Nexthink V6.18

Product Overview

Generated: 12/20/2020 2:01 pm

Table of Contents

Table of Contents

Table of Contents

Nexthink End-User IT Analytics

Software components

Nexthink is the innovator of End-user IT Analytics for security, ITSM, and workplace transformation. Nexthink maps all the IT services, how they are being consumed, and how the IT infrastructure is operating, from the only perspective that matters most, the end-users (workers). Nexthink provides essential visibility and insight into IT operations and security for IT Governance.

Nexthink Architecture

The architecture of Nexthink has been designed to simplify operations, ensure scaling and allow a rapid deployment. The system is composed of six main software components:

- The Collector captures information from all end-user desktops and laptops.
- The Mobile Bridge captures mobile device information from Microsoft Exchange.
- The Engine aggregates Collector and Mobile Bridge information and provides real-time IT analytics.
- The Finder is the rich client application for searching and analyzing data on Engines.
- The Portal aggregates Engine information and provides dashboarding, reporting, and long-term trending analytics.
- The Library is a cloud knowledge database.

Modular product structure

Nexthink offers a modular product structure that can grow with your needs. The product is licensed with respect to the number of monitored physical or virtual devices and, optionally, server users. On top of the basic product (Nexthink Analytics), the following modules can be purchased:

- **Nexthink Act** offers you a way to remotely act on the devices of the end-users for automated or assisted servicing.
- **Nexthink Engage** gives you the means to reach out to the end-users, gather their feedback regarding IT or other subjects, and notify them of relevant issues.
- **Nexthink Enhance** provides additional classification and security-related information to Analytics, including binary threat level and category, as well as web domain reputation, category and hosting country.
- **Nexthink Web & Cloud** grants access to analytics related to intranet and extranet HTTP and HTTPs web requests (now included in the Nexthink Analytics offer for new contracts).
- **Nexthink Integrate** enables the product API and access to continuously improved integration samples, reports, etc.

Nexthink Analytics as well as the modules grant access to investigations, widgets, dashboards, categories, etc. directly from the Nexthink Library, our cloud repository of content.

Collector

Introduction

The Collector is a light-weight agent based on patented technology. It captures and reports network connections, program executions, web requests, and many other activities and properties from the devices of the end-users on which it runs. It is implemented as a kernel driver and an accompanying service, offering remote and automated silent installations and negligible impact on the performance of local desktops, while minimizing network traffic.

CPU usage	Memory usage	Network traffic
• Less than 0.015% (in average)	• Kernel: Around 500 KB • User: Around 20 MB	• **UDP** 0.1 - 0.3 Kbps (in average) • **TCP** Depending on · Campaigns · Remote actions · Updates

The following figure depicts the role of the Collector within the Nexthink solution.

Collector components

The capabililty of the Collector for gathering user activity data is shared by the kernel driver and the helper service (or daemon) components. By running close to the operating system, the kernel driver detects some kinds of user activities that are only visible at this level.

Click to see the detailed list of components of the Collector.

Features

Multi-Platform

The Collector is available for both Windows and macOS operating systems. The present documentation states the platforms to which each feature applies. Likewise, the data model details the individual pieces of information collected for each platform.

Applies to platforms:

CrashGuard

Since the Windows Collector driver is a kernel-mode component, any error in its internals or its interaction with a misbehaving third-party driver can lead to system instabilities. Even with Nexthink putting as much attention as possible towards delivering bug-free software, the principle of precaution holds. The CrashGuard feature detects every system crash and, by default, it disables the Collector driver itself if the system crashes more than three times in a row after installation.

Applies to platforms:

Kernel traffic interception

Some applications may send and receive data to and from the network using kernel-mode components, actually hiding their network traffic from user-space monitoring applications. Being a kernel driver itself, the Windows Collector is nevertheless able to detect and report such traffic.

Applies to platforms:

Paths aliasing

The Collector identifies commonly used paths (e.g. C:\WINDOWS\, C:\Program Files\) and other special mount locations (removable mount points, network drives) with paths aliases. For example, if the DVD-Rom drive is mounted under D:, the Collector reports an application **setup.exe** being launched from this media as **%RemovableDrive%\setup.exe**.

Network switching

A change of network interface is transparent to the Collector, except when it invalidates the DNS resolution of the Engine. In the latter case, the process of adapting to a different network may take a few minutes and the Collector resends the whole context to the Engine.

Event logging

Main events and errors are written to either the standard Windows event logs or Mac OS logs.

On-the-fly configuration

The Collector driver parameters can be changed through the Collector Control Panel extension or the Collector Configuration tool. There is no need to restart the computer for the changes to become effective.

Related tasks

- Installing the Collector on Windows
- Installing the Collector on Mac OS

Related references

- Collector configuration tool
- Components of the Collector
- Data model

Mobile Bridge

The Mobile Bridge is a server software component that gathers information about the mobile devices which connect to your Microsoft Exchange mail servers through the ActiveSync protocol. The Mobile Bridge sends all the gathered data back to the Engine, where it is organized and stored along with the information sent by the Collectors.

Thanks to the Mobile Bridge, you can keep an eye on the access status and last synchronization time of all the mobile devices in your corporate network and establish links between your mobile users and desktop users. Nexthink offers you this information and much more from a single place in a uniform way, helping you keep your BYOD infrastructure under control. Query Nexthink about mobile devices and users by applying the same mechanisms that you would use for querying about desktop devices and users.

Related tasks

- Installing the Mobile Bridge

Finder

Nexthink Finder, built upon powerful visualization techniques, is the search and user interface to render visibility into your IT infrastructure. Analyze IT services and query what you need within seconds. Expand or drill-down the results in a few clicks to reveal swiftly, across the entire network, how many versions of a particular application are in use and on which workstations, the bandwidth consumed by the application, the servers and domains that the application accesses, the network response times, which users experienced issues, and much more.

Engine

Nexthink Engine is a high-performance analytics software capable of processing millions of endpoint activities in seconds. Events sent in real time by Collectors populate the Engine with activity data, furnishing a rich repository of historical and live IT infrastructure usage data from the end-user perspective. Engine leverages an in-memory database for rapid queries (via the Nexthink Finder) and flexible reporting (via the Nexthink Portal).

Related tasks

- Installing the Appliance

Portal

Nexthink Portal is the reporting tool, collaboration platform and centralized management platform of the Nexthink End-User IT analytics platform. A comprehensive set of dashboards are delivered out-of-the-box but it is possible in a matter of minutes to construct custom dashboards, valuable for anyone in the organization. Personalized metrics are simple to define as drag-and-drop widgets and can be quickly published and shared. Nexthink Portal front-end is a web application running inside a browser.

Related tasks

- Installing the Appliance

Nexthink Library

The Nexthink Library is an online knowledge database that gives you access to theme-based files, a large set of ready-to-use predefined investigations, templates, dashboards and application information accessible directly from the Finder and the Portal.

A separate component of the Nexthink Library is the Application Library. The Application Library helps you identify potential threats by submitting the digital footprint of any application found on a desktop or visited web domain to its reference databases. Thanks to the full integration between the Application Library and the Engine, your infrastructure information is always fully up-to-date, without the need for any manual interaction.

Related references

- Nexthink Library
- Nexthink Application Library

Licensing terms

License agreement

The links provided on this page hold the terms and conditions that govern the use of Nexthink software by customers who purchased a commercial license of Nexthink.

- If you purchased your license directly from Nexthink:
  License agreement V6.17 for direct customers
- If you purchased your license from an official reseller:
  License agreement V6.17 for indirect customers

Once the product is installed, find a copy of the licensing terms in any of the provisioned Nexthink Appliances under:

```
/var/nexthink/eula/license.txt
```

Open source software licenses

Nexthink software components make use of third-party software libraries that follow an open source licensing model. These libraries are redistributed in binary

form within selected Nexthink components.

Nexthink is grateful to the authors and contributors of all the high quality open source projects that make possible the development of our own product.

Find the full list of open source software libraries used by Nexthink, along with their corresponding licenses, in every deployed Nexthink Appliance under:

```
/var/nexthink/eula/Libraries_licenses.txt
```

Related references

- License agreement V6.17 for direct customers
- License agreement V6.17 for indirect customers

What's new in V6.18

New features

Improved Portal computations

The nightly computations of metrics in the Portal are particularly demanding in terms of hardware resources. Especially in big environments, nightly computations can take a very long time to finish.

Technical improvements and optimizations in V6.18 have boosted the performance of the Portal to significantly reduce the time and the resources needed to complete the nightly computations.

Find out more

Improved performance and increased entities in the Engine

New performance improvements in the Engine are largely responsible for reducing the time of nightly computations in the Portal. The Engine sensibly speeds up all types of Finder queries as well.

As a result of these improvements, each Engine supports now more entities to build your hierarchies. Up to 500 entities per Engine let you group your devices into more specific subsets. If your setup requires more than a total of 2 000 entities across all Engines, please contact Customer Success Services.

Find out more

Hardening for upgraded Appliances

Starting from V6.17, newly installed Nexthink Appliances are hardened by default. Hardening is however not automatically applied to existing installations during the upgrade.

In V6.18, find an option in the Web Console that lets you apply the same hardening settings of new installations to upgraded Appliances:

Find out more

Additional languages for Nexthink Engage

Campaigns are always more effective when you address the audience in their mother tongue. In addition to the custom translations that the creators of a campaign can provide, the user interface of notifications is now available in more languages than ever.

Find out more

Enabling Engage, Act and automatic updates on roaming devices

For roaming devices that are not connected to the corporate network through a VPN, there is a new documented procedure to let the Collector communicate with the Engine via the TCP channel by installing and configuring a reverse proxy on the Nexthink Appliance.

The TCP connection of the Collector enables Nexthink Engage, Nexthink Act and automatic updates on the devices where the Collector is installed.

Find out more

Deprecated features

Object identifiers

The ID of an object is a numeric attribute that distinguishes that object from any other object of the same type within a single Engine. To get ready for investigations across all Engines, the Finder neither accepts nor displays this single-Engine IDs. Instead, UIDs are now used to uniquely identify objects across all Engines.

The removal of single-Engine IDs for objects has the following consequences:

- Existing investigations using IDs of objects become invalid.
- The Finder uses UIDs when drilling-down or when using one-click investigations.
- You must provide UIDs when previewing a one-click investigation.
- Because UIDs are longer than IDs, the Investigations editor in the Finder can receive as input a single UID. If more than one UID is generated for an input box (as a result of a drill-down or one-click), the input is read only.

Contrary to objects, activities and events are still identified by single-Engine IDs, as an activity or event is bound to one Engine only.

V5 Widgets

After being deprecated in V6.17, V5 widgets have been completely removed from the product in V6.18.

Please ensure that you have replaced all your V5 widgets by V6 metrics and corresponding widgets before upgrading. V5 widgets are neither manageable nor visible in V6.18.

Data-model changes

Object identifiers

To get ready for investigations across all Engines, single-Engine IDs for all object types have been removed.

Click here for more information.

What's new in V6.17

New features

Compliance

GDPR

To ensure that your Nexthink setup complies with the *General Data Protection Regulation* (GDPR) that applies to personal data of all EU residents, Nexthink offers you additional anonymization and data retrieval options.

- Run the GDPR script in the Nexthink Appliance to:
  - Retrieve all the personal data relative to a particular user or device from either the Portal or the Engine.
  - Anonymize the personal data relative to a particular user or device in the Portal.
- Anonymized traffic Redirection
  - A couple of additional fields in the printing information can be anonymized with the help of the redirection service of Collector traffic to protect users' personal data.

Find out more

Appliance hardening

To reduce the exposure of the Appliance to external attacks and comply with the security policies of our customers, new installations of the Appliance are automatically secured out-of-the-box, whithout the need to apply the Security Hardening procedure.

Because of Appliance hardening, the real-time communication between the Portal and the Engine is established only if you federate your Appliances. Federation is thus *mandatory* from V6.17 on.

Find out more

Technology

Centralization of content

To provide users a unified experience with the Finder across all Engines in multi-Engine setups, virtually all the content that users create in the Finder is now centralized in the Portal and replicated in all Engines automatically.

Nexthink users find now their own set of investigations, one-clicks, and alerts at any time, regardless of the particular Engine to which they connect with the Finder.

Find out more

Rebootless Mac Collector

The new version of the Collector for macOS devices no longer needs to be rebooted after installation or upgrade. In addition, no explicit consent of the user is needed to install the last version of the Collector for macOS, because it completely works in user mode.

On the whole, these improvements simplify and speed-up the deployment and installation of the Collector on macOS devices.

Find out more

Improvements in AD authentication

To identify and authenticate Nexthink users with their standard Windows accounts, Nexthink provides an integration with Active Directory authentication services. Customers can additionally provision Nexthink users from groups of users in Active Directory and let users log in without the need to type in their password using Windows authentication.

Starting from V6.17, users whose *User Principal Name* (the currently preferred user logon name in Windows) does not match their *sAMAccountName* (the user logon name in pre-Windows 2000, but still in use today) in Active Directory will not have any trouble to log in to either the Finder or the Portal.

Find out more

Deprecated features

Web API V1 management

Web API V1 has been deprecated since Nexthink V6.9 in favour of the more flexible Web API V2 and its accompanying NXQL language.

Starting from 6.17, the tools to manage Web API V1 investigations are removed from the product. Therefore, Nexthink recommends you to convert all your Web API V1 investigations to their Web API V2 equivalents before upgrading to 6.17.

Nevertheless, to protect existing integrations, published Web API V1 investigations are not removed from the system when upgrading to 6.17 and they will still work after the upgrade. They will just be no longer accessible from the Finder.

Web API V1 functionality will be completely removed in a future version of Nexthink. Any integration relying on published investigations of the Web API V1 will stop working by then.

Until V6.16	After V6.17

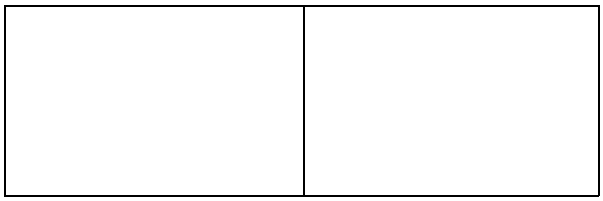

V5 widgets

Although the first version of Nexthink V6 appeared almost three years ago, V5 widgets were still supported for customers with V5 content in their Portal.

Starting from V6.17, V5 widgets are officially deprecated and will be removed from future versions of the product.

What's new in V6.16

New features

Quality of Service

Up to ten score tabs in the Device and User views

To increase the productivity of Finder users, the Device and User views double the quantity of available tabs for displaying scores.

Display up to ten score tabs in each one of the views to provide Finder users with checklists based on predigested key data, optionally including links to remote actions for assisted service scenarios.

Find out more

Detecting execution errors in Mac devices

To improve the visibility of issues in Mac devices and better measure the digital experience of Mac users, Nexthink now reports the applications with stability problems that crash while runnning on top of macOS.

Find out more

Allowing users to publish dashboards

A new profile option lets Nexthink normal users publish their own Portal modules. A task that, up to now, was reserved to administrators.

Democratizing the publication of dashboards enables users to participate more actively in the creation of useful content for the Portal, while still ensuring that administrators have control over the content published by other users.

Find out more

Compliance

Identifying macOS binaries with SHA-256

To securely identify binaries and enable an easier integration with modern security products, Nexthink applies now the cryptographic hash function SHA-256 to Mac binaries in order to get their digital fingerprint.

Previous hashing algorithms such as MD5 and SHA-1, although still available, are no longer considered secure enough against collision attacks.

Find out more

Technology

Replacing product certificates from the Web Console

To simplify the management of digital certificates that secure many of the communication channels within the product, the Web Console now provides a user interface to easily upload custom server certificates and replace the default certificates in both the Engine and the Portal appliances.

The Web Console interface removes the need for complex manual procedures through the command-line interface for replacing digital certificates in Nexthink Appliances, at least for the most common cases.

Find out more

Support for Windows 10 version 1803

Nexthink supports the latest update of Windows 10, namely version 1803. The Windows Collector has been extensively tested in this new version of the popular operating system and can be confidently installed on any device that runs Windows 10 version 1803.

Deprecated features

Simplification of account types

In order to simplify accounts and the management of rights and access, the concept of local administrators has been removed. Instead, it is now possible to grant users the ability to publish dashboards.

Data-model changes

New data for macOS

Application errors

Application crashes and all associated information are now available for macOS.

Additionally, the following aggregates are now available for macOS:

- Number of application crashes
- Application crash ratio

Field	Group	Type			
Number of application crashes	Errors	Aggregate			
	Number of application crashes				
	NXQL ID:	number_of_application_crashes			
Application crash ratio	Errors	Aggregate			
	Indicates the number of application crashes per 100 executions.				
	NXQL ID:	application_crash_ratio			

More information

SHA-256 hashes for binaries

The following field for binaries is now available on macOS:

- SHA-256 hash

Field	Group	Type			
SHA-256 hash	Properties	Field			
	Indicates the SHA-256 hash of the binary.				
	NXQL ID:	sha256			

More information

Last known connection status

The following fields are now available for macOS:

- Last known connection status
- Last seen on TCP
- Number of days since last seen on TCP

What's new in V6.15

New features

Quality of service

Run self-help scenarios as local system

In self-help scenarios, Nexthink Act remote actions launch Nexthink Engage campaigns to interact with the end users and guide them toward the resolution of an issue. Because these remote actions required the intervention of end users, they were always run in the context of the current interactive user. However, some actions require administrative privileges to carry out certain operations, while the current interactive user may not necessarily be an administrator.

To solve this problem of privileges, it is now possible to run self-help remote actions in behalf of the local system account, which always has administrative privileges.

Find out more

Application CPU and memory usage in macOS

To help you spot issues with the performance of applications in your macOS devices, Nexthink now reports the CPU and memory usage of applications.

Additionally, in case of high CPU or high memory consumption, warnings show up in the device and user views:

Find out more

Transformation

Windows 10 and Windows Server version and build number

Starting from Windows 10, Microsoft is delivering its popular Windows operating system as a service, rather than as a product. The new servicing model abandons the traditional version numbers and service packs for a new way of releasing new features without changing the major version number of the operating system. This new servicing model also implies however that the devices in your organization may be running different, but difficult to tell apart, flavors of Windows 10.

For its part, Windows Server is now offering a release model that aligns with that of Windows 10. You can choose between the more traditional Long-Term Servicing Channel (formerly called Long-Term Servicing Branch), whose current product is *Windows Server 2016* and has a release cycle of 2-3 years, and a Semi-Annual Servicing channel, which starts with *Windows Server, version 1709* and gets updated twice a year with new feature releases.

To help you distinguish the actual flavors of Windows that your devices are running, Nexthink now reports the new version number of all Windows devices in the enhanced **OS version and architecture** field as well as the build number in the new **OS build** field.

Find out more

Compliance

Identify binaries with SHA-256 hash

Nexthink identifies binaries by their hash; that is, by a fixed-length sequence of bytes that results from applying a cryptographic one-way function to the image of a binary file. The hash acts thus as a digital fingerprint of the binary file.

The Collector computes the MD5 and SHA-1 hashes of binary files and reports them to the Engine. Because these algorithms are no longer considered secure enough against collision attacks, the Collector now computes and reports the SHA-256 hash of binaries as well. As a side effect, reporting SHA-256 hashes makes the integration of Nexthink with third-party security tools easier, as most of these tools natively support SHA-256 nowadays.

Find out more

Technology

Viewing appliance certificates from the Web Console

The communication of client applications, such as the Finder or the web browser, with the server components of Nexthink, such as the Web Console, the Portal, or the Engine, as well as the communication among server components themselves is secured with digital certificates.

To help you keep the secure communication channels under control, the Web Console now displays the digital certificates that are in place in your appliances.

Find out more

Skipping Portal backup before upgrade

Because backing up the database of the Portal can be a lengthy process, it is now possible to avoid the automatic backup of the Portal before each upgrade when you consider that you already have a recent enough backup.

Find out more

Other changes

Matching lists of strings

When editing the conditions of an investigation, it is possible to look for matches of a given pattern with fields of type **String**, such as the name of devices.

Now it is also possible to set conditions that do pattern matching on fields which are lists of strings (fields of type **StringList**), such as the monitor models of a device.

Deprecated features

Windows 8 support

Windows 8 reached its end of life and Microsoft is releasing neither new patches nor security updates for Windows 8, pushing customers to upgrade to Windows 8.1 or Windows 10.

Therefore, Nexthink no longer supports Windows 8. On the other hand, Windows 8.1 is still fully supported.

Data-model changes

New data for macOS

The following new aggregates are available for MacOS:

- Total CPU time
- CPU usage ratio
- High application thread CPU time ratio
- Average memory usage per execution

Field	Group	Type			
Total CPU time	Activity	Aggregate			
	Indicates the sum of the CPU time of all executions on each device in scope and over all logical processors. Executions shorter than 30 seconds are ignored. • Example: if we consider two executions with the first one taking 50% of a logical processor during 30 minutes and the second one taking 100% of 2 logical processors during 60 minutes, the total CPU time is 135 minutes (= 50% * 30 min + 2 * 100% * 60 min).				
	NXQL ID:	total_cpu_time			
CPU usage ratio	Activity	Aggregate			
	Indicates the sum of the CPU time of all executions on each device in scope over all logical processors divided by their total duration. Executions shorter than 30 seconds are ignored. • Example: if we consider two executions with the first one taking 50% of a logical processor during 30 minutes and the second one taking 100% of 2 logical processors during 60 minutes, the CPU usage ratio is 150% (= [50% * 30 min + 2 * 100% * 60 min] / [30 min + 60 min]).				
	NXQL ID:	cpu_usage_ratio			
High application thread CPU time ratio	Warnings	Aggregate			
	Indicates the ratio between the time that the underlying executions are in high thread CPU usage and their execution				

	duration.				
	NXQL ID:	high_application_thread_cpu_time_ratio			
Average memory usage per execution	Activity	Aggregate			
	Indicates the average memory usage of all underlying executions before aggregation. The value is the average memory usage of all executions (calculated with a 5-minute resolution) multiplied by their cardinalities and divided by the total cardinality. • Example: if two tabs of the Chrome browser are opened at the same time, two distinct processes of chrome.exe are launched and they are aggregated by the Engine (i.e., event cardinality = 2). The average memory usage will be the average of the two processes before aggregation: it represents the average memory usage of a Chrome tab.				
	NXQL ID:	average_memory_usage_per_execution			

Additionally, the following fields are now available for each execution:

- Total CPU time
- Average memory usage

Field	Group	Type			
Total CPU time	Properties	Field			
	Indicates the sum of the CPU time of all executions (before aggregation by the Engine) over all logical processors. Executions shorter than 30 seconds are ignored. • Example: if we consider two executions that are launched at the same time (hence aggregated by the Engine), with the first one taking 50% of a logical processor during 30 minutes and the second one taking 100% of 2 logical				

	processors during 60 minutes, the total CPU time is 135 minutes (= 50% * 30 min + 2 * 100% * 60 min).					
	NXQL ID:	total_cpu_time				
Average memory usage	Activity	Field				
	Indicates the average memory usage of the underlying executions before aggregation with a sampling resolution of 5 minutes. • Example: if two tabs of the Chrome browser are opened at the same time, two distinct processes of chrome.exe are launched and they are aggregated by the Engine (i.e., event cardinality = 2). The average memory usage will be the average of the two processes before aggregation: it represents the average memory usage of a single Chrome tab.					
	NXQL ID:	average_memory_usage				

OS build

The operating system build is now reported for Windows devices.

Field	Group	Type			
OS build	Operating system	Field			
	Indicates the build number of the operating system: • '0.0.0.0': incompatible collector version or the data is not yet available				
	NXQL ID:	os_build			

Enhanced OS version and architecture

Due to a change in Microsoft's servicing model, the field "OS version and architecture" now includes the feature update release version of Windows 10.

As an example:

- "Windows 10 Pro (64 bits)" might become "Windows 10 Pro 1709 (64 bits)"

This would have the following impact on existing investigations and metrics.

The following conditions would still work:

- "Os version and architecture" matches "Windows 10*"
- "Os version and architecture" matches "Windows 10 Pro*"
- "Os version and architecture" matches "Windows 10* (64 bits)"
- "Os version and architecture" matches "*Pro*"

The following condition would become broken and should be updated:

- "Os version and architecture" is "Windows 10 Pro (64 bits)"

Additional reading:

- Windows 10 as a service (Microsoft)
- Windows 10 release information (Microsoft)
- Windows 10 version history (Wikipedia)

SHA-256 hashes for binaries

SHA-256 hashes are now available for binaries.

Field	Group	Type			
SHA-256 hash	Properties	Field			
	Indicates the SHA-256 hash of the binary.				
	NXQL ID:	sha256			

What's new in V6.14

New features

Device warnings for macOS

To help you spot user experience issues, the new version of the Mac Collector (which becomes 6.14, abandoning the 5.X numbering) now detects high memory and high CPU conditions on your macOS devices.

Find out more

Windows Server 2016 and version 1709 support

The Windows Collector now officially supports the latest versions of Windows Server:

- Windows Server 2016 (Long-Term Servicing Channel version)
- Windows Server, version 1709 (Semi-Annual Channel version)

Windows Server 2016 / 1709	Editions	Installation options
	• Datacenter • Standard • Essentials	• Desktop experience • Core server

Find out more

Manual triggering of remote actions restricted to one device by default

To prevent from misuse or from accidental massive triggering of remote actions that require special caution, such as those remote actions that reboot the device

or that may impact your infrastructure with a high network load, the manual triggering of remote actions is now restricted to one device at a time by default.

Find out more

White paper on scripting best practices for Nexthink Act

To help you develop your own PowerShell scripts for Nexthink Act, download from our Community site a white paper that describes the best practices for writing the scripts of remote actions.

Find out more

Fine-grained anonymization when redirecting Collector traffic

When redirecting Collector traffic to other Engines, anonymization enables the recording of Collector data while avoiding its association to a particular user or device. Now it is possible to exclude some fields from being anonymized, giving you more control over the data that is recorded.

Find out more

Security improvements in the Portal

To be compliant with the latest security standards, the security in the Portal has been improved in several areas; mainly, thanks to the addition of a reverse-proxy component that handles the connections to the Portal.

This has an impact in several procedures of the Portal:

- Drilling-down to the Finder.
- Portal backup and restore.
- Replacing the server certificate in the Portal.
- Specifying the secure protocols and ciphers in the Portal.
- New audit trail messages

Data-model changes

Device warnings for MacOS

Device warning (high CPU and memory) are now available for MacOS devices.

Field	Group	Type			
High device memory time ratio	Warnings	Aggregate			
	Indicates the ratio between the time the device is in high memory usage and its uptime.				
	NXQL ID:	high_device_memory_time_ratio			
High device overall CPU time ratio	Warnings	Aggregate			
	Indicates the ratio between the time the device is in high overall CPU usage and its uptime.				
	NXQL ID:	high_device_overall_cpu_time_ratio			

What's new in V6.13

New features

Free text answers in campaigns

To give end-users the chance to express their opinions in their own words, Nexthink Engage campaigns now allow the addition of free text to complement the responses to multiple choice and opinion scale questions.

End-users can now briefly justify a particular choice or give valuable insight into their opinions in a way that is just not possible with predefined responses only.

Find out more

Updated capabilities of Nexthink Act

Define up to 60 automatically triggered remote actions with more relaxed evaluation frequencies. Update the script execution policy with the Collector configuration tool (Nxtcfg) and measure the resource consumption of your scripts.

Find out more

Support for macOS High Sierra

Install the Collector on devices running the latest version of the operating system from Apple: macOS High Sierra.

Keep in mind that loading system extensions, such as some components of the Collector, requires explicit user approval in macOS High Sierra.

Find out more

Windows authentication and user provisioning in multi domain AD environments

Nexthink improves its interaction with Active Directory to provide support for implementations with multiple domains, such as those typically found in medium-sized to large organizations.

For both AD login of provisioned users and Windows authentication, Nexthink is able to recognize alternate domain names in user accounts (alternate UPN suffixes) through configuration.

For example, Nexthink can recognize the following two UPNs as different spellings of the same user:

- `john.smith@fr.acme.com`
- `john.smith@acme.com`

Find out more

User provisioning from multi-domain AD

Provision users to Nexthink from AD groups that belong to different domains in a tree.

Find out more

Windows authentication for multi-domain AD

Access both the Portal and the Finder with just a single click thanks to your Windows credentials, no matter what domain you belong to as a user.

Find out more

What's new in V6.12

New features

Acting on the devices of the end-users

The new Nexthink Act module lets you, for the first time with Nexthink, take action on the devices of the end-users.

Since its birth, Nexthink has provided you with invaluable information about your IT infrastructure from the perspective of the end-user. More recently, Nexthink has given you the means to involve the end-users themselves in the improvement of their overall experience with IT.

Now the Nexthink Act module completes the offer and opens up a whole new set of possibilities for improving your IT department in the areas of assisted support, automatic or guided remediation, and the gathering of on-demand data.

Find out more

New types of campaigns

The **Nexthink Engage** module brings in new types of campaigns to retrieve the opinion of the end-users either continuously or in repeating scenarios, without the need to retire and republish campaigns.

In addition, traditional campaigns (now called *one-off* campaigns) slightly change as well. Namely, users who declined to participate in a campaign or who fully answered the questions of a campaign are never notified again by the same campaign, instead of being notified again after 30 days if the campaign is republished.

For a better control of user rights, editing campaigns is now a distinct option from system configuration in the profile of Nexthink users.

Note that a new improved communication mechanism introduced in Nexthink 6.12 requires you to update all your devices to Collector 6.12 for campaigns to work. Devices with Collector 6.11 or previous are not able to display campaigns published with Nexthink 6.12, even when these campaigns are migrated from previous versions.

Find out more

Deploying Nexthink in the cloud

As more and more companies are moving part of their IT infrastructure to the cloud, there is a growing interest in deploying the server components of Nexthink on the most popular cloud platforms.

By providing a virtual hard disk image, installation packages and scripts, and

documentation, now Nexthink officially supports the installation of the Appliance components in Microsoft Azure.

For deploying Nexthink in Amazon, please contact our Customer Success Services.

Find out more

Renaming of modules

With the arrival of the new **Nexthink Act** module, other modules have had their name adapted to better reflect their purpose in a clear and concise way.

Find out more

Data-model changes

Device UID accessible via NXQL

The already existing unique identifier of devices is now accessible as the `device_uid` field of device objects in NXQL. This field is specially useful for integrations to trigger remote actions on specific devices.

Field	Group	Type			
UID	Properties	Field			
	NXQL ID:	device_uid			

	Indicates the universally unique identifier (based on Engine name and device ID).

What's new in V6.11

New features

Contact links in Campaigns

When encountering difficulties with a campaign, end-users may require a contact person to provide them support. In addition to web links, insert email and instant messaging links into campaigns so that end-users can easily contact the sender of the campaign or any other person who is able to help them in case of doubt.

Find out more

Branding the Portal

Customize the appearance of the Portal to harmonize it with the rest of your corporate IT environment. Brand the Portal with the logo of your company, which now shows up in the login dialog as well as in the top left corner of a Portal session, and replace the default background image in the login page. Display also your own logo in the email digests sent by the Portal.

Find out more

Finder usability improvements

Easily select multiple display fields for investigations, alerts, and metrics in the

Finder thanks to a new set of keyboard shortcuts that let you search, navigate, and pick the available fields faster and more conveniently.

As for categories, a new tooltip displays their description and all their possible keyword values when hovering the mouse over a category column in a list view of the Finder. The tooltip comes in handy when you are watching the list of results of an investigation and you quickly want to know about the meaning of a category column and all the tags which are applicable.

Find out more

Retrieving user logons and user interactions in virtualized and embedded environments

Because some virtualized and embedded environments start user sessions in a special way, the default mechanism of the Collector to detect user logons and user interactions does not work properly in these environments. When installing the Collector in a Citrix XenApp server, or in a Windows device that runs in kiosk mode, make sure to turn on the new *custom shells* feature, so that the Collector can report both user logon events and user interactions correctly.

Find out more

Other changes

Tracking Collector status

Tracking the status of the Collector is essential to know whether your devices are correctly connecting to your Engines. In addition to the status of the UDP connection between the Collector and the Engine, which lets you know whether a device is reporting activity data or not, you can now monitor the status of the TCP connection as well, thanks to a set of new device fields. Features such as the

automatic updates of the Collector and End-user Feedback rely on the TCP connection of the Collector for working properly.

When suspecting problems with the TCP connection of a particular Collector, run the Collector Configuration Tool for diagnosis.

Find out more

Sticky Customer Key

To simplify the deployment of the Collector, the Web Console now consistently returns the same Customer Key (instead of a new randomized equivalent key) each time that you download it from the master Appliance. In this way, if you lost access to the originally downloaded file, downloading the Customer Key again lets you compare it with the key used in the deployment of the Collectors and ensure that it is the right one.

Deprecated features

- Support for IE 9 and 10.

Data-model changes

Introducing Last Known Collector Status

Starting from V6.11, you can now see more information in Nexthink about the status of the connection between your Collectors and Engines. This change introduces the following new fields:

Field	Group	Type			
Last known connection status	Nexthink Collector	Field			
	Indicates the last known connection status of the device: • 'UDP': the device successfully connected via UDP but not TCP. • 'TCP: the device successfully connected via TCP but not UDP.				

	• 'UDP+TCP': the device successfully connected via both UDP and TCP. • '-': Collector version is below V6.6.				
	NXQL ID:	last_known_connection_status			
Last seen on TCP	Nexthink Collector	Field			
	Indicates the last time that the device was successfully connected through the TCP channel. • '-': the Collector is an older version that does not support TCP.				
	NXQL ID:	last_seen_on_tcp			
Number of days since last seen on TCP	Nexthink Collector	Field			
	Indicates the number of days since the last time that the device was successfully connected through the TCP channel. The field is updated every hour: • '-': the Collector is an older version that does not support TCP.				
	NXQL ID:	number_of_days_since_last_seen_on_tcp			

What's new in V6.10

New features

Branding campaigns

To give a professional look to your campaigns, brand the notifications received by the end-users with the corporate logo and a customized color palette for the buttons.

Branding a campaign helps catch the eye of the end-user. In addition, by putting the corporate image on display, customized notifications may give end-users the necessary confidence to readily answer the questions of the campaign.

Find out more

More flexibility for published campaigns

Even when carefully designing a campaign, it is possible to make mistakes. Right after publishing a campaign, you may realize that you omitted an option from a question, that the formulation of a question makes it difficult to understand, or that you simply misspelled a word in the title.

To let you remedy these and other situations, campaigns are now editable while in the published or retired states.

Find out more

Data-model changes

Domain Threat Level changed to Reputation

With V6.10, the field **Threat Level** for Domain objects has been deprecated and has been replaced by a more robust field called **Reputation**. With a reputation ranking, a domain owner's reaction when their web site is compromised is taken into account.

Field	Group	Type			
Reputation	Properties	Field			
	Indicates the reputation of the domain: • '-': internal domain or not yet tagged • 'trustworthy': clean domain which has not been connected to any security risks • 'low risk': benign domain which rarely delivers dangerous content • 'moderate risk': generally benign domain which has exhibited potentially risky behavior • 'high risk': potentially malicious domain which delivers dangerous content				
	NXQL ID:	threat_level			

Alongside this change, several new domain categories have been added and others removed.

- New domain categories: "Government", "Reference and research", "Training and tools", "Business and economy"
- Removed domain categories: "Instant messaging", "Business application"

What's new in V6.9

New features

Scores in user and device views

A score is a higher level indicator obtained by weighing together several different technical factors. Scores were introduced as part of V6.7 to allow you to understand the big picture and define KPIs that can be tracked over time. V6.9 enhances this concept by allowing you to understand the details of a score for a specific user or device. There are several use cases where this functionality is used, the most important one being in the context of your helpdesk.

When using Nexthink as part of the incident management process, this new functionality will allow you to:

- **Save time in L1 support by performing quicker analyses**: the red/green indicators enable your team to immediately diagnose common problems.
- **Improve the escalation process and have a more focused L2/L3 support**: by requiring that all escalated tickets comply with the L1 checklist you can ensure that key requirements have been met before escalation.
- **Decrease support costs by enabling shift-left**: the checklist can enable L1 to resolve incidents that would have previously been escalated.

The benefits are not limited to support use cases, however. These new visualizations are available for any score, for example enabling you to fix compliance and performance problems more easily.

To obtain more information about scores and how to customize them for your environment, please contact your Nexthink Customer Success Services representative.

Find out more

End-user feedback got international

Nexthink customers operate in multiple time and language zones, and when communicating with end-users (company employees) precision and clarity are essential.

For this reason we have adopted a standard format supported by numerous tools available on the market to allow you to translate your end-user feedback campaigns to different languages. Getting real-time, targeted feedback from your end-users, no matter where they are, has never been so easy!

Find out more

AD-based account provisioning is official!

Gone are the days when you needed to manually maintain user accounts in Nexthink. With V6.9 you can link Nexthink with Active Directory and provision users and permissions dynamically based on AD security groups. It's more secure, easier and faster!

Find out more

Other changes

Portal API beta

Do you wish to use Portal data in your favourite reporting tools? Get a slice of this feature by participating in our beta program. You will be able to try it out with a limited set of features and give us suggestions on how you would use it.

Did we interest you? To receive instructions on how to activate the feature directly in V6.9, contact us at support@nexthink.com using "Portal API beta program" as the subject of your email.

New aggregates and fields in NXQL

V6.9 contains three new aggregates in NXQL that allow you to better analyze performance problems. Several other fields and aggregates have also been added to NXQL.

Find out more

Visualize hierarchy as a tree

To aid in the configuration of complex hierarchies we have introduced a tool that allows Nexthink administrators to visually analyze the structure of the organization.

Data-model changes

Device performance aggregates

The following aggregates have been renamed in Finder:

- **high IO throughput time** becomes **high device IO throughput time ratio**
- **high page faults time** becomes **high device page faults time ratio**
- **high memory time** becomes **high device memory time ratio**

New aggregates and fields available in NXQL

The following aggregates were already available in the Finder, but are now also available in NXQL:

Name	Type				Properties
high_device_memory_time_ratio	permill				NU
	Indicates the ratio between the time the device is in high memory usage and its uptime.				
high_device_io_throughput_time_ratio	permill				NU
	Indicates the ratio between the time the device is in high IO throughput and its uptime.				
high_device_page_faults_time_ratio	permill				NU
	Indicates the ratio between the time the device is in high page faults and its uptime.				
number_of_printers	integer				
	Number of printers				

The following fields were already available in the Finder, but are now also available in NXQL:

Table	Name	Type				Properties
device	logical_cpu_number	integer				NU
		Indicates the number of cores multiplied by the number of threads that can run on each core through the use of hyperthreading.				
domain	storage	enum				
		Event storage policy for the domain (web request or none)				
printer	real_name	string				
		Most frequently seen display name				
user_activity	real_duration	millisecond				
		Indicates the time between the user logging on and the device being ready to use. Desktops and laptops are considered fully functional once the CPU usage drops below 15% and the disk usage drops below 80%, and servers once the CPU usage of all processes belonging to the corresponding user drops below 15%.				

Deprecated field

The field **Privileges** for users has been deprecated.

Deprecated features

Starting from V6.9, the Web API V1 is deprecated in favor of the more powerful and flexible Web API V2 (NXQL). Existing integrations will continue to work; Nexthink will communicate in advance before removing the Web API V1 from the product.

Q: What will happen to existing integrations using Web API V1?

- A: For the moment they will continue to work. However, we recommend to migrate them to the Web API V2 (NXQL) whenever possible.

Q: I?m planning a new integration. Can I still use the Web API V1?

- A: No. All new integrations must use the new Web API V2.

What's new in V6.8

New features

Workflows for questions

End-user feedback has been out for only one short release, but already we received a lot of feedback (no pun intended) and we are happy to propose the first enhancement to the solution. Workflows allow you define the flow of questions based on the user's answers. Let's consider a simple scenario: we want to understand why people are using Dropbox instead of the corporate document sharing solution. Depending on the answer we want to ask different follow up questions:

- I use Dropbox for personal files nothing left to ask
- Because the corporate solution suffers from technical issues which ones?
- I didn't know about the corporate solution here are some instructions

Find out more

Email digests just got an upgrade

The email digest gives you a concise update on what happened during the past week and it just got an upgrade! So what did we change? First of all we have included the ability to use scores as part of your digest metrics. Digests are meant to provide a quick summary of the most important data for you, and scores are a perfect fit for that. This means for instance that you will be able to receive each week a quick summary of your overall endpoint performance or share the latest evaluation of your environment compliance.

Secondly, based on feedback from our customers, we implemented two changes that have been requested several times. Number one: we have modified the smileys to be more *corporate*; meaning that we opted for a flat design. Number two: we now allow you to decide if users with access to a service module will receive or not a digest for it (previously the digest was automatically sent, without an option to opt-out).

Find out more

Express cost with scores, and more

Scores are an extremely powerful concept as they allow you to give business meaning to technical data. In this version we have included a number of changes for scores.

Ability to use sum, multiply, and weighted average operators

These simple functions allow you to define more complex scores, and even scores that express monetary notions. Imagine, for instance, users who are using a commercial software while a free alternative is available. Money saved = number of users x cost of the commercial software.

Ability to export scores through NXQL

What's better than a score in Nexthink? Well, of course, it is a score fully integrated into your ecosystem. Starting from now, it is very easy to automatically export scores to your CMDB or ITSM tool.

Find out more

Windows authentication, a.k.a. SSO, it's official!

Gone are the days when you needed to type your username and password to access your favorite dashboards. With this feature enabled, and provided that you logged in to your computer as a domain user (AD account), you can access both the Portal and the Finder with just one single click and without having to know your username or password.

Find out more

Azure support, in preview

The world is moving to the clouds, and because of that we are happy to announce our plan to support your Nexthink environment hosted in Azure. Hosting Nexthink in the cloud offers several advantages:

- No need to worry about the infrastructure required for Nexthink Appliances in your infrastructure;
- Activity for users working outside of the corporate network is automatically recorded, as long as their device is connected to the Internet;
- It is very easy to scale the environment if more computation power is needed;
- As usual maintenance of the environment is very easy, thanks to our automated update system;

- We did all the hard work of trying out all the different Azure packages, and we can recommend the best ones based on your environment.

Azure support is still in preview for now. If you are interested do not hesitate to contact support@nexthink.com.

What's next? Did anyone mention Amazon...?

Other changes

Welcome the Title widget

Having the right data at our disposal is essential. Having that data in a format that looks good and is easy to understand, makes it even better. For this reason we have introduced a new widget to help you find the optimal layout for your dashboard.

Find out more

More options for bar charts

Bar charts are the perfect visualization for comparing things. We have extended it by allowing you to display the ratio value of each group compared to the total value. Imagine that you have a metric that counts your devices with a grouping by OS. You have 100 devices: 60 on Windows 10, 30 on Windows 8.1 and 10 on Windows 7. What you want to see is a bar chart showing three bars, with 60%, 30% and 10% respectively. Today this is now possible.

Find out more

Device identification by name

Specially designed to overcome the shortcomings of the default device identification algorithm in virtualized environments, this already existing feature is no longer limited to Windows devices belonging to a domain. If you are currently using this feature, pay attention to the changes that might impact your setup.

Find out more

Data-model changes

Number of days since last seen for users

The number of days since a user was last seen by Engine is now a recorded value in Nexthink.

Field	Group	Type			
Number of days since last seen	Properties	Field			
	Indicates the number of days since the last time the user was seen by Nexthink. The field is updated every hour.				

What's new in V6.7

New features

End-User Feedback Module

The Nexthink End-User Feedback Module introduces a new way for IT to interact with the end-user, understand their needs and provide services that better align to business needs. By using this module, IT departments will be better able to:

- Discover user satisfaction levels with IT and the needs of end-users
- Understand how to improve the business productivity of end-users
- Quantify the true business impact of problems
- Get feedback during project, changes and transformation projects

Subjective end-user data can be analyzed together with the hundreds of metrics currently reported by Nexthink Collector to obtain true end-to-end visibility on both endpoint and infrastructure performance as well as end-user perception and satisfaction. This unique combination is key not only to understand what is happening but also to extract concrete action items to resolve problems and improve IT satisfaction.

If you are interested in learning more about the End-User Feedback Module and Nexthink's End-User Experience Improvement Solution please contact your Nexthink partner or account manager.

Find out more

Scores

Nexthink provides hundreds of metrics that can be used to measure endpoint and infrastructure performance at a technical level, but how can this information be used to understand the big picture and define KPIs that can be tracked over time? This is where the concept of score comes in. A score is a higher level indicator obtained by weighing several technical factors together.

With Nexthink you can build scores on endpoint performance, infrastructure health, service quality, and many more. Moreover, all scores can be customized to meet your needs. Scores are computed for individual devices or users and are aggregated across the whole environment to provide an understanding of the overall context. They can be broken down by region and analyzed as they evolve during time. To obtain more information about scores and customizing them for your environment, please contact your Nexthink Customer Success Services representative.

Other changes

Advanced tooltip

Our new Portal tooltip allows you to obtain more details about the indicators shown in your dashboards. The tooltip has been standardized for KPI, table and bar chart widgets. It shows the following data:

- Precise representation of the value, for instance, 1149 devices instead of 1.15k
- Details for ratios, for instance, to understand that 4.3% means 4 devices out of 93

- Variation with respect to the previous value
- Threshold, if one has been defined in the corresponding metric
- When applicable, links to access Finder and details

Find out more

More on bar charts

The bar charts introduced in the previous versions have been enhanced with support for quantity metrics, thresholds and fixed scale.

Find out more

Restrictions on default encryption protocols for Engine

We have introduced some restrictions to prevent the establishment of connections using obsolete encryption protocols to/from Engine. Starting with V6.7 by default Nexthink Engine will only support TLS1.1 or later. In most environments this should not have an impact, however we encourage you to validate that all external integrations - including LDAP - are still working correctly after upgrade.

Counting all users and devices - confirmed

Thanks to the very positive feedback, the new feature introduced as beta in the V6.6 Portal, allowing a count of all users and devices has been confirmed. It is now officially part of the product.

Data-model changes

Deprecated fields

In V6.7 there are two deprecated fields:

- Windows 7 (32-bit) compatibility
- Windows 7 (64-bit) compatibility

What's new in V6.6

New features

Bar charts

Comparing data is made trivial with the new bar chart! A bar chart can be processed pre-attentively; in other words it allows you to compare values subconsciously, without having to dwell on it. Hence, prefer bar charts to compare values and reveal the "shape" of your data, but stick to tables for looking up a specific value.

Find out more

Counting all users and devices (beta)

Nexthink Portal has been enhanced with a new way of computing metrics related to inventory use cases. Until now all values shown in Portal were computed by considering *active* objects. For instance a metric counting *Windows 10 devices* returns the number of devices running Windows 10 that were turned on during the period of time specified by the user (for instance a specific day). When it comes to inventory use cases however it might be more interesting to count all *known* devices with Windows 10 installed regardless of the fact that they were turned on or not. With V6.6 it is now possible to define such metrics for both users and devices.

This feature is available for all V6.6 installations, however we have decided to consider it in *beta* state while we evaluate users' response. We would love to hear your opinion at beta@nexthink.com!

Find out more

Finder display field selector

Finder just became more user friendly! When running investigations you need to pick out which fields to display among the hundreds that Nexthink provides. We have introduced a brand new selector that makes it possible to quickly search fields instead of manually navigating a complex menu. Moreover all fields are now better organized thanks to an improved categorization.

Find out more

New delivery model

Want to always have the latest features but find it difficult to keep up with the updates? We have drastically improved our delivery model and hope you find it easy to stay on the cutting edge.

Appliance update and federation

By default the product will remain automatically up-to-date, providing you with all the new features, benefits and security patches introduced in new releases. Moreover, advanced configuration settings allow you to define an update policy that matches internal requirements.

On top of the update policy, other configuration settings can now be automatically replicated across different appliances thus simplifying the

management of the Nexthink platform. The Nexthink Console has been completely reorganized to make the administration job even easier.

Find out more

Finder installer and auto-update

Introduced in beta in V6.5, the new Finder has now been confirmed. Installing the Finder has never been so easy thanks to our new install & update mechanism. Nexthink users with Finder access just need to sign in to the Portal and select **Install Nexthink Finder** after clicking on their username.

After that, it is just a question of launching the downloaded installer. This will start the Finder with pre-configured settings to connect to your Portal. From this point on, the Finder will update itself, seamlessly even if newer versions of the Portal and the Engine are installed.

Find out more about installing the Finder and updating the Finder.

Install and forget Collector (preview)

After the first installation, you will never have to think about Nexthink Collectors again. Starting from V6.6, Collectors are in fact able to update themselves by retrieving the latest version from their Engine. As for the Appliances, extensive configuration options allow you to apply the update policy that will best suit your company.

Find out more

New hardware requirements

The hardware requirements for the Nexthink Engine have been increased to accommodate the new delivery model. All V6.6 (and later) installations must now conform to the specification described in the hardware requirements page.

Other changes

NXQL limits

The 7-day limit imposed for certain NXQL queries has now been lifted. Developers however should pay attention to potential performance impact of running a large number of highly complex queries.

Finder and Windows Vista

Windows Vista is no longer supported to run Nexthink Finder V6.6. As Vista is no longer actively supported by Microsoft and its usage share is extremely low, we hope this should not be an issue to our users. In case you are still using this platform, we encourage you to move to a newer operating system to continue receiving the latest Nexthink versions and features.

Find out more

SHA-1

SHA-1 is now available on top of MD5 for binaries. This new field does not require a Collector update.

Data-model changes

SHA-1

SHA-1 is now available on top of MD5 for binaries. This new field does not require a Collector update.

Field	Group	Type			
SHA-1 hash	Properties	Field			
	Indicates the SHA-1 hash of the binary.				

New and deprecated Collector Update fields

Starting from V6.6, the Collector is now capable of updating itself by retrieving the latest version from the Engine. With this enhancement the Nexthink Updater is now redundant, and it has therefore been deprecated. As a result, there are some new fields relating to the Collector's ability to update itself, and some deprecated fields that related to the Updater.

New Collector update fields

The following fields have been added to support the Collector's ability to update itself:

Field	Group	Type			
Target version	Nexthink Collector	Field			
	Indicates the Collector package version that is targeted.				
Collector update group	Nexthink Collector	Field			
	Indicates the update group of Nexthink Collector: • manual: the Collector is manually updated • pilot: the Collector is updated as part of the pilot group • main: the Collector is updated as part of the main group.				

Collector status	Nexthink Collector	Field			
	Indicates the status of the Nexthink Collector package installed on the device: • unmanaged: the Collector is not automatically updated • up-to-date: the Collector is up-to-date • outdated: a newer Collector version is available.				
Last update	Nexthink Collector	Field			
	Indicates the last Collector update time.				
Last update status	Nexthink Collector	Field			
	Indicates the status of the last Collector update: • '-': the Collector was never updated • successful installation: the last Collector installation was successful • package download error: the Collector was not able to download the Collector package from Nexthink Appliance • package digital signature error: the Collector was not able to check the Collector package digital signature • device reboot required: the device needs to be rebooted to complete the Collector installation • package error: the Collector package				

	installation has failed • internal error: the Collector package installation has failed for an unexpected reason.

Deprecated Updater fields

The following fields relating to the Updater have been deprecated:

- Updater version
- Collector update status
- Last Updater request
- Updater error
- Collector installation log

You can still enable these fields to check the update status of your old Collectors.

Related tasks

- Viewing Collector deprecated fields

Deprecated features

Applies to platforms:

Updater

The existing Nexthink Updater is not able to upgrade Collectors to V6.6.

Since V6.6, Collectors have the intrinsic ability to update themselves. The Nexthink Updater has therefore been deprecated. Customers relying on the Updater to maintain their Collectors up-to-date are required to either:

- Use the executables generated by the Nexthink Collector Installer to deploy the Collector. These executables automatically remove the deprecated Updater when installing the new Collector.
- Remove the Updater from any device before deploying the Collector on it by means of the MSI. Uninstall the Updater via the Windows **Settings** (Control Panel) or using your favorite deployment tool. The MSI fails to install the Collector on those devices where the deprecated Updater is still present.

Force Collector installation on Servers

Every generated Collector installer is now able to install the Collector on any kind of device, be it a laptop, a desktop, or a server. The MSI parameter **DRV_FORCE_SERVER**, which was required to be set to 1 for installing the Collector on servers, has been deprecated and is no longer available.

Running multiple Engines on the same Appliance

Since the Engine is a very demanding application in terms of computing power, the possibility of running more than one Engine in the same Appliance has been ruled out.

Related tasks

- Updating the Collector

Related references

- Collector MSI parameters reference table

What's new in V6.5

New features

Improved boot and logon duration metrics

We have introduced several changes to improve the accuracy and usefulness of boot and logon duration metrics: boot values are now more precise and for user logon we now distinguish between an objective measurement representing the desktop being shown and a subjective measurement representing the device being optimally ready for use. More details are available in the Data-model changes page.

More granularity for visualizing CPU activity

In V6.4 we introduced a new device warning event called **High overall CPU usage** to better represent situations of high CPU consumption. This event is triggered when 70% of all logical processors are used, or in other words, when at least 70% of the total CPU capacity of the device is consumed.

In V6.5 we have modified the user and device views to include this event.

In the case of long lasting *blue* warnings: specific applications are consuming a large amount of CPU, but the overall device experience is not compromised. It can be interesting to investigate these situations because the high CPU usage might be abnormal and indicate an issue with specific applications. Moreover, battery life of laptops may be drastically compromised by applications that consume an elevated amount of CPU for long periods of time.

In the case of long lasting *yellow* warnings: the user experience is likely to be impacted and further analysis is recommended to identify and resolve the problem.

In both cases, the tooltips might offer additional insights into the applications causing the issues.

Find out more

Other changes

Metrics access rights

A new configuration setting allows you to more precisely control which Portal metrics users are allowed to access.

Migration considerations: the system will automatically choose a value to ensure that no user is given access to metrics that he could not access before the upgrade.

Find out more

New hardware requirements

The hardware requirements for the Nexthink Engine have been increased to accommodate upcoming features. Beginning with V6.6 (September 2016), up to 2GB of RAM and 2 CPUs should be added depending on your configuration.

Find out more

New connectivity requirements

Beginning with this version, both the Nexthink Engine and the Portal will connect to the Nexthink Application Library (in previous versions only the Engine was connecting).

Find out more

New Engine certificate

The default Engine certificate used to establish secure connections between the Web API, Finder, and Portal has been updated. For installations using the default certificates, users who saved the certificate in Finder or in the web-browsers will now again be asked to add an exception.

New Finder install and update mechanism, in beta

Installing the Finder has never been so easy thanks to our new install & update mechanism. Nexthink users with Finder access just need to log in to the Portal and select **Install Nexthink Finder** after clicking on their username.

After that, it is just a question of launching the downloaded installer, which will then directly start the Finder and configure it to connect to the right Portal.

From this point forward, the Finder will be automatically and transparently updated, even in case of new releases of the Portal and the Engine.

Find out more about installing the Finder and updating the Finder.

New HW requirements

Overview

The hardware requirements for the Nexthink Engine have been increased to accommodate upcoming features. Beginning with **V6.6 (September 2016)**, up to 2GB of RAM and 2 CPUs should be added depending on your configuration.

If you have an existing installation, you do not have to worry, as you will be able to upgrade and run the latest version of the product without upgrading your hardware. However, you should be aware that you will not be able to turn on the new features until the new requirements are met.

It is very important that new installations observe the new requirements and that you upgrade existing installations whenever you have the opportunity.

Check the new hardware requirements for multi-appliance set-up and for single-appliance.

FAQ

What are the new requirements?

We have two sets of new requirements:

- Minimum: these requirements must be implemented so the product runs smoothly. If customers are not able to implement the new minimum requirements, it will not be possible for Nexthink to provide technical support.
- Recommended: these requirements should be implemented to increase product performance and provide an extra buffer for future enhancements.

Should I use the minimum or the recommended requirements?

Whenever possible, we suggest that you implement the higher recommended requirements.

When are the requirements changing?

The new minimum and recommended requirements will change in V6.6 (September 2016), but new and existing installations should implement the new requirements as soon as possible.

What should I use for new installations?

You should use the new minimum requirements or (better) the recommended requirements. The goal is that the installation is ready for 6.6 and later.

What will happen if I upgrade existing installation to 6.6, but I do not meet the mimimum requirements?

No worries, the product will still work, but you will not be able to turn on the new features.

Why have we increased the minimum requirements?

New features will be released during the second half of 2016 and the beginning of 2017. These features will require more memory and CPU.

The requirements have not changed for many years, but the technology has evolved. It is necessary to implement the new requirements to provide a superior user experience and take advantage of the new features.

Data-model changes

Improved boot and logon duration metrics

Changes for event system boot

Boot duration accuracy has been improved by modifying the process used to define the completion of the boot process. Please note that this improvement does not require a new version of the Collector.

Field	Group	Type			
Duration	Properties	Field			
	Indicates the time between the kernel start and the launch of the 'logonui.exe' process				

Find out more

Changes for user logon

To improve the accuracy of the logon duration we have introduced two changes.

A new definition for logon duration: in previous versions of the product the logon duration represented the moment in time when the CPU usage dropped below a certain threshold. Although this could be used to judge the user experience, the value was too open to interpretation and discussion. For this reason we have modified the measurement to report a more objective value: the logon duration now represents the time elapsed between entering the credentials and the desktop being shown to the user.

Extended logon duration: a subjective measurement of the logon duration is still very useful to analyze the user perception and experience. For this reason we have introduced a second measurement called **Extended logon duration**. This measurement represents the time needed for the device to become optimally usable after a logon. To obtain this value we measure CPU and disk usage.

Please note that a new version of the collector is needed for part of this feature.

Find out more

Field	Group	Type			
Duration	Properties	Field			

	Indicates the time between the user logging on and the desktop being shown.
new Extended duration	Properties | Field Indicates the time between the user logging on and the device being ready to use. Desktops and laptops are considered fully functional once the CPU usage drops below 15% and the disk usage drops below 80%. Servers are considered fully functional once the CPU usage of all processes belonging to the corresponding user drops below 15%.

Find out more

New fields and aggregates for devices

Field	Group	Type
Last logon duration	Startup	Field
	Indicates the last recorded value for the time between the user logging on and when the desktop is displayed.	
new Last extended logon duration	Startup	Field
	Indicates the last recorded value for the time between the user logging on and the device is ready.	
Logon duration baseline	Startup	Field
	Indicates the logon duration averaged over the last logons. In the calculation, recent logons weigh more than older logons (exponentially weighted moving average).	
new Extended logon duration baseline	Startup	Field
	Indicates the extended logon duration averaged over the last logons. In the calculation, recent logons weigh more than older logons (exponentially weighted	

	moving average).				
Average logon duration	Startup	Aggregate			
	Indicates the average logon duration.				
new Average extended logon duration	Startup	Aggregate			
	Indicates the average extended logon duration.				

Migration considerations

After migrating to the new version, the following changes can be expected due to the new definition and calculation of boot and logon duration:

- Average boot duration values are expected to grow
- Average logon duration values are expected to drop significantly since desktops are always shown before the CPU drops

Customers wishing to continue to measure logons based on a subjective value (representing the device being ready) instead of an objective value (representing the desktop being shown) must modify metrics and investigations to use **Extended logon duration** instead of **Logon duration**.

What's new in V6.4

New features

Investigate in Finder

A lot of the power and ease of use of Nexthink relies on the fact that Finder users can always *drill-down* to find out more information about a specific topic. Drill-downs in fact allow us to understand who is impacted by an issue, discover the dependencies between CIs and observe the exact events that lead to a specific situation.

It's only logical to extend such functionality to the Portal as well. From now on you will be able to click on any value shown in Portal to **Investigate in Finder**. This feature is so seamlessly integrated that in a matter of hours you will ask yourself *How could I work without this before?*

Find out more

More scalability and speed for Engine

Scalability and speed are crucial topics for us and breaking existing limits has become one of our preferred obsessions. We are excited to announce that V6.4 Engines will support double the number of concurrent Finder users compared to V6.3, and we will do so with the same hardware! Moreover we have drastically improved the scalability of the product; if you want more speed you just need to add 1 CPU core for each 5 Finder users.

But this was not enough... scaling more is great but we also wanted to give you more speed. So, once again, without having to change any of your hardware, you will discover that V6.4 queries take in average 30-40% less time to complete!

Find out more about the Engine hardware requirements

Other changes

Collector tag for entities

The Collector tag (a value that can be set during Collector installation and is then reported by the Collector to the Engine) can now be used in Entity rules.

Find out more about setting collector tags and using collector tags

A new option for the nxt:// protocol

A new option for the nxt:// protocols enables you to open the device/user view with a specific time frame.

Find out more

User-Device views navigation

When navigating from the user view to the device view or vice versa, the time frame is kept.

Data-model changes

Average memory usage per execution

Introduced in V6.3 and previously restricted to objects of type Executable and Binary, this aggregate has now been extended to support the following objects:

- Users
- Devices
- Applications
- Executables
- Binaries

Field	Group	Type			
Average memory usage per execution	Activity	Aggregate			
	Indicates the average memory usage of all underlying executions before aggregation. The value is the average				

	memory usage of all executions (calculated with a 5-minute resolution) multiplied by their cardinalities and divided by the total cardinality. • Example: if two tabs of the Chrome browser are opened at the same time, two distinct processes of chrome.exe are launched and they are aggregated by the Engine (i.e., event cardinality = 2). The average memory usage will be the average of the two processes before aggregation: it represents the average memory usage of a Chrome tab.

Find out more

High CPU usage changes

A new device warning event called **High overall CPU usage** has been introduced to better represent situations of high CPU consumption. This event is triggered when 70% of all logical processors are used, or in other words when at least 70% of the total CPU capacity of the device is consumed. This value is configurable.

The aggregate **High device overall CPU time ratio** represents the ratio of time where the device suffers from high overall CPU usage.

Field	Group	Type			
High device overall CPU time ratio	Warnings	Aggregate			

	Indicates the ratio between the time that the device is in high overall CPU usage and its uptime.

Deprecated fields and aggregates

To simplify high CPU events following the introduction of the new **High overall CPU usage** event, the following events and aggregates have been deprecated:

- The event **High CPU usage** has been renamed to **High thread CPU usage** and has been deprecated.
- The aggregate **High CPU time** has been renamed to **High device thread CPU time ratio** and has been deprecated.

Please note that CPU warning events in the user and device views still refer to High thread CPU usage and have not yet been ported to High overall CPU usage)

Find out more

High application thread CPU time ratio

This new aggregate can be used to identify applications generating a large volume of **High thread CPU usage** events, or in other words applications with high CPU usage peaks. This aggregate is available for the following objects:

- Applications
- Executables
- Binaries

Field	Group	Type			
High application thread CPU time ratio	Warnings	Aggregate			
	Indicates the ratio between the time that the underlying executions are in high thread CPU usage and their execution duration.				

Find out more

What's new in V6.3

New features

Email digests

The email digest gives you a concise update on what happened during the past week. You certainly have several areas of responsibility; for instance you may be monitoring dashboards in both "Shadow IT" and "Malware protection" modules, or many more.

The digest gives you an overall summary of the status of your modules based on what you consider your most important metrics.

Find out more

Improved features for content creators

In this release we have included two features targeted especially to Nexthink content creators.

Run metric as investigation

It is now possible to run metrics just as if they were investigations. The Finder can automatically translate any metric into an investigation and run it, letting you quickly verify the data computed by the metric. This feature is available from both the metric context menu in the accordion and the metric designer.

Export nxt:// actions

The nxt application protocol provides you with the means to launch the Finder and perform some specific actions on it by just stating a URL. Creating nxt:// links is not always trivial however. For this reason, we have now added the ability to export contextual nxt:// links directly from the **Export** menu of the following accordion objects:

- Investigations: export a link that will execute the investigation, even if you have not imported it yet. This feature is incredibly useful to quickly share investigations or embed them in an email or document.
- Services: export a link that will open the service view.
- Metrics: export a link that will open the metric designer.
- Categories: export a link that will open the category designer.

Other changes

More history, now by default in the product

With V6.3 everyone will automatically benefit from the data history optimizations we introduced in V6.1. In fact the default aggregation policy will be automatically changed to **medium**, ensuring more history and an overall faster product. If you prefer to keep things unchanged, you can go back to **low** or **very low** directly from the Web Console.

Find out more

High DPI displays now supported by Finder

High DPI displays are becoming more and more common, especially among laptops. When your display packs 5M pixels on 13 inches however, chances are that you want to increase your zoom level in Windows. Previously, the Finder had some trouble to deal with zoom levels greater than 125%. All of that has been solved now, so you can enjoy the Finder at any high resolution and zoom factors.

SSO in beta

Active Directory SSO is now available as a beta feature for the Portal. Gone are the days when you needed to type your username and password to access your favorite dashboards. With this feature enabled, and provided that you logged in to your computer as a domain user (AD account), you can access the Portal with just one single click. Did we manage to interest you? You can contact beta@nexthink.com to receive detailed instructions on how to activate this feature.

Data-model changes

CPU of each execution

Starting from this release, Engine will store the CPU consumed by each program execution. The corresponding data can be extracted using the two following aggregates that apply to:

- Users
- Devices
- Applications
- Executables
- Binaries

Field	Group	Type			
Total CPU time	Activity	Aggregate			
	Indicates the sum of the CPU time of all executions on each device in scope and over all logical processors. • Example: if we consider two executions with the first one taking 50% of a logical processor during 30 minutes and the second one taking 100% of 2 logical processors during 60 minutes, the total CPU time is 135 minutes (= 50% * 30 min + 2 * 100% * 60 min).				
CPU usage ratio	Activity	Aggregate			
	Indicates the sum of the CPU time of all executions on each device in scope over all logical processors divided by their total duration. • Example: if we consider two executions with the first one taking 50% of a logical processor during 30 minutes and the second one taking 100% of 2 logical processors during 60 minutes, the CPU usage ratio is 150% (= [50% * 30				

	min + 2 * 100% * 60 min] / [30 min + 60 min]).

Moreover this data is also available in execution events:

Field	Group	Type			
Total CPU time	Activity	Aggregate			
	Indicates the sum of the CPU time of all executions on each device in scope and over all logical processors. • Example: if we consider two executions with the first one taking 50% of a logical processor during 30 minutes and the second one taking 100% of 2 logical processors during 60 minutes, the total CPU time is 135 minutes (= 50% * 30 min + 2 * 100% * 60 min).				

Memory of each execution

Starting from this release, Engine will store the memory consumed by each program execution. The corresponding data can be extracted using the following aggregates that applies to:

- Executables
- Binaries

Field	Group	Type			
Average memory usage per execution	Activity	Aggregate			
	Indicates the average memory usage of all underlying executions before aggregation. The value is the average				

	memory usage of all executions (calculated with a 5-minute resolution) multiplied by their cardinalities and divided by the total cardinality. • Example: if two tabs of the Chrome browser are opened at the same time, two distinct processes of chrome.exe are launched and they are aggregated by the Engine (i.e., event cardinality = 2). The average memory usage will be the average of the two processes before aggregation: it represents the average memory usage of a Chrome tab.

Moreover this data is also available in execution events:

Field	Group	Type			
Average memory usage	Properties	Field			
	Indicates the average memory usage of the underlying executions before aggregation with a sampling resolution of 5 minutes. • Example: if two tabs of the Chrome browser are opened at the same time, two				

	distinct processes of chrome.exe are launched and they are aggregated by the Engine (i.e., event cardinality = 2). The average memory usage will be the average of the two processes before aggregation: it represents the average memory usage of a single Chrome tab.

New field

The following field has been added:

Field	Group	Type			
Hard disks manufacturers	Local drives	Field			
	Indicates the list of hard disk manufacturers				

Changes in *packages*

Starting from Nexthink V6.3, those investigations retrieving packages or including a condition on packages have been simplified. The results take into account only those packages that are effectively installed, discarding uninstalled packages.

Find out more

Deprecated fields

The following fields of Binary have been deprecated in favor of the more expressive aggregates presented above:

- Average CPU usage
- Average memory usage

What's new in V6.2

New features

V6.2 comes with a wealth of new features aimed at simplifying and improving the use of the product. Moreover, we did substantial work on further optimizing Engine performance.

Dashboard description

John, the Nexthink administrator at Acme Corp. (a fictitious customer), has just finished creating a great dashboard that can be used to discover and track the usage of Shadow IT products in the organization. He wants to share this dashboard with several people in the IT team, but he's afraid that without some explanations not everyone will be able to fully understand the content and how to use it.

With V6.2 John can now write documentation directly inside of Portal and even create links to investigations in Finder. Thanks to this feature John is sure that everyone will be able to fully understand the risks posed by Shadow IT. Just like our user John, you too can now make sure everyone can fully understand the content of your dashboards.

Find out more

Microsoft DirectAccess support

DirectAccess is a technology from Microsoft that allows remote users to securely access internal network file shares, Web sites and applications without connecting to a VPN. DirectAccess works by creating a IPv6 tunnel from the remote PCs to the DirectAccess server. Starting from V6.2 all Nexthink components are able to communicate in a DirectAccess environment; moreover Collector will report network and web traffic transiting through a DirectAccess tunnel.

Find out more

Details in the past

Nexthink Portal allows you to track the evolution of your metrics for an unlimited period of time. Moreover, for metrics of type *count*, additional details about the involved objects are also available. For instance, if you click on a metric tracking the number of devices infected by malware, you will see the full list of infected machines. These details were, until today, limited to the current timeframes (yesterday, current week, current month, current quarter).

The latest version of Portal allows you to reserve additional disk space on the Portal appliance to store details for a longer period. If you want more data, you just need to add more disk space.

Find out more

Portal on your Operation Center big screens

Thanks to the real-time service overview dashboard introduced in V6.0, Portal is the ideal product to be displayed on your Ops Center big screen. To facilitate this use case, you can now configure a special account so that it's never signed out from Portal.

This gets even better when you want to display multiple dashboards in a slideshow. There are a number of free browser plugins that allow you to do just that!

Find out more

Improved Smart Search

The Finder search is getting even smarter. The system now provides suggestions based on services names and entities; for instance you can search for *users of SAP* or *devices in Dublin or Rome*.

In addition, we've added a set of new suggestions:

- New binaries/applications/executables
- Application Library fields
    - Domains classified as ? [e.g. Malicious domains]
    - Domains hosted in ? [e.g. Domains France]
    - Binaries classified in ? [e.g Binary virtualization]
- All servers
- Devices with low network availability
- Devices with high network response time
- Search user with full name (AD)

Find out more

Faster investigations

V6.2 comes with an Engine optimized for speed. Investigations will run up to 3.5 times faster thanks to increased parallelism during the computation of complex investigations and some code-level performance optimizations. You can accumulate this with the aggressive aggregation policies introduced in V6.1 for an even greater performance gain. Existing customers can contact Nexthink

Customer Success Services to discuss the best data optimization strategy for their infrastructure.

Other features

Improved access rights

We improved the way access rights are assigned. Now any *central administrator* can be given exactly the same rights as the *main central administrator*. Note that by default all *central administrators* will automatically gain the right to manage licenses. Central administrators with the *system configuration* right will automatically be able to publish Web API investigations and trigger a manual Engine AD sync. Find out more

Security improvements

When installing the product for the first time, HTTPS is the default Portal setting. Legacy HTTP access can still be activated in the Nexthink console.

nxt:// protocol

We've added two additional commands to the nxt:// protocol which allow you to edit metrics and categories. Find out more

Default aggregation policy

The default aggregation policy has been changed to *normal*. In general this increases the available Engine history by up to 10%.

Full traffic anonymization

Whether you need this for your pre-production environment or to comply with your privacy policy, you can now chose to completely anonymize Collector traffic, even before it reaches the Engine. Find out more

Data-model changes

Nexthink V6.2 comes with 10 new aggregates to get better and faster answers out of the product.

Application stability

These two aggregates can be used to identify your least stable applications, even if they are used by just a few users. These aggregates are available for the following objects:

- Users
- Devices
- Applications
- Executables
- Binaries

Field	Group	Type			
Application crash ratio	Errors	Aggregate			
	Indicates the number of application crashes per 100 executions.				
Application not responding event ratio	Errors	Aggregate			
	Indicates the number of application not responding events per 100 executions.				

Incoming and outgoing network traffic per device

These two aggregates can be used to identify applications that are generating a large amount of network traffic, even if they are used by just a few users. These aggregates are available for the following objects:

- Applications
- Executables
- Binaries
- Ports
- Destinations

Field	Group	Type			
Incoming network traffic per device	Volume	Aggregate			
	Indicates the incoming network traffic divided by the number of devices.				
Outgoing network traffic per device	Volume	Aggregate			
	Indicates the outgoing network traffic divided by the number of				

	devices.

Incoming and outgoing web traffic per device

These two aggregates can be used to identify applications that are generating a large amount of web traffic, even if they are used by just a few users. These aggregates are available for the following objects:

- Applications
- Executables
- Binaries
- Ports
- Destinations
- Domains

Field	Group	Type			
Incoming web traffic per device	Volume	Aggregate			
	Indicates the incoming web traffic divided by the number of devices.				
Outgoing web traffic per device	Volume	Aggregate			
	Indicates the outgoing web traffic divided by the number of devices.				

Total network and web traffic

These two aggregates can be used to compute the total web or network traffic. These aggregates are available for the following objects:

- Users
- Devices
- Applications
- Executables
- Binaries
- Ports
- Destinations
- Domains (only web traffic)

Field	Group	Type			
Total network traffic	Volume	Aggregate			
	Total network traffic (incoming and outgoing)				

Total web traffic	Volume	Aggregate			
	Total web traffic (incoming and outgoing)				

Changes in boot and logon duration

There are now two different ways to look at boot and logon duration.

Aggregate values

The following values represent the duration of boots and logons which happened during the timeframe of the investigation. If no boot or logon happened during this timeframe, then a dash (-) is reported.

Field	Group	Type			
Average system boot duration	Startup	Aggregate			
	Indicates the average system boot duration.				
Average user logon duration	Startup	Aggregate			
	Indicates the average user logon duration.				

Baseline values

The downside of the two aggregate values presented above is that if no boots or logon happened for a device during the investigation period, then no value is reported. For this reason we provide two additional values representing the moving average of boot and logon times. The values do not depend on the time frame specified in the investigation.

Field	Group	Type			
System boot duration baseline	Startup	Field			
	Indicates the system boot duration averaged over the last boots. In the calculation, recent boots weigh more than older boots (exponentially weighted moving average).				
User logon duration baseline	Startup	Field			
	Indicates the user logon duration averaged over the last logons. In the calculation, recent logons				

	weigh more than older logons (exponentially weighted moving average).

What's new in V6.1

New features

With V6.1, Nexthink fully supports migrations from earlier versions of the product. Moreover, V6.1 Engines can be optimized to store up to twice the amount of history with respect to V5.

Ready for migration

With this new release, Nexthink supports migrations from Nexthink V5.3. In order to simplify the migration process, V6.1 Portal can display, in read-only mode, legacy V5 dashboards. Existing customers can contact their account manager for a personalized migration offer.

Up to 2x history length in the Engine

Thanks to new compression algorithms, Engines can be configured to retain up to twice the amount of history, without any additional hardware requirements and with negligible loss of precision. Existing customers can contact Nexthink Customer Success Services to discuss the best data optimization strategy for their infrastructure.

A new anonymization mode

A new data anonymization mode has been introduced to make users and devices anonymous. This feature is in response to specific customer requests. For instance this mode can be applied to users who need to know if a service is functioning well, but do not need to know if any specific user has a problem. Find out more

Updater

The Nexthink Updater is again being shipped as part of the product. Please note that V6 Collector requires V6 Updater: existing customers relying on Nexthink Updater need to switch to version 6 in order to upgrade Collectors to V6. Find out more

Data-model changes

Metrics

Successful HTTP requests ratio

A new aggregate **Successful HTTP requests ratio** is now available in metrics. This aggregate can be used to track HTTP web services client and server errors.

Forbidden aggregates

Count metrics with a group-by referring to a different object no longer support aggregates conditions which include the value 0 (zero).

What's new in V6.0

New features

Whether you are CIO, IT Manager, Administrator, or an interested line of business manager, End-user Analytics is changing the way IT organizations are aligning their operations with the needs of the business and the end-user. With the V6 release, Nexthink is enabling organizations to accelerate and simplify the management and transformation of their complex IT infrastructure and amid rapidly changing business requirements and end-user work styles.

A brand new Portal

The simple, modern, flat look and feel of Portal V6 brings all focus on the data.

- The separation of the metric definition and UI presentation brings more power to you: now easily define the metrics that you want to compute and then combine them in your favorite visualizations. Find out more about Creating A Metric and Following The Evolution Of A Metric.

- Time and location have been unified in dashboards allowing you to compare data at a glance as you navigate. Find out more

- The new layout manager in Portal V6 based on award winning visual concepts allows you to easily arrange elements in a dashboard, any way you want and it always looks great! With new widgets, graph types, immediate previews and simplified steps designing and sharing custom and role-based dashboards is now a matter of minutes. Find out more

- The new service overview dashboard in Portal V6 helps you understand at a glance the status of all your IT services from the perspective of the end-users, in real-time. New service detail dashboards help you quickly understand how a service is used, where problems are located and identify users that are impacted. Find out more

User view

The new User View in Finder V6 presents all devices, information, activities, issues, changes and services related to an end-user, all in one place and against one timeline. In one click understand if an event or issue is reoccurring for a

specific user, since when and how often. New drill downs will accelerate problem identification and resolution by enabling you to check how many end-users are affected by similar patterns. Find out more

Server Collector

Extend your End-user Analytics with Windows Server Collector V6 to go beyond the first destination and start discovering, mapping and understanding end-to-end dependencies related to the end-user experience and service consumption while increasing overall security and compliance.

Content centralization

In the new V6 platform metrics, services, and categories are centralized and automatically synchronized across all Engines. Find out more

- Changes in categories and services are automatically reflected in dependent metrics and services to simplify the configuration.
- Metrics can be easily created in Finder starting from an investigation, and few click later you will be visualizing them in your Portal dashboard. Find out more
- Service thresholds are defined directly within Finder. Find out more
- Finder automatically proposes the list of available Engines during connection ? login once, and switch Engine in 3 seconds. Find out more

New system requirements

Portal hardware requirements

The number of cores required by the Portal appliance has been changed for large installations (starting from 20k devices). See Hardware Requirements for more information.

Connectivity requirements

V6 Finder connects to Portal using port 443 for authentication and managing centralized content. To support this, Engine connects to Portal using three additional ports: 7000, 7001 and 7002. See Connectivity Requirements for more information.

Data-model and API changes

Device

Device type

The field **Device type** now includes values **server** and **mobile**.

Number of logical processors

Added a new field **Number of logical processors** representing the total number of threads seen by the operating system.

Entity

The **Entity** field replaces the V5 ***Entity** category. Finder will automatically migrate investigations, one-click investigations and alerts.

Last system update

The semantic of **Last system update** has been modified to take into account only the last successful system update.

IO and page faults

The fields **High IO throughput time** and **High page faults time** can no longer be used with condition on Activities and Events.

NXT protocol

The syntax used to authorize and authenticate a user using the NXT protocol has been modified. See Bidirectional Integration With The Finder for more information.

Deprecated features

Data model

OS version

The field **OS version** has been deprecated in favor of **OS version and architecture**. The Finder automatically migrates those existing investigations, one-click investigations, and alerts that use the deprecated field.

Portal features

Types of widgets

Dashboards have been completely reworked to be visually more appealing and easier to create. In V6, the widgets included in dashboards are directly linked to the new concept of metrics. Therefore, all V5-style widgets have been deprecated, except for the software metering widget (at least partially).

Widget-related alerts

To unify the methods of alerting users, no widget has the ability to independently send email alerts to selected recipients anymore. That includes the software metering widget, even if this widget remains in the V5-style.

VDI assessment and capacity planning

The VDI assessment and capacity planning module is no longer included in the Portal. Corresponding features will be re-introduced in a later product release.

Portal reports

Reports in Microsoft Word format are no longer included in the Portal. An improved version will be included in a later product release.

Finder features

The **Compare with** tool in the **Timeline** tab of the device view has been deprecated. It is kept in the **Properties** tab of the device view, and it appears in the same tab of the new user view.