NOZOMI NETWORKS



Central Management Console
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**
- CMC overview 7
- Functionalities 8
- Alerts 10
- Updates 11

Chapter 2. Administration **13**
- Administration page 15
- Settings 17
  - Security control panel 17
  - Features 23
  - Users management 24
  - CLI 61
  - Dashboards 62
  - Threat Intelligence 70
  - Custom fields 76
  - Data integration 77
  - Credentials manager 118
  - Zone configurations 121
  - Synchronization settings 125
  - Alert playbooks 153
  - Update service contents 159
- System 165
  - General 165
  - Date and time 166
  - Updates and licenses 167
  - Export 171
  - Import 175
  - Health 188
  - Audit 192
  - Data 193
  - Migration tasks 195
  - Operations 197
  - Support 199
  - Backup and restore 200

Glossary 203



Chapter 1. Introduction



CMC overview

The Central Management Console (CMC) aggregates information from thousands of sites and assets to deliver instant awareness of your networks and their activity patterns.

Overview

The *Central Management Console (CMC)* makes it easy to have visibility of all your Nozomi Networks sensors. The *CMC* aggregates information from thousands of sites and assets to deliver instant awareness of your *operational technology (OT)*/*Internet of Things (IoT)* networks and their activity patterns, enabling users to quickly troubleshoot issues and accelerate incident response across your network.

CMC can be deployed:

- On-premises
- In a public cloud provider such as *Amazon Web Services (AWS)* or Microsoft Azure
- At the edge on physical or virtual appliances

Functionalities

The *CMC* has been designed to support complex deployments that cannot be addressed with a single sensor. The *CMC* shares the user experience of the sensor. In comparison to the sensor, some functionalities have been added, and some have been removed. The functionalities that have been added let you manage *hundreds* of sensors. The functionalities that have been removed increase efficiency for real-world, geographic deployments of the Nozomi Networks Solution architectures.

Sensors

The *CMC* **Sensors** page lets you see and manage all the connected sensors. A graphical representation of all the hierarchical structure of the connected sensors and the sensor Map is presented to allow a quick health check on a user-provided raster map.

Once sensors are connected, they are periodically synchronized with the *CMC*. In particular, the Environment of each sensor is merged into a global Environment and Alerts are received for a centralized overview of the system. Of course, Alerts can also be forwarded to a *security information and event management (SIEM)* directly from the *CMC*, thus enabling a simpler decoupling of components in the overall architecture.

Synchronization

To synchronize data, the sensors must be running the same major *Nozomi Networks Operating System (N2OS)* release, or one of the two prior major ones. For example, if the *CMC* is running version 23.0.x (the major release is 23.0), sensors can synchronize if running one of the following versions: 23.0.x, 22.5.x or 22.0.x.

Firmware update is also simpler with a *CMC*. Once the new firmware is deployed to it, all connected sensors can be automatically updated.

Functionalities

A description of the differences in functionalities between the Central Management Console (CMC) and sensors. The actions shown below are only applicable for CMCs that are in **All-In-One** *mode, not* **Multicontext** *mode.*

Node actions

Table 1. Node actions

Action	Sensor	CMC
Configure node	✓	✗
Show alerts	✓	✓
Show requested traces	✓	✗
Request a trace	✓	✗
Manage Learning	✓	✓
Navigate	✓	✓

Network links

Table 2. Network links

Action	Sensor	CMC
Configure node	✓	✗
Show alerts	✓	✓
Show requested traces	✓	✗
Request a trace	✓	✗
Show events	✓	✗
Show captured urls	✓	✗
Manage Learning	✓	✓

Table 2. Network links (continued)

Action	Sensor	CMC
Navigate	✓	✓

Process variables

Table 3. Process variables

Action	Sensor	CMC
Configure variable	✓	✗
Variable details	✓	✓
Add to favourites	✓	✓
Navigate	✓	✓

Alerts

A description of alerts management in a Central Management Console (CMC).

Alerts management in the *CMC* is similar to alerts management in a sensor. However, in the *CMC*, you can have all the alerts, from all the sensors, in one centralized place.

In a sensor, you can create a query, and therefore an assertion, that includes all the nodes, or links etc., for your complete infrastructure. You can create a *Global Assertion*. This is one or more groups of assertions that can be propagated to all the sensors. The *CMC* has control of these assertions, and sensors cannot edit or delete them.

It is possible to configure the *CMC* to forward alerts to a *SIEM* without the need to configure each sensor. For more details, see **Data integration**.

Updates

A description of the release update and rollback operations when you have a Central Management Console (CMC) and one or more sensors.

The Nozomi Networks update bundle applies to both the *CMC* and Guardian, except in the case of the Docker installation. It works for all physical and virtual sensors.

Once a sensor is connected to a *CMC*, the *CMC* will control all the updates. The software bundle is propagated from the *CMC* and, once the sensor receives the bundle, you can do the update manually, or automatically. You can configure this behavior in the **General settings** section of the **Synchronization settings** page.

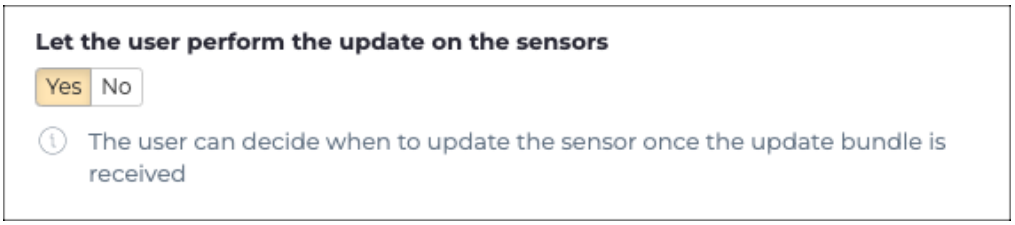


Figure 1. Let user perform updates section

If the *CMC* is configured to allow manual updates, the sensor's status bar displays a message that notifies you as soon as the sensor receives the update bundle.



Figure 2. Update available notification

The update process from the *CMC* can proceed as described in **Software update and rollback**. After the Central Management Console is updated, each sensor will receive the new software update.

If an error occurs during the update procedure, a message shows next to the related sensor's version number on the **Sensors** page.

To Rollback, first rollback the Central Management Console, and then proceed to rollback all the sensors as described in **Software update and rollback**.



Chapter 2. Administration



Administration page

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

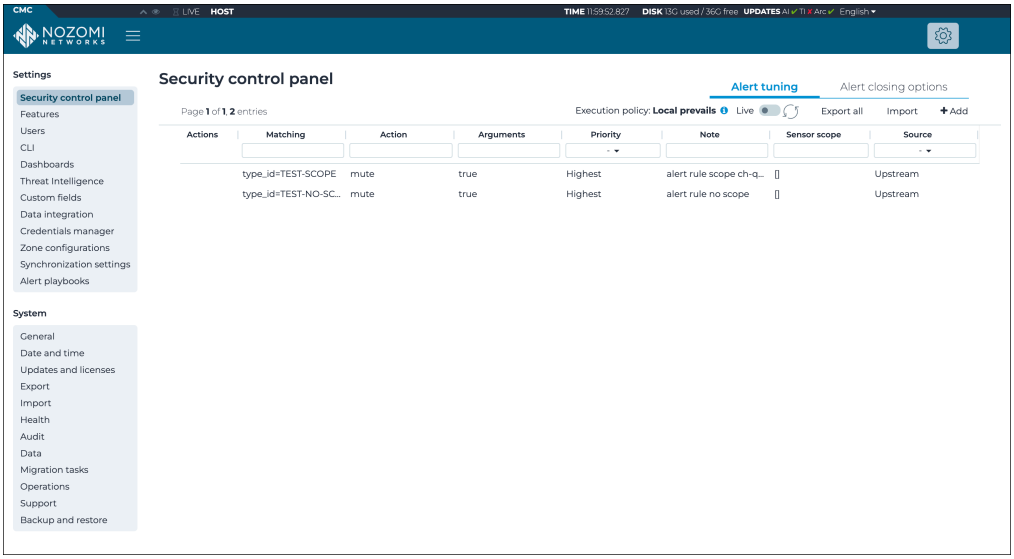


Figure 3. Administration page

Settings

The **Settings** section has these pages:

- Security control panel (on page 17)
- Features (on page 23)
- Users management (on page 24)
- CLI (on page 61)
- Dashboards (on page 62)
- Threat Intelligence (on page 70)
- Custom fields (on page 76)
- Data integration (on page 78)
- Credentials manager (on page 118)
- Zone configurations (on page 121)
- Synchronization settings (on page 125)
- Alert playbooks (on page 154)

System

The **System** section has these pages:

- General (on page 165)
- Date and time (on page 166)
- Updates and licenses (on page 167)
- Export (on page 171)
- Import (on page 175)
- Health (on page 188)
- Audit (on page 192)
- Data (on page 193)
- Migration tasks (on page 195)
- Operations (on page 197)
- Support (on page 199)
- Backup and restore (on page 200)

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.*

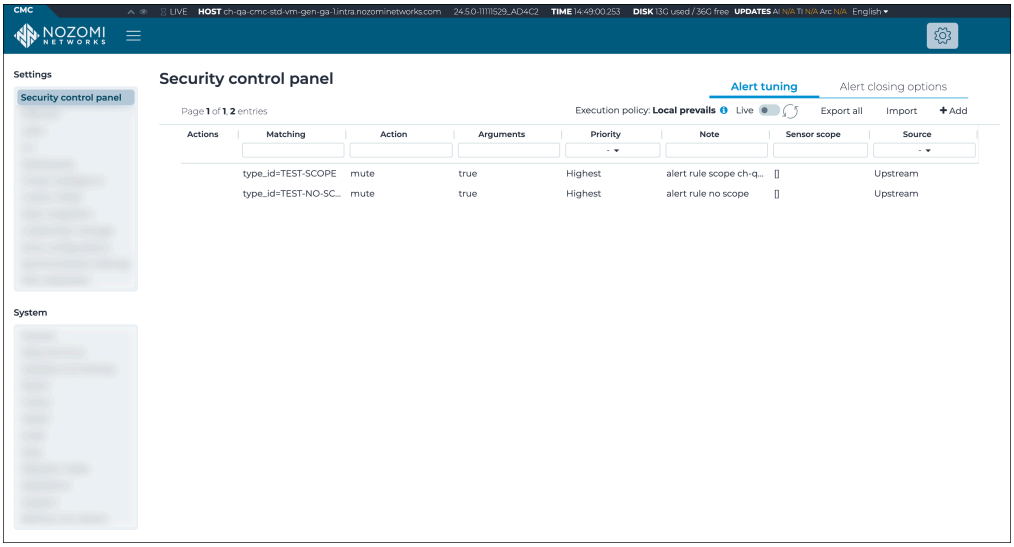


Figure 4. Security control panel page

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

- Alert tuning (on page 18)
- Alert closing options (on page 22)

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 5. 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 *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 all the items from 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.

Configure the alert tuning settings

Learn how to configure alert tuning settings to customize alert rules, actions, and sensor scopes for optimal system monitoring.

Procedure

1. In the top navigation bar, select 

   Result: The administration page opens.

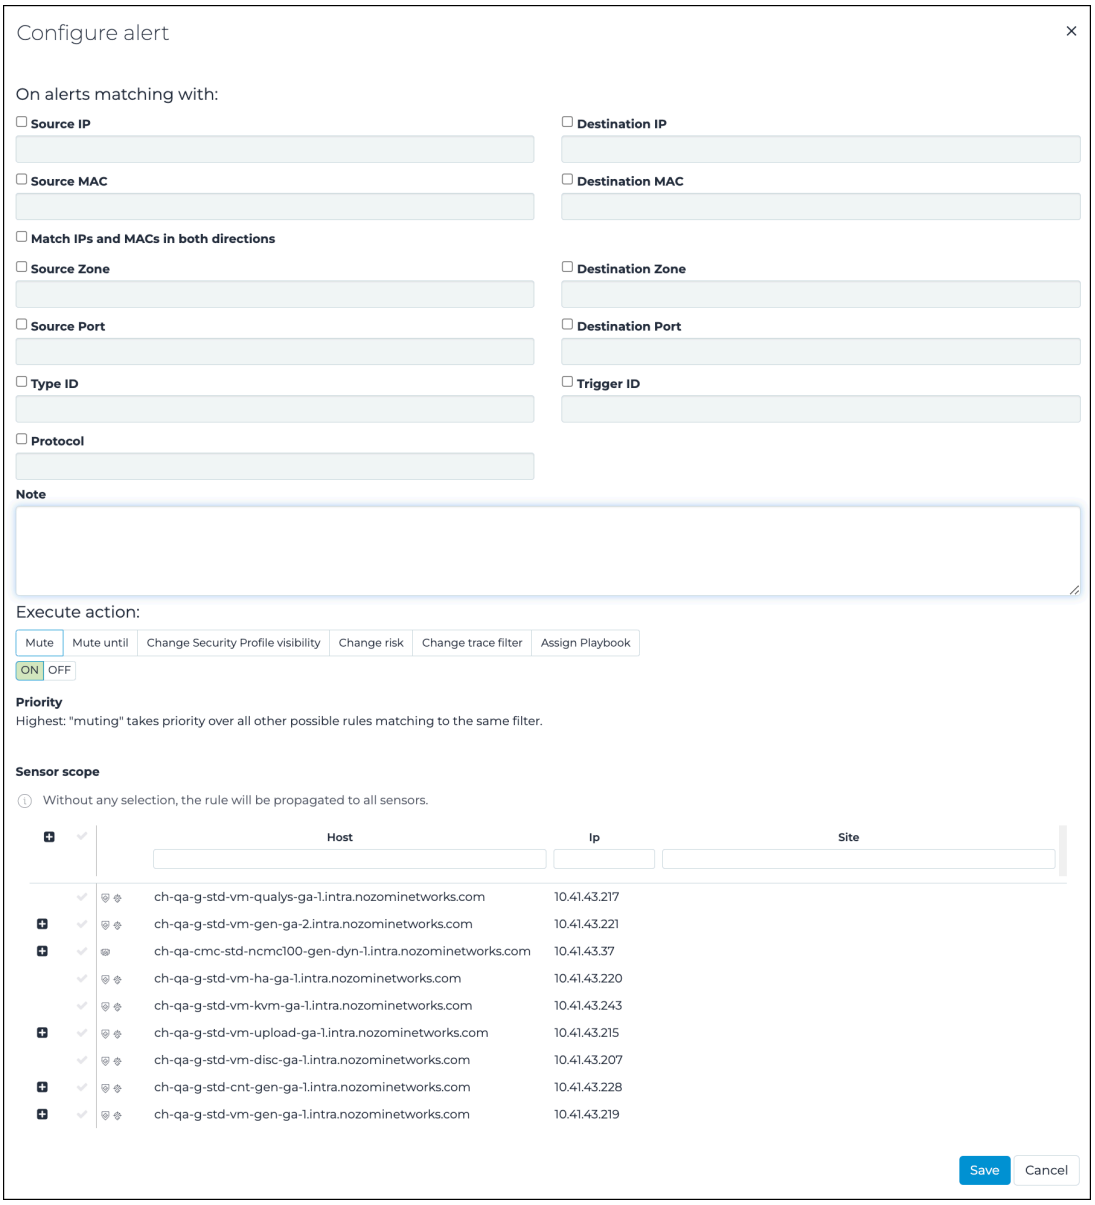
2. In the **Settings** section, select **Security control panel**.

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

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

   Result: A dialog shows.

4. In the **On alerts matching with:** section, configure the settings as necessary.



5. **Optional:** If necessary, add a note to describe the alert rule.
6. In the **Execute action:** section, select the applicable option.
7. In the **Sensor scope** section, select all the sensors that you want to propagate the alert rule to.
8. Select **Save**.

Results

The alert rule has been configured.

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**.

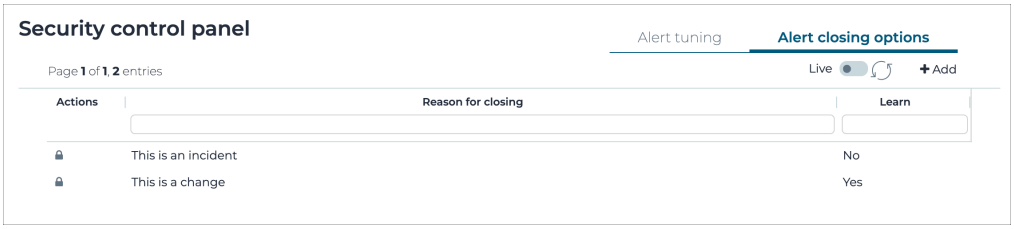


Figure 6. Alert closing options page

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.

Features

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

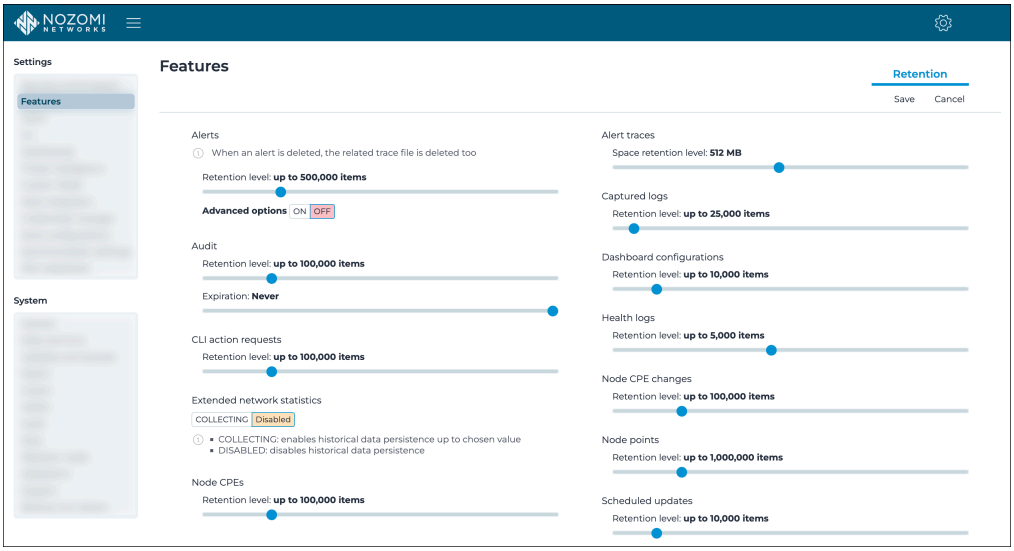


Figure 7. Features page

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.

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.*

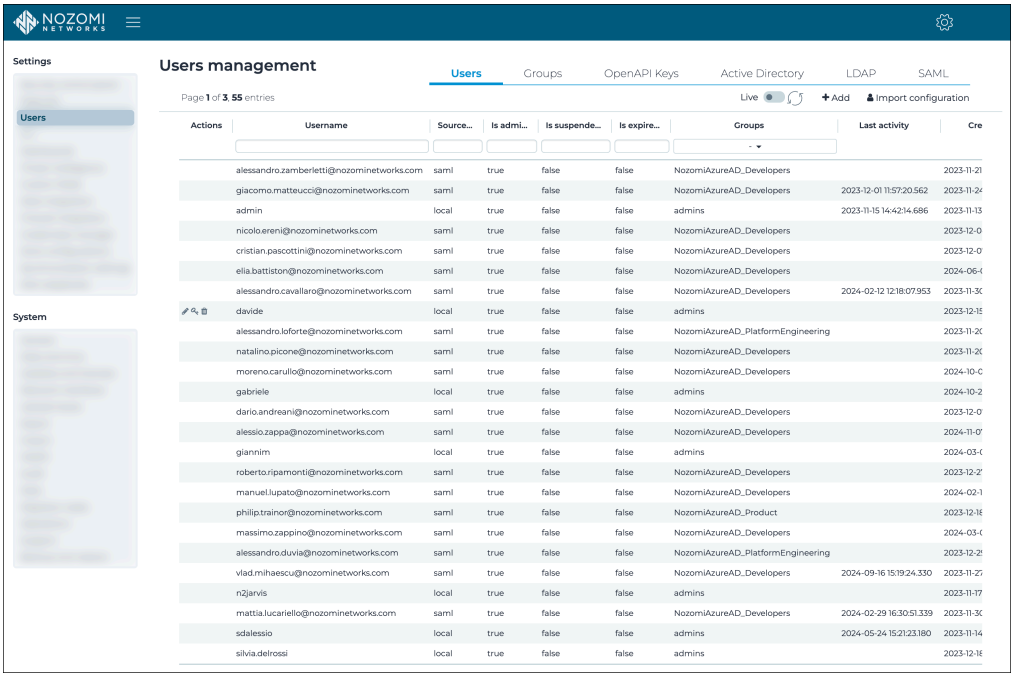


Figure 8. Users management page

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 24)
- Active Directory users (on page 25)
- LDAP users (on page 25)
- SAML users (on page 25)

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 47) 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 50) 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 55), or to Import SAML users from a CSV file (on page 30).

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.*

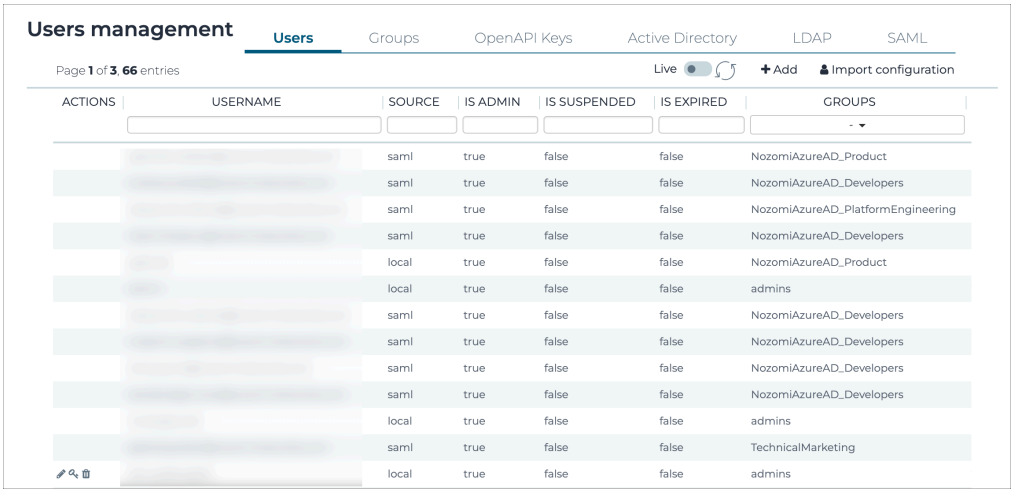


Figure 9. Users page

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.

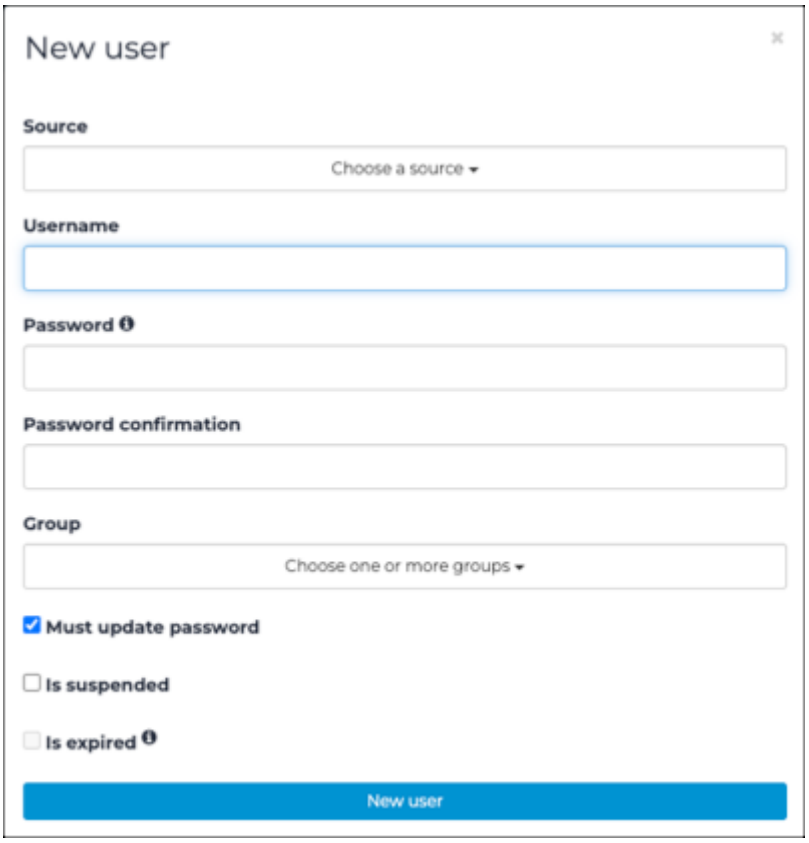
3. In the top right section, select **Users**.

   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**.

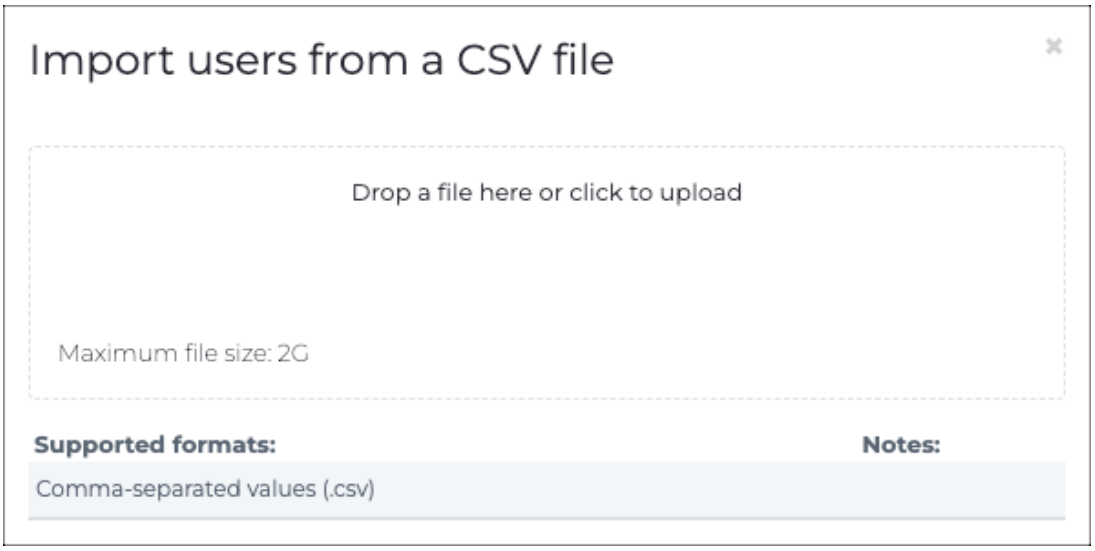
   Result: A dialog shows.

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.

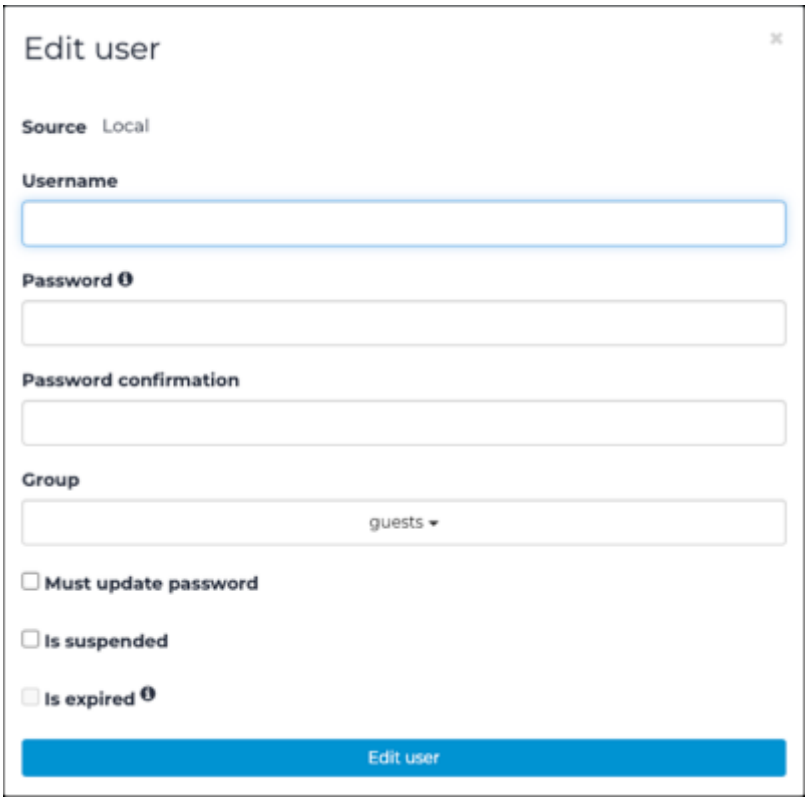
3. In the top right section, select **Users**.

   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.

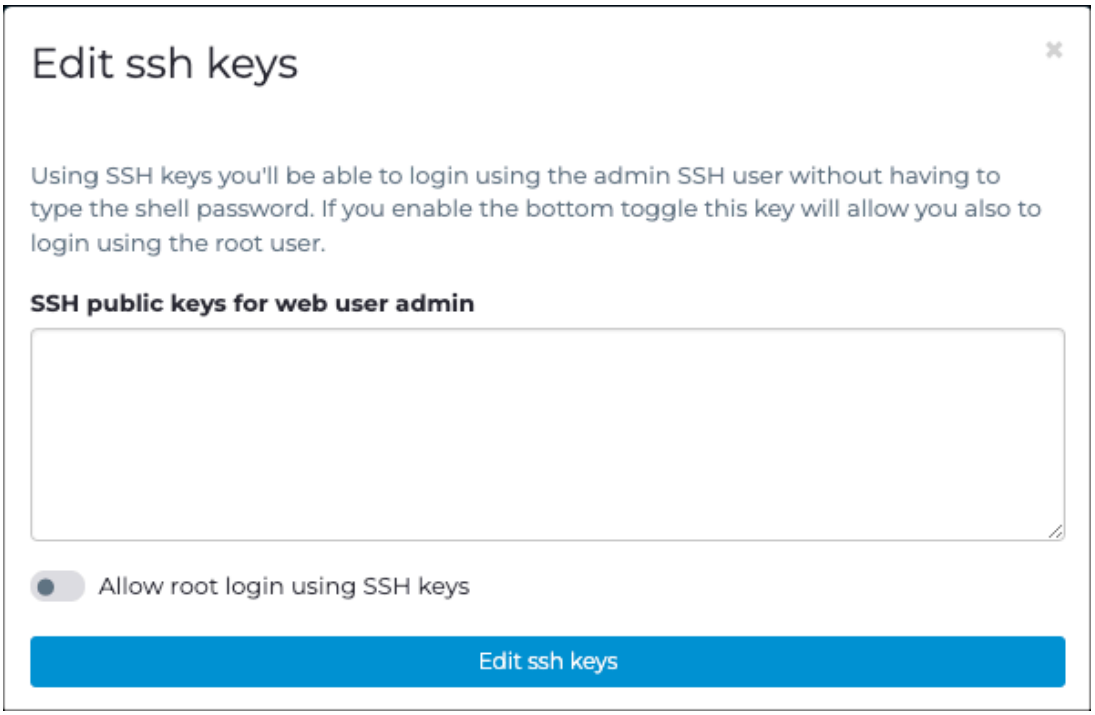
3. In the top right section, select **Users**.

   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 *secure shell (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.*

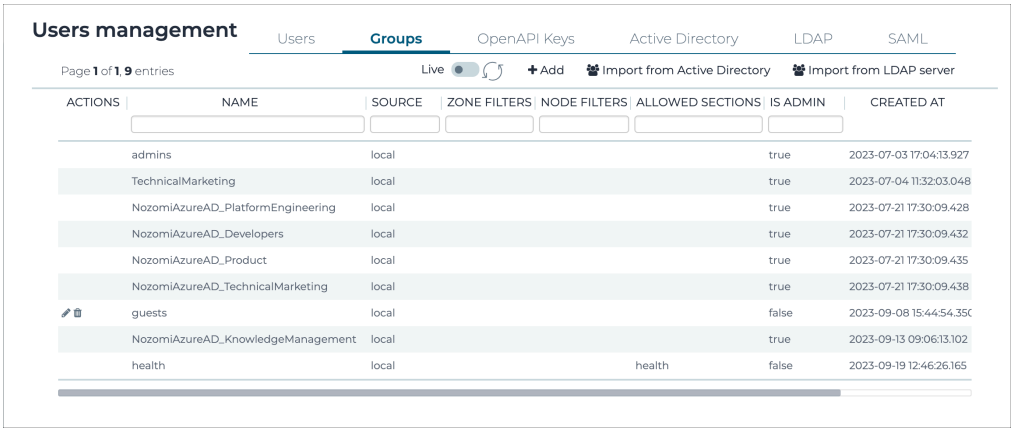


Figure 10. Groups page

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.

In a *CMC* in **Multicontext** mode, if a user belongs to multiple groups, where at least one of them is non-admin, and the non-admin groups have restrictions on sensors, nodes or zones, the most restrictive filter is applied to the user.

Two group types have predefined authorization policies:

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

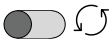
When a group is neither **Administrators** nor **Authentication Only**, the allowed features (sections) can be enabled/disabled individually.

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.

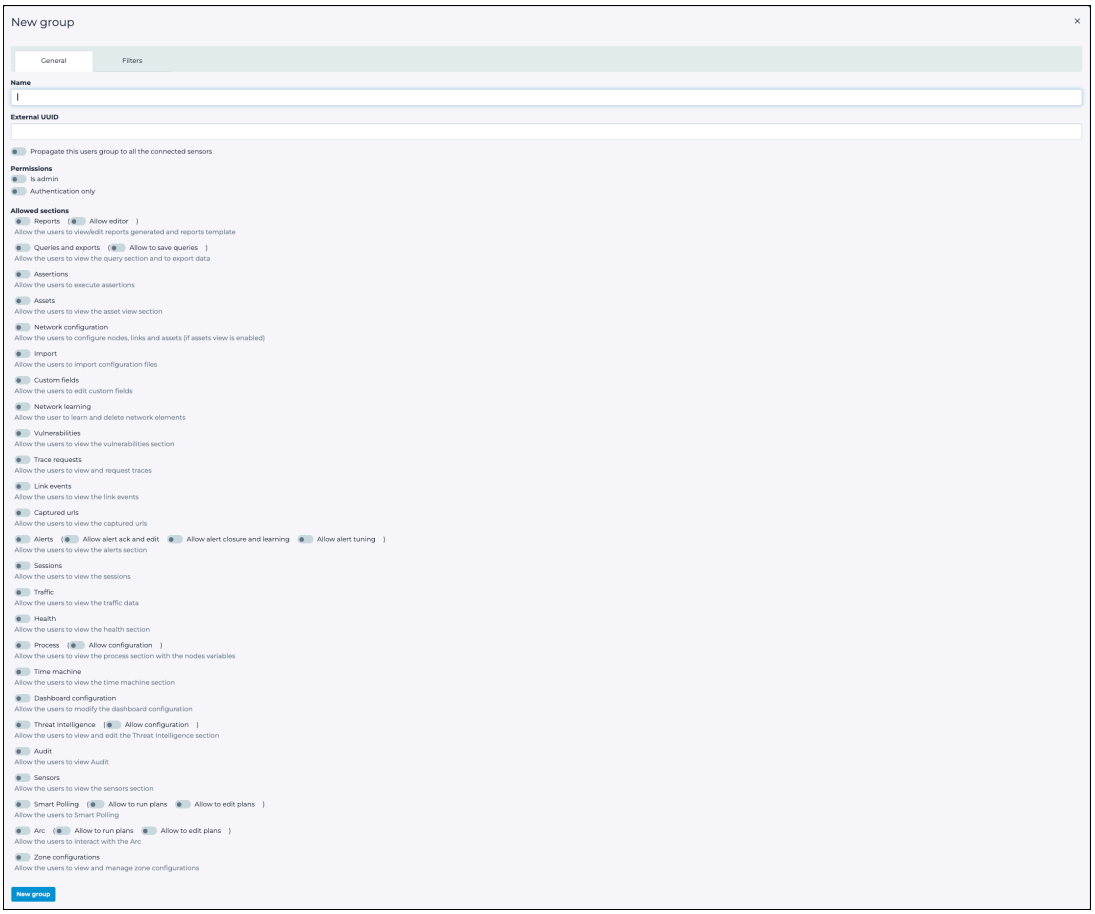
3. In the top right section, select **Groups**.

   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

    d. In the **Allowed sensors** field, select sensors that the users can access and see data coming from.

> **Note:**
> This feature is only available if:
> - The *CMC* is in multicontext mode
> - The **is admin** group permission is disabled

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.

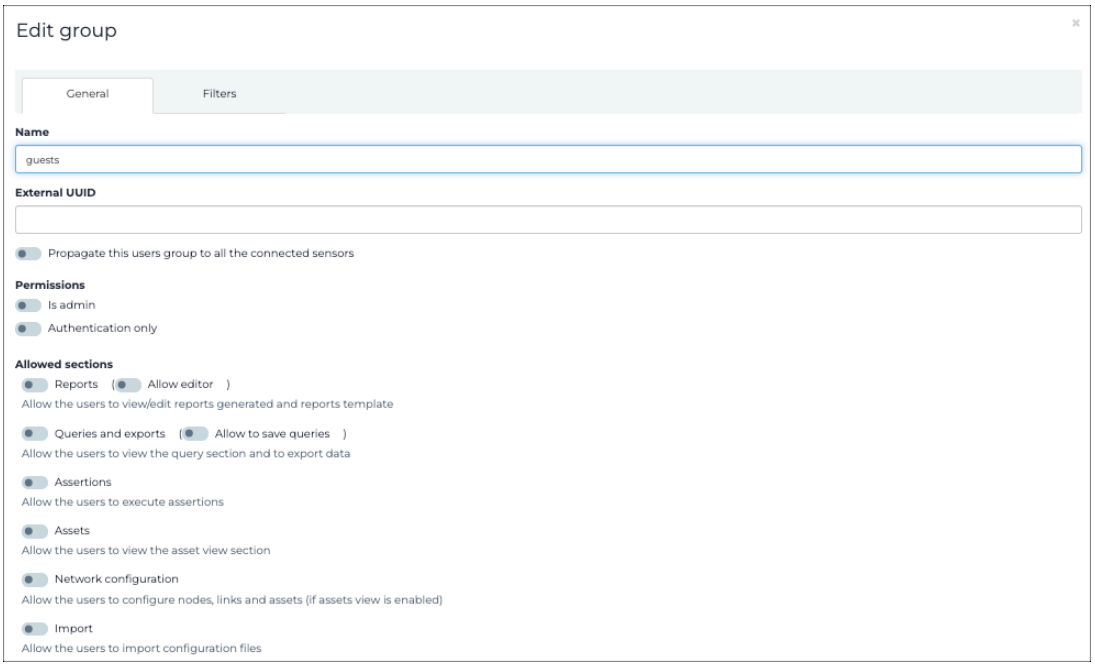
3. In the top right section, select **Groups**.

   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.*

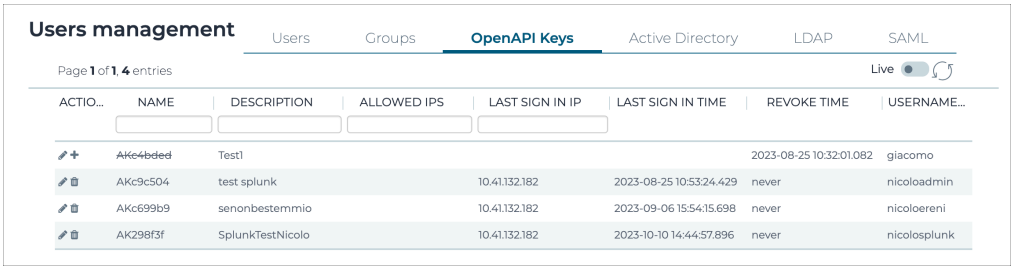


Figure 11. OpenAPI keys page

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 **CMC 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.

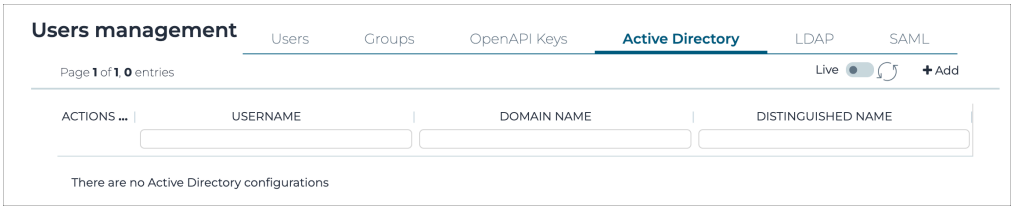


Figure 12. Active Directory page

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 CMC.

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**.

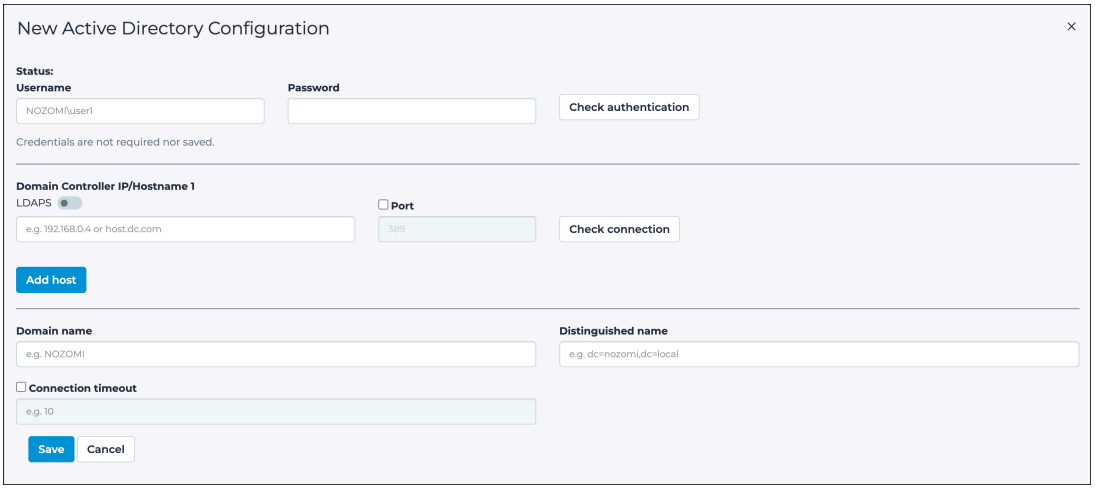
   Result: The **Active Directory** page opens.

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 *secure sockets layer (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.*

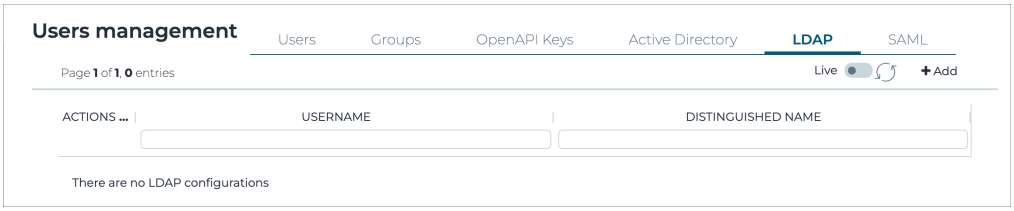


Figure 13. LDAP page

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.

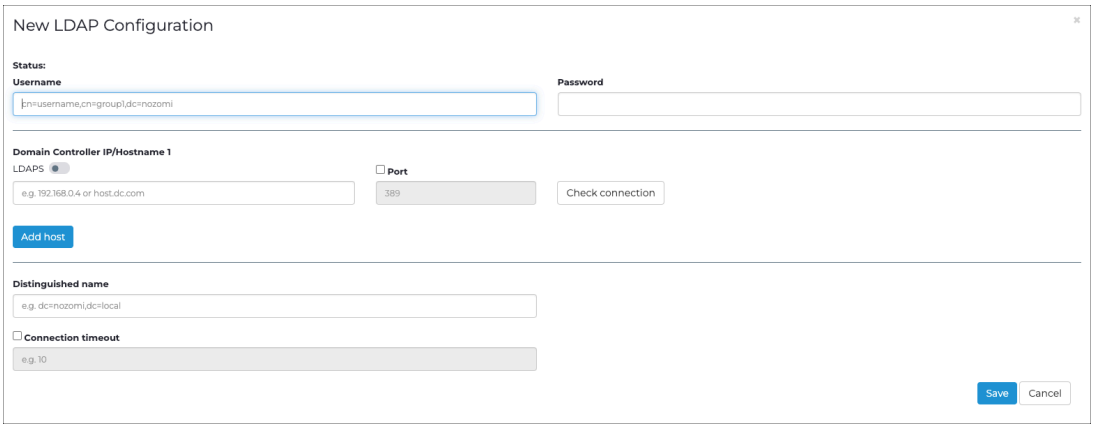
3. In the top right section, select **LDAP**.

   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.

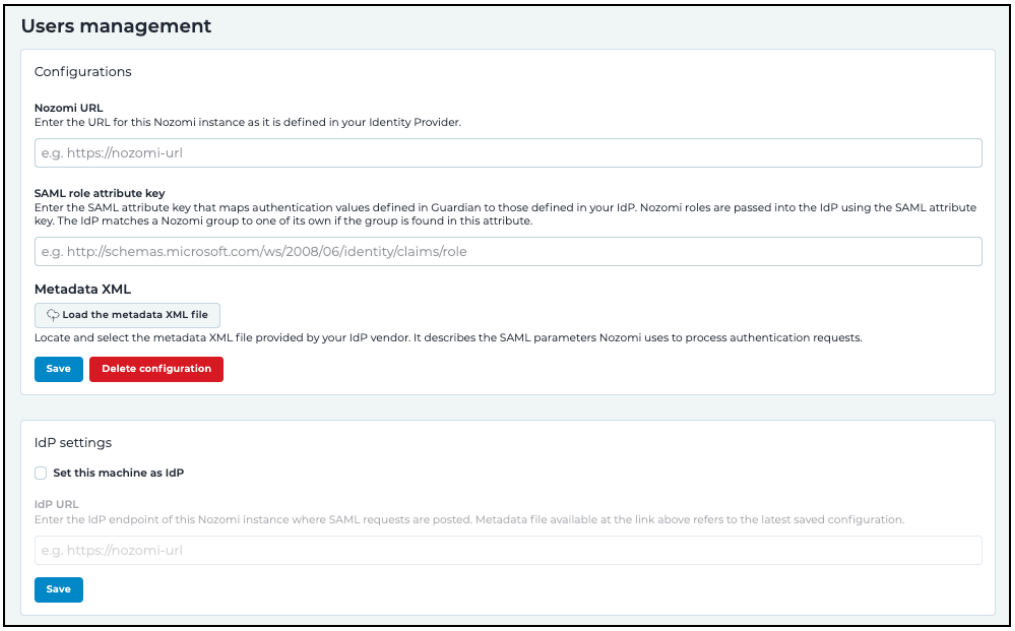


Figure 14. SAML page

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 *IdPs*, 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 CMC 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.

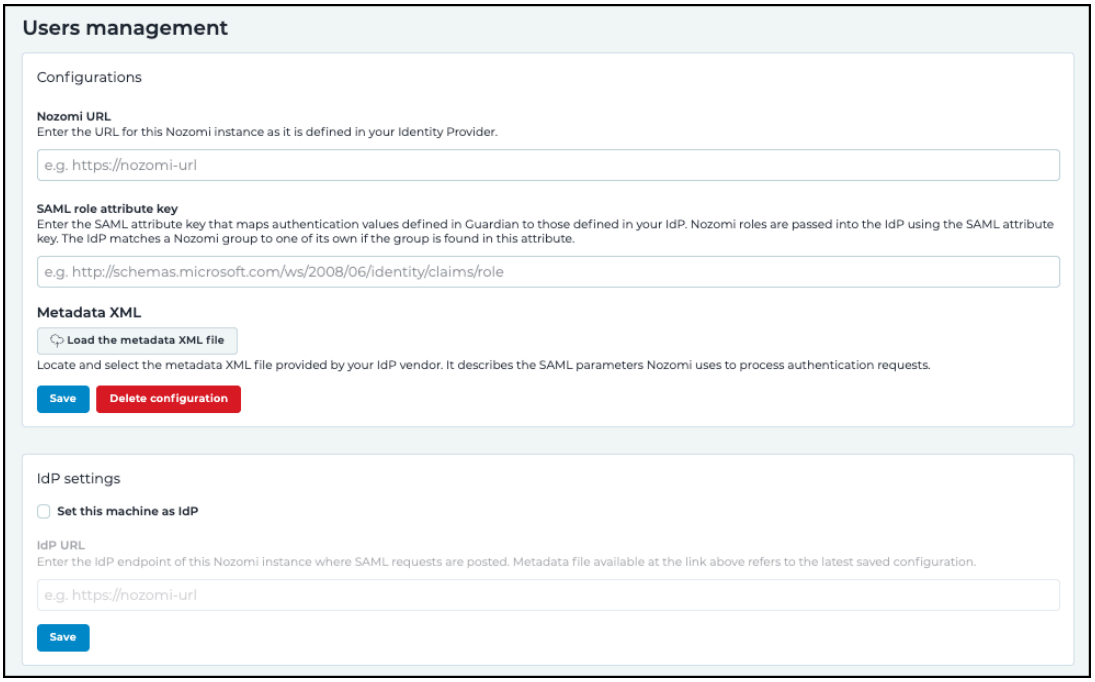
2. In the **Settings** section, select **Users**.

   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.

Single sign-on in the CMC

A description of how a Central Management Console (CMC) needs to be configured for single sign-on (SSO).

In installations with more than one *CMC* level:

- Each *CMC* should be configured as an *IdP*
- Each sensor should be configured to use the connected *CMC* as its *IdP*

For example, if a Guardian is connected to a mid-level *CMC*, that is connected to a top-level *CMC*, the following configurations allow *SSO* from the Guardian:

- The top-level *CMC* requires integration with the external *SSO* service on the **Users Management** page, and also requires configuration as an *IdP*
- The mid-level *CMC SAML* integration requires use of the top-level metadata file as its *SSO IdP*, and also requires configuration as an *IdP*
- The Guardian's *SAML* integration requires use of the mid-level *CMC*'s metadata file to allow the mid-level *CMC* to be used as an *IdP*

A login attempt from Guardian redirects the user through the mid-level *CMC* and the top-level *CMC* to the external *SAML* service.

Nozomi Networks recommends you use *hypertext transfer protocol secure (HTTPS)* as the protocol. However, you can also use *hypertext transfer protocol (HTTP)*. The address can also be a *fully qualified domain name (FQDN)*, or an *IP* address.

Examples:

```
cmc identity_provider_url https://192.168.1.8
cmc identity_provider_url https://cmc.example.com
```

Configure CMC for single sign-on (SSO)

How to configure Central Management Console (CMC) to use it as an identity provider for sensors that are connected to it.

About this task

Make sure that you have configured the *CMC* as an *IdP*.

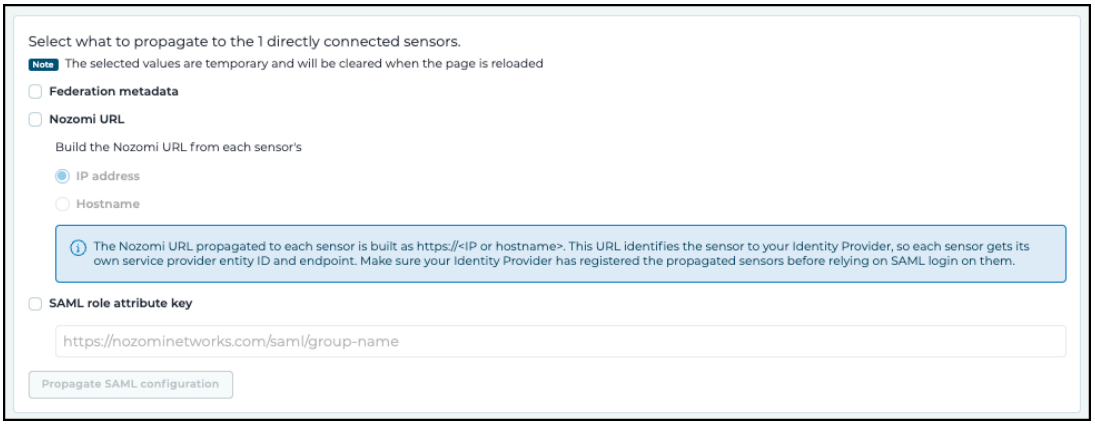
Procedure

1. In the *CMC*, add the following configuration rule to the `/data/cfg/n2os.conf.user` file. Replace `<ADDRESS>` with the address of the *CMC*.

   ```
   cmc identity_provider_url https://<ADDRESS>
   ```

2. To effect the change, restart all services and reboot the machine to effect the change.

3. Configure the sensors connected to the *CMC*. For more details, see Configure SAML integration (on page 55).
4. On the **SAML** page, customize the settings to be propagated to the connected sensors, and then select **Propagate SAML configuration**.
   You can push the *IdP* metadata, the *URL* for your Nozomi Networks instance, and the *SAML* role attribute key.



Note:
Use this option whenever the metadata changes, for example after renewing the *SSL* certificate on the *CMC*, instead of manually downloading and re-uploading the metadata file on each sensor.

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

CLI

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

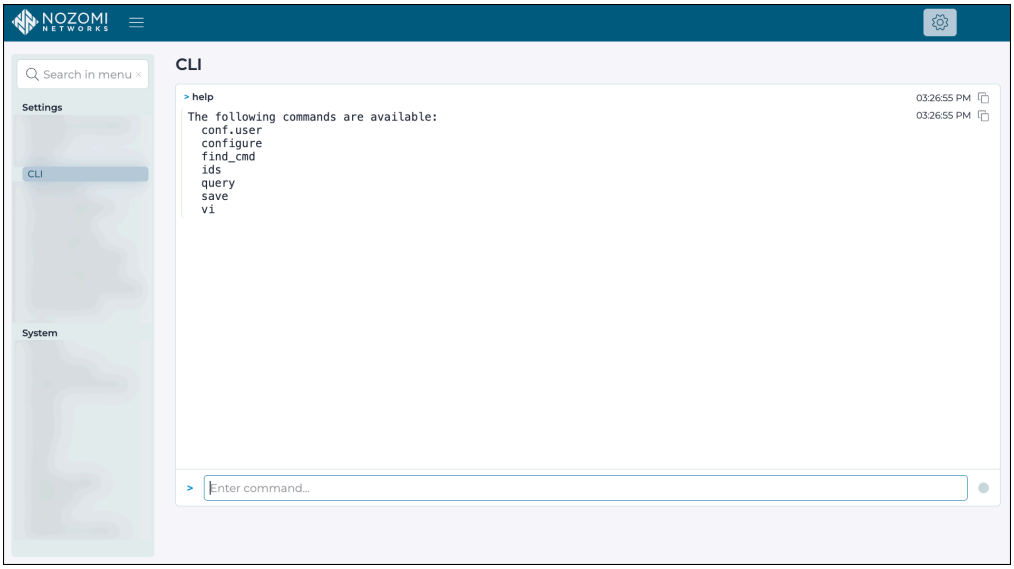


Figure 15. CLI page

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.

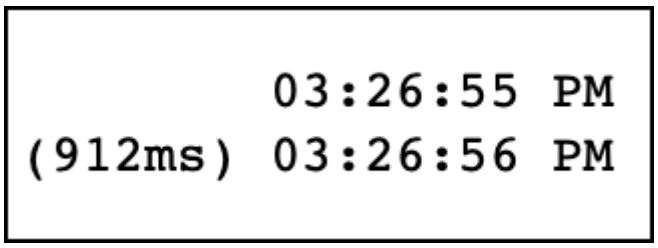


Figure 16. Duration shown next to timestamp

Table 4. 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.*

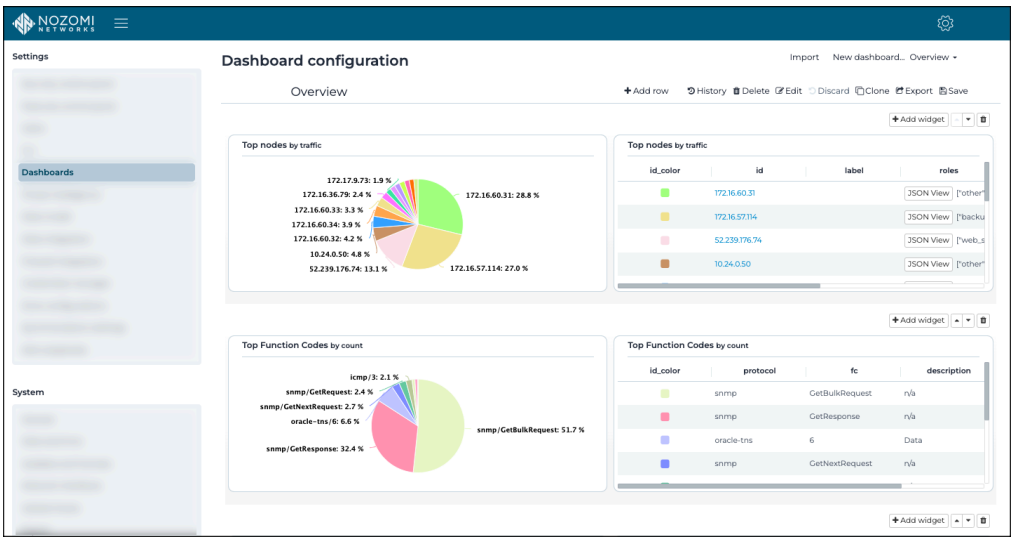


Figure 17. Dashboards page

Table 5. 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 68).

New dashboard
This lets you Create a dashboard (on page 64).

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.

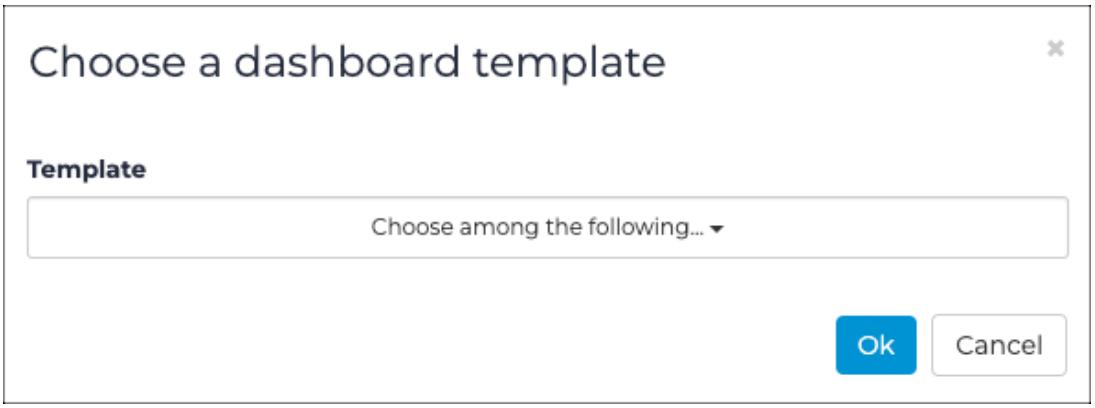
2. In the **Settings** section, select **Dashboards**.

   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.

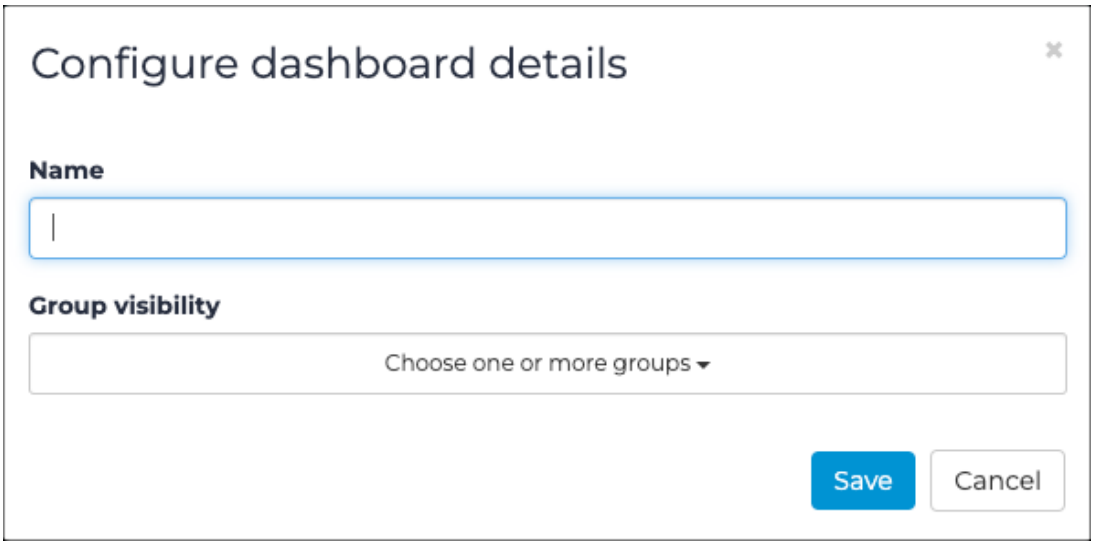


> **Note:**
> If you do not select a template, an empty dashboard will be created.

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 66)

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 64)
   - Select a dashboard (on page 68)
   - 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 *CMCs*.

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.*

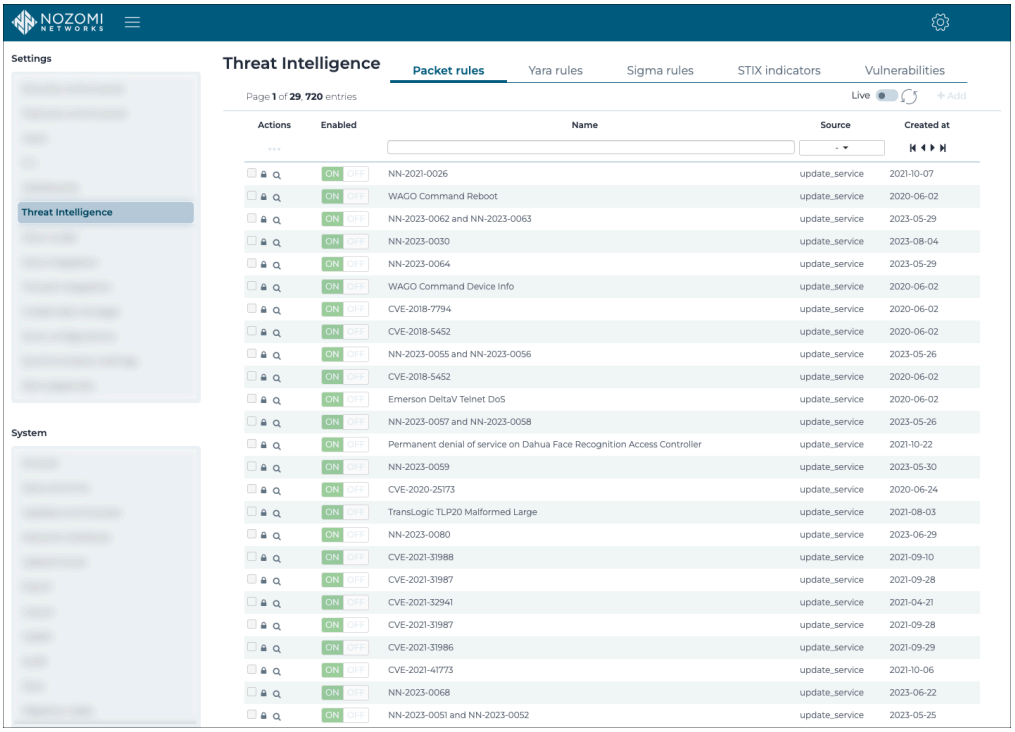


Figure 18. Threat Intelligence page

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

- Packet rules (on page 71)
- Yara rules (on page 72)
- Sigma rules (on page 73)
- STIX indicators (on page 74)
- Vulnerabilities (on page 75)

Packet rules

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

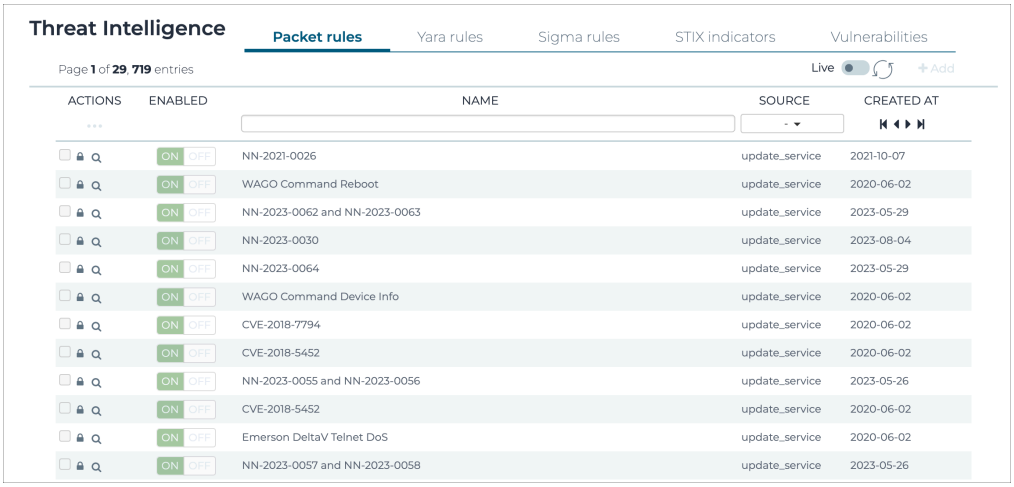


Figure 19. Packet rules page

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.*

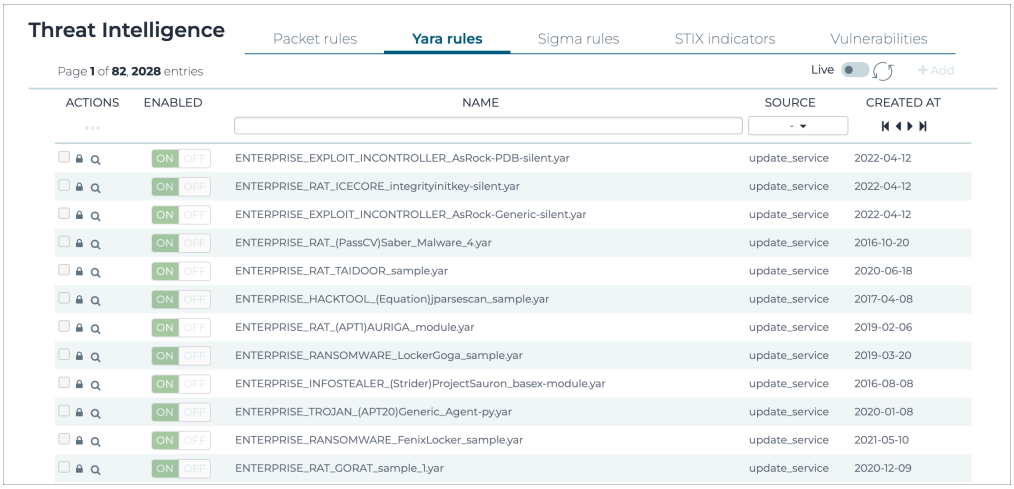


Figure 20. Yara rules page

YARA rules are executed on every file transferred over network protocols such as *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.*

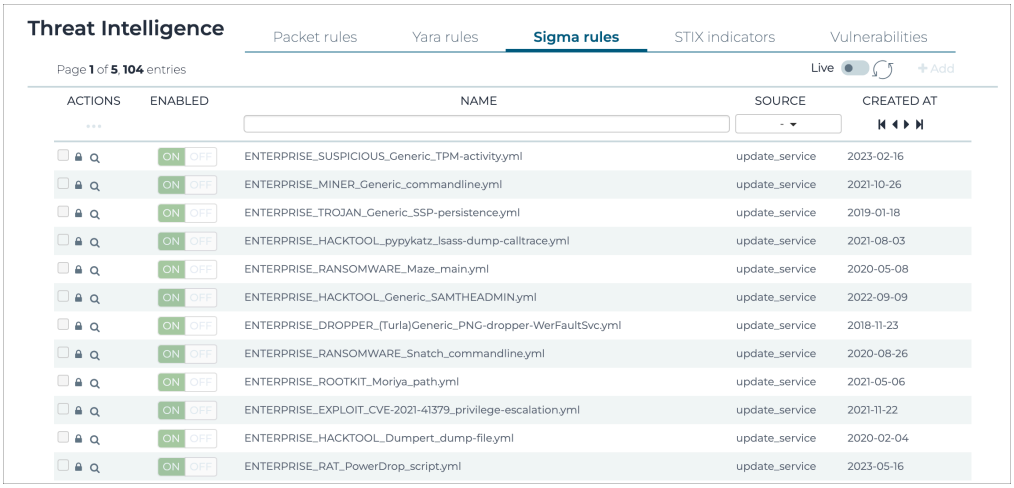


Figure 21. Sigma rules page

Sigma rules are versatile and generic rules written in the Sigma language. Primarily employed in threat detection and *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.*

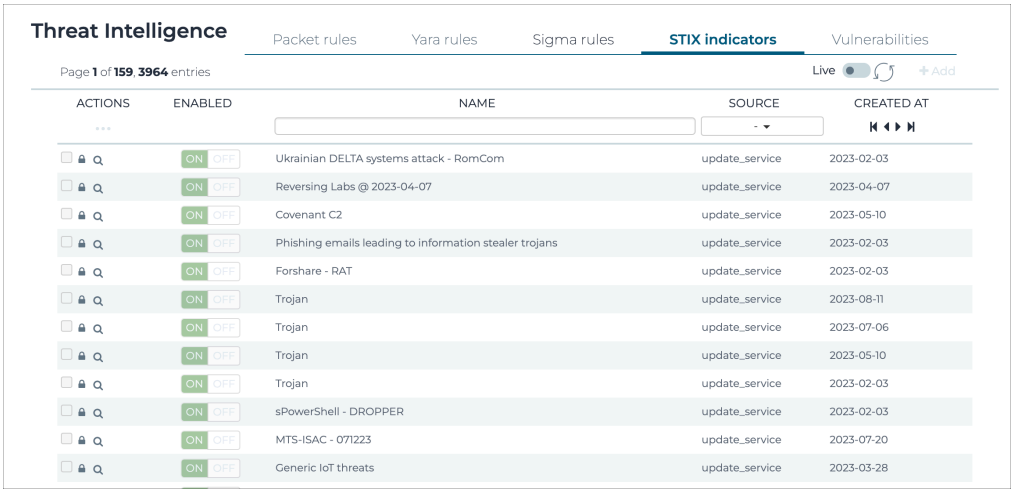


Figure 22. STIX indicators page

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.*

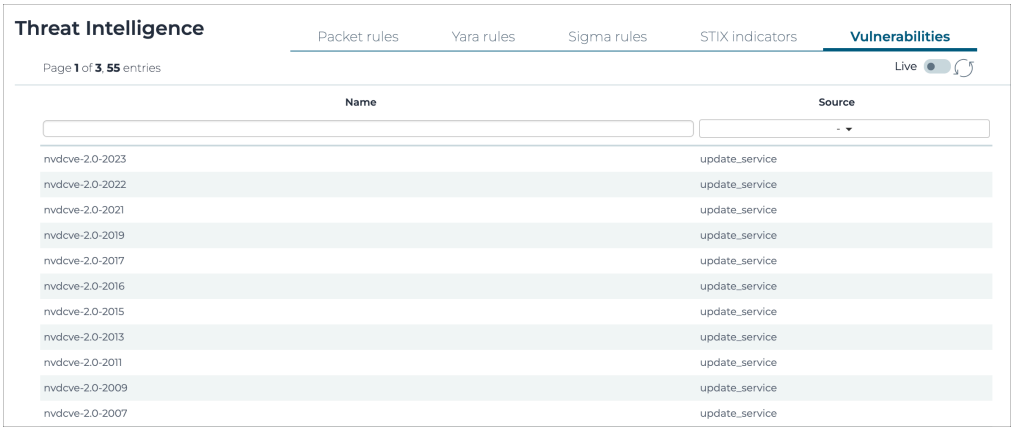
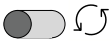


Figure 23. Vulnerabilities page

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.*

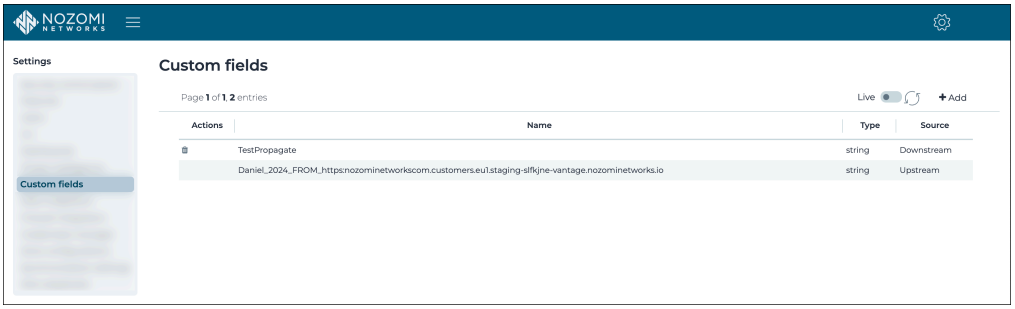
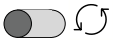


Figure 24. Custom fields page

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.

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
- Check Point IoT
- Kafka
- Cisco ISE
- NetWitness
- Aruba ClearPass
- Webhook

Data integration

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

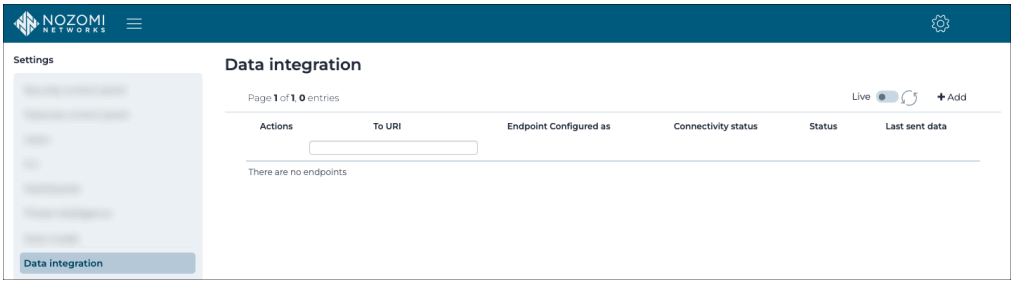


Figure 25. Data integration page

Note:

When you enable **Data integration** in the *CMC*, the data integrations of downstream sensors will not show in the *CMC*.

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*.

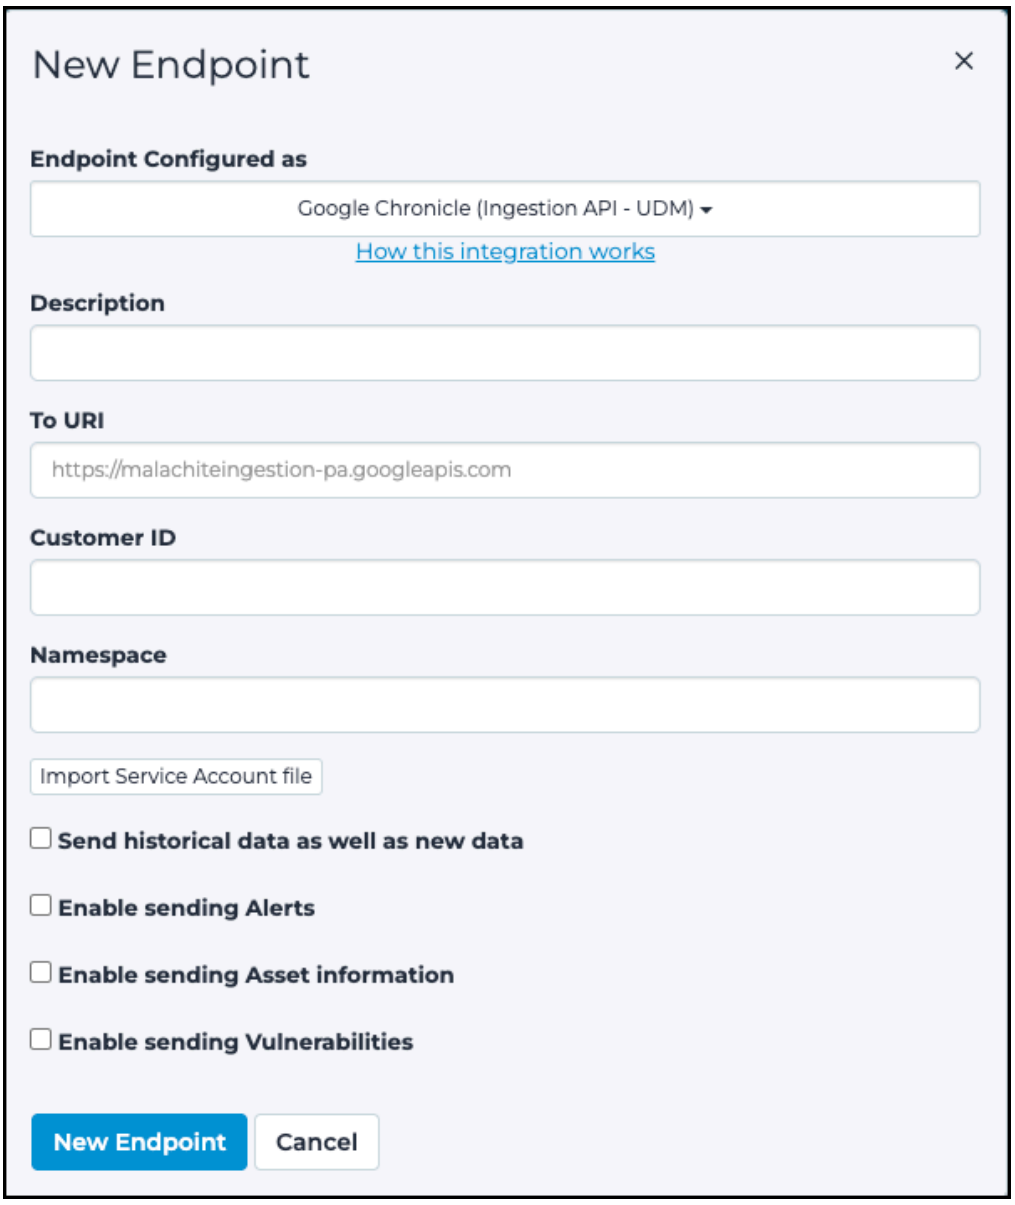


Figure 26. Google Chronicle dialog

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.

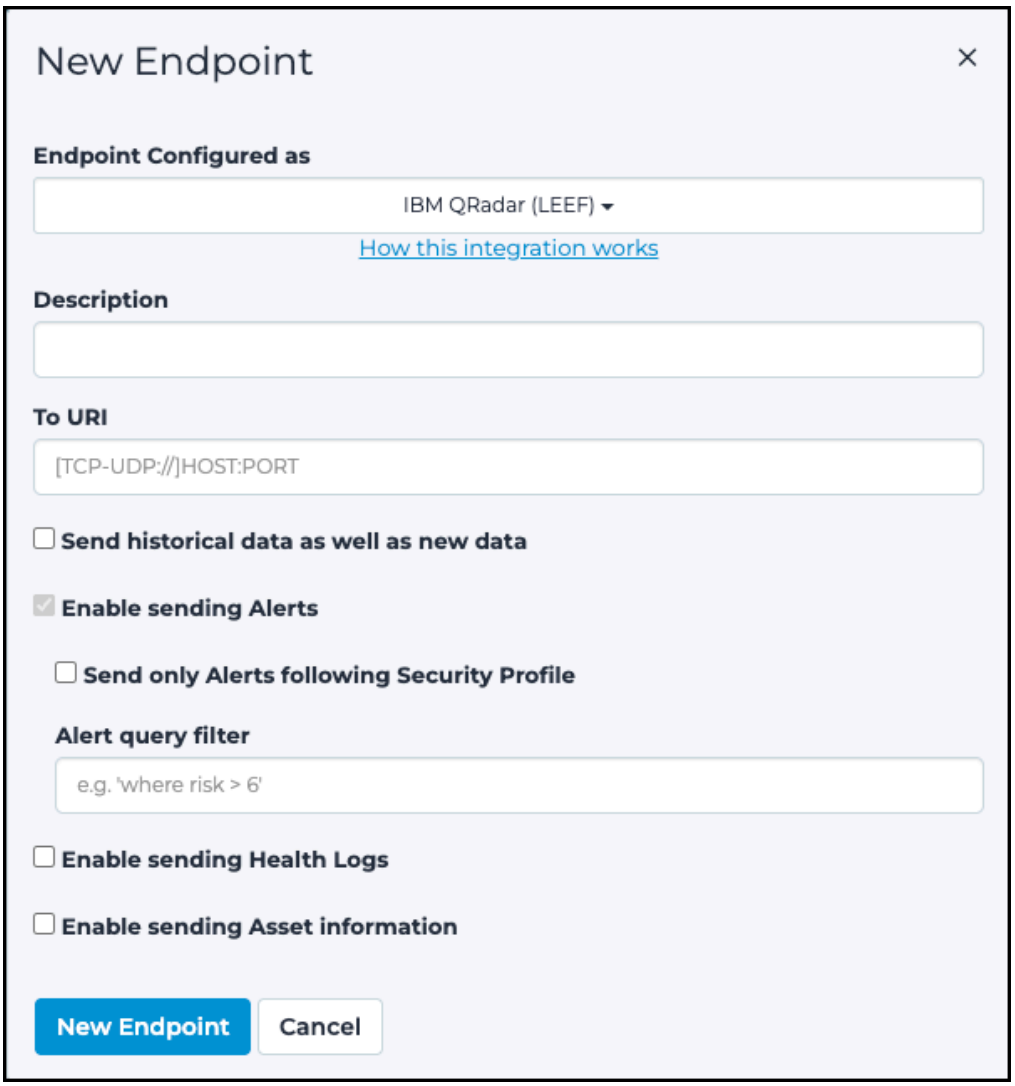


Figure 27. IBM QRadar (LEEF) dialog

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.

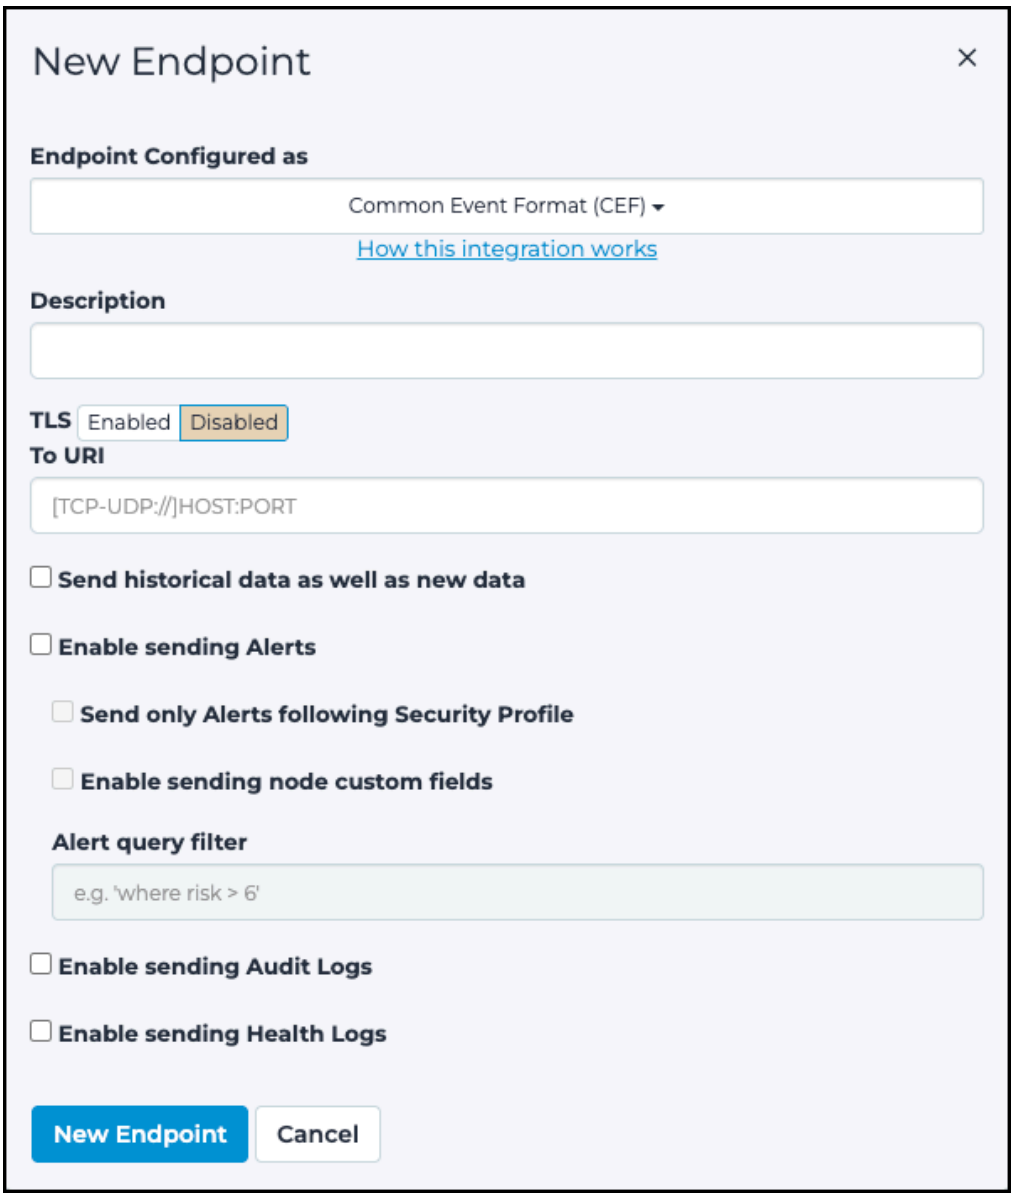


Figure 28. Common Event Format (CEF) dialog

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.

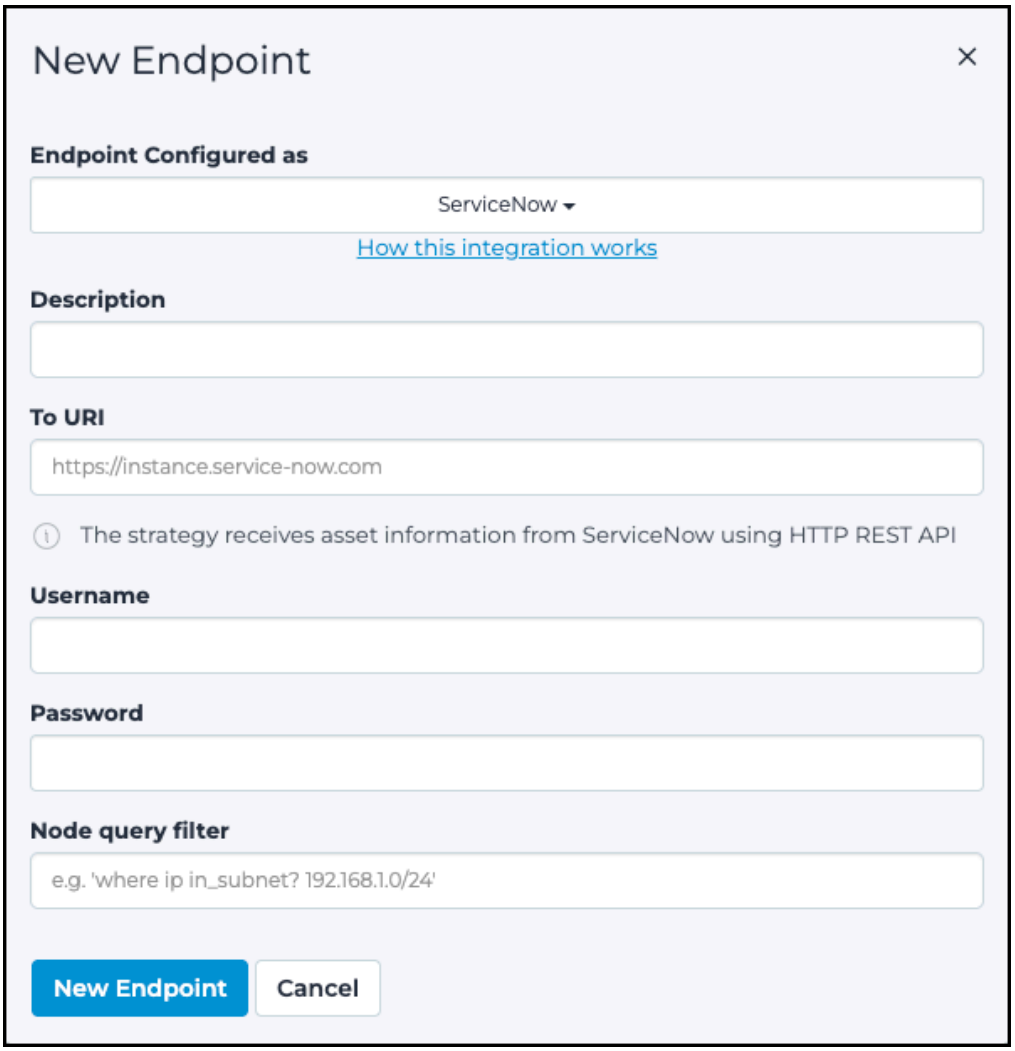


Figure 29. ServiceNow dialog

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.

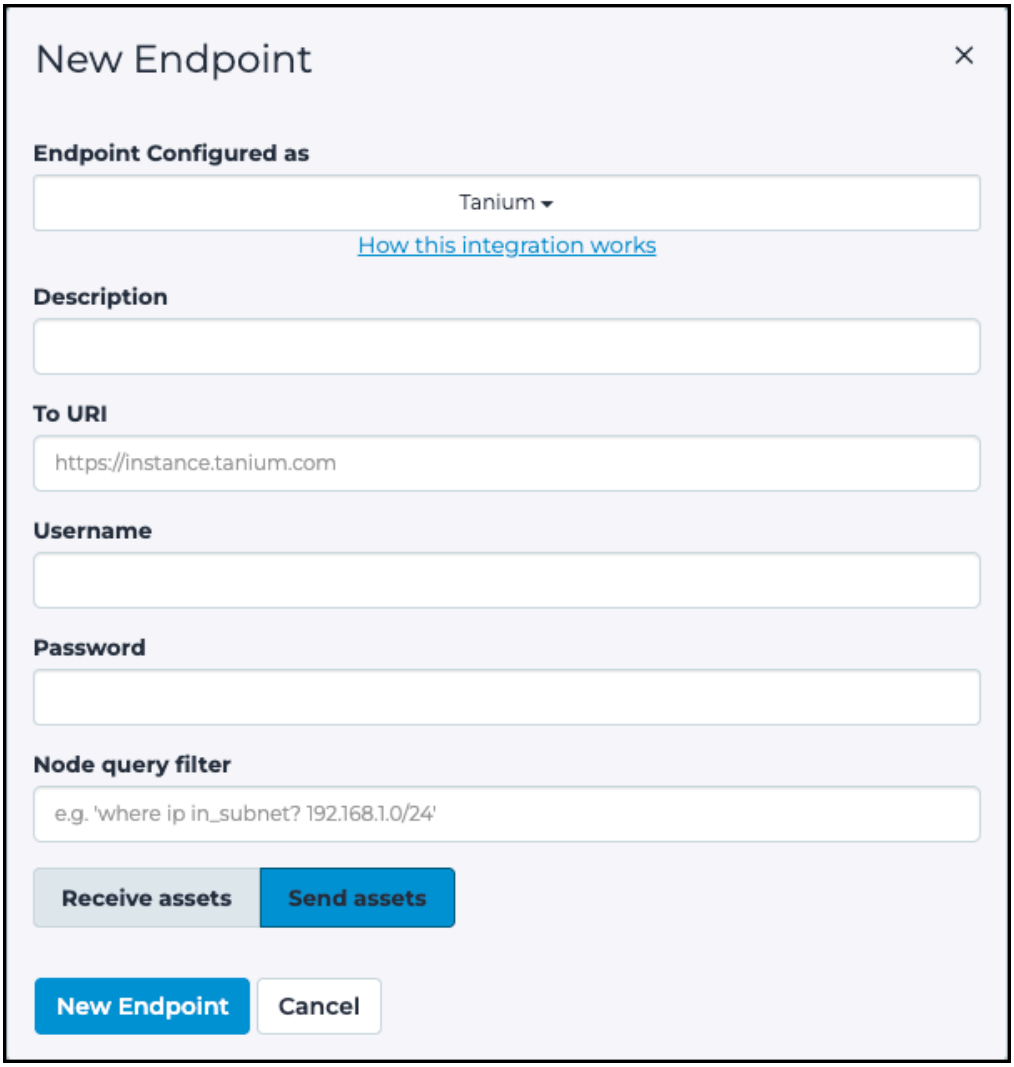


Figure 30. Tanium dialog

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.

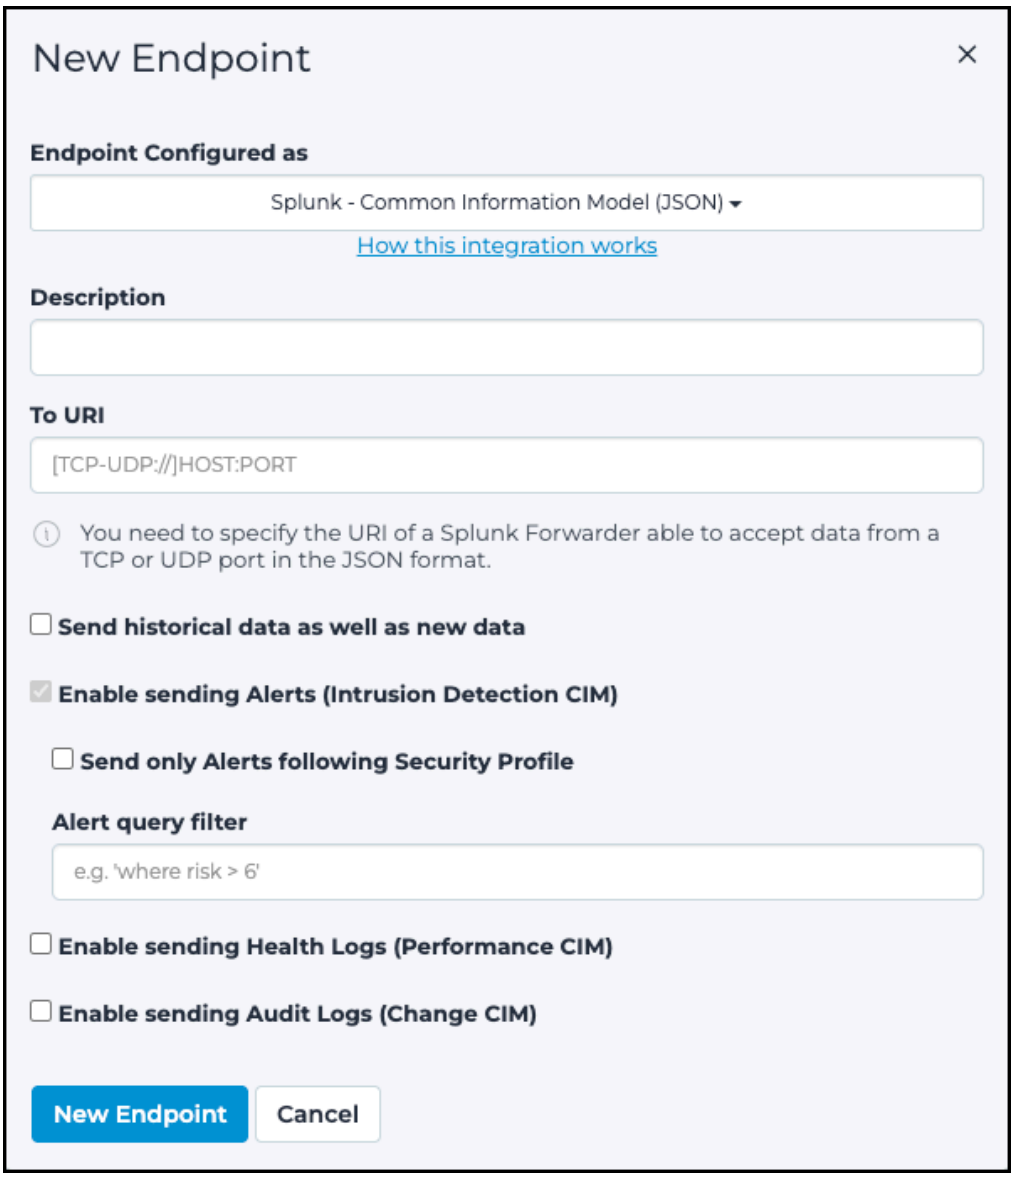


Figure 31. Splunk dialog

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.

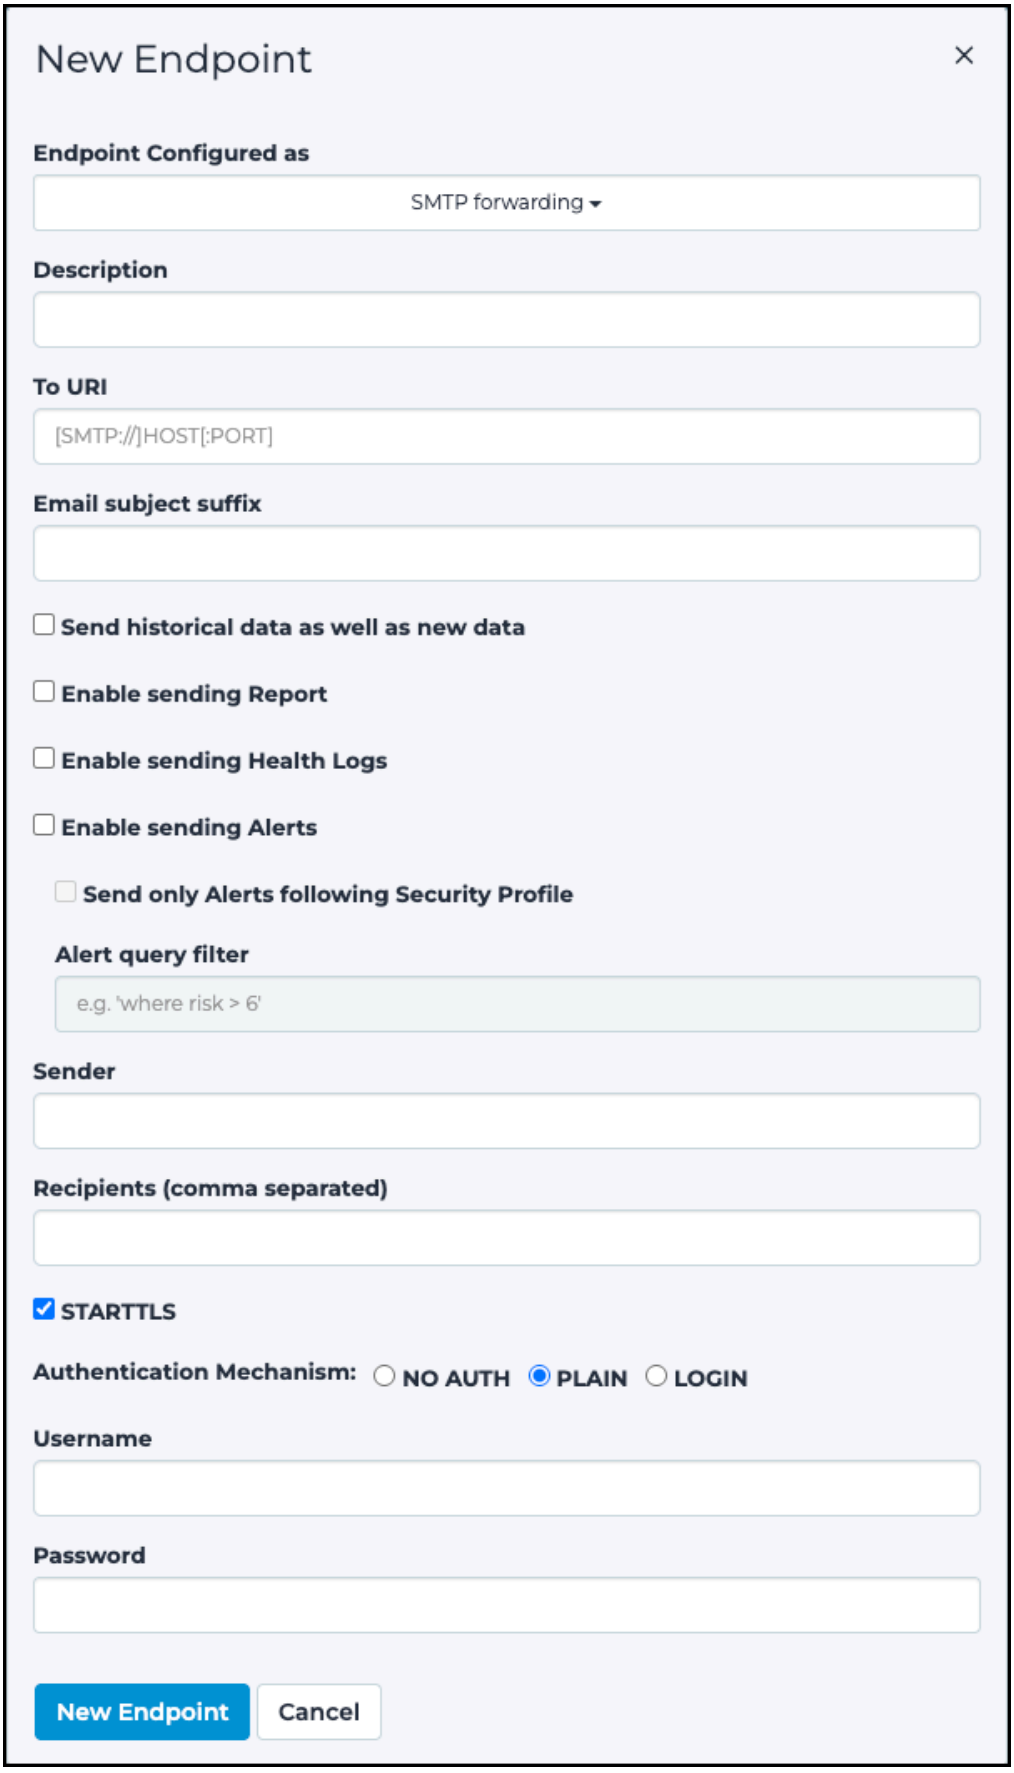


Figure 32. SMTP forwarding dialog

SNMP trap

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

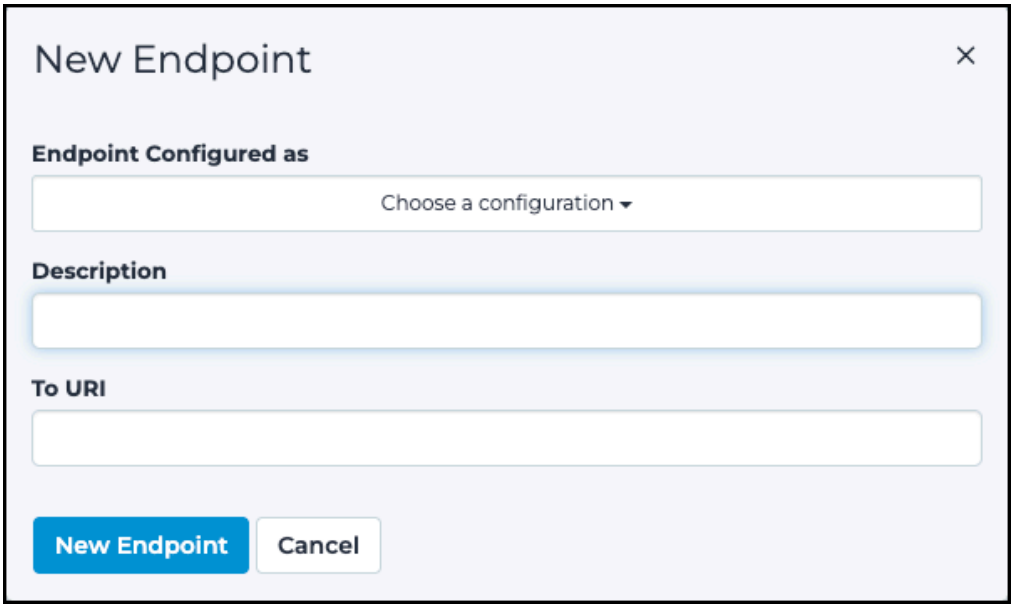


Figure 33. SNMP trap dialog

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.

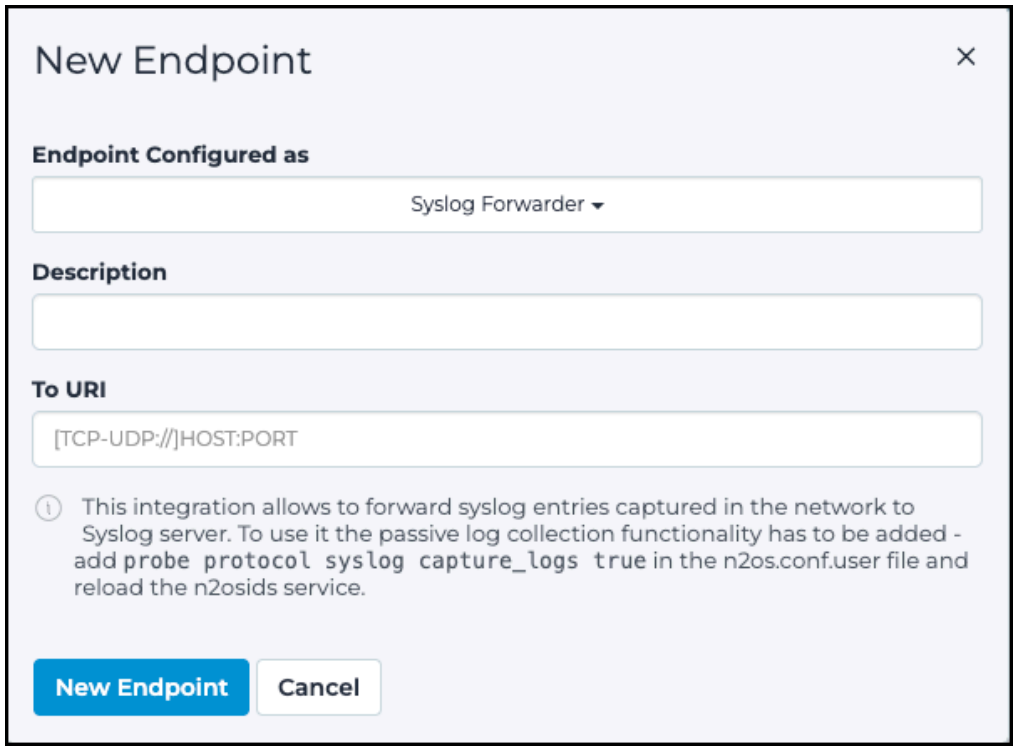


Figure 34. Syslog Forwarder dialog

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).

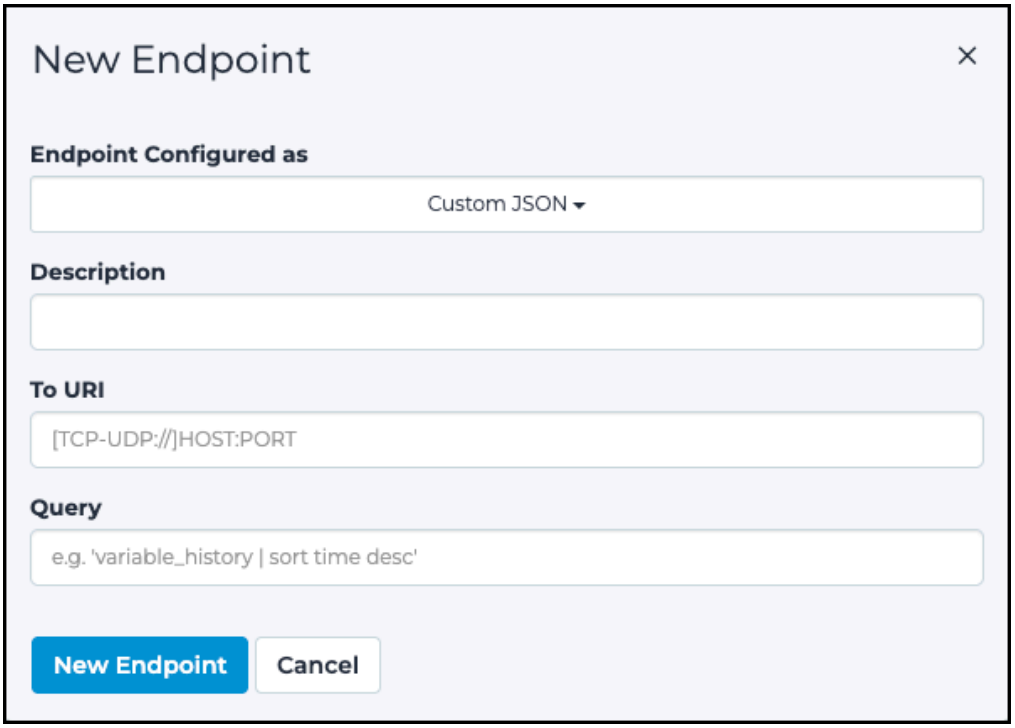


Figure 35. Custom JSON dialog

Custom CSV

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

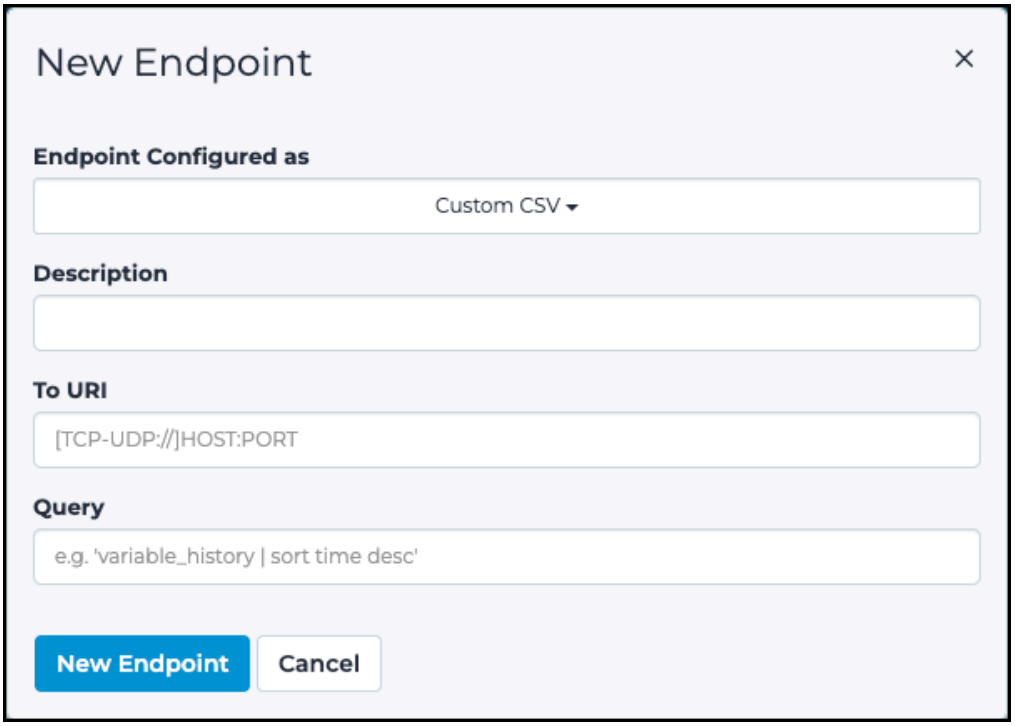


Figure 36. Custom CSV dialog

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.

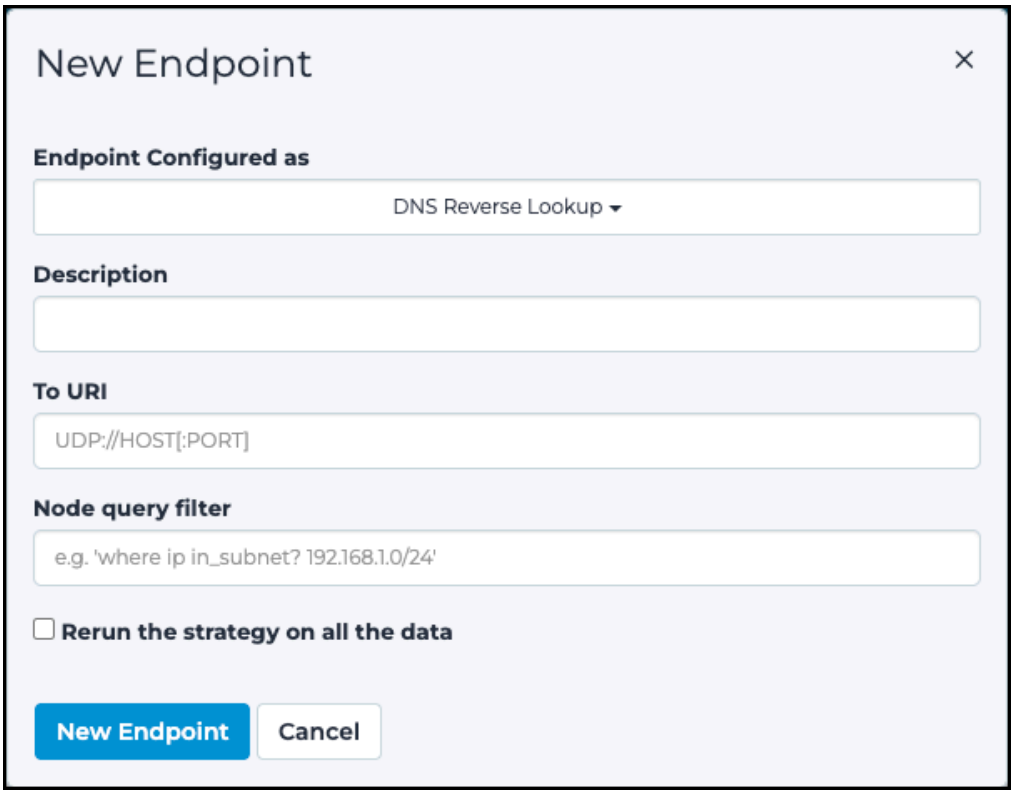


Figure 37. DNS Reserve Lookup dialog

Check Point IoT

This integration allows you to forward assets to Check Point IoT when Check Point SmartConsole asks for them. It also sends updated asset information to Check Point SmartConsole whenever an asset is updated on the appliance.

You can select **How this integration works** to view additional details. This integration is only available only on *CMCs*.

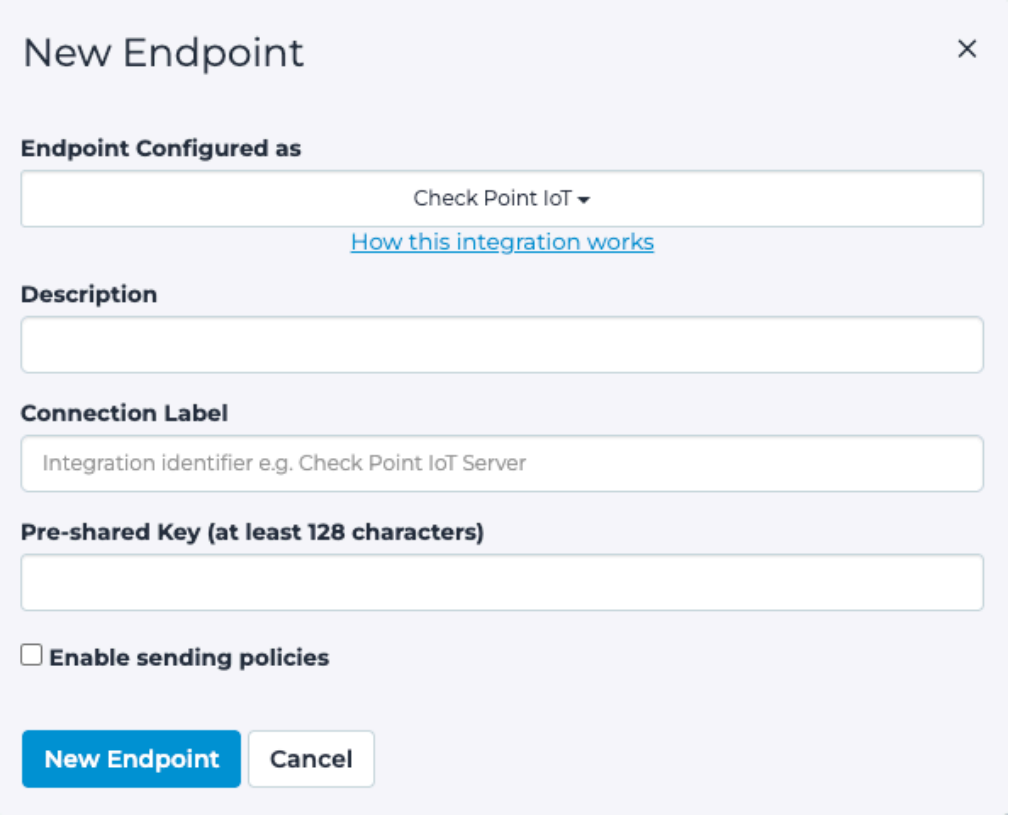


Figure 38. Check Point IoT dialog

Policies

If the **Enable sending policies** checkbox is selected, then the appliance will qualify itself as the policy provider for Check Point IoT. It will send nodes blocking policies when Check Point SmartConsole asks for them.

Configuration

You need to configure the Check Point SmartConsole with:

- Hostname: IPOfThisAppliance/websocket
- Port: 443
- Pre-shared Key: choose the same value when you configure this integration and the Check Point SmartConsole

Once you have created this configuration, you can test connectivity and authentication on the Check Point SmartConsole.

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.

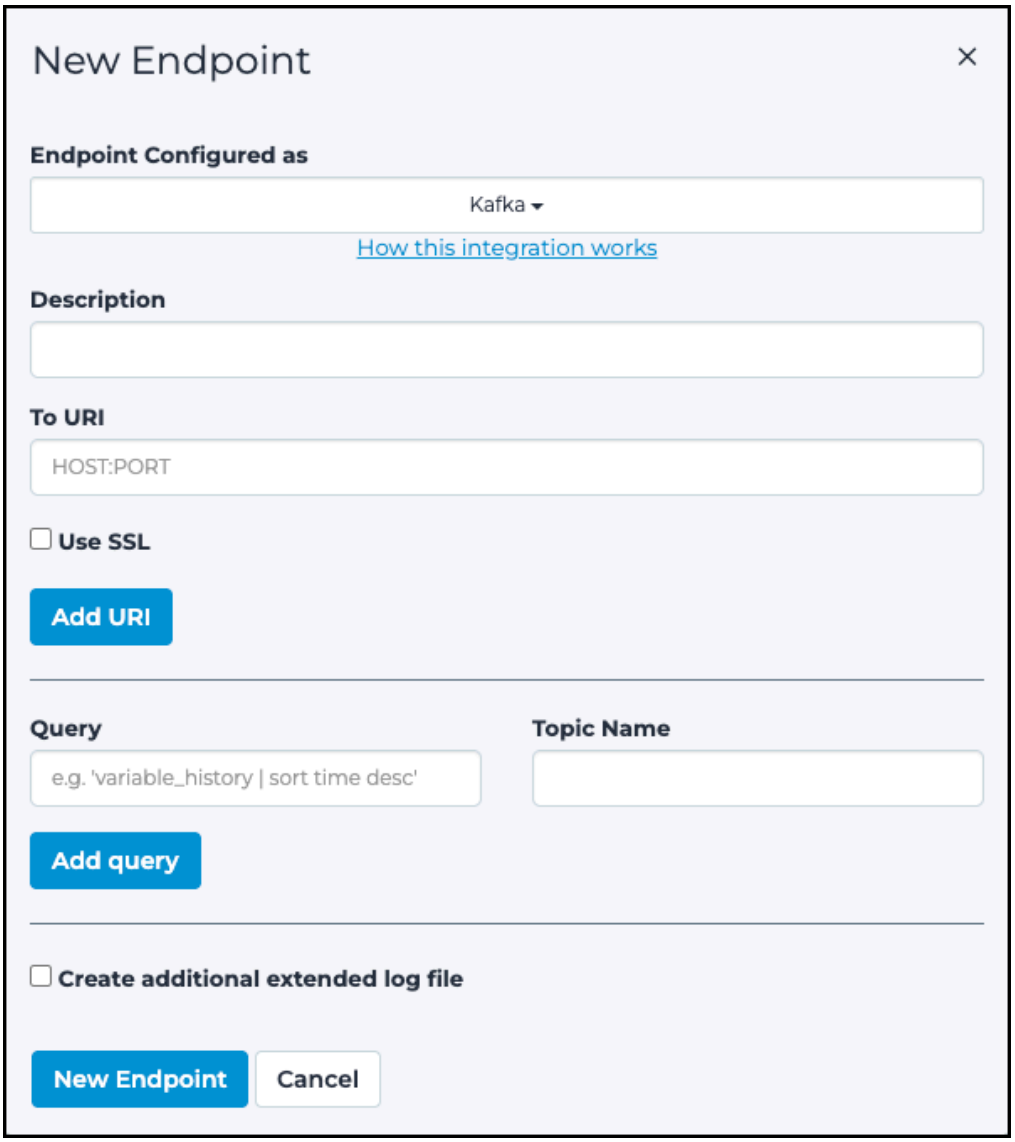


Figure 39. Kafka dialog

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

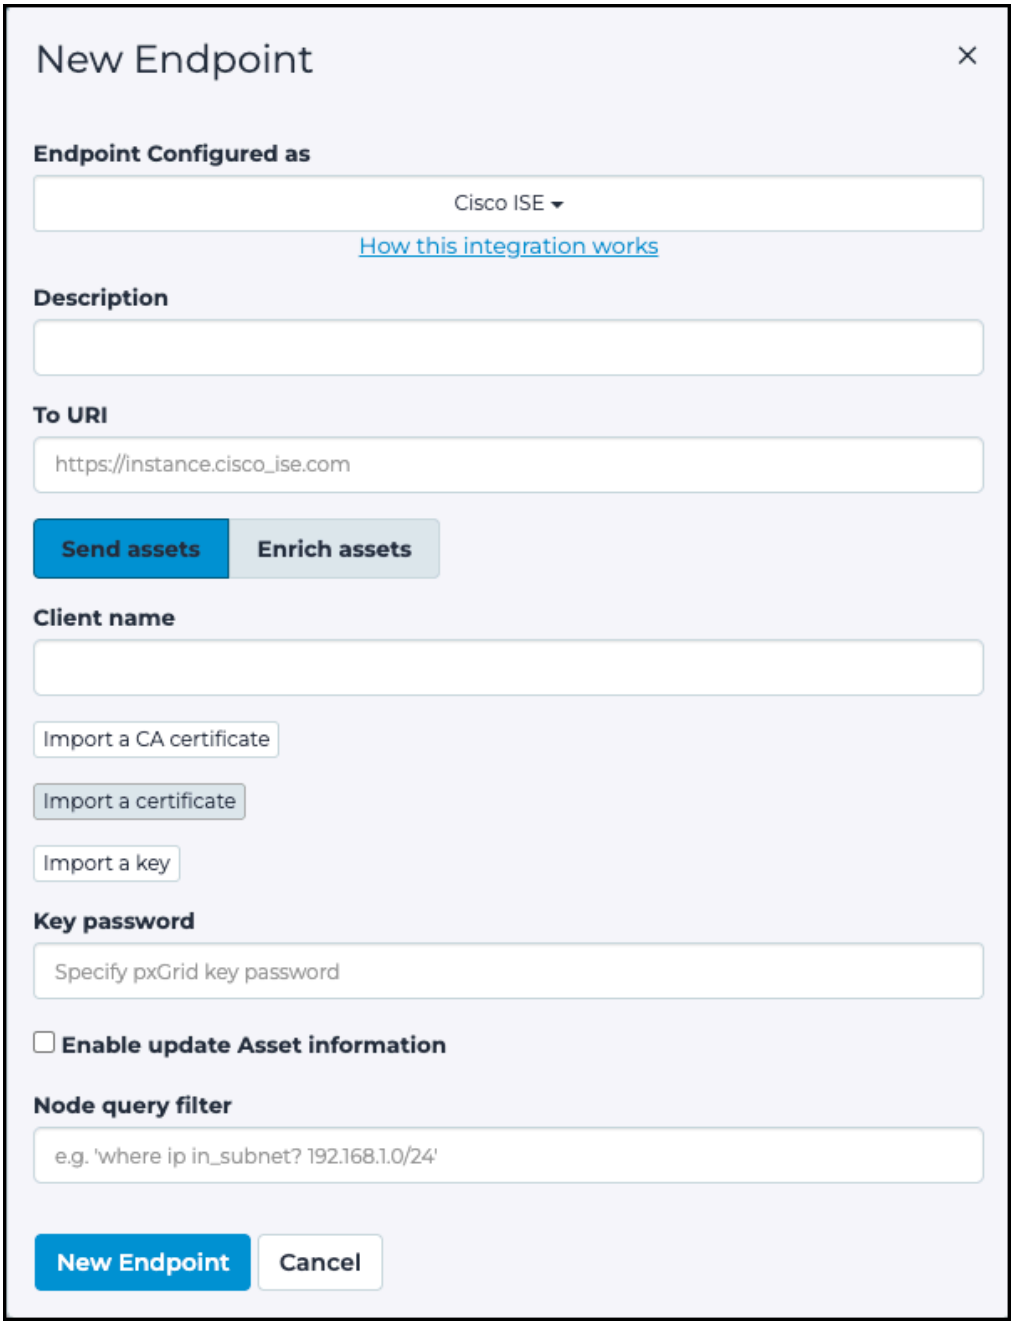


Figure 40. Send assets dialog

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

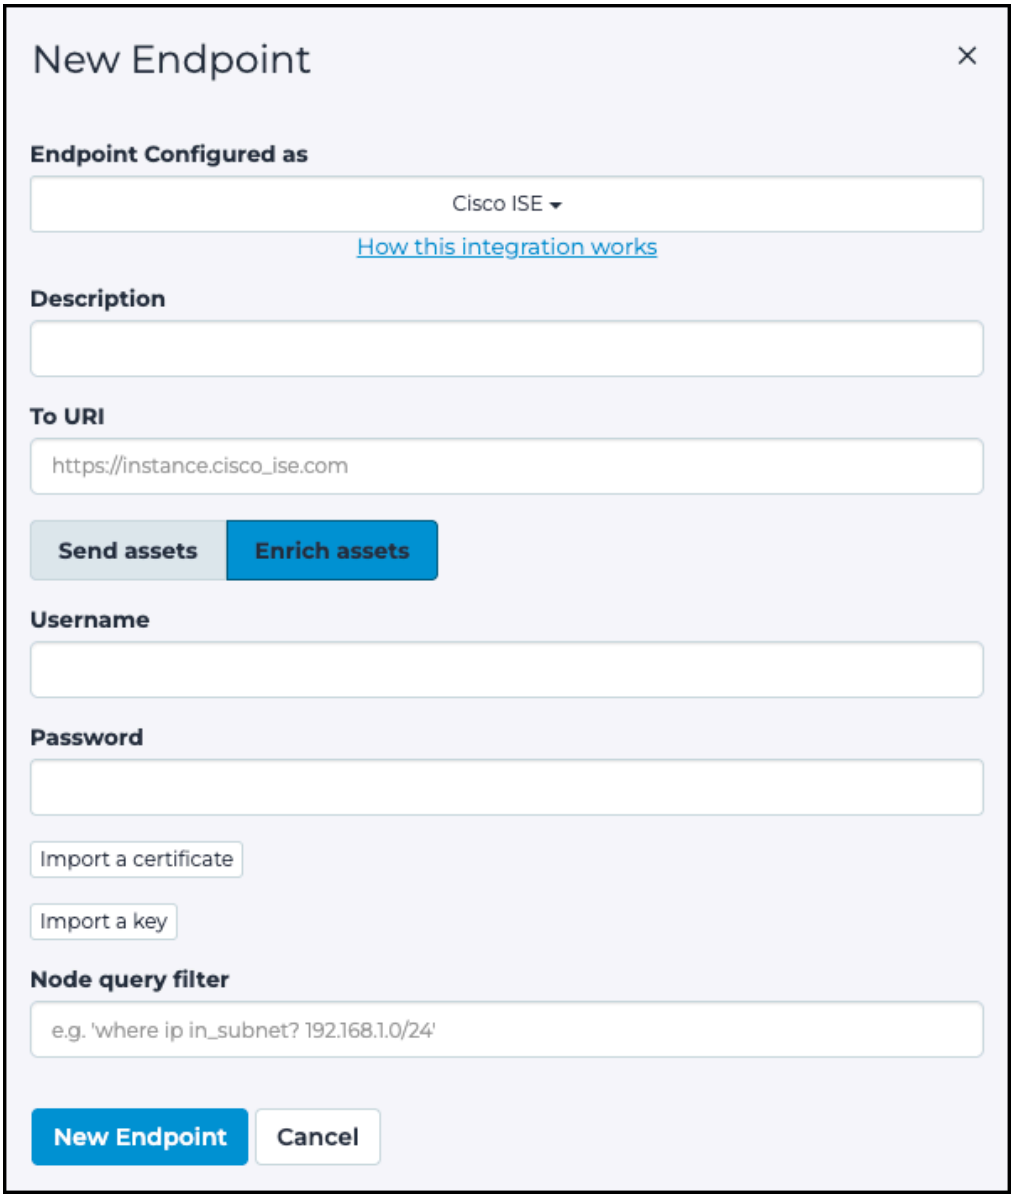


Figure 41. Enrich assets dialog

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.*

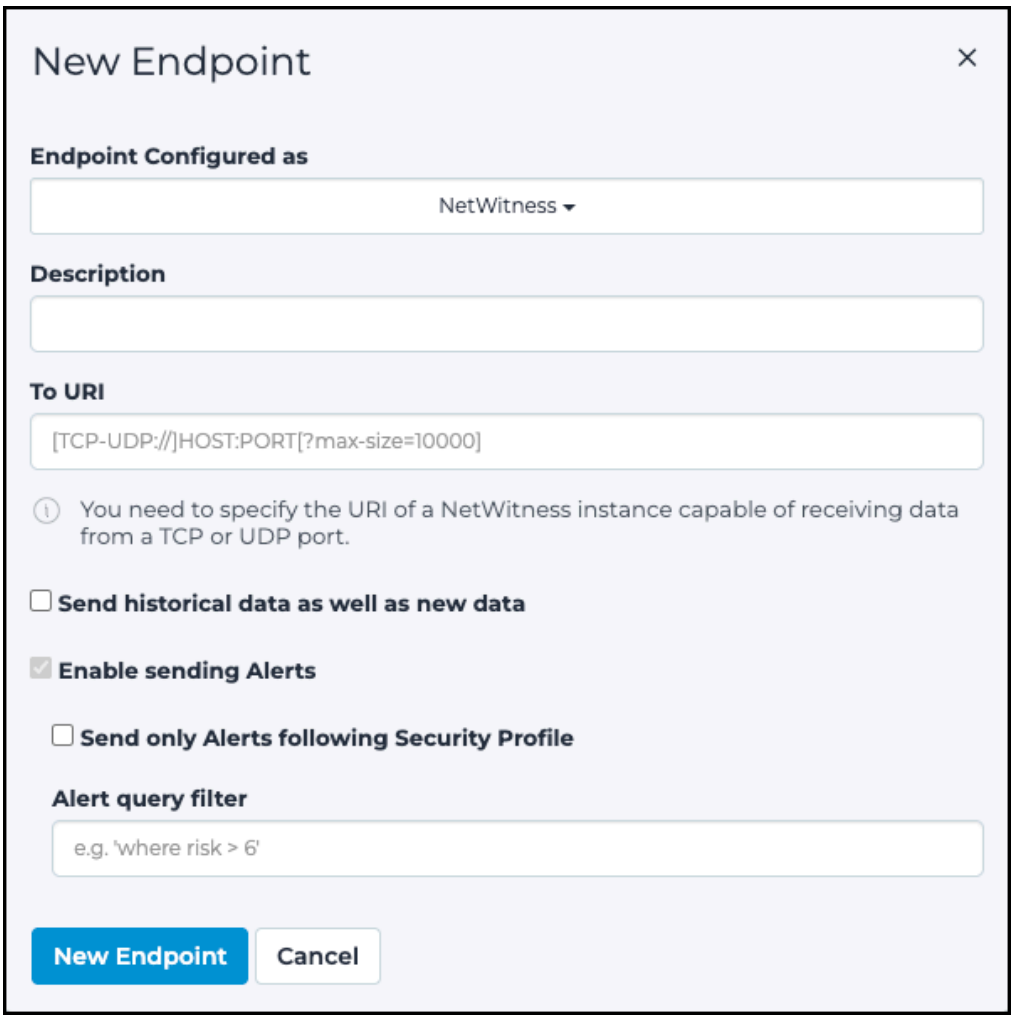


Figure 42. NetWitness dialog

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.*

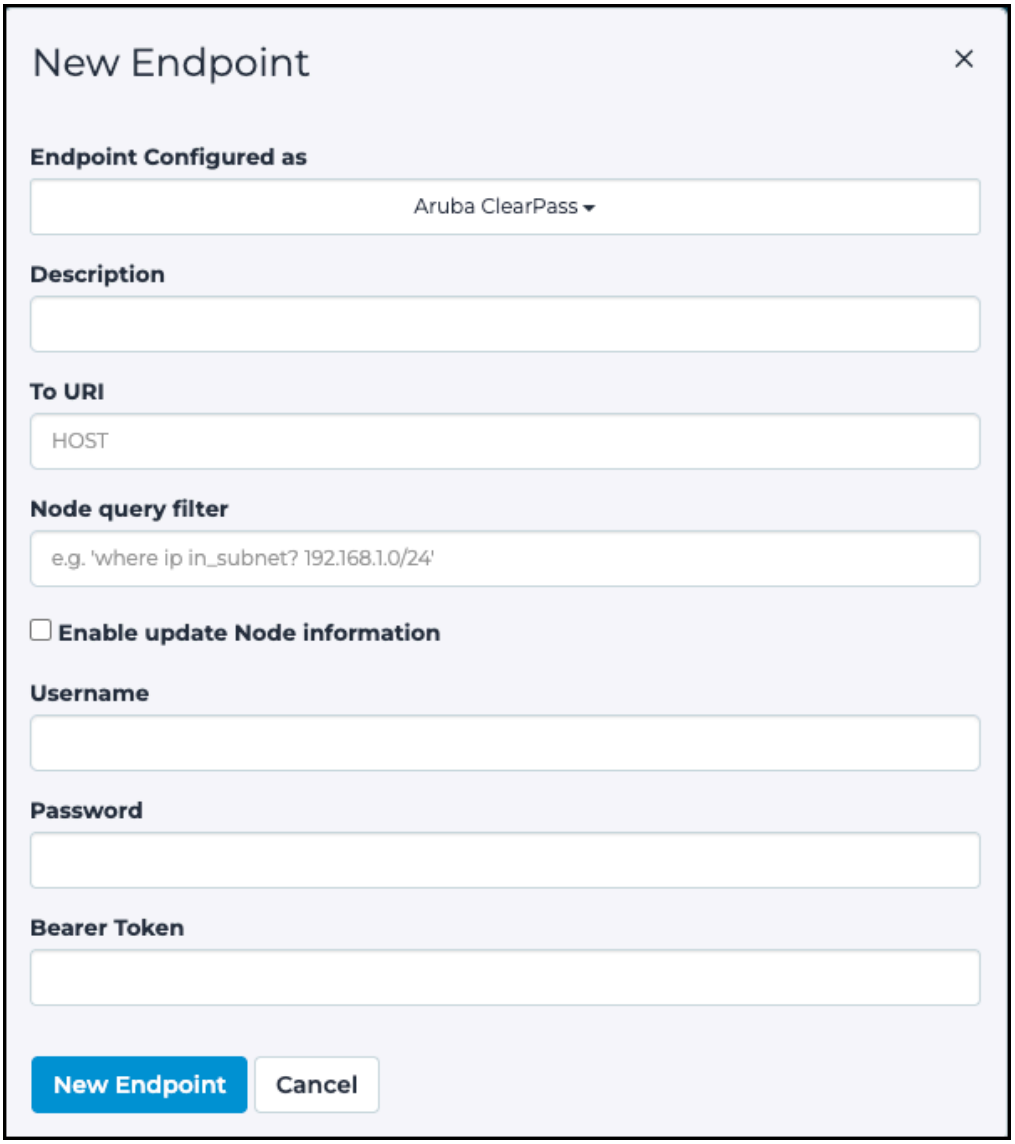


Figure 43. Aruba ClearPass dialog

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.

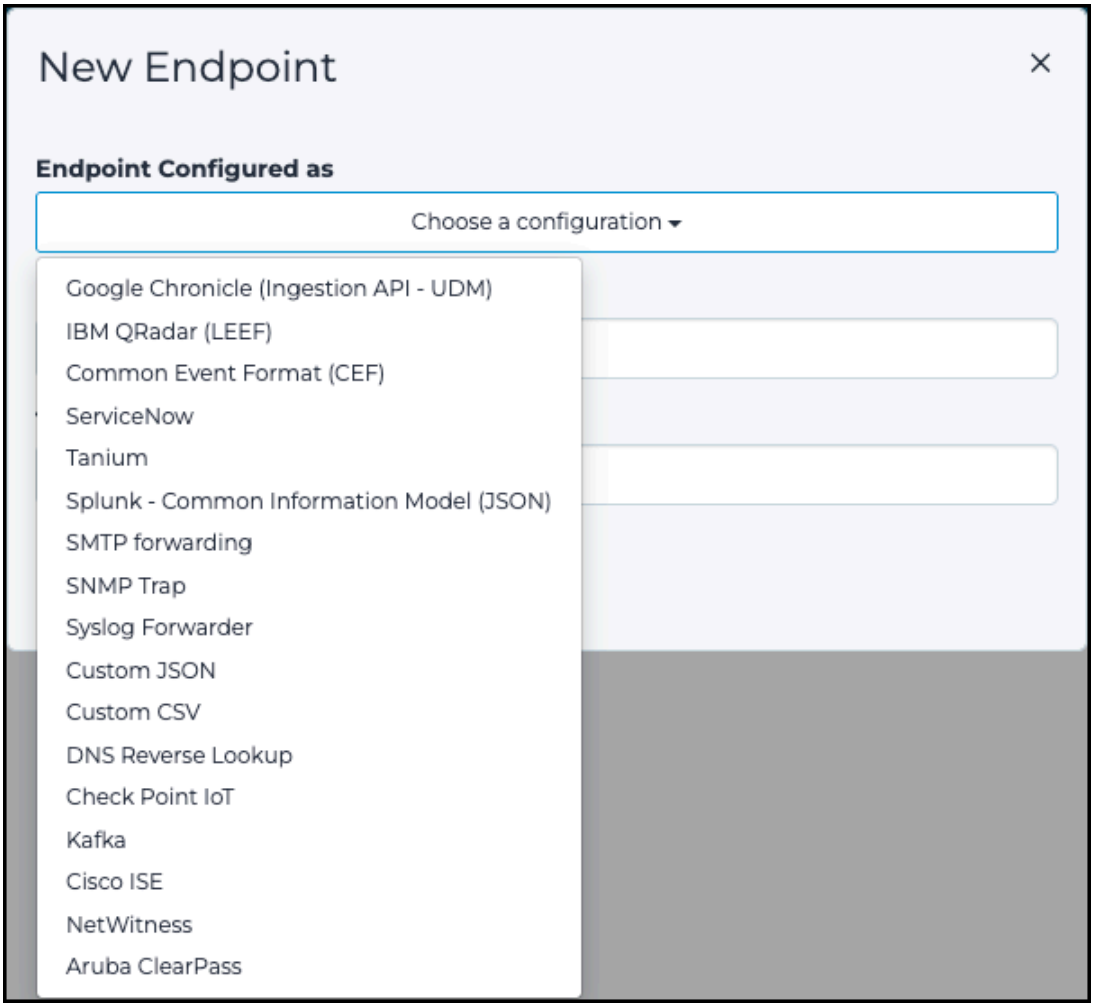
2. In the **Settings** section, select **Data integration**.

   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.

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.*

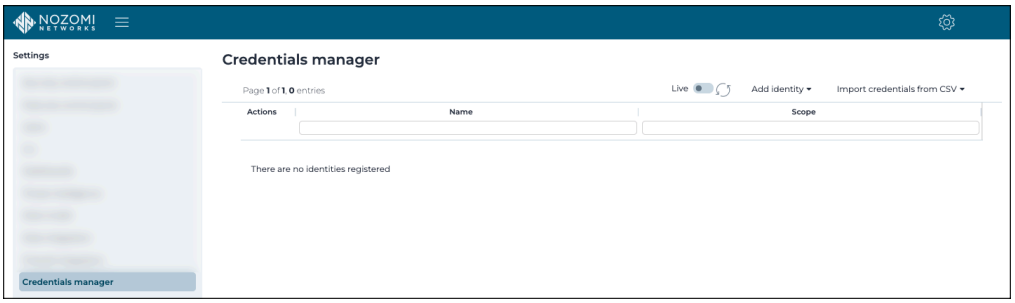


Figure 44. Credentials manager page

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.

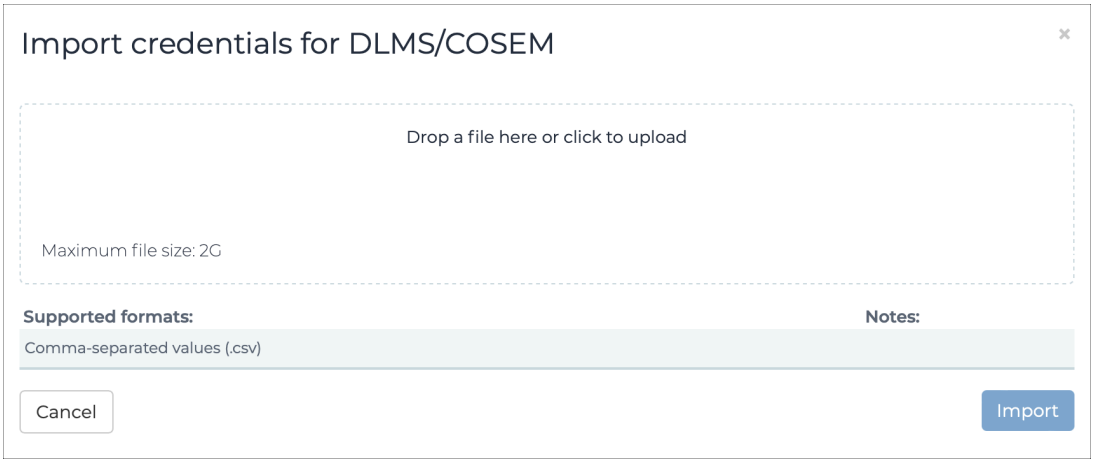
   Result: A dialog shows.

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.*

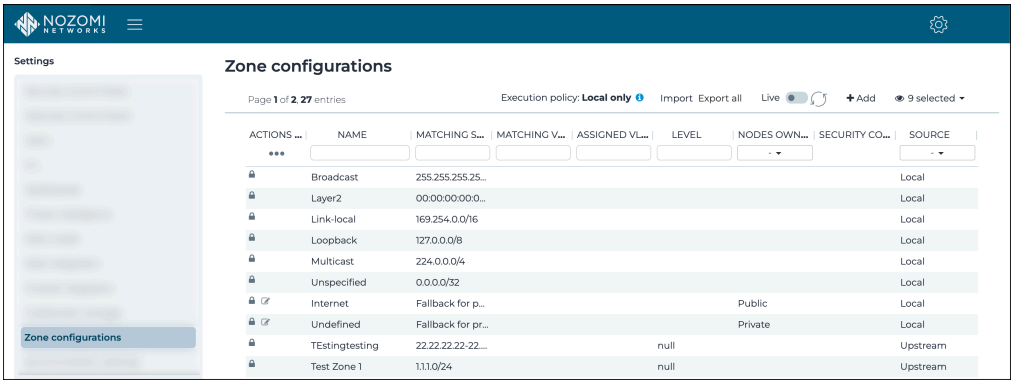


Figure 45. Zone configurations page

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 17).

Learning mode

Used to override the global settings from the **Learning** section of the Security control panel (on page 17).

Security profile

Used to override the global settings from the **Security profile** section of the Security control panel (on page 17).

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.

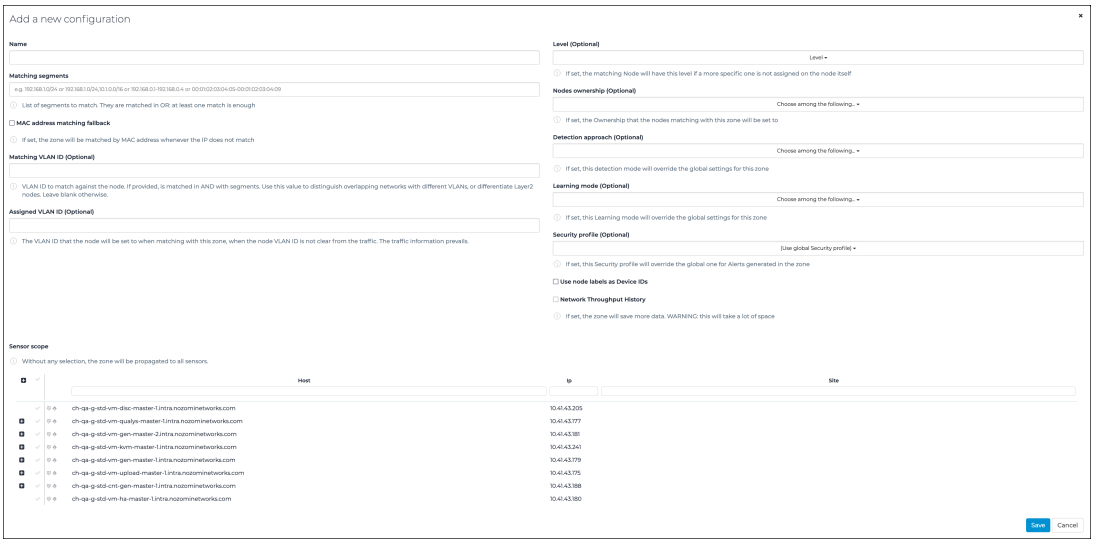
2. In the **Settings** section, select **Zone configurations**.

   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 upstream or downstream machines.*

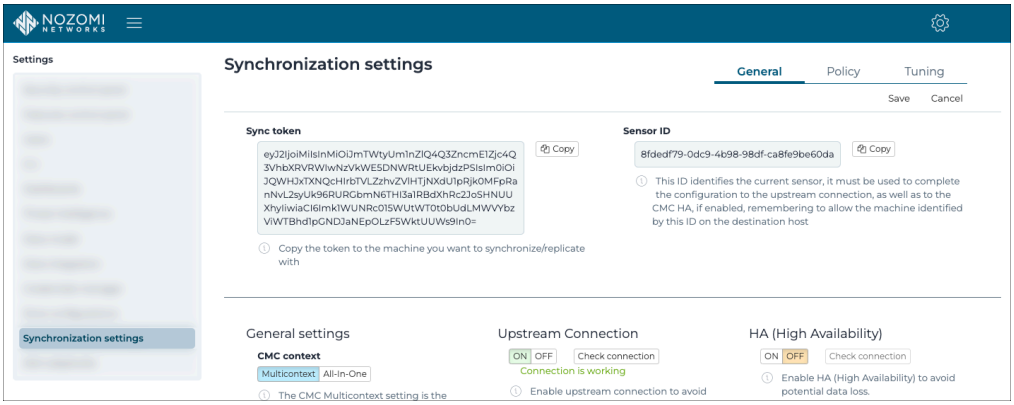


Figure 46. Synchronization settings page

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

- General (on page 126)
- Data Diode (on page 138)
- Policy (on page 149)
- Tuning (on page 151)

General

*The **General** page lets you configure the upstream connection, high availability (HA) modes, and other general settings.*

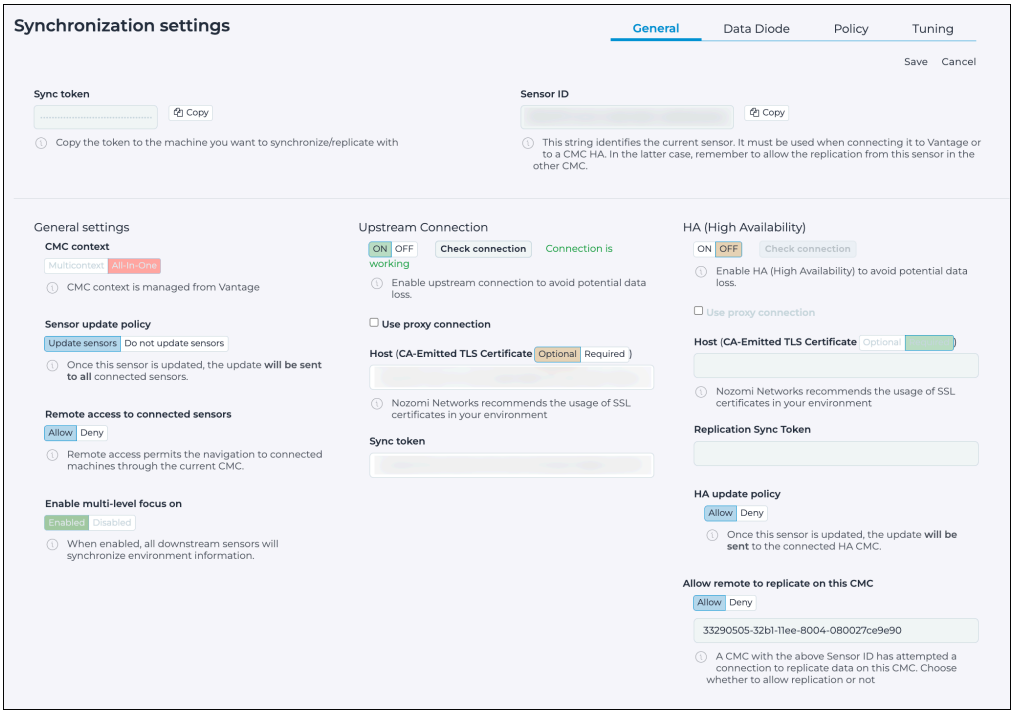


Figure 47. General page

The **General** page has these sections:

- General settings (on page 127)
- Upstream connection (on page 129)
- High Availability mode (on page 130)

General settings

*The **General settings** section lets you choose the context for the CMC, how sensors will be updated, and navigation settings.*

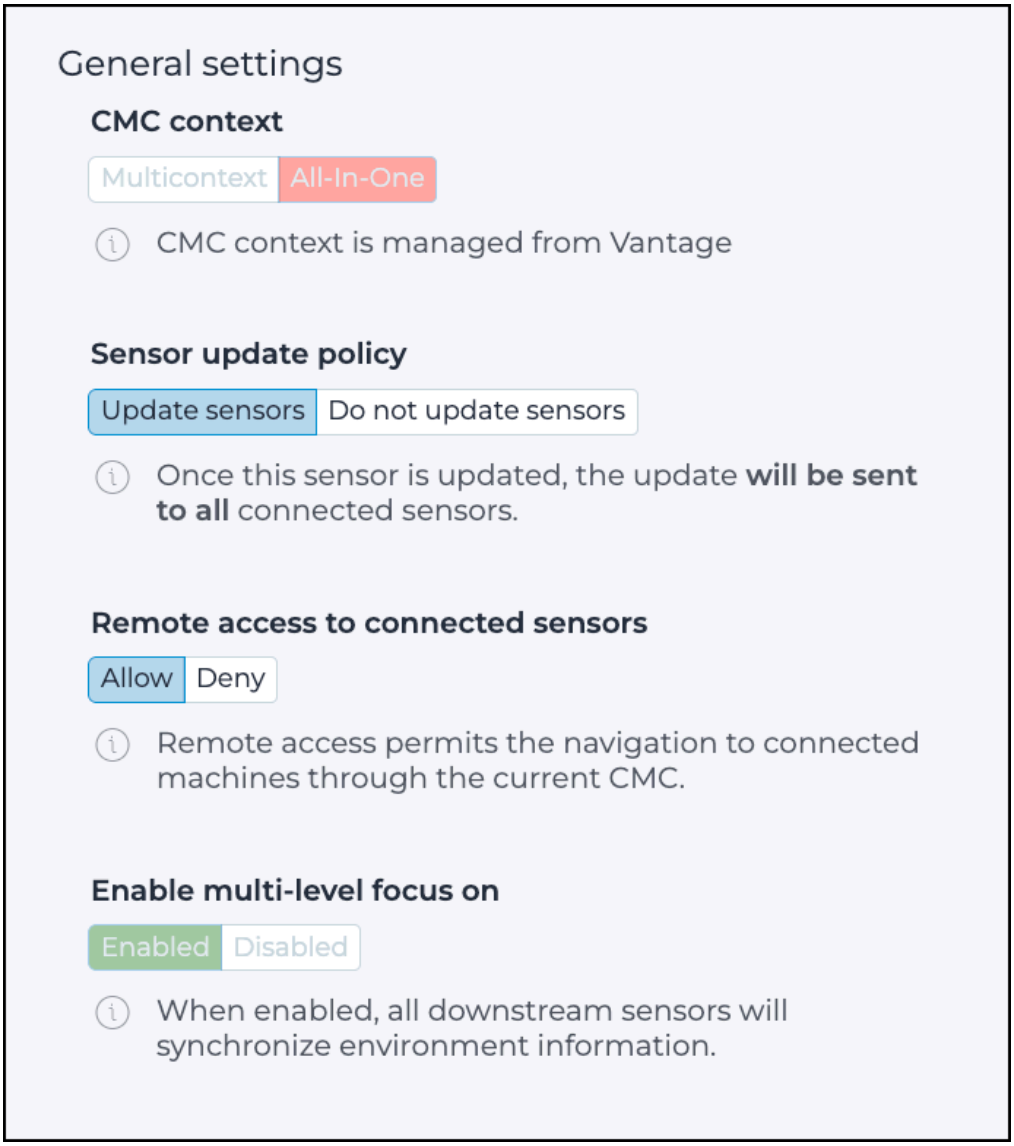


Figure 48. General settings

CMC context

This section has a toggle that lets you choose between two contexts:

- **Multicontext**
- **All-In-One**

Multicontext: is the default setting, and is suitable for most environments. **Multicontext** indicates that the data gathered from the sensors connected to the *CMC* will be collected and kept separately. In **Multicontext** mode, the user can focus on a single Guardian to access their data in their separate contexts. This is the default operational mode and it permits the highest scalability and supports multitenancy, which is ideal for *Managed Security Service Provider (MSSP)*s.

All-In-One: indicates that the information will be merged. In **All-In-One** mode, the user gets a unique, merged Environment section. This configuration is recommended for smaller and cohesive environments.

The **Alerts** and **Assets** pages are common to both modes.

Sensor update policy

This section has a toggle that lets you choose between:

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

This determines whether the sensors connected to the *CMC* will automatically receive updates when a new version of the software is available.

The **Let the user perform the update on the sensors** toggle lets you choose between:

- **Yes**, or
- **No**

Remote access to connected sensors

This section has a toggle that lets you choose between:

- **Allow**, or
- **Deny**

This lets you allow, or deny, the remote access to a sensor through the current *CMC*.

Enable multi-level focus on

A toggle lets you choose between:

- **Enabled**, or
- **Disabled**

When enabled, all downstream sensors synchronize environment information, regardless of their level in the *CMC* hierarchy. This lets you focus on those sensors directly from the current *CMC*.

Upstream connection

*A description of the **Upstream connection** section of the **General** page.*

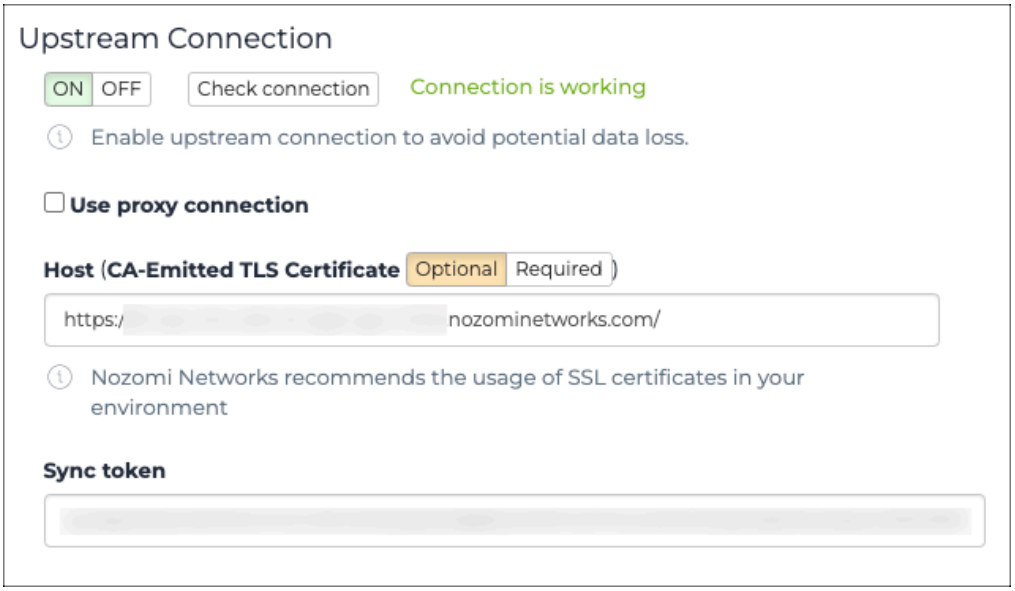


Figure 49. Upstream connection

Controls

The **Upstream connection** section has a toggle that lets you choose between:

- **ON**, or
- **OFF**

A **Check connection** button lets you check to see if the pairing between the *CMC* and the sensor is valid.

The **Use proxy connection** checkbox lets you connect to the *CMC* through a proxy server.

Host (CA-Emitted TLS Certificate)

This field is for the host address of the *CMC*. The protocol used will be *HTTPS*.

A toggle lets you choose between:

- **Optional**, or
- **Required**

You can set this toggle to **Required** if *CA*-emitted certificates are used.

Sync token

The sync(hronization) token is necessary to authenticate the connection. The pair of tokens can be generated from the *CMC*.

High Availability mode

High Availability mode (HA) lets you replicate data between two Central Management Consoles (CMC).

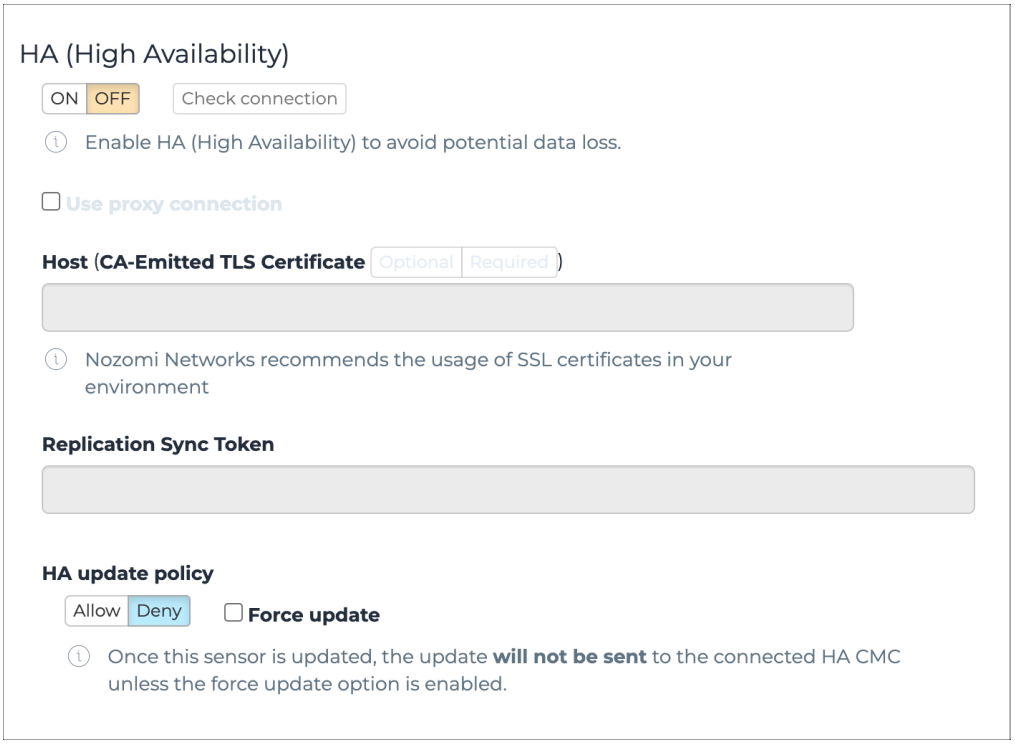


Figure 50. High Availability (HA) settings

General

To enable the highest level of resiliency, both *CMC*s must replicate each other. This is to ensure that when a *CMC* stops working, the connected sensors continue to send data to the replica *CMC*. This will also ensure that the data on each *CMC* will be kept up to date with the other *CMC*.

Passwords

When configuring *high availability (HA)*, we recommend that users choose the same admin password for both *CMC*s to avoid confusion. This is because during *HA* configuration, admin accounts are merged across both *HA* and *CMC*s and local users are synchronized. If each *CMC* has a different password for the admin account, then after *HA* configuration, only one of the passwords will work and it will be the same password for both *CMC*s.

Active Directory

Users from Active Directory are not replicated.

Threat Intelligence and Asset Intelligence

Threat Intelligence (TI) and *Asset Intelligence (AI)* contents are only available if the *CMC* has a valid license for those products (*TI* and *AI*, respectively).

Data integrations

Data integrations are not replicated. To avoid sending duplicated information, you must manually configure the same data integration on both machines, in the same exact manner. For example, with the same endpoint and options.

Replicated data

When two *CMC*s are configured to work together for *HA*, to avoid duplications and conflict, you need to perform some configurations only on one *CMC*. For more details, see Replicated data (on page 132). The *CMC* will take care of replicating the configuration options. For example, alert rules and zone configurations are replicated from one *CMC* to another.

Failover functionality

When one *CMC* fails, sensors will automatically fail over to the replicated *CMC*.

HA update policy

This section lets you control these settings:

- **Allow** the software update to propagate to the *HA* partner *CMC*

  Note:
  While this is set to **Allow**, the *CMC* running the newer version of *N2OS* will propagate the update bundle to the other *CMC*.

- **Deny** the software update from propagating to the *HA* partner *CMC*

  Note:
  While this is set to **Deny**, the *CMC* running the newer version of *N2OS* will not propagate the update bundle to the other *CMC*.

- **Force update** to propagate to the *HA* partner *CMC*

  Note:
  The *CMC* running the newer version of *N2OS* will propagate the update bundle to the other *CMC* just once. As soon as the update is completed, the **Force update** option will be automatically disabled.

Replicated data

A list of data types that are shared between CMCs in High Availability (HA) mode.

Alert
AlertEvent
AlertEventsAlert
AlertPlaybooks
AlertRule
Appliance
Assertion
AssertionGroup
Asset
AssetTypes
AuditItem
CapturedLog
DashboardConfiguration
HealthLog
MigrationTask
NodeCpe
NodeCpeChange
NodeCustomField
NodePoint
NodePointStatus
Report
ReportFile
ReportFolder
ScheduledUpdate
SpExecution
SpNodeExecution
SpNodeExecutionStatus
SpPlan
User
UserGroup
ZoneConfiguration

Configure high availability mode

To use high availability (HA) mode, you must first configure the two Central Management Consoles (CMC). To do this, you must copy the relevant details of each CMC to the other one, and then allow them to communicate with each other.

About this task

To successfully configure *HA* mode in the two *CMC*s, you need to:

- Set *HA* to **ON**
- Copy the *IP* address into the **Host** field of the second *CMC*
- Copy the **Sync token** into the **Replication Sync Token** field of the second *CMC*
- Do the above steps again in the second *CMC*
- **Allow** the two *CMC*s to communicate with each other

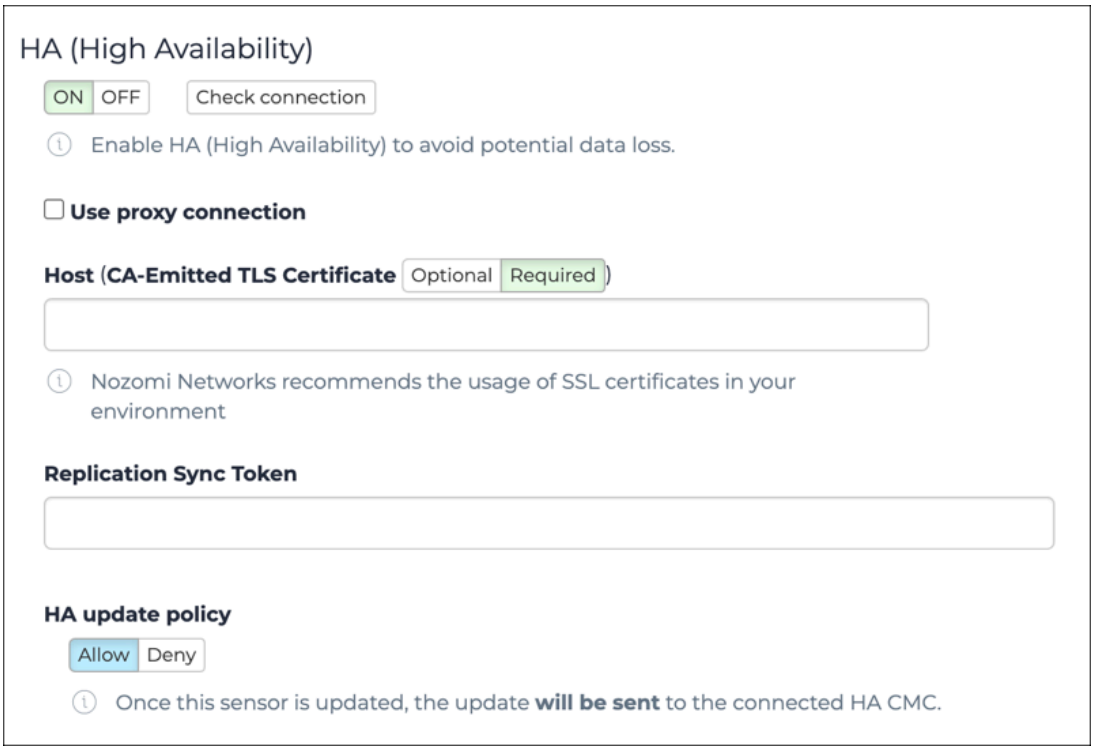
Procedure

1. Open the first *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 right section, select **General**.

   Result: The **General** page opens.
5. At the top of the **High Availability** section, select **ON**.



6. In the **Host** field, enter the *IP* address of the *CMC* that you want to share data with.

   > **Note:**
   > If a *CA*-signed *transport layer security (TLS)* certificate is not being used, you should select **Optional** instead of **Required**.

7. In the **Replication Sync Token** field, enter the **Sync token** of the *CMC* that you want to share data with.
8. At the top of the **High Availability** section, select **Check connection**.

   Result: If the connection is okay, a `Connection is working` message will show in green text.

9. Select **Save**.
10. Do the above steps again in the other *CMC*.
11. Wait for the two *CMC*s to communicate with each other.

    > **Note:**
    > This should take approximately five seconds.

12. In the **Allow remote to replicate on this CMC** field of the first *CMC*, the **Sensor ID** of the second *CMC* should show.
13. In this section, select **Allow**.
14. In the **Allow remote to replicate on this CMC** field of the second *CMC*, the **Sensor ID** of the first *CMC* should show.
15. In this section, select **Allow**.
16. Select **Save**.
17. **Optional:** If necessary, Check high availability mode is working correctly (on page 135)

Results

The *CMC*s have been configured in *HA* mode.

Check high availability mode is working correctly

Once you have configured high availability (HA) mode, you can do a check to make sure that it is working correctly.

Procedure

1. In the top navigation bar, select 

   Result: A menu shows.

2. In the **System** section, select **Health**.

   Result: The **Health** page opens.

3. In the top right section, select **Performance**.
4. In the **Replication status** section, make sure that all the entities are replicated correctly.

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*.

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.

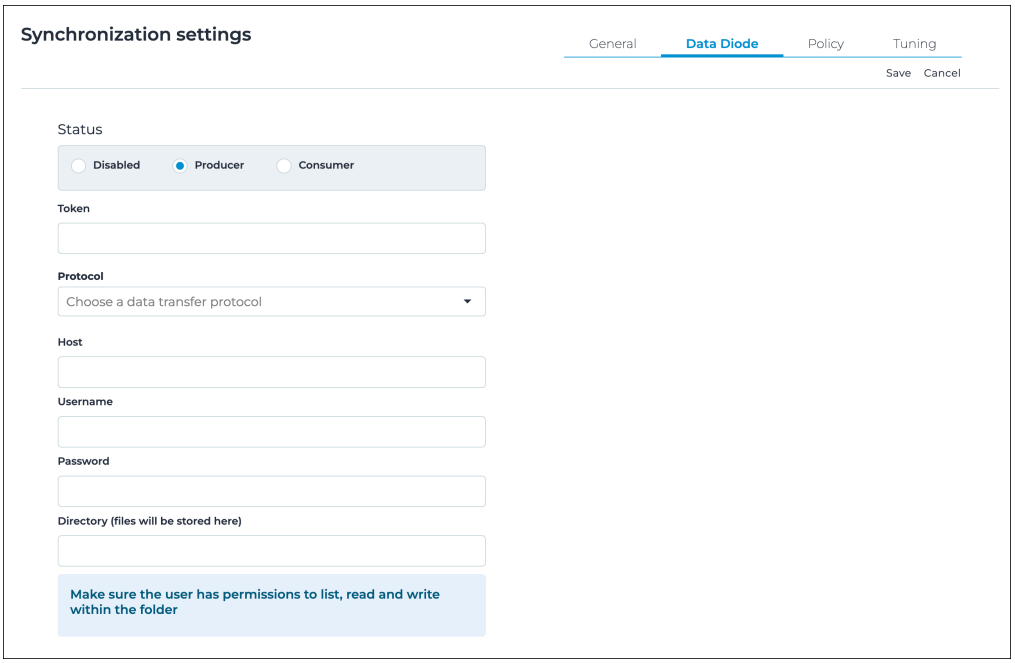


Figure 51. Data Diode page

Disabled

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

Producer

When this checkbox is enabled, it allows the *CMC* 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.

Consumer



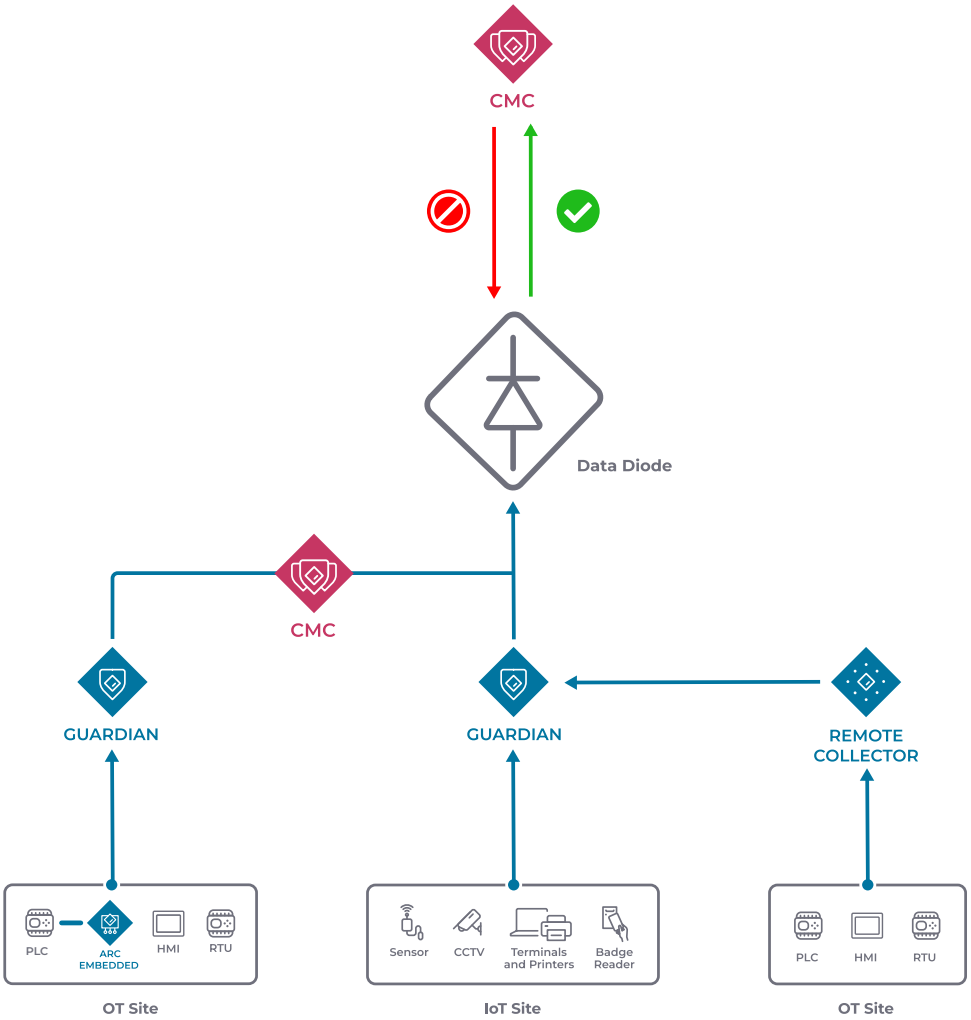
Figure 52. Consumer settings

When you select the **Consumer** checkbox, the **Token** field shows. This lets you copy the token to be used in the **Producer** settings of the downstream Guardian(s).

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 *industrial control systems (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 142), and
- Configure a CMC as a producer (on page 145), or
- Configure a Guardian as a producer (on page)

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 (*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 145) or the Configure a Guardian as a producer (on page) procedure.

Configure a CMC as a producer

Configure a Central Management Console (CMC) to communicate to an upstream data diode and 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 *CMC* as a data producer in a secure, one-way data transfer setup. This configuration enables secure, one-way communication with an upstream *CMC* through a data diode.

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 *CMC*.
7. In the top navigation bar, select ⚙

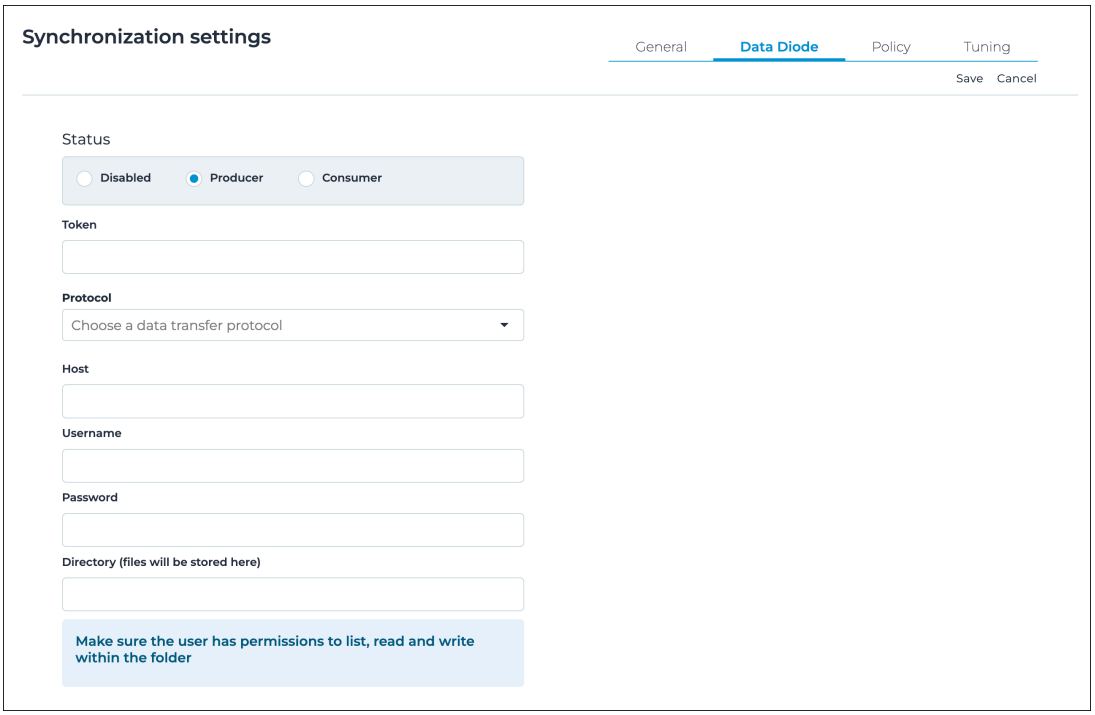
   Result: The administration page opens.
8. In the **Settings** section, select **Synchronization settings**.

   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.

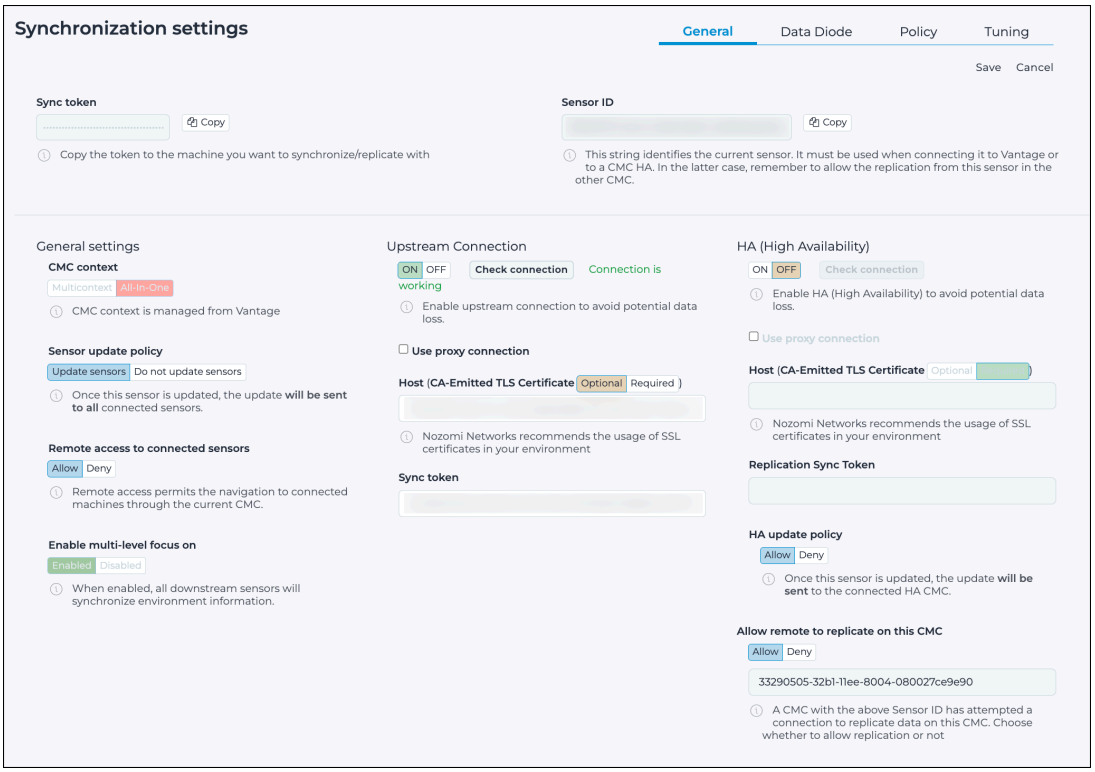
11. In the **Token** field, paste the consumer token that you copied in step 5 (on page 146).



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.

Policy

*The **Policy** page lets you centrally configure the data synchronization settings for Central Management Consoles (CMC) and sensors.*

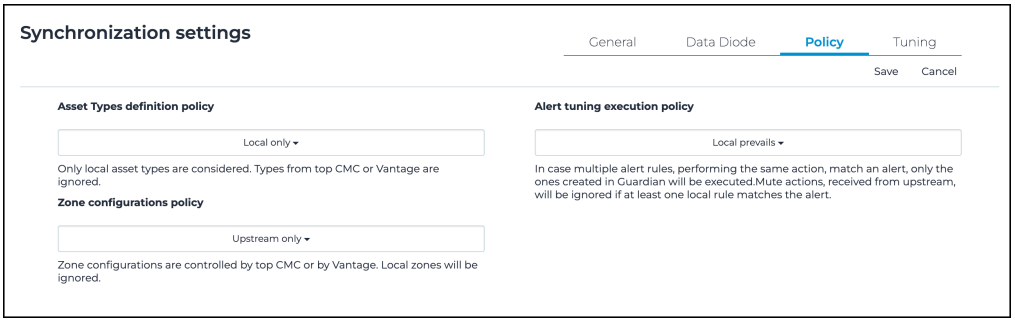


Figure 53. Policy page

General

Guardian and *CMC* deployments each have their own configurations. To simplify management of sensors that are connected to an upstream sensor, centralized configuration is available for:

- Asset types
- Alert rules
- Zone configurations
- Threat Intelligence

Asset Types definition policy

This dropdown lets you choose between these policies:

- Local only
- Upstream only

Local only: Only local asset types are considered. Types from Vantage, or the top *CMC*, as distributed to the connected sensors are ignored.

Upstream only: Asset types are imported by top *CMC* or by Vantage. Local asset types configured on any sensor will be ignored.

Alert Tuning execution policy

This section lets you specify a synchronization policy for alert rules from the *CMC*.

The dropdown lets you choose between these policies:

- Upstream only
- Upstream prevails
- Local prevails

Upstream only: Vantage, or the top *CMC* controls alert rules. Local alert rules configured on a sensor will be ignored.

Upstream prevails: In case of multiple overlapping alert rules, performing the same action match an alert, only the ones received from the upstream *CMC*/Vantage will be executed. Mute actions, created in Guardian, will be ignored if at least one rule, received from upstream, matches the alert.

Local prevails: In case of multiple overlapping alert rules, performing the same action, match an alert, only the ones created in Guardian will be executed. Mute actions, received from upstream, will be ignored if at least one local rule matches the alert.

Zone configurations policy

This section lets you specify a synchronization policy for zone configurations from the *CMC*.

The dropdown lets you choose between these policies:

- Local only
- Upstream only
- Apply all

Local only: Guardian controls zone configurations. Zones received from upstream will be ignored.

Upstream only: Vantage, or the top *CMC* controls the zone configurations. Local zones will be ignored.

Apply all: You can manage zone configurations at every level. All zones are applied.

Threat Intelligence contents management

This section lets you enable sensors to provide the user with the ability to enable and disable individual Threat Intelligence contents, and to define custom contents.

If the **Enable sensor contents** toggle is set to off (default), then customization of the Threat Intelligence contents can only be done from the upstream *CMC*/Vantage.

Tuning

A description of how you can enable, or disable, data synchronization for entities in the Central Management Console (CMC).

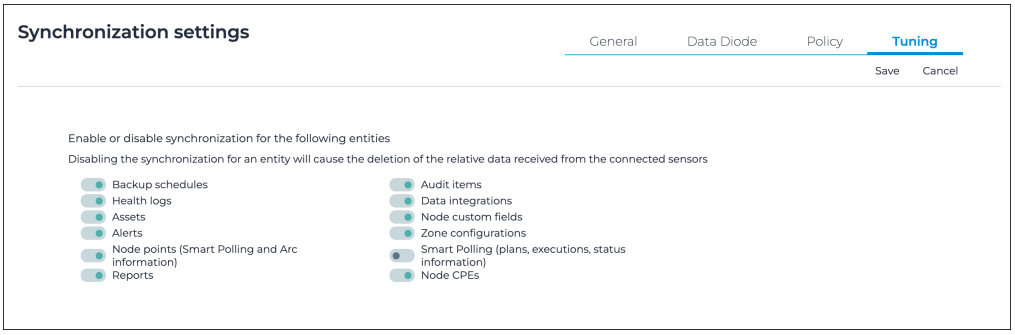


Figure 54. Tuning page (Multicontext mode)

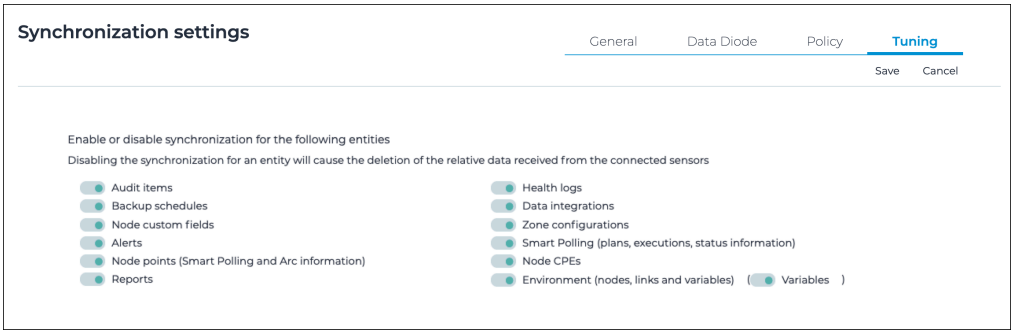


Figure 55. Tuning page (All-In-One mode)

You can configure synchronization in the *CMCs* in: **> Synchronization settings > Tuning**.

You can enable or disable the synchronization for the following entities:

Note:
Some entities are only available in one of the two modes.

- Assets (Multicontext mode only)
- Audit items
- Backup schedules

Note:
When you enable **Backup schedules** in the *CMC*, the backup schedules of downstream sensors will not show in the *CMC*. This feature is only available in Vantage.

- Node custom fields
- Alerts
- Node points (Smart Polling and Arc information)
- Reports
- Health logs
- Data integrations
- Zone configurations
- Smart Polling (plans, executions, status information)
- Node CPEs
- Environment (nodes, links and variables) (All-In-One mode only)
- Variables (All-In-One mode only)

The configuration is applied only to sensors directly connected to the *CMC* in which the configuration has been set. If the *CMC* has an *HA* connected, the tuning must be configured in both the *CMC*s.

Disabling synchronization for an entity will cause the deletion of all the items in the *CMC* already received from the respective sensors.

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.*

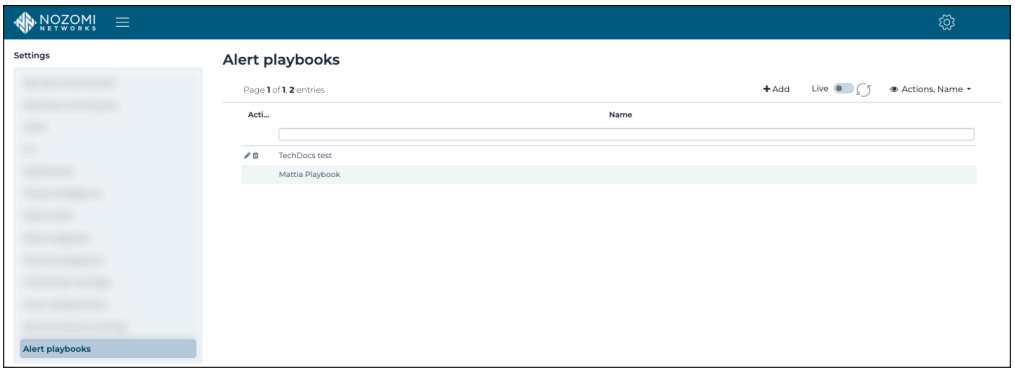


Figure 56. Alert playbooks page

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.

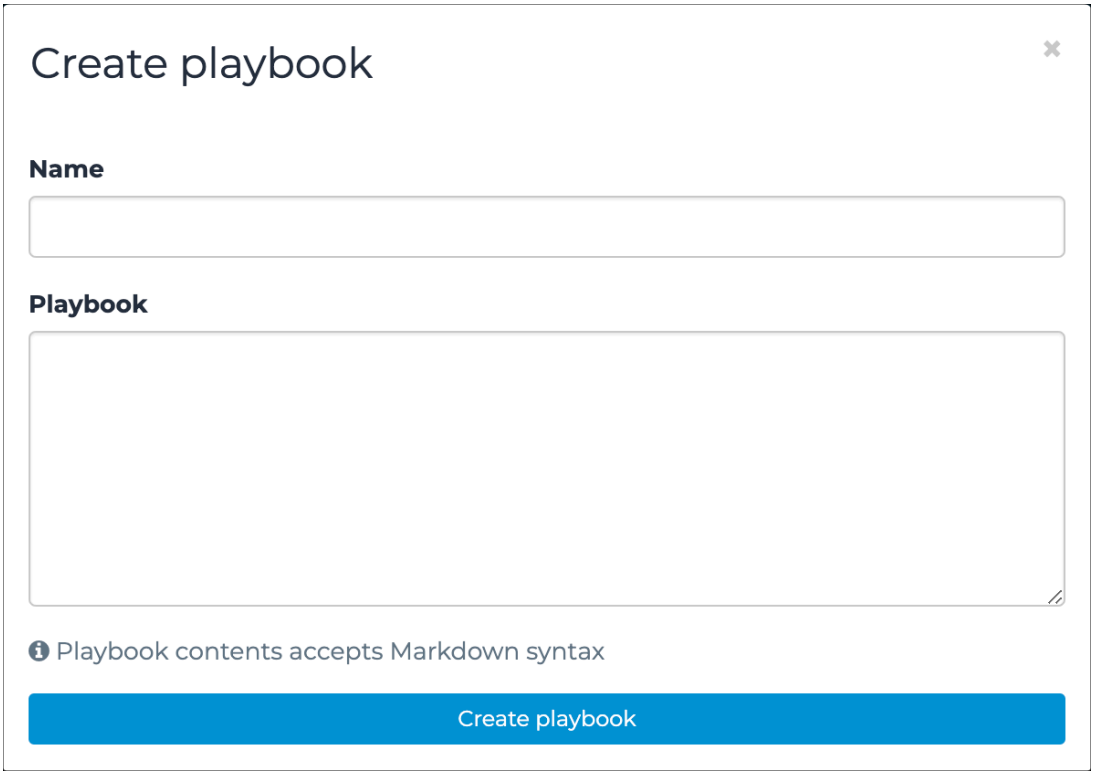
2. In the **Settings** section, select **Alert playbooks**.

   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.

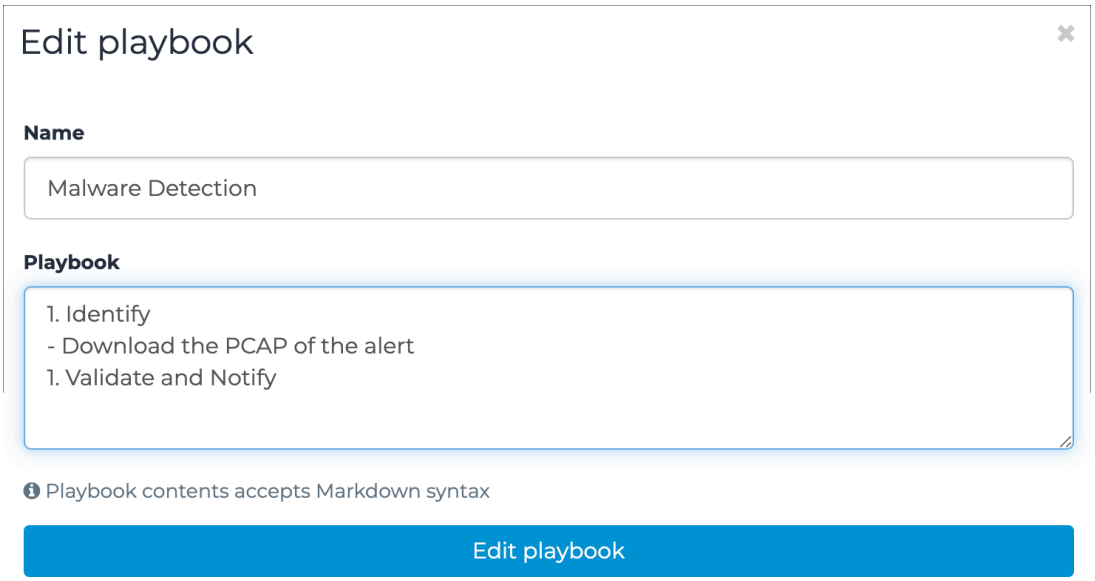
   Note:
   The list will be empty if:
   - No alert playbooks have been created
   - One or more alert playbooks have been created, but they were created in an upstream Vantage or *CMC*

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.*

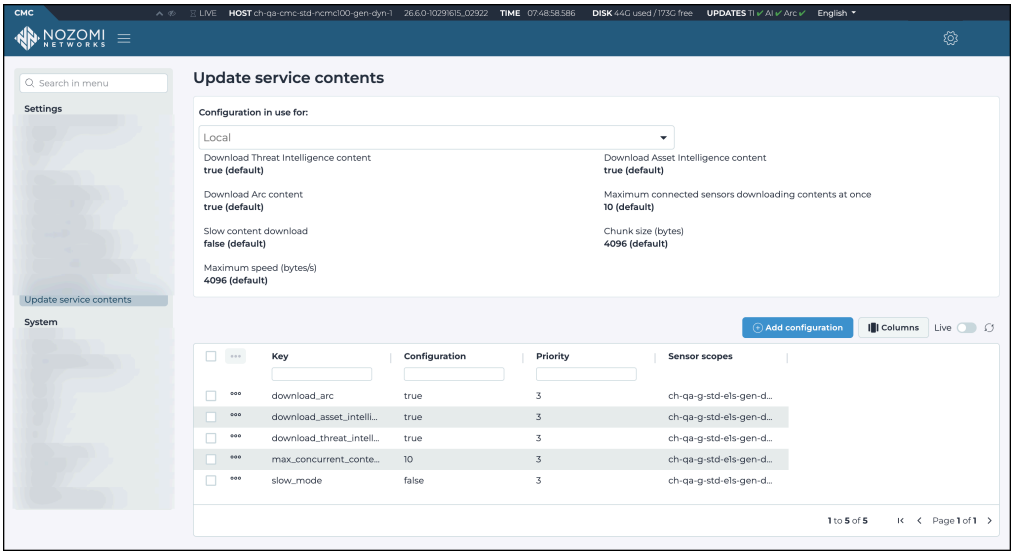


Figure 57. Update service contents page in CMC

The page has these sections:

- **Configuration in use**: The configuration currently applied to the appliance. On a *CMC*, select the appliance in the **Configuration in use for** list. When no appliance is selected, the configuration in use on the current machine, labeled **Local**, shows
- **Configurations table**: The custom configurations that override the default values for the appliances in their scopes. Custom configurations are managed from a *CMC*. For more details, see Manage update service contents configurations (on page 161)

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.



Note:

Configurations pushed by a parent *CMC* are read-only on the appliances it manages. They cannot be edited or deleted there.

Manage update service contents configurations

You can add, edit, or delete custom update service contents configurations from the CMC to override the default download behavior for specific appliances.

About this task

Only appliances and *CMC*s can be included in a configuration, as they are the machines that download the contents. For more details on the settings, see Update service contents (on page 159).

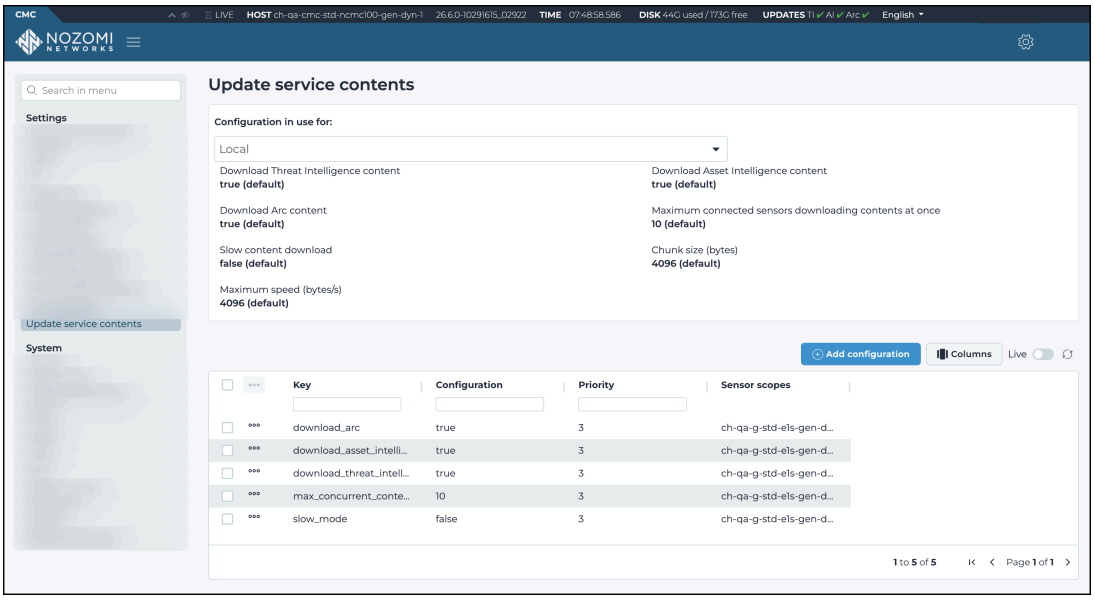
Procedure

1. In the top navigation bar, select ⚙

   Result: The administration page opens.

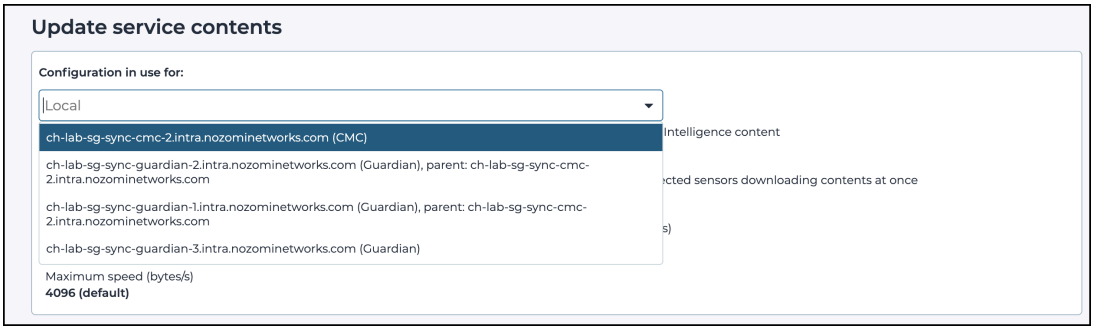
2. In the **Settings** section, select **Update service contents**.

   Result: The **Update service contents** page opens and shows the configuration in use on the current machine.



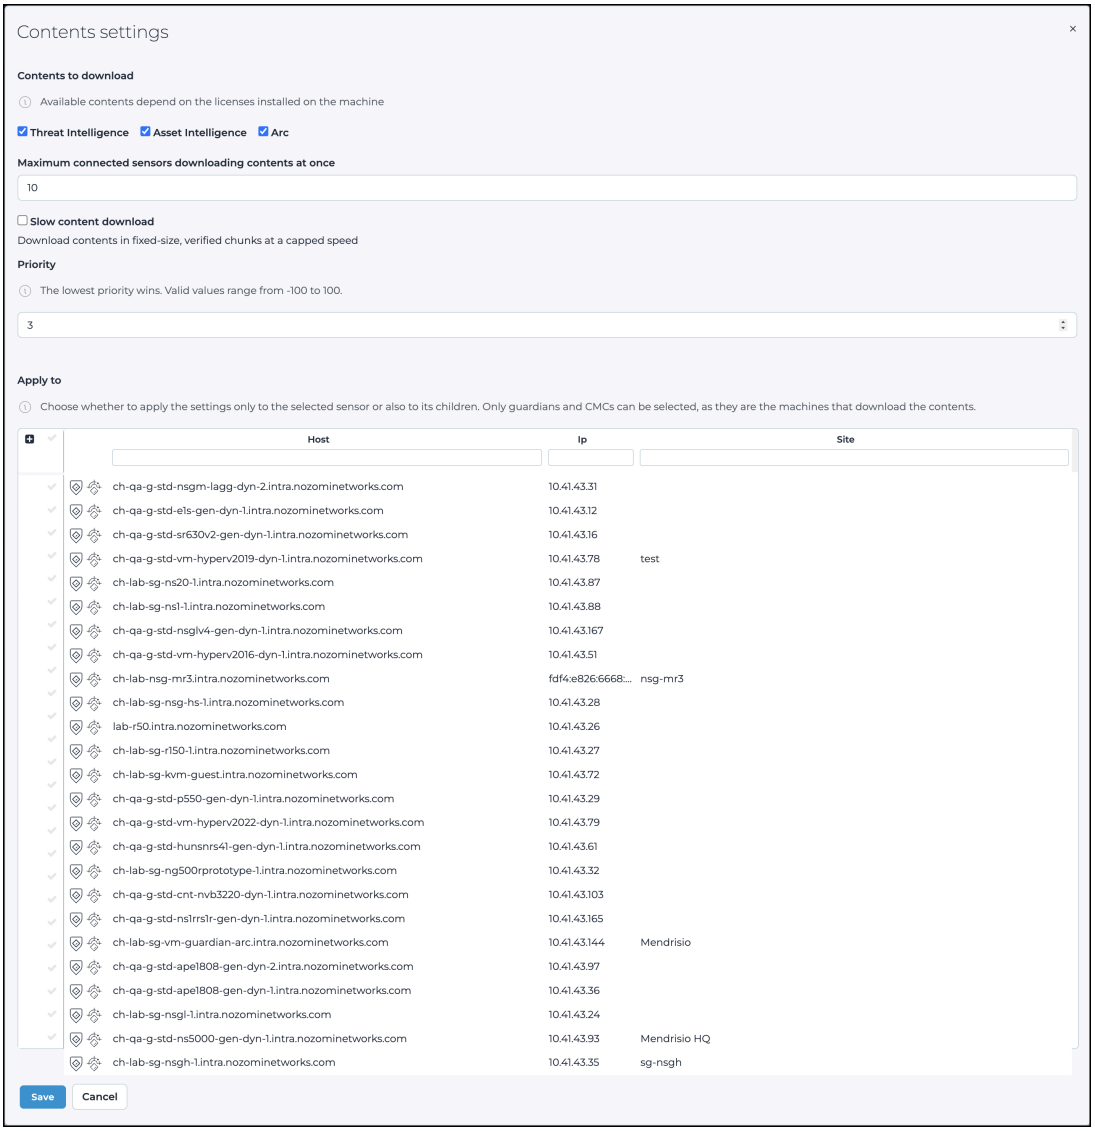
3. **Optional:** To open this page from the **Sensors** page, select **Contents settings** in the toolbar of the **List** tab.
   The action is available next to **Allow all filtered sensors**.

4. **Optional:** To see the configuration in use on a specific appliance, select it in the **Configuration in use for** list.
   The list shows the connected appliances with their type and parent. Select **Local** to go back to the configuration in use on the current machine.



5. Select **Add configuration**.

   Result: The **Contents settings** dialog shows.



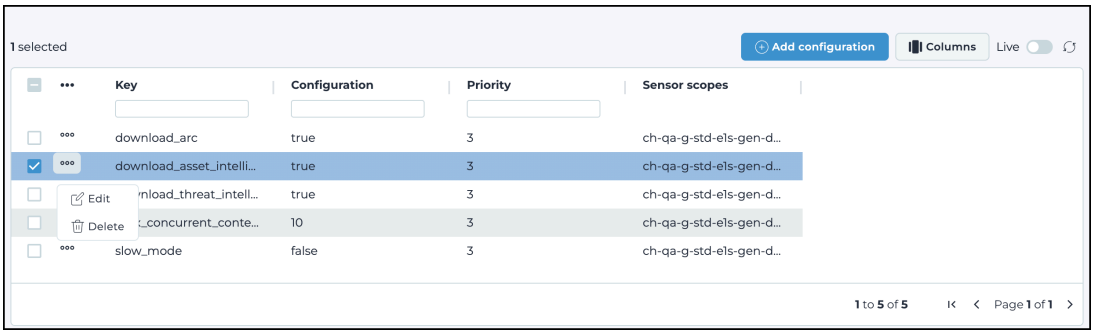
6. In **Contents to download**, select the types of content that the configuration downloads.
7. **Optional:** Select **Slow content download** and enter the **Maximum speed (bytes/s)** and the **Chunk size (bytes)**.
8. Enter the **Priority** of the configuration.
   The configuration with the lowest priority value applies to an appliance. Valid values range from -100 to 100.
9. In **Apply to**, select the appliances that the configuration applies to.
   If you do not select any appliance, the configuration applies to all the appliances.

10. Select **Save**.

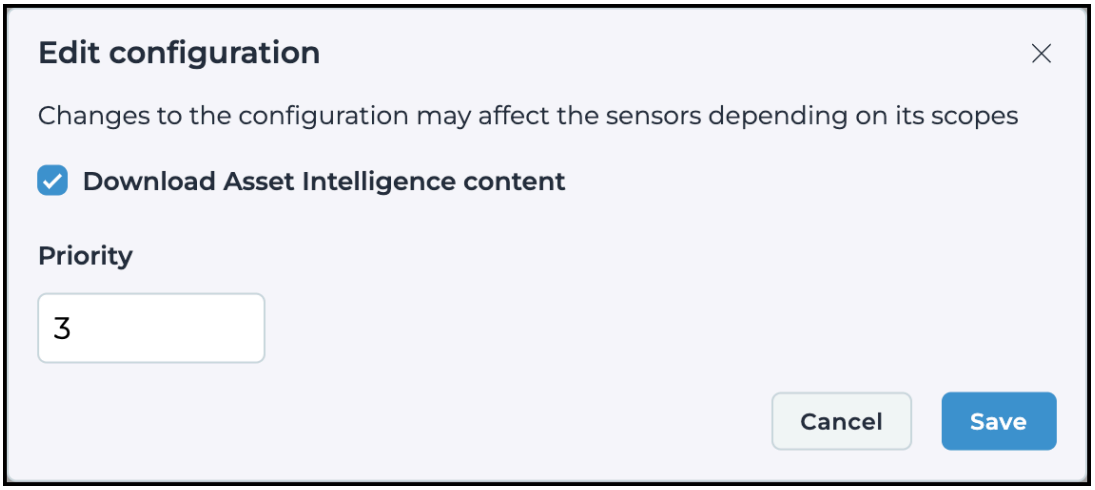
    Result: The configuration shows in the configurations table.

11. **Optional:**

    To edit or delete a configuration, select the three-dots menu next to its row and select **Edit** or **Delete**.



    If you select **Edit**, the **Edit configuration** dialog shows. You can change the value of the configuration and its priority.



    If you select **Delete**, confirm the deletion when prompted. The deletion is propagated to the machines that hold a copy of the configuration.
    To delete multiple configurations, select their rows and select **Bulk actions**. Confirm the deletion when prompted.

Results

The configuration applies to the appliances in its scopes. The update service applies the new values the next time it runs.

System

General

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

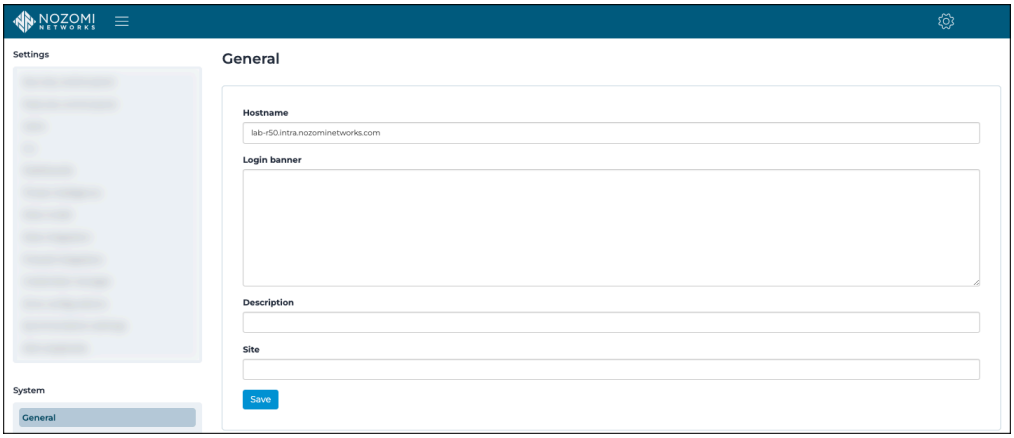


Figure 58. General page

Date and time

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

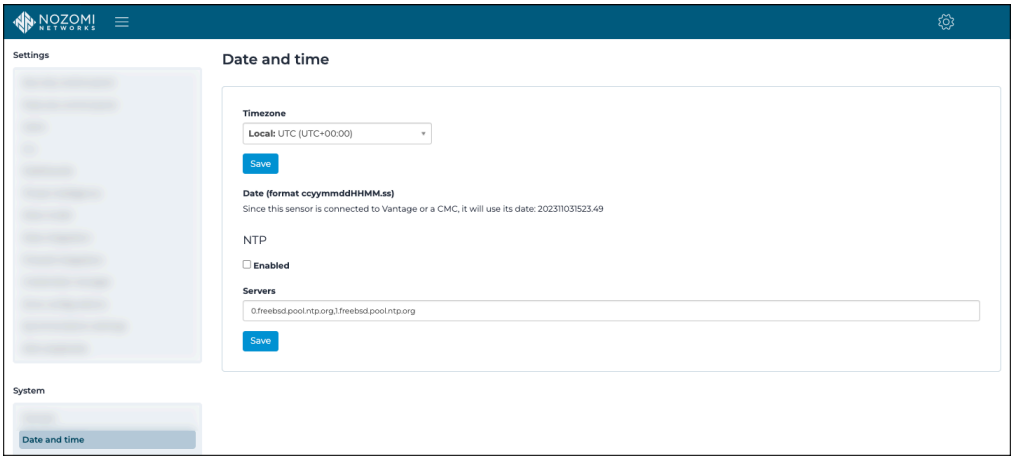


Figure 59. Date and time page

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 for this CMC and the sensors connected to it.*

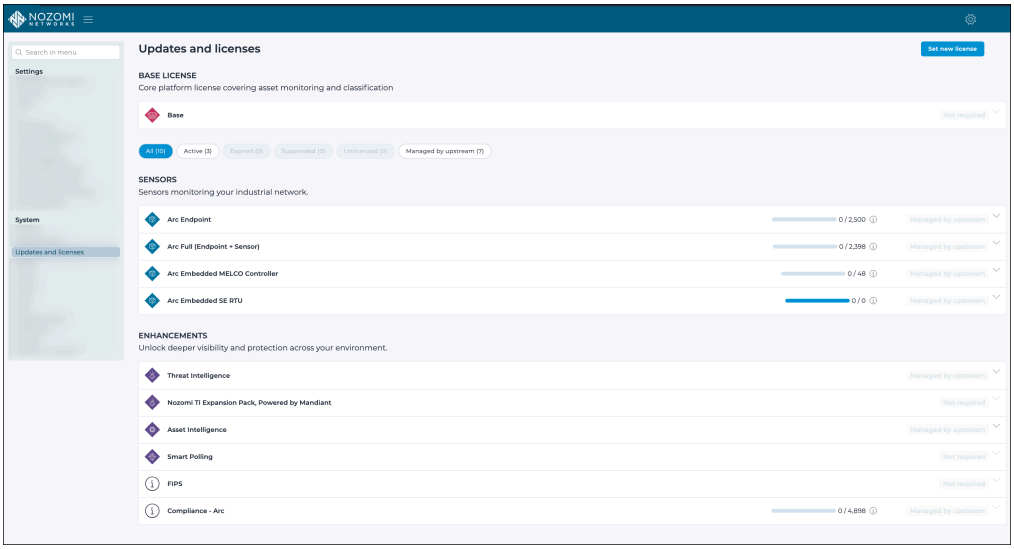


Figure 60. Updates and licenses page

For more details about the license types shown on this page, see License types (on page 168).

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.

Licenses managed by upstream

On a *CMC*, most licenses show a status of **MANAGED_BY_UPSTREAM**: the license is received from an upstream source, such as Vantage or a parent *CMC*, and is managed there. This *CMC* adopts the license as received, rather than managing it directly.

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 6. License status

License status	Description
UNLICENSED	Functionality is disabled.

Table 6. License status (continued)

License status	Description
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.

Export

*The **Export** page lets you export data in content packs.*

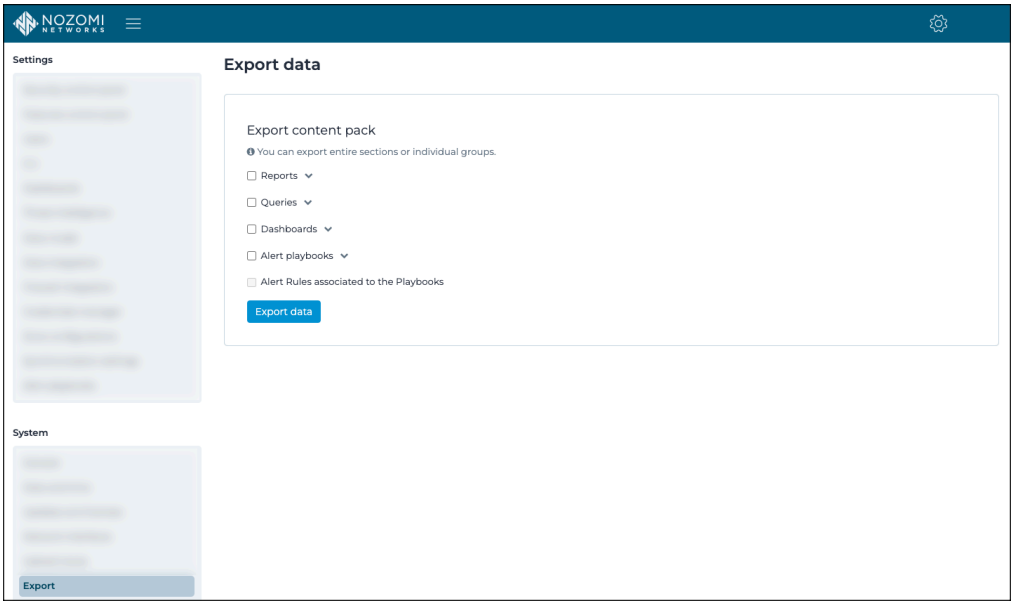


Figure 61. Export page

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.*

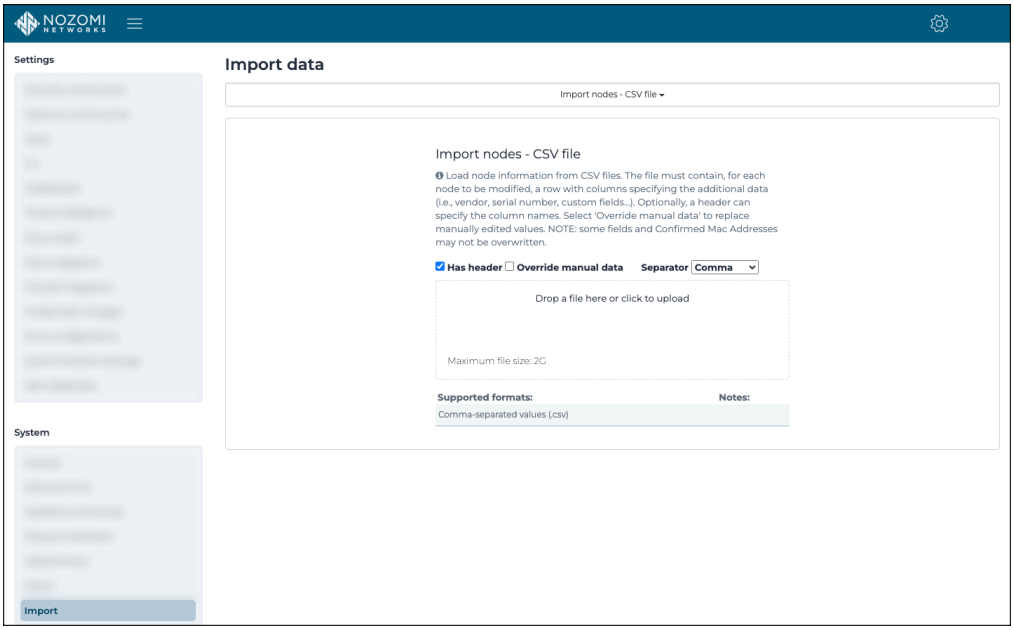


Figure 62. Import page

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 7. 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 171).

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**.

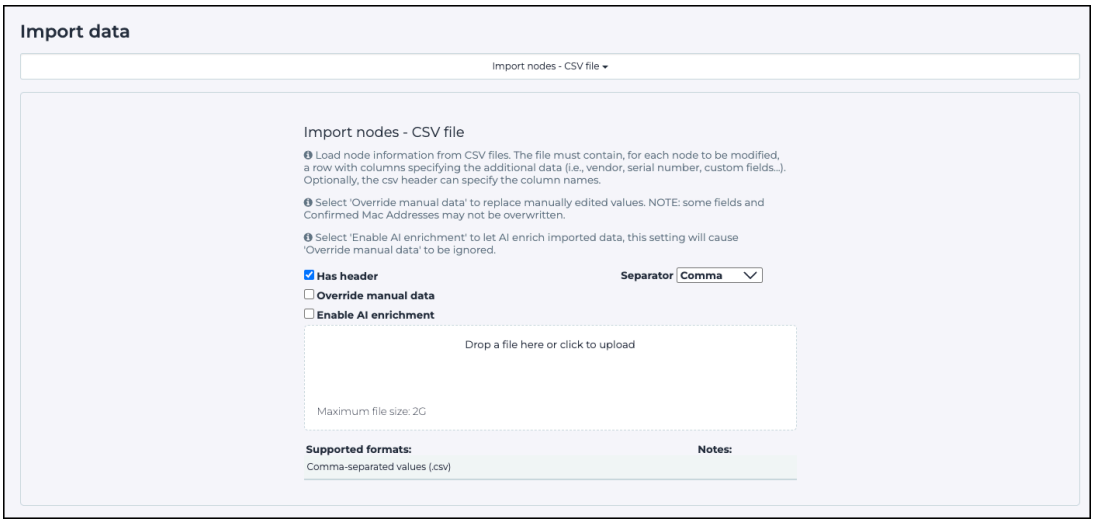
   Result: The **Import nodes - CSV file** section shows.

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**.

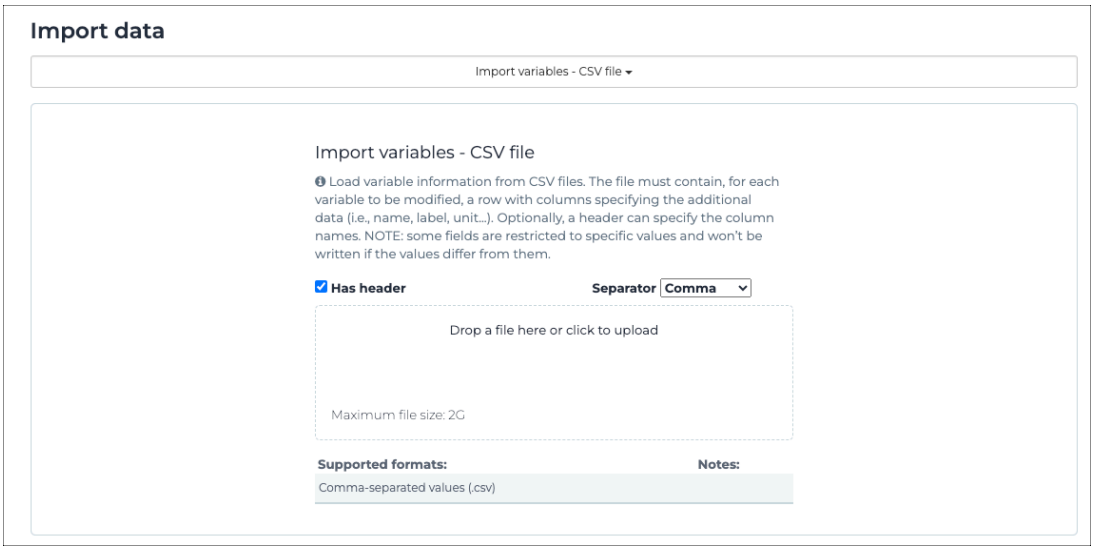
   Result: The **Import variables - CSV file** section shows.

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**.

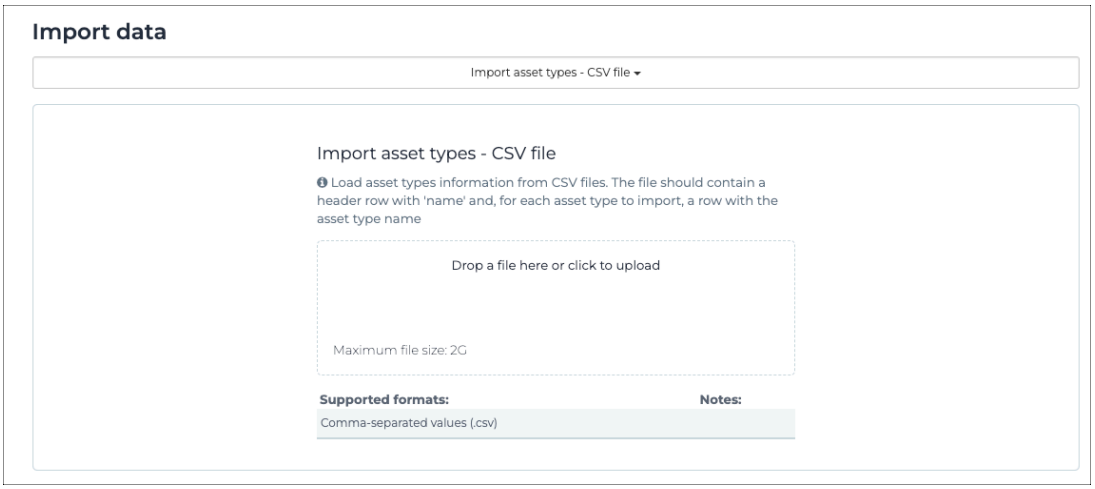
   Result: The **Import asset types - CSV file** section shows.

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**.

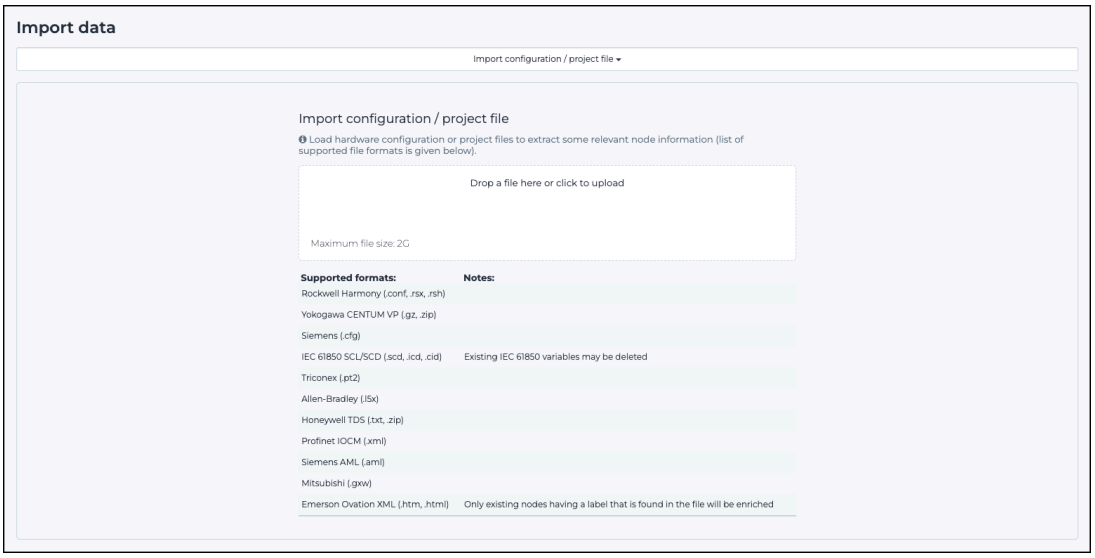
   Result: The **Import configuration / project file - CSV file** section shows.

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 184).

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 Central Management Console (CMC).

Table 8. 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 8. 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**.

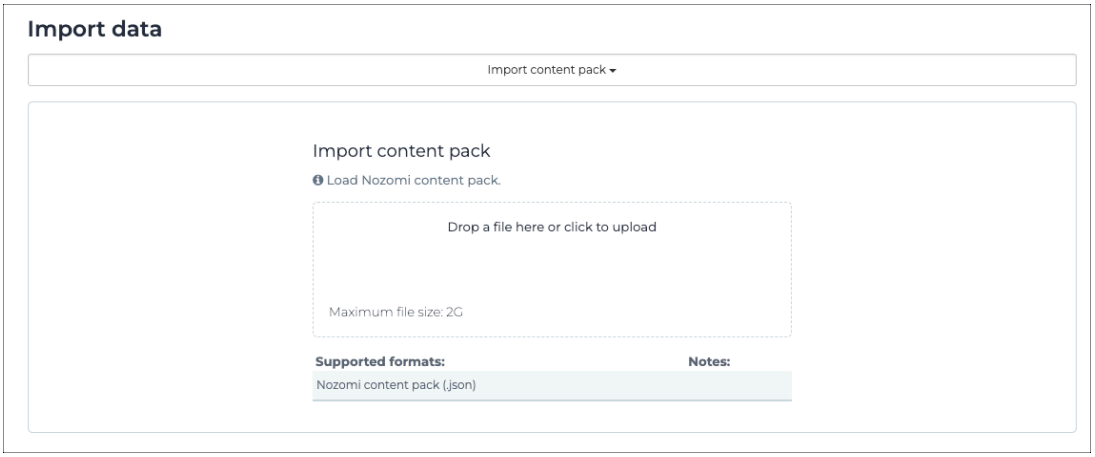
   Result: The **Import content pack** section shows.

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**.

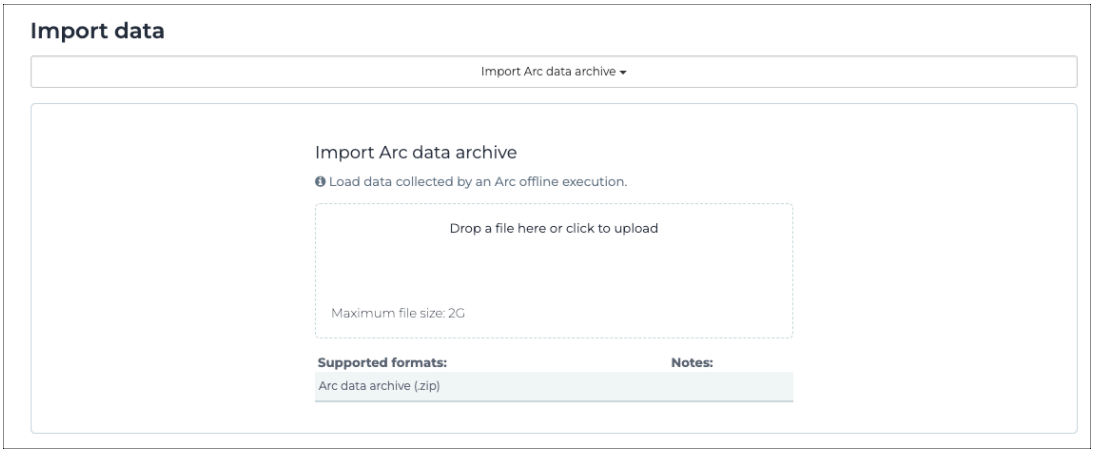
   Result: The **Import Arc data archive** section shows.

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.*

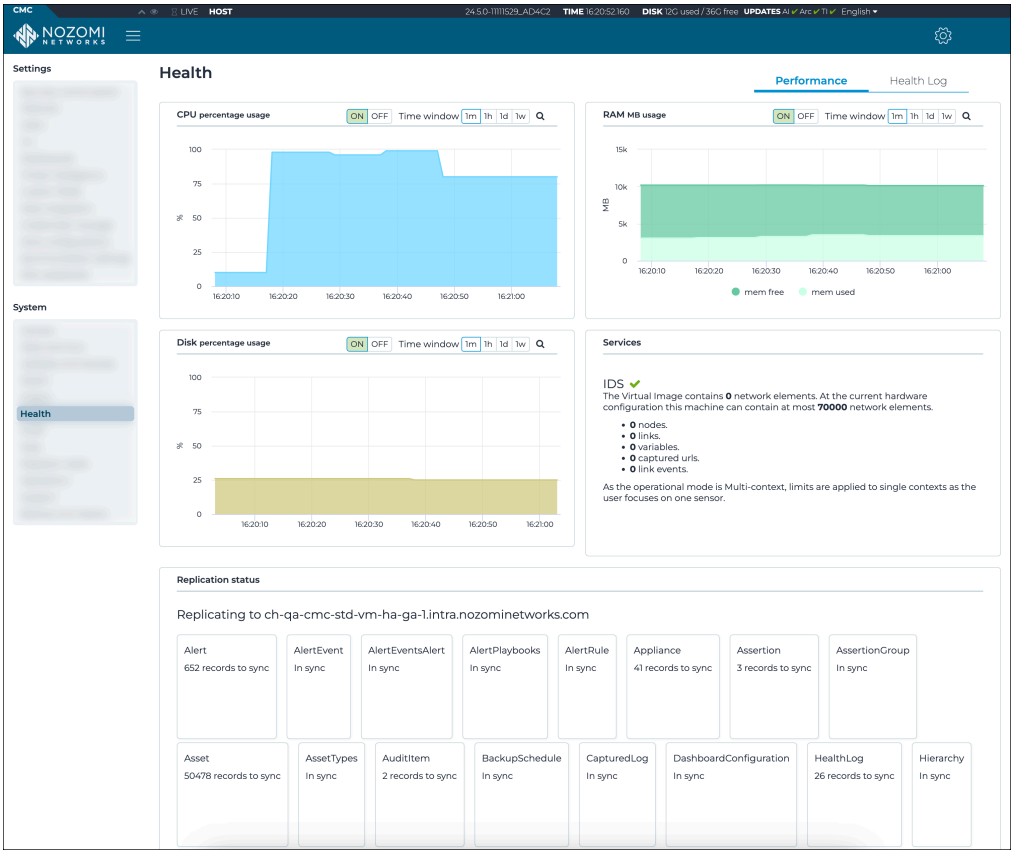


Figure 63. Health page

The Health page has these tabs:

- Performance (on page 189)
- Health log (on page 191)

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.*

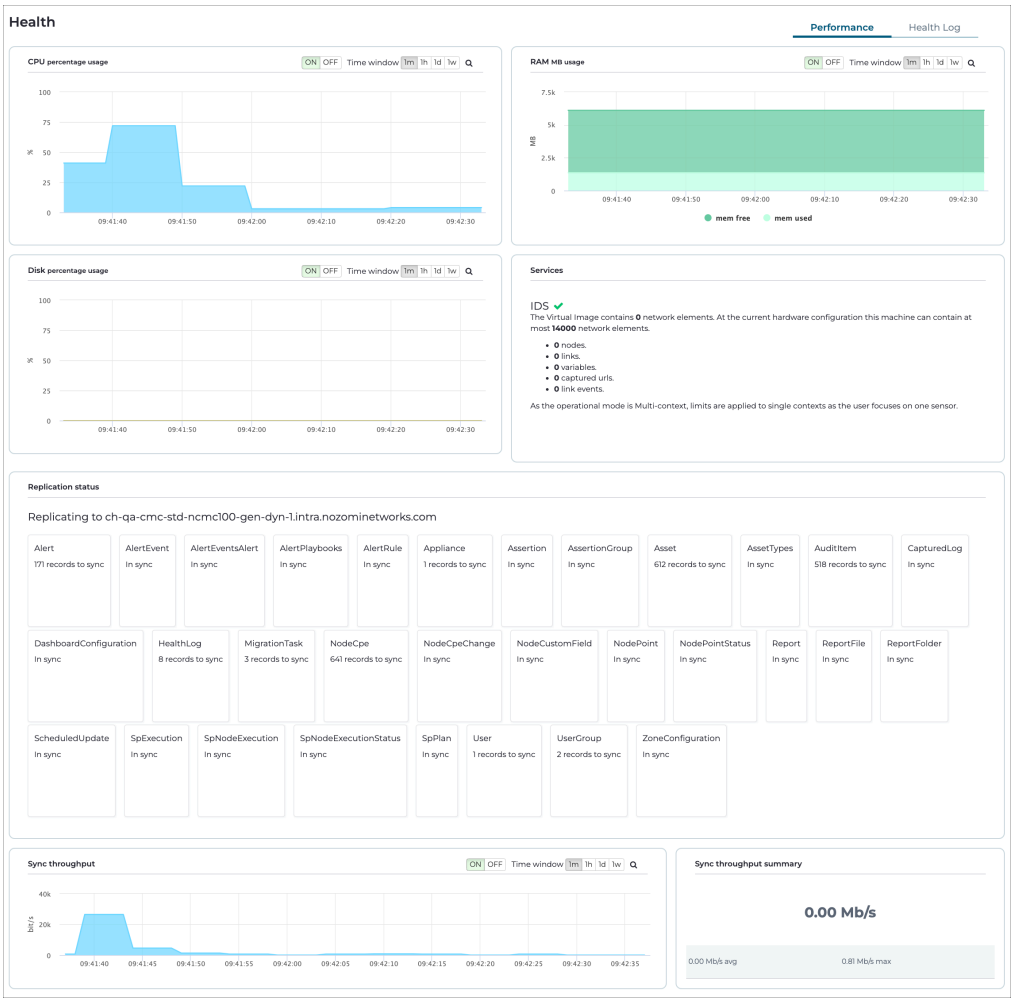


Figure 64. Performance page

Graphs

The page shows graph for:

- CPU percentage usage
- RAM MB usage
- Disk percentage usage
- Services
- Replication status
- Sync throughput
- Sync throughput summary

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.*

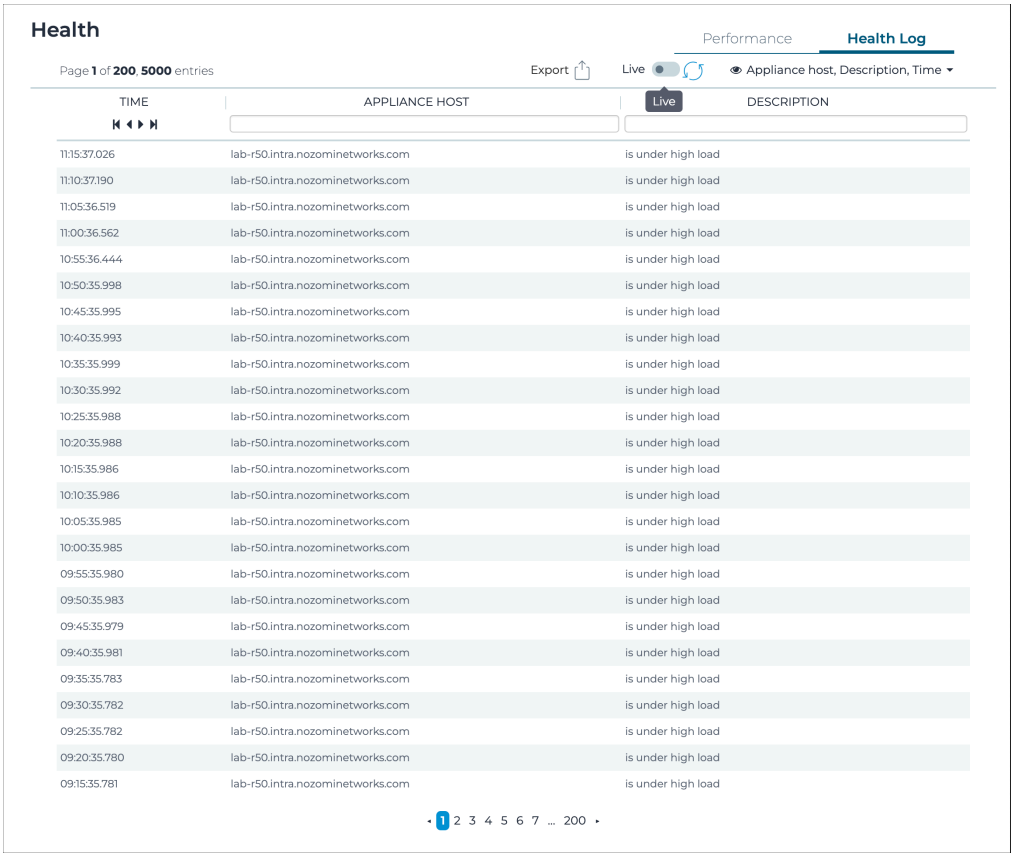


Figure 65. Health log page

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.*

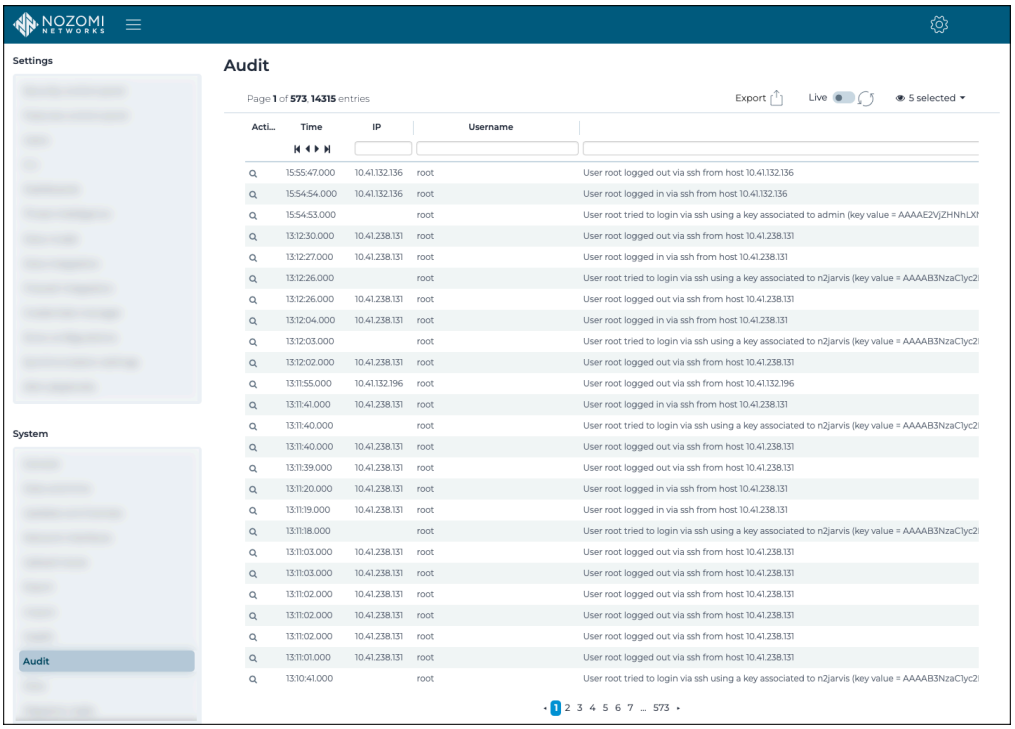


Figure 66. Audit page

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.*

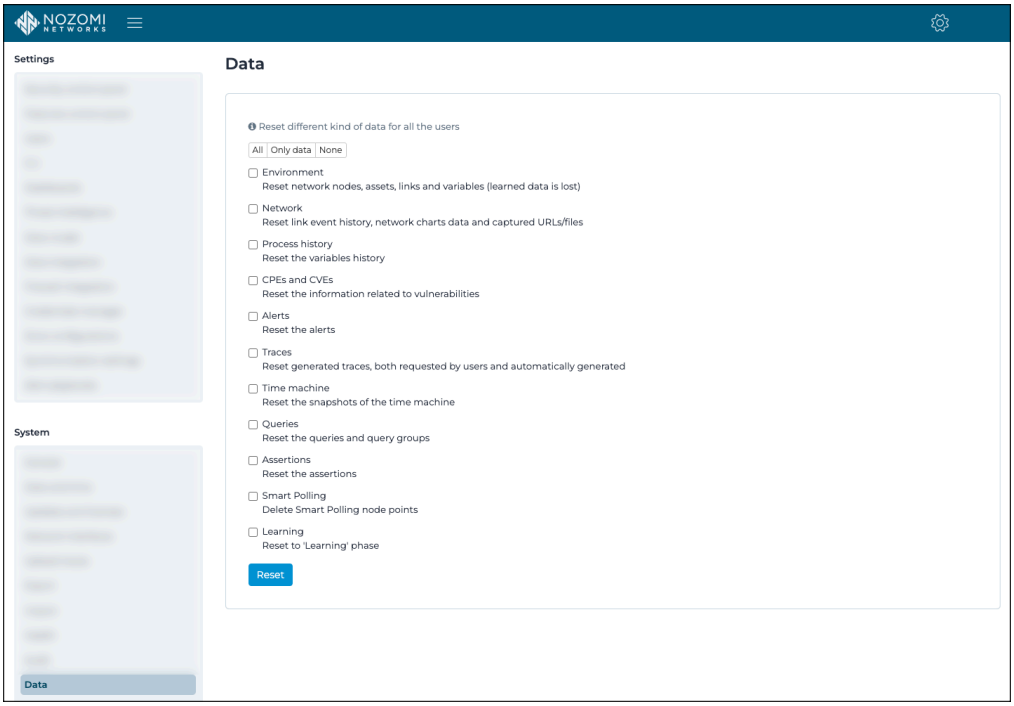


Figure 67. Data page

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.*

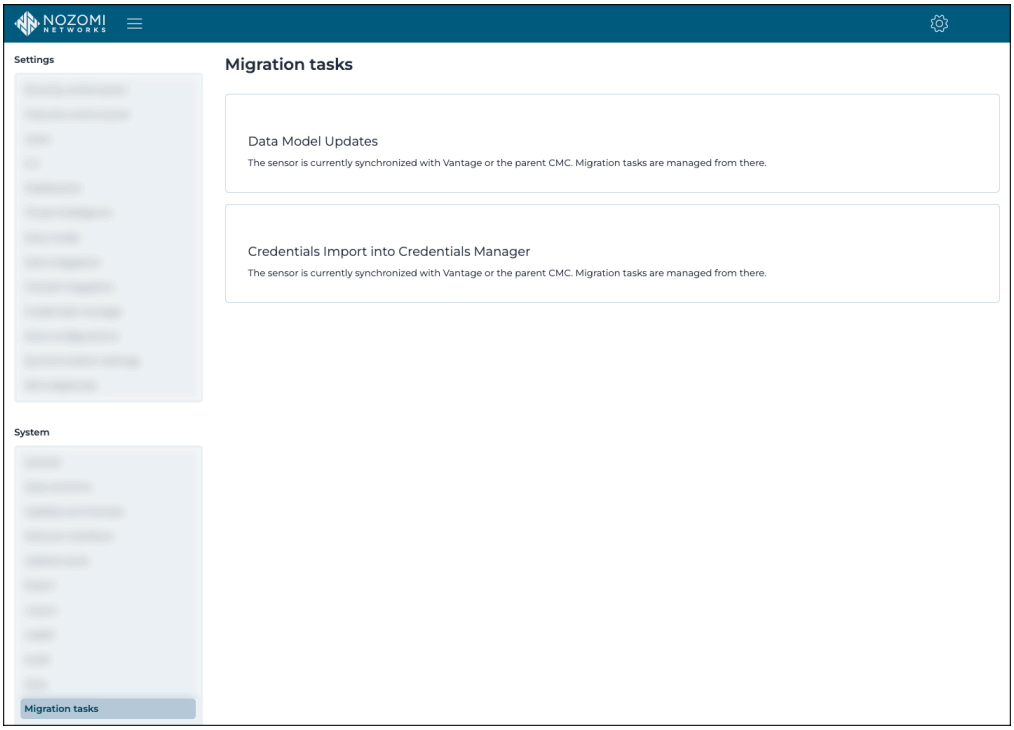


Figure 68. Migration tasks page

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.*

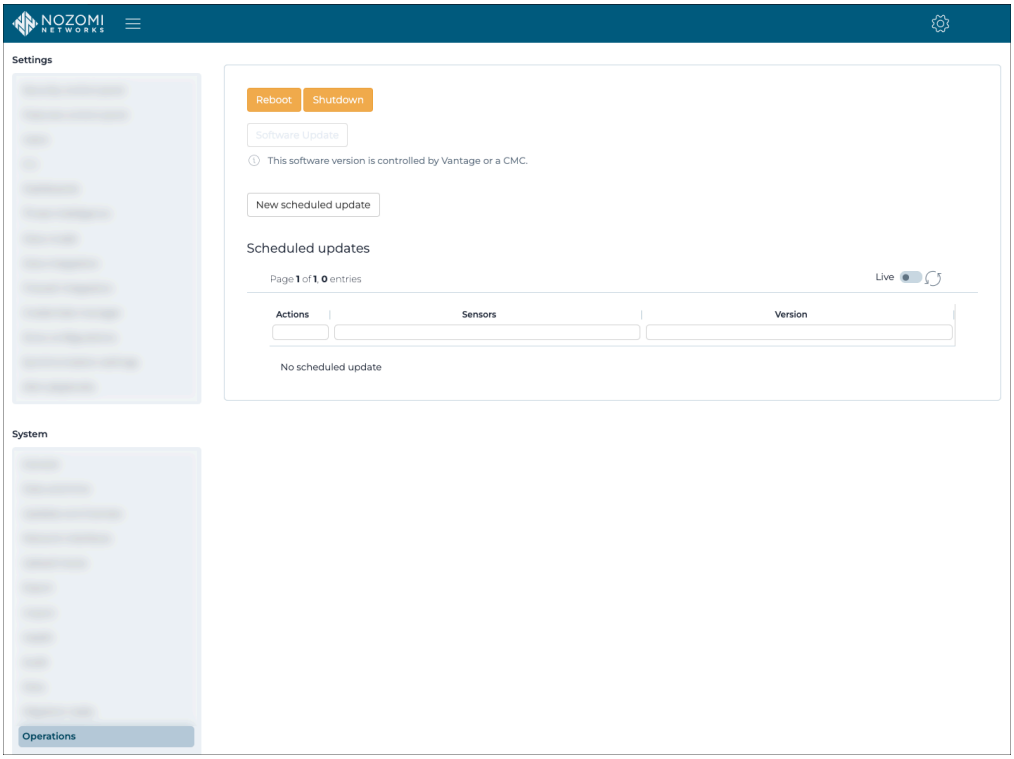


Figure 69. Operations page

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.

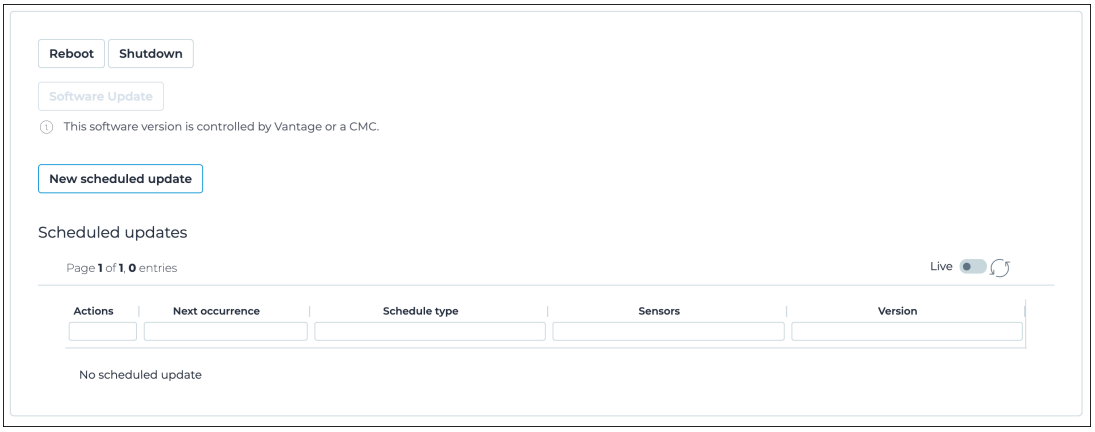
2. In the **System** section, select **Operations**.

   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.*

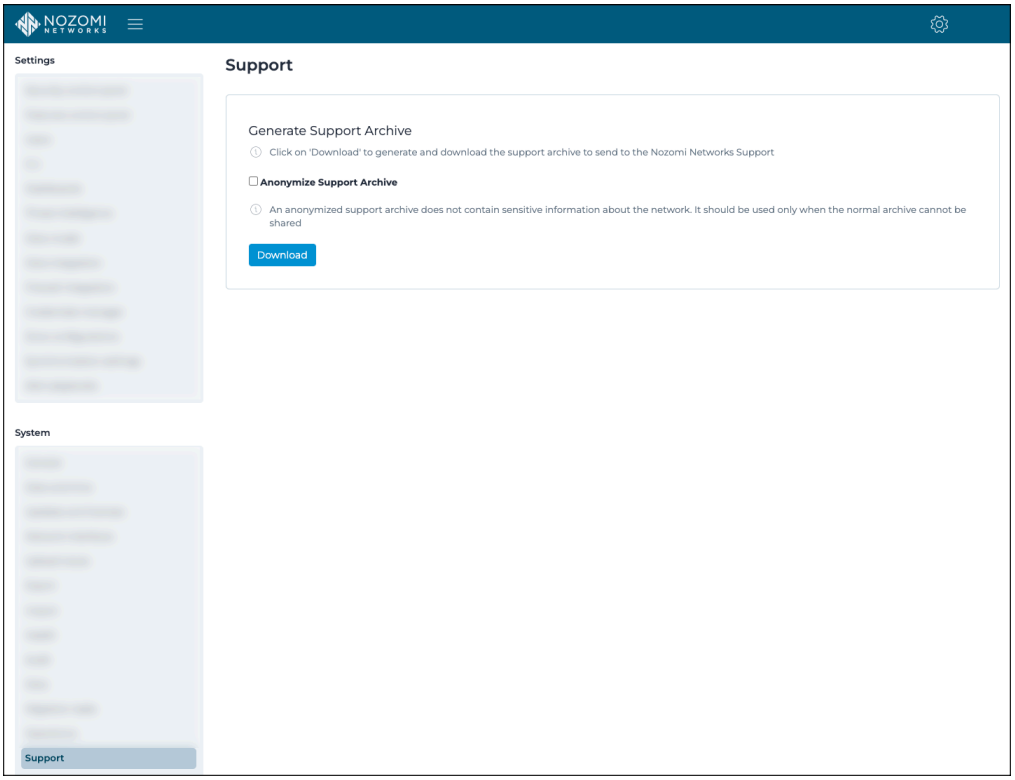


Figure 70. Support page

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.*

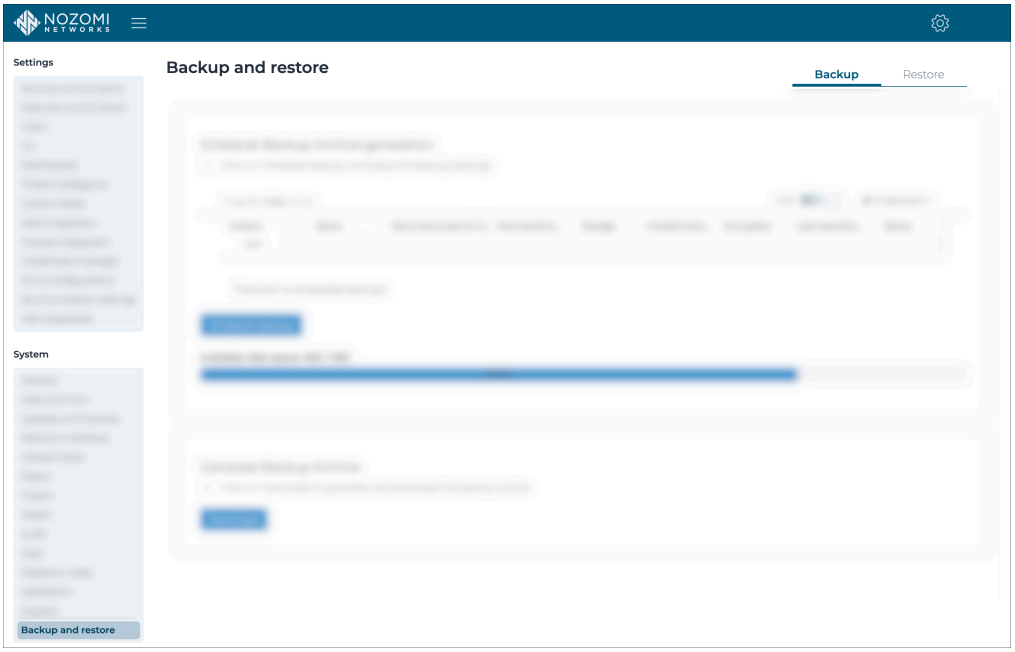


Figure 71. Backup and restore page

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.*

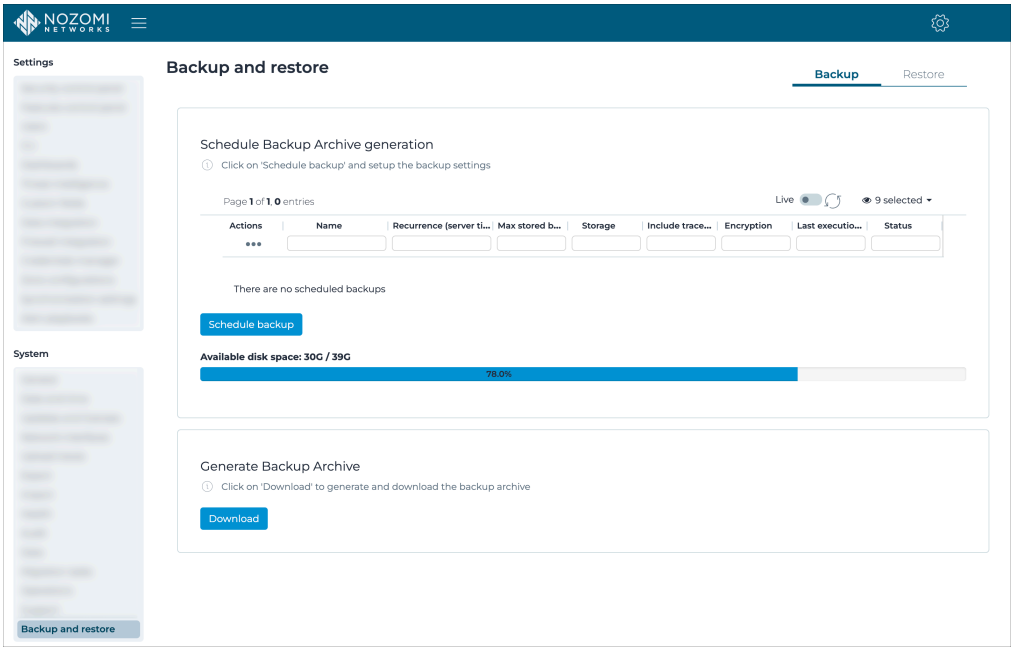


Figure 72. Backup page

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.*

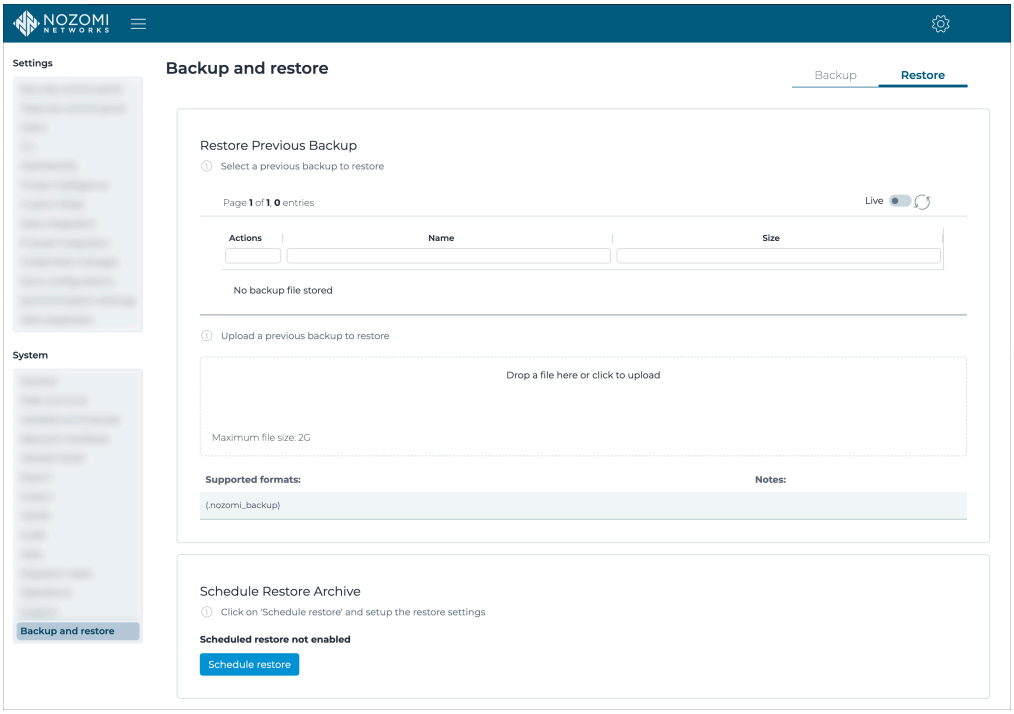


Figure 73. Restore page

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



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.

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.

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.

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.

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.

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 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 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 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.

