SPLUNK APP FOR NEXTCLOUD

Reasonably Quick Install Guide

23 August 2019

Bjorn Graabek

Table of Contents

The Splunk App for Nextcloud

The Splunk App for Nextcloud visualises data from a Nextcloud server. Although built by a Splunker, the app is not an official Splunk App nor is it supported by Splunk or Nextcloud, and neither company have any responsibility for the app.

What this guide is and what it isn't

This guide doesn't take you through the only way of performing the necessary steps to monitor a Nextcloud server with the "Splunk App for Nextcloud. Instead it attempts to make the process as simple and safe as possible. It usually doesn't attempt to explain why you should perform a particular step. What this means is that some steps will be performed via the UI of either Nextcloud or Splunk and other steps will be performed from the command-line even though most of the configurations could be performed either way. Ideally it should be possible, with this guide in hand, to install Splunk and configure it to monitor Nextcloud without understanding what a particular step is for, but nevertheless end out with a functioning system. I welcome feedback (bgraabek@gmail.com) if it doesn't, hopefully with some pointers as to what is confusing or misleading.

Assumptions

- You know very little or maybe even nothing about Splunk
- You have a functioning Nextcloud system v10 or higher
- Using the command-line interface on a Linux system doesn't scare you
- You have the relevant hardware (or you can create a VM) to run a Splunk server (more about that later)
- Splunk will be installed on a Linux host, either the same as the one your Nextcloud is running on, or a separate host. I recommend you install it on a separate host. You will need to decide which it will be.
- The Linux host on which Splunk will be installed uses one of the following file systems: ext2, ext3, ext4, btrfs, XFS

- Splunk will be installed as a single-instance deployment and not as a cluster. Not that the Splunk App for Nextcloud can't be used in a Splunk cluster, this quick install guide just doesn't describe how to set it up.
- Splunk itself is a commercial product, a free (but limited) version exists. Unless you already have a Splunk license, you are content with the features provided by Splunk Free. For more information about what you can and can't do with the free version, visit the "About Splunk Free page: https://docs.splunk.com/Documentation/Splunk/latest/Admin/MoreaboutSplunkFree

---

Nextcloud data sources

There are three (four if your Nextcloud server is v14 or above) data sources available, and the Splunk App for Nextcloud makes use of all of them if they are ingested into Splunk. The worst that will happen if not all data sources are ingested is that not all dashboards will "light up". It is not necessary to understand what data is provided by which source to monitor Nextcloud using the Splunk App for Nextcloud, but there may be some sources you care less about, and you might therefore decide not to ingest one or more of them.

Status.php

Calling the status.php file provides simple information as to whether the Nextcloud server is available, in maintenance mode, or not available. For more information see:

https://docs.nextcloud.com/server/12/admin_manual/operations/considerations_on_monitoring.html#status-php

The "Splunk Add-on for Nextcloud" includes a script which is run at 5-minute intervals to pull information in.

API Endpoint

An API endpoint exists through which tools can get information about the Nextcloud server. The "Server Information" app within Nextcloud display some of the information available through the API endpoint, but only in real-time. For more information see here:

https://nextcloud.com/monitoring/

The "Splunk Add-on for Nextcloud" includes a script which is run at 5-minute intervals to pull information in. The scripts can be run from a Splunk Universal Forwarder located on the host running the Nextcloud server or remotely, for example from a Splunk indexer. This guide will guide you through the steps to install the Add-on on a single-instance Splunk server.

Log file(s)

Prior to version 14, by default Nextcloud wrote all log entries to "nextcloud.log. Since version 14, log entries have been split between two log files, "audit.log" and "nextcloud.log". Most of the information used by the Splunk App for Nextcloud is from the "audit.log". For all of the dashboards to work, both log files are however needed. For example, Nextcloud writes log entries for successful logins in the audit.log whereas log entries for failed logins are written to Nextcloud.log

The default location of these log files is "/var/www/nextcloud/data", however, if you have changed your Nextcloud data directory or changed the location in the config/config.php file, the files will be elsewhere.

This guide will guide you through the steps to configure a Splunk Universal Forwarder to ingest both log files.

Configure your Nextcloud server

Logged in as an administrator in the Nextcloud UI, click on the icon for your user in the upper right corner, click on "Apps". Enable the Nextcloud "Auditing / Logging" and the "Monitoring" apps.

In the Nextcloud config/config.php file, you either need to change or add the 'loglevel' parameter. It should be set no higher than '1'. The "log_type" and "logfile" parameter defaults should be fine, but if you wish to ensure they are set correctly, set them as shown below. The default date and time format works well with Splunk, so I suggest you don't change it.

```
'log_type' => 'file',
'loglevel' => 1,
```

For more detailed information, see:
https://docs.nextcloud.com/server/14/admin_manual/configuration_server/logging_configuration.html

https://docs.nextcloud.com/server/14/admin_manual/configuration_server/config_sample_php_parameters.html#logging

If the config.php file is changed, Nextcloud must be restarted. This is done by restarting the web server.

Ensure you know the path to where the 'nextcloud.log' and 'audit.log' files are written to by Nextcloud, this information will be needed later.

Now log into your Nextcloud system as an administrator.

- Click on the letter representing the user id in the upper right corner and choose "Settings".
- Click on "Security" in the "Personal" section on the left side.
- Type "splunk" (or something else, it doesn't matter what name you use), and click on "Create new app password".
- Make a note of the username and password shown and click "Done".
- Click the three dots (…) and un-check "Allow filesystem access" (not strictly necessary, but it may contribute to a more secure Nextcloud server setup).

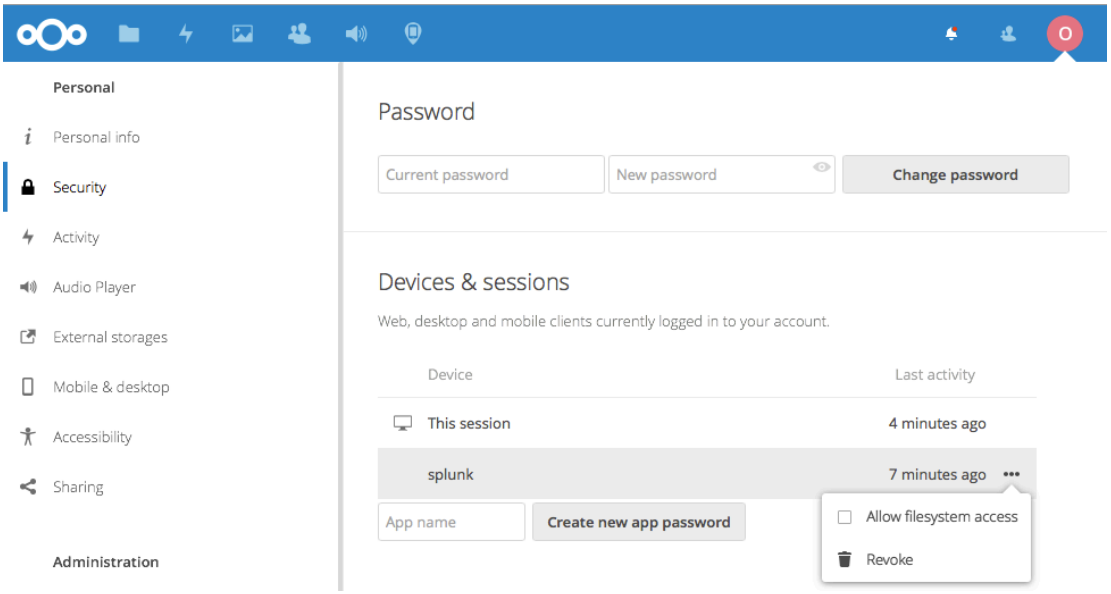

Splunk installation

Hardware requirements:

Splunk requires an x86 64-bit chip architecture and must be installed on a host running either Linux or Windows. Looking at the recommended hardware requirements for a Splunk server may be a bit shocking, but those requirements are meant to support an enterprise system ingesting 100's of GB of data on a daily basis. I have installed Splunk in VM's with no more than 2 GB of RAM and 2 CPU cores. I have also run Splunk on an HP N40L Microserver, though that did turn out to be a bit slow. In most cases, you also won't need a lot of disk capacity either. My home-based Nextcloud server services 5 active users. In one year, the nextcloud.log file has grown to 350 MB. A Splunk server compresses the data it stores and typically compresses the data by 50%. The software for a newly installed Splunk server (other than the OS) uses less than 1 GB of disk storage. For a small Nextcloud installation and a Splunk server which is only monitoring Nextcloud, you could get by with 10 GB. As they say, YMMV.

Software requirements:

The Splunk installation package is completely self-contained, other than the operating system there are no software prerequisites to make it work.

Download and install Splunk Enterprise or Splunk Free (there is no difference between the two download packages) from https://www.splunk.com/download. You will need to register before you can download Splunk. The Splunk App for Nextcloud is currently not supported on the product known as Splunk Cloud and Splunk Light doesn't support third-party apps at all.

Will Splunk be installed on the same server on which Nextcloud is running? If not, you also need to download the "Universal Forwarder". It must be installed on the same host running Nextcloud.

If you will be installing Splunk as root, read on, if not, follow the instructions here:

https://docs.splunk.com/Documentation/Splunk/latest/Installation/RunSplunkasadifferentornon-rootuser

To install Splunk in its default location (/opt/splunk) on Linux, run either of these commands (depending on which Linux distribution you use) where splunk_package_name is the name of the downloaded file:

rpm -i splunk_package_name.rpm

dpkg -i splunk_package_name.deb

For more detailed information on installing Splunk, visit this page:

http://docs.splunk.com/Documentation/Splunk/latest/Installation/InstallonLinux

You can of course install the tarball instead if you wish. Once installed, start Splunk. Assuming Splunk was installed in the default location on a Linux host, this means executing the following command:

```
sudo /opt/splunk/bin/splunk start
```

Read every single word of the license agreement. Or press the <ESC> key and then the "Y" key to agree. During this first start, you will be asked to "Create credentials for the administrator account". This is the password you will be using to log in to the Splunk server.

If you want Splunk to start automatically when its host is started, run this command:

```
sudo /opt/splunk/bin/splunk enable boot-start
```

Now log into Splunk using your browser. The default port for the Splunk UI is 8000, so if your host is called "mysplunkserver" you will access http://mysplunkserver:8000. You were asked for an

administrator username and for a password for the administrator when Splunk was started for the first time which you have hopefully remembered, use that to log in.

Encrypted Web UI (Optional)

You have logged in using http, so any data exchange between the Splunk server and your web browser is unencrypted. Read the Splunk documentation here for the steps to enable https:

https://docs.splunk.com/Documentation/Splunk/latest/Security/TurnonbasicencryptionwithSplunkWeb

---

Splunk App for Nextcloud installation and configuration

Once logged into Splunk:

- Click on "+ Find More Apps" on the left side of the UI. The "Browse More Apps" page loads.
- In the "Find apps by keyword, technology…" field, type "nextcloud".
- Click the green "Install" button for both the "Splunk App for Nextcloud" and the Splunk Add-on for Nextcloud". The user id and password you are requested to enter are the credentials you created when you first registered on the Splunk.com website, not the user id and password used to login to your Splunk server.

If the host on which you've installed Splunk does not have internet access, you must download the app and add-on from "Splunkbase" (the Splunk Appstore) using your choice of browser from these locations:

https://splunkbase.splunk.com/app/3397/

https://splunkbase.splunk.com/app/3398/

On the main page of the Splunk server, between the Splunk logo and the "Search & Reporting" button, click on the gear icon. On the page which now loads, click on the grey "Install app from file" button, click "Choose file", choose one of the files downloaded and click "Upload". Do this for both of the files downloaded from Splunkbase.

Whichever way the App and Add-on were downloaded and installed, once installed, the main dashboard in Splunk should look similar to this.

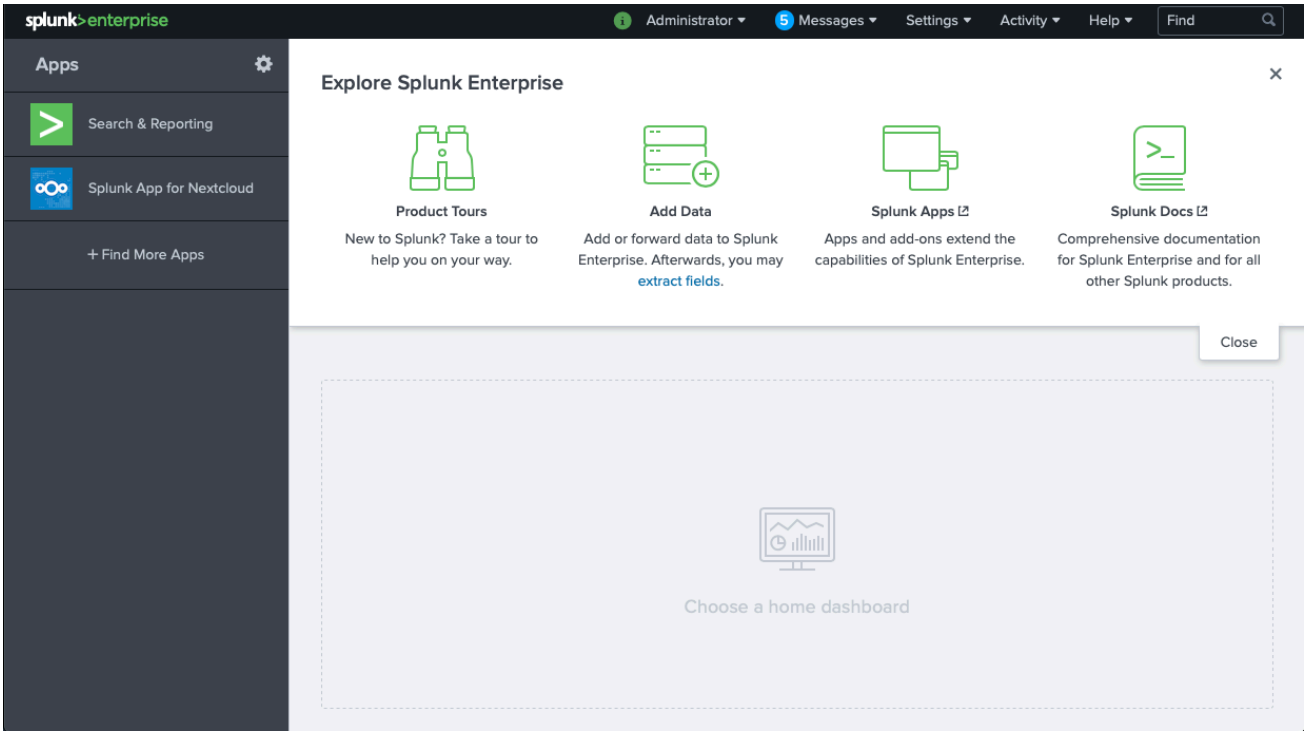

If Splunk has not been restarted after the app installation, the Nextcloud logo won't be shown, but that is purely cosmetics. To use the app, you click on it. But as there is no ingestion of Nextcloud data yet, none of the dashboards will light up with data.

Go back to the command line of the host on which Splunk is installed and run the following commands:

```
sudo -i
mkdir /opt/splunk/etc/apps/TA-nextcloud/local
cp /opt/splunk/etc/apps/TA-nextcloud/default/TA-nextcloud.conf.sample
/opt/splunk/etc/apps/TA-nextcloud/local/TA-nextcloud.conf
nano /opt/splunk/etc/apps/TA-nextcloud/local/TA-nextcloud.conf
```

On a Splunk server or forwarder, files in the 'default' directory contains, well, default information. It may be overwritten when upgrading the Nextcloud Add-on, whereas files located in the 'local' directory are not touched. This is why you are copying the file to the local directory and modifying it there instead of modifying the one in the default directory.

You can obviously use another editor than nano. The TA-nextcloud.conf file loaded in the editor should look something like this:

```
# USER is a Nextcloud user ID that is a member of the "admin" group.
USER=admin
PASSWORD=ICQDW-MHVXI-IBFBL-SUOMK
# NCSERVER should be set to the full base URL without the protocol.
# It would usually consist just of a hostname, but may also include a path.
# examples could be: “demo.nextcloud.com” or “www.example.com/nextcloud”
NCSERVER=nextcloud.example.com
```

```
# PROTOCOL should be either “http” or “https”.
# Do not add ://, just the protocol!
PROTOCOL=https
# If your Nextcloud server is accessed via HTTPS, the scripts may need
# to know where your local certificate store is.
# Uncomment the CA line below if this is the case.
# On my ubuntu 16.04 server that location is as shown below. Other
# OS’ may have them in a different location
CA="--capath /etc/ssl/certs/"
```

Add relevant values for USER, PASSWORD (you hopefully made a note of these as instructed above in the section about what to do on your Nextcloud server) NCSERVER and PROTOCOL. If PROTOCOL is set to “https” you will most likely have to uncomment (remove the # character) from the line starting with “CA=”. Now save the file.

Log into the Splunk UI by entering the URL in a browser.

- Click on “Settings” on the upper black menu bar.
- In the “Data” section, click “Data inputs”
- Click “Scripts”
- Click on “$SPLUNK_HOME/etc/apps/TA-nextcloud/bin/nextcloud-info.sh”
    - Click the “More settings” checkbox in the lower left corner
    - In the “Host field value” field enter the host name of the server hosting Nextcloud. Don’t enter the fully qualified domain name (FQDN), so if for example the server is addressed as “cloud.example.org” you should enter “cloud”. The host name entered should be the primary host name, not an alias.
    - Click “Save”
- Follow the three steps above for “$SPLUNK_HOME/etc/apps/TA-nextcloud/bin/nextcloud-status.sh” as well.
- Click “Enable” for “$SPLUNK_HOME/etc/apps/TA-nextcloud/bin/nextcloud-info.sh” and for “$SPLUNK_HOME/etc/apps/TA-nextcloud/bin/nextcloud-status.sh”.

Universal Forwarder installation and configuration

Log into the Splunk server UI by entering the URL in a browser.

- Click on “Settings” on the upper black menu bar.
- In the “Data” section, click “Forwarding and receiving”
- In the “Receive data” section, click “Configure receiving”

- If "9997" isn't already listed under the "Listen on this port" heading, click the green button "New Receiving Port"
- In the "Listen on this port" field, enter "9997" (without the quotes, just the number) and click the green "Save" button.

Now download and install the Universal Forwarder package on the server where the nextcloud.log and audit.log files are located. Choose the relevant package for your Linux distribution.

Download the Universal Forwarder here:

https://www.splunk.com/en_us/download/universal-forwarder.html

Once you've chosen the relevant installation package, it will start downloading, but the next page has an option to "Download via Command Line (wget)". Click on that link, copy the entire command from the web browser and then paste it into a terminal window on your Nextcloud server. Once the package has been downloaded, run either of the following commands, depending on your Linux distribution:

rpm -i <splunk_forwarder_package.rpm>

dpkg -i <splunk_forwarder_package.deb>

As with Splunk itself, you can of course install the tarball instead if you wish. If your Nextcloud server is hosted on a Raspberry Pi, a tarball is your only option (choose the 'ARMv6' package). Once installed, start the Splunk Universal Forwarder. Assuming Splunk was installed in the default location on a Linux host, this means executing the following command:

```
sudo /opt/splunkforwarder/bin/splunk start
```

Again, read every single word of the license agreement. During this first start, you will be asked to "Create credentials for the administrator account". Later when you change certain parameters of the Universal Forwarder you will be asked for this password, so remember it together with the username of the Universal Forwarder which is "admin". Now run the following commands where <splunk server> is either the name of the host (if you can ping the name, the Universal Forwarder should also be able to send data to the host) where Splunk is installed or its IP address. The <path to> is the directory where the Nextcloud log files are stored. The default is "/var/www/nextcloud/data", check if that directory contains a nextcloud.log file and a (depending on your Nextcloud server version) audit.log file:

```
sudo /opt/splunkforwarder/bin/splunk enable boot-start
sudo /opt/splunkforwarder/bin/splunk add forward-server <splunk server>:9997
sudo /opt/splunkforwarder/bin/splunk add monitor /<path to>/nextcloud.log -
sourcetype nextcloud-log
sudo /opt/splunkforwarder/bin/splunk add monitor /<path to>/audit.log -
sourcetype nextcloud-log
```

```
sudo /opt/splunkforwarder/bin/splunk restart
```

If you have followed the above instructions correctly and I haven't forgotten any necessary steps, the dashboards of the Splunk App for Nextcloud should start populating within a few minutes (assuming the nextcloud.log or audit.log file has already recorded Nextcloud activity).

Upgrading the Splunk App for Nextcloud

Assuming your Splunk server can access the internet, it should show on the main screen when the app is upgradable. Clicking on "Update" brings you to a screen from where the app can automatically be upgraded. The user id and password you need to enter is the one you registered on splunk.com when you first registered to be able to download the Splunk software.

An update for the add-on is only visible if you click on the gear icon visible on the left side of the main page of the Splunk server, below the Splunk logo and above the "Search & Reporting" button.

Troubleshooting

You can (obviously) use Splunk to figure out what is going on with Splunk. To get all sorts of information about how Splunk is performing, click on the "Settings" menu option and then click on the "Monitoring Console" icon.

Within the "Splunk App for Nextcloud", you can also choose "Other > App debugging" to view a dashboard that can help with troubleshooting if data is not showing up in the app dashboards.

Curl exit codes

The scripts used to ingest data from the Nextcloud monitoring endpoints use the "curl" command and in case of a problem, curl may return with an error message

If, for example, the scripts use HTTPS to access your Nextcloud server, and the scripts can't find the local certificate store, you would see this kind of message:

ERROR ExecProcessor - message from "/opt/splunk/etc/apps/TA-nextcloud/bin/nextcloud-status.sh" curl: (77) Problem with the SSL CA cert (path? access rights?)

Splunk Cloud Installation

Only Splunk Cloud vetted apps can be installed on Splunk Cloud instances. The Splunk App for Nextcloud has not been vetted for Splunk Cloud. To get the app vetted for Splunk Cloud would require that a customer make a request via the service portal for the vetting to take place.

Known issues

Some events from "nextcloud.log" are on not decoded properly, this is on purpose. The log file input is set to only ingest the first 10,000 characters. Some events are much larger than that. On my own Nextcloud systems I have seen the Nextcloud "appstoreFetcher" time out and create an event larger than 1 MB. As configured, only the first 10,000 characters of these events is ingested, and the event is therefore incomplete. The configuration can be changed to ingest the complete event, but you probably don't want these >1 MB events in Splunk.

Support

Although I work for Splunk, the Splunk Add-on and App for Nextcloud modules are not supported by Splunk, and support from myself is limited. I do welcome suggestions.

If you do have any issues, go here for the Add-on:

https://answers.splunk.com/app/questions/3397.html

And here for the App:

https://answers.splunk.com/app/questions/3398.html

Advanced Topics

The following topics are not necessary for the standard use of the Splunk App for Nextcloud. If you you do attempt to implement these topics, you might break the Splunk App for Nextcloud. As usual in IT, if you do, you get to keep both pieces. If you know what you are doing, any changes made on the basis of following these instructions are repairable.

Splunk indexes

The repository for data in Splunk is called an "index" (not to be confused with "indexers"). By default, all ingested data is stored in the default "index", its name is "main".

It is possible to create and store data in separate indexes. Reasons why you might want to do so are:

- Access management: By storing Nextcloud events in custom indexes you can control which users can access the Nextcloud data.
- Data retention: Different indexes can have different data retention policies.
- Data hygiene: for whatever reason you don't wish to mix different types of data in a particular index.

If you wish to store the Nextcloud events in a custom index you must:

- Create a custom index
- Configure event collection so events are sent to the custom index
- Configure Splunk so the Splunk App for Nextcloud dashboards can find the data in the custom index.

Access this link for information on how to create a custom index:
https://docs.splunk.com/Documentation/Splunk/latest/Indexer/Setupmultipleindexes

Just before the "Universal Forwarder Installation and Configuration" section of this document it is described how to configure the scripts that also pull in information about your Nextcloud server. To configure this information to be stored in a custom index, for each script select "More settings" and choose the destination index from the drop-down box.

In the "Universal Forwarder Installation and Configuration" section of this document are listed some commands to execute to configure the Universal Forwarder. The two "add monitor" commands (Example: `sudo /opt/splunkforwarder/bin/splunk add monitor /<path to>/audit.log -sourcetype nextcloud-log`) need an additional parameter to specify which index the events should be sent to. Add the "-index <newindex>" parameters to the two "add monitor" lines, replacing <newindex> with the name of the custom index you have created. You may wish to consult this link if you want to understand more about this and other options for configuring the Universal Forwarder to monitor logs:

https://docs.splunk.com/Documentation/Splunk/latest/Data/MonitorfilesanddirectoriesusingtheCLI

To access data stored in a custom index, a search would usually start with the "index=index_name" SPL command. Any data stored in a "default" index can be found without specifying which index it is stored in. For those familiar with the concept, this is similar to having a search path when accessing commands from the command line in Unix or Windows (if you don't know what I'm talking about, ignore this sentence). Splunk user roles can be configured to have multiple "default" indexes. The data in any index configured to be a default index will be searched without having to specify the "index=" command. To configure this:

- Navigate to "Settings > Access controls > Roles".
- Click the name of the role that should have multiple "default" indexes. Unless you wish to restrict who can access the Nextcloud data, choose the "User" role.
- Click on "Indexes"
- Put a check mark in the box for the index containing Nextcloud data in the column for default indexes. You can have multiple default indexes so don't remove the check mark from the "main" index.

See these screenshots for an example. Note: the index containing the Nextcloud data does not have to be named “nextcloud”.

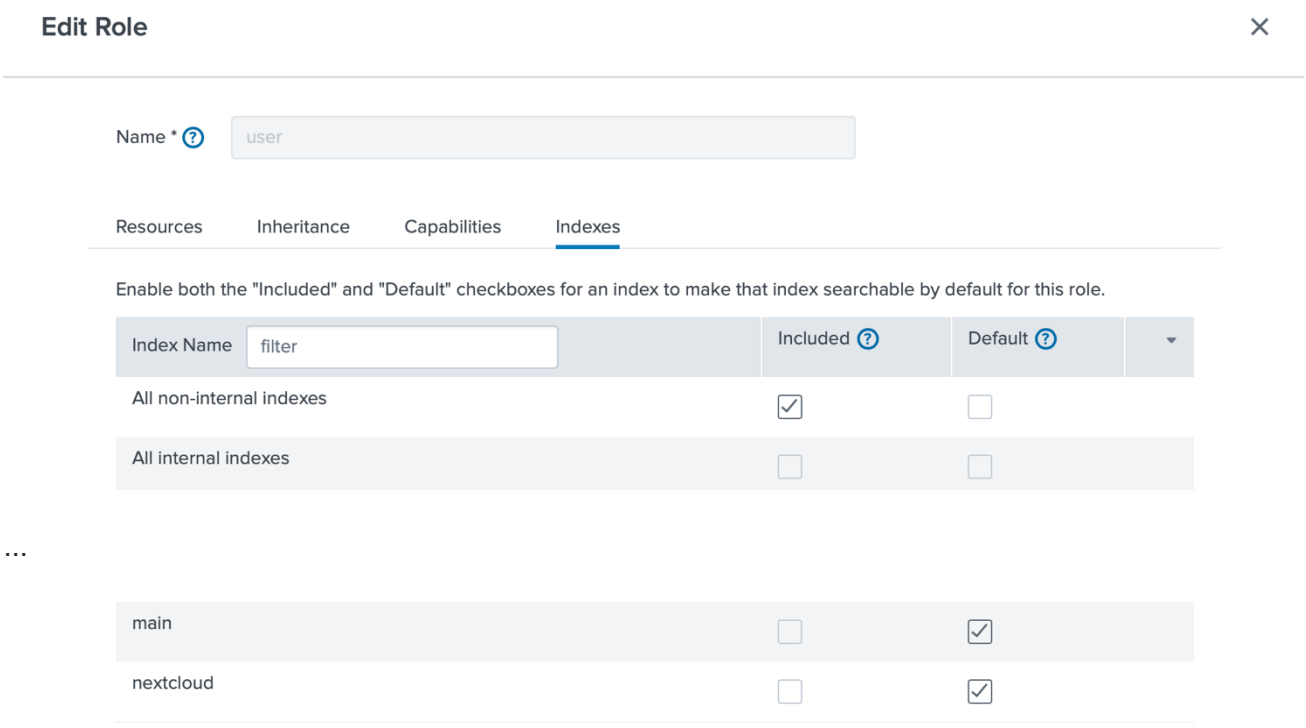

...

Customising dashboards

If the default time period in the dashboards doesn’t suit you (or the layout of the dashboard panels, or something else), change it. However, if you make any changes to a dashboard and the same dashboard is changed in a later version of the app, your changes take precedence and you will unfortunately not see the amazing changes made to the dashboard by me even though you update the app (did that make sense?). Basically, local changes have precedence over the default. One option is to clone the existing dashboards and then add the clone to the menu options. See this link for more information:
https://docs.splunk.com/Documentation/Splunk/latest/Viz/DashboardCloneHome

http://dev.splunk.com/view/webframework-developapps/SP-CAAAEP9

Changing the default time period for a dashboard

Follow this procedure on all dashboards with a time picker. Click the “Edit” button on the dashboard you want to change.

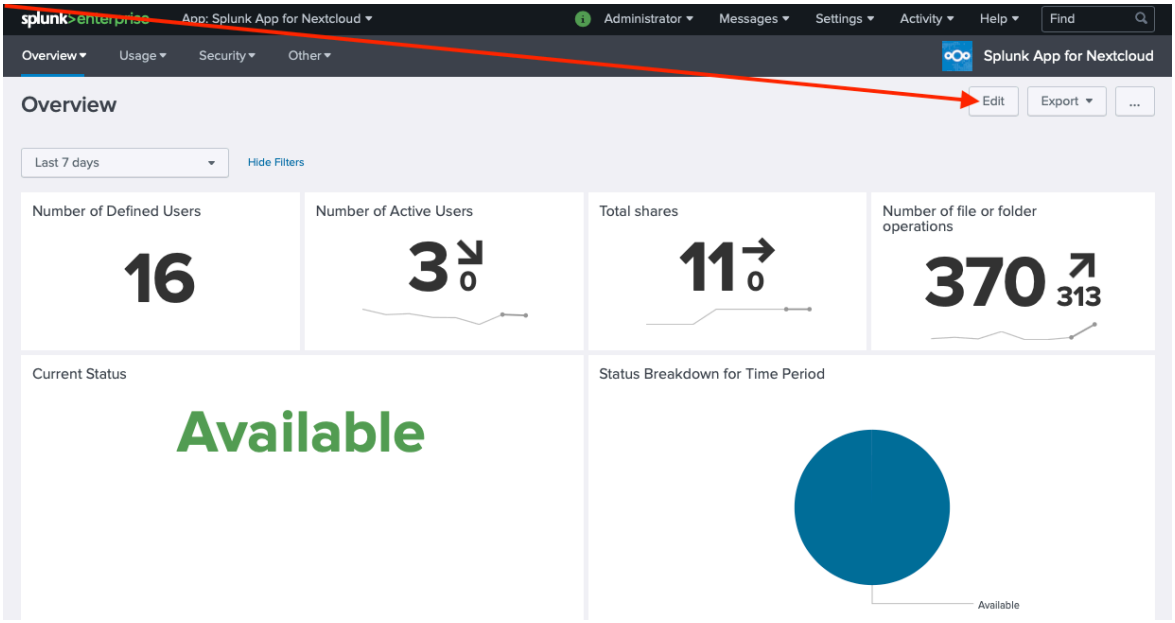

The dashboard will change into "edit" mode, and you should now click the pencil.

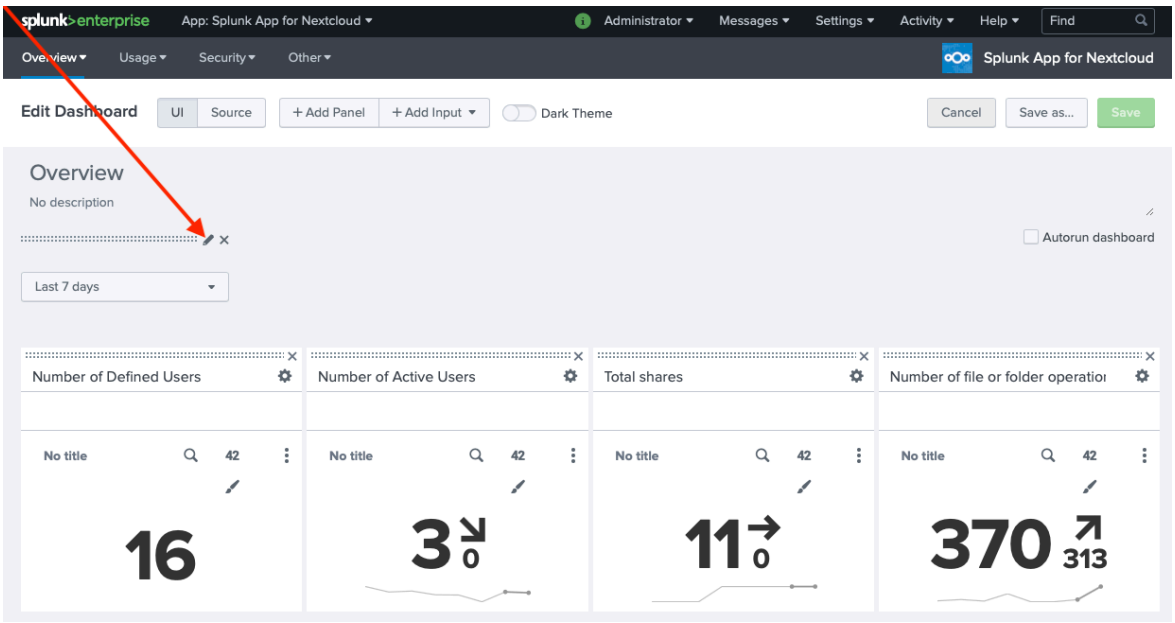

In the menu which appears, choose the time period you wish to use as the default, click "Apply" and then click on "Save" in the upper right corner.

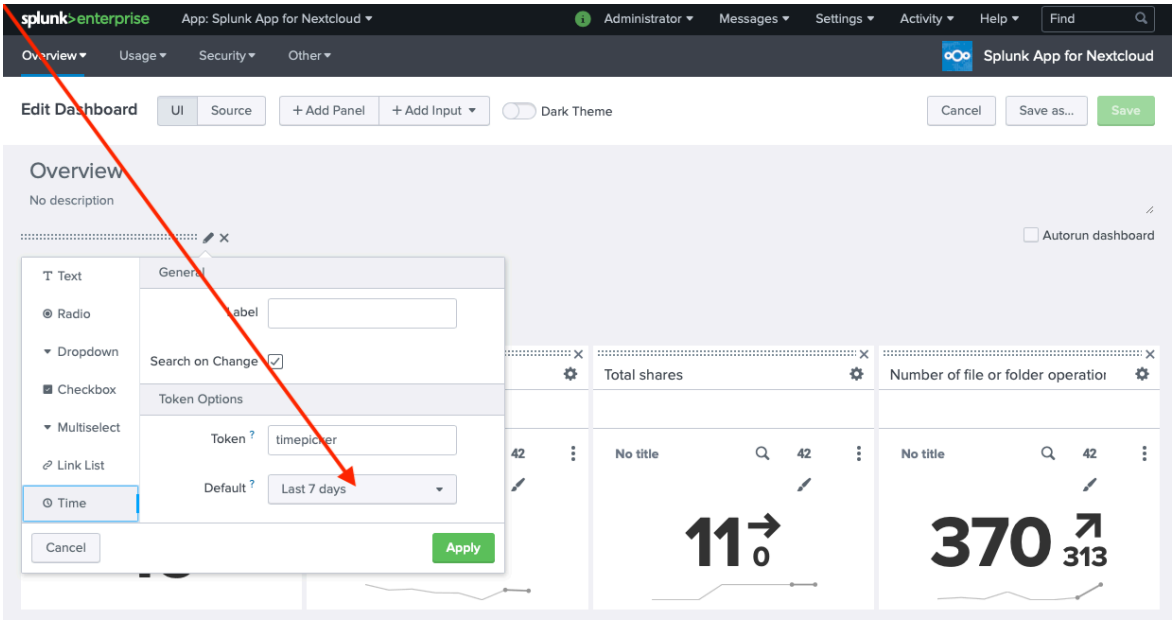

Changing the centre of location maps

If you don't live in Europe, you may want to change the latitude and longitude of the map centre. As shown in the section about changing the default time period, click edit on the dashboard you wish to change. Click on the pen within the panel holding the map. Change the latitude and longitude and maybe even the zoom factor to whatever suits your location. For the US for example, relevant values might be latitude: 39.828175, longitude: -98.5795, zoom: 4. After having changed the values, click the little x, and then click the green "Save" button in the upper right corner.

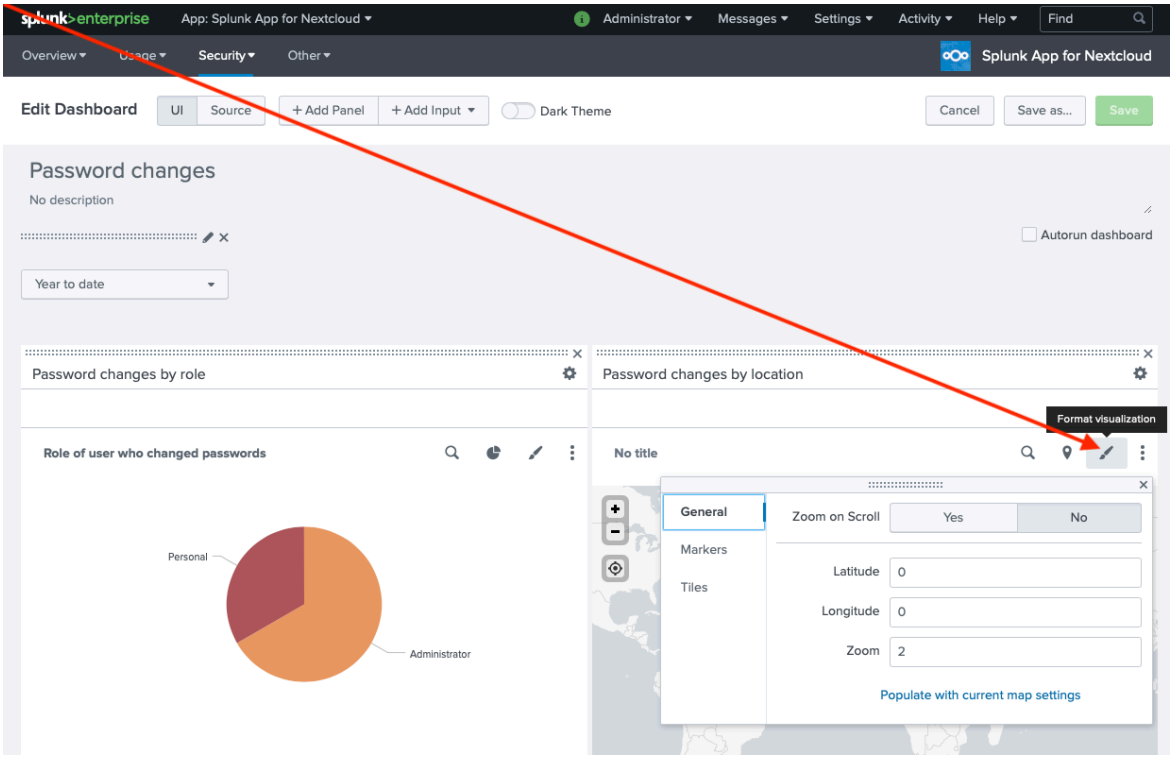

Changing the map

The default maps are built in to Splunk and they don't allow you to zoom terribly far in. I've found that the location data based on the IP addresses are not that accurate anyway, but if you do want to be able to zoom further in, you can change the maps used by Splunk to Openstreetmap. This does require that the Splunk server has an internet connection.

Do as above, click the "Edit" button, click on the pen for the map panel. In the menu, click on "Tiles", then click on the text "Populate from preset configuration". Chose "Open Street Map", click the little x, and then click the green "Save" button in the upper right corner.

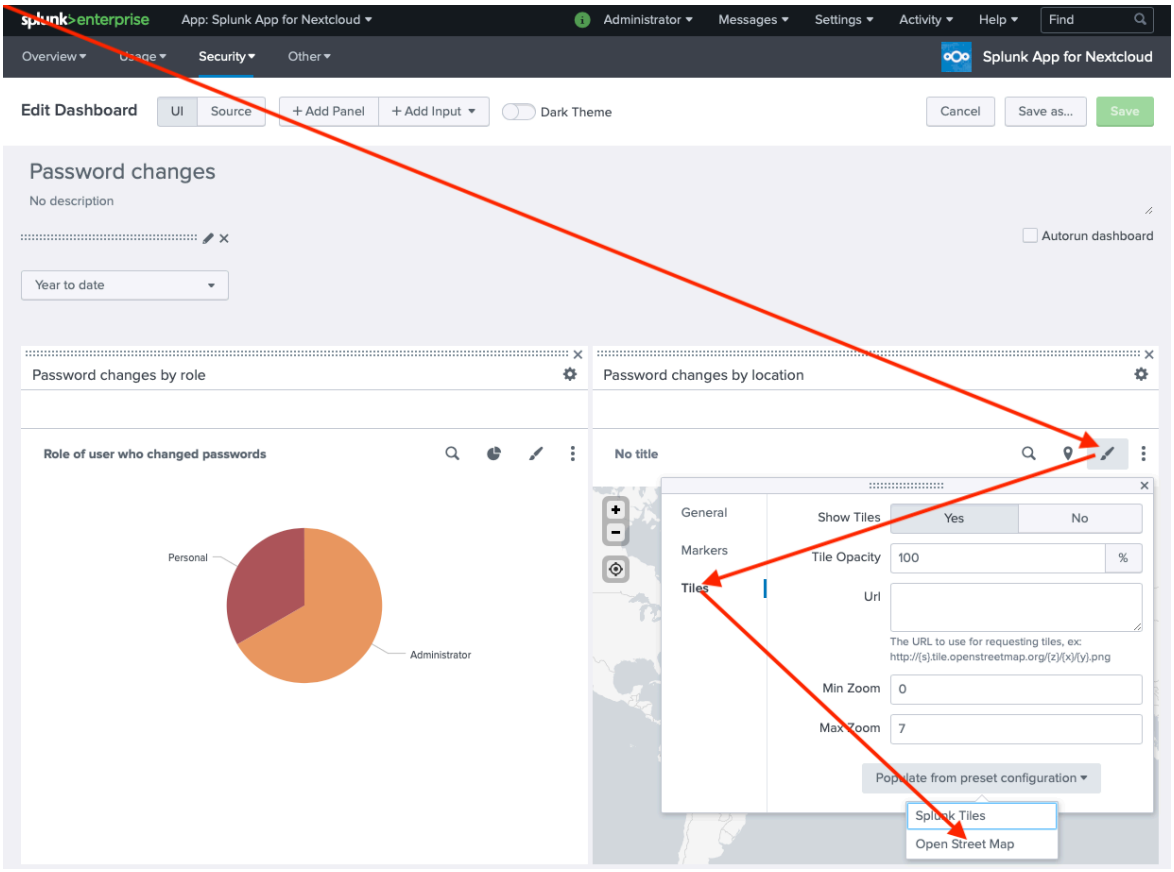

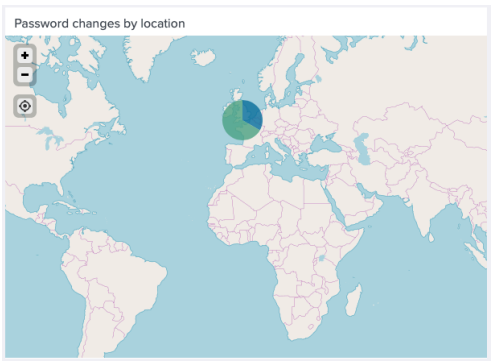

This is what the map should look like afterwards.