


Set up EclecticIQ Platform as a virtual appliance

Virtual appliance guide for system administrators

Last generated: May 26, 2017



©2017 EclecticIQ

All rights reserved. No part of this document may be reproduced in any form or by any electronic or mechanical means, including information storage and retrieval systems, without written permission from the author, except in the case of a reviewer, who may quote brief passages embodied in critical articles or in a review.

Trademarked names appear throughout this book. Rather than use a trademark symbol with every occurrence of a trademarked name, names are used in an editorial fashion, with no intention of infringement of the respective owner's trademark.

The information in this book is distributed on an "as is" basis, without warranty. Although every precaution has been taken in the preparation of this work, neither the author nor the publisher shall have any liability to any person or entity with respect to any loss or damage caused or alleged to be caused directly or indirectly by the information contained in this book.

©2017 by EclecticIQ BV. All rights reserved.
Last generated on May 26, 2017

Table of contents

Table of contents	2
Virtual appliance configuration guide	3
Scope	3
Goal	3
Audience	3
Feedback	3
Before you start	5
Hardware requirements	5
Single box	5
Software requirements	6
VM OS login credentials	7
EclecticIQ Platform login credentials	7
Install the platform	8
Install a VMWare Player VM image	8
Install Neo4j	9
Pre-installation checks	9
Install Neo4j	10
Rename the configuration files	10
Check file ownership	11
Reload Supervisor	12
Test Neo4j in the UI	13
Launch the platform	14
Start the VM	14
Get the VM IP address	14
Go to the platform	15
Upgrade the platform	16
Exit the platform	17
Back up your data	17
Shut down the platform	17
Check the prerequisites	19
Remove deprecated packages	19
Create the YUM repository configuration	20
Run the YUM install	21
Check the configuration	23
Check third-party configurations	23
Install poppler-utils	23
Migrate the database	24
Reindex Elasticsearch	24
Migrate the graph database	25
Check component versions	25
Run the fixtures	28
Run a final check	29
Check core processes and services	29
Check search indexing and graph	29
Check availability	31
Reload the Supervisor configuration	31
Access the platform	32

Virtual appliance configuration guide

This document guides you through the setup and the configuration of EclecticIQ Platform when you choose to install the product as a virtual appliance running in a virtual machine client.

Scope

This document guides you through the steps you need to carry out to complete the following tasks:

- Install — that is, decompress and save to a target location — the virtual machine containing the platform.
- Launch the platform..
- Upgrade the platform to a newer release.

Goal

After completing these tasks, you'll have achieved the following goals:

- The EclecticIQ Platform is set up and configured in the target system as a **virtual appliance** (https://en.wikipedia.org/wiki/virtual_appliance).
- The platform is ready for use as a virtual appliance inside a virtual machine client/player.

Audience

This document targets the following audience:

- DevOps
- System administrators

Feedback

No one reads manuals, ever. We know.

Yet, we strive to give you clear, concise, and complete documentation that helps you get stuff done neatly.

We are committed to crafting good documentation, because life is too short for bad doc.

We appreciate your comments, and we'd love to hear from you: if you have questions or suggestions, drop us a line and share your thoughts with us!

👉 The Product Team

Before you start

Review these system requirements before installing the platform.

This section covers the necessary hardware and software requirements to set up EclecticIQ Platform as a virtual appliance.

Hardware requirements

Hardware requirements for EclecticIQ Platform can vary depending on the target environment you plan to install the platform to. Therefore, the requirements outlined in this section are general guidelines that work in most cases, but they are not tailored to any specific situation.

Single box

Hardware requirement guidelines for EclecticIQ Platform and related dependencies installation on one target machine.

HW area	Minimum	Recommended	Notes
Environment	-	VM/Virtual appliance	
CPUs	4	8	Core count includes HT
CPU speed	2.5 GHz	2.5 GHz or faster	
Memory	16 GB	at least 32 GB	16 GB is unsuitable for production. A production environment should feature at least 32 GB memory. Consider expanding it to 64 GB when dealing with, for example, large data corpora ingestion or data-intensive graph visualizations. Monitor system memory usage to determine if your system may need more memory to operate smoothly.
Storage	SATA, 100 IOPS	SSD, 200 IOPS	Local attached storage is preferable to SAN or NAS; platform operations are write-intensive. Recommended IOPS range: 200-500

HW area	Minimum	Recommended	Notes
Drives	5	10	10 drives to set up 5 sets of mirrored drives (RAID 1)
Drive sizes (GB)	10, 10, 25, 50, 200	20, 20, 50, 75, 300	Each platform database should be allocated to a dedicated drive for data storage
Drive allocation (GB)	10	20	Root (EclecticIQ Platform + Redis)
	10	20	Log data storage
	25	50	Neo4j, graph database
	50	75	Elasticsearch, searching and indexing
	200	300	PostgreSQL, main data storage
Network	2 network interfaces	2 network interfaces	1 interface for production, the other for system management
Install size	~240 GB	~240 GB	Full install, based on VM image size

Software requirements

- **VMWare Player** (<https://www.vmware.com/products/player>).
VirtualBox (<https://www.virtualbox.org/>) is currently not supported.
- Internet connectivity.
- A web browser with JavaScript enabled.
Recommended: **Google Chrome** (<https://www.google.com/chrome/>).
- DNS name of the host you are going to use to access the platform.
Example: `platform.host`
- SSL certificate and key for the web server.
- EclecticIQ Platform login credentials.



Warning: The EclecticIQ Platform VM image does not include Neo4j. You need to install Neo4j manually, as described in [Install Neo4j](../_vm-install/vm_neo4j_install.html).

VM OS login credentials

SSH default login credentials for the VM OS	
user	packer
password	Packer123!

EclecticIQ Platform login credentials

EclecticIQ Platform default login credentials	
user	admin
password	EclecticIQ2015#

Install the platform

Download a VM image, import it into your VM client, install Neo4j, start working on cyber threat analysis.

EclecticIQ Platform is available as a virtual appliance in a downloadable VM image.

- Contact us to obtain a download link to a VM image. We offer VM images for **VMWare Player** (<https://www.vmware.com/products/player>). **VirtualBox** (<https://www.virtualbox.org/>) is currently not supported.
- Save the downloaded archive on your local machine and decompress its content.

Install a VMWare Player VM image

- Launch VMWare Workstation Player.
- The default VMWare Workstation Player start location is **Home**.
- On the right **Home** pane, select **Open a Virtual Machine**.
- On the **Open Virtual Machine** dialog, browse to the location where you saved the decompressed VM files.
- Select the appropriate .ova image file, and then click **Open**.
- The new VM becomes available in the left-hand pane in the player.
- On the right-hand pane, click **Edit virtual machine settings**.
- On the **Virtual Machine Settings** dialog, **Hardware** tab, click the **Add** button.
- On the **Add Hardware Wizard, Hardware Type** dialog, select **Network Adapter**, and then click **Next**.
- On the **Add Hardware Wizard, Network Adapter Type** dialog, select the **Host-only: A private network shared with the host** option, and then click **Finish**.
- Click **OK**.
- On the right-hand pane, click **Power On**.

Install Neo4j

Install Neo4j before starting working with EclecticIQ Platform.

Neo4j is not included in the EclecticIQ Platform VM image. To install Neo4j manually, do the following:

- Log into the platform OS.
- To perform several tasks in the procedure, you may need root-level access rights.
To obtain administration rights, run the following command(s):

```
$ sudo su -
```

- Go to the root directory:

```
$ cd /
```

Pre-installation checks

- Check if Neo4j is installed:

```
$ yum info neo4j
```

- Verify that no Neo4j-related services or processes are running.
If Neo4j is not installed, `neo4j-console` should either not be included in the returned list of services, or it should not be running if it is included:

```
$ supervisorctl status
```

- Verify that the `neo4j` user and the `neo4j` group own the `/media/neo4j/` directory, and therefore the `graph.db` graph database file it stores:

```
$ ls -l /media/neo4j
```

YUM checks if Neo4j exists on the system, and it looks for any available updates.

Install Neo4j

- Install Neo4j 2.3.1 Community:

```
$ yum install neo4j-2.3.1
```

Rename the configuration files

- After completing the Neo4j installation, check the original Neo4j configuration files:

```
$ cd /etc/neo4j
$ ls -ltr
```

- The returned list of files should include the following ones:

```
# configuration files installed by Neo4j
-rw-r--r--. 1 root  root  2284 Nov 10  2015 neo4j-wrapper.conf
-rw-r--r--. 1 root  root  3255 Nov 10  2015 neo4j-server.properties
-rw-r--r--. 1 root  root  1796 Nov 10  2015 neo4j.properties

# Platform backup configuration files for Neo4j to work with the platform
lrwxrwxrwx. 1 root  root    49 May 12 10:07 neo4j.properties.rpmshave ->
/opt/eclecticiq/etc-extras/neo4j/neo4j.properties
lrwxrwxrwx. 1 root  root    56 May 12 10:07 neo4j-server.properties.rpmshave ->
/opt/eclecticiq/etc-extras/neo4j/neo4j-server.properties
lrwxrwxrwx. 1 root  root    51 May 12 10:07 neo4j-wrapper.conf.rpmshave ->
/opt/eclecticiq/etc-extras/neo4j/neo4j-wrapper.conf
```

- In the /etc/neo4j directory remove the following configuration files installed by Neo4j:

```
$ rm neo4j-wrapper.conf

# Returns the following response. Type 'yes' to confirm deletion.
> rm: remove regular file 'neo4j-wrapper.conf'? yes
```

```
$ rm neo4j-server.properties

# Returns the following response. Type 'yes' to confirm deletion.
> rm: remove regular file 'neo4j-server.properties'? yes
```

```
$ rm neo4j.properties

# Returns the following response. Type 'yes' to confirm deletion.
> rm: remove regular file 'neo4j.properties'? yes
```

- Rename the Neo4j configuration files ending in `.rpmsave` by removing the `.rpmsave` suffix:

```
$ mv neo4j.properties.rpmsave neo4j.properties
$ mv neo4j-server.properties.rpmsave neo4j-server.properties
$ mv neo4j-wrapper.conf.rpmsave neo4j-wrapper.conf
```

- List the content of the `/etc/neo4j` directory to verify that renaming completed successfully:

```
$ ls -l
```

- In the response, the renamed files should still be symlinked to their target files, that is, the Neo4j configuration files provided with the platform:

```
lrwxrwxrwx. 1 root  root    49 May 12 10:07 neo4j.properties ->
/opt/eclecticiq/etc-extras/neo4j/neo4j.properties
lrwxrwxrwx. 1 root  root    56 May 12 10:07 neo4j-server.properties ->
/opt/eclecticiq/etc-extras/neo4j/neo4j-server.properties
lrwxrwxrwx. 1 root  root    51 May 12 10:07 neo4j-wrapper.conf ->
/opt/eclecticiq/etc-extras/neo4j/neo4j-wrapper.conf
```

Check file ownership

- List the content of the `/etc/neo4j` directory to verify that the Neo4j configuration file owners are `root` (user) and `root` (group):

```
$ ls -l /etc/neo4j
```

- In the response, the Neo4j configuration file owners should be `root root`:

```
lrwxrwxrwx. 1 root  root    49 May 12 10:07 neo4j.properties ->
/opt/eclecticiq/etc-extras/neo4j/neo4j.properties
lrwxrwxrwx. 1 root  root    56 May 12 10:07 neo4j-server.properties ->
/opt/eclecticiq/etc-extras/neo4j/neo4j-server.properties
lrwxrwxrwx. 1 root  root    51 May 12 10:07 neo4j-wrapper.conf ->
/opt/eclecticiq/etc-extras/neo4j/neo4j-wrapper.conf
```

- List the content of the `/media/neo4j` directory to verify that the directory owners are `neo4j` (user) and `neo4j` (group):

```
$ ls -l /media/neo4j
```

- In the response, the directory owners should be `neo4j neo4j`:

```
drwxr-xr-x. 4 neo4j neo4j  4096 May 12 10:08 graph.db
```

- List the content of the `/usr/share/neo4j/` directory to verify that the `/data` subdirectory owners are `neo4j` (user) and `neo4j` (group):

```
$ ls -l /usr/share/neo4j/
```

- In the response, the `/data` subdirectory owners should be `neo4j neo4j`:

```
drwxr-xr-x. 2 neo4j neo4j   16 May 12 10:03 data
```

Reload Supervisor

At this point Neo4j is installed, and it should use the correct configuration files to work with the platform. When you modify or update the Supervisor configuration, you need to run `$ supervisorctl reload` to update it and to reload the up-to-date configuration into the platform.

Alternatively, run first `$ supervisorctl reread` to reload the changed configuration, and then `$ supervisorctl update` to restart the managed applications whose configurations have changed.

- Restart Supervisor to load the changed configuration for Neo4j:

```
$ supervisorctl reload
```

- Verify that `neo4j-console` is running:

```
$ supervisorctl status
```

- The response should include the service, and it should flag it as running:

```
neo4j-console      RUNNING      pid 30726,      uptime 0:05:23
```

Test Neo4j in the UI

As a final check, sign into the web-based UI to make sure that the graph is up and running, and that it is working correctly. you can add entities to graph.

- Check the system health.
The status of the following processes should be green:
 - *graph-ingestion*
 - *intel-ingestion*
 - *neo4j-console*
- Add some entities to the graph to make sure that the graph loads them and displays them correctly.
- Create one or two new test entities using the entity editor, and then load them onto the graph to make sure it works as expected.

Launch the platform

Start the VM to make the platform available, and then access the platform through a web browser.

Start the VM



To access the VM, you may need to enter valid login credentials. If you do not have these details, contact us.

- Launch VirtualBox or VMWare Player, and then start/play the VM.
- If you are prompted for login credentials at startup, enter the provided user name and password.
- The VM should be up and running.

Get the VM IP address

By default, our VM images run **CentOS Linux 7 (1511)** (<https://lists.centos.org/pipermail/centos-announce/2015-december/021518.html>).

- Run the following command(s):

```
$ ifconfig
```

- Press **ENTER**.
- Look for the following entry to identify the VM IP address:

```
inet <IP_address>
```

The `inet` IP address is the one you need to use to access the platform.

Example: `inet 10.0.2.148`.

Go to <https://10.0.2.148> to sign in to the platform web-based UI.

Go to the platform

- In your host machine, launch a web browser (recommended: Google Chrome).
- In the address bar, enter the VM IP address.
Example: *https://10.0.2.148*
- The platform login screen is displayed.
- Sign in with the appropriate credentials.



Warning: The browser may display an untrusted connection warning: add it as an exception, and then proceed to the platform.

Upgrade the platform

When a new platform release is available, you can upgrade your existing installation to benefit from the latest features and enhancements.

When you download a new VM image containing a newer platform release than the currently installed one on your system, you can upgrade your platform installation to the latest available public version.

The upgrade procedure requires some housekeeping; once you are done, you can access new features, and you can enjoy the improvements we introduce in the product on a regular basis.



To perform several tasks in the procedure, you may need root-level access rights. To obtain administration rights, run the following command(s):

```
$ sudo su -
```

The upgrade procedure consists of the following steps:

Before the upgrade

- Exit the platform
- Back up your data
- Check the prerequisites

During the upgrade

- Create the YUM repository configuration
- Run the YUM install

After the upgrade

- Check the configuration
- Check third-party configurations
- Migrate the database
- Check component versions
- Run the fixtures
- Reload the Supervisor configuration
- Access the platform

Exit the platform

Sign out of the platform:

- Click the active user profile image on the top-right corner of the screen.
- From the drop-down menu select **Sign out**.
- You are signed out.

Back up your data



Warning: Before proceeding to upgrade the platform or any of its third-party components, always back up your data.

Shut down the platform

To gracefully shut down EclecticIQ Platform, do the following:

- Stop all running processes.
- Shut down the machine the platform is hosted on.

Generic shutdown

You can stop the platform without following any specific procedure. However, if you want to make sure the platform core services and processes gracefully shut down, manually execute the following commands:

```
$ supervisorctl stop all
```

```
$ systemctl stop postgresql-9.5 redis logstash elasticsearch
```

You can then proceed to shut down the virtual or physical machine hosting the platform.

Shutdown before a platform upgrade

If you are shutting down the platform before performing an upgrade, stop platform components in the order described below to make sure no Celery tasks are left over in the queue. This prevents hanging tasks in the queue from interfering with the upgrade procedure.

- Stop platform-api:

```
$ supervisorctl stop platform-api
```

- Stop the Celery beat:

```
$ supervisorctl stop task:beat
```

- Check Celery queues — they should be empty:

```
$ redis-cli llen enrichers  
  
$ redis-cli llen integrations  
  
$ redis-cli llen priority_enrichers  
  
$ redis-cli llen priority_providers  
  
$ redis-cli llen priority_utilities  
  
$ redis-cli llen providers  
  
$ redis-cli llen reindexing  
  
$ redis-cli llen utilities
```

Alternatively, check Redis keys:

```
$ redis-cli keys *
```

If the response *does not include* any keys related to any Celery queues, they are empty and you can continue.

- Stop the remaining Celery workers:

```
$ supervisorctl stop task:*
```

- Stop Supervisor-managed workers:

- kibana
- intel-ingestion
- opentaxii
- search-ingestion
- graph-ingestion
- neo4j

```
$ supervisorctl stop all
```

Alternatively, you can list them explicitly:

```
$ supervisorctl stop kibana intel-ingestion opentaxii search-ingestion graph-ingestion  
neo4j
```

Check the prerequisites

Before installing the new platform release, check the prerequisites, and verify that the following conditions are met:

- The required repositories are already installed on the target system.
- All required third-party components are already installed on the target system, and their versions match the recommended version information provided.
- To perform several tasks in the procedure, you may need root-level access rights.
To obtain administration rights, run the following command(s):

```
$ sudo su -
```

Remove deprecated packages

Before installing the new platform release, make sure you remove any deprecated packages.
To remove deprecated packages, run the following command(s):

```
$ yum remove <package_name>
```

or:

```
$ rpm -e <package_name>
```

In case uninstalling a deprecated package fails because of dependency-related errors, include the **--nodeps option** (<http://www.rpm.org/max-rpm-snapshot/s1-rpm-erase-additional-options.html#s2-rpm-erase-nodeps-option>):

```
$ rpm -e --nodeps <package_name>
```

The following packages are *deprecated*: remove them before upgrading.

Before upgrading to this version...	...remove this RPM package
0.17	eclecticiq-intel-search-0.16.0-1.x86_64
1.11.0	platform-opentaxii

Create the YUM repository configuration

If you are upgrading or performing a fresh install of EclecticIQ Platform using the **YUM** (<http://yum.baseurl.org/>) package manager, you need to define your YUM repository configuration for the platform:

- Start the host system and log into it with the appropriate credentials.
- Create a new file, and name it `/etc/yum.repos.d/eclecticiq-platform.repo`.
- Populate the file with the following content:

```
[eclecticiq-platform]
name=eclecticiq-platform
baseurl=https://downloads.eclecticiq.com/platform
gpgcheck=1
repo_gpgcheck=0
enabled=1
username=<UPDATE_USERNAME>
password=<UPDATE_PASSWORD>
gpgkey=https://downloads.eclecticiq.com/platform/repodata/repomd.xml.key
```



Warning: Make sure you update and define user name and password values as needed.

- The installation or upgrade procedure should normally take care of removing `/etc/yum.repos.d/private-eclecticiq.repo`.
In case it does not happen automatically, you can safely remove `/etc/yum.repos.d/private-eclecticiq.repo` manually:"

```
$ rm -fR /etc/yum.repos.d/private-eclecticiq.repo
```

- Run the following command to upgrade the YUM repository list, so that it includes only the `eclecticiq-platform` YUM repo:

```
$ yum upgrade --disablerepo=* --enablerepo=eclecticiq-platform eclecticiq-platform\*
```

If the current directory is `/etc/yum.repos.d` and the `yum upgrade` command returns an error, browse to a directory you have read and write access to, and then run the command from there.

Example:

```
$ cd tmp
$ yum upgrade --disablerepo=* --enablerepo=eclecticiq-platform eclecticiq-platform\*
```

Run the YUM install

To **upgrade** the EclecticlQ Platform packages, do the following:

- Start the VM and log in to the VM OS with the appropriate credentials
- Run the following command(s):

```
# Displays a list of installed platform packages
$ rpm -qa eclecti*

# Upgrade to the latest version with the YUM package manager
$ yum upgrade eclectici*

# Or:
# Upgrade to a specific version number with the YUM package manager
$ yum upgrade eclecticiq-platform-1.15.0
```

- The upgrade process begins.
- When it completes, you may wish to check if the newly installed version is the latest. To do so, Run the following command(s):

```
# Displays a list of installed platform packages
$ rpm -qa eclectic*
```

The package versions in the result list should have a higher version number, corresponding to the latest available public release.

Before starting the installation, a script carries out a version check to detect any previously installed platform versions.

It is possible to upgrade the platform only sequentially, it is not possible to skip versions.

Therefore, if you want to upgrade from release 1.10.0 to 1.14.1, you need to upgrade step by step by installing all intermediate releases until 1.14.1.

Example:

```
# Upgrade from 1.10.0 to 1.10.1
$ yum install -y eclecticiciq-platform-1.10.1

# Upgrade from 1.10.1 to 1.11.0
$ yum install -y eclecticiciq-platform-1.11.0

# Upgrade from 1.11.0 to 1.12.0
$ yum install -y eclecticiciq-platform-1.12.0

# Upgrade from 1.12.0 to 1.13.0
$ yum install -y eclecticiciq-platform-1.13.0

# Upgrade from 1.13.0 to 1.14.0
$ yum install -y eclecticiciq-platform-1.14.0

# Upgrade from 1.14.0 to 1.14.1
$ yum install -y eclecticiciq-platform-1.14.1
```

If the currently installed version details are the same as the corresponding information in the upgrade version, or if the former does not immediately precedes the latter, the upgrade procedure is aborted, and an error message is displayed:

```
You currently have version ${INSTALLED_VERSION}, but you need to have ${PREV_VERSION}
installed in order to upgrade ${PACKAGE_NAME}.
```

```
exit 99
```

<code>\${INSTALLED_VERSION}</code>	The currently installed version on your system. The upgrade fails because this version is too old.
<code>\${PREV_VERSION}</code>	Before starting the upgrade procedure again, install this platform version on your system. The upgrade procedure works only when the old and the new versions are contiguous in sequence.

<code>exit 99</code>	The exit code accompanying the error message.
----------------------	---

Check the configuration

After installing the platform, browse to `/opt/eclecticiq/etc/eclecticiq/`. Configuration files are stored here. You can find both the new/latest configuration files, as well as the ones belonging to the previous version of the platform.

<code>platform_settings.py</code>	The main configuration file for the platform
<code>opentaxii.yml</code>	The OpenTAXII (http://opentaxii.readthedocs.org/) configuration file

Verify that the platform configuration files reflect the new, upgraded environment. For example Neo4j has been upgraded from version 2.24 to 2.3.1 Community: make sure the environment described in the upgraded platform configuration file refers to the higher version number.

i You may need to carry out this task manually. In this case, you can diff the files with a tool like **Meld** (<http://meldmerge.org/>).

Check third-party configurations

After checking the platform configuration to make sure it correctly describes the upgraded environment, do the same with third-party product configurations.

i You may need to carry out this task manually. In this case, you can diff the files with a tool like **Meld** (<http://meldmerge.org/>).

Install poppler-utils

The platform requires **Poppler** (<https://poppler.freedesktop.org/>) to process PDF ([https://en.wikipedia.org/wiki/poppler_\(software\)#poppler-utils](https://en.wikipedia.org/wiki/poppler_(software)#poppler-utils)) file uploads. To install **poppler-utils** (<https://apps.fedoraproject.org/packages/poppler-utils>), do the following:


```
$ yum install poppler-utils
```

Migrate the database

When you upgrade the platform by downloading and decompressing a VM image, you may wish to import your existing database into the new platform version.

If you are upgrading the platform to a newer release, you need to migrate the existing database to the new platform release.

To perform the database migration, run the following script:

```
$ /opt/eclecticiq/migrations/db-migration.sh
```

Reindex Elasticsearch

To make sure you are applying the latest Elasticsearch schema, fully reindex Elasticsearch.

To do so, run the same command you would execute to create the Elasticsearch index.

- Configure the platform settings environment variable by prepending it to the `manage synchronize_elasticsearch` command.
- Run `manage synchronize_elasticsearch` to reindex all Elasticsearch indexes.
The command runs in the background. In case of an SSH disconnection, the process should keep running normally.

```
$ INTELWORKS_PLATFORM_SETTINGS=/opt/eclecticiq/etc/eclecticiq/platform_settings.py  
/opt/eclecticiq/platform/api/bin/manage synchronize_elasticsearch > ./reindexing.log  
2>&1 &
```

- Reindexing messages are logged to *reindexing.log*. You can open this file to inspect the scheduling of the reindexing workers.
Example:

```
{
  "avg_time_per_object_ms": 7,
  "chunk_time_ms": 57838,
  "event": "search_reindexing.processed_chunk",
  "level": "info",
  "logger": "eiq.platform.search_index",
  "n_processed": 7980,
  "request_body_size": 3349388,
  "timestamp": "2016-08-05T14:31:27.626829Z",
  "total_avg_time_per_object_ms": 17
}
```

Migrate the graph database



Make sure Neo4j is up and running before applying the graph database migrations.

- Before creating the graph schema, stop the `graph-ingestion` and any running `intel-ingestion` tasks:

```
$ supervisorctl stop graph-ingestion
$ supervisorctl stop intel-ingestion:*
```

- Then, run the following command to create the graph schema and to apply the necessary graph database migrations to the Neo4j database:

```
$ /opt/eclecticiq/platform/api/bin/manage do_graph_migrations
```

Check component versions

After migrating the database, check the versions of the upgraded components to verify that they are the correct ones.

Authenticate with the platform API to receive a bearer token, then pass it to the `/api/versions` API endpoint to obtain version information:

```
$ curl -X GET
-v
--insecure
-i
-H "Content-Type: application/json"
-H "Accept: application/json"
-H "Authorization: Bearer <token>"
https://platform.host/api/versions
```

The JSON response contains version information about the core platform components:

```
{
  "data": {

    "opentaxii": {
      "source": "devpi",
      "tag": "latest",
      "version": "1.12.0"
    },

    "opentaxii-platform-apis": {
      "branch": "master",
      "source": "github",
      "summary": "release/1.12.0",
      "tag": "latest",
      "version": "bc0478b3d10da0a3583563814e05d53116cfaba4"
    },

    "platform-api": {
      "branch": "master",
      "source": "github",
      "summary": "release/1.12.0",
      "tag": "latest",
      "version": "1b7d297394d716da79f9310f377ef72835b8427e"
    },

    "platform-database": {
      "current": "123a456bcd7",
      "head": "123a456bcd7 (head)",
      /*
        Allowed status values:
        - Up to date
        - Outdated
        - Head not available (when no manifest file is found)
      */
      "status": "Up to date"
    },

    "platform-ui": {
      "source": "artifactory",
      "tag": "latest",
      "version": "release/1.12.0"
    }
  }
}
```

This version information matches the corresponding details in the manifest files.
Manifest files are stored in the following directory:

```
$ cd /opt/eclecticiq/manifests/
```

The directory contains these manifest files:

```
opentaxii.mf
opentaxii-platform-apis.mf
platform-api.mf
platform-database.mf
platform-ui.mf
```

As for the (migrated) database, the current database version needs to match the one in the `platform-database.mf` manifest file:

- Request the current database version:

```
$ cd /opt/eclecticiq/migrations/ &&
INTELWORKS_PLATFORM_SETTINGS=/opt/eclecticiq/etc/eclecticiq/platform_settings.py
/opt/eclecticiq/platform/api/bin/manage alembic current
```

Example:

```
123a456bcd7 (head)
```

- Go to the manifest directory:

```
$ cd /opt/eclecticiq/manifests/
```

- Open the *platform-database.mf* manifest file:

```
$ nano platform-database.mf
```

- Verify that the database head hash in the manifest matches the returned current database version:

```
head: 123a456bcd7 (head)
```

Run the fixtures

To generate the default platform fixtures, run the following command(s):

```
$ /opt/eclecticiq/migrations/load-fixtures.sh
```

Run a final check

As a last step before launching the platform, it is good practice to check the following points:

- Core processes and services
- Search, indexing and graph
- Availability

Check core processes and services

Nginx

To check if Nginx is running, run the following command(s):

```
$ service nginx status
```

Supervisor

To check if the `supervisord` daemon is running, run the following command(s):

```
$ service supervisord status
```

PostgreSQL

To check if PostgreSQL is running, run the following command(s):

```
$ service postgresql-9.5 status
```

or:

```
$ systemctl list-units | grep -i postgre
```

Check search indexing and graph

Elasticsearch

To check if Elasticsearch is running, run the following command(s):

```
$ service elasticsearch status
```

Neo4j

Check the main connectivity properties to verify that Neo4j can communicate with other components.

Check neo4j-server.properties

- Open the *neo4j-server.properties* configuration file:

```
$ nano /etc/neo4j/neo4j-server.properties
```

- Make sure that the following properties are set to the corresponding values:

```
org.neo4j.server.webserver.address=0.0.0.0
org.neo4j.server.webserver.port=7474
org.neo4j.server.database.location=/<some_dir>/graph.db
dbms.security.auth_enabled=False
```

graph.db is the default name of the graph database file.

Make sure the directory where it is stored exists before setting it in this configuration file, and that Neo4j can read and write data to it.

Inspect the amount of allocated memory: Neo4j purrs when you give it plenty of memory to work with.

- Depending on your system resources, allocate at least 4 GB to Neo4j.
- Never set it below 1 GB.

To set the minimum and maximum memory values for Neo4j, do the following:

Check neo4j-wrapper.conf

- Open the *neo4j-wrapper.conf* configuration file:

```
$ nano /etc/neo4j/neo4j-wrapper.conf
```

- Make sure that the following properties are set to the corresponding values:

```
wrapper.java.initmemory=1024
wrapper.java.maxmemory=4096
```

- The recommended value for *wrapper.java.initmemory* is at least 4096 MB or more, if possible. The minimum value should never be lower than 1024 MB.
- The recommended value for *wrapper.java.maxmemory* is at least 4096 MB or more, if possible. This value should never be lower than *wrapper.java.initmemory*.

Check neo4j.properties

Set the size of the page cache to 5G:

- Open the *neo4j.properties* property file:

```
$ nano /etc/neo4j/neo4j.properties
```

- Make sure the `dbms.pagecache.memory` property is set to 5G:

```
dbms.pagecache.memory=5G
```

If this parameter is not configured, Neo4j defaults to using 50% of the available system RAM for page caching.

Check availability

Make sure Elasticsearch and Neo4j are available by making cURL calls to the corresponding endpoints:

```
# Check Elasticsearch availability
$ curl localhost:9200

# Check Neo4j availability
$ curl localhost:7474
```

If Neo4j is not running, restart the service by running the following command(s):

```
$ service neo4j start
```

or:

```
$ supervisorctl start neo4j
```

Reload the Supervisor configuration

To reload the Supervisor configuration and to restart all Supervisor-managed processes run the following command(s):

```
$ supervisorctl reload
```


To check the statuses of the tasks managed by Supervisor, run the following command(s):

```
$ supervisorctl status
```

The response should return `RUNNING` for all relevant tasks to confirm that all Supervisor tasks are being executed normally.

The following example serves as a guideline:

graph-ingestion	RUNNING	pid 3443,	uptime 1:10:56
intel-ingestion:0	RUNNING	pid 3323,	uptime 1:11:19
intel-ingestion:1	RUNNING	pid 3336,	uptime 1:11:14
intel-ingestion:2	RUNNING	pid 3363,	uptime 1:11:09
intel-ingestion:3	RUNNING	pid 3372,	uptime 1:11:04
kibana	RUNNING	pid 13837,	uptime 9 days, 0:38:17
neo4j-batching	RUNNING	pid 3590,	uptime 1:10:45
opentaxii	RUNNING	pid 3662,	uptime 1:10:39
platform-api	RUNNING	pid 2999,	uptime 1:12:47
search-ingestion	RUNNING	pid 3513,	uptime 1:10:49
task:beat	RUNNING	pid 3102,	uptime 1:12:39
task:enrichers	RUNNING	pid 3110,	uptime 1:12:26
task:integrations	RUNNING	pid 3135,	uptime 1:12:13
task:providers	RUNNING	pid 3204,	uptime 1:12:01
task:reindexing	RUNNING	pid 3223,	uptime 1:11:48
task:utilities	RUNNING	pid 3242,	uptime 1:11:35

Access the platform

- In your host machine, launch a web browser (recommended: Google Chrome).
- In the address bar, enter the VM IP address.
Example: `https://10.0.2.148`
- The platform login screen is displayed.
- Sign in with the appropriate credentials.



Warning: The browser may display an untrusted connection warning: add it as an exception, and then proceed to the platform.