

Nextthink V6.4

Installation and Configuration

Generated: 12/21/2019 9:10 am

Table of Contents

Planning your installation.....	1
Hardware requirements.....	1
Connectivity requirements.....	6
Data retention.....	11
Using Collector on the Internet.....	14
Installing Nexthink Components.....	19
Installing the Appliance.....	19
Installing the Collector on Windows.....	23
Installing the Collector on Mac OS.....	42
Installing the Mobile Bridge.....	46
Installing the Finder.....	52
Customer satisfaction program.....	54
Updating from V6.x.....	57
Updating the Appliance.....	57
Updating the Collector.....	59
Updating the Finder.....	65
Configuration.....	67
Setting up a software license.....	67
Specifying your internal networks and domains.....	69
Allocating resources for the Portal.....	71
Connecting the Portal to the Engines.....	73
Centralized Management of Appliances and Engines.....	74
Adding users.....	77
Hierarchizing your infrastructure.....	85
Setting the locale in the Portal.....	95
Changing the Time Zone of the Portal.....	99
Time Zones and data collection.....	100
Nightly task schedules timetable.....	103
Changing the data collection time of the Portal.....	104
Establishing a privacy policy.....	106
Security settings in the Appliance.....	119
Importing and replacing Certificates.....	121
Managing Appliance accounts.....	124
Sending email notifications from the Appliance.....	126
Controlling session timeouts in the Portal.....	127
Preventing password saving in the Finder.....	128

Table of Contents

Configuration

Establishing a data retention policy in the Engine.....	130
Special operation modes for the Engine and the Portal.....	131
Ignoring specific print ports.....	136
Enabling and Disabling the Engine Application Library Access.....	137
Importing data from Active Directory.....	138
Configuring the system log.....	142
Reporting the URL of HTTP web requests.....	147
Mobile Bridge configuration settings.....	151
Collector MSI parameters reference table.....	152
Nxtcfg - Collector configuration tool.....	157
Auditing logon events.....	159
Redirecting Collector traffic.....	160
Support for DirectAccess.....	163
Changing the thresholds of High CPU warnings.....	165

Maintenance.....167

Logging in to the CLI.....	167
Planning for disaster recovery.....	167
Web Console backup and restore.....	171
Engine backup and restore.....	172
Portal backup and restore.....	176
License backup and restore.....	179
Finding out unlicensed devices.....	180
Removing devices.....	182
Examining the logs in the Portal.....	183
Storing Engine data in a secondary disk drive.....	185
MSI Exec Returns 3010.....	186
Package Executable Mapping.....	187
Installing third-party software in the Appliance.....	188
Installing VMware Tools in the Appliance.....	190

Planning your installation

Hardware requirements

Nextthink Appliance

The Appliance consists of a Linux-based 64-bit operating system and all the packages needed to run one of the server components of Nextthink: the Portal or the Engine. The Portal and the Engine must be installed in separate physical or virtual machines, except for some specific cases (see below). When installed in virtual machines, the requirements are equivalent to those of a dedicated physical server.

Nextthink officially supports VMware and Hyper-V as virtualization platforms, but both the Portal and the Engine may run on any other virtualization platform of your choice. In all cases, your server must be powered by a 64-bit compatible processor (AMD64 or Intel 64 -not Itanium- architecture). The vast majority of AMD and Intel processors currently available in the market comply with this requirement.

Beware that some versions of popular virtualization platforms may impose particular limits on the number of CPUs and amount of RAM that you can assign to a virtual machine. In installations with many devices, the possible maximum values may not reach the specified requirements. Likewise, in virtualized environments with high load, the performance of IO operations may not be sufficient for the Portal or the Engine to write to disk normally. In case of doubt, please contact Nextthink Customer Success Services to validate your virtualized setup.

Definitions and remarks

Complexity

For Nextthink Portal appliance sizing a metric called "complexity" is defined; this metric is an indication on the amount of computation required for widgets:

`complexity = entities * hierarchies * (max_levels + 2)`

◇ entities is the total number of entities across all Engines;

◇ hierarchies is the total number of hierarchies;

◇ max_levels is the number of user-defined hierarchy levels for the hierarchy which has the largest number of levels (this does not take into account the root level or the entity level).

CPU cores

This value indicates the number of physical or virtual CPU cores that the appliance requires. In case of virtual environments these cores must be completely dedicated to the Nexthink appliance. The performance of each core must be equivalent to that of an Intel Xeon 2Ghz.

Memory

This value indicates the amount of RAM. Minimal requirement for all configurations is DDR3-1600 with data rate of 1600 MT/s.

Nexthink Portal

Max devices	Max complexity	Memory	Disk	Details (90 days)	CPU cores	Network
500	500	6 GB	40 GB	8 GB	2	100 Mbps
5k	500	8 GB	60 GB	60 GB	2	100 Mbps
10k	500	8 GB	100 GB	120 GB	4	100 Mbps
20k	1 000	12 GB	200 GB	220 GB	4	100 Mbps
50k	3 000	16 GB	300 GB	450 GB	6	100 Mbps
100k	10 000	32 GB	600 GB	700 GB	6	1 Gbps
150k	15 000	48 GB	1 TB	1 TB	8	1 Gbps
>150k	Ask	Ask	Ask	Ask	Ask	Ask

- The Portal requires at least 10 MB/s of disk throughput.
- The total maximum number of entities across all Engines is limited to 2 000.
- The maximum number of enabled metrics is limited to 500. If you define more than 500 metrics, those in excess are disabled (not computed).

The quantities in the **Details** column correspond approximately to the additional disk space required to store 90 days of historical details of count metrics. Add the value in the **Disk** column to the value in the **Details** column to get the total disk space required. For more information, see the article about data retention in the Portal.

Configure the Portal to make the most out of your hardware resources.

Nexthink Engine

			Memory	Disk		Network	
--	--	--	---------------	-------------	--	----------------	--

Max Events	Max devices / with Web & Cloud	Max entities			CPU cores		Disk throughput
20M	500 / 500	20	4 GB	80 GB	2	100 Mbps	5 MB/s
50M	3k / 2k	100	8 GB	100 GB	4	100 Mbps	10 MB/s
50M	5k / 3-4k	100	8 GB	120 GB	4	1 Gbps recommended	15 MB/s
100M	10k / 6-8k	100	12 GB	200 GB	8	1 Gbps recommended	25 MB/s

- The maximum number of supported devices for each configuration depends on the amount of web activity and the Web & Cloud privacy configuration.
- If you install the Collector in servers, take into account for the sizing of the Engine that a single server is roughly equivalent to 20 normal devices.
- The indicated number of cores include 20 simultaneous Finder users. If more than 20 users access Nextthink Engine simultaneously, 1 additional core is required for each 5 users (up to a maximum of 24 cores).
- Tests under controlled conditions have demonstrated that the Engine is capable of dealing with up to 100 *normalized* Finder users when run on a 24 cores appliance (20 users for the first 8 cores + 16 cores * 5 users per core).
- A *normalized* user is characterized for querying the Engine once every 25 seconds with a query that takes 10% of a core dedicated to the Engine. If Finder users deviate too much from this behavior, the number of supported users may vary drastically. Note as well that any other kind of query to the Engine (such as queries to the Web API) reduces the number of supported users.
- The maximum number of entities per Engine is described in the table above.
- The maximum number of supported mobile devices for all Engine configurations is 5 000.

Running multiple Engines on the same Appliance

Multiple Engine instances can run on the same physical or virtual appliance. In order to size the hardware, you should sum the memory, disk, cores and disk throughput required by each individual Engine.

Running Nexthink with a single appliance

For very small installations the Portal and the Engine can run on the same physical and virtual appliance.

Max devices	1 000
Max complexity	500
Events	20M
Memory	12 GB
Disk capacity	120 GB
Disk write speed	10 MB/s
CPU cores	4
Network	100 Mbps

External backups

The disk space requirements given for the Appliance already take into account the amount of space needed to keep up to ten internal backups of either the Portal or the Engine.

In the case that you activate external backups, Nexthink recommends you to reserve the following quantities of external storage, depending on the size of your setup. The figures indicate the file size for each individual backup.

Nexthink Portal

The backup size for the Portal depends on the number of devices, the complexity, the amount of history and the number of widgets and reports. We recommend regularly monitoring the used capacity and adapting it based on actual needs.

Max devices	External backup size
5k	3 GB
10k	5 GB
20k	10 GB
50k	15 GB
100k	30 GB
150k	50 GB

Nexthink Engine

The disk requirements for the backup of the Engine are more predictable than those of the Portal and only depend on the number of events stored in the Engine.

Max events	External backup size	Network throughput
20M	2 GB	5 MB/s
50M	4 GB	15 MB/s
100M	8 GB	25 MB/s

Mobile Bridge

To collect information from mobile devices synchronized via ActiveSync with Microsoft Exchange, the Mobile Bridge uses a Remote PowerShell connection to your Exchange Server.

Install the Mobile Bridge on a dedicated Windows Server 2008 R2 or later. The hardware requirements for the Mobile Bridge are those same ones recommended by Microsoft for installing their operating system. The Mobile Bridge is compatible with Exchange 2010 SP2 or 2013.

Nexthink Collector

	Without Web & Cloud	With Web & Cloud
Disk	35 MB	
Network card	Any, wireless or wired	
Network bandwidth	100-150 bps	150-250 bps

Nexthink Finder

Starting from Nexthink V6.3, the Finder supports high DPI screens. When setting DPI scaling in Windows, the Finder adapts its size properly.

Memory	4 GB system memory, at least 2 GB available
Disk capacity	50 MB
CPU	2 cores, 2 GHz
Network	100 Mbps recommended

Certified Hardware List

Nextthink V6 appliances include a Linux-based operating system that is derived from the freely distributed sources of a major North American Enterprise Linux vendor. This vendor maintains a list of supported hardware that has been tested and is certified to work with its Linux distribution. To help you choose your hardware for your appliances (the Portal and one or more Engines), verify that it is in the following list:

- Certified Hardware List (Red Hat link)

Related tasks

- Planning for disaster recovery
- Allocating resources for the Portal

Related references

- Server support
- Hardware requirements - Installing Windows Server 2008 (Microsoft link)
- Exchange 2013 System Requirements (Microsoft link)

Connectivity requirements

Find the connectivity requirements of every Nextthink product in the reference tables below. You can configure some of the products to use either a secure or a non secure channel for specific services (see the column **Reason**). Depending on their configuration, note that you may require to allow connections through a different port number. For each connection, the tables also indicate the transport protocol used. When an application protocol handles the connection over the transport layer, the name of the application protocol precedes the name of the transport protocol.

Engine

In the following table, we describe the different ports that must be open on the Engine appliance to communicate seamlessly with the other Nextthink components and with standard network services.

Port	Protocol	Direction	Reason	Domains
------	----------	-----------	--------	---------

Number		(IN/OUT)		
22	SSH / TCP	IN	Secure shell connection to the CLI	
25	SMTP / TCP	OUT	Mail server for notifications	
53	DNS / UDP	OUT	Resolving destination names by reverse IP	
80	HTTP / TCP	OUT	Connection to automatic updates (non secure)	updates v6.nexthink.com updates centos v6.nexthink.com
98	HTTP / TCP	IN	Download for Updaters (non secure)	
99	HTTPS / TCP	IN	Download for Updaters (secure)	
	HTTPS / TCP	IN	Administration through the Web Console	
123	NTP / UDP	OUT	Time synchronization	0.centos.pool.ntp.org 1.centos.pool.ntp.org 2.centos.pool.ntp.org
389	LDAP / TCP	OUT	Connection to Active Directory (non secure)	
443	HTTPS / TCP	OUT	Connection to the Application Library	application library v5.nexthink.com application library v6.nexthink.com
	HTTPS / TCP	OUT	Connection to automatic updates (secure)	updates v6.nexthink.com updates centos v6.nexthink.com
636	LDAPs / TCP	OUT	Connection to Active Directory (secure)	
999	UDP	IN	Traffic from Collector	
	TCP	IN	User connection from the Finder	

			or the Portal	
1671	HTTPS / TCP	IN	Access to the Web API	
7000 7001 7002	TCP	OUT	Communication channels with the Portal	
8888	HTTP / TCP	IN	Connection from Updater (non secure)	
	HTTPS / TCP	IN	Connection from Updater (secure)	
11031	HTTPS / TCP	OUT	Communication with the Mobile Bridge	

Portal

In the following table, we describe the different ports that must be open in the Portal appliance to communicate seamlessly with the other Nextthink components.

Port Number	Protocol	Direction (IN/OUT)	Reason	Domains
22	SSH / TCP	IN	Secure shell connection to the CLI	
25	SMTP / TCP	OUT	Mail server for notifications	
53	DNS / UDP	OUT	Lookup name of AD servers	
80	HTTP / TCP	IN	Access to the Portal (non secure)	
	HTTP / TCP	OUT	Connection for automatic updates (non secure)	updates v6.nextthink.com updates centos v6.nextthink.com
	HTTP / TCP	OUT	Connection to the Library	library.nextthink.com
88	TCP / UDP	OUT	Kerberos authentication	

			of AD users	
99	HTTPS / TCP	IN	Administration through the Web Console	
	HTTPS / TCP	OUT	Centralized administration of the Engine	
123	NTP / UDP	OUT	Time synchronization	0.centos.pool.ntp.org 1.centos.pool.ntp.org 2.centos.pool.ntp.org
389	LDAP / TCP	OUT	Connection to Active Directory (non secure)	
443	HTTPS / TCP	IN	Access to the Portal (secure)	
	HTTPS / TCP	IN	User connection from the Finder	
	HTTPS / TCP	OUT	Connection to the Online License mechanism	license.nextthink.com
	HTTPS / TCP	OUT	Connection to automatic updates (secure)	updates v6.nextthink.com updates centos v6.nextthink.com
636	LDAPs / TCP	OUT	Connection to Active Directory (secure)	
999	TCP	OUT	Connection to the Engine	
7000 7001 7002	TCP	IN	Communication channels with the Engine	
8100	HTTP / TCP	OUT	Send license information to Local License Manager	

Local License Manager

The Local License Manager resides in the same machine as the Portal.

Port	Protocol	Direction	Reason
------	----------	-----------	--------

Number		(IN/OUT)	
8100	HTTP / TCP	IN	Get license information from the Portal

Mobile Bridge

The Mobile Bridge needs to connect to the Exchange CAS to get mobile information. In turn, it offers a REST interface for the Engine to use to retrieve the collected information.

Port Number	Protocol	Direction (IN/OUT)	Reason
80	HTTP / TCP	OUT	Communication with Exchange (non secure)
443	HTTPS / TCP	OUT	Communication with Exchange (secure)
11031	HTTP / TCP	IN	REST interface for the Engine

Finder

In the following table, we describe the different ports that must be opened on the computers running the Finder to communicate seamlessly with the other Nexthink components.

Port Number	Protocol	Direction (IN/OUT)	Reason	Domains
80	HTTP / TCP	OUT	Connection to the Library	library.nexthink.com
	HTTP / TCP	OUT	Connection for the notification of new updates	doc.nexthink.com
	HTTP / TCP	OUT	Verification of security certificates	ocsp.verisign.com
	HTTP / TCP	OUT	Optional: Feedback for the customer experience program	report.nexthink.com
443	WebSocket / TCP	OUT	User connection to the Portal	
	HTTPS / TCP	OUT	Connection to download new updates	download.nexthink.com
	HTTPS / TCP	OUT	Connection to the customer improvement program site	finder analytics.nexthink.com

999	TCP	OUT	User connection to the Engine	
-----	-----	-----	-------------------------------	--

Collector

In the following table, we describe the different ports that must be opened on the computers running the Nexthink Collector to send data seamlessly with the Nexthink Engine.

Port Number	Protocol	Direction (IN/OUT)	Reason
999	UDP	OUT	Connection to the Nexthink Engine

Updater

If you install the Updater in order to deploy and update the Collector, here is the list of different network ports that need to be open in the computers running the Collector to communicate seamlessly with the Engine.

Port Number	Protocol	Direction (IN/OUT)	Reason
98	HTTP / TCP	OUT	Download from the Engine (non secure)
99	HTTPS / TCP	OUT	Download from the Engine (secure)
8888	HTTP / TCP	OUT	Connection to the Engine (non secure)
	HTTPS / TCP	OUT	Connection to the Engine (secure)

Data retention

Data retention in the Portal

Nexthink is able to keep historical data for several years in the database of the Portal. The Portal consolidates the data collected from the Engines and keeps them in permanent storage. The Portal does not store all the individual events, but the results of widget computation.

History	Number of devices
Several years (view by periods of 2 years max)	Up to 150 000 (consult for bigger setups)

Keeping historical details of count metrics

In the case of count metrics, the Portal stores the total number of objects satisfying a particular set of conditions. The Portal keeps these numbers of counted objects as regular historical data. In addition, for every count metric, the Portal stores the list of objects that contributed to the metric for the current day, week, month, and quarter. The list of objects that contributed to a count metric, along with their selected set of display fields, are known as the *details* of the metric. To see the details of a count metric, open a dashboard with KPI or table widgets representing the values of that count metric in the Portal, hover the mouse cursor over a particular value and select **Show details**.

When additional disk space is allocated, the Portal can store the details of count metrics not only for the current period (aggregated object lists for the current day, week, month, and quarter), but also for previous days, weeks, months, and quarters. To keep historical details of count metrics, make sure that you reserve some disk space for that purpose in the Portal appliance:

1. Log in to the Web Console as admin from a web browser:
`https://<appliance_address>:99`
2. Open the tab **Portal & Finder** and select **Parameters**.
3. In the **Portal parameters** sub-section, specify the quantity of disk space that you want to dedicate to the storage of history details for count metrics.

The number of days of stored historical details depend on the amount of disk space reserved, the number of enabled count metrics, the number of display fields included in each metric, and the actual metric data collected each day. When the disk space dedicated to store the details of count metrics is exhausted, the Portal launches a cleanup process that deletes one or a few full days of historical details, starting from the oldest day in the saved history.

Find orientative figures for setting the disk space for historical details in the hardware requirements of the Portal.

Data retention in the Engine

In its turn, the Engine stores real-time data with a greater level of detail than the Portal. All events are kept in volatile memory until the configured maximum is reached. You can browse all the data in the Engine down to the event level with the help of the Finder.

We present in this section several tables that contain an estimation of the number of weeks that correspond to a maximum number of events stored the

Engine. Once the maximum is reached, new events replace the oldest ones, which have already been consolidated in the Portal before being dismissed.

Please note that the amount of data required for mobile devices is negligible.

Without the Web Monitoring Feature

If the web monitoring feature is not enabled, the data retention periods are expressed as follow (in weeks):

	50M events	100M events
5 000 devices	2-3 weeks	4-5 weeks
10 000 devices	1-2 weeks	2-3 weeks

With Web Monitoring Feature but for Services Only

If only the services for web monitoring has been enabled, the following data retention periods are expected (in weeks):

	50M events	100M events
5 000 devices	2 weeks	4 weeks
10 000 devices	1 week	2 weeks

With Full Web Monitoring Feature

If the web monitoring feature is fully enabled, the following data retention periods are expected (in weeks):

	50M events	100M events
5 000 devices	1-2 weeks	2-3 weeks
10 000 devices	1 week	1-2 weeks

Increasing data retention in the Engine

If you reckon that your current data retention period is rather short, increase it by trading off detailed information for more capacity. Gain up to 50% more history by setting a more aggressive aggregation policy in the Engine.

Related tasks

- Establishing a data retention policy in the Engine

Using Collector on the Internet

The Nexthink solution provides support for many different network architectures. The typical scenario is always focused on the assumption that the computers that are monitored via the Nexthink Collector and, optionally kept updated via the Nexthink Updater, are always part of a corporate network. There are situations in which it is still possible to collect data from computers located in remote networks connected via the Internet.

This guide aims at highlighting the different issues that might surface when designing a Nexthink infrastructure that is using the Internet as transport, and the best practices that allow Nexthink to work in such cases. Designs where VPNs are used, either for site-to-site or client remote access, fall outside of the scope of this document, since they typically allow traffic to flow without restrictions from the remote machines.

A typical Nexthink network architecture including remote sites with machines on which the Collector is installed can be visualized in the following way:

While the part of the Nexthink infrastructure that is located inside of the boundary of the corporate network behaves in a normal fashion, there are different considerations related to the architecture on the Internet side, and how Nexthink behaves in such conditions. There are considerations to address about:

1. Location of the Nexthink Engine on the network
2. Considerations about firewalls
3. Considerations about Network Address Translation (NAT)
4. Impact of data loss due to using the Internet as transport for traffic

Location of the Nexthink Engine on the network

Three different architectures can be implemented:

1. The Engine is located on the DMZ (Demilitarized Zone);
2. The Engine is located on the Internal network;
3. Two Engines are installed, one on the DMZ and the other on the Internal network.

Security requirements might dictate the location for the Nexthink Engine.

In case a two Engines architecture is required, it is mandatory to correctly configure Domain Name Service (DNS) in order to supply the correct IP address to the Collector to those laptop computers which will be connected at different times to both the internal network and via the Internet. In such situations the so-called *split DNS* scenario will help, as it will always provide the correct IP address of the Engine to use depending on the location of the laptop computer at any given time.

Currently, the Collector refreshes DNS information every 60 minutes. This means that activities performed right after changing location might not be recorded by the Engine.

Here is a general network diagram of possible Nexthink architectures:

Considerations about firewalls

A firewall is typically situated on the boundary of the corporate network, and is used to separate at a logical level the internal network from the external network,

and the DMZ. The DMZ is where the servers that must be accessible from the Internet are typically located.

The Nextthink Engine can be located on the DMZ or in the corporate network.

Depending on the actual location of the Nextthink Engine, there will be a need to configure the firewall to allow incoming traffic from the remote Collectors located on the Internet to cross the firewall and flow towards the Engine.

If the Engine is located on the DMZ, appropriate security considerations must be made, even though the Nextthink Appliance exposes only the required services for its operation.

For example, a suitable best practice would allow only the UDP traffic coming from the remote Collectors to pass through the firewall and reach the Nextthink Engine on the DMZ.

The traffic flows from the remote endpoints towards the Engine happen on the protocols and ports described in the following articles:

- * [Connectivity Requirements for Collector](#)
- * [Connectivity Requirements for Updater](#)

Considerations about NAT

In most network designs, somewhere on the network there will be a device performing NAT, in order to translate the IP addresses used on the internal network to external IP addresses valid globally on the Internet. Most home routers perform NAT. The DMZ itself can have either public addresses, or be already subject to NAT. There is no standard way to design a network, so it will be important to work with the network team when designing the Nextthink infrastructure architecture.

If NAT is performed, all the devices sitting on the opposite side of the NAT with respect to the Engine will be reported with the same IP address. This does not result in any limitation with the Nextthink solution except for the fact that information about the actual IP address of these devices will not be available.

TCP traffic originated by the Nextthink Updater and crossing a NAT device will not be affected, and Updater operations will work normally.

Impact of possible data loss

While most corporate networks seldom have problems with data loss, when using the Internet to transport traffic from remote PCs towards the Nextthink Engine, there sometimes can be data loss.

The Collector provides duplication of important data to address this specific issue.

Another factor to take into consideration is the possibility that the UDP frames used by the Collector could become fragmented while in transit to the Nextthink Engine. This could be due to a mismatch between the Maximum Segment Size (MSS) allowed on the network path, and the amount of data present in the UDP frame generated by the Collector.

If the MSS exceeds the limit, then IP fragmentation will occur, and the UDP frames will be fragmented in more pieces. If one or more pieces of the UDP frame are lost in transit, then the whole UDP packet will be lost.

To avoid this issue with fragmentation, Nextthink recommends that the Collector be configured to create UDP frame of smaller size, down to a minimum value of 1 000 bytes. The configuration can be changed in the Collector Advanced properties; the parameter name is **Maximum payload size (bytes)**. Alternatively, use the Collector configuration tool to change the **mss** value of an already installed Collector.

The Maximum payload size value is stored in the Windows registry in the following key:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\nxtrdrv\params
The parameter name is ?mss? and its type is REG_DWORD, as shown below.

Should the reconfiguration of the MSS parameter be needed on a large number of machines on which the Collector is installed, the use of a Group Policy Object (GPO) is the recommended way to proceed. You should contact Nextthink Support in case more details are needed.

The MSS value can also be set when installing the Collector for the first time. The Collector MSI Parameters Reference Table contains detailed information about this topic.

Installing Nexthink Components

Installing the Appliance

Installing the Appliance on a Physical Server

If you are installing the Appliance on a physical server, you are required to enter the BIOS and modify some settings:

1. Power on the server and enter the BIOS setup. This is usually done by pressing down [F2] or [DEL] keys before the computer attempts to boot the operating system. Explore the exact method to enter BIOS setup in your user manual.
2. In the system settings, set the date and time of the system to match the current UTC (Coordinated Universal Time). Time precision is important to ensure consistency of data in the system. In order to keep time precision, you may configure the Appliance to use NTP servers as described in Network Parameters section below.
3. Insert your copy of the the Nexthink V6 Installation DVD into your DVD Drive.
4. Go to boot sequence (or boot order) settings and select the CD/DVD Drive to be the first device in the list of bootable devices. The system will boot from the Nexthink Installation DVD next time.
5. Exit BIOS setup saving your changes.

Installing the Appliance in Virtual Server

If you are installing the Appliance in a virtual machine, the exact installation steps will depend on your virtualization platform. Here, we assume that you are familiar with the creation and configuration of virtual machines on your virtualization platform. If this is not the case, please take your time to learn how to use it. Independently of the virtualization platform that you are using, you must perform the following operations:

1. Create a new virtual machine and select Linux 2.6 64-bit as the guest operating system.
2. Insert the Nexthink V6 Installation DVD into the DVD drive of the host machine or, alternatively, copy the ISO image of the Nexthink V6 Installation DVD to a file system accessible to the host machine.
3. Indicate your virtualization platform to use the DVD drive or the copied

ISO image for booting the your newly created virtual machine for the first time.

4. Start the virtual machine.
5. Explore the user manual of your virtualization platform to find out how to synchronize the clock of the virtual machine to the clock of the host, if the host has precise timing. Alternatively, use NTP servers as described in Network Parameters section below.

Finishing the Installation

When on of the previous steps have been completed, the installation is identical for both physical or virtual servers. If everything went well, your system will now boot from the Nextthink V6 Installation DVD and you should see the Nextthink V6 splash screen.

1. Press [ENTER] or wait for the 6 seconds timeout.
2. After the splash screen, the End-User License Agreement is displayed. Accept to proceed
3. In the following screen, you are warned that your hard disk will be erased during the installation procedure and that all data on it will be lost. Ensure that you do not have any valuable data on the hard disk and type Yes to proceed.
4. The installation procedure whether you want to configure the network. If you wish to configure the network, a small dialog box will open where you may change the default IP address, subnet mask and default gateway. If you choose not to configure your network, the Appliance will take these default values. These can be modified as described in Network Parameters section below.
5. Select the type of keyboard that is attached to your computer. By default, US keyboard is highlighted.
6. Choose a secure password for the root user, ensure that it is kept in a safe place or use another method to make sure that it is not forgotten.
7. The installation procedure is formating your hard disk and installing the necessary software packages.
8. Once the installation has finished, remove the Nextthink V6 Installation DVD from the DVD drive or detach the ISO image from your virtualization platform and reboot the system.

Appliance Configuration

Appliance Network Parameters

1. If nothing has been specified during the installation, the Appliance server default IP address is 192.168.0.99 with a netmask of 255.255.255.0. To access its web-based configuration page, configure your computer to have the static IP address on the same subnet as the Engine default IP address, that is, 192.168.0.0/24. If you modified the network configuration of the server during the installation of the Engine, use your modified configuration values instead of the defaults and then the following steps.
2. With a cross-over cable, connect the first Ethernet port on the server to the Network Interface Port (NIC) on the client computer. Alternatively, if there is a switch between the client PC and the server, connect both devices to it.
3. Test the network connectivity between the client computer and the server using ping command. Execute the ping test using the default IP address of the Engine: 192.168.0.99. If the ping test fails with the first Ethernet port, repeat the test by connecting to each available Ethernet port on the server. By default, the first Ethernet port is used for testing the network connections, but some servers may behave differently.
4. Using a web browser and type the following URL in the address field, using the HTTPS protocol: `https://192.168.0.99:99`. Note that if you just type the IP address 192.168.0.99:99, your web browser will use HTTP by default and you will get a blank page. Ensure that you specify the HTTPS protocol.
5. Accept the SSL certificate from Nexthink. Depending on the web browser, you may have a warning about the certificate. Accept it and you will arrive at the Console welcome screen:
6. Enter the Login Username as admin and Password also as admin to connect to the Appliance Web Console
7. Go to Appliance > Network parameters to enter or modify the Hostname, Domain name, DNS servers (IP addresses), Timezone and NTP time servers and then select Save.
8. Go to **Appliance > Network** interfaces, click on Edit and configure the IP address, the subnet mask and the default gateway as required (Note: Nexthink does not recommend DHCP).
9. Review the desired network settings before clicking on the Save button. Clicking on this button will save the settings, and the current connection will be lost. To reconnect to the web-console, use the Engine`s newly configured IP address (ensure you have accordingly modified the client computer network parameters to match the Appliance subnet).

Appliance Proxy configuration

Software installation/updates and access to the Nextthink Library are done using HTTP and HTTPS access. If those connections need to go through a proxy, go to **Appliance > Proxy**.

1. Check the Enable the proxy configuration.
2. Configure the server name or IP address and port number.
3. If required, set a username and password.
4. Choose the authentication method (Note: For software updates, only basic authentication method is supported).
5. To check that the Proxy chain is working and that the supplied parameters are correct, we recommend to test them by clicking on the Test button.
6. Click on Save to store the Proxy data into the configuration file permanently.

If you are required to manually configure the proxy to access the host where the Nextthink Application Library resides, the URL accessed by the Nextthink Engine is:

- application-library-v5.nextthink.com on TCP port 443 (HTTPS).
- application-library-v6.nextthink.com on TCP port 443 (HTTPS).

If you are required to manually configure the proxy to access the host where the Nextthink Updates reside, the URL accessed by the Nextthink Engine is:

- updates-v6.nextthink.com on TCP port 80 (HTTP) or 443 (HTTPS).
- updates?centos?v6.nextthink.com on TCP port 80 (HTTP) or 443 (HTTPS).

Nextthink Engine & Portal installation

Installation can be done either by using our online repository or, if the Appliance is not connected to the Internet, by uploading an offline installation package. When installing both the Engine and the Portal, especially during upgrading, always install or upgrade your Engines first and make sure that the update was successful. In that way, it is easier to roll-back to the original situation if something goes wrong.

From the Web Console, go to **Appliance > Install**:

1. Select if you want to use the Nextthink online repository or a file. In this second case, the file can be downloaded from the latest releases of our

- support portal. In the desired release page, look for the download of the Nextthink installation package (tgz file).
2. Choose to install the Engine, the Portal or both. Note that it is not recommended to install the Engine and the Portal on the same Appliance, except for small setups.
 3. Accept the license agreement.
 4. Click on install.

Please note that the installation process can be started only once. This page is no longer available after the first installation. Go to the **Update** page for upgrading your Appliance instead.

Related references

- Connectivity requirements

Installing the Collector on Windows

Overview

You must install the Collector in all the Windows devices from which you want to get end-user analytics. Depending on the number of devices and their geographical distribution in your corporate network, the installation of the Collector may be a technically challenging task.

For small setups you may opt for installing the Collector individually in each device. For medium to large setups, the installation of the Collector usually requires the use of automated deployment tools in practice.

To install the Collector and keep it up to date, Nextthink proposes a companion deployment tool to the Collector called the *Updater*. The Updater verifies the presence of the Collector in the devices of the end-users, checking its version and updating it when necessary. Control the deployment process from the Finder when deploying the Collector with the Updater.

Applies to platforms:

Prerequisites

You need:

- One or more Windows devices where to install the Collector.

- The Collector MSI or the Nextthink Collector Installer program.
- Optional: A third-party automated deployment tool. Instructions for installing through GPO (Active Directory) and with SCCM 2007 are provided.

You need to know:

- The IP address or DNS name of the Engine.
- UDP port number where the Engine is listening for the Collector (default 999).

Installing the Collector without the Updater

Installing the Collector on a single device

Use the Collector MSI package to install the Collector either in interactive mode or in silent (also known as *unattended*) mode. In the latter case, no user interaction is required once the installation process is started.

To install the Collector in interactive mode:

1. Double-click the Nextthink Collector MSI file (**NEXThink_Collector.msi**) and the following Screen is

- displayed.
2. Enter the IP address or DNS of the Engine and the UDP port open for

Collector communication.
3. Click **Install** to begin the installation.

4. Click **Finish** to close the dialog and end the installation.

Starting from V6, the Collector does not require rebooting the device after installation or upgrade. Once the MSI has successfully installed the Collector, you are requested to reboot the device only if you are upgrading from V5 or if you were running a Collector component during installation (usually, the Control Panel extension).

It is best practice not to install the Nextthink Collector Control Panel Extension when deploying the Collector on a production environment. The Collector Control Panel Extension is an optional tool for assistance during troubleshooting or testing phase, that is not required on the endpoint system for Collector to function correctly.

Alternatively, use *msiexec.exe* on the command-line to install the Collector in silent mode. The executable program **msiexec.exe** comes pre-installed with every Microsoft Windows operating system. Custom parameters are provided directly on the command-line and they are not saved from one installation to another. Therefore, this is not the recommended method to install the Collector, since commands are prone to typos. Instead, for a single installation, use the graphical installation method. For larger deployments without the Updater, we recommend to use the MSI Transform.

The mandatory parameters are:

- **DRV_IP**: this property must be set to the Engine IP address or DNS name.
- **DRV_PORT**: this property must be set to the Engine port number.

Here is an example of an unattended installation:

1. Type in the command line:

```
msiexec.exe /qn /i Nextthink_Collector.msi DRV_IP=192.168.84  
DRV_PORT=999 CPL_INSTALL=1
```

2. Wait for the installation process to complete.

Note that the MSI does not install in Windows Server devices by default. To install in a supported server device, set the option **DRV_FORCE_SERVER** to 1 in the MSI.

For a comprehensive list of available options for the Nextthink Collector, see the Collector MSI Parameters reference.

See also: Windows Installer (msiexec.exe) Command-Line options Reference.

Deploying the Collector using the Nextthink Collector Installer

The Nextthink Collector Installer is a tool that helps you deploy the Collector in situations where you cannot use the Updater. The tool is found in the archive of the Collector (**collector-<version number>.zip**) that you can find in the **Engine->Downloads** section of the Web Console. The name of the file is

NEXThink_Collector_Installer.exe.

The Collector Installer generates a standalone executable file that holds the MSI files for both the 32-bit and 64-bit versions of the Collector. Therefore, you can use the same executable to install the Collector on any machine that runs a supported version of Windows. When generating the executable, use the graphical interface of the Nexthink Collector Installer to set the installing options of the Collector:

Specify the configuration settings of the Engine that will receive Collector information:

- **Address:** DNS or IP address of the Engine.
- **Port:** Number of the UDP port where the Engine listens to data from the Collectors.
- **DNS:** Tick the option **Prefer IPv6** if you want the Collector to favor the use of IPv6 over IPv4 for communicating with the Engine when the name of the Engine resolves to both IPv6 and IPv4 addresses.

Among the other options, note that you can use the **Collector tag** to identify the installation. The tag is a fixed number that is sent to the Engine and is visible from the Finder.

If you check the option **Reuse configuration during upgrade**, the installer gets the settings **Address**, **Port**, and **Collector tag** from the configuration of any Collector previously installed in the device to upgrade. If no Collector is present in the device at the time of the upgrade, the installer reverts to the provided **Address**, **Port**, and **Collector tag** values.

There are some other self-explanatory installation options that you can check as well. In particular, check the option **Web and Cloud Data** if you have purchased the Web and Cloud module. Furthermore, click the button **Settings** to the right of this option to open a dialog where you can list the domains for which you want to store the full URL paths of web requests. That is, for every web request that falls under one of the specified domains, the Collector reports the full URL path to the Engine and not just the domain.

Click the option **Support Server Installations** for generating an executable that is able to install the Collector in supported Windows Server operating systems.

Finally, you specify a couple of directories:

- **Ouptut directory**: the place where the executable files to install or uninstall the Collector are generated.
- **Logs directory**: the network place where the installation logs of the Collectors deployed with this method are written.

When you click on **Create**, three files are generated:

- **NEXThink_Collector[engine_address].exe**: Executable file to install the Collector.
- **NEXThink_collector_Uninstaller[engine_address].exe**: Executable file to uninstall the Collector.
- **NEXThink_Collector[engine_address].exe.txt**: Text file with the list of the settings used to create the executable installer.

Nexthink Collector Installer error codes

The installer executable returns one of the following values:

- 0: Success; reboot is not required.
- -1: Success; reboot is required.
- 3: Failure; Collector installation has started but msixexec failed.
- other: Failure; the actual value corresponds to a Windows Internal error code.

The uninstaller executable returns one of the following values:

- 0: Success; reboot is not required.
- -1: Success; reboot is required.
- 1: Success; Collector was not found (nothing was uninstalled).
- 3: Failure; Collector uninstallation has started but msixexec failed.

- other: Failure; the actual value corresponds to a Windows Internal error code.

Remember that rebooting is usually not required when installing the V6 Collector. The installer requires you to reboot a device only if you are upgrading from the V5 Collector or if the Control Panel extension of the Collector was running during installation.

Installing the Collector with the Updater

Nexthink recommends to install and update the Collector with the help of a single component called the Updater, especially when you have to manage a large number of devices. By using the Updater, you can easily control the update process of the Collector from the Finder.

The Updater is a small Windows service that detects whether the Collector is installed in a single computer or not. If the Collector is indeed installed, the Updater can then detect if the Collector is up-to-date or if there is a new version of the Collector available. The Updater works in collaboration with the Engine to know if a Collector is outdated and, if this is the case, download and install the most recent version of the Collector from a centralized location. The configuration file of the Engine stores the installation options for the Collector and the configuration parameters of the Updater. Modify these settings following the instructions in the section Updating the Collector.

What you will need	One or more client computers
	Web access to the console
What you need to know	Engine IP address or DNS name
	TCP port number where the Engine is listening for the Updater

The installation of the Updater is very flexible. There are different possibilities to install the Updater in your end-user computers and you may combine them as you like, as long as they do not interfere with each other. Depending on your particular needs, you may want to take one of the following approaches:

1. **Individual installation:** you install the Updater on a single machine at a time, providing the configuration parameters of the Updater (Engine IP address and listening port number for Updater) to the setup program.
2. **Manual deployment:** instead of having to provide the configuration parameters of the Updater at every installation, you may want to generate a setup program holding pre-set configuration parameters with the help of Nexthink V5 Updater Installer tool. Then, you distribute the generated setup program to the end-user computers and execute it by whatever

means you prefer (for instance, using a remote execution program such as psexec).

3. **Deployment with external tools:** you can also use your preferred deployment tool to distribute and install the Updater in your set of end-user computers.

No matter what option you chose, if you installed the Updater in a client device with the correct configuration parameters, you are now able to control the installation of the Collector in that device from the Finder. The Finder not only lets you install the Collector, but it also interacts with the Engine and the Updater to let you see the update status of the Collector on every device. With the help of the Finder, you are able to manage the Collector deployment process and follow it step by step.

Installing the Updater on a single device

To install the Updater in a single machine, a setup program in the form of an MSI package is provided. Get the MSI package of the Updater from the Web Console:

1. Open the Web Console in your web browser.
2. Go to the **Engines** tab.
3. Select the **Downloads** section.
4. Download the Nextthink Collector Archive, a ZIP file that holds a set of tools to install the Collector.
5. Extract the file *NEXThink_Updater.msi* from the Collector Archive.

The MSI can work both in interactive or silent modes. In interactive mode, you set the configuration parameters of Updater as you are asked to do so by the setup program. Alternatively, you can install the Updater silently by providing the configuration parameters in the command line.

To install the Updater in interactive mode:

1. After retrieving the setup program of Updater from the Engine Downloads page, execute it in your client computer by double-clicking on the MSI file (its name is **NEXThink_Updater.msi**). The welcome screen of the setup program is displayed, including Updater version information. Click on

5. After a few seconds, the setup is completed and Updater is successfully installed. Click **Finish** to end the installation.

Alternatively, you can install the Updater silently by using the command line. You use the standard Windows Installer command line program `msiexec.exe` for that purpose. In silent mode, instead of supplying the Updater configuration parameters interactively, you specify them as properties within the same command that launches `msiexec.exe`. The public properties of Updater are:

- **UPD_IP**: property holding the Engine IP address or DNS name.
- **UPD_PORT**: property holding the port number where the Engine is listening for the Updater requests.
- **UPD_HTTPS**: this property is optional and it should not be normally changed. When set to 1, Updater uses the secure protocol HTTPS to communicate with the Engine. When set to 0, Updater uses HTTP instead. By default, this property is set to 0.

To install the Updater silently:

1. Open a Windows command line by pressing **Windows key + R** in your keyboard and then type **cmd** followed by [ENTER].
2. A command line interface window opens. Type the following command on it:

```
msiexec.exe /qn /i Nexthink_Updater.msi UPD_IP=192.168.84  
UPD_PORT=8888
```
3. Updater installs without giving any notification to the user.

Generate the Nexthink Updater Installer

For installations with a large number of end-user computers, specifying configuration settings for the Updater in every machine can be tedious and error prone. The Nexthink Updater Installer lets you create a silent setup program with the settings that you choose embedded in it.

1. Download the complete Nexthink Collector archive from Product Downloads: you get a zip file called **collector-<version number>.zip**.
2. Unzip the archive with your favorite decompression tool and extract the Nexthink Updater Installer. This is the file **NEXThink_Updater_Installer.exe** that you will find under the **installer** directory once you unzip the archive.
3. Execute the Nexthink Updater installer and set the configuration values for the Engine IP address, Engine port for the Updater and protocol (HTTPS

or HTTP, as explained above).

4. Specify additional options:

- ◆ **Hide from Add/Remove programs** prevents end-users from uninstalling Updater.
- ◆ **Support Server Installations (XenApp, RDS)** indicates whether the setup program should also install in server versions of Windows or not. Usually, you do not want to install Collector nor Updater in your servers.
- ◆ **Uninstall Collector V3/V4/V5** makes the generated installer to uninstall previous versions of the Updater if they are present in the device. The option is ticked by default and it is recommended to leave it ticked.
- ◆ **Output Directory** is the path where you want the setup program to be generated, that is, the executable that you will use to install Updater in your end-user computers.
- ◆ **Logs Directory** specifies the location where the Updater installer will upload error messages in case of problems with the installation of Updater. We recommend that you specify a shared folder in your network for this field. All Updater installers will then report errors to the same location. (By default, the error log goes into the %TEMP% folder)

5. Click on **Create** and the program generates an executable file and a text file with the configuration parameters that you just entered in the location that you specified. This executable is able to install Updater in all the versions of Windows supported by Updater.

Deploying the Updater or the Collector through GPO

In this section, learn how to deploy the Updater or the Collector over large groups of end-user devices using a standard Window technology such as Active Directory Group Policy. Nextthink recommends you to deploy the Updater rather than the Collector. The Updater, in its turn, installs the Collector and maintains it

up-to-date in successive upgrades. Additionally, it lets you follow and control the deployment process through the Finder. If you directly deploy the Collector instead, you need to redeploy it on every new update of the software. The installation policy in your company may however require you to directly deploy the Collector and avoid any installer agent such as the Updater. Choose whether to deploy the Updater or directly the Collector depending on your specific circumstances.

This section assumes that you are a systems administrator with a basic understanding of the Windows operating system and deploying enterprise software, and that you are familiar with Group Policy and Active Directory.

Creating a distribution point

1. Log on to the server as an Administrator user.
2. Create a shared network folder (this folder will contain the MSI package).
3. Set permissions on this folder in order to allow access to the distribution package.
4. Copy MSI(s) in the shared folder.
5. Copy MST(s) generated in the shared folder.

Creating a Group Policy Object

1. Click on the **Start** button, go to **All Programs**, select **Administrative Tools** and then select **Active Directory Users and Computers**.
2. Right-click your domain name in the console tree, select **New** and click **Organizational Unit**.
3. In the **New Object** dialog box, type a descriptive name for the new organizational unit, and then click **OK**.
4. In the right panel, select **Computers** and click on the computer that you want add to your Organizational Unit.
5. Drag and Drop these computers in the name of the Organizational Unit created. In the right panel, select **Nextthink_Collector_Deploy**, you will see all the computers tied to your Organizational Unit.
6. Click on the **Start** button, go to **All Programs**, select **Administrative Tools** and then select **Group Policy Management**.
7. Right-click your domain name in the console tree and select **Create a GPO in this domain**, and **Link it here....**
8. In the **New GPO** dialog box, type a descriptive name for the new policy, and then click **OK**.

Assigning a MSI package

1. Click on the **Start** button, go to **All Programs**, select **Administrative Tools** and then select **Group Policy Management**.
2. Right-click your GPO name and select **Edit....**
3. On this **Group Management Editor**, expand **Computer Policies, Software Settings and Software Installation**, select **New** and then click **Package....**
4. In the **Open** dialog box, browse to the distribution point you created for the Nextthink Collector during the distribution point.
5. Select the MSI file containing the Collector installer you want to deploy, and then click **Open**.
6. In the **Deploy Software** dialog box, select **Advanced**, and then click **OK**.
7. In the **Properties** dialog box for the package you created.
 1. Click the **Deployment** tab, and then select **Uninstall** this application when it falls out of the scope of management.
 2. Click **Advanced** on the **Deployment** tab, choose **Ignore language when deploying this package**, uncheck the option **Make this 32-bit X86 application available to Win64 machines**, and then click **OK**.
 3. On the **Modifications** tab, specify any modification transforms you want to apply when the package is installed by clicking **Add** and then opening each transform from its network location.
 4. On the **Security** tab, verify the name(s) of any computer(s) to which you are assigning software.
 5. Click **OK** to close the Properties dialog box.
8. In the **Group Policy** dialog box, expand **Computer Configuration, Administrative Templates, and Windows Components**.
 1. In the **Windows Components** folder, select **Windows Installer**.
 2. Select **Always install with elevated privileges**.
 1. Select **Properties**.
 2. In the **Always install with elevated privileges Properties** dialog box, click the **Setting** tab, select **Enabled**, and then click **OK**.
9. In the **Windows Installer** panel of the **Group Policy** dialog box, right-click **Logging**, and then select **Properties**.
 1. In the **Logging Properties** dialog box, on the **Setting** tab, select **Enabled**.
 2. Then, in the **Logging** text box, type **iweaprcv**.
 3. Click **OK** to close the **Logging Properties** dialog box.
10. In the **Group Policy** dialog box, click **File**, and then click **Exit**.

Note: The GPO must be propagated to the Active Directory Global Catalog and then to the individual computers. For this reason, allow 5-10 minutes before restarting the computers to which you are assigning the Nextthink Collector, or plan to restart the client computers twice before the system policies are synchronized.

Testing your results

1. Go to a the target PC that is member of the OU you tied the policy to.
2. Click **Start, Run** and type **gpupdate /force**.
3. A logoff or a restart message will appear: type **Y** and Enter.
4. When you restart, you should see the message **Installing Nextthink Collector...** for about a minute depending on the speed of your network and pc.
5. Click **Start, Run** and type **eventvwr.msc** to show the event viewer message. In the left pane, select **Application**, you will see some source events **Msiinstaller** logged as a Success Audit event.
6. If you have some errors, go to *C:Windows/Temp/Msi.log* and see the error log generated.

Redeploying a MSI package

Sometimes you may need to redeploy a package (for example when doing an upgrade). For redeploying a package you can follow these steps:

1. Click the **Start** button, go to **Programs**, select **Administrative Tools** and then select **Group Policy Management**.
2. Right-click your domain name in the console tree and select the **Properties** context menu.
3. Go to the **Group Policy** tab, select the object you used to deploy the package and click **Edit**.
4. Expand the **Software Settings** element (per-user or per-machine) which contains the deployed package.
5. Expand the **Software Installation** element which contains the deployed package.
6. Right-click the package in the right pane of the **Group Policy** window.
7. Select the **All Tasks** menu and click **Redeploy application**.
8. Click the **Yes** button for reinstalling the application wherever it is installed
9. Close the **Group Policy Management Editor**, click **OK** and exit the Group Policy Management.

Note that the GPO must be propagated to the Active Directory Global Catalog and then to the individual computers. For this reason, allow 5-10 minutes before

restarting the computers to which you are assigning the Collector, or plan to restart the client computers twice before the system policies are synchronized.

Removing a MSI package

1. Click on the **Start** button, go to **Programs**, select **Administrative Tools** and then select **Group Policy Management**.
2. Right-click your domain name in the console tree and select the **Properties** context menu.
3. Go to the **Group Policy** tab, select the object you used to deploy the package and click **Edit**.
4. Expand the **Software Settings** element (per-user or per-machine) which contains the deployed package.
5. Expand the **Software Installation** element which contains the deployed package.
6. Right-click the package in the right pane of the **Group Policy** window.
7. Select the **All Tasks** menu and click **Remove**.
8. Select from the following options:
 - ◆ Immediately uninstall the software from users and computers.
 - ◆ Allow users to continue to use the software but prevent new installations.
9. Click the **OK** button to continue.
10. Close the **Group Policy Management Editor**, click **OK** and exit the Group Policy Management.

Note that the GPO must be propagated to the Active Directory Global Catalog and then to the individual computers. For this reason, allow 5-10 minutes before restarting the computers to which you are assigning the Collector, or plan to restart the client computers twice before the system policies are synchronized.

Deploying the Updater or the Collector through SCCM 2007

In this section, learn how to deploy the Updater or the Collector over groups of end-user devices using Microsoft System Center Configuration Manager. As explained in the deployment through GPO, Nextthink recommends you to deploy the Updater rather than the Collector, but you can choose whether to deploy the Updater or directly the Collector depending on your specific circumstances.

This section assumes that you are a systems administrator with a basic understanding of the Windows operating system and deploying enterprise software, and that you are familiar with SCCM 2007. For other versions of SCCM, the procedure may be slightly different. Please refer to the section on deploying MSI packages in the user manual of your specific version to deploy the

Collector.

Creating a collection

1. Click on the **Start** button, go to **All Programs**, select **Microsoft System Center**, select **Configuration Manager 2007** and then select **ConfigMgr Console**.
2. In the **Configuration Manager Console**, navigate to **System Center Configuration Manager / Site Database / Computer Management**.
3. Right-click **Collections**, and then click **New Collection**.
4. On the **General** dialog box of the **New Collection Wizard**, enter a name for the collection.
5. Click **Next** and click the computer icon, which opens the **Create Direct Membership Rule Wizard**. Click **Next**.
6. On the **Search for Resources** dialog box:
 1. Click the **Resource** class drop-down menu and select **System Resource**.
 2. Click the **Attribute name** drop-down menu and select **Name**.
 3. In the **Value** field enter %, and then click **Next**.
7. On the **Collection Limiting** dialog box, click the **Browse...** button, select **All Windows Workstation or Professional Systems**, and then click **Next**.
8. On the **Select Resources** dialog box, select the check box for each of the targeted computer resources.
9. Click **Next**, and then on the **Finished** dialog box, click **Finish**.
10. On the **Membership Rules** dialog box of the **New Collection Wizard**, click **Next**.
11. On the **Advertisements** dialog box, for now, do not assign an advertisement because it is not yet created.
12. Click **Next**. On the **Security dialog** box, accept the defaults, click **Next**, and then click **Close**.

Creating a package source directory

1. Click on the **Start** button, go to **All Programs**, select **Microsoft System Center**, select **Configuration Manager 2007** and then select **ConfigMgr Console**.
2. Navigate to **System Center Configuration Manager / Site Database / Computer Management / Software Distribution**.
3. Right-click **Packages**, point to **New**, and then click **Package**.
4. On the **General** dialog box of the **New Package Wizard**, enter the:
 - ◆ Name
 - ◆ Version

- ◆ Manufacturer
 - ◆ Language
5. Click **Next** and do the the following:
 1. Select **This package contains source files**.
 2. Click the **Set** button and then enter the path for the location of the source files in the **Source directory** field.
 6. Click **Next** and accept the default settings on all of the following dialog boxes: Data Access, Distribution Settings, Reporting, and Security.
 7. On the **Wizard Completed** dialog box, click **Close**.

Creating a distribution point

To use a server as a distribution point for providing packages to distribute packages to your client computers, you must first designate a site system as a distribution point. To select a distribution point for the newly created package:

1. Right-click **Distribution Points**, click **New Distribution Points**, click **Next**, and then click the check box for the distribution point.
2. Click **Next**. Upon completion of the **New Distribution Points Wizard**, click **Close**.

Creating a program with setup and install parameters

1. Right-click **Programs**, point to **New**, and then click **Program**.
2. On the **General** dialog box:
 1. Enter a name for the package in the **Name** field.
 2. In the **Command line** field, browse and select the appropriate executable file.
 3. In the **Run** field, click the drop-down menu and select **Hidden**.
 4. In the **After running** field, verify the default of **No action required** is selected.
3. Click **Next** and accept the defaults on the **Requirements** dialog box.
4. On the **Environment** dialog box:
 1. Click the **Program can run** drop-down box and select **Whether or not a user is logged on**. This will enable Run with administrative rights for the Run mode.
 2. Leave the default for **Drive mode** to **Runs with UNC name**.
5. Click **Next**.
6. On the **Advanced** dialog box, select the check box for **Suppress program notifications**, and then click **Next**.
7. To view the **Summary** dialog box, click **Next**.
8. To finish the process of creating the new program, click **Next**, which will then display the **Wizard Completed** dialog box.

9. To exit from the **New Program Wizard**, click **Close**.

Verify that the package was installed on the distribution point, by:

- Navigating to System Center Configuration Manager, Site Database, Computer Management, Software Distribution, Packages, Collector package name, Package Status, Package Status.
- Checking the status changing from Install Pending to Installed.

Creating the advertisement

1. Right-click **Advertisements**, point to **New**, and then click **Advertisement**.
2. On the **General** dialog box of the **New Advertisement Wizard**:
 1. Enter a name for the advertisement in **Name** field.
 2. Click the **Browse** button for the **Package** field, and click on the package you want to advertise and then click **OK**.
 3. Click the **Browse** button for the **Collection** field, click on the collection.
3. Click **OK**, and then click **Next**.
4. On the **Schedule** dialog box, enter the date and time in the Advertisement start time fields for when the advertisement will begin, and then click the asterisk button for **Mandatory Assignments**.
5. On the **Assignment Schedule** dialog box, click the **Schedule** button and enter the same date and time that you previously entered in the **Advertisement start time** fields on the **Schedule** dialog box.
6. To return to the **Schedule** dialog box, click **OK**.
7. Accept the default values on the **Distribution Points**, **Interaction**, **Security**, and **Summary** dialog boxes.
8. Upon successful completion of the **New Advertisement Wizard**, click **Close** on the **Wizard Completed** dialog box.

On the client, wait for the next Machine Policy Retrieval & Evaluation Cycle.

Testing your results

1. Go to a the target PC that is member of the Collections you have created to deploy.
2. Click **Start, Run** and type **eventvwr.msc** to show the event viewer message.
3. In the left pane of the **Event Viewer**, select **Application**, you will see some source events **MsiInstaller** logged as a Success Audit event.
4. If you get any error, see the error log generated in **C:\Windows\Temp\Msi.log**.

Removing the package

1. Open the **Systems Management Server console** and expand the Package that contains the Collector.
2. Open the **Program Properties** dialog box, and then on the **General** tab:
 1. In the **Command line** field, browse and select the Collector file.

The package will be removed silently at the next Machine Policy Retrieval & Evaluation Cycle from the end-user device.

Installing the Collector with Windows Master image

WARNING	The Nexthink Collector and Nexthink Updater must not be included in a Windows Master image without checking the procedure with Nexthink Customer Success Services or Certified Partners.
---------	--

Interaction of the Collector with other software

To get valuable information from a device, part of the Collector needs to run in privileged mode as a kernel driver. Contrary to user applications, the programs that run in privileged mode can access the memory and the hardware of the device directly. Typically, these are the programs that control the peripherals in your device; such as the mouse, the keyboard, the hard disks or the network card. Other special programs, like antivirus software, may also need to run in privileged mode, at least partially. Errors in programs that run in privileged mode are not protected by the process isolation provided by the operating system and, therefore, they may result in system failure. Moreover, since all of these programs share the same memory space, a misbehaviour of one of them can destabilize all the others.

The Collector has been carefully designed and thoroughly tested to avoid any kind of program errors. It has also been engineered following the best practices for the development of kernel drivers, behaving as a good citizen with respect to the other drivers that are loaded into the system. Still, in some very rare cases, an elusive programming error may defeat our rigorous testing process or a misbehaving third-party driver can trigger a fault in the Collector. In these unfortunate situations, the Collector may become unstable and possibly lead to a system failure in the device of the end-user.

To protect you against driver misbehaviour, keep your Windows drivers up to date. Older versions of Windows drivers often contain more bugs that can lead to instabilities. If a third-party driver consistently destabilizes the Collector, the Collector can prevent the system from crashing again and again if you activate its *CrashGuard Protection*. This mechanism cancels the loading of the Collector at

system startup if the system crashes more than a specified number of times within a configurable security interval after system boot. See the definition of the parameters **DRV_CRASHGUARD** and **DRV_CGPI** of the Collector installer to configure the CrashGuard Protection during installation. Use the Control Panel extension of the Collector or the Collector configuration tool to modify the settings of the CrashGuard Protection after installation.

In any case, if you suspect that there is a compatibility problem between any of the drivers loaded into the end-user devices of your company and the Collector, please contact Nextthink Support.

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

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

Related tasks

- Create or delete a Group Policy object
- SCCM 2007 POC Installation Guidelines

Related references

- Updating the Collector
- Collector MSI Parameters reference
- Microsoft Windows Server 2003/2008 Group Policy home page
- System Center Configuration Manager

Installing the Collector on Mac OS

Overview

Nextthink distributes the Collector for Mac OS as a disk image (**.dmg**) file with the following contents:

- A package (**.pkg**) file for installing the Collector from a graphical user interface.
- The application **csi.app** for installing the Collector from the command line interface.

- A reporter shell application that gathers system information in the case that you find any issue when running the Collector on Mac OS.
- An uninstaller application to remove the Collector when it is no longer needed.

Starting from macOS 10.13 High Sierra, loading new kernel extensions into the system requires explicit user approval. As the Mac Collector requires the loading of kernel extensions, individual fresh installations of the Collector in macOS 10.13 High Sierra need thus to be approved by the end-user. Once it has been approved, later upgrades of the Collector do not need additional user approval for loading kernel extensions.

See below how to avoid the requirement of user approval when deploying the Collector in an enterprise environment.

In the case of an upgrade of the operating system to macOS 10.13 High Sierra, if the Collector was previously installed, the user approval to load kernel extensions is not required.

Remember to reboot your Mac computer every time that you install, upgrade, or uninstall the Collector.

Applies to platforms:

Graphical installation

To install the Collector on Mac OS using the graphical interface:

1. Double-click the provided disk image file to mount it into your filesystem and see its contents.
2. Double-click the package file **Nexthink_Collector_5.x.x.pkg** and the installer starts with the introduction.
3. Click **Continue** to proceed with the installation.
4. In the step **Personalization**, configure the settings of the Engine to which the Collector will send the gathered information:
 - ◆ Type in the host name or IP address of the appliance running the Engine under **Server name or IP address**.
 - ◆ Type in the port number where the Engine listens for connections from the Collector under **UDP port**.
5. Click **Continue** to go on.
6. In the step **Destination select**, the installer program shows the local paths in the system where it is going to install the different components of the Collector. Keep the default paths and click **Continue**.

7. The **Installation Type** step informs you about some details of the installation process, including the amount of space that the program is going to occupy on disk. Click **Install** to begin with the actual installation.
8. The installer shows the progress of the installation and it finishes with a summary message. If the installation was successful, click **Close** to terminate the procedure.
9. Reboot the computer to finish the installation.

Command line installation

The command line installation lets you install the Collector even when you have access to a computer only through the shell of Mac OS. Using the command line installation, you can thus install the Collector either locally or remotely through an *ssh* connection.

To install the Collector from the command line, choose between:

- Providing the installation options as arguments to the installation command.
- Writing the installation options in a configuration file.

In both cases, you execute the *csi* application provided with the disk image. To mount the disk image into the file system:

1. After downloading the image file from Product Downloads, pick one of the following options:
 - ◆ If you are installing the Collector in a remote computer:
 1. Copy the image file to the remote computer:

```
scp Nexthink_Collector_5.x.x.dmg  
<username>@<address>:
```
 2. Log in to the remote computer:

```
ssh <username>@<address>
```
 - ◆ If you are installing the Collector in the local computer:
 1. Change the directory to the one where you downloaded the image file.
2. Mount the image file:

```
hdiutil mount Nexthink_Collector_5.x.x.dmg
```

Providing installation options as arguments

Once with the image file mounted into the filesystem of the target Mac computer, install the Collector from the command line:

1. Change the directory to the path of the *csi* application:

```
cd /Volumes/Nextthink_Collector_5.x.x/csi.app/Contents/MacOS
```
2. Type in the following command, providing the DNS name or IP address of the Engine as the first argument and the port where the Engine listens as the second argument:

```
sudo ./csi -address <engine_address> -port  
<engine_udp_port>
```

 - ◆ If you are installing the Collector on macOS 10.13 High Sierra you need user approval.
3. Reboot your computer to finish the installation.

Using a configuration file

To install the Collector from the command line without having to pass any arguments to the *csi* application, create a configuration file named *config.plist* in the following path:

```
/Volumes/Nextthink_Collector_5.x.x/csi.app/Contents/Resources/config.plist
```

Inside the configuration file, specify the Engine DNS name or IP address, the log level of the installer program and the port where the Engine is listening to connections from the Collectors in the following way:

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE plist PUBLIC "-//Apple//DTD PLIST 1.0//EN"
"http://www.apple.com/DTDs/PropertyList-1.0.dtd">
<plist version="1.0">
  <array>
    <dict>
      <key>Address</key>
      <string>servername</string>
      <key>Log</key>
      <string>Informal</string>
      <key>Port</key>
      <string>serverport</string>
    </dict>
  </array>
</plist>
```

Possible values for the log level are:

Verbose

The installer logs every incidence, warning or error.

Informal

The installer logs only critical errors.

Silent

The installer produces no logs.

Once you have saved the configuration file, install the Collector without providing any argument:

1. Change the directory to the path of the *csi* application:

```
cd /Volumes/Nextthink_Collector_5.x.x/csi.app/Contents/MacOS
```

2. Execute the installer application without arguments:

```
sudo ./csi
```

- ◆ If you are installing the Collector on macOS 10.13 High Sierra you need user approval.

3. Reboot your computer to finish the installation.

Uninstalling the Collector

To uninstall the Collector, execute the *uninstaller* script that is provided with the image file. Assuming that you have mounted the image file into the filesystem of the computer where the Collector is installed:

1. From the shell, type in the following command:

```
sudo /Volumes/Nextthink_Collector_5.x.x/uninstaller
```

2. After the uninstaller script ends displaying the steps that it is performing, reboot your computer to finish the uninstallation. For instance, to reboot the computer 10 minutes after uninstallation, type in:

```
sudo shutdown -r +10
```

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

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

Related references

- Operating systems supported by the Collector

Installing the Mobile Bridge

The installation of the Mobile Bridge requires the configuration of two different machines:

- Configure the Exchange Server from which the Mobile Bridge will retrieve information about mobile devices.

- Set up and configure the Windows server on which the Mobile Bridge will run.

Creating the AD user for the Mobile Bridge

The installation of the Mobile Bridge requires the creation of a new user in your Active Directory. The Mobile Bridge impersonates this user to communicate with Exchange.

1. Log in as administrator to a Windows server and connect to Active Directory via the Microsoft Management Console.
2. Create a new user exclusively dedicated to the Mobile Bridge.
 - ◆ The user must belong to the **View-Only Organization Management** role group.
 - ◆ Verify that the user has the attribute **RemotePowerShellEnabled** set to **\$true** (this is the case by default when you create a new user).

Configuring the Exchange Server

Enabling Windows authentication in IIS for PowerShell

First of all, you must configure the IIS in your Exchange server to allow the Mobile Bridge user to connect via PowerShell to it using Windows authentication. If using Exchange 2010, we suggest to use a Client Access Server. To enable Windows authentication in IIS for PowerShell users:

1. Log in to the Exchange Server as administrator.
2. Open IIS Manager.
3. In the left hand-side pane, go to **Sites** -> *Name of your Exchange site* -> **PowerShell**.
4. Open the **Authentication** page. All authentication methods are disabled by default.
5. Right-click the status of **Windows Authentication** and choose **Enable**.

Advanced configurations

Configuring the Application Pool of IIS

To limit the performance hit of the Mobile Bridge when it connects to the Exchange Server, set a recycling interval for PowerShell and, optionally, limit the maximum amount of memory that it may use within IIS:

1. Log in to the Exchange Server as administrator.
2. Open IIS Manager.
3. In the **Connections** pane to the left, select the node **Application Pools** in the tree.
4. On the **Application Pools** page, select the application pool **MSExchangePowerShellAppPool**.
5. In the **Actions** pane to the right, click the option **Recycling...** under the section **Edit Application Pool**. This opens the **Recycling Conditions** dialog:
 1. Tick the option **Specific time(s)** and set its value to, for instance, 02:00 AM (choose an hour of low activity).
 2. Optional: Tick the option **Private memory usage** and set it to a sensible value according to the RAM available in your server (e.g. set it to 2 000 000 to limit the memory usage to ~2 GB).
6. Click **Next**.
7. Select the recycling events that you want to log, if any, and click **Finish**.

Disabling PowerShell logs in IIS

The installation of the Mobile Bridge generates a significant increase in the number of PowerShell requests to the IIS within your Exchange Server. This subsequently causes a significant increase in the size of the logs of IIS (up to 1 GB per day).

Therefore, we recommend that you disable PowerShell logging in IIS after installing the Mobile Bridge. To disable PowerShell logging:

1. Open IIS Manager and navigate to the PowerShell node.
2. In the **Features View**, double-click **Logging**.
3. In the **Actions** pane of the **Logging** page, click **Disable**.

Configuring the Mobile Bridge in the Windows server

Software requirements of the Mobile Bridge

Installing the Mobile Bridge requires the following software:

- Windows Server 2008 R2, 2012 or 2012 R2.
- .NET Framework 4.5
- PowerShell 4.0

Higher versions of this software may also be suitable for running the Mobile Bridge, but they have not been tested.

Hardware requirements can be found here.

Installing and running the Mobile Bridge service

To install the Mobile Bridge service with the user interface:

1. Double-click the installer file **Nexthink.Mobile.Bridge.msi** that you get from Product Downloads.
2. Accept the license agreement and click **Install**.
3. Once the wizard has ended the installation, click **Finish**.

The procedure above makes the Mobile Bridge use the default port number 11031 for communicating with the Engine. If you want to communicate with the Engine through a different port, install the Mobile Bridge from the command line:

1. Open a command line interface with elevated privileges.
2. Type in the following command to install the Mobile Bridge and, for instance, instruct it to use port number 12 000 for communicating with the Engine:
msiexec -i Nexthink.Mobile.Bridge.msi PORT=12000
3. Change the port by default in the Engine configuration as well.

To configure the Mobile Bridge service:

1. Open the command line interface with elevated privileges.
2. Configure the service with the address of the Exchange server and the user name(that of the dedicated user for the Mobile Bridge that you created in the Exchange server):
"c:\Program Files (x86)\Nexthink Mobile Bridge\Nexthink.Mobile.Bridge.exe" ^
-address myexchangeserver.example.com ^
-username nxtBridgeViewOnlyAdmin@example.com
3. Enter the password for the Mobile Bridge user. The password is encrypted and stored in the configuration file of the Mobile Bridge along with the address of the Exchange server and the name of the user.

To run the Mobile Bridge service:

1. From the command line interface, type in:
sc start NexthinkMobileBridge

At this point, the Mobile Bridge service validates your settings and attempts to connect to the Exchange Server.

Setting the Mobile Bridge velocity

If the Mobile Bridge takes too much time to retrieve the information from the Exchange server, or if at the opposite the Exchange server load due to the Mobile Bridge queries is too high, the *throttling* can be adjusted.

The *throttling* is the idle time between to queries of the Mobile Bridge. Its default value is 500 ms.

To change it, use the following configuration options:

```
<add key="Throttle" value="500" />
```

Filtering mobile devices based on AD security groups

If the Mobile Bridge should not report information of mobile users belonging to a specific group, use the **ExcludedGroupDn** option to specify the group whose users must not be monitored. On the other hand, to explicitly include a group for the Mobile Bridge to report information about its users, use the **IncludedGroupDn** option. The options translate to the following PowerShell filter when retrieving information:

```
-Filter {(MemberOfGroup -ne 'ExcludedGroupDn') -and  
        (MemberOfGroup -eq 'IncludedGroupDn')}
```

Troubleshooting the Mobile Bridge

The Mobile Bridge service does not start

To check the status of the service, type in the following command:

```
sc query NexthinkMobileBridge
```

If the service fails to start, look in the Windows Event logs for error messages indicating the possible reason and take appropriate action. Alternatively, check the log files in:

```
%ProgramData%\Nexthink Mobile Bridge\logs\  
  .\nexthink-mobile-bridge-global.txt  
  .\nexthink-mobile-bridge-powershell-errors.txt
```

Beware that the service may take a long time to start. The Mobile Bridge service needs to validate the connectivity to the Exchange server, the availability of the required PowerShell commands (*cmdlets*) as well as the version of the Exchange server before reporting a successful start to Windows.

The Mobile Bridge connection to Exchange fails

If PowerShell connectivity to the Exchange server is failing, carry out the following verifications:

1. Verify that the configured address for the Exchange server is reachable by typing in from the command line:

```
ping myexchangeserver.example.com
```

2. Verify that the provided credentials are valid by typing in the following two lines in a PowerShell window:

```
$session = New-PSSession -ConfigurationName  
Microsoft.Exchange -ConnectionUri  
"https://myexchangeserver.example.com/powershell/"  
Import-Session $session
```

The Mobile Bridge does not validate the Exchange server certificate for HTTPS

Before production, you may want to deactivate the validation of certificates when the Mobile Bridge connects to the Exchange server. To disable the validation of the certificate common name (CN), the certificate authority (CA) and the certificate status, respectively use the following configuration options:

```
<add key="SkipCNCheck" value="true"/>  
<add key="SkipCACheck" value="true"/>  
<add key="SkipRevocationCheck" value="true"/>
```

Related references

- Mobile Bridge configuration settings

Installing the Finder

Install the Finder by downloading and executing one of the setup programs available from the Product Downloads website (currently support.nexthink.com). To install the latest Finder:

1. In the **Product Downloads** page, select the first entry of the **Last V6 releases**.
2. In the **Download links** section, find the links for the Finder.
3. Choose one set of downloads for the Finder, depending on the architecture and memory of your computer:
 - ◆ 32-bit version: recommended for 32-bit versions of Windows or for computers with less than 4 GB of RAM (even when running a 64-bit version of Windows).
 - ◆ 64-bit version: suitable only for computers with more than 4 GB of RAM running a 64-bit version of Windows.
4. Choose one of the available Finder downloads:
 - ◆ MSI installer (*Nexthink_Finder.msi*): Select this option if it is the first time that you install the Finder and Microsoft .NET Framework 4.5 is already installed in your computer.
 - ◆ Bundled Installer, including the .NET Framework (*Nexthink_Finder.exe*): If you do not have the Microsoft .NET Framework 4.5 in your computer, select this download to install the Finder and the .NET Framework at the same time. The size of this download is significantly bigger than the size of the standalone installer.
 - ◆ Standalone executable (*nxfinder.exe*): This download is not really a setup program but an executable file of the Finder itself. Copy the file *nxfinder.exe* to the path of your preference in your computer and launch the Finder from there. Contrary to the two previous installation methods, this method lets you have several versions of the Finder available in a single computer at the same time. For normal Finder utilization, we recommend however to use one of the previous installation methods and have just a single version of Finder installed (the *nxt* protocol, for instance, works only with properly installed versions of the Finder).

Installation Procedure with the Standalone Installer

To install the Finder using the MSI package :

1. Launch the installation by double-clicking *Nexthink_Finder.msi*.

2. Read and acknowledge the End-User License Agreement of the Finder.
3. Provide a destination location where to install the Finder. If you do not have Administrator privileges, you must specify a user-writable directory.
4. Confirm the installation by clicking **Install**. The setup program notifies you when the installation is complete.

Installation Procedure with the .NET Framework 4.5

1. Launch the installation program by double-clicking *NexthinkFinder.exe*.
2. The installation of Microsoft .NET Framework 4.5 Client Profile starts.
3. Once .NET Framework 4.5 is installed, the actual Nexthink Finder V4.5 installation begins.
4. Read and accept the license agreement, provide a destination path and click on *Install in the same way as explained for the Standalone Installer*.
5. The setup program informs you when the installation is complete.

Customize the Finder Installer

The `msiexec.exe` command-line utility provides ways to install the Finder without using the UI of the installer.

For example, the command-line for a silent install of the Finder for all users will be:

```
msiexec /i NEXThink_Finder.msi ALLUSERS=1 /qn
```

To see all options, use:

```
msiexec ?help
```

Compatibility with the Portal and the Engine

To communicate properly with the rest of the Nexthink ecosystem, install a version of the Finder that is compatible with the versions of both the Portal and the Engine that you have in place. Please refer to the Release Notes of the Finder to verify that your particular version is compatible with your current Portal and Engine versions.

The Finder detects whether it is up to date with respect to the versions of the Portal and Engine in place when it tries to establish a connection after you log in. In the case that you are running an outdated version of the Finder, a dialog offers you the possibility to download a compatible release of the Finder from a given link.

Security certificates

Upon the first execution of the Finder, you may experience some warnings related to security certificates. Certificates ensure that the communication among Nextthink components is safe. Refer to the sections about logging in to the Finder and the replacement of security certificates for more information.

Related tasks

- Logging in to the Finder
- Importing and replacing Certificates

Customer satisfaction program

The following article outlines the data collected by Nextthink within the customer programs and how it functions.

Customer Experience Improvement Program (CEIP)

Nextthink is continually striving to understand and anticipate our customer needs in order to deliver world-class products and solutions. The Nextthink Customer Experience Improvement Program (CEIP) will deliver benefits to the customers by allowing us to understand how you use our software, so that we can provide you with a continuous enhancement of your Nextthink software experience. The program is voluntary and anonymous. Customers who choose to participate agree to share:

- Information such as operating system version, processor, and memory installed on the device running the Finder
- Nextthink product information such as version number
- The country where the software is running
- Nextthink feature usage information such as menu options or buttons selected
- Execution time for specific operations
- Error reports
- Engine performance and usage information

Frequently Asked Questions

What are the possible configuration settings of the CEIP program?

The participation in the program is enabled by default, but this can be modified in the Web Console. The options are:

- Enabled
- Disabled

What will Nexthink do with the information that is collected?

The information collected will be used to better understand how customers use Nexthink products, and how to improve Nexthink products by fixing issues and delivering the most useful new features in a much more streamlined manner.

Is the collected information anonymous?

Yes! Moreover, Nexthink takes many precautions in protecting the security of the information that is collected, transmitted and stored.

How does the Nexthink Customer Experience Improvement Program work?

This is an automated process that requires no effort to participate, it is transparent to the users. Customers simply choose to participate, granting Nexthink permission to securely receive anonymous data.

Will I receive spam if I participate in the program?

You will not receive any e-mail from Nexthink regarding this program, regardless of whether or not you participate. We do not collect personal identifiable information as part of this program that will be used to identify you or contact you.

Do I need an Internet connection?

An Internet connection is required to participate in this program. However, you do not need to be connected all the time. When an Internet connection becomes available, the information is automatically transmitted with minimal impact to your connection.

How long does the program last?

Information is collected as long as you use the product version for which you have agreed to participate or until you decide to stop participating in the program.

What is the anonymous installation ID used for? Upon the first startup of the Nexthink Finder, a random number is generated, the anonymous installation ID.

This installation ID can be used to help you if you experience issues with the Nextthink Finder. Should there be a specific need to interact with a user, with this specific ID we will be able to better understand which part of the Nextthink Finder is not working, do fine grained debugging, thus providing you with tailored customer support.

Which products support the Nextthink Customer Experience Improvement Program?

The Nextthink Finder and Engine are supported with the CEIP program.

Customer Success Program (CSP)

The goal of the Nextthink Customer Success Program (CSP) is to ensure the optimal operation and best possible experience with the Nextthink product. It reports data about appliances health, Nextthink product versions and usage.

The participation in the program is enabled by default. Customers wishing to opt-out should contact Nextthink Customer Success (customer.success@nextthink.com) to update their license. The customer will receive a license notification when the change is effective; after this notification the license has to be updated on the Portal to apply the changes. The opt-out will be visible on the Portal license management view.

Updating from V6.x

Updating the Appliance

When there is a system update or a new release of the Portal, the Engine or the Web Console, update all your Nexthink Appliances to take advantage of the new features, bug fixes and security updates.

Online update (recommended)

If your appliances have access to the Internet, this is the recommended method to do the update. Follow the procedure below to update each appliance:

1. Log in to the Web Console of the Appliance as administrator. In your browser, type the URL `https://<appliance.dns.or.ip>:99`.
2. In the section **Appliance**, select the tab **Update**.
3. Optional: Click the circular arrows in the **Last check for update** row to see if there is a new system update or any update of the installed Nexthink components: Portal, Engine or Web Console. If there is any update available, it is displayed in the cell on the right hand side. For each released component, find here a direct link to its release notes.
4. Optional: Check the box **Enable** of the **Automatic update** row to get the updates from the Nexthink repository as soon as they are published.
5. Optional: Press the button **Start connectivity test** to verify your connection to the Nexthink repository (`updates.nexthink.com`). If the repository is reachable, a message of success is displayed.
6. Click the button **Start update** to trigger the update procedure. You can follow the status of the update in the **Last update installation** row. At the end of the procedure, the message **Update complete** is displayed at the top of the tab.

Some updates require rebooting the Appliance to be complete. Refer to the chapter on rebooting the Appliance below for more information.

Offline update

The Appliance relies on *yum* to manage the upgrade of its components. When the appliance is connected to the Internet, the Web Console instructs the yum utility to get the upgrades from the Nexthink repository. In the case that your appliances are not connected to the Internet, you must download the offline

migration package and, if there is any system update, the Appliance ISO. You must then manually update the Appliance using yum from the command line.

The Appliance ISO contains the operating system, auxiliary packages and security updates of the Appliance; whereas the offline migration package is a tgz file that holds the Nexthink components: Portal, Engine and Web Console. For updating each one of your appliances offline, follow the steps below.

Applying system updates

To manually update the system packages of each Appliance, using yum and the Appliance ISO:

1. Attach the Appliance ISO to the physical or virtual system that hosts.
2. Log in to the command line interface (CLI) of the Appliance.
3. Mount the ISO with the following commands:
4. Update the system packages (ignore any message about already installed packages):

```
sudo mkdir -p /media/cdrom
sudo mount -t iso9660 /dev/cdrom /media/cdrom
sudo rpm -Uvh /media/cdrom/CentOS/centos-release-*.rpm
sudo yum --disablerepo=* --enablerepo=c7-media --nogpgcheck \
--exclude=nxconsole update
```

5. Wait for the operation to finish and then disconnect the ISO from the system using the following command:

```
sudo umount /media/cdrom
```

If the system updates include a modification of the kernel of the operating system, you need to reboot the Appliance to load the new kernel. Refer to the chapter on rebooting the Appliance below.

Updating Engine, Portal and Web Console

To manually update the Nexthink components of each Appliance:

1. Connect to the Appliance to update with your favorite SCP client and copy the offline update package (tgz file) to `/home/nexthink/`. Make sure that you copy the offline *update* package and not the offline *installation* package. The latter is designed for a clean install only, not for an update.
2. Untar the offline migration package:

```
tar -xzvf Nexthink-offline-update-6.x.tgz
```

3. Run the installation script:

```
sudo ./install_Nexthink_v6.sh
```

4. Log in to the Web Console as administrator.
5. Check that the update was correctly completed by verifying the versions of the installed components in the **Information** tab of the **Appliance** section.

Rebooting the Appliance

Usually, you do not need to reboot the Appliance after an update. In the case of system updates that install a new kernel for the operating system, however, it is necessary to reboot the Appliance to load the new kernel. This condition will be made clear in the release notes of the update.

To reboot the Appliance after an update:

1. Log in to the Web Console as administrator.
2. In the **Appliance** section, go to the **Reboot** tab.
3. To the question **Are you sure you want to reboot the Appliance?**, answer by clicking the **Yes** button.

Updating the Collector

It is possible to update the Collector on your Windows and Mac devices either manually or using your preferred deployment tool. Alternatively, Nextthink provides in addition the Updater tool for keeping your Collectors up to date on Windows devices.

The procedures below apply to updating the Collector on Windows devices with the help of the Updater.

Note that the Updater V6 is not backwards compatible with the Collector V5 or below.

Applies to platforms:

Configuring the Updater from the Web Console

To deploy a new version of the Collector using the Updater, upload the Collector package to the Appliance on which the Engine runs with the help of the Web Console:

1. Open the Web Console in a web browser:
`https://<appliance_ip_address>:99`

2. In the **Engine** tab, select the **Updater** entry from the left-hand side menu.
The configuration of the Updater parameters in the Engine shows up.
3. Tick the **State** check box and set it to **Active** to enable the deployed instances of the Updater to retrieve the Collector from the Engine.
4. Type in the port number configured in your Updaters to establish communication with the Engine in **Updater TCP port** (default 8888).
5. Optional: Tick the **SSL** check box and set it to **Active** if you want the communication between the Updater and the Engine to be encrypted.
Note that the use of SSL requires the instances of the Updater to have been deployed with the appropriate option.
6. The four following options are parameters that you can pass to the Collector MSI to control the way it is installed in the end-user devices:
 - ◆ Tick **Force loading** to load the Collector in memory as soon as it is installed (without waiting for a reboot).
 - ◆ Tick **Installation log** to send the logs to the Engine in case of error during the installation.
 - ◆ Tick **Add/Remove program entry** for the Collector to be displayed in the section Add/Remove Programs of Windows.
 - ◆ Tick **Control panel extension** to install a program for configuring the working parameters of the Collector in the Control Panel of Windows.
7. The next three display the versions of the software components related to the Collector that the Engine holds: **Driver version**, **Service version** and **Updater version**.
8. In **Collector selection**, click the button **Browse...** to select the zip package of the Collector. This is the zip file that you find in the Product Downloads page of Nextthink.
9. Click the button **Save** to make your changes permanent.

For keeping the Collector up to date by other means (individual install, or deployment via GPO or SCCM), see the article on installing the Collector.

Managing the deployment of the Collector from the Finder

When using the Updater to keep the Collector up-to-date, you can control the deployment process and watch its evolution from the Finder:

1. Log in to the Finder as a user with *system configuration* permissions.
2. In the left-hand side accordion, select the **Nextthink Collector** section.
3. Inside the **Nextthink Collector** section, find the two tabs **Manage** and **Discover**.
 - ◆ The **Manage** tab lets you control the deployment process and see the status of each device with respect to it.

- ◆ In the **Discover** tab, you can create collections of devices from the Active Directory or from imported CSV files. Use these collections to keep track of the devices that have the Collector installed and those that do not have it.

Keep in mind that a deployment with the Updater is a per Engine process. You deploy the Collector on the devices that lie under the supervision of a single Engine.

Discovering devices

A device that does not have the Updater or the Collector installed never sends information to the Engine. Therefore, the Engine ignores the existence of the device. To inform the Engine about all the devices in your network, including those that the Engine may not be aware of, create collections of devices in the **Discover** tab. You can create collections of devices based on:

- The information in Active Directory.
- The contents of a CSV file.

To create a collection of devices based on Active Directory, make sure first that you have configured the Active Directory server settings in the Engine. Then, in the **Discover** tab:

1. Right-click the title or the empty area of the tab.
2. Select the option **Create collection from AD...** in the context menu.
3. Review how to locate your devices in the Active Directory and fill in the blanks in the dialog:
 1. Set a name for the collection in the field **Name**.
 2. Write in the field **Include DN** a query pattern to retrieve all the devices whose Distinguished Name matches the pattern. You can use the wildcards *, to substitute for zero or more characters, and ?, to substitute for one character, in your query.
 3. Optional: If your query pattern above includes some devices (or other AD objects) that you want out of the collection, specify them in **Exclude DN** with another query and tick the check box to the left to activate the exclusion.
4. Click OK to create the collection.

If you do not have Active Directory available to your Engine, but you have other means to get a list of all the devices in your network, you can still create a collection of devices in the **Discover** tab by providing a CSV file. The CSV file must hold at least two values per entry:

- The NetBIOS name of the device.
- The IP address or DNS name of the device.

To create a collection from a CSV file:

1. Right-click the title or the empty area of the **Discover** tab.
2. Select the option **Create collection from CSV...** in the context menu.
3. Choose a CSV file from your filesystem in the dialog that opens. A wizard guides you through the import of the CSV.
4. In step 1 of the wizard:
 1. Select the encoding, the delimiter character and the text qualifier (character used to delimit text values) of the CSV file.
 2. Optional: Click **Show file** to see the actual CSV file and help you decide what are the correct options.
 3. Click **Next**.
5. In step 2 of the wizard:
 1. Give a name to the collection that you are creating in the field **Collection name**.
 2. In **Column selection**, pick the two columns from the CSV file that hold the Netbios name of the device and the IP address or DNS name (hostname). To guide you with the selection, the values of the first entry in your CSV file are displayed in the lists.
 3. Optional: Click **Back** to correct the options that you chose in step 1 of the wizard if you realize that you set something wrong.
 4. Click **Import**.
6. The wizard reports the number of devices successfully added to the collection from the CSV file. In case of error, click **Show details** to see the reasons for not importing all the entries from the file.
7. Click **OK** to end the wizard.

In the **Discover** tab, every collection of devices displays its total number of devices to the right of its name. Additionally, each collection is divided into two disjoint groups of devices that also show their number of devices:

- **Without Collector**: those devices that do not have the Collector installed.
- **With Collector**: those devices that have the Collector installed.

To get a list of the devices in the collection or in any of the groups, double-click the collection or the group in the **Discover** tab. The groups get updated at the same time as the Engine detects if the Collector is installed in or uninstalled from the devices in the collection.

Controlling Collector deployment

In the **Manage** tab, you can control the update of the Collector and follow the progress of the installation for the machines that have the Updater installed. The **Manage** tab is divided into two sections: **Devices** and **Settings**".

The **Devices** section groups devices according to their deployment status. To see the list of devices that form the group, double-click the name of the group. Click the plus sign at the left of each main group to expand it into subgroups and the minus sign to collapse subgroups back. Find below a list with all possible status and their meaning:

- **Manageable**: All devices that have the Updater installed.
 - ◆ **Up-to-date**: Devices with the last version of the Collector (the last uploaded to the Engine) installed.
 - ◆ **Outdated**: Devices whose version of the Collector is older than the last uploaded.
 - ◇ **Out-of-date**: Outdated devices for which no action has been taken.
 - ◇ **Scheduled for update**: Outdated devices that have been marked for update, but whose Updater has not started yet the process of upgrading the Collector.
 - ◇ **Updating**: Outdated devices whose Updater is either downloading the last Collector, upgrading it or waiting for a reboot to complete the installation.
 - ◆ **Uninstalling**: Devices for which you want to remove both the Collector and the Updater.
 - ◇ **Scheduled for uninstall**: Devices marked for unistallation of the Collector and the Updater, but whose Updater has not started yet the process.
 - ◇ **Uninstalling**: Devices that are currently uninstalling the Collector or the Updater, or which are waiting for a reboot to complete the uninstallation.
 - ◆ **Repairing**: Devices in which you want to reinstall the Collector, usually because there was an error during its installation.
 - ◇ **Scheduled for repair**: Devices that are marked for reinstalling the Collector, but whose Updater has not started the repairing process yet.
 - ◇ **Repairing**: Devices whose Updater is uninstalling the Collector or waiting for a reboot to complete the uninstallation.
 - ◆ **With error**: Devices that had trouble downloading the Collector or installing it.

- **Unmanageable:** Devices that have the Collector installed, but not the Updater. You cannot control the Collector deployment process from the Finder on these devices.

Right-click the name of a group to trigger appropriate actions on all the devices of that group. For instance, right-click the group **Out-of-date** and select **Update all** from the context menu to start the process of updating the Collector in all the members of the group. As a result, all devices in the **Out-of-date** group move immediately to the **Scheduled for update** group.

To perform an individual deployment action on a particular device, double-click a group where the device can be found to open the list of all the member devices of the group in a tab of the Finder. From the list of devices, right-click the entry of the specific device and select **Manage Collector** from the context menu. A list of the possible deployment actions for the device appears to the right of the menu. Select the desired option. For instance, you can choose **Schedule for repair** if the device is in the group **With error**.

In the **Settings** section, find two options that modify the way of deploying the Collector:

- Auto-update
- Reboot policy

Tick the **Auto-update** option if you want the system to trigger the update of the Collector automatically for all the manageable devices whose Collector is out of date. Using this option, you no longer have to manually schedule for update any of your out-of-date devices. Nevertheless, devices that had an error during Collector installation are not automatically updated. You still have to schedule for repair devices with errors by hand, after possibly doing some research on the causes of the error.

Click the button **Set reboot policy...** to decide what the devices should do after a Collector update:

1. Choose one of the options in the list **After update**:
 - ◆ **no automatic reboot:** Wait until the user reboots its computer to complete the installation.
 - ◆ **reboot automatically:** Order a reboot of the device, but wait until all open applications are closed.
 - ◆ **reboot automatically forcibly closing applications:** Reboot the device, forcing open applications to close.

2. If you select one of the two latter options, specify the moment for rebooting the devices in the section **Reboot time**. Choose between:
 - ◆ **Immediate**: Reboot the devices right after they update the Collector.
 - ◆ **Scheduled at xx:xx:xx (local time)**: Reboot the devices at the specified time of the day, once they have updated the Collector.
3. Click **OK** to apply the reboot policy.

Limitation when uninstalling the Updater

When an Updater sends a Collector update request to its associated Engine, the device which is running that Updater automatically becomes **Manageable**. That is, the Engine knows that the Updater is installed in a device when it receives at least one Collector update request from that particular device. You can see that device belonging to the group of **Manageable** devices when controlling the deployment of the Collector from the Finder.

However, if the Updater is later uninstalled from the device, the Engine has no way to realize that the device is no longer manageable. Therefore, a device can indefinitely stay in the **Manageable** state long after the Updater has been removed.

Related tasks

- Installing the Collector
- Importing data from Active Directory

Related references

- Collector MSI parameters reference table

Updating the Finder

Whenever you log in to the Finder, the Finder detects if there is an incompatibility between its own version and the version of the Portal to which it connects. In such a case, the Finder notifies the availability of a new update with a pop up dialog at the moment of establishing the connection. The dialog includes an appropriate message and two links: one link points to the release notes and the other link to the installer file (MSI) of the new version of the Finder.

To update the Finder:

1. Log in to the Finder with your own account. If a new update is ready for download, a dialog shows up.
2. Optional: Click the link **Click here to find out what's new...** to read the release notes of the new version.
3. Click the link **Click here to upgrade...** to start the update procedure. This opens the Finder downloads page in your default web browser.
4. Choose the appropriate version of the Finder for your computer (32 or 64-bit):
 - ◆ Click the **exe** installer if you also need to install the .NET Framework 4.5.
 - ◆ Click the **MSI** installer if you already have the .NET Framework 4.5 installed.
 - ◆ In case of doubt, download the MSI. If the Finder does not start because you do not have the .NET Framework 4.5 installed in your computer, download and install the .NET Framework 4.5 later.
5. Save the installer file to your disk.
6. Optional: Uninstall the current version of the Finder from your Windows system. If you omit this step, running the new installer removes the old version anyway.
7. Run the MSI file to install the new version of the Finder.
8. Follow the steps in the installation wizard.

If the update is not strictly necessary, because the currently installed Finder is still compatible with the Portal to which you are connecting, you can click one of the following buttons to continue working without doing the update:

- **Remind me later**, to display the dialog again the next time that you try to connect.
- **Do not show again**, to prevent the update dialog from showing up.

In the case that you choose not to go on with the update, you can still bring back the dialog from the main help menu of the Finder:

1. Click the question mark at the top-right corner of the Finder window.
2. Select **Upgrade Nextthink Finder...** and the update dialog shows up again.

Configuration

Setting up a software license

Once you have installed all the appliances (the Portal and one or more Engines), you require a software license to make the whole system work properly. Before requesting a license, make sure that you have the following information readily available:

- Total number of devices to monitor (Windows and Mac OS).
- Total number of Mobile devices to survey.
- Number of servers.
- Validity period of the license:
 - ◆ Start date.
 - ◆ Expiration date.
- Desired optional modules, along with their own validity period (start and end dates):
 - ◆ Integration toolkit.
 - ◆ Web and Cloud.
 - ◆ Security.
- Type of license required:
 - ◆ Online.
 - ◆ Offline.

The validity period of the optional modules cannot exceed the validity period of the license.

To request more information about the licensing process or for any question related to the license of your specific setup, please send an email to:

Determining the activation mode of your license

The activation of the license depends on the connectivity of your appliances to the Internet.

- Request an **online** license if your appliances connect to the Internet. Online licenses are easier to set up and more flexible for updating than offline licenses. Nextthink recommends to use an online license whenever possible.

- Request an **offline** license only when your appliances cannot connect to the Internet.

Ordering and activating a new license

Once you issue a Purchase Order, you will receive a new license activation key from the Sales operations department of Nextthink by email.

To activate your new license:

1. Open the Portal and log in as central administrator.
2. In the **ADMINISTRATION** menu, select **Licenses** under the **SYSTEM CONFIGURATION** section.
3. Enter the activation key.
4. Select your mode of activation for the license:
 - ◆ Online activation.
 - ◆ Offline activation.
5. Allocate the number of licensed end-user devices (Windows and Mac), servers and Mobile devices among your Engines.
6. Finish the activation of the license:
 - ◆ If you requested an online activation of the license:
 1. Click **Apply** and you are done.
 - ◆ If you requested an offline activation of the license:
 1. Click **next** to go on with the activation.
 2. Click **Download license file** to get an encrypted file holding your license information.
 3. Go to <https://sign-license.nextthink.com/> to get your signed license file
 4. Upload the signed license file in Portal

Updating a modified license

When a license is modified in the Central License Manager, an automatic email notification is sent:

- For online license, the process is automatic.
- For offline license, the new license file has to be manually uploaded to the Portal.

More information can be found with the description of email notification for license modification.

Concurrent management of the license

Beware that more than one central administrator may access the license management dashboard in the Portal at the same time. In this case, the Portal displays appropriate warning messages if the concurrent modification of the license might lead to inconsistencies.

Specifying your internal networks and domains

Specify the fully qualified domain name of the Engine and the address where the Engine can find the Portal.

Additionally, to help the Engine make the difference between network traffic inside your organization and network traffic destined to external entities, specify your internal networks and domains from the Web Console.

Specifying the DNS name of the Engine

To specify the fully qualified domain name (DNS name) of the Engine:

1. Log in to the Web Console as admin.
2. Click the **Engine** tab at the top right corner and select **Parameters** from the left-hand side menu.
3. Type in the DNS name of the Engine in the entry **Engine DNS name** (e.g. *myengine.example.com*).
4. Click **Save** to store your changes.

Connections to the Engine through the Web API use the configured DNS name of the Engine for communication. For the Web API to work, this value must be correctly set. If the Engine does not have a DNS name, type in its IP address instead.

The Updater also needs the DNS name of the Engine to be correctly set for retrieving new versions of the Collector when available.

Specifying the address of the Portal

The Engine needs to know where it can find the Portal in order to get licensing information and send real-time services data to it. To specify address of the Portal:

1. Log in to the Web Console as admin.
2. Click the **Engine** tab at the top right corner and select **Parameters** from the left-hand side menu.
3. Type in the DNS name or IP address of the Portal in the entry **Portal address**.
4. Click **Save** to store your changes.

Specifying the internal networks

To specify the subnetworks that the Engine must recognize as belonging to your organization:

1. Log in to the Web Console as admin.
2. Click the **Engine** tab at the top right corner of the page and select **Internal networks & domains** from the left-hand side menu.
3. Click the plus button at the right of the table entitled **Internal network configuration** to add a new internal network to the table.
4. For each one of your internal IP networks, specify:
 - ◆ The subnetwork base address in the column **Network**.
 - ◆ The subnetwork mask in the column **Mask**.
5. Repeat the operation for as many internal networks as you need to specify.
6. Optional: Click the button with an **e** inside at the right of the network entry in the table to edit its contents.
7. Optional: Click the button with an **x** inside at the right of the network entry in the table to remove the entry.
8. Click **Save** to make your changes permanent and restart the Engine (or wait until you have finished configuring your internal domains).

Specifying the internal domains

Specifying the internal domains is only useful if you have purchased the Web and Cloud module. You need to write down only those domains that are hosted in servers outside your internal networks, so they are still considered *internal* web traffic even though they can be managed by an external organization. Domains served from your internal network are naturally considered internal.

The Engine never compacts domains identified as internal and it never sends these domains to the Application Library for detecting threats, since they are trusted.

To specify your internal domains:

1. Log in to the Web Console as admin.
2. Click the **Engine** tab at the top right corner and select **Internal networks & domains** from the left-hand side menu.
3. Write down the list of domains inside the text box under the title **Engine internal domains** at the bottom of the page. Use the wildcards **?** and ***** to replace one or several characters of the domain name and separate each domain in the list by a space. For instance:
***.example.com *.nextthink.com *.nextthink.ch**
4. Click **Save** to make your changes permanent and restart the Engine.

Related tasks

- Reporting the URL of HTTP web requests

Allocating resources for the Portal

Adapt the configuration of the Portal to your available hardware resources in order to maximize their utilization and optimize performance. To that end, edit the configuration file `startup.properties` in the Portal and set the appropriate running mode and memory options depending on your hardware resources and size of your installation.

Find the suitable running mode and memory settings for your installation size in the table below. The **Total memory** size corresponds to the actual memory installed, according to the specified hardware requirements for the Portal appliance. If you are using a single appliance for both the Engine and the Portal, divide the memory installed by two (hardware requirements for a single appliance) to get an estimation of the **Total memory** for the Portal:

Max devices	Total Memory	Configuration
500	6 GB	MODE=SMALL SMALL_MEMORY=4G
5k	8 GB	MODE=SMALL SMALL_MEMORY=6G
10k	8 GB	MODE=SMALL SMALL_MEMORY=6G

20k	12 GB	MODE=MEDIUM MEDIUM_UI_MEMORY=8G MEDIUM_INFRA_MEMORY=2G
50k	16 GB	MODE=MEDIUM MEDIUM_UI_MEMORY=12G MEDIUM_INFRA_MEMORY=2G
100k	32 GB	MODE=MEDIUM MEDIUM_UI_MEMORY=27G MEDIUM_INFRA_MEMORY=3G
150k	48 GB	MODE=MEDIUM MEDIUM_UI_MEMORY=40G MEDIUM_INFRA_MEMORY=4G
>150k	Ask	MODE=LARGE Ask

To edit the configuration file of the Portal and set the values that fit your hardware:

1. Log in to the CLI of the Portal appliance.
2. Stop the Portal:

```
sudo systemctl stop nxportal
```
3. Make a copy of the sample configuration file to use it as the current configuration file:

```
sudo cp /var/nexthink/portal/conf/startup.properties.sample \
/var/nexthink/portal/conf/startup.properties
```
4. Edit the configuration file with the appropriate values from the table above:

```
sudo vi /var/nexthink/portal/conf/startup.properties
```
5. Restart the Portal:

```
sudo systemctl start nxportal
```

For example, in an installation with 45 000 devices, look up in the table above and find that, for a maximum of 50k devices, you must set the running mode of the Portal to MEDIUM, and allocate 12 GB of memory for the user interface, in addition to 2 GB more for the infrastructure.

In that case, change the values of those parameters in the `startup.properties` configuration file of the Portal. The file should look like this:

```
# allowed modes are SMALL, MEDIUM, LARGE
MODE=MEDIUM
```

```
SMALL_MEMORY=2G
```

```
MEDIUM_UI_MEMORY=12G
MEDIUM_INFRA_MEMORY=2G
```

```
LARGE_UI_MEMORY=16G
LARGE_INFRA_MEMORY=4G
LARGE_COMM_MEMORY=4G
```

For LARGE installations, please contact Nextthink for instructions on how to properly allocate resources for the Portal. You may also need to increase the number of connections to the Portal database.

Related references

- Hardware requirements
- Support FAQ: Maximum number of connections for PostgreSQL

Connecting the Portal to the Engines

For the Portal to compute and display data in its widgets, you must connect it to the Engines that receive and organize the end-user data coming from the Collectors.

To connect the Portal to an Engine:

1. Log in to the Portal as central administrator.
2. In the top menu **ADMINISTRATION**, select the **Engines** dashboard under the section **SYSTEM CONFIGURATION**.
3. Click the plus sign that is located in the top right corner of the widget **Engines Management**. The dialog to add a new Engine shows up.
4. Type in the IP address or DNS name of the appliance that hosts the Engine in the **Address (IP or hostname)** field.
5. In the **Port** field, type in the port number that the Engine uses to communicate with the Finder and the Portal.
6. Optional: In the field **Description**, write down a brief sentence to help you distinguish the new Engine that you want to connect to the Portal from other Engines.
7. Click **Ok**.

After completing the procedure, the **Engines** dashboard displays the new Engine as a row in the table of connected Engines. The row displays the name, address, description, version and timezone of the Engine. However, since the connection is not yet established, a red dot appears in the first column of the row and the actual name, version and timezone of the Engine are not available yet.

To establish the connection and get information from the Engine:

1. In the table of Engines find the chain and pencil icons that are placed to the right of the **Name** column.
2. Click on the chain icon to establish the connection with the Engine. The red dot turns to yellow and then to green, to indicate that the connection is now established. The widget fills in the name, version and timezone of the Engine with the information that the Engine itself sends.

Repeat the operation described above for any other Engine that you want to connect to the Portal.

Once the connection is established, the Portal collects information from the Engine in a regular basis. While the Engine connection is working, you cannot edit the parameters of the Engine and so the pencil icon in the row that holds the information about the Engine is disabled. Otherwise, if the Portal cannot establish a connection to the Engine, the dot in the beginning of the row stays red, which means that you probably did not set the parameters of the Engine correctly. In this case, click the pencil icon, edit the connection parameters of the Engine as explained above and try to establish the connection again by clicking the chain icon.

Similarly, if the appliance of an Engine changes its configuration and the modifications make the connection with the Portal fail, the dot will turn to red as well. To recover the connection:

1. Click the chain icon (displayed as a broken chain now) to unlink the Engine from the Portal and be able to edit the modified parameters.
2. Edit the parameters of the Engine as we just explained above.
3. Click again the chain icon and wait for the red dot to turn to green.

Centralized Management of Appliances and Engines

Appliances and Engines can be managed individually from the Appliance Web Console. If you have several appliances, the Centralized Management solution will allow you to configure them centrally.

To access the Centralized Management of Appliances in the Portal, in the **ADMINISTRATION** top menu, select the **Appliances** dashboard under the section **SYSTEM CONFIGURATION**.

Configure Appliances and Engines from here, performing an action on all the Appliances or Engines that you select.

Centralized Appliance Management

Enable Appliance for centralized management

Connect to the Appliance Web Console with the admin user as shown in Installing the Appliance. Enable the Appliance Remote Account of the Appliance by checking the **Enable** option of the Account status. The **Set new password** field is not mandatory. If no value is entered, the default password is **api**.

Adding an Appliance for centralized management

In the Appliances section:

1. Click on the plus icon
2. Enter the Appliance **Name**, **IP Address** and **Port**, the **Remote management password** and **Description**. Note that in the **IP address** field, one should enter either the IP address or the DNS name in order to match what was entered in the **Address** field of the Engines dashboard in the Nextthink Portal section (refer to Connecting The Portal To The Engines). Validate your entry to add the appliance to the centralized management.

Central configuration of Appliances

Select the Appliances to configure and press on the configuration icon. You have several configuration options:

Enable Web Console

Allow the connection of users to the Web Console in the Appliance.

Disable Web Console

Prevent users from connecting to the Web Console in the Appliance.

Configure SMTP

Nextthink uses the Simple Mail Transfer Protocol (SMTP) for sending reports and alerts to a mail server for relaying to the Nextthink administrator. Enter or update the SMTP settings requested.

Configure NTP

Nextthink synchronizes with NTP servers using the Network Time Protocol. Enter the list of NTP server addresses separated by a space. One option is to use the NTP servers offered by <http://www.pool.ntp.org/en/>, such as 0.pool.ntp.org, 1.pool.ntp.org, etc.

Update

Check and update the Nextthink Appliance.

Reboot

Reboot the machine that hosts the Appliance.

Add or remove Engines related to an Appliance

1. Click on the information icon corresponding to the Appliance of interest.
2. Click on the chain icon to start or stop the central management of an Engine. If the central management is started, the Engine will show up in the Engines section of the Appliances dashboard.

Edit or remove an Appliance from centralized management

- To edit an Appliance, click on the corresponding pencil icon
- To delete an Appliance, click on the corresponding trashcan icon

Centralized Engine Management

The Engines section shows the Engines associated to the Appliances managed centrally. Select the Engine to configure and press on the configuration icon. You have several configuration options:

Configure LDAP

Set up your LDAP servers to get Active Directory information for Nextthink objects.

Information

Displays a table with the different configuration settings and general information on the Engine.

Refresh DNS

Use this option to refresh DNS information on the Engine. This can be necessary if you wish to reflect changes on a DNS server (configuration changes or updates in the resolution of a particular destination). The Engine resolves new Destinations, but it does not refresh their DNS automatically if it changes.

Refresh LDAP

This option is generally used in a scenario where the LDAP server integration is performed after Engine installation. In this case, launch this option to trigger the Engine refresh of its LDAP information. The Engine gets information from the configured LDAP servers on every new user detected.

Restart

Stop and restart the Engine. Note that restarting the Engine results in a temporal loss of data received from Collectors during the time of starting up.

Adding users

Overview

Right after installation, the only user that exists in the system is the first and main central administrator or *admin* user. The admin user has unrestricted access to all data available in both the Portal and the Finder. Moreover, the admin user is able to create and modify all kinds of content in the system, including dashboards, investigations, categories, alerts and user accounts.

Incidentally, you may want to give other people the chance to log in to the system and use it without necessarily having all the capabilities of the admin user. The admin user can thus create accounts for other users, restrict their views on the data and limit their ability to alter content. In this section, learn how to add users to the system and control their access to the data recorded.

Prerequisites

Before defining new profiles and users, ensure that you have installed a license for the product. Otherwise, some configuration pages will not show up.

Account update considerations

Beware that changes to accounts and their permissions may not take immediate effect on logged in users.

For users logged in to the Finder or to the Portal, the user keeps the permissions before the change during the session lifetime. For users making use of Web API (NXQL), the old permissions are still in force up to five minutes after the change, until the Engine synchronizes account information with the Portal.

Defining user roles

The *roles* attributed to a user determine how the user interacts with the system. The tasks of the users of the system depend on their responsibilities. Roles let you group the elements that allow users to carry out the tasks that are assigned to them. Using roles, you can specify the modules that the users playing that role can see in the Portal, the investigations that they are able to run in the Finder, and the alerts that they must be aware of.

To incorporate elements into a role, first create these elements in the Finder. It is not essential to have all the elements ready before defining a role. You can start by creating the role and edit it later to add the missing elements.

To define a new role:

1. Log in to the Portal as administrator .
2. Click the **ADMINISTRATION** drop-down menu at the top of the window.
3. Select the option **Roles** to open the dashboard for editing roles.
4. Click the plus sign at the top right hand side of the dashboard to open the wizard for adding a new role.

Step 1: Adding modules

1. Type in the name of the new role in the **Name** field.
2. Optional: Click **Add module** to add an existing module of the Portal to the role. A dialog to choose the module pops up.
 1. Select a module from the list labeled **Module**.
 2. Click **Add**. The dialog closes and the selected module is added to the **Modules** list of the role.
3. Repeat the previous step to add as many modules as the role needs.
4. Click **Next** to go on with the next step of the wizard.

Step 2: Adding service-based alerts

1. Optional: Click **Add alert** to include service-based alerts to the role. A dialog to specify the alerts pops up.
 1. Select a service-based alert from the list labeled **Alert**.
 2. Optional: Click **yes** in the **Mandatory** section to force the subscription to the alert of all users with the current role. By default, the alert is not mandatory.
 3. Click **Ok**.
2. Repeat the previous step to add as many alerts as the role needs.
3. Click **Next**.

Step 3: Adding investigations

1. Optional: Click **Add investigation** to share existing investigations with all users who have the current role assigned. A dialog to specify the investigation pops up.
 1. Export an investigation or a folder of investigations from the Finder to the clipboard.
 2. Paste the contents of the clipboard on the dialog of the wizard.
 3. Click **Add**. The dialog to paste the investigation closes and the investigation is added to the **Investigations** list of the role.
2. Repeat the previous step to add as many investigations as the role needs.

Step 4: Adding one-click investigations

1. Optional: Export a pack with all the one-click investigations that you want to add to the role from the Finder.
 1. Paste the pack of one-click investigations on the dialog of the wizard.
2. Click **Next**.

Step 5: Adding investigation-based alerts

1. Optional: Click **Add alert** to include investigation-based (Finder) alerts to the role. A dialog to specify the alert pops up.
 1. Export an alert or a folder of alerts from the Finder to the clipboard.
 2. Paste the contents of the clipboard on the dialog of the wizard.
 3. Click **Add**. The dialog to paste the alert closes and the alert is added to the **Alerts** list of the role.
 - ◇ The syslog notification mechanism of global alerts is local to the Engine where the global alert was created and, therefore, not propagated to other Engines via roles. If you add a global alert with syslog notification enabled to a role, only the email notification mechanism is propagated to the users with that role.
2. Repeat the previous step to add as many alerts as the role needs.
3. Click **Finish** to end the wizard. The new role is added to the list of the **Roles** dashboard.

Defining user profiles

The *profile* of a user defines the type of user, the access rights of the user to the different domains of a hierarchy (both as a viewer and as administrator, if applicable) and to the functions of the Finder. Moreover, you can associate one

or multiple roles to a profile. Thus, users are able to play any of the roles associated to their profile, along with any other possible role that you may additionally assign to them.

Profile types

There are three main types of profiles:

User

This profile is intended for users that only have the right to view the information; both in the Portal and, optionally, in the Finder. They are able to see only the data that belongs to their view domain (a subset of the available hierarchies), possibly limited by privacy settings as well.

Administrator

In addition to viewing information, users with an Administrator profile can publish modules and manage Portal components. The view domains and administrative domains of an Administrator may be different for each hierarchy. Optionally, Administrators can have the right to create other user accounts.

Central administrator

This profile defines a special kind of Administrator who has access to all hierarchies and the right to create and modify profiles and hierarchies. Only a central administrator can configure the license and the connections to the Engines.

See here the complete matrix of access rights and permissions.

To create a new profile:

1. Log in to the Portal as central administrator.
2. Click the **ADMINISTRATION** drop-down menu at the top of the window.
3. Select the option **Profiles** to open the dashboard for editing profiles.
4. Click the plus sign at the top right hand side of the dashboard to add a new profile. The wizard to add a new profile opens.

Step 1: Choosing the type of account

1. Type in a name for the new profile in the field labeled **Profile name**.
2. Select one of the three types of accounts from the choice **Account type**.
 - ◆ Select **User** if the profile is intended for users without administrative tasks.
 - ◇ Optional: Uncheck the box **Allow creation of personal dashboards** to prevent users with the current profile from

creating their own modules and dashboards. By default, the box is checked, allowing the users to create Portal content.

- ◆ Select **Administrator** if the profile is intended for users with administrative tasks.
 - ◇ Optional: Uncheck the box **Allow creation of user accounts** to prevent the administrative user from creating new accounts. By default, the box is checked, allowing administrators to create new accounts.
- ◆ Select **Central administrator** to create users that can administer the whole system in the same way as the admin user, except for the fact that you can restrict what they see in their data privacy settings.

3. Click **Next** to go on with the next step of the wizard.

Step 2 (Administrator profile only): Select administration domain

If you selected to create a profile of the type Administrator in the previous step, set now its administration rights. If you chose to create a profile of the type Central administrator instead, this step is skipped, since the administration rights of central administrators are not limited.

1. In the field **Administration domain**, select a hierarchy in the first field, a level in the second and a node in the third. The profile will have administrative rights over the selected node and all the nodes below it in the hierarchy.
 - ◆ Leave the top node of the administration domain undefined by choosing **---parameter---** from the list. Define the top node of the administration domain individually for each user when creating their user account.
2. Optional: If you checked the option **Accounts creation** in the previous step, select now the profiles that an Administrator with this profile can assign to the user accounts that he creates. Use the **Ctrl** key while clicking the names of the profiles in the list **Profiles available for accounts creation** to select more than one profile. If no profile is selected, the Administrator will not be able to create user accounts.
3. Click **Next**.

Step 3: Set privacy settings, roles and view domain

This is step 2 if you chose a user or central administrator profile.

1. Select the **Data privacy** settings for the profile:

- ◆ **anonymous users, devices, destinations and domains**: user accounts with this profile cannot see the names of users, devices, destinations, or domains.
 - ◆ **anonymous users and devices**: user accounts with this profile can see neither the names of users nor of devices.
 - ◆ **anonymous users**: user accounts with this profile cannot see the names of users.
 - ◆ **none (full access)**: user accounts with this profile have full access to the collected data.
2. Select the roles of the profile by clicking their name in the **Role(s)** list. Use the **Ctrl** key to select several roles at the same time. The investigations, alerts, modules, etc attributed to the selected roles are inherited by the profile.
 3. Specify the view domain of the profile for each defined hierarchy. Users with the current profile can only view the objects grouped in the specified domain:
 1. In the **from** field, select the highest level in the hierarchy that belongs to the view domain.
 2. In the **Node** field, either:
 - ◇ Choose the top node of the view domain from the available nodes of the level. This node and all the nodes below it belong to the view domain, down to the level specified in the next step.
 - ◇ Leave the top node undefined by choosing **--parameter--** from the list. Define the top node of the view domain individually for each user when creating their user account.
 3. In the **to** field, select the lowest level in the hierarchy that belongs to the view domain.
 4. Click **Next**.

Step 4: Set Finder access

This is step 3 if you chose a user or central administrator profile. If you want the users with the current profile to be able to access the Finder:

1. Check the box **Finder access**.
2. Select the time zone of the user.
3. Optional: Check the box **Allow edition of application and object tags** if you want the users with the current profile to be able to manually modify the tags of objects in the Finder.
4. Optional: Check the box **Allow system configuration** if you want the users with the current profile to be able to edit categories, services, metrics, and global alerts, as well as import and export content, or

manually synchronize users and devices with Active Directory. You can only select this option if you gave full access to the profile in the privacy settings of the previous step.

5. Optional: Check the box **Allow management of web API** for users with the profile to be able to publish and edit web API investigations in the Finder. The profile must have full access in the privacy settings, a top view in at least one of the hierarchies defined, and full Web & Cloud visibility (see below) to be eligible for the management of the web API.
6. Optional: Check the box **Allow management of Collectors** if you want the users with the current profile to be able to follow and control the deployment of the Collector from the Finder. Again, you can only select this option if you gave full access to the profile in the privacy settings of the previous step.
7. Optional: If you have purchased the Web & Cloud module, set the visibility level of the users with the current profile to **restricted** or **full** in the list under **Web & Cloud visibility**.
8. Click **Finish** to end the creation of the profile. The profile is added to the list of profiles in the dashboard.

Creating a user

After creating roles and profiles for users, finally create user accounts that make use of them:

To create a user account:

1. Log in to the Portal as central administrator, or as administrator with account creation rights.
2. Click the **ADMINISTRATION** drop-down menu at the top of the window.
3. Select the option **Accounts** to open the dashboard for editing accounts.
4. Click the plus sign in the top right corner of the dashboard. The wizard to create a new user account shows up.

Step 1: Setting personal data and profile

Nexthink supports user authentication both internally or through Active Directory.

1. Type in the name of the user:
 - ◆ To use internal authentication, type in the desired account (login) name of the user in the field **Username**.
 - ◆ To authenticate users through Active Directory, type in the **sAMAccountName** of the user followed by the @ character and the DNS domain name (e.g. `jwick@example.com`) in the field

Username. Note that this field is case sensitive. Therefore, the name of the Nextthink account must exactly match the sAMAccountName name in Active Directory.

2. Type in the complete name of the user in the field **Full name**.
3. Configure the email address for sending notifications to the user in the field **Email address**.
4. Type in a password for the user in the field **Password** and retype it in **Password confirmation**.
5. Select the profile of the user from the list **Profile**. The user gets all the permissions, default content and roles associated to the profile.
 - ◆ If the selected profile does not define a particular top node for the administration or view domains of the users with that profile (because one of the two domains or both are parameterized), select now the top nodes of those domains individually for the current user.
6. Optional: tick the check box **Never automatically sign out this account from Portal when active** if you want to override the session timeout control configured in the Portal and never log out the user from the Portal while active. Note that having a live view on a service keeps a user active even without actual user interaction.
7. Click **Next**.

Step 2: Setting additional roles

1. Optional: If you want the user account to inherit content from one or more roles that do not belong to its assigned profile, select the desired roles from the list **Additional roles**. Use the **Ctrl** key to select more than one.
2. Click **Ok** to end the creation of the user account. The account is added to the list of accounts in the dashboard.

Related tasks

- Controlling session timeouts in the Portal

Related references

- Access rights and permissions
- Active Directory Authentication

Hierarchizing your infrastructure

Overview

To manage the complexity of a big company or organization, you usually divide it into a set of hierarchical levels. You can build hierarchies according to different criteria. For instance, if a company is spread throughout several countries, it is possible to group parts of the organization according to their geographical location. You can then arrange the locations in a hierarchy of cities, regions, countries and even continents. Other possibility is to divide the company into functional departments, such as Research and Development, Human Resources, etc. and then divide each department into units, each unit into sub-units and so forth, until you are satisfied with the decomposition. Several hierarchies may be built for the same company and coexist within it at the same time.

Nexthink hierarchies let you arrange the devices in your IT infrastructure in a way that reflects the structure of your company, with the advantage of getting results from Nexthink that directly map into the existing structure. For instance, you can quickly detect if a problem impacted every device in your company or just the computers in the department of Human Resources. Break down results from investigations, dashboard widgets and IT services according to the defined hierarchies. In addition, use hierarchies to delimit scopes of visibility for users (view domains) and administration rights over parts of the company (administration domains).

Example of a hierarchy built with mixed functional and location criteria

Specifying entities

To organize your set of devices into a hierarchy, group your devices by *entities*. Entities are logical groups of devices that make up the first level of all hierarchies. Each device belongs to at most one entity, whose name is displayed in the special device field **Entity**. To assign entities to each device, write a Comma Separated Value (CSV) file that specifies the entity names and the rules

to assign an entity to groups of devices. The format of the CSV file is described in the next section.

To assign entities to sets of devices:

1. Log in to the Portal as a central administrator.
2. In the **ADMINISTRATION** menu, select **Hierarchies** under **SYSTEM CONFIGURATION**.
3. Click the pencil icon in the top left corner of the **Hierarchies** panel, next to the total number of entities.
4. In the dialog that shows up, click the button **Choose file** to pick the **CSV file to import**. Once you have chosen the file, the dialog displays a **Preview** field below that shows how your CSV file will be imported. If columns are not correctly detected, modify the fields described in the next steps.
5. Specify the character that separates the columns in the CSV file in the **Delimiter** field. By default, the delimiter is the comma character.
6. Choose the text encoding of the CSV file in the field **Encoding**. If you choose a UTF encoding, do not use an editor that creates a BOM header at the beginning of the file (e.g. Notepad). You can select one of the following text encodings:
 - ◆ ISO-8859-1 (Latin 1).
 - ◆ UTF-8.
 - ◆ UTF-16.
7. In the field **Text qualifier**, specify the type of quotes that you used to delimit text in the CSV file, if necessary.
8. Click **Ok** to import the CSV file and modify the entities. A summary of the changes carried out appears in a new dialog.
9. Click **Ok** again in the summary dialog to finish the import.

Format of the CSV file for defining entities

The CSV file that defines the entities must have five columns per line, or six columns if you add an optional comment as the last item. Either all of the lines or none must provide a comment, although the comment of a line may be an empty string in the former case. Each line in the CSV file defines an entity that is assigned to a set of devices in a particular Engine. The entities that you specify here are the basic building blocks of the hierarchies that you will build later; therefore, they are placed at the lowest level of the hierarchies, called the *Entity* level. The Entity field of devices gets a value according to the specified rules. Each line in the CSV file holds the following items, ordered below by their position:

1. Engine name
2. Entity name
3. Entity assignment rule
4. Type of rule
5. Platform
6. Optional comment

The rules for assigning entities to devices in the CSV file are simpler than the rules for categories that you can specify in the Finder. The CSV file supports four types of rules. Choose one of them in the column **Type of rule**. Each type of rule refers to the field of the device that must match the pattern specified in the column **Entity assignment rule** for the device to belong to the entity. See below the list of types of rules and their corresponding device field:

ip

Last IP address of the device.

name

The name of the device.

dn

The distinguished name of the device (an Active Directory value).

collector_tag

The custom tag that identifies the Collector installed in the device.

The format of the pattern in the column **Entity assignment rule** depends on the type of rule that you specified to select devices:

- For an **ip** rule, specify either a single IP address in dot-decimal notation, for example 192.168.0.10, or a subnet in CIDR notation, for example 192.168.0.10/24.
- For a **name** or **dn** rule, give the name or the distinguished name of the device. You can use the wildcards **?** and ***** as substitutes for one or several characters.
- For a **collector_tag** rule, indicate the exact number used to tag the installation. Note that several Collectors can be installed in different devices using the same tag.

In the fifth column (the **Platform**), specify the kind of devices to which the rule applies. You can set it to ***** for the rule to apply to every kind of device. Otherwise, you can use the values **windows**, **mac_os** and **mobile** for the rule to apply only to Windows, Mac or mobile devices, respectively. If you want to apply a rule to a couple of platforms only, repeat the same rule using different platform values.

In a fresh installation of Nexthink, the default rule for assigning entities is the following:

- "Nexthink";"other";"";"name";"";"Automatically generated default entity"

That is, the default rule assigns the entity *other* to every device of the Engine called **Nexthink**, which is the default name of the Engine. From there, you can replace the content of the entity rules as explained above.

Devices that do not match any entity assignment rule are assigned the empty entity, which is represented by a dash sign (-) in both the Finder and the Portal.

Priority of the entity assignment rules

The order of the definitions of entities in the CSV file determines the priority of their assignment rules. Devices that match the rules of several entities are assigned to that entity whose rule appears first in the CSV file.

This is similar to the auto-tagging order of keywords when editing categories in the Finder.

One entity per Engine limitation

A single entity cannot spread among different Engines. In the CSV file, you cannot have the entity **GE** on two Engines, so the following is not valid:

```
"Engine1";"GE";"172.16.1.0/24";"ip"; ""  
"Engine2";"GE";"172.16.4.0/24";"ip"; ""
```

Limit on the number of rules per entity

The maximum number of rules that you can specify in the CSV file for a single entity is 1000.

If more than 1000 rules are specified for one entity, the rules for that particular entity are invalid and thus ignored. All the devices that do not match any subsequent valid rule of another entity are assigned the empty entity, represented by the dash sign (-).

Creating a hierarchy

Once you have specified the entities that form the base of the hierarchies, you can start building your own hierarchies by adding new levels on top of the

entities.

To create a new hierarchy:

1. Log in to the Portal as a central administrator.
2. In the **ADMINISTRATION** menu, select **Hierarchies** under **SYSTEM CONFIGURATION**.
3. Click the plus sign the icons displayed at the top right of the panel. The dialog to add a new hierarchy shows up.
4. Type in a name for the new hierarchy in the **Name** field.
5. Add levels to your hierarchy. See the next section for details.
6. In the choice group **Base hierarchy on**, choose between **all Engines** to create a global hierarchy or **selected Engines** to create a hierarchy that applies to a set of Engines. Note that if you create a hierarchy that applies to a set of selected Engines, you can later promote it to a global hierarchy. On the other hand, if you create a global hierarchy, it is impossible to downgrade it to a hierarchy based on a group of selected Engines.
 - ◆ If you decided to create the hierarchy for a group of **selected Engines**, select your Engines as follows:
 1. Click the **Add** button below the table of Engines. A small dialog with a list of Engines shows up.
 2. Pick an Engine from the list and click **Ok**. Repeat from the previous step until you have selected all the Engines that you wish. The selected Engines are displayed in the table.
7. Click **Ok** to finish the creation of the hierarchy.

Adding hierarchy levels

The levels of the hierarchy indicate the depth of the tree that graphically represents the hierarchy. In the example figure of the hierarchy above, there are three levels defined:

1. Entity level: The lowest level in the hierarchy. It is composed of the names of entities. Each name represents the set of the devices assigned to the entity, according to the rules in the CSV file.
2. Region level: Groups entities into different regions named after the four cardinal points (North, South, East and West).
3. Department level: Divides the company into several departments that are located in one or several regions.

The Entity level is mandatory for all hierarchies. When you create a new hierarchy, you add levels on top of the Entity level. The root node of the hierarchy is always at the central administration level, which is never defined explicitly.

To add levels to a hierarchy from the dialog to create a new hierarchy:

1. Click the **Add** button below the table of levels. A small dialog to edit the level shows up.
2. Enter the name of the level.
3. Click **Ok** to add the level to the table.
4. Repeat from the first step to create as many levels as you need.
5. Optional: Move the created levels up or down in the table by clicking the arrows that appear in the next column, to the right of the name of the level. Note that the Entity level is always the lowest level and that you cannot move it inside the table.

There is a special level that you can use directly above the Entity level called the Engine level. This level makes a first groupment of entities per Engine. To create the Engine level, click the icon with the small Nextthink logo and the plus sign that is placed to the right of the Entity level in the table of levels of the dialog to create hierarchies. The Engine level is automatically filled by the system, which detects the entities (keywords) that are present in each Engine. For that reason, keywords must not be repeated in different Engines. At the end of the process, a new node is created at the Engine level for each Engine found in your system. Similarly to the Entity level, this level cannot be moved upwards or downwards inside the hierarchy.

To manually create the nodes for the other non-special levels, read the following section.

Building the hierarchy tree by editing the entities

Once you have finished creating a hierarchy and its levels, you need to specify nodes for every level. Nodes in one level are used to group the elements of the level below to form the hierarchy. You add nodes to a level by editing the entities of the hierarchy.

To add nodes to the levels of a hierarchy:

1. In the **Hierarchies** panel, select the entities that you want to group from the **Entity** table. Click the row that represents an entity in the table while holding the **Ctrl** or **Shift** keys down to select multiple entities.
2. Click the button **Edit selected entities** below the **Entity** table. A dialog appears with a set of text fields, where each field holds the name of the node to which the set of selected entities belong. Since this is the first time that you edit the entities, the text fields are displayed empty.
3. Type in node names for every level displayed in the dialog.

4. Click **Ok** to group the selected entities below the specified nodes in the hierarchy.
5. Click the floppy disk icon in the top right part of the **Hierarchies** panel to save your work on hierarchies.

Editing a hierarchy

To edit a hierarchy, click the pencil icon that you see at the top right of the **Hierarchies** panel. The dialogs and options for editing the hierarchy are identical to those used when you created the hierarchy.

When you edit the entities of an existing hierarchy, they may already belong to some of the nodes in the hierarchy. You can see the names of the nodes in the columns of the different levels in the **Entity** table. After selecting a group of entities and clicking the button **Edit selected entities**, you find the names of the nodes in the dialog that displays the levels of the hierarchy for the selected entities:

- If the selected entities belong to only one node at a particular level, the text field for that level displays the name of the node.
- If the selected entities belong to different nodes at a particular level, the text field for that level displays the value **[multiple]**.

With the edition of entities, you can add or remove branches from your hierarchy tree or modify it in any other way you choose. Find below a couple of examples:

Example of creating a branch

Example of moving a branch

Be careful when editing a hierarchy that has been already used for aggregating results or for defining user domains. After the edition of an existing hierarchy, a dialog called **Impact of changes** displays all the elements in the Portal that got their associated domains invalidated because of the changes in the hierarchy. Click **Continue** to carry on with the changes anyway. Alternatively, click **Cancel** to revert the changes or to re-edit the hierarchy for reducing the impact.

If you edit a hierarchy, do not forget to save your changes by clicking the floppy disk icon at the top right of the **Hierarchies** panel.

Cleaning up the hierarchy

Eventually, a hierarchy may be based on entities that are no longer used. A couple of cases may bring up this situation:

- The CSV file that defines the entities got some rules removed.
- All the devices assigned to a particular entity were removed from an Engine.
- An Engine became temporarily or definitively unreachable.

The entities that are no longer in use are not automatically removed from the system. Instead, they are respresented in the **Entity** table with an exclamation mark ! at the beginning of the row. This indicates that the entity was not present in any Engine. You can redefine the entities and add the corresponding keywords to enable these entities again, or you can remove them if you no longer need them. To erase the unused entities:

1. Click the broom icon in the top right part of the **Hierarchies** panel. A check list of the unused entities shows up.
2. Check the box of every entity that you want to delete.
3. Click the button **Delete selected entities**.

Note that if an entity is removed and then is detected in an Engine, it will appear again in the **Entity** table, though without any values for the nodes up in the hierarchy.

Viewing hierarchies

If you have created multiple hierarchies, the **Hierarchies** panel lets you select the hierarchy that you want to view. Just pick the desired hierarchy from the list that is placed as the first element in the top heading of the widget, labeled by the word **Hierarchy**, before the other icons.

Once you select a hierarchy, you see the levels of the hierarchy with the list of nodes for each level in the upper part of the widget. In the lower part, you see the **Entity** table, with the names of the entities and the nodes that they belong to. The entities shown in the entity table are filtered by the nodes that you select in the list of nodes of the hierarchy levels. To view all the entities, select the special keyword **All** from the list of nodes of every level. The keyword **All** means that you want to see the entities of all the nodes at that level.

Additionally, you can select the **Overview** mode. In this mode, you just see a big **Entity** table where the columns include the levels of all the hierarchies at the same time. This mode lets you quickly view all the nodes to which an entity belongs in any of the defined hierarchies.

Renaming levels and nodes

When viewing a particular hierarchy in the **Hierarchies** panel, note that there is a clickable text to the right of every level labeled (**rename**). This text also appears to the right of the Entity level in the Entity table. To rename a level in your hierarchy:

1. Click the (**rename**) word to the right of the level. A small dialog to edit the name of the level shows up.
2. Type the new name for the level. The new name must not conflict with the name of any other level in the hierarchy.
3. Click **Ok** to actually rename the level.

Below the list of nodes of every level, you also find a piece of clickable text labeled **rename node** (except for the nodes of an Engine level, because these have the names of the Engines and you are not allowed to change them). To change the name of a node:

1. Select the name of the node inside the list of the level.
2. Click **rename node**. A small dialog to edit the name of the node shows up.
3. Type the new name for the node. The new name must not conflict with the name of any other node in the same level.
4. Click **Ok** to actually rename the node. Only the nodes that are part of the filter to view the hierarchy are renamed (see previous section).

Note that renaming levels and nodes is not the same as editing a hierarchy. Although you can edit a hierarchy to change the names of its levels and nodes, the effect of editing a hierarchy is much stronger to that of just renaming a level or a node. For example, if you change the name of a node by editing the entities of the hierarchy, you are actually creating a new node. The hierarchy itself and its

associated results are modified. On the other hand, renaming a node just changes its text. The node is still the same, but with a different representation in text, so the structure of the hierarchy does not change.

Renaming nodes may affect nevertheless to the results of widgets or investigations grouped by hierarchies. Renaming levels does not modify any result.

Exporting and importing hierarchies

To backup and restore a hierarchy, you can export it to a CSV file or import it from a CSV file from the **Hierarchies** panel.

To export a hierarchy to a CSV file:

1. Select the hierarchy that you want to export in the list of hierarchies of the widget (the list at the top part labeled **Hierarchy**).
2. Click the icon with the arrow down and the initials **CSV** at the top right part of the widget to download the hierarchy as a CSV file.
3. Follow the instructions of your web browser to save the CSV file in the local filesystem.

To import a hierarchy from a CSV file:

1. Click the icon with the plus sign and the initials **CSV** at the top right part of the widget. The dialog to import the hierarchy shows up.
2. Click on the button **Browse** to select the CSV file to import from your local filesystem. A preview of the CSV to import is displayed according to your import options.
3. For the other options in the dialog, select the semicolon as separator character, UTF-8 as text encoding and the double quotes as text identifier if your file was generated by the Portal. Otherwise, use your own custom settings.
4. Click **Ok** to import the hierarchy.

Deleting a hierarchy

Deleting a hierarchy has a direct impact on all objects that depend on that hierarchy. Be sure to know what you are doing before deleting a established hierarchy. The following may happen when you remove a hierarchy from the system (not an exhaustive list):

- Administrators whose administration domain is based on the hierarchy are not be able to log in to the Portal.
- Objects in a view domain based on the hierarchy are visible to central administrators only.
- User accounts with a view domain based on the hierarchy see nothing because they no longer have access rights.

Related tasks

- Creating categories and keywords

Related concepts

- Hierarchy
- Category

Setting the locale in the Portal

Overview

The user interface of the Portal is available in two languages: English and French. Change the locale settings in the configuration file of the Portal to choose the language for the user interface. The locale settings also determine the format of date and time expressions in the Portal.

Weeks are numbered in the Portal to identify weekly periods along the year. Depending on your location, weeks may start on a different day of the week. For instance, in some countries the week starts on Monday, whereas in other countries the week starts on Sunday. In different regions, there are different conventions as well to specify which week is the first week of the year. Configure the Portal to specify both the first day of the week and the first week of the year depending on your local conventions.

Language and date-time format

Basically, there are three possible configurations: International English, US English, and French. By default, the Portal is set to international English, which is different from US English only in the format of dates and time. In international English, days come first in dates and time is expressed in 24 hours format;

whereas in US English, months come first in dates and time is expressed in a 12 hours format with the AM or PM suffix. Find examples of the differences among the three formats in the table below.

	International English	US English	French
Locale settings	en_CH en_UK	en_US	fr fr_CH
Date format	Jan '14 7 Sep 21.09.14	Jan '14 Sep 7 09/21/14	janv. 14 7 sept. 21.09.14
Time Format	14:45:12 15:00 today	02:45:12pm 3pm today	14:45:12 15:00 aujourd'hui

To set the locale in the Portal:

1. Log in to the CLI of the Portal appliance.
2. Optional: If the Portal has no configuration file yet, that is, if `portal.conf` does not exist in folder `/var/nexthink/portal/conf`, create it by copying the defaults from the sample configuration file:


```
sudo -u nxportal cp
/var/nexthink/portal/conf/portal.conf.sample \
/var/nexthink/portal/conf/portal.conf
```
3. Edit the configuration file of the Portal:


```
sudo vi /var/nexthink/portal/conf/portal.conf
```
4. Set the default locale option by typing in the following line. For example, to set the locale to French:


```
globalconfig.portal.user.default-locale = "fr"
```
5. Save your changes and quit the editor by typing:


```
:wq
```
6. Restart the Portal to apply your settings:


```
sudo systemctl restart nxportal
```

Setting the first day of the week

Days are numbered from 0 (Sunday) to 6 (Saturday). To specify the first day of the week, set it as the first element of the **week-days** array, followed by the next four days, in the configuration file:

1. Log in to the CLI of the Portal appliance.
2. Optional: If the Portal has no configuration file yet, that is, if `portal.conf` does not exist in folder `/var/nexthink/portal/conf`, create it by copying the defaults from the sample configuration file:


```
sudo cp -u nxportal
/var/nexthink/portal/conf/portal.conf.sample \
/var/nexthink/portal/conf/portal.conf
```
3. Edit the configuration file of the Portal:


```
sudo vi /var/nexthink/portal/conf/portal.conf
```
4. Add the following line to set the first day of the week:
 - ◆ For example, to set day 1 (Monday) as the first day of the week, as it is common in most of Europe and other parts of the world that follow the ISO standard, type in:


```
globalconfig.portal.portal.week-days = [1, 2, 3, 4,
5]
```
 - ◆ Alternatively, to set the day 0 (Sunday) as the first day of the week, as it is custom in the UK and the USA, type in:


```
globalconfig.portal.portal.week-days = [0, 1, 2, 3,
4]
```
 - ◆ And to set day 6 (Saturday) as the first day, as it is usual in islamic countries, type in:


```
globalconfig.portal.portal.week-days = [6, 0, 1, 2,
3]
```
5. Save your changes and quit the editor by typing:


```
:wq
```
6. Restart the Portal to apply your settings:


```
sudo systemctl restart nxportal
```

Specifying the first week of the year

Because the Portal numbers weeks to let you navigate through weekly periods, it is important for the Portal to know which week is considered to be the first week of the year in your region. The configuration setting for determining the first week of the year is related to the convention for choosing the first day of the week. It is expressed by a number that, when subtracted by the number which represents the first day of the week, indicates the latest day of the week that must belong to the new year (it has to lie in January) for the whole week to be regarded as the first week of the new year.

There are three standard values for indicating the first week of the year for different regions of the globe:

- ISO: 4
When subtracted by 1 (Monday), it yields 3. Three days after Monday is **Thursday**.

- North American: **6**
When subtracted by 0 (Sunday), it yields 6. Six days after Sunday is **Saturday**.
- Islamic: **12**
When subtracted by 6 (Saturday), it yields 6. Six days after Saturday is **Friday**.

As an example, let us look at the transition from 2015 to 2016 for each one of the standard regions:

- In a region following the ISO standard, the first days of the new year fall in the week from Monday, Dec 28th 2015 till Sunday, Jan 3rd 2016. Since Thursday of that week is on Dec 31st 2015, it is not in January of the new year. So that is not the first week of 2016, but the last week of 2015. The first week of 2016 goes from Jan 4th till Jan 10th 2016.
- In a North American region, the week with days in both years goes from Sunday, Dec 27th 2015 to Saturday, Jan 2nd 2016. Because Saturday lies in January 2016, this week is reckoned to be the first week of the year.
- In an Islamic region, the week that marks the transition between the two years goes from Saturday, Dec 26th 2015 to Friday, Jan 1st 2016. Since Friday lies in January 2016, this week is then regarded as the first week of the year.

Note that in regions that follow either the North American or the Islamic conventions, it is enough that the last day of the week falls into January of the new year for the whole week to be the first one of the year.

To set the value for determining the first week of the year:

1. Log in to the CLI of the Portal appliance.
2. Optional: If the Portal has no configuration file yet, that is, if `portal.conf` does not exist in folder `/var/nexthink/portal/conf`, create it by copying the defaults from the sample configuration file:


```
sudo -u nxportal cp
/var/nexthink/portal/conf/portal.conf.sample \
/var/nexthink/portal/conf/portal.conf
```
3. Edit the configuration file of the Portal:


```
sudo vi /var/nexthink/portal/conf/portal.conf
```
4. Add the following line depending on the convention followed in your region to compute the first week of the year:
 - ◆ ISO:


```
globalconfig.portal.portal.first-week-of-year-contains
= 4
```
 - ◆ North America:

```
globalconfig.portal.portal.first-week-of-year-contains
= 6
```

◆ **Islamic:**

```
globalconfig.portal.portal.first-week-of-year-contains
= 12
```

5. Save your changes and quit the editor by typing:

```
:wq
```

6. Restart the Portal to apply your settings:

```
sudo systemctl restart nxportal
```

Changing the Time Zone of the Portal

Overview

Because of the distributed nature of the Nextthink solution, the time zone of the Portal may refer to either:

- The time zone of the machine where the Portal itself is installed.
- The time zone of the Portal account in each Engine.

The local time of the Portal

Use the Web console to change the time zone of the appliance that is running the Portal:

1. Connect to the Web Console over your web browser:
`https://<appliance_ip_address>:99`
2. Select the **Network Parameters** option in the Appliance settings.
3. Choose the appropriate time zone from the list, according to the place where the Portal is located.

The Portal uses the time zone of the machine where it is installed in combination with the time zone of the Portal account on each Engine to schedule the collection of data from the Engines. For more information, see Time Zones and data collection.

The time zone of the Portal account

The time zone of the Portal account determines the time shift between the Portal and each Engine and it influences both the time of data collection and the results of the computation of dashboards.

The time zone of the Portal account is set to the same value as the time zone of the admin account in all Engines. As a result, data collected from different Engines coincide in real-time, although it may correspond to different local times.

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

If you need help or assistance, please contact your Nextthink Certified Partner.
Related tasks

- Time Zones and data collection

Time Zones and data collection

Overview

The Portal collects data from the Engine once every day to compute the metrics for its dashboards and build up its history. Because collecting data from the Engine is a costly operation, the Portal is programmed by default to get the data during the night, when the activity of the Engine is supposed to be low. By default, at one o'clock in the morning, the Portal starts collecting information about the events that occurred during the last day, that is, the 24 hours that went by from past midnight to midnight one hour ago. The whole computation process can take up to several hours, depending on the quantity of data collected and the number and complexity of the metrics to compute.

Special care has to be taken when the Portal and the Engine are placed in different time zones, in particular when the Portal is connected to multiple Engines. A setup with Engines placed in distant locations may lead to surprising results in the Portal if the data collection process is not well understood. One o'clock in the morning in one time zone may be two in the afternoon in another. Thus, data collection may not be triggered during the night for all Engines.

This document explains how the Portal determines when to start collecting data from the connected Engines and other issues that arise when the Portal and the Engines are placed in different time zones.

The time zone of the Portal account

The Portal connects to the Engine by means of a dedicated account. This account is unique to each Engine and is similar to the accounts of the users of

the Finder. The time zone of the Portal account matches the time zone of the admin user in every Engine.

Default behavior

By default, the time zone of the admin user (and, therefore, that of the Portal account) is configured in every Engine to have the time zone of Europe/Zurich, which corresponds to Central European Time (CET, UTC +1 hour) during the winter and Central European Summer Time (CEST, UTC +2 hours) during the summer. Therefore, from the point of view of the Portal, all Engines share the same time zone (Europe/Zurich), even when this is actually not the case.

To schedule the collection of data, the Portal computes the local time that is equivalent to 01:00 Zurich time. When the scheduled time is reached, the Portal begins to collect data from all Engines.

If you change the time zone of the admin account, a similar scenario occurs. All the Engines automatically set the time zone of their Portal accounts to be the same as the time zone of the admin account. As a result, the Portal starts collecting data from all Engines at 01:00 according to the time zone of the admin account. As explained in the previous default case, the Portal computes the equivalent local time for scheduling the data collection.

Example

Let us illustrate the influence of time zones in the data collection with an example involving one Portal connected to two Engines. Imagine that we have a Portal installed in London, one Engine in New York and another Engine in Paris. For the sake of simplicity, we are not going to deal with daylight savings. Therefore, we assume that the Portal in London has UTC time, that the Engine in New York has UTC -5 hours and that the Engine in Paris has UTC +1 hour as their respective time zones.

Suppose that most of the devices with the Collector installed are located in Paris. It makes sense thus to have the time zone of the admin account set to Paris. This ensures that the computation occurs during the night in Paris, when most of the devices are inactive. Since the Portal account shares the same time zone of the admin account, both the Engine in New York and the Engine in Paris have the time zone of the Portal account set to Paris time.

The Portal in London triggers the computation at 01:00 Paris time, that is 00:00 London time. The Engine in Paris has its data collected as usual, from midnight one day ago to midnight one hour ago. However, for the Engine in New York the situation is different. Since its time zone has been centralized to Paris, data collection is performed from 18h last day to 18h today, coinciding in real-time with the collection of data in Paris.

Impact on users

As we said at the beginning, data collection is a costly operation. It increases sensibly the load of the Portal and the Engines while it is going on. To impact the fewer users possible, the Portal collects data during the night. However, in scenarios with multiple time zones involved, the night is not simultaneous for everyone. More users may be impacted as a result of the Portal performing data collection during local working hours.

For instance, In the previous example, where the Portal adapts to the time zone of Paris, users of the Portal in New York may experience poor response time if they try to connect to the Portal late in the evening, because data collection was started at 19:00 New York time and it can go on for a few hours.

Similarly, users of the Finder may experience a decrease in the performance of their connection to an Engine, if the Engine is being solicited by the Portal because of the data collection process.

Therefore, it is recommended to use the time zone of the Engine where most of the users of both the Portal and the Finder are located. In this way, you reduce the impact of data collection on the majority of your users.

Interpreting the results

Be careful with metrics that compute values for particular intervals of time in a day. For instance, let us consider a metric *Number of desktops with nightly activity* that is based on a *between* hours condition. The metric is supposed to return the number of desktops which had any kind of activity during the night, but we have seen that the night is not simultaneous for everybody in setups with multiple time zones.

In the example, the Engine in New York is computing from 18:00 yesterday to 18:00 today, but the Portal makes the computation with respect to the centralized time zone, which is Paris time. Therefore, the widget reports the desktops with nightly activity according to Paris time and not to New York time, even for desktops placed in New York.

Remember that the widgets in the Portal display their results with respect to the time zone used to launch the computation:

- By default, the time zone of Europe/Zurich.
- The time zone of the admin account, if you change it from Europe/Zurich to any other value.

The users of the Portal see time information in their web browser according to one of these possible time zones and it is the same time zone for all users. You should therefore not confuse the time zone of the results in the Portal with the time zone configured in the profile of the user. The time zone in the profile of the user exclusively serves to present information in the Finder, if the user of the Portal is allowed access to the Finder.

Related tasks

- Changing the nightly computation time of the Portal
- Changing the Time Zone of the Portal

Nightly task schedules timetable

This table summarizes the time of execution of those tasks that the different Nextthink components perform during the night, when the activity in your IT infrastructure is supposed to be low.

Some of them are configurable so you can adapt their activation to the time that suits you best.

Local time	Task	Affects	Indicative duration	Defined in
22:15	Portal backup	Portal	< 3 minutes	<code>/etc/cron.d/portal-crontab</code>
01:00	License check	Engine	< 5 minutes	non-configurable
01:00	Data collection	Portal and Engine	minutes to hours	Parameter <code>globalconfig.portal.collector.time-t</code> in file <code>/var/nextthink/portal/conf/portal.com</code>
01:10	Web Console backup	Web Console	< 3 minutes	<code>/etc/cron.d/nxconsole-crontab</code>
03:45	Engine cleaning and maintenance	Engine	15 - 30 minutes	non-configurable
04:15	Engine backup	Engine	< 5 minutes	<code>/etc/cron.d/nxengine-crontab</code>

Related tasks

- Web Console backup and restore
- Portal backup and restore
- Engine backup and restore

Changing the data collection time of the Portal

Changing the starting time of data collection

To change the default time of data collection in the Portal:

1. Log in to the Appliance that hosts the Portal.

2. Optional: If the Portal has no configuration file yet, that is, if `portal.conf` does not exist in folder `/var/nexthink/portal/conf`, create it by copying the defaults from the sample configuration file:

```
sudo -u nxportal cp
/var/nexthink/portal/conf/portal.conf.sample \
/var/nexthink/portal/conf/portal.conf
```

3. Edit the configuration file of the Portal:

```
sudo vi /var/nexthink/portal/conf/portal.conf.
```

4. Add the following lines to the configuration file of the Portal, or modify their values if they are already present. For example, to start the data collection at 2h20:

```
# Time (hour) at which collection for the previous day takes
place
globalconfig.portal.collector.time-to-collect = 2
# Time (minutes) at which collection for the previous day
takes place
globalconfig.portal.collector.time-to-collect-minutes = 20
```

5. Save your changes and exit the vi editor:

```
:wq
```

6. Restart the Portal:

```
sudo systemctl restart nxportal
```

Note that the actual time for triggering the nightly computation depends on how you configure the time zone of the Portal and the Engines.

Changing the maximum number of days collected

Every night, the Portal usually collects data of metrics for the past day only. However, for those metrics with their last days empty of data (because they could not be computed or because their history was cleared), the Portal computes not only the past day, but the number of days configured (up to the maximum number of days available in each Engine).

To set a different number of days to go back and compute metrics with no history, add the following line to the configuration file of the Portal. For instance, to compute five days of history, type in:

```
globalconfig.portal.collector.nb-of-days = 5
```

By default, the Portal goes back **3 days** in the past to compute metrics when the data for their last days are missing. Set the configuration variable to **-1** for the historical computation to go back up to the maximum number of days available in each Engine. Remember that computing metrics for dates in the past has some limitations.

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

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

Related references

- Time Zones and data collection

Establishing a privacy policy

Overview

Nextthink privacy is built around five pillars:

Security of information: The information is collected via encrypted channels and the access to all databases is restricted.

User privileges: Users have privileges that restrict the access to only a subset of devices or locations (domain view). Edition and configuration of the system require special access rights (administration privilege and edition rights). Access to external web domains and web requests need special privileges.

Anonymization: The view of users, devices, destinations and web domains is by default anonymized. Users need special privileges to access the full data.

Storage policy: The full set of information is collected and stored by default. However, it is possible to remove and prevent collecting devices and other information from the dataset. There is also a special policy for Web & Cloud storage that can prevent the collection of web domains.

Audit trails: Every change in the configuration settings is audited, including account edition.

Security of information

Overview of communication channels

The following schema describes the communication architecture from a high level point of view.

The table describes the communication channels used to access or transport sensitive information:

Core components			Protocol or encryption
Collector	-->	Engine	UDP encrypted
Finder	<-->	Engine	TLS
Portal	<-->	Engine	HTTPS by default
Portal	<-->	Nexthink Central License Manager	HTTPS
Optional			
Shell	<-->	Appliance (Engine or Portal)	SSH
API	<-->	Engine	REST HTTPS
Active directory	<-->	Engine	SSL
Application Library	<-->	Engine	HTTPS
Investigation Library	<-->	Portal	HTTP
Investigation Library	<-->	Finder	HTTP
DB backup	<-->	Engine	SMB
Email	<-->	Engine	SMTP
Nexthink updates	<-->	Finder, Appliance	HTTPS, HTTP
Nexthink customer improvement program	<-->	Finder	HTTPS

All the channels that transport sensitive information are encrypted. All optional channels have to be activated or configured, apart from the shell that is set-up by default.

Collected data

Nexthink does not collect any information about the content of files, e-mail, web sites or any other content. Nexthink collects the following data:

Objects (represent real life items recognized by Nexthink)
<ul style="list-style-type: none"> • User • Device • Package • Application

<ul style="list-style-type: none"> • Executable • Binary • Port • Destination • Printer • Domains
Activities (represent actions performed by Objects)
<ul style="list-style-type: none"> • Installation • Execution • Connection • Print job • System boot • User logon • Web request
Events (are warning or errors)
<ul style="list-style-type: none"> • Device warning • Device error • Execution warning • Execution error

User privileges

Accounts are based on *profiles* and *roles*.

Profiles determine the access rights of a user:

- Access to the Portal, eventually limited to a domain (see below for the meaning of domain), with or without the right to administer.
- Access related to domains (Web & Cloud visibility) in the Finder. By default, users can only see domains that are configured in web-based services.
- Access to the Finder and, eventually, the rights to edit applications, objects tags, categories, services and global alerts.

Roles define the default content that the user can access:

- For non-administrator users, this restricts the content that can be accessed in the Portal

Limiting the administration and the view to a domain

Devices can be grouped along a hierarchical tree. This allows to group devices by (e.g.) Department / Region / Entities:

View Domains

A View domain represents which data a user has the right to see. It is defined by a node of the hierarchy and optionally by a limit in the depth. Based on the previous example, a view domain could limit the view to a specific Department and allow the user to drill-down on the underlying Region but prevent to see the details by Entities.

Admin Domain

An Administration domain is a part of a hierarchy that an administrator could manage. In this case, management means that the administrator could, for example, create users with view domains included in its administration domain or create content focused on it. An administration domain is defined by a node of the hierarchy; the domain is the sub-tree of this node.

Privileges for users of Nextthink Finder

For users of the Finder, select their privileges when creating the user profiles (step 4).

The privileges are related to the edition and application of object tags, the modification of the system configuration (categories, metrics, etc), and other features for system management.

Anonymization

Access rights to data

There are four levels of data privacy, defined in the profile of the account, that specify the access rights of each account to particular pieces of information:

Access rights	Description
Anonymous users, devices destination, and web domains	The names of users, devices, destinations, and web domains are not visible to the account
Anonymous users and devices	The names of users and devices are not visible to the account
Anonymous users	Only the names of users are not visible to the account
None (Full access)	No restrictions: all names are visible

The following table explains what is visible for **users, devices , destinations and domains** relative to the data privacy level.

Data Privacy Level	Users	Devices	Destinations	Domains
None (Full Access)	Username	Computer name	Destination name	Domain name
	Distinguished Name	Windows SID	IP address	Nextthink ID
	Full Name	IP address	Nextthink ID	
	Nextthink ID	Nextthink ID		
Anonymous Users	<i>Anonymous Users</i>	Computer name	Destination name	Domain name
		Windows SID	IP address	Nextthink ID
		IP address	Nextthink ID	
		Nextthink ID		
Anonymous Users and Devices	<i>Anonymous Users</i>	<i>Anonymous Devices</i>	Destination name	Domain name
			IP address	Nextthink ID
			Nextthink ID	

Anonymous users, devices destination, and web domains	<i>Anonymous users</i>	<i>Anonymized device</i>	<i>Anonymized destination</i>	<i>Anonymized domain</i>
--	------------------------	--------------------------	-------------------------------	--------------------------

Display - anonymized User ID

When the data privacy level enforces anonymous users, users are displayed as in the screenshot below. Investigation using the name of the user is not possible. But if an authorized user provides the user ID, it will be possible to make an investigation and retrieve data.

Display - anonymized devices

When the data privacy level enforces anonymous devices, devices are displayed as in the screenshot below. As for the user ID, it is not possible to make any direct investigation without knowing the device ID.

Display - anonymized destinations

When the data privacy level enforces anonymous destinations, destinations are displayed as in the screenshot below. Direct investigations without knowing the destination ID are not possible.

Display - anonymized domains

When the data privacy level enforces anonymous domains, domains are displayed as in the screenshot below. Direct investigations without knowing the domain ID are not possible.

Categories

Categories also support data privacy: a level can be set for a category so that only accounts with the same or a higher data privacy level will be able to see and use a given category. For example, if a category is created with a Data Privacy level set to "none (full access)", only Finder user accounts having a "none (full access)" level will be able to see and use this category. The privacy settings on categories applies only to the Finder.

Examples of user profiles

This is an example of some of the user profiles and privileges that can be configured with current Nextthink privacy features:

Nextthink administrator
He is the administrator of Nextthink products within the enterprise and therefore has full access rights.

<p>User privileges</p> <p>Portal:</p> <p>Administrator: central</p> <p>Reader: all domains</p> <p>Finder:</p> <p>Allow access, allow edition</p>	<p>Anonymization</p> <p>Portal & Finder:</p> <p>none (full access)</p>
CIO	
<p>He needs high level information. Therefore he will mainly use Portal as a Reader.</p>	
<p>User privileges</p> <p>Portal:</p> <p>Administrator: no</p> <p>Reader: all domains</p> <p>Finder:</p> <p>No access, No edition</p>	<p>Anonymization</p> <p>Portal & Finder:</p> <p>(anonymous users)</p>
Privacy officer	
<p>He has the full access regarding data anonymization and can provide the User ID to other co-worker if needed.</p>	
<p>User privileges</p> <p>Portal:</p> <p>Administrator: no</p> <p>Reader: all domains</p> <p>Finder:</p> <p>Allow access, No edition</p>	<p>Anonymization</p> <p>Portal & Finder:</p> <p>none (full access)</p>
Security engineer	
<p>He needs full access to all data such that he can investigate any issues.</p>	

<p>User privileges</p> <p>Portal:</p> <p>Administrator: no</p> <p>Reader: all domains</p> <p>Finder:</p> <p>Allow access, allow edition</p>	<p>Anonymization</p> <p>Portal & Finder:</p> <p>none (full access)</p>
Network & system engineer	
<p>He needs access regarding connection and destination but does not need to access user information.</p>	
<p>User privileges</p> <p>Portal:</p> <p>Administrator: no</p> <p>Reader: all domains</p> <p>Finder:</p> <p>No access, No edition</p>	<p>Anonymization</p> <p>Portal & Finder:</p> <p>anonymous users</p>
Support engineer	
<p>He only needs to access user information when required and needs to ask the privacy officer for User ID.</p>	
<p>User privileges</p> <p>Portal:</p> <p>Administrator: no</p> <p>Reader: all domains</p> <p>Finder:</p> <p>Allow access, No edition</p>	<p>Anonymization</p> <p>Portal & Finder:</p> <p>anonymous users</p>
IT project manager (transformation)	

He is only accessing information related to a specific project and only needs anonymous information.	
<p>User privileges</p> <p>Portal:</p> <p>Administrator: limited domains</p> <p>Reader: limited domains</p> <p>Finder:</p> <p>Allow access, allow edition</p>	<p>Anonymization</p> <p>Portal & Finder:</p> <p>anonymous users, devices, destinations and domains</p>

Storage policy

Database

The following databases are used in Nextthink product:

Engine	Portal
Database (in memory)	Database
Database <ul style="list-style-type: none"> • Internal (automatic) • External (not configured by default) 	Database backup <ul style="list-style-type: none"> • Internal (automatic) • External (not configured by default)

Ignoring fields

In addition to the anonymization of data, it is possible to configure the system to ignore certain data that is delivered by the collector. In this case, data are not recorded at all:

ignore_username	If this is set to true, engine will no longer store the user names and Finder will show 'Unknown' for all usernames.
user_interaction	If set to false, user interaction information will no longer be recorded (it will not be displayed in the device view and the "interaction time" aggregate will be always 0%).
ignore_windows_license	If set to true, windows license key will no longer be stored.

ignore_print_jobs	If set to true, all print jobs will be ignored.
ignore_external_ip	If set to true, destination IP address in connections will be set to 0.0.0.0
ignore_external_domains	If set to true, domains that are external will not be recorded.

Retention time

By default, a device is removed automatically from the Engine Database after 3 months of no activity. The retention time can be configured.

Ignoring specific devices

For each device, it is possible to restrain the collected information at the level of the Engine. The possible settings are:

- Web requests, connections and executions (by default, everything is stored)
- Connections and executions
- Executions only
- None
- Remove

For the latter case, this means that the device will be removed from Engine database if there is no activity for more than one day (i.e. the Collector was uninstalled).

In the Finder, right-click a particular device in the list view results of an investigation or in the top-left icon of its own device view and select **Edit...** :

Ignoring specific application, executables, binaries and domains

The same is possible for applications, executables and binaries. The only difference is that it is not possible to remove them, but only to stop storing the related information.

Web & Cloud

There are three storage policies, that can be applied to every engine and that applies to all domains and web requests. This can be set up in the Webconsole:

Web & Cloud storage policy	Use cases	Web domains
1 None	I don't want to store any information related to web domains.	Domains and web requests is discarded.
2 Services only	+ I want to monitor internal or external web services like salesforce.com, office365.	Storage is discarded unless related to a configured web-based service. (*)
3 All	+ I want to discover all web applications used in my company. + I want to see if there are any security breach in my company	Every domains and web requests are stored. But the visibility can be restricted and depends on user privileges. (*) (**)

(*) If a web service is created, the underlying web request and domains are **stored** and there are **no restriction** on visibility.

(**) If a web request is NOT defined in a service, its access will be **restricted**.

Portal account visibility

Finder users need special privileges to view web domains and web requests that are not part of a web-based service (see here above). The same setting is available for the Portal account. If the visibility is "restricted" it will prevent Widget to show data that are not part of a web-based service. This can be set up in the Webconsole.

Engine internal domains

Internal domains are never sent to the Application Library. The following rules allow to identify internal domains:

- Domains with non-official TLD (top level domain)
- Domains with name corresponding to IP addresses belonging to Engine internal network.
- Domains with names matching custom rules (e.g. *.nexthink.com). These rules can be set up in the Webconsole.

Blacklisted domains

For privacy reasons, you may want to avoid storing web requests to particular domains. For instance, a web application that collects opinions and complaints of employees about their peers and superiors requires the anonymity of the participants. However, with the right level of permissions, a user of the Finder can easily discover who connected to the application and when, just by investigating the web requests that are addressed to the domain of the web application. To make the system ignore web requests to specific domains, add the domains to the *blacklist* found in the Web Console.

To add a domain to the blacklist:

1. Log in to the Web Console as administrator.
2. Go to the **Engine** settings and select the **Privacy** tab.
3. Add the domain to the list **Blacklisted domains**:
 - ◆ Separate the names of the domains with a single space character (e.g. *anonymize.nexthink.com *.example.com*).
 - ◆ You can use wildcards in the names of the domains:
 - ◇ The question mark **?** may be replaced by any single character.
 - ◇ The asterisk ***** may be replaced by any number of characters.

Audit trails

Auditing Nexthink is performed using the syslog framework. It captures actions performed with administrator rights that may impact the system. It is not a logging facility.

Only the action and who performs it is audited. The values that are set are not logged.

The complete list of audit point is available [here](#).

Customer improvement program

The Nextthink Customer Experience Improvement Program will deliver benefits to the customers by allowing us to understand how customers use Nextthink software, so that continuous enhancement can be provided. The program is voluntary and anonymous and can also be disabled by default for all users.

[Find out more](#)

Nextthink library

Nextthink Library is a cloud-based knowledge database that gives customers access to a large set of ready-to-use predefined investigations, reports, templates and application information. The Nextthink Library is not mandatory and its access has to be enabled.

When enabled, anonymized data are collected and send to the library. This allow the tagging of binaries with threat level and categorization, and hardware and software compatibility assessment.

The details of collected attributes are described in a dedicated document available on the partner portal.

Related tasks

- [Adding Users](#)

Related references

- [Customer Experience Improvement Program](#)
- [Nextthink Library](#)

Security settings in the Appliance

Overview

The Appliance uses standard mechanisms for authentication and security. Connections to the CLI of the Appliance are established through OpenSSH, which is the SSH implementation installed in the operating system of the

Appliance, and connections to the Portal are managed by the security layer of the underlying Java implementation.

Some of the encryption algorithms allowed by these technologies may be considered weak and relatively easy to break, according to current technology standards. Ciphers that use short keys may compromise the security of the Appliance. To protect you against attacks that aim to break the ciphers used, you can control the allowed ciphers in the Appliance and disable those that you consider too weak. Just make sure that your SSH clients and browsers support the encryption methods that are not disabled.

SSH configuration

Starting from Nexthink V5.1, the default configuration of SSH in the Appliance is set to exclusively use ciphers and hashes that are considered strong. However, this configuration is automatically set only for fresh installations of Nexthink V5.1. If you upgraded to Nexthink V5.1 or if you work on a previous version, you can set the same ciphers allowed by default in Nexthink V5.1:

1. Log in to the CLI of the Appliance.
2. Edit the SSH configuration document:
sudo vi /etc/ssh/sshd_config
3. Add the following two lines at the end of the configuration file:
Ciphers
aes256-ctr, aes192-ctr, aes128-ctr, arcfour256, arcfour128, arcfour
MACs
hmac-sha2-512, hmac-sha2-256, hmac-sha1, umac-64@openssh.com, hmac-ripemd160
4. Restart the SSH daemon:
sudo systemctl restart sshd

Portal configuration

By default, the Portal supports TLS 1.1 and TLS 1.2 as security protocols. Most modern browsers and operating systems are able to use these protocols to secure their communications over the Internet. Associated to these protocols, the Portal also supports a default set of cipher suites (considered strong) to negotiate the security settings of a connection.

However, users of Internet Explorer in either Windows Vista or Windows XP, for instance, are limited to TLS 1.0. This limit also applies if you use the Finder in Windows Vista. Therefore, if you want the Portal to support TLS 1.0, you must add it to the list of supported protocols in the configuration file of the Portal.

To change the supported protocols and cipher suites:

1. Log in to the CLI of the Appliance hosting the Portal.
2. Edit the configuration file of the Portal:

```
sudo vi /var/nexthink/portal/conf/portal.conf
```
3. Type in the protocols supported in the entries:
 - ◆ `globalconfig.jetty.ssl-included-protocols`
 - ◆ `globalconfig.jetty.ssl-included-ciphersuites`
4. Save the file and exit by typing:

```
:wq
```
5. Restart the Portal:

```
sudo systemctl restart nxportal
```

For instance, type in the following lines in the configuration file of the Portal to reproduce the protocols and cipher suites supported by default:

```
globalconfig.jetty.ssl-included-protocols = ["TLSv1.1", "TLSv1.2"]

globalconfig.jetty.ssl-included-ciphersuites = [
  "TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384",
  "TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256",
  "TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA",
  "TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA",
  "TLS_DHE_RSA_WITH_AES_256_GCM_SHA384",
  "TLS_DHE_RSA_WITH_AES_128_GCM_SHA256",
  "TLS_DHE_RSA_WITH_AES_256_CBC_SHA",
  "TLS_DHE_RSA_WITH_AES_128_CBC_SHA",
  "TLS_RSA_WITH_AES_128_CBC_SHA" ]
```

Note that these entries are actually absent from the default configuration file of the Portal (you will not see them when editing the file). In reality, the contents shown above are internal to the Portal and they are used only if you do not override the settings in the configuration file.

For example, to support the protocol TLS 1.0 in addition to the default protocols TLS 1.1 and TLS 1.2, create or substitute the entry of included protocols for:

```
globalconfig.jetty.ssl-included-protocols = ["TLSv1", "TLSv1.1",
"TLSv1.2"]
```

Importing and replacing Certificates

Nextthink encrypts the communications among its components to protect the information that travels through the network. The Appliance components (Engine, Portal and Web Console) and the Finder all use TLS to communicate with each

other.

By default, Nextthink includes a set of self-signed certificates and a cryptographic key to secure network communications via TLS. However, the security policy in your company may require you to replace the default Nextthink certificates by your own; whether they are signed by a third-party certification authority or self-signed.

Learn here how to replace the certificates in the Appliance and in the Finder.

Format of the Certificates

- The certificate and key are in PEM format.
- If needed, use OpenSSL to convert your own certificates to the PEM format.

Replacing the Certificate on the Engine

The key is stored at `/var/nextthink/engine/common/etc/key.pem`. The certificate is stored at `/var/nextthink/engine/common/etc/certificate.pem`.

1. Connect to the Engine via the command line interface.
2. Stop the Engine:
sudo systemctl stop nxlanch
3. Back up the default key and certificate:
**sudo mv /var/nextthink/engine/common/etc/key.pem \
/var/nextthink/engine/common/etc/key.pem.default**
**sudo mv /var/nextthink/engine/common/etc/certificate.pem \
/var/nextthink/engine/common/etc/certificate.pem.default**
4. Upload the new key and certificate files to the *nextthink* account (`/home/nextthink`) of the Appliance using, for instance, SCP.
5. Copy your key and certificate files to the Engine:
**sudo cp /home/nextthink/key.pem \
/var/nextthink/engine/common/etc/key.pem**
**sudo cp /home/nextthink/certificate.pem \
/var/nextthink/engine/common/etc/certificate.pem**
6. Restart the Engine:
sudo systemctl start nxlanch

Replacing the Certificate on the Console

The key and the certificate are concatenated in single file at

/var/nexthink/console/etc/certificate.pem.

1. Connect to the Appliance via the command line interface.
2. Stop the Console:
sudo systemctl stop nxconsole
3. Back up the default certificate:
**sudo cp /var/nexthink/console/etc/certificate.pem \
/var/nexthink/console/etc/certificate.pem.default**
4. Concatenate your key and certificate:
**cat /path/to/your/key.pem /path/to/your/certificate.pem >
full_certificate.pem**
5. Copy the concatenated certificate to:
**sudo cp full_certificate.pem
/var/nexthink/console/etc/certificate.pem**
6. Restart the Console:
sudo systemctl start nxconsole

Replacing the Certificate on the Portal

The key and the certificate are stored in the keystore of the Portal:
/var/nexthink/portal/keystore/keystore.jks

1. Connect to the Portal via the command line interface.
2. Stop the Portal:
sudo systemctl stop nxportal
3. Back up the default keystore:
**sudo cp /var/nexthink/portal/keystore/keystore.jks \
/var/nexthink/portal/keystore/keystore.jks.default**
4. Delete the Nexthink Certificate from the keystore:
**sudo /usr/java/default/bin/keytool \
-keystore /var/nexthink/portal/keystore/keystore.jks \
-delete -storepass nexthink -alias portal**
5. Convert the new key to DER format using OpenSSL:
**openssl pkcs8 -topk8 -nocrypt -in /path/to/your/key.pem \
-inform PEM -out key.der -outform DER**
6. Convert the new certificate to DER format using OpenSSL:
**openssl x509 -in /path/to/your/certificate.pem \
-inform PEM -out certificate.der -outform DER**
7. Download the following Java class to import certificates:
<http://download.nexthink.com/doc/ImportKeyV6.class>
8. Import the certificate into the keystore:
sudo java ImportKeyV6 key.der certificate.der
9. Restart the Portal

sudo systemctl start nxportal

Configuring the Finder

If you install a certificate in the Portal and in the Engine that is signed by a recognized Certification Authority, the Finder is able to connect to the Portal and to the Engine out-of-the-box.

If you install a self-signed certificate in the Portal and in the Engine, import this certificate into all the computers where the Finder is installed. Otherwise, the connection of the Finder displays the message **The security certificate of Nexthink Portal / Engine could not be validated**. You may decide nevertheless to continue with the connection.

To manually import the certificate into the Finder:

1. Type **WinKey+R** to open the Run dialog.
2. Type in *certmgr.msc* and press **OK**.
3. Right-click **Trusted Root Certification Authority**, and select **All Tasks > Import**.
4. The **Certificate Import Wizard** starts. Click **Next**.
5. Click **Browse** and select the *certificate.pem* file.
6. Click **Next**.
7. In the dialog **Place all certificates in the following store**, click the **Browse** button.
8. Tick the box **Show physical stores**, and select **Trusted Root Certificate Authority\Local computer**.
9. Click **Next**.
10. End the wizard by clicking **Finish**.

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

- Logging in to the CLI

Managing Appliance accounts

1. Connect as Admin to the Console over a Web Browser.

2. Select the Appliance Tab and then Accounts on the left side.

There are three Accounts managed from the Console.

- Appliance administration account used to administer the Appliance.
- Appliance remote account used for the central management from the Portal.
- Nextthink support account used for the command line interface.

Appliance administration account

The default values are:

- User = **admin**
- Password = **admin**

Appliance remote account

Set the Account status to **Enable**. Set a new password.

By default the User is managed by Nextthink.

Nextthink support account

Set the Account status to **Enable**.

The default values are:

- User = **nextthink**
- Password = **123456**

For all three accounts, Nextthink recommends to change the password after first login.

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

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

Sending email notifications from the Appliance

For the Engine and the Portal to send alert notifications and dashboard digests via email, configure the SMTP settings of the Appliance through the Web Console:

1. Open a web browser and type in the following URL to open the Web Console:
`https://<IP_address_of_Appliance>:99`
2. Log in to the Web Console as admin.
3. In the **Appliance** tab, select **SMTP** from the menu on the left-hand side. The Web Console displays the **Appliance SMTP configuration**.
 - ◆ If no mail account is present, click the circled plus sign in the bottom right corner of the widget to add a new account.
 - ◆ If a mail account is already set up, click the circled **e** to edit the account or the circled cross to remove it and then proceed from the previous step.
4. Set the account data:
 1. Type in the IP address and port of the **SMTP server** to use.
 2. Check the box **TLS** if your mail server requires encrypted communication. The Appliance only supports STARTTLS as the mechanism to establish an encrypted mail channel.
 3. Type in the **Sender email** account that will send emails on behalf of your Engine or Portal.
 4. Type in the user **Login** and **Password** for the email account if the SMTP server requires authentication.
 5. Optional: Type in your email address in the field **Test e-mail** and click the button **Send test e-mail** to check your settings. If they are correct, you should receive an email message from the Appliance in a short time.
5. Click **Save** to make your changes permanent.

Set the address of the Portal for the links in email digests

Users of email digests can click on some parts of the digest to open the Portal and display the appropriate dashboard; that is, the dashboard that contains detailed info about the clicked part (metric or service).

For the links in the digest to correctly point to the Portal, set the base address of the Portal (DNS or IP) in the Web Console. By default, this is the *hostname* of the Portal appliance. If no hostname is configured, the Web Console will propose to use *https://localhost*. For external reachability or other specific reasons, you may

set a different DNS name or IP address to access the Portal from the email digests.

To set the base address of the Portal for the links in the email digests:

1. Open a web browser and type in the following URL to open the Web Console:
`https://<IP_address_of_Appliance>:99`
2. Log in to the Web Console as admin.
3. In the **Portal & Finder** tab, select **Parameters** from the menu on the left-hand side.
4. Under the section **Portal parameters**, type in the address of the Portal in **Portal address used for email digest**.
5. Click **Save** and wait for the Portal to restart.

Related tasks

- Managing Appliance accounts
- Installing the Appliance
- Receiving email digests
- Receiving alerts

Controlling session timeouts in the Portal

Overview

To prevent Cross-Site Request Forgery (CSRF), sessions in the Portal are time-limited and protected by secure tokens.

By default, a token remains valid for 8 hours. If you are inactive for more than 8 hours while in a Portal session, your next action in the Portal will redirect you to the login page.

In turn, a session is valid for 24 hours by default. After continuously using the Portal for 24 hours without interruption, the session expires and you are forced to log in again to renew the session.

Setting the value of session timeouts

The validity time for both tokens and sessions is configurable. Remember that the longer the interval, the more vulnerable the Portal is to CSRF attacks.

1. Log in to the CLI of the Portal appliance.
2. Optional: If the Portal has no configuration file yet, that is, if `portal.conf` does not exist in folder `/var/nexthink/portal/conf`, create it by copying the defaults from the sample configuration file:


```
sudo -u nxportal cp
/var/nexthink/portal/conf/portal.conf.sample \
/var/nexthink/portal/conf/portal.conf
```
3. Edit the Portal configuration file:


```
sudo vi /var/nexthink/portal/conf/portal.conf
```
4. Type in the following line to set the value for the validity time of tokens (minimum value is 2 minutes). Use the suffix **h** to specify the time interval in hours and **m** to express it in minutes. For example, to set the period to its default value of 8 hours:


```
globalconfig.portal.session.token-validity-period = 8 h
```
5. Type in the following line to set the value for the validity time of sessions. For example, to set the period to its default value of 24 hours:


```
globalconfig.portal.session.maximum-session-lifetime = 24 h
```

 - ◆ Optional: Express it in minutes:


```
globalconfig.portal.session.maximum-session-lifetime = 1440
m
```
6. Save your changes and exit:


```
:wq
```
7. Restart the Portal to apply your settings:


```
sudo systemctl restart nxportal
```

Overriding session timeouts

Note that, when creating a user, the user may be granted the privilege of never being timed out. In that case, the values configured for session timeouts do not apply to that user.

Related tasks

- Adding users

Preventing password saving in the Finder

Overview

Saving the password of login sessions in the Finder may be a convenient feature for users to avoid typing their password again and again. However, for security reasons, you may want to enforce a policy of making password input mandatory, especially if the users share the workstations that they use to log in to the Finder.

Procedure

The Finder reads a key in the Windows registry to know whether to allow users to save their password or not. If the value of the key is set to 1, the Finder hides the options **Remember password** and **Sign me in automatically** in the login dialog.

To prevent users from saving their password in Finder sessions:

1. In the computer where the Finder is installed, press **Win(key)+R** to display the run dialog.
2. Type in **regedit** as the program to open in the dialog and press **Enter**. The Registry Editor opens.
3. Browse the Windows registry in the Registry Editor and select the key **HKEY_LOCAL_MACHINE\SOFTWARE\Nexthink**.
 - ◆ If the key does not exist, create it by right-clicking the **SOFTWARE** folder:
 1. Select **New -> Key** from the context menu.
 2. Type in *Nexthink* as the name of the new key.
 3. Right-click the area on the right-hand side of the Registry Editor that holds the list of values for the key.
 4. Select **New -> DWORD (32-bit) Value** from the context menu.
 5. Type in **preventUsersFromSavingPassword** as the name of the value.
4. Right-click the value with the name **preventUsersFromSavingPassword** to change its data.
5. Select **Modify...** in the context menu. The dialog to edit the value shows up.
6. Set the value of the field **Value data** to 1 in the dialog.
7. Click **OK**.

This method changes the value of the registry key in one computer only. Alternatively, you can use GPO to impose the same value for the registry key in all the computers where the Finder is installed.

Related tasks

- Logging in to the Finder

Establishing a data retention policy in the Engine

Overview

The Engine stores the real-time data that it receives from the the Collectors in the form of events. Events are very numerous and they usually take most of the memory reserved to the Engine. The types of events that occupy most of the space in memory are executions, connections, and web requests. When two or more of these events are very similar to each other and they occur in sequence, the Engine may consider that they are actually the same event. In that case, the Engine combines the data of the events and stores only one event in its database. We say then that the Engine *aggregates* the information of several events into one; thus saving memory space and resulting in a larger history for the Engine.

When you have the web monitoring feature fully enabled, you usually collect a huge number of web domains. In the same spirit of event aggregation, when two or more domain names share their highest level domains, the Engine may group them into one generalized domain by obeying specific rules. This process is known as domain *compaction* or domain *compression* and it consists in replacing one or more of the lower level domains in the domain name by the wildcard character *. For instance, the Engine might compact the domains **one.example.com** and **two.example.com** into ***.example.com**. Note however that those domains declared as internal or included in the definition of a web-based service are considered of special interest to you and, therefore, they are never compacted.

Learn here how to set up the policies for both the aggregation of events and the compaction of domains in the Engine.

Setting the aggregation policy for events

Choose among four strategies of aggregation for an optimal trade-off between detailed event information and history length. The more aggressive the policy, the fewer individual (non aggregated) events are visible from the Finder.

1. Log in to the Web Console as admin.
2. Go to the **Engine** tab.
3. Select the **Parameters** section.
4. Select one of the following aggregation policies from the list labeled **Aggregation policy for executions, connections and web requests:**
 - **very low**, for the traditional minimal aggregation.

- **low**, for increasing the history 10% approx. while keeping most of the individual events.
 - **medium**, for a more aggressive aggregation policy to increase history in the Engine up to 40%. This is the recommended setting.
 - **high**, for the most aggressive aggregation policy to increase history in the Engine up to 50%
5. Click **Save**. Note that the Engine is restarted after saving the changes.

Setting the compaction level for domains

1. Log in to the Web Console as admin.
2. Go to the **Engine** tab.
3. Select the **Parameters** section.
4. Select one of the following domain compression policies from the list labeled **Domain compression**:
 - **medium**, the default compression policy for domains with more than five levels or with repetitive (or randomly generated) subdomains. This is the recommended setting.
 - **high**, to apply a compression method to all the stored domain names according to a public list of domain suffixes.
5. Click **Save**. Note that the Engine is restarted after saving the changes.

For a detailed explanation of compaction policies, see the section about compaction in the definition of domain.

Related concepts

- Event
- Domain

Related references

- Data retention in the Engine
- Public suffix list (external)

Special operation modes for the Engine and the Portal

When operating normally, the Engine receives and processes all information coming from the Collectors and sends data to the Portal over time. For demo

purposes or other special reasons, you may want to alter the normal functioning of the Engine and the Portal. In this chapter, learn how to freeze the time in the Engine and in the Portal (demo mode) or how to make the Engine store device information only and filter all other events (zero config mode).

In addition, know how to deal with the compatibility mode of Internet Explorer when browsing the Portal. Lastly, if you have the Web and Cloud module activated, learn how to configure the Engine for recording HTTP connections with extended status codes from a proxy.

Setting up demo mode

While working with the Portal, imagine that you detect an interesting occurrence in your network, such as a high rate of failures in a service at a particular time of the day. You may want to share your findings with other people in your team or with management. Ideally, you would like to replay the same situation at a later time to analyze what happened at that point in time with the help of all Nextthink products. To that end, you can back up the databases of the Engine and the Portal and restore them later in other instances in demo mode.

Demo mode consists in freezing the time of the Engine and the Portal, so they do not evolve with the passing of time. Therefore, you consistently find the same data that was present when you made the backup in both the Engine and the Portal. To prevent data loss in your production environment, you must not use the production Engine and Portal to play your demos, but dedicated instances of the Engine and the Portal that you have installed elsewhere; for instance, a virtual machine in your personal desktop.

An Engine in demo mode does not process any packet coming from the Collector nor performs any kind of activity: it does not create new events in the database, it does not notify new alerts, it does not send or retrieve information from the application library, etc.

To set up the demo mode in the Engine:

1. Log in to the CLI of the appliance that hosts the demo Engine.
2. Edit the configuration file of the Engine that is found in `/var/nextthink/engine/01/etc/nxengine.xml` and set the **mode** tag to

static_time:

```
<config>
  <engine>
    <mode>static_time</mode>
  </engine>
</config>
```

3. Restart the Engine:

```
nxinfo launch --restart
```

The keyword **static_time** forces the Engine to freeze its internal date and time to the moment right after the end of the last event included in its database. Since the time is frozen, the Engine no longer sends real-time service information to the Portal. For the Portal to work in sync with your demo Engine, the time set in the Portal must match the time in the Engine and the Portal must receive real-time services data from the Engine.

To get the time settings from the Engine and send the data of real-time services to the Portal, take these additional steps in the Engine appliance:

1. Call the function **now** in the Engine and note down the result. The function gives you the frozen time:

```
nxinfo shell -e "call now()"
```

2. Schedule a cron job to send real-time service data to the Portal every 10 minutes:

1. Execute in the CLI of the Engine:

```
sudo crontab -e
```

2. In the vi text editor that opens, type in the following line:

```
*/10 * * * * /usr/bin/nxinfo lua --command  
"monitor:send_data_to_portal()"
```

3. Save your changes and quit the editor with the command:

```
:wq
```

After Engine configuration, set the demo mode in the Portal:

1. Log in to the CLI of the appliance that hosts the demo Portal.
2. Optional: If the Portal has no configuration file yet, that is, if `portal.conf` does not exist in folder `/var/nexthink/portal/conf`, create it by copying the defaults from the sample configuration file:

```
sudo cp /var/nexthink/portal/conf/portal.conf.sample \  
/var/nexthink/portal/conf/portal.conf
```

3. Edit the configuration file of the Portal:

```
sudo vi /var/nexthink/portal/conf/portal.conf
```

4. Add the following lines, where **EngineTime** is the frozen time in the Engine that you noted down previously:

```
# Demo mode  
globalconfig.portal.special.demo = true  
globalconfig.portal.special.static.time = EngineTime
```

5. Save your changes and quit the editor:

```
:wq
```

6. Restart the Portal:

```
sudo systemctl restart nxportal
```


Now you have your Engine and Portal ready in demo mode. You may have to wait up to ten minutes for real-time services to receive data from the Engine though.

Stopping the time in the Engine

With the **static_time** option, the Engine selects the optimal point in time to freeze the time in the Engine for a demo. This time corresponds to the instant right after the occurrence of the last event recorded in the database of the Engine. In the case that you want to freeze the time of the Engine to a different point in time, you can do it by setting the following option in the configuration file of the Engine (*/var/nexthink/engine/01/etc/nxengine.xml*):

```
<config>
<engine>
<tweak>
<static_now>time</static_now>
</tweak>
</engine>
</config>
```

Where **time** is in the format YYYY-MM-DDTHH:MM:SS (e.g. 2014-01-01T18:00:00).

This option should be used with care because it can leave events that were originally in the database out of the time range of the Engine or make them too old. Use preferably the **static_time** option for your demos unless you have a very specific requirement.

Storing only device information in the Engine

This mode of operation can be used to deploy a large number of Collectors in a setup with several Engines. The deployment is done in two phases. During the first phase, all Collectors send information to one special Engine that is configured to store device information only. Then, in the second phase, Collectors are classified and definitively configured to send data to a normally operating Engine. For the details on the procedure, please contact Nextthink Customer Success Services.

This special mode of operation of the Engine is known as *zero config* mode. An Engine in zero config mode shows the following properties:

- The Engine processes and stores only device information coming from the Collectors, namely, the MAC address, IP address and SID of the devices.

- All activities and information related to other objects are discarded.
- Devices are created with a special storage policy called **inventory**. A device with this storage policy is never removed from the database in spite of having no events associated.
 - The Engine accepts an unlimited number of devices.
 - The Engine rejects any connection from the Portal.
 - The communication of the Engine with the application library is disabled.

To set up zero config mode, add the following option to the configuration file of the Engine (`/var/nexthink/engine/01/etc/nxengine.xml`):

```
<config>
<engine>
<mode>zero_config</mode>
</engine>
</config>
```

Dealing with compatibility mode of IE in the Portal

The Portal is usually best displayed with the latest rendering capabilities of modern browsers. When working with Internet Explorer, though, you or your organization may have set the browser to *compatibility mode* because of some legacy web applications which are key to your business and are best rendered in older versions of IE.

Even with *compatibility mode* on, web sites can still indicate Internet Explorer to use its latest rendering engine instead. To that end, they use a particular HTTP header when serving web pages. To make the Portal work in this mode, so it tells Internet Explorer to use its most recent render version:

1. Log in to the CLI of the appliance hosting the Portal.
2. Optional: If the Portal has no configuration file yet, that is, if `portal.conf` does not exist in folder `/var/nexthink/portal/conf`, create it by copying the defaults from the sample configuration file:

```
sudo cp /var/nexthink/portal/conf/portal.conf.sample \
/var/nexthink/portal/conf/portal.conf
```
3. Edit the configuration file of the Portal:

```
sudo vi /var/nexthink/portal/conf/portal.conf
```
4. Add the following line,:

```
globalconfig.portal.http.compatibility-mode = true
```
5. Save your changes and quit the editor:

```
:wq
```
6. Restart the Portal:

```
sudo systemctl restart nxportal
```

Recording web requests with extended connection status codes

During normal operation, the Engine ignores web requests with connection status codes between 300 and 499 by default. These extended status codes may be issued by proxies when establishing a secure connection with a server on a client request.

Starting from Engine 5.2.8, you can tell the Engine to record these connections by logging in to the CLI and typing the following command:

```
sudo nxinfo config --set \  
web_monitoring_accept_proxy_extended_status_codes=true
```

Restart the Engine for the new configuration to take effect and beware that acknowledging this kind of connections may significantly increase the number of recorded web request events and, therefore, decrease your time interval for data retention.

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

If you need help or assistance, please contact your Nextthink Certified Partner.
Related tasks

- Logging in to the CLI
- Engine backup and restore
- Portal backup and restore

Ignoring specific print ports

To prevent the Engine from recording print jobs that use specific print ports, list the print protocol prefixes of the ports to be ignored under the **ignored_print_ports** item of the configuration file of the Engine. Along with the print jobs, the Engine also discards the printers that are associated with them.

By default, when the element **ignored_print_ports** is not specified, this option is set in the Engine to ignore the ports with prefixes **TS** and **CLIENT**. Popular virtual environments use these print protocols to print on redirected printers. In this way, the Engine avoids recording duplicate print jobs and printers in virtual environments where the Collector is installed in both client devices and remotely

accessible virtual machines.

To set the prefixes of the print protocols that the Engine must ignore:

1. Log in to the CLI of the appliance that hosts the Engine.
2. Open the configuration file of the Engine for editing:
sudo vi /var/nexthink/engine/01/etc/nxengine.xml
3. Under **config / local / aggregation** add the following lines:

```
<ignored_print_ports>  
  <port_prefix>PREFIX_1</port_prefix>  
  ...  
  <port_prefix>PREFIX_N</port_prefix>  
</ignored_print_ports>
```
4. Save your changes and exit with the following command:
:wq
5. To make your changes effective, restart the Engine:
nxinfo launch -e

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

- Logging in to the CLI

Related references

- Information on printers and printing

Related concepts

- Printer
- Print job

Enabling and Disabling the Engine Application Library Access

Enabling the Engine Application Library Access

Purchase the Security module and enable Application Library Access in the Engine to determine the threat level and category of binaries and web domains:

1. Log in to the Web Console.
2. Select **Engine** and then **Application Library**.
3. Enable **Application Library access** by ticking the check box.
4. Optional: Enable **SSL certificate check** by ticking the check box.
5. Optional: Check the connectivity test by pressing the **Start Connectivity test** button.
6. Click **Save**.

Disabling the Application Library Access

1. Disable the **Application Library access** by unticking the check box.
2. Optional: Disable the **SSL certificate check** by unticking the check box.
3. Click **Save**.

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

If you need help or assistance, please contact your Nextthink Certified Partner.
Related references

- Nextthink Application Library

Importing data from Active Directory

The Engine provides an out the box integration with Active Directory to retrieve the following information via the Lightweight Directory Access Protocol (LDAP):

- **User**: Distinguished Name, Full name, Department, Job title.
- **Device**: Distinguished Name.

The Engine retrieves as well the following information through DNS resolution (DNS namespaces mirrors the AD domains used by an organization):

- **Printer**: Host name.
- **Destination**: Name.

This article discusses data integration from Active Directory and should not be confused with Active Directory Authentication.

LDAPv3 and Active Directory

Reference document: Active Directory LDAP Compliance provided by Microsoft.

Windows Server 2000

The Windows 2000 implementation of Active Directory is an LDAP-compliant directory supporting the core LDAPv3 RFCs available.

Windows Server 2003

Building on the foundation established in Windows 2000 Server, the Active Directory service in Windows Server 2003 is offering new LDAPv3 capabilities:

- **Transport Layer Security (TLS)** - Connections to Active Directory over LDAP can now be protected using the TLS security protocol.
- **Digest Authentication Mechanism** - Connections to Active Directory over LDAP can now be authenticated using the DIGEST-MD5 Simple Authentication and Security Layer (SASL) authentication mechanism. The Windows Digest Security Support Provider (SSP) provides an interface for using Digest Authentication as an SASL mechanism.

Windows Server 2008 and 2012

Both Windows Server 2008 and Windows Server 2012 support LDAPv3.

Other implementations

Although Nexthink officially supports Active Directory based on Windows Servers only, other LDAPv3 compliant implementations (such as OpenLDAP) should work as long as the schema in use is the same as in Active Directory.

Setting Up Active Directory Authentication

LDAP servers require an authenticated connection before they will allow queries (searches). This authenticated connection is called a bind. Most LDAPs allow an anonymous bind where no username or password is submitted; however, others restrict searches to its members and require an authenticated username and password. An Active Directory server requires authenticated access for read-only

searches, and you need to have a bind DN and the corresponding bind password. The syntax for the bind DN depends on the LDAP server itself:

NetBIOS logon name

<domain name>\<username>

Active Directory User Principal Name (UPN)

username@domain.name

Distinguished Name

CN=username, OU=users, DC=domain, DC=name

The Engine supports the authenticated method using the **Distinguished Name** syntax only.

Configuring the Engine through the Web-console

1. Connect to the Web Console (by default <https://engine.yourcompany.com:99>).
2. On the left menu, go to **Engine > Active Directories**.
3. Click the plus symbol on the right to add a new Active Directory server.
4. Complete the LDAP Server Connection fields as follows:
 - ◆ **LDAP server name:** The generic name for your LDAP server. Example: if you write ?nextthink.ch?, the usernames in the Finder will be shown as user@nextthink.ch.
 - ◆ **LDAP Server:** Enter here the IP address of your Active Directory server (we currently do not support the DNS or Netbios name) and the TCP server port (usually 389).
 - ◆ **LDAP Bind DN:** The Distinguished Name. Example: CN=reflexengine, CN=applications, OU=servers, DC=company, DC=local.
 - ◆ **LDAP Bind Password:** Enter the password corresponding to the LDAP Bind DN account.
 - ◆ **LDAP Base DN:** The Base DN to be used as a starting point for directory searches. Base DN is usually derived from the Bind DN by removing the user name and specifying the group where users are located. Example: If Distinguished Name = ?CN=reflexengine, CN=applications, OU=servers, DC=company, DC=local?, you can choose the Base DN as ?DC=company, DC=local?.
 - ◆ **LDAP Scope:** The SCOPE setting is the starting point of an LDAP search and the depth from the base DN to which the search should occur. There are three options (values) that can be assigned to the LDAP SCOPE parameter (we strongly recommend the SUBTREE scope option):

- ◇ **BASE:** This value is used to indicate searching only the entry at the base DN, resulting in only that entry being returned (keeping in mind that it also has to meet the search filter criteria!).
- ◇ **ONELEVEL:** This value is used to indicate searching all entries one level under the base DN - but not including the base DN and not including any entries under that one level under the base DN.
- ◇ **SUBTREE:** This value is used to indicate searching of all entries at all levels under and including the specified base

DN.

5. Click on **Save** (the Engine reboots).

Trusted Domains

Due to the technology used to query Active Directory, the Engine retrieves information from those objects belonging to the domain specified in the configuration only (see **LDAP Base DN** above). It does not follow referrals nor retrieve any information from objects in other domains, even when these other domains share a trust relationship with the configured domain.

Add as many Active Directory servers to the configuration as needed to retrieve objects from several domains.

Querying Active Directory to obtain a User's Distinguished Name

For testing purposes, we advise you to use a powerful tool from Microsoft called Active Directory Explorer. Download it from [here](#).

Here is an example on how you can retrieve a user's DN using this tool :

1. Connect to your AD using your windows username.
2. Click on **Search** > "**class = User -- user**" > "**Attribute = sAMAccountname**" > "**relation = is**" > "**value = YOUR Windows username**", then click on **Add**.
3. Click on **Search** to retrieve the corresponding user's DN.

Active Directory data retrieval

The Engine queries its configured LDAP servers each time that it discovers a new user or a new device.

Engines do not automatically refresh LDAP information once they have retrieved it for a particular user or device. It is however possible to force a manual update via the Finder:

1. Log in to the Finder as a user with *system configuration* permissions.
2. Click the sprocket icon in the top right corner of the Finder window.
3. Select the option **Synchronize with Active Directory...**

The Finder schedules a synchronization with Active Directory data.

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

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

Configuring the system log

Overview

Syslog is a de facto standard solution for logging messages in UNIX-derived systems, such as the operating system of the Appliance. Programs use the syslog system call to send arbitrary messages to the syslog service. In addition to the message itself, two parameters are provided to the syslog system call: the facility and the level. The facility refers to the type of program that required the

logging of a message. Facilities are named after typical UNIX services such as mail, ftp, or cron, subsystems such as the kernel, the printer, or the clock, and others are reserved for local use. The level indicates the importance or seriousness of the message. Possible values for the level are critical, warning, notice, etc.

The Nextthink components in the Appliance use the system log service to keep a record of significant occurrences, including:

- Audit trail events
- System alerts
- Investigation-based global alerts
- Internal state of the Engine

For writing to the system log, the Appliance relies on the *rsyslog* package, which has become the default logging service in many Linux distributions. Although it adds new advanced features, rsyslog still keeps backwards compatibility with the configuration files of the original syslog daemon. If you are familiar with the configuration of rsyslog, you may easily customize the output of the logs written by the Nextthink components and adapt them to your needs.

From this point on, we may refer to rsyslog as syslog when we talk about the logging service in general and not about specific features of rsyslog.

Default configuration and log files

The configuration file for rsyslog is found in `/etc/rsyslog.conf`. For the sake of clarity, the specific modifications of Nextthink to the configuration of rsyslog are stored in a separate file, which is found in `/etc/nextthink/nx_rsyslog.conf`. This file is applied to the main configuration file by means of an include directive in `/etc/rsyslog.conf`.

The default configuration of Nextthink dispatches log messages to different files depending on their content. Find these log files under `/var/log/nextthink`:

File	Purpose
<code>alert.log</code>	<ul style="list-style-type: none">• System and investigation-based global alerts• Debug info from logger (rsyslog)
<code>audit.log</code>	Audit trail events

engine.log	Internal state of the Engine
------------	------------------------------

By default, the Portal, the Engine, and the Web Console write their audit events to the `audit.log` file of the Appliance that hosts each one of them. Only the Engine writes to the `alert.log` and the `engine.log` files. In turn, the Portal does not use the `syslog` service to write information about its internal state, but its own logging tools. The internal logs of the Portal are found under `/var/nexthink/portal/log`.

The audit and alert logs are suitable for automatic processing, since their format is well-defined and stable. However, the format of the internal logs of the Engine is not guaranteed and may be subject to change. Therefore, do not rely on the contents of the Engine log for automating your processes.

Nextthink uses UTF-8 encoding for its log messages. Rsyslog preserves the encoding.

Configuration of alerts

In addition to email, you can use the system log as a notification mechanism for both system and global alerts. For global alerts, you need to enable `syslog` notification when creating the alert.

The part of the `syslog` configuration file `/etc/nexthink/nx_rsyslog.conf` which is relevant for alerts is shown below:

```
$template
RFC5424format,"<%pri%>1 %timestamp:::date-rfc3339% %hostname%
%programname% %procid%msg%\n"
...
# alerts
local5.=notice -/var/log/nexthink/alert.log;
...
# alerts
local6.=notice -/var/log/nexthink/alert.log; RFC5424format
```

The first line defines an output format for `syslog` messages by means of a template. The template is named *RFC5424format* because it follows the recommended format for `syslog` messages which is described in the most recent Internet standard about the `syslog` protocol: RFC 5424. The template defines the output to be composed of a priority number followed by the timestamp, the host name, the program name, the id of the process which issued the `syslog` message and the message itself. Once defined in this way, a template can be applied to

one or several message filters.

For alerts, you can see that we declare two filters in the syslog configuration file, depending on the facility specified to log the alert. Both filters are instructed to write their output to the same file: `/var/log/nexthink/alert.log`. The minus sign before the file name is there to improve the performance of the syslog daemon. It indicates that syslog output to the file is buffered, so the syslog system will not directly write to the filesystem but to a buffer in memory and then really write to the disk once the buffer is full. The two filters however accept messages from different facilities. If the facility used is `local5`, rsyslog will use the default syslog output format. On the other hand, if the facility used is `local6`, rsyslog will use the output format defined by the template *RFC5424format* for every logged alert.

To choose between legacy (`local5`) or modern (`local6`) format for the log messages of global alerts, set the following parameter in the main section of the configuration file of the Engine (`/var/nexthink/engine/01/etc/nxengine.xml`):

```
<syslog>
  <legacy_alert_format>true</legacy_alert_format>
</syslog>
```

For details on the formatting of alerts, see the article on integrating alerts.

Logging to a remote server

The syslog protocol lets you send log messages through the network to be consumed by syslog servers other than the local Appliance.

To send log messages to a remote syslog server, modify each line in the syslog configuration file of Nexthink by substituting the name of the log file for the name or IP address of the receiving server. The name of the server must be preceded by a single or a double at-sign (`@` or `@@`), depending on whether you want to send the log messages via UDP or TCP, respectively. Follow the name or IP address of the remote server by a colon (`:`) and the port number where the server is listening for syslog messages. For example, to send different types of log messages to remote servers.

Remember to use either `local5` or `local6` entries in slave Appliances, depending on the setting for the Engine `legacy_alert_format` to be `true` or `false`, respectively. For master Appliances, recall that the Portal always uses syslog `local5` facility and exclusively for audit events:

```
# Send general log to a server listening to UDP port 514
local5.=debug;local5.=info;local5.=error @udp-server.example.com:514;
nxFormat

# Send audit logs to a server listening to UDP port 514
local5.=warning @udp-server.example.com:514; nxAuditFormat

# Send alert logs to a server listening to TCP port 10514
local5.=notice @@tcp-server.example.com:10514;
```

Note that you do not have to choose between saving the logs in a file and send them to a remote server. It is possible to do both by repeating the same line in the syslog configuration changing the destination of the logs. Check the rsyslog documentation for options when sending log messages through the network, specially when using TCP.

Logging accesses to the CLI

Besides user access to the Nextthink components such as the Finder and the Portal, the access to the command line interface of the Appliance is an event of interest in the audit trail.

To filter the syslog messages related to accesses to the CLI of the Appliance and send them to a destination of your choice, specify the programs that control the command line inside conditional statements in the configuration file of syslog:

```
# Log access to the CLI
if $programname == 'sshd' then -/var/log/nexthink/audit.log
if $programname == 'sudo' then @udp-server.example.com:514
if $programname == 'login' then @@tcp-server.example.com:10514
```

These programs control remote access (**sshd**) to the Appliance, logging in (**login**) to the Appliance, and execute as the superuser (**sudo**) in the Appliance. In the example above, each program is sending its output to a different destination, but you can send the output of all programs to the same destination.

Restarting the Engine and the syslog service

Restart the Engine if its configuration file required any change:

```
sudo systemctl restart nxengine@1
```

After any modification to the configuration file of syslog, restart the service for the changes to be effective:

```
sudo systemctl restart rsyslog
```

Related tasks

- Creating an investigation-based alert
- Integrating alerts
- Examining the logs in the Portal

Related references

- System alerts
- Audit trail
- Rsyslog (external link)

Reporting the URL of HTTP web requests

If you have purchased the Web and Cloud module, you may set up the Collector to send the URLs of those HTTP web requests that the end-users address to a selected group of domain names. By default, for every web request, the Collector only reports the domain name inside the request to the Engine (and not the full URL) to keep the amount of generated network traffic low and avoid flooding the Engine with lots of URLs. Nevertheless, when the Collector is allowed to report the URLs of just a few web requests, the generated traffic still remains reasonably low, while you may benefit from this additional information to define services based on particular URL paths or investigations that include conditions on URLs of web requests.

Learn in this chapter how to specify the list of domain names for which the Collector must report the URLs of the HTTP requests that are addressed to them from the devices of the end-users.

Accepted syntax for the list of domains

Independently of the method chosen to configure the Collector, the accepted syntax for specifying domains is the same. The allowed characters to write

domain names are a subset of the ASCII character set that comprises:

- The range of letters from **a** to **z** and from **A** to **Z**.
- The digits from **0** to **9**.
- The symbols **.** (dot) and **-** (hyphen).
- The symbols **:** (colon) and **/** (slash).
- The symbol ***** (star) to substitute zero or more characters.

Let us see some examples of domain names and how are they interpreted by the Collector:

www.example.com	Matches all HTTP requests addressed to www.example.com
http://www.example.com	Same as above: matches HTTP requests to www.example.com
example.com	Matches all HTTP requests to example.com
http://example.com/index.html	Matches the same as example.com (the URL path after the host name is ignored)
*.example.com	Matches any prefix before the first dot (e.g. www.example.com and ftp.example.com, but not example.com)
*example.com	Matches any prefix (e.g. www.example.com, ftp.example.com, example.com, another-example.com)
***example.com	Same as above (multiple consecutive stars count as one)
ftp.example.com	Matches all HTTP requests addressed to ftp.example.com (Note that the protocol is HTTP and not FTP)
ftp://ftp.example.com	Error: only HTTP scheme is allowed
https://example.com	Error: only HTTP scheme is allowed
-example.com	Error: domain names cannot begin or end with a hyphen
*	Error: the <i>match all</i> star pattern is not allowed alone

Configuring the list of domains in the Collector

Specify the list of the domains for which the Collector reports the URLs of web requests either before or after deploying the Collector:

- Before deploying the Collector:
 - ◆ Passing parameters to the MSI.
 - ◆ Using the Nextthink Collector Installer.
- After deploying the Collector:
 - ◆ Using the Nextthink Collector Configuration Tool.
 - ◆ Changing the value of a registry key.

Beware that if you use the Updater to deploy the Collector, many parameters of the MSI, and the list of domains in particular, cannot be set at installation time and are not saved between updates. For every automatic update of the Collector, you must reapply the settings after deployment.

Passing parameters to the MSI

Specify the list of domain names by setting the value of the parameter **DRV_WEB_AND_CLOUD_HOSTS** when you install the Collector using its MSI file. The value supplied must be a comma separated list of the domains with the syntax defined in the previous section.

This option requires the parameter **DRV_WEB_AND_CLOUD_DATA** to be set to 1 (its default value) for the Collector to gather web related information.

Using the Nexthink Collector Installer

If you use the Nexthink Collector Installer to deploy the Collector, specify the list of domains for which you want to get the full URLs in the **Web And Cloud Settings** dialog that appears when you click the **Settings** button:

In the case that you are updating the Collector, the new settings replace any previously configured list of domains.

Using the Nexthink Collector Configuration Tool

If you have already deployed the Collector, use the Nexthink Collector Configuration Tool to modify the list of domains for which to report full URLs accessed from a particular device. This requires the presence of the Nexthink Collector Configuration Tool in the device; which is installed along with the Collector by default, unless you set the MSI option **CFG_INSTALL** to 0.

Execute the tool with administrator privileges and specify the list of domains as a parameter in the command line with domains separated by commas:

```
C:\Windows\System32\nxtcfg.exe /s wm_domains="csv_list_of_domains"
```

Setting the value of a registry key

The list of domains for which to report full URLs is saved in the registry under:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\nxtrdrv\params\hosts
```

If you change the value of this variable, the Collector detects its modification and applies the changes accordingly. If an error is detected in the syntax of a domain, the error is logged but the service just skips to the next domain in the list. Under high load, the Collector can miss the modification of the environment variable and you must reboot to force the change. For this reason, this method is recommended only for testing in pre-production environments.

For debugging purposes, it is allowed in this case to use the *match all* star pattern: *. This is the only exception to the rule and it may help you detect connectivity problems in a particular device.

Technical and security limits

By using any of the described methods, you can specify up to a maximum of 20 domains. The Collector limits the length of a URL to a maximum of 1024 characters. In the rare case of processing a URL longer than 1024 characters, the Collector truncates it to the first 1024 characters.

Note that the feature is only available for HTTP and not for HTTPS web requests. Due to TLS encryption, it is not possible to get the URLs of HTTPS requests. Moreover, reporting the exact URL of an HTTPS request might incur in a security or privacy breach.

In the same sense, the Collector never reports the *query string* part of a URL, that is, the optional list of parameters used by web applications that is placed at the end of the URL after a question mark. Query strings often carry sensitive information such as login names and passwords.

Related tasks

- Creating a service
- Specifying URL paths of web-based services

- Installing the Collector

Related references

- Collector MSI parameters reference table
- Nextthink Collector Configuration tool

Mobile Bridge configuration settings

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

```
<!-- The Address, Username and Password settings must be configured  
via the command line:  
Nextthink.Mobile.Bridge -username <username@domain>  
                        -address <myserver.example.com>  
you will be prompted for the password -->
```

```

<appSettings>
  <add key="Address" value="example.com" />
  <add key="Username" value="bridge@example.com" />
  <add key="Password" value="HASH" />
  <add key="UseSsl" value="true" />
  <add key="AuthenticationMechanism" value="Default" />
  <add key="SkipCACheck" value="true" />
  <add key="SkipCNCheck" value="true" />
  <add key="SkipRevocationCheck" value="true" />
  <add key="Secret" value="SECRET" />
  <add key="Port" value="11031" />
  <add key="NumberOfRequestPerSession" value="20000" />
  <add key="Timeout" value="60000" />
  <add key="Throttle" value="0" />
  <add key="MaxAgeInMinutes" value="60" />
  <add key="InLoopWaitInSeconds" value="60" />
  <add key="FailureDelayInSeconds" value="600" />
  <add key="UserRefreshPeriodInHours" value="24" />
  <add key="NumberOfRequests" value="64" />
  <add key="ExcludedGroupDn" value="" />
  <add key="IncludedGroupDn" value="" />
</appSettings>
</configuration>

```

Related tasks

- Installing the Mobile Bridge

Collector MSI parameters reference table

Mandatory parameters

Option Name	Default value	Description
DRV_IP	-	Set the Engine IP or DNS name
DRV_PORT	-	Set the Engine port number

Optional parameters

Option Name	Default value	Description
CFG_INSTALL	1	Install the Nxtcfg tool for changing the configuration of the Collector from the command line. 1: install, 0: do not install
CPL_INSTALL	0	Install the Collector Control Panel extension. 1: install, 0: do not install
DRV_ACTIVATE_DMP	0	<p>Specifies whether the target system should be configured for generating memory dumps in case of STOP message (System crash). Its value can be 0 (disabled), 1 (full memory dump), 2 (kernel memory dump) and 3 (memory minidump). The recommended value is 2 (kernel memory dump).</p> <ul style="list-style-type: none"> • This is a non-reversible setting: it will not be rolled back to its initial value after uninstalling Collector 3. • The MSI package will not change the system setting for a less verbose memory dump setting (e.g. if current setting is to generate kernel memory dumps and DRV_ACTIVATE_DMP is set to 3 (memory minidump), no action will be performed)
DRV_BFBD	0	Delay in seconds during initialization of the driver before we start sending UDP packets to the Engine. Maximum value: 240 (4 min)
DRV_CRASHGUARD	3	Specifies the maximum CrashGuard count can reach before the Collector driver loading is being cancelled at boot-time. If set to 0, the CrashGuard feature will be disabled
DRV_DESC	0	Delay Engine communication Socket Creation : To avoid having the traffic blocked by certain firewalls, the Collector

		socket layer is created during the last initialization steps. [1: enable, 0: disable]
DRV_LOGSIZE	32	Addition of log rotation when enabling DRV_LOGMODE for the logging [Range for value: 1 -> 512 (MB)].
DRV_REACTIVATION	96	Reactivate the collector after a given time. The max value is 8766 --1 year.
DRV_TAG	0	Assign to any installer to help you organize and remember the creator of the installation. Possibles values are 0 to 2147483647.
DRV_LOGMODE	0	Specifies the logging mode. Possible values are 0, 1 and 2, meaning Silent, Verbose and Debug, respectively. 2 (Debug) is not recommended.
DRV_FORCE_SERVER	0	Specifies the installation on a server. Possible values are 0, 1. To prevent the installation of the Collector on Windows Server, set to 0. To allow installation on Windows Server, set to 1.
DRV_DWEF	0	Disables Windows enumerate functionality. Possible values are 0, 1. If set to 1, the Collector does not report any Windows freeze or hung problems. (This will result in the Finder not displaying any information about "application not responding".)
DRV_CGPI	0	?CrashGuard Protection Interval Value?. It is the time since boot in minutes after which we save the CrashGuard info to the registry.
DRV_MSS	1224	Maximum size of the UDP packet for transfers between the Collector and the Engine. Allowed values range from 1000 to 16384.
DRV_WEB_AND_CLOUD_DATA	1	Gather Web and Cloud information. Default value is 1 to gather and send the data (only if you have purchased the Web and Cloud module). Set to 0 for not recording the web connections of devices.
DRV_WEB_AND_CLOUD_HOSTS	-	List of comma separated host names for which to send the full URL of each web request. Requires the Web and Cloud module and the parameter

		DRV_WEB_AND_CLOUD_DATA to be set to 1.
DRV_DSPS	1	Disable SMB print notifications. Starting from version 5.2.8.0, the Collector does not report SMB prints by default. Set the option to 0 to enable SMB print reporting. Set to 1 to disable it.
DRV_PREFERIPV6	0	Favor IPv6 over IPv4 (or viceversa) for communicating with the Engine. When the DNS lookup of the name of the Engine resolves to both IPv6 and IPv4 addresses, prefer IPv6 when set to 1 and IPv4 when set to 0.

Windows parameters

Option Name	Default value	Description
ARNOREMOVE	-	Setting the ARPNOREMOVE property disables the Add or Remove Programs functionality in Control Panel that removes the product. For Windows 2000, this disables the Remove button for the product from the Add or Remove Programs in Control Panel. For earlier operating systems, this has the effect of removing the product from the list of installed products on the Add or Remove Programs in Control Panel.
ARNOREPAIR	-	Set the ARPNOREPAIR property to disable the Repair button in the Programs Wizard.
ARPSYSTEMCOMPONENT	-	Setting the ARPSYSTEMCOMPONENT property to 1 using the command line or a transform prevents the application from being displayed in the Add or Remove Programs list of Control Panel.
ARNOMODIFY	1	Setting the ARPNOMODIFY property disables Add or Remove Programs functionality in Control Panel that modifies the product. For Windows 2000, this disables the Modify button for the product in Add or Remove Programs in Control Panel. On earlier operating systems, clicking the Add or Remove Programs button uninstalls the product rather than entering the maintenance mode wizard. Note: the Collector MSI package does not support this feature. ARPNOMODIFY must be set to 1.
REBOOT	-	The REBOOT property suppresses certain prompts for a restart of the system. An administrator typically uses

this property with a series of installations to install several products at the same time with only one restart at the end.

The ForceReboot and ScheduleReboot actions inform the installer to prompt the user to restart the system. The installer can also determine that a restart is necessary whether there are any ForceReboot or ScheduleReboot actions in the sequence. For example, the installer automatically prompts for a restart if it needs to replace any files in use during the installation.

You can suppress certain prompts for restarts by setting the REBOOT property as follows.

REBOOT = Force Always prompt for a restart at the end of the installation. The UI always prompts the user with an option to restart at the end. If there is no user interface, and this is not a multiple-package installation, the system automatically restarts at the end of the installation. If this is a multiple-package installation, there is no automatic restart of the system and the installer returns ERROR_SUCCESS_REBOOT_REQUIRED.

REBOOT = Suppress Suppress prompts for a restart at the end of the installation. The installer still prompts the user with an option to restart during the installation whenever it encounters the ForceReboot action. If there is no user interface, the system automatically restarts at each ForceReboot. Restarts at the end of the installation (for example, caused by an attempt to install a file in use) are suppressed.

REBOOT = ReallySuppress Suppress all restarts and restart prompts initiated by ForceReboot during the installation. Suppress all restarts and restart prompts at

		the end of the installation. Both the restart prompt and the restart itself are suppressed. For example, restarts at the end of the installation, caused by an attempt to install a file in use, are suppressed.
--	--	--

Starting from V6, the Collector is usually able to upgrade without the need to reboot the device. Only when migrating from V5 or when the target device interferes with the installation process (for instance, by running the Collector Control Panel extension during installation), a reboot is necessary. Set the REBOOT option in these cases to specify your choice.

For instance, if you do not want your devices to reboot right away after a V5 to V6 migration, set REBOOT=ReallySuppress. As a drawback, if you set this option, the upgrade to V6 will not be complete until the end-users reboot their devices.

In unattended execution mode, all choices are silently accepted. For example, if REBOOT=Force, the computer will automatically be rebooted after the MSI package installation.

Casing of properties

Always specify the names of the parameters (the properties) of the MSI in capital letters. If you include the properties with lower case letters in an MST, they will be considered private properties and you will not be able to modify them later from the command-line.

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.

Nxtcfg - Collector configuration tool

Nxtcfg is a small console application that allows to read and modify the configuration parameters of the Collector. Please, make sure to run Nxtcfg with administrator privileges.

Installation

By default, the Nxtcfg tool is installed along with the Collector when installing the Collector MSI. Once the Collector is installed, the Nxtcfg tool is located under

C:\Windows\System32\nxtcfg.exe.

The Collector MSI version used determines the nxtcfg version installed (Windows 32-bit or 64-bit system).

If not required, add the option CFG_INSTALL=0 to the MSI command line, when installing the Collector.

Options

Option	Description	Example
/disable	Turn off the Nextthink Collector (the driver is kept in memory in idle state).	nxtcfg.exe /disable
/enable	Turn on the Nextthink Collector.	nxtcfg.exe /enable
/g	Get the value of a particular configuration parameter from the Collector.	nxtcfg.exe /g ip
/s	Set the value of one or more configuration parameters of the Collector.	nxtcfg.exe /s ip=192.168.0.1 port=999
/l	List all the configuration parameters of the Collector with their current values.	nxtcfg.exe /l
/d	Dump all the configuration parameters of the Collector and their corresponding values to a file.	nxtcfg.exe /d C:\temp\collector.cfg

Configuration parameters

Parameter	Description	Default value	Range
ip	IP address or DNS name of the Engine.	-	-
port	UDP port number where the Engine is listening.	-	[1 - 65535]
tag	Optional number to identify the installation.	0	[0 - 2147483647]
cgpi	<i>CrashGuard Protection Interval Value</i> . It is the time interval since boot (in minutes) after which a dirty reboot does not increase the CrashGuard.	0 min	-
logmode	Logging mode <ul style="list-style-type: none">• 0 - Silent• 1 - Verbose• 2 - Debug (not recommended for production)	0	[0 - 2]

logsize	Maximum size of log file when logging is enabled. Logs are rotated after the maximum is reached.	32 MB	[1 - 512] MB
dsps	Disable (1) or enable (0) SMB print monitoring functionality	1	[0 - 1]
iops	Enable (1) or disable (0) IOPS monitoring functionality	0	[0 - 1]
dwef	When set, the Collector does not report application freezes nor hungs.	0	[0 - 1]
mss	Maximum size, in bytes, of the UDP packets sent from the Collector to the Engine.	1224 B	[1000 - 16384] B
wme	When set, the Collector reports Web and Cloud data.	1	[0 - 1]
wm_domains	List of domains for which to report the URL of web requests	-	Comma separated domain names
prefer_ipv6	When set, the Collector prefers IPv6 to communicate with the Engine when the name of the Engine resolves to both IPv6 and IPv4 addresses.	0	[0 - 1]

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

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

Auditing logon events

For Nextthink to report accurate logon times and logon durations, especially in the case that you use roaming user profiles in your Windows setup, configure the audit of logon events in all your devices. You can do so with the help of Active Directory by applying a GPO to the domain of your devices.

Enabling the audit of logon events

To enable the audit of logon events:

1. Open the **Group Policy Management Console**.
2. Right-click the domain node of your devices and select the option **Create a GPO in this domain, and Link it here....** A dialog to create the new GPO shows up.
3. Type in the name of the GPO. For example, *Logon Audit Policy*.
4. Click **OK** and the new GPO appears in the tree.

5. Right-click the newly created GPO and select the option **Edit...** The console displays the settings for the GPO.
6. Expand the node **Computer Configuration** and navigate to **Windows Settings / Security Settings / Local Policies / Audit Policy**.
7. Double-click the policy **Audit logon events**.
8. Check the **Success** and, optionally, the **Failure** options.
9. Click **OK** to save your changes.
10. Run the command **gupdate /force** to update the GPO.

The devices in the specified domain now record the logon events in the Security log.

Overwriting or clearing events from the Security log

After you activate the audit of logon events, make sure that the Security log of Windows always has enough space to save new logon events. Set the properties of the Security log to perform an appropriate action when the maximum size of the log is reached:

- **Overwrite events as needed (oldest events first).** *Recommended.*
- **Archive the log when full, do not overwrite events.**
- **Do not overwrite events (Clear logs manually).**

Use the preferred first option to avoid problems with the size of the Security log.

If you choose the last option and the Security log runs out of space, you may no longer be able to log in to the device. Indeed, if the Security log is full and events are not overwritten, trying to write an audit logon event to the log fails, making the whole login procedure fail as well.

Related references

- [Boot and logon duration](#)

Redirecting Collector traffic

Overview

For testing or redundancy purposes, redirect the Collector traffic that reaches one Engine to other Engines.

Configure the redirection service *nxredirect* that runs on the Engine appliance to forward the traffic received from the Collectors to other Engines of your choice.

Configuring the redirection service

For the redirection service to automatically start after every system boot:

1. Log in to CLI of the Engine.
2. Enable the redirection service:

```
sudo systemctl enable nxredirect
```

To configure the redirection service:

1. Log in to the CLI of the Engine.
2. Open or create the configuration file of the redirection service:

```
sudo vi /etc/nexthink/nx_redirect.conf
```
3. Write some redirection rules (see below for examples).
4. Save your changes and exit:

```
:wq
```
5. Restart the service:

```
sudo systemctl restart nxredirect
```

Writing redirection rules

The following lines are a sample configuration of the redirection service:

```
listenraw port=999  
[dst=192.168.0.25:997,192.168.0.26:997 send]
```

1. The first line tells the *nxredirect* service to listen to the traffic received by all interfaces on port 999.
2. The second line sends the received Collector packets to port 997 of the Engines with IP addresses 192.168.0.25 and 192.168.0.26.

Anonymizing redirected traffic

For generic data analysis purposes, you may want to have access to all the data in an Engine related to services, connections, executions, etc. without necessarily associating them to a particular person or group of people. That is, you may want to analyze the data collected while keeping users, devices, and printers anonymous.

To have a redundant Engine that holds all significant data while hiding sensitive information about users, devices, and printers, redirect traffic to that Engine with anonymization turned on. To anonymize Collector traffic, configure the redirection service as in the following example:

```
listenraw port=999
[anon=encryption_key dst=192.168.0.27:998 send]
```

Note the addition of the **anon** keyword, followed by the encryption key of your choice. For the sake of efficiency, use it preferably before specifying the destinations, specially if you have many. In that way, anonymization takes place only once before replicating and splitting the traffic.

When anonymizing Collector traffic, some fields of the device, the user, and the printer objects are encrypted, other fields are randomized, and others are removed.

Device anonymization

Device	Field	Action
Properties	SID	Randomized
	Name	Encrypted
Network	Last IP address	Replaced by Engine IP
	IP addresses	Replaced by Engine IP
	MAC	Randomized
	Group name	Encrypted
Operating system	Windows license key	Removed
Hardware	BIOS serial number	Removed
	Chassis serial number	Removed
	Device UUID	Removed
Active Directory	Distinguished name	Not retrieved

Note that a change in the encryption key implies a duplication of the devices. If you are redirecting to an existing Engine, remember to erase the database to avoid duplications.

User anonymization

User	Field	Action
Properties	SID	Randomized
	Name	Encrypted
Active Directory	Distinguished name	Not retrieved
	Full name	Not retrieved
	Department	Not retrieved
	Job title	Not retrieved

Printer anonymization

Printer	Field	Action
Properties	Name	Encrypted

Support for DirectAccess

Overview

Microsoft DirectAccess is a technology that provides remote connectivity to devices equipped with Windows 7 and higher operating systems. Similar in concept to a traditional virtual private network (VPN), DirectAccess allows users to securely access network resources inside the intranet of their organization when connected to the Internet. Unlike traditional VPN connections, which usually require explicit user action to be initiated and terminated, DirectAccess is transparent to the end user and automatically connects to the intranet of the company when needed.

DirectAccess relies on clients and applications that support the IPv6 stack. It encapsulates the traffic to route it through the Internet and, once it reaches the intranet, a companion technology transforms the IPv6 addresses into IPv4 if needed; that is, if the intranet uses IPv4 internally, which is usually the case.

Impact on Nexthink

Since DirectAccess requires client applications to use IPv6, three Nexthink products are impacted when a set of devices in your organization connect to the corporate network via DirectAccess: the Collector, the Engine, and the Finder.

Collector

The Collector must be able to send information to the Engine from devices that connect to the intranet of their organization through DirectAccess. Therefore, the Collector must use IPv6 to send its information. In addition, the Collector must be able to capture network information of those applications running on devices connected through DirectAccess, which also use the IPv6 stack.

When installing the Collector in a DirectAccess environment, check the option **Prefer IPv6** when running the Collector installer, or the MSI parameter **DRV_PREFERIPV6**, for the Collector to use IPv6 rather than IPv4 to send information. You can equally modify the value of this setting when the Collector is already installed with the help of the Collector configuration tool by adjusting the value of the parameter **prefer_ipv6**.

Engine

The Engine must be able to detect Collector traffic coming from DirectAccess and translate the received IPv6 addresses to their IPv4 counterparts within the intranet. To identify Collector traffic, the Engine needs to know the IPv6 subnetwork used by DirectAccess.

By default, the Engine identifies and translates IPv6 addresses in the subnet fda9:11e5:84fa::/48. If you use a different subnetwork, configure the Engine as in the following example, substituting the DirectAccess prefix given for your own:

1. Stop the Engine

```
nxinfo launch -r
```

2. Configure the IPv6 subnet:

```
sudo nxinfo config -s  
"direct_access.prefix=fda9:11e5:84fa::/48"
```

3. Restart the Engine

```
nxinfo launch -a
```

Finder

The Finder must be able to connect to both the Portal and the Engine even when run from a device connected to the corporate network via DirectAccess. In the case of the Finder, no additional configuration is needed, but you must use DNS names in the login dialog to resolve the address of the Portal, because the dialog does not support IPv6 addresses.

Changing the thresholds of High CPU warnings

Overview

High CPU warnings for devices and executions are triggered when the CPU load exceeds some default values. The default values have been chosen to detect both significant high CPU loads in a device and the particular applications that cause high CPU load during their execution.

If you receive too many high CPU warnings in your setup, up to the point that they stop being meaningful, raise the default thresholds. To change the default thresholds, edit the configuration file of the Engine:

1. Log in to the CLI of the Engine.
2. Edit the configuration file:

```
sudo vi /var/nexthink/engine/01/etc/nxengine.xml
```
3. Change the high CPU settings inside the tag **<aggregation>** (under **<config>**, **<engine>**). See below each possible individual setting.

Repeat this operation in every Engine of your setup.

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.

Device warnings

The deprecated device warning **High thread CPU usage (deprecated)** is triggered when the CPU load in a device exceeds 80% of a single logical processor. To change that threshold, modify the value of the following setting:

```
<machine_high_cpu>80</machine_high_cpu>
```

The device warning **High overall CPU usage** is triggered when the CPU load is above 70%, taking into account all the logical processors of the device. This threshold is independently modifiable for each type of device (laptop, desktop, or server):

```
<normalized_high_cpu_laptop>70</normalized_high_cpu_laptop>  
<normalized_high_cpu_desktop>70</normalized_high_cpu_desktop>  
<normalized_high_cpu_server>70</normalized_high_cpu_server>
```


Execution warnings

By default, for any process to trigger a **High thread CPU usage** warning, it has to take more than 50% of CPU load. The threshold is controlled by the following setting:

```
<process_high_cpu>50</process_high_cpu>
```

In the case of the system process, the threshold is lowered to 40%. Change the default with the following setting:

```
<system_high_cpu>40</system_high_cpu>
```

Related references

- Errors and warnings for devices and executions

Maintenance

Logging in to the CLI

The command line interface (CLI) of the Nextthink Appliance gives you access to a terminal where you can inspect and control every aspect of the system by using all the power of the Linux shell.

To log in to the CLI, connect to the Appliance with the help of an SSH client as the user **nextthink**.

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

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

Planning for disaster recovery

The Nextthink Appliance provides you with different backup techniques that allow you to recover from either a partial or a full disaster:

- A partial disaster is a failure that affects one or several of the server components of Nextthink (Web Console, Engine or Portal), while the Appliance is still accessible.
- A full disaster is a complete system failure that prevents any further access to the Appliance.

The mechanisms for partial disaster recovery are automatically put in place after the installation of the Appliance. Each one of the server components in the Appliance generates a daily backup of its data for its own recovery. In this way, if any of the components crashes, you can at least get the component back to the state it had the day before the crash.

Full disaster recovery, on the other hand, requires you to save the backups to an external storage device outside the Appliance before total breakdown. You can automate this process by activating the provided mechanism to [save backup files externally](#). If you want to install your own backup tool, first read and follow the recommendations of the article on installing third-party software in the Appliance. Beware that a serious hardware issue in your Appliance can make your data unrecoverable if you do not save it elsewhere.

Partial disaster recovery

In case of a server component malfunction, use its daily backup files for recovery. In addition to the daily backups, the server components make an automatic backup of their data before migration as well. That is useful in the case that the software upgrade process goes wrong.

To learn about the information that is saved during the backup process and how to recover from a partial disaster, read the corresponding documentation for each component:

- Web Console automatic backup and Web Console restore
- Engine automatic backup and Engine restore
- Portal automatic backup and Portal restore

Full disaster recovery

In case of a total failure of the Appliance, you need to be ready to start from anew. As a prerequisite, you must have previously saved the backups of all the server components in the Appliance to an external storage device. Remember that you can automate this process by [activating external backups](#) from the Web Console.

It is also recommended that you manually save the license file that is stored in the appliance that hosts the Portal once you activate the product. You must save it only when you receive a new license, so it is not included in the automatic external backups.

To perform full recovery:

1. Download an Appliance ISO with the same version of the Appliance that failed.
2. Install the Appliance following the steps described in [Installing the Appliance](#).
3. Choose to install either the Portal or the Engine as described in [Engine & Portal Installation](#), depending on the main server component that your Appliance was running.
4. Copy the backups to the new Appliance using any SCP client.
5. Restore the Web Console first as described in [Restoring the Web Console](#) to set the general parameters of the Appliance.
6. Restore the installed server component: Engine or Portal, as documented in [Restoring the Engine](#) or [Restoring the Portal](#).

7. In the case of a complete failure of the appliance that hosts the Portal, restore the license file.

Activating external backups

The Appliance provides a mechanism to automate the saving of backup files to an external SMB share. This mechanism makes a copy of the daily backup of every server component (Web Console, Engine or Portal, Console) to the SMB share right after the backup file is created.

Before activating external backups, you must set up the SMB share:

1. Configure the user account
2. Set the permissions on the destination folder
3. Share the folder

To activate external backups in the Appliance:

1. Log in to the Web Console as admin from a web browser:
`https://<IP_address_of_Appliance>:99`
2. Select the **External backup** section in the **Appliance** tab.
3. Tick the Active box
4. Fill in the rest of the fields in the form as shown in the screenshot below, according to the settings of your SMB share:

The files saved in the SMB share for the different components have the following format:

- Web Console:
console-<timestamp>.tgz
- Engine:

- **nxengine-<instance>-<hostname>-<timestamp>.tgz**
- Portal (main backup and history details of count metrics):
 - **portal-<hostname>-<timestamp>.tgz**
 - **portal-<hostname>-history_YYYYMMDD-<timestamp>.backup**

For advanced users, it is possible to customize the mount options of the SMB share for external backups. These are the options found after the **-o** flag of the **mount** command. By default, the Appliance mounts the SMB share using the options **guest** and **credentials**. After activating external backups via the Web Console, set additional mount options for the SMB share by editing the global config file:

1. Log into the CLI of the Appliance.
2. Edit the global configuration file of the Appliance:


```
sudo vi /etc/nexthink/nexthink-config.xml
```
3. Inside the section **smb_backup** add a new entry to specify one or more additional options, separated by commas:


```
<additional_mount_options>option</additional_mount_options>
```
4. Save your changes and exit:


```
:wq
```

The resulting configuration file should look like this:

```
<?xml version="1.0"?>
<config>
  <global>
    <smb_backup>
      ...
      <additional_mount_options>option=value</additional_mount_options>
    </smb_backup>
  </global>
</config>
```

Related tasks

- Web Console backup and restore
- Engine backup and restore
- Portal backup and restore
- License backup and restore
- Installing third-party software in the Appliance

Web Console backup and restore

Manual Backup

To manually back up the Web Console:

1. Log in to Command Line Interface of the Appliance.
2. Execute the following command to create a new backup. Optionally specify a different name for the backup file without the extension (tgz is automatically added):

```
sudo /var/nexthink/console/helpers/backup-console.sh  
[backup-file]
```

The backup file contains the full database of the Web Console (**console-db.backup**) and the content of the following files:

- /etc/nexthink/nexthink-config.xml
- /etc/yum/pluginconf.d/proxy.conf
- /var/nexthink/console/etc/certificate.pem

Find the backup file in the directory:

```
/var/nexthink/console/backup
```

Automatic Backup

Every day at 01:10 an automatic backup is triggered using a crontab entry. Up to 10 backup files are used to keep history, all located in the directory:

```
/var/nexthink/console/backup
```

A link file named console-backup.tgz is also created in that directory and points to the last backup.

Restoring the Web Console

To completely restore the Web Console settings and account configuration, log in to the shell of the Appliance, get your backup file, and follow the next steps:

1. Stop the Web Console:
sudo systemctl stop nxconsole

2. Untar your backup file (suppose that it is named console-backup.tgz) in a directory in your home:
mkdir console-bk
tar xvzf console-backup.tgz -C console-bk
3. Copy the configuration files in the backup to their intended location:
cd console-bk
sudo cp etc/nexthink/nexthink-config.xml /etc/nexthink
sudo cp etc/yum/pluginconf.d/proxy.conf
/etc/yum/pluginconf.d
sudo cp var/nexthink/console/etc/certificate.pem
/var/nexthink/console/etc
4. Drop the database of the Web Console:
dropdb -U postgres console
5. Drop the *console* user of the database:
dropuser -U postgres console
6. Create an empty database:
/var/nexthink/console/helpers/create-db.sh
7. Restore the database of the Console (console-db.backup file from the backup):
pg_restore -U postgres -d console console-db.backup
8. Restart the Console:
sudo systemctl start nxconsole

The Web Console is now restored with all its users and settings in place.

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

- Logging in to the CLI

Engine backup and restore

Manual backup

To make a complete backup of the Engine, execute manually the same script that is executed automatically during the daily backup of the Appliance:

1. Log in to the CLI of the Appliance running the Engine.

2. Execute the command:
sudo /var/nexthink/engine/common/bin/nightly_backup.sh
3. Optionally: If you want to keep the logs, copy the log files stored under:
/var/log/nexthink/
 - ◆ **engine.log**
 - ◆ **audit.log**
 - ◆ **alert.log**
 - ◆ All the compressed older logs stored in gz files.

For every instance of the Engine in the Appliance, the script creates a tgz file (GZIP compressed Tar archive format) with the contents of the Engine database and its configuration. Find the backup files under:

/var/nexthink/engine/<instance_number>/backups/nxengine-backup-<id>.tgz

Alternatively, you can make a backup of the Engine database only. Copying the database file while the Engine is running is not a good idea, because the Engine is continuously modifying the file, and the result could be a corrupted file. Instead, make a safe backup while your Engine is running by following these steps:

1. Log in to Command Line Interface of the Appliance.
2. Execute the following command to generate a compressed gz file with the database of the Engine:
nxinfo backup [--engine <instance_number>] --name <name_of_backup_file>

The file is copied to your current directory.

Automatic backup

The Appliance automatically makes a backup of all the Engines running on it via a cron job. The job is executed every day at 04h15 by default. Find the cron job specification under:

/etc/cron.d/nxengine-crontab

And the script executed in here:

/var/nexthink/engine/common/bin/nightly_backup.sh

The script makes a copy of each one of the the databases and the configuration of the Engines that are present in the Appliance and compresses them in separate tgz files. Backup files are stored in the instance backup directory:

`/var/nexthink/engine/<instance_number>/backups/nxengine-backup-<id>.tgz`

Specify the maximum number of local backups that are kept in the directory by editing the first variable in the script **nightly_backup.sh**. By default, up to ten backups are kept simultaneously:

```
NUM_OF_BACKUPS=10
```

In a test environment, you may want to disable automatic backups to save disk space in the Appliance. To that end, comment out the line that executes the nightly backup in the crontab file by prepending a hash sign (**#**) to it.

Each local backup file gets assigned a number from one to the maximum number of simultaneous backups in the directory. When the maximum number is reached, the count begins again and backup files are progressively overwritten. In order to get the most recent backup file, there is a symbolic link to the latest backup (note the absence of identifier):

`/var/nexthink/engine/<instance_number>/backups/nxengine-backup.tgz`

If external backups have been activated, the automatic script copies the daily backup to external storage right after generating it.

Restoring the Engine

Restore the Engine either in the same Appliance from which you made the backups or in a different Appliance. In the case that you are restoring your backups in another Appliance, make sure that its network configuration is the same as the configuration of the original Appliance. Otherwise, you may no longer receive data from the Collectors and have the wrong internal networks configured. In addition, the new Appliance requires you to reallocate the devices assigned to the original Appliance from the Portal. In case that you are using a license with online activation, this process should be transparent. If you are using a license with offline activation, you must repeat the procedure to get your license signed.

To restore a complete backup of the Engine:

1. Log in to CLI of the Appliance where you want to restore the Engine.
2. Stop the running Engine. Optionally, specify the instance of the Engine if running more than one:
`nxinfo launch -r [-g <instance_number>]`
3. Copy the backup file into the Engine directory:

- ```
sudo cp nxengine-backup-<id>.tgz
/var/nexthink/engine/<instance number>
```
4. Extract the database and configuration files from the backup file:

```
cd /var/nexthink/engine/<instance number>
sudo tar -xvzf nxengine-backup-<id>.tgz
```
  5. Remove the database of the Engine in place:

```
sudo nxinfo remove -r [-g <instance number>]
```
  6. Restore the database of the Engine backup:

```
sudo nxinfo restore -n /var/nexthink/engine/<instance
number>/data/nxengine-db.gz [-g <instance number>]
```
  7. Restart the Engine:

```
nxinfo launch -a [-g <instance_number>]
```
  8. Validate that the Engine is running properly:

```
nxinfo info
```

While the Engine is starting, the last command displays the message:

```
nxengine is booting...
```

After a few minutes, once the Engine has finished loading the database, the execution of this command displays the basic configuration and some statistics of the Engine. This means that the restore process was successful.

If you made a backup of the database only and you want to restore it in the current Appliance, you just need to restore the database of the Engine, and not any of the configuration files, which are already in place:

1. Log in to CLI of the Appliance where you want to restore the Engine.
2. Stop the running Engine. Optionally, specify the instance of the Engine if running more than one:

```
nxinfo launch -r [-g <instance_number>]
```
3. Remove the database of the Engine in place:

```
sudo nxinfo remove -r [-g <instance number>]
```
4. Restore the database of the Engine backup:

```
sudo nxinfo restore -n nxengine-db.gz [-g <instance number>]
```
5. Restart the Engine:

```
nxinfo launch -a [-g <instance_number>]
```
6. Validate that the Engine is running properly:

```
nxinfo info
```

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 Nextthink Certified Partner.  
Related tasks

- Planning for disaster recovery
- Setting up a software license
- Logging in to the CLI

## Portal backup and restore

### Manual Backup

To manually back up the Portal:

1. Log in to Command Line Interface of the Appliance that hosts the Portal.
2. Execute the following script, noting that you must not add any extension to the name of the target file. The script automatically appends the **.tgz** extension to the name of the backup file:

```
sudo /var/nextthink/portal/backup/backup-portal.sh
target-filename
```

In addition, if you want to take a backup of the history details of count metrics, you must have configured the Portal to automatically keep these history details day by day. See in the next section the directory where the Portal stores the backup files of history details. If the Portal has not been configured to store the history details, it is not possible to recompute them afterwards manually.

Copy the contents of the history directory to another location (e.g. to a USB key) to make a manual backup of the history details of count metrics:

```
cp -r /var/nextthink/portal/backup/history/ target-folder
```

### Automatic Backup

Every day at 22h15, a cron job triggers an automatic backup of the Portal. The automatic backup system keeps a history of up to ten backup files. The backup files are located in:

```
/var/nextthink/portal/backup
```

The file named **portal-backup.tgz** is a symbolic link that points to the last backup file in the history. The backup file holds the main database of the Portal and the content of the configuration folder:

```
/var/nexthink/portal/conf
```

In addition, if you have configured your Portal to store the history details of count metrics, that is, the lists of objects that contributed to the count metric on a particular day, these are stored under:

```
/var/nexthink/portal/backup/history
```

The name of history detail files has the format *history\_YYYYMMDD.backup*. The number of files kept for the history details depend on the disk space reserved for this purpose.

### ***On upgrade backup***

In addition to the automatic nightly backup of the Portal, the appliance automatically makes a new backup of the Portal before each upgrade. The file is placed in the same directory as the nightly backups and its name has the following format (where **X.X.X.X** indicates the version to which the Portal is upgrading):

```
/var/nexthink/portal/backup/portal-backup_before-X.X.X.X.tgz
```

## **Restoring the Portal**

To restore the Portal state from a backup file:

1. Log in to Command Line Interface of the Appliance that hosts the Portal.
2. Execute the restore script:

```
sudo /var/nexthink/portal/backup/restore-portal.sh \
[-d history_details_directory] <backup-filename>
```

If you saved the history details of count metrics, use the **-d** option to specify the directory that holds these files. The history files are expected to have the same name format specified above (*history\_YYYYMMDD.backup*). In the case that you configured the Portal to save the backups and history details to an external share, this name format is changed to *portal-<hostname>-history\_YYYYMMDD-<timestamp>.backup*. To restore the history files with the script, they must have their original name format. To rename all the history detail files stored in an external share, copy them to a directory in the Appliance and then type in the following command:

```
reg="(history_[0-9]+)"; \
for file in *.backup; do if [[${file} =~ ${reg}]];\
then mv $file ${BASH_REMATCH[1]}.backup;\
fi; done
```

```
fi; done
```

In the external share, you may have stored a set of details files whose total size exceeds the reserved disk size for history details configured in the Portal. Remember to manually select only the more recent files whose total size is within the configured limit. Use the command `du -h` in the folder containing the files with history details to get their total size, compare it to the value that you have configured in the Portal data retention, and remove the oldest files in the set until the total size of the files matches or is below the configured value. Failing to do so results in the Portal taking more time to restore history details that must be removed afterwards anyway, because there is no disk space left reserved for them.

The script only restores the database of the Portal, that is, the state of your dashboards. It does not restore the configuration files though, because you may want to keep your current configuration. If you need to restore the configuration of the Portal:

1. Stop the Portal:

```
sudo systemctl stop nxportal
```

2. Untar the backup file:

```
tar -xvzf portal-backup.tgz
```

3. Copy the contents of the **conf** directory to the Portal configuration directory:

```
sudo cp -r conf/ /var/nexthink/portal/
```

4. Restart the Portal:

```
sudo systemctl start nxportal
```

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

- Logging in to the CLI
- Planning for disaster recovery

#### Related references

- Data retention
- Nightly task schedules timetable

## License backup and restore

In the case of a complete failure of the appliance that hosts the Portal, the locally cached license may be lost as well. To avoid this, you can manually save a copy of your license file into the same shared folder that you use for your external Portal backups, for example.

### Manual backup

To save the cached license file:

1. Log in to the CLI of the appliance that hosts the Portal.
2. Copy the cached license files to an external storage medium, for instance, the external share for the Portal configured in the Web Console:

```
mkdir -p /<external_share>/LicenseRestore
sudo cp /var/nexthink/llm/data \
/{license.file, llm_private_key.txt, llm_public_key.txt} \
/<external_share>/LicenseRestore
```

### Restoring the license

Before restoring the license, restore first the Portal in the new appliance and configure the access to the storage medium (typically the external share) where you stored a copy of the license file.

To restore the cached license file:

1. Log in to the new appliance that hosts the Portal.
2. Copy the the license file backups to the correct folder in the appliance:

```
sudo cp /<external_share>/LicenseRestore \
/{license.file, llm_private_key.txt, llm_public_key.txt} \
/var/nexthink/llm/data/
sudo chown nxllicense:nxllicense /var/nexthink/llm/data/*
```
3. Restart the local license manager service:

```
sudo systemctl restart nxllicense
```
4. Check that the LLM works correctly:

```
sudo /var/nexthink/llm/bin/check.sh
```

### Recreating the license

In the case of a full disaster where you do not have an external backup of the cached license file, deactivate the product from the Portal:

1. Log in to the Portal as admin.

2. In the **ADMINISTRATION** menu, click **Licenses** under the **SYSTEM CONFIGURATION** section.
3. Click the button **Deactivate product** at the top right corner of the **Licenses** panel.
4. Ask Nexthink for a new license and reactivate the product.

#### Related tasks

- Planning for disaster recovery
- Portal backup and restore
- Logging in to the CLI
- Setting up a software license

## Finding out unlicensed devices

### Overview

When ordering a license for Nexthink, you must specify the number of end-user devices on which you want to deploy the Collector. From that total, you allocate the maximum number of devices to each Engine. If the number of devices reporting to an Engine is actually higher than the maximum number of allocated devices for that Engine, the Engine discards the devices in excess. No data are stored for these *unlicensed* devices and, therefore, they are not visible in the Finder or the Portal.

Learn here how to find out unlicensed devices from the Engine log, so you can adapt your license or your allocation strategy accordingly.

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.

### Reaching the maximum number of devices

When the Collector is installed in a device, it soon starts sending information relative to the device to its configured Engine. At that point, the Engine realizes that the device exists. When the Engine identifies a new device in the network, it increases its count of devices and checks its maximum number of allocated devices, according to the distribution of the license. If the count is within the limit, the Engine keeps the device and accepts the data coming from the Collector. On the other hand, if the count is above the limit, the Engine discards the device and

any other data coming from the Collector.

When the Engine discards a device because it exceeds the maximum number of allocated devices, it logs the following message:

```
machine <Name>|<MAC>[|<MAC>]* out of license
```

The message states the name of the device followed by its detected MAC addresses. While the Engine is running, it logs the message only once per device, even if the Collector in the device may keep sending information. If the Engine is restarted, it logs the message again as soon as it receives data from the offending Collector.

Note that if you choose to ignore a particular device by setting its storage policy to **none** or **remove**, the device will never appear in the log as unlicensed. In fact, those devices whose storage policy is set to **none** are still counted by the licensing mechanism as active, so they are not really unlicensed. On the other hand, devices with their storage policy set to **remove** are discarded immediately by the licensing mechanism. They never appear in the log of unlicensed devices though, because they no longer send any information to the Engine when it actually removes them from its database (they must have the Collector uninstalled for at least one day).

## Looking for unlicensed devices in the log

To extract the list of unlicensed devices from the Engine log, look for messages matching the format shown above.

If the Engine has been running for a long time, the information about unlicensed devices in the log may be outdated (you may have removed some Collectors and installed some others in the meantime or the log file may have rotated). To make sure that you get up-to-date information regarding unlicensed devices from the Engine log, restart the Engine previous to examining the log file:

1. Log in to the CLI of the appliance hosting the Engine.
2. Stop the Engine:

```
nxinfo launch -r
```
3. Make a backup of the log file:

```
sudo mv /var/log/nexthink/engine.log
/var/log/nexthink/engine.log.bk
```
4. Start the Engine:

```
nxinfo launch -a
```
5. Wait for the Collectors to send information.



After waiting for a reasonable amount of time (one full day, for instance), examine the Engine log as described below:

1. Type in the following command to retrieve the list of *out of license* messages:

```
sudo grep -E ".+machine.+out of license"
/var/log/nexthink/engine.log
```

2. Optional: Count the lines of the previous result to get the total number of devices in excess:

```
sudo grep -E ".+machine.+out of license"
/var/log/nexthink/engine.log \ | wc -l
```

## Related tasks

- Setting up a software license
- Establishing a privacy policy
- Removing devices

# Removing devices

## Manually removing devices

To manually remove a device from the Finder:

1. Log in to Finder with administrative rights.
2. Type the name of the device in the Search field.
3. Right-click the device in the results of the search and select **Drill-down**.
4. Right-click the device listed and select **Edit...** (or type **Ctrl+Alt+E**). The **Edit device** dialog shows up.
5. Select the option **remove** from the list **Storage** at the bottom of the dialog.
6. Click **Apply**. The device is marked for removal.

The Finder still displays the device until it is removed from the Engine database the next day, during the nightly cleanup. You must uninstall the Collector from the device to be removed. Failing to do so results in the Engine recreating the removed device as soon as the Collector sends data from it.

Once the device is completely removed from the Engine, the system increases the number of available device licenses by one unit.

Applies to platforms:

## Automatic removal of inactive devices

If a device is inactive for more than 90 days, the Engine purges its data and automatically frees one license from the pool.

Note that Mobile licenses are counted separately from Windows and Mac OS licenses.

### *Changing the maximum inactivity period of devices*

Modify the maximum inactivity period of devices for the Engine to identify a device as inactive either more quickly or more slowly and, accordingly, remove it from its database. Note that modifying the maximum inactivity period of devices is local to each Engine.

To modify the maximum inactivity period of devices:

1. Log in to the CLI of the Appliance that hosts the Engine.
2. Open the configuration file of the Engine for edition:  

```
sudo vi /var/nexthink/engine/01/etc/nxengine.xml
```
3. Inside the limit section, set the new inactivity period in seconds (default value is 7776000 seconds, that is, 90 days):  

```
<limit>
<max_inactivity_period>7776000</max_inactivity_period>
</limit>
```
4. Save your changes and exit by typing in:  

```
:wq
```
5. Restart the Engine:  

```
nxinfo launch --restart
```

Beware that setting the maximum inactivity period too low may result in an inefficient removal and recreation of devices with regular inactivity intervals, when these intervals are longer than the specified maximum period.

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.

## Examining the logs in the Portal

The log files of the Portal are located in the Appliance that hosts it under:

**`/var/nexthink/portal/log/`**

The names of the all the log files of the Portal are prefixed with the word **portal\_**.

## Log files

The names of the log files reflect the current running mode of the Portal. Note that, for medium and large modes, subsystems of the Portal write to different log files:

- **portal\_<running mode>.log**: Standard log file.
- **portal\_activity\_<running mode>.log**: Extract more high levels information such as switch in engine connection state, Size of data stored in database, Memory consumption of the JVM.
- **portal\_<running mode>.err** : Standard error stream with low-level error messages (for support and unexpected cases).
- **portal\_<running mode>.out** : Standard output stream with low-level information (for support and unexpected cases).

## Running Modes

Depending on the size of the Portal database, there are different running modes.

To know the current running mode, take a look at the file:

**`/var/nexthink/portal/conf/startup.properties`**

The name and the number of log files depend on the running mode, as listed below:

### ***Small mode***

- **SMALL**: Single node when running in single JVM mode.

### ***Medium mode***

- **MEDIUM\_UI**: Portal UI, Portal compute and HTTP server when running in dual JVM mode.
- **MEDIUM\_INFRA**: Content manager, login manager, communication layer, real-time layer when running in dual JVM mode.

## ***Large mode***

- **LARGE\_UI**: Portal UI, Portal compute and HTTP server when running in Three JVM mode.
- **LARGE\_INFRA**: Content manager, login manager, when running in Three JVM mode.
- **LARGE\_COMM**: Communication layer, real-time layer when running in Three JVM mode.

## Related tasks

- Allocating resources for the Portal

## **Storing Engine data in a secondary disk drive**

In some situations, you may want the Engine to store its data in a disk drive different from the system drive:

- Little space available in the system drive.
- Faster secondary drive.

The following procedure shows you the recommended way for storing the data of the Engine in a secondary disk drive:

1. Log in to the CLI of the Appliance that hosts the Engine.
2. Create a new partition in the secondary disk using **fdisk**. For this part of the procedure, we assume that your secondary disk is a second SCSI or SATA device in the Appliance named **/dev/sdb**. If this is not the case, you make have to adapt the commands below to suit your specific needs. Type the following commands to create the first primary partition in the secondary disk:

```
sudo fdisk /dev/sdb
n (for creating a new partition)
p (for creating a primary partition)
1 (create the first partition)
1 (default number for the first cylinder)
2610 (default number for the last cylinder)
w (write the partition info to the disk)
```

3. Format your newly created partition with the ext3 filesystem:

```
sudo mkfs -t ext3 /dev/sdb1
```
4. For the second part of the procedure, we assume that you want to move the first instance of the Engine to a secondary drive. If this is not the case,

substitute the 1 by the appropriate Engine number in the following commands. Stop the Engine:

```
nxinfo launch -r (use the option -g <engine_number> to stop a different instance of the Engine)
```

5. Rename the data folder of the Engine to keep its contents:

```
cd /var/nexthink/engine/
sudo mv 01/ 01-old/
```

6. Recreate the data folder of the Engine:

```
sudo mkdir 01/
```

7. Mount the folder on the recently created partition of the secondary disk:

```
sudo mount /dev/sdb1 /var/nexthink/engine/01
```

8. Edit the **/etc/fstab** file for the system to automatically mount the secondary drive while booting:

```
sudo vi /etc/fstab
```

9. Add the following line to the end of the file:

```
/dev/sdb1 /var/nexthink/engine/01 ext3 defaults 1 2
```

10. Save your changes and quit the text editor:

```
:wq
```

11. Copy the contents of the old data folder of the Engine to the new data folder:

```
sudo cp -r /var/nexthink/engine/01-old/*
/var/nexthink/engine/01
```

12. Restart the Engine

```
nxinfo launch -a (use the option -g <engine_number> to start a different instance of the Engine)
```

Now the Engine is using the secondary disk drive as storage medium.

You can use a similar method to store the logs of the Engine in a secondary disk drive. Just mount the directory **/var/log/nexthink** on a partition of the secondary disk in much the same way as explained above for **/var/nexthink/engine/01**.

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.

## MSI Exec Returns 3010

When successfully installing Nexthink Collector using its MSI package, the return code of the Windows Installer process (i.e. msiexec.exe) may be 3010 instead of 0.

Some deployment tools may consider that a non-zero return code corresponds to an installation failure, hence hindering a clean installation of the Collector.

Windows Installer does not always return 0 upon successful MSI installation. It can also return 3010 if a computer restart is required for completing the installation. As the installation of the Collector may require a computer restart in some cases, the Windows Installer process can return 3010.

If your deployment tool does not consider 3010 a valid return code, a possible workaround is to run the installation process in the context of a batch script that checks the return code of the installer. If it is 0 or 3010, the script must then itself return 0, making the deployment tool explicitly understand that everything went fine.

## Package Executable Mapping

Finding out which package an executable belongs to is not an trivial task and is not 100% accurate, an executable may even belong to no package. To do so, use the heuristic described below.

Let's define an executable as the tuple path, hash and name/size i.e. [PATH,HASH,FILE].

An MSI package contains both an installation and uninstallation scripts linked to embedded resources, usually binaries. Once installed, an MSI is stored on the machine but its resources are striped out to save disk space. However most embedded binaries are listed either by name or by size. In addition, an MSI defines an installation directory.

So for each MSI we have the tuple [{HASH},{FILE},DIR] even if some installed binaries may not be present neither {HASH} nor {FILE}.

Other type of packages are treated as black box and we take only the installation directory if present or by the path of its uninstallation program if not. so we have the tuple [{},{},DIR].

An executable [PATH,HASH,FILE] is associated to a package [{HASH},{FILE},{DIR}] whenever one of those conditions is met:

- HASH is contained in {HASH}
- DIR is equal to {DIR} \*

- DIR parent is equal to {DIR} \*
- FILE is contained in {FILE}

If no specific package can be associated to a executable, it is associated to the default "unknown" package.

The following directories are excluded:

- WINDOWS e.g. C:\WINDOWS
- SYSTEM e.g. C:\WINDOWS\system32
- PROGRAM\_FILES\_COMMON e.g. C:\Program Files\Common Files\Common Files
- PROGRAM\_FILES e.g. C:\Program Files\Common Files
- COMMON\_STARTMENU e.g. C:\Documents and Settings\LeeT\Start Menu
- COMMON\_PROGRAMS e.g. C:\Documents and Settings\LeeT\Start Menu\Programs
- COMMON\_STARTUP e.g. C:\Documents and Settings\gjaunin\Start Menu\Programs\Startup
- COMMON\_MUSIC e.g. C:\Documents and Settings\LeeT\Favorites
- COMMON\_FAVORITES e.g. C:\Documents and Settings\LeeT\Favorites
- COMMON\_DOCUMENTS e.g. C:\Documents and Settings\LeeT\My Documents
- COMMON\_DESKTOPDIRECTORY e.g. C:\Documents and Settings\LeeT\Desktop
- COMMON\_APPDATA e.g. C:\Documents and Settings\LeeT\Application Data

## Installing third-party software in the Appliance

The Appliance consists of a Linux-based operating system on which you can install the Portal or the Engine. The software packages included in the Appliance have been carefully selected and fine tuned to work together with both Nextthink products in order to deliver the best performance possible. Both the Portal and the Engine are very demanding in terms of computing resources and they usually require the full dedication of the hardware specified to run them.

Therefore, the installation of third-party software that competes for computing resources with the Nextthink products in the Appliance can degrade the overall performance of the Appliance or hinder the proper functioning of the Portal or the Engine.

As an exception, Nexthink recommends the installation of VMware Tools in those virtualized Appliances that run on VMware products.

## **Supported third-party software**

Nexthink only supports third-party software in any of the following two cases:

- The installation procedure of the software is described in the official Nexthink documentation.
- An engineer from Customer Success Services, the Presales team, or the MSP team in Nexthink undertakes the installation of the software.

Nexthink cannot provide support to customers or partners who do not comply with the statements above. To regain access to Nexthink Support, you must remove all non-conforming third-party software from the Appliance.

## **Installing typical third-party tools**

Usually, you may want to install third-party software in the Appliance to perform any of the following tasks:

- Backup the Appliance
- Monitor the Appliance
- Protect the Appliance against computer viruses

The tools that typically perform these tasks may have a major impact in the performance of the system; therefore, Nexthink recommends not to install any additional tool. Should you choose to go ahead and install third-party software (because it is mandated by the security policy of your company, for example), we strongly recommend that you first test your setup in a pre-production environment.

### ***Backup the Appliance***

Starting from Nexthink V4.1, the Appliance includes an automatic backup mechanism that lets you push all the database and configuration files to a shared directory. Configure the automatic backup of the Appliance from the Web Console to recover from a full or partial data loss.

If you are compelled to install a third-party backup tool, schedule it to perform the backup when the Appliance is less active and always test it first in a pre-production environment. Depending on the product that you installed in the Appliance, follow the corresponding piece of advice:



## Engine

The Engine is less active during the night, when it receives less data from Collectors and it has finished the cleanup of its database. Schedule the backup at around 04h30.

## Portal

The Portal is less active when fewer users are connected to it and it is not collecting data from the Engine. Since data collection starts at 01h00 and it can last for several hours, schedule the backup of the Portal between the end of the working hours and 01h00.

## ***Monitor the Appliance***

Currently, Nexthink does not provide any specific tool to monitor the status of the Appliance. Advanced users may take advantage however of the standard tools of Linux installed in the Appliance, such as *ping* or *SSH*, to test the connectivity of the Appliance, or use the command line shell to enquire about the status of the Engine or the Portal.

If you really need to install a third-party monitoring tool in the Appliance, be specially careful if it is the Engine that is running on the Appliance. A monitoring tool can greatly interfere with the Engine during periods of high activity.

## ***Protect the Appliance against computer viruses***

The Appliance is always delivered with the latest security updates of the underlying Linux-based operating system. The risk of vulnerabilities is thus reduced to a minimum. Still, if you have any particular requirements in terms of protection of the operating system, create a feature request for Nexthink Support or contact your Nexthink Account Manager to initiate a discussion.

## Related references

- Nexthink Appliance (hardware requirements)
- Planning for disaster recovery (backup)

## **Installing VMware Tools in the Appliance**

Nexthink recommends installing VMware Tools in any Appliance that runs on top of VMware virtualization products such as vSphere. VMware Tools significantly improves the performance and manageability of virtualized Appliances.

Starting from Nextthink V6, the Appliance is distributed with the **open-vm-tools** package already pre-installed. Therefore, no action is required on your part. When you deploy the Appliance in a VMware environment, it directly benefits from the features provided by the package. In addition, the package is automatically updated via the Appliance updates whenever a new version is available.

If for some reason you need to install the commercial version of VMware Tools, uninstall the open-vm-tools package first and then proceed as follows. Note however that VMware recommends the use of open-vm-tools on those platforms where the package is available, so **do not install the commercial version of VMware Tools** unless you really know what you are doing.

To install the commercial version of VMware Tools in the Appliance:

1. Open the vSphere Web Client and log in to connect to your vCenter Server.
2. On the left-side pane, click **vCenter** and select **Virtual Machines** from the **Inventory Lists** section.
3. Click the name of the virtual machine that runs the Appliance.
4. In the **Summary** tab, a yellow warning box displays the message **VMware Tools is not installed on this virtual machine**.
5. Click the link to the right of the warning message that reads **Install VMware Tools**.
6. Click **Mount** in the pop up dialog. A virtual CD with VMware Tools is now attached to your VM.
7. Open a terminal connection to the Nextthink Appliance (e.g. click **Launch Console** or connect to it via ssh) and log in to its CLI.
8. Type the following commands to mount the virtual CD:

```
sudo mkdir /mnt/cdrom
sudo mount -t iso9660 /dev/cdrom /mnt/cdrom
```

9. Check whether the mount was successful by listing the contents of the cdrom folder. The file **VMwareTools-<version>.tar.gz** must appear in the list:

```
ls /mnt/cdrom/
```

Copy the VMware Tools file to the tmp folder and extract its contents:

```
cp /mnt/cdrom/VMwareTools-*.tar.gz /tmp/
cd /tmp
tar -xvzf VMwareTools-*.tar.gz
cd vmware-tools-distrib
```

10. Install the VMware tools executing the following script:

```
sudo ./vmware-install.pl
```

11. Press **Enter** to accept the default option whenever asked during the installation process.
12. Reboot the Appliance after install:

```
sudo reboot
```

After installing VMware Tools, you should be able to see the IP addresses of the VM hosting the Appliance in the **Summary** tab. The warning message about the installation of VMware Tools disappears.

#### Related references

- [Open-VM-Tools project on GitHub](#)