

Nexthink V6.23

User Manual

Generated: 4/28/2020 12:23 am

Table of Contents

Getting started.....	1
Logging in to the Finder.....	1
Logging in to the Portal.....	7
Querying the system.....	9
Searching the subject of interest.....	9
Executing an investigation.....	12
Creating an investigation.....	13
Editing the options of an investigation.....	16
Combining logical conditions in investigations.....	21
Navigating through the results of an investigation.....	23
Properties of users and devices.....	31
Visualizing system activity in the Finder.....	36
Getting a quick overview.....	36
Graphically observing the activity of users and devices.....	37
Observing service performance.....	48
Viewing network connections.....	55
Viewing web requests.....	58
Viewing executions.....	62
Monitoring IT custom metrics.....	66
Creating a metric.....	66
Examples of metrics.....	73
Following the evolution of a metric.....	78
Finding the visuals of a metric.....	81
Monitoring IT services.....	83
Analyzing service quality.....	83
Creating a service.....	85
Following the evolution of a service.....	89
Specifying URL paths of web-based services.....	90
Engaging with the end user.....	92
Getting feedback from the end users.....	92
Types of campaigns.....	96
Creating a campaign.....	101
Editing a campaign.....	107
Types of questions.....	109

Table of Contents

Engaging with the end user	
Controlling the flow of questions.....	111
Translating a campaign.....	113
Triggering a campaign manually.....	117
Scrutinizing the results of a campaign.....	118
Rating devices and users with scores.....	122
Computing scores.....	122
Creating a score.....	124
Checking and comparing ratings.....	126
Computing potential savings.....	129
Score XML Reference.....	131
Documenting scores.....	141
Remotely acting on devices.....	145
Scenarios for remote actions.....	145
Creating a remote action.....	147
Executing remote actions.....	151
Triggering a remote action manually.....	154
Writing scripts for remote actions.....	156
Example of self-healing scenario.....	163
Example of self-help scenario.....	167
Organizing objects with categories.....	175
Classifying objects of the same type.....	175
Creating categories and keywords.....	177
Tagging objects manually.....	179
Tagging objects automatically.....	180
Importing tags from text files.....	182
Getting notified by the system.....	185
Receiving Engage campaigns.....	185
Receiving email digests.....	187
Receiving alerts.....	190
Creating a service-based alert.....	195
Creating an investigation-based alert.....	197

Table of Contents

Building web-based dashboards.....	201
Introducing dashboards in the Portal.....	201
Creating a dashboard.....	202
Examining metrics in depth.....	203
Documenting dashboards.....	208
Assessing license use.....	211
Computing dashboard data.....	217
Reusing dashboard content.....	225
 Importing and Exporting authored content.....	 229
Methods for reusing authored content.....	229
Manually sharing Finder content.....	230
Importing a content pack.....	232
Conflict resolution.....	234
Exporting a content pack.....	238

Getting started

Logging in to the Finder

To start working with the Finder, authenticate first as a valid user. Take advantage of single sign-on solutions if your administrator has enabled either SAML authentication or Windows authentication of users. Otherwise, if plain authentication via Active Directory is available in your setup, log in with your Windows credentials. Alternatively, if no external authentication mechanism is available, use the credentials of your own dedicated Nexthink account.

The login process starts by connecting the Finder to the Portal. From the Portal, the Finder retrieves first the list of available Engines within your view domain and finally connects to the Engine of your choice. Once connected, the Finder has access to the data stored both in the Engine and in the centralized content manager of the Portal; letting you visualize, organize, and query end-user data in a variety of ways.

Quick connect

When you run the Finder for the first time, the login dialog appears in *Quick connect* mode, as indicated by its title at the top left corner of the dialog. Quick connect mode lets you specify all the necessary credentials to connect to any Portal. To log in to the Finder in Quick connect mode:

1. Fill in the **Host : Port** field with the DNS name or IP address of the Portal to which you want to connect, optionally followed by a colon and the port number where the Portal is listening for connections (by default, 443).

2. Depending on the authentication mechanisms available, choose one of the following:
 - ◆ If your administrator has enabled either SAML or Windows authentication of users, tick the box **Use single sign-on**.
 - ◆ If only internal or plain Active Directory authentication are available, enter your credentials manually:
 1. Type in your **Username**, which is the name of the user as registered in the Portal.
 2. Type in your **Password**.
3. Optional: Check the option **Remember me** to have the **Host : Port** and **Username** fields prefilled with the same data that you just typed the next time that you log in to the Finder from the same computer. The next login, you will just have to retype your password (if not externally managed).
4. Optional: Check the option **Sign me in automatically** for the login dialog to remember your password as well and skip the login step altogether the next time that you run the Finder from the same computer. Checking this option implies that the previous **Remember me** option is also checked. Administrators can disable this option.
5. Click **Connect**.
 - ◆ If SAML authentication is enabled and you have not logged in to your corporate account yet, the Finder will open your corporate login page on your default web browser. Log in with your corporate account as usual.
6. If more than one Engine within your view domain is available in the Portal, a list with all connected Engines shows up in the dialog **Engine selection**:
 - ◆ Click the name of an Engine to connect to it.
 - ◆ Click the star to the right of the name of an Engine to make that Engine your favorite. The next time that you log in to the Finder from the same computer, the step to select the Engine is skipped and you connect directly to your favorite Engine. You can later change your favorite Engine once you have logged in.

If the Finder cannot connect to the Portal or to the selected Engine, it aborts the login process and displays an error message with the reason for the failure. For warnings related to security certificates during the connection, see the section on

certificate issues below.

Creating a session

In a multi-user or multi-Appliance environment (e.g. an environment with test and production Appliances), you may have to log in to the Finder with distinct user accounts and connect to different Appliances from a single post. In these cases, to save you from typing the credentials every time that you have to log in to a different Appliance, store the credentials for distinct users and associated Appliances into *sessions*. Later, log in to the Finder faster by accessing your stored sessions.

To create a new session in the login dialog:

1. Click the **+New** button found at the top right corner of the login dialog. The login dialog turns into *session creation mode*, copying the information that you typed previously in Quick connect mode, if any, or from a previously selected session into the fields **Host : Port** and **Username** of the new session. Note that sessions do not store passwords by default. You can later specify to remember your password if you frequently use the same session.
2. Type in the DNS name or IP address of the Portal appliance, followed by a colon and the port number where the Portal is listening for connections in the field **Host : Port**. You can keep the copied value, if any.
3. Type in the name of the user to store in the session in the field **Username**. Again, you can keep the copied value, if the name of the user is not empty.
4. Optional: Change the name of the session that is displayed at the top of the dialog by clicking on it and typing an alternative name. By default, the name of the session is built from the name of the user and the Portal appliance in the form: *Username on Host : Port*.
5. Optional: If the system supports either SAML or Windows authentication, tick the box **Use single sign-on** and the Finder will let the external authentication mechanism manage the username and password. When using SAML or Windows authentication, specify a proper DNS name for the Portal (not an IP address) under **Host:Port**.

6. Click **Create** to save the session. The login dialog switches now to *session mode*.

Note that sessions are created locally in your instance of the Finder. Therefore, the sessions that you create in one computer are not automatically available when you try to log in to the Finder from another computer.

Using a session to log in

Once you have created one or more sessions, you can use them to quickly log in to different Engines. You only have to select the appropriate session and, eventually, enter your password. To log in from a saved session from the login dialog:

1. Click the down arrow in the top left tab of the login dialog, to the right of its title. A drop down list appears with the names of the saved sessions and the **Quick connect** option at the top.
2. Select one of the saved sessions. The user and Engine information stored in the session are displayed in the dialog. If you select **Quick connect** instead, you go back to Quick connect mode. Read the previous section on using the login dialog in Quick connect mode.
3. Type in your password, if needed (not needed if you previously told the Finder to remember or if using SAML or Windows authentication).
4. Optional: To store your password with the session information, check the option **Remember password**.
5. Optional: Check the option **Sign me automatically** to skip the login step altogether and start the connection to the Engine of the selected session as soon as you open the Finder. This option requires the **Remember password** option to be checked.
6. Click **Connect** or press **Enter** and the Finder starts to establish a connection to the chosen Portal.
7. If more than one Engine within your view domain is available in the Portal, a list with all connected Engines shows up in the dialog **Engine selection**:
 - ◆ Click the name of an Engine to connect to it.

- ◆ Click the star to the right of the name of an Engine to make that Engine your favorite. The next time that you log in to the Finder from the same computer, the step to select the Engine is skipped and you connect directly to your favorite Engine. You can later change your favorite Engine once you have logged in.

If the Finder cannot connect to the Portal or to the selected Engine, it aborts the login process and displays an error message with the reason for the failure. For warnings related to security certificates during the connection, see the section on certificate issues below.

Editing a session

If one of your saved sessions has wrong data or there was a change in user or Appliance settings, you may want to edit the session. To edit the values stored in a session:

1. Click the down arrow placed in the top left tab of the login dialog (to the right of the title) and select the session that you want to edit from the drop down list.
2. Click the **Edit** button that you find at the top right corner of the login dialog. You enter session edition mode.
3. Edit the fields **Username**, **Host : Port** and the name of the session in the same way as you did when creating the session (see previous section).
4. Click **Save** to save your changes and go back to session mode.

Deleting a session

When you do not need a session anymore, remove it from your list of sessions. To delete a session from the login dialog:

1. Click the down arrow placed in the top left tab of the login dialog (to the right of the title) and select the session that you want to delete from the drop down list.
2. Click the **Delete** button that you find at the top right corner of the login dialog. A dialog appears asking you for confirmation on deleting the session.
3. Click **Yes** to confirm that you really want to delete the session. The login dialog removes the session from the list and goes back to Quick connect mode.

Certificate issues

While connecting to the Portal or to the Engine of your choice, the Finder may display a warning message about security certificates on a dialog with the title:

There is a problem with the Portal / Engine security certificate

Security certificates ensure that the connections among Nexthink components are safe. A problem with the certificates implies that there is a potential risk of impersonation. In particular, if you use the default self-signed certificates from Nexthink (or any other certificate not signed by a trusted CA), you can read the following message in the dialog:

The security certificate of Nexthink Portal / Engine could not be validated.

If you are testing the solution or you are sure that the certificate is correct:

1. Optional: Click **Show certificate** to display detailed information about the current certificate.
2. Optional: Tick the option **Do not notify me again for this certificate** if you want to accept the current certificate as valid and avoid the warning in subsequent logons.
3. Click **Continue anyway** to go on with the login process despite the warning message.

Otherwise, contact your administrator for replacing the certificates. Log in only after your administrator has finished installing the new security certificates.

Related tasks

- Adding users
- Enabling Windows authentication of users
- Enabling SAML authentication of users
- Preventing password saving in the Finder
- Importing and replacing Certificates

Related references

- Active Directory authentication

Logging in to the Portal

To log in to the Portal, take advantage of single sign-on solutions if your administrator has enabled either SAML authentication or Windows authentication of users. Otherwise, if plain authentication via Active Directory is available in your setup, log in with your Windows credentials. Alternatively, if no external authentication mechanism is available, use the credentials of your own dedicated Nextthink account.

To log in to the Portal:

1. Open your favorite web browser (check the list of supported web browsers in the release notes of the Portal).
2. Type in the name or IP address of the Portal in the address bar of the browser.
 - ◆ For external authentication via SAML or Windows, typing the DNS name of the Portal is the only valid option.
3. On the login dialog, choose any of the following:
 - ◆ Click **use corporate single sign-on**, to authenticate through SAML (only available if SAML authentication is enabled in your setup).
 - ◇ If you have not logged in to your corporate single sign-on solution yet, the Portal redirects your web browser to the corporate login page.
 1. Type in your corporate credentials.
 - ◆ Click **use Windows authentication**, to authenticate with your Windows credentials (only available if Windows authentication is enabled in your setup).
 - ◆ Enter your Portal credentials manually:
 1. Type in your user name under **Username**.
 2. Type in your **Password**.
 3. Optional: Tick the box **Remember me** for the Portal to automatically fill in the Username field the next time that you

- log in.
4. Click **Sign in**.

Optionally, brand the logo and the background image displayed in the login page.

Related tasks

- Branding the Portal
- Enabling Windows authentication of users
- Enabling SAML authentication of users

Related references

- Active Directory authentication

Querying the system

Searching the subject of interest

Searching in the Finder

The Finder offers you many different possibilities to get information about your IT infrastructure. To help you find your way through, the Search box in the Start page is the best place to get started with the Finder. The Search box is a valuable tool for both beginners and experienced users and it can save you lots of time by directly bringing useful information to your fingertips. To look for a particular piece of information in the Finder, just type what you are looking for in the Search box.

To search in the Finder with the Search box:

1. Locate the Search Box at the top of the Start page, placed to the right of the Nexthink logo and labelled with the legend *What are you looking for?*. The Search box has a distinctive magnifying glass on its right side.
2. Type related terms to the subject of interest in the Search box (minimum three characters).
3. Press **Enter** or click the **Search** button.
4. Optional: Click the cross to cancel the current search and erase both the Search box and the list of results.

The results of a search are all items and investigations that are related in some way to the words that you typed in. For instance, you can look for users, devices or applications by typing in the first letters of their names. In this way, if you type **nexthink** in the Search box, you get a list with the applications from this company that are installed in your system, such as the Finder and the Collector, along with their executables and binaries, the packages that installed those applications, the devices where they were installed or even the visited web domains that are related to Nexthink, such as `doc.nexthink.com`. Similarly, you can look for investigations related to generic objects by typing the name of the object. For example, if you type in the keyword *printers* in the search box, you get investigations related to printers under the section **Investigations**. You can also search by platform. For example, type `Mac OS applications` to search for applications on the Mac OS platform.

Search results

As an example, let us examine the search results when looking for user *gdaniel*

in the Finder:

The results are divided into two sections:

- The section **Investigations** on the left-hand side of the window.
- The related items section on the right-hand side of the window, which includes the objects, services, metrics, categories, campaigns, scores, and remote actions that match the search terms.

The **Investigations** section usually shows both suggested and existing investigations:

- *Existing* investigations are those that have been previously created in the Finder and are available to the current Finder user.
- *Suggested* investigations are those inferred from the search terms and proposed by the Finder on the fly. All suggested investigations in the list are marked with the keyword **Suggested**.

In this specific example, only suggested investigations show up, as there are no existing investigations that include the term *gdaniel*. Each investigation is preceded by an icon that indicates the object on which the investigation is based. Click an investigation on the list to run it.

Result items on the right-hand side of the window are grouped by type. Click the result item to drill-down to it or right-click the item to trigger different actions.

Cross-Engine search results

When the Cross-Engine search features are fully enabled (including search and auto-complete), the Finder generates the suggested investigations on the left-hand side of the search window by matching the search terms with data

retrieved from all Engines:

Items related to the search are also retrieved from all Engines and displayed on the right-hand side of the search window, regardless of the Cross-Engine features being fully or just partially enabled.

Retrieved users and devices are limited by your view domain. Click the name of a user or device to respectively open the user or device view, regardless of the Engine where the user or device is located.

The search results of objects other than users or devices are not limited by your view domain. However, if you drill-down from the search results to an object that is not in your view domain, you get an empty list with the following highlighted message at the top: **Either there are no results or you do not have permission to view them.**

[Click here](#) for more information on how to effectively search in the Finder.

Searching devices in the Portal

The Portal includes a search tool, called *Device locator*, that looks for devices across all connected Engines. To know which Engine receives information from a particular device, enter the name of the device or the name of a user of the device in the Device locator.

To look for a particular device in the Device locator:

1. Click the magnifying glass in the top right corner of the Portal. The Portal brings up the **Device locator**.
2. Type the name of the device or the name of a user that has interacted with the device in the search box labelled *Enter a device or user name*.

Incomplete names are allowed.

3. Click **Search** or press **Enter**.
4. Results show up below up to a maximum of ten entries. If necessary, refine your search to reduce the number of matching results. Each entry displays the name of the device, the user that last logged in to the device, the Engine to which the device sends information, and the last time that the device has been seen in the network.
5. Optional: Click the name of the device in the list of results to open the device view of that device in the Finder.
6. Optional: Click the name of the last logged in user in the list of results to open the user view of that user in the Finder.

Related references

- Search in Finder

Related concepts

- Object
- Investigation
- Service
- Category

Executing an investigation

The basic unit to query the Nexthink database is the *investigation*. To execute an investigation in the Finder:

1. Find the investigation that you want to execute by either:
 - ◆ Browsing the **Investigations** section on the left-hand side panel of the main window.
 - ◆ Searching a term in the Search box and looking under the **Investigations** or the **Suggested investigations** sections in the results of the search.
2. Double-click the name of the investigation.
3. The results of the investigation are displayed as a table on a new tab in the Finder.

Running a metric as an investigation

Besides investigations, metrics also express queries in their definitions. For testing the definition of a metric, it is useful to execute it as an investigation and check out the results. To run the query defined in a metric as a quick investigation:

1. Find the metric that you want to execute as an investigation by either:
 - ◆ Browsing the **Metrics** section on the left-hand side panel (the accordion) of the main window.
 - ◆ Searching a term in the Search box and looking under the **Metrics** section in the right-hand side of the search results.
2. Right-click the name of the metric.
3. Select **Run as investigation** from the context menu.

The generated quick investigation shares the conditions specified for the metric, ignoring the settings for the computation of a ratio and the thresholds, when defined. The time frame of the investigation generated for a metric is the current day.

Related tasks

- Navigating through the results of an investigation
- Searching the subject of interest
- Creating an investigation
- Creating a metric

Related concepts

- Investigation

Creating an investigation

Investigations let you query the real-time database and look for specific information about objects, activities or events.

To create a new investigation in the Finder, you have different choices:

- [Create a quick investigation](#) to run it immediately. Optionally, save the investigation later.
- [Create a named investigation](#) and save it at the time of creation for later reuse.
- [Duplicate an existing investigation](#) that is similar to the investigation that you want to create.

- [Save or edit a drill-down or a one-click investigation](#) from its list of results.

Creating a quick investigation

To create a quick (unnamed) investigation in the Finder:

1. Click the **New Investigation** button at the bottom of the start page.
2. Edit the options of the investigation.
3. Click the **Run** button to execute the investigation that you just created.
4. Optional: if you want to save the investigation, press **Ctrl+S** or click the floppy disk icon at the top right corner of the window that displays the list of results.

Creating a named investigation

To create a named investigation in the Finder:

1. Go to the **Investigations** section on the left-hand side of the main window.
 - ♦ If you want to place the investigation at the top level of the investigations tree, right-click the header or the empty area of the sub-section **My investigations**.
 - ♦ If you want to place the new investigation inside a particular folder, right-click one of the folders of the sub-section **My investigations**.
2. Select **Create new investigation** from the context menu.
3. A new investigation with the temporary name **Untitled investigation n**, where *n* is a growing number to avoid name collisions, shows up in a new tab. Optional: Click the name of the investigation and replace it with an appropriate name.
4. Optional: Click the field below the name of the investigation labeled **Enter description here...** and give a short description of the purpose of the investigation.
5. Edit the options of the investigation.
6. Click **Save & Run** to save the new investigation and execute it right after.

Duplicating an existing investigation

Creating an investigation from scratch can be a lengthy process because of the many options and settings available for editing the investigation. If you already have at your disposal an investigation that is similar to the investigation that you want to create, you can save a lot of time by duplicating the existing investigation and adapting it to your needs, instead of going through the whole process of

creating a new investigation.

To duplicate an existing investigation:

1. Find the investigation to duplicate in the **Investigations** section of the accordion.
2. Right-click the investigation name and select **Duplicate** from the context menu. A new investigation with the name **Copy of A**, where *A* is the name of the existing investigation, is created. A new tab opens the investigation for editing.
3. Optional: Click the name of the investigation and change it to something more meaningful than **Copy of A**.
4. Optional: Change the description of the investigation to reflect the purpose of the new investigation.
5. Modify the investigation options. Because you selected a similar investigation, there should be only a few options to change.
6. Click on **Save & Run** to save the investigation and execute it right away.

Saving after navigation

When you navigate through results by drilling-down or executing a one-click, you are effectively executing a new investigation that the Finder creates on the fly. The investigation is discarded as soon as you close the tab that displays the list of results of the drill-down or one-click. However, if you think that the investigation is useful and that you might need it later, you can save it before closing the list of results.

To save an investigation generated by a drill-down or a one-click from its list of results, either:

- Click the floppy disk (save) icon in the top right corner of the window and give a name to the investigation.
- Click the pencil and paper (edit) icon placed to the right of the save icon. You are asked to save the investigation with a name. Additionally, you can edit the investigation settings.

Related tasks

- Editing the options of an investigation

Editing the options of an investigation

Overview

To edit the options of an existing investigation, either:

- Right-click the investigation name in the **Investigations** section of the left and select **Edit**.
- Execute the investigation and click the pencil and paper icon that appears in the top right corner of the list of results.

When you create a new investigation or edit an existing investigation, the Finder opens a dialog that lets you set all the options of the investigation.

The first thing that you find at the top of the dialog is the name of the investigation and an optional description of what it does. Click the name or the description to modify their contents.

Below the name and the description, you find three distinct sections that let you design the investigation to get the desired results:

- **Retrieve**
- **Matching**
- **Display**

Retrieve section

In the **Retrieve** section, choose the object, activity or event of interest. The execution of the investigation returns a list of results with items of the selected type.

Platform selection

In the upper-right part of the **Retrieve** section, find three check boxes to select the platforms that are applicable to the investigation. The conditions and display fields that you are able to edit in the investigation depend on the platforms that you select here.

- If you choose one platform, you can use conditions and display fields available for that platform.

- If you select multiple platforms, only those conditions and display fields shared by all the selected platforms are available.

For instance, if you select to retrieve devices of the Mobile platform, you can only set conditions on devices or user fields, because all other objects are not available for Mobile.

In a similar way, if you choose to retrieve an object type that is not available in all platforms, the check boxes of the platforms in which the object is not available are ineligible.

For example, if you choose to retrieve domains, which are only available for the Windows platform, the check boxes of both Mac OS and Mobile platforms are disabled.

By default, when you create a new investigation, only the Windows platform is ticked in this section.

Matching section

In the **Matching** section, you select the criteria that the objects, activities or events of the type that you chose in the **Retrieve** section must fulfill to appear in the list of results. The **Matching** section is divided into two subsections: **Conditions** and **Time Frame**.

Conditions

The matching **Conditions** are a set of rules that apply to any type of item related to the one selected in the **Retrieve** section. You can set constraints on the properties or categories of objects, activities or events to filter the results of your investigation.

To add a new condition:

1. In the **Conditions** subsection, click the link **Click here to add a new condition**. The placeholders for the condition fields show up.
2. Set the object, activity or event to which the condition applies.
3. Set the attribute or category that you want to constraint.
4. Set the operator for comparison (e.g. *is*, *is not*, *starts with*, etc).
5. Set the matching value, if you selected an attribute constraint, or the matching keyword, if you selected a category constraint.
 - ♦ As you type, auto-complete looks in the Engine for values that match the written characters whenever possible (e.g. when setting

conditions on names and not on numerical values). If the appropriate Cross-Engine features are enabled, auto-complete looks for matching values in all Engines.

Some combinations of conditions and display settings are incompatible. If you add a condition and a red exclamation mark appears on its right side, the condition may conflict with another condition or with one of the chosen attributes to display. Hovering the mouse over the exclamation icon will tell you the reason for the conflict. Investigations with conflicting conditions cannot be saved. Deselect the conflicting display attributes or delete the conflicting condition before saving the investigation.

To delete a condition:

1. Click the trash icon to the right of the condition fields.

To make a template investigation:

1. Instead of providing a matching value in the last condition field, click the question mark to its right to transform the investigation into a template investigation. The actual matching value is provided as a parameter when executing the investigation.

By default, the results of an investigation must fulfill all the expressed conditions. That is, the resulting filter is a logical AND of all the conditions. If you want to combine the conditions in a different way:

1. Click the **Advanced** area to expand it.
2. Combine the conditions in the **Logical expression** field using the numbers of the conditions and the Boolean operators AND and OR. For instance: 1 AND (2 OR 3).

The final **and** in the **Conditions** section allows you to specify a condition on an aggregate of the object selected in the **Retrieve** section. Activities and events do not have associated aggregate values.

Read the article on how to combine conditions to exactly express your intention when writing an investigation.

Time frame

To limit the results of the investigation to a particular range of time, use one of the following options:

Full available period (*start date to end date*)

Do not limit the results. The investigation uses the full range of time available in the Engine, which is stated in the start and end dates. If Cross-Engine features are enabled, the start and end dates are adapted to the maximum span available across all Engines within the view domain of the user. This option is not available for investigations based on activities or events nor for any investigation based on objects that needs to go through activities or events.

On *date*

Limit the results of the investigation to a particular day. The available dates to pick are either those of the current Engine or, if Cross-Engine features are enabled, those of any Engine within the view domain of the user.

During the last *x* days / hours.

Get the most recent matching results, that is, those that occurred less than the specified number of days or hours ago. Note that, when expressed in days, the time is partitioned in natural days, going from 0h to 23h59. As a consequence, it is not the same to restrict the time frame to the last day (from midnight today until now) than to the last 24 hours.

From *start date and hour to end date and hour*

Specify the period limit manually. Again, available dates are either those of the current Engine or those of any Engine within the view domain of the user, if Cross-Engine features are enabled.

Additionally, for specified time frames that span through several days (with the exception of the Full period choice), you can optionally specify a range of hours of interest:

Between *start hour* and *end hour*

Choose a period of interest inside every single day included in the investigation.

To avoid long computation times in the Engine, the time frame of investigations that need to go through activities or events is limited by default to a maximum of 7 days. It is possible to remove this 7-day limit and launch investigations whose time frame spans up to the maximum number of days available in the Engine.

Display section

In the **Display** section, determine how the Finder presents the results of the investigation. Choose between showing all the available results or just a fixed number of entries, according to some sorting criterion. In addition, select the *fields* (attributes and categories) of the retrieved objects that will be arranged as

columns in the list of results.

Optionally restricting the number of results

To either display all the results of the investigation or restrict their number, use the option that you find at the top of the **Display** section. Choose between:

All results

Display all retrieved items without limit.

The top *x items* ordered by *field* ascending / descending

Limit the list of results to the first *x items* in ascending or descending order, according to the specified field.

Selecting the columns

Under **Columns**, specify the fields whose values you wish to see as columns in the list of results of the investigation. Select the fields by means of a label selector, where each label holds the name of a field. The Finder pre-populates the label selector with a set of default fields that depend on the type of item to retrieve and the previously specified options for the investigation.

To add a column to the list of results:

1. Click the label selector to place the cursor on it. A selection menu exhibits all available fields organized by sections.
2. Select the field either by clicking or by typing its name:
 - ◆ Click the name of the field that you want to add as column. The field must not have been already added to the label selector (in which case, it is disabled in the menu).
 - ◆ Start typing in the name of the desired field. The selection menu pops up, showing only those fields whose name includes the characters entered.
 1. Optional: Click the name of the desired field in the selection menu to add it directly. As indicated above, the field must not have been already added.
 2. Optional: Press **Tab** to auto-complete the name of the field if it is the only field left in the menu.
 - ◇ Optional: Learn how to use the keyboard for an even faster selection of columns.

To be eligible, fields must be compatible with the options specified for the investigation (e.g. some aggregates are not available if the time frame selected is the full period available). Hover the mouse cursor over a disabled field to know

about the reasons for the incompatibility.

To remove a column from the list of results, either:

- Click the cross sign on the right side of the label that holds the field name.
- Place the cursor to the left of the field label and press **Delete** or to the right of the field label and press **Backspace**. To remove all the labels at once, press **Ctrl+A** to select them all and press **Delete**.

Note that if you have restricted the number of results according to the value of a field, that field is mandatory and it cannot be removed from the label selector.

In any case, the set of labels in the label selector must never be empty. If you remove all the labels from the selector, then the a label with the unique identifier of the object (**UID** field) is automatically added.

Related tasks

- Combining logical conditions in investigations
- Expanding the time frame of investigations in the Finder
- Enabling Cross-Engine Finder features

Related references

- Keyboard shortcuts for column display selection

Combining logical conditions in investigations

Overview

Because of their strict logical interpretation, conditions may be difficult to get right in some kinds of investigations. Beware especially of literal translations from natural language to logical conditions, because they may lead to unexpected results.

To express what you really mean in an investigation, keep in mind that relations between objects in Nexthink are all related to events in some way. For example, if a device is linked to an application, it is because the application was executed at some point in that device. The system iterates through these relations to return the results of an investigation.

Combining logical AND, OR, and aggregate conditions

Consider a case where you want to get a list of the devices that have executed two different applications (e.g. *Internet Explorer* and *Firefox*) within a particular time frame, whose duration is irrelevant to our discussion. You may be tempted to use a logical AND (the default) to combine the conditions of an investigation on devices:

- Condition 1: Application name is *Internet Explorer*
- Condition 2: Application name is *Firefox*
- Logical expression: 1 AND 2

If you run this investigation, the list of devices that you get is always empty. The reason is that any relation between a device and an application ultimately relies on executions and no execution may simultaneously satisfy the two conditions: it is either an *Internet Explorer* execution or a *Firefox* execution, but not both. When the system iterates through these relations, it discards them all because none is matching the two conditions at the same time, as required by the logical operator AND. Hence the empty result.

To properly state the desired query, keep the same conditions but modify a couple of inputs. First, change the logical expression from AND to OR:

- Logical expression: 1 OR 2

In this way, the system keeps the relations that match either the first condition or the second. That is, you get all the devices that executed either *Internet Explorer*, *Firefox*, or both; although you are only interested in the last group.

Last, restrict the output to only those devices that executed the two applications. Add the following aggregate condition at the end of the conditions section:

- Aggregate condition: Number of applications is 2

If you change your mind afterwards and decide to ask for the devices that executed either *Internet Explorer* or *Firefox*, but not both, use the following aggregate condition instead:

- Aggregate condition: Number of applications is 1

Although not useful in the example, the logical AND is still the most common operator to combine conditions. Use it when you want to enforce two compatible conditions at the same time. For instance, to know the devices that use *Internet*

Explorer for browsing a particular domain, say *www.example.com*, create an investigation on devices with the following conditions:

- Condition 1: Application name is *Internet Explorer*
- Condition 2: Domain name is *www.example.com*
- Logical expression: 1 AND 2

Negating with an aggregate condition

Consider now the case of getting the list of devices that have not executed Firefox. By translating from natural language, you may wrongly assume that the following condition in an investigation on devices yields the correct result:

- Condition: Application name is not *Firefox*

Instead, this investigation returns all the devices that have executed any application, except for those few devices that have exclusively executed Firefox (probably none, as devices typically execute multiple applications simultaneously). Most likely, the result will thus be the full list of active devices.

The correct way to negate a condition is to first state the condition in a positive way and then negate the statement with an aggregate condition that looks for zero instances. For our example with Firefox, that would be:

- Condition: Application name is *Firefox*
- Aggregate condition: Number of applications is 0

This way of negating conditions applies to many types of objects and events, such as executables, binaries, connections, packages, etc.

Related references

- Investigation with packages

Navigating through the results of an investigation

Overview

After executing an investigation, you are presented with a list of all the items that matched your query conditions. This is the **List** view of the Finder.

The list displays all the fields and aggregates that you selected when you edited the options of the investigation.

Cross-Engine List view

If the Cross-Engine features in the Finder are enabled, click the button **List (all entities)** to get results from all Engines; more specifically, from all the entities that belong to the view domain of the Finder user. In consequence, if a user's domain view is limited to entities within a single Engine, the user will not get results from other Engines.

If the button **List (all entities)** is disabled, hover the mouse cursor over the button and wait for a tooltip to show up. The tooltip describes the first encountered reason for the incompatibility of the investigation with the Cross-Engine feature. Because of particular investigation options usually related to aggregates (e.g., conditions on aggregates or ordering of results based on aggregates), some investigations cannot be executed across multiple Engines and are thus incompatible with the List (all entities) view.

On the other hand, the List (all entities) view is still available for investigations that specify display fields which cannot be computed across multiple Engines, as long as the investigation is not invalidated for other reasons. In this case, the List (all entities) view resolves the incompatibility by not displaying the conflicting columns.

The List (all entities) view displays up to a maximum of 10 000 results by default, although this number is configurable and you can overcome this limit by exporting the results of the investigation. The **Network**, **Web**, and **Local** activity views are not available when displaying Cross-Engine results.

The results are ordered by the values of the first sortable column. A message in the status bar at the bottom of the window indicates whether the condition of maximum displayable results is reached. For instance, when displaying the

results of an investigation based on devices that returns more than ten thousand objects, the status bar states the following:

Only the first 10 000 devices (ordered by "Name") are displayed

In turn, if any of the connected Engines does not reply to the Cross-Engine query, the Finder displays the following warning message at the top of the window:

**Partial results are displayed because some Engines didn't respond.
Contact your administrator if the problem persists.**

Moreover, when displaying the results of an investigation based on devices, the Finder adds a special column to the List (all entities) view called **Engine**. The Engine column holds the name of the Engine that stores the data of the related device. A single device that reports to several Engines gets a different UID from each Engine and is therefore seen as a set of distinct objects in the list of results.

Sorting the results

Order the results in the **List** view according to the value of one of the displayed fields by clicking its corresponding column header. The arrow to the right of the column name indicates if the sorting is made in ascending (arrow up) or descending (arrow down) order. Click the column header again to change the direction of the arrow.

By default, results are sorted according to the values of the first column in ascending order; unless otherwise specified in the investigation options, where you can look for a number of top objects ordered in either ascending or descending fashion. You can click on any other column header of the List view to sort the results in a different way.

Note that ordering the results in the List view applies to the already displayed results only; that is, ordering the results by clicking on a column does not trigger a new investigation. For instance, if an investigation shows the top 10 devices with the highest TCP traffic, clicking the **Name** header in the List view reorders the same 10 devices according to their name. The same behavior reproduces in a **List (all entities)** view that reaches its configured maximum number of results: ordering the results by clicking on a column does not trigger a new Cross-Engine investigation.

Changing the time frame


The **List** and **List (all entities)** views display the time frame that applies to the given results of the investigation in the top-center part of its own tab. To change the time frame of the investigation, click the calendar icon that appears to the right of the current time frame. A dialog very similar to the Time frame section of the investigation designer shows up. If Cross-Engine features are enabled, the dates available for selection are those available across all Engines within the domain view of the user, as explained in the article about the Time frame section. Otherwise, the available dates correspond to the history of the current Engine. Set the new desired time frame and click **Apply**.

To come back to the original time frame of your investigation, click the calendar icon and then push the button **Reset**.

If you selected a limited time frame such as a particular day, you can also navigate easily with the arrows you find in both sides of the calendar icon. Just press the arrow to the right to move to the next available day, or the arrow to the left to move to the previous day.

Setting the platform

From the **List** view, filter the results of an investigation according to platforms at any time:

1. Click the platform icons at the top of the **List view**, , and a dialog to select the platforms shows up.
2. Tick the check box for every platform that you want to include in the results.
3. Optional: To go back to the platforms originally selected by the investigation, click **Reset**.
4. Click **Apply** to filter the results according to the selected platforms.

Adding and removing display fields

To quickly add or remove fields displayed as columns in the **List** view:

1. Right-click anywhere in the column headers (the top part with the names of the columns). A label selector shows up.
2. Use the label selector to add or remove columns in the same way as you select the columns when creating the investigation.
3. Click **Apply**.

To quickly remove a single column, right-click the column header and select **Remove column** from the context menu.

Drilling-down

Drilling-down to other items from your list of result items is one of the most powerful tools that you have for navigating through the results of your investigations. Drilling-down lets you get items related to the items in the list of results while keeping the context of your investigation, that is, enforcing the time frame and the conditions of the original investigation.

A drill-down is actually a quick investigation on objects, activities, or events that are related to a selection of the results of a previous investigation. For instance, imagine that you execute an investigation on devices that looks for those devices that executed the Nexthink Finder yesterday. You get a list of devices as a result. Imagine now that you want to know the users that executed the Finder yesterday from one or several of those devices. You can get the list of those particular users by drilling-down from the results of your previous investigation. Note that drilling-down keeps the conditions and the time frame of the original investigation, that is, the execution of the Nexthink Finder yesterday.

The drill-down keeps the Cross-Engine context as well. When selecting items of a **List** view to drill-down, the results lie within the limits of the current Engine. On the other hand, if your selection belongs to a **List (all entities)** view, the results of drilling-down extend to all available Engines.

To drill-down from a list of results of an investigation:

1. Execute the investigation of your choice.
2. Select one or more of the items in the **List** or **List (all entities)** views.
3. Right-click the items selected. A context menu shows up.
4. Select the option **Drill-down to** and choose a type of item. Items are classified into:
 - ◆ **Objects**

- ◆ **Activities**

- ◆ **Events**

5. Choose one class of items and then a particular type of object, activity or event. Only those types of items that can be related in some way to the items in the list of results are eligible for drilling-down.
 - ◆ If the items in the list of results are filtered by platform, the drilling-down shows only those items which are compatible with the selected platform.
 - ◆ In the case that you selected multiple platforms, the drilling-down shows all those items which are compatible with any of the selected platforms.
6. A new tab with the list of results for the drill-down opens.

The items that you can select for drilling-down depend also on the platform of the item you drill-down from. For instance, you cannot drill-down to printers from a Mac OS device, because the Mac platform in Nexthink does not support printers.

One-click investigations

One-click investigations, or *one-clicks* for short, are similar to drill-down investigations, except for the fact that they do not keep the context of the previous investigation.

For instance, to go on with our previous example, imagine that you are navigating the **List** view of an investigation that returns all the devices that executed the Nexthink Finder yesterday, and that you want to know all the users of a particular device. Drilling-down to users returns only those users who executed the Finder yesterday on that device. On the other hand, a one-click investigation on users returns all the users who have ever been seen in the device, regardless of what they were doing or when.

One-clicks do keep the Cross-Engine context though. When executing a one-click from a **List** view, the results are limited to the current Engine; whereas a one-click from a **List (all entities)** view returns items from all available Engines.

To perform a one-click investigation from the list of results of a previous investigation:

1. Execute the investigation of your choice.
2. Select one or more of the items in the **List** or **List (all entities)** view.
3. Right-click the items selected. A context menu shows up.

4. Select the option **One-click investigation** and retrieve all the items of a particular class. Choose among:
 - ◆ **Retrieve all objects**
 - ◆ **Retrieve all activities**
 - ◆ **Retrieve all events**

Note: for binary objects, specify first if you want to retrieve items related to the binary itself, or to the executable or the application to which the binary belongs. Similarly, for executable objects, choose first if you want to retrieve items related to the executable itself or to the application to which the executable belongs.
5. Select a particular type of object, activity or event. Only those types of items that can be related in some way to the items in the list of results are eligible for a one-click investigation.
 - ◆ If the items in the list of results are filtered by platform, the one-click investigation shows only those items which are compatible with the selected platform.
 - ◆ In the case that you selected multiple platforms, the one-click investigation shows all those items which are compatible with any of the selected platforms.
6. The Finder opens a new tab with the list of results for the one-click investigation.

Again, similarly to what happens with drill-downs, the items that you can select when you do a one-click investigation depend also on the platform of the *one-clicked* object. For example, you cannot retrieve all events related from a Mobile device because Mobile devices do not support events.

Exporting the results of a Cross-Engine investigation

Cross-Engine investigations that display **All results** are limited to show up to a maximum of 10 000 entries in the **List (all entities)** view by default. Although this limit is configurable, it is often impractical to deal with a large number of objects in the list view of the Finder. To overcome this limit and actually get all the possible outcomes, export the results of Cross-Engine investigations. From the **List (all entities)** view in the Finder:

1. Click the menu icon at the top right corner of the view.
2. Select **Export results** from the list.

The Finder opens your default web browser with the message **Export in progress** while the Portal generates a CSV file that gathers the results from all Engines. Once the CSV file is ready and zipped, the browser automatically initiates its download from the Portal. After the download is complete, the Portal removes its local copy of the file and the URL used for the download becomes invalid.

Saving your modifications

When you change the time frame or the displayed fields, or you drill-down, or do a one-click from the **List** view of an investigation, the system is actually executing a different investigation from the original one.

To save the new investigations that you create by applying modifications to the **List** view, click the menu icon at the top right of the view and select **Save investigation...** or press **Ctrl+S**.

Getting a graphical representation of the data

The **List** view gives you a plain text representation of the data stored in the Nexthink database. While this is perfect if you want to have a list with the exact values, it can be difficult for a human to get an insight of what is actually happening inside your IT infrastructure with just a textual representation.

To get a graphical representation of the results in the List view, click one of the buttons in the top-left corner of the List view:

Network activity

To visualize network connections.

Web activity

To visualize web requests.

Local activity

To visualize local program executions.

The visualizations are computed within the context of your investigation. Therefore, not all three visualizations are present for all investigations. A

visualization is available only if the context contains relevant information for it.

Related tasks

- Enabling the Cross-Engine features in the Finder
- Executing an investigation
- Editing the options of an investigation
- Creating an investigation
- Viewing network connections
- Viewing web requests
- Viewing executions

Properties of users and devices

Overview

One option to see and compare the properties of users or devices is to create appropriate investigations and display as many of their properties as desired, visually comparing the values in the list of results.

A more compact way to see the complete list of properties of a single user or device is to open the **Properties** tab of either the user or the device views in the Finder. To that end, open the user view or device view of a particular user or device and click the **Properties** button. On a single page, the **Properties** tab presents all the relevant attributes of a user or device conveniently grouped.

Device sections	User sections
Hardware Model, SN, CPU, RAM, and graphics information.	Properties Active Directory information about the user.
Network NIC, domain or workgroup, and Active Directory information.	Last user activity Last time that the user was seen on each accessed device.
Startup Last boot and logon and their average duration.	Categories Keywords assigned to the user for each user category.
Operating system Version, updates, and lists of device users by privilege levels and groups.	
Local drives	

List of hard disks on the device.	
Security	
Antivirus, antispyware, and firewall settings.	
Categories	
Keywords assigned to the device for each user category.	

Comparing the properties of devices on an Engine

The **Properties** tab lets you compare some of the properties of a device with the properties of other devices. The comparison is limited to devices on the same Engine (no Cross-Engine comparisons) that share the same platform, as the properties of a device depend on its platform.

To know if a property is comparable, look for the small bars icon to the left of the property name. Clicking the name of a comparable property opens a bar chart in which the heading of each bar is one (or a range) of the possible values for that property and the length of the bar indicates the number of total devices that share that same value. The actual number of devices that share the value is also shown as a figure to the right of the bar. The bar whose value matches the value of the property of the currently selected device in the device view is highlighted in the chart.

To show a list of the or devices that contribute to the length a bar in the bar chart, double-click the bar or right-click on it and select **Drill-down**. The drill-down is equivalent to launching an investigation on devices with that property value.

The chart shows thus the distribution of certain properties among the devices in your network. This is very useful to assess particular problems of a device. For instance, if a user complains of a laptop that takes too long to boot, click the property **Average boot duration** in the properties page. Compare the boot duration of the complaining user's laptop with the average boot duration of all other devices and find out whether the user has a good reason to complain or not. If so, check the hardware properties of the laptop and consider a possible hardware upgrade to reduce its boot duration.

By default, clicking a small bars icon compares a property of a device with the same property of all other devices on the same Engine. To limit the group of devices that take part in the comparison, use the **Compare with** tool. This tool is found in the top right corner of the **Properties** tab in the device view. If the Cross-Engine features are enabled, the name of the tool is **Compare with devices on Engine** to emphasized that the comparison is on the Engine where the device is located. The name of the Engine shows up to the right of the tool name.

To select a group of devices for comparison with the current device, the **Compare with devices on Engine** tool makes use of entities or categories. Thus, build up comparison groups with all devices that share the same category

or entity keyword.

For example, to limit the group of devices used in a comparison:

1. In the **Compare with devices on Engine** tool, select **devices with category** from the list instead of the default **all devices**. A selection box appears to the right holding the list of all device categories. Another selection box appears below with the text **and same keyword**. To the right of this last box, find the value of the keyword for the selected category and for the current user within parentheses.
2. Select a category from the list.
3. Select the category keyword that defines the group.
 - ◆ If you want to compare with the group of devices that share the same keyword as the current user, leave the default value **and same keyword** and you are done.
 - ◆ If you want to compare with another set of devices, choose the option **and keyword** instead of the default **and same keyword** in the selection box. Another selection box appears to the right, holding the list of all available keywords for the selected category.
 1. Select the keyword that defines the comparison group of devices from this last box.

Limiting the comparison group by entity is very similar:

1. In the **Compare with devices on Engine** tool, select either the option **devices with same entity** or **devices with entity**.
 - ◆ If you choose **devices with same entity**, the comparison group is formed by those devices that share the same entity as the current device. The name of the entity is displayed within parentheses to the right of the selection box.
 - ◆ If you choose **devices with entity**, another selection box appears to the right, holding the list of all available entities.
 1. Select an entity from the list to define the group of devices for comparison.

Remember that, in any case, the Finder compares the properties of a device with devices that share the same platform only.

Comparing the properties of users on an Engine

If Cross-Engine features are disabled, it is possible to compare the properties of users on the same Engine in a similar way that you compare the properties of devices. The **Compare with** tool in the user view is unavailable when

Cross-Engine features are turned on.

In the case of users, only the category keywords are subject to comparison:

Related tasks

- Graphically observing the activity of users and devices
- Enabling Cross-Engine Finder features

Visualizing system activity in the Finder

Getting a quick overview

Get a glimpse of recent happenings in your setup thanks to the *Quick views* of the Start page. Quick views summarize important information and let you perform frequent tasks, such as executing the most recently used investigations or inspecting services in detail. By default, the Start page displays the Quick views at the bottom of its main space, which is shared with the Search tool. Thus, launching a search effectively hides the Quick views, which are replaced by the results of the search. To get the Quick views back, erase the contents of the Search box.

Quick views are disposed horizontally in the screen as individual boxes. Scroll the Quick views with the help of the arrows placed on the left and right sides of the views to see all the boxes.

Recent alerts

The **Recent alerts** box highlights the alerts on devices that were triggered during the last 24 hours. The box displays the name of the alert along with the number of devices that generated the alert inside a red circle.

To analyze the alert in detail, click the name of the alert in the box for getting the list of impacted devices.

Services with threshold violations

The **Services with threshold violations** box shows a list of services ordered by the gravity of the violation (depending on whether the service crossed either the error or the warning threshold) and the number of devices impacted during the last 24 hours. For each service, the list displays the name of the service and, inside a red circle, the number of devices that attempted to connect to the service but found an issue.

To analyze the service in detail, click the name of the service in the list for opening the Service view.

Recently used investigations

The **Recently used investigations** box holds a list of the last five investigations that you executed in the Finder. To execute again any of the investigations in the list, click the name of the investigation.

Only named investigations are stored in the list of recently used investigations.

Graphically observing the activity of users and devices

Overview

To see at a glance the recent activities, scores, or detailed properties of a particular user or device, respectively open the *user view* or the *device view* in the Finder. Both the user and the device views include a **Timeline** tab and a **Properties** tab, as well as up to ten score tabs:

- Select the **Timeline** tab to explore the activities of a user or device in chronological order.
- Select the **Properties** tab to display detailed information about a user or device.
- Select a score tab to see the breakdown of the ratings of a user or device, according to the aspects covered by the main score. Optionally launch remote actions on devices with a low score when the documentation of the score includes links to remote actions.

By default, the device and user views open the **Timeline** tab. Note however that mobile devices do not have a **Timeline** tab, so the device view of mobile devices display only the **Properties** tab.

To open the user view or the device view of a particular user or device, either:

- Look for the user or device in the search box of the **Start** page and click the name of the user or device in the results of the search.
- From the list of results of an investigation based on users or devices, right-click the entry of the user or device and select **Display user view** or **Display device view**, or double-click the entry of the user or device, or select it and press **Enter**.
- From any of the other graphical views of the results of an investigation (**Network**, **Web** or **Local** activity views) that display users or devices,

right-click the name or the icon of a user or device and select **Display user view** or **Display device view**.

- From the user view itself, open the device view of any of the devices listed in the **Devices** section of the **Timeline** tab, or listed in the **Last user activity** section of the **Properties** tab, by clicking their name.
- Likewise, from the device view, open the user view of any of the users that interacted with the device, displayed in the section **Users** of the **Timeline** tab, by clicking their name.

At the top of the view, get basic information about the selected object:

User view	Device view
Name The name of the user.	Name The name of the device.
Type The class of user: local, domain or system.	Platform icon A pictorial representation of the platform of the device: Windows, Mac or mobile.
First Seen The first time of recorded user activity.	Entity The leaf node in the hierarchy, followed by the name of the Engine between parentheses, to which the device belongs.
Last Seen The last time of recorded user activity.	Last IP address The IP address of the device during its last recorded connection.
	Last Seen The last time of recorded device activity.

Below this basic object information, find the buttons that let you switch between the **Timeline**, **Properties**, and scores tabs. When selecting the **Properties** tab, a comparison tool appears to the right of the basic object information. Because this section focuses on visualizing system activity, the rest of this article is dedicated to the **Timeline** tab. To know more about the **Properties** tab, see the article on properties of users and devices. For more information about the score tabs, see how to check the ratings of users and devices.

To refresh the view, click the button with a circular arrow placed to the far right of the tab selection buttons. Refreshing the view is particularly useful when it is open for a long time and you want to see the last activity of a user or device.

Exploring activities in the timelines

The **Timeline** tab displays in fact several timelines grouped by sections. While the actual sections and their content depend on the type of object observed (user or device), the techniques to explore the timelines remain essentially the same.

To know the time scale of the timelines, find a ruler at the top of the view that divides the horizontal space in equal parts. Each subdivision of the ruler corresponds to a time interval of the recent history of the user or device under examination. Date and time labels in the ruler indicate the precise moment associated to a subdivision mark. In accordance with the ruler, an activity or event in the timelines found by following down a vertical line from a particular subdivision occurred during the time interval associated to that subdivision.

Hover the mouse cursor over a timeline with data and keep it there for a moment. A kind of structured tooltip eventually shows up. The tooltip summarizes the activities and events related to the timeline that happened during the time slot under the mouse cursor. A vertical and a horizontal dashed lines, crossing at the timeline slice pointed by the mouse cursor, show up shortly after the tooltip to help you locate the time interval in the ruler and the title of the timeline.

To investigate further what happened during a timeline slice, right-click the timeline at the point of interest. A context menu displays a list of options that let you open different views or drill down to related items, depending on the particular timeline. To directly drill down to the related main objects or events instead, double-click the timeline .

By default, the **Timeline** tab displays the last 24 hours in the history of a user or device.

Navigating through history

To the left of the date and time ruler, click the button with a triangle pointing to the left to go back in time. Likewise, click the button to the right of the ruler that depicts a triangle pointing to the right to go forward in time. For displaying data further in the past or closer to the present, the ruler and the timelines scroll right or left accordingly, following the opposite direction of the arrow clicked.

Alternatively, hover the mouse pointer over the ruler. The pointer turns into a double-headed horizontal arrow. Click and drag the pointer to the left to go forward in time. To go back in time, click and drag the pointer to the right. The ruler and the timelines scroll as you drag the mouse pointer.

The available history is limited by the amount of events recorded in the in-memory database of the Engine. The Finder stops scrolling to the past once you reach the time of the oldest event in the database. For completeness, the Finder lets you scroll a few hours into the future. It does not make much sense to go beyond the present time though, as the future is naturally empty of data.

Zooming

The default settings of the **Timeline** tab let you see the last 24 hours of a user or device. At that zoom level, every subdivision in a timeline represents a time interval of 30 minutes. With this granularity, two events separated by ten minutes, for instance, may reside in the same time slot, giving the appearance of simultaneity.

To know which event happened first, select an area surrounding the apparently simultaneous events and zoom in:

1. Click the part of the timeline located immediately before the events of interest and keep the mouse left button pressed.
2. Drag the mouse cursor over the events of interest and release the mouse button as soon as you have covered them with a rectangular selection area.
3. Click the magnifying glass with the plus sign that is placed in the top right corner of the timelines or press **Enter**.

This *zoom in* button is enabled only when you have selected an area in the timelines. It also gets disabled when you reach the maximum allowed resolution (one second per subdivision).

Some timelines related to events also propose an option to zoom in in their context menu. As an alternative to the zooming method proposed above, right-click the timeline and select **Zoom in on events** when available.

To zoom out to the previous level, click the magnifying glass with the minus sign in the top right corner of the timelines or press **Backspace**. The *zoom out* button is enabled until you reach the maximum time span allowed (7 days).

To go back to the default 24 hours view, click the house icon placed to the left of the two magnifying glasses.

Timeline sections of the user view

In the timelines of the user view, find events and activities related to the devices with which the user interacted, the print jobs that the user started and the services that the user accessed.

Remember that timelines are actionable. Right-clicking a point in the timeline brings up a context menu with drill-downs and other options to jump to information related to the data in the timeline.

Applies to platforms:

Devices

For every active device linked to the user, find one or several timelines associated to it. If Cross-Engine features are enabled in the Finder, the section includes those devices located on any Engine; otherwise, the section only displays the devices located on the Engine to which the Finder is currently connected. The information displayed in the timelines depends on the platform of the device. For Windows or Mac devices, a main timeline groups all the information available. Click the plus icon to the left of the name of the Windows or Mac device to expand the main timeline into its individual components. Mobile devices display the times of synchronization with the server.

Windows or Mac	Mobile
Device alerts Occurrences of investigation-based alerts. Errors Applications not responding or crashing, bluescreens and hard resets. Warnings Notifications of high cpu load, high memory usage, or a big number input and output operations or page faults. Interaction Times when the user was active on that device (with the keyboard or the mouse), in addition to system boots and user logons.	Exchange ActiveSync synchronization Synchronization of the mobile device with the Exchange server.

Although Windows and Mac devices share the same timelines, note that all the information available for Windows devices is not yet available in Mac devices. Namely, Mac devices display only *hard resets* as errors and lacks IO output and page faults warnings.

For device events to appear in the user view, they must be related to some user interaction with the machine.

Printers

For each printer, find the print jobs that the user has sent. Click the plus icon to the left of the name of the printer to break down the print jobs by device. Each print job appears then on a different timeline depending on the device that the user employed to send the print job.

Services

See the activity of the user in relation to the services that you have defined. Click the plus button to the left of the name of the service to break down the activity by device. Again, if Cross-Engine features are enabled, devices located on a different Engine from the Engine to which the Finder is currently connected are also shown on the list.

Depending on how you defined the service, you can further break down to the activity of the executables that compose the service.

Timeline sections of the device view

In the timeline, you can quickly detect whether the computer generated any alert, experienced any error or warning, had new software installed, connected properly to networked services, etc. This information is presented in different sections.

Note that the timeline is not available for the Mobile platform and that not all its sections are available or complete for the Mac platform. Namely, for Mac, the Errors section displays only hard reset errors, whereas the Web services section does not exist.

Note as well that a same device that has connected to several Engines because it changed its assignment, is seen as a different device by each Engine. Therefore, all the data in the device view comes from a single Engine and there is no merged data, even when the Cross-Engine features are enabled.

From top to bottom, the timeline of the device view displays the sections detailed below.

Applies to platforms:

Alerts

There are two separate sections:

- Global alerts.
- My alerts (user-defined alerts).

Each defined alert has its own timeline. Occurrences of the alert are marked in the timeline, graphically showing their start time and the duration. For the sake of clarity, only alerts that have been triggered during the selected time frame are displayed.

To see the exact time of triggering and the duration of an alert, hover the mouse cursor over the occurrence of the alert. If more than one occurrences of the alert overlap, the hovering tooltip gives you a list of all the occurrences.

To see a list of all the devices that triggered an alert, right-click the mark of the alert in the time-line, choose an occurrence if more than one is available and select **Show Alert**.

Errors

Signal errors in the device, such as application or system crashes. The error is shown in the timeline as a red circle with a number inside. The number inside the circle is bigger than one if more than one error condition overlap in the timeline. Hovering the mouse over the circle gives you a summary of the reason for the error (or the reasons, in the case of overlapping errors).

Warnings

Warnings are represented in the timeline as small boxes. The intensity of the color that fills the box indicates the severity of the warning. The more intense the color is, the more severe is the warning. High memory usage, high IO operations, and high page faults warnings use a yellow shade to signal the condition in the timeline.

On their turn, high CPU warnings signal their condition with two different colors, depending on the particular cause for issuing the warning:

- Yellow, if the overall load in the CPU of the device is high, regardless of the load being caused by the execution of a few or a lot of applications.
- Blue, if some specific applications have a high CPU consumption, but this load is not enough to signal an overall warning for the device.

Hovering the mouse cursor over a warning displays a summary of the reasons for the warning. For example, when hovering over warnings on applications using too much CPU or memory, a tooltip gives you a list of the applications that contributed the most to consumption of these resources.

Activity

In the **Activity** section, you find information about momentary activities, such as the detection of new binaries, print jobs, system boots, user logons and package and patch installations and uninstallations. You find as well information on lasting activities such as executions and connections.

Momentary activities are shown in their own timeline as blue circles with a number inside that indicates the number of overlapping events, similar to the red circles used for displaying errors. Lasting activities, in turn, are shown as blue squared boxes in the timeline, where the brightness of the color indicates the level of the activity (number of executions or connection traffic), similar to the boxes that are used to display warnings. As usual, if the system did not perform any activity of a certain type the activity is not shown at all, instead of displaying an empty timeline.

For every momentary activity, hovering the mouse cursor over the blue circle gives you a summary list of the causes for displaying the activity. For instance, hovering over a **New binaries** occurrence in the timeline displays a list of the binaries whose execution has been detected for the first time at that precise moment. Right-clicking in a blue circle of a momentary activity lets you choose among different options depending on the type of activity.

For lasting activities, that is **Connections** and **Executions**, hovering the mouse over a blue box yields:

- For **Connections**, the amount of traffic registered during the time span of the box.
- For **Executions**, the number of processes run on the time span of the box.

You can drill-down from a box of a lasting activity to the list of individual connections or executions that make it up by right-clicking in the box and selecting **Show connections** or **Show executions**. Connections have an additional option **Show network activity** that lets you navigate directly to a Network activity view and specify the metric to see in it (traffic in, traffic out, failed connections, etc).

In the **Activity** section, yellow color in the timeline warns you about administrator activity. A warning message notifies the use of administration privileges when you hover the mouse cursor over an activity timeline with yellow color. Two kinds of activities use a yellow display when they are carried out by users with administration privileges: **User logons** and **Executions**.

- When a user logs in to a device with administrator privileges, the circle representing the user logon activity is no longer blue, but yellow.
- When a program is run with administrative privileges, the blue boxes that show the executions are crossed by a yellow line to warn that at least one had admin privileges.

Network services

For every defined network-based service, you see a timeline indicating the status of the connections of the selected device to the service. Network connections to the service are displayed again as blue boxes. If any connection problem is detected, the blue boxes are crossed by a yellow line to indicate a warning and by a red line to indicate an error.

To see a summary with the statistics of the connections to the service (total traffic, number of connections, failed connections, response time, etc), hover the mouse over the desired box in the timeline. Additionally, you get a summary list of the errors and warnings that happened during the period delimited by the box, if any.

To open the Service view, click the name of the service at the beginning of the timeline or double-click a box in the timeline. There you find detailed information about the service for the last 24 hours.

Finally, you can also navigate to the Network activity view of the connections to the service from the timeline by right-clicking on any box and selecting **Show network activity**. Double-click in the box, as with connections in the **Activity** section.

Web services

If you installed the **Web & Cloud** module as an addition to the Nexthink Platform, you find a **Web services** section in the device view dedicated to web-based services. This section is very similar to the one dedicated to network-based services.

By hovering the mouse cursor over the boxes in the timeline, you get the statistics about the web-based service: traffic, requests, type of responses, average response time, etc.

To open the Service view, click the name of the web-based service at the beginning of the timeline or double-click a box in the timeline. To navigate to the Web activity view, right-click a box in the timeline and select **Show web activity**.

Users

At the lower part of the device view, you find the timelines that measure the interaction of users with the selected device. There is a timeline for each user that interacted with the machine and the account name of the user is displayed to the left of the timeline. Hovering over each box in the timeline gives you the total duration of the interaction.

For privacy reasons, measurement of the interaction time of the user with the computer can be disabled. If user interaction measurement is disabled, the **Users** section is omitted in the device view.

Click the name of a user to open the corresponding *User view*.

Related tasks

- Comparing the properties of users and devices
- Checking and comparing ratings
- Enabling Cross-Engine Finder features

Related concepts

- Device

Related references

- Errors and warnings for devices and executions
- Alerts tooltips
- Errors tooltips
- Warnings tooltips
- Activity tooltips
- Services tooltips

Observing service performance

Analyzing service status in real-time

To analyze the status of a service in real-time for the last 60 minutes and the last 24 hours, open the Service View in the Finder:

1. Go to the **Services** section in the accordion of the Finder.
2. Double-click the name of the service.

An overview of the service is displayed in a new tab. The overview is divided into two differentiated parts:

- The upper half part, which shows the metrics that measure the performance of the service:

In turn, the metrics that measure the performance of the service are divided into two rows. The first row gives a short-term view of the service for the last 60 minutes and the second row gives a longer-term view for the last 24 hours. For each metric, you get the current value over the period of 60 minutes or 24 hours and the difference with respect to the previous period. For instance, if the time now is 16h40 (measured with a 10 minute resolution), the last 60 minutes span from 15h40 to 16h40 and the difference is computed with respect to the time span from 14h40 to 15h40.

An indicator arrow accompanies the difference value. The arrow points up if the current value is bigger than the previous value and it points down if the current value is smaller than the previous value. In addition, the color of the arrow indicates if the variation on the value of the metric means an improvement or a degradation of the service, because bigger does not necessarily mean better. For instance, an increasing number of failed connections is indeed a bad sign for a service. The arrow looks therefore green when the service improves, red when the service degrades and blue if the metric is merely informative and neutral with respect to the quality of the service.

- The lower half part, which shows the evolution of the selected metric over time:

The evolution of service metrics is shown as a bar chart, which is in turn composed of two parts as well. By default, the top bar chart shows the total values of the metric during the last 24 hours and the charts below break down the total into the partial values according to your defined hierarchies.

Similarly to the indicator arrows that compare the current value of a metric to its previous value, the bar charts also use a color code for quick visual recognition. A red bar indicates an error condition in one or more of the concerned devices, yellow means warning and blue is used for neutral metrics. The metric **Devices with issues** may combine error and warning conditions in a same bar. The amount of red and yellow in the bar is then proportional to the number of devices with errors and warnings respectively. A single device may also have yellow and red combined on its own bars, meaning that the device generated both errors and warnings during the specified period.

Service overview in one or all Engines

When the Cross-Engine features are enabled in the Finder, the option **Overview (all entities)** at the top-left of the service view displays results for all Engines. Select the option **Overview** instead, to display results for the current Engine only.

Breaking down results by hierarchy levels

When the **Overview (all entities)** button is pressed, you can break down the results by the levels defined in your hierarchy in a similar way as you would do in

the Portal:

To focus on the results of a particular node in the hierarchy:

- Select the desired node from the hierarchy navigation bar at the top-left corner of the service view.
- Click the name of the node in the bar chart at the lower part of the service view. The bar chart displays the results of the nodes located one level below the node selected on the navigation bar.

When the Cross-Engine features are disabled or you selected the **Overview** button, you are limited to navigate through the entities of the current Engine. To focus on the results of a particular entity:

- Select the entity from the list placed at the top left corner of the tab window, preceded by the legend **Display data for**.

- Click the name of the entity at the head of its corresponding bar chart.

When you focus on an entity, the top bar chart shows the values of the metric for that entity only. The bar charts below, when present, show the individual devices belonging to the entity and whether they contributed to the value of the metric or not during a particular hour. Therefore, the bar charts for individual devices are only present when the selected metric counts the number of impacted devices. In this case, bars do not denote a quantity level, but just membership of a particular device to the group of impacted devices. If the selected metric computes anything other than a number of devices, such as an average of connections or network traffic, there are no bar charts for the individual devices. Instead, you find the text **no device data in this mode**.

The metrics that count devices and, consequently, show individual device data are the same for both connection-based and web-based services:

- Devices with issues.
- Devices using service.
- Application crashes (this metric does not count the number of crashes but the number of devices with application crashes).

Looking for an entity or device

Look for a particular entity or device in the bar chart by typing its name.

To look for a particular entity or device:

1. Find the text box with the magnifying glass placed to the left of the bar chart, on top of the names of the entities or devices.
2. Type in a part of the name of the entity or the device that you wish to find.

The bar chart displays only the entities or devices containing the exact sequence of characters typed in.

To the right of the text box, there are a couple of arrows that are very similar to the arrows in the **totals** column. You can click the up arrow to arrange the displayed entries in alphabetical order or the down arrow to display the entries in reverse alphabetical order.

Evolution of service metrics over time

Right after you create a service, all its metrics have a zero value and the two rows that display the metrics (the first for short-term metrics and the second for long-term metrics) are both labeled **Last 10 minutes**. This is a consequence of the minimum time resolution of service monitoring. Both rows evolve incrementally every 10 minutes until they reach their maximum window of **Last 60 minutes** for the first row and **Last 24 hours** for the second row.

In the lower half part of the service overview, below the figures of the metrics, you find a bar chart that shows the evolution of the service with time. The bar chart graphically displays the values of the selected metric, as measured in intervals of one hour for the last 24 hours. To change the selected metric, click the name or the value of the metric that you wish to display. A pointer arrow zooms in the selected metric and singles it out from all the other metrics.

In the bar chart, the rightmost interval corresponds to the current hour. The interval of the current hour is divided into sub-intervals of 10 minutes each. The bar that shows the value of the selected metric for the current hour displays the partial results of the metric from the beginning of the hour up to the last 10 minutes sub-interval. The included sub-intervals of 10 minutes are highlighted in the time axis.

Placed to the right of the current hour interval, you find a bar that groups the total value of the metric for the last 24 hours. Contrary to all other bars in the chart, the bar with the total value grows horizontally. Each row in the bar chart displays its own totals, so you can visually compare their values. Rows can be ordered with respect to their total value for the metric. To sort the rows by their total value in the last 24 hours:

1. Find the arrows below the header of the column **totals**.
2. Click the arrow pointing up to sort the rows in ascending order.
Alternatively, click the arrow pointing down to sort the rows in descending order.

Bar chart tooltips

Hovering the mouse cursor over a one hour interval in the bar chart, even when no bar is present in that interval, shows up a small window with the complete summary of the values of the metrics for the selected hour.

The currently selected metric is highlighted in the window. In addition to the metrics, the summary breaks down the number of devices with issues into devices with specific errors and warnings. For instance, if there is an application crash in two different devices during the selected hour, you see the message *2 with application crashes*. Error messages are colored in red and warnings in yellow.

See below the list of possible errors and warnings:

Type	Common	Connection-based	Web-based
Errors	with application crashes	with connection failures lasting more than a minute	with number of 5xx requests over 3 times last week average
Warnings		with response time over 3 times last week average	with average request duration over 3 times last week average with number of 4xx requests over 3 times last week average with number of 3xx requests over 3 times last week average

For more information about service errors and warnings see [Service errors and warnings](#).

Network activity and Web activity

The **Network** and **Web** views are available in the context of the current Engine only; that is, when the **Overview** button is pressed or when the Cross-Engine features are disabled.

To see the individual connections of a connection-based service, click the **Network** activity button, placed on top of the main window of the service and to the right of the **Overview** button, which is pressed by default. The main window of the service displays then a graphical view of all the service connections for the last 24 hours instead of the service overview. To learn more about how to interpret the graphical view of connections, see the chapter on [Viewing network connections](#).

Similarly, to see the web requests of a web-based service, click the **Web** button that replaces the **Network** activity button of connection-based services. You get a graphical view of all the web traffic related to the service for the last 24 hours. See the chapter on [Viewing web requests](#) for an explanation on how to interpret the graphical view of web traffic.

Drill-downs

From the service overview, you can launch contextual investigations related to the service. Contextual investigations and their visualizations let you examine in detail the connections or web requests accessing the service and the activity of devices using the service.

In addition to the buttons of **Network** and **Web** activity, you can launch investigations and visualizations of connections and web requests from the bar chart of the service overview. Double-click a bar in the bar chart to launch a contextual investigation or graphical visualization that depends on the metric selected and is restricted to the entity or device at the head of the row where the bar is placed and the interval of one hour that delimits the graphical bar. For example, if you double-click a bar of the default entity *empty* whose level indicates the number of devices using the service from 12h00 to 13h00, you get a list with all the individual devices belonging to the entity *empty* and making use of the service during that period. As another example, if you double-click a bar that indicates failed connections, you get a visualization of the individual failed connections.

Double-clicking a bar gives you a default investigation or visualization. Alternatively, you can right-click a bar to choose the investigation or visualization that you want to see. When double-clicking or right-clicking a bar in the bar chart of the service overview, results are displayed in a new tab window.

When you focus on an entity by clicking its name in the bar chart or choosing it in the **Display data for** list, you see data of individual devices in the bar chart. Clicking the name of a device opens the device view, which shows detailed information on the activity of the device and is explained in the chapter about observing the activity of users and devices.

Related tasks

- Analyzing service quality
- Creating a service
- Following the evolution of a service
- Hierarchizing Your Infrastructure
- Viewing web requests
- Viewing network connections
- Enabling Cross-Engine Finder features
- Introducing dashboards in the Portal

Related concepts

- Service
- Hierarchy
- Entity

Related references

- Real-time and consolidated service data

Viewing network connections

To visualize network connections in a graphical way, use the *Network activity view*. The Network activity view relates all the objects that participate in a network connection.

To open the Network activity view:

1. Execute an investigation based on connections or on any of the objects that participate directly or indirectly in a network connection: device, user, application, executable, binary, port or destination.
2. From the list of results, click the **Network** button placed at the top of the list.

You can also open the Network activity view from other contexts such as the device view or the services view. In any case, an underlying investigation is generated for you, displaying the Network activity that corresponds to that investigation.

Interpreting the Network activity view

The Network activity view arranges objects in five columns, one per class of object that participates in a connection: device, user, binary, port and destination (applications and executables participate indirectly in a connection through binaries). Straight lines connect objects in contiguous columns whenever the linked objects took part in a network connection. The thickness of the lines indicate a quantity that depends on the kind of information that you selected in the **Display** choice list that is found in the top left corner of the panel. By default, lines display the total traffic, so the thicker the line is, the higher the amount of traffic exchanged during the connection. But you can select to display the number of failed connections, the average response time, the bitrate of the connection, etc. Hovering the mouse over a line gives you the exact quantity of the information selected. A dashed line indicates a zero amount of the kind of

information selected; for instance, you may select to display the total traffic and have a very thick line joining a set of objects and then select to display the number of failed connections and have a dashed line instead, because all connections were successful.

Grouping objects in the columns of the diagram

When a column contains many objects, the Finder collapses them into groups of objects. You distinguish collapsed objects from single objects by the small plus sign that shows on the icon of the collapsed objects. To expand a group of collapsed objects, click the plus sign. If a dotted line appears between expanded objects, that means that they share common traits and they are suitable for collapsing. Click the dotted line and the objects linked by the line will collapse. Alternatively, right-click on an object and select **Collapse** or **Expand**. The right-clicking option allows you to force a collapse of objects even when no dotted line links them. As a special case, binary objects can collapse into executables and applications and be displayed as such in the column dedicated to binaries. Groups of devices and ports can additionally display a star on their respective icons. For a group of devices, the star denotes that they all belong to a single entity, which is a special kind of category. In the case of ports, the star indicates that the icon represents a group of scanned ports. You cannot expand scanned ports to analyze them individually. The group of external destinations is represented by an icon with the shape of a cloud in the column dedicated to destinations. Similarly to a group of scanned ports, the group of external destinations cannot be expanded to individual destinations.

If the number of objects in a column exceeds the vertical space available to show them, an arrow above the column and an arrow below the column help you reach the objects that lie out of bounds. Hover the mouse cursor over the arrows to navigate up or down the objects in the column or click the arrows to navigate quickly.

Navigating through paths and objects

When you click on a line of the Network activity joining two objects, a full path from device to destination is highlighted based on your selection. Right-click the path to drill-down to related objects or activities or execute a related one-click investigation, as you would do from the list results of an executed investigation. You can select several paths at the same time by pressing the **Ctrl** key while you click the paths.

If the Network activity corresponds to an investigation based on objects and not directly based on connections, a list of objects from the results of the

investigation appears to the left of the diagram. For instance, if you execute an investigation on devices and then select the Network activity view, the left hand side of the diagram displays a list of the devices included in the results of the investigation. The list of objects interacts with the displayed lines of the Network activity diagram. If you select a path in the Network activity diagram, the objects that took part in the selected connections are highlighted in the list. The reverse is also true: if you select a specific object from the list, paths representing connections in which the object took part are highlighted in the Network activity diagram. Again, you can select several paths or several objects at the same time by pressing the **Ctrl** key while clicking the lines or the names of the objects. Right-click the name of an object to get the usual drill-down and one-click investigation options associated to the object.

The bar chart of time limited investigations

If the Network activity view relates to an investigation limited in time (full period investigations and investigations specifying **Between** hours are excluded), a bar chart spanning the period of the investigation appears below the diagram of columns. The height of a bar represents a quantity that depends on the type of information selected in the **Display** choice list, in the same way as the thickness of a line in the diagram does. The value of a bar is valid within the time that corresponds to its width. Hover the mouse over a bar to display the numeric value represented and the time interval that the bar spans. The Finder automatically computes the width of the bars and scales them to fit the time frame of the underlying investigation:

- For a maximum time frame of 7 days, a bar represents 2 hours of data.
- For a minimum time frame of 30 minutes, a bar represents 30 seconds of data.

The bars in the chart also interact with the path lines of the Network activity diagram. Click a bar and the associated paths will be highlighted in the diagram. Click a line of the Network activity diagram and the corresponding sections of the bars will be highlighted in the bar chart. Once again, you can select several lines or several bars by clicking them while you press the **Ctrl** key. Right-click a bar or a group of selected bars to drill-down to related objects or activities or to execute one-click investigations.

Zooming in and out

To limit the number of lines in a diagram to those that correspond to one or more bars in the bar chart, use the **zoom in** icon (the magnifying glass with a plus sign) placed on the top right corner of the Network activity diagram. Selecting

one or more bars in the bar chart enables the **zoom in** icon. Click the **zoom in** icon and only the lines that relate to the bars selected will remain displayed in the diagram. After zooming in, come back to the original time frame by clicking the **zoom out** icon that lies to the right of the **zoom in** icon.

In a similar way, you can reduce the number of paths in a diagram to those selected using the zoom. Select one or more paths in the diagram and click the **zoom in** icon. Only the selected paths and their related objects remain displayed in the diagram. Click the **zoom out** icon to come back to the previous zoom level.

Limits of the diagram

If the investigation involves a big amount of connections, the Network activity view may not be able to display all the corresponding paths. When the limit of ten thousand paths is exceeded, a warning icon appears in the top right corner of the diagram, to the left of the zoom icons, meaning that only partial results are shown in the diagram.

To see the Network activity diagram in full screen mode, click the growing square icon that is placed to the right of the zoom icons. This is specially useful in diagrams with lots of connections to better distinguish the different paths. To come back to the original view of the diagram, click the shrinking square icon that replaces the growing square icon in full screen mode.

Related concepts

- Connection

Related references

- Incoming traffic measurement

Viewing web requests

If you added the **Web & Cloud** module to your installation of the Nexthink Platform, you can visualize web requests in a graphical way with the *Web activity view*. The Web activity view relates all the objects that participate in a web request in much the same way that the Network activity view displays network connections, since a web request is a special kind of connection.

To open the Web activity view:

1. Execute an investigation based on web requests or on any of the objects that participate directly or indirectly in a web request: device, user, application, executable, binary, port or destination.
2. From the list of results, click the **Web** button placed at the top of the list.

You can also open the Web activity view from other contexts such as the device view or the web services view. In any case, an underlying investigation is generated for you, displaying the Web activity that corresponds to that investigation.

Interpreting the Web activity view

The Web activity view arranges the objects that participate in a web request in six columns. The five first columns are the same columns of the Network activity view: device, user, binary, port and destination. The Web activity view adds a new column: domain. The domain column exposes the visited web domains and relates them to internal or external destinations. As in the Network activity view, straight lines link the objects of contiguous columns. The lines represent the kind of information selected in the **Display** choice list, found in the top left corner of the diagram, and their thickness is an indication of the amount of data displayed. You can choose to display ingoing or outgoing web traffic, successful or failed HTTP or TLS connections and their duration and response times.

Grouping domains

The objects in the columns of the Web activity view can be collapsed into groups or expanded exactly in the same way as the objects in the Network activity view. In addition, there are some special groups of domains that are created when collapsing individual domains and that you can see when you hover the mouse cursor over the icon of the domain group:

Aliases

Groups domains whose name is not a fully qualified domain name. Aliases usually appear when end-users configure automatic search for local domains. For example, the domain *svn* can be an alias for *svn.intranet.example.com*.

Unnamed

Groups domains that do not have a name. Individual unnamed domains are represented by their IP address.

Named

Fully qualified domain names are grouped by their top-level domain name (com, org, country code, etc).

If there are many domains and they do not fit in their column, you can reach all domains by scrolling with the arrows that appear above and below the column, exactly as you would do with any other column that is also in the Network activity view.

Navigating through paths and objects

The navigation in the Web activity view follows the same rules as in the Network activity view. You can select paths or objects and drill-down or execute one-click investigations in the same way as in the Network activity view. The difference is that the Web activity view does not show information about connections in general, but only about web requests. On the other hand, you have additional information about domains.

When you click on a line of the Web activity joining two objects, a full path from device to destination is highlighted based on your selection. Right-click the path to drill-down to related objects or activities or execute a related one-click investigation, as you would do from the list results of an executed investigation. You can select several paths at the same time by pressing the **Ctrl** key while you click the paths.

If the Web activity corresponds to an investigation based on objects and not directly based on web requests, a list of objects from the results of the investigation appears to the left of the diagram. For instance, if you execute an investigation on devices and then select the Web activity view, the left hand side of the diagram displays a list of the devices included in the results of the investigation. The list of objects interacts with the displayed lines of the Web activity diagram. If you select a path in the Web activity diagram, the objects that took part in the selected connections are highlighted in the list. The reverse is also true: if you select a specific object from the list, paths representing connections in which the object took part are highlighted in the Web activity diagram. Again, you can select several paths or several objects at the same time by pressing the **Ctrl** key while clicking the lines or the names of the objects. Right-click the name of an object to get the usual drill-down and one-click investigation options associated to the object.

The bar chart of time limited investigations

The functionality of the bar chart in the Web activity view is also very similar to that of the bar chart in the Network activity view. Thanks to the bar chart, you can see a timeline of the web activity and establish the correspondence between bars of activity and paths in the diagram.

If the Web activity view relates to an investigation limited in time (full period investigations and investigations specifying **Between** hours are excluded), a bar chart spanning the period of the investigation appears below the diagram of columns. The height of a bar represents a quantity that depends on the type of information selected in the **Display** choice list, in the same way as the thickness of a line in the diagram does. The value of a bar is valid within the time that corresponds to its width. Hover the mouse over a bar to display the numeric value represented and the time interval that the bar spans. The Finder automatically computes the width of the bars and scales them to fit the time frame of the underlying investigation:

- For a maximum time frame of 7 days, a bar represents 2 hours of data.
- For a minimum time frame of 30 minutes, a bar represents 30 seconds of data.

The bars in the chart also interact with the path lines of the Web activity diagram. Click a bar and the associated paths will be highlighted in the diagram. Click a line of the Web activity diagram and the corresponding sections of the bars will be highlighted in the bar chart. Once again, you can select several lines or several bars by clicking them while you press the **Ctrl** key. Right-click a bar or a group of selected bars to drill-down to related objects or activities or to execute one-click investigations.

Zooming in and out

Limit the number of web requests displayed in the diagram to those that you select using the zoom tool. You can equally limit them to the selected bars in the bar chart.

To limit the number of lines in a diagram to those that correspond to one or more bars in the bar chart, use the **zoom in** icon (the magnifying glass with a plus sign) placed on the top right corner of the Network activity diagram. Selecting one or more bars in the bar chart enables the **zoom in** icon. Click the **zoom in** icon and only the lines that relate to the bars selected will remain displayed in the diagram. After zooming in, come back to the original time frame by clicking the **zoom out** icon that lies to the right of the **zoom in** icon.

In a similar way, you can reduce the number of paths in a diagram to those selected using the zoom. Select one or more paths in the diagram and click the **zoom in** icon. Only the selected paths and their related objects remain displayed in the diagram. Click the **zoom out** icon to come back to the previous zoom level.

Limits of the diagram

The diagram is not able to show more than ten thousand paths. When the maximum number of paths is exceeded, a yellow warning icon shows up in the top right corner of the diagram to inform you that only partial results are displayed.

Viewing executions

To visualize program executions in a graphical way, use the *Local activity view*. The Local activity view relates all the objects that take part in the execution of a program.

To open the Local activity view:

1. Execute an investigation based on executions or on any object that takes part in the execution of a program: device, user, application, executable or binary.
2. From the list of results, click the **Local** button placed at the top of the list.

You can also open the Local activity view from other contexts such as the executions section of the device view. In any case, the Local activity view is a graphical representation of an underlying investigation. When you open the Local activity view from a context other than an investigation, the Finder automatically creates the investigation for you.

Interpreting the Local activity view

The Local activity view arranges objects in three columns, one per class of object that participates in an execution: device, user and binary. The Local activity view considers applications and executables as objects that hold binaries. They can appear in the same column as binaries when these are visually collapsed into groups.

The lines that join the object in the columns can represent the number of executions or the duration of the executions in which the connected objects took part, depending on the type of information that you selected in the **Display** choice list in the top left corner of the diagram. The thicker the line that links the objects, the higher the number of executions or the longer the duration of the executions that it represents. You can also make the lines display information about network traffic aggregated during an execution using the **Display** choice

list. For detailed information on individual connections, use the Network activity view that is documented in the chapter on Viewing network connections. Alternatively, use the Web activity view, documented in the chapter on Viewing web requests, if you are interested in details about web traffic.

Hovering the mouse cursor over a line in the diagram gives you the exact quantity of the kind of data that the line represents, that is, the kind that you selected in the **Display** choice list. A dashed line indicates zero data.

Grouping objects in the columns of the diagram

The Local activity view shares with the Network activity and the Web activity views the same mechanism for collapsing and expanding objects in the column diagram. When a column contains many objects, the Finder collapses them into groups. You distinguish collapsed objects from single objects by the small plus sign that shows on the icon of the collapsed objects. To expand a group of collapsed objects, click the plus sign. If a dotted line appears between expanded objects, that means that they share common traits and they are suitable for collapsing. Click the dotted line and the objects linked by the line will collapse. Alternatively, right-click on an object and select **Collapse** or **Expand**. The right-clicking option allows you to force a collapse of objects even when no dotted line links them. As a special case, binary objects can collapse into executables and applications and be displayed as such in the column dedicated to binaries. A collapsed group of devices that displays an additional star on the icon denotes that all members of the group belong to a single entity, which is a special kind of category.

If the number of objects in a column exceeds the vertical space available to show them, an arrow above the column and an arrow below the column help you reach the objects that lie out of bounds. Hover the mouse cursor over the arrows to navigate up or down the objects in the column or click the arrows to navigate quickly.

Navigating through paths and objects

When you click on a line of the Local activity diagram, a full path from device to binary is highlighted based on your selection. Right-click the path to drill-down to related objects or activities or execute a related one-click investigation, as you would do from the list results of an executed investigation. You can select several paths at the same time by pressing the **Ctrl** key while you click the paths.

If the Local activity corresponds to an investigation based on objects and not directly based on executions, a list of objects from the results of the investigation

appears to the left of the diagram. For instance, if you execute an investigation on devices and then select the Local activity view, the left hand side of the diagram displays a list of the devices included in the results of the investigation. The list of objects interacts with the displayed lines of the Local activity. If you select a path in the Local activity diagram, the objects that took part in the selected executions are highlighted in the list. The reverse is also true: if you select a specific object from the list, paths representing executions in which the object took part are highlighted in the Local activity diagram. Again, you can select several paths or several objects at the same time by pressing the **Ctrl** key while clicking the lines or the names of the objects. Right-click the name of an object to get the usual drill-down and one-click investigation options associated to the object.

The bar chart of time limited investigations

If the Local activity view relates to an investigation limited in time (full period investigations and investigations specifying **Between** hours are excluded), a bar chart spanning the period of the investigation appears below the diagram of columns. The height of a bar represents a quantity that depends on the type of information selected in the **Display** choice list, in the same way as the thickness of a line in the diagram does. The value of a bar is valid within the time that corresponds to its width. Hover the mouse over a bar to display the numeric value represented and the time interval that the bar spans. The Finder automatically computes the width of the bars and scales them to fit the time frame of the underlying investigation:

- For a maximum time frame of 7 days, a bar represents 2 hours of data.
- For a minimum time frame of 30 minutes, a bar represents 30 seconds of data.

The bars in the chart also interact with the path lines of the Local activity diagram. Click a bar and the associated paths will be highlighted in the diagram. Click a line of the Local activity diagram and the corresponding sections of the bars will be highlighted in the bar chart. Once again, you can select several lines or several bars by clicking them while you press the **Ctrl** key. Right-click a bar or a group of selected bars to drill-down to related objects or activities or to execute one-click investigations.

Zooming in and out

To limit the number of lines in a diagram to those that correspond to one or more bars in the bar chart, use the **zoom in** icon (the magnifying glass with a plus sign) placed on the top right corner of the Network activity diagram. Selecting

one or more bars in the bar chart enables the **zoom in** icon. Click the **zoom in** icon and only the lines that relate to the bars selected will remain displayed in the diagram. After zooming in, come back to the original time frame by clicking the **zoom out** icon that lies to the right of the **zoom in** icon.

In a similar way, you can reduce the number of paths in a diagram to those selected using the zoom. Select one or more paths in the diagram and click the **zoom in** icon. Only the selected paths and their related objects remain displayed in the diagram. Click the **zoom out** icon to come back to the previous zoom level.

Limits of the diagram

If the underlying investigation involves a big amount of executions, the Local activity view may not be able to display all the corresponding paths. When the limit of ten thousand paths is exceeded, a warning icon appears in the top right corner of the diagram, to the left of the zoom icons, meaning that only partial results are shown in the diagram.

To see the Local activity diagram in full screen mode, click the growing square icon that is placed to the right of the zoom icons. This is specially useful in diagrams with lots of executions, to help you better distinguish the different paths. To come back to the original view of the diagram, click the shrinking square icon that replaces the growing square icon in full screen mode.

Related references

- Incoming traffic measurement

Monitoring IT custom metrics

Creating a metric

Overview

Use metrics to measure and survey over time the characteristics of those elements in your IT infrastructure that interest you the most. Metrics reduce complex queries to a single figure, table, or chart that lets you quickly evaluate the state of key components in your IT ecosystem.

To create metrics in the Finder, you must be a user with granted access to the Finder and have the right to edit system content: categories, services, alerts, and metrics.

Creating a metric in the Finder

To create a new metric in the Finder:

1. Log in to the Finder as a user with the appropriate permissions.
2. Select the **Metrics** section from the accordion on the left-hand side of the main window.
3. Right-click the header of the **Metrics** section or the empty area below it.
4. Select **Create new metric** from the menu.
5. Complete the form to define the metric:
 1. Type in the name of the metric.
 2. Optional: Give a short description on the function of the metric.
 3. Fill in each section of the form as described below.
6. Optional: Click **Run as investigation** to execute the metric as if it were a quick investigation. This option gives you an instant preview of the results that you can expect from the metric.
7. Click **Save** to store permanently the definition of the metric.

Retrieve section

Specify the kind of objects used to compute the metric In the **Retrieve** section:

1. In the upper-right part of the section, tick the boxes of the platforms to which the metric applies: Windows, Mac OS or Mobile.

2. Select the type of the objects on which to base the computation of the metric: users, devices, applications, etc.

The choice of a particular kind of object has a logical impact on the grouping options and aggregated values that you can use to define the metric.

Compute daily section

The Portal computes the value of each metric once every day, hence the name of the section. For longer periods of time, the Portal combines the values obtained for a metric over several days in the way that you specify.

In the **Compute daily** section, choose the type of metric that you want to create. Depending on the type of metric, specify how to compute aggregated values for the metric or how to group the retrieved objects when seeing the details of the metric in the Portal.

1. Choose the type of metric from the following options. The word *objects* is replaced by the kind of objects that you selected in the **Retrieve** section:
 - ◆ Count metric
 1. Select **the total number of [active] objects** to create a metric that counts the selected type of objects. The **active** keyword, when present, indicates that only those objects that were seen during the last day are taken into account for the calculating the value of the metric for that day. Otherwise, all objects in the full history of the Engine are taken into account. The presence of the **active** keyword depends on the type of objects being counted:
 - The **active** keyword is mandatory for applications, executables, binaries, ports, destinations, domains, and printers.
 - The **active** keyword is necessarily absent for packages.
 - For users and devices, choose to count either all the users or devices or only those that were active the last day.
 - ◆ Quantity metric (only available for *device* and *user* objects)
 - ◇ For device objects:
 1. Select **the aggregate of active devices** to create a metric that computes a single value by combining the values of the chosen aggregate in the retrieved devices.

2. Choose the particular device activity, score, or output of remote action on which to base the metric from the list. Only outputs that yield a numeric value are available for creating quantity metrics.
 - If you choose to combine the values of either a score or the output of a remote action, select between **active devices** or **all devices** depending on whether you want to limit the retrieved devices to those that were active during the last day or not.
- ◇ For user objects:
 1. Select **the score / EUF opinion scale of all users** to create a metric that computes a single value by combining user scores or the answers given by the users to an opinion scale question.
 2. Choose the particular user score or opinion scale question on which to base the metric from the list.
- ◆ Top metric
 1. Select **the top N objects with highest / lowest aggregate** to create a metric that keeps a top list of objects with the highest or lowest value of the specified aggregate.
 2. Choose between 10 to 100 (in increments of 10) the number N of objects to display in the top list.
 3. Set either *highest* or *lowest* to order the top list of objects by descending or ascending value of the aggregate respectively.
 4. Select the particular aggregate on which to base the top list.
2. Select the grouping options or aggregate computation when applicable:
 - ◆ For count and quantity metrics, optionally select how the Portal groups the results when you observe the details of the metric. In the **Group by** list:
 - ◇ Select - **none** - if you do not need to group the results.
 - ◇ Select a field or a category to group objects by the value of that particular field or category.
 - Optional: Add another field or category after the keyword **and** to define a second criterium to *group by* the objects.
 - ◆ For quantity and top metrics, select how to compute the aggregated values for the metric in **Aggregated by**.
 - ◇ For quantity metrics, describe how to combine the individual values associated to each one of the retrieved devices. The Portal applies the selected aggregation method to compute the metric for groups of devices (such as those that share a

same location in the hierarchy) and for time frames that span longer than one day. Depending on the aggregate (only the options that make sense for the particular aggregate are available), choose among:

- **sum over all devices and the whole timeframe**
- **maximum value per device per day**
- **average value per device per day**
- **minimum value per device per day**

◇ For top metrics, describe how the aggregate is computed individually for each object in the top list for periods of time longer than one day. Depending on the aggregate (only the options that make sense for the particular aggregate are available), choose among:

- **sum over the whole timeframe**
- **maximum value per day**
- **average value per day**
- **minimum value per day**

Matching section

In the **Matching** section, specify a set the conditions that the objects in the **Retrieve** section must meet for taking part in the computation of the metric and, optionally, the hours of the day when the retrieved objects must match those conditions.

The procedure to define these conditions in the **Conditions** part of the definition of a metric is very similar to setting the conditions of an investigation. Conditions are a set of rules that apply to any item (object, activity, or event) that is related to the objects selected in the **Retrieve section**:

To add the conditions:

1. In the **Conditions** subsection, click the link **Click here to add a new condition**. The placeholders for the condition fields show up.
2. Set the object, activity, or event to which the condition applies.
3. Set the attribute or category that you want to constraint.
4. Set the operator for comparison (e.g. *is*, *is not*, *starts with*, etc).
5. Set the matching value, if you selected an attribute constraint, or the matching keyword, if you selected a category constraint.
6. Optional: When setting more than one condition, click on **Advanced** to edit the logical expression that combines the conditions. By default, conditions are combined by a logical AND, so all must be fulfilled.

7. Optional: Specify a condition on an aggregate value after the **and** keyword at the bottom of the conditions section. The value of the aggregate is computed for the whole day unless you specify a **Between** clause (see below).

Any type of metric accepts the values of opinion scale questions and scores as constraints for conditions. In that way, you can create a metric that counts the number of users who gave an answer with value greater than 5 to a particular opinion scale question, for example.

On the other hand, some combinations of metrics, conditions, and additional display fields are incompatible. In particular, count metrics that take into account inactive objects are incompatible with setting a condition on activities or events, or when using a *between* clause. Under the hood, the computation of these metrics is based on investigations over the full available period in the Engine; therefore, they have similar limitations. An error icon (a small white cross inside a red circle) appears at the selection for counting the **total number of objects** to indicate the mismatch. There is an exception to this rule: you can set conditions on installations when retrieving the total number of packages. In general, if you add a condition and an error icon appears on its right side, the condition is in conflict with another condition or with one of the chosen attributes to display. By hovering the mouse over the error icon, see the reason for the conflict in a tooltip.

To delete a condition:

1. Click the trash icon to the right of the condition fields.

For count metrics on active objects, an additional setting specifies how objects are counted when the time frame in the Portal is longer than one day:

1. Choose either **at least one day** or **the last active day** from:
For periods longer than one day, count *objects* that meet conditions on *at least one day* / *the last active day*

As a guideline, select **at least one day** when you want the metric to account for event occurrences during the period. For instance, a metric that counts the number of devices with hard resets. For a device to be counted, the hard reset can happen any time during the interval. On the other hand, select **the last active day** for metrics that have an inventory function. For example, a metric that counts the number of devices with two monitors. Indeed, if a device does not have two monitors on the last day of the period, it must not be added to the inventory.

Note that this setting is not available for count metrics that count both active and inactive objects. Indeed, it does not make much sense to see the values of this kind of metrics for periods longer than one day. Widgets based on metrics that count all objects are disabled in the Portal when a period longer than one day is selected.

Optionally, restrict the time of the day to which the computation of the metric applies by ticking the check box and setting the start time and the end time of the computation in the last line of the **Matching** section:

- **Between *HH:MM* and *HH:MM*.**

Options section

The contents of the **Options** section depend on the type of metric that you selected.

For count metrics, the following options are available:

1. Tick **Include ratio** to get a percentage in addition to a cardinal number as the displayed value of the count metric in the Portal.
 - ◆ The percentage is computed as the number of objects that match the conditions defined in the metric divided by the number of objects that match another set of conditions.
 - ◆ Define the new set of conditions in the same way as you defined the conditions for the metric with the tools that show up when you tick the **Include ratio** check box.
2. Tick **Include variation and threshold indicator** for the Portal to display changes in the value of the metric with respect to its last value and indicate whether a value increase or decrease is a positive or a negative thing for you. Select one from these three options:
 - ◆ **variation indicator only**
 1. By default, a green arrow up and a red arrow down are displayed below the statement, meaning that an increase in the value of the metric is a positive thing. For instance, an increase in the number of devices with antivirus is a good thing.
 2. To change the default and indicate that a value increase of the metric is a negative thing, click **invert** to switch the two arrows. For example, an increase in the number of bluescreens is a bad thing.
 - ◆ **variation indicator and 1 threshold**

1. Choose the direction of the variation indicator by clicking **invert** if appropriate (see first option).
2. Set the value of the threshold on top of the horizontal arrow. Only when the increase in the value of the metric is higher than the threshold, the indicator changes color (to green if increase is good, to red if it is bad).
3. If you ticked the option **Include ratio**, you may want to define the threshold as a ratio:
 1. To the right of the variation message, choose between **absolute** or **ratio** values (Include *variation message* based on *absolute / ratio* values).
 - Choose **absolute** to express the threshold in number of objects, as in the case described above.
 - Choose **ratio** to express the threshold in a percentage related to the ratio value computed for the metric.
 - ◆ If you choose *ratio* you can ignore the threshold value when grouping results by hierarchy in the Portal if the number of devices in a node of the hierarchy is too small. Tick the option **Ignore threshold if there are less than N devices in a hierarchy node** and set a value for *N*.

◆ **variation indicator and 2 thresholds**

1. Choose the direction of the variation indicator by clicking **invert** if appropriate (see first option).
2. Set the thresholds in the same way as explained in the 1 threshold case. You can also express the thresholds in percentage values if you ticked the **Include ratio** option. The only difference with the 1 threshold option is the addition of the yellow color:
 - If the increase in value of the metric is higher than the first threshold, the variation indicator turns yellow.
 - Only when the increase in the value of the metric is higher than the second threshold, the variation indicator changes color (to green if increase is good, to red if it is bad).
3. Set the **Additional display fields** using the label selector in the same way as you select the columns to display in an investigation. The fields selected as criteria to group the results in the **Group by** list are automatically added as display fields and they cannot be removed.

1. Later, select in the Portal which of these fields to display when seeing the details of the metric.

For quantity metrics, only the options for the variation indicator and thresholds are available. Refer to step 2 of count metrics for instructions. Ignore the suboptions related to ratio, since these are not supported by quantity metrics.

In the case of top metrics, only the option to set additional display fields is available. Refer to step 3 of count metrics for instructions.

Related tasks

- Editing the options of an investigation
- Executing an investigation

Related concepts

- Metric

Examples of metrics

Overview

Creating a new metric may be a daunting task for beginners because of the many options available. To help you with the creation of metrics, let us walk through a first example that covers the creation of three metrics, each one of a different type: count, quantity, and top. Then let us explore how to create a quantity metric based on the output of a remote action.

The first example gets information on binaries that are considered dangerous. To that end, we propose the creation of three metrics, which the reader can later refine and expand:

- Devices executing dangerous binaries (Count metric)
Count the number of devices that execute dangerous binaries.
- Cumulated execution time of dangerous binaries (Quantity metric)
Measure how long your devices were exposed to the execution of dangerous binaries.
- Top most executed dangerous executables (Top metric)
List the top ten executables associated to dangerous binaries by number of executions.

In this example, we consider a binary to be dangerous when its **Threat level** field is set to *high threat*. Nexthink automatically sets the value of this field via the application library. You may later come up with your own definition of a *dangerous* binary and adapt the conditions in the example metrics accordingly.

The second example of a quantity metric gets information about the battery status of laptop devices. As this is not an information that Nexthink retrieves by default, we need a remote action that returns on demand data:

- Average battery health (Quantity metric)
With the help of a remote action, get the average battery health of all devices and compare by manufacturer.

For every step in the creation of the metrics that requires the choice of an option, we explain our decision in detail. We assume however that you know the basics of creating a metric.

Count metric

The first metric reflects the number of devices impacted by the execution of dangerous binaries.

Create a count metric in the Finder and edit its options:

1. Type in the name of the metric: **Devices executing dangerous binaries**.
2. Optional: Type in a description for the metric.
3. In the **RETRIEVE** section, click **devices**.
4. In the **COMPUTE DAILY** section:
 1. Select the option **the total number of all devices** to create a *count* metric.
 2. Assuming that we might be interested in the status of the antivirus of those devices executing dangerous binaries, choose **Group by *antivirus up-to-date* and *antivirus RTP*** to classify the devices by the update status of their antivirus and their activation of the real-time protection.
5. In the **MATCHING** section:
 1. Add the condition **Binary *Threat level* is *high***.
 2. Leave the default option **Count devices that meet conditions on *at least one day* in period**. We want to count the devices that executed a dangerous binary anytime within the observed interval (that is, the period that you set in the navigation tool of the Portal when watching the results of the metric). We do not select thus the option *the last active day*, which is intended for metrics that have an

- inventory function.
6. In the **OPTIONS** section:
 1. Tick the box **Include ratio** without including any new condition. In that way, you compare the number of impacted devices with the total number of devices.
 2. Tick the box and select the option **Include variation indicator only**. We do not need to set any threshold and we keep the default option for the sense of the variation: an increase in the value of the metric is bad (red arrow up) and a decrease of its value is good (green arrow down).
 3. Optional: Tick any of the **Additional display fields** that you want to add.

Quantity metric

From aggregate values

As second metric, let us measure for how long dangerous binaries have been executing on the devices.

Create a new metric and edit its options:

1. Type in the name of the metric: **Cumulated execution of dangerous binaries**.
2. Optional: Type in a description for the metric.
3. In the **RETRIEVE** section, click **devices**, since quantity metrics can only be selected for devices.
4. In the **COMPUTE DAILY** section:
 1. Select the second option to create a *quantity* metric and build the sentence: **the cumulated execution duration of active devices**.
 2. In the **Group by** option, keep the default - **none** -, as we do not need to break down the results.
 3. In the **Aggregate by** option, select **sum over all devices and the whole timeframe**. We are interested in the total execution time over all devices and not in the average execution time per device, which is the other available option.
5. In the **MATCHING** section:
 1. Add the condition **Binary Threat level is high**.
6. In the **OPTIONS** section:
 1. Tick the box and select the option **Include variation indicator and two thresholds**. We want to set warning and error conditions if the cumulated execution time of dangerous binaries exceeds some values.

2. In the bar to indicate the thresholds, keep the sense of variation (red arrow up, green arrow down) and set the first threshold to **10 min** and the second to **1 hours**.

From numerical outputs of remote actions

Build a third metric directly based on the output of a remote action that yields a numerical value. In general, scores let you build quantity metrics for devices based on any kind of output from a remote action, but only numerical outputs are allowed as a direct value in quantity metrics.

To create a quantity metric that measures the average health of the battery on all your devices and compares by manufacturer:

1. Install the required Library pack for the example.
 1. Open the Nexthink Library page of the Battery Status pack from your favorite web browser.
 2. Click the **Install** button.
2. Log in to the Finder as a user with the right to create metrics.
3. Right-click the **Metrics** section and a context menu shows up.
4. Select **Create new metric** from the menu.
5. Type in the name of the metric: **Average battery health**.
6. Optional: Type in a description for the metric.
7. In the **RETRIEVE** section, click **devices**, since quantity metrics can only be selected for devices.
8. In the **COMPUTE DAILY** section:
 1. Select the second option to create a *quantity* metric and build the sentence: **the *get Battery Status / Battery1 Health* of all devices**.
 2. In the **Group by** option, select **device manufacturer** as first grouping criterion and **- none -** as the second, to make it possible to break down the results by the manufacturer of the device.
 3. In the **Aggregate by** option, select **average value per device**.
9. In the **MATCHING** section:
 1. Add the condition **Get Battery Status / Battery1 Health is greater or equal to 0.1** to exclude devices that return zero instead of a valid value.

Top metric

Finally, let us add a metric that retrieves the top 10 most executed executables whose binary representations are considered dangerous. Remember that an executable in Nexthink groups the different versions (binary images) of a program in a single object. In this case, a metric retrieving executables is probably more convenient than a top metric retrieving the individual binaries. Indeed, having a list of different executables is preferable to seeing different binary versions of the same executable repeated in a list.

Create a new metric and edit its options:

1. Type in the name of the metric: **Top most executed dangerous executables**.
2. Optional: Type in a description for the metric.
3. In the **RETRIEVE** section, click **executables**.
4. In the **COMPUTE DAILY** section:
 1. Select the second option to create a *top* metric and build the sentence: **the top 10 executables with *highest number of executions***.
 2. In the **Aggregate by** option, select **maximum value per day**. A perfectly valid option as well would be **sum over the whole timeframe** to see the total number of executions of each executable. For this time, however, we want to classify the executables by their maximum burst of executions in one day and, in that way, find out the dangerous executables which are run more aggressively. We are not much interested either in the other available aggregation option **average value per day**, because we want to detect the extreme cases.

5. In the **MATCHING** section:
 1. Add the condition **Binary Threat level is high**.
6. In the **OPTIONS** section:
 1. Optional: Tick any of the **Additional display fields** that you want to add.

Conclusion

We hope that these examples have helped you clarify some of the concepts behind the creation of a metric. Keep on reading to know how to create widgets in the Portal to display the values of the metrics in the Portal. For more information on how the Portal computes and presents metric data, read this article on aggregation and grouping.

Related tasks

- Creating a metric

Related references

- Portal aggregation and grouping
- Nextthink Application Library
- Nextthink Library

Following the evolution of a metric

Overview

After defining a metric in the Finder, follow the evolution of its value from the Portal. Choose among different visualizations (*widgets*) for your metric.

To add metrics to the Portal, you must be a user with the right to create personal dashboards. If you do not have the right to create your own dashboards, you can still see the dashboards attributed to your roles.

Adding a metric to the Portal

To add a metric to the Portal, insert it into a dashboard. You may either reuse an existing dashboard or create a new one. In any case, the dashboard must belong to a module of type *Basic*.

Dashboards in Basic modules are initially empty, displaying just a blank page where you can add and organize widgets for your metrics. Combine several types of widgets in a dashboard to see your metrics from different points of view.

To create a new Basic module and dashboard:

1. Click the menu icon (three bars) on the right-hand side of the dark blue ribbon.
2. Select **Create new module...** at the bottom of the menu.
3. Choose **Basic** as the type of module to create.
4. A new Basic module with a default empty dashboard appear in your Portal.
5. Optional: Rename the module and the dashboard (by default **Untitled module** and **Untitled dashboard**) by clicking the menu icon again and selecting **Rename....**
 1. Type in the new names for both the module and the dashboard under **Module name** and **Dashboard Name**.
 2. Click **Done**.

Alternatively, choose an existing dashboard from the module navigation tool that you find on the left-hand side of the dark blue ribbon:

1. Click the module navigation tool (it displays the names of the current module and dashboard).
 - ◆ If you have the permissions to see published content, and there actually are some modules published, you may see the rubrics **My content** and **All published**.
 1. Click **My content** if this is the case.
2. Select a Basic module from the section **PERSONAL**. Only those modules that belong to you are available for editing.
 - ◆ If you do not see the section **PERSONAL**, you do not have permissions to create dashboards. Ask your administrator.
 - ◆ If the **PERSONAL** section is empty of modules, it displays the message **No personal module. Click here to create one.:**
 - ◇ Click the word **here** in the message to create your first personal module as an alternative to the method seen above.
 - ◇ Select a dashboard from any of the other sections (dashboards included in your roles or under the **All published** rubric, if available) and copy it to your personal section to be able to modify it:
 1. Click the menu icon on the right-hand side of the blue ribbon once the dashboard is open.

2. Select **Copy module to my content** from the menu.
3. Select one of the available dashboards in the module.

Once you have either the new or the existing dashboard in your screen, add the metric to it:

1. Click the menu icon on the right-hand side of the dark blue ribbon.
2. Select **Edit content** in the **DASHBOARD** section. The dashboard is now in edit mode.
 - ◆ If the dashboard was empty, a plus sign appears on it.
 - ◆ If the dashboard has content, plus signs appear while you hover the mouse over the limits of existing widgets.
3. Click one of the plus signs to add a new widget for your metric at that location of the dashboard.
 - ◆ Note that widgets can hold more than one metric. If you prefer to add your metric to an existing widget, click the sprocket icon that appears when you hover the mouse over the top-right corner of a widget and select **Edit**.
4. Choose the type of widget that you want to create:
 - ◆ KPI.
 - ◆ Table.
 - ◆ Line chart.
 - ◆ Bar chart.
5. Fill-in the dialog to add the widget:
 1. Optional: Type in a title for the widget.
 2. Click the button **Add metrics**.
 3. Select a metric from the list of available metrics.
 4. Click the button **Add**. The list of metrics turns into a preview of the widget.
 5. Depending on the type of widget that you chose, set the **DISPLAY** options.
 6. Optional: Click **Add metrics** again to add as many metrics as you want to the widget.
 7. Click **Done** to finish editing the widget and come back to dashboard editing.
6. To finish editing the dashboard, click the check mark that replaces the menu icon on the right-hand side of the dark blue ribbon while in edit mode.

Related references

- Types of widgets

Finding the visuals of a metric

Overview

Right after you add a metric to a widget in a dashboard, the link between the metric and its visual components in the Portal —widget and dashboard— is clear for you. After adding multiple metrics to several widgets in different dashboards though, it is easy to lose track of which metrics relate to which widgets and dashboards.

Learn here how to navigate from a metric in the Finder to the dashboards that hold the metric in the Portal and find out the widgets within a dashboard that display the values of the metric.

Finding the dashboards and widgets linked to a metric

To find out the dashboards and widgets that make use of a particular metric:

1. Log in to the Finder.
2. In the **Metrics** section, right-click the name of the desired metric.
3. Select **Show in Nextthink Portal** from the menu.

- ◆ Select **Personal** to look for the metric on dashboards that you created yourself, located under **DASHBOARDS My content PERSONAL** in the Portal.
- ◆ Select the name of a module (if any available) that is part of your roles to look for the metric on dashboards included in the module. In the example above, you can find the modules **Support Level 1** and **Windows 10 migration**, for instance.

- ◆ Select **All published** to look for the metric on dashboards that you have the right see under **DASHBOARDS All published** in the Portal.
- ◆ If the metric is not used on any dashboard that you have the right to see in the Portal, the option to show in Nexthink Portal is disabled. Even if you are an administrator with access to all published dashboards, beware that this does not necessarily imply that the metric is not used in the Portal at all, because it may be linked to an unpublished Personal dashboard of another user. The tooltip message clearly indicates this fact: **This metric is not used in any of your dashboards; it might be used by another user**

4. Click any of the listed dashboards to open the dashboard in the Portal. The Finder opens a new tab on your default browser to show the dashboard.
 - ◆ If you are not logged in yet, the Portal asks you to log in first.
5. Once the dashboard is open in the Portal, find the widgets that use the metric thanks to the distinctive dotted blue border that highlights them. For instance, in the dashboard below, only the widget **Devices with high CPU** holds the metric of interest:

Monitoring IT services

Analyzing service quality

Overview

Today, companies strongly depend on a set of IT Services to carry out their daily work. Although the particular set of deployed IT services may vary from one company to another, the correct functioning of services is vital for the day-to-day business of every organization.

With Nexthink, define the IT services that you want to survey by specifying their characteristic resources and see how well those IT services perform inside your corporate network. Thanks to the user-centered approach of Nexthink, you get a glimpse on the quality of the service as it is perceived by the end-users. See in real-time the status of a service and find out who are the impacted users when the service degrades. In the case of service degradation, Nexthink may help you determine where the problem lies. For long term analysis, follow the evolution of the service throughout extended time intervals.

Nexthink divides IT services into two types:

Connection-based services

Services that rely on TCP connections for providing their solutions to the users. Both Windows and Mac devices report information on connection-based services.

Applies to platforms:

Web-based services

Web sites or web applications that use the HTTP protocol or its encrypted variant with TLS (HTTPS) to provide a service to their users. Web-based services also rely on TCP for transport, but defining the service at the HTTP level lets you filter the service by the accessed web domains or URLs. For the moment, only Windows devices report information specific to web-based services.

Applies to platforms:

Service metrics

The type of service determines the exact metrics that are used to measure its performance. The set of metrics is nevertheless very similar for both connection-based and web-based services. Most of the metrics are shared or serve a similar purpose in both types of services. See below a list of all the metrics where the shared metrics are emphasized:

Connection-based service metrics	Web-based service metrics
<ul style="list-style-type: none">• <i>Devices with issues</i>• <i>Devices using service</i>• Failed connections• <i>Application crashes</i>• <i>Average response time</i>• Incoming traffic• Outgoing traffic• Incoming bitrate• Outgoing bitrate	<ul style="list-style-type: none">• <i>Devices with issues</i>• <i>Devices using service</i>• Failed web requests• <i>Application crashes</i>• <i>Average response time</i>• Incoming web traffic• Outgoing web traffic

Analyzing services in the Finder and the Portal

To analyze service quality, Nexthink offers you two possibilities:

The Service View in the Finder

The Service View provides you with a real-time view of the service, including the values of the service metrics for the last 60 minutes and the last 24 hours. Break-down results according to entities and, if Cross-Engine features are enabled in the Finder, according to the nodes in your hierarchies. Drill-down to individual connections or devices, and quickly switch to the Network or Web views directly from the Service View.

The service dashboards in the Portal

Service modules summarize the status of your services in a special overview dashboard. Moreover, for each individual service in the module, a detailed service dashboard provides you with a full view of the service key metrics through a comprehensive set of widgets. In addition to a real-time view for the last 60 minutes and the last 24 hours, the Portal lets you select longer periods of time to see historical values and trends in your service-related data.

Related tasks

- Creating a service

- Observing service performance
- Following the evolution of a service
- Enabling Cross-Engine Finder features

Related concepts

- Service

Related references

- Real-time and consolidated service data

Creating a service

Overview

Define services in the Finder, based either on web connections or on lower-level TCP connections, by specifying one or more of the following:

- The destinations (servers) that provide the service.
- The web domains that host the service (only for web-based services).
- The executable files of the applications that access the service.
- The ports that applications use to connect to the service.
- The devices that consume the service (useful for limiting the devices monitored rather than for defining the service).

Nexthink automatically associates the network connections that match the given specifications to the service thus defined, helping you keep track of the status and utilization of your IT services.

In addition, to survey the health status of services more easily, set thresholds on the number of devices with issues that access a service. Color codes -green, yellow, and red- help you distinguish whether the service is doing fine or not, depending on whether the thresholds are respected or exceeded.

Procedure

To create a new service in the Finder:

1. Select the **Services** section from the accordion on the left-hand side of the main window.

2. Right-click the header of the **Services** section or the empty area below it.
3. Select **Create new service** from the menu. The Finder displays the form to define the service.
4. Give a name to the service.
5. Optional: Give a short description of the service.
6. Select the type of service:
 - ◆ Check **Network connections** to create a connection-based service.
 - ◆ Check **HTTP and TLS web requests** to create a web-based service.
 - ◆ Check **HTTP web requests with URL path** to create a web-based service by specifying a URL path. Remember to configure the Collector to send path information.
7. Specify the devices that connect to the service.
 - ◆ **any**: Record the connections from any device to the service.
 - ◆ **with name**: Record only the connections to the service of the device with the given name.
 1. Type in the name of the device. You can use the substitution characters * and ? as indicated in the tooltip.
 - ◆ **with IP network**: Record only the connections of the devices with at least one of their IP addresses belonging to the given subnetwork.
 - ◆ **with keyword**: Record only the connections of the devices tagged with the given keyword.
 1. Select a category of devices from the list.
 2. Select a keyword of the chosen category from the list.
8. Specify the executables that connect to the service.
 - ◆ **any**: No matter what executable can establish a connection to the service.
 - ◆ **with name**: Only the executable with the given name connects to the service.
 1. Type in the name of the executable. You can use the substitution characters * and ? as indicated in the tooltip.
 - ◆ **with keyword**: Only the executables tagged with the given keyword connect to the service.
 1. Select a category of executables from the list.
 2. Select a keyword of the chosen category from the list.
9. Specify the TCP ports reserved for the service.
 - ◆ **any**: No particular TCP port is reserved to connect to the service.
 - ◆ **with number**: The service is bound to the TCP port with the given number.
 1. Type in the port number.

- ◆ **with keyword:** Only the TCP ports tagged with the given keyword are characteristic of the service.
 1. Select a category of ports from the list.
 2. Select a keyword of the chosen category from the list.
10. Specify the destinations that provide the service.
- ◆ **any:** The destination of the connection is not relevant for defining the service.
 - ◆ **with IP address:** Only the connections to the destination with the given IP address are associated to the service.
 1. Type in the IP address of the destination (e.g. 192.168.3.1).
 - ◆ **with IP network:** The service is offered by a group of destinations belonging to the given subnet.
 1. Type in the IP subnetwork in CIDR notation (e.g. 192.168.3.0/24).
 - ◆ **external:** The service is offered by destinations outside the monitored networks.
 - ◆ **with keyword:** Only the destinations tagged with the given keyword offer the service.
 1. Select a category of destinations from the list.
 2. Select a keyword of the chosen category from the list.
11. When creating a web-based service, specify the domains that belong to the service:
- ◆ **any:** The service is not restricted to any domain in particular. This option is not available if you are creating a web-based service with URL path.
 - ◆ **with name:** Only the domain with the specified name belongs to the service.
 1. Type in the name of the domain. You can use the substitution characters * and ? as indicated in the tooltip.
 - ◆ **with keyword:** Only the domains tagged with the given keyword belong to the service.
 1. Select a category of domains from the list.
 2. Select a keyword of the chosen category from the list.
 - ◆ Additionally, if you are creating a service with URL path:
 1. Type in the **URL Path** that the service will recognize as its own.
12. Recommended: Tick the box **Set threshold on devices with issues** to specify the minimum quantity of impacted devices needed for considering that the service is degraded or down. The message changes to **Set 1 threshold / 2 thresholds on devices with issues, considering errors only / errors and entity warnings**.
- ◆ Select **1 threshold** to define two statuses for the service: ok (green), or error (red).

In the colored bar below the message, set the required percentage of devices that must have issues for the service to turn its status from ok to error.

- ◆ Select **2 thresholds** to define three statuses for the service: ok (green), warning (yellow), or error (red).

In the colored bar below the message, set the required two percentages of devices with issues for the service to turn its status from ok to warning, and from warning to error, respectively.

- ◆ Select **errors only** to count as devices with issues only those devices in the error state.
- ◆ Select **errors and entity warnings** to also count in those devices that belong to entities in the warning state, in addition to devices with errors or in an entity in the error state (see Service errors and warnings).

1. Optional: Tick the box **Ignore unless at least *N* devices are impacted in a hierarchy node** to take into account the percentages indicated as thresholds only if *N* or more devices within a node of the hierarchy have issues. This is useful to avoid false alarms (services in error state) in nodes with few devices.

13. Click on **Save** to store your new service permanently.

Follow the evolution of services either from the service view of the Finder or from the Portal. Setting thresholds helps you quickly distinguish those services that are working as expected from those which may require intervention.

Limit on the number of services

Using this procedure, define as many services as you wish. However, for performance reasons, a maximum of **100 services** may be simultaneously enabled. If you have defined more than 100 services, those in excess are disabled; meaning that they do not receive any connection nor, in consequence, display any data..

Related tasks

- Following the evolution of a service
- Specifying URL paths of web-based services

Related references

- Service errors and warnings

Related concepts

- Service
- Device
- Executable
- Port
- Destination
- Category
- Keyword

Following the evolution of a service

Overview

In Observing service performance you have learned how to interpret the service view in the Finder. The Portal offers you a different perspective on services, allowing you to see not only real-time data, but also longer time intervals for historical data and trend analysis.

Adding a service dashboard to the Portal

To add a service to the Portal, create a dashboard for it. You can either add the dashboard to an existing service module or create a new service module. Only modules of type *service* can hold service-based dashboards.

To create a new service module:

1. Log in to the Portal as administrator or as a user with the right to create personal dashboards.
2. Click the menu (three bars) icon on the right-hand side of the dark blue ribbon.
3. Select the last option **Create new module...**
4. Select **Service Monitoring** to create a new module for hosting services.

Alternatively, navigate to an existing service module using the navigation tool on the left-hand side of the dark blue ribbon. The tool displays the names of the module and dashboard that are currently active in the Portal:

1. Click the names of the current module and dashboard.
2. Select a service module from the section **My content**, since only those modules that belong to you are available for editing.

3. Open the **Overview** dashboard or any other dashboard of a service in the module.
4. Click the menu icon on the right-hand side of the dark blue ribbon.
5. Select **Configure module....**

Once you have created a new service module or entered the configuration of an existing service module, add the desired service (or services) to it:

1. From the list of available services, select those that you want to see in the module and click the button with the arrow pointing to the right. Services move from the **AVAILABLE** list to the **SELECTED** list.
2. Optional: If you selected some services by mistake, select them again and click the button with the arrow pointing to the left. Services move from the **SELECTED** list back to the **AVAILABLE** list.
3. Optional: All service modules have an **Overview** dashboard where you can see the list of all other dashboards in the module, each one representing a different service. Select how you want to arrange this list of services and their status break down, whether in rows or in columns, by setting the option **In overview dashboard, display services as rows / columns**.
4. To end adding the service (or services), click **Create**, if you are creating a new module, or **Done**, if you are modifying an existing module.

Specifying URL paths of web-based services

If you have purchased the Web and Cloud module, you can define services in Nexthink for measuring the connectivity of the end-users to web applications and web sites by stating the domain name of the particular provider of the service. In some cases though, you may want to differentiate among several services that lie within the same domain. For instance, you may want to make the difference between:

- www.example.com/agenda
- www.example.com/share

To treat these two addresses as different services, you must first configure your Collectors to report the URLs of the web requests addressed to the domain www.example.com.

Then create two services in the Finder of the type **HTTP web requests with URL path**. In addition to the name of the domain, you can specify now a URL path to distinguish services within the same domain. Note however that the URL

path option is only available for HTTP and not for TLS web-based services. Following with our example, you would specify as URL paths:

- agenda
- share

All subpaths are naturally included in the definition. You may also use the star pattern to substitute for zero or more characters in the URL path.

Domain	URL path	Matches
www.example.com	agenda	<ul style="list-style-type: none">• www.example.com/agenda/mycal.php
	share	<ul style="list-style-type: none">• www.example.com/share/myshare.html• www.example.com/share/groupshare.html
	sh*	<ul style="list-style-type: none">• www.example.com/share/myshare.html• www.example.com/sharing/index.html

Related tasks

- Creating a Service
- Reporting the URL of HTTP web requests

Engaging with the end user

Getting feedback from the end users

Overview

Nexthink offers you deep visibility into your IT infrastructure from the perspective of the end users. With Nexthink, you get extensive information about the hardware of the end-user devices, the software that is installed and run on these devices, and the usage that end users make of the network. All this objective information is very useful for your IT department to know what is actually happening within your corporate network.

To complement this objective view of your IT infrastructure, it is important to know as well the subjective opinion of the end users. Measure the IT satisfaction level of the end users by asking them directly about their perception on the different IT services and pieces of equipment that are provided to them. Target only the concerned users by designing pertinent investigations in the Finder and keep track of their responses with dedicated dashboards in the Portal.

Analyze the results and compare objective information with subjective IT satisfaction to get a better and more complete understanding of your IT environment as a whole.

Planning a campaign

In Nexthink, the act of starting a survey to retrieve the opinion of the end users with respect to some aspect of their IT environment is known as *publishing a campaign*.

Campaigns are a powerful tool to let you know the opinion of your end users about a particular subject. The power of campaigns comes to some extent from the fact that you can address questions to potentially all of the end users with a single action. Contrary to the creation of an investigation or a metric in the Finder, the publication of a campaign has a direct impact on the end users. When you publish a campaign, users receive either notifications that prompt them to take the survey (if notifications are enabled) or the survey itself. Because answering the questions of a campaign may disrupt the daily work routine of the end users, you must plan your campaigns with special care.

To design a successful campaign, consider the following points:

- The goal of the campaign.
 - ♦ State what you want to achieve with the help of all stakeholders.
 - ♦ Define the actions to take once the campaign has ended.
- The target audience.
 - ♦ Clearly define the group of users targeted by the campaign.
 - ♦ Adapt the language of the questions in the campaign to the target audience.
- The length and duration of the campaign.
 - ♦ Ask only a few questions per campaign.
 - ♦ Choose the type of campaign according to your previously defined goals:
 - ♦ One-off campaign.
 - ♦ Recurring campaign.
 - ♦ Continuous Satisfaction Measurement campaign.
 - ♦ Control the time for collecting answers in order to take action based on feedback.

When you publish a campaign, you want to maximize:

- The number of targeted users who answer your questions.
- The quality of their answers.

And minimize:

- The burden on the users.

Implicating the end user

People will tend to respond to a campaign if they feel concerned by its subject matter. If you define well the target audience of your campaign, users who are part of it should be willing to give their opinion on the subject. Remember, however, that this might not always be the case and that, by asking the users to participate in a survey, you are somehow asking them for a favor. Therefore, try to create a trust relation with the end user: use greetings and personalize the notifications of your campaigns, provide the reasons why you are asking about a particular matter, and take action based on the answers received.

Write the questions of a campaign ensuring that the targeted audience understands your questions. IT technicians tend to use a terminology that is often obscure to most of the non-technical users. Therefore, it is important to avoid technical jargon and adapt the language of your questions to the group of end

users that you are addressing. For instance, when you are addressing to a group of generic users, ask *How happy are you with the sending speed of your email?* instead of *How satisfied are you with the performance of the SMTP server?*

Defining the target audience

To define the target audience of a campaign, create or reuse an investigation based on users. Remember that you can create new investigations by saving the results of drill-downs or one-clicks. These are often very useful for clearly defining the target audience of campaigns which are related to a particular issue.

For instance, imagine that you find out that a group of devices are having high response times when accessing the mail server thanks to Nexthink. You may want to drill-down from devices to users and save the resulting investigation. Then, you can use this investigation in a campaign to get feedback from those users only who have been impacted by the high response times.

Services defined in Nexthink are usually a good starting point to define a campaign as well. From the **Services** tab in the Finder, it is easy to get the users that had access to the service during the last day by right-clicking the name of a service and selecting **Display > Objects > Users**. You can also get those devices that accessed the service (or had trouble accessing the service) from the Service view. Similarly to the previous case, drilling-down from devices to users gives you the target audience for a service-related campaign.

To refine the investigations that determine the target audience of your campaigns, filter out the users who should not receive it. For instance, system users will never answer the questions of a campaign; therefore, it is useless to include them in the investigation. You may also want to leave a group of special users out of the loop of notifications. Use categories and keywords to tag these special users and add appropriate conditions to the investigations of the campaign to filter them out.

Because the investigation on users that is linked to a campaign is periodically re-evaluated while the campaign is active, the time frame of investigations used in campaigns must be either of type *full period* or *during last days / hours*, and not a *between* fixed interval. Regardless of the results of the last re-evaluation of the linked investigation and while the campaign is still active, users who got a notification from a campaign at some point keep receiving notifications (with the frequency specified in the definition of the campaign) until they either fully answer the questions of the campaign or decline to participate. That is, once a user has been targeted, the user remains targeted.

Timing considerations

To make your campaigns more effective, do not overwhelm the end users with too many questions. A few precise questions are often enough, helping users stay focused and provide better quality answers; whereas lengthy campaigns may discourage users to finish the survey and not read attentively the last questions. Typically ask from one to five questions per campaign.

To communicate urgent or important matters that all targeted users should read, consider skipping the notification step of a campaign. Instead of displaying a notification, the first (and usually the only) question of the campaign is displayed right away on the screen of the end user. Thus, by saving the end user an unnecessary click to accept the campaign, the campaign cannot be rejected and its message is more likely to be read.

Depending on the type of campaign, regularly control the period dedicated to collecting feedback in order to take action based on the answers of the end users. For one-off campaigns, the moment to take action usually corresponds to the end of the campaign, when it is retired. For recurring and continuous campaigns, take action periodically before the same groups of users are targeted again.

Phases of a campaign

The life-cycle of a campaign comprises the following phases:

Created

The campaign has been saved but not published yet. You can edit the campaign and make any kind of modifications to it.

Published

Once you publish a campaign, the campaign becomes active and end users start receiving the set of questions associated to the campaign. You can edit a campaign that has been published with some restrictions.

Retired

When the time to end the campaign comes, stop sending questions to the end users by retiring the campaign. After the campaign is retired, its results are still available in the Engine and in the dashboards that you may have created in the Portal. You may later republish the same campaign or modify the campaign, with some restrictions, before republishing.

Republishing a campaign that was retired makes exactly the same campaign active again. Users may be retargeted depending on the type of campaign and on the status of the users with respect to the campaign (whether they fully

answered, declined to participate, or other).

To republish an old campaign with new questions or other major modifications, you need to create a new campaign. A quick way for reusing an old campaign is to create the new campaign by duplication: right-click the name of the old campaign in the **Campaigns** tab and select **Duplicate**. Modify the new campaign to your taste and publish it when you are ready.

Before publishing a campaign, it is convenient to get a glance of how the end users will be notified (if the campaign enables notifications) and how they will see the questions of the campaign in their own computers. Click the **Preview** button at the bottom of the window where you edit the campaign to see your recently edited questions in the same form as the end users will see them. Every time that you make a change to the campaign, click **Preview** again to avoid bad surprises once you finally publish the campaign.

Remember to take action based on the results of a campaign. This will ensure that the end users see the value of providing quality answers and motivate them to keep answering future campaigns.

Related concepts

- Campaign

Related references

- Types of Campaigns

Types of campaigns

Overview

There are a variety of situations in which you may want to know the opinion of the end users about a particular subject. Depending on the nature of the subject, specially on its temporal characteristics, asking for feedback usually requires different approaches. For instance, you may want to know the reaction of the end users to a single change in the IT infrastructure and ask them about it only once. In the case of a recurring event, however, it may be appropriate to repeatedly ask end users for feedback at every occurrence of the event. In other cases, such as when willing to know how end users perceive the quality of an IT service, it makes sense to continuously measure the evolution of the end-user satisfaction.

Three different types of campaigns let you adapt the measurement of the end-user experience to these usage patterns:

One-off campaign

Ask for feedback to the end users only once.

Recurring campaign

Repeatedly ask the end users for feedback about a matter that happens time and time again.

Continuous Satisfaction Measurement campaign (or continuous campaign for short)

Continuously estimate the overall satisfaction of end users with respect to a particular subject.

Specifying the type of campaign

When creating a campaign, indirectly specify its type by choosing the appropriate options in the **RECIPIENTS** section.

The first option lets you choose how the campaign targets users, which may already determine the type of the campaign:

- When targeted **manually** or **with a remote action**, the result is a special kind of recurring campaign, as the campaign can target users again manually, through the Engage API, or through a remote action. The way to prevent resending the campaign to the same user in a short period of time depends on the chosen option (see quiet period below).
- When targeted **with an investigation**, the type of campaign is determined by two additional options that appear just below the placeholder for the investigation that targets users.

The first option of a campaign that targets users with an investigation reads:

- **The campaign will be sent to (...)**
 - ◆ To create either a one-off or a recurring campaign, choose **all target users**.

One-off and recurring campaigns are sent to all the users retrieved by the investigation associated to the campaign when it is evaluated.
 - ◆ To create a continuous satisfaction measurement campaign, choose **a different sample of target users every day (continuous measurement)**

Only a small fraction of the total target population actually receives the campaign each day. The system computes a

sample of the target users from the total set of users retrieved by the associated investigation with the aim of maximizing the representativeness of the sample, while minimizing the annoyance of the end users (not all users need to answer the campaign).

Because all targeted users receive one-off or recurring campaigns, the second option for these types of campaigns reads:

- **Target users will receive the campaign**

- ♦ To create a one-off campaign, select **only once** from the list.
- ♦ To create a recurring campaign, select **again** from the list. A second part of the sentence shows up, letting you define the recurring or *quiet* period; that is, the time between a user answers a campaign and the same user receives the same campaign again:
 - ◊ Enter a number followed by a time unit (to choose from a list) as the quiet period.

Because only a sample of the target users receive a campaign each day in continuous campaigns and because continuous campaigns always define a quiet period, the second option for continuous campaigns switches to:

- **Target users can receive the campaign again after**

- ♦ Enter a number followed by a time unit (to choose from a list) as the quiet period.

Characteristics of a campaign

Target evaluation and re-evaluation

In campaigns that target users manually, a Finder user selects the users from a list view or from a user view and then triggers the campaign. Therefore, there is no automatic evaluation of targets. To manually trigger a campaign, the campaign must be published beforehand.

In campaigns targeted with remote actions, the targets are determined by the remote action itself. The users who interact with the devices targeted by the remote action receive the campaign, which must have been published beforehand.

In campaigns targeted with investigations, the evaluation of targets starts at the moment of publishing the campaign. When you publish a campaign, the campaign executes its associated investigation to retrieve target users for the

first time. In the case of a continuous campaign, it additionally computes a sample of the target users from the results of the investigation.

After publication, the set of target users is periodically re-evaluated by executing the investigation again and again (and taking a sample, if the campaign is continuous). The frequency of re-evaluation depends on the type of campaign, but it is always relative to the time of the first evaluation:

- For one-off and recurring campaigns: 10 minutes.
- For continuous campaigns: 24 hours.

That is, for one-off and recurring campaigns, re-evaluation happens frequently. For continuous campaigns, it happens once every day. Re-evaluation ends when the campaign is retired.

Sending the campaign to the end users

An end user that is retrieved by the investigation associated to a one-off or a recurring campaign, or which is part of the sample of target users selected by a continuous campaign, receives the campaign:

- The process of sending the campaign to end users starts immediately after the campaign is published.
- The sending process stops as soon as the campaign is retired.
- If the campaign is republished, the sending process restarts straightaway.

Users who are offline (that is, not logged in to a device) when the campaign is sent cannot receive it immediately. Instead, the moment when offline users receive the campaign depends on the type of campaign:

- For one-off and recurring campaigns: users receive the campaign the next time that they are targeted.
- For continuous campaigns: users receive the campaign as soon as they log in to a device.

Reminder of notifications

An end user who got a notification from a campaign that has notifications enabled is reminded to provide the answers to the questions of the campaign with a new notification if:

- The user did not answer all the questions of the campaign.
- The user did not decline to participate in the campaign.

Once a user gets a notification from a campaign, the user keeps receiving reminder notifications to finish answering the questions even if that particular user is no longer part of the selected target users (after the re-evaluation of the investigation associated to the campaign). Reminder notifications appear thus in the device of the end user with the frequency specified in the definition of the campaign until one of the following conditions is met:

- The user fully answers the campaign.
- The user declines to participate in the campaign.
- The campaign is retired.

Reminder notifications work with all types of campaigns, except for those launched from remote actions, and are independent of the quiet period.

Quiet period

The quiet period is a method of recurring and continuous campaigns to avoid bothering users more than is strictly necessary. Because of the repetitive nature of these types of campaigns, users may have to respond to the questions of a campaign more than once. The quiet period defines the minimum time interval between the moment when a user answers all the questions of a campaign and the moment when the same user receives the same campaign again.

- The quiet period of recurring and continuous campaigns that are targeted with an investigation is specified in the definition of the campaign.
- The quiet period of campaigns that are targeted with a remote action is controlled by the triggering period of the associated remote action.
- The quiet period of manually targeted campaigns is fixed to two hours to prevent accidental re-triggering.

Effects of retiring and republishing a campaign

When a campaign targeted with an investigation is retired, it stops evaluating the target users and stops being sent. If the campaign is republished, it must be emphasized that it is still the same campaign that is relaunched again, which is different from starting anew. The history of the campaign is kept and, in particular, the status of the end users in relation to the campaign. The effects of republishing a campaign on the end users depends thus on their status and on the type of campaign:

- On users who declined to participate
 - ◆ Republishing the campaign has no effect on these users, regardless of the type of campaign. They already declined to

- participate and will never receive the same campaign again.
- On users who fully answered the campaign
 - ◆ For one-off campaigns, republishing has no effect, since the users already answered the campaign and will not receive it any more.
 - ◆ For recurring and continuous campaigns, users who already answered the campaign may be retargeted again, but always respecting their quiet period.
- On users who were notified, but neither answered completely nor declined
 - ◆ For one-off campaigns, users are retargeted immediately, as long as they are still part of the results of the investigation that retrieves target users. Otherwise, users do not receive reminder notifications right away, but only when they become target users again.
 - ◆ For recurring and continuous campaigns, users may be retargeted respecting their quiet period, which is enforced in this case even though users did not fully answer the campaign in its previous incarnation. In the same way as in one-off campaigns, until users are not selected again as targets, they do not receive reminder notifications.

To target again users who declined to participate or users that fully answered a one-off campaign, you must create and publish a new campaign.

Retiring a manually targeted campaign implies that no Finder user can trigger the campaign, unless it is republished again. In the latter case, the quiet period of two hours applies to the retargeting of users who already answered the campaign.

Retiring a campaign targeted with a remote action implies that the remote action will not be able to display the campaign. Republish the campaign for the remote action to work properly.

Creating a campaign

Procedure

To create a new campaign:

1. Log in to the Finder as a user with the right to create campaigns.
2. Select the **Campaigns** section on the left-hand side of the main window.
3. Right-click the header of the **Campaigns** section or the empty area below it.

4. Select **Create new campaign** from the context menu. The form to create the campaign shows up in a new tab.
5. Click the temporary name **Untitled campaign n** and replace it by an appropriate name for the campaign.
6. Optional: Briefly describe the purpose of the campaign in the field with the text **Enter optional description here....**
7. Provide personal information about the owner of the campaign in the **SENDER** section. The information about the sender is displayed both in the notification pop-up and when end-users take the campaign, to help them identify who is requesting their attention:
 1. As **Name**, type in the name of the sender.
 2. As **Title**, type in the job title of the sender.
 3. Optional: Click the icon under **Picture** and choose an image that identifies the sender.
8. In the **NOTIFICATION** section, specify how your users are notified when the campaign is published, in addition to the sender information:
 1. Optional: Tick **Skip notification** for the campaign to jump over the notification step when being shown to the end users (useful for communicating important matters). Note that you still need to provide the rest of the fields in the **NOTIFICATION** section for backwards compatibility with Collectors V6.20 and lower, which still show the notification regardless of the value of this setting.
 2. As **Subject**, enter the main theme of the campaign.
 3. Optional: As **Footer**, type in additional text that appears below the subject of the campaign in the notification pop-up.
 4. From the **Reminder Frequency** list, select how often the notification is repeated for a particular user, until the user completes the survey. The possible values are:
 - ◇ Twice per day
 - ◇ Once per day
 - ◇ Twice per week
 - ◇ Once per week
 5. Optional: Choose the persistence of the notification on the screen of the end user under **Notification is always visible**:
 - ◇ **until user closes it**: the notification stays indefinitely on the screen until the end user either accepts or declines the campaign.
 - ◇ **for 60 seconds**: the notification hides after 60 seconds of no action from the user.
9. In the **RECIPIENTS** section, select first how to target users by completing the sentence **Users are targeted**:
 - ◆ Choose **with an investigation** to select the target users by means of an investigation.

1. Drag and drop an investigation on users into the **RECIPIENTS** section.
2. Optional: Right-click the investigation and select **Run...** to see the list of users that will compose the initial audience of the campaign (depending on the settings of the investigation, the audience may vary with time). Beware that you are connecting to a single Engine with the Finder, so you may get less results than the campaign itself in a multi-Engine setup.
3. Select the desired type of campaign by combining the two available targeting policies. In the option **The campaign will be sent to:**
 1. Choose to send the campaign to **all target users** to create a one-off or recurring campaign. In the option **Target users will receive the campaign:**
 - Select **only once** to create a one-off campaign.
 - Select **again** and a quiet period to create a recurring campaign.
 2. Choose to send the campaign to **a different sample of the target users every day** to create a continuous satisfaction measurement campaign.
 - Select a quiet period for the option **Target users can receive the campaign again after x [time unit]**.
- ◆ Choose **with a remote action** to use the campaign in a self-help scenario of Nexthink Act.
- ◆ Choose **manually** to later trigger the campaign either manually from the Finder or through the Engage API.
10. In the **QUESTIONS** section, compose the questions to be answered by the targeted users.
 1. Type in an identifier for the question in **Question name**. This is the name displayed in the Finder and the Portal when referring to the question.
 2. Choose the type of question from the **Type** list:
 - ◇ **Single answer:** The user must select only one out of all the available options.
 - ◇ **Multiple answer:** The user may tick several options from the list of all those possible. At least one must be selected for the user to be able to continue.
 - ◇ **Opinion scale:** This is basically a single answer question where each choice gets assigned a numerical value.
 3. Type in the actual question in **Text displayed**.

4. Specify the available choices for the question in **Choices**. For each choice:
 1. Type in the text that the user reads as an option to choose in the **Text displayed** column.
 2. Type in an identifier for the choice in the **Name** column. This is the name displayed in the Finder and the Portal when referring to this particular choice.
 3. Optional: If the question was of the opinion scale type, assign a numeric value to it in the **Value** column. By default, any new opinion scale choice gets the maximum value from those assigned to the previous choices plus one.
 4. Optional: If the question was of the single answer or opinion scale types, click the *step to* icon that displays a turning arrow to alter the natural order of the questions in the campaign when the user selects the current choice as answer. The *step to* icon is not available for choices of multiple answer questions, as the user may select more than one choice.
 1. From the drop-down list that shows up, choose the question to which the user that selects the current choice will jump.
 5. Optional: Click **Add choice** to create an additional choice and repeat the previous steps.
 6. Optional: Click the up and down arrows to the right of a choice to reposition it in the desired order.
 7. Optional: Click the bin icon to the right of a choice to remove the choice from the list.
5. Optional: Click the general *step to* icon for the question if the next question to be asked is not the next question defined.
 1. From the drop-down list that shows up, choose the question to which the user will jump after answering the current question. Note that the *step to* options of individual choices take precedence over the general *step to* option for the question.
6. Optional: If the question is of the multiple answer or the opinion scale types, tick **Optional text** to add text field to the question where users can freely express themselves.
 1. Type in a **Label** for the text field.
7. Optional: Under **Optional content**, write introductory text to the survey (to be displayed on the first question) and some final words (to be displayed after the last question).
 1. Under **On first question**, enter the text of the different introductory elements to the campaign:

- **Greetings:** Brief salutation formula. To personalize your greetings, use a word between braces to refer to the user (e.g. {user}, {colleague}, {co-worker}, etc). In the greetings, the placeholder is replaced by the actual name of the user who receives the notification, if the system is able to retrieve it. Otherwise, the notification displays the word without the braces.
- **Header:** Message displayed above the first question. Useful to clarify the purpose of the campaign to the user.
- **Footer:** Message displayed below the first question. Useful to give additional info about the campaign process (estimated duration, deadline, contact information, etc).

2. Under **After last question**, type in some acknowledgment text as **Final message**.

11. Optional: Under **Translations**, click the link **Click here to add a new translation** for translating the campaign to a different language.
12. Optional: Before saving your campaign, click **Preview** to inspect how the end-users will see the campaign in their own devices. The preview function relies on the Collector being installed in the same machine that runs the Finder.
13. To save your new campaign, either:
 - ◆ Click **Save** to keep your changes without publishing the campaign. The name of the campaign appears grayed out to indicate that the campaign remains in the idle state.
 - ◆ Click **Save & Publish** to store the changes and publish immediately your campaign. Targeted users start receiving notifications to take the survey. Once the campaign is published, you can edit only a restricted set of its features.

Inserting links into campaigns

To provide additional information about a campaign and give end-users the possibility to contact support either by email or instant messaging, insert appropriate links into any of the following sections:

Sender	On first question	Question	After last question
<ul style="list-style-type: none"> • Name • Title 	<ul style="list-style-type: none"> • Greetings • Header • Footer 	<ul style="list-style-type: none"> • Text displayed 	<ul style="list-style-type: none"> • Final message

The links must have the following formats depending on their type:

Web links

[Link text] (<http://www.example.com/campaign-info>)

Email links

[Link text] (<mailto:someone@example.com?subject=Optional%20subject>)

Instant messaging links

[Link text] (<im:sip:username>)

When a user clicks a link in a campaign, the OS opens the application associated to the protocol.

Brand image on campaigns

Further customize the appearance of campaigns with your own logo and colors to reflect your brand identity and encourage end-users to answer.

Applies to platforms:

Maximum number of published campaigns

To gather the answers from the end users, campaigns add custom fields to the user object. In addition, the traffic between the Collectors and their corresponding Engine also increases as a result of asking end-user feedback.

To keep the impact of publishing campaigns on both the network traffic and the performance of the Engine within bounds, the maximum number of simultaneously published campaigns is limited. According to the type of campaign, create up to:

- 15 one-off or recurring campaigns.
- 15 continuous satisfaction measurement campaigns.
- 100 manually triggered campaigns.
- 1000 campaigns embedded in remote actions.

Note that the limit on the number of available custom fields may actually lower the maximum number of campaigns that you can publish.

Related tasks

- Translating a campaign
- Creating an investigation
- Triggering a campaign manually

- Triggering campaigns via their API
- Hierarchizing your infrastructure
- Branding of campaigns

Related concepts

- Campaign

Related references

- Campaign display compatibility
- Maximum supported values

Editing a campaign

Overview

Once a campaign is created, you can reopen it in the Finder for editing. If the campaign has not been published yet, you can modify it entirely. On the other hand, if the campaign has already been published or retired, there are some restrictions with respect to the features of the campaign that you can modify.

Editing a campaign

To edit a campaign:

1. Log in to the Finder as a user with the right to create campaigns.
2. Select the **Campaigns** section on the left-hand side panel of the main window.
3. Double-click the name of an existing campaign, or right-click and select **Edit**.
4. Modify the desired fields. If the campaign is already published or retired, only a subset of the campaign features are available for editing.
5. Click **Save** to activate your changes and make them permanent.

Editable features of a published or retired campaign

To prevent you from modifications that could falsify the results of a campaign or set the campaign in an inconsistent state, not all the features of a campaign are available for editing when the campaign has already been published or retired. Beware that modifying an editable feature of a campaign may also have

consequences.

List of editable features

Find below the list of editable features with the description of the consequences if you change them.

Campaign name

Changing the name of a campaign invalidates existing scores and NXQL queries linked to the campaign.

Displayable fields

SENDER (Name, Title, and Picture), NOTIFICATION (Subject and Footer), Question text, Choice text, Greetings, Header, Footer, Final message, Translations.

Make sure that the changes in these fields do not modify the intended original meaning. Use preferently to correct typos or misunderstandings.

Question name and Choice name

The labels in the Portal are automatically updated to agree with the new names. However, it breaks existing scores and NXQL queries related to the question.

RECIPIENTS

A change in the investigation on users associated to the campaign immediately redefines its target audience.

Reminder frequency

Immediately changes the rate at which users receive notifications.

Choice value (of opinion scale questions)

The answers received up to the moment of the change are unaffected, meaning that the Portal will not recompute them with the new values.

Choice order

Moving choices up or down in the list causes them to be displayed in a different order.

List of non-editable features

The following features are not editable because their modification would certainly lead to inconsistencies in the campaign.

- The type of a question.
- The flow of questions.
- Add or remove questions.
- Add or remove choices of questions.

Types of questions

Overview

There are three types of questions that you can use in your campaigns: single answer, multiple answer, and opinion scale. Learn here which type of question is most suitable for each situation.

Single answer

Use a *single answer* question when the end-user must select only one option among all possible choices. This situation happens when choices exclude each other. A typical example are *yes* or *no* questions, such as *Do you like our new web portal?* Possible answers might be:

- *Yes.*
- *No.*
- *I do not know about our new web portal.*

Note that the answers include an option for the user to indicate lack of knowledge about the subject matter being asked. This type of option gives you valuable information about the awareness of your end-users regarding new features or changes in the IT environment.

Each single option is displayed as a clickable button in the notification window received by the end-users. If many choices are available, buttons may span multiple lines across the notification window.

Multiple answer

Use a *multiple answer* question when the end-user may select one or more of the possible choices simultaneously. Contrary to a single choice question, in a multiple answer question the choices are not mutually exclusive. For example, to the question *What do you like about our new web portal?*, the possible answers might be:

- It loads fast.
- It is visually appealing.
- I quickly find what I need.
- I have no opinion.

Each answer is preceded by a check box that allows end-users to tick the desired options. Again, if many choices are available, the possible answers may span multiple lines across the notification window.

End-users must at least select one option to be able to go on with the survey. An **OK** button appears after the end-user makes its first selection. Therefore, it is convenient to include a default option that end-users can tick if none of the others apply (for instance, the last option in our example).

Opinion scale

Use an *opinion scale* question when the end-users must evaluate a product or service and give it a mark. A numerical value is associated to each possible answer. By default, the value attributed to the first specified answer is 1 and it increases automatically with each new added answer. However, it is possible to assign different values to each answer. Use these values later for computing scores and metrics related to end-user feedback.

You can have at most one answer without any value attributed to it. It is usually reserved for users who do not have an opinion about the asked subject matter. To avoid biasing, answers without a value assigned do not take part in the computation of averages for scores and metrics related to end-user feedback.

For instance, to the question *How easy to use is the new web portal?*, the possible answers might be:

- Not easy at all
- Not that easy
- Easy
- Very easy
- I have not tried the new portal (choice without attributed value)

Each choice of an opinion scale question is displayed on top of a clickable radio button. Since the order is important, choices are displayed along one line only in the notification window. It is thus recommended to limit the choices up to a maximum of seven per question for the notification window to fit in the screen.

Controlling the flow of questions

Overview

The order in which you define the questions of a campaign determine the natural order in which the user answers them. However, it does not always make sense to ask all the defined questions. This is usually the case of questions that depend on how the user answered a previous question.

For instance, consider a campaign that holds the two following questions:

1. A single answer question *Do you need a cloud storage service for sharing your work?*, with possible outcomes *Yes* or *No*.
2. A multiple answer question *Select the cloud storage services that you prefer*, with possible choices *Dropbox*, *Google Drive*, *Microsoft OneDrive*, etc.

It is obvious that users who answered *No* to the first question will find the second question superfluous or even annoying. In this particular example, you can avoid the situation simply by removing the first question and adding a default choice to the second question: *I do not need a cloud storage service*. In more complex

scenarios, however, you may really want to jump over a question or two when the situation requires it.

Specifying your question flow

There are two ways to alter the order of the questions in a campaign and both are based on the same mechanism: the *step to* option. A *step to* declaration indicates the campaign that the normal flow of questions is modified. Instead of asking the user the next question in the defined order, the campaign shall jump to the question specified in the *step to* option. It is also possible to specify a direct jump to the end of the campaign, if there are no pertinent questions left.

The two ways of altering the order of the questions in a campaign refer to the places where you can apply the *step to* option. These are the following:

- A particular choice of either a single answer or opinion scale question.
- The end of the question.

By applying a *step to* option to individual choices, you create branches in the question flow. Note however that you cannot apply a *step to* option to the choices of multiple answer questions. Indeed, it would be difficult for a campaign to decide which question to jump to if a user selects two or more choices of a multiple answer question and each choice specifies a different *step to* option.

Step to options at the end of a question are generally used to merge the different branches generated by individual choices. For single answer and opinion scale questions, the *step to* options of individual choices take precedence over the *step to* option declared at the end of the question. This is explicitly stated in the Finder with the message **If not specified above, step to** when declaring the *step to* option at the end of a single answer or opinion scale question. For multiple answer questions, the message just reads **Step to**, since you cannot have individual *step to* options for each choice.

Thus, when creating a campaign, define your questions in the order that you find the most natural. Then, for each choice that should create a path which deviates from the natural order, set a *step to* option. Finally, merge different paths back into one by jumping from the end of questions in distinct paths to a common question (or to the end of the campaign).

You will end up with a workflow that, if depicted, should resemble the following

diagram.

Note that workflows can help you reduce the number of questions that you ask in a campaign. In the example workflow above, for instance, users will answer at most three questions out of the five questions defined.

Displaying workflows in devices with outdated Collectors

For end-users with versions of the Collector prior to 6.8 installed in their devices, campaigns may be displayed only if they have no *step to* options defined or not be displayed at all. See the article on the display compatibility of the Collector with campaigns for specific support information by a particular version of the Collector.

Related tasks

- Creating a campaign

Related references

- Campaign display compatibility

Translating a campaign

Overview

To get actionable feedback from your end-users when you publish a campaign, it is essential that the end-users can clearly understand the questions being asked, as well as all the possible answers. The level of clarity and precision required is often only achieved if the campaign is written in a language that the end-user speaks fluently.

In multilingual environments, translate your campaigns to the different languages spoken by your community of end-users to communicate with them in their own mother tongue whenever possible. When an appropriate translation is available, end-users receive the content of the campaign in their language of choice, according to the language settings of their operating system.

Providing translations

When you create a new campaign, you write the questions and possible answers in a particular language. This language is called the *base language* of the campaign.

To provide a new translation for a campaign, find the **Translations** section at the bottom of the campaign editor:

1. Click the link that reads **Click here to add a new translation**.
2. Depending on other translations being available:
 - ◆ If this is the first translation of the campaign, specify the base language:
 1. In the dialog, select the language in which you originally wrote the campaign from the **Base language** list.
 2. Click **OK**.
 - ◆ If there is at least another translation available, the base language was already defined. No dialog to choose the base language shows up.
3. In the dialog, select the target language of the translation from the **Additional language** list.
4. Click **OK**. The list of available translations gets updated under **Translations**.
5. Optional: Repeat from step one to add as many languages as necessary.

Translating campaign text

Once you have created a new translation, it is displayed in the list in an error state because the translation is *incomplete*. Indeed, you need to provide actual translations for all text items in the campaign that are visible to the end-user; that is:

- The title of the sender.
- The header and the footer of the campaign notification.
- The displayed text of questions and answers.
- The greetings, header, and footer of the first question.
- The final message after the last question.

For translating the text of a campaign, Nexthink has adopted the *PO* format. The translation cycle consists thus in exporting the untranslated text of the campaign to a PO file, provide the translation of all text items to the target language with the help of an editor, and import the edited PO file back into the Finder.

To translate the text items in a campaign:

1. Under **Translations**, find the entry in the list of available translations that corresponds to the target language.
2. Click **Export**, the first option in the fourth column, to export the translation to a PO file.
 1. Select a destination folder for exporting the translation.
 2. Click **OK**. The Finder creates a folder with the name of the campaign in the chosen destination folder and a PO file inside it. The name of the PO file has the format
`Campaign_name.language_code.po`.
3. Edit the PO file. You can use any of the specialized PO editors available in the market (recommended) or a plain text editor that supports UTF-8 character encoding.
 1. Provide a translation for every text item in the PO file.
 2. Save your changes.
4. Back to the **Translations** section in the Finder, click **Import**, the second option in the fourth column of the corresponding entry in the list of available translations, to get the translated version of the campaign.
 1. Select the PO file with the translated text.
 2. Click **OK**. The state of the translation in the Finder should be updated to *complete* and show no errors.
 - ◊ If the translation still shows as *incomplete*, go back to the previous main step and ensure that you provide translations for all the text items in the PO file.
 - ◊ If the campaign changed while you were editing the PO file, go back to step one.
5. Optional: Click **Preview**, the option in the sixth column of an entry in the list of available translations, to inspect how the translated campaign will display in the devices of the end-users.
6. Optional: Click **Remove**, the third option in the fourth column of an entry in the list of available translations, to delete it.
7. Click **Save** to store your changes.

Before a campaign is published, if translatable text items are created or modified after you provided a translation for the campaign, the translation enters the *incomplete* error state again. Re-export the translation and provide the appropriate wording in the target language for the missing or modified items.

Default language

When displaying a campaign to an end-user, the Collector tries to find an appropriate available translation that matches the language settings of the operating system in the device of the end-user. If the Collector is not able to find such a translation, it falls back to the *default language* of the campaign.

By default, the default language of a campaign is the base language. To change the default language, open the campaign editor:

1. Under **Translations**, find the entry in the list of available translations that corresponds to the language that you want to set as default.
2. Click **Set as default**, the option in the fifth column, to make the language of the translation the default language of the campaign.
3. Click **Save** to store your changes.

If the translation selected to be the default is later removed, the editor resorts to the base language as the default language of the campaign.

Language of buttons in campaign notifications

The messages displayed on the buttons in the notifications of campaigns to accept or reject the campaign, ask for a reminder, etc. are also translated according to the language settings of the operating system running on the device of the end-user.

For these messages, however, the set of languages available is predefined by Nexthink and it is not possible to manually provide a translation. This is the list of currently supported languages:

• Afrikaans	• Finnish	• Korean	• Swedish
• Arabic	• French	• Norwegian	• Thai
• Chinese	• German	• Persian	• Turkish
• Danish	• Hindi	• Portuguese	• Vietnamese?
• Dutch	• Italian	• Russian	
• English	• Japanese	• Spanish	

Displaying translations in devices with outdated Collectors

For end-users with versions of the Collector prior to 6.9 installed in their devices, campaigns may be displayed only in their base language or not be displayed at all. See the article on the display compatibility of the Collector with campaigns for

specific support information by a particular version of the Collector.

Related tasks

- Creating a campaign

Related references

- Campaign display compatibility

Triggering a campaign manually

Overview

To manually trigger a campaign from the Finder, the following conditions must be met:

- The profile of the user must:
 - ◆ Allow editing of campaigns.
- The campaign must:
 - ◆ Target users manually.

To trigger a campaign from an external product, use the API of campaigns.

Applies to platforms:

From a list view

To manually trigger a campaign from the list of results of an investigation on users or the list view of a user category:

1. Select a user or a set of users from the list.
 - ◆ Drag the mouse cursor or use the **Ctrl** or **Shift** keys to select multiple users.
2. Right-click your selection.
3. Choose **Campaigns** from the menu. A list with the available campaigns shows up.
4. Click the name of the desired campaign.
 - ◆ A warning message shows up, displaying the number of users selected and asking for confirmation. Click **Yes** to proceed.

From the user view

In the user view of the Finder, manually trigger campaigns from:

- The user icon on the top-left corner of the user view.

Manually triggering a campaign from the user icon in the top-left corner of the device view is very similar to triggering a campaign from a list view:

1. Right-click the device icon.
2. Choose **Campaigns** from the menu. A list with the available campaigns shows up.
3. Click the name of the desired campaign.
 - ♦ A warning message shows up, stating that the campaign will be sent to 1 user and asking for confirmation. Click **Yes** to proceed.

From the Start page

When searching in the Start page in the Finder, a list of users may show up in the right-hand side of the window as a result of the search. You can directly send a campaign to any of the users that appear in the list of results:

1. Right-click the name of the user.
2. Choose **Campaigns** from the menu. A list with the available campaigns shows up.
3. Click the name of the desired campaign.
 - ♦ A warning message shows up, stating that the campaign will be sent to 1 user and asking for confirmation. Click **Yes** to proceed.

Related tasks

- Triggering campaigns via their API
- Graphically observing the activity of users and devices

Scrutinizing the results of a campaign

Overview

After you publish a campaign, the targeted end-users start answering the questions that constitute the campaign. Learn here how to examine in detail the answers that users provide to your campaigns, either by creating investigations

in the Finder or by building new dashboards in the Portal.

Quick campaign results

To quickly display the results of an ongoing campaign so far:

1. Log in to the Finder as a user with the right to create campaigns.
2. Select the **Campaigns** section on the left-hand side of the main window.
3. Right-click the name of a published campaign.
4. Select **Display results** from the menu.

You get a list of all the users along with their campaign-related fields. Refine the obtained results by adding or removing columns and conditions to this automatically generated investigation.

User fields related to campaigns

Once you create a campaign, a new set of custom fields related to the campaign are available in all objects of type user. These fields let you see the responses of every user to each one of the questions in the campaign, as well as additional information according to the status of the user in relation to the campaign. These fields are also available as conditions on user objects in investigations.

The names of all campaign-related fields start with the name of the campaign and have the following values:

Campaign name - **Last action time**

The last time that the user interacted with a campaign (dealt with a notification or answered any question).

Campaign name - **Number of days since last action**

How long since the user interacted with a campaign for the last time.

Campaign name - **Status**

How far the user has gone through the campaign.

Campaign name - **Question name**

There is one additional field per question defined in the campaign. The value of this field corresponds to the answer given by the user:

- ◇ For single answer questions, the value is the name of the answer selected by the user.
- ◇ For multiple answer questions, the value is a list with the names of all the answers selected by the user.
- ◇ For opinion scale questions, the value is the name of the answer selected by the user, followed by the numerical value attributed to the answer between parentheses.

Campaign name - Question name (optional text)

For multiple answer or opinion scale questions that include an optional free text field, the value is the text typed in by the user.

Status of a user

The status of a user in relation to a campaign depends on whether the user accepts to participate in it or not and, in the former case, on how far the user went into answering the questions of the campaign. The status of the user regarding the campaign changes:

Unknown

The campaign did not target the end-user or the end-user did not receive the notification yet, either because of problems communicating with the user's device or because the device is turned off.

Notified

The user received the notification from the campaign but ignored it.

Opened

The user clicked on the notification to answer the campaign but did not answer any question.

Partially answered

The user started to answer the questions of the campaign but did not go to the end of it.

Fully answered

The user answered all the questions of the campaign.

Postponed

After receiving the notification, the user decided not to immediately take the survey but be reminded about it in the future (the user clicked the button **Later**).

Declined

After receiving the notification, the user decided not to take the survey (the user clicked the button **Never**).

You can run investigations with conditions on users using the created campaign-related Custom Fields.

Metric computations related to campaigns

Campaigns also let you define metrics around them. Build dashboards in the Portal that display campaign-related metrics to analyze the responses of your end-users and measure the success of your campaigns.

Count metrics

Metrics that count the number of users may benefit from the campaign-related custom fields of users to group the results. When displaying the results of a count metric in the Portal, either in a table or in bar chart widget, you can group the users by their response to a particular question or by their status in connection with a campaign.

In this way, follow the progress of a campaign by creating a count metric that groups users by their status in relation to the campaign. Add the metric to a bar chart widget in a dashboard and you will get a breakdown with the number of users who fully answered the campaign, those who only answered a few questions, those who answered none, etc.

In addition to grouping results, campaign-related custom fields of users let you define ratios in count metrics. For instance, if you define a metric that gets the number of users who both fully answered the campaign and answered positively to a particular question and then compares it to the total number of users who fully answered the campaign, you get a count metric that computes the ratio of users who answered positively.

It is recommended that you always use the modifier **all users** and not **active users** when defining count metrics that relate to end-user feedback. The reason is that you usually want to take into account the responses of all users who answered the campaign, no matter if they were not active the last day.

Quantity metrics

Similarly to the generated custom fields for users that were mentioned above, creating a campaign also adds new options to create quantity metrics based on users. Since quantity metrics require a numerical value, only opinion scale questions make these options available. Thus, for every opinion scale question included in a campaign, there is a corresponding option to create a quantity metric for objects of type user. The label of the name has the following format:

- *Campaign name - Opinion scale question*

Quantity metrics related to opinion scale questions make easy to compute average response values in the Portal.

Rating devices and users with scores

Computing scores

Overview

A *score* applies either to devices or to users and it offers an evaluation of each object from one or several points of view. The value of a score depends thus on the status of the object with respect to a particular aspect or the combination of different aspects.

In this way, scores reduce the complexity of low level analytics to a single numerical value that has business significance. For instance, instead of dealing with multiple measures of network connectivity and their technicalities, gain instant insight into the connectivity status of a device with the help of a single score that combines all of those measures.

Scores are organized into tree structures of up to five levels. Scores at the lowest level of the tree are called *leaf* scores. A leaf score directly depends either on the computation of an aggregate or on the value of a device or user field. In turn, scores at higher levels of the hierarchy are called *composite* scores. A composite score depends on lower level scores, which may themselves be either composite or leaf scores.

Leaf scores

To compute a leaf score, provide one of the following inputs:

- **Field input:** The value of a field that belongs either to a device or to a

user object.

- **Computation input:** The value of an aggregate that results from an NXQL investigation on devices or users.

The final value of a leaf score is computed by applying a normalization function to the raw input value (field or computation) in order to get a number within a predefined range. For example, a score based on the number of system failures which ranges from 0 to 10 (the higher the better) might have a normalization function that maps 0 failures to value 10, more than 5 failures to value 0, and a number failures between 1 and 4 to a list of decreasing intermediate values.

Update of leaf scores

Leaf scores whose input is a field of user or device objects are updated every minute. For their part, leaf scores with a computation input are updated either once per day at a specific time or periodically every 15 minutes, 1 hour, or 6 hours, depending on their configuration.

Composite scores

A composite score is computed by combining the values of the scores at its immediately lower level. To combine the lower level scores, use one of the following operations:

- **Average:** Add the scores at the lower level and divide the result by their total number.
- **Weighted Average:** Similar to the average operation, but with factors (*weights*) that multiply each lower level score to reflect their importance.
- **Min:** Return the minimum value from the available lower level scores.
- **Max:** Return the maximum value from the available lower level scores.
- **Sum:** Add the lower level scores.
- **Multiply:** Multiply the lower level scores that make up the composite score. Useful for combining closely related scores, for example: the *total logon time* of a device is the *average logon time* multiplied by the *number of logons*.

Update of composite scores

Composite scores are updated when any of the scores on which the composite score depends is modified. Note thus that the modification of a leaf score at the lowest level of the hierarchy may trigger the recalculation of all scores above it.

Scope of a score

A score may apply to a particular group of devices or users, and not to all of them. This group is called the *scope* of the score. Define the scope of a score by filtering the targeted objects with a condition on a property. For instance, for a score to apply only to laptop computers, put a filter as the scope of the score that returns those devices whose type is set to **laptop**.

Creating a score

Overview

Scores are described in the form of XML files. To create a new score, either:

- Use the Nexthink Scores Creator tool from your web browser (recommended).
- Directly write the XML of the score with your favorite text editor or with a specialized XML editor to define a new score and import it into the solution.

Contact Nexthink Customer Success Services to help you build your own scores or fine-tune the available scores to suit your needs. Check the Nexthink Library for the latest published packs that include scores.

Score definitions

A score *definition* holds the detailed description of either one composite score or one leaf score in a single XML file. Note however that a composite score usually comprises several *subscores*. Therefore, in practice, XML files that hold the definition of a composite score contain in fact the definition of all the subscores on which the main composite score depends.

Procedure

To write a new score or update an existing score:

1. Open the Nexthink Scores Creator tool from your web browser:
 - ♦ If you want to write a new score, click the button **Start a new score** at the top of the window.
 - ♦ If you want to modify an existing score, click the button **Load score...** and choose the XML file to upload.

2. Optional: Click **Settings** to configure the Web API settings of a reachable Engine.
 1. **Engine IP**: Type in the address of the engine.
 2. **Port**: Type in the port number of the Web API in the Engine.
 3. **Version**: Select the product version to generate a compatible score definition.

The Engine settings let you test the scope and NXQL queries used in your scores thanks to the buttons:

 - **PREVIEW SCOPE**
 - **PREVIEW QUERY.**
3. Fill in the forms for the main score and the subscores.
4. Click **Save score...** to save your score as an XML file.
 - ♦ If any error is detected in your score, a warning message shows up before saving the file.

Alternatively, if you are fluent in XML or want to make a quick local change to one of your scores from your computer:

1. Open your favorite text editor or specialized XML editor.
2. Write the definition of the score in an XML file, according to the XML schema of scores.
3. Save the XML file.

Once you have the XML file with the contents of your score:

1. Open the Finder as a user with the right to edit scores (the profile of the user must allow system configuration).
2. Select the **Scores** section on the left-hand side panel of the main window.
3. Right-click the header of the **Scores** section or the empty area below it.
4. Select **Import > Scores from file...** from the menu.
5. In the dialog, choose the XML file that holds the definition of your score.
6. Click **Open** to import the score.

Limit on the number of scores

There are two limiting factors that put a constraint in the maximum number of scores that can be defined:

- **Storage**: 500 scores in total (either composite or leaf scores, including subscores in the count).
- **Computation power**: 50 leaf scores of the **Computation input** type. Be careful if you use complex NXQL queries with **except**, **union**, or **intersect**

in leaf scores, since they might sensibly increase the computation power needed.

If you need to go over these limits, contact Nexthink Customer Success Services for an evaluation of your setup and a configuration customization.

The operations described in this article should only be performed by a Nexthink Engineer or a Nexthink Certified Partner.

If you need help or assistance, please contact your Nexthink Certified Partner.

Related references

- [Nexthink Scores Creator \(external\)](#)
- [Score XML Reference](#)
- [Nexthink Library](#)

Checking and comparing ratings

Overview

Once you have a set of scores defined, see the actual ratings of devices or users by either:

- Opening respectively the device or user view of a particular device or user.
- Adding a score column to the displayed results of an investigation on devices or users.

Note however that some scores (or subscores) may be configured to be displayed exclusively in the Portal and not in the Finder. These scores are visible neither from the device view nor from the user view, and they are not eligible for being displayed in the results of an investigation. Some other scores may be configured to be visible in the Finder only in the context of defining a quantity metric.

To learn about the different ways to open the device and user views in the Finder, see the article on observing the activity of users and devices.

To know how to add columns to the list of results of an investigation, see how to edit the options of an investigation or modify its displayed results.

Checking scores in the device and user views

In the device and user view of the Finder, up to 10 score tabs are placed to the right of the **Timeline** and **Properties** tabs, depending on the number of scores that you have defined. Each tab displays the name of the score and its current value.

In addition:

- If the score defines thresholds, a colored line at the top of the tab indicates the status of the score according to its current value.
- If the score defines no thresholds, the scores has no status and it does not show any special color in the tab.
- If the Cross-Engine features are enabled in the Finder, a user score is averaged over the individual scores of the user on each Engine where the user was active.

The Finder arranges score tabs by the names of the scores in alphabetical order. If more than 10 scores are available for devices or users, the Finder ignores those in excess. Additionally, the Finder ignores those scores whose current value is null.

Breakdown of a score

Click on a score tab to see the list of leaf scores that make up the main score. If the main score has composite subscores, the leaf scores are organized into groups under the names of the composite subscores to which they belong. The names of the composite subscores act therefore as section headers.

Leaf scores let you reason about the causes for a main score to have a particular value. For instance, in the case of main score having a mediocre value, its decomposition into leaf scores may let you find out whether the final value is caused by a generally medium value of leaf subscores or if there are just a couple of leaf scores with very low values that take the overall value down. For a score that measures *Device performance*, for example, that would be the difference between having to upgrade the whole device (if all leaf scores perform poorly), or upgrading just the hard disk (if only the leaf scores related to storage are really low).

For each leaf score, the Finder displays:

- The value of the leaf score with a background color indicating its status (if any).

- The name of the leaf score.
- The actual value of the field or computation on which the score is based.

Hover the mouse cursor over the value or the name of the leaf score to get its description in a tooltip.

Status of a score

Scores that define thresholds have a status. The status of a score is indicated by a color which depends on its value. There are three possible statuses:

- Green, which means *good*.
- Yellow, which means *bad*.
- Red, which means *critical*.

The thresholds that determine the ranges of values for each status, as well as their labels, are set in the definition of the score and hence they are configurable.

The status of a score is indicated by the color of the line at the top of its corresponding tab. In its turn, the status of a leaf subscore is indicated by the background color used to display its value. Scores and subscores with no status show no special color.

Viewing score values in investigations

When selecting the display columns of either devices or users, find the **Scores** section at the bottom of the selection dialog. For each defined score, select either the main score or any subscore that you wish to see in the list of results of the investigation.

Whereas the device and user views let you see all the scores of a particular device or user, investigations let you compare the value of a particular score among many devices or users at the same time.

When examining the scores of users on the **List (all entities)** view (i.e. in a Cross-Engine context), the score of a user is averaged over the individual scores of the user on each Engine where the user was active. Note that a same device that is active on several Engines is considered a different device by each Engine. Thus, on the **List (all entities)** view, a single device has a different entry per Engine and displays the local score on each entry.

Scores with a *null* value are represented with a dash sign (-) in the lists of results of an investigation.

Related tasks

- Graphically Observing The Activity Of Users And Devices
- Editing the options of an investigation
- Navigating the results of an investigation
- Enabling Cross-Engine Finder features

Computing potential savings

Overview

In combination with metrics, scores help you find out potential money and time savings within your IT infrastructure through the calculation of costs related to software and hardware inventory or through the estimation of the time spent on certain activities.

Let us illustrate how to discover possible time and money savings with a couple of examples. If you are interested in these or similar examples, please contact Nexthink Customer Success Services.

Time savings - Startup duration

The startup time of a device is unproductive as well as annoying for the end-user, specially when the waiting time is too long. In Nexthink, the Collector takes two measures that add up to the total startup time of a device: the boot time and the logon time. A score to find out potential time savings in the startup time of a device may focus either on the boot time, on the logon time, or on both.

For instance, to compute the total boot time of a device during one day, create a composite score that multiplies the boot duration of the device by the number of times that it was booted that day. That is, create two leaf scores based on the aggregates **Average boot duration** and **Number of system boots** and compose them with the *multiply* operation to get a composite score that yields the total boot time. Optionally, do the same for the logon time and compose the boot and logon time scores with the *sum* operation to get a single composite score that holds the total startup time. To combine the score with metrics, it must be computed once every day at midnight in order for the metric to get the value of the score for the full day.

Now that the score provides the startup time of each device, let us use metrics to get the sum over all devices. To that end, create a quantity metric that computes

daily the startup time score of **active devices** and aggregate it by the **sum over all devices and the whole timeframe**.

When displayed in the Portal, the metric may be used to analyze the startup time of devices and help you determine the reason why some devices have longer startup times. You may use this information to reduce time wasted in the startup process by taking appropriate action over those devices whose startup time is the longest: update the operating system, the system memory, replace old models, etc.

Cost savings - Underutilized software

Software that is installed but never or seldom used incurs in unjustified licensing costs. To help you estimate the cost of licensed software that is not used, start by creating a score on devices. Use the score to individually assign cost to each device that has a particular software application installed, but that did not execute it during -for instance- the last month. Compute the score every day at midnight to combine it with a metric.

To calculate the total cost of underutilizing a particular software application, create a quantity metric that computes the cost score of **all devices** daily (you want to include those devices that were not necessarily active the last day) and aggregate the cost by the **sum over all devices**.

Display this metric in the Portal to know how much money you can potentially save by removing unused installed programs. As a complement to scores for cost savings, remember that you can create software metering metrics for assessing license use as well.

Related tasks

- Creating a metric
- Assessing license use

Related references

- Portal aggregation and grouping

Score XML Reference

Overview

Scores are defined in XML documents. This article details the format of XML files that hold the definition of a score.

The root element ScoreDef

As any other XML file, a score definition must begin with the XML header. Use **utf-8** encoding:

```
<?xml version="1.0" encoding="utf-8"?>
```

The XML header is followed by the root element. The name of the root element for XML documents that hold a score definition is **ScoreDef**:

```
<ScoreDef UID="[<unique identifier>]" Name="[<name of the score>]"
  Object="[device|user]" InObjectView="[true|false]"
  Status="[enabled|disabled]" SyntaxVersion="1"
  DataModelVersion="12">
  <Platforms ...
  <ScopeQuery ...
  <ComputationSchedule ...
  <Thresholds ...
  <CompositeScore|LeafScore ...
</ScoreDef>
```

Along with the **ScoreDef** element, find the following attributes:

UID

A universally unique identifier for the score. It is a 128-bit integer number expressed in hexadecimal format with separators and available tools exist to generate them with low collision probability . Example:

f54972e7-0f9c-46ce-8931-bbe31b06f7b4.

Name

The name of the score as shown in the Finder.

Object

The type of object to which the score applies. It can either be **user** or **device**.

InObjectView

Whether the score should appear inside the User or Device view in the Finder or not. Note that the User and Device views can display a maximum of five scores each one. Possible values are **true** or **false**.

Status

Whether the score should be calculated or not. Possible values are **enabled** or **disabled**.

SyntaxVersion

The version of the schema for the XML file. Currently fixed to **1**.

DataModelVersion

The version of the Nextthink Data Model on which the score relies for its queries. Currently fixed to **12**.

After the attributes, find the following elements inside **ScoreDef**:

Platforms

The platforms to which the score applies. Enumerate each supported platform within a **Platform** element. Possible values for platform are **windows**, **mac_os**, or **mobile**. Example:

```
<Platforms><Platform>windows</Platform></Platforms>
```

ScopeQuery

Define the scope of the query to retrieve users or devices by filtering the objects. Write a *where* clause in NXQL inside the **Filtering** element. For example, to filter out servers from a score on devices:

```
<ScopeQuery><Filtering>(where device (ne device_type (enum server)))</Filtering></ScopeQuery>
```

ComputationSchedule

The moment to compute all the leaf scores that have a computation input. It can be either an hour of the day, for computing the score every day at that hour, or a period, for computing the score more often. It applies only to those leaf scores whose input is a computation and not a field, as field-based scores are updated every minute. The main score is updated whenever a leaf score changes, be it of the computation input or of the field input type. To specify an hour of the day, insert the element **At** with a value between 0 and 23. To specify a period, insert the element **Every** with values 15min, 1h, or 6h. Example:

```
<ComputationSchedule><At>2</At></ComputationSchedule>
```

Thresholds

Optional: The limit values of the score that determine its status. Define up to three thresholds that the score has to exceed to be in a particular status. The Finder displays the score with a particular color depending on its status: red, yellow, or green. Specify the thresholds from green to red or from red to green (yellow is optional) always with the values in ascending order. Each **Threshold** element holds the **Color** attribute, as well as a **Keyword** element that defines the limit value of the threshold in its **From** attribute and the name of the corresponding status in the **Label** attribute. For a score to define no thresholds, declare exactly one threshold with the color *none*. Example: `<Threshold color="green"><Keyword From="9" Label="good"/></Threshold>`

Once the general elements of the score definition are laid out, add a single composite or leaf score element to the definition. This last element describes how to compute the score from the values stored in the Nexthink database. If you add a **LeafScore** element to your **ScoreDef**, the score will depend on one computation or on the value of one field only. On the other hand, if you add a **CompositeScore** element to the **ScoreDef**, a combination of several scores is used to compute the main score.

As a matter of fact, a composite score is itself composed of other composite or leaf scores, forming a tree of up to five levels. Ultimately, all the nodes at the lowest level of the tree must be leaf scores.

Composite and leaf scores

CompositeScore and **LeafScore** are the elements in the score definition that are used to compute individual scores.

The following attributes are common to the **CompositeScore** and **LeafScore** elements:

UID

A universally unique identifier for the score (similar to that of the **ScoreDef** element).

Name

The name of the score.

Description

A textual description of the score.

Visibility

Optional: Whether the score should be visible in the Finder, nowhere, or only in quantity metrics in the Portal. Possible values are **visible**, **hidden**,

and **visible only in quantity metrics**. By default, a score is **visible** everywhere.

Weight

Optional: A floating point number with up to two decimal places that acts as multiplier of the computed score when its immediate composite parent score performs a *weighted average* operation (see how a composite score is computed below).

In addition to the common attributes, both the **CompositeScore** and the **LeafScore** elements optionally include a **Document** element. The **Document** element holds the documentation of the score that is displayed in the corresponding score tab of either the device view or the user view, depending on the object to which the score applies. Besides the detailed description of the score, the documentation of a score may also contain:

- Links to external HTTP resources (e.g. knowledge base, external documentation, etc).
- Links to remote actions which can be manually triggered, to take appropriate action on devices that display a poor score.

See how to document scores for the complete reference of the **Document** element.

Computing a composite score

A composite score results from the combination of its direct child scores by means of an *operation*. Each composite scores declares thus an **Operation** element that specifies how to combine its child scores. After the operation, find the list of child scores that contribute to the composite score:

```
<CompositeScore UID="[<unique identifier>]" Name="[<name of the
score>]"
  Description="[<score description>]" Visibility="[visible|hidden]"
  Weight="[<float>]">
  <Operation>[Average|Min|Max|Sum|WeightedAverage|Multiply]</Operation>

  <LeafScore ...
  <CompositeScore ...
  <LeafScore ...
  ...

  <Document ...
</CompositeScore>
```

Find below the list of possible values for the **Operation** element:

Average

Compute the arithmetic mean of the direct child scores.

Min

Get the minimum direct child score.

Max

Get the maximum direct child score.

Sum

Compute the addition of the direct child scores.

WeightedAverage

Compute the arithmetic mean of the direct child scores after multiplying each child score by the quantity specified in its **Weight** attribute.

Multiply

Compute the multiplication of the direct child scores.

The list of child scores is limited by the maximum number of scores that you can simultaneously enable. Remember however that this limit includes the nested scores and that the maximum nesting of scores is 5 levels.

Computing a leaf score

A leaf score is the result of applying a normalization procedure to an input value coming from the Nextthink database. The input is either the value of a field that belongs to the user or the device objects or the result of a computation on devices or users expressed in the NXQL language.

Thus, a **LeafScore** is composed of an **Input** element and a **Normalization** element:

```
<LeafScore UID="[<unique identifier>]" Name="[<name of the score>]"
  Description="[<score description>]" Visibility="[visible|hidden]"
  Weight="[float]">
  <Input ...
  <Normalization ...

  <Document ...
</LeafScore>
```

In its turn, the **Input** element encloses either a **Field** or a **Computation** element.

For an input of the field type, specify the name of the field to retrieve inside the **Name** attribute of the **Field** element. Type in the name of the field as stated in the NXQL data model. For example, to retrieve the WMI status of a device, write:

```
<Input><Field Name="wmi_status" /></Input>
```

An input of the computation type is slightly more complex to define. First give a name, a description, and a unique identifier to the computation as attributes of the **Computation** element. Then write a NXQL query to extract information from the Nexthink database inside a child **Query** element. As attributes of the **Query** element, set the **Output** to be the aggregate or field in the query that should be passed as input to the leaf score and the **DefaultOutputValue** to the value returned when the previously specified output is undefined. The structure of a computation input is thus as follows:

```
<Input>
  <Computation Name="[<name>]" Description="[<desc>]"
    UID="[<identifier>]">
    <Query Output="[<name_of_aggregate>|<name_of_field>]"
      DefaultOutputValue="[<value>]">

      (select (id ...) (from [user|device] ...
      </Query>
    </Computation>
  </Input>
```

Note that the embedded NXQL query must select the id of the objects retrieved (users or devices), as well as the field declared as output, if any. If an aggregate is declared as output instead, include the calculation of the aggregate in the query within a **compute** clause. For example, to compute the ratio of successful HTTP requests of a device during the last week for using it as the input to a leaf score, write the following **Query** element:

```
<Query Output="successful_http_requests_ratio"
DefaultOutputValue="NULL">
  (select (id) (from device
    (with web_request
      (compute successful_http_requests_ratio)
      (between now-7d now))))
</Query>
```

Note the use of the keyword **NULL** as the default output value in the previous query. Whenever it does not make sense to return a value for the score if the

underlying field or aggregate is undefined, write **NULL** as the default output value. In the previous example query, the ratio of HTTP requests is undefined if the device made no HTTP requests. Rather than force an artificial value for the score, it is preferable to return no value.

A leaf score with no value that is part of a composite score is discarded for the computation of the composite score. If a composite score or a main score cannot be computed because of lack of underlying values, they have no value as a result. The Finder displays a dash sign (-) for scores with an undefined value.

In the case that the input to a leaf score actually gets a proper value, transform it to give it business significance by means of a normalization function. The purpose of scores is indeed to make sense out of the detailed measures in the Nexthink database. Normalize quantities, ratios, enumerations, or even strings to a numerical range, usually the range from 0 to 10, that makes it easier for you to understand the status of a user or device with respect to the measured input.

The **Normalization** element defines thus how to map the input values to the desired range. Depending on the type of the input, apply a different normalization function:

Type of input	Available normalizations
Numeric type	<ul style="list-style-type: none">• Step function• Piecewise linear function
Enumerated type	Direct mapping
String type	Direct mapping (with optional placeholders)

To normalize numerical input values by means of a step function, define a list of **Ranges** with input values in ascending order (attributed scores do not need to follow any particular order), as in the following example:

```
<Normalization>
  <Ranges>
    <Range>
      <From Value="0" Score="0" />
    </Range>
    <Range>
      <From Value="0.6" Score="5" />
    </Range>
    <Range>
      <From Value="0.8" Score="10" />
    </Range>
  </Ranges>
```

</Normalization>

The example defines a step function for what it looks like a ratio input that returns:

- A score of 0 for an input value between 0% and 60% (not included).
- A score of 5 for an input value between 60% and 80% (not included).
- A score of 10 for an input value equal to or greater than 80%.

As mentioned previously, scores do not necessarily grow with the input value. For instance, consider this alternative normalization where the score drops to zero for input values higher than 80%:

```
<Normalization>
  <Ranges>
    <Range>
      <From Value="0" Score="5" />
    </Range>
    <Range>
      <From Value="0.6" Score="10" />
    </Range>
    <Range>
      <From Value="0.8" Score="0" />
    </Range>
  </Ranges>
</Normalization>
```

The resulting step function returns:

- A score of 5 for an input value between 0% and 60% (not included).
- A score of 10 for an input value between 60% and 80% (not included).
- A score of 0 for an input value equal to or greater than 80%.

To normalize numerical input values by means of a piecewise linear function, define a list of **Ranges** similar to those found in the step function, but this time

with an additional **To** element that delimits a linear interval and specifies the final value of the score within the interval:

```
<Normalization>
  <Ranges>
    <Range>
      <From Value="0" Score="0" />
      <To Value="0.6" Score="2" />
    </Range>
    <Range>
      <From Value="0.6" Score="2" />
      <To Value="0.8" Score="9" />
    </Range>
    <Range>
      <From Value="0.8" Score="9" />
      <To Value="1.0" Score="10" />
    </Range>
    <Range>
      <From Value="1.0" Score="10" />
    </Range>
  </Ranges>
</Normalization>
```

The example defines a piecewise linear function that maps our ratio input to a score that evolves linearly within each interval:

- For an input value between 0% and 60%, the score ranges from 0 to 2.
- For an input value between 60% and 90%, the score ranges from 2 to 9.
- For an input value between 90% and 100%, the score ranges from 9 to 10.
- For an input value equal to or greater than 100%, the score is 10.

As shown in the figure, linear interpolation is performed within each interval. For instance, an input value of 30% would receive a score of 1.

Note that the **Value** and **Score** attributes of a **To** element must be equal to those of the **From** element defined in the next range for the piecewise linear function to be continuous. The last range does not define a **To** element to avoid imposing a limit to the input value (in the case that our hypothetical ratio input could be higher than 100%).

For inputs of an enumerated type, the normalization function must directly map each possible input to a score, as in the following example:

```
<Normalization>
  <Enums>
    <Enum Value="ok" Score="10" />
    <Enum Value="failure" Score="0" />
  </Enums>
</Normalization>
```

Similarly, for inputs of the string type, the normalization function must map the possible inputs to a score. An inputs matches the value specified when it contains the string. You can use the wildcards:

- *, as placeholder for 0 or more characters.
- ?, as placeholder for a single character.

For example, this is a possible normalization function for an input that returns the Windows License Key of a device:

```
<Normalization>
  <Strings>
    <String Value="Windows is not activated" Score="0" />
    <String Value="?????-?????-?????-?????-?????" Score="10" />
  </Strings>
</Normalization>
```

In the case of Enums and Strings, beware that the values are **case sensitive**. Therefore, ensure that your definitions match the expected values as, for instance:

- Value="Ok" is not the same as Value="ok".

When displaying a leaf score in the device or user views, the Finder shows the payload of the score; that is, the input value that resulted in a particular score. To

improve the readability of the scores in the device and user views, when the input to a score has an obscure meaning, tag each normalization rule with a label that substitutes the raw input as the payload of the score. Use the **Label** attribute within the **Range**, **Enum**, or **String** elements. For example:

```
<Normalization>
  <Enums>
    <Enum Label="Service working" Value="ok" Score="10" />
    <Enum Label="Service not available" Value="failure" Score="0" />
  </Enums>
</Normalization>
```

Score XML validation

The structure and datatypes of XML files that describe scores are defined in an XML Schema Definition (XSD) file that Nextthink provides. To get the schema file (`score.xsd`):

1. Log in to the Finder as a user with the right to manage scores.
2. Select **Scores** in the left-hand side panel of the main window.
3. Right-click anywhere in the **Scores** space to bring up a context menu.
4. Choose **Export > Score schema to file...** from the menu.
5. Select the location where to store the file in the dialog and press **Save**.

To write your own scores, take the existing scores that you can find in the Library as example and validate the XML files that you create against this `score.xsd` schema file.

Related tasks

- Computing scores
- Creating a score
- Documenting scores

Documenting scores

Overview

Each composite and leaf score may include a **Document** element in its definition whose main purpose is to provide detailed documentation for the score when it is displayed in a score tab in either the user view or the device view.

The documentation of a score optionally holds:

- Links to other HTTP resources.
- Links to remote actions (must allow manual triggering).

This article assumes that you are acquainted with writing definitions of scores in XML format and specifically focuses in documenting your scores.

Document structure

The **Document** element is an optional XML element that you can find at the end of the definition of a composite or a leaf score.

The **Document** element has the following structure:

- A **Header**.
- A list of **Sections**.

```
<Document>
  <Header>[<Main documentation header>]</Header>

  <Sections>
    <Section>...</Section>
    ...
  </Sections>
</Document>
```

Both the header and the list of sections are optional, but the documentation of a score is only as useful as the content provided in these elements.

Document sections

In turn, each section in the list includes a title, followed by an optional sequence of descriptions, HTTP links, and links to remote actions. Typically, precede each link or group of links by a description that explains their purpose.

```
<Section>
  <Title>[Title for the section]</Title>
  <Description>[Text describing the contents]</Description>
  <RemoteAction ... />
  <HTTP ... />
  ...
</Section>
```

</Section>

Add as many descriptions, remote actions, and HTTP links to the section as you need, in the order that suits you best.

Adding HTTP links

To add an HTTP link, provide the destination URL in the **href** attribute and, optionally, a text to show in place of the URL:

```
<HTTP href="http(s)://url">[Text for the link]</HTTP>
```

Adding links to remote actions

To add a link to a remote action that allows manual triggering, provide the unique identifier (UID) of the remote action:

```
<RemoteAction UID="[identifier]" />
```

To get the UID of a remote action:

1. Log in to the Finder as a user with the permission to edit remote actions.
2. Locate the desired remote action in the **Remote actions** section of the left-hand side panel.
3. Right-click the remote action name.
4. Select **Export > Remote action to clipboard**.
5. Open your favorite XML or plain text editor.
6. Press **Ctrl+V** to paste the contents of the clipboard on the editor.
7. Find the **UID** attribute of the **Action** element in the XML of the remote action.

The Finder uses the name of the remote action as the text for the link and precedes the link with the characteristic rocket icon of remote actions.

It only makes sense to add remote actions to the documentation of scores that apply to devices and not to users.

Rendering example

The Finder displays the documentation elements of a score as follows:

Related references

- [Score XML Reference](#)

Remotely acting on devices

Scenarios for remote actions

Overview

The Nexthink Act module lets you execute remote actions on the devices of the end-users. Remote actions open up many possibilities for the prevention and remediation of issues encountered by the end-users as well as for gathering additional information from devices. A remote action is basically a PowerShell script that can accept parameters and optionally return outputs to be stored in the Engine.

Depending on the purpose of a remote action and how it is triggered, Nexthink has identified four main scenarios or use cases for remote actions:

- Assisted service
- Self-help
- Self-healing
- On-demand data

In addition, the Nexthink Act module exposes an API to programmatically launch remote actions, extending their availability to third-party products, such as self-service portals or ticketing systems.

Applies to platforms:

Assisted Service

In the assisted service scenario, the goal is to help support teams provide assistance to end-users.

When support agents with access to the Finder receive an incident or detect an issue on a device, they may manually trigger an appropriate remote action to fix the issue. Support agents may trigger remote actions from:

- The list of results of an investigation on devices or a category on devices.
 - ◆ Context menu on selected devices.
- The Device view.
 - ◆ Context menu on the selected device.
 - ◆ Links to remote actions in the documentation of score tabs.

Administrators control who can trigger the remote actions by assigning them to roles. Only those Finder users who have the right to edit remote actions or who have a remote action included in the roles of their profile can manually trigger the execution of a remote action.

Self-help

In the self-help scenario, the goal is to help end-users solve issues by themselves, without necessarily requiring external intervention from support.

The self-help scenario requires the activation of the Nexthink Engage module, as it relies on campaigns to notify the end-users about the issues found and ask them to take action.

Define remote actions for self-help to be automatically triggered when a known issue is detected on the device of the end-user. Once triggered, the remote action is connected to a campaign that guides the end-user to the resolution of the issue. The remote action should take into account the responses of the end-user and act accordingly.

Depending on whether the script payload of a remote action requires administrative privileges or not, run the remote action for self-help:

- In the context of the current interactive user, if the action requires no special privileges.
- In the context of the local system account, if the action requires administrative privileges.

Self-healing

In the self-healing scenario, the goal is to fully automate the remediation of specific issues in the devices of the end-users while avoiding human intervention.

Define remote actions for self-healing to be automatically triggered when a known issue is detected on a device and execute a remediation script that does not need the interaction of the end-user. Carefully consider the frequency of execution of this kind of remote actions.

On-demand data

In the on-demand data scenario, the goal is to get additional data from devices and incorporate it into Nexthink.

Remote actions are scheduled to be periodically executed on the devices of interest to gather the additional information that Nexthink does not collect by default. This additional information is then accessible through both the Finder and NXQL.

Related tasks

- Triggering remote actions via their API

Creating a remote action

Overview

Remote actions let you execute PowerShell scripts on the devices of the employees.

Applies to platforms:

Procedure

To create a new remote action:

1. Log in to the Finder as a user with the right to edit remote actions.
2. Select the **Remote actions** section on the left-hand side panel of the main window.
3. Right-click the header of the section or the empty area below it.
4. Select **Create new remote action** from the menu. The form to define the new remote action shows up in a new tab.
5. Type in an appropriate name for the action by replacing the temporary name **Untitled remote action n**.
6. Optional: Briefly describe the purpose of the remote action in the field labeled **Enter description here....**
7. In the **SCHEDULE** section, choose how to trigger the execution of the remote action:
8. Optional: Tick **Automatically run the remote action** to let the system trigger the remote action on a selected group of devices. Once the box is ticked, the Finder displays completely the sentence as **Automatically run the remote action on the following devices (evaluated every [period])** and opens up the controls to specify the devices on which to act:
 1. Choose the evaluation period of the investigation that retrieves the devices on which to act. There are three possible choices (see the maximum number of remote actions that you can define):

- ◇ **1 hour (default)**: for normal evaluation frequency.
 - ◇ **10 minutes**: for medium evaluation frequency.
 - ◇ **1 minute**: for fast evaluation frequency.
2. Specify the target devices by dragging an investigation on devices and dropping it into the area labeled **Drag & Drop device investigation here**.
 - ◇ The time frame of the investigation must be the full available period or during the last n days, hours, or minutes.
 - ◇ For fast evaluated investigations (frequency of 1 minute), the investigation must not compute aggregates and its time frame must be shorter than or equal to 1 hour.
 3. Optional: Right-click the dropped investigation and select **Run...** to see the devices that would be targeted. Keep in mind that the Finder displays the results of a single Engine only, so the actual number of target devices may be much bigger in a multi-Engine setup. Get more information at the moment of saving the remote action.
 4. Define the triggering period of the remote action by selecting a value and a unit from the lists inside the sentence **Trigger the remote action [value] [unit] on devices that are still returned by the investigation**. A device that executes a remote action and responds to it will not execute the same remote action again until the triggering period has elapsed, and only if the device is still a target (part of the results of the investigation). The triggering period must be bigger than or equal to the evaluation period.
 - ◇ For remote actions whose investigations to select target devices are based on activities or events, ensure that every evaluation of a device as a target triggers the remote action once and only once by choosing a triggering period that matches the time frame of the investigation.
 9. Optional: Tick **Allow manual triggering of the script on these devices** to let users with the right to edit remote actions or with the remote action included in their roles to manually trigger the execution of the remote action on particular devices. Once the box is ticked, the Finder displays a list for you to choose the devices on which the remote action may be manually triggered:
 - ◆ Choose **any** to let users manually trigger the remote action on any of all the available devices.
 - ◆ Choose **with keyword** to let users manually trigger the remote action on devices tagged with a particular keyword. If you choose this option, two additional lists show up to let you choose one category on devices and the exact keyword.

- ◆ Optional: Tick **The remote action can be triggered on multiple devices at once** so that users can manually run the remote action on a selection of multiple devices. By default, to prevent misuse of remote actions, users can manually run a remote action on a single device only.
 - ◇ As a rule of thumb, do not allow for potentially dangerous remote actions to be triggered manually on multiple devices. For example, do not allow multiple manual triggering of remote actions that require rebooting a device or that are network intensive.
- 10. In the **SCRIPT** section, click **Import...** to open the script that the remote action will execute.
 1. Select a PowerShell script from the dialog. The dedicated text area displays the script with syntax highlighting. Below the text area, the Finder shows a security message about the script signature:
 - ◇ **The script is unsigned**, for scripts that are not signed.
 - ◇ **The script is signed by: [author]**, for scripts that include a digital signature.
 1. Optional: Click the button **Show certificate** to open a standard dialog that shows the cer.
- 11. Optional: Click **Export...** to export the displayed script to an external file. This is useful if you want to modify a script that was provided as part of a content pack.
- 12. If the script includes any formal parameters, these are listed under the **Parameters** subsection.
 1. Type in the actual values to the right of each of the parameters listed. The actual values substitute the named parameters in the script during its execution.
 - **Warning:** Do not enclose the actual values for the parameters within any kind of quotation marks. Some scripts may require though the use of two consecutive double-quotes to indicate an empty value for a parameter ("").
- 13. If the script provides any outputs, these are listed under the **Outputs** subsection. For each output, the Finder displays its name in the script and its data type. Add the following:
 1. Type in a **Label** for the output. Use this label to refer to the output in the Finder, when defining investigations or metrics that relate to the remote action. In NXQL, use always the output name instead.
 2. Optional: Type in a brief **Description** of the output. The description is displayed as a column tooltip when the output is added to the results of an investigation.

14. In the **Advanced configuration** subsection, give additional details about the execution of the script:
 1. Choose the user of the device that the remote action must impersonate to run the script. Under **Run the script as**, choose:
 - ◇ **local system (default)**, to run the script as the system user.
 - ◇ **current interactive user**, to run the script in the context of the user who is logged in to the device at the point when the remote action is triggered.
 2. Set the maximum duration of execution of the script within the sentence **The script will time out after [n] seconds**. If the script is not finished by the specified time, it is stopped with a failure status. The default time out value is 120 seconds (2 minutes). Allowed values range from 10 to 604800 seconds.
15. Click **Save** to permanently store the remote action.
 - ◆ If you created an automatically triggered remote action, a dialog with the number of impacted devices across all Engines (calculated at the moment of saving the remote action) shows up:

1. Press **Yes, target N devices** to proceed.

Maximum number of remote actions

Using the procedure described above, create as many remote actions as needed provided that the total number of generated custom fields (the sum of the outputs specified in each action plus 6 default fields per action) does not exceed **500 fields**.

This is a hard limit that applies to both manually and automatically triggered remote actions. Additional limits apply to the maximum number of automatically triggered remote actions that you can enable simultaneously. See the details below.

Maximum number of enabled remote actions with automatic triggering

Enable up to **60 automatically triggered remote actions** in total.

- Out of these 60 remote actions, up to **30 remote actions** may have an evaluation period of less than 1 hour.
- Out of these latter 30 remote actions, up to **10 remote actions** may have an evaluation period of 1 minute.

For example, to fully utilize the maximum number of remote actions with the highest frequency possible, distribute the 60 available remote actions in the following way:

- 10 fast frequency remote actions (1 minute).
- 20 medium frequency remote actions (10 minutes).
- 30 normal frequency remote actions (1 hour).

Related tasks

- Creating an investigation

Related references

- Remote actions custom fields (NXQL Tutorial)

Executing remote actions

Automatic and manual triggering

The way of triggering a remote action, whether manually or automatically, usually depends on the particular usage scenario.

Manual triggering is basically used in *Assisted service* scenarios, where the a support team member executes a remote action to help an end-user solve an issue. The rest of scenarios typically schedule the execution of a remote action either periodically or in response to the detection of a specific issue:

Manual triggering	Automatic triggering
<ul style="list-style-type: none">• Assisted service	<ul style="list-style-type: none">• Self-help• Self-healing• On-demand data

Delayed first execution of automatically triggered actions

Once an automatically triggered remote action is saved, it is scheduled for execution. To avoid excess network and processing load, the Engine distributes the execution of remote actions so that not all actions run simultaneously on the target devices.

Thus, due to this distributed scheduling, the first execution of an automatically triggered remote action can be delayed up to its triggering period or, if the triggering period is longer than an hour, up to a maximum of one hour after it was created.

Controlling script execution on the device

When a remote action is triggered either automatically or manually on a device, the payload of the remote action is executed on the device only if the *script execution policy* allows it. Define the script execution policy of a device when installing the Collector on that device. Optionally, modify later the script execution policy of a device with the help of `Nxtcfg` (the Collector configuration tool).

For security reasons, the script execution policy controls the remote action scripts that are allowed to execute on a particular device.

- **Signed by a trusted publisher or by Nexthink (default):** the Collector runs on the device only those remote actions with a PowerShell script that is signed either by Nexthink or by a company whose certificate is listed in the Trusted Publishers certificate store.
- **Signed by a trusted publisher:** the Collector runs on the device only those remote actions with a PowerShell script that is signed by a company whose certificate is listed in the Trusted Publishers certificate store.
- **Disabled:** the Collector runs no remote action on the device.
- **Unrestricted:** the Collector runs any remote action on the device, regardless of the digital signature of its script. Useful for testing purposes, but not recommended in production.

Executing your own signed scripts

The default policy lets official remote actions from the Nexthink Library execute on the device without requiring any additional configuration. If you choose to create and sign your own scripts for remote actions, add the signing certificate to the Trusted Publishers certificate store in Windows.

In addition, if your certificate to sign scripts was generated by a private CA (that is, a CA whose root certificate is not already present in the Trusted Root Certification Authorities certificate store of Windows), add the root certificate of the CA to the Trusted Root Certification Authorities certificate store. Similarly, if you used an intermediate certificate to sign your scripts, include the full chain of intermediate certificates in the Intermediate Certification Authorities certificate store:

1. Log in to Windows as an administrator.
2. Press **Win+R** keys to open the Run dialog:
 1. Type in **certlm.msc**.
 2. Click **OK**.
3. Click **Yes** to allow the program to make changes to your device.
4. Right-click the name of the desired certification store (e.g. **Trusted Publishers**) in the left-hand side list.
 1. Select **All-tasks > Import...** from the context menu to start the Certificate Import Wizard.
5. Click **Next** to start the wizard.
6. Click **Browse** and select your certificate file.
7. Click **Next**.
8. In the dialog **Place all certificates in the following store**, click **Next** to accept the proposed certificate store.
9. Verify the certificate to be imported and click **Finish**.

It is recommended though not to import the certificates individually on every device, but to use a administration tool to deploy the certificates on all devices at the same time. Use Group Policy Objects (GPOs) of Active Directory to this end, for instance.

Results of remote actions

The outputs declared by a remote action are displayed in the Finder. Each output defines a label which is used by the Finder to name the column in the list of results of the remote action.

In addition to their defined outputs, remote actions return their status in a supplementary field called **Execution status details**. If difficulties were found during the execution of the remote action, the status field indicates it with an error message.

Related tasks

- Triggering a remote action manually
- Installing the Collector on Windows
- Nxtcfg - Collector Configuration Tool

Related references

- Errors in the execution of remote actions

Triggering a remote action manually

Overview

To manually trigger the execution of a remote action on a device from the Finder, the following conditions must be met:

- The profile of the user must either:
 - ◆ Allow editing of remote actions.
 - ◆ Include the remote action in one of its roles.
- The remote action must:
 - ◆ Allow manual triggering.
- The device must be:
 - ◆ A Windows device.
 - ◆ Part of the set of devices on which the remote action allows manual triggering.

To quickly identify the remote actions that allow manual triggering, the Finder adds a small hand on top of their characteristic rocket icon in the list of **Remote actions** in the accordion.

To trigger a remote action from an external product, use the API of remote actions.

Applies to platforms:

From a list view

To manually trigger the execution of a remote action from the list of results of an investigation on devices or the list view of a device category:

1. Select a device or a set of devices from the list.
 - ◆ Drag the mouse cursor or use the **Ctrl** or **Shift** keys to select multiple devices.

- ◆ By default, remote actions can be manually triggered on a single device only. Explicitly enable a remote action to be triggered on multiple devices when creating the remote action.
- 2. Right-click your selection.
- 3. Choose **Remote actions** from the menu. A list with the available remote actions shows up.
- 4. Click the name of the desired remote action.
 - ◆ If you selected multiple devices, a warning message shows up asking for confirmation. Click **Yes** to proceed.

From the device view

In the device view of the Finder, manually trigger remote actions from:

- The device icon on the top-left corner of the device view.
- The links to remote actions in the score tabs.

Manually triggering a remote action from the device icon in the top-left corner of the device view is very similar to triggering a remote action from a list view:

1. Right-click the device icon.
2. Choose **Remote actions** from the menu. A list with the available remote actions shows up.
3. Click the name of the desired remote action.

To manually trigger a remote action from a score tab in the device view, the score tab must include proper documentation with links to remote actions:

1. Click the score tab to select it.
2. Look for an appropriate link to a remote action (a link preceded by a rocket icon) in the documentation of the score.
 - ◆ If needed, click a subscore that provides its own documentation. Only documented subscores are clickable.
3. Click the link to the remote action.

From the Start page

When searching in the Start page in the Finder, a list of devices may show up on the right-hand side of the window as a result of the search. You can directly trigger a remote action on any of the devices that appear in the list of results:

1. Right-click the name of device.

2. Choose **Remote actions** from the menu. A list with the available remote actions shows up.
3. Click the name of the desired remote action.

Related tasks

- Triggering remote actions via their API
- Graphically observing the activity of users and devices
- Documenting scores

Writing scripts for remote actions

Overview

The payload of remote actions are PowerShell scripts that run on the devices of the end-users. PowerShell is a scripting language from Microsoft that is suited for task automation and configuration management and it is built on top of the .NET Framework in Windows devices. PowerShell is therefore ideal to get on-demand data from devices, perform self-healing tasks, or modifying the configuration of a device, which are typical use cases for remote actions.

Learn here the specifics of how to write PowerShell scripts for remote actions. This article assumes familiarity of the reader with PowerShell scripting.

Find more information about writing scripts for remote actions on Community:

- Nexthink Act Scripting Best Practices
- Nexthink Act Security Whitepaper
- Discussion forum on Nexthink Act

Applies to platforms:

Encoding

To write your own PowerShell scripts for remote actions, encode the files that hold the text of your scripts in UTF-8 with BOM. The BOM (Byte Order Mark) is a Unicode character that must be present at the beginning of the file and whose representation in UTF-8 documents is the following sequence of three bytes (in hexadecimal): EF BB BF.

Failing to provide the right encoding to your script files will result in the inability of remote actions to run on the devices of the end-users.

Signing your own scripts

For security reasons, Nexthink strongly recommends that you sign all the scripts that you write. Unsigned scripts are only suitable for testing in pre-production environments.

Ensure that you have deployed the appropriate signing certificates to the devices that run the Collector, so they can run your signed scripts as well.

Writing generic scripts

Scripts may be written in a generic way so they can be adapted to particular use cases. To make your script generic, declare formal parameters at the beginning of the PowerShell script. Provide actual values to the parameters in the Finder, when editing the remote action that holds the script.

Genericity is specially useful in the case of digitally signed scripts that require some customization. When a script is signed, any modification to its text content breaks the signature. By making a signed script generic, you allow its customization via parameters. With parameters, the text of the script remains unchanged and thus the digital signature remains valid. The values for the parameters that a remote action passes to the script determine its actual behavior.

For example, to make a script that reads from a registry key generic, declare a parameter in the script that holds the path to the key in the registry. In this way, several remote actions may use the same script to read from different registry keys by supplying a different path to the parameter in the script.

Declare parameters at the beginning of a script in the usual way for PowerShell scripts:

```
param(  
    [string]$filePath,  
    [string]$regPath  
)
```

The editor of remote actions recognizes the parameters of any imported PowerShell script and lists them in the **Parameters** section, below the script text. Provide actual values to the parameters in the text input boxes displayed to the right of each parameter name. Note that the actual values are always passed to the script as text: if the script declares parameters whose type is other than the **string** type, ensure that the values that you provide can be properly converted to

the type of their corresponding parameter.

Output variables

The execution of a script may generate some outputs that you want to store as on-demand data in the Engine. To that end, Nexthink provides a .NET assembly (`nxtremoteactions.dll`) that is installed in the device of the end-user at the same time as the Collector. The assembly includes a class called **Nxt** that provides the methods to write results to the Engine.

To use the **Nxt** class, add the following line at the beginning of a PowerShell script for remote actions:

```
Add-Type -Path $env:NEXTHINK\RemoteActions\nxtremoteactions.dll
```

Use the methods of the class to write the desired outputs. All write methods accept two arguments: the name of the output and the value to write. For instance, to write the size of a file to the Engine:

```
[Nxt]::WriteOutputSize("FileSize", $fileSize)
```

The editor of remote actions recognizes the calls to write outputs in the script and lists the output variables under the **Outputs** section below the script text. Set the label of the output to indicate how to refer to it in investigations and metrics.

The ending of each write method indicates the type of the output. Find the list of available methods and the corresponding PowerShell type of the value to be written in the table below:

Nxt write method	PowerShell type	Constraints
WriteOutputString	[string]	0 - 1024 characters (output truncated if bigger)
WriteOutputBool	[bool]	true / false
WriteOutputUInt32	[uint32]	<ul style="list-style-type: none">• Min: 0• Max: 4 294 967 295
WriteOutputFloat	[float]	<ul style="list-style-type: none">• Min: -3.4E+38• Max: 3.4E+38
WriteOutputSize	[float]	<ul style="list-style-type: none">• Min: 0• Max: 3.4E+38
WriteOutputRatio	[float]	

WriteOutputBitRate	[float]	
WriteOutputDateTime	[DateTime]	DD.MM.YYYY@HH:MM
WriteOutputDuration	[TimeSpan]	<ul style="list-style-type: none"> • Min: 0 s • Max: 49 days • Precision in seconds
WriteOutputStringList	[string[]]	Same as string

Interacting with end-users in self-help scenarios

To help end-users solve issues by themselves, combine remote actions with campaigns. In combination with remote actions, campaigns let you inform end-users about the detection of an issue and guide them through its resolution.

Because self-help scenarios make use of remote actions and campaigns, they require the purchase of both the Nexthink Act and the Nexthink Engage modules.

To display a campaign in the desktop of an end-user that is interacting with a device:

- The campaign must target users with a remote action and be published.
- The script of the remote action can be executed either:
 - ◆ In the context of the end-user, if the action does not need any special privileges.
 - ◆ In the context of the local system account, if the action needs administrative privileges.

Obtaining the UID of a campaign

The methods to run a campaign from a remote action that are detailed in the section below require the UID of the campaign to be passed as an argument. Thus, to run a campaign from a remote action, first you need to get the UID of the campaign.

To pass this UID to a remote action, it is convenient to declare a parameter in the script of the remote action for each campaign that it has to run and use the UID as the actual value for the parameter when editing the remote action.

To get the UID of a campaign and pass it to a remote action as a parameter:

1. Log in to the Finder as a user with the right to edit campaigns and remote actions.
2. In the **Campaigns** section of the left-hand side accordion, right-click the name of a campaign that you wish to launch from a remote action.
3. Select **Export > Campaign Uid to clipboard** to copy the UID of the campaign.
4. In the **Remote actions** section of the left-hand side accordion, double-click the name of the remote action that should run the campaign to edit it. The remote action should include a script that declares parameters for storing the UID of the campaign.
5. In the **Parameters** section under the script, select the parameter that should hold the UID of the campaign.
6. Press **Ctrl+V** to paste the actual UID of the campaign and assign it to the corresponding parameter.

Running a campaign from the script of a remote action

To interact with campaigns from a remote action, the script of the remote action requires to load a .NET assembly (`nxtcampaignaction.dll`) that is installed in the device of the end-user along with the Collector. The assembly includes the class **Nxt.CampaignAction** that provides the methods to control the execution of a campaign and get the responses from the user.

To load the assembly, add the following line at the beginning of the script:

```
Add-Type -Path $env:NEXTHINK\RemoteActions\nxtcampaignaction.dll
```

The methods of **Nxt.CampaignAction** to control campaigns are the following:

```
NxTrayResp RunCampaign(string campaignUid)
    Run the campaign identified by the campaignUid and block until the user
    finishes answering. Return the response in an object of type
    NxTrayResp.
NxTrayResp RunCampaign(string campaignUid, int timeout)
    Run the campaign identified by the campaignUid and block until the user
    finishes answering or after the specified timeout (in seconds) elapses.
    Return the response in an object of type NxTrayResp.
RunStandAloneCampaign(string campaignUid)
    Run the campaign identified by the campaignUid without blocking nor
    expecting results.
string GetResponseStatus(NxTrayResp response)
    Given a response object of type NxTrayResp, the method returns a string
    that reflects the status of the campaign. Possible values for the status:
```

- ◇ **fully**, the user has fully answered the questions of the campaign.
- ◇ **declined**, the user declined to participate in the campaign.
- ◇ **postponed**, the user accepted to participate in the campaign.
- ◇ **timeout**, the campaign was timed out before the user finished answering.
- ◇ **connectionfailed**, the script was unable to connect to the Collector component that controls campaign notifications.
- ◇ **notificationfailed**, the script was unable to notify the Collector component that controls campaign notifications.

```
string[] GetResponseAnswer(NxTrayResp response, string
questionLabel)
```

Given a response object of type `NxTrayResp` and the label that identifies a question in the campaign, the method returns the response given by the user.

- ◇ In the case of a single answer question, the returned string array has only one element.
- ◇ In the case of a multiple answer question, the returned string array has as many elements as answers selected by the user. Optional free text is ignored.
- ◇ If the user has not fully answered the campaign (i.e. status is not **fully**), the returned string array is empty. Optional free text is ignored.

Note that, for security reasons, remote actions for self-help scenarios ignore the optional free text answers of multiple answer or opinion scale questions. Therefore, it is useless to include optional free text answers in campaigns to be used exclusively in self-help scenarios.

Example code

Find below a simple example of how to run a campaign from a remote action. The example assumes that you have created a single-answer campaign where one of the possible answers is *yes*:

```
Add-Type -Path "$env:NEXTHINK\RemoteActions\nxtcampaignaction.dll"

$CampaignUid = "<UID of a single-answer campaign>"
$maxWaitTimeInSeconds = 60

$result = [Nxt.CampaignAction]::RunCampaign($CampaignUid,
$maxWaitTimeInSeconds)
$status = [Nxt.CampaignAction]::GetResponseStatus($result)

if ($status -eq "fully") {
```

```

$questionName = "question1"
$choiceName = [Nxt.CampaignAction]::GetResponseAnswer($result,
$questionName)
if ($choiceName -eq "yes") {
    # user has confirmed - let's do some actions:
    Start-Process notepad.exe
}
}

```

See a more elaborated example of self-help scenario with two campaigns [here](#).

Error handling

Nexthink detects whether the execution of a remote action was successful or not based on the return value of the PowerShell process that ran the script:

- An exit code value 0 indicates a successful execution.
- A non-zero value indicates some kind of error.

However, unhandled exceptions in PowerShell may cause the termination of a script without returning an appropriate exit code. To handle unexpected errors, Nexthink recommends you to begin the body of all your scripts with the following code snippet:

```

trap {
    $host.ui.WriteLine($_.ToString())
    exit 1
}

```

Place this default error handler right below the addition of the necessary DLL dependencies, which follow the optional declaration of formal parameters.

Performance measurement

To measure the performance and resource consumption of your scripts, turn on the logs of the Collector in debug mode with the help of the Collector Configuration tool:

```
nxtcfg.exe /s logmode=2
```

The output is stored in the `nxtcod.log` file.

The operations described in this article should only be performed by a Nexthink Engineer or a Nexthink Certified Partner.

If you need help or assistance, please contact your Nexthink Certified Partner.

Related tasks

- Executing remote actions
- Example of self-help scenario

Related references

- Nexthink Act Scripting Best Practices (Community)
- Nexthink Act Security Whitepaper (Community)
- Discussion forum on Nexthink Act (Community)

Example of self-healing scenario

Overview

For their daily operation, the devices of a public company must continuously run an enterprise application called `MustRun`. This application is not very stable though and it crashes intermittently. When `mustrun.exe` crashes, it is unable to shut down properly. As a result, the application usually leaves one of its log files in an inconsistent state. The corruption of this log file prevents `MustRun` from starting up again, until the file is deleted.

As soon as an end user realizes that `MustRun` is no longer running on a device, the user must restart the application by first deleting the corrupted log file and then relaunching `mustrun.exe`. This procedure is inconvenient for end users and has a negative impact on their productivity. In addition, inexperienced end users who are not familiar with the problem yet are very likely to require support from the help desk in order to solve it.

In this example, learn how to leverage Nexthink Act to automatically detect the crash of `mustrun.exe`, delete the log file, and restart the application. Automating the process results thus in increased productivity of the end users and a reduction on the number of reported incidents.

Creating the remote action

To remediate the problem of the `MustRun` application with Nexthink Act, create and schedule a remote action. The remote action shall include:

- An investigation that selects those devices on which `MustRun` crashed recently.
- A PowerShell script that deletes the corrupted log file and restarts the application on the selected devices.

Defining the target investigation

Start by creating an investigation that detects the crashes of the executable file `mustrun.exe` and returns the devices on which it happened. Since the Engine can insert events that lie up to 30 minutes in the past, set the time frame of the investigation to span the last 30 minutes. Combined with an appropriate scheduling of the remote action, this time frame ensures that no reported application crash is missed. This is a very conservative choice. Because of the speed at which Collector reports application crash events, time frames of around 10 minutes should be equally valid:

Because Nexthink detects application crashes, the investigation in the remote action of this example is already able to return the specific devices on which the problem occurs. On the other hand, in cases where Nexthink does not retrieve the information needed to detect the issue by default (e.g. a change in the value of a registry key), check the faulty condition in the script of the remote action itself. When the script checks the occurrence of a problem, the associated investigation must target all potentially impacted devices.

Therefore, we can classify the problem detection mechanism of a remote action as either:

- Investigation-based.
- Script-based.

Scheduling the remote action

After saving the investigation, create a remote action that periodically evaluates the investigation and runs the remediation script on the selected devices.

To properly schedule the remote action, configure the two periods:

Evaluation period

The time interval between two evaluations of the associated investigation. In our example, this value indicates how often the remote action checks for `MustRun` crashes. This period should be lower than or equal to the time frame of the associated investigation to not miss any application crash event. The smaller the evaluation period, the more responsive is the remote action, but also the more load is put into the system. To detect application crashes, an evaluation period of 10 minutes should be responsive enough. For critical applications, select a fast evaluation period as low as 1 minute.

Triggering period

The time interval between two consecutive triggerings of the remote action. For a remote action that detects issues by means of an investigation on events, such as application crashes, set the triggering period to be equal to the time frame of the investigation (30 minutes, in the example). This ensures that the execution of the script is not triggered more than once for the same event.

To associate the previously created investigation to the remote action, drag and drop the investigation onto the appropriate area of the editor of remote actions.

Adding the PowerShell script

Open your favorite text editor and type in the remediation script. Remember to encode the text of the script in UTF-8 with BOM when saving the file.

The script does the following:

1. Add the Nexthink dynamic library that deals with remote actions (`nxtremoteactions.dll`) by means of the **Add-Type** cmdlet.
2. Initialize the result of the script to the empty string.
3. Initialize a couple of variables with:
 - ♦ The path of the executable `mustrun.exe`.
 - ♦ The path of the corrupted log file to delete.
4. Try to remove the log file with the **Remove-item** cmdlet and set the result accordingly.
5. Restart the `MustRun` application with the **Start-Process** cmdlet.
6. Send the result to the Engine with the **WriteOutputString** function of the object **NXT** which was imported from the remote actions library.

```
Add-Type -Path $env:NEXTHINK\RemoteActions\nxtremoteactions.dll

[string] $result = ""

# The paths to the MustRun application and its log file
$mrexex = "$env:ProgramFiles\MustRun\mustrun.exe"
$logfile = "$env:ProgramFiles\MustRun\log.txt"

# Delete the log file if it is present
try {
    Remove-item $logfile -ErrorAction Stop
    $result = "The corrupted log file was deleted"
} catch {
    $result = "The log file does not exist"
}

# Restart the application
Start-Process -FilePath $mrexex

[NXT]::WriteOutputString("Result", $result)
```

For security reasons, Nexthink recommends that you sign your scripts. For testing purposes, it is safe though to use unsigned scripts in pre-production environments only.

In the editor of the remote actions, click **Import...** to link the script to the remote action. The Finder interprets the source of the script and lists the **Result** output under the **Outputs** section.

Adapt the previous script to your own use cases for profiting of the self-healing

capabilities of Nexthink Act.

Related references

- Creating an investigation
- Creating a remote action
- Timestamping of events

Example of self-help scenario

Overview

Albeit the storage capacity of hard drives is very large nowadays, it is still finite. With the daily use of a device, the hard drive tends to fill up not only with useful documents and programs, but also with many temporary files, downloads, images, etc. that consume ever more disk space. All this junk files leave less and less free space for new documents, applications, or even updates. When the amount of free disk space in a device goes below a certain threshold, the device becomes unusable because nothing else can be stored in it.

In this example, learn how to leverage both Nexthink Act and Nexthink Engage to warn end users of a low free disk space condition on their devices before it is too late and guide them through the process of liberating disk space by removing junk.

Creating the Campaigns

To warn end users of a particular issue and guide them through its resolution, self-help scenarios require the use of campaigns, which are available through the Nexthink Engage module. Remote actions are able to control the launch of campaigns and get the answers from the end users. Drive campaigns through remote actions to:

- Inform the end user of a particular issue.
- Propose a series of steps (possibly with different options) to resolve the issue.
- Notify the user once the issue has been solved.

To perform the actions described above, a self-help scenario typically demands two campaigns. In the case of our low disk space example, create the following two campaigns:

- **Propose disk cleanup**, a campaign that warns end users that they are running low on disk free space and proposes either:
 - ◆ A light clean of the disk.
 - ◆ A deep clean of the disk.
 - ◆ Not to clean the disk, with the following options:
 - ◇ Ask again later.
 - ◇ Never ask again.
- **Confirm disk cleanup completed**, a campaign that notifies the user when the cleaning process ends; as long as the user did not refuse to clean the disk in the previous campaign.

The **Propose disk cleanup** campaign offers a third choice that does not trigger any disk cleanup. This choice is suitable for end users who might decide to clean their disk later, never clean the disk despite the warning, or clean the disk themselves without the help of the remote action.

To indicate that the campaign is launched from a remote action, complete the sentence **Users are targeted** with the option **with a remote action** in the **RECIPIENTS** section of the editor. This is how the campaign **Propose disk cleanup** looks like in the editor of campaigns in the Finder:

And this is how the campaign **Confirm disk cleanup completed** is displayed:

Creating the remote action

To fix low free disk space issues, the remote action in our example shall include:

- An investigation that selects those devices with low system drive free space.
- A script that gathers the answers of end users to the questions of the first campaign, **Propose disk cleanup**, and decide whether to clean the disk of the device or not based on those answers.

Defining the target investigation

Before creating the remote action itself, create an investigation that returns the devices with 5 GB of disk space available or less on their system drive. The investigation shall filter out those devices whose owners declined the proposal to clean the disk. Contrary to standalone campaigns, find the answers to campaigns launched from remote actions as fields of the **Device** object and not of the **User** object.

Scheduling the remote action

Create now the remote action and link the previous investigation to it. In the remote action editor:

1. Tick the option to automatically run the remote action.
2. Drag and drop the investigation onto the appropriate area within the **SCHEDULE** section.

As evaluation period, the remote action executes the investigation every 10 minutes for more responsiveness. To reduce the Engine load, setting an evaluation period of 1 hour should be enough as well.

As triggering period, the remote action specifies 1 day. In that way, the remote action will not shortly bother the same user again if the device still suffers from low free disk space, but it will wait until the next day.

Adding the PowerShell script

In your favorite text editor, type in the remediation script. Remember to encode the text of the script in UTF-8 with BOM when saving the file. Failing to do so results in the script not running on the target devices when the remote action is triggered.

The script does the following:

1. Define two parameters of the type string to hold the identifiers (Uid) of the associated campaigns.
2. Add the Nexthink dynamic library that deals with remote actions (`nxtremoteactions.dll`) by means of the **Add-Type** cmdlet.
3. Add the Nexthink dynamic library that deals with campaigns launched from remote actions (`nxtcampaignaction.dll`) by means of the **Add-Type** cmdlet. This library includes the object `[Nxt.CampaignAction]` to handle the launching and get the answers of campaigns.
4. Define a trap for error handling.
5. Define a function **HandleFullResponse** to get the answer of a user to the questions of the first campaign and take action accordingly: clean the disk or handle a negative answer. Note the use of the function `[Nxt.CampaignAction]::GetResponseAnswer` to get the actual value of the answers to the campaign.
6. Define a function **CleanupDisk** that removes all files in the Recycle Bin (note that the cmdlet **Clear-RecycleBin** is available from PowerShell 5.0 only). For simplicity, this example script does not make a difference between a light and a deep cleanup although it accepts the type as an argument to the function. Note the use of the function `[Nxt.CampaignAction]::RunStandaloneCampaign` to launch the second campaign after the cleaning up. It is launched as standalone because it is just an informative campaign and there is no answer to collect.
7. Define a helper function **GetRecycleBinSize** to get the total size of all the files in the Recycle Bin for reporting the amount of space cleaned to the

- Engine as an output.
8. Define a function **HandleNegativeResponse** to ask users who do not want their disk to be cleaned if they want just to postpone the action or never be asked again.
 9. Define a helper function **WriteOuptus** that sends back to the Engine the response of the user and the number of bytes erased, if any.
 10. Run the first campaign to propose users the cleaning of their disk. Note the use of the function `[Nxt.CampaignAction]::RunCampaign` with the Uid of the first campaign and a timeout of 120 seconds.
 11. Get the response status with the help of the function `[Nxt.CampaignAction]::GetResponseStatus`.
 12. Depending on the status of the response:
 - ◆ Call the function **HandleFullResponse**, if the user fully responded to the questions of the first campaign.
 - ◆ Postpone the campaign and return this state to the Engine, if the user answered to postpone or if the campaign timed out.
 - ◆ Decline the campaign and return this state to the Engine, if the user declined the proposal to clean the disk.

```
#
# Clean Recycle Bin
# Version 1.0.0.0
# Copyright (C) 2017 Nexthink SA, Switzerland
#

param(
    [string]$ProposeCampaignId,
    [string]$ConfirmCampaignId
)

Add-Type -Path $env:NEXTHINK\RemoteActions\nxtremoteactions.dll
Add-Type -Path $env:NEXTHINK\RemoteActions\nxtcampaignaction.dll

trap {
    $host.ui.WriteLine($_.ToString())
    exit 1
}

function HandleFullResponse($response) {
    $responseAnswer =
    [Nxt.CampaignAction]::GetResponseAnswer($response, "Clean disk")[0]
    switch($responseAnswer) {
        "Light" { CleanupDisk($responseAnswer) }
        "Deep"  { CleanupDisk($responseAnswer) }
        "No"    { HandleNegativeResponse($response) }
        default { throw "Invalid answer, please ensure the
first question name is set to 'Clean disk'" }
    }
}
```

```

    }
}

function CleanupDisk($type) {
    $path="$($env:SystemDrive)\`$recycle.bin"
    $bytesCleaned = GetRecycleBinSize($path)
    Clear-RecycleBin -Force -ErrorAction SilentlyContinue
    WriteOutputs $type $bytesCleaned
    [Nxt.CampaignAction]::RunStandAloneCampaign($ConfirmCampaignId)
}

function GetRecycleBinSize($path) {
    $measure=0
    Try {
        $measure=(Get-ChildItem $path -force -recurse
-ErrorAction SilentlyContinue) | measure length -sum
    }
    Catch {
    }
    return $measure.sum
}

function HandleNegativeResponse($response) {
    $responseAnswer =
    [Nxt.CampaignAction]::GetResponseAnswer($response, "Ask again")
    switch($responseAnswer) {
        "Postponed" { WriteOutputs("Postponed") }
        "Declined" { WriteOutputs("Declined") }
        default { throw "Invalid answer, please ensure the
second question name is set to 'Ask again'" }
    }
}

function WriteOutputs($userReply, $bytesCleaned){
    [Nxt]::WriteOutputString("UserReply", $userReply)
    [Nxt]::WriteOutputSize("SpaceCleaned", $bytesCleaned)
}

$response = [Nxt.CampaignAction]::RunCampaign($ProposeCampaignId, 120)
$responseStatus = [Nxt.CampaignAction]::GetResponseStatus($response)
switch($responseStatus) {
    "fully" { HandleFullResponse($response) }
    "timeout" { WriteOutputs("Postponed") }
    "postponed" { WriteOutputs("Postponed") }
    "declined" { WriteOutputs("Declined") }
    default { throw "Failed to handle campaign response: "
+ $responseStatus }
}

```

For security reasons, Nexthink recommends that you sign your scripts. For

testing purposes, it is safe though to use unsigned scripts in pre-production environments only.

In the editor of the remote actions, click **Import...** to link the script to the remote action. The Finder interprets the source of the script and lists both the detected parameters and the outputs under the script. The **Parameters** of the script are the Uids of the two campaigns previously created and the **Outputs** are the response of the user and the size of disk space cleaned.

To pass the Uid of a campaign to the remote action:

1. In the left-hand side menu of the Finder, select the **Campaigns** section.
2. Right-click the desired campaign.
3. From the menu select **Export > Campaign Uid to clipboard**.
4. In the left-hand side menu of the Finder, select the **Remote actions** section.
5. Double-click the name of the remote action that will launch the campaign to open the editor of remote actions.
6. In the **Parameters** section under the text of the script, right-click the text input box that corresponds to the string parameter that holds the Uid of the campaign in the script.
7. Select **Paste** to insert the Uid of the desired campaign.

Advanced configuration

Since the remote action does not require administrative privileges, tick the option to run the script as the **current interactive user**. Remote action that require administrative privileges must be run as the **local system** account instead.

Allocate enough time so that the user can answer the campaign and the script is able to finish before being timed out. The example defines a timeout value of 300 seconds (5 minutes).

Organizing objects with categories

Classifying objects of the same type

Overview

Nexthink distinguishes ten types of objects to describe your IT infrastructure:

• User	• Package	• Executable	• Port	• Domain
• Device	• Application	• Binary	• Destination	• Printer

When you create an object-based investigation or widget, you target a particular type of object. Object types are predefined and their number is fixed. However, you might want not to target all the objects of a type, but only a custom-defined subset. To that end, Nexthink introduces the concept of *Categories*. Categories work one level below the object type, defining classes of objects within objects of the same type. For instance, you can create different classes of users, devices or applications according to the processes and policies in your own organization. The ability to define your own subtypes of objects lets you create investigations and widgets specifically targeted to those subtypes.

Each category applies thus to only one type of object. The category defines a set of classes, each one identified by a *keyword*, that partitions the group of all the objects of the given type into disjoint subsets. Therefore, at most one keyword per category can be assigned to an object in an Engine. If an object does not fit into any of the classes defined by the category, the object gets assigned the *empty* keyword for that category.

Tagging objects

The process of assigning keywords to objects is called *tagging*. A tag is a label consisting of a pair **Category:Keyword** that is attached to the object. The word *tag* is therefore commonly used as synonym for keyword in Nexthink parlance.

Tagging can be manual, automatic (based on a set of auto-tagging rules) or semi-automatic (based on tags loaded manually from a CSV file).

Displaying the tags of objects

For object-based investigations, optionally display a column for each available category. Add the desired category-related columns either when editing the options of the investigation or while navigating through the list of results. The columns of categories show the keywords assigned to each object.

If a keyword is assigned to an object as a result of automatic tagging, a small icon with the letter **A** appears to the right of the name of the keyword. If you tag the object manually or you import the tag from a file, no special indication is given.

Displaying multiple tags in Cross-Engine scenarios

In an installation with multiple Engines, an object that is present in several Engines can be tagged differently on each Engine. This is seldom for keywords assigned with auto-tagging rules, but it is common if an automatically assigned keyword is manually overridden in one Engine only.

Thus, when examining the results of a Cross-Engine investigation on the **List (all entities)** view, an object may display more than one keyword in a category column. All the keywords assigned to the object are listed in sequence, separated by a comma from one another. In this case, an icon with the letter **A** appears to the right of the list of keywords only if all the keywords have been automatically assigned on every Engine.

Device objects are special though because a single device is identified as a different object by each Engine. Therefore, in the case of devices, at most one keyword per category is allowed, even when considering Cross-Engine results.

Platforms and automatic tagging of devices

When you automatically tag devices based on the value of one of their attributes, remember that not all attributes are present in devices of different platforms.

Related tasks

- Editing the options of an investigation
- Navigating through the results of an investigation
- Tagging objects manually
- Tagging objects automatically
- Importing tags from text files

Related references

- Local and shared content

Related concepts

- Category
- Keyword

Creating categories and keywords

Categories and keywords let you classify objects of the same type into user-defined groups. Only users with the appropriate permissions can create or edit categories.

Create a new category

To create a new category using the Finder:

1. Go to the **Categories** section on the left-hand side of the main window.
 - ◆ If you want to create the new category at the top-level, right-click the header or the empty area of the **Categories** section.
 - ◆ If you want to create the new category inside a particular folder, right-click an existing folder in the **Categories** section.
2. Select **Create new category**.
3. From the choice list, select the type of object on which the new category is based.
4. A new category, with the temporary name **Untitled Category n**, where *n* is a growing number to avoid name clashes, shows up in a new tab.
Optional: Click the name of the category and replace it with an appropriate name.
5. Optional: Click the field below the name of the category labeled **Enter description here...** and give a short description on how the new category classifies objects.
6. Select the level of privacy associated to the category. For more information, see the chapter on Privacy.
7. Add new keywords to the category (follow the set of instructions below).
8. Optional: Depending on the number of the keywords that you define, creating a category can be a lengthy process. Click the **Save** button after a worthy modification to make sure you do not lose your job.

9. Click the **Save and close button** to save your last changes and close the edition tab. Alternatively, click the **Close** button if you already saved your work.

Edit an existing category

To edit a category:

1. Find the category to edit. You can either:
 - ◆ Search the name of the category in the Search box of the welcome screen
 1. Click the desired result of the search.
 - ◆ Browse the **Categories** section on the left-hand side of the main window.
 1. Right-click the name of the category.
 2. Select **Edit...**
2. Modify the name, description or keywords of the category.
3. Click **Save and close**.

Add and edit keywords

To add new keywords to a category:

1. Create a new category or edit an existing category.
2. Locate the buttons for editing keywords at the bottom left corner of the panel.
3. Click the button with the plus sign to add a new keyword.
4. Type a name for the new keyword and click **OK**. You cannot repeat the same name for a keyword in a category.
5. Optional: Set the auto-tagging conditions for the keyword. See the chapter on Tagging objects automatically.
6. Optional: To save you some work, click the button picturing two documents to make a copy of an existing keyword if you need to create similar classes of objects. The auto-tagging conditions that you defined for one keyword are copied to the other.
7. Optional: Click the button with a pencil to rename a keyword.
8. Optional: Click the button with the trash can to remove a keyword.

All the functions of the buttons for editing keywords are also available in a context menu by right-clicking the **Keywords** panel.

Related tasks

- Adding users (create a profile that allows the edition of categories)

Tagging objects manually

Once you have defined a category, you can manually tag the objects of the type associated to the category. *Tagging* is applying a keyword of a category to an object. Thus the object becomes a member of the class of objects identified by the keyword.

Some category keywords include auto-tagging conditions in their definition for tagging objects automatically. Other keywords may not define any auto-tagging conditions. To manually tag objects you can use both types of keywords: with or without auto-tagging conditions.

In any case, note that when you tag an object manually, you override the automatic tagging based on conditions for that object. Despite the override, manually tagging an object does not change the definition of a category keyword. The modification applies to the object locally in the Engine and not to the definition of categories, which are themselves centralized and apply to all Engines.

To manually tag an object with a keyword in the Finder:

1. Find the object to tag in the list of results of an investigation. For that purpose, you can use the default investigation that executes after you edit the category that defines the keyword.
2. Right-click the name of the object and select **Edit....** A dialog with the list of all the categories that you can apply to the object pops up.
 - ♦ Note that the option to edit an object is not available in Cross-Engine contexts; that is, when examining the results of an investigation from the **List (all entities)** view.
3. Next to each category, you have a drop down list to choose a keyword for the object. By default, the system tags objects automatically, so the option **auto** appears selected in the drop down. Click the down arrow of the drop down list to see all the available keywords for that category.
4. Select the desired keyword to tag the object. Otherwise, you can assign the *empty* keyword or let the system automatically tag the object by choosing the *auto* option.
 - ♦ Optional: If you do not find a suitable keyword for the object, you can add a new keyword to the category by selecting **Add new keyword...** from the drop down list.

5. If the operation was successful, the dialog finishes with a message displaying the number of objects that were updated. Click **OK** to close the dialog.

You may apply the same keyword to a group of objects if you select several objects instead of just one. Use the **Ctrl** key to click multiple objects or the **Shift** key to select a range of objects in the list. Then right-click and proceed as explained above.

Related references

- Local and shared content

Tagging objects automatically

Depending on the number of objects to classify, tagging objects manually can be a very long and inefficient process. Alternatively, you can let the system automatically tag objects for you, as long as you instruct the system on how to classify objects according to the values of their attributes. This is done by assigning *auto-tagging conditions* to the keywords of a category.

Defining auto-tagging conditions

To define auto-tagging conditions for a keyword in the Finder:

1. Create or edit a category (see the chapter on Creating categories and keywords).
2. In the **Keywords** panel, select the desired keyword. The right hand side of the panel shows the auto-tagging conditions for the keyword. If no condition has been defined yet, this part of the panel is empty the keyword is never assigned automatically to any object.
3. Click the link **Click here to add a new condition** to specify an auto-tagging. The way to express auto-tagging conditions is similar to the way you specify the conditions of an investigation. A row with three drop down lists appears.
4. Choose the attribute of the object whose value will be compared in the first drop down list. For device objects, beware that all attributes are available for selection, but not all of them are necessarily available for all of the platforms.
5. Select a comparison operator in the second drop down list.

6. In the last drop down list or combo box, select or type the value to compare with the selected attribute. Once you define a condition for a keyword, a small icon with the letter *A* appears to the left of the name of the keyword in the panel, indicating that the keyword is used in automatic tagging.
7. Optional: Click the trash can placed to the right of the drop down lists to remove a condition.
8. Optional: Go back to point 3 to add new conditions. If you create more than one auto-tagging condition for one keyword, you can specify how the conditions combine using the **Logical expression** field that appears below the conditions. By default, conditions are combined by a logical or. Therefore, it is enough for an object to fulfil one of the conditions to be tagged with the keyword. You can use the logical operators *OR* and *AND* to combine the auto-tagging conditions of a keyword.
9. Click **Save** to permanently save your work and continue editing the category or **Save and Close** to save your work and end the edition of auto-tagging conditions.

Setting the precedence of automatic keywords

An object can be tagged by at most one keyword per category. If an object satisfies the auto-tagging conditions imposed by two or more different keywords, the system tags it with only one of the keywords. You can decide what keywords take precedence over others by establishing a ranking of automatic keywords. Remember that tagging an object manually overrides any automatic assignment of keywords.

To set the precedence of automatic keywords in the Finder:

1. Create or edit a category that holds automatic keywords, that is, keywords that specify auto-tagging conditions (see above).
2. Click the button **Set auto-tagging order...** located below the list of keywords. A dialog pops up with the list of all the automatic keywords, ordered by precedence. By default, keywords are ordered in alphabetic order.
3. Click the name of a keyword in the dialog to change its order of precedence.
4. Click the button **Move up** to promote the selected keyword to a higher rank or the button **Move down** to lower the precedence of the keyword.
5. Click **OK** once you are satisfied with the ordering of keywords or **Cancel** to ignore the changes made.

Performance considerations

Nexthink triggers the automatic tagging process for all objects of a given type when you create or modify a category that applies to that type. The modification of a category can indeed imply a modification of the auto-tagging conditions, so every object must be rechecked against the new conditions. At a lower scale, the modification of the attributes of an object may also trigger the automatic retagging of the object. If the modified attribute is suitable for comparison in auto-tagging conditions, the new value of the attribute can make the object fall into a different class and, therefore, be tagged with a different keyword.

Although the system does not impose a hard limit on the number of categories and keywords that you can define, automatic tagging can become a costly operation depending on the number of computations required. Specifying too many keywords or auto-tagging conditions may have a significant impact in the overall performance of the system. The maximum recommended values for keeping the system responsive are listed below:

- 25 categories per type of object (device, user, application, etc).
- 200 automatic keywords per category.
- 20 auto-tagging conditions per keyword.

Importing tags from text files

Overview

Importing tags from a file is a semi-automatic way of tagging objects. Tagging with the help of a text file with CSV format is particularly useful in situations where the classification of objects depends on an external source of information. In such cases, you cannot use automatic tagging because the classification is independent of the attributes of the object and using manual tagging can take too long if the number of objects to tag is high.

The result of importing tags from a file is the same as if the objects were tagged manually. That is, importing tags overrides automatic tagging for the objects and the categories involved. As in manual tagging, importing tags from a text file changes the objects locally in a particular Engine and not the objects in other Engines nor the definition of the categories. Categories are centralized, but importing tags from text files is not.

Again, because tagging from text files is equivalent to manual tagging, the last tag that is applied to an object (within a category) is the tag that prevails. In your CSV file, if multiple tagging rules that belong to the same category match a particular object, the import process successively tags the object for every matching rule. Therefore, the last matching rule in the file gives the object its definitive tag. Beware that the list of rules may appear differently in the import dialog. The order of the rules in the file determines their order of application.

Format of the CSV file

The file to import must be a text file organized as a table, where each line of text represents a row of the table and a delimiter character (usually a comma) separates values into columns within a row. This is commonly known as a *Comma-Separated Value* file or CSV file. For the purpose of importing tags, the CSV file is organized into columns of keywords. The first element of each column is the name of the category to which the keywords in the column belong. A special column is used to match the lines of the CSV file with the objects to be tagged. Unlike other columns, this column does not hold keywords of a category, but values of an attribute that designates the object to be tagged. Only attributes that identify an object, such as the UID or the name, can be used to do the match. The allowed attributes for matching depend on the type of the object. Each line of the CSV file can therefore be seen as a row of a table, where one value identifies the object to be tagged and the rest of the values are the keywords that will be applied to the object.

The CSV file can have other columns with information that is not related to categories. These columns can be ignored during the import process.

Limit on the size of the CSV file

The maximum size of the CSV file to import tags is **140 KB**, which corresponds approximately to 2'500 rows.

If the number of objects to tag is bigger, split the CSV file into two or more files and import them separately.

Import process

To import tags from a CSV file:

1. In the Finder, click the sprocket icon located on the right hand side of the header of the main window.

2. Select **Import tags from file...** from the menu.
3. Choose the CSV file in the standard dialog to open files. A wizard to set the import options shows up.
4. Select the text encoding (e.g. Western European or UTF), the character that delimits columns (e.g. comma or semicolon) and the character that qualifies a value as text (e.g. single or double quotes).
5. Optional: Click the button **Show file** to see the file that you are importing. This is useful to validate your previous selection of text encoding and special characters.
6. Click **Next** to go on with the import process.
7. Select the type of the objects that you want to tag in the drop down list labeled **Apply on**.
8. Select the attribute of the object that will be used to do the matching in the drop down list **Match objects with**.
9. Choose the column of the CSV file whose values match those of the attribute of the object in the drop down **using column**.
10. Check the columns of the CSV file that correspond to the categories that you want to import in the list entitled **Create or update categories with columns**.
11. Optional: Click the button **Select All** to select all the columns at once.
12. Optional: Click **Unselect** to remove the selection from all the columns.
13. Click **Import** to proceed with the importation of tags. A final dialog shows you the number of objects that have been tagged and the problems found during tagging. If needed, repeat the import until you are satisfied with the result.

Related references

- Local and shared (centralized) content

Getting notified by the system

Receiving Engage campaigns

Overview

As soon as a campaign is published, the set of end users that constitute the target audience starts receiving the campaign on all their devices that run the Collector.

As a first step, campaigns optionally display a desktop notification that, in addition to inform end users about both the subject matter and the sender of the campaign, lets them decide whether to participate in the campaign or not. Alternatively, for short, urgent, or important campaigns that users should not decline, configure campaigns to skip the notification step and directly jump into the first question.

The Windows Collector includes a tray application that displays campaigns with a unified look-and-feel across all types of Windows devices. In particular, campaign notifications look similar to standard Windows 10 notifications in all supported versions of Windows. Brand your campaigns for Windows users with your own corporate logo and colors.

Applies to platforms:

For its part, the Mac Collector displays campaign notifications through the standard macOS notification mechanism. By default, users receive *alert* type notifications, so the notification stays on the screen until the user clicks either the button to answer the campaign or to postpone it.

For each individual application, macOS 10.15 Catalina and later prompts users to either allow notifications or not. In the case of the Mac Collector, a dialog shows up after installation to let the user decide whether to allow notifications from the Collector or not. To prevent users from blocking campaign notifications, configure the profile of users to enable notifications for the Engage component of the Collector.

Applies to platforms:

Conditions to receive campaigns

For a user to receive a campaign, the following conditions must be met:

- The device is turned on.
- The user is logged in to the device.
- The Collector is running on the device and has an active TCP connection with the Engine.

Campaign delivery is retried every 15 minutes on each device that does not fulfill the conditions.

In addition, the Collector does not display the campaign in the following cases:

- The device is running an application in full-screen mode, as displaying the campaign might annoy the user.
- The user has not been active for the last 15 seconds. The campaign might go unnoticed if the user is absent.
- The user did not allow notifications from the Mac Collector.

Accepting or rejecting to participate

If the campaign includes a notification When receiving a notification from a campaign, the end user sees:

- The sender of the campaign, including an optional picture.
- The title of the campaign.
- A text inviting to take the campaign (the configured *footer*).
- A couple of buttons to decide whether to answers the questions of the campaign. The user has two choices:
 - ◆ Click **Show** to take the survey of the campaign immediately.
 - ◆ Click **Not now** to skip the participation on the campaign.

If the user clicks the **Not now** button to skip the campaign, another dialog pops up with the question *When would you like to be reminded?*. The user is presented with two choices:

- Click **Never** not to receive any other notification regarding this campaign ever again in the current device. This choice is remembered as long as the campaign is active. The user gets notified again in the same device only if the campaign is retired and then republished.
- Click **Later** to answer the questions of the campaign at some point in the future. The value in the configuration setting of the campaign **Reminder**

frequency determines the time when the user will receive again a notification for participating in the campaign.

Note that the participation of a user in a campaign is local to each device that the user utilizes. For instance, if a user declines to participate in a campaign from one particular device, the user may still receive notifications about the same campaign in other devices.

Related tasks

- Scrutinizing the results of a campaign
- Enabling Engage notifications in macOS
- Branding of campaigns

Receiving email digests

Overview

Email digests provide you with a concise update on what happened during the past week. Receive individualized digests that cover your areas of responsibility, as regards the modules attributed to you in the Portal. Digests give you an overall summary of the status of your modules based on the metrics that interest you the most and on key figures of the defined services.

Setting up the reception of email digests

Before using email digests, make sure that the main administrator has properly configured email notifications in the Portal appliance via the Web Console, including:

- The configuration of email notifications for alerts and digests in the Appliance (SMTP).
- The configuration of the base address of the Portal for the links in the email digest.

By default, once the Portal is configured to send email digests, users receive weekly digests with information of all the modules of the types Basic or Service within their own Portal content (that is, personal modules or modules included in their roles). Therefore, if you have access to published modules and you want to get information about them in your email digests, you must make a copy of the modules to your local content first.

Users may configure their accounts to stop receiving email digests or specify the kinds of modules that they shall include (personal, role-based, or both):

1. Log in to the Portal with your user account.
2. Click the label that displays your user name (followed by a downwards arrow) in the top right corner of the main window.
3. Select **My account** from the drop-down list.
4. In the **SETTINGS** section of the **My account** dialog, set your preferences under **WEEKLY EMAIL DIGEST**.
5. Tick the box **Send me the digest for** to receive emails digests for the modules selected in the drop-down list below:
 - ◆ Select **all my modules**, to get an email digest that includes information from all the modules that are part of your Portal content.
 - ◆ Select **only modules in my roles**, to get an email digest with information coming only from those modules which are attributed to you through your roles.
 - ◆ Select **only modules in my personal content**, to get an email digest with information coming only from those modules that you created yourself (or that you copied locally to your personal content).
6. If there is more than one hierarchy defined in your Portal, select the hierarchy according to which the email digests must be computed from the list labelled **and computed for the hierarchy**.
7. Click **Save**.

Optionally, it is also possible to brand the logo displayed in email digests.

Configuring the contents of email digests

For each module of type Basic within your Portal content (that is, for each module based on metrics), choose the metrics and scores that you want to include in the digest. In the case of service modules, the contents to be included in the digest are automatically generated out of the services overview and the number of reported issues.

By default, a module adds no content to the digest. Therefore, configure your modules based on metrics to add elements to your email digest. Since service modules do not allow such fine grained control, their configuration just lets you individually include them or exclude them from the email digest.

To configure the digest of a module based on metrics:

1. Log in to the Portal with your user account.

2. Browse any of the dashboards of a module based on metrics by selecting it from the navigation tool that is placed on the left-hand side of the top blue ribbon.
3. Click the menu icon (three bars) on the right-hand side of the top blue ribbon.
4. Select **Configure email digest...** from the drop-down list. A dialog to choose metrics from the current module shows up.
5. Click one of the three entries with a plus sign to add a new metric to the digest.
 1. Select a metric from the list **SELECT METRIC TO ADD**. Note that you can only select metrics from the current module.
 2. Optional: Search the metric by name with the help of the search box on the top right of the list.
 3. Click the **Add** to include the selected metric in the email digest.
6. Optional: Repeat the previous step to add more metrics to the digest from the same module. You can add up to three metrics from the same module in an email digest.
7. Optional: Click **Send me a preview by email** to get a glance of how the email digest looks like with the added metrics.
8. Optional: Click the cross icon to the right of an already filled entry to empty it and remove the metric from the digest.
9. Click **Done** to finish the configuration of the digest.

To configure the digest of a module based on services:

1. Log in to the Portal with your user account.
2. Browse any of the dashboards of a service module by selecting it from the navigation tool that is placed on the left-hand side of the top blue ribbon.
3. Click the menu icon (three bars) on the right-hand side of the top blue ribbon.
4. Select **Configure email digest...** from the drop-down list. A dialog shows up.
5. Tick **Include this module in the email digest** to summarize the contents of the module in the email digest.
6. Optional: Click **Send me a preview by email** to get a glance of how the email digest looks like with the added service module.
7. Click **Done**.

Interpreting email digests

The digest displays an emoticon followed by a message about the general status of the module during the past week. The facial expression of the emoticon reflects the status of the module, offering you instant insight into how well the

module is doing. Below the emoticon and the general status message, find individual information about the services or metrics that make up the dashboards in the module.

Email digests have been designed to present essential information in a clear and appealing manner. The emoticons and messages in the digest are concise, easy to understand, and they get right to the point. They try to ensure that issues do not go unperceived by highlighting them in the first place.

Color codes also help you quickly identify where issues are detected. For figures in the digest on services, or on metrics for which you have defined thresholds, the digest displays them in color following the usual convention: green indicates good news, red is bad news, and yellow something in between (when available). Blue is used for those metrics whose variation is not necessarily good or bad; that is, those metrics for which you have not indicated whether growth is positive or negative in their definition.

Opening the Portal from an email digest

Each figure in the digest, be it from a Service module or from a Basic module, is clickable. When configured properly (see the section on setting up the reception of email digests above), clicking on a figure in the digest opens the Portal in the correct dashboard; that is, the dashboard that holds the metric or service associated to the selected figure.

If you click on a figure of the digest but you have not started a Portal session yet, the Portal will display the login page. After logging in, you will be redirected to the appropriate dashboard.

Related tasks

- Creating a dashboard
- Sending email notifications from the Appliance
- Branding the Portal

Receiving alerts

Overview

To get notified by the system in the case that a particular situation happens in your IT infrastructure, use alerts. You can receive alerts via email or, in some

cases, through the system log of the Appliance. Additionally, you can see the occurrences of service-based alerts in the Portal or those of investigation-based alerts on devices in the Device view of the Finder. To receive alerts via email, configure first the SMTP settings of the Appliance.

Find below a table that summarizes the types of alerts, the potential receivers of the alert and the possible mechanisms to send the alert:

Alert type	Recipient	Sending mechanism
System alert	<ul style="list-style-type: none"> • administrator 	<ul style="list-style-type: none"> • email • syslog
Investigation-based alert	<ul style="list-style-type: none"> • Email addresses specified in the definition of the alert • Users that have the alert assigned to one of their roles 	<ul style="list-style-type: none"> • email • syslog (only global alerts) • Device view (for investigation-based alerts on devices)
Service-based alert	<ul style="list-style-type: none"> • Users that add the alert to their Portal • Users that have the alert assigned to one of their roles 	<ul style="list-style-type: none"> • email • Dashboards My alerts and History in the Portal

Receiving system alerts

Configure the email address of the admin account to receive system alerts. The Engine sends system alerts by email to inform the administrator of the status of the system. Remember that system alerts are predefined alerts that notify special circumstances detected by the Engine, such as internal errors or the expiration of your software license. No user can modify them.

To send emails, the Engine uses the SMTP configuration of the Appliance where it is installed. Only the main administrator receives emails from system alerts. In addition to sending system alerts via email, the Engine logs each occurrence of a system alert in the system log of the Appliance.

Receiving service-based alerts from the Portal

The Portal can generate service-based alerts, which are always related to service dashboards. The triggering of a service-based alerts depends therefore on the results of their corresponding service, as seen by each particular user (their view domain). When a service-based alert is triggered, it appears as active in the special dashboard **My Alerts** of the Portal. In addition, the Portal sends emails to subscribers of the alert both when the alert is triggered and when the alert terminates.

To determine when a service-based alert is triggered and terminated, it is important to know about the *sliding window* mechanism of services. The values of a service are sampled each 10 minutes, but these values are added along one hour intervals for the purpose of generating alerts. Thus, an alert will be triggered or terminated depending on the values collected by its associated service during the last hour.

For example, let us suppose that you set the ratio of devices with errors in a service to 10% and that you create an alert associated to this service. In the next 10 minutes, the service reports that 15 out of 100 devices are in error. Therefore, since the ratio of devices in error is 15% and there was no other previous data available for the last one hour, the Portal triggers the alert associated to the service. Now let us consider the status of the alert depending on the results of the service after the next interval of 10 minutes:

The service reports that no device used it

The accumulated value of the service for the last hour is still 15 devices in error out of 100; the same as in the previous 10 minutes. Therefore, the alert is still active. Following the same reasoning, if no one uses the service for one hour, the alert remains active until the end of the hour, when the first use reported (15%) lies behind the last hour window. At that point, the alert terminates.

The service reports that 5 out of 100 devices are in error

Supposing that the 100 devices that accessed the service during this interval are the same devices that accessed the service during the previous interval, the accumulated number of devices with error for the last hour are at least 15 (more if the 5 newly reported devices are not part of the 15). The alert is therefore still active.

The service reports that 15 out of 300 devices are in error

Supposing that the 15 devices with error are the same 15 devices with error of the previous interval, and that the 300 devices that accessed the service include the 100 devices of the previous interval, the ratio of devices in error to devices accessing the service becomes now 5%. Since

5% is lower than the threshold of 10% configured in the real-time service widget, the alert terminates immediately.

Another factor to take into account is the minimum number of devices impacted that you configured in the alert. If you set the option **ignore until x or more devices are impacted**, beware that **impacted device** in this context means *device accessing the service* and not *device with error*. In our previous example, the number of impacted devices during the first 10 minutes interval is therefore 100, and not 15.

Finally, it is important to note that the triggering of an alert also depends on the **view domain** of its subscribers. In the same way that the results of a service dashboard depend on the view domain of the user who watches them, a service-based alert is sent to one of its subscribers only if the alert conditions are met within the view domain of that subscriber. That is, service-based alerts are triggered individually per each subscriber and not globally. Using again our previous example, if the value of 15 devices in error out of 100 (15%) was seen by a user with a complete view of the system and another user with a restricted view domain can only see 1 device in error out of 25 (4%), only the first user will receive the alert (>10%).

Moreover, if you have defined multiple **hierarchies** in your setup, each user has a view domain per hierarchy. An alert can be triggered when its conditions are met within any of these different view domains. The actual view domains in which the alert was triggered and their corresponding hierarchy are detailed in the email that the subscriber of the alert receives. Therefore, only one email is sent per alert and per subscriber even if the alert is triggered in multiple domains.

Viewing service-based alerts in the Portal

To see information about the occurrences of service-based alerts in the Portal:

1. Log in to the Portal as a particular user.
2. Click the bell icon in the top right of the Portal window and choose either:
 - ◆ **My alerts**, to see a table with the list of alerts to which you are subscribed and information on their last occurrence.
 - ◇ If an alert is still active, a bell icon is displayed in the first column of the **My alerts** table for that alert.
 - ◆ **History**, to see the past occurrences of the alerts to which you are subscribed.
 - ◇ It is possible to filter the results of the **History** table by hierarchy, alert name, and date by editing the controls on the top right of the dashboard.

In both dashboards **My alerts** and **History**, it is possible to look at the details of an alert, which enable you to break down the results by hierarchy levels:

1. Hover the mouse over the name of the alert in an entry of the table and an info icon (a little **i** with blue background) appears to its right.
2. Click the info icon to see the details of the alert.
3. In the dialog with the details, select the desired hierarchy from the list **Display impact data for hierarchy**.
4. Select a level of the hierarchy from the list **Hierarchy level** to get the results of the alert for a particular level only.
5. Optional: Click the right pointing arrow with the label **CSV** at the bottom right corner of the table to export the details of the alert to a CSV file.

Receiving alerts linked to roles

Even when you do not create alerts yourself, you can receive alerts that are assigned to you by the roles attributed to your user account. You cannot modify or delete any of the alerts that are linked to your roles. The alerts that can be linked to roles are of two types:

- Investigation-based alerts
- Service-based alerts

To see the investigation-based alerts that are attributed to you according to your roles:

1. Log in to the Finder with your user account.
2. Click the **Settings** section of the accordion on the left-hand side of the Finder.
3. Inside **Settings**, select **Role-based alerts** from the drop down list labeled **Section**. For each one of your roles, you see a folder holding the alerts that correspond to the role.

Moreover, in the Device view of the Finder, there is a separate alert timeline for each one of your assigned roles. Alerts on devices that are attributed to you because of your roles are displayed there.

To see the service-based alerts that got assigned to you because of your roles:

1. Log in to the Portal with your user account.
2. Click the bell icon in the top right corner of the window.
3. Select the dashboard **My alerts**.

4. There you have the list of the alerts that you added to your Portal or that were automatically added because they are mandatory for at least one of your roles. You can see if the alert is linked to a role in the last column.

Related tasks

- Sending email notifications from the Appliance (Configuration)
- Creating a service-based alert
- Creating an investigation-based alert
- Defining user roles

Related concepts

- Alert

Creating a service-based alert

Service-based alerts rely on the thresholds of the services that you define in the Finder. The prerequisites for a service to accept the creation of an alert are thus the following:

- The service must be active (i.e. not disabled).
- The service must specify thresholds in its definition.
- There is no other service-based alert defined for the same service.

To create a service-based alert in the Portal:

1. Log in to the Portal as admin.
2. In the **ADMINISTRATION** top menu, under **CONTENT MANAGEMENT**, select the **Alerts** dashboard.
3. Click the plus sign on the top right corner of the **Alerts** dashboard. The dialog to configure a new service-based alert pops up.
 1. Type in the **Name** of the alert.
 2. Optional: Describe the function of the alert briefly in the **Description** field.
 3. Assign the alert to a category or create a new category for alerts in the **Category** section.
 4. In the **Service** list selection, pick the service that you want to use as a base for the alert. Only the active services that specify thresholds in their definition and that are not used in any other alert

are available. The dialog displays the thresholds configured for the service as well as the box **ignore until x or more devices are impacted**, since they are used to determine when to trigger the alert.

5. If the service defines two thresholds, choose the one that triggers the alert in **Trigger alert when x threshold is crossed**:
 - ◇ Select **yellow** to trigger the alert when the ratio of the devices with an issue over the devices using the service exceeds the warning threshold.
 - ◇ Select **red** to trigger the alert when the ratio of the devices with an issue over the devices using the service exceeds the error threshold.
6. Tick the **Send email when alert is triggered** check box to send the details of the alert to its subscribed users via email.
4. Click **Create** and your new alert is added to the list of alerts under administration.

The alert is now created, but it is not enabled until users subscribe to it. Users may subscribe individually to an alert if the alert is part of their roles. If a role includes an alert as mandatory, all the users who have that role assigned are automatically subscribed to the alert.

To subscribe to a service-based alert:

1. Log in to the Portal with your user account.
2. Click the bell icon in the top right corner of the Portal.
3. Click the entry **My alerts** from the drop-down list.
4. Click the plus sign icon followed by the label **Subscribe to alert** in the top right of the special dashboard **My alerts**.
5. Select one or more alerts from the list **Alert**. The list shows only the alerts that are not yet added to your Portal. If there are no alerts that you can subscribe to, the dialog displays an appropriate message and finishes.
6. Optional: Filter the list of alerts by role with the help of the drop-down list **Filter by role**. This is useful for locating a particular alert in the list when many alerts are defined for your roles.
7. Click the button **OK**. The alert is added to the special dashboard **My alerts**.

Related tasks

- Receiving Alerts from the Portal
- Adding users
- Creating a service

Creating an investigation-based alert

To define your own alerts on any kind of object, use investigation-based alerts. As their name indicates, investigation-based alerts express their triggering condition in the form of an investigation. The Engine periodically executes the investigations associated to alerts, depending on the frequency specified for each alert.

You can create an investigation-based alert by either:

- Defining the alert from scratch.
- Using an existing investigation as starting point.

The dialog to create alerts in the Finder is very similar to the dialog for designing investigations. There are a couple differences though:

- Investigations associated to an alert must be based on objects. You cannot associate an investigation based on activities or events to an alert.
- The time frame of the investigation depends on the frequency of the alert. You cannot specify a different time frame in the dialog for designing the investigation.
- An additional section at the end of the dialog for alerts let you specify the criticality, the frequency and the action to take when the alert is triggered.

Creating an investigation-based alert from scratch

To create an investigation-based alert from scratch:

1. Log in to the Finder.
2. Go to the **Settings** section in the accordion.
3. Inside the **Settings** panel, click the drop down list **Section** and either:
 - ♦ Select **Global alerts** to create an alert visible to every user. This option is only available if your account has the right privileges to create global alerts. Currently, every user can see global alerts only in the timeline of the Device view. The investigation of the alert must therefore be based on devices for the alert to be visible in the Finder.
 - ♦ Select **My alerts** to create an alert that is visible to you only in the Finder. This option is available to every user. If the alert is based on devices, it is displayed in the timeline of the Device view.
4. Right-click the area of the section and select **Create new alert** or type **Ctrl+N**. The dialog for designing a new alert shows up.

5. Enter a name for the alert by replacing the default **Untitled alert x** at the top of the dialog.
6. Optional: Type a brief description of the alert below the name.
7. Edit the investigation part of the alert as any other investigation, with the restrictions that you must retrieve an object and not an activity or an event and that you do not have to specify the time frame. After specifying the attributes to display, you reach the specific **ALERT** section.
8. Set the level of the alert in the drop down list **Criticality**:
 - ◆ Select **Normal** for non-critical alerts. Normal alerts based on devices are displayed in yellow in the timeline of the Device view.
 - ◆ Select **High** for critical alerts. Critical alerts based on devices are displayed in red in the timeline of the Device view.
9. Specify the frequency with which the system will check the conditions to trigger the alert:
 - ◆ Select **Immediate** for the system to check the conditions that trigger the alert almost continuously (every 30 seconds). Due to the nature of immediate alerts, you cannot select many display attributes that are usually available for the investigation. The Finder warns you when you select a display attribute that is incompatible with immediate alerts by showing a red cross to the right of the **Immediate** keyword. Hover the mouse cursor over the red cross to see the list of incompatible attributes that you selected. To avoid flooding, one hour must elapse between two consecutive immediate alerts for the same object.
 - ◆ Select **Hourly** to evaluate the conditions 15 minutes after the end of every hour and get the results for the whole hour that just passed by.
 - ◆ Select **Daily** to evaluate the conditions 15 minutes after midnight every day and return the results for the past day.
 - ◆ Select **Weekly** to evaluate the conditions 15 minutes after midnight every Monday and send the results of the investigation for the last week.
10. Choose the action to take when the alert is triggered. Note that the results of an investigation-based alert are limited to a maximum of 250 objects with 15 attributes per object. Results exceeding these values are cut down to avoid sending too much data.
 - ◆ Check **Send syslog** to write the results of the investigation associated to the alert to the system log of the Appliance. This option is only supported by global alerts.
 - ◆ Check **Send e-mail** to send the results of the alert by email to selected recipients. Note that any recipient can receive both global and non-global alerts.

11. Click **Save & Preview** to save the new alert and run the associated investigation.

Creating an alert from an existing investigation

To create an investigation-based alert from an existing investigation:

1. Log in to the Finder.
2. Find the desired investigation in the accordion.
3. Right-click the name of the investigation and select **Add to My alerts...** (or **Add to Global alerts...** if you have the right privileges). The dialog to design the alert shows up with the data of the investigation prefilled, so you only need to fill the **ALERT** section.
4. Optional: Change the name of the alert at the top of the dialog. By default, the alert borrows its name from the investigation.
5. Optional: Change the description of the alert that you find below the name. By default, the description of the alert is inherited from the investigation as well.
6. Set the criticality, frequency and actions of the alert as described above.
7. Optional: Modify the investigation settings to meet your needs. The original investigation is not modified.
8. Click **Save & Preview** to save the alert and execute the associated investigation.

Limit on the number of alerts

Using the methods described above, you can create and enable up to a maximum of 50 global alerts. For their part, users can create (or receive from their roles) and enable up to 10 additional local alerts per account on each Engine.

The total number of enabled alerts in an Engine, including the global alerts and the local or role-based alerts of all users, is limited to 150 alerts.

Timing considerations for immediate alerts on Application Library fields

Because fetching data from the Application Library is not instantaneous, immediate alerts that rely on the value of Application Library fields may fail to fulfill their purpose.

For instance, an immediate alert that is configured to detect the appearance of new binaries with high threat level in the system will fail to detect the first execution of such binaries. In effect, soon after the execution of a new binary and its creation in the Engine, the conditions of the immediate alert are evaluated. At that point, however, the Engine has not updated the threat level of the binary yet. It is only around 5 minutes later that the Engine connects to the Application Library and sets the value for the threat level.

Thus, the best practice in this case is to create an hourly alert in addition to the immediate alert and with the same conditions of the immediate alert. In this way, the first execution of a binary with a high threat level may escape the immediate alert, but will be detected by the hourly alert at the end of the current hour interval, once the Engine has set the correct threat level of the binary.

Related tasks

- Integrating alerts
- Connecting to Nexthink Cloud Services

Building web-based dashboards

Introducing dashboards in the Portal

Overview

Build web-based dashboards in the Portal to readily present the information about your IT infrastructure that interests you the most. Use indicators, tables, line charts, and bar charts to display real-time and historical data related to activities, objects, license use, and IT services within your corporate network.

Thanks to the web-based dashboards of the Portal, quickly find out those parts of your IT infrastructure that are working well and those that may require your attention. Hardware and software issues, network issues, or service degradation are all instantly apparent in properly designed Portal dashboards. The historical data stored in the Portal make dashboards an ideal tool to follow the progress of *change and transformation* projects as well.

Types of dashboards

Dashboards are grouped into modules. The type of a module determines the type of dashboards that it can hold. There are three types of modules:

Basic module

A basic module holds dashboards based on metrics. The dashboards of basic modules are highly customizable, letting you mix different metrics and distinct visualizations (*widgets*) for those metrics in the same dashboard. Combine carefully the widgets and metrics in a dashboard to get the most complete and consistent view possible on a topic of interest.

Service module

A service module holds dashboards based on services. Each dashboard in a service module represents a different service. Because service-based dashboards are not as customizable as metric-based dashboards, they are even easier to create. A predefined set of widgets gives you a variety of points of view on the service: the number of devices using the service, the network traffic generated, the issues related to connectivity or software malfunction, etc.

Software metering module

A software metering module holds dashboards that assess the use of software licenses within your company. Each dashboard presents you with

information about the number of installations and actual use of a particular application, relating this information with the number of licenses available.

Dashboards and hierarchies

Take advantage of your existing hierarchies when working with dashboards. The Portal automatically compute dashboard results for every level in your hierarchies. When viewing a dashboard, choose a node of the current hierarchy from the hierarchy navigation tool on the ribbon to display information that only applies to that particular node:

A well-designed hierarchy can greatly help you locate the source of problems within your network.

Creating a dashboard

Overview

Create a new dashboard in the Portal whenever you wish to get a visual representation of a part or an aspect of your IT infrastructure. Dashboards summarize a lot of data and complex investigations in a single web page; therefore, they are most useful when built with a consistent set of graphical items.

Administrators can create content in the Portal and share it with other users in the form of dashboards. Normal users whose profile lets them create personal dashboards can equally create their own content. In addition, normal users whose profile lets them publish dashboards can also share their dashboards with other users.

Portal dashboards are organized into modules, which act as containers for dashboards. The type of module determines the kind of dashboards that the module may hold.

Creating modules and dashboards

To create a dashboard, you need a module first. Depending on the kind of content that you want to add to the dashboard, the module must be of an appropriate type.

Follow the instructions in the corresponding section of the documentation to create or reuse a module of a particular type and add a new dashboard to it:

- To follow the evolution of *metrics*, choose a module of type **Basic**.
- To follow the evolution of *services*, choose a module of type **Service Monitoring**.
- To keep the usage of *software licenses* inside your company under control, choose a module of type **Software Metering**.

Related concepts

- Dashboard

Examining metrics in depth

Overview

Hover the mouse cursor over a figure in a widget to get the value of the represented metric with improved precision and, in some cases, with extra information inside a pop-up speech bubble. In the same speech bubble, find up to two options for examining the metric in depth:

- **Show details** (count metrics only)
- **Investigate in Finder**

When displaying the values of count metrics in a dashboard, you may be interested in knowing which objects are behind the numbers. For any figure in a dashboard that represents a count metric, it is possible to see the objects that added up to the total. The list of these objects is called the *details* of the metric. If you selected grouping options or additional display fields when creating your metric, they are displayed in the details.

If you want to go deeper in your investigations and you are a user with access to the Finder, you have the possibility to drill-down directly from the widget representation of a metric to the Finder. This works with any kind of metric and

not just with count metrics. When drilling-down from the Portal to the Finder, the Finder generates and runs a new investigation on the fly that keeps as much contextual information as possible to recreate the same conditions found in the Portal. Once in the Finder, you can modify the generated investigation, find relations with other objects and activities, and analyze the results with the full potential of the Finder.

Getting improved precision

Hover the mouse cursor over the displayed values in a KPI, table, or bar chart widget to bring up a speech bubble that displays the same value, but with higher precision whenever possible. In addition, get the variation with respect to its previous value, even if the widget was configured not to display the variation. For instance, the speech bubble below displays an additional decimal digit and the variation of the following KPI widget:

In the case of values that represent a count ratio, the speech bubble displays not only the ratio with higher precision, but also the number of involved objects. For example, in the bar chart below, see the precise ratio of Windows laptops, the variation of the ratio with respect to its previous value, and the number of Windows laptops with respect to the total number of devices in the speech bubble:

Displaying the details of count metrics

To display the details of a count metric from a dashboard:

1. Find in the dashboard a widget (KPI, table, or bar chart) that displays the metric. It is not possible to show the details of a metric from a line chart.
2. Hover the mouse over the figure. A small speech bubble shows up.

3. Click **Show details** in the bubble. The list of the objects that added up to the count metric is displayed.
4. Optional: If the list of objects is too big to fit in the dialog, use the controls at the bottom left of the dialog to move through pages. There is a limit to the maximum number of objects that the Portal can display in the details. If the limit is reached, the dialog indicates it with an appropriate message.
5. Optional: Click the arrow link with the label **CSV** at the bottom right corner of the window to export the list as a CSV file.
6. Click **OK** once you are done to close the dialog and return to the dashboard.

In table widgets, if you have grouped the results of the count metric, you get exactly the list of objects that correspond to the clicked figure, according to the grouping options.

Note that service dashboards do not have the option to display details. Showing the details is only available in dashboards based on metrics.

By default, the Portal stores the details of count metrics for the current day, week, month, and quarter. To keep historical data for previous periods, reserve additional disk space for the Portal.

Drilling-down to the Finder

As a prerequisite for drilling-down to the Finder, ensure that the address of the Portal is correctly configured, in the same way as needed for updating the Finder.

To drill-down from the Portal to the Finder for investigating the results of a metric in detail:

1. Log in to the Portal as a user with Finder access from a machine that has the Finder properly installed.
2. Locate in a dashboard the KPI, table, or bar chart widget that holds the value of the metric that you want to analyze with the help of the Finder. Line charts do not support drilling-down to the Finder.
3. Hover the mouse over the value of interest. A speech bubble shows up.
 - ◆ Alternatively, display the details of a count metric and hover the mouse over the name of an object in the list. A speech bubble also appears.
4. Click the link **Investigate in Finder** inside the bubble.
 - ◆ If the value of the metric in the Portal is aggregated from several Engines, a dialog displays the list of related Engines.

1. Click the name of the Engine to which the Finder must connect.
5. The Finder opens a view related to the metric, based on the contextual information taken from the Portal.

The contextual information partially determines what you see in the Finder when it opens. The context includes:

- The type of metric that you are drilling-down from.
- The *group by* options of the metric.
- The current node of the hierarchy that you are viewing in the Portal.
- The current time frame in the Portal.

These are the results that you get in the Finder, according to the type of metric from which you drill-down:

- When drilling-down from the value of a count metric (total or ratio), the Finder runs an investigation that returns the list of counted objects.
- When drilling-down from the value of an quantity metric, the Finder runs an investigation on devices and displays the results in a list. The list is ordered by the aggregate that defines the quantity.
- When drilling-down from the name of an object in either a top metric or in the details of a count metric, the behavior of the Finder depends on the kind of object selected:
 - ♦ If the object is a device, the Finder opens its device view.
 - ♦ If the object is a user, the Finder opens its user view.
 - ♦ For any other kind of object, the Finder runs an investigation that returns a list with the object as the single result.

The results in the Finder include the context of the Portal:

Group options

If any *group by* options are specified in the definition of the metric, the Finder includes them in the results of the executed investigation (not applicable to drilling-down to the device or user views).

Hierarchy

The results of the investigations are limited to the node of the hierarchy viewed in the Portal: they are transformed into a condition on a set of entities (not applicable to the device or user views).

Date

The date of the investigation run in the Finder matches the date viewed in the Portal (the time frame must be bounded to one single day, see limits below).

Limits of drilling-down to the Finder

In some cases, when drilling-down from the Portal to the Finder, it is possible to get inconsistencies between the results displayed in the Finder and the values observed in the Portal. To deal with those situations, the Finder highlights a warning message in the top of the view that explains the reason for the possible differences.

No data available

The data retention period in the Engine is shorter than in the Portal. When drilling-down to the Finder, the Engine may have already lost the history of the data requested if the date selected in the Portal lies too far in the past. The Finder displays the following message:

There is no data or only partial data in the Engine for the time frame you have requested (*time frame*). The oldest historical data in the Engine dates from *date*.

Data from multiple Engines

This case occurs when the values displayed in the Portal are aggregated from multiple Engines. When displaying the results in the Finder, only the results of one Engine are available; therefore, inconsistencies may arise. The Finder displays the warning:

The data displayed in the Portal is taken from multiple Engines. The results shown below are only for the *Engine name* Engine.

Time frame longer than one day

If the selected time frame in the Portal is longer than one day (that is, the time frame is a week, a month, or a quarter), the Finder ignores the time frame and displays the results for today. The reason lies in that the Portal aggregates the results of every metric daily, adding them up to make the total for the time frame selected. That is different from launching an investigation that aggregates results over the whole time frame. The warning message in the Finder reads:

The weekly/monthly/quarterly data displayed in the Portal cannot be retrieved in the Finder. The results shown below are for today.

Not historicized data

Some metrics are based on values for which the Engine does not keep a historical record. The Portal, however, keeps track of what their value actually was when it made its daily compute. This is in fact similar to the limitation that you get when you want to compute metrics for dates in the past. The Finder displays the message:

The data displayed below may differ from what is available in Portal as object properties and categories may have changed since the metric was computed.

Time frame of count metrics that count all objects

When viewing a widget that displays a count metric that counts all objects (active and inactive), the time frame in the the Portal must be of one day (choosing any other time frame disables the widget). If you drill down to the Finder, the time frame in the Finder changes to the full available period for better reflecting the results seen in the Portal. The Finder warns you about the change of time frame with the following message:

Time frame changed: these results go from the oldest available date in the Engine (*date*) to now (*date*).

Related tasks

- Computing dashboard data

Related references

- Types of widgets
- Data retention

Documenting dashboards

Overview

Write relevant information for users of your dashboards within the dashboards themselves. For instance, write dashboard descriptions to explain how to:

- Interpret the dashboard content.
- Incorporate the dashboard into a business process.
- Configure Nexthink to get proper results in the dashboard.

Only dashboards of basic modules, that is, those which are built with widgets based on metrics, may have accompanying descriptions.

Reading descriptions

The description of a dashboard is found on the right-hand side of the dashboard itself as a retractable panel. When retracted, the panel is reduced to a small

square, similar to a page marker, with the typical letter **i** (for *info*) inside.

Click the letter **i** to expand the panel and read the description of the dashboard. Once you are done reading, you can retract the panel back by clicking either the **i** letter again, or the cross sign at the top right of the description panel.

If the dashboard has no description yet, neither a panel nor an **i** letter are visible in the dashboard. Edit the dashboard to add a description as explained below.

Writing a description

To write a description of a dashboard:

1. Log in to the Portal as a user with the right to create dashboards.
2. Open or create a dashboard based on metrics.
3. Click the menu icon on the right-hand side of the dark blue ribbon.
4. Under the **DASHBOARD** section, select **Edit content**.
5. Click the letter **i** on the top right-hand side of the dashboard to expand the dashboard description panel. Even if the **i** was not visible during reading mode (because of the absence of a description), it becomes visible when editing the dashboard.
6. Write a description for the dashboard inside the description panel. To format the output, see the section below.
7. Click the check mark that appears on the right-hand side of the blue ribbon (in place of the menu icon) to finish editing the dashboard.

A description can contain up to 10 000 characters.

Formatting your description

To make your dashboard descriptions visually more appealing, give them format according to the following markup rules:

Formatting	Type in	Rendered
Normal text	Plain text	Plain text
Title	# Big title	<i>Big title</i>
Italic text	*Text in italics*	<i>Text in italics</i>
Bold text	**Text in bold**	Text in bold
Bulleted list	* item1 * item2	• item1

		• item2
Numbered list	1. item1 1. item2	1. item1 2. item2
Hyperlink	[Link text] (http://example.com)	Link text
Referenced hyperlink	[Link text][1] ... [1]: http://example.com	Link text
Horizontal rule	Text before --- Text after	Text before <hr/> Text after

Blocks

To emphasize specific parts of your description, four types of predefined blocks are at your disposal: **configuration**, **process**, **tip**, and **warning**.

To define a block, enclose your text between two lines of a three-colon sequence (:::), where the first sequence is followed by the name of the block, which must be one of the four predefined block types. See below an example of a block, where you have to substitute **[block_name]** for one name of the predefined block types:

```
::: [block_name]
Text within the block.
:::
```

The content of a block is highlighted from normal text, being displayed in the dashboard description panel inside a box with lighter background. The name of the block, preceded by an associated icon, appears as the title of the block:

Type in	Rendered
::: configuration :::	
::: process :::	

<pre> ::: tip ::: </pre>	
<pre> ::: warning ::: </pre>	

Related tasks

- Creating a dashboard
- Following the evolution of a metric

Assessing license use

Overview

Check the license compliance for a set of programs by measuring their real utilization. Compare the number of installed programs with the number of purchased licenses and save license costs for programs that are rarely used or, in the opposite case, allow some budget to purchase additional licenses.

To examine the use of program licenses within your company, first create a *software metering metric* for each product of interest. Then add to your dashboards one or more software metering modules to display the previously defined metrics. Software metering modules help you automate the assessment of software licenses across your hierarchies.

Applies to platforms:

Creating software metering metrics

Unlike other types of metrics, which are created in the Finder, software metering metrics are created in the Portal.

To create a software metering metric:

1. Log in to the Portal as administrator.
2. From the top drop-down menu **ADMINISTRATION**, select **Software metering metrics** under **CONTENT MANAGEMENT**.
3. Click the plus sign at the top right corner of the dashboard. The dialog to create a new software metering metric shows up.
4. Type in a name for the metric under **Name**.
5. Optional: Describe the purpose of the metric under **Description**.

6. Click **Continue**. The dialog starts retrieving the list of available programs from the Engines linked to the Portal while displaying the message **Fetching programs from Engines**. Once the process is finished, the dialog is ready for you to configure the programs to assess in the **Programs** table.
7. Type the number of days in the past to measure program utilization in the box **for the last x days**. By default, the number of days is 30.
8. Click **Add program** to define how to audit a new program. A new dialog appears to let you set up the program and its licenses.
 1. Step 1 of adding a program: identify the program
 1. Type a name to identify the program in the entry **Program label**. Each program name must be unique in the metric.
 2. Optional: Enter a web address with additional information about the program in the entry **Add an external link to program's informations**. Use, for instance, the official web page of the vendor of the program.
 3. Choose how to select in the next step the packages that correspond to the program:
 - To pick the packages from a list proposed by the widget, select **from the list**.
 - To provide an investigation which will pick the packages for you, select **corresponding to investigation**.
 2. Step 2 of adding a program: identify the packages
 1. Depending on the mode that you selected in step 1 to select the packages, you either:
 - Choose one or more packages from a list. If you have a mixed installation of Windows and Mac OS devices, the list displays packages from both platforms. Only installed packages are shown in the list:
 1. Filter the list of packages by choosing the initial letter of their name in **Available packages starting by**.
 2. Click the name of one package (or more while maintaining the **Ctrl** key pressed).
 3. Click **Add to selected** to add the packages to your selection.
 4. Optional: By default, all versions of the packages selected are accounted for license use. If needed, choose a specific version of a package by clicking the pencil icon to the left of the name of the package in the list of selected packages. Type in a particular version in the

Edit package version dialog and click **Edit**.

5. Click **Next** to move on to step 3 of the wizard.

- Alternatively, use an investigation on packages:

1. Export an investigation based on packages from the Finder to the clipboard.

2. Paste the investigation in the text area under the label **Paste an investigation**.

3. In the **Investigation** root element of the XML, manually change the value of the attribute **SyntaxVersion** from 3 to 2 (i.e. the attribute should read **SyntaxVersion="2"**).

4. Click **Next** to move on to step 3 of the wizard.

3. Step 3 of adding a program: configure licenses

1. Choose the level of the hierarchy through which you want to distribute the licenses in **Set licenses on level**. Remember that all the licenses must be distributed over nodes at the same level in the hierarchy. Therefore, the chosen level remains valid for all the licenses that you add to the metric.

2. Click **Add license** to configure a new program license. The dialog for adding a license pops up.

3. Type a name to identify the license in the entry under the label **Reference**.

4. Optional: Type a brief description of the license in the text area **Description**.

5. Optional: Type in the number of purchased licenses in the entry labeled **Quantity**. Leave it blank if you have an unlimited set of licenses.

6. Optional: Enter a date in the section **Expires on** to specify the time validity of your license. If you leave it blank, the license never expires.

7. Optional: Pick a node to which licenses apply from the **Node** list. The nodes that you can pick belong to the hierarchy level that you selected before.

8. Click **Add** to finish adding the license.

4. Step 4 of adding a program: specify binaries

1. For the metric to compute statistics about program usage, you may pick a set of binaries which are related to the selected packages in the **Significant binaries** section:

- Choose **no significant binary** to avoid this step altogether.

- Choose **based on a list** to create a list of binaries. If you have a mixed installation of both Windows and Mac OS devices, you can add binaries of both

platforms to the list.

1. Click **Add binary** to add a new binary to the list. A dialog for adding binaries appears.
2. Type in the exact name of the executable program in **Executable name** (include the .exe extension if it is a Windows binary).
3. Type the first numbers of the versions of the executable that will be taken into account for computing the statistics in **Version starts with**.
4. Click **Add** to add the defined binary to the list.
- Choose **based on investigation** to automatically select significant binaries from an investigation.
 1. Export an investigation based on binaries from the Finder to the clipboard. The investigation can be multi-platform, but it cannot have a time frame of *During last x minutes (hours or days are valid)*.
 2. Paste the investigation to the text area below the option **based on investigation**.
 3. In the **Investigation** root element of the XML, manually change the value of the attribute **SyntaxVersion** from 3 to 2 (i.e. the attribute should read **SyntaxVersion="2"**).
2. Set the threshold of daily average usage of binaries per device to consider only those binaries that a device actually executes for a significant amount of time. Select the number of hours and minutes in **Under usage threshold**.
3. Click **Add** to end the wizard for adding programs.
9. Click **Finish** to end up the edition of the metric.

Right after adding a new software metering metric, optionally launch its computation in the Portal. Otherwise, software metering metrics are processed during the nightly computation of the Portal, as any other metric.

Creating a Software Metering Module

Once you have defined one or more software metering metrics, create a software metering module to visualize them in your dashboards.

To create a software metering module:

1. Log in to the Portal as administrator.

2. Click the menu icon in the right-hand side of the blue ribbon at the top of the window.
3. Select the option **Create new module....** The dialog to create a new module shows up.
4. Choose **Software Metering** as the type for your new module.
5. Select the software metering metrics that you wish to include in your module from the list of **AVAILABLE** metrics. Use the **Ctrl** or **Maj** keys while you click with your mouse to select more than one metric at the same time.
6. Click the left-to-right arrow to move the selected metrics from the **AVAILABLE** list to the **SELECTED** list.
7. Click the button **Create** to generate a new software metering module with the selected metrics.

Interpreting the results of software metering metrics

To view the results of a software metering metric, navigate to the software metering module that holds the metric (the module must lie within your available dashboards) and select the dashboard that displays the metric of interest.

The dashboard of the metric shows a table with license compliance info that can be broken down both by software product and by hierarchy level. By default, the table displays an overall view of all the products located at the hierarchy level that corresponds to the view domain of the user (for instance, if the user is a central administrator, that would be the root level). In the overall view, you get the following columns in each row of the table:

- A circular bullet whose color indicates the status of license usage:
 - ◆ **Green** when the number of installations of a particular software product are equal or lower to the number of valid licenses for that level of the hierarchy.
 - ◆ **Red** when the number of installations of a software product exceed the number of valid licenses available for the level.
 - ◆ **Gray** is only used when you are viewing the results on nodes that do not belong to the hierarchy levels that you chose in the configuration of the licenses. For instance, if you configure the number of valid licenses only at the root (top) level of the hierarchy, you will see gray circles when you break down the results to lower levels, because you have not provided license information specifically for those levels.
- The name that identifies the program in the column **Product**.
- The number of devices that have the program installed in the column **Installed**.

- The number of valid licenses for the program in the column **Nb valid licenses**.
- The ratio of devices with the program installed divided by the number of valid licenses in the column **Compliance**.
- The number of devices that underutilize their licensed products in the column **Nb devices in under usage**.

Navigating through products and hierarchy levels

If you configured more than one software product in your metric, you can select one of the products to see the license information broken down by hierarchy. At the top left of the dashboard, choose a program from the **Product** drop-down list to show the results for that particular product. Alternatively, hover the mouse cursor over the name of the product in the table and click the down arrow that appears to its right.

The **Product** column of table is then replaced by a column with the name of current hierarchy level, filled with the names of the nodes in that level. For each hierarchy node, the table displays the license status of the product with the color codes seen above.

To go down in the hierarchy, use the drop down list for the hierarchy that is placed to the right of the **Product** drop-down list. Select one node from the drop-down and the table will adapt to present the results in the chosen hierarchy level. Alternatively, you can hover the mouse cursor over the name of a node in the table and click the down arrow that appears to its right.

To go back to a higher level in the hierarchy, there is a navigation tool over the **Product** and hierarchy drop-down lists that displays the whole hierarchy path of the current level. Click the name of the desired level/node in this navigation tool to move to that level and update the table results accordingly. If you defined more than one hierarchy, the navigation tool also lets you choose the particular hierarchy according to which you want to assess the use of licenses. To switch from one hierarchy to another, click the name of the current hierarchy in the navigation tool and select the hierarchy to replace it from the drop-down list. Of course, if you just defined one hierarchy, no drop-down list shows up.

Analyzing binary usage

If you configured the metric to survey significant binaries, you can get detailed information about the use of those binaries on the monitored devices. Only those devices with the appropriate packages installed that have been active for the last 30 days are included in the computation.

To get statistics about binary usage:

1. Hover the mouse cursor over the results of the **Installed** column. A small info icon appears to the right of the value.
2. Click the info icon and the widget displays the **Details** table, with the list of devices where the program is installed.
3. Select the type of information that you want to see from the list in the top left of the widget above the table. Choose between:
 - ◆ **Binary usage**: for each configured significant binary, the table shows whether the binary was actually used in the device or not.
 - ◆ **Statistics**: the table displays the cumulated average usage and the average executions per day of all significant binaries, as well as the time of their last activity for each one of the involved devices. For devices that did not execute any significant binary during the last 30 days, the average usage displays:
 - ◇ **0 ms**, if the device was active during the last day (the day of the computation).
 - ◇ - (a dash sign), if the device was not active during the last day.
4. Optional: Click the link **CSV** at the bottom right corner of the **Details** dialog to export the results to a CSV file.

The info icon also shows up in the column of underused devices. As expected, clicking on it gives you the details of underused devices only.

Computing dashboard data

Overview

The Portal employs different techniques to retrieve and process dashboard data from the connected Engines. Depending on their mode of activation, the techniques for computing dashboard data can be classified as automatic or manual:

- Automatic computations
 - ◆ Computation after metric creation
 - ◆ Nightly computation
 - ◆ Real-time computation (live service data)
- Manual computations
 - ◆ Launch computation of metric from the Finder
 - ◆ Launch computation of software metering metric from the Portal

In this article learn about the different types of computations, how to manually launch the computation of metrics from either the Finder or the Portal, and how to track the status of computations from the Portal. Learn also about the consequences of computing the values of a metric for dates in the past.

Automatic computations

Computation after metric creation

When you create and save a new metric in the Finder, it automatically launches its computation for the past day. A newly created metric has thus one day of history unless you clear the metric manually.

Nightly computation

The nightly computation is the main source of dashboard information for the Portal. It usually takes place during the night (hence its name), when the load on the Portal and on the Engines is lower because users seldom connect at night. At the scheduled time, the Portal starts collecting data related to the enabled services and metrics in all connected Engines and processes it for being displayed in the dashboards. This automatic computation is scheduled by default at 1 am local time, but this time is configurable.

Thus, every night the Portal collects the accumulated data in the Engines for the past day. In the case of metrics whose history is missing for some consecutive days including the last day, the Portal computes not only the past day, but the number of days configured.

Real-time computation

When displaying live data in service dashboards, the Portal continuously receives data from the Engines about the current status of the services. This automatic computation requires no special configuration.

Manual computations

Manual computations do not replace the nightly computation of the Portal, but are a complement to it. Manually computing a metric is specially useful when you create or modify a metric. Instead of waiting for the nightly computation to get the results for the metric, you can manually trigger its computation to see how it looks in the Portal.

Metric computation from the Finder

To trigger the computation of metrics from the Finder:

1. Right-click the name of a metric or of a folder holding metrics in the left-hand side accordion.
2. Select **Compute** from the context menu:
 - ◆ Choose **For the last day**, to compute or recompute the results of the metric (or metrics) for the day before.
 1. A dialog warns you that the operation will clear the history of the metric (or metrics) for the last day. Click **Yes** to proceed.
 2. A final dialog informs you that the computation of the metric (or metrics) will start shortly. Click **OK**.
 - ◆ Choose **Over the maximum available period**, to compute the metric (or metrics) for all the past days available in each connected Engine. This option is not available for count metrics that take into account all objects (active and inactive).
 1. A dialog warns you that the operation will clear **all the history** of the selected metric. If you selected a folder, the dialog lists all the metrics in the selected folder, indicating those count metrics in the folder that take into account all objects, which are ignored for the computation. Click **Yes** to proceed.
 2. A final dialog informs you that the computation of the metric (or metrics) will start shortly. Click **OK**.

Beware of the option to recompute a metric over the maximum available period. It really clears **all the history** of a metric. For instance, if you have three months of historical data in the Portal for a particular metric and two weeks of data available in the Engine, asking to recompute the metric over the full period will erase the three months of history and recreate only two weeks. Therefore, use this option with care and only when the modification of a metric completely invalidates previous results. Keep in mind as well that the recomputation of metrics for past days suffers from some limitations.

Software metering computation from the Portal

When creating a new software metering metric in the Portal, you are given the option to compute the metric for the last day immediately after its creation.

To compute software metering metrics that are already created:

1. Log in to the Portal as administrator.

2. Click the **ADMINISTRATION** item in the top menu.
3. Select **Software metering metrics** under **CONTENT MANAGEMENT** in the drop-down menu.
4. In the list of available metrics, tick the box in the first column for each metric that you want to compute. Alternatively, tick the box at the top of the list to select all available metrics at once.
5. Click the sprocket icon placed in the top right corner of the dashboard. A dialog informs you that the computation will be done for the last day.
6. Optional: Tick the box in the dialog to overwrite current data; that is, to clear and recompute as well those metrics that already have results for the last day.
7. Click **OK** to schedule the computation. Another dialog informs you that the computation is scheduled and that you can track it.
8. Click **Done**.

As any other metric, software metering metrics are also computed during the nightly computation of the Portal.

Tracking computations

To track the computations of metrics in the Portal:

1. Log in to the Portal as administrator.
2. Click the **ADMINISTRATION** item in the top menu.
3. Select **Computations** under **CONTENT MANAGEMENT** in the drop-down menu.

The **Computations** dashboard displays a table that includes both the manual and the automatic computations launched today by default. Filter the results by selecting other time frames and the user who initiated the computation from the two lists interleaved in the following sentence that you find at the top of the table:

- **Display computations for -time frame- created by -user.**

As time frame, choose among:

- **today** (default)
- **last week**
- **last month**

To filter by the user who launched the computation, choose among three options:

- **everybody** (default): list all the computations, no matter who initiated them.
- **myself**: list only the computations initiated by the current user.
- **Nexthink Portal**: list only the nightly computations initiated by the system itself.

The table shows the current status of each computation within the selected time frame, including (when applicable) its start time, end time, progress ratio, duration, and owner. The dashboard content is automatically refreshed every 10 seconds.

It is possible to stop a scheduled or running computation. Administrators can stop only those computations initiated by themselves. Central administrators, on the other hand, can stop computations initiated by anyone. If you stop a computation during its execution, the results for those dates that have been fully computed are saved and the rest are discarded.

The following table describes the possible computation statuses:

Icon	Name	Description
	Scheduled	When a computation is scheduled, the operation has been submitted and put in the queue. There is no ongoing computation yet; therefore, there are no details about progress, time, or duration, but the computation can be cancelled.
	Running	The computation is currently being performed. The start time and progress fields are available. The computation can be cancelled.
	Completed (success)	All metrics were computed without error. Details about the computed metrics are available.
	Completed (failure)	Errors occurred during the computation. Some metrics may have been computed correctly but at least one had an error. Details show the completed and failed computations of the metrics.

	Cancelled	The computation has been manually cancelled. Details may show some metrics that have been completely computed before cancellation (with or without errors), while other metrics have not been computed.
--	-----------	---

Display the details of all finished computations (either completed or cancelled) by clicking the information button.

Computation details

To get details about a specific computation in the **Computations** dashboard:

- Click the info button in the table, when the display mode is **all computations**.
- Click any point in the line chart (or the info button to the right of a value), when the display mode is **automatic computations**.

The details dialog displays the list of metrics with their individual computation status. You can only see the details of those metrics that fall into your administration domain.

In the **all computations** mode, the table displays individual computations, including manual computations and separate attempts of nightly computations. Therefore, when you open the computation details from this mode, the date is fixed and there is no selection of attempts. On the other hand, when opening the details of a nightly computation from the **automatic computations** mode, it is possible to select a different date for convenience and see a different. In addition,

ef multiple attempts were needed for the nightly computation of the selected date, you can also select the specific attempt. Thus, when opening the details dialog from the **automatic computations** mode, find these additional selectors:

- **Display details for** *-date*: to select a particular date.
- **This computation required multiple executions. Display** *-attempt*: to select a particular attempt.

In both modes, the dialog lets you choose the metrics for which you want to see the details, depending on the status of their computation:

- Completed (successfully or not, but not skipped)
- Skipped
- Successful
- Unsuccessful

For unsuccessfully computed metrics, click the cross icon to display the error that stopped the computation.

Widgets with no data

If the Portal has not computed data for a metric on a particular date, the widgets that display the value of the metric show it in different ways depending on their visual appearance:

- In KPI or table widgets, the value of a metric with no data is represented with a dash (-) character.
- In line charts, when hovering the mouse over the points where the metric was not computed display the message **no data computed**.

In the widgets of service dashboards, you find the same behavior when there is no data for the service. Similarly, in the overview dashboard of a service module, a dash character representing the status of a service indicates a lack of data for that particular service.

Note that you still may see live data in a service dashboard that received no data from nightly computations. For instance, the first day that you create a service, you may see live data for that service, but you will get no data for the service if you navigate to a previous date.

Computing metrics for dates in the past

There are a few cases where you may want to compute the value of metric for a date in the past:

- You create a new metric and you want to compute its value for some days before its creation.
- You modify a metric in such a way that it completely invalidates previously computed values. You may want to clear the history of the metric and recompute it for some days in the past.
- The nightly computation of the Portal failed for some reason at a specific date and you want to have the value of a metric for that date.

You can either trigger the computation of a metric manually from the Finder or wait for a new nightly computation of the Portal. In the case of newly created metrics, only manually triggering the computation will do, because they automatically compute the last day at the time of creation, preventing thus the Portal from going further to the past. Remember that the nightly computation automatically recomputes the metrics up to the configured number of past days, if the data of the metrics are missing for some consecutive days including the last day. The nightly computation will not automatically fill in the gaps though if the failed computations are placed between successful computations.

Regardless of computing the metrics manually or automatically, you must be aware of the limitations of computing metrics for dates in the past.

Metrics may include in their computation values of objects (or conditions on these values) for which the Engine does not keep a historical record. Examples of these values may be the free disk space or the antivirus status of a device. During the nightly computation, the Portal takes the current value of those fields to compute the metrics for that day. In this way, the Portal always uses the most recent values to compute the metrics for the last day.

When computing the metrics for past dates, however, the Portal does not know the values of these fields in the past, so it still uses their most recent values. Since these fields usually do not change quickly over time, the computed values are often a good approximation, but if the value has a significant change (for instance, the antivirus real-time protection goes from *on* to *off*) the value of the metric may be wrong.

Therefore, beware of computing metrics for dates in the past when the metrics depend on values of fields for which the Engine does not keep a historical record.

Related tasks

- Changing the nightly computation time of the Portal

Related references

- Widget Compute State In Charts

Reusing dashboard content

Overview

Reuse dashboard content by either:

- Copying and pasting widgets and dashboards (within basic modules).
- Publishing your modules.

All users with the right to create personal dashboards can copy and paste either individual widgets or full dashboards to accelerate the creation of new dashboards.

In addition, users with the right profile and administrators can publish Portal content for other users to share the same dashboards. The unit of sharing is the module; therefore, when a module is published, it is shared with all the dashboards and widgets that it holds.

Nevertheless, not all the users are able to see all the published modules. Users can see a particular published module only according to specific rules that depend on the users' type and their roles.

Copying and pasting a widget

Only Nexthink users with the right to create personal dashboards can copy and paste widgets. Moreover, users can only copy widgets from and paste widgets to dashboards in the **PERSONAL** section, because editing permissions are required to access the individual widgets and modify the dashboards. If a widget is not part of a personal dashboard, copy the dashboard to your **PERSONAL** section first, as indicated in the section about publishing a module below. Only those widgets that represent metrics, that is, all widgets in dashboards of the basic type of module, are copiable.

To copy a widget:

1. Log in to the Portal as administrator or as a normal user with the right to create personal dashboards.
2. Locate the dashboard that holds the desired widget with the help of the navigation tool on the left-hand side of the dark blue ribbon.
3. Click the menu icon on the right-hand side of the dark blue ribbon.
4. Select **Edit content** to set the dashboard into edit mode.
5. Hover the mouse cursor over the widget that you wish to copy.
6. Click the sprocket icon that shows up at the top-right corner of the widget.
7. Select **Copy**.

The copied widget is kept in the clipboard until the user logs out or copies another widget. Therefore, the copied widget can be pasted multiple times.

To paste a previously copied widget:

1. Log in to the Portal as administrator or as a normal user with the right to create personal dashboards.
2. Locate a dashboard in the **PERSONAL** section.
3. Click the menu icon on the right-hand side of the dark blue ribbon.
4. Select **Edit content** to set the dashboard into edit mode.
5. Hover the mouse cursor over the place where you want to add the widget.
A plus sign shows up in the available spots.
6. Click the plus sign to add a widget.
7. Click the link at the end of the sentence that is displayed at the top of the dialog to add a widget:
Please choose the type of widget you would like to add or **Paste widget**.

Copying and pasting a dashboard

Only Nextthink users with the right to create personal dashboards can copy dashboards of the basic type from any section (either from their own content or from published content) and paste them to modules within their **PERSONAL** section.

To copy a dashboard:

1. Log in to the Portal as administrator or as a normal user with the right to create personal dashboards.
2. Locate the desired dashboard with the help of the navigation tool on the left-hand side of the dark blue ribbon.

3. Click the menu icon on the right-hand side of the dark blue ribbon.
4. Select **Copy dashboard**.

The copied dashboard is kept in the clipboard until the user logs out or copies another dashboard. Therefore, the copied dashboard can be pasted multiple times.

To paste a previously copied dashboard:

1. Log in to the Portal as administrator or as a normal user with the right to create personal dashboards.
2. Locate the desired module where to paste the copied dashboard with the help of the navigation tool on the left-hand side of the dark blue ribbon.
3. Click the menu icon on the right-hand side of the dark blue ribbon.
4. Select **Paste dashboard**.
 - ◆ If a dashboard with the same name already exists in the module, the dashboard name gets the suffix *Copy*. Subsequent copies add an identifying number to the suffix, starting from number 2.
 - ◆ Note that this option is only enabled if there is a dashboard ready to be pasted in the clipboard.

Publishing a module

To publish a module, you need to be an administrator or a normal user whose profile explicitly allows the publication of dashboards. To publish a module, first create the module or modify an already published module by copying it to your personal content. To make a local copy of a published module in the **PERSONAL** section:

1. Log in to the Portal as administrator or as a normal user with the right to create and publish Portal content.
2. Click the module navigation tool on the left-hand side of the dark blue ribbon.
3. Open the published module. Either:
 - ◆ Click the name of a module under a section corresponding to one of your assigned roles in **My content**.
 - ◆ Click the name of a module in **All published**.
4. Click the menu icon on the right-hand side of the dark blue ribbon.
5. Select the option **Copy module to my content**.

Once you have created the new module or modified an existing module, publish it:

1. While viewing the module content in the Portal, click the menu icon on the right-hand side of the dark blue ribbon.
2. Under the section **MODULE**, select the option **Publish module....** A dialog shows up.
3. Publish your module as new or substitute an existing module. Choose between:
 - ♦ **Create new module.** Type in a name for the module that is not repeated. Note that, in the list of existing modules, normal users do not see the names of those modules that were created by administrators. Therefore, there can be a name conflict even if you do not see the name that you chose for your module in the list.
 - ♦ **Replace existing module.** Note that only administrators can replace the modules created by other administrators.
 1. Choose the module to replace from the list.
4. Optional: Tick the box **Keep a copy in my personal content** if you do not want the module to disappear from your **PERSONAL** section after publishing.
5. Click **Done** to actually publish the module.

After publishing the module, administrators can assign the module to a role.

Seeing a published module

Normal users see only those published modules that are included in their roles inside the **My content** group, or modules that were created by other users within the **All published** group.

For their part, administrators see, in addition to the published modules included in their roles, the modules published by themselves and by other administrators or normal users in the **All published** group.

Find the summary of permissions and access rights for Portal users.

Related tasks

- Adding users

Related references

- Access rights and permissions

Importing and Exporting authored content

Methods for reusing authored content

Overview

There are some situations when you want to share with others the content that you have created either in the Finder or in the Portal. Within a normal setup (that is, an installation with one Portal and several connected Engines) you can create some type of content that is instantly shared by other users of the system. This is the case, for instance, of *centralized* content created with the Finder or of content associated to roles.

Manual import/export in the Finder

Other kind of content that you create in the Finder is local to each Engine. To share local content, first connect to the Engine holding the elements to share with the Finder and export them manually, then connect to the target Engine to import them back. This manual method is also valid for exporting local or centralized content to Engines in different setups, as long as dependencies are respected. For instance, to export an investigation that uses in its conditions a category that is not present in the target Engine, you must import the category beforehand.

Publishing modules

The Portal has its own way of sharing content among different users: administrators can publish the modules that they create, and optionally assign them to roles, so that other users can profit from them. This method however does not let you take content from one Portal and import it into another Portal. This latter scenario typically occurs when you have separate pre-production and production environments.

Content packs

To export a coherent set of Finder and Portal elements at the same time, use content packs. A content pack groups in a single file multiple items of Finder content and Portal modules. For instance, recalling our dependency example above, you can include in a same content pack an investigation and the categories on which the investigation depends. Similarly, when exporting Portal modules, you definitely want to include in the same content pack the metrics and

the services on which the modules depend. See in the following articles how to import content packs or make your own content packs.

Nextthink Library

Ready-made content packs are available from the Nextthink Library. The packs in the Library are designed to provide you with customizable solutions for different IT areas: security, transformation, compliance, etc. The method to import a library pack is different from importing a content pack from a file.

Related tasks

- Reusing dashboard content
- Adding users (Defining user roles)
- Manually sharing Finder content

Related references

- Local and shared content
- Nextthink Library

Manually sharing Finder content

Overview

The Finder is able to export content either to the clipboard or to a file in XML format, regardless of the content being local to an Engine or centralized. The XML can be later imported into another system, again with the help of the Finder.

We exclude from this discussion the content that is local to the Finder itself (sessions and custom actions) or non-sharable (tags). See the reference article on local and shared content content for more information.

Exporting content

To share an item, such as an investigation:

1. Log in to the Finder.
2. Locate the investigation on the left-hand side accordion panel of the Finder.
3. Right-click the name of the investigation and select **Export**:

- ◆ Click **Investigation to clipboard** to copy the contents of the investigation to the clipboard.
- ◆ Click **Investigation to file...** to save the contents of the investigation to an XML file.

If you have organized your investigations in folders, right-click the name of a folder and select **Export** to actually export all the investigations in the folder. To export all the investigations that are visible from the Finder at once, right-click the header of the section **Investigations** (or in the empty area within the section) and select the **Export** option. Again, when exporting multiple investigations, choose whether you want to export the **Investigations to clipboard** or the **Investigations to file....**

Importing content

To import the previously exported investigation into another Engine:

1. Open a Finder session in the other Engine.
2. Right-click the header of the **Investigations** section or in the empty area of the section and select **Import** from the menu.
3. Depending on how you exported the investigations:
 - ◆ Select **Investigation from clipboard**, if you copied the contents of the investigation to the clipboard.
 - ◆ Select **Investigation from file...**, if you saved the contents of the investigation to an XML file.

If any of the imported investigation already exists in the destination Engine, the Finder displays a dialog for resolving the conflicts.

Other sharable content

The method described above using investigations as an example works with other types of content as well. Apply the same method to export and import:

- Investigations
- Investigation-based alerts
- One-click investigations
- Categories
- Metrics
- Services
- Campaigns
- Scores
- Remote actions

Note that you can manually export and import not only content that is local to an Engine, but also centralized content. Although centralized content is shared by definition, its sharing capabilities are limited to the set of Engines managed by one Portal. Therefore, you may still want to export centralized content to Engines managed by different Portals. For example, this is specially useful for bringing content from a pre-production environment to a production environment.

Related tasks

- Conflict resolution

Related references

- Local and shared content

Importing a content pack

Importing a pack from the Library

Via the Nexthink Library website

Install a Library Pack on your setup while browsing the catalog of the Nexthink Library. This procedure requires that you have the Finder installed on your computer:

1. Open your favorite web browser.
2. Navigate to the Nexthink Library website.
3. Optional: combine the search tools to filter and sort the list of Library packs.
4. Select a Library pack from the list.
5. Click the **Install** button to the right of the description of the pack. A dialog shows up.
 1. Read attentively the modules required by the Library pack.
 2. Tick the box **I understand** to acknowledge the content of the dialog.
 3. Click **Install**. The Finder opens automatically.
6. Log in to the Finder, if not already logged in.
7. Finish the installation of the pack from the Finder by configuring the pack, if needed.

Via the Finder

Within the Finder, directly import content from the Nexthink Library as follows:

1. Log in to the Finder.
2. In the left-hand side accordion, select **Nexthink Library**
3. Look for the pack that you want to import:
 - ◆ In the **Browse** tab, look for the pack in a list that you can refine by type of solution and sort using different criteria.
 - ◆ In the **Search** tab, type in some keywords to find the appropriate pack.
4. Click the plus icon to the right of the name of the pack to import it.
5. Resolve conflicts, if any.

Importing a pack from a file

To import a content pack from a file:

1. Log in to the Finder.
2. Click the sprocket icon in the top right corner of the Finder window.
3. Select **Import content pack from file...** from the menu.
4. Choose the content pack file.
5. Review the content to be imported. Each item is imported into an appropriate place depending on its type:
 - ◆ Centralized content (categories, metrics, and services) are shared among Engines in the same setup.
 - ◆ Alerts are locally imported into the **global alerts** section.
 - ◆ Investigations and one-clicks are locally imported into their corresponding sections.
 - ◆ Portal modules are imported into your **PERSONAL** section in the Portal.
6. Resolve conflicts, if any.
7. When the message **Pack <name> successfully imported** shows up, click **OK** to finish the import process.

Sharing the imported modules

When importing a content pack, note that the imported Portal modules stay in your **PERSONAL** section of the Portal. Remember to publish these modules if you want to share them with the rest of users.

Related tasks

- Conflict resolution
- Reusing dashboard content

Related references

- Local and shared content

Conflict resolution

Overview

When importing content, either manually, from the Library, or from a content pack, it is possible that the content to be imported conflicts with content that is already present in the target setup. Investigations, alerts, modules, etc. with the same name as those found among the imported items may already exist in the same locations where the items are copied. Learn here how to deal with these conflicts and resolve them.

Options for resolving conflicts

The resolution of conflicts is a step in the import process, regardless of the method used for importing content. A dialog displays a list of the offending items along with a choice of possible correction actions to remediate the different issues.

The options available for resolving conflicts depend on the type of content to be imported. The default option proposed by the system to resolve a conflict follows this conservative rule:

- Existing content must not be broken.

See below the list of all possible actions with an explanation of their meaning. Keep in mind that not all of the options are eligible for any kind of item to be imported and that they may have different effects on different kinds of items:

Skip

The item will not be imported.

Replace

The item to be imported takes the place of the conflicting existing item.

Import a copy and Keep both

Both options keep a copy of both the existing item and the item to be imported by automatically renaming this latter, adding a numerical index to its name. The system displays one option or the other for selection depending on whether renaming the item to be imported may break dependencies or not.

Merge

The item to be imported is combined with the existing item to produce an item that is a mixture of both.

Conflicts in metrics

While all other items are identified solely by their name, metrics can also be identified by their *unique identifier* (UID). Although not visible to the user, the system can verify the UIDs of two conflicting metrics that have the same name to determine if they are different versions of the same original metric (same UIDs) or if they are actually two completely different metrics that happen to have the same name (different UIDs).

The options for resolving conflicts between metrics are different when they have the same UID and when they have distinct UIDs.

Choosing the appropriate option

Skip

When the system detects that the item to be imported has exactly the same definition as the existing item with which it is in conflict, it automatically skips the import, notifying the user by displaying the action **Skip (identical)** in the list of conflicting items.

For the cases when the content is not identical, choose **Skip** if you want to preserve the definition of the existing item.

Replace

Replacing the definition of an existing item by that of an imported item has consequences that depend on the type of item imported:

Metrics

Replacing a metric clears the history of the metric in the Portal only if its new definition introduces major changes (for instance, it has different grouping options). You can only replace a metric by another when both have the same UID.

Services

Replacing a service resets its corresponding Service View in the Finder and its previously computed baselines. Nevertheless, the history of the service is preserved in the Portal.

Categories

Replacing a category triggers an automatic retagging of objects when the auto-tagging rules change. In addition, all manual tags are lost.

Import a copy and Keep both

When keeping a copy of both the existing item and the imported item, the latter is renamed with an index. For instance, if you have an item called **Compliance** and you import another item with the same name (therefore causing a conflict with the existing item), the system assigns it the name **Compliance (2)** to indicate that it is a different version of the item.

There are two formulations for this option depending on the type of the items to be imported. The system suggests either one of:

Import a copy, for categories and services

Since conditions in investigations and metrics may depend on the names of categories and services, renaming the imported categories and services may result in broken conditions.

Keep both, for investigations, folders of investigations, alerts, one-clicks, metrics with different UIDs, and modules

The system proposes this formulation for all other types of items that may be renamed safely. The case of modules is special in the sense that keeping both the existing and the imported modules is the only available action for resolving the conflict; therefore, the system automatically takes this action without prompting the user.

Merge

Two types of items may merge with others of their kind:

- Folders (of investigations, alerts, or one-clicks)
- Categories

When folders are merged, the resulting folder combines the contents of the existing and the imported folders. If individual items inside the folder conflict, you must resolve those conflicts separately.

When merging two categories, all the keywords that are present in the category to be imported and not in the existing category are added to the resulting category. The result of the merge, understood as a set operation, is thus the union of the keywords in the two categories.

To determine what to do with the keywords that are present both in the category to be imported and in the existing category, but defining different auto-tagging rules, the merge operation comes in two flavors:

Merge - preserving auto-tagging rules

To preserve the auto-tagging rules of the keyword in the existing category against the auto-tagging rules defined by the keyword to be imported.

Merge - replacing auto-tagging rules

To substitute the auto-tagging rules of the keyword in the existing category for the auto-tagging rules of the keyword to be imported.

For categories, it is recommended to choose first **Import a copy** before committing on a **Merge**. Decide later whether you really want to merge the two categories. If so, re-import and choose **Merge**. After merging, check the order of the keywords in the resulting category.

Conflict resolution table

Find below a table with all the types of objects and the available conflict resolution options for each one of them:

Items	Conflict	Skip	Replace	Import a copy	Keep both	Merge	Merge - preserving auto-tagging rules	Merge - replacing auto-tagging rules
Investigations / alerts / one-clicks	Same name (in same folder)	X			Default			
Folder of investigations / alerts / one-clicks	Same name (in same folder)	X	X		Default	X		
Metrics	Same name (globally) and same UID	Default	X					
	Same name (globally) and same UID				Default			

Services	Same name (globally)	Default	X	X				
Categories	Same name (globally)	X	X	X			Default	X
Modules	Same name (in PERSONAL section)				Default			

Exporting a content pack

Overview

In a previous article, we have already seen how to manually share content generated in the Finder. To complete the picture, learn here how to manually export content from the Portal and, optionally, how to create a content pack for combining both Finder and Portal content into a single file.

The operations described in this article should only be performed by a Nexthink Engineer or a Nexthink Certified Partner.

If you need help or assistance, please contact your Nexthink Certified Partner.

Exporting modules from the Portal

There is currently no graphical interface to exporting modules from the Portal; therefore you need to use the command line. The export process takes the specified module and produces an XML file as output. This file can later be imported into another setup from the Finder in the same way as you would import any other content pack. Note that the tool only lists and exports published modules by default.

To export a published module:

1. Log in to the CLI of the appliance hosting the Portal.
2. Go to the directory holding the export script:

```
cd /var/nexthink/portal/rsquery
```
3. Optional: Display the usage of the export script:

```
./exportLibraryV6.py --help
```
4. Optional: List all the published modules available for export:

```
./exportLibraryV6.py --list
```
5. Export a module to an XML file:

```
./exportLibraryV6.py --export --moduleName="name" \
--exportfilename="/tmp/module.xml"
```

You can export more than one module at a time by specifying several modules by their name or their UIDs in the export command.

Creating a single content pack file

With the methods for exporting modules from the Portal and other content from the Finder, you have now at your disposal the tools to export any kind of content from one setup to another. However, for your convenience, you may want to put all the content that you want to export from both the Portal and the Finder in a single content pack file.

To that end, edit your own content pack XML file and add all the content to it. If you want to include Portal modules, it is recommended to reuse the XML file generated by the Portal as a starting point to build your own content pack. In any case, the XML of a content pack has the following structure:

```
<?xml version="1.0" encoding="UTF-8"?>
<Pack Description="brief description" Name="name" SyntaxVersion="1">
  <Contents>
    <Content Type="type of content">

      </Content>

    ...
  </Contents>
</Pack>
```

Inside each **Content** tag, copy the Finder or Portal content that you have previously exported as XML files, after removing their own XML header. In the case of Finder content, the XML must contain a tree of objects and not a single object. To export a tree of objects from the Finder, group the content in a folder and export the folder. Alternatively, export a tree of objects by exporting a whole section: right-click the header of a section in the accordion and select **Export**. For instance, to export all your categories at once from the Finder, right-click the header of the **Categories** section in the left-hand side accordion and select **Export**.

Indicate the type of content that you have copied inside the **Content** tag by setting the value of the **Type** attribute. See below the list of all possible values for the **Type** attribute:

PortalModuleCollection

Indicates that the content are modules exported from the Portal.

investigations-tree

Indicates that the content is a set of investigations.

services-tree

Indicates that the content is a set of services.

metrics-tree

Indicates that the content is a set of metrics.

fields-tree

Indicates that the content is a set of categories.

ObjectOneClicksTree

Replace *Object* by the actual name of an object to indicate that the content is a set of one-clicks relative to that object. *Object* may be one of User, Device, Package, Application, Executable, Binary, Port, Destination, Domain, Printer, Execution, Connection, or Web_request.

scores-tree

Indicates that the content is a set of scores.

publications-tree

Indicates that the content is a set of campaigns.

actions-tree

Indicates that the content is a set of remote actions.

Note that the **Settings** section in the accordion of the Finder is special when exporting content, because it exports only the content that it is displaying at the moment, be it alerts, or one-clicks.

One-clicks in particular are also special on their own. When exporting all the one-clicks at once and not only the one-clicks associated to one kind of object, the Finder actually generates a content pack (not a simple tree) which itself includes all the trees of one-clicks associated to each kind of object. Keep this in mind when editing your own content pack.

Related tasks

- Manually sharing Finder content
- Importing a content pack