

EclectiQ Platform requirements

Hardware and software requirements for system administrators

Last generated: March 06, 2018



©2018 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.

©2018 by EclecticIQ BV. All rights reserved.
Last generated on Mar 6, 2018

Table of contents

Table of contents	2
Before you start	3
About EclecticIQ Platform	4
Hardware requirements	4
Single box	4
Data mount points	5
Scaling out	6
Software requirements	6
Credentials and host name	6
Operating systems	7
Web browsers	7
Encoding	8
Time zone	8
Third-party products	8
Bundled third party software	9
SELinux	14
Check SELinux status	14
Check SELinux mode	14
Set SELinux to permissive mode	15
Post-installation check	15
SELinux is not installed	16
SELinux is installed but it is not enabled	17
Whitelist URLs	19
Repositories	19
Enrichers and feeds	19
Other URLs	21
Open ports	22

Before you start

Review these system requirements before proceeding to install the platform from an RPM package or a tarball.

Command examples

At times throughout this document, you may need to enter commands in the terminal, in the console, or in the command line. The commands we ask you to enter are prefixed by `$`, and they look like this:

```
$ run command
```

Code examples

Inline code comments, especially for bash and shell command examples, are usually prefixed by `#`:

```
# This is a self-explanatory inline code comment  
  
$ run command
```

Code and configuration examples look like this:

```
{  
  "this" : [  
    {  
      "key_name" : "key_value"  
    }  
  ]  
}
```

The actual format of each code snippet depends on the corresponding programming language.

Text editors

When you need to open a file to edit it, code examples use the Vim text editor.

Example:

```
$ vi /etc/nginx/conf.d/platform.conf
```

Feel free to replace `vi` with your preferred weapon of choice, be it **Vim** (<http://www.vim.org/>), **Emacs**, or (<https://www.gnu.org/software/emacs/>) **Nano** (<https://www.nano-editor.org/>).
(https://imgs.xkcd.com/comics/real_programmers.png)

About EclecticIQ Platform

EclecticIQ Platform is powered by **STIX** (<https://stixproject.github.io/>) and **TAXII** (<http://taxiiproject.github.io/about/>) open standards.

It enables ingesting, consolidating, analyzing, integrating, and collaborating on intelligence from multiple sources.

EclecticIQ Platform	Key features
Feed management	Manage multiple cyber threat intelligence feeds from any source, in many different formats.
Enrichment	Enrich existing intelligence with external data sources, and refine it with de-duplication and pattern recognition.
Sharing	Investigate and identify threats together with partners as part of an information ecosystem.
Collaboration	Analyze and create intelligence in collaboration with other teams and departments.
Insights	Generate insight thanks to a high-fidelity, normalized view into your intelligence.
Integration	Understand how cyber intelligence relates to and how it can affect your organization and your environment.

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	-	Physical machine/ <i>rpm</i> and <i>deb</i> installs	

HW area	Minimum	Recommended	Notes
CPUs	4	8	Core count includes HT
CPU speed	2.5 GHz	2.5 GHz or faster	
Memory	32 GB	64 GB or more	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. Operations and tasks carried out through the web-based UI may be memory-intensive: the web browser can use ~1 GB or more, occasionally. 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
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

Data mount points

When installing and configuring platform components such as PostgreSQL, Elasticsearch, and Neo4j you need to specify dedicated locations where these products store their data.

The table below shows the recommended mount point paths for each data store.

Component	Mount point	Minimum size (GB)	Recommended size (GB)
Logs	<code>/var/log</code>	10	20
Neo4j	<code>/media/neo4j</code>	25	50
Elasticsearch	<code>/media/elasticsearch</code>	50	75
PostgreSQL	<code>/media/pgsql</code>	200	300

Redis is installed to the root partition where the platform is installed to. During the configuration step, you can set the Redis data location in the *redis.conf* configuration file. The recommended target directory for Redis data is */media/redis*. This is not a mount point on a separate partition, it is a subdirectory in the root partition.

Scaling out

The easiest approach to scaling out involves allocating dedicated machines to the databases. In this scenario, you install each of the following components on a separate machine:

- EclecticIQ Platform
- PostgreSQL
- Redis
- Elasticsearch
- Neo4j

To optimize read-write operations and to ensure that the storage drives are fast, set up dedicated drives per partition.

Software requirements

Credentials and host name

To correctly configure the system after installing the required dependencies and third-party products, ensure you have the following information available:

- DNS name of the host machine you are going to use to access the platform.
Example: `${platform_host}`
- SSL certificate and key for the web server.
- EclecticIQ Platform login credentials.

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

EclecticIQ Platform default login credentials	
user name	admin
password	EclecticIQ2015#

Operating systems

Supported operating systems:

- **CentOS Linux 7 (1511)** (<https://lists.centos.org/pipermail/centos-announce/2015-december/021518.html>)
- **Red Hat Enterprise Linux 7** (<https://www.redhat.com/>)
- **Ubuntu Server 16.04 LTS (Xenial Xerus)** (<https://wiki.ubuntu.com/xenialxerus/releasenotes>)

Web browsers

EclecticIQ Platform web-based GUI supports the following web browsers:

- Google Chrome (latest)

Encoding

The platform default character encoding is **UTF-8** (<http://www.utf-8.com/>). Dependencies and components that exchange data with the platform need to apply the same encoding.

Time zone

The global time zone configuration needs to be **UTC**.

While you can set a local or a custom time zone value for the platform, the host environment needs to be consistently on **UTC time**. This includes OS, databases, as well as any other products or components that allow setting a time zone, and that interact/interoperate with the platform.

Third-party products

Third-party software includes required dependencies for EclecticIQ Platform to operate correctly.

Make sure that the following software products are *already installed* on the target system *before* installing the platform. During installation, the platform checks for these dependencies. If they are missing, the installation procedure aborts.

Dependency	Version	Reference
Oracle Java JDK	1.8.0	Oracle Java SE web site (https://www.oracle.com/java/technologies/java-se.html)
PostgreSQL	10.1	PostgreSQL web site (https://yum.postgresql.org/repopackages.php)
Redis	3.2.10	Redis web site (http://redis.io)

Dependency	Version	Reference
Nginx	1.12.x	Nginx web site (https://nginx.org)
Neo4j	3.3.3 Community	Neo4j web site (http://neo4j.com)
Elasticsearch	5.6.7	Elasticsearch web site (https://www.elastic.co)
elasticdump	3.0.0	elasticdump Node js module page (https://www.npmjs.com/package/elasticdump)
Logstash	5.6.8	Logstash reference documentation (https://www.elastic.co/guide/en/logstash/5.6/index.html)
Kibana	5.6.8	Kibana reference documentation (https://www.elastic.co/guide/en/kibana/current/getting-started.html)
unzip	n/a	Compress and decompress file archives ()
Node.js	6.x.x or later	Node.js web site (https://nodejs.org/)
poppler-utils	n/a	poppler-utils reference page (https://apps.fedoraproject.org/packages/poppler-utils)
Postfix	n/a	Postfix web site (http://www.postfix.org/)
Supervisor	n/a	Supervisor web site (http://supervisord.org/)
StatsD	n/a	Metrics aggregator for the dashboard (https://github.com/etsy/statsd)
statsd- elasticsearch- backend	n/a	Backend for Elasticsearch to work with StatsD (https://github.com/markkimsal/statsd-elasticsearch-backend)

Warning: About Elasticsearch

During complex index upgrades and reindexing operations, Elasticsearch may require additional disk space to store temporary working files and temporary copies of the existing indices.

Monitor your Elasticsearch partition usage. Before it reaches 50% of the available space in the partition, extend it, so that the new partition size is at least twice as large as the sum of the existing Elasticsearch indices.

Example: if Elasticsearch currently uses 43 GB of disk space, extend the partition where Elasticsearch lives, so that it is at least 86 GB.

Bundled third party software

EclectiQ Platform is bundled with the following third party software.

Each product on the list abides by its own terms and conditions and its own license.

Dependency	Version	Reference
alembic	0.9.6	Bitbucket
amqp	1.4.9	GitHub
anyjson	0.3.3	Bitbucket
apispec	0.4.1	GitHub
appnope	0.1.0	GitHub
argswargs	1.0.3	GitHub
asn1crypto	0.23.0	GitHub
attrs	17.3.0	attrs
bcrypt	3.1.4	GitHub
beautifulsoup4	4.6.0	Crummy
billiard	3.3.0.23	GitHub
blinker	1.4	PythonHosted
boto3	1.4.7	GitHub
botocore	1.7.45	GitHub
cabby	0.1.18	GitHub
cachetools	2.0.1	GitHub
celery	3.1.18	Celery
certifi	2017.11.5	Certifi
cfffi	1.11.2	CFFI
chardet	3.0.4	GitHub
click	6.7	GitHub
colorama	0.3.9	GitHub
colorlog	3.1.0	GitHub

Dependency	Version	Reference
cryptography	2.1.3	GitHub
datauri	1.0.0	GitHub
decorator	4.1.2	GitHub
defusedxml	0.5.0	GitHub
docutils	0.14	Docutils
elasticsearch	1.7.0	GitHub
fancycompleter	0.8	Bitbucket
feedparser	5.2.1	GitHub
flamegraph	0.1	GitHub
Flask	0.10.1	GitHub
Flask-Classy	0.6.10	GitHub
Flask-JWT	0.2.0	GitHub
flask-marshmallow	0.8.0	GitHub
Flask-Redis	0.3.0	GitHub
Flask-SQLAlchemy	2.3.2	GitHub
furl	1.0.1	GitHub
future	0.16.0	Python-Future
gunicorn	19.7.1	Gunicorn
idna	2.6	GitHub
inflect	0.2.5	pyPI
ipdb	0.10.3	GitHub
ipython	6.2.1	IPython
ipython_genutils	0.2.0	IPython
itsdangerous	0.24	GitHub
jedi	0.11.0	GitHub
Jinja2	2.10	Jinja
jmespath	0.9.3	GitHub

Dependency	Version	Reference
kombu	3.0.37	Kombu
libtaxii	1.1.111	TAXII
lxml	4.1.1	lxml
Mako	1.0.7	mako
MarkupSafe	1.0	GitHub
marshmallow	2.6.1	GitHub
marshmallow-sqlalchemy	0.13.2	GitHub
newrelic	2.96.0.80	New Relic
orderedmultidict	0.7.11	GitHub
paramiko	2.4.0	GitHub
parso	0.1.0	GitHub
pdbpp	0.9.2	GitHub
pexpect	4.3.0	Pexpect
pickleshare	0.7.4	GitHub
prompt_toolkit	1.0.15	GitHub
psycpg2	2.7.3.2	init.d
ptyprocess	0.5.2	GitHub
pyasn1	0.3.7	GitHub
pycparser	2.18	GitHub
Pygments	2.2.0	Pygments
PyJWT	1.4.0	GitHub
pyldap	2.4.25.1	GitHub
pymisp	2.4.80	GitHub
PyNaCl	1.2.0	GitHub
pyOpenSSL	17.3.0	pyOpenSSL
pysaml2	4.5.0	GitHub

Dependency	Version	Reference
python-dateutil	2.4.2	dateutil
python-editor	1.0.3	GitHub
python-gnupg	0.4.1	PythonHosted
python-magic	0.4.13	GitHub
python-slugify	1.2.4	GitHub
pytz	2017.3	PythonHosted
PyYAML	3.12	PyYAML
rarfile	3.0.0	GitHub
redis	2.10.6	GitHub
requests	2.18.4	Requests
requests-futures	0.9.7	GitHub
retrying	1.3.3	GitHub
rfc3986	1.1.0	RFC 3986
s3transfer	0.1.11	GitHub
simplegeneric	0.8.1	CheeseShop
six	1.11.0	PyPI
SQLAlchemy	1.1.15	SQLAlchemy
statsd	3.2.1	GitHub
structlog	16.0.0	structlog
tld	0.7.9	GitHub
traitlets	4.3.2	IPython
unrar	0.3	GitHub
urllib3	1.22	urllib3
wcwidth	0.1.7	GitHub
Werkzeug	0.12.2	The Pallets Projects
wmctrl	0.3	Bitbucket

SELinux

EclecticIQ Platform supports **SELinux** (<http://selinuxproject.org/>).

- If you are using or plan to use SELinux in the environment where the platform is installed, you should carry out this check.
- If you are not using SELinux and are not planning to implement it in the environment where the platform is installed, you do not need to do anything and you can safely disregard this section.

Check SELinux status

If SELinux is installed, check if it is enabled or disabled. Run the following command(s):

```
$ sestatus -v
```

If SELinux is disabled, the response includes the following line:

```
SELinux status: disabled
```

Check SELinux mode

You can check also which SELinux mode is currently active. Run the following command(s):

```
$ getenforce
```

The allowed modes are `enforcing`, `permissive`, and `disabled`.

The active mode may not be the same as the `SELINUX` value defined in the SELinux global configuration file:

```
# This file controls the state of SELinux on the system.
# SELINUX= can take one of these three values:
#     enforcing - SELinux security policy is enforced.
#     permissive - SELinux prints warnings instead of enforcing.
#     disabled - No SELinux policy is loaded.
SELINUX=permissive
# SELINUXTYPE= can take one of three two values:
#     targeted - Targeted processes are protected,
#     minimum - Modification of targeted policy. Only selected processes are
protected.
#     mls - Multi Level Security protection.
SELINUXTYPE=targeted
```

This can happen after changing and saving **SELinux global configuration file**

(<https://selinuxproject.org/page/configurationfiles>), and before executing a system reboot for the changes to become effective.

Set SELinux to permissive mode

The recommended SELinux mode to offload complexity during installation and configuration is `permissive`. To set SELinux to work permissively, run the following command(s):

```
$ setenforce permissive
```

Post-installation check

The platform post-install script included in the RPM package attempts to automatically set the appropriate SELinux security labels for the files that are deployed to `/opt/eclecticiq`.

- If SELinux is not installed or if it is disabled, the post-install script included in the RPM install package does not attempt to configure any SELinux file security labels for the files that are deployed to `/opt/eclecticiq`.
- If SELinux is installed and it is enabled, and if the platform post-install script does not set the SELinux security labels to the applicable platform files, run the following command(s):

```
$ semanage fcontext -a -t var_log_t -f d "/opt/eclecticiq"
```

- If SELinux policy-related errors occur, the command returns a response that can be similar to this example:

```
SELinux: Could not downgrade policy file /etc/selinux/targeted/policy/policy.29,
searching for an older version.
SELinux: Could not open policy file <= /etc/selinux/targeted/policy/policy.29: No
such file or directory
/sbin/load_policy: Can't load policy: No such file or directory
libsemanage.semanage_reload_policy: load_policy returned error code 2.
```

The response provides more context about the affected files and the reasons why it was not possible to set the security labels.

SELinux is not installed

If SELinux is not installed on the target system, do the following:

- After completing the platform installation, install and enable SELinux.
- To set the correct security contexts, execute the following script:

```
BASE_PATH="/opt/eclecticiq"

if [ -x "$(command -v semanage)" ]; then

    SELINUX_MODE=$(getenforce)

    if ! [ $SELINUX_MODE == "Disabled" ]; then

        semanage fcontext -a -t etc_t "$BASE_PATH/etc(/.*)?"
        semanage fcontext -a -t etc_t "$BASE_PATH/etc-extras(/.*)?"

        semanage fcontext -a -t httpd_config_t "$BASE_PATH/etc/nginx(/.*)?"
        semanage fcontext -a -t httpd_config_t "$BASE_PATH/etc-extras/nginx(/.*)?"

        # By default, newly created files and directories inherit the SELinux type
        # of the corresponding parents, so that log files have the correct type.
        # However, we do not want to relabel existing logs.
        semanage fcontext -a -t var_log_t -f d "$BASE_PATH/logs"

        restorecon -RF $BASE_PATH

        echo "SELinux security labels configured."
    else
        echo "SELinux is not enabled. Security labels won't be configured."
    fi
else
    echo "SELinux is not installed. Security labels won't be configured."
fi
```

- You may need to reboot the system for the changes to become effective.

SELinux is installed but it is not enabled

If SELinux is installed on the target system but it is not enabled, do the following:

- Enable SELinux, either by editing its configuration file, and then by rebooting the system, or by running one of the following commands:

```
# Set SELinux to permissive mode
$ setenforce 0

# Set SELinux to enforcing mode
$ setenforce 1
```

- Create the following bash script:

```
BASE_PATH="/opt/eclecticiq"

if [ -x "$(command -v semanage)" ]; then

    SELINUX_MODE=$(getenforce)

    if ! [ $SELINUX_MODE == "Disabled" ]; then

        semanage fcontext -a -t etc_t "$BASE_PATH/etc(/.*)?"
        semanage fcontext -a -t etc_t "$BASE_PATH/etc-extras(/.*)?"

        semanage fcontext -a -t httpd_config_t "$BASE_PATH/etc/nginx(/.*)?"
        semanage fcontext -a -t httpd_config_t "$BASE_PATH/etc-extras/nginx(/.*)?"

        # By default, newly created files and directories inherit the SELinux type
        # of the corresponding parents, so that log files have the correct type.
        # However, we do not want to relabel existing logs.
        semanage fcontext -a -t var_log_t -f d "$BASE_PATH/logs"

        restorecon -RF $BASE_PATH

        echo "SELinux security labels configured."
    else
        echo "SELinux is not enabled. Security labels won't be configured."
    fi
else
    echo "SELinux is not installed. Security labels won't be configured."
fi
```

- Save it, make it executable, and then run it.

- You may need to reboot the system for the changes to become effective.

Whitelist URLs

The platform needs to access external data sources to ingest intel, as well as to enrich entities and observables. You may want to whitelist these URLs, domains and addresses, so that the platform can communicate with the external intel and service providers.

Repositories

When installing or upgrading the platform and its dependencies, the system needs to access the following source repositories.

Repository URL	Belongs to	Repo type
https://downloads.eclecticiq.com/	EclectiQ Platform	deb, rpm
https://dl.fedoraproject.org/pub/epel/7/x86_64/	EPEL	rpm

Enrichers and feeds

Feeds and enrichers access data sources through these URLs. Whitelist the domains and allow traffic to and from them.

Domain	Belongs to	Type
http://\${elasticsearch_instance_url}:9200/\${schema_resource}	Elasticsearch sightings	enricher
https://cybercrime-portal.fox-it.com/	Fox-IT INTELL Portal	enricher, incoming feed
https://api.intel471.com/v1/	Intel 471	enricher, incoming feed
http://api.openresolve.com/{}/{	OpenDNS OpenResolve	enricher
http://\${pydat_instance_url}:8000/	PyDat	enricher
https://stat.ripe.net/data/geoloc/{	RIPEstat GeolIP	enricher

Domain	Belongs to	Type
https://stat.ripe.net/data/whois/{}	RIPEstat Whois	enricher
https://panacea.threatgrid.com/api/v3/	Cisco Threat Grid	enricher, incoming feed
https://www.virustotal.com/vtapi/v2/{}	VirusTotal	enricher
https://endlesstunnel.info/v3	Flashpoint AggregINT	enricher
https://endlesstunnel.info/v3	Flashpoint Blueprint	enricher
https://endlesstunnel.info/v3	Flashpoint Thresher	enricher
https://api.passivetotal.org/v2	PassiveTotal Whois	enricher
https://api.passivetotal.org/v2	PassiveTotal Passive DNS	enricher
https://api.passivetotal.org/v2	PassiveTotal IP/Domain	enricher
https://api.passivetotal.org/v2	PassiveTotal Malware	enricher
http://\${splunk_instance_url}:8089/	Splunk sightings	enricher
http://api.domaintools.com/v1/{}/host-domains	DomainTools Hosted Domains	enricher
http://api.domaintools.com/v1/reputation	DomainTools Reputation	enricher
https://api.domaintools.com/v1/{}/host-domains	DomainTools Suspicious Domains	enricher
https://api.isightpartners.com/search/{}	FireEye iSIGHT	enricher
https://api.recordedfuture.com/live/sc/entity/{}	Recorded Future	enricher
https://unshorten.me/s/{}	Unshorten-URL	enricher
https://api.dnsdb.info/{}	Farsight DNSDB	enricher
https://www.threatcrowd.org/{}	ThreatCrowd	enricher
https://censys.io/api/v1/search/ipv4	Censys	enricher

Domain	Belongs to	Type
http://api.domaintools.com/v1/{}/name-server-domains/	DomainTools Malicious Server Domains	enricher
http://api.domaintools.com/v1/{}/whois/parsed	DomainTools Parsed Whois	enricher
https://intelapi.crowdstrike.com/indicator/v1/search/{}	CrowdStrike Falcon Intelligence Indicator	enricher
http://api.domaintools.com/v1/reverse-whois/{}	DomainTools Reverse Whois	enricher
https://cve.circl.lu/api/cve/	CVE Search	enricher
https://www.circl.lu/v2pssl/cquery/{}	CIRCL IPs related to SSL certificate	enricher
https://www.circl.lu/v2pssl/cfetch/{}	CIRCL SSL Certificate Fetcher	enricher
https://api.shodan.io/shodan/	Shodan	enricher
https://api.spycloud.io/sp-v1/breach	SpyCloud Breach Data	enricher
https://api.emaildefense.proofpoint.com/v1	Proofpoint Email Threat	enricher
http://\${misp_instance_url}/	MISP API	enricher
/absolute/path/to/GeoLite2-City.mmdb	MaxMind GeoIP	enricher
https://nti.nsfocusglobal.com/api/v1/search/{}	NSFocus Intelligence	enricher

Other URLs

Domain	Belongs to	Type
http://\${variable_subdomain}.cyberfeed.net:\${port_number}	AnubisNetworks	incoming feed

Domain	Belongs to	Type
https://www.threathq.com/	PhishMe Intelligence	incoming feed
https://api.threatrecon.co/	Threat Recon	incoming feed
http://hailataxii.com	Hail a TAXII	open source cyber threat intelligence source
https://test.taxiistand.com/	TAXII Stand	public OpenTAXII test server

Open ports

The platform components communicate with the platform and with each other through these ports. Make sure they are open within the platform network.

Port	Belongs to
9200	elasticsearch
5601	kibana
7474; 7473	neo4j
4008	neo4j-batching
8008	platform-api
5432	postgresql-10
6379	redis
6755	logstash
80; 443	nginx
25; 587	postfix
9001	supervisor
8125	statsd

