Nexthink V6.5

Product Overview

Generated: 10/14/2019 10:11 am

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.

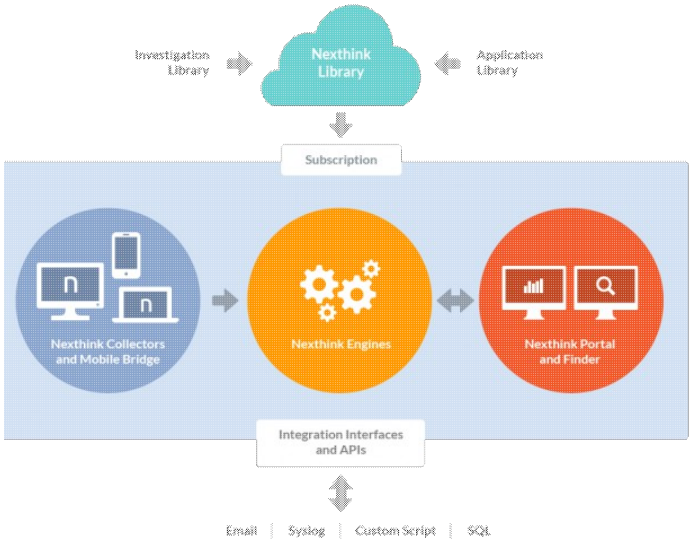

Modular product structure

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

- **Security** provides security related features including binary threat level and category, web threat level, category and hosting country.
- **Web & Cloud** grants access to analytics related to intranet and extranet HTTP and HTTPs web requests.
- **VDI Transformation** includes the analytics and Portal dashboards to ensure a successful VDI transformation project (coming soon for V6).
- **Integration toolkit** enables the product API and access to continuously improved integration samples, reports, etc.

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

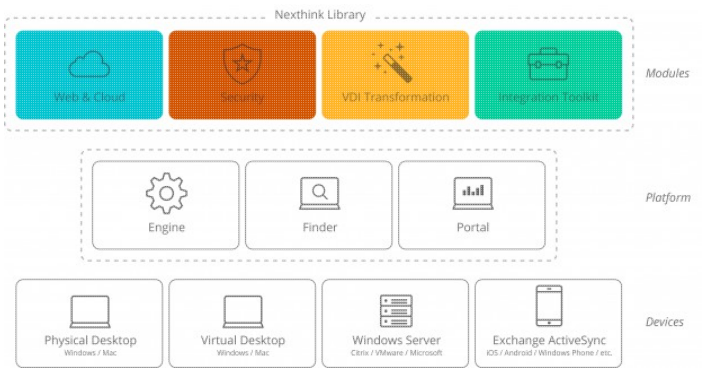

Collector

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%	• Kernel: Around 500 KB • User: Around 20 MB	• Between 0.1 - 0.3 Kbps

The following figure depicts the functioning of the Collector as part of the whole Nexthink solution.

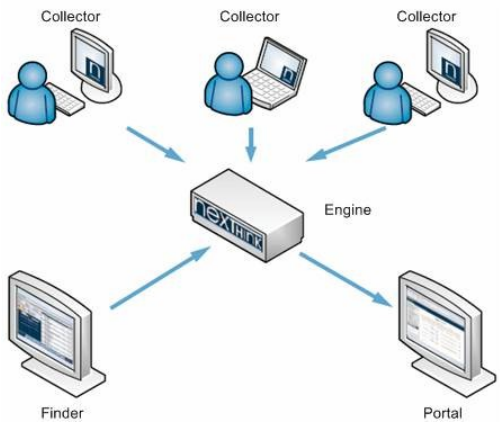

Collector components

The Collector is mainly a kernel driver, but it is made up of several components:

Driver
: The core part of the Collector. It is responsible for gathering information about the devices and the activities of the end-users and for sending them to the Engine.
Applies to platforms:

Updater
: An optional companion service that detects whether there is a new version of the Collector available for installation. If so, the Updater downloads and installs the new version in the device of the end-user.
Applies to platforms:

Control Panel extension
: Add this optional component to see the configuration of the Collector and be able to change it from the Control Panel of Windows.
Applies to platforms:

Features

Multi-Platform

The Collector is available for both Windows and Mac OS operating systems. Originally developed for Windows, the Mac OS version of the Collector has some limitations with respect to its Windows counterpart. Besides Windows specific data, information on web requests and printing is not yet available in the Mac OS version of the Collector.

The automatic deployment of the Collector with the Updater is only available in the Windows platform.

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 Crash Guard feature detects every system crash and 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**.

Detection of Engine

The Collector driver is able to detect when the Engine is not reachable in the local network. In this case, the Collector disables itself for 10 minutes.

Network interfaces supervision

The Collector dectects if a network interface appears on or disappears from the device where it is installed. In this case, the Collector driver resends the whole context to the Engine. The process of adapting to a different network interface may take the Collector a few minutes.

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 changes to become effective.

Related tasks

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

Related references

- Components of the Collector

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

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.

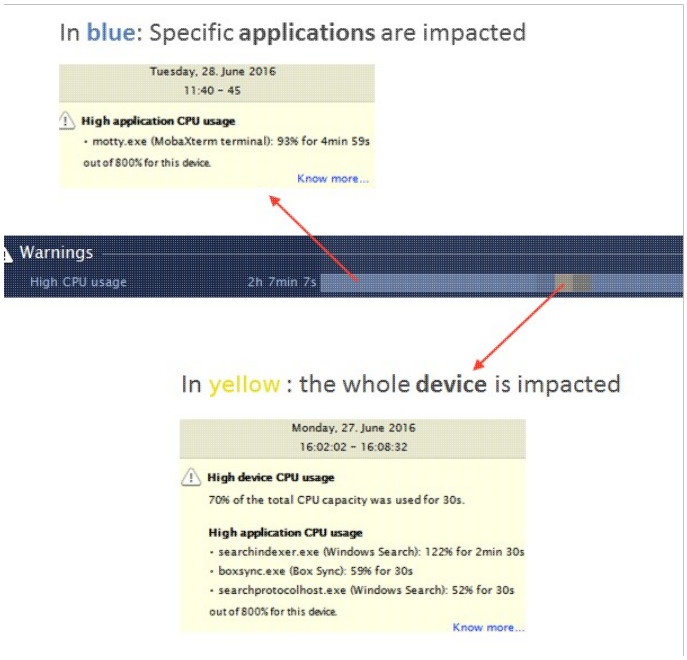

In the case of long lasting *blue* warnings: specific applications are consuming a large amount of CPU, but the overall device experience is not compromised. It can be interesting to investigate these situations because the high CPU usage might be abnormal and indicate an issue with specific applications. 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.

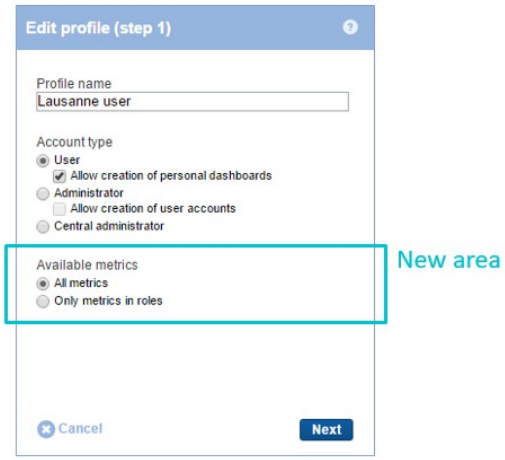

Migration considerations: the system will automatically choose a value to ensure that no user is given access to metrics that he could not access before the upgrade.

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.

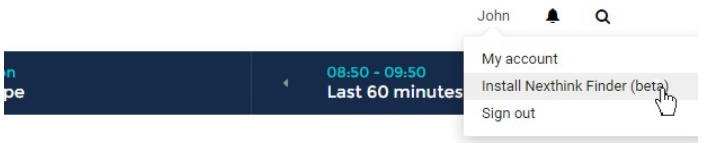

After that, it is just a question of launching the downloaded installer, which will then directly start the Finder and configure it to connect to the right Portal.

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.

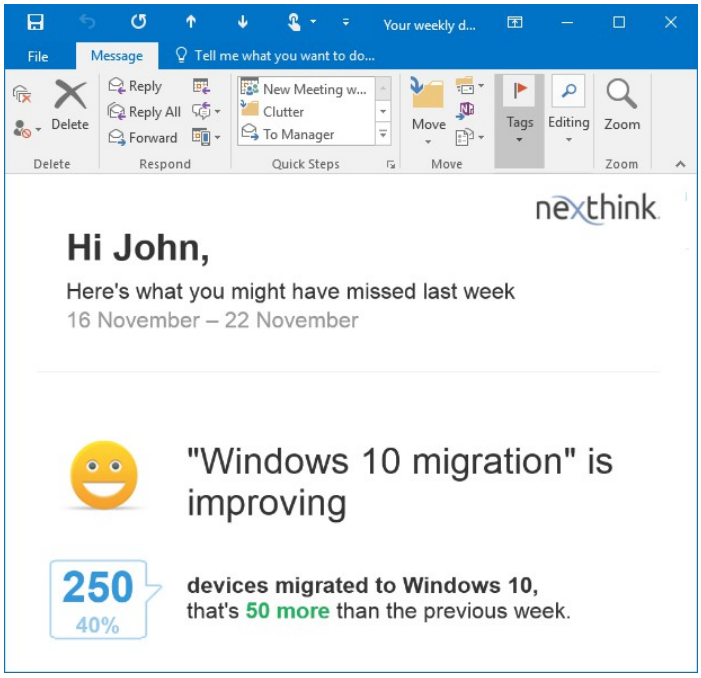

The digest gives you an overall summary of the status of your modules based on what you consider your most important metrics.

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.

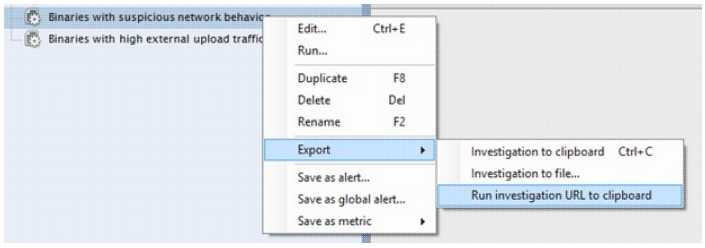

Other changes

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.

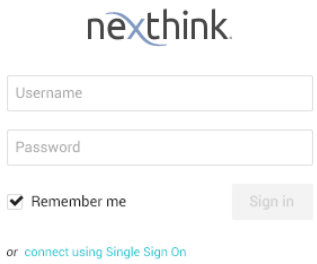

Data-model changes

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.

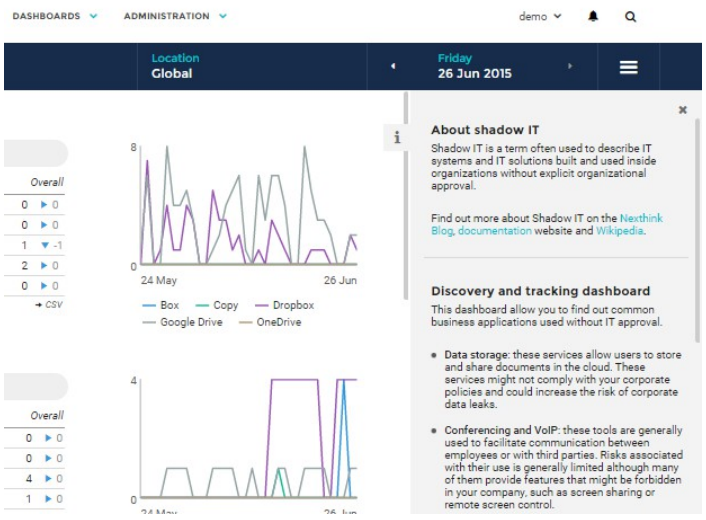

With V6.2 John can now write documentation directly inside of Portal and even create links to investigations in Finder. Thanks to this feature John is sure that everyone will be able to fully understand the risks posed by Shadow IT. Just like our user John, you too can now make sure everyone can fully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

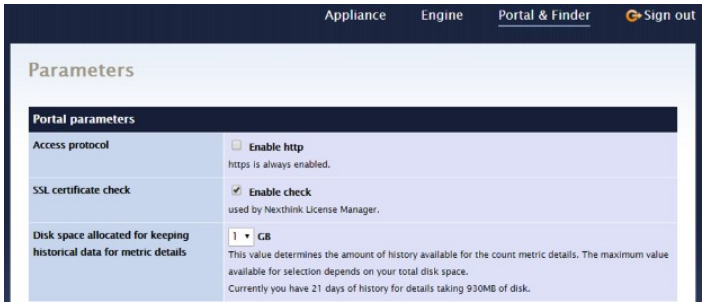

The latest version of Portal allows you to reserve additional disk space on the Portal appliance to store details for a longer period. If you want more data, you just need to add more disk space.

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; moreover the value is now updated even when other tools (such as SCCM) are used to deploy Windows updates.

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.