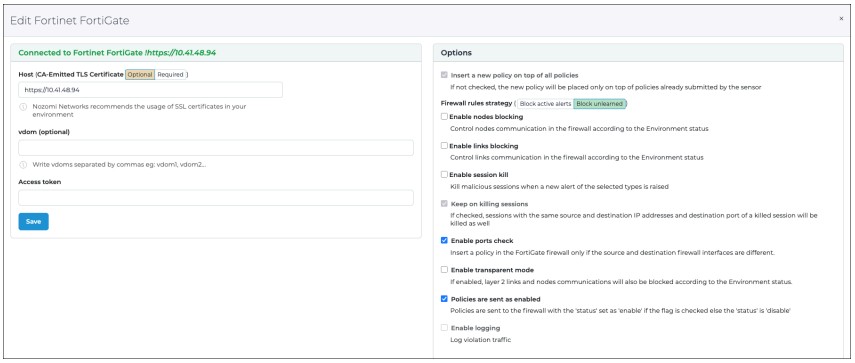


Figure 52. Firewall rules strategy toggle

- **Block active alerts**: Guardian creates link-blocking policies when specific alerts are triggered. Guardian generates rules based on real-time detection events, not on node learning state. For more details, see Alert-based firewall rule creation (on page 136).
- **Block unlearned**: Guardian creates block policies for unlearned nodes and links. This is the default behavior.

If you selected **Block active alerts**, select one or more alert types that should trigger firewall rule creation:

- `SIGN:MITM`
- `SIGN:SCADA-INJECTION`
- `SIGN:PACKET-RULE`
- `SIGN:OUTBOUND-CONNECTIONS`
- `SIGN:MALWARE-DETECTED`
- `SIGN:MULTIPLE-ACCESS-DENIED`
- `SIGN:MULTIPLE-UNSUCCESSFUL-LOGINS`
- `SIGN:NETWORK-SCAN`

If `SIGN:PACKET-RULE` is selected, enter a **Severity threshold** value and a **Risk threshold** value. Guardian creates a firewall rule for `SIGN:PACKET-RULE` alerts only when both the alert severity and the risk score meet or exceed the configured thresholds.

   c. Select **Enable nodes blocking**.
   d. Select **Enable links blocking**.
   e. Select **Enable session kill**. Then select the specific alert type(s).
   f. Select **Keep on selecting sessions**.
   g. Select **Enable ports check**.
   h. Select **Enable transparent mode**.
   i. Select **Policies are sent as enabled**.

9. Select **Save**.

Results

The firewall integration has been configured.

Configure Palo Alto Networks v10.1

Configure Guardian firewall integration with the Palo Alto Networks v10.1 firewall.

Before you begin

Make sure that you have administrator privileges.

About this task

Starting with version 10.0, PAN-OS provides a REST *API*. The Guardian integration that relies on this new *API* supports the same features as the previous Palo Alto integration, plus these:

- **Commit by user**: Commits the current changes required by the user, which are represented by the credentials used for the *API*. Global commits are no longer performed
- **Dynamic Access Groups for Node Blocking**: Dynamic Access Group references a tag, which is then assigned to a new *IP* address for objects that are created on the firewall. This will automatically apply the global Guardian denylist rule to each new address without modifying the firewall ruleset

> **Note:**
> This firewall integration supports IPv6 addresses.

Procedure

1. In the top navigation bar, select 

   Result: The administration page opens.

2. In the **Settings** section, select **Firewall integration**.

   Result: The **Firewall integration** page opens.

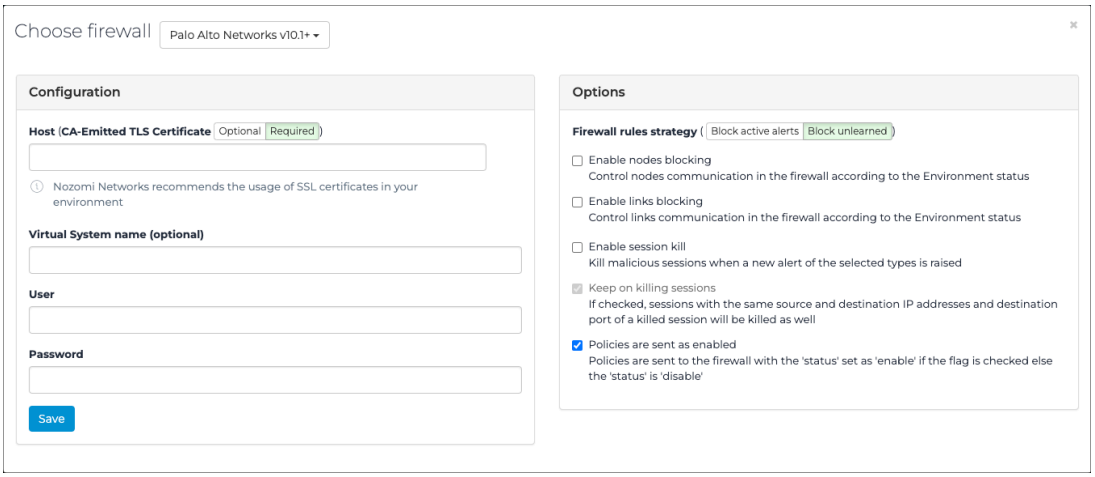
3. In the top right section, select **+**

   Result: A dialog shows.

4. From the **Choose firewall** dropdown, select **Palo Alto Networks v10.1+**.

   Result: A dialog shows.

5. If it is not populated already, in the **Host (CA-Emitted TLS Certificate)** field, enter the host *IP* address.



6. **Optional:** In the **Virtual System name (optional)** field, enter a name.
7. In the **User** field, enter your user name.
8. In the **Password** field, enter your password.

9. **Optional:** If necessary, in the **Options** section, select one or more of these options:
    a. For **Firewall rules strategy**, select a strategy.
        - **Block active alerts**: Guardian creates link-blocking policies when specific alerts are triggered. Guardian generates rules based on real-time detection events, not on node learning state. For more details, see Alert-based firewall rule creation (on page 136).
        - **Block unlearned**: Guardian creates block policies for unlearned nodes and links. This is the default behavior.

    If you selected **Block active alerts**, select one or more alert types that should trigger firewall rule creation:

        - `SIGN:MITM`
        - `SIGN:SCADA-INJECTION`
        - `SIGN:PACKET-RULE`
        - `SIGN:OUTBOUND-CONNECTIONS`
        - `SIGN:MALWARE-DETECTED`
        - `SIGN:MULTIPLE-ACCESS-DENIED`
        - `SIGN:MULTIPLE-UNSUCCESSFUL-LOGINS`
        - `SIGN:NETWORK-SCAN`

    If `SIGN:PACKET-RULE` is selected, enter a **Severity threshold** value and a **Risk threshold** value. Guardian creates a firewall rule for `SIGN:PACKET-RULE` alerts only when both the alert severity and the risk score meet or exceed the configured thresholds.

    b. Select **Enable nodes blocking**.

    c. Select **Enable links blocking**.

    d. Select **Enable session kill**. Then select the specific alert type(s).

    e. Select **Keep on selecting sessions**.

    f. Select **Policies are sent as enabled**.

10. Select **Save**.

Results

The firewall integration has been configured.

Configure Stormshield SNS

Configure Guardian firewall integration with the Stormshield SNS firewall.

Before you begin

Make sure that you have administrator privileges.

Procedure

1. In the top navigation bar, select ⚙

   Result: The administration page opens.

2. In the **Settings** section, select **Firewall integration**.

   Result: The **Firewall integration** page opens.

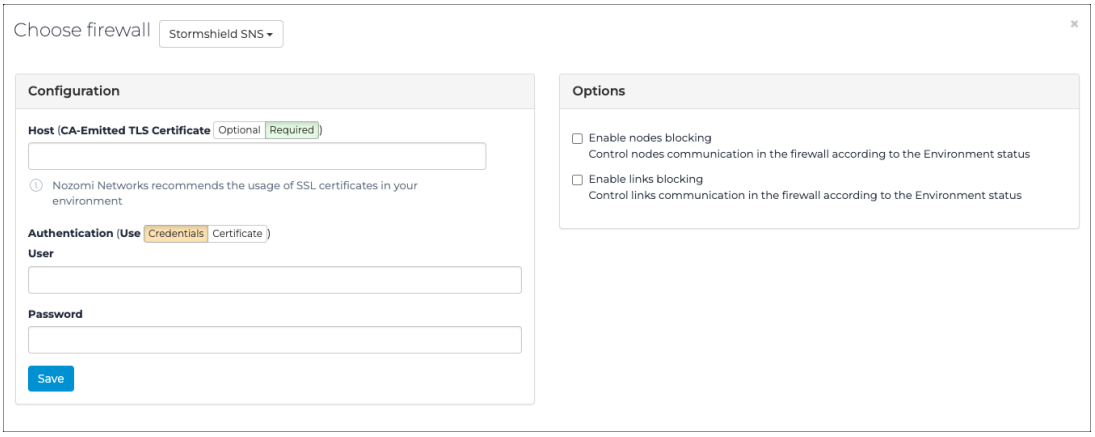
3. In the top right section, select **+**

   Result: A dialog shows.

4. From the **Choose firewall** dropdown, select **Stormshield SNS**.

   Result: A dialog shows.

5. If it is not populated already, in the **Host (CA-Emitted TLS Certificate)** field, enter the host *IP* address.



6. In the **Authentication** section, choose an option:

   Choose from:

   - **Credentials**
   - **Certificate**

7. If you chose **Credentials**, continue from do step 7.a (on page 151). If you chose **Certificate**, continue from step 8 (on page 152).

   a. In the **User** field, enter your user name.
   b. In the **Password** field, enter your password.
   c. Go to step 12 (on page 152).

8. If you chose **Certificate**, choose an option:

   Choose from:
   - **Import the certificate**
   - **Import the key**

9. If you chose, **Import the certificate**, continue from step 10 (on page 152). If you chose **Import the key**, continue from step 11 (on page 152).
10. Import the certificate.
    a. Select **Import the certificate**.
       Result: a dialog shows.
    b. Select the file and import it.
11. Import the key.
    a. Select **Import the key**.
    b. Select the file and import it.
12. **Optional:** If necessary, tune the integration's behavior in the **Options** section.
    a. If necessary, select **Enable nodes blocking**.
    b. If necessary, select **Enable links blocking**.
13. Select **Save**.

Results

The firewall integration has been configured.

Configure TXone EdgeOne / OT Defense Console

Configure Guardian firewall integration with the TXone EdgeOne / OT Defense Console firewall.

Before you begin

Make sure that you have administrator privileges.

About this task

TXOne EdgeOne / OT Defense Console (ODC) provides a REST *API* v1.1. The Guardian integration relying on this *API* supports the same features as previous integrations from Trend Micro.

Procedure

1. In the top navigation bar, select 

   Result: The administration page opens.

2. In the **Settings** section, select **Firewall integration**.

   Result: The **Firewall integration** page opens.

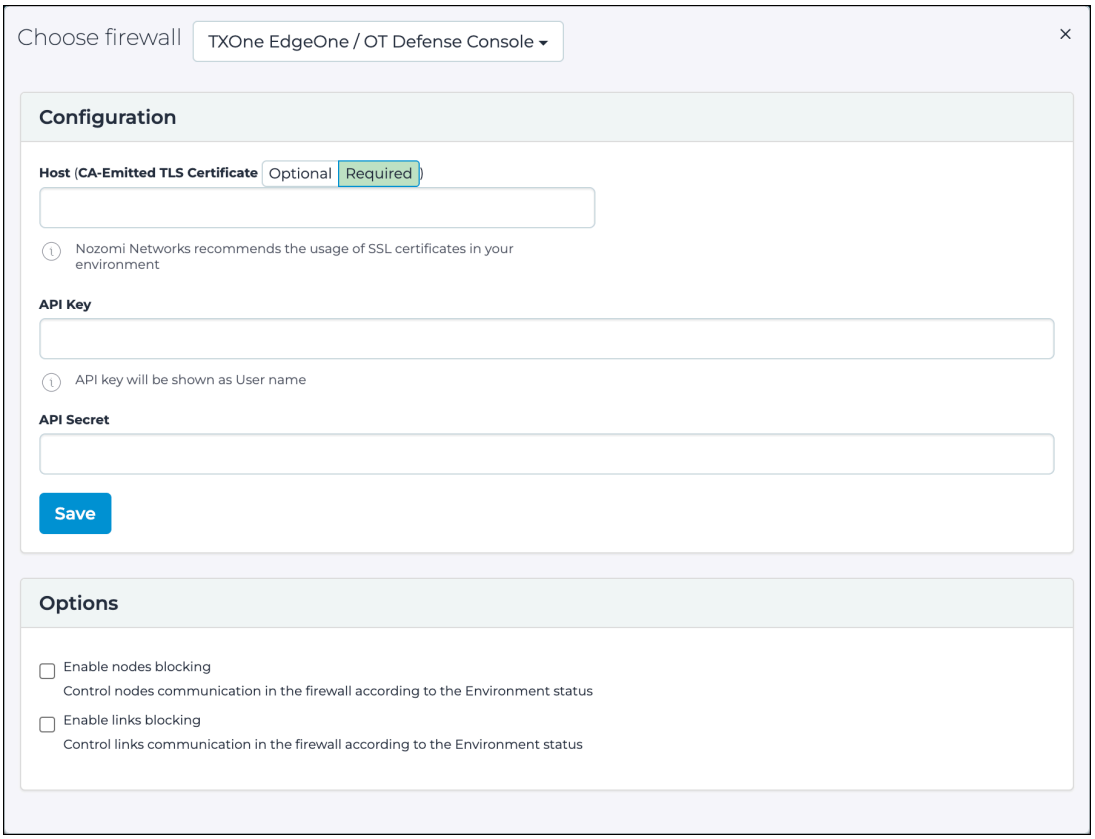
3. In the top right section, select **+**

   Result: A dialog shows.

4. From the **Choose firewall** dropdown, select **TXone EdgeOne / OT Defense Console**.

   Result: A dialog shows.

5. If it is not populated already, in the **Host (CA-Emitted TLS Certificate)** field, enter the host *IP* address.



6. In the **API Key** field, enter the *API* key.
7. In the **API Secret** field, enter the *API* secret.
8. **Optional:** If necessary, tune the integration's behavior in the **Options** section.
    a. If necessary, select **Enable nodes blocking**.
    b. If necessary, select **Enable links blocking**.
9. Select **Save**.

Results

The firewall integration has been configured.

Configure Check Point R81.20

Configure Guardian firewall integration with the Check Point R81.20 firewall.

Before you begin

Make sure that you have:

- Administrator privileges
- A valid username and password for authentication
- Access to the Check Point R81.20 management host
- The name of the gateway where the rules will be installed

About this task

Check Point R81.20 side, this firewall integration:

- Creates a Check Point R81.20 layer (Nozomi layer) containing the rules
- Adds the rules to block links and nodes that are unlearned

> **Note:**
> If a node or link changes its status to learned, the corresponding rule will be removed.

- Creates a session named **Nozomi Guardian** in the firewall to manage rule insertion and removal
- Generates Check Point R81.20 Service objects for link ports that are not already mapped in the firewall
- Creates host objects to be used in the block link and nodes rules

Procedure

1. In the top navigation bar, select ⚙

   Result: The administration page opens.

2. In the **Settings** section, select **Firewall integration**.

   Result: The **Firewall integration** page opens.

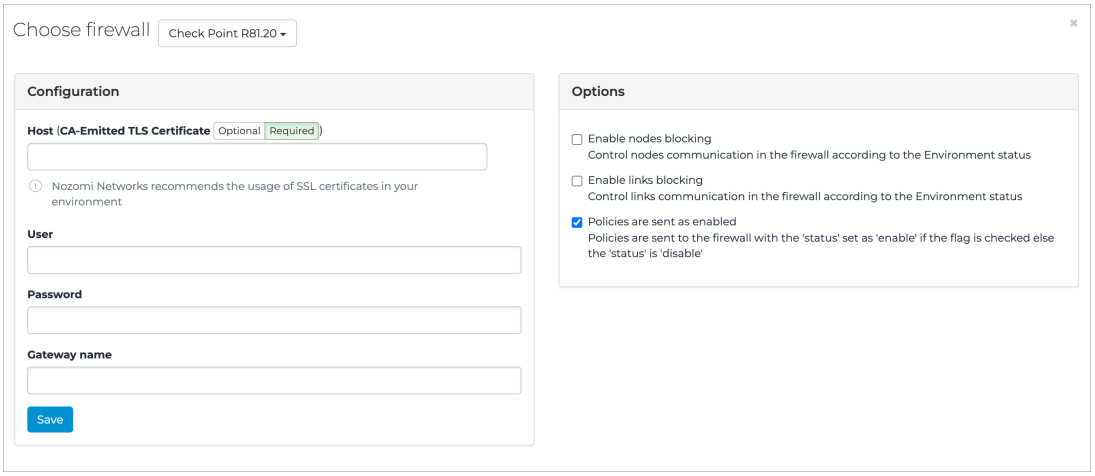
3. In the top right section, select **+**

   Result: A dialog shows.

4. From the **Choose firewall** dropdown, select **Check Point R81.20**.

   Result: A dialog shows.

5. If it is not populated already, in the **Host (CA-Emitted TLS Certificate)** field, enter the host *IP* address.



6. In the **User** field, enter your user name.
7. In the **Password** field, enter your password.
8. In the **Gateway name** field, enter a name.
9. **Optional:** Select **Enable nodes blocking**.
   Enable this option if you want to block nodes.
10. **Optional:** Select **Enable links blocking**.
    Enable this option if you want to block links.
11. **Optional:** Select **Policies are sent as enabled**.
    Make sure that this option is selected so that policies are active upon deployment.
12. Select **Save**.

Results

The firewall integration has been configured.

Credentials manager

*The **Credentials manager** is a tool that lets you centralize the management of monitored endpoints to make it easier to create, delete, or update the credentials that are needed to access those endpoints.*

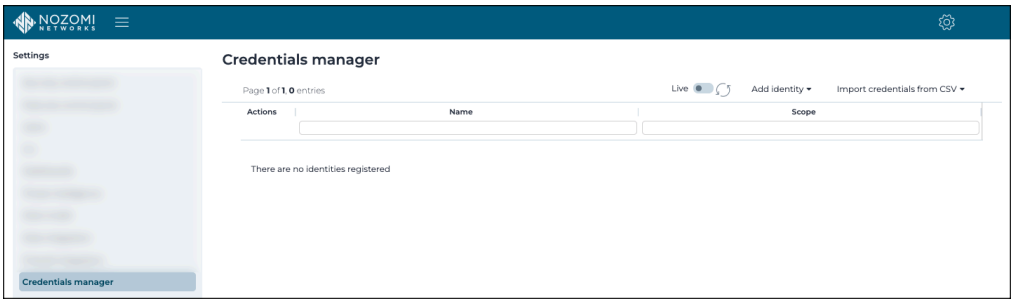


Figure 53. Credentials manager page

Credentials manager is a feature that lets you securely store passwords, and other sensitive information, that Guardian uses to:

- Access hosts through Smart Polling or Arc
- Decrypt encrypted transmissions that are passively detected

Before Credentials manager, Smart Polling managed device credentials. Users migrating their system over to a version with Credentials manager will have their existing device credentials automatically migrated over to Credentials manager as part of the execution of the migration tasks after the upgrade.

Add an identity

*You can use **Credentials manager** to add identities.*

About this task

Each identity has a unique name to distinguish it from other identities. To insert a node into the applicability list of an identity, you can:

- Enter an *IP* address, or a subnet mask, into the applicable field, or
- Select a set of nodes from a list

> **Note:**
> A node cannot belong to multiple identities for the same scope.

Procedure

1. In the top navigation bar, select 

   Result: The administration page opens.

2. In the **Settings** section, select **Credentials manager**.

   Result: The **Credentials manager** page opens.

3. In the top right section, select **Add identity**.

   Result: A dropdown shows.

4. From the dropdown, select the applicable option.

   Result: A dialog shows.

5. Enter the details as necessary.
6. **Optional:** Select **Select nodes from list**, or use **Add node ID / subnet**.

   Result: A dialog shows.

7. **Optional:** Select the nodes desired for the scope of applicability.
8. Select **Apply** to selected nodes.

   Result: The nodes are associated to the credentials in the new identity.

9. Select **Save**.

Results

The identity has been added.

Import credentials from a CSV file

*The **Credentials manager** lets you import credentials from a comma-separated value (CSV) file.*

Procedure

1. In the top navigation bar, select ⚙

   Result: The administration page opens.

2. In the **Settings** section, select **Credentials manager**.

   Result: The **Credentials manager** page opens.

3. In the top right section, select **Import credentials from CSV**.

   Result: A dropdown shows.

4. From the dropdown, select the applicable option.

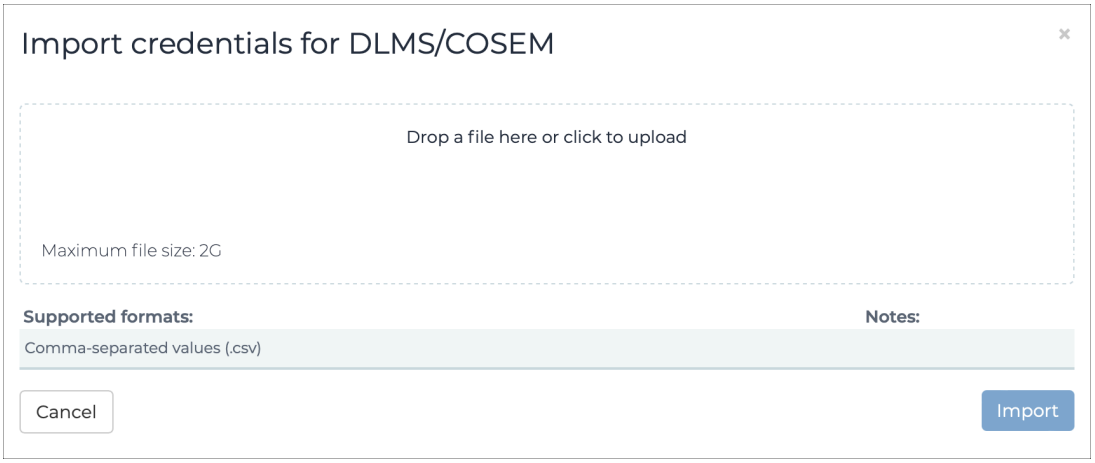
   Result: A dialog shows.

5. Choose a method to upload a file:

   Choose from:

   - Drag your file into the **Drop a file here or click to upload** field
   - Click in the **Drop a file here or click to upload** field

6. If you chose the second method, select the correct file to upload.



> **Note:**
> The only supported format is *CSV*.

7. Wait for the file to upload.
8. Select **Import**.

Zone configurations

*The **Zone configurations** page shows all the zone configurations in your environment and lets you add new zone configurations and edit them.*

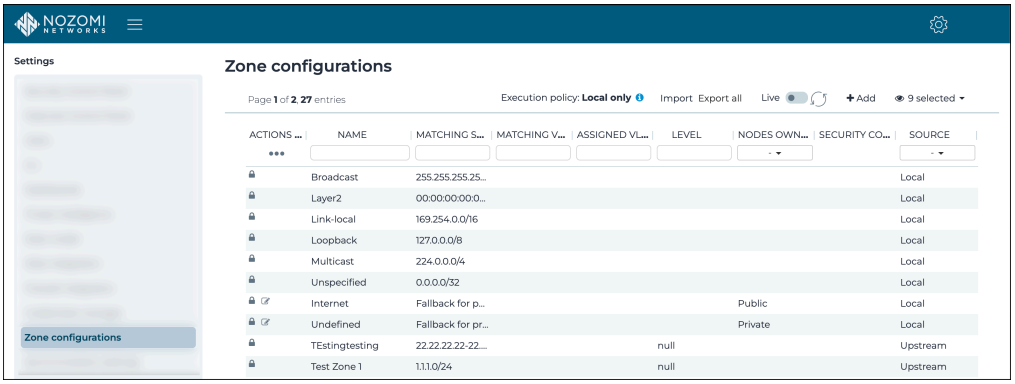


Figure 54. Zone configurations page

Execution policy

This shows information about the currently configured data synchronization policy.

Import

This button lets you import a *configuration file (CFG)*.

Export all

This button lets you download a *CFG* of all the zone configurations.

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.

+ Add

This button lets you add a new zone configuration.

Column selection

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

Name

The zone must be given a name without spaces. It must include at least one network segment.

Matching segments

This shows the information such as *IP*s, *IP* address range, or *MAC* address, matched by the node discovered within the range (only the IP range is shown in the column, not the nodes).

MAC address matching fallback

The node *ID* must match the zone network segments to make the node part of zone. There are cases where this matching strategy is not enough, for example if you want to have nodes without an *IP* as node *ID* match a zone defined with *MAC* address ranges. In those cases we can enable this fallback matching strategy in order to match against the *MAC* address of the node whenever the node *IP* does not match a segment.

Matching VLAN ID

This only lists nodes that belong to the related *virtual local area network (VLAN)*. This needs the *VLAN* tag to be extracted from the traffic.

For example, if a zone has been configured as `192.168.4.0/24`, with a *VLAN ID* set to 5:

- There is a node `192.168.4.2`, that belongs to the *VLAN*
- There is a node `192.168.4.3`, that does not belongs to the *VLAN*

When filtering the view with this zone, only the node `192.168.4.2` will show.

Assigned VLAN ID

Nodes that belong to this zone are assigned this *VLAN ID*.

Level

The level defines the position of the nodes pertaining to the given zone within the Purdue model. Once a level has been set for a zone, all nodes included in that zone are assigned the same level, unless a per-node configuration has been specified as well. This means that, if two or more zones overlap, a node that belongs to all of them will inherit the level of the most restrictive zone.

For example, if `10.1.1.1/32` belongs to Zone 1 (Level 1) and `10.1.0.0/16` (Level 2) belongs to Zone 2, then `10.1.1.1` will be assigned.

Nodes ownership

Ownership of the nodes belonging to the given zone. Once the ownership has been set for a zone, all nodes included in that zone inherit such ownership, overwriting the single nodes' ownership.

Detection approach

Used to override the global settings from the **Learning** section of the Security control panel (on page 13).

Learning mode

Used to override the global settings from the **Learning** section of the Security control panel (on page 13).

Security profile

Used to override the global settings from the **Security profile** section of the Security control panel (on page 13).

Use node labels as Device IDs

This lets you use a node label, such as the computer name, as a device *ID*.

Network Throughput History

If enabled, nodes pertaining to the zone will have an extended history for bytes sent and received, and all links for bytes transferred. The fields, whose default setting is 0:

- last_1hour_bytes
- last_1day_bytes and
- last_1week_bytes

will typically work like their counterparts for 5, 15, and 30 minutes. These fields are evaluated every 5 minutes and their time span is:

- **last_1hour_bytes =>** the last hour at granularity of 5 minutes. For example, if it is 15:32 the field will cover the time span from 14:30 to 15:30
- **last_1day_bytes =>** the last day at granularity of 1 hour, but updated every 5 minutes. For example, if it is 15:32 on Tuesday, the field will cover the time span from 16:00 on Monday to 16:00 on Tuesday and the data is updated at 15:30 on Tuesday
- **last_1week_bytes =>** the last week at granularity of 1 day, but updated every 5 minutes. For example, if it is 15:32 on Tuesday, the field will cover the time span from 00:00 on Wednesday of the previous week to 24:00 on Tuesday of the current week, and the data is updated at 15:30 on Tuesday this week



Note:
The default setting is for **Network Throughput History** to be disabled and needs to be explicitly enabled in the **Retention** tab of the **Features Control Panel**.
It is important to note that when it is activated, it will quickly consume extra disk space. The default setting for disk consumption is 512 *megabyte (MB)*.
You can configure this from 64 *MB* to 5 *GB* in the **Features** page. When the disk consumption limit is reached, older data is erased to make room for more recent samples.

Add a zone configuration

You can configure and control zones in the Central Management Console (CMC) and then propagate them to all the Guardian sensors that are connected. You can specify an execution policy in the CMC to resolve zone conflicts.

Procedure

1. In the top navigation bar, select ⚙

   Result: A menu shows.

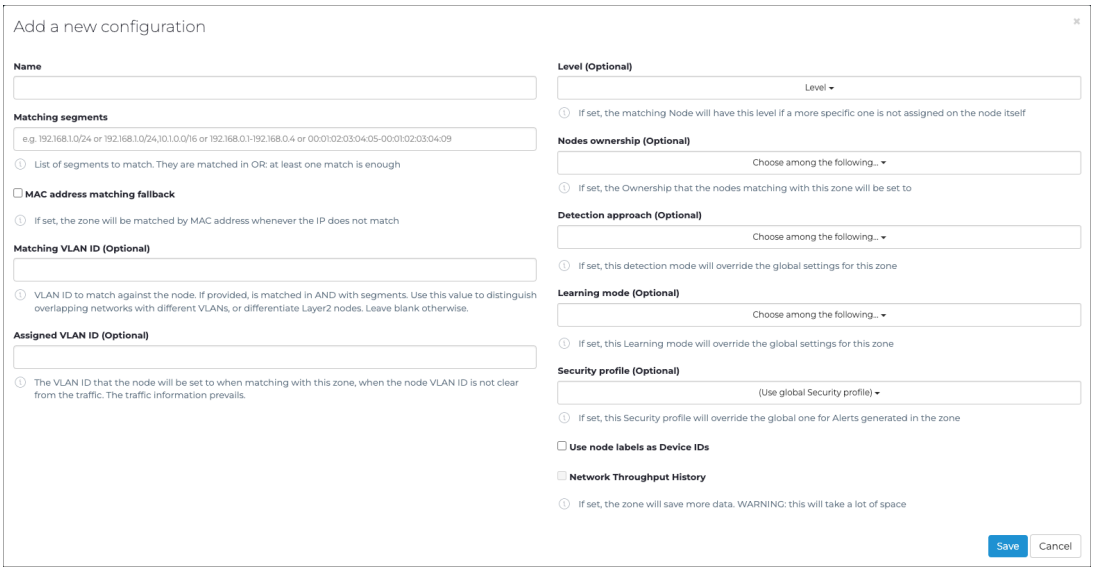
2. In the **Settings** section, select **Zone configurations**.

   Result: The **Zone configurations** page opens.

3. In the top right section, select **+Add**.

   Result: A dialog shows.

4. Configure the settings as necessary. For more details, see **Zone configurations**.



5. Select **Save**.

Results

The zone configuration has been added.

Synchronization settings

*The **Synchronization settings** page lets you customize the parameters related to Vantage or Central Management Console (CMC).*

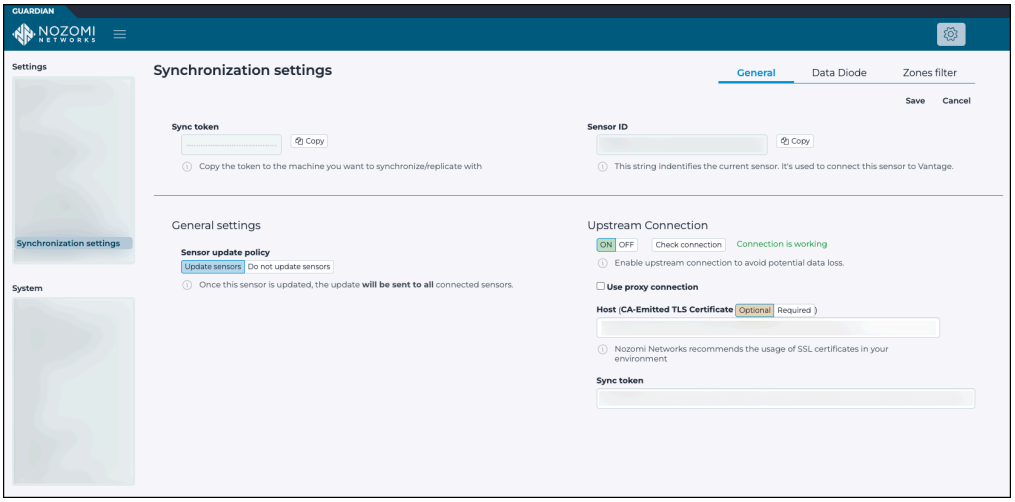


Figure 55. Synchronization settings page

The Synchronization settings page has these tabs:

- General (on page 165)
- Data Diode (on page 166)
- Zones filter (on page 176)

General

*The **General** page lets you customize the parameters related to Vantage or Central Management Console (CMC).*

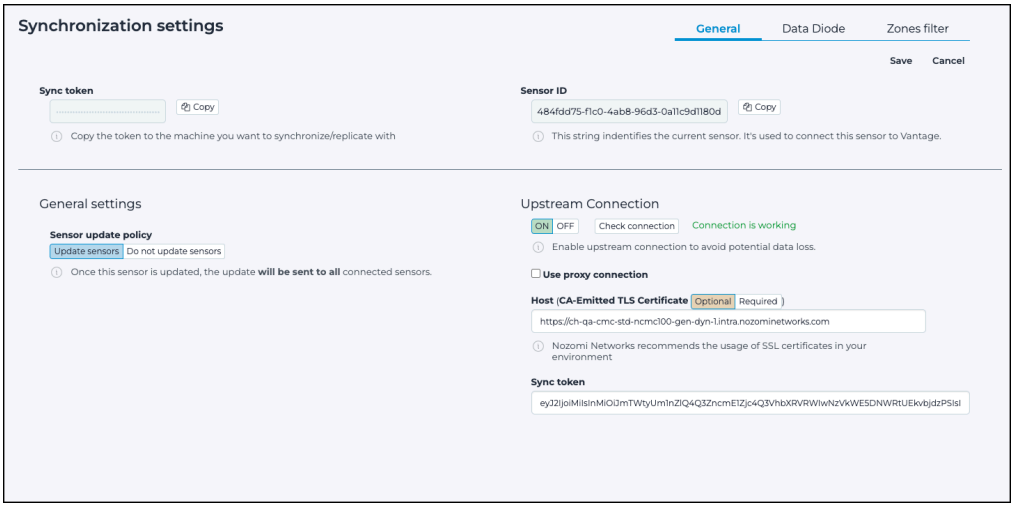


Figure 56. General page

Sync token

To synchronize/replicate with another machine, you will need to copy the **Sync token** to the machine you want to synchronize with.

Sensor ID

This is the *ID* of the current sensor.

General settings

The **Sensor update policy** has two options:

- **Update sensors**
- **Do not update sensors**

Upstream Connection

This section has a toggle that can be set to **ON** or **OFF**. When this is set to **ON**, you can select **Check connection** to make sure that the connection is okay.

The **Use proxy connection** checkbox lets you use an auxiliary system that provides a gateway between the sensor and the upstream sensor or connection.

The **Host (CA-Emitted TLS Certificate)** field is where you need to enter the *IP* address of the machine that you want to connect to. This can be set to either:

- **Optional**
- **Required**

The **Sync token** field lets you enter the synchronization token for the upstream machine.

Data Diode

Data diodes provide unidirectional network communication to securely transfer data from high-security environments, ensuring critical infrastructure and industrial control systems are protected from inbound cyber threats.



Figure 57. Data Diode page

Disabled

When this checkbox is selected, the data diode functionality is disabled.

Producer

When this checkbox is enabled, it allows the Guardian to send data out to the data diode and the upstream *CMC*.

When **Producer** is enabled, the fields below show.

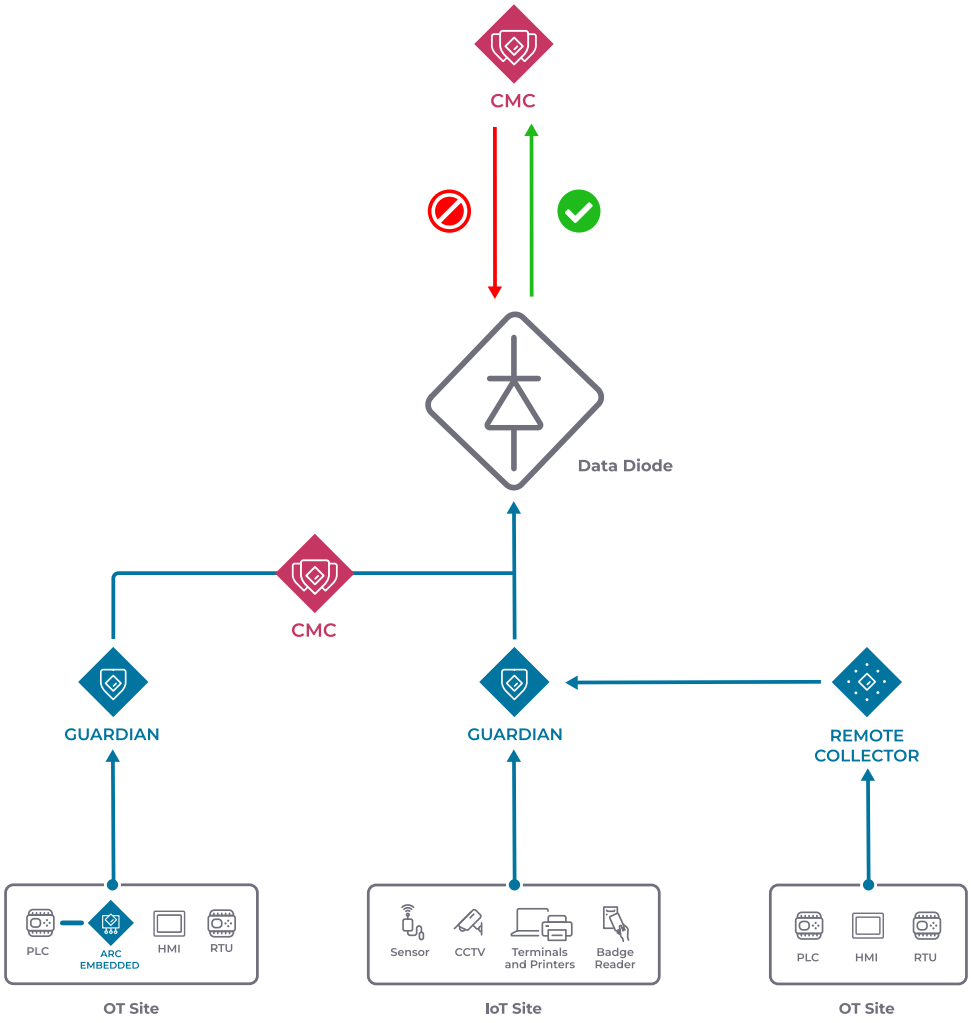
- Protocol (communication protocol to be used)
- Token (of the upstream *CMC*)
- Host (address of the data diode)
- Username (for the data diode)
- Password (for the data diode)
- Directory (the location on the data diode to store the data)

This enables one-way data transfer to the data-diode and the upstream *CMC* system.

Data diode overview

Data diodes provide unidirectional network communication to securely transfer data from high-security environments, ensuring critical infrastructure and industrial control systems (ICS) are protected from inbound cyber threats.

In cybersecurity software, data diodes are used to establish unidirectional network communication, which allows data to flow in only one direction. This ensures that sensitive networks can send out information, such as system logs or alerts, without risking inbound connections that attackers could exploit. Data diodes prevent a reverse flow of data to effectively block potential cyberattacks and data leaks.



Data diodes are particularly crucial in securing critical infrastructure and *ICS*, such as power grids and water supply systems. They allow operational data to be monitored externally without exposing these systems to outside threats. Additionally, they isolate high-security networks from less secure ones, maintaining strict separation to prevent unauthorized access while enabling necessary data sharing.

This technology secures data transfer to protect the integrity of logs and alerts, which supports compliance with regulatory requirements. As they use a physical barrier against cyber threats, data diodes enhance the security posture of organizations in sectors where safeguarding data and maintaining system integrity are critical.

Unlike firewalls, which allow controlled bidirectional communication and require ongoing management, data diodes provide a hardware-based, one-way data flow that is inherently more secure against breaches. This makes them ideal for environments that require absolute isolation.

Data diodes complement firewalls, and offer simpler, fail-safe protection for high-risk, unidirectional data transfer scenarios.

Data diode configuration

A general overview of the requirements to configure data diode architecture.

The data diode architecture has a:

- Producer (*CMC*, Guardian)
- Consumer (*CMC*)
- Data diode

The producer sends data to the consumer and the diode ensures secure one-way communication between them.

The producer can be either a *CMC*, or a Guardian.

The consumer must always be a *CMC*. The data diode is between the producer and the consumer.

To configure a data diode architecture, you must:

- Configure a CMC as a consumer (on page 169), and
- Configure a CMC as a producer (on page), or
- Configure a Guardian as a producer (on page 172)

Configure a CMC as a consumer

Configure a Central Management Console (CMC) to be a consumer of data transmitted through a data diode. This configuration ensures secure, one-way communication from a producer to a consumer.

Before you begin

Make sure that you have:

- *N2OS* v25.2.0 or later installed on both the Guardian and the *CMC*
- A data diode installed between the producer and the consumer
- Access to both systems
- This information for the data diode:
  - Host (*fully qualified domain name (FQDN)*)
  - An account with the correct privileges
  - Account username and password
  - Location to store the data

About this task

To allow secure file transfer, the data diode must authenticate to the consumer *CMC* via *SSH*.

> **Note:**
> Different types of data diodes are supported. This procedure does not go into detail about the data diode configuration.

Procedure

1. Open the consumer *CMC*.
2. In the top navigation bar, select ⚙

   Result: The administration page opens.
3. In the **Settings** section, select **Synchronization settings**.

   Result: The **Synchronization settings** page opens.
4. Select **Data Diode**.

   Result: The **Data Diode** page shows.

5. In the **Status** section, select **Consumer**.



Result: The **Token** field shows.

6. Select **Save**.
7. If an *SSH* key pair does not already exist on the data diode, create one.
   On the data diode, generate a user-owned RSA 4096 or ECDSA key for the user who will manage file transfers to the *CMC*. For example, you can use `ssh-keygen` to do this.
   a. On the data diode, use this command: `ssh-keygen -t rsa -b 4096 -f ~/.ssh/id_rsa`
   b. Copy the content of the public key (`id_rsa.pub`) to the file `/data/ssh_n2os-datadiode/ssh_authorized_keys` in the *CMC*.

> **Attention:**
> Copy as raw text, without formatting. Do not use copy-paste from Microsoft Word or other WYSIWYG text editors.

> **Note:**
> The `n2os-datadiode` user is preconfigured on the *CMC* for data diode file transfers. The directory and authorized key file should already exist.

8. Transfer files via *SSH*
   a. Make sure that the correct permissions are set for the transfer of the files.
   b. To send files from the producer to the consumer, enter the command: `scp <yourfile.ext> n2os-datadiode@<consumer-cmc-ip>:~`
   The ~ symbol is assigned to the `n2os-datadiode` target folder: `/var/data_diode_in/spool`

> **Note:**
> The *CMC* will automatically collect and process the files.

What to do next

Do the Configure a CMC as a producer (on page) or the Configure a Guardian as a producer (on page 172) procedure.

Configure a Guardian as a producer

Configure a Guardian sensor to communicate to an upstream data diode and Central Management Console (CMC).

Before you begin

Make sure that you have:

- Make sure that you have done the **Configure a CMC as a consumer** procedure
- *N2OS* v25.2.0 or later installed on both the Guardian and the *CMC*
- A data diode installed between the producer and the consumer
- Access to both systems
- This information for the data diode:
  - Host (*FQDN*)
  - An account with the correct privileges
  - Account username and password
  - Location to store the data

About this task

Use this task to configure a Guardian as a data producer for secure, one-way data transfer through a data diode. You'll set up communication with the *CMC* and establish the necessary connection for a secure data flow.

Note:
Different types of data diodes are supported. This procedure does not go into detail about the data diode configuration.

Procedure

1. Open the consumer *CMC*.
2. In the top navigation bar, select ⚙

   Result: The administration page opens.
3. In the **Settings** section, select **Synchronization settings**.

   Result: The **Synchronization settings** page opens.
4. Select **Data Diode**.

   Result: The **Data Diode** page shows.

5. To copy the token, select **Copy**.



6. Open the producer Guardian.
7. In the top navigation bar, select ⚙

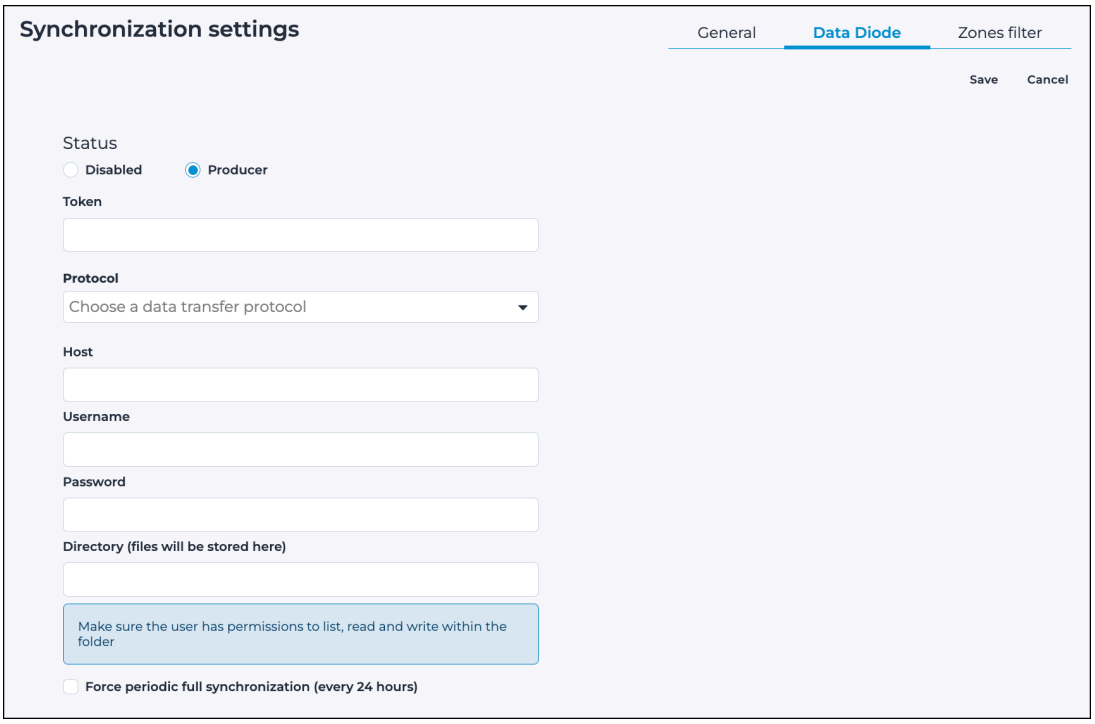
   Result: The administration page opens.
8. In the **Settings** section, select **Synchronization settings**.

   Result: The **Synchronization settings** page opens.
9. Select **Data Diode**.

   Result: The **Data Diode** page shows.
10. In the **Status** section, select **Producer**.

    Result: The producer fields show.

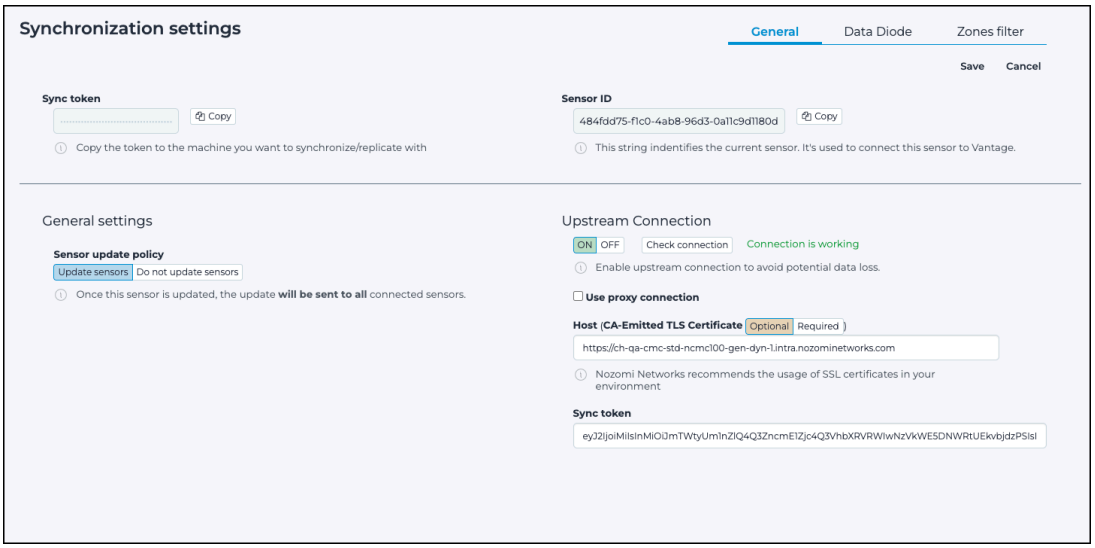
11. In the **Token** field, paste the consumer token that you copied in step 5 (on page 173).



12. From the **Protocol** dropdown, select the applicable protocol.
13. In the **Host** field, enter the *FQDN* of the data diode.
14. In the **Username** field, enter the username for the data diode.
15. In the **Password** field, enter the password for the data diode.
16. In the **Directory** field, enter the data diode location in which the data will be stored.
17. Select **Save**.
18. Select **General**.

    Result: The **General** page shows.
19. In the **Upstream Connection** section, select **ON**.
20. In the **Upstream Connection** section, select **Optional**.

21. In the **Upstream Connection** section, enter some text in the **Host** field.
    You can enter any text in the **Host** field, but you cannot leave it empty.



22. In the **Sync token** field, enter some text.
    You can enter any text in the **Sync token** field, but you cannot leave it empty.
23. Select **Save**.

Results

Communication between the producer and the consumer has been configured.

Zones filter

*Filter and manage zones to control which zones are marked as **Filtered** for upstream data transfer.*



Figure 58. Zones filter page

Columns

The **Columns** button lets you select which of the available columns for the current page will show.

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 graph.

Add a zone for upstream synchronization

*Set a zone to be **Filtered** so that entities in that zone can be updated with data from an upstream sensor. Use synchronization settings to select and enable the desired zones.*

About this task

To synchronize entities based on their zone value, you can specify a `true` | `false` value in a filter. For entities that have more than one zone, only one zone needs to be set to `true` for the entity to be synchronized.

Procedure

1. In the top navigation bar, select ⚙

   Result: A menu shows.

2. In the **Settings** section, select **Synchronization settings**.

   Result: The **Synchronization settings** page opens.

3. Choose a method to select one, or more, items.

   Choose from:

   - Select the top checkbox to select all the items in the current table view
   - Select multiple checkboxes for the items that you want to choose
   - Select the checkbox for the item that you want to choose

4. In the table, select the ••• icon.
5. Select **Add to filtered**.

Results

The value in the **Filtered** column now shows as `true` and the zone will now be synchronized with the upstream sensor.

Remove a zone from upstream synchronization

Remove a zone from synchronization so entities in that zone stop receiving updates from an downstream sensor.

About this task

For entities that have more than one zone, only one zone needs to be set to `true` for the entity to be synchronized.

Procedure

1. In the top navigation bar, select ⚙

   Result: A menu shows.

2. In the **Settings** section, select **Synchronization settings**.

   Result: The **Synchronization settings** page opens.

3. Choose a method to select one, or more, items.

   Choose from:

   - Select the top checkbox to select all the items in the current table view
   - Select multiple checkboxes for the items that you want to choose
   - Select the checkbox for the item that you want to choose

4. In the table, select the ••• icon.
5. Select **Remove from filtered**.

Results

The value in the **Filtered** column now shows as `false` and the zone will no longer be synchronized with the upstream sensor.

Connect a sensor to CMC

Connect one or more sensors to a Central Management Console (CMC).

Procedure

1. Open the *CMC*.
2. In the top navigation bar, select ⚙

   Result: A menu shows.
3. In the **Settings** section, select **Synchronization settings**.

   Result: The **Synchronization settings** page opens.
4. In the top left **Sync token** section, select **Copy**.
5. Open the sensor that you want to connect to the *CMC*.
6. In the top navigation bar, select ⚙

   Result: A menu shows.
7. In the **Settings** section, select **Synchronization settings**.

   Result: The **Synchronization settings** page opens.
8. Go to the **Upstream Connection** section, in the **Sync token** field, paste the sync token that you copied from the *CMC*.
9. In the **Host** field, paste the *IP* of the *CMC*. The *HTTPS* protocol will be used.
10. **Optional:** If no *CA*-emitted certificates are used, you can select Optional to make the verification of certificates optional.
11. To make sure that the connection has worked correctly, select **Check connection**.

    Result: If the connection is working, green text that states `Connection is working` shows.
12. Select **Save**.
13. Go back to the *CMC*.
14. In the top navigation bar, select **Sensors**.

    Result: The **Sensors** page opens.
15. In the list, make sure that the sensor you just connected shows.

> **Note:**
> When a sensor is connected for the first time, it will notify its status and receive firmware updates. However, it will not be permitted to perform additional actions. To enable a complete integration of the sensor you will need to allow the sensor.

Results

The sensor has been connected to the *CMC*.

Alert playbooks

Alert playbooks overview

An explanation of alert playbooks.

Alert playbooks are a set of instructions that guide you on how to take the correct action when an alert is raised.

An alert playbook describes the actions, tasks and other guidelines that users should follow when an alert is raised. You use an alert rule to assign an alert playbook to a specific alert. When that type of alert is raised, the alert rule that matches will insert a copy of the alert playbook into the alert.

Alert rules can use different matching criteria to assign the same alert playbook to multiple different alerts types. For more details, see **Configure an alert**, in the **User Guide**.

Once an alert playbook is visible on a triggered alert (from the Alerts panel), you can modify it without affecting the original. This is typically used to add notes for the specific alert, or to mark completed actions.

Note:

When you create an alert playbook, or an alert rule in *CMC*/Vantage, it will be propagated to all the connected sensors.

Alert playbooks

*The **Alert playbooks** page shows a list of all the alert playbooks.*

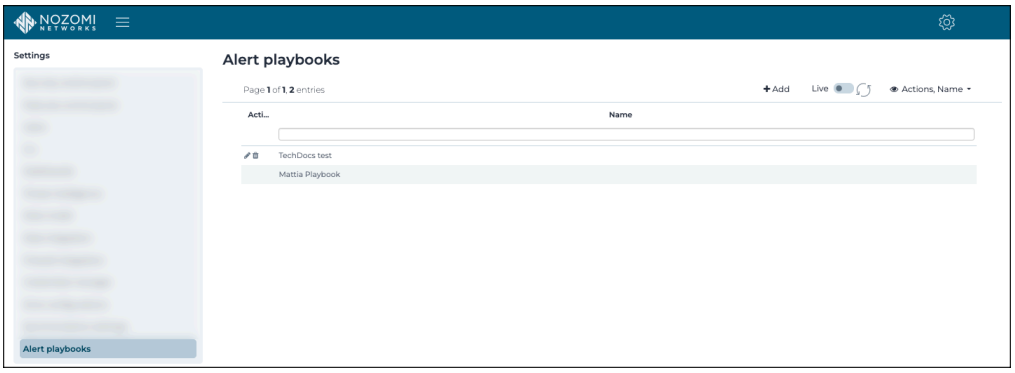


Figure 59. Alert playbooks page

Add

This lets you add new alert playbooks.

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.

Create an alert playbook

*You can use the **Alert playbooks** page to create a new alert playbook.*

Procedure

1. In the top navigation bar, select ⚙

   Result: The administration page opens.

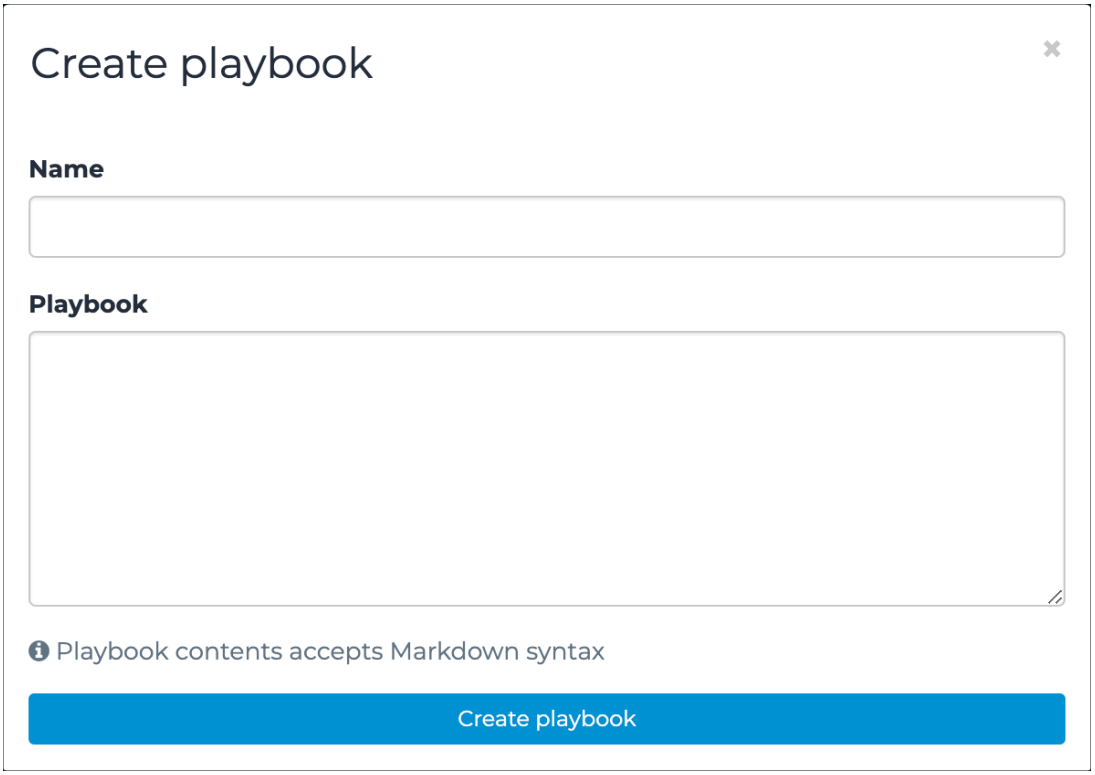
2. In the **Settings** section, select **Alert playbooks**.

   Result: The **Alert playbooks** page opens and shows a list of available alert playbooks.

3. In the top right, select **+ Add**.

   Result: A dialog shows.

4. In the **Name** field, enter a name for the alert playbook.



5. In the **Playbook** field, enter the relevant steps to follow when an alert related to this playbook is raised.
6. Select **Create playbook**.

Results

The alert playbook has been created.

Edit an alert playbook

After an alert playbook has been created, you can edit the details.

Procedure

1. In the top navigation bar, select ⚙

   Result: The administration page opens.

2. In the **Settings** section, select **Alert playbooks**.

   Result: The **Alert playbooks** page opens and shows a list of available alert playbooks.

3. From the list, select an alert playbook template.

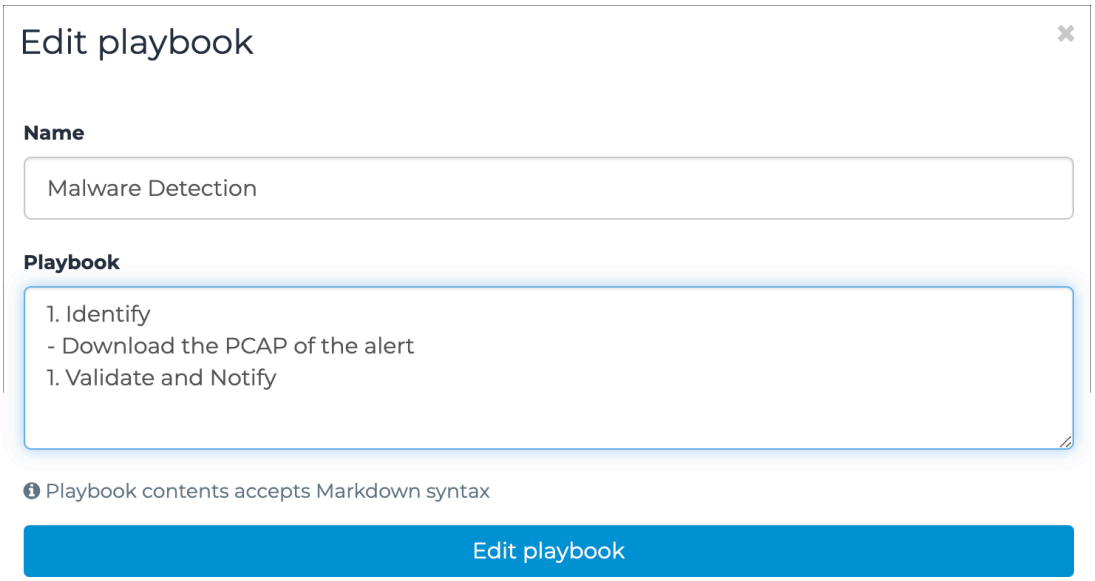
   Note:
   The list will be empty if:
   - No alert playbooks have been created
   - One or more alert playbooks have been created, but they were created in an upstream Vantage or *CMC*

4. To the left of the applicable alert playbook, select the ✎ icon.

   Result: A dialog shows.

5. **Optional:**

   If necessary, in the **Name** field, edit the text.



6. **Optional:** If necessary, in the **Playbook** field, edit the text.
7. Select **Edit playbook**.

Results

The alert playbook has been edited.

Delete an alert playbook

*You can use the **Alert playbooks** page to delete an alert playbook.*

Procedure

1. In the top navigation bar, select 

   Result: The administration page opens.

2. In the **Settings** section, select **Alert playbooks**.

   Result: The **Alert playbooks** page opens and shows a list of available alert playbooks.

3. To the left of the applicable alert playbook, select the  icon.

   > **Note:**
   > If an alert rule(s) that is related to the alert playbook, you will be prompted to confirm the deletion. If you proceed with the deletion, all alert rules associated with the alert playbook are also deleted.

   Result: A dialog shows if the alert playbook is related to an alert rule(s).

4. To confirm the deletion, select **Yes**.



Results

The alert playbook has been deleted.

Update service contents

*The **Update service contents** page lets you manage how the appliance downloads Threat Intelligence, Asset Intelligence, and Arc content, to reduce bandwidth consumption on slow or unstable network connections.*

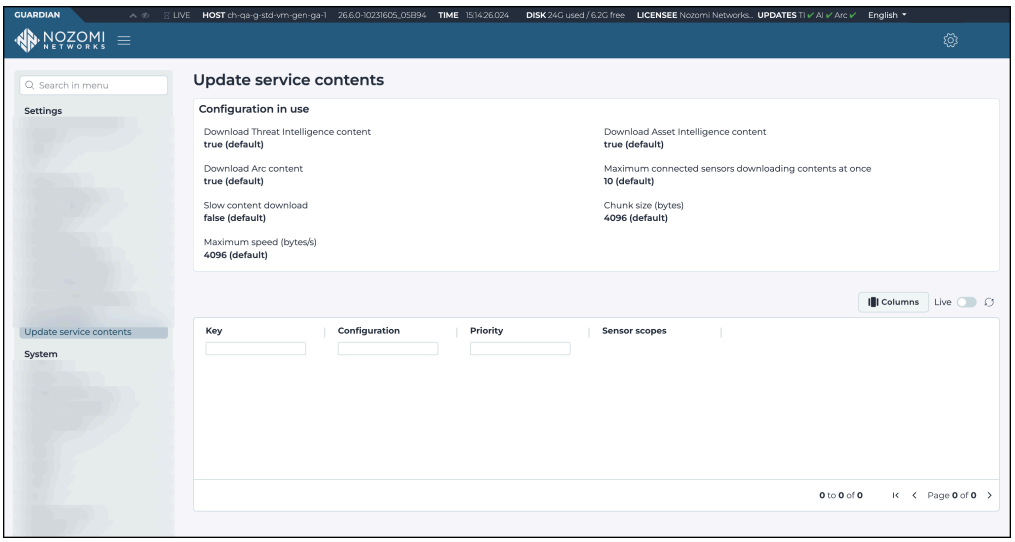


Figure 60. Update service contents page in Guardian

The page has these sections:

- **Configuration in use**: The configuration currently applied to the appliance
- **Configurations table**: The custom configurations that override the default values for the appliances in their scopes. The configurations table is read-only. Custom configurations are managed from a *CMC*

The configuration has these settings:

- **Download Threat Intelligence content**, **Download Asset Intelligence content**, and **Download Arc content**: Enable or disable the download of each type of content. Disable a type of content when the corresponding functionality is not in use, to prevent unnecessary data from being transferred to the appliance
- **Slow content download**: Enable a slower, more controlled download strategy. When it is enabled, these settings are available:
  - **Maximum speed (bytes/s)**: Limits the bandwidth used by the update service, preventing content downloads from saturating the available network connection
  - **Chunk size (bytes)**: Controls the size of individual content chunks. Instead of downloading a large content package as a single transfer, the content is split into smaller chunks, and the download process briefly releases bandwidth between chunks. Splitting content into smaller chunks also improves resilience on unstable connections: if the connection is interrupted, only the affected chunk needs to be retransmitted
- **Maximum connected sensors downloading contents at once**: Controls how many downstream appliances that use the current appliance as an upstream source are allowed to download content concurrently. Limiting the number of simultaneous downloads reduces the workload on the upstream appliance
- **Priority**: The priority of a custom configuration. When several configurations apply to the same appliance, the one with the lowest priority value applies. If they have the same priority, the one created first applies

The settings have these default values and valid ranges:

- **Download Threat Intelligence content**, **Download Asset Intelligence content**, and **Download Arc content**: Enabled by default
- **Slow content download**: Disabled by default
- **Maximum speed (bytes/s)**: Default 4096. Valid values range from 1024 to 1073741824
- **Chunk size (bytes)**: Default 4096. Valid values range from 128 to 10485760
- **Maximum connected sensors downloading contents at once**: Default 10. Valid values range from 1 to 100
- **Priority**: Default 3. Valid values range from -100 to 100



Note:

The available contents depend on the licenses installed on the machine.

System

General

*The **General** page lets you change the hostname and specify a login banner on a sensor.*

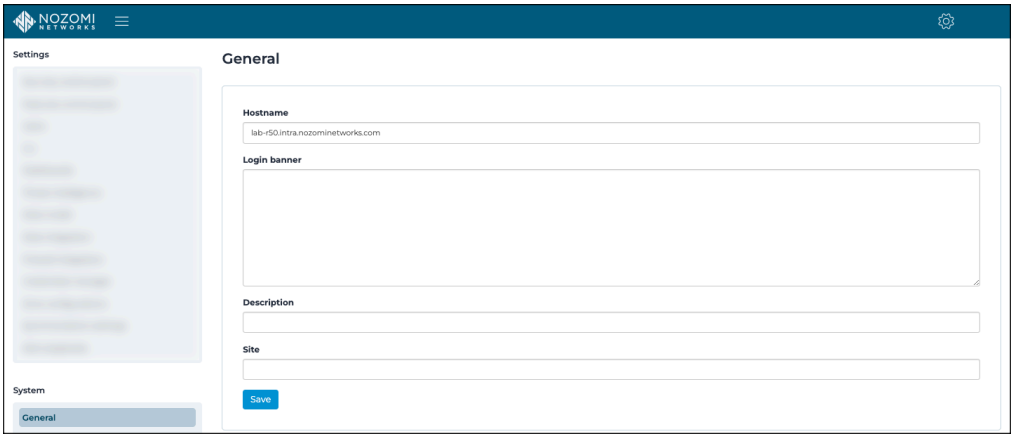


Figure 61. General page

Date and time

*The **Date and time** page lets you configure the date and time of a sensor.*

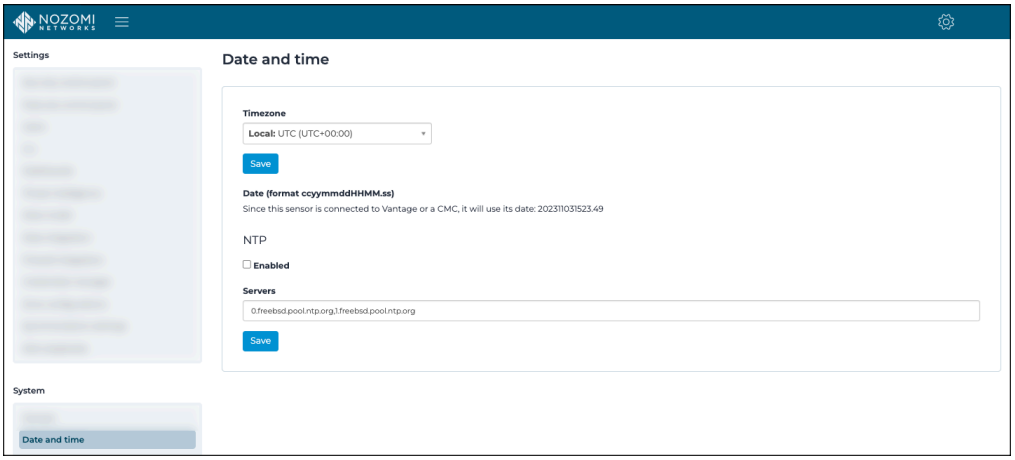


Figure 62. Date and time page

Timezone

This dropdown lets you select a timezone.

Date

This section lets you **Pick a date** and, optionally set the sensor as a client. The date settings are disabled if the sensor is connected to an upstream Vantage or *CMC*.

NTP (Network Time Protocol)

The **Enabled** checkbox lets you synchronize the servers in the network to the computer clock time. To do this you must enter a list of comma-separated server addresses.

Note:
When a sensor is attached to Vantage or a *CMC*, you cannot manually set its date and time. Sensors that are connected to Vantage or a *CMC* (with no *network time protocol (NTP)* configured) will automatically get time synchronization from Vantage, or the parent *CMC*.

Configure the date and time

You can configure the timezone and the current date and time of the sensor. You can also write a list of comma-separated server addresses to enable time synchronize with a network time protocol (NTP).

Procedure

1. In the top navigation bar, select 

   Result: The administration page opens.

2. In the **System** section, select **Date and time**.

   Result: The **Date and time** page opens.

3. From the **Timezone** dropdown, select the correct option.
4. Select **Save**.
5. In the **Date** field, select **Pick a date**.

   > **Note:**
   > If the sensor is connected to Vantage or a *CMC*, the date is automatically taken from upstream.

6. **Optional:** If applicable, select **Set as client**.
7. Select **Save**.
8. **Optional:** To synchronize the servers in the network to the computer clock time, in the **NTP** section, select **Enabled**.
9. **Optional:** In the **Servers** field, enter the applicable list of comma-separated server addresses.
10. Select **Save**.

Results

The date and time settings have been configured.

Updates and licenses

*The **Updates and licenses** page lets you view and manage the licenses installed on this Guardian.*

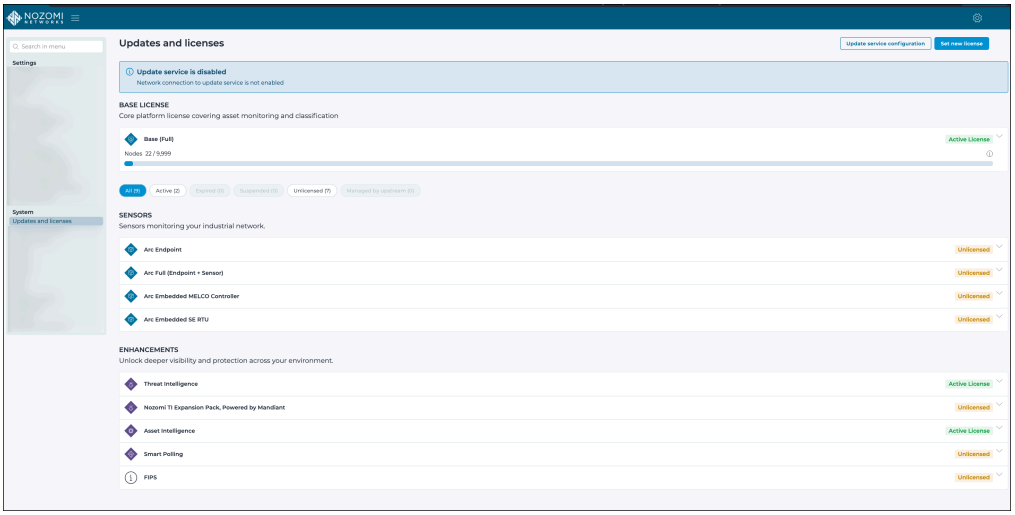


Figure 63. Updates and licenses page

For more details about the license types shown on this page, see License types (on page 193).

Filtering licenses

Select a filter to show only the licenses in that **License status**: **All**, **Active**, **Expired**, **Suspended**, **Unlicensed**, or **Managed by upstream**. Each filter shows a count of the licenses in that status.

Node quotas

Each license shows the number of nodes it currently covers against the maximum allowed by that license, for example `Nodes 22 / 9,999`. If a license is issued with node limits and those limits are exceeded, functionality is only enabled for the nodes within the covered limit; new nodes are not analyzed.

Update service

This section only shows on a Guardian that is not connected to a *CMC*. It controls how Threat Intelligence and Asset Intelligence content updates are delivered to this Guardian. If the Guardian has no network connection to the Nozomi Networks update service, a banner shows at the top of the page: **Update service connection issue**.

Select **Update service configuration** to open a panel where you can:

- Enable or disable an automatic network connection to the Nozomi Networks update service, including proxy settings, and check whether that connection is working.
- Force an immediate content update over that connection.
- Switch to manual mode to upload a content update package directly instead of connecting to the update service.

License types

Separate licenses control the base software and optional sensors and enhancements, such as Arc, Smart Polling, Threat Intelligence, Asset Intelligence, the Nozomi TI Expansion Pack Powered by Mandiant, and Federal Information Processing Standards (FIPS).

Base license

This section contains the base license, which is mandatory and includes passive monitoring.

Sensors

This section contains optional licenses for sensor-related products.

Arc Endpoint

Lets you enable an optional license for Arc Endpoint.

Arc Full (Endpoint + Sensor)

Lets you enable an optional license for Arc Full (Endpoint + Sensor).

Arc Embedded Melco Controller

Lets you enable an optional license for Arc Embedded Melco Controller.

Arc Embedded SE RTU

Lets you enable an optional license for Arc Embedded Schneider Electric *remote terminal unit (RTU)*.

Enhancements

This section contains optional licenses for enhancement features.

Smart Polling

Lets you enable an optional license for Smart Polling.

Threat Intelligence

Lets you enable an optional license for Threat Intelligence.

Asset Intelligence

Lets you enable an optional license for Asset Intelligence.

Nozomi TI Expansion Pack, Powered by Mandiant

Lets you enable an optional license for the Nozomi TI Expansion Pack, Powered by Mandiant.

FIPS

Information-only panel that shows details about an optional *Federal Information Processing Standards (FIPS)* license. To install *FIPS*, you need to enable *FIPS* mode in a console.

Compliance - Arc

Lets you enable an optional license for compliance scanning on Arc sensors.

License status

Each section that has an active license will show the **License status**.

Table 3. License status

License status	Description
UNLICENSED	Functionality is disabled.
ACTIVE	Functionality is enabled. Be aware of the expiration date to allow time for renewals. If a license is issued with node limits, and the limits are exceeded, functionality is only enabled for the covered nodes within the limits. New nodes are not analyzed.
EXPIRING	30 days before the expiration date, notifications will show in the *UI* to remind users of the impending expiration of their current license and to take actions towards renewing their license.
EXPIRED	Starting from the expiration date, notifications will show in the *UI* to notify users of the expiration of their current license and to take actions towards renewing their license. Nozomi Networks offers a grace period after the official expiration date. In this grace period, the license still functions as it would in the **ACTIVE** status, providing time for an emergency renewal of the license.
SUSPENDED	The grace period for an **EXPIRED** license has ended, and functionality is disabled. The content that was analyzed or imported before the license expired remains, however no new analyses are performed, and no new signatures are imported.
MANAGED_BY_UPSTREAM	The license is managed by an upstream sensor, such as *CMC* or Vantage. This status is only visible when the sensor is connected and synced to an upstream sensor.
NOT_REQUIRED	A license is not required for this feature on this device. For example, *FIPS* on a *CMC* does not require a separate license.

Install a license

Before you can use a product, you must install the applicable license.

Before you begin

If you are installing an additional license, make sure that you have installed a base license first.

Procedure

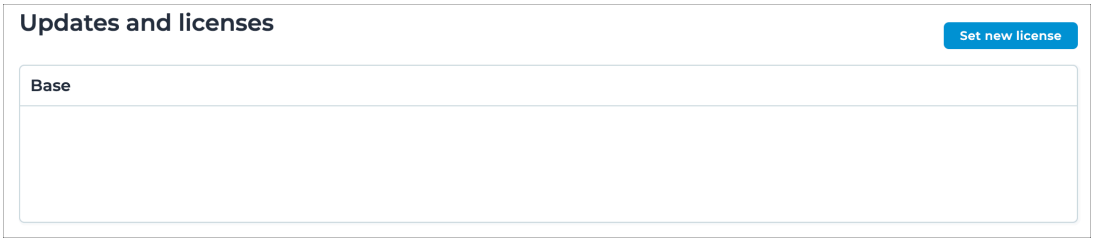
1. In the top navigation bar, select ⚙

   Result: The administration page opens.

2. In the **System** section, select **Updates and licenses**.

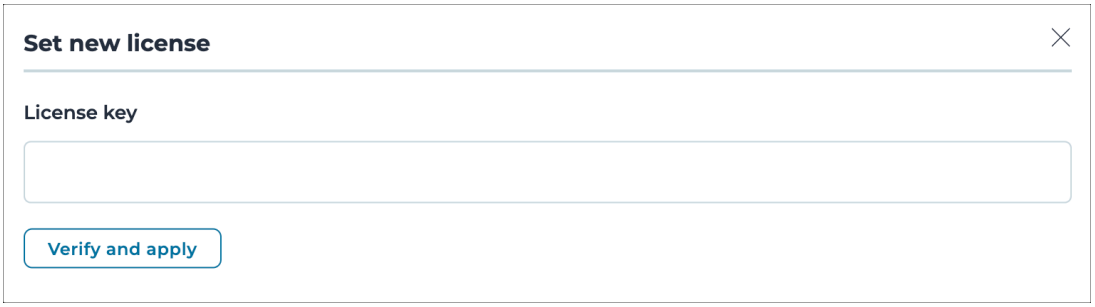
   Result: The **Updates and licenses** page opens.

3. In the top right of the section, select **Set new license**.



   Result: A dialog shows.

4. To copy the **Machine ID**, select **Copy**.



5. Send the machine *ID* to Nozomi Networks with your license request.
6. Wait to receive your license key from Nozomi Networks.
7. In the **License key** field, paste the license key.
8. Select **Verify and apply**.

Results

The license has been installed.

Network interfaces

The Network interfaces page shows information on how much network traffic is being ingested/captured.

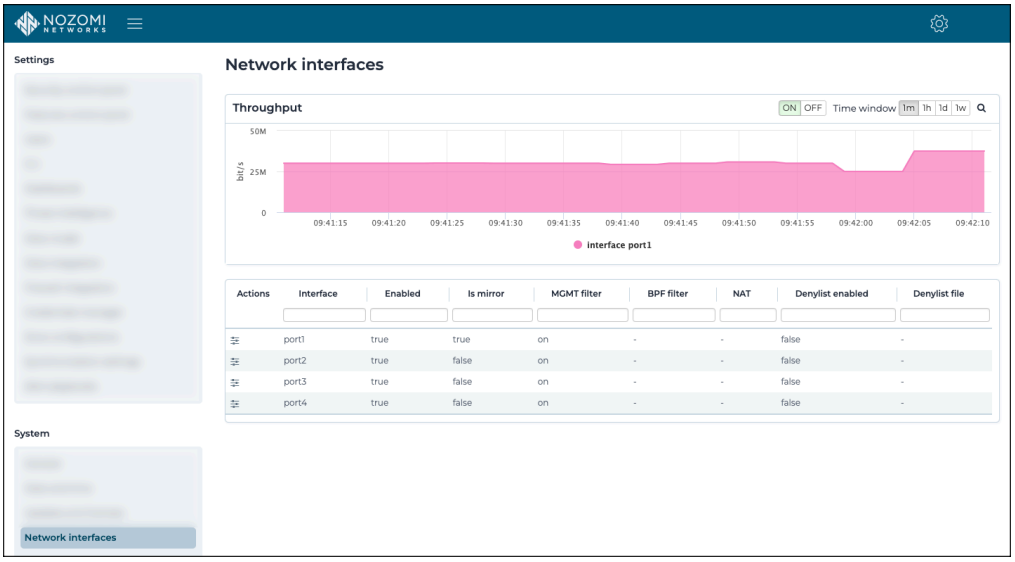


Figure 64. Network interfaces page

Throughput

This shows the amount of information that is being shown.

ON/OFF

This toggle lets you set the automatic update of the diagram.

Time window

This lets you choose the timeframe of the window. You can choose between:

- One minute
- One hour
- One day
- One week

Configure a NAT rule

*The **Network interfaces** page lets you configure network address translation (NAT) rules.*

Procedure

1. In the top navigation bar, select 

   Result: The administration page opens.

2. In the **System** section, select **Network interfaces**.

   Result: The **Network interfaces** page opens.

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

   Result: A dialog shows.

4. In the **Label** field, enter a label for the interface.



Note:
The label will show instead of the network interface name in all areas of the user interface.

5. **Optional:** To disable the network interface from sniffing traffic, select the toggle to **OFF**.
6. In the **NAT** section, configure the settings as necessary for the *network address translation (NAT)* rule.



Note:
The *NAT* rule lets you rewrite the source and destination *IP* addresses of packets sniffed on this interface. For example, to translate `192.168.1.100` in `10.1.1.100` you have to configure the rule: `192.168.0.0` `10.1.0.0 /16`.

   a. In the **Original subnet** field, enter a value.
   b. In the **Translated subnet** field, enter a value.
   c. In the **CIDR mask** field, enter a value.
7. Select **Save**.

Results

The *NAT* rule has been configured.

Configure a BPF filter

*The **Network interfaces** page lets you configure Berkeley Packet Filters (BPF).*

Procedure

1. In the top navigation bar, select ⚙

   Result: The administration page opens.

2. In the **System** section, select **Network interfaces**.

   Result: The **Network interfaces** page opens.

3. To the left of the applicable interface, select the ⚙ icon.

   Result: A dialog shows.

4. In the **Label** field, enter a label for the interface.



Note:
The label will show instead of the network interface name in all areas of the user interface.

5. **Optional:** To disable the network interface from sniffing traffic, select the toggle to **OFF**.
6. In the **BPF filter** section, choose a method to use to configure the *Berkeley Packet Filter (BPF)* settings.

   Choose from:

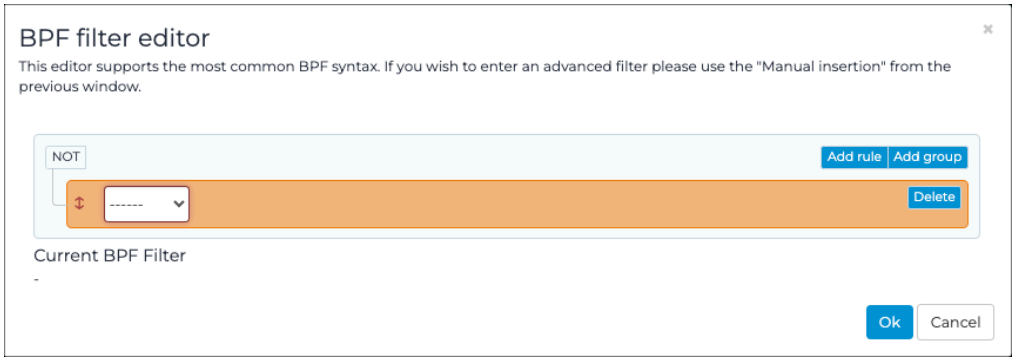
   - To use the visual editor, go to step 7 (on page 201)
   - To manually enter a filter, go to step 8 (on page 201)

7. Use the visual editor.
    a. Select **BPF filter editor**.



    Result: A dialog shows.

    b. Edit the settings as necessary.



    c. Select **Ok**.

8. Manually enter a *BPF*.



    a. Set the **Manual insertion of a custom expression** toggle to on.
    b. Manually enter the *BPF*.
    c. Select **Save**.

Results

The *BPF* has been configured.

Configure a denylist

*The **Network interfaces** page lets you configure denylists.*

About this task

In the **Denylist** section, you can upload a text file that contains a denylist. This can contain a list of *IP* addresses that are explicit, or with netmasks, or wildcards, that Guardian will not process. A wildcard in digit 2, 3 or 4 is equivalent to a /8, /16 or /24 netmask. The effect is similar to that of a *BPF*, however a denylist can handle tens of thousands of *IP* addresses, which is beyond the capabilities of a *BPF*.

Procedure

1. In the top navigation bar, select 

   Result: The administration page opens.

2. In the **System** section, select **Network interfaces**.

   Result: The **Network interfaces** page opens.

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

   Result: A dialog shows.

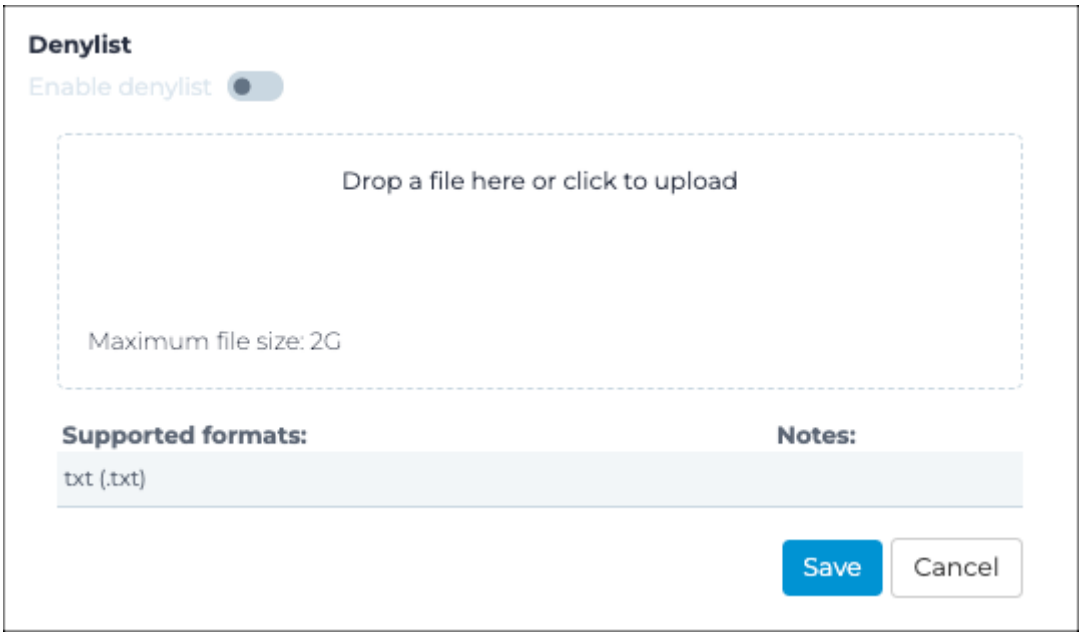
4. In the **Label** field, enter a label for the interface.



Note:
The label will show instead of the network interface name in all areas of the user interface.

5. **Optional:** To disable the network interface from sniffing traffic, select the toggle to **OFF**.

6. In the **Denylist** section, set the **Enable denylist** toggle to on.



7. Choose a method to upload the denylist file:

   Choose from:

   - Drag the denylist file into the **Drop a file here or click to upload** field
   - Click in the **Drop a file here or click to upload** field and located the file to be uploaded

> **Note:**
>
> A denylist must contain:
>
> - One entry per line
> - A dash (-) followed by a space and an *IP* address. Optionally, it can contain a wild card or a netmask.

> **Note:**
>
> The maximum file size is 2 *GB*. The supported file type is text files (.txt).

Results

The denylist has been configured.

Denylist examples

Table 4. Denylist examples

`- 192.168.1.*`
`- 192.168.2.1`
`- 192.168.2.6...192.168.3.90`
`- 192.168.4.1/20`
`- 192.168.5.1 tcp 80`
`- 192.168.5.2 tcp 80 100-110`
`- 192.168.5.3 udp 11 22 2000-30000`
	A wildcard in the IP in position 2, 3 or 4 is equivalent to a subnet ending in /8, /16 or /24. The last lines block only packets from a specific host/port (or set of ports).

`- 192.168.2.*` `- 192.168.2.1`	These entries are not disjoint, but the Guardian is able to fix automatically the conflict.
`- 192.168.3.2 tcp 80` `- 192.168.3.1/24 tcp 100`	A logical error: 192.168.3.2 is an ambiguous match.

`- 192.168.3.1 tcp 100`
`- 192.168.3.2 tcp 80`
`- 192.168.3.3...192.168.3.255 tcp 100`
	As the entries do not intersect, this is the right way to fix the error above.

`- 192.168.3.3 tcp 80` `- 192.168.3.3 udp 100`	These lines are in conflict: an ip can be associated with at most one port filter.

Upload traces

*The **Upload traces** page lets you upload trace files and select the options for the uploaded file.*

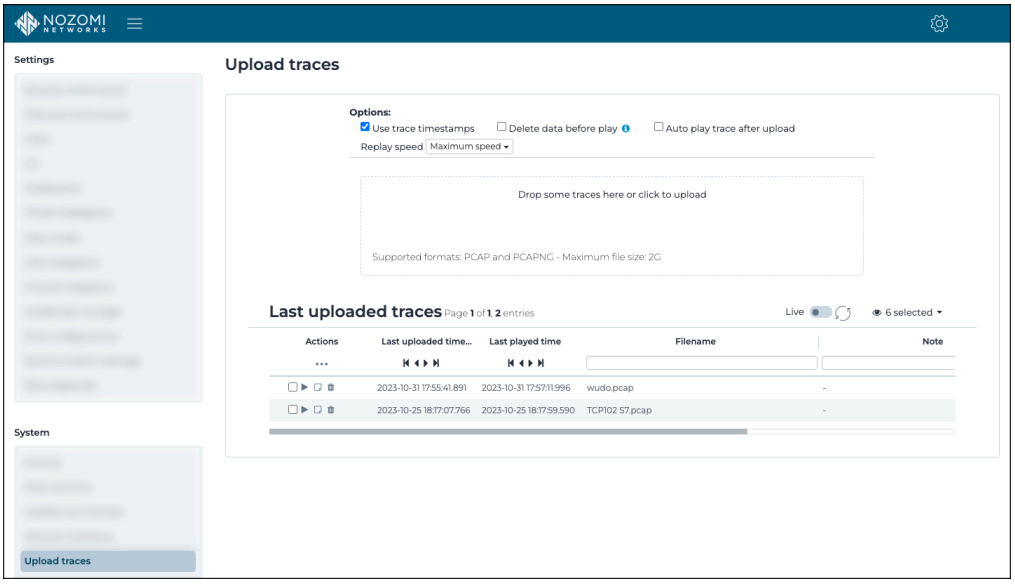


Figure 65. Upload traces page

Use trace timestamps

You can select this option to use the time captured in the trace file. Otherwise, the current time is used.

Important:

We recommend not selecting this checkbox because it causes data to be hidden due to time filters.

Delete data before play

Select this option to delete the data in the sensor before you run the play action. When multiple traces are played at once, deletion is applied only before running the first trace.

Auto play trace after upload

Select this option to play the trace immediately after the upload is complete.

Replay speed

You can choose from the available options to customize the maximum throughput at which the *pcap* is played.

Upload field

You can click in the upload field, or drag a file into the upload field to upload a trace. The recommended formats are:

- *pcap*
- *packet capture next generation (pcapNg)*

The maximum file size that you can upload is 2 *GB*.

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.

Actions

The Actions column lets you:

- Replay a trace file (on page 209)
- Edit a note for a trace file (on page 209)
- Delete a trace file (on page 211)

Last uploaded traces

Before you take an action, you can use this element to filter and sort the traces. You can select the text in the header to sort in ascending or descending order. You can then select the header again to toggle between the two.

Last played time

Before you take an action, you can use this element to filter and sort the traces. You can select the text in the header to sort in ascending or descending order. You can then select the header again to toggle between the two.

Upload a trace file

*The **Upload traces** page lets you upload a trace file.*

Procedure

1. In the top navigation bar, select ⚙

   Result: The administration page opens.

2. In the **System** section, select **Upload traces**.

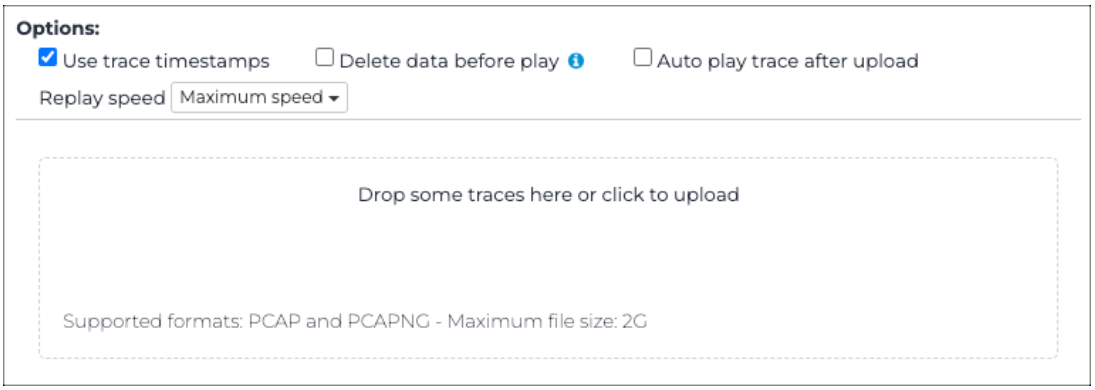
   Result: The **Upload traces** page opens.

3. Choose a method to upload a trace file.

   Choose from:

   - Drag your image file into the **Drop some traces here or click to upload** field
   - Click in the **Drop some traces here or click to upload** field

4. If you chose the second method, select the correct file to upload.



> **Note:**
> Supported formats are:
>
> - *pcap*
> - *pcapNg*
>
> The maximum permitted file size is 2 *GB*.

5. Wait for the file to upload.

Results

The trace file has been uploaded.

Replay a trace file

After you have uploaded a trace file, you can replay it whenever you want.

About this task

As a result of playing traces, alerts can be generated. If the played file is artificial, the system might not recognize the alert timestamp. In this case, a value containing `InvalidDate` shows in the time column of the alert table.

Procedure

1. In the top navigation bar, select ⚙

   Result: The administration page opens.

2. In the **System** section, select **Upload traces**.

   Result: The **Upload traces** page opens.

3. Choose a method.

   Choose from:

   - To replay a single file, go to step 4 (on page 209)
   - To replay multiple files, go to 5 (on page 209)

4. To replay a single file, to the left of the applicable trace file, select the ▶ icon.
5. Replay multiple files.
   a. Select multiple checkboxes for the items that you want to choose.
   b. In the Actions column, select the ••• icon.
   Result: A menu shows.
   c. From the menu, select **Play selected**.

Results

The trace file(s) has (have) been replayed.

Edit a note for a trace file

After you upload a trace file, you can write a note for it, or edit an existing one.

Procedure

1. In the top navigation bar, select ⚙

   Result: The administration page opens.

2. In the **System** section, select **Upload traces**.

   Result: The **Upload traces** page opens.

3. To the left of the applicable trace file, select the icon.
4. Write or edit the details as necessary.
5. Select **Save**.

Results

The note has been edited.

Delete a trace file

After you have uploaded a trace file, you can delete it.

Procedure

1. In the top navigation bar, select 

   Result: The administration page opens.

2. In the **System** section, select **Upload traces**.

   Result: The **Upload traces** page opens.

3. To the left of the applicable trace file, select the  icon.

Results

The trace file has been deleted.

Export

*The **Export** page lets you export data in content packs.*

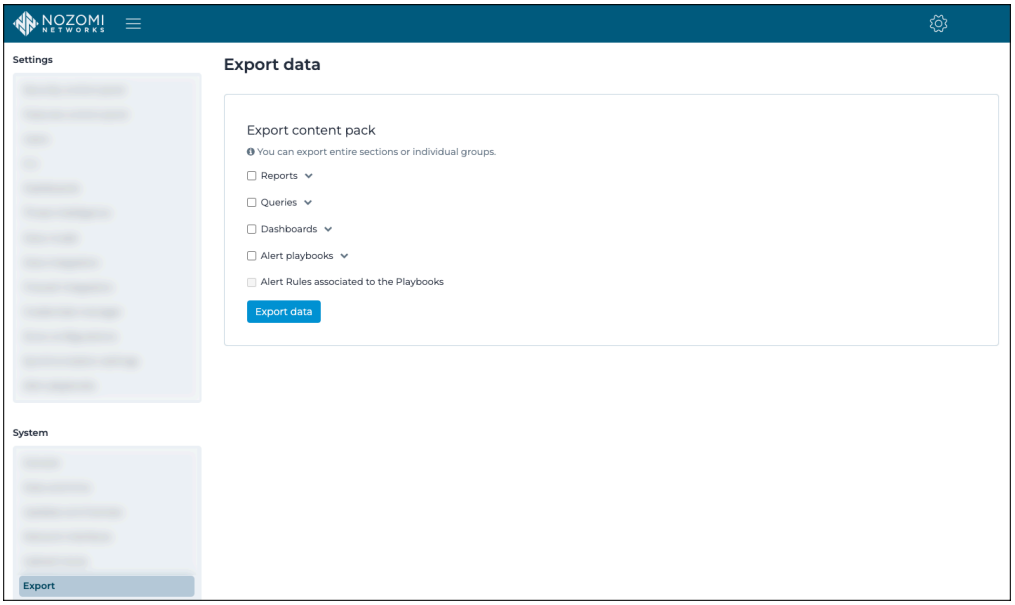


Figure 66. Export page

General

You can export data in content packs, which are files in *JSON* format that you can open and read with a text reader.

Content packs

Content packs are a reporting and query feature that packages multiple templates into a single file for team collaboration. A single content pack can contain one or more:

- Reports
- Queries
- Dashboards
- Alert playbooks

> **Note:**
> When you select **Alerts Playbooks**, you get the option to choose **Alert Rules associated with the Playbooks**.

Content packs package data in a single file, which can then be distributed to a team of people. They can then be edited and used across multiple systems. This is especially useful in complex reporting arrangements, such as compliance with government regulations, or hunting for a specific threat.

The file expands so you can add other information in the *JSON* format to the content pack. The Nozomi Networks product ignores data that it doesn't understand and continues to parse the file. This lets you add data for other systems to the content pack.

Note:
When you export a dashboard and insert the related content pack into a new Guardian instance, the dashboards load and function the same as those from the original Guardian. Imported dashboards include all the queries and widgets associated with the original saved dashboards.

Export a content pack

*You can use the **Export** page to export data in content packs which can contain report, queries, dashboards, or alert playbooks.*

About this task

This procedure shows you how to export a content pack from the **Export** page. To export directly from a dashboard, see **Settings > Dashboards > Export a dashboard** .

Procedure

1. In the top navigation bar, select 

   Result: The administration page opens.

2. In the **System** section, select **Export**.

   Result: The **Export data** page opens.

3. In the Export content pack section, select options that you want to export.
4. **Optional:** Export reports.
   a. Select the **Reports** checkbox.
   b. Select the  icon.
   Result: The section expands.
   c. **Optional:** If necessary, select individual items as necessary.
5. **Optional:** Export queries.
   a. Select the **Queries** checkbox.
   b. Select the  icon.
   Result: The section expands.
   c. **Optional:** If necessary, select individual items as necessary.
6. **Optional:** Export dashboards.
   a. Select the **Dashboards** checkbox.
   b. Select the  icon.
   Result: The section expands.
   c. **Optional:** If necessary, select individual items as necessary.

7. **Optional:** Export alert playbooks.
    a. Select the **Alert Playbooks** checkbox.
    b. Select the ⌄ icon.
    Result: The section expands.
    c. **Optional:** If necessary, select individual items as necessary.
    d. **Optional:** If necessary, select **Alert Rules associated to the Playbooks**.
8. Select **Export data**.

Import

*The **Import** page lets you import data from multiple different sources, and in multiple formats.*

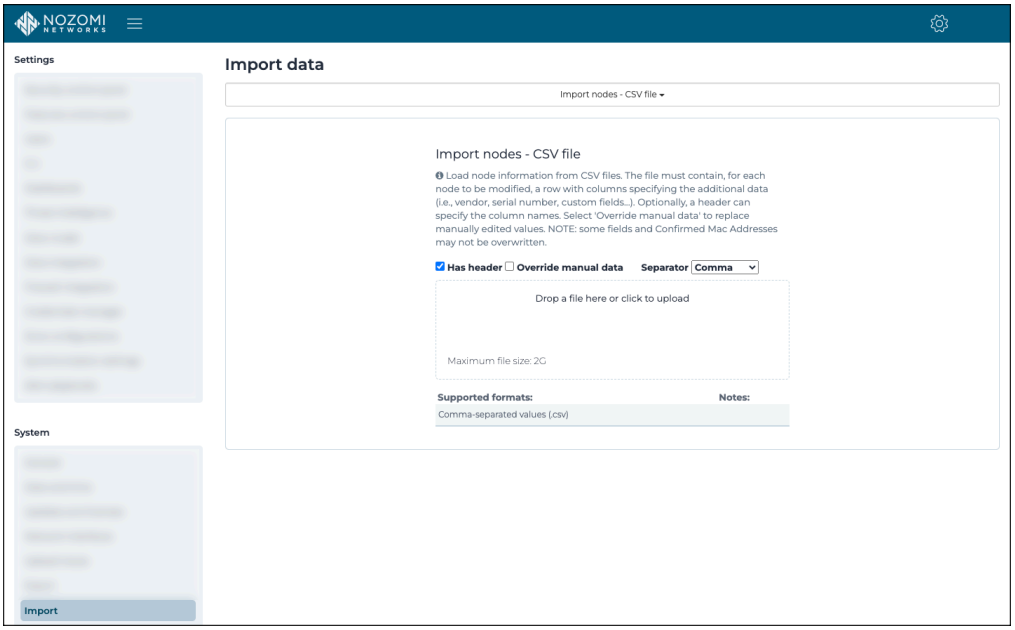


Figure 67. Import page

Import nodes

This feature lets you bind fields from a *CSV* files to Nozomi Networks fields so that you can add nodes (flag **create non-existing nodes**), or enrich existing ones. For each modified node, the file must include a row with columns that specify the additional data.

Note:
Some special fields and confirmed *MAC* addresses are restricted to specific values and so they might not be overwritten.

Import variables

This feature lets you add new variables (flag **create non-existing variables**), or to enrich existing variables. For each modified variable, the file must include a row with columns that specify the additional data.

Import asset types

This feature lets you enlarge the built-in set of asset types with a set of new custom types. The *CSV* file should include a header row labeled **Name** and the list of asset type names in the following rows, one per row. A name identifies each asset type. This means that, during the import process, duplicate names are ignored and notified.

Table 5. Default asset types

actuator	inverter	printer_scanner
audio_video	IO_module	radio_transmitter
AVR	IOT_device	robot
barcode_reader	light_bridge	router
camera	media_converter	RTU
computer	medical_imager	sensor
controller	meter	server
digital_io	mobile_phone	switch
drone	network_security_appliance	tablet
DSL_modem	OT_device	time_appliance
firewall	other	UPS
gateway	PDU	VOIP_phone
HMI	PLC	WAP
IED	power_generator	
infusion_system	power_line_carrier	

Import configuration / project file

This feature lets you import a project file. The information written in the project file is added to the asset data.

Note:
For XML Profinet GSDML, you can import an *XML* Profinet GSDML file that describes a physical device. Information in the file lets Guardian extract and display information about the device configuration when device configuration packets are displayed in traffic.

Import content pack

You can use the **Export** page to export data in content packs and then you can use the **Import** page of a different machine to import them. For more details, see Export (on page 212).

Import Arc data archive

When Arc is in **Offline** mode, the data is collected locally and then exported in an archive file from the machine. The archive file can then be manually imported into *CMC*, Guardian, or Vantage.

Import nodes from a CSV file

*You can use the **Import** page to import nodes from a comma-separated value (CSV) file.*

About this task

This feature lets you bind fields from a *CSV* files to Nozomi Networks fields so that you can add nodes, or enrich existing ones. For each modified node, the file must include a row with columns that specify the additional data.

Note:
Some special fields and confirmed *MAC* addresses are restricted to specific values and so they might not be overwritten.

Procedure

1. In the top navigation bar, select 

   Result: The administration page opens.

2. In the **System** section, select **Import**.

   Result: The **Import data** page opens.

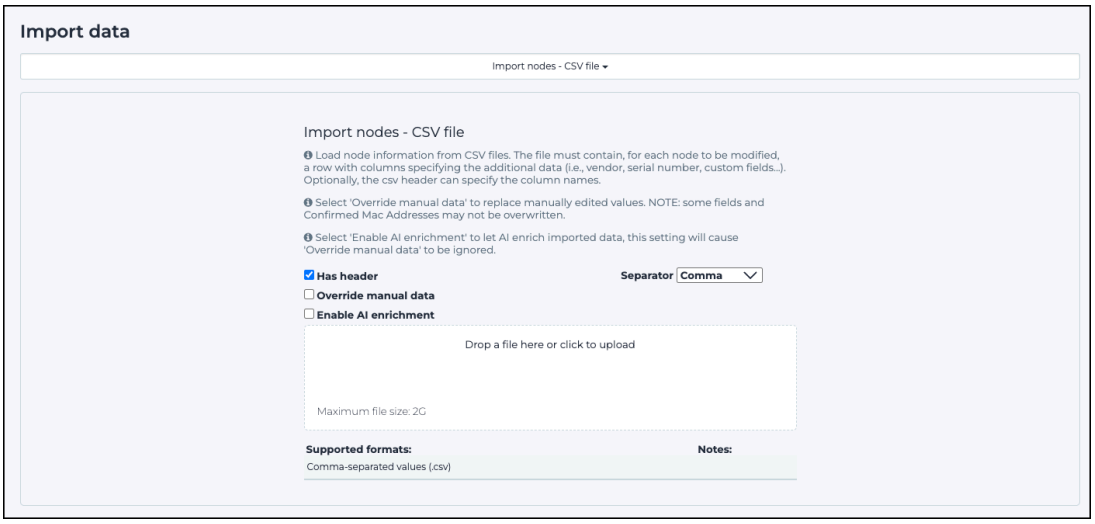
3. From the dropdown at the top of the page, select **Import nodes - CSV file**.

   Result: The **Import nodes - CSV file** section shows.

4. **Optional:** If the *CSV* file has headers in the first line of the file, select **Has header** to view the column titles.
5. **Optional:** If you want the file data to replace field values that were previously imported or manually overridden, select **Override manual data**.
6. From the **Separator** dropdown, select the correct option.
7. Choose a method to upload a file.

   Choose from:
   - Drag your file into the **Drop a file here or click to upload** field
   - Click in the **Drop a file here or click to upload** field

8. If you chose the second method, select the correct file to upload.



Note:

The supported format is *CSV*.

The maximum permitted file size is 2 *GB*.

9. Wait for the file to upload.

Results

The *CSV* file has been uploaded.

Import variables from a CSV file

*You can use the **Import** page to import variables from a comma-separated value (CSV) file.*

Procedure

1. In the top navigation bar, select 

   Result: The administration page opens.

2. In the **System** section, select **Import**.

   Result: The **Import data** page opens.

3. From the dropdown at the top of the page, select **Import variables - CSV file**.

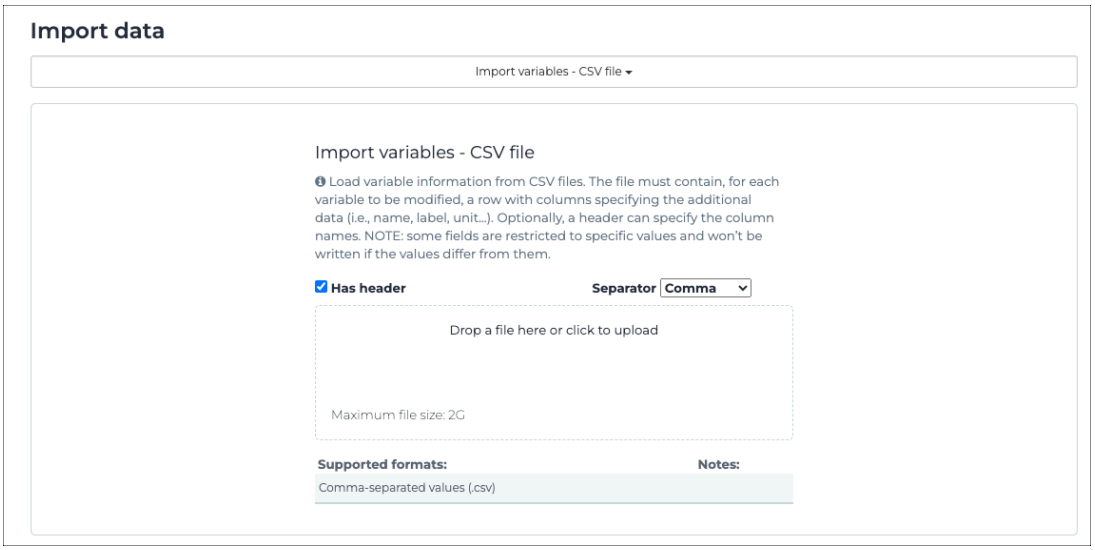
   Result: The **Import variables - CSV file** section shows.

4. **Optional:** If the *CSV* file has headers in the first line of the file, select **Has header** to view the column titles.
5. From the **Separator** dropdown, select the correct option.
6. Choose a method to upload a file.

   Choose from:

   - Drag your file into the **Drop a file here or click to upload** field
   - Click in the **Drop a file here or click to upload** field

7. If you chose the second method, select the correct file to upload.



Note:

The supported format is *CSV*.

The maximum permitted file size is 2 *GB*.

8. Wait for the file to upload.

Results

The *CSV* file has been uploaded.

Import asset types from a CSV file

*You can use the **Import** page to import asset types from a comma-separated value (CSV) file.*

Procedure

1. In the top navigation bar, select ⚙

   Result: The administration page opens.

2. In the **System** section, select **Import**.

   Result: The **Import data** page opens.

3. From the dropdown at the top of the page, select **Import asset types - CSV file**.

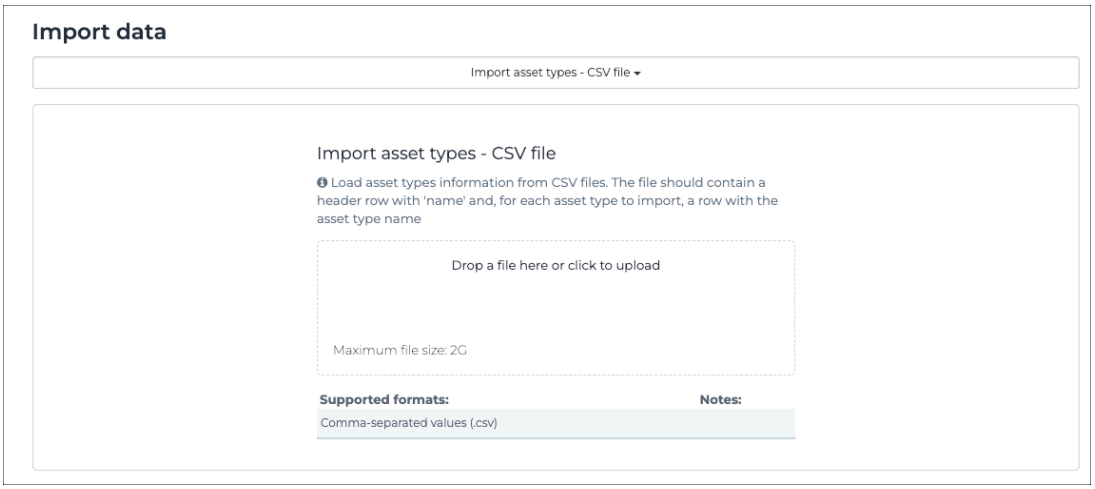
   Result: The **Import asset types - CSV file** section shows.

4. Choose a method to upload a file.

   Choose from:

   - Drag your file into the **Drop a file here or click to upload** field
   - Click in the **Drop a file here or click to upload** field

5. If you chose the second method, select the correct file to upload.



> **Note:**
> The supported format is *CSV*.
> The maximum permitted file size is 2 *GB*.

6. Wait for the file to upload.

Results

The *CSV* file has been uploaded.

Import a configuration or project file

*You can use the **Import** page to import configuration or project files.*

Procedure

1. In the top navigation bar, select ⚙

   Result: The administration page opens.

2. In the **System** section, select **Import**.

   Result: The **Import data** page opens.

3. From the dropdown at the top of the page, select **Import configuration / project file**.

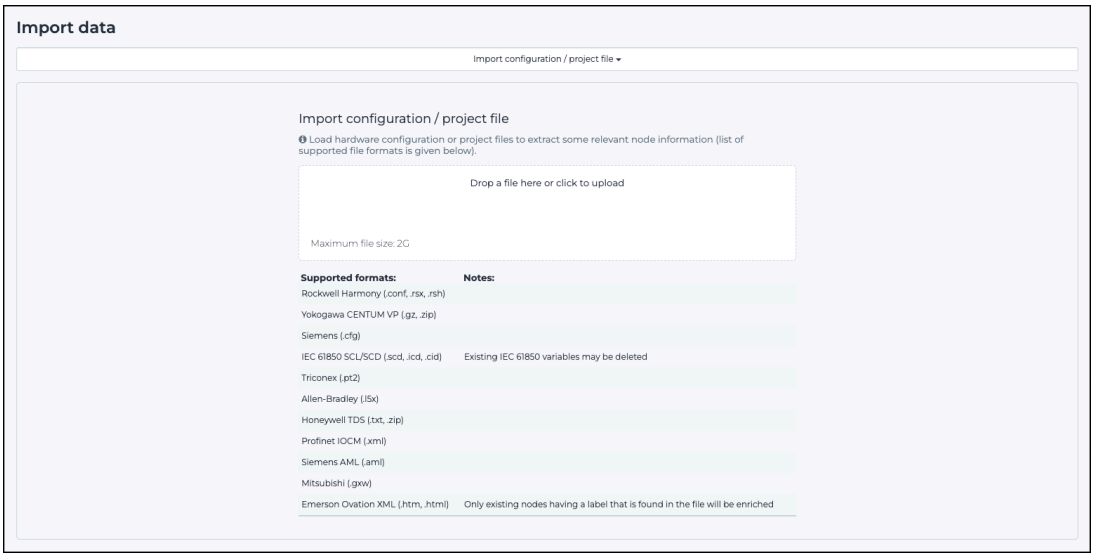
   Result: The **Import configuration / project file - CSV file** section shows.

4. Choose a method to upload a file.

   Choose from:

   - Drag your file into the **Drop a file here or click to upload** field
   - Click in the **Drop a file here or click to upload** field

5. If you chose the second method, select the correct file to upload.



   For more details on supported file types, see Supported file types (on page 225).

6. Wait for the file to upload.

Results

The file has been uploaded.

Supported file types

Supported configuration and project file types for import in Guardian.

Table 6. Supported file types

Supported file type	Notes
Allen-Bradley (.l5x)	Imported fields are: `ip`, `firmware version`, `product name`, `modules` (i.e. `port`, `address`, `product code`, `firmware version`, `product name`, `vendor`).
Emerson Ovation XML (.htm, .html)	Imported fields are: `Vendor`, `Product Name`, `Serial Number`, `OS`, `FW`, `Installed software`, `hotfixes`, `Ovation specific hotfixes`, `Hardware components`, and `redundancy partner information`.
Honeywell TDS (.txt, .zip)	Imported fields are: `label`, `vendor`
IEC 61850 SCL/SCD (.scd, .icd, .cid)	Imported file types are: `ip`, `VLAN`, `product name`, `asset type`, `vendor`, `AppID`, `data model` **Note**: Importing this file changes the knowledge of the IEC 61850 data model in use in the system, thus improving Guardian's ability to accurately extract variables. Existing variables related to this protocol are deleted to avoid inconsistencies with those extracted from traffic after the import.
Mitsubishi Electric (.gxw, .gx3)	Imported fields are: `ip`, `vendor`, `type`, `hardware components`
Profinet IOCM (.xml)	Imported fields are: `modules` (i.e. `slot`, `subslot`, `vendor`, `software release`, `hardware release`). **Note**: It might also extract variables.
Rockwell Harmony (.conf, .rsx, .rsh)	Imported fields are: `ip`, `product name`
Siemens AML (.aml)	Imported fields are: `ip`, `mac address`, `vendor`, `product name`, `label`, `modules` (i.e. `rack`, `slot`, `subslot`, `product code`, `module code`).
Siemens (.cfg)	Imported fields are: `ip`, `mac address`, `vendor`, `product name`, `label`, `modules` (i.e. `rack`, `slot`, `subslot`, `product code`, `module code`).
Triconex (.pt2)	Imported fields are: `ip`, `label`, `vendor`

Table 6. Supported file types (continued)

Supported file type	Notes
Yokogawa CENTUM VP (.gz, .zip)	Imported fields are: `ip`, `label`, `vendor`, `product name`, `modules` (i.e. `slots`, each with `vendor`, `product name`, `firmware version`). **Note**: You need to create a *ZIP* file that includes: `*.edf`, `/ETC/SystemRevisionInf.txt`, `project.atr`

Import a content pack

*You can use the **Import** page to import content packs.*

Procedure

1. In the top navigation bar, select ⚙

   Result: The administration page opens.

2. In the **System** section, select **Import**.

   Result: The **Import data** page opens.

3. From the dropdown at the top of the page, select **Import content pack**.

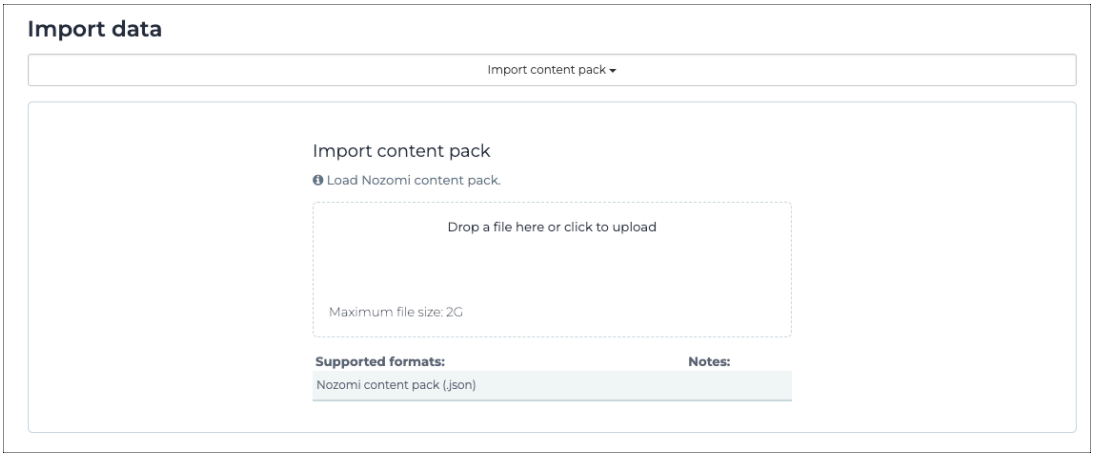
   Result: The **Import content pack** section shows.

4. Choose a method to upload a file.

   Choose from:

   - Drag your file into the **Drop a file here or click to upload** field
   - Click in the **Drop a file here or click to upload** field

5. If you chose the second method, select the correct file to upload.



> **Note:**
> The supported format is *JSON*.
> The maximum permitted file size is 2 *GB*.

6. Wait for the file to upload.

Results

The *JSON* file has been uploaded.

Import an Arc data archive

*You can use the **Import** page to import an Arc data archive.*

Procedure

1. In the top navigation bar, select ⚙

   Result: The administration page opens.

2. In the **System** section, select **Import**.

   Result: The **Import data** page opens.

3. From the dropdown at the top of the page, select **Import Arc data archive**.

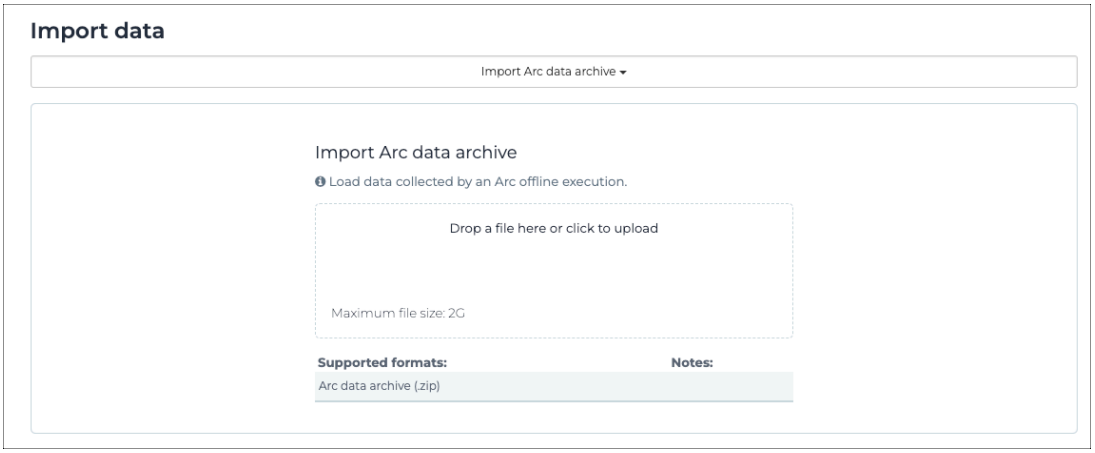
   Result: The **Import Arc data archive** section shows.

4. Choose a method to upload a file.

   Choose from:

   - Drag your file into the **Drop a file here or click to upload** field
   - Click in the **Drop a file here or click to upload** field

5. If you chose the second method, select the correct file to upload.



> **Note:**
> The supported format is *ZIP*.
> The maximum permitted file size is 2 *GB*.

6. Wait for the file to upload.

Results

The *ZIP* file has been uploaded.

Health

*The **Health** page lets you monitor the health of your sensors.*

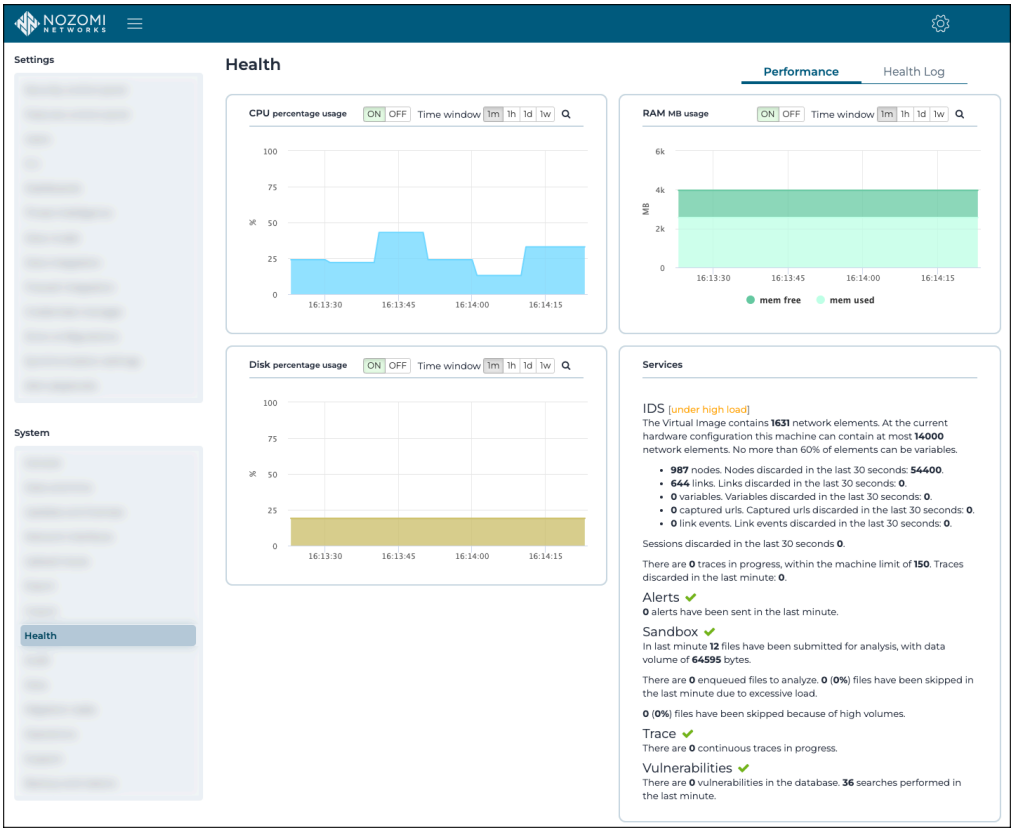


Figure 68. Health page

The Health page has these tabs:

- Performance (on page 230)
- Health log (on page 232)

Performance

*The **Performance** page has graphical sections that show percentage usage for central processing unit (CPU), random-access memory (RAM), and disk usage. It also has a **Services** section.*

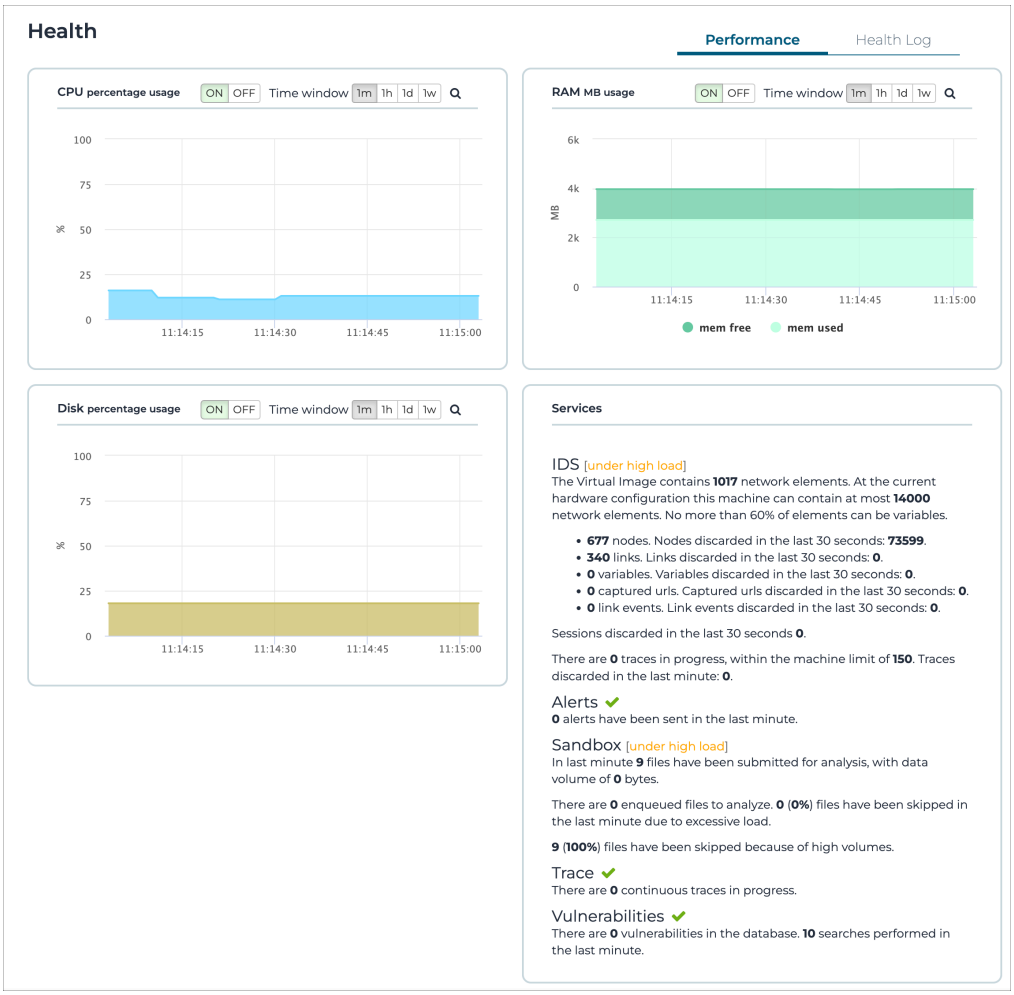


Figure 69. Performance page

Graphs

The page shows graph for:

- CPU percentage usage
- RAM MB usage
- Disk percentage usage

The vertical axis show percentage usage and the horizontal axis shows time.

These sections have an **ON/OFF** toggle and time controls that let you choose the timeframe of the window. You can choose between:

- One minute
- One hour
- One day
- One week

Services

This section shows information about:

- The *intrusion detection system (IDS)*
- Alerts
- Sandbox
- Trace
- Vulnerabilities

Health log

*The **Health log** page shows the details of performance issues that the sensor experiences. The information is for such as central processing unit (CPU), random-access memory (RAM), and disk usage, interface status, stale sensors, or generic high load.*

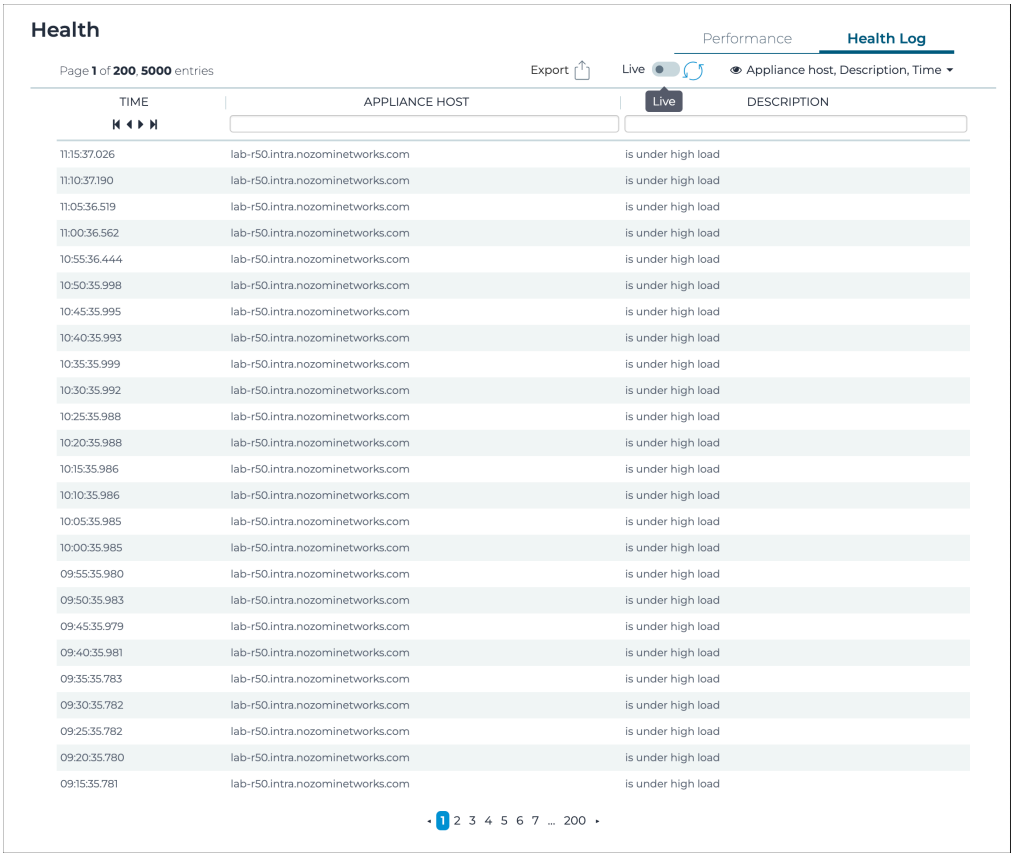


Figure 70. Health log page

The *central processing unit (CPU)* percentage usage, *random-access memory (RAM) MB* usage and disk percentage usage will show as:

- Good
- Average, or
- Poor

Stale or **unreachable** describes the status of the communication between Remote Collector, Guardian, and *CMC* (sync). It means that the last time the sensor communicated back to the upstream, the configured threshold was exceeded.

Audit

*The **Audit** page shows a list of all relevant user actions, from login/ logout to configuration operations, such as manually learning or deleting objects from the environment. This includes all recorded user actions based on the internet protocol (IP) address and username of the user who performed the action.*

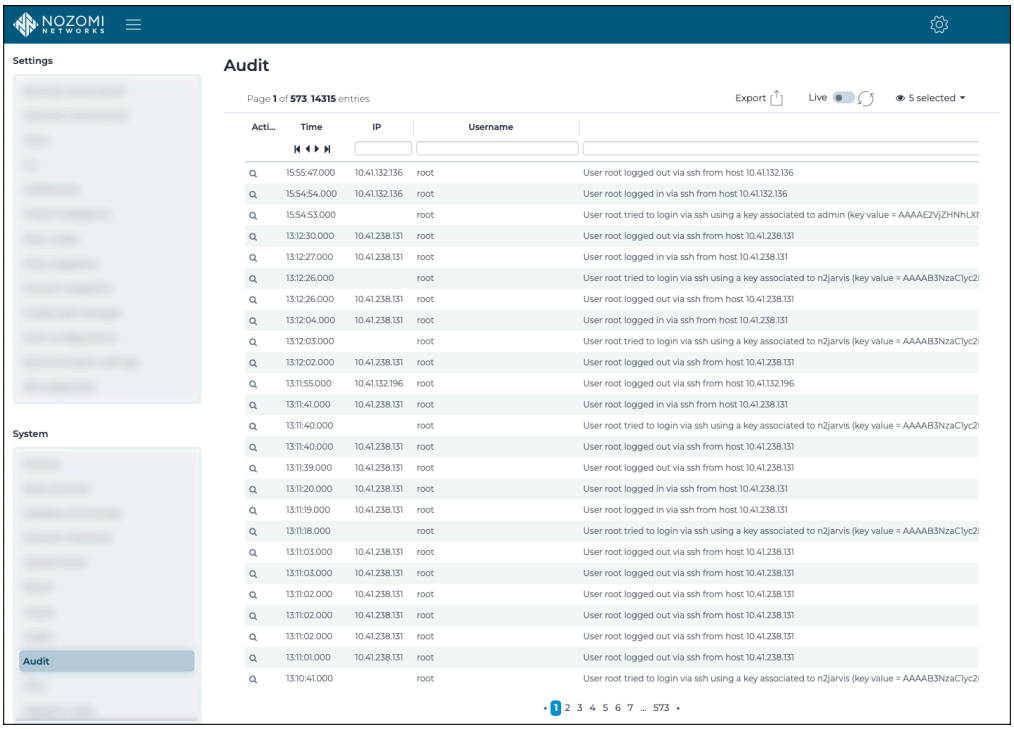


Figure 71. Audit page

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.

Data

*The **Data** page lets you selectively reset several kinds of data. You can select **All**, **Only data**, or **None**, depending on the type of user data that you want to reset.*

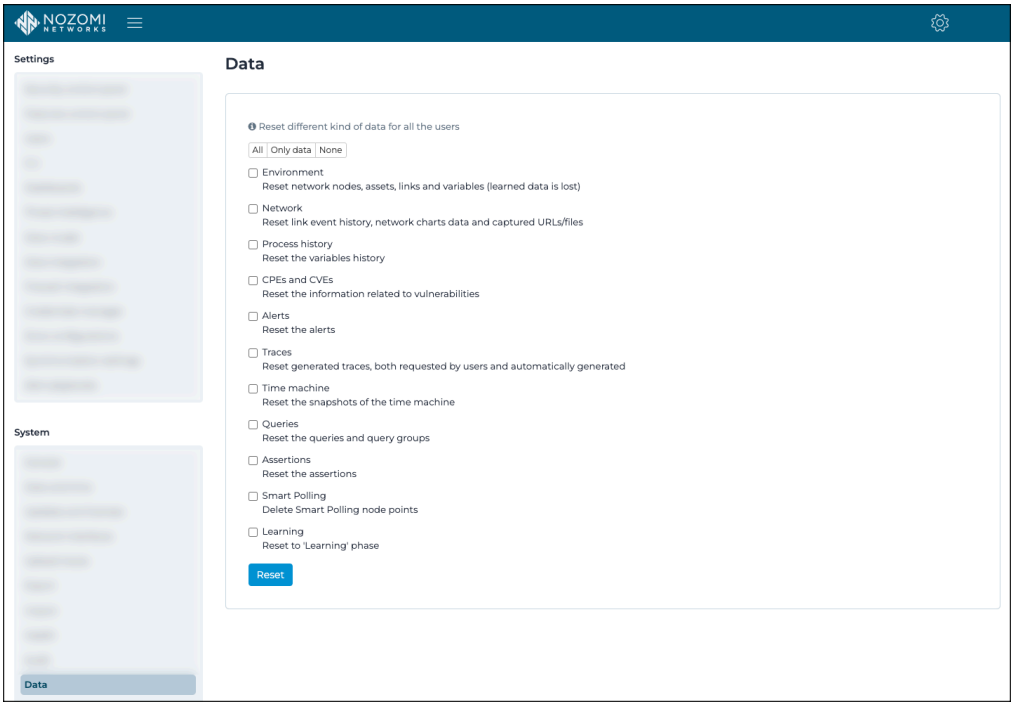


Figure 72. Data page

Environment

Reset network nodes, assets, links and variables (learned data is lost).

Network

Reset link event history, network charts data and captured *URL*s/files

Process history

Reset the variables history.

CPEs and CVEs

Reset the information related to vulnerabilities.

Alerts

Reset the alerts.

Traces

Reset user-requested and automatically-generated traces.

Time machine

Reset the snapshots of the time machine.

Queries
Reset the queries and query groups.

Assertions
Reset the assertions.

Smart Polling
Delete Smart Polling node points.

Learning
Reset to the **Learning** phase.

Migration tasks

*The **Migration tasks** page gives you access to automated tasks that help you migrate the system configuration and settings to newer standards.*

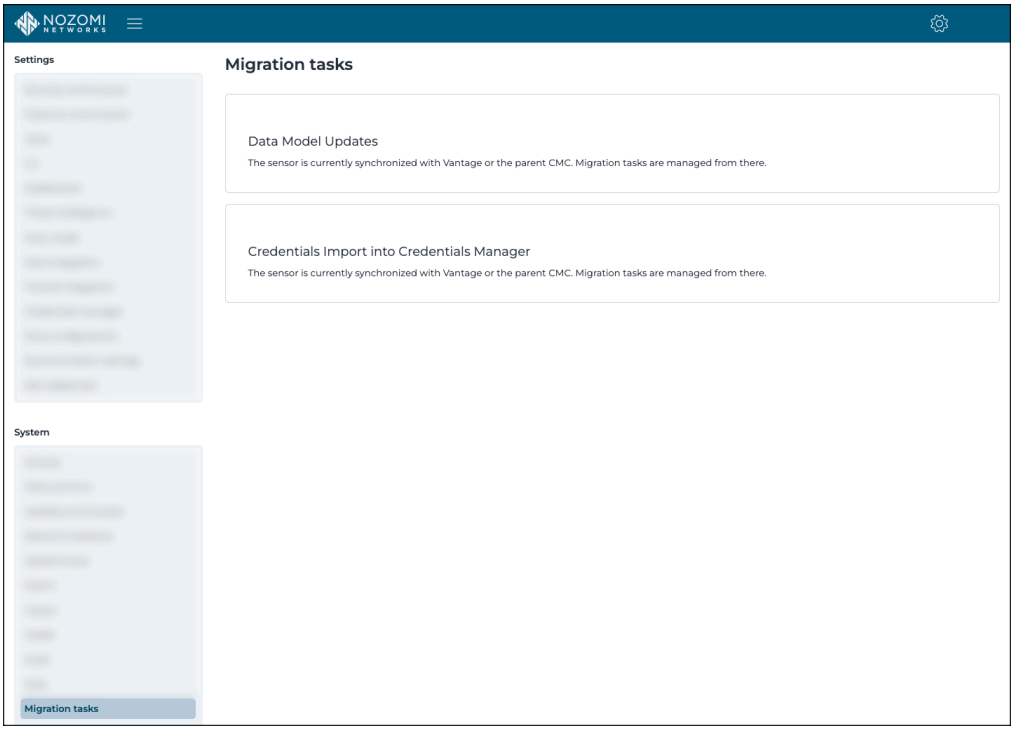


Figure 73. Migration tasks page

Migration tasks are a helper tool that can be used to apply specific changes to configuration, settings and data to make use of new features, or adapt to model changes. These tasks become available upon the installation of new versions of the software, and are described in the related release notes.

Migration tasks can only be executed from the top-level *CMC* of an installation. Guardians that are not connected to a *CMC* can also run migration tasks. If the installation is connected to Vantage, then Vantage can run the migration tasks.

Each migration task is presented separately, giving an overview of the changes that will be applied on each connected sensor. Tasks can be executed on individual Guardians, or on *CMC*s, When a *CMC* executes the migration tasks, all the sensors that are connected downstream, either directly or indirectly, receive the instruction to execute the task. You can select **Execute all** to apply the migration tasks globally. With this approach, migration tasks can be ignored, which disables the execution of the corresponding task on the chosen sensor, or sensors.

When the migration task is executed, a spinning wheel shows next to the sensors that are executing it. Since the execution of a task is an asynchronous process, it can take up to several minutes to complete. During this time, it is safe to leave the page, or disconnect from the Web *UI*. The **Migration tasks** page will report the result of each execution.

You can select **Hide permanently** to hide migration tasks. This hides the migration task and it cannot be executed again.

Operations

*The **Operations** page lets you reboot or shutdown the sensor, or manage the software updates for your sensor.*

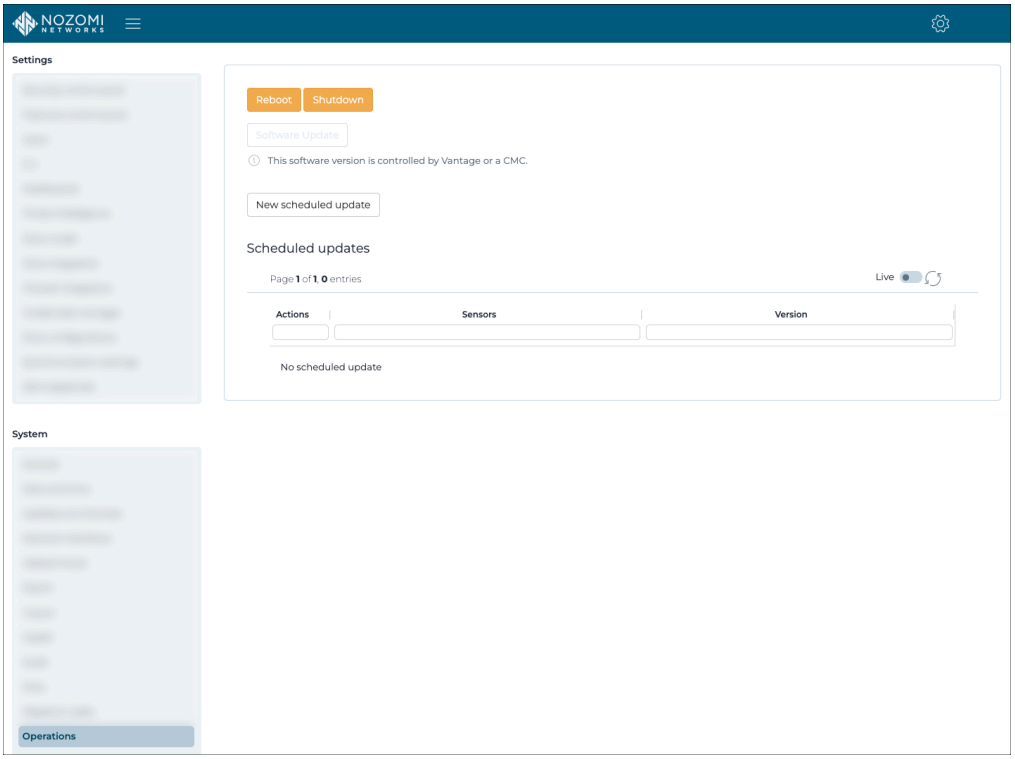


Figure 74. Operations page

Reboot

This lets you reboot the system.

Shutdown

This lets you shutdown the system.

Software update

This lets you update the software on the sensor. If there is an upstream Vantage or *CMC* that is in control of the sensor, this button will be inactive.

New scheduled update

This lets you set a schedule for the updates.

Reboot or shutdown the system from the web UI

You can reboot or shutdown the system from the web UI.

About this task

The reboot and shutdown commands are performed from the web UI. Alternatively, you can perform both commands from the shell console (inside an *SSH* session).

Procedure

1. In the top navigation bar, select ⚙

   Result: The administration page opens.

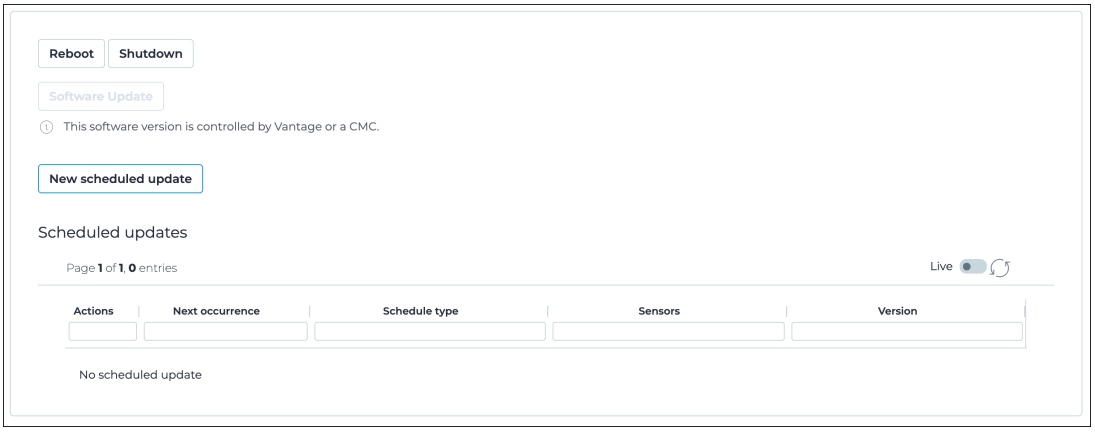
2. In the **System** section, select **Operations**.

   Result: The **Operations** page opens.

3. Choose the action that you want to do:

   Choose from:

   - In the top left corner, select **Reboot**
   - In the top left corner, select **Shutdown**



Results

The system reboots or shuts down.

Support

*The **Support** page section lets you generate a support archive that you can then send to Customer Support when you need technical support.*

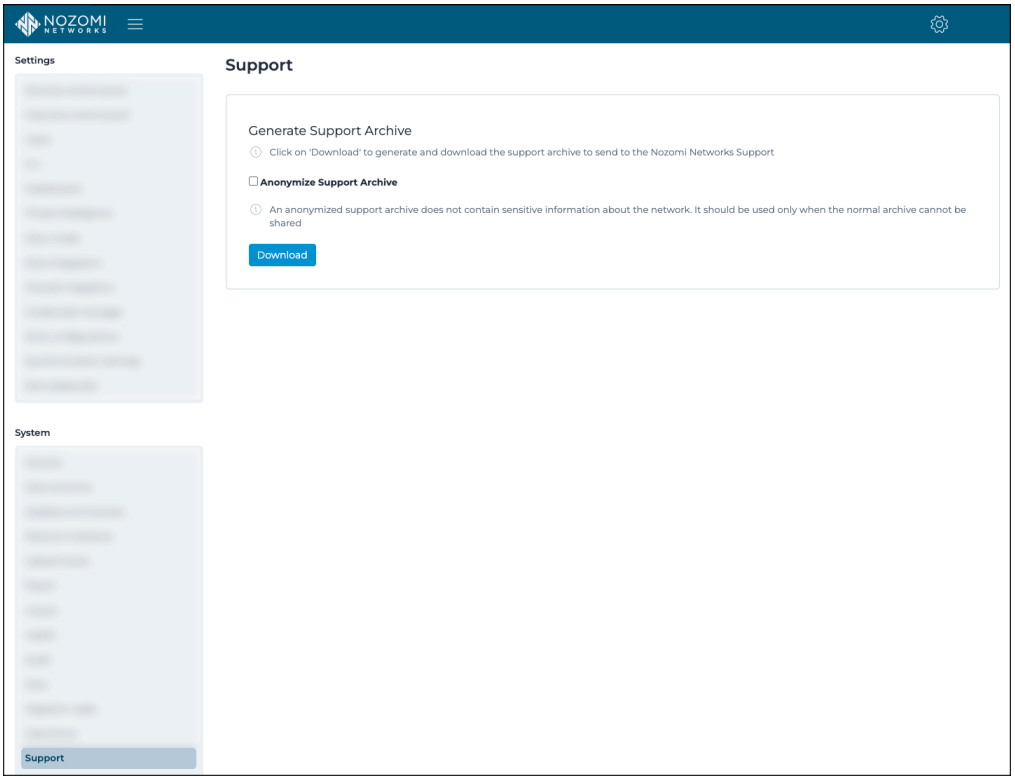


Figure 75. Support page

General

You can generate a support archive that can then be uploaded to a support case in the Customer Support Portal.

Anonymize Support Archive

This option removes sensitive network information from the generated archive. This option should only be used when a normal archive cannot be shared.

Backup and restore

*The **Backup and restore** page lets you generate and schedule backup archives as well as restore the software from a backup.*

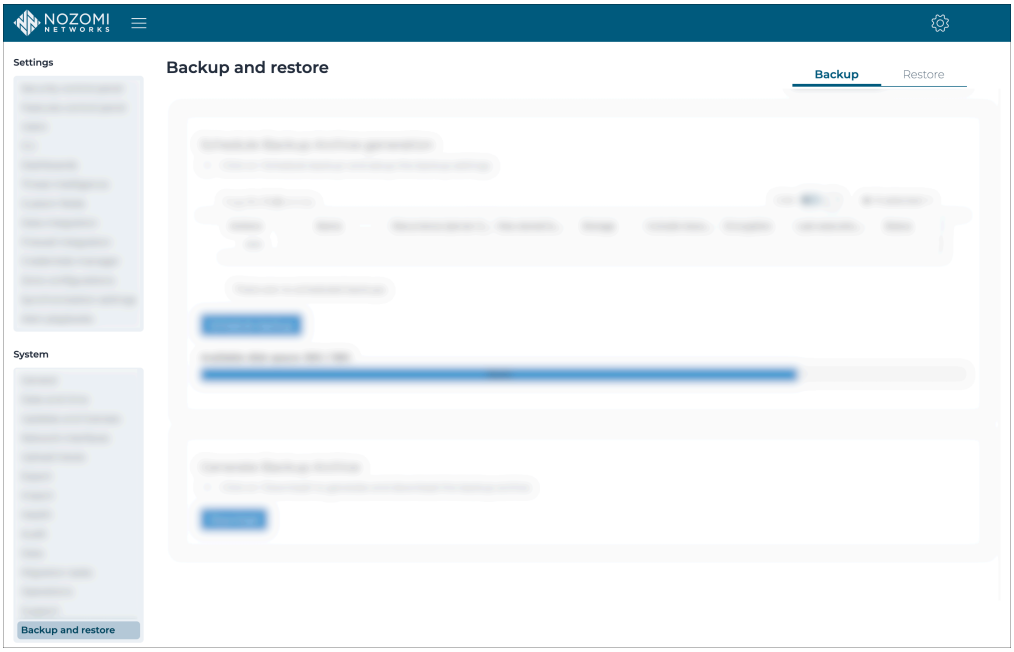


Figure 76. Backup and restore page

The **Backup and restore** page has these tabs:

- Backup
- Restore

Important:

Backups do not contain data that relates to vulnerabilities or schedules.

Therefore, restored machines will preserve the old tables with the previous:

- Vulnerabilities
- Report schedules
- Backup and restore schedules

Backup

*The **Backup** page lets you generate and schedule backup archives as well as restore the software from a backup.*

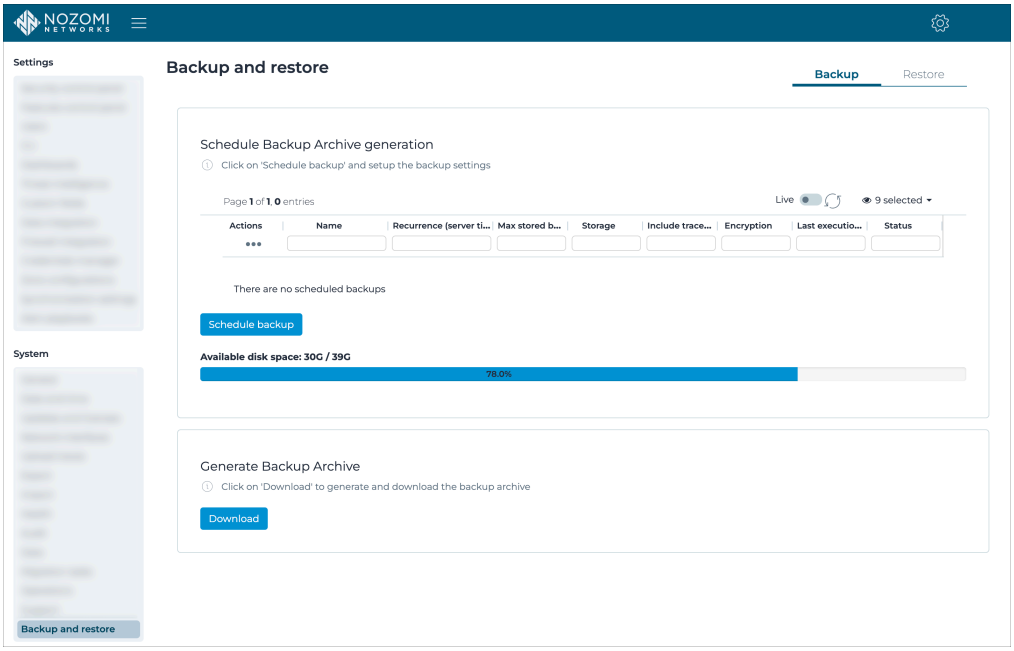


Figure 77. Backup page

Schedule Backup Archive generation

This lets you schedule the generation of a backup archive. For more details, see the **Guardian and CMC Maintenance Guide**.

Generate Backup Archive

This section lets you generate a backup archive. For more details, see the **Guardian and CMC Maintenance Guide**.

The supported format is **nozomi_backup**.

Restore

*The **Restore** page lets you restore a previous backup, and to schedule the restoration of a backup archive.*

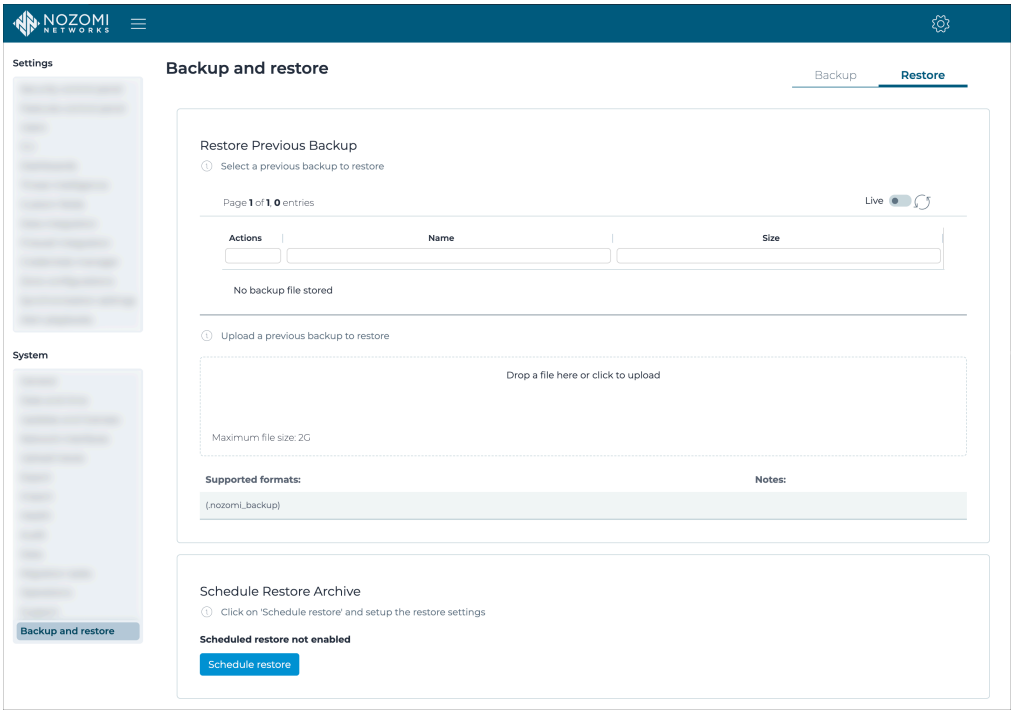


Figure 78. Restore page

Restore Previous Backup

This section lets you restore a previously created backup. For more details, see the **Guardian and CMC Maintenance Guide**.

Schedule Restore Archive

This lets you schedule the restoration of a backup archive. For more details, see the **Guardian and CMC Maintenance Guide**.



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.